TEL3214
Computer Communication Networks

Lecture 7a

Routing – IPv4 with OSPFv2

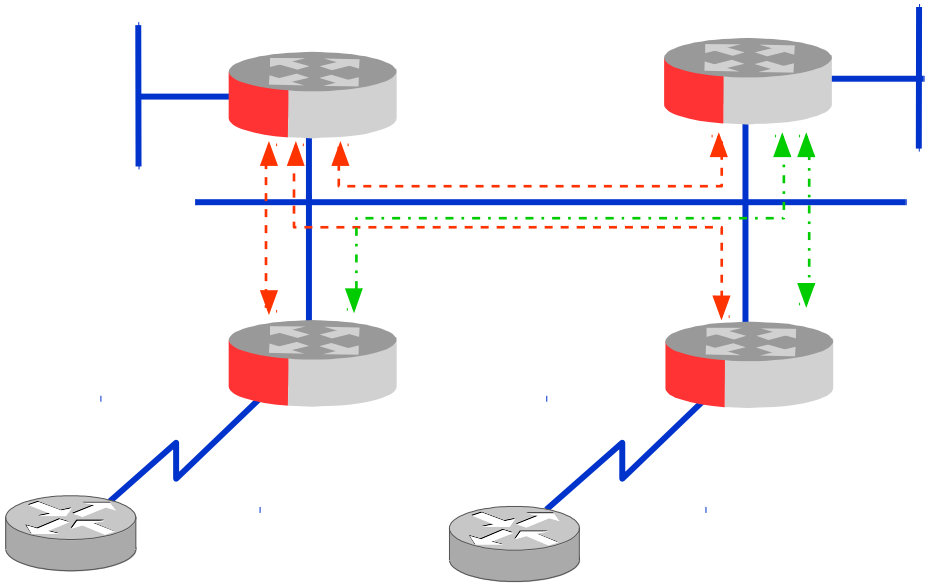

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

Routing

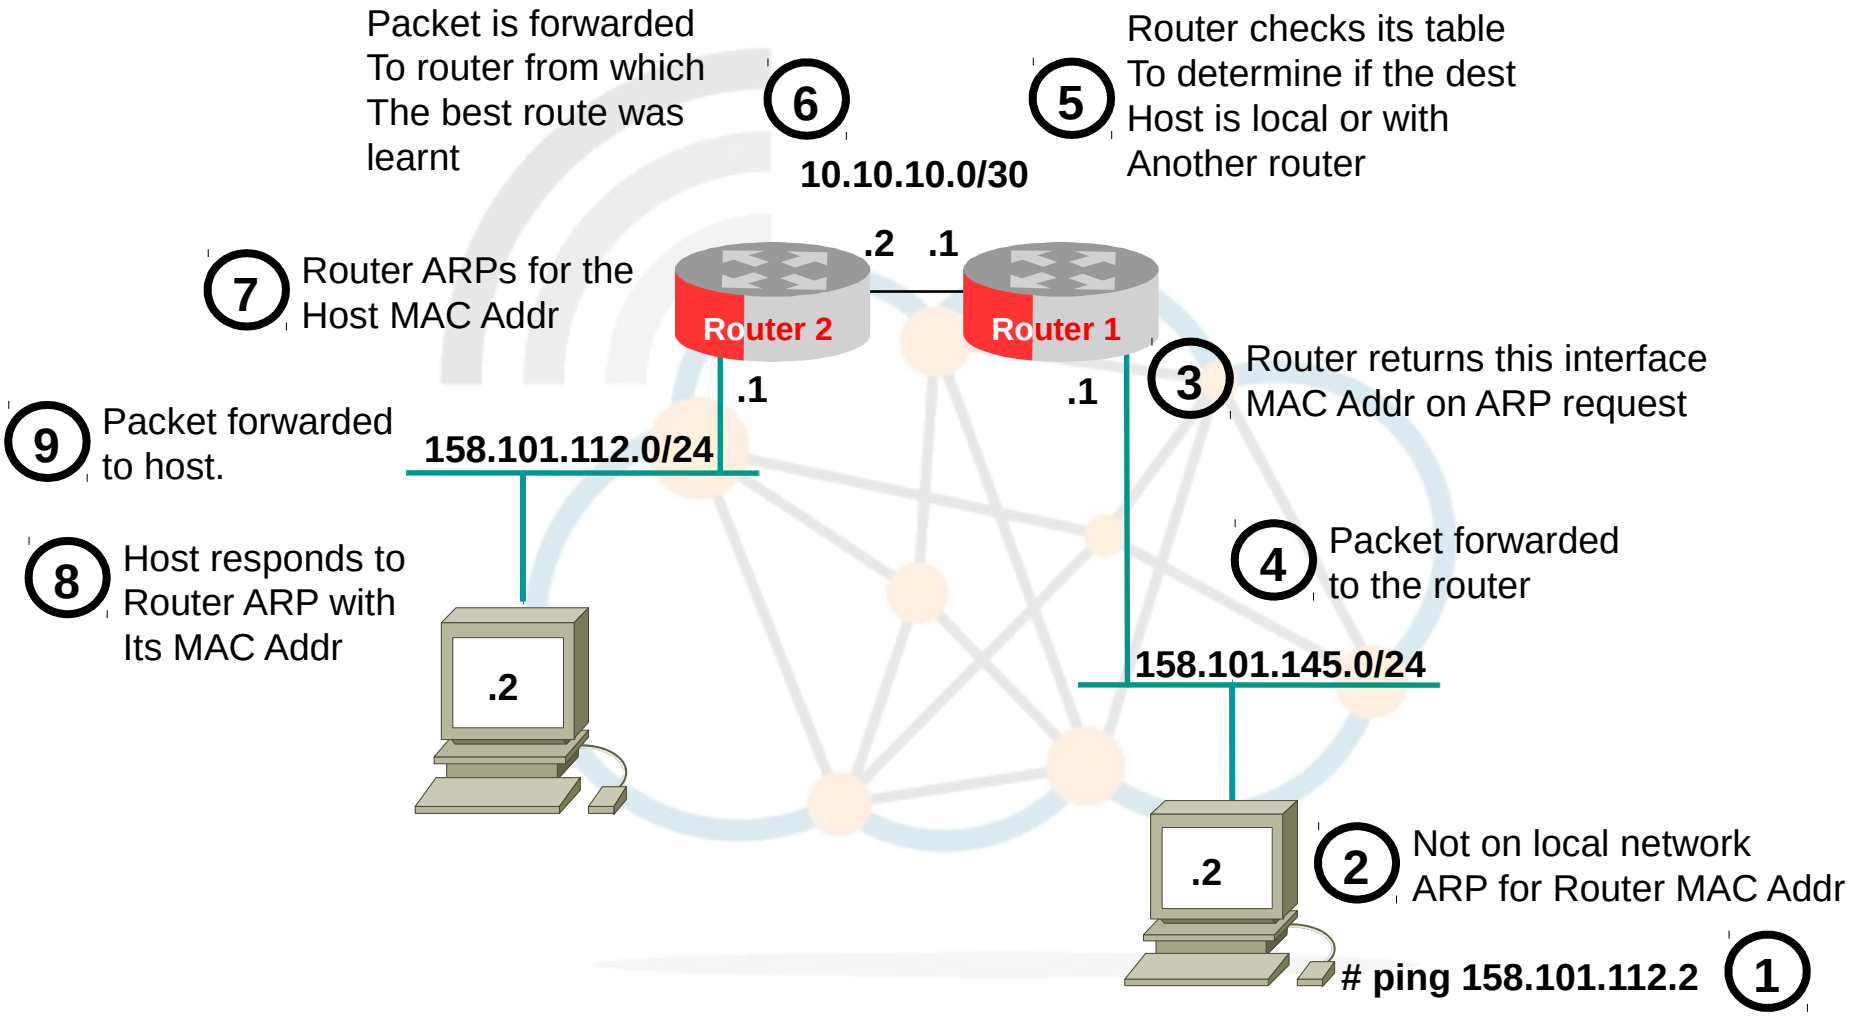
Packet is forwarded
To router from which
The best route was
learnt
6
10.10.10.0/30
5
Router checks its table
To determine if the dest
Host is local or with
Another router
.2
.1
7
Router ARPs for the
Host MAC Addr
Router 2
Router 1
.1
.1
3
Router returns this interface
MAC Addr on ARP request
9
Packet forwarded
to host.
158.101.112.0/24
8
Host responds to
Router ARP with
Its MAC Addr
4
Packet forwarded
to the router
158.101.145.0/24
.2
.2
2
Not on local network
ARP for Router MAC Addr
ping 158.101.112.2
1

Quagga

- Quagga Routing Software Suite is a GPL advanced routing software package that provides a suite of TCP/IP based routing protocols:
  - RIP
  - RIPng
  - OSPFv2
  - OSPFv3
  - Babel
  - BGP4

- Basis for many routing products.
- CLI practically idential to that used by Cisco.

