6LOWPAN WG
Internet-Draft
Expires: January 15, 2009

S. Chakrabarti
IP Infusion
E. Nordmark
Sun Microsystems, Inc.
July 14, 2008

LowPan Neighbor Discovery Extensions
draft-chakrabarti-6lowpan-ipv6-nd-05.txt

Status of this Memo

Abstract

IETF 6LowPan working group defines IPv6 over low-power personal area network (IEEE 802.15.4). IEEE 802.15.4 link layer does not have multicast support, although it supports broadcast. Due to the nature of LowPan network or sensor networks, broadcast messages should be minimized. This document suggests some optimizations to IPv6 Neighbor Discovery related multicast messages in order to reduce signaling in the low-cost low-powered network.

Table of Contents

1. Introduction

The IPv6-over-IEEE 802.15.4 [LOWPAN] document has specified IPv6 headers carrying over IEEE 802.15.4 network with the help of a adaptation layer which sits between the MAC layer and the network layer. The LowPan network is characterized by low-powered, low bit-rate, short ranged, low cost network. Thus, all-node multicast defined in Neighbor Discovery [ND] is not often desirable in the LowPan network. But IEEE 802.15.4 does not have multicast support, however, it supports broadcast. Broadcast messages could be used in some cases to represent all-node multicast messages, but periodic broadcast messages should be minimized in the LowPan network in order to conserve energy. The goal of this document is to minimize periodic multicast signals used by Neighbor Discovery [ND], minimize total number of Neighbor Discovery related signaling messages without loosing generality of Neighbor Discovery and autoconfiguration. It also aims to identify the default values for periodic advertisements, router and prefix lifetime values that are suitable for LowPan networks.

The IPv6-over-IEEE 802.15.4 [LOWPAN] document provides mesh routing capability at the link layer. Yet each node is configured with IPv6 addresses. Thus a IEEE 802.15.4 may look like one single IPv6 subnet to the IP layer. It may be possible that routing advertisements are used only for prefix advertisement purpose for auto-configuration of IPv6 addresses. Yet, Neighbor Solicitation, Neighbor Advertisements, Neighbor Unreachability Detection (NUD) take place as usual for neighbor to neighbor communication. Also, some LowPan networks may use IPv6 routing (for example, star topology). Hence minimizing periodic router signaling messages are required for efficient use of IPv6 in the LowPan network.

Please note that this version of draft is not complete in determining a solution for reducing the Neighbor Discovery signaling messages; the work is in progress by the authors.

2. Definition Of Terms

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].

PAN co-ordinator:
A PAN co-ordinator is usually an IPv6 router with higher processing and electrical power in one 6lowpan network. This document assumes that there is one PAN co-ordinator per one 6lowpan. PAN co-ordinator may perform other administrative duties

in the network.

Co-ordinator:
Co-ordinators are usually Full functional IEEE 802.15.4 devices (FFD). The co-ordinators can receive IPv6 packets and forward them using L2 layer Mesh routing protocols.

3. Goals

This document aims to reduce signaling messages due to IPv6 Neighbor Discovery protocol, and in particular those messages that are multicast.

A LowPan network does not have multicast capability in the layer 2. However, the link-layer provides broadcast functionality, supporting IPv6 for broadcast would defeat its purpose. Also the Neighbor Discovery document [ND] is designed for regular Internet Network where nodes are sufficiently powered and they are configured in a stable infrastructure of network, unlike LowPan or sensor networks. The strawman idea for this section is to attempt to minimize the multicast Neighbor Discovery signaling packets.

The following are the goals:

1. Minimizing Multicast messages such as:
   * Reduce or avoid multicast for Duplicate Address Detection.
   * Limit multicast Router Solicitations and Router Advertisements.
   * Avoid/Reduce Multicast Neighbor Solicitations.
2. Reduce or avoid (unicast) Neighbor Unreachability Detection messages.

4. Background on IPv6 Neighbor Discovery

IPv6 Neighbor Discovery [ND] provides several important functions such as Router Discovery, Address Resolution, Duplicate Address Detection, Redirect, Prefix and Parameter Discovery. In order to get a feel for how this works we provide a short walk-through of a device initializing on a network. The walk-through uses the "normal" type of usage of Neighbor Discovery on an Ethernet.

During power-on and initialization the following steps are common on Ethernet networks:

1. The host picks a link-local IPv6 address, and checks whether it is unique on the link by multicasting a Neighbor Solicitation as part of the Duplicate Address Detection mechanism. Before it does this, the host joins the solicited-node multicast address on the interface. This joining of the solicited-node group requires

sending a MLD join message in order to handle links where there are MLD-snooping switches.

2. When the host initializes its network it multicasts one or a few Neighbor Solicitation messages to the all-router IPv6 multicast address, until it receives a Router Advertisement.
3. The router(s) that receives the Router Solicitation, responds with a Router Advertisement. Per the ND specification the routers can unicast the RA (if the RS contained a Sender Link-layer Address option), or the router can multicast the RA to the all-nodes IPv6 multicast address.
4. Once the host has received a RA with one or more Prefix options that have the "A" flag set, the host performs stateless address autoconfiguration. For each IPv6 address it forms as part of this, it does Duplicate Address Detection by multicasting a Neighbor Solicitation message to the solicited-node multicast address.

