Cabrillo College

CCNP – Advanced Routing

Ch. 5 OSPF - Multi-areas (Part II)

Rick Graziani, Instructor

Mar 16, 2002

1

This Week

OSPF

- E1 and E2 routes
- Default Routes
- Route Summarization
- NSSA (Not So Stubby Areas)
- Virtual Links
- Load Balancing
- show commands

2

Quick Review

- Areas
- LSAs
- Stub Area
- Totally Stubby Area

3

Area Types

- Standard or Normal Areas
  - Backbone
  - Non-Backbone
- Stub
  - Stub Area
  - Totally Stubby Area (TSA)
  - Not-so-stubby-area (NSSA)

4

Area Types

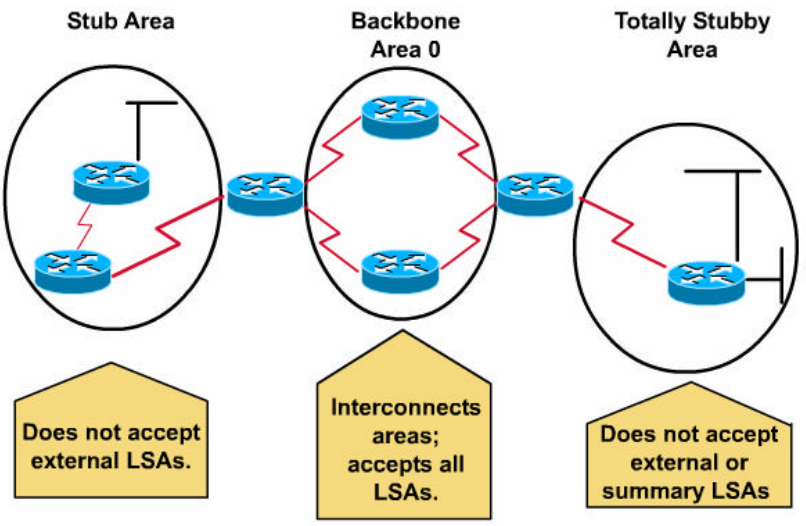

5

LSA 1 Router LSA

Multi Area OSPF

Normal Areas

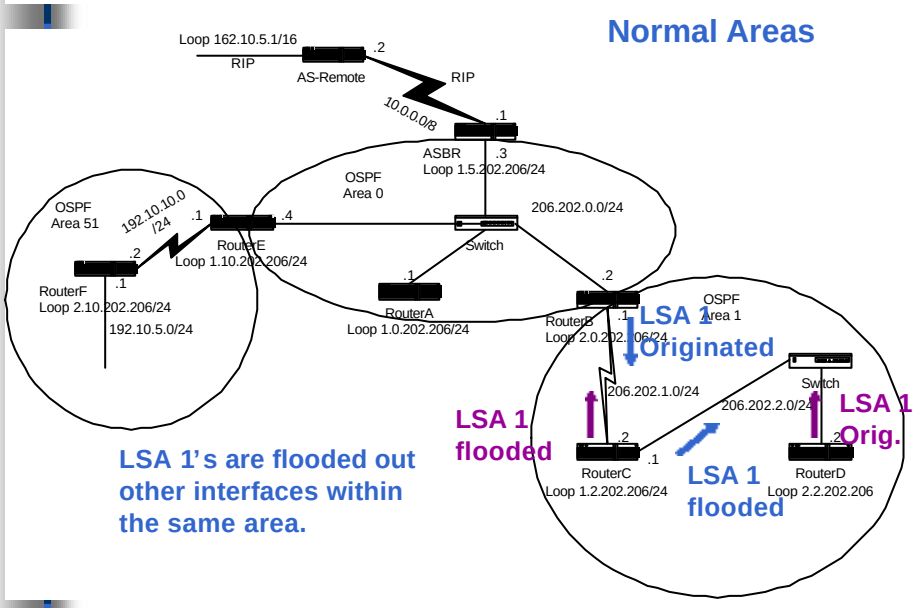

LSA 1's are flooded out other interfaces within the same area.

6

LSA 2 Network LSA

Multi Area OSPF

Normal Areas

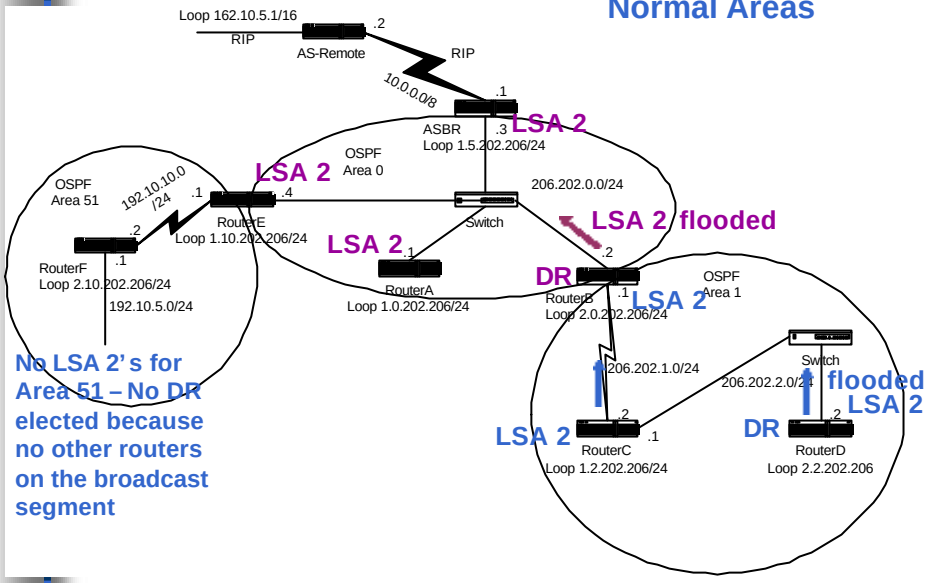

No LSA 2's for Area 51 – No DR elected because no other routers on the broadcast segment

LSA 3 Summary LSA

Multi Area OSPF

Normal Areas

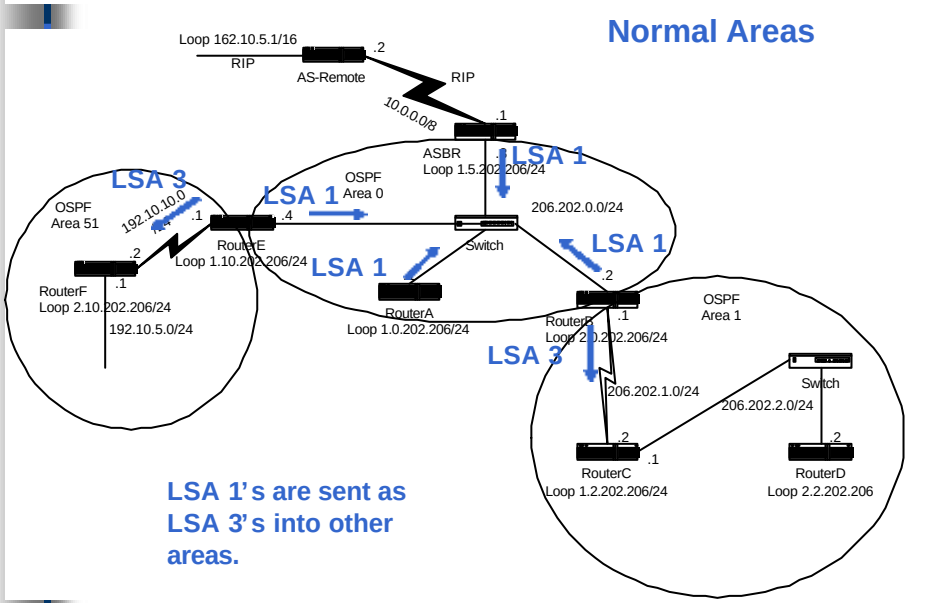

LSA 1's are sent as LSA 3's into other areas.

LSA 3 Summary LSA

Multi Area OSPF

Normal Areas

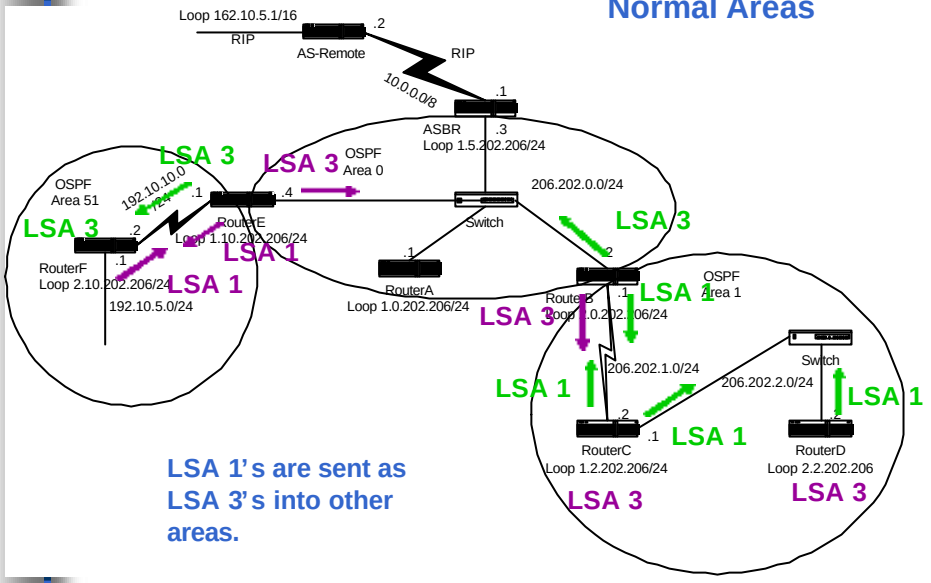

LSA 1's are sent as LSA 3's into other areas.

