

Network Working Group  
Internet-Draft  
Expires: April 1, 2005

A. Lindem  
Redback Networks  
A. Oswal  
Cisco Systems  
October 2004

OSPFv3 Destination Address Filter  
draft-lindem-ospfv3-dest-filter-03.txt

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of [section 3 of RFC 3667](#). By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she become aware will be disclosed, in accordance with [RFC 3668](#).

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>.

This Internet-Draft will expire on April 1, 2005.

Copyright Notice

Copyright (C) The Internet Society (2004).

## Abstract

OSPFv2 has been criticized for its vulnerability to Denial of Service (DOS) attacks. With OSPFv3, it is a simple matter to filter on the destination address at an implementation dependent level in order to limit the scope of DOS attacks to directly attached routers. Unlike hop limit checking mechanisms, it is compatible with the existing OSPFv3 behavior. However, this level of protection will preclude the

deployment of virtual links in topologies where the filtering is applied.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">3</a>
<a href="#">2.</a>	Proposed Solution . . . . .	<a href="#">4</a>
<a href="#">2.1</a>	Virtual Links . . . . .	<a href="#">4</a>
<a href="#">2.2</a>	Tunnels . . . . .	<a href="#">4</a>
<a href="#">3.</a>	Implementation Placement and Granularity of Filter . . . . .	<a href="#">6</a>
<a href="#">4.</a>	RIPng Applicability . . . . .	<a href="#">7</a>
<a href="#">5.</a>	Compatibility . . . . .	<a href="#">8</a>
<a href="#">6.</a>	Security Considerations . . . . .	<a href="#">9</a>
<a href="#">7.</a>	IANA Considerations . . . . .	<a href="#">10</a>
<a href="#">8.</a>	References . . . . .	<a href="#">11</a>
<a href="#">8.1</a>	Normative References . . . . .	<a href="#">11</a>
<a href="#">8.2</a>	Informative References . . . . .	<a href="#">11</a>
	Authors' Addresses . . . . .	<a href="#">12</a>
<a href="#">A.</a>	Acknowledgments . . . . .	<a href="#">13</a>
	Intellectual Property and Copyright Statements . . . . .	<a href="#">14</a>

## 1. Introduction

OSPFv2 [[OSPFv2](#)] and OSPFv3 [[OSPFv3](#)] both have been criticised for their vulnerability to Denial of Service attacks [[VULNER](#)]. Both support cryptographic authentication to prevent an attacker from being able to spoof an OSPFv2 or OSPFv3 packet ([[OSPFv2](#)] and [[AUTHv3](#)]). However, in many cases the MD5 or IPSEC protection actually exacerbates the attack due to the computational overhead involved. For OSPFv3, this document proposes limiting accepted OSPFv3 packets to those that are not routable. Doing so allows these packets to be filtered at a low level for a relatively small computational cost.

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 [[RFC2119](#)].

## [2.](#) Proposed Solution

In order to limit the vulnerability to DOS attacks to directly attached routers, OSPFv3 packets are only accepted if the destination address in the packet header is a link-local unicast address or

link-local scoped multicast address. Both these address types are never forwarded more than one hop. Unlike hop limit checking mechanisms [\[GTSM\]](#), this technique is fully backward compatible with the OSPFv3 which doesn't specify that OSPFv3 packets be sent with a hop limit of 255. The only hop limit specification is that the link-scoped multicast packets are sent with a hop limit of 1. Hence, this mechanism can be deployed on one OSPFv3 router at a time.

In order to make the checking simple and low cost, this document suggests checking the first two octets of the IPv6 destination address for a valid link local unicast or link-local scoped multicast address. Based on the IPv6 Address Architecture [\[IP6ADDR\]](#), this would equate to:

```
if (((first-two-octets & 0xffc0) != 0xfe80) &&
    ((first-two-octets & 0xff0f) != 0xff02)) {
    drop the packet;
}
```

Alternately, an implementation may also check the multicast address flags to assure they are 0x0 since the OSPFv3 specification explicitly uses multicast addresses ff02::5 (AllSPFRouters) and ff02::6 (AllDRouters) [\[OSPFv3\]](#).

```
if (((first-two-octets & 0xffc0) != 0xfe80) &&
    ((first-two-octets & 0xffff) != 0xff02)) {
    drop the packet;
}
```

## [2.1](#) Virtual Links

Virtual links make use of a global IPv6 unicast destination address. Hence, the proposed destination address filter and virtual links are incompatible. Depending on the granularity of the filtering, virtual links may still be used (See [Section 3.0](#)).

