Routing Introduction

Direct vs. Indirect Delivery

Static vs. Dynamic Routing

Distance Vector vs. Link State

“The most simple way to accelerate a Router is at 9.8 m/sec/sec.”

Seen on Usenet

Routing Basics

- Routing Introduction
  - Direct Delivery
  - Indirect Delivery
  - Static Routing
  - Default Routing
- Dynamic Routing
  - Distance Vector Routing
  - Link State Routing

What is routing?

- ***Finding a path to a destination address***
- **Direct delivery performed by host**
  - **Destination network = local network**
- **Indirect delivery performed by router**
  - **Destination network ≠ local network**
  - **Packet is forwarded to default gateway**

Direct Delivery

- IP host checks if packet's destination network is identical with local network
  - By applying the configured subnet mask of the host's interface
- If destination network = local network then the L2 address of the destination is discovered using ARP
  - Not necessary on point-to-point connections

IP Host Facts

- Also IP hosts have routing tables !
  - But typically only a static route to the default gateway is entered
- ARP cache aging timer: 20 minutes

Indirect Delivery

- Default gateway delivers packet in behalf of its host using a routing table
- Routing table components
  - Destination network (+ subnet mask)
  - Next hop (+ outgoing interface)
  - Metric (+ Administrative Distance)

Router

- **Initially Unix workstations with several network interface cards**
- **Today specialized hardware**

Cisco 3600 Router

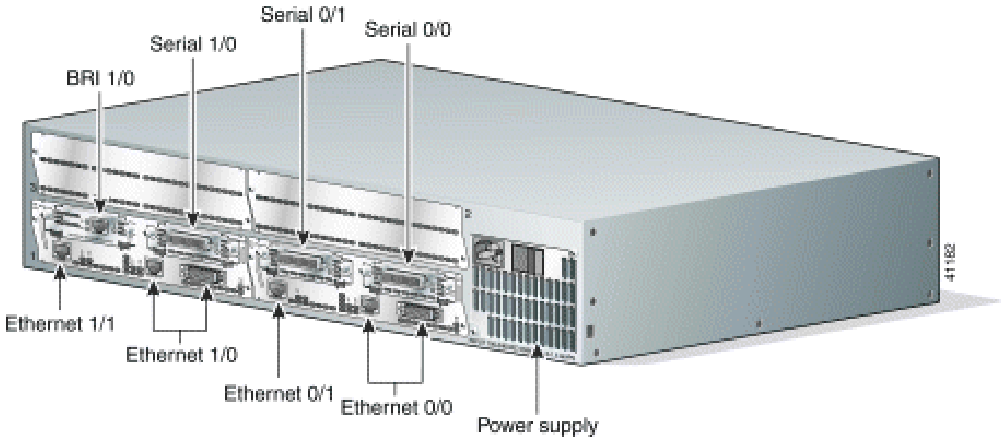

Routing Table Example

```
Gateway of last resort is 175.18.1.2 to network 0.0.0.0

     10.0.0.0 255.255.0.0 is subnetted, 4 subnets
C       10.1.0.0 is directly connected, Ethernet1
R       10.2.0.0  [120/1] via 10.4.0.1, 00:00:05, Ethernet0
R       10.3.0.0  [120/5] via 10.4.0.1, 00:00:05, Ethernet0
C       10.4.0.0 is directly connected, Ethernet0
R    192.168.12.0  [120/3] via 10.1.0.5, 00:00:08, Ethernet1
S    194.30.222.0 [1/0] via 10.4.0.1
S    194.30.223.0 [1/0] via 10.1.0.5
C    175.18.1.0 255.255.255.0 is directly connected, Serial0
S*   0.0.0.0 0.0.0.0 [1/0] via 175.18.1.2
```

