Routing Basics

ISP/IXP Workshops

Routing Concepts

- **IPv4**

- **Routing**

- **Forwarding**

- **Some definitions**

- **Policy options**

- **Routing Protocols**

IPv4

- **Internet uses IPv4**

  addresses are 32 bits long

  range from 1.0.0.0 to 223.255.255.255

  0.0.0.0 to 0.255.255.255 and 224.0.0.0 to 255.255.255.255 have "special" uses

- **IPv4 address has a network portion and a host portion**

IPv4 address format

- **Address and subnet mask**

    written as

    12.34.56.78 255.255.255.0 *or*

    12.34.56.78/24

    mask represents the number of network bits in the 32 bit address

    the remaining bits are the host bits

What does a router do?

A day in a life of a router

find path

forward packet, forward packet, forward packet, forward packet...

find alternate path

forward packet, forward packet, forward packet, forward packet…

repeat until powered off

Routing versus Forwarding

- **Routing = building maps and giving directions**

- **Forwarding = moving packets between interfaces according to the "directions"**

IP Routing – finding the path

- **Path derived from information received from a routing protocol**

- **Several alternative paths may exist**

  best next hop stored in forwarding table

- **Decisions are updated periodically or as topology changes (event driven)**

- **Decisions are based on:**

  topology, policies and metrics (hop count, filtering, delay, bandwidth, etc.)
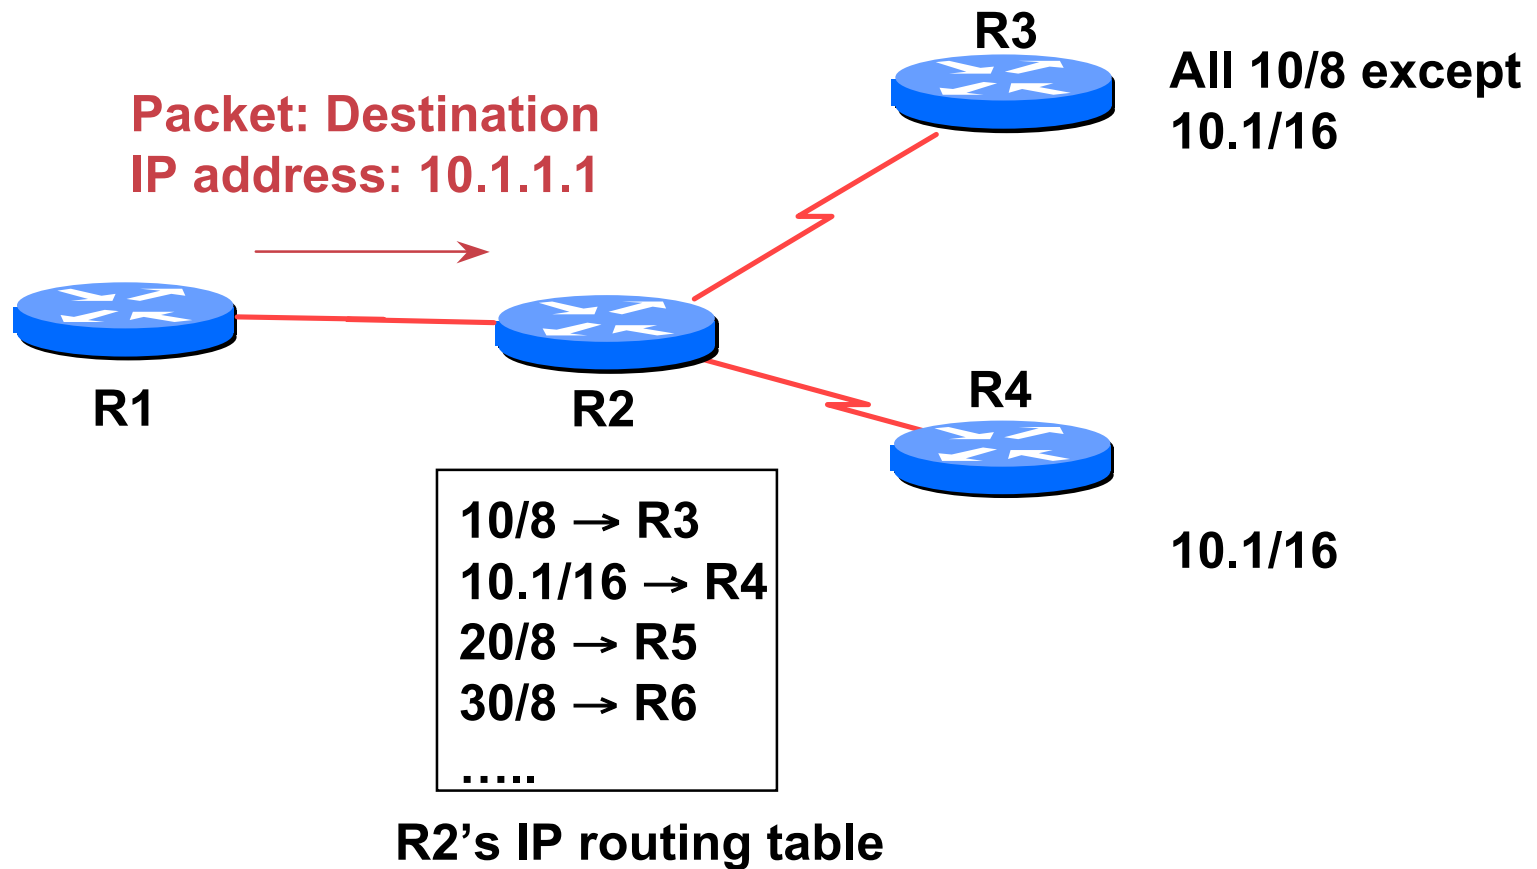
IP route lookup

