

Internet Engineering Task Force (IETF)
Request for Comments: 6860
Updates: 2328, 5340
Category: Standards Track
ISSN: 2070-1721

Y. Yang
A. Retana
A. Roy
Cisco Systems, Inc.
January 2013

Hiding Transit-Only Networks in OSPF

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 are 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 reduce vulnerability to remote attacks.

In the context of this document, 'hiding' implies that the prefixes are not installed in the routing tables on OSPF routers. In some cases, IP addresses may still be visible when using OSPFv2.

This document updates RFCs 2328 and 5340.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6860>.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 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	3
2.1. Point-to-Point Networks	3
2.1.1. Advertising Point-to-Point Networks	4
2.1.2. Hiding Point-to-Point Networks	4
2.2. Broadcast Networks	5
2.2.1. Advertising Broadcast Networks	5
2.2.2. Hiding Broadcast Networks	5
2.2.2.1. Sending Network-LSA	5
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 ...	7
2.3.2.2. Hiding Point-to-Multipoint Networks	8
3. Hiding IPv6 Transit-Only Networks in OSPFv3	9
3.1. Hiding AF-Enabled Transit-Only Networks in OSPFv3	9
4. Operational Considerations	9
4.1. Forwarding Address	10
4.2. Virtual Links	10
4.3. Unnumbered Interfaces	10
5. Security Considerations	11
6. Acknowledgments	11
7. References	12
7.1. Normative References	12
7.2. Informative References	12

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 reduce vulnerability to remote attacks.

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

In the context of this document, 'hiding' implies that the prefixes are not installed in the routing tables on OSPF routers. In [OSPFv2], the IPv4 interface addresses are still visible in the Router-LSA links and the network-LSA Link-State ID (LSID). In [OSPFv3], the router-LSAs and network-LSAs do not contain IPv6 addresses and are not visible.

This document updates [OSPFv2] and [OSPFv3] by specifying a mechanism that can be used to hide transit-only networks.

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.

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.

```

+---+.1      198.51.100.0/30      .2+---+
|RT1|-----|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 two link descriptions in its router-LSA: one Type 1 link (point-to-point) describing the neighboring router, and one Type 3 link (stub) describing the assigned IPv4 subnet.

An example of a router-LSA originated by RT1 would look like the following:

```
LS age = 0                ;newly (re-)originated
LS type = 1               ;router-LSA
Link State ID = 192.0.2.1 ;RT1's Router ID
Advertising Router = 192.0.2.1 ;RT1's Router ID
#links = 2
  Link ID = 192.0.2.2      ;RT2's Router ID
  Link Data = 198.51.100.1 ;Interface IP address
  Type = 1                 ;connects to RT2
  Metric = 10

  Link ID= 198.51.100.0    ;IP network/subnet number
  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 to install a route to the corresponding subnet in the Routing Information Base (RIB).

2.1.2. Hiding Point-to-Point Networks

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

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

```
LS age = 0                ;newly (re-)originated
LS type = 1               ;router-LSA
Link State ID = 192.0.2.1 ;RT1's Router ID
Advertising Router = 192.0.2.1 ;RT1's Router ID
#links = 1
  Link ID = 192.0.2.2      ;RT2's Router ID
  Link Data = 198.51.100.1 ;Interface IP address
  Type = 1                 ;connects to RT2
  Metric = 10
```

2.2. Broadcast Networks

A broadcast network 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 routers RT3, RT4, and RT5.




Figure 2. Broadcast Network

2.2.1. Advertising Broadcast Networks

A Designated Router (DR) describes a broadcast network in a network-LSA. Assuming that 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 = 198.51.100.3 ;IP address of the DR (RT3)
Advertising Router = 192.0.2.3 ;RT3's Router ID
Network Mask = 255.255.255.0
    Attached Router = 192.0.2.3 ;RT3's Router ID
    Attached Router = 192.0.2.4 ;RT4's Router ID
    Attached Router = 192.0.2.5 ;RT5's 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 the two-way connectivity check.

2.2.2. Hiding Broadcast Networks

2.2.2.1. Sending Network-LSA

A special subnet mask value of 255.255.255.255 MUST be used in the network-LSA to hide a transit-only broadcast network. 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 two-way connectivity check would proceed normally.

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

LS age = 0	;newly (re-)originated
LS type = 2	;network-LSA
Link State ID = 198.51.100.3	;IP address of the DR (RT3)
Advertising Router = 192.0.2.3	;RT3's Router ID
Network Mask = 255.255.255.255	;special subnet mask
Attached Router = 192.0.2.3	;RT3's Router ID
Attached Router = 192.0.2.4	;RT4's Router ID
Attached Router = 192.0.2.5	;RT5's Router ID

2.2.2.2. Receiving Network-LSA

It is 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 calculate the routing table normally [OSPFv2]. 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 router that has not yet been upgraded 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 it is better than the current result, which exposes the whole subnet.

