

Network Working Group
Internet Draft
Expiration Date: December 1998
File name: [draft-ietf-ospf-mlinks-00.txt](#)

P. Murphy
US Geological Survey
July 1998

OSPF Multiple Area Links

Status of this Memo

This document is an Internet-Draft. 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".

To learn the current status of any Internet-Draft, please check the "1id- abstracts.txt" listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), nic.nordu.net (Europe), munnari.oz.au (Pacific Rim), ftp.ietf.org (US East Coast), or ftp.isi.edu (US West Coast).

Table Of Contents

1.0	Abstract	1
2.0	Overview	2
2.1	Requirement for OSPF Multiple Area Links	2
2.2	Proposed Solution	3
3.0	Implementation Details	6
3.1	SecondaryAreas Interface parameter	6
3.2	Advertising Secondary Areas	6
3.3	Forming Secondary Adjacencies	7
3.4	Advertising Secondary Adjacencies	7
3.5	Secondary Area Link State Data Base Exchange and Update ..	8
3.6	The Shortest Path First Calculation	9
3.7	Flushing Secondary Adjacencies	9
4.0	Acknowledgments	10
5.0	References	10
6.0	Security Considerations	10
7.0	Authors' Addresses	10
Appendix A	: Router LSAs	11
Appendix B	: mlink Opaque LSAs	13
Appendix C	: Configuration Parameters	14

[1.0](#) Abstract

This memo describes an option to the OSPF Version 2 specification which allows multiple areas to share the same link. This option adds no additional link state flooding over shared links other than the normal link state database exchange and update originating from the link's configured primary area. The option applies to standard areas, stub areas, and NSSAreas. It eliminates the excess Area 0 link state baggage that accompanies the use of virtual links as currently practiced when configuring similar transits for standard OSPF areas. Routers with this option configured are backward compatible with routers running the standard OSPF Version 2 compliant implementation as defined in [RFC 2328](#), and can be restricted to a subset of the OSPF routing domain. This option is applied only on OSPF border routers.

Implementation of OSPF multiple area links requires a modification to the OSPF interface data structure which allows each interface of a border router to be connected to multiple areas. One area is always configured as the interface's primary area. Any additional areas which are configured for an interface are called the interface's secondary areas. Two adjacent border routers with mutually shared secondary areas may transit a secondary area's intra-area traffic over the adjacency. A typical application is a stub area or NSSArea which is dual homed to the Area 0 backbone and loosely joined by an internal slow speed link. If there exists a high speed Area 0 link between two of the stub area's border routers, this option allows the

stub area's intra-area traffic to route across this high speed area 0 link. The current specification forces traffic to prefer the intra-area path over the slow speed link.

Another not so common application makes Area 0 the secondary area over a local high speed LAN link with the primary area a local stub or NSSArea. Here Area 0 is not primary due to topological limitations which restrict its applicability. E.g. the local LAN link could be a campus backbone with dozens of routers on it, all part of the same NSSArea.


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

2.0 Overview

2.1 Requirement for OSPF Multiple Area Links

Large corporate networks in today's modern Internet invest tremendous human and equipment resources, coupled with sizable budget, into 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 high speed and high reliability network links. Performance over a T1 link can pale by comparison to performance over a DS3 or OC3 backbone link. Indeed, even a high speed 100 megabit LAN link can carry an order of magnitude more traffic than a standard 10 megabit link. Large corporate networks have sizable backbone routing tables and routinely use stub areas and NSSAreas. 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 NSSArea 1. A1 and B1 are internal routers in NSSArea 1. N1 and M1 are standard ethernet networks in NSSArea 1 directly attached to B0 and B1 respectively. Assume all T1 links have an OSPF cost of 28 in both directions, N1 and M1 are attached with interface costs 2. The DS3 link has an OSPF cost of 1. All costs are symmetric. Under the current OSPF specification the preferred path between A0 and M1 is

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

with OSPF cost 58, even though there exists a more preferred path through B0 namely

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

with OSPF cost 31. Indeed the DS3 link between A0 and B0 has the potential of making the performance between A0 and B1 approach that of a single T1.

