

Open Shortest Path First  
Internet-Draft  
Updates: [2328](#), [5340](#) (if approved)  
Intended status: Standards Track  
Expires: August 17, 2014

Z. Zhang  
L. Wang  
Juniper Networks, Inc.  
D. Dubois  
General Dynamics C4S  
V. Julka  
T. McMillan  
L3 Communications, Linkabit  
February 13, 2014

**OSPF Two-part Metric**  
**draft-zzhang-ospf-two-part-metric-01.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 August 17, 2014.

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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">3</a>
<a href="#">2.</a>	Proposed Enhancement . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Spetications . . . . .	<a href="#">4</a>
<a href="#">4.</a>	IANA Considerations . . . . .	<a href="#">6</a>
<a href="#">5.</a>	Security Considerations . . . . .	<a href="#">6</a>
<a href="#">6.</a>	References . . . . .	<a href="#">7</a>
<a href="#">6.1.</a>	Normative References . . . . .	<a href="#">7</a>
<a href="#">6.2.</a>	Informative References . . . . .	<a href="#">7</a>



## **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 does 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 networks, with which no Network LSA is used and routers simply includes 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 on which 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 go 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 communication capability degrades significantly. Through the radio control protocol all other routers determine that the metric to this particular one 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 need 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 become inhibiting.

## **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 [RFC6845](#) hybrid interface type, the network is just treated as a regular broadcast one. A router on the network no longer needs to list individual metrics to each neighbors in its Router LSA. In case of symmetric metric to/from the satellite, it is represented by the transit link's metric in the Router LSA. In case of asymmetric metric, it is rerepresented by a special MT Metric ([Section 3](#)).

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 radio level (layer 2), at layer 3 the satellite does not play any role. It does not need to be running layer 3 protocol at all. 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 here. This document does not currently specify a way to detect a router's capability of supporting this, and relies on operator's due diligence in provisioning. A protocol mechanism may be developed in the future.

The "Router interface parameters" has the following additions:

- o Two-part metric: TRUE if the interface connects to a multi-access network that uses two-part metric.
- o Interface input cost: Link state metric from the network to this router. Defaulted to "Interface output cost". May be configured or dynamically adjusted to a value different from the "Interface output cost". If different from the output cost, it MUST be advertised in addition to the link (output) cost for this interface in the router's Router LSA.

To signal that a network is using Two-part Metric, a new Link Type X (value TBD) is added. It is similar to Type 2, except that the network uses Two-part Metric, and there may be an optional network-



to-router metric (interface input cost).

Link type	Description	Link ID
-----		
X	Link to transit network that uses the Two-part Metric	Interface of Designated Router

For backward compatibility, if Two-Part Metric is to be used in an area, all routers in the area must support it. To ensure this, a new bit 0x00000004 (pending WG consensus and IANA assignment) in the LLS EOF-TLV [[RFC5613](#)] is used:

Bit	Name	Reference
0x00000001	LSDB Resynchronization (LR)	<a href="#">RFC4811</a>
0x00000002	Restart Signal (RS-bit)	<a href="#">RFC4812</a>
0x00000004	Two-Part Metric (TM-bit)	[this document]

Each router must be configured to include in its Hello packets an EOF-TLV with the bit set, and any Hello packet without the EOF-TLV or without the bit set must be ignored. With this, a router that does not support Two-Part Metric will never receive a Router LSA containing the new Link Type X.

For the network-to-router metric (interface input cost) that is different from the router-to-network metric (interface output cost), with OSPFv2 it is carried in an MT-ID field ([RFC4915](#)):

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Link ID                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Link Data                             |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Type  | # MT-ID |          metric          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  MT-ID |          |1|          MT-ID metric |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               ...                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

With OSPFv3, it is carried in a Router Multi-Topology sub-TLV:





[illegible]

As illustrated above for both cases, the least significant bit of the field between MT-ID and MT-ID metric fields can be set to indicate that this is the network-to-router metric for the corresponding topology. It is clear that this scheme is compatible with multi-topology.

During intra-area SPF calculation, when a vertex W corresponding to a Network LSA is added to the candidate list because of a Type X link in a Router LSA for vertex V (that was just added to the shortest-path tree), W is marked that it uses Two-part Metric. Later, when a vertex V marked with Two-part Metric (which must correspond to a Network LSA) is added to the shortest-path tree, for the vertex W that is reached via a link in V's corresponding LSA, the exact reverse link (of Type X) from W to V is located from W's corresponding Router LSA. If the reverse link does not exist, W is not considered and the next link in V is checked. If the reverse link has a network-to-router metric, that metric is used as the link cost between V and W. Otherwise, that reverse link's (default) metric is used as the link cost between V and W.

#### 4. IANA Considerations

This document makes no request of IANA (unless the Link Type values in Router LSAs are assigned by IANA).

Note to RFC Editor: this section may be removed on publication as an RFC.

## 5. Security Considerations

This document does not introduce new security risks.

## 6. References



### **6.1. Normative References**

- [I-D.ietf-ospf-mt-ospfv3] Mirtorabi, S. and A. Roy, "Multi-topology routing in OSPFv3 (MT-OSPFv3)", [draft-ietf-ospf-mt-ospfv3-03](#) (work in progress), July 2007.
- [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-01](#) (work in progress), February 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.

### **6.2. 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: lililw@juniper.net

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

