ECE/COMPSCI 356 Computer Network Architecture

Lecture 15: Border Gateway Protocol

Neil Gong

neil.gong@duke.edu

Slides credit: Xiaowei Yang, PD

Today

- Border Gateway Protocol (BGP)
- IP tunnels

The Internet

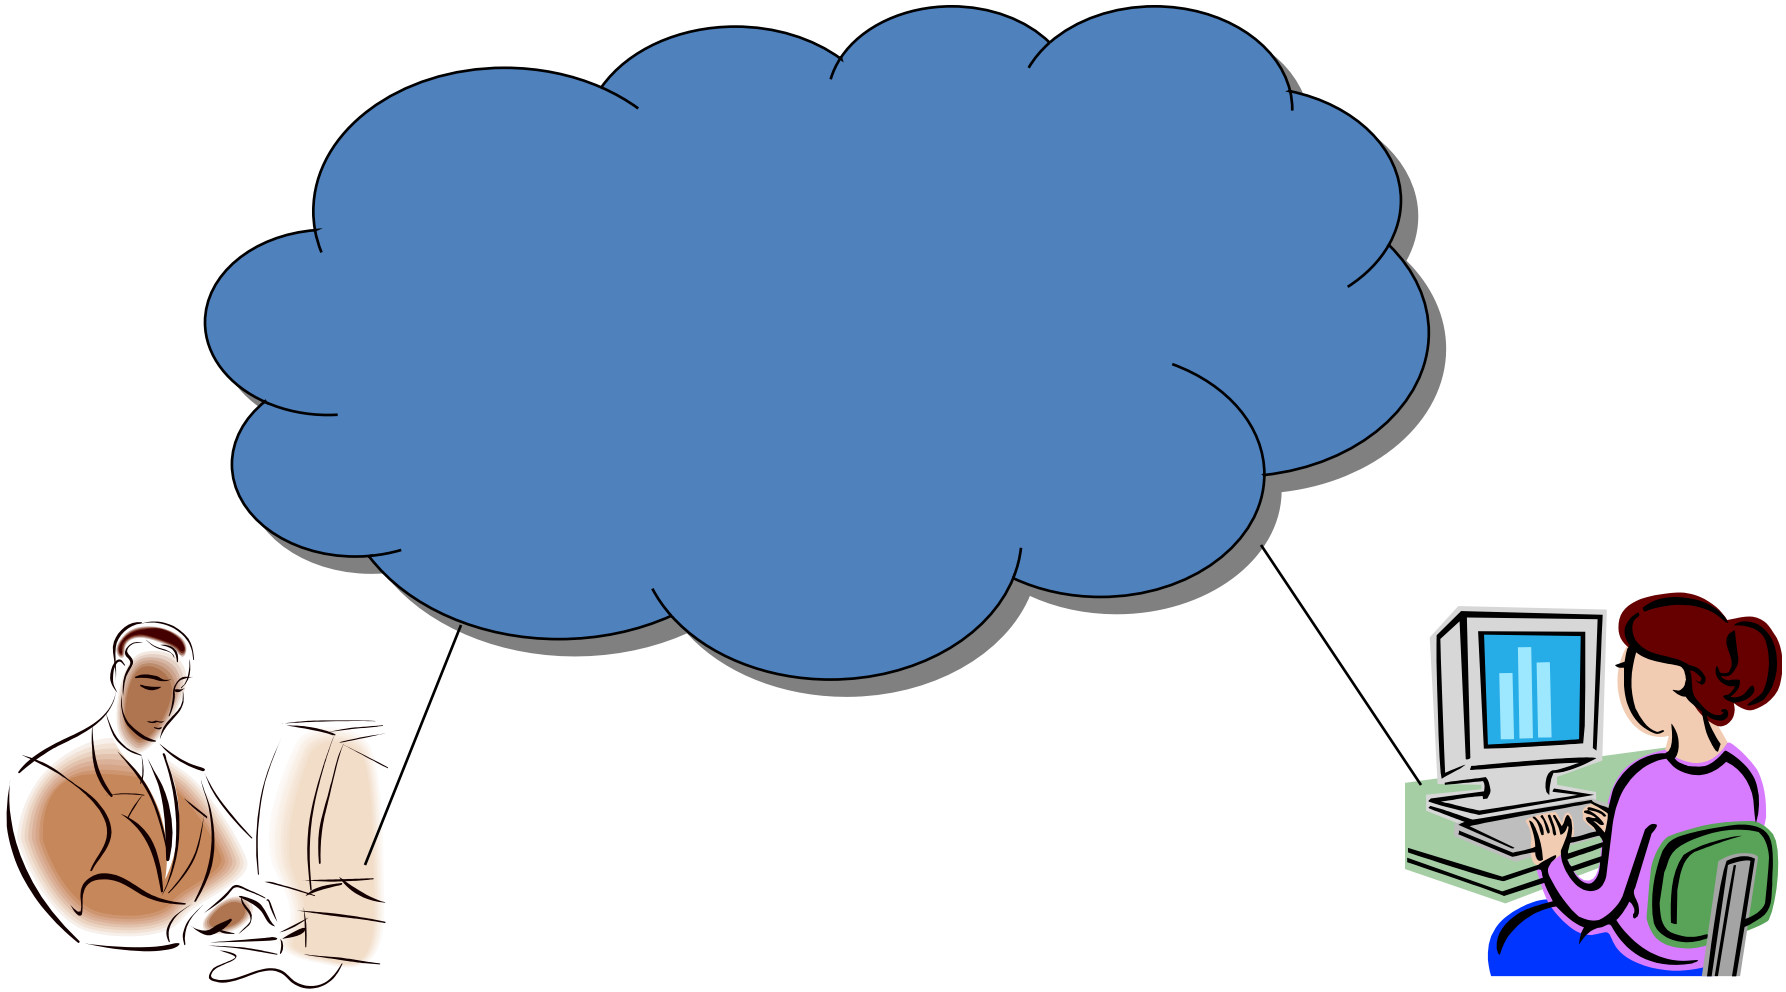