In addition, periodically the routers multicast Router Advertisements to the all-nodes IPv6 multicast address. This is done so that the hosts can determine when a router has been removed (without relying on Neighbor Unreachability Detection), find out about the extension of the lifetimes for the prefixes, and also be informed about new and removed prefixes (i.e. in the case of renumbering) or other configuration changes. By default the periodic router advertisements are approximately at every 7 minutes on an average, and the default router lifetimes is 90 minutes. By default the lifetime for the prefixes are 7 days, thus this doesn't require very frequent advertisements.

When a host attached to the link wants to communicate with another host, if that host's address is part of the on-link prefixes, then the host will multicast a Neighbor Solicitation to find the peer's link layer address. This also applies when the router receives a packet that it needs to forward to a host in the on-link prefix. The response to the NS is a Neighbor Advertisement which is unicast.

If the destination is not part of the on-link prefixes, then the host just sends the packet to one of the default routers; the host has the L2 address of the routers from the Router Advertisement messages it has received.

The Neighbor Discovery standard allows a configuration where the routers do not advertise any prefixes with the "on-link" ("L") flag set, but it can still advertise them with the "stateless address config" ("A") flag. This configuration isn't typically used on Ethernets. Should it be used, the above description still applies; since no destination is part of the on-link prefixes, the packet would be sent to one of the default routers. This would cause the

routers to forward the packet to the destination. Should the destination be on the same link as the source, then the routers can send a Redirect message back to the sender informing it of this fact. Unlike in IPv4, the Redirect message can include the link-layer address of the destination. If this is done, then subsequence messages between the hosts will be sent without involving the router. Please note that RFC4903 section 4.3 discusses isssues with Router-advertisement and cleared "on-link" bit when nodes are present both on-link and off-link.

The information about the neighbors and their L2 addresses are cached in a neighbor cache. The entries in this cache do not need to be timed out, but a node can discard them if it runs low on memory. This means that the information can be stale, in two different ways:

- The layer 2 address has changed, for instance due to somebody replacing the Ethernet NIC, resulting in a different Ethernet address for the node.
- The node is no longer reachable; it could be down.

The nodes detect such unreachability using Neighbor Unreachability Detection, which consists of sending unicast Neighbor Solicitations to the currently known L2 address of the neighbor and expect a unicast Neighbor Advertisement in response. If there is no response after 3 unicast NSes the node will declare the neighbor unreachable and fall back to sending multicast NSes to try to discover a potentially new L2 address for the peer.

The first case can in general appear for both a host and a router; what matters is whether the L2 address changes of the node. The second case can also happen for both a host and a router, but the recovery is different in the two cases. If it is a router that is no longer reachable, and there are multiple default routers on the link, then the unreachability will result in the host trying to use one of the other default routers. But if a host is unreachable there is no alternate path to use. Thus in that case the only benefit of NUD is to be able to generate ICMP errors.

5. Assumptions about Topology and Address Mapping

In order to minimize the multicast packets we need to make some assumptions that tie together the L2 functional elements and the L3 functional elements. We also state our understanding of how IPv6 would map to a LowPan network:

- One PAN-ID defines one LowPan Network.
- Each LowPan corresponds to one IPv6 subnet as PAN-ID may be used to determine a subnet.

- There is one PAN co-ordinator or PAN group leader per one LowPan.
- The IPv6 router which sits in the boundary of the LowPan is a PAN co-ordinator.
- There can be multiple co-ordinators in the LowPan.
- When a device connects to the LowPan at layer 2, it finds out a (unicast) layer 2 address for its co-ordinator.
- By recursive construction, the previous item implies that a co-ordinator knows its co-ordinator from when it connected to the LowPan, hence there is distributed knowledge of unicast addresses that lead all the way to the PAN co-ordinator.
- The IPv6 router assigns prefixes through the prefix advertisement. The nodes are auto-configured with the advertised prefixes.
- The other FFDs in the network do not act as a IPv6 router, but they generally route data packet in the L2 layer (Mesh layer routing).
- Star topology assumes that each node is one hop away from the PAN co-ordinator.
- This document defines a mesh topology (see diagram below). In mesh topology, each node is capable of forwarding. Thus it can be assumed as a network of full functional devices (FFDs) with one PAN co-ordinator and multiple co-ordinators.
- FFDs may or may not be more than one hop away from the PAN co-ordinator.
- We assume that the 64-bit EUI-64 addresses are used as link-layer address in the lowpan, since these addresses never change for a given node. Appendix A discusses some additional considerations should we apply this to the 16-bit addresses.

