Review Question

Examine the routing table below

```
Gateway of last resort is not set
     10.0.0.0/27 is subnetted, 1 subnets
S      10.1.10.64 [1/0] via 192.168.1.206
     172.16.0.0/25 is subnetted, 1 subnets
O      172.16.5.128 [110/65] via 192.168.1.202, 00:01:10, Serial0/0
     192.168.1.0/30 is subnetted, 5 subnets
C      192.168.1.192 is directly connected, Serial0/1
O      192.168.1.196 [110/128] via 192.168.1.193, 00:01:10, Serial0/1
                     [110/128] via 192.168.1.202, 00:01:10, Serial0/0
C      192.168.1.200 is directly connected, Serial0/0/0
C      192.168.1.204 is directly connected, FastEthernet0/1
O      192.168.1.208 [110/65] via 192.168.1.202, 00:01:00, Serial0/0
O      192.168.10.0/24 [110/66] via 192.168.1.202, 00:00:04, Serial0/0
C      198.18.0.0/24 is directly connected, FastEthernet0/0
```

- What action will the router take when it receives a packet destined for 192.168.10.25?
- Assuming that the next router has no routes configured, describe the action that will be taken by the next router in relation to this packet

Murdoch UNIVERSITY

Review Question

Examine the routing table below

```
Gateway of last resort is not set
    10.0.0.0/27 is subnetted, 1 subnets
S      10.1.10.64 [1/0] via 192.168.1.206
    172.16.0.0/25 is subnetted, 1 subnets
O      172.16.5.128 [110/65] via 192.168.1.202, 00:01:10, Serial0/0
    192.168.1.0/30 is subnetted, 5 subnets
C      192.168.1.192 is directly connected, Serial0/1
O      192.168.1.196 [110/128] via 192.168.1.193, 00:01:10, Serial0/1
                     [110/128] via 192.168.1.202, 00:01:10, Serial0/0
C      192.168.1.200 is directly connected, Serial0/0/0
C      192.168.1.204 is directly connected, FastEthernet0/1
O      192.168.1.208 [110/65] via 192.168.1.202, 00:01:00, Serial0/0
O      192.168.10.0/24 [110/66] via 192.168.1.202, 00:00:04, Serial0/0
C      198.18.0.0/24 is directly connected, FastEthernet0/0
```

- What action will the router take when it receives a packet destined for 192.168.10.25?

Destination will match the route for 192.168.10.0/24 and be forwarded via Serial0/0

Murdoch UNIVERSITY

Review Question

Examine the routing table below

```
Gateway of last resort is not set
     10.0.0.0/27 is subnetted, 1 subnets
S      10.1.10.64 [1/0] via 192.168.1.206
     172.16.0.0/25 is subnetted, 1 subnets
[Output omitted]
```

- Assuming that the next router has no routes configured, describe the action that will be taken by the next router in relation to this packet (destined for 192.168.10.25)

As the router has no routes configured, it will only have knowledge of directly connected networks. Therefore, one of three actions is possible:

1. The router is directly connected to the 192.168.10.0/24 network and will forward the packet
2. The router is not directly connected to the 192.168.10.0/24 and the packet will be dropped

Murdoch University

Murdoch
UNIVERSITY

OSPF, Exterior Routing and the Internet

ICT169

Foundations of Data Communications

Admin

- Mid-Semester Test results released
  - Review your marks and feedback ASAP
  - Ask questions if you don’t understand (of me or your tutor)
- Practical Exam
  - Worth 25% of your final grade
  - Final lab session (semester week 14) in your regular lab
  - Requirements listed on LMS (Topic 12)
  - Start preparing now!
- PASS is still running
  - Sessions on Monday and Thursday

Murdoch UNIVERSITY

Last Week

- We took a closer look at the packet forwarding and routing process
- Also looked at two different methods for installing IP routes: static and dynamic routing
    - Distance vector and link-state routing

Lecture Overview

- A closer look at Open Shortest Path First (OSPF)
- Border Gateway Protocol and Routing over the Internet
- Content Distribution Networks and Net Neutrality

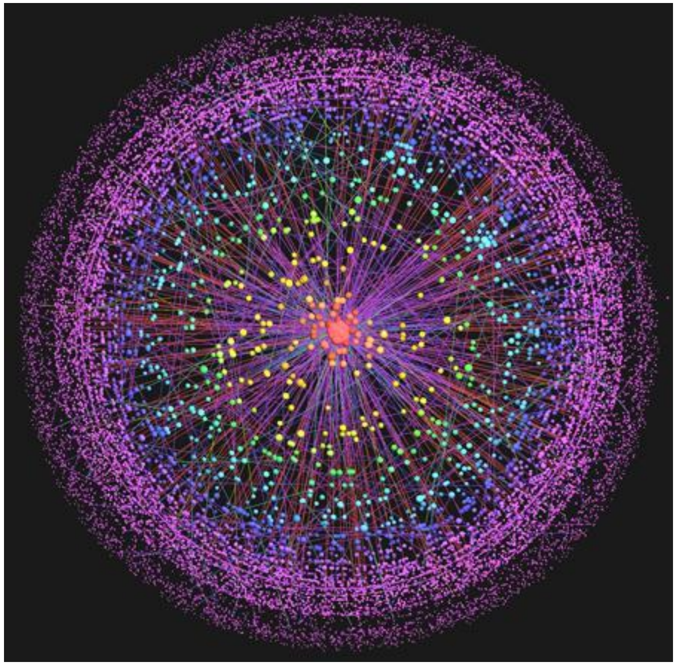

Packet Forwarding Revisited

- When a device is readying a packet for submission, it first checks whether the destination is in its local subnet
- If the destination is in the local subnet, the devices will use Layer 2 addressing
- Otherwise, the packet will be forwarded to the default gateway (its local router)

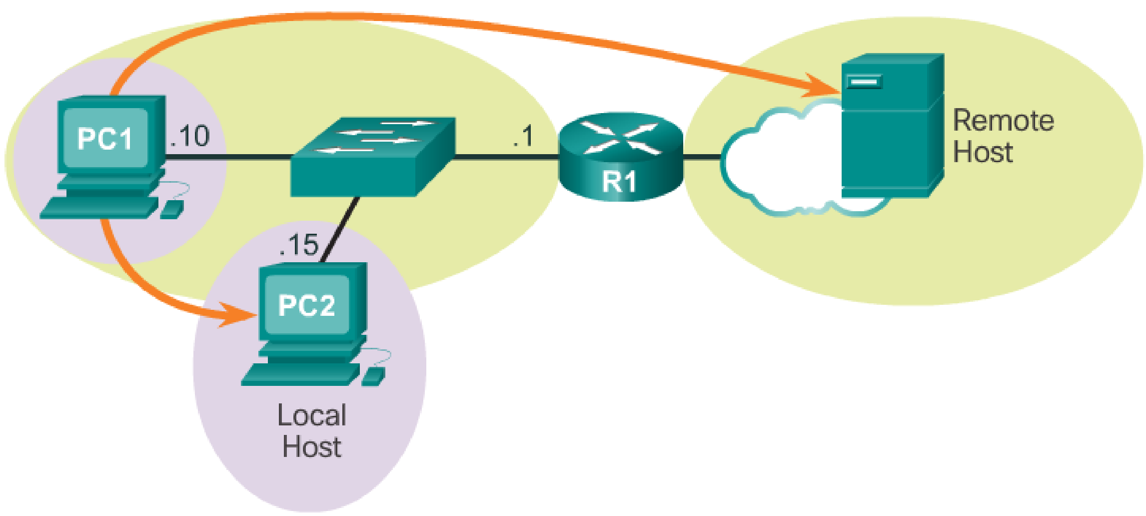

Murdoch UNIVERSITY

Routers and the Routing Process

- Routers are used to forward packets from one IP network to another
- Routers interfaces must belong to different IP networks
- Decision to forward the packet made based on the destination IP address

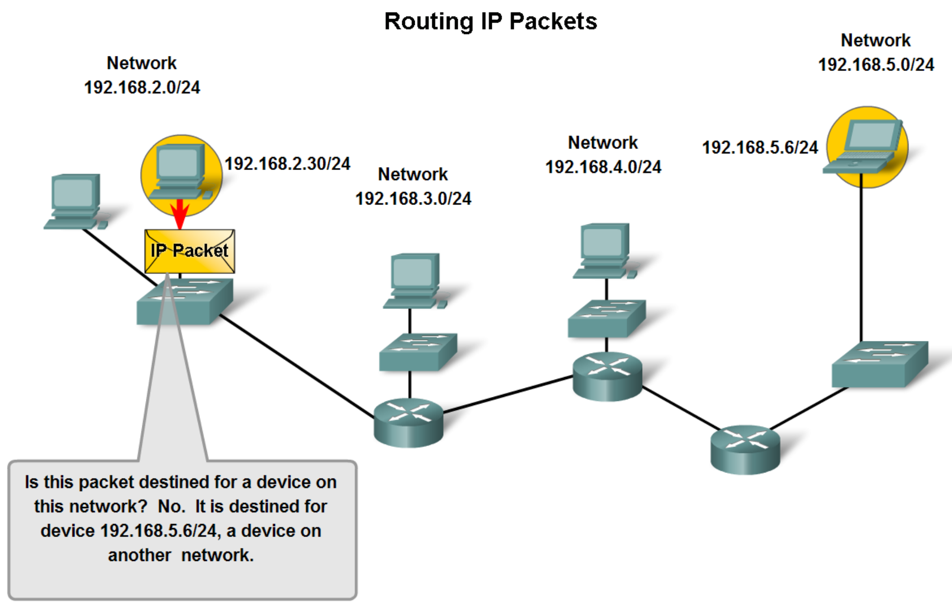

Routing Principles

Three key principles to routing:

1. Every router makes its decision alone, based on the information it has in its own routing table.
2. The fact that one router has certain information in its routing table does not mean that other routers have the same information.
3. Routing information about a path from one network to another does not provide routing information about the reverse, or return, path.

