

Network Working Group
Internet Draft
Expiration Date: February 2000
File name: [draft-ietf-ospf-mlinks-01.txt](#)

P. Murphy
US Geological Survey
August 2000

Unnumbered OSPF Multiple Area Links

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#). Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

Table Of Contents

1.0	Abstract	1
2.0	Overview	2
2.1	Requirement for Unnumbered OSPF Multiple Area Links	2
2.2	Proposed Solution	3
2.2.1	Secondary Adjacencies	4
2.2.2	Secondary Neighbor Discovery	4
2.2.3	Advertising Secondary Links	5
2.3	Application	5
3.0	Secondary Interfaces	5
3.1	The SecondaryAreas Interface parameter	5
3.2	The Secondary Interface Data Structure	6
3.3	The Secondary Interface State Machine	6
3.4	OSPF Protocol Packet Processing	7
3.5	Designated Router Selection and Function	7
4.0	Synchronizing Secondary Areas Using Mlink Type-9 LSAs	8
5.0	Secondary Neighbors	9
5.1	The Secondary Neighbor Data Structure	9
5.2	The Secondary Neighbor State Machine	10
5.3	Events Triggered by the Primary Neighbor State Machine....	11
5.4	Receiving Hello Packets	11
5.5	Receiving Mlink Type-9 Neighbor Discover LSAs	12
6.0	Advertising Secondary Adjacencies	13
7.0	The Shortest Path First Calculation	14
8.0	Flushing Secondary Adjacencies	14
9.0	Security Considerations	15
10.0	Acknowledgments	15
11.0	References	15
12.0	Authors' Addresses	15
Appendix A	Router-LSAs	16
Appendix B	Mlink Type-9 Neighbor Discover LSAs	20
Appendix C	Configuration Parameters	23

[1.0](#) Abstract

This memo describes an option to the OSPF Version 2 specification which allows multiple areas to share the same OSPF interface and links. One area is always configured as an interface's primary area. The interface's remaining areas are termed secondary and view it as unnumbered. Two border routers adjacent across the same secondary area may forward the area's intra-area traffic across its secondary links. This option applies to standard areas, stub areas, and NSSAs. It works over any OSPF interface. Routers with this option configured are backward compatible with routers running the standard OSPF Version 2 compliant implementation as defined in [RFC 2328](#).


Please send comments to ospf@discuss.microsoft.com.

2.0 Overview

2.1 Requirement for Unnumbered OSPF Multiple Area Links

Large corporate networks in today's modern Internet have tremendous human and equipment resources coupled with sizable budget invested in their backbone infrastructure. Their regional networks are normally not so fortunate and must multi-home to the backbone in order to take advantage of the backbone's faster and more reliable network links. Performance over a T1 link can pale by comparison to performance over a DS3 or OC3 backbone link. Large corporate networks have sizable backbone routing tables and routinely use stub areas and NSSAs. Under the current OSPF specification intra-area routes are always preferred over inter-area routes even when the traffic is sourced from and destined to the same non-backbone OSPF area.

Consider the following typical OSPF example:



where A0 and B0 are border routers attached to NSSA 1. A1 and B1 are internal routers in NSSA 1. N1 and M1 are standard ethernet networks in NSSA 1 which are directly attached to B0 and B1 respectively. The cost of each link is shown in parentheses. All OSPF costs are symmetric. Under the current OSPF specification the preferred path between A0 and M1 is

A0<->A1<->B1<->M1

with an OSPF cost of 86, even though there exists a more preferred path through B0 namely

A0<->B0<->B1<->M1

with an OSPF cost of 31.

In addition the NSSA 1 OSPF preferred path between A1 and N1 is

A1<->B1<->B0<->N1,

with an OSPF cost of 86, even though there exists a more preferred path through the link between A0 and B0, namely

A1<->A0<->B0<->N1

with an OSPF cost of 31. Under the current OSPF specification NSSA 1's router A1 does not even see this path to N1 since it has no knowledge of Area 0's topology. A0, on the other hand, would at least see the summary LSA of N1, but still cannot take advantage of it due to OSPF intra-area path preference.

What is needed is a tool which makes the Area 0 link between A0 and B0 visible inside NSSA 1. A common practice is to use multiple interface subnets. However this is not an option when the path between A0 and B0 is unnumbered or when one desires that the NSSA 1 path between A0 and B0 be unnumbered. Moreover when configuring multiple interface subnets over multi-access networks, inverse-arp limitations come into play and additional layer 2 PVCs may be required imposing potential resource and budgetary ramifications. The existing tools for the multiple area usage of the same physical link are awkward to configure, implementation dependent, unnecessarily complex and unwieldy to maintain. These tools also burden the physical link with the simultaneous database exchange of multiple OSPF interfaces during adjacency formation.

