Routing Protocols

Static and Dynamic Routes

- **Static routes**
  - The network administrator manually enters the routing information in the router.
  - based upon the final packet destination.
- **Dynamic routes**
  - Routers can learn the information from each other on the fly.
  - Using **routing protocol** to update routing information.
    - EGP, BGP, RIP, IGRP, EIGRP, OSPF …

Static vs. dynamic routes

- **Static routes**
  - When there is only one path to a destination network.
  - For hiding parts of an internetwork.
  - To test a particular link in a network.
- **Dynamic routes**
  - Automatic maintenance of routing table.
  - Relies on routing protocol to share knowledge.
  - Timely distribution of routing information and updates.
  - Routers can adjust to changing network conditions.

Static routes

- route add 192.168.1.0/24 dev eth0
- ip route 10.0.0.0 255.0.0.0 serial 3
- ip route 131.108.0.0 255.255.0.0 171.65.3.4

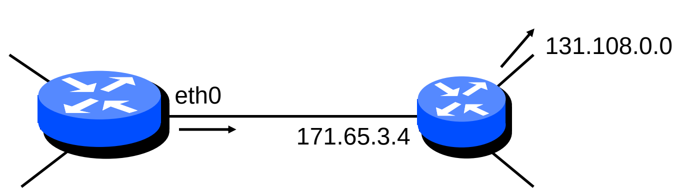

IGP and EGP

Two groups of Routing Protocols:

- **igp: Interior gateway protocols**
  - Used **within** an autonomous system.

  Examples: RIP, IGRP, EIGRP, OSPF

- **egp: Exterior gateway protocols**
  - Used to define routes **between** autonomous systems.
  - Examples: EGP, BGP

Autonomous System

- A set of networks sharing the same routing policy.
- Identified by "AS number"
- Internal connectivity
- One contiguous unit

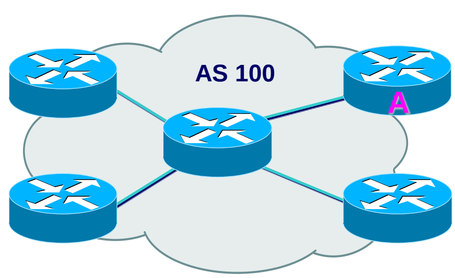

- Examples
  - service provider
  - multi-homed customer
  - anyone needing policy discrimination

