




Campus Networking Workshop

Introduction to OSPF

Modified from originals by Philip Smith



These materials are licensed under the Creative Commons *Attribution-Noncommercial 3.0 Unported* license
(<http://creativecommons.org/licenses/by-nc/3.0/>)

IPv4

- Internet Protocol version 4
 - Addresses are 32 bits long
 - An IPv4 address has a network portion and a host portion
 - “Subnet mask” used to mark the separation
 - Represents the number of network bits
 - The remaining bits are “host” bits
 - Written as
 - 12.34.56.78 255.255.255.0 or
 - 12.34.56.78/24

IPv6

- Internet Protocol version 6
 - Addresses are 128 bits long
 - An IPv6 address also has a network portion and a host portion
 - “Prefix Length” used to mark the separation
 - Represents the number of network bits
 - The remaining bits are “host” bits
 - Written as
 - 2001:db8::/32

What do routers do?

- An IP packet whose source address has different subnet bits than its destination address, has to be forwarded by a router
- A “route” is a pointer that says: “to reach this subnet, send the packet to this router”
 - 12.23.45.0/24 -> 11.22.33.44
- Routes can be “static” if they are configured manually, or dynamic if they are learned from other routers

Routing vs Forwarding

- Routing is not the same as Forwarding
- Routing is the building of maps
 - Each routing protocol usually has its own routing database
 - Routing protocols populate the forwarding table
- Forwarding is passing the packet to the next hop device
 - Forwarding table contains the best path to the next hop for each prefix
 - There is only ONE forwarding table

Routing Protocols

- Mechanisms that routers follow to exchange routes
 - Can be standard (RIP, OSPF, ISIS, BGP) or proprietary (EIGRP)
 - Categories: Interior (IGP), or exterior (EGP)
 - Categories: *Link state* or *Distance Vector*

OSPF Background

- Developed by IETF – RFC1247
 - Designed for Internet TCP/IP environment
- OSPF v2 described in RFC2328/STD54
- Link state/Shortest Path First Technology
- Dynamic Routing
- Fast Convergence
- Route authentication


Link State Algorithm

- Each router contains a database containing a map of the whole topology
 - Links
 - Their state (including cost)
- All routers have the same information
- All routers calculate the best path to every destination
- Any link state changes are flooded across the network
 - “Global spread of local knowledge”

Link State Routing

- Automatic neighbour discovery
 - Neighbours are physically connected routers
- Each router constructs a Link State Packet (LSP)
 - Distributes the LSP to neighbours...
 - ...using an LSA (Link State Advertisement)
- Each router computes its best path to every destination
- On network failure
 - New LSPs are flooded
 - All routers recompute shortest path tree

