

Open Shortest Path First IGP
Internet-Draft
Intended status: Standards Track
Expires: August 27, 2017

S. Hegde
Juniper Networks, Inc.
P. Sarkar
H. Gredler
Individual
M. Nanduri
Microsoft Corporation
L. Jalil
Verizon
February 23, 2017

OSPF Link Overload
draft-ietf-ospf-link-overload-05

Abstract

When a link is being prepared to be taken out of service, the traffic needs to be diverted from both ends of the link. Increasing the metric to the highest metric on one side of the link is not sufficient to divert the traffic flowing in the other direction.

It is useful for routers in an OSPFv2 or OSPFv3 routing domain to be able to advertise a link being in an overload state to indicate impending maintenance activity on the link. This information can be used by the network devices to re-route the traffic effectively.

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

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 August 27, 2017.

Copyright Notice

Copyright (c) 2017 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.	Motivation	3
3.	Flooding Scope	4
3.1.	Area scope flooding	4
3.2.	Link scope flooding	4
4.	Link-Overload sub-TLV	4
4.1.	OSPFv2 Link-overload sub-TLV	4
4.2.	OSPFv3 Link-Overload sub-TLV	5
5.	Elements of procedure	5
5.1.	Point-to-point links	6
5.2.	Broadcast/NBMA links	6
5.3.	Point-to-multipoint links	7
5.4.	Unnumbered interfaces	7
5.5.	Hybrid Broadcast and P2MP interfaces	7
6.	Backward compatibility	8
7.	Applications	8
7.1.	Pseudowire Services	8
7.2.	Controller based Traffic Engineering Deployments	9
7.3.	L3VPN Services and sham-links	10
7.4.	Hub and spoke deployment	10
8.	Security Considerations	11
9.	IANA Considerations	11
10.	Acknowledgements	11
11.	References	11
11.1.	Normative References	11
11.2.	Informative References	12

Authors' Addresses	13
------------------------------	----

1. Introduction

When a node is being prepared for a planned maintenance or upgrade, [RFC6987] provides mechanisms to advertise the node being in an overload state by setting all outgoing link costs to MAX-METRIC (0xffff). These procedures are specific to the maintenance activity on a node and cannot be used when a single link attached to the node requires maintenance.

In traffic-engineering deployments, LSPs need to be diverted from the link without disrupting the services. It is useful to be able to advertise the impending maintenance activity on the link and to have LSP re-routing policies at the ingress to route the LSPs away from the link.

Many OSPFv2 or OSPFv3 deployments run on overlay networks provisioned by means of pseudo-wires or L2-circuits. Prior to devices in the underlying network going offline 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, the existing stub router mechanism described in [RFC6987] cannot be used. An application specific to this use case is described in Section 7.1

This document provides mechanisms to advertise link-overload state in the flexible encodings provided by OSPFv2 Prefix/Link Attribute Advertisement ([RFC7684]) and RI LSA ([RFC7770]). 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. Motivation

The motivation of this document is to reduce manual intervention during maintenance activities. The following objectives help to accomplish this in a range of deployment scenarios.

1. Advertise impending maintenance activity so that traffic from both directions can be diverted away from the link.
2. Allow the solution to be backward compatible so that nodes that do not understand the new advertisement do not cause routing loops.
3. Advertise the maintenance activity to other nodes in the network so that LSP ingress routers/controllers can learn of the

impending maintenance activity and apply specific policies to re-route the LSPs for traffic-engineering based deployments.

4. Allow the link to be used as last resort link to prevent traffic disruption when alternate paths are not available.

3. Flooding Scope

The link-overload information can be flooded in area scoped extended link LSA [RFC7684] or a link scoped RI LSA [RFC7770] or both based on the needs of the application. Section 7 describes applications requiring area scope as well as link scope link-overload information.

3.1. Area scope flooding

For OSPFv2, Link-Overload sub-TLV is carried in the extended Link TLV as defined in [RFC7684].

3.2. Link scope flooding

The link local scope RI LSA MAY carry the Link-Overload sub-TLV as defined in Section 4. The link local scope RI-LSA corresponds to the link on which the LSA arrives and there is no need to explicitly specify the remote IPv4 address. The remote IPv4 address field MAY be zero when the Link-Overload sub-TLV is carried in the link local RI LSA. The Link-Overload sub-TLV MAY appear in any instance of the link local RI-LSA. The Link-Overload sub-TLV is carried in the RI-LSA for both OSPFv2 and OSPFv3.

4. Link-Overload sub-TLV

4.1. OSPFv2 Link-overload sub-TLV

The Link-Overload sub-TLV identifies the link being in overload state. It is carried in extended Link TLV as defined in [RFC7684] or link local scope RI LSA as defined in [RFC7770].

OSPFv2. The Link-Overload information is carried as a property of the link and is flooded across the area. This information can be used by ingress routers or controllers to take special actions. An application specific to this use case is described in Section 7.2.

The precise action taken by the remote node at the other end of the link identified as overloaded depends on the link type.

