CIS 4930: Secure IoT

Prof. Kaushal Kafle

Lecture 18

Internet Routing

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

Routing outside of the local subnet

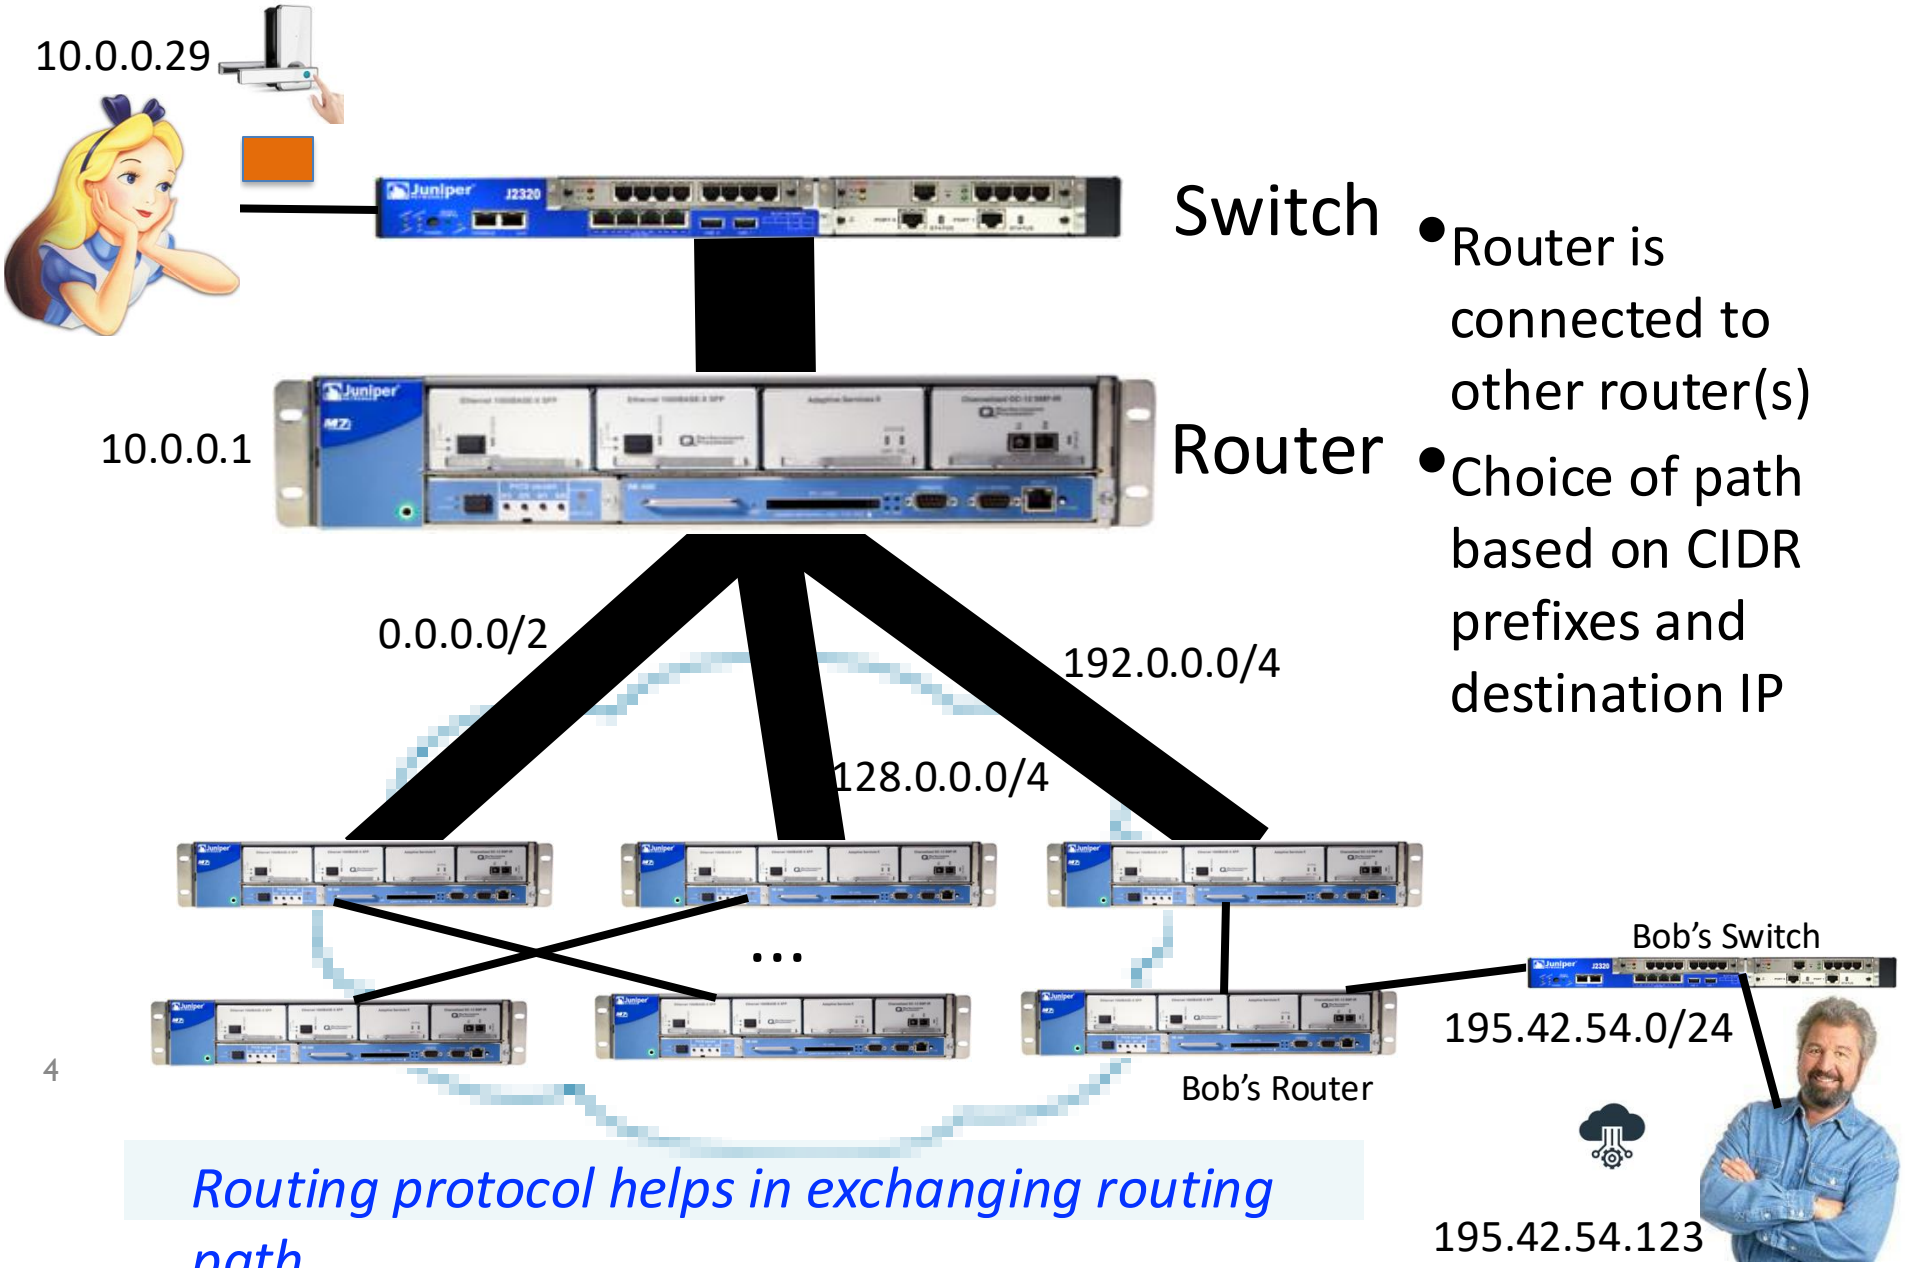

- Router is connected to other router(s)
- Choice of path based on CIDR prefixes and destination IP

Routing protocol helps in exchanging routing path

Routing in a nutshell

- The Internet...

