Routing & Protocols

Today's Talk

1. Terminology
2. Routing
3. Static Routes
4. Interior Gateway Protocols
5. Exterior Gateway Protocols
6. Building an ISP network

Terminology

- network number
- prefix
- mask (or length)

Static routes

hand configured routing

1. tell the router which way to send packets
2. based upon final packet destination

Static routes

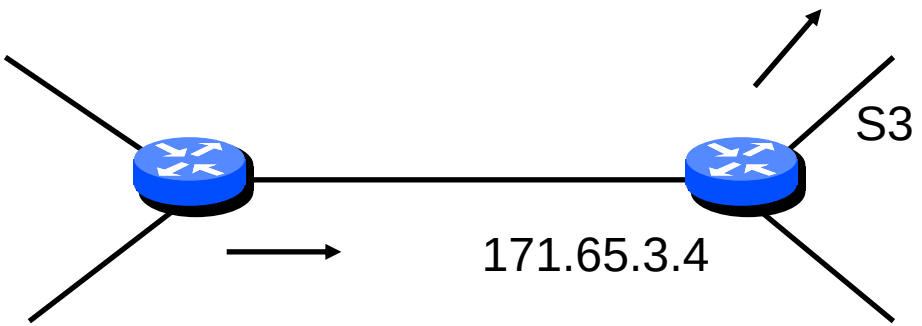


```
1. ip route 10.0.0.0
   255.0.0.0 serial 3
2. ip route 131.108.0.0
   255.255.0.0 171.65.3.4
```

Terminology

Interior Gateway Protocol (IGP)

1. RIP, IGRP, HELLO, OSPF
2. Primary goal is optimal connectivity
3. Strong distance metrics
4. May not have good administrative controls

Terminology

Distance vector protocols

1. listen to neighboring routers
2. install routes in table, lowest distance wins
3. advertise all routes in table
4. very simple
5. very stupid

Terminology

Distance vector protocols

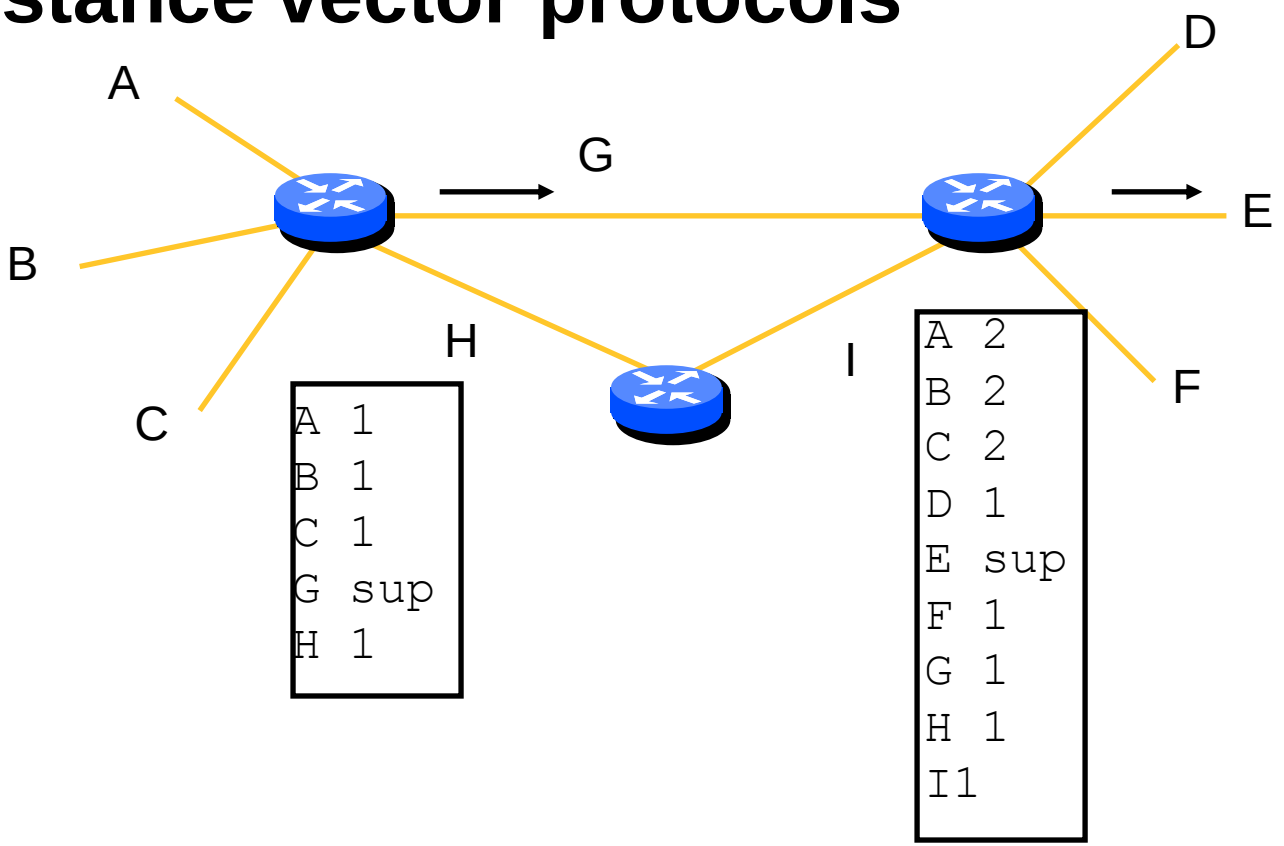

Terminology

Link state protocols

1. information about adjacencies sent to all routers
2. each router builds a topology database
3. a "shortest path" algorithm is used to find best route
4. converge as quickly as databases can be updated

Terminology

Link state protocols

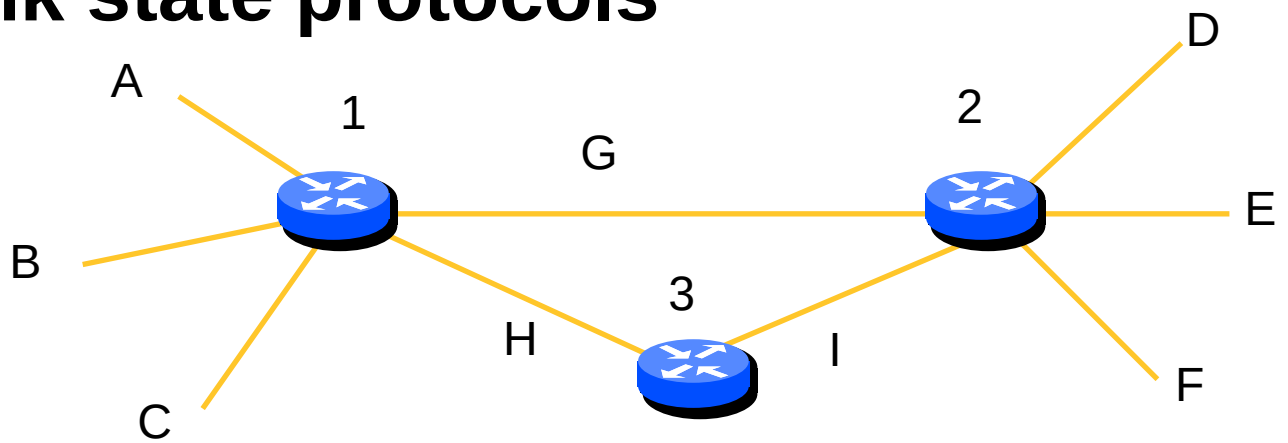

router 1
A, B, C, G, H

router 3
H, I

router 2
D, E, F, G, I

A - 1 - G - 2 - D

Interior Gateway Protocols

Routing Information Protocol (RIP)

