CS 114: Network Security

Lecture 14 - Routing

Prof. Daniel Votipka
Spring 2023

(some slides courtesy of Prof. Micah Sherr)

Tufts UNIVERSITY

Plan for today

- Administrivia
- Review DNS
- Secure Routing
  - Overview
  - Protocols
  - Attacks
  - Defenses

Administrivia

- Upcoming Cybersecurity Talks:
  - 3/16 - Bailey Kacsmar, Waterloo
    - 3pm in JCC 270
  - 3/30 - Nirvan Tyagi, Cornell
    - 3pm in JCC 270

DNS Review

But what if Alice doesn't know Bob's (bob.com) IP address?

Naive Recursive Query

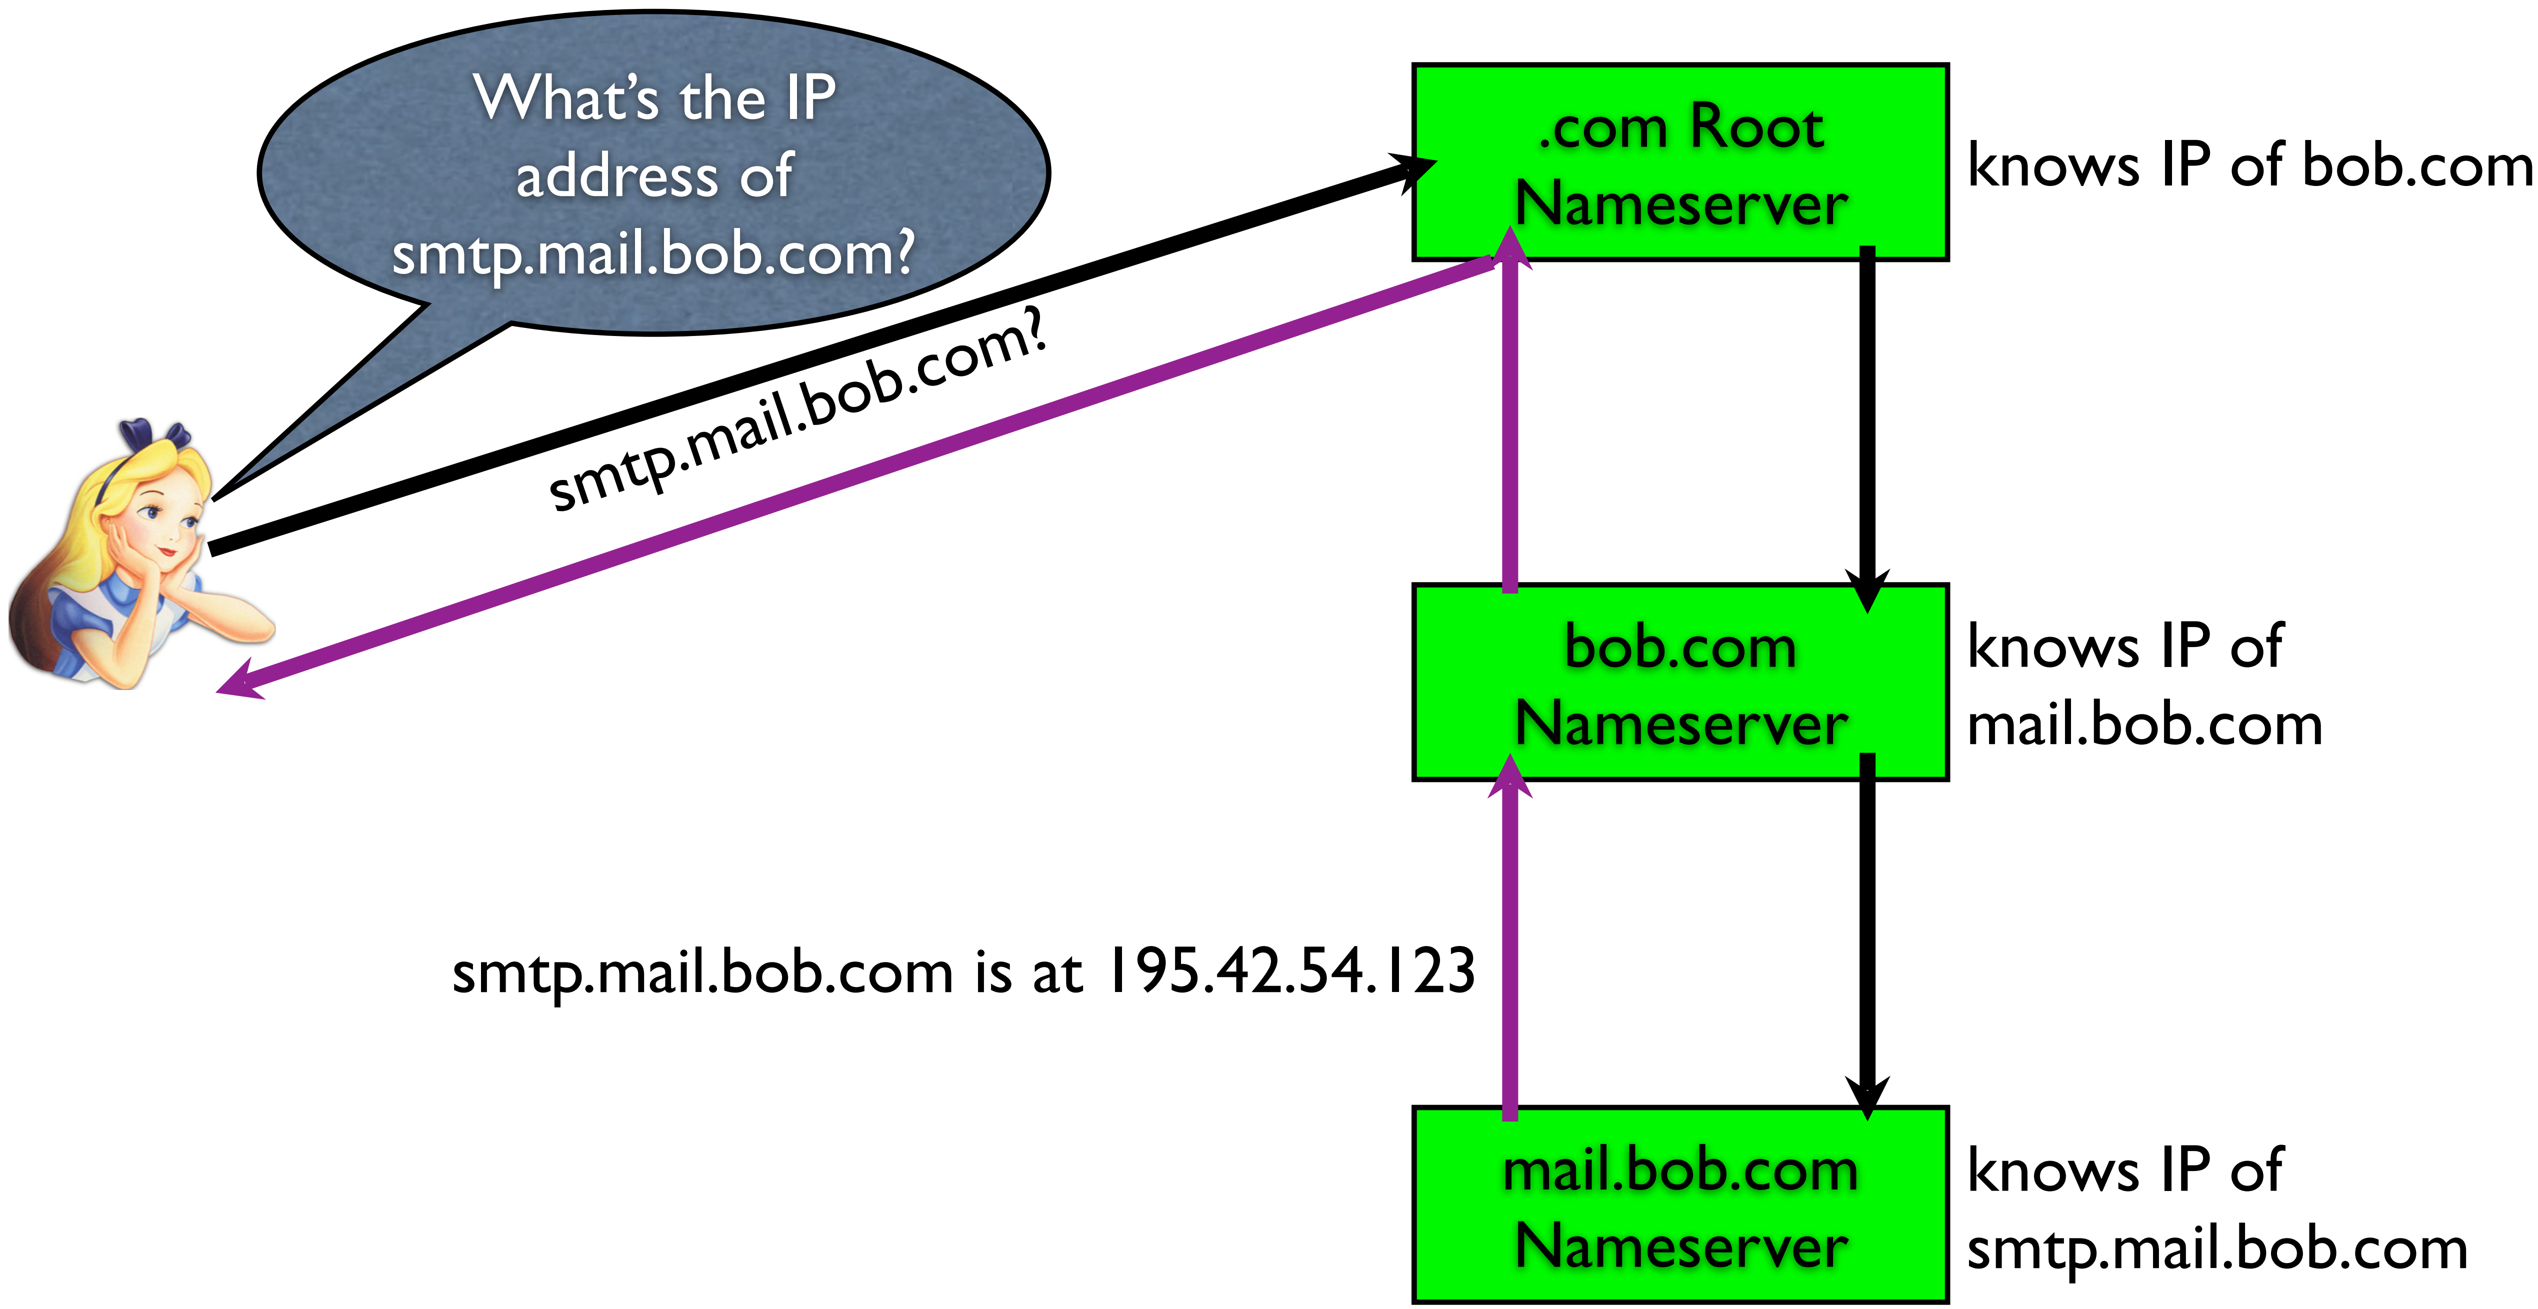

Naive Iterative Query

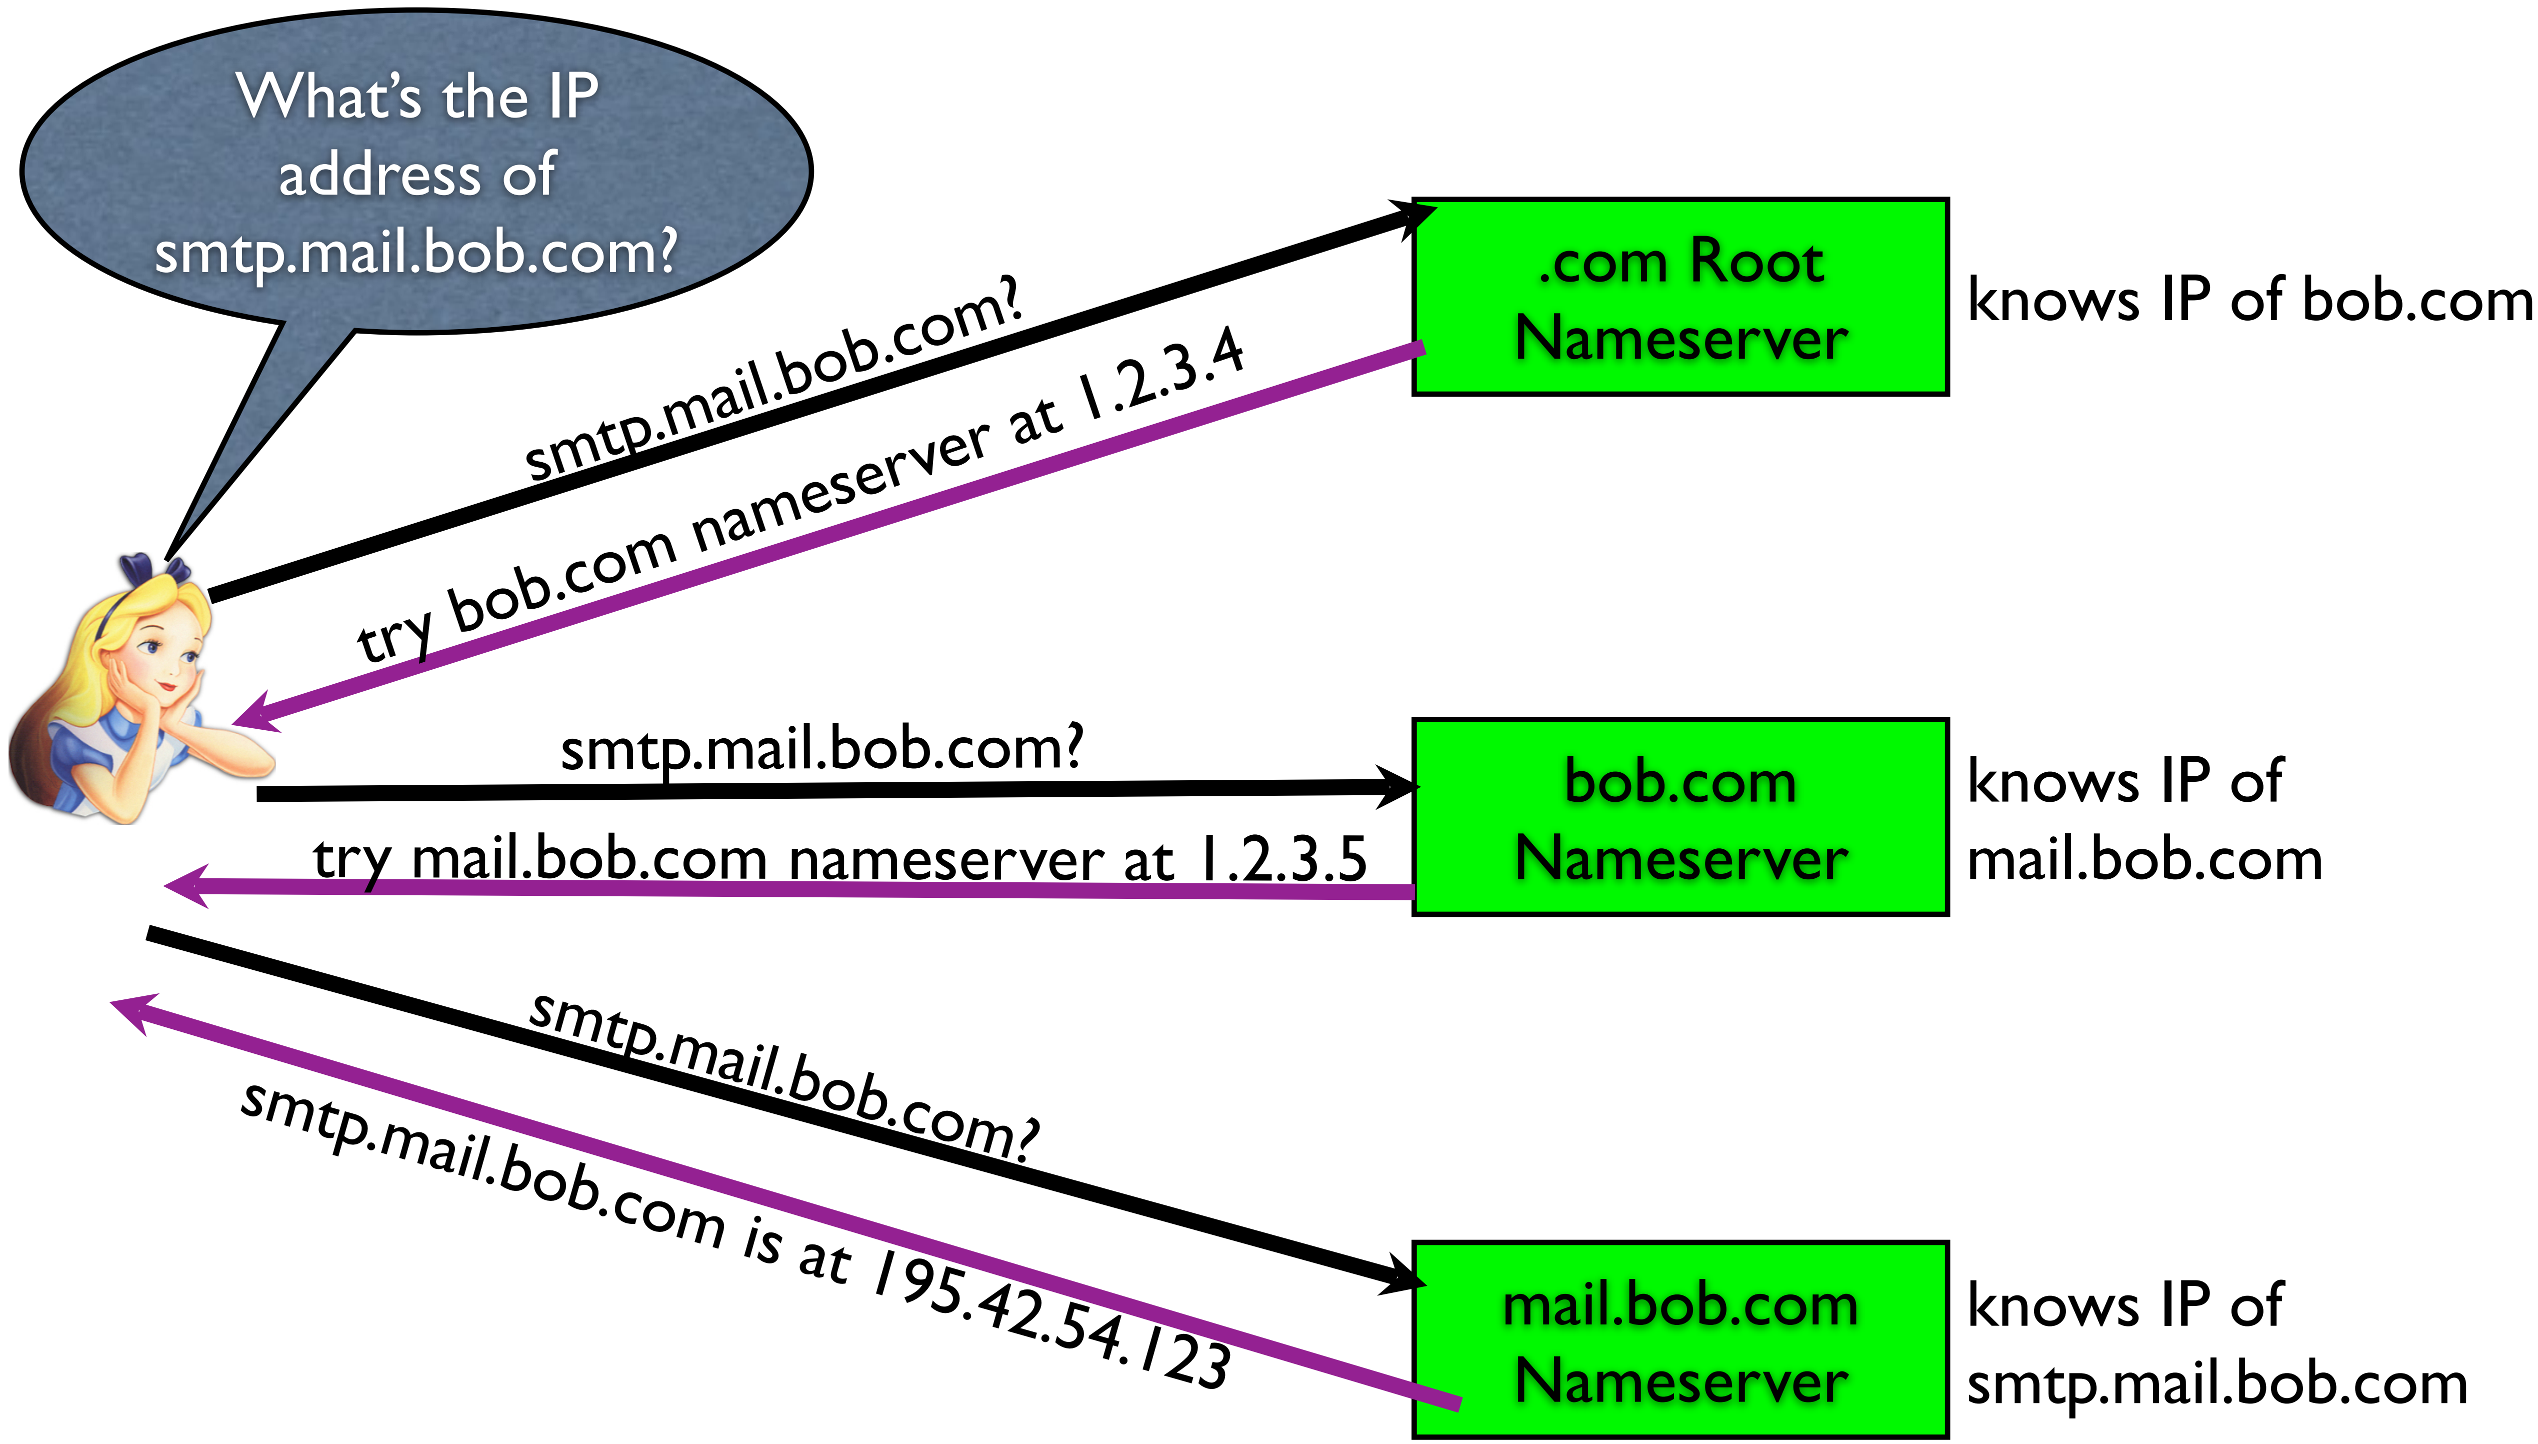

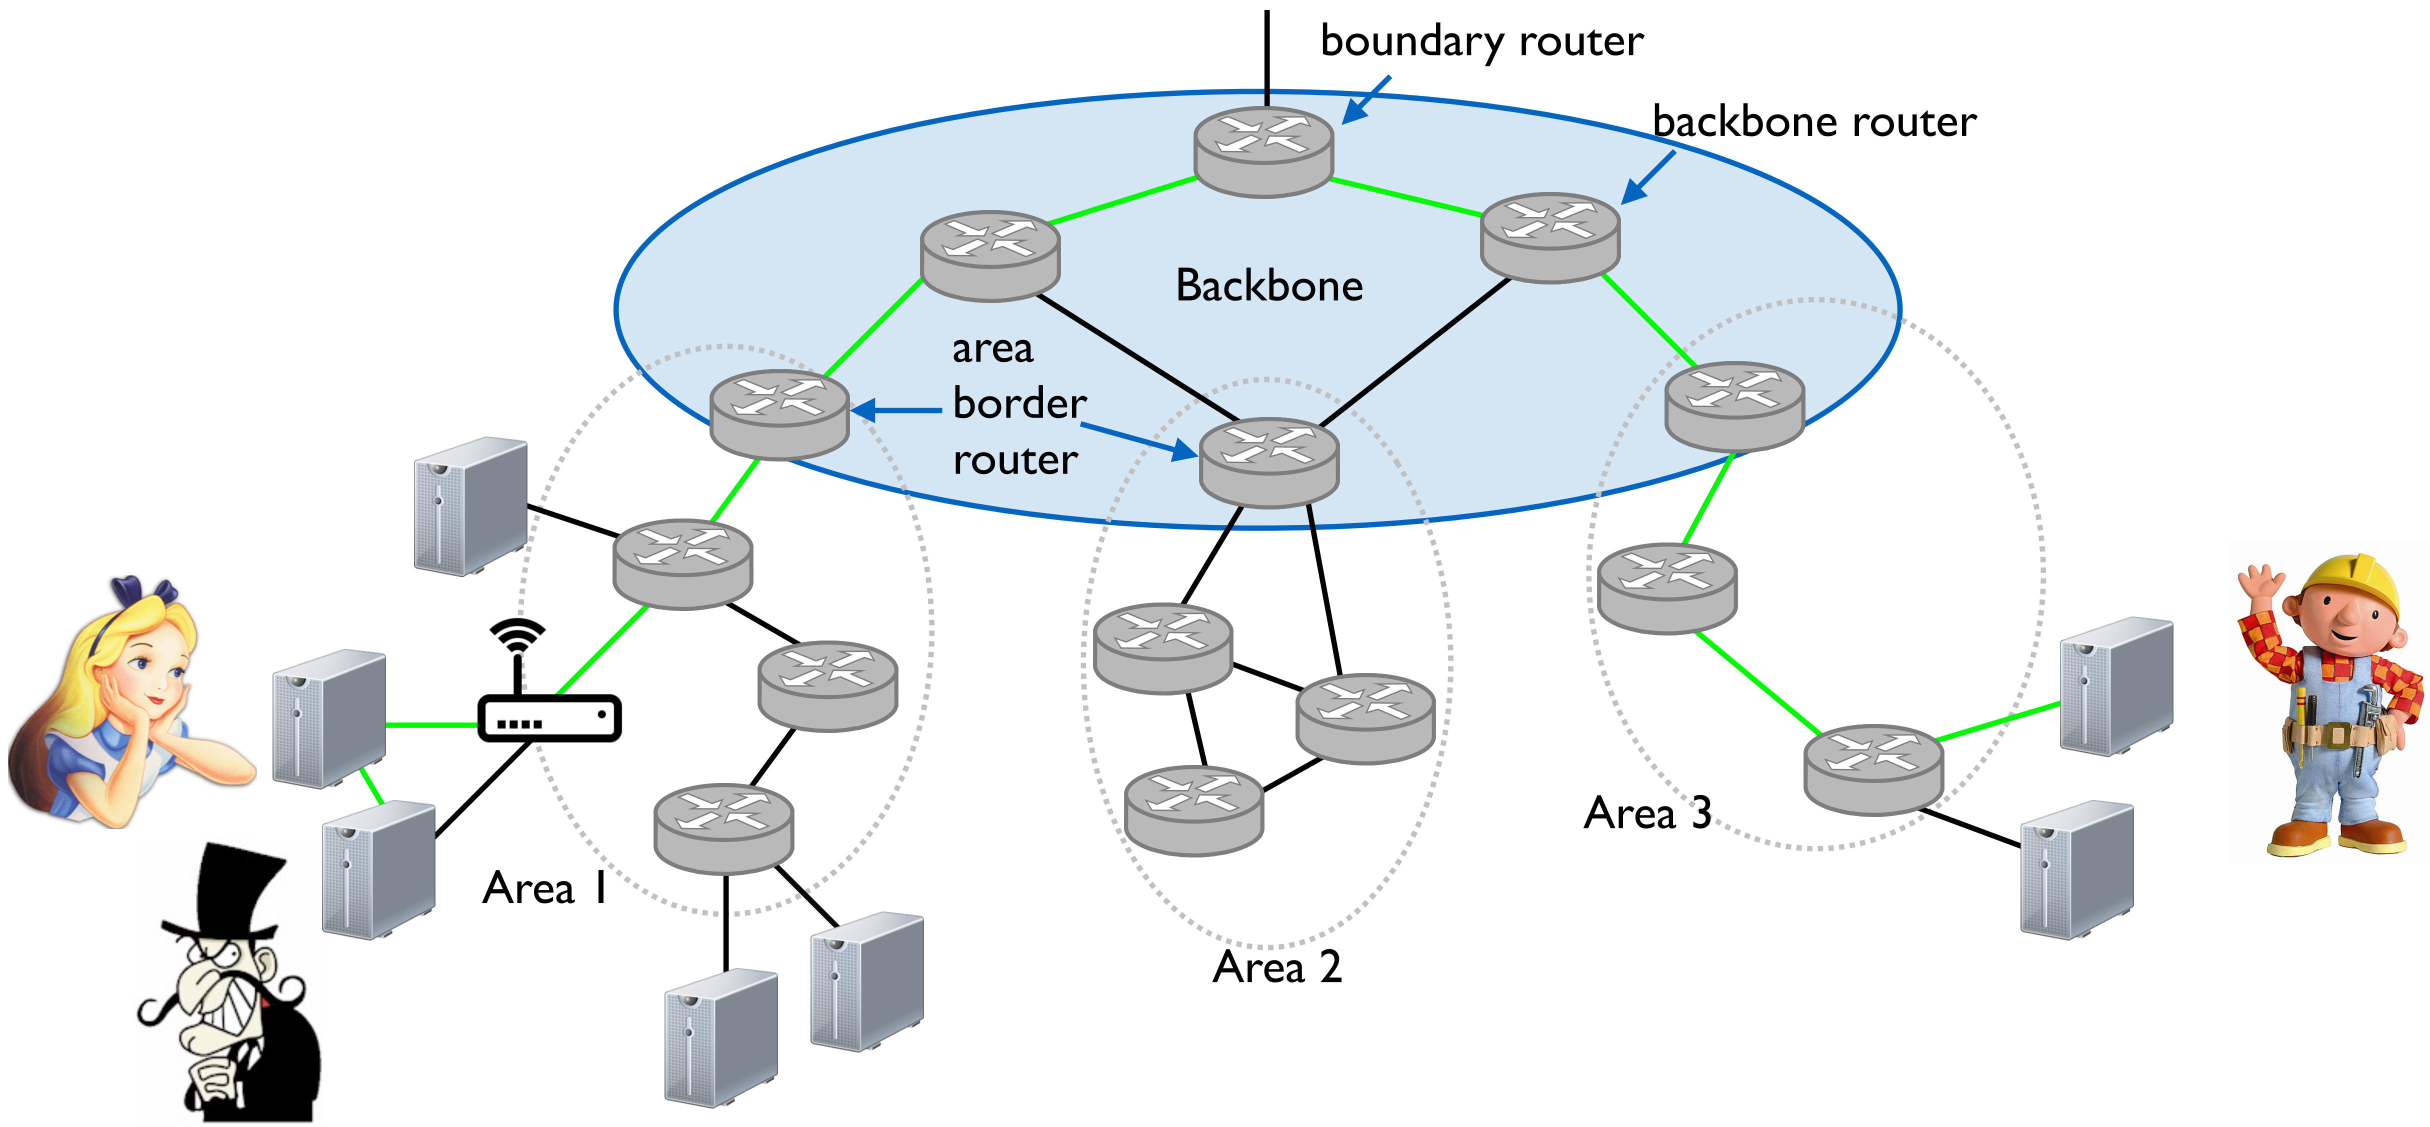
boundary router
backbone router
Backbone
area
border
router
Area 1
Area 2
Area 3

DNSSEC Mechanisms

- Each domain signs their “zone” with a private key