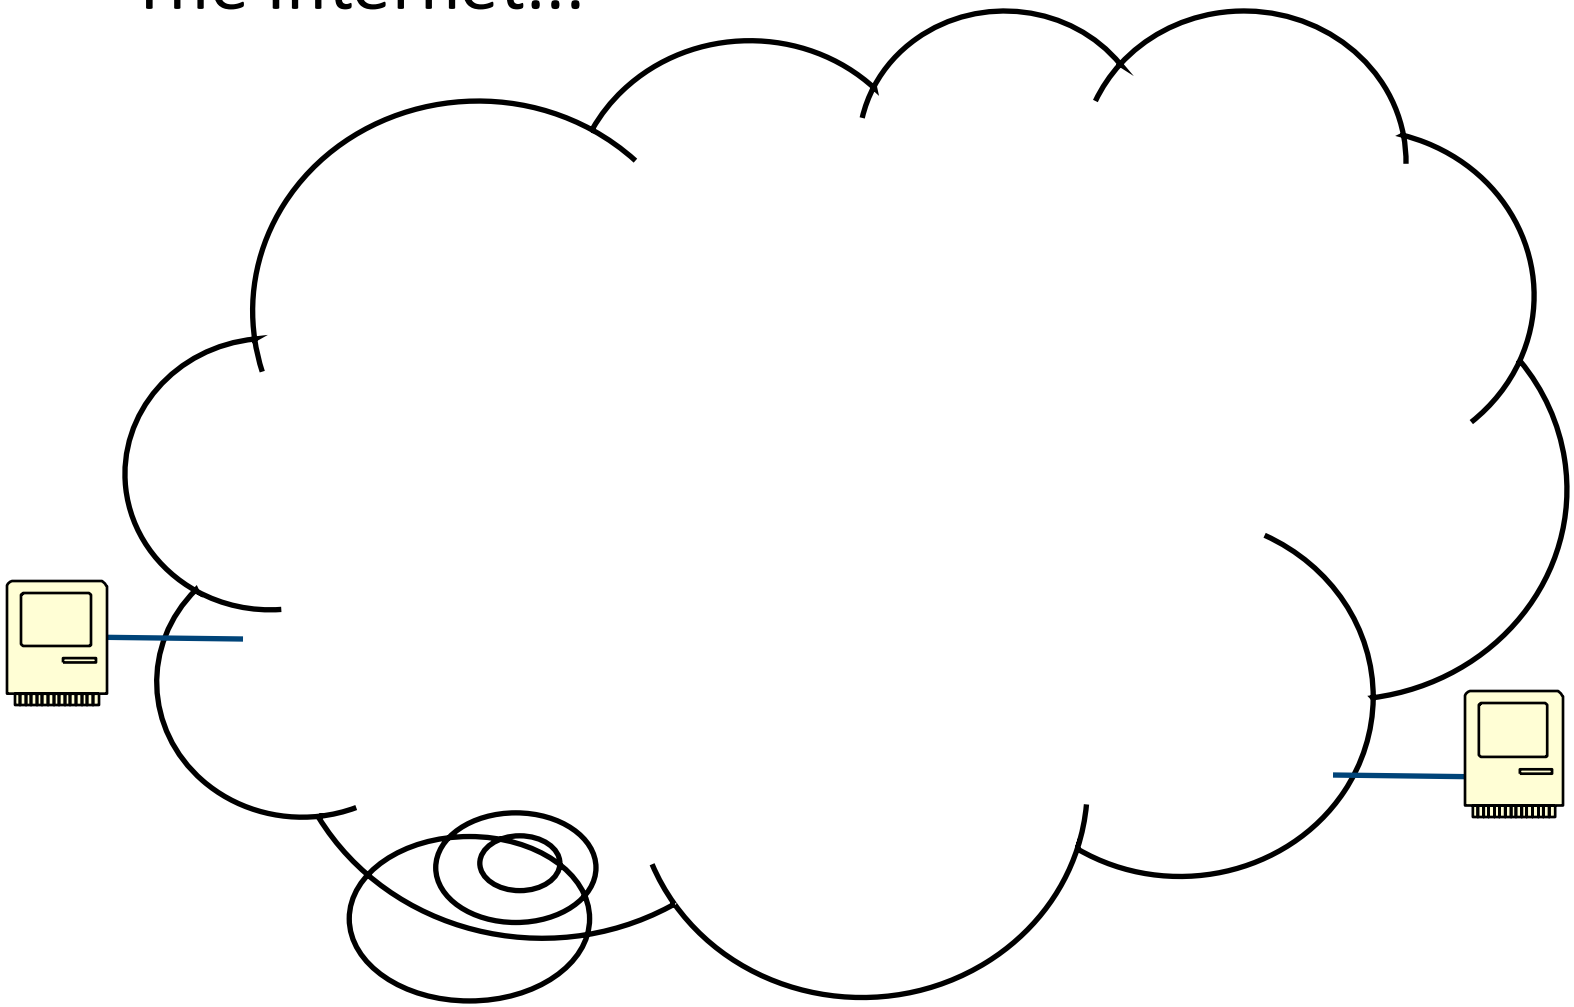

Routing in a nutshell

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

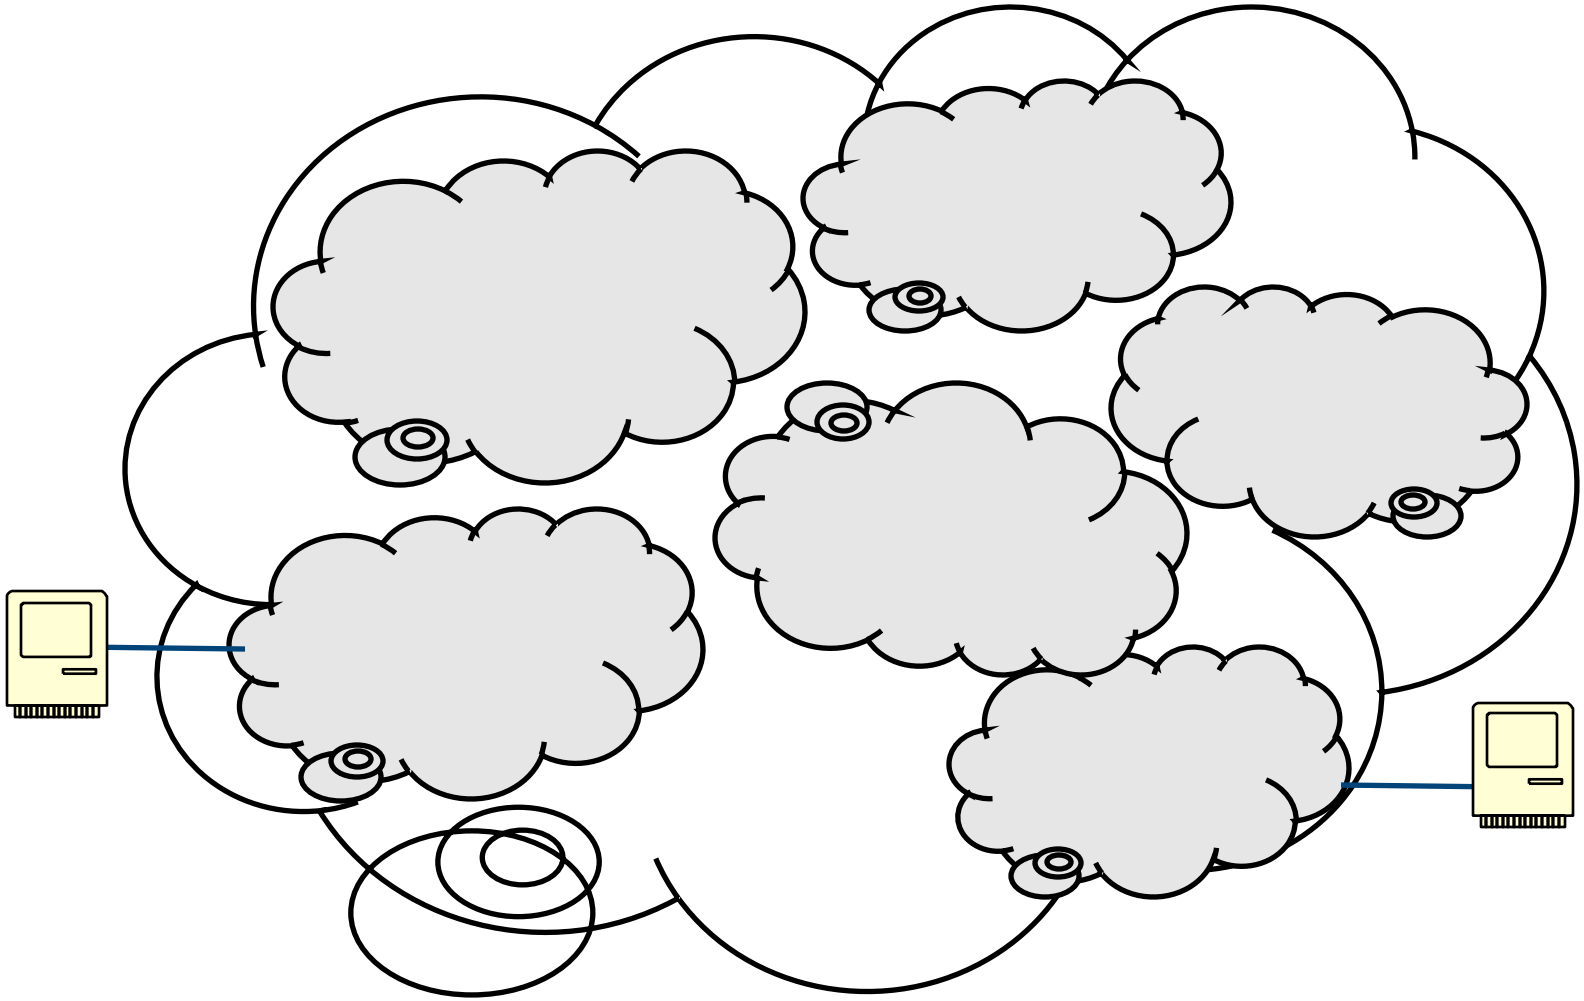

…linked at Border Routers.

