                        6LoWPAN Backbone Router
                  draft-thubert-6lowpan-backbone-router-03

Abstract

   Some LLN subnets are expected to scale up to the thousands of nodes
   and hundreds of routers. This paper proposes an IPv6 version of the
   Backbone Router concept that enables such a degree of scalability
   using a high speed network as a backbone to the subnet.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

In order to meet industrial requirements for non-critical monitoring,
alerting, supervisory control, open loop control and some closed loop
control applications, both [ISA100.11a] and wireless [HART] standards
leverage advanced technology at every layer, including the Time
Synchronized Channel Hopping (TSCH) Medium Access Control (MAC)
operation that enables deterministic behaviours for time-sensitive
flows. Additionally, [ISA100.11a] endorsed the 6LoWPAN Header
Compression [RFC6282] format for the network header, making it
possible to utilize IPv6 based protocols such as BACnet IP, Profibus
IP and Modbus TCP without significant changes to those protocols.

[ISA100.11a] has introduced the concept of a Backbone Router that
would interconnect small LLNs over a high speed Backbone network and
scale a single ISA100.11a network up to the thousands of nodes. In
that model the LLNs and the backbone form a single subnet in which
nodes can move freely without the need of renumbering, and the
Backbone Router is a special kind of Border Router designed to manage
the interaction between the LLNs and the backbone at layer 3.
Similar scalability requirements exist in the metering and monitoring
industries. In a network that large, it is impossible for a node to
register to all Border Routers as suggested for smaller topologies in
Neighbor Discovery Optimization for Low-power and Lossy Networks
[RFC6775].

This paper specifies IP layer functionalities that are required to
implement the concept of a Backbone Router with IPv6, in particular
the application of the "IP Version 6 Addressing Architecture"
[RFC4291], " the Neighbor Discovery Protocol" [RFC4861] and "IPv6
Stateless Address Autoconfiguration" [RFC4862].

The use of EUI-64 based link local addresses, Neighbor Discovery
Proxying [RFC4389], Neighbor Discovery Optimization for Low-power
and Lossy Networks [RFC6775], the IPv6 Routing Protocol for Low power
and Lossy Networks [RFC6550] , the mixed mode of Efficiency aware
IPv6 Neighbor Discovery Optimizations
[I-D.chakrabarti-nordmark-6man-efficient-nd] and Optimistic Duplicate
Address Detection [RFC4429] are discussed. Also, the concept of
Backbone Link is introduced to implement the backbone network that
was envisioned by ISA100.11a.

This operation of the Backbone Router requires that some protocol
operates over the LLNs from which node registrations can be obtained,
and that can disseminate the location of the backbone Router over the
LLN. Further expectations will be detailed.

The way the PAN IDs and 16-bit short addresses are allocated and

distributed in the case of an 802.15.4 network is not covered by this
specification. Similarly, the aspects of joining and securing the
network are out of scope. The way the nodes in the LLN learn about
their Backbone Router depends on the protocol used in the LLN. In
the case of RPL, a Border Router is the root of the DODAG that it
serves and represents all nodes attached to that DODAG.

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].

Readers are expected to be familiar with all the terms and concepts
that are discussed in "Neighbor Discovery for IP version 6"
[RFC4861], "IPv6 Stateless Address Autoconfiguration" [RFC4862],
"IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs):
Overview, Assumptions, Problem Statement, and Goals" [RFC4919],
Neighbor Discovery Optimization for Low-power and Lossy Networks
[RFC6775] and "Transmission of IPv6 Packets over IEEE 802.15.4
Networks" [RFC4944].

Readers would benefit from reading "Mobility Support in IPv6"
[RFC3775], "Neighbor Discovery Proxies (ND Proxy)" [RFC4389] and
"Optimistic Duplicate Address Detection" [RFC4429] prior to this
specification for a clear understanding of the art in ND-proxying and
binding.

Additionally, this document uses terminology from
[I-D.ietf-roll-terminology], and introduces the following
terminology:

Backbone   This is an IPv6 Backbone link that interconnects 2 or more
      Backbone Routers. It is expected to be deployed as a high
      speed backbone in order to federate a potentially large set of
      LLNS. Also referred to as a LLN backbone or Backbone network.

