

Network Working Group
Request for Comments: 3509
Category: Informational

A. Zinin
Alcatel
A. Lindem
Redback Networks
D. Yeung
Procket Networks
April 2003

Alternative Implementations of OSPF Area Border Routers

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

Open Shortest Path First (OSPF) is a link-state intra-domain routing protocol used for routing in IP networks. Though the definition of the Area Border Router (ABR) in the OSPF specification does not require a router with multiple attached areas to have a backbone connection, it is actually necessary to provide successful routing to the inter-area and external destinations. If this requirement is not met, all traffic destined for the areas not connected to such an ABR or out of the OSPF domain, is dropped. This document describes alternative ABR behaviors implemented in Cisco and IBM routers.

1 Overview

1.1 Introduction

An OSPF routing domain can be split into several subdomains, called areas, which limit the scope of LSA flooding. According to [Ref1] a router having attachments to multiple areas is called an "area border router" (ABR). The primary function of an ABR is to provide its attached areas with Type-3 and Type-4 LSAs, which are used for describing routes and AS boundary routers (ASBRs) in other areas, as well as to perform actual inter-area routing.

1.2 Motivation

In OSPF domains the area topology is restricted so that there must be a backbone area (area 0) and all other areas must have either physical or virtual connections to the backbone. The reason for this star-like topology is that OSPF inter-area routing uses the distance-vector approach and a strict area hierarchy permits avoidance of the "counting to infinity" problem. OSPF prevents inter-area routing loops by implementing a split-horizon mechanism, allowing ABRs to inject into the backbone only Summary-LSAs derived from the intra-area routes, and limiting ABRs' SPF calculation to consider only Summary-LSAs in the backbone area's link-state database.

The last restriction leads to a problem when an ABR has no backbone connection (in OSPF, an ABR does not need to be attached to the backbone). Consider a sample OSPF domain depicted in the Figure 1.




Figure 1. ABR dropping transit traffic

In this example R1, R2, and R3 are ABRs. R1 and R2 have backbone connections, while R3 doesn't.

Following the section 12.4.1 of [Ref1], R3 will identify itself as an ABR by setting the bit B in its router-LSA. Being an ABR, R3 can only consider summary-LSAs from the backbone when building the routing table (according to section 16.2 of [Ref1]), so it will not have any inter-area routes in its routing table, but only intra-area routes from both Area 1 and Area 2. Consequently, according to section 12.4.3 of [Ref1], R3 will originate into Areas 1 and 2 only summary-LSAs covering destinations in the directly attached areas, i.e., in Area 2---the summary-LSAs for Area 1, and in Area 1---the summary-LSAs for Area 2.


At the same time, router R2, as an ABR connected to the backbone, will inject into Area 2 summary-LSAs describing the destinations in Area 0 (the backbone), Area 1 and other areas reachable through the backbone.

This results in a situation where internal router R4 calculates its routes to destinations in the backbone and areas other than Area 1 via R2. The topology of Area 2 itself can be such that the best path from R4 to R2 is via R3, so all traffic destined for the backbone and backbone-attached areas goes through R3. Router R3 in turn, having only intra-area routes for areas 1 and 2, will drop all traffic not destined for the areas directly attached to it. The same problem can occur when a backbone-connected ABR loses all of its adjacencies in the backbone---even if there are other normally functioning ABRs in the attached areas, all traffic going to the backbone (destined for it or for other areas) will be dropped.

In a standard OSPF implementation this situation can be remedied by use of Virtual Links (see section 15 of [Ref1] for more details). In this case, router R3 will have a virtual backbone connection, will form an adjacency over it, will receive all LSAs directly from a backbone-attached router (R1 or R2, or both in our example) and will install intra- or inter-area routes.

While being an unavoidable technique for repairing a partitioned backbone area, the use of virtual links in the described situation adds extra configuration headaches and system traffic overhead.

Another situation where standard ABR behavior does not provide acceptable results is when it is necessary to provide redundant connectivity to an ASBR via several different OSPF areas. This would allow a provider to aggregate all their customers connecting through a single access point into one area while still offering a redundant connection through another access point in a different area, as shown in Figure 2.



Customer Networks (CN1--CNx) Advertised as AS External or NSSA External Routes

Figure 2. Dual Homed Customer Router

This technique is already used in a number of networks including one of a major provider.

The next section describes alternative ABR behaviors, implemented in Cisco and IBM routers. The changes are in the ABR definition and inter-area route calculation. Any other parts of standard OSPF are not changed.