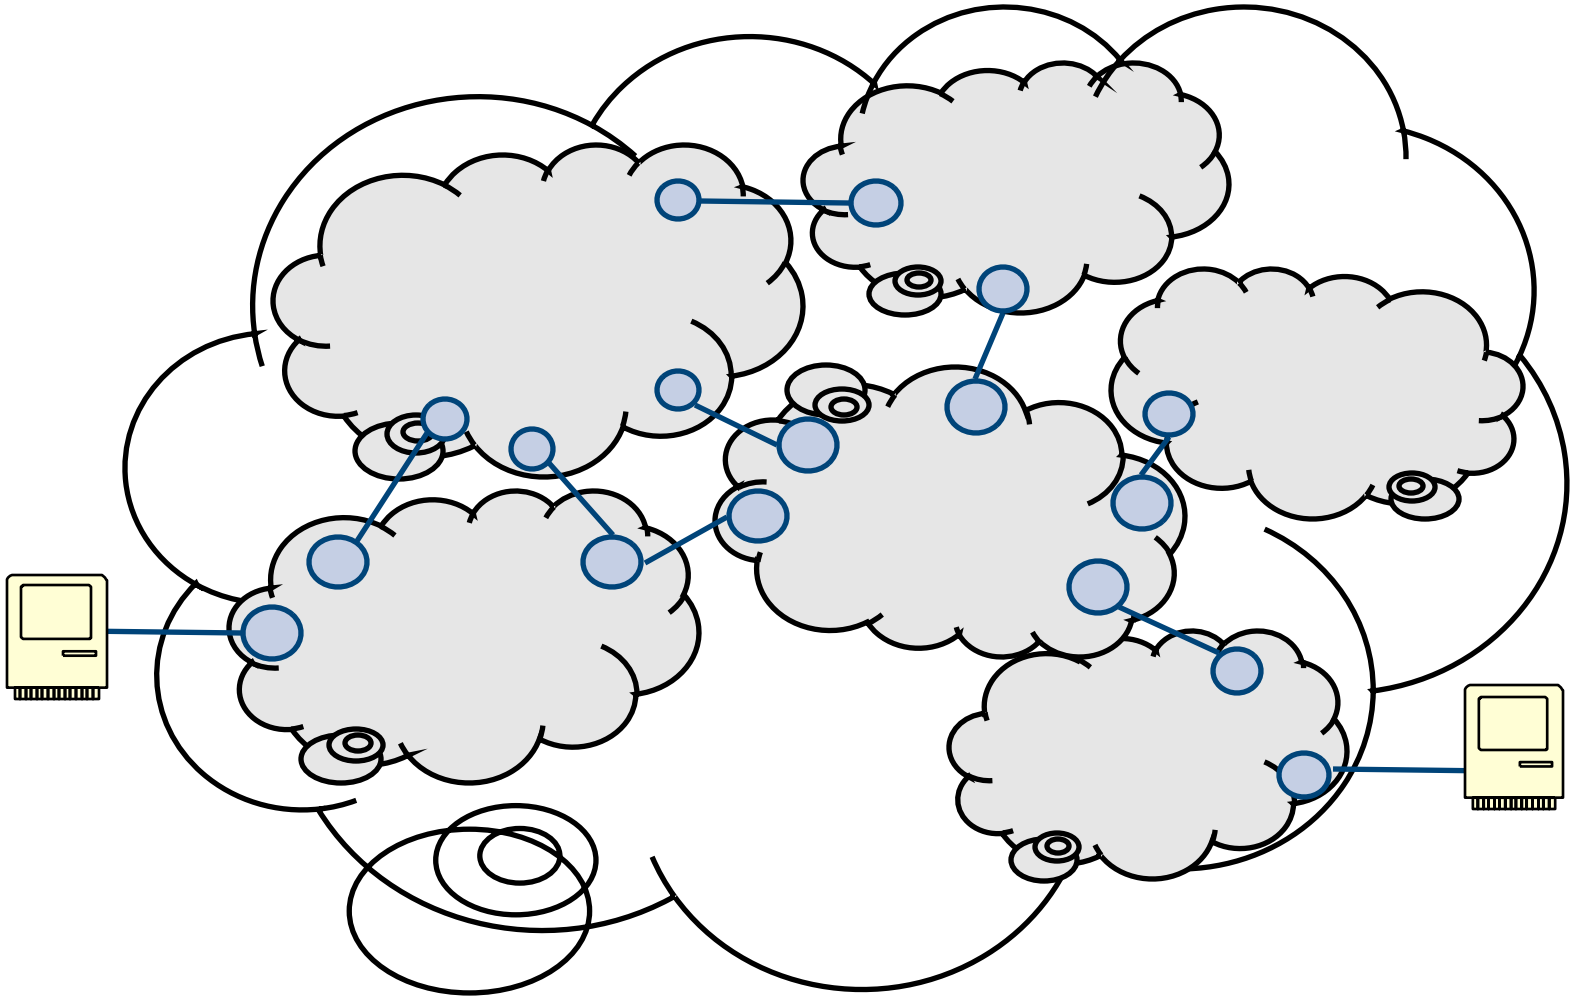

BGP determines which ASes to follow from source to destination

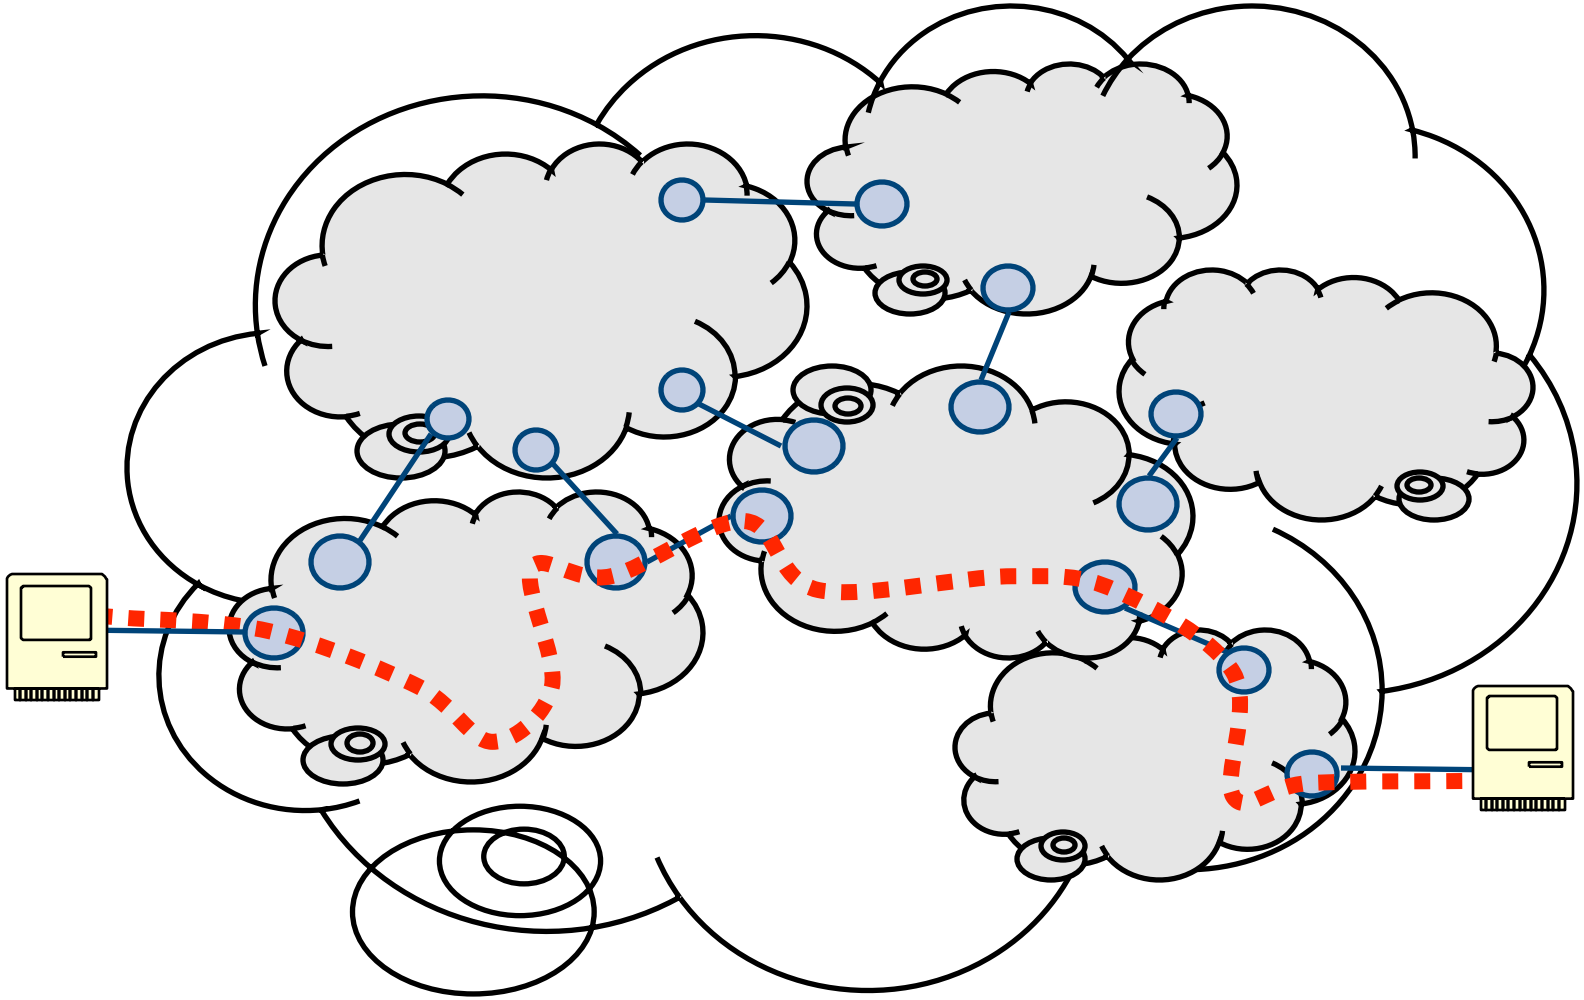

- Each AS is responsible for moving packets inside it.
- Intra-AS routing is (mostly) independent from Inter-AS routing. This is done through OSPF.

