DATA COMMUNICATOIN NETWORKING

Instructor: Ouldooz Baghban Karimi

Course Book & Slides:
Computer Networking, A Top-Down Approach
By: Kurose, Ross

Course Overview

- **Basics of Computer Networks**
  - Internet & Protocol Stack
  - Application Layer
  - Transport Layer
  - Network Layer
  - Data Link Layer

- **Advanced Topics**
  - Case Studies of Computer Networks
  - Internet Applications
  - Network Management
  - Network Security

Routing in the Internet

- **Intra-AS Routing**
  - Also known as Interior Gateway Protocols
    - **RIP**: Routing Information Protocol
    - **OSPF**: Open Shortest Path First
    - **IGRP**: Interior Gateway Routing Protocol (Cisco proprietary)
- **Inter-AS Routing**
  - **BGP**: Border Gateway Protocol

Routing Information Protocol

- **Distance Vector Algorithm**
  - Distance Vector
    - Distance metric: Hope Count
      - Max: 15 hops
      - Each link costs 1
    - DVs exchanged with neighbors every 30/sec in advertisement
    - Each advertisement lists of up to 25 destination subnets

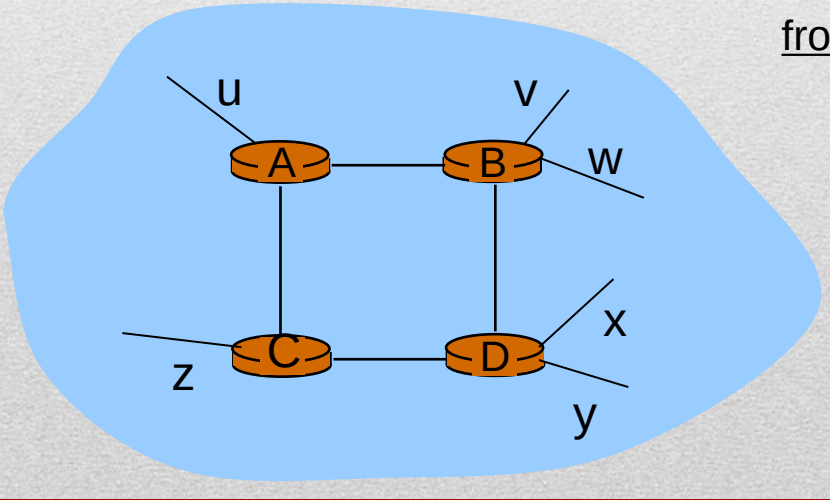

from router A to destination *subnets:*

subnet	hops
u	1
v	2
w	2
x	3
y	3
z	2

Routing Information Protocol

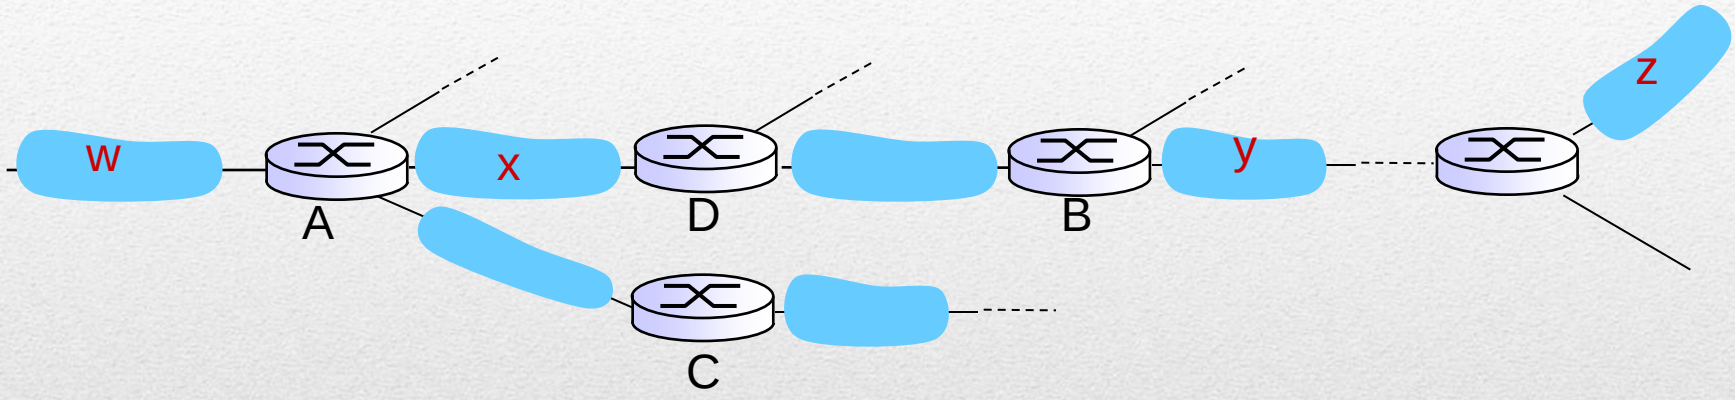

routing table in router D

destination subnet	next router	# hops to dest
w	A	2
y	B	2
z	B	7
x	--	1
….	….	….

Routing Information Protocol

A-to-D advertisement

dest	next	hops
w	-	1
x	-	1
z	C	4
....	...	...

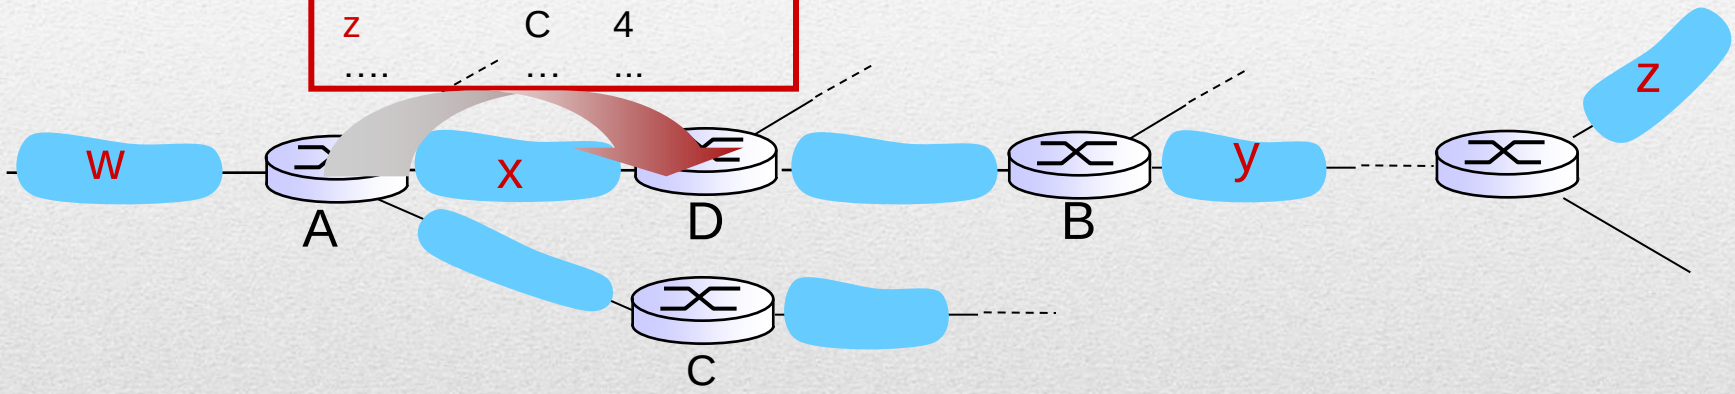

routing table in router D

destination subnet	next router	# hops to dest
w	A	2
y	B	2
z	B	7
x	--	1
....	....	....

A

5

RIP Link Failure & Recovery

- If no advertisement heard after 180 Sec
  - Neighbor/link declared dead
  - New advertisement sent to neighbors
  - Neighbors in turn send out new advertisements if their tables changed
  - Link Failure information propagates to entire net
  - **Poison reverse** is used to prevent Ping-Pong loops

RIP Table Processing

- RIP routing tables managed by application-level process called route-d (daemon)
- Advertisements sent in UDP packets, periodically repeated