This document currently supports the following topologies:

1) Star Topology

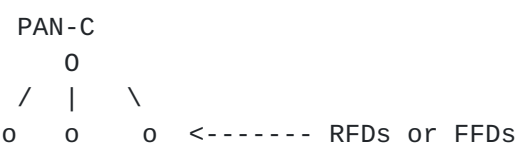

2) Mesh Topology

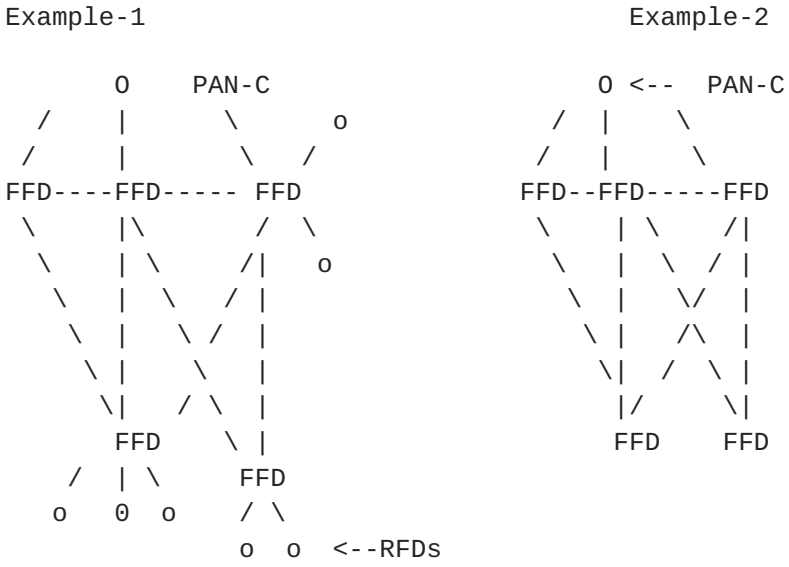

Each FFD may act as simple co-ordinators but there is only one PAN co-ordinator in the mesh LowPan topology.

6. Minimizing Multicast of Router Solicitations and Advertisements

We have potential issues with multicast of RS and RA during the initialization of a host, and also due to the periodic RAs.

6.1. Avoiding L2 broadcast of initial RS and RA

Since we assume that the PAN co-ordinator is also the IPv6 router we can completely avoid multicast RS and RA during the initialization of a lowpan node as follows.

When a lowpan node initializes and needs to send a Router Solicitation addressed to the IPv6 all-routers multicast address, it should only send the request to the node's co-ordinator. Thus the L2 destination address will be that of the co-ordinator. In a star

topology the co-ordinator is also the PAN co-ordinator hence the IPv6 router. Thus this will result in the RS being delivered to the router. In a mesh topology, when such a packet is received by a co-ordinator, it will look at the IPv6 header and if that is destined to the all-routers IPv6 multicast address, then it will relay that packet (without decrementing the IPv6 hop limit) to its co-ordinator. This will deliver the RS towards the PAN co-ordinator.

The lowpan nodes MUST include a Sender Link-layer Address option in the Router Solicitation, since this then allows the router to unicast a Router Advertisement in response.

Since every host sends a RS when it attaches to the PAN and there is a single router for the PAN (the PAN co-ordinator), this router will observe the arrival of all the IPv6 link local addresses. If the host uses auto-configuration, then the IPv6 router can form the host's IPv6 global address with the prefix it is advertising. The IPv6 router can also update the global address of the node, when it receives the first neighbor solicitation from the node. However, for the single PAN communication, link-local address will be sufficient. This allows the router to have a complete list of all the nodes on the link with their L2 addresses, which can be useful for other optimizations below.

The above behavior is also consistent with the approach in DNA, which relies on every host sending a RS when it attaches to a new link.

6.2. Avoiding L2 broadcast of periodic RAs

It isn't clear whether a periodic RA sent every 7 minutes to all-nodes will be a major problem. But, in mesh topology, an periodic Router Advertisement(RA) from the IPv6 router means flooding in the network. There are a few alternatives discussed in this document. Firstly, it might very well be possible to increase the default time significantly as long as the hosts can use some other mechanism to detect when the router (the PAN co-ordinator) disappears. We assume that there will be a backup router incase the primary router becomes unavailable. For experimental purpose, the maximum router advertisement interval in the lowpan network SHOULD NOT be less than 1500 sec. The default router advertisement lifetime then SHOULD be set to no less than 4500 seconds. The suggested default value for the default router advertisement lifetime is 7200 sec.

Secondly, instead of broadcasting such router advertisements at layer 2, since the router knows all the layer 2 addresses of the nodes in the PAN from the Router Solicitations it has received, it can instead do replicated unicast at layer 2. Thus sending a packet to the all-nodes IPv6 multicast address would result in sending one copy to each

unicast address. This requires more energy for the router, but might overall be a benefit for the nodes in the PAN.

Alternately, there could be a hybrid mechanism of the above two proposals. In this case, the PAN co-ordinator aka IPv6 router sends periodic RA to the co-ordinators in its PAN by sending unicast messages to each of them. The RA announcement interval SHOULD be longer than the standard default RA interval defined in the IPv6 Neighbor Discovery document. This is done because of power saving of the 6lowpan devices; note that even co-ordinators could be battery powered. Each co-ordinator can act as proxy IPv6 router advertiser and they can broadcast the RA on behalf of the IPv6-router in their own domain periodically (interval is TBD). In this method, the co-ordinators could use the beacon enabled L2 slots for router advertisements to avoid collisions and power wastage. But, in this method, a new mechanism needs to be developed for co-ordinator to co-ordinator message exchange. Thus, the first method is recommended as a default method by this document.

