

Open Shortest Path First
Internet-Draft
Updates: [2328](#), [5340](#) (if approved)
Intended status: Standards Track
Expires: April 17, 2015

Z. Zhang
L. Wang
Juniper Networks, Inc.
A. Lindem
Cisco Systems
D. Dubois
General Dynamics C4S
V. Julka
T. McMillan
L3 Communications, Linkabit
October 14, 2014

OSPF Two-part Metric
draft-zzhang-ospf-two-part-metric-05.txt

Abstract

This document specifies an optional extension to the OSPF protocol, to represent the metric on a multi-access network as two parts: the metric from a router to the network, and the metric from the network to the router. The router to router metric would be the sum of the two.

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](#).

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 April 17, 2015.

Copyright Notice

Copyright (c) 2014 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.	Proposed Enhancement	3
3.	Specifications	4
3.1.	Router Interface Parameters	4
3.2.	Advertising Network-to-Router metric in OSPFv2	5
3.3.	Advertising Network-to-Router metric in OSPFv3	5
3.4.	SPF Calculation	5
3.5.	Backward Compatibility	6
4.	IANA Considerations	6
5.	Security Considerations	6
6.	Acknowledgements	6
7.	References	7
7.1.	Normative References	7
7.2.	Informative References	8

1. Introduction

For a broadcast network, a Network-LSA is advertised to list all routers on the network, and each router on the 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. [RFC 6845](#) 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 p2p links to all routers on the same network with individual metrics. Broadcast capability is still utilized to optimize database synchronization and adjacency maintenance.

That works well for broadcast networks where the metric between different pair of routers are really independent. For example, VPLS networks.

With certain types of broadcast networks, further optimization can be made to reduce the size of the Router-LSAs and 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 (routers being or in tandem with the terminals), [RFC 6845](#) 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. The router in question also determines that its metric to reach all others also changed and it also 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/frequency of updates of those large Router-LSAs could inhibit network scaling.

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 from all other terminals to it 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.

Now instead of using the [RFC 6845](#) hybrid interface type, the network is just 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 metric can be calculated.

With the proposed enhancement, the size of 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), at layer-3, the satellite does not participate in packet forwarding. In fact, the satellite does not need to be running 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 protocol 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 3.5](#).

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 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. Defaulted to "Interface output cost" but 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 [ietf-ospf-lsa-extend], defined in this document as the Network-to-Router Metric Sub-TLV. The type of the Sub-TLV is TBD. 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           |           0           |           MT metric           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Multiple such Sub-TLVs can exist in a single OSPF Extended Link TLV, one for each topology. 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 metric in OSPFv3

For OSPFv3, the same Network-to-Router Metric Sub-TLV definition is used, though it is part of the Router-Link TLV of E-Router-LSA [ietf-ospf-ospfv3-lsa-extend]. Currently OSPFv3 Multi-Topology is not defined so the only valid value for the MT field is 0 and only one such Sub-TLV SHOULD be included in the Router-Link TLV. Received Sub-TLVs with non-zero MT field MUST be ignored.

Similarly, the Sub-TLV MUST ONLY appear in Router-Link TLVs for Link Type 2 (connection to a transit network) and MUST be ignored if received for other link types.

3.4. SPF Calculation

During the first stage of shortest-path tree calculation for an area, when a vertex V corresponding to a Network-LSA is added to the shortest-path tree and its adjacent vertex W (joined by a link in V's corresponding Network LSA), the cost from V to W, which is 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 For OSPFv3, if vertex W has a corresponding E-Router-LSA with a Router-Link TLV for the link from W to V, and the Router-Link TLV has a Network-to-Router Metric Sub-TLV, then the cost from V to W is the metric in the Sub-TLV. If not, the cost is 0.

3.5. 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 [acee-ospf-rfc4970bis] that includes the following Router Functional Capability Bit:

Bit	Capabilities
0	OSPF Two-part Metric [TPM]

Upon detecting the presence of a reachable Router-LSA without a companion RI LSA that has the bit set, all routers MUST disable the two-part metric functionalities and take the following actions:

- o If this router currently advertises network-to-router costs, remove the Network-to-Router Metric Sub-TLVs. This may lead to removal of parent TLVs and even withdrawal of the parent LSAs.
- o Recalculate routes w/o considering any network-to-router costs.

4. IANA Considerations

This document requests IANA to assign a new bit in the Router Functional Capabilities TLV to indicate the capability of supporting two-part metric, a new Sub-TLV in the OSPF Extended-Link TLV Sub-TLV Registry, and a new Sub-TLV in the The OSPFv3 Extend-LSA Sub-TLV registry.

5. Security Considerations

This document does not introduce new security risks.

6. Acknowledgements

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

7. References

7.1. Normative References

- [I-D.acee-ospf-rfc4970bis] Lindem, A., Shen, N., Vasseur, J., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", [draft-acee-ospf-rfc4970bis-00](#) (work in progress), July 2014.
- [I-D.ietf-ospf-lsa-extend] Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 LSA Extendibility", [draft-ietf-ospf-lsa-extend-00](#) (work in progress), August 2014.
- [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-04](#) (work in progress), September 2014.
- [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.
- [RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", [RFC 4915](#), June 2007.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", [RFC 5340](#), July 2008.
- [RFC5613] Zinin, A., Roy, A., Nguyen, L., Friedman, B., and D. Yeung, "OSPF Link-Local Signaling", [RFC 5613](#), August 2009.

7.2. Informative References

[RFC6845] Sheth, N., Wang, L., and J. Zhang,
"OSPF Hybrid Broadcast and Point-
to-Multipoint Interface Type",
[RFC 6845](#), January 2013.

Authors' Addresses

Jeffrey Zhang
Juniper Networks, Inc.
10 Technology Park Drive
Westford, MA 01886

EMail: zzhang@juniper.net

Lili Wang
Juniper Networks, Inc.
10 Technology Park Drive
Westford, MA 01886

EMail: liliw@juniper.net

Acee Lindem
Cisco Systems
301 Midenhall Way
Cary, NC 27513

EMail: acee@cisco.com

David Dubois
General Dynamics C4S
400 John Quincy Adams Road
Taunton, MA 02780

EMail: dave.dubois@gdc4s.com

Vibhor Julka
L3 Communications, Linkabit
9890 Towne Centre Drive
San Diego, CA 92121

EMail: vibhor.julka@l-3Com.com

Tom McMillan
L3 Communications, Linkabit
9890 Towne Centre Drive
San Diego, CA 92121

EMail: tom.mcmillan@l-3com.com