

Networking Working Group	A. Tavakoli	
Internet-Draft	S. Dawson-Haggerty	
Intended status: Informational	UC Berkeley	
Expires: September 11, 2009	J. Hui	
	Arch Rock Corporation	
	D. Culler	
	UC Berkeley	
	March 10, 2009	

[TOC](#)

HYDRO: A Hybrid Routing Protocol for Lossy and Low Power Networks draft-tavakoli-hydro-01

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on September 11, 2009.

Copyright Notice

Copyright (c) 2009 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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

HYDRO is a hybrid routing protocol for Lossy and Low power Networks (L2Ns) that embraces centralized and distributed routing mechanisms. Through the use of standard ICMP Route Advertisements and Route Solicitations, Node Routers build Default Routes to Border Routers. These routes, which maintain multiple options per each Node Router when available, are maintained through data-driven link estimation. Node Routers periodically report a high-quality subset of their Default Route Table to Border Routers, which then form a global view of the topology. When a Node Router attempts to route to another Node Router in the network, if no matching entry exists in the Node Router's Flow Table, it forwards the packet to a Border Router, which then installs the correct Flow Table Entries in the network to enable more efficient subsequent routing.

Table of Contents

- [1. Terminology](#)
- [2. Introduction](#)
- [3. Overview](#)
- [4. HYDRO Terminology](#)
- [5. Data Structures](#)
 - [5.1. Default Route Table](#)
 - [5.2. Flow Table](#)
 - [5.3. Link Database](#)
- [6. Message Formats](#)
 - [6.1. Additional Adaptations](#)
 - [6.2. Routing Header](#)
 - [6.3. Route Install Header](#)
 - [6.4. Router Advertisement Extension](#)
 - [6.5. Topology Report Message](#)
 - [6.6. Topology Report Package](#)
- [7. Detailed Operation](#)
 - [7.1. Link Estimation](#)
 - [7.2. Router Discovery](#)
 - [7.3. Default Route Formation](#)
 - [7.4. Global Topology Construction](#)
 - [7.5. Node Forwarding](#)
 - [7.6. Border Router Forwarding](#)
 - [7.7. Route Installation](#)
 - [7.8. Multicast Forwarding](#)
 - [7.9. Multiple Border Routers](#)
- [8. Extensions](#)
- [9. Configuration Parameters and Other Administrative Options](#)
- [10. Applicability Statement](#)
- [11. Security Considerations](#)

12.	IANA Considerations
13.	Acknowledgements
14.	References
14.1.	Normative References
14.2.	Informative References
§	Authors' Addresses

1. Terminology

[TOC](#)

LLN: Lossy and Low power Network

MTU: Maximum Transmission Unit

ROLL: Routing in Low power and Lossy Networks

TDMA: Time Division Multiple Access

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 [RFC 2119 \(Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.\)](#) [RFC2119].

2. Introduction

[TOC](#)

The Hybrid Routing Protocol (HYDRO) for Lossy and Low Power Networks utilizes the two-tiered hierarchy that is used in many deployments. It provides point-to-point routing using a combination of centralized and distributed mechanisms.

L2N networks are often deployed in two tiers: Node Routers (or Nodes from now on) in the network which gather data but are resource constrained, and Border Routers, which provide connectivity between the L2N subnet and another network. The Nodes have constrained resources, in the form of limited energy, memory, and bandwidth. Consequently, a routing protocol MUST ascribe to a set of minimal criteria, as laid out in [the IETF Roll Protocols Survey \(Tavakoli, A., Dawson-Haggerty, S., and P. Levis, "Overview of Existing Routing Protocols for Low Power and Lossy Networks," April 2009.\)](#) [I-D.ietf-roll-protocols-survey]. These requirements were culled from a set of application requirement documents, with the targeted domains including [Industrial Monitoring \(Networks, D., Thubert, P., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low Power and Lossy Networks," June 2009.\)](#) [I-D.ietf-roll-indus-routing-reqs], [Building Automation \(Martocci, J., Riou, N., Mil, P., and W. Vermeylen, "Building Automation Routing Requirements in Low Power and Lossy Networks," January 2010.\)](#) [I-D.ietf-roll-building-routing-reqs], [Urban Sensing \(Dohler, M., Watteyne, T., Winter, T., Barthel, D., Jacquenet, C., Madhusudan, G.,](#)

[and G. Chegaray, "Urban WSNs Routing Requirements in Low Power and Lossy Networks," March 2009.](#)) [I-D.ietf-roll-urban-routing-reqs], and [Home Automation \(Brandt, A. and J. Buron, "Home Automation Routing Requirements in Low Power and Lossy Networks," January 2010.\)](#) [I-D.ietf-roll-home-routing-reqs], with the requirements intended to provide a necessary, but not necessarily sufficient baseline. They apply to low-power Node Routers, but not to the Border routers. The Border Routers, however, are assumed to have a relatively abundant supply of resources, leading to HYDRO's goal of maintaining as much state and complexity at the Border Routers as possible, and only pushing necessary functionality to the Nodes.

3. Overview

[TOC](#)

HYDRO provides point-to-point connectivity between all Nodes in a network, but is optimized for situations in which the number of network destinations in the network is far smaller than the number of Nodes. Additionally, much of the traffic is typically between Node routers in the network and hosts on other networks. Thus, our design optimizes first for this traffic pattern. It also not uncommon for there to be some unicast traffic between different Node Routers; for instance, for control or debugging, and so we provide a second mechanism for optimizing this traffic.

An underlying set of routing trees are created using IPv6 Route Advertisement and Route Solicitation Messages. Each Border Router advertises a route with '0' cost. Nodes process Router Advertisements from both Border Routers and Nodes, and add their Link Cost Estimate to the advertised Route Cost Estimate, and then themselves advertise this new cumulative Overall Route Cost.

HYDRO allows for multiple Border Routers to manage a network. We detail the mechanisms for setting up and managing such a topology in later sections, but for ease of understanding frame our discussion in the context of networks with only a single Border Router. Such a simplification does not alter the functionality of Nodes.

