

6lo Working Group
Internet-Draft
Intended status: Informational
Expires: October 6, 2016

R. Droms
P. Duffy
Cisco
April 4, 2016

Assignment of an Ethertype for IPv6 with LoWPAN Encapsulation
draft-droms-6lo-ethertype-request-01

Abstract

When carried over layer 2 technologies such as Ethernet, IPv6 datagrams using LoWPAN encapsulation as defined in RFC 4944 must be identified so the receiver can correctly interpret the encoded IPv6 datagram. This document requests the assignment of an Ethertype for that purpose.

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 October 6, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

1. Introduction

The IETF has defined a format for IPv6 [RFC2460] datagram encapsulation [RFC4944] ("LoWPAN encapsulation"). This document regards any IPv6 datagram using the Dispatch octet as defined in section 5.1 of RFC 4944 to be using LoWPAN encapsulation. LoWPAN encapsulation as defined in RFC 4944 has been updated by [RFC6282], and may be extended and modified by future IETF standards document. The intended layer 2 technology for IPv6 datagrams using LoWPAN encapsulation as originally defined is [IEEE.802.15.4_2011], which does not provide for a protocol switch in its layer 2 headers.

There is interest in carrying IPv6 datagrams over layer 2 technologies that do include a protocol switch field:

- o Usage of LoWPAN encapsulation in conjunction with IEEE 802.15.9 Multiplexed Data Service [IEEE802159], which provides the ability to perform upper layer protocol dispatch for IEEE 802.15.4 networks. Wi-SUN Alliance intends to use the 15.9 Multiplexed Data Information Element to dispatch LoWPAN encapsulation frames to upper stack layers. As specified in IEEE 802.15.9, dispatch of LoWPAN encapsulation frames will require an Ethertype be assigned for LoWPAN encapsulation.
- o LoWPAN encapsulation will likely be needed for WiFi Alliance's HaLow [HALOW] standard (low power operation in the 900 MHz band)
- o Other layer 2 technologies such as Ethernet and debugging tools such as Wireshark require a unique protocol type field for LoWPAN encapsulation to properly interpret IPv6 datagrams that use LoWPAN encapsulation.

2. Request to IEEE for assignment of an Ethertype

When this document is published, the IETF will formally submit a request to IEEE for assignment of an Ethertype for IPv6 datagrams using LoWPAN encapsulation.

3. IANA Considerations

This memo includes no request to IANA.

4. Security Considerations

This document is intended only to request assignment of an Ethertype for IPv6 datagrams using LoWPAN encapsulation. It has no incremental implications for security beyond those in the relevant protocols.

5. Normative References

- [HALLOW] Wi-Fi Alliance, "Wi-Fi HaLow",
<http://www.wi-fi.org/discover-wi-fi/wi-fi-halow> .
- [IEEE.802.15.4_2011]
IEEE, "IEEE Standard for Local and metropolitan area networks--Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs)", IEEE 802.15.4-2011, DOI 10.1109/ieeestd.2011.6012487, September 2011, <<http://ieeexplore.ieee.org/servlet/opac?punumber=6012485>>.
- [IEEE802159]
IEEE, "IEEE Draft Recommended Practice for Transport of Key Management Protocol (KMP) Datagrams", IEEE P802.15.9/D04, May 2015.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.

Authors' Addresses

Ralph Droms
Cisco
55 Cambridge Parkway
Cambridge, Massachusetts
US

Phone: +1 617 621 1904
Email: rdroms.ietf@gmail.com

Paul Duffy
Cisco
1414 Massachusetts Ave.
Boxborough, Massachusetts 01719
US

Phone: +1 978 204 9993
Email: paduffy@cisco.com

6Lo Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 22, 2016

Y-G. Hong
ETRI
C. Gomez
UPC/i2cat
Y-H. Choi
ETRI
D-Y. Ko
SKtelecom
March 21, 2016

Use cases for IPv6 over Networks of Resource-constrained Nodes
draft-hong-6lo-use-cases-01

Abstract

This document describes the characteristics of link layer technologies that are used at constrained node networks and typical use cases of IPv6 over networks of resource-constrained nodes. In addition to IEEE 802.15.4, various link layer technologies such as BLE, ITU-T G.9959 (Z-Wave), DECT-ULE, MS/TP, NFC, and LTE MTC are widely used at constrained node networks for typical services. Based on these link layer technologies, IPv6 over networks of resource-constrained nodes has various and practical use cases. To efficiently implement typical services, the applicability and consideration of several design spaces are described.

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 22, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Conventions and Terminology	4
3. 6lo Link layer technologies	4
3.1. ITU-T G.9959	4
3.2. Bluetooth Low Energy	4
3.3. DECT-ULE	4
3.4. Master-Slave/Token-Passing	5
3.5. NFC	6
3.6. LTE MTC	6
4. Design Space	7
5. 6lo Use Cases	8
5.1. Use case of NFC: Alternative Secure Transfer	8
5.2. Use case of ITU-T G.9959: Smart Home	10
5.3. Use case of Bluetooth Low Energy: Smartphone-Based Interaction with Constrained Devices	12
5.4. Use case of DECT-ULE: Smart Home	13
5.5. Use case of LTE MTC	14
6. IANA Considerations	16
7. Security Considerations	16
8. Acknowledgements	16
9. References	16
9.1. Normative References	17
9.2. Informative References	18
Authors' Addresses	19

1. Introduction

Running IPv6 on constrained node networks has different features due to the characteristics of constrained node networks such as small packet size, short link-layer address, low bandwidth, network topology, low power, low cost, and large number of devices [RFC4919].

For example, because some IEEE 802.15.4 link layers have a frame size of 127 octets and IPv6 requires an MTU of 1280 bytes, an appropriate fragmentation and reassembly adaptation layer must be provided at the layer of below IPv6. Also, the limited size of IEEE 802.15.4 frame, the length shortage of data delivery, and low energy consumption requirements make the need for header compression. IETF 6lowpan (IPv6 over Low powerWPAN) working group published [RFC4944], an adaptation layer for sending IPv6 packets over IEEE 802.15.4, [RFC6282], compression format for IPv6 datagrams over IEEE 802.15.4-based networks, and [RFC6775], Neighbor Discovery Optimization for 6lowpan.

As IoT (Internet of Things) services become more popular, various link layer technologies such as Bluetooth Low Energy (Bluetooth LE), ITU-T G.9959 (Z-Wave), Digital Enhanced Cordless Telecommunications - Ultra Low Energy (DECT-ULE), Master-Slave/Token Passing (MS/TP), and Near Field Communication (NFC) are actively used. And the need of transmission of IPv6 packets over these link layer technologies is required. A number of IPv6-over-foo documents have been developed in the IETF 6lo (IPv6 over Networks of Resource-constrained Nodes) and 6tisch (IPv6 over the TSCH mode of IEEE 802.15.4e) working group.

In the 6lowpan working group, the [RFC6568], "Design and Application Spaces for 6LoWPANs" was published and it describes potential application scenarios and use cases for low-power wireless personal area networks. In this document, various design space dimension such as deployment, network size, power source, connectivity, multi-hop communication, traffic pattern, security level, mobility, and QoS were analyzed. And it described a fundamental set of 6lowpan application scenarios and use cases: Industrial monitoring-Hospital storage rooms, Structural monitoring-Bridge safety monitoring, Connected home-Home Automation, Healthcare-Healthcare at home by tele-assistance, Vehicle telematics-telematics, and Agricultural monitoring-Automated vineyard.

Even though the [RFC6568] describes some potential application scenarios and use cases and it lists the design space in the context of 6lowpan, it does not cover the different use cases and design space in the context of the 6lo working group. This document provides the use cases of 6lo, considering the following:

- o 6lo use cases MAY be uniquely different to the 6lowpan use cases.
- o 6lo use cases SHOULD cover various IoT related wire/wireless link layer technologies providing practical information of such technologies.

- o 6lo use cases MAY describe characteristics and typical use cases of each link layer technology, and then 6lo use cases's applicability.

2. Conventions and 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. 6lo Link layer technologies

3.1. ITU-T G.9959

The ITU-T G.9959 recommendation [G.9959] targets low-power Personal Area Networks (PANs). G.9959 defines how a unique 32-bit HomeID network identifier is assigned by a network controller and how an 8-bit NodeID host identifier is allocated to each node. NodeIDs are unique within the network identified by the HomeID. The G.9959 HomeID represents an IPv6 subnet that is identified by one or more IPv6 prefixes [RFC7428].

3.2. Bluetooth Low Energy

Bluetooth LE was introduced in Bluetooth 4.0, enhanced in Bluetooth 4.1, and developed even further in successive versions. Bluetooth SIG has also published Internet Protocol Support Profile (IPSP), which includes Internet Protocol Support Service (IPSS). The IPSP enables discovery of IP-enabled devices and establishment of link-layer connection for transporting IPv6 packets. IPv6 over Bluetooth LE is dependent on both Bluetooth 4.1 and IPSP 1.0 or newer.

Devices such as mobile phones, notebooks, tablets and other handheld computing devices which will include Bluetooth 4.1 chipsets will also have the low-energy variant of Bluetooth. Bluetooth LE will also be included in many different types of accessories that collaborate with mobile devices such as phones, tablets and notebook computers. An example of a use case for a Bluetooth LE accessory is a heart rate monitor that sends data via the mobile phone to a server on the Internet [RFC7668].

3.3. DECT-ULE

DECT ULE is a low power air interface technology that is designed to support both circuit switched services, such as voice communication, and packet mode data services at modest data rate.

The DECT ULE protocol stack consists of the PHY layer operating at frequencies in the 1880 - 1920 MHz frequency band depending on the region and uses a symbol rate of 1.152 Mbps. Radio bearers are allocated by use of FDMA/TDMA/TDD technics.

In its generic network topology, DECT is defined as a cellular network technology. However, the most common configuration is a star network with a single Fixed Parts (FP) defining the network with a number of PP attached. The MAC layer supports traditional DECT as this is used for services like discovery, pairing, security features etc. All these features have been reused from DECT.

The DECT ULE device can switch to the ULE mode of operation, utilizing the new ULE MAC layer features. The DECT ULE Data Link Control (DLC) provides multiplexing as well as segmentation and re-assembly for larger packets from layers above. The DECT ULE layer also implements per-message authentication and encryption. The DLC layer ensures packet integrity and preserves packet order, but delivery is based on best effort.

The current DECT ULE MAC layer standard supports low bandwidth data broadcast. However the usage of this broadcast service has not yet been standardized for higher layers [I-D.ietf-6lo-dect-ule].

3.4. Master-Slave/Token-Passing

MS/TP is a contention-free access method for the RS-485 physical layer, which is used extensively in building automation networks. This specification defines the frame format for transmission of IPv6 [RFC2460] packets and the method of forming link-local and statelessly autoconfigured IPv6 addresses on MS/TP networks. The general approach is to adapt elements of the 6LoWPAN [RFC4944] specification to constrained wired networks.

An MS/TP device is typically based on a low-cost microcontroller with limited processing power and memory. Together with low data rates and a small address space, these constraints are similar to those faced in 6LoWPAN networks and suggest some elements of that solution might be leveraged. MS/TP differs significantly from 6LoWPAN in at least three respects: a) MS/TP devices typically have a continuous source of power, b) all MS/TP devices on a segment can communicate directly so there are no hidden node or mesh routing issues, and c) recent changes to MS/TP provide support for large payloads, eliminating the need for link-layer fragmentation and reassembly.

MS/TP is designed to enable multidrop networks over shielded twisted pair wiring. It can support a data rate of 115,200 baud on segments up to 1000 meters in length, or segments up to 1200 meters in length

at lower baud rates. An MS/TP link requires only a UART, an RS-485 transceiver with a driver that can be disabled, and a 5ms resolution timer. These features make MS/TP a cost-effective field bus for the most numerous and least expensive devices in a building automation network [I-D.ietf-6lo-6lobac].

3.5. NFC

NFC technology enables simple and safe two-way interactions between electronic devices, allowing consumers to perform contactless transactions, access digital content, and connect electronic devices with a single touch. NFC complements many popular consumer level wireless technologies, by utilizing the key elements in existing standards for contactless card technology (ISO/IEC 14443 A&B and JIS-X 6319-4). NFC can be compatible with existing contactless card infrastructure and it enables a consumer to utilize one device across different systems.

Extending the capability of contactless card technology, NFC also enables devices to share information at a distance that is less than 10 cm with a maximum communication speed of 424 kbps. Users can share business cards, make transactions, access information from a smart poster or provide credentials for access control systems with a simple touch.

NFC's bidirectional communication ability is ideal for establishing connections with other technologies by the simplicity of touch. In addition to the easy connection and quick transactions, simple data sharing is also available [I-D.ietf-6lo-nfc].

3.6. LTE MTC

LTE category defines the overall performance and capabilities of the UE(User Equipment). For example, the maximum down rate of category 1 UE and category 2 UE are 10.3 Mbit/s and 51.0 Mbit/s respectively. There are many categories in LTE standard. 3GPP standards defined the category 0 to be used for low rate IoT service in release 12. Since category 1 and category 0 could be used for low rate IoT service, we call LTE MTC[LTE_MTC].

LTE MTC have the advantages compared to above category 2 to be used for low rate IoT service such as low power and low cost.

The below figure shows the primary characteristics of LTE MTC.

Category	Max. Date Rate Down	Max. Date Rate Up
Category 0	1.0 Mbit/s	1.0 Mbit/s
Category 1	10.3 Mbit/s	5.2 Mbit/s

Table 1: Primary characteristics of LTE MTC

4. Design Space

The [RFC6568] lists the dimensions used to describe the design space of wireless sensor networks in the context of the 6LoWPAN working group. The design space is already limited by the unique characteristics of a LoWPAN (e.g., low power, short range, low bit rate). In the RFC 6558, the following design space dimensions are described; Deployment, Network size, Power source, Connectivity, Multi-hop communication, Traffic pattern, Mobility, Quality of Service (QoS).

The design space dimensions of 6lo are a little different to those of the RFC 6558 due to the different characteristics of 6lo link layer technologies. The following design space dimensions can be considered.

- o Deployment/Bootstrapping: 6lo nodes can be connected randomly, or in an organized manner. The bootstrapping has different characteristics of each link layer technologies.
- o Topology: Topology of 6lo networks may inherently follow the characteristics of each link layer technology. Point-to-point, star, tree or mesh topologies can be configured.
- o L2-Mesh or L3-Mesh: L2-mesh and L3-mesh may inherently follow the characteristics of each link layer technologies. Some link layer technologies may support L2-mesh and some may not support.
- o Multi-link subnet, single subnet: The selection of multi-link subnet and single subnet depends on connectivity and the number of 6lo nodes.
- o Data rate: Originally, the link layer technologies of 6lo have low rate of data transmission. But, by adjusting the MTU, it can deliver higher data rate.

- o Buffering requirements: Some 6lo use case may require more data rate than the link layer technology support. In this case, a buffering mechanism to manage the data is required.
- o Security Requirements: Some 6lo use case can transfer some important and personal data between 6lo nodes. In this case, high-level security support is required.
- o Mobility across 6lo networks and subnets: The movement of 6lo nodes is dependent on the 6lo use case. If the 6lo nodes can move or moved around, it requires the mobility management mechanism.
- o Time synchronization requirements: The requirement of time synchronization of the upper layer service is dependent on the 6lo use case. For some 6lo use case related to health service, the measured data must be recorded with exact time and must be transferred with time synchronization.
- o Reliability and QoS: Some 6lo use case requires high reliability, for example real-time service or health-related services.
- o Traffic patterns: 6lo use case may various traffic patterns. Some 6lo use case may require short data length and randomly. Some 6lo use case may require continuous data and periodic data transmission.
- o Security Bootstrapping: Without the external operations, 6lo nodes must have the security bootstrapping mechanism.
- o Power use strategy: to enable certain use cases, there may be requirements on the class of energy availability and the strategy followed for using power for communication [RFC7228]. Each link layer technology defines a particular power use strategy which may be tuned [I-D.ietf-lwig-energy-efficient].
- o Energy limitation: The energy limitation class [RFC7228] is specific to each use case, and may or may not be related to the power use strategy.

5. 6lo Use Cases

5.1. Use case of NFC: Alternative Secure Transfer

According to applications, various secured data can be handled and transferred. Depending on security level of the data, methods for transfer can be alternatively selected. The personal data having serious issues should be transferred securely, but data transfer by using Wi-Fi and Bluetooth connections cannot always be secure because

of their a little long radio frequency range. Hackers can overhear the personal data transfer behind hidden areas. Therefore, methods need to be alternatively selected to transfer secured data. Voice and video data, which are not respectively secure and requires long transmission range, can be transferred by 3G/4G technologies, such as WCDMA, GSM, and LTE. Big size data, which are not secure and requires high speed and broad bandwidth, can be transferred by Wi-Fi and wired network technologies. However, the person data, which are serious issues so requires secure transfer in wireless area, can be securely transferred by NFC technology. It has very short frequency range ? nearly single touch communication.

Example: Secure Transfer by Using NFC in Healthcare Services with Tele-Assistance

A senior citizen who lives alone wears one to several wearable 6lo devices to measure heartbeat, pulse rate, etc. The 6lo devices are densely installed at home for movement detection. An LoWPAN Border Router (LBR) at home will send the sensed information to a connected healthcare center. Portable base stations with LCDs may be used to check the data at home, as well. Data is gathered in both periodic and event-driven fashion. In this application, event-driven data can be very time-critical. In addition, privacy also becomes a serious issue in this case, as the sensed data is very personal.

While the senior citizen is provided audio and video healthcare services by a tele-assistance based on LTE connections, the senior citizen can alternatively use NFC connections to transfer the personal sensed data to the tele-assistance. At this moment, hidden hackers can overhear the data based on the LTE connection, but they cannot gather the personal data over the NFC connection.

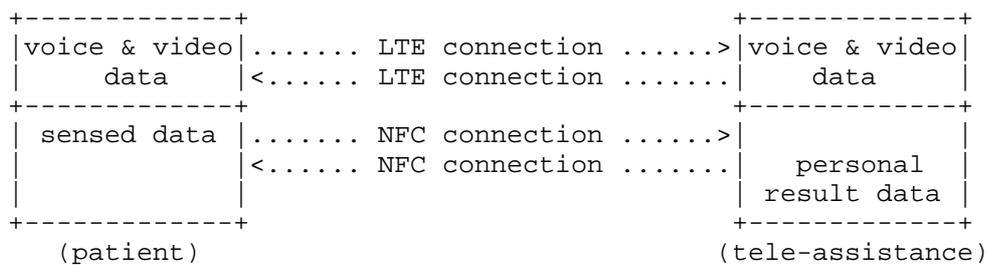


Figure 1: Alternative Secure Transfer in Healthcare Services

Dominant parameters in secure transfer by using NFC in healthcare services:

- o Deployment/Bootstrapping: Pre-planned. MP2P/P2MP (data collection), P2P (local diagnostic).
- o Topology: Small, NFC-enabled device connected to the Internet.
- o L2-mesh or L3-mesh: NFC does not support L2-mesh, L3-mesh can be configured.
- o Multi-link subnet, single subnet: a Single-hop for gateway; patient's body network is mesh topology.
- o Data rate: Small data rate.
- o Buffering requirements: Low requirement.
- o Security requirements: Data privacy and security must be provided. Encryption is required.
- o Mobility: Moderate (patient's mobility).
- o Time Synchronization: Highly required.
- o Reliability and QoS: High level of reliability support (life-or-death implication), role-based.
- o Traffic patterns: Short data length and periodic (randomly).
- o Security Bootstrapping: Highly required.
- o Other Issues: Plug-and-play configuration is required for mainly non-technical end-users. Real-time data acquisition and analysis are important. Efficient data management is needed for various devices that have different duty cycles, and for role-based data control. Reliability and robustness of the network are also essential.

5.2. Use case of ITU-T G.9959: Smart Home

Z-Wave is one of the main technologies that may be used to enable smart home applications. Born as a proprietary technology, Z-Wave was specifically designed for this use case. Recently, the Z-Wave radio interface (physical and MAC layers) has been standardized as the ITU-T G.9959 specification.

Example: Use of ITU-T G.9959 for Home Automation

Variety of home devices (e.g. light dimmers/switches, plugs, thermostats, blinds/curtains and remote controls) are augmented with

ITU-T G.9959 interfaces. A user may turn on/off or may control home appliances by pressing a wall switch or by pressing a button in a remote control. Scenes may be programmed, so that after a given event, the home devices adopt a specific configuration. Sensors may also periodically send measurements of several parameters (e.g. gas presence, light, temperature, humidity, etc.) which are collected at a sink device, or may generate commands for actuators (e.g. a smoke sensor may send an alarm message to a safety system).

The devices involved in the described scenario are nodes of a network that follows the mesh topology, which is suitable for path diversity to face indoor multipath propagation issues. The multihop paradigm allows end-to-end connectivity when direct range communication is not possible. Security support is required, specially for safety-related communication. When a user interaction (e.g. a button press) triggers a message that encapsulates a command, if the message is lost, the user may have to perform further interactions to achieve the desired effect (e.g. a light is turned off). A reaction to a user interaction will be perceived by the user as immediate as long as the reaction takes place after less than 0.5 seconds [RFC5826].

Dominant parameters in home automation scenarios with ITU-T G.9959:

- o Deployment/Bootstrapping: Pre-planned.
- o Topology: Mesh topology.
- o L2-mesh or L3-mesh: ITU-T G.9959 provides support for L2-mesh, and L3-mesh can also be used (the latter requires an IP-based routing protocol).
- o Multi-link subnet, single subnet: Multi-link subnet.
- o Data rate: Small data rate, infrequent transmissions.
- o Buffering requirements: Low requirement.
- o Security requirements: Data privacy and security must be provided. Encryption is required.
- o Mobility: Most devices are static. A few devices (e.g. remote control) are portable.
- o Time Synchronization: TBD.
- o Reliability and QoS: Moderate to high level of reliability support. Actions as a result of human-generated traffic should occur after less than 0.5 seconds.

- o Traffic patterns: Periodic (sensor readings) and aperiodic (user-triggered interaction).
- o Security Bootstrapping: Required.

5.3. Use case of Bluetooth Low Energy: Smartphone-Based Interaction with Constrained Devices

The key feature behind the current high Bluetooth LE momentum is its support in a large majority of smartphones in the market. Bluetooth LE can be used to allow the interaction between the smartphone and surrounding sensors or actuators. Furthermore, Bluetooth LE is also the main radio interface currently available in wearables. Since a smartphone typically has several radio interfaces that provide Internet access, such as Wi-Fi or 4G, the smartphone can act as a gateway for nearby devices such as sensors, actuators or wearables. Bluetooth LE may be used in several domains, including healthcare, sports/wellness and home automation.

Example: Bluetooth LE-based Body Area Network for fitness

A person wears a smartwatch for fitness purposes. The smartwatch has several sensors (e.g. heart rate, accelerometer, gyrometer, GPS, temperature, etc.), a display, and a Bluetooth LE radio interface. The smartwatch can show fitness-related statistics on its display. However, when a paired smartphone is in the range of the smartwatch, the latter can report almost real-time measurements of its sensors to the smartphone, which can forward the data to a cloud service on the Internet. In addition, the smartwatch can receive notifications (e.g. alarm signals) from the cloud service via the smartphone. On the other hand, the smartphone may locally generate messages for the smartwatch, such as e-mail reception or calendar notifications.

The functionality supported by the smartwatch may be complemented by other devices such as other on-body sensors, wireless headsets or head-mounted displays. All such devices may connect to the smartphone creating a star topology network whereby the smartphone is the central component.

Dominant parameters in home automation scenarios with Bluetooth LE:

- o Deployment/Bootstrapping: Pre-planned.
- o Topology: Star topology.
- o L2-mesh or L3-mesh: No.
- o Multi-link subnet, single subnet: Multi-link subnet.

- o Data rate: TBD.
- o Buffering requirements: Low requirement.
- o Security requirements: For health-critical information, data privacy and security must be provided. Encryption is required. Some types of notifications sent by the smartphone may not need.
- o Mobility: Low.
- o Time Synchronization: the link layer, which is based on TDMA, provides a basis for time synchronization.
- o Reliability and QoS: a relatively low ratio of message losses is acceptable for periodic sensor readings. End-to-end latency of sensor readings is not subject to stringent requirements. The latency of should be low for critical notifications or alarms, generated by either the smartphone or an Internet cloud service.
- o Traffic patterns: periodic (sensor readings) and aperiodic (smartphone-generated notifications).
- o Security Bootstrapping: Required.

5.4. Use case of DECT-ULE: Smart Home

DECT is a technology widely used for wireless telephone communications in residential scenarios. Since DECT-ULE is a low-power variant of DECT, DECT-ULE can be used to connect constrained devices such as sensors and actuators to a Fixed Part, a device that typically acts as a base station for wireless telephones. Therefore, DECT-ULE is specially suitable for the connected home space in application areas such as home automation, smart metering, safety, healthcare, etc.

Example: use of DECT-ULE for Smart Metering

The smart electricity meter of a home is equipped with a DECT-ULE transceiver. This device is in the coverage range of the Fixed Part of the home. The Fixed Part can act as a router connected to the Internet. This way, the smart meter can transmit electricity consumption readings through the DECT-ULE link with the Fixed Part, and the latter can forward such readings to the utility company using Wide Area Network (WAN) links. The meter can also receive queries from the utility company or from an advanced energy control system controlled by the user, which may also be connected to the Fixed Part via DECT-ULE.

Dominant parameters in smart metering scenarios with DECT-ULE:

- o Deployment/Bootstrapping: Pre-planned.
- o Topology: Star topology.
- o L2-mesh or L3-mesh: No.
- o Multi-link subnet, single subnet: Multi-link subnet.
- o Data rate: Small data rate, infrequent transmissions.
- o Buffering requirements: Low requirement.
- o Security requirements: Data privacy and security must be provided. Encryption is required.
- o Mobility: No.
- o Time Synchronization: TBD.
- o Reliability and QoS: bounded latency, stringent reliability service agreements [I-D.ietf-roll-applicability-ami].
- o Traffic patterns: Periodic (meter reading notifications sent by the meter) and aperiodic (user- or company-triggered queries to the meter, and messages triggered by local events such as power outage or leak detection [I-D.ietf-roll-applicability-ami]).
- o Security Bootstrapping: required.

5.5. Use case of LTE MTC

Wireless link layer technologies can be divided short range connectivity and long range connectivity. BLE, ITU-T G.9959 (Z-Wave), DECT-ULE, MS/TP, NFC are used for short range connectivity. LTE MTC is used for long range connectivity. And there is another long range connectivity technology. It is LPWAN(Low Power Wide Area Network) technology such as LoRa, Sigfox and etc. Therefore, the use case of LTE MTC should be compared to LPWAN.

Example: Use of wireless backhaul for LoRa gateway

LoRa is the most promising technology of LPWAN. LoRa network architecture has a star of star topology. LoRa gateway relay the messages from LoRa end device to application server and vice versa. LoRa gateway can has two types of backhaul, wired and wireless backhaul.

If LoRa gateway has wireless backhaul, it should have LTE modem. Since the modem cost of LTE MTC is cheaper than the modem cost of above LTE category 2, it is helpful to design to use LTE MTC. Since the maximum data rate of LoRa end device is 50kbps, it is sufficient to use LTE MTC without using category 2.

Dominant parameters in LoRa gateway scenarios with above example:

- o Deployment/Bootstrapping: Pre-planned.
- o Topology: Star topology.
- o L2-mesh or L3-mesh: No.
- o Multi-link subnet, single subnet: Single subnet.
- o Data rate: depends on 3GPP specification.
- o Buffering requirements: High requirement.
- o Security requirements: No, because data security is already provided in LoRa specification.
- o Mobility: Static.
- o Time Synchronization: Highly required.
- o Reliability and QoS: TBD.
- o Traffic patterns: Random.
- o Security Bootstrapping: required.

Example: Use of controlling car

Car sharing service becomes more popular. Customers wish to control the car with smart phone application. For example, customers wish to lock/unlock the car door with smart phone application, because customers may not have a car key. Customers wish to blow with smart phone application to locate the car easily.

Therefore, rental car should have a long range connectivity capable modem such as LoRa end device and LTE UE. However, LoRa may not be used because LoRa has low reliability and may not be supported in indoor environment such as basement parking lot. And since the message of controlling car is very small, it is sufficient to use LTE MTC but above category 2.

Dominant parameters in controlling car scenarios with above example:

- o Deployment/Bootstrapping: Pre-planned.
- o Topology: Star topology.
- o L2-mesh or L3-mesh: No.
- o Multi-link subnet, single subnet: Single subnet.
- o Data rate: depends on 3GPP specification.
- o Buffering requirements: High requirement.
- o Security requirements: High requirement.
- o Mobility: Always dynamic .
- o Time Synchronization: Highly required.
- o Reliability and QoS: TBD.
- o Traffic patterns: Random.
- o Security Bootstrapping: required.

6. IANA Considerations

There are no IANA considerations related to this document.

7. Security Considerations

[TBD]

8. Acknowledgements

Carles Gomez has been funded in part by the Spanish Government (Ministerio de Educacion, Cultura y Deporte) through the Jose Castillejo grant CAS15/00336. His contribution to this work has been carried out during his stay as a visiting scholar at the Computer Laboratory of the University of Cambridge.

9. References

9.1. 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>>.
- [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, DOI 10.17487/RFC4919, August 2007, <<http://www.rfc-editor.org/info/rfc4919>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [RFC5826] Brandt, A., Buron, J., and G. Porcu, "Home Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5826, DOI 10.17487/RFC5826, April 2010, <<http://www.rfc-editor.org/info/rfc5826>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.
- [RFC6568] Kim, E., Kaspar, D., and JP. Vasseur, "Design and Application Spaces for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6568, DOI 10.17487/RFC6568, April 2012, <<http://www.rfc-editor.org/info/rfc6568>>.
- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<http://www.rfc-editor.org/info/rfc6775>>.
- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<http://www.rfc-editor.org/info/rfc7228>>.

- [RFC7428] Brandt, A. and J. Buron, "Transmission of IPv6 Packets over ITU-T G.9959 Networks", RFC 7428, DOI 10.17487/RFC7428, February 2015, <<http://www.rfc-editor.org/info/rfc7428>>.
- [RFC7668] Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "IPv6 over BLUETOOTH(R) Low Energy", RFC 7668, DOI 10.17487/RFC7668, October 2015, <<http://www.rfc-editor.org/info/rfc7668>>.

9.2. Informative References

- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [I-D.ietf-6lo-dect-ule]
Mariager, P., Petersen, J., Shelby, Z., Logt, M., and D. Barthel, "Transmission of IPv6 Packets over DECT Ultra Low Energy", draft-ietf-6lo-dect-ule-03 (work in progress), September 2015.
- [I-D.ietf-6lo-6lobac]
Lynn, K., Martocci, J., Neilson, C., and S. Donaldson, "Transmission of IPv6 over MS/TP Networks", draft-ietf-6lo-6lobac-02 (work in progress), July 2015.
- [I-D.ietf-6lo-nfc]
Youn, J. and Y. Hong, "Transmission of IPv6 Packets over Near Field Communication", draft-ietf-6lo-nfc-01 (work in progress), July 2015.
- [I-D.ietf-lwig-energy-efficient]
Gomez, C., Kovatsch, M., Tian, H., and Z. Cao, "Energy-Efficient Features of Internet of Things Protocols", draft-ietf-lwig-energy-efficient-04 (work in progress), February 2016.
- [I-D.ietf-roll-applicability-ami]
Popa, D., Gillmore, M., Toutain, L., Hui, J., Salazar, R., Monden, K., and N. Cam-Winget, "Applicability Statement for the Routing Protocol for Low Power and Lossy Networks (RPL) in AMI Networks", draft-ietf-roll-applicability-ami-11 (work in progress), August 2015.

- [G.9959] "International Telecommunication Union, "Short range narrow-band digital radiocommunication transceivers - PHY and MAC layer specifications", ITU-T Recommendation", January 2015.
- [LTE_MTC] "3GPP TS 36.306 V13.0.0, 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; Evolved Universal Terrestrial Radio Access (E-UTRA); User Equipment (UE) radio access capabilities (Release 13)", December 2015.

Authors' Addresses

Yong-Geun Hong
ETRI
161 Gajeong-Dong Yuseung-Gu
Daejeon 305-700
Korea

Phone: +82 42 860 6557
Email: yghong@etri.re.kr

Carles Gomez
Universitat Politecnica de Catalunya/Fundacio i2cat
C/Esteve Terradas, 7
Castelldefels 08860
Spain

Email: carlesgo@entel.upc.edu

Younghwan Choi
ETRI
218 Gajeongno, Yuseong
Daejeon 305-700
Korea

Phone: +82 42 860 1429
Email: yhc@etri.re.kr

Deoknyong Ko
SKtelecom
9-1 Byundang-gu Sunae-dong, Seongnam-si
Gyeonggi-do 13595
Korea

Phone: +82 10 3356 8052
Email: engineer@sk.com

6Lo Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 25, 2016

K. Lynn, Ed.
Verizon Labs
J. Martocci
Johnson Controls
C. Neilson
Delta Controls
S. Donaldson
Honeywell
February 22, 2016

Transmission of IPv6 over MS/TP Networks
draft-ietf-6lo-6lobac-04

Abstract

Master-Slave/Token-Passing (MS/TP) is a medium access control method for the RS-485 physical layer, which is used extensively in building automation networks. This specification defines the frame format for transmission of IPv6 packets and the method of forming link-local and statelessly autoconfigured IPv6 addresses on MS/TP networks.