Murdoch UNIVERSITY

From Zinin, A. (2002). *Cisco IP routing: packet forwarding and intra-domain routing protocols*. Boston: Addison-Wesley.

Establishing Routes

- Routers are able to acquire knowledge of remote networks using two techniques:
  - A network engineer may manually enter a **static route**
  - **Routing protocols** may be used to propagate information between routers

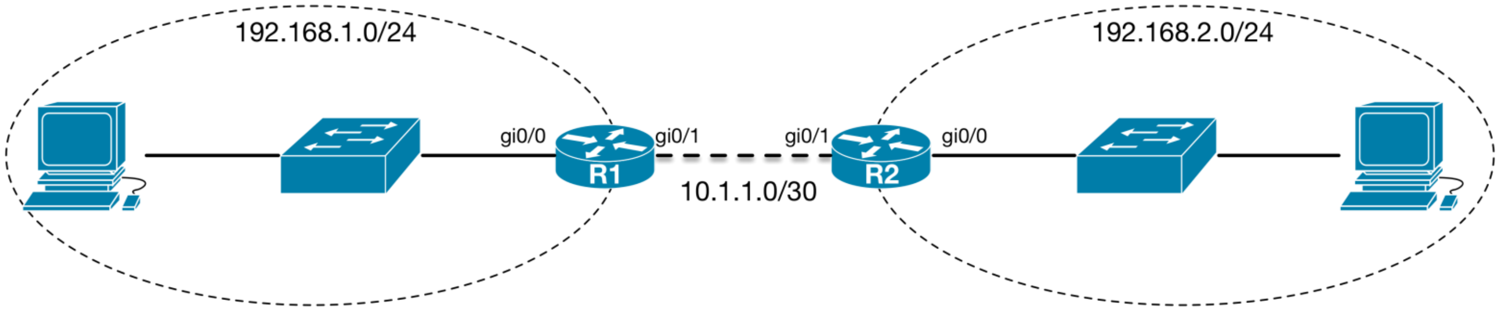

Dynamic Routing

- More commonly, dynamic routing protocols are used to propagate routing information throughout the network
    - Advertise networks the router knows how to reach, rather than specify networks you want to get to
- Determines best path to a given destination without human intervention
- Routers **automatically update routing table** when topology changes

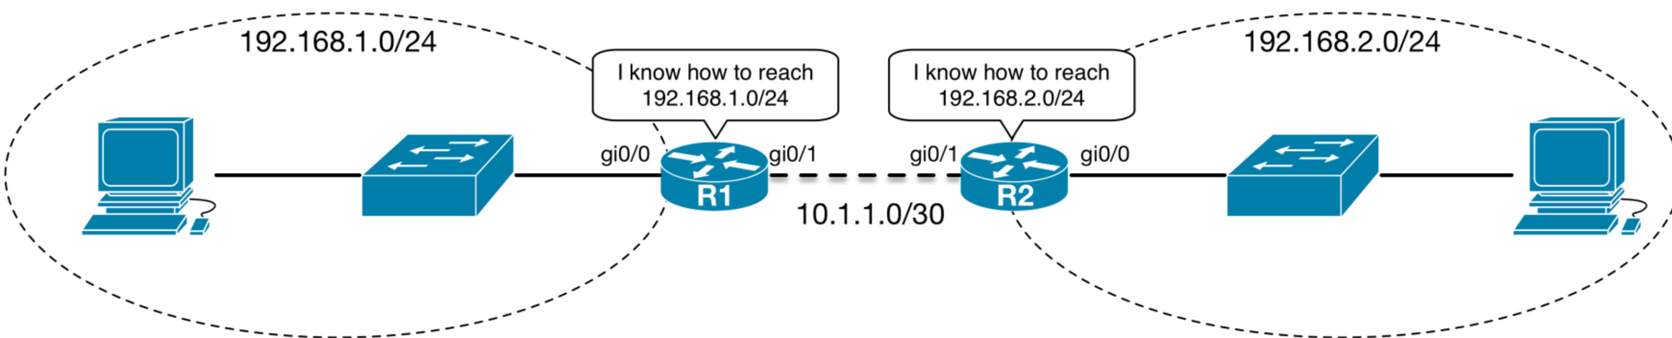

Murdoch UNIVERSITY

Dynamic Routing Protocols

- Routing protocols can be divided into two categories
- **Interior Gateway Protocols (IGPs)** are used to route within an administrative domain (eg. a single organisation)
  - Some examples include: Router Information Protocol (RIP), Enhanced Interior Gateway Routing Protocol (EIGRP), and Open Shortest Path First (OSPF)
- **Exterior Gateway Protocols (EGPs)** are used for routing packets between autonomous systems (eg. between organisations)
  - Only one EGP in common use: Border Gateway Protocol (BGP)
  - Used to route Internet traffic

Open Shortest Path First (OSPF)

- Standardised link-state routing protocol used in enterprise networks
  - Described in RFC2328 (ratified in 1998)
- Developed to address limitations of RIP
  - Inability to scale beyond 15 routers
  - Unable to consider link speed in routing decisions
- Implements Dijkstra's algorithm
- Two versions of OSPF exist:
  - OSPF version 2 for IPv4
  - OSPF version 3 for IPv6 (RFC2740)

Open Shortest Path First (cont.)

- Implements the key properties of link-state routing protocols:
    - Routers construct individual shortest path trees
    - Updates are triggered when network topology changes
        - Less bandwidth use (no periodic updates)
        - Faster notification of changes to network
- Uses cost as a routing metric to make more intelligent decisions about routing
    - Link cost takes speed into account
- Allows for routers to be grouped into multiple areas for scalability

OSPF Operation

So how does this all work?

1. Routers form neighbour adjacencies with neighbouring routers
2. Flood **Link State Advertisements (LSAs)** to all neighbouring routers
3. Add information from received LSAs to **Link-State Database (LSDB)**
4. Compute best routes based on Dijkstra’s algorithm
5. Install best routes into the **routing table**

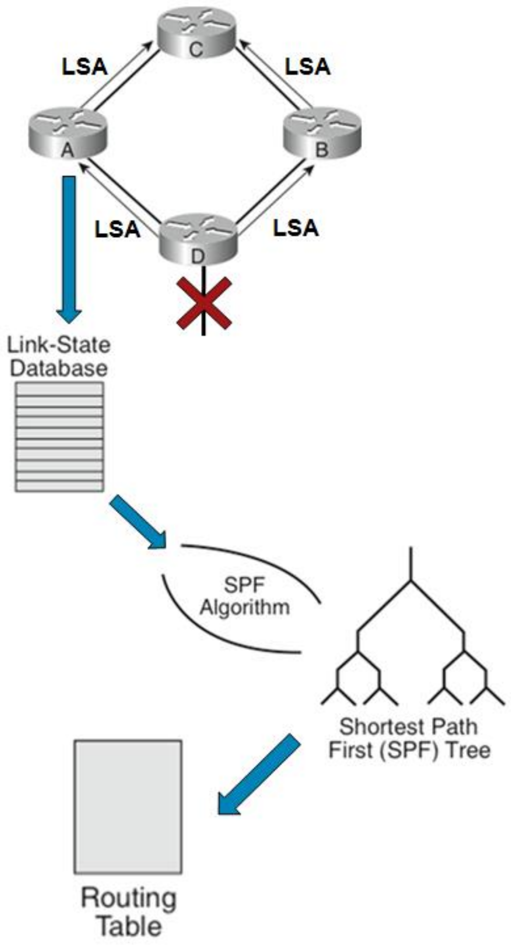

OSPF Neighbour Adjacencies

- Routers with OSPF configured transmit Hello packets on all interfaces participating in OSPF
- When a router running OSPF receives a Hello on an OSPF-enabled interface it forms a **neighbour adjacency**
- Routing updates are only sent to links where a neighbour relationship exists
- If no Hellos are received for some time, the link is assumed to have failed and the adjacency is broken
    - Triggers update to notify other routers of the change

```
R1#show ip ospf neighbors

Neighbor ID      Pri    State              Dead Time    Address          Interface
192.168.10.14      1    2WAY/DROTHER       00:00:36     192.168.10.10
GigabitEthernet0/1
192.168.10.13      0    FULL/  -           00:00:35     192.168.10.6     Serial0/0/0
```

OSPF Network Statements

- Network statements are used to specify which interfaces will participate in the dynamic routing process
- Interfaces participating in the routing process will be advertised to other routers
- Network statements in OSPF take up the form:

```
network 172.16.1.1 0.0.0.0 area 0
```