- Public keys published via DNS
- Zones signed by parent zones
- Ideally, you only need a self-signed root, and follow keys down the hierarchy

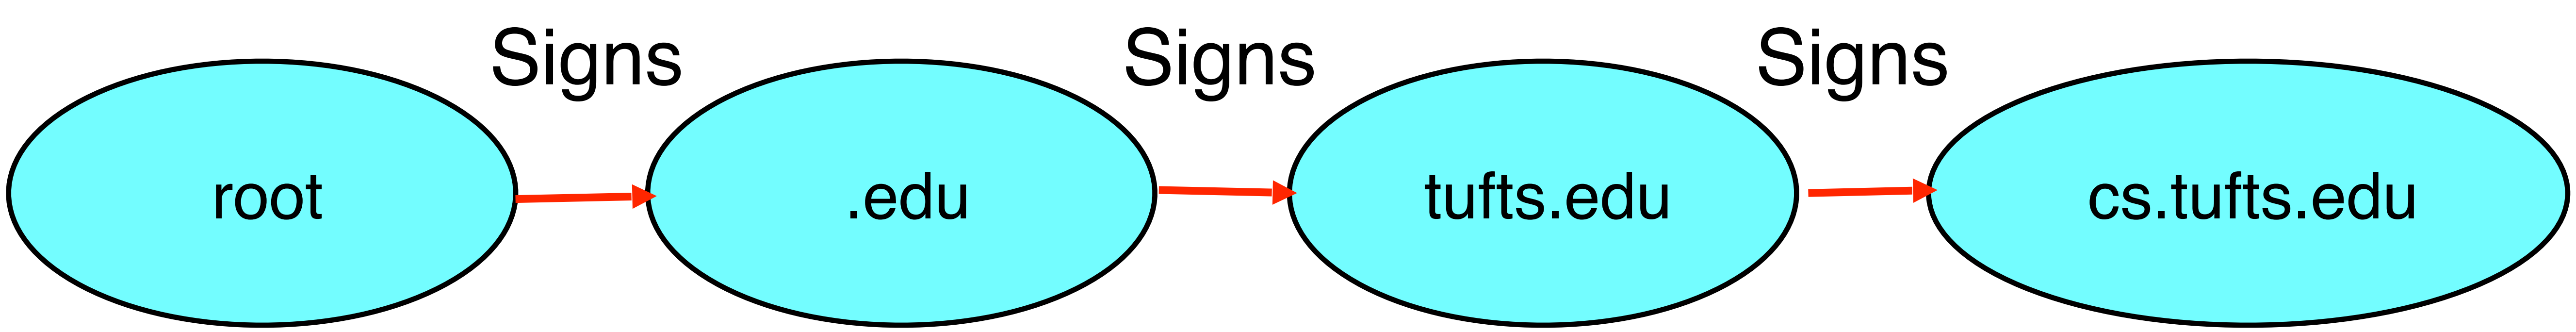

Plan for today

- Administrivia
- Review DNS
- Secure Routing
  - Overview
  - Protocols
  - Attacks
  - Defenses

Routing outside of the local subnet

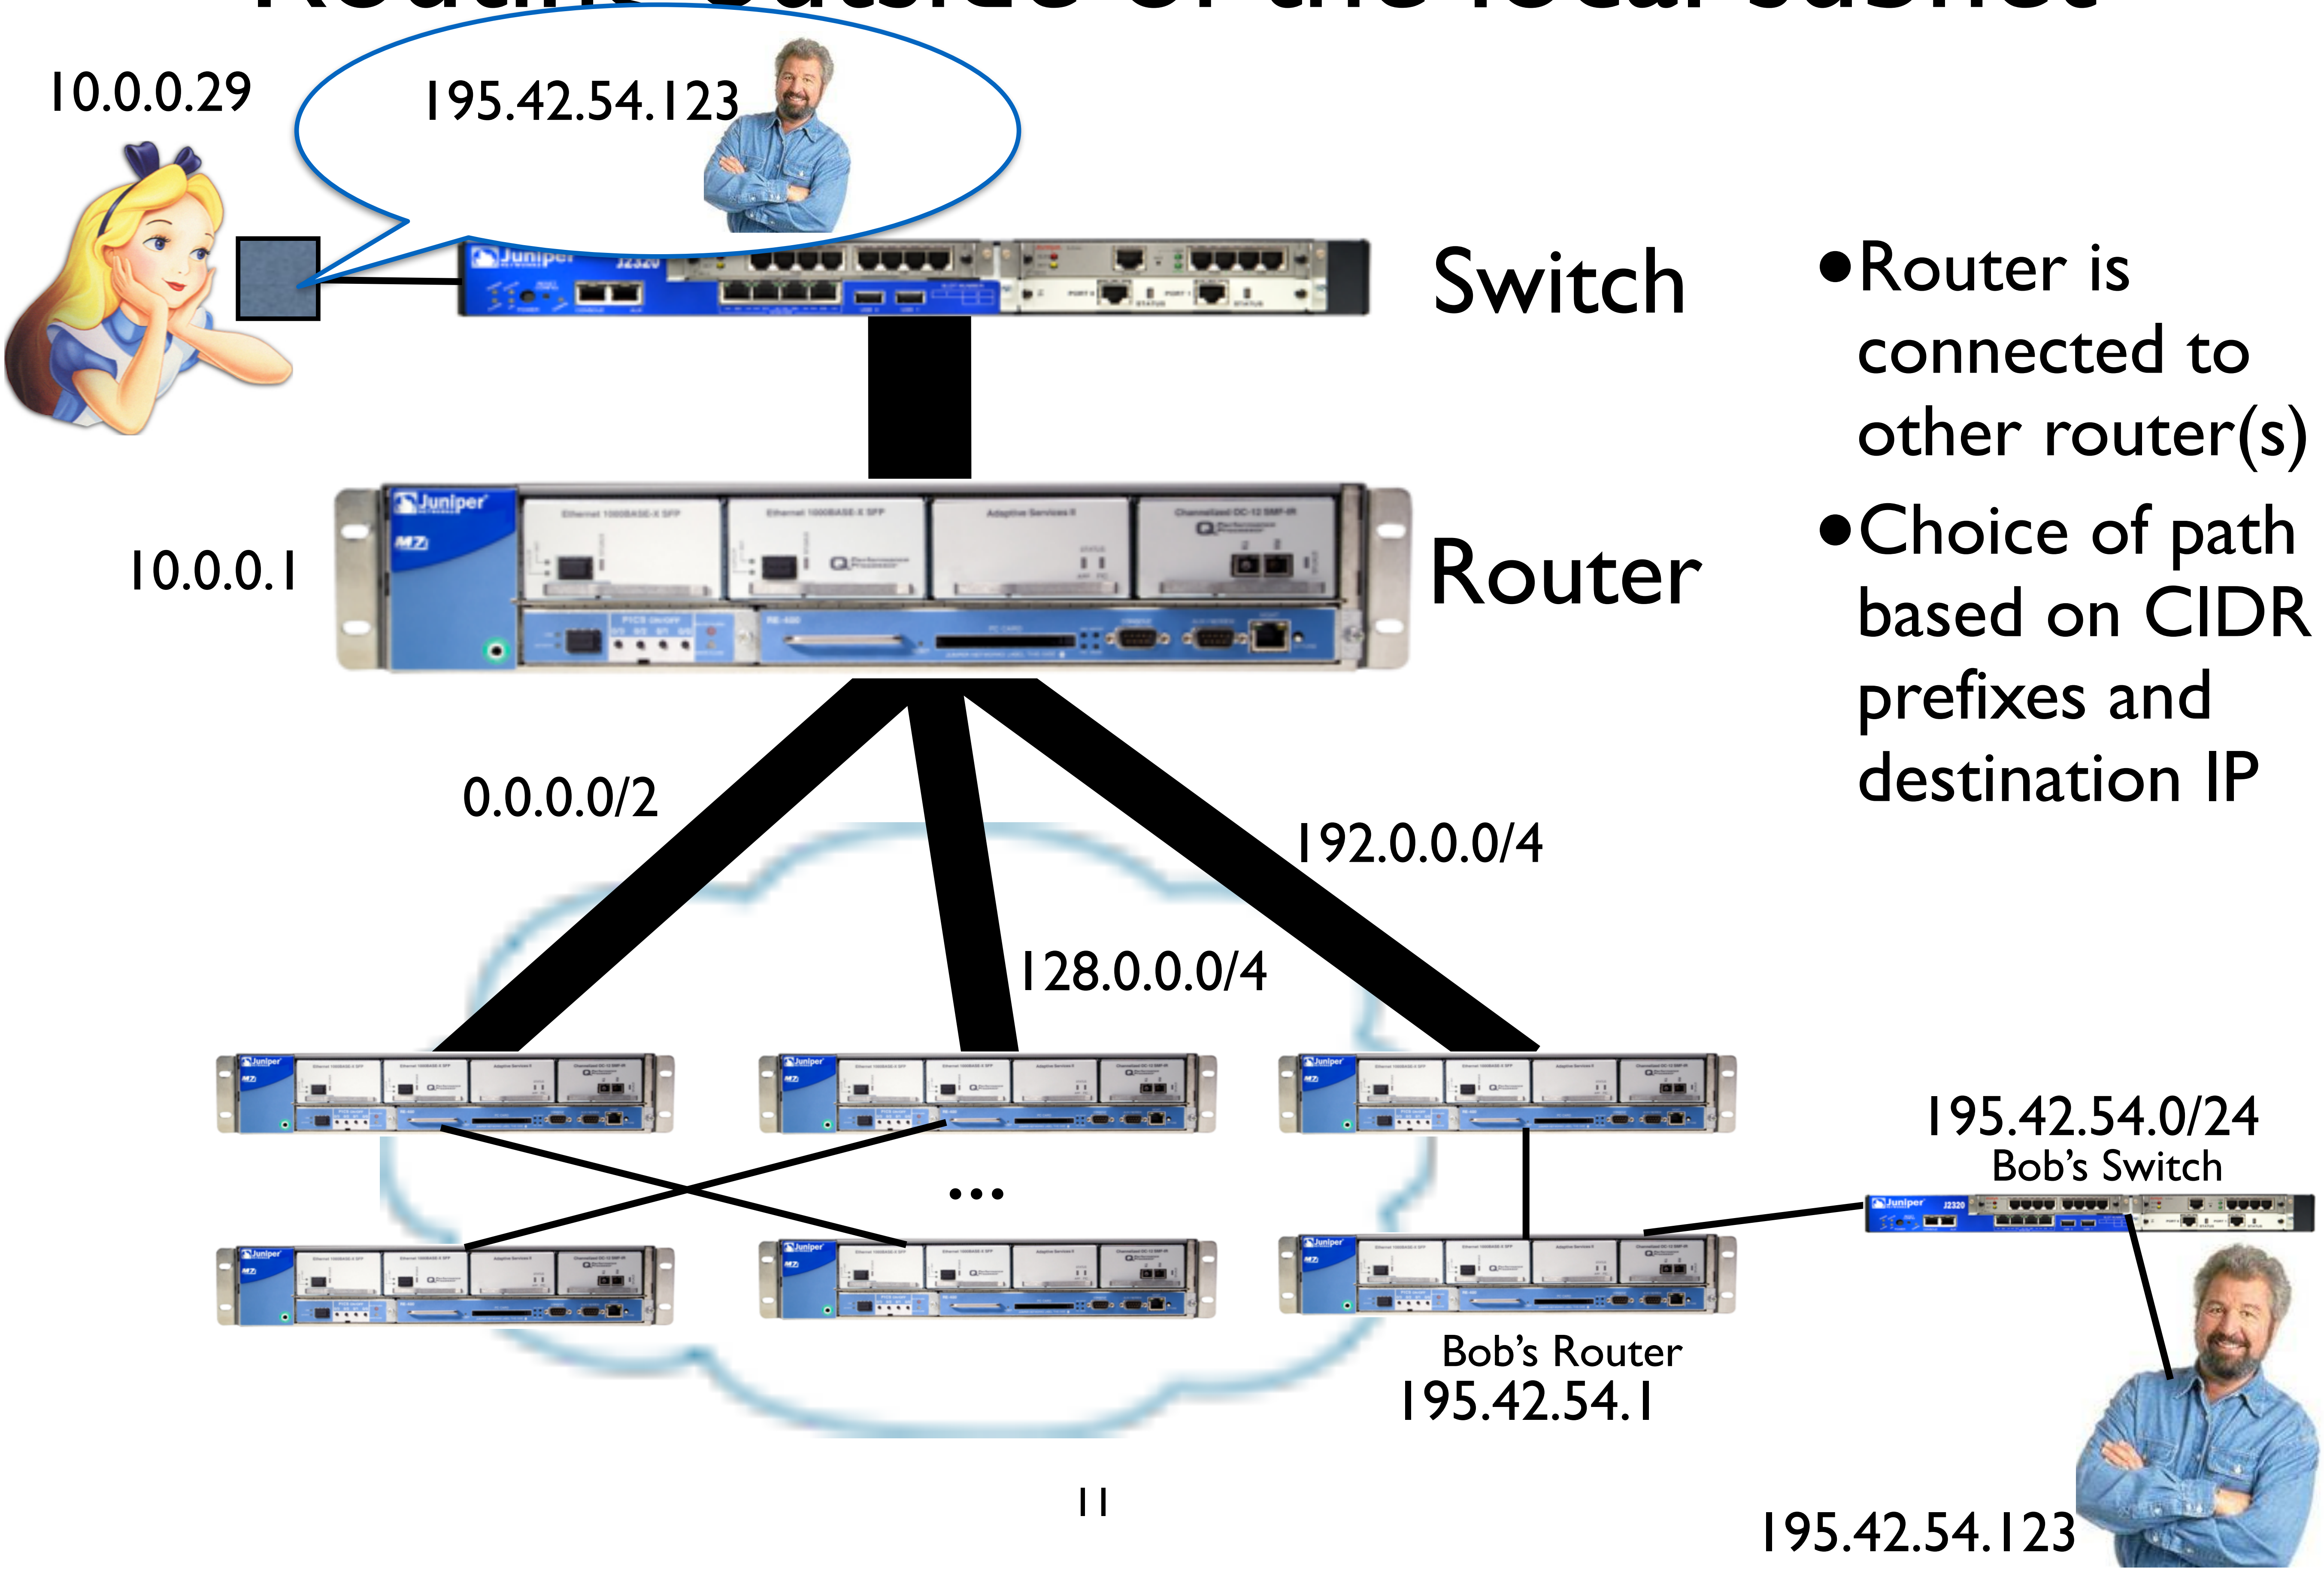

- Router is connected to other router(s)
- Choice of path based on CIDR prefixes and destination IP

Routing Security

- Bad guys/gals/Internet-enabled toaster ovens play games with routing protocols.
- Implications for diverted traffic:
  - Enemy can see the traffic.
  - Enemy can easily modify the traffic.
  - Enemy can drop the traffic.

- Routing security in a nutshell: Cryptography can mitigate effects, but not stop them.

