

OSPF Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 5, 2012

Y. Yang
Cisco Systems
A. Retana
Hewlett-Packard Co.
A. Roy
Cisco Systems
February 2, 2012

Hiding Transit-only Networks in OSPF
<[draft-ietf-ospf-prefix-hiding-02.txt](#)>

Abstract

A transit-only network is defined as a network connecting routers only. In OSPF, transit-only networks are usually configured with routable IP addresses, which are advertised in Link State Advertisements (LSAs) but not needed for data traffic. In addition, remote attacks can be launched against routers by sending packets to these transit-only networks. This document presents a mechanism to hide transit-only networks to speed up network convergence and minimize remote attack vulnerability.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on August 5, 2012.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/bcp78) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Requirements notation	3
2.	Hiding IPv4 Transit-only Networks in OSPFv2	4
2.1.	Point-to-Point Networks	4
2.1.1.	Advertising Point-to-Point Networks	4
2.1.2.	Hiding Point-to-Point Networks	5
2.2.	Broadcast Networks	5
2.2.1.	Advertising Broadcast Networks	5
2.2.2.	Hiding Broadcast Networks	6
2.2.2.1.	Sending Network-LSA	6
2.2.2.2.	Receiving Network-LSA	6
2.2.2.2.1.	Backward Compatibility	6
2.3.	Non-Broadcast Networks	7
2.3.1.	NBMA	7
2.3.2.	Point-to-MultiPoint	7
2.3.2.1.	Advertising Point-to-MultiPoint Networks	8
2.3.2.2.	Hiding Point-to-MultiPoint Networks	8
3.	Hiding IPv6 Transit-only Networks in OSPFv3	9
4.	Hiding AF Enabled Transit-only Networks in OSPFv3	9
5.	Operational Considerations	10
6.	Security Considerations	10
7.	IANA Considerations	10
8.	References	10
8.1.	Normative References	10
8.2.	Informative References	11
Appendix A.	Acknowledgments	11
	Authors' Addresses	11

1. Introduction

A transit-only network is defined as a network connecting routers only. In OSPF, transit-only networks are usually configured with routable IP addresses, which are advertised in LSAs but not needed for data traffic. In addition, remote attacks can be launched against routers by sending packets to these transit-only networks. This document presents a mechanism to hide transit-only networks to speed up network convergence and minimize remote attack vulnerability.

Hiding transit-only networks will not impact reachability to the end hosts.

1.1. Requirements notation

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 [KEYWORD].

2. Hiding IPv4 Transit-only Networks in OSPFv2

In [OSPFv2], networks are classified as point-to-point, broadcast, or non-broadcast. In the following sections, we will review how these OSPF networks are being advertised and discuss how to hide them consequently.

2.1. Point-to-Point Networks

A point-to-point network joins a single pair of routers. Figure 1 shows a point-to-point network connecting routers RT1 and RT2.




Figure 1 Physical point-to-point network

2.1.1. Advertising Point-to-Point Networks

For each numbered point-to-point network, a router has 2 link descriptions in its router-LSA, one Type 1 link (point-to-point) regarding the neighboring router, and one Type 3 link (stub) regarding the assigned IPv4 address.

An example of router-LSA originated by RT1 would look like

```

LS age = 0                      ;newly (re)originated
LS type = 1                      ;router-LSA
Link State ID = 1.1.1.1          ;RT1's Router ID
Advertising Router = 1.1.1.1     ;RT1's Router ID
#links = 2
  Link ID = 2.2.2.2              ;RT2's Router ID
  Link Data = 10.1.1.1           ;Interface IP address
  Type = 1                      ;connects to RT2
  Metric = 10

  Link ID= 10.1.1.0              ;Interface IP address
  Link Data = 255.255.255.252    ;Subnet's mask
  Type = 3                      ;Connects to stub network
  Metric = 10

```

The Type 1 link will be used for SPF calculation while the Type 3 link will be used for Routing Information Base (RIB) installation.

2.1.2. Hiding Point-to-Point Networks

To hide a transit-only point-to-point network, the Type 3 link MUST be omitted from the router-LSA.

An example of router-LSA originated by RT1, hiding the point-to-point network depicted in Figure 1, would look like

```
LS age = 0                ;newly (re)originated
LS type = 1                ;router-LSA
Link State ID = 1.1.1.1    ;RT1's Router ID
Advertising Router = 1.1.1.1 ;RT1's Router ID
#links = 1
  Link ID = 2.2.2.2        ;RT2's Router ID
  Link Data = 10.1.1.1     ;Interface IP address
  Type = 1                 ;connects to RT2
  Metric = 10
```

2.2. Broadcast Networks

A broadcast networks joins many (more than two) routers, and supports the capability to address a single physical message to all of the attached routers. Figure 2 shows a broadcast network connecting router RT3, RT4, and RT5.




