detnet N. Finn
Internet-Draft P. Thubert
Intended status: Standards Track Cisco
Expires: April 26, 2015 October 23, 2014

Deterministic Networking Problem Statement
draft-finn-detnet-problem-statement-01

Abstract

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

Table of Contents

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 guaranteed 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.

In order to serve this extended requirement, the IETF and the IEEE must collaborate and define an abstract model that can be applicable both at Layer-2 and Layer-3, and along segments of different technologies. 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 IEEE802.1 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 IEEE802.15.4e MAC.

The proposed model should enable a fully scheduled operation orchestrated by a central controller, as well as 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.

Once the abstract model is agreed upon, the IETF will need to 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 IETF will also need to specify the necessary protocols, or protocol additions, based on relevant IETF technologies such as PCE, MPLS and 6TiSCH, to implement the selected model. As a result of this work, it will be possible to establish a multi-hop path over the IP network, for a particular flow with precise timing and throughput requirements, and carry this particular flow along the multi-hop path with such characteristics as low latency and ultra-low jitter, duplication and elimination of packets over non-congruent paths for a higher delivery ratio, and/or zero congestion loss. 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, and an existing path may be moved or removed.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

3. 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 that 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.

Considerable experience ([ODVA],[AVnu], [Profinet],[HSR-PRP], etc...) has shown that these applications need a 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 critical packet flows that:

   * Can be unicast or multicast;

   * Need absolute guarantees of minimum and maximum latency end-to-end across the network;

   * Need a packet loss ratio in the range of 1.0e-9 to 1.0e-12, or better;

   * Can, in total, absorb more than half of the network's available bandwidth (that is, 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.

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.

When 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 set of queuing, shaping, and scheduling algorithms (:::reference to section, below :::) that enable each bridge or router to compute the exact number of buffers to be allocated for each flow or class of flows. The present authors assume that these techniques will be used by the DetNet Working Group.

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 unthrottled (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.

4. Related IETF work

4.1. 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].

4.2. 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 flows, and Traffic Engineering capabilities. Typical applications are open and closed control loops, as well as supervisory control flows and management.

The 6TiSCH Working Group focuses only on the TSCH mode of the IEEE802.15.4e 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.

5. Problem Statement

5.1. 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.

5.2. Centralized Path Computation and Installation

A centralized routing model, such as provided with a PCE, enables global and per-flow optimizations. The model is attractive but a number of issues are left to be solved. In particular:

- 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.

5.3. 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.

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.

Security must cover:

- the protection of the signaling protocol
- the authentication and authorisation of the controlling nodes
- the identification and shaping of the flows

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. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

9.2. 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.", .

[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.", .

[I-D.ietf-6tisch-architecture]
Thubert, P., Watteyne, T., and R. Assimiti, "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4e", draft-ietf-6tisch-architecture-03 (work in progress), July 2014.

[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-02 (work in progress), October 2014.

[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-02 (work in progress), September 2014.

[IEEE802.1AS-2011]
IEEE, "Timing and Synchronizations (IEEE 802.1AS-2011)", 2011, <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>.

[IEEE802.1Q-2011]
IEEE, "MAC Bridges and VLANs (IEEE 802.1Q-2011", 2011, <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>.

[IEEE802.1Qav]
IEEE, "Forwarding and Queuing (IEEE 802.1Qav-2009)", 2009, <http://standards.ieee.org/getieee802/download/802.1Qav-2009.pdf>.

[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.", .

[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.

[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.

[WirelessHART]
www.hartcomm.org, "Industrial Communication Networks - Wireless Communication Network and Communication Profiles - WirelessHART - IEC 62591", 2010.

Authors' Addresses

Norm Finn
Cisco Systems
510 McCarthy Blvd
SJ-24
Milpitas, California 95035
USA

Phone: +1 925 980 6430
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