Each Node maintains information about multiple potential next-hops to the Border Router, which form a set of Default Routes. The ordering of these choices, is based on a combination of the Overall Route Cost, the Confidence in the Link Cost Estimate (which is a component of that Overall Route Cost), and a Node's Willingness To Route (Willingness). Confidence MAY be measured by the number of packets over which the Link Cost Estimate is formed, both of which are provided by the link layer. Each Node periodically reports a high-quality subset of its Default Route Table, as well as any requested Node Attributes, whether static or dynamic, in the form of Topology Reports to a Border Router. Using this information, the Border Routers create a global view of the topology of the subnet.

When a Node needs to route a packet, either as its originator or a forwarder, it attempts to classify a packet against its Flow Table (which is initially empty). If an entry matches the packet, then the packet is forwarded according to the rules specified by the entry. Otherwise, the Node forwards the packet along a Default Route to the Border Router.

When a data packet destined for a Node Router, with a destination address within the subnet reaches a Border Router, the Border Router determines the best route, according to its global topology view and a user-specified routing policy (e.g. lowest cost). It then forwards the packet to the destination by inserting a source routing header into the packet and forwarding it to the first hop in the source route. It SHOULD additionally generate a Route Install Header for either the source or destination of the packet, instructing it to install the necessary Flow Table Entries to route future similar packets more efficiently. This header may be set in a new packet or added to the original data packet.

HYDRO mechanisms require the frequent transmission (or exchange) of IPv6 addresses. Because full 128-bit addresses in this setting are not practical, we assume that all Nodes are associated with a locally unique 16-bit short identifier, and furthermore that all Border Routers have access to the mapping between 16-bit and 64-bit interface identifiers. This 16-bit identifier should be equivalent to the 16-bit form of compressed address specified by [\[RFC4944\] \(Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks," September 2007.\)](#). The mapping may be static or dynamic, using mechanisms like DHCPv6 but is beyond the scope of this document. All routing takes place using 16-bit short identifiers

4. HYDRO Terminology

[TOC](#)

Border Router: A device with relatively abundant resources and multiple interfaces, including an interface (the 'L2N Interface') that shares a link type with Nodes.

L2N Network: A network with a single IPv6 prefix encompassing both a set of Border Routers and Node. Multiple types of links connect these routers, and the Border Routers are all in a single link-local multicast domain. All interfaces on this network share a single prefix (the L2N prefix).

L2N Interface: The interface which the Border Router uses to communicate with devices on the subnet.

Secondary Interface: A second interface the Border Router uses to communicate with other Border Routers. Examples include an Ethernet or 802.11 mesh.

Default Route: A route leading to a Border Router from a Node, composed of a locally maintained next-hop link.

Default Route Table: A Node's table which contains information about a set of Default Routes and their associated next-hop Nodes, including the cost of communicating with a Border Router using that particular Default Route.

Flow Table: A table with a (Classification, Action) format, in which 'Classification' helps categorize which packets an entry applies to, and 'Action' describes how to forward a matching packet.

Neighbor: A Node or Border Router within a single IP-hop.

Node Attributes: Information about a Node, either static (e.g. bandwidth, resources), or dynamic (e.g. energy left).

Node Router (Node): Any node in the network that serves as a router, but not as a Border Router. Typically a constrained device.

Overall Route Cost: The cumulative cost of a particular Default Route option, obtained by summing the Route Cost Estimate advertised by the associated next-hop Node, and the Link Cost Estimate of communicating with that Node.

Primary Default Route: The Default Route that a Node attempts to use before all other potential Default Routes. Typically the top entry in the Default Route Table, but we detail mechanisms that allow other entries in the Default Route Table to also periodically serve as the Primary Default Route.

Route Install Message: A message containing the bidirectional path between two Nodes. The message also dictates installation type.

Topology Report: Includes a subset of information about a subset of Default Route next-hops from the Default Route Table, as well as requested Node Attributes.

Willingness To Route (Willingness): An advertised value that indicates a Node's willingness to route packets for other Nodes. Setting the value of this metric is policy-driven and outside the scope of this document.

5. Data Structures

[TOC](#)

5.1. Default Route Table

[TOC](#)

Each Node maintains a Default Route Table, whose entries contain the following fields:

Neighbor Address: 16-bit short IPv6 address of the neighboring Node or Border Router.

Route Hops: The distance to a Border Router in hops using the specified Neighbor as the next-hop.

Route Cost Estimate: Route cost advertised by the next-hop Node for this entry.

Willingness: The Willingness to Route advertised by the next-hop Node for this entry.

There are also certain fields associated with a Default Route Entry that MUST be provided by the Link Layer, but are used by HYDRO for Default Route Table seeding and maintenance. These fields are:

Link Cost Estimate: Estimate of the link cost to communicate with the next-hop Node for this entry. The link layer MAY need multiple packet transmissions and acknowledgements to provide this estimate.

Link Quality: Last recorded link quality measurement. The link layer MUST be able to provide this value based on a single packet reception. The metric and scale of this field is likely hardware-dependent.

Confidence: The Confidence of the link layer in the Link Cost Estimate it provides.

The size of the Default Route Table is dictated by NUM_DEFAULT_ENTRIES.

5.2. Flow Table

[TOC](#)

Each Node maintains a Flow Table, whose entries are composed of two parts:

Flow Match: The Flow Match contains information against which packets can be checked. Flow Match MUST include a Packet Destination field, and MAY contain certain additional fields, such as Packet Source.

Flow Path: The Flow Path indicates the path a packet should take. This can either be a full source route, or simply the next-hop to which the packet should be sent to.

5.3. Link Database

[TOC](#)

Each Border Router maintains a Link Database, whose entries have several parts:

Link <n1,n2>: a link between Node 'n1' and 'n2', reported in a Topology Report by 'n1'.

Sequence Number: The sequence number contained in the Topology Report which last reported this edge.