IP Routing Basics

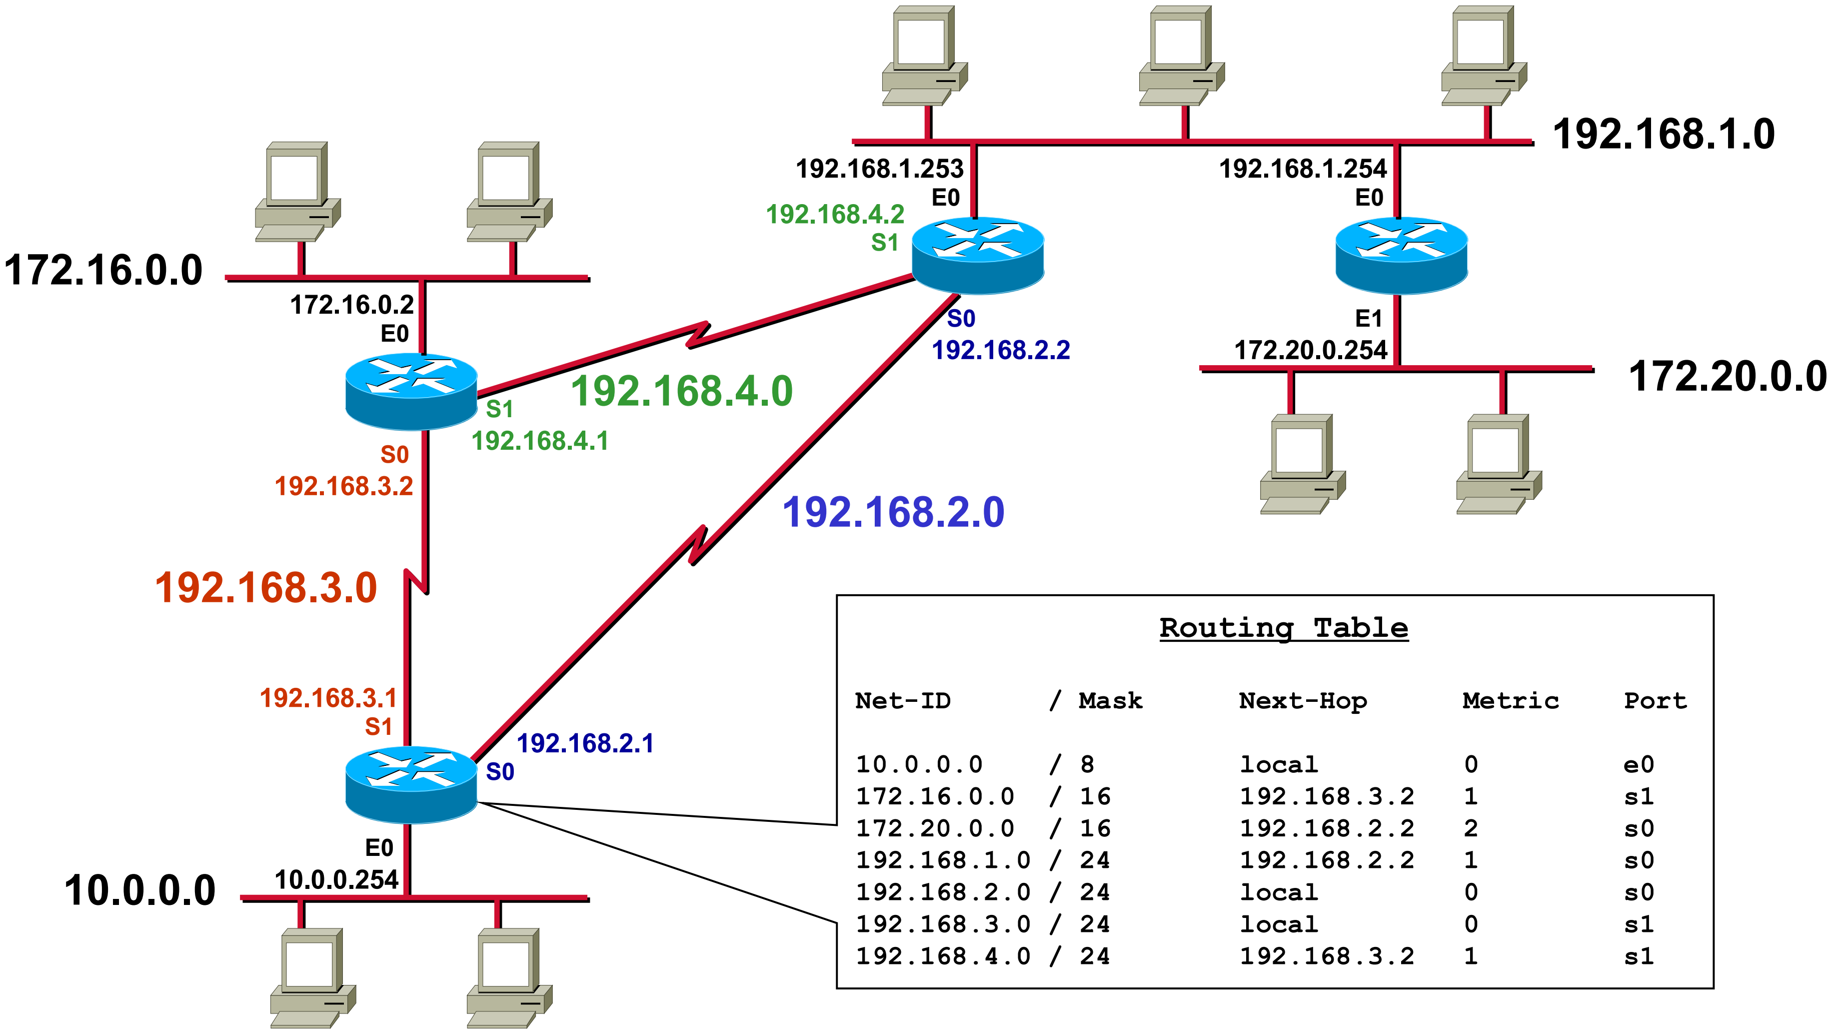

Routing Table

Net-ID	/ Mask	Next-Hop	Metric	Port
10.0.0.0	/ 8	local	0	e0
172.16.0.0	/ 16	192.168.3.2	1	s1
172.20.0.0	/ 16	192.168.2.2	2	s0
192.168.1.0	/ 24	192.168.2.2	1	s0
192.168.2.0	/ 24	local	0	s0
192.168.3.0	/ 24	local	0	s1
192.168.4.0	/ 24	192.168.3.2	1	s1

Static or Dynamic

- **Static routing entries are configured manually**
  - **Override routes learned via dynamic routing**
  - ***Can* be set as permanent (will not be removed if interface goes down)**
  - **Only way for certain technologies (DDR)**
- **Dynamic routing entries are learned by routing protocols**
  - **Adapts to topology changes**
  - **But additional routing-traffic overhead**

Reasons for Static Routing

- Very low bandwidth links (e. g. dialup links)
- Administrator needs control over the link
- Backup links
- Link is the only path to a stub network
- Router has very limited resources and cannot run a routing protocol

ip route *prefix mask* **{***ip-address* **|** *interface-type interface-number***}** **[***distance***]** **[tag** *tag***]** **[permanent]**

Tag value that can be used as a "match" value for controlling redistribution via route maps

Specifies that the route will not be removed, even if the interface shuts down

Routing Paradigm

- Destination Based Routing
  - Source address is not taken into account for the forward decision
- Hop by Hop Routing
  - IP datagram's follow the signposts given by routing table entries
  - Network's routing state must be loop-free and consistent
- Least Cost Routing
  - Typically only the best path is entered into routing table

Static Routing (1)