IGP vs EGP

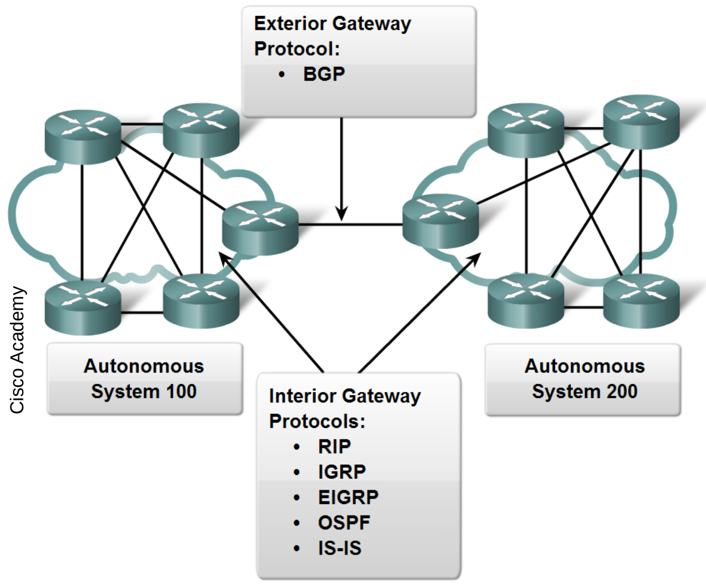

Dynamic Routing Protocols

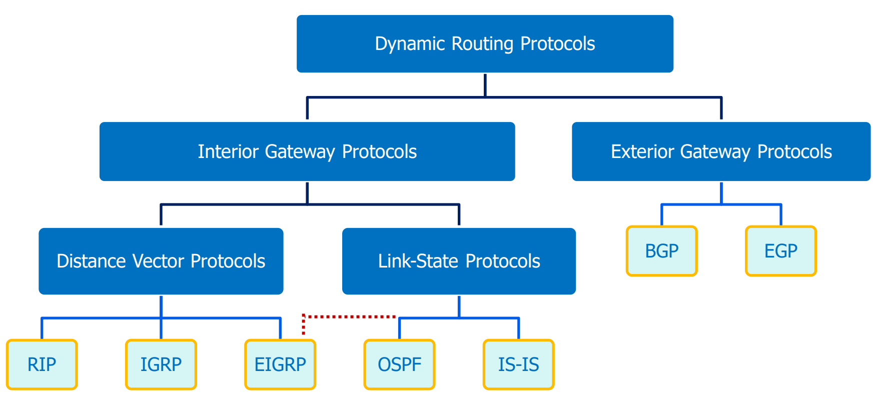

Dynamic Routing Protocols

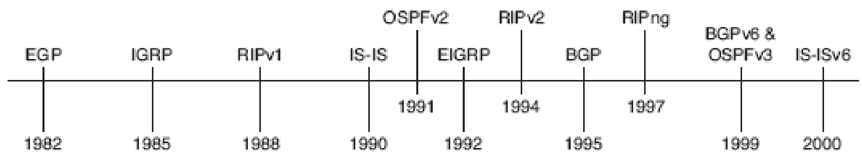

	Interior Gateway Protocols				Exterior Gateway Protocols
	Distance Vector Routing Protocols		Link State Routing Protocols		Path Vector
Classful	RIP	IGRP			EGP
Classless	RIPv2	EIGRP	OSPFv2	IS-IS	BGPv4
IPv6	RIPng	EIGRP for IPv6	OSPFv3	IS-IS for IPv6	BGPv4 for IPv6

Routing Protocols Metrics

- Metric: A value used to evaluate routes.
  - Examples: Number of Hops

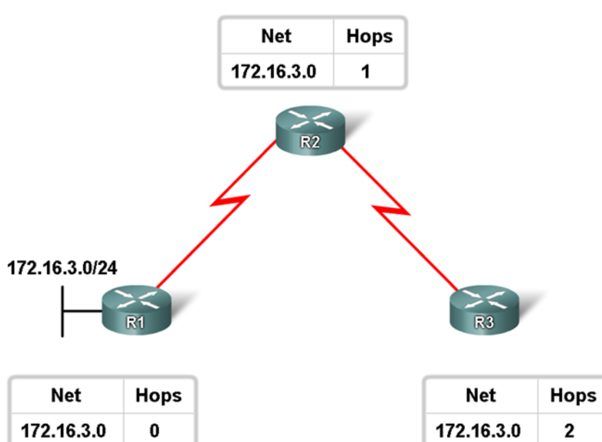

Routing Protocols Metrics

- Metrics used in IP routing protocols
  - Bandwidth
  - Cost
  - Delay
  - Hop count
  - Load
  - Reliability

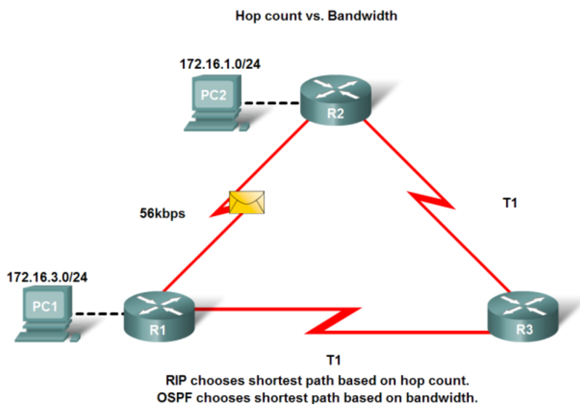

isep Instituto Superior de Engenharia do Porto

Interior Gateway Protocol (IGP)

Interior Gateway Protocol (IGP)

- Primary goal is optimal connectivity
- Strong distance metrics
- May not have good administrative controls
- Examples
  - RIP, IGRP, HELLO, OSPF
- Two basic methods
  - **Distance Vector** protocols
  - **Link State** Protocols

Distance Vector Protocols

- Usually apply the Bellman-Ford algorithm

- listen to neighboring routers
- install routes in table,
- use lowest distances
- advertise all routes in table

- very simple

Cold Start

Network	Interface	Hop
10.1.0.0	Fa0/0	0
10.2.0.0	S0/0/0	0

Network	Interface	Hop
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/1	0

Network	Interface	Hop
10.3.0.0	S0/0/1	0
10.4.0.0	Fa0/0	0

- Router powers up:
  - Knows nothing about the network topology
  - Does not know devices on the other end of its links.
  - Knows only information saved in NVRAM (startup-config).
- Network discovery
  - part of the process of the routing protocol algorithm
  - enables routers to learn about remote networks.