Introduction to Dynamic Host Configuration Protocol (DHCP)

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

DHCP

- DHCP is a mechanism for the automatic configuration of hosts.
  - Bootstrap Protocol (BOOTP).
- DHCP Central server maintains a list of IP addresses which may be assigned on one or more subnets.
- DHCP client request
  - IP address
  - Gateway router
  - DNS Server
- DHCPv4
- DHCPv6
  - StateLess Address Auto Configuration (SLAAC).

DHCP Flow

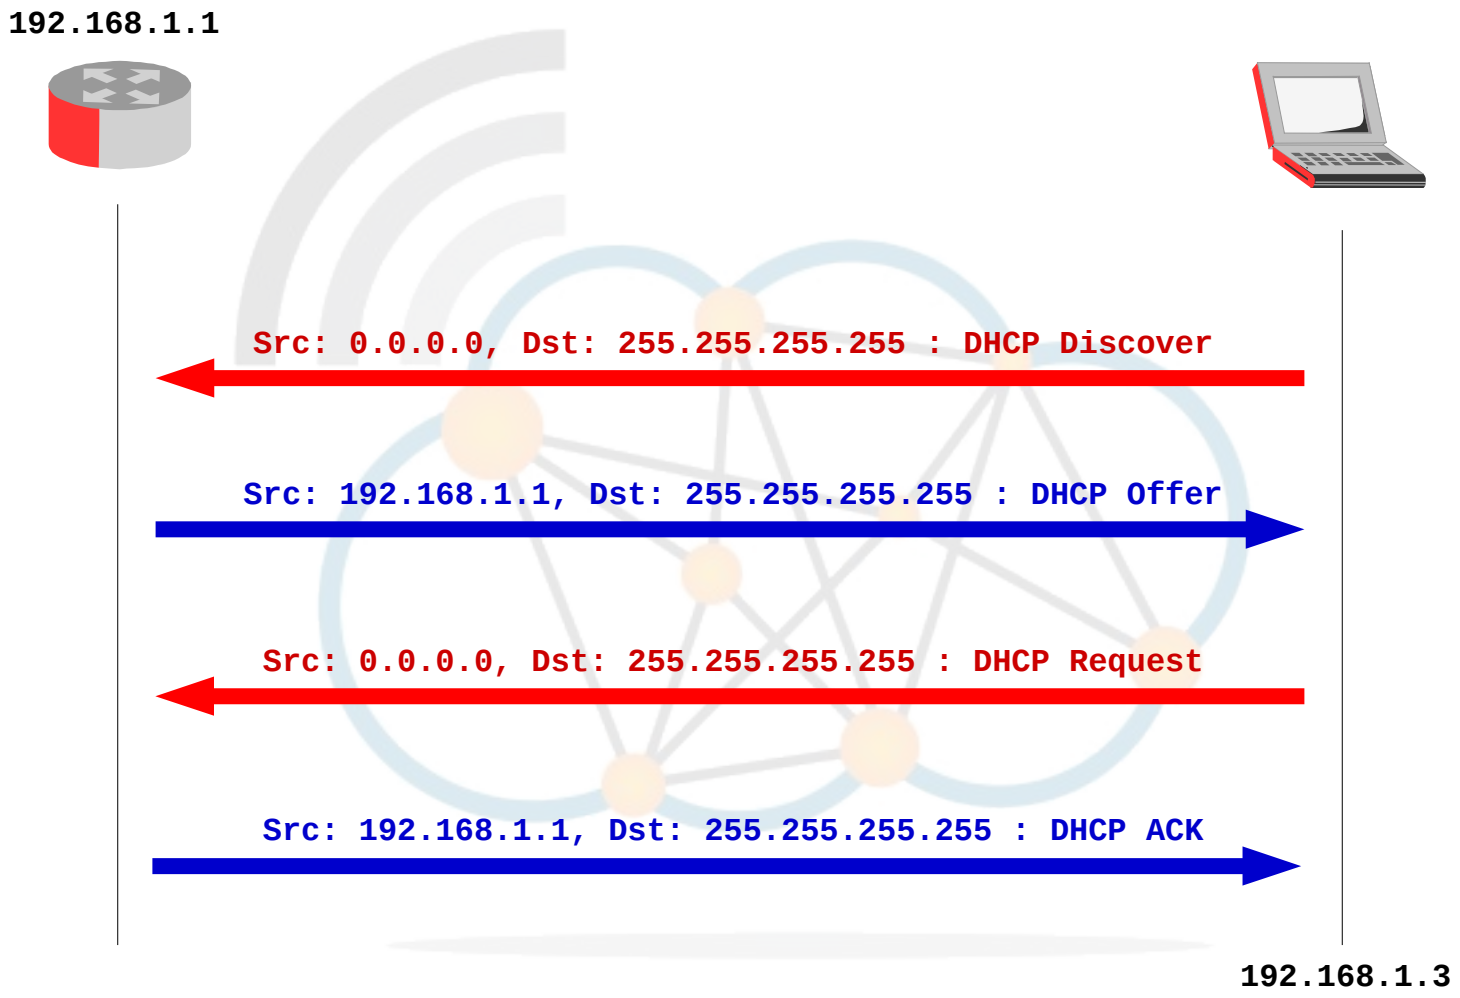
192.168.1.1
Src: 0.0.0.0, Dst: 255.255.255.255 : DHCP Discover
Src: 192.168.1.1, Dst: 255.255.255.255 : DHCP Offer
Src: 0.0.0.0, Dst: 255.255.255.255 : DHCP Request
Src: 192.168.1.1, Dst: 255.255.255.255 : DHCP ACK
192.168.1.3

DHCP Request

```
# dhclient -v eth0
Internet Systems Consortium DHCP Client 4.3.1
Copyright 2004-2014 Internet Systems Consortium.
All rights reserved.
For info, please visit
https://www.isc.org/software/dhcp/

Listening on LPF/eth0/00:00:00:aa:00:01
Sending on   LPF/eth0/00:00:00:aa:00:01
Sending on   Socket/fallback
DHCPREQUEST on eth0 to 255.255.255.255 port 67
DHCPACK from 196.33.41.1
bound to 196.33.41.3 -- renewal in 293 seconds.
```

DHCP Lab

- Run the network:
  - **TEL3214-DHCP-Example.imn**
- Run Wireshark on Host n4.
- From the bash prompt on Host n2:

```
root@n2:/tmp/pycore.57892/n2.conf# dhclient -v eth0
Internet Systems Consortium DHCP Client 4.3.1
Copyright 2004-2014 Internet Systems Consortium.
All rights reserved.
For info, please visit https://www.isc.org/software/dhcp/

Listening on LPF/eth0/00:00:00:aa:00:01
Sending on   LPF/eth0/00:00:00:aa:00:01
Sending on   Socket/fallback
DHCPREQUEST on eth0 to 255.255.255.255 port 67
DHCPACK from 192.168.1.1
bound to 192.168.1.3 -- renewal in 293 seconds.
```