Figure 2 Broadcast network

2.2.1. Advertising Broadcast Networks

For each broadcast network, a designated router (DR) describes it in its network-LSA. Assuming RT3 is elected as the DR in Figure 2, an example of the network-LSA originated by RT3 would look like


```
LS age = 0                ;newly (re)originated
LS type = 2               ;network-LSA
Link State ID = 10.2.2.3  ;IP address of the DR (RT3)
Advertising Router = 3.3.3.3 ;RT3's Router ID
Network Mask = 255.255.255.0
    Attached Router = 3.3.3.3 ;Router ID
    Attached Router = 4.4.4.4 ;Router ID
    Attached Router = 5.5.5.5 ;Router ID
```

OSPF obtains the IP network number from the combination of the Link State ID and the Network Mask. In addition, the Link State ID is also being used for 2-way connectivity check.

2.2.2. Hiding Broadcast Networks

2.2.2.1. Sending Network-LSA

To hide a transit-only broadcast network, a special network mask value 255.255.255.255 MUST be used in the network-LSA. While a broadcast network connects more than routers, using 255.255.255.255 will not hide an access broadcast network accidentally.

As there is no change of the Link State ID, the 2-way connectivity check would proceed normally.

An example of network-LSA originated by RT3, hiding the broadcast network depicted in Figure 2, would look like

```
LS age = 0                ;newly (re)originated
LS type = 2               ;network-LSA
Link State ID = 10.2.2.3  ;IP address of the DR (RT3)
Advertising Router = 3.3.3.3 ;RT3's Router ID
Network Mask = 255.255.255.255 ;special subnet mask
    Attached Router = 3.3.3.3 ;Router ID
    Attached Router = 4.4.4.4 ;Router ID
    Attached Router = 5.5.5.5 ;Router ID
```

2.2.2.2. Receiving Network-LSA

It's RECOMMENDED that all routers in an area be upgraded at the same time to process the modified network-LSA correctly and consistently.

When a router receives a network-LSA, it MUST check the 2-way connectivity as normal. However, if the network mask in the network-LSA is 255.255.255.255, the router MUST NOT install the route in the RIB.

2.2.2.2.1. Backward Compatibility

When a not-yet-upgraded router receives a modified network-LSA, as specified in [section 2.2.2.1](#), a host route to the originating DR will be installed. This is not ideal but better than the current result, which exposes the whole subnet.

In a partial deployment scenario, upgraded routers and not-yet-upgraded routers may coexist. The former do not have the host routes aforementioned, while the latter do have. Such inconsistencies create routing black holes, which should normally be avoided. In this case, however, as packets destined for the transit-only networks are dropped somewhere in the network, the black holes actually help DRs defend from the remote attacks.

In summary, the modification of the network-LSA, as specified in [section 2.2.2.1](#), is backward compatible with the current specification of [OSPFv2], even in a partial deployment scenario.

[2.3.](#) Non-Broadcast Networks

A non-broadcast networks joins many (more than two) routers, but does NOT support the capability to address a single physical message to all of the attached routers. As mentioned in [OSPFv2], OSPF runs in one of two modes over non-broadcast networks: Non-Broadcast Multi-Access (NBMA) or Point-to-MultiPoint.

[2.3.1.](#) NBMA

In NBMA mode, OSPF emulates operation over a broadcast network: a Designated Router is elected for the NBMA network, and the Designated Router originates an LSA for the network.

To hide a NBMA transit-only network, OSPF adopts the same modification over the broadcast transit-only network, as defined in [section 2.2.2](#).

[2.3.2.](#) Point-to-MultiPoint

In point-to-MultiPoint mode, OSPF treats the non-broadcast network as a collection of point-to-point links.

Figure 3 shows a non-broadcast network connecting router RT6, RT7, RT8, and RT9. In this network, all routers can communicate directly, except for routers RT7 and RT8.




Figure 3 Non-Broadcast network

[2.3.2.1. Advertising Point-to-MultiPoint Networks](#)

For a point-to-multipoint network, a router has multiple link descriptions in its router-LSA, one Type 1 link (point-to-point) for EACH directly communicable router, and one Type 3 link (stub) advertising its interface IPv4 address with a subnet mask of 255.255.255.255.

An example of router-LSA originated by RT7 would look like

```

LS age = 0                ;newly (re)originated
LS type = 1               ;router-LSA
Link State ID = 7.7.7.7   ;RT7's Router ID
Advertising Router = 7.7.7.7 ;RT7's Router ID
#links = 3
  Link ID = 6.6.6.6       ;RT6's Router ID
  Link Data = 10.3.3.7    ;Interface IP address
  Type = 1               ;connects to RT6
  Metric = 10

  Link ID = 9.9.9.9       ;RT9's Router ID
  Link Data = 10.3.3.7    ;Interface IP address
  Type = 1               ;connects to RT9
  Metric = 10

  Link ID= 10.3.3.7       ;Interface IP address
  Link Data = 255.255.255.255 ;Subnet's mask
  Type = 3               ;Connects to stub network
  Metric = 0
  
```

[2.3.2.2. Hiding Point-to-MultiPoint Networks](#)

To hide a transit-only point-to-multipoint network, the Type 3 link MUST be omitted from the router-LSA.

