                 Deterministic Networking Problem Statement
                    draft-finn-detnet-problem-statement-05

Abstract

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

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 18, 2016.

Table of Contents

1. Introduction

The Deterministic Networking Use Cases
[I-D.grossman-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, and

while such media must be encompassed by any useful 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.

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, 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 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, duplication 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. 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],[IEC62439],
   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; 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 1.0e-9 to 1.0e-12, or
          better, on Ethernet, and in the order of 1.0e-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;

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.

[4](#). Problem Statement

[4.1](#). 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 and 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 of 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
seu th epath up, whether it is centralized or distributed.

[4.2](#). 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.

[4.3](#). 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:

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 following classical Traffic Engineering (TE)
   models.

To enable a centralized model, DetNet should produce the complete SDN
architecture with describes at a high level the interaction and data
models to:

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

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

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

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

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

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

o provide additional reliability through redundancy, in particular
   with packet replication and elimination;

o indicate the flows and packet sequences in-band with the flows;

4.4. 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 CSPF
    that is aware of DetNet.

3. The path is installed using RSVP-TE, associated with flow
    identification, per-hop behavior such as replication and
    elimination, blocked resources, and flow timing information.

4. Traffic flows are transported through the MPLS-TE tunnel, using
    the reserved resources for this flow at each hop.

4.5. 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.

5. 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.

Security must cover:

o the protection of the signaling protocol

o the authentication and authorization of the controlling nodes

o the identification and shaping of the flows

o the isolation of flows from leakage and other influences from any
   activity sharing physical resources.

6. IANA Considerations

This document does not require an action from IANA.

7. Acknowledgments

The authors wish to thank Lou Berger, 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, Patrick
Wetterwald, Subha Dhesikan, Rudy Klecka and Pat Thaler for their
various contribution to this work.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
           Requirement Levels", BCP 14, RFC 2119,
           DOI 10.17487/RFC2119, March 1997,
           <http://www.rfc-editor.org/info/rfc2119>.

8.2. 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 IEEE Audio Video Bridging
           (AVB) and Time-Sensitive Networking (TSN) standards.".

[CCAMP]   IETF, "Common Control and Measurement Plane",
           <https://datatracker.ietf.org/doc/charter-ietf-ccamp/>.

   [EIP]   http://www.odva.org/, "EtherNet/IP provides users with the
                network tools to deploy standard Ethernet technology (IEEE
                802.3 combined with the TCP/IP Suite) for industrial
                automation applications while enabling Internet and
                enterprise connectivity data anytime, anywhere.",
                <http://www.odva.org/Portals/0/Library/
                Publications_Numbered/
                PUB00138R3_CIP_Adv_Tech_Series_EtherNetIP.pdf>.

   [HART]   www.hartcomm.org, "Highway Addressable Remote Transducer,
                a group of specifications for industrial process and
                control devices administered by the HART Foundation".

   [I-D.finn-detnet-architecture]
                Finn, N., Thubert, P., and M. Teener, "Deterministic
                Networking Architecture", draft-finn-detnet-
                architecture-03 (work in progress), March 2016.

   [I-D.grossman-detnet-use-cases]
                Grossman, E., Gunther, C., Thubert, P., Wetterwald, P.,
                Raymond, J., Korhonen, J., Kaneko, Y., Das, S., and Y.
                Zha, "Deterministic Networking Use Cases", draft-grossman-
                detnet-use-cases-01 (work in progress), November 2015.

   [I-D.ietf-6tisch-architecture]
                Thubert, P., "An Architecture for IPv6 over the TSCH mode
                of IEEE 802.15.4", draft-ietf-6tisch-architecture-09 (work
                in progress), November 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.ietf-teas-yang-te-topo]
                Liu, X., Bryskin, I., Beeram, V., Saad, T., Shah, H., and
                O. Dios, "YANG Data Model for TE Topologies", draft-ietf-
                teas-yang-te-topo-02 (work in progress), October 2015.

   [I-D.svshah-tsvwg-deterministic-forwarding]
                Shah, S. and P. Thubert, "Deterministic Forwarding PHB",
                draft-svshah-tsvwg-deterministic-forwarding-04 (work in
                progress), August 2015.

[I-D.zhao-pce-pcep-extension-for-pce-controller]
          Zhao, Q., Li, Z., Dhody, D., and C. Zhou, "PCEP Procedures
          and Protocol Extensions for Using PCE as a Central
          Controller (PCECC) of LSPs", draft-zhao-pce-pcep-
          extension-for-pce-controller-03 (work in progress), March
          2016.

[IEC62439]
          IEC, "Industrial communication networks - High
          availability automation networks - Part 3: Parallel
          Redundancy Protocol (PRP) and High-availability Seamless
          Redundancy (HSR) - IEC62439-3", 2012,
          <https://webstore.iec.ch/publication/7018>.

[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".

[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>.

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

[MPLS]   IETF, "Multiprotocol Label Switching",
          <https://datatracker.ietf.org/doc/charter-ietf-mpls/>.

[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/>.

[RFC2547] Rosen, E. and Y. Rekhter, "BGP/MPLS VPNs", RFC 2547,
          DOI 10.17487/RFC2547, March 1999,
          <http://www.rfc-editor.org/info/rfc2547>.

