

Individual Submission
Internet Draft
[draft-haberman-ipngwg-auto-prefix-00.txt](#)
November 2000

B. Haberman
J. Martin
Nortel Networks

Automatic Prefix Delegation Protocol for
Internet Protocol Version 6 (IPv6)

<[draft-haberman-ipngwg-auto-prefix-00.txt](#)>

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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/lid-abstracts.txt>.

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

Abstract

The expansion of the IP address space provided by IPv6 makes it both possible and reasonable to allocate entire subnets to environments that had been previously limited to a few individual IP addresses. Other protocols such as Neighbor Discovery and Stateless Address Autoconfiguration allow hosts within those subnets to be automatically configured. The router between this subnet and the upstream world requires just one more piece to make this process automatic, a network prefix.

This document describes a mechanism for the automated delegation of an IPv6 network prefix. It allows routers to request a specific size prefix and inform the upstream router of the routing protocols of which it is capable. Upon authorizing the request the delegating router then returns a prefix, the desired routing protocol, and a

lifetime for the use of the prefix.

Haberman, Martin

1

Internet Draft <Automatic Prefix Delegation>

May 2001

1. Introduction

This specification defines the Prefix Delegation (PD) protocol for Internet Protocol Version 6 (IPv6). Routers use Prefix Delegation to request a network prefix for use on directly attached networks. Prefix Delegation also allows the requesting router to specify the routing protocols in which it is capable of participating. Upon receipt of the request, the delegating router may authenticate the request, and will establish if the requested prefix size is acceptable. The delegating router then specifies the prefix for use, the length of time for which that prefix is delegated, and the routing protocol to be used.

Unless specified otherwise (in a document that covers operating IP over a particular link type) this document applies to all link types. However, because PD uses link-layer multicast, it is possible that on some link types (e.g., NBMA links) alternative mechanisms to implement PD must be specified (in the appropriate document covering the operation of IP over a particular link type).

2. Terminology

2.1 General

This document uses the terminology defined in [[RFC 2460](#)] and [RFC 2461] and in addition:

Requesting Router - The router that is requesting that a prefix be assigned

Delegating Router - The router that is responding to the prefix request

2.2 Addresses

Prefix Delegation makes use of a number of different addresses defined in [ADDR-ARCH], including:

Global address – A unicast address having global scope

2.3 Requirements

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 [[RFC 2119](#)].

3. Scope of Work

This proposal is meant to give a singly homed leaf router the ability to obtain an IPv6 prefix that can be used within its leaf

Haberman, Martin

2

Internet Draft <Automatic Prefix Delegation> May 2001
network. Future revisions of this document may support a more generic approach to dynamic prefix delegation.

It is also assumed that the delegating server/router shares a network connection with the requesting router. Future revisions may remove this restriction and allow for either multi-hop messages or a relay function.

4. Protocol Overview

The Prefix Delegation protocol defines two new ICMP message types, the Prefix Request and the Prefix Delegation. The Prefix Request is used by the Requesting Router to communicate requests to the Delegating Router. Conversely, the Prefix Delegation is used by the Delegating Router to communicate prefix and error information with the Requesting Router.

4.1 Delegator Query

The Requesting Router begins the Prefix Delegation process by sending a Prefix Request message of type [DELEGATOR QUERY] to the ALL-DELEGATORS link-local multicast address (XXXX::XX).

Upon receipt of the Delegator query, a Delegating Router determines if it is configured to provide prefixes of the specified scope. If so, it unicasts a Prefix Delegation of type Prefix Delegator to the Requestor. If not, the message is silently discarded.

After sending the query, the Requestor waits for Query Interval (Default: 5) seconds for one or more Delegating Routers to respond. If there is no response, the Delegator Query is sent again up to Max Query times (Default: 3). If no response is received, there are no Prefix Delegation services available, and Prefix Delegation has failed.

If more than one response is received to the query, the response with the numerically highest source IP address is used.

4.2 Initial Request

Once a Delegating Router is chosen, the Requestor sends a Prefix Request message of type Initial Request to the unicast IP address of the Delegating Router.