7. Minimizing Multicast of Neighbor Solicitations

There are two cases to consider when it comes to reducing or eliminating the multicast Neighbor Solicitation messages. One is finding the actual neighbors and their L2 addresses, whether initiated from nodes in the lowpan network or from the outside. Another case is the resource spent on looking for non-existent neighbors. For example a non-existent or unreachable IPv6 address may have same subnet prefix(es) assigned to the 6lowpan network. A neighbor solicitation for a non-existing IPv6 address might happen by accident, or as a DoS attack from outside the lowpan network.

7.1. Avoiding L2 broadcast of NS messages for existing nodes

Sending Neighbor Solicitation messages to resolve a layer 2 address for a IPv6 address is also a waste of bandwidth and energy in the LowPan network. Thus the following proposal attempts to distribute the link-layer information of nodes to the PAN co-ordinator which has higher possibility of being a full-functional device. Even if the co-ordinator lives a couple of L2 hops away from the enquiring node (assuming mesh routing at the L2 layer), each neighbor solicitation in this proposal will involve only a few nodes in the path of traversal of the packet. This is in effect, a unicast solicitation for resolving an address.

This part of the proposal does not require any protocol changes but instead relies on the existing support in Neighbor Discovery. The prefix advertising router would not set the "on-link" ("L") flag in

the prefixes it advertises, even though these prefixes are on-link. This would result in the hosts in the PAN initially sending packets to all through the router for on-link destination nodes (The hosts don't know the destination is on-link). This causes the router to both forward the packet back to the PAN and send a Redirect message back to the sender. The redirect can include the L2 address for the target, since it is likely to have that information in its neighbor cache.

The above works well when the router knows the L2 addresses for each IPv6 address (link-local and global) in the PAN. If the router doesn't know this (e.g., for a global address), the router would need to multicast Neighbor Solicitation before sending the redirect. More thoughts are needed to specify the protocol behavior when the router's cache entry becomes stale or when a lowpan node moves away from the PAN. What if the router itself looses power? Should there be multiple routers/co-ordinators who keep the neighbor information in distributed fashion - so that absence or failure of one router will not affect the neighbor solicitation negatively? Or will a PAN co-ordinator failure be noticed by all the lowpan nodes, so that they all can send a new RS to the router (which would inform the new PAN co-ordinator (the IPv6 router) of their L2 addresses.

7.2. Avoiding L2 broadcast of NS messages for non-existent nodes

If we apply the above configuration with no advertised on-link prefixes, the packet for an unknown or non-existent node would end up at the router for the lowpan. This is true whether the packet was originated by a node in the lowpan, or originated somewhere else in the Internet and forwarded to the lowpan router.

In the standard Neighbor Discovery behavior, this would result in the router multicasting a Neighbor Solicitation message. If the node doesn't exist or is unreachable from the router, then repeated packets would result in repeated multicast NS messages (limited to one per second for a single IPv6 address).

We can avoid this resource consumption if the router (the PAN co-ordinator) can keep an authoritative list of all the IPv6 addresses that are present in the lowpan - and their L2 addresses. Please note that it is not guaranteed that the IPv6-router can have complete list of "all" nodes present in the 6lowpan; it can only keep track of nodes from which it received the Router Solicitation. If we can have the router also maintain this for global IPv6 addresses, then we can avoid any NS messages for unknown destinations.

Then if a node sends a packet for a destination which does not exist in the LowPan network, the PAN co-ordinator will not be able to find

the node in its lookup table. The PAN co-ordinator can immediately send a ICMPv6 unknown destination error message to the originator in response to the packet. This will avoid any multicast Neighbor Solicitation messages.

Note that the above assumes that if there are multiple IPv6 routers attached to the PAN, they all can somehow efficiently find out the IPv6 address and L2 addresses that are in the PAN.

8. Minimizing Multicast for Duplicate Address Detection

Duplicate Address Detection (DAD) [ND] is sent to the solicited-node multicast address which is derived from lower 24 bits of the target IPv6 address. Should nodes in LowPan network use duplicate address detection Avoiding duplicate address detection will save broadcast signaling in the PAN-since 802.15.4 does not have multicast capabilities. Besides, each duplicate address detection message is capable of waking up sleeping reduced function devices in the network and making the devices less and less energy efficient.

In a star topology, it might be OK to broadcast the DAD message, but in a multi-hop LowPan network, it means flooding. In the following optimization case, the assumption is that the IPv6 Router keeps a table of IPv6 to EUI-64bit MAC address mapping of each IPv6 configured node in the PAN.

When a node (RFD or FFD) boots up or configures its LowPan interface, it can send the Duplicate Address Detection message the same way we suggested sending the Router Solicitations. In a multi-hop network the message is relayed to the PAN co-ordinator (IPv6 router) by the nodes co-ordinator. The IPv6 router can check in its complete list of neighbor whether the address is a duplicate, and respond appropriately. (If we want to allow Secure Neighbor Discovery [SEND], then it makes sense for the router to relay the message (without decrementing the hop limit) to the "owner" of the address, so that this node can answer with the appropriate SeND signature etc.)

