OSPFv2/v3 (IPv4/IPv6)

PROTOCOL HIGHLIGHTS

Type	Link State (Internal)
Algorithm	Dijkstra's Shortest Path First (SPF)
Standard	IETF RFC — OSPFv2 2328, OSPFv3 5340
Metric	Cost (10^8 ÷ link bandwidth in bits per second)
Timers	10 / 30 s hello · 40 / 120 s dead
Administrative distance	110
Network protocols	IPv4, IPv6
Authentication	Plain text, Message-Digest 5 (MD5)
Multicast	`224.0.0.5 / .6` · `FF02::5 / ::6` (all / DR)

OSPF METRIC

$$\text{cost} = \frac{\text{100,000,000 bits per second (reference bandwidth)}}{\text{outbound interface bandwidth (in bits)}}$$

T1 (1.544)	64
10 Mb	10
100 Mb	1
1 Gb	1
10 Gb	1

- Values in the metric are those seen with **`show interface`**
- Only values on egress interfaces are considered in the calculation
- No floating point maths — all calculated values are rounded down
- Default reference bandwidth is 10^8 bits per second

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

bw from <1–4294967>, 100 is the default

- **Type 1** (`O E1 / O N1`) — redistributed cost *plus* the cost to the ASBR
- **Type 2** (`O E2 / O N2`) — redistributed cost only, fixed (default)

OSPF BASIC CONFIGURATION

IPv4 configuration steps

```
R1(config)#router ospf 1
    Locally significant process ID, <1–65535>
R1(config-router)#default-information originate
    Distribute a default route into OSPF
R1(config-router)#network 172.16.10.0 0.0.0.3 area 1
    Enable OSPF on interfaces, advertise networks, put the interface into an area
R1(config-router)#passive-interface Ethernet 0
    Suppress sending and receiving OSPF packets
R1(config-router)#area 1 stub [no-summary]
    Configure OSPF stub features — see terminology
R1(config-router)#router-id 1.1.1.1
    Statically define the router ID in 32-bit dotted decimal format
```

IPv6 configuration steps

```
R1(config)#ipv6 router ospf 1
    Locally significant process ID, <1–65535>
R1(config-router)#router-id 1.1.1.1
    OSPF router ID in 32-bit dotted decimal format
R1(config-router)#passive-interface Ethernet 0
    Suppress sending and receiving OSPF packets
R1(config-router)#area n.n.n.n stub [no-summary]
    Configure OSPF stub area features — see terminology
R1(config)#interface FastEthernet 0/0
R1(config-if)#ipv6 ospf 1 area n.n.n.n
    Enable OSPF on the interface, advertise the network, set the area
R1(config-router)#default-information originate
    Distribute a default route into OSPF
```

HEADER FORMAT

`OSPFv2 = 0x02` (IPv4)
`OSPFv3 = 0x03` (IPv6)

1 Hello 4 LS Update
2 DB Descriptor 5 LS Ack
3 LS Request

Applies to the entire OSPF packet, excluding the IP header

Version	8	Type	8	Packet Length	16
Router ID					32
Area ID					32
Checksum	16			Authentication Type	16
Authentication					64
Data — Hello, DBD, LS Request, LS Update, LS Ack					

The OSPFv2 header is **24 bytes** — version, type, packet length, router ID, area ID, checksum, authentication type, and 64 bits of authentication. RFC 2328 §A.3.1.

FRAME PAYLOAD

Frame header	IPv4 header (protocol 89) IPv6 header (next header 89)	OSPF packet payload	Frame footer