The Requestor may or may not have a Security Association with the Delegating Router, however if Authentication is required and no SA is present, the Delegator will reject the request with an error response indicating that Authentication is required. The Requestor then builds a Security Association with the Delegator and sends another Initial Request including the SA information.

Internet Draft <Automatic Prefix Delegation> May 2001
If no response is heard within Request Timeout seconds (Default: 5), the Initial Request should be sent again, up to Max Initial Request (Default: 3) tries. If no response is heard, a Delegator Query is sent and the process restarted. If this cycle is repeated Max Delegation Attempts times (Default: 3), Prefix Delegation has failed.

4.3 Authentication and Authorization

Upon receipt of the Prefix Request of any type, the Delegating Router establishes if there is a need for Authentication, based upon local policy. If Authentication is required and none is provided, the Delegator will return a Prefix Delegation message, with a code of Authentication Required.

Once the authentication credentials of the requestor, if present,

are established, any request that requires the allocation of a prefix must be checked for Authorization. Authorization is established by verifying that the requested prefix length for the specific Requestor is acceptable by locally configured policy. If the prefix length requested falls outside of policy, a Prefix Delegation error message of type Not Authorized is returned.

4.4 Prefix Delegation

After the request is verified to be acceptable, the Delegating Router allocates the requested prefix size from its pool of available addresses. The creation and management of that pool is beyond the scope of this document, but it can be supposed that minimalistically a Delegating Router will be statically configured with a fixed pool. If no acceptable prefix is available, a Prefix Delegation message with a code of Prefix Unavailable is returned.

The Delegating Router then compares the list of available routing protocols in the Request against its own capabilities and the local policies regarding routing. For the purposes of this comparison, static routes are considered a routing protocol. If no acceptable match is found, static routes are used.

The Delegating Router then sends a Prefix Delegation message to the Requesting Router containing a code of Prefix Delegation and all of the prefix and routing information. The Delegating router then activates the negotiated protocol on the interface to which the Requestor is attached.

4.5 Prefix Refresh

All Prefix Delegations have a finite lifetime. Upon receiving a Prefix Delegation, the requesting router initiates a timer such that before the lifetime is expired, the Requesting Router sends a Prefix Request with a Renewal Refresh code directly to the Delegating router.

A Requesting Router receiving the Prefix Unavailable code, or no response at all, has not had the prefix renewed. It will expire at the end of the initial lifetime. To acquire a new prefix, the Requesting Router must begin anew as described in [Section 4.1](#).

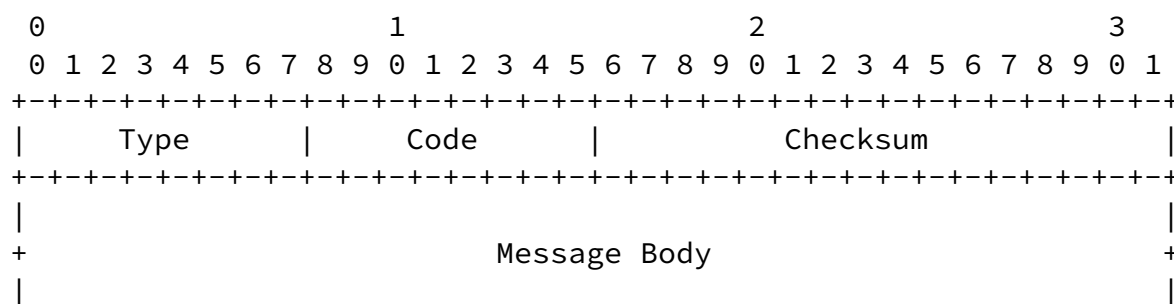
4.6 Prefix Return

If the Requesting Router no longer requires the use of a prefix, it can return that prefix to the control of the Delegating Router through the use of the Prefix Return code in a Prefix Request. The requesting router sends a Prefix Request directly to the Delegating Router.

Upon receipt and verification (if needed) of this message, the Delegating Router returns the prefix to the pool and issues a Prefix Delegation with a code of Prefix Returned.

5. Messages

All messages have the following general format:



5.1 Prefix Request Message

The Prefix Request Message is sent to request, renew, or release a prefix.