- **Based on destination IP address**

- **"longest match" routing**

    more specific prefix preferred over less specific prefix

    example: packet with destination of 10.1.1.1/32 is sent to the router announcing 10.1/16 rather than the router announcing 10/8.
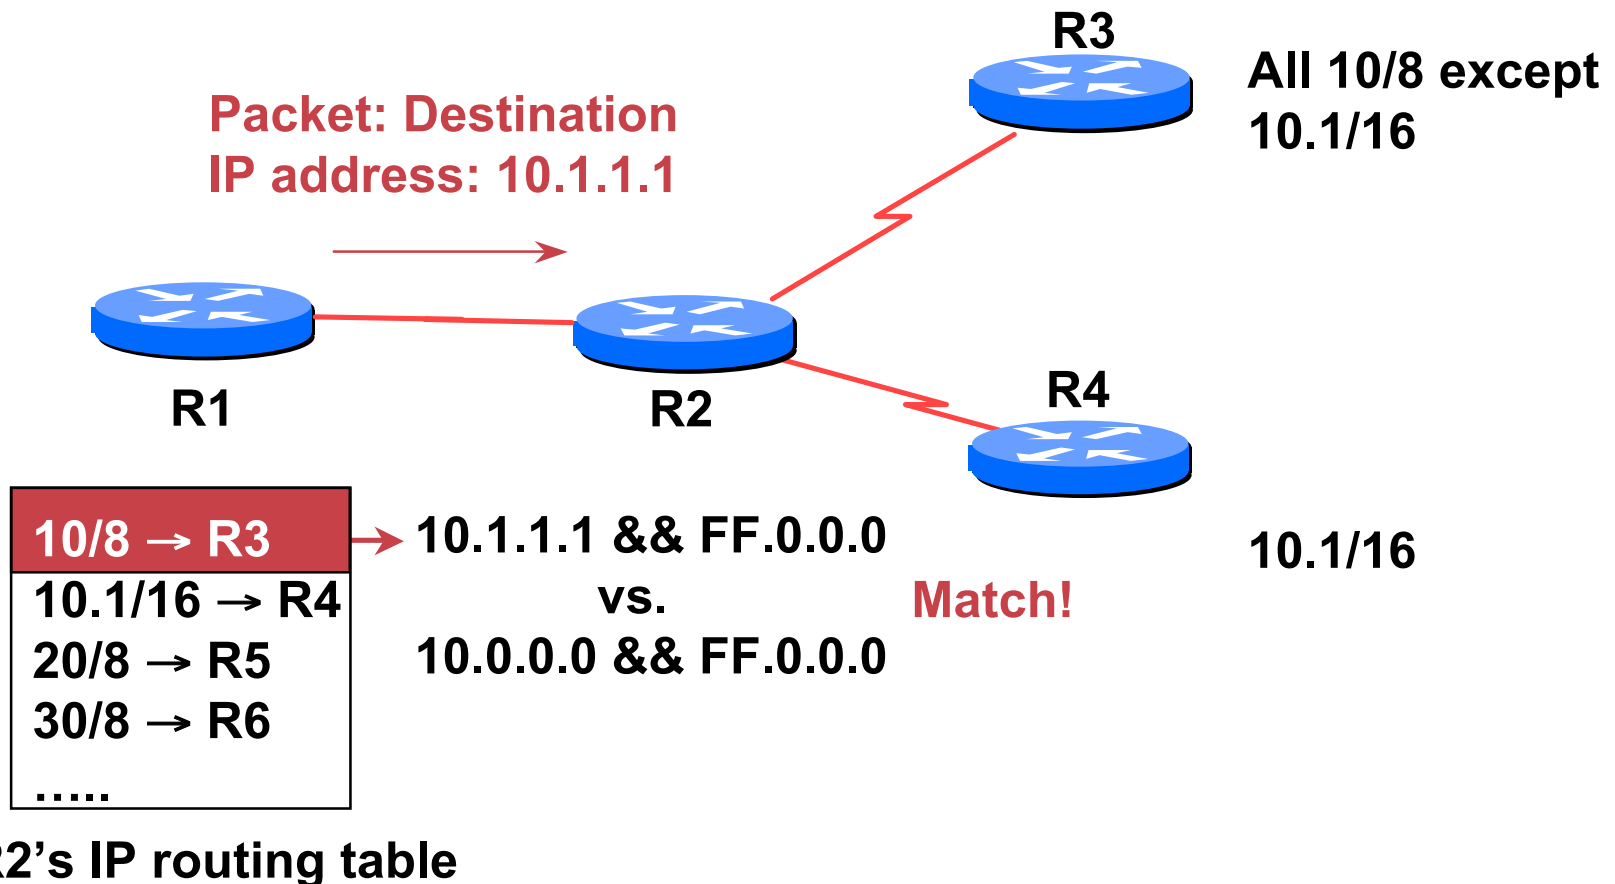
IP route lookup
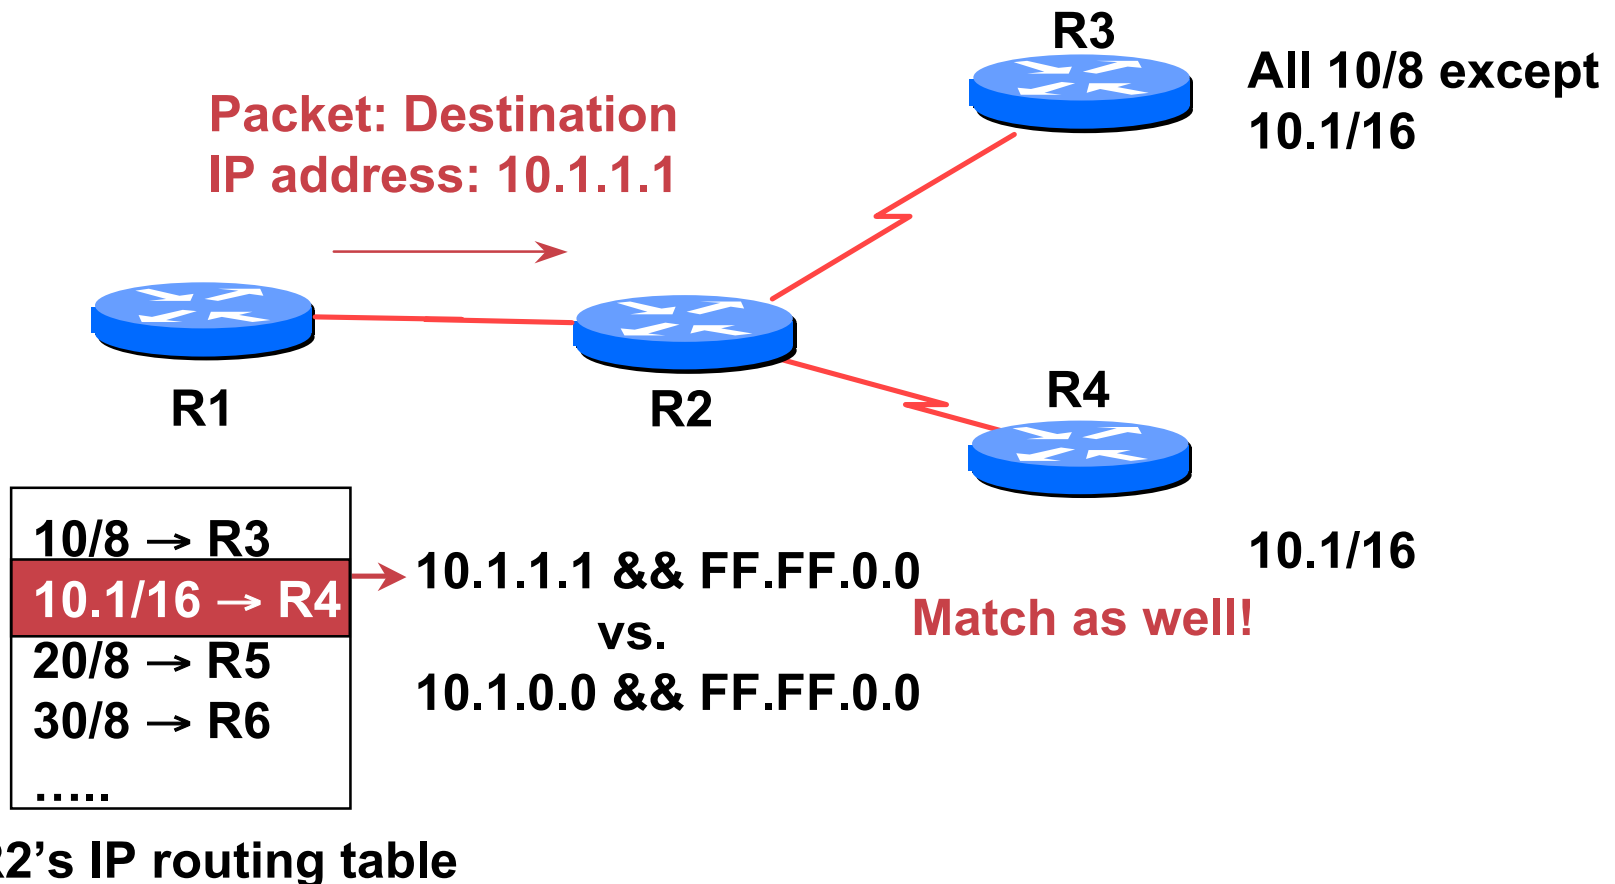
- **Based on destination IP address**

