

Network Working Group
Request for Comments: 5185
Category: Standards Track

S. Mirtorabi
Nuova Systems
P. Psenak
Cisco Systems
A. Lindem, Ed.
A. Oswal
Redback Networks
May 2008

OSPF Multi-Area Adjacency

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This document describes an extension to the Open Shortest Path First (OSPF) protocol to allow a single physical link to be shared by multiple areas. This is necessary to allow the link to be considered an intra-area link in multiple areas. This would create an intra-area path in each of the corresponding areas sharing the same link.

Table of Contents


1.	Introduction	3
1.1.	Motivation	3
1.2.	Possible Solutions	3
1.3.	Proposed Solution	4
1.4.	Requirements Notation	4
2.	Functional Specifications	4
2.1.	Multi-Area Adjacency Configuration and Neighbor Discovery	4
2.2.	Multi-Area Adjacency Packet Transmission	5
2.3.	Multi-Area Adjacency Control Packet Reception Changes	5
2.4.	Interface Data Structure	6
2.5.	Interface FSM	6
2.6.	Neighbor Data Structure and Neighbor FSM	6
2.7.	Advertising Multi-Area Adjacencies	6
3.	Compatibility	7
3.1.	Adjacency Endpoint Compatibility	7
4.	OSPFv3 Applicability	7
5.	Security Considerations	7
6.	References	8
6.1.	Normative References	8
6.2.	Informative References	8
	Appendix A. Acknowledgments	9

1. Introduction

1.1. Motivation

It is often a requirement to have an Open Shortest Path First (OSPF) [OSPF] link in multiple areas. This will allow the link to be considered as an intra-area path in each area and be preferred over higher cost links. A simple example of this requirement is to use a high-speed link between two Area Border Routers (ABRs) in multiple areas.

Consider the following topology:



Multi-Link Topology

The backbone area link between R1 and R2 is a high-speed link, and it is desirable to forward Area 1's traffic between R1 and R2 over that link. In the current OSPF specification [OSPF], intra-area paths are preferred over inter-area paths. As a result, R1 will always route traffic to R4 through Area 1 over the lower speed links. R1 will even use the intra-area Area 1 path through R3 to get to Area 1 networks connected to R2. An OSPF virtual link cannot be used to solve this problem without moving the link between R1 and R2 to Area 1. This is not desirable if the physical link is, in fact, part of the network's backbone topology.

The protocol extension described herein will rectify this problem by allowing the link between R1 and R2 to be part of both the backbone area and Area 1.

1.2. Possible Solutions

For numbered interfaces, the OSPF (Open Shortest Path First) specification [OSPF] allows a separate OSPF interface to be configured in each area using a secondary address. The disadvantages of this approach are that it requires additional IP address configuration, it doesn't apply to unnumbered interfaces, and advertising secondary addresses will result in a larger overall routing table.

Allowing a link with a single address to simply be configured in multiple areas would also solve the problem. However, this would result in the subnet corresponding to the interface residing in multiple areas that is contrary to the definition of an OSPF area as a collection of subnets.

Another approach is to simply allow unnumbered links to be configured in multiple areas. Section 8.2. of the OSPF specification [OSPF] already specifies that the OSPF area ID should be used to demultiplex received OSPF packets. One limitation of this approach is that multi-access networks are not supported. Although this limitation may be overcome for LAN media with support of "Point-to-Point operation over LAN in link-state routing protocols" [P2PLAN], it may not be acceptable to configure the link as unnumbered due to network management policies. Many popular network management applications individually test the path to each interface by pinging its IP address.

1.3. Proposed Solution

ABRs will simply establish multiple adjacencies belonging to different areas. Each multi-area adjacency is announced as a point-to-point link in the configured area. However, unlike numbered point-to-point links, no type 3 link is advertised for multi-area adjacencies. This point-to-point link will provide a topological path for that area. The first or primary adjacency using the link will operate and advertise the link in a manner consistent with RFC 2328 [OSPF].

1.4. 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 RFC 2119 [RFC-KEYWORDS].

2. Functional Specifications

2.1. Multi-Area Adjacency Configuration and Neighbor Discovery

Multi-area adjacencies are configured between two routers having a common interface. On point-to-point interfaces, there is no need to configure the neighbor's address since there can be only one neighbor. For all other network types, the neighbor address of each multi-area adjacency must be configured or automatically discovered via a mechanism external to OSPF.

2.2. Multi-Area Adjacency Packet Transmission

On point-to-point interfaces, OSPF control packets are sent to the AllSPFRouters address. For all other network types, OSPF control packets are unicast to the remote neighbor's IP address.

