Individual Submission
Internet Draft
draft-haberman-ipngwg-auto-prefix-02.txt
February 2002
Expires August 2002

B. Haberman
J. Martin

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

<draft-haberman-ipngwg-auto-prefix-02.txt>

Status of this Memo

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 either a specific prefix or any prefix. Upon authorizing the request the delegating router then returns a prefix and a lifetime for the use of the prefix. Optionally, the delegating and requesting routers can exchange routing protocol information.

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. 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 and the length of time for which that prefix is delegated.

The Prefix Delegation protocol supports extensible options. These options may be used to negotiate additional operational parameters, such as routing protocol information.

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 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 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-Routers link-local multicast address (FF02::2).

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 within the Query Interval, 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 chosen 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.

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 Message Security

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

The building of a Security Association between the Delegator and the Requestor is based on the Authentication and/or Encapsulated Security Payload extension headers defined in [RFC 2402] and [RFC 2406].

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 sends a Prefix Delegation message to the Requesting Router containing a code of Prefix Delegation and all of the prefix information. The Requesting Router then activates the prefix on its interface of choice.

4.5 Prefix Refresh

All Prefix Delegations have a lifetime that MUST follow the rules defined in Section 4.6.2 of [RFC 2461]. Upon receiving a Prefix Delegation, the requesting router initiates a timer such that before the lifetime expires, the Requesting Router sends a Prefix Request

with code=REFRESH directly to the Delegating router.

If the Requestor receives no response within [RENEWAL TIMEOUT] seconds (Default: 5), the Renewal Request should be sent again, up to [MAX RENEWAL REQUEST] (Default: 3) tries. If no response is heard the previously allocated prefix is not renewed.

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:

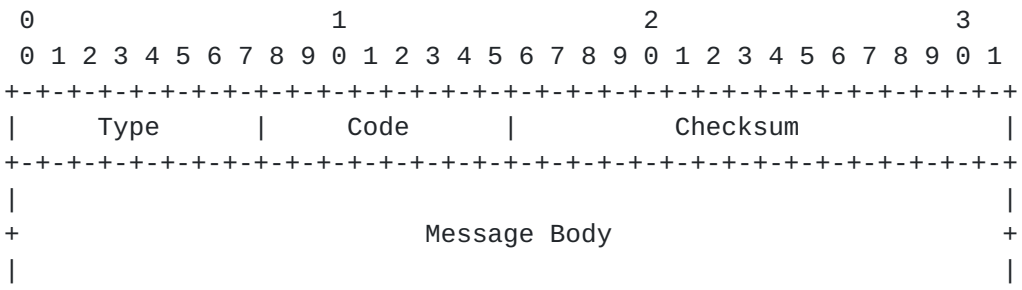

The following sections describe the specific messages and options used in delegating IPv6 prefixes.

5.1 Prefix Request Message

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

prefix.

IP Fields

Source Address
An IP address assigned to the sending interface.

Destination Address
The All-Routers link-local multicast address (FF02::2)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-Routers link-local multicast address with no Authentication Header.

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 Header. Unused fields MUST be set to zero. An Initial Request code MAY contain a Prefix Option.

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. A Renewal Request code MUST contain at least one Prefix Option.

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. A Prefix Return code MUST contain at least one Prefix Option.

Checksum
The ICMP checksum as defined in [RFC 2463].

5.2 Prefix Delegation Message Format

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

IP Fields

Source Address
An IP address assigned to the sending interface.

Destination Address
The IP address of the Requestor as specified by the IP Source Address in 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:

Prefix Delegator (0)

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. 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. 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. 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. 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, a Prefix Option MUST be included.

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 Option MUST be included.

Checksum

The ICMP checksum.

5.3 Prefix Option

The Subnet Prefix Option is used to relay prefix information between Requestors and Delegators. It has the following format:

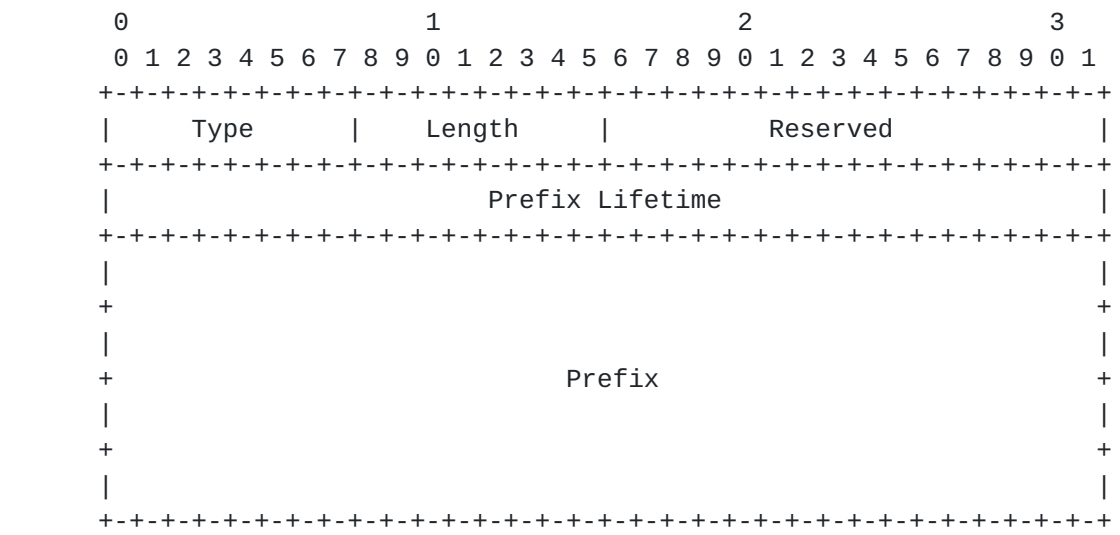

Prefix Option Fields

Type = 0x01
The field identifies the presence of a subnet prefix. This option MUST follow either a Prefix Request header or a Prefix Delegation header.

Length
The length of the prefix contained in the option.

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

Prefix Lifetime
The lifetime of the prefix contained in the option.

IPv6 Prefix
The Prefix field is used to carry a subnet prefix. The host portion of the IP address MUST be padded with zeros.

6. Security Considerations

The ability to automate the delegation of prefixes opens several security vulnerabilities. Rogue delegators can issue bogus prefixes to requestors. This may cause denial of service due to unreachability. Rogue requestors may consume valuable resources from legitimate delegators, thus denying others the use of the prefixes. For these reasons, the use of IPSec-based Authentication and/or Encryption is suggested.

7. To Do's

- Additional security discussion
- Relay functionality
- Routing capabilities option

8. Acknowledgements

We would like to acknowledge and thank Jun-ichiro itojun Hagino, Dave Thaler, Yamasaki Toshi, Ole Troan, and Kazuaki Tsuchiya for their feedback and suggestions on this document.

9. 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
haberman@lorien.sc.innovationslab.net

Jim Martin
jim@interop.net