Open Shortest Path First

ROUTE Module 3

Agenda

- Introduction
- Basic configuration
- Authentication
- LSA
- Area Types
- OSPF in NBMA
- Virtual Links
- Securing OSPF

Open Shortest Path First

- OSPF is currently the most commonly used link-state protocol
- Classless (supporting VLSM), summarization, authentication, fast convergence
- Open standard existing in two versions:
  - OSPFv2 for IPv4 networks – RFC 2328
  - OSPFv3 for IPv6 network – RFC 5340
- In IPv4 it has own IP protocol number 89 and using multicast for communication:
  - 224.0.0.5: address of all OSPF routers on segment
  - 224.0.0.6: address of DR/BDR on segment
- Cost is metric derived from the speed of link
- Administrative distance of the OSPF network is 110, but it’s possible to redefine AS for intra-area, inter-area or external routes

OSPF Terms ①

- **Link**
  - Interface of router
- **Link-state**
  - Properties of link – IP address/subnet mask, cost, neighbors
- **Link-state ID (LSID)**
  - Unique ID under which link-state record is accessible in link-state database
- **Router ID (RID)**
  - 4B long unique identifier for router in the routing domain
  - It might or might not be same as an IP address of some router's interfaces

OSPF Terms ②

- **Area**
  - The set of networks and routers sharing topology knowledge
  - Area is identified by 4B long number
  - Every area must be directly connected to backbone (area with id 0)
  - Area borders are on routers not on links
- **Area Border Router (ABR)**
  - Router with interfaces residing in different areas
  - In OSPF design every ABR must be connected with at least one interface to backbone
  - ABR acts as point of summarization, filtering or just passing routing information between areas
- **Autonomous System Boundary Router (ASBR)**
  - Router on border between OSPF domain and the rest of the network (either different routing protocol domain or other AS)
  - ASBR acts as the point for redistribution (import), filtering or summarization of the routing information outside of OSPF domain

OSPF Terms ③

- **Designated Router (DR)**
  - Router on multi-access segment which serves as the central point for exchanging routing information on the segment
  - One DR is elected for every multi-access segment
  - DR is also responsible for representing link-state of multi-access segment
- **Backup Designated Router (BDR)**
  - Router on the multi-access segment backing up functionality of DR – it could take over DR role in case of original DR failure
  - BDR is not necessarily elected (there COULD be none on multi-access segment)
- **Link-state Advertisement (LSA)**
  - Abstract data structure sent in OSPF packets which describes topology information
  - Every LSA has its own header and body
  - *It's not just single packet!!!*

OSPF Terms ④

- **Neighborhood**
  - Communication relationship between two neighbor routers
  - It is established immediately when routers agree on mandatory parameters
  - No routing information are exchanged through neighborhood – it just confirms ability of routers to mutually communicate
  - *In short words neighborhood is established between any two correctly configured OSPF routers*
- **Adjacency**
  - Also communication relationship between two neighbor routers
  - Adjacency is closer – it allows to exchange routing information between the two adjacent routers
  - It is established only between some of OSPF routers

Packet Types ①

- OSPF Header

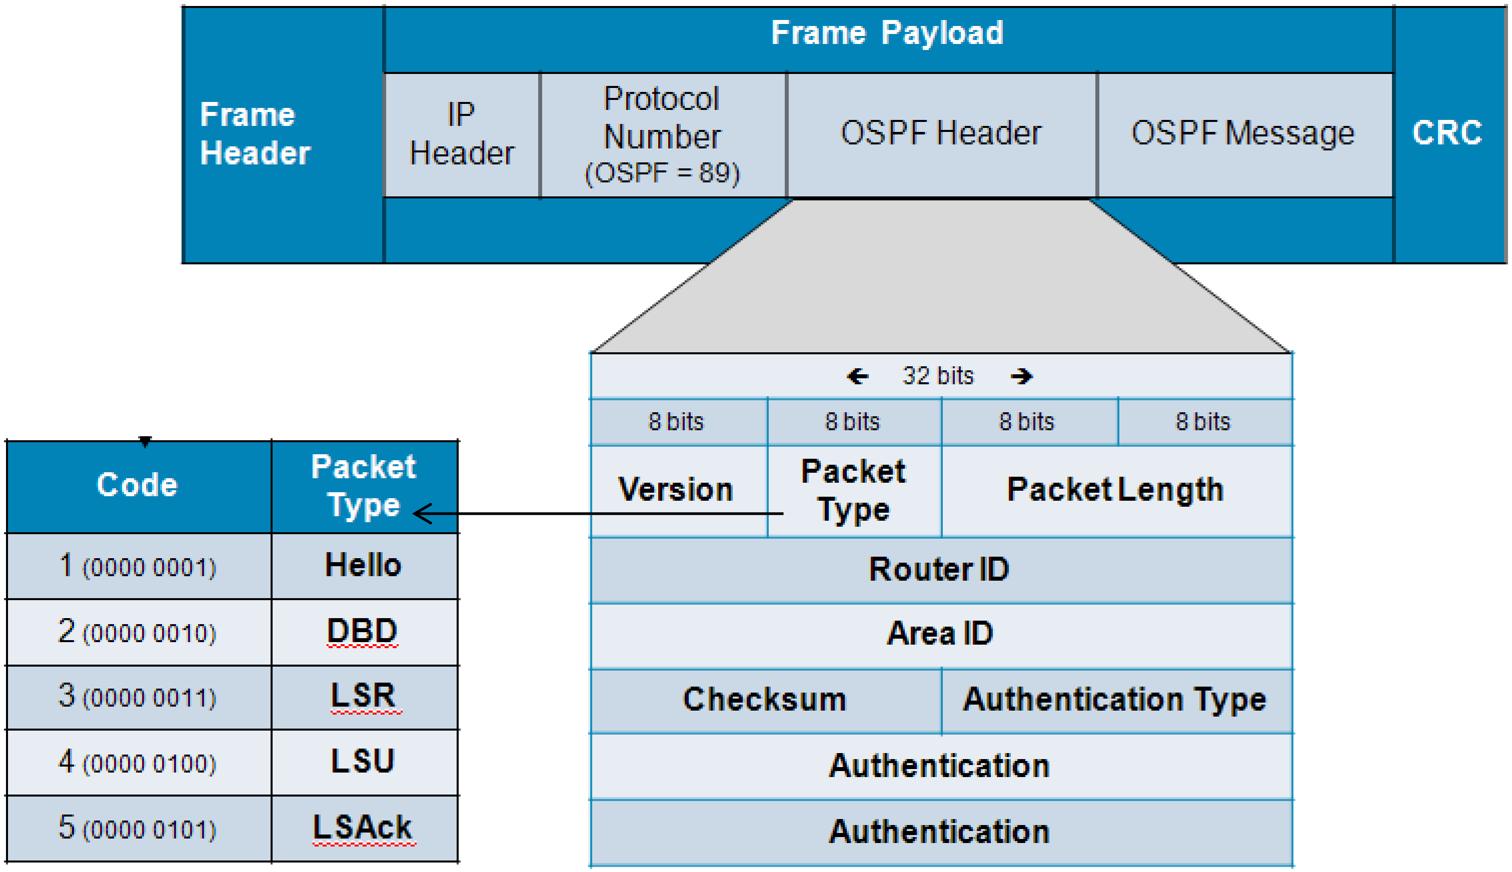

Code	Packet Type
1 (0000 0001)	Hello
2 (0000 0010)	DBD
3 (0000 0011)	LSR
4 (0000 0100)	LSU
5 (0000 0101)	LSAck

- OSPF has 5 messages

Packet Types ②

- **Hello packet**
  - Is used for discovering and checking neighborhood relationship with neighbor routers or to elect DR/BDR on multi-access router
  - It carries mandatory parameters
  - **Hello Interval** affects period of sending Hello router
    - every 10 seconds on broadcast and Point-to-Point links
    - every 30 seconds on NBMA and Poin-to-Multipoint links
  - When **Dead Interval** is times out neighborhood is considered down
    - By default is 4× larger than Hello interval

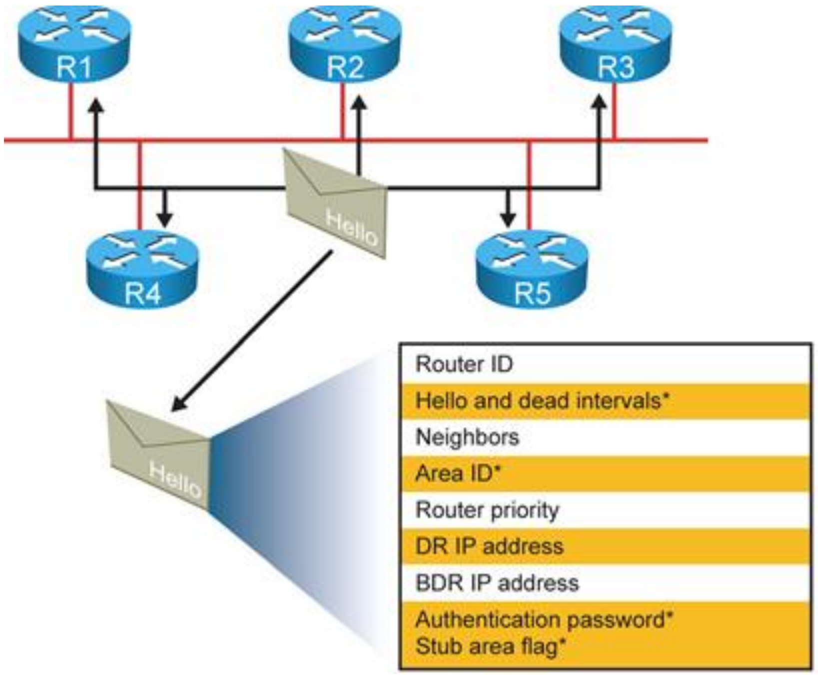

Packet Types ③

- **Database Description Packet (DDP or DBD)**
  - This packet is sent during initial synchronization of topological databases between the pair of routers
  - *It carries only "headers" of link-states not the all routing information*
    - More precisely only LSIDs of LSA-records in topological database
  - DBD packets are sent during phase of synchronizing topological databases – the goal is to compare LSAs between routers
  - In the header of DBD packets is carried also MTU – it MUST be same between communicating routers, otherwise router with the lower MTU wouldn't accept DBDs from router with higher MTU

Packet Types ④

- **Link-state Request (LSR)**
  - With LSRs router is able to ask for target LSA record from the neighbor’s topology database
  - It contains LID of requested LSA
- **Link-state Update (LSU)**
  - LSU carries topological information
  - Inside LSU are transmitted one or more LSAs
- **Link-state Acknowledgement (LSAck)**
  - Surprisingly used for acknowledgement of correctly received LSA
  - Multiple LSAs are acknowledged with just one LSAck

State Transitions

- Every OSPF process transits through 5 states:
  1. Discovery of neighbors and building communication relationships
  2. DR/BDR election (if necessary)
  3. Synchronization of topological databases
  4. Shortest Path Tree computation and populating of routing tables
  5. Keeping topology table updated

State 1

- Neighbors discover each other thanks to Hello packets that are sent on multicast address 224.0.0.5 or directly via unicast
- Routers check parameters inside Hello packet weather they are as expected – `IF true THEN` routers consider themselves as neighbors (neighborhood)
- Parameters that MUST be same between neighbors:
  - Network and subnet mask
  - Area number
  - Authentication
  - Hello and Dead Interval
- `IF` DR/BDR are already elected `THEN` their IP address would be inside Hello packets to allow other OSPF routers established adjacencies
- `IF` DR/BDR are not yet elected and network type needs election `THEN` election process is started

State 2 ①

- DR and BDR MUST be elected on every multi-access segment
- Every OSPF router has priority set in range 0 to 255 on each interface connected to multi-access segment
- Priority influences election process:
  - Routers with the priority 0 don't participate in election
  - Router with the highest priority on segment becomes DR
  - Router with the second highest priority becomes BDR
  - `IF` multiple priority values are same `THEN` the highest Router ID is used as tiebreaker
- Router waits with election process for so called "Wait Interval" (same as Dead Interval)
  - *The reason is to postpone process until enough Hello packets are received to see objective topology state*
- DR/BDR election is not preemptive – elected DR/BDR keeps this role until failure, it CAN NOT be push back by router with even higher priority

State 2 ②

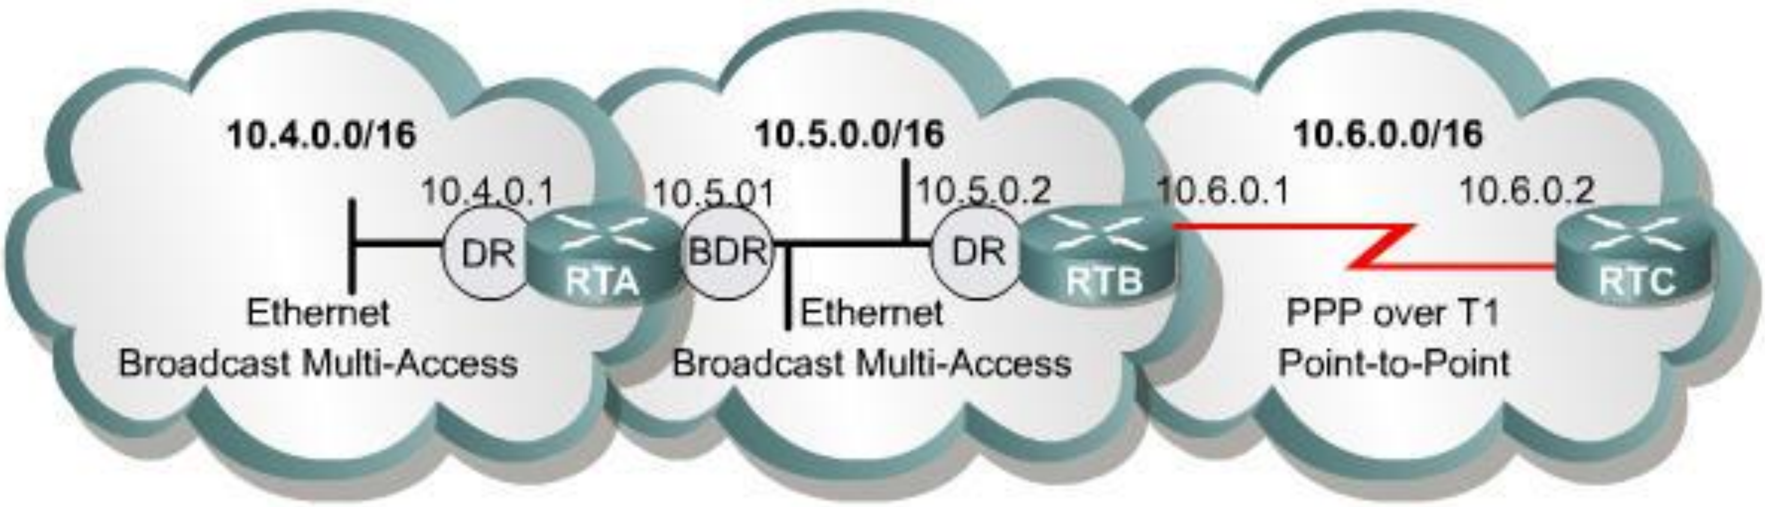

OSPF routers perform DR and BDR elections only on multiaccess IP networks.

