OSPFv2 (Part 3)

Configuration Details - and full details

Agenda

- Survey: how many people feel confident with booking & using NetLab?
- Wk03Day2 republished with correction: added a couple of extra config cmds
- In the News: draft-ietf-lsr-ospf-terminology-03 from July 9, 2022 (see p. 2)
- New material:
  - Overview & review of OSPF
  - OSPF Configuration

Assignments and Lab work

- Lab 3 post-lab due **before** your lab session this week
- Lab 4 pre-lab: due **before** your lab session this week
- Lab 4: Single area OSPFv2, this week
- Readings: ENSA Ch 1, 2 due **by** for Mon's lecture (today!)

References

- NetAcad ENSA course: Ch 1, 2 ("modules" 1, 2)

A Vague History of Routing Protocols

IGP = Interior Gateway Protocol – *within* a network controlled by a single entity
EGP = Exterior Gateway Protocol – *between* networks controlled by different entities

	IGP	EGP
Dawn of networking	Few sites: static routing	
Early days of networking	RIP Point-to-Point paths (OSPF being designed)	BGP developed 1989
Middle days of networking	IGRP (Cisco) ISIS - RFC1195 - 1990 EIGRP (Cisco) 1993 OSPF RFC2328 - 1998	BGP (since 1994)
Present time	OSPFv3 / ISIS	BGP invading ← BGP

DR/BDR elections & operation

Elections & operation follow a set of hard & fast rules. **Memorize** the following:

1. Final state (upon completion of a nifty algorithm which prevents glitches):
   Active router with *highest* interface priority becomes **DR**;
   Active router with 2nd *highest* interface priority becomes **BDR**;
   all other routers on the same segment are **DR Other**
2. A tie is decided based on Router ID (unique so no further tie-breakers needed)
3. If the DR fails or becomes disconnected, the BDR is *immediately* promoted to DR without a new election for DR.
4. If the BDR fails, gets disconnected, or becomes DR, then a new election is held for BDR.

– Priority range is 0-255; 0 means can *never* be DR or BDR; 1 is the default value
– Election duration is ~30+ secs (WaitTimer) and varies depending on vendor.

Consider the implications for a "better" router that's a late-comer to the segment: there must be *at least two failures* (including one DR, one BDR) for it to become DR. How well does that work for newer, more powerful routers that take always seem to take soooo much longer to boot?

... And who's the genius who suggested Leader/Follower terminology in last Wed's class? You'd think they're doing some serious networking outside of class time:
https://www.ietf.org/id/draft-ietf-lsr-ospf-terminology-03.html
from 9/Jul/2022. (Version 0 of this document dates from 4/Jun/2022.)
... Please amend your textbooks accordingly!

Neighbour vs Full Adjacency

When reading about OSPF, it's sometimes necessary to pay close attention to the terms "neighbour" vs "adjacency" (or fully adjacent).

Only for those interested: DR Others form a full adjacency with the DR & BDR, but stay as Neighbour amongst themselves.

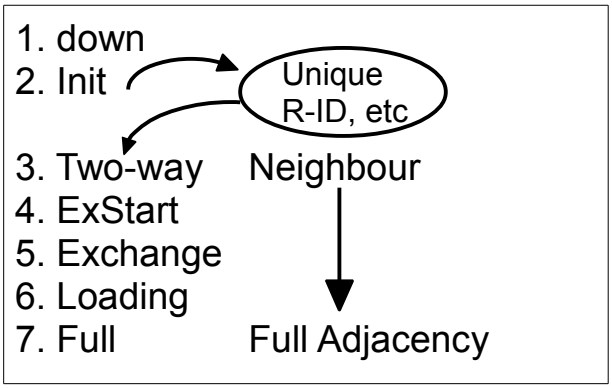

Determination of Router ID

Precedence is determined by the network vendor, so steps vary. Cisco's order is:

1. Specifically configured by administrator (BEST option)
2. If no R-ID configured, then highest IPv4 address on an ***active*** loopback (Good)
3. If no IPv4 loopbacks, then highest IPv4 address on an ***active*** physical interface

So #1 & #2 are fine, but what are the implications for #3 with a slow link partner?!!

Cisco's alternatives for configuring subnets advertised in OSPF

Compared to other vendors, Cisco's OSPF config is messy, plain and simple. We'll choose one for NET2000 based on the list below of pros and cons, but you'll need to recognize and understand both. Sorry! (But don't worry, you'll be soooo relieved when you see another vendor's *clean* implementation in NET3012.)

Note: in the table below, the word "subnet" is used to signify the network ID, so that the word "network" can be reserved for the special statement used by Cisco.