Metric: The Metric associated with this edge in the last Topology Report

Confidence: The Confidence associated with this edge in the last Topology Report

Attributes: Any Node Attributes contained in the last Topology Report, such as Willingness to route

6.1. Additional Adaptations

The 6lowpan working group has defined an adaptation layer for IPv6 datagrams, when carried over Low Power and Lossy links. This is necessary due to the small link MTU. We suggest an additional function of the adaptation layer involving extension headers, which is necessary to use these header efficiently. The generalized IPv6 extension header format is defined in [\[RFC2460\] \(Deering, S. and R. Hinden, "Internet Protocol, Version 6 \(IPv6\) Specification," December 1998.\)](#), and is reproduced below.

[illegible]

The "Hdr Ext Len" is specified in units of 8-octet chunks, not including the first 8 octets. A similar layout is used for TLV-encoded options, except the 'Next Header' field is named the 'Option Type' field. We suggest that the 6lowpan working group may wish to define, as part of the adaptation layer, lengths in units of single octets; we will assume for the rest of this document that such a modification is in effect. Although this limits the length of extension headers and options to 255 bytes, we feel that this is not a significant limitation and the payload savings resulting from removing unnecessary padding are significant.

This recommendation applies specifically to Hop-by-hop, Destination, and Routing extension headers (with 'next header' values 0, 60, and 43, respectively). It also applies to TLV-encoded sub-headers as specified in [\[RFC2460\] \(Deering, S. and R. Hinden, "Internet Protocol, Version 6 \(IPv6\) Specification," December 1998.\)](#).

6.2. Routing Header

IPv6 Routing extension headers are defined in [\[RFC2460\] \(Deering, S. and R. Hinden, "Internet Protocol, Version 6 \(IPv6\) Specification," December 1998.\)](#). However, the loose source routing they define is not appropriate for this setting since (a) it uses 128-bit addresses, and (b) it is deprecated by [\[RFC5095\] \(Abley, J., Savola, P., and G. Neville-Neil, "Deprecation of Type 0 Routing Headers in IPv6," December 2007.\)](#). Source routing is a critical facility in this class of network, and so we define a new routing type. Alternatively, we may consider this a routing type 0 header, with adaptations.

```
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Next Header | Hdr Ext Len | Routing Type=0 | Segments Left |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                Address[1]                |                Address[2]                |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
.                                           .
.                                           .
.                                           .
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                Address[n-1]                |                Address[n]                |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

The addresses in this header are stored in compressed 16-bit form, as specified by [\[RFC4944\] \(Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks," September 2007.\)](#) (or later).

6.3. Route Install Header

[TOC](#)

A Route Install Header may be placed either as a Destination option or a Hop-by-hop option, depending on the method of install. Processing of this header will depend on which of these is in effect.

```

+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Option Type | Option Len | M Len | |R| M | Path Len |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           Flow Match (D)           |           Address[1]           |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           Address[2]           |           Address[3]           |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
.                                     .                                     .
.                                     .                                     .
.                                     .                                     .
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           Address[Path Len - 1]           |           Address[Path Len]           |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Option Type: the TLV dispatch value TBD.

M Len : 'match length'. The length of the flow match in octets. MUST be at least 2, but MAY be extended in future drafts.

R : 'reverse'. if set, indicates that the reverse path should also be installed

M : 'method'. 00 is HOP_BY_HOP and 01 is FULL_PATH. Other values are reserved

Path Len: the number of addresses contained in the list. If zero, indicates that an IPv6 routing header should be consulted for the path to D.

Flow Match (D): the destination address of this flow.

Address[...]: a list of intermediate Nodes on a path from a source address to D. The source must be the destination of the IPv6 datagram the header is carried in; that source address must not be included in the path, and Address[n] = D.

6.4. Router Advertisement Extension

[TOC](#)

In order to carry the Total Path Cost in a router advertisement, we define a new extension header to the Router Advertisement message defined in [\[RFC2461\] \(Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 \(IPv6\)," December 1998.\)](#).

```

+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Option Type | Option Len |           Metric           |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Willingness |
+-+--+--+--+--+--+--+

```

Option type: the TLV value for this header TBD.

Option length: The option length, in octets (4).

Metric: the Overall Route Cost from a Node to the original advertising router, the Border Router. If this extension is not present, a cost of zero should be assumed.

Willingness: a value indicating the Node's willingness or capacity to route traffic for other Nodes.

6.5. Topology Report Message

[TOC](#)