Cold Start

Network	Interface	Hop
10.1.0.0	Fa0/0	0
10.2.0.0	S0/0/0	0

Network	Interface	Hop
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/1	0

Network	Interface	Hop
10.3.0.0	S0/0/1	0
10.4.0.0	Fa0/0	0

Only knows about it's own networks

- **R1:**
  - 10.1.0.0 available through interface FastEthernet 0/0
  - 10.2.0.0 available through interface Serial 0/0/0
- **R2:**
  - 10.2.0.0 available through interface Serial 0/0/0
  - 10.3.0.0 available through interface Serial 0/0/1
- **R3:**
  - 10.3.0.0 available through interface Serial 0/0/0
  - 10.4.0.0 available through interface FastEthernet 0/0

Initial Exchange of Routing Information

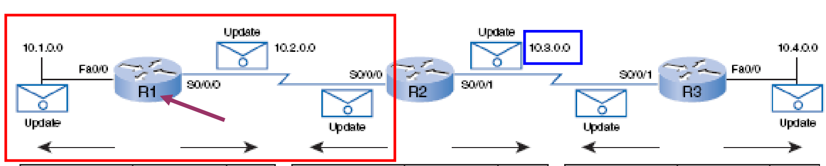

Network	Interface	Hop
10.1.0.0	Fa0/0	0
10.2.0.0	S0/0/0	0

Network	Interface	Hop
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/1	0

Network	Interface	Hop
10.3.0.0	S0/0/1	0
10.4.0.0	Fa0/0	0

sends receives

- **R1:**
  - **Sends** an update about network 10.1.0.0 out the Serial 0/0/0 interface with a metric of 1
  - **Sends** an update about network 10.2.0.0 out the FastEthernet 0/0 interface with a metric of 1
  - **Receives** an update from R2 about network 10.3.0.0 on Serial 0/0/0 with a metric of 1
  - **Stores** network 10.3.0.0 in the routing table with a metric of 1

Initial Exchange of Routing Information

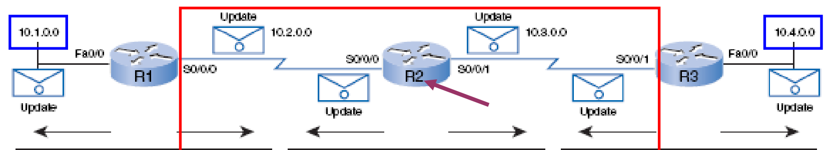

Network	Interface	Hop
10.1.0.0	Fa0/0	0
10.2.0.0	S0/0/0	0

Network	Interface	Hop
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/1	0

Network	Interface	Hop
10.3.0.0	S0/0/1	0
10.4.0.0	Fa0/0	0

receives sends receives

- **R2:**
  - **Sends** an update about network 10.3.0.0 out the S0/0/0 interface with a metric of 1
  - **Sends** an update about network 10.2.0.0 out the S0/0/1 interface with a metric of 1
  - **Receives** an update from R1 about network 10.1.0.0 on S0/0/0 with a metric of 1
  - **Stores** network 10.1.0.0 in the routing table with a metric of 1
  - **Receives** an update from R3 about network 10.4.0.0 on S0/0/1 with a metric of 1
  - **Stores** network 10.4.0.0 in the routing table with a metric of 1

Initial Exchange of Routing Information

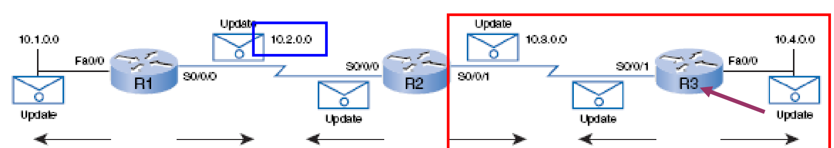

Network	Interface	Hop
10.1.0.0	Fa0/0	0
10.2.0.0	S0/0/0	0

Network	Interface	Hop
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/1	0

Network	Interface	Hop
10.3.0.0	S0/0/1	0
10.4.0.0	Fa0/0	0

receives sends

- **R3:**
  - **Sends** an update about network 10.4.0.0 out the S0/0/1 interface with a metric of 1
  - **Sends** an update about network 10.3.0.0 out the Fa0/0 interface with a metric of 1
  - **Receives** an update from R2 about network 10.2.0.0 on S0/0/1 with a metric of 1
  - **Stores** network 10.2.0.0 in the routing table with a metric of 1

Initial Exchange of Routing Information

Network	Interface	Hop
10.1.0.0	Fa0/0	0
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/0	1

Network	Interface	Hop
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/1	0
10.1.0.0	S0/0/0	1
10.4.0.0	S0/0/1	1

Network	Interface	Hop
10.3.0.0	S0/0/1	0
10.4.0.0	Fa0/0	0
10.2.0.0	S0/0/1	1

- First round of update exchanges:
  - Each router knows about the networks connected to its direct neighbors.
  - R1 does not yet know about 10.4.0.0
  - R3 does not yet know about 10.1.0.0.
- Another exchange of routing information is needed.

