

Introduction to OSPF

ISP Workshops



These materials are licensed under the Creative Commons Attribution-NonCommercial 4.0 International license (<http://creativecommons.org/licenses/by-nc/4.0/>)

Last updated 23rd June 2021

Acknowledgements

- This material originated from the Cisco ISP/IXP Workshop Programme developed by Philip Smith & Barry Greene
- Use of these materials is encouraged as long as the source is fully acknowledged and this notice remains in place
- Bug fixes and improvements are welcomed
 - Please email *workshop (at) bgp4all.com*

Philip Smith

BGP Videos

- NSRC has made a video recording of this presentation, as part of a library of BGP videos for the whole community to use:
 - https://learn.nsrc.org/bgp#intro_to_routing

The screenshot shows the NSRC (Network Startup Resource Center) website. The navigation bar includes 'Home', 'About', 'BGP for All', 'perfSONAR', 'ScienceDMZ', 'FedIdM', and 'Contact Us'. A search bar is located on the right. The main content area is divided into three columns:

- BGP for All:** A section with a description of BGP and a 'Video Topics' sidebar. The sidebar has four items: 'BGP for All' (orange), 'perfSONAR' (yellow), 'ScienceDMZ' (blue), and 'FedIdM' (green).
- Introduction to Routing:** A list of 15 topics including 'Internet Routing', 'Routing Protocols', 'Introduction to IS-IS', 'IS-IS Levels', 'IS-IS Adjacencies', 'Best Configuration Practices for IS-IS on Cisco IOS', 'IS-IS Authentication, Default Routes and IPv6', 'Introduction to OSPF', 'OSPF Areas', 'OSPF Adjacencies', 'Best Configuration Practices for OSPF on Cisco IOS', 'OSPF Authentication, Default Routes and IPv6', 'Comparing OSPF and IS-IS', 'Choosing between OSPF and IS-IS', 'Migrating from OSPF to IS-IS', 'Migration Plan', and 'Finalizing Migration'.
- Introduction to BGP:** A list of 6 topics including 'Introduction to Border Gateway Protocol', 'Transit and Peering', 'Autonomous Systems', 'How BGP works', 'Supporting Multiple Protocols', 'IBGP versus EBGP', 'Setting up EBGP', and 'Setting up IBGP'.
- BGP Case Studies:** A list of 5 case studies including 'Peering Priorities', 'Transit Provider Peering at an IXP', 'Customer Multihomed between two IXP members', 'Traffic Engineering for an ISP connected to two IXes', and 'Traffic Engineering for an ISP with two interfaces on one IX LAN'.
- Communities:** A list of 4 communities including 'Communities: RFC 1998 Traffic Engineering', 'Communities: Simplifying Traffic Engineering', 'How to Apply Communities to Originated Routes', and 'How to Use Communities for Service Identification'.

A video player is embedded in the center-right, showing a video titled 'BGP for All' with a play button and a 'Watch on YouTube' button. The video player interface includes 'Watch later', 'Share', and '1/158'.

OSPF

- Open Shortest Path First
- Open:
 - Meaning an Open Standard
 - Developed by IETF (OSPF Working Group) for IP – RFC1247
 - Current standard is OSPFv2 (RFC2328)
- Shortest Path First:
 - Edsger Dijkstra's algorithm for producing shortest path tree through a graph
 - Dijkstra, E. W. (1959). "A note on two problems in connexion with graphs". *Numerische Mathematik* **1**: 269–271
 - <http://www-m3.ma.tum.de/foswiki/pub/MN0506/WebHome/dijkstra.pdf>

OSPF

- Known as a Link State Routing Protocol
 - The other link state routing protocol is IS-IS
 - Each node in the network computes the map of connectivity through the network
- The other type of Routing Protocol is Distance Vector
 - Like EIGRP or RIP
 - Each node shares its view of the routing table with other nodes

OSPF

- Routers with OSPF enabled on them look for neighbouring routers also running OSPF
 - Using the "Hello" protocol
 - The "Hello" packet includes the subnet mask, list of known neighbours, and details such as "hello interval" and "router dead interval"
 - Hello interval – how often the router will send Hellos
 - Router dead interval – how long to wait before deciding router has disappeared
 - The values of "hello interval", "router dead interval" and subnet mask must match on both neighbours
 - When a neighbouring router responds with matching details, a neighbour relationship is formed