Routing

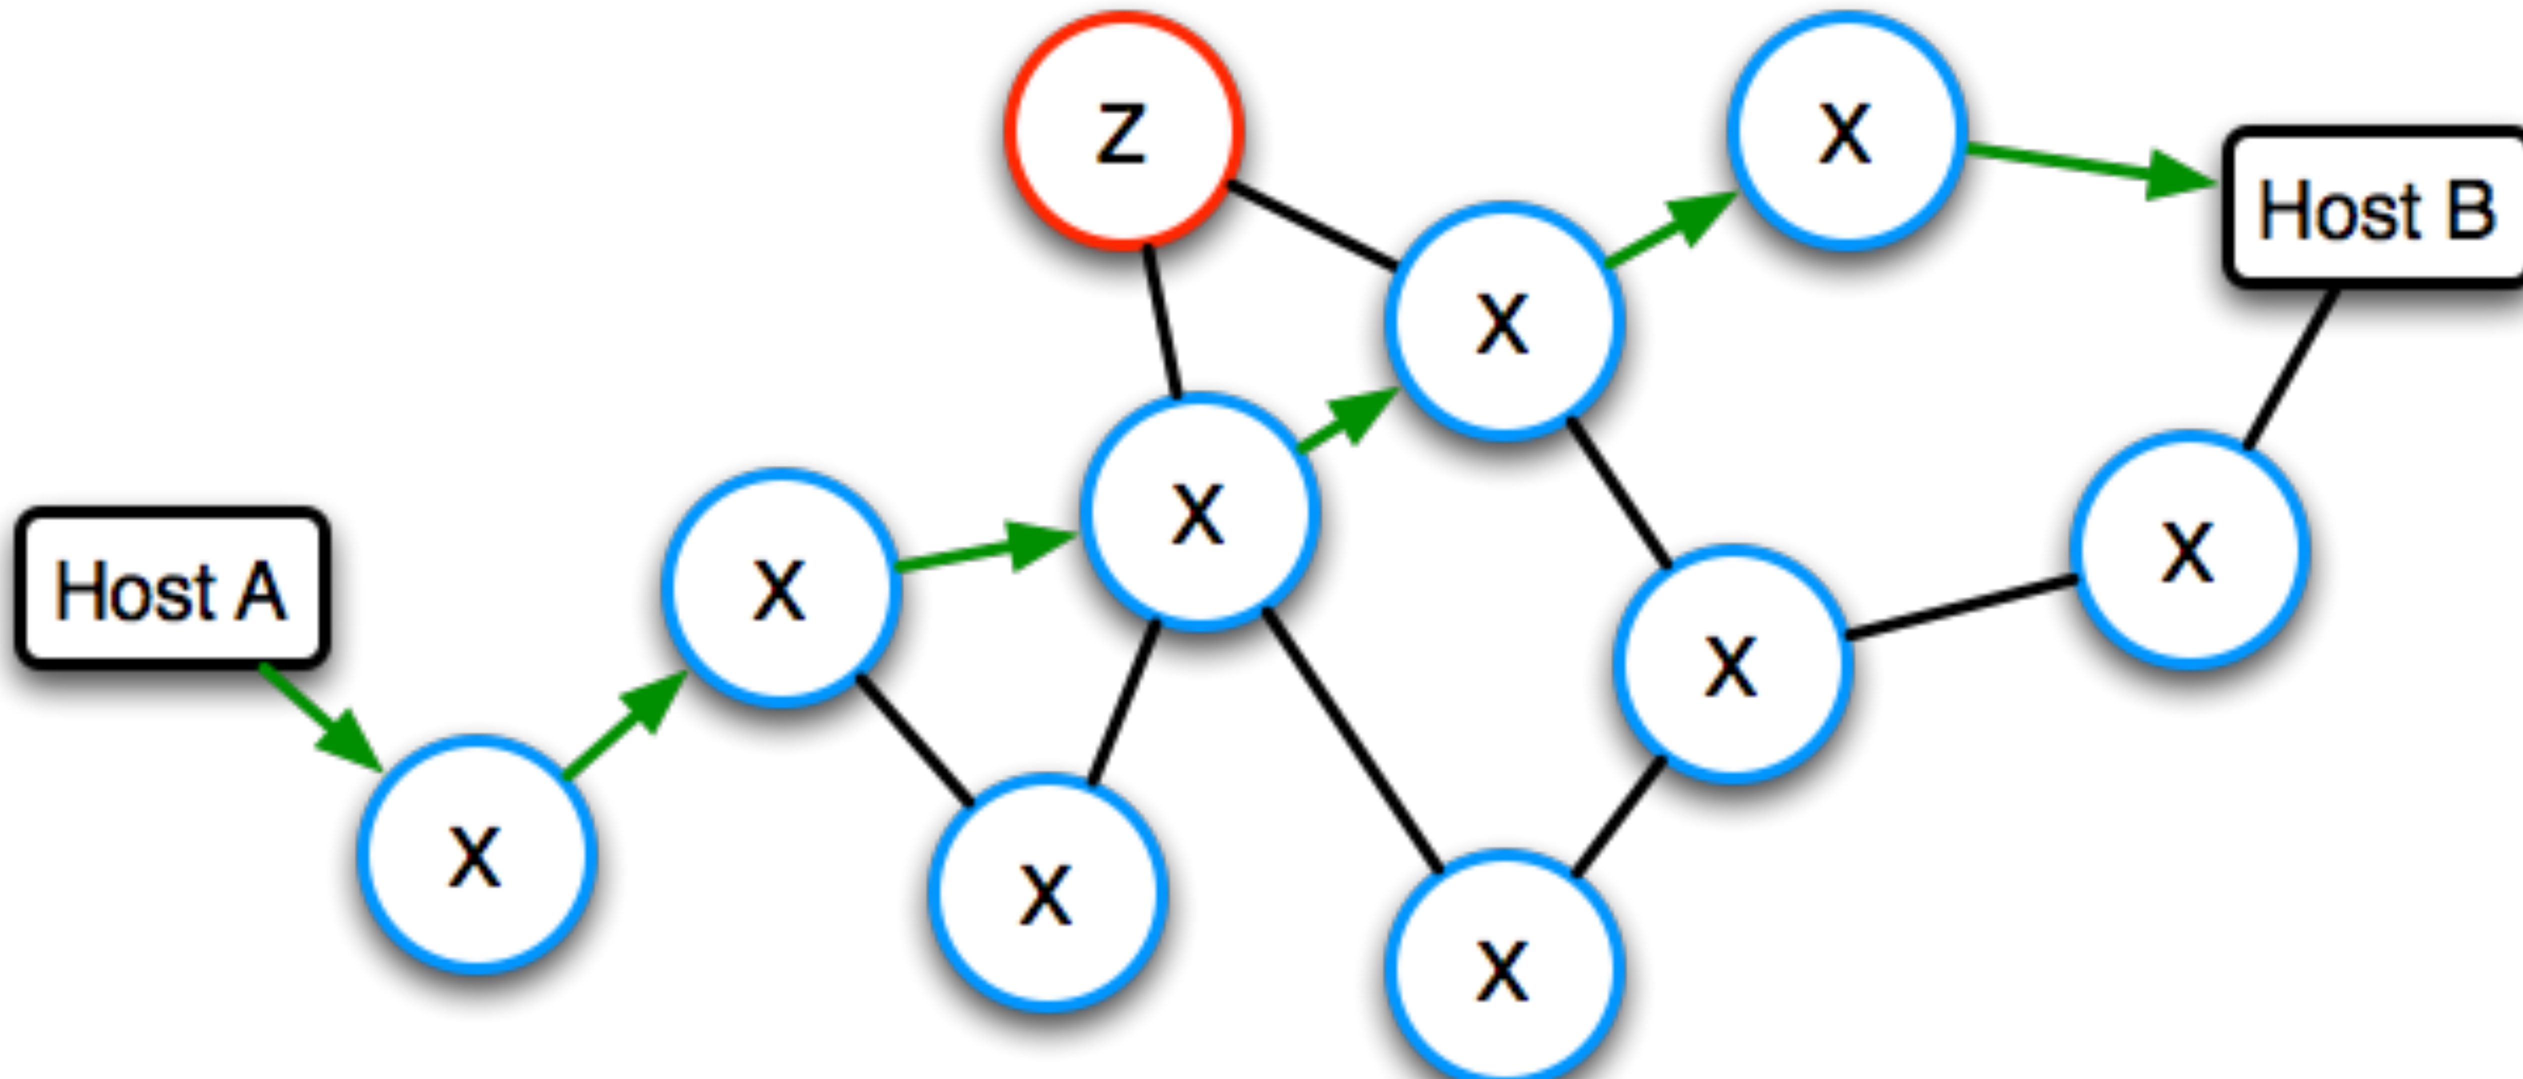

The Enemy's Goal

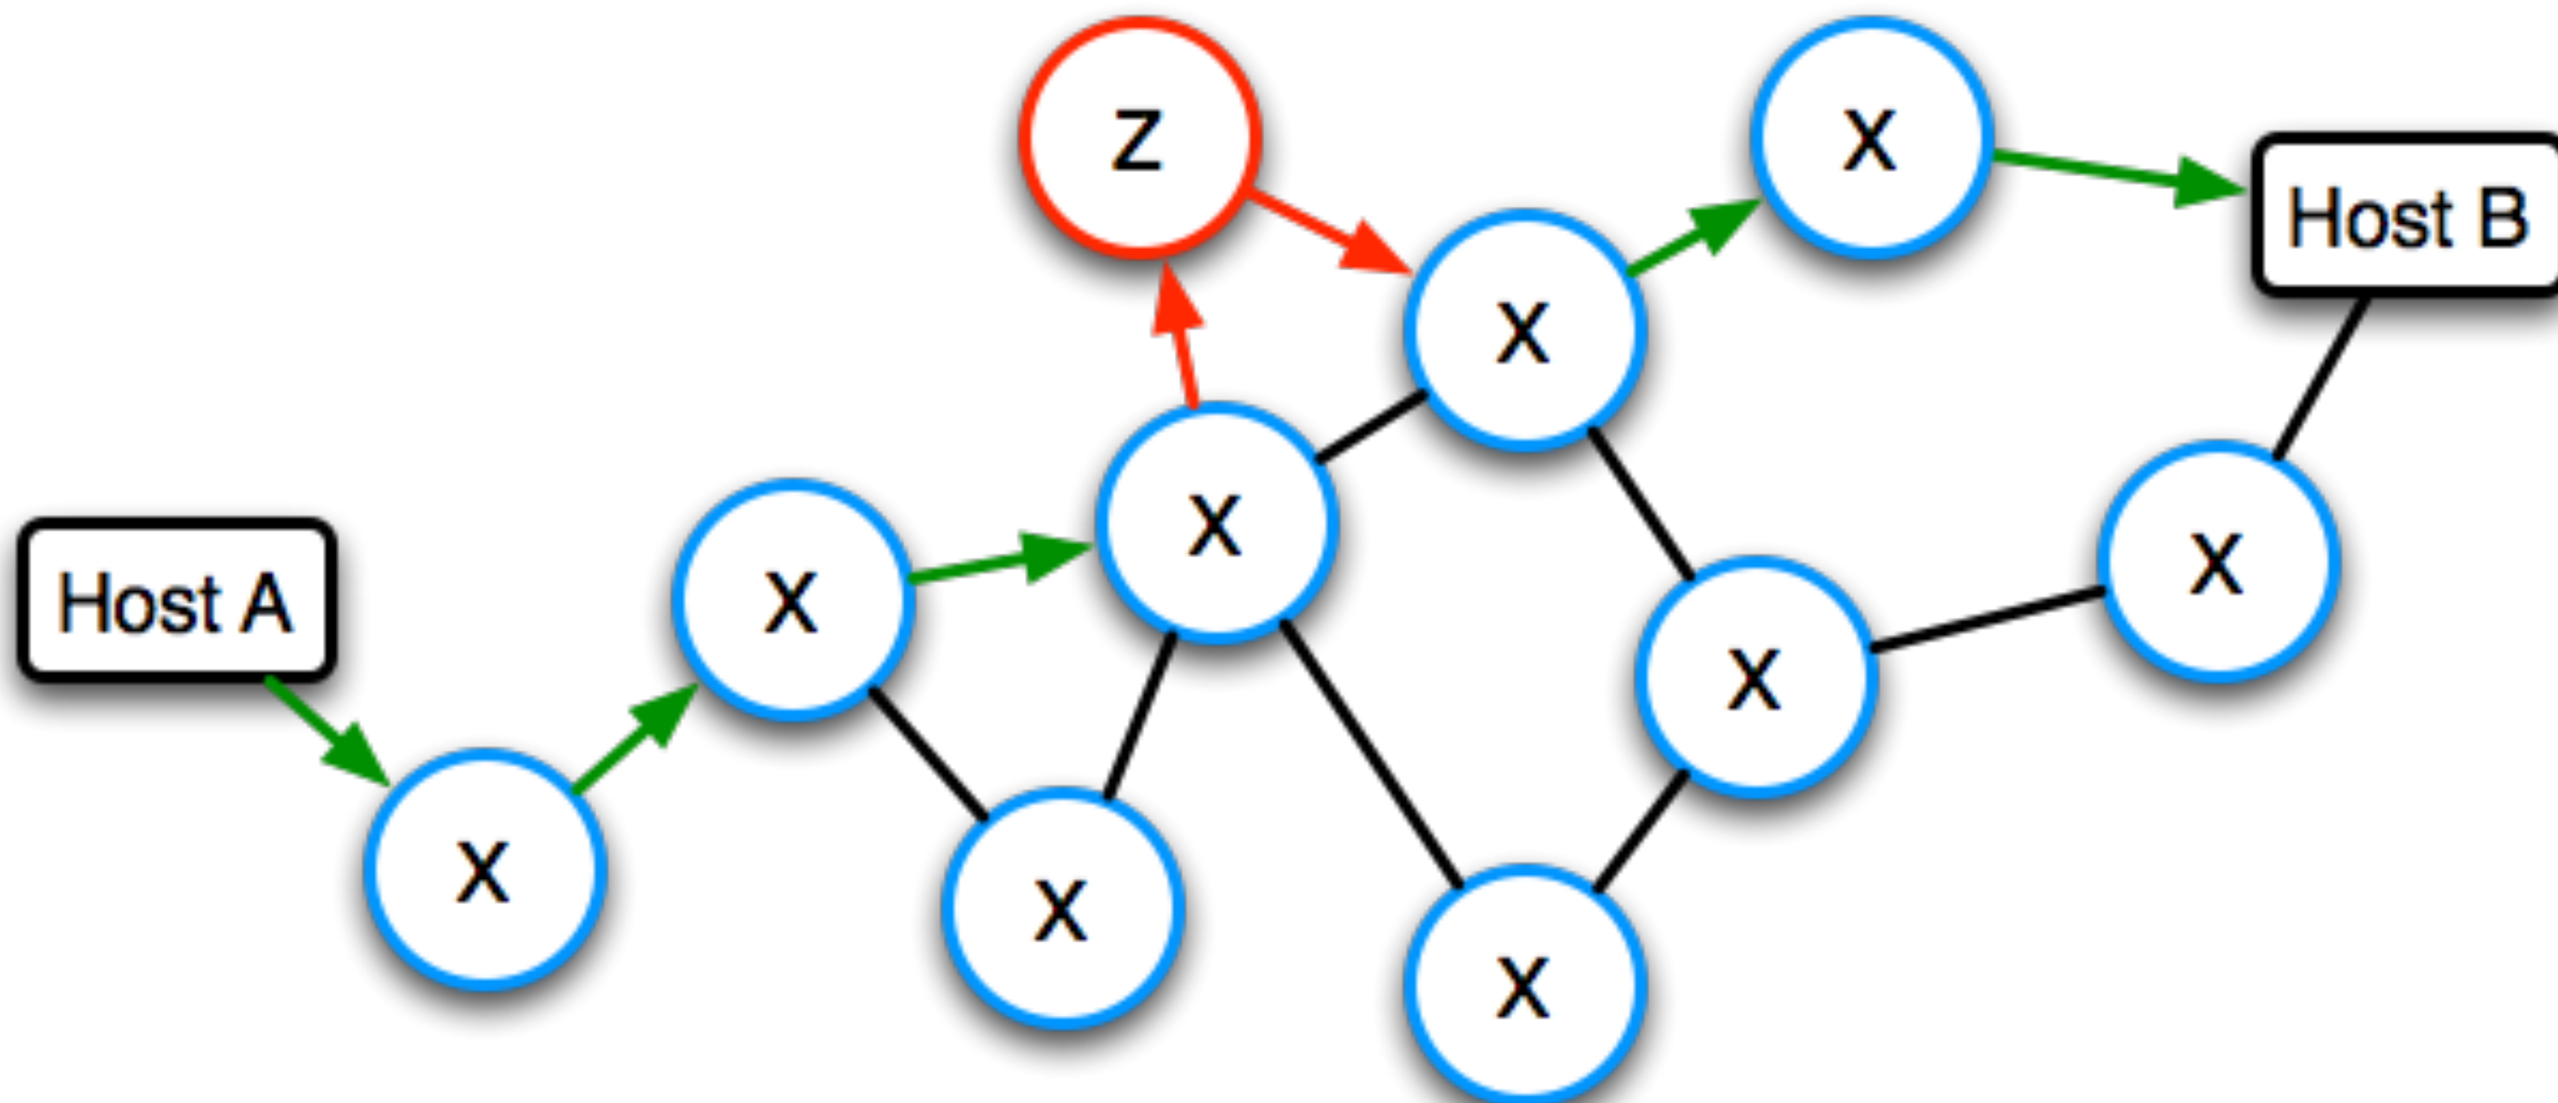

Plan for today

- Administrivia
- Review DNS
- Secure Routing
  - Overview
  - Protocols
  - Attacks
  - Defenses

Routing Protocols

- Routers speak to each other
- They exchange topology and cost information
- Each router calculates the shortest path to each destination
- Routers forward packets along locally shortest path
- Attacker can lie to other routers

Normal Behavior

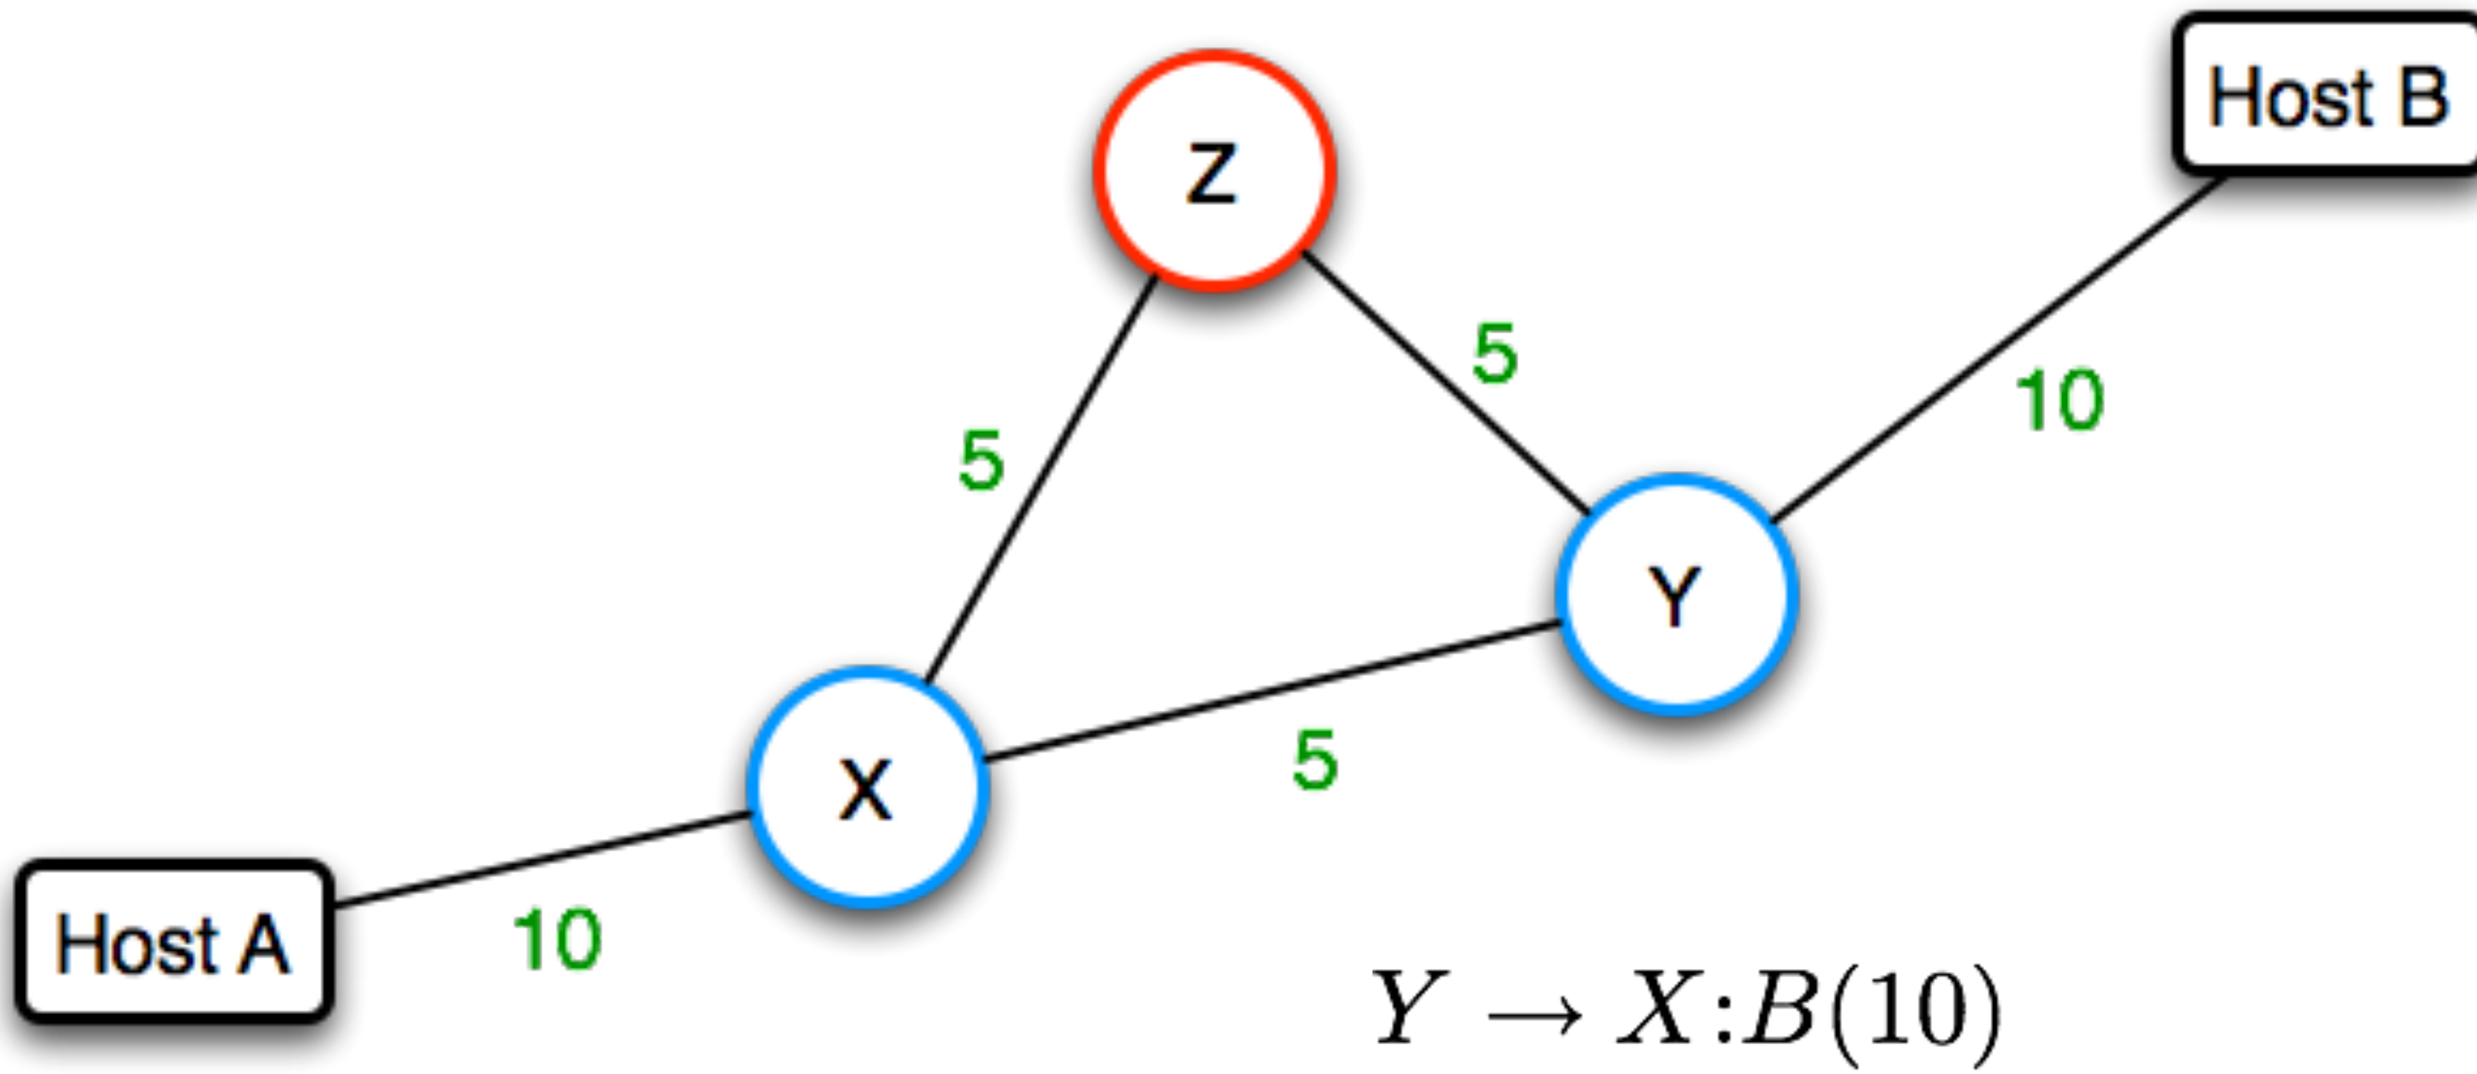

$$Y \rightarrow X : B(10)$$
$$Y \rightarrow Z : B(10)$$
$$Z \rightarrow X : Y(5), B(15)$$
$$X \rightarrow A : Z(5), Y(5), B(15)$$