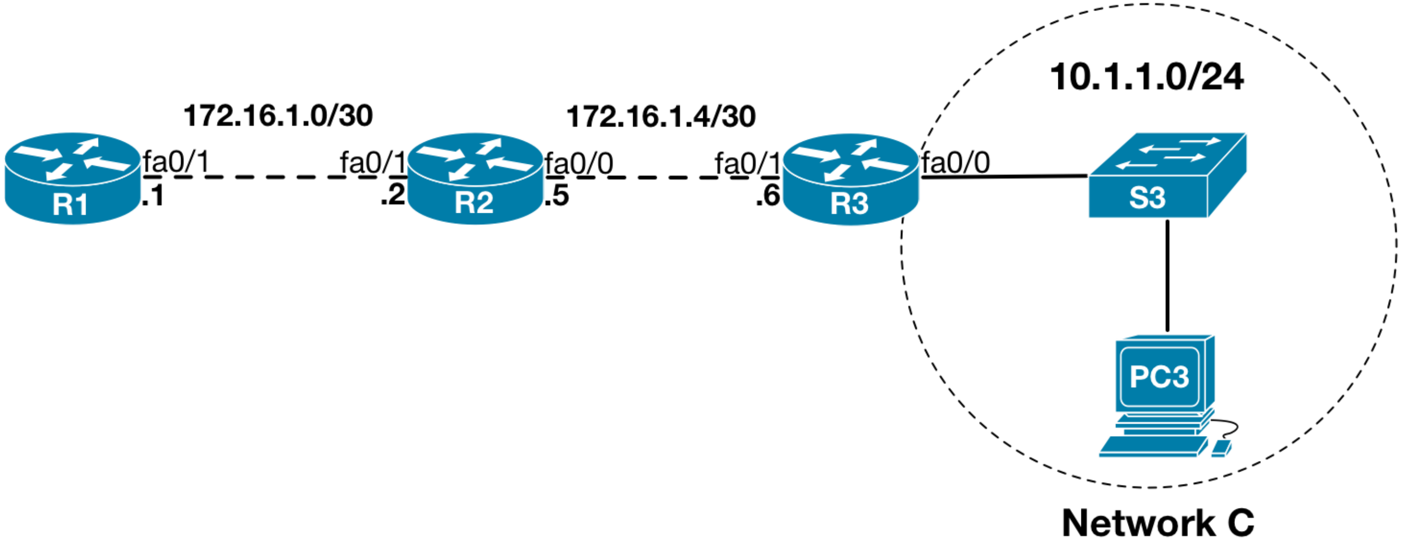

Link State Advertisements

- Recall that Link-State Advertisements (LSAs) are used to propagate information about the network topology
- Routers flood Link State Advertisements (LSAs) to all neighbouring routers
- Routers add information from received LSAs to their Link-State Database

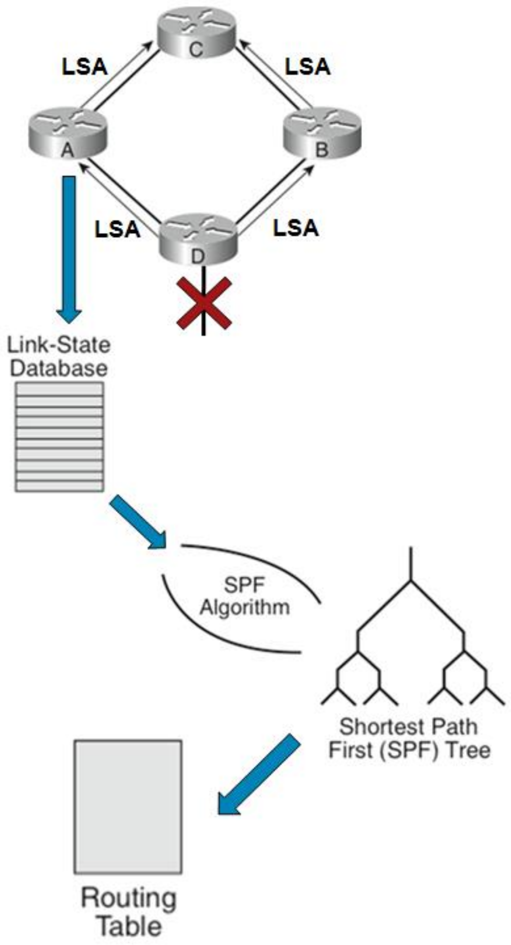

Murdoch UNIVERSITY

The Link-State Database (LSDB)

- A representation of the network topology as known by OSPF
  - Sometimes referred to as the topology table
- Contains a list of all LSAs that the router has received
- Used to calculate the best routes to each destination network using SPF algorithm
- Best routes installed into routing table

```
R1# show ip ospf database
            OSPF Router with ID (192.168.10.9) (Process ID 1)

                          Router Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum Link count
192.168.10.13   192.168.10.13   30          0x80000005 0x00b9b0 5
192.168.10.9    192.168.10.9    29          0x80000005 0x00473c 4
192.168.10.14   192.168.10.14   30          0x80000005 0x00ac0e 4
<Output omitted>
```

Murdoch UNIVERSITY

OSPF Cost Metric

- OSPF calculates the best path based on cost
- Cost of using a link is inversely proportional to its speed (bandwidth available)
  - Higher bandwidth (speed) means lower cost
  - Lower bandwidth means higher cost
  - Can manually alter the bandwidth as needed

OSPF Cost Metric (cont.)

- Calculation of cost is based on the formula:

Cost = 100,000,000 / Bandwidth (bps)

  - 10Mbps Ethernet would have a cost of:

    100,000,000 / 10,000,000 = 10

  - 1.5Mbps ADSL would have a cost of:

    100,000,000 / 1,544,000 = 64

- By default OSPF, can't distinguish between 100Mbps (cost = 1) and higher
- Can be fixed by altering the reference bandwidth to a higher value

OSPF Areas

- Even though OSPF is more efficient than predecessors, routers would struggle to compute SPF trees as network becomes larger
- Network can be broken up into **areas** which are semi-independent routing domains
  - Most LSAs are contained within their area
  - Each area maintains a separate LSDB
- Each area is then summarised for the rest of the network

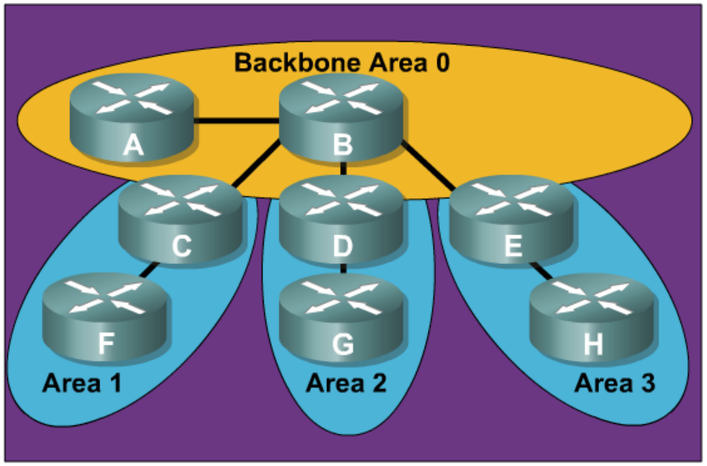

Murdoch UNIVERSITY

OSPF Area Types

- Every OSPF network must have a **backbone area** (referred to as Area 0)
  - Also known as the transit area
- All other areas must connect to the backbone area through an **Area Border Router (ABR)**
- ABRs summarise the LSAs for both areas they are connected to and propagate these summaries into the opposite area

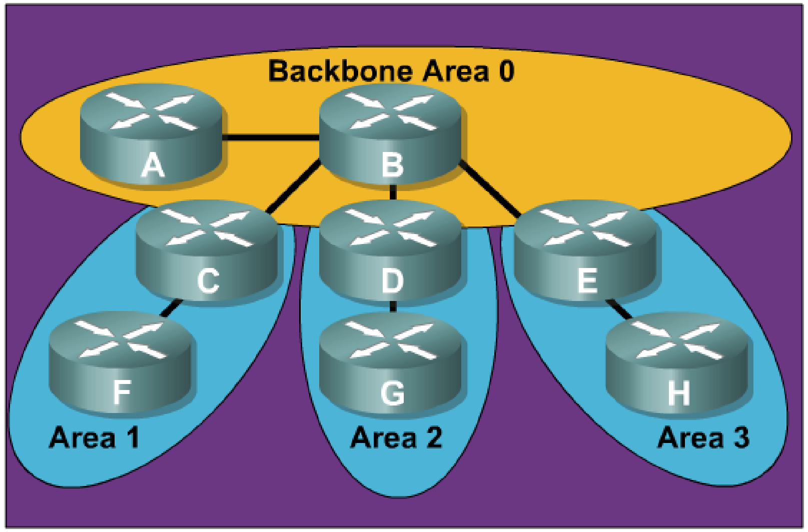

Murdoch University

OSPF Designated Routers

- In multi-access networks OSPF elects a **Designated Router (DR)** and a **Backup Designated Router (BDR)**
- All routers form adjacencies with the DR only
- LSAs are sent to, and propagated by, the DR
    - Reduces the number of adjacencies and amount of LSA flooding needed