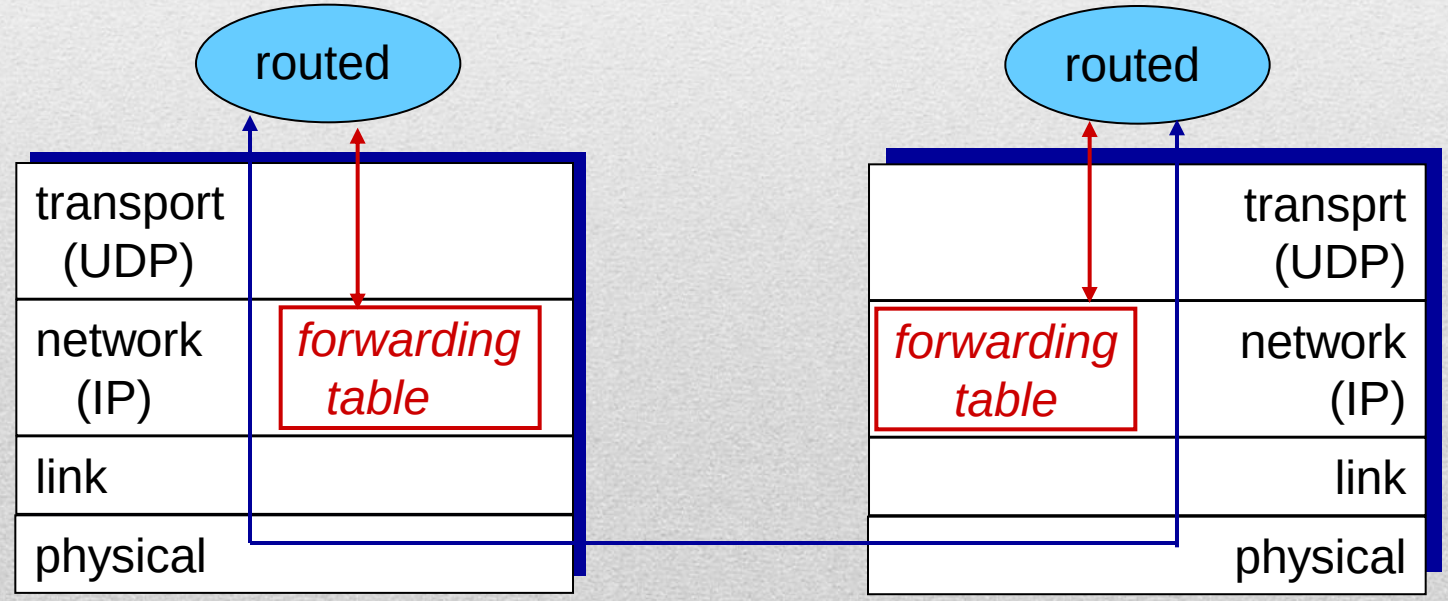

Open Shortest Path First

- Open: publicly available
- Link State Algorithm
  - LS packet dissemination
  - Topology map at each node
  - Route computation using Dijekstra's algorithm
- OSPF advertisement carries one entry per neighbor
  - Advertisements flooded to entire AS
    - Carried in OSPF messages directly over IP
- ISIS routing protocol: Nearly identical to OSPF

OSPF Advanced Features

- Security
  - All OSPF messages authenticated (to prevent malicious intrusion)
- Multi-path
  - Multiple same-cost paths allowed (only one path in RIP)
- TOD
  - For each link, multiple cost metrics for different TOS (e.g., satellite link cost set "low" for best effort ToS; high for real time ToS)
- Integrated uni- and multicast support:
- Integrated Unicast and Multicast Support
  - Multicast OSPF (MOSPF) uses same topology data base as OSPF
- Hierarchical OSPF in large domains

