

Network Working Group  
Internet Draft

Acee Lindem (Redback Networks)  
Anand Oswal (Redback Networks)

Expiration Date: Sept 2004

File name: [draft-lindem-ospfv3-dest-filter-01.txt](#)

Apr 2004

OSPFv3 Destination Address Filter  
draft-lindem-ospfv3-dest-filter-01.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

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>	Overview .....	<a href="#">2</a>
<a href="#">2</a>	Proposed Solution .....	<a href="#">2</a>
<a href="#">2.1</a>	Virtual Links .....	<a href="#">3</a>
<a href="#">2.2</a>	Tunnels .....	<a href="#">3</a>
<a href="#">3</a>	Implementation and Granularity of Filter .....	<a href="#">3</a>
<a href="#">4</a>	RIPng Applicability .....	<a href="#">3</a>
<a href="#">5</a>	Security Considerations .....	<a href="#">4</a>
<a href="#">6</a>	Intellectual Property .....	<a href="#">4</a>

<a href="#">7</a>	Normative References .....	<a href="#">4</a>
<a href="#">8</a>	Informative References .....	<a href="#">5</a>
<a href="#">9</a>	Acknowledgments .....	<a href="#">5</a>
<a href="#">10</a>	Authors' Addresses .....	<a href="#">5</a>

## [1.](#) Overview

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 [[RFC-2119](#)].

## [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 [[ADDR-ARCH](#)], 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 AllSPFRouters (FF02::5).

## [3.0](#) 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.

Lindem, Oswal

[Page 3]

---

Internet Draft      OSPFv3 Destination Address Filter

Apr 2004

#### [5.](#) 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.

#### [6.](#) Intellectual Property

The IETF takes no position regarding the validity or scope of any intellectual property 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; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication 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 implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive

Director.

## 7. Normative References

- [RFC-2119] Bradner, S., "Key words for use in RFC's to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1977.
- [OSPFv2] Moy, J., "OSPF Version 2", [RFC 2328](#), April 1998.
- [OSPFv3] Coltun, R., Ferguson, D. and Moy, J., "OSPF for IPv6", [RFC 2740](#), December 1999.
- [ADDR-ARCH] Hinden, R. and Deering, S., "IP Version 6 Addressing Architecture", [RFC 3513](#), April 2003.
- [SOCKET] Gilligan, B., Thomson, S., Bound, J., McCann, J. and Stevens, R., "Basic Socket Interface Extensions for IPv6", [RFC 3493](#), February 2003.
- [RIPng] Malkin, G. and Minnear, R., "RIPng for IPv6", [RFC 2080](#), January 1997.

Lindem, Oswal

[Page 4]

---

Internet Draft      OSPFv3 Destination Address Filter

Apr 2004

## 8. Informative References

- [GTSM] Gill, V., Heasley, J., and Meyer, D., "The Generalized TTL Security Mechanism (GTSM) drdraft-gill-gtsh-04.txt, Work in progress.
- [AUTHv3] Gupta, M. and Melam, N., "Authentication/Confidentiality for OSPFv3", [draft-ietf-ospf-ospfv3-auth-04.txt](#), Work in progress.
- [VULNER] Jones, E. and Le Moigne, O., "OSPF Security Vulnerabilities Analysis", [draft-jones-ospf-vuln-01.txt](#), Work in progress.
- [GRE] Farinacci, D., Li, T., Hanks, S., Meyer, D. and Traina, P., "Generic Routing Encapsulation (GRE)", [RFC 2784](#), March 2000.

## 9. Acknowledgments

The authors wish to acknowledge Enke Chen and George Apostolopoulos for their thorough review.

10. Authors' Addresses

Acee Lindem  
Redback Networks  
102 Carric Bend Court  
Cary, NC 27519  
Email: [acee@redback.com](mailto:acee@redback.com)

Anand Oswal  
Redback Networks  
300 Holger  
San Jose, CA  
Email: [aoswal@redback.com](mailto:aoswal@redback.com)

Lindem, Oswal

[Page 5]