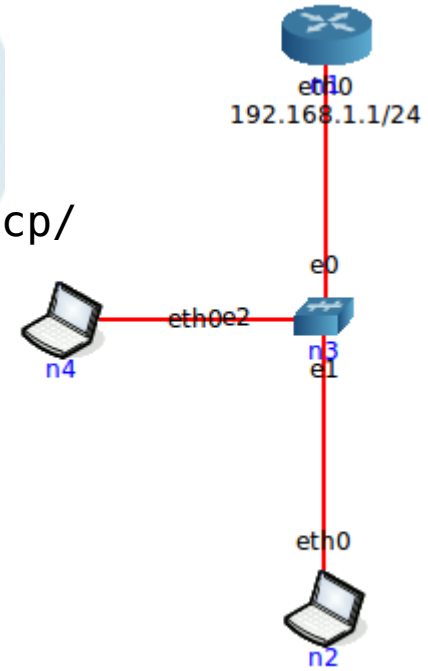

DHCP Lab

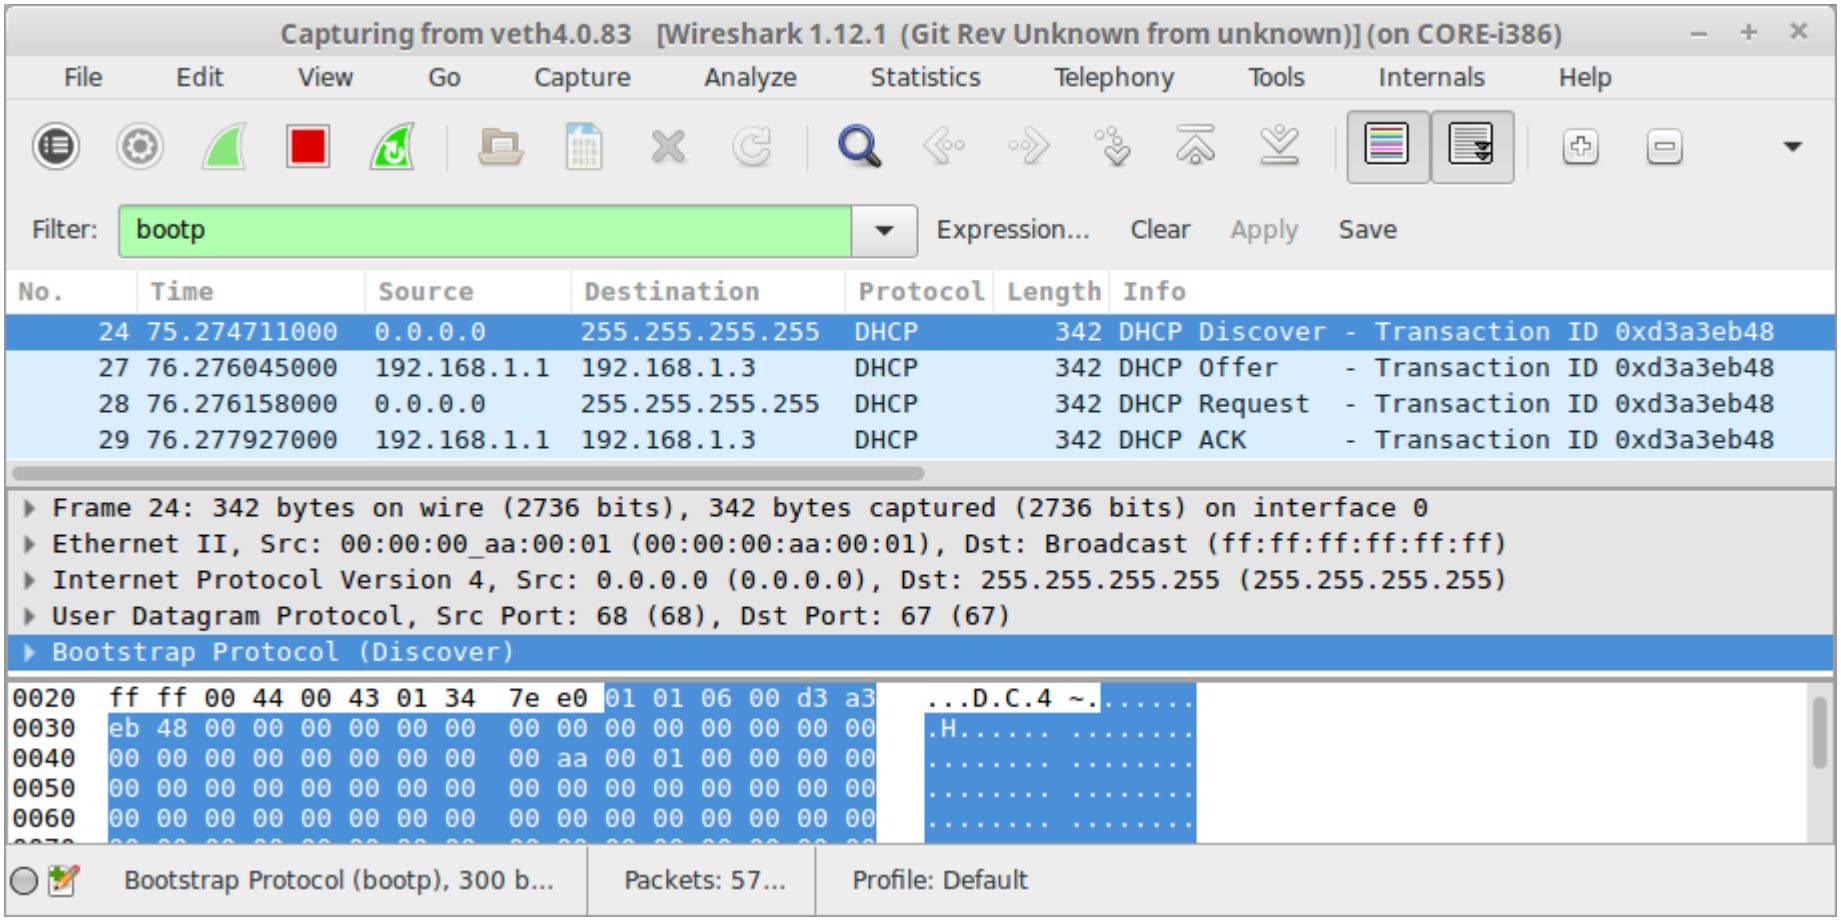
Capturing from veth4.0.83 [Wireshark 1.12.1 (Git Rev Unknown from unknown)] (on CORE-i386)
File Edit View Go Capture Analyze Statistics Telephony Tools Internals Help
Filter: bootp
Expression... Clear Apply Save
No. Time Source Destination Protocol Length Info
24 75.274711000 0.0.0.0 255.255.255.255 DHCP 342 DHCP Discover - Transaction ID 0xd3a3eb48
27 76.276045000 192.168.1.1 192.168.1.3 DHCP 342 DHCP Offer - Transaction ID 0xd3a3eb48
28 76.276158000 0.0.0.0 255.255.255.255 DHCP 342 DHCP Request - Transaction ID 0xd3a3eb48
29 76.277927000 192.168.1.1 192.168.1.3 DHCP 342 DHCP ACK - Transaction ID 0xd3a3eb48
Frame 24: 342 bytes on wire (2736 bits), 342 bytes captured (2736 bits) on interface 0
Ethernet II, Src: 00:00:00_aa:00:01 (00:00:00:aa:00:01), Dst: Broadcast (ff:ff:ff:ff:ff:ff)
Internet Protocol Version 4, Src: 0.0.0.0 (0.0.0.0), Dst: 255.255.255.255 (255.255.255.255)
User Datagram Protocol, Src Port: 68 (68), Dst Port: 67 (67)
Bootstrap Protocol (Discover)
0020 ff ff 00 44 00 43 01 34 7e e0 01 01 06 00 d3 a3 ...D.C.4 ~.....
0030 eb 48 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .H..... ....
0040 00 00 00 00 00 00 00 00 00 aa 00 01 00 00 00 00 ....
0050 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ....
0060 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ....
Bootstrap Protocol (bootp), 300 b...
Packets: 57...
Profile: Default

IPv4 Routed network on NTE

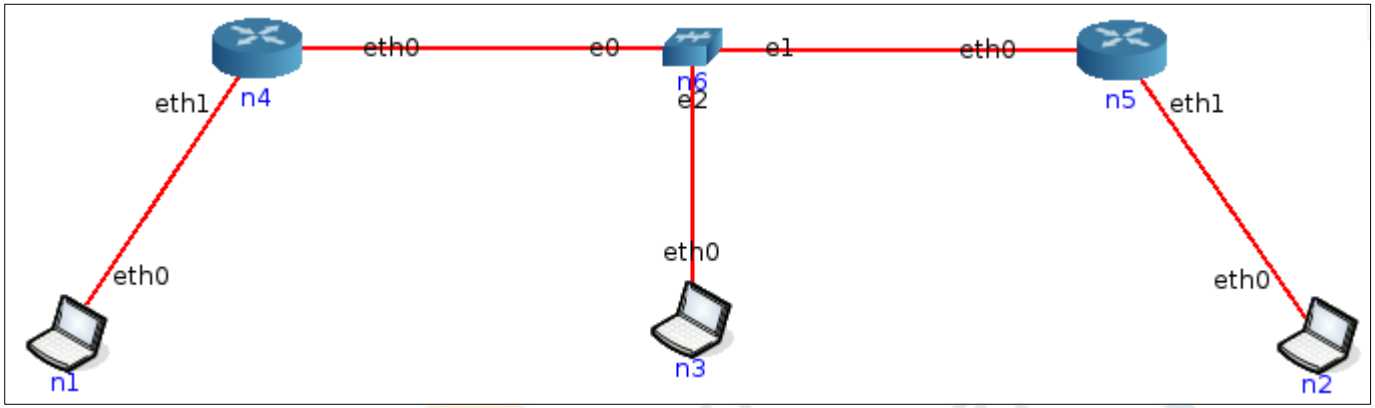

- DHCP assigned IP addresses to n1 and n2
- 10.10.10.0/30 for backbone
- Configure Static routing
- Configure OSPF routing
- TEL3214-Routing-Example.imn

