Border Gateway Protocol

The Global Internet

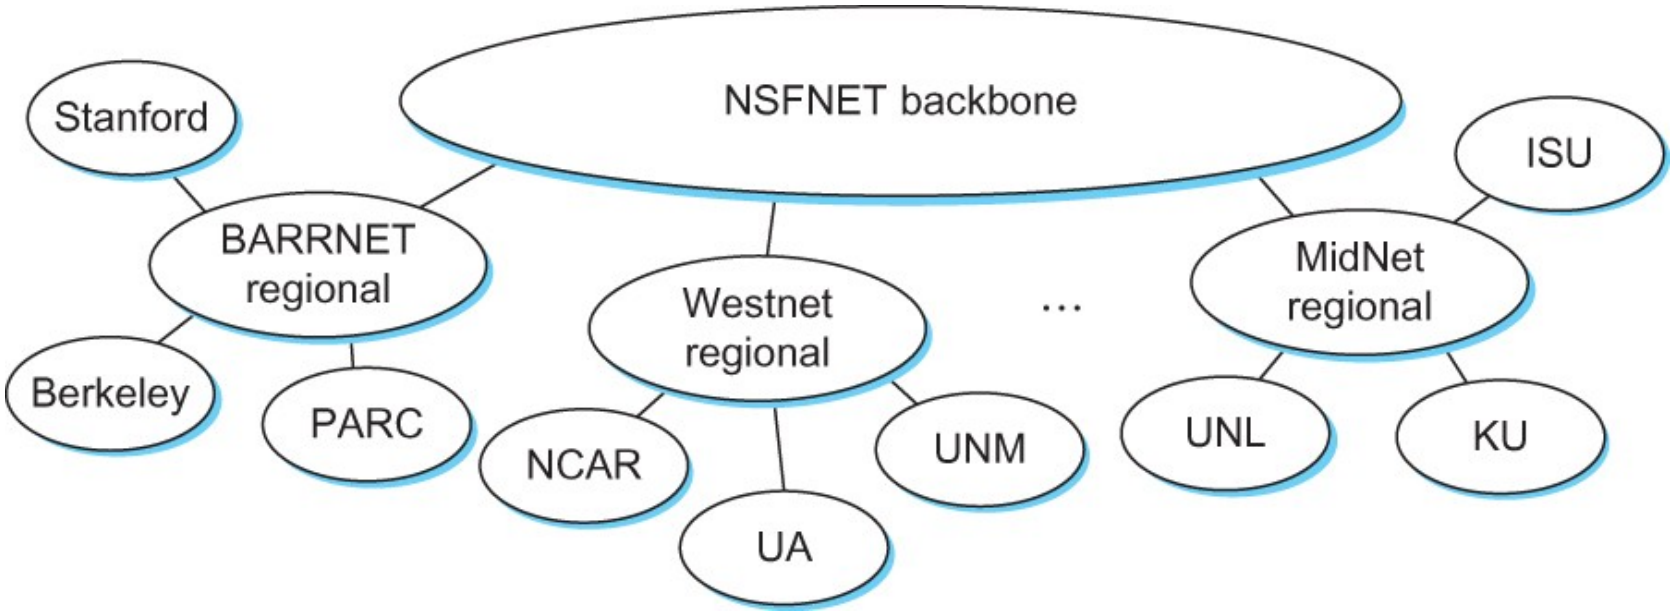

The tree structure of the Internet in 1990

The Global Internet

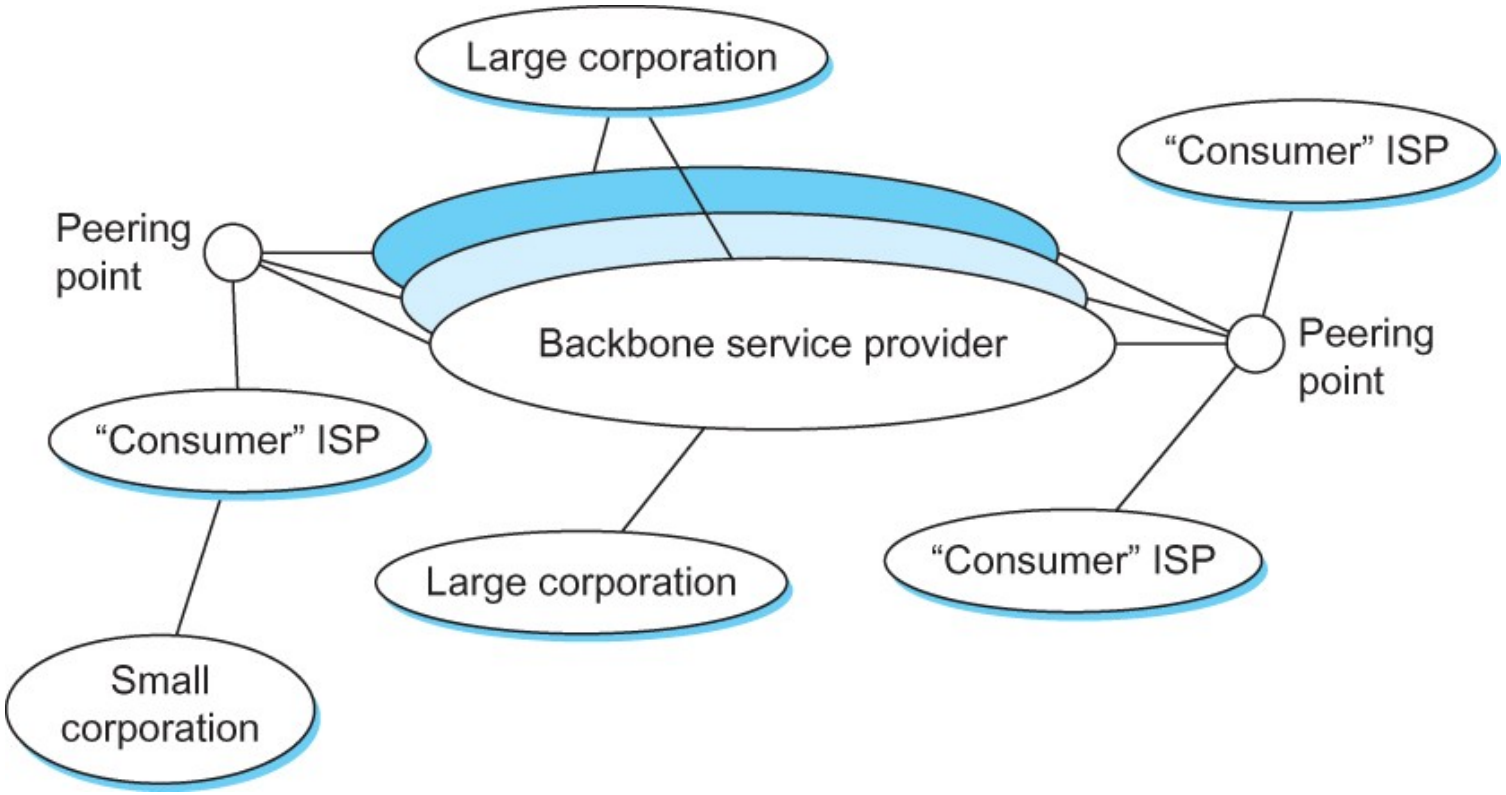

A simple multi-provider Internet

Interdomain Routing (BGP)

- Internet is organized as autonomous systems (AS) each of which is under the control of a single administrative entity

- Autonomous System (AS)
  - corresponds to an administrative domain
  - examples: University, company, backbone network

- A corporation's internal network might be a single AS, as may the network of a single Internet service provider