The Internet: Zooming In 2x

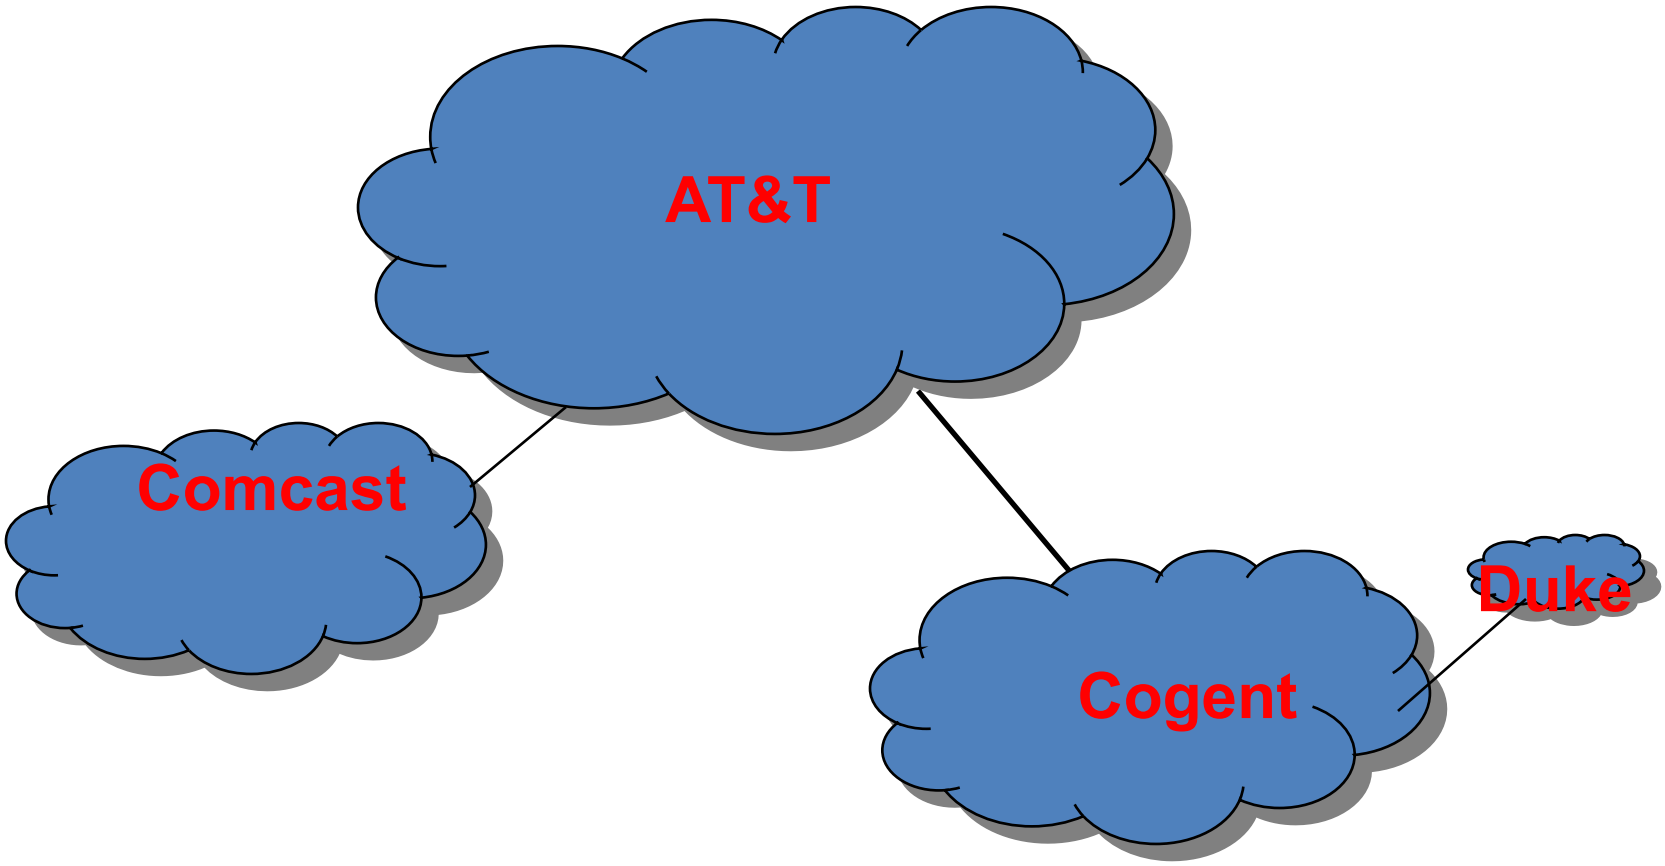

Intra-domain vs. inter-domain routing

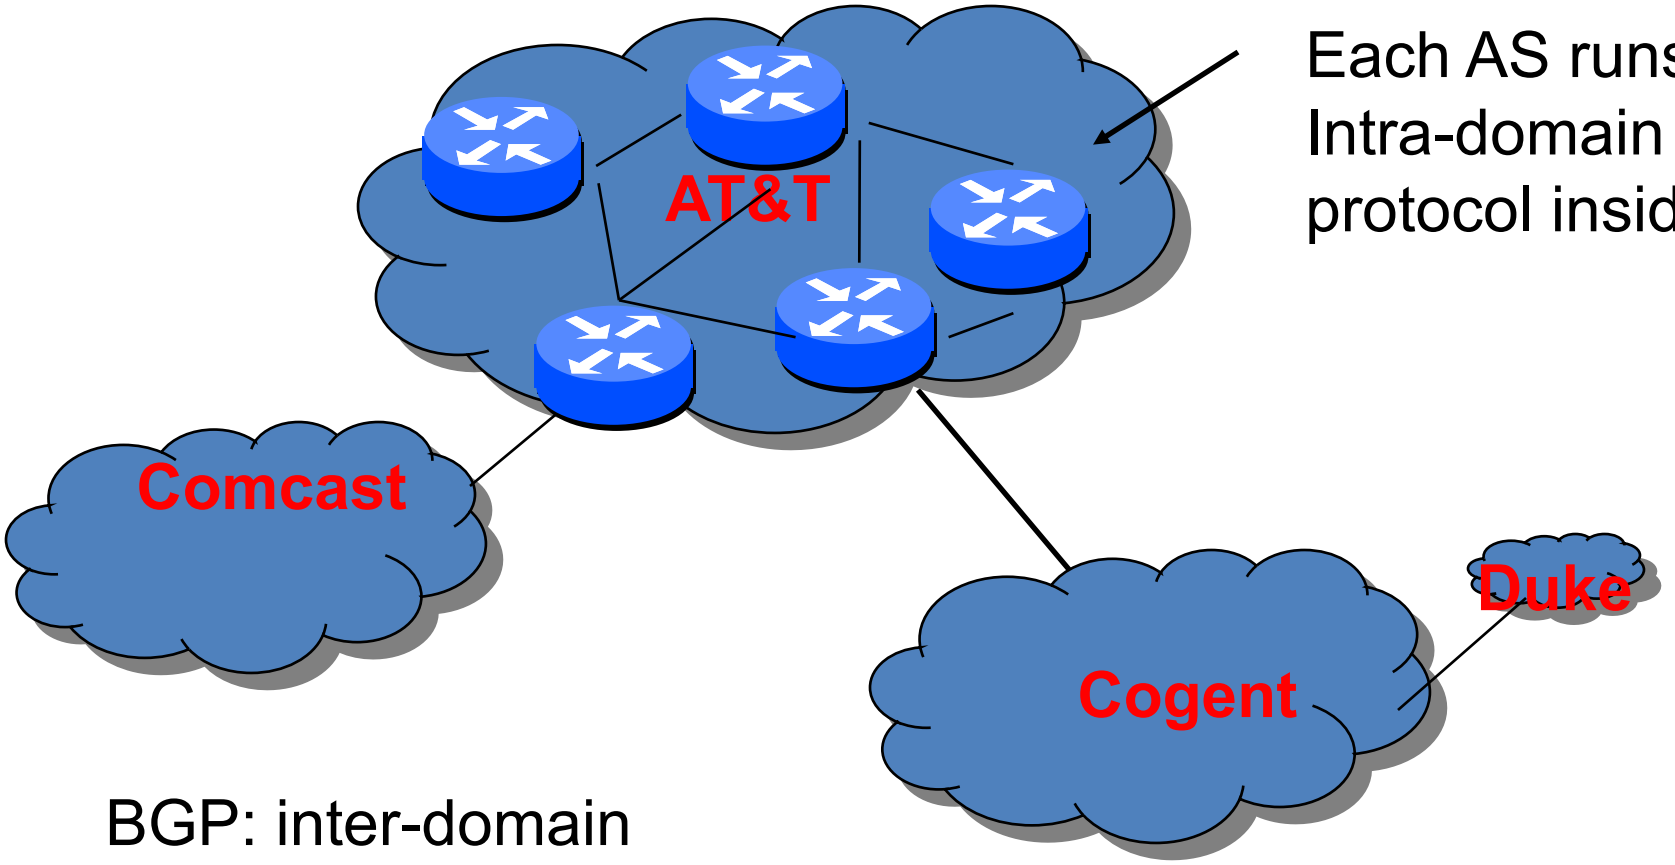

Each AS runs an Intra-domain routing protocol inside

BGP: inter-domain

The Global Internet