State 3 ①

- Each pair of routers goes through multiple states during establishment of adjacent relationship:
- Phase **Down**
  - Starting phase – router is sending Hello packets but it doesn't received any Hello packets from neighbor
- Phase **Attempt**
  - Phase accessible only on NBMA networks – same as Down but router is unicasting Hello packets directly to neighbor IP address
- Phase **Init**
  - Router has received a Hello packet from its neighbor, but the receiving router's ID is not included in the Hello packet
  - When a router receives a Hello packet from a neighbor, it SHOULD list the sender's router ID in its Hello packet as an acknowledgment that it received a valid Hello packet (with correct parameters)

State 3 ②

- Phase **2-Way**
  - Router received valid Hello packet with neighbor’s Router ID inside

- All neighbor routers pass through all previous states
- In 2-Way phase all conditions for establishing adjacent (adjacency) relationship between the pair of routers are satisfied
- Ordinary neighbors stays at 2-Way phase – they are not exchanging topological information
- Through other phases continue only pairs of routers where:
  - either one of them is DR/BDR
  - or both are on the network segment where DR/BDR are not elected

State 3 ③

- Phase **ExStart**
  - Routers exchange empty DBD packets to discover who will be Master (router with higher Router ID) and who will be Slave
  - Master can increment sequence numbers in DBD packets
  - Slave has to reply to Master by repeating of current Masters sequence number
- Phase **Exchange**
  - Routers are exchanging DBD packets with descriptions of their topology databases
  - Router is building list of more recent LSAs from neighbors topology table which it requests from neighbor in following phase

State 3 ④

- Phase **Loading**
  - The list of LSAs built in previous phase is used and router starts updating its topology database
  - LSRs are sent, LSUs are replied and router acknowledge reception with LSAcks
- Phase **Full**
  - Router enters this phase when topology databases between neighbors are synchronized
  - Topology database is same as database of other routers in OSPF domain

State 3 ⑤

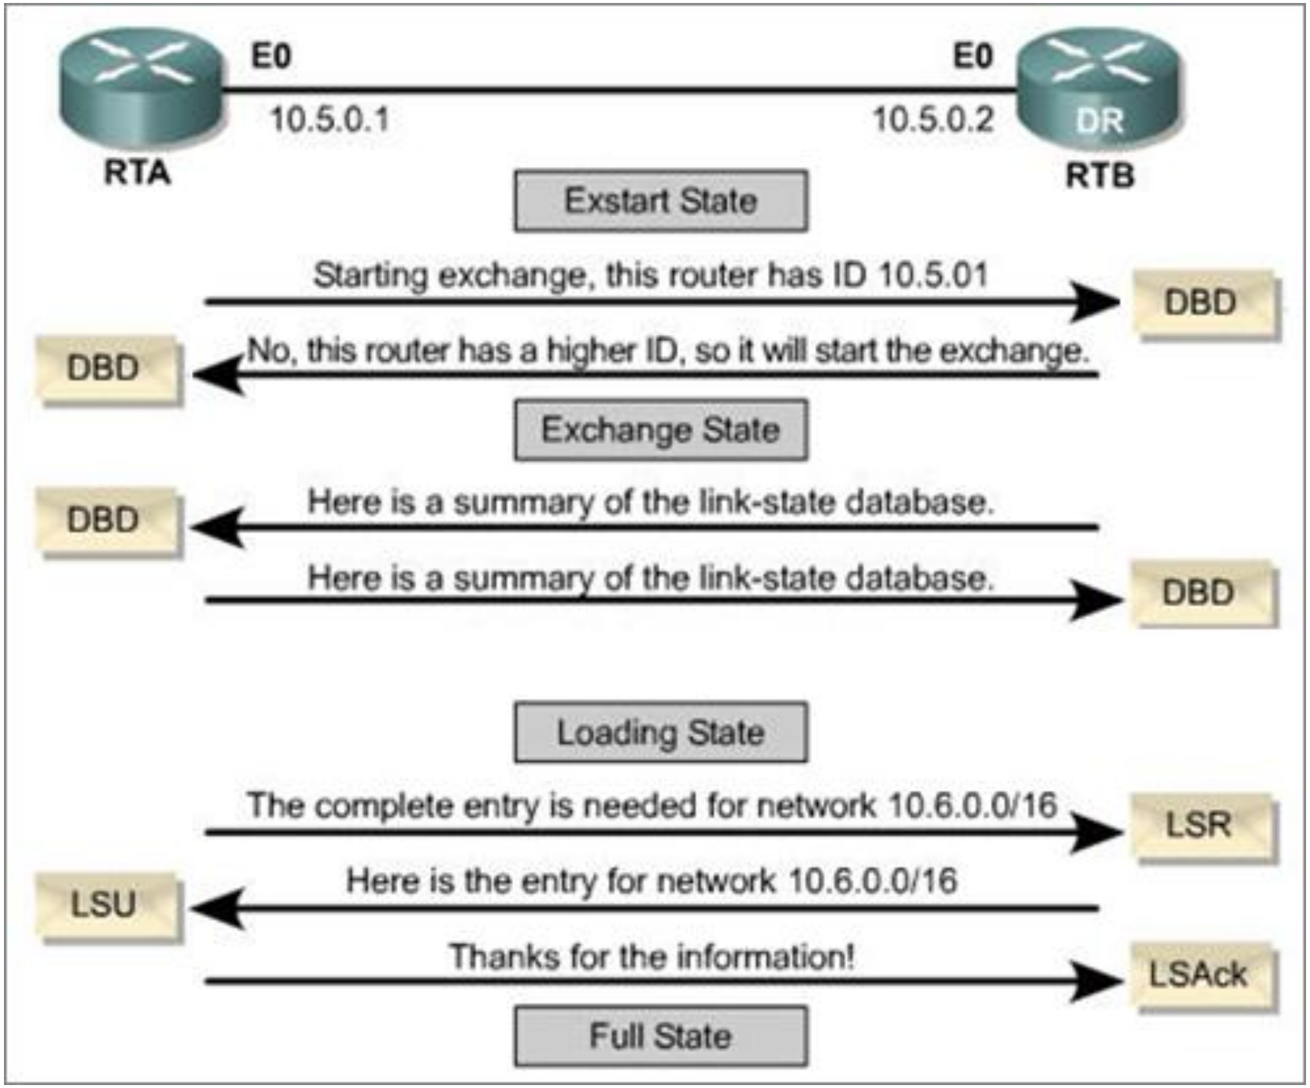

E0
10.5.0.1
RTA
E0
10.5.0.2
DR
RTB
Exstart State
Starting exchange, this router has ID 10.5.01
DBD
No, this router has a higher ID, so it will start the exchange.
DBD
Exchange State
Here is a summary of the link-state database.
DBD
Here is a summary of the link-state database.
DBD
Loading State
The complete entry is needed for network 10.6.0.0/16
LSR
Here is the entry for network 10.6.0.0/16
LSU
Thanks for the information!
LSAck
Full State

States 4 and 5

- Router with synchronized topological database SHOULD start Dijsktra's Shortest Path Tree (SPT) algorithm
- Every topological change will initiate:
  1. Advertising of link-state change to neighbors
  2. Computation of the new SPT
- Routers in Full phase send LSU immediately
- On network with elected DR/BDR
  - Router (that recognized link-state change) send LSU on address 224.0.0.6 of DR
  - DR pass this information to all other OSPF routers on segment via address 224.0.0.5
  - Recipients acknowledge with LSAck

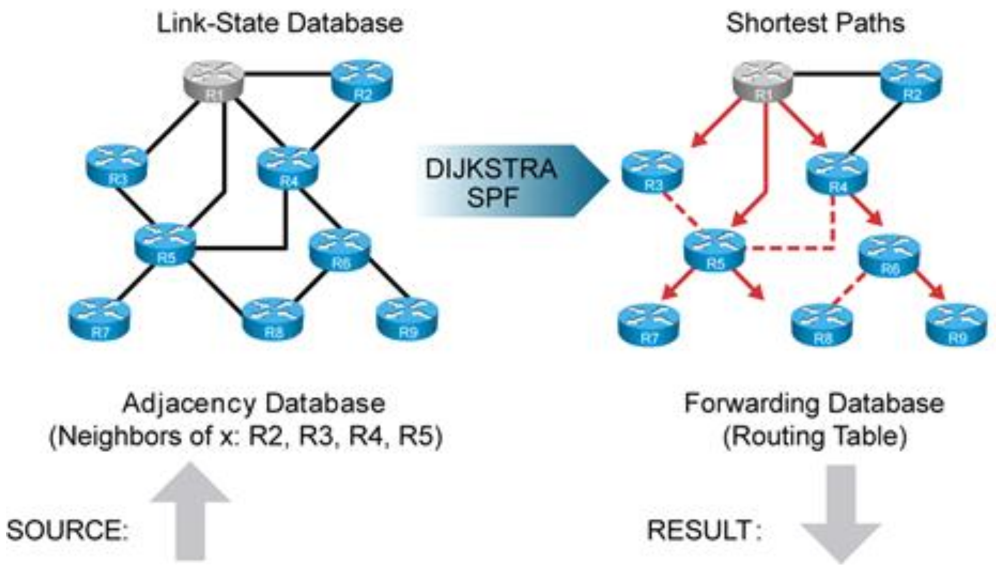

State Transitions Diagram

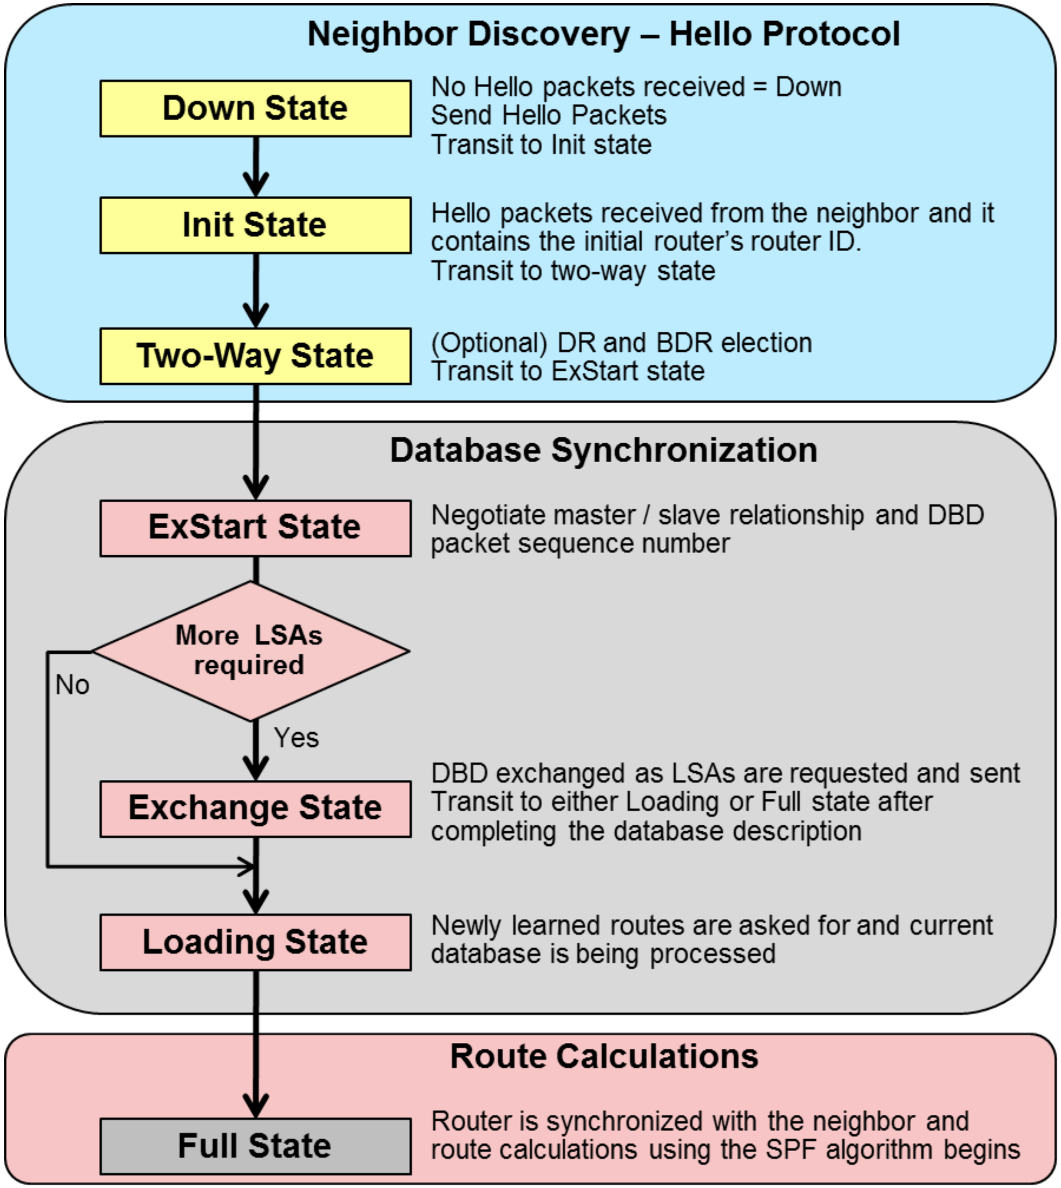

Configuration

Basic Configuration

```
Router(config)# router ospf process-id [vrf vpn-name]
```

- Starts OSPF process. Number must be in range from 1 to 65535 and its just locally relevant

```
Router(config-router)#
  network ip-address wildcard-mask area area-id
```

- Defines the list of interfaces (with appropriate networks) that will participate in OSPF routing process

```
Router(config-if)#
  ip ospf process-id area area-id [secondaries none]
```

- Alternative way to enable OSPF on per-interface basis
- **secondaries none** forbids sending of secondary addresses

Changing of Metric ①

- OSPF metric is called **cost** and lower cost is better/preferred

$$Cost = \frac{100\ \text{Mbps}}{\text{Bandwidth}}$$

Interface Type	10^8/bps = Cost
Fast Ethernet and faster	10^8/100,000,000 bps = 1
Ethernet	10^8/10,000,000 bps = 10
E1	10^8/2,048,000 bps = 48
T1	10^8/1,544,000 bps = 64
128 kbps	10^8/128,000 bps = 781
64 kbps	10^8/64,000 bps = 1562
56 kbps	10^8/56,000 bps = 1785

- *This evidentially does not reflect current situations!*

Changing of Metric ②

```
RouterA(config-router)#
  auto-cost reference-bandwidth ref-bw
```

- With this command 100 Mbps referential bandwidth could be changed to anything in range from 1 to 4 294 967 Mbps
- All routers MUST use same reference bandwidth to work properly – it MUST be statically configured because it is not carried in any OSPF message

```
RouterA(config-if)# ip ospf cost interface-cost
```

