River STONE NETWORKS TM

Support

Riverstone OSPF Support Page

RS OSPF Background

- Riverstone Routing Model
- Route Selection in Riverstone Routers
- RIB Memory Management and OSPF
- RS OSPF Implementation
- OSPF Metrics and Load Sharing
- Externals and Route Redistribution from/to OSPF
- Router ID, loopback interface
- OSPF Area Support
- OSPF Not-so-Stubby Areas (NSSA)

OSPF Configuration Tasks

- Router ID, add stub host
- Intra-area router in a normal area
- ABR in a normal area
- ASBR in a normal area
- Intra-area router in a stub area
- ABR in a stub area
- Intra-area router in a TSA
- ABR in a TSA
- Intra-area router in a NSSA
- ABR in a NSSA
- ASBR in a NSSA
- Configuring authentication
- Configuring summarization on area borders
- Configuring route filtering on area borders
- Configuring aggregation at the ASBR
- Configuring route filtering at the border of NSSA
- Passive interfaces – why to configure them
- Configuring Virtual Links

Design Topics

- How to generate a default route in OSPF?
- OSPF metrics. Equal cost multipath routing
- OSPF Network types
  - On Ethernet
  - On PPP
  - On POS
  - On ATM
- How to aggregate external routes on an ASBR?
- Controlling DR Election
- Why sometimes Ethernet interfaces are configured as Point-to-Multipoint?
- How to tag OSPF routes?
- What is an Alternative ABR and how does it work?

Troubleshooting OSPF

- Understanding the "ospf show interface" command
- Undestanding the "ospf show neighbor" command

Privacy Statement   Terms of Use   Feedback

Riverstone Routing Model

Riverstone Routing Model

RIB Model

RS Preference and Route Selection

OSPF Route Selection

Related Commands
Displaying the RIB and FIB
Displaying Currently Configured RS Preference Values

RIB Model

The RS router follows routing model similar to that of GateD (it is a derivative of GateD). Each routing protocol can be source of routing information. That routing information can be subjected to **Import Policies**, which affect whether or not the information will enter the **Routing Information Base (RIB).** Each BGP peer is considered to be a separate routing process (and a separate import policy can be applied to it).

Import policies may not be applied to routes representing directly connected interfaces, static routes, and aggregate routes. These will be in the RIB, as long as they are valid.

The RIB contains all routes that are valid and are not rejected by an import policy. It is normal for the RIB to contain multiple routes to the same prefix, but from different protocol sources. In this case, RS Preference is used as a measure of preference between different routing protocols.

The RS Route Selection Process, with the help of RS Route Preference, chooses the **Active Routes** from the RIB, and copies them into the **Forwarding Information Base (FIB)**. The FIB is used for packet forwarding, and contains a simple mapping between prefixes and next hops to be used for those prefixes.
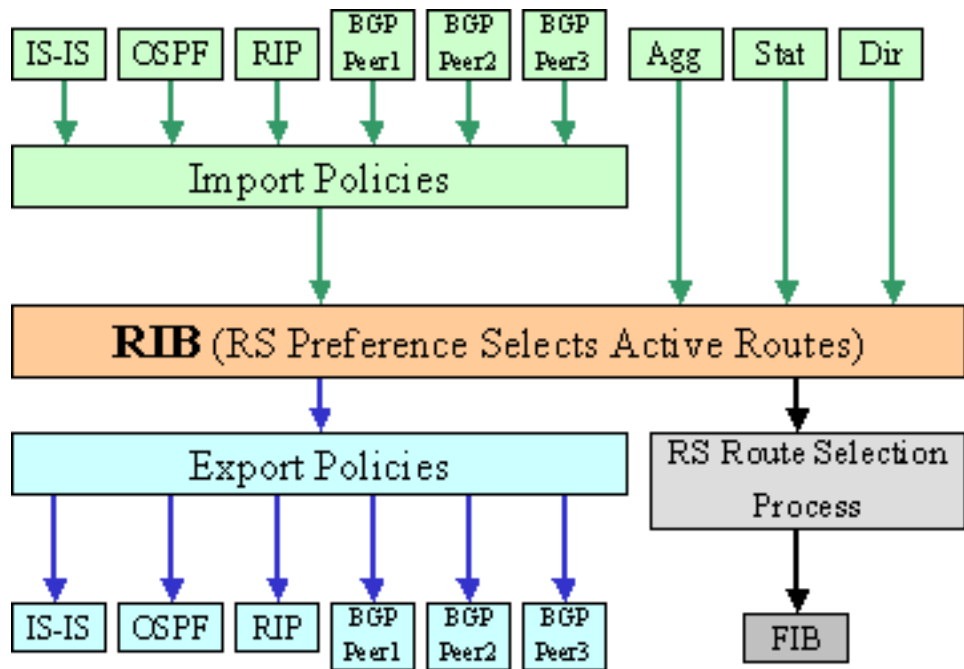
Figure 1: The RS RIB Model

Export Policies can be applied to the active routes in the FIB to control which of those will be exported (redistributed) to other routing protocols. Unlike Import policies, Export policies can be applied to prefixes from any source (including connected, static, and aggregate routes).

RS Preference and Route Selection

When multiple routes to the same prefix are present in the RIB, the router needs to decide which source (of the same information) will be considered more "trustworthy" than others. For that, the **RS Preference** value is used. Each routing protocol is assigned a default RS Preference value, which can be configurationally modified.

Prefix Type	Default Preference	Description
DIRECT interfaces	0	Routes to directly attached
STATIC	5	Static routes

OSPF	10	OSPF Internal route
ISIS_L1	15	IS-IS level 1 route
ISIS_L2	18	IS-IS level 2 route
RIP	100	Routing Information Protocol (RIP)
AGGREGATE	130	Aggregate route
OSPF_ASE	150	OSPF External route
ISIS_L1_EXT	160	IS-IS level 1 external route
ISIS_L2_EXT	165	IS-IS level 2 external route
BGP	170	Border Gateway Protocol

Table 1: Default RS Preferences (ROS 7.0 and above)

Lower preference values are more preferred. The default RS preference for internal OSPF routes can be changed with the following command:

RS(config)#ospf set preference <number>

The default RS preference for OSPF External routes can be changed with:

RS(config)#ospf set ase-defaults preference <number>

The default RS preference can be overridden by an Import Policy for specific prefixes. The valid values for RS Preference are 0-255.

OSPF Route Selection

The RS preference is only used to select an active routes when multiple routes are present **from different protocol sources**. Each routing protocol has its own mechanism to select routes and offer those for inclusion in the RIB. OSPF will choose the best route based on the Shortest Path Algorithm (SPF), which runs against the Link State Database (LSDB). SPF may select up to four Equal Cost Paths (for the same prefix) to be offered for inclusion in the RIB. RS Preference has no role in this selection, only OSPF Metric does.

Related Commands

Displaying the RIB and FIB

RS#ip-router show rib

Displays the contents of the RIB in a brief format (one line per entry).

```
RS1# ip-router show rib
Routing Tables:
Generate Default: no
Destinations: 8     Routes: 13
Holddown: 0     Delete: 3     Hidden: 4
Codes: Network - Destination Network Address
       S - Status + = Best Route, - = Last Active, * = Both
       Src - Source of the route :
       Ag - Aggregate, B - BGP derived, C - Connected
       R - RIP derived, St - Static, O - OSPF derived
       OE - OSPF ASE derived, D - Default
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2
       Next hop - Gateway for the route ; Next hops in use: 4
       Netif - Next hop interface
       Prf1 - Preference of the route, Prf2 - Second Preference of
the route
       Metrc1 - Metric1 of the route, Metrc2 - Metric2 of the
route
       Age - Age of the route
Network/Mask        S Src Next hop                 Netif Prf1 Metrc1
Metrc2      Age
-----------         - --- --------                 ----- ---- ------ ---
---         ---
1.1.1.1/32          *   C 1.1.1.1                     lo     0
0        0     19:31
1.1.1.1/32              O                                   -10
1        0      8:51
2.2.2.2/32          *   O 10.0.0.2                   TO-RS2   10
3        0      6:47
2.2.2.2/32              O 10.0.0.2                   TO-RS2  -10
2        0      6:47
10/30               *   C 10.0.0.1                   TO-RS2    0
```

```
0        0          7
10/30                        O 10.0.0.1                     TO-RS2     10
2        0      14:47
63/8                   *     C 63.228.71.1                   LAN1     0
0        0       9:41
63/8                         O 63.228.71.1                   LAN1    10
200          0       9:14
70.70.70/24            *     O 10.0.0.2                      TO-RS2    10
22        0       7:02
127/8                  *     St 127.0.0.1                     lo      0
0         0      19:30
127.0.0.1/32           *     C 127.0.0.1                      lo      0
0         0      19:31
```

Several prefixes are listed twice because they are heard from different protocol sources. The Active Route is selected based on the RS Route Selection process (using RS Route Preference).

RS#ip-router show route x.x.x.x/y detail

Displays all RIB entries for prefix x.x.x.x/y in detail. This is useful when one tries to understand the reasons for a particular Active Route selection.

```
RS1# ip-router show route 10.0.0.0/30 detail

    10                    mask 255.255.255.252
            entries 2     announce 1
            TSI:
                      Instability Histories:
        *Direct      Preference:    0
            *NextHop: 10.0.0.1             Interface: 10.0.0.1(TO-
RS2) !
            State: <Int Active Retain>
            Age: 10     Metric: 0     Metric2: 0     Tag: 0
            Task: IF
            Announcement bits(1):
                2-KRT
            AS Path: IGP (Id 1)
        OSPF     Preference:   10          Source: 0.0.0.0
            *NextHop: 10.0.0.1              Interface: 10.0.0.1(TO-
RS2) !
            State: <Int Gateway>
```

```
            Age: 20:40      Metric: 2      Metric2: 0      Tag: 0
            Task: OSPF
            No announcement bits set
            AS Path: IGP (Id 1)
```

RS#ip show routes
Displays the FIB.

```
RS1# ip show routes

Destination              Gateway                    Owner      Netif
-----------              -------                    -----      -----
1.1.1.1                  1.1.1.1                    -          lo0
2.2.2.2                  10.0.0.2                   OSPF       TO-RS2
10.0.0.0/30              directly connected         -          TO-RS2
63.0.0.0/8               directly connected         -          LAN1
70.70.70.0/24            10.0.0.2                   OSPF       TO-RS2
127.0.0.1                127.0.0.1                  -          lo0
```

Displaying Currently Configured RS Preference Values

RS#ip-router show route-preferences
Displays the current default RS Preference values for all protocols.

```
RS1# ip-router show route-preferences

Route              Prefrences:
-------            ----------

Direct              0
Static              5
OSPF                10
ISIS L1             15
ISIS L2             18
ISIS L1 External  160
ISIS L2 External  165
Default             20
Router Discovery  55
RIP                100
DIRECT AGGREGATE  0
```

```
AGGREGATE          130
OSPF ASE           150
IBGP/EBGP          170
```

Go back to OSPF Support Home

Last Updated June 3, 2002

RIB Memory Management

RIB Memory Management

Advanced Memory Management

Related Commands
Show Threshold Levels and Drops
Change Thresholds/Disable Advanced Memory Management

Advanced Memory Management

An advanced memory management technique was introduced in ROS 7.0 and above. Its goal is to prevent the RIB from ever running out of memory because of excessive routes received from routing protocols. Each routing protocol has its own means of limiting the number of routes it will offer to the RIB. Those are protocol-specific. The Advanced Memory Management provides a "second line of defense" overall.

Four memory threshold levels are introduced. When those levels of memory utilization are reached, the resulting action is protocol-specific. For OSPF and IS-IS (as link-state routing protocols), the Advanced Memory Management does not affect how their routes are populated in the RIB. The different thresholds have different actions with regard to BGP routes.

The default threshold levels are as follows:

Low Memory Threshold 1   : 60%
Low Memory Threshold 2   : 65%
Low Memory Threshold 3   : 70%
Low Memory Threshold 4   : 75%

Advanced memory management is enabled by default. Be aware that otherwise eligible prefixes may not be installed in the RIB because memory thresholds were reached.   It can be configurationally disabled, in which case a misbehaving routing process (e.g. a BGP peer announcing excessive amount of routes) can starve memory for other routing processes.

Related Commands

Show Threshold Levels and Drops

ip-router show summary drops

Display:
- ❍ Summary of routes in the RIB
- ❍ Information on routes that we NOT installed in the RIB because thresholds we reached
- ❍ The percentages for the current Level 1/2/3/4 thresholds.
- ❍ The current threshold value
- ❍ Memory statistics for GateD (the routing process)

```
RS1# ip-router show summary drops

Summary of routes in RIB
------- -- ------ -- ---
Number of Unique routes : 8
Number of routes         : 13
Kernel routes            : 0
Direct routes            : 4
Static routes            : 1
RIP routes               : 0
OSPF routes              : 8
OSPF ASE routes          : 0
ISIS level 1 routes      : 0
ISIS level 2 routes      : 0
BGP routes               : 0
Other Protocol routes    : 0
Hidden routes            : 4


Routes not added due to low memory              : 0
BGP Routes removed to accomodate new routes : 0
BGP Routes removed due to low memory        : 0
RIP Routes removed to accomodate new routes : 0
RIP Routes removed due to low memory        : 0
Low Memory Threshold 1                          : 60
```

```
Low Memory Threshold 2                          : 65
Low Memory Threshold 3                          : 70
Low Memory Threshold 4                          : 75
Current Threshold value                         : 0
Malloc Percentage Used                          : 6
Gated Free Memory                               : 32612
Gated Free Memory(Global counter)               : 32612
```

Change Thresholds/Disable Advanced Memory Management

ip-router global set memory-threshold …
In CONFIG mode. Disables Advanced Memory Management, or changes the threshold levels for it.

```
RS1(config)# ip-router global set memory-threshold disable
RS1(config)# save active
```

To disable Advanced Memory Management

OR

```
RS1(config)# ip-router global set memory-threshold level-4 80
RS1(config)# save active
```

To change the Level-4 threshold percentage from its default (75%) to 80%.

Go back to
Last Updated June 3, 2002

Riverstone OSPF Implementation

Riverstone's OSPF implementation is compliant with RFC 2328 – OSPF Version 2, and RFC 1587 – NSSA Option. Following is a list of the key features:

❍ Area Support – Normal, Stub, Totally Stubby, and Not-So-Stubby Areas (NSSA) are supported.

❍ NSSA option – as specified in RFC 1587

❍ Virtual Links – Virtual links are supported

❍ Authentication – plain text or MD5 area authentication is supported

❍ Route Redistribution

   ❍ External routes can be **redistributed in** OSPF from the following sources
      ■ Directly connected interfaces
      ■ Static routes
      ■ Aggregate Routes
      ■ RIP
      ■ IS-IS
      ■ BGP
      ■ Virtual IP addresses associated with load-balancing configuration

   ❍ OSPF routes can be **redistributed from** OSPF to the following destination protocols
      ■ RIP
      ■ IS-IS
      ■ BGP

❍ Flexible interface parameters – manual/automatic OSPF interface cost, timers, interface type, priority, etc.

❍ Opaque LSA Support – as defined in RFC 2370

❍ OSPF Traffic Engineering (OSPF-TE) – as defined in draft-katz-yeung-ospf-traffic-06.txt

Go back to OSPF Support Home

Last Updated June 3, 2002

OSPF Metric and Equal Cost Multi-Path Forwarding

OSPF Metric and Equal Cost Multi-Path Forwarding

OSPF Interface Cost

Default OSPF Costs

Every OSPF interface is assigned a cost in the range 1-65535. It represents the cost for traffic to **EXIT** the router from the interface. The default value of the cost is calculated from the following formula:

OSPF Cost = $2\text{x}10^8$/<Interface Speed in bps>

Based on that formula, the following OSPF costs will be assigned to these interfaces, once OSPF is enabled on them.

Interface Type (Speed)	Default OSPF Cost
Gigabit Ethernet	2
10/100 Mbps Ethernet	20
T3	44
T1	1333

When an IP interface is assigned to a VLAN that includes physical ports with different speeds, the fastest port is considered when calculating the OSPF cost for that interface. Similarly, when OSPF is enabled on an IP interface running over a SmartTrunk, the fastest link of that SmartTrunk is taken into account. For ports that may have variable speed (e.g. 10/100 Ethernet ports), the speed of the port at the time of its addition to OSPF is taken as a default cost.

Manual Configuration of OSPF Costs

In some cases, accepting the default OSPF costs is inappropriate. OSPF interface costs can be overridden on a per-interface basis with the command:

RS(config) ospf set interface <Name> cost <Number>

Reference Bandwidth

By default, the OSPF interface cost is calculated as $2\text{x}10^9$/Interface speed (in bps). This will result in calculating OSPF cost of 1 for all interfaces faster than 1 Gbps. In some cases that might be undesirable.   The reference bandwidth feature allows flexibility in the OSPF cost assignments, without forcing you to manually assign OSPF costs to every interface. The reference bandwidth replaces the 10^9 component in the above formula. The reference bandwidth is specified in Mbps, i.e. the line:

*RS(config)# ospf set ref-bwdth 10000*

Specifies the reference bandwidth to be 10Gbps, resulting in the following calculation of OSPF interface costs:

Interface Type (Speed)	Default OSPF Cost
10 GE	2
1 GE	20
100 Mbps	200
Etc.	

The reference bandwidth is a box-wide setting, and affects the default OSPF cost calculation for all interfaces of the router.

OSPF Route Metric

Every OSPF route has a metric associated with it. It is calculated as the sum of the interface costs of the routers in the path from the router calculating the metric to the destination represented by the route. The metric of routes may be asymmetrical, even though the same physical path is traversed. This is particularly true when manual cost assignment is done, or the network is comprised of multi-vendor equipment, which calculates costs differently. The following figure is an example of asymmetrical OSPF metric.
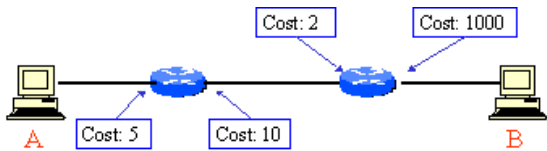

- A to B – Metric 1010
- B to A – Metric 7

Equal Cost Multi-Path Forwarding

OSPF can maintain up to 4 equal-cost paths (ECMPs) for the same destination (technically, those should be called "equal-metric" paths, since it is the "Metric", rather than the "Cost" that is used here). Those ECMPs will be installed as a single prefix with multiple next hops in the RIB, and if all of the next hops are viable, they will be installed in the FIB. See the Riverstone Routing Model document for more information on RIB/FIB.

The example below illustrates the existence of 3 ECMPs from A to D.
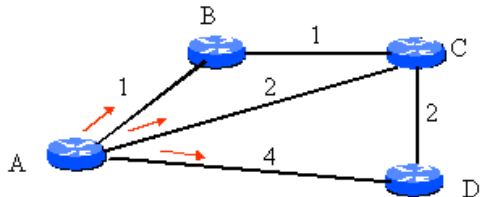
There are two aspects of Equal Cost Multi-Path Forwarding:
1. Calculating ECMP routes and maintaining them in the FIB
2. Actually forwarding traffic across those ECMPs.

As described above, the Riverstone router will calculate up to four ECMPs in the FIB using any routing protocol (including static routes). It will also **forward** based on up to four next-hop gateways for the same destination.

The load-sharing that results from ECMP forwarding is **flow-based**. Packet sequencing within a flow is preserved when load-sharing.

ECMP Forwarding before ROS 8.0.0.0

In firmware releases before ROS 8.0.0.0 flows are distributed in a round-robin without regard of the contents of the IP header.

ECMP Forwarding in ROS 8.0.0.0 and later

In ROS 8.0.0.0 the default load-distribution method was changed from round-robin to hash-based. This results in better distribution of flows, especially in very bursty traffic patterns.

A hash value is used to choose the next-hop that will be used for a particular flow. This hash value is derived from the Source and Destination IP addresses of the flow. The default hash algorithm (Hash 1) performs the following function:

(Source-IP XOR Destination-IP) MOD (Number of Gateways)

in order to derive which gateway to use for a flow. The 3 hash functions perform permutations with different parts of the Source and Destination IP addresses. The hash function used can be changed with the command:

```
ip set multipath-hash-variant <1-4>
```

If the ECMP forwarding based on hash algorithms is not desired, you can revert back to round-robin load sharing by issuing the command

```
ip set multipath-hash-variant disable
```

OSPF Route Types and Metric

OSPF distinguishes between four types of routes:

1. Intra-area – routes representing destinations within the same OSPF area
2. Inter-area – routes representing destinations within another OSPF area
3. ASE Type 1 (External Type 1) – routes from external sources redistributed as Type 1
4. ASE Type 2 (External Type 2) – routes from external sources redistributed as Type 2

OSPF has its own internal preference mechanism (not to be confused with RS Preference). This mechanism dictates that if certain prefix is heard from different OSPF routers as different types of routes, the following hierarchy will be followed to make a choice:

Intra-area routes are always preferred over Inter-area routes, which are always preferred over Type 1 Externals, which in turn are always preferred over Type 2 Externals. This choice overrides the OSPF Metric for a route, meaning that an Intra-area route with higher metric will be preferred over an Inter-area route with lower metric value.

The following diagram illustrates the different preference and metric values and the places they are used in the route selection process.
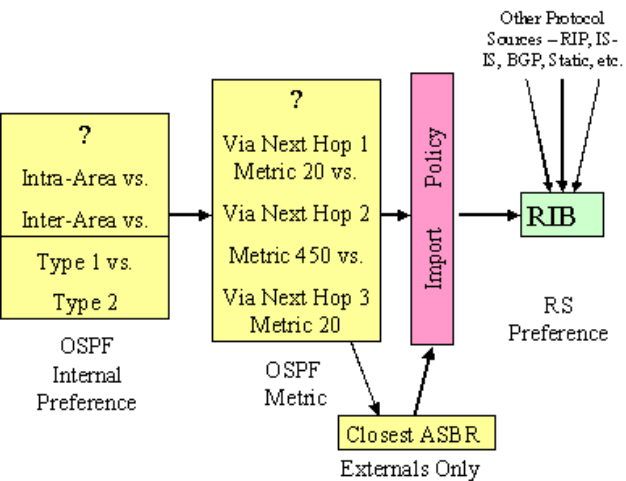
The OSPF internal preference will be used to select between competing routes (to the same prefix) of different OSPF type. Once that choice is made, OSPF Metric will be used to select the best route within certain type (up to 4 ECMPs can be chosen here). An optional import policy can be applied to control if the selected routes are going to make it in the RIB. Once the routes are in the RIB, RS Preference will be used to prefer the same route that arrives possibly from other protocol sources.

For OSPF Equal Cost Multi-Path Routing to occur, the multiple routes must be (1) of the same type and (2) with the same cost. Routes with the same OSPF cost but different OSPF type (e.g. intra-area and inter-area) will not be used for load balancing, instead the route with the more preferred type (see above) will be chosen.

An Additional Consideration for External Routes

In order for ECMP to occur across two or more external routes, an additional requirement must be met – the distances to the respective ASBRs that injected those Externals must be the same (in addition to the type of External and its cost).

Example of Using OSPF Equal Cost Multi-path Routes

Equal cost routes are calculated by OSPF for all route types – Intra-area, Inter-area, and External. The following example illustrates the usage of an Inter-area Equal Cost Multiipath route.

All routers are configured with router ID, shown on the figure by the name of the router. All router IDs are injected in OSPF by the command "ospf add stub-host x.x.x.x/32 to-area <AREA_ID> cost 33". The cost of 33 has been chosen arbitrarily.
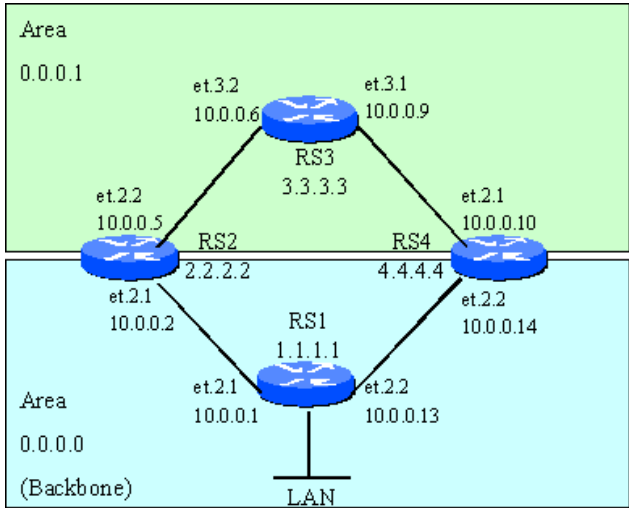
RS1 Configuration
```
interface create ip TO-RS2 address-netmask 10.0.0.1/30 port et.2.1
interface create ip TO-RS4 address-netmask 10.0.0.13/30 port et.2.2
interface create ip LAN address-netmask 63.228.71.1/24 port et.2.8
interface add ip lo0 address-netmask 1.1.1.1/32
ip-router global set router-id 1.1.1.1
ospf create area backbone
ospf add interface TO-RS2 to-area backbone
ospf add interface TO-RS4 to-area backbone
ospf add stub-host 1.1.1.1 to-area backbone cost 33
ospf start
system set name RS1
```

RS2 Configuration
```
interface create ip TO-RS1 address-netmask 10.0.0.2/30 port et.2.1
interface create ip TO-RS3 address-netmask 10.0.0.5/30 port et.2.2
interface add ip lo0 address-netmask 2.2.2.2/32
ip-router global set router-id 2.2.2.2
ospf create area backbone
ospf create area 0.0.0.1
ospf add interface TO-RS1 to-area backbone
ospf add interface TO-RS3 to-area 0.0.0.1
ospf add stub-host 2.2.2.2 to-area backbone cost 33
ospf start
system set name RS2
```

RS3 Configuration
```
interface create ip TO-RS2 address-netmask 10.0.0.6/30 port et.3.2
interface create ip TO-RS4 address-netmask 10.0.0.9/30 port et.3.1
interface add ip lo0 address-netmask 3.3.3.3/32
ip-router global set router-id 3.3.3.3
ospf create area 0.0.0.1
ospf add interface TO-RS2 to-area 0.0.0.1
ospf add interface TO-RS4 to-area 0.0.0.1
ospf add stub-host 3.3.3.3 to-area 0.0.0.1 cost 33
ospf start
system set name RS3
```

RS4 Configuration
```
interface create ip TO-RS1 address-netmask 10.0.0.14/30 port et.2.2
interface create ip TO-RS3 address-netmask 10.0.0.10/30 port et.2.1
interface add ip lo0 address-netmask 4.4.4.4/32
ip-router global set router-id 4.4.4.4
ospf create area backbone
ospf create area 0.0.0.1
ospf add interface TO-RS1 to-area backbone
ospf add interface TO-RS3 to-area 0.0.0.1
ospf add stub-host 4.4.4.4 to-area backbone cost 33
ospf start
system set name RS4
```

Router RS3 has a router ID of 3.3.3.3. From RS1, this route (3.3.3.3/32) is reachable via the two ABR routers – RS2 and RS4. Furthermore, the cost via both paths is calculated at 73:

73 =
33 +   (the cost we assigned the router ID route)
20 +   (the default OSPF cost of a Fast Ethernet interface)
20   (the default OSPF cost of a Fast Ethernet interface)

Since we have not assigned any specific costs to OSPF interfaces, defaults were accepted. Additionally, the 3.3.3.3/32 route is advertised via both paths as an OSPF Inter-Area route (OSPF_IA). Since both the OSPF metric and type of the two routes are the same, equal cost multi-path routing will commence:

```
RS1# ip show routes

Destination         Gateway             Owner    Netif
-----------         -------             -----    -----
1.1.1.1             1.1.1.1             -        lo0
2.2.2.2             10.0.0.2            OSPF     TO-RS2
3.3.3.3             10.0.0.2            OSPF_IA  TO-RS2
                    10.0.0.14           OSPF_IA  TO-RS4
4.4.4.4             10.0.0.14           OSPF     TO-RS4
10.0.0.0/30         directly connected  -        TO-RS2
10.0.0.4/30         10.0.0.2            OSPF_IA  TO-RS2
10.0.0.8/30         10.0.0.14           OSPF_IA  TO-RS4
10.0.0.12/30        directly connected  -        TO-RS4
63.228.71.0/24      directly connected  -        LAN
127.0.0.1           127.0.0.1           -        lo0
RS1#
```

Last Updated June 6, 2002

External Route Handling and Redistribution

External Route Handling and Redistribution

External Routes

The OSPF protocol allows for routes originating from non-OSPF sources to be redistributed into OSPF as OSPF External Routes (sometimes called OSPF ASE routes). Such redistribution can only occur on an ASBR (Autonomous System Boundary Router). An ASBR can be configured in Normal and NSSA areas, but **not** in Stub and Totally Stubby areas.