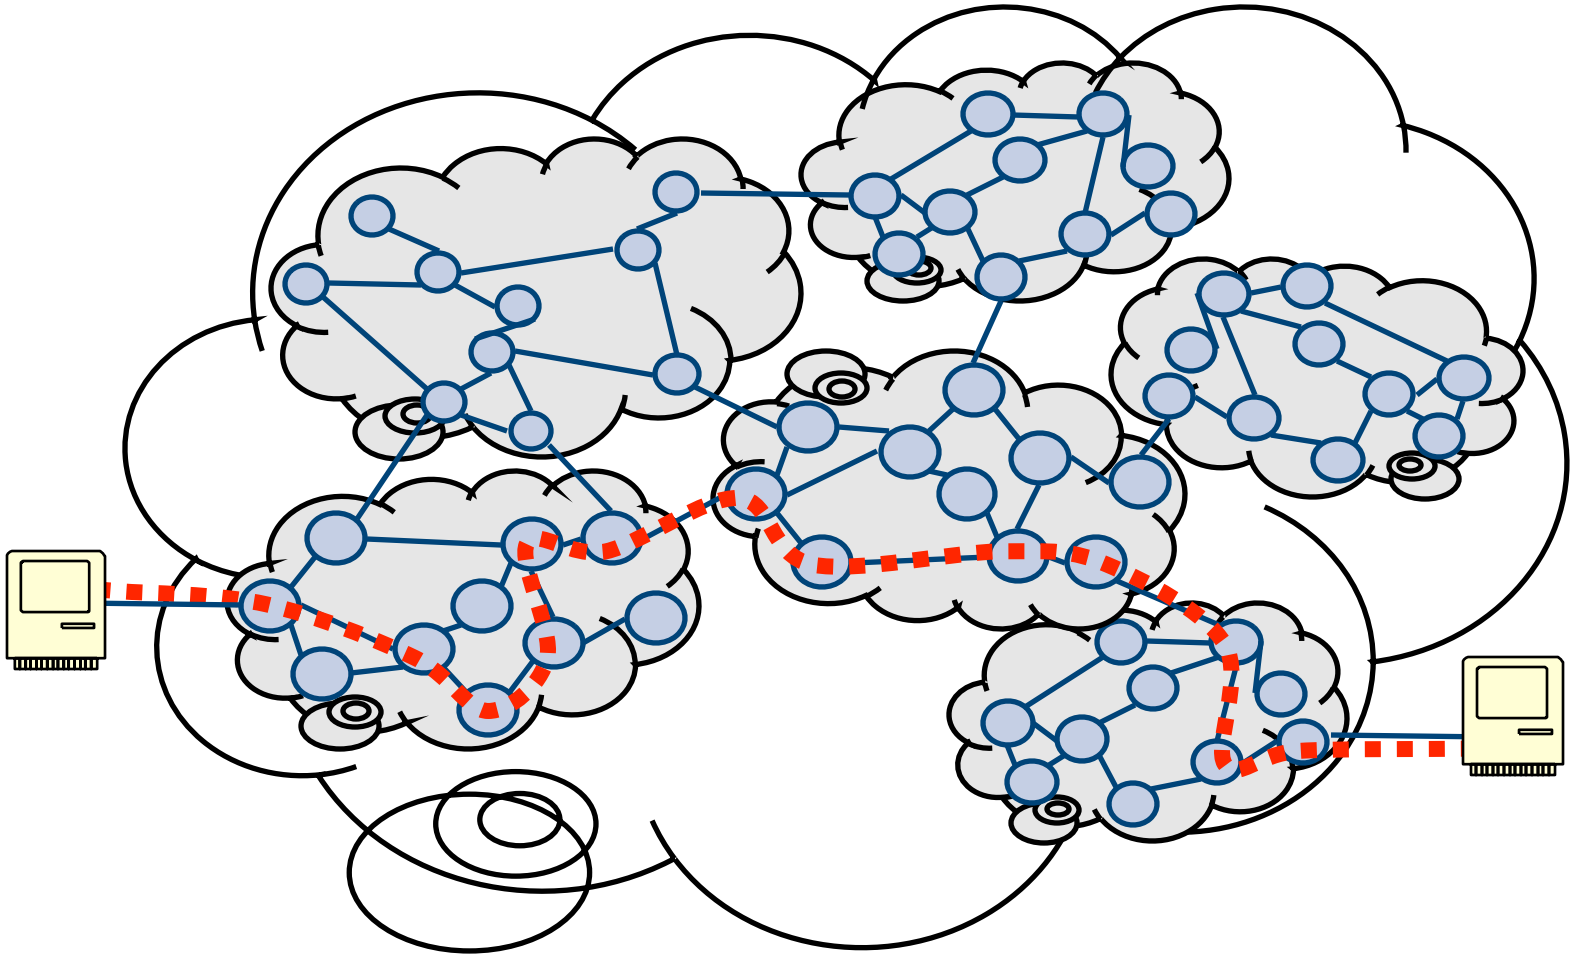

Internal Networks

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

Well-defined admin (e.g., ISPs)

Better control over the network

IP address

- IP(v4) addresses (32-bit) are divided into prefix and suffix.
  - Prefix: Network address
  - Suffix: Host address
- Depending on the routing protocol, no. of bits that are prefix or suffix can vary.

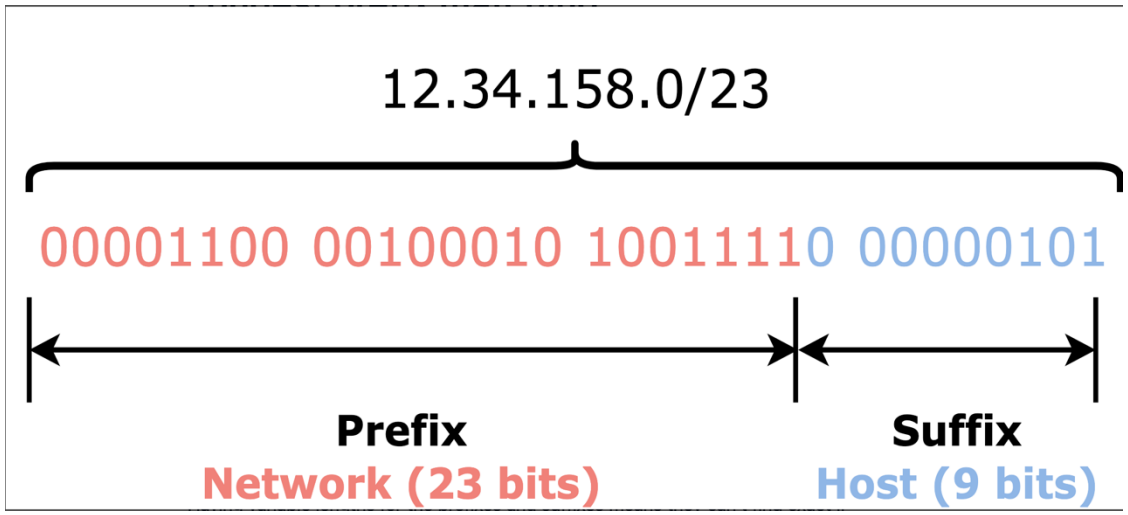

OSPF (Open Shortest Path First)

- 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.

