

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: April 17, 2014

G. Yan  
Y. Liu  
X. Zhang  
Huawei Technologies  
October 14, 2013

OSPF Extensions for Link State Database Synchronization Group  
draft-ylz-ospf-lsdb-sync-group-01

## Abstract

This document describes a special scheme of OSPF Database synchronization by dividing OSPF Routers into different groups and preventing OSPF Routers from different groups to synchronize, this scheme can help to enhance the number of OSPF adjacencies and the capability of deploying OSPF on large network.

## 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 April 17, 2014.

## Copyright Notice

Copyright (c) 2013 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

Internet-Draft

OSPF SYNC GROUP

October 2013

(<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="#">2</a>
<a href="#">2.</a>	Terminology . . . . .	<a href="#">2</a>
<a href="#">3.</a>	Requirement . . . . .	<a href="#">3</a>
<a href="#">3.1.</a>	Hub and Spoke Scenario . . . . .	<a href="#">3</a>
<a href="#">4.</a>	Solution . . . . .	<a href="#">4</a>
<a href="#">4.1.</a>	Synchronization Group . . . . .	<a href="#">4</a>
<a href="#">4.2.</a>	Synchronization Group Adjacency . . . . .	<a href="#">4</a>
<a href="#">4.3.</a>	LSA Synchronization by Group . . . . .	<a href="#">5</a>
<a href="#">4.4.</a>	Route Calculation . . . . .	<a href="#">6</a>
<a href="#">4.5.</a>	Change of Group Member's Group ID . . . . .	<a href="#">6</a>
<a href="#">5.</a>	Protocol Extension . . . . .	<a href="#">6</a>
<a href="#">6.</a>	Protocol Process . . . . .	<a href="#">8</a>
<a href="#">6.1.</a>	Sending . . . . .	<a href="#">8</a>
<a href="#">6.2.</a>	Receiving . . . . .	<a href="#">8</a>
<a href="#">7.</a>	IANA Considerations . . . . .	<a href="#">8</a>
<a href="#">8.</a>	Security Considerations . . . . .	<a href="#">8</a>
<a href="#">9.</a>	Normative References . . . . .	<a href="#">8</a>
	Authors' Addresses . . . . .	<a href="#">9</a>

## [1.](#) Introduction

The OSPF routing protocol has been used more and more widely in IP radio access networks and enterprise networks. The routers in these networks are usually small devices which have limited memory and CPU, and can not be deployed on large network or large number of adjacencies. This document describes a new mechanism for improving this issue.

## [2.](#) Terminology

**Synchronization Group:** Subset of one area in which the OSPF Routers can exchange LSAs with each other. The OSPF Routers from different Synchronization Groups MUST NOT exchange LSAs with each other.

Synchronization Group ID: Identity of Synchronization Group.

Group Member: One role of OSPF in Synchronization Group which has a unique Synchronization Group ID carried in its LSA. Only those

members who are carrying the same Synchronization Group ID SHOULD establish adjacency. Otherwise no adjacency SHOULD be established. The interface of Group Member is called Group Member interface.

Group Master: One role of OSPF in Synchronization Group which should establish adjacencies with any other OSPF Routers ignoring Group IDs. The interface of Group Master is called Group Master interface.

### [3.](#) Requirement

The OSPF routing protocol [[RFC2328](#)] has a mechanism to make sure the LSA database is synchronized. If the LSA database is not synchronized, there may be a loop of routes or a black hole of routes. One OSPF router in the network must exchange link status information with its adjacencies. In some scenarios that there are thousands of LSAs and more than one thousand adjacencies in the network, lots of LSAs need to be transferred on all adjacencies for synchronization. The physical device resources including memory/CPU/line card are the critical factor to extend the size of network.

#### [3.1.](#) Hub and Spoke Scenario




Figure 1 Hub and Spoke Network

Suppose Hub A establishes  $n$  adjacencies and there are  $m$  LSAs in the network. Hub A will transfer (including receiving and sending) these LSAs on its  $n$  interfaces. The total LSAs to transfer is  $O(n*m)$  without considering the number of LSAs which need to be retransmitted. The time to synchronize LSAs to  $n$  adjacencies for Hub

A is critical for route convergence. Hub A need support high performance of transferring packets. In some scenarios, hub and spoke devices are poor hardware condition with limited memory and CPU, especially the spoke devices can not store all LSAs in the network because of less memory. In such case, there are some special requirements:

- 1) The spoke devices are not necessary to learn routes with each other.
- 2) The spoke devices learn default route or gate way route from hub device.