**Packet: Destination
IP address: 10.1.1.1**

R3
**All 10/8 except
10.1/16**

R1

R2

R4

10/8 → R3
10.1/16 → R4
20/8 → R5
30/8 → R6
…..

10.1/16

R2's IP routing table

IP route lookup:
Longest match routing

- ## Based on destination IP address

R3

All 10/8 except 10.1/16

Packet: Destination IP address: 10.1.1.1

R1

R2

R4

10/8 → R3
10.1/16 → R4
20/8 → R5
30/8 → R6
…..

R2's IP routing table

10.1.1.1 && FF.0.0.0
vs.
10.0.0.0 && FF.0.0.0

Match!

10.1/16

IP route lookup:
Longest match routing

- **Based on destination IP address**

R3

All 10/8 except 10.1/16

Packet: Destination IP address: 10.1.1.1

R1

R2

R4

10/8 → R3	
10.1/16 → R4	
20/8 → R5	
30/8 → R6	
…..	

10.1.1.1 && FF.FF.0.0
vs.
10.1.0.0 && FF.FF.0.0

10.1/16

Match as well!

R2's IP routing table

IP route lookup:
Longest match routing

- ## Based on destination IP address

R3

All 10/8 except 10.1/16

Packet: Destination IP address: 10.1.1.1

R1

R2

R4

10.1/16

10/8 → R3	
10.1/16 → R4	
20/8 → R5	
30/8 → R6	
…..	

10.1.1.1 && FF.0.0.0
vs.
20.0.0.0 && FF.0.0.0

Does not match!

R2's IP routing table

IP route lookup:
Longest match routing

- **Based on destination IP address**

Packet: Destination IP address: 10.1.1.1

R3

All 10/8 except 10.1/16

R1   R2

R4

10.1/16

```
10/8 → R3
10.1/16 → R4
20/8 → R5
30/8 → R6
…..
```

10.1.1.1 && FF.0.0.0
vs.
30.0.0.0 && FF.0.0.0

Does not match!