Topology Reports are carried as an IPv6 hop-by-hop option inserted in outgoing packets.

```
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Option Type | Option Len | AL | Sequence Number          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     (Attributes)                                     |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Metric[1]   | Confidence[1] | Address[1]          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
.                                                     .
.                                                     .
.                                                     .
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Metric[N]   | Confidence[N] | Address[N]          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

Sequence Number: a monotonically increasing identifier of this topology report

AL: "Attribute Length": the length of the "Attributes" field in octets. If the length is zero, no Node Attributes are carried. If the length is one, a single octet follows the Sequence Number field, which is the "Willingness" value. Other values of AL are reserved.

Attributes: Node Attributes which may be used to make routing decisions; defined outside of the HYDRO protocol

Metric[i] : a Link Cost Estimate

Confidence[i] : the link estimator's Confidence in Metric[i]

Address[i] : a compressed IPv6 address indicating the Node or Border Router Metric[i] and Confidence[i] refer to.

There is no explicit field specifying the number of entries; implementations should infer this value from the Option Length field and the Attribute Length field.

[TOC](#)

6.6. Topology Report Package

A Topology Report Package is used for transmitting a collection of Topology Report Messages between Border Routers. This message is formed by suffixing the Topology Report with the 2-octet address of its originator. Multiple Topology Report Messages may be concatenated into a single Topology Report Package.

7. Detailed Operation

[TOC](#)

This section describes in greater detail the HYDRO protocol. We begin by describing the interaction of HYDRO with the link layer, and the Router Advertisement mechanism. We then explain how these primitives are used to form Default Routes, report topology, and route back into the network. Finally, we detail how Border Routers may insert Flow Table Entries into the network to improve inefficient routes.

7.1. Link Estimation

[TOC](#)

A critical design element for good performance in L2N networks is good link estimation; low power wireless links exhibit significant temporal and spatial variations in quality due to a number of physical and environmental phenomenon. This poses a challenge to traditional IP link abstraction since any protocol must account for this behavior.

The link layer MUST provide an additive link metric for potential links to the routing layer, such that the cost of a path $C[A \rightarrow B \rightarrow C] = C[A \rightarrow B] + C[B \rightarrow C]$. Examples of such a metric are Expected Transmissions (ETX) or Expected Transmission Time (ETT). Additionally, the link layer MUST provide an estimate of its Confidence in each Link Cost Estimate. These estimates MAY be binary or integer valued.

The link layer must also provide a mechanism for 'local broadcast'; the 802.15.4 broadcast address (0xffff) is such a mechanism. It MAY be advantageous to replace this mechanism with a dedicated discovery component in the future.

7.2. Router Discovery

[TOC](#)

Router Advertisement and Router Solicitation messages and the associated protocol have been standardized for IPv6 in [\[RFC2461\]](#) (Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)," December 1998.). While many of the mechanisms can be

applied directly in L2Ns, some require modification for efficiency or performance. The 6lowpan working group has created a team to investigate these necessary modifications; once their work is complete HYDRO SHOULD be modified to use those mechanisms.

We do not alter the Router Solicitation or Router Advertisement packet formats. We do define a new Router Advertisement extension, the Multihop extension, which allows Nodes to advertise their Overall Route Cost value. Router Advertisements without this extension are assumed to have a Total Route Cost of zero. In addition, another extension allows Router Advertisements to specify a Node's Willingness value.

When an event triggers the generation of Router Solicitation or Router Advertisement messages, a binary exponential timer is started. Each time the timer fires, a new Router Solicitation or Advertisement is generated and sent to the IPv6 all-routers link-local multicast address (ff02::2), using the link broadcast facility.

There are two conditions which cause a reset of the Router Solicitation timer, and thus the transmission of new Router Solicitations:

1. Boot
 2. No valid Default Route is present in the Default Route Table.
- There are two conditions which cause the Router Advertisement timer to be started, if it is not running, and thus the transmission of Router Advertisements:

1. Reception of a Router Solicitation, and the Primary Default Route is valid.
2. The Total Route Cost of the Primary Default Route changes by more than a threshold, ROUTE_COST_NOTIF_DIFF.

7.3. Default Route Formation

[TOC](#)

The Border Router is manually configured with an IPv6 address, including a full 64-bit network prefix on its L2N Interface; all Border Routers on the same L2N Subnet have L2N Interfaces with addresses on the same prefix. A Node automatically transmits Router Solicitations on boot, and nearby Nodes with a valid Primary Default Route will reply with Router Advertisements. This section specifies processing rules for these Router Advertisements. The result is an ordered table of potential next-hops towards the Border Router.

A Node that receives a Router Advertisement from a Potential Default Route Neighbor (PNeighbor) first checks if a Default Route entry for PNeighbor exists in its Default Route Table. If an entry for PNeighbor does not exist, it begins processing the Route Advertisement to determine if an entry for PNeighbor should be inserted.

The Default Route Table is sorted based on a combination of Overall Route Cost, which is the sum value of the Route Cost Estimate advertised by the next-hop Neighbor and the Link Cost Estimate, the Confidence, and the Willingness advertised by an entry's next-hop Node.

The Confidence and Link Cost Estimate MUST be provided by the link layer.

Processing begins by checking whether the link layer Link Quality Metric exceeds the necessary threshold (LINK_ADMIT_THRESH) to be considered. The exact metric is implementation specific, and MAY NOT be the Link Cost Estimate ultimately used for Default Route maintenance. For instance, implementations may use a link measurement like RSSI or LQI, along with a static or dynamic threshold to reject poor links. The rationale is that estimators like ETX require multiple packets before they can be trusted, whereas the routing layer must evaluate a flood of Router Advertisements for possible use as Default Routes.