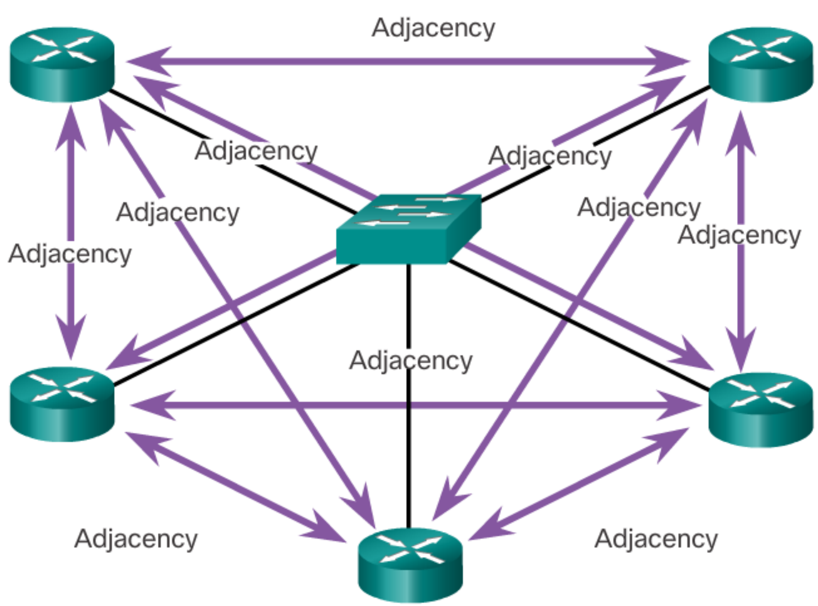

Murdoch UNIVERSITY

External Gateway Routing

- To date, we've focused on interior routing (routing within an organisation)
- **Exterior Gateway Protocols (EGPs)** are used for routing packets between autonomous systems
- Routers participating in exterior routing will usually be owned by different organisations
- Have knowledge of the whole Internet routing table
- Only one EGP used: Border Gateway Protocol (BGP)

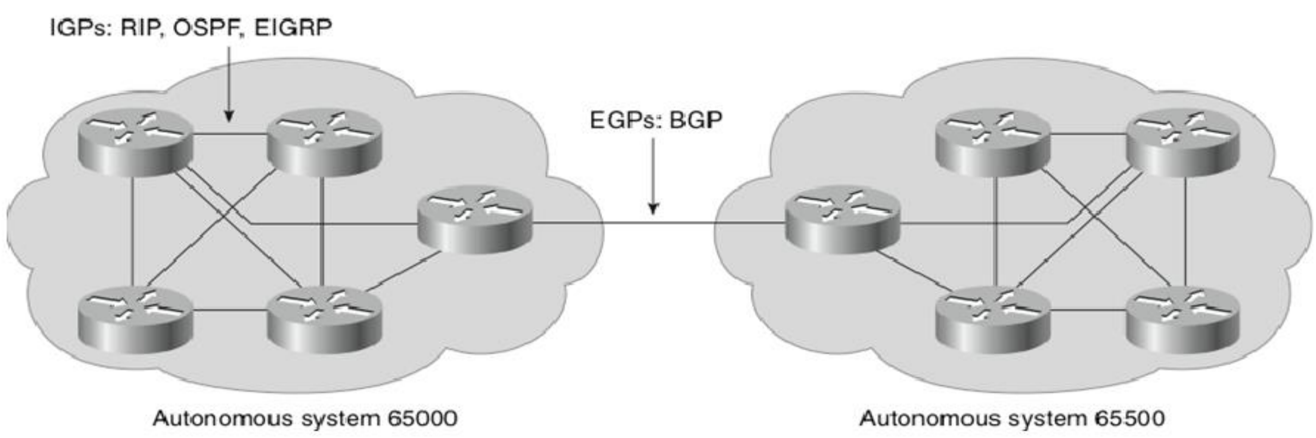

Autonomous Systems

- An Autonomous System (AS) is a group of routers operate within a single administrative domain (an organisation)
  - Routers within an AS should also share similar routing policies
  - Viewed as a single entity by the outside world
  - Will usually run one or more IGPs for routing internally
- Only large organisations would be considered to be their own AS
  - Examples: ISPs, large enterprise, universities, federal governments
  - Smaller companies may group together to form an AS
- Each AS is assigned a unique number
  - Assignments handled by IANA and the RIRs

Autonomous System Numbers

- Each AS is assigned a unique number
  - Assignments handled by IANA and the RIRs
- 32 bit number written as a single number
  - Can also be divided into two 16 bit groups written as x.y
- Private AS numbers which can be used internally also exist but must not be advertised externally
  - 64,512 – 65,535
  - 4,200,000,000 – 4,294,967,294
- Organisations that connect to a single ISP must use an private AS number

Internet Assigned Numbers Authority

Murdoch University

Border Gateway Protocol (BGP)

- The only exterior gateway protocol currently available and in use
- Designed to allow routing **between autonomous systems** and manage large numbers of routes
- Can also be used internally (usually only for very large networks)
- Path vector routing protocol
    - Similar to distance-vector routing algorithm
    - Includes a list of AS along the path to destination
- Only routing protocol to use TCP (most IGPs use IP)

Routing Between AS

- Like interior routing protocols, BGP is designed to form loop-free routes between AS (referred to as external BGP)
- Can also distribute information within an AS (using internal BGP)

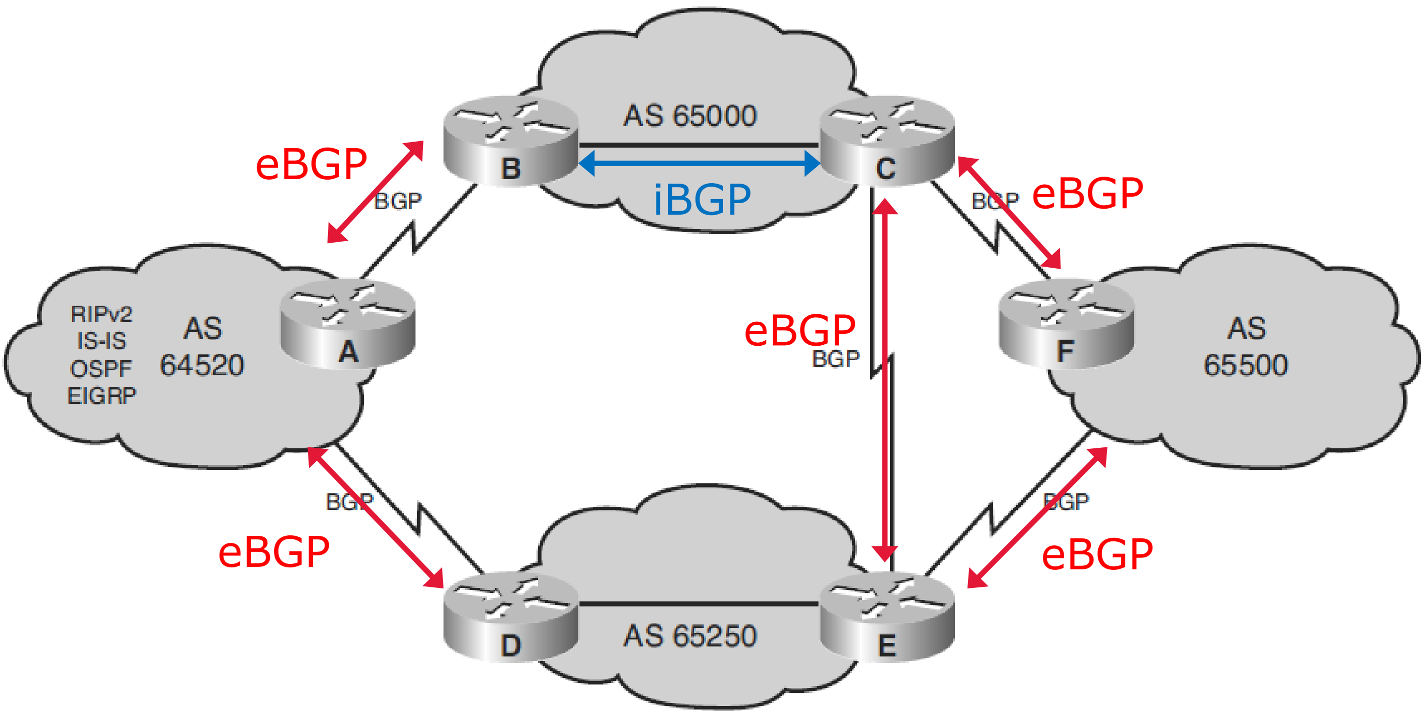

Policy-based Routing

- IGPs like OSPF seek to identify the best path (using metrics like cost)
- BGP allows for routing to be further manipulated by modifying a number of path attributes
  - Next-Hop
  - Weight
  - Local Preference
  - AS_PATH
- Routers exchange these attributes as part of BGP routing updates
- These path attributes are usually set based on organisational policies and agreements

Murdoch UNIVERSITY

Policy-based Routing Example

- Imagine that your organisation is AS 65500 and has two ISPs (ISP A and B)
- ISP A is the primary ISP, while ISP B should be used as a backup
- Set **local preference** for ISP A (AS 65000)

