

Open Shortest Path First IGP
Internet-Draft
Intended status: Standards Track
Expires: January 17, 2016

S. Hegde
P. Sarkar
H. Gredler
Juniper Networks, Inc.
M. Nanduri
Microsoft Corporation
L. Jalil
Verizon
July 16, 2015

OSPF Link Overload
draft-hegde-ospf-link-overload-01

Abstract

Many OSPFv2 or OSPFv3 deployments run on overlay networks provisioned by means of pseudo-wires or L2-circuits. When the devices in the underlying network go for maintenance, it is useful to divert the traffic away from the node before the maintenance is actually scheduled. Since the nodes in the underlying network are not visible to OSPF, existing stub router mechanism described in [RFC3137] cannot be used. It is useful for routers in OSPFv2 or OSPFv3 routing domain to be able to advertise a link being in overload state to indicate impending maintenance activity in the underlying network devices.

This document describes the protocol extensions to disseminate link overload information in OSPFv2 and OSPFv3 protocol.

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].

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 January 17, 2016.

Copyright Notice

Copyright (c) 2015 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	4
2.	Link overload sub TLV	4
2.1.	OSPFv2 Link overload sub TLV	4
2.2.	OSPFv3 Link overload sub TLV	4
3.	Elements of procedure	5
3.1.	Point-to-point links	5
3.2.	Broadcast/NBMA links	6
3.3.	Point-to-multipoint links	6
4.	Backward compatibility	6
5.	Applications	6
5.1.	Pseudowire Services	7
5.2.	Controller based Traffic Engineering Deployments	7
6.	Security Considerations	8
7.	IANA Considerations	8
8.	Acknowledgements	8
9.	References	8
9.1.	Normative References	8
9.2.	Informative References	9
	Authors' Addresses	9

1. Introduction

It is useful for routers in OSPFv2 or OSPFv3 routing domain to be able to advertise a link being in overload state to indicate impending maintenance activity on the link. This document provides mechanisms to advertise link overload state in the flexible encodings provided by OSPFv2 Prefix/Link Attribute Advertisement ([I-D.ietf-ospf-prefix-link-attr]) and OSPFv3 Extended LSA ([I-D.ietf-ospf-ospfv3-lsa-extend]). Throughout this document, OSPF is used when the text applies to both OSPFv2 and OSPFv3. OSPFv2 or OSPFv3 is used when the text is specific to one version of the OSPF protocol.

2. Link overload sub TLV

2.1. OSPFv2 Link overload sub TLV

Link overload sub TLV is carried as part of the Extended link TLV as defined in [I-D.ietf-ospf-prefix-link-attr] for OSPFv2.




Figure 1: Link overload sub TLV for OSPFv2

Type : TBA

Length: 4

Value: Remote IPv4 address. The remote IP4 address is used to identify the particular link that is in overload state when there are multiple parallel links between two nodes.

2.2. OSPFv3 Link overload sub TLV

Link overload sub TLV is carried in the Router-link TLV as defined in the [I-D.ietf-ospf-ospfv3-lsa-extend] for OSPFv3. The Router-Link

TLV contains the neighbor interface-id and can uniquely identify the link on the remote node.




Figure 2: Link overload sub TLV for OSPFv3

Type : TBA

Length: 0

3. Elements of procedure

The Link Overload sub TLV indicates that the Link which carries the sub TLV is overloaded. The node that has the link going for maintenance, sets metric of the link to MAX-METRIC and re-originates the router LSA. The metric in the reverse direction also need to change to divert the traffic from reverse direction. The node SHOULD originate Link overload sub TLV and include it in Extended link TLV and originate the Extended Link Opaque LSA as defined in [I-D.ietf-ospf-prefix-link-attr] for OSPFv2 and E-Router-LSA as defined in [I-D.ietf-ospf-ospfv3-lsa-extend] for OSPFv3 and flood in the OSPF area.

when the originator of the Link Overload sub TLV, purges the extended link opaque LSA or re-originates without the Link Overload sub TLV, the metric on the remote node SHOULD be changed back to the original value.

Based on the link type of the overloaded link below actions MAY be taken by the receiver.

3.1. Point-to-point links

When a link overload TLV is received for a point-to-point link the receiver SHOULD identify the local link which corresponds to the overloaded link and set the metric to MAX-METRIC (0xffff). Receiver node MUST re-originate the router-LSA with the changed metric and flood into the OSPF area.