OSPF floods all external routes into all Normal areas. External routes are not propagated into Stub, Totally Stubby, and NSSA areas. External routes are injected using Type 5 and Type 7 LSAs. Type 5 LSAs are only used in Normal areas, and Type 7 LSAs are only used in NSSA. An ABR of and NSSA can translate Type 7 to Type 5 LSAs into the backbone.

All external routes are redistributed into the OSPF domain with an initial metric (cost). Depending on how that cost is handled in OSPF, all external routes are classified in two groups: Externals with Type 1 Metric and Externals with Type 2 Metric.

Type 1 Externals

Type 1 Externals have cumulative cost, i.e. the cost of a Type 1 External route, as observed on Router A within the OSPF network, is derived as the sum of two components:

1.   The cost assigned to the route at the redistribution point (ASBR)
2.   The OSPF cost from Router A to the ASBR that redistributed the External.

It is appropriate to use Type 1 Externals when the external cost of the route is comparable with the OSPF costs. For example, if we are redistributing from RIP, where the cost is hop count, and we have configured OSPF to also use hop count as cost, then it will make sense to redistribute routes from RIP as Type 1 Externals.

Type 2 Externals

For Type 2 Externals, the external metric is ignored and the cost of the route is equal to the cost to reach the ASBR that redistributes it

Type 1 Externals are always preferred over Type 2 Externals, regardless of their metric. By default, External routes are redistributed with Type 2 metric.

Redistribution into OSPF from Other Sources

Redistribution is a shortcut way of configuring an export policy (for more details on export policies, check the Users Guide, as well as the page on [Riverstone Routing Model](#)). Redistribution (as an export policy) takes active routes from the RIB that originated from a given source protocol, and advertises them to a target protocol. When redistributing routes into OSPF, the target protocol can be either OSPF, or OSPF-NSSA. This depends on where the redistributing ASBR is located – in a Normal or NSSA area. In the former case, external routes will be injected as Type 5 LSAs, while in the latter case they will be injected as Type 7 LSAs. The generic syntax of the redistribution command is:

```
ip-router policy redistribute from-proto <SOURCE> to-proto <DESTINATION> metric
<#> [ase-type 1|2] [filter xxx | network y.y.y.y/z]
```

The "filter" or "network" statements can be optionally used to define only a subset of routes from the source to be redistributed. The redistribution command is illustrated with the following examples:

Redistribution from directly connected interfaces

```
ip-router policy redistribute from-proto direct to-proto ospf network
10.10.10.0/30 metric 20
```

Redistributes into OSPF as a Type 5 external the route associated with one particular directly connected interface. The external will be redistributed with a Type 2 external metric (default); its metric will be 20 and will remains 20 as the route is propagated throughout the OSPF routing domain.

Redistribution from aggregate

```
ip-router policy redistribute from-proto aggregate to-proto ospf metric 10 ase-
type 1
```

```
ip-router policy summarize 12.12.12.0/24
ip-router policy summarize 124.34.0.0/16
ip-router policy summarize 33.0.0.0/8
```

Redistributes into OSPF as a Type 5 externals all routes of type Aggregate (as seen by the ASBR). In this example, three Aggregate routes have been defined, so all of them will be redistributed. The externals will be redistributed with a Type 1 external metric; at the ASBR their cost will be 10, but it will increase as the routes are propagated within the OSPF routing domain

Redistribution from static routes

```
ip-router policy redistribute from-proto static to-proto ospf-nssa metric 10
filter SELECTED_FEW

ip-router create filter SELECTED_FEW 10.10.0.0/16
ip-router add filter SELECTED_FEW 22.13.2.0/24
ip-router add filter SELECTED_FEW 200.0.0.0/19
```

Redistributes into OSPF as a Type 7 externals all routes of type Static (as seen by the ASBR), that match filter SELECTED_FEW. This is a command that will be issued on an ASBR in a NSSA area. In this example, the filter SELECTED_FEW is defined to contain three prefixes.

If any aggregation of the static routes is to be performed, it must be done at the ASBR, and instead of redistributing static routes, aggregates should be distributed.

Redistribution from RIP

The example below illustrates the default redistribution of RIP routes into OSPF. By default, the RIP routes are redistributed into OSPF as Type 2 Externals, and their OSPF cost is assigned to be equal to the cost to the redistributing ASBR (RS2)
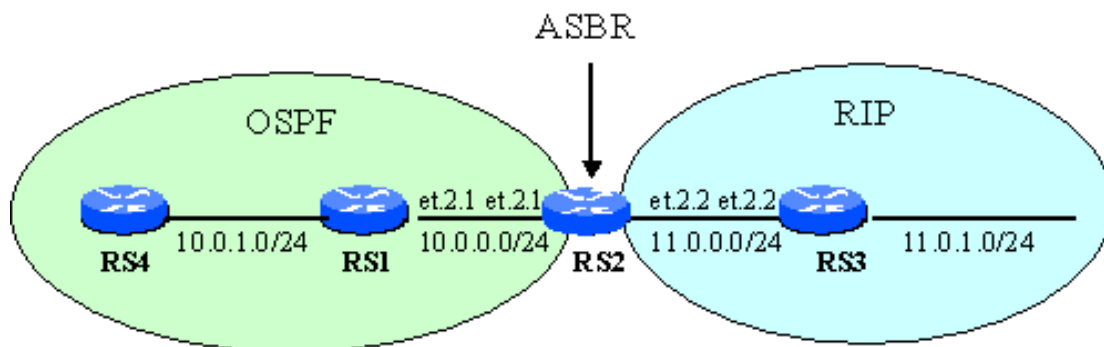
RS2 Configuration (ASBR)
```
interface create ip TO-RS1 address-netmask 10.0.0.2/24 port et.2.1
interface create ip TO-RS3 address-netmask 11.0.0.2/24 port et.2.2
```

```
interface add ip lo0 address-netmask 2.2.2.2/32
ip-router global set router-id 2.2.2.2
ip add route default gateway 127.0.0.1 no-install
ip-router policy redistribute from-proto rip to-proto ospf
rip add interface TO-RS3
rip set interface TO-RS3 version 2
rip start
ospf create area backbone
ospf add interface TO-RS1 to-area backbone
ospf add stub-host 2.2.2.2 to-area backbone cost 1
ospf add interface TO-RS3 to-area backbone
ospf set interface all cost 1
ospf start
system set name RS2
```

With this configuration, RIP routes are redistributed as Type 2 Externals, and their RIP cost at RS2 is ignored.

Redistribution from IS-IS

```
ip-router policy redistribute from-proto isis-level-1|isis-level-2 to-proto
ospf
```

Redistribution from BGP

```
ip-router policy redistribute from-proto source-as <#> to-proto ospf
```

Redistribution from OSPF into Other Protocols

When redistributing OSPF routes into other routing protocols, the source protocol is OSPF and the target protocol can be RIP, IS-IS, or BGP.

```
ip-router policy redistribute from-proto ospf to-proto rip|isis|bgp
[<additional modifiers>]
```

The additional modifiers here will be specific for each destination protocol.

Go back to OSPF Support Home

Last Updated June 3, 2002

Basic OSPF Configuration Tasks - Examples

Basic OSPF Configuration Tasks - Examples

Router ID

Requirements and Configuration

Loopback Interface

Recommended Configuration

Normal Area

An Intra-Area Router

An Area Border Router (ABR)

Autonomous System Boundary Router (ASBR)

Stub Area

Intra-Area Router

ABR

ASBR

Totally Stubby Area (TSA)

Intra-Area Router

ABR

ASBR

Not-So-Stubby Area (NSSA)

Intra-Area Router

ABR

ASBR

Router ID

Requirements and Configuration

Every OSPF router must be identified by a unique Router ID. The Router ID has the format of an IP address (dotted decimal). The router ID is configured with the following command:

ip-router global set router-id <x.x.x.x>
where x.x.x.x is the selected router ID.

Loopback Interface

The RS router (similarly to UNIX hosts) has a loobpack interface, which is always active, even though it is not configured. Its name and IP address are **lo0** and **127.0.0.1** respectively. The loopback interface is very useful for management of the router, because it is always. In order for the loopback interface to be reached, a secondary (routable) IP address needs to be added to it. The command for that is:

interface add ip lo0 address-netmask x.x.x.x/y
where x.x.x.x/y are the secondary IP address and mask for the loopback interface.

Recommended Configuration

When the Router ID is not configured, it is selected in the following manner:
1. If the loopback interface is configured, it is selected for Router ID
2. If the loopback interface is NOT configured, the IP address of the first active interface is selected as Router ID. If that IP inteface has multiple IP addresses assigned to it, then the numerically LOWEST one of them is selected for Routr ID.

Leaving the Router ID unconfigured and therefore assigned to a physical IP interface is not advisable since physical interfaces are unstable and their IP addresses can be configurationally changed, which will cause the OSPF router to drop and recreate its neighbor adjacencies, leading to unnecessary instability. ***Riverstone Networks recommends that Router ID be always manually configured and be the same as the loopback interface.***

Example of recommended configuration:

interface add ip lo0 address-netmask 21.21.21.21/32
ip-router global set router-id 21.21.21.21
ospf add stub-host 21.21.21.21. to-area backbone cost 1

The last command from this example will add the host route associated with the Router ID to the Link State Database of Area Backbone and will assign it a cost of 1. This line is not mandatory part of the Router ID configuration, but is very helpful when using the Router ID for management, BGP peering, or other purposes. On an ABR, this route can be added to either area configured on the ABR. The route, when added in this manner, is subject to normal OSPF summarization at the ABR.
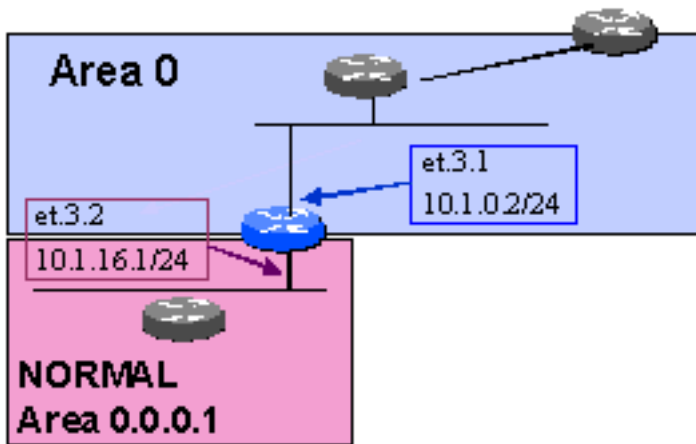
Normal Area

An Intra-Area Router
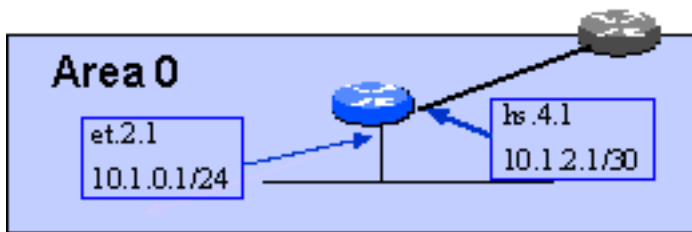
port set hs.4.1 wan-encapsulation ppp speed 45000000
interface create LAN ip address-netmask 10.1.0.1/24 port et.2.1
interface create WAN ip address-netmask 10.1.2.1/30 port hs.4.1
ospf create area backbone
ospf add interface LAN to-area backbone
ospf add interface WAN to-area backbone
ospf start

An Area Border Router (ABR)

The example below illustrates an ABR configuration between two Normal areas – Backbone and 0.0.0.1.

interface create In-Area-0 ip address-netmask 10.1.0.2/24 port et.3.1
interface create In-Area-1 ip address-netmask 10.1.16.1/24 port et.3.2
ospf create area backbone
ospf create area 0.0.0.1
ospf add interface In-Area-0 to-area backbone
ospf add interface In-Area-1 to-area 0.0.0.1
ospf start

Autonomous System Boundary Router (ASBR)

This example illustrates the configuration of an ASBR that attaches to a Normal area.
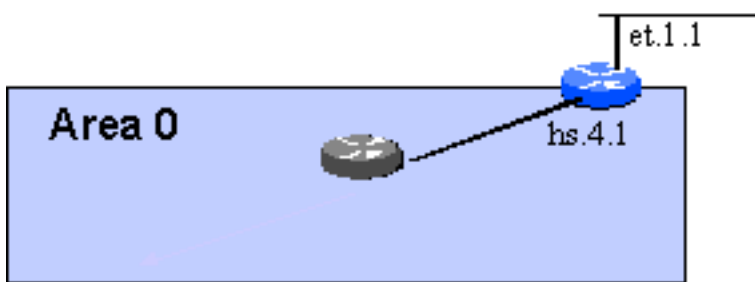
port set hs.4.1 wan-encapsulation ppp speed 45000000
interface create LAN ip address-netmask 192.168.0.1/24 port et.1.1
interface create WAN ip address-netmask 10.1.2.2/30 port hs.4.1
ospf create area backbone
ospf add interface WAN to-area backbone
ospf start
ip-router policy redistribute from-proto direct to-proto ospf

ASBRs always introduce Externals for the OSPF domain. The above example will inject 192.68.0.0/24 as an External route. It will be flooded throughout all normal areas of the OSPF network.

Stub Area

Intra-Area Router

interface create LAN ip address-netmask 10.1.16.2/24 port et.1.1
ospf create area 0.0.0.1
ospf set area 0.0.0.1 stub
ospf add interface LAN to-area 0.0.0.1
ospf start
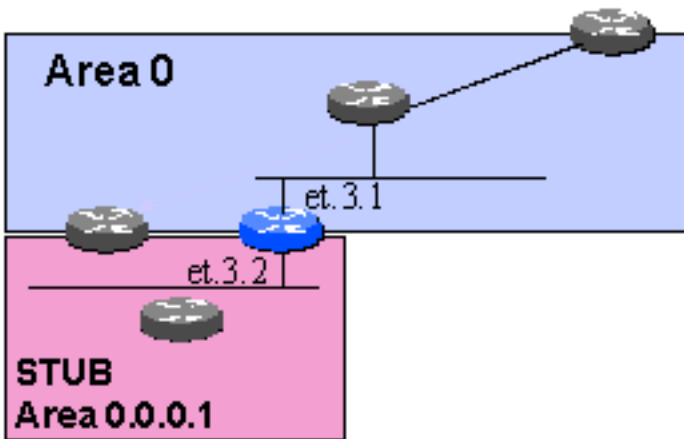
ABR

interface create In-Area-0 ip address-netmask 10.1.0.2/24 port et.3.1
interface create In-Area-1 ip address-netmask 10.1.16.1/24 port et.3.2
ospf create area backbone
ospf create area 0.0.0.1
ospf set area 0.0.0.1 stub stub-cost 22
ospf add interface In-Area-0 to-area backbone
ospf add interface In-Area-1 to-area 0.0.0.1
ospf add summary-range 10.1.0.0/20 to-area backbone
ospf add summary-range 10.1.16.0/20 to-area 0.0.0.1
ospf start

NOTE: The "stub-cost <NUMBER>" clause is a mandatory element in the configuration of the ABR to a stub area. Without it, no default route will be injected into the Stub area.

ASBR

ASBR cannot be configured in a Stub Area. Only Normal Areas and NSSA support ASBRs.

Totally Stubby Area (TSA)

Intra-Area Router

interface create LAN ip address-netmask 10.1.16.2/24 port et.1.1
ospf create area 0.0.0.1
ospf set area 0.0.0.1 stub no-summary
ospf add interface LAN to-area 0.0.0.1
ospf start

ABR

interface create In-Area-0 ip address-netmask 10.1.0.2/24 port et.3.1
interface create In-Area-1 ip address-netmask 10.1.16.1/24 port et.3.2
ospf create area backbone
ospf create area 0.0.0.1
ospf set area 0.0.0.1 stub stub-cost 22 no-summary
ospf add interface In-Area-0 to-area backbone
ospf add interface In-Area-1 to-area 0.0.0.1

ospf add summary-range 10.1.0.0/20 to-area backbone
ospf add summary-range 10.1.16.0/20 to-area 0.0.0.1
ospf start

ASBR

ASBR cannot be configured in a Totally Stubby Area. Only Normal Areas and NSSA support ASBRs.
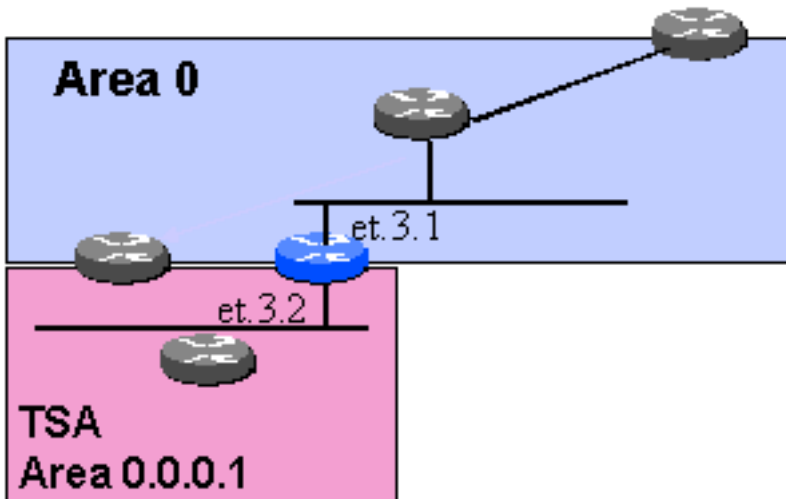

Not-So-Stubby Area (NSSA)

Intra-Area Router

interface create LAN ip address-netmask 10.1.16.2/24 port et.1.1
ospf create area 0.0.0.1
ospf set area 0.0.0.1 nssa
ospf add interface LAN to-area 0.0.0.1
ospf start
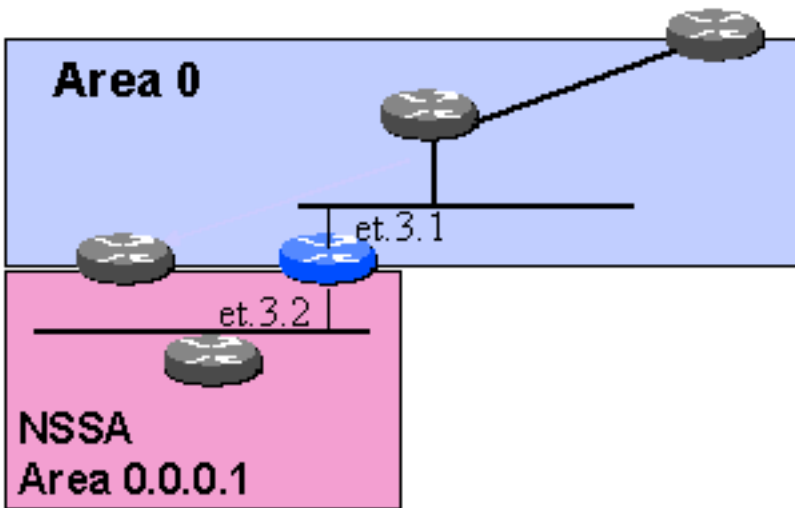
ABR

interface create In-Area-0 ip address-netmask 10.1.0.2/24 port et.3.1
interface create In-Area-1 ip address-netmask 10.1.16.1/24 port et.3.2
ospf create area backbone
ospf create area 0.0.0.1
ospf set area 0.0.0.1 nssa nssa-metric 4
ospf add interface In-Area-0 to-area backbone
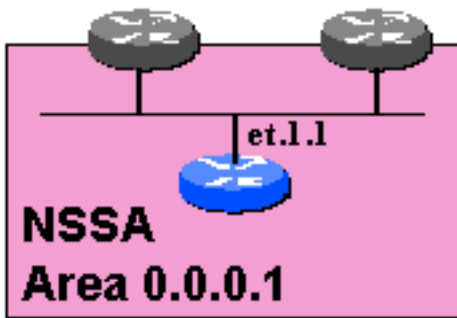*ospf add interface In-Area-1 to-area 0.0.0.1*
ospf start

The "nssa-metric" clause will cause the ABR to inject a default route into the NSSA with a cost of 4.

ASBR
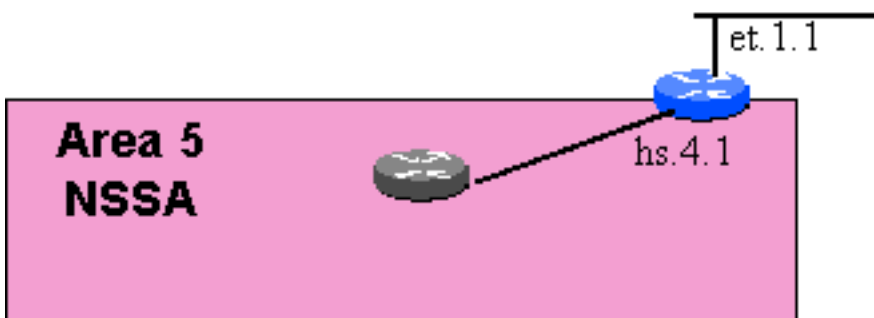
port set hs.4.1 wan-encapsulation ppp speed 45000000
interface create LAN ip address-netmask 192.168.0.1/24 port et.1.1
interface create WAN ip address-netmask 10.1.2.2/30 port hs.4.1
ospf create area backbone
ospf add interface WAN to-area backbone
ospf start
ip-router policy redistribute from-proto direct to-proto ospf-nssa

Note that in NSSA, the destination protocol for the redistribution command must be "ospf-nssa". This will cause the external routes to be redistributed as Type 7 LSAs. NSSA does not support Type 5 LSAs, therefore redistribution to protocol "ospf" will not work here.

To find more information on NSSA support on the RS, read here.

Go back to OSPF Support Home
Last updated June 3, 2002

OSPF Area Support

OSPF Area Support

OSPF Areas and Area Terms
   Area Border Router (ABR)
   Autonomous System Boundary Router (ASBR)

OSPF Area Types
   Normal Area
   Stub Area
   Totally Stubby Area (TSA)
   Not-So-Stubby Area (NSSA)

OSPF Areas and Area Terms

Every OSPF network can be subdivided into regions called **Areas**. A router maintains a topology database for each area it participates in. There are several reasons to divide an OSPF network into areas.

- ❍ **Conserving router resources** – an OSPF router does NOT maintain a topology database of areas that is does not participate in. Therefore by subdividing the network into area hierarchy, router resources are conserved.

- ❍ **Hiding information** – Route summarization can be implemented at the area borders. This allows for reduction in the amount of detailed routing information that is advertised outside of an area. If a specific prefix that belongs to a summary changes, that information will not be distributed outside of the area where the prefix exists, which promotes stability.

- ❍ **Handling of External Routes** – different area types handle external routes differently.

Stub, NSSA, and Totally Stubby areas help reducing the amount of external routes in the OSPF network.

Each area is uniquely identified by its **Area ID**. The Area ID is a 32-bit number that uses the dotted decimal notation of an IP address (it is not an IP address however). All routers within a given area must agree on the Area ID identifying the area. Area ID **MUST** be configured in each router manually – there is no default. The following command is used to create an area:

ospf create area <AREA-ID>

"AREA-ID" can be replaced with the area ID in the format x.x.x.x, or by the word "backbone", which is a replacement for area ID 0.0.0.0.

INTEROPERABILITY NOTE:

> It is acceptable (but bad form) to define the AREA-ID with a single number, such as "ospf create area 1". The system will expand this shortcut entry with **trailing** zeros, as if the operator had actually entered "ospf create area 1.0.0.0". Note that this is a non-standardized behavior, and while some vendors may choose to expand such a command in the same way, others might expand it with **leading** zeros, resulting in "ospf create area 0.0.0.1". Since all four numbers in the AREA-ID" are significant, 1.0.0.0 and 0.0.0.1 refer to completely different areas.
> ***Riverstone Networks recommends that all AREA-IDs be explicitly configured in full form.***

An OSPF network may consist of a single area or multiple areas. If a single area is present, it should be AREA-ID 0.0.0.0 (backbone). If multiple areas are configured, they should all be attached to the backbone. In most cases that attachment is physical (via an Area Border Router (ABR) that connects Area 0.0.0.0 to the non-zero area. The attachment can also be logical via a Virtual Link.

The Backbone (Area 0.0.0.0) must attach to all other areas. Therefore it must be contiguous.

Area Border Router (ABR)

The ABR is a router that connects the backbone to a non-zero area. Each non-zero area therefore has one or more ABRs. To configure an ABR, create more than one area on it.