Interdomain Routing

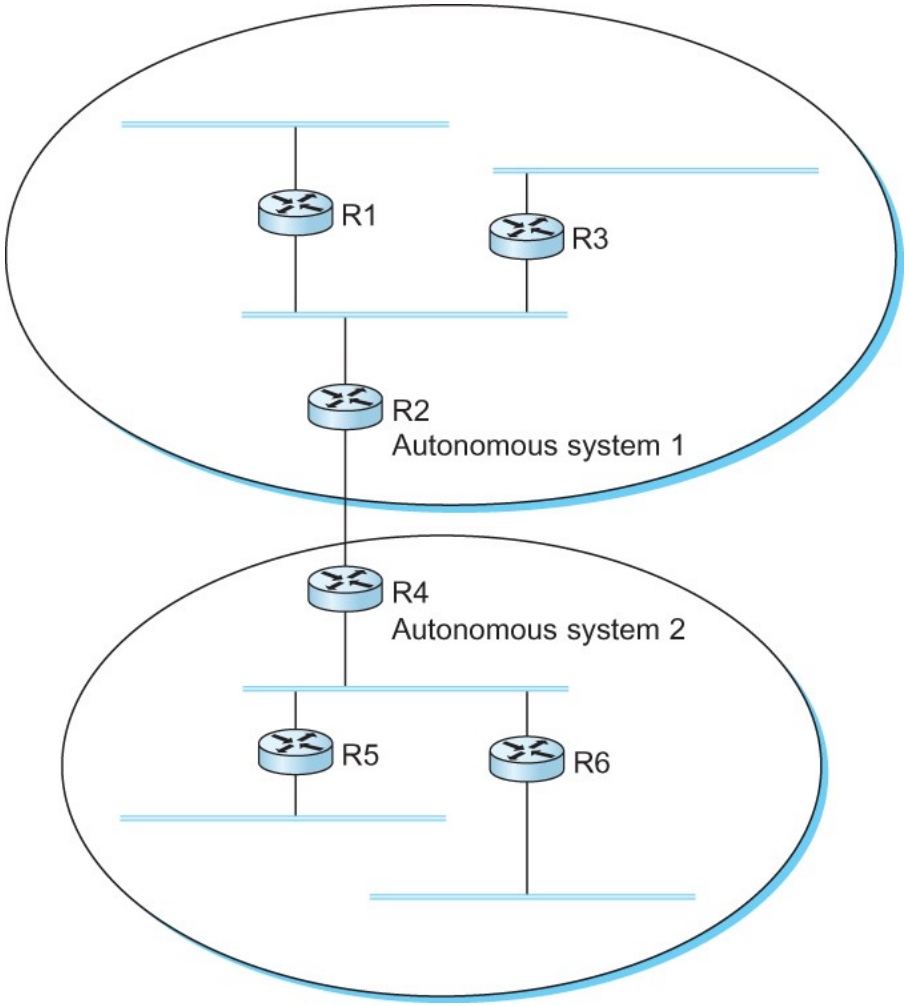

A network with two autonomous system

Route Propagation

- Idea: Provide an additional way to hierarchically aggregate routing information is a large internet.
  - Improves scalability
- Divide the routing problem in two parts:
  - Routing within a single autonomous system
  - Routing between autonomous systems
- Another name for autonomous systems in the Internet is routing domains
  - Two-level route propagation hierarchy
    - Inter-domain routing protocol (Internet-wide standard)
    - Intra-domain routing protocol (each AS selects its own)

EGP and BGP

- Inter-domain Routing Protocols
  - Exterior Gateway Protocol (EGP)
    - Forced a tree-like topology onto the Internet
    - Did not allow for the topology to become general
      - Tree like structure: there is a single backbone and autonomous systems are connected only as parents and children and not as peers
  - Border Gateway Protocol (BGP)
    - Assumes that the Internet is an arbitrarily interconnected set of ASs.
    - Today's Internet consists of an interconnection of multiple backbone networks (they are usually called service provider networks, and they are operated by private companies rather than the government)
    - Sites are connected to each other in arbitrary ways

BGP

- Some large corporations connect directly to one or more of the backbone, while others connect to smaller, non-backbone service providers.
- Many service providers exist mainly to provide service to “consumers” (individuals with PCs in their homes), and these providers must connect to the backbone providers
- Often many providers arrange to interconnect with each other at a single “peering point”

BGP-4: Border Gateway Protocol

- Assumes the Internet is an arbitrarily interconnected set of AS's.
- 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*).

BGP

- The goal of Inter-domain routing is to 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
- Why?

BGP

- Scalability: An Internet backbone router must be able to forward any packet destined anywhere in the Internet
  - Having a routing table that will provide a match for any valid IP address
- 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 provid
- 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 (transit AS only)
  - gives *path* information
- In addition to the BGP speakers, the AS has one or more border “gateways” which need not be the same as the speakers
- The border gateways are the routers through which packets enter and leave the AS