Alternatively, if we assume that each IPv6 configured node has EUI-64bit MAC address and the nodes do not use IPv6 temporary addresses [RFC 3041] or CGA addresses [SEND-CGA], then all the IPv6 addresses would be directly derived from the EUI-64, hence will be unique. In that case IPv6 Duplicate Address Detection mechanism is not needed in the LowPan network.

9. Neighbor Unreachability Detection

If we used the EUI-64 MAC addresses in the lowpan and an IPv6 address is never reassigned to a different node in the lowpan, then there is no need to perform NUD towards a host. This is because the IPv6->MAC address mapping would never change.

In addition, if we had some other mechanism, e.g. a layer 2 detecttion of a PAN co-ordinator failure and recovery or replacement by a different PAN co-ordinator, then we do not need to do NUD towards the router.

[How does a node find out about unreachability (link-down, host-down) of a node in the link? Should this be restricted only when data-delivery happens? Should this be left to the L2-routing errors? Note, for star topology, NUD to router is not an issue, but perhaps redundant.] In IPv6, NUD is used to detect unreachability between host-host, router-router and host to router. In lowpan network, a node can always do NUD for its neighbor using its unicast IPv6 address. It is also possible that when a node becomes unavailable, the respective co-ordiantor node can record that and propagate the unreachability through the network to the IPv6 router. In this case, the IPv6 router should implement an L2-to-L3 interface to handle the L2 message and take the node out of its cache information. The details L2 mechansim is out of scope of this document.

10. Open Issues

The following are the known open issues:

- For LowPan nodes, should we say that all NS messages must contain the source link-layer address option to avoid triggering a followup neighbor solicitation in reverse direction?
- What is the best way to keep the neighbor cache information distributed among different routers/co-ordinators for fault-tolerance? Should we develop a inter co-ordinator message exchange protocol for distributed information, proxy and caching?
- Should a node de-register itself from router/co-ordinator's neighbor cache if it decides to move away?
- What is the default lifetime and interval values for routers and prefixes? (assigned 1800sec and 900 sec for experimental purpose)
- We have a mechanism for the PAN co-ordinator to find out all the IPv6 link-local addresses in the lowpan. We need to extend this to the global addresses that are used as well.

11. Fault Tolerant IPv6-routers

Neighbor Discovery Registration Extension [ND-REG] defines a new ND registration option and a registration message for nodes to register with one or more multiple IPv6 routers. Using this mechanism, the primary IPv6-router may optionally include the ND registration option in the RA message and includes the backup router's IP-address. Thus the node can then register with the both primary and secondary rotuers. In case, the primary router goes down the secondary router can take over. It is assumed that both primary and secondary routers are on the same IPv6-link and advertise the same prefixes.

The ND registration [ND-REG] mechanism also supports a 'N' bit with RA which can be optionally used by the IPv6-router when it wants the node to register immediately, say after it becomes operational again. In order to avoid flood of registration messages, the IPv6-router SHOULD introduce delay between sending consecutive RAs with 'N' bit.

12. Sequence of Operantions between 6lowpan hosts and IPv6-router

This section discusses the sequence of operations of IP address assignment and router-discovery process in the 6lowpan network. It assumes that the L2-address assignment has already taken place and the 6lowpan network is capable of routing data and control IP packets at the link-layer.

