Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 27, 2007

S. Mirtorabi
Force10 Networks
P. Psenak
Cisco Systems
A. Lindem (Editor)
A. Oswal
Redback Networks
March 26, 2007

OSPF Multi-Area Adjacency
draft-ietf-ospf-multi-area-adj-07.txt

Status of this Memo

Copyright Notice

Abstract

This document describes an extension to OSPF 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 to the corresponding areas sharing the same link.

Table of Contents

1. Introduction

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 RFC2119 [RFC-KEYWORDS].

1.2. Motivation

There could be a requirement to have a link in multiple areas in order to allow the link to be considered as an intra-area link in multiple areas and be preferred over high cost intra-area paths. A simple example is to use a high-speed backbone link between two Area Border Routers (ABRs) to create multi-area adjacencies belonging to different areas.

1.3. Possible Solutions

For numbered interfaces, the OSPF 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, 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 already specifies that the OSPF area ID should be used to de-multiplex 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.

1.4. Proposed Solution

ABRs will simply establish multiple adjacencies belonging to different areas. Each multi-area adjacency is announced as a point-to-point unnumbered link in the configured area. 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 consistent with RFC 2328 [OSPF].

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 physical point-to-point networks, 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 physical point-to-point networks, 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] and is changed as follow:

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:

- The version number field must specify protocol version 2.

- 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 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.

- For virtual links, the receiving router must be an area border router, 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.

- 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.

- 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.

- 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 [OSPF]).

- [...] The remainder of section 8.2 [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 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 unnumbered 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).

This will announce a topological path through the corresponding area. While advertising the neighbor's IP address in the link data isn't consistent with the unnumbered link model, it is required to eliminate ambiguity when there are parallel point-to-point 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 unnumbered 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, it will be cleaner from a deployment standpoint for both neighbors to be configured as multi-area adjacencies.

4. OSPFv3 Applicability

All mechanisms defined in this document also apply to OSPFv3 [OSPFV3]. As in OSPF, a multi-area adjacency is advertised as an unnumbered point-to-point link in the advertising router's router-LSA. Since no prefixes are associated with a multi-area adjacency, it will have no impact on the advertising 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.

5. Security Considerations

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

6. IANA Considerations

This document does not require any IANA assignments or action.

7. References

7.1. Normative References

[OSPF] Moy, J., "OSPF Version 2", 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 RFC's to Indicate Requirement Levels", RFC 2119, March 1997.

7.2. Informative References

[P2PLAN] Shen, N. and A. Zinin, "Point-to-point operation over LAN in link-state routing protocols", draft-ietf-isis-igp-p2p-over-lan-06.txt (work in progress).

Appendix A. Acknowledgments

The authors wish to acknowledge Pat Murphy for bringing focus to 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.

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

Authors' Addresses

Sina Mirtorabi
Force10 Networks
350 Holger Way
San Jose, CA 95134
USA

Email: sina@force10networks.com

Peter Psenak
Cisco Systems
Apollo Business Center
Mlynaki nivy 43
821 09 Bratislava, Sovakia
Slovakia

Email: ppsenak@cisco.com

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

Email: acee@redback.com

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

Email: aoswal@redback.com

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).