R2's IP routing table

IP route lookup:
Longest match routing

- **Based on destination IP address**

R3

All 10/8 except 10.1/16

Packet: Destination IP address: 10.1.1.1

R1

R2

R4

10.1/16

10/8 → R3	
10.1/16 → R4	← Longest match, 16 bit netmask
20/8 → R5	
30/8 → R6	
…..	

R2's IP routing table

IP Forwarding

- **Router makes decision on which interface a packet is sent to**

- **Forwarding table populated by routing process**

- **Forwarding decisions:**

  destination address

  class of service (fair queuing, precedence, others)

  local requirements (packet filtering)

- **Can be aided by special hardware**

Routing Tables Feed the Forwarding Table

Forward Table (FIB)

Routing Information Base (RIB)

BGP 4 Routing Table

OSPF – Link State Database

Static Routes
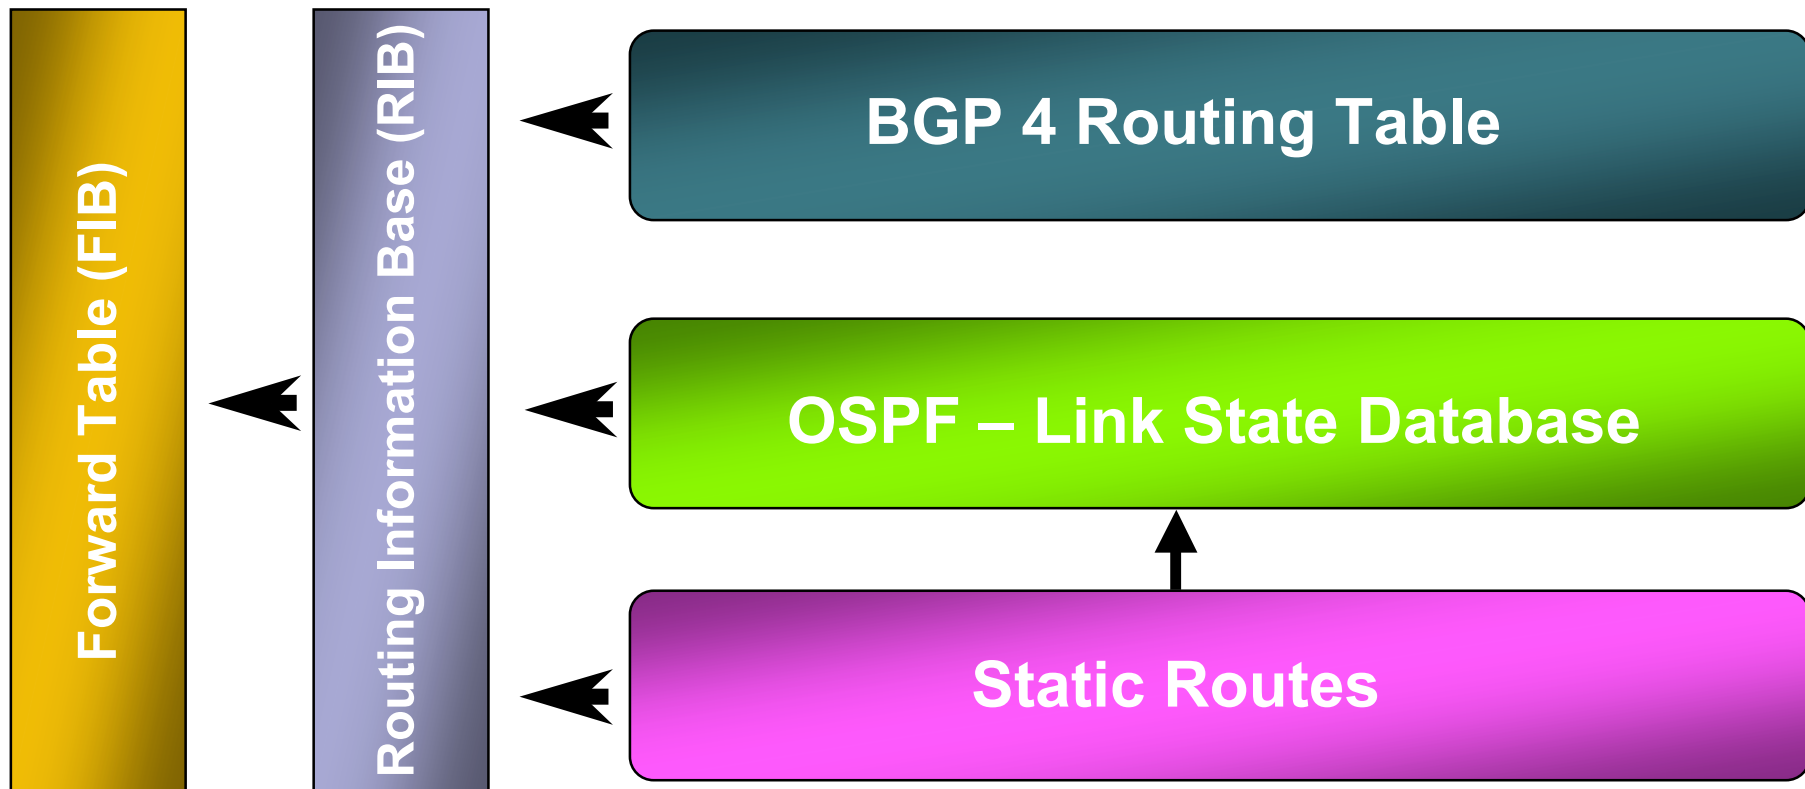
RIBs and FIBs

- ## FIB is the Forwarding Table

    It contains destinations and the interfaces to get to those destinations

    Used by the router to figure out where to send the packet

    Careful! Some people call this a route!

- ## RIB is the Routing Table

    It contains a list of all the destinations and the various next hops used to get to those destinations – and lots of other information too!

    One destination can have lots of possible next-hops – only the best next-hop goes into the FIB

Explicit versus Default Routing

- **Default:**

  simple, cheap (cycles, memory, bandwidth)

  low granularity (metric games)

- **Explicit (default free zone)**

  high overhead, complex, high cost, high granularity

- **Hybrid**

  minimise overhead

 **provide useful granularity**

  requires some filtering knowledge

Egress Traffic

- **How packets leave your network**

- **Egress traffic depends on:**

    route availability (what others send you)

    route acceptance (what you accept from others)

    policy and tuning (what you do with routes from others)

    Peering and transit agreements

