

MULTICAST EXTENSIONS TO OSPF (MOSPF)

Version 2 of the Open Shortest Path First (OSPF) routing protocol is defined in [RFC-1583](#). It is an Interior Gateway Protocol (IGP) specifically designed to distribute unicast topology information among routers belonging to a single Autonomous System. OSPF is based on link-state algorithms which permit rapid route calculation with a minimum of routing protocol traffic. In addition to efficient out calculation, OSPF is an open standard that supports hierarchical routing, load balancing, and the import of external routing information.

The Multicast Extensions to OSPF (MOSPF) are defined in [RFC-1584](#). MOSPF routers maintain a current image of the network topology through the unicast OSPF link-state routing protocol. MOSPF enhances the OSPF protocol by providing the ability to route multicast IP traffic. The multicast extensions to OSPF are built on top of OSPF Version 2 so that a multicast routing capability can be easily introduced into an OSPF Version 2 routing domain. The enhancements that have been added are backwards compatible so that routers running MOSPF will interoperate with non-multicast OSPF routers when forwarding unicast IP data traffic. MOSPF, unlike DVMRP, does not provide support for tunnels.

Intra-Area Routing with MOSPF

Intra-Area Routing describes the basic routing algorithm employed by MOSPF. This elementary algorithm runs inside a single OSPF area and supports multicast forwarding when the source and all destination group members reside in the same OSPF area, or when the entire Autonomous System is a single OSPF area. The following discussion assumes that the reader is familiar with the basic operation of the OSPF routing protocol.

Local Group Database

Similar to DVMRP, MOSPF routers use the Internet Group Management Protocol (IGMP) to monitor multicast group membership on directly attached subnetworks. MOSPF routers are required to implement a "local group database" which maintains a list of directly attached group members and determines the local router's responsibility for delivering multicast datagrams to these group members.

On any given subnetwork, the transmission of IGMP Host Membership Queries is performed solely by the Designated Router (DR). Also, the responsibility of listening to IGMP Host Membership Reports is performed only by the Designated Router (DR) and the Backup Designated Router (BDR). This means that in a mixed environment containing both MOSPF and OSPF routers, an MOSPF router must be elected the DR for the subnetwork if IGMP Queries are to be generated. This can be achieved by simply assigning all non-MOSPF routers a RouterPriority of 0 to prevent them from becoming the DR or BDR, thus allowing an MOSPF router to become the DR for the subnetwork.

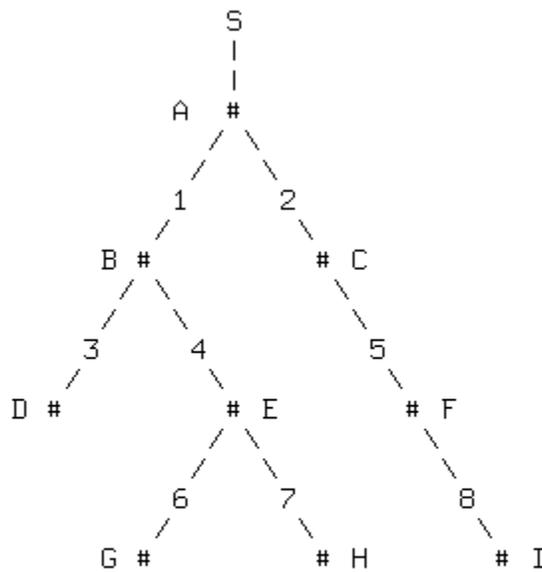
The DR is responsible for communicating group membership information to all other routers in the OSPF area by flooding Group-Membership LSAs. The DR originates a separate Group-Membership LSA for each multicast group having one or more entries in the DR's local group database. Similar to

Router-LSAs and Network-LSAs, Group Membership-LSAs are flooded throughout a single area only. This ensures that all remotely-originated multicast datagrams are forwarded to the specified subnetwork for distribution to local group members.

Datagram's Shortest Path Tree

The datagram's shortest path tree describes the path taken by a multicast datagram as it travels through the internetwork from the source subnetwork to each of the individual group members. The shortest path tree for each (source, group) pair is built "on demand" when a router receives the first multicast datagram for a particular (source, group) pair.