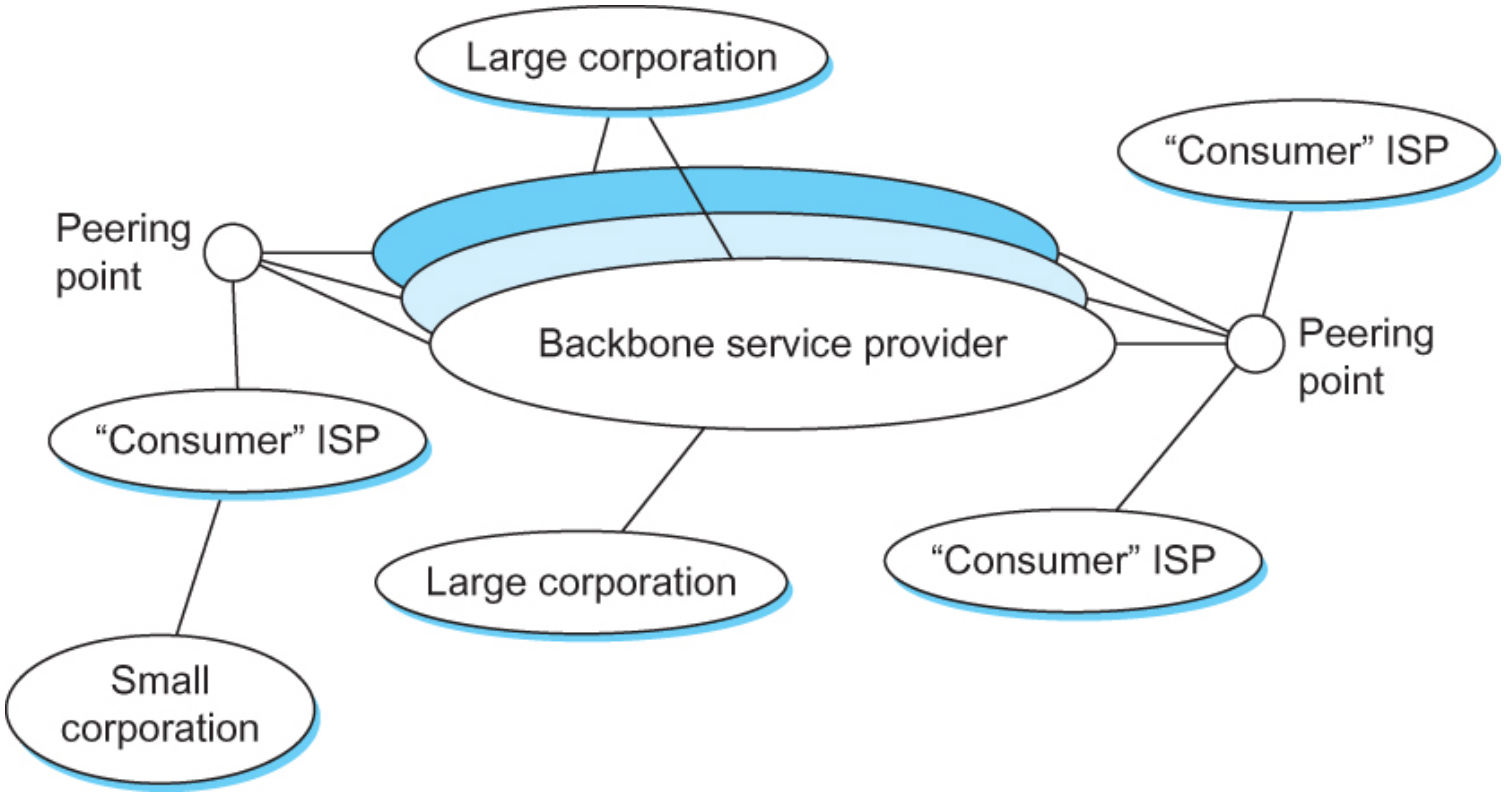

A simple multi-provider Internet

Classify AS

- Define *local traffic* as traffic that originates at or terminates on nodes within an AS, and *transit traffic* as traffic that passes through an AS.