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 August 25, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents

carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. MS/TP Mode for IPv6	6
3. Addressing Modes	6
4. Maximum Transmission Unit (MTU)	6
5. LoBAC Adaptation Layer	7
6. Stateless Address Autoconfiguration	8
7. IPv6 Link Local Address	8
8. Unicast Address Mapping	9
9. Multicast Address Mapping	9
10. Header Compression	10
11. IANA Considerations	10
12. Security Considerations	10
13. Acknowledgments	11
14. References	11
Appendix A. Abstract MAC Interface	13
Appendix B. Consistent Overhead Byte Stuffing [COBS]	16
Appendix C. Encoded CRC-32K [CRC32K]	20
Appendix D. Example 6LoBAC Packet Decode	22
Authors' Addresses	27

1. Introduction

Master-Slave/Token-Passing (MS/TP) is a medium access control (MAC) protocol for the RS-485 [TIA-485-A] physical layer, which is used extensively in building automation networks. This specification defines the frame format for transmission of IPv6 [RFC2460] packets and the method of forming link-local and statelessly autoconfigured IPv6 addresses on MS/TP networks. The general approach is to adapt elements of the 6LoWPAN specifications [RFC4944], [RFC6282], and [RFC6775] to constrained wired networks.

An MS/TP device is typically based on a low-cost microcontroller with limited processing power and memory. Together with low data rates and a small MAC address space, these constraints are similar to those faced in 6LoWPAN networks and suggest some elements of that solution might be leveraged. MS/TP differs significantly from 6LoWPAN in at least three respects: a) MS/TP devices typically have a continuous source of power, b) all MS/TP devices on a segment can communicate directly so there are no hidden node or mesh routing issues, and c) recent changes to MS/TP provide support for large payloads,

eliminating the need for fragmentation and reassembly below IPv6.

The following sections provide a brief overview of MS/TP, then describe how to form IPv6 addresses and encapsulate IPv6 packets in MS/TP frames. This document also specifies a header compression mechanism, based on [RFC6282], that is REQUIRED in order to reduce latency and make IPv6 practical on MS/TP networks.

1.1. Requirements Language

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].

1.2. Abbreviations Used

ASHRAE: American Society of Heating, Refrigerating, and Air-Conditioning Engineers (<http://www.ashrae.org>)

BACnet: An ISO/ANSI/ASHRAE Standard Data Communication Protocol for Building Automation and Control Networks

CRC: Cyclic Redundancy Check

MAC: Medium Access Control

MSDU: MAC Service Data Unit (MAC client data)

MTU: Maximum Transmission Unit

UART: Universal Asynchronous Transmitter/Receiver

1.3. MS/TP Overview

This section provides a brief overview of MS/TP, which is specified in ANSI/ASHRAE 135-2012 (BACnet) Clause 9 [Clause9] and included herein by reference. BACnet [Clause9] also covers physical layer deployment options.

MS/TP is designed to enable multidrop networks over shielded twisted pair wiring. It can support network segments up to 1000 meters in length at a data rate of 115,200 baud, or segments up to 1200 meters in length at lower baud rates. An MS/TP link requires only a UART, an RS-485 [TIA-485-A] transceiver with a driver that can be disabled, and a 5ms resolution timer. These features make MS/TP a cost-effective field bus for the most numerous and least expensive devices in a building automation network.

The differential signaling used by [TIA-485-A] requires a contention-free MAC. MS/TP uses a token to control access to a multidrop bus. A master node may initiate the transmission of a data frame when it holds the token. After sending at most a configured maximum number of data frames, a master node passes the token to the next master node (as determined by MAC address). Slave nodes do not support the frame format required to convey IPv6 over MS/TP and therefore SHALL NOT be considered part of this specification.

MS/TP COBS-encoded* frames have the following format:

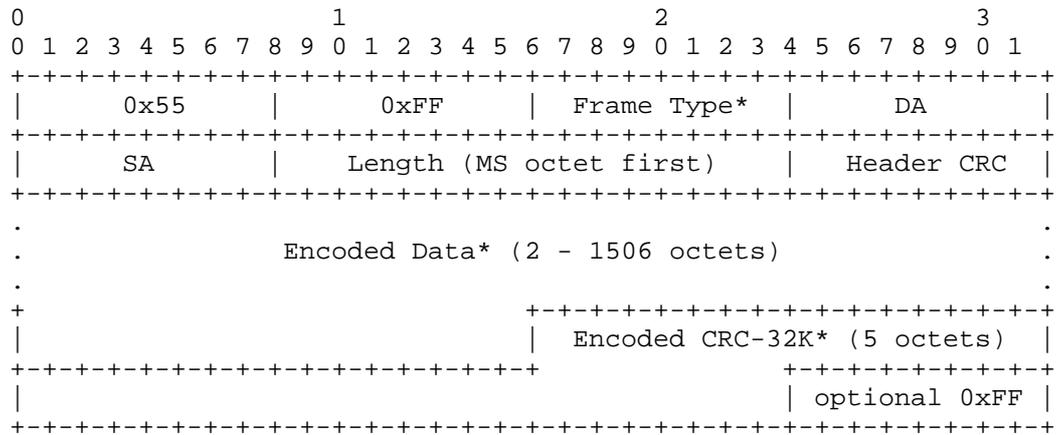


Figure 1: MS/TP COBS-Encoded Frame Format

*NOTE: BACnet Addendum 135-2012an [Addendum_an] defines a range of Frame Type values to designate frames that contain data and data CRC fields encoded using Consistent Overhead Byte Stuffing [COBS] (see Appendix B). The purpose of COBS encoding is to eliminate preamble sequences from the Encoded Data and Encoded CRC-32K fields. The maximum length of an MSDU as defined by this specification is 1500 octets (before encoding). The Encoded Data is covered by a 32-bit CRC [CRC32K] (see Appendix C), which is itself then COBS encoded.

MS/TP COBS-encoded frame fields have the following descriptions:

Preamble	two octet preamble: 0x55, 0xFF
Frame Type	one octet
Destination Address	one octet address
Source Address	one octet address
Length	two octets, most significant octet first
Header CRC	one octet
Encoded Data	2 - 1506 octets (see Appendix B)
Encoded CRC-32K	five octets (see Appendix C)
(pad)	(optional) at most one octet of trailer: 0xFF

The Frame Type is used to distinguish between different types of MAC frames. The types relevant to this specification (in decimal) are:

- 0 Token
- 1 Poll For Master
- 2 Reply To Poll For Master
- ...
- 34 IPv6 over MS/TP (LoBAC) Encapsulation

Frame Types 8 - 31 and 35 - 127 are reserved for assignment by ASHRAE. Frame Types 32 - 127 designate COBS-encoded frames and MUST convey Encoded Data and Encoded CRC-32K fields. All master nodes MUST understand Token, Poll For Master, and Reply to Poll For Master control frames. See Section 2 for additional details.

The Destination and Source Addresses are each one octet in length. See Section 3 for additional details.

For COBS-encoded frames, the Length field specifies the combined length of the [COBS] Encoded Data and Encoded CRC-32K fields in octets, minus two. (This adjustment is required for backward compatibility with legacy MS/TP devices.) See Section 4 and Appendices for additional details.

The Header CRC field covers the Frame Type, Destination Address, Source Address, and Length fields. The Header CRC generation and check procedures are specified in BACnet [Clause9].

1.4. Goals and Constraints

The primary goal of this specification is to enable IPv6 directly on wired end devices in building automation and control networks by leveraging existing standards to the greatest extent possible. A secondary goal is to co-exist with legacy MS/TP implementations. Only the minimum changes necessary to support IPv6 over MS/TP were specified in BACnet [Addendum_an] (see Note in Section 1.3).

In order to co-exist with legacy devices, no changes are permitted to the MS/TP addressing modes, frame header format, control frames, or Master Node state machine as specified in BACnet [Clause9].

2. MS/TP Mode for IPv6

ASHRAE has assigned an MS/TP Frame Type value of 34 to indicate IPv6 over MS/TP (LoBAC) Encapsulation. This falls within the range of values that designate COBS-encoded data frames.

All MS/TP master nodes (including those that support IPv6) must understand Token, Poll For Master, and Reply to Poll For Master control frames and support the Master Node state machine as specified in BACnet [Clause9]. MS/TP master nodes that support IPv6 must also support the Receive Frame state machine as specified in [Clause9] and extended by BACnet [Addendum_an].

All MS/TP nodes that support IPv6 MUST support a data rate of 115,200 baud and MAY optionally support lower data rates as defined in BACnet [Clause9].

3. Addressing Modes

MS/TP node (MAC) addresses are one octet in length. The method of assigning MAC addresses is outside the scope of this specification. However, each MS/TP node on the link MUST have a unique address in order to ensure correct MAC operation.

BACnet [Clause9] specifies that addresses 0 through 127 are valid for master nodes. The method specified in Section 6 for creating a MAC-layer-derived Interface Identifier (IID) ensures that an IID of all zeros can never result.

A Destination Address of 255 (all nodes) indicates a MAC-layer broadcast. MS/TP does not support multicast, therefore all IPv6 multicast packets SHOULD be broadcast at the MAC layer and filtered at the IPv6 layer. A Source Address of 255 MUST NOT be used.

This specification assumes that at most one unique local and/or global IPv6 prefix is assigned to each MS/TP segment. Hosts learn IPv6 prefixes via router advertisements according to [RFC4861].

4. Maximum Transmission Unit (MTU)

BACnet [Addendum_an] supports MSDUs up to 2032 octets in length. This specification defines an MSDU length of at least 1280 octets and at most 1500 octets (before encoding). This is sufficient to convey the minimum MTU required by IPv6 [RFC2460] without the need for link-

layer fragmentation and reassembly. Support for an MSDU length of 1500 octets is RECOMMENDED.

5. LoBAC Adaptation Layer

The relatively low data rates of MS/TP indicate header compression as a means to reduce latency. This section specifies an adaptation layer to support compressed IPv6 headers and the compression format is specified in Section 10.

Implementations MAY also support Generic Header Compression (GHC) [RFC7400] for transport layer headers. A node implementing [RFC7400] MUST probe its peers for GHC support before applying GHC.

The encapsulation format defined in this section (subsequently referred to as the "LoBAC" encapsulation) comprises the MSDU of an IPv6 over MS/TP frame. The LoBAC payload (i.e., an IPv6 packet) follows an encapsulation header stack. LoBAC is a subset of the LOWPAN encapsulation defined in [RFC4944] and extended by [RFC6282], therefore the use of "LOWPAN" in literals below is intentional. The primary difference between LOWPAN and LoBAC is omission of the Mesh, Broadcast, Fragmentation, and LOWPAN_HC1 headers.

All LoBAC encapsulated datagrams transmitted over MS/TP are prefixed by an encapsulation header stack consisting of a Dispatch value followed by zero or more header fields. The only sequence currently defined for LoBAC is the LOWPAN_IPHC header followed by payload, as shown below:

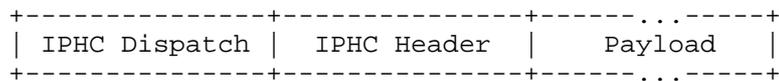


Figure 2: A LoBAC Encapsulated LOWPAN_IPHC Compressed IPv6 Datagram

The Dispatch value may be treated as an unstructured namespace. Only a single pattern is used to represent current LoBAC functionality.

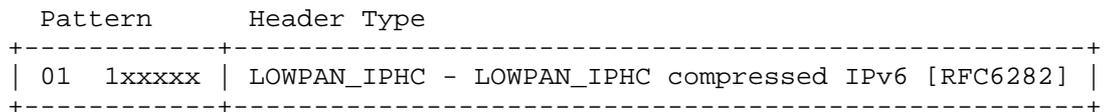


Figure 3: LoBAC Dispatch Value Bit Pattern

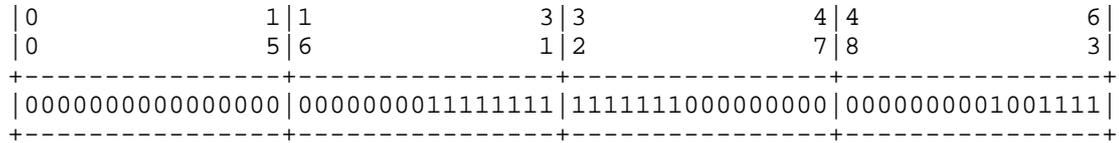
Other IANA-assigned 6LoWPAN Dispatch values do not apply to this specification.

6. Stateless Address Autoconfiguration

This section defines how to obtain an IPv6 Interface Identifier. The general procedure for creating a MAC-address-derived IID is described in [RFC4291] Appendix A, "Creating Modified EUI-64 Format Interface Identifiers", as updated by [RFC7136].

The IID SHOULD NOT embed an [EUI-64] or any other globally unique hardware identifier assigned to a device (see Section 12).

The Interface Identifier for link-local addresses SHOULD be formed by concatenating a node's 8-bit MS/TP MAC address to the seven octets 0x00, 0x00, 0x00, 0xFF, 0xFE, 0x00, 0x00. For example, an MS/TP MAC address of hexadecimal value 0x4F results in the following IID:



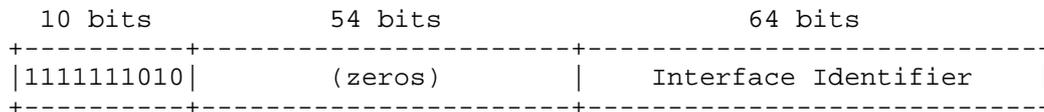
This is the RECOMMENDED method of forming an IID for use in link-local addresses, as it affords the most efficient header compression provided by the LOWPAN_IPHC [RFC6282] format specified in Section 10.

A 64-bit privacy IID is RECOMMENDED for routable addresses and SHOULD be locally generated according to one of the methods cited in Section 12. A node that generates a 64-bit privacy IID MUST register it with its local router(s) by sending a Neighbor Solicitation (NS) message with the Address Registration Option (ARO) and process Neighbor Advertisements (NA) according to [RFC6775].

An IPv6 address prefix used for stateless autoconfiguration [RFC4862] of an MS/TP interface MUST have a length of 64 bits.

7. IPv6 Link Local Address

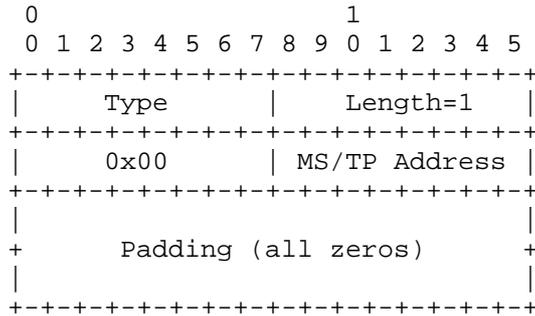
The IPv6 link-local address [RFC4291] for an MS/TP interface is formed by appending the Interface Identifier, as defined above, to the prefix FE80::/64.



8. Unicast Address Mapping

The address resolution procedure for mapping IPv6 non-multicast addresses into MS/TP MAC-layer addresses follows the general description in Section 7.2 of [RFC4861], unless otherwise specified.

The Source/Target Link-layer Address option has the following form when the addresses are 8-bit MS/TP MAC-layer (node) addresses.



Option fields:

Type:

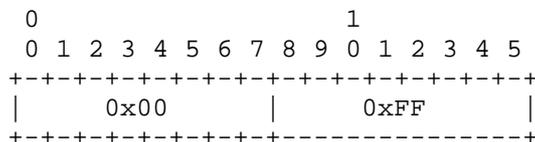
- 1: for Source Link-layer address.
- 2: for Target Link-layer address.

Length: This is the length of this option (including the type and length fields) in units of 8 octets. The value of this field is 1 for 8-bit MS/TP MAC addresses.

MS/TP Address: The 8-bit address in canonical bit order [RFC2469]. This is the unicast address the interface currently responds to.

9. Multicast Address Mapping

All IPv6 multicast packets SHOULD be sent to MS/TP Destination Address 255 (broadcast) and filtered at the IPv6 layer. When represented as a 16-bit address in a compressed header (see Section 10), it MUST be formed by padding on the left with a zero:



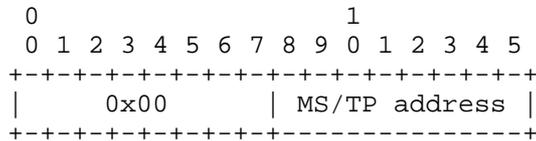
10. Header Compression

LoBAC uses LOWPAN_IPHC IPv6 compression, which is specified in [RFC6282] and included herein by reference. This section will simply identify substitutions that should be made when interpreting the text of [RFC6282].

In general the following substitutions should be made:

- Replace instances of "6LoWPAN" with "MS/TP network"
- Replace instances of "IEEE 802.15.4 address" with "MS/TP address"

When a 16-bit address is called for (i.e., an IEEE 802.15.4 "short address") it MUST be formed by padding the MS/TP address to the left with a zero:



If LOWPAN_IPHC compression [RFC6282] is used with context, the border router(s) directly attached to the MS/TP segment MUST disseminate the 6LoWPAN Context Option (6CO) according to [RFC6775], Section 7.2.

11. IANA Considerations

This document uses values previously reserved by [RFC4944] and [RFC6282] and makes no further requests of IANA.

Note to RFC Editor: this section may be removed upon publication.

12. Security Considerations

Routable addresses that contain IIDs generated using MS/TP node addresses may expose a network to address scanning attacks. For this reason, it is RECOMMENDED that a different (but stable) IID be generated for each routable address in use according to, for example, [RFC3315], [RFC3972], [RFC4941], [RFC5535], or [RFC7217].

MS/TP networks are by definition wired and not susceptible to casual eavesdropping. By the same token, MS/TP nodes are stationary and correlation of activities or location tracking of individuals is unlikely.

13. Acknowledgments

We are grateful to the authors of [RFC4944] and members of the IETF 6LoWPAN working group; this document borrows liberally from their work. Ralph Droms and Brian Haberman provided indispensable guidance and support from the outset. Peter van der Stok, James Woodyatt, and Carsten Bormann provided detailed reviews. Stuart Cheshire invented the very clever COBS encoding. Michael Osborne made the critical observation that seperately encoding the data and CRC32K fields would allow the CRC to be calculated on-the-fly. Alexandru Petrescu, Brian Frank, Geoff Mulligan, and Don Sturek offered valuable comments.

14. References

14.1. Normative References

[Addendum_an]

ASHRAE, "ANSI/ASHRAE Addenda an, at, au, av, aw, ax, and az to ANSI/ASHRAE Standard 135-2012, BACnet - A Data Communication Protocol for Building Automation and Control Networks", July 2014, <https://www.ashrae.org/File%20Library/docLib/StdAddenda/07-31-2014_135_2012_an_at_au_av_aw_ax_az_Final.pdf>.

[Clause9] American Society of Heating, Refrigerating, and Air-Conditioning Engineers, "BACnet - A Data Communication Protocol for Building Automation and Control Networks", ANSI/ASHRAE 135-2012 (Clause 9), March 2013.

[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>>.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.

[RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <<http://www.rfc-editor.org/info/rfc3315>>.

[RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <<http://www.rfc-editor.org/info/rfc3972>>.

- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<http://www.rfc-editor.org/info/rfc4291>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<http://www.rfc-editor.org/info/rfc4861>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<http://www.rfc-editor.org/info/rfc4862>>.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007, <<http://www.rfc-editor.org/info/rfc4941>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [RFC5535] Bagnulo, M., "Hash-Based Addresses (HBA)", RFC 5535, DOI 10.17487/RFC5535, June 2009, <<http://www.rfc-editor.org/info/rfc5535>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.
- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<http://www.rfc-editor.org/info/rfc6775>>.
- [RFC7136] Carpenter, B. and S. Jiang, "Significance of IPv6 Interface Identifiers", RFC 7136, DOI 10.17487/RFC7136, February 2014, <<http://www.rfc-editor.org/info/rfc7136>>.
- [RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <<http://www.rfc-editor.org/info/rfc7217>>.

- [RFC7400] Bormann, C., "6LoWPAN-GHC: Generic Header Compression for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 7400, DOI 10.17487/RFC7400, November 2014, <<http://www.rfc-editor.org/info/rfc7400>>.

14.2. Informative References

- [COBS] Cheshire, S. and M. Baker, "Consistent Overhead Byte Stuffing", IEEE/ACM TRANSACTIONS ON NETWORKING, VOL.7, NO.2 , April 1999, <<http://www.stuartcheshire.org/papers/COBSforToN.pdf>>.
- [CRC32K] Koopman, P., "32-Bit Cyclic Redundancy Codes for Internet Applications", IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2002) , June 2002, <http://www.ece.cmu.edu/~koopman/networks/dsn02/dsn02_koopman.pdf>.
- [EUI-64] IEEE, "Guidelines for 64-bit Global Identifier (EUI-64) Registration Authority", March 1997, <<http://standards.ieee.org/regauth/oui/tutorials/EUI64.html>>.
- [IEEE.802.3]
"Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications", IEEE Std 802.3-2012, December 2012, <<http://standards.ieee.org/getieee802/802.3.html>>.
- [RFC2469] Narten, T. and C. Burton, "A Caution On The Canonical Ordering Of Link-Layer Addresses", RFC 2469, DOI 10.17487/RFC2469, December 1998, <<http://www.rfc-editor.org/info/rfc2469>>.
- [TIA-485-A]
Telecommunications Industry Association, "TIA-485-A, Electrical Characteristics of Generators and Receivers for Use in Balanced Digital Multipoint Systems (ANSI/TIA/EIA-485-A-98) (R2003)", March 2003.

Appendix A. Abstract MAC Interface

This Appendix is informative and not part of the standard.

BACnet [Clause9] defines support for MAC-layer clients through its

SendFrame and ReceivedDataNoReply procedures. However, it does not define a network-protocol independent abstract interface for the MAC. This is provided below as an aid to implementation.

A.1. MA-DATA.request

A.1.1. Function

This primitive defines the transfer of data from a MAC client entity to a single peer entity or multiple peer entities in the case of a broadcast address.

A.1.2. Semantics of the Service Primitive

The semantics of the primitive are as follows:

```
MA-DATA.request (
    destination_address,
    source_address,
    data,
    priority,
    type
)
```

The 'destination_address' parameter may specify either an individual or a broadcast MAC entity address. It must contain sufficient information to create the Destination Address field (see Section 1.3) that is prepended to the frame by the local MAC sublayer entity. The 'source_address' parameter, if present, must specify an individual MAC address. If the source_address parameter is omitted, the local MAC sublayer entity will insert a value associated with that entity.

The 'data' parameter specifies the MAC service data unit (MSDU) to be transferred by the MAC sublayer entity. There is sufficient information associated with the MSDU for the MAC sublayer entity to determine the length of the data unit.

The 'priority' parameter specifies the priority desired for the data unit transfer. The priority parameter is ignored by MS/TP.

The 'type' parameter specifies the value of the MS/TP Frame Type field that is prepended to the frame by the local MAC sublayer entity.

A.1.3. When Generated

This primitive is generated by the MAC client entity whenever data shall be transferred to a peer entity or entities. This can be in response to a request from higher protocol layers or from data generated internally to the MAC client, such as a Token frame.

A.1.4. Effect on Receipt

Receipt of this primitive will cause the MAC entity to insert all MAC specific fields, including Destination Address, Source Address, Frame Type, and any fields that are unique to the particular media access method, and pass the properly formed frame to the lower protocol layers for transfer to the peer MAC sublayer entity or entities.

A.2. MA-DATA.indication

A.2.1. Function

This primitive defines the transfer of data from the MAC sublayer entity to the MAC client entity or entities in the case of a broadcast address.

A.2.2. Semantics of the Service Primitive

The semantics of the primitive are as follows:

```
MA-DATA.indication (  
    destination_address,  
    source_address,  
    data,  
    priority,  
    type  
)
```

The 'destination_address' parameter may be either an individual or a broadcast address as specified by the Destination Address field of the incoming frame. The 'source_address' parameter is an individual address as specified by the Source Address field of the incoming frame.

The 'data' parameter specifies the MAC service data unit (MSDU) as received by the local MAC entity. There is sufficient information associated with the MSDU for the MAC sublayer client to determine the length of the data unit.

The 'priority' parameter specifies the priority desired for the data unit transfer. The priority parameter is ignored by MS/TP.

The 'type' parameter is the value of the MS/TP Frame Type field of the incoming frame.

A.2.3. When Generated

The MA_DATA.indication is passed from the MAC sublayer entity to the MAC client entity or entities to indicate the arrival of a frame to the local MAC sublayer entity that is destined for the MAC client. Such frames are reported only if they are validly formed, received without error, and their destination address designates the local MAC entity. Frames destined for the MAC Control sublayer are not passed to the MAC client.

A.2.4. Effect on Receipt

The effect of receipt of this primitive by the MAC client is unspecified.

Appendix B. Consistent Overhead Byte Stuffing [COBS]

This Appendix is informative and not part of the standard.

BACnet [Addendum_an] corrects a long-standing issue with the MS/TP specification; namely that preamble sequences were not escaped whenever they appeared in the Data or Data CRC fields. In rare cases, this resulted in dropped frames due to loss of frame synchronization. The solution is to encode the Data and 32-bit Data CRC fields before transmission using Consistent Overhead Byte Stuffing [COBS] and decode these fields upon reception.

COBS is a run-length encoding method that nominally removes '0x00' octets from its input. Any selected octet value may be removed by XOR'ing that value with each octet of the COBS output. BACnet [Addendum_an] specifies the preamble octet '0x55' for removal.

The minimum overhead of COBS is one octet per encoded field. The worst-case overhead in long fields is bounded to one octet in 254, or less than 0.4%, as described in [COBS].

Frame encoding proceeds logically in two passes. The Encoded Data field is prepared by passing the MSDU through the COBS encoder and XOR'ing the preamble octet '0x55' with each octet of the output. The Encoded CRC-32K field is then prepared by calculating a CRC-32K over the Encoded Data field and formatting it for transmission as described in Appendix C. The combined length of these fields, minus two octets for compatibility with existing MS/TP devices, is placed in the MS/TP header Length field before transmission.

Example COBS encoder and decoder functions are shown below for illustration. Complete examples of use and test vectors are provided in BACnet [Addendum_an].

```
#include <stddef.h>
#include <stdint.h>

#define CRC32K_INITIAL_VALUE (0xFFFFFFFF)
#define MSTP_PREAMBLE_X55 (0x55)

/*
 * Encodes 'length' octets of data located at 'from' and
 * writes one or more COBS code blocks at 'to', removing any
 * 'mask' octets that may present be in the encoded data.
 * Returns the length of the encoded data.
 */

size_t
cobs_encode (uint8_t *to, const uint8_t *from, size_t length,
            uint8_t mask)
{
    size_t code_index = 0;
    size_t read_index = 0;
    size_t write_index = 1;
    uint8_t code = 1;
    uint8_t data, last_code;

    while (read_index < length) {
        data = from[read_index++];
        /*
         * In the case of encountering a non-zero octet in the data,
         * simply copy input to output and increment the code octet.
         */
        if (data != 0) {
            to[write_index++] = data ^ mask;
            code++;
            if (code != 255)
                continue;
        }
        /*
         * In the case of encountering a zero in the data or having
         * copied the maximum number (254) of non-zero octets, store
         * the code octet and reset the encoder state variables.
         */
        last_code = code;
        to[code_index] = code ^ mask;
        code_index = write_index++;
        code = 1;
    }
}
```

```
    }
    /*
    * If the last chunk contains exactly 254 non-zero octets, then
    * this exception is handled above (and returned length must be
    * adjusted). Otherwise, encode the last chunk normally, as if
    * a "phantom zero" is appended to the data.
    */
    if ((last_code == 255) && (code == 1))
        write_index--;
    else
        to[code_index] = code ^ mask;

    return write_index;
}
```

```
#include <stddef.h>
#include <stdint.h>

#define CRC32K_INITIAL_VALUE (0xFFFFFFFF)
#define MSTP_PREAMBLE_X55 (0x55)

/*
 * Decodes 'length' octets of data located at 'from' and
 * writes the original client data at 'to', restoring any
 * 'mask' octets that may present in the encoded data.
 * Returns the length of the encoded data or zero if error.
 */
size_t
cobs_decode (uint8_t *to, const uint8_t *from, size_t length,
             uint8_t mask)
{
    size_t read_index = 0;
    size_t write_index = 0;
    uint8_t code, last_code;

    while (read_index < length) {
        code = from[read_index] ^ mask;
        last_code = code;
        /*
         * Sanity check the encoding to prevent the while() loop below
         * from overrunning the output buffer.
         */
        if (read_index + code > length)
            return 0;

        read_index++;
        while (--code > 0)
            to[write_index++] = from[read_index++] ^ mask;
        /*
         * Restore the implicit zero at the end of each decoded block
         * except when it contains exactly 254 non-zero octets or the
         * end of data has been reached.
         */
        if ((last_code != 255) && (read_index < length))
            to[write_index++] = 0;
    }
    return write_index;
}
```

Appendix C. Encoded CRC-32K [CRC32K]

This Appendix is informative and not part of the standard.

Extending the payload of MS/TP to 1500 octets required upgrading the Data CRC from 16 bits to 32 bits. P.Koopman has authored several papers on evaluating CRC polynomials for network applications. In [CRC32K], he surveyed the entire 32-bit polynomial space and noted some that exceed the [IEEE.802.3] polynomial in performance. BACnet [Addendum_an] specifies the CRC-32K (Koopman) polynomial.

The specified use of the `calc_crc32K()` function is as follows. Before a frame is transmitted, 'crc_value' is initialized to all ones. After passing each octet of the [COBS] Encoded Data through the function, the ones complement of the resulting 'crc_value' is arranged in LSB-first order and is itself [COBS] encoded. The length of the resulting Encoded CRC-32K field is always five octets.

Upon reception of a frame, 'crc_value' is initialized to all ones. The octets of the Encoded Data field are accumulated by the `calc_crc32K()` function before decoding. The Encoded CRC-32K field is then decoded and the resulting four octets are accumulated by the `calc_crc32K()` function. If the result is the expected residue value 'CRC32K_RESIDUE', then the frame was received correctly.