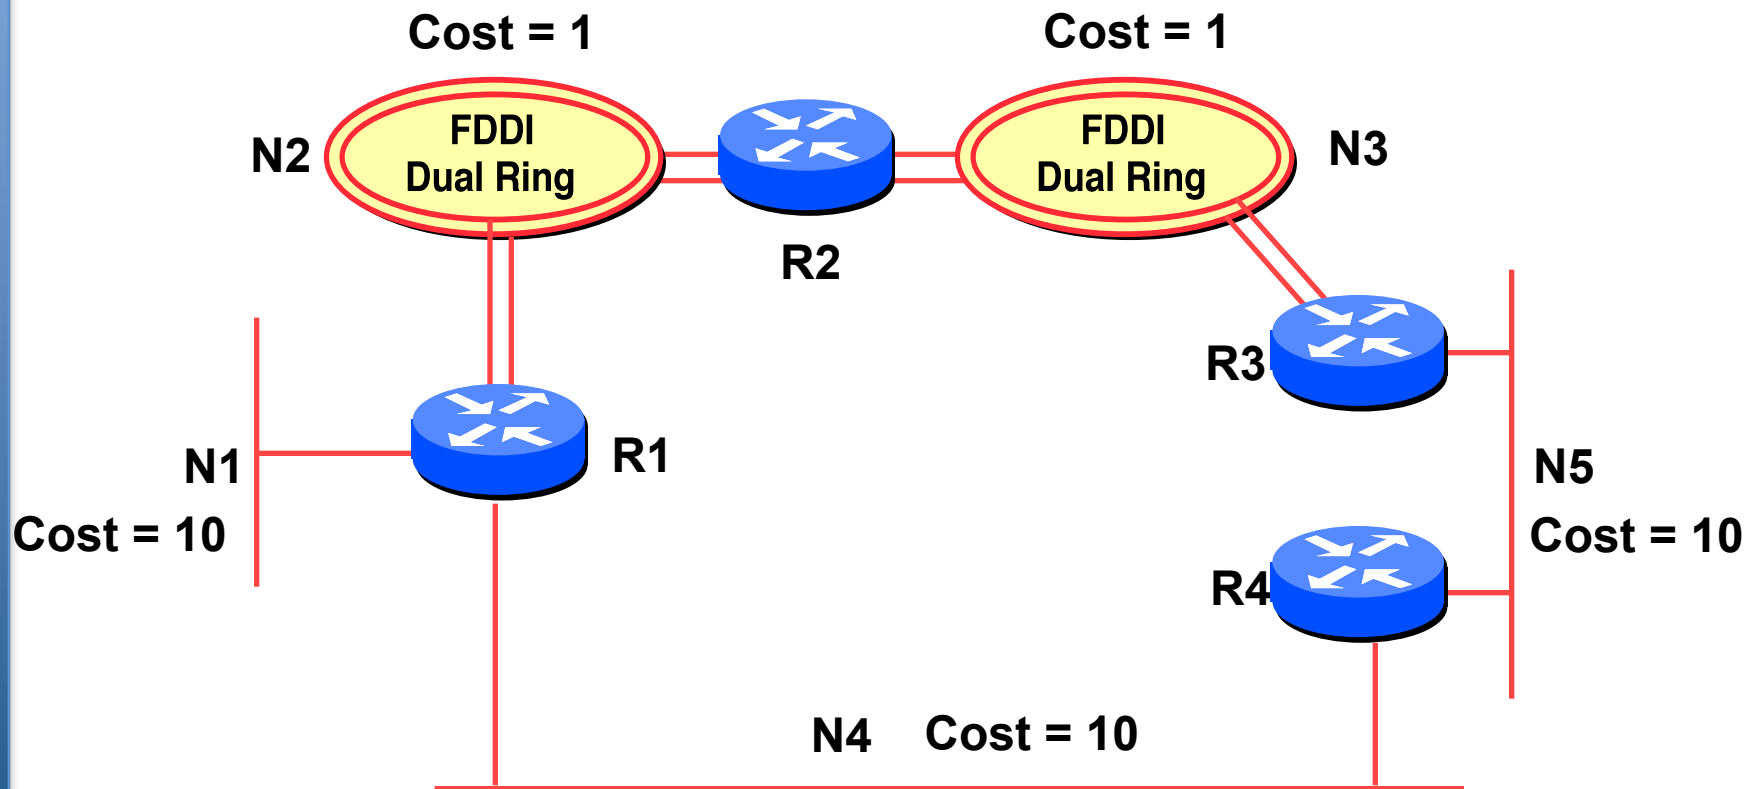
Low Bandwidth Requirements



- Only changes are propagated
- Multicast used on multi-access broadcast networks
 - 224.0.0.5 used for all OSPF speakers
 - 224.0.0.6 used for DR and BDR routers

“Shortest Path First”


The optimal path is determined by the sum of the interface costs: $\text{Cost} = 10^8 / \text{bandwidth}$



OSPF: How it works

Hello Protocol

- Responsible for establishing and maintaining neighbour relationships
- Elects Designated Router on broadcast networks



OSPF: How it works

- Hello Protocol
 - Hello Packets sent periodically on all OSPF enabled interfaces
 - Adjacencies formed between **some** neighbours
- Hello Packet
 - Contains information like Router Priority, Hello Interval, a list of known neighbours, Router Dead Interval, and the network mask

OSPF: How it works

- Trade Information using LSAs
 - LSAs are added to the OSPF database
 - LSAs are passed on to OSPF neighbours
- Each router builds an identical link state database
- SPF algorithm run on the database
- Forwarding table built from the SPF tree

OSPF: How it works

When change occurs:

- Announce the change to all OSPF neighbours
- All routers run the SPF algorithm on the revised database
- Install any change in the forwarding table


