

Network Working Group
Internet-Draft
Obsoletes: RFC3137
(if approved)
Intended status: Informational
Expires: August 29, 2011

A. Retana
L. Nguyen
R. White
A. Zinin
Cisco Systems, Inc.
D. McPherson
Verisign, Inc.
February 25, 2011

OSPF Stub Router Advertisement
draft-retana-ospf-rfc3137bis-00

Abstract

This memo describes a backward-compatible technique that may be used by OSPF (Open Shortest Path First) implementations to advertise unavailability to forward transit traffic or to lower the preference level for the paths through such a router. In some cases, it is desirable not to route transit traffic via a specific OSPF router. However, OSPF does not specify a standard way to accomplish this.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 29, 2011.

Copyright Notice

Copyright (c) 2011 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. Motivation | 3 |
| 2. Requirements Language | 3 |
| 3. Proposed Solution | 3 |
| 4. Compatibility issues | 4 |
| 5. Security Considerations | 4 |
| 6. Acknowledgements | 4 |
| 7. References | 5 |
| 7.1. Normative References | 5 |
| 7.2. Informative References | 5 |
| Appendix A. Changes from RFC 3137 | 5 |
| Authors' Addresses | 5 |

1. Motivation

In some situations, it may be advantageous to inform routers in a network not to use a specific router as a transit point, but still route to it. Possible situations include the following:

- o The router is in a critical condition (for example, has very high CPU load or does not have enough memory to store all LSAs or build the routing table).
- o Graceful introduction and removal of the router to/from the network.
- o Other (administrative or traffic engineering) reasons.

Note that the proposed solution does not remove the router from the topology view of the network (as could be done by just flushing that router's router-LSA), but prevents other routers from using it for transit routing, while still routing packets to the router's own IP addresses, i.e., the router is announced as a stub.

It must be emphasized that the proposed solution provides real benefits in networks designed with at least some level of redundancy so that traffic can be routed around the stub router. Otherwise, traffic destined for the networks reachable through such a stub router will be still routed through it.

2. Requirements Language

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 [RFC2119].

3. Proposed Solution

The solution described in this document solves two challenges associated with the outlined problem. In the description below, router X is the router announcing itself as a stub.

- 1) Making other routers prefer routes around router X while performing the Dijkstra calculation.
- 2) Allowing other routers to reach IP prefixes directly connected to router X.

Note that it would be easy to address issue 1) alone by just flushing

router X's router-LSA from the domain. However, it does not solve problem 2), since other routers will not be able to use links to router X in Dijkstra (no back link), and because router X will not have links to its neighbors.

To address both problems, router X announces its router-LSA to the neighbors with the costs of all non-stub links (links of the types other than 3) set to LSInfinity (16-bit value 0xFFFF, rather than 24-bit value 0xFFFFFFFF used in summary and AS-external LSAs).

The solution above applies to both OSPFv2 [RFC2328] and OSPFv3 [RFC5340].

4. Compatibility issues

Some inconsistency may be seen when the network is constructed of the routers that perform intra-area Dijkstra calculation as specified in RFC 1247 [RFC1247] (discarding link records in router-LSAs that have LSInfinity cost value) and routers that perform it as specified in RFC 1583 [RFC1583] and higher (do not treat links with LSInfinity cost as unreachable). Note that this inconsistency will not lead to routing loops, because if there are some alternate paths in the network, both types of routers will agree on using them rather than the path through the stub router. If the path through the stub router is the only one, the routers of the first type will not use the stub router for transit (which is the desired behavior), while the routers of the second type will still use this path.

5. Security Considerations

The technique described in this document does not introduce any new security issues into the OSPF protocol.

6. Acknowledgements

The authors of this document do not make any claims on the originality of the ideas described. Among other people, we would like to acknowledge Henk Smit for being part of one of the initial discussions around this topic.

We would also like to thank Shishio Tsuchiya, Gunter Van de Velde and Tomohiro Yamagata for reminding us of the need to document the OSPFv3 behavior.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

7.2. Informative References

- [RFC1247] Moy, J., "OSPF Version 2", RFC 1247, July 1991.
- [RFC1583] Moy, J., "OSPF Version 2", RFC 1583, March 1994.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.

Appendix A. Changes from RFC 3137

- o Edited in support for OSPFv3.
- o Updated references and author information.
- o Miscellaneous edits.

Authors' Addresses

Alvaro Retana
Cisco Systems, Inc.
7025 Kit Creek Rd.
Research Triangle Park, NC 27709
USA

Email: aretana@cisco.com

Liem Nguyen
Cisco Systems, Inc.
3750 Cisco Way
San Jose, CA 95134
USA

Phone: +1 408 527 0670
Email: lhnguyen@cisco.com

Russ White
Cisco Systems, Inc.
7025 Kit Creek Rd.
Research Triangle Park, NC 27709
USA

Email: russwh@cisco.com

Alex Zinin
Cisco Systems, Inc.
Capital Tower, 168 Robinson Rd.
Singapore, Singapore 068912
Singapore

Email: azinin@cisco.com

Danny McPherson
Verisign, Inc.
21345 Ridgetop Circle
Dulles, VA 20166
USA

Email: dmcpherson@verisign.com