Example:

ospf create area backbone
ospf create area 0.0.0.1
ospf create area 0.0.0.17

The above configuration creates an ABR between three areas 0 backbone, 0.0.0.1, and 0.0.0.17. Sometimes, the AREA-IDs are colloquially called backbone, area 1, and area 17, respectively, please see the above Interoperability Note.

You may want to also read about the concept of Alternative ABR relaxes the requirement for an ABR to be connected to the Backbone, and allows the ABR to connect two non-zero areas.

Autonomous System Boundary Router (ASBR)

A router that redistributes prefixes from other sources into OSPF is an ASBR. Routes from non-OSPF sources are referred to as External routes. An ASBR can be configured in Normal and NSSA areas, bit not in Stub and Totally-Stubby Areas. To define an ASBR, configure route redistribution (or export) on the router, as illustrated on the example below. It defines the router as an ASBR which redistributes 1.1.2.0/24 as an External route in OSPF.
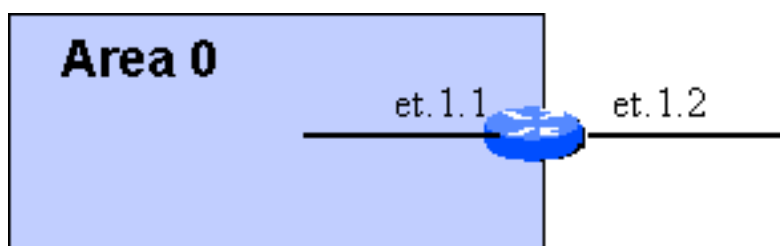
interface create ip ONE address-netmask 1.1.1.1/24 port et.1.1
interface create ip TWO address-netmask 1.1.2.1/24 port et.1.2
ospf create area backbone
ospf add interface ONE to-area backbone
ospf start
ip-router policy redistribute from-proto direct to-proto ospf

OSPF Area Types

OSPF handles External Routes by flooding them throughout the OSPF network. That flooding occurs even across area borders, therefore implementing hierarchy with OSPF areas provides little benefit when large number of external routes are redistributed into OSPF. In order to handle these

situations, various area types were created, which differ in the way they handle external routes, and in their ability to support an ASBR.

There are four area types supported by the RS. These are Normal , Stub, Totally Stubby, and Not-So-Stubby.

Normal Area

When an OSPF area is created by the "ospf create area <AREA-ID>" command, it is created as a normal area. Normal areas handle externals by flooding them bi-directionally across the Area Border Router (ABR). If an OSPF network is comprised of Normal areas only, then all External routes will be flooded throughout the whole network. This behavior is illustrated on Figure 1, where both Area 0 and Area 1 are configured as normal areas and have ASBRs that redistribute externals into their respective areas. Externals from Area 0 will be flooded throughout Area 1 and vice versa.
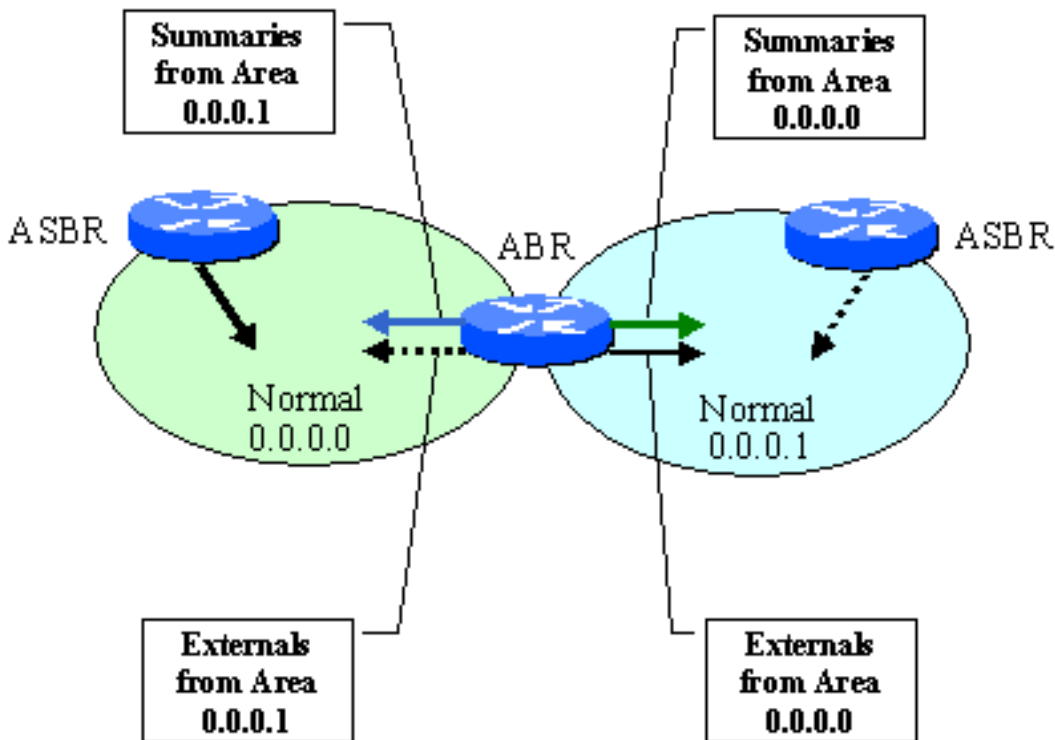
Figure 1: Handling Routes in Normal Areas

on complete configuration examples for routers in a Normal area.

Stub Area

When many Externals are redistributed into OSPF, the LSDB of all routers will be burdened by processing them. To limit the effects of this, Stub Areas can be configured. Stub areas handle externals in the following manner:

- ❍ The ABR will not flood Externals coming from Area 0 into the Stub Area

- ❍ Instead, the ABR will generate a default route into the stub area (as a Type 3 Summary LSA).

- ❍ ASBRs are NOT permitted in the STUB area (no router can redistribute externals into OSPF within the Stub area).

These rules are illustrated on Figure 2 below.
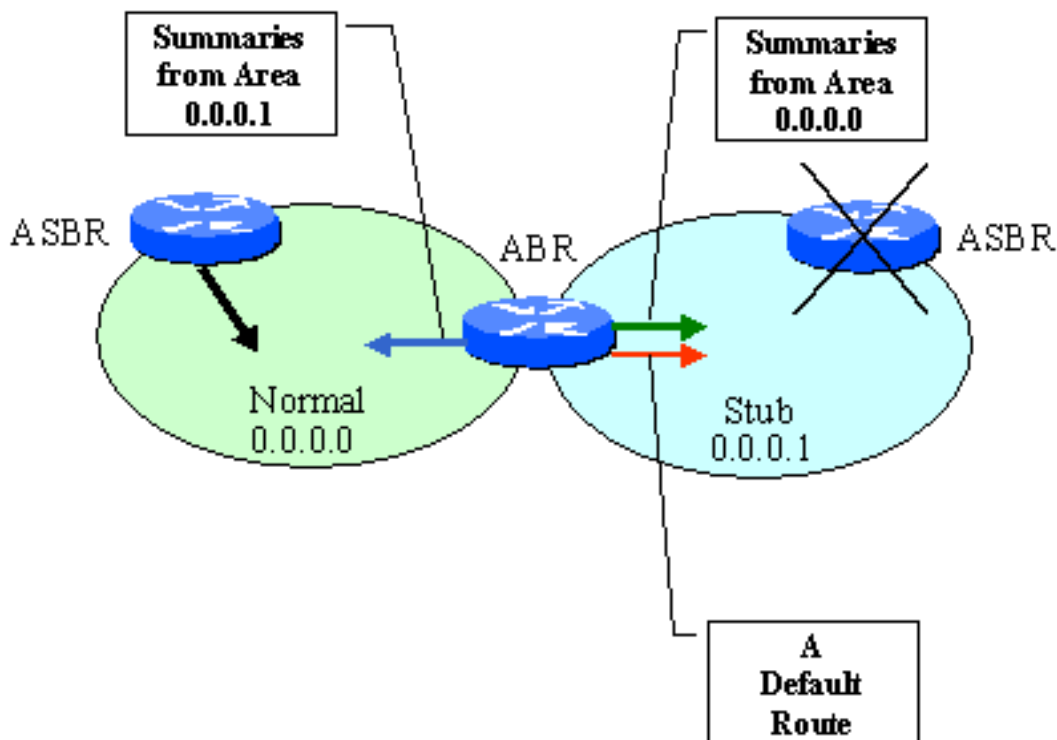
Figure 2: Handling Routes in a Stub Area

To configure an ABR to a Stub area, issue the following command:

ospf set area 0.0.0.1 stub stub-cost <number>

The stub-cost parameter determines what the OSPF cost will be for the default route that will be injected into the Stub area. The value can be from 0 to 65535.

To configure any other (intra-area) router within a Stub area, issue the following command:

ospf set area 0.0.0.1 stub

See more on complete configuration examples for routers in a Stub area.

Totally Stubby Area (TSA)

TSA differ from Stub areas in that there are not even summary routes injected into the TSA. The only route that is injected by the ABR is the default route. All intra-area routes follow the default for all destinations – both internal and external to the OSPF domain.
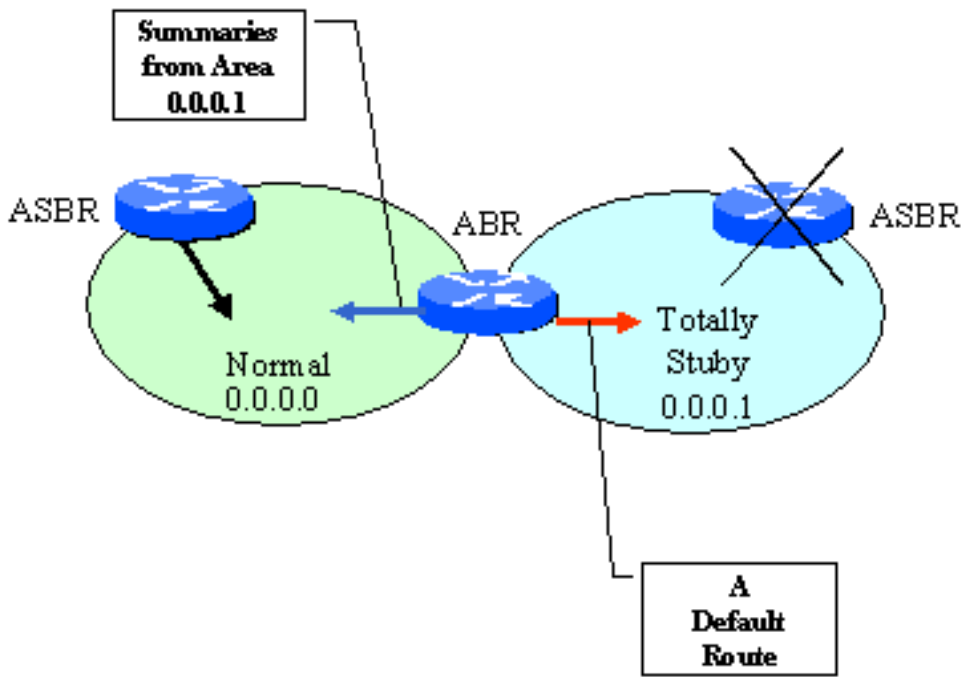
Figure 3: Handling routes in a Totaly Stubby Area (TSA)

To configure an ABR to a Totally Stubby area, issue the following command:

ospf set area 0.0.0.1 stub no-summary stub-cost <number>

The stub-cost parameter determines what the OSPF cost will be for the default route that will be injected into the TSA. The value can be from 0 to 65535.

To configure any other (intra-area) router within a TSA, issue the following command:

ospf set area 0.0.0.1 stub no-summary

See more on complete configuration examples for routers in a Totally Stubby Area.

Not-So-Stubby Area (NSSA)

To add some flexibility in the handling of Externals, and still preserve the flood-prevention capability of Stub areas, NSSA can be configured. NSSA is defined in RFC 1587, and allows for

a Externals to be injected directly into the NSSA.   However the NSSA still will not accept Externals originated in any other area.   This is accomplished by introducing a new LSA type, Type 7 LSA. An ASBR in the NSSA will inject Externals using Type 7 LSA. When crossing the ABR, those Type 7 LSAs will be translated into regular Type 5 LSAs, and will be handled like regular Externals throughout the rest of the network.

As a result, the NSSA will maintain only Externals that originate locally (in the NSSA) in addition to the Summary routes that come from the ABR. This behavior can be seen on Figure 4.
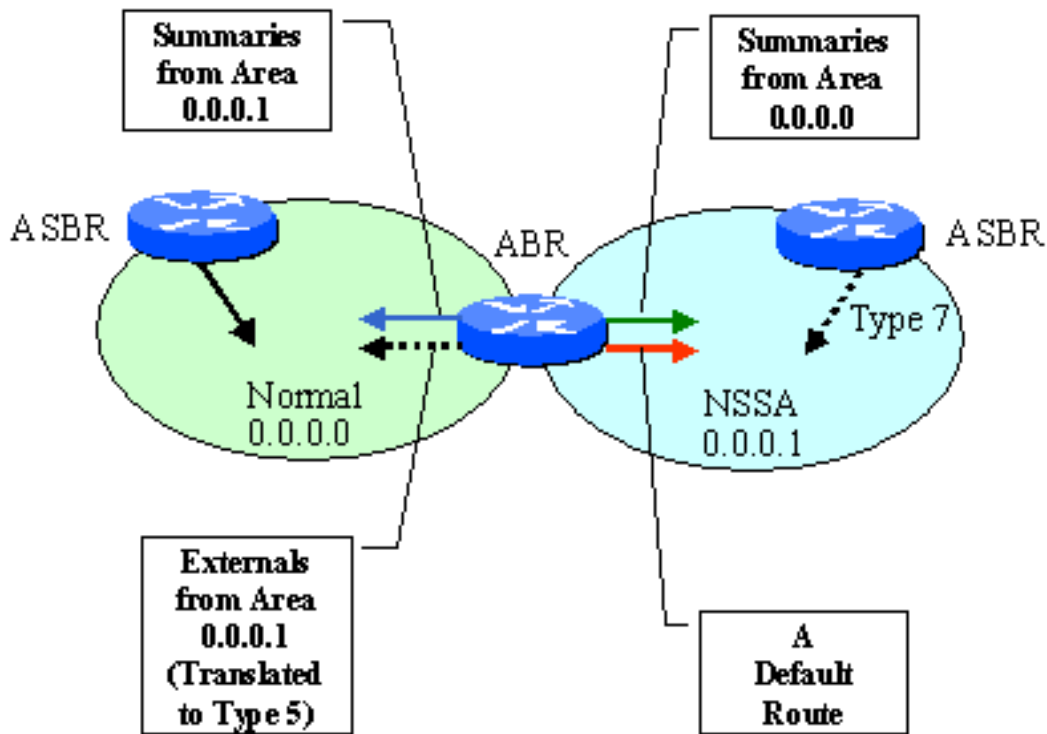
Figure 4: Route Handling in a NSSA.

To configure an ABR to a NSSA, issue the following command:

ospf set area 0.0.0.1 nssa nssa-cost <number>

The nssa-cost parameter determines what the OSPF cost will be for the default route that will be injected into the NSSA. The value can be from 0 to 65535.

To configure any other (intra-area) router within a NSSA, issue the following command:

ospf set area 0.0.0.1 nssa

[See more](#) on complete configuration examples for routers in a NSSA.

Go back to [OSPF Support Home](#)
Last Updated June 3, 2002

Not-So-Stubby Area (NSSA)

Not-So-Stubby Area (NSSA)

Introduction

OSPF Normal areas flood all external routes across area borders. In the presence of large number of external routes, this may be a problem, as external routes cannot be summarized at the ABRs. Stub areas, and Totally Stubby Areas (TSA), are designed to alleviate the problem by preventing external routes from being injected into the Stub/TSA area, and instead a default routes is injected. Stub/TSA areas are incapable of carrying external routes (Type 5 LSAs), and hence are incapable of supporting ASBRs.

Not-so-Stubby Areas allow for supporting ASBRs within the NSSA, while maintaining the same behavior as Stub/TSA of NOT injecting external (Type 5) routes coming from the backbone. Thus NSSA routers benefit from the significant reduction of external routes coming from the backbone, while having the capability to carry a limited number of externals that originate in the NSSA.

Many service providers opt for carrying their customer routes via BGP, rather than a dynamic protocol, such as OSPF. They accomplish this by redistributing into IBGP customer routes at the device they originate. This, although preferred, is not always possible, as many customer-facing devices (such as RAS servers) do not support BGP. Redistributing such routes as OSPF externals may be the only option. NSSA is particularly well suited for such situations.

OSPF NSSA is defined in [RFC 1587, "The OSPF NSSA Option".](#) and further enhanced in [http://www.ietf.org/internet-drafts/draft-ietf-ospf-nssa-update-10.txt](#) ?

How Does it Work

Type 7 LSA
To provide the ability of carrying external routes originated in the NSSA, a new LSA type was defined, Type 7 LSA. It has the structure and semantics of a Type 5 (External) LSA, with a two differences:
   ❍ Type 7 LSAs can be originated and propagated within the NSSA, they do not cross area borders like Type 5 LSAs do.
   ❍ Type 5 LSAs are not supported in NSSA; they can be neither originated nor propagated in NSSA.

```
Routing Bit Set on this LSA
LS Age: 806
Options: E NP
LS Type: NSSA
Link State ID: 50.50.1
Advertising Router: 3.3.3.3
LS Seq Number: 80000001
Checksum: f84a
```

```
Length: 36
Network Mask: /24
    Metric Type: 2
    TOS: 0
    Metric: 1
    Forward Address: 10.0.0.9
    External Route Tag: 0
```

- Bit P - This bit is used to tell the NSSA ABR whether to translate type 7 into type 5.
- No Type 7/5 translation means bit P = 0.
- Type 7/5 translation means bit P = 1.
- If bit P = 0, then the NSSA ABR must not translate this LSA into Type 5. This happens when NSSA ASBR is also an NSSA ABR.
- If bit P = 1, then the NSSA ABR (if there are multiple NSSA ABRs, the one with lowest router ID) must translate this type 7 LSA into a type 5 LSA.

Translation of Type 7 to Type 5 LSA at the ABR

An ASBR in a NSSA injects its external routes as Type 7 LSAs (The NSSA does not support Type 5 Externals). The ABR can (optionally) translate those to regular Type 5 LSAs and inject them into the backbone as Type 5 Externals. At this point the these LSAs are indistinguishable from other Type 5 LSAs that may have come from other source. As shown on the figure below, these Type 5 LSAs, along with those from other sources, are flooded on all other non-stub areas.

Figure 1: Translation of Type 7 LSAs into Type 5
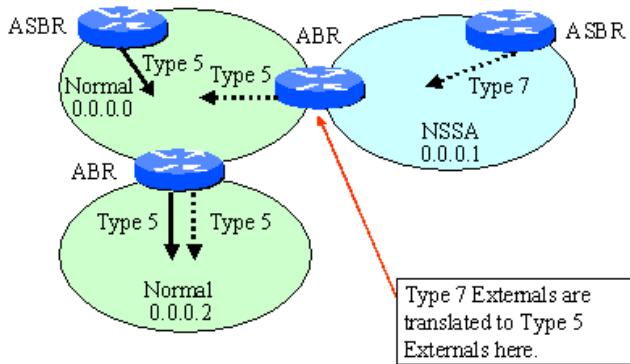
When the NSSA is connected to the backbone via more than one ABR, the ABR with the highest Router ID is selected to perform the Type 7 à Type 5 translation. If that ABR fails, the ABR with next-highest Router ID is selected.

Example NSSA Configuration

The following figure illustrates the configuration of an NSSA with two ABRs.

Figure 2: Example configuration of NSSA with redundant ABRs

RS1 Configuration	RS2 Configuration
| ```
interface create ip TO-RS2 address-netmask 10.0.0.1/30 port et.2.1
interface create ip TO-RS4 address-netmask 10.0.0.13/30 port et.2.2
interface create ip LAN address-netmask 63.228.71.1/24 port et.2.8
interface add ip lo0 address-netmask 1.1.1.1/32
ip-router global set router-id 1.1.1.1
ospf create area backbone
ospf add interface TO-RS2 to-area backbone
ospf add interface TO-RS4 to-area backbone
ospf add stub-host 1.1.1.1 to-area backbone cost 33
ospf start
system set name RS1
``` | ```
interface create ip TO-RS1 address-netmask 10.0.0.2/30 port et.2.1
interface create ip TO-RS3 address-netmask 10.0.0.5/30 port et.2.2
interface add ip lo0 address-netmask 2.2.2.2/32
ip-router global set router-id 2.2.2.2
ospf create area backbone
ospf create area 0.0.0.1
ospf set area 0.0.0.1 nssa nssa-cost 1
ospf add interface TO-RS1 to-area backbone
ospf add interface TO-RS3 to-area 0.0.0.1
ospf add stub-host 2.2.2.2 to-area backbone cost 33
ospf start
system set name RS2
``` |
| RS3 Configuration | RS4 Configuration |
| ```
interface create ip TO-RS2 address-netmask 10.0.0.6/30 port et.3.2
interface create ip TO-RS4 address-netmask 10.0.0.9/30 port et.3.1
interface create ip NSSA-1 address-netmask 50.50.0.1/24 port et.3.5
interface create ip NSSA-2 address-netmask 50.50.1.1/24 port et.3.6
interface create ip NSSA-3 address-netmask 50.50.2.1/24 port et.3.7
interface create ip NSSA-4 address-netmask 50.50.3.1/24 port et.3.8
interface add ip lo0 address-netmask 3.3.3.3/32
ip-router global set router-id 3.3.3.3
ospf create area 0.0.0.1 nssa
ospf add interface TO-RS2 to-area 0.0.0.1
ospf add interface TO-RS4 to-area 0.0.0.1
ospf add stub-host 3.3.3.3 to-area 0.0.0.1 cost 33
ospf start
ip-router policy redistribute from-proto direct to-proto ospf-nssa
system set name RS3
``` | ```
interface create ip TO-RS1 address-netmask 10.0.0.14/30 port et.2.2
interface create ip TO-RS3 address-netmask 10.0.0.10/30 port et.2.1
interface add ip lo0 address-netmask 4.4.4.4/32
ip-router global set router-id 4.4.4.4
ospf create area backbone
ospf create area 0.0.0.1 nssa
ospf set area 0.0.0.1 nssa nssa-cost 1
ospf add interface TO-RS1 to-area backbone
ospf add interface TO-RS3 to-area 0.0.0.1
ospf add stub-host 4.4.4.4 to-area backbone cost 33
ospf start
system set name RS4
``` |
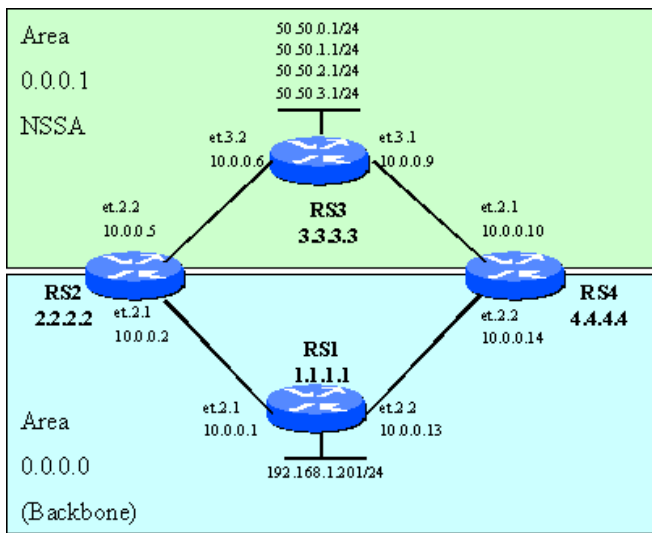
RS2 and RS4 are the ABRs, and RS3 is an ASBR in the NSSA, which redistributes direct routes (50.50.0.0/24, 50.50.1.0/24, 50.50.2.0/24, and 50.50.3.0/24) as an external Type 7 LSA into the NSSA. That Type 7 external LSA is translated by router RS4 (as its Router ID is higher than that of RS2), and injected in the backbone area as an equivalent Type 5 LSA.

Configuration Notes:

1. The redistribution statement in RS3's configuration causes the Type 7 LSA to be injected with a Type 2 External metric (see Externals for more information). This is the default behavior, which can be overridden with the "nssa-type 1" clause of the command.

2. The command "ospf set area 0.0.0.1 nssa nssa-cost 1" in RS2 and RS4's configuration causes the default route injected into the NSSA by each of those ABRs to have a cost of 1. Note that this is a Type 7 External default route.

3. There is no configuration necessary on the ABR(s) to translate Type 7 to Type 5 LSAs.

4. Because a single router is selected to perform Type 7 à Type 5 translation, there is no load sharing for destinations that originated as Type 7 LSAs and were translated into Type 5 LSAs. In the example of Figure 2, this translation is performed at RS4 (higher router ID), and therefore all Type 7 destinations will be reachable via RS4.

```
RS1# ip show routes

Destination          Gateway              Owner      Netif
-----------          -------              -----      -----
```

```
1.1.1.1                 1.1.1.1                 -         lo0
2.2.2.2                 10.0.0.2                OSPF      TO-RS2
3.3.3.3                 10.0.0.2                OSPF_IA   TO-RS2
                        10.0.0.14               OSPF_IA   TO-RS4
4.4.4.4                 10.0.0.14               OSPF      TO-RS4
10.0.0.0/30             directly connected      -         TO-RS2
10.0.0.4/30             10.0.0.2                OSPF_IA   TO-RS2
10.0.0.8/30             10.0.0.14               OSPF_IA   TO-RS4
10.0.0.12/30            directly connected      -         TO-RS4
50.50.0.0/24            10.0.0.14               OSPF_ASE  TO-RS4
50.50.1.0/24            10.0.0.14               OSPF_ASE  TO-RS4
50.50.2.0/24            10.0.0.14               OSPF_ASE  TO-RS4
50.50.3.0/24            10.0.0.14               OSPF_ASE  TO-RS4
127.0.0.1               127.0.0.1               -         lo0
192.168.1.0/24          directly connected      -         LAN
RS1#
```

Observe that all highlighted destinations (3.3.3.3/32, and 50.50.x.0/24) originate at router RS3, but there is load sharing occurring only for 3.3.3.3/32.   This is because 50.50.x.0/24 are injected as NSSA Type 7 LSAs, and 3.3.3.3/32 isn't.