Next Exchange of Routing Information

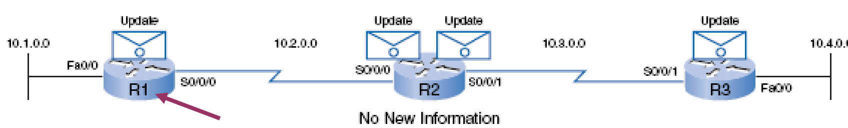

Network	Interface	Hop
10.1.0.0	Fa0/0	0
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/0	1
10.4.0.0	S0/0/0	2

Network	Interface	Hop
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/1	0
10.1.0.0	S0/0/0	1
10.4.0.0	S0/0/1	1

Network	Interface	Hop
10.3.0.0	S0/0/1	0
10.4.0.0	Fa0/0	0
10.2.0.0	S0/0/1	1
10.1.0.0	S0/0/0	2

sends

R1 and R3 now have complete routing tables.

receives

- **R1:**
  - **Sends** an update about network 10.1.0.0 out the S0/0/0 interface with a metric of 1.
  - **Sends** an update about networks 10.2.0.0 with a metric of 1 and 10.3.0.0 with a metric of 2 out the Fa0/0 interface.
  - **Receives** an update from R2 about network **10.4.0.0 on S0/0/0 with a metric of 2.**
  - **Stores** network 10.4.0.0 in the routing table with a metric of 2.
  - **Same update from R2** contains information about network 10.3.0.0 on S0/0/0 with a metric of 1. There is no change; therefore, the routing information remains the same.

Next Exchange of Routing Information

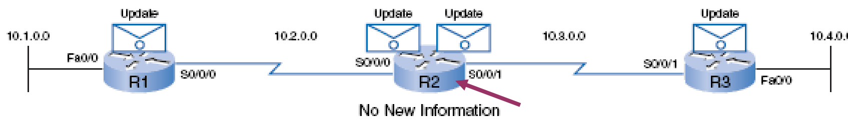

Network	Interface	Hop
10.1.0.0	Fa0/0	0
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/0	1
10.4.0.0	S0/0/0	2

Network	Interface	Hop
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/1	0
10.1.0.0	S0/0/0	1
10.4.0.0	S0/0/1	1

Network	Interface	Hop
10.3.0.0	S0/0/1	0
10.4.0.0	Fa0/0	0
10.2.0.0	S0/0/1	1
10.1.0.0	S0/0/0	2

receives — sends — receives

R1 and R3 now have complete routing tables.

- **R2:**
  - **Sends** an update about networks 10.3.0.0 with a metric of 1 and **10.4.0.0 with a metric of 2 out the S0/0/0 interface**.
  - **Sends** an update about networks **10.1.0.0 with a metric of 2 (new)** and 10.2.0.0 with a metric of 1 out the S0/0/1 interface.
  - **Receives** an update from R1 about network 10.1.0.0 on S0/0/0. There is no change; therefore, the routing information remains the same.
  - **Receives** an update from R3 about network 10.4.0.0 on S0/0/1. There is no change; therefore, the routing information remains the same.

Next Exchange of Routing Information

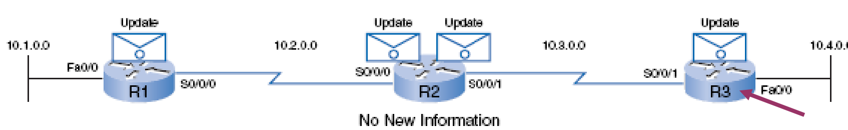

Network	Interface	Hop
10.1.0.0	Fa0/0	0
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/0	1
10.4.0.0	S0/0/0	2

Network	Interface	Hop
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/1	0
10.1.0.0	S0/0/0	1
10.4.0.0	S0/0/1	1

Network	Interface	Hop
10.3.0.0	S0/0/1	0
10.4.0.0	Fa0/0	0
10.2.0.0	S0/0/1	1
10.1.0.0	S0/0/0	2

receives — sends

R1 and R3 now have complete routing tables.