When the initial datagram arrives, the source subnetwork is located in the MOSPF link state database. The MOSPF link state database is simply the standard OSPF link state database with the addition of Group-Membership LSAs. Based on the Router-LSAs and Network-LSAs in the MOSPF link state database, a source-rooted shortest-path tree is constructed using Dijkstra's algorithm. After the tree is built, Group-Membership LSAs are used to prune those branches that do not lead to subnetworks containing individual group members. The result of the Dijkstra calculation is a pruned shortest-path tree rooted at the datagram's source.



LEGEND

Router

Shortest Path Tree for (S, G)

To forward a multicast datagram to downstream members of the group, each router must determine its position in the datagram's shortest path delivery tree. Assume that figure above illustrates the shortest path tree for a particular (source, group) pair. Router E's upstream node is Router B and there are two

downstream interfaces: one connecting to Subnetwork 6 and another connecting to Subnetwork 7.

Note the following properties of the basic MOSPF routing algorithm:

- For a given multicast datagram, all routers within an OSPF area calculate the same source-rooted shortest path delivery tree. Tie-breakers have been defined to guarantee that if several equal- cost paths exist, all routers agree on a single path through the area. Unlike unicast OSPF, MOSPF does not support the concept of equal-cost multipath routing.
- Synchronized link state databases containing Group-Membership LSAs allow an MOSPF router to perform the Reverse Path Multicasting (RPM) computation "in memory." Unlike DVMRP, this means that the first datagram of a group transmission does not have to be forwarded to all routers in the area.
- The "on demand" construction of the shortest-path delivery tree has the benefit of spreading calculations over time, resulting in a lesser impact for participating routers.

Forwarding Cache

Each MOSPF router makes its forwarding decision based on the contents of its forwarding cache. The forwarding cache is built from the source-rooted shortest-path tree for each (source, group) pair and the router's local group database. After the router discovers its position in the shortest path tree, a forwarding cache entry is created containing the (source, group) pair, the upstream node, and the downstream interfaces. At this point, the Dijkstra shortest path tree is discarded releasing all resources associated with the creation of the tree. From this point on, the forwarding cache entry is used to forward all subsequent datagrams for the (source, group) pair.

```

=====
Destination      Source          Upstream      Downstream    TTL
-----
224.1.1.1        128.1.0.2      11            12  13          5
224.1.1.1        128.4.1.2      11            12  13          2
224.1.1.1        128.5.2.2      11            12  13          3
224.2.2.2        128.2.0.3      12            11
=====

```

■ MOSPF Forwarding Cache

```

=====

```

Above figure displays the forwarding cache for a typical MOSPF router. The elements in the display include the following items:

Destination-- The destination group address to which matching datagrams are forwarded.

Source-- The datagram's source subnetwork. Each Destination/Source pair identifies a separate forwarding cache entry.

Upstream-- The interface from which a matching datagram must be received

Downstream-- The interfaces over which a matching datagram should be forwarded to reach Destination

group members

TTL-- The minimum number of hops a datagram will travel to reach the multicast group members. This allows the router to discard datagrams that do not have a chance of reaching a destination group member.

The information in the forwarding cache is not aged or periodically refreshed. It is maintained as long as there are system resources available (i.e., memory) or until the next topology change. In general, the contents of the forwarding cache will change when:

- The topology of the OSPF internetwork changes forcing all of the datagram shortest-path trees to be recalculated.
- There is a change in the Group-Membership LSAs indicating that the distribution of individual group members has changed.

Mixing MOSPF and OSPF Routers

MOSPF routers can be combined with non-multicast OSPF routers. This permits the gradual deployment of MOSPF and allows experimentation with multicast routing on a limited scale. When MOSPF and non-multicast OSPF routers are mixed within an Autonomous System, all routers will interoperate in the forwarding of unicast datagrams.