A Node next checks if an empty spot exists in the Default Route table. If a spot exists, then HYDRO determines where in the table to insert an entry for PNeighbor. The algorithm begins by creating a new Default Route entry at the bottom of the Default Route table. It then repeatedly compares the new Default Route entry with the entry in the spot above it. If the higher spot is empty, or if the entry also has a '0' Confidence value and its advertised Route Cost Estimate is greater than the PNeighbor's advertised Route Cost Estimate, the two entries' positions are swapped. The iteration stops when this condition fails, or the algorithm reaches the top of the Default Route Table.

If the Default Route Table is full, then HYDRO determines whether the bottom Default Route Table entry should be evicted to create room for PNeighbor. If the bottom entry is not Mature, meaning that its Confidence value is less than CONF_EVICT_THRESHOLD, or its Route Hops value is less than PNeighbor's Route Hops value, we leave the Default Route Table as is and discard the Route Advertisement.

Otherwise, if PNeighbor's Route Cost Estimate is lower than the existing Default Route Entry's Route Cost Estimate by at least PATH_COST_DIFF_THRESH, or if the absolute value of the difference between the two Route Cost Estimates is less than PATH_COST_DIFF_THRESH and PNeighbor's Link Quality Estimate is greater than the existing Default Route Entry's Link Quality Estimate by at least LINK_QUALITY_DIFF_THRESH, then the existing bottom Default Route Table Entry is evicted to make room for PNeighbor's entry.

If an entry slot is found for PNeighbor, the field values are set according to the Router Advertisement. If an entry previously existed for PNeighbor, then the field values are simply updated.

If the Route Cost in the Route Advertisement is set to MAX_ROUTE_COST and PNeighbor was already present in the Default Route Table, the existing PNeighbor entry is evicted.

The top entry in the Default Route Table is the Primary Default Route. We define the Overall Node Route Cost at any given time to be the Overall Route Cost of the Primary Default Route, and Overall Route Hops to be the Route Hops Value of the Primary Default Route. If no Primary Default Route exists, these are set to MAX_ROUTE_COST and MAX_HOPS, respectively.

At the end of each period, as determined by PERIOD_LENGTH, each Node reevaluates its Default Route Table. If the Default Route Table is

empty (implying that no Primary Default Route exists), a Router Solicitation is sent out, requesting Router Advertisements from neighboring Nodes. In addition, if the absolute difference between the Node's Overall Node Route Cost at the end of the last period, and the current Overall Node Route Cost is greater than `ROUTE_COST_NOTIF_DIFF`, or the Overall Route Hops has changed since the end of the last period, the Node sends a Routing Advertisement.

After each Router Solicitation, if after `SOLICITATION_PERIOD` time, no Primary Default Route exists, another Router Solicitation is sent out. This SHOULD be modified to use a Trickle or Exponential Backoff Timer.

7.4. Global Topology Construction

[TOC](#)

Each Border Router maintains a view of the topology of the entire L2N Subnet using information obtained from the Topology Reports created by each Node.

Each Node prepares a Topology Report periodically, with period length determined by `TOP_REPORT_PERIOD`. The Node then holds off on sending the report for `TOP_REPORT_WAIT` time, in an attempt to piggyback the report on a data packet. This Topology Report is added to an IP datagram as an IPv6 hop-by-hop options header using a new TLV dispatch value.

For each entry within the top `DEFAULT_TOP_THRESH` slots (inclusive), if the slot is not empty, the entry is flagged as Mature (`Confidence >= CONF_EVICT_THRESH`), or it is set as the Primary Default Route, then its Address, Link Cost Estimate, and Confidence are inserted into the Topology Report.

In addition to Default Route information, the Node SHOULD include requested Node Attribute values, which are defined outside of HYDRO by the application, and SHOULD include its Willingness value

The Topology Report is then sent to the Border Router. The report MAY be added to an outgoing data packet destined to the Border Router, or in a new message generated by the Node. Topology reports SHOULD NOT be inserted into messages where a Flow Match is available in the Flow Entry Table.

The Border Router maintains a graph of the entire network. When it receives a Topology Report from a Node, it first compares the Sequence Number contained in the Topology Report against the last received sequence number (LSN) from that Node. If the Sequence Number is greater than the LSN, or if it is less than `LSN - SEQ_ROLLOVER_THRESH`, or if this is the first report from this Node Router, the Border Router removes the edges created by previous Topology Report from the same Node, and inserts edges based on the information in the latest Topology Report. Edge costs are determined by policy. One option is to use the Link Cost Estimates from the Topology Report. In addition, the Node Attributes can be associated with the appropriate graph vertices.

7.5. Node Forwarding

[TOC](#)

For any given packet at a Node, whether it originated there or is simply being forwarded, a series of steps are taken in order to determine how to send the packet. HYDRO SHOULD specify multiple next-hops for any particular packet when available. If transmission to the first next-hop fails, HYDRO will attempt to use the next next-hop, until all choices provided have failed. The link layer MUST use link-layer acknowledgements in order to determine transmission success or failure, and SHOULD retry failed transmissions. This section describes how the ordered list of potential next-hops is constructed for a particular packet.

1. If the packet contains an IPv6 Routing Header, the address of the next-hop is determined using that route, and the 'Segments remaining' field of the routing header is decremented. The address from the routing header is provided as the first choice next-hop. If a router receives a packet containing a routing header where the segments remaining field is zero but the router is not the destination, the router SHOULD ignore the routing header.
2. If no routing header exists, HYDRO examines the Flow Table to determine if a matching entry exists. The matching process depends on the particular instantiation of the Flow Match data structure.
 - 2.1. If the matching entry indicates the Flow Path Type is FULL_PATH, then this indicates the Flow Path is a full Source Route. If the packet was locally generated, a new routing header is inserted into the packet and the next hop of that path provided as the first choice next-hop.
 - 2.2. If the matching entry indicates the Flow Path Type is NEXT_HOP_PATH, then this indicates the Flow Path is a next hop address, with each Node along the path expected to have a similar entry. This Node address is provided as the first next-hop choice. The NEXT_HOP_PATH data structure MAY store multiple potential next hops, in an ordered list, specified by NUM_HOP_CHOICES. More than one potential next hop MAY be provided using this list.