Show running-config

- Right-click on the icon
- **Shell window > bash**.
- Type **vtysh**.

```
n4# show running-config
Building configuration...

Current configuration:
!
!
service integrated-vtysh-config
!
interface eth0
 ipv6 nd suppress-ra
!
interface eth1
 ipv6 nd suppress-ra
!
interface lo
!
router ospf
!
ip forwarding
ipv6 forwarding
!
line vty
!
end
```

Basic configuration

- Configure:
  - Hostname
  - IP addresses

```
n4# conf t
n4(config)#  hostname RTR_n4
RTR_n4(config)#

RTR_n4(config)# interface eth0
RTR_n4(config-if)# ip address 10.10.10.1/30
RTR_n4(config-if)# no shut

RTR_n4(config-if)# int eth1
RTR_n4(config-if)# ip addr 192.168.1.1/24
RTR_n4(config-if)# no shut

RTR_n4(config-if)# int lo
RTR_n4(config-if)# ip addr 10.0.0.1/32
RTR_n4(config-if)# exit
```

Show running-config

- Check config

```
RTR_n4(config)# exit

RTR_n4# show run
Building configuration...
```

```
Current configuration:
!
hostname RTR_n4
!
service integrated-vtysh-config
!
interface eth0
 ip address 10.10.10.1/24
 ipv6 nd suppress-ra
!
interface eth1
 ip address 192.168.1.1/24
 ipv6 nd suppress-ra
!
interface lo
 ip address 10.0.0.1/32
!
router ospf
!
ip forwarding
ipv6 forwarding
!
line vty
!
end
```

DHCP Server configuration

- Configure DHCP Server.
- Right click on Router n4 and select **Services...**
- Under the **Utility** column click **DHCP** and the **spanner symbol** to the right of it.
- Add the following to the end of the text and click **Apply**.

```
subnet 192.168.1.0 netmask 255.255.255.0 {
    pool {
        range 192.168.1.2 192.168.1.254;
        Default-lease-time 600;
        option routers 192.168.1.1;
        option domain-name-servers 8.8.8.8;
    }
}
```

Configure Router n5

```
n5(config)# hostname RTR_n5
RTR_n5(config)# int eth0
RTR_n5(config-if)# ip addr 10.10.10.2/30
RTR_n5(config-if)# no shut
RTR_n5(config-if)# int eth1
RTR_n5(config-if)# ip addr 192.168.2.1/24
RTR_n5(config-if)# no shut
RTR_n5(config-if)# int lo
RTR_n5(config-if)# ip addr 10.0.0.2/32
RTR_n5(config-if)# no shut
```

- Add the following to the DHCP Services file.

```
subnet 192.168.1.0 netmask 255.255.255.0 {
    pool {
        range 192.168.2.2 192.168.2.254;
        default-lease-time 600;
        option routers 192.168.2.1;
        option domain-name-servers 8.8.8.8;
    }
}
```

Show running-config on Router n5

- Check config

```
RTR_n5(config)# exit

RTR_n5# show run
Building configuration...
```

```
Current configuration:
!
hostname RTR_n5
!
service integrated-vtysh-config
!
interface eth0
 ip address 10.10.10.2/30
 ipv6 nd suppress-ra
!
interface eth1
 ip address 192.168.2.1/24
 ipv6 nd suppress-ra
!
interface lo
 ip address 10.0.0.2/32
!
router ospf
!
ip forwarding
ipv6 forwarding
!
line vty
!
end
```

Confirm DHCP Service is running

- Confirm DHCP Service.

```
root@n4:/tmp/pycore.39622/n4.conf# service isc-dhcp-server status
Status of ISC DHCP server: dhcpd is running.
```

```
root@n4:/tmp/pycore.39622/n4.conf# service isc-dhcp-server status
Status of ISC DHCP server: dhcpd is not running.
```

```
root@n4:/tmp/pycore.39622/n4.conf# service isc-dhcp-server status
Starting ISC DHCP server: dhcpd.
```

Note: **`systemctl status isc-dhcp-server`**
command fails with a D-Bus connection error
This is due to the way CORE builds LinuX Containers (LXC)

Confirm DHCP Client is getting an IP address

- Use the ***dhclient*** to refresh IP address.

```
root@n1:/tmp/pycore.39622/n1.conf# dhclient -v eth0
Internet Systems Consortium DHCP Client 4.3.1
Copyright 2004-2014 Internet Systems Consortium.
All rights reserved.
For info, please visit https://www.isc.org/software/dhcp/

Listening on LPF/eth0/00:00:00:aa:00:06
Sending on   LPF/eth0/00:00:00:aa:00:06
Sending on   Socket/fallback
DHCPREQUEST on eth0 to 255.255.255.255 port 67
DHCPACK from 192.168.1.1
RTNETLINK answers: File exists
bound to 192.168.1.3 -- renewal in 221 seconds.
```

Review routes in each router

- What's missing ?

```
RTR_n4# show ip route
Codes: K - kernel route, C - connected, S - static, R - RIP,
       O - OSPF, I - IS-IS, B - BGP, A - Babel,
       > - selected route, * - FIB route

C>* 10.0.0.1/32 is directly connected, lo
C>* 10.10.10.0/30 is directly connected, eth0
C>* 127.0.0.0/8 is directly connected, lo
C>* 192.168.1.0/24 is directly connected, eth1


RTR_n5# show ip route
Codes: K - kernel route, C - connected, S - static, R - RIP,
       O - OSPF, I - IS-IS, B - BGP, A - Babel,
       > - selected route, * - FIB route

C>* 10.0.0.2/32 is directly connected, lo
C>* 10.10.10.0/30 is directly connected, eth0
C>* 127.0.0.0/8 is directly connected, lo
C>* 192.168.2.0/24 is directly connected, eth1
```

Static route Router n4

```
RTR_n4(config)# ip route 192.168.2.0/24 10.10.10.2

RTR_n4(config)# exit

RTR_n4# show ip route
Codes: K - kernel route, C - connected, S - static, R - RIP,
       O - OSPF, I - IS-IS, B - BGP, A - Babel,
       > - selected route, * - FIB route

C>* 10.0.0.1/32 is directly connected, lo
C>* 10.10.10.0/30 is directly connected, eth0
C>* 127.0.0.0/8 is directly connected, lo
C>* 192.168.1.0/24 is directly connected, eth1
S>* 192.168.2.0/24 [1/0] via 10.10.10.2, eth0
```

Static route Router n5

```
RTR_n5(config)# ip route 192.168.1.0/24 10.10.10.1

RTR_n5(config)# exit

RTR_n5# show ip route
Codes: K - kernel route, C - connected, S - static, R - RIP,
       O - OSPF, I - IS-IS, B - BGP, A - Babel,
       > - selected route, * - FIB route

C>* 10.0.0.2/32 is directly connected, lo
C>* 10.10.10.0/30 is directly connected, eth0
C>* 127.0.0.0/8 is directly connected, lo
S>* 192.168.1.0/24 [1/0] via 10.10.10.1, eth0
C>* 192.168.2.0/24 is directly connected, eth1
```

Test from Host n1 to Host n2

```
root@n1:/tmp/pycore.48245/n1.conf# ping -c1 192.168.2.3
PING 192.168.2.3 (192.168.2.3) 56(84) bytes of data.
64 bytes from 192.168.2.3: icmp_seq=1 ttl=62 time=0.036 ms

--- 192.168.2.3 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.036/0.036/0.036/0.000 ms


root@n1:/tmp/pycore.48245/n1.conf# traceroute 192.168.2.3
traceroute to 192.168.2.3 (192.168.2.3), 30 hops max, 60 byte packets
 1  192.168.1.1 (192.168.1.1)  0.035 ms  0.008 ms  0.005 ms
 2  10.10.10.2 (10.10.10.2)  0.022 ms  0.011 ms  0.010 ms
 3  192.168.2.3 (192.168.2.3)  0.022 ms  0.014 ms  0.014 ms
```

OSPF

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

Open Shortest Path First (OSPF)

- OSPF is an Interior Gateway Protocol (IGP) most suited for use in large networks.
- OSPF uses a link-state algorithm to exchange routing information between routers in an autonomous system (AS).
- An AS is a collection of routers and networks administratively configured to belong to a single organisation.
- OSPF enables the routers to quickly synchronize their topological databases, topology information for the AS only floods in response to topological change.