It is important to note that an MOSPF router is required to eliminate all non-multicast OSPF routers when it builds its source-rooted shortest-path delivery tree. An MOSPF router can easily determine the multicast capability of any other router based on the setting of the multicast bit (MC-bit) in the Options field of each router's link state advertisements. The omission of non-multicast routers can create a number of potential problems when forwarding multicast traffic:

- Multicast datagrams may be forwarded along suboptimal routes since the shortest path between two points may require traversal of a non-multicast OSPF router.
- Even though there is unicast connectivity to a destination, there may not be multicast connectivity. For example, the network may partition with respect to multicast connectivity since the only path between two points requires traversal of a non-multicast OSPF router.
- The forwarding of multicast and unicast datagrams between two points may follow entirely different paths through the internetwork. This may make routing problems somewhat more difficult to debug.
- The Designated Router for a multi-access network must be an MOSPF router. If a non-multicast OSPF router is elected the DR, the subnetwork will not be selected to forward multicast datagrams since a non-multicast DR cannot generate Group- Membership LSAs for its subnetwork.

Inter-Area Routing with MOSPF

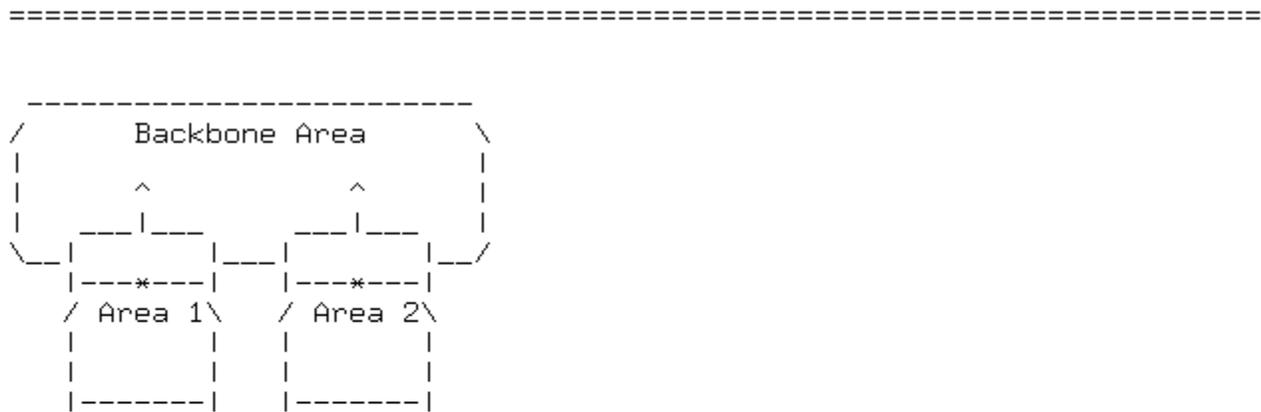
Inter-area routing involves the case where a datagram's source and some of its destination group members reside in different OSPF areas. It should be noted that the forwarding of multicast datagrams continues to be determined by the contents of the forwarding cache which is still built from the local group database and the datagram shortest-path trees. The major differences are related to the way that

group membership information is propagated and the way that the inter-area shortest-path tree is constructed.

Inter-Area Multicast Forwarders

In MOSPF, a subset of an area's Area Border Routers (ABRs) function as "inter-area multicast forwarders." An inter-area multicast forwarder is responsible for the forwarding of group membership information and multicast datagrams between areas. Configuration parameters determine whether or not a particular ABR also functions as an inter-area multicast forwarder.

Inter-area multicast forwarders summarize their attached areas' group membership information to the backbone by originating new Group-Membership LSAs into the backbone area. It is important to note that the summarization of group membership in MOSPF is asymmetric. This means that group membership information from non-backbone areas is flooded into the backbone. However, the backbone does not readvertise either backbone group membership information or group membership information learned from other non-backbone areas into any non-backbone areas.