9

LSA 4 ASBR Summary LSA

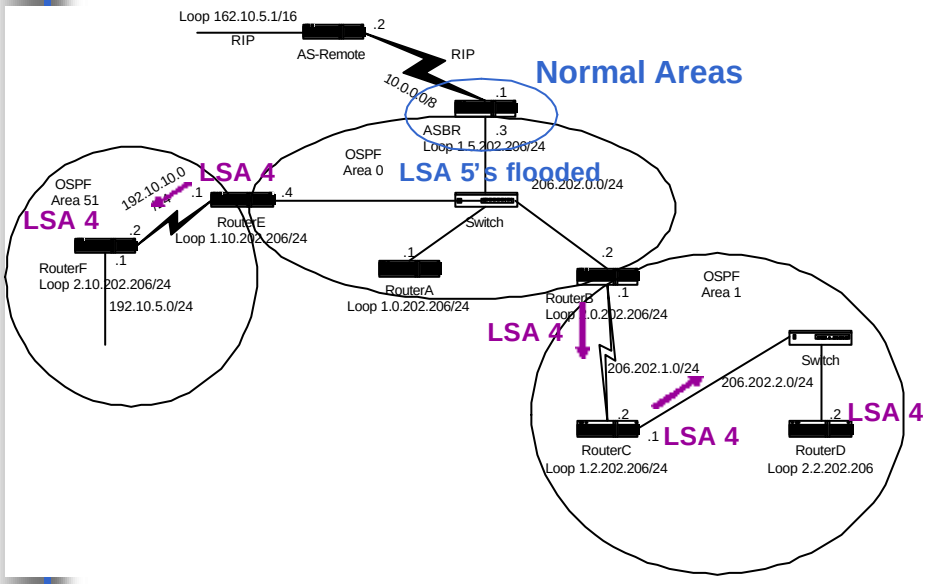

10

LSA 5 AS External LSA

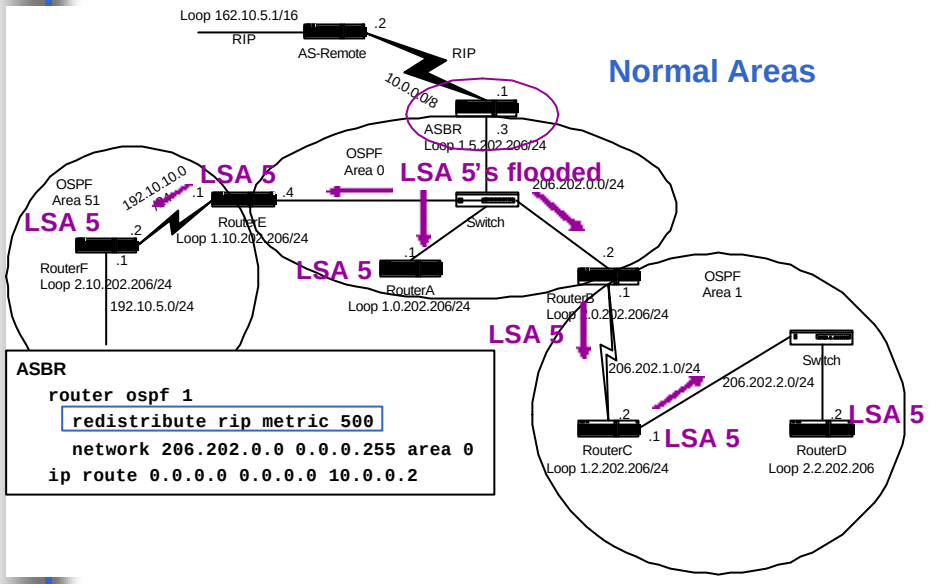

ASBR

```
router ospf 1
 redistribute rip metric 500
  network 206.202.0.0 0.0.0.255 area 0
ip route 0.0.0.0 0.0.0.0 10.0.0.2
```

11

Stub Areas

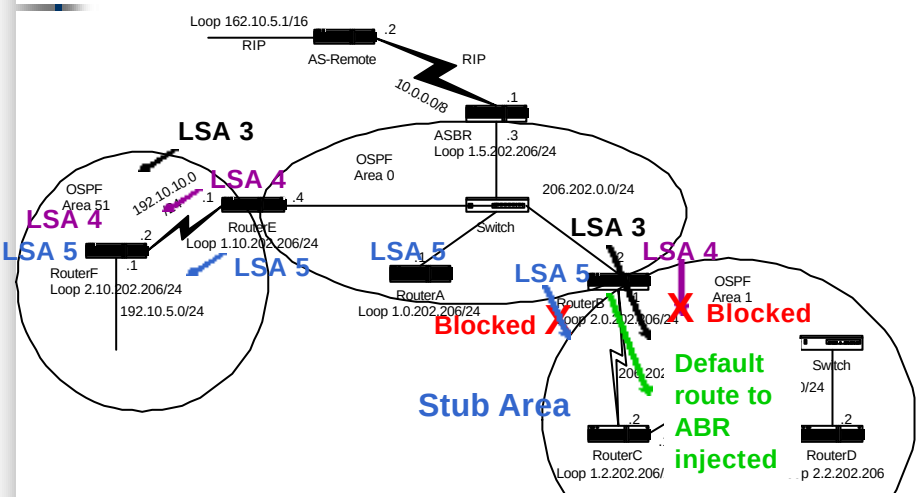

- All routers in the area must be configured as "stub" including the ABR:

```
router ospf 1
   network 206.202.0.0 0.0.0.255 area 0
   network 206.202.1.0 0.0.0.255 area 1
   area 1 stub
```

Totally Stubby Areas

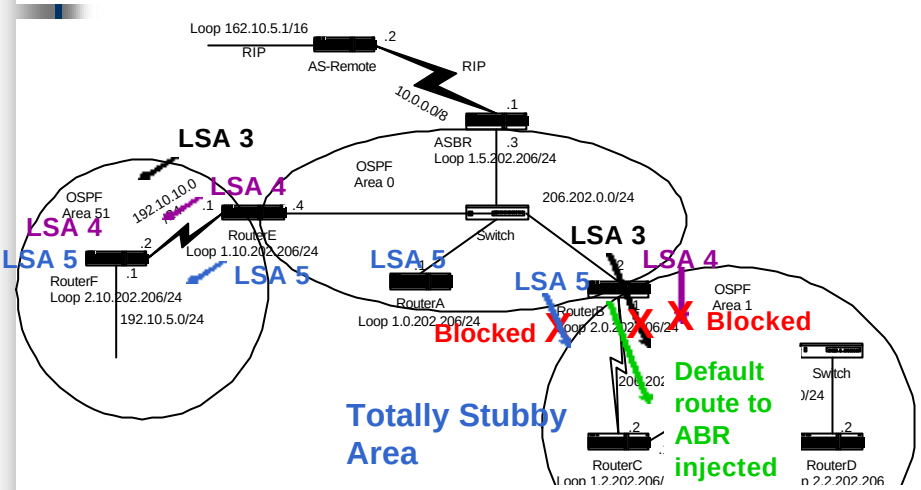

- All routers in the area must be configured as "stub" except the ABR "stub no summary":

```
router ospf 1
   network 206.202.0.0 0.0.0.255 area 0
   network 206.202.1.0 0.0.0.255 area 1
   area 1 stub no-summary
```

NSSA Example

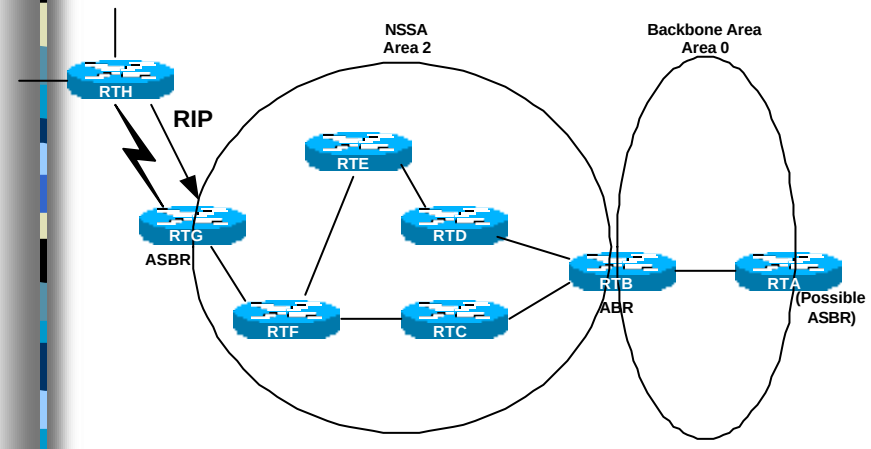

14

NSSA

- Relatively new, standards based OSPF enhancement, RFC 1587.
- NSSA allows an area to remain a stub area, but carry external routing information (Type 7 LSAs) from its stubby end back towards the OSPF backbone.
- ASBR in NSSA injects external routing information into the backbone and the NSSA area, but rejects external routing information coming from the ABR.
- The ABR does not inject a default route into the NSSA.
  - This is true for a NSSA Stub, but a default route is injected for a NSSA Totally Stubby area.
- Note: RFC 1587, " A default route must not be injected into the NSSA as a summary (type-3) LSA as in the stub area case."
- What???
- Following scenario is only example of how NSSA works. For the purposes of learning about NSSAs, don' t get hung up on the why' s and what if s.