An example CRC-32K function is shown below for illustration. Complete examples of use and test vectors are provided in BACnet [Addendum_an].

```
#include <stdint.h>

/* See BACnet Addendum 135-2012an, section G.3.2 */
#define CRC32K_INITIAL_VALUE (0xFFFFFFFF)
#define CRC32K_RESIDUE (0x0843323B)

/* CRC-32K polynomial, 1 + x**1 + ... + x**30 (+ x**32) */
#define CRC32K_POLY (0xEB31D82E)

/*
 * Accumulate 'data_value' into the CRC in 'crc_value'.
 * Return updated CRC.
 *
 * Note: crcValue must be set to CRC32K_INITIAL_VALUE
 * before initial call.
 */
uint32_t
calc_crc32k (uint8_t data_value, uint32_t crc_value)
{
    int b;

    for (b = 0; b < 8; b++) {
        if ((data_value & 1) ^ (crc_value & 1)) {
            crc_value >>= 1;
            crc_value ^= CRC32K_POLY;
        } else {
            crc_value >>= 1;
        }
        data_value >>= 1;
    }
    return crc_value;
}
```

Appendix D. Example 6LoBAC Packet Decode

This Appendix is informative and not part of the standard.

```

No.      Time          Source          Destination
 5161  8.816048      aaaa::1       aaaa::ff:fe00:1
Protocol Length Info
ICMPv6   547      Echo (ping) request id=0x2ee5, seq=2,
                hop limit=63 (reply in 5165)

Frame 5161: 547 bytes on wire (4376 bits), 547 bytes captured
            (4376 bits) on interface 0
  Interface id: 0 (/tmp/pipe)
  Encapsulation type: BACnet MS/TP (63)
  Arrival Time: Sep  3, 2015 19:46:44.377881000 EDT
  [Time shift for this packet: 0.000000000 seconds]
  Epoch Time: 1441324004.377881000 seconds
  [Time delta from previous captured frame: 0.050715000 seconds]
  [Time delta from previous displayed frame: 0.050715000 seconds]
  [Time since reference or first frame: 8.816048000 seconds]
  Frame Number: 5161
  Frame Length: 547 bytes (4376 bits)
  Capture Length: 547 bytes (4376 bits)
  [Frame is marked: False]
  [Frame is ignored: False]
  [Protocols in frame: mstp:6lowpan:ipv6:ipv6.nxt:icmpv6:data]
  [Coloring Rule Name: ICMP]
  [Coloring Rule String: icmp || icmpv6]
BACnet MS/TP, Src (2), Dst (1), IPv6 Encapsulation
  Preamble 55: 0x55
  Preamble FF: 0xff
  Frame Type: IPv6 Encapsulation (34)
  Destination Address: 1
  Source Address: 2
  Length: 537
  Header CRC: 0x1c [correct]
    [Good: True]
    [Bad: False]
  Extended Data CRC: 0x9e7259e2 [correct]
6LoWPAN
  IPHC Header
    011. .... = Pattern: IP header compression (0x03)
    ...1 1... .... .... = Traffic class and flow label:
                          Version, traffic class, and flow label
                          compressed (0x0003)
    .... .0.. .... .... = Next header: Inline
    .... ..00 .... .... = Hop limit: Inline (0x0000)

```

```

..... 1... .. = Context identifier extension: True
.... 1... .. = Source address compression: Stateful
.... ..01 .. = Source address mode:
                    64-bits inline (0x0001)
..... 0... .. = Multicast address compression: False
..... 1... .. = Destination address compression:
                    Stateful
..... ..10 = Destination address mode:
                    16-bits inline (0x0002)
0000 .... = Source context identifier: 0x00
.... 0000 = Destination context identifier: 0x00
[Source context: aaaa:: (aaaa::)]
[Destination context: aaaa:: (aaaa::)]
Next header: ICMPv6 (0x3a)
Hop limit: 63
Source: aaaa::1 (aaaa::1)
Destination: aaaa::ff:fe00:1 (aaaa::ff:fe00:1)
Internet Protocol Version 6, Src: aaaa::1 (aaaa::1),
                    Dst: aaaa::ff:fe00:1 (aaaa::ff:fe00:1)
0110 .... .. = Version: 6
.... 0000 0000 .... .. = Traffic class:
                    0x00000000
.... 0000 00.. .... .. = Differentiated
                    Services Field:
                    Default (0x00000000)
.... ..0. .... .. = ECN-Capable Transport
                    (ECT): Not set
.... ..0 .... .. = ECN-CE: Not set
.... ..0000 0000 0000 0000 0000 = Flowlabel: 0x00000000
Payload length: 518
Next header: ICMPv6 (58)
Hop limit: 63
Source: aaaa::1 (aaaa::1)
Destination: aaaa::ff:fe00:1 (aaaa::ff:fe00:1)
Internet Control Message Protocol v6
Type: Echo (ping) request (128)
Code: 0
Checksum: 0x783f [correct]
Identifier: 0x2ee5
Sequence: 2
[Response In: 5165]
Data (510 bytes)
    Data: e4dbe8553ba0040008090a0b0c0d0e0f1011121314151617...
    [Length: 510]

```

Frame (547 bytes):

```

55 ff 22 01 02 02 19 1c 56 2d 83 56 6f 6a 54 54 U.".....V-.VojTT
54 54 54 54 57 54 56 54 d5 50 2d 6a 7b b0 5c 57 TTTTWTVT.P-j{.\W
b1 8e bd 00 6e f5 51 ac 5d 5c 5f 5e 59 58 5b 5a ....n.Q.]\_^YX[Z
45 44 47 46 41 40 43 42 4d 4c 4f 4e 49 48 4b 4a EDGFA@CBMLONIHKJ
75 74 77 76 71 70 73 72 7d 7c 7f 7e 79 78 7b 7a utwvqpsr}|.~yx{z
65 64 67 66 61 60 63 62 6d 6c 6f 6e 69 68 6b 6a edgfa`cbmlonihkj
15 14 17 16 11 10 13 12 1d 1c 1f 1e 19 18 1b 1a .....
05 04 07 06 01 00 03 02 0d 0c 0f 0e 09 08 0b 0a .....
35 34 37 36 31 30 33 32 3d 3c 3f 3e 39 38 3b 3a 54761032=<?>98;:
25 24 27 26 21 20 23 22 2d 2c 2f 2e 29 28 2b 2a %$'&! #-,/.)(+*
d5 d4 d7 d6 d1 d0 d3 d2 dd dc df de d9 d8 db da .....
c5 c4 c7 c6 c1 c0 c3 c2 cd cc cf ce c9 c8 cb ca .....
f5 f4 f7 f6 f1 f0 f3 f2 fd fc ff fe f9 f8 fb fa .....
e5 e4 e7 e6 e1 e0 e3 e2 ed ec ef ee e9 e8 eb ea .....
95 94 97 96 91 90 93 92 9d 9c 9f 9e 99 98 9b 9a .....
85 84 87 86 81 80 83 82 8d 8c 8f 8e 89 88 8b 8a .....
b5 b4 b7 b6 b1 b0 b3 b2 bd bc bf be b9 b8 bb ba .....
a5 a4 a7 a6 a1 a0 a3 a2 ad ac af ae a9 a8 ab aa .....
ab 54 57 56 51 50 53 52 5d 5c 5f 5e 59 58 5b 5a .TWVQPSR]\_^YX[Z
45 44 47 46 41 40 43 42 4d 4c 4f 4e 49 48 4b 4a EDGFA@CBMLONIHKJ
75 74 77 76 71 70 73 72 7d 7c 7f 7e 79 78 7b 7a utwvqpsr}|.~yx{z
65 64 67 66 61 60 63 62 6d 6c 6f 6e 69 68 6b 6a edgfa`cbmlonihkj
15 14 17 16 11 10 13 12 1d 1c 1f 1e 19 18 1b 1a .....
05 04 07 06 01 00 03 02 0d 0c 0f 0e 09 08 0b 0a .....
35 34 37 36 31 30 33 32 3d 3c 3f 3e 39 38 3b 3a 54761032=<?>98;:
25 24 27 26 21 20 23 22 2d 2c 2f 2e 29 28 2b 2a %$'&! #-,/.)(+*
d5 d4 d7 d6 d1 d0 d3 d2 dd dc df de d9 d8 db da .....
c5 c4 c7 c6 c1 c0 c3 c2 cd cc cf ce c9 c8 cb ca .....
f5 f4 f7 f6 f1 f0 f3 f2 fd fc ff fe f9 f8 fb fa .....
e5 e4 e7 e6 e1 e0 e3 e2 ed ec ef ee e9 e8 eb ea .....
95 94 97 96 91 90 93 92 9d 9c 9f 9e 99 98 9b 9a .....
85 84 87 86 81 80 83 82 8d 8c 8f 8e 89 88 8b 8a .....
b5 b4 b7 b6 b1 b0 b3 b2 bd bc bf be b9 b8 bb ba .....
a5 a4 a7 a6 a1 a0 a3 a2 ad ac af ae a9 a8 50 cb .....P.
27 0c b7 '...'

```

Decoded Data and CRC32K (537 bytes):

```

78 d6 00 3a 3f 00 00 00 00 00 00 01 00 01 80 x...:?......
00 78 3f 2e e5 00 02 e4 db e8 55 3b a0 04 00 08 .x?.....U;....
09 0a 0b 0c 0d 0e 0f 10 11 12 13 14 15 16 17 18 .....
19 1a 1b 1c 1d 1e 1f 20 21 22 23 24 25 26 27 28 ..... !"#$$%&'(
29 2a 2b 2c 2d 2e 2f 30 31 32 33 34 35 36 37 38 )*+,-./012345678
39 3a 3b 3c 3d 3e 3f 40 41 42 43 44 45 46 47 48 9:;<=>?@ABCDEFGH
49 4a 4b 4c 4d 4e 4f 50 51 52 53 54 55 56 57 58 IJKLMNOPQRSTUVWXYZ
59 5a 5b 5c 5d 5e 5f 60 61 62 63 64 65 66 67 68 YZ[\]^_`abcdefghijklm
69 6a 6b 6c 6d 6e 6f 70 71 72 73 74 75 76 77 78 ijklmnopqrstuvwxyz
79 7a 7b 7c 7d 7e 7f 80 81 82 83 84 85 86 87 88 yz{|}~.....
89 8a 8b 8c 8d 8e 8f 90 91 92 93 94 95 96 97 98 .....
99 9a 9b 9c 9d 9e 9f a0 a1 a2 a3 a4 a5 a6 a7 a8 .....
a9 aa ab ac ad ae af b0 b1 b2 b3 b4 b5 b6 b7 b8 .....
b9 ba bb bc bd be bf c0 c1 c2 c3 c4 c5 c6 c7 c8 .....
c9 ca cb cc cd ce cf d0 d1 d2 d3 d4 d5 d6 d7 d8 .....
d9 da db dc dd de df e0 e1 e2 e3 e4 e5 e6 e7 e8 .....
e9 ea eb ec ed ee ef f0 f1 f2 f3 f4 f5 f6 f7 f8 .....
f9 fa fb fc fd fe ff 00 01 02 03 04 05 06 07 08 .....
09 0a 0b 0c 0d 0e 0f 10 11 12 13 14 15 16 17 18 .....
19 1a 1b 1c 1d 1e 1f 20 21 22 23 24 25 26 27 28 ..... !"#$$%&'(
29 2a 2b 2c 2d 2e 2f 30 31 32 33 34 35 36 37 38 )*+,-./012345678
39 3a 3b 3c 3d 3e 3f 40 41 42 43 44 45 46 47 48 9:;<=>?@ABCDEFGH
49 4a 4b 4c 4d 4e 4f 50 51 52 53 54 55 56 57 58 IJKLMNOPQRSTUVWXYZ
59 5a 5b 5c 5d 5e 5f 60 61 62 63 64 65 66 67 68 YZ[\]^_`abcdefghijklm
69 6a 6b 6c 6d 6e 6f 70 71 72 73 74 75 76 77 78 ijklmnopqrstuvwxyz
79 7a 7b 7c 7d 7e 7f 80 81 82 83 84 85 86 87 88 yz{|}~.....
89 8a 8b 8c 8d 8e 8f 90 91 92 93 94 95 96 97 98 .....
99 9a 9b 9c 9d 9e 9f a0 a1 a2 a3 a4 a5 a6 a7 a8 .....
a9 aa ab ac ad ae af b0 b1 b2 b3 b4 b5 b6 b7 b8 .....
b9 ba bb bc bd be bf c0 c1 c2 c3 c4 c5 c6 c7 c8 .....
c9 ca cb cc cd ce cf d0 d1 d2 d3 d4 d5 d6 d7 d8 .....
d9 da db dc dd de df e0 e1 e2 e3 e4 e5 e6 e7 e8 .....
e9 ea eb ec ed ee ef f0 f1 f2 f3 f4 f5 f6 f7 f8 .....
f9 fa fb fc fd fe ff 9e 72 59 e2 .....rY.

```

```

Decompressed 6LoWPAN IPHC (558 bytes):
60 00 00 00 02 06 3a 3f aa aa 00 00 00 00 00 00  \.....:~?.....
00 00 00 00 00 00 00 01 aa aa 00 00 00 00 00 00  .....
00 00 00 ff fe 00 00 01 80 00 78 3f 2e e5 00 02  .....x?....
e4 db e8 55 3b a0 04 00 08 09 0a 0b 0c 0d 0e 0f  ...U;.....
10 11 12 13 14 15 16 17 18 19 1a 1b 1c 1d 1e 1f  .....
20 21 22 23 24 25 26 27 28 29 2a 2b 2c 2d 2e 2f  !"#%&'()*+,-./
30 31 32 33 34 35 36 37 38 39 3a 3b 3c 3d 3e 3f  0123456789:;<=>?
40 41 42 43 44 45 46 47 48 49 4a 4b 4c 4d 4e 4f  @ABCDEFGHIJKLMNO
50 51 52 53 54 55 56 57 58 59 5a 5b 5c 5d 5e 5f  PQRSTUVWXYZ[\]^_
60 61 62 63 64 65 66 67 68 69 6a 6b 6c 6d 6e 6f  `abcdefghijklmnop
70 71 72 73 74 75 76 77 78 79 7a 7b 7c 7d 7e 7f  pqrstuvwxyz{|}~.
80 81 82 83 84 85 86 87 88 89 8a 8b 8c 8d 8e 8f  .....
90 91 92 93 94 95 96 97 98 99 9a 9b 9c 9d 9e 9f  .....
a0 a1 a2 a3 a4 a5 a6 a7 a8 a9 aa ab ac ad ae af  .....
b0 b1 b2 b3 b4 b5 b6 b7 b8 b9 ba bb bc bd be bf  .....
c0 c1 c2 c3 c4 c5 c6 c7 c8 c9 ca cb cc cd ce cf  .....
d0 d1 d2 d3 d4 d5 d6 d7 d8 d9 da db dc dd de df  .....
e0 e1 e2 e3 e4 e5 e6 e7 e8 e9 ea eb ec ed ee ef  .....
f0 f1 f2 f3 f4 f5 f6 f7 f8 f9 fa fb fc fd        .....

```

Authors' Addresses

Kerry Lynn (editor)
Verizon Labs
50 Sylvan Rd
Waltham , MA 02451
USA

Phone: +1 781 296 9722
Email: kerlyn@ieee.org

Jerry Martocci
Johnson Controls, Inc.
507 E. Michigan St
Milwaukee , WI 53202
USA

Phone: +1 414 524 4010
Email: jerald.p.martocci@jci.com

Carl Neilson
Delta Controls, Inc.
17850 56th Ave
Surrey , BC V3S 1C7
Canada

Phone: +1 604 575 5913
Email: cneilson@deltacontrols.com

Stuart Donaldson
Honeywell Automation & Control Solutions
6670 185th Ave NE
Redmond , WA 98052
USA

Email: stuart.donaldson@honeywell.com

6Lo Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 22, 2016

Y-G. Hong
Y-H. Choi
ETRI
J-S. Youn
DONG-EUI Univ
D-K. Kim
KNU
J-H. Choi
Samsung Electronics Co.,
March 21, 2016

Transmission of IPv6 Packets over Near Field Communication
draft-ietf-6lo-nfc-03

Abstract

Near field communication (NFC) is a set of standards for smartphones and portable devices to establish radio communication with each other by touching them together or bringing them into proximity, usually no more than 10 cm. NFC standards cover communications protocols and data exchange formats, and are based on existing radio-frequency identification (RFID) standards including ISO/IEC 14443 and FeliCa. The standards include ISO/IEC 18092 and those defined by the NFC Forum. The NFC technology has been widely implemented and available in mobile phones, laptop computers, and many other devices. This document describes how IPv6 is transmitted over NFC using 6LowPAN techniques.

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 22, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Conventions and Terminology	4
3. Overview of Near Field Communication Technology	4
3.1. Peer-to-peer Mode of NFC	4
3.2. Protocol Stacks of NFC	5
3.3. NFC-enabled Device Addressing	6
3.4. NFC MAC PDU Size and MTU	6
4. Specification of IPv6 over NFC	8
4.1. Protocol Stacks	8
4.2. Link Model	9
4.3. Stateless Address Autoconfiguration	10
4.4. IPv6 Link Local Address	10
4.5. Neighbor Discovery	11
4.6. Dispatch Header	11
4.7. Header Compression	12
4.8. Fragmentation and Reassembly	12
4.9. Unicast Address Mapping	13
4.10. Multicast Address Mapping	13
5. Internet Connectivity Scenarios	14
5.1. NFC-enabled Device Connected to the Internet	14
5.2. Isolated NFC-enabled Device Network	15
6. IANA Considerations	15
7. Security Considerations	15
8. Acknowledgements	16
9. References	16
9.1. Normative References	16
9.2. Informative References	17
Authors' Addresses	17

1. Introduction

NFC is a set of short-range wireless technologies, typically requiring a distance of 10 cm or less. NFC operates at 13.56 MHz on ISO/IEC 18000-3 air interface and at rates ranging from 106 kbit/s to 424 kbit/s. NFC always involves an initiator and a target; the initiator actively generates an RF field that can power a passive target. This enables NFC targets to take very simple form factors such as tags, stickers, key fobs, or cards that do not require batteries. NFC peer-to-peer communication is possible, provided both devices are powered. NFC builds upon RFID systems by allowing two-way communication between endpoints, where earlier systems such as contactless smart cards were one-way only. It has been used in devices such as mobile phones, running Android operating system, named with a feature called "Android Beam". In addition, it is expected for the other mobile phones, running the other operating systems (e.g., iOS, etc.) to be equipped with NFC technology in the near future.

Considering the potential for exponential growth in the number of heterogeneous air interface technologies, NFC would be widely used as one of the other air interface technologies, such as Bluetooth Low Energy (BT-LE), Wi-Fi, and so on. Each of the heterogeneous air interface technologies has its own characteristics, which cannot be covered by the other technologies, so various kinds of air interface technologies would be existing together. Therefore, it is required for them to communicate each other. NFC also has the strongest point (e.g., secure communication distance of 10 cm) to prevent the third party from attacking privacy.

When the number of devices and things having different air interface technologies communicate each other, IPv6 is an ideal internet protocols owing to its large address space. Also, NFC would be one of the endpoints using IPv6. Therefore, This document describes how IPv6 is transmitted over NFC using 6LoWPAN techniques with following scopes.

- o Overview of NFC technologies;
- o Specifications for IPv6 over NFC;
 - * Neighbor Discovery;
 - * Addressing and Configuration;
 - * Header Compression;
 - * Fragmentation & Reassembly for a IPv6 datagram;

RFC4944 [1] specifies the transmission of IPv6 over IEEE 802.15.4. The NFC link also has similar characteristics to that of IEEE 802.15.4. Many of the mechanisms defined in the RFC4944 [1] can be applied to the transmission of IPv6 on NFC links. This document specifies the details of IPv6 transmission over NFC links.

2. Conventions and 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 [2].

3. Overview of Near Field Communication Technology

NFC technology enables simple and safe two-way interactions between electronic devices, allowing consumers to perform contactless transactions, access digital content, and connect electronic devices with a single touch. NFC complements many popular consumer level wireless technologies, by utilizing the key elements in existing standards for contactless card technology (ISO/IEC 14443 A&B and JIS-X 6319-4). NFC can be compatible with existing contactless card infrastructure and it enables a consumer to utilize one device across different systems.

Extending the capability of contactless card technology, NFC also enables devices to share information at a distance that is less than 10 cm with a maximum communication speed of 424 kbps. Users can share business cards, make transactions, access information from a smart poster or provide credentials for access control systems with a simple touch.

NFC's bidirectional communication ability is ideal for establishing connections with other technologies by the simplicity of touch. In addition to the easy connection and quick transactions, simple data sharing is also available.

3.1. Peer-to-peer Mode of NFC

NFC-enabled devices are unique in that they can support three modes of operation: card emulation, peer-to-peer, and reader/writer. Peer-to-peer mode enables two NFC-enabled devices to communicate with each other to exchange information and share files, so that users of NFC-enabled devices can quickly share contact information and other files with a touch. Therefore, a NFC-enabled device can securely send IPv6 packets to any corresponding node on the Internet when a NFC-enabled gateway is linked to the Internet.

3.2. Protocol Stacks of NFC

The IP protocol can use the services provided by Logical Link Control Protocol (LLCP) in the NFC stack to provide reliable, two-way transport of information between the peer devices. Figure 1 depicts the NFC P2P protocol stack with IPv6 bindings to the LLCP.

For data communication in IPv6 over NFC, an IPv6 packet SHALL be received at LLCP of NFC and transported to an Information Field in Protocol Data Unit (I PDU) of LLCP of the NFC-enabled peer device. Since LLCP does not support fragmentation and reassembly, upper layers SHOULD support fragmentation and reassembly. For IPv6 addressing or address configuration, LLCP SHALL provide related information, such as link layer addresses, to its upper layer. LLCP to IPv6 protocol Binding SHALL transfer the SSAP and DSAP value to the IPv6 over NFC protocol. SSAP stands for Source Service Access Point, which is 6-bit value meaning a kind of Logical Link Control (LLC) address, while DSAP means a LLC address of destination NFC-enabled device.

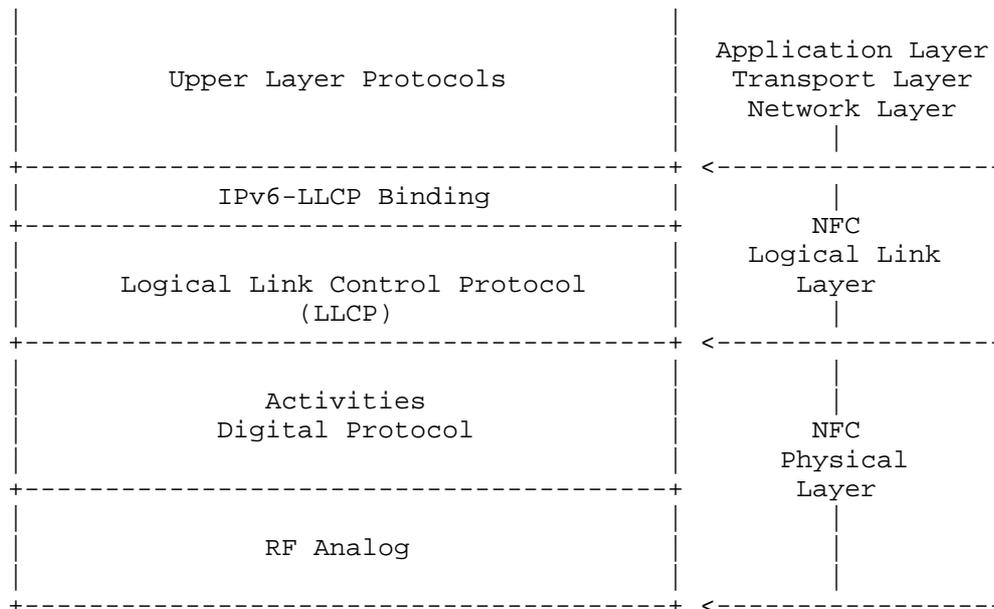


Figure 1: Protocol Stacks of NFC

The LLCP consists of Logical Link Control (LLC) and MAC Mapping. The MAC Mapping integrates an existing RF protocol into the LLCP architecture. The LLC contains three components, such as Link Management, Connection-oriented Transport, and Connection-less

Transport. The Link Management component is responsible for serializing all connection-oriented and connectionless LLC PDU (Protocol Data Unit) exchanges and for aggregation and disaggregation of small PDUs. This component also guarantees asynchronous balanced mode communication and provides link status supervision by performing the symmetry procedure. The Connection-oriented Transport component is responsible for maintaining all connection-oriented data exchanges including connection set-up and termination. The Connectionless Transport component is responsible for handling unacknowledged data exchanges.

3.3. NFC-enabled Device Addressing

NFC-enabled devices are identified by 6-bit LLC address. In other words, Any address SHALL be usable as both an SSAP and a DSAP address. According to NFCForum-TS-LLCP_1.1 [3], address values between 0 and 31 (00h - 1Fh) SHALL be reserved for well-known service access points for Service Discovery Protocol (SDP). Address values between 32 and 63 (20h - 3Fh) inclusively, SHALL be assigned by the local LLC as the result of an upper layer service request.

3.4. NFC MAC PDU Size and MTU

As mentioned in Section 3.2, an IPv6 packet SHALL be received at LLCP of NFC and transported to an Unnumbered Information Protocol Data Unit (UI PDU) and an Information Field in Protocol Data Unit (I PDU) of LLCP of the NFC-enabled peer device. The format of the UI PDU and I PDU SHALL be as shown in Figure 2 and Figure 3.

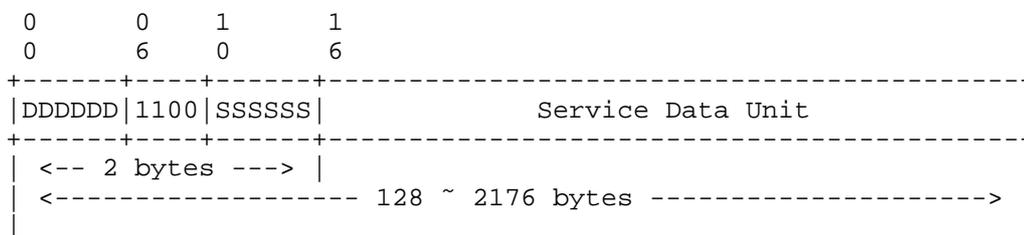


Figure 2: Format of the UI PDU in NFC

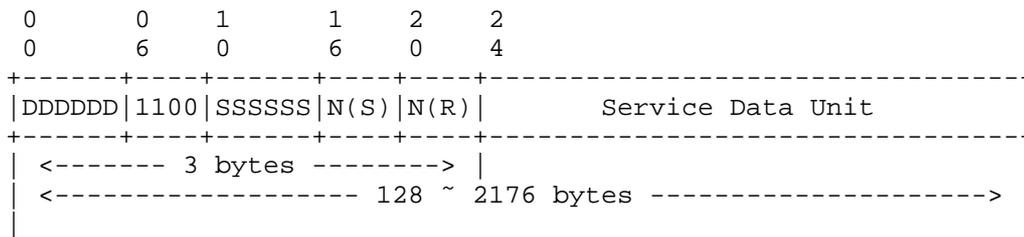


Figure 3: Format of the I PDU in NFC

The I PDU sequence field SHALL contain two sequence numbers: The send sequence number N(S) and the receive sequence number N(R). The send sequence number N(S) SHALL indicate the sequence number associated with this I PDU. The receive sequence number N(R) value SHALL indicate that I PDUs numbered up through N(R) - 1 have been received correctly by the sender of this I PDU and successfully passed to the senders SAP identified in the SSAP field. These I PDUs SHALL be considered as acknowledged.

The information field of an I PDU SHALL contain a single service data unit. The maximum number of octets in the information field SHALL be determined by the Maximum Information Unit (MIU) for the data link connection. The default value of the MIU for I PDUs SHALL be 128 octets. The local and remote LLCs each establish and maintain distinct MIU values for each data link connection endpoint. Also, An LLC MAY announce a larger MIU for a data link connection by transmitting an MIUX extension parameter within the information field. If no MIUX parameter is transmitted, the default MIU value of 128 SHALL be used. Otherwise, the MTU size in NFC LLCP SHALL calculate the MIU value as follows:

$$MIU = 128 + MIUX.$$

According to NFCForum-TS-LLCP_1.1 [3], format of the MIUX parameter TLV is as shown in Figure 4.

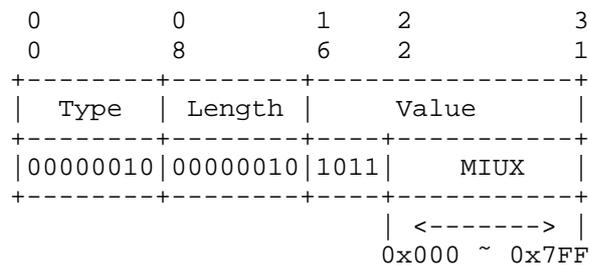


Figure 4: Format of the MIUX Parameter TLV

When the MIUX parameter is encoded as a TLV, the TLV Type field SHALL be 0x02 and the TLV Length field SHALL be 0x02. The MIUX parameter SHALL be encoded into the least significant 11 bits of the TLV Value field. The unused bits in the TLV Value field SHALL be set to zero by the sender and SHALL be ignored by the receiver. However, a maximum value of the TLV Value field can be 0x7FF, and a maximum size of the MTU in NFC LLCP SHALL calculate 2176 bytes.

4. Specification of IPv6 over NFC

NFC technology sets also has considerations and requirements owing to low power consumption and allowed protocol overhead. 6LoWPAN standards RFC4944 [1], RFC6775 [4], and RFC6282 [5] provide useful functionality for reducing overhead which can be applied to BT-LE. This functionality comprises of link-local IPv6 addresses and stateless IPv6 address auto-configuration (see Section 4.3), Neighbor Discovery (see Section 4.5) and header compression (see Section 4.7).

One of the differences between IEEE 802.15.4 and NFC is that the former supports both star and mesh topology (and requires a routing protocol), whereas NFC can support direct peer-to-peer connection and simple mesh-like topology depending on NFC application scenarios because of very short RF distance of 10 cm or less.

4.1. Protocol Stacks

Figure 5 illustrates IPv6 over NFC. Upper layer protocols can be transport protocols (TCP and UDP), application layer, and the others capable running on the top of IPv6.

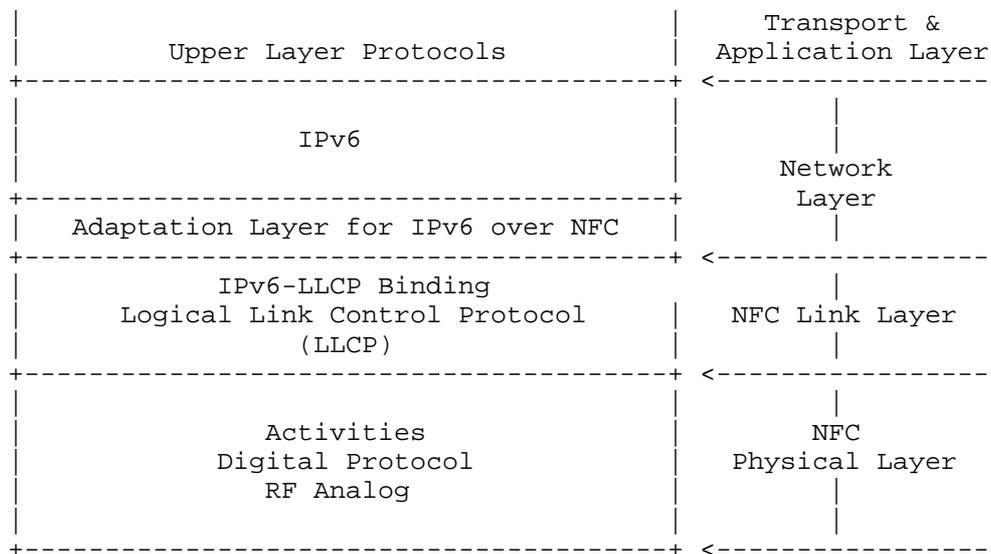


Figure 5: Protocol Stacks for IPv6 over NFC

Adaptation layer for IPv6 over NFC SHALL support neighbor discovery, address auto-configuration, header compression, and fragmentation & reassembly.

4.2. Link Model

In the case of BT-LE, Logical Link Control and Adaptation Protocol (L2CAP) supports fragmentation and reassembly (FAR) functionality; therefore, adaptation layer for IPv6 over BT-LE does not have to conduct the FAR procedure. The NFC LLCP, by contrast, does not support the FAR functionality, so IPv6 over NFC needs to consider the FAR functionality, defined in RFC4944 [1]. However, MTU on NFC link can be configured in a connection procedure and extended enough to fit the MTU of IPv6 packet.

The NFC link between two communicating devices is considered to be a point-to-point link only. Unlike in BT-LE, NFC link does not consider star topology and mesh network topology but peer-to-peer topology and simple multi-hop topology. Due to this characteristics, 6LoWPAN functionality, such as addressing and auto-configuration, and header compression, is specialized into NFC.

4.3. Stateless Address Autoconfiguration

A NFC-enabled device (i.e., 6LN) performs stateless address autoconfiguration as per RFC4862 [6]. A 64-bit Interface identifier (IID) for a NFC interface is formed by utilizing the 6-bit NFC LLCP address (i.e., SSAP or DSAP) (see Section 3.3). In the viewpoint of address configuration, such an IID MAY guarantee a stable IPv6 address because each data link connection is uniquely identified by the pair of DSAP and SSAP included in the header of each LLC PDU in NFC.

Following the guidance of RFC7136 [10], interface Identifiers of all unicast addresses for NFC-enabled devices are formed on the basis of 64 bits long and constructed in a modified EUI-64 format as shown in Figure 6.

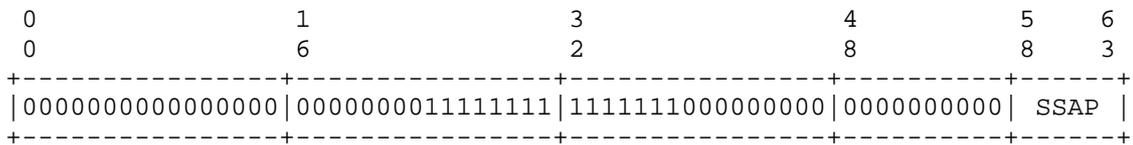


Figure 6: Formation of IID from NFC-enabled device address

In addition, the "Universal/Local" bit in the case of NFC-enabled device address MUST be set to 0 RFC4291 [7].

4.4. IPv6 Link Local Address

Only if the NFC-enabled device address is known to be a public address the "Universal/Local" bit can be set to 1. The IPv6 link-local address for a NFC-enabled device is formed by appending the IID, to the prefix FE80::/64, as depicted in Figure 7.

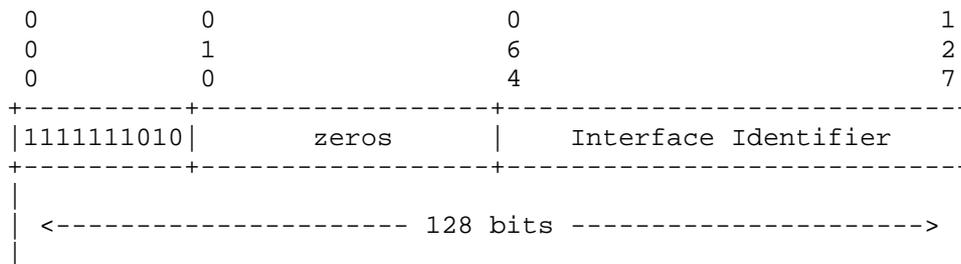


Figure 7: IPv6 link-local address in NFC

The tool for a 6LBR to obtain an IPv6 prefix for numbering the NFC network is can be accomplished via DHCPv6 Prefix Delegation (RFC3633 [8]).

4.5. Neighbor Discovery

Neighbor Discovery Optimization for 6LoWPANs (RFC6775 [4]) describes the neighbor discovery approach in several 6LoWPAN topologies, such as mesh topology. NFC does not consider complicated mesh topology but simple multi-hop network topology or directly connected peer-to-peer network. Therefore, the following aspects of RFC6775 are applicable to NFC:

1. In a case that a NFC-enabled device (6LN) is directly connected to 6LBR, A NFC 6LN MUST register its address with the 6LBR by sending a Neighbor Solicitation (NS) message with the Address Registration Option (ARO) and process the Neighbor Advertisement (NA) accordingly. In addition, DHCPv6 is used to assigned an address, Duplicate Address Detection (DAD) is not required.
2. For sending Router Solicitations and processing Router Advertisements the NFC 6LNs MUST follow Sections 5.3 and 5.4 of the RFC6775.