These solutions are targeted to the situation when an ABR has no backbone connection. They imply that a router connected to multiple areas without a backbone connection is not an ABR and should function as a router internal to every attached area. This solution emulates a situation where separate OSPF processes are run for each area and supply routes to the routing table. It remedies the situation described in the examples above by not dropping transit traffic. Note that a router following it does not function as a real border router---it doesn't originate summary-LSAs. Nevertheless such a behavior may be desirable in certain situations.

Note that the proposed solutions do not obviate the need of virtual link configuration in case an area has no physical backbone connection at all. The methods described here improve the behavior of a router connecting two or more backbone-attached areas.

2 Changes to ABR Behavior

2.1 Definitions

The following definitions will be used in this document to describe the new ABR behaviors:

Configured area:

An area is considered configured if the router has at least one interface in any state assigned to that area.

Actively Attached area:

An area is considered actively attached if the router has at least one interface in that area in the state other than Down.

Active Backbone Connection:

A router is considered to have an active backbone connection if the backbone area is actively attached and there is at least one fully adjacent neighbor in it.

Area Border Router (ABR):

Cisco Systems Interpretation:

A router is considered to be an ABR if it has more than one area Actively Attached and one of them is the backbone area.

IBM Interpretation:

A router is considered to be an ABR if it has more than one Actively Attached area and the backbone area Configured.

2.2 Implementation Details

The following changes are made to the base OSPF, described in [Ref1]:

1. The algorithm for Type 1 LSA (router-LSA) origination is changed to prevent a multi-area connected router from identifying itself as an ABR by the bit B (as described in section 12.4.1 of [Ref1]) until it considers itself as an ABR according to the definitions given in section 2.1.
2. The algorithm for the routing table calculation is changed to allow the router to consider the summary-LSAs from all attached areas if it is not an ABR, but has more than one attached area, or it does not have an Active Backbone Connection. Definitions of the terms used in this paragraph are given in section 2.1.

So, the paragraph 1 of section 16.2 of [Ref1] should be interpreted as follows:

"The inter-area routes are calculated by examining summary-LSAs. If the router is an ABR and has an Active Backbone Connection, only backbone summary-LSAs are examined. Otherwise (either the router is not an ABR or it has no Active Backbone Connection), the router should consider summary-LSAs from all Actively Attached areas..."

3. For Cisco ABR approach, the algorithm for the summary-LSAs origination is changed to prevent loops of summary-LSAs in situations where the router considers itself an ABR but doesn't have an Active Backbone Connection (and, consequently, examines summaries from all attached areas). The algorithm is changed to allow an ABR to announce only intra-area routes in such a situation.

So, the paragraph 2 of subsection 12.4.3 of [Ref1] should be interpreted as follows:

"Summary-LSAs are originated by area border routers. The precise summary routes to advertise into an area are determined by examining the routing table structure (see Section 11) in accordance with the algorithm described below. Note that while only intra-area routes are advertised into the backbone, if the router has an Active Backbone Connection, both intra-area and inter-area routes are advertised into the other areas; otherwise, the router only advertises intra-area routes into non-backbone areas."

For this policy to be applied we change steps 6 and 7 in the summary origination algorithm to be as follows:

Step 6:

"Else, if the destination of this route is an AS boundary router, a summary-LSA should be originated if and only if the routing table entry describes the preferred path to the AS boundary router (see Step 3 of Section 16.4). If so, a Type 4 summary-LSA is originated for the destination, with Link State ID equal to the AS boundary router's Router ID and metric equal to the routing table entry's cost. If the ABR performing this algorithm does not have an Active Backbone Connection, it can originate Type 4 summary-LSA only if the type of the route to the ASBR is intra-area. Note: Type 4 summary-LSAs should not be generated if Area A has been configured as a stub area."

Step 7:

"Else, the Destination type is network. If this is an inter-area route and the ABR performing this algorithm has an Active Backbone Connection, generate a Type 3 summary-LSA for the destination, with Link State ID equal to the network's address (if necessary, the Link State ID can also have one or more of the network's host bits set; see Appendix E for details) and metric equal to the routing table cost."

The changes in the ABR behavior described in this section allow a multi-area connected router to successfully route traffic destined for the backbone and other areas. Note that if the router does not have a backbone area Configured it does not actively attract inter-area traffic, because it does not consider itself an ABR and does not originate summary-LSAs. It still can forward traffic from one attached area to another along intra-area routes in case other routers in corresponding areas have the best inter-area paths over it, as described in section 1.2.

