

Internet Engineering Task Force (IETF)
Request for Comments: 8042
Updates: 2328
Category: Standards Track
ISSN: 2070-1721

Z. Zhang
L. Wang
Juniper Networks, Inc.
A. Lindem
Cisco Systems
December 2016

OSPF Two-Part Metric

Abstract

This document specifies an optional OSPF protocol extension to represent router metrics in a multi-access network in two parts: the metric from the router to the network and the metric from the network to the router. For such networks, the router-to-router metric for OSPF route computation is the sum of the two parts. This document updates RFC 2328.

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 <http://www.rfc-editor.org/info/rfc8042>.

Copyright Notice

Copyright (c) 2016 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	2
1.1.	Requirements Language	3
2.	Proposed Enhancement	3
3.	Specifications	4
3.1.	Router Interface Parameters	4
3.2.	Advertising Network-to-Router Metric in OSPFv2	4
3.3.	Advertising Network-to-Router Traffic Engineering (TE) Metric	5
3.4.	Advertising Network-to-Router Metric in OSPFv3	5
3.5.	OSPF Stub Router Behavior	5
3.6.	SPF Calculation	5
3.7.	Backward Compatibility	6
4.	IANA Considerations	6
5.	Security Considerations	6
6.	References	7
6.1.	Normative References	7
6.2.	Informative References	7
	Acknowledgements	8
	Contributors	8
	Authors' Addresses	9

1. Introduction

With Open Shortest Path First (OSPF) [RFC2328] [RFC5340]), a Network-LSA (Link State Advertisement) is advertised to list all routers on a broadcast network. Additionally, each router on the broadcast network includes a link in its Router-LSA to describe its connection to the network. The link in the Router-LSA includes a metric but the listed routers in the Network-LSA do not include a metric. This is based on the assumption that from a particular router, all others on the same network can be reached with the same metric.

With some broadcast networks, different routers can be reached with different metrics. [RFC6845] extends the OSPF protocol with a hybrid interface type for that kind of broadcast network, where no Network-LSA is advertised and Router-LSAs simply include point-to-point links to all routers on the same network with individual metrics. Broadcast capability is still used to optimize database synchronization and adjacency maintenance.

This works well for broadcast networks where the metric between different pairs of routers are really independent, for example, Virtual Private LAN Service (VPLS) networks.

With certain types of broadcast networks, further optimization can be made to reduce the size of Router-LSAs and the number of updates.

Consider a satellite radio network with fixed and mobile ground terminals. All communication goes through the satellite. When the mobile terminals move about, their communication capability may change. When OSPF runs over the radio network, [RFC6845] hybrid interface can be used, but with the following drawbacks.

Consider that one terminal/router moves into an area where its communication capability degrades significantly. Through the radio control protocol, all other routers determine that the metric to this particular router changed and they all need to update their Router-LSAs accordingly. In addition, the router in question determines that its metric to reach all others also changed and it needs to update its Router-LSA. Consider that there could be many terminals and many of them can be moving fast and frequently. The number and frequency of updates of those large Router-LSAs could inhibit network scaling.

1.1. 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 [RFC2119].

2. Proposed Enhancement

Notice that in the above scenario, when one terminal's communication capability changes, its metric to all other terminals and the metric to it from all other terminals will all change in a similar fashion. Given this, the above problem can be easily addressed by breaking the metric into two parts: the metric to the satellite and the metric from the satellite. The metric from terminal R1 to R2 would be the sum of the metric from R1 to the satellite and the metric from the satellite to R2.

Instead of using the hybrid interface type described in [RFC6845], the network is treated as a regular broadcast network. A router on the network no longer lists individual metrics to each neighbor in its Router-LSA. Instead, each router advertises the metric from the network to itself in addition to the normal metric for the network. With the normal Router-to-Network and additional Network-to-Router metrics advertised for each router, individual Router-to-Router metrics can be calculated.

With the proposed enhancement, the size of the Router-LSA will be significantly reduced. In addition, when a router's communication capability changes, only that router needs to update its Router-LSA.

Note that while the example uses the satellite as the relay point at the radio level (layer 2), the satellite does not participate in packet forwarding at layer 3. In fact, the satellite does not need to run any layer-3 protocol. Therefore, for generality, the metric is abstracted as to/from the "network" rather than specifically to/from the "satellite".

3. Specifications

The following specifications are added to or modified from the base OSPF protocol. If an area contains one or more two-part metric networks, then all routers in the area MUST support the extensions specified herein. This is ensured by procedures described in Section 3.7.

3.1. Router Interface Parameters

The "Router interface parameters" have the following additions:

- o Two-part metric: TRUE if the interface connects to a multi-access network that uses a two-part metric. All routers connected to the same network SHOULD have the same configuration for their corresponding interfaces.
- o Interface input cost: Link-state metric from the two-part-metric network to this router. Defaults to "Interface output cost" but is not valid for normal networks using a single metric. May be configured or dynamically adjusted to a value different from the "Interface output cost".

