

Network Working Group
Request for Comments: 3137
Category: Informational

A. Retana
L. Nguyen
R. White
Cisco Systems
A. Zinin
Nexsi Systems
D. McPherson
Amber Networks
June 2001

OSPF Stub Router Advertisement

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 (2001). All Rights Reserved.

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.

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 router's own IP addresses, i.e., the router is announced as 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. 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 as follows.

- o costs of all non-stub links (links of the types other than 3) are set to LSInfinity (16-bit value 0xFFFF, rather than 24-bit value 0xFFFFFFFF used in summary and AS-external LSAs).
- o costs of stub links (type 3) are set to the interface output cost.

This addresses issues 1) and 2).

3. Compatibility issues

Some inconsistency may be seen when the network is constructed of the routers that perform intra-area Dijkstra calculation as specified in [RFC1247] (discarding link records in router-LSAs that have LSInfinity cost value) and routers that perform it as specified in [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.

4. 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.

5. Security Considerations

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

6. References

- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [RFC1247] Moy, J., "OSPF Version 2", RFC 1247, July 1991.
- [RFC1583] Moy, J., "OSPF Version 2", RFC 1583, March 1994.

7. Authors' Addresses

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

EEmail: aretana@cisco.com

Liem Nguyen
7025 Kit Creek Rd.
Research Triangle Park, NC 27709
USA

EEmail: lhnguyen@cisco.com

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

EEmail: riw@cisco.com

Alex Zinin
Nexsi Systems
1959 Concourse Drive
San Jose, CA 95131

EEmail: azinin@nexsi.com

Danny McPherson
Amber Networks
48664 Milmont Drive
Fremont, CA 94538

EEmail: danny@ambernetworks.com

8. Full Copyright Statement

Copyright (C) The Internet Society (2001). 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.