By processing all summaries when the backbone is not active, we prevent the ABR, which has just lost its last backbone adjacency, from dropping any packets going through the ABR in question to another ABR and destined towards the backbone or other areas not connected to the ABR directly.

3 Virtual Link Treatment

The Cisco ABR approach described in this document requires an ABR to have at least one active interface in the backbone area. This requirement may cause problems with virtual links in those rare situations where the backbone area is purely virtual, as shown in Figure 3, and the state of the VL is determined as in [Ref1].




Figure 3. Purely Virtual Backbone

If R1 and R2 treat virtual links as in [Ref1], their virtual links will never go up, because their router-LSAs do not contain the B-bit, which is, in turn, because the routers do not have active interfaces (virtual links) in the backbone and do not consider themselves ABRs.

Note that this problem does not appear if one of the routers has a real interface in the backbone, as it usually is in real networks.

Though the situation described is deemed to be rather rare, implementations supporting Cisco ABR behavior may consider changing VL-specific code so that a virtual link is reported up (an InterfaceUp event is generated) when a router with corresponding router-ID is seen via Dijkstra, no matter whether its router-LSA indicates that it is an ABR or not. This means that checking of configured virtual links should be done not in step 4 of the algorithm in 16.1 of [Ref1] when a router routing entry is added, but every time a vertex is added to the SPT in step 3 of the same algorithm.

4 Compatibility

The changes of the OSPF ABR operations do not influence any aspects of the router-to-router cooperation and do not create routing loops, and hence are fully compatible with standard OSPF. Proof of compatibility is outside the scope of this document.

5 Deployment Considerations

This section discusses the deployments details of the ABR behaviors described in this document. Note that this approach is fully compatible with standard ABR behavior, so ABRs acting as described in [Ref1] and in this document can coexist in an OSPF domain and will function without problems.

Deployment of ABRs using the alternative methods improves the behavior of a router connected to multiple areas without a backbone attachment, but can lead to unexpected routing asymmetry, as described below.

Consider an OSPF domain depicted in Figure 4.




Figure 4. Inter-area routing asymmetry

Assume that R3 uses the approach described in this document. In this case R2 will have inter-area routes to network M via ABR R1 only. R5 in turn will have its inter-area route to network N via R4, but as far as R4 is only reachable via R3, all traffic destined to network N will pass through R3. R3 will have an intra-area route to network N via R2 and will, of course, route it directly to it (because intra-area routes are always preferred over inter-area ones). Traffic going back from network N to network M will pass through R2 and will be routed to R1, as R2 will not have any inter-area routes via R3. So, traffic from N to M will always go through the backbone while traffic from M to N will cross the areas directly via R3 and, in this example, will not use a more optimal path through the backbone.

Note that this problem is not caused by the fact that R3 uses the alternative approach. The reason for attracting the attention to it is that R3 is not really functioning as an ABR in case this new behavior is used, i.e., it does not inject summary-LSAs into the attached areas, but inter-area traffic can still go through it.

6 Security Considerations

The alternative ABR behaviors specified in this document do not raise any security issues that are not already covered in [Ref1].

7 Acknowledgements

Authors would like to thank Alvaro Retana, Russ White, and Liem Nguyen for their review of the document.

8 Disclaimer

This document describes OSPF ABR implementations of respective vendors "as is", only for informational purposes, and without any warranties, guarantees or support. These implementations are subject to possible future changes. For the purposes of easier deployment, information about software versions where described behavior was integrated is provided below.

Initial Cisco ABR implementation (slightly different from the one described in this memo, requiring non-backbone areas to be configured, and not necessarily actively attached in the ABR definition) was introduced in Cisco IOS (tm) version 11.1(6). Cisco ABR behavior described in this document was integrated in Cisco IOS (tm) in version 12.1(3)T.

The ABR behavior described as IBM ABR approach was implemented by IBM in IBM Nways Multiprotocol Routing Services (MRS) 3.3.

Note that the authors do not intend to keep this document in sync with actual implementations.

10 References

[Ref1] Moy, J., "OSPF version 2", STD 54, RFC 2328, April 1998.

11 Authors' Addresses

Alex Zinin
Alcatel

E-Mail: zinin@psg.com

Derek M. Yeung
Procket Networks
1100 Cadillac Ct
Milpitas, CA 95035

Phone: 408-635-7911
E-Mail: myeung@procket.com

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

Phone: 919-387-6971
E-Mail: acee@redback.com

12 Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.