Broadcast Networks

- These are network technologies such as Ethernet
- Introduces Designated and Backup Designated routers (DR and BDR)
 - Only DR and BDR form full adjacencies with other routers
 - The remaining routers remain in a “2-way” state with each other
 - If they were adjacent, we’d have n-squared scaling problem
 - If DR or BDR “disappear”, re-election of missing router takes place

Designated Router

One per multi-access network


- Generates network link advertisements for the multi-access network
- Speeds database synchronisation



Designated Router

- All routers are adjacent to the DR
 - All routers are adjacent to the BDR also
- All routers exchange routing information with DR
 - BDR also stays synchronized with the DR
- DR updates the database of all its neighbours
 - BDR waits silently and only takes over if DR dies
- This scales!
 - $2n$ problem rather than having an n -squared problem.


Designated Router



- Adjacencies only formed with DR and BDR
- LSAs propagate along the adjacencies

Designated Router Priority

- Determined by interface priority
- Otherwise by highest router ID
 - (For Cisco IOS, this is address of loopback interface, otherwise highest IP address on router)






More Advanced OSPF

- OSPF Areas
- Virtual Links
- Router Classification
- OSPF route types
- External Routes
- Route authentication
- Equal cost multipath

OSPF Areas

- Group of contiguous hosts and networks
- Per area topological database
 - Invisible outside the area
 - Reduction in routing traffic
- Backbone area contiguous
 - All other areas must be connected to the backbone
- Virtual Links