OSPF Vs the Distance Vector

- Compared to other distance vector protocols like RIP and IGRP, OSPF:
  - Chooses the least costly path as the best path
  - Can calculate equal cost multiple paths to a destination
  - Distributes external information independently
  - Propagates routing information quickly and stably
  - Handles variable length subnet masks (VLSM)
  - Supports multicasting
  - Responds quickly to topological changes by utilizing “reliable flooding” to minimize routing traffic
  - Is loop free
  - Supports large metrics, external route tags and authentication of protocol exchanges

OSPF Adjacencies with DR & BDR

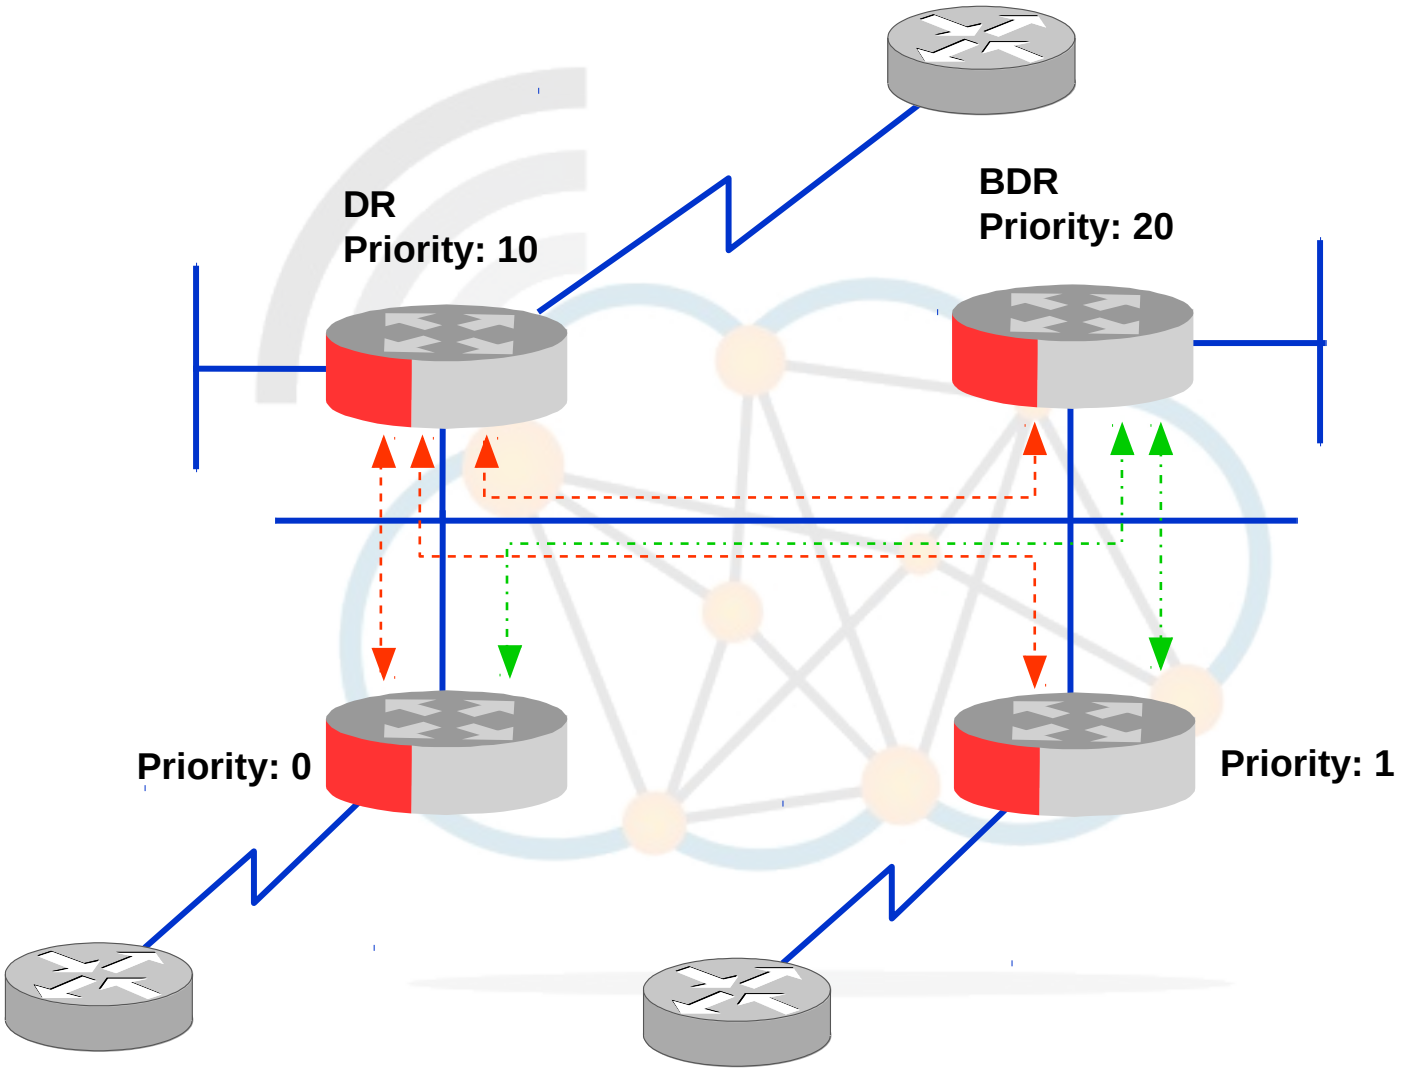

Authentication

- All OSPF protocol exchanges are authenticated.
- OSPF authentication ensures routers exchange information only with trusted neighbours.
  - Simple password
    - clear text
    - case sensitive
    - not encrypted.
  - Cryptographic
    - MD5 authentication.
    - Each key is identified by the combination of an interface and Key ID. A default key ID of 0 is automatically set when an interface is configured for cryptographic authentication.
    - An interface can have multiple active keys.
    - Each key has four time constants associated with it, governing the use of the key during specific time periods.