The Unnumbered OSPF Multiple Area Links option is a simpler tool for configuring multiple area links, requiring just a simple list of the areas sharing an OSPF link. Furthermore it prioritizes database exchange, giving preference to the primary area and Area 0 over other non-backbone secondary areas. If, in the above example, the link between A0 and B0 were part of NSSA 1, path preference would be optimal as the DS3 path would be intra-area for NSSA 1.

2.2 Proposed Solution

The Unnumbered OSPF Multiple Area Links option lets multiple areas use the same link between two border routers for intra-area traffic. Traffic may originate from inside each of the configured areas, as every router in the configured areas sees the link as part of its link state topology. Routers with configured unnumbered OSPF multiple area links must be in Area 0, although the connection to Area 0 may be an unnumbered OSPF multiple area link.

The current OSPF specification only allows one area per OSPF interface. Thus, should an Area L dual-home to Area 0 via two border routers which are adjacent over another Area K's router<->router link or router<->network link, the adjacent link over Area K is not seen inside Area L, even though the two border routers exist physically

adjacent within Area L. This, coupled with intra-area route preference, prevents Area L from utilizing Area K's adjacent path.

2.2.1 Secondary Adjacencies

The softening of this restriction is facilitated by the addition of a new interface configuration parameter called `SecondaryAreas`. This parameter specifies a list of additional areas in which an OSPF interface also belongs (See [Appendix C](#)). The areas listed in this parameter are called the interface's secondary areas, as opposed to the interface's configured area, hereafter called the primary area. A separate interface data structure is generated for each of the secondary areas. For each secondary area listed in the `SecondaryAreas` parameter a router can form OSPF adjacencies with the primary area's neighboring routers across the interface's directly attached network or router link. These adjacencies are called secondary adjacencies.

Secondary adjacencies are formed after the primary area's link state data base exchange has synchronized matching secondary areas with any of its primary neighbors (See [Section 4.0](#)). Link state database exchange occurs across a secondary adjacency along with normal flooding. Note that a secondary adjacency for area 0 may not be configured over an interface which is linked to a transit area for a configured virtual link.

2.2.2 Secondary Neighbor Discovery

A Type-9 opaque LSA is used to form secondary adjacencies over a primary link with adjacent opaque capable border routers. Until an opaque type is assigned for this option experimental opaque type 224 will be used. The option's LSA is referred to as an mlink Type-9 LSA. A Type-9 LSA is used for the exchange/update of an interface's secondary areas because the flooding scope of this opaque LSA needs to be restricted to routers which are directly attached to a common network or router link. A separate mlink Type-9 LSA is originated for each of the primary interface's secondary areas. Included in a secondary area's mlink Type-9 LSA is the secondary area's configured optional capabilities, its authentication parameters, plus, if required, its configured Designated Router priority. Mlink Type-9 LSAs are propagated during the exchange/update of the primary area's link state data base (LSDB) with its adjacent primary neighbors.

If a router A receives an mlink Type-9 LSA for area L originating from a router B during the exchange/update of primary area K's LSDB, then router A may form a secondary adjacency in area L with router B provided the LSA's optional capabilities and authentication parameters match those configured for the secondary area in the `SecondaryAreas` parameter. See [Section 5.2](#) for details about secondary

neighbor adjacency formation. LSDB exchange and update across a secondary adjacency proceed as defined in [OSPF] Sections 10 and 13.

2.2.3 Advertising Secondary Links

Secondary adjacencies are advertised as point-to-point or point-to-multipoint links. Any intervening transit network always belongs to the interface's primary area. Moreover, there is no network LSA for a secondary adjacency's intervening transit network in the secondary area. Hence all secondary adjacencies must be advertised in a router's router-LSA with unnumbered point-to-point type.

During the Dykstra shortest path first (SPF) calculation the LSAs for secondary adjacencies look like unnumbered point-to-point links. A Border router never includes in a secondary area's SPF tree any network across which one of its secondary adjacencies span. This ensures synchronization of the SPF tree amongst the secondary area's routers. However border routers may use the IP addresses of neighboring routers for determining a destination's next hop across a secondary link over a point-to-multipoint or transit network.