- **R3:**
  - **Sends** an update about network 10.4.0.0 out the S0/0/1 interface.
  - **Sends** an update about networks 10.2.0.0 with a metric of 2 and 10.3.0.0 with a metric of 1 out the Fa0/0 interface.
  - **Receives** an update from R2 about network **10.1.0.0 with a metric of 2.**
  - **Stores** network 10.1.0.0 in the routing table with a metric of 2.
  - **Same update from R2** contains information about network 10.2.0.0 on S0/0/1 with a metric of 1. There is no change; therefore, the routing information remains the same.

Problems with classic IGP

- slow convergence
- count to infinity
- no mask information
  - no CIDR
  - no VLSM
  - no subnet 0

Distance Vector Protocols

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

- **Slow convergence**

Count to infinity problem

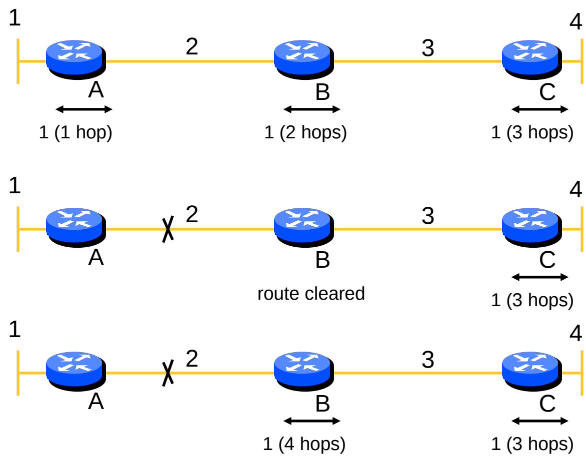

Count to infinity: split-horizon

- Distance Vector Protocols suffer from **Count to Infinity** problem.
- Solution: **Split-Horizon**

- Don't feed selected route back to source
  - no feedback on source interface
  - no feedback to source neighbor

Count to infinity: split-horizon

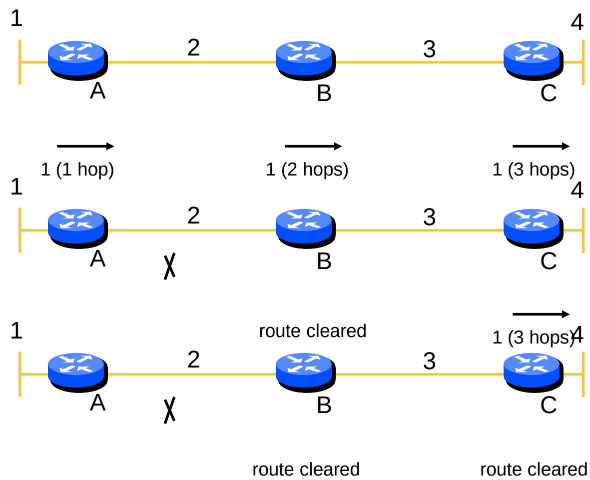

Split Horizon

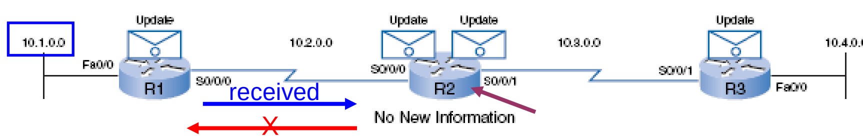

Network	Interface	Hop
10.1.0.0	Fa0/0	0
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/0	1
10.4.0.0	S0/0/0	2

Network	Interface	Hop
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/1	0
10.1.0.0	S0/0/0	1
10.4.0.0	S0/0/1	1

Network	Interface	Hop
10.3.0.0	S0/0/1	0
10.4.0.0	Fa0/0	0
10.2.0.0	S0/0/1	1
10.1.0.0	S0/0/0	2

R1 and R3 now have complete routing tables.

- Distance vector routing protocols typically apply the **split horizon**.
  - **Information is not sent to the interface from which it was received.**
- Example
  - R2 would not send an update out S0/0/0 containing the network 10.1.0.0 because R2 learned about that network through S0/0/0.

RIP

RIP: Routing Information Protocol

- Most popular Routing Protocol.
- Interior Gateway Protocol.
- Distance Vector Protocol.
- Only metric is number of hops.
- Maximum number of hops is 15.
- Updates every 30 seconds.
- Doesn’t always select fastest path.
- Uses UDP, port 520.

RIP Versions

RIPv1

- No network mask.

RIPv2

- Classless (CIDR)

RIPng

- Supports IPv6
- UDP port 521, multicast group FF02::9.