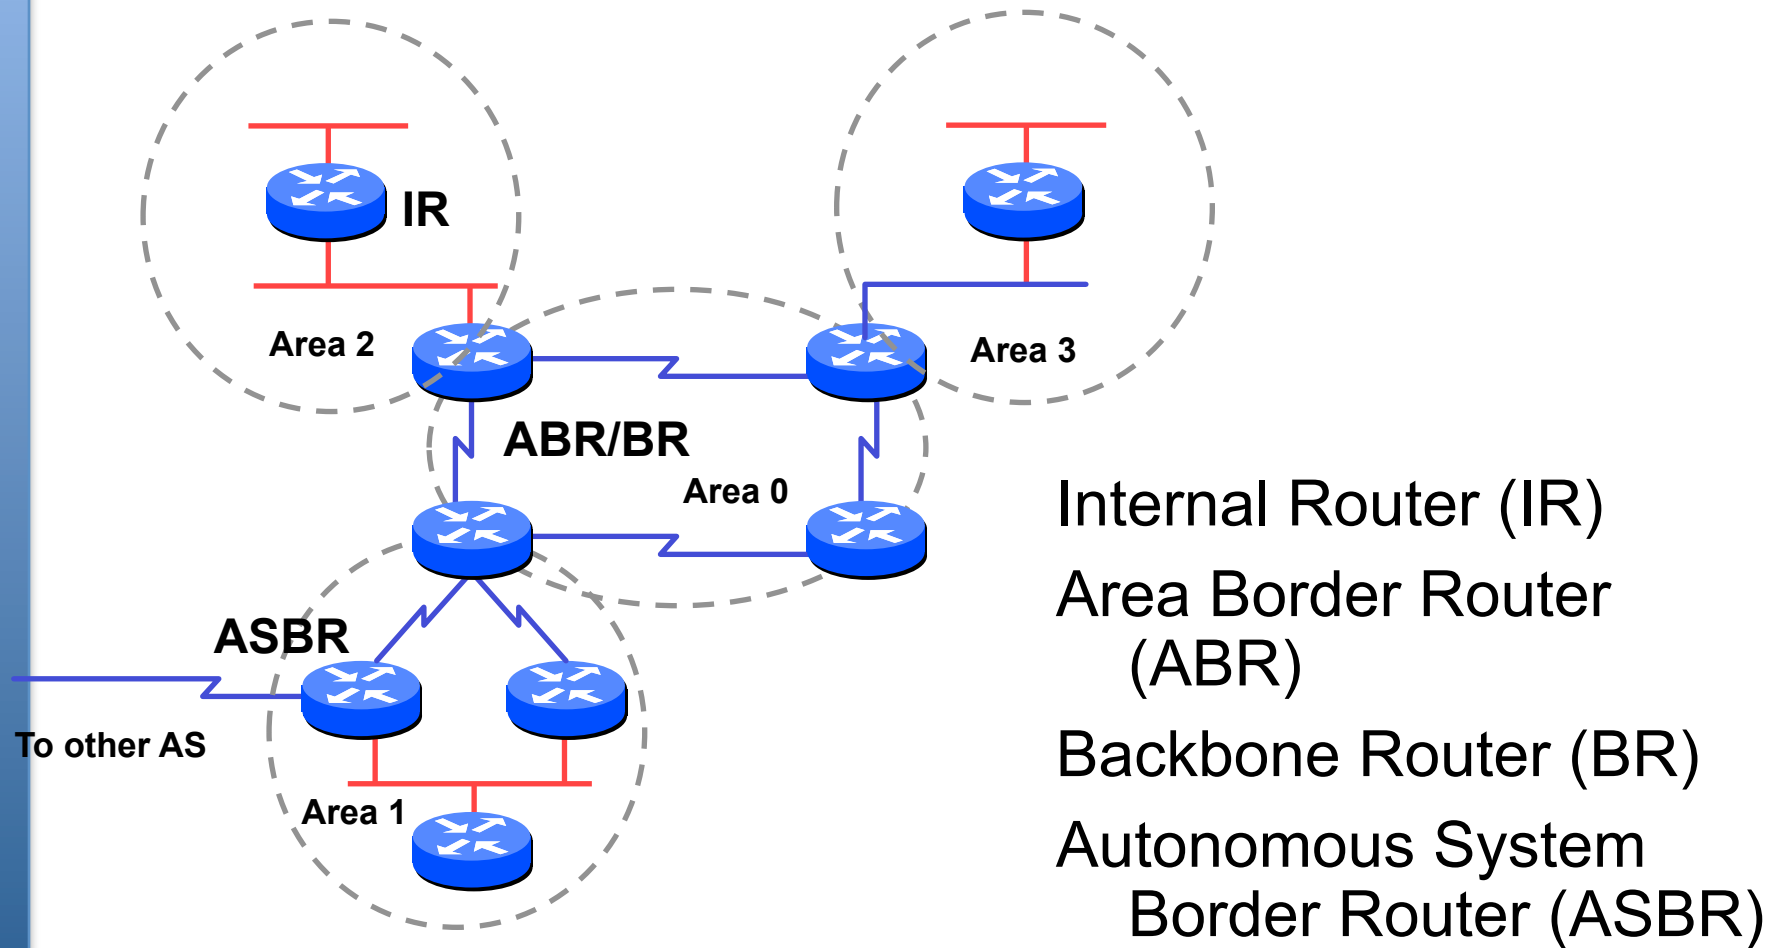
OSPF Areas

- Reduces routing traffic in area 0
- Consider subdividing network into areas
 - Once area 0 is more than 10 to 15 routers
 - Once area 0 topology starts getting complex
- Area design often mimics typical ISP core network design
- Virtual links are used for “awkward” connectivity topologies