Route Manipulation at the NSSA ABR

Summarization of multiple Type 7 into a single Type 5 LSA

At the ABR of an NSSA, Type 7 Externals can be summarized into a single Type 5 External. Based on the configuration of Figure 2, the following line will need to be added to the ABR(s) configuration:

```
ospf add nssa-network 50.50.0.0/22 to-area 0.0.0.1
```

```
RS1# ip show routes

Destination             Gateway                 Owner     Netif
-----------             -------                 -----     -----
1.1.1.1                 1.1.1.1                 -         lo0
2.2.2.2                 10.0.0.2                OSPF      TO-RS2
3.3.3.3                 10.0.0.2                OSPF_IA   TO-RS2
                        10.0.0.14               OSPF_IA   TO-RS4
4.4.4.4                 10.0.0.14               OSPF      TO-RS4
10.0.0.0/30             directly connected      -         TO-RS2
10.0.0.4/30             10.0.0.2                OSPF_IA   TO-RS2
10.0.0.8/30             10.0.0.14               OSPF_IA   TO-RS4
10.0.0.12/30            directly connected      -         TO-RS4
50.50.0.0/22            10.0.0.14               OSPF_ASE  TO-RS4
127.0.0.1               127.0.0.1               -         lo0
192.168.1.0/24          directly connected      -         LAN
```

The four entries for 50.50.x.0/24 are now replaced with a single 50.50.0.0/22. This summary will continue to be injected in the backbone as long as one of the four contributing Type 7 LSAs is received at the advertising ABR.

If there are Type 7 External routes that do not fall within a summarization range, those are translated into Type 5 Externals and propagated into the backbone.

Route Filtering from the NSSA into the Backbone (Type 7 Externals)

Type 7 Externals can also be filtered at the ABR, so they can be selectively translated into Type 5. Filtering is accomplished with the command:

```
ospf add nssa-network 50.50.0.0/23 to-area 0.0.0.1 restrict
```

When applied on the ABRs of Figure 2, it will result in filtering 50.50.0.0/24 and 50.50.1.0/24, and translating into Type 5 50.50.2.0/24 and 50.50.3.0/24, as evident from the forwarding table of RS1 below:

```
RS1# ip show routes

Destination             Gateway                 Owner     Netif
-----------             -------                 -----     -----
```

```
1.1.1.1               1.1.1.1             -        lo0
2.2.2.2               10.0.0.2            OSPF     TO-RS2
3.3.3.3               10.0.0.2            OSPF_IA  TO-RS2
                      10.0.0.14           OSPF_IA  TO-RS4
4.4.4.4               10.0.0.14           OSPF     TO-RS4
10.0.0.0/30           directly connected  -       TO-RS2
10.0.0.4/30           10.0.0.2            OSPF_IA  TO-RS2
10.0.0.8/30           10.0.0.14           OSPF_IA  TO-RS4
10.0.0.12/30          directly connected  -       TO-RS4
50.50.2.0/24          10.0.0.2            OSPF_ASE TO-RS2
50.50.3.0/24          10.0.0.2            OSPF_ASE TO-RS2
127.0.0.1             127.0.0.1           -        lo0
192.168.1.0/24        directly connected  -       LAN
RS1#
```

Route Filtering from the NSSA into the Backbone (Inter-area Routes)

This filtering is configured in the same manner as filtering for Normal areas. See Route Filtering on Area Borders for details.

Route Filtering from the Backbone into the NSSA (Inter-area Routes)

This filtering is configured in the same manner as filtering for Normal areas. See Route Filtering on Area Borders for details.

Special Translation Behavior When the NSSA ABR is Also an ASBR

It is possible that the ABR of the NSSA is also an ASBR. There are two possibilities to inject the externals:

1. Inject them as Type 5 LSAs
`ip-router policy redistribute from-proto direct to-proto ospf`
2. Inject them as Type 7 LSAs into the NSSA
`ip-router policy redistribute from-proto direct to-proto ospf-nssa`
3. Inject them as both Type 5 and Type 7 LSAs in the respective areas.
`ip-router policy redistribute from-proto direct to-proto ospf`
`ip-router policy redistribute from-proto direct to-proto ospf-nssa`

In the first case, only a Type 5 LSA will be injected in the LSDB.

In the second case, a Type 7 LSA will be injected in the LSDB **and it will NOT be translated to Type 5**.

In the third case, both a Type 7 and a Type 5 LSAs will be injected.   This is probably what most people want, which emulates the case where the ASBR is within the NSSA (and is different than the ABR).

Default Route Handling in NSSA

There are two possible sources of default route in NSSA:

1. A default injected by the ABR with the `"ospf set area 0.0.0.1 nssa [nssa-cost x]  [nssa-type 1|2]"`
2. A default injected by an ASBR within the NSSA with the `"ip-router policy redistribute from-proto <NAME> to-proto ospf-nssa …"`

From both of these sources, the default route is injected as a Type 7 External. Injection from the ABR can be suspended simply by NOT specifying the "nssa-cost x" clause in the command. In such case, the only source of default in the NSSA would be an ASBR.

Both of these sources may have multiple routers injecting default routes, e.g. several ABRs and several ASBRs. The usual rules of OSPF External routes selection apply; Type 1 metric is preferred over Type 2, and within the same metric type, the metric value is compared.

Injecting a Conditional Default Route in NSSA

In some cases, all connectivity out of the NSSA may be dependent upon a default route, injected by an ASBR located in the Backbone. A typical such case is illustrated on Figure 3. Assume that the ASBR injecting the default (R1) is a router that does BGP peering with the outside world. If that router looses its outside

connection and stops injecting the default, there is no point in carrying traffic outside of the NSSA area, only to drop it at R1.

In this case it is useful to stop injecting the default in NSSA as soon as the default in the backbone area disappears.

Figure 3: Conditional Default in NSSA

Conditional default injection is configured on the ABR:

```
ospf set area 1 nssa conditionally nssa-cost 4
```

The result is that a default route will be injected in the NSSA byt the ABR if and only if another default (from another source) is present in the ABR forwarding table.

Conditional default route injection is ONLY supported in NSSA.

Go back to OSPF Support Home

Last Updated June 3, 2002

Configuring Authentication

Configuring Authentication

Introduction

By default, OSPF packets between routers are not authenticated and could be forged. Two types of authentication can be configured – simple and MD5. With simple authentication, a password (key) is configured on each router and is included in plaintext in each OSPF packet originated by that router. Simple authentication is therefore not secure. MD5 authentication is based on shared secret keys that are configured in all routers in the area. Each router computes an MD5 hash for each packet based on the content of the packet and the configured secret key. Then it includes the resulting hash value in the OSPF packet. The receiving router, using the pre-configured secret key, will compute an MD5 hash of the packet and compare it with the hash value that the packet carries thus verifying its authenticity. Sequence numbers are also employed with MD5 authentication to protect against replay attacks.

The RS is compliant with [RFC 2328, Appendix D](#), which requires authentication parameters to be configurable on a per-interface basis. However, in most cases, OSPF authentication is configured on a per-area basis. For demonstration purposes, we will use per-area configuration in our examples below.

Configuring Simple (Plain Text) Authentication

There are two steps to configure simple authentication:

1. Configure a key (on a keychain).
   ```
   ip-router authentication create key-chain <key-chain-name>
   id <key ID> key <Key-String>
   ```

2. Configure OSPF to use the key.
   ```
   ospf set area|interface <name> authentication-method
   simple key-chain <key-chain-name>
   ```

Note: Only one key per keychain can be configured for simple authentication

RS1 Configuration	RS2 Configuration
version 7.0 interface create ip TO-RS2 address-netmask 10.0.0.1/24 port et.2.1 interface add ip lo0 address-netmask 1.1.1.1/32 ip-router global set router-id 1.1.1.1 ip-router authentication create key-chain 1 id 1 key <encrypted>ee74b8e975f037c6 ospf create area backbone ospf add stub-host 1.1.1.1 to-area backbone cost 1 ospf add interface TO-RS2 to-area backbone	version 7.0 interface create ip TO-RS1 address-netmask 10.0.0.2/24 port et.2.1 interface add ip lo0 address-netmask 2.2.2.2/32 ip-router global set router-id 2.2.2.2 ip-router authentication create key-chain 22 id 1 key <encrypted>ee74b8e975f037c6 ospf create area backbone ospf add interface TO-RS1 to-area backbone ospf add stub-host 2.2.2.2 to-area backbone cost 1

```
ospf set area backbone authentication-      ospf set area backbone authentication-
method simple key-chain 1                    method simple key-chain 22
ospf start                                   ospf start
system set name RS1                          system set name RS2
```

Configuring MD5 Authentication

There are two steps to configure MD5 authentication:

3.   Configure a key (on a keychain) – same as in simple authentication.

```
ip-router authentication create key-chain <key-chain-name>
id <key ID> key <Key-String>
```

4.   Configure OSPF to use the keychain.

```
ospf set area|interface <name> authentication-method md5
key-chain <key-chain-name>
```

Note: Multiple keys per keychain can be configured for MD5 authentication (useful when gradual change of the authentication key is needed). However in the current software release (7.0.0.0) only one key is used by OSPF.

RS1 Configuration	RS2 Configuration
| ```
version 7.0
interface create ip TO-RS2 address-
netmask 10.0.0.1/24 port et.2.1
interface add ip lo0 address-netmask
1.1.1.1/32
ip-router global set router-id 1.1.1.1
ip-router authentication create key-
chain 1 id 1 key
<encrypted>ee74b8e975f037c6
ospf create area backbone
ospf add stub-host 1.1.1.1 to-area
backbone cost 1
ospf add interface TO-RS2 to-area
backbone
ospf set area backbone authentication-
``` | ```
version 7.0
interface create ip TO-RS1 address-
netmask 10.0.0.2/24 port et.2.1
interface add ip lo0 address-netmask
2.2.2.2/32
ip-router global set router-id 2.2.2.2
ip-router authentication create key-
chain 22 id 1 key
<encrypted>ee74b8e975f037c6
ospf create area backbone
ospf add interface TO-RS1 to-area
backbone
ospf add stub-host 2.2.2.2 to-area
backbone cost 1
ospf set area backbone authentication-
``` |

```
method md5 key-chain 1              method md5 key-chain 22
ospf start                          ospf start
system set name RS1                 system set name RS2
ospf trace hello receive detail
```

Copying and Pasting Encrypted Passwords

The router stores password in encrypted form, to prevent anyone reading the config from realizing the clear text of the password. The encryption algorithm used for this has two useful properties:

§   The same clear-text passwords result in the same encrypted string
§   Copy-and-paste the line containing the encrypted password *in its entirety* results in the same password (as if typed in clear text).

Please note that the correct method to copy-and-paste encrypted passwords is by pasting the entire line. Example:

Correct: `ip-router authentication create key-chain 1 id 1 key <encrypted>561bb244554a58f1`

Incorrect: `ip-router authentication create key-chain 1 id 1 key 561bb244554a58f1`

(this will result in double-encrypting the string "`561bb244554a58f1`" yielding a completely different password`.`

Monitoring and Troubleshooting

OSPF Hello tracing can be utilized to troubleshoot any issues with OSPF authentication:

`ospf trace hello detail`
`ip-router global set trace-state on`

Here is the tracing output when all is well. Areas of interest in this output are highlighted.

`RS1# -04-24 12:54:46 OSPF SEND: 10.0.0.1 -> 224.0.0.5: Version 2, Type Hello (1), Length 48`

```
-04-24 12:54:46          Router ID 1.1.1.1, Area 0.0.0.0,
Authentication <MD5> (2)
-04-24 12:54:46             Zero: 0, Key ID: 1, Length: 16, Sequence:
0
-04-24 12:54:46             Mask 255.255.255, Options <E> (2),
Priority 1, Neighbours 1
-04-24 12:54:46             Intervals: Hello 10s, Dead Router 40s,
Designated Router 10.0.0.2, Backup 10.0.0.1
-04-24 12:54:46             Neighbours: 2.2.2.2
-04-24 12:54:50 OSPF RECV: 10.0.0.2 -> 224.0.0.5: Version 2, Type
Hello (1), Length 48
-04-24 12:54:50          Router ID 2.2.2.2, Area 0.0.0.0,
Authentication <MD5> (2)
-04-24 12:54:50             Zero: 0, Key ID: 1, Length: 16, Sequence:
0
-04-24 12:54:50             Mask 255.255.255, Options <E> (2),
Priority 1, Neighbours 1
-04-24 12:54:50             Intervals: Hello 10s, Dead Router 40s,
Designated Router 10.0.0.2, Backup 10.0.0.1
-04-24 12:54:50             Neighbours: 1.1.1.1
```

The following is the tracing output when the authentication keys between the two routers are mismatched.

```
-04-24 12:59:26 OSPF SEND: 10.0.0.1 -> 224.0.0.5: Version 2, Type
Hello (1), Length 44
-04-24 12:59:26          Router ID 1.1.1.1, Area 0.0.0.0,
Authentication <MD5> (2)
-04-24 12:59:26             Zero: 0, Key ID: 1, Length: 16, Sequence:
0
-04-24 12:59:26             Mask 255.255.255, Options <E> (2),
Priority 1, Neighbours 0
-04-24 12:59:26             Intervals: Hello 10s, Dead Router 40s,
Designated Router 10.0.0.1, Backup 0.0.0.0
-04-24 12:59:30 OSPF RECV: 10.0.0.2 -> 224.0.0.5: Version 2, Type
Hello (1), Length 44
-04-24 12:59:30          Router ID 2.2.2.2, Area 0.0.0.0,
Authentication <MD5> (2)
-04-24 12:59:30             Zero: 0, Key ID: 1, Length: 16, Sequence:
0
```

```
-04-24 12:59:30 OSPF: invalid packet: Invalid authentication key
(24)
```

Go back to [OSPF Support Home](#)

Last updated June 4, 2002

Configuring Route Summarization

The most significant benefit of introducing hierarchy (areas) in an OSPF network is the ability to perform route summarization at the area borders (on the ABR). Summarization is the ability to replace a number of OSPF route advertisements with a single one, thus reducing memory and computational requirements for the routers that receive the summary.

For summarization to occur, two things are needed:

    o   Contiguous IP addressing within areas, such that a number of adjacent longer prefixes can be expressed with a single shorter prefix.

    o   A configuration at the ABR to cause summarization. The configuration is done with the command "ospf add summary range x.x.x.x/y to-area <AREA_ID>".

We will illustrate this with the example below:

The Backbone area consists of 4 subnets, 10.0.0.0/24, 10.0.1.0/24, 10.0.2.0/24, and 10.0.3.0/24. Area 0.0.0.11 consists of 4 other subnets, also contiguously assigned, 11.0.0.0/24, 11.0.1.0/24, 11.0.2.0/24, and 11.0.3.0/24. For both areas, the first condition above is satisfied; the prefixes are

contiguously assigned, and can be expressed with summary ranges. For the Backbone area, this summary range is 10.0.0.0/22, and for area 0.0.0.11 the range is 11.0.0.0/22.

In order for the summarization to occur, we also need to configure it on the ABR. We will illustrate this by configuring summarization ONLY for area 0.0.0.11, and not for the Backbone. We will observe that the result of this is that routes from area 0.0.0.11 are summarized when injected into the backbone, and routes from the backbone are NOT summarized when injected into area 0.0.0.11.

RS1 Configuration and Forwarding Table	RS1# ip show routes
| `interface create ip TO-RS2 address-netmask`<br>`10.0.0.1/24 port et.2.1`<br>`interface create ip LAN1 address-netmask`<br>`10.0.1.1/24 port et.2.6`<br>`interface create ip LAN2 address-netmask`<br>`10.0.2.1/24 port et.2.7`<br>`interface create ip LAN3 address-netmask`<br>`10.0.3.1/24 port et.2.8`<br>`interface add ip lo0 address-netmask`<br>`1.1.1.1/32`<br>`ip-router global set router-id 1.1.1.1`<br>`ospf create area backbone`<br>`ospf add interface all to-area backbone`<br>`ospf add stub-host 1.1.1.1 to-area`<br>`backbone cost 1`<br>`ospf start`<br>`system set name RS1` | ```
Destination
Gateway                  Owner     Netif
-----------              -------
                 -----    -----
1.1.1.1
1.1.1.1                   -         lo0
2.2.2.2
10.0.0.2                  OSPF      TO-RS2
3.3.3.3
10.0.0.2                  OSPF_IA   TO-RS2
10.0.0.0/24              directly
connected    -           TO-RS2
10.0.1.0/24              directly
connected    -           LAN1
10.0.2.0/24              directly
connected    -           LAN2
10.0.3.0/24              directly
connected    -           LAN3
11.0.0.0/22              10.0.0.2
     OSPF_IA    TO-RS2
127.0.0.1
127.0.0.1                -          lo0
``` |

RS2 Configuration and Forwarding Table

```
interface create ip TO-RS1 address-netmask 10.0.0.2/24 port et.2.1
interface create ip TO-RS3 address-netmask 11.0.0.2/24 port et.2.2
interface add ip lo0 address-netmask 2.2.2.2/32
ip-router global set router-id 2.2.2.2
ospf create area backbone
ospf create area 0.0.0.11
ospf add interface TO-RS1 to-area backbone
ospf add interface TO-RS3 to-area 0.0.0.11
ospf add stub-host 2.2.2.2 to-area backbone cost 1
ospf add summary-range 11.0.0.0/22 to-area 0.0.0.11
ospf start
system set name RS2
```

RS2# ip show routes

Destination	Gateway	Owner	Netif
1.1.1.1	10.0.0.1	OSPF	TO-RS1
2.2.2.2	2.2.2.2	-	lo0
3.3.3.3	11.0.0.1	OSPF	TO-RS3
10.0.0.0/24	directly connected	-	TO-RS1
10.0.1.0/24	10.0.0.1	OSPF	TO-RS1
10.0.2.0/24	10.0.0.1	OSPF	TO-RS1
10.0.3.0/24	10.0.0.1	OSPF	TO-RS1
11.0.0.0/24	directly connected	-	TO-RS3
11.0.1.0/24	11.0.0.1	OSPF	TO-RS3
11.0.2.0/24	11.0.0.1	OSPF	TO-RS3
11.0.3.0/24	11.0.0.1	OSPF	TO-RS3
127.0.0.1	127.0.0.1	-	lo0

RS3 Configuration and Forwarding Table

```
interface create ip TO-RS2 address-netmask 11.0.0.1/24 port et.2.2
interface create ip LAN1 address-netmask 11.0.1.1/24 port et.2.6
interface create ip LAN2 address-netmask 11.0.2.1/24 port et.2.7
interface create ip LAN3 address-netmask 11.0.3.1/24 port et.2.8
interface add ip lo0 address-netmask 3.3.3.3/32
ip-router global set router-id 3.3.3.3
ospf create area 0.0.0.11
ospf add interface all to-area 0.0.0.11
ospf add stub-host 3.3.3.3 to-area 0.0.0.11 cost 1
ospf start
system set name RS3
```

RS3# ip show routes

Destination	Gateway	Owner	Netif
1.1.1.1	11.0.0.2	OSPF_IA	TO-RS2
2.2.2.2	11.0.0.2	OSPF_IA	TO-RS2
3.3.3.3	3.3.3.3	-	lo0
10.0.0.0/24	11.0.0.2	OSPF_IA	TO-RS2
10.0.1.0/24	11.0.0.2	OSPF_IA	TO-RS2
10.0.2.0/24	11.0.0.2	OSPF_IA	TO-RS2
10.0.3.0/24			

```
11.0.0.2                OSPF_IA   TO-RS2
11.0.0.0/24             directly
connected     -          TO-RS2
11.0.1.0/24             directly
connected     -          LAN1
11.0.2.0/24             directly
connected     -          LAN2
11.0.3.0/24             directly
connected     -          LAN3
127.0.0.1
127.0.0.1               -          lo0
```

The following observations and comments can be made:

1. Only routes from area 0.0.0.11 are summarized when injected into the Backbone area. This can be observed in the forwarding table of RS1. Routes from the backbone area are NOT summarized, as evident from the forwarding table of RS3.

2. RS2, the ABR that performs the summarization, does NOT benefit from it. It maintains detailed routing knowledge of both areas in which it participates.

3. When a summary is configured, it will be injected in the neighboring area as long as at least one contributing prefix is reachable from the ABR doing the summary.

Go back to [OSPF Support Home](#)
Last Updated June 4, 2002

Configuring Route Filtering

OSPF routes that are injected from one area into another are called "Summary Routes" and are represented by Type 3 LSAs. The term "Summary" routes is used regardless of whether or not actual summarization is performed with regard to those routes.

You can also configure filtering of the Summary routes on the ABR. You **may not** configure filtering of OSPF External routes (OSPF ASE) on the ABR, those are flooded in all areas capable of handling Externals.

To filter Summary (Type 3) routes on the ABR, issue the following command:

```
ospf add summary-range x.x.x.x/y to-area <AREA_ID> restrict
```

All Type 3 LSAs (routes) that fall within the range of x.x.x.x/y and originate in area "AREA_ID" will not be propagated in the adjacent area.

In the example below we illustrate both summarization and filtering. We will summarize the first two subnets of area 0.0.0.11 into 11.0.0.0/23, and will filter the second two of them, 11.0.2.0/24 and 11.0.3.0/24, when these are injected into the Backbone.

RS1 Configuration and Forwarding Table	RS1# ip show routes

`interface create ip TO-RS2 address-netmask`
`10.0.0.1/24 port et.2.1`
`interface create ip LAN1 address-netmask`
`10.0.1.1/24 port et.2.6`
`interface create ip LAN2 address-netmask`
`10.0.2.1/24 port et.2.7`
`interface create ip LAN3 address-netmask`
`10.0.3.1/24 port et.2.8`
`interface add ip lo0 address-netmask`
`1.1.1.1/32`
`ip-router global set router-id 1.1.1.1`
`ospf create area backbone`
`ospf add interface all to-area backbone`
`ospf add stub-host 1.1.1.1 to-area`
`backbone cost 1`
`ospf start`
`system set name RS1`
	`Destination`
	`Gateway              Owner    Netif`
	`-----------          -------`
	`                     -----    -----`
	`1.1.1.1`
	`1.1.1.1              -        lo0`
	`2.2.2.2`
	`10.0.0.2             OSPF     TO-RS2`
	`3.3.3.3`
	`10.0.0.2             OSPF_IA  TO-RS2`
	`10.0.0.0/24          directly`
	`connected   -        TO-RS2`
	`10.0.1.0/24          directly`
	`connected   -        LAN1`
	`10.0.2.0/24          directly`
	`connected   -        LAN2`
	`10.0.3.0/24          directly`
	`connected   -        LAN3`
	==`11.0.0.0/23`==
	==`10.0.0.2             OSPF_IA  TO-RS2`==
	`127.0.0.1`
	`127.0.0.1            -        lo0`

RS2 Configuration and Forwarding Table	RS2# ip show routes

`interface create ip TO-RS1 address-netmask`
`10.0.0.2/24 port et.2.1`
`interface create ip TO-RS3 address-netmask`
`11.0.0.2/24 port et.2.2`
`interface add ip lo0 address-netmask`
`2.2.2.2/32`
`ip-router global set router-id 2.2.2.2`
`ospf create area backbone`
`ospf create area 0.0.0.11`
`ospf add interface TO-RS1 to-area backbone`
`ospf add interface TO-RS3 to-area 0.0.0.11`
`ospf add stub-host 2.2.2.2 to-area`
`backbone cost 1`
==`ospf add summary-range 11.0.0.0/23 to-area`==
==`0.0.0.11`==
==`ospf add summary-range 11.0.2.0/23 to-area`==
==`0.0.0.11 restrict`==
`ospf start`
`system set name RS2`
	`Destination`
	`Gateway              Owner    Netif`
	`-----------          -------`
	`                     -----    -----`
	`1.1.1.1`
	`10.0.0.1             OSPF     TO-RS1`
	`2.2.2.2`
	`2.2.2.2              -        lo0`
	`3.3.3.3`
	`11.0.0.1             OSPF     TO-RS3`
	`10.0.0.0/24          directly`
	`connected   -        TO-RS1`
	`10.0.1.0/24`
	`10.0.0.1             OSPF     TO-RS1`
	`10.0.2.0/24`
	`10.0.0.1             OSPF     TO-RS1`
	`10.0.3.0/24`
	`10.0.0.1             OSPF     TO-RS1`
	`11.0.0.0/24          directly`
	`connected   -        TO-RS3`

11.0.1.0/24			
11.0.0.1		OSPF	TO-RS3
11.0.2.0/24			
11.0.0.1		OSPF	TO-RS3
11.0.3.0/24			
11.0.0.1		OSPF	TO-RS3
127.0.0.1			
127.0.0.1		–	lo0

RS3 Configuration and Forwarding Table	RS3# ip show routes

interface create ip TO-RS2 address-netmask 11.0.0.1/24 port et.2.2
interface create ip LAN1 address-netmask 11.0.1.1/24 port et.2.6
interface create ip LAN2 address-netmask 11.0.2.1/24 port et.2.7
interface create ip LAN3 address-netmask 11.0.3.1/24 port et.2.8
interface add ip lo0 address-netmask 3.3.3.3/32
ip-router global set router-id 3.3.3.3
ospf create area 0.0.0.11
ospf add interface all to-area 0.0.0.11
ospf add stub-host 3.3.3.3 to-area 0.0.0.11 cost 1
ospf start
system set name RS3
	Destination
	Gateway Owner Netif
	----------- -------
	                          ----- -----
	1.1.1.1
	11.0.0.2 OSPF_IA TO-RS2
	2.2.2.2
	11.0.0.2 OSPF_IA TO-RS2
	3.3.3.3
	3.3.3.3 – lo0
	10.0.0.0/24
	11.0.0.2 OSPF_IA TO-RS2
	10.0.1.0/24
	11.0.0.2 OSPF_IA TO-RS2
	10.0.2.0/24
	11.0.0.2 OSPF_IA TO-RS2
	10.0.3.0/24
	11.0.0.2 OSPF_IA TO-RS2
	11.0.0.0/24 directly connected – TO-RS2
	11.0.1.0/24 directly connected – LAN1
	11.0.2.0/24 directly connected – LAN2
	11.0.3.0/24 directly connected – LAN3
	127.0.0.1
	     127.0.0.1 –
	lo0

Related Topic:

[Configuring Route Summarization on the ABR.](#)

Go back to [OSPF Support Home](#)

Last Updated June 4, 2002

Configuring Aggregation at an ASBR

Configuring Aggregation at an ASBR

Why Aggregate Externals

When external routes are injected into OSPF, they are flooded across areas without an ability to summarize them at Area Border Routers (ABRs). For this reason, in the presence of large number of externals, the benefit of hierarchical design with areas diminishes. Generally, one of the main design goals of an OSPF network is to handle as few as possible external routes.

The number of external routes can be controlled by two complimentary techniques:
1. Configuring areas as Stub, TSA, or NSSA.
2. Aggregating the external routes before they are injected by the ASBR.

The first method will simply prevent Externals from entering other areas. Externals will still exist within the area in which they were injected originally. The second method will reduce the number of Externals at the point of injection.

Where to Aggregate Externals

ASBR Located in a Normal Area

When the ASBR is located in a Normal area, Externals can only be aggregated on the ASBR itself. Once injected into OSPF, they cannot be aggregated downstream. The implication of that rule is that all routes to be aggregated must have adjacent IP prefixes, as they exist on the ASBR that aggregates them.

RS1 Configuration	RS1# ip show routes			
interface create ip TO-RS2 address-netmask 10.0.0.1/24 port et.2.1				
interface create ip LAN1 address-netmask 10.0.1.1/24 port et.2.6	Destination	Gateway	Owner	Netif
interface create ip LAN2 address-netmask 10.0.2.1/24 port et.2.7	-----------	-------	-----	-----
interface create ip LAN3 address-netmask 10.0.3.1/24 port et.2.8	1.1.1.1	1.1.1.1	-	lo0
interface add ip lo0 address-netmask 1.1.1.1/32	2.2.2.2	10.0.0.2	OSPF	TO-RS2
ip-router global set router-id 1.1.1.1	3.3.3.3	10.0.0.2	OSPF_IA	TO-RS2
ospf create area backbone	10.0.0.0/24	directly connected	-	TO-RS2
ospf add interface all to-area backbone	10.0.2.0/24	directly connected	-	LAN2
ospf add stub-host 1.1.1.1 to-area backbone cost 1	10.0.3.0/24	directly connected	-	LAN3
ospf start	11.0.0.0/24	10.0.0.2	OSPF_IA	TO-RS2
system set name RS1	12.0.0.0/22	10.0.0.2	OSPF_ASE	TO-RS2
	127.0.0.1	127.0.0.1	-	lo0

