

Internet Engineering Task Force
Internet-Draft
Updates: [3137](#) (if approved)
Intended status: Standards Track
Expires: July 9, 2011

S. Tsuchiya, Ed.
G. Van de Velde
Cisco Systems
T. Yamagata
KDDI Corporation
January 5, 2011

OSPFv3 Stub Router Advertisement
draft-shishio-ospf-ospfv3-stub-03

Abstract

OSPFv3 accommodates for the possibility to indicate through the R-bit if a router is an active router and should be taken into consideration as a transit device. Another method available is the v6-bit indicating if a router or link should be excluded from IPv6 routing calculations.

A direct result is that OSPFv3 has "no transit capability" potentially based upon the setting of R-bit and V6-bit, unlike the stub OSPFv2 router functionality. This feature proposal has as purpose to re-introduce existing OSPFv2 stub router behavior into OSPFv3 to keep the operational service provider experience used to deploy, troubleshoot and be familiar with OSPFv2 stub routing.

OSPFv3 has similar metric field information field of all of LSAs, with exception of the Link-LSA, so [RFC3137](#) method can be re-utilized in OSPFv3.

To drive consistency between OSPFv2 and OSPFv3, there should be next to supporting both R-bit and v6-bit be support for "max-metric".

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."

Internet-Draft

OSPFv3 Stub Router Advertisement

January 2011

This Internet-Draft will expire on July 9, 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.

Internet-Draft

OSPFv3 Stub Router Advertisement

January 2011

Table of Contents

| | | |
|-----------------------------|---------------------------------------|-------------------|
| 1. | Motivation | 4 |
| 2. | Requirements Language | 4 |
| 3. | OSPFv2 operation | 4 |
| 3.1. | Wait for BGP during booting | 5 |
| 3.2. | LDP Synchronization | 5 |
| 3.3. | Configuration change | 6 |
| 4. | OSPFv3 operation | 6 |
| 4.1. | R-bit and v6-bit | 6 |
| 4.2. | OSPFv2 compatibility mode | 6 |
| 5. | Acknowledgements | 7 |
| 6. | IANA Considerations | 7 |
| 7. | Security Considerations | 7 |
| 8. | References | 7 |
| 8.1. | Normative References | 7 |
| 8.2. | Informative References | 8 |
| Appendix A. | Additional Stuff | 8 |
| | Authors' Addresses | 8 |

1. Motivation

OSPF Stub Router Advertisement [[RFC3137](#)] describes a set of situations when the Service provider has a desire to utilize this functionality.

- o The router is in a critical condition resulting in either a very high CPU load, or not enough memory to store all LSAs, or doesn't succeed to build the routing table
- o Graceful introduction or removal of a router to or from the network
- o Other (administrative or traffic engineering) reasons

Even if the network will be moved or migrated towards from IPv4 in combination with OSPFv2 towards IPv6 using OSPFv3 [[OSPFv3](#)] technology, it remains important that the operational behaviour remains similar between routing protocols.

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. OSPFv2 operation

[RFC3137](#) [[RFC3137](#)] describes in [section 2](#) in detail the behavior of link cost metrics. i.e. Router X announces its Router LSA to the neighbor with costs of all non-stub links which are set to LSInfinity(0xFFFF), while stub links are announced with interface cost.

Often the operator will check interface metric of ospf database assuming he would like to confirm whether the router is announcing LSInfinity.

Many service provider operators are using OSPF stub router advertisement [[RFC3137](#)] for OSPFv2 [[OSPFv2](#)]. This feature is supported by majority of OSPFv2 implementations. Use cases are below;

[3.1.](#) Wait for BGP during booting




Figure 1

In this example R2 can be assumed it is the best path towards the Internet from R1. When R2 reloads it would result that R3 would be best path for going towards Internet. From the moment R2 is reloaded and OSPF has converged, there may be a situation when BGP is still not converged to the full. If in this situation R1 should not send traffic towards R2 just yet. R2 should send LSInfinity(0xFFFF) to indicate that R1 should wait for R2's BGP converge. Once BGP is fully converged, R2 LSA's out with correct interface metric value in OSPFv2 area which will result in R2 being reintroduced as the primary path.

[3.2.](#) LDP Synchronization




Figure 2

Assume that from R1 the best path to R4 is via R2 in this MPLS network. When R2 is reloaded, then R3 is the only and hence also the best path. At some point in time R2 is fully successfully reloaded resulting that OSPF has converged also. This does not necessary mean LDP has fully converged either. In this situation R1 should not send traffic to R2 immediately. In that case R2 could send LSInfinity(0xFFFF) resulting in a situation where R1 must wait for R2 to be fully be available and transit states have been passed. From the moment LDP converged on R2, it can distribute the traditional Interface OSPF metric value. This operation will result that OSPF and LDP have a converged behaviour. The importance and the description of this behaviour can be found in [[LDP-Sync](#)].