Consider also the NSSArea 1 OSPF preferred path between A1 and N1,

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

with cost 58. Again there exists a more preferred path through the link between A0 and B0, namely

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

with cost 31. This link also has the potential of performing at single hop T1 speeds. Under the current OSPF specification NSSArea 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 in NSSArea 1. Virtual links are not an option since virtual links don't work over NSSAreas. If the link between A0 and B0 were part of NSSArea 1, path preference would be optimal as the DS3 path would be intra-area for NSSArea 1. This cannot be the non-backbone equivalent of an Area 0 virtual link. Excessive use of such a feature would negate the purpose of the Area 0 backbone.

2.2 Proposed Solution

OSPF multiple area links provide the capability of allowing multiple areas to use the same link/path 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/path as part of its link state topology. Only areas which are dual homed to Area 0 may take advantage of this option, since any router with multiple directly attached areas must be in Area 0 and it takes at least two such border routers to create a multiple area link.

The current OSPF specification only allows one area per interface. Thus, should an Area L dual-home to two Area 0 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 link.

The softening of this restriction is facilitated by the addition of a new parameter to the interface data structure called SecondaryAreas. The SecondaryAreas interface parameter is a self-defining structure containing a list of Area Ids for additional areas in which the

interface also belongs, along with their associated OSPF costs. The areas listed in this parameter are called the interface's secondary areas, as opposed to the interface's configured Area ID, hereafter called the primary area. For each secondary area listed in the SecondaryAreas parameter there is a list of neighboring Router Ids which have formed adjacencies for the secondary area across the interface's directly attached network or router link. These adjacencies are called secondary adjacencies. Secondary adjacencies are formed during the primary area's link state data base exchange and update as defined in Sections [3.2](#) and [3.3](#).

A type 9 Opaque LSA is used to exchange/update the interface's secondary areas with adjacent Opaque capable border routers. Until an opaque type is assigned for this option, it will be simply referred to as an mlink type 9 opaque LSA. A type 9 opaque LSA is used because the flooding scope of the mlink type 9 opaque LSA needs to be restricted to those routers directly attached to a common network or router link. Each router will list its configured secondary addresses as defined by the interface's SecondaryAreas parameter in its mlink type 9 opaque LSA. The mlink type 9 opaque LSA is propagated during the primary area's link state data base exchange and update with its fully adjacent neighbors. If a router A receives an mlink type 9 opaque LSA from an adjacent router B during primary area K's link state data base exchange, then router A and B form a secondary adjacency in area L if and only if

- 1) Area L is listed in the interface's SecondaryArea parameter.
- 2) Area L is listed in router B's received mlink type 9 opaque LSA.
- 3) Router A contains a type 1 (router) LSA from router B in Area L's link state data base.

If routers A and B form a secondary adjacency then the router ID of router B is added to Area L's list of secondary adjacencies for the interface.

Secondary adjacencies are either point-to-point or point-to-multipoint. Any intervening transit networks always belong to the interface's primary area. There is no network LSA for a secondary adjacency's intervening transit network in the secondary area. Hence all secondary adjacencies for Area L must be advertised in router A's type 1 (router) LSA with unnumbered point-to-point type.

During the Dykstra SPF calculation the LSAs for secondary adjacencies look like unnumbered point-to-point links, except on border routers where they originate. 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 do use the IP addresses stored in neighboring border router LSAs for determining link paths across a multi-access secondary adjacency.

There is no secondary link state data base exchange or update over secondary adjacencies. Instead, these adjacencies rely on the secondary area's base topology (minus the secondary links) for flooding link state information regarding an area's secondary adjacencies. The reasons for not attempting a secondary link state data base exchange/update are two fold. First, the goal of this document is to establish a transit for multiple areas over a single link without imposing the additional traffic load caused by passing the secondary area's link state database over that link. Second the added complexity of synchronizing multiple secondary area link state databases over the link is unnecessary when the primary area is the Area 0 backbone. Should the secondary area partition, traffic between the two parts of the partitioned secondary area would still use the secondary adjacent link via OSPF inter-area routing.

It is true that routing between the different parts of a secondary area partitioned by a defunct secondary adjacency over an Area 0 link would be based on the inter-area route computation results from processing summary links. Summary-range aggregation may over simplify routing between the different parts of a partitioned area. However this is no different that what exists under the current OSPF base standards. Furthermore, the proper partitioning of a secondary area's assigned address space not only allows the summary-range aggregates to effect proper routing during topological failures, but it also promotes more cost effective and efficient routing even when secondary areas are whole.