```
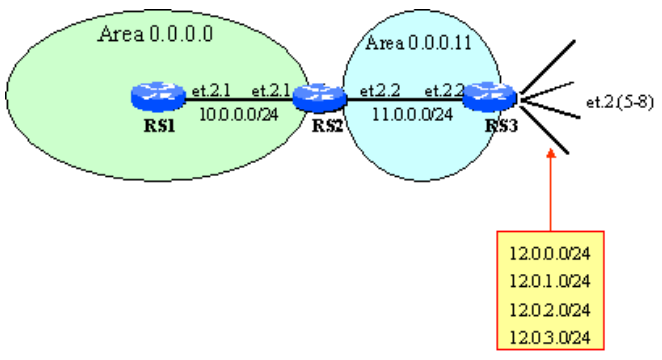RS2 Configuration
interface create ip TO-RS1 address-netmask 10.0.0.2/24 port et.2.1
interface create ip TO-RS3 address-netmask 11.0.0.2/24 port et.2.2
interface add ip lo0 address-netmask 2.2.2.2/32
ip-router global set router-id 2.2.2.2
ospf create area backbone
ospf create area 0.0.0.11
ospf add interface TO-RS1 to-area backbone
ospf add interface TO-RS3 to-area 0.0.0.11
ospf add stub-host 2.2.2.2 to-area backbone cost 1
ospf start
system set name RS2
```

```
RS2# ip show routes

Destination        Gateway             Owner        Netif
-----------        -------             -----        -----
1.1.1.1            10.0.0.1            OSPF         TO-RS1
2.2.2.2            2.2.2.2             -            lo0
3.3.3.3            11.0.0.1            OSPF         TO-RS3
10.0.0.0/24        directly connected  -            TO-RS1
10.0.2.0/24        10.0.0.1            OSPF         TO-RS1
10.0.3.0/24        10.0.0.1            OSPF         TO-RS1
11.0.0.0/24        directly connected  -            TO-RS3
12.0.0.0/22        11.0.0.1            OSPF_ASE     TO-RS3
127.0.0.1          127.0.0.1           -            lo0
```

```
RS3 Configuration
interface create ip TO-RS2 address-netmask 11.0.0.1/24 port et.2.2
interface create ip LAN1 address-netmask 12.0.0.1/24 port et.2.5
interface create ip LAN2 address-netmask 12.0.1.1/24 port et.2.6
interface create ip LAN3 address-netmask 12.0.2.1/24 port et.2.7
interface create ip LAN4 address-netmask 12.0.3.1/24 port et.2.8
interface add ip lo0 address-netmask 3.3.3.3/32
ip-router global set router-id 3.3.3.3
ip-router policy redistribute from-proto aggregate to-proto ospf
                                                      metric 55
ospf create area 0.0.0.11
ospf add stub-host 3.3.3.3 to-area 0.0.0.11 cost 1
ospf add interface TO-RS2 to-area 0.0.0.11
ospf start
system set name RS3
ip-router policy summarize route 12.0.0.0/22
```

```
RS3# ip show routes

Destination        Gateway             Owner        Netif
-----------        -------             -----        -----
1.1.1.1            11.0.0.2            OSPF_IA      TO-RS2
2.2.2.2            11.0.0.2            OSPF_IA      TO-RS2
3.3.3.3            3.3.3.3             -            lo0
10.0.0.0/24        11.0.0.2            OSPF_IA      TO-RS2
10.0.2.0/24        11.0.0.2            OSPF_IA      TO-RS2
10.0.3.0/24        11.0.0.2            OSPF_IA      TO-RS2
11.0.0.0/24        directly connected  -            TO-RS2
12.0.0.0/24        directly connected  -            LAN1
12.0.0.0/22        127.0.0.1           Aggregate    lo0
12.0.1.0/24        directly connected  -            LAN2
12.0.2.0/24        directly connected  -            LAN3
12.0.3.0/24        directly connected  -            LAN4
127.0.0.1          127.0.0.1           -            lo0
```

The highlighted lines of the RS3 config accomplish the aggregation and redistribution of the aggregate route as a Type 5 External with metric 55. The aggregate is generated based on the presence of any route in the range 12.0.0.0/22. It is placed in RS3's forwarding table, along with all specific 12.x.x.x routes. The aggregate is however the only one that is redistributed as External, as evident form the forwarding tables of RS1 and RS2.

ASBR Located in a NSSA

Unlike Type 5 Externals injected in Normal areas, Type 7 Externals injected into an NSSA have two opportunities for aggregation:

1. At the ASBR injecting them (similar to aggregation in Normal area)

2. At the ABR

We will illustrate this with the example below, where RS3 and RS4 are both ASBRs in area 0.0.0.11, which is an NSSA. RS4 has four directly connected (and adjacently addressed interfaces, which it will redistribute as one aggregate, while RS3 has two interfaces and two static routes, but all of them are adjacently addressed. RS3 will aggregate all four into one aggregate route.

As a result of this, two aggregate Externals will arrive at the ABR (RS2), which will summarize those into one single Type 5 External to be injected in the backbone.

Go back to
Last Updated June 4, 2002

Passive Interfaces

In some cases it is desirable to include a subnet into the OSPF routing process (and Link State Database), without actually running OSPF on the interface of the router connected to that subnet. This is particularly useful for interfaces that are used as BGP peering links or for customer connectivity.

On the figure above, router RS1 is using port et.1.2 for BGP peering with RS2. Let's also assume that AS1 is using OSPF as its IGP. The IP subnet representing the link between the two peering routers, must be available to OSPF, as it contains the NEXT_HOP value for all BGP routes that are received from AS2.

There are two choices for incorporating this subnet into OSPF.

1. Redistribute it as an external into OSPF.
2. Include the interface of port et.1.2 into OSPF.

The second approach is typically preferred, as it injects the route as "native" OSPF route, subject to address summarization at an ABR, while the first approach will inject it as a non-summarizable External. The only problem with the first approach is that it is a security "hole", as it opens the possibility of RS1 establishing an OSPF adjacency with RS2 and exchanging OSPF routes between the two, which is clearly not intended.

This is handled by defining the interface as "passive", which injects it as an interior OSPF route, while at the same time the interface does NOT run OSPF, and the router is NOT capable of establishing adjacency over it.

Below are the configurations for the two approaches mentioned above.

RS1 Configuration – with REDISTRIBUTION	RS1 Configuration – with PASSIVE INTERFACE

`interface create ip INSIDE address-netmask 1.1.1.1/24 port et.1.1`
`interface create ip OUTSIDE address-netmask 1.1.2.1/24 port et.1.2`
`!`
`ospf create area backbone`
`ospf add interface INSIDE to-area backbone`
`ospf start`
`!`
`ip-router policy redistribute from-proto direct to-proto ospf`
	`interface create ip INSIDE address-netmask 1.1.1.1/24 port et.1.1`
	`interface create ip OUTSIDE address-netmask 1.1.2.1/24 port et.1.2`
	`!`
	`ospf create area backbone`
	`ospf add interface INSIDE to-area backbone`
	`ospf add interface OUTSIDE to-area backbone`
	`ospf set interface OUTSIDE passive`
	`ospf start`

Last Updated June 4, 2002

Virtual Links

Virtual Links

Background

The OSPF standard (RFC 2328) requires that all non-zero areas attach to the backbone area. In most cases this attachment is physical via an ABR (Area Border Router). In some cases however the physical connectivity does not allow for all areas to attach directly to the backbone. Instead it is only possible for an area (say, Area 2) to attach to a different, non-zero area (say, Area 1). This is common when two companies' OSPF networks are merged. In this case, a Virtual Link must be used across Area 1 (the Transit Area) in order to connect Area 2 to the Backbone.

There are a few design limitations and recommendations regarding Virtual Links:
- o The Transit Area must be Normal or NSSA. It cannot be Stub, or TSA.
- o Virtual Links reduce the intuitiveness of the network and should be avoided. They are appropriate for temporary solutions to inadequate physical connectivity.

To configure a Virtual Link, the following two commands are necessary:

```
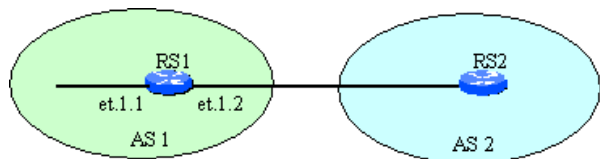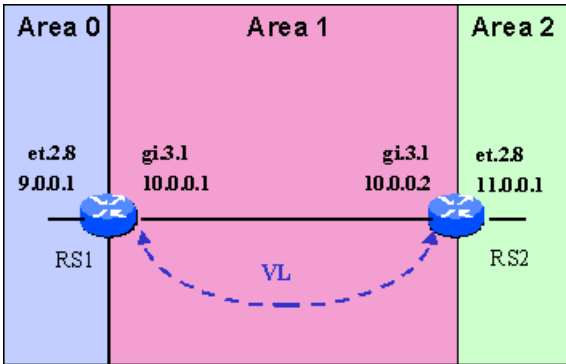ospf add virtual-link <NAME> neighbor <Router ID> transit-area <Area ID>
ospf set virtual-link <NAME> state enable
```

Typical Use

The most common use of a Virtual Link is to connect an area to the Backbone, when no physical connectivity is available. In the figure below, Area 1 is a **Transit Area**, and the virtual link between RS1 and RS2 connects the remote Area 2 to the Backbone.

RS1	RS2
interface create ip IN-BB address-netmask 9.0.0.1/28 port et.2.8 interface create ip IN-AREA1 address-netmask 10.0.0.1/28 port gi.3.1 interface add ip lo0 address-netmask 1.1.1.1/32 ip-router global set router-id 1.1.1.1 ospf create area backbone ospf create area 1 ospf add interface IN-BB to-area backbone ospf add interface IN-AREA1 to-area 1 ospf add virtual-link A2 neighbor 2.2.2.2 transit-area 1 ospf add stub-host 1.1.1.1 to-area backbone cost 1 ospf set virtual-link A2 state enable ospf start system set name RS1	interface create ip IN-AREA2 address-netmask 11.0.0.1/28 port et.2.8 interface create ip IN-AREA1 address-netmask 10.0.0.2/28 port gi.3.1 interface add ip lo0 address-netmask 2.2.2.2/32 ip-router global set router-id 2.2.2.2 ospf create area 1 ospf create area 2 ospf add interface IN-AREA2 to-area 2 ospf add interface IN-AREA1 to-area 1 ospf add virtual-link A2 neighbor 1.1.1.1 transit-area 1 ospf add stub-host 2.2.2.2 to-area 1 cost 1 ospf set virtual-link A2 state enable ospf start system set name RS2

```
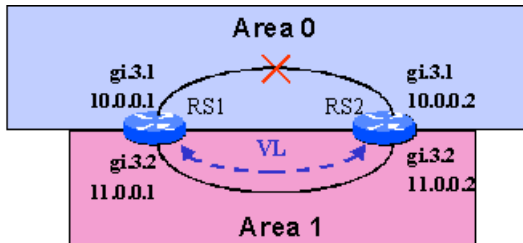RS2# ospf show virtual-links
Internet Address 1.1.1.1, Transit Area 1
    Router ID 2.2.2.2, Network Type Virtual, Cost: 2
    Transmit Delay is 1 sec, State Pt2Pt, Priority 1
    Designated Router (ID) 0.0.0.0, Interface address 1.1.1.1
    Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
        Hello due in 13:27:03
    Neighbor Count is 1
```

Using Virtual Links to Maintain Contiguous Backbone

OSPF requires that the Backbone area be contiguous. OSPF has no mechanism to cope with a situation where Area 0 is split in two or more discontinuous parts. Therefore it is imperative to design the Backbone (Area 0) with sufficient physical redundancy such that it is never split.

In situations where the physical connectivity in the backbone is insufficient, a Virtual Link across non-zero area can be used to provide backup for the physical links within the backbone. The idea is illustrated in the figure below:

In this instance the Backbone area is comprised of a single link between RS1 and RS2. If this link was to become severed, the Backbone will be split, and both RS1 and RS2 will claim that they are connected to "Area 0", but will not communicate to each other over Area 0. To prevent this from happening, a Virtual Link can be configured between RS1 and RS2 **across Area 1**. This link will serve to provide Backbone connectivity between RS1 and RS2 in case the physical link connecting their gi.3.1 ports is disconnected.

RS1	RS2

`interface create ip IN-AREA0 address-netmask 10.0.0.1/24 port gi.3.1`
`interface create ip IN-AREA1 address-netmask 11.0.0.1/24 port gi.3.2`
`interface add ip lo0 address-netmask 1.1.1.1/32`
`ip-router global set router-id 1.1.1.1`
`ospf create area backbone`
`ospf create area 1`
`ospf add stub-host 1.1.1.1 to-area backbone cost 1`
`ospf add interface IN-AREA0 to-area backbone`
`ospf add interface IN-AREA1 to-area 1`
`ospf add virtual-link BB-BACKUP neighbor 2.2.2.2 transit-area 1`
`ospf set virtual-link BB-BACKUP state enable`
`ospf start`
`system set name RS1`
	`interface create ip IN-AREA0 address-netmask 10.0.0.2/24 port gi.3.1`
	`interface create ip IN-AREA1 address-netmask 11.0.0.2/24 port gi.3.2`
	`interface add ip lo0 address-netmask 2.2.2.2/32`
	`ip-router global set router-id 2.2.2.2`
	`ospf create area backbone`
	`ospf create area 1`
	`ospf add stub-host 2.2.2.2 to-area backbone cost 1`
	`ospf add interface IN-AREA0 to-area backbone`
	`ospf add interface IN-AREA1 to-area 1`
	`ospf add virtual-link BB-BACKUP neighbor 1.1.1.1 transit-area 1`
	`ospf set virtual-link BB-BACKUP state enable`
	`ospf start`
	`system set name RS2`

With the virtual link up, there is an additional adjacency that is formed between RS1 and RS2 across the Virtual Link:

```
RS1# ospf show neighbor
Neighbor 2.2.2.2, interface address 10.0.0.2 [mem: 81b28748]
    In the area 0.0.0.0 via interface address 10.0.0.1
    Neighbor priority is 1, State is Full
    Options 0
    Dead timer due in 16:49:46
Neighbor 2.2.2.2, interface address 11.0.0.2 [mem: 81b28500]
    In the area 0.0.0.0 via interface address 2.2.2.2
    Neighbor priority is 1, State is Full
    Options 0
    Dead timer due in 16:49:46
Neighbor 2.2.2.2, interface address 11.0.0.2 [mem: 81b54528]
    In the area 1 via interface address 11.0.0.1
    Neighbor priority is 1, State is Full
    Options 0
```

```
    Dead timer due in 16:49:53
```

Troubleshooting Virtual Links

The most common misconfiguration with Virtual Links is specifying the remote router's **Interface Address** rather than **Router ID** when defining the Virtual Link, i.e. on RS1's configuration above, specifying

```
ospf add virtual-link BB-BACKUP neighbor 10.0.0.2 transit-area 1
```

instead of the correct

```
ospf add virtual-link BB-BACKUP neighbor 2.2.2.2 transit-area 1
```

In this case the Virtual Link will not come up, which is evident from the following output:

```
RS1# ospf show virtual-links
Internet Address 10.0.0.2, Transit Area 1
    Router ID 1.1.1.1, Network Type Virtual, Cost: 1
    Transmit Delay is 1 sec, State Down, Priority 1
    Designated Router (ID) 0.0.0.0, Interface address 10.0.0.2
    No backup designated router on this network
    Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Neighbor Count is 0
```

Go back to OSPF Support Home

Last Updated June 4, 2002

How to Generate a Default Route in OSPF

How to Generate a Default Route in OSPF

Default Routes in Normal Areas

Default Routes in NSSA

Injecting a Conditional Default Route in NSSA

Default Routes in a Stub or Totally Stubby Area (TSA)

Depending on the area type, default routes can originate from several sources in an OSPF network. In general, a default route is injected by ABR routers into Stub/TSA/NSSA areas. In addition, a default route can be injected by an ASBR located in a Normal or NSSA area.   Because of all these possibilities, it is important to understand how the interaction between the different defaults works.

Default Routes in Normal Areas

In Normal areas, the only source of default route can be an ASBR. An ASBR can simply redistribute into OSPF a default route that is already in its forwarding table, or if there is none, generate one and then redistribute it. When generation is necessary, the "no-install" clause is usually used, which causes the ASBR to generate the default route and make it eligible for redistribution, but NOT use it for its own forwarding. This is typically useful on an ASBR that receives full routes from a BGP peer, so it does not use a default route for its own forwarding. The "no-install" clause is optional.

A typical application of this type of default route is the case of multi-homing to different ISPs. RS2 will receive full BGP routes from two upstream ISPs, and will select the optimal BGP route for every destination. There is no need however to propagate that knowledge within the OSPF network and burden the interior routers with the detailed knowledge of the world. Instead, injecting a default route that leads everyone to RS2 is more appropriate.

On the ASBR (RS2) , configure:

```
ip add route default gateway 127.0.0.1 no-install
ip-router policy redistribute from-proto static to-proto ospf
network default
```

The above lines will generate a static default route, which will be injected in OSPF as a Type 5 External. As a Type 5 External, this default route will be flooded throughout all Normal areas of the OSPF network. All routers within the OSPF network will use that default route to reach external destinations. Please note that the "no-install" option will cause the injecting router NOT to use this route for forwarding, instead we want it to forward based on specific routes received from the upstream peers.

Default Routes in NSSA

There are two possible sources of default route in NSSA:

1. A default injected by the ABR with the "ospf set area 0.0.0.1 nssa [nssa-cost x] [nssa-type 1|2]"

2. A default injected by an ASBR within the NSSA with the `"ip-router policy`
   `redistribute from-proto <NAME> to-proto ospf-nssa …"`

From both of these sources, the default route is injected as a Type 7 External. Injection from the ABR can be suspended simply by NOT specifying the "nssa-cost x" clause in the command. In such case, the only source of default in the NSSA would be an ASBR.

Both of these sources may have multiple routers injecting default routes, e.g. several ABRs and several ASBRs. The usual rules of OSPF External routes selection apply; Type 1 metric is preferred over Type 2, and within the same metric type, the metric value is compared.

Injecting a Conditional Default Route in NSSA

Beginning with firmware version ROS7.0.0.3 the ability was added to inject a default route from the ABR of a NSSA based on the presence of a default route in the FIB or the ABR (which would be received from the backbone).   In environments that depend on default route to forward traffic, this ability is useful since it allows for traffic to be dropped at the edge as soon as the default route of the network disappears. If this feature is not used, then the ABR of the NSSA would continue to inject a default into the NSSA thus attracting traffic to itself, even though it has no way of forwarding that traffic (since its own default route has disappeared).

To inject a default route from the ABR into the NSSA based on the presence of an active default route in the ABR, issue the following command:

`ospf set area <Area ID> nssa nssa-cost <cost> conditionally`

For more detailed example of conditional default injection, see the NSSA page.

Default Routes in a Stub or Totally Stubby Area (TSA)

The only possible source of default route in a Stub are or TSA is the ABR of that area (or the ABRs, if multiple are present). The default is injected (just like all the other routes that come from the ARB, as Type 3 Summary LSA. To inject a default in Stub/TSA, configure the following on the ABR(s):

`ospf set area 0.0.0.1 stub stub-cost 22`

"stub-cost 22" will set the OSPF cost of the default route to 22. If multiple ABRs are present and they all inject a default, intra-area routers will follow the default with the cheapest cost.

Go back to [OSPF Support Home](#)

OSPF Network Types

OSPF Network Types

Introduction

Broadcast

NBMA (Non-Broadcast Multi-Access)

Point-to-Point

Point-to-Multipoint

Introduction

In order to correctly build the network topology, OSPF maintains network types, i.e. network segments are described not only with an IP address and mask, but also with the type of network they represent. Four network types are supported by OSPF:

- o   Broadcast

- o   NBMA (Non-Broadcast Multiple Access)

- o   Point-to-Point

- o   Point-to-Multipoint

Based on the Layer 2 interface type (Ethernet, PPP, Frame Relay, ATM, etc.) there are defaults for the OSPF network types. For example, the RS defaults to OSPF type Broadcast for Ethernet interfaces.

In order for two OSPF routers to establish adjacency over a common link, they must agree on the OSPF network type of that link. Different default values for OSPF Network types is the most common reason for interoperability problems. When the two routers have a mismatch in their default setting, one end must be re-configured to match the other. Here we document the default settings and possible values for the RS platform.

Broadcast

For the purposes of OSPF, Broadcast interface is an interface that has native Layer 2 broadcast/multicast capabilities.

The typical example of a broadcast/multicast network is an Ethernet, in which IP Multicast packets are transmitted as Ethernet frames with multicast destination MAC addresses.

On a broadcast network, OSPF will send its HELLOs to IP multicast address of 224.0.0.5. Those HELLOs will be heard by every OSPF router on the network. Therefore no special configuration is required for neighbor discovery – neighbor discovery is automatic.

```
RS1(config)# ip-router global set trace-state on
RS1(config)# ospf trace hello detail
-04-24 14:27:59 OSPF SEND: 10.0.0.1 -> 224.0.0.5: Version 2, Type Hello (1),
Length 48
-04-24 14:27:59          Router ID 1.1.1.1, Area 0.0.0.0, Authentication <None>
(0)
-04-24 14:27:59          Authentication data: 00000000 00000000
-04-24 14:27:59          Mask 255.255.255, Options <E> (2), Priority 1,
Neighbours 1
-04-24 14:27:59          Intervals: Hello 10s, Dead Router 40s, Designated
Router 10.0.0.1, Backup 10.0.0.2
-04-24 14:27:59          Neighbours: 2.2.2.2
```

Designated Router (DR) and Backup Designated Router (BDR) are selected on Broadcast networks.

On the RS router, Broadcast network type is by default assigned to the following interfaces:

- o Ethernet
- o ATM
- o Frame Relay

NBMA (Non-Broadcast Multi-Access)

The NBMA network type was created so that OSPF can correctly model Frame Relay and X.25 environments, where the network is multi-access, but has no native Layer 2 broadcast/multicast capabilities. Instead of multicasting its HELLOs, OSPF uses unicast to send them: In addition, the HELLO timer is extended to 30 seconds and the Dead Router timer is extended to 120 sec.

```
RS1(config)# ip-router global set trace-state on
RS1(config)# ospf trace hello detail
-04-24 15:13:03 OSPF SEND: 10.0.0.1 -> 10.0.0.2: Version 2, Type Hello (1),
Length 44
-04-24 15:13:03          Router ID 1.1.1.1, Area 0.0.0.0, Authentication <None>
(0)
-04-24 15:13:03          Authentication data: 00000000 00000000
-04-24 15:13:03          Mask 255.255.255, Options <E> (2), Priority 1,
Neighbours 0
-04-24 15:13:03          Intervals: Hello 30s, Dead Router 120s, Designated
Router 10.0.0.1, Backup 0.0.0.0
```

Because NBMA networks have no multicast/broadcast capabilities, all OSPF packets are sent as unicast, and therefore OSPF neighbors must be manually configured:

```
ospf add nbma-neighbor 10.0.0.2 to-interface TO-RS2
```

Designated Router (DR) and Backup Designated Router (BDR) are selected on Broadcast networks.

NBMA networks have one significant drawback – they require full-mesh connectivity between all OSPF routers. This is in most cases impractical and prone to unpredictable failure modes, and therefore NBMA network type is usually avoided. Point-to-Multipoint is used instead.

On the RS router, none of the interfaces default to NBMA network type.

Point-to-Point

The point-to-point interface type is reserved for point-to-point links, i.e. links that can only have two routers attached to them. Typical examples of those are PPP, Cisco HDLC, or POS. The neighbor address of the link is usually discovered via a Layer 2 specific mechanism, such as IPCP or Inverse ARP. OSPF however uses a multicast destination IP address for its HELLOs, and leaves it up to the specific Data Link layer to handle those. These multicasts are typically transmitted as Layer 2 unicasts, using the known mapping mentioned above.