- Configuration of cost on per-interface basis in range from 1 to 65535

Router ID ①

- OSPF router MUST be uniquely identified among all other routers in topology – Router ID (RID)
  - RID is most notably used in the header of many LSAs inside **link-state database**
- RID is chosen at time of initiation of OSPF process – otherwise error message is generated
- RID is chosen in following order:
  - By command `router-id` in OSPF process
  - `IF` this command isn’t present `THEN` the highest IP address of any active Loopback interface is used
  - `IF` no Loopbacks are active `THEN` the highest IP address of any other active physical interface is used
  - *The first two options are considered best practice*

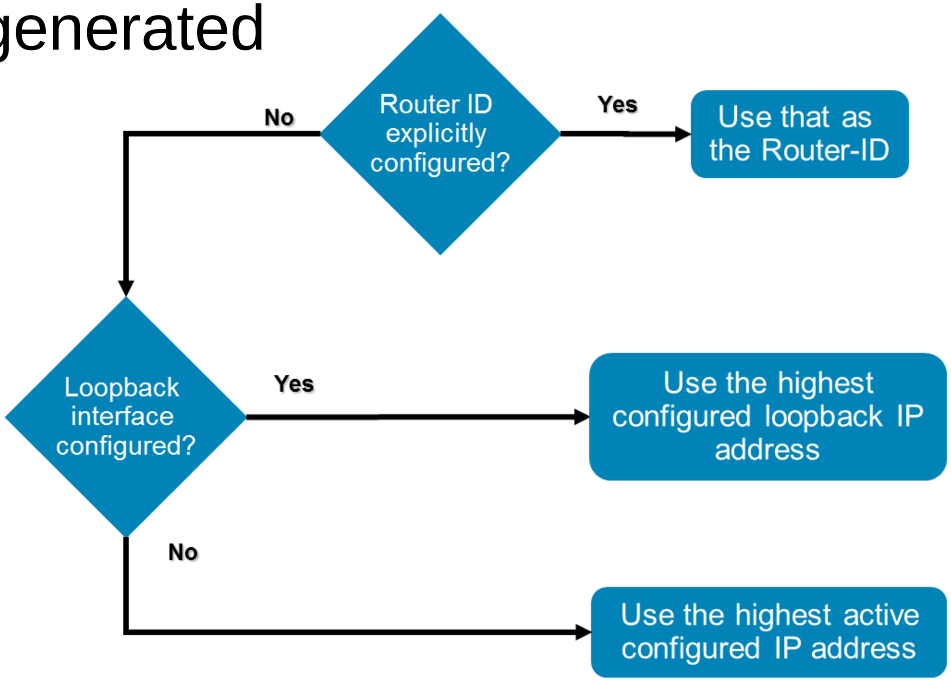

Router ID ②

```
Router(config-router)# router-id A.B.C.D
```

- Command is used inside OSPF process configuration context
- RID is 32-bit long unsigned integer
- Once the OSPF RID is set, it does not change, even if the interface that router is using for the router ID goes down!!!

```
Router# clear ip ospf process
```

- Suggested way of resetting OSPF process

```
Router(config)# router ospf 1
Router(config-router)# router-id 172.16.1.1

Router# clear ip ospf process
Clear all OSPF processes? [no] yes
```

Default Route in OSPF

```
Router(config-router)#
  default-information originate [always] [metric cost]
```

- Generates default route as one of LSAs
  - Without optional **always**: Router generates default route IFF it already has default route inside own routing table
  - With **always**: Router is generating default route no matter what (even thou it is possible not aware of it)
- It is not possible to redistribute default route into OSPF

Passive Interface

- The sending and receiving of routing updates could be disabled by following command:

```
Router(config-router)# passive interface {default | IFACE}
```

- The specified interface address appears as a stub network in the OSPF domain

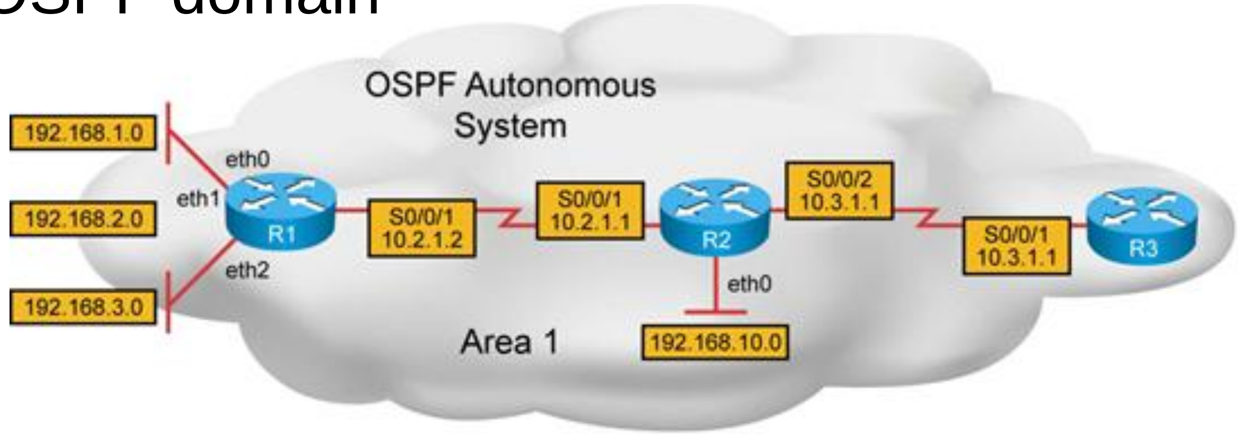

R1#

```
router ospf 100
 network 192.168.0.0 0.0.255.255 area 1
 network 10.2.0.0 0.0.255.255 area 1
 passive-interface default
 no passive-interface Serial0/0/1
```

R2#

```
router ospf 100
 network 192.168.0.0 0.0.255.255 area 1
 network 10.2.0.0 0.0.255.255 area 1
 network 10.3.0.0 0.0.255.255 area 1
 passive-interface Ethernet0
```

Authentication

Simple Password Authentication ①

```
Router(config-if)# ip ospf authentication-key password
```

- Configures plaintext password on interface

```
Router(config-router)# area area-id authentication
```

- For backward compatibility with older IOSes also area authentication mode is supported

```
Router(config-if)# ip ospf authentication [null]
```

- In newer IOSes each interface could support different authentication method
- **`null`** optional argument deactivates authentication on target interface

Simple Password Authentication ②

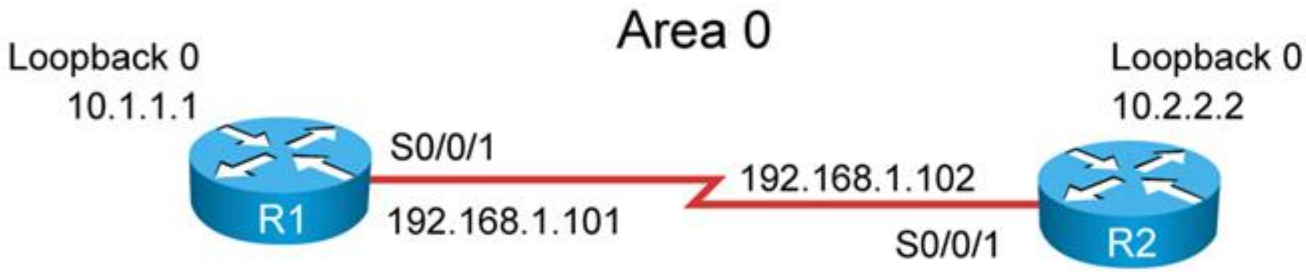

R1#

```
<output omitted>
interface Loopback0
 ip address 10.1.1.1 255.255.255.0

<output omitted>
interface Serial0/0/1
 ip address 192.168.1.101 255.255.255.224
 ip ospf authentication
 ip ospf authentication-key plainkey

<output omitted>
router ospf 10
 log-adjacency-changes
 network 10.1.1.1 0.0.0.0 area 0
 network 192.168.1.0 0.0.0.255 area 0
```

R2#

```
<output omitted>
interface Loopback0
 ip address 10.2.2.2 255.255.255.0

<output omitted>
interface Serial0/0/1
 ip address 192.168.1.102 255.255.255.224
 ip ospf authentication-key plainkey

<output omitted>
router ospf 10
 log-adjacency-changes
 network 10.2.2.2 0.0.0.0 area 0
 network 192.168.1.0 0.0.0.255 area 0
 area 0 authentication
```

MD5 Authentication ①

```
Router(config-if)#ip ospf message-digest-key keyid md5 key
```

- Creates key and binds it with ID
  - Pair (KeyID)-Key must be same between neighbors
  - `IF` multiple keys are present on interface `THEN` the last added key is used for signing outgoing message
  - All of present keys are used when accepting message
- Older IOS whole area MD5 authentication

```
Router(config-router)#
  area area-id authentication message-digest
```

- It turns MD5 authentication on target interface. And as in previous case optional argument **`null`** deactivates it.

```
Router(config-if)#
  ip ospf authentication {message-digest | null}
```

MD5 Authentication ②

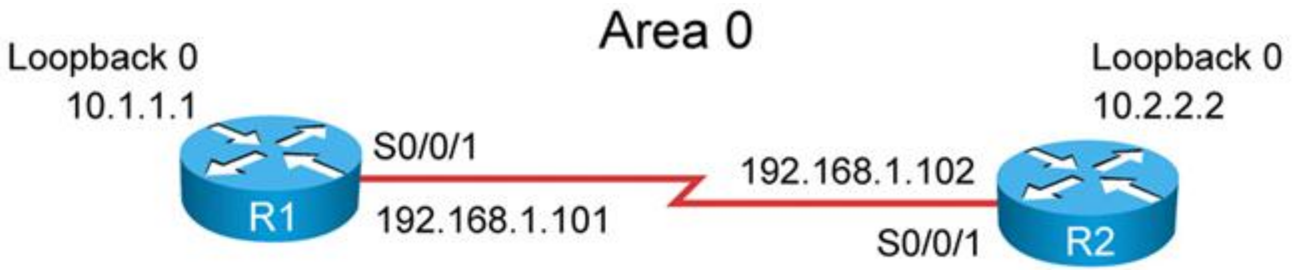

R1#

```
<output omitted>
interface Loopback0
 ip address 10.1.1.1 255.255.255.0

<output omitted>
interface Serial0/0/1
 ip address 192.168.1.101 255.255.255.224
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 mysecret

<output omitted>
router ospf 10
 log-adjacency-changes
 network 10.1.1.1 0.0.0.0 area 0
 network 192.168.1.0 0.0.0.255 area 0
```

R2#

```
<output omitted>
interface Loopback0
 ip address 10.2.2.2 255.255.255.0

<output omitted>
interface Serial0/0/1
 ip address 192.168.1.102 255.255.255.224
 ip ospf message-digest-key 1 md5 mysecret

<output omitted>
router ospf 10
 log-adjacency-changes
 network 10.2.2.2 0.0.0.0 area 0
 network 192.168.1.0 0.0.0.255 area 0
 area 0 authentication message-digest
```

Verification of Authentication

```
R1#show ip ospf interface
Serial2/0 is up, line protocol is up
  Internet Address 192.168.1.101/27, Area 0
  Process ID 10, Router ID 10.1.1.1, Network Type POINT_TO_POINT, Cost: 64
  Transmit Delay is 1 sec, State POINT_TO_POINT
Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.2.2.2
  Suppress hello for 0 neighbor(s)
  Message digest authentication enabled
    Youngest key id is 1
Loopback0 is up, line protocol is up
  Internet Address 10.1.1.1/24, Area 0
  Process ID 10, Router ID 10.1.1.1, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host

R1#ping 10.2.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/32 ms
```

```
R1#
*Feb 17 18:51:31.242: OSPF: Rcv pkt from 192.168.1.102, Serial0/0/1 :
Mismatch Authentication type. Input packet specified type 0, we use type 1

R2#
*Feb 17 18:50:43.046: OSPF: Rcv pkt from 192.168.1.101, Serial0/0/1 :
Mismatch Authentication type. Input packet specified type 1, we use type 0
```

```
R1#
*Feb 17 18:54:01.238: OSPF: Rcv pkt from 192.168.1.102, Serial0/0/1 :
Mismatch Authentication Key - Clear Text

R2#
*Feb 17 18:53:13.050: OSPF: Rcv pkt from 192.168.1.101, Serial0/0/1 :
Mismatch Authentication Key - Clear Text
```

LSAs

OSPF Databases

- OSPF is working with three databases
  - **Adjacency Database** (**`show ip ospf neighbor`**)
    - Databases of neighbor routers showing also relationships
  - **Link-state Database (LSDB)** (**`show ip ospf database`**)
    - Topological database consisting of oriented graph model of network built up with information from LSAs
    - All routers in the one area have identical LSDB
  - **Forwarding Database** (**`show ip route`**)
    - Routing table consists of information about every reachable destination network
    - OSPF theoretically knows whole path from itself to destination network but in routing table is only the first next-hop

Adjacency Database

The interface on which this router has formed adjacency with the neighbor.

```
R2# show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address       Interface
10.64.0.1        0    FULL/DROTHER    00:00:30    10.64.0.1     FastEthernet0/0
10.2.1.1         0    FULL/  -        00:00:34    10.2.1.1      Serial0/0/1
```

Lists the neighbors in the order they were learned.

The OSPF state of the interface. FULL state means that the router and its neighbor have identical OSPF link-state databases.

The IP address of the neighbor's interface to which this router is directly connected.

The OSPF priority of the interface.

The amount of time remaining that the router will wait to receive an OSPF Hello packet from the neighbor before declaring the neighbor down.