5.1. Point-to-point links

The node that has the link to be taken out of service SHOULD set metric of the link to MAX-METRIC (0xffff) and re-originate the Router-LSA. The TE metric SHOULD be set to MAX-TE-METRIC -1 (0xffffffe) and the node SHOULD re-originate the TE Link Opaque LSAs. When a Link-Overload sub-TLV is received for a point-to-point link, the remote node SHOULD identify the local link which corresponds to the overloaded link and set the metric to MAX-METRIC (0xffff) and the remote node SHOULD re-originate the router-LSA with the changed metric. The TE metric SHOULD be set to MAX-TE-METRIC -1 (0xffffffe) and the TE opaque LSA for the link SHOULD be re-originated with new value.

Extended link opaque LSAs and the Extended link TLV are not scoped for multi-topology [RFC4915]. In multi-topology deployments [RFC4915], the Link-Overload sub-TLV carried in an Extended Link opaque LSA corresponds to all the topologies the link belongs to. The receiver node SHOULD change the metric in the reverse direction corresponding to all the topologies to which the reverse link belongs and re-originate the Router LSA as defined in [RFC4915].

When the originator of the Link-Overload sub-TLV purges the Extended Link Opaque LSA or re-originate it without the Link-Overload sub-TLV, the remote node must re-originate the appropriate LSAs with the metric and TE metric values set to their original values.

5.2. 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 logically. 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 remote link cannot be changed since all the neighbours are on same link. Setting the link cost to MAX-METRIC would impact paths going via all neighbours.

The node that has the link to be taken out of service SHOULD set metric of the link to MAX-METRIC(0xffff) and re-originate the Router-LSA. The TE metric SHOULD be set to MAX-TE-METRIC -1(0xfffffff) and the node SHOULD re-originate the TE Link Opaque LSAs. For a broadcast link, the two part metric as described in [RFC8042] is used. The node originating the Link-Overload sub-TLV MUST set the metric in the Network-to-Router Metric sub-TLV to MAX-METRIC 0xffff for OSPFv2 and OSPFv3 and re-originate the LSAs the TLV is carried-in. The nodes that receive the two part metric should follow the procedures described in [RFC8042]. The backward compatibility procedures described in [RFC8042] should be followed to ensure loop free routing.

5.3. Point-to-multipoint links

Operation for the point-to-multipoint links is similar to the point-to-point links. When a Link-Overload sub-TLV is received for a point-to-multipoint link the remote node SHOULD identify the neighbour which corresponds to the overloaded link and set the metric to MAX-METRIC (0xffff). The remote node MUST re-originate the Router-LSA with the changed metric and flood into the OSPF area.

5.4. Unnumbered interfaces

Unnumbered interface do not have a unique IP addresses and borrow address from other interfaces. [RFC2328] describes procedures to handle unnumbered interfaces in the context of the Router LSA. We apply a similar procedure to the Extended Link TLV carrying the Link-Overload sub-TLV in to handle unnumbered interfaces. The link-data field in the Extended Link TLV carries the interface-id instead of the IP address. The Link-Overload sub-TLV carries the remote interface-id in the remote-ip-address field if the interface is unnumbered. Procedures to obtain interface-id of the remote side are defined in [RFC4203].

5.5. Hybrid Broadcast and P2MP interfaces

Hybrid Broadcast and P2MP interfaces represent a broadcast network modeled as P2MP interfaces. [RFC6845] describes procedures to handle these interfaces. Operation for the Hybrid interfaces is similar to the P2MP interfaces. When a Link-Overload sub-TLV is received for a hybrid link the remote node SHOULD identify the neighbour which corresponds to the overloaded link and set the metric to MAX-METRIC (0xffff). All the remote nodes connected to originator MUST re-originate the Router-LSA with the changed metric and flood into the OSPF area.

6. Backward compatibility

The mechanism described in the document is fully backward compatible. It is required that the originator of the Link-Overload sub-TLV as well as the node at the remote end of the link identified as overloaded understand the extensions defined in this document. In the case of broadcast links, the backward compatibility procedures as described in [RFC8042] are applicable.

7. Applications

7.1. 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 run over the pseudo-wire to create 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 an alternate path. There can be large number of customers attached to a PE node and the remote end-points for these pseudo-wires are spread across the service provider's network. It is a tedious and error-prone process to change the metric for all pseudo-wires in both directions. The link-overload feature simplifies the process by increasing the metric on the link in the reverse direction as well so that traffic in both directions is diverted away from the PE undergoing maintenance. The Link-Overload feature allows the link to be used as a last resort link so that traffic is not disrupted when alternative paths are not available.




Figure 3: Pseudowire Services

In the example shown in Figure 3, when the PE1 node is going for maintenance, service providers set the PE1 to overload state. The PE1 going in overload state triggers all the CEs (In this example CE1) connected to the PE to set their pseudowire links passing via PE1 to link-overload state. The mechanisms used to communicate between PE1 and CE1 is outside the scope of this document. CE1 sets the link-overload state on its private VLAN connecting CE3, CE2 and CE4 and modifies the metric to MAX_METRIC and floods the information, the remote end of the link at CE3, CE2, and CE4 also set the metric on the link to MAX-METRIC and the traffic from both directions gets diverted away from the link.