2.3 Application

Consider now the example mentioned in [Section 2.1](#) and assume NSSA 1 is configured as a secondary area over the Area 0 link between A0 and B0. Since the routers A0, A1, B0, and B1 possess NSSA 1's router-LSAs originating from A0 and B0 which define a secondary link between A0 and B0 in NSSA 1, the NSSA 1 intra-area SPF tree computed by the Dykstra calculation now includes this DS3 link with a cost of 1. Thus the path between A0 and M1 is the more preferred path

A0<->B0<->B1<->M1

and the path between A1 and N1 is the more preferred path

A1<->A0<->B0<->N1

3.0 Secondary Interfaces

3.1 The SecondaryAreas Interface Parameter

A new configuration parameter called SecondaryAreas has been added to the OSPF interface data structure (See [Appendix C](#)). The SecondaryAreas interface parameter lists a set of areas which may form unnumbered secondary adjacencies across the interface. If the SecondaryAreas interface parameter has a null list, then no secondary areas are configured for this interface and no secondary adjacencies can be formed over the interface. For each secondary area listed in

the SecondaryAreas interface parameter the interface cost, authentication parameters, and Designated Router priority are also configurable. The SecondaryAreas' interface costs and the Designated Router priorities default to the values assigned to the primary interface. Note that if a virtual link is configured across a transit area which is linked to an interface, then Area 0 may not be configured in the interface's SecondaryAreas parameter.

3.2 The Secondary Interface Data Structure

An OSPF interface data structure should be built for each of an OSPF interface's secondary areas. These interface data structures are called secondary interface data structures and are bound to the primary interface. The configured primary area of an OSPF interface generates its primary interface data structure. Recall [[OSPF](#)] Sections [8.2](#) and [10.5](#) imply that a point-to-point layer 2 link between two routers may have only one OSPF interface per area. Additional areas may be added as secondary areas, but they must not duplicate areas already configured for the layer 2 link.

A secondary interface's list of neighboring routers is generated by examining the primary interface's received mlink Type-9 LSAs as defined in [Section 4.0](#) below. Before a neighboring router may be added to the secondary interface's list of neighboring routers it must also occur in the primary interface's list of neighboring routers.

Secondary interfaces copy their interface type from the primary interface's data structure and still compute a Designated Router and a Backup Designated Router over broadcast and NBMA networks. This promotes efficient flooding across a primary interface's transit network. However, these network links are always advertised as unnumbered router<->router links and behave exactly like point-to-multipoint interfaces. As such a secondary link's Designated Router does not originate a type-2 network LSA into the secondary area for the primary interface's transit network.

3.3 The Secondary Interface State Machine

The interface states of a secondary interface are the same as the interface states of a primary interface. However in some cases state transitions are triggered by the primary interface's state machine. The InterfaceUp event is redefined for secondary links:

InterfaceUp

This event is triggered by the primary interface when it transitions to either Point-to-point, DR Other, DR, or Backup

state. Its effect on the secondary interface's state is similar to the effect the InterfaceUp event has on the primary interface's state with one exception. The periodic sending of Hello packets is suppressed for a secondary interface. If the primary interface network is a physical point-to-point or point-to-multipoint network, then the secondary interface transitions to Point-to-point state. Else if the router is not eligible to become the Designated Router, the interface state transitions to DR Other.

Otherwise the attached primary network is a broadcast or NBMA network and the router is eligible to become a Designated Router for the secondary area. In this case, an attempt is made to discover the attached network's Designated Router for the secondary area. The interface state is set to Waiting and the single shot Wait Timer is started.

The events BackupSeen and NeighborChange for a secondary interface are triggered during the processing of the primary interface's received mlink Type-9 LSAs in very much the same way these events are triggered for a standard OSPF interface during the processing of received Hello Packets (See [Section 5.5](#)).

The events InterfaceDown, LoopInd, UnloopInd, and Wait Timer are unchanged for secondary interfaces. Note that since the secondary interface's InterfaceUp event occurs only after the primary interface has transitioned to DR other, DR or Backup state, a secondary interface's Wait Timer will never start before the primary interface's Wait Timer has fired.

[3.4](#) OSPF Protocol Packet processing

Since secondary interfaces are unnumbered interfaces, OSPF protocol packet processing needs a minor adjustment. For point-to-point networks there are no changes. For broadcast, NBMA, and point-to-multipoint links, the packet's IP source address is required to be on the same network as the receiving OSPF primary interface. This can be verified by comparing the packet's IP source address to the primary interface's source address, after masking both addresses with the interface mask (See [[OSPF](#)] [Section 8.2](#)).