[3.3.](#) Configuration change




Figure 3

When operator needs R2 configuration change, R2 sends LSInfinity(0xFFFF) for traffic engineering. R2 configuration completed, R2 sends correct metric value in OSPFv2 area.

[4.](#) OSPFv3 operation

[4.1.](#) R-bit and v6-bit

[OSPFv3] explains at [section 2.7](#) the following: If the "R-bit" is clear, an OSPF speaker can participate in OSPF topology distribution without being used to forward transit traffic. The V6-bit specializes the R-bit; if the V6-bit is clear, an OSPF speaker can participate in OSPF topology distribution without being used to forward IPv6 datagrams. If the R-bit is set and the V6-bit is clear, IPv6 datagrams are not forwarded but datagrams belonging to another protocol family may be forwarded.

This protocol implementation is useful in multi address family environment such as [[OSPFv3-AF](#)]. But Service Provider operators have to check both the "R-bit" and "v6-bit" during their operation and introduce both training and operational changes to make this a true usable technology. Operators have believe that a useful approach would be to rely upon successful IPv4 OSPFv2 behaviour and to add a "OSPFv2 compatibility mode" in IPv6 only environment to mimic OSPFv2 behavior in this environment.

The functionality of the R-bit and v6-bit operations is described in [[OSPFv3](#)]'s Errata 2078 more detail.

A Router should support a "R-bit" know with a clear wait for BGP or waiting-before-becoming-active time on start-up same as [[RFC3137](#)] indicates.

[4.2.](#) OSPFv2 compatibility mode

An OSPFv3 routing device has through the area scope LSAs metric information of all of devices. As result the router can announce the interface metric LSInfinity(0xFFFF). This is simple implementation

model not requiring operational service provider changes.

[5.](#) Acknowledgements

This document is based on [RFC3137](#) [[RFC3137](#)] and [RFC5340](#) [[OSPFv3](#)]Errata 2078. [RFC3137](#) [[RFC3137](#)] done by Alvaro Retana,Liem Nguyen,Russ White,Alex Zinin and Danny McPherson. [RFC5340](#) [[OSPFv3](#)]

done by Rob Coltun, Dennis Ferguson, John Moy and Acee Lindem. Balaji Ganesh pointed out problem of [Section 4.8.1](#) RFC5340 [OSPFv3] on Errata 2046. Michael Barnes brought up the fact, Acee Lindem reported on Errata 2078. Tsuyoshi Momose advised us about how to write internet-draft. Fred Baker was reviewed this document in initial stage. Cisco OSPF coder's are gave comments about this document. Especially Peter Psenak did deep discussion about both modes. Michael Barnes pointed out Errata 2078 exists. The authors would like to thank all of them for their activity, comments and help.

[6.](#) IANA Considerations

This document has no actions for IANA.

[7.](#) Security Considerations

The technique described in this document does not introduce any new security issues into OSPFv3 protocol.

[8.](#) References

[8.1.](#) Normative References

- [OSPFv2] Moy, J., "'OSPF Version 2'", [RFC 2328](#), STD 54, April 1998, <<http://tools.ietf.org/html/rfc2328>>.
- [OSPFv3] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "'OSPF for IPv6'", [RFC 5340](#), July 2008, <<http://tools.ietf.org/html/rfc5340>>.
- [OSPFv3-AF] Lindem, A., Mirtorabi, S., Roy, A., Barnes, M., and R. Aggarwal, "'Support of Address Families in OSPFv3'", [RFC 5838](#), April 2010, <<http://tools.ietf.org/html/rfc5838>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[8.2.](#) Informative References

[LDP-Sync]

Jork, M., Atlas, A., and L. Fang, "'LDP IGP Synchronization'", [RFC 5443](http://tools.ietf.org/html/rfc5443), March 2009, <<http://tools.ietf.org/html/rfc5443>>.

[RFC3137] Retana, A., Nguyen, L., White, R., Zinin, A., and D. McPherson, "'OSPF Stub Router Advertisement'", [RFC 3137](http://tools.ietf.org/html/rfc3137), June 2001, <<http://tools.ietf.org/html/rfc3137>>.

[Appendix A](#). Additional Stuff

This becomes an Appendix.

Authors' Addresses

Shishio Tsuchiya (editor)
Cisco Systems
Shinjuku Mitsui Building, 2-1-1, Nishi-Shinjuku
Shinjuku-Ku, Tokyo 163-0409
Japan

Phone: +81 3 6434 6543
Email: shtsuchi@cisco.com

Gunter Van de Velde
Cisco Systems
Pegasus Parc
De kleetlaan 6a, DIEGEM, BRABANT 1831
BELGIUM

Phone: +32 2 704 5473
Email: gunter@cisco.com

Tomohiro Yamagata
KDDI Corporation
Garden Air Tower, 3-10-10, Iidabashi
Chiyoda-Ku, Tokyo 102-8460
Japan

Phone: +81 3 6678 3089
Email: to-yamagata@kddi.com