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

BGP Example

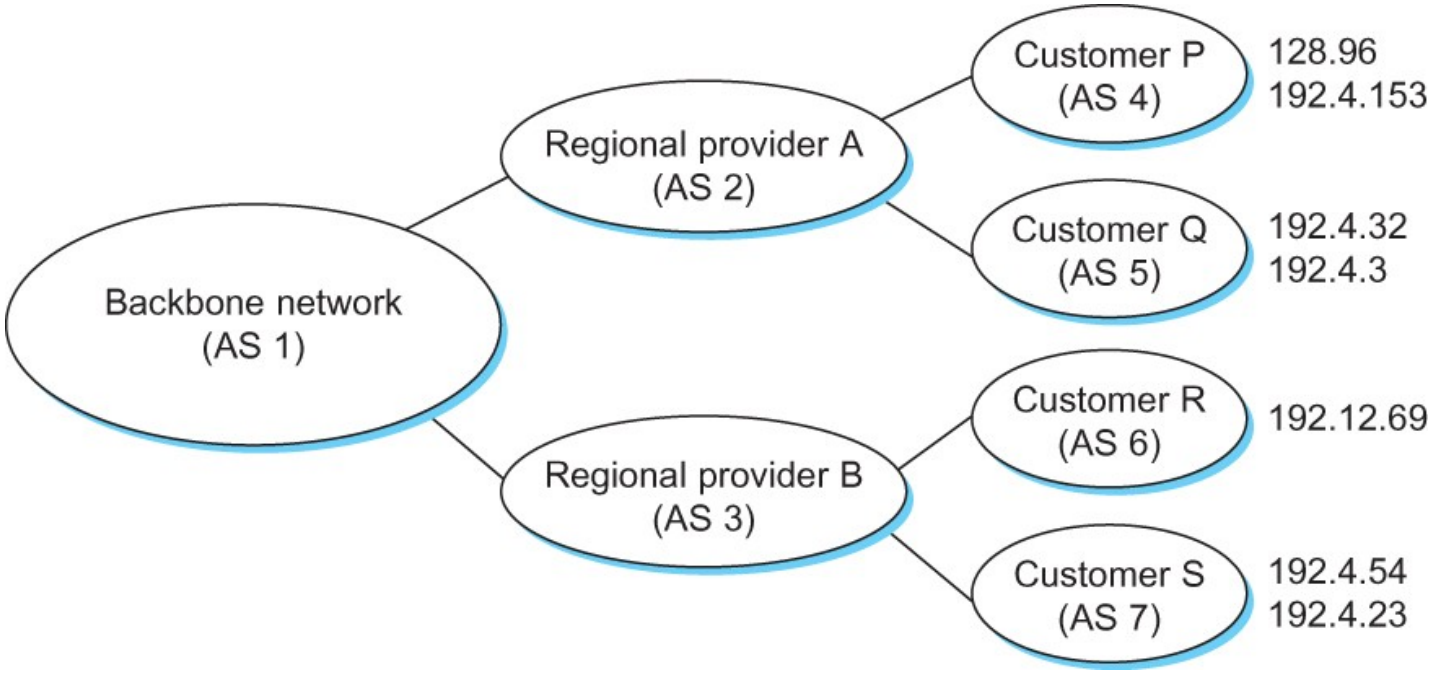

Example of a network running BGP

BGP 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

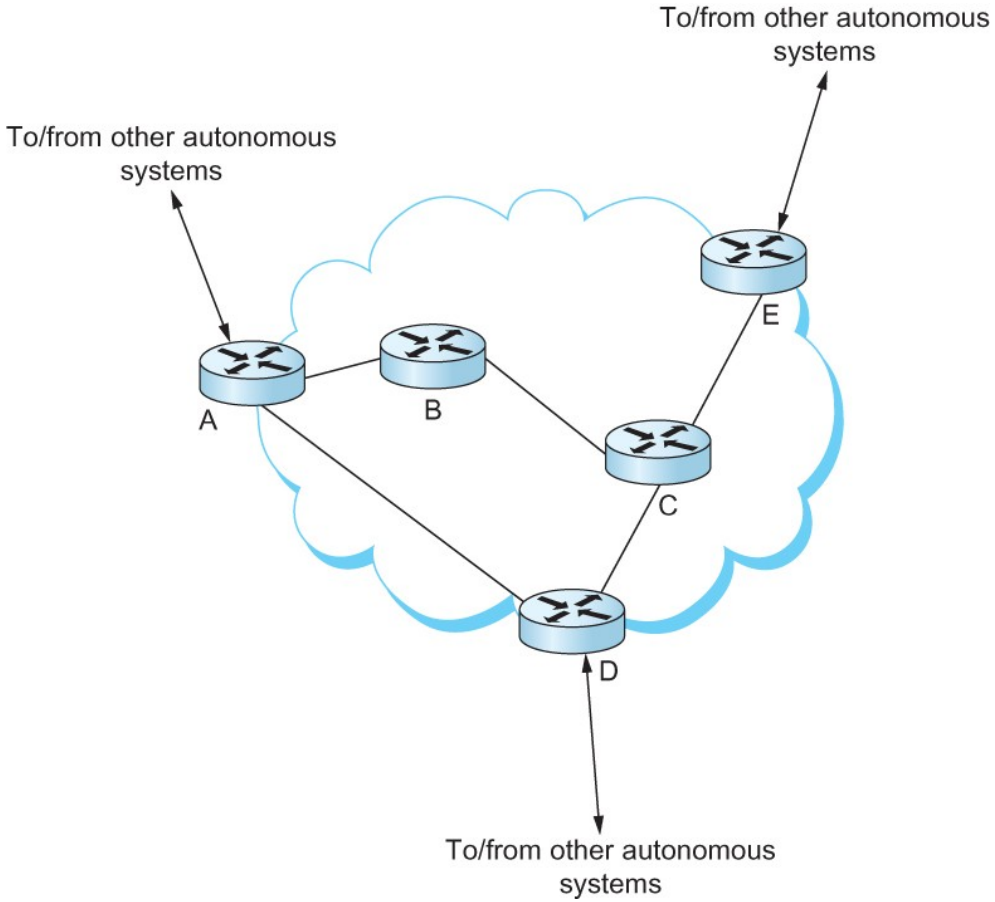

All routers run iBGP and an intradomain routing