

www.howtonetwork.net

CCNA: OSPF CHEAT SHEET

Key Characteristics

Type: Link State

Algorithm: Dijkstra's (Shortest Path First) Algorithm

Standard: RFC 2328

Administrative Distance: 110

Metric: Cost

Protocol/Protocol Number: IP/89

Authentication: Yes (MD5 and Plain Text)

Supports VLSM and Route Summarization

Support for IPv6 (RFC 2740)

Fast Convergence

Metric Calculation

Cost = 100 Mbps / Link Speed

OSPF cost can be modified in three ways:

1. (config-if)#ip ospf cost
2. (config-if)#bandwidth
3. (config-router)#auto-costreference-bandwidth

Neighbor Process

To become neighbors OSPF routers must agree on the following parameters

1. Area ID
2. Same subnet
3. Authentication (if used)
4. Hello Interval and Dead Interval
5. Area Type (Stub, NSSA)
6. Router IDs must be unique

OSPF Neighbor States

Down: Previously known neighbor has failed

Init: an interim state in which Hello has been heard from the neighbor but that Hello does not list the local router's RID

Two-way: the neighbor has sent a Hello that lists the local router's RID in the list of seen routers

Full: Both routers complete the database exchange process and have identical LSDB. Fully adjacent

Areas

OSPF runs SPF algorithm and requires a lot of processing power and memory. If the size of network is too large this could cause slower convergence and can lead to following problems:

1. more memory is required to maintain the link state database
2. more processing power is required to process the link state database
3. the links state database grows exponentially with the size of OSPF domain
4. a single change in network topology (for example: link up/down) would trigger all routers to re-run the SPF (again) to calculate the shortest path

To cope with these problems, areas are configured. There are two basic types:

1. Backbone Area or Area 0: All other area must be connected to area 0
2. Non-backbone Area: any other area with area-id other than zero

Timers

Hellos are sent to multicast address: 224.0.0.5 (ALLSPFRouters)

Broadcast Multi-access = 10 seconds

Point-to-Point & NBMA = 30 seconds

Dead Timer = Four Times the hello interval

Broadcast = 40 seconds

NBMA and P2P = 120 seconds

To change hello and dead intervals use the command 'config-if)# ip ospf hello-interval seconds' and 'config-if)# ip ospf dead-interval seconds'

Router ID (RID)

Router ID must be configured before an OSPF process could be started.

Cisco Routers uses the following criteria to select a router ID:

1. RID configured with "router-id" command
2. If manual RID not configured, select the highest number IP address on any loopback interface in "up/up" state
3. If loopback interfaces not configured, select the highest number IP address on any non-loopback interface in "up/up" state

Router Types

Internal Router: whose (all) interfaces resides within the same area

Backbone Router: A router that resides in the backbone area

Area Border Router: an ABR connect two or more Areas

ASBR: Autonomous System Boundary Router or an ASBR connects an external routing domain to an OSPF routing domain

Route Types and Preference

Intra-Area Routes: A route to a network in the same area as the router. Denoted by "O" in the routing table.

Inter-Area Routes: A route to a network in another area as the router. Denoted by "O IA" in the routing table

External Route: A route to network that is external to the OSPF routing domain. Denoted by 'E1' or 'E2' in the routing table.

Routes Preference:

1. Intra-Area (O) > Inter-Area (O IA)
2. Inter-Area (O IA) > External Type-1 (E1)
3. External Type-1 (E1) > External Type-2 (E2)

Designated Router (DR) / Backup DR (BDR) Election

There are two problems with multi-access networks:

1. For "N" routers, it requires " $N(N-1)/2$ " adjacencies
2. Flooding of this excess LSAs would be chaotic itself for the network

DR/BDR addresses the challenge of adjacency creation and LSA flooding on multi-access networks only

No election on P2P and P2MP network type

The following criteria is used for DR/BDR election:

1. Router with highest interface priority is elected as DR
2. Any other router with second highest priority is elected as BDR
3. If priority is equal, highest RID is used as tie-breaker
4. The DR/BDR election is held between two or more neighbors who reach the TWO-WAY state

The priority ranges from 0-to-255 and default value is 1

Priority of 0 means that router will not take part in DR and BDR election

DR is never preempted even if a router with better priority is present. Manual reset is required for preemption

If a router becomes active and it checks for an active DR and BDR on the network. If there already is an active DR and BDR on the segment, the new router simply accepts them. If there is not, then an election is held for DR/BDR selection

After the DR/BDR have been elected, the other router known as DROthers establish adjacencies with DR and BDR only

Neighbors are still tracked on multicast address: 224.0.0.5 but DROthers multicast updates to AllDRRouters address: 224.0.0.6.

Only DR and BDR listen to this address and DR in-turn flood updates to DROthers on 224.0.0.5

DR/BDR is property of a router's interface not the router itself

Virtual Links

It is link through non-backbone area to backbone area.