3. Finally, HYDRO uses the Default Route Table to choose a next-hop, and forwards the packet towards the Border Router. HYDRO will attempt to send to the Primary Default Route first, followed by the top entries of the Default Routing Table.

The number of next-hops that HYDRO attempts to use when forwarding a particular packet is determined by NUM_NEXT_CHOICES. In addition, a packet that is being forwarded MUST NOT select its previous-hop Node as the next-hop Node. If an entry specifies such a case, it is ignored and the process described above continues.

HYDRO receives feedback from the link layer following each transmission about which Default Route next-hops the link was able to successfully use. If the number of consecutive failures of the Primary Default Route is greater than MAX_CONSEC_FAILURES, HYDRO initiates a search for a new Primary Default Route. If another Default Route Table entry exists

whose Route Hops is less than the Primary Default Route's Route Hops, and its Route Cost Estimate is less than the Primary Default Route's Route Cost Estimate, then one of these candidates is selected at random (if multiple choices exist) to be the new Primary Default Route (although the table is not reordered). If no such candidate is found, then the Route Hops comparison is dropped and the same search is performed again with only Route Cost Estimate comparisons. If again no candidate is found, then the Primary Default Route remains unchanged. In addition, at the end of each period, with probability `NEW_PRIMARY_ROUTE_PROB`, this same exploration for an alternate Primary Default Route occurs, allowing HYDRO an opportunity to increase the Confidence in other Default Route Entry Link Cost Estimates.

As a final task at the end of each period, if a Node that is a single hop away from a Border Router determines that it has failed to successfully transmit a single packet to a Border Router (while at least making one attempt), it determines that the link to the Border Router is gone, and SHOULD broadcast a Router Advertisement advertising an Overall Route Cost of `MAX_ROUTE_COST`, and Overall Route Hops of `MAX_HOPS`. Not broadcasting such a message MAY lead to longer convergence times in networks with multiple Border Routers that may fail or be mobile.

After each forwarding attempt, HYDRO reorders the Default Route Table using updated link statistics from the recent attempt. If the packet was transmitted successfully (using Node 'A' as the next-hop) and the entry (A) is not the top entry in the Default Route Table, it swaps positions with the entry above it (B) if its Confidence is greater than `CONF_PROM_THRESHOLD` and one of the following three conditions hold:

1. Overall Route Cost of A + `WILLINGNESS_COST_THRESH` < Overall Route Cost of B.
2. Overall Route Cost of A is less than the Overall Route Cost of B, or greater than it by less than `PATH_COST_DIFF_THRESH`, and the absolute value of the difference between the Willingness of A and the Willingness of B is less than or equal to `WILLINGNESS_THRESH`.
3. A's Overall Route Cost is within `WILLINGNESS_COST_THRESH` of B's Overall Route Cost, and A's Willingness > B's Willingness + `WILLINGNESS_THRESH`.

7.6. Border Router Forwarding

[TOC](#)

Border Routers will receive packets from Nodes in the L2N in the course of normal operation, since Nodes will forward packets to a Border Router when they do not have a matching Flow Table Entry for a packet. The Border Router should process packets which originated on its L2N interface as follows:

NOTE: Some of the following steps assume multiple Border Routers exist in a network. We describe the setup and maintenance of such a topology in a later section.

1. If the destination address is a global unicast address, normal IPv6 processing rules apply; the packet will be routed according to the IPv6 Routing Table Entry for that destination.
2. If the packet contains an IPv6 Routing Header and the first hop of that route is another Border Router, the packet is dropped.
3. If the destination is a unicast address within the L2N subnet and the Border Router does not have a route in its route database, the packet is dropped. It MAY generate an ICMPv6 Host Unreachable message in this case, but implementers SHOULD take care to limit the rate of message generation; it MAY be advantageous in many cases not to generate this message.
4. Otherwise, the Border Router has a stored route to the destination (in the L2N subnet):
 - 4.1 If the Border Router has a route to the destination and the next-hop is another Border Router, the packet is forwarded to that Border Router without modification.
 - 4.2. If the Border Router has a route to the destination and the next-hop is a Node, the Border Router inserts an IPv6 Routing Header containing the source route to the destination, and transmits the packet to the first hop of this route.

7.7. Route Installation

[TOC](#)

In the course of forwarding a packet as described in the 'Border Router Forwarding' section, a Border Router may determine that it is not on the best path between the source and destination Node (where 'best' is defined by a user-specified routing policy, such as lowest cost). Alternatively, there may be administratively configured routes between Nodes in the network. HYDRO SHOULD install routes to enable more efficient routing, but in any case, the exact policy for determining when and what routes are installed is not part of the HYDRO protocol; what is required is that HYDRO Nodes process installation messages as described.

Two types of routes may be installed: hop-by-hop and full path routes. For a hop-by-hop route from Node A to Node B, each intermediate Node N_i on the path $P=[A, N_1, N_2 \dots N_i, B]$ from A to B adds a Flow Table Entry indicating that packets destined to B (and also potentially originating from A) have next hop N_{i+1} . A full path install of this same route would place the entire path P into the Flow Table of Node A; outgoing packets from A destined to B use the path to insert an IPv6 Routing Header containing the route. If the reverse option is selected, intermediate Nodes would have to create Flow Table entries for the reverse direction as well for hop-by-hop route installs, and the full

reverse path would have to be installed at B for full path route installs.

Suppose a Border Router R wishes to install a route between Nodes A and B. A Route Install Header is created and placed into the IPv6 Destination Options header. The 'method' field is set to either HOP_BY_HOP or FULL_PATH, and the 'reverse' field is set to indicate whether or not the path from B to A should also be installed. The Flow Match is initialized to indicate that the destination address of the Flow is B. This header is placed either in a new packet addressed to Node A, or in the header of an existing packet which is destined to Node A.