LEGEND

- ^
| Group Membership LSAs
- |-----| Area Border Router and
 Inter-Area Multicast Forwarder
- * Wild-Card Multicast
 Receiver Interface

■ Inter-Area Routing Architecture

To permit the forwarding of multicast traffic between areas, MOSPF introduces the concept of a "wild-card multicast receiver." A wild-card multicast receiver is a router that receives all multicast traffic generated in an area, regardless of the multicast group membership. In non-backbone areas, all inter-area multicast forwarders operate as wild-card multicast receivers. This guarantees that all multicast traffic originating in a non-backbone area is delivered to its inter-area multicast forwarder, and then if

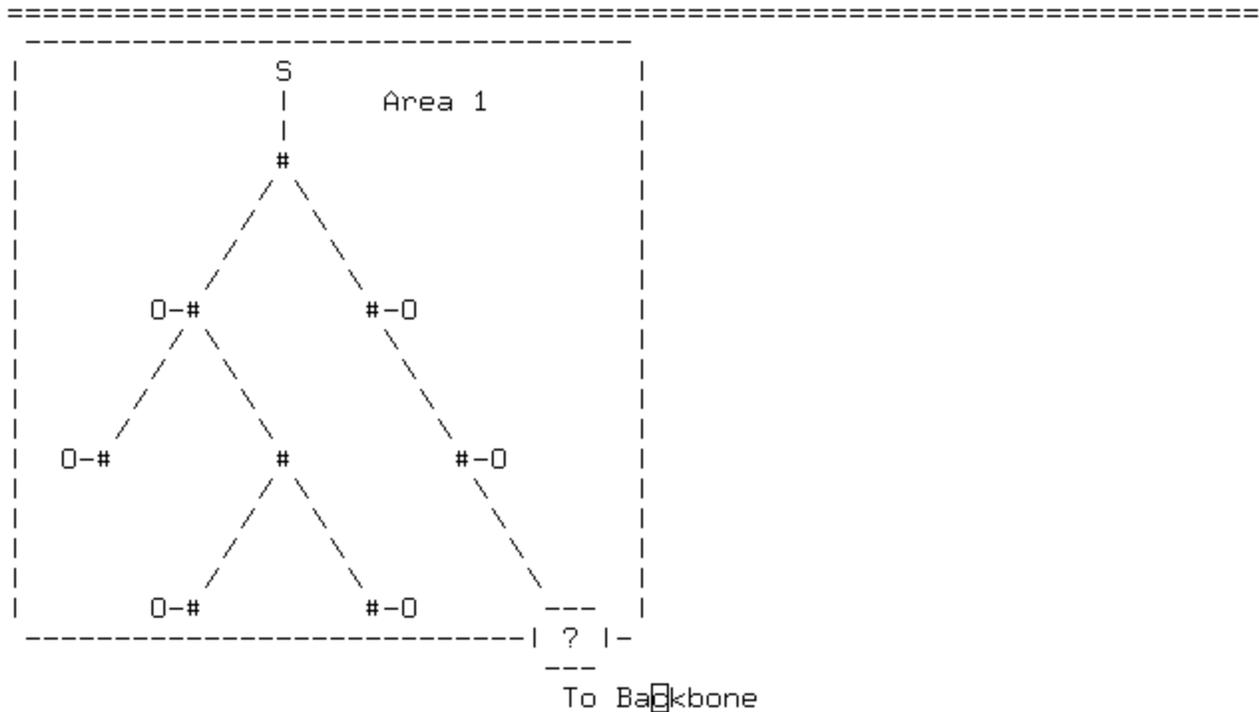
necessary into the backbone area. Since the backbone has group membership knowledge for all areas, the datagram can then be forwarded to group members residing in the backbone and other non-backbone areas. The backbone area does not require wild-card multicast receivers because the routers in the backbone area have complete knowledge of group membership information for the entire OSPF system.

Inter-Area Datagram Shortest-Path Tree

In the case of inter-area multicast routing, it is often impossible to build a complete datagram shortest-path delivery tree. Incomplete trees are created because detailed topological and group membership information for each OSPF area is not distributed to other OSPF areas. To overcome these limitations, topological estimates are made through the use of wild-card receivers and OSPF Summary-Links LSAs.