Internet Draft										<Automatic Prefix Delegation>										May 2001																			
0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
Type										Code										Checksum																			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
S Prefix Length										Routing Capabilities													Reserved																
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
+																																							
+																																							
+																																							
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							

IP Fields

Source Address

An IP address assigned to the sending interface.

Destination Address

The ALL-DELEGATORS link-local multicast address (XXXX::XX) for Delegator Query messages. All other Prefix Request messages should be sent to a unicast address of the Delegating Router.

Authentication Header

If a Security Association for the IP Authentication Header exists between the sender and the destination address, then the sender SHOULD include this header. No such header is required for the initial prefix request that is multicast, but may be required for further progress.

ICMP Fields

Type

XXX (Where XXX is assigned by IANA)

Code

They Type of Request Code:

Delegator Query (0)

The Delegator Query is used by the Requestor to identify potential Delegating Routers. It is sent to the all-delegators link-local multicast address with no Authentication Header. For this Query, the Scope field is required. Unused fields MUST be set to zero.

Initial Request (1)

The Initial Request is used to initiate the request process. It is sent to the unicast IP address of the Delegating Router, and may carry an Authentication

Haberman, Martin

6

Internet Draft <Automatic Prefix Delegation> May 2001
Header. For this initial request, the Scope, Prefix Length, and Routing Capabilities fields are required. Unused fields MUST be set to zero.

Renewal Request (2)

The Renewal Request is used to renew a prefix that has been previously allocated. It is sent to a unicast IP address of the Delegating Router and may carry an Authentication Header. For the renewal, the Scope, Prefix Length, Routing Capabilities, and Prefix fields are required.

Prefix Return (3)

The Prefix Return is used to return an unused prefix, or portion of a prefix to the control of the Delegating Router. It is sent to a unicast IP address of the Delegating Router and may carry an Authentication Header. For the Return, the Scope, Prefix Length, and Prefix fields are required. Unused fields MUST be set to zero.

Checksum

The ICMP checksum as defined in [[RFC 2463](#)].

Prefix Request Fields

S

A one bit Scope Flag. A value of zero (0) indicates that the request is for a prefix of Global Scope, a one (1) indicates site-local.

Prefix Length

The length of the prefix being requested, renewed, or released.

Routing Capabilities

This bit-field allows the requestor to specify the routing protocols in which it is capable of participating. For the purposes of this field, a static route is considered a routing protocol.

At this time, the only defined value is zero (0), indicating that the requestor is capable of static routes. This is an area for further work.

Reserved

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

IPv6 Prefix

Haberman, Martin

7

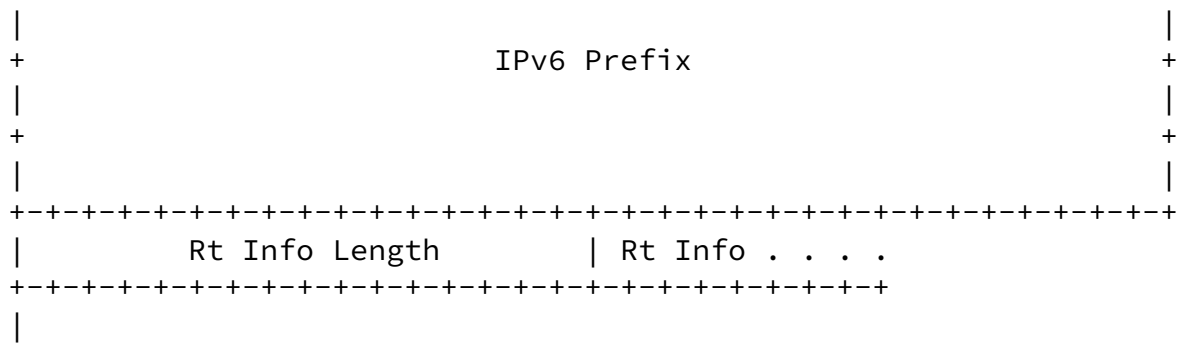
Internet Draft <Automatic Prefix Delegation> May 2001

The Prefix field is used to carry a previously assigned prefix. The host portion of the IP address MUST be padded with zeros.