[3.5](#) Designated Router Selection and Function

The election of the Designated Router and the Backup Designated router for a secondary link over a broadcast or NBMA network proceeds as described in [[OSPF](#)] [Section 9.4](#). Eligibility is determined from the Designated Router priorities defined in the SecondaryAreas parameter and the received mlink Type-9 LSAs, as well as its neighbors' views of the Designated Router and the Backup Designated

Router for the secondary link, which is also found in the received mlink Type-9 LSAs. Any change in the calculating router's Designated Router or Backup Designated Router for a secondary link will result in the reorigination of the corresponding secondary area's mlink Type-9 LSA.

Note that the Designated Router for a secondary link does not originate a network-LSA (Type-2) into its secondary area for the broadcast or NBMA network over which the secondary link bridges. All secondary links remain unnumbered point-to-point or point-to-multipoint (see [Section 6.0](#)). The only function of a secondary link's Designated Router is flooding optimization.

4.0 Synchronizing Secondary Areas using Mlink Type-9 LSAs

Mlink Type-9 LSAs are used to discover and maintain secondary neighbor relationships and to elect the Designated Router and Backup Designated Router for multi-access secondary adjacencies. If the primary interface is of broadcast or NBMA type then all of its candidate Designated Routers must be opaque capable, even if these routers have no secondary areas configured for their link to the broadcast or NBMA network. Otherwise mlink Type-9 LSAs may not propagate to all potential routers capable of forming secondary adjacencies over the network. Note that these candidate Designated Routers do not have to support unnumbered OSPF multiple area links, but they do have to be opaque capable in order to flood mlink Type-9 LSAs to their adjacent primary neighbors who may or may not support unnumbered OSPF multiple area links.

The format of the mlink Type-9 LSA is detailed in [Appendix B](#). Any router which has an interface with a non-empty SecondaryAreas interface parameter must originate an mlink Type-9 LSA for each configured secondary area. If an interface's SecondaryAreas parameter is null, then no mlink Type-9 LSAs should be advertised. A secondary area's mlink Type-9 LSA minimally lists as opaque information the area's secondary area ID, its optional capabilities, its authentication parameters, and, if required, its Designated Router priority.

In addition, to ensure proper secondary neighbor state transition, a secondary area's mlink Type-9 LSA also lists as opaque information those primary neighbors from which an mlink Type-9 LSA has been received for this area with matching optional capabilities and authentication parameters (See [Section 5.5](#)). Any change in this list will result in the reorigination of a new instance of the secondary area's mlink Type-9 LSA. Also for broadcast and NBMA networks the mlink Type-9 LSA lists the router's current choice for the area's Designated Router and Backup Designated Router. A value of 0.0.0.0 in

these fields means that one has not yet been selected.

Unlike Hello Packets, a new instance of a secondary area's mlink Type-9 LSA is only originated when the LSA's opaque information content has changed or when the LSRefreshInterval has expired. The LSA's content may change during the processing of received Hellos and received mlink Type-9 LSAs, or when an mlink Type-9 LSA ages out or when the primary interface parameters of a secondary area are reconfigured. A new mlink Type-9 LSA is reoriginated following expiration of its LSRefreshInterval or when changes occur in its secondary area's interface cost, authentication parameters, router priority, Designated Router, Backup Designated Router, or the area's list of secondary neighbors at state greater than Init. Like other LSAs, the reorigination of a mlink Type-9 LSA is subject to the MinLSInterval. New mlink Type-9 LSAs are built from the list of configured secondary areas plus their corresponding secondary interface state machines and secondary neighbor state machines. Mlink type-9 LSAs are only flooded at the routers fully adjacent primary neighbors. If a secondary area is removed from a primary interface's configured secondary areas, its locally originated mlink type-9 LSA should be flushed.

The mlink Type-9 LSA of each of the primary interface's secondary areas must have a unique opaque ID. The last 24 bits of the Secondary Area ID will normally produce unique Opaque IDs. When it doesn't, alter the least significant bits until uniqueness is achieved.

5.0 Secondary Neighbors

An OSPF router converses with its secondary neighbors. Each separate conversation is described by a "secondary neighbor data structure". Conversations between secondary neighbors are bound to the secondary interface and identified by both the Area ID and either the router's OSPF Router ID or its Neighbor IP address (see below). Unless discussed below details for the creation and maintenance of secondary neighbors and secondary adjacencies are the same as discussed in [OSPF] Sections [9](#) and [10](#).

5.1 The Secondary Neighbor Data Structure

The neighbor data structure for secondary adjacencies is identical to the neighbor data structure for standard OSPF adjacencies. Secondary neighbors are discovered by comparing the contents of the primary interface's received mlink Type-9 LSAs with its configured list of secondary areas and then comparing the secondary areas' configured optional capabilities (see [Section 5.5](#)) with those listed for them in the neighboring router's mlink Type-9 LSAs. A separate neighbor data structure is built for each matching secondary area.