```
RS1(config)# ip-router global set trace-state on
RS1(config)# ospf trace hello detail
-04-24 16:07:46 OSPF RECV: 10.0.0.2 -> 224.0.0.5: Version 2, Type Hello (1),
Length 48
-04-24 16:07:46          Router ID 2.2.2.2, Area 0.0.0.0, Authentication <None>
(0)
-04-24 16:07:46          Authentication data: 00000000 00000000
-04-24 16:07:46          Mask 255.255.255.252, Options <E> (2), Priority 0,
Neighbours 1
-04-24 16:07:46          Intervals: Hello 10s, Dead Router 40s, Designated
Router 0.0.0.0, Backup 0.0.0.0
-04-24 16:07:46          Neighbours: 1.1.1.1
```

Because the Layer 3 multicasts are transmitted with Layer 2 unicasts using Layer 2-specific mechanisms, neighbor discovery is automatic.

Designated Router (DR) and Backup Designated Router (BDR) are NOT selected on Point-to-Point networks.

On the RS router, the following interfaces default to Point-to-Point network type:

- o   WAN (Frame Relay and PPP)
- o   POS

Point-to-Multipoint

The Point-to-Multipoint interface type was originally created to be configured on networks that otherwise would have been NBMA, but cannot satisfy the full-mesh connectivity requirement of NBMA. Point-to-Multipoint networks are modeled as a collection of Point-to-Point networks, and just like Point-to-Point, don't have DR/BDR

On the RS router, none of the interfaces default to Point-to-Multipoint network type, however all can be configured as Point-to-Multipoint:

ospf add interface \<NAME> type point-to-multipoint

To read more on the specifics of Point-to-Multipoint interface type, click here.

Go back to OSPF Support Home

Last Updated June 5, 2002

Ethernet Interfaces and OSPF Network Types

Ethernet Interfaces and OSPF Network Types

- Default Network Type for Ethernet Interfaces - Broadcast
- Can be Configured as:
  - NBMA (Non-Broadcast Multi-Access)
  - Point-to-Multipoint

Default Network Type for Ethernet Interfaces - Broadcast

The default OSPF network type for Ethernet interfaces is "Broadcast".

RS1 Configuration	RS2 Configuration
version 7.0 interface create ip TO-RS2 address-netmask 10.0.0.1/24 port et.2.1 interface add ip lo0 address-netmask 1.1.1.1/32 ip-router global set router-id 1.1.1.1 ospf create area backbone ospf add stub-host 1.1.1.1 to-area backbone cost 1 ospf add interface TO-RS2 to-area backbone ospf start system set name RS1	version 7.0 interface create ip TO-RS1 address-netmask 10.0.0.2/24 port et.2.1 interface add ip lo0 address-netmask 2.2.2.2/32 ip-router global set router-id 2.2.2.2 ospf create area backbone ospf add interface TO-RS1 to-area backbone ospf add stub-host 2.2.2.2 to-area backbone cost 1 ospf start system set name RS2

```
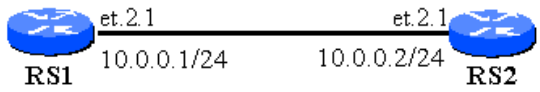RS1# ospf show interfaces
Internet Address 10.0.0.1/24, Area 0.0.0.0
    Router ID 1.1.1.1, Network Type Broadcast, Cost: 20
    Transmit Delay is 1 sec, State DR, Priority 1
    Designated Router (ID) 1.1.1.1, Interface address 10.0.0.1
    Backup Designated Router 2.2.2.2
    Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
        Hello due in 14:28:48
    Neighbor Count is 1

RS1#
```

Can be Configured as:

NBMA (Non-Broadcast Multi-Access)

We mention this here only for completeness. Configuring Ethernet interfaces as NBMA network type has limited usefulness (unless it is done to match the settings of a neighbor that is configured that way).

RS1 Configuration	RS2 Configuration

```
version 7.0                                                      version 7.0
interface create ip TO-RS2 address-netmask 10.0.0.1/24 port et.2.1   interface create ip TO-RS1 address-netmask 10.0.0.2/24 port et.2.1
interface add ip lo0 address-netmask 1.1.1.1/32                  interface add ip lo0 address-netmask 2.2.2.2/32
ip-router global set router-id 1.1.1.1                           ip-router global set router-id 2.2.2.2
ospf create area backbone                                        ospf create area backbone
ospf add stub-host 1.1.1.1 to-area backbone cost 1              ospf add stub-host 2.2.2.2 to-area backbone cost 1
ospf add interface TO-RS2 to-area backbone type non-broadcast   ospf add interface TO-RS1 to-area backbone type non-broadcast
ospf add nbma-neighbor 10.0.0.2 to-interface TO-RS2             ospf add nbma-neighbor 10.0.0.1 to-interface TO-RS1
ospf start                                                       ospf start
system set name RS1                                              system set name RS2
```

```
RS1# ospf show interfaces
Internet Address 10.0.0.1/24, Area 0.0.0.0
    Router ID 1.1.1.1, Network Type Non-Broadcast Multi-Access, Cost: 20
    Transmit Delay is 1 sec, State DR, Priority 1
    Designated Router (ID) 1.1.1.1, Interface address 10.0.0.1
    Backup Designated Router 2.2.2.2
    Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
        Hello due in 22:02:34
    Neighbor Count is 1

RS1#
```

Point-to-Multipoint

Default P-MP – Unicast

RS1 Configuration	RS2 Configuration
version 7.0 interface create ip TO-RS2 address-netmask 10.0.0.1/24 port et.2.1 interface add ip lo0 address-netmask 1.1.1.1/32 ip-router global set router-id 1.1.1.1 ospf create area backbone ospf add stub-host 1.1.1.1 to-area backbone cost 1 ospf add interface TO-RS2 to-area backbone type point-to-multipoint ospf add pmp-neighbor 10.0.0.2 to-interface TO-RS2 ospf start system set name RS1	version 7.0 interface create ip TO-RS1 address-netmask 10.0.0.2/24 port et.2.1 interface add ip lo0 address-netmask 2.2.2.2/32 ip-router global set router-id 2.2.2.2 ospf create area backbone ospf add stub-host 2.2.2.2 to-area backbone cost 1 ospf add interface TO-RS1 to-area backbone type point-to-multipoint ospf add pmp-neighbor 10.0.0.1 to-interface TO-RS1 ospf start system set name RS2

```
RS1# ospf show interfaces
Internet Address 10.0.0.1/24, Area 0.0.0.0
    Router ID 1.1.1.1, Network Type Point-To-Multipoint, Cost: 20
    Transmit Delay is 1 sec, State Pt2Pt, Priority 1
    Designated Router (ID) 0.0.0.0, Interface address 10.0.0.1
    Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
        Hello due in 21:54:54
    Neighbor Count is 1
```

P-MP with Multicast OSPF Packets

RS1 Configuration	RS2 Configuration
version 9.0 interface create ip TO-RS2 address-netmask 10.0.0.1/24 port et.2.1 interface add ip lo0 address-netmask 1.1.1.1/32 ip-router global set router-id 1.1.1.1 ospf create area backbone ospf add stub-host 1.1.1.1 to-area backbone cost 1 ospf add interface TO-RS2 to-area backbone type point-to-multipoint ospf set interface TO-RS2 do-multicast ospf set interface TO-RS2 strict-routers off ospf start system set name RS1	version 9.0 interface create ip TO-RS1 address-netmask 10.0.0.2/24 port et.2.1 interface add ip lo0 address-netmask 2.2.2.2/32 ip-router global set router-id 2.2.2.2 ospf create area backbone ospf add stub-host 2.2.2.2 to-area backbone cost 1 ospf add interface TO-RS1 to-area backbone type point-to-multipoint ospf set interface TO-RS1 do-multicast ospf set interface TO-RS1 strict-routers off ospf start system set name RS2

For more discussion on the reasoning for configuring Ethernet interfaces as Point-to-Multipoint network type, see here.

Ethernet interfaces CANNOT be configured as Point-to-Point OSPF network type.

Last Updated June 4, 2002

PPP Interfaces and OSPF Network Types

PPP Interfaces and OSPF Network Types

- Default Network Type: Point-to-Point
- Can be Configured as:
  - Broadcast
  - NBMA (Non-Broadcast Multiple Access)
  - Point-to-Multipoint

Default Network Type: Point-to-Point

The default OSPF network type for PPP (or Cisco HDLC) interfaces is "Point-to-Point". When the interface is configured and added to OSPF without any special modifiers, it is of point-to-point type.

RS1 Configuration	RS2 Configuration
version 7.0 port set se.4.3 wan-encapsulation ppp speed 56000 interface create ip TO-RS2 address-netmask 10.0.0.1/30 port se.4.3 interface add ip lo0 address-netmask 1.1.1.1/32 ip-router global set router-id 1.1.1.1 ospf create area backbone ospf add stub-host 1.1.1.1 to-area backbone cost 1 ospf add interface TO-RS2 to-area backbone ospf start system set name RS1	version 7.0 port set se.4.3 wan-encapsulation ppp speed 56000 interface create ip TO-RS1 address-netmask 10.0.0.2/30 port se.4.3 interface add ip lo0 address-netmask 2.2.2.2/32 ip-router global set router-id 2.2.2.2 ospf create area backbone ospf add stub-host 2.2.2.2 to-area backbone cost 1 ospf add interface TO-RS1 to-area backbone ospf start system set name RS2

```
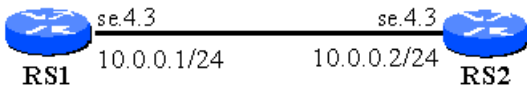RS1# ospf show interfaces
Internet Address 10.0.0.1/30, Area 0.0.0.0
    Router ID 1.1.1.1, Network Type Point-To-Point, Cost: 35714
    Transmit Delay is 1 sec, State Pt2Pt, Priority 1
    Designated Router (ID) 0.0.0.0, Interface address 10.0.0.1
    Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
        Hello due in 22:01:29
    Neighbor Count is 1
RS1#
```

Can be Configured as:

Sometimes it is necessary to override the default OSPF network type in order to match the settings of equipment on the other end of the link.

Broadcast

Simply specify "type broadcast" when you create the IP interface.

RS1 Configuration	RS2 Configuration

```
version 7.0                                                    version 7.0
port set se.4.3 wan-encapsulation ppp speed 56000             port set se.4.3 wan-encapsulation ppp speed 56000
interface create ip TO-RS2 address-netmask 10.0.0.1/30 port se.4.3   interface create ip TO-RS1 address-netmask 10.0.0.2/30 port se.4.3
type broadcast                                                 type broadcast
interface add ip lo0 address-netmask 1.1.1.1/32               interface add ip lo0 address-netmask 2.2.2.2/32
ip-router global set router-id 1.1.1.1                        ip-router global set router-id 2.2.2.2
ospf create area backbone                                     ospf create area backbone
ospf add stub-host 1.1.1.1 to-area backbone cost 1           ospf add stub-host 2.2.2.2 to-area backbone cost 1
ospf add interface TO-RS2 to-area backbone                   ospf add interface TO-RS1 to-area backbone
ospf start                                                    ospf start
system set name RS1                                           system set name RS2
```

```
RS1# ospf show interfaces
Internet Address 10.0.0.1/30, Area 0.0.0.0
    Router ID 1.1.1.1, Network Type Broadcast, Cost: 35714
    Transmit Delay is 1 sec, State DR, Priority 1
    Designated Router (ID) 1.1.1.1, Interface address 10.0.0.1
    No backup designated router on this network
    Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
        Hello due in 22:10:19
    Neighbor Count is 1
```

NBMA (Non-Broadcast Multiple Access)

Create the IP interface as usual, but when adding it to OSPF, add it as "non-broadcast" type. Also, you will have to define NBMA neighbors.

RS1 Configuration	RS2 Configuration
| ```
version 7.0
port set se.4.3 wan-encapsulation ppp speed 56000
interface create ip TO-RS2 address-netmask 10.0.0.1/30 port se.4.3
interface add ip lo0 address-netmask 1.1.1.1/32
ip-router global set router-id 1.1.1.1
ospf create area backbone
ospf add stub-host 1.1.1.1 to-area backbone cost 1
ospf add interface TO-RS2 to-area backbone type non-broadcast
ospf add nbma-neighbor 10.0.0.2 to-interface TO-RS2
ospf start
system set name RS1
``` | ```
version 7.0
port set se.4.3 wan-encapsulation ppp speed 56000
interface create ip TO-RS1 address-netmask 10.0.0.2/30 port se.4.3
interface add ip lo0 address-netmask 2.2.2.2/32
ip-router global set router-id 2.2.2.2
ospf create area backbone
ospf add stub-host 2.2.2.2 to-area backbone cost 1
ospf add interface TO-RS1 to-area backbone type non-broadcast
ospf add nbma-neighbor 10.0.0.1 to-interface TO-RS1
ospf start
system set name RS2
``` |

```
RS1# ospf show interfaces
Internet Address 10.0.0.1/30, Area 0.0.0.0
    Router ID 1.1.1.1, Network Type Non-Broadcast Multi-Access, Cost: 35714
    Transmit Delay is 1 sec, State Back DR, Priority 1
    Designated Router (ID) 2.2.2.2, Interface address 10.0.0.1
    Backup Designated Router 1.1.1.1
    Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
        Hello due in 22:14:53
    Neighbor Count is 1
```

Point-to-Multipoint

Create the IP interface as usual, but when adding it to OSPF, add it as "point-to-multipoint" type.

RS1 Configuration	RS2 Configuration
| ```
version 7.0
port set se.4.3 wan-encapsulation ppp speed 56000
interface create ip TO-RS2 address-netmask 10.0.0.1/30 port se.4.3
interface add ip lo0 address-netmask 1.1.1.1/32
ip-router global set router-id 1.1.1.1
ospf create area backbone
ospf add stub-host 1.1.1.1 to-area backbone cost 1
ospf add interface TO-RS2 to-area backbone type point-to-multipoint
ospf start
system set name RS1
``` | ```
version 7.0
port set se.4.3 wan-encapsulation ppp speed 56000
interface create ip TO-RS1 address-netmask 10.0.0.2/30 port se.4.3
interface add ip lo0 address-netmask 2.2.2.2/32
ip-router global set router-id 2.2.2.2
ospf create area backbone
ospf add stub-host 2.2.2.2 to-area backbone cost 1
ospf add interface TO-RS1 to-area backbone type point-to-multipoint
ospf start
system set name RS2
``` |

```
RS1# ospf show interfaces
```

```
Internet Address 10.0.0.1/30, Area 0.0.0.0
    Router ID 1.1.1.1, Network Type Point-To-Multipoint, Cost: 35714
    Transmit Delay is 1 sec, State Pt2Pt, Priority 1
    Designated Router (ID) 0.0.0.0, Interface address 10.0.0.1
    Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
        Hello due in 22:28:28
    Neighbor Count is 1
```

 For more details on configuring Point-to-Multipoint interface type see here.

Go back to OSPF Support Home
Last Updated June 4, 2002

POS Interfaces and OSPF Network Types

POS Interfaces and OSPF Network Types

[Default Network Type for POS Interfaces – Point-to-Point](#)
[Can be Configured as:](#)
  [NBMA (Non-Broadcast Multi-Access)](#)
  [Point-to-Multipoint](#)
  [Broadcast](#)

Default Network Type for POS Interfaces – Point-to-Point

The default OSPF network type for POS interfaces is "Point-toPoint".

RS1 Configuration	RS2 Configuration
version 7.0 interface create ip TO-RS2 address-netmask 10.0.0.1/24 port so.5.1 interface add ip lo0 address-netmask 1.1.1.1/32 ip-router global set router-id 1.1.1.1 ospf create area backbone ospf add interface TO-RS2 to-area backbone ospf start system set name RS1	version 7.0 interface create ip TO-RS1 address-netmask 10.0.0.2/24 port so.2.1 interface add ip lo0 address-netmask 2.2.2.2/32 ip-router global set router-id 2.2.2.2 ospf create area backbone ospf add interface TO-RS1 to-area backbone ospf start system set name RS2

```
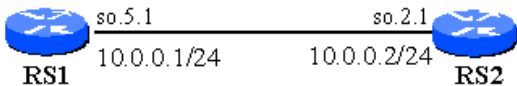RS1#  ospf show interfaces
Internet Address 10.0.0.1/24, Area 0.0.0.0
    Router ID 1.1.1.1, Network Type Point-To-Point, Cost: 12
    Transmit Delay is 1 sec, State Pt2Pt, Priority 1
    Designated Router (ID) 0.0.0.0, Interface address 10.0.0.1
    Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
        Hello due in 11:18:21
    Neighbor Count is 1

RS1#
```

Can be Configured as:

NBMA (Non-Broadcast Multi-Access)

We mention this here only for completeness. Configuring ATM interfaces as NBMA network type has limited usefulness (unless it is done to match the settings of a neighbor that is configured that way, or the PVC is part of a full mesh of PVCs). Consider Point-to-Multipoint configuration as an alternative.

RS1 Configuration	RS2 Configuration

```
version 7.0                                                    version 7.0
interface create ip TO-RS2 address-netmask 10.0.0.1/24 port so.5.1    interface create ip TO-RS1 address-netmask 10.0.0.2/24 port so.2.1
interface add ip lo0 address-netmask 1.1.1.1/32                interface add ip lo0 address-netmask 2.2.2.2/32
ip-router global set router-id 1.1.1.1                         ip-router global set router-id 2.2.2.2
ospf create area backbone                                      ospf create area backbone
ospf add interface TO-RS2 to-area backbone type non-broadcast  ospf add interface TO-RS1 to-area backbone type non-broadcast
ospf add nbma-neighbor 10.0.0.2 to-interface TO-RS2            ospf add nbma-neighbor 10.0.0.1 to-interface TO-RS1
ospf start                                                     ospf start
system set name RS1                                            system set name RS2
```

```
RS1# ospf show interfaces
Internet Address 10.0.0.1/24, Area 0.0.0.0
    Router ID 1.1.1.1, Network Type Non-Broadcast Multi-Access, Cost: 12
    Transmit Delay is 1 sec, State DR, Priority 1
    Designated Router (ID) 1.1.1.1, Interface address 10.0.0.1
    Backup Designated Router 2.2.2.2
    Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
        Hello due in 11:28:46
    Neighbor Count is 1

RS1#
```

Point-to-Multipoint

RS1 Configuration	RS2 Configuration
version 7.0 interface create ip TO-RS2 address-netmask 10.0.0.1/24 port so.5.1 interface add ip lo0 address-netmask 1.1.1.1/32 ip-router global set router-id 1.1.1.1 ospf create area backbone ospf add interface TO-RS2 to-area backbone type point-to-multipoint ospf add pmp-neighbor 10.0.0.2 to-interface TO-RS2 ospf start system set name RS1	version 7.0 interface create ip TO-RS1 address-netmask 10.0.0.2/24 port so.2.1 interface add ip lo0 address-netmask 2.2.2.2/32 ip-router global set router-id 2.2.2.2 ospf create area backbone ospf add interface TO-RS1 to-area backbone type point-to-multipoint ospf add pmp-neighbor 10.0.0.1 to-interface TO-RS1 ospf start system set name RS2

```
RS1# ospf show interfaces
Internet Address 10.0.0.1/24, Area 0.0.0.0
    Router ID 1.1.1.1, Network Type Point-To-Multipoint, Cost: 12
    Transmit Delay is 1 sec, State Pt2Pt, Priority 1
    Designated Router (ID) 0.0.0.0, Interface address 10.0.0.1
    Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
        Hello due in 22:05:14
    Neighbor Count is 1

RS1#
```

Broadcast

By default, a POS port uses routed PPP encapsulation. With it, it is impossible to configure the interface as a "Broadcast" type. In order to do that, the PPP encapsulation must be changed to "Bridged", using the "ppp set ppp-encaps-bgd ports x.x.x" command. Once that is done, the interface on the POS port can now be created as a "Broadcast" type and added to OSPF as such.

RS1 Configuration	RS2 Configuration

```
RS1# ospf show interfaces
Internet Address 10.0.0.1/24, Area 0.0.0.0
     Router ID 1.1.1.1, Network Type Broadcast, Cost: 12
     Transmit Delay is 1 sec, State Back DR, Priority 1
     Designated Router (ID) 2.2.2.2, Interface address 10.0.0.1
     Backup Designated Router 1.1.1.1
     Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
         Hello due in 11:49:49
     Neighbor Count is 1

RS1#
```

Go back to OSPF Support Home

Last Updated June 3, 2002

ATM Interfaces and OSPF Network Types

ATM Interfaces and OSPF Network Types

    Default Network Type for ATM Interfaces - Broadcast
    Can be Configured as:
        NBMA (Non-Broadcast Multi-Access)
        Point-to-Point
        Point-to-Multipoint

Default Network Type for ATM Interfaces - Broadcast

The default OSPF network type for ATM interfaces is "Broadcast". When an IP interface is assigned on top of an ATM PVC, and the interface is added to OSPF, the resulting network type is "Broadcast".

RS1 Configuration	RS2 Configuration
version 7.0 atm create vcl port at.7.1.0.1 interface create ip TO-RS2 address-netmask 10.0.0.1/24 port at.7.1.0.1 interface add ip lo0 address-netmask 1.1.1.1/32 ip-router global set router-id 1.1.1.1 ospf create area backbone ospf add interface TO-RS2 to-area backbone ospf start system set name RS1	version 7.0 atm create vcl port at.5.1.0.1 interface create ip TO-RS1 address-netmask 10.0.0.2/24 port at.5.1.0.1 interface add ip lo0 address-netmask 2.2.2.2/32 ip-router global set router-id 2.2.2.2 ospf create area backbone ospf add interface TO-RS1 to-area backbone ospf start system set name RS2

```
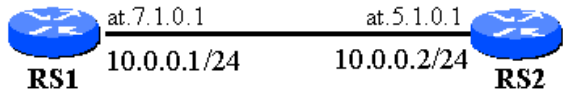RS1# ospf show interfaces
Internet Address 10.0.0.1/24, Area 0.0.0.0
    Router ID 1.1.1.1, Network Type Broadcast, Cost: 12
    Transmit Delay is 1 sec, State DR, Priority 1
    Designated Router (ID) 1.1.1.1, Interface address 10.0.0.1
    No backup designated router on this network
    Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
        Hello due in 21:49:10
    Neighbor Count is 1

RS1#
```

Can be Configured as:

NBMA (Non-Broadcast Multi-Access)

We mention this here only for completeness. Configuring ATM interfaces as NBMA network type has limited usefulness (unless it is done to match the settings of a neighbor that is configured that way, or the PVC is part of a full mesh of PVCs). Consider Point-to-Multipoint configuration as an alternative.

RS1 Configuration	RS2 Configuration

```
version 7.0                                        version 7.0
atm create vcl port at.7.1.0.1                     atm create vcl port at.5.1.0.1
interface create ip TO-RS2 address-netmask 10.0.0.1/24 port   interface create ip TO-RS1 address-netmask 10.0.0.2/24 port
at.7.1.0.1                                         at.5.1.0.1
interface add ip lo0 address-netmask 1.1.1.1/32    interface add ip lo0 address-netmask 2.2.2.2/32
ip-router global set router-id 1.1.1.1             ip-router global set router-id 2.2.2.2
ospf create area backbone                          ospf create area backbone
ospf add interface TO-RS2 to-area backbone type non-broadcast   ospf add interface TO-RS1 to-area backbone type non-broadcast
ospf add nbma-neighbor 10.0.0.2 to-interface TO-RS2   ospf add nbma-neighbor 10.0.0.1 to-interface TO-RS1
ospf start                                         ospf start
system set name RS1                                system set name RS2
```

```
RS1# ospf show interfaces
Internet Address 10.0.0.1/24, Area 0.0.0.0
    Router ID 1.1.1.1, Network Type Non-Broadcast Multi-Access, Cost: 12
    Transmit Delay is 1 sec, State DR, Priority 1
    Designated Router (ID) 1.1.1.1, Interface address 10.0.0.1
    Backup Designated Router 2.2.2.2
    Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
        Hello due in 22:02:34
    Neighbor Count is 1

RS1#
```

Point-to-Point

Configuring an ATM interface as Point-to-Point is common; this is the default setting for some router vendors. If you need to configure an ATM interface of the RS to function as a Point-to-Point OSPF interface, you must do so at the line that creates the interface.

RS1 Configuration	RS2 Configuration
| ```
version 9.0
atm create vcl port at.4.1.0.1
interface create ip TO-RS2 address-netmask 10.0.0.1/24 port
at.4.1.0.1 type point-to-point
interface add ip lo0 address-netmask 1.1.1.1/32
ip-router global set router-id 1.1.1.1
ospf create area backbone
ospf add interface TO-RS2 to-area backbone
ospf start
system set name RS1
``` | ```
version 9.0
atm create vcl port at.4.1.0.1
interface create ip TO-RS1 address-netmask 10.0.0.2/24 port
at.4.1.0.1 type point-to-point
interface add ip lo0 address-netmask 2.2.2.2/32
ip-router global set router-id 2.2.2.2
ospf create area backbone
ospf add interface TO-RS1 to-area backbone
ospf start
system set name RS2
``` |

```
RS1# ospf show interfaces
Internet Address 10.0.0.1/24, Area 0.0.0.0
    Router ID 1.1.1.1, Network Type Point-To-Point, Cost: 12
    Transmit Delay is 1 sec, State Pt2Pt, Priority 1
    Designated Router (ID) 0.0.0.0, Interface address 10.0.0.1
    Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
        Hello due in 14:42:24
    Neighbor Count is 1
    Authentication not enabled
```

In order for the above configuration to work, it is assumed that both ends of the connection support Inverse ARP. If one end does not support Inverse ARP, then the other end must be configured with an explicit Point-to-Point Peer address:

```
RS1(config)# interface create ip TO-RS2 address-netmask 10.0.0.1/24 port at.4.1.0.1 type point-to-point peer-address 10.0.0.2
```

Point-to-Multipoint

RS1 Configuration	RS2 Configuration
| ```
version 7.0
atm create vcl port at.7.1.0.1
interface create ip TO-RS2 address-netmask 10.0.0.1/24 port
at.7.1.0.1
interface add ip lo0 address-netmask 1.1.1.1/32
ip-router global set router-id 1.1.1.1
ospf create area backbone
ospf add interface TO-RS2 to-area backbone type point-to-multipoint
ospf add pmp-neighbor 10.0.0.2 to-interface TO-RS2
ospf start
system set name RS1
``` | ```
version 7.0
atm create vcl port at.5.1.0.1
interface create ip TO-RS1 address-netmask 10.0.0.2/24 port
at.5.1.0.1
interface add ip lo0 address-netmask 2.2.2.2/32
ip-router global set router-id 2.2.2.2
ospf create area backbone
ospf add interface TO-RS1 to-area backbone type point-to-multipoint
ospf add pmp-neighbor 10.0.0.1 to-interface TO-RS1
ospf start
system set name RS2
``` |

```
RS1# ospf show interfaces
Internet Address 10.0.0.1/24, Area 0.0.0.0
    Router ID 1.1.1.1, Network Type Point-To-Multipoint, Cost: 12
    Transmit Delay is 1 sec, State Pt2Pt, Priority 1
    Designated Router (ID) 0.0.0.0, Interface address 10.0.0.1
    Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
        Hello due in 22:05:14
    Neighbor Count is 1

RS1#
```

Go back to OSPF Support Home

Last Updated June 3, 2002

How to Control DR/BDR Election

How to Control DR/BDR Election

Background

Frame Relay or ATM Hub-and-Spoke Network

A Router with Many Interfaces

Background

The OSPF protocol provides for automatic election of Designated Router (DR) and Backup Designated Router (BDR) on Broadcast and NBMA segments. There is no DR/BDR election on Point-to-Point and Point-to-Multipoint networks. In most cases the default DR/BDR election is fine, but in some situation it must be controlled. Below are two typical scenarios where controlling the DR/BDR is appropriate.

The OSPF interface priority is the parameter that controls DR election. Its default value is "1". When OSPF interface priority is "0", the router is not eligible to become a DR. If a router comes up on a network segment and there are no other routers there already, it will become a DR. If there is an existing DR (or BDR), the new router will NOT attempt to preempt them. If several routers come up at roughly the same time (or the DR fails), a new DR election process will be triggered. The router with the highest priority value then will become a DR.

In short, the DR/BDR election process is not deterministic and depends on the sequence of events. Therefore it is important to be able to prevent routers from EVER becoming DR/BDR, when that is appropriate for the topology in which they are connected. This si accomplished by setting their OSPF interface priority to "0":

```
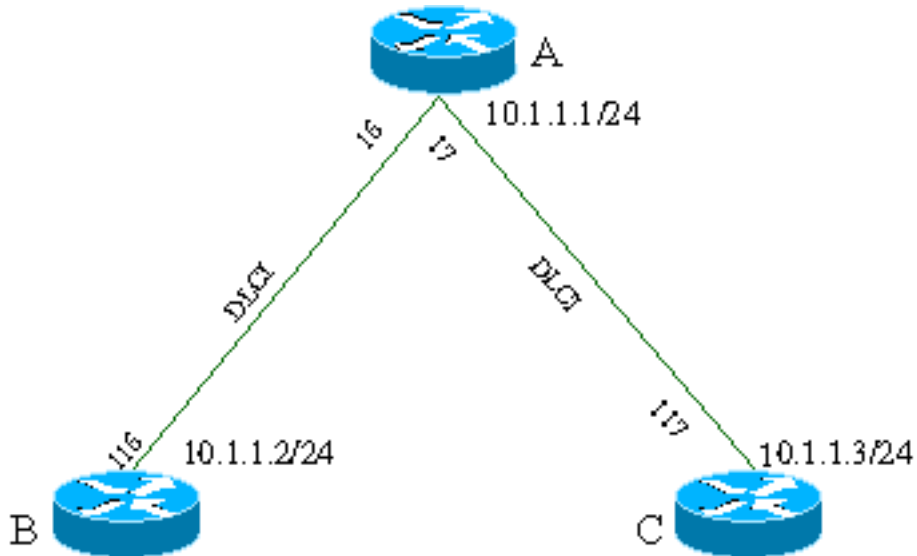ospf set interface <NAME> priority 0
```

Frame Relay or ATM Hub-and-Spoke Network

In a hub-and-spoke topology, such as the one illustrated below (for the case of Frame Relay), only one router (A) has connectivity to all the other routers. Routers B and C do not have direct connectivity to each other. Because of that, they should not be allowed to become DR/BDR. Only router A is eligible to become DR. Below are the router' configurations to accomplish that.