There are two cases that need to be considered when constructing an inter-area shortest-path delivery tree. The first involves the condition when the source subnetwork is located in the same area as the router performing the calculation. The second situation occurs when the source subnetwork is located in a different area than the router performing the calculation.

If the source of a multicast datagram resides in the same area as the router performing the calculation, the pruning process must be careful to ensure that branches leading to other areas are not removed from the tree. Only those branches having no group members nor wild-card multicast receivers are pruned. Branches containing wild-card multicast receivers must be retained since the local routers do not know if there are group members residing in other areas.

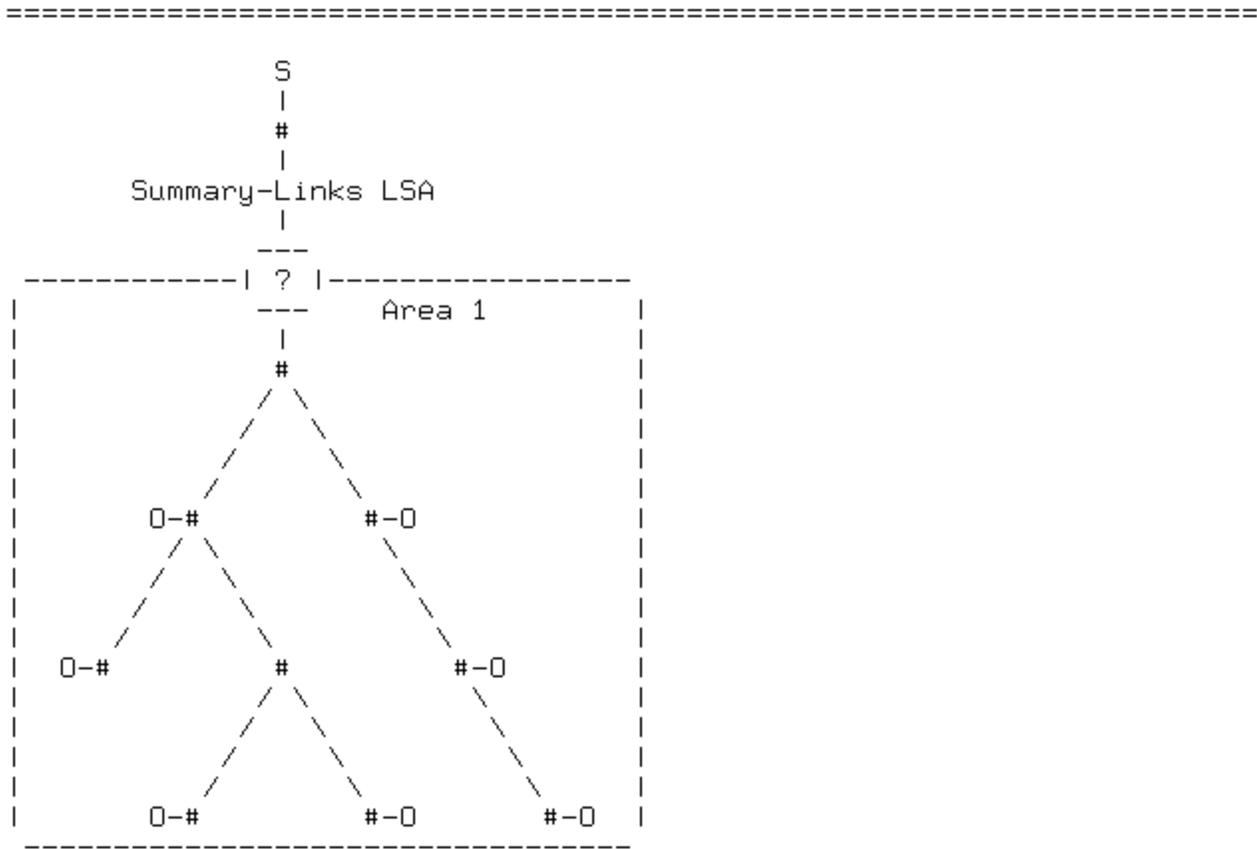


LEGEND

- S Source Subnetwork
- 0 Subnet Containing Group Members
- # Intra-Area MOSPF Router
- ? WildCard Multicast Receiver