Diagram of Processing LSA

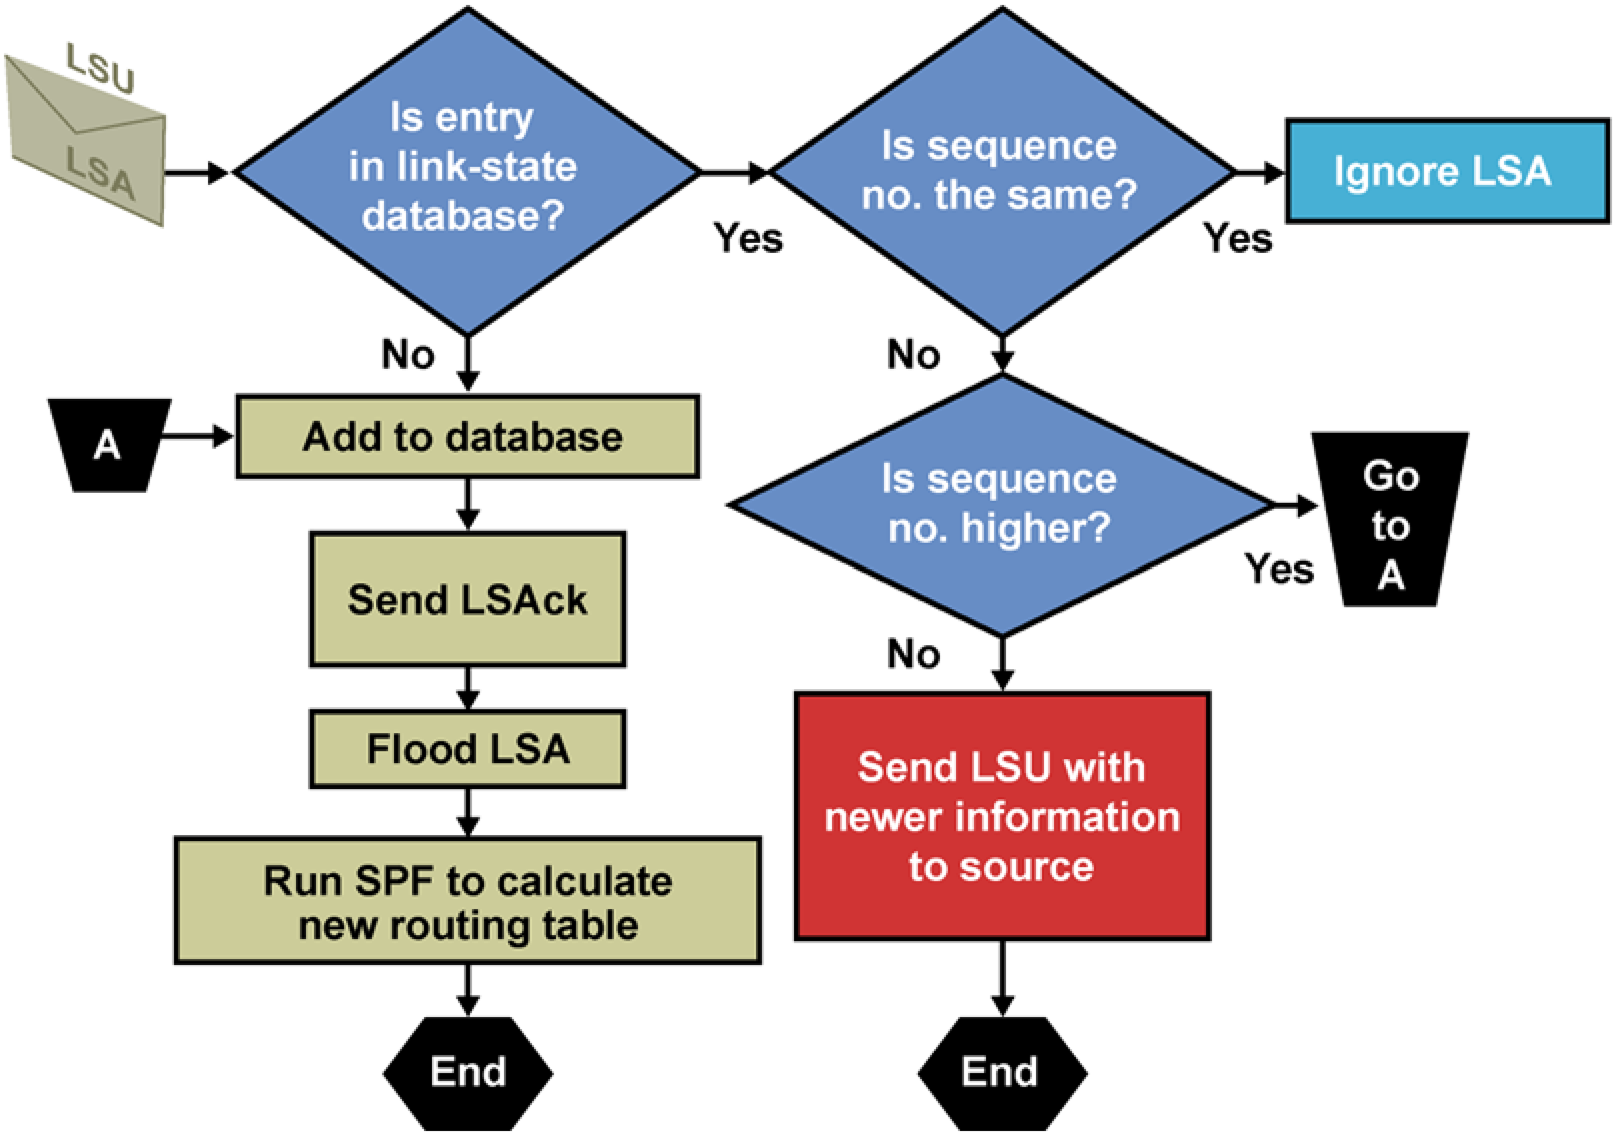

LSA Sequencing

- Each LSA in LSDB has its own sequence number (4B long)
- *LSA with a higher sequence number is considered newer!*
- Numbering starts with initial value 0x80000001 and ends with 0x7FFFFFFF
  - Value 0x80000000 is reserved and never used
- Router increments sequence number of all LSAs by 1 every 30 minutes and sends them all away
- After some time sequence reach value 0x7FFFFFFF
  - `IF` the next increment would be value 0x80000000 `THEN` the router lets this LSA expire by setting its Maxage Timer to 60 minutes
  - LSA is flushed and sequencing starts again from 0x80000001

Link-state Advertisements Type

LSA Type	Description
1	Router LSAs
2	Network LSAs
3 or 4	Summary LSAs
5	Autonomous System External LSAs
6	Multicast OSPF LSAs
7	Defined for Not-So-Stubby Areas
8	External Attributes LSA for Border Gateway Protocol (BGP)
9, 10, 11	Opaque LSAs

Common LSA Header

- Every LSA has header with following fields:
  - Link State Age: Age of LSA in seconds (maximally 3600)
  - Options: Bit flags for extended OSPF operation
  - Link State Type: Type of LSA
  - Link State ID (LSID): 4B long uniquely identifying this LSA in LSDB
  - Advertising Router: RID of router which generated this LSA
  - LS Sequence Number: Sequence number described above
  - LS Checksum: Control checksum
  - Length: Length of LSA in bytes

LSA Type 1 (LSA1): Router LSA O

- Generated by every router in an area to describe Intra-area routes
  - LSA includes the list of directly attached links and is identified by originating router ID
- Floods only within its area and CAN NOT cross an ABR
- LSID = RID of router which generated this LSA1
- LSA1 consists of 1 or more Link ID (LID) which identify to whom (network) and how (**link type**) is router connected

Link Types for LSA1

Link Type	Description	Link ID
1	Point-to-point connection with neighbor router	RID of neighbor router
2	Interface in transit network	IP address of DR in target network
3	Interface in stub network	Net ID/subnet mask
4	Virtual link	RID of virtual neighbor router

LSA Type 2 (LSA2): Network LSA O

- Generated by DR for every transit multi-access network
  - LSA lists each RID of the attached routers that make up the area, including the DR itself, as well as the subnet mask used on the link
- Floods within its area only; does not cross ABR
- LSID = IP address of DR (pointer to LSA1 LID)

LSA Type 3 (LSA3): Network Summary LSA O IA

- Generated by ABR
  - LSA3 contains possibly summarized network but without any additional topology information (just like distance-vector)
  - Without summarization it contains list of all networks in target area (one LSA3 per every network
- Regenerated by subsequent ABRs to flood throughout the autonomous system
- LSID = network (subnet) that’s advertising

LSA Type 4 (LSA4): ASBR Summary LSA

- Generated by the ABR of the originating area to advertise an ASBR to all other areas in the OSPF domain
- They are regenerated by all subsequent
- Link-state ID = RID of target ASBR

LSA Type 5 (LSA5): External LSA O E1/E2

- Generated by the ASBR to advertise networks from other autonomous systems
  - For every external network one LSA5
  - LSA5 are flooded to the entire AS
- LSID = external network number

Conclusion About LSA Types ①

- **LSA1 (Router LSA)**
  - **Router**
  - Connected Point-to-Point links with neighbors
  - Connected multi-access segment links with neighbors
  - Connected stub areas (non-transit networks without any other router)
- **LSA2 (Network LSA)**
  - **Multi-access network**
  - Connection to all neighbor routers
  - LSA1 and LSA2 together completely describe topology in one target area
- **LSA3 (Network Summary LSA)**
  - **IP addresses of network in different areas** (could be summarized)
  - Backward pointer on LSA1 of ABR in Advertising Router field
- **LSA4 (ASBR Summary LSA)**
  - **Existence of ASBR router in different areas**
  - Backward pointer on LSA1 of ABR in Advertising Router field
- **LSA5 (AS External LSA)**
  - **External network IP** behind ASBR
  - Backward pointer on LSA1/LSA4 of ASBR in Advertising Router field

Conclusion About LSA Types ②

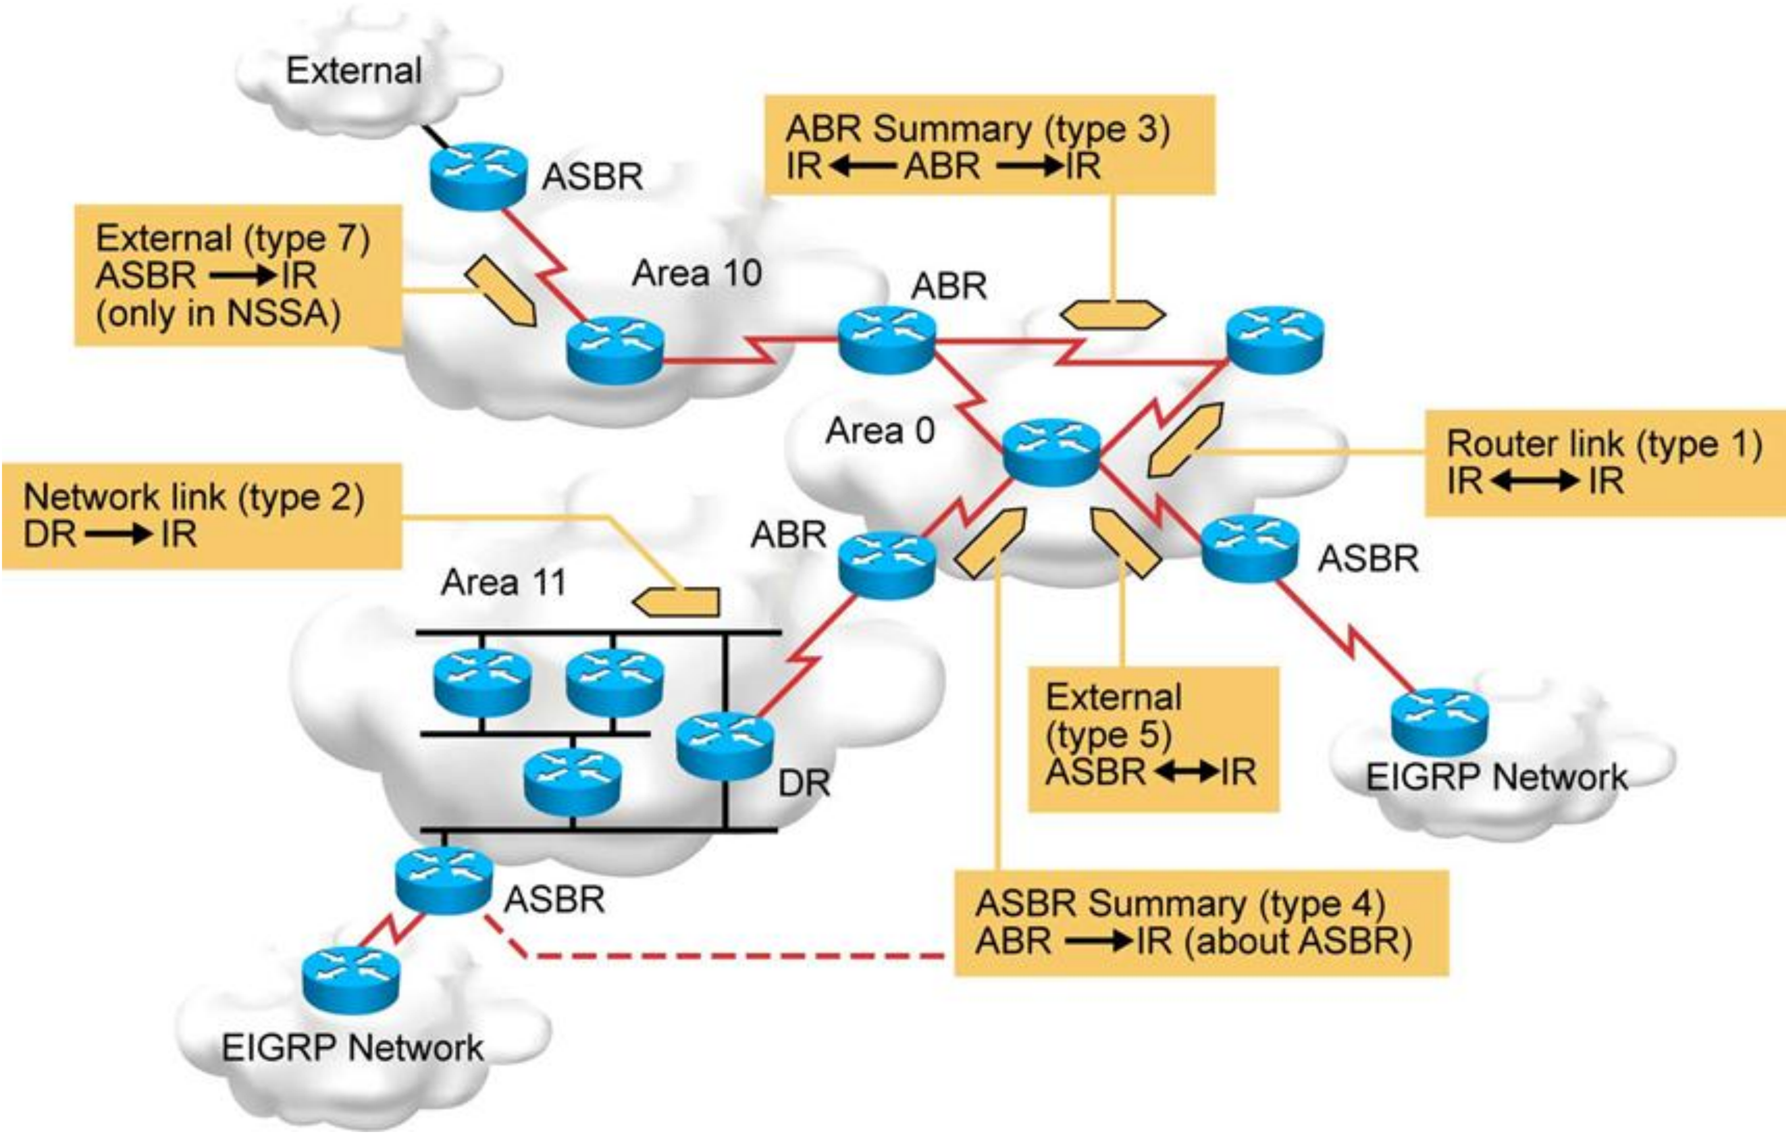

Meaning of LSDB Columns

```
RouterA# show ip ospf database
     OSPF Router with ID (10.0.0.11) (Process ID 1)

                Router Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum Link count
10.0.0.11       10.0.0.11       548         0x80000002 0x00401A 1
10.0.0.12       10.0.0.12       549         0x80000004 0x003A1B 1
100.100.100.100 100.100.100.100 548         0x800002D7 0x00EEA9 2

                Net Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum
172.31.1.3      100.100.100.100 549         0x80000001 0x004EC9

                Summary Net Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum
10.1.0.0        10.0.0.11       654         0x80000001 0x00FB11
10.1.0.0        10.0.0.12       601         0x80000001 0x00F516
<output omitted>
```

Total number of directly attached links, used only on router LSAs. Each point-to-point serial link counts as two; all other links count as one, including Ethernet links

Identifies each LSA. Even thou it's called Link ID it's Link-state ID actually!!!