In a partial-deployment scenario, upgraded routers and routers that have not yet been upgraded may coexist. The former do not install the host route to the DR's interface, while the latter install it. 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 the DRs defend themselves from 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 network 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 an NBMA transit-only network, OSPF adopts the same modification as that used over the broadcast transit-only network (see 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 routers 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 a router-LSA originated by RT7 would look like the following:

```

LS age = 0                ;newly (re-)originated
LS type = 1               ;router-LSA
Link State ID = 192.0.2.7 ;RT7's Router ID
Advertising Router = 192.0.2.7 ;RT7's Router ID
#links = 3
  Link ID = 192.0.2.6      ;RT6's Router ID
  Link Data = 198.51.100.7 ;Interface IP address
  Type = 1                 ;connects to RT6
  Metric = 10

  Link ID = 192.0.2.9      ;RT9's Router ID
  Link Data = 198.51.100.7 ;Interface IP address
  Type = 1                 ;connects to RT9
  Metric = 10

  Link ID = 198.51.100.7    ;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 is omitted from the router-LSA.

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

```

LS age = 0                ;newly (re-)originated
LS type = 1               ;router-LSA
Link State ID = 192.0.2.7 ;RT7's Router ID
Advertising Router = 192.0.2.7 ;RT7's Router ID
#links = 2
  Link ID = 192.0.2.6      ;RT6's Router ID
  Link Data = 198.51.100.7 ;Interface IP address
  Type = 1                 ;connects to RT6
  Metric = 10

  Link ID = 192.0.2.9      ;RT9's Router ID
  Link Data = 198.51.100.7 ;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. Instead, two new LSA types, link-LSA and intra-area-prefix-LSA, have been introduced. A link-LSA associates a list of IPv6 addresses to a link and has local-link flooding scope, and an intra-area-prefix-LSA either associates a list of IPv6 addresses with a router by referencing a router-LSA or associates a list of IPv6 addresses with a broadcast/NBMA network by referencing a network-LSA. In the latter case, the prefixes in the link-LSAs from adjacent neighbors are copied into the intra-area-prefix-LSA by the Designated Router.

To hide a transit-only network in [OSPFv3], the IPv6 address prefixes are omitted from the router-LSA. Consequently, when a Designated Router builds an intra-area-prefix-LSA referencing a network-LSA, these IPv6 address prefixes will be omitted.

In addition, when a router builds an intra-area-prefix-LSA that is referencing a router-LSA, the associated IPv6 address prefixes from the transit-only network MUST also be omitted from the intra-area-prefix-LSA.

3.1. 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, as link-LSAs and intra-area-prefix-LSAs are still being used, the same mechanism explained in Section 3 can be used to hide those AF-enabled transit-only networks as well.

4. Operational Considerations

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

Note that the extension defined in [UNNUMBERED] may provide the user with the IP address of an interface. If that address was hidden, as specified in this document, then even though the address is assigned to the interface, it will not be reachable.

4.1. Forwarding Address

A non-zero forwarding address can be advertised in AS-external-LSAs and Not-So-Stubby Area LSAs (NSSA-LSAs) [NSSA] to achieve optimal routing to Autonomous System (AS) external routes. The matching routing table entry for the forwarding address must exist to facilitate the SPF calculation.

In other words, when prefix-hiding is configured on the next-hop interface, the next-hop address MUST NOT be advertised as a forwarding address.

Consequently, sub-optimal routing to these AS external routes may exist when prefix-hiding is configured.

4.2. Virtual Links

Virtual links are used to connect physically separate components of the backbone. The virtual link's viability is determined by the existence of an intra-area path between two endpoints. The matching routing table entries of the endpoints must exist to ensure the virtual link's operation.

In other words, if prefix-hiding is configured on an interface, the virtual link endpoint MUST NOT use that interface's IP address as the virtual interface's IP address.

4.3. Unnumbered Interfaces

Note that no host route is generated for, and no IP packets can be addressed to, interfaces to unnumbered point-to-point networks [OSPFv2]. In other words, these addresses are already hidden.

However, for manageability purposes, it may be common practice to manually include the numbered interface (for example, a loopback interface to which the unnumbered interface points) in routing updates. If needed, the numbered interface's address can be hidden by using the mechanisms described in this document or by simply not advertising it.

Before deciding to hide (or suppress the advertisement of) a numbered interface, it is very important to consider other uses that interface may have. Examples of common uses may include virtual link endpoint, inter-domain routing peering point, etc. In other words, it may not be possible to hide the address associated to an unnumbered interface due to other applications in the network.

5. Security Considerations

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

The mechanisms described above result in reachability information from transit-only networks not being installed in the routers' forwarding tables. The effect is that even if the address of a transit-only network is known, the forwarding information is not present in the routers to reach the destination. Also, in some cases, the address information is completely omitted from the LSA.

Some information in the LSA (such as the OSPF Router ID) cannot be omitted. Even though the Router ID may be taken from an IPv4 address on the router, the configuration can be easily changed. Note again that having an address doesn't guarantee reachability if the information is hidden from the forwarding tables.

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

6. Acknowledgments

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, Shraddha Hegde, Rajesh Shetty, Marek Karasek, Michael Barnes, Paul Wells, Adrian Farrel, and Stephen Farrell for their feedback on the document.

7. References

7.1. Normative References

- [KEYWORD] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [NSSA] Murphy, P., "The OSPF Not-So-Stubby Area (NSSA) Option", RFC 3101, January 2003.
- [OSPFv2] Moy, J., "OSPF Version 2", STD 54, 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., Ed., Mirtorabi, S., Roy, A., Barnes, M., and R. Aggarwal, "Support of Address Families in OSPFv3", RFC 5838, April 2010.

7.2. Informative References

- [UNNUMBERED] Atlas, A., Ed., Bonica, R., Ed., Pignataro, C., Ed., Shen, N., and JR. Rivers, "Extending ICMP for Interface and Next-Hop Identification", RFC 5837, April 2010.

Authors' Addresses

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

EMail: yiya@cisco.com

Alvaro Retana
Cisco Systems, Inc.
7025 Kit Creek Road
RTP, NC 27709
USA

EMail: aretana@cisco.com

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

EMail: akr@cisco.com