When Node A receives the Route Install Message, it processes it as part of its processing of the Destination Options extension headers.

Regardless of the method of install, the router searches its Flow Table using the Flow Match in the routing header. If one is present, it replaces it with the new information (if the 'method' is FULL_PATH). If the 'method' is HOP_BY_HOP, since there can be NUM_FLOW_CHOICES for each Flow Table entry, if the same next-hop Node already exists, it is brought to the top of the list. Otherwise, the bottom next-hop in the list is evicted, and the new next-hop is placed at the top of the list. If the Flow Table is full, HYDRO evicts an entry using a Least Recently Used policy. The Node then adds the route contained in the routing header to the Flow Table: it copies into the Flow Path either the entire route, if the method is FULL_PATH, or the next hop, if the method is HOP_BY_HOP.

If the method is HOP_BY_HOP, or if the method is FULL_PATH and the reverse bit is set, the router generates a new packet. The new packet contains an IPv6 Routing Header with the route specified in the Route Install Header. If the method was HOP_BY_HOP, the router adds an IPv6 hop-by-hop option extension containing a new routing header. This new hop-by-hop routing header does not contain the route since that information is present in the IPv6 Routing Header, and so the path_len field is zero.

If a Node receives a Route Install Header in a hop-by-hop option, the method MUST be HOP_BY_HOP. If this is the case, it inserts a new Flow Table entry using the Flow Match contained in the Route Install Header, and consults the routing header to determine the next-hop address to install. The next hop address is the one obtained by following normal processing rules for the source header when forwarding the packet.

7.8. Multicast Forwarding

[TOC](#)

In this so-called "route-over" design, the link-local multicast address is mapped directly to link-local broadcast at the link layer. Thus any packet destined to a link-local multicast group should be sent by the

originator using the link local broadcast facility; these packets clearly MUST NOT be forwarded or retransmitted by any non-origin Node. Additionally, it is often desirable for Nodes to send packets to all Nodes in a subnet, regardless of whether or not they are within a link-local broadcast domain. For this purpose, we use the IPv6 site-local scope. First, we define the well-known address ff05::XXXX to correspond to a subnet-wide flood. Nodes receiving packets destined to this address SHOULD forward these packets using the link-local broadcast address; this may be done using a mechanism like Trickle to reduce dependency on network density; more details about this forwarding is elided from this document.

7.9. Multiple Border Routers

[TOC](#)

It is often desirable for multiple Border Routers to be present on single L2N subnet to provide a diversity of routing options and failover connectivity. If more than one Border Router is present, layer 2 connectivity between them will be assumed, for example using Ethernet, 802.11 mesh, or virtual tunneled IP links (something similar to [\[RFC2003\] \(Perkins, C., "IP Encapsulation within IP," October 1996.\)](#)). The interface the Border Routers use to connect directly to the other Border Routers is named the Secondary Interface. The multiple Border Routers assign themselves unique address on the L2N subnet prefix using their Secondary Interface via manual configuration. On startup, the Border Routers join a link-local multicast group ff02::XXXX on their secondary interface and begin listening for messages on UDP port XXXX.

When a Border Router receives a Topology Report from a Node Router, the Topology Report should be appended to a Topology Report Package header and transmitted to the link-local multicast group. Border Routers receiving Topology Report Packages should process the contained Topology Reports identically to locally received reports.

When Border Routers receive a Topology Report Package, they must annotate the previous hop of the message in their Link Database to indicate if it is a Border Router; that way, they can determine when the next hop to a Node is another Border Router. Furthermore, the transmission of Topology Report messages to other Border Routers ensures that all Border Routers maintain a consistent view of the topology.

8. Extensions

[TOC](#)

The protocol specified in this draft addresses many of the unique demands of this space and the routing requirements spelled out in [the](#)

[IETF Roll Protocols Survey \(Tavakoli, A., Dawson-Haggerty, S., and P. Levis, "Overview of Existing Routing Protocols for Low Power and Lossy Networks," April 2009.\)](#) [I-D.ietf-roll-protocols-survey]. However, there are certain limitations of the design that are left for future revisions of the draft.

First, routing from a Border Router to a Node relies on source routing. Also, the FULL_PATH route install method uses source routing. For paths longer than 5-10 hops, the per-packet overhead of this mechanism becomes unacceptable, as does the state at each Node that must maintain these paths. There are many methods of fixing or ameliorating this problem. Using the current design, deploying additional Border Routers to decrease the network depth may be an option in some situations. Alternatively, a later protocol specification could include mechanisms for loose source routing between landmarks, which might be determined by Border Routers. This would serve as a hybrid of the source routing and hop-by-hop route install options. The need to test the efficiency and functionality of these adaptations, particularly in the face of multiple link failures, preclude their full specification at this point in the draft.

Another issue is that of network numbering. The current draft requires that Border Routers with two interfaces to assign themselves different address on the same L2N prefix to both of these interfaces. This also occurs in the (experimental) design for IPv6 Neighbor Discovery Proxies [RFC4389] (Thaler, D., Talwar, M., and C. Patel, "Neighbor Discovery Proxies (ND Proxy)," April 2006.), but seems to the authors to be somewhat inelegant. It may be possible to devise a better solution.

9. Configuration Parameters and Other Administrative Options

[TOC](#)

CONF_EVICT_THRESHOLD: Setting of this MAY be dependent on the data rate (because that determines how long it takes to get a Mature entry).

Initial implementations use '5'.

CONF_PROM_THRESHOLD: Threshold for being able to promote an entry in the Default Route Table. Dependent on link layer.

DEFAULT_TOP_THRESH: The maximum number of Default Route Entries (from the top) that MAY be included in a Topology Report. Recommended value is 4.

LINK_ADMIT_THRESH: This value depends on the availability of hardware based link quality metrics, (such as LQI for CC2420-Based Hardware Platforms), and the setting of the threshold may depend on the actual deployment environment as well.