Ingress Traffic

- **How packets get to your network and your customers' networks**

- **Ingress traffic depends on:**

    what information you send and to whom

    based on your addressing and AS's

    based on others' policy (what they accept from you and what they do with it)

Autonomous System (AS)

AS 100

- **Collection of networks with same routing policy**

- **Single routing protocol**

- **Usually under single ownership, trust and administrative control**

Definition of terms

- **Neighbours**

  AS's which directly exchange routing information

  Routers which exchange routing information

- **Announce**

  send routing information to a neighbour

- **Accept**

  receive and use routing information sent by a neighbour

- **Originate**

  insert routing information into external announcements (usually as a result of the IGP)

- **Peers**

  routers in neighbouring AS's or within one AS which exchange routing and policy information

Routing flow and packet flow

For networks in AS1 and AS2 to communicate:

AS1 must announce to AS2

AS2 must accept from AS1

AS2 must announce to AS1

AS1 must accept from AS2

Routing flow and Traffic flow

- **Traffic flow is always in the opposite direction of the flow of Routing information**

    Filtering outgoing routing information inhibits traffic flow inbound

    Filtering inbound routing information inhibits traffic flow outbound

# Routing Flow/Packet Flow:
With multiple ASes

For net N1 in AS1 to send traffic to net N16 in AS16:

- **AS16 must originate and announce N16 to AS8.**

- **AS8 must accept N16 from AS16.**

- **AS8 must announce N16 to AS1 or AS34.**

- **AS1 must accept N16 from AS8 or AS34.**

For two-way packet flow, similar policies must exist for N1.

Routing Flow/Packet Flow: With multiple ASes

AS 1

N1

AS 34

AS 8

AS16

N16

As multiple paths between sites are implemented it is easy to see how policies can become quite complex.

