

Internet Engineering Task Force (IETF)  
Request for Comments: 8379  
Category: Standards Track  
ISSN: 2070-1721

S. Hegde  
Juniper Networks, Inc.  
P. Sarkar  
Arrcus, Inc.  
H. Gredler  
RtBrick Inc.  
M. Nanduri  
ebay Corporation  
L. Jalil  
Verizon  
May 2018

## OSPF Graceful Link Shutdown

### 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 value on one side of the link is not sufficient to divert the traffic flowing in the other direction.

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

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

### Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8379>.

Copyright Notice

Copyright (c) 2018 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 (https://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
  - 1.1. Requirements Language . . . . . 3
- 2. Motivation . . . . . 3
- 3. Flooding Scope . . . . . 4
- 4. Protocol Extensions . . . . . 4
  - 4.1. OSPFv2 Graceful-Link-Shutdown Sub-TLV . . . . . 4
  - 4.2. Remote IPv4 Address Sub-TLV . . . . . 4
  - 4.3. Local/Remote Interface ID Sub-TLV . . . . . 5
  - 4.4. OSPFv3 Graceful-Link-Shutdown Sub-TLV . . . . . 6
  - 4.5. BGP-LS Graceful-Link-Shutdown TLV . . . . . 6
  - 4.6. Distinguishing Parallel Links . . . . . 7
- 5. Elements of Procedure . . . . . 8
  - 5.1. Point-to-Point Links . . . . . 8
  - 5.2. Broadcast/NBMA Links . . . . . 9
  - 5.3. Point-to-Multipoint Links . . . . . 10
  - 5.4. Unnumbered Interfaces . . . . . 10
  - 5.5. Hybrid Broadcast and P2MP Interfaces . . . . . 10
- 6. Backward Compatibility . . . . . 10
- 7. Applications . . . . . 11
  - 7.1. Overlay Network . . . . . 11
  - 7.2. Controller-Based Deployments . . . . . 12
  - 7.3. L3VPN Services and Sham Links . . . . . 13
  - 7.4. Hub and Spoke Deployment . . . . . 13
- 8. Security Considerations . . . . . 13
- 9. IANA Considerations . . . . . 14
- 10. References . . . . . 14
  - 10.1. Normative References . . . . . 14
  - 10.2. Informative References . . . . . 16
- Acknowledgements . . . . . 16
- Authors’ Addresses . . . . . 17

## 1. Introduction

This document describes a mechanism for gracefully taking a link out of service while allowing it to be used if no other path is available. It also provides a mechanism to divert the traffic from both directions of the link.

Many OSPFv2 or OSPFv3 deployments run on overlay networks provisioned by means of pseudowires 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 maintenance is actually performed. Since the nodes in the underlying network are not visible to OSPF, the existing stub-router mechanism described in [RFC6987] cannot be used. In a service provider's network, there may be many CE-to-CE connections that run over a single PE. It is cumbersome to change the metric on every CE-to-CE connection in both directions. This document provides a mechanism to change the metric of the link on the remote side and also use the link as a last-resort link if no alternate paths are available. An application specific to this use case is described in detail in Section 7.1.

This document provides mechanisms to advertise graceful-link-shutdown state in the flexible encodings provided by "OSPFv2 Prefix/Link Attribute Advertisement" [RFC7684] and the E-Router-LSA [RFC8362] for OSPFv3. 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.

### 1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

## 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 Label Switched Path (LSP) ingress routers/controllers can learn about the impending maintenance activity and apply specific policies to reroute the LSPs for deployments based on Traffic Engineering (TE).
- 4. Allow the link to be used as a last-resort link to prevent traffic disruption when alternate paths are not available.

3. Flooding Scope

The graceful-link-shutdown information is flooded in an area-scoped Extended Link Opaque LSA [RFC7684] for OSPFv2 and in an E-Router-LSA for OSPFv3 [RFC8362]. The Graceful-Link-Shutdown sub-TLV MAY be processed by the head-end nodes or the controller as described in the Section 7. The procedures for processing the Graceful-Link-Shutdown sub-TLV are described in Section 5.

4. Protocol Extensions

4.1. OSPFv2 Graceful-Link-Shutdown Sub-TLV

The Graceful-Link-Shutdown sub-TLV identifies the link as being gracefully shutdown. It is advertised in the Extended Link TLV of the Extended Link Opaque LSA as defined in [RFC7684].