- We can classify AS's into three types:
  - *Stub AS*: an AS that has only a single connection to one other AS; such an AS will only carry local traffic (*small corporation in the figure of the previous page*).
  - *Multihomed AS*: an AS that has connections to more than one other AS, but refuses to carry transit traffic (*large corporation at the top in the figure of the previous page*).
  - *Transit AS*: an AS that has connections to more than one other AS, and is designed to carry both transit and local traffic (*backbone providers in the figure of the previous page*).

Interdomain Routing

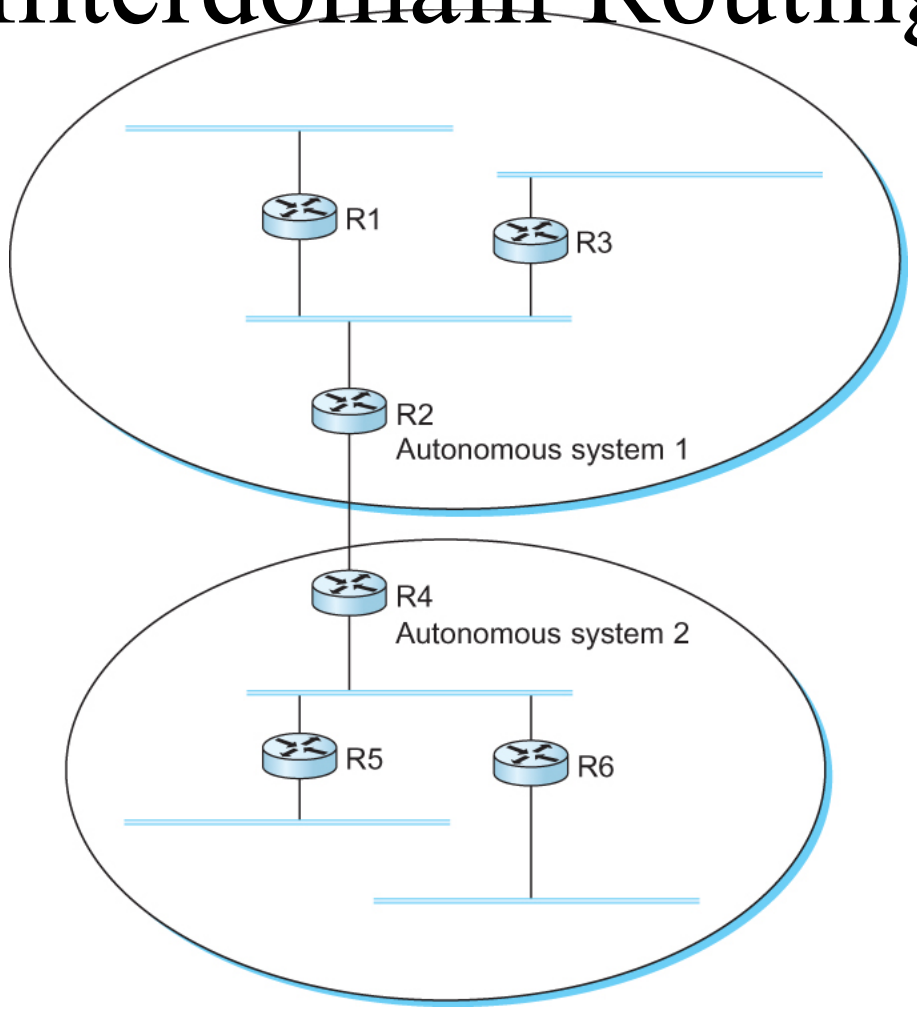

A network with two autonomous system

Design goals of BGP

- Applicable to arbitrarily interconnected set of ASs
- Find any path to the intended destination that is loop free
  - We are concerned with reachability than optimality
  - Finding path anywhere close to optimal is considered to be a great achievement

Reachability instead of optimality

- Autonomous nature of the domains
  - It is impossible to calculate meaningful path costs for a path that crosses multiple ASs
  - A cost of 1000 across one provider might imply a great path but it might mean an unacceptable bad one from another provider
  - Each AS can have its own policies
- Issues of trust
  - Provider A might be unwilling to believe certain advertisements from provider B

BGP

Each AS has:

- One BGP *speaker* that advertises:
  - local networks
  - other reachable networks
  - gives *path* information
- In addition to the BGP speakers, the AS has one or more border routers which need not be the same as the speakers
- packets enter and leave the AS through border routers

Path vector algorithm in BGP

- BGP does not belong to either of the two main classes of routing protocols (distance vectors and link-state protocols)

- BGP advertises *complete paths* as an enumerated lists of ASs to reach a particular network

Path Vector

- Records what ASes a route goes through
- Loop avoidance: Immediately discard
- Shortest path heuristics