Routing Policy

- **Used to control traffic flow in and out of an ISP network**

- **ISP makes decisions on what routing information to accept and discard from its neighbours**

  Individual routes

  Routes originated by specific ASes

  Routes traversing specific ASes

  Routes belonging to other groupings

  Groupings which you define as you see fit

Routing Policy Limitations

- **AS99 uses red link for traffic to the red AS and the green link for remaining traffic**

- **To implement this policy, AS99 has to:**

    Accept routes originating from the red AS on the red link

    Accept all other routes on the green link

Routing Policy Limitations

red — Internet — AS22 — AS99

red

green — green

packet flow

- **AS99 would like packets coming from the green AS to use the green link.**

- **But unless AS22 cooperates in pushing traffic from the green AS down the green link, there is very little that AS99 can do to achieve this aim**

Routing Policy Issues

- **191000 prefixes (not realistic to set policy on all of them individually)**

- **22500 origin AS's (too many)**

- **routes tied to a specific AS or path may be unstable regardless of connectivity**

- **groups of AS's are a natural abstraction for filtering purposes**

Routing Protocols

We now know what routing means…

…but what do the routers get up to?

And why are we doing this anyway?

1: How Does Routing Work?

- **Internet is made up of the ISPs who connect to each other's networks**

- **How does an ISP in Kenya tell an ISP in Japan what customers they have?**

- **And how does that ISP send data packets to the customers of the ISP in Japan, and get responses back**

  After all, as on a local ethernet, two way packet flow is needed for communication between two devices

2: How Does Routing Work?

- **ISP in Kenya could buy a direct connection to the ISP in Japan**

    But this doesn't scale – thousands of ISPs, would need thousands of connections, and cost would be astronomical

- **Instead, ISP in Kenya tells his neighbouring ISPs what customers he has**

    And the neighbouring ISPs pass this information on to their neighbours, and so on

    This process repeats until the information reaches the ISP in Japan

3: How Does Routing Work?

- **This process is called "Routing"**

- **The mechanisms used are called "Routing Protocols"**

- **Routing and Routing Protocols ensures that the Internet can scale, that thousands of ISPs can provide connectivity to each other, giving us the Internet we see today**

4: How Does Routing Work?

- **ISP in Kenya doesn't actually tell his neighbouring ISPs the names of the customers**

    (network equipment does not understand names)

- **Instead, he has received an IP address block as a member of the Regional Internet Registry serving Kenya**

    His customers have received address space from this address block as part of their "Internet service"

    And he announces this address block to his neighbouring ISPs – this is called announcing a "route"

Routing Protocols

- **Routers use "routing protocols" to exchange routing information with each other**

  IGP is used to refer to the process running on routers inside an ISP's network

  EGP is used to refer to the process running between routers bordering directly connected ISP networks

What Is an IGP?

- **Interior Gateway Protocol**
- **Within an Autonomous System**
- **Carries information about internal infrastructure prefixes**
- **Examples – OSPF, ISIS, EIGRP**

Why Do We Need an IGP?

- **ISP backbone scaling**

    Hierarchy

    Limiting scope of failure

    Only used for ISP's infrastructure addresses, not customers or anything else

    Design goal is to minimise number of prefixes in IGP to aid scalability and rapid convergence

What Is an EGP?

- **Exterior Gateway Protocol**

- **Used to convey routing information between Autonomous Systems**

- **De-coupled from the IGP**

- **Current EGP is BGP**

Why Do We Need an EGP?

- **Scaling to large network**

  Hierarchy

  Limit scope of failure

- **Define Administrative Boundary**

- **Policy**

  Control reachability of prefixes

  Merge separate organizations

  Connect multiple IGPs

Interior versus Exterior Routing Protocols

- **Interior**

  automatic neighbour discovery

  generally trust your IGP routers

  prefixes go to all IGP routers

  binds routers in one AS together

- **Exterior**

  specifically configured peers

  connecting with outside networks

  set administrative boundaries

  binds AS's together

Interior versus Exterior Routing Protocols

- **Interior**

  Carries ISP infrastructure addresses only

 **ISPs aim to keep the IGP small for efficiency and scalability**

- **Exterior**

  Carries customer prefixes

  Carries Internet prefixes

  EGPs are independent of ISP network topology

Hierarchy of Routing Protocols

Other ISPs

BGP4

**BGP4
and OSPF/ISIS**

BGP4

IXP

Static/BGP4

Customers

FYI: IOS Default Administrative Distances

Route Source	Default Distance
Connected Interface	0
Static Route	1
Enhanced IGRP Summary Route	5
External BGP	20
Internal Enhanced IGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
EGP	140
External Enhanced IGRP	170
Internal BGP	200
Unknown	255

Routing Basics

ISP/IXP Workshops