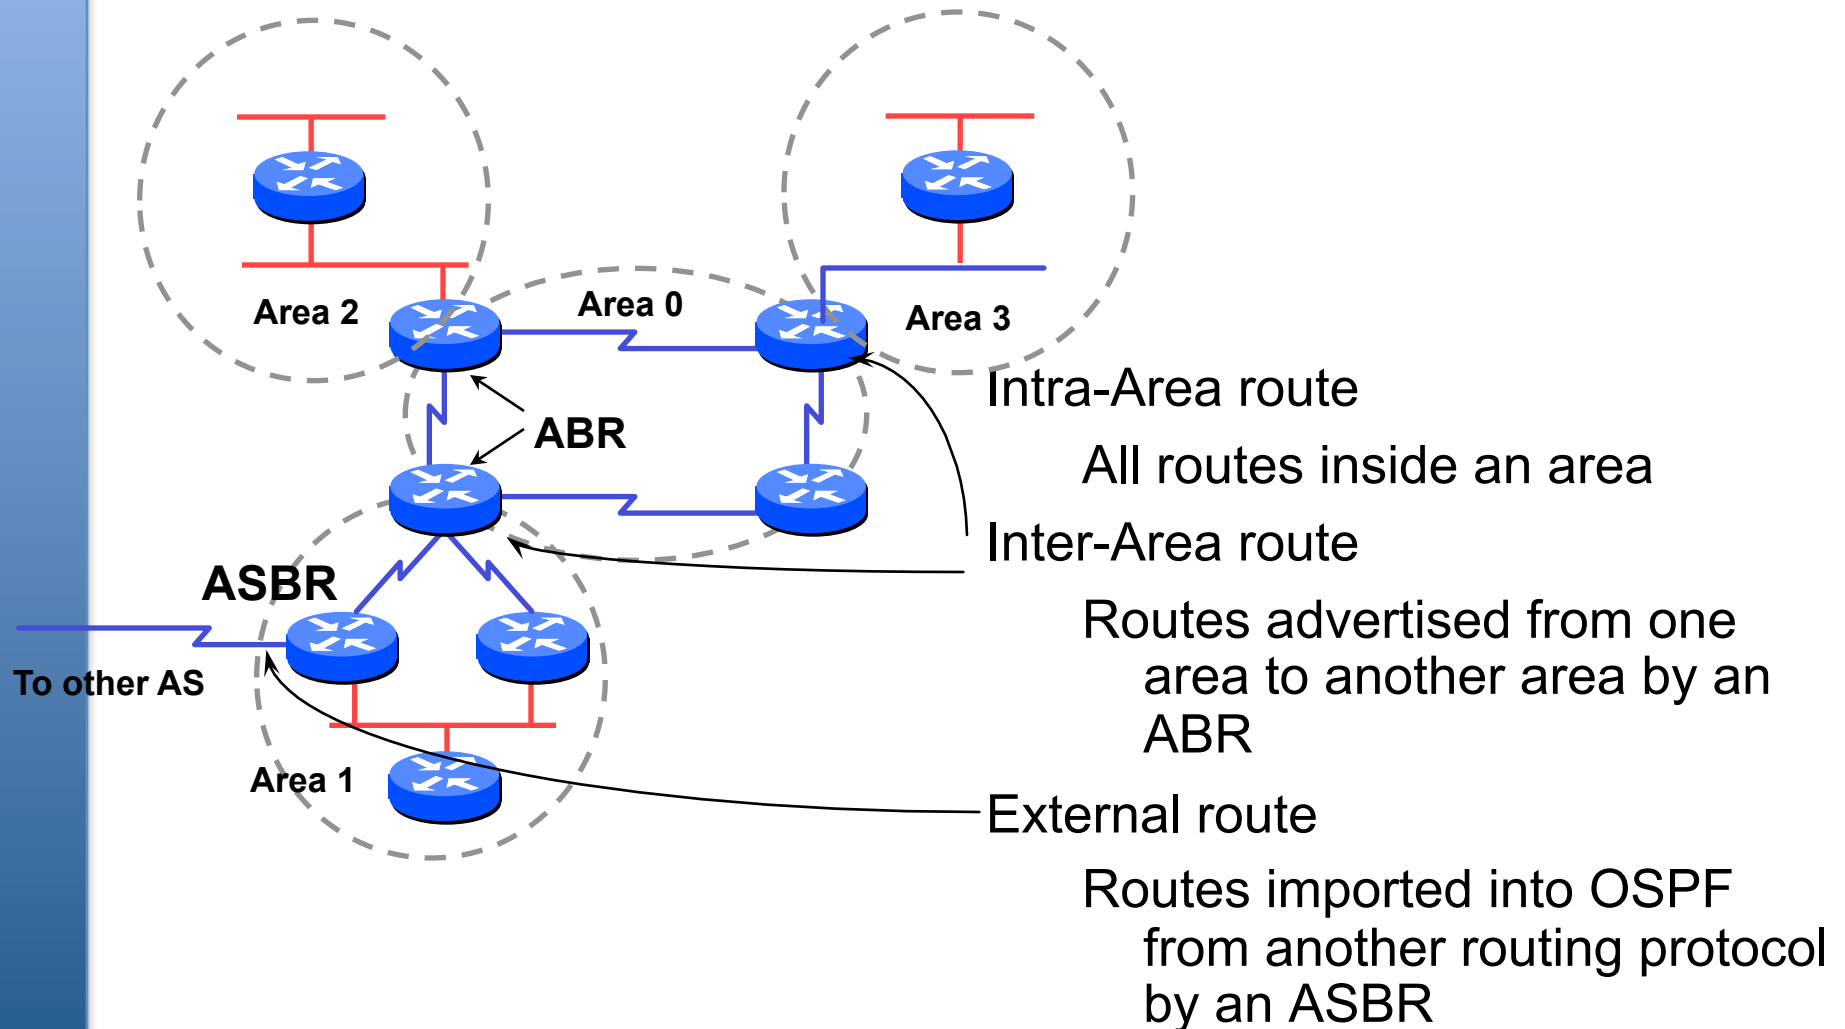
Virtual Links

- OSPF requires that all areas **MUST** be connected to area 0
- If topology is such that an area cannot have a physical connection to a device in area 0, then a virtual link must be configured
- Otherwise the disconnected area will only be able to have connectivity to its immediately neighbouring area, and not the rest of the network