## [4.](#) Solution

This document introduces a new mechanism which supports LSAs to synchronize in part of one area.

### [4.1.](#) Synchronization Group

One area is divided into several parts called Synchronization Groups which are identified by unique ID (Synchronization Group ID).

Two roles of OSPF Routers are defined here: Group Member OSPF Router and Group Master OSPF Router.

Group Member OSPF Routers from the same Synchronization Group SHOULD establish adjacencies and exchange LSAs;

Group Member OSPF Routers from different Synchronization Group MUST NOT establish adjacencies and MUST NOT exchange LSAs with each other.

The interface of Group Member OSPF Router is defined as Group Member Interface; The interface of Group Master OSPF Router is defined as Group Master Interface.

### [4.2.](#) Synchronization Group Adjacency

The procedure of establishing Synchronization Group Adjacency:

- 1) A new TLV called Synchronization Group sub-TLV (see [section 5](#)) can be carried in Router Information Opaque LSA [[RFC4970](#)] and Link-Local

Signaling (LLS for short) [[RFC4813](#)]. Group Member Interface sends Router Information Opaque LSA including Synchronization Group TLV with its related Synchronization Group ID and M bit set to 0; Group Master Interface sends Router Information Opaque LSA including Synchronization Group TLV with its Synchronization Group ID set to 0 and M bit set to 1. (Note: Group Member Routers/Interface SHOULD set Synchronization Group ID as configured value which is not 0. Group Master Routers/Interface SHOULD set Synchronization Group ID as 0.)

2) For Group Member whose Synchronization Group ID is not 0,

a) If Group Member OSPF Router has one or more neighbors of Group Member, none of them don't support Synchronization Group feature, they MUST NOT form normal adjacency/adjacencies.

b) If Group Member OSPF Router has one or more neighbors of Group Member which have the same Synchronization Group ID, they SHOULD form Synchronization Group adjacency/adjacencies; otherwise, they MUST NOT form normal adjacency/adjacencies.

3) For Group Master OSPF Router whose Synchronization Group ID is 0,

a) If Group Master OSPF Router has one or more neighbors of OSPF Routers, but none of them supports Synchronization Group feature, then they can form normal adjacency/adjacencies.

b) If Group Master OSPF Router has one or more neighbors of Group Member OSPF Routers which have the same Synchronization Group ID, they can form Synchronization Group adjacency/adjacencies. If some neighbors have different Synchronization Group IDs, they MUST NOT form normal adjacency/adjacencies.

c) If Group Master OSPF Router has more than one neighbor of OSPF Routers and some of them support Synchronization Group feature while others don't, then they MUST NOT form normal adjacency/adjacencies.

d) Group Master OSPF Routers SHOULD form normal adjacencies with each other on the LAN where no Group Member OSPF exists.

e) If a common LAN is shared by two or more Group Master OSPF Routers and some Group Member OSPF Routers whose Synchronization Group ID are the same, they can form Synchronization Group adjacency/adjacencies.

f) If a common LAN is shared by two or more Group Master OSPF Routers and some Group Member OSPF Routers whose Synchronization Group ID are not the same, they MUST NOT form adjacency/adjacencies.

Synchronization Group Adjacency SHOULD be advertised as normal adjacency.

#### 4.3. LSA Synchronization by Group

Group Member OSPF MUST carry Synchronization Group TLV (see [section 5](#)) in Router Information Opaque LSA and LLS. Group Member OSPF Router just holds the LSAs which are generated by OSPF Routers which in the same Synchronization Group, or by Group Master OSPF Routers, or by OSPF Routers which don't support Synchronization Group feature.

The procedure of packets:

- 1) Group Member Interface and Group Master Interface which has Synchronization Group Adjacency SHOULD send/receive special LSA packets which just contain the entries of the same group and non-group LSAs.
- 2) When Group Member Interface receives some LSAs from packet which is not belong the same group,

- a) If the LSA entry with non-maxage, the LSA entry SHOULD be discard.
- b) If the LSA entry with maxage, the LSA entry SHOULD be processed as described in OSPF [[RFC2328](#)].