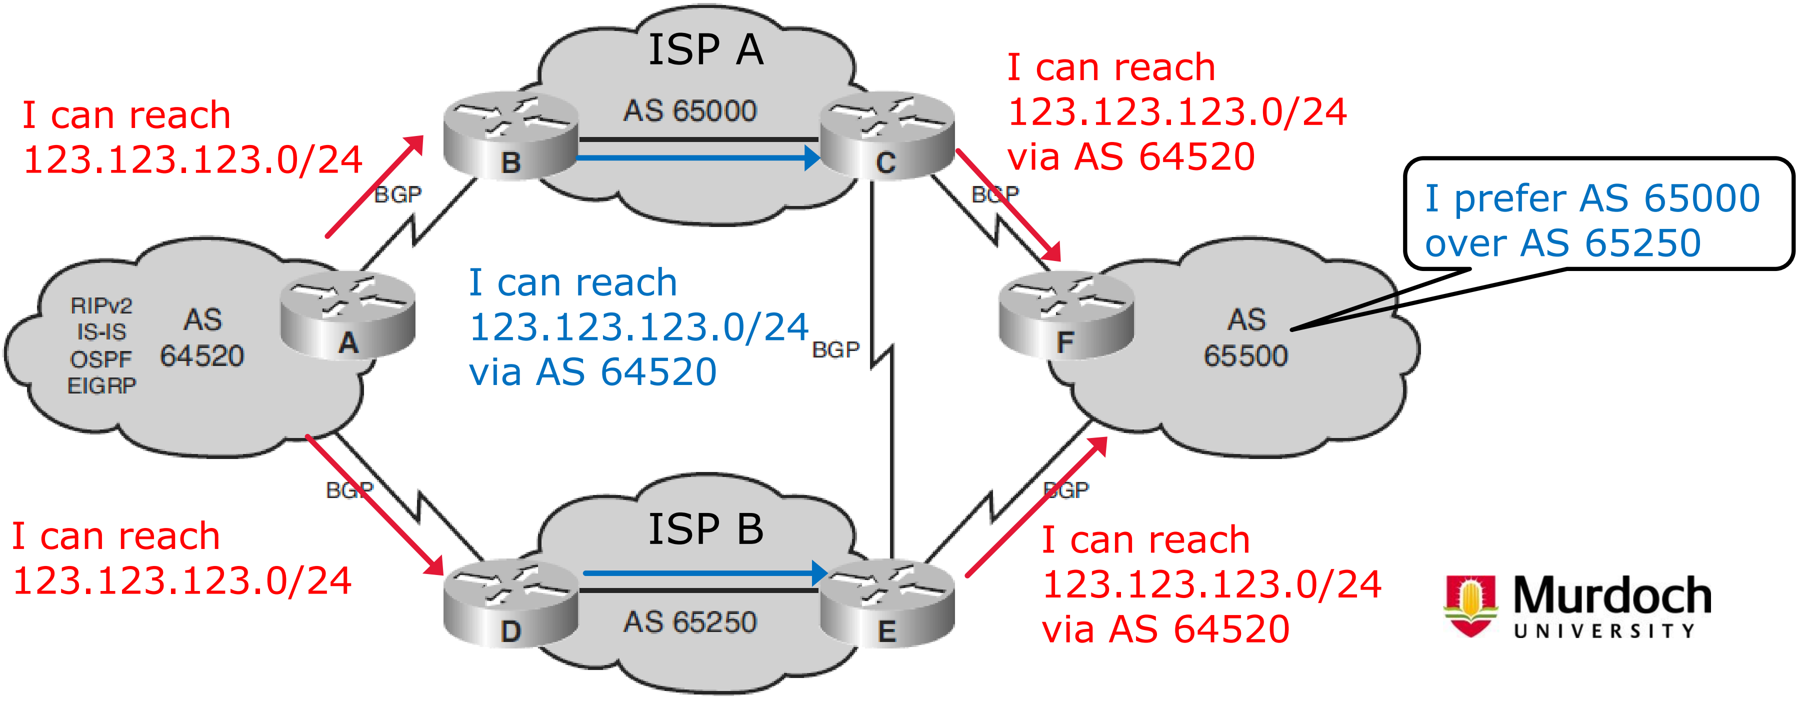

Murdoch University

Policy-based Routing (cont.)

- Policy-based routing can be manipulated to achieve many goals but can increase complexity of routing
  - Load balancing – prioritise between different ISPs or links between branches
  - Performance management – avoid congesting certain bottleneck links, or route around congested ones
- However, it can also be misused (deliberately or inadvertently)
  - In 2010, China advertised ‘better’ paths for several networks causing traffic to be diverted through Chinese routers
  - In 2012, Dodo advertised ‘better’ paths than its upstream provider, causing traffic to be routed through itself (and couldn’t cope with the load)

Break

When we return: The Internet

Structure of the Internet

- Recall that the Internet is global mesh of interconnected autonomous systems
- Hierarchical structure of service providers

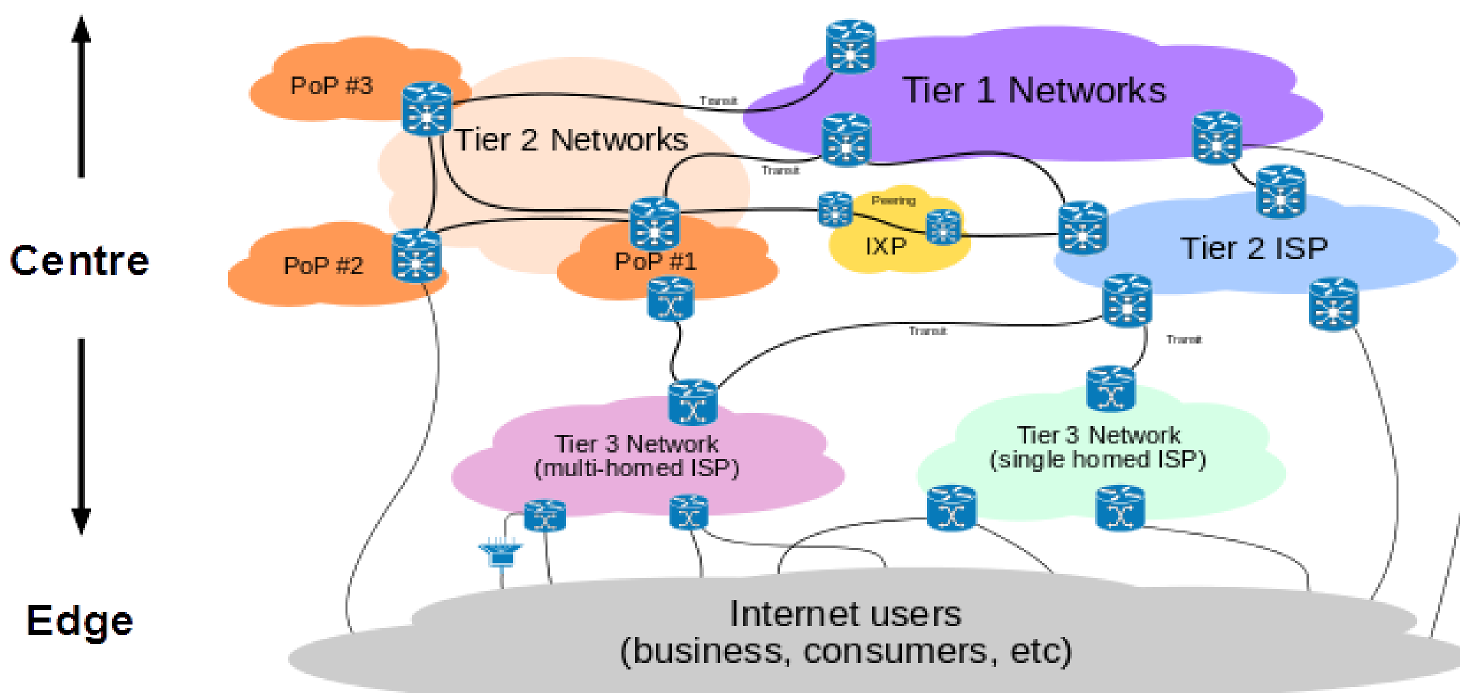

https://en.wikipedia.org/wiki/Tier_1_network

Service Provider Tiering

- Tier 1 – Internet backbone
    - Large network providers that connect to other Tier 1 providers
    - May be – but are not necessarily – a retail provider
    - Examples: CenturyLink, AT&T, Telstra*
- Tier 2 – National providers
    - Purchase capacity and routing services from Tier 1 providers
    - Peer with other Tier 2 providers (and some Tier 1)
    - Operate on a national level with Points of Presence (POPs)
    - Examples: Comcast, AAPT

Murdoch UNIVERSITY

Service Provider Tiering (cont.)