OSPF Hello messages

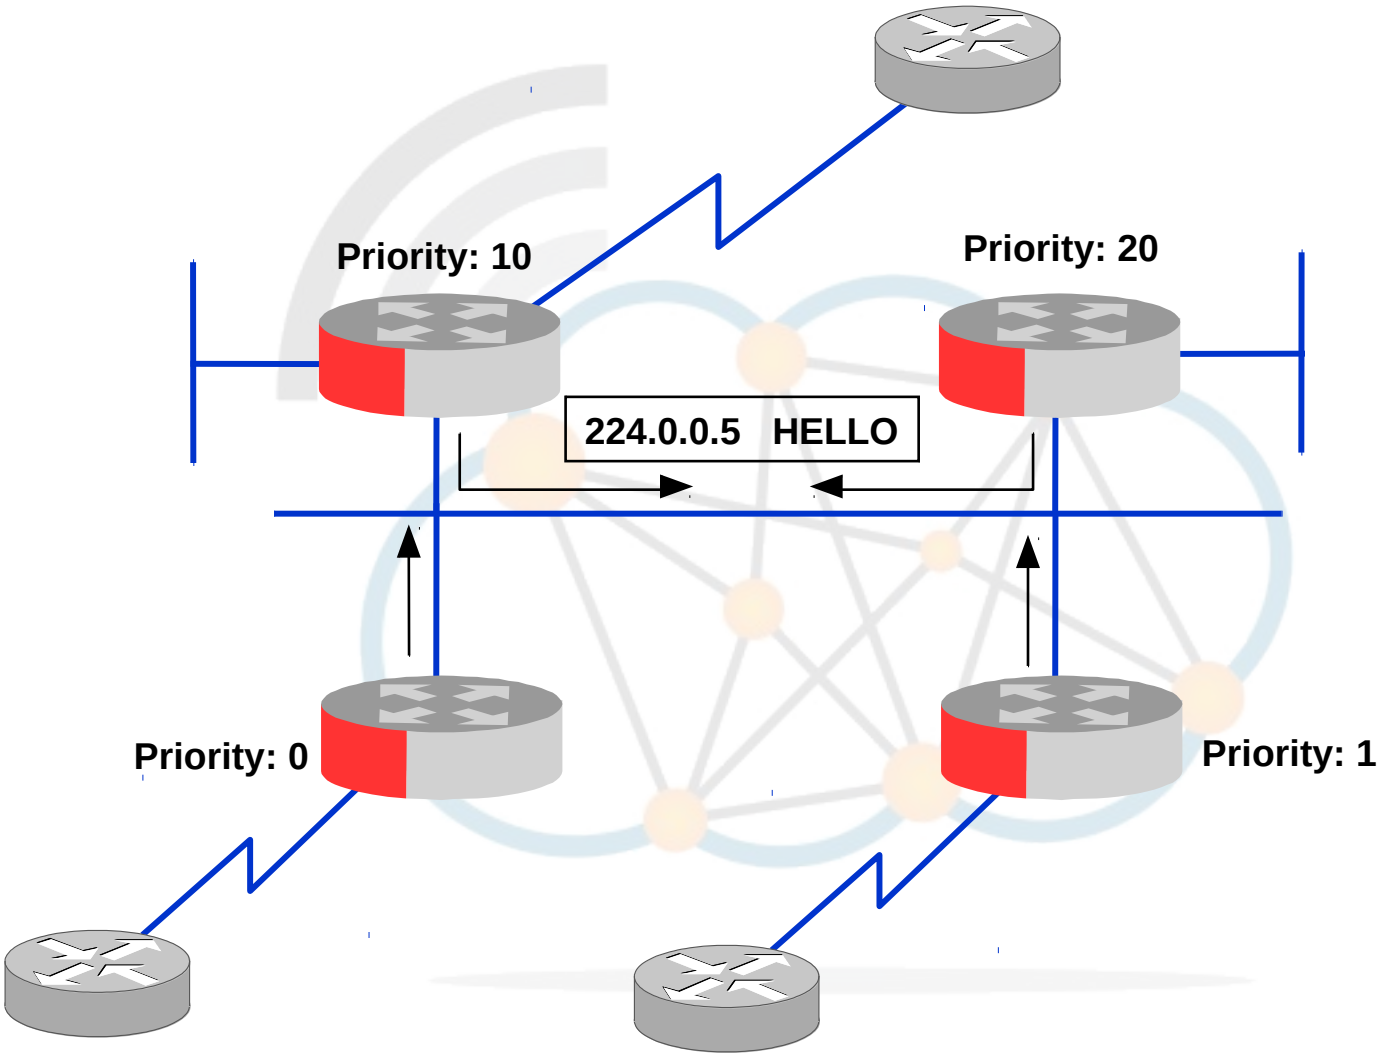

DR BDR Election

- All routers create list of eligible routers
  - Priority greater than 0
  - OSPF State of 2 way
  - DR or BDR IP Address in same network as interface
- The BDR is chosen first which is the router with the highest priority
- The DR is chosen from the remaining routers again the one with the highest priority
- If there were not enough routers to have a BDR and a DR then the BDR becomes the DR
- If the priorities are equal the Router ID is used as a tie-breaker.

OSPF Adjacencies with DR & BDR

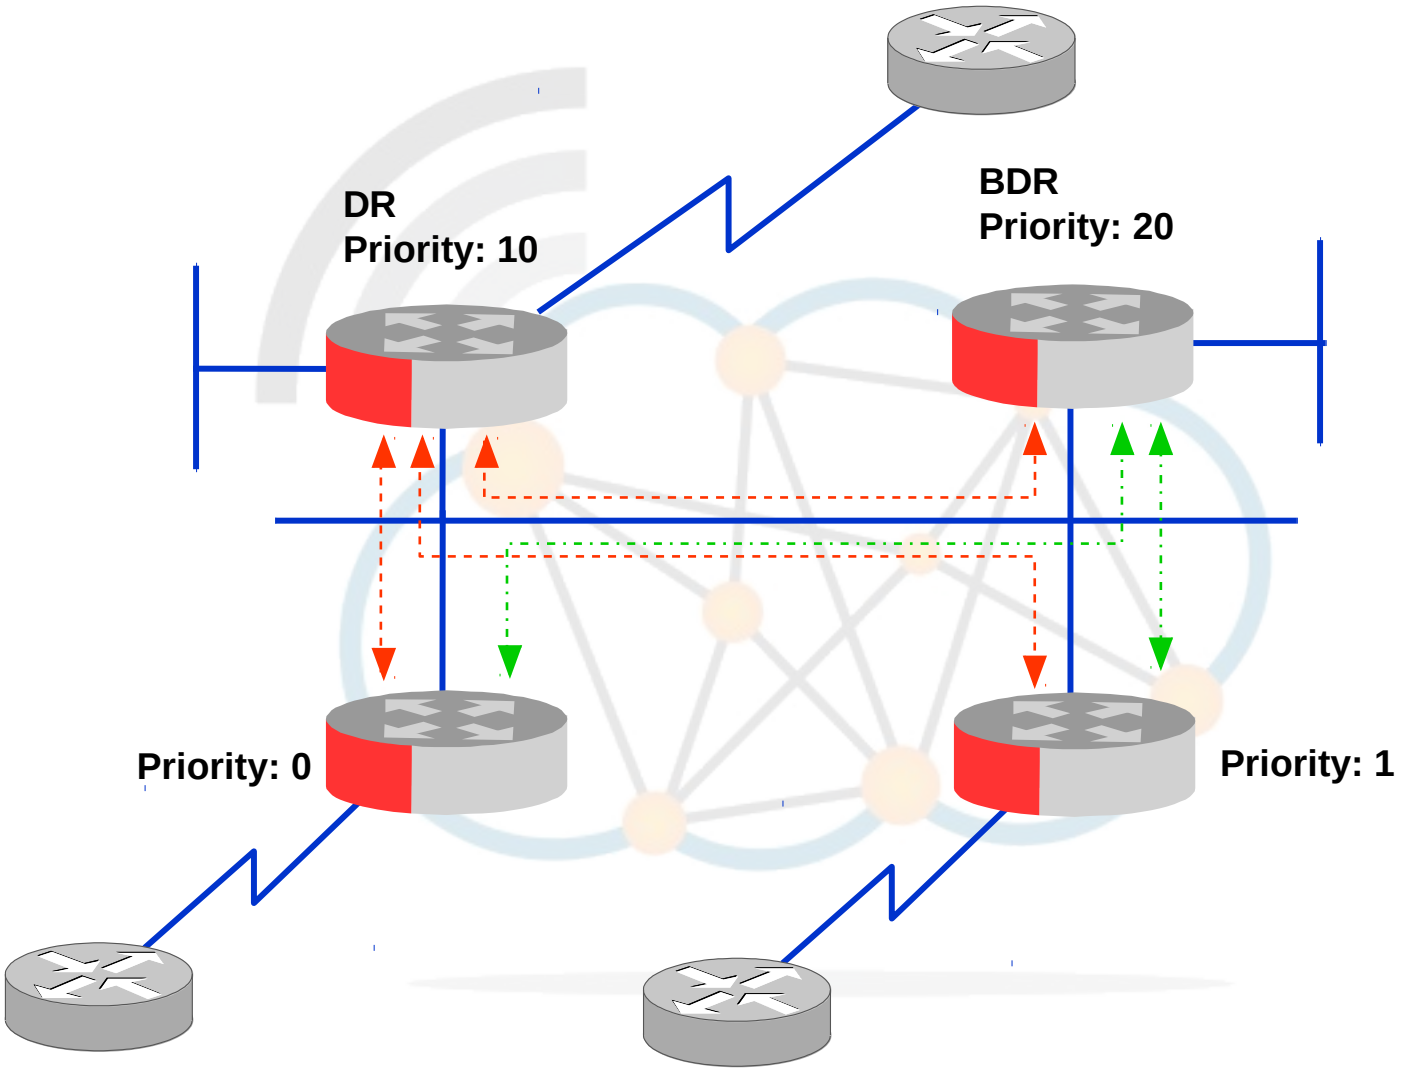

Router LSAs (Type 1)