15

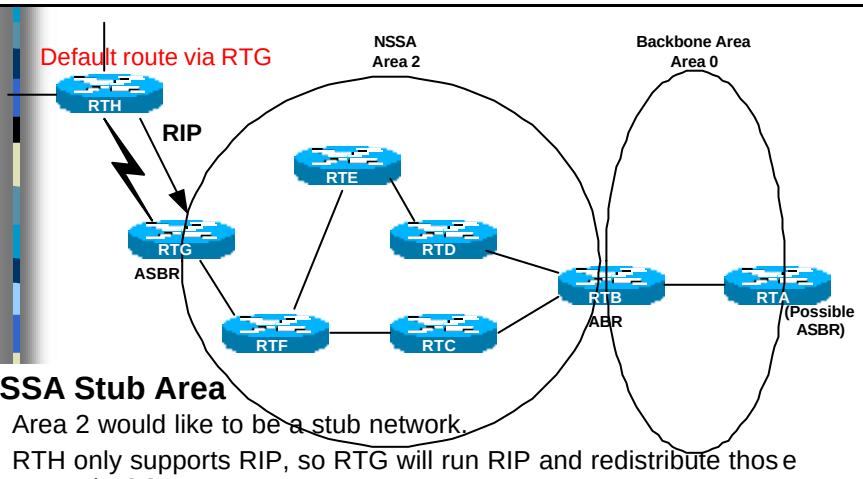

NSSA Stub Area

- Area 2 would like to be a stub network.
- RTH only supports RIP, so RTG will run RIP and redistribute those routes in OSPF.
- Unfortunately, this makes the area 2 router, RTG, an ASBR and therefore area 2 can no longer be a stub area.
- RTH does not need to learn routes from OSPF, a default route to RTG is all it needs.
- But all OSPF routers must know about the networks attached to the RIP router, RTH, to route packets to them.

.6

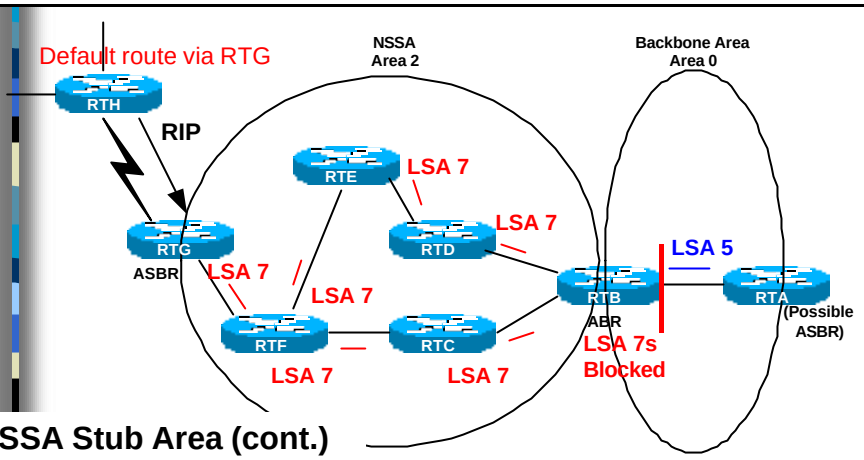

NSSA Stub Area (cont.)

- NSSA allow external routes to be advertised into the OSPF AS while retaining the characteristics of a stub area to the rest of the OSPF AS.
- ASBR RTG will originate Type-7 LSAs to advertise the external destinations.
- These LSA 7s are flooded through the NSSA but are blocked by the NSSA ABR.
- The NSSA ABR translates LSA 7s into LSA 5s and flood other areas.

LSA Types (con' t)

Type 7 LSA NSSA External Link Entry

- Originated by an ASBR connected to an NSSA.
- Type 7 messages can be flooded throughout NSSAs and translated into LSA Type 5 messages by ABRs.
- Routes learned via Type-7 LSAs are denoted by either a default "N1" or an "N2" in the routing table. (Relative to E1 and E2).

18

NSSA Generic

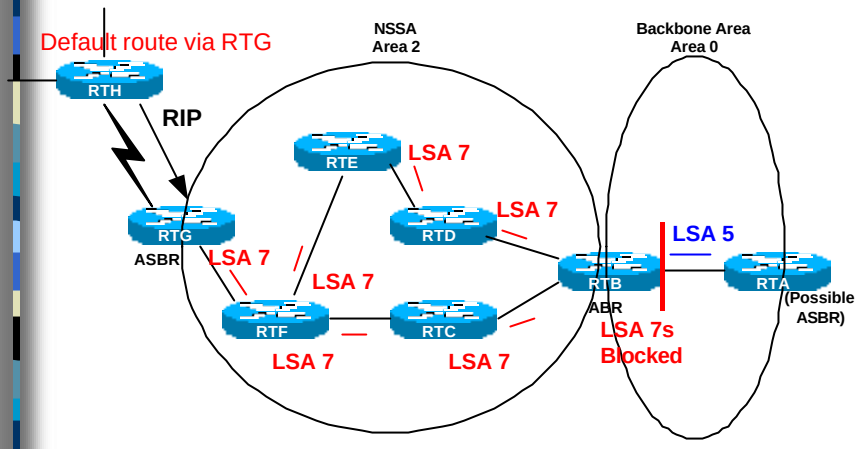

Configuring NSSA Stub Area

Configured for all routers in Area 2:

```
router ospf 1
  network 172.16.2.0 0.0.0.255 area 2
  area 2 nssa
```

19

NSSA Stub and NSSA Totally Stubby

- There are two flavors in NSSA:
  - stub
  - totally stubby
- Area 2 routers may or may not receive Inter-area routes from RTA, depending upon NSSA configuration
- NSSA areas have take on the same characteristics as stub and totally stubby areas, along with the characteristics of NSSA areas.

NSSA stub areas:

- NSSAs that block type 4 and 5, but allow type 3.
- To make a stub area into an NSSA, use the following command under the OSPF configuration.
- This command must be configured on all routers in area 2.

```
router ospf 1
   area 2 nssa
```

20

NSSA Stub Areas

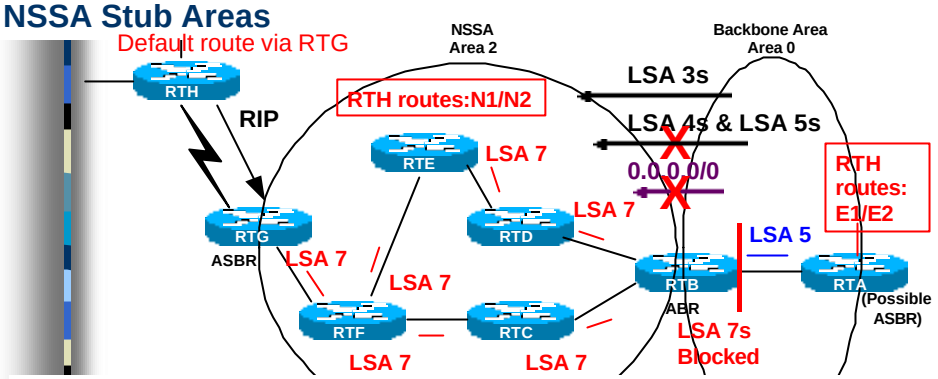

NSSA Stub Area Routing Tables:

- **RTG:** Area 2 routes, Area 0 routes (IA), RTH RIP routes
  - **No** 0.0.0.0/0 (IA) route from RTB (ABR), despite documentation
- **Area 2 Internal Routers:** Area 2 routes, RTH routes (N1/N2), Area 0 routes (IA)
  - **No** 0.0.0.0/0 (IA) route from RTB (ABR), despite documentation
- **RTB:** Area 2 routes, Area 0 routes, RTH routes (N1/N2), External routes if redistributed from RTA ASBR (E1/E2)
- **RTA:** Area 0 routes, Area 2 routes, RTH routes (E1/E2), External routes if redistributed from RTA (E1/E2)
- **Note**: Area 2 routers may or may not receive E1/E2 routes from RTA, depending upon NSSA configuration (next).

NSSA Stub Areas

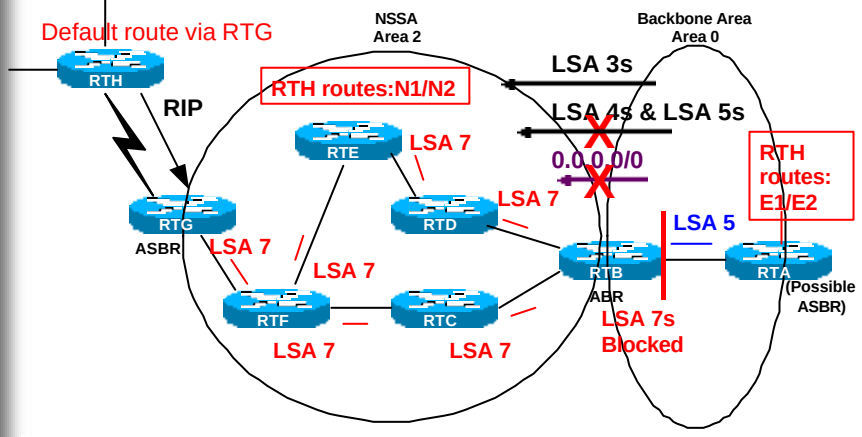

Area 2 routers:

```
router ospf 1
     network 172.16.2.0 0.0.0.255 area 2
     area 2 nssa
```

22

NSSA Totally Stubby Area

- **NSSA totally stub areas**: Allow only summary default routes and filters everything else.
- To configure an **NSSA totally stub area**, use the following command under the OSPF configuration on the **NSSA ABR**:

```
router ospf 1
   area 2 nssa no-summary
```

- Configure this command on **NSSA ABRs only**.
- All other routers in area 2 (internal area 2 routers):

```
router ospf 1
   area 2 nssa
```

- After defining the NSSA totally stub area, area 2 has the following characteristics (in addition to the above NSSA characteristics):
  - No type 3 or 4 summary LSAs are allowed in area 2. This means no inter-area routes are allowed in area 2.
  - A default route is injected into the NSSA totally stub area as a type 3 summary LSA by the ABR.