- IP only
- distance vector protocol
- slow convergence
- does not carry mask information
- reasonably simple design & configuration
- does not scale (maximum 15 hops)
- poor metrics (hop-count)

Interior Gateway Protocols

Interior Gateway Routing Protocol (IGRP)

1. IP only
2. distance vector protocol
3. slow convergence (like RIP)
4. does not carry mask information (like RIP)
5. very simple design & configuration
    1. powerful proprietary metric
    2. load sharing across diverse links

Interior Gateway Protocols

The IGRP metric

1. always get optimal routing
2. metric vector, not single value
    1. bandwidth
    2. delay
    3. hops
    4. reliability
    5. loading

Interior Gateway Protocols

Enhanced IGRP

1. multi-protocol (IP, IPX, Appletalk)
2. fast convergence (like OSPF)
3. very simple design & configuration (like IGRP)
    1. IGRP metric
    2. allows load sharing across diverse links

Interior Gateway Protocols

Enhanced IGRP

1. distance vector based protocol
2. NOT a Bellman-Ford protocol
   Uses "dual" algorithm
3. alternative to OSPF & I-ISIS
4. can be bandwidth intensive on slow links

Interior Gateway Protocols

Integrated IS-IS (I-ISIS)

1. multi-protocol (CLNP, IP, IPX, ...)
2. link state protocol
3. fast convergence
4. design and architecture moderately complex
5. configuration may be simple

Interior Gateway Protocols

Open Shortest Path First (OSPF)

1. IS - IS = 0

Interior Gateway Protocols

Open Shortest Path First (OSPF)

1. IP only
2. link state protocol
3. fast convergence
4. design and architecture very complex
5. configuration can be simple

Interior Gateway Protocols

Which to use?

1. Your interior network is actually VERY simple.
2. Your IGP should only carry your routes and your direct customers'

Interior Gateway Protcols

Problems with "classic" protocols

1. slow convergence
2. count to infinity
3. no mask information
    1. no CIDR
    2. no VLSM
    3. no subnet 0

Interior Gateway Protocols

Slow convergence

1. **advertisement period**
    1. **entire routing table dumped every n seconds**
2. **timeout period**
    1. **usually 3 times advertisement period**
3. **RIP values are normally 30 and 90 seconds!**