Malicious Behavior

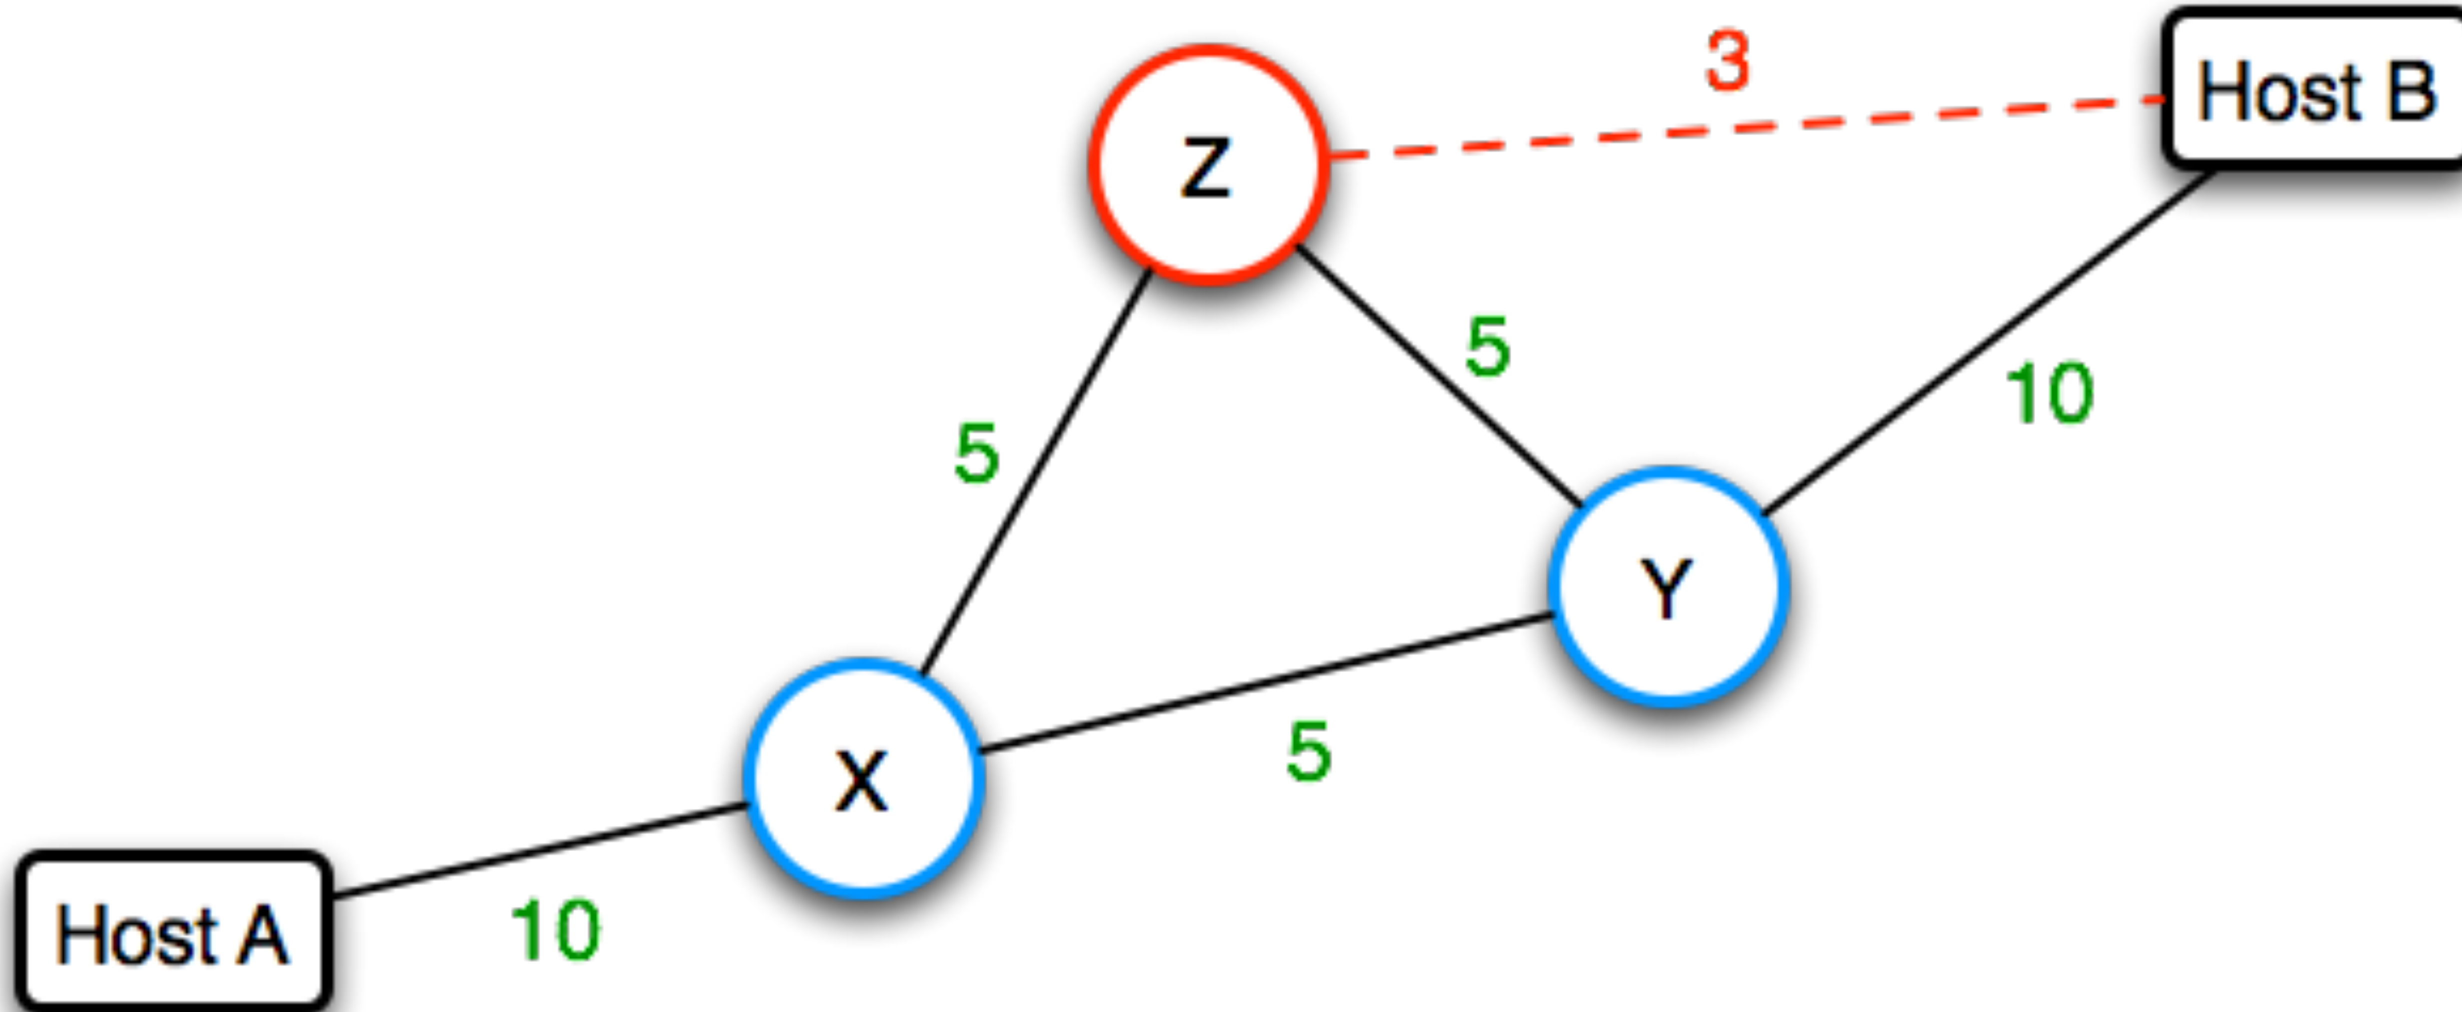

$$Y \rightarrow X : B(10)$$
$$Y \rightarrow Z : B(10)$$
$$Z \rightarrow X : Y(5), B(3)$$
$$X \rightarrow A : Z(5), Y(5), B(8)$$

Why is this difficult?

- X (or Y) has no knowledge of Z's real connectivity.
- The problem isn't the link from X to Z:
  - The problem is the lack of integrity of the info being sent
  - Non-trivial complexity: Z might be deceived by some other neighbor Q

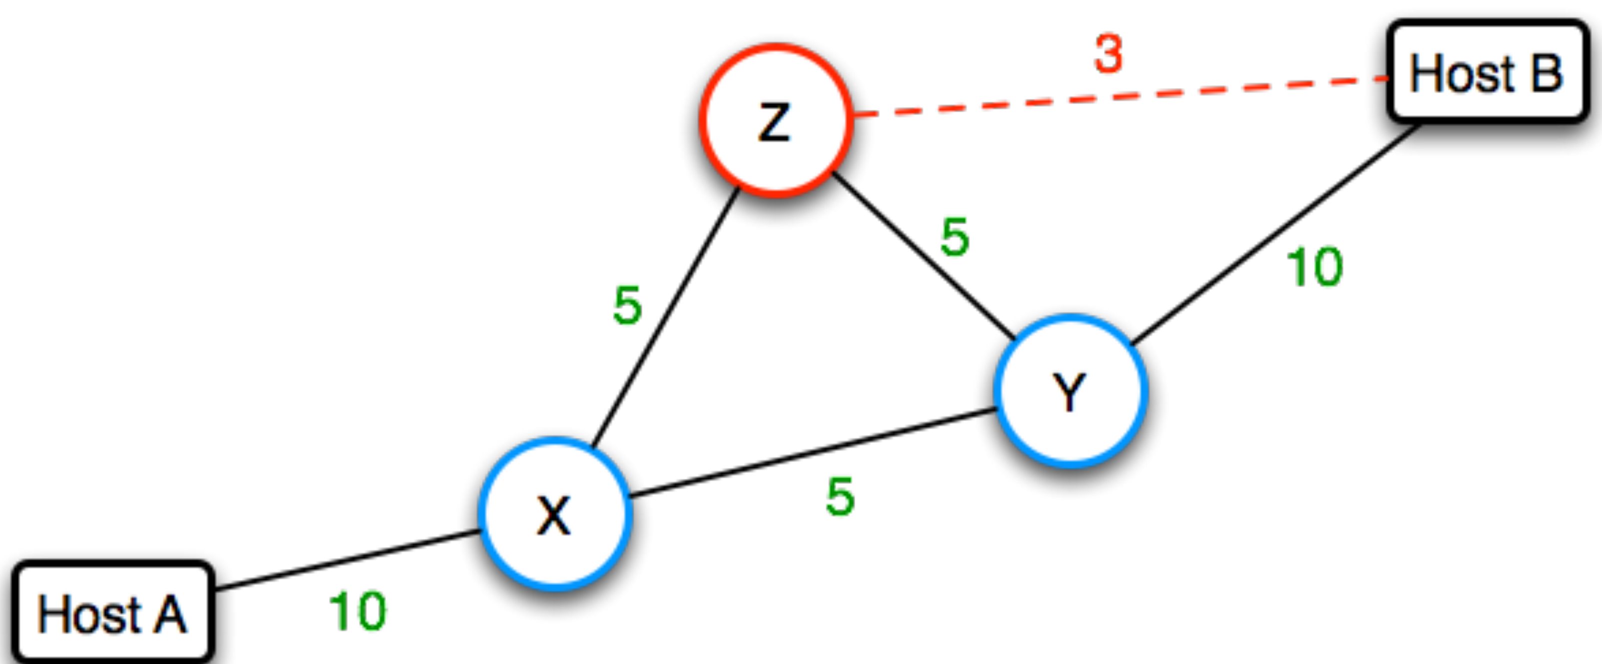

$$Y \rightarrow X{:}B(10)$$
$$Y \rightarrow Z{:}B(10)$$
$$Z \rightarrow X{:}Y(5), B(3)$$
$$X \rightarrow A{:}Z(5), Y(5), B(8)$$

Internet Routing

- Two flavors: internal and external
  - **Intradomain** - Internal (within ISP, company): primarily OSPF.
  - **Interdomain** - External (between ISPs, and some customers): BGP.

Internal Networks

- Common management
- Common agreement on cost metrics
- ISPs have very specialized topologies and well-controlled networks

OSPF (Open Shortest Path First)

- Each node announces its own connectivity.
- Announcements include link cost
  - Each node re-announces all information received from peers.
  - Every node learns the full map of the network.
  - Each node calculates the shortest path to all destinations (e.g., via Dijkstra's).
  - *Scalability*: limited to a few thousand nodes at most.

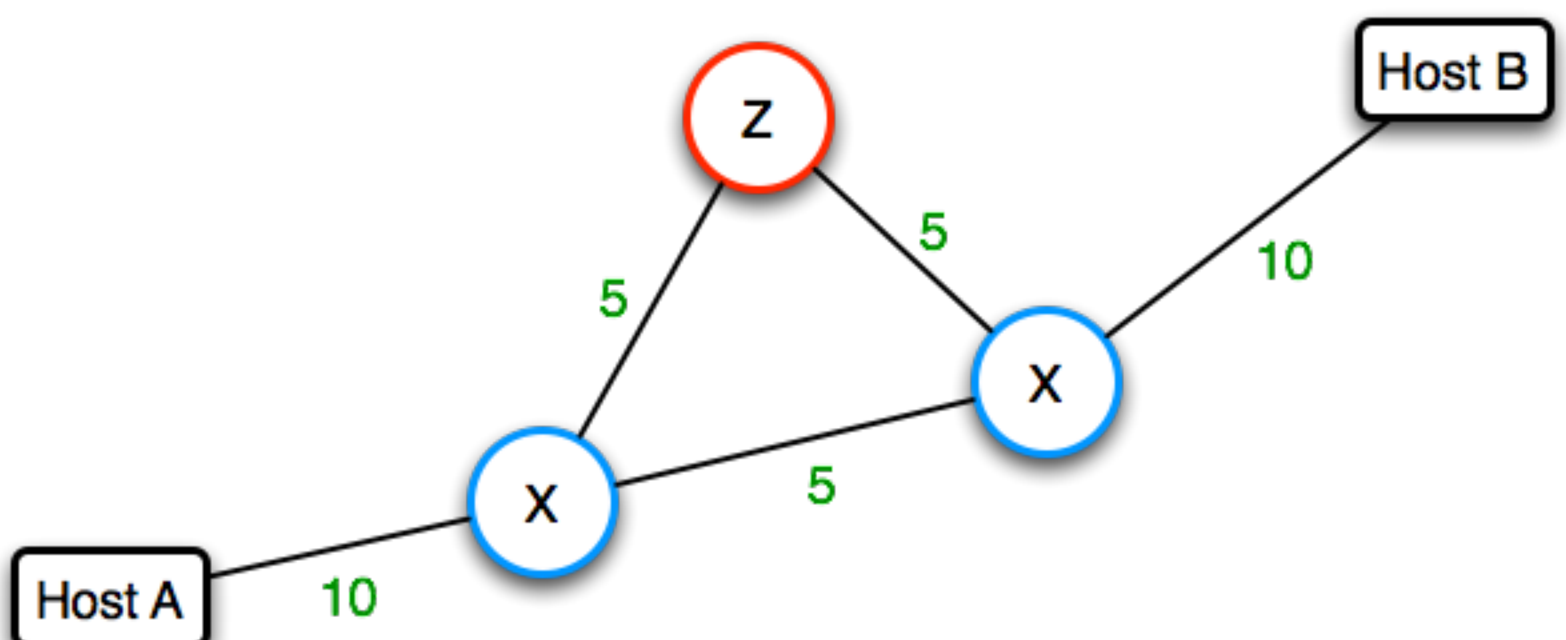

Border Gateway Protocol (BGP)

- BGP routes information at the **autonomous system** level
- BGP is (mostly) a **path vector protocol**
  - Routing tables include path necessary to reach destination
  - Vectors communicated amongst routers

Routing in a nutshell

- The Internet...