- Static routes to and from stub networks

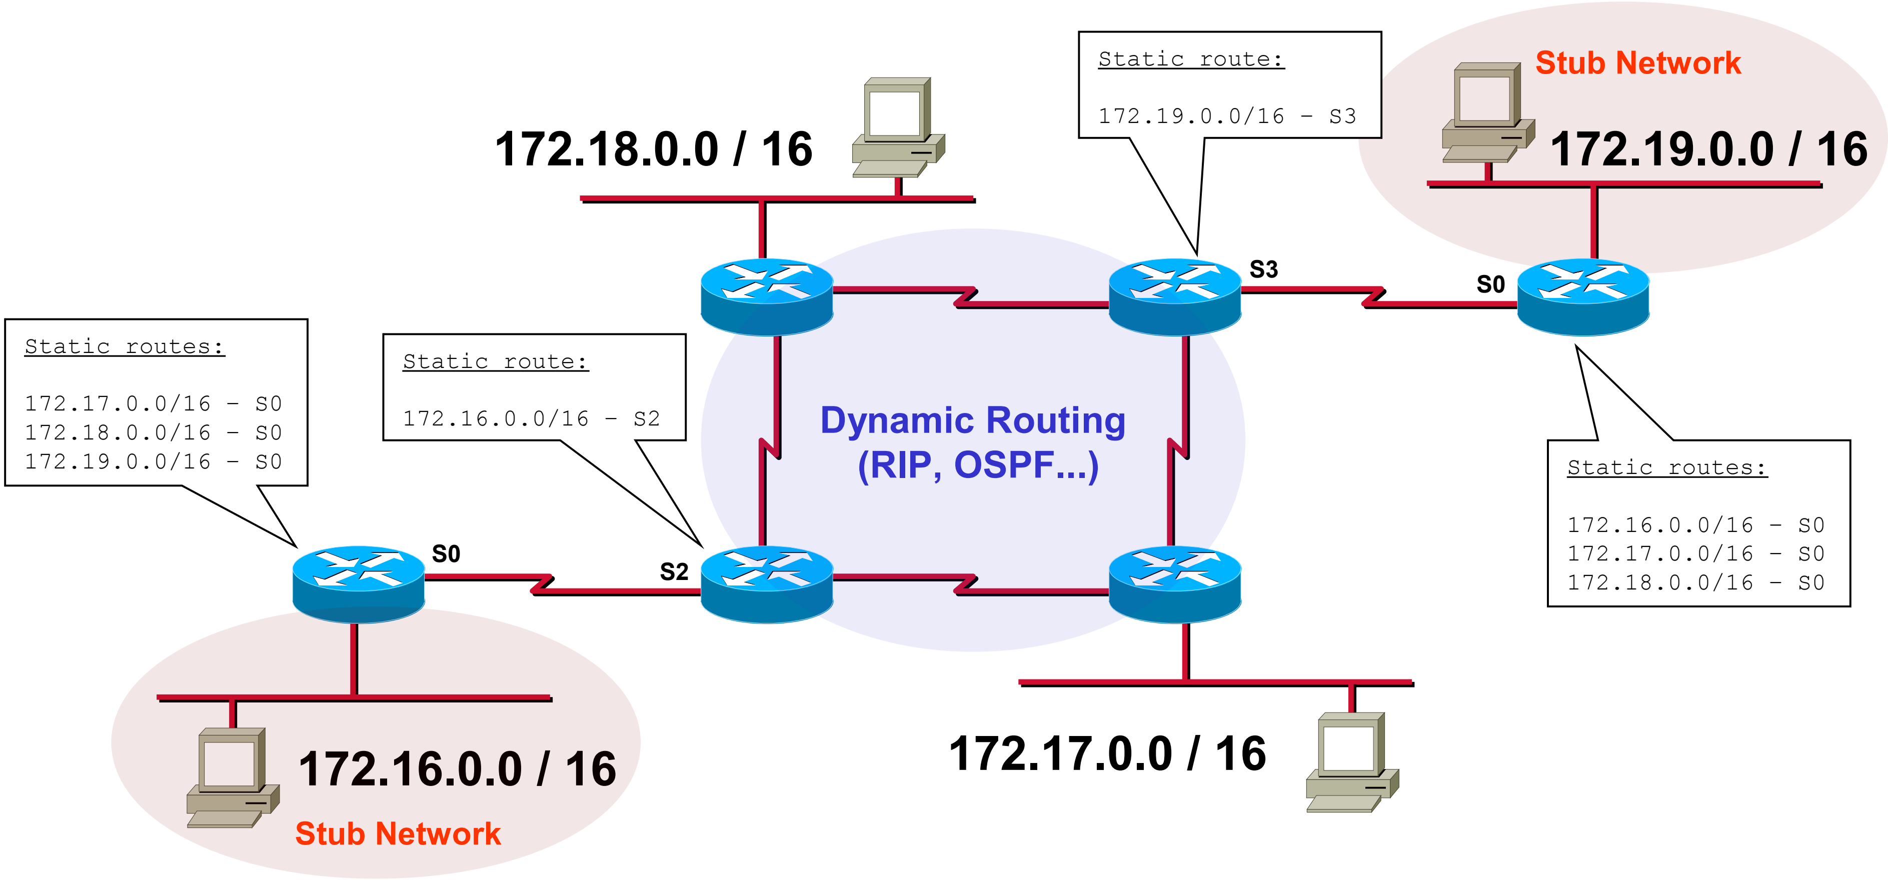

Static Routing (2)