The procedure of LSA, When Group Member OSPF A receives OSPF B's LSA x:

- 1) if x is maxage LSA, A SHOULD synchronize LSA x to adjacencies of Group Member. maxage LSA SHOULD be synchronized as described in OSPF [[RFC2328](#)].
- 2) if x is not maxage LSA and A does not receive the Synchronization Group sub-TLV of Router Information Opaque LSA of B before, LSA x cannot make sure which group it belongs to, so LSA x SHOULD not be

sent to adjacencies of Group Member.

3) if x is not maxage LSA and A has received the Router Information Opaque LSA of B,

a) If B does not support Synchronization Group feature, LSA x SHOULD be synchronized to all adjacencies of A.

b) If B supports Synchronization Group feature and has the same Synchronization Group ID, LSA x SHOULD be synchronized to all adjacencies of A; otherwise, LSA x SHOULD NOT be sent to adjacencies of Group Member.

#### [4.4.](#) Route Calculation

OSPF calculates routes based on its LSA database. No change in this document.

#### [4.5.](#) Change of Group Member's Group ID

When a Group Member OSPF wants to change its group ID to another one, it is RECOMMENDED to change system ID of OSPF for a new one at the same time. Group member will generate new LSAs and be synchronized in a new group. The old LSAs will remain in the old group, but it will not be used in route calculation. The remain life of the old LSAs will be reduced until LSAs reach Maxage to flush.

### [5.](#) Protocol Extension

A new TLV called Synchronization Group TLV is defined to be included in Router Information Opaque LSA and LLS. The Opaque LSA transmitted by a interface that belongs to one Synchronization Group MUST include this TLV.

Type TBD

Length # of octets in the value field (2 octets)

Value

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

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Type           |           Length           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Flags   | Reserved   | Synchronization Group ID |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Sub-TLV Type(*)   |   Sub-TLV Length(*)   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Sub-TLV values(*)           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
* - if present

```

Flags (1 octet):

```

  0 1 2 3 4 5 6 7
+---+---+---+---+---+---+
| Reserved | S|M|
+---+---+---+---+---+---+

```

M - Group Master/Member Role bit  
S - subtlv present bit

(Note: Group Master sets M bit to 1 and Group Member sets M bit to 0)

Synchronization Group ID (2 octets): The range of value of Synchronization Group ID is from 0 to 65535. The value of 0 SHALL be used by Group Master only. Synchronization Group ID of Group Member MUST NOT be 0.

Sub-TLV is not defined in this document which could be extended in future.



Synchronization Group TLV is introduced to indicate whether support Synchronization Group feature or not.

#### [6.1.](#) Sending

Synchronization Group TLV MUST be carried in the Router Information Opaque LSA and LLS, and MUST encode just one time in LSA; This TLV is optional to carry for Group Master OSPF Router.

Group Member Interface MUST carry this TLV in its Router Information Opaque LSA and LLS with its related Synchronization Group ID and M bit set to 0;

Group Master Interface sends Router Information Opaque LSA and LLS including synchronization Group TLV with its Synchronization Group ID set to 0 and M bit set to 1.

#### [6.2.](#) Receiving

If there are more than one Synchronization Group TLVs in Router Information Opaque LSA, the first one is selected, others are ignored.

if Synchronization Group ID and M bit are both set to 0 in TLV, the TLV is illegal and SHOULD be ignored.

if Synchronization Group ID is not 0 and M bit is set to 1 in TLV, the TLV is illegal and SHOULD be ignored.

### [7.](#) IANA Considerations

This document requests that IANA allocate from the OSPF TLV Codepoints Registry a new TLV, referred to as the "Synchronization Group" TLV

### [8.](#) Security Considerations

This document does not introduce any new security concerns to OSPF or any other specifications referenced in this document.

### [9.](#) Normative References

[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.

- [RFC4813] Friedman, B., Nguyen, L., Roy, A., Yeung, D., and A. Zinin, "OSPF Link-Local Signaling", [RFC 4813](#), March 2007.
- [RFC4970] Lindem, A., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", [RFC 4970](#), July 2007.

#### Authors' Addresses

Gang Yan  
Huawei Technologies  
Huawei Bld., No.156 Beiqing Rd.  
Beijing 100095  
China

Email: yangang@huawei.com

Yuanjiao Liu  
Huawei Technologies  
Huawei Bld., No.156 Beiqing Rd.  
Beijing 100095  
China

Email: liuyuanjiao@huawei.com

Xudong Zhang  
Huawei Technologies  
Huawei Bld., No.156 Beiqing Rd.  
Beijing 100095  
China

Email: zhangxudong@huawei.com