Backbone Router   An IPv6 router that federates the LLN using a
      Backbone link as a backbone. A BBR acts as a 6LoWPAN Border
      Routers (6LBR) and an Energy Aware Default Router (NEAR).

Extended LLN   This is the aggregation of multiple LLNs as defined in
      [RFC4919] interconnected by a Backbone Link via Backbone
      Routers and forming a single IPv6 link.

Energy-Constrained Node   An IPv6 node that operates ND registration
      in order to save energy. This can be a LLN node, or an
      Efficiency-Aware Host(EAH) on the backbone.

Binding   The association of the Energy-Constrained Node IPv6 address
      and Interface ID with associated proxying states including the
      remaining lifetime of that association.

Registration The process during which a Energy-Constrained Node
     injects its address in a protocol through which the Border
     Router can learn the address and proxy ND for it.

Primary BBR

     The BBR that will defend a registered address for the purpose
     of DAD over the backbone

Secondary BBR

     A BBR to which the address is registered. A Secondary Router
     MAY advertise the address over the backbone and proxy for it.

3. Overview

   The scope of this draft is a Backbone Link that federates multiple
   LLNs as a single IPv6 subnet. Each LLN in the subnet is anchored at
   a Backbone Router (BBR). The Backbone Routers interconnect the LLNs
   over the Backbone Link and emulate that the LLN nodes are present on
   the Backbone by proxy-ND operations. An LLN node can move freely
   from an LLN anchored at a Backbone Router to an LLN anchored at
   another Backbone Router on the same backbone and conserve any of the
   IPv6 addresses that it has formed. In a same fashion, an Efficiency-
   Aware Host (EAH) residing on the Backbone may change its BBR - acting
   as IPv6 ND-efficiency-aware Router (NEAR) - and conserve its
   addresses with no disruption.

   Energy-Constrained Nodes are often associated with radios and
   therefore may change their point of attachment in the network.
   Virtual devices - typically virtual machines in a datacenter - also
   move though in a different fashion, from a physical device to the
   next. In the case if a movement, it might be difficult for Stateful
   devices in the network such as the NEAR and SAVI switches to
   differentiate a duplication from a movement. And if indeed it is a
   movement, then it might be difficult to select to freshest
   information to know where the device actually is.


```
              ---+------------------------
                 |            Plant Network
                 |
            +-----+
            |     | Gateway
            |     |
            +-----+
               |
               |        Backbone Link
        +--------------------+------------------+
        |                    |                  |
      +-----+             +-----+            +-----+
      |     | Backbone    |     | Backbone   |     | Backbone
      |     | router      |     | router     |     | router
      +-----+             +-----+            +-----+
         o                 o   o   o             o o
      o o    o  o        o o   o  o  o       o  o  o  o o
     o  o o  o o          o   o  o  o  o       o  o  o o o
     o   o  o  o            o     o  o          o  o   o
      o    o o              o   o                 o o

        LLN                 LLN                LLN
```
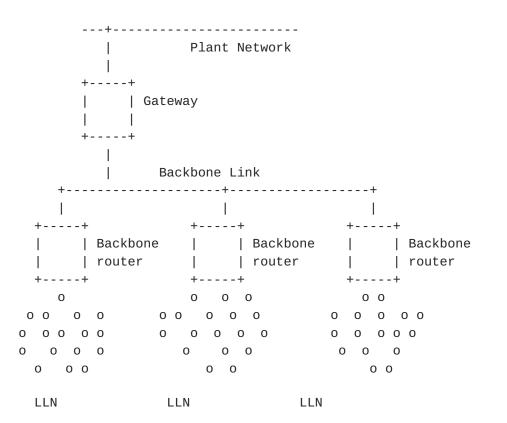
                  Figure 1: Backbone Link and Backbone Routers

   The Backbone Link is used as reference for Neighbor Discovery
   operations, by extending the concept of a Home Link as defined in
   [RFC3775] for Mobile IPv6. In particular, Backbone Routers perform
   ND proxying for the Energy-Constrained Nodes in the LLNs they own
   through a node registration.