3.2. Advertising Network-to-Router Metric in OSPFv2

For OSPFv2, the Network-to-Router metric is encoded in an OSPF Extended Link TLV Sub-TLV [RFC7684], defined in this document as the Network-to-Router Metric Sub-TLV. The type of the sub-TLV is 4. The length of the sub-TLV is 4 (for the value part only). The value part of the sub-TLV is defined as follows:

```

      0                               1                               2                               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
      +-----+-----+-----+-----+-----+-----+-----+-----+
      |           MT-ID           |           0           |           MT Metric           |
      +-----+-----+-----+-----+-----+-----+-----+-----+

```

Multiple such sub-TLVs can exist in a single OSPF Extended Link TLV, one for each topology [RFC4915]. Each sub-TLV will have a unique Multi-Topology Identifier (MT-ID) and will adhere to the advertisement rules defined in Section 3.4 of [RFC4915]. The OSPF

Extended Link TLV identifies the transit link to the network and is part of an OSPFv2 Extended-Link Opaque LSA. The sub-TLV MUST ONLY appear in Extended-Link TLVs for Link Type 2 (link to transit network) and MUST be ignored if received for other link types.

3.3. Advertising Network-to-Router Traffic Engineering (TE) Metric

A Traffic Engineering Network-to-Router Metric Sub-TLV is defined, similar to the Traffic Engineering Metric Sub-TLV defined in Section 2.5.5 of [RFC3630]. The only difference is the TLV type, which is 35. The sub-TLV MUST only appear in Type 2 Link TLVs (Multi-access) of Traffic Engineer LSAs (OSPF2) or Intra-Area-TE-LSAs (OSPFv3) [RFC5329], and MUST appear at most once in such a Link TLV.

3.4. Advertising Network-to-Router Metric in OSPFv3

Network-to-Router metric advertisement in OSPFv3 Extended Router-LSA [OSPFV3-EXTENDED-LSA] will be described in a separate document.

3.5. OSPF Stub Router Behavior

When an OSPF router with interfaces to multi-access networks using two-part metrics is advertising itself as a stub router [RFC6987], only the Router-to-Network metric in the stub router's OSPF Router-LSA links for those networks is set to the MaxLinkMetric. This is fully backward compatible and will result in the same behavior as described in [RFC6987].

3.6. SPF Calculation

The first stage of the shortest-path tree calculation is described in Section 16.1 of [RFC2328]. With a two-part metric, when a vertex V corresponding to a Network-LSA has just been added to the Shortest Path Tree (SPT) and an adjacent vertex W (joined by a link in V's corresponding Network-LSA) is being added to the candidate list, the cost from V to W (W's network-to-router cost) is determined as follows:

- o For OSPFv2, if vertex W has a corresponding Extended-Link Opaque LSA with an Extended Link TLV for the link from W to V, and the Extended Link TLV has a Network-to-Router Metric Sub-TLV for the corresponding topology, then the cost from V to W is the metric in the sub-TLV. Otherwise, the cost is 0.
- o OSPFv3 [RFC5340] Shortest Path First (SPF) changes will be described in a separate document.

3.7. Backward Compatibility

Due to the change of procedures in the SPF calculation, all routers in an area that includes one or more two-part metric networks must support the changes specified in this document. To ensure that, if an area is provisioned to support two-part metric networks, all routers supporting this capability must advertise a Router Information (RI) LSA with a Router Functional Capabilities TLV [RFC7770] that includes the following Router Functional Capability Bit:

Bit	Capabilities
6	Two-Part Metric support

Upon detecting the presence of a reachable Router-LSA without a companion RI LSA that has the bit set, all routers MUST recalculate routes without considering any network-to-router costs.

4. IANA Considerations

IANA has made the following assignments per this document:

- o Two-Part Metric support (6) was added to the "OSPF Router Informational Capability Bits" registry.
- o Network-to-Router Metric Sub-TLV (4) has been added to the "OSPFv2 Extended Link TLV Sub-TLVs" registry.
- o Network-to-Router TE Metric Sub-TLV (35) has been added to the "Types for sub-TLVs of TE Link TLV (Value 2)" registry.

5. Security Considerations

This document does not introduce new security risks. Existing security considerations in OSPFv2 and OSPFv3 apply.

6. References

6.1. 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, <<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>>.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <<http://www.rfc-editor.org/info/rfc3630>>.
- [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>>.
- [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, <<http://www.rfc-editor.org/info/rfc5329>>.
- [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>>.
- [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>>.

6.2. Informative References

- [OSPFV3-EXTENDED-LSA] Lindem, A., Mirtorabi, S., and A. Roy, "OSPFv3 LSA Extendibility", Work in Progress, draft-ietf-ospf-ospfv3-lsa-extend-13.txt, October 2016.
- [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>>.

- [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>>.
- [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>>.

Acknowledgements

The authors would like to thank Abhay Roy, Hannes Gredler, Peter Psenak, and Eric Wu for their comments and suggestions.

Contributors

David Dubois
General Dynamics C4S
400 John Quincy Adams Road
Taunton, MA 02780
United States of America
Email: dave.dubois@gd-ms.com

Vibhor Julka
Individual Contributor

Email: vjulka1@yahoo.com

Tom McMillan
L3 Communications, Linkabit
9890 Towne Centre Drive
San Diego, CA 92121
United States of America
Email: tom.mcmillan@l-3com.com

Authors' Addresses

Zhaohui Zhang
Juniper Networks, Inc.
10 Technology Park Drive
Westford, MA 01886
United States of America

Email: z Zhang@juniper.net

Lili Wang
Juniper Networks, Inc.
10 Technology Park Drive
Westford, MA 01886
United States of America

Email: liliw@juniper.net

Acee Lindem
Cisco Systems
301 Midenhall Way
Cary, NC 27513
United States of America

Email: acee@cisco.com