4.6. Dispatch Header

All IPv6-over-NFC encapsulated datagrams transmitted over NFC are prefixed by an encapsulation header stack consisting of a Dispatch value followed by zero or more header fields. The only sequence currently defined for IPv6-over-NFC is the LOWPAN_IPHC header followed by payload, as depicted in Figure 8.

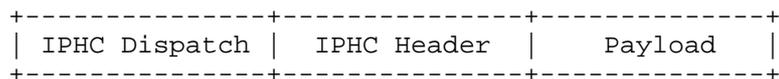


Figure 8: A IPv6-over-NFC Encapsulated 6LOWPAN_IPHC Compressed IPv6 Datagram

The dispatch value may be treated as an unstructured namespace. Only a single pattern is used to represent current IPv6-over-NFC functionality.

Pattern	Header Type	Reference
01 1xxxxx	6LOWPAN_IPHC	[RFC6282]

Figure 9: Dispatch Values

Other IANA-assigned 6LoWPAN Dispatch values do not apply to this specification.

4.7. Header Compression

Header compression as defined in RFC6282 [5] , which specifies the compression format for IPv6 datagrams on top of IEEE 802.15.4, is REQUIRED in this document as the basis for IPv6 header compression on top of NFC. All headers MUST be compressed according to RFC6282 encoding formats.

Therefore, IPv6 header compression in RFC6282 [5] MUST be implemented. Further, implementations MAY also support Generic Header Compression (GHC) of RFC7400 [11]. A node implementing GHC MUST probe its peers for GHC support before applying GHC.

If a 16-bit address is required as a short address of IEEE 802.15.4, it MUST be formed by padding the 6-bit NFC link-layer (node) address to the left with zeros as shown in Figure 10.

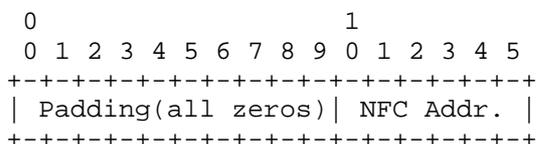


Figure 10: NFC short adress format

4.8. Fragmentation and Reassembly

NFC provides fragmentation and reassembly (FAR) for payloads from 128 bytes up to 2176 bytes as mention in Section 3.4. The MTU of a general IPv6 packet can fit into a sigle NFC link frame. Therefore, the FAR functionality as defined in RFC4944, which specifies the fragmentation methods for IPv6 datagrams on top of IEEE 802.15.4, is NOT REQUIRED in this document as the basis for IPv6 datagram FAR on top of NFC. The NFC link connection for IPv6 over NFC MUST be configured with an equivalent MIU size to fit the MTU of IPv6 Packet. However, the default configuration of MIUX value is 0x480 in order to fit the MTU (1280 bytes) of a IPv6 packet.

4.9. Unicast Address Mapping

The address resolution procedure for mapping IPv6 non-multicast addresses into NFC link-layer addresses follows the general description in Section 7.2 of RFC4861 [9], unless otherwise specified.

The Source/Target link-layer Address option has the following form when the addresses are 6-bit NFC link-layer (node) addresses.

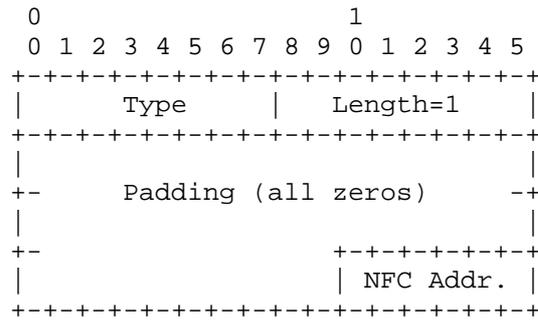


Figure 11: Unicast address mapping

Option fields:

Type:

- 1: for Source Link-layer address.
- 2: for Target Link-layer address.

Length:

This is the length of this option (including the type and length fields) in units of 8 octets. The value of this field is 1 for 6-bit NFC node addresses.

NFC address:

The 6-bit address in canonical bit order. This is the unicast address the interface currently responds to.

4.10. Multicast Address Mapping

All IPv6 multicast packets MUST be sent to NFC Destination Address, 0x3F (broadcast) and filtered at the IPv6 layer. When represented as a 16-bit address in a compressed header, it MUST be formed by padding

on the left with a zero. In addition, the NFC Destination Address, 0x3F, MUST not be used as a unicast NFC address of SSAP or DSAP.

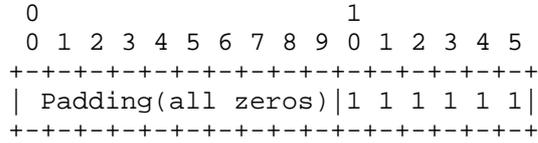


Figure 12: Multicast address mapping

5. Internet Connectivity Scenarios

As two typical scenarios, the NFC network can be isolated and connected to the Internet.

5.1. NFC-enabled Device Connected to the Internet

One of the key applications by using adaptation technology of IPv6 over NFC is the most securely transmitting IPv6 packets because RF distance between 6LN and 6LBR SHOULD be within 10 cm. If any third party wants to hack into the RF between them, it MUST come to nearly touch them. Applications can choose which kinds of air interfaces (e.g., BT-LE, Wi-Fi, NFC, etc.) to send data depending characteristics of data. NFC SHALL be the best solution for secured and private information.

Figure 13 illustrates an example of NFC-enabled device network connected to the Internet. Distance between 6LN and 6LBR SHOULD be 10 cm or less. If there is any of close laptop computers to a user, it SHALL becomes the 6LBR. Additionally, When the user mounts a NFC-enabled air interface adapter (e.g., portable small NFC dongle) on the close laptop PC, the user's NFC-enabled device (6LN) can communicate the laptop PC (6LBR) within 10 cm distance.



Figure 13: NFC-enabled device network connected to the Internet

5.2. Isolated NFC-enabled Device Network

In some scenarios, the NFC-enabled device network may transiently be a simple isolated network as shown in the Figure 14.

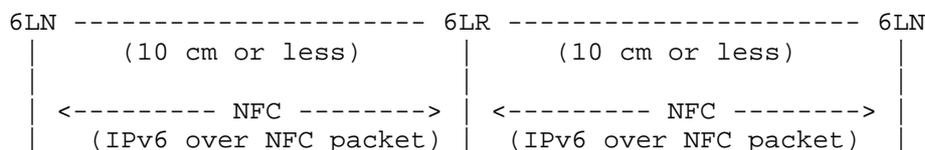


Figure 14: Isolated NFC-enabled device network

In mobile phone markets, applications are designed and made by user developers. They may image interesting applications, where three or more mobile phones touch or attach each other to accomplish outstanding performance. For instance, three or more mobile phones can play multi-channel sound of music together. In addition, attached three or more mobile phones can make an extended banner to show longer sentences in a concert hall.

6. IANA Considerations

There are no IANA considerations related to this document.

7. Security Considerations

When interface identifiers (IIDs) are generated, devices and users are required to consider mitigating various threats, such as correlation of activities over time, location tracking, device-specific vulnerability exploitation, and address scanning.

IPv6-over-NFC is, in practice, not used for long-lived links for big size data transfer or multimedia streaming, but used for extremely short-lived links (i.e., single touch-based approaches) for ID verification and mobile payment. This will mitigate the threat of correlation of activities over time.

IPv6-over-NFC uses an IPv6 interface identifier formed from a "Short Address" and a set of well-known constant bits (such as padding with '0's) for the modified EUI-64 format. However, the short address of NFC link layer (LLC) is not generated as a physically permanent value but logically generated for each connection. Thus, every single touch connection can use a different short address of NFC link with an extremely short-lived link. This can mitigate address scanning as well as location tracking and device-specific vulnerability exploitation.

However, malicious tries for one connection of a long-lived link with NFC technology are not secure, so the method of deriving interface identifiers from 6-bit NFC Link layer addresses is intended to preserve global uniqueness when it is possible. Therefore, it requires to protect from duplication through accident or forgery and to define a way to include sufficient bit of entropy in the IPv6 interface identifier, such as random EUI-64.

8. Acknowledgements

We are grateful to the members of the IETF 6lo working group.

Michael Richardson, Suresh Krishnan, Pascal Thubert, Carsten Bormann, and Alexandru Petrescu have provided valuable feedback for this draft.

9. References

9.1. Normative References

- [1] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [2] 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>>.
- [3] "Logical Link Control Protocol version 1.1", NFC Forum Technical Specification , June 2011.
- [4] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<http://www.rfc-editor.org/info/rfc6775>>.
- [5] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.
- [6] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<http://www.rfc-editor.org/info/rfc4862>>.

- [7] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<http://www.rfc-editor.org/info/rfc4291>>.
- [8] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, DOI 10.17487/RFC3633, December 2003, <<http://www.rfc-editor.org/info/rfc3633>>.
- [9] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<http://www.rfc-editor.org/info/rfc4861>>.
- [10] Carpenter, B. and S. Jiang, "Significance of IPv6 Interface Identifiers", RFC 7136, DOI 10.17487/RFC7136, February 2014, <<http://www.rfc-editor.org/info/rfc7136>>.
- [11] Bormann, C., "6LoWPAN-GHC: Generic Header Compression for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 7400, DOI 10.17487/RFC7400, November 2014, <<http://www.rfc-editor.org/info/rfc7400>>.

9.2. Informative References

- [12] "Near Field Communication - Interface and Protocol (NFCIP-1) 3rd Ed.", ECMA-340 , June 2013.

Authors' Addresses

Yong-Geun Hong
ETRI
161 Gajeong-Dong Yuseung-Gu
Daejeon 305-700
Korea

Phone: +82 42 860 6557
Email: yghong@etri.re.kr

Younghwan Choi
ETRI
218 Gajeongno, Yuseong
Daejeon 305-700
Korea

Phone: +82 42 860 1429
Email: yhc@etri.re.kr

Joo-Sang Youn
DONG-EUI University
176 Eomgwangno Busan_jin_gu
Busan 614-714
Korea

Phone: +82 51 890 1993
Email: joosang.youn@gmail.com

Dongkyun Kim
Kyungpook National University
80 Daehak-ro, Buk-gu
Daegu 702-701
Korea

Phone: +82 53 950 7571
Email: dongkyun@knu.ac.kr

JinHyouk Choi
Samsung Electronics Co.,
129 Samsung-ro, Youngdong-gu
Suwon 447-712
Korea

Phone: +82 2 2254 0114
Email: jinchoe@samsung.com

6lo
Internet-Draft
Updates: 4944 (if approved)
Intended status: Standards Track
Expires: July 18, 2016

P. Thubert, Ed.
Cisco
January 15, 2016

6LoWPAN Paging Dispatch
draft-ietf-6lo-paging-dispatch-01

Abstract

This specification introduces a new context switch mechanism for 6LoWPAN compression, expressed in terms of Pages and signaled by a new Paging Dispatch.

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 July 18, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Terminology	3
3. Updating RFC 4944	3
4. Page 1 Paging Dispatch	4
5. Security Considerations	4
6. IANA Considerations	4
6.1. Consuming Dispatch Types	5
6.2. New Per-Page Dispatch Type registries	5
7. Acknowledgments	5
8. References	5
8.1. Normative References	5
8.2. Informative References	6
Author's Address	6

1. Introduction

The design of Low Power and Lossy Networks (LLNs) is generally focused on saving energy, which often is a very constrained resource. Other constraints, such as memory capacity and duty cycle restrictions on LLN devices, usually derive from that primary concern. Energy is often available only from primary batteries that are expected to last for years, or is scavenged from the environment in very limited amounts. Any protocol that is intended for use in LLNs must be designed with a primary focus on saving energy, which is a strict requirement.

Controlling the amount of data transmission is one possible means of saving energy. In a number of LLN standards, the frame size is limited to much smaller values than the IPv6 maximum transmission unit (MTU) of 1280 bytes. In particular, an LLN that relies on the classical Physical Layer (PHY) of IEEE 802.15.4 [IEEE802154] is limited to 127 bytes per frame. The need to compress IPv6 packets over IEEE 802.15.4 led to the 6LoWPAN Header Compression [RFC6282] work (6LoWPAN-HC).

As more and more protocols need to be compressed, the encoding capabilities of the original dispatch defined in the 6lo adaptation layer framework ([RFC4944],[RFC6282]) becomes saturated. This specification introduces a new context switch mechanism for 6LoWPAN compression, expressed in terms of Pages and signaled by a new Paging Dispatch mechanism.

2. Terminology

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

The Terminology used in this document is consistent with and incorporates that described in Terms Used in Routing for Low-Power and Lossy Networks [RFC7102] and Terminology for Constrained-Node Networks [RFC7228].

3. Updating RFC 4944

This draft adapts 6LoWPAN while maintaining backward compatibility with IPv6 over IEEE 802.15.4 [RFC4944] by introducing a concept of a "parsing context" in the 6LoWPAN parser, a context being identified by a Page Number. This specification defines 16 Pages.

Pages are delimited in a 6LoWPAN packet by a Paging Dispatch value that indicates the next current Page. The Page Number is encoded in a Paging Dispatch with the Value Bit Pattern of 1111xxxx where xxxx is the Page Number, 0 to 15, as described in Figure 1:

```

0
0 1 2 3 4 5 6 7
+---+---+---+---+
|1|1|1|1|Page Nb|
+---+---+---+---+

```

Figure 1: Paging Dispatch with Page Number Encoding.

Values of the Dispatch byte defined in [RFC4944] are considered as belonging to the Page 0 parsing context, which is the default and does not need to be signaled explicitly at the beginning of a 6LoWPAN packet. This ensures backward compatibility with existing implementations of 6LoWPAN.

The Dispatch bits defined in Page 0 by [RFC4944] are free to be reused in Pages 1 to 15. This specification allocates some values in Page 1 in Section 4 and leaves the rest open for future allocations.

Note: This specification does not use the Escape Dispatch, which extends Page 0 to more values, but rather allocates another Dispatch Bit Pattern (1111xxxx) for a new Paging Dispatch, that is present in all Pages, including Page 0 and Pages defined in future specifications, to indicate the next parsing context represented by its Page Number. The rationale for avoiding that approach is that

there can be multiple occurrences of a new header indexed by this specification in a single frame and the overhead on an octet each time for the Escape Dispatch would be prohibitive.

A Page (say Page N) is said to be active once the Page N Paging Dispatch is parsed, and as long as no other Paging Dispatch is parsed.

4. Page 1 Paging Dispatch

This specification defines some special properties for Page 1, detailed below:

The Dispatch bits defined for LOWPAN_IPHC by the Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks [RFC6282] are defined with the same values in Page 1 so there is no need to switch context from Page 1 to Page 0 to decode a packet that is encoded per [RFC6282].

Mesh Headers represent Layer-2 information and are processed before any Layer-3 information that is encoded in Page 1. If a 6LoWPAN packet requires a Mesh header, the Mesh Header MUST always be placed in the packet before the first Page 1 Paging Dispatch, if any.

For the same reason, Fragment Headers as defined in [RFC4944] MUST always be placed in the packet before the first Page 1 Paging Dispatch, if any.

The NALP Dispatch Bit Pattern as defined in [RFC4944] is only defined for the first octet in the packet. Switching back to Page 0 for NALP inside a 6LoWPAN packet does not make sense.

As a result, there is no need so far for restoring the Page 0 parsing context after a context was switched to Page 1, so the value for the Page 0 Paging Dispatch of 11110000 may not actually occur in those packets that adhere to 6LoWPAN specifications available at the time of writing this specification.

5. Security Considerations

The security considerations of [RFC4944] and [RFC6282] apply.

6. IANA Considerations

6.1. Consuming Dispatch Types

This document allocates 16 values from the Dispatch type field registry that was created for [RFC4944]. The allocated values are from 11 110000 through 11 111111 and represent Page Numbers 0 through 15 as discussed in this document.

6.2. New Per-Page Dispatch Type registries

This document creates 15 new IANA registries for the Per-Page Dispatch type fields, indexed by Page Number, 1 to 15. Each Registry corresponds to a bit-field of one octet.

Future assignments in these registries are to be coordinated via IANA under the policy of "Specification Required" [RFC2434]. It is expected that this policy will allow for other (non-IETF) organizations to more easily obtain assignments.

These registries extend the Dispatch type field registry that was created for [RFC4944], which is considered as the registry for Page 0.

As described above, this document allocates in the registry associated to Page 1 the Per-Page Dispatch type field values that are allocated in the Dispatch type field for LOWPAN_IPHC by [RFC6282]. Those values are from 01 100000 through 01 111111 and they have the same definition in Page 1 as they do in Page 0, meaning that the registries for Page 0 and Page 1 are an exact overlap in this range.

7. Acknowledgments

The authors wish to thank Thomas Watteyne, Tengfei Chang, Martin Turon, James Woodyatt, Samita Chakrabarti, Jonathan Hui, Gabriel Montenegro and Ralph Droms for constructive reviews to the design in the 6lo Working Group.

8. References

8.1. Normative References

[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", 2015.

- [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>>.
- [RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 2434, DOI 10.17487/RFC2434, October 1998, <<http://www.rfc-editor.org/info/rfc2434>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.

8.2. Informative References

- [RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102, DOI 10.17487/RFC7102, January 2014, <<http://www.rfc-editor.org/info/rfc7102>>.
- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<http://www.rfc-editor.org/info/rfc7228>>.

Author's Address

Pascal Thubert (editor)
Cisco Systems
Building D - Regus
45 Allee des Ormes
BP1200
MOUGINS - Sophia Antipolis 06254
FRANCE

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

roll
Internet-Draft
Intended status: Standards Track
Expires: September 22, 2016

P. Thubert, Ed.
Cisco
C. Bormann
Uni Bremen TZI
L. Toutain
IMT-TELECOM Bretagne
R. Cragie
ARM
March 21, 2016

6LoWPAN Routing Header
draft-ietf-roll-routing-dispatch-00

Abstract

This specification introduces a new 6LoWPAN dispatch type for use in 6LoWPAN Route-Over topologies, that initially covers the needs of RPL (RFC6550) data packets compression. Using this dispatch type, this specification defines a method to compress RPL Option (RFC6553) information and Routing Header type 3 (RFC6554), an efficient IP-in-IP technique and is extensible for more applications.

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 22, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Terminology	6
3.	Using the Page Dispatch	6
3.1.	New Routing Header Dispatch (6LoRH)	6
3.2.	Placement Of 6LoRH headers	6
3.2.1.	Relative To Non-6LoRH Headers	7
3.2.2.	Relative To Other 6LoRH Headers	7
4.	6LoWPAN Routing Header General Format	8
4.1.	Elective Format	8
4.2.	Critical Format	9
4.3.	Compressing Addresses	9
4.3.1.	Coalescence	10
4.3.2.	DODAG Root Address Determination	10
5.	The SRH 6LoRH Header	11
5.1.	Encoding	11
5.2.	SRH-6LoRH General Operation	13
5.2.1.	Uncompressed SRH Operation	13
5.2.2.	6LoRH-Compressed SRH Operation	13
5.2.3.	Inner LOWPAN_IPHC Compression	14
5.3.	The Design Point of Popping Entries	14
5.4.	Compression Reference for SRH-6LoRH header entries	15
5.5.	Popping Headers	16
5.6.	Forwarding	17
6.	The RPL Packet Information 6LoRH	17
6.1.	Compressing the RPLInstanceID	19
6.2.	Compressing the SenderRank	19
6.3.	The Overall RPI-6LoRH encoding	19
7.	The IP-in-IP 6LoRH Header	22
8.	Security Considerations	23
9.	IANA Considerations	23
9.1.	Reserving Space in 6LoWPAN Dispatch Page 1	23
9.2.	New 6LoWPAN Routing Header Type Registry	24
10.	Acknowledgments	24
11.	References	24
11.1.	Normative References	24
11.2.	Informative References	25
Appendix A.	Examples	26
A.1.	Examples Compressing The RPI	26
A.2.	Example Of Downward Packet In Non-Storing Mode	28

A.3. Example of SRH-6LoRH life-cycle	30
Authors' Addresses	32

1. Introduction

The design of Low Power and Lossy Networks (LLNs) is generally focused on saving energy, a very constrained resource in most cases. The other constraints, such as the memory capacity and the duty cycling of the LLN devices, derive from that primary concern. Energy is often available from primary batteries that are expected to last for years, or is scavenged from the environment in very limited quantities. Any protocol that is intended for use in LLNs must be designed with the primary concern of saving energy as a strict requirement.

Controlling the amount of data transmission is one possible venue to save energy. In a number of LLN standards, the frame size is limited to much smaller values than the IPv6 maximum transmission unit (MTU) of 1280 bytes. In particular, an LLN that relies on the classical Physical Layer (PHY) of IEEE 802.15.4 [IEEE802154] is limited to 127 bytes per frame. The need to compress IPv6 packets over IEEE 802.15.4 led to the 6LoWPAN Header Compression [RFC6282] work (6LoWPAN-HC).

Innovative Route-over techniques have been and are still being developed for routing inside a LLN. In a general fashion, such techniques require additional information in the packet to provide loop prevention and to indicate information such as flow identification, source routing information, etc.

For reasons such as security and the capability to send ICMP errors back to the source, an original packet must not be tampered with, and any information that must be inserted in or removed from an IPv6 packet must be placed in an extra IP-in-IP encapsulation. This is the case when the additional routing information is inserted by a router on the path of a packet, for instance a mesh root, as opposed to the source node. This is also the case when some routing information must be removed from a packet that flows outside the LLN. When to use RFC 6553, 6554 and IPv6-in-IPv6 [I-D.robles-roll-useofrplinfo] details different cases where RFC 6553, RFC 6554 and IPv6-in-IPv6 encapsulation is required to set the bases to help defining the compression of RPL routing information in LLN environments.

When using [RFC6282] the outer IP header of an IP-in-IP encapsulation may be compressed down to 2 octets in stateless compression and down to 3 octets in stateful compression when context information must be added.

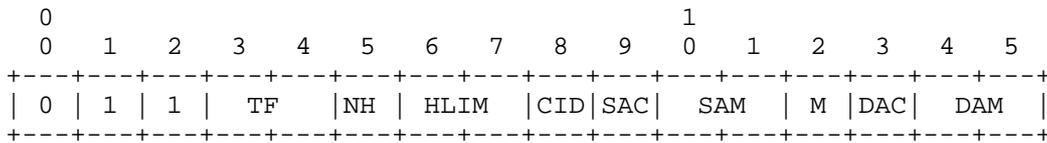


Figure 1: LOWPAN_IPHC base Encoding (RFC6282).

The Stateless Compression of an IPv6 addresses can only happen if the IPv6 address can be deduced from the MAC addresses, meaning that the IP end point is also the MAC-layer endpoint. This is generally not the case in a RPL network which is generally a multi-hop route-over (i.e., operated at Layer-3) network. A better compression, which does not involve variable compressions depending on the hop in the mesh, can be achieved based on the fact that the outer encapsulation is usually between the source (or destination) of the inner packet and the root. Also, the inner IP header can only be compressed by [RFC6282] if all the fields preceding it are also compressed. This specification makes the inner IP header the first header to be compressed by [RFC6282], and keeps the inner packet encoded the same way whether it is encapsulated or not, thus preserving existing implementations.

As an example, the Routing Protocol for Low Power and Lossy Networks [RFC6550] (RPL) is designed to optimize the routing operations in constrained LLNs. As part of this optimization, RPL requires the addition of RPL Packet Information (RPI) in every packet, as defined in Section 11.2 of [RFC6550].

The RPL Option for Carrying RPL Information in Data-Plane Datagrams [RFC6553] specification indicates how the RPI can be placed in a RPL Option (RPL-OPT) that is placed in an IPv6 Hop-by-Hop header.

This representation demands a total of 8 bytes, while in most cases the actual RPI payload requires only 19 bits. Since the Hop-by-Hop header must not flow outside of the RPL domain, it must be inserted in packets entering the domain and be removed from packets that leave the domain. In both cases, this operation implies an IP-in-IP encapsulation.

Additionally, in the case of the Non-Storing Mode of Operation (MOP), RPL requires a Source Routing Header (SRH) in all packets that are routed down a RPL graph. For that purpose, the [IPv6 Routing Header for Source Routes with RPL] (#RFC6554) specification defines the type 3 Routing Header for IPv6 (RH3).

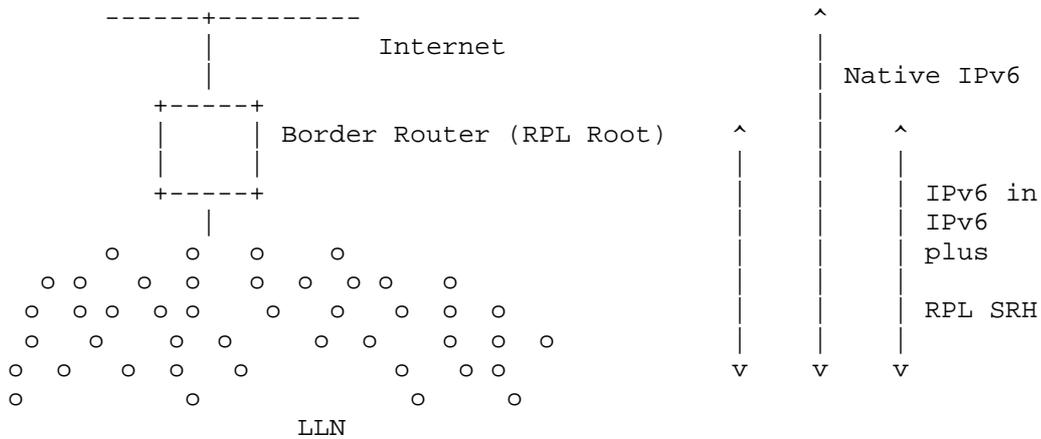


Figure 2: IP-in-IP Encapsulation within the LLN.

With Non-Storing RPL, even if the source is a node in the same LLN, the packet must first reach up the graph to the root so that the root can insert the SRH to go down the graph. In any fashion, whether the packet was originated in a node in the LLN or outside the LLN, and regardless of whether the packet stays within the LLN or not, as long as the source of the packet is not the root itself, the source-routing operation also implies an IP-in-IP encapsulation at the root in order to insert the SRH.

6TiSCH [I-D.ietf-6tisch-architecture] specifies the operation of IPv6 over the TimeSlotted Channel Hopping [RFC7554] (TSCH) mode of operation of IEEE 802.15.4. The architecture requires the use of both RPL and the 6lo adaptation layer over IEEE 802.15.4. Because it inherits the constraints on frame size from the MAC layer, 6TiSCH cannot afford to allocate 8 bytes per packet on the RPI. Hence the requirement for 6LoWPAN header compression of the RPI.

An extensible compression technique is required that simplifies IP-in-IP encapsulation when it is needed, and optimally compresses existing routing artifacts found in RPL LLNs.

This specification extends the 6lo adaptation layer framework ([RFC4944],[RFC6282]) so as to carry routing information for route-over networks based on RPL. The specification includes the formats necessary for RPL and is extensible for additional formats.

2. Terminology

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

The Terminology used in this document is consistent with and incorporates that described in 'Terminology in Low power And Lossy Networks' [RFC7102] and [RFC6550].

The terms Route-over and Mesh-under are defined in [RFC6775].

Other terms in use in LLNs are found in [RFC7228].

The term "byte" is used in its now customary sense as a synonym for "octet".

3. Using the Page Dispatch

The 6LoWPAN Paging Dispatch [I-D.ietf-6lo-paging-dispatch] specification extends the 6lo adaptation layer framework ([RFC4944], [RFC6282]) by introducing a concept of "context" in the 6LoWPAN parser, a context being identified by a Page number. The specification defines 16 Pages.

This draft operates within Page 1, which is indicated by a Dispatch Value of binary 11110001.

3.1. New Routing Header Dispatch (6LoRH)

This specification introduces a new 6LoWPAN Routing Header (6LoRH) to carry IPv6 routing information. The 6LoRH may contain source routing information such as a compressed form of SRH, as well as other sorts of routing information such as the RPI and IP-in-IP encapsulation.

The 6LoRH is expressed in a 6LoWPAN packet as a Type-Length-Value (TLV) field, which is extensible for future use.

This specification uses the bit pattern 10xxxxxx in Page 1 for the new 6LoRH Dispatch. Section 4 describes how RPL artifacts in data packets can be compressed as 6LoRH headers.

3.2. Placement Of 6LoRH headers

3.2.1. Relative To Non-6LoRH Headers

In a zone of a packet where Page 1 is active (i.e., once a Page 1 Paging Dispatch is parsed and no subsequent Paging Dispatch has been parsed, the parsing of the packet MUST follow this specification if the 6LoRH Bit Pattern Section 3.1 is found.

With this specification, the 6LoRH Dispatch is only defined in Page 1, so it MUST be placed in the packet in a zone where the Page 1 context is active.

Because a 6LoRH header requires a Page 1 context, it MUST always be placed after any Fragmentation Header and/or Mesh Header [RFC4944].

A 6LoRH header MUST always be placed before the LOWPAN_IPHC as defined in 6LoWPAN Header Compression [RFC6282]. It is designed in such a fashion that placing or removing a header that is encoded with 6LoRH does not modify the part of the packet that is encoded with LOWPAN_IPHC, whether there is an IP-in-IP encapsulation or not. For instance, the final destination of the packet is always the one in the LOWPAN_IPHC whether there is a Routing Header or not.

3.2.2. Relative To Other 6LoRH Headers

IPv6 [RFC2460] defines chains of headers that are introduced by an IPv6 header and terminated by either another IPv6 header (IP-in-IP) or an Upper Layer Protocol (ULP) header. When an outer header is stripped from the packet, the whole chain goes with it. When one or more header(s) are inserted by an intermediate router, that router normally chains the headers and encapsulates the result in IP-in-IP.

With this specification, the chains of headers MUST be compressed in the same order as they appear in the uncompressed form of the packet. This means that if there is more than one nested IP-in-IP encapsulations, the first IP-in-IP encapsulation, with all its chain of headers, is encoded first in the compressed form.

In the compressed form of a packet that has SRH or HbH headers after the inner IPv6 header (e.g. if there is no IP-in-IP encapsulation), these headers are placed in the 6LoRH form before the 6LoWPAN-IPHC that represents the IPv6 header Section 3.2.1. If this packet gets encapsulated and some other SRH or HbH headers are added as part of the encapsulation, placing the 6LoRH headers next to one another may present an ambiguity on which header belong to which chain in the uncompressed form.

In order to disambiguate the headers that follow the inner IPv6 header in the uncompressed form from the headers that follow the

outer IP-in-IP header, it is REQUIRED that the compressed IP-in-IP header is placed last in the encoded chain. This means that the 6LoRH headers that are found after the last compressed IP-in-IP header are to be inserted after the IPv6 header that is encoded with the 6LOWPAN-IPHC when decompressing the packet.

With regards to the relative placement of the SRH and the RPI in the compressed form, it is a design point for this specification that the SRH entries are consumed as the packet progresses down the LLN Section 5.3. In order to make this operation simpler in the compressed form, it is REQUIRED that in the compressed form, the addresses along the source route path are encoded in the order of the path, and that the compressed SRH are placed before the compressed RPI.

4. 6LoWPAN Routing Header General Format

The 6LoRH uses the Dispatch Value Bit Pattern of 10xxxxxx in Page 1.

The Dispatch Value Bit Pattern is split in two forms of 6LoRH:

Elective (6LoRHE) that may be skipped if not understood

Critical (6LoRHC) that may not be ignored

4.1. Elective Format

The 6LoRHE uses the Dispatch Value Bit Pattern of 101xxxxx. A 6LoRHE may be ignored and skipped in parsing. If it is ignored, the 6LoRHE is forwarded with no change inside the LLN.

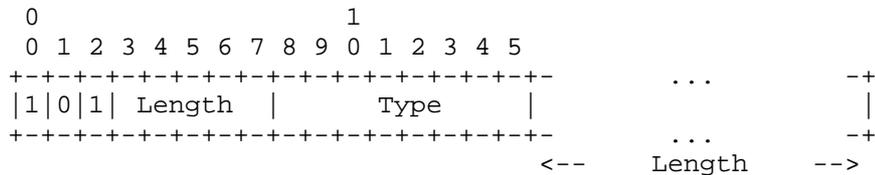


Figure 3: Elective 6LoWPAN Routing Header.

Length:

Length of the 6LoRHE expressed in bytes, excluding the first 2 bytes. This enables a node to skip a 6LoRHE header that it does not support and/or cannot parse, for instance if the Type is not recognized.

Type:

Type of the 6LoRHE

4.2. Critical Format

The 6LoRHC uses the Dispatch Value Bit Pattern of 100xxxxx.

A node which does not support the 6LoRHC Type MUST silently discard the packet.

Note: The situation where a node receives a message with a Critical 6LoWPAN Routing Header that it does not understand is a critical administrative error whereby the wrong device is placed in a network. It makes no sense to overburden the constrained device with code that would send an ICMP error to the source. Rather, it is expected that the device will raise some management alert indicating that it cannot operate in this network for that reason. As a result, there is no provision for the exchange of error messages for this situation, so it should be avoided by judicious use of administrative control and/or capability indications by the device manufacturer.