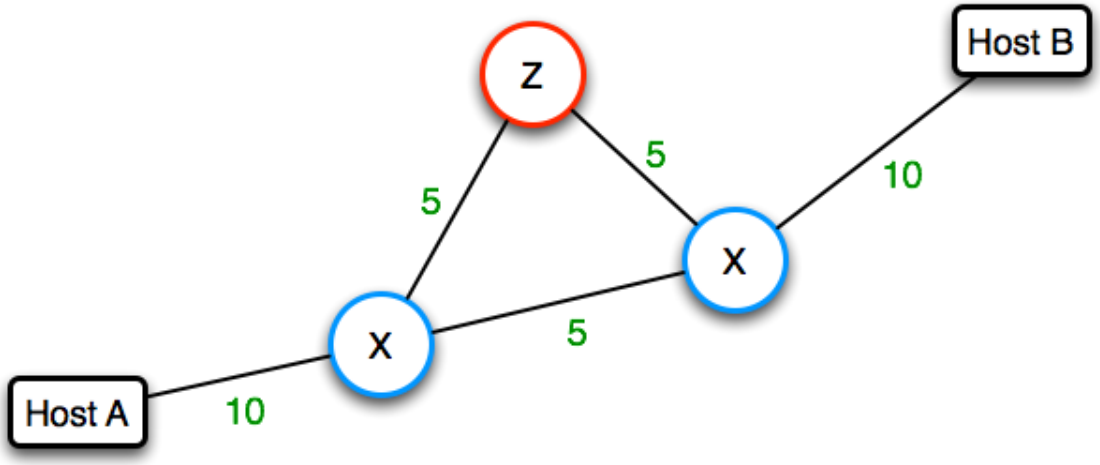

Border Gateway Protocol (BGP)

- 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
    - Bunch of attributes describing the *route* that should be taken
    - Contain the list of ASes in the route

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

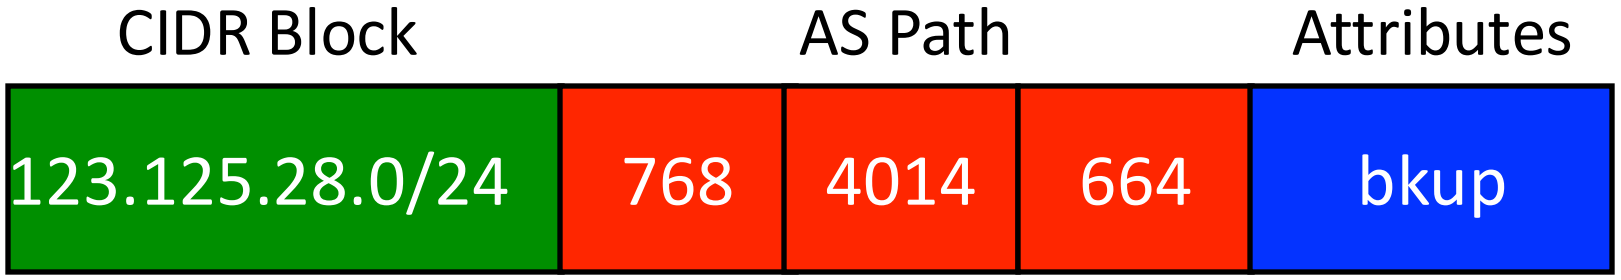

BGP Attacks

Later: Defenses

Attack: Prefix Hijacking

- An attacker can claim to originate a known prefix
- For example, my organization 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:

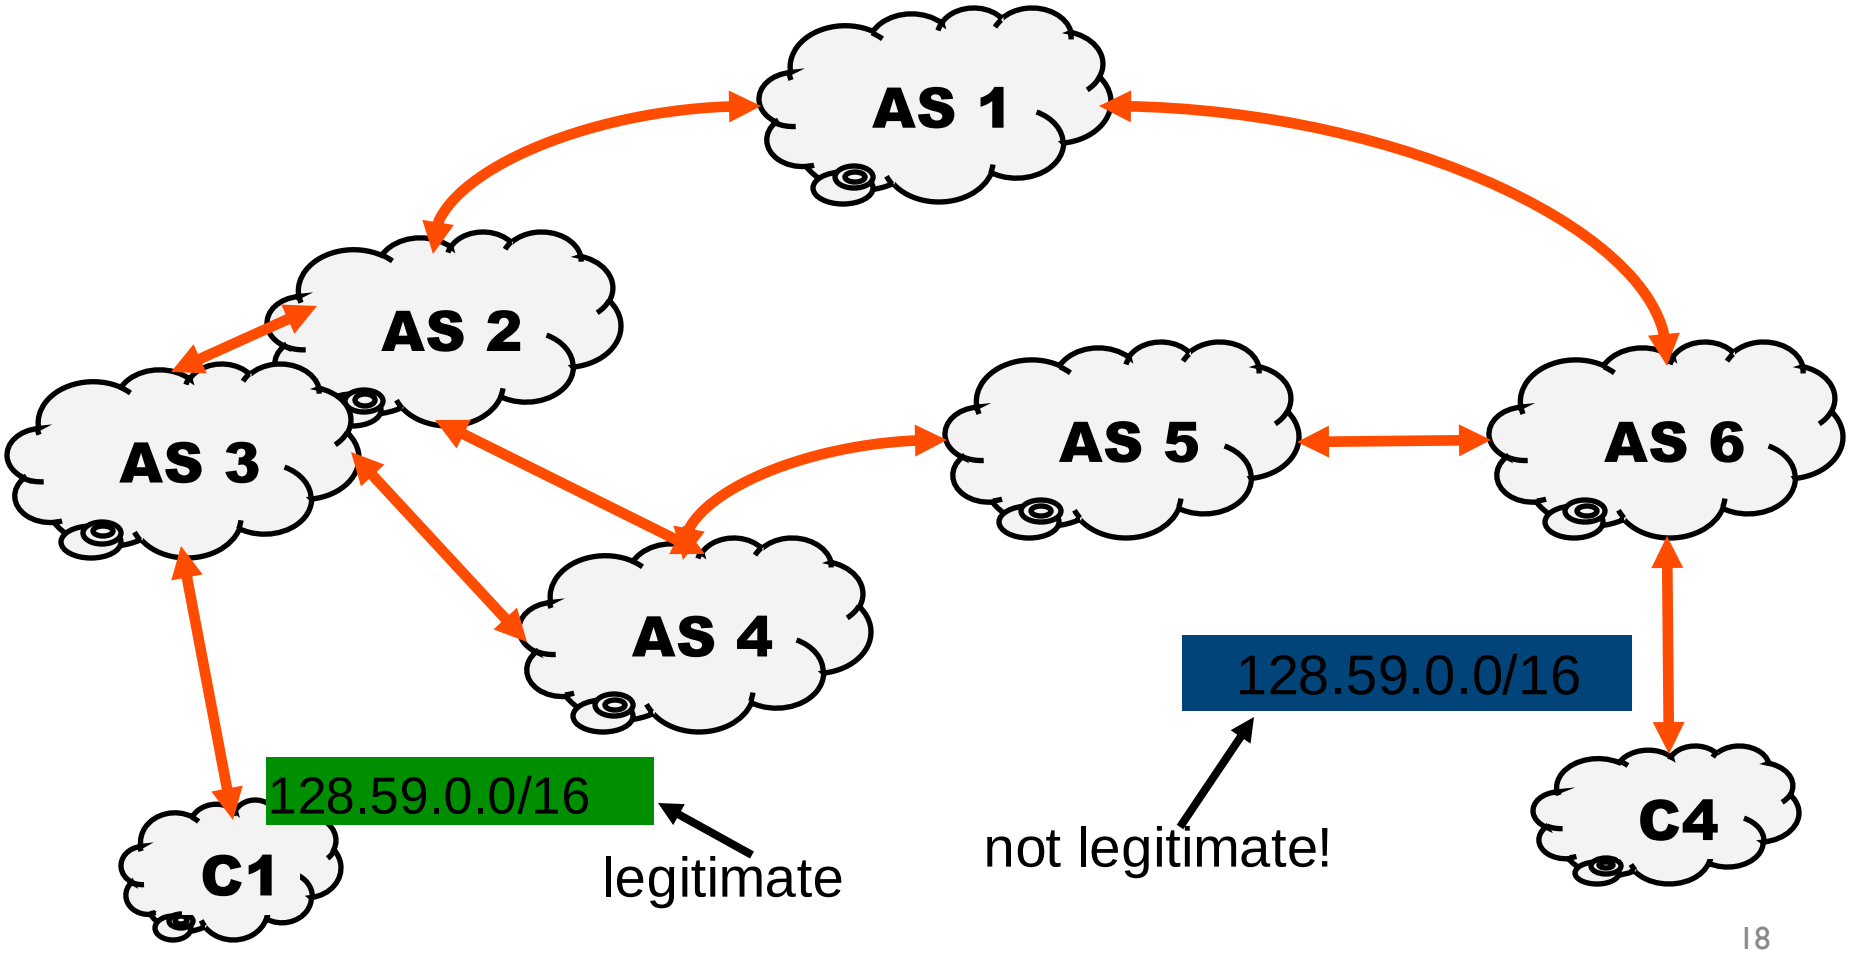

- Prefix becomes unreachable from the part of the net believing C4's announcement.