Datagram Shortest Path Tree -Source in Same Area

If the source of a multicast datagram resides in a different area than the router performing the calculation, the details describing the local topology surrounding the source station are not known. However, this information can be estimated using information provided by Summary Links LSAs for the source subnetwork. In this case, the base of the tree begins with branches directly connecting the source subnetwork to each of the local area's inter-area multicast forwarders. The inter-area multicast forwarders must be included in the tree since any multicast datagrams originating outside the local area will enter the area via an inter-area multicast forwarder.



LEGEND

- S Source Subnetwork
- 0 Subnet Containing Group Members
- # Interarea MOSPF Router
- ? Intra-Area Multicast Forwarder

Fig 5. Shortest Path Tree -Source in Different Area

Since each inter-area multicast forwarder is also an ABR, it must maintain a separate link state database for each attached area. This means that each inter-area multicast forwarder is required to calculate a separate forwarding tree for each of its attached areas. After the individual trees are calculated, they are merged into a single forwarding cache entry for the (source, group) pair and then the individual trees are discarded.

Inter-Autonomous System Multicasting with MOSPF

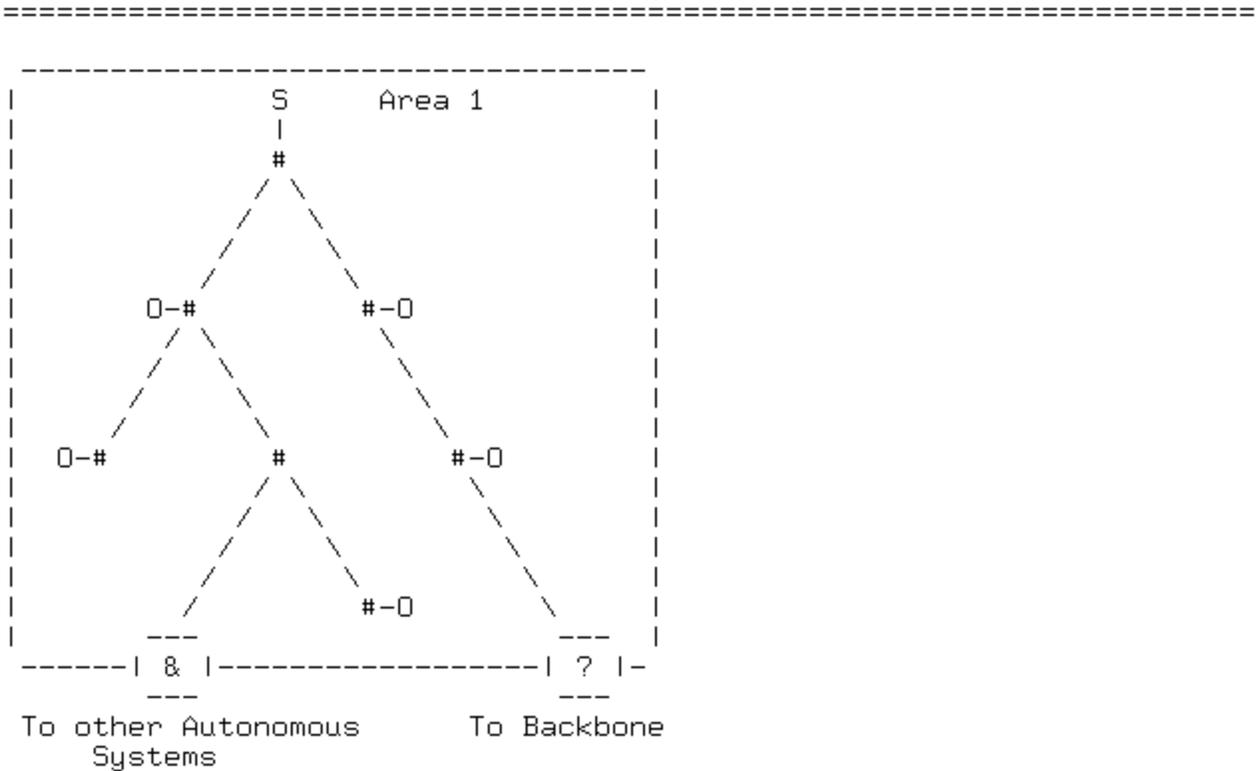
Inter-Autonomous System Multicasting involves the situation where a datagram's source and at least some of its destination group members reside in different Autonomous Systems. It should be emphasized that in OSPF terminology "inter-AS" communication also refers to connectivity between an OSPF domain and another routing domain which could be within the same Autonomous System.