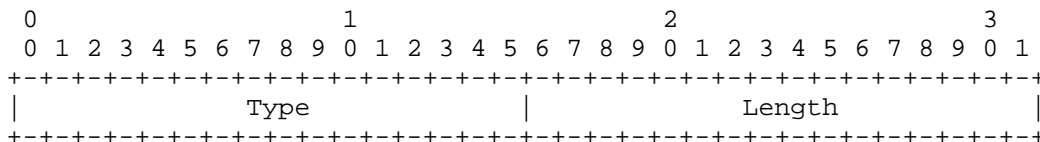


Figure 1: Graceful-Link-Shutdown Sub-TLV for OSPFv2

Type: 7

Length: 0

4.2. Remote IPv4 Address Sub-TLV

This sub-TLV specifies the IPv4 address of the remote endpoint on the link. It is advertised in the Extended Link TLV as defined in [RFC7684]. This sub-TLV is optional and MAY be advertised in an area-scoped Extended Link Opaque LSA to identify the link when there are multiple parallel links between two nodes.

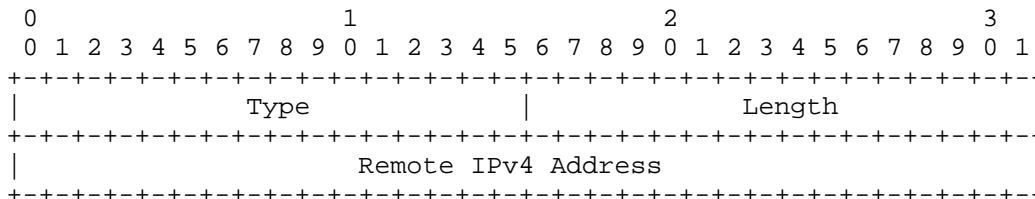


Figure 2: Remote IPv4 Address Sub-TLV

Type: 8

Length: 4

Value: Remote IPv4 address. The remote IPv4 address is used to identify a particular link on the remote side when there are multiple parallel links between two nodes.

4.3. Local/Remote Interface ID Sub-TLV

This sub-TLV specifies Local and Remote Interface IDs. It is advertised in the Extended Link TLV as defined in [RFC7684]. This sub-TLV is optional and MAY be advertised in an area-scoped Extended Link Opaque LSA to identify the link when there are multiple parallel unnumbered links between two nodes. The Local Interface ID is generally readily available. One of the mechanisms to obtain the Remote Interface ID is described in [RFC4203].

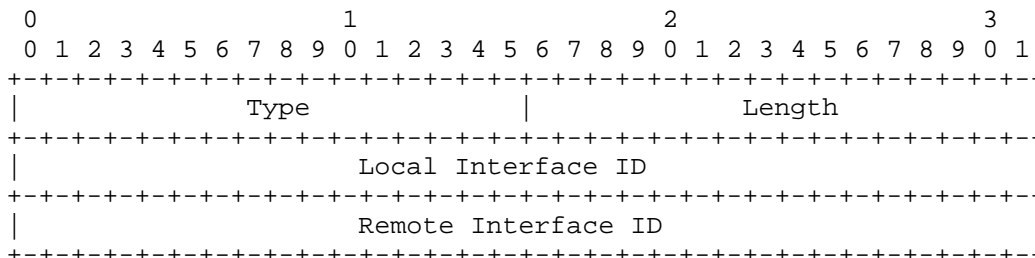


Figure 3: Local/Remote Interface ID Sub-TLV

Type: 9

Length: 8

Value: 4 octets of the Local Interface ID followed by 4 octets of the Remote Interface ID.

4.4. OSPFv3 Graceful-Link-Shutdown Sub-TLV

The Graceful-Link-Shutdown sub-TLV is carried in the Router-Link TLV as defined in [RFC8362] for OSPFv3. The Router-Link TLV contains the Neighbor Interface ID and can uniquely identify the link on the remote node.