Interior Gateway Protocols

Count to infinty problem

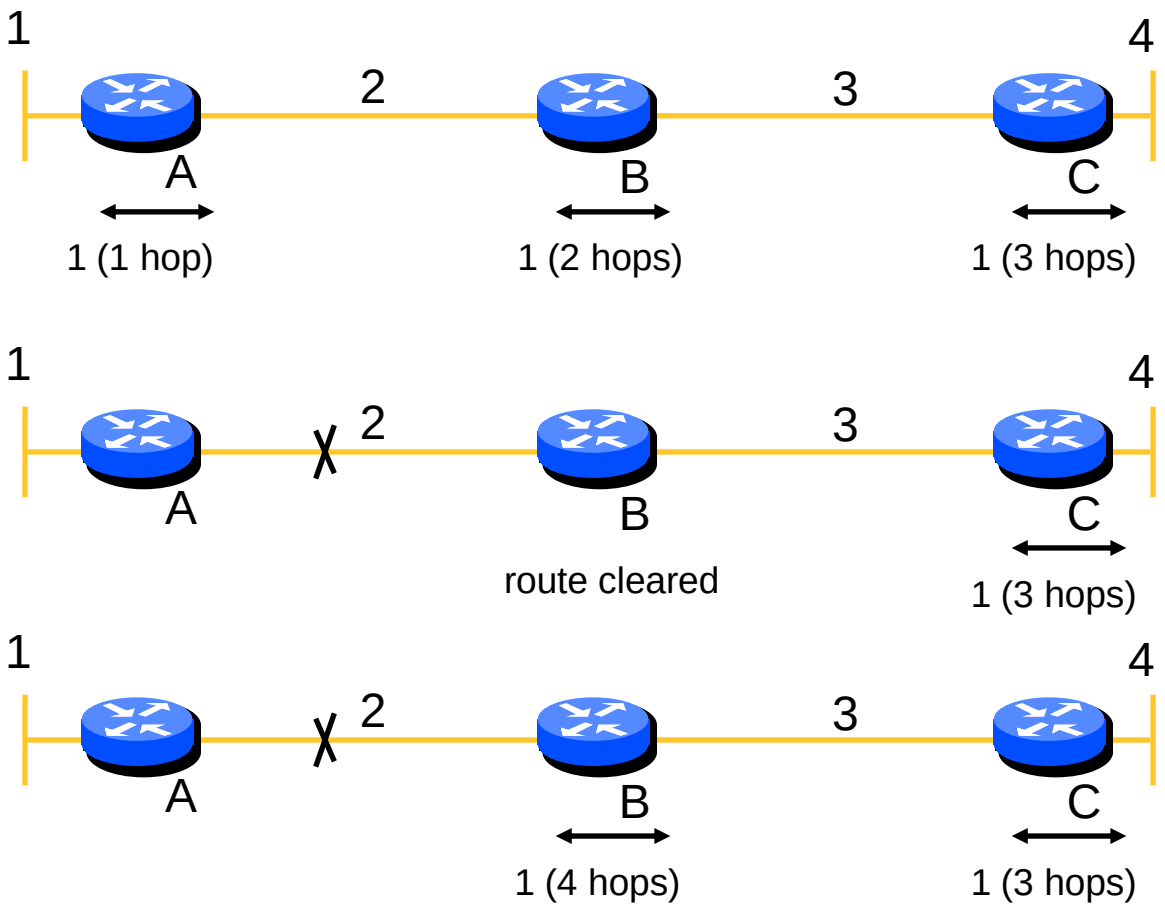

Interior Gateway Protocols

Count to infinity: split-horizon

1. Don't feed selected route back to source
   1. no feedback on source interface
   2. no feedback to source neighbor

Interior Gateway Protocols

Count to infinity: split-horizon

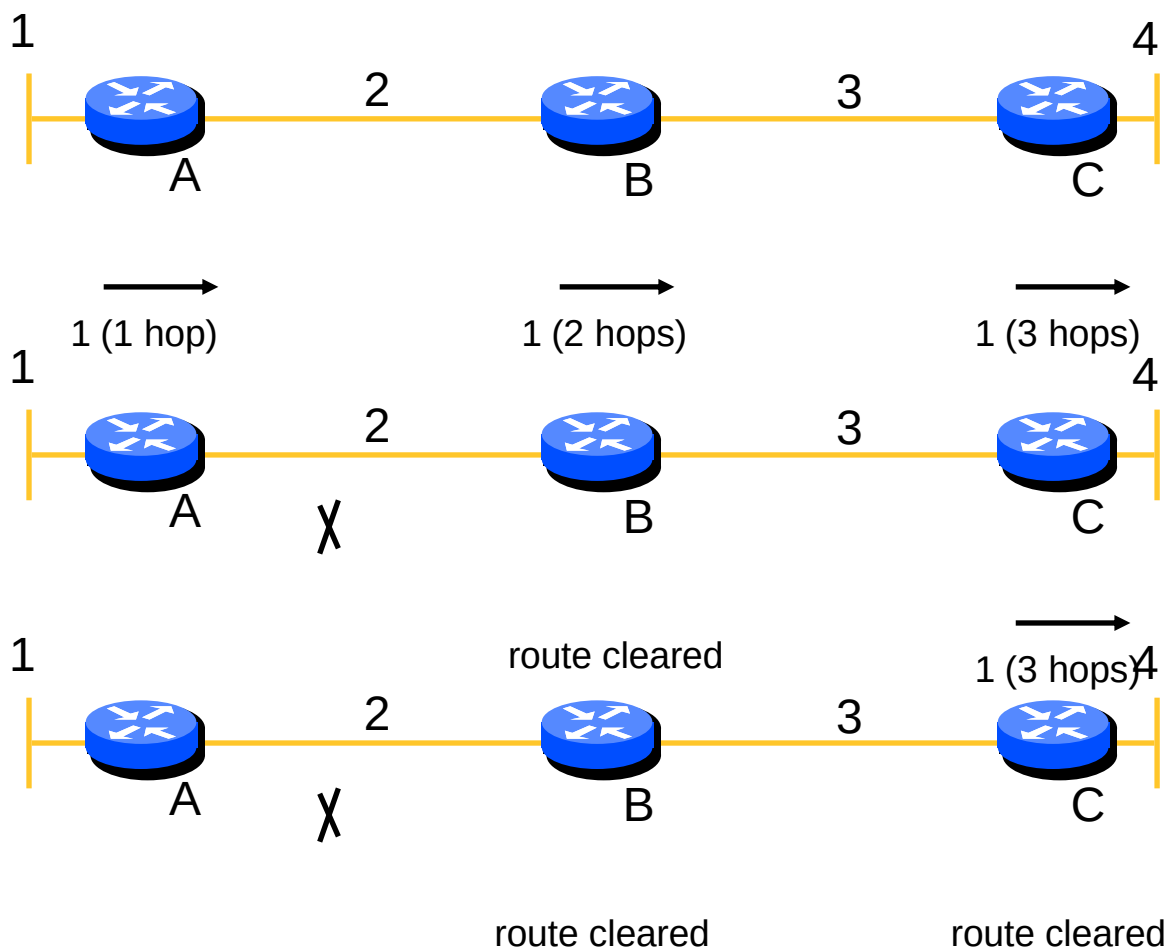

Interior Gateway Protocols

Count to infinity: hold-down

1. Split horizon not sufficient!
2. Holddown period
    1. interval during which "less attractive" updates are ignored

Interior Gateway Protocols

Count to infinity: hold-down

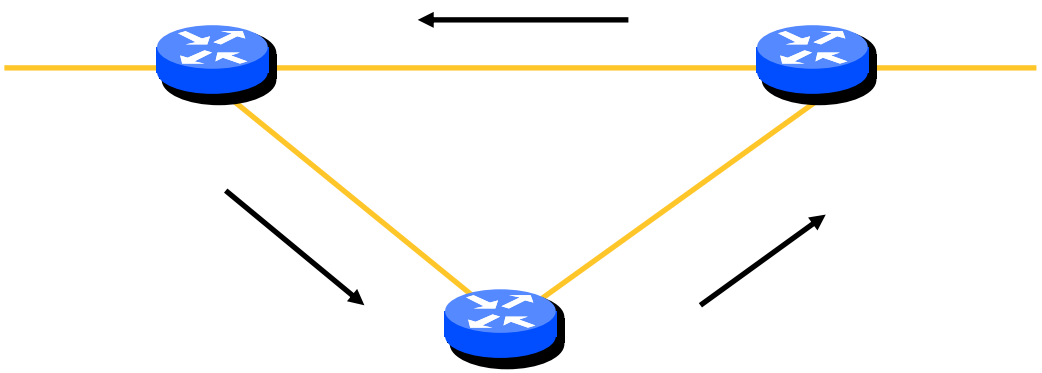

Interior Gateway Protocols

The universal rule

- You will always trade bandwidth for speed of convergence

Interior Gateway Protocols

OSPF configuration

1. myth
   1. OSPF is hard to use
2. reality:

```
1. router ospf 1
   network 192.111.107.0 0.0.0.255 area 0
```

Interior Gateway Protocols

OSPF operation

1. every OSPF router sends out 'hello' packets
2. hello packets used to determine if neighbor is up
3. hello packets are small easy to process packets
4. hello packets are sent periodically (usually short interval)

Interior Gateway Protocols

OSPF operation

1. once an adjacency is established, trade information with your neighbor
2. topology information is packaged in a "link state announcement"
3. announcements are sent ONCE, and only updated if there's a change
    1. (or every 45mins...)

Interior Gateway Protocols

OSPF operation

1. change occurs
2. broadcast change
3. run SPF algorithm
4. install output into forwarding table

Interior Gateway Protocols

making OSPF scale

1. each link transition causes a broadcast and SPF run
2. OSPF can group routers to appear as one single router
3. OSPF areas

Interior Gateway Protocols

OSPF areas (before)

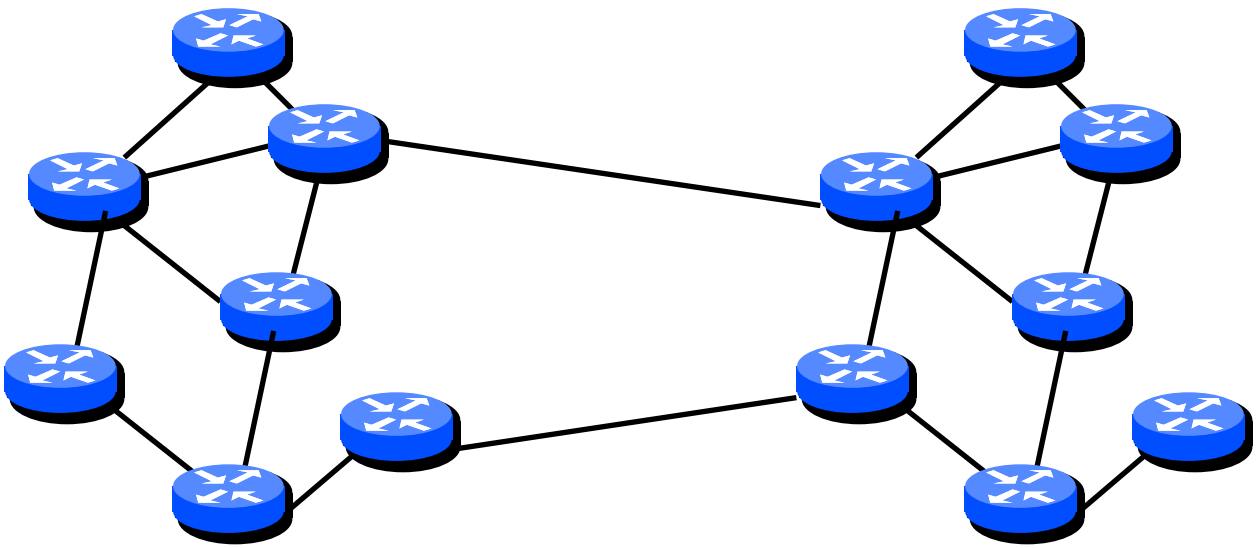

Interior Gateway Protocols

OSPF areas (after)