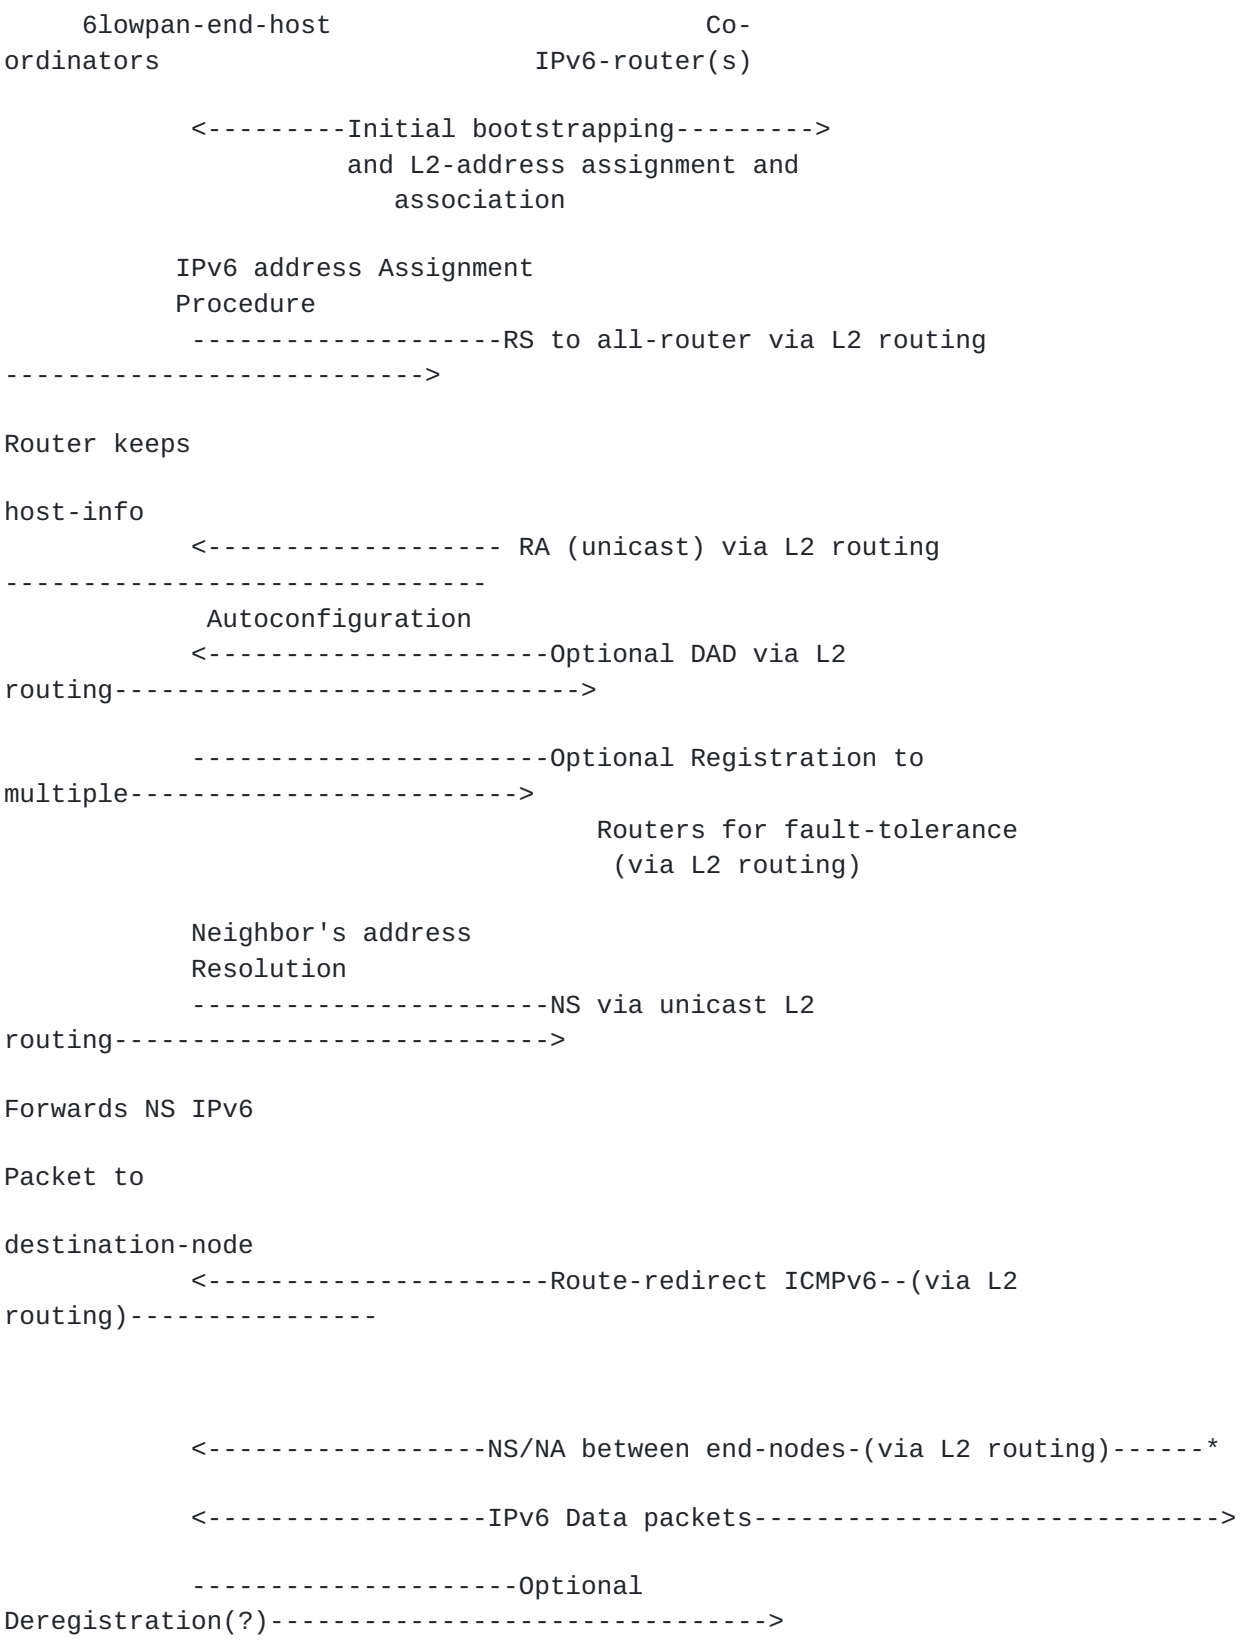
6lowpan-end-host
Co-ordinators
IPv6-router(s)
<---------Initial bootstrapping--------->
and L2-address assignment and association
IPv6 address Assignment Procedure
--------------------RS to all-router via L2 routing
--------------------------->
Router keeps
host-info
<------------------- RA (unicast) via L2 routing