Hierarchical OSPF

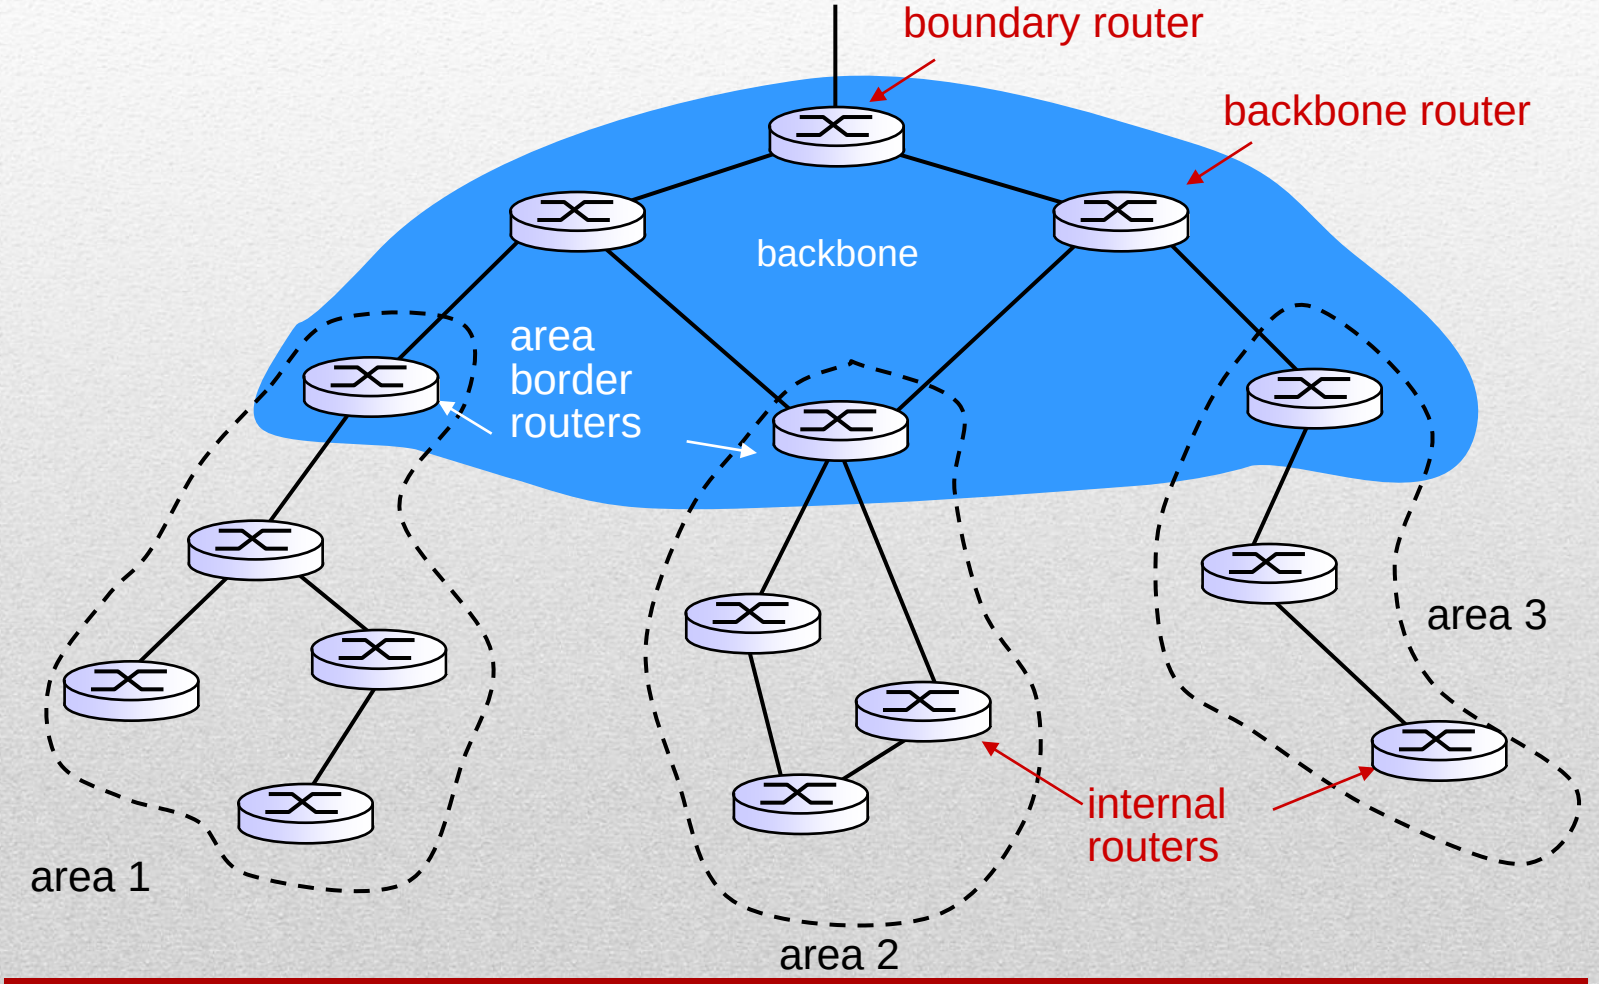

Hierarchical OSPF

- Two-level hierarchy: local area, backbone.
  - Link-state advertisements only in area
  - Each nodes has detailed area topology; only know direction (shortest path) to nets in other areas.
- Area border routers: summarize distances to nets in own area, advertise to other Area Border routers.
- Backbone routers: run OSPF routing limited to backbone.
- Boundary routers: connect to other AS's.