## [2.2](#) Tunnels

In order to support OSPF over tunnels, e.g. GRE [\[GRE\]](#), it is necessary for the destination filter to be applied after OSPF packets are delivered to the tunnel endpoint and decapsulated. Furthermore, the encapsulated OSPFv3 packet's destination address should be

Internet-Draft

OSPFv3 Destination Address Filter

October 2004

AlllSPFRouters (FF02::5).

### 3. Implementation Placement and Granularity of Filter

The placement and granularity of the destination address filter is an engineering decision that must be made for each implementation. Obviously, the sooner it is done after packet reception the less resource that is consumed processing packets that will be dropped. However, since the checking has to be confined to OSPFv3 packets that are delivered locally it may be easier to delay the checking until the packets have been identified as such. A convenient place in an implementation using the BSD socket model [[SOCKET](#)] is the point at which an inbound packet is added to the OSPFv3 socket.

The granularity of the check will limit the usage of virtual links at the granularity which it is applied. For example, if it is applied at the BSD socket level, virtual links may not be used by any OSPF instance utilizing that socket. Alternately, additional configuration and checking could be added at the socket level so that the destination filter is only applied to certain instances, areas, or interfaces. Implementations will need to balance their market requirements for virtual link deployment. In any case, the use of virtual link SHOULD be allowed either by configuration or the filter should be automatically disabled when a virtual link is configured.

#### [4.](#) RIPng Applicability

The destination filter described herein is also applicable to RIPng [RIPNG]. The filter simply needs to be applied to UDP port 521. In RIPng there is no concept of a virtual link and no requirement to send to IPv6 global addresses.

## [5.](#) Compatibility

All mechanisms described in this document are backward-compatible with standard OSPFv3 implementations [[OSPFv3](#)].



## 6. Security Considerations

This document recommends a mechanism that can be used to limit OSPFv3 Denial of Service (DOS) attacks to directly attached networks. Hence, the entire document deals with security.

## [7.](#) IANA Considerations

This document does not require any IANA assignments or action.

## [8.](#) References

### [8.1](#) Normative References

- [IP6ADDR] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 3513](#), April 2003.
- [OSPFv2] Moy, J., "OSPF Version 2", [RFC 2328](#), April 1998.
- [OSPFv3] Moy, J., Ferguson, D. and R. Colton, "OSPF for IPv6", [RFC 2740](#), December 1999.
- [RFC2119] Bradner, S., "Key words for use in RFC's to Indicate Requirement Levels", [RFC 2328](#), March 1977.
- [RIPng] Malkin, G. and R. Minnear, "RIPng for IPv6", [RFC 2080](#), January 1997.
- [SOCKET] Gilligan, B., Thomson, S., Bound, J. and J. McCann, "Basic Socket Interface Extensions for IPv6", [RFC 3493](#), February 2003.

## [8.2](#) Informative References

- [AUTHv3] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", [draft-ietf-ospf-ospfv3-auth-04.txt](#) (work in progress).
- [GRE] Farinacci, D., Li, T., Hanks, S., Meyer, D. and P. Traina, "Generic Routing Encapsulation (GRE)", [RFC 2784](#), March 2000.
- [GTSM] Gill, V., Heasley, J. and D. Meyer, "The Generalized TTL Security Mechanism (GTSM)", [draft-ietf-rtgwg-rfc3682bis-02.txt](#) (work in progress).
- [VULNER] Jones, E. and O. Le Moigne, "OSPF Security Vulnerabilities Analysis", [draft-ietf-rpsec-ospf-vuln-00.txt](#) (work in progress).

### Authors' Addresses

Acee Lindem  
Redback Networks  
102 Carric Bend Court  
Cary, NC 27519  
USA

E-Mail: [acee@redback.com](mailto:acee@redback.com)

Anand Oswal  
Cisco Systems  
225 West Tasman Drive  
San Jose, CA 95134  
USA

E-Mail: [sina@cisco.com](mailto:sina@cisco.com)

## [Appendix A](#). Acknowledgments

The authors wish to acknowledge Mukesh Gupta, Venkata Naidu, Enke Chen, and George Apostolopoulos for their review and comments.

The RFC text was produced using Marshall Rose's `xml2rfc` tool.

## Intellectual Property Statement

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](mailto:ietf-ipr@ietf.org).

## Disclaimer of Validity

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.

#### 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.

#### Acknowledgment

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