To facilitate inter-AS multicast routing, selected Autonomous System Boundary Routers (ASBRs) are configured as "inter-AS multicast forwarders." MOSPF makes the assumption that each inter-AS

multicast forwarder executes an inter-AS multicast routing protocol (such as DVMRP) which forwards multicast datagrams in a reverse path forwarding (RPF) manner. Each inter-AS multicast forwarder functions as a wild-card multicast receiver in each of its attached areas. This guarantees that each inter-AS multicast forwarder remains on all pruned shortest-path trees and receives all multicast datagrams, regardless of the multicast group membership.

Three cases need to be considered when describing the construction of an inter-AS shortest-path delivery tree. The first occurs when the source subnetwork is located in the same area as the router performing the calculation. For the second case, the source subnetwork resides in a different area than the router performing the calculation. The final case occurs when the source subnetwork is located in a different AS (or in another routing domain within the same AS) than the router performing the calculation.

The first two cases are similar to the inter-area examples described in the previous section. The only enhancement is that inter-AS multicast forwarders must also be included on the pruned shortest path delivery tree. Branches containing inter-AS multicast forwarders must be retained since the local routers do not know if there are group members residing in other Autonomous Systems. When a multicast datagram arrives at an inter-AS multicast forwarder, it is the responsibility of the ASBR to determine whether the datagram should be forwarded outside of the local Autonomous System. Figure 6 illustrates a sample inter-AS shortest path delivery tree when the source subnetwork resides in the same area as the router performing the calculation.