protocol. Border routers (A, D, E) also run eBGP to other ASs

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 routing table, and combined table at router B

Routing Areas

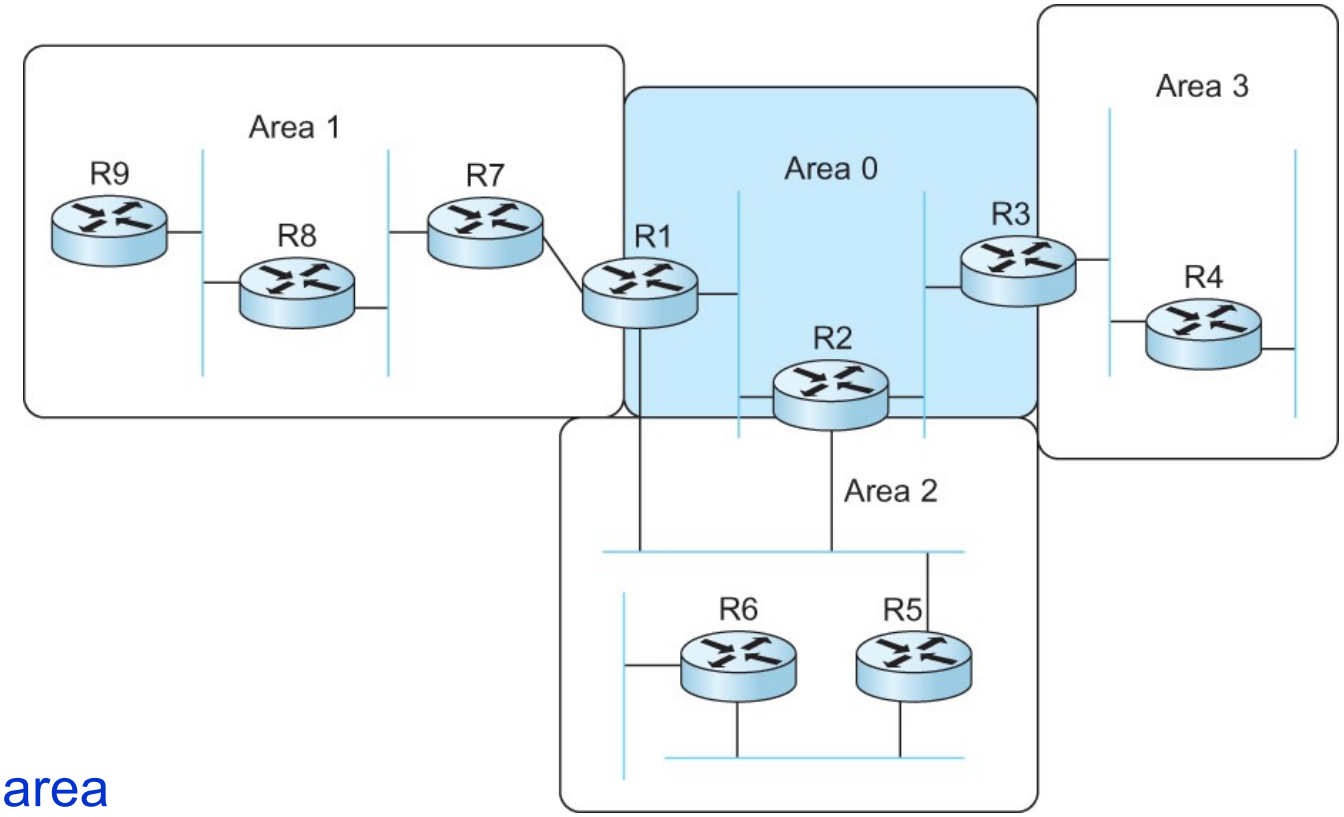

Backbone area

Area border router (ABR)

A domain divided into area