The Inactivity Timer used for a primary neighbor data structure is synchronized with the Inactivity Timer for all of its configured secondary neighbor data structures. The secondary link's Designated Router and the Backup Designated Router, as viewed by a secondary neighbor, are specified as opaque information stored in the neighbor's corresponding mlink Type-9 LSA and received over the primary interface. Also included are the neighbor's Designated Router priority for the secondary link, its optional capabilities, and a list of secondary neighbors it sees over the secondary link.

The secondary Neighbor ID and Neighbor IP address are copied from the corresponding primary area neighbor data structure. The OSPF optional capabilities which are supported by the neighboring router for the secondary area are learned from opaque information stored in the neighbor's mlink Type-9 LSA for this area and must match the local router's optional capabilities configured for the secondary area.

5.2 The Secondary Neighbor State Machine

While secondary neighbor states are identical to OSPF's neighbor states, there are some important distinctions in their actions and the events which trigger them. Every router which is a secondary neighbor for a secondary interface is also a primary neighbor for its primary interface. The two neighbor data structures are loosely bound together so that events which happen to the primary neighbor may trigger events for the secondary neighbor.

A secondary neighbor's Init, 1-way, and 2-way state transitions are controlled by the reception across the primary interface of Hello Packets and mlink Type-9 LSAs (See Sections [5.4](#) and [5.5](#)). The combination of the two eliminates the requirement for the periodic sending of Hello Packets for each secondary interface. When a primary neighbor state machine receives a new mlink Type-9 LSA from a primary neighbor via link state update, or confirms the status of an existing one during database exchange, the LSA's content may create a new secondary neighbor or effect the state of an existing secondary neighbor. A router must clear its own mlink Type-9 LSAs from the database summary list during database exchange with its primary neighbors to enable its secondary neighbors to transition past Init state. Any existing mlink Type-9 LSAs previously received from other primary neighbors must be cleared from the database summary list during a primary neighbor's database exchange before their content can create new secondary neighbors or change the state of existing secondary neighbors.

If an mlink Type-9 LSA of a primary neighbor ages out, the KillNbr event is executed for the primary neighbor's corresponding secondary adjacency, and the LSA should be flushed. A newly received mlink

Type-9 LSA from a primary neighbor may effect the creation of a new secondary adjacency for the neighbor. It may also cause the corresponding secondary neighbor to drop to 2-Way state or lower or be destroyed altogether. (See [Section 5.5](#))

A secondary neighbor state machine never enters Attempt state for a NBMA secondary interface. Instead it relies solely on its corresponding primary neighbor state machine to start communications over an NBMA. Hence there is no Start event for a secondary neighbor state machine (See [Section 5.4](#)).

[5.3](#) Events Triggered by the Primary Neighbor State Machine

When a primary neighbor state machine transitions down to ExStart due to either a SeqNumberMismatch or BadLSReq event, it triggers a 1-WayReceived Event for any secondary neighbors at state 2-way or greater.

The events KillNbr, LLDnwn, and Inactivity Timer of a primary neighbor state machine simultaneously trigger the same events for its loosely bound secondary neighbor state machines.

Since the formation of secondary neighbors is linked with the database exchange and the link state update of the mlink Type-9 LSAs received from their loosely bound primary neighbors, the timing of this exchange effects when a secondary neighbor transitions to ExStart state. The mlink Type-9 LSA of a secondary Area 0 should be listed at the top of the database summary list. The mlink type-9 LSAs of non-backbone secondary areas should be listed at the bottom of the database summary list. These tools mitigate the effect of the database exchange by non-backbone secondary areas on the primary neighbor state machine as it transitions to Full state, while at the same time allowing an Area 0 secondary neighbor state machine to proceed to Full state roughly in parallel with the primary neighbor state machine.

[5.4](#) Receiving Hello Packets

If a Hello Packet is received from one of the interface's existing primary neighbors which has loosely bound secondary neighbors then the corresponding secondary neighbor state machines are executed in line as follows:

- o Each Hello Packet causes each of the neighbor's existing secondary neighbor state machines at state greater than Down to be executed with the event HelloReceived.
- o Then the list of neighbors contained in the Hello Packet is

examined. If the router itself does not appear in this list, then each of the neighbor's existing secondary neighbor state machines at state greater than Init should be executed with the event 1-WayReceived.