I can reach d via B,D

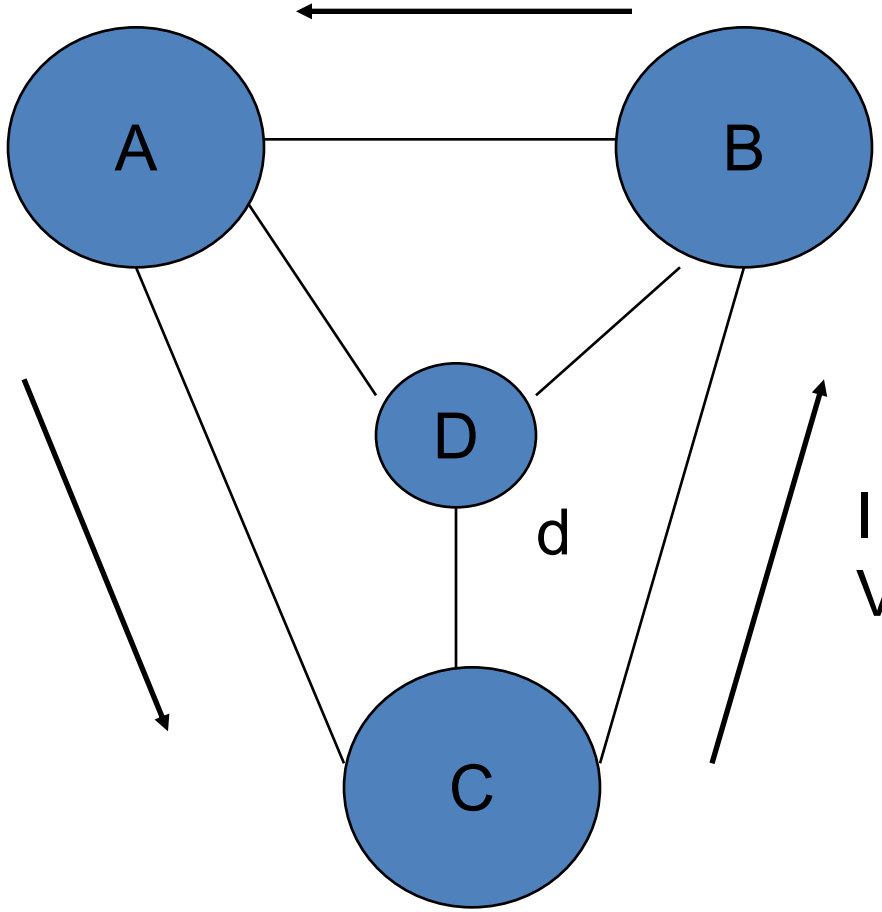

I can reach d
Via A,B,D

I can reach d
Via C,A,B,D

More Example

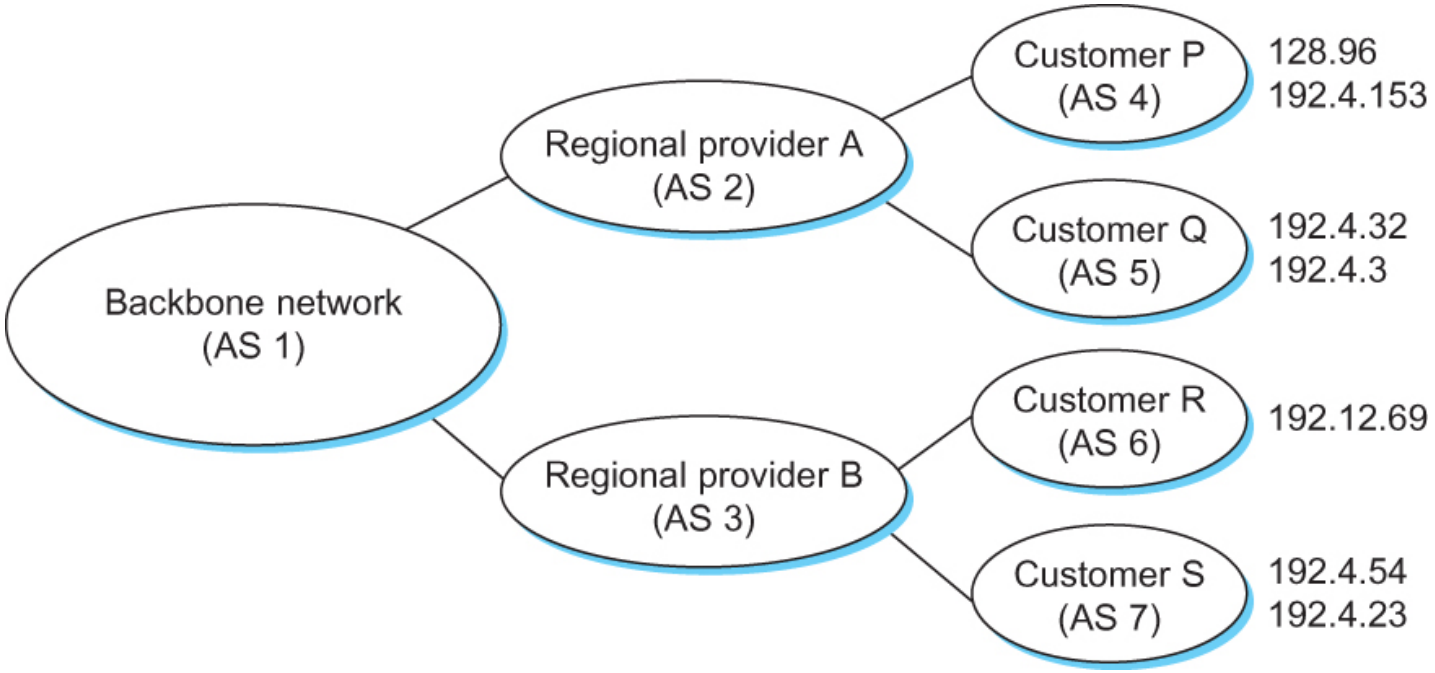

Example of a network running BGP

More Example

- Speaker for AS 2 advertises reachability to P and Q
  - Network 128.96, 192.4.153, 192.4.32, and 192.4.3, can be reached directly from AS 2.
- Speaker for backbone network then advertises
  - Networks 128.96, 192.4.153, 192.4.32, and 192.4.3 can be reached along the path <AS 1, AS 2>.
- Speaker can also cancel previously advertised paths

BGP Issues

- It should be apparent that the AS numbers carried in BGP need to be unique
- For example, AS 2 can only recognize itself in the AS path in the example if no other AS identifies itself in the same way
- AS numbers are 16-bit numbers assigned by a central authority