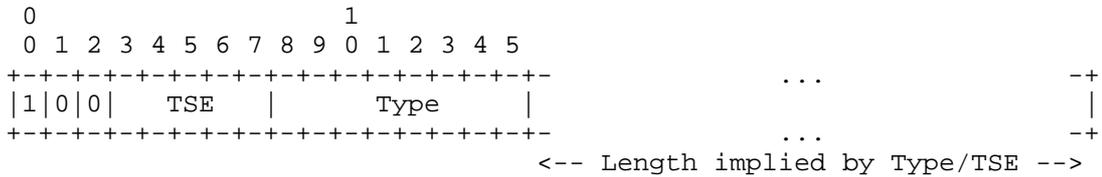


Figure 4: Critical 6LoWPAN Routing Header.

TSE:

Type Specific Extension. The meaning depends on the Type, which must be known in all of the nodes. The interpretation of the TSE depends on the Type field that follows. For instance, it may be used to transport control bits, the number of elements in an array, or the length of the remainder of the 6LoRHC expressed in a unit other than bytes.

Type:

Type of the 6LoRHC

4.3. Compressing Addresses

The general technique used in this draft to compress an address is first to determine a reference that as a long prefix match with this address, and then elide that matching piece. In order to reconstruct the compress address, the receiving node will perform the process of coalescence described in section Section 4.3.1.

One possible reference is the root of the RPL DODAG that is being traversed. It is used by 6LoRH as the reference to compress an outer

messages. In the compressed form of the packet, the context can be signaled in a Context Identifier Extension.

With this specification, the compression information is provided to the stack by RPL, and RPL exchanges it through the DODAGID field in the DAG Information Object (DIO) messages, as described in more details below. In the compressed form of the packet, the context can be signaled in by the RPLInstanceID in the RPI.

With RPL [RFC6550], the address of the DODAG root is known from the DODAGID field of the DIO messages. For a Global Instance, the RPLInstanceID that is present in the RPI is enough information to identify the DODAG that this node participates to and its associated root. But for a Local Instance, the address of the root MUST be explicit, either in some device configuration or signaled in the packet, as the source or the destination address, respectively.

When implicit, the address of the DODAG root MUST be determined as follows:

If the whole network is a single DODAG then the root can be well-known and does not need to be signaled in the packets. But since RPL does not expose that property, it can only be known by a configuration applied to all nodes.

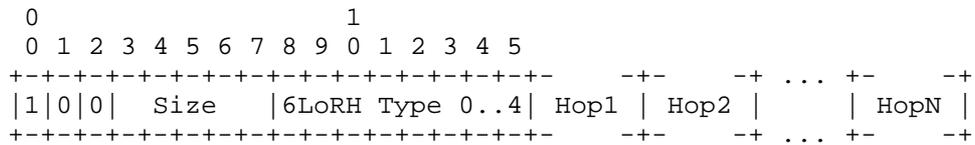
Else, the router that encapsulates the packet and compresses it with this specification MUST also place an RPI in the packet as prescribed by [RFC6550] to enable the identification of the DODAG. The RPI must be present even in the case when the router also places an SRH header in the packet.

It is expected that the RPL implementation maintains an abstract context table, indexed by Global RPLInstanceID, that provides the address of the root of the DODAG that this nodes participates to for that particular RPL Instance.

5. The SRH 6LoRH Header

5.1. Encoding

The Source Routing Header 6LoRH (SRH-6LoRH) header is a Critical 6LoWPAN Routing Header that provides a compressed form for the SRH, as defined in [RFC6554] for use by RPL routers. Routers that need to forward a packet with a SRH-6LoRH are expected to be RPL routers and are expected to support this specification. If a non-RPL router receives a packet with a SRH-6LoRH, this means that there was a routing error and the packet should be dropped so the Type cannot be ignored.



Size indicates the number of compressed addresses

Figure 6: The SRH-6LoRH.

The 6LoRH Type indicates the compression level used in a given SRH-6LoRH header.

One or more 6LoRH header(s) MAY be placed in a 6LoWPAN packet.

It results that all addresses in a given SRH-6LoRH header MUST be compressed in an identical fashion, down to using the identical number of bytes per address. In order to get different degrees of compression, multiple consecutive SRH-6LoRH headers MUST be used.

Type 0 means that the address is compressed down to one byte, whereas Type 4 means that the address is provided in full in the SRH-6LoRH with no compression. The complete list of Types of SRH-6LoRH and the corresponding compression level are provided in Figure 7:

6LoRH Type	Length of compressed IPv6 address (bytes)
0	1
1	2
2	4
3	8
4	16

Figure 7: The SRH-6LoRH Types.

In the case of a SRH-6LoRH header, the TSE field is used as a Size, which encodes the number of hops minus 1; so a Size of 0 means one hop, and the maximum that can be encoded is 32 hops. (If more than 32 hops need to be expressed, a sequence of SRH-6LoRH elements can be employed.) It results that the Length in bytes of a SRH-6LoRH header is:

$$2 + \text{Length_of_compressed_IPv6_address} * (\text{Size} + 1)$$

5.2. SRH-6LoRH General Operation

5.2.1. Uncompressed SRH Operation

In the non-compressed form, when the root generates or forwards a packet in non-Storing Mode, it needs to include a Source Routing Header [RFC6554] to signal a strict source-route path to a final destination down the DODAG.

All the hops along the path, but the first one, are encoded in order in the SRH. The last entry in the SRH is the final destination and the destination in the IPv6 header is the first hop along the source-route path. The intermediate hops perform a swap and the Segment-Left field indicates the active entry in the Routing Header [RFC2460].

The current destination of the packet, which is the termination of the current segment, is indicated at all times by the destination address of the IPv6 header.

5.2.2. 6LoRH-Compressed SRH Operation

The handling of the SRH-6LoRH is different: there is no swap, and a forwarding router that corresponds to the first entry in the first SRH-6LoRH upon reception of a packet effectively consumes that entry when forwarding. This means that the size of a compressed source-routed packet decreases as the packet progresses along its path and that the routing information is lost along the way. This also means that an SRH encoded with 6LoRH is not recoverable and cannot be protected.

When compressed with this specification, all the remaining hops MUST be encoded in order in one or more consecutive SRH-6LoRH headers. Whether or not there is a SRH-6LoRH header present, the address of the final destination is indicated in the LoWPAN_IPHC at all times along the path. Examples of this are provided in Appendix A.

The current destination (termination of the current segment) for a compressed source-routed packet is indicated in the first entry of the first SRH-6LoRH. In strict source-routing, that entry MUST match an address of the router that receives the packet.

The last entry in the last SRH-6LoRH is the last router on the way to the final destination in the LLN. This router can be the final destination if it is found desirable to carry a whole IP-in-IP encapsulation all the way. Else, it is the RPL parent of the final destination, or a router acting at 6LR [RFC6775] for the destination host, and advertising the host as an external route to RPL.

If the SRH-6LoRH header is contained in an IP-in-IP encapsulation, the last router removes the whole chain of headers. Otherwise, it removes the SRH-6LoRH header only.

5.2.3. Inner LOWPAN_IPHC Compression

6LoWPAN ND [RFC6282] is designed to support more than one IPv6 address per node and per Interface Identifier (IID), an IID being typically derived from a MAC address to optimize the LOWPAN-IPHC compression.

Link local addresses are compressed with stateless address compression (S/DAC=0). The other addresses are derived from different prefixes and they can be compressed with stateful address compression based on a context (S/DAC=1).

But stateless compression is only defined for the specific link-local prefix as opposed to the prefix in an encapsulating header. And with stateful compression, the compression reference is found in a context, as opposed to an encapsulating header.

It results that in the case of an IP-in-IP encapsulation, it is possible to compress an inner source (respectively destination) IP address in a LOWPAN_IPHC based on the encapsulating IP header only if stateful (context-based) compression is used. The compression will operate only if the IID in the source (respectively the destination) IP address in the outer and inner headers match, which usually means that they refer to the same node. This is encoded as S/DAC = 1 and S/AM=11. It must be noted that the outer destination address that is used to compress the inner destination address is the last entry in the last SRH-6LoRH header.

5.3. The Design Point of Popping Entries

In order to save energy and to optimize the chances of transmission success on lossy media, it is a design point for this specification that the entries in the SRH that have been used are removed from the packet. This creates a discrepancy from the art of IPv6 where Routing Header are mutable but recoverable.

With this specification, the packet can be expanded at any hop into a valid IPv6 packet, including a SRH, and compressed back. But the packet as decompressed along the way will not carry all the consumed addresses that packet would have if it had been forwarded in the uncompressed form.

It is noted that:

The value of keeping the whole RH in an IPv6 header is for the receiver to reverse it to use the symmetrical path on the way back.

It is generally not a good idea to reverse a routing header. The RH may have been used to stay away from the shortest path for some reason that is only valid on the way in (segment routing).

There is no use of reversing a RH in the present RPL specifications.

P2P RPL reverses a path that was learned reactively, as a part of the protocol operation, which is probably a cleaner way than a reversed echo on the data path.

Reversing a header is discouraged by [RFC2460] for RH0 unless it is authenticated, which requires an Authentication Header (AH). There is no definition of an AH operation for SRH, and there is no indication that the need exists in LLNs.

It is noted that AH does not protect the RH on the way. AH is a validation at the receiver with the sole value of enabling the receiver to reversing it.

A RPL domain is usually protected by L2 security and that secures both RPL itself and the RH in the packets, at every hop. This is a better security than that provided by AH.

In summary, the benefit of saving energy and lowering the chances of loss by sending smaller frames over the LLN are seen as overwhelming compared to the value of possibly reversing the header.

5.4. Compression Reference for SRH-6LoRH header entries

In order to optimize the compression of IP addresses present in the SRH headers, this specification requires that the 6LoWPAN layer identifies an address that is used as reference for the compression.

With this specification, the Compression Reference for the first address found in an SRH header is the source of the IPv6 packet, and then the reference for each subsequent entry is the address of its predecessor once it is uncompressed.

With RPL [RFC6550], an SRH header may only be present in Non-Storing mode, and it may only be placed in the packet by the root of the DODAG, which must be the source of the resulting IPv6 packet [RFC2460]. In this case, the address used as Compression Reference

is that the address of the root, and it can be implicit when the address of the root is.

The Compression Reference MUST be determined as follows:

The reference address may be obtained by configuration. The configuration may indicate either the address in full, or the identifier of a 6LoWPAN Context that carries the address [RFC6775], for instance one of the 16 Context Identifiers used in LOWPAN-IPHC [RFC6282].

Else, and if there is no IP-in-IP encapsulation, the source address in the IPv6 header that is compressed with LOWPAN-IPHC is the reference for the compression.

Else, and if the IP-in-IP compression specified in this document is used and the Encapsulator Address is provided, then the Encapsulator Address is the reference.

5.5. Popping Headers

Upon reception, the router checks whether the address in the first entry of the first SRH-6LoRH one of its own addresses. In that case, router MUST consume that entry before forwarding, which is an action of popping from a stack, where the stack is effectively the sequence of entries in consecutive SRH-6LoRH headers.

Popping an entry of an SRH-6LoRH header is a recursive action performed as follows:

If the Size of the SRH-6LoRH header is 1 or more, indicating that there are at least 2 entries in the header, the router removes the first entry and decrements the Size (by 1).

Else (meaning that this is the last entry in the SRH-6LoRH header), and if there is no next SRH-6LoRH header after this then the SRH-6LoRH is removed.

Else, if there is a next SRH-6LoRH of a Type with a larger or equal value, meaning a same or lesser compression yielding same or larger compressed forms, then the SRH-6LoRH is removed.

Else, the first entry of the next SRH-6LoRH is popped from the next SRH-6LoRH and coalesced with the first entry of this SRH-6LoRH.

At the end of the process, if there is no more SRH-6LoRH in the packet, then the processing node is the last router along the source route path.

5.6. Forwarding

When receiving a packet with a SRH-6LoRH, a router determines the IPv6 address of the current segment endpoint.

If strict source routing is enforced and thus router is not the segment endpoint for the packet then this router MUST drop the packet.

If this router is the current segment endpoint, then the router pops its address as described in Section 5.5 and continues processing the packet.

If there is still a SRH-6LoRH, then the router determines the new segment endpoint and routes the packet towards that endpoint.

Otherwise the router uses the destination in the inner IP header to forward or accept the packet.

The segment endpoint of a packet MUST be determined as follows:

The router first determines the Compression Reference as discussed in Section 4.3.1.

The router then coalesces the Compression Reference with the first entry of the first SRH-6LoRH header as discussed in Section 5.4. If the type of the SRH-6LoRH header is type 4 then the coalescence is a full override.

Since the Compression Reference is an uncompressed address, the coalesced IPv6 address is also expressed in the full 128bits.

An example of this operation is provided in Appendix A.3.

6. The RPL Packet Information 6LoRH

[RFC6550], Section 11.2, specifies the RPL Packet Information (RPI) as a set of fields that are placed by RPL routers in IP packets to identify the RPL Instance, detect anomalies and trigger corrective actions.

In particular, the SenderRank, which is the scalar metric computed by a specialized Objective Function such as [RFC6552], indicates the Rank of the sender and is modified at each hop. The SenderRank field is used to validate that the packet progresses in the expected direction, either upwards or downwards, along the DODAG.

RPL defines the RPL Option for Carrying RPL Information in Data-Plane Datagrams [RFC6553] to transport the RPI, which is carried in an IPv6 Hop-by-Hop Options Header [RFC2460], typically consuming eight bytes per packet.

With [RFC6553], the RPL option is encoded as six octets, which must be placed in a Hop-by-Hop header that consumes two additional octets for a total of eight octets. To limit the header's range to just the RPL domain, the Hop-by-Hop header must be added to (or removed from) packets that cross the border of the RPL domain.

The 8-byte overhead is detrimental to LLN operation, in particular with regards to bandwidth and battery constraints. These bytes may cause a containing frame to grow above maximum frame size, leading to Layer 2 or 6LoWPAN [RFC4944] fragmentation, which in turn leads to even more energy expenditure and issues discussed in LLN Fragment Forwarding and Recovery [I-D.thubert-6lo-forwarding-fragments].

An additional overhead comes from the need, in certain cases, to add an IP-in-IP encapsulation to carry the Hop-by-Hop header. This is needed when the router that inserts the Hop-by-Hop header is not the source of the packet, so that an error can be returned to the router. This is also the case when a packet originated by a RPL node must be stripped from the Hop-by-Hop header to be routed outside the RPL domain.

For that reason, this specification defines an IP-in-IP-6LoRH header in Section 7, but it must be noted that removal of a 6LoRH header does not require manipulation of the packet in the LOWPAN_IPHC, and thus, if the source address in the LOWPAN_IPHC is the node that inserted the IP-in-IP-6LoRH header then this situation alone does not mandate an IP-in-IP-6LoRH header.

Note: A typical packet in RPL non-storing mode going down the RPL graph requires an IP-in-IP encapsulation of the SRH, whereas the RPI is usually (and quite illegally) omitted, unless it is important to indicate the RPLInstanceID. To match this structure, an optimized IP-in-IP 6LoRH header is defined in Section 7.

As a result, a RPL packet may bear only an RPI-6LoRH header and no IP-in-IP-6LoRH header. In that case, the source and destination of the packet are specified by the LOWPAN_IPHC.

As with [RFC6553], the fields in the RPI include an 'O', an 'R', and an 'F' bit, an 8-bit RPLInstanceID (with some internal structure), and a 16-bit SenderRank.

The remainder of this section defines the RPI-6LoRH header, which is a Critical 6LoWPAN Routing Header that is designed to transport the RPI in 6LoWPAN LLNs.

6.1. Compressing the RPLInstanceID

RPL Instances are discussed in [RFC6550], Section 5. A number of simple use cases do not require more than one RPL Instance, and in such cases, the RPL Instance is expected to be the Global Instance 0. A global RPLInstanceID is encoded in a RPLInstanceID field as follows:

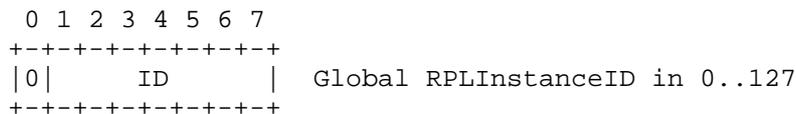


Figure 8: RPLInstanceID Field Format for Global Instances.

For the particular case of the Global Instance 0, the RPLInstanceID field is all zeros. This specification allows to elide a RPLInstanceID field that is all zeros, and defines a I flag that, when set, signals that the field is elided.

6.2. Compressing the SenderRank

The SenderRank is the result of the DAGRank operation on the rank of the sender; here the DAGRank operation is defined in [RFC6550], Section 3.5.1, as:

$$\text{DAGRank}(\text{rank}) = \text{floor}(\text{rank}/\text{MinHopRankIncrease})$$

If MinHopRankIncrease is set to a multiple of 256, the least significant 8 bits of the SenderRank will be all zeroes; by eliding those, the SenderRank can be compressed into a single byte. This idea is used in [RFC6550] by defining DEFAULT_MIN_HOP_RANK_INCREASE as 256 and in [RFC6552] that defaults MinHopRankIncrease to DEFAULT_MIN_HOP_RANK_INCREASE.

This specification allows to encode the SenderRank as either one or two bytes, and defines a K flag that, when set, signals that a single byte is used.

6.3. The Overall RPI-6LoRH encoding

The RPI-6LoRH header provides a compressed form for the RPL RPI. Routers that need to forward a packet with a RPI-6LoRH header are expected to be RPL routers that support this specification. If a

non-RPL router receives a packet with a RPI-6LoRH header, there was a routing error and the packet should be dropped. Thus the Type field MUST NOT be ignored.

Since the I flag is not set, the TSE field does not need to be a length expressed in bytes. In that case the field is fully reused for control bits that encode the O, R and F flags from the RPI, as well as the I and K flags that indicate the compression format.

The Type for the RPI-6LoRH is 5.

The RPI-6LoRH header is immediately followed by the RPLInstanceID field, unless that field is fully elided, and then the SenderRank, which is either compressed into one byte or fully in-lined as two bytes. The I and K flags in the RPI-6LoRH header indicate whether the RPLInstanceID is elided and/or the SenderRank is compressed. Depending on these bits, the Length of the RPI-6LoRH may vary as described hereafter.

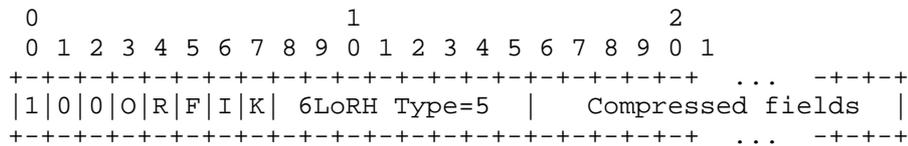


Figure 9: The Generic RPI-6LoRH Format.

O, R, and F bits: The O, R, and F bits are defined in [RFC6550], section 11.2.

I bit: If it is set, the RPLInstanceID is elided and the RPLInstanceID is the Global RPLInstanceID 0. If it is not set, the octet immediately following the type field contains the RPLInstanceID as specified in [RFC6550], section 5.1.

K bit: If it is set, the SenderRank is compressed into one octet, with the least significant octet elided. If it is not set, the SenderRank, is fully inlined as two octets.

In Figure 10, the RPLInstanceID is the Global RPLInstanceID 0, and the MinHopRankIncrease is a multiple of 256 so the least significant byte is all zeros and can be elided:

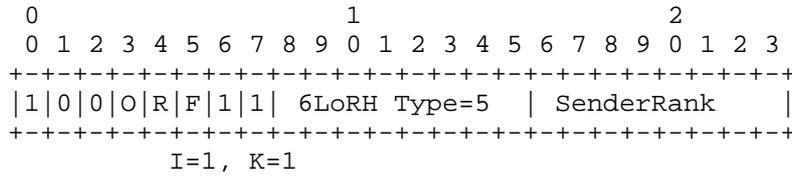


Figure 10: The most compressed RPI-6LoRH.

In Figure 11, the RPLInstanceID is the Global RPLInstanceID 0, but both bytes of the SenderRank are significant so it can not be compressed:

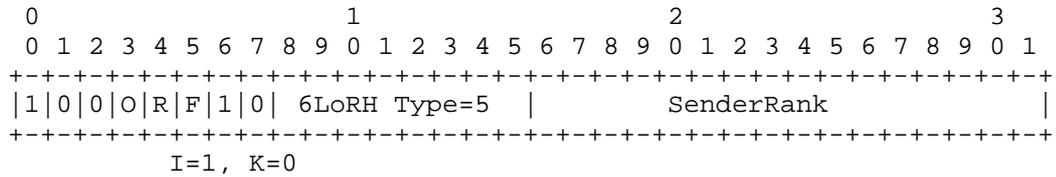


Figure 11: Eliding the RPLInstanceID.

In Figure 12, the RPLInstanceID is not the Global RPLInstanceID 0, and the MinHopRankIncrease is a multiple of 256:

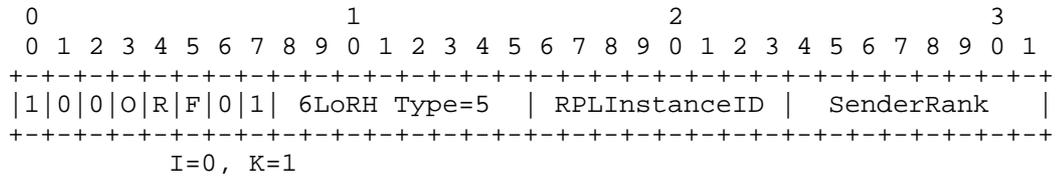


Figure 12: Compressing SenderRank.

In Figure 13, the RPLInstanceID is not the Global RPLInstanceID 0, and both bytes of the SenderRank are significant:

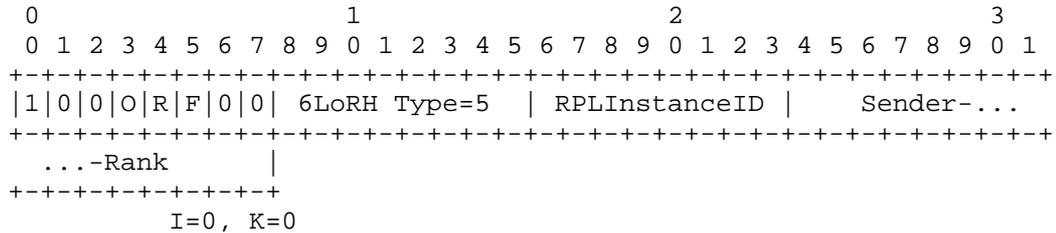


Figure 13: Least compressed form of RPI-6LoRH.

7. The IP-in-IP 6LoRH Header

The IP-in-IP 6LoRH (IP-in-IP-6LoRH) header is an Elective 6LoWPAN Routing Header that provides a compressed form for the encapsulating IPv6 Header in the case of an IP-in-IP encapsulation.

An IP-in-IP encapsulation is used to insert a field such as a Routing Header or an RPI at a router that is not the source of the packet. In order to send an error back regarding the inserted field, the address of the router that performs the insertion must be provided.

The encapsulation can also enable the last router prior to Destination to remove a field such as the RPI, but this can be done in the compressed form by removing the RPI-6LoRH, so an IP-in-IP-6LoRH encapsulation is not required for that sole purpose.

This field is not critical for routing so the Type can be ignored, and the TSE field contains the Length in bytes.

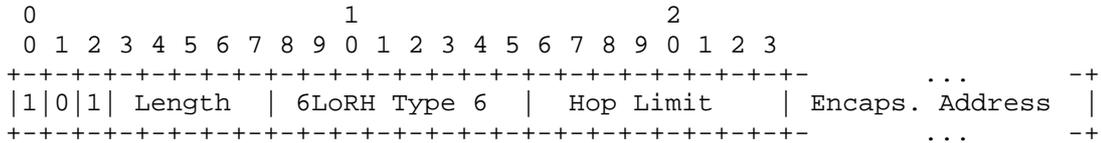


Figure 14: The IP-in-IP-6LoRH.

The Length of an IP-in-IP-6LoRH header is expressed in bytes and MUST be at least 1, to indicate a Hop Limit (HL), that is decremented at each hop. When the HL reaches 0, the packet is dropped per [RFC2460].

If the Length of an IP-in-IP-6LoRH header is exactly 1, then the Encapsulator Address is elided, which means that the Encapsulator is a well-known router, for instance the root in a RPL graph.

The most efficient compression of an IP-in-IP encapsulation that can be achieved with this specification is obtained when an endpoint of the packet is the root of the RPL DODAG associated to the RPL Instance that is used to forward the packet, and the root address is known implicitly as opposed to signaled explicitly in the data packets.

If the Length of an IP-in-IP-6LoRH header is greater than 1, then an Encapsulator Address is placed in a compressed form after the Hop Limit field. The value of the Length indicates which compression is performed on the Encapsulator Address. For instance, a Size of 3

indicates that the Encapsulator Address is compressed to 2 bytes. The reference for the compression is the address of the root of the DODAG. The way the address of the root is determined is discussed in Section 4.3.2.

With RPL, the destination address in the IP-in-IP header is implicitly the root in the RPL graph for packets going upwards, and, in storing mode, it is the destination address in the IPHC for packets going downwards. In non-storing mode, there is no implicit value for packets going downwards.

If the implicit value is correct, the destination IP address of the IP-in-IP encapsulation can be elided. Else, the destination IP address of the IP-in-IP header is transported in a SRH-6LoRH header as the first entry of the first of these headers.

If the final destination of the packet is a leaf that does not support this specification, then the chain of 6LoRH headers must be stripped by the RPL/6LR router to which the leaf is attached. In that example, the destination IP address of the IP-in-IP header cannot be elided.

In the special case where a 6LoRH header is used to route 6LoWPAN fragments, the destination address is not accessible in the IPHC on all fragments and can be elided only for the first fragment and for packets going upwards.

8. Security Considerations

The security considerations of [RFC4944], [RFC6282], and [RFC6553] apply.

Using a compressed format as opposed to the full in-line format is logically equivalent and is believed to not create an opening for a new threat when compared to [RFC6550], [RFC6553] and [RFC6554].

9. IANA Considerations

9.1. Reserving Space in 6LoWPAN Dispatch Page 1

This specification reserves Dispatch Value Bit Patterns within the 6LoWPAN Dispatch Page 1 as follows:

101xxxxx: for Elective 6LoWPAN Routing Headers

100xxxxx: for Critical 6LoWPAN Routing Headers.

9.2. New 6LoWPAN Routing Header Type Registry

This document creates an IANA registry for the 6LoWPAN Routing Header Type, and assigns the following values:

- 0..4: SRH-6LoRH [RFCthis]
- 5: RPI-6LoRH [RFCthis]
- 6: IP-in-IP-6LoRH [RFCthis]

10. Acknowledgments

The authors wish to thank Tom Phinney, Thomas Watteyne, Tengfei Chang, Martin Turon, James Woodyatt, Samita Chakrabarti, Jonathan Hui, Gabriel Montenegro and Ralph Droms for constructive reviews to the design in the 6lo Working Group. The overall discussion involved participants to the 6MAN, 6TiSCH and ROLL WGs, thank you all. Special thanks to the chairs of the ROLL WG, Michael Richardson and Ines Robles, and Brian Haberman, Internet Area A-D, and Adrian Farrel, Routing Area A-D, for driving this complex effort across Working Groups and Areas.

11. References

11.1. Normative References

- [I-D.ietf-6lo-paging-dispatch]
Thubert, P., "6LoWPAN Paging Dispatch", draft-ietf-6lo-paging-dispatch-01 (work in progress), January 2016.
- [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", 2015.
- [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>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.

- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., 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, DOI 10.17487/RFC6550, March 2012, <<http://www.rfc-editor.org/info/rfc6550>>.
- [RFC6552] Thubert, P., Ed., "Objective Function Zero for the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6552, DOI 10.17487/RFC6552, March 2012, <<http://www.rfc-editor.org/info/rfc6552>>.
- [RFC6553] Hui, J. and JP. Vasseur, "The Routing Protocol for Low-Power and Lossy Networks (RPL) Option for Carrying RPL Information in Data-Plane Datagrams", RFC 6553, DOI 10.17487/RFC6553, March 2012, <<http://www.rfc-editor.org/info/rfc6553>>.
- [RFC6554] Hui, J., Vasseur, JP., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6554, DOI 10.17487/RFC6554, March 2012, <<http://www.rfc-editor.org/info/rfc6554>>.
- [RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102, DOI 10.17487/RFC7102, January 2014, <<http://www.rfc-editor.org/info/rfc7102>>.
- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<http://www.rfc-editor.org/info/rfc7228>>.

11.2. Informative References

[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.robles-roll-useofrplinfo]

Robles, I., Richardson, M., and P. Thubert, "When to use RFC 6553, 6554 and IPv6-in-IPv6", draft-robles-roll-useofrplinfo-02 (work in progress), October 2015.

[I-D.thubert-6lo-forwarding-fragments]

Thubert, P. and J. Hui, "LLN Fragment Forwarding and Recovery", draft-thubert-6lo-forwarding-fragments-02 (work in progress), November 2014.

[RFC6775]

Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<http://www.rfc-editor.org/info/rfc6775>>.

[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>>.

Appendix A. Examples

A.1. Examples Compressing The RPI

The example in Figure 15 illustrates the 6LoRH compression of a classical packet in Storing Mode in all directions, as well as in non-Storing mode for a packet going up the DODAG following the default route to the root. In this particular example, a fragmentation process takes place per [RFC4944], and the fragment headers must be placed in Page 0 before switching to Page 1:

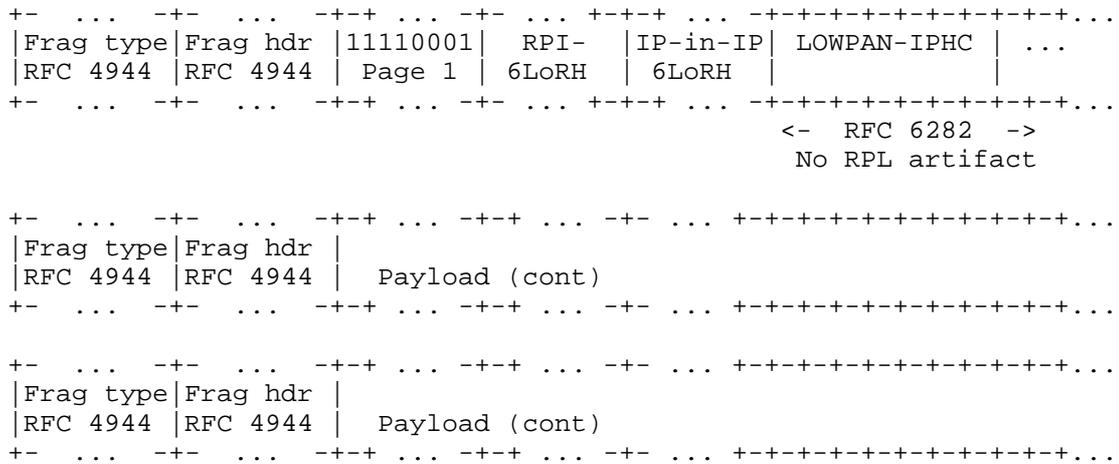


Figure 15: Example Compressed Packet with RPI.

In Storing Mode, if the packet stays within the RPL domain, then it is possible to save the IP-in-IP encapsulation, in which case only the RPI is compressed with a 6LoRH, as illustrated in Figure 16 in the case of a non-fragmented ICMP packet:

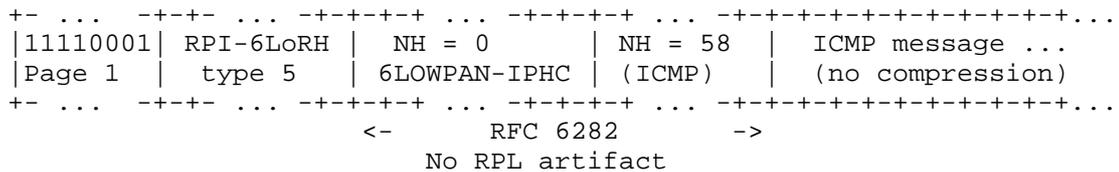


Figure 16: Example ICMP Packet with RPI in Storing Mode.

The format in Figure 16 is logically equivalent to the non-compressed format illustrated in Figure 17:

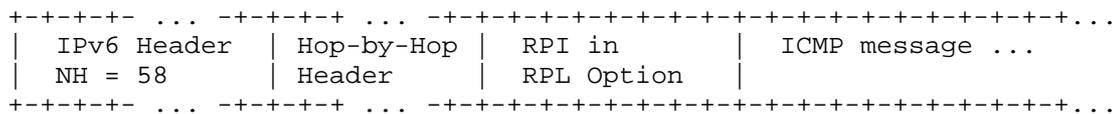


Figure 17: Uncompressed ICMP Packet with RPI.

For a UDP packet, the transport header can be compressed with 6LoWPAN HC [RFC6282] as illustrated in Figure 18:

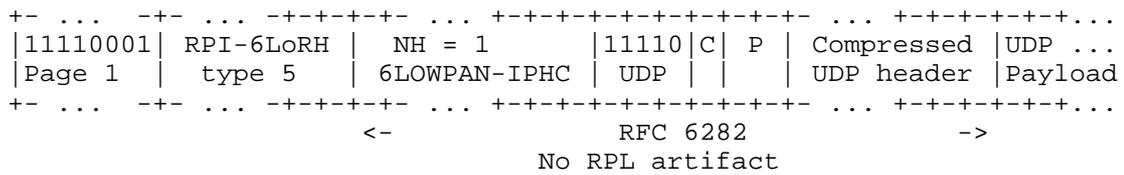


Figure 18: Uncompressed ICMP Packet with RPI.