Used to connect:

1. an area to backbone area through non-backbone area
2. a partitioned backbone area through non-backbone area

Rules:

1. A virtual link can only be configured between ABRs
2. The transit area must have full routing information and it cannot be stub

• Configuration Example

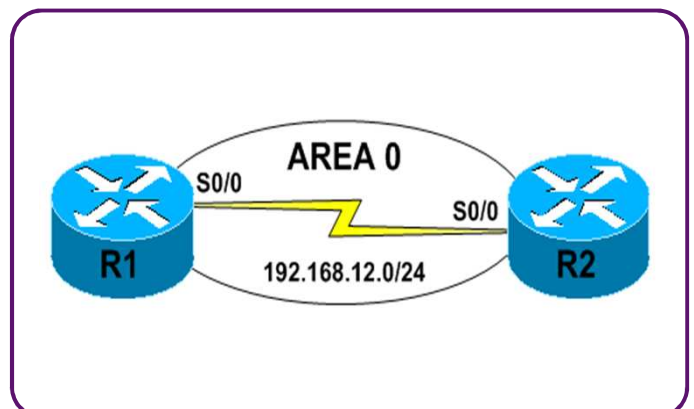
Single Area

Router R1:

```
interface loopback 0
ip address 10.1.1.1 255.255.255.255
!
interface serial 0/0
ip address 192.168.12.1 255.255.255.0
!
router ospf 100
router-id 1.1.1.1
network 192.168.12.0 0.0.0.255 area 0
network 10.1.1.1 0.0.0.0 area 0
```

Router R2:

```
interface loopback 0
ip address 10.2.2.2 255.255.255.255
!
interface serial 0/0
ip address 192.168.12.2 255.255.255.0
!
router ospf 100
router-id 2.2.2.2
network 192.168.12.0 0.0.0.255 area 0
network 10.2.2.2 0.0.0.0 area 0
```



```
R2# show ip route | begin Gateway
Gateway of last resort is not set
```

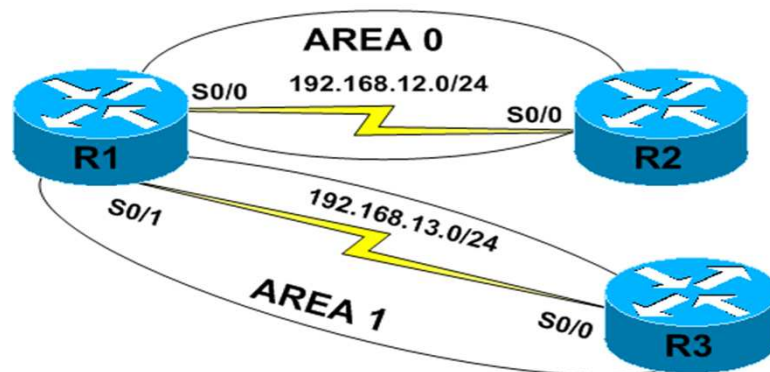
```
C 192.168.12.0/24 is directly connected, Serial0/0
```

```
10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
```

```
C 10.2.2.0/24 is directly connected, Loopback0
```

```
O 10.1.1.1/32 [110/65] via 192.168.12.1, 00:00:02, Serial0/0
```

• Configuration Example



Multi-Area

```

Router R1:
interface loopback 0
ip address 10.1.1.1 255.255.255.255
!
interface serial 0/0
ip address 192.168.12.1 255.255.255.0
!
interface serial 0/1
ip address 192.168.13.1 255.255.255.0
!
router ospf 100
router-id 1.1.1.1
network 192.168.12.1 0.0.0.0 area 0
network 192.168.13.1 0.0.0.0 area 1
network 10.1.1.1 0.0.0.0 area 0
  
```

```

Router R2:
interface loopback 0
ip address 10.2.2.2 255.255.255.255
!
interface serial 0/0
ip address 192.168.12.2 255.255.255.0
!
router ospf 100
router-id 2.2.2.2
network 192.168.12.2 0.0.0.0 area 0
network 10.2.2.2 0.0.0.0 area 0
  
```

```

Router R3:
interface serial 0/0
ip address 192.168.13.3 255.255.255.0
!
interface loopback 0
ip address 10.3.3.3 255.255.255.255
!
router ospf 100
router-id 3.3.3.3
network 192.168.13.3 0.0.0.0 area 1
network 10.1.1.3 0.0.0.0 area 1
  
```

```

R2# show ip route | begin Gateway
Gateway of last resort is not set
  
```

```

C 192.168.12.0/24 is directly connected, Serial0/0
O IA 192.168.13.0/24 [110/128] via 192.168.12.1, 00:00:03, Serial0/0
  10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O IA 10.3.3.3/32 [110/129] via 192.168.12.1, 00:00:12, Serial0/0
C 10.2.2.0/24 is directly connected, Loopback0
O 10.1.1.1/32 [110/65] via 192.168.12.1, 00:00:12, Serial0/0
  
```

OSPF Troubleshooting Command

1. show ip protocols
2. show ip ospf [<process-id>]
3. show ip route [ospf]
4. show ip ospf interface [brief | <interface-id>]
5. show ip ospf neighbor
6. show ip ospf database
7. debug ip ospf [hello | adjacency | events]