Hello packets do not cause secondary neighbor structures to be created and do not effect the election of Designated Routers and Backup Designated Routers by the secondary interface state machines. That is triggered by the processing of the primary interface's received mlink Type-9 LSAs. All primary neighbors are considered secondary neighbors for each configured secondary area. They transition from Down to Init state with their corresponding primary neighbors and freeze at Init state until the creation of their secondary neighbor data structures is triggered when the primary interface receives their mlink Type-9 LSAs. (See [Section 5.5](#))

5.5 Receiving Mlink Type-9 Neighbor Discovery LSAs

This section explains the detailed processing of a received mlink type-9 LSA. When an mlink Type-9 LSA is acknowledged during a primary neighbor's database exchange or received via link state update, its Secondary Area ID is checked against those listed in the primary interface's SecondaryAreas parameter. If it is not listed processing of this LSA stops. If it is listed and its optional capabilities or authentication parameters no longer match those stored in the SecondaryAreas parameter, then execute a KillNbr event for any existing corresponding secondary neighbor state machine belonging to this neighbor and terminate processing of this LSA.

Else check for the existence of a secondary neighbor state machine for the Secondary Area ID that corresponds to the originating primary neighbor. If one does not exist, create one. Initialize its state to Init and copy the Neighbor ID, the Inactivity Timer, and the Neighbor IP address from the corresponding primary neighbor state machine.

When the received mlink Type-9 LSA originated from a primary neighbor over a broadcast, point-to-multipoint or NBMA network set the corresponding secondary neighbor data structure's Router Priority field, Neighbor's Designated Router field, and the Neighbor's Backup Designated Router field from the corresponding fields in the LSA's opaque data. Changes in these fields should be noted for possible use in the steps below.

Now the rest of the mlink Type-9 LSA is examined generating events to be given to the corresponding secondary neighbor and secondary interface state machines. These state machines are specified either to be executed or scheduled (see [\[OSPF\] Section 4.4](#)). For example, by specifying below that the secondary neighbor state machine be

executed in line, several neighbor state transitions may be effected by a single received mlink Type-9 LSA:

- o The list of secondary neighbors contained in the LSA's opaque data is examined. If the router itself appears in this list, then the corresponding secondary neighbor state machine for this neighbor should be executed with the event 2-WayReceived. Otherwise, the secondary neighbor state machine should be executed with the event 1-WayReceived, and the processing of the LSA stops.
- o Next, if a change in the neighbor's Secondary Router Priority field in the LSA was noted, the corresponding secondary interface state machine is scheduled with the event NeighborChange.
- o If the neighbor is both declaring itself to be the Designated Router in the LSA (LSA's Designated Router field = Neighbor IP address), and the Backup Designated Router field in the LSA's opaque information is equal to 0.0.0.0, and the corresponding secondary interface is in state Waiting, then schedule the secondary interface state machine with the event BackupSeen. Otherwise, if the neighbor is declaring itself to be the Designated Router in the LSA and it had not previously, or the secondary neighbor is not declaring itself Designated Router in the LSA where it had previously, then schedule the secondary interface state machine with the event NeighborChange.
- o If the secondary neighbor is declaring itself to be the Backup Designated Router in the LSA (LSA's Backup Designated Router field = Neighbor IP address) and the corresponding secondary interface is in state Waiting, then schedule the secondary interface state machine with the event BackupSeen. Otherwise, if the neighbor is declaring itself to be the Backup Designated Router in the LSA and it had not previously, or the neighbor is not declaring itself Backup Designated Router where it had previously, the secondary interface state machine is scheduled with the event NeighborChange.

6.0 Advertising Secondary Adjacencies

Border routers advertise their secondary adjacencies in their router-LSAs as unnumbered point-to-point links even though they may be numbered point-to-point links, point-to-multipoint links, or broadcast or NBMA network links in the OSPF interface's primary area. This allows their interconnecting networks to retain a single area identity. As unnumbered point-to-point links, all secondary adjacencies have link type 1. The building of the router-LSA, as described in [[OSPF](#)] [Section 12.4.1](#), is virtually unchanged.

Even though secondary adjacencies are considered unnumbered point-to-point links, for the purpose of defining Link Data in the router-LSA (see [Appendix A](#)) they use the primary interface's IP address. For secondary adjacencies [\[OSPF\] Section 12.4.1.1](#) is reduced to simply

If a neighboring router has formed a secondary adjacency then add a type 1 (point-to-point) link for this adjacency. Its Link ID should be set to the Router ID of the neighboring router. If the primary interface is numbered, the Link Data should specify the primary interface's IP address. For unnumbered point-to-point networks, the Link Data field should specify the interface's MIB-II [\[MIB\]](#) ifIndex value. The cost should be set to the secondary area's cost specified in the primary interface's SecondaryAreas parameter.