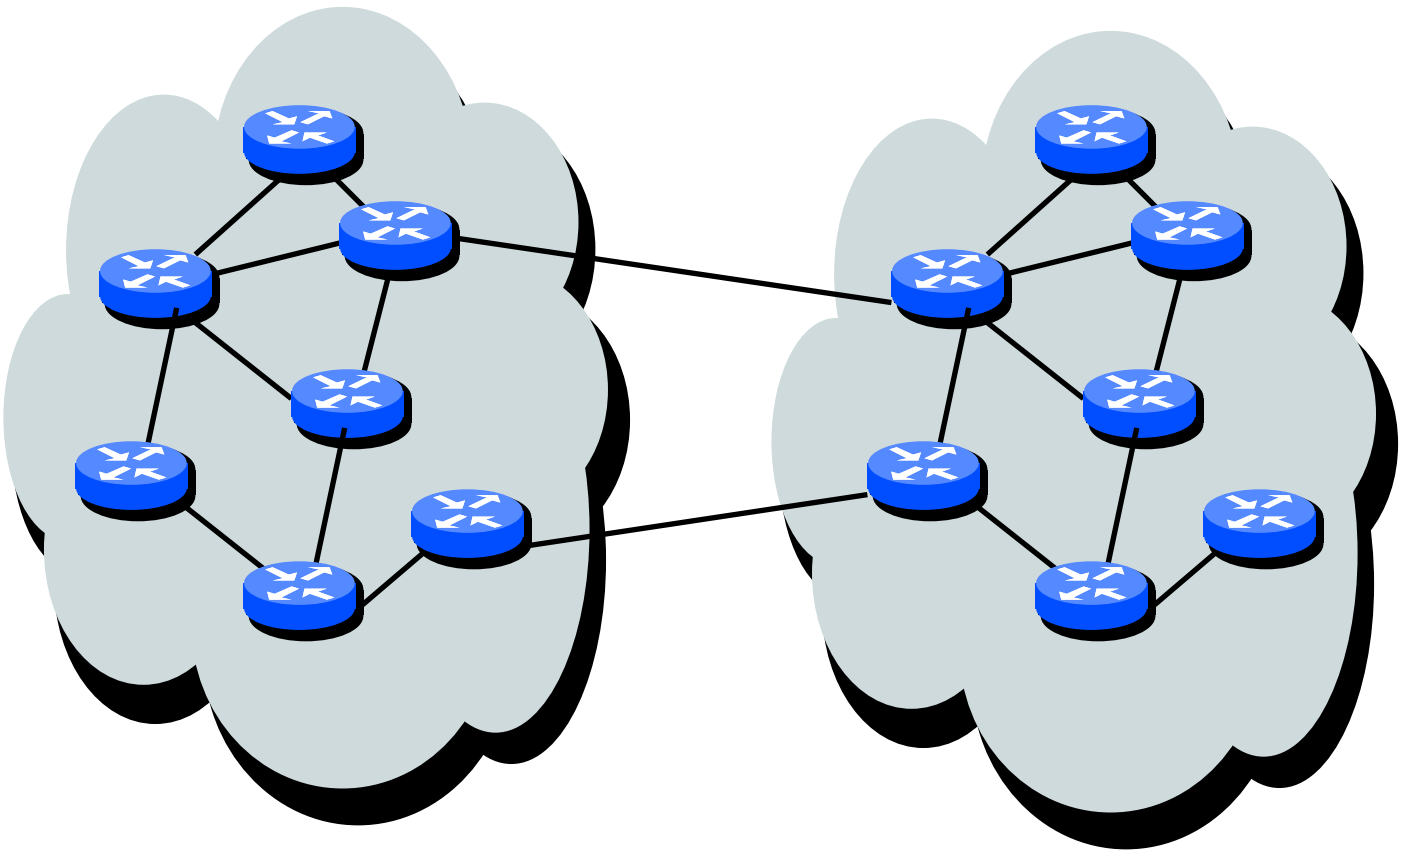

Interior Gateway Protocols

OSPF areas - partitioning

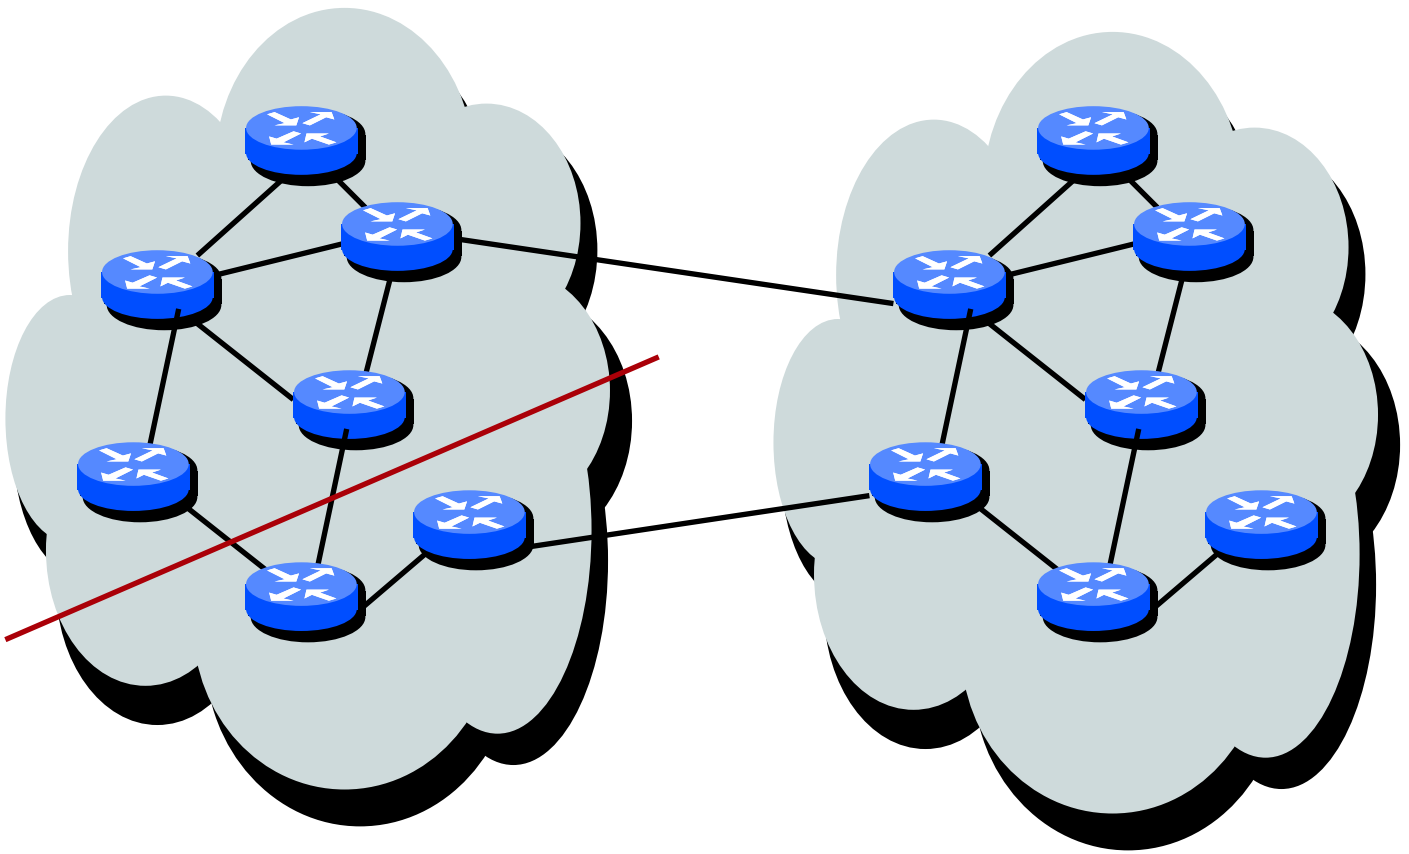

Interior Gateway Protocols

OSPF areas - partition repair

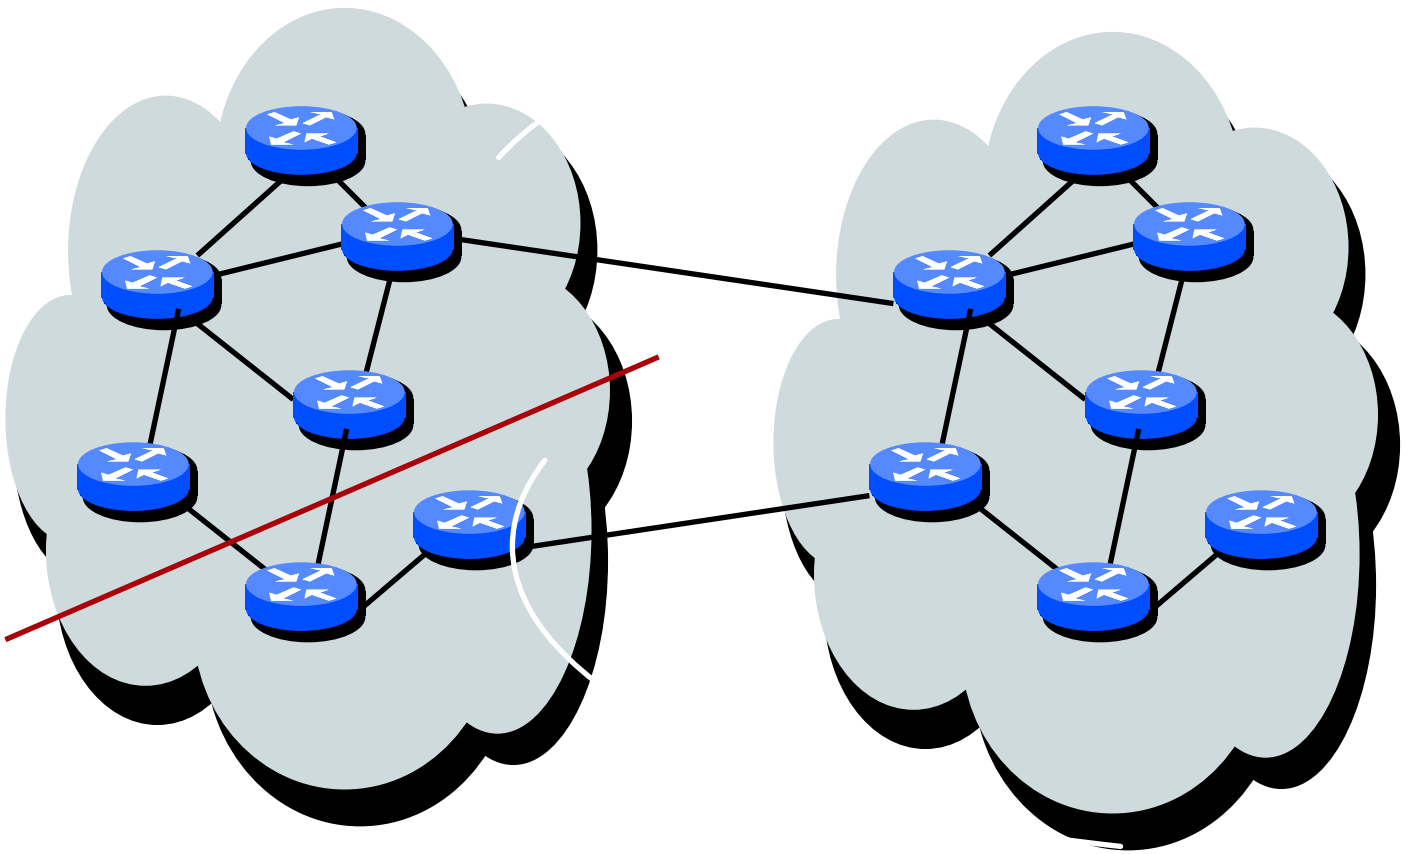

Interior Gateway Protocols

OSPF areas

1. rule of thumb:
   no more than 150 routers/area
2. reality:
   no more than 500 routers/area
3. backbone "area" is an area
4. proper use of areas reduce bandwidth & CPU utilization

Interior Gateway Protocols

EIGRP operation

1. design goals were
    1. make it as fast as OSPF & IS-IS
    2. make it trivial to configure
    3. easy migration from IGRP

Interior Gateway Protocols

EIGRP operation

```
1. router eigrp 1
   network 192.108.0.0 mask 255.255.0.0
```

Interior Gateway Protocols

EIGRP operation - caveats

1. nothing is for free
2. EIGRP works best on high speed links
3. EIGRP doesn't scale well in high-meshed frame-relay networks
    1. star networks OK