- Tier 3 – Local providers
    - Purchase capacity and routing services from higher tier providers
    - Operate in a small geographical region
    - Examples: Dodo, Mate Communicate, Verizon

Murdoch UNIVERSITY

Relationships Between Service Providers

- Peering
  - Exchange of traffic between two ISPs (usually of the same tier) without costs involved
- Transit
  - Pay another ISP (usually of a higher tier) to carry your traffic to another network

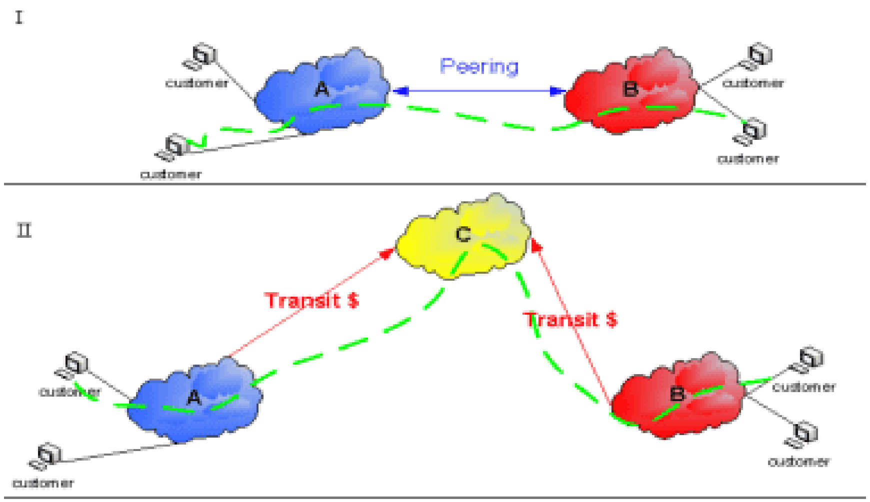

Murdoch UNIVERSITY

Peering Agreements

- Peering is usually the preferred option for ISPs as it allows them to extend reach without incurring costs
  - Additional peers will add redundancy to the network
  - Can also improve performance (alternate routes out of the network)
- Usually mutually beneficial (equal benefit to both service providers)
  - Otherwise the side that doesn't benefit would probably favour a Transit arrangement
- Connection between service providers occurs at an Internet Exchange Point

Murdoch UNIVERSITY

Internet Exchange Points (IXPs)

- Connection point between two or more ISPs
- Usually housed in colocation facilities
- Can be non-profit, government run, or private enterprise

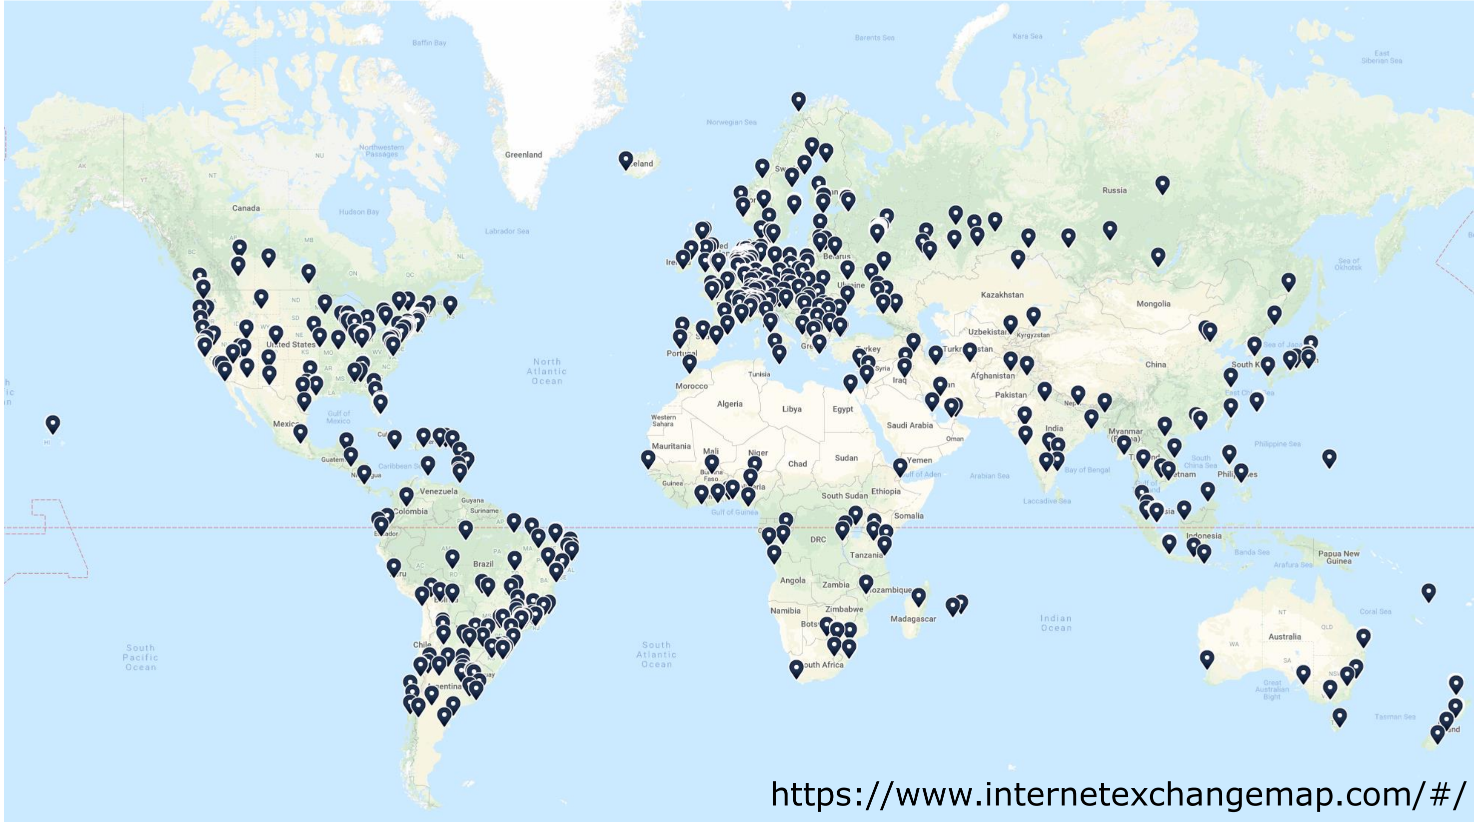

https://www.internetexchangemap.com/#/

Murdoch University

Point of Presence (POP)

- Demarcation point at which ISPs deploy equipment to connect customers (either end-users or transit ISPs) to their network
- Facility is often owned by a telecommunications carrier
- Larger ISPs often maintain a large network of POPs
  - Usually at least one per major city
  - Smaller ISPs may not have resources to build multiple POPs; performance is degraded

BGP and The Internet

- Routing over the Internet is managed by BGP
    - Large numbers of routes would make IGPs unsuitable
- Each ISP, government body, and (large) organisation is assigned an AS number
- Traffic will pass through numerous AS to get to its final destination

Autonomous Systems and Routing over the Internet

- Traffic to Google passes through at least two AS to reach its destination
  - AS7575 – AARNET (Murdoch University's upstream provider)
  - AS15169 – Google
- Note that Murdoch doesn't have its own AS number

```
traceroute to google.com (216.58.199.46), 30 hops max, 60 byte packets
 1  southcore1.murdoch.edu.au (134.115.64.250) [AS7575]
 2  134.115.1.21 (134.115.1.21) [AS7575]
 3  vl219.xe-5-1-3.pe1.knsg.wa.aarnet.net.au (138.44.176.98) [AS7575]
 4  et-1-3-0.pe1.prka.sa.aarnet.net.au (113.197.15.44) [AS7575]
 5  et-0-1-0.pe1.adel.sa.aarnet.net.au (113.197.15.36) [AS7575]
 6  et-7-3-0.pe1.wmlb.vic.aarnet.net.au (113.197.15.28) [AS7575]
 7  ae2.vic-wmlb-bdr1.aarnet.net.au (113.197.15.139) [AS7575]
 8  as15169-gw1.aarnet.net.au (202.158.210.41) [AS7575]
 9  * * *
10  209.85.255.174 (209.85.255.174) [AS15169]
11  209.85.243.145 (209.85.243.145) [AS15169]
12  syd09s12-in-f14.1e100.net (216.58.199.46) [AS15169]
```

Murdoch UNIVERSITY

Policy-based Routing and Routing over the Internet

- Policy-based routing is used to enact business rules
  - Some peering agreements may be more favourable than others
  - Avoid using transit arrangements; they cost money
- Consider the Local ISP with two routes to Netflix

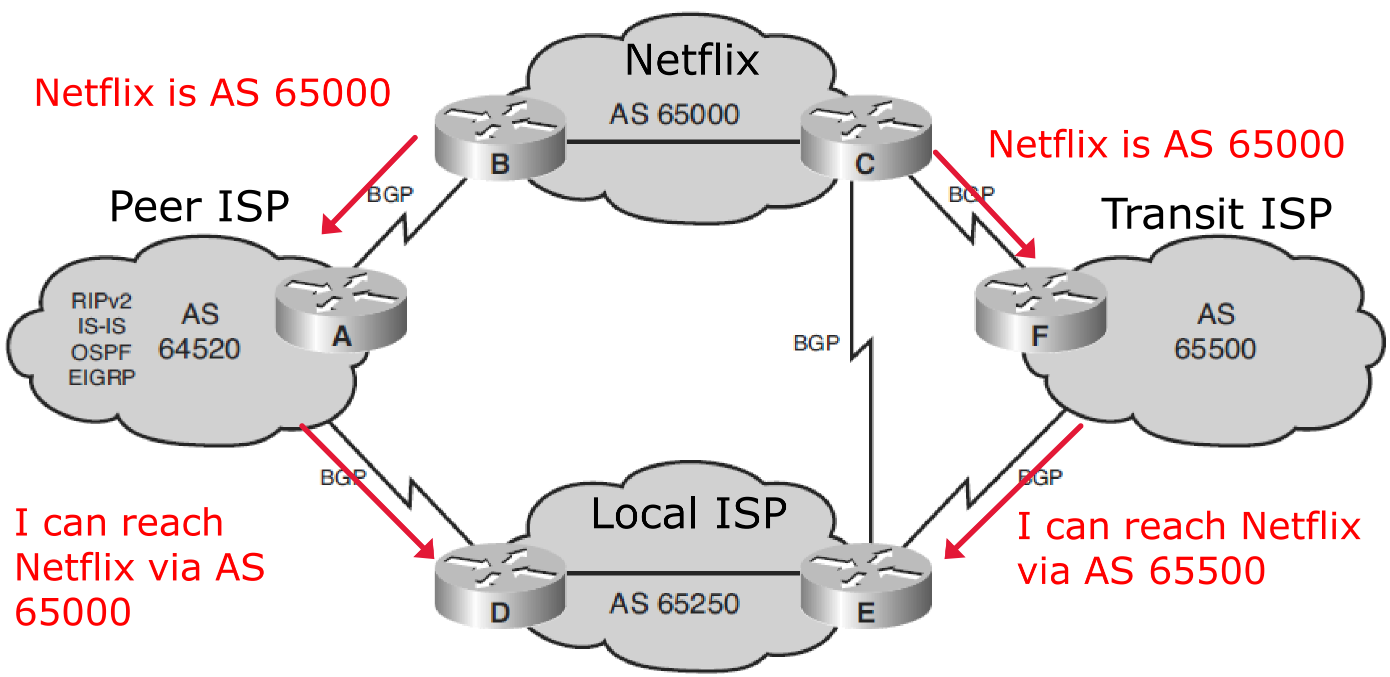

Policy-based Routing and Routing over the Internet (cont.)