   The Backbone Router operation is compatible with that of a Home
   Agent. This enables mobility support for LLN devices that would move
   outside of the network delimited by the Backbone link. This also
   enables collocation of Home Agent functionality within Backbone
   Router functionality on the same interface of a router.

   A Energy-Constrained Node registers and claims ownership of its
   addresse(s) using proactive acknowledged registration exchanges with
   a neighboring router. In case of a complex LLN topology, the router
   might be an intermediate LLN Router that relays the registration to
   the BBR (acting as LBR) as described in [RFC6775]. In turn, the
   Backbone Routers operate as a distributed database of all the Energy-
   Constrained Nodes whether theyreside on the LLNs or the backbone, and
   use the Neighbor Discovery Protocol to share that information across
   the Backbone in the classical ND reactive fashion.

   For the purpose of Neighbor Discovery proxying, this specification
   documents the LLN Master Neighbor Registry, a conceptual data
   structure that is similar to the MIP6 binding cache. The Master
   Neighbor Registry is fed by redistributing addresses learnt from the
   registration protocol used over the LLN.

   Another function of the Backbone Router is to perform 6LoWPAN
   compression and expansion between the LLN and the Backbone Link and
   ensure MTU compatibility. Packets flow uncompressed over the
   Backbone Link and are routed normally towards a Gateway or an
   Application sitting on the Backbone link or on a different link that
   is reachable over the IP network.

[4](#). **New types and formats**

   The specification expects that the protocol running on the LLN can
   provide a sequence number called Transaction ID (TID) that is
   associated to the registration. When a node registers to multiple
   BBRs, it is expected that the same TID is used, to enable the BBR to
   correlate the registrations as being a single one, and differentiate
   that situation from a movement. Otherwise, the resolution makes it
   so that only the most recent registration was perceived from the
   highest TID is kept.

   The specification expects that the protocol running on the LLN can
   provide a unique ID for the owner of the address that is being
   registered. The Owner Unique ID enables to differentiate a duplicate
   registration from a double registration. In case of a duplicate, the
   last registration looses. The Owner Unique ID can be as simple as a
   EUI-64 burnin address, if the device manufacturor is convinced that
   there can not be a manuf error that would cause duplicate EUI64
   addresses. Alternatively, the unique ID can be a hash of supposedly
   unique information from multiple orthogonal sources, for instance:

   o Burn in address.

   o configured address, id, security keys...

   o (pseudo) Random number, radio link metrics ...

   In any fashion, it is recommended that the device stores the unique
   Id in persistent memory. Otherwise, it will be prevented to
   reregister after a reboot that would cause a loss of memory until the
   Backbone Router times out the registration.

   The unique ID and the sequence number are placed in a new ND option
   that is used by the Backbone Routers over the Backbone link to detect
   duplicates and movements. The option format is as follows:

[4.1](#). **The Enhanced Address Registration Option (EARO)**

   This option is designed to be used with standard NS and NA messages
   between backbone Routers over a backbone link and may be used between
   LRs and LBRs over the LLN. By using this option, the binding in
   question can be uniquely identified and matched with the Master
   Neighbor Registry entries of each Backbone Router.

```
    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |    Type      |  Length = 2   |    Status     |   Reserved    |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |T|N| Reserved |      TID      |      Registration Lifetime    |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                                                               |
   +          OUI   ( EUI-64 or equivalent )                      +
   |                                                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

                         Figure 2: EARO

   Option Fields

   Type:

   Length: 2

   T: One bit flag. Set if the next octet is a used as a TID.

   N: One bit flag. et if the device moved. If not set, the router will
      refrain from sending gratuitous NA(O) over the backbone, for
      instance after the DAD operation upon entry creation.

   Reserved: This field is unused. It MUST be initialized to zero by
      the sender and MUST be ignored by the receiver.

   TID: 1-byte integer; a transaction id that is maintained by the
      device and incremented with each transaction. it is recommended
      that the device maintains the TID in a persistent storage.

   Owner Unique Identifier: A globally unique identifier for the host's
      interface associated with the binding for the NS/NA message in
      question. This can be the EUI-64 derived IID of an interface,
      which can be hashed with other supposedly unique information from
      multiple orthogonal sources.

5. Backbone Router Operations

5.1. Backbone Link and Router

   The Backbone Router is a specific kind of Border Router that performs
   proxy Neighbor Discovery on its backbone interface on behalf of the
   nodes that it has discovered on its Low Power Lossy Network
   interfaces. On the LLN side, the Backbone Router acquires its states
   about the nodes by terminating protocols such as RPL [RFC6550] or
   6LoWPAN ND [RFC6775] as a LLN Border Router. It is expected that the
   backbone is the medium used to connect the subnet to the rest of the
   infrastructure, and that all the LBRs are connected to that backbone
   and support the Backbone Router feature as specified in this
   document.

   The backbone is expected to be a high speed, reliable Backbone link,
   with affordable multicast capabilities, such as an Ethernet Network
   or a fully meshed NBMA network with multicast emulation, which allows
   a full support of classical ND as specified in [RFC4861] and
   subsequent RFCs. In other words, the backbone is not a LLN. Still,
   some restrictions of the attached LLNs will apply to the backbone.
   In particular, it is expected that the MTU is set to the same value
   on the backbone and all attached LLNs.

5.2. ND Proxy Operations

   This specification enables a Backbone Router to proxy Neighbor
   Discovery operations over the backbone on behalf of the nodes that
   are registered to it, allowing any device on the backbone to reach a
   Energy-Constrained Node as if it was on-link.

   In the context of this specification, proxy ND means:

   o defending a registered address over the backbone using NA messages
      with the Override bit set

   o advertising a registered address over the backbone using NA
      messages, asynchronously or as a response to a Neighbor
      Solicitation messages.

   o Looking up a destination over the backbone in order to deliver
      packets arriving from the LLN using Neighbor Sollicitation
      messages.

   o Forwarding packets from the LLN over the backkbone, and hte other
      way around.

   o Eventually triggering a look up for a destination over the LLN
      that would not be registered at a given point of time, or as a
      verification of a registration.

   The draft introduces the concept of primary and secondary BBRs. The
   concept is defined with the granularity of an address, that is a
   given BBR can be primary for a given address and secondary or another
   one, regardess on whether the addresses belong to the same node or
   not. The primary Backbone Router is in charge of protecting the
   address for DAD over the Backbone. Any of the Primary and Secondary
   BBR may claim the address over the backbone, since they are all
   capable to route from the backbone to the LLN device.

   When the protocol used to register the nodes over the LLN is RPL
   [RFC6550], it is expected that one BBR acts as virtual root
   coordinating LLN BBRs (with the same DODAGID) over the non-LLN
   backbone. In that case, the virtual root may act as primary BBR for
   all addresses that it cares to support, whereas the physical roots to
   which the node is attached are secondary BBRs. It is also possible
   in a given deployment that the DODAGs are not coordinated. In that