Interior Gateway Protocols

summarization

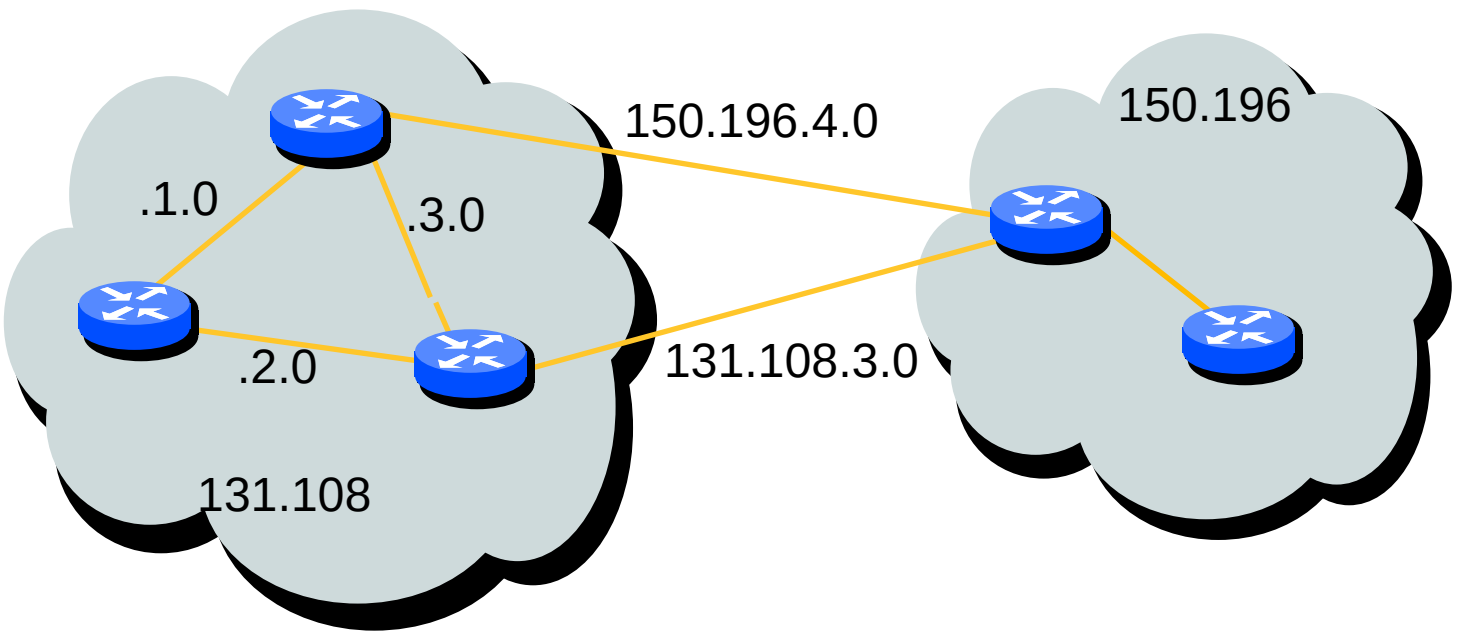

1. **classful routing protocols naturally summarize to network numbers at boundaries**

Interior Gateway Protocols

summarization

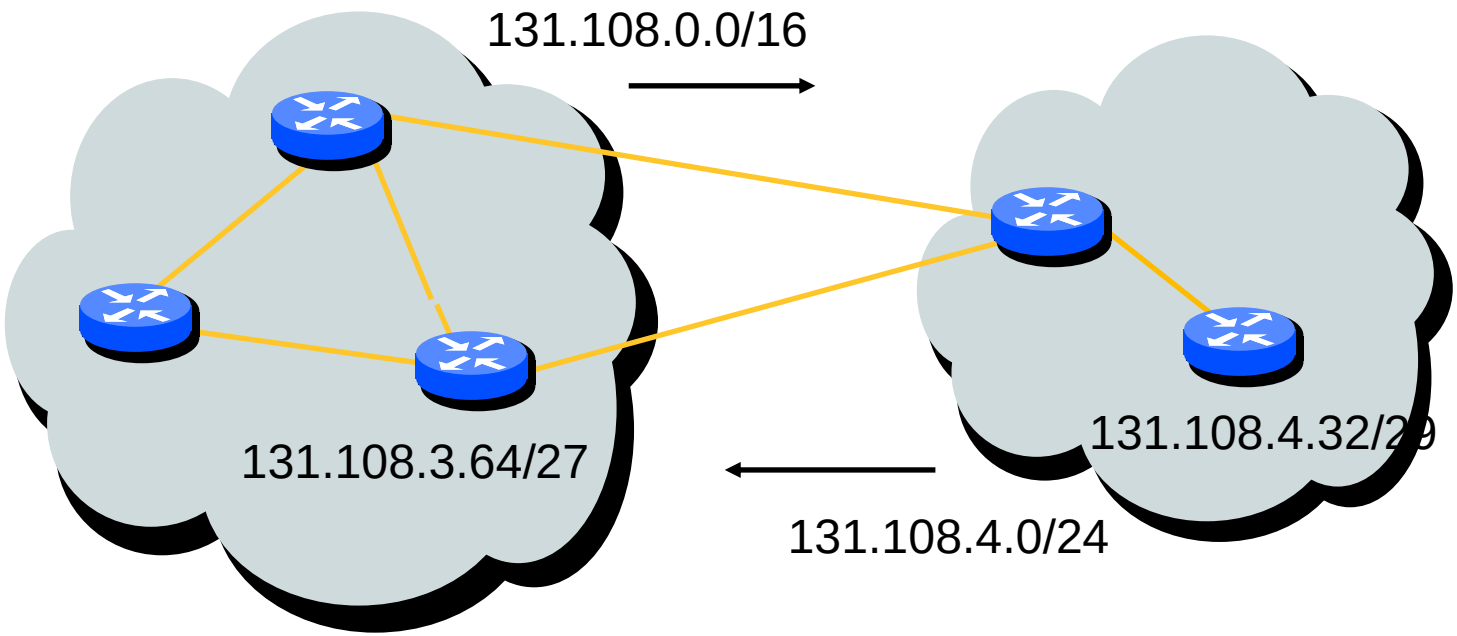

1. **classless routing protocols summarize at arbitrary bit boundaries**

Interior Gateway Protocols

route filtering

1. pseudo-security (bad idea!)
2. low bandwidth links
3. eliminate unnecessary information