- Static routes in "Hub and Spoke" topologies

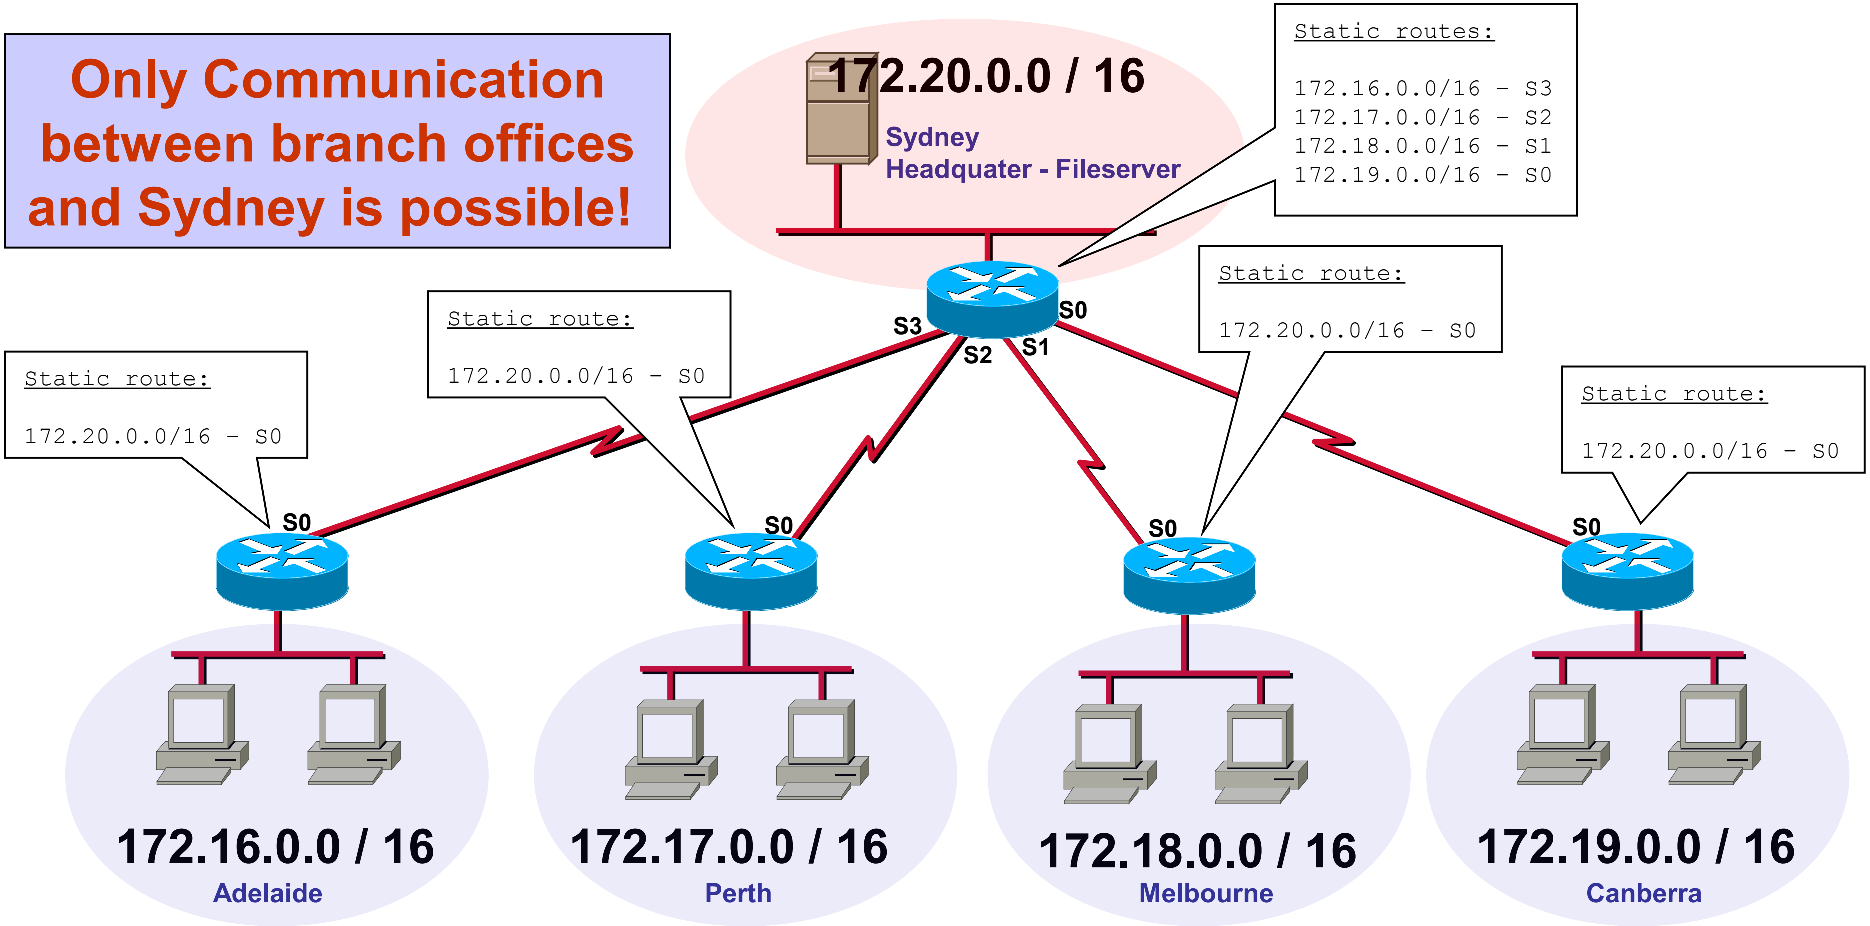

Default Routing

- Special static route
  - Traffic to unknown destinations are forwarded to default router ("Gateway of Last Resort")
- Routing table entry "0.0.0.0 0.0.0.0"
- Hopefully, default gateway knows more destination networks
- Advantage: Smaller routing tables!

Default Routing (1)