   case, there is no virtual root and no secondary BBR; the DODAG root
   is primary all the nodes registered to it over the backbone.

   When the protocol used to register the nodes over the LLN is 6LoWPAN
   ND [RFC6775], the Backbone Routers act as a distributed DAD table,
   using classical ND over the backbone to detect duplication. This
   specification requires that:

   1. Registrations for all addresses that can be required to reach the
       device over the backbone, including registrations for IPv6
       addresses based on burn-in EUI64 addresses are passed to the DAD
       table.

   2. Nodes include the EARO in their NS used for registering those
       addresses and the LRs propagate that option to the LBRs.

   A false positive duplicate detection may arise over the backbone, for
   instance if the node registers to more than one LBR, or if the node
   has moved. Both situations are handled gracefully unbeknownst to the
   node. In the former case, one LBR becomes primary to defend the
   address over the backbone while the others become secondary and may
   still forward packets back and forth. In the latter case the LBR
   that receives the newest registration wins and becomes primary.

5.3. Claiming and Defending Addresses

   Upon a new or an updated registration, the BBR performs a DAD
   operation. If either a TID or a OUI is available, the BBR places

them in a EARO in all its ND messages over the backbone. If content
is not available for a given field, it is set to 0.

If a primary already exists over the backbone, it will answer the DAD
with an RA.

o If a RA is received with the O bit set, the primary rejects the
   DAD and the DAD fails. the BBR needs to inform the LLN protocol
   that the address is a duplicate.

o If a RA is received with the O bit reset, the primary accepts the
   BBR as secondary and DAD succeeds. The BBR may install or
   maintain its proxy states for that address. This router MAY
   advertise the address using a NA. during a registration flow, it
   MAY set the O bit.

o If no RA is received, this router assumes the role of primary and
   DAD succeeds. The BBR may install or maintain its proxy states
   for that address. This router advertises the address using a NA
   with the O bit reset.

When the BBR installs or maintains its proxy states for an address,
it sends an NA with a SLLA option for that address. The Primary BR
MAY set the O bit if it wished to attract the traffic for that
address.

5.4. Conflict Resolution

A conflict arise when multiple BBRs get a registration from a same
address. This situation might arise when a node moves from a BBR to
another, when a node registers to multiple BBRs, or in the RPL case
when the BRs belong to a single coordinated DODAG.

The primary looks up the EARO in ND messages with a SLLA option.

o If there is no option, normal ND operation takes place and the
   primary defends the address with a NA with the O bit set, adding
   the EARO with its own information.

o If there is a EARO and the OUIs are different, then the conflict
   apparently happens between different nodes, and the the primary
   defends the address with a NA with the O bit set, adding the EARO
   with its own information. If the TID in the EARO is in the
   straight part of the lollipop, it is possible that the request
   comes from the same node that has rebooted and formed a new OUI
   The primary BBR may assess its registered entry prior to
   answering.

   o If there is a EARO and the OUIs are the same:

      * If the TID in the ND message is newer than the most recent one
         known by the primary router, this is interpreted as a node
         moving. The primary cleans up its states and stops defending
         the address.

      * If the TID in the ND message is the same as the most recent one
         known by the primary router, this is interpreted as a double
         registration. In case of a DAD, the promary responds with a NA
         with the O bit reset, to confirm its position as primary,
         including the EARO.

      * If the TID in the ND message is older than the most recent one
         known by the primary router, this is interpreted as a stale
         information. The primary defends the address with a NA with
         the O bit set, adding the EARO with its own information.

      * If the TIDs are very different (more than 16 apart, discounting
         the straight part of the lollipop), it is impossible to resolve
         for sure. The primary BBR should assess its registered entry
         prior to answering.

5.5. Assessing an entry

   In a number of cases, it might happen that the information at the
   primary BBR is stale and obsolete. In particular, a node with no
   permanent storage might reboot and form a different OUI, in which
   case the information at the BBR is inaccurate and should be removed.
   In another case, the BBR and the node have been out of reach for too
   long and the TID that the BBR maintains is so far out that it is
   impossible to compare it with that stored at the BBR.