IGRP

IGRP: Interior Gateway Routing Protocol

- Cisco proprietary, free use.
- Interior Gateway Protocol.
- Distance Vector Protocol.
- Only IP
- Metric can be composed from:
  - hops,
  - bandwidth,
  - load,
  - delay,
  - reliability.
- Maximum number of hops is 255.
- Updates every 90 seconds.

Link-State IGP

- Interior Gateway Protocols
- Using Link-State Information sharing
- Examples:
  - OSPF: Open Shortest Path First
  - IS–IS: Intermediate System to Intermediate System

Link state protocols

- information about adjacencies sent to **all** routers
- each router builds a topology database
- a "shortest path" algorithm (*Dijkstra*) is used to find best route
- converge as quickly as databases can be updated

Link state protocols

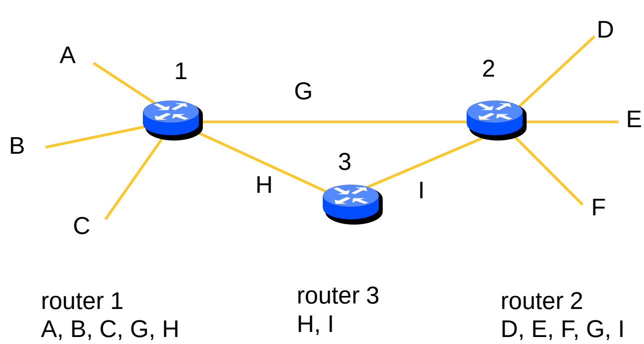

A - 1 - G - 2 - D

OSPF

OSPF: Open Shortest Path First

- Interior Gateway Protocol.
- Link State Protocol.
- IP only
- Fast convergence
- Metric can be composed from:
  - speed,
  - traffic,
  - reliability,
  - security.
- Event-triggered updates.

OSPF operation

- routers send 'hello' packets to neighbors.
- used to maintain a list of active links.
- hello packets are small easy to process packets
- hello packets are sent periodically (usually short interval)

OSPF Operation

- topology information is packaged in a "link state announcement"
- announcements are broadcasted
  - During initialization
  - When there's a change
  - Periodically (every 45mins)

OSPF Operation

- Each router maintains an own copy of the **Link State Database (LSDB)**
- Used to store link states
  - Found during the init phase (Hello packets)
  - Announcements collected from other routers.
- When a change occurs
  - run SPF (Dijkstra's algorithm)
  - install result into routing table

OSPF scalling

- Each link transition causes a broadcast.
- Each update causes a SPF run.
  - Hi processing requirement

- OSPF can group routers into areas
- Each area appears as a single router.

OSPF Areas

Without areas,
Every router must known about all links.

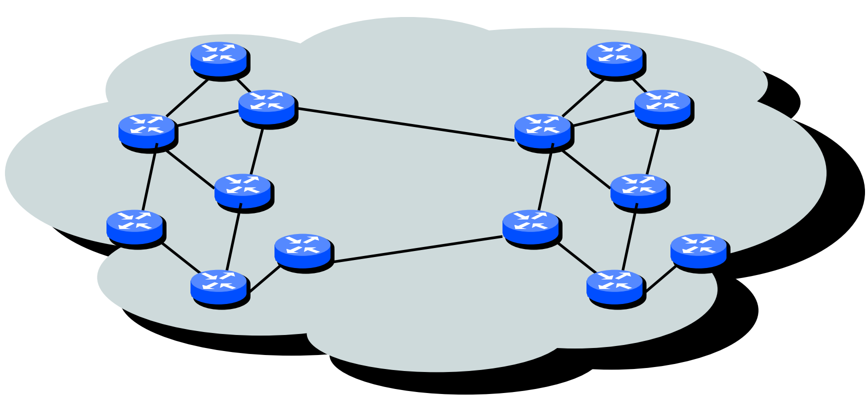

OSPF Areas

Using Areas

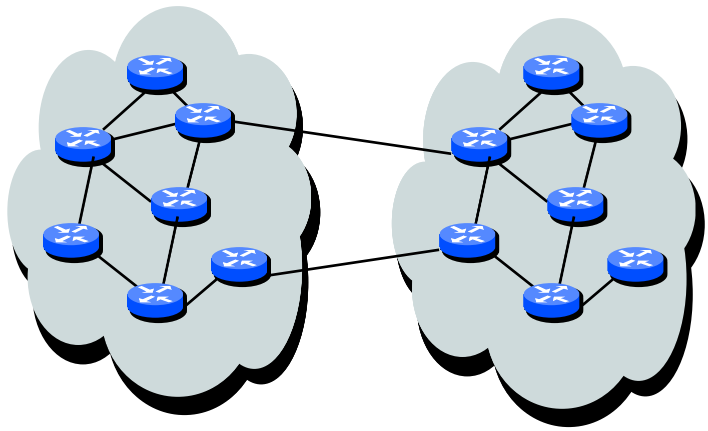

OSPF areas

- rule of thumb:
  no more than 150 routers/area
- reality:
  no more than 500 routers/area
- backbone "area" is an area
- proper use of areas reduce
  - Bandwidth
  - CPU utilization