Internet Inter-AS Routing: BGP

- Border Gateway Protocol: the de facto inter-domain routing protocol
  - Glue that holds the Internet together
- BGP provides each AS a means to
  - eBGP: Obtain subnet reachability information from neighboring Ass.
  - iBGP: Propagate reachability information to all AS-internal routers.
  - Determine good routes to other networks based on reachability information and policy.
- Allows subnet to advertise its existence to rest of Internet

BGP Basics

- BGP session: two BGP routers (peers) exchange BGP messages:
  - Advertising paths to different destination network prefixes (path vector protocol)
  - Exchanged over semi-permanent TCP connections

- When AS3 advertises a prefix to AS1
  - AS3 promises it will forward datagrams towards that prefix
  - AS3 can aggregate prefixes in its advertisement

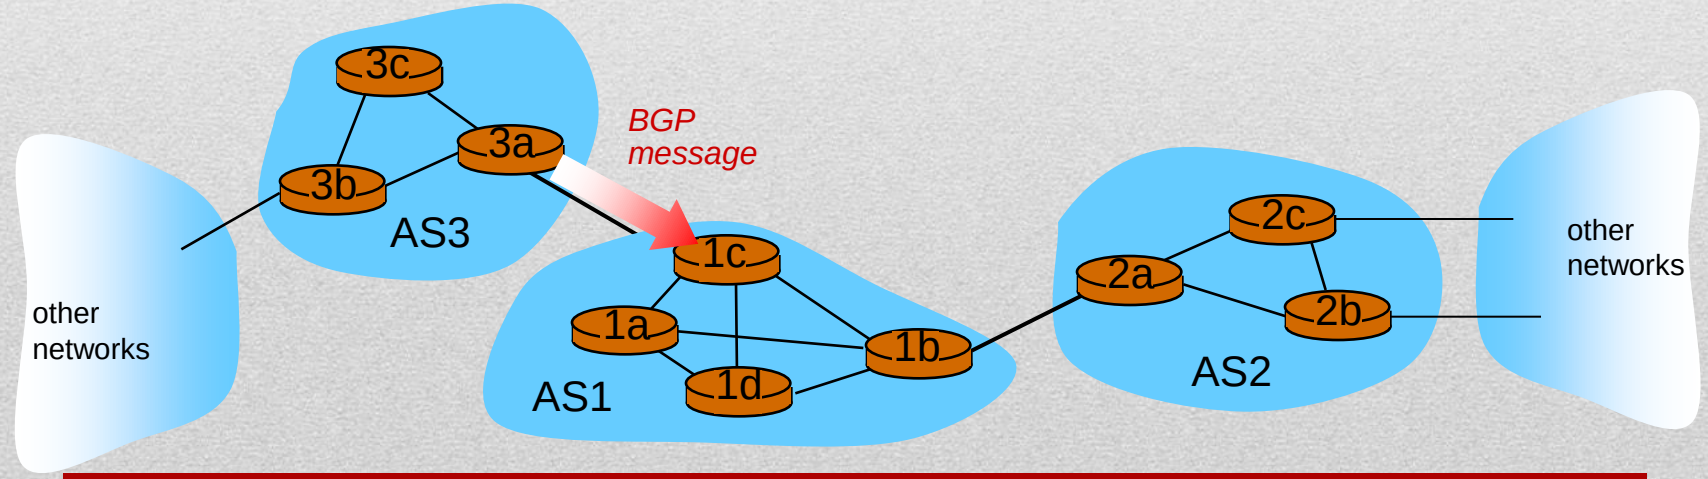

BGP: Distributing Path Information

- Using eBGP session between 3a and 1c, AS3 sends prefix reachability info to AS1.
  - 1c can then use iBGP do distribute new prefix info to all routers in AS1
  - 1b can then re-advertise new reachability info to AS2 over 1b-to-2a eBGP session
- When router learns of new prefix, it creates entry for prefix in its forwarding table.