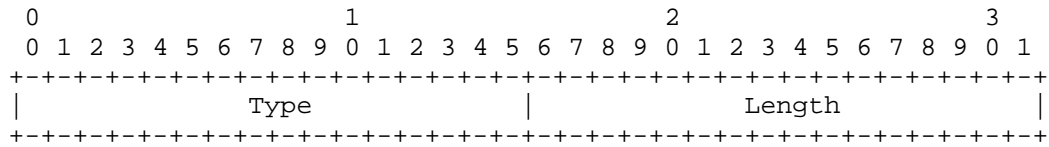


Figure 4: Graceful-Link-Shutdown Sub-TLV for OSPFv3

Type: 8

Length: 0

4.5. BGP-LS Graceful-Link-Shutdown TLV

BGP-LS as defined in [RFC7752] is a mechanism that distributes network information to the external entities using the BGP routing protocol. Graceful link shutdown is important link information that the external entities can use for various use cases as defined in Section 7. BGP Link Network Layer Reachability Information (NLRI) is used to carry the link information. A new TLV called "Graceful-Link-Shutdown" is defined to describe the link attribute corresponding to graceful-link-shutdown state. The TLV format is as described in Section 3.1 of [RFC7752]. There is no Value field, and the Length field is set to zero for this TLV.

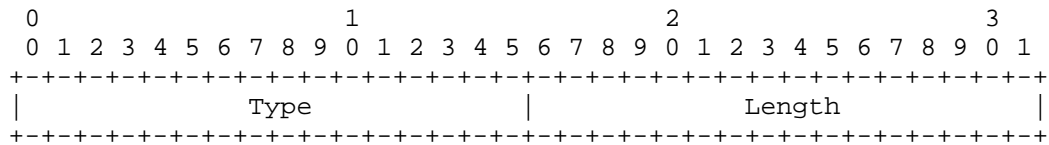


Figure 5: Graceful-Link-Shutdown TLV for BGP-LS

Type: 1121

Length: 0

## 4.6. Distinguishing Parallel Links




Figure 6: Parallel Links

Consider two routers, A and B, connected with two parallel point-to-point interfaces. I.w and I.x represent the interface address on Router A's side, and I.y and I.z represent interface addresses on Router B's side. The Extended Link Opaque LSA as defined in [RFC7684] describes links using Link Type, Link ID, and Link Data. For example, a link with the address I.w is described as below on Router A.

Link Type = Point-to-point

Link ID = Router ID of B

Link Data = I.w

A third node (controller or head-end) in the network cannot distinguish the interface on Router B, which is connected to this particular Interface on Router A based on the link information described above. The interface with address I.y or I.z could be chosen due to this ambiguity. In such cases, a Remote IPv4 Address sub-TLV should be originated and added to the Extended Link TLV. The use cases as described in Section 7 require controller or head-end nodes to interpret the graceful-link-shutdown information and hence the need for the Remote IPv4 Address sub-TLV. I.y is carried in the Extended Link TLV, which unambiguously identifies the interface on the remote side. The OSPFv3 Router-Link TLV as described in [RFC8362] contains an Interface ID and a neighbor's Interface ID, which can uniquely identify connecting the interface on the remote side; hence, OSPFv3 does not require a separate remote IPv6 address to be advertised along with the OSPFv3 Graceful-Link-Shutdown sub-TLV.

## 5. Elements of Procedure

As defined in [RFC7684], every link on the node will have a separate Extended Link Opaque LSA. The node that has the link to be taken out of service MUST advertise the Graceful-Link-Shutdown sub-TLV in the Extended Link TLV of the Extended Link Opaque LSA for OSPFv2, as defined in [RFC7684], and in the Router-Link TLV of E-Router-LSA for OSPFv3. The Graceful-Link-Shutdown sub-TLV indicates that the link identified by the sub-TLV is subjected to maintenance.

For the purposes of changing the metric OSPFv2 and OSPFv3 Router-LSAs need to be reoriginated. To change the Traffic Engineering metric, TE Opaque LSAs in OSPFv2 [RFC3630] and Intra-area-TE-LSAs in OSPFv3 [RFC5329] need to be reoriginated.

The graceful-link-shutdown information is advertised as a property of the link and is flooded through 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.

When a link is ready to carry traffic, the Graceful-Link-Shutdown sub-TLV MUST be removed from the Extended Link TLV/Router-Link TLV, and the corresponding LSAs MUST be readvertised. Similarly, the metric MUST be set to original values, and the corresponding LSAs MUST be readvertised.

The procedures described in this document may be used to divert the traffic away from the link in scenarios other than link-shutdown or link-replacement activity.