23

NSSA Totally Stubby Areas

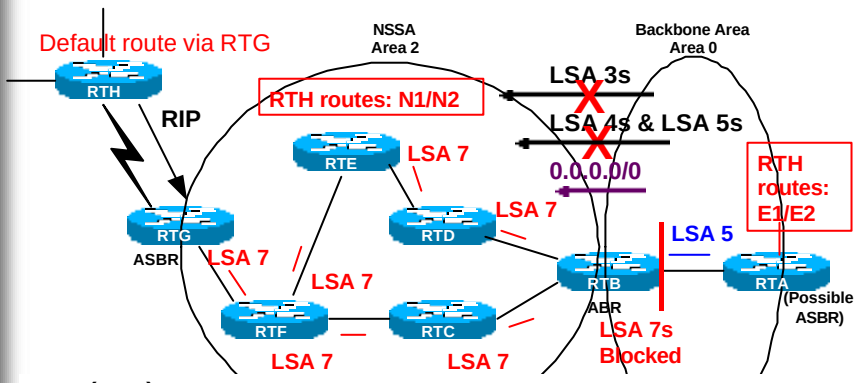


```
RTB (ABR):
router ospf 1
        network 172.16.1.0 0.0.0.255 area 0
        network 172.16.2.0 0.0.0.255 area 2 ...
        area 2 nssa no-summary
Area 2 routers:
router ospf 1
        network 172.16.2.0 0.0.0.255 area 2
        area 2 nssa
```

24

NSSA Totally Stubby Areas

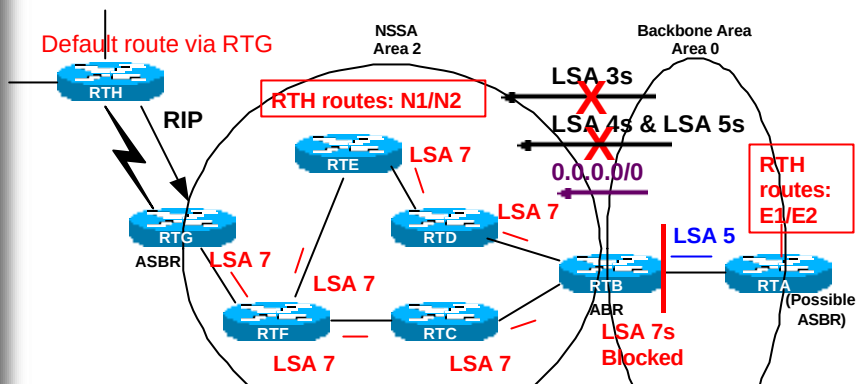

NSSA Totally Stubby Area Routing Tables:

- **RTG:** Area 2 routes, RTH RIP routes, 0.0.0.0/0 (IA) route from RTB (ABR)
  - Totally Stubby: **No** Area 0 routes or external routes from RTA
- **Area 2 Internal Routers:** Area 2 routes, RTH routes (N1/N2), 0.0.0.0/0 (IA) route from RTB (ABR)
  - Totally Stubby: **No** Area 0 routes or external routes from RTA
- **RTB:** Area 2 routes, Area 0 routes, RTH routes (N1/N2), External routes from RTA ASBR (E1/E2) if redistributed by ASBR
- **RTA:** Area 0 routes, Area 2 routes, RTH routes (E1/E2), External routes (E1/E2)

More on NSSA

- Examples
- See NSSA document on my web site for more info.

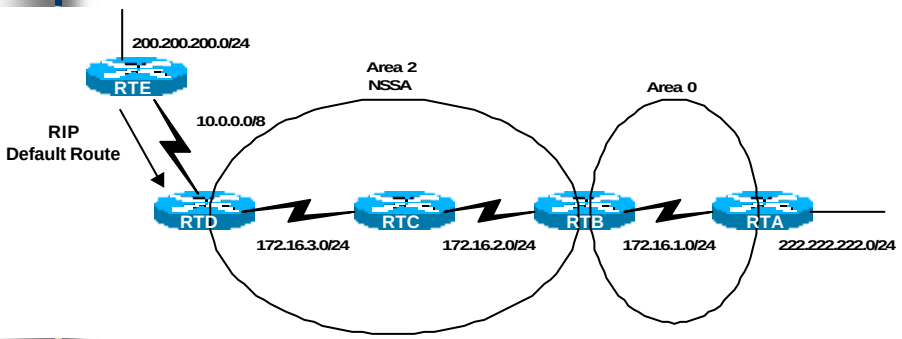

Virtual Links

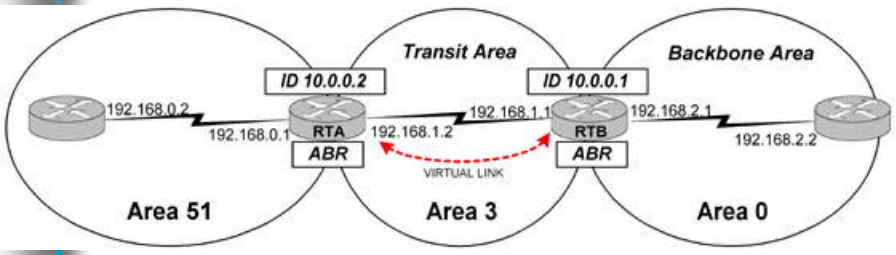

27

Virtual Links

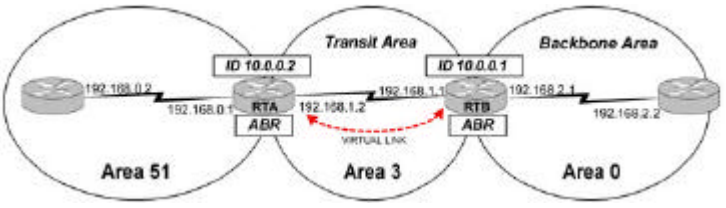

- All areas in an OSPF autonomous system must be physically connected to the backbone area (area 0).
- In some cases where this is not possible, you can use a virtual link to connect to the backbone through a non-backbone area.
- As mentioned above, you can also use virtual links to connect two parts of a partitioned backbone through a non-backbone area.
- The area through which you configure the virtual link, known as a transit area, must have full routing information.
- Must be configured between two ABRs.
- The transit area **cannot** be a stub area.

28

Virtual Links

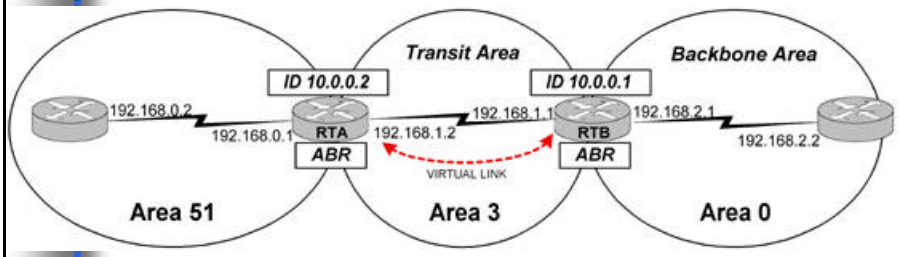

- A virtual link has the following two requirements:
  - It must be established between two routers that share a common area and are both ABRs.
  - One of these two routers must be connected to the backbone.
- Doyle, "should be used only as a temporary fix to an unavoidable topology problem."

29

Virtual Links

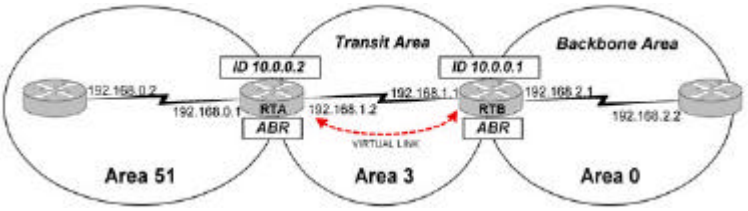

The command to configure a virtual link is as follows:

`area <area-id> virtual-link <remote-router-id>`

```
RTA(config)#router ospf 1
RTA(config-router)#network 192.168.0.0 0.0.0.255 area 51
RTA(config-router)#network 192.168.1.0 0.0.0.255 area 3
RTA(config-router)#area 3 virtual-link 10.0.0.1
...

RTB(config)#router ospf 1
RTB(config-router)#network 192.168.1.0 0.0.0.255 area 3
RTB(config-router)#network 192.168.2.0 0.0.0.255 area 0
RTB(config-router)#area 3 virtual-link 10.0.0.2
```

30

Virtual Links

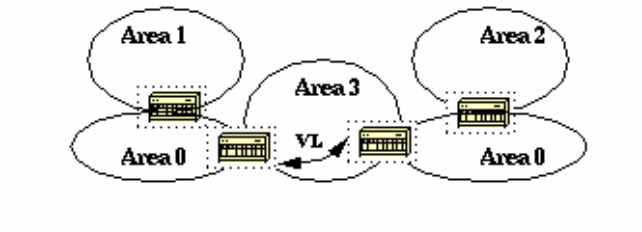

- OSPF allows for linking discontinuous parts of the backbone using a virtual link.
- In some cases, different area 0s need to be linked together. This can occur if, for example, a company is trying to merge two separate OSPF networks into one network with a common area 0.
- In other instances, virtual-links are added for redundancy in case some router failure causes the backbone to be split into two. (CCO)
- Whatever the reason may be, a virtual link can be configured between separate ABRs that touch area 0 from each side and having a common area.