- Default Routes from stub networks

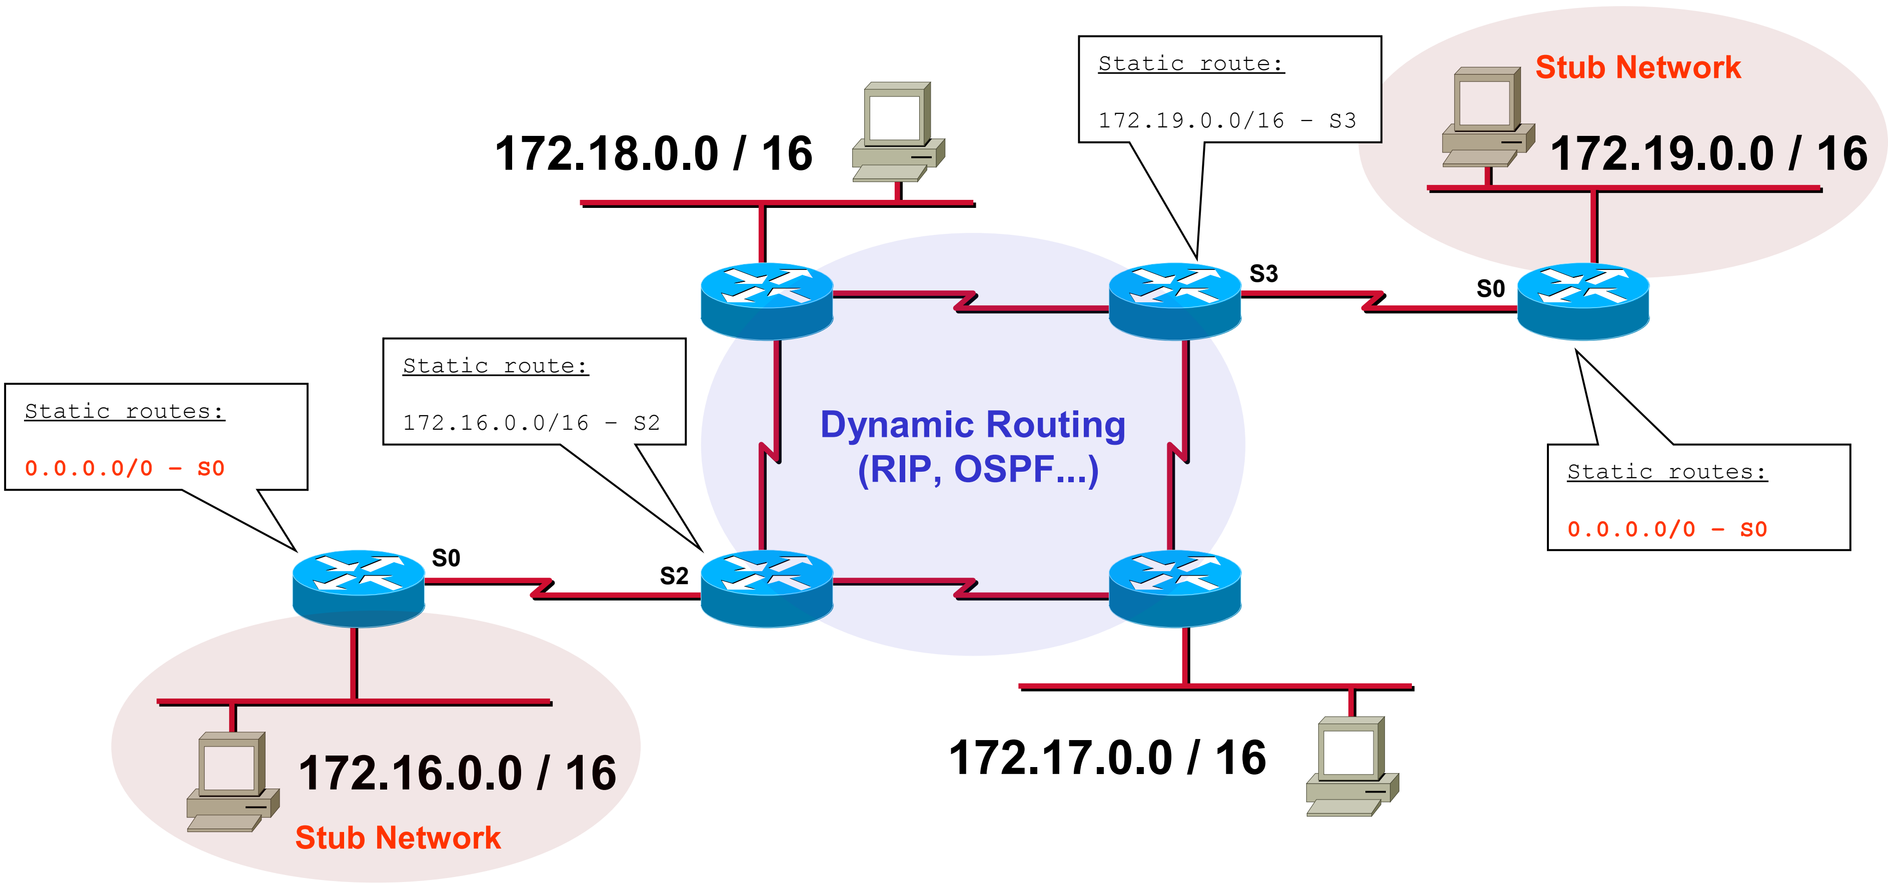

Default Routing (2)