Classification of Routers

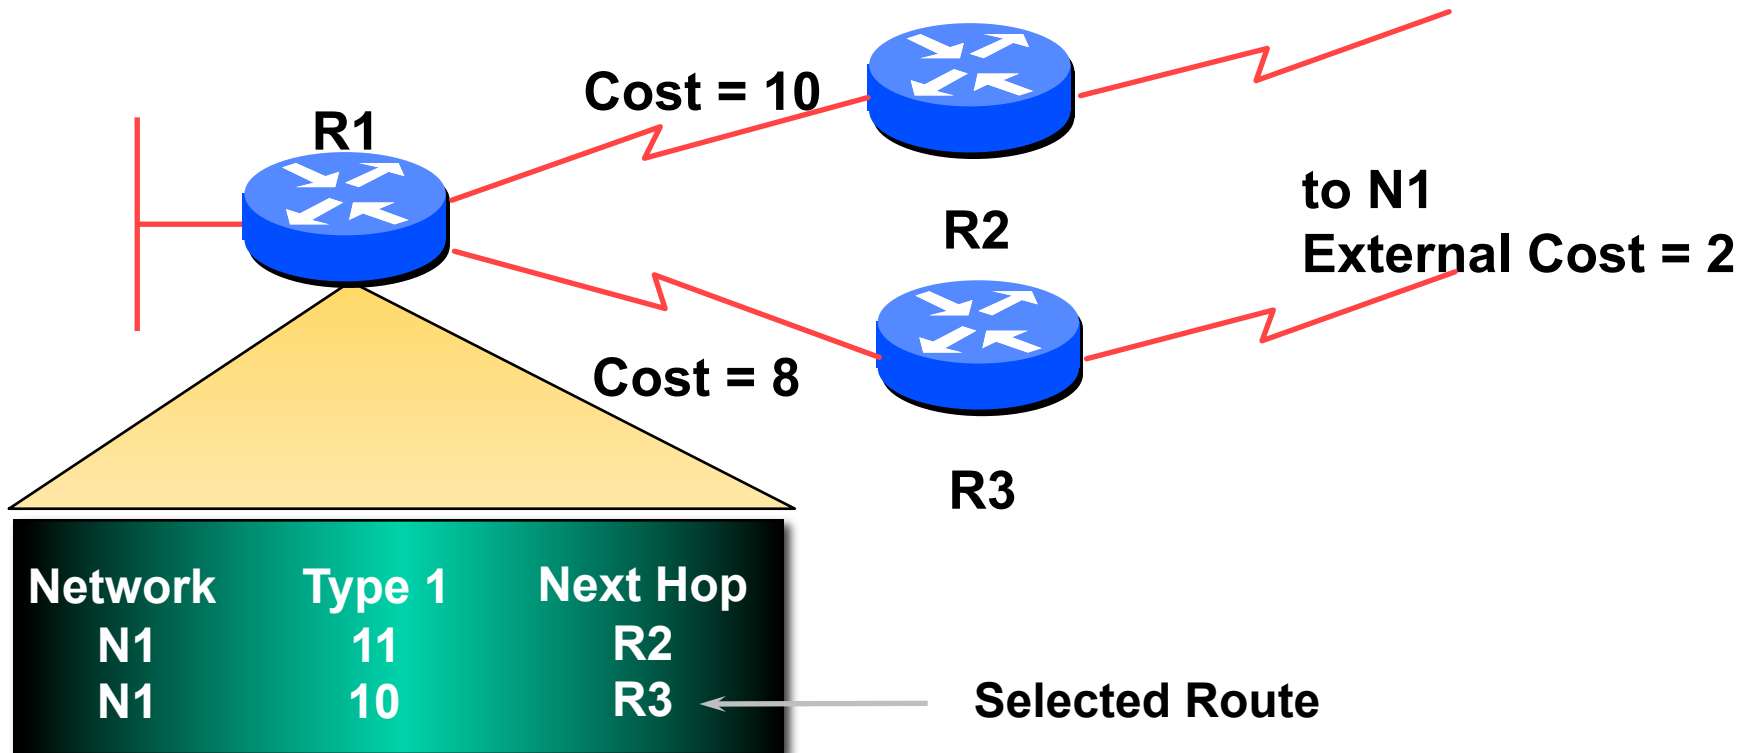


OSPF Route Types



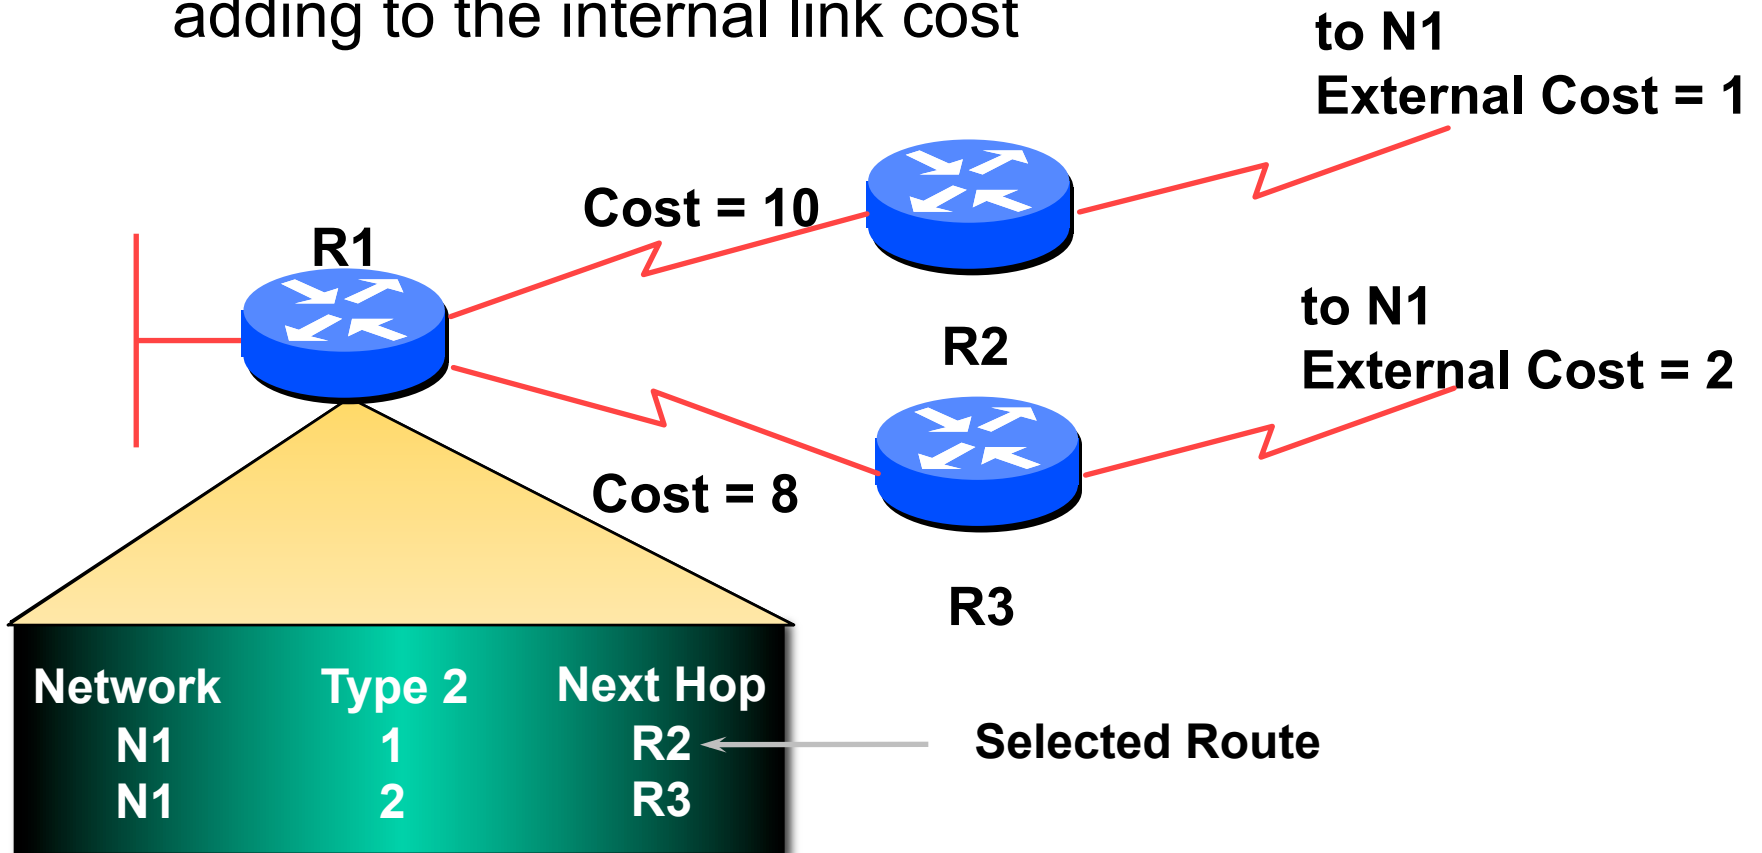
External Routes

Type 1 external metric: metrics are added to the summarised internal link cost



External Routes

Type 2 external metric: metrics are compared without adding to the internal link cost



Route Authentication

It is recommended to use route authentication for OSPF

...and all other routing protocols

Susceptible to denial of service attacks

OSPF runs on TCP/IP

Automatic neighbour discovery

Equal Cost Multipath

If n paths to same destination have equal cost, OSPF will install n entries in the forwarding table

Loadsharing over the n paths

Useful for expanding links across an ISP backbone

Don't need to use hardware multiplexors

Don't need to use static routing

OSPFv3

- OSPFv2 only supports IPv4
- OSPFv3 developed for IPv6 only
 - Dual stack networks need to run both protocols
 - They run independently of each other

OSPFv2 vs. OSPFv3

- Very similar, with a few differences