31

Route Summarization

Inter-Area Route Summarization - Area Range

- By default ABRs do not summarize routes between areas.
- Route summarization is the consolidation of advertised addresses.
- This feature causes a single summary route to be advertised to other areas by an ABR.
- In OSPF, an ABR will advertise networks in one area into another area.
- If the network numbers in an area are assigned in a way such that they are contiguous, you can configure the ABR to advertise a summary route that covers all the individual networks within the area that fall into the specified range.

On the ABR (Summarizes routes before injecting them into different area)

```
Router(config-router)# area area-id range
  network-address subnet-mask
```

- ***area-id*** **- Identifier of the area about which routes are to be summarized. (From area)**

32

Route Summarization

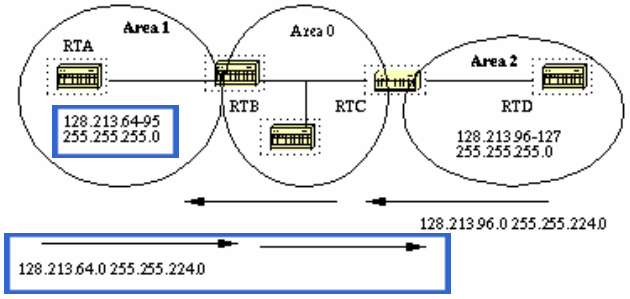

- RTB is summarizing the range of subnets from 128.213.64.0 to 128.213.95.0 into one range: 128.213.64.0 255.255.224.0.
- This is achieved by masking the first three left most bits of 64 using a mask of 255.255.224.0.

128.213.64.0/24 - 01000000

128.213.95.0/24 – 01011111

128.213.64.0/19 - 01000000

33

Route Summarization

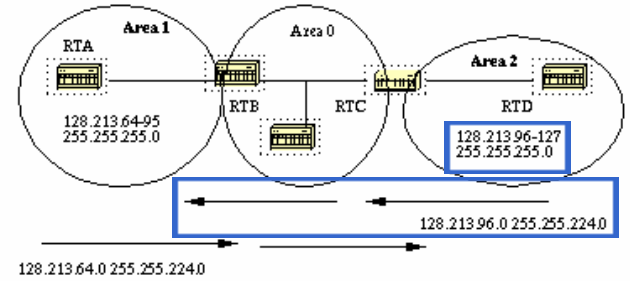

- In the same way, RTC is generating the summary address 128.213.96.0 255.255.224.0 into the backbone.
- Note that this summarization was successful because we have two distinct ranges of subnets, 64-95 and 96-127.

128.213.96.0/24 - 01100000

128.213.127.0/24 – 01111111

128.213.96.0/19 - 01100000

34

Route Summarization

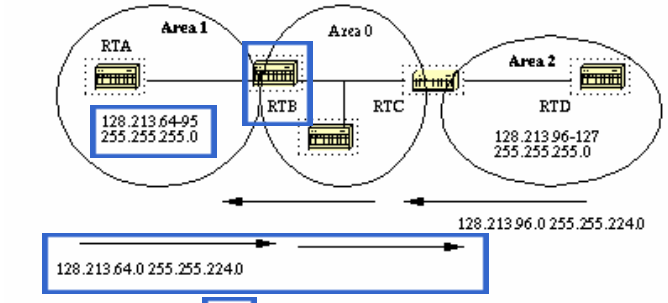

128.213.64.0/24 - 01000000

128.213.95.0/24 – 01011111

128.213.64.0/19 - 01000000

RTB

```
router ospf 100
 area 1 range 128.213.64.0 255.255.224.0
```

35

Route Summarization

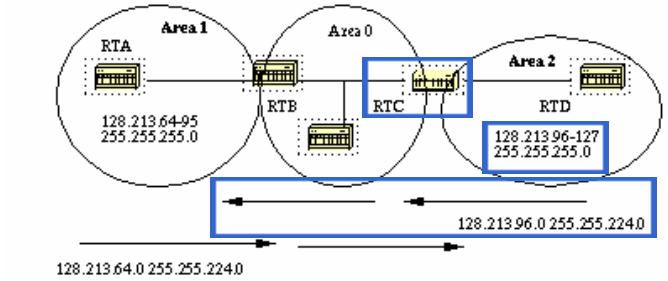

128.213.96.0/24 - 01100000

128.213.127.0/24 – 01111111

128.213.96.0/19 - 01100000

RTC

```
router ospf 100
 area 2 range 128.213.96.0 255.255.224.0
```

36

Route Summarization

External Route Summarization - summary-address

- When redistributing routes from other protocols into OSPF (later), each route is advertised individually in an external link state advertisement (LSA).
- However, you can configure the Cisco IOS software to advertise a single route for all the redistributed routes that are covered by a specified network address and mask.
- Doing so helps decrease the size of the OSPF link state database.

On the ASBR only (Summarizes external routes before injecting them into the OSPF domain.)

```
Router(config-router)# summary-address network-address subnet-mask
```

37

Route Summarization

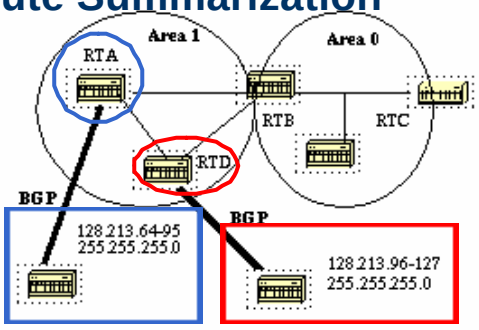

RTA

```
router ospf 100
summary-address 128.213.64.0 255.255.224.0
redistribute bgp 50 metric 1000 subnets (later)
```

RTD

```
router ospf 100
summary-address 128.213.96.0 255.255.224.0
redistribute bgp 20 metric 1000 subnets (later)
```

38

Injecting Default Routes into OSPF

- By default, 0.0.0.0/0 route is not propagated from the ASBR to other routers.
- An autonomous system boundary router (ASBR) can be forced to generate a default route into the OSPF domain.
- As discussed earlier, a router becomes an ASBR whenever routes are redistributed into an OSPF domain.
- However, an ASBR does not, by default, generate a default route into the OSPF routing domain.

39

How Does OSPF Generate Default Routes?

The way that OSPF generates default routes (0.0.0.0) varies depending on the type of area the default route is being injected into.

Stub and Totally Stubby Areas

- For stub and totally stubby areas, the area border router (ABR) to the stub area generates a summary link-state advertisement (LSA) with the link-state ID 0.0.0.0.
- This is true even if the ABR doesn't have a default route.
- In this scenario, you don't need to use the **default-information originate** command.

40

Stub Areas

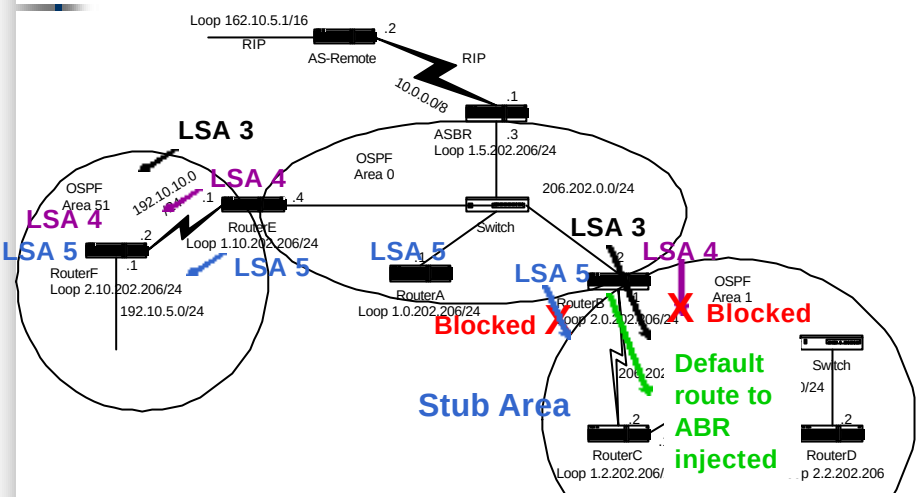

- All routers in the area must be configured as "stub" including the ABR:

```
router ospf 1
   network 206.202.0.0 0.0.0.255 area 0
   network 206.202.1.0 0.0.0.255 area 1
   area 1 stub
```

Totally Stubby Areas

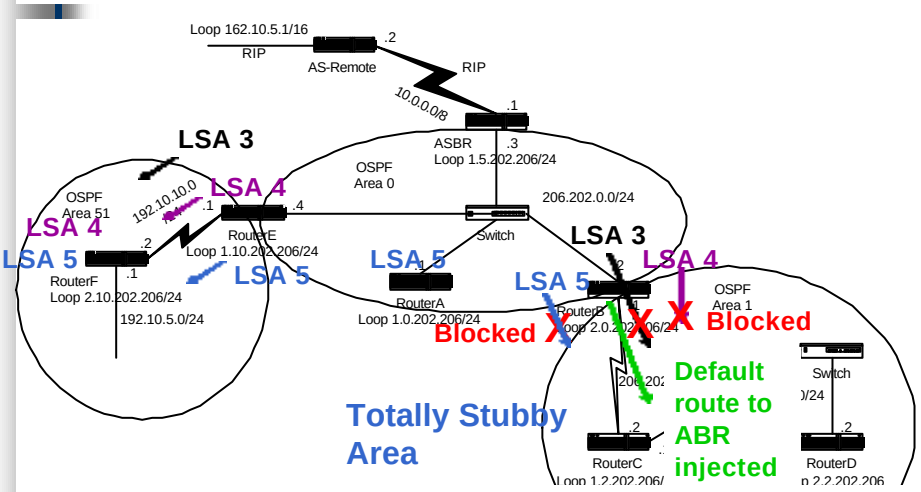

- All routers in the area must be configured as "stub" except the ABR "stub no summary":

```
router ospf 1
   network 206.202.0.0 0.0.0.255 area 0
   network 206.202.1.0 0.0.0.255 area 1
   area 1 stub no-summary
```

How Does OSPF Generate Default Routes?