2.3. Multi-Area Adjacency Control Packet Reception Changes

Receiving protocol packets is described in Section 8.2 of [OSPF]. The text starting with the second paragraph and continuing through the third bullet beneath that paragraph is changed as follows:

Next, the OSPF packet header is verified. The fields specified in the header must match those configured for the receiving interface. If they do not, the packet should be discarded:

- o The version number field must specify protocol version 2.
- o The Area ID found in the OSPF header must be verified. If all of the following cases fail, the packet should be discarded. The Area ID specified in the header must either:
 1. Match the Area ID of the receiving interface. In this case, the packet has been sent over a single hop. Therefore, the packet's IP source address is required to be on the same network as the receiving interface. This can be verified by comparing the packet's IP source address to the interface's IP address, after masking both addresses with the interface mask. This comparison should not be performed on point-to-point networks. On point-to-point networks, the interface addresses of each end of the link are assigned independently, if they are assigned at all.
 2. Indicate a non-backbone area. In this case, the packet has been sent over a multi-area adjacency. If the area-id matches the configured area for a multi-area adjacency, the packet is accepted and is from now on associated with the multi-area adjacency for that area.
 3. Indicate the backbone. In this case, the packet has been sent over a virtual link or a multi-area adjacency.
- o For virtual links, the receiving router must be an ABR, and the Router ID specified in the packet (the source router) must be the other end of a configured virtual link. The receiving interface must also attach to the virtual link's configured transit area. If all of these checks succeed, the packet is accepted and is from now on associated with the virtual link.

- o For multi-area adjacencies, if the area-id matches the configured area for the multi-area adjacency, the packet is accepted and is from now on associated with the multi-area adjacency for that area.
- o Note that if there is a match for both a virtual link and a multi-area adjacency then this is a configuration error that should be handled at the configuration level.
- o Packets whose IP destination is AllDRouters should only be accepted if the state of the receiving interface is DR or Backup (see Section 9.1 of [OSPF]).
- o [...] The remainder of Section 8.2 of [OSPF] is unchanged.

2.4. Interface Data Structure

An OSPF interface data structure is built for each configured multi-area adjacency as specified in Section 9 of [OSPF]. The interface type will always be point-to-point.

2.5. Interface FSM

The interface Finite State Machine (FSM) will be the same as a point-to-point link irrespective of the underlying physical link.

2.6. Neighbor Data Structure and Neighbor FSM

Both the neighbor data structure and neighbor FSM are the same as for standard OSPF, specified in Section 10 of [OSPF].

2.7. Advertising Multi-Area Adjacencies

Multi-area adjacencies are announced as point-to-point links. Once the router's multi-area adjacency reaches the FULL state, it will be added as a link type 1 to the Router Link State Advertisement (LSA) with:

Link ID = Remote's Router ID

Link Data = Neighbor's IP Address or IfIndex (if the underlying interface is unnumbered).

Unlike numbered point-to-point links, no type 3 link is advertised for multi-area adjacencies.

3. Compatibility

All mechanisms described in this document are backward compatible with standard OSPF implementations [OSPF].

3.1. Adjacency Endpoint Compatibility

Since multi-area adjacencies are modeled as point-to-point links, it is only necessary for the router at the other end of the adjacency to model the adjacency as a point-to-point link. However, the network topology will be easier to represent and troubleshoot if both neighbors are symmetrically configured as multi-area adjacencies.

4. OSPFv3 Applicability

The mechanisms defined in this document also apply to OSPFv3 [OSPFV3]. As in OSPF, a multi-area adjacency is advertised as a point-to-point link in the advertising router's router-LSA. Since OSPFv3 router-LSA links are independent of addressing semantics and unambiguously identify OSPFv3 neighbors (refer to Section 3.4.3.1 of [OSPFV3]), the change to router-LSA links described in Section 2.7 is not applicable to OSPFv3. Furthermore, no prefixes corresponding to the multi-area adjacency are advertised in the router's intra-area-prefix-LSA.

A link-LSA SHOULD NOT be advertised for a multi-area adjacency. The neighbor's IPv6 link local address can be learned in other ways, e.g., it can be extracted from the IPv6 header of Hello packets received over the multi-area adjacency. The neighbor IPv6 link local address is required for the OSPFv3 route next-hop calculation on multi-access networks (refer to Section 3.8.1.1 of [OSPFV3]).

5. Security Considerations

This document does not raise any security issues that are not already covered in [OSPF] or [OSPFV3].

6. References

6.1. Normative References

- [OSPF] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [OSPFV3] Coltun, R., Ferguson, D., and J. Moy, "OSPF for IPv6", RFC 2740, December 1999.
- [RFC-KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

6.2. Informative References

- [P2PLAN] Shen, N. and A. Zinin, "Point-to-point operation over LAN in link-state routing protocols", Work in Progress.

Appendix A. Acknowledgments

The authors wish to acknowledge Pat Murphy for convincing the OSPF WG to address the requirement.

Thanks to Mitchell Erblich's for his last call review and comments.

Thanks to Padma Pillay-Esnault for her last call review and comments. Also, thanks to Padma for comments on the OSPFv3 applicability section that was last called separately.

Thanks to Nischal Seth for pointing out that the document inadvertently precluded point-to-point over LAN interfaces.

Thanks to Ben Campbell for performing the General Area review.

Thanks to Jari Arkko for comments during the IESG review.

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

Authors' Addresses

Sina Mirtorabi
Nuova Systems
3 West Plumeria Drive
San Jose, CA 95134
USA

EEmail: sina@nuovasystems.com

Peter Psenak
Cisco Systems
Apollo Business Center
Mlynske nivy 43
821 09 Bratislava
Slovakia

EEmail: ppsenak@cisco.com

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

EEmail: acee@redback.com

Anand Oswal
Redback Networks
300 Holger Way
San Jose, CA 95134
USA

EEmail: aoswal@redback.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

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.

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, THE IETF TRUST 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.

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.