Autoconfiguration
<----------------------Optional DAD via L2 routing------------------------------>
----------------------Optional Registration to multiple------------------------->
Routers for fault-tolerance (via L2 routing)
Neighbor's address Resolution
----------------------NS via unicast L2 routing---------------------------->
Forwards NS IPv6
Packet to
destination-node
<----------------------Route-redirect ICMPv6--(via L2 routing)----------------
<------------------NS/NA between end-nodes-(via L2 routing)------*
<------------------IPv6 Data packets------------------------------>
---------------------Optional Deregistration(?)------------------------------->

13. Security Considerations

These optimizations are not known to introduce any new threats against Neighbor Discovery beyond what is already document in [ND-Threat].. However, the effect of a rogue router is more severe in Low-power wireless network than in the network of powered systems. The 6lowpan security analysis [6lowpan-threat] discusses possible threats. In 6lowpan network it is important the IPv6 router advertisement is authenticated. Since IPv6 router is the central

point of neighbor solicitation, the solicited response from the IPv6 router should be authenticated. SEND [SEND] recommends CGA option, RSA option and time-stamp and nonce options for various neighbor discovery messages. Certificates works best for a fixed network router; it is not sure if certificate authentication is appropriate choice in this situation. RSA keys are large, thus router advertisement with RSA option may not be a good idea for lowpan network. RSA may be replaced with Elyptic Curve Cryptographic (ECC) signature because ECC is known to work for small low-power devices. However, the sender of neighbor solicitation messages should include a nonce option for a matching response for replay protection and weak authentication. However, in most cases, the layer 2 authentication, if applied, may be sufficient for many deployments. Those situations where external threats are severe, the subset of SEND [SEND] should be used between the IPv6 router and the end nodes as a layer 3 security option.

14. Applicability of Neighbor Discovery Optimization in non-multicast networks

IPv6 Neighbor Discovery optimization described in this document can be applicable to other new technolgies which like to offer IPv6 functionality, but does not support multicast or broadcast. It is similarly applicable to the other wireless technolgies where minimum signaling is desired and unicast communication is preferable. The optimization solution, presented in this document, assumes that all nodes are connected to an onlink router. In most cases, the technolgies where multicast or broadcast might be prohibitively expensive, require one-to-one connection to a central node.

Thus, at the network layer, following Neighbor Discovery optimizations are suggested:

o Avoiding multicast/broadcast of RS by the node at the startup

It is assumed that the node knows the router's IPv6 address on the link. The node sends a Router Solicitation message directly to the router's link-local address. This RS message MUST carry sender's link-layer address (SLLA) option. The router, upon receipt of the unicast RS messages, updates its address-cache table with requestor's IP-address and link-layer address.

o Avoiding multicast periodic Router Advertisements by the router

Assume that the router does not send periodic unsolicited advertisements to the multicast or broadcast address. But since it keeps a table of solicitor's list (as mentioned above), the router

periodically sends RA to the solicitor's unicast IPv6 address. The default periodic interval should be set by the link-layer technology specific requirements.

o Minimizing/avoiding multicast or broadcast for Neighbor Solicitation

IPv6 neighbor discovery suggests using solicited node multicast destination address for neighbor solicitation for existing and non-existing nodes. Network technologies that do not support multicast but only broadcast, it is an expensive and energy-consuming operation. On the other hand, other L2 technologies like 802.16 which does not support multicast, require an alternative solution. In both cases the proposed optimization is useful. We divide the solution in two cases 1) existing nodes 2) non-existing nodes.

The first part of the proposal does not require any protocol changes but instead relies on the existing support in Neighbor Discovery. The prefix advertising router would not set the "on-link" ("L") flag in the prefixes it advertises, even though these prefixes are on-link. This would result in the hosts in the link initially sending packets to all through the router for on-link destination nodes (The hosts don't know the destination is on-link). This causes the router to both forward the packet to the target and send a Redirect message back to the sender. The redirect can include the L2 address for the target, since it is likely to have that information in its neighbor cache. This solution assumes that the router has knowledge about all the link-local nodes IP-address and corresponding L2-address.

For non-existing nodes, the router will receive the solicitation message, as described above. But in this case, no entry will be found in router's address cache corresponding to the target IPv6 address. The router will discard the solicitation message and responds to the sender with ICMPv6 destination unreachable error message.

Note that the optimization assumes that the router reliably can keep information about every node on the link as they appear and disappear, across router booting and crashing.

o Address Resolution using Neighbor solicitation messages

Our assumption is that all IPv6 packets are going through the IPv6 router due to the off-link flag in the IPv6 router advertisements. Thus a node A initially always sends data packet to the IPv6 router for another node B, which might be on the same link. The router sends router redirects for the target on-link IP address to the sender node A. The router redirect carries the target link-layer

address option. Once the originator receives the redirect message, it can update the neighbor cache entry for target IPv6 address. The sender node then communicates with node B directly. Note that this procedure is already defined in RFC2461.

o Avoiding Duplicate Address Detection (DAD)

Duplicate address detection may be avoided if no privacy or temporary address is used in the onlink nodes, and the uniqueness of the IPv6 address is assured. In an IPv6 network, DAD is performed by sending a solicitated node multicast address derived from the sender's IPv6 address. In a non-multicast network, the DAD message should be sent to the default router. The router is responsible for responding to the DAD messages. But this method of DAD optimization only works when the router reliably keeps track of all the IPv6 nodes on the link at any point of time. However, this document does not specify how to store nighbor cache table in the IPv6 router during a system failure.