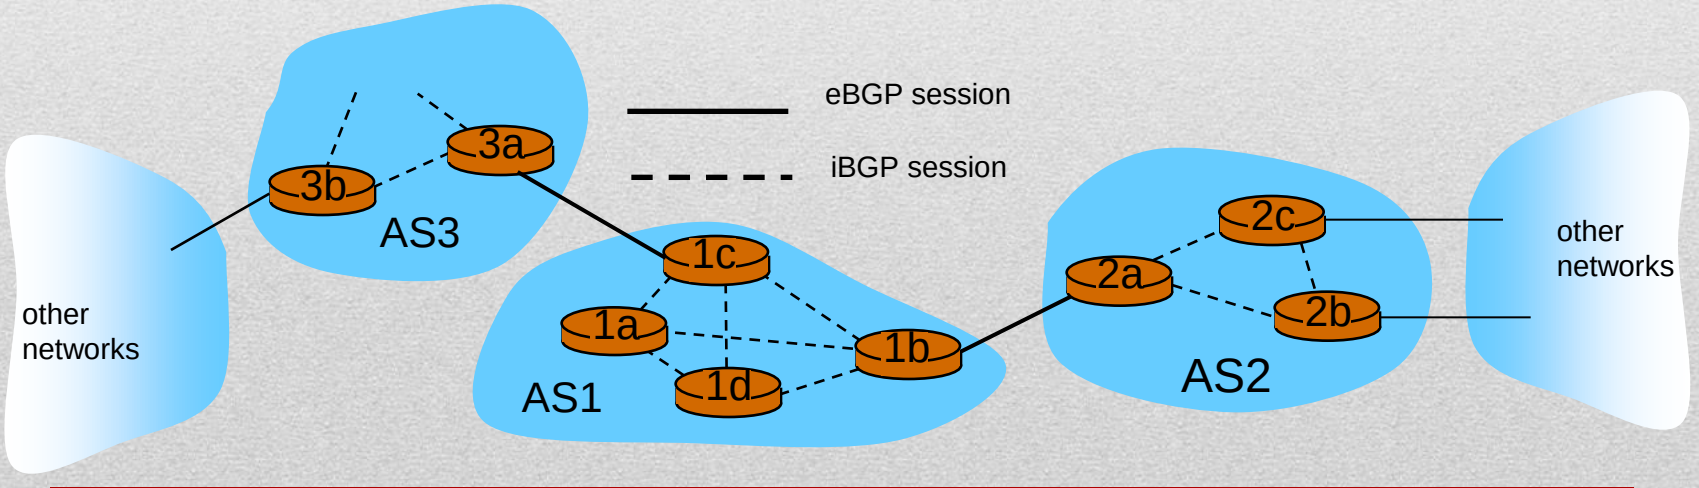

Path Attributes & BGP Routes

- Advertised prefix includes BGP attributes
  - Prefix + attributes = route
- Two important attributes:
  - AS-PATH: contains ASs through which prefix advertisement has passed: e.g., AS 67, AS 17
  - NEXT-HOP: indicates specific internal-AS router to next-hop AS. (may be multiple links from current AS to next-hop-AS)
- Gateway router receiving route advertisement uses import policy to accept/decline
  - e.g., never route through AS x
  - Policy-based routing

BGP

- **Route Selection**
  - Router may learn about more than 1 route to destination AS, selects route based on:
    - Local preference value attribute: policy decision
    - Shortest AS-PATH
    - Closest NEXT-HOP router: hot potato routing
    - Additional criteria

- **Messages**
  - BGP messages exchanged between peers over TCP connection
  - BGP messages:
    - OPEN: opens TCP connection to peer and authenticates sender
    - UPDATE: advertises new path (or withdraws old)
    - KEEPALIVE: keeps connection alive in absence of UPDATES; also ACKs OPEN request
    - NOTIFICATION: reports errors in previous msg; also used to close connection