An example of router-LSA originated by RT7, hiding the point-to-point network depicted in Figure 3, would look like

```
LS age = 0                      ;newly (re)originated
LS type = 1                     ;router-LSA
Link State ID = 7.7.7.7         ;RT7's Router ID
Advertising Router = 7.7.7.7    ;RT7's Router ID
#links = 2
  Link ID = 6.6.6.6             ;RT6's Router ID
  Link Data = 10.3.3.7          ;Interface IP address
  Type = 1                     ;connects to RT6
  Metric = 10

  Link ID = 9.9.9.9             ;RT9's Router ID
  Link Data = 10.3.3.7          ;Interface IP address
  Type = 1                     ;connects to RT9
  Metric = 10
```

3. Hiding IPv6 Transit-only Networks in OSPFv3

In [[OSPFv3](#)], addressing semantics have been removed from the OSPF protocol packets and the main LSA types, leaving a network-protocol-independent core.

More specifically, router-LSAs and network-LSAs no longer contain network addresses, but simply express topology information. A new LSA called the intra-area-prefix-LSA has been introduced. This LSA carries all IPv6 prefix information that in [[OSPFv2](#)] is included in router-LSAs and network-LSAs.

Such changes simplify the process to hide the IPv6 addresses of the transit-only networks in [[OSPFv3](#)] -- simply omitting the correspondent IPv6 unicast prefixes from the intra-area-prefix-LSA will hide these prefixes.

4. Hiding AF Enabled Transit-only Networks in OSPFv3

[OSPF-AF] supports multiple Address Families (AFs) by mapping each AF to a separate Instance ID and OSPFv3 instance.

In the meantime, each prefix advertised in OSPFv3 has a prefix Length field [OSPFV3], which facilitates advertising prefixes of different lengths in different AFs. The existing LSAs defined in OSPFv3 are used for prefix advertising and there is no need to define new LSAs.

In other words, intra-area-prefix-LSAs are still being used to

advertise the attached networks, and same method explained in [section 3](#) can also be used to hide those AF enabled transit-only networks.

5. Operational Considerations

By eliminating the ability to reach transit-only networks, the ability to manage these interfaces may be reduced. In order to not reduce the functionality and capability of the overall network, it is recommended that extensions such as [RFC5837](#) be also implemented.

6. Security Considerations

One motivation for this document is to reduce remote attack vulnerability by hiding transit-only networks. The result should then be that fewer OSPF core networks will be exposed to unauthorized access.

While the steps described in this document are meant to be applied to transit-only networks ONLY, they could be used to hide other networks as well. It is expected that the same care that users put on the configuration of other routing protocol parameters is used in the configuration of this extension.

7. IANA Considerations

No actions are required from IANA as result of the publication of this document.

8. References

8.1. Normative References

- [KEYWORD] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [OSPFv2] Moy, J., "OSPF Version 2", [RFC 2328](#), April 1998.
- [OSPFv3] Coltun, R., Ferguson, D., Moy, J., and A. Lindem , "OSPF for IPv6", [RFC 5340](#), July 2008.
- [OSPF-AF] Lindem, A., Mirtorabi, S., Roy, A., Barnes, M., and R. Aggarwal, "Support of Address Families in OSPFv3", [RFC5838](#), April 2010.

8.2. Informative References

[RFC5837] Atlas, A., Bonica, R., Pignataro, C., Shen, N., and JR. Rivers, "Extending ICMP for Interface and Next-Hop Identification", [RFC5837](#), April 2010.

Appendix A. Acknowledgments

The draft text was produced using Stefan Santesson's NroffEdit application.

The idea of using a special subnet mask to hide broadcast networks in OSPF was originally introduced in the US patent "Apparatus and method to hide transit only multi-access networks in OSPF" (patent number: 7,929,524), by Yi Yang, Alvaro Retana, James Ng, Abhay Roy, Alfred Lindem, Sina Mirtorabi, Timothy Gage, and Khalid Raza.

The authors would like to thank Acee Lindem for his feedback on the document.

Authors' Addresses

Yi Yang
Cisco Systems
7025 Kit Creek Road
RTP, NC 27709
USA

Email: yiya@cisco.com

Alvaro Retana
Hewlett-Packard Co.
2610 Wycliff Road
Raleigh, NC 27607
USA

Email: alvaro.retana@hp.com

Abhay Roy
Cisco Systems
225 West Tasman Drive
San Jose, CA 95134
USA

Email: akr@cisco.com