   In such situation, the primary Backbone Router has the possibility to
   assess the registration. this is performed by the protocol in use to
   register the node over the LLN.

   When the protocol used to register the nodes over the LLN is RPL
   [RFC6550], the BBR sends a targetted DIS to the registered address
   over the registered path. A DAO back indicates that the current
   registration is still valid and provides the adequate data to resolve
   the conflict.

   When the protocol used to register the nodes over the LLN is 6LoWPAN
   ND [RFC6775].

6. Security Considerations

   This specification expects that the link layer is sufficiently
   protected, either by means of physical or IP security for the
   Backbone Link or MAC sublayer cryptography. In particular, it is
   expected that the LLN MAC provides secure unicast to/from the
   Backbone Router and secure BBRoadcast from the Backbone Router in a
   way that prevents tempering with or replaying the RA messages.

   The use of EUI-64 for forming the Interface ID in the link local
   address prevents the usage of Secure ND ([RFC3971] and [RFC3972]) and
   address privacy techniques. Considering the envisioned deployments
   and the MAC layer security applied, this is not considered an issue
   at this time.

7. IANA Considerations

   A new type is requested for an ND option.

8. Acknowledgments

   TBD

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.

   [RFC2460]   Deering, S. and R. Hinden, "Internet Protocol, Version 6
               (IPv6) Specification", RFC 2460, December 1998.

   [RFC3775]   Johnson, D., Perkins, C., and J. Arkko, "Mobility Support
               in IPv6", RFC 3775, June 2004.

   [RFC4291]   Hinden, R. and S. Deering, "IP Version 6 Addressing
               Architecture", RFC 4291, February 2006.

   [RFC4429]   Moore, N., "Optimistic Duplicate Address Detection (DAD)
               for IPv6", RFC 4429, April 2006.

   [RFC4443]   Conta, A., Deering, S., and M. Gupta, "Internet Control
               Message Protocol (ICMPv6) for the Internet Protocol
               Version 6 (IPv6) Specification", RFC 4443, March 2006.

   [RFC4861]   Narten, T., Nordmark, E., Simpson, W., and H. Soliman,
               "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861,
               September 2007.

   [RFC4862]   Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless
               Address Autoconfiguration", RFC 4862, September 2007.

   [RFC4944]   Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler,
               "Transmission of IPv6 Packets over IEEE 802.15.4
               Networks", RFC 4944, September 2007.

   [RFC6282]   Hui, J. and P. Thubert, "Compression Format for IPv6
               Datagrams over IEEE 802.15.4-Based Networks", RFC 6282,
               September 2011.

   [RFC6550]   Winter, T., Thubert, P., Brandt, A., Hui, J., Kelsey, R.,
               Levis, P., Pister, K., Struik, R., Vasseur, JP., and R.
               Alexander, "RPL: IPv6 Routing Protocol for Low-Power and
               Lossy Networks", RFC 6550, March 2012.

   [RFC6775]   Shelby, Z., Chakrabarti, S., Nordmark, E., and C. Bormann,
               "Neighbor Discovery Optimization for IPv6 over Low-Power
               Wireless Personal Area Networks (6LoWPANs)", RFC 6775,
               November 2012.

9.2. Informative References

   [I-D.chakrabarti-nordmark-6man-efficient-nd]
              Chakrabarti, S., Nordmark, E., and M. Wasserman,
              "Efficiency aware IPv6 Neighbor Discovery Optimizations",
              draft-chakrabarti-nordmark-6man-efficient-nd-01 (work in
              progress), November 2012.

   [I-D.ietf-roll-terminology]
              Vasseur, J., "Terminology in Low power And Lossy
              Networks", draft-ietf-roll-terminology-11 (work in
              progress), February 2013.

   [I-D.van-beijnum-multi-mtu]
              Beijnum, I., "Extensions for Multi-MTU Subnets",
              draft-van-beijnum-multi-mtu-03 (work in progress),
              July 2010.

   [RFC3963] Devarapalli, V., Wakikawa, R., Petrescu, A., and P.
              Thubert, "Network Mobility (NEMO) Basic Support Protocol",
              RFC 3963, January 2005.

   [RFC3971] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure
              Neighbor Discovery (SEND)", RFC 3971, March 2005.

   [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)",
              RFC 3972, March 2005.

   [RFC4389] Thaler, D., Talwar, M., and C. Patel, "Neighbor Discovery
              Proxies (ND Proxy)", RFC 4389, April 2006.

   [RFC4919] Kushalnagar, N., Montenegro, G., and C. Schumacher, "IPv6
              over Low-Power Wireless Personal Area Networks (6LoWPANs):
              Overview, Assumptions, Problem Statement, and Goals",
              RFC 4919, August 2007.

9.3. External Informative References

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

   [ISA100.11a]
              ISA, "ISA100, Wireless Systems for Automation", May 2008,
              <http://www.isa.org/Community/
              SP100WirelessSystemsforAutomation>.

Author's Address

   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