

Network Working Group
Request for Comments: 3769
Category: Informational

S. Miyakawa
NTT Communications Corporation
R. Droms
Cisco
June 2004

Requirements for IPv6 Prefix Delegation

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2004).

Abstract

This document describes requirements for how IPv6 address prefixes should be delegated to an IPv6 subscriber's network (or "site").

1. Introduction

With the deployment of IPv6 [1], several Internet Service Providers are ready to offer IPv6 access to the public. In conjunction with widely deployed "always on" media such as ADSL and the expectation that customers will be assigned a /48 IPv6 unicast address prefix (see RFC 3513 [3] and section 3 of RFC 3177 [2]), an efficient mechanism for delegating address prefixes to the customer's sites is needed. The delegation mechanism will be intended to automate the process of informing the customer's networking equipment of the prefixes to be used at the customer's site.

This document clarifies the requirements for IPv6 address prefix delegation from the ISP to the site.

2. Scenario and terminology

The following figure illustrates a likely example for the organization of a network providing subscription IPv6 service:

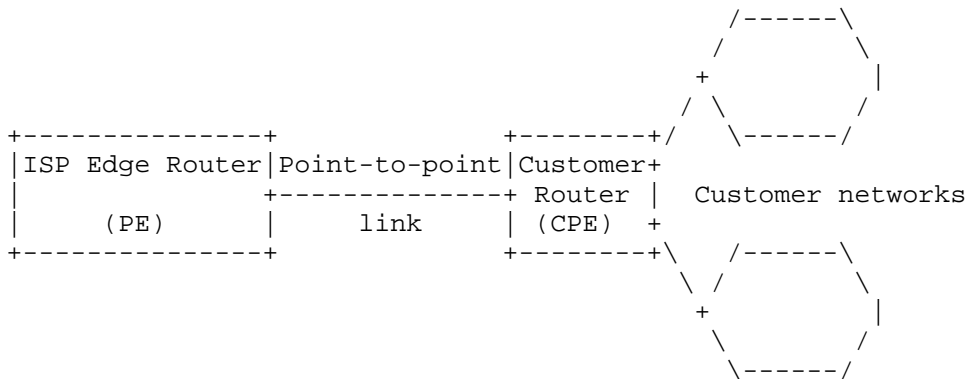


Figure 1: Illustration of ISP-customer network architecture

Terminology:

PE: Provider edge device; the device connected to the service provider's network infrastructure at which the link to the customer site is terminated

CPE: Customer premises equipment; the device at the customer site at which the link to the ISP is terminated

3. Requirements for Prefix Delegation

The purpose of the prefix delegation mechanism is to delegate and manage prefixes to the CPE automatically.

3.1. Number and Length of Delegated Prefixes

The prefix delegation mechanism should allow for delegation of prefixes of lengths between /48 and /64, inclusively. Other lengths should also be supported. The mechanism should allow for delegation of more than one prefix to the customer.

3.2. Use of Delegated Prefixes in Customer Network

The prefix delegation mechanism must not prohibit or inhibit the assignment of longer prefixes, created from the delegated prefixes, to links within the customer network. The prefix delegation mechanism is not required to report any prefix delegations within the customer's network back to the ISP.

3.3. Static and Dynamic Assignment

The prefix delegation mechanism should allow for long-lived static pre-assignment of prefixes and for automated, possibly short-lived, on-demand, dynamic assignment of prefixes to a customer.

3.4. Policy-based Assignment

The prefix delegation mechanism should allow for the use of policy in assigning prefixes to a customer. For example, the customer's identity and type of subscribed service may be used to determine the address block from which the customer's prefix is selected, and the length of the prefix assigned to the customer.

3.5. Expression of Requirements or Preferences by the CPE

The CPE must be able to express requirements or preferences in its request to the PE. For example, the CPE should be able to express a preference for a prefix length.

3.6. Security and Authentication

The prefix delegation mechanism must provide for reliable authentication of the identity of the customer to which the prefixes are to be assigned, and must provide for reliable, secure transmission of the delegated prefixes to the customer.

The prefix delegation should provide for reliable authentication of the identity of the service provider's edge router.

3.7. Accounting

The prefix delegation mechanism must allow for the ISP to obtain accounting information about delegated prefixes from the PE.

3.8. Hardware technology Considerations

The prefix delegation mechanism should work on any hardware link technology between the PE and the CPE and should be hardware technology independent. The mechanism must work on shared links.

The mechanism should work with all hardware technologies with either an authentication mechanism or without, but ISPs would like to take advantage of the hardware technology's authentication mechanism if it exists.

4. Security considerations

Section 3.6 specifies security requirements for the prefix delegation mechanism. For point to point links, where one trusts that there is no man in the middle, or one trusts layer two authentication, authentication may not be necessary.

A rogue PE can issue bogus prefixes to a requesting router. This may cause denial of service due to unreachability.

A rogue CPE may be able to mount a denial of service attack by repeated requests for delegated prefixes that exhaust the PE's available prefixes.

5. Acknowledgments

The authors would like to express thanks to Randy Bush, Thomas Narten, Micheal Py, Pekka Savola, Dave Thaler, as well as other members of the IPv6 working group and the IESG for their review and constructive comments. The authors would also like to thank the people in the IPv6 operation group of the Internet Association of Japan and NTT Communications IPv6 project, especially Toshi Yamasaki and Yasuhiro Shirasaki for their original discussion and suggestions.

6. Informative References

- [1] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [2] IAB and IESG, "IAB/IESG Recommendations on IPv6 Address", RFC 3177, September 2001.
- [3] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, April 2003.

7. Authors' Addresses

Shin Miyakawa
NTT Communications Corporation
Tokyo
Japan

Phone: +81-3-6800-3262
EMail: miyakawa@nttv6.jp

Ralph Droms
Cisco
1414 Massachusetts Avenue
Boxborough, MA 01719
USA

Phone: +1 978.936.1674
EMail: rdroms@cisco.com

8. Full Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.