Advertising router; the source router of the LSA

The maximum age counter in seconds; the maximum age is 1 hour or 3,600 seconds

Sequence number of the LSA; this number begins at 0x80000001 and increases with each update of the LSA

Checksum of the individual LSA to ensure reliable receipt of that LSA

Area Types

Motivation Behind Areas

- Benefits of using areas:
  - Reduce the size of routing tables
  - Isolate topology changes as much as possible to the area in which they occur
  - Allow only summary LSA updates to cross area boundaries and
  - Reap all the benefits of using a hierarchical addressing scheme

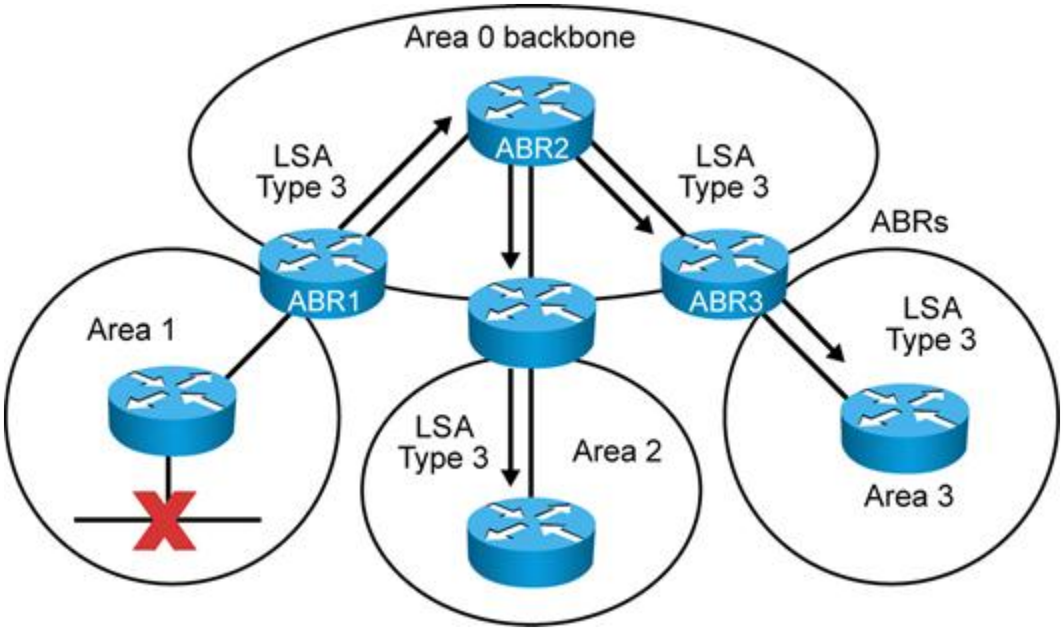

Area Types ①

- **Backbone Area**
  - Referred to as Area 0
  - A.k.a. Transit Area
- **Regular Areas**
  - A.k.a. Non-backbone areas
  - All regular areas must connect to the backbone area
  - *Why?*
  - Strong prevention against routing loops

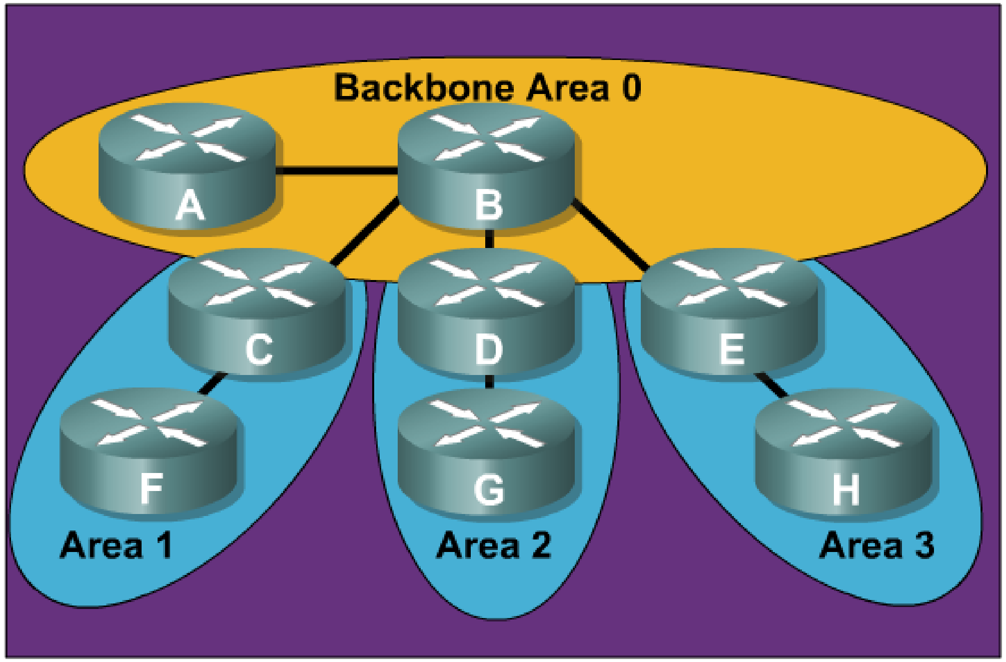

Area Types ②

- It is definitely not necessary for a router in one area to have a list of all networks in all areas (even redistributed ones)
- Reducing size of LSDB and routing table could be accomplished by a choosing appropriate **area type**
- Hence, OSPF regular area can be further divided into 4 **stub area** subtypes:
  - Stub area
  - Totally Stubby area
  - Not-So-Stubby Area (NSSA)
  - Totally Not-So-Stubby Area (Totally NSSA)

Area Types ③

- Special area types differ by filtering and alternating LSAs coming from backbone to area through ABR
- Special area types could be configured only on non-backbone areas
- Area 0
  - is always transit this area type CAN NOT be changed
  - Is always aware of all intra-area, inter-area and external routes
- Reduction and simplification of LSDB is visible only on internal routers of special area type

Transit/Regular Area

- Transit/Regular area is standard area type
  - It has information about networks in different areas (but not about their topology)
  - It has information about ASBR and external networks
  - It could contain some ASBR router(s)
- Configuration of router which is part of multiples area is trivial
  - In command **`network`** argument **`area`** is used
  - Remember that ABR has at least one interface in area 0

Stub Area

- **Stub area** rejects LSA4 and LSA5, they are not sent inside
  - Only O and O IA routes exist in routing table of stubby area router
  - There MUST be no ASBR inside the stub area and it does know nothing about external networks
  - The area MUST not be used as a transit area for virtual links
  - Preferably there is a single exit point from that area
  - `IF` there are multiple ABRs `THEN` one or more SHOULD inject a default route
- Every LSA5 is converted to default route through LSA3 when sent to stubby area
- Configuration on every router in stubby area:

```
Router(config-router)# area area-id stub
```

Totally Stubby Area

- **Totally Stubby area** rejects LSA3, LSA4 and LSA5, LSA4 and LSA5 are not sent inside it
  - Only O routes exist in routing table of totally stubby area router
  - Same qualities as Stub area with difference that Totally Stubby area relies on additional functionality of ABR (providing default gateway)
  - Hence, suitable for areas with only one ABR
- Every LSA3 and LSA5 is converted as default route through LSA3 when sent to stubby area
- Configuration for inner routers is same as in Stub area case and only on ABR following command must be entered

```
Router(config-router)# area area-id stub no-summary
```

NSSA and Totally NSSA ①

- In some cases it is required to have area with Stub/(Totally Stubby) characteristic BUT with ASBR inside it
  - Redistribution of static routes or routes from other routing protocols
- **NSSA** and **Totally NSSA** were invented for these purposes – *they have same qualities as their siblings but they allow existence of ASBR inside and external networks*
  - Among O and O IA routes in routing table could be present also external routes injected by local ASBR in (Totally) NSSA routers
- When comparing Stubby and NSSA the major difference is that NSSA area filters LSA4 and LSA5 upon receiving but ABR does not convert them to default route through LSA3 automatically
  - This default route injection must be additionally configured

NSSA and Totally NSSA ②

- Configuration snippet for NSSA or Totally NSSA
  - On every inner router of (Totally) NSSA except ABR issue:

```
Router(config-router)# area area-id nssa
```

  - ABR is configured with optional parameters in following manner:

```
Router(config-router)# area area-id nssa [no-summary]
[default-information-originate] [no-redistribution]
```

    - **`no-summary`**: LSA3 are not propagated, area is than considered as Totally NSSA
    - **`default-information-originate`**: IF area is NSSA THEN filtered LSAs are converted to default route
    - **`no-redistribution`**: External network aren’t redistributed into NSSA

LSA Type 7 (LSA7): NSSA LSA O N1/N2

- External networks are carried in NSSA with **LSA7** which is on ABR converted to LSA5 and then flooded further

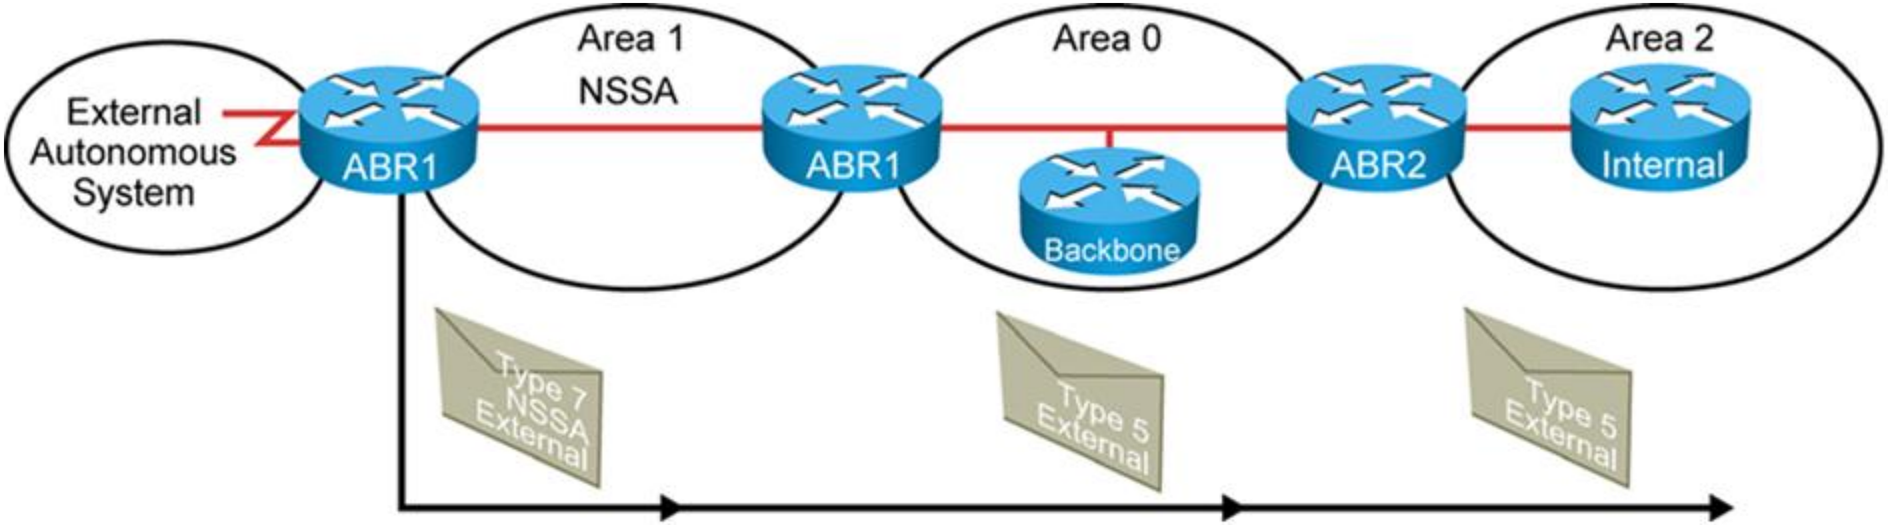

- Generated by an ASBR inside a Not-so-stubby area (NSSA) to describe routes redistributed into the NSSA
  - LSA 7 is translated into LSA 5 as it leaves the NSSA
- LSID = external network number