Network statements (Original method)	**Interface config** (Added later)
– One place to determine included subnets (examine only OSPF process config)	– Must scan all interfaces to find subnets
– Config **not** collected all in one place; must still config some params in I/F e.g. priority or point-to-point (faster!)	– Config **not** collected all in one place: must still config some params in process e.g. reference bandwidth
– Can activate **all** interfaces with quad-0 statement, but that really means **all** (no sub-set selection possible!)	– Can activate **many** (all?) interfaces simultaneously using "int range ..." cmd (so sub-set selection **IS** possible)
– Must know brand new *wildcard* mask (exact opposite of subnet mask)	N/A so no problem
– Four possible formats for network + mask (quad-0, wildcard, subnet, exact)	N/A so no problem
– Must type I/F subnet and network identically (no typos!) or else won't work	N/A so no problem
– Asymmetry between IPv4 and IPv6: *must* use interface method for IPv6	N/A so no problem
– *More tedious* to copy & modify: must change BOTH IP addr and network stmt	– *Easy* to copy & modify a config for a similar router: only change I/F address
– To check/tshoot why a I/F isn't active in OSPF must interpret "network" address, double-check address of I/F, but at least "passive" statements are the same sect	– To check/tshoot why a I/F isn't active in OSPF must check both I/F config and process config for "passive" statements
Note: Both "network" statements and interface config can be combined; the results are a combination in a logical "OR" fashion.	

Configuring Advertised Subnets in OSPF (Cisco)

Courtesy of Bob Vachon.

```
255.255.255.255
255.255.255.252
0.   0.   0.   3
```

Subnet Mask	Wildcard Mask
255.255.255.0	0.0.0.**255**
255.255.255.128	0.0.0.**127**
255.255.255.192	0.0.0.**63**
255.255.255.224	0.0.0.**31**
255.255.255.240	0.0.0.**15**
255.255.255.248	0.0.0.**7**
255.255.255.252	0.0.0.**3**

