

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 blue router icon with a white question mark in the center, set against a white background.

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 packet
- "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 packet




IP route lookup: Longest match routing

- Based on destination IP packet




IP route lookup: Longest match routing

- Based on destination IP packet




IP route lookup: Longest match routing

- Based on destination IP packet




IP route lookup: Longest match routing

- Based on destination IP packet



IP route lookup: Longest match routing

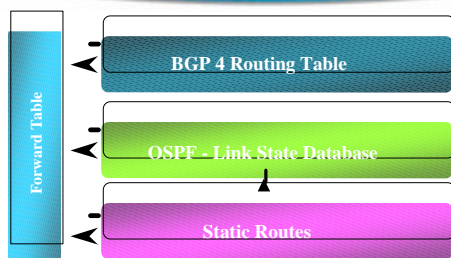
- Based on destination IP packet



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



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)




- 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
- **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 flowing in
 - filtering incoming routing information inhibits traffic flowing out

Routing policy limitations




AS99 uses red link for traffic going to the red AS and green link for traffic going to the green AS

To implement this policy for AS99:

- accept routes originating in the red AS on the red link
- accept all other routes on the green link


Routing policy limitations



For packets flowing *toward* AS 99:

Unless AS 22 and all other intermediate AS's co-operate in pushing *green* traffic to the *green* link then some reasonable policies can not be implemented.

Routing policy with multiple AS's




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 policy with multiple AS's



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

Granularity of routing policy

- What to announce/accept
- Preferences between multiple accepts
 - single route
 - routes originated by single AS
 - routes originated by a group of AS's
 - routes traversing specific path
 - routes traversing specific AS
 - routes belonging to other groupings (including combinations)

Routing Policy Issues

- 80000+ prefixes (not realistic to set policy on all of them individually)
- 7500+ 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

What Is an IGP?

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

ENP1327 Workshop © 2006 Cisco Systems, Inc. www.cisco.com

31

Why Do We Need an IGP?

- ISP backbone scaling
 - Hierarchy
 - Modular infrastructure construction
 - Limiting scope of failure
 - Healing of infrastructure faults using dynamic routing with fast convergence

ENP1327 Workshop © 2006 Cisco Systems, Inc. www.cisco.com

32

What Is an EGP?

- Exterior Gateway Protocol
- Used to convey routing information between Autonomous Systems
- De-coupled from the IGP
- Current EGP is BGP

ENP1327 Workshop © 2006 Cisco Systems, Inc. www.cisco.com

33

Why Do We Need an EGP?

- Scaling to large network
 - Hierarchy
 - Limit scope of failure
- Policy
 - Control reachability to prefixes
 - Merge separate organizations
 - Connect multiple IGPs

ENP1327 Workshop © 2006 Cisco Systems, Inc. www.cisco.com

34

Interior versus Exterior Routing Protocols

- | | |
|---|---|
| <ul style="list-style-type: none"> • Interior <ul style="list-style-type: none"> automatic neighbour discovery generally trust your IGP routers prefixes go to all IGP routers binds routers in one AS together | <ul style="list-style-type: none"> • Exterior <ul style="list-style-type: none"> specifically configured peers connecting with outside networks set administrative boundaries binds AS's together |
|---|---|

ENP1327 Workshop © 2006 Cisco Systems, Inc. www.cisco.com

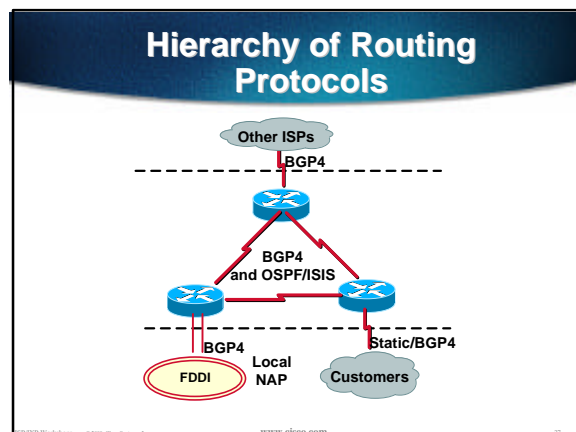
35

Interior versus Exterior Routing Protocols

- | | |
|--|--|
| <ul style="list-style-type: none"> • Interior <ul style="list-style-type: none"> Carries ISP infrastructure addresses only ISPs aim to keep the IGP small for efficiency and scalability | <ul style="list-style-type: none"> • Exterior <ul style="list-style-type: none"> Carries customer prefixes Carries Internet prefixes EGPs are independent of ISP network topology |
|--|--|

ENP1327 Workshop © 2006 Cisco Systems, Inc. www.cisco.com

36



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