Integrating Interdomain and Intradomain Routing

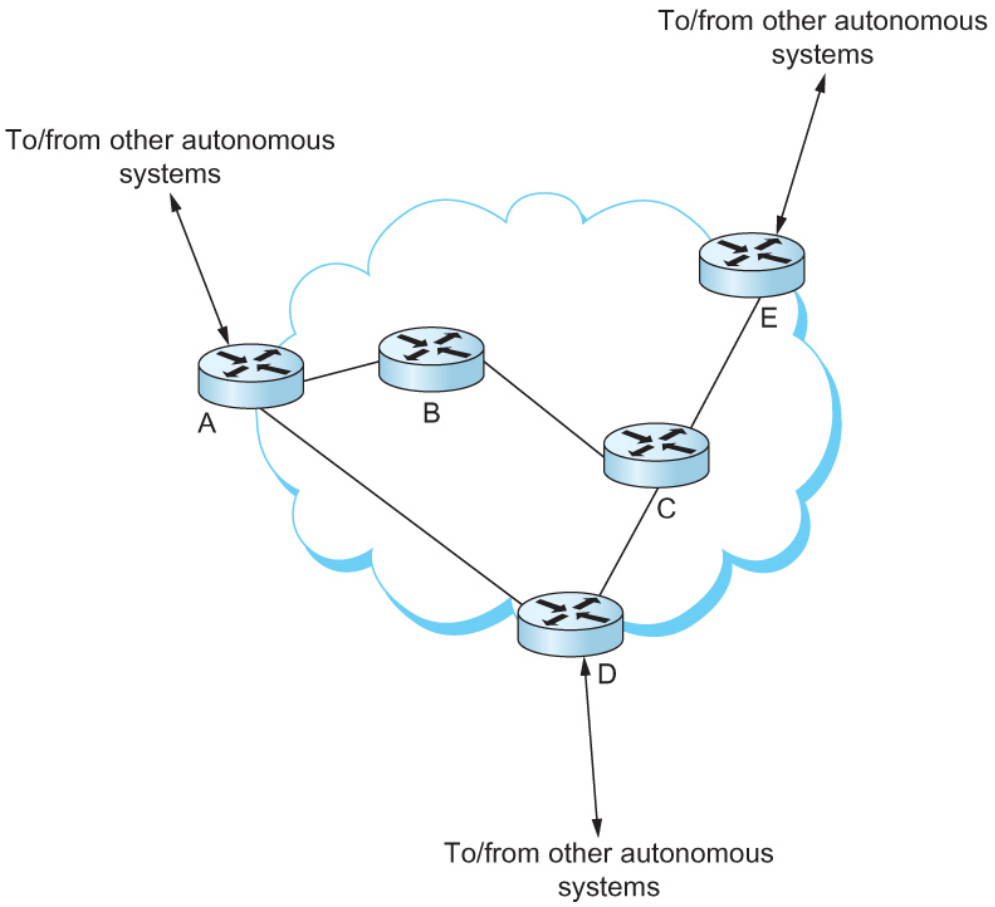

Two components of BGP

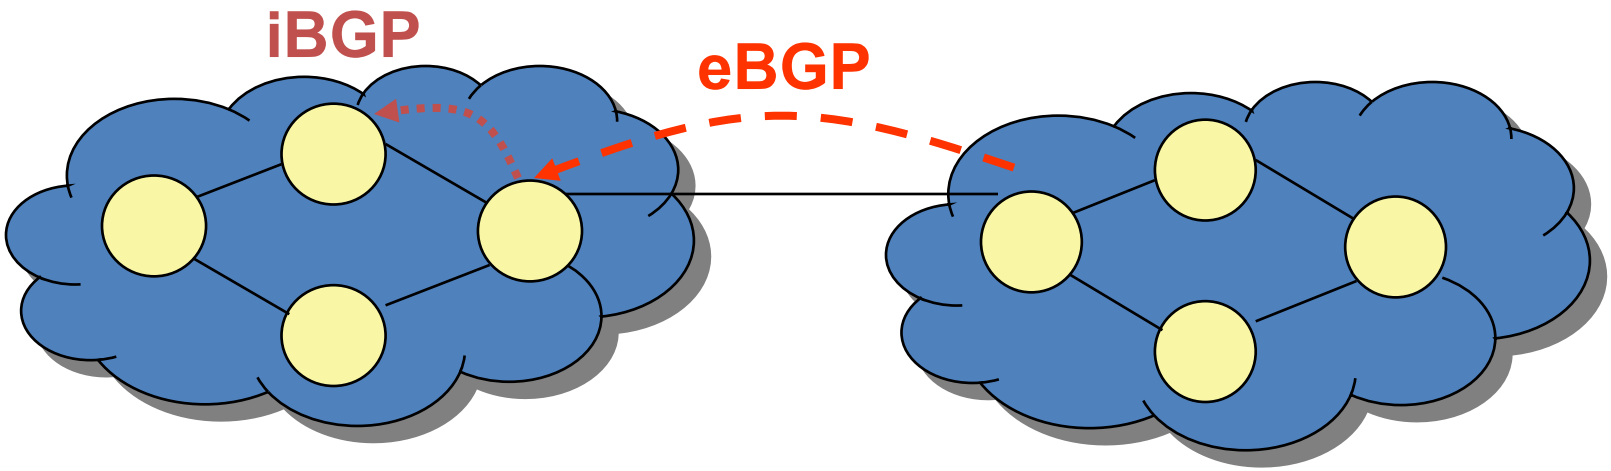

- eBGP : between two routers in different ASes
- iBGP : between internal routers of an AS.

Integrating Interdomain and Intradomain Routing

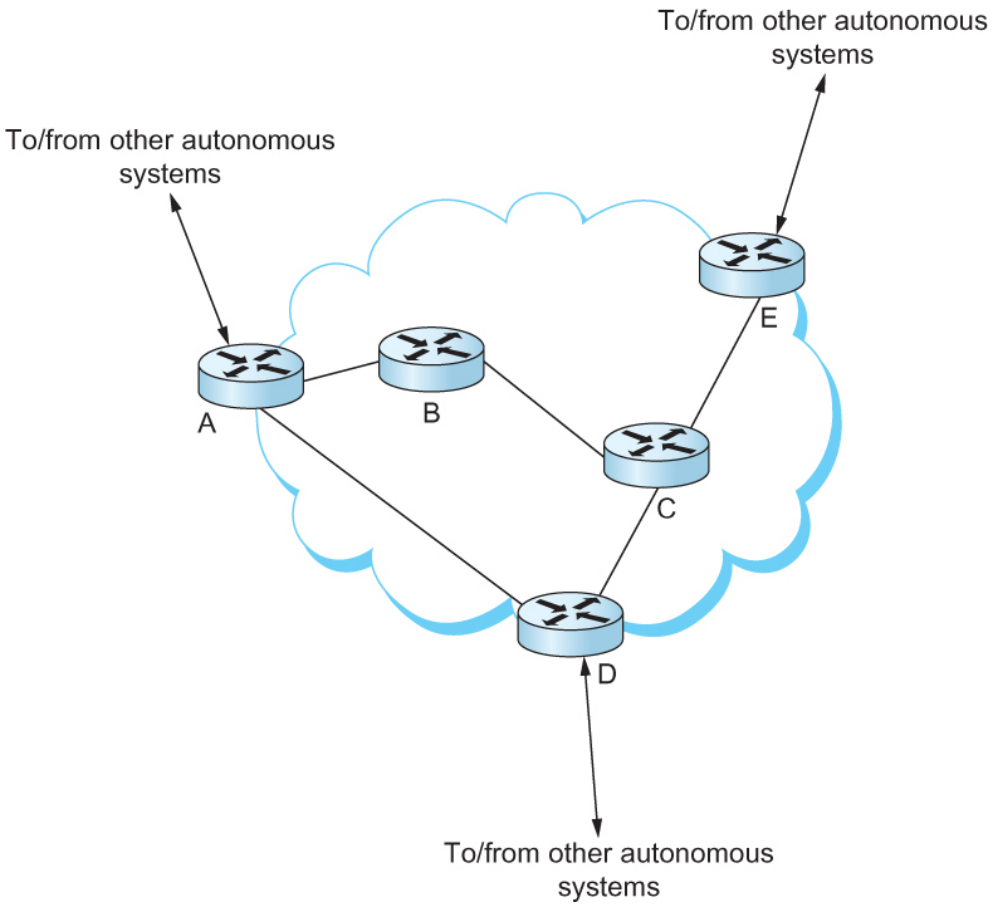

All routers run iBGP and an intradomain routing protocol. Border routers (A, D, E) also run eBGP to other ASs