OSPF Neighbour Relationships

- A relationship is formed between selected neighbouring routers for the purpose of exchanging routing information
 - This is called an **ADJACENCY**
- Not every pair of neighbouring routers become adjacent
 - On multi-access networks (e.g. ethernet), only selected routers form adjacencies

OSPF Adjacencies

- ❑ Once an adjacency is formed, neighbours share their link state information
 - Information goes in a **Link State Packet (LSP)**
 - LSPs sent to a neighbour are known as **Link State Announcements (LSA)**
- ❑ New information received from neighbours is used to compute a new view of the network
- ❑ On a link failure
 - New LSPs are flooded
 - The routers recompute the routing table

OSPF across a network

- ❑ All routers across the network form neighbour relationships with their directly attached neighbours
- ❑ Each router computes the routing table
- ❑ Once each router has the same view of the network, the network has **converged**
- ❑ The IGP design for a network is crucially important to ensure **scalability** and **rapid convergence**
- ❑ Generally: **the fewer the prefixes, the faster the convergence**

OSPF Areas

- OSPF has the concept of areas
 - All networks must have an area 0, the “default” area
 - Areas are used to scale OSPF for large networks
 - There are many types of areas, to suit many different types of infrastructure and topologies
 - Most small to medium networks (up to ~300 routers) can happily use a single area

OSPF


- OSPFv2 is for IPv4
 - For carrying IPv4 prefixes only
- OSPFv3 is for IPv6
 - For carrying IPv6 prefixes only
 - Based on OSPFv2 but is specifically for IPv6
 - Documented in RFC5340
 - Is totally independent of OSPFv2
- Configuration concepts and syntax are very similar
 - (There are subtle differences/improvements)

Links in OSPF

- Two types of links in OSPF:
 - Point-to-point link
 - Only one other router on the link, forming a point-to-point adjacency
 - Multi-access network (e.g. ethernet)
 - Potential for many other routers on the network, with several other adjacencies
- OSPF in multi-access networks has optimisations to aid scaling
 - Two routers are elected to originate the LSAs for the whole multi-access network
 - Called “**Designated Router**” and “**Backup Designated Router**”
 - Other routers on the multi-access network form adjacencies with the DR and BDR


Designated Router

- There is ONE designated router per multi-access network
 - Generates network link advertisements
 - Assists in database synchronization
 - Scales OSPF for multi-access (ethernet) networks




Selecting the Designated Router

- ❑ Configured priority (per interface)
 - Configure high priority on the routers to be the DR/BDR
- ❑ Else priority determined by highest router ID
 - Router ID is a 32-bit integer
 - Set manually, otherwise derived from the loopback interface IPv4 address, otherwise the highest IPv4 address on the router




Adjacencies on multi-access networks

- DR and BDR form FULL adjacencies:
 - With each other
 - With all other routers on the multi-access network
 - Databases are synchronised
 - LSAs propagate along adjacencies



Adjacencies on multi-access networks

- Neighbour relationships between routers which are not DR or BDR are called 2WAY
 - They see each other in HELLO packets but do not exchange topology information
 - The neighbours then are not adjacent



Adjacencies: Examples

- To find adjacency state, use:

```
show ip[v6] ospf neighbor
```

- Point-to-Point link

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.10.15.236	0	FULL/ -	00:00:35	10.10.15.16	Serial1/0

- FULL: other router to DR/BDR

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.10.15.225	1	FULL/BDR	00:00:35	10.10.15.2	FastEth0/0
10.10.15.226	1	FULL/DR	00:00:35	10.10.15.3	FastEth0/0

- 2WAY: other router to other router

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.10.15.227	1	2WAY/DROTHER	00:00:35	10.10.15.4	FastEth0/0

OSPF on Cisco IOS

- ❑ Starting OSPFv2 (IPv4) in Cisco's IOS

```
router ospf 42
```

- Where "42" is the process ID

- ❑ Starting OSPFv3 (IPv6) in Cisco's IOS

```
ipv6 router ospf 42
```

- Where "42" is the process ID

- ❑ OSPF process ID is local to the router

- Allows the possibility of running multiple instances of OSPF on one router
- The process ID is not passed between routers
- The process ID must be a number – no option for alphanumeric ID
- **Some ISPs configure the process ID to be the same as their BGP Autonomous System Number**

Adding interfaces to OSPF

□ OSPF interface configuration:

- When OSPF is configured for a subnet or on an interface, the router will automatically attempt to find neighbours on that subnet or interface

```
router ospf 42
  passive-interface default
```

- ISP Best Practice is to disable this behaviour:
And then explicitly enable the interface to allow OSPF to search for neighbours as required:

```
router ospf 42
  no passive-interface Gigabit 4/0
```

OSPF on Cisco IOS

- Enabling OSPF on an interface does **two** things:
 1. Enables the Hello protocol for forming neighbour relationships and adjacencies with other routers connected to that interface
 2. Announces the interface subnet(s) into OSPF
- Care needed
 - Must avoid enabling the Hello protocol on untrusted networks
 - (e.g. those outside your Autonomous System)

OSPFv2 on Cisco IOS

- Forming neighbour relationships
 - OSPFv2 needs to be activated on the interface the neighbour relationship is desired on:

```
interface Gigabit 4/0
  ip address 192.168.1.1 255.255.255.252
  ip ospf 42 area 0
!
router ospf 42
  passive-interface default
  no passive-interface Gigabit 4/0
!
```

OSPFv3 on Cisco IOS

- Forming neighbour relationships
 - OSPFv3 needs to be activated on the interface the neighbour relationship is desired on:


```
interface Gigabit 4/0
  ipv6 address 2001:DB8:10:FE::4/64
  ipv6 ospf 42 area 0
  !
ipv6 router ospf 42
  passive-interface default
  no passive-interface Gigabit 4/0
  !
```

OSPF interface costs

- Cisco IOS sets the interface cost automatically
 - Formula used: $\text{cost} = 10^8 / \text{interface bandwidth}$
 - Which is fine for interfaces up to 100Mbps
- Many operators develop their own interface cost strategy
 - ```
ip ospf cost 100
```
  - Sets interface cost to 100
  - Care needed as the sum of costs determines the best path through the network
- OSPF chooses lowest cost path through a network
- OSPF will load balance over paths with equal cost to the same destination

# OSPF Metric Calculation


- Best path/lowest cost = 60






# OSPF Metric Calculation

- Best path/lowest cost = 60



- Equal cost paths = 70



# OSPF Neighbour Authentication

---

- Neighbour authentication is highly recommended
  - Prevents unauthorised routers from forming neighbour relationships and potentially compromising the network
- OSPFv2 – Authentication is built-in
  - There are two types:
    - Plain text password
    - MD5 hash
- OSPFv3 – uses standard IP security header
  - There are two types:
    - MD5 hash
    - SHA1

# OSPFv2 – Neighbour Authentication

---

- ❑ Configuring authentication for area 0
  - Interfaces still need the authentication key, e.g. POS4/0

```
router ospf 42
 area 0 authentication message-digest
 !
 interface Gigabit 4/0
 ip ospf message-digest-key <key-no> md5 <passwd>
 !
```

- ❑ Configuring authentication per interface:

```
interface Gigabit 4/0
 ip ospf authentication message-digest
 ip ospf message-digest-key <key-no> md5 <passwd>
 !
```

# OSPFv3 – Neighbour Authentication

---

- ❑ Configuring authentication for all interfaces in area 0
  - The key is included in the command turning on authentication for area 0:

```
ipv6 router ospf 42
 area 0 authentication ipsec spi 256 md5 <passwd>
!
```

- ❑ Configuring authentication per interface:

```
interface Gigabit 4/0
 ipv6 ospf authentication ipsec spi 256 md5 <passwd>
!
```

# Originating a Default Route

---

- ❑ Originating a default route into OSPF:

```
router ospf 42
 default-information originate
```

- Which will originate a default route into OSPF only if a default route exists in the RIB

```
router ospf 42
 default-information originate always
```

- Which will always originate a default route into OSPF even if no default is in the RIB
- ❑ There are equivalent commands for OSPFv3

# OSPF on Point-to-Point Ethernet

---

- ❑ OSPF on point-to-point ethernet:
  - DR and BDR election is not needed on a point to point link – so it is disabled, which is more efficient

```
interface gigabit 2/0
ip ospf network point-to-point
```

- ❑ There are equivalent commands for OSPFv3

# Conclusion

---

- OSPF is a Link State Routing Protocol
- Quick and simple to get started
  - But has a myriad of options and features to cover almost all types of network topology
  - ISPs keep their OSPF design **SIMPLE**
  - >400 routers in a single area is entirely feasible

# Introduction to OSPF



ISP Workshops