[RFC2702] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and J.
          McManus, "Requirements for Traffic Engineering Over MPLS",
          RFC 2702, DOI 10.17487/RFC2702, September 1999,
          <http://www.rfc-editor.org/info/rfc2702>.

[RFC3031]   Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol
            Label Switching Architecture", RFC 3031,
            DOI 10.17487/RFC3031, January 2001,
            <http://www.rfc-editor.org/info/rfc3031>.

[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,
            <http://www.rfc-editor.org/info/rfc3209>.

[RFC3272]   Awduche, D., Chiu, A., Elwalid, A., Widjaja, I., and X.
            Xiao, "Overview and Principles of Internet Traffic
            Engineering", RFC 3272, DOI 10.17487/RFC3272, May 2002,
            <http://www.rfc-editor.org/info/rfc3272>.

[RFC3630]   Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering
            (TE) Extensions to OSPF Version 2", RFC 3630,
            DOI 10.17487/RFC3630, September 2003,
            <http://www.rfc-editor.org/info/rfc3630>.

[RFC3945]   Mannie, E., Ed., "Generalized Multi-Protocol Label
            Switching (GMPLS) Architecture", RFC 3945,
            DOI 10.17487/RFC3945, October 2004,
            <http://www.rfc-editor.org/info/rfc3945>.

[RFC3985]   Bryant, S., Ed. and P. Pate, Ed., "Pseudo Wire Emulation
            Edge-to-Edge (PWE3) Architecture", RFC 3985,
            DOI 10.17487/RFC3985, March 2005,
            <http://www.rfc-editor.org/info/rfc3985>.

[RFC4203]   Kompella, K., Ed. and Y. Rekhter, Ed., "OSPF Extensions in
            Support of Generalized Multi-Protocol Label Switching
            (GMPLS)", RFC 4203, DOI 10.17487/RFC4203, October 2005,
            <http://www.rfc-editor.org/info/rfc4203>.

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

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

[RFC5127]   Chan, K., Babiarz, J., and F. Baker, "Aggregation of
            Diffserv Service Classes", RFC 5127, DOI 10.17487/RFC5127,
            February 2008, <http://www.rfc-editor.org/info/rfc5127>.

[RFC5151]   Farrel, A., Ed., Ayyangar, A., and JP. Vasseur, "Inter-
            Domain MPLS and GMPLS Traffic Engineering -- Resource
            Reservation Protocol-Traffic Engineering (RSVP-TE)
            Extensions", RFC 5151, DOI 10.17487/RFC5151, February
            2008, <http://www.rfc-editor.org/info/rfc5151>.

[RFC5305]   Li, T. and H. Smit, "IS-IS Extensions for Traffic
            Engineering", RFC 5305, DOI 10.17487/RFC5305, October
            2008, <http://www.rfc-editor.org/info/rfc5305>.

[RFC5329]   Ishiguro, K., Manral, V., Davey, A., and A. Lindem, Ed.,
            "Traffic Engineering Extensions to OSPF Version 3",
            RFC 5329, DOI 10.17487/RFC5329, September 2008,
            <http://www.rfc-editor.org/info/rfc5329>.

[RFC5440]   Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation
            Element (PCE) Communication Protocol (PCEP)", RFC 5440,
            DOI 10.17487/RFC5440, March 2009,
            <http://www.rfc-editor.org/info/rfc5440>.

[RFC5673]   Pister, K., Ed., Thubert, P., Ed., Dwars, S., and T.
            Phinney, "Industrial Routing Requirements in Low-Power and
            Lossy Networks", RFC 5673, DOI 10.17487/RFC5673, October
            2009, <http://www.rfc-editor.org/info/rfc5673>.

[RFC7384]   Mizrahi, T., "Security Requirements of Time Protocols in
            Packet Switched Networks", RFC 7384, DOI 10.17487/RFC7384,
            October 2014, <http://www.rfc-editor.org/info/rfc7384>.

[RFC7426]   Haleplidis, E., Ed., Pentikousis, K., Ed., Denazis, S.,
            Hadi Salim, J., Meyer, D., and O. Koufopavlou, "Software-
            Defined Networking (SDN): Layers and Architecture
            Terminology", RFC 7426, DOI 10.17487/RFC7426, January
            2015, <http://www.rfc-editor.org/info/rfc7426>.

[RFC7554]   Watteyne, T., Ed., Palattella, M., and L. Grieco, "Using
            IEEE 802.15.4e Time-Slotted Channel Hopping (TSCH) in the
            Internet of Things (IoT): Problem Statement", RFC 7554,
            DOI 10.17487/RFC7554, May 2015,
            <http://www.rfc-editor.org/info/rfc7554>.

[TEAS]   IETF, "Traffic Engineering Architecture and Signaling",
            <https://datatracker.ietf.org/doc/charter-ietf-teas/>.

[TiSCH]   IETF, "IPv6 over the TSCH mode over 802.15.4",
            <https://datatracker.ietf.org/doc/charter-ietf-6tisch/>.

   [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: +1 408 526 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