The precise action taken by the remote node at the other end of the link identified for graceful-shutdown depends on the link type.

### 5.1. Point-to-Point Links

The node that has the link to be taken out of service MUST set the metric of the link to MaxLinkMetric (0xffff) and reoriginate its Router-LSA. The Traffic Engineering metric of the link SHOULD be set to (0xffffffff), and the node SHOULD reoriginate the corresponding TE Link Opaque LSAs. When a Graceful-Link-Shutdown sub-TLV is received for a point-to-point link, the remote node MUST identify the local link that corresponds to the graceful-shutdown link and set its metric to MaxLinkMetric (0xffff), and the remote node MUST reoriginate its Router-LSA with the changed metric. When TE is enabled, the Traffic Engineering metric of the link SHOULD be set to (0xffffffff) and follow the procedures in [RFC5817]. Similarly, the



remote node SHOULD set the Traffic Engineering metric of the link to 0xffffffff and SHOULD reoriginate the TE Link Opaque LSA for the link with the new value.

The Extended Link Opaque LSAs and the Extended Link TLV are not scoped for multi-topology [RFC4915]. In multi-topology deployments [RFC4915], the Graceful-Link-Shutdown sub-TLV advertised in an Extended Link Opaque LSA corresponds to all the topologies that include the link. The receiver node SHOULD change the metric in the reverse direction for all the topologies that include the remote link and reoriginate the Router-LSA as defined in [RFC4915].

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

## 5.2. Broadcast/NBMA Links

Broadcast or Non-Broadcast Multi-Access (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 logically connect. As a result, routers on the broadcast or NBMA network advertise only their adjacency to the DR. Routers that do not act as DRs do not form or advertise adjacencies with each other. For the broadcast links, the MaxLinkMetric on the remote link cannot be changed since all the neighbors are on same link. Setting the link cost to MaxLinkMetric would impact all paths that traverse any of the neighbors connected on that broadcast link.

The node that has the link to be taken out of service MUST set the metric of the link to MaxLinkMetric (0xffff) and reoriginate the Router-LSA. The Traffic Engineering metric of the link SHOULD be set to (0xffffffff), and the node SHOULD reoriginate the corresponding TE Link Opaque LSAs. For a broadcast link, the two-part metric as described in [RFC8042] is used. The node originating the Graceful-Link-Shutdown sub-TLV MUST set the metric in the Network-to-Router Metric sub-TLV to MaxLinkMetric (0xffff) for OSPFv2 and OSPFv3 and reoriginate the corresponding LSAs. 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 (P2MP) links is similar to the point-to-point links. When a Graceful-Link-Shutdown sub-TLV is received for a point-to-multipoint link, the remote node MUST identify the neighbor that corresponds to the graceful-shutdown link and set its metric to MaxLinkMetric (0xffff). The remote node MUST reoriginate the Router-LSA with the changed metric for the corresponding neighbor.

### 5.4. Unnumbered Interfaces

Unnumbered interfaces do not have a unique IP address and borrow their 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 advertising the Graceful-Link-Shutdown sub-TLV in order to handle unnumbered interfaces. The Link-Data field in the Extended Link TLV includes the Local Interface ID instead of the IP address. The Local/Remote Interface ID sub-TLV MUST be advertised when there are multiple parallel unnumbered interfaces between two nodes. One of the mechanisms to obtain the Interface ID of the remote side is 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 operation for the P2MP interfaces. When a Graceful-Link-Shutdown sub-TLV is received for a hybrid link, the remote node MUST identify the neighbor that corresponds to the graceful-shutdown link and set its metric to MaxLinkMetric (0xffff). All the remote nodes connected to the originator MUST reoriginate the Router-LSA with the changed metric for the neighbor.

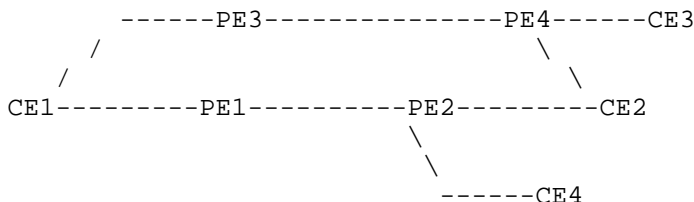
## 6. Backward Compatibility

The mechanisms described in the document are fully backward compatible. It is required that the node advertising the Graceful-Link-Shutdown sub-TLV as well as the node at the remote end of the graceful-shutdown link support the extensions described herein for the traffic to be diverted from the graceful-shutdown link. If the remote node doesn't support the capability, it will still use the graceful-shutdown link, but there are no other adverse effects. In the case of broadcast links using two-part metrics, the backward-compatibility procedures as described in [RFC8042] are applicable.

