

Network Working Group
Internet Draft
Expiration Date: April 2000
File name: [draft-ietf-ospf-abr-alt-01.txt](#)

Alex Zinin
AMT Group
Acee Lindem
IBM
Derek Yeung
Cisco Systems
October 1999

Alternative OSPF ABR Implementations
draft-ietf-ospf-abr-alt-01.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet Drafts are working documents of the Internet Engineering Task Force (IETF), its Areas, and its Working Groups. Note that other groups may also distribute working documents as Internet Drafts.

Internet Drafts are draft documents valid for a maximum of six months. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a "working draft" or "work in progress".

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

OSPF [[Ref1](#)] is a link-state intra-domain routing protocol used for routing in IP networks. Though the definition of the ABR in the current 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 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" technique. OSPF prevents inter-area routing loops by implementing a split-horizon mechanism, permitting 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 the [section 12.4.3](#) of [Ref1], R3 will originate for 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 the 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 effectively drop all traffic not destined for the areas directly attached to it. The same problem can be seen 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 the 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 summary-LSAs directly from a backbone-attached router (R1 or R2, or both in our example) and will install 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 is not applicable 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 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.

Described solutions are targeted to the situation when an ABR has no backbone connection. It implies that a router connected to multiple areas without a backbone connection is not an ABR and must 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 in the meaning of 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.

Though this document is initially oriented to processing Type-3 LSAs and, consequently, is targeted to improving OSPF inter-area routing, it's acceptable to apply described methods to Type-4 LSAs, which will

lead to improvement of external routing in an OSPF domain.

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 configured and the backbone area actively attached.

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 of 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 of 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 the backbone area connection is not available. 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 read 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 must consider summary-LSAs from all actively attached areas..."

3. The algorithm of 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 announce only intra-area routes in such a situation.

So, the paragraph 2 of subsection 12.4.3 of [[Ref1](#)] should be read 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 only intra-area routes are advertised into the backbone, both intra-area and inter-area routes are advertised into the other areas, provided that the router has an active backbone connection. Otherwise the router is allowed to advertise only 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 Handling Transitions

The behavior described in this document requires special processing when the router's ABR status or backbone connection status changes.

3.1 ABR status change

The router's ABR flag can change due to the following events:

For Cisco ABR definition:

1. A new area has been configured. The router becomes an ABR if the new area is the first non-backbone area and the router has the backbone area actively attached.

2. An area has been deconfigured. The router stops being an ABR if only one configured area has left or the area in question is the backbone.
3. The state machine of an OSPF-enabled interface in the backbone area has transitioned to a state higher than Down. The router becomes an ABR if there is at least one non-backbone area configured.
4. The state machine of an OSPF-enabled interface in the backbone area has transitioned to the state Down. The router stops being an ABR if the interface in question was the last one in the backbone area.

For IBM ABR definition:

1. A new area has been configured. The router becomes an ABR if the new area is the backbone and the router is actively attached to two or more areas.
2. An area has been deconfigured. The router stops being an ABR if only one configured area has left or the area in question is the backbone.
3. The state machine of an OSPF-enabled interface transitions into a state higher than Down, an area becomes actively attached. If this is the second actively attached area and the backbone area is configured the router becomes an ABR.
4. The state machine of an OSPF-enabled interface transitions into Down state. If this is the last active interface in the area, it is no more considered actively attached. The router stops being an ABR if there are no longer multiple actively attached areas.

If after one of the events listed above the router's ABR flag has changed to the positive state, the router must do the following:

- a) reconstruct the router-LSAs for all area databases, setting the bit B to 1.
- b) flood the new router-LSA to all neighbors according to [\[Ref1\]](#).
- c) if the router has an active backbone connection---schedule the routing table recalculation for all areas to get rid of the inter-area routes through non-backbone areas.

If the router's ABR flag has changed to the negative state, the router must do the following:

- a) reconstruct the router-LSAs for all areas, setting the bit B to 0.
- b) flood the new router-LSA to all neighbors according to [\[Ref1\]](#).
- c) schedule the routing table recalculation to install the inter-area routes via non-backbone areas.
- d) flush all locally originated summary-LSAs from all non-backbone areas.

3.2 Backbone connection state change

A change of the backbone connection state can be a result of the following events:

1. The state machine for a neighbor in the backbone area has transitioned to the state Full. If the neighbor in question is the first fully adjacent neighbor, the backbone connection becomes active.
2. The state machine for a neighbor in the backbone area has transitioned to a state other than Full or a neighbor has been deleted from the neighbor list of an interface in the backbone area. The backbone connection becomes inactive if the neighbor in question was the last fully adjacent neighbor in the backbone.

If the state of the backbone connection has changed due to one of the events listed above, the router must schedule the routing table calculation according to [\[Ref1\]](#) and the changes described in [section 2.2](#) (even if the router is still an ABR). This will make the router flush old inter-area routes from the routing table and install new ones.

[4](#) Virtual Link Treatment

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.




Figure 3. Purely Virtual Backbone

If R1 and R3 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 described situation is deemed to be rather rare, it is recommended to change the implementation supporting Cisco ABR behavior 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 must 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.

5 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.

6 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 get into 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 get into 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.

7 Security Considerations

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

8 References

[Ref1] J. Moy. OSPF version 2. Technical Report [RFC 2328](http://ftp.isi.edu/in-notes/rfc2328.txt), Internet Engineering Task Force, 1998. [ftp://ftp.isi.edu/in-notes/rfc2328.txt](http://ftp.isi.edu/in-notes/rfc2328.txt).

9 Authors' Address

Alex D. Zinin
AMT Group
8 Maly Znamensky per., bld. 1
Office 3b
121019 Moscow, Russia
E-mail: zinin@amt.ru

Acee Lindem
IBM Corporation
Research Triangle Park, NC USA
E-mail: acee@raleigh.ibm.com

Derek M. Yeung
Cisco Systems Inc.
San Jose, CA, USA
E-mail: myeung@cisco.com