- **Default routes in "Hub and Spoke" topologies**

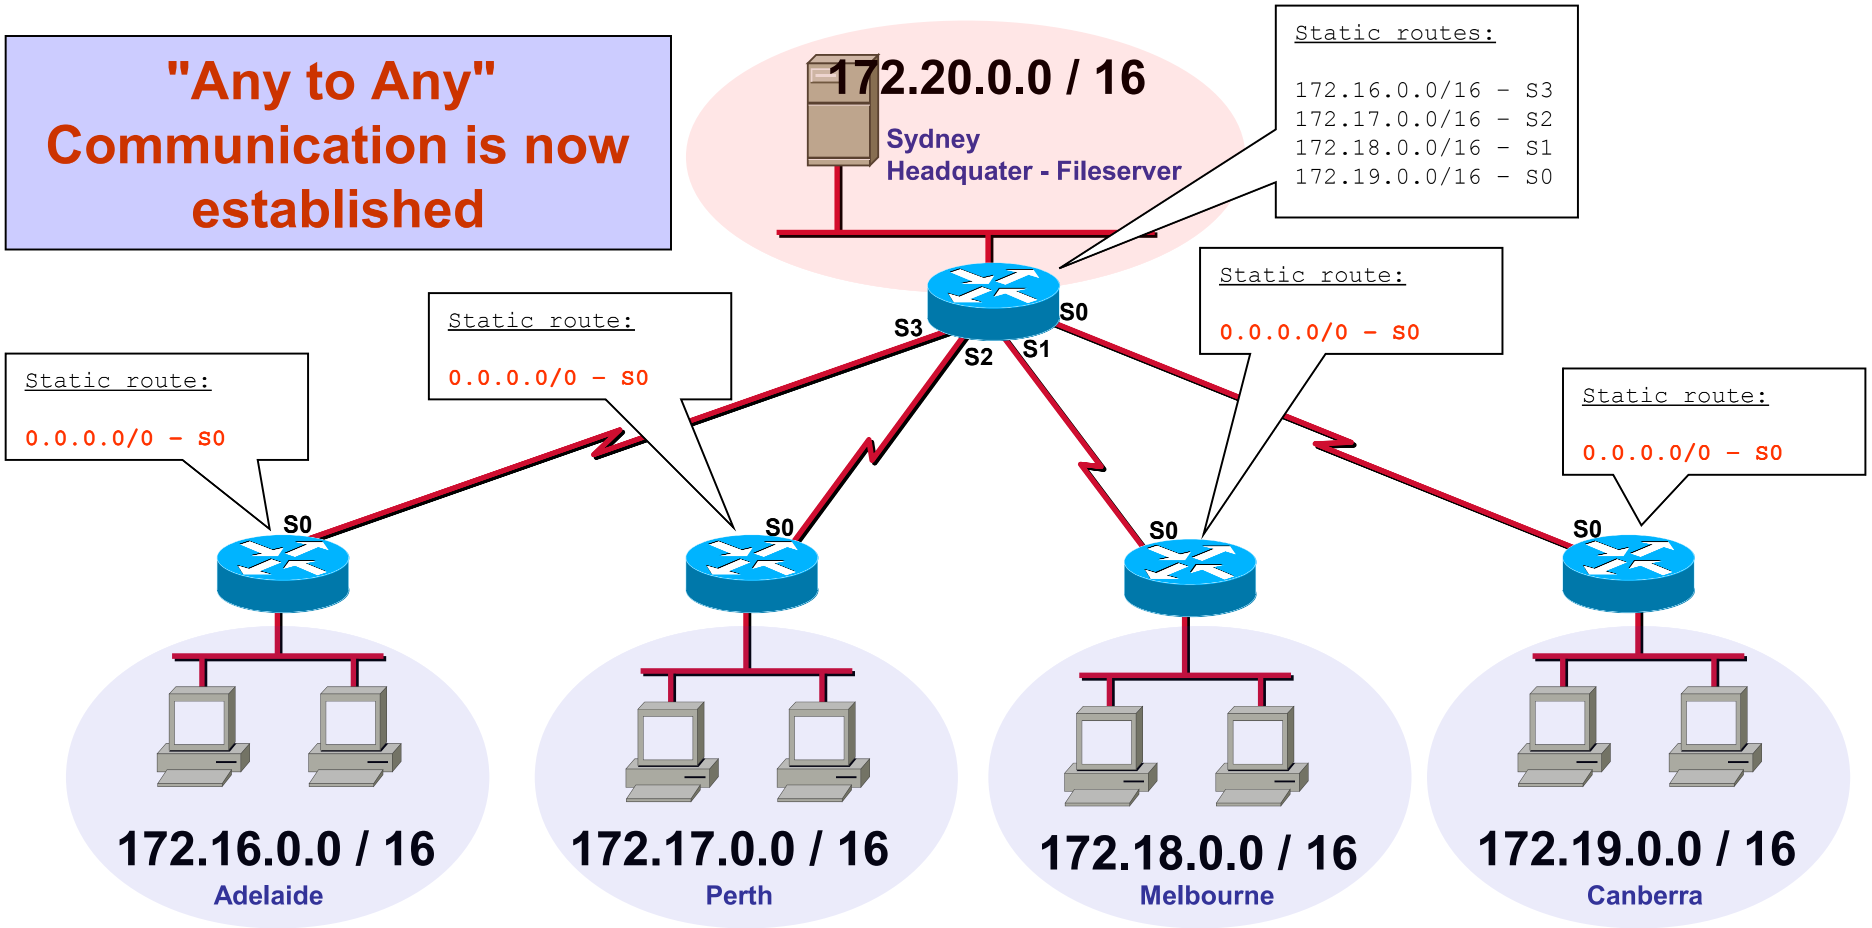

Default Routing (3)