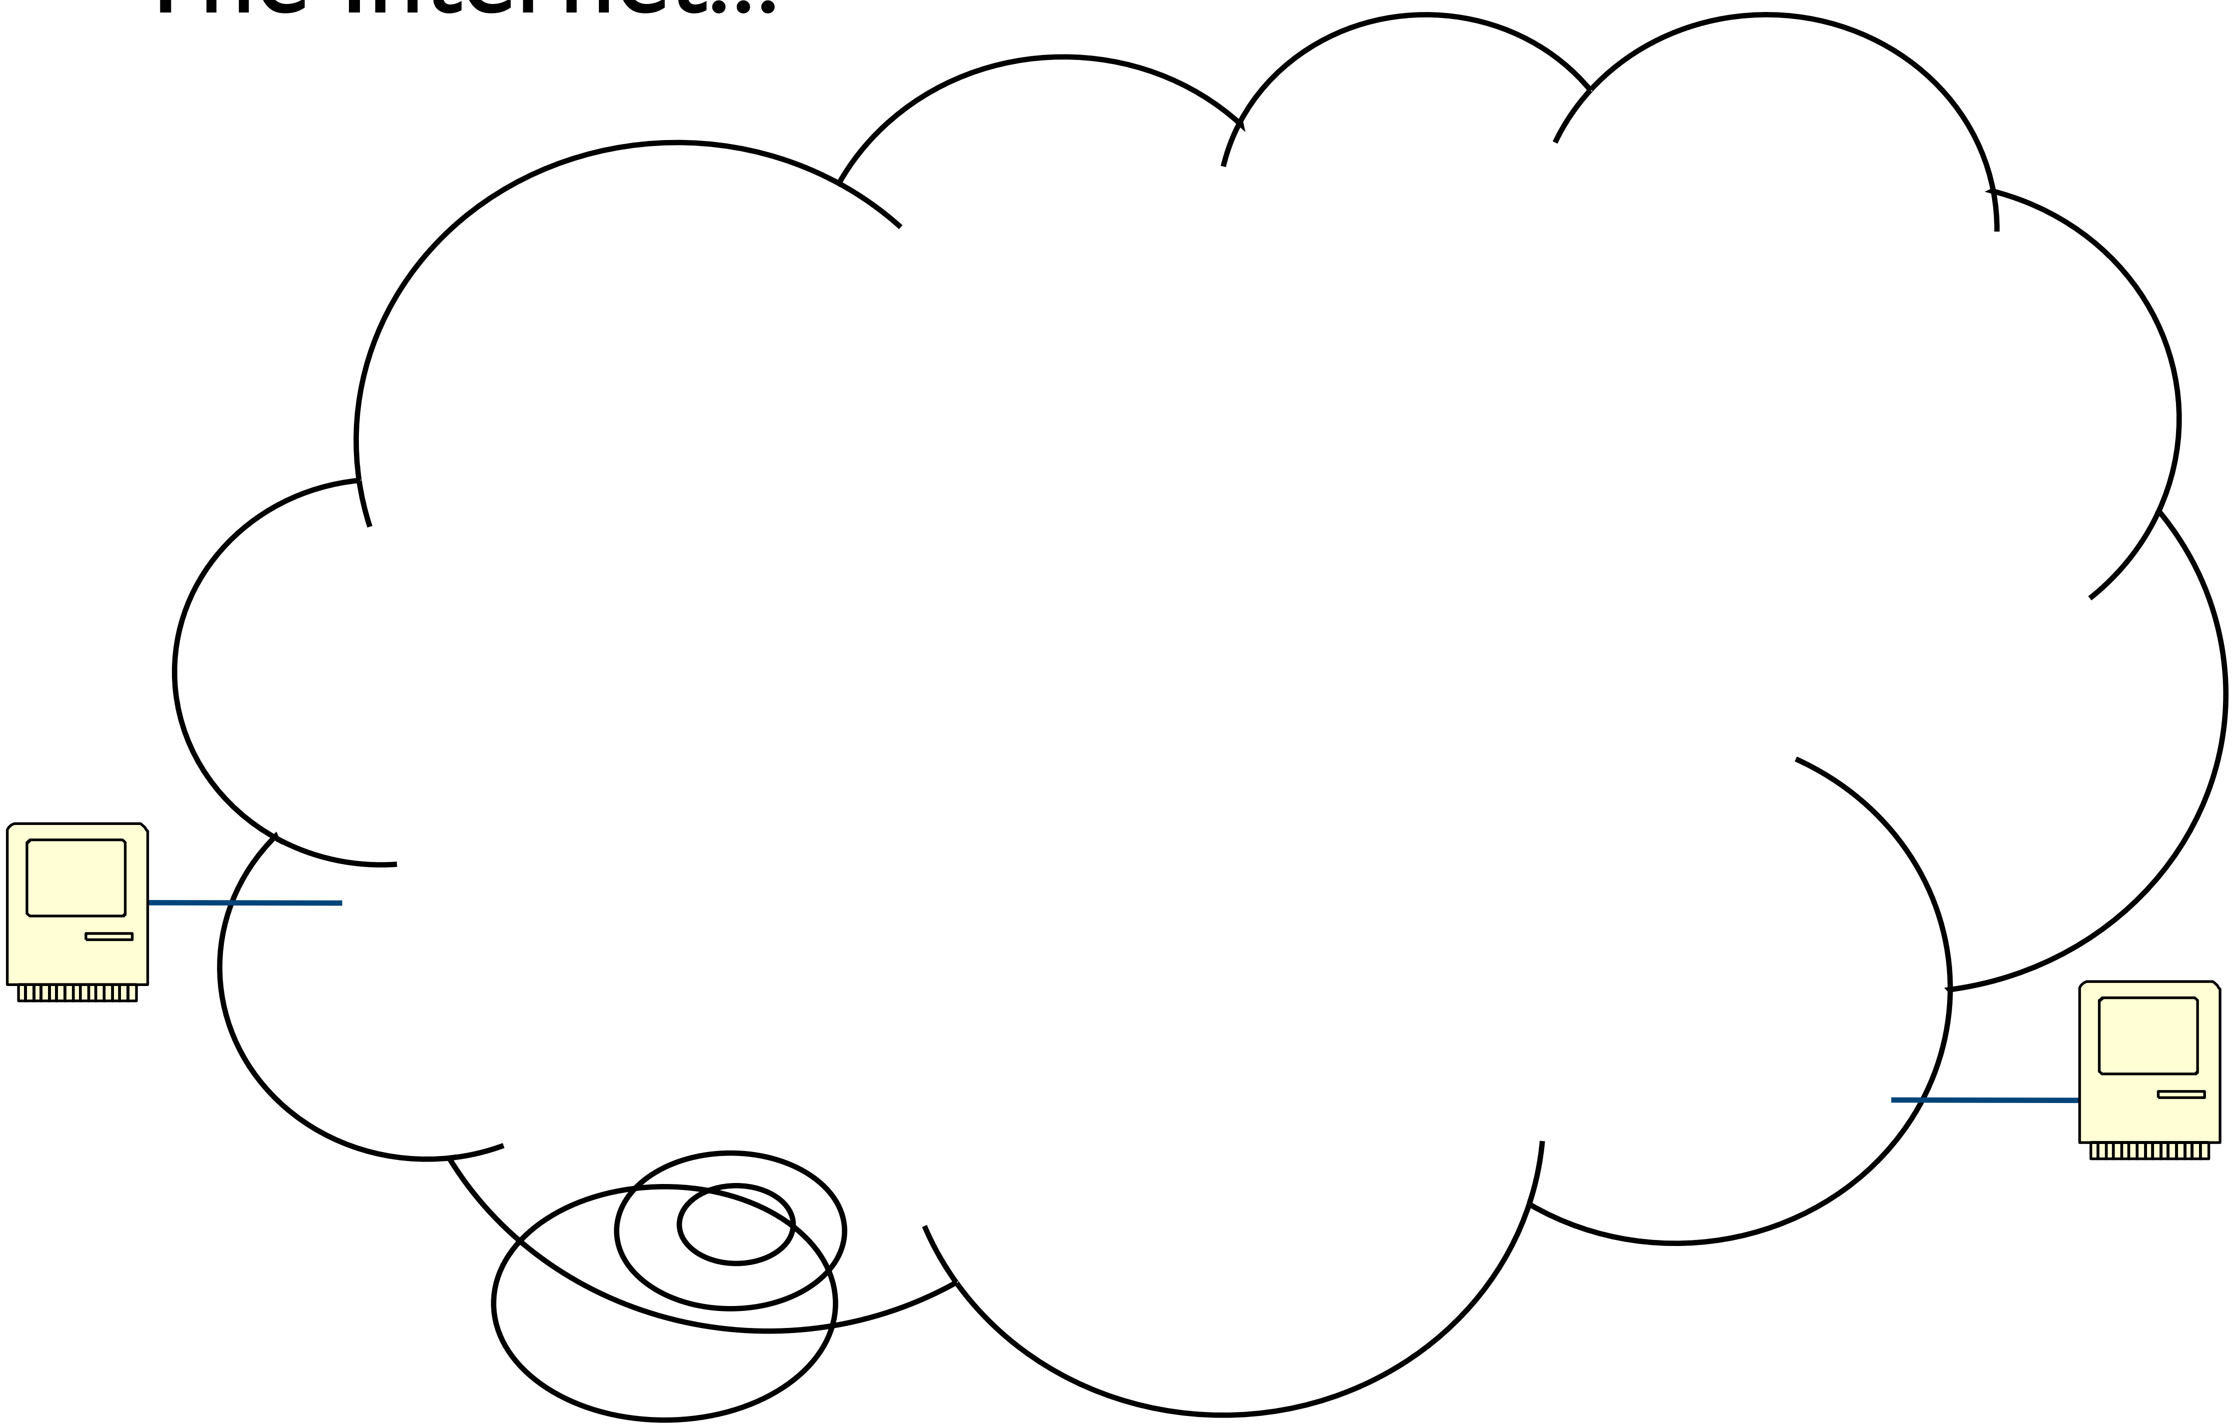

Routing in a nutshell

- …is made up of Autonomous Systems (ASes)…

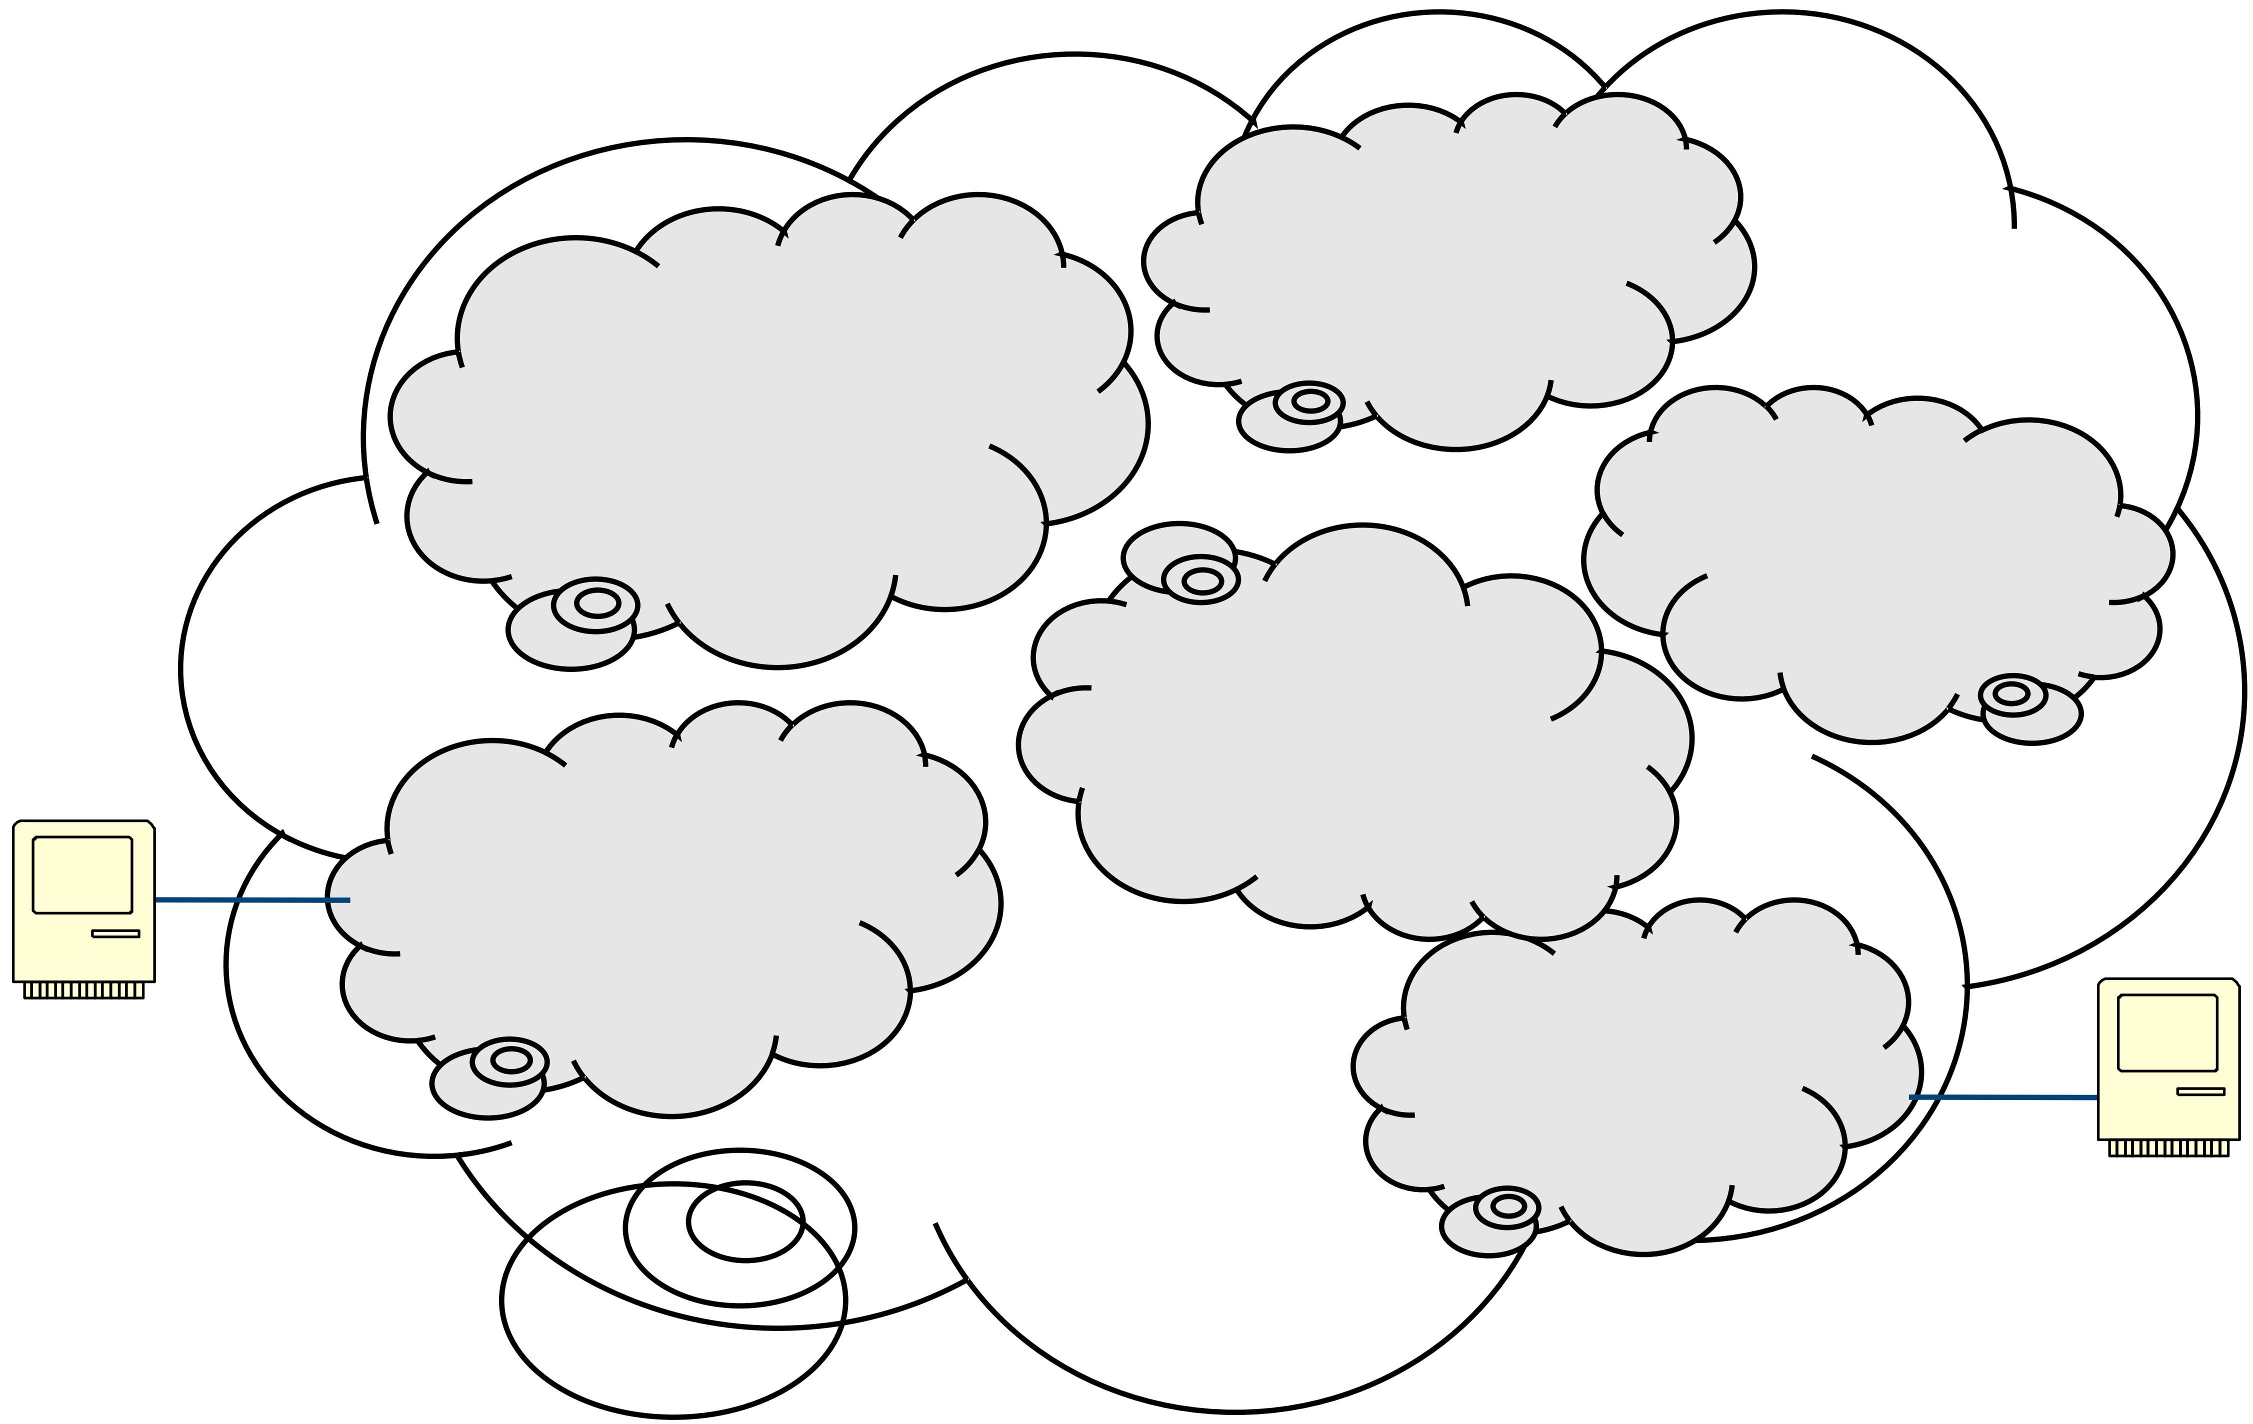

…linked at Border Routers.