Integrating Interdomain and Intradomain Routing

Prefix	BGP Next Hop
18.0/16	E
12.5.5/24	A
128.34/16	D
128.69./16	A

BGP table for the AS

Router	IGP Path
A	A
C	C
D	C
E	C

IGP table for router B

Prefix	IGP Path
18.0/16	C
12.5.5/24	A
128.34/16	C
128.69./16	A

Combined table for router B

BGP routing table, IGP (Interior Gateway Protocols) routing table, and combined table at router B

Today

- Border Gateway Protocol (BGP)
- IP tunnels

IP tunnels

- Tunnels
  - A technique used in many scenarios
    - VPN, IPv4-v6 transition, Mobile IP, Multicast, Non-IP forwarding, IPsec

What is a tunnel

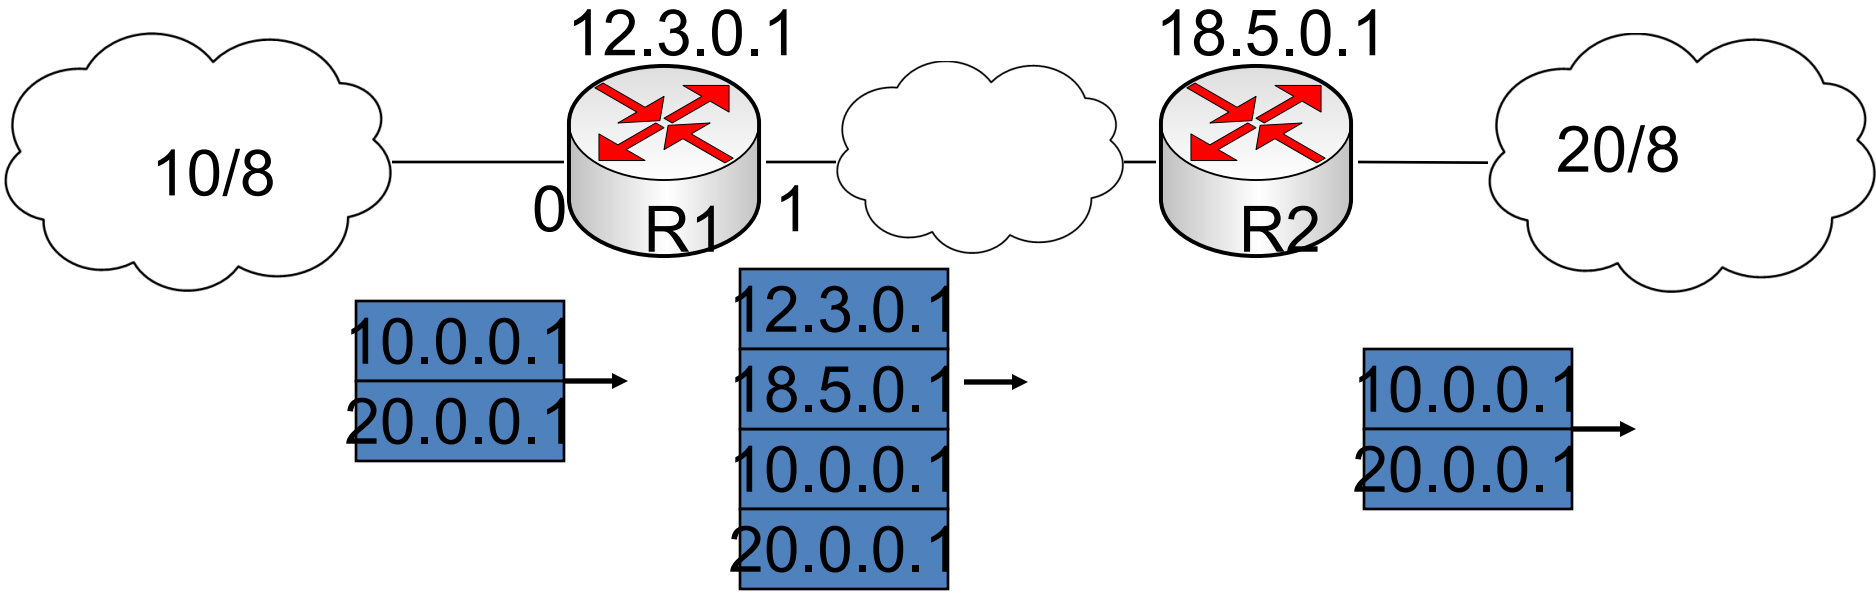

- A “pseudowire”, or a virtual point-to-point link
- The head router encapsulates a packet in an outer header destined to the tail router

Virtual interface

- A router adds a tunnel header for packets sent to a virtual interface

NetworkNum	nextHop
10/8	ether0
20/8	tun0
0/0	ether1

Tunnel applications

- Traversing a region of network with a different addressing format or with insufficient routing knowledge
- Building virtual private networks

FIGURE 3.26 An example of virtual private networks: (a) two separate private networks; (b) two virtual private networks sharing common switches.