- **Default Routes to the Internet**

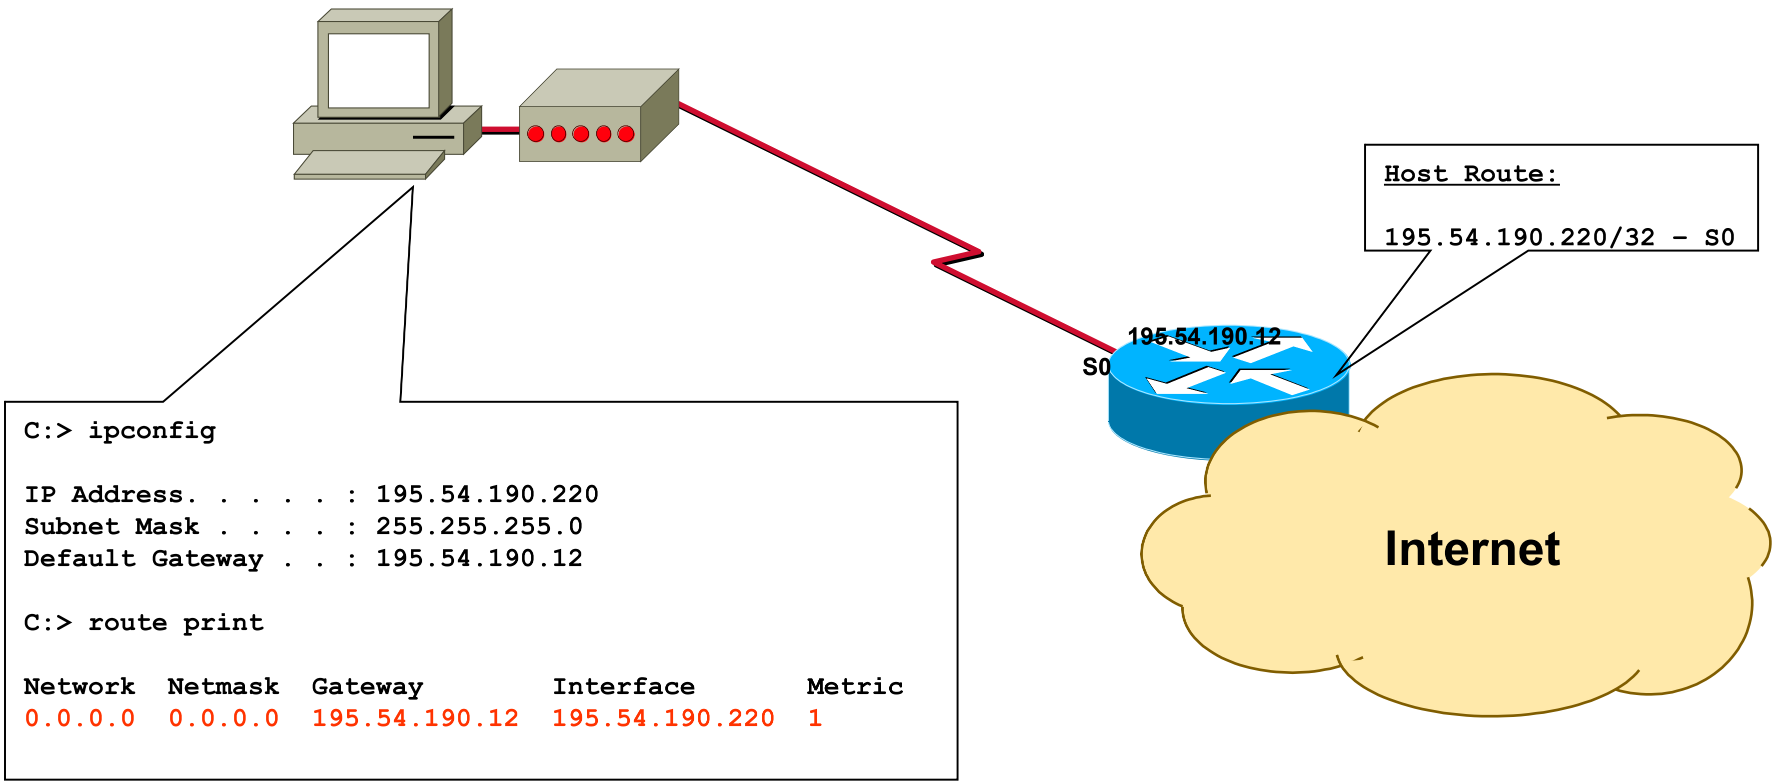

On Demand Routing (ODR)

- **Efficient for hub-and-spoke topologies**
  - **Same configuration at each router**
- **Uses CDP to send the prefixes of attached networks from the spokes, or stub networks, to the hub or core router**
  - **CDP does this automatically (!)**
- **The hub router sends its interface address of the shared link as the default route for the stub router**
- **Note:**
  - **Don't enable routing protocols on spoke routers**
  - **CDP must be enabled (don't forget e. g. ATM interfaces)**
  - **Every 60 sec a CDP message is sent per default (change with "cdp timer" command)**

```
(config)# router odr                                   ! Only on hub router
```

Dynamic Routing

- **Each router can run one or more routing protocols**
- **Routing protocols are information sources to create routing table**
- **Routing protocols differ in convergence time, loop avoidance, network size, complexity**

Routing Protocol Comparison

Routing Protocol	Complexity	Max. Size	Convergence Time	Reliability	Protocol Traffic
RIP	very simple	16 Hops	Up to 480 secs	Not absolutely loop-safe	High
RIPv2	very simple	16 Hops	Up to 480 secs	Not absolutely loop-safe	High
IGRP	simple	x	x	medium	medium
EIGRP	complex	x	x	x	x
OSPF	very complex	Thousands of Routers	Fast	High	low/ depends
IS-IS	complex	Thousands of Routers	Fast	High	x
BGP-4	complex	more than 100,000 networks	Fast	Very High	x

Metric

- Routing protocols typically find out more than one route to the destination
- Metrics help to decide which path to use
  - Hop count
  - Cost (reciprocal value of bandwidth)
  - Load, Reliability, Delay, MTU

Administrative Distance

- Several routing protocols independently find out different routes to same destination
  - Which one to choose?
- "Administrative Distance" is a trustiness-value associated to each routing protocol
  - The lower the better
  - Can be changed

Administrative Distances Chart

Unknown	255
I-BGP	200
E-EIGRP	170
EGP	140
RIP	120
IS-IS	115
OSPF	110
IGRP	100
I-EIGRP	90
E-BGP	20
EIGRP Summary Route	5
Static route to next hop	1
Static route through interface	0
Directly Connected	0

Remember

1) Using the METRIC *one* routing protocol determines the best path to a destination.

2) A router running multiple routing protocols might be told about multiple possible paths to one destination.

3) Here the METRIC cannot help for decisions because *different type of METRICS* cannot be compared with each other.

4) A router chooses the route which is proposed by the routing protocol with the *lowest* ADMINISTRATIVE DISTANCE

AD with Static Routes

- Each static route can be given a different administrative distance
- This way fall-back routes can be configured