7. Applications

7.1. Overlay Network

Many service providers offer L2 services to a customer connecting different locations. The customer’s IGP protocol creates a seamless private network (overlay network) across the locations for the customer. Service providers want to offer graceful-shutdown functionality when the PE device is taken out for maintenance. There can be large number of customers attached to a PE node, and the remote endpoints for these L2 attachment circuits are spread across the service provider’s network. Changing the metric for all corresponding L2 circuits in both directions is a tedious and error-prone process. The graceful-link-shutdown feature simplifies the process by increasing the metric on the CE-CE overlay link so that traffic in both directions is diverted away from the PE undergoing maintenance. The graceful-link-shutdown feature allows the link to be used as a last-resort link so that traffic is not disrupted when alternate paths are not available.



CE: Customer Edge  
 PE: Provider Edge

Figure 7: Overlay Network

In the example shown in Figure 7, when the PE1 node is going out of service for maintenance, a service provider sets the PE1 to stub-router state and communicates the pending maintenance action to the overlay customer networks. The mechanisms used to communicate between PE1 and CE1 is outside the scope of this document. CE1 sets the graceful-link-shutdown state on its links connecting CE3, CE2, and CE4, changes the metric to MaxLinkMetric, and reoriginates the corresponding LSA. The remote end of the link at CE3, CE2, and CE4 also set the metric on the link to MaxLinkMetric, and the traffic from both directions gets diverted away from PE1.

7.2. Controller-Based Deployments

In controller-based deployments where the controller participates in the IGP protocol, the controller can also receive the graceful-link-shutdown information as a warning that link maintenance is imminent. Using this information, the controller can find alternate paths for traffic that uses the affected link. The controller can apply various policies and reroute the LSPs away from the link undergoing maintenance. If there are no alternate paths satisfying the constraints, the controller might temporarily relax those constraints and put the service on a different path. Increasing the link metric alone does not specify the maintenance activity as the metric could increase in events such as LDP-IGP synchronization. An explicit indication from the router using the Graceful-Link-Shutdown sub-TLV is needed to inform the controller or head-end routers.

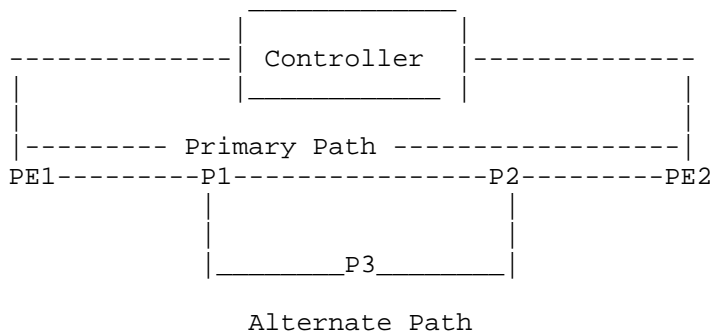


Figure 8: Controller-Based Traffic Engineering

In the above example, the 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 10 Gbps. When the P1->P2 link is being prepared for maintenance, the controller receives the graceful-link-shutdown information, as there is no alternate path available that satisfies the constraints, and the 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 Layer 3 Virtual Private Network (L3VPN) services to customers, and CE-PE links run OSPF [RFC4577]. When the PE is taken out of service for maintenance, all the links on the PE can be set to graceful-link-shutdown state, which will guarantee that the traffic to/from dual-homed CEs gets diverted. The interaction between OSPF and BGP is outside the scope of this document. A mechanism based on [RFC6987] with summaries and externals that are advertised with high metrics could also be used to achieve the same functionality when implementations support high metrics advertisement for summaries and externals.