Interior Gateway Protocols

route filtering

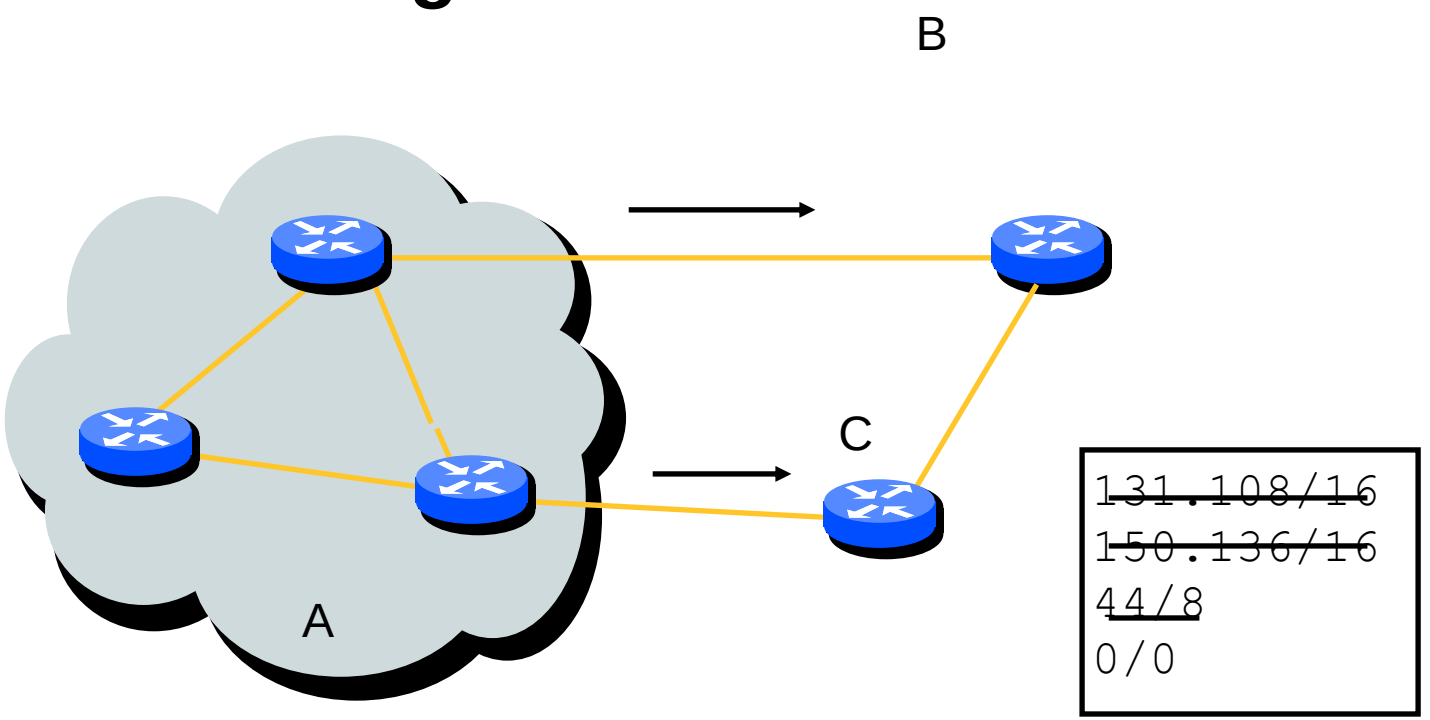

Interior Gateway Protocols

redistribution

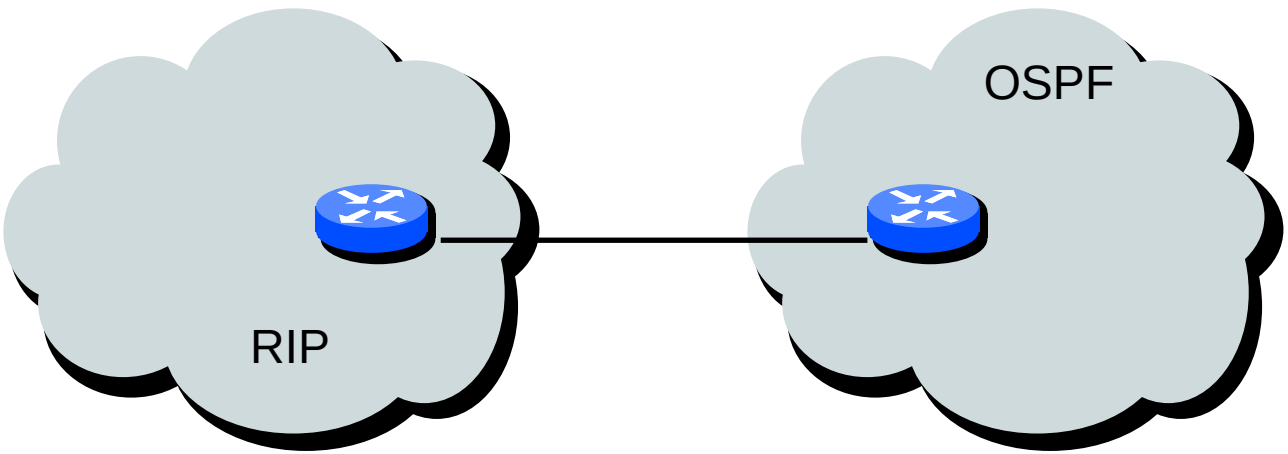

1. you run OSPF
2. your neighbor runs RIP

Interior Gateway Protocols

redistribution

1. run RIP on their interface
2.
```
router rip
  network 192.111.107.0
```
3. configure OSPF to redistribute RIP
4.
```
router ospf 1
  network 135.111.104.0
  0.0.0.255 area 0
  redistribute rip metric 10
```