5.2 Prefix Delegation Message

The Prefix Delegation Messages are sent to provide the addresses of available Prefix Delegates, to provide prefix and routing data, and for error returns.

[illegible]



IP Fields

Source Address

An IP address assigned to the sending interface.

Destination Address

The IP address of the Requestor as specified by the Source Address of the Prefix Request message.

Authentication Header

If a Security Association for the IP Authentication Header exists between the sender and the destination address, then the sender SHOULD include this header.

ICMP Fields

Type

XXX+1 (Where XXX+1 is assigned by IANA)

Code

The Type of Response Code:

The Prefix Delegator is used by the Delegator to inform the Requestor that it is available to provide prefixes of the desired type. It is sent to the unicast IP address in the Source Address portion of the Prefix Request packet. For this Response, the Scope field is required. Unused fields MUST be set to zero.

Authentication Required (1)

The Authentication Required message indicates to the Requestor that a Security Association must be

established before a prefix can be delegated. It is sent to the unicast IP address in the Source Address portion of the Prefix Request packet. For this message, no additional fields are required. Unused fields MUST be set to zero.

Authorization Failed (2)

The Authorization Failed message indicates to the Requestor that either it is not authorized to request a prefix, or that the prefix requested fell outside of local policy. It is sent to the unicast IP address in the Source Address portion of the Prefix Request packet. For this message, no additional fields are required. Unused fields MUST be set to zero.

Prefix Unavailable (3)

The Prefix Unavailable indicates that the Prefix Request was acceptable, but the Delegator does not have sufficient available address space to fulfill the request. It is sent to the unicast IP address in the Source Address portion of the Prefix Request packet. For this message, no additional fields are required. Unused fields MUST be set to zero.

Prefix Delegated (4)

The Prefix Delegated message actually provides the prefix information that the Requestor has requested. It is sent to the unicast IP address in the Source Address portion of the Prefix Request packet. For this message, all fields are required.

Prefix Returned (5)

The Prefix Return is used to confirm the return of a prefix. It is sent to the unicast IP address in the Source Address portion of the Prefix Request packet. For this message, the Prefix Length and IPv6 Prefix fields are required.

Checksum

The ICMP checksum.

S

A one bit Scope Flag. A value of zero (0) indicates that the response is for a prefix of Global Scope, a one (1) indicates site-local.

Prefix Length

The length of the prefix being provided.

Lifetime

The time (in seconds) that the Requestor is permitted to use the allocated prefix. At the end of this period, the Delegator assumes control of the prefix. This lifetime can be extended through the renewal process.

Rt Proto

This field specifies the routing protocol in which the Requestor should participate.

At this time, the only defined value is zero, for Static Routes. This is an area for further development of this document.

IPv6 Prefix

The Prefix field is used to carry the assigned prefix. The host portion of the IP address MUST be padded with zeros.

Rt Info Length

The length, in octets, of the Routing Information field. At this time, since Static is the only defined protocol; this field should have a value of zero.

Routing Information

This field carries protocol specific information to allow the Requesting router to configure itself to participate in routing.

This field will be described in later versions of this document. At this time, since Static is the only defined protocol; this field should be zero length.

6. To Do's

- Additional security discussion
- Expand routing protocol negotiation
- Multiple hops between requestor and delegator
- Cascading delegations
- Removal of leaf network restriction

Internet Draft <Automatic Prefix Delegation>

May 2001

7. References

- [RFC 2460] S. Deering and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.
- [RFC 2461] T. Narten, E. Nordmark, and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", [RFC 2461](#), December 1998.
- [RFC 2119] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), [BCP 14](#), March 1997.
- [RFC 2463] A. Conta and S. Deering, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", [RFC 2463](#), December 1998.

Authors' Addresses

Brian Haberman
Nortel Networks
4309 Emperor Blvd.
Suite 200
Durham, NC 27703

Phone : 1-919-992-4439
Email : haberman@nortelnetworks.com

Jim Martin
Nortel Networks
4401 Great America Parkway
Santa Clara, Ca 95054

Phone : 1-408-495-3792
Email : jrm@nortelnetworks.com

Haberman, Martin

11