IS-IS

IS-IS: Intermediate System to Intermediate System

- Digital Equipment Corporation
- Developed to be used in DECnet
- OSI, ISO/IEC 10589:2002
- IETF, RFC 7142
- Extended to IP as: **I-ISIS**

I-ISIS

I-ISIS: **Integrated IS-IS**

- Interior Gateway Protocol
- multi-protocol (CLNP, IP, IPX, ...)
- link state protocol
- fast convergence
- design and architecture moderately complex
- configuration may be simple

EIGRP

EIGRP: Enhanced IGRP

- Interior Gateway Protocol.
- multi-protocol (IP, IPX, Appletalk)
- hybrid (dual) protocol
  - based on distance vector algorithm
  - fast convergence (like OSPF)
- IGRP metric
- allows load sharing across diverse links
- can be bandwidth intensive on slow links

EIGRP

- Design goals
  - as fast as OSPF & IS-IS
  - easy migration from IGRP
  - trivial to configure
- Results
  - works best on high speed links
  - doesn't scale well in high-meshed networks
  - star networks OK

Distance Vector vs Link State

	Distance Vector	Link State
Shared Information	Whole routing table	Local links
Sharing targets	Neighbords	All routers
Convergence Time	Long	Short
Processing	Light	Heavy
Algorithm	Bellamn Ford	Dijkstra (on each node)
Examples	RIP, IGRP	OSPF, ISIS

Exterior Routing

egp

egp: Exterior Gateway Protocol

- Examples: EGP, BGP, IDRP
  EGP $\in$ egp
- Primary goal is to provide reachability information outside administrative domain.
- Secondary goal is administrative control.
- Metrics may be arbitrary or weak.

Exterior Routing

can be done by **Static Routes**

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

Multiple IGPs with route leaking

- Run an instance of an IGP at each site for local routing
- Run a backbone IGP at each border router.
- redistribute local IGP into backbone IGP
- redistribute backbone IGP into local IGP (or default)
- backbone routers share common administration

Multiple IGPs with route leaking

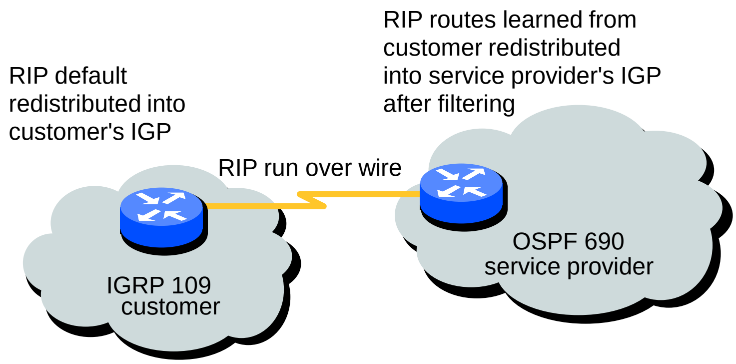

EGP

EGP: Exterior Gateway Protocol

- An example of an **egp**
- historical protocol
- obsolete
- assumes a central core
- no transit service except via core
- periodic updates of topologic information
- no metrics

IDRP

IDRP: Inter-Domain Routing Protocol

- ISO 10747
- Almost identical to BGP-4
- Multi-protocol
  - IP
  - CLNP
  - IPX

BGP-4

BGP: Border Gateway Protocol version 4

- carries external routes only
- uses reliable transport mechanism (TCP)
- not a periodic routing protocol
- allows limited policy selection
- AS path insures loop free routing
- "best path" determined at AS granularity