OSPF TERMINOLOGY

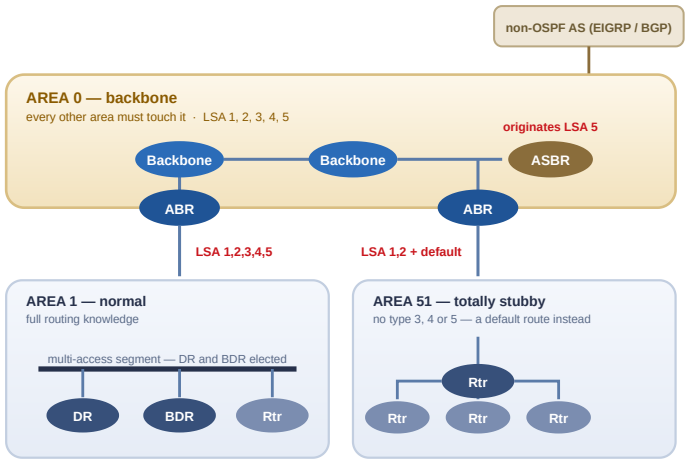

Backbone area (0) — the transit area interconnecting all non-backbone areas.

Backbone router — a router with all interfaces in area 0.

Normal area — any non-backbone area not configured as a stub area.

Area border router (ABR) — a router with interfaces in the backbone and in other areas.

Autonomous system boundary router (ASBR) — a router with an interface in OSPF and one in a non-OSPF autonomous system, performing route redistribution.

Designated router (DR) — elected on multi-access networks to manage adjacencies. Election is on highest priority (0–255, 1 default), then highest router ID. The DR forms full adjacencies with all neighbors; the **BDR** backs it up.

Stub area — allows LSA type 3 and a default route only. ABR and area routers are configured with **`area n.n.n.n stub`**.

Not-so-stubby area (NSSA) — a stub area with an ASBR, which allows type 7 LSAs.

Totally stubby area — allows only a default route from the ABR. The ABR is configured with **`area n.n.n.n stub no-summary`**, internal routers with **`area n.n.n.n stub`**.

Totally not-so-stubby area — a totally stubby area with an ASBR.

OSPFv2/v3 (IPv4/IPv6)

OSPF INTERFACE CONFIGURATIONS

```
R1(config-if)#bandwidth <1-10000000>
```
Adjusts the interface bandwidth in kbps, used in the metric calculation

```
R1(config-if)#ip ospf cost <1-65535>
```
Statically configure the interface cost, ignoring the bandwidth calculation

```
R1(config-if)#ip ospf hello-interval <1-65535>
```
Configure the hello time in seconds

```
R1(config-if)#ip ospf dead-interval [<1-65535> | minimal]
```
Configure the dead time in seconds, or `minimal` for one second

```
R1(config-if)#ip ospf network [broadcast | non-broadcast |
    point-to-multipoint | point-to-point]
```
Change how OSPF interprets the network type of the interface

```
R1(config-if)#ip ospf priority <0-255>
```
Manually set the DR priority of the interface — 1 is the default, 0 means the router will never become DR or BDR

IPv6 uses the same options under **`ipv6 ospf`** — e.g. **`ipv6 ospf cost`**, **`ipv6 ospf priority`**.

OSPF AUTHENTICATION (IPV4)

```
R1(config-router)#area <1-65535> authentication [message-digest]
```
Sets plain text or MD5 authentication for the specified area

```
R1(config-if)#ip ospf authentication-key WORD
```
Sets the plain text key used for neighbor packet exchange

```
R1(config-if)#ip ospf message-digest-key 1 md5 WORD
```
Sets the MD5 key ID and key string used for packet exchange

Type 0 — null. No authentication in the packet header. The default.

Type 1 — plain text. Clear-text passwords, readable on the wire.

Type 2 — MD5. Cryptographic. The key ID and the key must match on neighbors.

OSPF VERIFICATION COMMANDS

```
R1#show ip ospf neighbor
Neighbor ID  Pri State           Dead Time Address    Interface
2.2.2.2        1 FULL/BDR        00:00:34  10.1.1.2   FastEthernet0/0
3.3.3.3        1 FULL/DROTHER    00:00:31  10.1.1.3   FastEthernet0/0
```