By not adding the host route type 3 links noted in [\[OSPF\] Section 12.4.1](#), secondary adjacencies retain the look and feel of unnumbered point-to-point links to the remaining routers in the secondary area, even though they may configure their link data with the primary interface's IP address.

[7.0](#) The Shortest Path First Calculation

Since secondary links appear in the link state data base of an area as unnumbered point-to-point links there is no change required in the Shortest Path First (SPF) calculation, except on those border routers where they are configured. Border routers do not include in a secondary area's SPF tree any network which one of its secondary adjacencies transit. This ensures synchronization of the SPF tree amongst the secondary area's routers. However border routers do use the IP addresses stored in their secondary neighbors' router-LSAs for determining a destination's next hop across a secondary link when the associated interface connects to a point-to-multipoint or transit network. In this case the outgoing interface is derived directly from the destination router's next hop IP address.

[8.0](#) Flushing Secondary Adjacencies

Secondary adjacencies are flushed from the link state database of a secondary area when their neighbor states transition down from Full status, just as with normal OSPF adjacencies. A new router-LSA is built for the secondary area and flooded out all of the secondary area's primary and secondary interfaces. Finally the SPF calculation is performed to reflect the new link state topology.

9.0 Security Considerations

Security issues are not discussed in this memo.

10.0 Acknowledgments

This document was produced by the OSPF Working Group, chaired by John Moy. In addition, comments from the following individual are also acknowledged:

Acee Lindem IBM

11.0 References

- [MIB] McCloghrie, K., and M. Rose, "Management Information Base for network management of TCP/IP-based internets: MIB-II", STD 17, [RFC 1213](#), Hughes LAN Systems, Performance Systems International, March 1991.

- [OPAQ] Coltun, Rob, "The OSPF Opaque LSA Option", FORE Systems, August 1998.

- [OSPF] Moy, J., "OSPF Version 2", [RFC 2328](#), Ascend Communications, Inc., April 1998.

12.0 Author's Address

Pat Murphy
US Geological Survey
345 Middlefield Road, Mail Stop 870
Menlo Park, California 94560

Phone: (650)329-4044
EMail: pmurphy@usgs.gov

Appendix A. Router-LSAs

Router-LSAs are Type 1 LSAs. Each router in an area originates a router-LSA. The router-LSA describes the state and cost of the router's links (i.e., interfaces) to the area. All of the router's links to an area, including secondary links, must be described in a single router-LSA. For details concerning the construction of router-LSAs see this document's [Section 6.0](#) and [\[OSPF\] Section 12.4.1](#).