If the packet is received from the Internet in Storing Mode, then the root is supposed to encapsulate the packet to insert the RPI. The resulting format would be as represented in Figure 19:

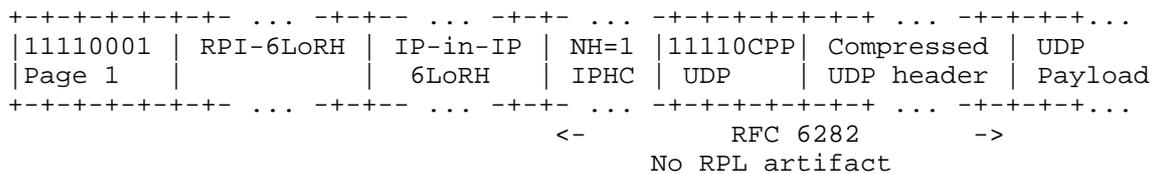


Figure 19: RPI inserted by the root in Storing Mode.

A.2. Example Of Downward Packet In Non-Storing Mode

The example illustrated in Figure 20 is a classical packet in non-Storing mode for a packet going down the DODAG following a source routed path from the root. Say that we have 4 forwarding hops to reach a destination. In the non-compressed form, when the root generates the packet, the last 3 hops are encoded in a Routing Header type 3 (SRH) and the first hop is the destination of the packet. The intermediate hops perform a swap and the hop count indicates the current active hop [RFC2460], [RFC6554].

When compressed with this specification, the 4 hops are encoded in SRH-6LoRH when the root generates the packet, and the final destination is left in the LOWPAN-IPHC. There is no swap, and the forwarding node that corresponds to the first entry effectively consumes it when forwarding, which means that the size of the encoded packet decreases and that the hop information is lost.

If the last hop in a SRH-6LoRH is not the final destination then it removes the SRH-6LoRH before forwarding.

In the particular example illustrated in Figure 20, all addresses in the DODAG are assigned from a same /112 prefix and the last 2 octets encoding an identifier such as a IEEE 802.15.4 short address. In that case, all addresses can be compressed to 2 octets, using the

root address as reference. There will be one SRH_6LoRH header, with, in this example, 3 compressed addresses:

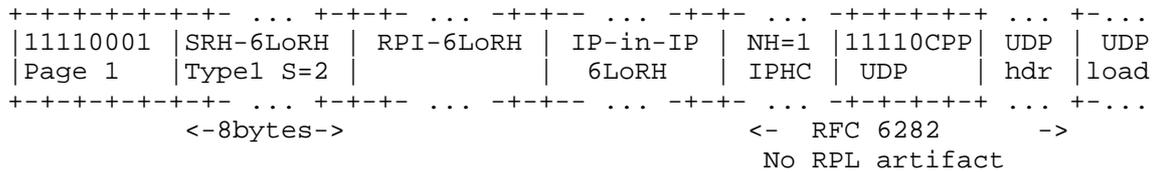


Figure 20: Example Compressed Packet with SRH.

One may note that the RPI is provided. This is because the address of the root that is the source of the IP-in-IP header is elided and inferred from the RPLInstanceID in the RPI. Once found from a local context, that address is used as Compression Reference to expand addresses in the SRH-6LoRH.

With the RPL specifications available at the time of writing this draft, the root is the only node that may incorporate a SRH in an IP packet. When the root forwards a packet that it did not generate, it has to encapsulate the packet with IP-in-IP.

But if the root generates the packet towards a node in its DODAG, then it should avoid the extra IP-in-IP as illustrated in Figure 21:

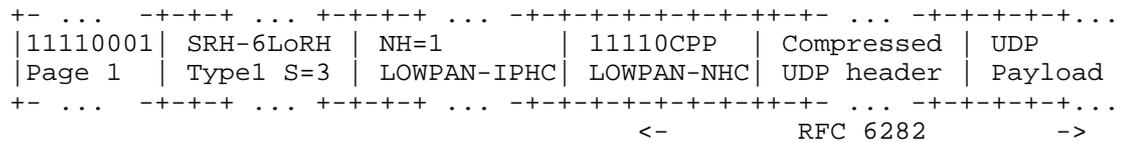


Figure 21: compressed SRH 4*2bytes entries sourced by root.

Note: the RPI is not represented though RPL [RFC6550] generally expects it. In this particular case, since the Compression Reference for the SRH-6LoRH is the source address in the LOWPAN-IPHC, and the routing is strict along the source route path, the RPI does not appear to be absolutely necessary.

In Figure 21, all the nodes along the source route path share a same /112 prefix. This is typical of IPv6 addresses derived from an IEEE802.15.4 short address, as long as all the nodes share a same PAN-ID. In that case, a type-1 SRH-6LoRH header can be used for encoding. The IPv6 address of the root is taken as reference, and only the last 2 octets of the address of the intermediate hops is encoded. The Size of 3 indicates 4 hops, resulting in a SRH-6LoRH of 10 bytes.

A.3. Example of SRH-6LoRH life-cycle

This section illustrates the operation specified in Section 5.6 of forwarding a packet with a compressed SRH along an A->B->C->D source route path. The operation of popping addresses is exemplified at each hop.

Packet as received by node A

```
-----
Type 3 SRH-6LoRH Size = 0   AAAA AAAA AAAA AAAA
Type 1 SRH-6LoRH Size = 0                               BBBB
Type 2 SRH-6LoRH Size = 1                               CCCC CCCC
                                                         DDDD DDDD
```

Step 1 popping BBBB the first entry of the next SRH-6LoRH
 Step 2 next is if larger value (2 vs. 1) the SRH-6LoRH is removed

```
Type 3 SRH-6LoRH Size = 0   AAAA AAAA AAAA AAAA
Type 2 SRH-6LoRH Size = 1                               CCCC CCCC
                                                         DDDD DDDD
```

Step 3: recursion ended, coalescing BBBB with the first entry

```
Type 3 SRH-6LoRH Size = 0   AAAA AAAA AAAA BBBB
```

Step 4: routing based on next segment endpoint to B

Figure 22: Processing at Node A.

Packet as received by node B

```
-----
Type 3 SRH-6LoRH Size = 0   AAAA AAAA AAAA BBBB
Type 2 SRH-6LoRH Size = 1           CCCC CCCC
                                   DDDD DDDD
```

Step 1 popping CCCC CCCC, the first entry of the next SRH-6LoRH
 Step 2 removing the first entry and decrementing the Size (by 1)

```
Type 3 SRH-6LoRH Size = 0   AAAA AAAA AAAA BBBB
Type 2 SRH-6LoRH Size = 0           DDDD DDDD
```

Step 3: recursion ended, coalescing CCCC CCCC with the first entry
 Type 3 SRH-6LoRH Size = 0 AAAA AAAA CCCC CCCC

Step 4: routing based on next segment endpoint to C

Figure 23: Processing at Node B.

Packet as received by node C

```
-----
Type 3 SRH-6LoRH Size = 0   AAAA AAAA CCCC CCCC
Type 2 SRH-6LoRH Size = 0           DDDD DDDD
```

Step 1 popping DDDD DDDD, the first entry of the next SRH-6LoRH
 Step 2 the SRH-6LoRH is removed

```
Type 3 SRH-6LoRH Size = 0   AAAA AAAA CCCC CCCC
```

Step 3: recursion ended, coalescing DDDD DDDDD with the first entry
 Type 3 SRH-6LoRH Size = 0 AAAA AAAA DDDD DDDD

Step 4: routing based on next segment endpoint to D

Figure 24: Processing at Node C.

Packet as received by node D

Type 3 SRH-6LoRH Size = 0 AAAA AAAA DDDD DDDD

Step 1 the SRH-6LoRH is removed.
Step 2 no more header, routing based on inner IP header.

Figure 25: Processing at Node D.

Authors' Addresses

Pascal Thubert (editor)
Cisco Systems
Building D - Regus
45 Allee des Ormes
BP1200
MOUGINS - Sophia Antipolis 06254
FRANCE

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

Carsten Bormann
Universitaet Bremen TZI
Postfach 330440
Bremen D-28359
Germany

Phone: +49-421-218-63921
Email: cabo@tzi.org

Laurent Toutain
Institut MINES TELECOM; TELECOM Bretagne
2 rue de la Chataigneraie
CS 17607
Cesson-Sevigne Cedex 35576
France

Email: Laurent.Toutain@telecom-bretagne.eu

Robert Cragie
ARM Ltd.
110 Fulbourn Road
Cambridge CB1 9NJ
UK

Email: robert.cragie@gridmerge.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 17, 2016

T. Kivinen
INSIDE Secure
P. Kinney
Kinney Consulting LLC
March 16, 2016

IEEE 802.15.4 Information Element for IETF
draft-kivinen-802-15-ie-00.txt

Abstract

IEEE Std. 802.15.4-2015 has Information Elements (IE) that can be used to extend the 802.15.4 in interoperable manner. IEEE 802.15 Assigned Numbers Authority (ANA) manages the registry of the Information Elements, and this document requests ANA to allocate a number for IETF and provides the information how the IE is formatted to provide sub types.

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 17, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- 1. Introduction 2
- 2. Terminology 2
- 3. Users of the IETF IE 3
- 4. IETF IE Subtype Format 3
- 5. Request to allocate IETF IE 4
- 6. Security Considerations 4
- 7. IANA Considerations 4
- 8. References 4
 - 8.1. Normative References 4
 - 8.2. Informative References 4
- Authors' Addresses 5

1. Introduction

The IEEE Std. 802.15.4-2015 [IEEE-802-15-4] has Information Elements (IE) that can be used to extended the 802.15.4 in interoperable manner. There are two different IE types, Header IE and Payload IE. The Header IEs are part of the Medium Access Control (MAC) header, and they are never encrypted, but they may be authenticated. Most of the Header IE processing is done by the MAC, and IETF protocols should not need to extend up with them. The Payload IEs are part of the MAC payload and they may be encrypted and authenticated.

IETF protocols will need to include information in the 802.15.4 frames, and standard 802.15.4 way of doing that is to include payload IE in the frame that will contain the information. Because of this the IETF needs to obtain a dedicated Payload IE.

The 802.15.4 operations manual provides information on how a standardization organization may request an allocation of the one IE to them. To make this request the standardization organization needs to: provide the reason for the request; a description of the protocol format that shows there is sufficient subtype capability; a statement that the external organization understands that only one ID number will be issued.

This document provides the information needed for the request.

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. Users of the IETF IE

There are several IETF working groups such as 6tisch, 6lo, core etc, which could benefit from the IETF IE. The 6tisch working group has already expressed the need for the IE, and this allocation should provide them a way forward.

4. IETF IE Subtype Format

The maximum length of the Payload IE content is 2047 octets, and 802.15.4 frame contains a list of payload IEs, i.e. a single frame can have multiple payload IEs, terminated with the payload IE terminator, and may be followed by the payload.

Because the frame contains a list of the payloads, there is no need to provide internal structure inside the IETF IE, and the Payload IE format of the 802.15.4 contains the Length field. The length of the Sub-Type Content can be calculated from the Length field of the IETF IE.

The format of the IETF IE is as follows:

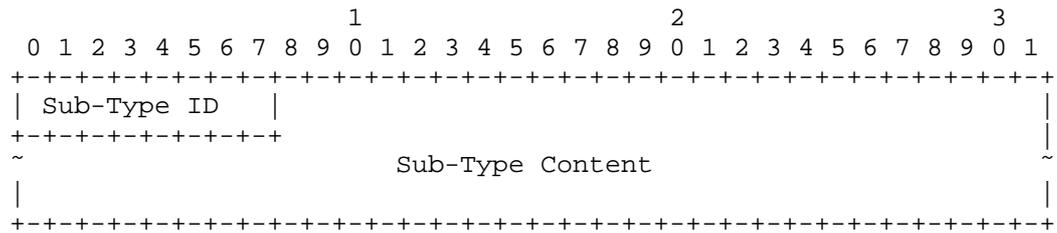


Figure 1: IETF IE Subtype Format

- o Sub-Type ID is the IANA allocated number specifying the sub-type of the IETF IE. Value 0 is reserved for future extensibility, i.e., in case a longer Sub-Type ID field is needed.
- o Sub-Type Content is the actual content of the information element, and its length can be calculated from the Length field of the IETF IE.

One IEEE 802.15.4 frame can contain multiple IETF IEs for same or different sub types.

5. Request to allocate IETF IE

IETF would request the 802.15.4 Working Group to allocate a Payload IE for IETF use. Furthermore IETF understands that only one ID will be issued to it.

6. Security Considerations

This document creates an IANA registry for IETF IE Sub-type IDs, and the security of the protocols using the IEs needs to be described in the actual documents allocating values from this registry.

The IEEE Std. 802.15.4-2015 [IEEE-802-15-4] contains methods where security of the IE can be enforced when a frame is received, but this is only per IE type, thus all IETF IEs will have same security level requirements regardless of the Sub-Type ID used. This can cause issues if different security processing would be needed and any of those IEs would need to be processed in the MAC level. Fortunately everything IETF does should be in a higher level than the MAC level, thus the higher layer processing for these IEs needs to perform separate security policy checking based on the IETF IE Sub-Type ID in addition to the checks done by the MAC.

7. IANA Considerations

This document creates a new registry for IETF IE Sub-type IDs registry:

Value	Sub-type ID
0	Reserved
1-200	Unassigned
201-255	Private Use

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

[IEEE-802-15-4] "IEEE Standard for Low-Rate Wireless Personal Area Networks (WPANs)", IEEE Standard 802.15.4, 2015.

Authors' Addresses

Tero Kivinen
INSIDE Secure
Eerikinkatu 28
HELSINKI FI-00180
FI

Email: kivinen@iki.fi

Pat Kinney
Kinney Consulting LLC

Email: pat.kinney@kinneyconsultingllc.com

6lo
Internet-Draft
Updates: 6775 (if approved)
Intended status: Standards Track
Expires: February 23, 2017

M. Sethi, Ed.
Ericsson
P. Thubert
Cisco
B. Sarikaya, Ed.
Huawei USA
August 22, 2016

Address Protected Neighbor Discovery for Low-power and Lossy Networks
draft-sarikaya-6lo-ap-nd-04

Abstract

This document defines an extension to 6LoWPAN Neighbor Discovery. This extension is designed for low-power and lossy network environments and it supports multi-hop operation. Nodes supporting this extension compute a Cryptographically Unique Interface ID and associate it with one or more of their Registered Addresses. The Cryptographic ID (Crypto-ID) uniquely identifies the owner of the Registered Address. It is used in place of the EUI-64 address that is specified in RFC 6775. Once an address is registered with a Cryptographic ID, only the owner of that ID can modify the state information of the Registered Address in the 6LR and 6LBR.

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 February 23, 2017.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Terminology	4
3. Requirements	4
4. Protocol Interactions	5
4.1. Overview	5
4.2. Updating RFC 6775	7
4.2.1. Crypto-ID Calculation	10
4.3. Multihop Operation	13
5. Security Considerations	14
6. IANA considerations	14
7. Acknowledgements	14
8. References	14
8.1. Normative References	14
8.2. Informative references	16
Authors' Addresses	17

1. Introduction

Neighbor discovery for IPv6 [RFC4861] and stateless address autoconfiguration [RFC4862] are together referred to as neighbor discovery protocols (NDP). They are defined for regular hosts that have sufficient memory and computation capabilities. These protocols are however not suitable for resource-constrained devices. Therefore, they require adaptation to work on resource-constrained hosts operating over a low-power and lossy network (LLN). Neighbor Discovery optimizations for 6LoWPAN networks include simple optimizations such as a host address registration feature. This feature uses the address registration option (ARO) which is sent in the unicast Neighbor Solicitation (NS) and Neighbor Advertisement (NA) messages [RFC6775].

With 6LoWPAN ND [RFC6775], the ARO option includes a EUI-64 interface ID to uniquely identify the interface of the Registered Address on the registering device, so as to correlate further registrations for the same address and avoid address duplication. The EUI-64 interface ID is not secure and its ownership cannot be verified. Consequently,

any device claiming the same EUI-64 interface ID may take over an existing registration and attract the traffic for that address. The address registration mechanism in [RFC6775] is limited as it does not require a node to prove its ownership of the EUI-64 Interface ID. Therefore, any node connected to the subnet and aware of the registered address to EUI-64 interface ID mapping may effectively fake the same interface ID and steal an address.

In this document, we extend 6LoWPAN ND to protect the address ownership with cryptographic material, but as opposed to Secure Neighbor Discovery (SEND) [RFC3971] and Cryptographically Generated Addresses (CGAs) [RFC3972], the cryptographic material generated is not embedded in the Interface ID (IID) as an IPv6 address. Instead, the generated cryptographic ID is used as a correlator associated with the registration of the IP address. This approach is made possible with 6LoWPAN ND [RFC6775], where the 6LR and the 6LBR maintain state information for each Registered Address. If a cryptographic ID is associated with the first 6LoWPAN ND registration, then it can be used to validate any future updates to the registration.

In order to achieve this ownership verification, in this extension specification, the EUI-64 interface ID used in 6LoWPAN ND is replaced with cryptographic material whose ownership can be verified. The extension also provides new means for the 6LR to validate ownership of the registration, and thus, the ownership of registered address. The resulting protocol is called Protected Address Registration protocol (ND-PAR).

In ND-PAR, a node typically generates one 64-bit cryptographic ID (Crypto-ID) and uses it as Unique Interface ID in the registration of one (or more) of its addresses with the 6LR, which it attaches to and uses as default router. The 6LR validates ownership of the cryptographic ID typically upon creation or update of a registration state, for instance following an apparent movement from one point of attachment to another. The ARO option is modified to carry the Unique Interface ID, and through the DAR/DAC exchange.

Compared with SeND, this specification saves ~1Kbyte in every NS/NA message. Also SeND requires one cryptographic address per IPv6 address. This specification separates the cryptographic identifier from the IPv6 address so that a node can have more than one IPv6 address protected by the same cryptographic identifier. SeND forces the IPv6 address to be cryptographic since it integrates the CGA as an IID. 6LoWPAN derives the IPv6 address from other things like a short address in 802.15.4 to enable a better compression.

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 [RFC3971], [RFC3972], [RFC4861], [RFC4919], [RFC6775], and [I-D.ietf-6lo-backbone-router] which proposes an evolution of [RFC6775] for wider applicability.

This document defines Crypto-ID as an identifier of variable size which in most cases is 64 bits long. It is generated using cryptographic means explained later in this document.

The document also conforms to the terms and models described in [RFC5889] and uses the vocabulary and the concepts defined in [RFC4291] for the IPv6 Architecture.

This document uses [RFC7102] for Terminology in Low power And Lossy Networks.

3. Requirements

In this section we state requirements of a secure neighbor discovery protocol for low-power and lossy networks.

- o The protocol MUST be based on the Neighbor Discovery Optimization for Low-power and Lossy Networks protocol defined in [RFC6775]. RFC6775 utilizes optimizations such as host-initiated interactions for sleeping resource-constrained hosts and elimination of multicast address resolution.
- o New options to be added to Neighbor Solicitation messages MUST lead to small packet sizes, especially compared with existing protocols such as SEcure Neighbor Discovery (SEND). Smaller packet sizes facilitate low-power transmission by resource-constrained nodes on lossy links.
- o The support for this registration mechanism SHOULD be extensible to more LLN links than IEEE 802.15.4 only. Support for at least the LLN links for which a 6lo "IPv6 over foo" specification exists, as well as Low-Power Wi-Fi SHOULD be possible.
- o As part of this extension, a mechanism to compute a unique Identifier should be provided with the capability to form a Link Local Address that SHOULD be unique at least within the LLN connected to a 6LBR.

- o The Address Registration Option used in the ND registration SHOULD be extended to carry the relevant forms of Unique Interface Identifier.
- o The Neighbour Discovery should specify the formation of a site-local address that follows the security recommendations from [RFC7217].

4. Protocol Interactions

Protected address and registration neighbor discovery protocol (ND-PAR) modifies Neighbor Discovery Optimization for Low-power and Lossy Networks [RFC6775] as explained in this section.

4.1. Overview

The scope of the present work is a 6LoWPAN Low Power Lossy Network (LLN), typically a stub network connected to a larger IP network via a Border Router called a 6LBR per [RFC6775].

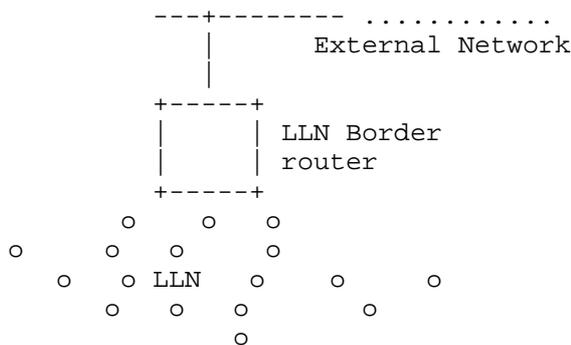


Figure 1: Basic Configuration

The 6LBR maintains a registration state for all devices in the attached LLN, and, in conjunction with the first-hop router (the 6LR), is in a position to validate uniqueness and grant ownership of an IPv6 address before it can be used in the LLN. This is a fundamental difference with a classical network that relies on IPv6 address auto-configuration [RFC4862], where there is no guarantee of ownership from the network, and any IPv6 Neighbor Discovery packet must be individually secured [RFC3971].

In a mesh network, the 6LR is directly connected to the host device. This specification expects that the peer-wise layer-2 security is deployed so that all the packets from a particular host are securely identifiable by the 6LR. The 6LR may be multiple hops away from the

6LBR. Packets are routed between the 6LR and the 6LBR via other 6LRs. This specification expects that a chain of trust is established so that a packet that was validated by the first 6LR can be safely routed by the next 6LRs to the 6LBR.

[I-D.ietf-6tisch-architecture] suggests to use of RPL [RFC6550] as the routing protocol between the 6LRs and the 6LBR, and leveraging a backbone router [I-D.ietf-6lo-backbone-router] to extend the LLN in a larger multilink subnet [RFC4903]. In that model, a registration flow happens as shown in Figure 2. Note that network beyond the 6LBR is out of scope for this document.

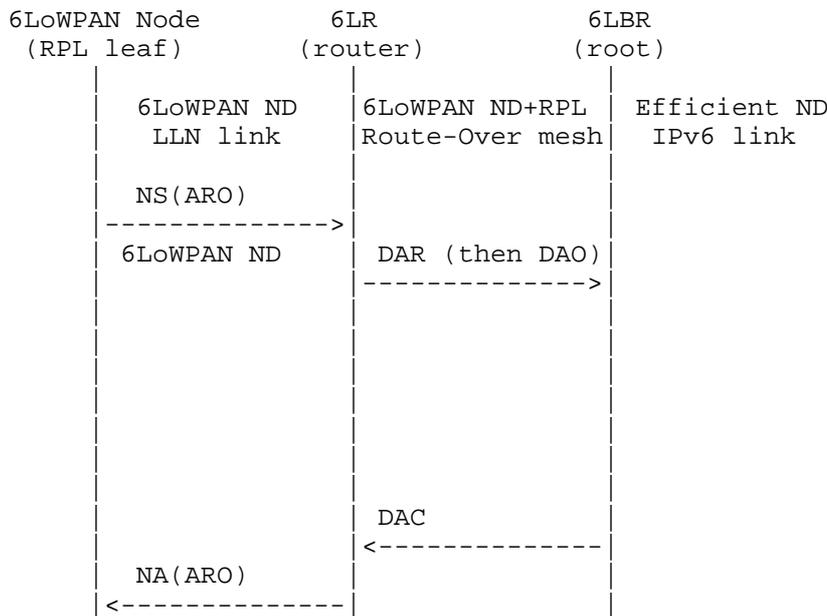


Figure 2: (Re-)Registration Flow over Multi-Link Subnet

A new device that joins the network auto-configures an address and performs an initial registration to an on-link 6LR with an NS message that carries a new Address Registration Option (ARO) [RFC6775]. The 6LR validates the address with the central 6LBR using a DAR/DAC exchange, and the 6LR confirms (or denies) the address ownership with an NA message that also carries an Address Registration Option.

The registration mechanism in [RFC6775] was created for the original purpose of Duplicate Address Detection (DAD), whereby use of an address would be granted as long as the address is not already present in the subnet. But [RFC6775] does not require that the 6LR use the registration for source address validation (SAVI) [RFC7039].

Protected address registration protocol proposed in this document enforces SAVI. With this we ensure that only the correct owner uses the registered address in the source address field. Therefore a destination node can trust that the source is the real owner without using SeND. All packets destined for a node go through the 6LR to which it is attached. The 6LR maintains state information for the registered address along with the MAC address, and link-layer cryptographic key associated with that node. The 6LR therefore only delivers packets to the real owner based on its state information.

In order to validate address ownership, the registration mechanism (that goes all the way to the 6LBR with the DAR/DAC) enables the 6LBR to correlate further claims for a registered address from the device to which it is granted, based on a Unique Interface Identifier (UID). This UID is derived from the MAC address of the device (EUI-64).

This document uses a randomly generated value as an alternate UID for the registration. Proof of ownership of the UID is passed with the first registration to a given 6LR, and enforced at the 6LR, which validates the proof. With this new operation, the 6LR allows only packets from a connected host if the connected host owns the registration of the source address of the packet.

In a multihop 6LoWPAN, the registration with Crypto-ID is propagated to 6LBR as described in Section 4.3. If a chain of trust is present between the 6LR and the 6LBR, then there is no need to propagate the proof of ownership to the 6LBR. All the 6LBR needs to know is that this particular UID is randomly generated, so as to enforce that any update via a different 6LR is also random.

4.2. Updating RFC 6775

Protocol interactions are as defined in Figure 2. The Crypto-ID is calculated as described in Section 4.2.1.

The Target Address field in NS message is set to the prefix concatenated with the node's address. This address does not need duplicate address detection as Crypto-ID is globally unique. So a host cannot steal an address that is already registered unless it has the key used for generating the Crypto-ID. The same Crypto-ID can thus be used to protect multiple addresses e.g. when the node receives a different prefix.

Local or on-link protocol interactions are shown in Figure 3. Crypto-ID and ARO are passed to and stored by the 6LR/6LBR on the first NS and not sent again in the next NS. The operation starts with 6LR sending a Router Advertisement (RA) message to 6LN.

The 6LR/6LBR ensures first-come/first-serve by storing the ARO and the Crypto-ID correlated to the node being registered. The node is free to claim any address it likes as long as it is the first to make such a claim. The node becomes owner of that address and the address is bound to the Crypto-ID in the 6LR/6LBR registry. This procedure avoids the constrained device to compute multiple keys for multiple addresses. The registration process allows the node to tie all the addresses to the same Crypto-ID and have the 6LR/6LBR enforce first-come first-serve after that.

A condition where a 6LN uses multiple IPv6 addresses may happen when the node moves at a different place and receives a different prefix. In this scenario, the node uses the same Crypto-ID to protect its new IPv6 address. This prevents other nodes from stealing the address and trying to use it as their source address.

Note that if the device that moves always forms new MAC and IP address [RFC6775], then this new address can be used for registration. In case of a collision of the new MAC and therefore IP address, the node can easily form a new IPv6 address. This is one case where the use of Crypto-ID would not be needed. Crypto-ID or ND-PAR should be activated when the IP address is claimed at another place, or for a different MAC address at the same place, e.g. for MAC address privacy [I-D.ietf-6man-ipv6-address-generation-privacy].

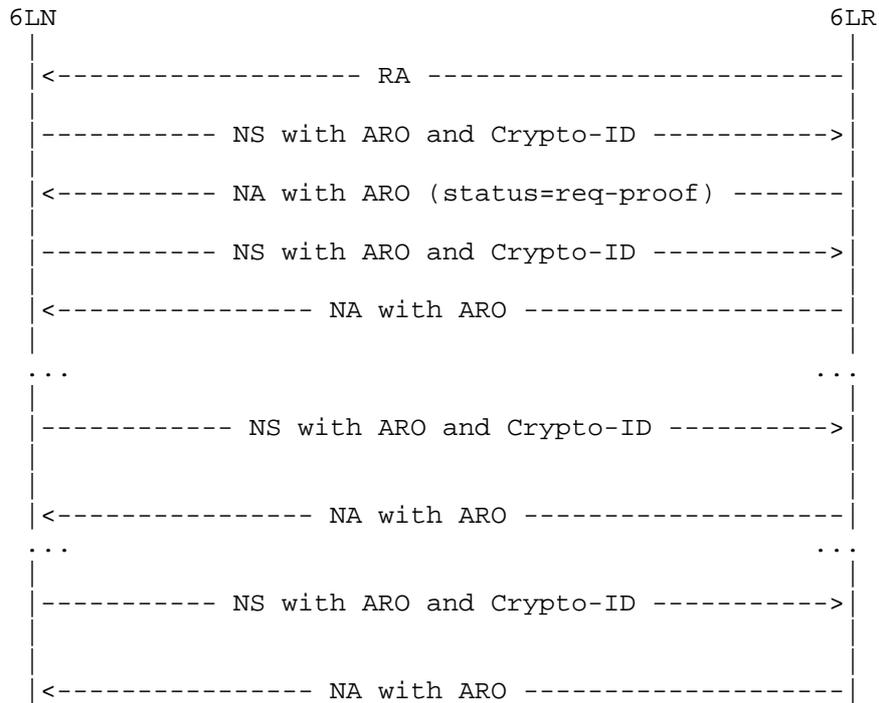


Figure 3: On-link Protocol Operation

Elliptic Curve Cryptography (ECC) is used in the calculation of cryptographic identifier (Crypto-ID). The digital signature is constructed by using the 6LN’s private key over its EUI-64 (MAC) address. The signature value is computed using the ECDSA signature algorithm and the hash function used is SHA-256 [RFC6234]. Public Key is the most important parameter in CGA Parameters (sent by 6LN in an NS message). ECC Public Key could be in uncompressed form or in compressed form where the first octet of the OCTET STRING is 0x04 and 0x02 or 0x03, respectively. Point compression can further reduce the key size by about 32 octets.

After calculating its Crypto-ID, a 6LN sends it along with the CGA parameters in the first NS message, see Figure 3. In order to send Crypto-ID, a modified address registration option called Enhanced Address Registration Option (EARO) is defined in Figure 4. As defined in the figure this ID is variable length, varying between 64 to 128 bits. This ID is 128 bits long only if it is used as IPv6 address. This may happen when some application uses one IP address of the device as device ID. It would make sense in that case to build a real CGA IPv6 address. The prefix of the address would be obtained from prefix information option (PIO in RA) [RFC4861].

6LN also sends some other parameters to enable 6LR or 6LBR to verify the Crypto-ID. The option shown in Figure 5 can be used. In the figure, CGA Parameters field contains the public key, prefix and some other values. It is a simplified form of CGA Option defined in [RFC3971].

4.2.1. Crypto-ID Calculation

First, the modifier is set to a random or pseudo-random 128-bit value. Next, concatenate from left to right the modifier, 9 zero octets and the ECC public key. SHA-256 algorithm is applied on the concatenation. The 112 leftmost bits of the hash value is taken. Concatenate from left to right the modifier value, the subnet prefix and the encoded public key. NIST P-256 is executed on the concatenation. The leftmost bits of the result is used as the Crypto-ID. The length is normally 64 bits, however it could be 128 bits.

In respecting the cryptographical algorithm agility [RFC7696], Curve 25519 [RFC7748] can also be used instead of NIST P-256. This is indicated by 6LN by setting the Crypto Type field in CGA Parameters Option to a value of 1. If 6LBR does not support Curve 25519, it will set Crypto Type field to zero. This means that the default algorithm (NIST P-256) will be used.

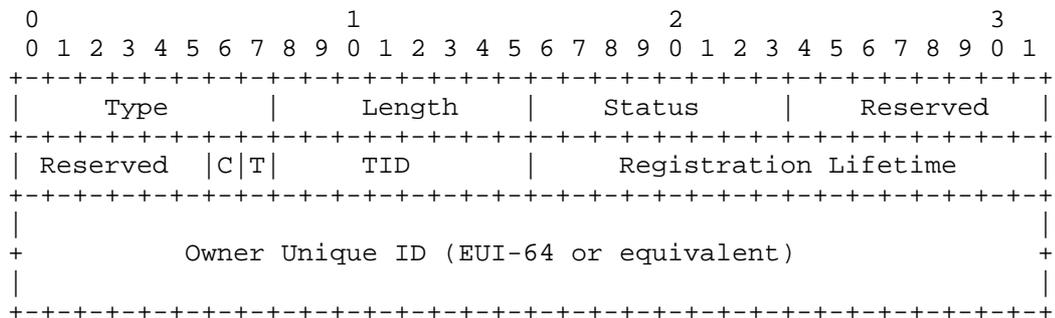


Figure 4: Enhanced Address Registration Option

Type:

TBA1

Length:

8-bit unsigned integer. The length of the option (including the type and length fields) in units of 8 bytes. The value 0 is invalid. A value of 3 with the C flag set indicates a Crypto-ID of 128 bits.

Status:

8-bit unsigned integer. Indicates the status of a registration in the NA response. MUST be set to 0 in NS messages. See below.

Reserved:

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

C:

C bit when set is used to indicate that Owner Unique ID fields contains Crypto-ID.

T and TID:

Defined in [I-D.ietf-6lo-backbone-router].