 - New LSA types to separate links from their prefixes
 - Avoids SPF recalculations when only the link prefix changes
 - Removes OSPF-specific authentication
 - Relies on underlying IPv6 security headers
 - Supports multiple instances

OSPF Configuration – Start

- Start the OSPF process

```
router ospf 100
```

- “100” is the process ID
 - Process ID is unique to the router
 - You can run multiple OSPF processes on the same router
 - It is common to set the process ID to autonomous system (AS) number

OSPF Configuration

Advertising Networks Option 1

```
router OSPF 100
```

```
  redistribute connected subnets
```

- Includes all the subnets connected to the router, but announces them as external type-2 LSAs, which are not summarized
 - Not a good option when using multiple areas
- It also does not give you control over which networks you want to announce

OSPF Configuration

Advertising Networks Option 2

- **Per-link configuration (IOS 12.4 and later)**
 - OSPF configured on each interface
 - Passive interfaces do not establish adjacencies

```
router ospf 100
  passive-interface default
  no passive interface GigabitEthernet0
interface GigabitEthernet0
  ip address 10.10.0.1 255.255.255.0
  ip ospf 100 area 0
interface FastEthernet0
  ip address 192.168.10.1 255.255.255.0
  ip ospf 100 area 0
```

OSPF Configuration

Advertising Networks Option 3

- **Network statements**
 - Announce subnets included in given network blocks
 - Add one network statement per subnet, or use wildcard masks to include multiple subnets

```
ospf router 100
  passive-interface default
  no passive-interface GigabitEthernet0
  network 10.10.0.0 0.0.0.255 area 0
  network 192.168.0.0 0.0.0.3 area 0
```

OSPF Configuration

Authenticating adjacencies

- **Very important to control which devices can establish adjacencies**
 - **A malicious user could inject routes and steal your traffic!**

```
router ospf 100
  area 0 authentication message-digest
interface GigabitEthernet0/0
  ip ospf authentication-key <key>
```



Questions?