```
Router A:
port set se.4.1 wan-encapsulation frame-relay speed 1536000
frame-relay set lmi state enable ports se.4.1
frame-relay create vc port se.4.1.16
frame-relay create vc port se.4.1.17

vlan create WAN port-based id 10
interface create ip WAN_INT address-netmask 10.1.1.1/24 vlan WAN
vlan add ports se.4.1.16-17 to WAN

ospf create area backbone
ospf add interface WAN_INT to-area backbone type non-broadcast
ospf add nbma-neighbor 10.1.1.2 to-interface WAN_INT
ospf add nbma-neighbor 10.1.1.3 to-interface WAN_INT
ospf set interface WAN_INT priority 1
ospf start
```

```
Router B:
port set se.4.1 wan-encapsulation frame-relay speed 1536000
frame-relay set lmi state enable ports se.4.1
frame-relay create vc port se.4.1.116

vlan create WAN port-based id 10
interface create ip WAN_INT address-netmask 10.1.1.2/30 vlan WAN
vlan add ports se.4.1.116 to WAN

ospf create area backbone
ospf add interface WAN_INT to-area backbone type non-broadcast
ospf set interface WAN_INT priority 0
ospf start
```

```
Router C:
port set se.4.1 wan-encapsulation frame-relay speed 1536000
frame-relay set lmi state enable ports se.4.1
frame-relay create vc port se.4.1.117

vlan create WAN port-based id 10
interface create ip WAN_INT address-netmask 10.1.1.3/30 vlan WAN
vlan add ports se.4.1.117 to WAN

ospf create area backbone
ospf add interface WAN_INT to-area backbone type non-broadcast
ospf set interface WAN_INT priority 0
ospf start
```

A Router with Many Interfaces

The role of a DR/BDR places additional demand on the routers selected to do it. If a router has many interfaces, it is a good practice to configure some of them as ineligible to become DR/BDR. On those interfaces, it will be expected that other routers will take this role. This can be controlled by selectively configuring interfaces with zero and non-zero priority.

Go back to OSPF Support Home

Last Updated June 4, 2002

Point-to Multipoint Interfaces

Point-to Multipoint Interface Configuration

Why Configure Point-to-Multipoint
  Replacement for NBMA
  Replacement for Broadcast

Methods for Configuring Point-to-Multipoint Interfaces
  Default – Unicast
  Multicast

Why Configure Point-to-Multipoint

Replacement for NBMA

The Point-to Multipoint interface type was originally conceived as a replacement for interfaces that otherwise would have been NBMA. The P-MP type has the advantage that it is modeled as a collection of Point-to-Point links, and therefore poses no requirements for full-mesh connectivity that NBMA does.

Replacement for Broadcast

In many cases Ethernet interfaces today are used as "point-to-point" connections, i.e., there are only two routers connected on a common Ethernet segment. This is especially true when Gigabit Ethernet links are used as a POS replacement for short- to medium-range distances – a very common practice in Metro networks.

However even though there are only two routers on such Ethernet segments, the DR/BDR election process still occurs, as the Ethernet interfaces default to Broadcast network type. This election is sometimes undesirable for two reasons:

- It causes a mandatory waiting interval before the adjacency is established. This waiting interval is mandated by the standards for Broadcast (as well as NBMA) interfaces, and is equal to the Dead Interval timer (default 40 seconds). The purpose of this waiting interval is to ensure that there is no existing DR on the segment, before a router decides to promote itself to be DR.
- Since one of the routers is selected as a DR, and the other as a BDR, both routers incur additional overhead (compared to the case when they are neither DR nor BDR).

For these reasons, Point-to-Multipoint interface type is frequently configured on Ethernet segments that have only two (or just a few) routers on them. This achieves faster establishment of adjacency over those interfaces, as well as lower overhead on the routers for maintaining those adjacencies.

See the section below for the two methods of configuring P-MP interfaces.

Methods for Configuring Point-to-Multipoint Interfaces

Default – Unicast

By default, if you simply configure a Point-to-Multipoint interface, OSPF packets out that interface will be sent as Unicast packets. For that reason, it is also necessary to configure the IP address of the other end(s) of the P-MP interface.

```
ospf add interface <NAME> type point-to-multipoint
ospf add pmp-neighbor <IP Address> to-interface <NAME>
```

Unless the Point-to-Multipoint neighbor is configured in this fashion, no HELLOs will be sent, and HELLOs received from neighbors will be ignored.

You can observe the unicast HELLOs sent/received by enabling OSPF tracing:

```
RS#(config) ip-router global set trace-state on
RS#(config) ospf trace hello
RS#(config) save active

RS1# -06-04 13:40:40 OSPF RECV: 10.0.0.2 -> 10.0.0.1: Version 2, Type Hello (1),
Length 48
-06-04 13:40:40          Router ID 2.2.2.2, Area 0.0.0.0, Authentication <MD5>
(2)
-06-04 13:40:40          Zero: 0, Key ID: 1, Length: 16, Sequence: 0
-06-04 13:40:40          Mask 255.255.255, Options <E> (2), Priority 0,
Neighbours 1
-06-04 13:40:40          Intervals: Hello 10s, Dead Router 40s, Designated Router
0.0.0.0, Backup 0.0.0.0
-06-04 13:40:40          Neighbours: 1.1.1.1
-06-04 13:40:45 OSPF SEND: 10.0.0.1 -> 10.0.0.2: Version 2, Type Hello (1),
Length 48
-06-04 13:40:45          Router ID 1.1.1.1, Area 0.0.0.0, Authentication <MD5>
(2)
-06-04 13:40:45          Zero: 0, Key ID: 1, Length: 16, Sequence: 0
-06-04 13:40:45          Mask 255.255.255, Options <E> (2), Priority 0,
Neighbours 1
-06-04 13:40:45          Intervals: Hello 10s, Dead Router 40s, Designated Router
0.0.0.0, Backup 0.0.0.0
-06-04 13:40:45          Neighbours: 2.2.2.2
```

Multicast

To enable dynamic neighbor discovery on Point-to-Multipoint interfaces, and to ensure interoperability with some

vendors, you can configure a Point-to-Multipoint interface to send and receive Multicast OSPF packets. In that case, there is no need to configure explicitly every Point-to-Multipoint neighbor.

To enable multicast packet sending/receipt on Point-to-Multipoint interfaces, two commands are necessary:

```
RS2(config)# ospf set interface TO-RS1 do-multicast on
RS2(config)# ospf set interface TO-RS1 strict-routers off
```

The first command causes the router to send multicast OSPF packets. The second one enables it to receive HELLO packets from unknown P-MP neighbors. If the second command is omitted, the router will be receiving multicast HELLOs from a neighbor, but will reject them, since the default behavior of P-MP interfaces is to ONLY acknowledge HELLOs from known neighbors. Here is what HELLO Tracing will show in this instance:

```
RS1# -06-04 13:46:40 OSPF RECV: 10.0.0.2 -> 224.0.0.5: Version 2, Type Hello
(1), Length 44
-06-04 13:46:40          Router ID 2.2.2.2, Area 0.0.0.0, Authentication <MD5>
(2)
-06-04 13:46:40          Zero: 0, Key ID: 1, Length: 16, Sequence: 0
-06-04 13:46:40          Mask 255.255.255, Options <E> (2), Priority 0,
Neighbours 0
-06-04 13:46:40          Intervals: Hello 10s, Dead Router 40s, Designated Router
0.0.0.0, Backup 0.0.0.0
-06-04 13:46:40 OSPF: invalid packet: Unknown neighbour (12)
-
```

Go back to OSPF Support Home

Last Updated June 6, 2002

How to Tag OSPF Routes

OSPF can carry a 32-bit tag value for routes that are redistributed into OSPF. This tag value can then be used for selective redistribution out of OSPF. The original intention of the tag value was to carry limited BGP attribute information when BGP routes are redistributed into OSPF. However most service providers today avoid redistributing BGP routes into OSPF and vice versa.

Nevertheless, the tagging capability is there and can be used via the redistribution command:

```
ip-router policy redistribute from-proto bgp source-as <AS-Number> network x.x.x.x/y to-proto ospf tag <Value> exact
```

Last Updated June 4, 2002

Go back to OSPF Support Home

Alternative Area Border Router (ABR)

Traditional ABR

Example 1

Traditional ABR Redefined

Alternative ABR

Example 1 (Revisited)

Alternative ABRs Should Only Be Used To Connect Backbone-Connected Areas

Example 2
Example 3

Caveat: Routing Asymmetry

The RS automatically supports the Alternative ABR implementation, as defined in the IETF "Alternative OSPF ABR Implementations" Internet Working Draft. This feature improves the behavior of a router connected to multiple areas without a direct backbone attachment itself. Behavior modifications allow the Alternative ABR to successfully forward routes to the backbone and other areas despite not being actively connected to the backbone.

> **Note: The RS implements the Alternative ABR feature automatically. No configuration changes are necessary.**

To understand the motivation for this feature, first consider the traditional ABR.

The Traditional ABR

The OSPF routing domain consists of subdomains, called areas, that limit the scope of LSA flooding. According to RFC 2328, a router having attachments to multiple areas is called an area border router, or ABR. The RFC places no restrictions on the type of areas an ABR connects.

The IETF *Alternative OSPF ABR Implementations Internet Working Draft* describes the primary role of an ABR as "to provide its attached areas with Type-3 and Type-4 LSAs (which are used for describing routes and ASBRs in other areas) as well as to perform actual inter-area routing."

In doing so, an ABR observes the following behavioral restrictions:

- An ABR injects into the backbone (Area 0) and other areas only those Summary LSAs (inter-area routes) derived from its connected areas' intra-area routes.
- An ABR only considers those Summary LSAs (inter-area routes) that are present in the backbone's link-state database when performing SPF routing calculations.

Example 1

We can examine these restrictions using the following topology, which is frequently deployed in a technique called "dual homing". Dual homing uses several different OSPF areas to provide redundant connectivity to an external network. In the following topology, R3 is the portal to the external network. Areas 1 and 2 provide redundant connectivity.

Example 1

In this example, R1 is an ABR that connects Area 1 to Area 0. R4 is an ABR that connects Area 2 to Area 0. By traditional definition, R3, because it connects Area 1 to Area 2, is also an ABR, even though it has no direct backbone connection. Since R3 must act in accordance with the above ABR restrictions, we make the following observations:

R3's Routing Table

> An ABR injects into the backbone (Area 0) only those Summary LSAs (inter-area routes) derived from its connected areas' intra-area routes.

- Since R3 must inject intra-area routes derived from its connected areas (1 and 2) into the backbone, its routing table contains intra-area routes from those two areas.

  An ABR only considers those Summary LSAs (inter-area routes) that are present in the backbone's link-state database when performing SPF routing calculations.

- Since R3 is not connected to the backbone, its routing table contains no Summary LSAs.

Area 1

- R3 injects intra-area routes for its connected area, Area 2, into Area 1 as Summary LSAs.
- R1 does this as well, in addition to injecting Summary LSAs from the backbone into Area 1.

Area 2

- R3 injects intra-area routes for its connected area, Area 1, into Area 2 as Summary LSAs.
- R4 does this as well, in addition to injecting Summary LSAs from the backbone into Area 2.

From the last observation, we know that R5 learns of networks in Area 0 through R4. To reach R4, the topology requires that R5 forward the packet to R3. (Even if R5 has other means of reaching R4, OSPF path selection may still mandate, for cost reasons, that it go through R3.)

When the packet from R5 that is destined for Area 0 arrives at R3, however, R3 does not know how to reach Area 0. (Recall from observations 1 and 2 that because R3 is not connected to the backbone, its routing table contains no Summary LSAs for Area 0. Only intra-area routes from directly-connected areas, Areas 1 and 2, exist in R3's routing table.) Therefore, R3 drops the packet.

In this situation, R3 fails as a traditional ABR because it is not connected to the backbone and the rules governing a traditional ABR restrict it from learning Summary LSAs from any other area.

The Traditional ABR Redefined

In a standard OSPF implementation, this situation can be remedied using virtual links. However, virtual links can be difficult to configure and impose significant system overhead.

The Alternative ABR is less costly and was designed to fix this problem. A part of the Alternative ABR specification is refining the definition of a traditional ABR.

Riverstone's OSPF implementation considers a router to be an ABR if it satisfies three requirements:

- Has one or more non-backbone areas *actively attached*. As defined in the IETF working draft, "An area is considered actively attached if the router has at least one interface in that area in the state other than Down."
- Has Area 0 configured.
- Has an interface in the Up state in Area 0. This is not equivalent to requiring that the interface in Area 0 be actively attached. Unlike active attachment, this requirement is satisfied even if the adjacent interface on the Area 0 peer is in the Down state. As long as the ABR's interface in Area 0 has not been administratively shut down (ex. using the **interface down** or **port disable** command), it will continue to function as an ABR.

The Alternative ABR

The definition of an Alternative ABR stems from that of the ABR.

If an ABR that is actively attached to more than one non-backbone area ceases to satisfy the above Area 0 requirements (configured and an interface in the Up state), it begins to function as an Alternative ABR, provided that its non-backbone areas are connected to the backbone themselves. (We examine the latter point in detail in Examples 2 and 3.)

The Alternative ABR obeys the traditional ABR definition and behavior with the following modifications:

- It is *not* allowed to originate Summary LSAs. (Only traditional ABRs can generate Summary LSAs), and
- It *is* allowed to consider Summary LSAs from *all* actively attached areas when performing SPF calculations.

Example 1 (Revisited)

Using these modifications, we revisit Example 1 and observe the following changes about R3:

- R3 is no longer an ABR because it does not have an interface in the backbone in the Up state.
- Since it is no longer an ABR, R3 cannot generate Summary LSAs.
- Because it is no longer an ABR, R3 is now permitted to consider Summary LSAs from all attached areas in routing decisions. In this case, R3 can now learn routes to Area 0 from both R1 and R4, and has no trouble forwarding R5's packet to the correct next hop.

The Alternative ABR Should Only Be Used To Connect Backbone-Connected Areas

The Alternative ABR should only be used to connect backbone-connected areas. To understand why, consider the following example:

Example 2

Example 2

In this example, R4 is an ABR that connects Areas 0, 1, and 2. Ignoring the backbone-connected areas requirement for the moment, consider what happens when R4 loses its connection to the backbone.

Case 1: R1's Backbone Interface Is Administratively Shut Down

If R4 loses its connectivity to the backbone because R1's backbone interface is administratively shut down, R4 continues to act as a traditional ABR. As discussed earlier, an ABR is only required to have the backbone configured and an Up interface in the backbone. In this case, since R4 still has the backbone configured and its backbone interface is not shut down, R4 continues to satisfy both requirements.

Case 2: R4's Backbone Interface Is Administratively Shut Down

If R4 loses its connectivity to the backbone because its backbone interface is administratively shut down, it qualifies to become an Alternative ABR. (Ignoring the backbone-connected areas requirement.) Becoming an Alternative ABR means that R4 ceases to generate any Summary LSAs. It only generates intra-area routes for its connected areas: Areas 1 and 2. Because of this, R5 and R6 cease to learn about networks in Area 1 from R4. R1, R2, and R3 also cease to learn about networks in Area 2 from R4. Since it is prohibited from generating Summary LSAs, R4 fulfills no better purpose in this topology than that of an intra-area router in Area 1 and Area 2, providing no connectivity between them.

Example 3

Now let's revisit Case 2 in the situation where Area 2 is connected to the backbone.

Example 3

If Area 2 is connected to the backbone, R6, acting as a traditional ABR, would inject summaries about Area 2 into the backbone. It would also take the summaries about Area 1 from the backbone and inject those into Area 2.

R1 would do the same thing. It injects summaries about Area 1 into the backbone and injects Area 2 summaries from Area 0 into Area 1.

All of the routers in this topology learn about all of the networks through the backbone via R1 and R6. So R4 losing its backbone connection and becoming an Alternative ABR does not impact the connectivity of the network.

The key to achieving meaningful routing through an Alternative ABR is ensuring that the areas the Alternative ABR connects are connected to the backbone themselves. They must not rely solely on the

Alternative ABR, which cannot generate any Summary LSAs, for information about other areas. As the IETF draft reiterates, "[This feature does] not obviate the need of virtual link configuration in case an area has no physical backbone connection at all. The methods described here improve the behavior of a router connecting two or more backbone-attached areas."

Caveat: Routing Asymmetry

The IETF working draft also states that "The changes of the OSPF ABR operations do not influence any aspects of the router-to-router cooperation and do not create routing loops." Nor do they pose any additional security concerns. Hence, they are "fully compatible with standard OSPF."

However, be aware that deploying Alternative ABRs can lead to routing asymmetry.

[Example 1](#) illustrates this. In that topology, note that because it loses its connection to the backbone, R3 functions as an Alternative ABR. It does not generate any Summary LSAs, but learns Summary LSAs from Areas 1 and 2.

R5 learns of network 1.0.0.0/24 through the Summary LSAs that R4 injects from the backbone into Area 2. Similarly, R2 learns of network 2.0.0.0/24 through the Summary LSAs that R1 injects from the backbone into Area 1.

When R5 has a packet destined for network 1.0.0.0/24, it sends the packet to R4 through R3. R3, however, has a closer route to network 1.0.0.0/24, through R2. So when R3 receives R5's packet, instead of sending the packet to R4, it forwards the packet to R2 directly.

When R2 has a packet destined for network 2.0.0.0/24, however, it does not send the packet to R3. Instead, it sends the packet to R1 because it learned about the 2.0.0.0/24 network through the routes that R1 injects from the backbone into its area. From R1, R2's packet moves to R4 and R3 before arriving at R5.

R2's packet traverses this longer path due to the fact that because R3 does not meet the definition of a traditional ABR, it does not generate Summary LSAs for Area 2 into Area 1. Therefore, R2 does not know that network 2.0.0.0/24 is reachable, even closer, through R3.

So despite their topological similarities, packets from R2 to network 2.0.0.0/24 take a less optimal path than packets from R5 to network 1.0.0.0/24. This routing asymmetry results from the fact that even though R3 cannot inject Summary LSAs into its attached areas, it can still use them.

Go back to [OSPF Support Home](#)

Last Updated June 3, 2002

The *"ospf show interfaces"* and *"ospf monitor interfaces"* Commands

These two commands will display the state of all interfaces that are running OSPF. "ospf show interfaces I more verbose and detailed, while "ospf monitor interfaces" is more compact and convenient when a quick glance at the status of all OSPF interfaces is needed. In the example below, there is only one interface running OSPF; we will use it for illustration here.


```
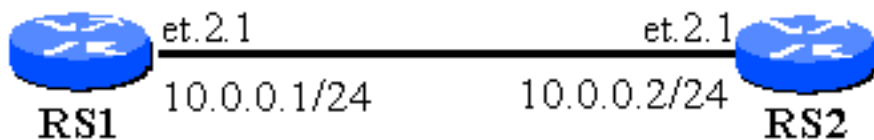RS1# ospf show interfaces
Internet Address 10.0.0.1/24, Area 0.0.0.0
    Router ID 1.1.1.1, Network Type Broadcast, Cost: 20
    Transmit Delay is 1 sec, State DR, Priority 1
    Designated Router (ID) 1.1.1.1, Interface address 10.0.0.1
    Backup Designated Router 2.2.2.2
    Timer intervals configured, Hello 10, Dead 40, Wait 40,
Retransmit 5
        Hello due in 15:09:28
    Neighbor Count is 1

RS1#


RS1# ospf  monitor  interfaces
```

```
Area: 0.0.0.0
IP Address        Type  State     Cost Pri DR               BDR
--------------------------------------------------------------------
----
10.0.0.1          Bcast Back DR   20   1    2.2.2.2          1.1.1.1
```

Internet Address

> The IP address of that interface.

Area

> The area ID in this the interface is configured.

Router ID

> The router ID of the router. If one is configured, it should be here, otherwise the highest interface IP address will be in this field.

Network Type

> The OSPF network type assigned to this interface. Make sure these match on both routers that share a link.

Cost

> The OSPF cost assigned to the interface. By default it will be 2*10^9/Interface speed, but can be overridden manually.

Transmit Delay

> The amount of time the router waits before it floods an LSA over this link.

State

One of the following possible states of an interface:

DR: The router is the designated router on the interface. Applicable for broadcast and NBMA interfaces.

    Back DR: The router is the backup designated router (BDR) on the interface. Applicable for broadcast and NBMA interfaces.

    Other: The router is neither the DR nor the BDR on this interface.

    Waiting: The interface is in a waiting state before deciding if it will be DR/BDR/Other. This state is applicable for Broadcast and NBMA interface type. The length of the wait should never be longer than 40 seconds for Broadcast interfaces and 120 seconds for NBMA interfaces.

    `Pt2Pt`: This interface is either Point-to-Point or Point-to-Multipoint.

Priority

Router priority. The router with the highest priority on a subnet becomes DR (Designated Router). Router priority defaults to 1. When the router priority is set to 0 on an interface, the router is not eligible to become DR/BDR on that network.

Designated Router (ID)

The router ID of the DR for this network.

Interface Address

The interface address of the DR with which it connects to this network.

Backup Designated Router

The router ID for the BDR on this network.

Timer Intervals

The standard 4 OSPF timer values for this interface:

Hello: Time in seconds between HELLO messages. Default for Broadcast, Point-to-Point, and Point-to-Multipoint networks, is 10 seconds, and for NBMA networks, is 30 seconds.

Dead: Number of seconds without hearing a HELLO message before declaring a neighbor down. This timer bounds the worst-case scenario for falure detection by OSPF. The default is 4 timer the HELLO interval.

Wait: When an OSPF router first comes up on a segment, it waits this many seconds before it initiates a DR election.

Retransmit: Number of seconds to wait for acknowledgement of a database description (DD) packet before retransmitting.

Hello due in: When will the router send the next HELLO out this interface.

Neighbor Count

Number of neighbors on this interface from which HELLOs have been received.

Go back to OSPF Support Home
Last Updated June 4, 2002

The *"ospf show neighbor"* and *"ospf monitor neighbors"* Commands

These two commands provide similar, but differently formatted information regarding the status of the OSPF neighbors. "ospf show neighbor" is more detailed and verbose, while "ospf monitor neighbor" is more compact and useful when we want to take a glance at the status of many neighbors in a tabulated format. Both commands are illustrated using the following topology.

"ospf show neighbor"

```
RS1# ospf show neighbor
Neighbor 2.2.2.2, interface address 10.0.0.2 [mem: 81697150]
    In the area 0.0.0.0 via interface address 10.0.0.1
    Neighbor priority is 1, State is Full
    Options 0
    Dead timer due in 22:11:30
```

"ospf monitor neighbors"

```
RS1# ospf monitor neighbors
```

```
Codes: E  - Interface to neighbor does not belong to stub/NSSA
area
       NP - Interface to Neighbor belongs to NSSA area
       MS - Neighbor is the Master during Database exchange
       I  - Initial packet being exchanged during DB exchange
       M  - More packets to be exchanged during DB exchange
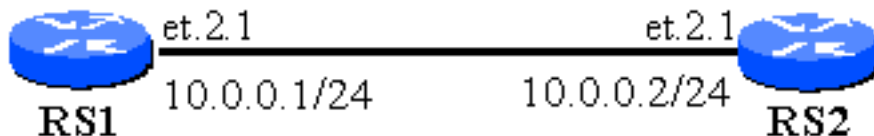