Normal Areas

- By default, in normal areas routers don't generate default routes.
- To have an OSPF router generate a default route, use the **default-information originate** command.
- This generates an external type-2 link with link-state ID 0.0.0.0 and network mask 0.0.0.0.
- This command should only be used on the ASBR.
  - Some documentation states this command works only on an ASBR while other documentation states this command turns a router into an ASBR.

43

Injecting Default Routes into OSPF

To have OSPF generate a default route use the following:

```
router ospf 10
default-information originate [always] [metric
  metric-value] [metric-type type-value] [route-
  map map-name]
```

44

There are two ways to generate a default.

1) default-information originate

- If the ASBR already has the default route (ip route 0.0.0.0 0.0.0.0), you can advertise 0.0.0.0 into the area.

2) default-information originate always

- If the ASBR doesn't have the route (ip route 0.0.0.0 0.0.0.0), you can add the keyword **always** to the **default-information originate** command, and then advertise 0.0.0.0.
- You should be careful when using the **always** keyword. If your router advertises a default (0.0.0.0) inside the domain and does not have a default itself or a path to reach the destinations, routing will be broken.

45

Injecting Default Routes into OSPF

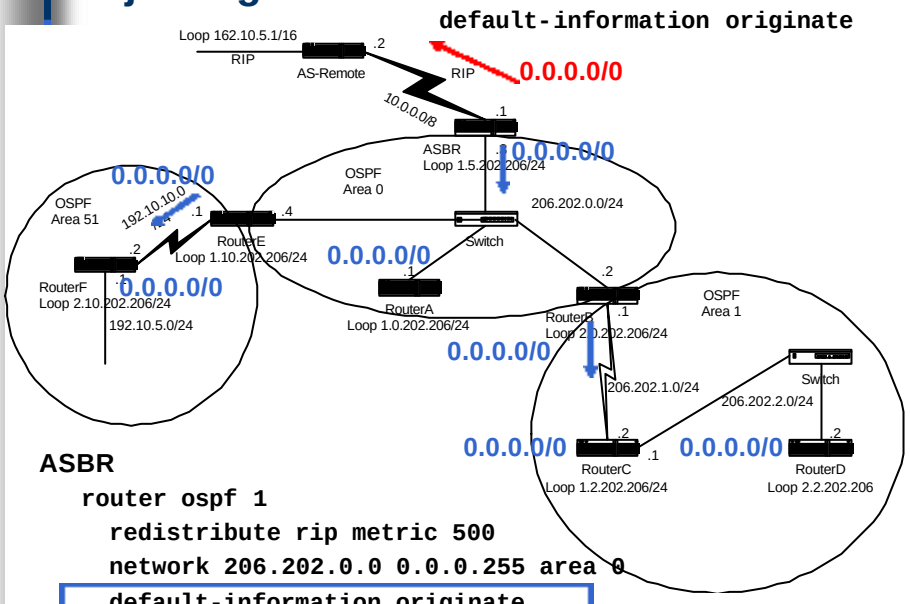

ASBR

```
router ospf 1
  redistribute rip metric 500
  network 206.202.0.0 0.0.0.255 area 0
  default-information originate
ip route 0.0.0.0 0.0.0.0 10.0.0.2
```

46

Injecting Default Routes into OSPF

`default-information originate always`

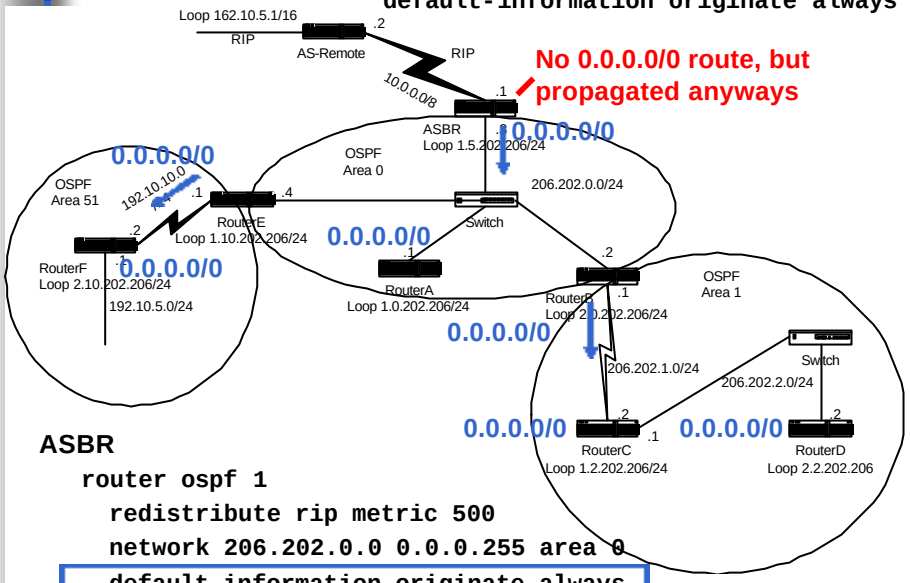

ASBR

```
router ospf 1
  redistribute rip metric 500
  network 206.202.0.0 0.0.0.255 area 0
  default-information originate always
```


Redistributing External Routes

E1 vs. E2 External Routes

- External routes fall under two categories, external type 1 and external type 2.
- The difference between the two is in the way the cost (metric) of the route is being calculated.
- A **type 1 (E1)** cost is the addition of the external cost and the internal cost used to reach that route.
- The cost of a **type 2 (E2)** route is always the external cost, irrespective of the interior cost to reach that route.
- Type 2 (E2) is the default!

Redistributing External Routes

```
router ospf 1
  redistribute routing-protocol metric-type [1|2]
```

- **metric-type 1 -** A **type 1** cost is the addition of the external cost and the internal cost used to reach that route.

  `redistribute rip metric-type 1`

- **metric-type 2 -** The cost of a **type 2** route is always the external cost, irrespective of the interior cost to reach that route.

  `redistribute rip metric-type 2`

- We will look at this command, along with internal/external costs, later in the chapter discussion route redistribution.

Redistributing External Routes

`metric-type 1`

RIP routes redistributed with a metric (cost) of 500 plus the outgoing cost of the interface and a metric-type 1

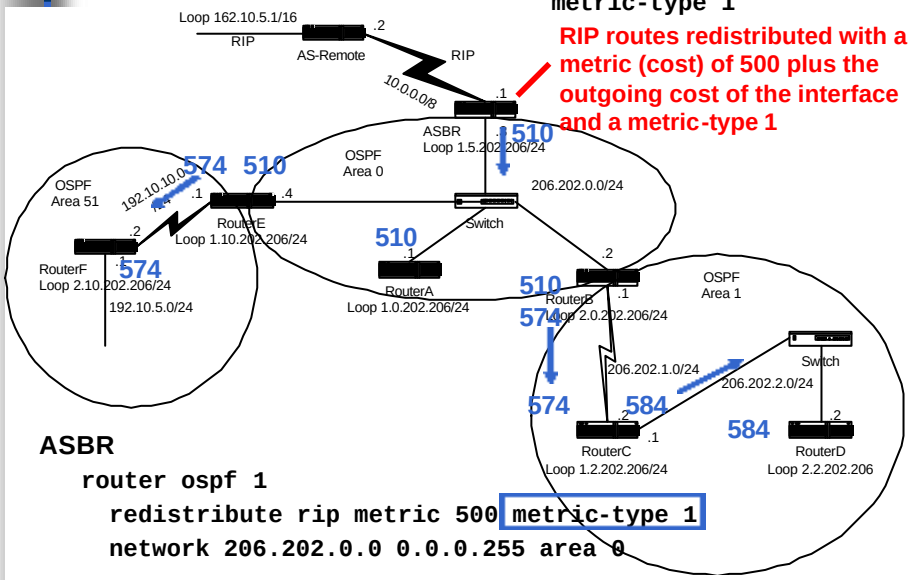

ASBR

```
router ospf 1
  redistribute rip metric 500 metric-type 1
  network 206.202.0.0 0.0.0.255 area 0
```