Another useful use case is when ISPs provide sham-link services to customers [RFC4577]. When the PE goes out of service for maintenance, all sham links on the PE can be set to graceful-link-shutdown 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 large 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. The mechanism defined in [RFC6987] can be used to divert the Spoke-to-Spoke traffic from the overloaded Hub router. The traffic that flows from Spokes via the Hub into an external network may not be diverted in certain scenarios. When a Hub node goes down for maintenance, all links on the Hub can be set to graceful-link-shutdown state, and traffic gets diverted from the Spoke sites as well without having to make configuration changes on the Spokes.

## 8. Security Considerations

This document utilizes the OSPF packets and LSAs described in [RFC2328], [RFC3630], [RFC5329], and [RFC5340]. The authentication procedures described in [RFC2328] for OSPFv2 and [RFC4552] for OSPFv3 are applicable to this document as well. This document does not introduce any further security issues other than those discussed in [RFC2328] and [RFC5340].

## 9. IANA Considerations

IANA has registered the following in the "OSPFv2 Extended Link TLV Sub-TLVs" registry:

- 7 - Graceful-Link-Shutdown Sub-TLV
- 8 - Remote IPv4 Address Sub-TLV
- 9 - Local/Remote Interface ID Sub-TLV

IANA has registered the following value in the "OSPFv3 Extended-LSA Sub-TLVs" registry:

- 8 - Graceful-Link-Shutdown sub-TLV

IANA has registered the following value in the "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" registry [RFC7752]":

- 1121 - Graceful-Link-Shutdown TLV

## 10. References

### 10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <<https://www.rfc-editor.org/info/rfc2328>>.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <<https://www.rfc-editor.org/info/rfc3630>>.
- [RFC5329] Ishiguro, K., Manral, V., Davey, A., and A. Lindem, Ed., "Traffic Engineering Extensions to OSPF Version 3", RFC 5329, DOI 10.17487/RFC5329, September 2008, <<https://www.rfc-editor.org/info/rfc5329>>.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <<https://www.rfc-editor.org/info/rfc5340>>.

- [RFC5817] Ali, Z., Vasseur, JP., Zamfir, A., and J. Newton, "Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks", RFC 5817, DOI 10.17487/RFC5817, April 2010, <<https://www.rfc-editor.org/info/rfc5817>>.
- [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, <<https://www.rfc-editor.org/info/rfc6845>>.
- [RFC6987] Retana, A., Nguyen, L., Zinin, A., White, R., and D. McPherson, "OSPF Stub Router Advertisement", RFC 6987, DOI 10.17487/RFC6987, September 2013, <<https://www.rfc-editor.org/info/rfc6987>>.
- [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, <<https://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, <<https://www.rfc-editor.org/info/rfc7752>>.
- [RFC8042] Zhang, Z., Wang, L., and A. Lindem, "OSPF Two-Part Metric", RFC 8042, DOI 10.17487/RFC8042, December 2016, <<https://www.rfc-editor.org/info/rfc8042>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8362] Lindem, A., Roy, A., Goethals, D., Reddy Vallem, V., and F. Baker, "OSPFv3 Link State Advertisement (LSA) Extensibility", RFC 8362, DOI 10.17487/RFC8362, April 2018, <<https://www.rfc-editor.org/info/rfc8362>>.

## 10.2. Informative References

- [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, <<https://www.rfc-editor.org/info/rfc4203>>.
- [RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", RFC 4552, DOI 10.17487/RFC4552, June 2006, <<https://www.rfc-editor.org/info/rfc4552>>.
- [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, <<https://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, <<https://www.rfc-editor.org/info/rfc4915>>.

## Acknowledgements

Thanks to Chris Bowers for valuable input and edits to the document. Thanks to Jeffrey Zhang, Acee Lindem, and Ketan Talaulikar for their input. Thanks to Karsten Thomann for careful review and input on the applications where graceful link shutdown is useful.

Thanks to Alia Atlas, Deborah Brungard, Alvaro Retana, Andrew G. Malis, and Tim Chown for their valuable input.



## Authors' Addresses

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

Email: shraddha@juniper.net

Pushpasis Sarkar  
Arrcus, Inc.

Email: pushpasis.ietf@gmail.com

Hannes Gredler  
RtBrick Inc.

Email: hannes@rtbrick.com

Mohan Nanduri  
ebay Corporation  
2025 Hamilton Avenue  
San Jose, CA 98052  
United States of America

Email: mnanduri@ebay.com

Luay Jalil  
Verizon

Email: luay.jalil@verizon.com