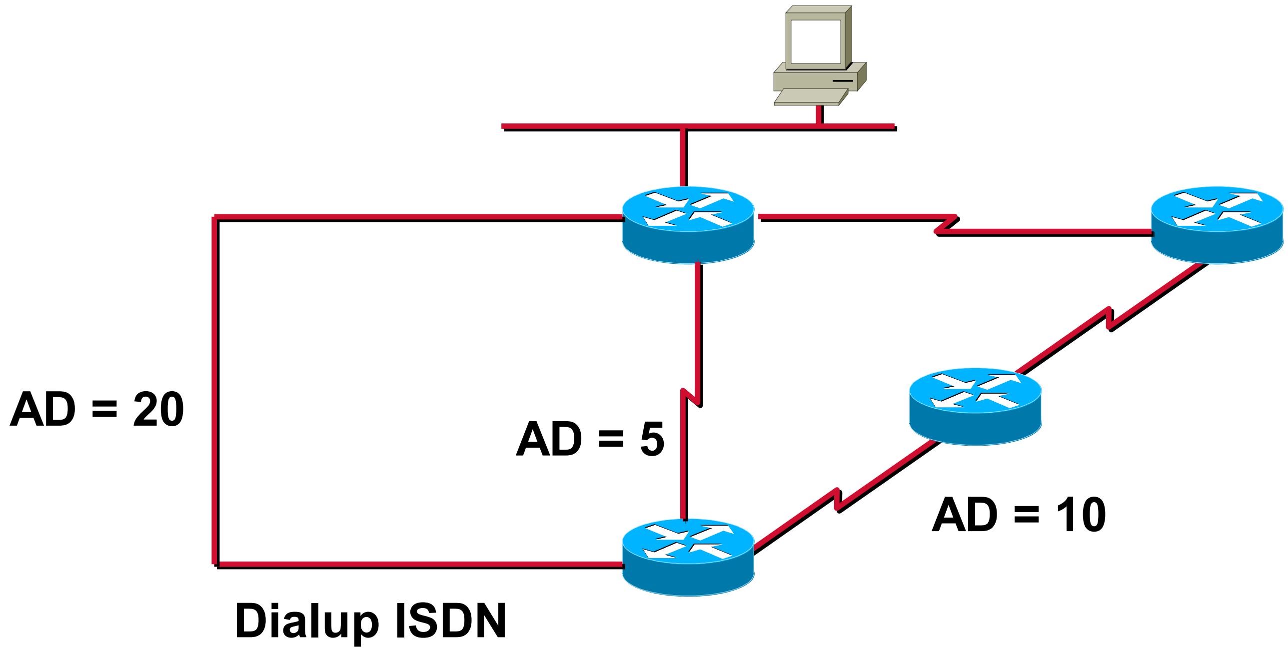

Classification

- Depending on age:
  - Classful (no subnet masks)
  - Classless (VLSM/CIDR supported)
- Depending on scope:
  - IGP (Inside an Autonomous System)
  - EGP (Between Autonomous Systems)
- Depending on algorithm:
  - Distance Vector (Signpost principle)
  - Link State (Roadmap principle)

Classful route match (1/2)

- 1) If the classful network number is NOT listed in the routing table, use the default route if available (otherwise discard the packet)
- 2) If the classful network number is listed in the routing table:
  - If the listed network number is NOT subnetted and matches the IP-packet's destination address then use this route
  - If this network is subnetted, then lookup the corresponding subnet; if no subnet matches then *discard the packet* (even if a default route exists!)

Example:

IP Packet
DA = 10.35.72.26
SA = …

Routing Table:

```
10.0.0.0/8 is subnetted, 4 subnets:
10.22.0.0/16   via 172.17.7.19
10.31.0.0/16   via 172.17.8.31
10.34.0.0/16   via 172.18.1.254
10.35.0.0/16   via 192.186.176.254
0.0.0.0/0      via 172.19.41.254
```

Classful route match (2/2)

- 1) If the classful network number is NOT listed in the routing table, use the default route if available (otherwise discard the packet)
- 2) If the classful network number is listed in the routing table:
  - If the listed network number is NOT subnetted and matches the IP-packet's destination address then use this route
  - **If this network is subnetted, then lookup the corresponding subnet; if no subnet matches then *discard the packet* (even if a default route exists!)**

Example:

IP Packet
DA = 10.35.72.26
SA = …

Routing Table:

```
10.0.0.0/8 is subnetted, 4 subnets:
10.22.0.0/16   via 172.17.7.19
10.31.0.0/16   via 172.17.8.31
10.34.0.0/16   via 172.18.1.254
0.0.0.0/0      via 172.19.41.254
```

DISCARD THE PACKET (!)

Classless routing: Longest match

- **The subnet mask of each route entry tells how many bits must be compared with the IP-packet's destination address**
- **The router takes the route with the longest match**

Example:

IP Packet
DA = 10.35.72.26
SA = ...

→

Routing Table:

10.0.0.0/8	via 172.16.1.1
10.22.0.0/16	via 172.17.7.19
10.31.0.0/16	via 172.17.8.31
10.34.0.0/16	via 172.18.1.254
10.35.0.0/16	via 192.186.176.254
<u>10.35.64.0/19</u>	<u>via 192.186.177.254</u>
10.35.192.0/19	via 172.19.54.1

Distance Vector (1)

- After powering-up each router only knows about directly attached networks
- Routing table is sent periodically to all neighbor-routers
- Received updates are examined, changes are adopted in own routing table
- Metric information (originally) is number of hops
- "Bellman-Ford" algorithm

Distance Vector (2)

- Next hop is always originating router
  - Topology behind next hop unknown
  - Signpost principle
- Loops can occur!
  Additional mechanisms needed:
  - Maximum hop count
  - Split horizon (with poison reverse)
  - Triggered update
  - Hold down
- Examples: RIP, RIPv2, IGRP (Cisco)

Link State (1)

- **Each two neighbored routers establish adjacency**
- **Routers learn real topology information**
  - **Through "Link State Advertisements"**
  - **Stored in database (Roadmap principle)**
- **Updates only upon topology changes**
  - **Propagated by *flooding* (very fast convergence)**

Link State (2)

- Routing table entries are calculated by applying the Shortest Path First (SPF) algorithm on the database
  - Loop-safe
  - Alternative paths immediately known
  - CPU and memory greedy
- Large networks can be split into areas
- Examples: OSPF, Integrated IS-IS

Summary

- Routing is the "art" of finding the best way to a given destination
- Can be static or dynamic
  - Static means: YOU are defining the way packets are going
  - Dynamic means: A routing protocol is "trying" to find the best way to a given destination
- In today's routers the route with the longest match is used
- Routing protocols either implement the principle *Distance Vector* or *Link State*