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               LS age               |   Options   |   1   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Link State ID         |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Advertising Router    |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               LS sequence number    |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               LS checksum            |               length   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   0   |V|E|B|   0   |               # links                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Link ID               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Link Data              |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type   |   # TOS   |               metric                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               ...               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   TOS   |   0   |               TOS metric                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Link ID               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Link Data              |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               ...               |

```

In router-LSAs, the Link State ID field is set to the router's OSPF Router ID. Router-LSAs are flooded throughout a single area only.

bit V

When set, the router is an endpoint of one or more fully adjacent virtual links having the described area as its transit area (V is for virtual link endpoint).

bit E

When set, the router is an AS boundary router (E is for external).

bit B

When set, the router is an area border router (B is for border).

bit W

When set, the router is a wild-card multicast receiver (W is for wild).

bit Nt

When set, the router is a NSSA border router which translates type-7 LSAs into type-5 LSAs (Nt is for NSSA translation).

links

The number of router links described in this LSA. This must be the total collection of router links (i.e., interfaces) to the area as well as a separate entry for each fully adjacent secondary neighbor across a broadcast or NBMA network link. Secondary interfaces to broadcast and NBMA networks are considered point-to-multipoint networks.

The following fields are used to describe each router link (i.e., interface). Each router link is typed (see the below Type field). The Type field indicates the kind of link being described. It may be a link to a transit network, to another router or to a stub network. The values of all the other fields describing a router link depend on the link's Type. For example, each link has an associated 32-bit Link Data field. For links to stub networks this field specifies the network's IP address mask. For other link types the Link Data field specifies the router interface's IP address.

Type

A quick description of the router link for one of the following. Note that host routes are classified as links to stub networks with network mask of 0xffffffff.

Type	Description
1	Point-to-point connection to another router
2	Connection to a transit network
3	Connection to a stub network
4	Virtual link

Link ID

Identifies the object that this router link connects to. Its value depends on the link's Type. When connecting to an object that also originates an LSA (i.e., another router or a transit network) the Link ID is equal to the neighboring LSA's Link State ID. Secondary links are always type 1 point-to-point regardless of the type of their matching primary link. This provides the key for looking up the neighboring LSA in the link state database during the routing table calculation. See [\[OSPF\] Section 12.2](#) for more details.

Type	Link ID
1	Neighboring router's Router ID
2	IP address of Designated Router
3	IP network/subnet number
4	Neighboring router's Router ID

Link Data

Value again depends on the link's Type field. For connections to stub networks, Link Data specifies the network's IP address mask. For unnumbered point-to-point connections, it specifies the interface's MIB-II [\[MIB\]](#) ifIndex value. For the other link types it specifies the router interface's IP address. This latter piece of information is needed during the routing table build process, when calculating the IP address of the next hop. Although secondary links are considered unnumbered point-to-point links, they do evaluate Link Data as if they were numbered whenever their interface has an IP address associated with its primary area. See [Section 6.0](#) of this document and [\[OSPF\] Section 16.1.1](#) for more details.

TOS

The number of different TOS metrics given for this link, not counting the required link metric (referred to as the TOS 0 metric in [\[1583\]](#)). For example, if no additional TOS metrics are given, this field is set to 0. Secondary Areas do not use TOS.

metric

The cost of using this router link. For secondary links the cost is specified in the primary interface's SecondaryAreas parameter.

Additional TOS-specific information may also be included, for backward compatibility with previous versions of the OSPF specification ([\[1583\]](#)). Within each link, and for each desired TOS,

TOS-specific link information may be encoded as follows:

TOS

IP Type of Service that this metric refers to. The encoding of TOS in OSPF LSAs is described in [1583] [Section 12.3](#).

TOS metric

TOS-specific metric information.

The link-state ID of an mlink Type-9 LSA is divided into an Opaque Type field (the first 8 bits), which has value mlink, and the Opaque ID (the remaining 24 bits). The mlink Type-9 LSA of each of the primary interface's secondary areas must have a unique opaque ID. The last 24 bits of the Secondary Area ID will normally produce unique Opaque IDs. When it doesn't, alter the least

significant bits until uniqueness is achieved.

Options

The optional capabilities supported by this router for the primary area, as documented in [\[OSPF\]](#) Section A.2.

Secondary Area ID

The area ID of the area across which a neighboring router may form a secondary adjacency.

Sec AuType

Identifies the authentication procedure to be used. Authentication is discussed in [\[OSPF\]](#) [Appendix D](#) of the specification. Consult [\[OSPF\]](#) [Appendix D](#) for a list of the currently defined authentication types.

Sec Authentication

A 64-bit field for use by the authentication scheme. See [\[OSPF\]](#) [Appendix D](#) for details.

Neighbors Cnt

This 16 bit field describes the number of secondary neighbors listed for the Secondary Area ID. If there are no secondary neighbors listed, then Neighbor Cnt should be set to 0.

Sec Options

The optional capabilities supported by this router for this secondary area, as documented in [\[OSPF\]](#) Section A.2.

Sec Rtr Pri

This router's Designated Router priority for a secondary link across a transit network. This parameter is used during the secondary links (Backup) Designated Router election. If set to 0, the router will be ineligible to become the (Backup) Designated Router for this secondary link.

Designated Router

The identity of the Designated Router for a secondary link, in the view of the originating router. The Designated Router is identified here by the IP address across the primary link. Its value is set to 0.0.0.0 if there is no Designated Router for the secondary link.

Backup Designated Router

The identity of the Backup Designated Router for a secondary link, in the view of the originating router. The Backup Designated Router is identified here by its IP address across the primary link. Its value is set to 0.0.0.0 if there is no

Backup Designated Router for the secondary link.

Secondary Neighbor

The Router IDs of each router from whom valid mlink Type-9 LSAs have been received across the primary link for the secondary area with matching optional capabilities and authentication parameters.

Appendix C. Interface Data Structure

Chapter 9 in the OSPF specification documents the interface data structure and the data items which are included in it. [Section 3.1](#) of this document defines a new interface configuration parameter called SecondaryAreas which supports unnumbered OSPF multiple area links. The SecondaryAreas parameter's default setting is null. The SecondaryAreas parameter in a secondary interface data structure is always null.

For each secondary area listed in an interface's SecondaryAreas parameter, the secondary link's interface cost, authentication parameters and Designated Router priority must be configurable. If the interface cost and Designated Router priority are not configured they default to their corresponding values for the OSPF primary interface. If a virtual link is configured across a transit area linked to an interface, then Area 0 may not be configured in the interface's SecondaryAreas parameter.