Conclusion on Area Types ①

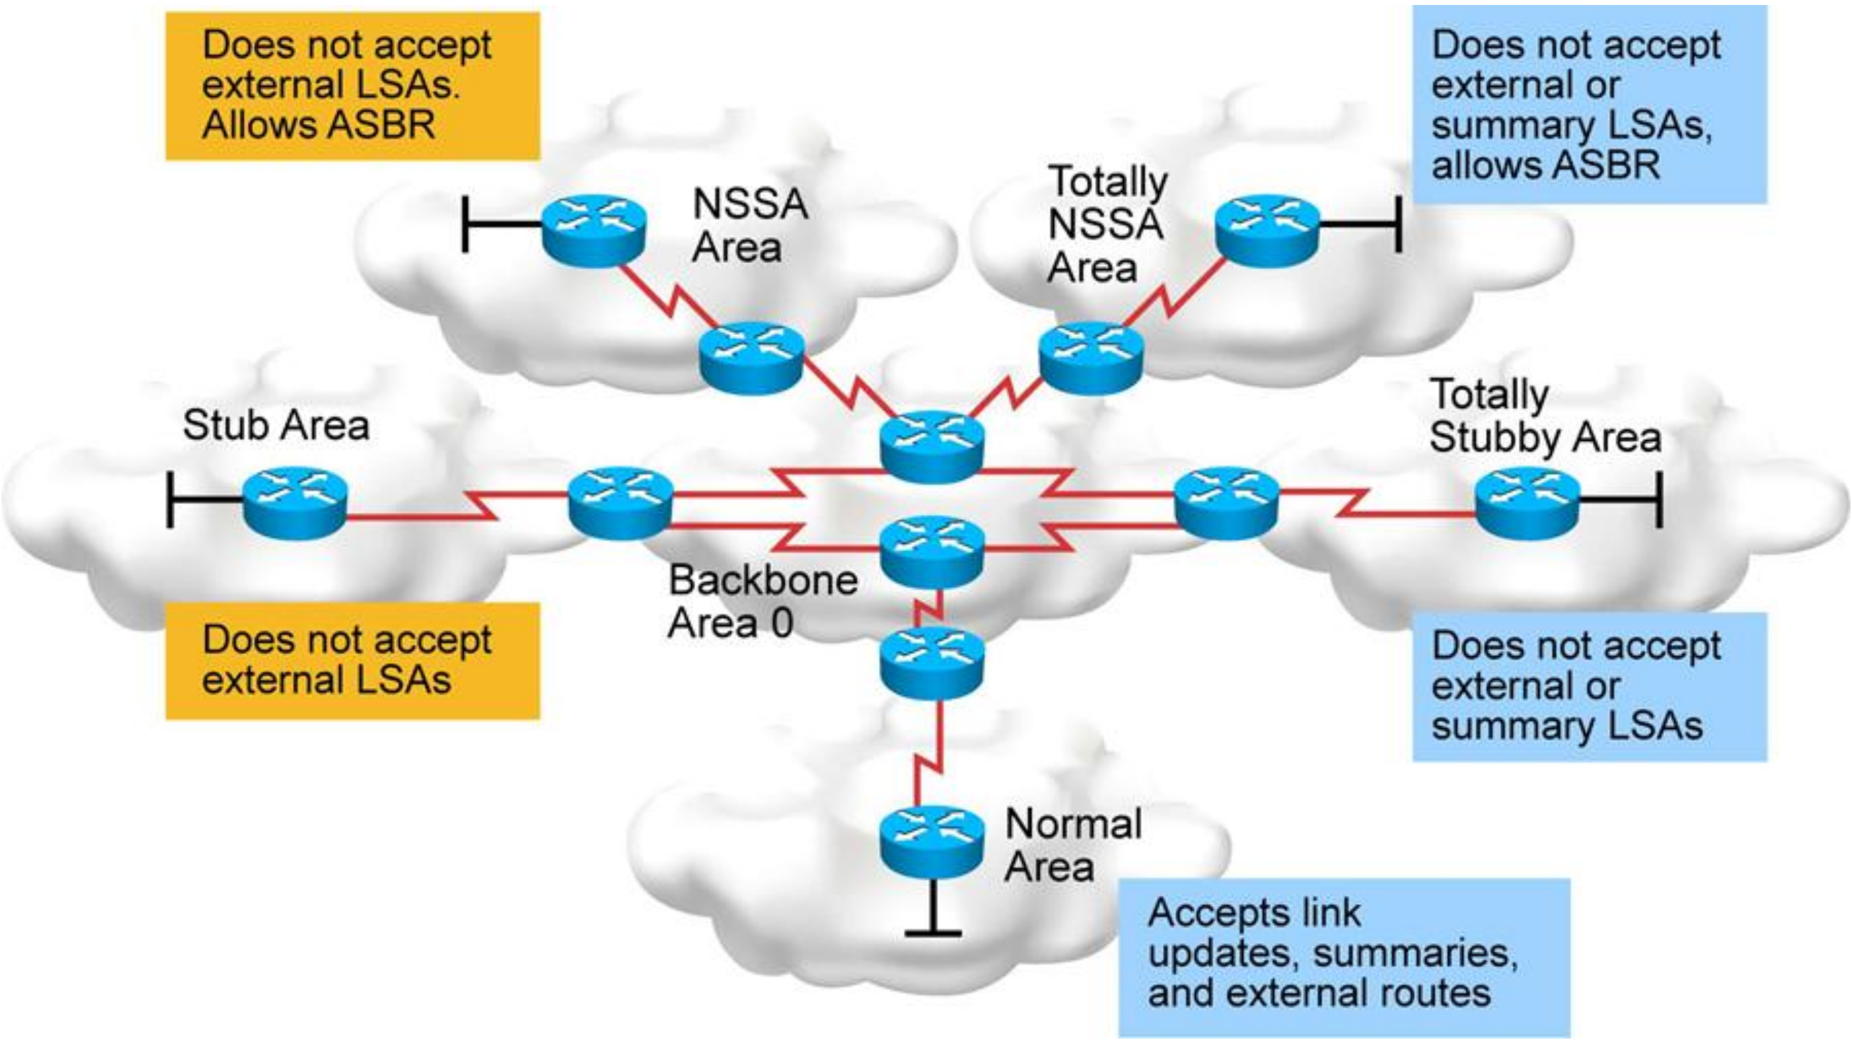

Conclusion on Area Types ②

Area Type	Accepts routes within area O	Accepts routes from other areas O IA	Accepts external routes O E1 and O E2	Allows ASBR	Cisco proprietary
Standard	Yes	Yes	Yes	Yes	No
Backbone	Yes	Yes	Yes	Yes	No
Stub	Yes	Yes	No (uses default route)	No	No
Totally stubby	Yes	No (uses default route)	No (uses default route)	No	Yes
NSSA	Yes	Yes	No (uses default route)	Yes **O N1/N2**	No
Totally NSSA	Yes	No (uses default route)	No (uses default route)	Yes **O N1/N2**	Yes

Best Path Selection: External Routes ①

- Based on **external type** network configuration on ASBR the metric of an external route is computed differently
- O E1 routes to external networks – *comparable*
  - = (to ABR) + (from ABR to ASBR) + (from ASBR to external network)
  - Use this packet type when there are multiple ASBRs advertising a route to the same autonomous system
- O E2 routes to external networks – *incomparable*
  - = (from ASBR to external network)
  - Use this packet type preferably if only one ASBR is advertising a route to the autonomous system
  - `IF` multiple routes have same metric `THEN` the route via the closest ASBR is used
- O N1/N2 are equivalents of O E1/E2 but redistributed into OSPF in NSSA

Best Path Selection: External Routes ②

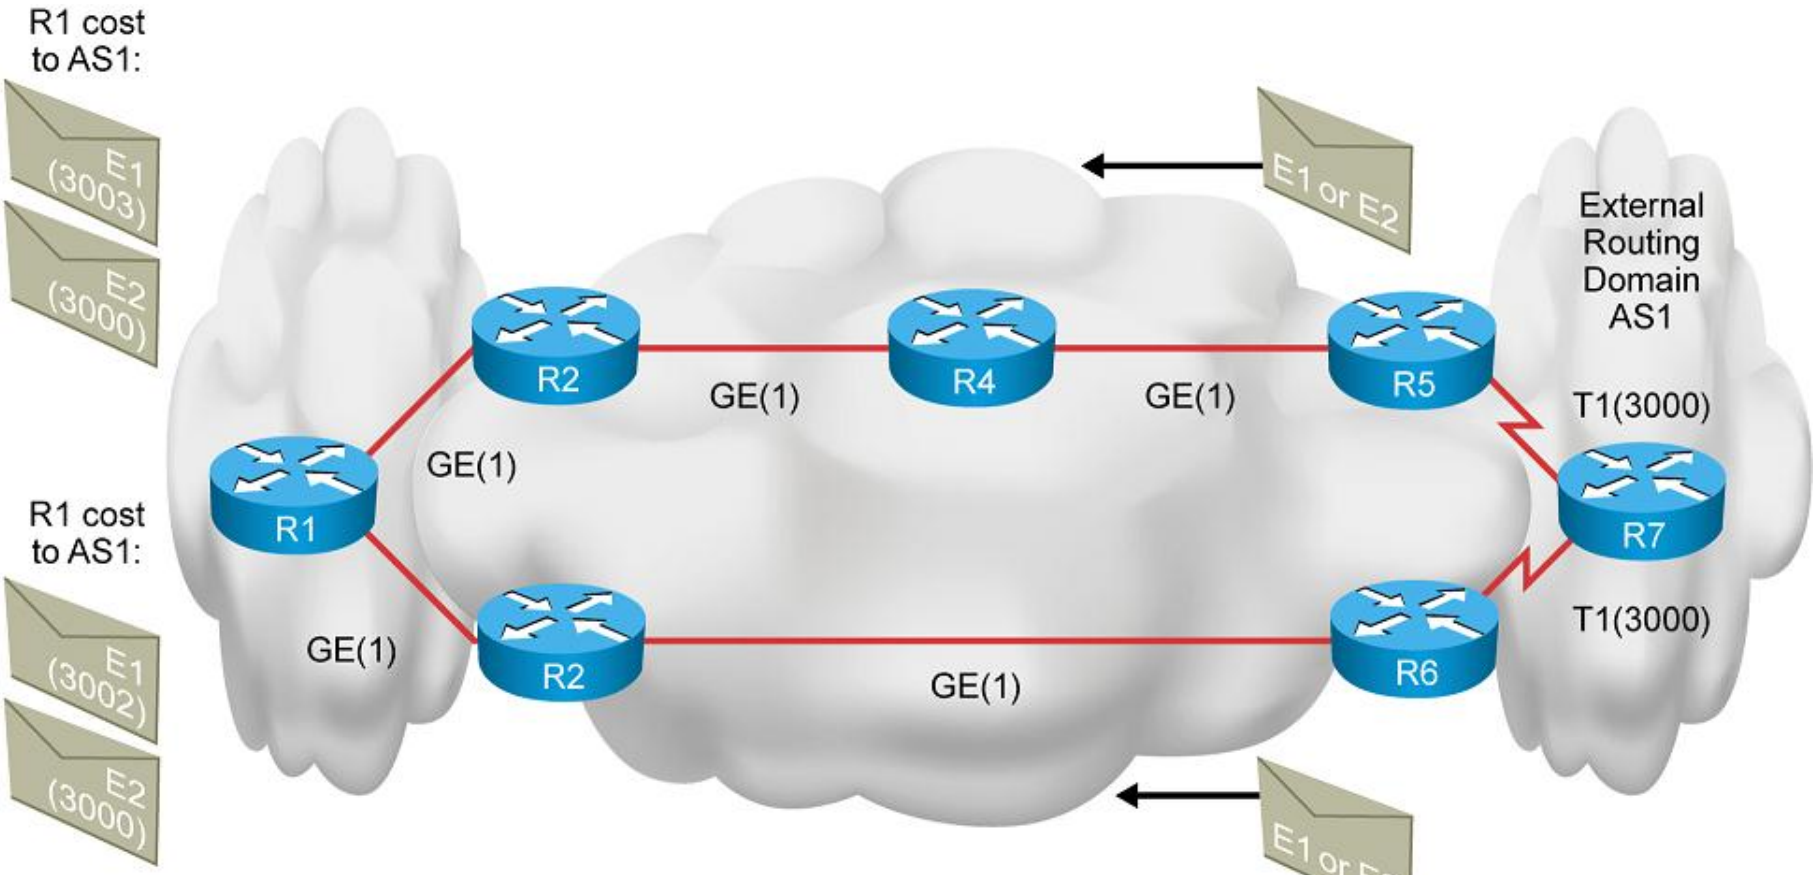

Best Path Selection Algorithm

- Best path is calculated by following rules:
  1. All routers uses intra-area (O) routes with best metric to target destination network
  2. All routers uses inter-area (O IA) routes with best metric to target area (= cost to ABR + cost from ABR to target area)
  3. Routers except ones inside Stub area uses external routes with best (O E1/E2) metric to target external routing domain
- `IF` multiple O IA/E1/E2/N1/N2 routes share same metric `THEN` all of them are used

Summarization

Summarization ①

- In OSPF exists two kinds of summarization:
  1. Summarization of inter-area networks
  2. Summarization of external networks redistributed into OSPF

- Inter-area summarization is configured on ABR:

```
Router(config-router)#
  area area-id range NET MASK [not-advertise] [cost cost]
```

  - **not-advertise**: optional argument that stops advertising of network and its components to other areas
  - **cost**: optional argument that sets cost of summary network

Summarization ②

- External networks summarization is configuring only on ASBR:

```
Router(config-router)#
 summary-address NET MASK [not-advertise] [cost cost] [tag tag]
```

  - Usage of optional arguments **not-advertise** and **cost** is analogical to previous summarization command
  - **tag**: tag number could be used in “match” clauses in route-maps, more about this topic in next Module

- Whenever summarization is not configured correctly and there are multiple ASBRs, or multiple ABRs in an area, suboptimal routing is possible
  - For example, summarizing overlapping ranges from two different routers can cause packets to be sent to the wrong destination.

Summarization ③

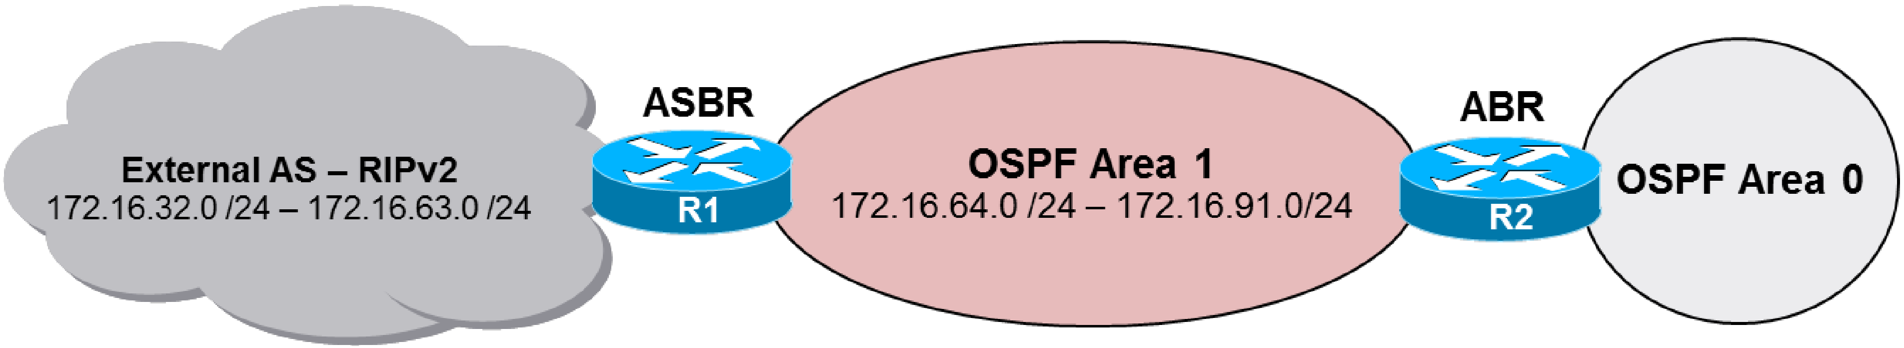


```
R1(config)# router ospf 100
R1(config-router)# network 172.16.64.1 0.0.0.0 area 1
R1(config-router)# summary-address 172.16.32.0 255.255.224.0
```

```
R2(config)# router ospf 100
R2(config-router)# network 172.16.64.1 0.0.0.0 area 1
R2(config-router)# area 1 range 172.16.64.0 255.255.224.0
```

Configuration of OSPF in NBMA

Understanding OSPF Network Types ①

- OSPF recognizes these types of networks
  - **Broadcast**
    - Spreading of multicast and broadcast is done by network itself. All routers on the segment could see each other (e.g. Ethernet).
  - **Non-Broadcast Multi-Access**
    - Spreading of multicast and broadcast is done by sender. Full connectivity of neighbor routers is required, otherwise there could exist problems in reachability of next-hop routers (e.g. ATM, X.25)
    - It could be used in Frame Relay with proper IP/DLCI mapping
  - **Point-to-Point**
    - Broadcast and multicast as well as full connectivity is no problem
  - **Point-to-Multipoint**
    - In principle it’s collection of separate Point-to-Point links over the one physical interface (e.g. Frame Relay)

Understanding OSPF Network Types ②

- Network type influences:
  - weather it is necessary to configure neighbors manually or weather they are discovered automatically by Hello packets
  - how are relationships and connections between routers modeled in LSDB
- **Broadcast and Nonbroadcast**
  - Broadcast: neighbors discovered automatically
  - Nonbroadcast: neighbors must be configured manually
- **Point-to-* and Multi-access**
  - Point-to-*: LSDB models connection between each pair of routers as separate LSA1
  - Multiaccess: LSDB models shared network as one node interconnecting all routers via LSA2