3.2. Broadcast/NBMA links

Broadcast or NBMA networks in OSPF are represented by a star topology where the Designated Router (DR) is the central point to which all other routers on the broadcast or NBMA network connect. As a result, routers on the broadcast or NBMA network advertise only their adjacency to the DR. Routers that do not act as DR do not form or advertise adjacencies with each other. For the Broadcast links, the MAX-METRIC on the outgoing link cannot be changed since all the neighbors are on same link. Setting the link cost to MAX-METRIC would impact paths going via all neighbors.

When a link-overload TLV is received by the remote end for a broadcast/NBMA link

- If it's DROther or BDR for that link, SHOULD not take any action.
- If receiving node is DR for the link, it MUST remove the originator of the link overload TLV from the list of connected neighbors and MUST re-originate the network LSA and flood into the OSPF area.

3.3. Point-to-multipoint links

Operation for the point-to-multipoint links is similar to the point-to-point links. When a link overload TLV is received for a point-to-multipoint link the receiver SHOULD identify the neighbor which corresponds to the overloaded link and set the metric to MAX-METRIC (0xffff). Receiver node MUST re-originate the router-LSA with the changed metric and flood into the OSPF area.

4. Backward compatibility

The mechanism described in the document is fully backward compatible. It is required that the originator and receiver of link-overload sub TLV understand the extensions defined in this document and in case of broadcast links the originator and the DR need to understand the extensions. Other nodes in the network compute based on increased metric and hence the feature is backward compatible.

5. Applications

5.1. Pseudowire Services




Figure 3: Pseudowire Services

Many service providers offer pseudo-wire services to customers using L2 circuits. The IGP protocol that runs in the customer network would also to run over the pseudo-wire to get seamless private network for the customer. Service providers want to offer overload kind of functionality when the PE device is taken-out for maintenance. The provider should guarantee that the PE is taken out for maintenance only after the service is successfully diverted on the alternate path. Link overload feature provides facilities to achieve this service by increasing the metric on the link but still allowing the traffic to use the link when there is no alternate path available.

5.2. Controller based Traffic Engineering Deployments




Figure 4: Controller based Traffic Engineering

Controller based deployments where the controller participates in the IGP protocol gets the link-overload information when the link maintenance is impending. Using this information controller finds an

alternate path. If there are no alternate paths satisfying the traffic engineering constraints, controller might temporarily relax the constraints and put the service on different path. In the above example when P1->P2 link goes for maintenance, controller gets the link-overload information and sets up an alternate path via P1->P3->P2. Once the traffic is diverted, P1->P2 link can be taken out for maintenance/upgrade.

6. Security Considerations

This document does not introduce any further security issues other than those discussed in [RFC2328] and [RFC5340].

7. IANA Considerations

This specification updates one OSPF registry:

OSPF Extended Link TLVs Registry

i) TBD - Link Overload TLV OSPFV3 Router Link TLV Registry

i) TBD - Link Overload TLV

8. Acknowledgements

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [RFC3137] Retana, A., Nguyen, L., White, R., Zinin, A., and D. McPherson, "OSPF Stub Router Advertisement", RFC 3137, June 2001.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.

9.2. Informative References

[I-D.ietf-ospf-ospfv3-lsa-extend]

Lindem, A., Mirtorabi, S., Roy, A., and F. Baker, "OSPFv3 LSA Extendibility", draft-ietf-ospf-ospfv3-lsa-extend-06 (work in progress), February 2015.

[I-D.ietf-ospf-prefix-link-attr]

Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", draft-ietf-ospf-prefix-link-attr-03 (work in progress), February 2015.

Authors' Addresses

Shraddha Hegde
Juniper Networks, Inc.
Embassy Business Park
Bangalore, KA 560093
India

Email: shraddha@juniper.net

Pushpasis Sarkar
Juniper Networks, Inc.
Embassy Business Park
Bangalore, KA 560093
India

Email: psarkar@juniper.net

Hannes Gredler
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
US

Email: hannes@juniper.net

Mohan Nanduri
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
US

Email: mmanduri@microsoft.com

Luay Jalil
Verizon

Email: luay.jalil@verizon.com