AS-Path

- Sequence of AS a route has traversed
- Loop detection
- Apply policy

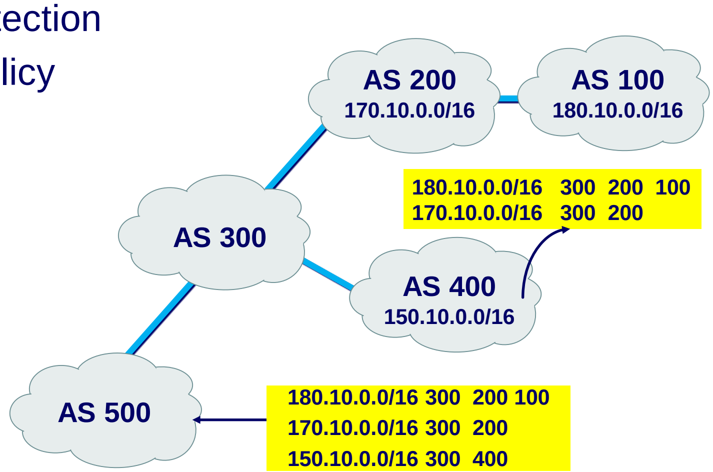

iBGP

iBGP: Internal BGP

- Same protocol as BGP;
- Used when different BGP peers share same AS_Number
- iBGP nodes are fully meshed;
- No re-advertisement of route updates to prevent looping;

iBGP / eBGP

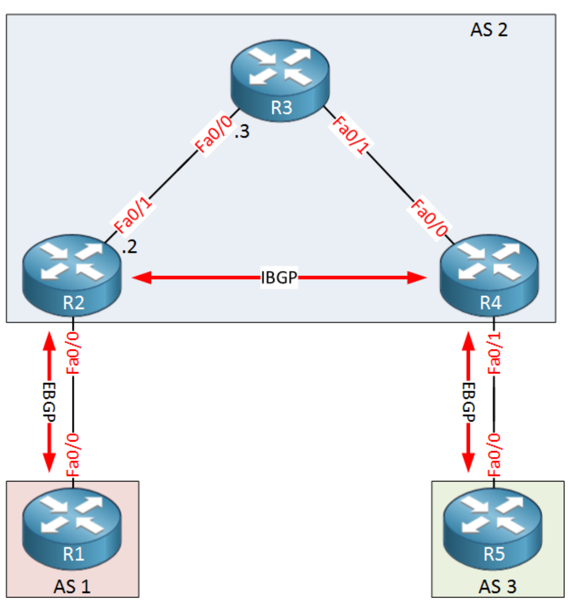

isep Instituto Superior de Engenharia do Porto

BGP - Local Preference

- Specifies preference of links
- Local to an AS
- Used to influence BGP path selection
- Path with highest local preference wins

isep Instituto Superior de Engenharia do Porto

Local Preference

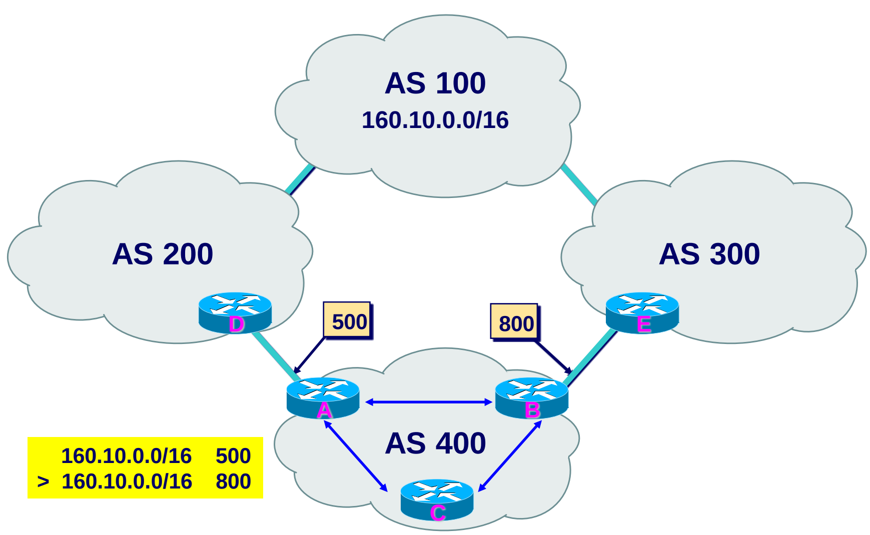

isep Instituto Superior de Engenharia do Porto

BGP route selection

- do not consider path if no next hop route
- largest weight
  - local to router
- highest local preference
  - global within AS
- shortest AS path

isep Instituto Superior de Engenharia do Porto

Technical information on BGP

- RFC-1772
  - application of the Border Gateway Protocol
- RFC-1771
  - BGP-4 protocol reference document
- RFC-1745
  - BGP <-> OSPF interaction