Operational Modes over NBMA

- RFC 2328 specifies for NBMA this modes:
  - Nonbroadcast (NBMA)
  - Point-to-Multipoint
- Cisco implements additional proprietary modes:
  - Point-to-Multipoint Nonbroadcast
  - Broadcast
  - Point-to-Point
- Operational mode is configured on per-interface basis:

```
Router(config-if)#
  ip ospf network [{broadcast | non-broadcast |
  point-to-multipoint [non-broadcast] | point-to-point}]
```

RFC-compatible Mode Nonbroadcast

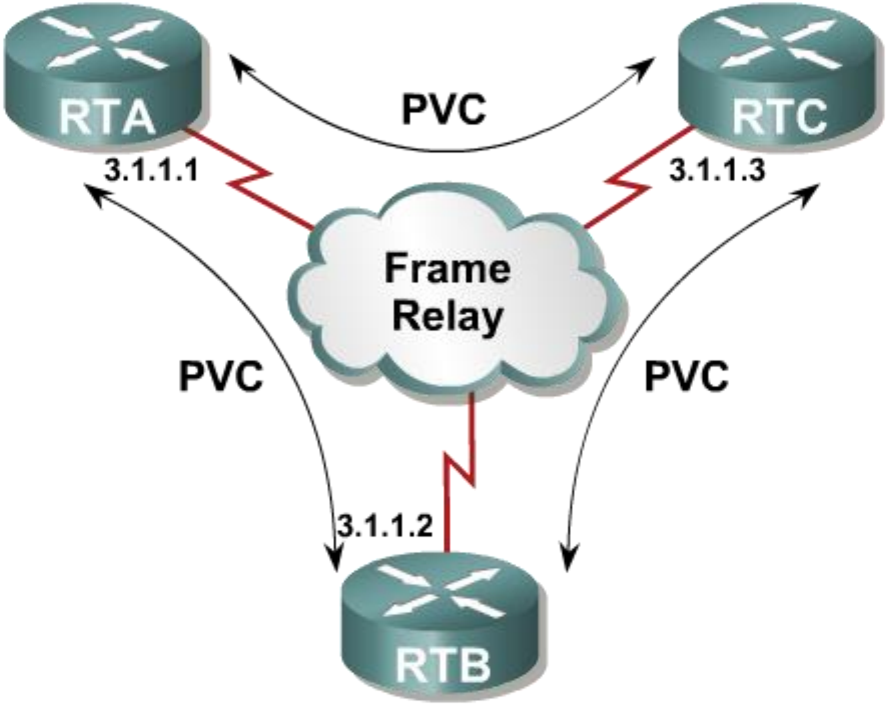

- All routers are interconnected by shared network
- Neighbors must be configured manually – OSPF does not even try to send multicasts
- DR and BDR are elected
- All routers must have full connectivity or IP/DLCI mapping must be tuned to accomplish it
- Mode suitable for full-mesh networks

```
RTB(config-if)# ip ospf network non-broadcast

RTB(config-router)# network 3.1.1.0 0.0.0.255 area 0
RTB(config-router)# neighbor 3.1.1.1
RTB(config-router)# neighbor 3.1.1.3
```

RFC-compatible Mode Point-to-Multipoint

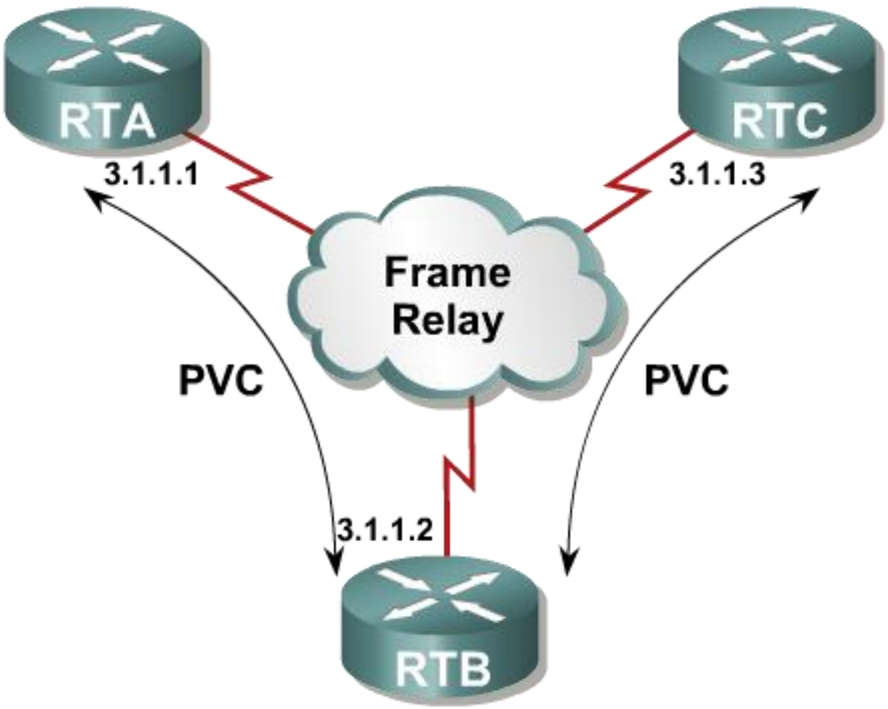

- All routers are interconnected by shared network
- Neighbors discover each other dynamically via multicast Hello packets
- DLCI mapping must have flag broadcast configured
- DR and BDR are not elected
- Usually used in partial-mesh or hub-and-spoke topologies

```
RTB(config-if)# ip ospf network point-to-multipoint

RTB(config-router)# network 3.1.1.0 0.0.0.255 area 0
```

Cisco Mode Point-to-Multipoint Non-broadcast

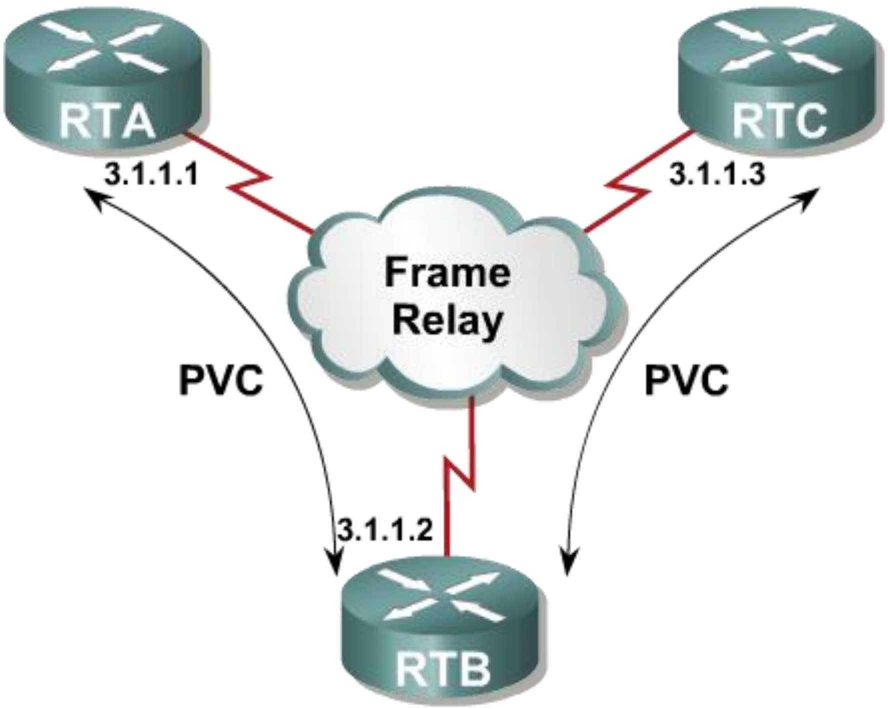

- All routers are interconnected by shared network
- Neighbors must be configured manually just like in Non-broadcast mode
- DR and BDR are not elected just like in Point-to-Multipoint Mode
- Used in very rare cases where is not possible or suitable to discover neighbors automatically (e.g. ATM or CLIP)

```
RTB(config-if)#ip ospf network point-to-multipoint non-broadcast

RTB(config-router)# network 3.1.1.0 0.0.0.255 area 0
RTB(config-router)# neighbor 3.1.1.1 cost 10
RTB(config-router)# neighbor 3.1.1.3 cost 20
```

Cisco Mode Broadcast

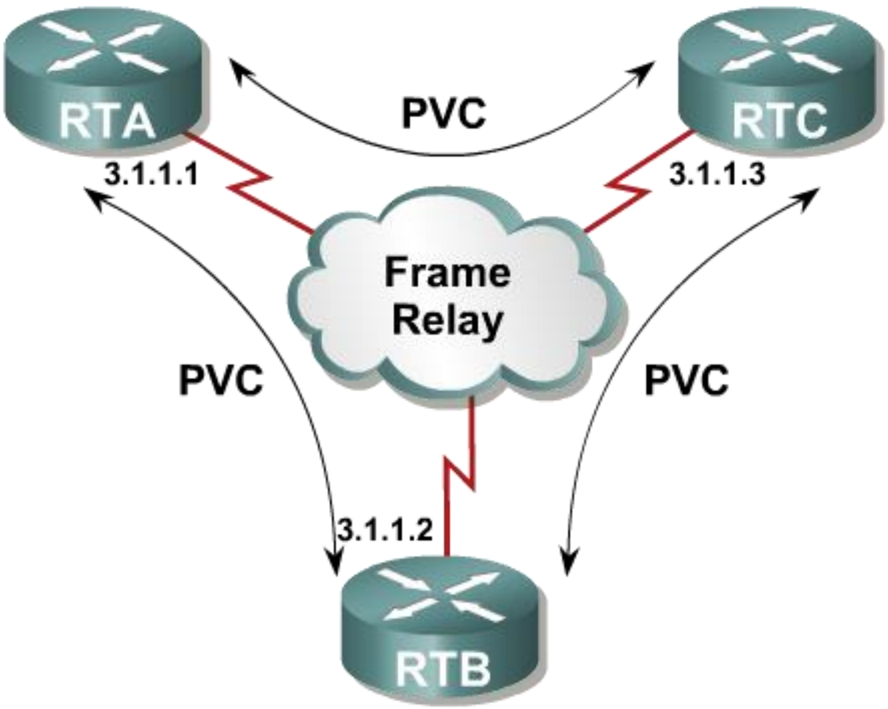

- It forces WAN interface to behave like LAN
- All routers are interconnected by shared network
- Neighbors discover each other dynamically via multicast Hello packets
- DLCI mapping must have flag broadcast configured
- DR and BDR are elected
- Requires full-mesh topology

```
RTB(config-if)# ip ospf network broadcast

RTB(config-router)# network 3.1.1.0 0.0.0.255 area 0
```

Cisco Mode Point-to-Point

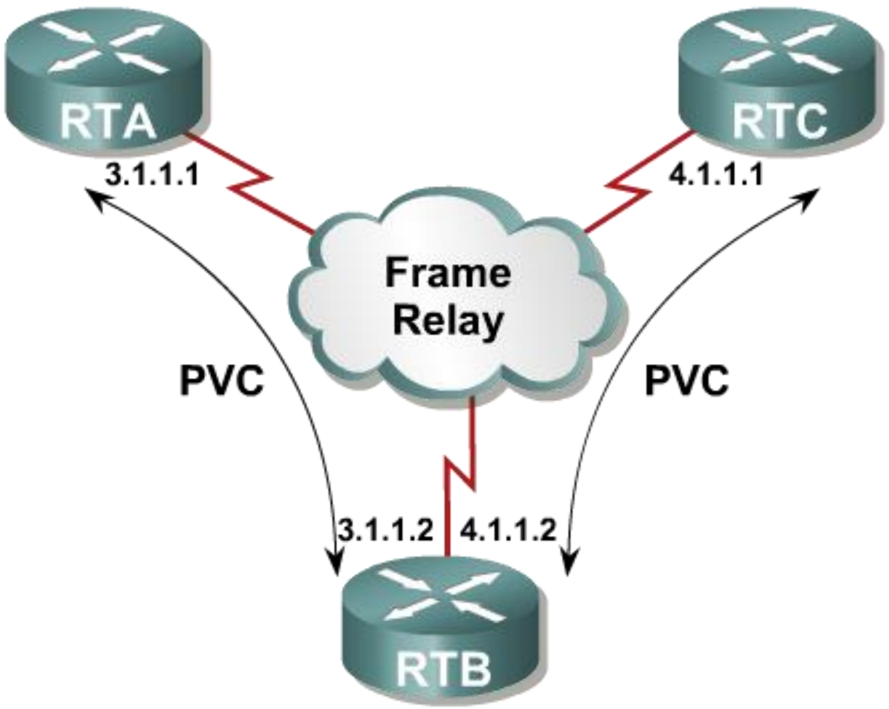

- Every pair of subinterfaces is connected via VC with separate IP network
- DR and BDR are not elected
- Used only in Point-to-Point topologies

```
RTB(config)# interface serial 0/0.1
RTB(config-subif)# ip address 3.1.1.2 255.255.255.0
RTB(config-subif)# interface serial 0/0.2
RTB(config-subif)# ip address 4.1.1.2 255.255.255.0

RTB(config-router)# network 3.1.1.0 0.0.0.255 area 0
RTB(config-router)# network 4.1.1.0 0.0.0.255 area 0
```

Neighbor Command

```
Router(config-router)#
  neighbor ip-address [poll-interval number] [cost number]
  [database-filter all]
```

- Above command statically define neighbor for OSPF routing process
  - **`poll-interval`**: Specifies interval of sending Hello packets in case that neighbor is down (after Dead Interval timed out)
  - **`cost`**: Argument valid only for Point-to-Multipoint links to define metric for target neighbor because router can’t differentiate on one physical interface
  - **`database-filter all`**: Filter sending of all types of LSAs to neighbor
- *Memo technical help: “Nonbroadcast Needs Neighbors”*

Priority Problem

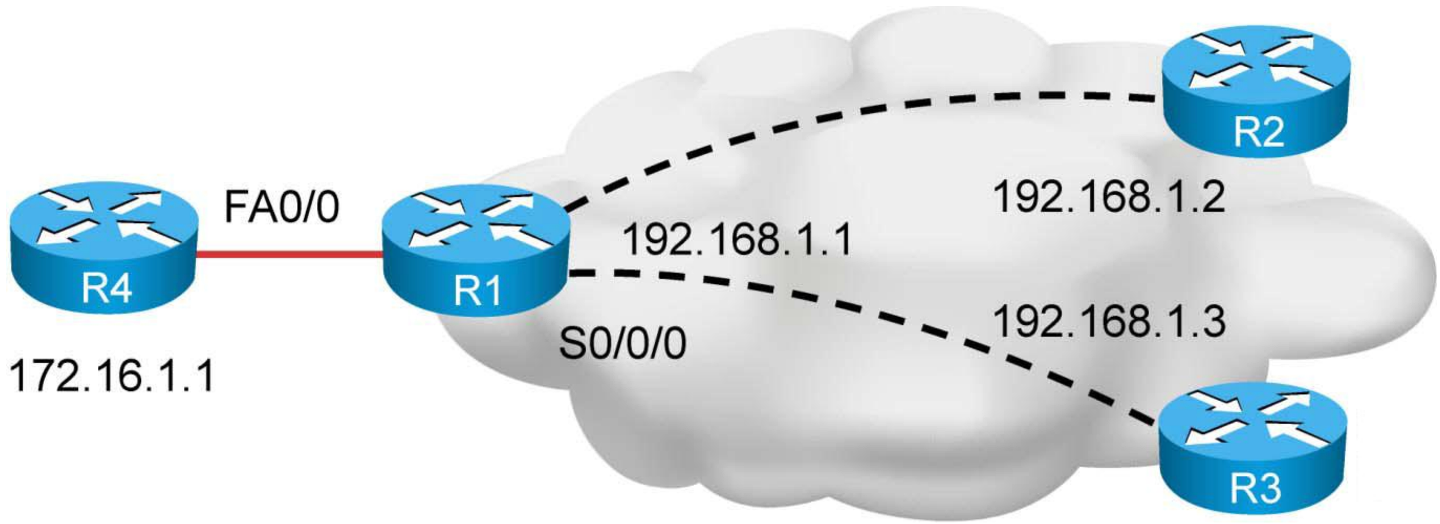

- It is necessary to modify (even default) priority in case of hub-and-spoke topology where DR and BDR are elected
- Spokes SHOULD be always configured with priority 0
  - Otherwise they could consider themselves as DR or BDR on segment because they can not see no other router on segment except the hub

```
Router(config-if)# ip ospf priority 0
```