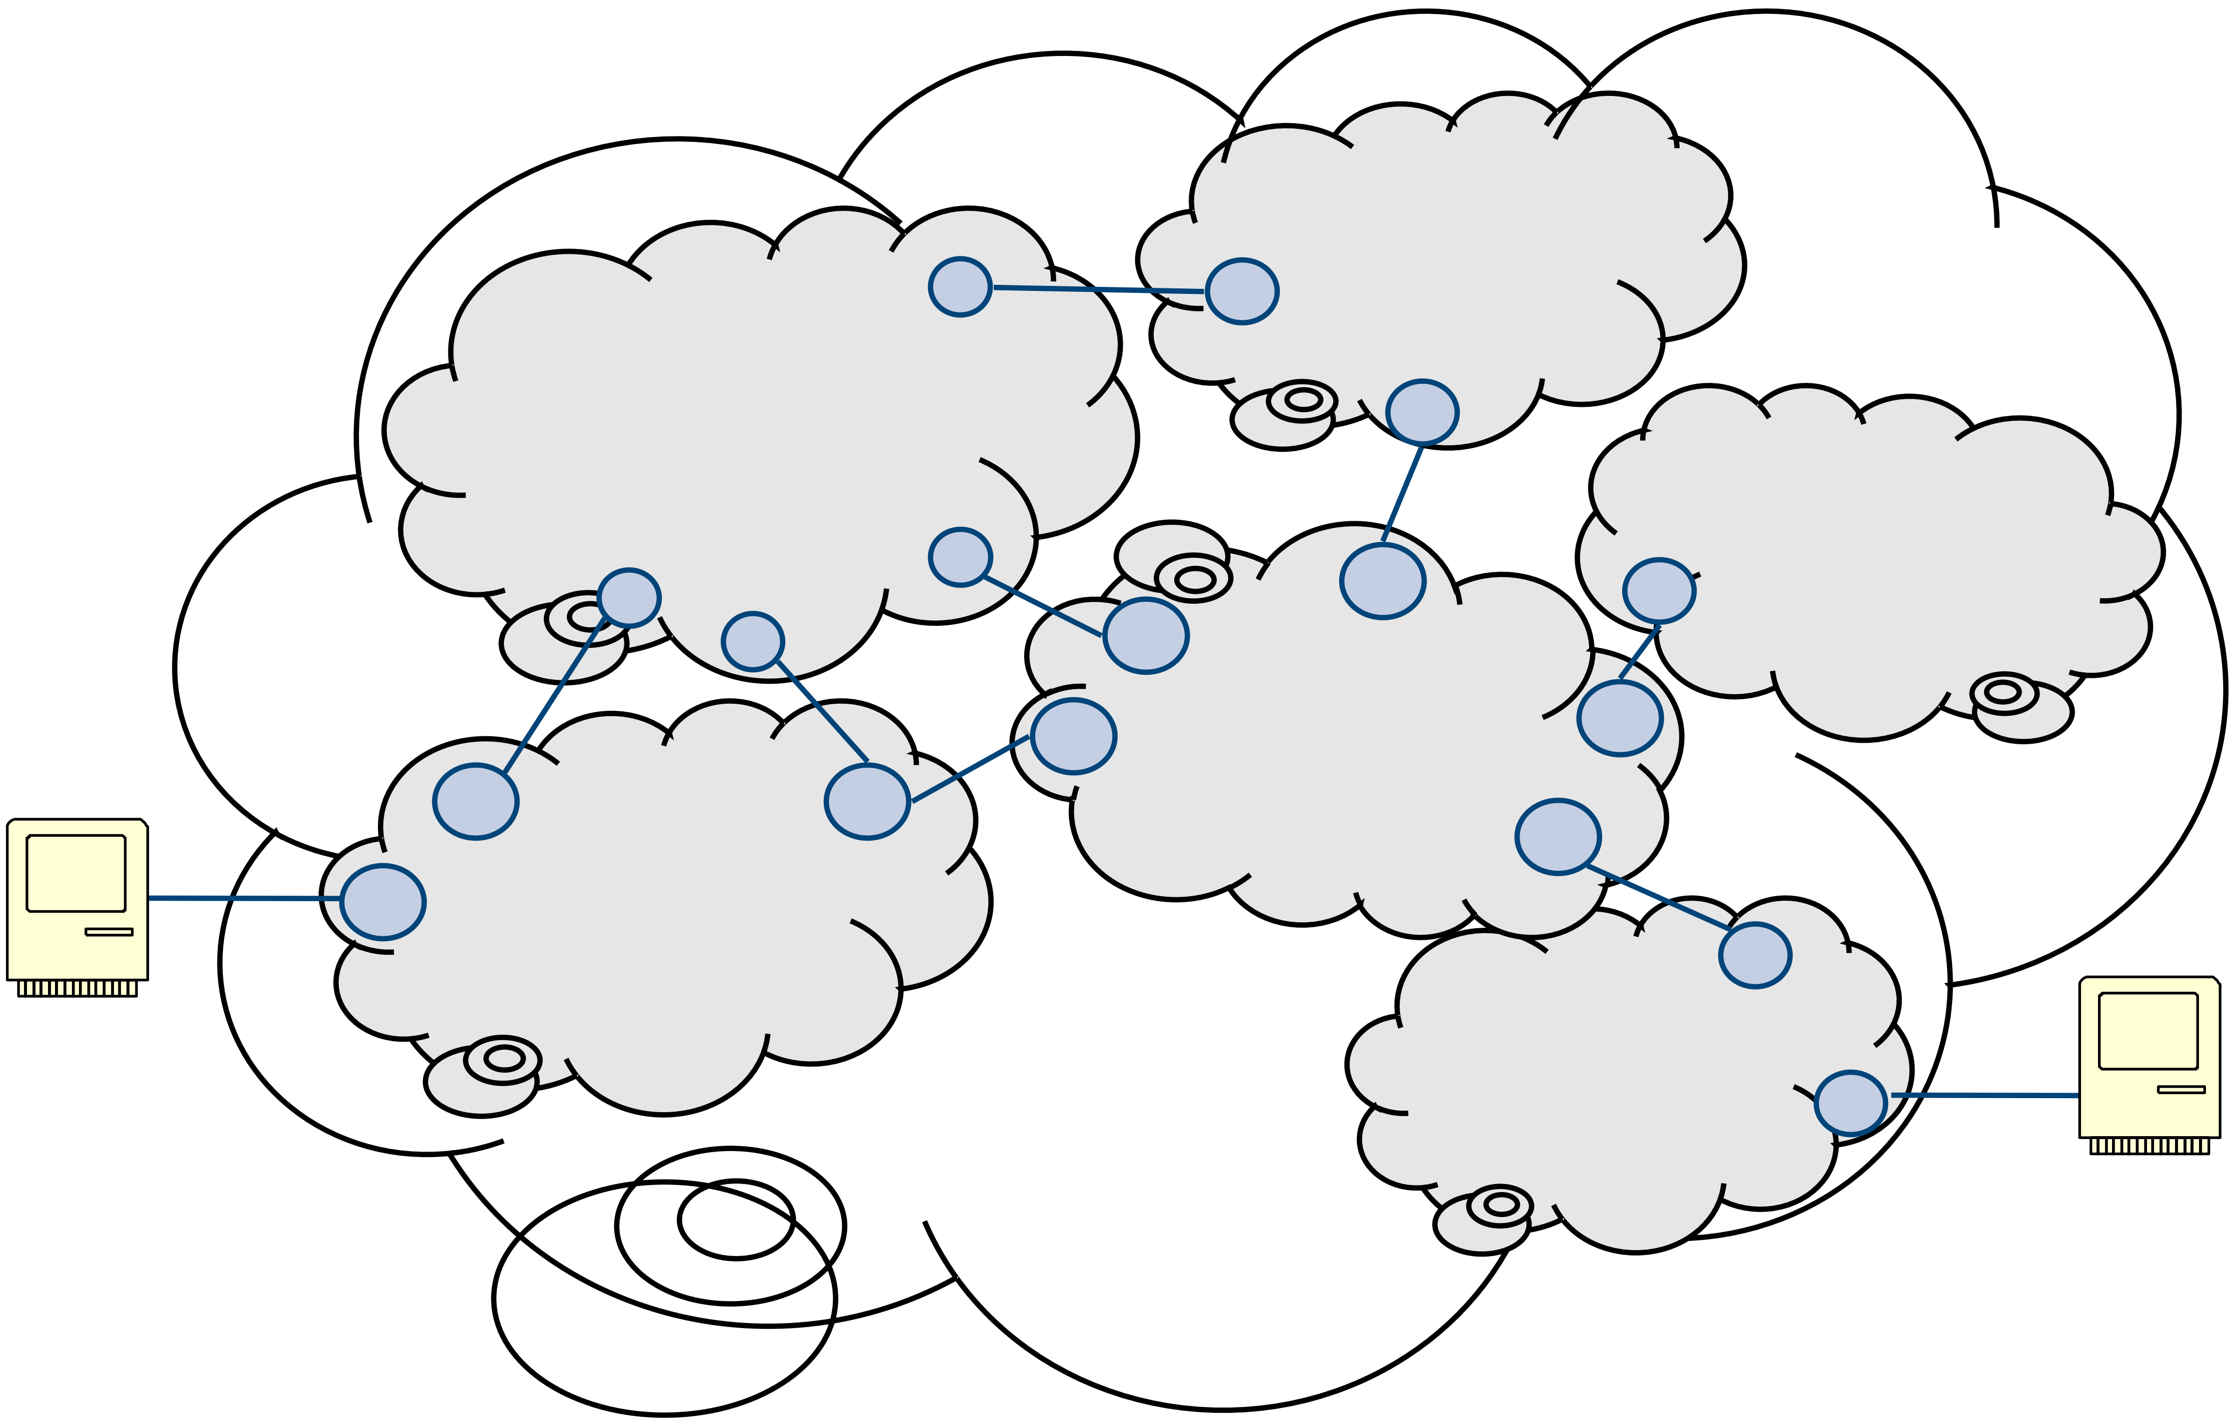

BGP determines which ASes to follow from source to destination

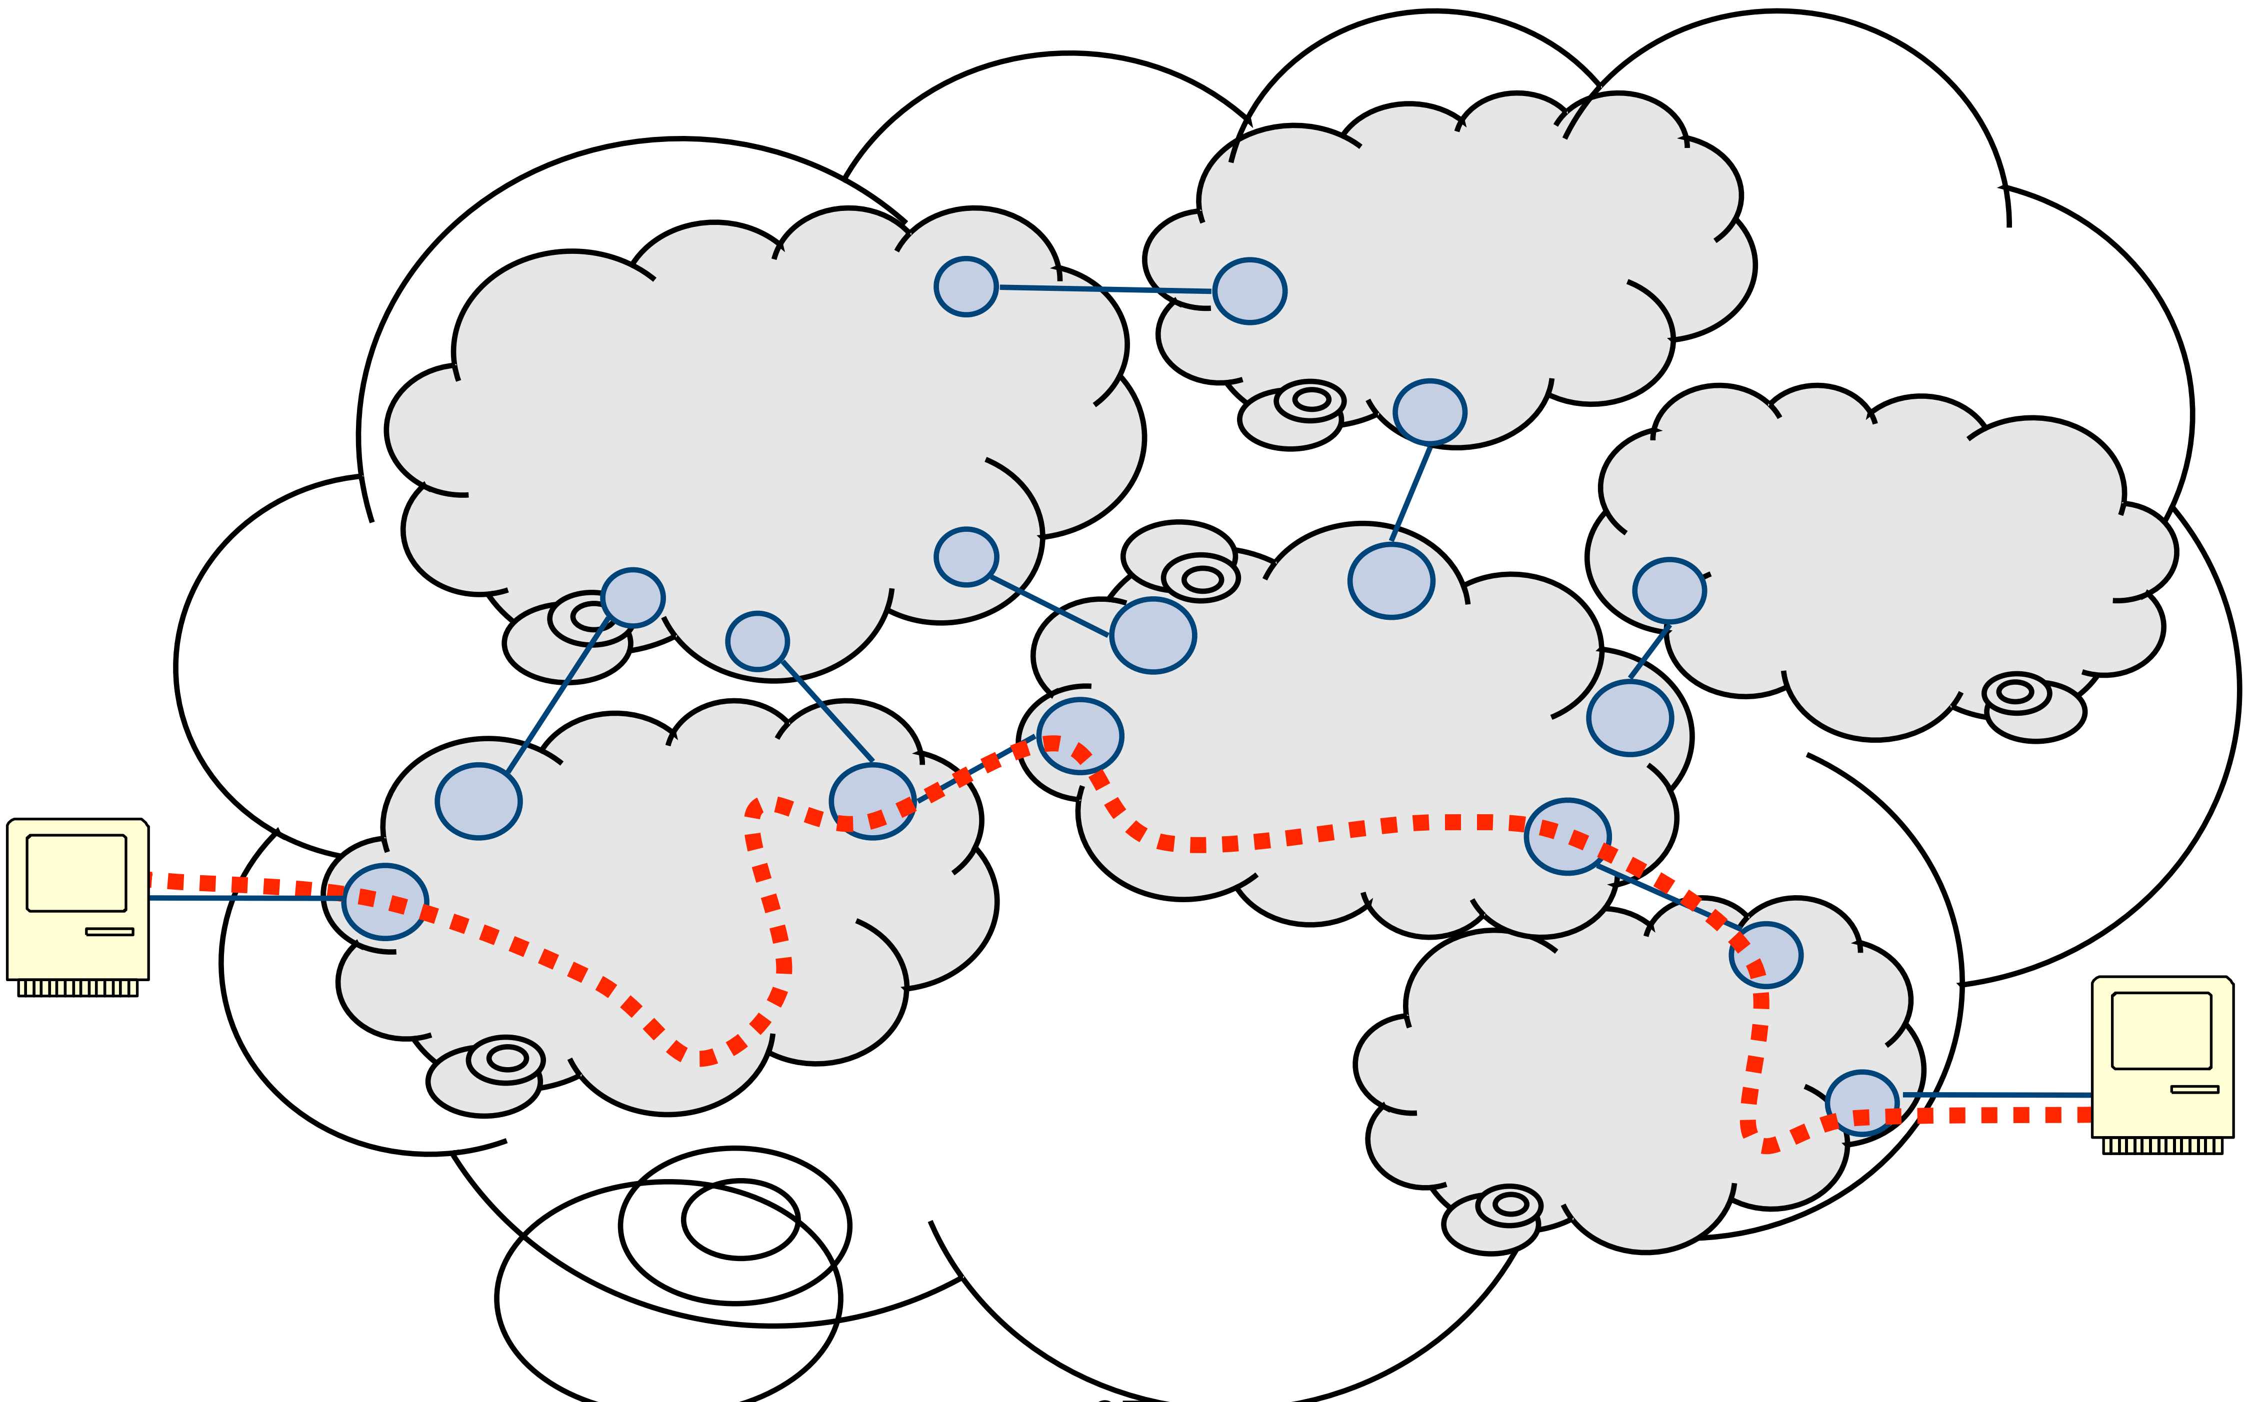

- Each AS is responsible for moving packets inside it.
- Intra-AS routing is (mostly) independent from Inter-AS routing.

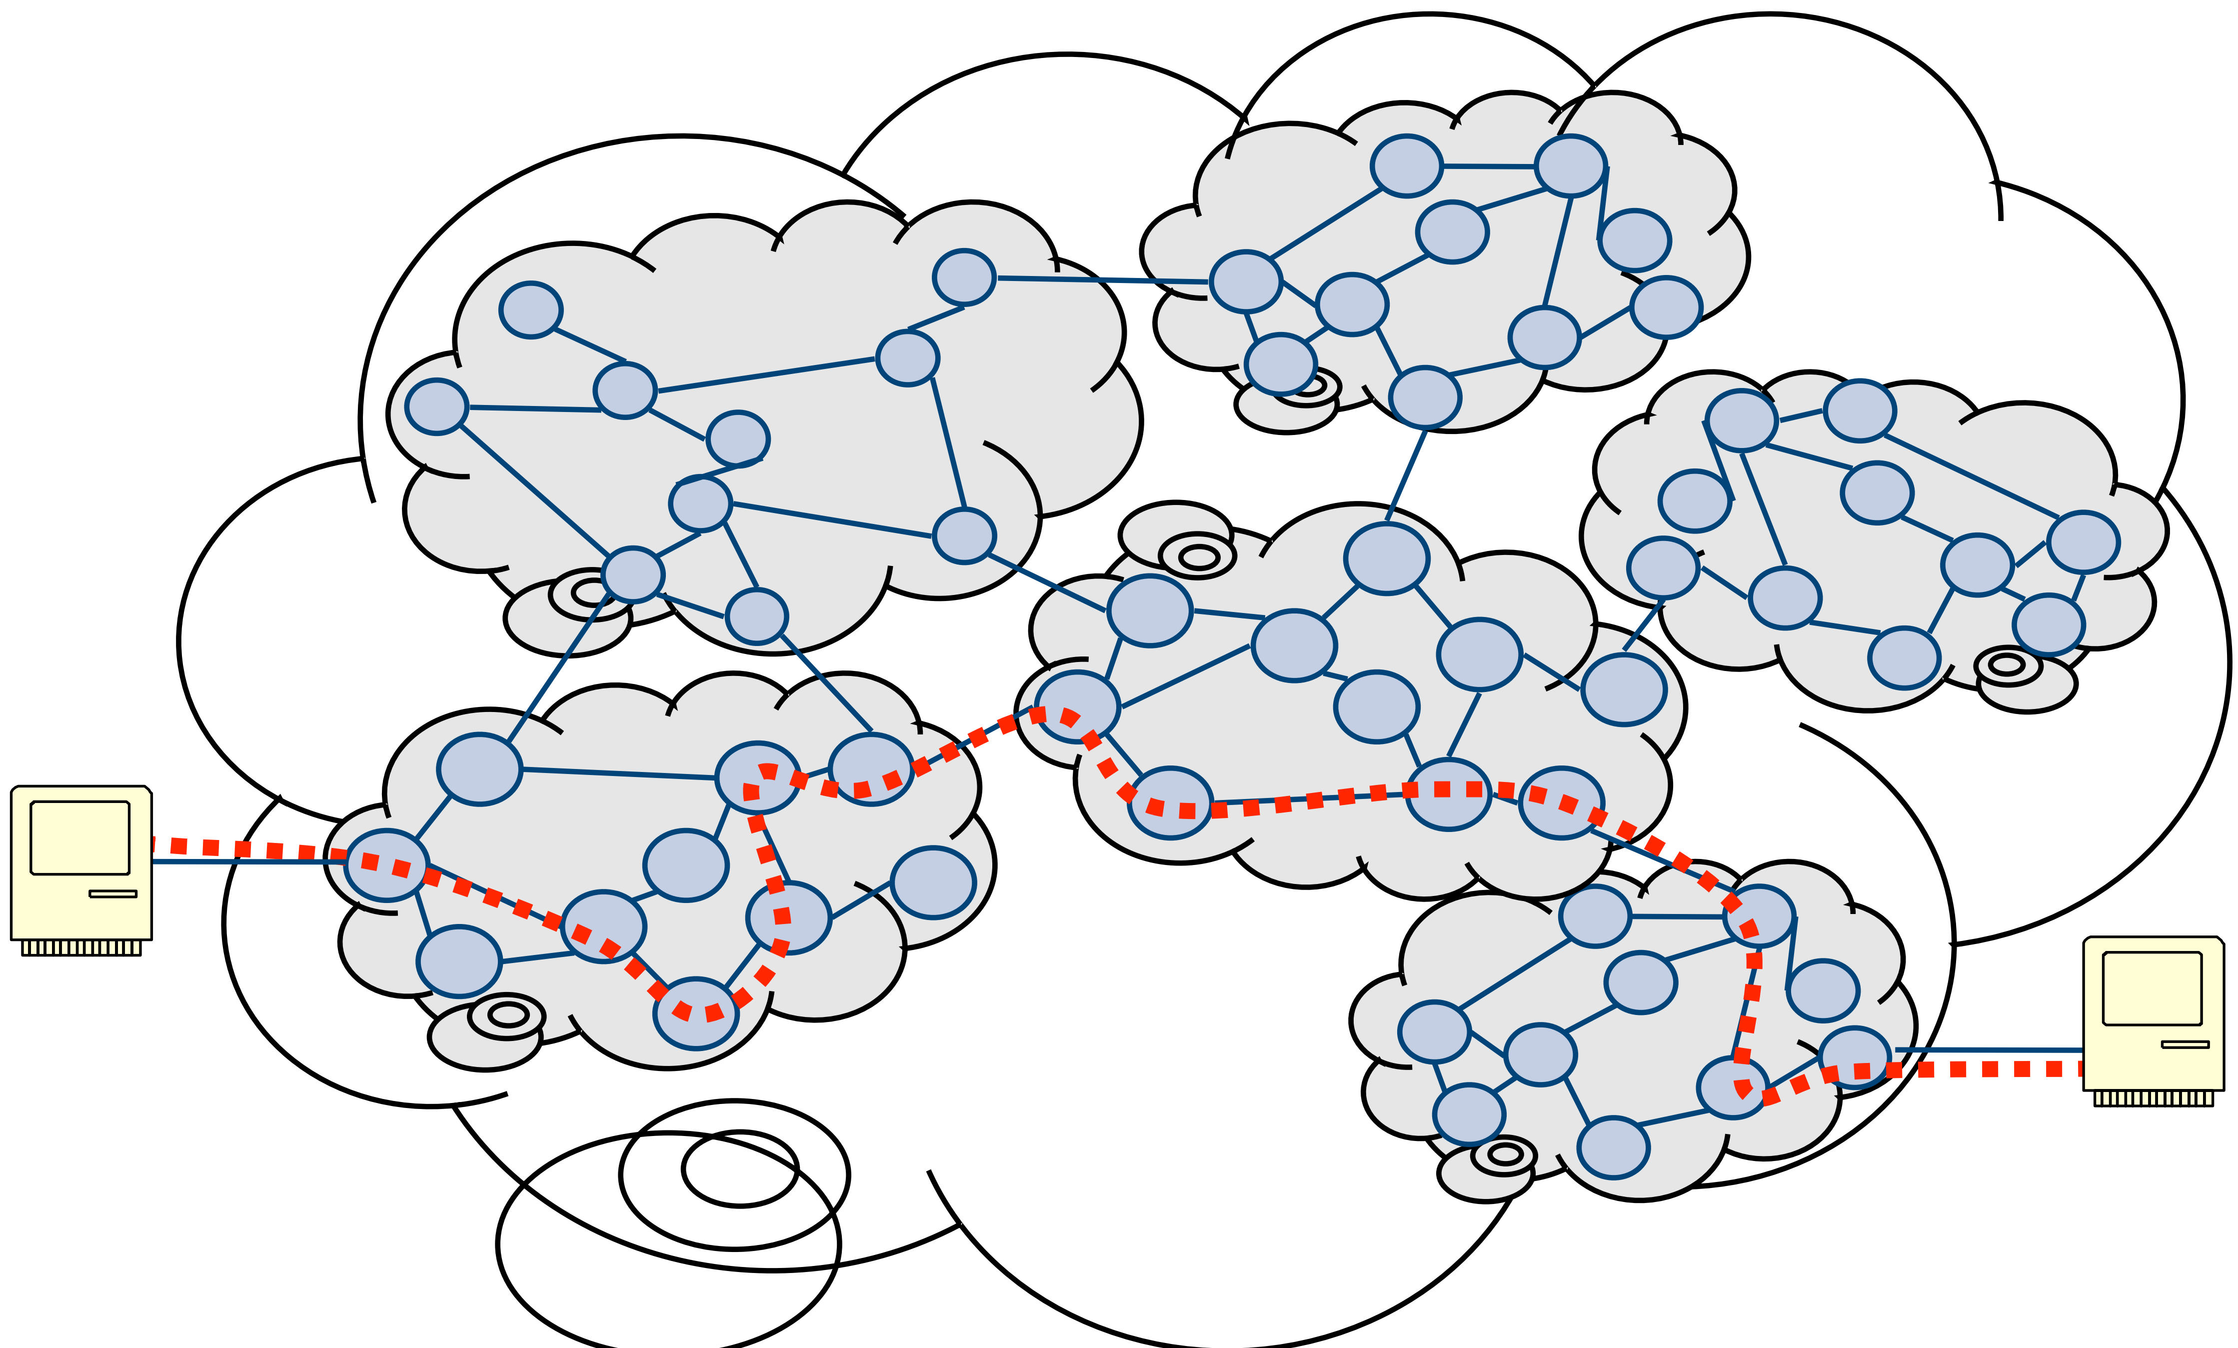

The BGP Protocol

- BGP messages
  - **Origin** announcements:
    - "I own this block of addresses"
  - Route **advertisements**:
    - "To get to this address block, send packets destined for it to me. And by the way, here is the path of ASes it will take"
  - Route **withdrawals**:
    - "Remember the route to this address block I told you about, that path of ASes no longer works"
- Route decisions
  - Border routers receive origin announcements/route advertisements from their peers
  - They choose the "best" path and send their selection downstream
- BGP Attributes
  - BGP messages have additional attributes to help routers choose the "best" path

https://engineering.fb.com/2021/10/05/networking-traffic/outage-details/

CIDR Block	Path			Attributes
123.125.28.0/24	768	4014	664	bkup

How Internet Routing Works

- Select next-hop based on longest-prefix match
- In case of ties, in order of preference (most to least):
  - ASes tend to prefer **customers** or **peers** over **providers** (because $$$)
  - In case of tie (see previous bullet), ASes tend to prefer shortest AS path

Plan for today

- Administrivia
- Review DNS
- Secure Routing
  - Overview
  - Protocols
  - Attacks
  - Defenses

BGP Attacks

Later: Defenses

Attack: Prefix Hijacking

- An attacker can claim to originate a known prefix
- For example, Tufts could decide to be AT&T for a day, and advertise 12.0.0.0/8
- **Route filtering** (where does route advertisement come from?) should catch this, but many operators do not perform proper filtering policy within their AS

- If another AS advertises one of our prefixes, bad things happen:

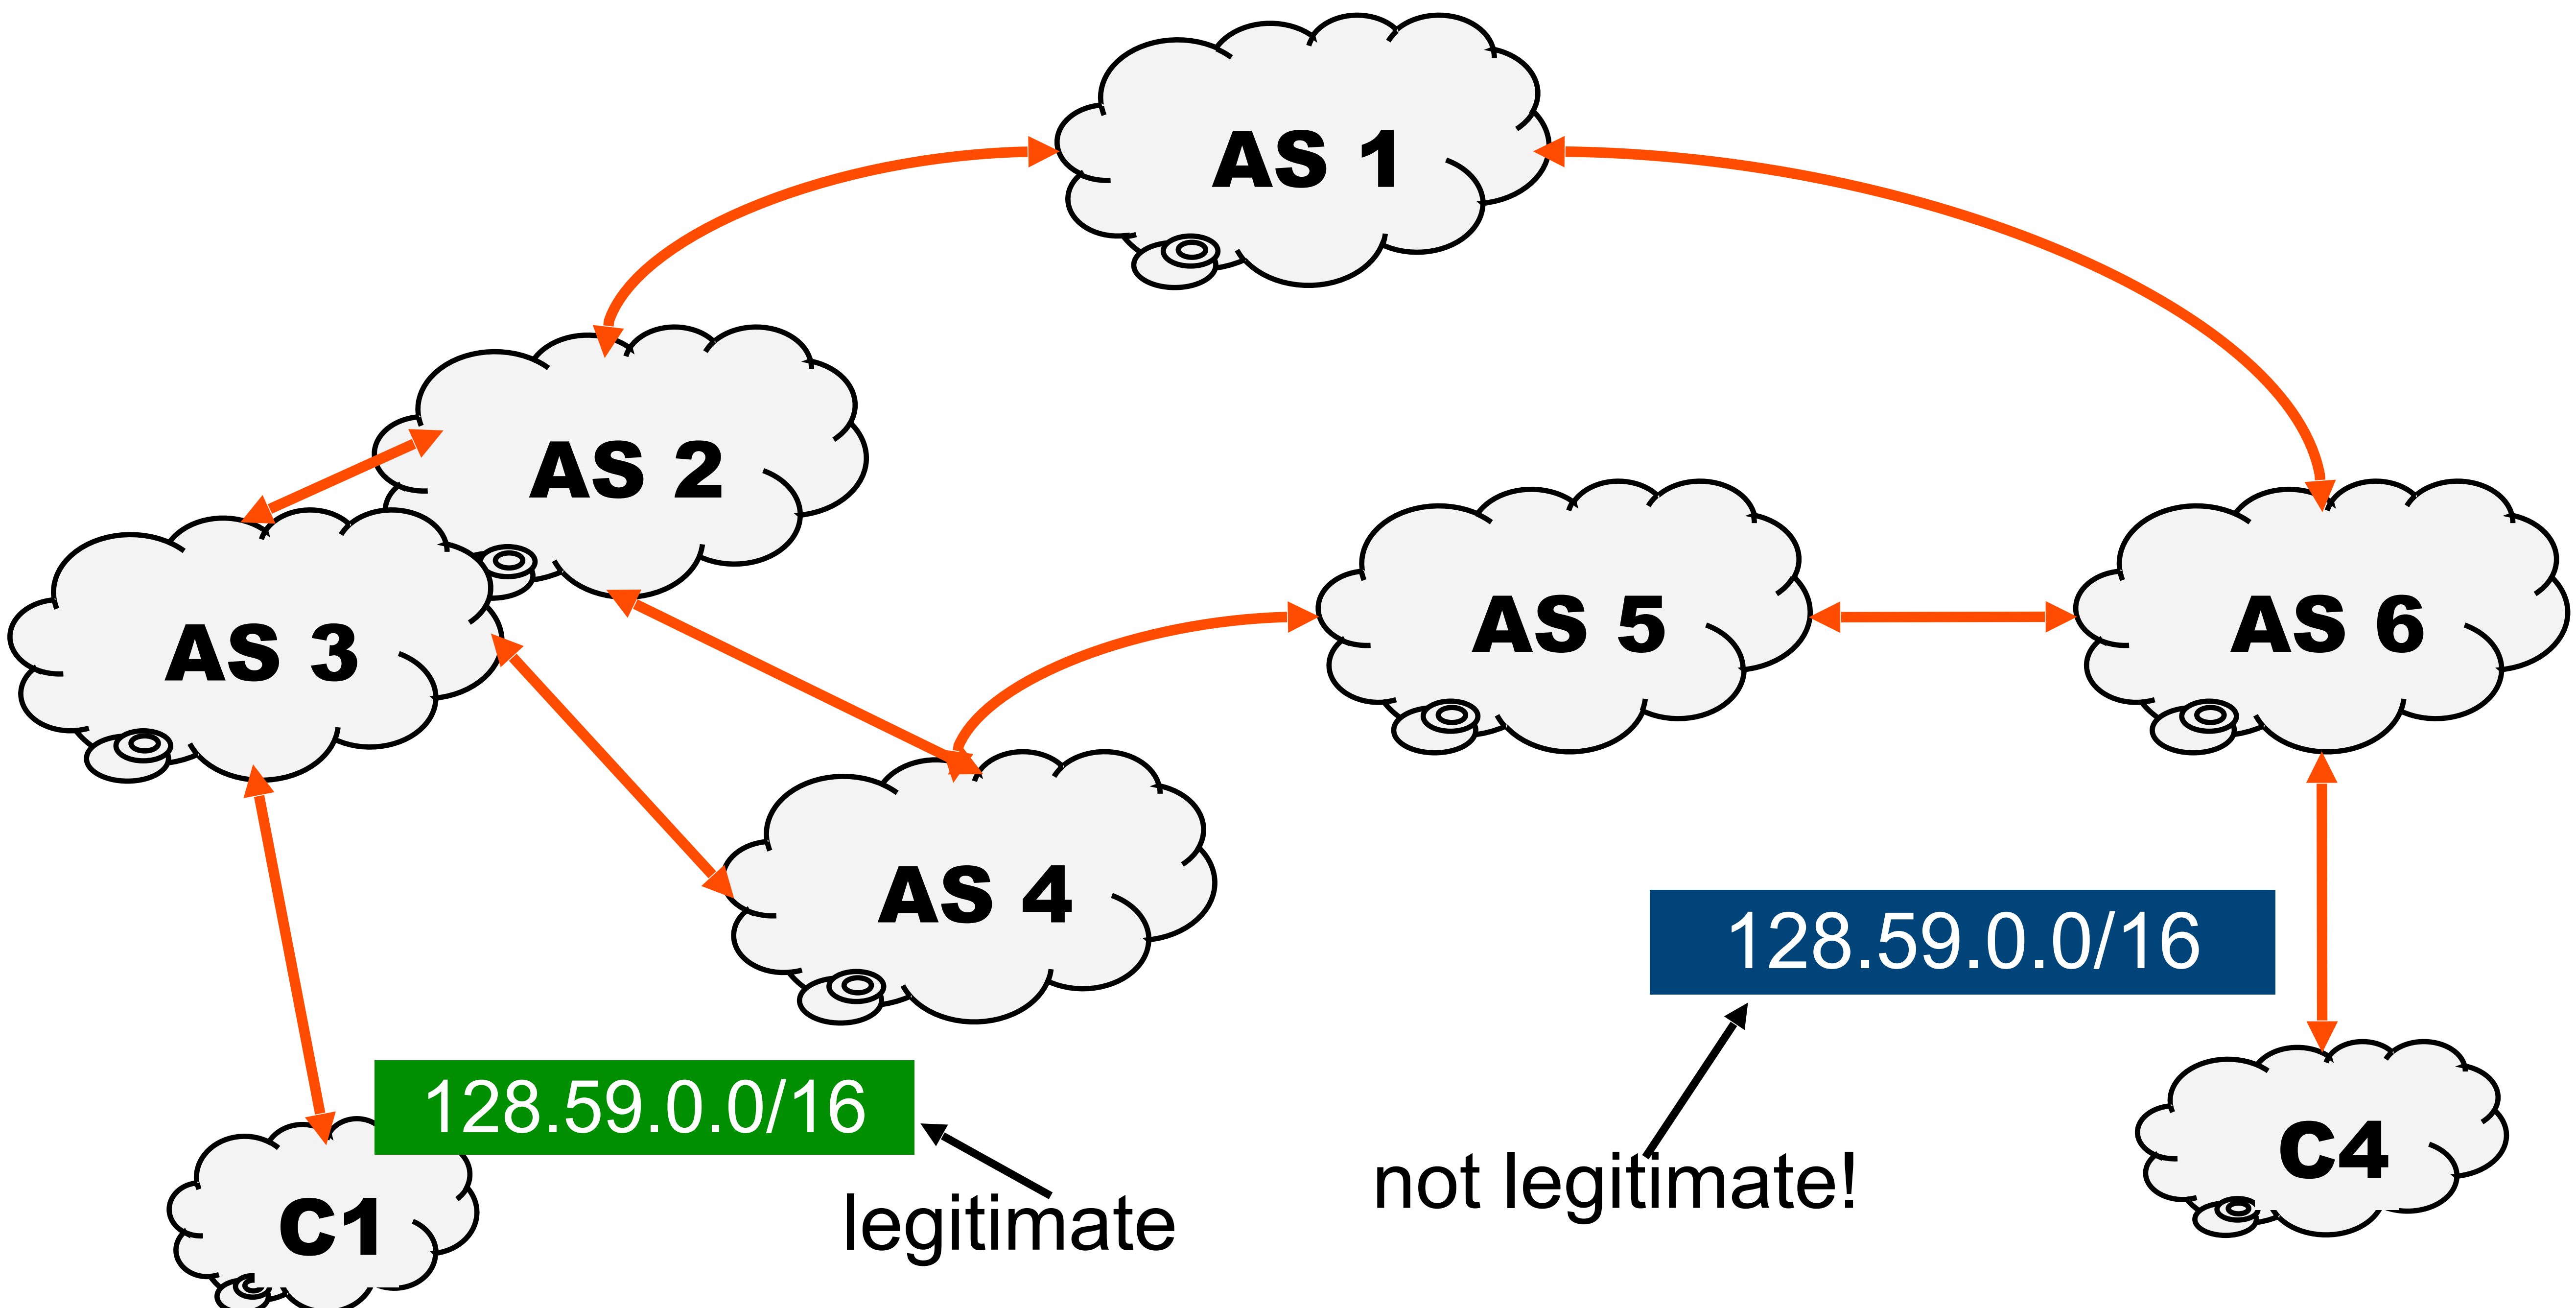

- Prefix becomes unreachable from the part of the net believing C4's announcement.

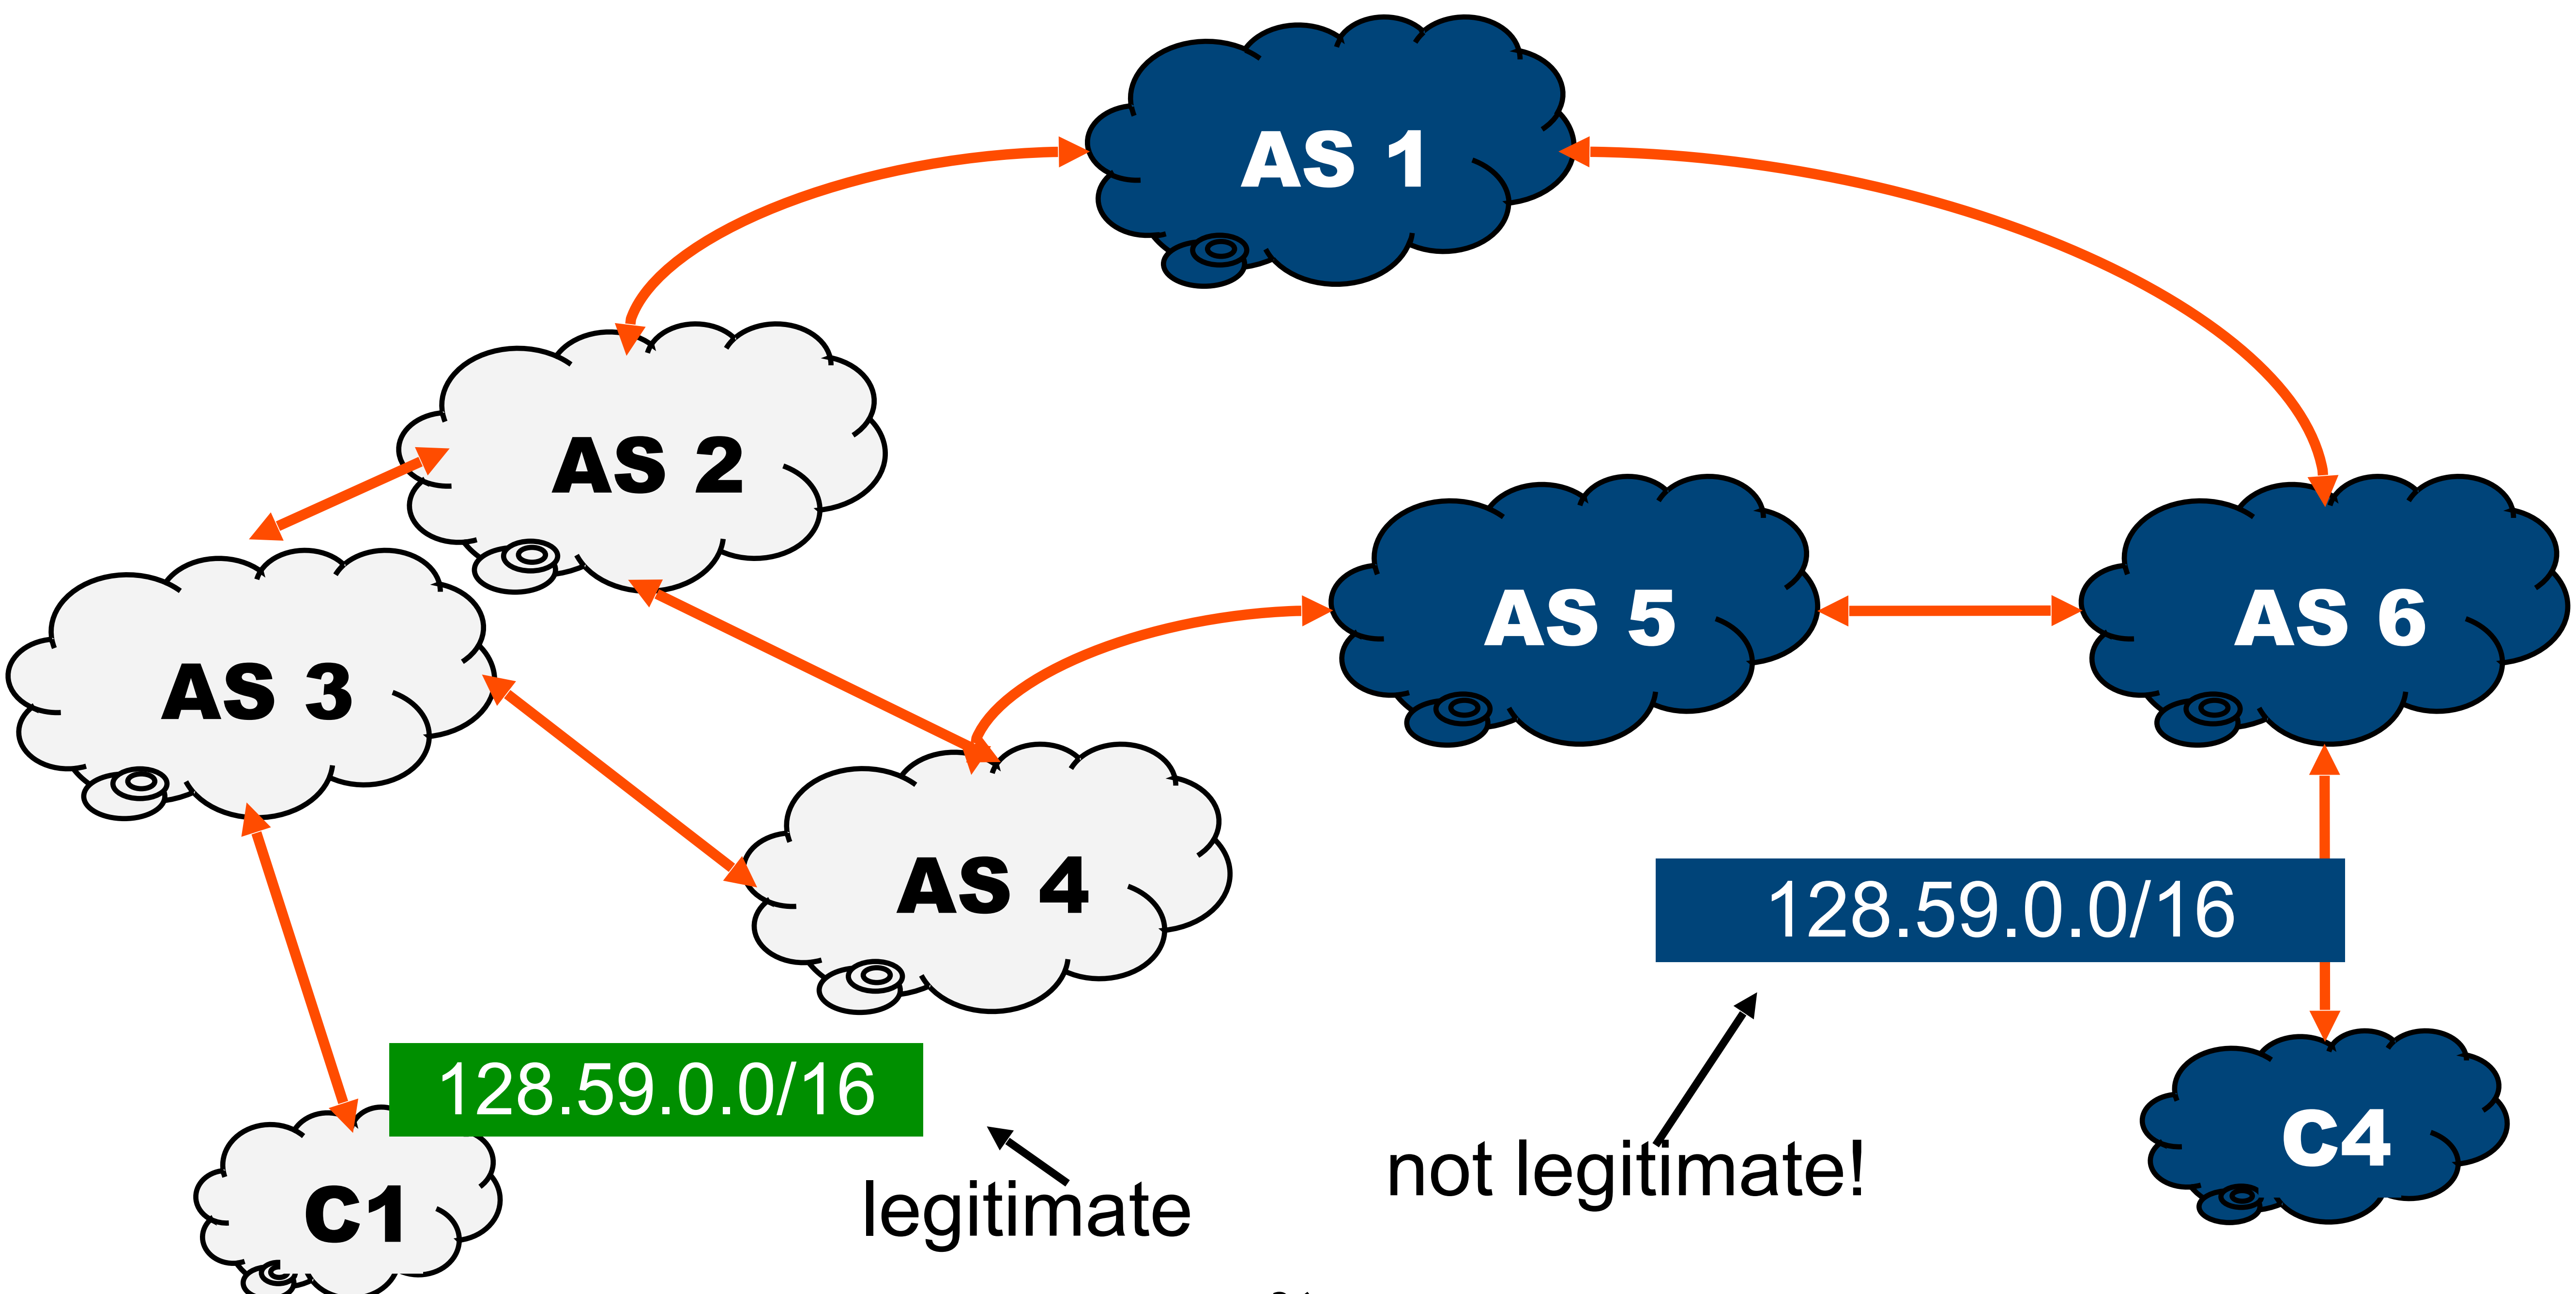

Attack: Longest-Prefix Matching

- Within the AS, a prefix can be broken into smaller blocks and advertised as such
- Because of **longest-prefix matching**, these will be preferred (eg. 12.10.8.0/24 is preferred over 12.0.0.0/8 because it is more specific)
- Attacker can get clever (say 100.200.0.0/16 is targeted IP block)
  - Attacker sends origin announcement for 100.200.0.0/17 and 10.200.129.0/17 (covers all of 100.200.0.0/16!)
  - Attack has limits: most ASes won't propogate announcements more specific than /24