Redistributing External Routes

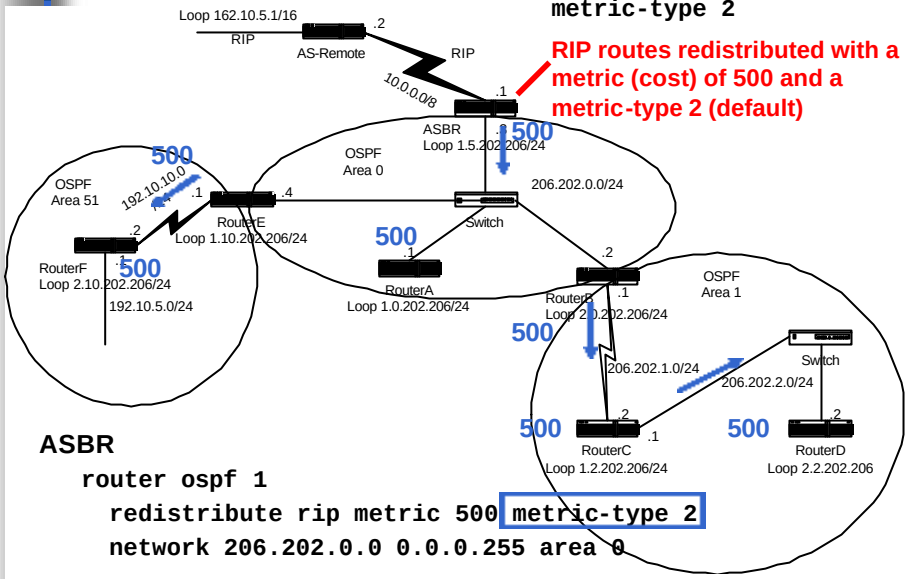

metric-type 2

RIP routes redistributed with a metric (cost) of 500 and a metric-type 2 (default)

ASBR

```
router ospf 1
  redistribute rip metric 500 metric-type 2
  network 206.202.0.0 0.0.0.255 area 0
```

51

Redistributing External Routes

- **Instructors**: Due to a temporary problem I am having with Visio, I was unable to create the diagrams for this example. Look for a newer version of this presentation to be out shortly or contact me at graziani@cabrillo.cc.ca.us

So when should you redistribute a Type-1 (E1) External route?

- If there is more than one ABR for the area and the area is not a stub or totally stubby area.
  - In this case one of the ABRs may provide a shorter path for certain non-area 0 internal routers, than other ABRs.
  - E1 routes will include all internal costs from the internal router to the ABR and to the ASBR, allowing each router to choose which ABR provides the shorter path.
- Multiple ASBRs redistributing the same networks.
  - In this case the routers' cost to each ASBR can be used to choose the shortest path to the destination.

52

Know your outputs

- show ip route
- show ip ospf
- show ip ospf neighbor
- show ip ospf border-router
- show ip database
- show ip interface

53

show ip route

```
RouterC#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B -
   BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate
   default
       U - per-user static route, o - ODR

Gateway of last resort is 206.202.1.1 to network 0.0.0.0

C    206.202.2.0/24 is directly connected, Ethernet0
     1.0.0.0/24 is subnetted, 1 subnets
C       1.2.202.0 is directly connected, Loopback0
O IA 206.202.0.0/24 [110/74] via 206.202.1.1, 00:09:34, Serial0
C    206.202.1.0/24 is directly connected, Serial0
O*IA 0.0.0.0/0 [110/65] via 206.202.1.1, 00:09:34, Serial0
```

54

show ip ospf

```
RouterC#show ip ospf
 Routing Process "ospf 1" with ID 1.2.202.206
 Supports only single TOS(TOS0) routes
 SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
 Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
 Number of external LSA 0. Checksum Sum 0x0
 Number of DCbitless external LSA 0
 Number of DoNotAge external LSA 0
 Number of areas in this router is 1. 0 normal 1 stub 0 nssa
    Area 1
        Number of interfaces in this area is 2
        It is a stub area
        Area has no authentication
        SPF algorithm executed 14 times
        Area ranges are
        Number of LSA 6. Checksum Sum 0x1F204
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
```

55

show ip ospf neighbor

- Displays a list of neighbors and their link state status

```
RouterC#show ip ospf neig

Neighbor ID     Pri   State      Dead Time   Address         Interface
2.2.202.206       1   FULL/DR    00:00:33    206.202.2.2     Ethernet0
2.0.202.206       1   FULL/  -   00:00:32    206.202.1.1     Serial0
```

56

show ip ospf border-router

- To display the internal OSPF routing table entries to an Area Border Router (ABR) and Autonomous System Boundary Router (ASBR), use the **show ip ospf border-routers** privileged EXEC command.
- LSA 4's (routes to ASBRs) are not installed in the main IP routing table but in the special internal OSPF routing table.

```
Router# show ip ospf border-routers
OSPF Process 109 internal Routing Table

Destination      Next Hop        Cost   Type     Rte Type Area         SPF No

160.89.97.53     144.144.1.53    10     ABR      INTRA    0.0.0.3        3
160.89.103.51    160.89.96.51    10     ABR      INTRA    0.0.0.3        3
160.89.103.52    160.89.96.51    20     ASBR     INTER    0.0.0.3        3
160.89.103.52    144.144.1.53    22     ASBR     INTER    0.0.0.3        3
```

Destination - Router ID of the destination.
Next Hop - Next hop toward the destination.
Cost - Cost of using this route.
Type - The router type of the destination; it is either an ABR or ASBR or both.
Rte Type - The type of this route; it is either an intra-area or interarea route.
Area - The area ID of the area from which this route is learned.
SPF No - The internal number of the shortest path first (SPF) calculation that installs this route.

show ip ospf database

- Displays a summary of the topological, link-state database

```
RouterC#show ip ospf data
       OSPF Router with ID (1.2.202.206) (Process ID 1)
                Router Link States (Area 1)
Link ID         ADV Router      Age         Seq#       Checksum Link
   count
1.2.202.206     1.2.202.206     590         0x80000009 0xB787   3
2.0.202.206     2.0.202.206     591         0x80000004 0x1880   2
2.2.202.206     2.2.202.206     797         0x80000006 0x3B7E   1
               Net Link States (Area 1)
Link ID         ADV Router      Age         Seq#       Checksum
206.202.2.2     2.2.202.206     797         0x80000001 0x1DA9
                Summary Net Link States (Area 1)
Link ID         ADV Router      Age         Seq#       Checksum
0.0.0.0         2.0.202.206     598         0x80000001 0x3B67
206.202.0.0     2.0.202.206     598         0x80000004 0x8D6F
```

58

show ip ospf interface

- Displays OSPF information regarding a specific interface or interfaces
- (next slide)

59

```
RouterC#show ip ospf inter
Ethernet0 is up, line protocol is up
  Internet Address 206.202.2.1/24, Area 1
  Process ID 1, Router ID 1.2.202.206, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 2.2.202.206, Interface address 206.202.2.2
  Backup Designated router (ID) 1.2.202.206, Interface address
    206.202.2.1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:07
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 2.2.202.206  (Designated Router)
  Suppress hello for 0 neighbor(s)
Serial0 is up, line protocol is up
  Internet Address 206.202.1.2/24, Area 1
  Process ID 1, Router ID 1.2.202.206, Network Type POINT_TO_POINT,
    Cost: 64
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:07
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 2.0.202.206
  Suppress hello for 0 neighbor(s)
```

Do you know these?

60

OSPF Extra' s, FAQs, and FYIs

- The following sections contain information to help you understand OSPF.
- This information is not necessarily on the CCNP Advanced Routing Exam.

61

Extra: OSPF over ISDN

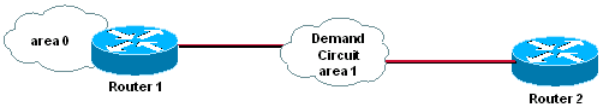

- OSPF Hello traffic can keep an ISDN link up indefinitely.
- By entering the command "**ip ospf demand-circuit**" on one side of a BRI, adjacencies will be formed and:
  - Ongoing OSPF Hellos will be suppressed
  - The DNA (Do-Not-Age) bit is set in the LSA so that this entry is not aged in the router' s LSDB.
    - LSA is not flooded when reaching LSRefresh
    - LSA is not flooded if there is a new version but the contents are the same

`show ip ospf interface bri 0`

- "Run as demand circuit"
- "(Hello Suppressed)"

`show ip ospf neighbor`

- Dead Time: "-"

62

Extra: OSPF over ISDN

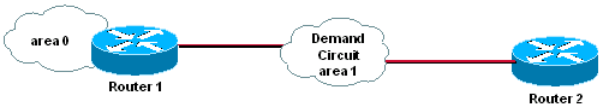

Router1

```
interface BRI1/1
  ip address 192.158.254.13 {/30}
  ip ospf demand-circuit

router ospf 20
  network 192.158.254.0 0.0.0.255 area 0
```

Router2

```
interface BRI1/0
  ip address 192.158.254.14 {/30}

router ospf 20
  network 192.158.254.0 0.0.0.255 area 0
```

Note: You need to configure the demand circuit at one end of the link only. However, if you configure this command on both ends it does not cause any harm.

Suggestion: To reduce the affect of link flaps on the demand circuit, configure the area that contains the demand circuit as totally stub.

- In this case configure Area 1 to be a totally stubby area.
- Summarizing routes on Router 1 can also help if the flapping link is within the summarized range.

63

Extra: OSPF and Load Balancing

- OSPF only supports equal-cost load balancing.
- By default, four equally good routes to the same destination are kept in the routing table for load balancing.
- This can be increased up to six with the **maximum-paths** command.
- The bandwidth and/or ip ospf cost (or in the case of serial links [1.544 Mbps] the lack of) commands can be used to make unequal-cost links look like equal-cost links to OSPF for load balancing.
  - This should be done with caution, as it may burden slower links and/or not make efficient use of faster links.