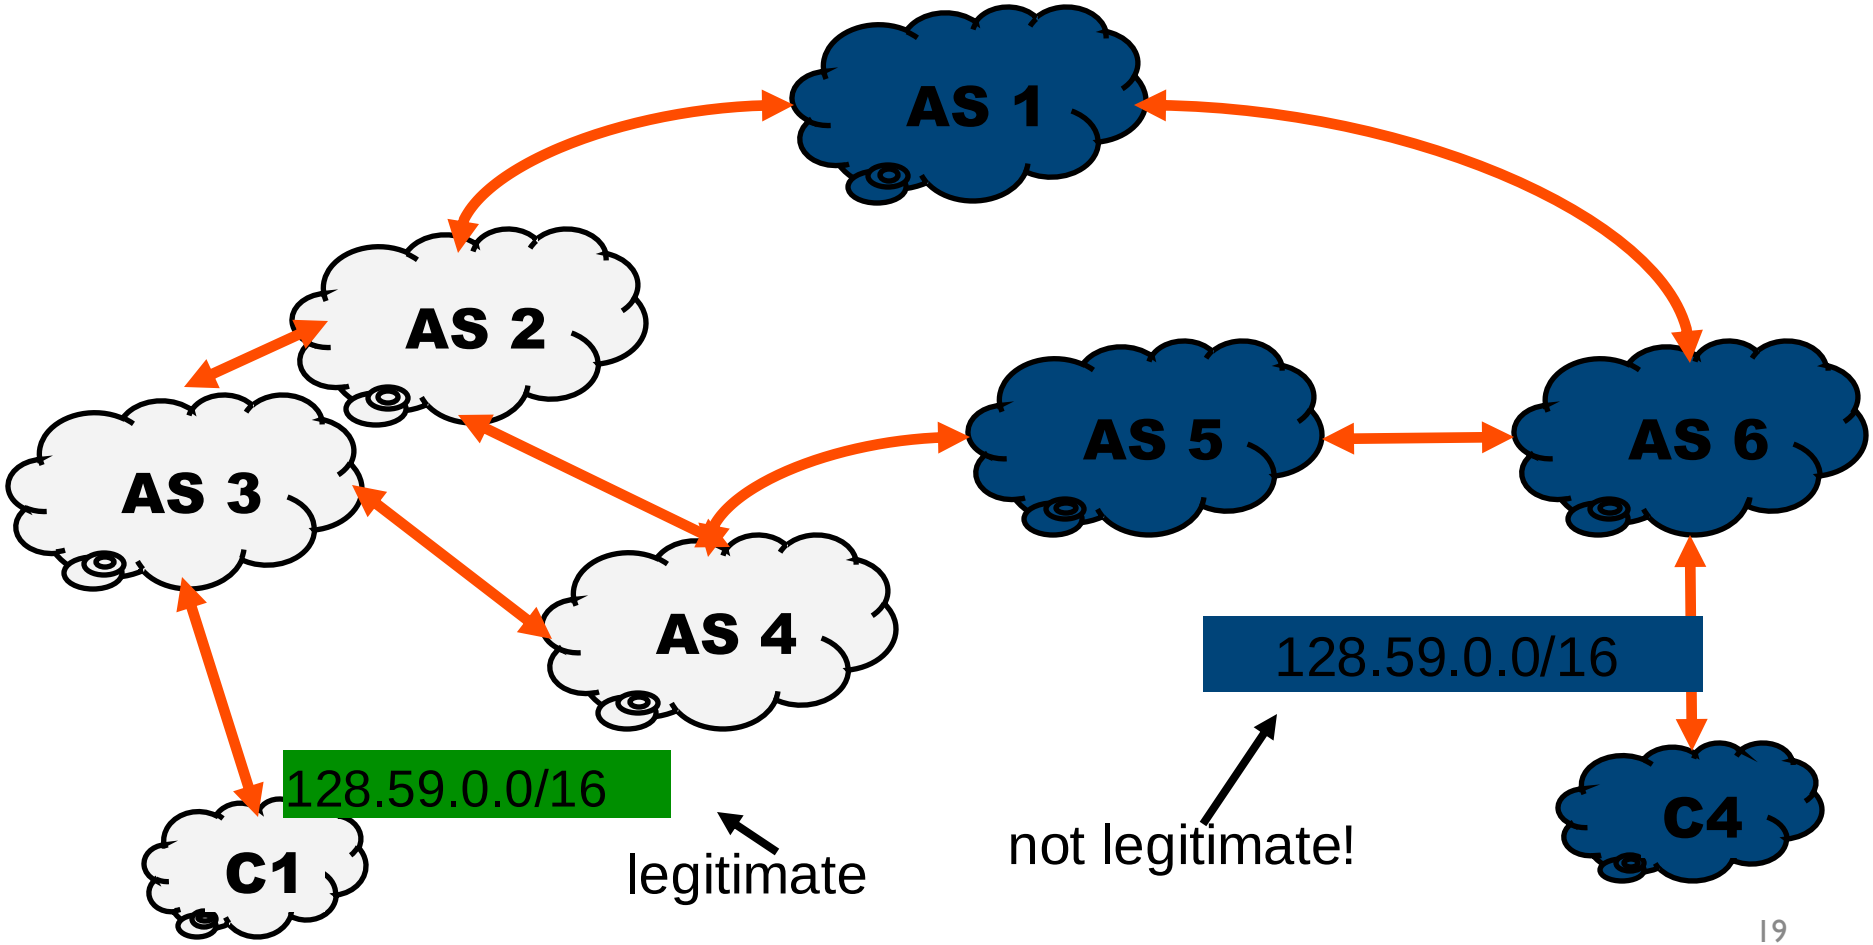

Longest-Prefix Matching

- IP(v4) addresses (32-bit) are divided into prefix and suffix.
  - Prefix: Network address
  - Suffix: Host address
- Depending on the routing protocol, no. of bits that are prefix or suffix can vary.

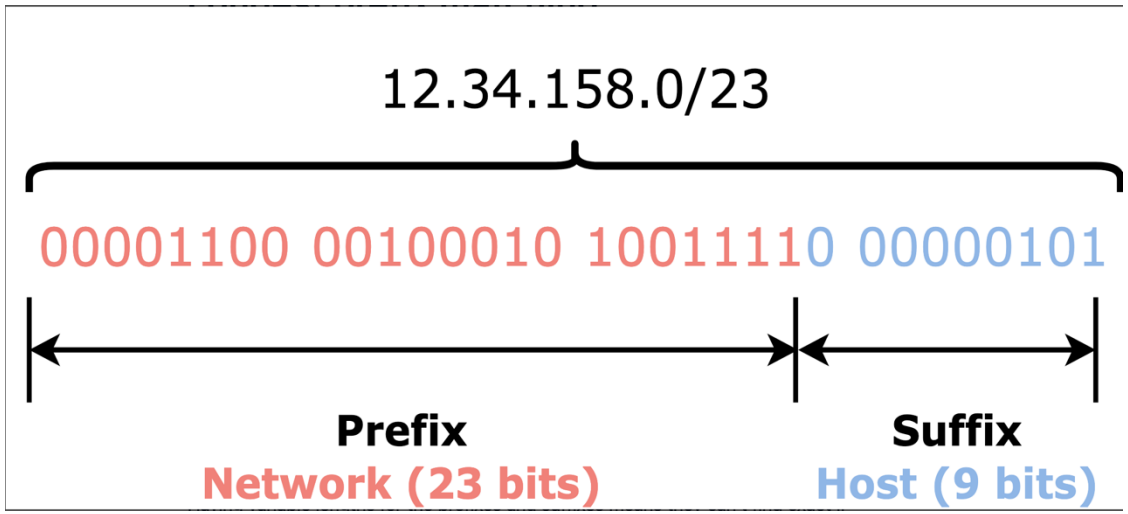

Longest-Prefix Matching

- 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)

Attack: Sub-Prefix Hijacking

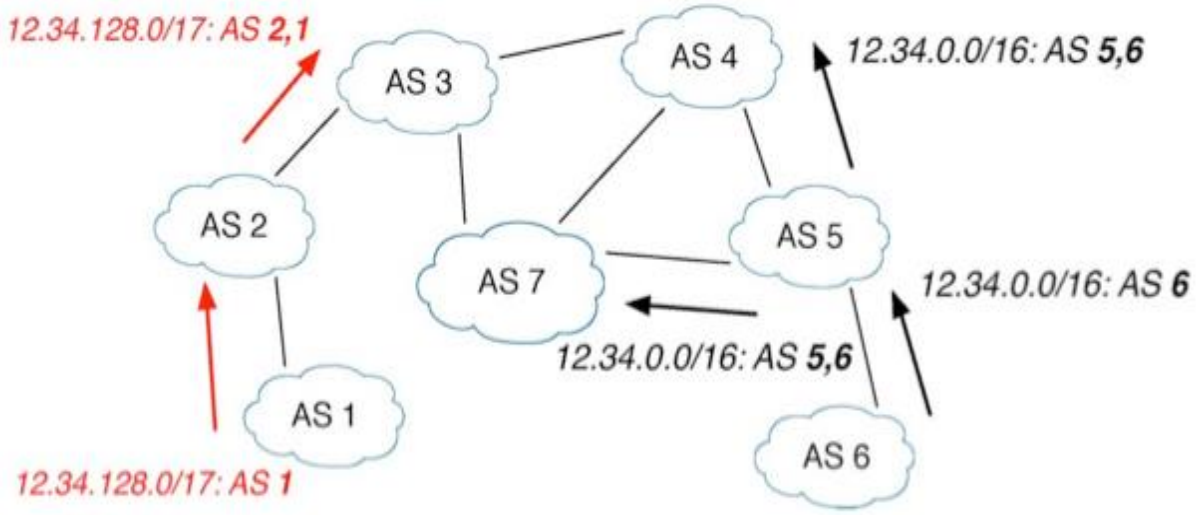

Fig. 3. ***An example of deaggregation. Because AS 1 advertises a longer prefix for the address block 12.34.128.0/17, it will be preferred over the larger advertised block 12.34.0.0/16 even if it is invalid.***

- Much more dangerous than prefix hijacking
  - Why?