BGP Routing Policy

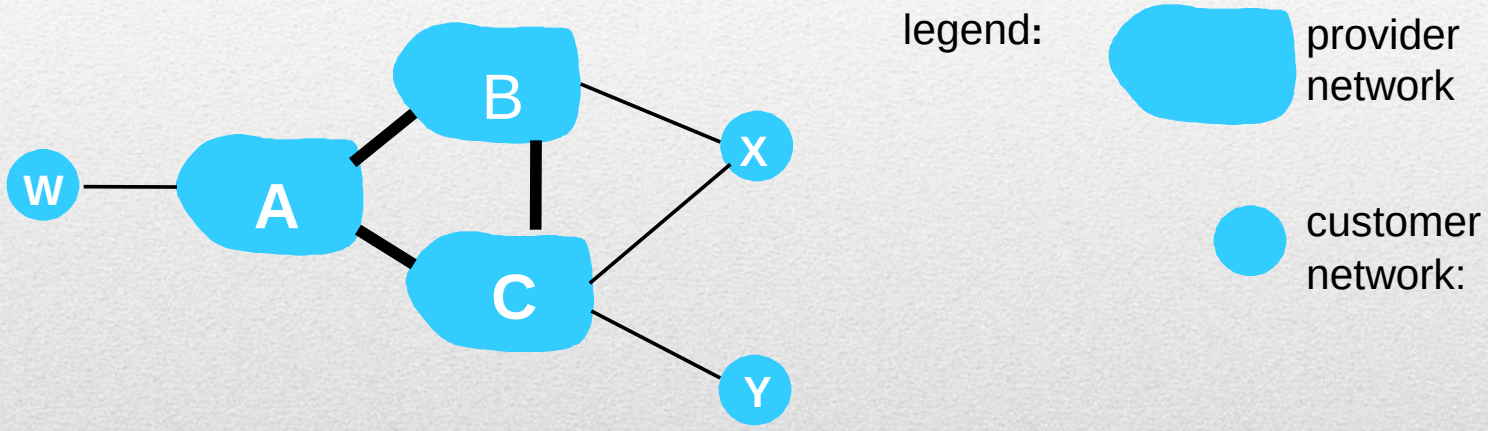

- A,B,C are *provider networks*
- X,W,Y are customer (of provider networks)
- X is *dual-homed:* attached to two networks
  - X does not want to route from B via X to C
  - .. so X will not advertise to B a route to C

BGP Routing Policy

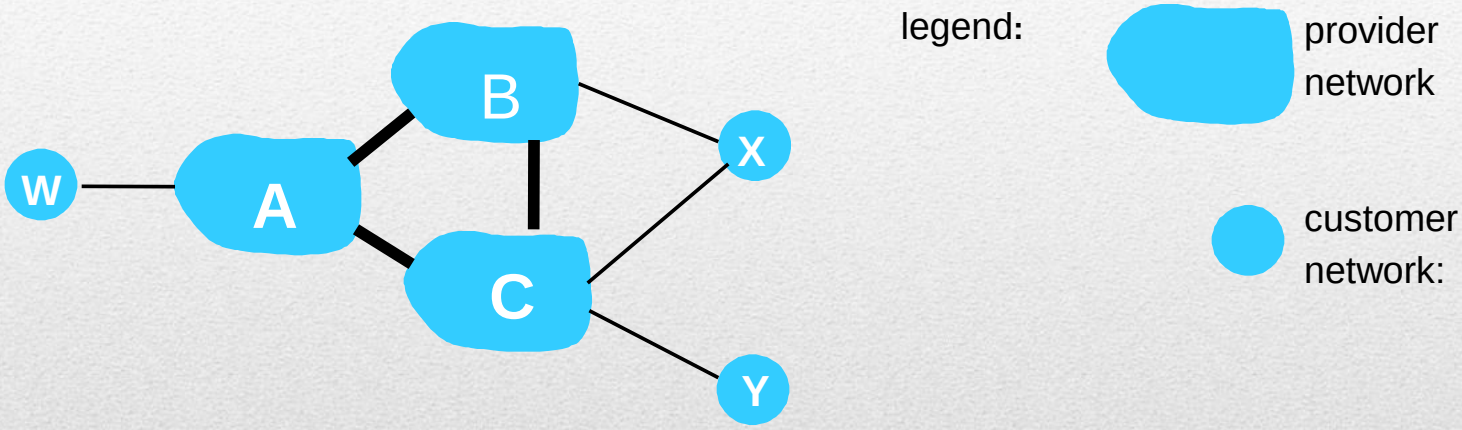

- A advertises path AW to B
- B advertises path BAW to X
- Should B advertise path BAW to C?
  - No way! B gets no "revenue" for routing CBAW since neither W nor C are B's customers
  - B wants to force C to route to w via A
  - B wants to route *only* to/from its customers!

Why different Intra-AS & Inter-AS?

- **Policy**
  - Inter-AS: admin wants control over how its traffic routed, who routes through its net.
  - Intra-AS: single admin, so no policy decisions needed
- **Scale**
  - Hierarchical routing saves table size, reduced update traffic
- **Performance**
  - Intra-AS: can focus on performance
  - Inter-AS: policy may dominate over performance