Consider now the example mentioned in [Section 2.3](#) and assume Area 1 is configured as a secondary area over the Area 0 link between A0 and B0. The NSSArea 1 intra-area shortest path first tree computed by the Dykstra calculation now includes the DS3 link between A0 and B0 with a cost of 1. Given that all four routers possess link state advertisements for the link between A0 and B0 in NSSArea 1, the path between A0 and M1 is now

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

and the path between A1 and N1 is

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

Furthermore, should the link between A1 and B1 fail terminating the

NSSArea 1's secondary adjacency between A0 and B0, the path between A0 and M1 would still be

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

as the summary LSA for M1 originating from B0 forms an inter-area path to M1.

3.0 Implementation Details

3.1 SecondaryAreas Interface parameter

A new parameter called SecondaryAreas is added to the OSPF interface data structure. Implementations must support the manual configuration of this parameter. The SecondaryAreas interface parameter lists the set of Areas which can form secondary adjacencies over the interface, complete with their secondary area interface costs. By default a secondary area's interface cost is set to the value of the interface's primary area cost. Only those Areas listed in the SecondaryAreas interface parameter are candidate areas for the interface's secondary adjacencies. 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, there is a list of Router Ids which have formed secondary adjacencies to the router over the interface. These Routers Ids should also be listed in the interface's list of neighboring routers. If the list of Router Ids is null, then the interface has no secondary adjacencies for this secondary area. This list is populated as defined in [Section 3.3](#) below.

3.2 Advertising Secondary Areas

All routers which support secondary areas must also be opaque capable. If the interface is of broadcast type, then all candidate designated routers must be opaque capable, even if they have no secondary areas configured for the interface which connects to the broadcast transit network. Otherwise type 9 opaque LSAs, which are used by this option, will not propagate to all potential routers capable of forming secondary adjacencies over the broadcast transit network. Note that candidate designated routers do not have to support OSPF multiple area links, but they do have to be opaque capable in order to flood the type 9 opaque LSA to their adjacent neighbors who may or may not support OSPF multiple area links.

Any router which has an interface with a non-empty SecondaryAreas

interface parameter must advertise a type 9 opaque LSA with opaque type mlink to the interface's fully adjacent neighbors. The mlink type 9 opaque LSA contains the Area Ids listed in the interface's SecondaryAreas parameter as Opaque Information. See [Appendix B](#) for details. If an interface's SecondaryAreas parameter is null, then the mlink type 9 opaque LSA should not be advertised.

By default, the mlink opaque type sets its opaque ID to the last 24 bits of the interface's IP address, if it has one. In the case of unnumbered point-to-point connections, the mlink opaque ID is set to the last 24 bits of the interface's MIB-II [[MIB](#)] ifIndex value. All implementations should also support the manual configuration of an interface's mlink opaque ID in the rare instance this default schema results in duplication.

[3.3](#) Forming Secondary Adjacencies

Two border routers A and B in area L may form a secondary adjacency across an interface, whose primary area is K, under the following conditions:

- (a) Routers A and B are fully adjacent across the interface's primary area K or they share a common area K full adjacency with the interface's designated router.
- (b) Routers A and B see router LSAs (type-1) for each other in Area L's link state data base.
- (c) Routers A and B have exchanged/updated mlink type 9 opaque LSAs across the interface both of which contain Area L as a candidate secondary area.

When all three conditions are met, a secondary adjacency is formed between routers A and B across the interface's primary area. There is no particular ordering of events. Either (a), (b), or (c) can trigger the event which forms the secondary adjacency, provided the other two conditions have already been met.

[3.4](#) 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 transit network links in their associated interface's primary area. This allows their interconnecting networks to retain a single area identity, thus avoiding the inevitable problems which would occur with duplicate summary links advertisements and aggregation. It also makes configuration and deployment seamless.

As unnumbered point-to-point links, all secondary adjacencies have link type 1. The building of the router LSA is described in [\[OSPF\] Section 12.4.1](#). For the purpose of building the router LSA, an interface belongs to both its primary area and its secondary areas. With this in mind, no change is required to [\[OSPF\] Section 12.4.1](#).