IPv6 — **`show ipv6 ospf neighbor`**

1. The neighbor's **router ID** — not its interface address
2. Its OSPF priority on that link. 0 means it will never be DR or BDR
3. Adjacency state, then its role on the segment after the slash
4. Time until the neighbor is declared down if no hello arrives
5. The neighbor's address on this link — the next hop
6. The local interface it is reachable through

```
R1#show ip protocols
Routing Protocol is "ospf 1"          process ID
  Router ID 1.1.1.1                   router ID
  Number of areas in this router is 2. 2 normal 0 stub 0 nssa
  Maximum path: 4                     load balance paths
  Routing for Networks:
    10.1.1.0 0.0.0.255 area 0         network statements
    172.16.10.0 0.0.0.3 area 1
  Routing Information Sources:
    Gateway     Distance    Last Update
    2.2.2.2          110    00:04:21
  Distance: (default is 110)          AD
```

IPv6 — **`show ipv6 protocols`**

INTERFACE DETAIL AND THE LINK-STATE DATABASE

```
R1#show ip ospf interface Fa0/0
FastEthernet0/0 is up, line protocol is up
  Internet Address 10.1.1.1/24, Area 0
  Process ID 1, Router ID 1.1.1.1,
    Network Type BROADCAST, Cost: 1    cost
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 1.1.1.1, addr 10.1.1.1
  Backup Designated router (ID) 2.2.2.2, addr 10.1.1.2
  Hello 10, Dead 40, Wait 40, Retransmit 5    timers
  Neighbor Count is 2, Adjacent neighbor count is 2
```

Hello and dead must match on both ends or the adjacency never forms. IPv6 — **`show ipv6 ospf interface`**

```
R1#show ip ospf database
    OSPF Router with ID (1.1.1.1) (Process ID 1)

        Router Link States (Area 0)       type 1
Link ID    ADV Router   Age  Seq#       Checksum Link count
1.1.1.1    1.1.1.1      421  0x80000003 0x00A1B2 2
2.2.2.2    2.2.2.2      418  0x80000004 0x0043C1 3

        Net Link States (Area 0)          type 2
Link ID    ADV Router   Age  Seq#       Checksum
10.1.1.1   1.1.1.1      421  0x80000001 0x009C41
```

IPv6 — **`show ipv6 ospf database`**

ADJACENCY STATES AND LSA TYPES

Neighbor states, in the order they happen

State	What it means
Down	No hello received
Attempt	NBMA only — unicast hello sent, no reply
Init	Hello arrived, but our router ID was not in it
2-Way	Bidirectional. **DR and BDR elected here**
ExStart	Master/slave chosen, sequence number agreed
Exchange	Database descriptors exchanged
Loading	LS requests sent for anything missing
Full	Databases synchronised — adjacency complete

A neighbor stuck in **ExStart** is very often an MTU mismatch. Stuck in **2-Way** between DROTHERs is normal, not a fault.

LSA types — the ones on the area diagram

Type	Name	Originated by, and how far it floods
1	**Router**	Every router, for its own links. Its area only
2	**Network**	The DR on a multi-access segment. Its area only
3	**Summary (network)**	An ABR, networks from one area into another
4	**Summary (ASBR)**	An ABR, advertising the route to an ASBR
5	**External**	An ASBR, redistributed routes. Whole AS
7	**NSSA external**	An ASBR in an NSSA. ABR converts it to type 5

A stub area blocks types 4 and 5. A totally stubby area also blocks type 3, leaving a default route from the ABR.

BASIC OSPF DEBUG COMMANDS

```
R1#debug ip ospf ?
  adj              OSPF adjacency events
  database-timer   OSPF database timer
  events           OSPF events
  flood            OSPF flooding
  hello            OSPF hello events
  lsa-generation   OSPF LSA generation
  packet           OSPF packets
  retransmission   OSPF retransmission events
  spf              OSPF SPF
```

- **`adj`** — reach for this first. Shows adjacencies forming and, more usefully, why they stop
- **`hello`** — mismatched timers, area IDs, authentication or subnet masks all show up here
- **`spf`** — frequent runs mean something is flapping
- IPv6 — **`debug ipv6 ospf`**, same options

On a production router use **`debug ip ospf adj`**, not a bare debug, and keep **`undebug all`** ready. Better still, turn on **`logging buffered`** and read it with **`show logging`**.