Conclusion on NBMA Network Types ①

- **Multi-access/Point-to-(Multi)Point**
  - MA: DR and BDR are elected, more memory intensive because of LSA2 pseudonode
  - P-t-(M)P: DR and BDR are not elected, also more memory intensive
- **Broadcast/Non-Broadcast**
  - B: Neighbors discovered automatically
  - NB: Neighbors configured with command `neighbor`
- Suitability for use:
  - BMA: Ethernet, WiFi
  - NBMA: ATM, X.25, not very good for FR
  - PtMP: FR
  - PtMP NB: Rare cases
  - PtP: Only for point-to-point links (tunnels or serial interfaces)

Conclusion on NBMA Network Types ②

OSPF Mode	NBMA Preferred Topology	Subnet Address	Hello Timer	Adjacency	RFC or Cisco
Broadcast	Full or Partial Mesh	Same	10 sec	Automatic DR/BDR Elected	Cisco
Nonbroadcast (NBMA)	Full or Partial Mesh	Same	30 sec	Manual Configuration DR/BDR Elected	RFC
Point-to-Multipoint	Partial-Mesh or Star	Same	30 sec	Automatic No DR/BDR	RFC
Point-to-Multipoint Nonbroadcast	Partial-Mesh or Star	Same	30 sec	Manual Configuration No DR/BDR	Cisco
Point-to-Point	Partial-Mesh or Star, Using Subinterface	Different for Each Subinterface	10 sec	Automatic No DR/BDR	Cisco

Virtual Link

Virtual Link

- Temporary OSPF workaround **virtual link** is usually used when:
  1. area doesn't have ABR directly connected to backbone and we want to connect this **discontiguous area** to area 0
  2. backbone itself is partitioned
- A logical connection is built between router A and router B – no tunnel just OSPF messages are unicasted

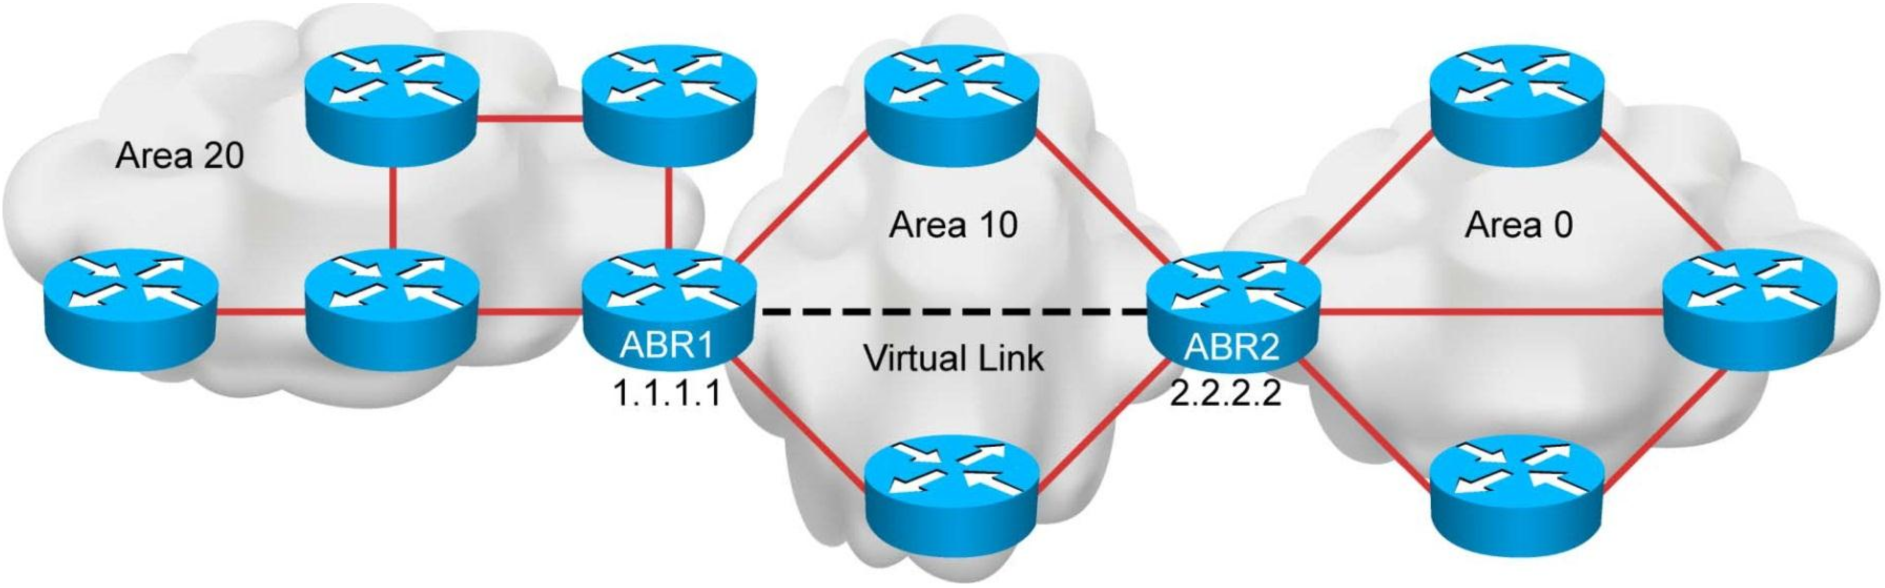

Virtual Link: LSAs

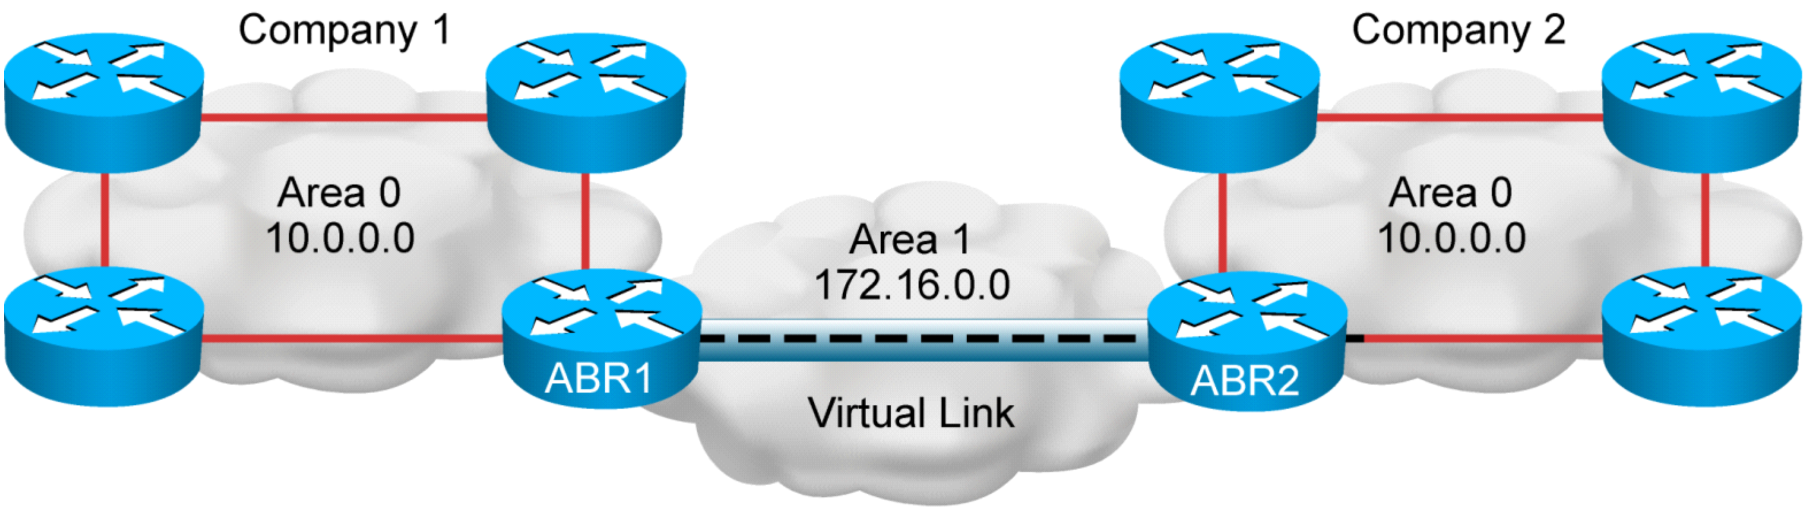

- LSAs are usually regenerated every 30 minutes and expire after max age of 60 minute
- LSA carried through virtual link has set DoNotAge flag hence they are not aging

Virtual Link: Configuration

- Virtual link configuration snippet:

```
Router(config-router)#
  area area-id virtual-link router-id [authentication
  [message-digest | null]] [hello-interval seconds]
  [retransmit-interval seconds] [transmit-delay seconds]
  [dead-interval seconds] [[authentication-key key] |
  [message-digest-key key-id md5 key]]
```

  - Both endpoints of virtual link are identified by theirs RID and they have to be in same area
    - More then good idea to use `router-id`
  - Area through which Virtual link is traversing MUST NOT be Stub area
  - Both endpoints become members of area 0 and serve ABR functionality
    - Be aware of authentication in area 0, it applies also on virtual link!
  - Because virtual link does not have interface, all relevant parameters are configured in one command ☹

Virtual Link: Example

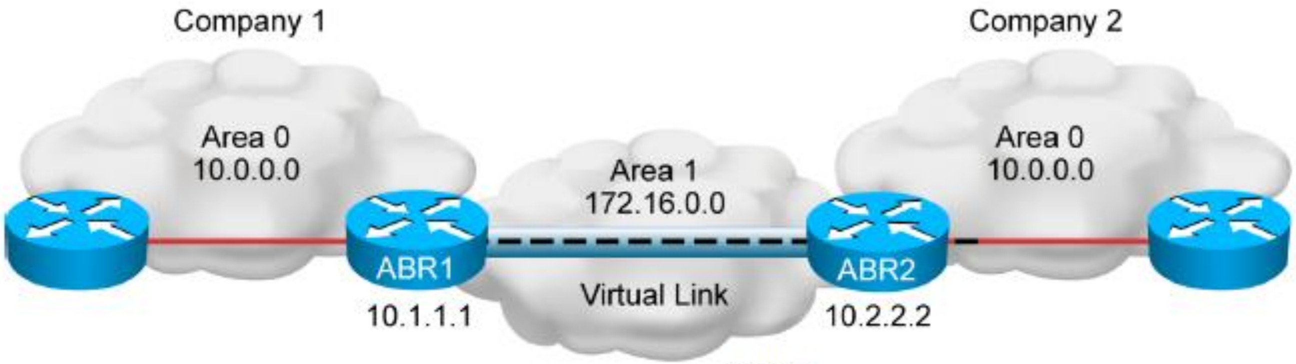

ABR1#

```
router ospf 100
 network 172.16.0.0 0.0.255.255 area 1
 network 10.0.0.0 0.255.255.255 area 0
 area 1 virtual-link 10.2.2.2
```

ABR2#

```
router ospf 100
 network 172.16.0.0 0.0.255.255 area 1
 network 10.0.0.0 0.255.255.255 area 0
 area 1 virtual-link 10.1.1.1
```

```
ABR1#show ip ospf virtual-links

Virtual Link OSPF_VL0 to router 10.2.2.2 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 1, via interface Serial0/0/1, Cost of using 781
  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
    Adjacency State FULL (Hello suppressed)
    Index 1/2, retransmission queue length 0, number of retransmission 1
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 1, maximum is 1
    Last retransmission scan time is 0 msec, maximum is 0 msec
```

LSDB Overload Protection

Overload Protection

- Excessive LSAs generated by other routers can drain local router resources

```
Router(config-router)#
  max-lsa maximum-number [threshold-percentage]
  [warning-only] [ignore-time minutes]
  [ignore-count count-number] [reset-time minutes]
```

Parameter	Description
maximum-number	Maximum number of received LSAs which will be processed to LSDB.
threshold-percentage	(Optional) IF this percentage of maximum number of processed LSA is exceeded THEN warning message is generated. By default 75 percent.
warning-only	(Optional) IF maximal number is exceeded THEN router generates only warning but doesn’t transit to Ignore stavu. By default it’s turn off.
ignore-time *minutes*	(Optional) Interval during router will ignore LSAs from neighbor routers after maximum number of LSAs is exceeded. By default 5 minutes.
ignore-count *count-number*	(Optional) How many times router could transit to Ignore state before OSPF process will be shutdown. By default 5 times.
reset-time *minutes*	(Optional) Timeout after counter of transits to Ignore state is reset. By default 10 minutes.

Useful Command

```
show ip protocols
show ip ospf
show ip ospf neighbor
show ip ospf database
show ip route ospf
show ip ospf interface
debug ip ospf events
debug ip ospf adjacency
debug ip ospf packet
```

Link-State Database

- Command **show ip ospf database** shows whole content of LSDB of router

```
RTC# show ip ospf database

            OSPF Router with ID (192.168.1.253) (Process ID 3)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
192.168.1.249   192.168.1.249   1705        0x80000005 0x00D5B0 5
192.168.1.253   192.168.1.253   1578        0x80000006 0x009F91 5
```

- After 30 minutes LSAs are automatically updated

```
RTC# show ip ospf database

            OSPF Router with ID (192.168.1.253) (Process ID 3)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
192.168.1.249   192.168.1.249   106         0x80000006 0x00D3B1 5
192.168.1.253   192.168.1.253   58          0x80000007 0x009D92 5
```

Reference Literature

- OSPF Design Guide
  - http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094e9e.shtml
- Configuring OSPF
  - http://www.cisco.com/en/US/docs/ios/12_0/np1/configuration/guide/1cospf.html
- OSPF Technology Page
  - http://www.cisco.com/en/US/tech/tk365/tk480/tsd_technology_support_sub-protocol_home.html
- OSPF Authentication
  - http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a0080094069.shtml

Slides adapted by Vladimír Veselý partially from official course materials but the most of the credit goes to CCIE#23527 Ing. Peter Palúch, Ph.D.

Last update: 2012-09-07