o Neighbor Unreachability Detection (NUD) optimization

NUD messages are sent to host-host, host-router and router-router for reachability detection. NUD is alwyas unicast solicitation message. If all messages are supposed to go through the on-link router, then it is assumed that the router sends periodic NS messages to each node to keep its neighbor cache up-to-date. The router must be informed when a node leaves the network. It is not clear at this point, whether a Layer 3 solution is needed for detecting absence of node in the neighbor cache, since the router may receive Layer 2 indication for the loss of a node in the link. If uniqueness of the link-layer address of a node is guaranteed, and the IPv6 router can reliably keep track of all the nodes in its neighbor cache at all times, then it is possible to avoid NUD messages from the router to the nodes.

o Node Deregistration

This model of optimization suggests that the neighbor cache information must be deleted at the on-link default IPv6 router(s) when a node moves away from the IPv6 link. It is not clear at this point if a L3 solution is needed for loss of node detection. Informing all the routers in the link about departure of a node from the network, may be a function of L2 indication.

15. IANA Considerations

There are no IANA considerations for this document.

16. Acknowledgements

Many thanks to Dave Thaler and Jim bound for their input to the document, and to Eunsook Eunah Kim for her comments on the document.

17. References

17.1. Normative References

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

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

[LOWPAN] Montenegro, G. and N. Kushalnagar, "Transmission of IPv6 Packets over IEEE 802.15.4 networks", RFC 4944, September 2007.

[LOWPANG] Kushalnagar, N. and G. Montenegro, "6LoWPAN: Overview, Assumptions, Problem Statement and Goals", RFC 4919, August 2007.

17.2. Informative References

[IPV6] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6), Specification", RFC 2460, December 1998.

[STATELESS] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.

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

[ND-Threat] Nikander, P., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery(ND) Trust Models and Threats", RFC 3756, May 2004.

[ND-REG] Nordmark, E., "IPv6 Neighbor Discovery(ND) Registration Extension", draft-nordmark-6lowpan-reg-00.txt (work in progress), June 2008.

[AUTOCONF]

Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.

[SEND-CGA] Aura, T., "Cryptographically Generated Addresses", RFC 3972, March 2005.

[IEEE] IEEE Computer Society, "IEEE Std. 802.15.4-2003", , October 2003.

[16NG] Jee, J., "IP over 802.16 Problem Statements and Goals", draft-jee-16ng-ps-goals-00.txt (work in progress), February 2006.

[6lowpan-threat] Park, D., Kim, K., Seo, E., and S. Chakrabarti, "IPv6 over Low Power WPAN Security Analysis", draft-daniel-6lowpan-security-analysis-02.txt (work in progress), January 2007.

Appendix A. Supporting short addresses?

[Discuss issues with supporting 16bit addresses with lowpan-ND.]

We would at least need NUD to handle 16 bit MAC addresses, since they can change. Depending on how frequently they change, supporting 16 bit MAC addresses might make some other proposed optimization techniques (in this document) more difficult or impossible. In other situation of unreachability, the sender node (if it's a router) sends a multicast neighbor soliciation following a NUD failure or if the sender is a host it deletes the unreachable neighbor's entry and then sends the solicitation packet to the IPv6-router(PAN co-ordinator).

Also, when the L2 address change, usually the affected node sends unsolicited Neighbor Advertisements to the neighbors to note the L2 address change in their neighbor cache. Since 16bit short addresses are not unique, there is higher probability that these addresses will change and thus we need to optimize unsolicited neighbor advertisments in the LowPan network. Perhaps the IPv6-router or the PAN co-ordinator can be used as the central contact point. Hence the changed node may notify the PAN co-ordinator about the change. But then either the PAN co-ordinator will have to notify all nodes about that change or other nodes need to find out the change through neighbor solicitation following a NUD. It seems, IEEE EUI-64 addresses are simple to use in case of IPv6.

Appendix B. Summary of proposed optimizations

Scenarios	Optimization
Initial RS to all-routers	Initial RS to PAN co-ord address (unicast) Must use SLLA option for L2 addr
Initial RA to all-nodes	Avoid
Periodic RA to all-nodes	See 5.2 for proposals
Neighbor Solicitation	Send unicast message to the IPv6-router
Unsolicited Neighbor Adv	Avoid by using IEEE EUI-64 MAC
Neighbor Unreachability Detection	Avoided if IEEE EUI-64 MAC addr used
Duplicate address Detection	Same as above
ND Constants	Suggested default maxRAadvtime=1500 sec default RouterAdvlife=7200 sec

Authors' Addresses

Samita Chakrabarti
IP Infusion
1188 Arquest Street
Sunnyvale, CA,USA.

Email: samitac@ipinfusion.com

Erik Nordmark
Sun Microsystems, Inc.
17 Network Circle
Menlo Park, CA 94025
USA

Email: Erik.Nordmark@Sun.COM

Full Copyright Statement

Intellectual Property