Interface: 10.0.0.1        Area: 0.0.0.0
Neighbor ID      Nbr IP Addr      State       Mode    Options      Pri
-----------------------------------------------------------------------
2.2.2.2          10.0.0.2         Full                E            1
```

Neighbor (Neighbor ID)

The router ID of the neighbor.

Interface Address (Nbr IP Addr)

The IP address of the interface the neighbor uses to connect to this network.

Neighbor Priority (Pri)

The OSPF interface priority of the neighbor. Relevant to the selection process for DR
(Designated Router).

State

The state in the adjacency we have reached with this neighbor. The two normal states are
"Two-way" and "Full". All others are transitional states and of a router is stuck in them for
a prolonged period of time, this is an indication of a problem.

Options

Explained in the command (ospf monitor neighbor).

Mode

Before adjacency is established, this field will be populated with either Master (M) or Slave (S). These determine the roles of the neighbors in the connection. After the adjacency is established, this field is depopulated.

Go back to [OSPF Support Home](#)

Last Updated June 4, 2002

OSPF Tracing

[OSPF Tracing](#)

[Introduction](#)
[Configuration](#)
[Examples](#)

Introduction

Tracing is a very useful method to troubleshoot routing protocols, including OSPF, on GateD implementations, including that of Riversotne Networks. Tracing allows protocol packets to be displayed (and decoded) on the console, or redirected to a TELNET session. One should use caution when enabling tracing as it may be very chatty and overwhelm the terminal it displays on, without giving the operator the ability to stop it.

If you are unsure about the traffic you expect tracing to generate, make sure that you have two sessions to the RS – one on which you display the tracing output, and the other which can be used to issue a command to stop tracing in case it overwhelms the first session.

By default, the output of the tracing appears on the system console. It can be copied on a TELNET session, by issuing the following command for the TELNET session:

```
cli terminal monitor on
```

Configuration

To enable tracing, two steps are necessary:

1.   Enable tracing capabilities for GateD (in general):

```
RS(config)# ip-route global set trace-state on
```

2.   Configure what exactly to be traced. In the case of OSPF, this is done with the command

```
RS(config)# ospf trace <parameters>
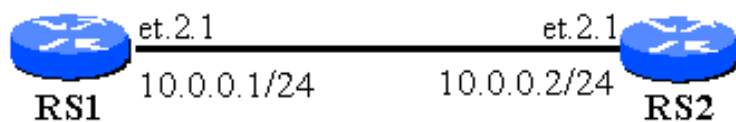```

The possible parameters of the above command are listed below:

```
RS1(config)# ospf trace ?
 ack                    - Traces OSPF Link State Ack packets used in
                           synchronizing OSPF databases.
```

```
db                        - Traces changes to the OSPF LSA database
dd                        - Traces OSPF Database Description packets used in
                            synchronizing OSPF databases.
debug                     - Traces OSPF at the debugging level of detail
drelect                   - Traces OSPF designated router election process
flood                     - Traces OSPF flooding algorithm
hello                     - Traces OSPF HELLO packets used to determine
                            neighbour reachability
local-options             - Sets various trace options for this protocol only.
                            By default, these trace-options are inherited from
                            those specified by the 'ip-router global set
                            trace-options' command.
packets                   - Traces OSPF packets
request                   - Traces OSPF Link State Request packets used in
                            synchronizing OSPF databases.
spf                       - Traces Shortest Path First (SPF) calculations
trap                      - Traces OSPF trap packets
update                    - Traces OSPF Link State Update packets used in
                            synchronizing OSPF databases.
```

The options are self-explanatory. One of the most useful ones is HELLO tracing. It is helpful when troubleshooting mismatched router entries

Examples

This example illustrates how OSPF tracing allows for monitoring the HELLO messages sent and received by a router.

```
RS1(config)# ospf trace hello detail
RS1(config)# ip-router global set trace-state on
RS1(config)# save active
RS1# cli terminal monitor on
%CLI-I-MONITOR_ON, Terminal monitor enabled
RS1# -05-04 22:59:40 OSPF SEND: 10.0.0.1 -> 224.0.0.5: Version 2, Type Hello (1), Length 48
-05-04 22:59:40         Router ID 1.1.1.1, Area 0.0.0.0, Authentication <None> (0)
-05-04 22:59:40         Authentication data: 00000000 00000000
-05-04 22:59:40         Mask 255.255.255, Options <E> (2), Priority 1, Neighbours 1
-05-04 22:59:40         Intervals: Hello 10s, Dead Router 40s, Designated Router 10.0.0.2, Backup
10.0.0.1
-05-04 22:59:40         Neighbours: 2.2.2.2
-05-04 22:59:40 OSPF RECV: 10.0.0.2 -> 224.0.0.5: Version 2, Type Hello (1), Length 48
-05-04 22:59:40         Router ID 2.2.2.2, Area 0.0.0.0, Authentication <None> (0)
-05-04 22:59:40         Authentication data: 00000000 00000000
-05-04 22:59:40         Mask 255.255.255, Options <E> (2), Priority 1, Neighbours 1
-05-04 22:59:40         Intervals: Hello 10s, Dead Router 40s, Designated Router 10.0.0.2, Backup
10.0.0.1
-05-04 22:59:40         Neighbours: 1.1.1.1
```

```
-05-04 22:59:50 OSPF SEND: 10.0.0.1 -> 224.0.0.5: Version 2, Type Hello (1), Length 48
-05-04 22:59:50          Router ID 1.1.1.1, Area 0.0.0.0, Authentication <None> (0)
-05-04 22:59:50          Authentication data: 00000000 00000000
-05-04 22:59:50          Mask 255.255.255, Options <E> (2), Priority 1, Neighbours 1
-05-04 22:59:50          Intervals: Hello 10s, Dead Router 40s, Designated Router 10.0.0.2, Backup
10.0.0.1
-05-04 22:59:50          Neighbours: 2.2.2.2
-05-04 22:59:50 OSPF RECV: 10.0.0.2 -> 224.0.0.5: Version 2, Type Hello (1), Length 48
-05-04 22:59:50          Router ID 2.2.2.2, Area 0.0.0.0, Authentication <None> (0)
-05-04 22:59:50          Authentication data: 00000000 00000000
-05-04 22:59:50          Mask 255.255.255, Options <E> (2), Priority 1, Neighbours 1
-05-04 22:59:50          Intervals: Hello 10s, Dead Router 40s, Designated Router 10.0.0.2, Backup
10.0.0.1
-05-04 22:59:50          Neighbours: 1.1.1.1

RS1# cli terminal monitor off
%CLI-I-MONITOR_OFF, Terminal monitor disabled
RS1#
```

Check out the following examples that illustrate using the OSPF tracing capabilities to troubleshoot specific problems.

Why are neighbors stuck in "INIT" state?
Why is a neighbor stuck in "Exchange/ExStart" state?

Go back to OSPF Support Home

Last Updated June 4, 2002

Understanding OSPF Neighbor States

When two OSPF neighbors establish an adjacency, they go through a finite state machine (FSM). For each neighbor, the router maintains a separate, independent from the other neighbors, FSM. Each FSM goes through several states before establishing adjacency. The two normal long-term states are **"2-way"** and **"Full"**. Depending on the role a router (and its neighbors) play on a particular network, we expect the router to reach either "Two-way" or "Full" state with its neighbors. However it should never maintain any other state with a neighbor for a prolonged period of time.

Full

Loading

Exchange

ExStart

2-Way

Init

Attempt

Down

In the example below RS1 has one neighbor on its interface et.2.1, and 8 neighbors on interface et.2.2. We will expect RS1 to always be in "Full" state with RS2. It may be in "Full" or Two-way" state with routers RS3-10, depending on who the DR/BDR for that network segment is.

Below we discuss each neighbor state in more details.

Down

This is the initial stat of the neighbor relationship. It is typically very short-lived for broadcast networks. If a router has a neighbor in a "Down" state that means that the router has not received any information from that neighbor, yet it know about it. This state can occur for NBMA or Point-to-Multipoint neighbors, which are configured (the router knows about them), but have not sent any packets to the receiving router. This state cannot occur on Broadcast interface types, because on those the only way to know about a neighbor is to receive their HELLOs. Therefore on Broadcast interface types, a neighbor is typically non-existent or is in "Init" state.

Attempt

This state is also only valid for neighbors reachable via NBMA interfaces. It indicates that the neighbor has not sent us any information, but we are trying to contact it via sending regular unicast HELLOs to it.

Init

When we receive a HELLO packet from a neighbor, but it does NOT include our router ID in it, we will place that neighbor in an "Init" state. This means that the neighbor has most likely NOT received a HELLO packet from us and therefore does not know our router ID. Once the neighbor

receives a HELLO packet from us and recognizes our router ID, it will send us HELLO with it included, which will cause us to advance to "2-way" state with that neighbor.

2-Way

In this state the communication between us and the neighbor is bi-directional. This means that we have seen our own routr ID in the HELLOs that arrive from the neighbor. On Point-to-Point and Point-to-Multipoint interfaces, we expect a neighbor to always advance from this to "Full" state. On Broadcast and NBMA interfaces, only the DR/BDR will advance to "Full" state with their neighbors, there remaining neighbors will remain in "2-Way" state.

ExStart

This is the beginning of adjacency formation between two neighbors. At this state they decide who is the Master and who is the Slave between them. The Master gets to increment Sequence Numbers in the database description packets to follow. The Master/Slave election is on a per-neighbor basis. A router could be Master for some adjacencies, and Slave for others. The Master/Slave election is irrelevant to DR/BDR election for a network.

Exchange

During this state the router describes it Link State Database (LSDB) by sending a series of Database Description (DD) Packets to the neighbor. These packets are acknowledged using sequence numbers.

Loading

During this stage, the router explicitly requests LSAs from the neighbor that might have been missed or corrupted during the Exchange process.

Full

This is the normal state of an adjacency (except in some cases in a Broadcast network – see the discussion under 2-Way). Adjacency is fully formed in this state.

Go back to [OSPF Support Home](#)

Go back to OSPF Support Home

Why are Neighbors Stuck in "Init" State?

[Why are Neighbors Stuck in "Init" State?](#)

[Background](#)
[Possible Causes:](#)
    [Access Control List](#)
    [Authentication Mismatch](#)
[Troubleshooting](#)

Background

A router places an OSPF neighbor in "Init" state when it receives a HELLO packet from it that does not contain its own router ID. In the diagram below, router RS1 will place RS2 in "Init" state when it receives a HELLO from RS2, but that HELLO does not contain RS1's router ID.

"Init" should be a short-lived state, until the two neighbors "meet" each other. If the neighbor is in "Init" for more that the duration of "Router Dead" interval, it indicates a problem.

RS1 Configuration	RS2 Configuration

```
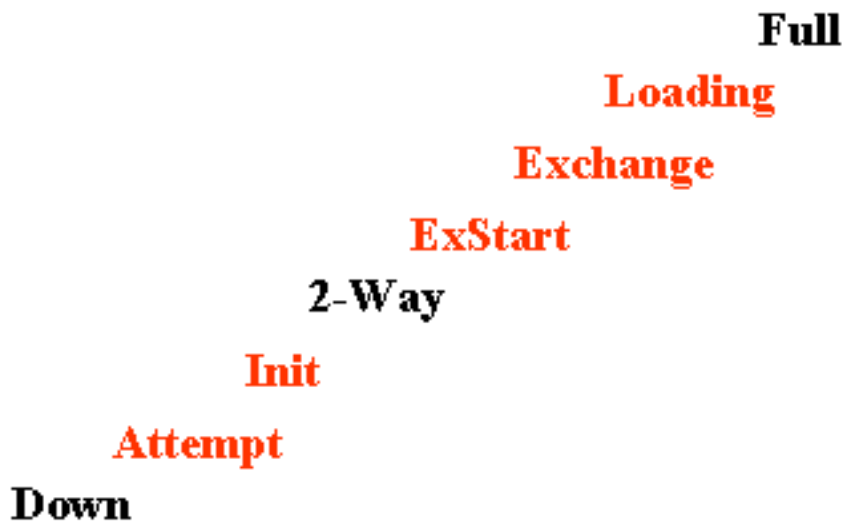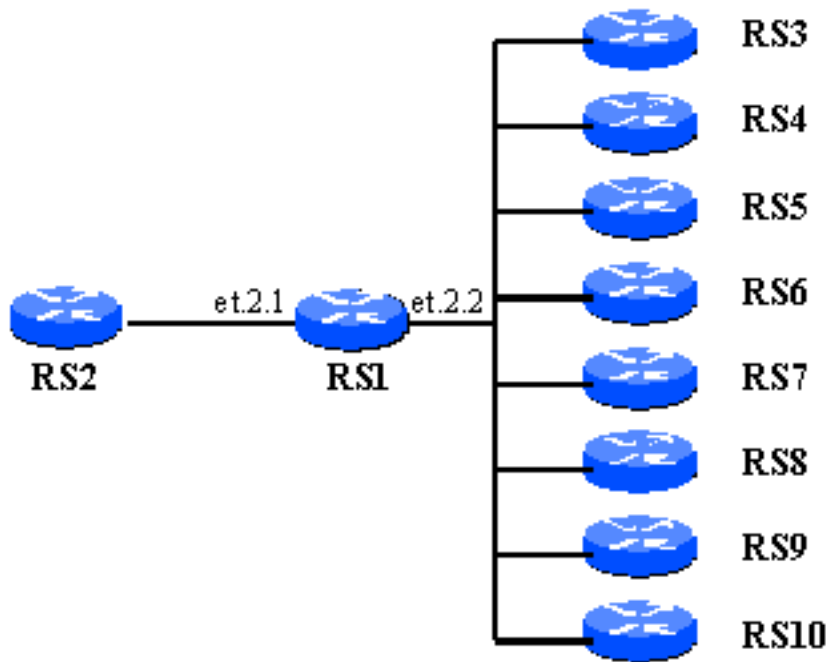version 7.0                                      version 7.0
interface create ip TO-RS2 address-netmask       interface create ip TO-RS1 address-netmask
10.0.0.1/24 port et.2.1                           10.0.0.2/24 port et.2.1
interface add ip lo0 address-netmask 1.1.1.1/32  interface add ip lo0 address-netmask 2.2.2.2/32
ip-router global set router-id 1.1.1.1           ip-router global set router-id 2.2.2.2
ospf create area backbone                        ospf create area backbone
ospf add stub-host 1.1.1.1 to-area backbone      ospf add interface TO-RS1 to-area backbone
cost 1                                            ospf add stub-host 2.2.2.2 to-area backbone
ospf add interface TO-RS2 to-area backbone       cost 1
ospf start                                        ospf set area backbone authentication-method
system set name RS1                               simple key-chain 3
                                                  ospf start
                                                  system set name RS2
```

```
RS1# ospf show neighbor
Neighbor 2.2.2.2, interface address 10.0.0.2 [mem: 81694760]
    In the area 0.0.0.0 via interface address 10.0.0.1
    Neighbor priority is 1, State is Init
    Options 0
    Dead timer due in 11:20:36
RS1#
```

Possible Causes:

Access Control List

There might be an access control list that prevents RS1 to send multicast packets to RS2, and therefore RS2 is unaware of the existence of RS1. Ensure that multicast packets (with destination IP addresses of 224.0.0.5 – AllSPFRouters) can reach from RS1 to RS2.

Authentication Mismatch

Mismatch in the OSPF authentication settings for the two neighbors is frequently the reason for this problem. The router that is configured with no authentication will see its neighbor in "Init" state, while the router configured with authentication will not its neighbor at all (because it will ignore the non-authenticated packets coming from it).

Troubleshooting

As a first step, check the configuration on both ends to make sure there is no obvious misconfiguration.

Secondly, enabling OSPF HELLO tracing on one or both ends will frequently reveal the problem. In the example shown above, we enable OSPF tracing on RS1 and look for a mismatch:

```
RS1(config)# ip-router global set trace-state on
RS1(config)# ospf trace hello
RS1(config)# save active
```

```
RS1(config)# Exit
RS1# cli terminal monitor on
%CLI-I-MONITOR_ON, Terminal monitor enabled
RS1# -05-07 11:02:56 OSPF SEND: 10.0.0.1 -> 224.0.0.5: Version 2, Type Hello (1), Length 48
-05-07 11:02:56        Router ID 1.1.1.1, Area 0.0.0.0, Authentication <None> (0)
-05-07 11:02:56        Authentication data: 00000000 00000000
-05-07 11:02:56        Mask 255.255.255, Options <E> (2), Priority 1, Neighbours 1
-05-07 11:02:56        Intervals: Hello 10s, Dead Router 40s, Designated Router 10.0.0.1, Backup
0.0.0.0
-05-07 11:02:56        Neighbours: 2.2.2.2
-05-07 11:02:56 OSPF RECV: 10.0.0.2 -> 224.0.0.5: Version 2, Type Hello (1), Length 44
-05-07 11:02:56        Router ID 2.2.2.2, Area 0.0.0.0, Authentication <Simple> (1)
-05-07 11:02:56        Authentication data: 00000000 00000000
-05-07 11:02:56        Mask 255.255.255, Options <E> (2), Priority 1, Neighbours 0
-05-07 11:02:56        Intervals: Hello 10s, Dead Router 40s, Designated Router 10.0.0.2, Backup
0.0.0.0
```

The authentication settings mismatch is evident. So is the fact that RS1 recognizes RS2 as an neighbor (because it accepts all OSPF packets with no authentication), while RS2 does NOT recognize RS1 as a neighbor, because it drops RS1's packets which fail the authentication check.

On the RS2 side, the "ospf show neighbors" command reveals nothing, as all OSPF packets from RS1 are dropped as non-authenticated.

Go back to OSPF Support Home
Last Updated June 4, 2002

Why are Neighbors Stuck in "Exchange" or "ExStart" State?

Why are Neighbors Stuck in "Exchange" or "ExStart" State?

- Background
- Possible Causes
  - Mismatched MTU sizes between the neighbors
  - Change in the Options Field From the Neighbor

Background

The "Exchange" phase is the beginning of adjacency formation between two neighbors. It follows the "2-Way" state for neighbors that have decided to proceed and establish full adjacency. During the "ExStart" state, the two neighbors decide who will be the "Master" and who will be the "Slave" in the conversation between them. During the "Exchange" state the two neighbors exchange their databases by sending and acknowledging Database Description (DD) packets.

"ExStart" and "Exchange" are short-lived states. Once a router enters one of these states with a neighbor, it should proceed to Full state within several seconds (depending on the size of the LSDB and the network type).

Possible Causes

Mismatched MTU sizes between the neighbors

If one of the neighbors has a large LSDB to transmit to the other neighbor, when it reaches the Exchange state with that neighbor, it will pack an MTU-worth of LSAs in a DD packet and will attempt to send them to the neighbor. If the neighbor is incapable of processing a packet this large, it will drop it and fail to acknowledge it. Which will lead to infinite repetitive attempts from the originating router to re-send the DD packet. While doing that the routers will be in "Exchange" state.

There are some subtleties in how neighbors form different vendors handle mismatched MTU sizes. Depending on these, some neighbors may be stuch in "Exchange" while other will be stuck in "ExStart" state. Some vendors will allow neighbors to progress from "ExStart" to "Exchange" state even though there is an MTU mismatch, which will result in the behavior described in the previous paragraph. Others will NOT progress from "ExStart" as soon as the MTU mismatch is discovered. The RS behaves this way. Regardless of the differences, the troubleshooting is the same.

In the example below we have simulated an MTU mismatch situation, in which RS1 is receiving large number of BGP prefixes and is redistributing those to OSPF, which results in a large LSDB that it needs to transmit to RS2 once they reach "Exchange" state. However RS2 has been configured with an MTU of 1450 on its gi.3.2 port, while RS1 uses the default 1500. In real world, this type of mismatch typically occurs on Frame Relay and ATM PVCs, on which different vendors have adopted different default values.

RS1 Configuration	RS2 Configuration

```
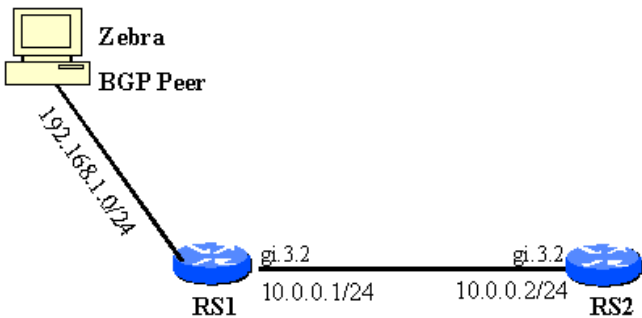version 7.0                                                    version 7.0
interface create ip PEER address-netmask 192.168.1.201/24 port  port set gi.3.2 mtu 1450
et.2.2                                                          port mirroring monitor-port et.2.2 target-port gi.3.2
interface create ip TO-RS2 address-netmask 10.0.0.1/24 port gi.3.2  interface create ip TO-RS1 address-netmask 10.0.0.2/24 port gi.3.2
interface add ip lo0 address-netmask 1.1.1.1/32                interface add ip lo0 address-netmask 2.2.2.2/32
ip-router global set router-id 1.1.1.1                         interface add ip en0 address-netmask 192.168.1.202/24
ip-router global set autonomous-system 65057                   ip-router global set router-id 2.2.2.2
ip-router policy redistribute from-proto bgp source-as 65002 to-  ospf create area backbone
proto ospf                                                     ospf add stub-host 2.2.2.2 to-area backbone cost 1
ospf create area backbone                                      ospf add interface TO-RS1 to-area backbone
ospf add interface all to-area backbone                        ospf start
ospf add stub-host 1.1.1.1 to-area backbone cost 1             system set name RS2
ospf start                                                     ospf trace local-options state normal
bgp create peer-group ZEBRA type external autonomous-system 65002  ospf trace dd detail
bgp add peer-host 192.168.1.101 group ZEBRA                    RS2#
bgp set peer-host 192.168.1.101 group ZEBRA hold-time 10000
bgp start
system set name RS1
```

```
RS1# ospf show neighbor
Neighbor 2.2.2.2, interface address 10.0.0.2 [mem: 81cab800]
    In the area 0.0.0.0 via interface address 10.0.0.1
    Neighbor priority is 1, State is Exch Start
    Options 7
    Dead timer due in 15:15:18


RS1#

RS2# ospf show neighbor
Neighbor 1.1.1.1, interface address 10.0.0.1 [mem: 81694c98]
    In the area 0.0.0.0 via interface address 10.0.0.2
    Neighbor priority is 1, State is Exch Start
    Options 7
    Dead timer due in 12:51:04


RS2#
```

Troubleshooting

As a first step, check the configuration on both ends to make sure there is no obvious misconfiguration.

Secondly, enabling OSPF Database Description Packet tracing on one or both ends will frequently reveal the problem.
The error message in the tracing output typically refers to a confusing Master/Initial flag. In the example shown above,
we enable OSPF DD tracing on RS2 and look for the error message:

```
RS2(config)# ip-router global set trace-state on
RS2(config)# ospf trace dd detail
RS2(config)# save active
RS2(config)# Exit
RS2# cli terminal monitor on

%CLI-I-MONITOR_ON, Terminal monitor enabled

RS2# -05-08 12:43:49 OSPF SEND: 10.0.0.2 -> 224.0.0.5: Version 2, Type Hello (1), Length 48
-05-08 12:43:49           Router ID 2.2.2.2, Area 0.0.0.0, Authentication <None> (0)
-05-08 12:43:49           Authentication data: 00000000 00000000
-05-08 12:43:49 OSPF SEND: 10.0.0.2 -> 10.0.0.1: Version 2, Type Database Desciption (2), Length 32
-05-08 12:43:49           Router ID 2.2.2.2, Area 0.0.0.0, Authentication <None> (0)
-05-08 12:43:49           Authentication data: 00000000 00000000
-05-08 12:43:49           Options <E> (2), Flags <MS M I> (7), Sequence 9
-05-08 12:43:49 OSPF RECV: 10.0.0.1 -> 224.0.0.5: Version 2, Type Hello (1), Length 48
-05-08 12:43:49           Router ID 1.1.1.1, Area 0.0.0.0, Authentication <None> (0)
-05-08 12:43:49           Authentication data: 00000000 00000000
-05-08 12:43:49 OSPF RECV: 10.0.0.1 -> 10.0.0.2: Version 2, Type Database Desciption (2), Length 32
-05-08 12:43:49           Router ID 1.1.1.1, Area 0.0.0.0, Authentication <None> (0)
-05-08 12:43:49           Authentication data: 00000000 00000000
-05-08 12:43:49           Options <E> (2), Flags <MS M I> (7), Sequence 10
-05-08 12:43:49 OSPF: invalid DD packet: Confusing Master/Initial flags (21)
-05-08 12:43:54 OSPF SEND: 10.0.0.2 -> 10.0.0.1: Version 2, Type Database Desciption (2), Length 32
-05-08 12:43:54           Router ID 2.2.2.2, Area 0.0.0.0, Authentication <None> (0)
-05-08 12:43:54           Authentication data: 00000000 00000000
-05-08 12:43:54           Options <E> (2), Flags <MS M I> (7), Sequence 9
-05-08 12:43:54 OSPF RECV: 10.0.0.1 -> 10.0.0.2: Version 2, Type Database Desciption (2), Length 32
-05-08 12:43:54           Router ID 1.1.1.1, Area 0.0.0.0, Authentication <None> (0)
```

```
-05-08 12:43:54        Authentication data: 00000000 00000000
-05-08 12:43:54        Options <E> (2), Flags <MS M I> (7), Sequence 10
-05-08 12:43:54 OSPF: invalid DD packet: Confusing Master/Initial flags (21)
```

The MTU sizes on both OSPF neighbors that share a link must be the same in order for the adjacency to proceed. To change the port MTU size on the RS, issue the following command:

`"port set <port-name> mtu <size> "`

Note also that on the RS, an MTU size can be assigned to an IP interface when the interface is created:

`"Interface create ip <NAME> address-netmask x.x.x.x/y mtu <size>"`

This MTU size is used by the router with regards to packets that it is originating (routing protocols, ICMP echos, etc). However if the PORT MTU is smaller than the Interface MTU, the Interface MTU will be ignored and the maximum size of packets originated by the router will still be controlled by the PORT MTU.

For packets that are passing through the router, the **port MTU** is the only one that applies.

Change in the Options Field From the Neighbor

Some vendors change the Options field between the HELLO and DD stage of the adjacency formation. While this behavior is not advisable, it is not specifically prohibited by the standard. Before ver. 7.0.0.3 an RS router would reject DD packets which have an Option field that is different from that negotiated during HELLO. In these situations, the neighbor state will remain at Exchange. In ver. 7.0.0.3, and enhancement was added to allow for such changes form the neighbor and proceed with the adjacency establishment.

Go back to OSPF Support Home
Last Updated June 4, 2002