LINK_QUALITY_DIFF_THRESH: Will vary based on link quality metric used. For CC2420 LQI, 10 is the recommended value.

MAX_CONSEC_FAILURES: Maximum number of consecutive failures after which a new Primary Default Route SHOULD be used. Recommended value is 20 (if link layer retransmissions are included).

MAX_HOPS: Upper bound value to denote invalid entry.

MAX_ROUTE_COST: Upper bound value to denote invalid entry.

NEW_PRIMARY_ROUTE_PROB: Probability of looking for another Primary Default Route at the end of a period. Recommended value is 25%.

NUM_DEFAULT_ENTRIES: The size of the Default Route Table, recommended to be '8'.

NUM_FLOW_CHOICES: The maximum number of next-hop choices for a HOP_BY_HOP Flow Table Entry. Recommended value is 1.

NUM_NEXT_CHOICES: The maximum number of next-hop choices to be used when attempting to send a packet.

PATH_COST_DIFF_THRESH: Threshold for determining whether two Route Costs are similar. Depends on Metric used for Route Costs, but if ETX (Expected Transmissions) are used, the recommended value is 1.

PERIOD_LENGTH: Application dependent, will affect control traffic rate and convergence rate, potentially.

ROUTE_COST_NOTIF_DIFF: Difference in the route cost that merits a Route Advertisement. Route Cost is metric dependent, but if ETX is used, then a recommended value is 0.5.

SOLICITATION_PERIOD: Tradeoff between control traffic generated and route formation.

TOP_REPORT_PERIOD: To meet the ROLL requirements, this should be lower bounded by the inverse of the data rate.

TOP_REPORT_WAIT: This should be set to the variance in data reporting periods.

WILLINGNESS_COST_THRESH: The buffer in Overall Route Cost that SHOULD be traded for a more willing Node to route packets. This is dependent on a particular application's tolerance for routing stretch to reduce load on certain Nodes.

WILLINGNESS_THRESH: The difference in Willingness that is necessary for promoting a Default Route Entry that gives up less Overall Route Cost than WILLINGNESS_COST_THRESH but more than PATH_COST_DIFF_THRESH. Is dependent on the range in Willingness.

10. Applicability Statement

[TOC](#)

HYDRO is designed for use in L2Ns that are composed of a two-tiered hierarchy in which a set of Nodes are supported by a much smaller set of more capable Border Routers. In other words, HYDRO was designed for scenarios laid out by the ROLL working group, as detailed in the [Protocols Survey \(Tavakoli, A., Dawson-Haggerty, S., and P. Levis, "Overview of Existing Routing Protocols for Low Power and Lossy Networks," April 2009.\)](#) [I-D.ietf-roll-protocols-survey]. As such, we now provide a brief analysis of HYDRO's performance relative to the five quantified metrics, demonstrating that it passes all five criteria.

Table Scalability: The size of the Default Route Table is limited by a constant, NUM_DEFAULT_ENTRIES, and as such is $O(1)$. The Flow Table scales with the number of destinations, and so the total amount of routing state is $O(D)$, which meets the criteria.

Loss Response: In the case that a link fails, this is only detected when trying to send a data packet over the link. In such a case, the data packet is then forwarded over another path, potentially leading to the installation of a new route in the network. The updated quality of the link is reported with the next Topology Report. This limits the scope of the loss response only to the active path, achieving a pass.

Control Cost: The three forms of control cost are route formation, Topology Reports, and Route Install Messages. Route Advertisement messages initiated by the Border Router can be turned down to arbitrarily low frequencies, and Route Solicitation messages by Nodes are only of single-hop scope and triggered by failure of all Default Routes. Route Install Messages are a product of data traffic, and so subsequently HYDRO meets the requirement of being bound by the data rate plus a small constant.

Link Cost: Link qualities play an integral part of routing in HYDRO.

Node Cost: Each Node can report the desired Node Attributes as part of its Topology Reports, as well as its Willingness, and the objective function for calculating routes at the Border Router can incorporate these.

11. Security Considerations

[TOC](#)

12. IANA Considerations

[TOC](#)

This document includes no request to IANA.

13. Acknowledgements

[TOC](#)

Special thanks to members of the ROLL working group for helpful comments and discussion that aided in the design of HYDRO.

14. References

[TOC](#)

14.1. Normative References

[TOC](#)

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

14.2. Informative References

[TOC](#)

[I-D.ietf-roll-building-routing-reqs]	Martocci, J., Riou, N., Mil, P., and W. Vermeylen, " Building Automation Routing Requirements in Low Power and Lossy Networks ," draft-ietf-roll-building-routing-reqs-09 (work in progress), January 2010 (TXT).
[I-D.ietf-roll-home-routing-reqs]	Brandt, A. and J. Buron, " Home Automation Routing Requirements in Low Power and Lossy Networks ," draft-ietf-roll-home-routing-reqs-11 (work in progress), January 2010 (TXT).
[I-D.ietf-roll-indus-routing-reqs]	Networks, D., Thubert, P., Dwars, S., and T. Phinney, " Industrial Routing Requirements in Low Power and Lossy Networks ," draft-ietf-roll-indus-routing-reqs-06 (work in progress), June 2009 (TXT).
[I-D.ietf-roll-protocols-survey]	Tavakoli, A., Dawson-Haggerty, S., and P. Levis, " Overview of Existing Routing Protocols for Low Power and Lossy Networks ," draft-ietf-roll-protocols-survey-07 (work in progress), April 2009 (TXT).
[I-D.ietf-roll-urban-routing-reqs]	Dohler, M., Watteyne, T., Winter, T., Barthel, D., Jacquenet, C., Madhusudan, G., and G. Chegaray, " Urban WSNs Routing Requirements in Low Power and Lossy Networks ," draft-ietf-roll-urban-routing-reqs-05 (work in progress), March 2009 (TXT).
[RFC2003]	Perkins, C. , " IP Encapsulation within IP ," RFC 2003, October 1996 (TXT , HTML , XML).
[RFC2460]	Deering, S. and R. Hinden , " Internet Protocol, Version 6 (IPv6) Specification ," RFC 2460, December 1998 (TXT , HTML , XML).
[RFC2461]	Narten, T. , Nordmark, E. , and W. Simpson , " Neighbor Discovery for IP Version 6 (IPv6) ," RFC 2461, December 1998 (TXT , HTML , XML).
[RFC4389]	Thaler, D., Talwar, M., and C. Patel, " Neighbor Discovery Proxies (ND Proxy) ," RFC 4389, April 2006 (TXT).
[RFC4944]	Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, " Transmission of IPv6 Packets over IEEE 802.15.4 Networks ," RFC 4944, September 2007 (TXT).
[RFC5095]	

Abley, J., Savola, P., and G. Neville-Neil,
“[Deprecation of Type 0 Routing Headers in IPv6](#),”
RFC 5095, December 2007 ([TXT](#)).

Authors' Addresses

[TOC](#)

	Arsalan Tavakoli
	UC Berkeley
	Soda Hall, UC Berkeley
	Berkeley, CA 94720
	USA
Email:	arsalan@cs.berkeley.edu
	Stephen Dawson-Haggerty
	UC Berkeley
	Soda Hall, UC Berkeley
	Berkeley, CA 94720
	USA
Email:	stevedh@cs.berkeley.edu
	Jonathan W. Hui
	Arch Rock Corporation
	501 2nd St. Ste. 410
	San Francisco, CA 94107
	USA
Email:	jhui@archrock.com
	David Culler
	UC Berkeley
	Soda Hall, UC Berkeley
	Berkeley, CA 94720
	USA
Email:	culler@cs.berkeley.edu