- Consider the Local ISP with two routes to Netflix
- The Peer ISP allows traffic to pass through their network without cost
- The Transit ISP will charge for handling traffic
- How would you route traffic?

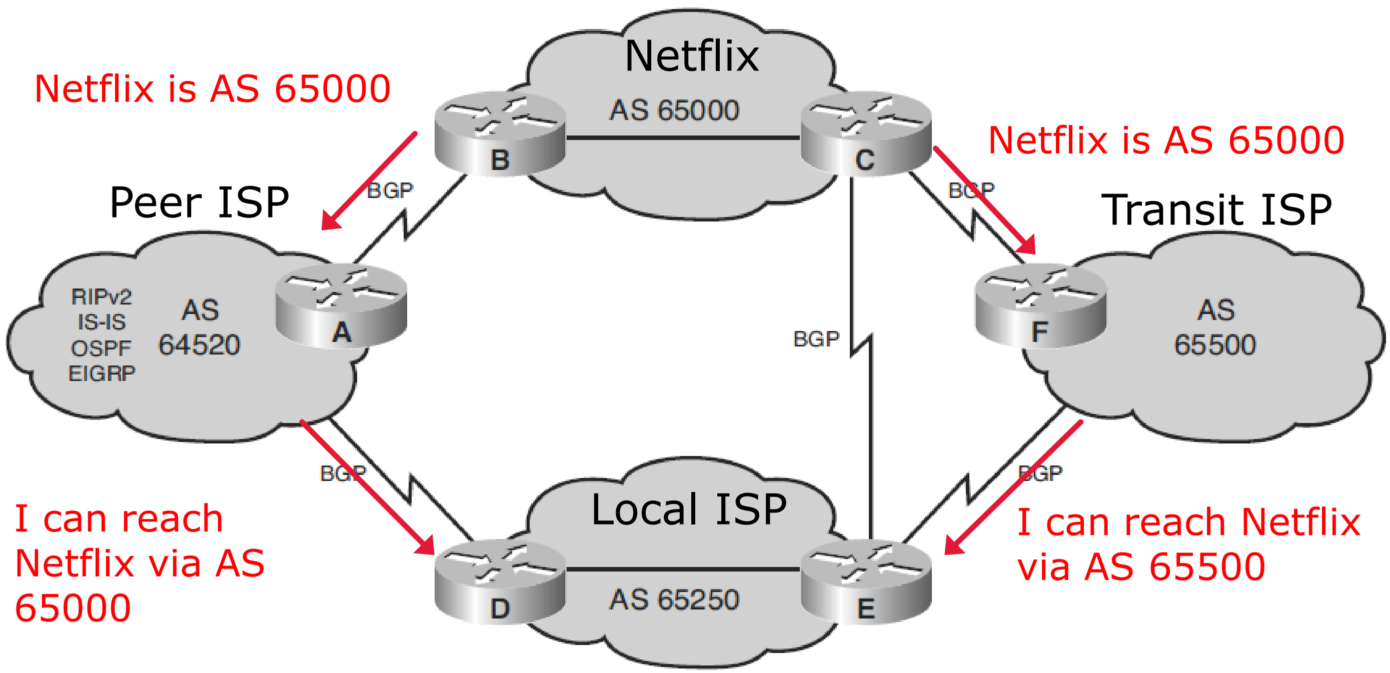

Murdoch University

Content Delivery over the Internet

- High speed Internet has given rise to a range of content providers (eg. streaming video and music)
- Delivering large amounts of data across long distances – such as between continents – is undesirable
    - High latency – packets are transmitted long distances
    - Added transit costs – multiple peering / transit arrangements needed
- Large content providers now typically use **Content Distribution Networks (CDNs)**

Murdoch UNIVERSITY

Content Distribution Networks (CDNs)

- Also known as Content Delivery Networks
- Geographically distributed networks of servers specifically designed to deliver content to end-users
- Each server in the network holds a copy of hosted content to ensure that users are served by a nearby server
    - Minimises delay to end-user
    - Reduces transit costs
    - Reduces load on original content server (and provides some redundancy)
- Some CDNs are run by content providers (eg. Google, Netflix) while others are commercial services (eg. Akamai, Cachefly, Cloudflare)

CDNs via HTTP Redirect

- User sends HTTP GET to the main content server
- Server provides a redirect to the nearest CDN cache server
- Pros
  - Fine-grained control
  - Server selection is based on client IP or location information
- Cons
  - Main content server is still involved
  - Additional TCP connection required

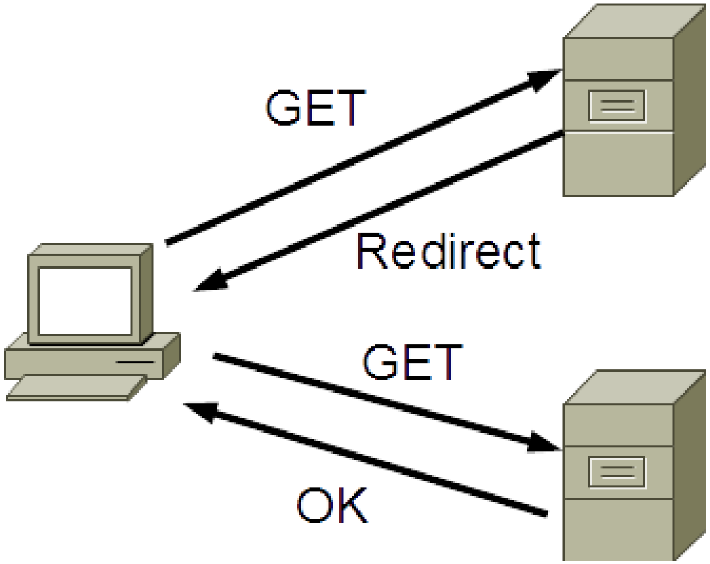

Murdoch UNIVERSITY

DNS-based Server Selection

- User requests cdn.site.com
- DNS request will reach authoritative DNS server of the CDN provider
- Authoritative DNS server will reply with IP address of CDN cache server closest to the user
- Pros
  - No added overheads
  - Response cached after first request
- Cons
  - Based on IP address of ISP DNS
  - DNS caching increases update time
- Akamai Technologies uses this approach

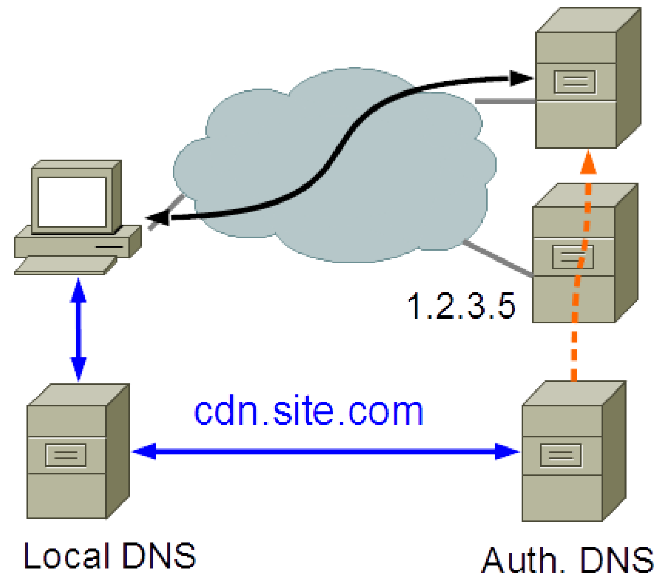

Anycast Routing Server Selection

- Packets will be routed to the closest server based on BGP routing information
- Pros
  - No dependency on DNS or TCP
- Cons
  - Complex implementation (requires multiple peers and consistent transit providers)
- Cloudflare uses this approach