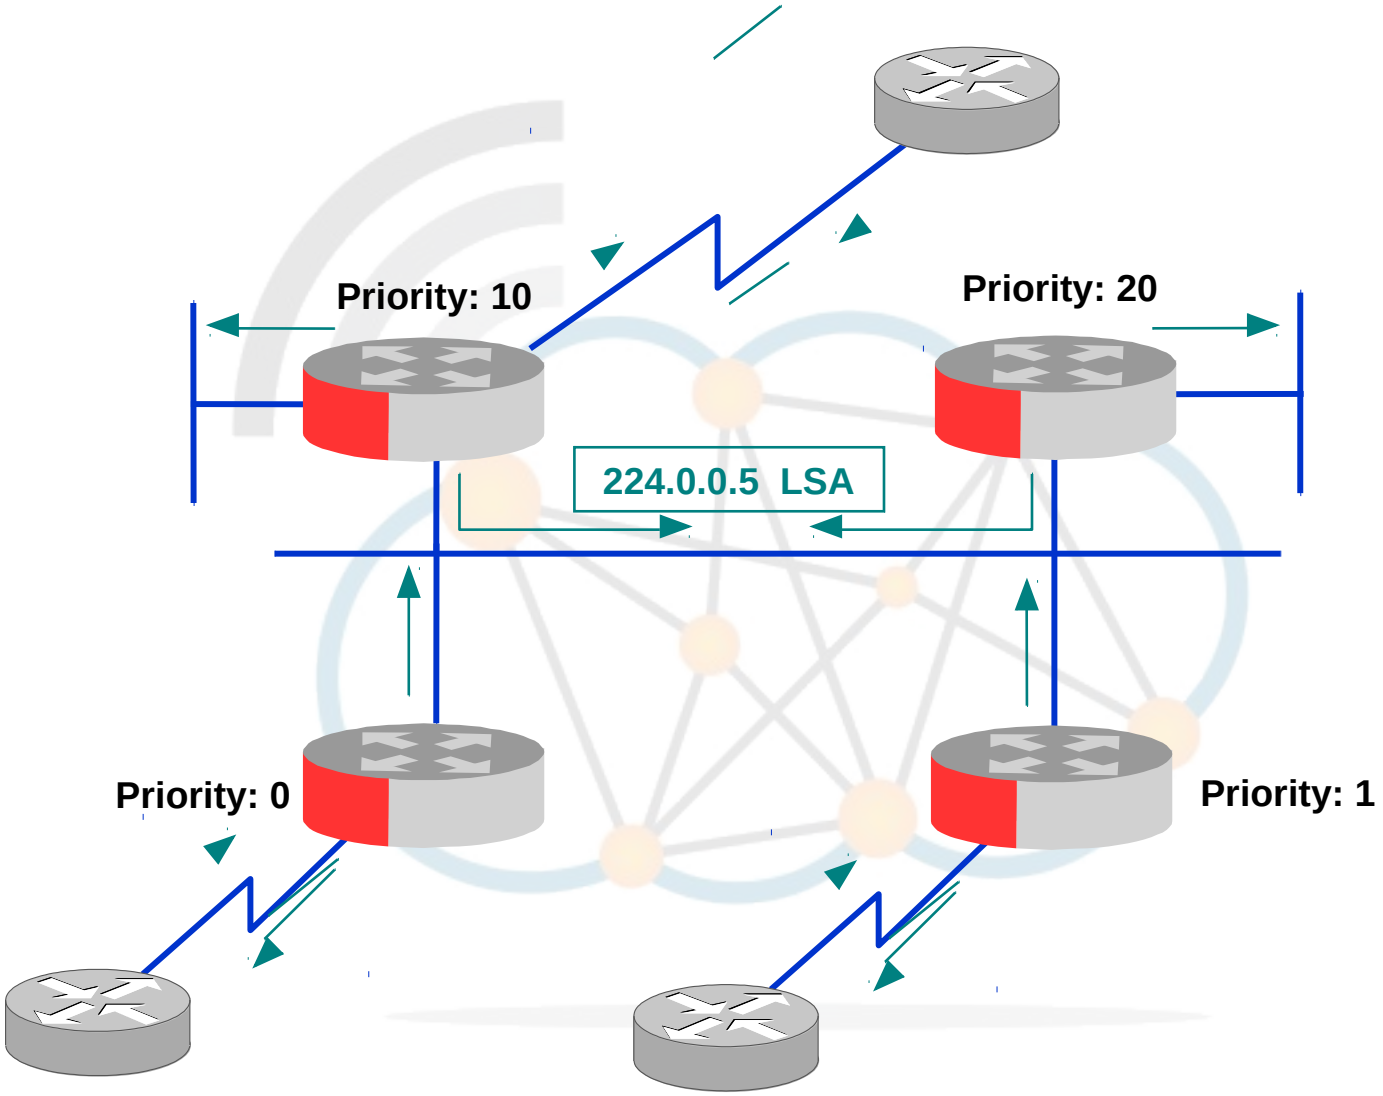

Network LSAs (Type 2)

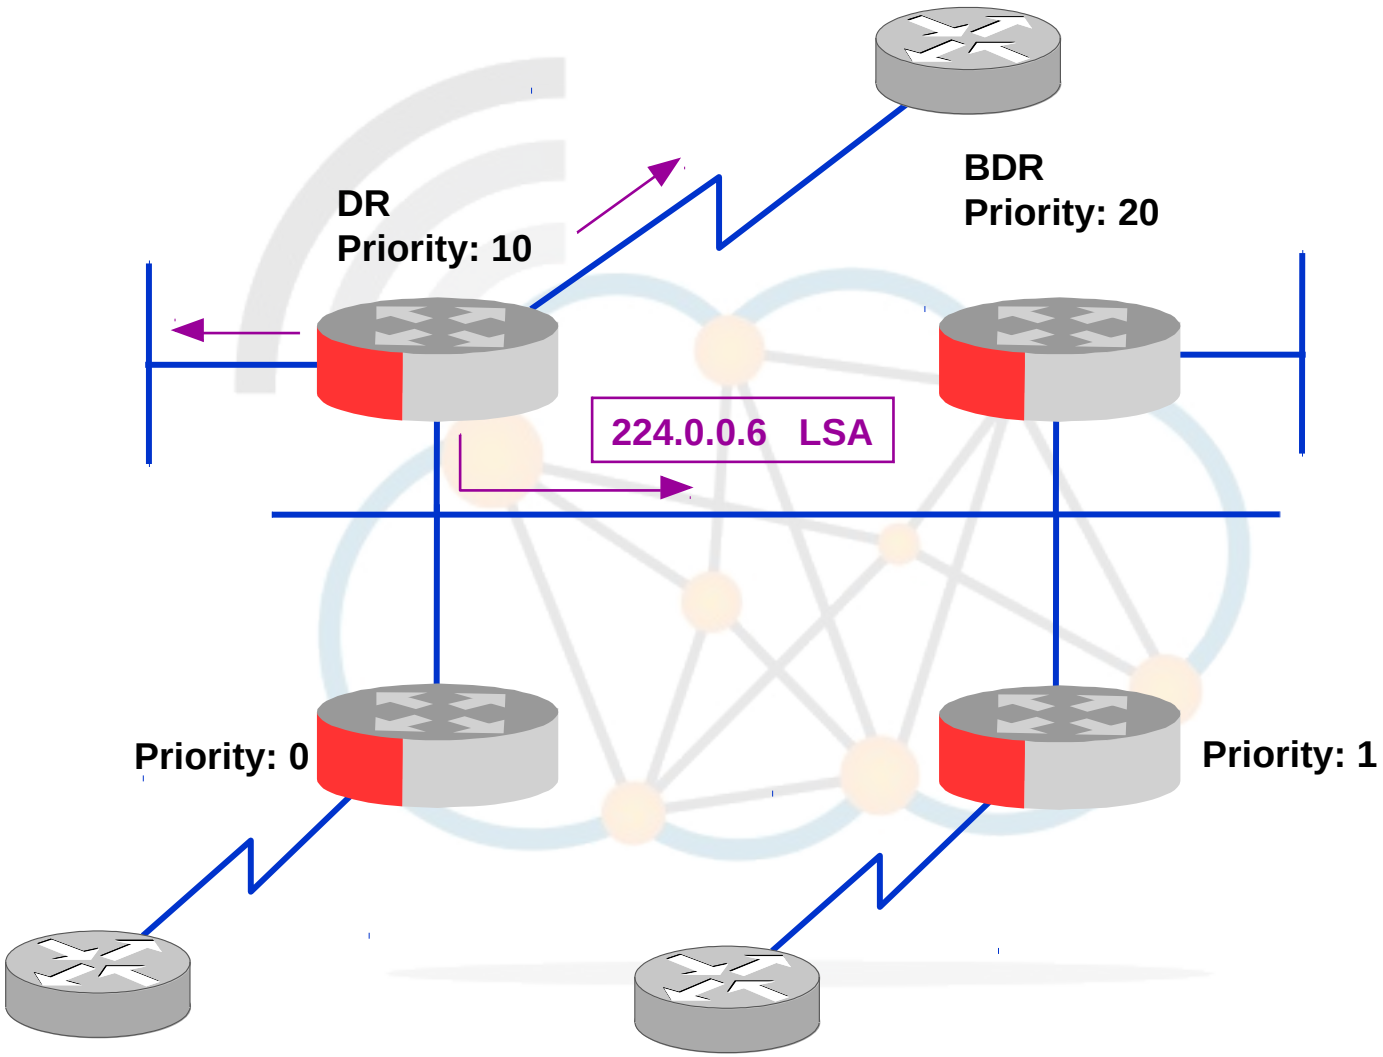

OSPF Timers

- OSPFs various timers interact as follows:
  - If a link goes down for twenty seconds, then comes back up, OSPF doesn't notice.
  - If a link flaps constantly, but at least one of every four Hello packets make it across, OSPF doesn't notice.
  - If a link goes down for anywhere from a minute to half an hour, OSPF floods an LSA when it goes down, and another LSA when it comes back up.
  - If a link stays down for more than half an hour, LSAs originated by remote routers (that have become unreachable) begin to age out. When the link comes back up, all these LSAs will be re-flooded.