Attack: Link Cutting

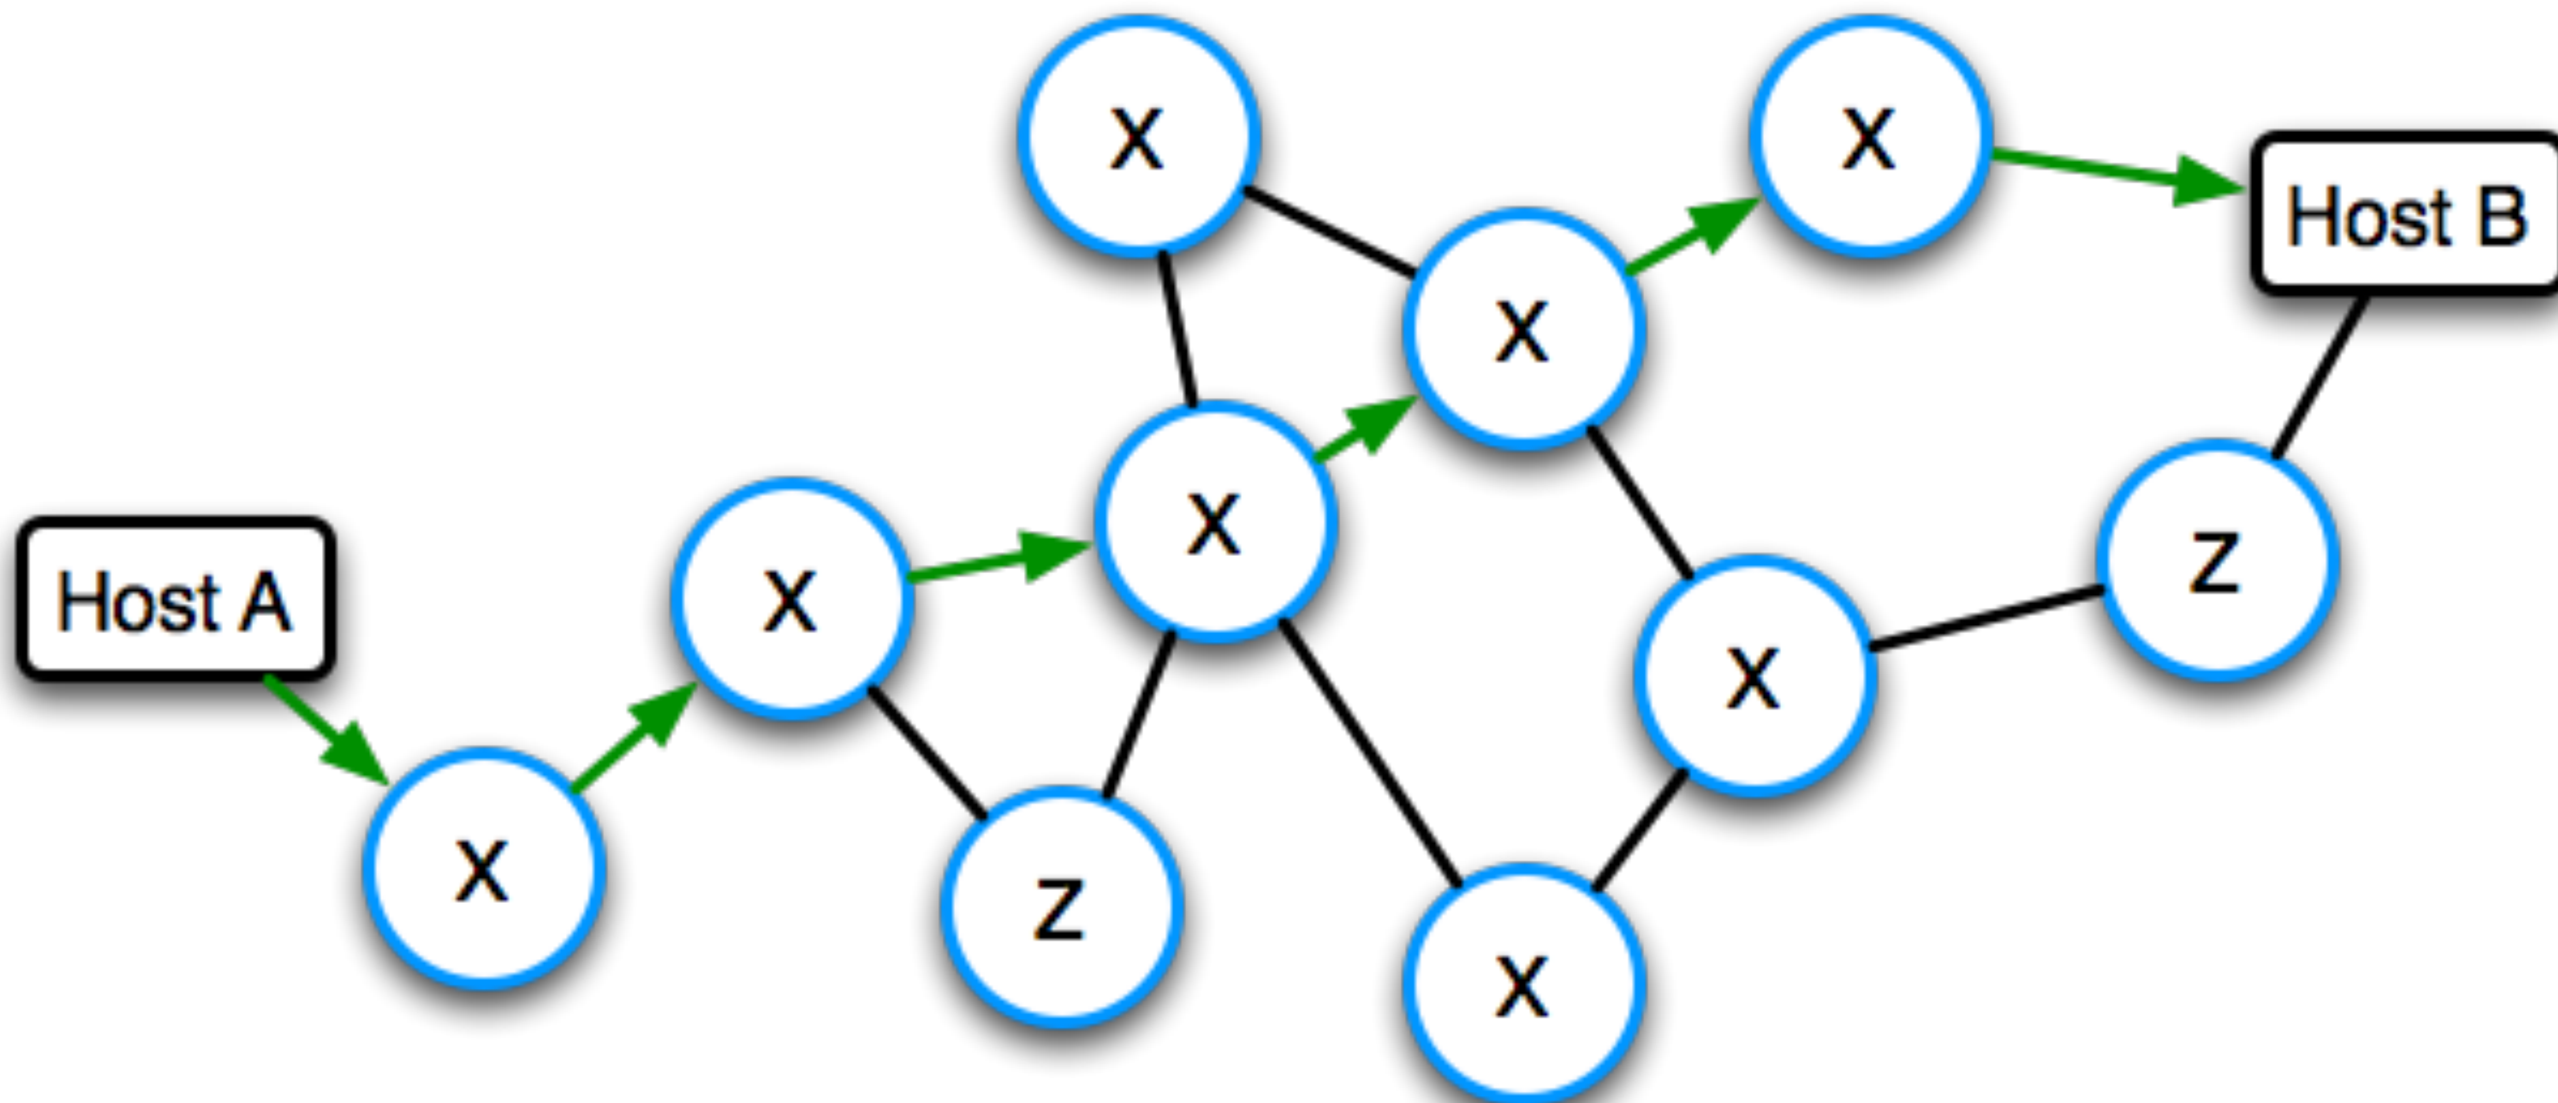

Attack: Link Cutting

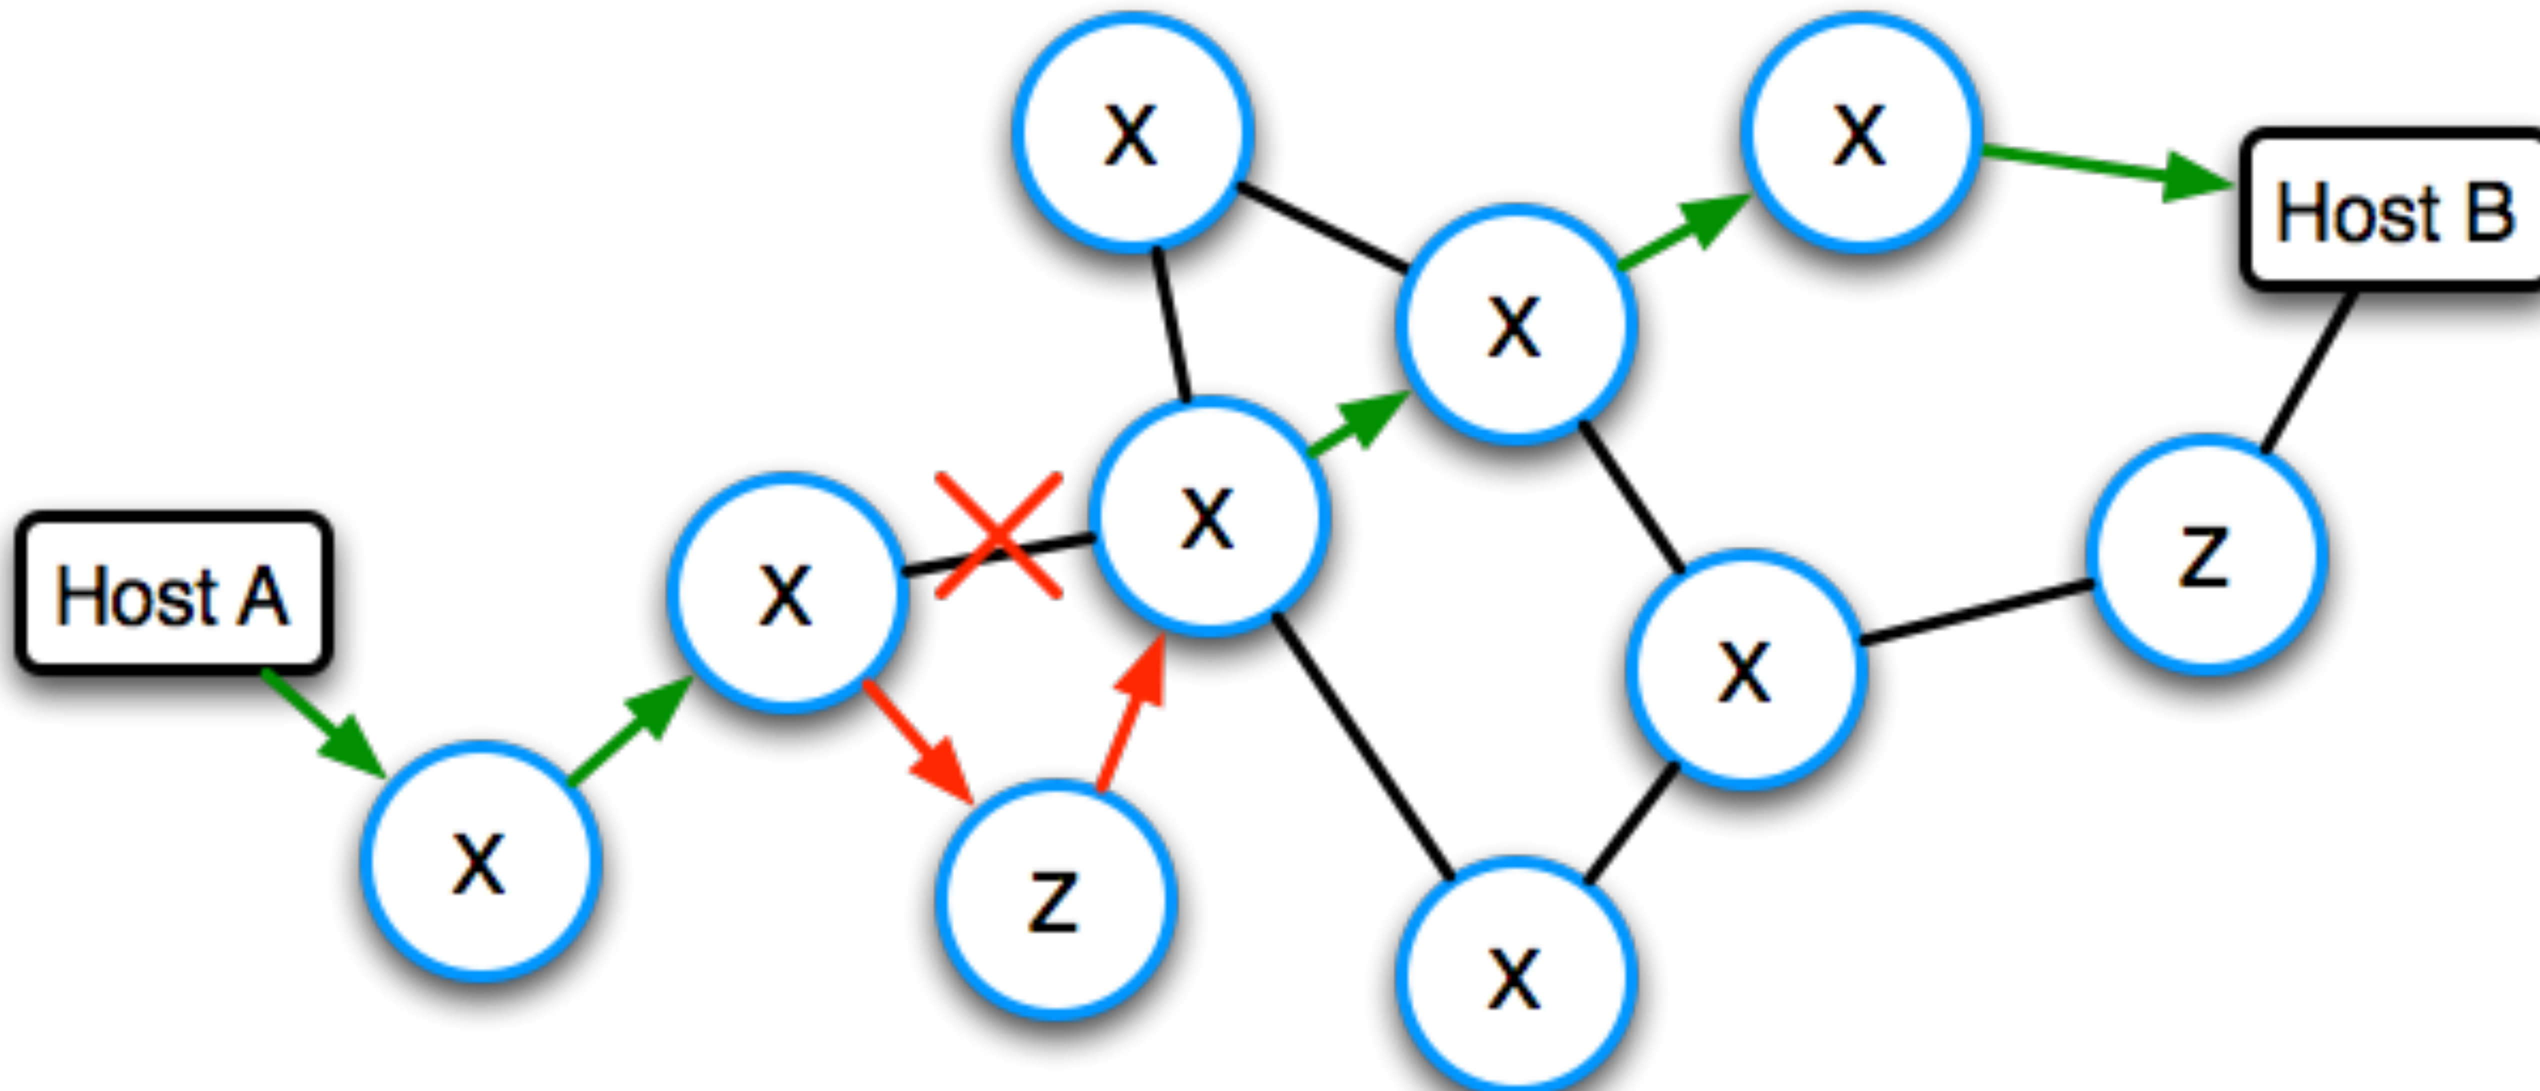

Attack: Link Cutting

- **Link cutting**
    - If the attacker knows the network topology, bringing down certain links (through DoS attacks or a backhoe) can force traffic into the pattern they desire
- Taking control of the router
    - For example, exploiting a buffer overflow
- Physical destruction of the router
    - As always, network security is dependent on physical security

Plan for today

- Administrivia
- Review DNS
- Secure Routing
  - Overview
  - Protocols
  - Attacks
  - Defenses

Later: Defenses

Solving BGP Security

- Reality: most deployed techniques for securing BGP have been at the local level
  - Filtering
  - Securing BGP peering
- Future: a number of complex protocols have been proposed to solve some or all BGP security issue
  - E.g., sBGP, soBGP, IRV, SPV

Filtering

- Filtering just drops BGP message (typically advertisements) as they are passed between ASes
  - Ingress filtering (as it is received)
  - Egress filtering (as it is sent)
- Types of filtering
  - By prefix (e.g., *bogon*/*martian* list)
  - By path
  - By policy
- ISP ASes aggressively filter (this is the main security mechanism)

sBGP

- sBGP was the first leading candidate for routing security
  - Model: routing and origination announcements are signed
  - Signatures are validated based on shared trust associations (CAs)
- It all begins with the keys (really two parallel PKIs)

1. *Binding routers and organizations to ASes.*
2. *Origin authentication PKI*

Route Attestations

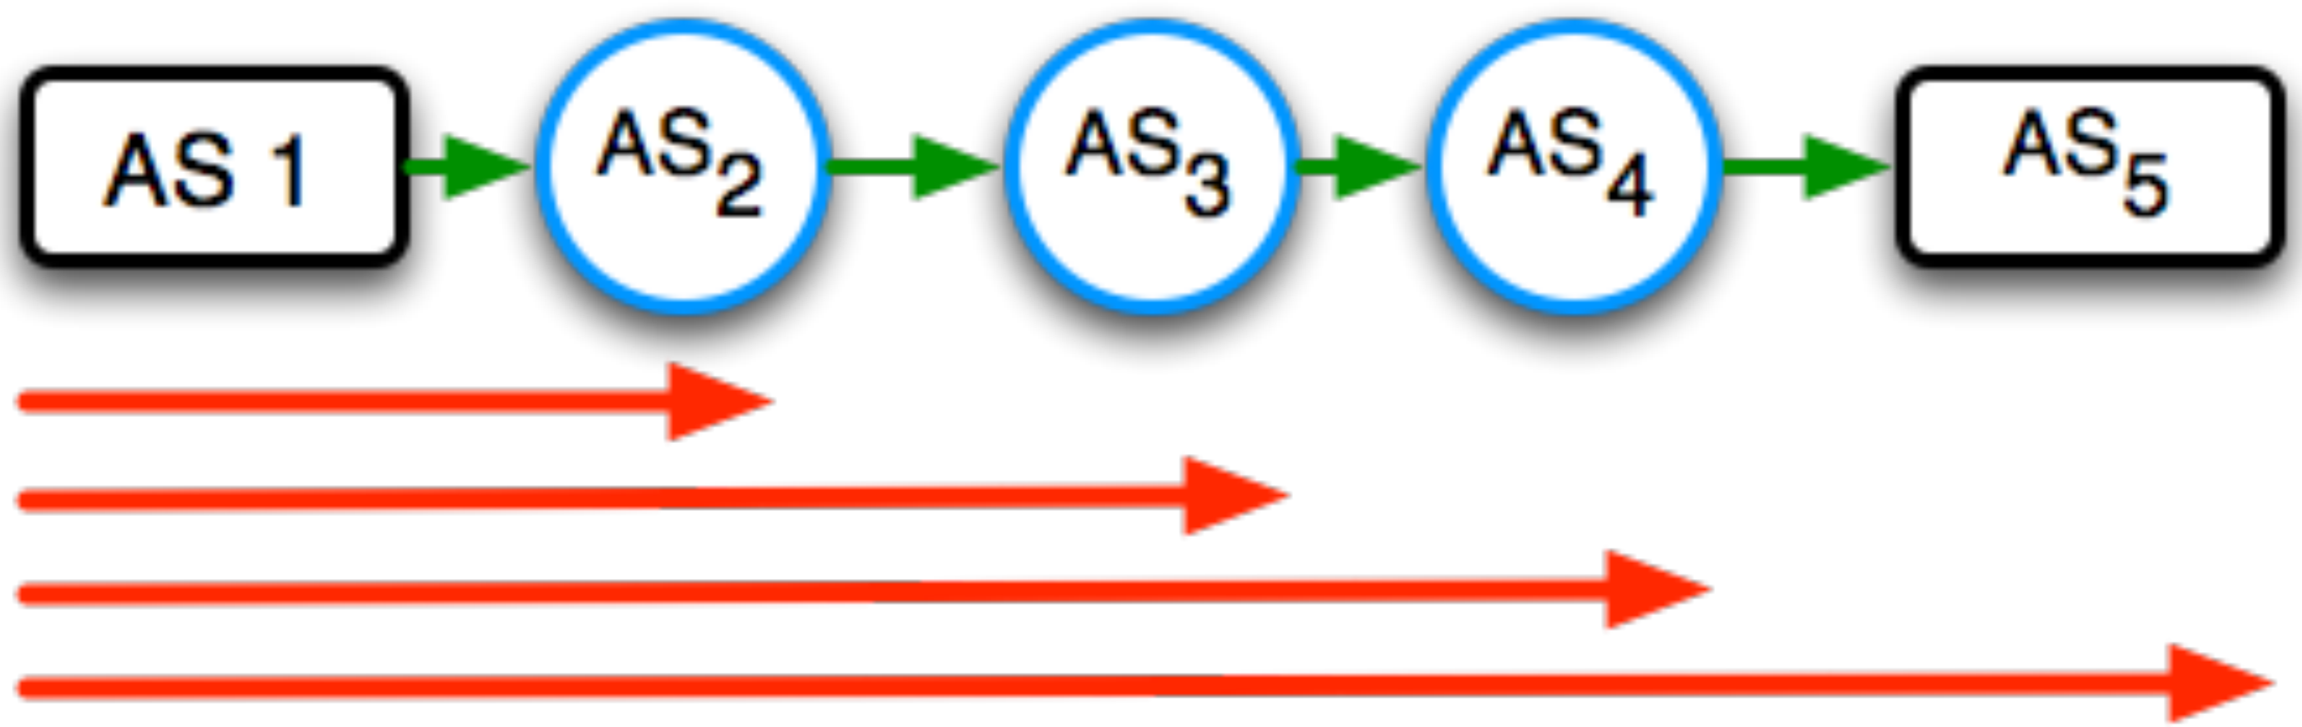

- Signing recursively: each advertisement signs everything it receives, plus the last hop.

$$(5, (4, (3, (2, 1)_{k_{AS_1}})_{k_{AS_2}})_{k_{AS_3}})_{k_{AS_4}}$$

sBGP Issues

- *Single point of trust*: is there an authority that everyone will trust to provide address/path certification?
  - Chinese Military vs. NSA?
- *Cost*: validating signatures is very computationally expensive
  - Can a router sustain the load?
- *Incremental deployability*: requires changes to sBGP message formats
  - All implementations must change

BGP Security

- After almost two decades of work, we are not much closer to a global security solution ...
  - Problems are often not technical ...
    - Cost of building routers
    - Backward compatibility
    - Incremental deployment
- In the future, we will likely move from a border filtering to more and more cryptographically aided solutions.
  - Mining past advertisements and understanding "expected" routing advertisements will also be key where crypto is not appropriate or feasible.

Summary

- Administrivia
- Review DNS
- Secure Routing
  - Overview
  - Protocols
  - Attacks
  - Defenses