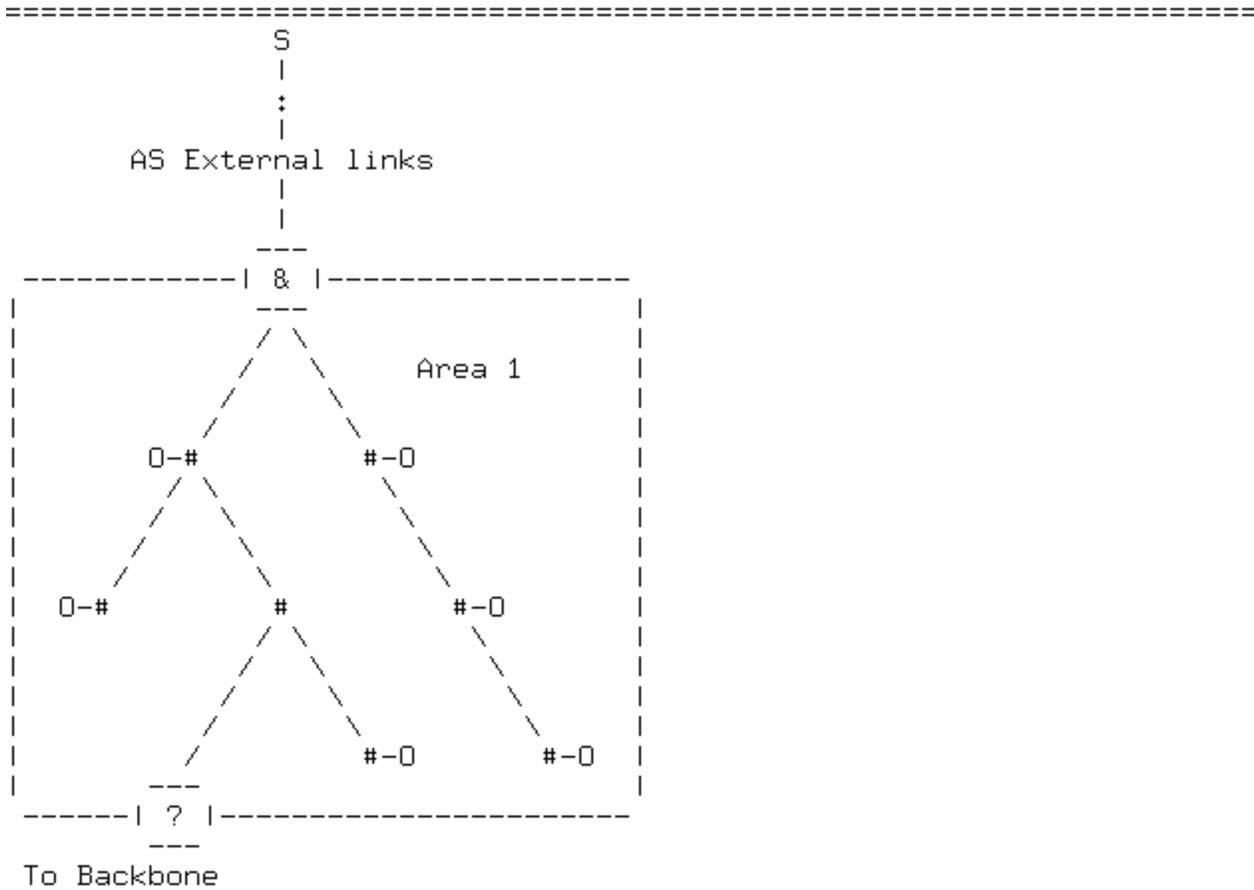
LEGEND

- S Source Subnetwork
- 0 Subnet Containing Group Members
- # Intra-Area MOSPF Router
- ? Inter-Area Multicast Forwarder
- & Inter-AS Multicast Forwarder

Figure 6 . Inter-AS Datagram Shortest Path Tree -Source in Same Area

If the source of a multicast datagram resides in a different Autonomous System than the router performing the calculation, the details describing the local topology surrounding the source station are not known. However, this information can be estimated using the multicast-capable AS External links describing the source subnetwork. In this case, the base of the tree begins with branches directly connecting the source subnetwork to each of the local area's inter-AS multicast forwarders.

Figure 7 shows a sample inter-AS shortest-path delivery tree when the inter-AS multicast forwarder resides in the same area as the router performing the calculation. If the inter-AS multicast forwarder is located in a different area than the router performing the calculation, the topology surrounding the source is approximated by combining the Summary-ASBR Link with the multicast capable AS External Link.



LEGEND

- S Source Subnetwork
- 0 Subnet Containing Group Members
- # Intra-Area MOSPF Router
- ? Inter-Area Multicast Forwarder
- & Inter-AS Multicast Forwarder

Figure 7 Inter-AS Datagram Shortest Path Tree -Source in Different AS

Version 2 of the Open Shortest Path First (OSPF) routing protocol is defined in RFC-1583. It is an Interior Gateway Protocol (IGP) specifically designed to distribute unicast topology information among routers belonging to a single Autonomous System. OSPF is based on link-state algorithms which permit rapid route calculation with a minimum of routing protocol traffic. In addition to efficient oute calculation, OSPF is an open standard that supports hierarchical routing, load balancing, and the import of external routing information.

The Multicast Extensions to OSPF (MOSPF) are defined in RFC-1584. MOSPF routers maintain a current image of the network topology through the unicast OSPF link-state routing protocol. MOSPF enhances the OSPF protocol by providing the ability to route multicast IP traffic. The multicast extensions to OSPF are built on top of OSPF Version 2 so that a multicast routing capability can be easily introduced into an OSPF Version 2 routing domain. The enhancements that have been added are

backwards compatible so that routers running MOSPF will interoperate with non-multicast OSPF routers when forwarding unicast IP data traffic. MOSPF, unlike DVMRP, does not provide support for tunnels.