Interior Gateway Protocols

redistribution

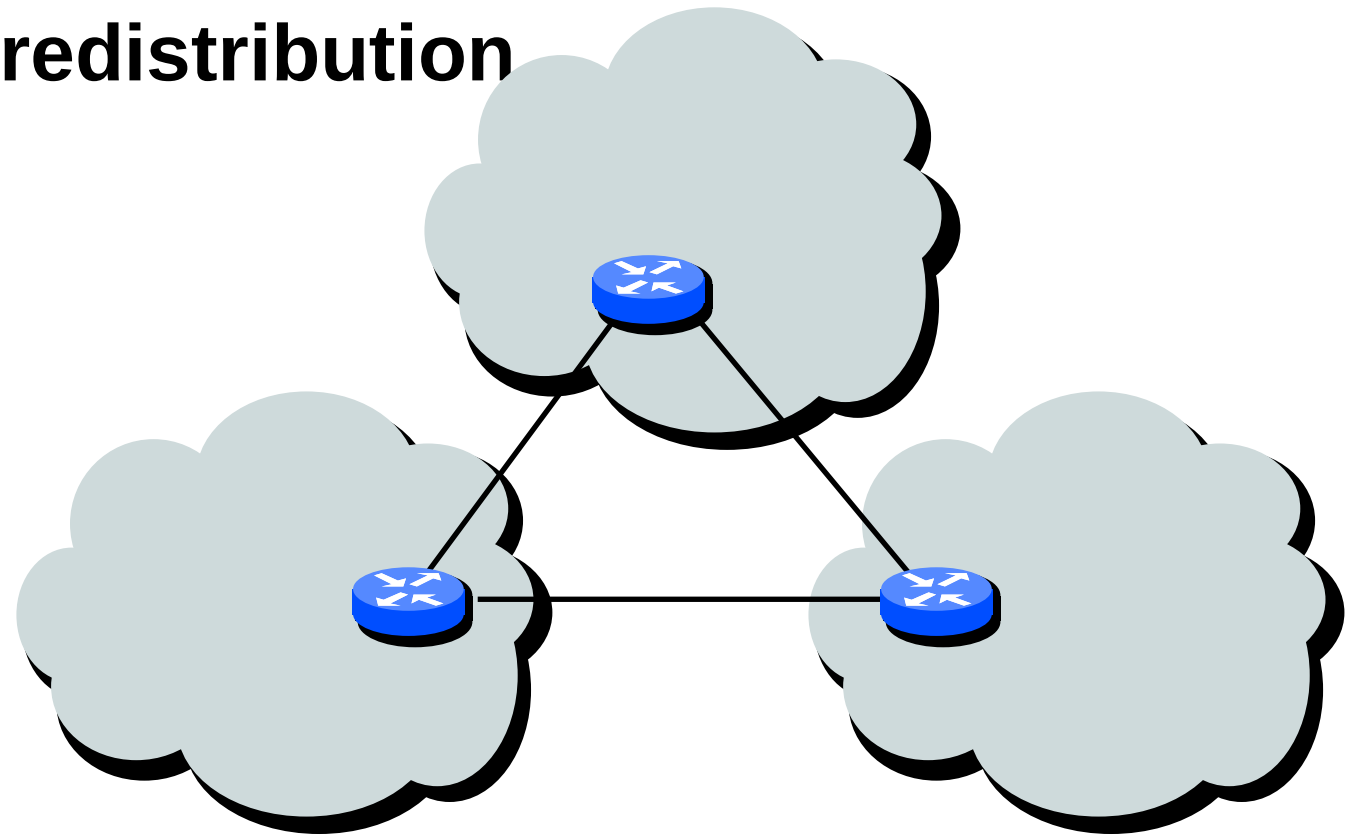

1. bi-directional redistribution MUST be filtered!

Interior Gateway Protocols

redistribution

```
1. router rip
   network 192.111.107.0
2. router ospf 1
   network 135.111.104.0 0.0.0.255
   area 0
   redistribute rip metric 10
   distribute-list 1 out rip
3. access-list 1 permit 192.111.107.0
   0.0.0.255
```

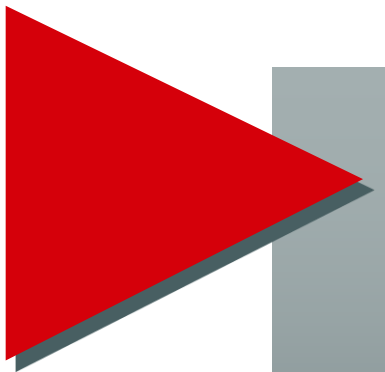

Exterior routing

Exterior routing

1. Terminology
2. What is exterior routing?
3. Routing protocols
4. Overview of BGP
5. Putting it all together
6. Further information

Terminology

Autonomous System

1. **A set of networks sharing the same routing policy.**
2. **Internal connectivity**
3. **One contiguous unit**
4. **Identified by "AS number"**
5. **Examples**
    1. **service provider**
    2. **multi-homed customer**
    3. **anyone needing policy discrimination**

Terminology

Exterior routes

- Routes learned from other autonomous systems

Terminology

Exterior Gateway Protocol

- egp vs EGP
- EGP, BGP, IDRP
- Primary goal is to provide reachability information outside administrative domain
- Secondary goal is administrative control
- Metrics may be arbitrary or weak

Terminology

Natural network mask

- Classful mask
  - Class A = 8 bits
    - networks 1...127
  - Class B = 16 bits
    - networks 128.0...191.255
  - Class C = 24 bits
    - networks 192.0.0...223.255.255

Terminology

DMZ network

- de-militarised zone
- area between North and South Korea
- shared network between ASs
  - before, neither AS carried it in IGP
  - now, both carry it in IGP

Terminology

DMZ network

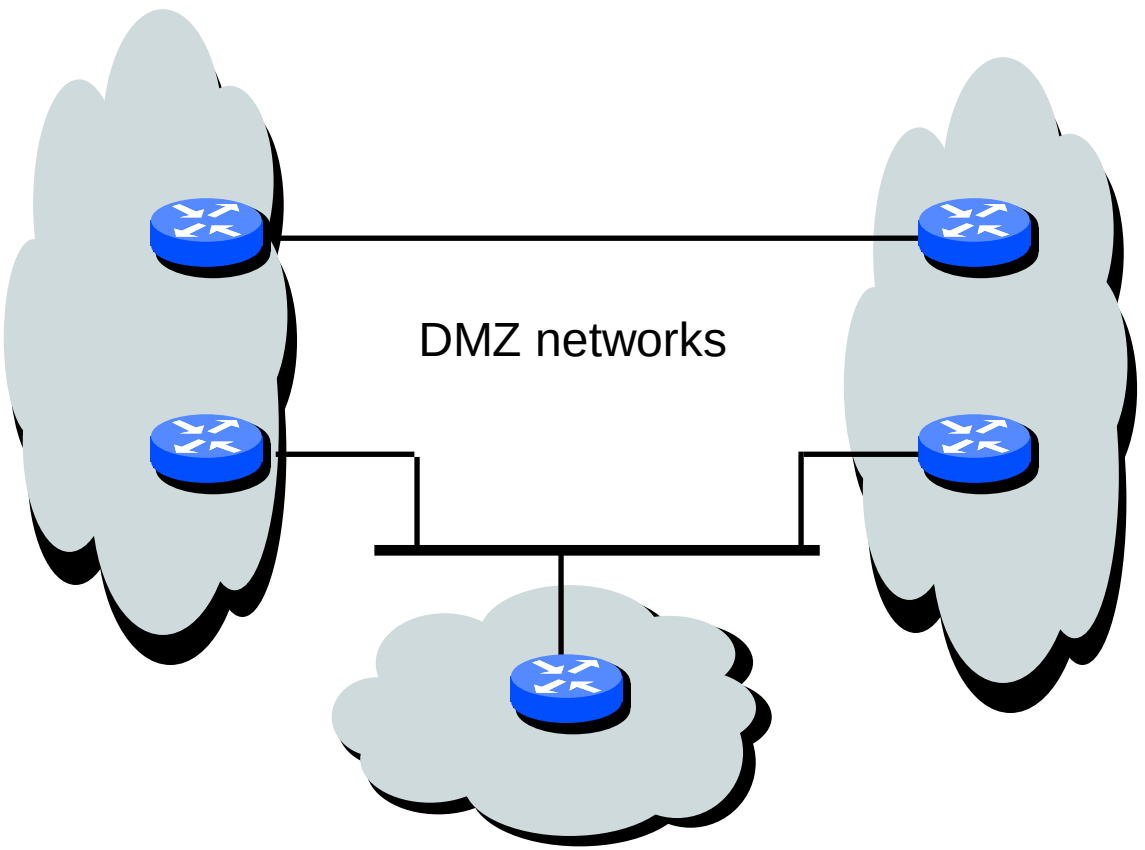

Why do we need exterior routing?

Why not make entire internet a single cloud?

- separate policy control
- filtering on networks doesn't scale well
- service provider selection given multiple choices
- everything must scale to hundreds of thousands of routes

Exterior Routing

1. static routes
2. multiple IGP instances
3. OSPF inter-domain routing
4. EGP
5. IDRP
6. BGP version 4

Exterior Routing

Static routes

- no path information
- very versatile
- low protocol overhead
- high maintenance overhead
- very very very bad convergence time
  - requires manual configuration

Thank you