Even though secondary adjacencies are considered unnumbered point-to-point links, for the purpose of defining Link Data in the type 1 (router) LSA (see [Appendix A](#)) we allow them to use the 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 link (point-to-point). The Link ID should be set to the Router ID of the neighboring router. If the interface is numbered, the Link Data should specify the IP interface 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 of the point-to-point interface as defined in the interface's SecondaryAreas parameter.

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

[3.5](#) Secondary Area Link State Data Base Exchange and Update

There is no link state data base exchange or update across a secondary adjacency. Only the interface's primary area synchronizes its link state data base with the interface's neighboring routers. Instead, secondary areas rely on their base topology (i.e. the topology without the secondary links), for complete flooding and synchronization of the link state data base throughout the secondary area. If the primary area is the area 0 backbone this is not a problem, since the routing table entries generated by the secondary areas summary links will forward packets when a secondary area's base topology partitions in such a way that its secondary adjacencies are dropped.

If the primary area is non-zero, then a simple IP over IP tunnel may be configured for the secondary across the physical link. This tunnel will perform link state data base synchronization for the secondary area across the interface. The IP tunnel is not part of this specification, but clearly it is advantageous to any OSPF multiple area links implementation if tunneling of IP over IP is supported. Note that if the tunnel is suitably configured with a high OSPF cost,

it would never be used for forwarding packets other than those required for link state data base exchange and update. Moreover the tunnel will always exist as long as the secondary adjacency's physical link is active.

3.6 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' type-1 (router) LSAs for determining the destination next hop across a secondary adjacency when the associated interface connects the destination router via a point-to-multipoint link or transit network link. In this case the outgoing interface is derived directly from the destination router's next hop IP address.

3.7 Flushing Secondary Adjacencies

The dropping and flushing of secondary adjacencies from the link state database of a secondary area is triggered by any of the following events:

- The secondary area is manually purged from the interface's SecondaryAreas parameter.

- An adjacent router loses its primary area adjacency. Note that over a multi-access network termination happens when a neighboring router ceases to be listed in the Designated Router's list of neighbors.

- The secondary area is no longer listed in the adjacent router's mlink type 9 opaque LSA.

- The adjacent router's mlink type 9 opaque LSA is flushed from the primary area or exceeds MAXAge.

- The adjacent router's type 1 (router) LSA is flushed from the secondary area or exceeds MAXAge. In other words, the secondary area has partitioned with the two formerly adjacent routers now in separate parts.

When any of these events occur the interface's SecondaryAreas parameter is updated to purge the adjacency. A new router LSA is

built for the secondary area and flooded out those interfaces for which the secondary area is primary. The SPF calculation is done to reflect the new link state topology.

4.0 Acknowledgments

This document was produced by the OSPF Working Group, chaired by John Moy.

In addition, the comments of the following individuals are also acknowledged:

Derek Yeung
Rob Colton

Cisco Systems
Fore Systems

5.0 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.
- [1583] Moy, J., "OSPF Version 2", [RFC 1583](#), Proteon, Inc., March 1994.

6.0 Security Considerations

Security issues are not discussed in this memo.

7.0 Author's Address

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

Phone: (650)329-4044
EMail: pmurphy@noc.doi.net

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, including secondary links, to an area must be described in a single router-LSA. For details concerning the construction of router-LSAs, see this document's [Section 3.6](#) 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 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 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.

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. 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 area are always type 1 point-to-point regardless of the type of their associated primary area. 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 3.6](#) 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 may use TOS.

metric

The cost of using this router link.

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 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 mlink Opaque LSA is used directly by OSPF to distribute list of candidate secondary areas among neighboring routers. The flooding scope of the mlink type 9 Opaque LSA is link-local, which means mlink LSAs are never forwarded beyond the local (sub)network.

[illegible]

The link-state ID of the mlink Opaque 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 Opaque ID is set to the last 24 bits of the interface's primary area IP address, if it has one. In the case of unnumbered point-to-point connections, the mlink opaque ID is set to the last 24 bits of the interface's MIB-II [MIB] ifIndex value. All implementations should also support the manual configuration of an interface's mlink opaque ID in the rare instance this default schema results in duplication.

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 configuration parameter called SecondaryAreas which is to be included in support of OSPF multiple area links. Sections [3.2](#) and [3.4](#) describe the application of the parameter.

In addition, for each secondary area listed in an interface's SecondaryAreas parameter, the interface data structure must maintain a list of Router IDs which have formed secondary adjacencies for this secondary area.