Two options for enabling OSPF on an interface

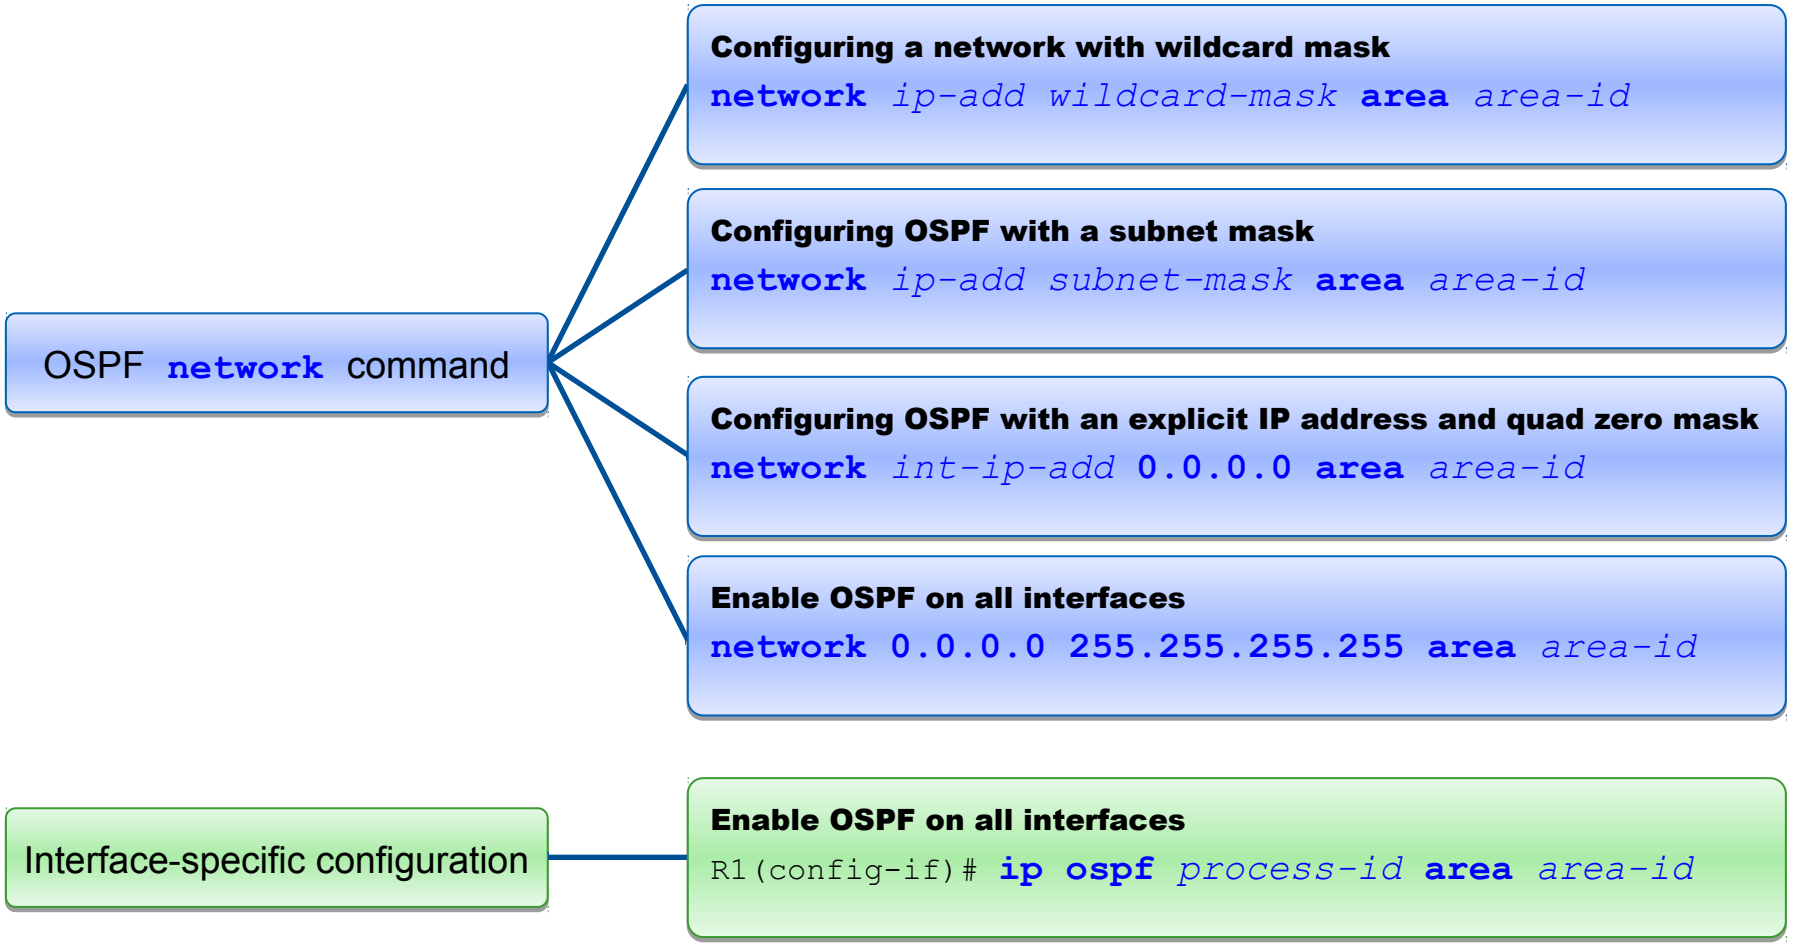

Configuring Advertised Subnets in OSPF (Cisco)

Courtesy of Bob Vachon.

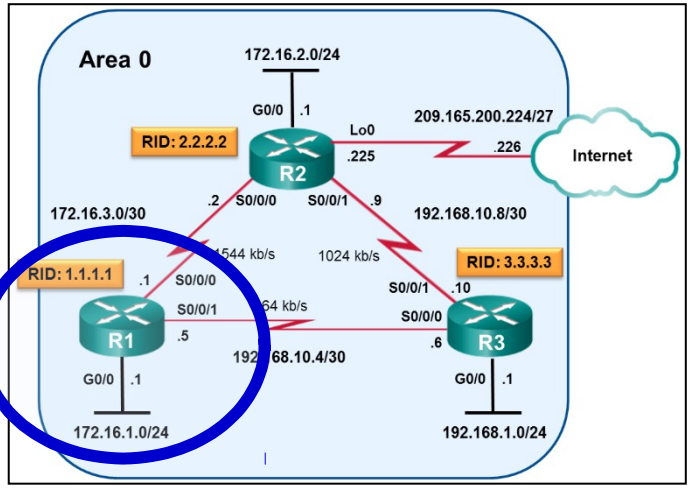

Network with wildcard mask

```
R1(config)# router ospf 10
R1(config-router)# route-id 1.1.1.1
R1(config-router)# network 172.16.1.0 0.0.0.255 area 0
R1(config-router)# network 172.16.3.0 0.0.0.3 area 0
R1(config-router)# network 192.168.10.4 0.0.0.3 area 0
R1(config-router)# end
R1#
```

Network with subnet mask

```
R1(config)# router ospf 10
R1(config-router)# route-id 1.1.1.1
R1(config-router)# network 172.16.1.0 255.255.255.0 area 0
R1(config-router)# network 172.16.3.0 255.255.255.252 area 0
R1(config-router)# network 192.168.10.4 255.255.255.252 area 0
R1(config-router)# end
R1#
```

IP address and quad zero mask

```
R1(config)# router ospf 10
R1(config-router)# route-id 1.1.1.1
R1(config-router)# network 172.16.1.1 0.0.0.0 area 0
R1(config-router)# network 172.16.3.1 0.0.0.0 area 0
R1(config-router)# network 192.168.10.5 0.0.0.0 area 0
R1(config-router)# end
R1#
```

Enable OSPF on all interfaces

```
R1(config)# router ospf 10
R1(config-router)# route-id 1.1.1.1
R1(config-router)# network 0.0.0.0 255.255.255.255 area 0
R1(config-router)# end
R1#
```

Enable OSPF on interfaces

```
R1(config)# int s0/0/0
R1(config-if)# ip ospf 10 area 0
R1(config-if)# exit
R1(config)#
R1(config)# int s0/0/1
R1(config-if)# ip ospf 10 area 0
R1(config-if)# exit
R1(config)#
R1(config)# int g0/0
R1(config-if)# ip ospf 10 area 0
R1(config-if)# end
R1#
```

Interface config is nice because it's simple.

However:

Configuration is not centralized

Complexity increases as the number of router interfaces increase.

Note: Interface-specific settings take precedence over the **network** statement.