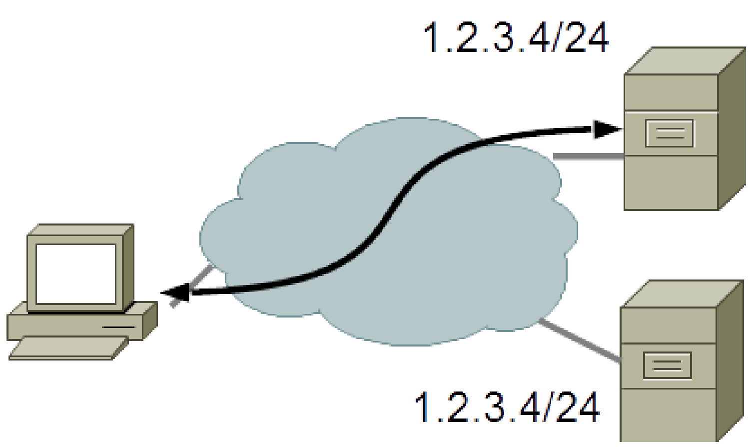

Murdoch UNIVERSITY

CDNs in Content Delivery

- CDNs are now in very common use on the Internet
  - Facebook, Netflix, Google, Instagram, Twitter, iTunes, and Spotify (amongst many others) all use CDNs
- Some ISPs have opted to host CDN nodes as a differentiator
  - Improves performance for their customers
  - Reduces bandwidth costs (content is inside their network)
  - Allows opportunity for zero rating (use of service without counting against data caps)
- This leads us to the debate of Net Neutrality

Net Neutrality

Describes “the idea that Internet service providers (ISPs) should treat all data that travels over their networks fairly, without improper discrimination in favor of particular apps, sites or services” – Electronic Frontier Foundation

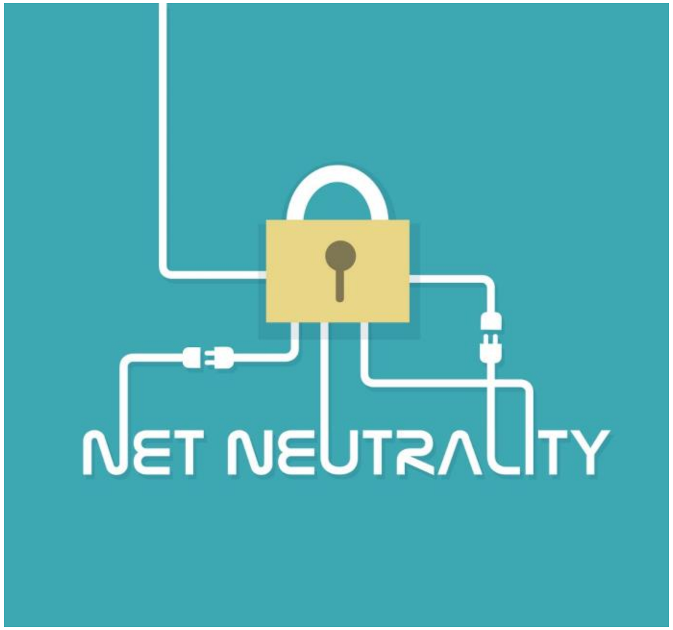

https://blog.prototypr.io/net-neutrality-and-design-f622a7aed604

Murdoch UNIVERSITY

Net Neutrality (cont.)

- Differential treatment of traffic could lead to:
  - Peer-to-Peer traffic being blocked
  - Throttling (slowing) of encrypted traffic
  - Content providers paying more for bandwidth (which will most likely be passed on to end-users)
  - ISPs subsiding preferred services through zero-rating

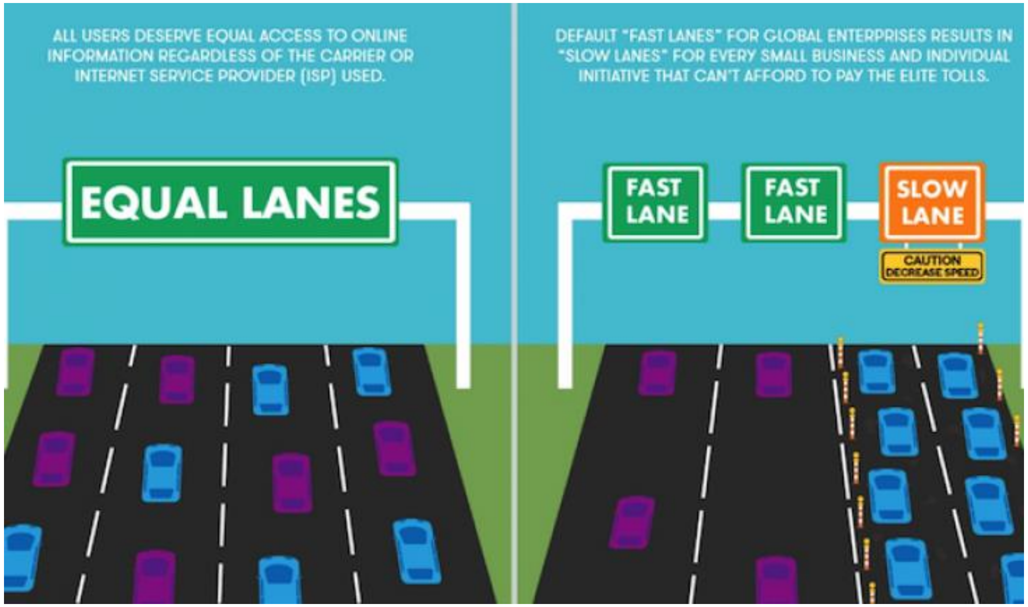

https://www.nextbigfuture.com/2017/11/net-was-already-unfair-but-how-will-it-become-more-unfair-without-net-neutrality.html

Murdoch UNIVERSITY

Net Neutrality Debate

- Content providers like Amazon, Netflix, and Google have a vested interest in ensuring Net Neutrality is enacted
  - Network access costs are currently externalised
  - Would increase (possibly significantly) if Net Neutrality isn't enforced
- Many ISPs are against Net Neutrality
  - Impediment to offering additional services and perks to differentiate themselves
  - Fearful of becoming a utility for data delivery

Murdoch UNIVERSITY

Net Neutrality Debate (cont.)

- Users are likely conflicted
  - Offers like zero-rating and fast lanes are enticing
  - But what happens if services they use isn't part of the offer (or they need to pay more)?
  - New services that don't get preferential treatment may be less appealing

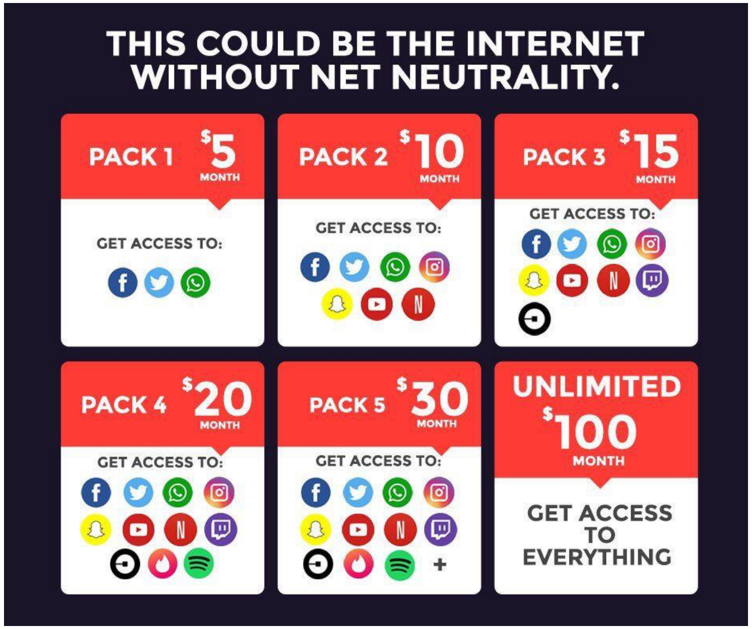

Murdoch UNIVERSITY

Net Neutrality – Current Status

- In 2015, the US Federal Communications Commission reclassified the Internet as a basic telecommunications service (enabling Net Neutrality)
    - Internet access treated as a utility
    - Regulation repealed in 2017 by new administration
- Australia has no specific Net Neutrality laws
    - Many ISPs have zero-rating (also known as freezones or unmetered data)
    - Australian Competition and Consumer Commission may enforce some neutrality

Murdoch UNIVERSITY

Lecture Objectives

You should now be able to:

- Describe the operation of Open Shortest Path First (OSPF)
- Describe the purpose of Neighbour Adjacencies in OSPF
- Describe the purpose of Link-State Advertisements
- Describe the role of the Link-State Database in computing routes
- Describe the metric used by OSPF in determining the best path
- Identify the role of the Designated and Backup Designated Routers
- Describe the purpose of areas in OSPF
- Describe the role of the backbone area in OSPF operation
- Differentiate between interior and exterior gateway routing
- Describe the operation of the Border Gateway Protocol
- Differentiate between the BGP and interior gateway routing protocols
- Differentiate between the three classes of Internet Service Providers in relation to the structure of the Internet
- Describe the role of Internet Exchange Points in routing over the Internet
- Describe the role of a Point of Presence in the structure of the Internet
- Describe the purpose of a Content Distribution Network
- Define the concept of Net Neutrality

Lecture Summary and the Week Ahead

- We've taken a closer look at link-state routing through OSPF, as well as a high level look at BGP, and how it facilitates routing over the Internet
- Examined the structure of the Internet, and some
- Readings for this week are Routing and Switching Essentials Chapter 8, as well as links available via LMS
- In the labs: configuring OSPF in Packet Tracer

Murdoch UNIVERSITY

Next Week

- All about wireless networks: Applications, different standards, security and more