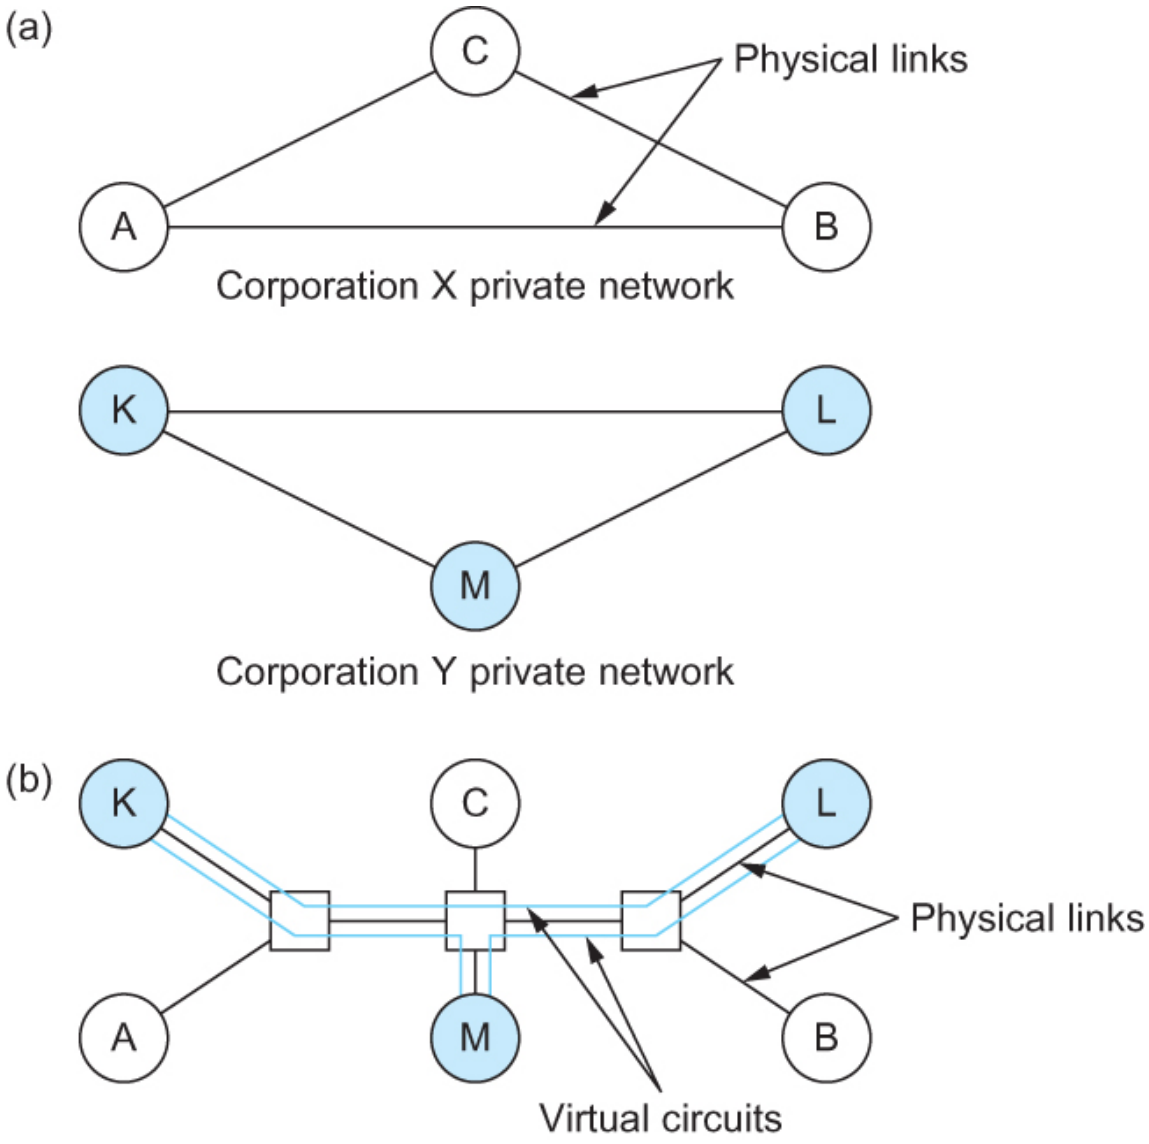

IPv4-v6 transition

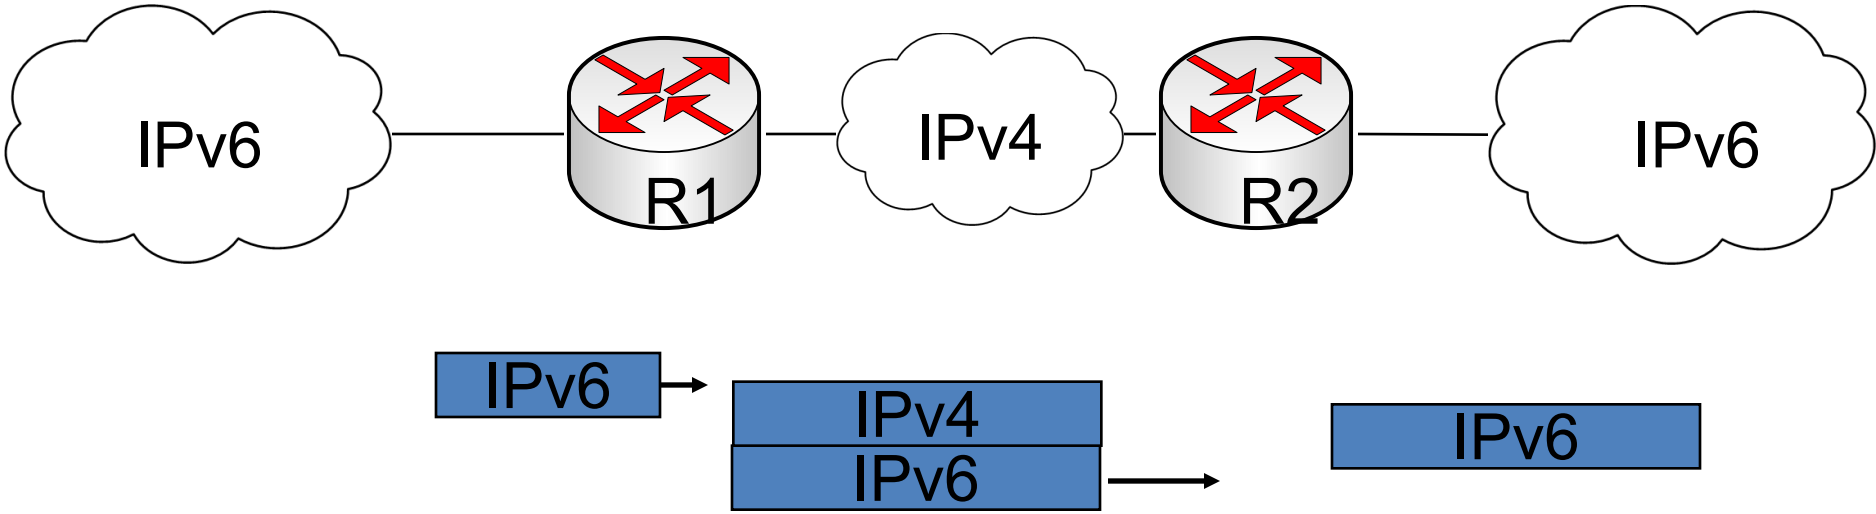

Today

- Border Gateway Protocol (BGP)
- IP tunnels