Owner Unique ID:

In this specification, this field contains Crypto-ID, a variable length field to carry the Crypto-ID or random UID. This field is normally 64 bits long. It could be 128 bits long if IPv6 address is used as the Crypto-ID.

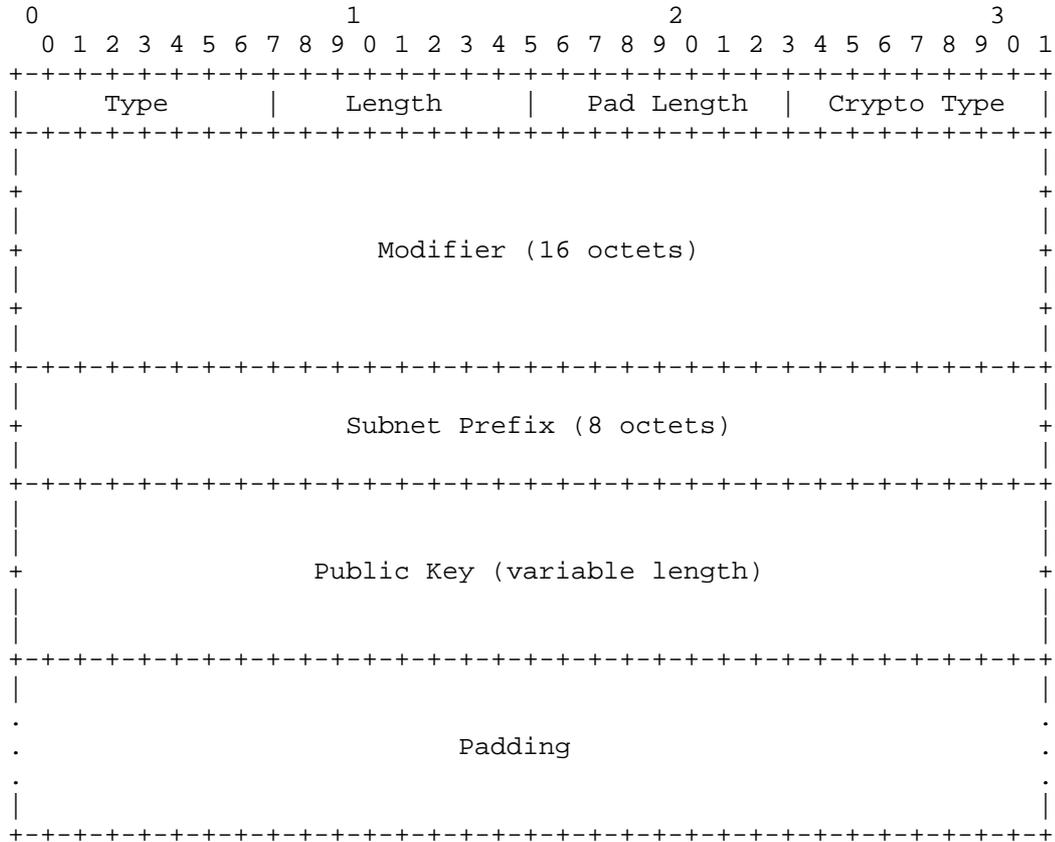


Figure 5: CGA Parameters Option

Type:

TBA2

Length:

The length of the option in units of 8 octets.

Pad Length:

The length of the Padding field.

Crypto Type:

The type of cryptographic algorithm used in calculation Crypto-ID. Default value of all zeros indicate NIST P-256. A value of 1 is assigned for Curve 25519. New values may be defined later.

Modifier:

128 bit random value.

Subnet Prefix:

64 bit subnet prefix.

Public Key:

ECC public key of 6LN.

Padding:

A variable-length field making the option length a multiple of 8, containing as many octets as specified in the Pad Length field.

4.3. Multihop Operation

In multihop 6LoWPAN, 6LBR sends RAs with prefixes downstream and it is the 6LR that receives and relays them to the nodes. 6LR and 6LBR communicate with the ICMPv6 Duplicate Address Request (DAR) and the Duplicate Address Confirmation (DAC) messages. The DAR and DAC use the same message format as NS and NA with different ICMPv6 type values.

In ND-PAR we extend DAR/DAC messages to carry cryptographically generated UID. In a multihop 6LoWPAN, the node exchanges the messages shown in Figure 2. The 6LBR must be aware of who owns an address (EUI-64) to defend the first node if there is an attacker on another 6LR. Because of this the content that the source signs and the signature needs to be propagated to the 6LBR in DAR message. For this purpose the DAR message sent by 6LR to 6LBR MUST contain CGA Parameters and Digital Signature Option carrying the CGA that the node calculates and its public key. DAR message also contains ARO.

It is possible that occasionally, 6LR may miss the node's UID (that it received in ARO). 6LR should be able to ask for it again. This is done by restarting the exchanges shown in Figure 3. The result enables 6LR to refresh the information that was lost. 6LR MUST send DAR message with ARO to 6LBR. 6LBR as a reply forms a DAC message with the information copied from the DAR and the Status field is set to zero. With this exchange, the 6LBR can (re)validate and store the information to make sure that the 6LR is not a fake.

In some cases 6LBR may use DAC message to signal to 6LR that it expects Crypto-ID from 6LR also asks 6LR to verify the EUI-64 6LR received from 6LN. This may happen when a 6LN node is compromised and a fake node is sending the Crypto-ID as if it is the node's EUI-64. Note that the detection in this case can only be done by 6LBR not by 6LR.

5. Security Considerations

The same considerations regarding the threats to the Local Link Network covered in [RFC3971] apply.

The threats discussed in Section 9.2 of [RFC3971] are countered by the protocol described in this document as well.

Collisions of Crypto-ID is a possibility that needs to be considered. The formula for calculating probability of a collision is $1 - e^{-k^2/(2n)}$. If the Crypto-ID is 64-bit long, then the chance of finding a collision is 0.01% when the network contains 66 million nodes. It is important to note that the collision is only relevant when this happens within one stub network (6LBR). A collision of ID in ND-PAR is a rare event. However, when such a collision does happen, the protocol operation is not affected, although it opens a window for a node to hijack an address from another. The link-layer security ensures that the nodes would normally not be aware of a collision on the subnet. If a malicious node is able to gain knowledge of a collision through other means, the only thing that it could do is to steal addresses from the other honest node. This would be no different from what is already possible in a 6lo network today.

6. IANA considerations

IANA is requested to assign two new option type values, TBA1 and TBA2 under the subregistry "IPv6 Neighbor Discovery Option Formats".

7. Acknowledgements

We are grateful to Rene Struik and Robert Moskowitz for their comments that lead to many improvements to this document.

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>>.
- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, DOI 10.17487/RFC3971, March 2005, <<http://www.rfc-editor.org/info/rfc3971>>.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <<http://www.rfc-editor.org/info/rfc3972>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<http://www.rfc-editor.org/info/rfc4291>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<http://www.rfc-editor.org/info/rfc4861>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<http://www.rfc-editor.org/info/rfc4862>>.
- [RFC4903] Thaler, D., "Multi-Link Subnet Issues", RFC 4903, DOI 10.17487/RFC4903, June 2007, <<http://www.rfc-editor.org/info/rfc4903>>.
- [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, DOI 10.17487/RFC4919, August 2007, <<http://www.rfc-editor.org/info/rfc4919>>.
- [RFC5889] Baccelli, E., Ed. and M. Townsley, Ed., "IP Addressing Model in Ad Hoc Networks", RFC 5889, DOI 10.17487/RFC5889, September 2010, <<http://www.rfc-editor.org/info/rfc5889>>.
- [RFC6234] Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)", RFC 6234, DOI 10.17487/RFC6234, May 2011, <<http://www.rfc-editor.org/info/rfc6234>>.

- [RFC6550] Winter, T., Ed., Thubert, P., Ed., 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, DOI 10.17487/RFC6550, March 2012, <<http://www.rfc-editor.org/info/rfc6550>>.
- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<http://www.rfc-editor.org/info/rfc6775>>.
- [RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102, DOI 10.17487/RFC7102, January 2014, <<http://www.rfc-editor.org/info/rfc7102>>.
- [RFC7039] Wu, J., Bi, J., Bagnulo, M., Baker, F., and C. Vogt, Ed., "Source Address Validation Improvement (SAVI) Framework", RFC 7039, DOI 10.17487/RFC7039, October 2013, <<http://www.rfc-editor.org/info/rfc7039>>.
- [RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <<http://www.rfc-editor.org/info/rfc7217>>.
- [RFC7696] Housley, R., "Guidelines for Cryptographic Algorithm Agility and Selecting Mandatory-to-Implement Algorithms", BCP 201, RFC 7696, DOI 10.17487/RFC7696, November 2015, <<http://www.rfc-editor.org/info/rfc7696>>.
- [RFC7748] Langley, A., Hamburg, M., and S. Turner, "Elliptic Curves for Security", RFC 7748, DOI 10.17487/RFC7748, January 2016, <<http://www.rfc-editor.org/info/rfc7748>>.

8.2. Informative references

- [I-D.ietf-6lo-backbone-router]
Thubert, P., "IPv6 Backbone Router", draft-ietf-6lo-backbone-router-01 (work in progress), March 2016.
- [I-D.ietf-6tisch-architecture]
Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", draft-ietf-6tisch-architecture-10 (work in progress), June 2016.

[I-D.ietf-6man-ipv6-address-generation-privacy]
Cooper, A., Gont, F., and D. Thaler, "Privacy
Considerations for IPv6 Address Generation Mechanisms",
draft-ietf-6man-ipv6-address-generation-privacy-08 (work
in progress), September 2015.

Authors' Addresses

Mohit Sethi (editor)
Ericsson
Hirsalantie
Jorvas 02420

Email: mohit@piuha.net

Pascal Thubert
Cisco Systems, Inc
Building D
45 Allee des Ormes - BP1200
MOUGINS - Sophia Antipolis 06254
FRANCE

Phone: +33 497 23 26 34
Email: pthubert@cisco.com

Behcet Sarikaya (editor)
Huawei USA
5340 Legacy Dr. Building 3
Plano, TX 75024

Email: sarikaya@ieee.org

610
Internet-Draft
Intended status: Standards Track
Expires: May 12, 2016

P. Thubert, Ed.
cisco
November 9, 2015

IPv6 Backbone Router
draft-thubert-610-backbone-router-03

Abstract

This specification proposes an update to IPv6 Neighbor Discovery, to enhance the operation of IPv6 over wireless links that exhibit lossy multicast support, and enable a large degree of scalability by splitting the broadcast domains. A higher speed backbone federates multiple wireless links to form a large MultiLink Subnet. Backbone Routers acting as Layer-3 Access Point route packets to registered nodes, where an classical Layer-2 Access Point would bridge. Conversely, wireless nodes register to the Backbone Router to setup routing services in a fashion that is essentially similar to a classical Layer-2 association.

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 May 12, 2016.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents

carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Applicability and Requirements Served	5
3. Terminology	6
4. Overview	9
5. New Types And Formats	10
5.1. Transaction ID	10
5.2. Owner Unique ID	11
5.3. The Enhanced Address Registration Option (EARO)	11
6. Backbone Router Routing Operations	14
6.1. Over the Backbone Link	14
6.2. Over the LLN Link	15
7. Backbone Router Proxy Operations	17
7.1. Registration and Binding State Creation	19
7.2. Defending Addresses	21
8. Security Considerations	22
9. Protocol Constants	22
10. IANA Considerations	22
11. Acknowledgments	23
12. References	23
12.1. Normative References	23
12.2. Informative References	24
12.3. External Informative References	28
Appendix A. Requirements	28
A.1. Requirements Related to Mobility	28
A.2. Requirements Related to Routing Protocols	29
A.3. Requirements Related to the Variety of Low-Power Link types	30
A.4. Requirements Related to Proxy Operations	31
A.5. Requirements Related to Security	31
A.6. Requirements Related to Scalability	32
Author's Address	33

1. Introduction

Though in most cases, including Low-Power ones, IEEE802.11 [IEEE80211] is operated as a wireless extension to an Ethernet bridged domain, the impact of radio broadcasts for IPv6 [RFC2460] multicast operations, in particular related to the power consumption of battery-operated devices, lead the community to rethink the plain layer-2 approach and consider splitting the broadcast domain between

the wired and the wireless access links. To that effect, the current IEEE802.11 specifications require the capability to perform ARP and ND proxy [RFC4389] functions at the Access Points (APs), but rely on snooping for acquiring the related state, which is unsatisfactory in a lossy and mobile environments.

Without a proxy, any IP multicast that circulates in the bridged domain ends up broadcasted by the Access Points to all STAs, including Low-Power battery-operated ones. With an incorrect or missing state in the proxy, a packet may not be delivered to the destination, which may have operational impacts depending on the criticality of the packet.

Some messages are lost for the lack of retries, regardless of their degree of criticality; it results for instance that Duplicate Address Detection (DAD) as defined in [RFC4862] is mostly broken over Wi-Fi [I-D.yourtchenko-6man-dad-issues].

On the other hand, IPv6 multicast messages are processed by most if not all wireless nodes over the fabric even when very few if any of the nodes is effectively listening to the multicast address. It results that a simple Neighbor Solicitation (NS) message [RFC4861], that is supposedly targeted to a very small group of nodes, ends up polluting the whole wireless bandwidth across the fabric [I-D.vyncke-6man-mcast-not-efficient].

It appears that in a variety of Wireless Local Area Networks (WLANs) and Wireless Personal Area Networks (WPANs), the decision to leverage the broadcast support of a particular link should be left to Layer-3 based on the criticality of the message and the number of interested listeners on that link, for the lack of capability to indicate that criticality to the lower layer. To achieve this, the operation at the Access Point cannot be a Layer-2 bridge operation, but that of a Layer-3 router; the concept of MultiLink Subnet (MLSN) must be reintroduced, with IPv6 backbone routers (6BBRs) interconnecting the various links and routing within the subnet. For link-scope multicast operations, a 6BBR participates to MLD on its access links and a multicast routing protocol is setup between the 6BBRs over the backbone of the MLSN.

As the network scales up, none of the approaches of using either broadcast or N*unicast for a multicast packet is really satisfying and the protocols themselves need to be adapted to reduce their use of multicast.

One degree of improvement can be achieved by changing the tuning of the protocol parameters and operational practices, such as suggested in Reducing energy consumption of Router Advertisements

[I-D.ietf-v6ops-reducing-ra-energy-consumption] (RA). This works enables to lower the rate of RA messages but does not solve the problem associated with multicast NS and NA messages, which are a lot more frequent in large-scale radio environments with mobile devices which exhibit intermittent access patterns and short-lived IPv6 addresses.

In the context of IEEE802.15.4 [IEEE802154], the more drastic step of considering the radio as a medium that is different from Ethernet because of the impact of multicast, was already taken with the adoption of Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs) [RFC6775]. This specification applies that same thinking to other wireless links such as Low-Power IEEE802.11 (Wi-Fi) and IEEE802.15.1 (Bluetooth) [IEEE802151], and extends [RFC6775] to enable proxy operation by the 6BBR so as to decouple the broadcast domain in the backbone from the wireless links. The proxy operation can be maintained asynchronous so that low-power nodes or nodes that are deep in a mesh do not need to be bothered synchronously when a lookup is performed for their addresses, effectively implementing the ND contribution to the concept of a Sleep Proxy [I-D.nordmark-6man-dad-approaches].

DHCPv6 [RFC3315] is still a viable option in Low power and Lossy Network (LLN) to assign IPv6 global addresses. However, the IETF standard that supports address assignment specifically for LLNs is 6LoWPAN ND [RFC6775], which is a mix of IPv6 stateless autoconfiguration mechanism (SLAAC) [RFC4862] and a new registration process for ND. This specification introduces a Layer-3 association process based on 6LoWPAN ND that maintains a proxy state in the 6BBR to keep the LLN nodes reachable and protect their addresses through sleeping periods.

A number of use cases, including the Industrial Internet, require a large scale deployment of monitoring sensors that can only be realized in a cost-effective fashion with wireless technologies. Mesh networks are deployed when simpler hub-and-spoke topologies are not sufficient for the expected size, throughput, and density. Meshes imply the routing of packets, operated at either Layer-2 or Layer-3. For routing over a mesh at Layer-3, the IETF has designed the IPv6 Routing Protocol over LLN (RPL) [RFC6550]. 6LoWPAN ND was designed as a stand-alone mechanism separately from RPL, and the interaction between the 2 protocols was not defined. This specification details how periodic updates from RPL can be used by the RPL root to renew the association of the RPL node to the 6BBR on its behalf so as to maintain the proxy operation active for that node.

This document suggests a limited evolution to [RFC6775] so as to allow operation of a 6LoWPAN ND node while a routing protocol (in first instance RPL) is present and operational. It also suggests a more generalized use of the information in the ARO option of the ND messages outside the strict LLN domain, for instance over a converged backbone.

2. Applicability and Requirements Served

Efficiency aware IPv6 Neighbor Discovery Optimizations [I-D.chakrabarti-nordmark-6man-efficient-nd] suggests that 6LoWPAN ND [RFC6775] can be extended to other types of links beyond IEEE802.15.4 for which it was defined. The registration technique is beneficial when the Link-Layer technique used to carry IPv6 multicast packets is not sufficiently efficient in terms of delivery ratio or energy consumption in the end devices, in particular to enable energy-constrained sleeping nodes. The value of such extension is especially apparent in the case of mobile wireless nodes, to reduce the multicast operations that are related to classical ND ([RFC4861], [RFC4862]) and plague the wireless medium.

This specification updates and generalizes 6LoWPAN ND to a broader range of Low power and Lossy Networks (LLNs) with a solid support for Duplicate Address Detection (DAD) and address lookup that does not require broadcasts over the LLNs. The term LLN is used loosely in this specification to cover multiple types of WLANs and WPANs, including Low-Power Wi-Fi, BLUETOOTH(R) Low Energy, IEEE802.11AH and IEEE802.15.4 wireless meshes, so as to address the requirements listed in Appendix A.3

The scope of this draft is a Backbone Link that federates multiple LLNs as a single IPv6 MultiLink Subnet. Each LLN in the subnet is anchored at an IPv6 Backbone Router (6BBR). The Backbone Routers interconnect the LLNs over the Backbone Link and emulate that the LLN nodes are present on the Backbone using proxy-ND operations. This specification extends IPv6 ND over the backbone to discriminate address movement from duplication and eliminate stale state in the backbone routers and backbone nodes once a LLN node has roamed. This way, mobile nodes may roam rapidly from a 6BBR to the next and requirements in Appendix A.1 are met.

This specification can be used by any wireless node to associate at Layer-3 with a 6BBR and register its IPv6 addresses to obtain routing services including proxy-ND operations over the backbone, effectively providing a solution to the requirements expressed in Appendix A.4.

The Link Layer Address (LLA) that is returned as Target LLA (TLA) in Neighbor Advertisements (NA) messages by the 6BBR on behalf of the

Registered Node over the backbone may be that of the Registering Node, in which case the 6BBR needs to bridge the unicast packets (Bridging proxy), or that of the 6BBR on the backbone, in which case the 6BBRs needs to route the unicast packets (Routing proxy). In the latter case, the 6BBR may maintain the list of correspondents to which it has advertised its own MAC address on behalf of the LLN node and the IPv6 ND operation is minimized as the number of nodes scale up in the LLN. This enables to meet the requirements in Appendix A.6 as long as the 6BBRs are dimensioned for the number of registration that each needs to support.

In the context of the the TimeSlotted Channel Hopping (TSCH) mode of [IEEE802154], the 6TiSCH architecture [I-D.ietf-6tisch-architecture] introduces how a 6LoWPAN ND host could connect to the Internet via a RPL mesh Network, but this requires additions to the 6LoWPAN ND protocol to support mobility and reachability in a secured and manageable environment. This specification details the new operations that are required to implement the 6TiSCH architecture and serves the requirements listed in Appendix A.2.

In the case of Low-Power IEEE802.11, a 6BBR may be collocated with a standalone AP or a CAPWAP [RFC5415] wireless controller, and the wireless client (STA) leverages this specification to register its IPv6 address(es) to the 6BBR over the wireless medium. In the case of a 6TiSCH LLN mesh, the RPL root is collocated with a 6LoWPAN Border Router (6LBR), and either collocated with or connected to the 6BBR over an IPv6 Link. The 6LBR leverages this specification to register the LLN nodes on their behalf to the 6BBR. In the case of BTLE, the 6BBR is collocated with the router that implements the BTLE central role as discussed in section 2.2 of [I-D.ietf-6lo-btle].

3. 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 "Multi-link Subnet Support in IPv6" [I-D.ietf-ipv6-multilink-subnets].

Readers would benefit from reading "Multi-Link Subnet Issues" [RFC4903], "Mobility Support in IPv6" [RFC6275], "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 [I-D.ietf-6tisch-terminology], and introduces the following terminology:

LLN Low Power Lossy Network. Used loosely in this specification to represent WLANs and WPANs. See [RFC4919]

Backbone This is an IPv6 transit 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 MultiLink Subnet.

Registration The process during which a wireless Node registers its address(es) with the Border Router so the 6BBR can proxy ND for it over the backbone.

Binding The state in the 6BBR that associates an IP address with a MAC address, a port and some other information about the node that owns the IP address.

Registered Node The node for which the registration is performed, which owns the fields in the EARO option.

Registering Node The node that performs the registration to the 6BBR, either for one of its own addresses, in which case it is Registered Node and indicates its own MAC Address as SLLA in the NS(ARO), or on behalf of a Registered Node that is reachable over a LLN mesh. In the latter case, if the Registered Node is reachable from the 6BBR over a Mesh-Under mesh, the Registering Node indicates the MAC Address of the Registered Node as SLLA in the NS(ARO). Otherwise, it is expected that the Registered Device is reachable over a Route-Over mesh from the Registering Node, in which case the SLLA in the NS(ARO) is that of the Registering Node, which causes it to attract the packets from the 6BBR to the Registered Node and route them over the LLN.

Registered Address The address owned by the Registered Node node that is being registered.

Sleeping Proxy A 6BBR acts as a Sleeping Proxy if it answers ND Neighbor Solicitation over the backbone on behalf of the Registered Node whenever possible. This is the default mode for this specification but it may be overridden, for instance by configuration, into Unicasting Proxy.

Unicasting Proxy As a Unicasting Proxy, the 6BBR forwards NS messages to the Registering Node, transforming Layer-2 multicast into unicast whenever possible.

Routing proxy A 6BBR acts as a routing proxy if it advertises its own MAC address, as opposed to that of the node that performs the registration, as the TLLA in the proxied NAs over the backbone. In that case, the MAC address of the node is not visible at Layer-2 over the backbone and the bridging fabric is not aware of the addresses of the LLN devices and their mobility. The 6BBR installs a connected host route towards the registered node over the interface to the node, and acts as a Layer-3 router for unicast packets to the node. The 6BBR updates the ND Neighbor Cache Entries (NCE) in correspondent nodes if the wireless node moves and registers to another 6BBR, either with a single broadcast, or with a series of unicast NA(O) messages, indicating the TLLA of the new router.

Bridging proxy A 6BBR acts as a bridging proxy if it advertises the MAC address of the node that performs the registration as the TLLA in the proxied NAs over the backbone. In that case, the MAC address and the mobility of the node is still visible across the bridged backbone fabric, as is traditionally the case with Layer-2 APs. The 6BBR acts as a Layer-2 bridge for unicast packets to the registered node. The MAC address exposed in the S/TLLA is that of the Registering Node, which is not necessarily the Registered Device. When a device moves within a LLN mesh, it may end up attached to a different 6LBR acting as Registering Node, and the LLA that is exposed over the backbone will change.

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.

4. Overview

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, transparently.

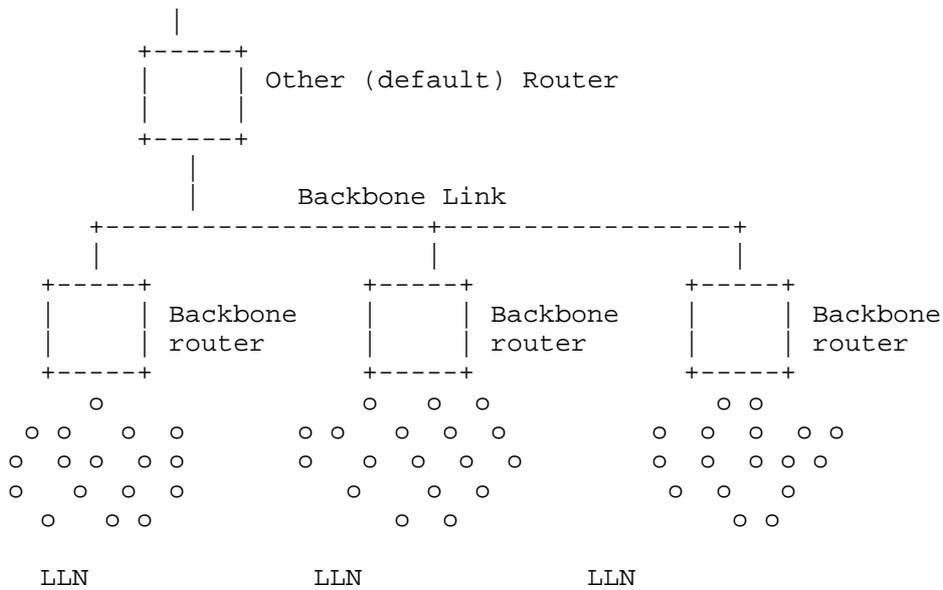


Figure 1: Backbone Link and Backbone Routers

The Backbone Routers maintain an abstract Binding Table of their Registered Nodes. The Binding Table operates as a distributed database of all the wireless Nodes whether they reside on the LLNs or on the backbone, and use an extension to the Neighbor Discovery Protocol to exchange that information across the Backbone in the classical ND reactive fashion.

The Address Registration Option (ARO) defined in [RFC6775] is extended to enable the registration for routing and proxy Neighbor Discovery operations by the 6BBR, and the Extended ARO (EARO) option is included in the ND exchanges over the backbone between the 6BBRs to sort out duplication from movement.

Address duplication is sorted out with the Owner Unique-ID field in the EARO, which is a generalization of the EUI-64 that allows different types of unique IDs beyond the name space derived from the MAC addresses. First-Come First-Serve rules apply, whether the

duplication happens between LLN nodes as represented by their respective 6BBRs, or between an LLN node and a classical node that defends its address over the backbone with classical ND and does not include the EARO option.

In case of conflicting registrations to multiple 6BBRs from a same node, a sequence counter called Transaction ID (TID) is introduced that enables 6BBRs to sort out the latest anchor for that node. Registrations with a same TID are compatible and maintained, but, in case of different TIDs, only the freshest registration is maintained and the stale state is eliminated.

With this specification, Backbone Routers perform ND proxy over the Backbone Link on behalf of their Registered Nodes. The Backbone Router operation is essentially similar to that of a Mobile IPv6 (MIPv6) [RFC6275] Home Agent. This enables mobility support for LLN nodes that would move outside of the network delimited by the Backbone link attach to a Home Agent from that point on. This also enables collocation of Home Agent functionality within Backbone Router functionality on the same backbone interface of a router. Further specification may extend this by allowing the 6BBR to redistribute host routes in routing protocols that would operate over the backbone, or in MIPv6 or the Locator/ID Separation Protocol (LISP) [RFC6830] to support mobility on behalf of the nodes, etc...

The Optimistic Duplicate Address Detection [RFC4429] (ODAD) specification details how an address can be used before a Duplicate Address Detection (DAD) is complete, and insists that an address that is TENTATIVE should not be associated to a Source Link-Layer Address Option in a Neighbor Solicitation message. This specification leverages ODAD to create a temporary proxy state in the 6BBR till DAD is completed over the backbone. This way, the specification enables to distribute proxy states across multiple 6BBR and co-exist with classical ND over the backbone.

5. New Types And Formats

5.1. Transaction ID

The specification expects that the Registered Node can provide a sequence number called Transaction ID (TID) that is incremented with each re-registration. The TID essentially obeys the same rules as the Path Sequence field in the Transit Information Option (TIO) found in RPL's Destination Advertisement Object (DAO). This way, the LLN node can use the same counter for ND and RPL, and a 6LBR acting as RPL root may easily maintain the registration on behalf of a RPL node deep inside the mesh by simply using the RPL TIO Path Sequence as TID for EARO.

When a Registered Node is registered to multiple BBRs in parallel, it is expected that the same TID is used, to enable the 6BBRs to correlate the registrations as being a single one, and differentiate that situation from a movement.

If the TIDs are different, the resolution inherited from RPL sorts out the most recent registration and other ones are removed. The operation for computing and comparing the Path Sequence is detailed in section 7 of [RFC6550] and applies to the TID in the exact same fashion.

5.2. Owner Unique ID

The Owner Unique ID (OUID) enables to differentiate a real duplicate address registration from a double registration or a movement. An ND message from the 6BBR over the backbone that is proxied on behalf of a Registered Node must carry the most recent EARO option seen for that node. A NS/NA with an EARO and a NS/NA without a EARO thus represent different nodes and if they relate to a same target then they reflect an address duplication. The Owner Unique ID can be as simple as a EUI-64 burn-in address, if duplicate EUI-64 addresses are avoided.

Alternatively, the unique ID can be a cryptographic string that can be used to prove the ownership of the registration as discussed in Address Protected Neighbor Discovery for Low-power and Lossy Networks [I-D.sarikaya-6lo-ap-nd].

In any fashion, it is recommended that the node stores the unique Id or the keys used to generate that ID in persistent memory. Otherwise, it will be prevented to re-register after a reboot that would cause a loss of memory until the Backbone Router times out the registration.

5.3. The Enhanced Address Registration Option (EARO)

With the ARO option defined in 6LoWPAN ND [RFC6775], the address being registered and its owner can be uniquely identified and matched with the Binding Table entries of each Backbone Router.

The Enhanced Address Registration Option (EARO) is intended to be used as a replacement to the ARO option within Neighbor Discovery NS and NA messages between a LLN node and its 6LoWPAN Router (6LR), as well as in Duplicate Address Request (DAR) and the Duplicate Address Confirmation (DAC) messages between 6LRs and 6LBRs in LLNs meshes such as 6TiSCH networks.

An NS message with an EARO option is a registration if and only if it also carries an SLLAO option. The AERO option also used in NS and NA messages between Backbone Routers over the backbone link to sort out the distributed registration state, and in that case, it does not carry the SLLAO option and is not confused with a registration.

The EARO extends the ARO and is recognized by the setting of the TID bit. A node that supports this specification MUST always use an EARO as a replacement to an ARO in its registration to a router. This is harmless since the TID bit and fields are reserved in [RFC6775] are ignored by a legacy router. A router that supports this specification answers to an ARO with an ARO and to an EARO with an EARO.

This specification changes the behavior of the peers in a registration flows. To enable backward compatibility, a node that registers to a router that is not known to support this specification MUST behave as prescribed by [RFC6775]. Once the router is known to support this specification, the node MUST obey this specification.

When using the EARO option, the address being registered is found in the Target Address field of the NS and NA messages. This differs from 6LOWPAN ND [RFC6775] which specifies that the address being registered is the source of the NS.

The reason for this change is to enable proxy-registrations on behalf of other nodes in Route-Over meshes, for instance to enable that a RPL root registers addresses on behalf LLN nodes that are deeper in a 6TiSCH mesh. In that case, the Registering Node MUST indicate its own address as source of the ND message and its MAC address in the Source Link-Layer Address Option (SLLAO), since it still expects to get the packets and route them down the mesh. But the Registered Address belongs to another node, the Registered Node, and that address is indicated in the Target Address field of the NS message.

One way of achieving all the above is for a node to first register an address that it owns in order to validate that the router supports this specification, placing the same address in the Source and Target Address fields of the NS message. The node may for instance register an address that is based on EUI-64. For such address, DAD is not required and using the SLLAO option in the NS is actually more amenable with older ND specifications such as ODAD [RFC4429].

Once that first registration is complete, the node knows from the setting of the TID in the response whether the router supports this specification. If this is verified, the node may register other addresses that it owns, or proxy-register addresses on behalf some another node, indicating those addresses being registered in the

Target Address field of the NS messages, while using one of its own, already registered, addresses as source.

The format of the EARO option is as follows:

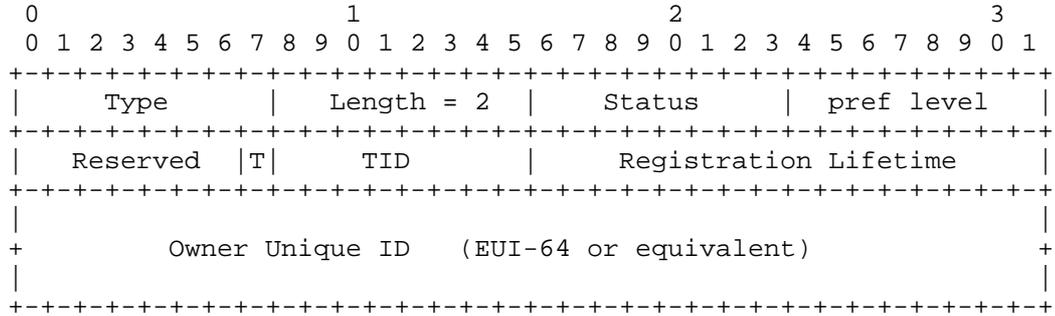


Figure 2: EARO

Option Fields

Type:

Length: 2

Status: OK=0; Duplicate=1; Full=2; Moved=3; Removed=4;

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

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

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

Registration Lifetime: 1-byte integer; expressed in minutes. 0 means that the registration has ended and the state should be removed.