64

Extra: OSPF and DNS Lookups

- Loopback interfaces simplify the management and troubleshooting of OSPF routing domains by providing predictable Router Ids.
- This can be taken one step further by recording the Router Ids in a Domain Name Service (DNS) database.
- The router can then be configured to consult the server address-to-name mappings, or Reverse DNS lookups, and then display the routers by name instead of by Router ID.

65

Extra: OSPF and DNS Lookups

For example:

```
SanJose2#show ip ospf database
       OSPF Router with ID (1.0.202.206) (Process ID 1)
                Router Link States (Area 0)
Link ID         ADV Router      Age    Seq#   Checksum  Lkcnt
1.0.202.206     SanJose1        69     0x80000005 0xA733   1
1.5.202.206     Baypoint        357    0x80000005 0x8329   1
```

- SanJose2 was configured to perform DNS lookups as follows:

```
ip name-server 206.202.1.10
ip ospf name-lookup
```

- The first command specifies the DNS server.
- The second command enables the OSPF process to perform DNS lookups.
- This can also be used for identifying router interfaces such as SanJose-e0.

66

Extra: IOS 12.01(T) – router-id

`router-id`

- To use a fixed router ID, use the **router-id** router configuration command.
- To force OSPF to use the previous OSPF router ID behavior, use the **no** form of this command.
- Takes precedence over Loopback address

```
router ospf 1
    router-id ip-address
```

67

OSPF FAQs and FYIs

Q: Why are loopbacks advertised as /32 host routes in OSPF?

A: Loopbacks are considered host routes in OSPF, and they're advertised as /32. For more information, see section 9.1 of RFC 2328. In Cisco IOS ® version 11.3T and 12.0, if the **ip ospf network point-to-point** command is configured under loopbacks, then OSPF advertises the loopback subnet as the actual subnet configured on loopbacks.

http://www.cisco.com/warp/public/104/9.html

Q: Can a virtual link cross more than one area.

A: No.

68

OSPF FAQs and FYIs

Q: What happens within OSPF if there is more than one route to a destination? What is the preference of OSPF in choosing a best route?

A: Here is the OSPF preference rules:

- Intra-area routes area always most preferred.
- Inter-area routes are preferred over AS or NSSA external routes.
- AS-external routes and NSSA-external routes are of equal preference. Within these routes, preferences are as follows:
  - External Type-1 routes are always preferred
    - If equal, route-metric (cost) is the tie-breaker
  - External Type-2 routes
    - If equal, route metric and distance to the originating router are used as tie-breakers.
  - If still a tie (Type-1 or Type-2), AS-external (LSA 5) routes are preferred over NSSA external (LSA 7) routes.
- If these rules do not solve the tie, routes are installed as parallel routes.

69

OSPF FAQs and FYIs

OSPF Packet Pacing

- Introduced in Cisco IOS 11.3
- Helps avoid packet drops at the receiving side, caused by uncontrolled bursts of link-state updates.
- The receiving router may not be able to queue and process all of the packets so some packets are dropped.
- To make matters worse, when the sending router does not receive LSAcks for all of the LSAs sent, so retransmits along with other LSAs needed to be sent.
- Currently Cisco IOS Packet Pacing, every 33 milliseconds (non-configurable) the router builds a link-state update and sends it to its neighbors.
- The next group of LSAs is transmitted after another 33 milliseconds.
- This speeds up convergence and decreases the length of the transition period.

70

OSPF FAQs and FYIs

OSPF Group Pacing

- Introduced in Cisco IOS 11.3
- Every LSA is aged whiled stored in the LSDB.
- ALL LSAs are aged independently of one another.
- When an LSA reaches LSRefreshTime (30 minutes) the router that originated the it floods the LSA.
- When an LSA reaches MaxAge (60 minutes) the router floods the LSA, even if it did not originate the LSA.
- If a router has a lot of LSAs, maintaining a separate timer can be expensive.
- With Cisco OSPF Group Pacing, LSAs are collected into groups by their ages, with ages within 4 minutes by default (can be configured).
- The router maintains timers for LSA groups instead of individual LSAs.
- This is used for all LSA operations including LSA aging and LSA refreshing.

71

OSPF FAQs and FYIs – know this one!

Cisco SPF Scheduling (Review)

- SPF algorithm is CPU intensive and takes some time depending upon the size of the area (coming next week), the number of routers, the size of the link state database.
- A flapping link can cause an OSPF router to keep on recomputing a new routing table, and never converge.
- To minimize this problem:
  - SPF calculations are **delayed by 5 seconds** after **receiving an LSU** (Link State Update)
  - **Delay between consecutive SPF calculations** is **10 seconds**
- You can configure the delay time between when OSPF receives a topology change and when it starts a shortest path first (SPF) calculation (*spf-delay*).
- You can also configure the hold time between two consecutive SPF calculations (*spf-holdtime*).

```
Router(config-router)#timers spf spf-delay spf-holdtime
```

72

OSPF Design Issues

- **Number of Routers per Area**
- **Number of Neighbors**
- **Number of Areas per ABR**
- **Full Mesh vs. Partial Mesh**
- **Memory Issues**

73

OSPF Design Issues - FYI

- The following information is taken from Cisco CCO.
- http://www.cisco.com/warp/public/104/3.html
- The OSPF RFC (1583) did not specify any guidelines for the number of routers in an area or number the of neighbors per segment or what is the best way to architect a network.
- Different people have different approaches to designing OSPF networks.
- The important thing to remember is that any protocol can fail under pressure.
- The idea is not to challenge the protocol but rather to work with it in order to get the best behavior.
- The following are a list of things to consider.
  - **Number of Routers per Area**
  - **Number of Neighbors**
  - **Number of Areas per ABR**
  - **Full Mesh vs. Partial Mesh**
  - **Memory Issues**

74

OSPF Design Issues

Number of Routers per Area

The maximum number of routers per area depends on several factors, including the following:

- What kind of area do you have?
- What kind of CPU power do you have in that area?
- What kind of media?
- Will you be running OSPF in NBMA mode?
- Is your NBMA network meshed?
- Do you have a lot of external LSAs in the network?
- Are other areas well summarized?
- For this reason, it's difficult to specify a maximum number of routers per area.

75

OSPF Design Issues

Number of Neighbors

- The number of routers connected to the same LAN is also important.
- Each LAN has a DR and BDR that build adjacencies with all other routers.
- The fewer neighbors that exist on the LAN, the smaller the number of adjacencies a DR or BDR have to build.
- That depends on how much power your router has. You could always change the OSPF priority to select your DR.
- Also if possible, try to avoid having the same router be the DR on more than one segment.
- If DR selection is based on the highest RID, then one router could accidentally become a DR over all segments it is connected to.
- This router would be doing extra effort while other routers are idle.

76

OSPF Design Issues

Number of Areas per ABR

- ABRs will keep a copy of the database for all areas they service.
- If a router is connected to five areas for example, it will have to keep a list of five different databases.
- The number of areas per ABR is a number that is dependent on many factors, including type of area (normal, stub, NSSA), ABR CPU power, number of routes per area, and number of external routes per area.
- For this reason, a specific number of areas per ABR cannot be recommended.
- Of course, it's better not to overload an ABR when you can always spread the areas over other routers.
- The following diagram shows the difference between one ABR holding five different databases (including area 0) and two ABRs holding three databases each.
- Again, these are just guidelines, the more areas you configure per ABR the lower performance you get. In some cases, the lower performance can be tolerated.

77

OSPF Design Issues

Full Mesh vs. Partial Mesh

- Non Broadcast Multi-Access (NBMA) clouds such as Frame Relay or X.25, are always a challenge.
- The combination of low bandwidth and too many link-states is a recipe for problems.
- A partial mesh topology has proven to behave much better than a full mesh.
- A carefully laid out point-to-point or point-to-multipoint network works much better than multipoint networks that have to deal with DR issues.

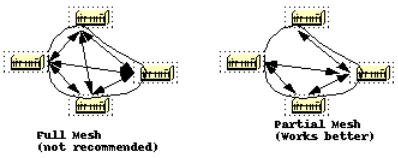

78

OSPF Design Issues

Memory Issues

- It is not easy to figure out the memory needed for a particular OSPF configuration. Memory issues usually come up when too many external routes are injected in the OSPF domain. A backbone area with 40 routers and a default route to the outside world would have less memory issues compared with a backbone area with 4 routers and 33,000 external routes injected into OSPF.
- Memory could also be conserved by using a good OSPF design. Summarization at the area border routers and use of stub areas could further minimize the number of routes exchanged.
- The total memory used by OSPF is the sum of the memory used in the routing table (**show ip route summary**) and the memory used in the link-state database. The following numbers are a rule of thumb estimate. Each entry in the routing table will consume between approximately 200 and 280 bytes plus 44 bytes per extra path. Each LSA will consume a 100 byte overhead plus the size of the actual link state advertisement, possibly another 60 to 100 bytes (for router links, this depends on the number of interfaces on the router). This should be added to memory used by other processes and by the IOS itself. If you really want to know the exact number, you can do a **show memory** with and without OSPF being turned on. The difference in the processor memory used would be the answer (keep a backup copy of the configs).
- Normally, a routing table with less than 500K bytes could be accommodated with 2 to 4 MB RAM; Large networks with greater than 500K may need 8 to 16 MB, or 32 to 64 MB if full routes are injected from the Internet.

Whew!

80

Cabrillo College

CCNP – Advanced Routing

Ch. 5 OSPF - Multi-areas (Part II)

Rick Graziani, Instructor

81