7.2. Controller based Traffic Engineering Deployments

In controller-based deployments where the controller participates in the IGP protocol, the controller can also receive the link-overload information as a warning that link maintenance is imminent. Using this information, the controller can find alternate paths for traffic which use the affected link. The controller can apply various policies and re-route the LSPs away from the link undergoing maintenance. If there are no alternate paths satisfying the traffic engineering constraints, the controller might temporarily relax those constraints and put the service on a different path.




Figure 4: Controller based Traffic Engineering

In the above example, PE1->PE2 LSP is set-up to satisfy a constraint of 10 Gbps bandwidth on each link. The links P1->P3 and P3->P2 have only 1 Gbps capacity and there is no alternate path satisfying the bandwidth constraint of 10GB. When P1->P2 link is being prepared for maintenance, the controller receives the link-overload information, as there is no alternate path available which satisfies the constraints, controller chooses a path that is less optimal and temporarily sets up an alternate path via P1->P3->P2. Once the traffic is diverted, the P1->P2 link can be taken out of service for maintenance/upgrade.

7.3. L3VPN Services and sham-links

Many service providers offer L3VPN services to customers and CE-PE links run OSPF [RFC4577]. When PE goes for maintenance, all the links on the PE can be set to link-overload state which will guarantee that the traffic from CEs also gets diverted. The interaction between OSPF and BGP is outside the scope of this document.

Another useful usecase is when ISPs provide sham-link services to customers [RFC4577]. When PE goes for maintenance, all sham-links on the PE can be set to link-overload state and traffic can be diverted from both ends without having to touch the configurations on the remote end of the sham-links.

7.4. Hub and spoke deployment

OSPF is largely deployed in Hub and Spoke deployments with a number of spokes connecting to the Hub. It is a general practice to deploy multiple Hubs with all spokes connecting to these Hubs to achieve redundancy. When a Hub node goes down for maintenance, all links on the Hub can be set to link-overload state and traffic gets diverted

from spoke sites as well without having to make configuration changes on the spokes.

8. Security Considerations

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

9. IANA Considerations

This specification updates one OSPF registry:

OSPF Extended Link TLVs Registry

i) TBD - Link-Overload sub-TLV

OSPFV3 Router Link TLV Registry

i) TBD - Link-Overload sub-TLV

OSPF RI TLV Registry

i) TBD - Link-Overload sub-TLV

BGP-LS Link NLRI Registry [RFC7752]

i) TBD - Link-Overload sub-TLV

10. Acknowledgements

Thanks to Chris Bowers for valuable inputs and edits to the document. Thanks to Jeffrey Zhang and Acee Lindem for inputs. Thanks to Karsten Thomann for careful review and inputs on the applications where link-overload is useful.

11. References

11.1. Normative References

[RFC6845] Sheth, N., Wang, L., and J. Zhang, "OSPF Hybrid Broadcast and Point-to-Multipoint Interface Type", RFC 6845, DOI 10.17487/RFC6845, January 2013, <<http://www.rfc-editor.org/info/rfc6845>>.

[RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", RFC 7684, DOI 10.17487/RFC7684, November 2015, <<http://www.rfc-editor.org/info/rfc7684>>.

- [RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", RFC 7752, DOI 10.17487/RFC7752, March 2016, <<http://www.rfc-editor.org/info/rfc7752>>.
- [RFC7770] Lindem, A., Ed., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", RFC 7770, DOI 10.17487/RFC7770, February 2016, <<http://www.rfc-editor.org/info/rfc7770>>.
- [RFC8042] Zhang, Z., Wang, L., and A. Lindem, "OSPF Two-Part Metric", RFC 8042, DOI 10.17487/RFC8042, December 2016, <<http://www.rfc-editor.org/info/rfc8042>>.

11.2. Informative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <<http://www.rfc-editor.org/info/rfc2328>>.
- [RFC4203] Kompella, K., Ed. and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, DOI 10.17487/RFC4203, October 2005, <<http://www.rfc-editor.org/info/rfc4203>>.
- [RFC4577] Rosen, E., Psenak, P., and P. Pillay-Esnault, "OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4577, DOI 10.17487/RFC4577, June 2006, <<http://www.rfc-editor.org/info/rfc4577>>.
- [RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", RFC 4915, DOI 10.17487/RFC4915, June 2007, <<http://www.rfc-editor.org/info/rfc4915>>.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <<http://www.rfc-editor.org/info/rfc5340>>.

[RFC6987] Retana, A., Nguyen, L., Zinin, A., White, R., and D. McPherson, "OSPF Stub Router Advertisement", RFC 6987, DOI 10.17487/RFC6987, September 2013, <<http://www.rfc-editor.org/info/rfc6987>>.

Authors' Addresses

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

Email: shraddha@juniper.net

Pushpasis Sarkar
Individual

Email: pushpasis.ietf@gmail.com

Hannes Gredler
Individual

Email: hannes@gredler.at

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

Email: mnanduri@microsoft.com

Luay Jalil
Verizon

Email: luay.jalil@verizon.com