

IPv6 Working Group
INTERNET-DRAFT
May 4, 2004
Expires November 2004

R. Hinden
Nokia
D. Thaler
Microsoft

IPv6 Host to Router Load Sharing
<[draft-ietf-ipv6-host-load-sharing-02.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/1id-abstracts.txt>

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

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

The original IPv6 conceptual sending algorithm does not do load-sharing among equivalent IPv6 routers, and suggests schemes which can be problematic in practice. This document updates the conceptual sending algorithm so that traffic to different destinations can be distributed among routers in an efficient fashion.

1. Introduction

In the conceptual sending algorithm in [ND] and in the optional extension in [ROUTERSEL], a next hop is chosen when no destination cache entry exists for an off-link destination or when communication through an existing router is failing. Normally a router is selected the first time traffic is sent to a specific destination IP address. Subsequent traffic to the same destination address continues to use the same router unless there is some reason to change to a different router (e.g., a redirect message is received, or the router is found to be unreachable).

In both the base algorithm and in the optional extension, sometimes a host has a choice of multiple equivalent routers for a destination. That is, all other factors are equal and a host must break a tie via some implementation-specific means.

It is typically desirable when there is more than one equivalent router that hosts distribute their outgoing traffic among these routers. This shares the load among multiple routers and provides better performance for the host's traffic.

[ND] does not require any particular behavior in this respect. It specifies that an implementation may always choose the same router (e.g., the first in the list) or may cycle through the routers in a round-robin manner. Both of these suggestions are problematic.

Clearly, always choosing the same router does not provide load sharing. Some problems with load sharing using naive tie-breaking techniques such as round-robin and random are discussed in [MULTIPATH]. While the destination cache provides some stability since the determination is not per-packet, cache evictions or timeouts can still result in unstable or unpredictable paths over time, lowering the performance and making it harder to diagnose problems. Round-robin selection may also result in

Expires November 2004

[Page 2]

synchronization issues among hosts, where in the worst case the load is concentrated on one router at a time.

In the remainder of this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in [[RFC2119](#)].

2. Load Sharing

When a host chooses from multiple equivalent routers, it SHOULD support choosing using some method which distributes load for different destinations among the equivalent routers rather than always choosing the same router (e.g., the first in the list). Furthermore, a host that does attempt to distribute load among routers SHOULD use a hash-based scheme, such as those described in [[MULTIPATH](#)], which takes the destination IP address into account.

Note that traffic for a given destination address will use the same router as long as the Destination Cache Entry for the destination address is not deleted. With a hash-based scheme, traffic for a given destination address will use the same router over time even if the Destination Cache Entry is deleted, as long as the list of equivalent routers remains the same.

3. Acknowledgments

The authors of this document would like to thank Erik Nordmark, Brian Haberman, Steve Deering, Aron Silverton, and Christian Huitema.

4. Security Considerations

As mentioned in [[MULTIPATH](#)], when next-hop selection is predictable, an application can synthesize traffic that will all hash the same, making it possible to launch a denial-of-service attack against the load sharing algorithm, and overload a particular router. This can even be done by a remote application that can cause a host to respond to a given destination address. A special case of this is when the same (single) next-hop is always selected, such as in the algorithm allowed by [[ND](#)]. Introducing hashing can make such an attack more difficult; the

more unpredictable the hash is, the harder it becomes to conduct a denial-of-service attack against any single router.

5. Normative References

[ND] Narten, T., Nordmark, E. and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", [RFC 2461](#), December 1998.

[RFC2119]

Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), [BCP0014](#), March 1997.

6. Informative References

[MULTIPATH]

Thaler, D. and C. Hopps, "Multipath Issues in Unicast and Multicast Next-Hop Selection", [RFC 2991](#), November 2000.

[ROUTERSEL]

Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", Work in progress, [draft-ietf-ipv6-router-selection-03.txt](#), December 2003.

7. Authors' Addresses

Robert Hinden
Nokia
313 Fairchild Drive
Mountain View, CA 94043
Phone: +1 650 625-2004
Email: bob.hinden@nokia.com

Dave Thaler
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
Phone: +1 425 703 8835
EMail: dthaler@microsoft.com

8. Revision History

(This section to be removed before publication as an RFC)

Changes from [draft-ietf-ipv6-load-sharing-01.txt](#):

- o Changed load sharing from a MUST to a SHOULD.
- o Added standard IETF intellectual property notice.

Changes from [draft-ietf-ipv6-router-selection-02.txt](#):

- o Split load sharing back into its own document.
- o Made hash-based, rather than random, the rule.

9. Full Copyright Statement

Copyright (C) The Internet Society (2004). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

10. 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.