Generating the Shortest Path Tree

- Dijkstra's SPF algorithm
  - Compute the Shortest Path Tree
- Populate the routing table

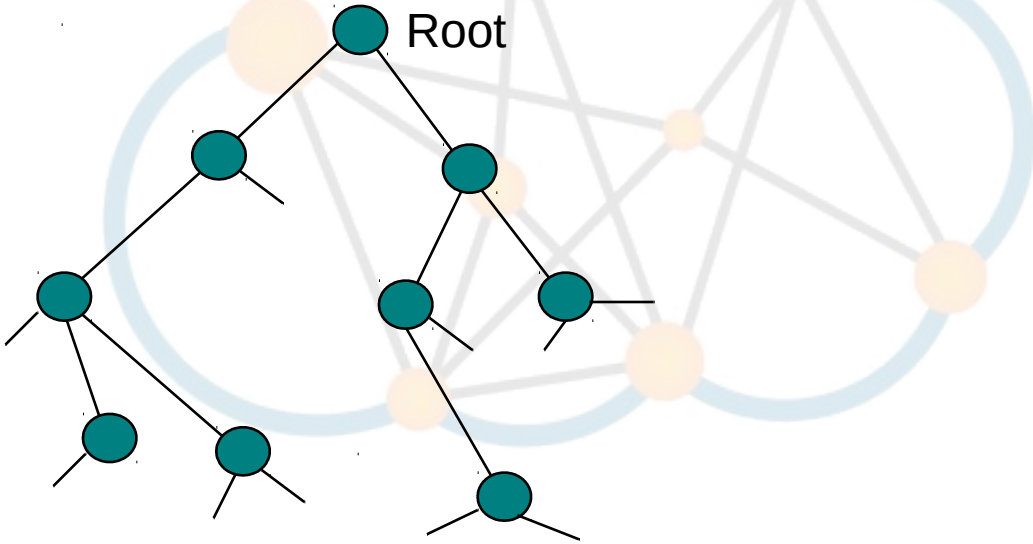

Administrative Distance

- Routes to the same location can be learned from numerous sources
- Mechanism to determine the best sources

Route Source	Default Distance
Local Interface	0
Static Route	1
EIGRP	90
IGRP	100
OSPF	110
RIP	120
External EIGRP	170
Unknown	255

Configure OSPFv2

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

Remove Static routes

```
RTR_n4(config)# no ip route 192.168.2.0/24 10.10.10.2
```

```
RTR_n5(config)# no ip route 192.168.1.0/24 10.10.10.1
```

OSPFv2 Options

```
RTR_n4# conf t
RTR_n4(config)# router ospf
RTR_n4(config-router)# ?
  area                    OSPF area parameters
  auto-cost               Calculate OSPF interface cost according to bandwidth
  capability              Enable specific OSPF feature
  compatible              OSPF compatibility list
  default-information     Control distribution of default information
  default-metric          Set metric of redistributed routes
  distance                Define an administrative distance
  distribute-list         Filter networks in routing updates
  end                     End current mode and change to enable mode
  exit                    Exit current mode and down to previous mode
  list                    Print command list
  log-adjacency-changes   Log changes in adjacency state
  max-metric              OSPF maximum / infinite-distance metric
  mpls-te                 MPLS-TE specific commands
  neighbor                Specify neighbor router
  network                 Enable routing on an IP network
  no                      Negate a command or set its defaults
  ospf                    OSPF specific commands
  passive-interface       Suppress routing updates on an interface
  quit                    Exit current mode and down to previous mode
  redistribute            Redistribute information from another routing protocol
  refresh                 Adjust refresh parameters
  router-id               router-id for the OSPF process
  timers                  Adjust routing timers
```

Configure OSPFv2

Router n4

```
RTR_n4(config-router)# router-id 10.0.0.1

RTR_n4(config-router)# network 192.168.1.0/24 area 0.0.0.0
RTR_n4(config-router)# network 10.10.10.0/30 area 0.0.0.0
```

Router n5

```
RTR_n5(config-router)# router-id 10.0.0.2

RTR_n5(config-router)# network 192.168.2.0/24 area 0.0.0.0
RTR_n5(config-router)# network 10.10.10.0/30 area 0.0.0.0
```

Review OSPFv2 in Router n4

```
RTR_n4# show ip ospf neighbor

Neighbor ID Pri State    Dead Time Address   Interface        RXmtL RqstL DBsmL
10.0.0.2     1  Full/DR  32.163s 10.10.10.2  eth1:10.10.10.1   0     0     0

RTR_n4# show ip ospf database

       OSPF Router with ID (10.0.0.1)

                Router Link States (Area 0.0.0.0)

Link ID         ADV Router      Age  Seq#       CkSum  Link count
10.0.0.1        10.0.0.1         274 0x8000000c 0x174e 2
10.0.0.2        10.0.0.2         405 0x80000006 0x2443 2

                Net Link States (Area 0.0.0.0)

Link ID         ADV Router      Age  Seq#       CkSum
10.10.10.2      10.0.0.2         315 0x80000002 0x64b4
```

Review OSPFv2 in Router n4

```
RTR_n4# show ip ospf route
============ OSPF network routing table ============
N    10.10.10.0/30         [10] area: 0.0.0.0
                           directly attached to eth1
N    192.168.1.0/24        [10] area: 0.0.0.0
                           directly attached to eth0
N    192.168.2.0/24        [20] area: 0.0.0.0
                           via 10.10.10.2, eth1

============ OSPF router routing table =============

============ OSPF external routing table ===========


RTR_n4# sh ip route
Codes: K - kernel route, C - connected, S - static, R - RIP,
       O - OSPF, I - IS-IS, B - BGP, A - Babel,
       > - selected route, * - FIB route

C>* 10.0.0.1/32 is directly connected, lo
O   10.10.10.0/30 [110/10] is directly connected, eth1, 00:45:50
C>* 10.10.10.0/30 is directly connected, eth1
C>* 127.0.0.0/8 is directly connected, lo
O   192.168.1.0/24 [110/10] is directly connected, eth0, 00:45:50
C>* 192.168.1.0/24 is directly connected, eth0
O>* 192.168.2.0/24 [110/20] via 10.10.10.2, eth1, 00:45:40
```

Review OSPFv2 in Router n5

```
RTR_n5# sh ip o n

Neighbor ID Pri State       Dead Time Address     Interface         RXmtL RqstL DBsmL
10.0.0.1     1  Full/Backup 32.315s   10.10.10.1  eth0:10.10.10.2     0     0     0


RTR_n5# show ip o d

        OSPF Router with ID (10.0.0.2)

                Router Link States (Area 0.0.0.0)

Link ID         ADV Router      Age  Seq#       CkSum  Link count
10.0.0.1        10.0.0.1        468 0x8000000c 0x174e 2
10.0.0.2        10.0.0.2        597 0x80000006 0x2443 2

                Net Link States (Area 0.0.0.0)

Link ID         ADV Router      Age  Seq#       CkSum
10.10.10.2      10.0.0.2        507 0x80000002 0x64b4
```

Review OSPFv2 in Router n5

```
RTR_n5# show ip o r
============ OSPF network routing table ============
N    10.10.10.0/30          [10] area: 0.0.0.0
                            directly attached to eth0
N    192.168.1.0/24         [20] area: 0.0.0.0
                            via 10.10.10.1, eth0
N    192.168.2.0/24         [10] area: 0.0.0.0
                            directly attached to eth1

============ OSPF router routing table =============

============ OSPF external routing table ===========


RTR_n5# sh ip ro
Codes: K - kernel route, C - connected, S - static, R - RIP,
       O - OSPF, I - IS-IS, B - BGP, A - Babel,
       > - selected route, * - FIB route

C>* 10.0.0.2/32 is directly connected, lo
O   10.10.10.0/30 [110/10] is directly connected, eth0, 00:51:33
C>* 10.10.10.0/30 is directly connected, eth0
C>* 127.0.0.0/8 is directly connected, lo
O>* 192.168.1.0/24 [110/20] via 10.10.10.1, eth0, 00:47:48
O   192.168.2.0/24 [110/10] is directly connected, eth1, 00:51:52
C>* 192.168.2.0/24 is directly connected, eth1
```

Testing the link from Host n1 to Host n2

```
root@n1:/tmp/pycore.41200/n1.conf# ping -c1 192.168.2.3
PING 192.168.2.3 (192.168.2.3) 56(84) bytes of data.
64 bytes from 192.168.2.3: icmp_seq=1 ttl=62 time=0.104 ms

--- 192.168.2.3 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.104/0.104/0.104/0.000 ms


root@n1:/tmp/pycore.41200/n1.conf# traceroute 192.168.2.3
traceroute to 192.168.2.3 (192.168.2.3), 30 hops max, 60 byte packets
 1  192.168.1.1 (192.168.1.1)  0.031 ms  0.006 ms  0.005 ms
 2  10.10.10.2 (10.10.10.2)  0.022 ms  0.010 ms  0.009 ms
 3  192.168.2.3 (192.168.2.3)  0.018 ms  0.013 ms  0.012 ms
```

Thank you

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com