Owner Unique Identifier: A globally unique identifier for the node associated. This can be the EUI-64 derived IID of an interface, or some provable ID obtained cryptographically.

6. Backbone Router Routing Operations

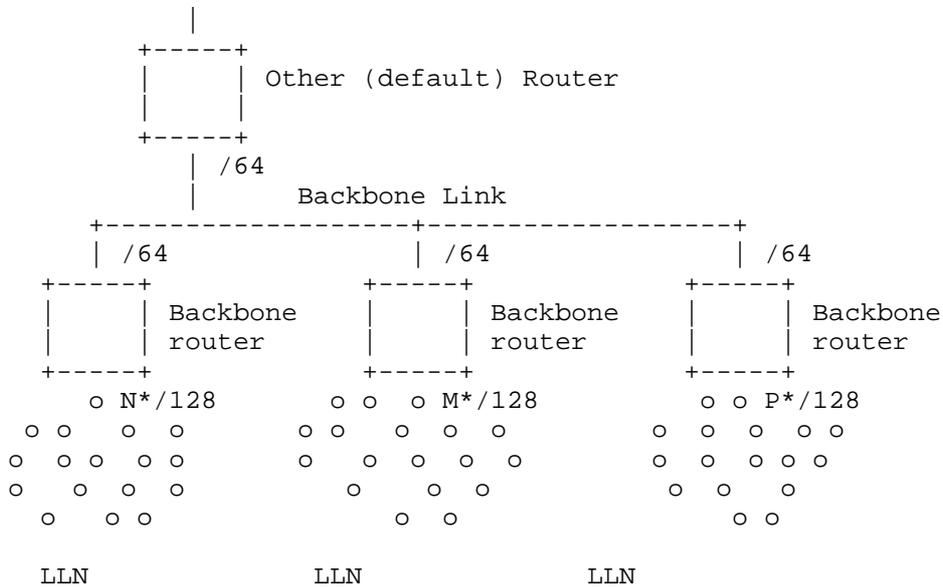


Figure 3: Routing Configuration in the ML Subnet

6.1. Over the Backbone Link

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 LLN interfaces.

The backbone is expected to be a high speed, reliable Backbone link, with affordable and reliable multicast capabilities, such as a bridged Ethernet Network, and to allow 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, and the scalability of the whole subnet requires that broadcast operations are avoided as much as possible on the backbone as well. Unless configured otherwise, the Backbone Router MUST echo the MTU that it learns in RAs over the backbone in the RAs that it sends towards the LLN links.

As a router, the Backbone Router behaves like any other IPv6 router on the backbone side. It has a connected route installed towards the

backbone for the prefixes that are present on that backbone and that it proxies for on the LLN interfaces.

As a proxy, the 6BBR uses an EARO option in the NS-DAD and the multicast NA messages that it generates on behalf of a Registered Node, and it places an EARO in its unicast NA messages if and only if the NS/NA that stimulates it had an EARO in it.

When possible, the 6BBR SHOULD use unicast or solicited-node multicast address (SNMA) [RFC4291] to defend its Registered Addresses over the backbone. In particular, the 6BBR MUST join the SNMA group that corresponds to a Registered Address as soon as it creates an entry for that address and as long as it maintains that entry, whatever the state of the entry. The expectation is that it is possible to get a message delivered to all the nodes on the backbone that listen to a particular address and support this specification - which includes all the 6BBRs in the MultiLink Subnet - by sending a multicast message to the associated SNMA over the backbone.

The support of Optimistic DAD (ODAD) [RFC4429] is recommended for all nodes in the backbone and followed by the 6BBRs in their proxy activity over the backbone. With ODAD, any optimistic node MUST join the SNMA of a Tentative address, which interacts better with this specification.

This specification allows the 6BBR in Routing Proxy mode to advertise the Registered IPv6 Address with the 6BBR Link Layer Address, and attempts to update Neighbor Cache Entries (NCE) in correspondent nodes over the backbone, using gratuitous NA(Override). This method may fail if the multicast message is not properly received, and correspondent nodes may maintain an incorrect neighbor state, which they will eventually discover through Neighbor Unreachability Detection (NUD). Because mobility may be slow, the NUD procedure defined in [RFC4861] may be too impatient, and the support of [RFC7048] is recommended in all nodes in the network.

Since the MultiLink Subnet may grow very large in terms of individual IPv6 addresses, multicasts should be avoided as much as possible even on the backbone. Though it is possible for plain hosts to participate with legacy IPv6 ND support, the support by all nodes connected to the backbone of [I-D.nordmark-6man-rs-refresh] is recommended, and this implies the support of [RFC7559] as well.

6.2. Over the LLN Link

As a router, the Nodes and Backbone Router operation on the LLN follows [RFC6775]. Per that specification, LLN Hosts generally do not depend on multicast RAs to discover routers. It is still

generally required for LLN nodes to accept multicast RAs [I-D.ietf-v6ops-reducing-ra-energy-consumption], but those are rare on the LLN link. Nodes are expected to follow the Simple Procedures for Detecting Network Attachment in IPv6 [RFC6059] (DNA procedures) to assert movements, and to support the Packet-Loss Resiliency for Router Solicitations [RFC7559] to make the unicast RS more reliable.

The Backbone Router acquires its states about the addresses on the LLN side through a registration process from either the nodes themselves, or from a node such as a RPL root / 6LBR (the Registering Node) that performs the registration on behalf of the address owner (the Registered Node).

When operating as a Routing Proxy, the router installs hosts routes (/128) to the Registered Addresses over the LLN links, via the Registering Node as identified by the Source Address and the SLLAO option in the NS(EARO) messages.

In that mode, the 6BBR handles the ND protocol over the backbone on behalf of the Registered Nodes, using its own MAC address in the TLLA and SLLA options in proxied NS and NA messages. It results that for each Registered Address, a number of peer Nodes on the backbone have resolved the address with the 6BBR MAC address and keep that mapping stored in their Neighbor cache.

The 6BBR SHOULD maintain, per Registered Address, the list of the peers on the backbone to which it answered with its MAC address, and when a binding moves to a different 6BBR, it SHOULD send a unicast gratuitous NA(O) individually to each of them to inform them that the address has moved and pass the MAC address of the new 6BBR in the TLLAO option. If the 6BBR can not maintain that list, then it SHOULD remember whether that list is empty or not and if not, send a multicast NA(O) to all nodes to update the impacted Neighbor Caches with the information from the new 6BBR.

The Bridging Proxy is a variation where the BBR function is implemented in a Layer-3 switch or an wireless Access Point that acts as a Host from the IPv6 standpoint, and, in particular, does not operate the routing of IPv6 packets. In that case, the SLLAO in the proxied NA messages is that of the Registering Node and classical bridging operations take place on data frames.

If a registration moves from one 6BBR to the next, but the Registering Node does not change, as indicated by the S/TLLAO option in the ND exchanges, there is no need to update the Neighbor Caches in the peers Nodes on the backbone. On the other hand, if the LLAO changes, the 6BBR SHOULD inform all the relevant peers as described above, to update the impacted Neighbor Caches. In the same fashion,

if the Registering Node changes with a new registration, the 6BBR SHOULD also update the impacted Neighbor Caches over the backbone.

7. Backbone Router 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 node on the backbone to reach a Registered Node as if it was on-link. The backbone and the LLNs are considered different Links in a MultiLink subnet but the prefix that is used may still be advertised as on-link on the backbone to support legacy nodes; multicast ND messages are link-scoped and not forwarded across the backbone routers.

ND Messages on the backbone side that do not match to a registration on the LLN side are not acted upon on the LLN side, which stands protected. On the LLN side, the prefixes associated to the MultiLink Subnet are presented as not on-link, so address resolution for other hosts do not occur.

The default operation in this specification is Sleeping proxy which means:

- o creating a new entry in an abstract Binding Table for a new Registered Address and validating that the address is not a duplicate over the backbone
- o defending a Registered Address over the backbone using NA messages with the Override bit set on behalf of the sleeping node whenever possible
- 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 Solicitation messages.
- o Forwarding packets from the LLN over the backbone, and the other way around.
- o Eventually triggering a liveness verification of a stale registration.

A 6BBR may act as a Sleeping Proxy only if the state of the binding entry is REACHABLE, or TENTATIVE in which case the answer is delayed.

In any other state, the Sleeping Proxy operates as a Unicasting Proxy.

As a Unicasting Proxy, the 6BBR forwards NS messages to the Registering Node, transforming Layer-2 multicast into unicast whenever possible. This is not possible in UNREACHABLE state, so the NS messages are multicasted, and rate-limited to protect the medium with an exponential back-off. In other states, The messages are forwarded to the Registering Node as unicast Layer-2 messages. In TENTATIVE state, the NS message is either held till DAD completes, or dropped.

The draft introduces the optional concept of primary and secondary BBRs. The primary is the backbone router that has the highest EUI-64 address of all the 6BBRs that share a registration for a same Registered Address, with the same Owner Unique ID and same Transaction ID, the EUI-64 address being considered as an unsigned 64bit integer. The concept is defined with the granularity of an address, that is a given 6BBR can be primary for a given address and secondary or another one, regardless 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 6BBR may claim the address over the backbone, since they are all capable to route from the backbone to the LLN node, and the address appears on the backbone as an anycast address.

The Backbone Routers maintain a distributed binding table, using classical ND over the backbone to detect duplication. This specification requires that:

1. All addresses that can be reachable from the backbone, including IPv6 addresses based on burn-in EUI64 addresses MUST be registered to the 6BBR.
2. A Registered Node MUST include the EARO option in an NS message that used to register an addresses to a 6LR; the 6LR MUST propagate that option unchanged to the 6LBR in the DAR/DAC exchange, and the 6LBR MUST propagate that option unchanged in proxy registrations.
3. The 6LR MUST echo the same EARO option in the NA that it uses to respond, but for the status filed which is not used in NS messages, and significant in NA.

A false positive duplicate detection may arise over the backbone, for instance if the Registered Address is registered 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.

The expectation in this specification is that there is a single Registering Node at a time per Backbone Router for a given Registered Address, but that a Registered Address may be registered to Multiple 6BBRs for higher availability.

Over the LLN, and for any given Registered Address, it is REQUIRED that:

- de-registrations (newer TID, same OUID, null Lifetime) are accepted and responded immediately with a status of 4; the entry is deleted;

- newer registrations (newer TID, same OUID, non-null Lifetime) are accepted and responded with a status of 0 (success); the entry is updated with the new TID, the new Registration Lifetime and the new Registering Node, if any has changed; in TENTATIVE state the response is held and may be overwritten; in other states the Registration-Lifetime timer is restarted and the entry is placed in REACHABLE state.

- identical registrations (same TID, same OUID) from a same Registering Node are not processed but responded with a status of 0 (success); they are expected to be identical and an error may be logged if not; in TENTATIVE state, the response is held and may be overwritten, but it MUST be eventually produced and it carries the result of the DAD process;

- older registrations (not(newer or equal) TID, same OUID) from a same Registering Node are ignored;

- identical and older registrations (not-newer TID, same OUID) from a different Registering Node are responded immediately with a status of 3 (moved); this may be rate limited to protect the medium;

- and any registration for a different Registered Node (different OUID) are responded immediately with a status of 1 (duplicate).

7.1. Registration and Binding State Creation

Upon a registration for a new address with an NS(EARO), the 6BBR performs a DAD operation over the backbone placing the new address as target in the NS-DAD message. The EARO from the registration MUST be

placed unchanged in the NS-DAD message, and an entry is created in TENTATIVE state for a duration of TENTATIVE_DURATION. The NS-DAD message is sent multicast over the backbone to the SNMA address associated with the registered address. If that operation is known to be costly, and the 6BBR has an indication from another source (such as a NCE) that the Registered Address was present on the backbone, that information may be leveraged to send the NS-DAD message as a Layer-2 unicast to the MAC that was associated with the Registered Address.

In TENTATIVE state:

- o the entry is removed if an NA is received over the backbone for the Registered Address with no EARO option, or with an EARO option with a status of 1 (duplicate) that indicates an existing registration for another LLN node. The OUID and TID fields in the EARO option received over the backbone are ignored. A status of 1 is returned in the EARO option of the NA back to the Registering Node;
- o the entry is also removed if an NA with an ARO option with a status of 3 (moved), or a NS-DAD with an ARO option that indicates a newer registration for the same Registered Node, is received over the backbone for the Registered Address. A status of 3 is returned in the NA(EARO) back to the Registering Node;
- o when a registration is updated but not deleted, e.g. from a newer registration, the DAD process on the backbone continues and the running timers are not restarted;
- o Other NS (including DAD with no EARO option) and NA from the backbone are not responded in TENTATIVE state, but the list of their origins may be kept in memory and if so, the 6BBR may send them each a unicast NA with eventually an EARO option when the TENTATIVE_DURATION timer elapses, so as to cover legacy nodes that do not support ODAD.
- o When the TENTATIVE_DURATION timer elapses, a status 0 (success) is returned in a NA(EARO) back to the Registering Node(s), and the entry goes to REACHABLE state for the Registration Lifetime; the DAD process is successful and the 6BBR MUST send a multicast NA(EARO) to the SNMA associated to the Registered Address over the backbone with the Override bit set so as to take over the binding from other 6BBRs.

7.2. Defending Addresses

If a 6BBR has an entry in REACHABLE state for a Registered Address:

- o If the 6BBR is primary, or does not support the concept, it MUST defend that address over the backbone upon an incoming NS-DAD, either if the NS does not carry an EARO, or if an EARO is present that indicates a different Registering Node (different OUID). The 6BBR sends a NA message with the Override bit set and the NA carries an EARO option if and only if the NS-DAD did so. When present, the EARO in the NA(O) that is sent in response to the NS-DAD(EARO) carries a status of 1 (duplicate), and the OUID and TID fields in the EARO option are obfuscated with null or random values to avoid network scanning and impersonation attacks.
- o If the 6BBR receives an NS-DAD(EARO) that reflect a newer registration, the 6BBR updates the entry and the routing state to forward packets to the new 6BBR, but keeps the entry REACHABLE. In that phase, it MAY use REDIRECT messages to reroute traffic for the Registered Address to the new 6BBR.
- o If the 6BBR receives an NA(EARO) that reflect a newer registration, the 6BBR removes its entry and sends a NA(AERO) with a status of 3 (moved) to the Registering Node, if the Registering Node is different from the Registered Node. If necessary, the 6BBR cleans up ND cache in peers nodes as discussed in Section 6.1, by sending a series of unicast to the impacted nodes, or one broadcast NA(O) to all-nodes.
- o If the 6BBR received a NS(LOOKUP) for a Registered Address, it answers immediately with an NA on behalf of the Registered Node, without polling it. There is no need of an EARO in that exchange.
- o When the Registration-Lifetime timer elapses, the entry goes to STALE state for a duration of STABLE_STALE_DURATION in LLNs that keep stable addresses such as LWPANs, and UNSTABLE_STALE_DURATION in LLNs where addresses are renewed rapidly, e.g. for privacy reasons.

The STALE state is a chance to keep track of the backbone peers that may have an ND cache pointing on this 6BBR in case the Registered Address shows back up on this or a different 6BBR at a later time. In STALE state:

- o If the Registered Address is claimed by another node on the backbone, with an NS-DAD or an NA, the 6BBR does not defend the address. Upon an NA(O), or the stale time elapses, the 6BBR

removes its entry and sends a NA(AERO) with a status of 4 (removed) to the Registering Node.

- o If the 6BBR received a NS(LOOKUP) for a Registered Address, the 6BBR MUST send an NS(NUD) following rules in [RFC7048] to the registering Node targeting the Registered Address prior to answering. If the NUD succeeds, the operation in REACHABLE state applies. If the NUD fails, the 6BBR refrains from answering the lookup. The NUD expected to be mapped by the Registering Node into a liveness validation of the Registered Node if they are in fact different nodes.

8. 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 Broadcast 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. This specification RECOMMENDS the use of additional protection against address theft such as provided by [I-D.sarikaya-6lo-ap-nd], which guarantees the ownership of the OUID.

When the ownership of the OUID cannot be assessed, this specification limits the cases where the OUID and the TID are multicasted, and obfuscates them in responses to attempts to take over an address.

9. Protocol Constants

This Specification uses the following constants:

TENTATIVE_DURATION:	800 milliseconds
STABLE_STALE_DURATION:	24 hours
UNSTABLE_STALE_DURATION:	5 minutes
DEFAULT_NS_POLLING:	3 times

10. IANA Considerations

This document requires the following additions:

Address Registration Option Status Values Registry

Status	Description
3	Moved: The registration fails because it is not the freshest.
4	Removed: The binding state was removed

IANA is required to change the registry accordingly

Table 1: New ARO Status values

11. Acknowledgments

Kudos to Eric Levy-Abegnoli who designed the First Hop Security infrastructure at Cisco.

12. References

12.1. 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>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<http://www.rfc-editor.org/info/rfc4291>>.
- [RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, DOI 10.17487/RFC4429, April 2006, <<http://www.rfc-editor.org/info/rfc4429>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<http://www.rfc-editor.org/info/rfc4861>>.

- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<http://www.rfc-editor.org/info/rfc4862>>.
- [RFC6059] Krishnan, S. and G. Daley, "Simple Procedures for Detecting Network Attachment in IPv6", RFC 6059, DOI 10.17487/RFC6059, November 2010, <<http://www.rfc-editor.org/info/rfc6059>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., 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, DOI 10.17487/RFC6550, March 2012, <<http://www.rfc-editor.org/info/rfc6550>>.
- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<http://www.rfc-editor.org/info/rfc6775>>.

12.2. Informative References

- [I-D.chakrabarti-nordmark-6man-efficient-nd]
Chakrabarti, S., Nordmark, E., Thubert, P., and M. Wasserman, "IPv6 Neighbor Discovery Optimizations for Wired and Wireless Networks", draft-chakrabarti-nordmark-6man-efficient-nd-07 (work in progress), February 2015.
- [I-D.delcarpio-6lo-wlanah]
Vega, L., Robles, I., and R. Morabito, "IPv6 over 802.11ah", draft-delcarpio-6lo-wlanah-01 (work in progress), October 2015.
- [I-D.ietf-6lo-6lobac]
Lynn, K., Martocci, J., Neilson, C., and S. Donaldson, "Transmission of IPv6 over MS/TP Networks", draft-ietf-6lo-6lobac-03 (work in progress), October 2015.
- [I-D.ietf-6lo-btle]
Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "IPv6 over BLUETOOTH(R) Low Energy", draft-ietf-6lo-btle-17 (work in progress), August 2015.

- [I-D.ietf-6lo-dect-ule]
Mariager, P., Petersen, J., Shelby, Z., Logt, M., and D. Barthel, "Transmission of IPv6 Packets over DECT Ultra Low Energy", draft-ietf-6lo-dect-ule-03 (work in progress), September 2015.
- [I-D.ietf-6lo-nfc]
Youn, J. and Y. Hong, "Transmission of IPv6 Packets over Near Field Communication", draft-ietf-6lo-nfc-02 (work in progress), October 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-08 (work in progress), May 2015.
- [I-D.ietf-6tisch-terminology]
Palattella, M., Thubert, P., Watteyne, T., and Q. Wang, "Terminology in IPv6 over the TSCH mode of IEEE 802.15.4e", draft-ietf-6tisch-terminology-06 (work in progress), November 2015.
- [I-D.ietf-bier-architecture]
Wijnands, I., Rosen, E., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast using Bit Index Explicit Replication", draft-ietf-bier-architecture-02 (work in progress), July 2015.
- [I-D.ietf-ipv6-multilink-subnets]
Thaler, D. and C. Huitema, "Multi-link Subnet Support in IPv6", draft-ietf-ipv6-multilink-subnets-00 (work in progress), July 2002.
- [I-D.ietf-roll-terminology]
Vasseur, J., "Terms used in Routing for Low power And Lossy Networks", draft-ietf-roll-terminology-13 (work in progress), October 2013.
- [I-D.ietf-v6ops-reducing-ra-energy-consumption]
Yourtchenko, A. and L. Colitti, "Reducing energy consumption of Router Advertisements", draft-ietf-v6ops-reducing-ra-energy-consumption-03 (work in progress), November 2015.
- [I-D.nordmark-6man-dad-approaches]
Nordmark, E., "Possible approaches to make DAD more robust and/or efficient", draft-nordmark-6man-dad-approaches-02 (work in progress), October 2015.

- [I-D.nordmark-6man-rs-refresh]
Nordmark, E., Yourtchenko, A., and S. Krishnan, "IPv6 Neighbor Discovery Optional Unicast RS/RA Refresh", draft-nordmark-6man-rs-refresh-01 (work in progress), October 2014.
- [I-D.popa-6lo-6loplc-ipv6-over-ieee19012-networks]
Popa, D. and J. Hui, "6LoPLC: Transmission of IPv6 Packets over IEEE 1901.2 Narrowband Powerline Communication Networks", draft-popa-6lo-6loplc-ipv6-over-ieee19012-networks-00 (work in progress), March 2014.
- [I-D.sarikaya-6lo-ap-nd]
Sarikaya, B. and P. Thubert, "Address Protected Neighbor Discovery for Low-power and Lossy Networks", draft-sarikaya-6lo-ap-nd-01 (work in progress), October 2015.
- [I-D.vyncke-6man-mcast-not-efficient]
Vyncke, E., Thubert, P., Levy-Abegnoli, E., and A. Yourtchenko, "Why Network-Layer Multicast is Not Always Efficient At Datalink Layer", draft-vyncke-6man-mcast-not-efficient-01 (work in progress), February 2014.
- [I-D.yourtchenko-6man-dad-issues]
Yourtchenko, A. and E. Nordmark, "A survey of issues related to IPv6 Duplicate Address Detection", draft-yourtchenko-6man-dad-issues-01 (work in progress), March 2015.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <<http://www.rfc-editor.org/info/rfc3315>>.
- [RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <<http://www.rfc-editor.org/info/rfc3810>>.
- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "Secure Neighbor Discovery (SEND)", RFC 3971, DOI 10.17487/RFC3971, March 2005, <<http://www.rfc-editor.org/info/rfc3971>>.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <<http://www.rfc-editor.org/info/rfc3972>>.

- [RFC4389] Thaler, D., Talwar, M., and C. Patel, "Neighbor Discovery Proxies (ND Proxy)", RFC 4389, DOI 10.17487/RFC4389, April 2006, <<http://www.rfc-editor.org/info/rfc4389>>.
- [RFC4903] Thaler, D., "Multi-Link Subnet Issues", RFC 4903, DOI 10.17487/RFC4903, June 2007, <<http://www.rfc-editor.org/info/rfc4903>>.
- [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, DOI 10.17487/RFC4919, August 2007, <<http://www.rfc-editor.org/info/rfc4919>>.
- [RFC5415] Calhoun, P., Ed., Montemurro, M., Ed., and D. Stanley, Ed., "Control And Provisioning of Wireless Access Points (CAPWAP) Protocol Specification", RFC 5415, DOI 10.17487/RFC5415, March 2009, <<http://www.rfc-editor.org/info/rfc5415>>.
- [RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <<http://www.rfc-editor.org/info/rfc6275>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.
- [RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, DOI 10.17487/RFC6830, January 2013, <<http://www.rfc-editor.org/info/rfc6830>>.
- [RFC7048] Nordmark, E. and I. Gashinsky, "Neighbor Unreachability Detection Is Too Impatient", RFC 7048, DOI 10.17487/RFC7048, January 2014, <<http://www.rfc-editor.org/info/rfc7048>>.
- [RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <<http://www.rfc-editor.org/info/rfc7217>>.

- [RFC7428] Brandt, A. and J. Buron, "Transmission of IPv6 Packets over ITU-T G.9959 Networks", RFC 7428, DOI 10.17487/RFC7428, February 2015, <<http://www.rfc-editor.org/info/rfc7428>>.
- [RFC7559] Krishnan, S., Anipko, D., and D. Thaler, "Packet-Loss Resiliency for Router Solicitations", RFC 7559, DOI 10.17487/RFC7559, May 2015, <<http://www.rfc-editor.org/info/rfc7559>>.

12.3. External Informative References

- [IEEE80211]
IEEE standard for Information Technology, "IEEE Standard for Information technology-- Telecommunications and information exchange between systems Local and metropolitan area networks-- Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications".
- [IEEE802151]
IEEE standard for Information Technology, "IEEE Standard for Information Technology - Telecommunications and Information Exchange Between Systems - Local and Metropolitan Area Networks - Specific Requirements. - Part 15.1: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Wireless Personal Area Networks (WPANs)".
- [IEEE802154]
IEEE standard for Information Technology, "IEEE Standard for Local and metropolitan area networks-- Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs)".

Appendix A. Requirements

This section lists requirements that were discussed at 6lo for an update to 6LoWPAN ND. This specification meets most of them, but those listed in Appendix A.5 which are deferred to a different specification such as [I-D.sarikaya-6lo-ap-nd].

A.1. Requirements Related to Mobility

Due to the unstable nature of LLN links, even in a LLN of immobile nodes a 6LoWPAN Node may change its point of attachment to a 6LR, say 6LR-a, and may not be able to notify 6LR-a. Consequently, 6LR-a may still attract traffic that it cannot deliver any more. When links to

a 6LR change state, there is thus a need to identify stale states in a 6LR and restore reachability in a timely fashion.

Req1.1: Upon a change of point of attachment, connectivity via a new 6LR MUST be restored timely without the need to de-register from the previous 6LR.

Req1.2: For that purpose, the protocol MUST enable to differentiate between multiple registrations from one 6LoWPAN Node and registrations from different 6LoWPAN Nodes claiming the same address.

Req1.3: Stale states MUST be cleaned up in 6LRs.

Req1.4: A 6LoWPAN Node SHOULD also be capable to register its Address to multiple 6LRs, and this, concurrently.

A.2. Requirements Related to Routing Protocols

The point of attachment of a 6LoWPAN Node may be a 6LR in an LLN mesh. IPv6 routing in a LLN can be based on RPL, which is the routing protocol that was defined at the IETF for this particular purpose. Other routing protocols than RPL are also considered by Standard Defining Organizations (SDO) on the basis of the expected network characteristics. It is required that a 6LoWPAN Node attached via ND to a 6LR would need to participate in the selected routing protocol to obtain reachability via the 6LR.

Next to the 6LBR unicast address registered by ND, other addresses including multicast addresses are needed as well. For example a routing protocol often uses a multicast address to register changes to established paths. ND needs to register such a multicast address to enable routing concurrently with discovery.

Multicast is needed for groups. Groups MAY be formed by device type (e.g. routers, street lamps), location (Geography, RPL sub-tree), or both.

The Bit Index Explicit Replication (BIER) Architecture [I-D.ietf-bier-architecture] proposes an optimized technique to enable multicast in a LLN with a very limited requirement for routing state in the nodes.

Related requirements are:

Req2.1: The ND registration method SHOULD be extended in such a fashion that the 6LR MAY advertise the Address of a 6LoWPAN Node over the selected routing protocol and obtain reachability to that Address using the selected routing protocol.

Req2.2: Considering RPL, the Address Registration Option that is used in the ND registration SHOULD be extended to carry enough information to generate a DAO message as specified in [RFC6550] section 6.4, in particular the capability to compute a Path Sequence and, as an option, a RPLInstanceID.

Req2.3: Multicast operations SHOULD be supported and optimized, for instance using BIER or MPL. Whether ND is appropriate for the registration to the 6BBR is to be defined, considering the additional burden of supporting the Multicast Listener Discovery Version 2 [RFC3810] (MLDv2) for IPv6.

A.3. Requirements Related to the Variety of Low-Power Link types

6LoWPAN ND [RFC6775] was defined with a focus on IEEE802.15.4 and in particular the capability to derive a unique Identifier from a globally unique MAC-64 address. At this point, the 6lo Working Group is extending the 6LoWPAN Header Compression (HC) [RFC6282] technique to other link types ITU-T G.9959 [RFC7428], Master-Slave/Token-Passing [I-D.ietf-6lo-6lobac], DECT Ultra Low Energy [I-D.ietf-6lo-dect-ule], Near Field Communication [I-D.ietf-6lo-nfc], IEEE802.11ah [I-D.delcarpio-6lo-wlanah], as well as IEEE1901.2 Narrowband Powerline Communication Networks [I-D.popa-6lo-6loplc-ipv6-over-ieee19012-networks] and BLUETOOTH(R) Low Energy [I-D.ietf-6lo-btle].

Related requirements are:

Req3.1: The support of the registration mechanism SHOULD be extended to more LLN links than IEEE 802.15.4, matching at least the LLN links for which an "IPv6 over foo" specification exists, as well as Low-Power Wi-Fi.

Req3.2: As part of this extension, a mechanism to compute a unique Identifier should be provided, with the capability to form a Link-Local Address that SHOULD be unique at least within the LLN connected to a 6LBR discovered by ND in each node within the LLN.

Req3.3: The Address Registration Option used in the ND registration SHOULD be extended to carry the relevant forms of unique Identifier.

Req3.4: The Neighbour Discovery should specify the formation of a site-local address that follows the security recommendations from [RFC7217].

A.4. Requirements Related to Proxy Operations

Duty-cycled devices may not be able to answer themselves to a lookup from a node that uses classical ND on a backbone and may need a proxy. Additionally, the duty-cycled device may need to rely on the 6LBR to perform registration to the 6BBR.

The ND registration method SHOULD defend the addresses of duty-cycled devices that are sleeping most of the time and not capable to defend their own Addresses.

Related requirements are:

Req4.1: The registration mechanism SHOULD enable a third party to proxy register an Address on behalf of a 6LoWPAN node that may be sleeping or located deeper in an LLN mesh.

Req4.2: The registration mechanism SHOULD be applicable to a duty-cycled device regardless of the link type, and enable a 6BBR to operate as a proxy to defend the registered Addresses on its behalf.

Req4.3: The registration mechanism SHOULD enable long sleep durations, in the order of multiple days to a month.

A.5. Requirements Related to Security

In order to guarantee the operations of the 6LoWPAN ND flows, the spoofing of the 6LR, 6LBR and 6BBRs roles should be avoided. Once a node successfully registers an address, 6LoWPAN ND should provide energy-efficient means for the 6LBR to protect that ownership even when the node that registered the address is sleeping.

In particular, the 6LR and the 6LBR then should be able to verify whether a subsequent registration for a given Address comes from the original node.

In a LLN it makes sense to base security on layer-2 security. During bootstrap of the LLN, nodes join the network after authorization by a Joining Assistant (JA) or a Commissioning Tool (CT). After joining nodes communicate with each other via secured links. The keys for the layer-2 security are distributed by the JA/CT. The JA/CT can be part of the LLN or be outside the LLN. In both cases it is needed that packets are routed between JA/CT and the joining node.

Related requirements are:

Req5.1: 6LoWPAN ND security mechanisms SHOULD provide a mechanism for the 6LR, 6LBR and 6BBR to authenticate and authorize one another for

their respective roles, as well as with the 6LoWPAN Node for the role of 6LR.

Req5.2: 6LoWPAN ND security mechanisms SHOULD provide a mechanism for the 6LR and the 6LBR to validate new registration of authorized nodes. Joining of unauthorized nodes MUST be impossible.

Req5.3: 6LoWPAN ND security mechanisms SHOULD lead to small packet sizes. In particular, the NS, NA, DAR and DAC messages for a re-registration flow SHOULD NOT exceed 80 octets so as to fit in a secured IEEE802.15.4 frame.

Req5.4: Recurrent 6LoWPAN ND security operations MUST NOT be computationally intensive on the LoWPAN Node CPU. When a Key hash calculation is employed, a mechanism lighter than SHA-1 SHOULD be preferred.

Req5.5: The number of Keys that the 6LoWPAN Node needs to manipulate SHOULD be minimized.

Req5.6: The 6LoWPAN ND security mechanisms SHOULD enable CCM* for use at both Layer 2 and Layer 3, and SHOULD enable the reuse of security code that has to be present on the device for upper layer security such as TLS.

Req5.7: Public key and signature sizes SHOULD be minimized while maintaining adequate confidentiality and data origin authentication for multiple types of applications with various degrees of criticality.

Req5.8: Routing of packets should continue when links pass from the unsecured to the secured state.

Req5.9: 6LoWPAN ND security mechanisms SHOULD provide a mechanism for the 6LR and the 6LBR to validate whether a new registration for a given address corresponds to the same 6LoWPAN Node that registered it initially, and, if not, determine the rightful owner, and deny or clean-up the registration that is duplicate.

A.6. Requirements Related to Scalability

Use cases from Automatic Meter Reading (AMR, collection tree operations) and Advanced Metering Infrastructure (AMI, bi-directional communication to the meters) indicate the needs for a large number of LLN nodes pertaining to a single RPL DODAG (e.g. 5000) and connected to the 6LBR over a large number of LLN hops (e.g. 15).

Related requirements are:

Req6.1: The registration mechanism SHOULD enable a single 6LBR to register multiple thousands of devices.

Req6.2: The timing of the registration operation should allow for a large latency such as found in LLNs with ten and more hops.

Author's Address

Pascal Thubert (editor)
Cisco Systems, Inc
Building D
45 Allee des Ormes - BP1200
MOUGINS - Sophia Antipolis 06254
FRANCE

Phone: +33 497 23 26 34
Email: pthubert@cisco.com