

Network Working Group
Internet-Draft
Obsoletes: [3137](#) (if approved)
Intended status: Informational
Expires: July 21, 2013

A. Retana
L. Nguyen
Cisco Systems, Inc.
A. Zinin
Cinarra Systems
R. White
D. McPherson
Verisign, Inc.
January 17, 2013

OSPF Stub Router Advertisement
draft-ietf-ospf-rfc3137bis-03

Abstract

This document 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.

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 July 21, 2013.

Copyright Notice

Copyright (c) 2013 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.	Introduction	3
2.	Solutions	3
2.1.	OSPFv3-only Solution	4
3.	Maximum Link Metric	4
4.	Deployment Considerations	4
5.	Security Considerations	5
6.	IANA Considerations	5
7.	Acknowledgements	5
8.	Informative References	5
Appendix A.	Change Log	5
A.1.	Changes between the -00 and -01 versions.	5
A.2.	Changes between the -01 and -02 versions.	6
A.3.	Changes between the -02 and -03 versions.	6
	Authors' Addresses	6

1. Introduction

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 solution introduced in this document 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 discourages 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 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 may still be routed through it.

2. Solutions

The solution introduced 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 MaxLinkMetric.

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

2.1. OSPFv3-only Solution

OSPFv3 [[RFC5340](#)] introduced additional options to provide similar, if not better, control of the forwarding topology; the R-bit provides a more granular indication of whether a router is active and should be used for transit traffic.

It is left to network operators to decide which technique to use in their network.

3. Maximum Link Metric

[Section 2](#) refers to the cost of all non-stub links as MaxLinkMetric, which is a new fixed architectural value introduced in this document.

MaxLinkMetric

The metric value indicating that the link described by an LSA should not be used as transit. Used in router-LSAs (see [Section 2](#)). It is defined to be the 16-bit binary value of all ones: 0xffff.

4. Deployment Considerations

When using MaxLinkMetric, some inconsistency may be seen if the network is constructed of routers that perform intra-area Dijkstra calculation as specified in [[RFC1247](#)] (discarding link records in router-LSAs that have a MaxLinkMetric cost value) and routers that perform it as specified in [[RFC1583](#)] and higher (do not treat links with MaxLinkMetric 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.

On the other hand, clearing the R-bit will consistently result in the router not being used as transit.

5. Security Considerations

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

6. IANA Considerations

This document has no actions for IANA.

7. 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, Tomohiro Yamagata, Faraz Shamim and Acee Lindem who provided significant input for the latest version of this document.

8. 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. Change Log

A.1. Changes between the -00 and -01 versions.

- o Defined a new architectural constant (MaxLinkMetric) to eliminate any confusion about the interpretation of LSInfinity.
- o Added a section to reference the R-bit and V6-bit in OSPFv3.
- o Updated acks and contact information.

A.2. Changes between the -01 and -02 versions.

- o Took out references to not having a standard solution and incorporated the R-bit solution as part of the (renamed) "Solutions" section.
- o Various minor edits and reordered sections.

A.3. Changes between the -02 and -03 versions.

- o Updated contact information.
- o Renamed the 'Motivation' section to 'Introduction' because of an error in idnits.
- o Took out the [rfc2119](#) references as none of the keywords are used in the text.
- o Added an 'IANA Considerations' section to indicate that there are no actions required.

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

Email: lhnguyen@cisco.com

Alex Zinin
Cinarra Systems
Menlo Park, CA
USA

Email: alex.zinin@gmail.com

Russ White
Verisign, Inc.
12061 Bluemont Way
Reston, VA 20190
USA

Email: riwhite@verisign.com

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

Email: dmcpherson@verisign.com