Attack: Path Forgery

- If an AS_PATH attribute is completely forged, the attacker has control over traffic
- This can allow for traffic analysis since traffic is engineered in the direction the attacker desires

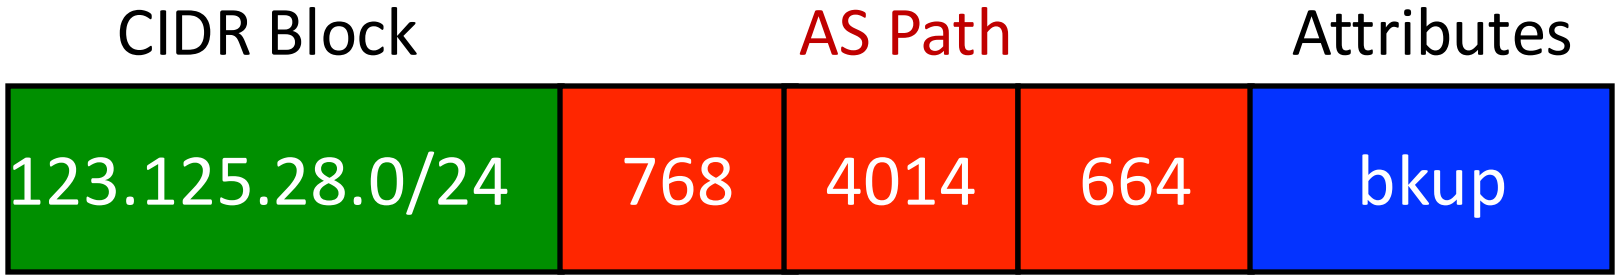

Other Attacks

- **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

Solutions (?)

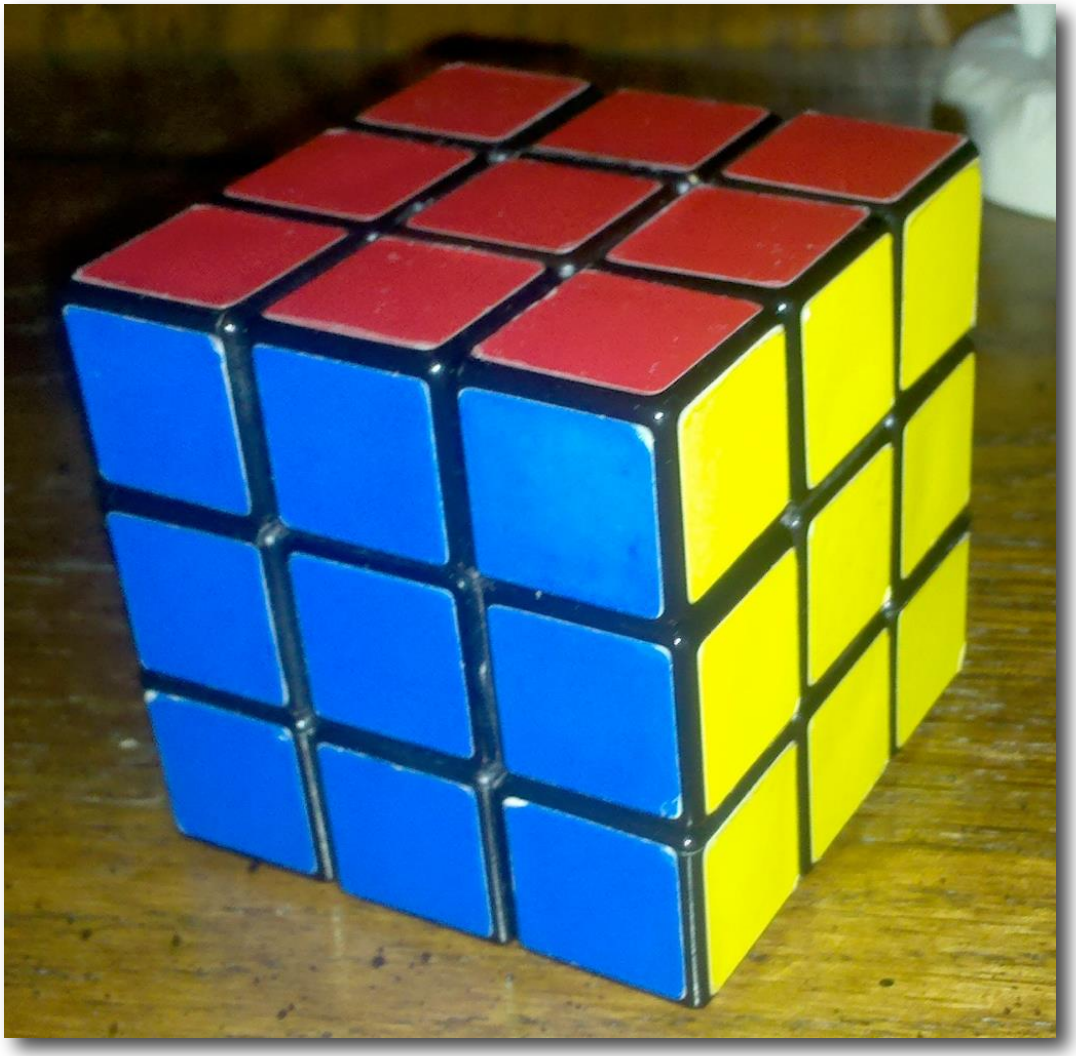

Solving BGP Security

- Reality: most deployed techniques for securing BGP have been at the local level
  - Filtering
  - Securing BGP peering (i.e., securing the connection between neighboring BGP routers)
- Future: a number of complex protocols have been proposed to solve some or all BGP security issue
  - E.g., S-BGP, 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
  - By path
  - By policy

- ISP ASes aggressively filter (this is the main security mechanism)

Prefix Filtering Intuition

- AS's have business relationships that influence the cost of sending traffic
  - Customer, provider, peer
- *Rule of thumb*:
  AS *a* will typically announce a route to a neighbor AS *n* only if
  - *n* is a customer of *a*
  - The route is for a prefix originated by *a*
  - The route is through a customer of *a*
- Provides a basis for defining prefix filters
  - If an AS has no incentive, treat it as sus!

Prefix Filtering

- Benefits: Simple and effective
- Challenges:
  - Prefix filtering works only on customer links
  - Lopsided incentives (e.g., the one filtering is often not the victim)

RPKI

- The Resource Public Key Infrastructure (RPKI) uses cryptography for *origin validation*
- Goal: Only allow legitimate ASes to advertise specific IP blocks.
  - i.e., *can't advertise new origin unless AS signs it.*
- Establishes a hierarchy based on the allocation of CIDR addresses defined by **Regional Internet Registries** (RIRs)
  - RIRs oversee allocation of IP-blocks to ASes
- Does not do path validation
- Benefits:
  - Offline cryptography (verify updates once per day)
  - Protection from hijacks
- Challenges:
  - RPKI takedowns (i.e., withdraw invalid advertisements) and misconfigurations
  - Does not work for route leaks or path shortening attacks

sBGP

- sBGP was the first leading candidate for routing security
  - Provides both origin validation and path validation
  - Still under consideration, but somewhat limited
- 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)
    - 2 keys:
      1. *Binding routers and organizations to ASes.*
      2. *Origin authentication (as done in RPKI)*

Route Attestations

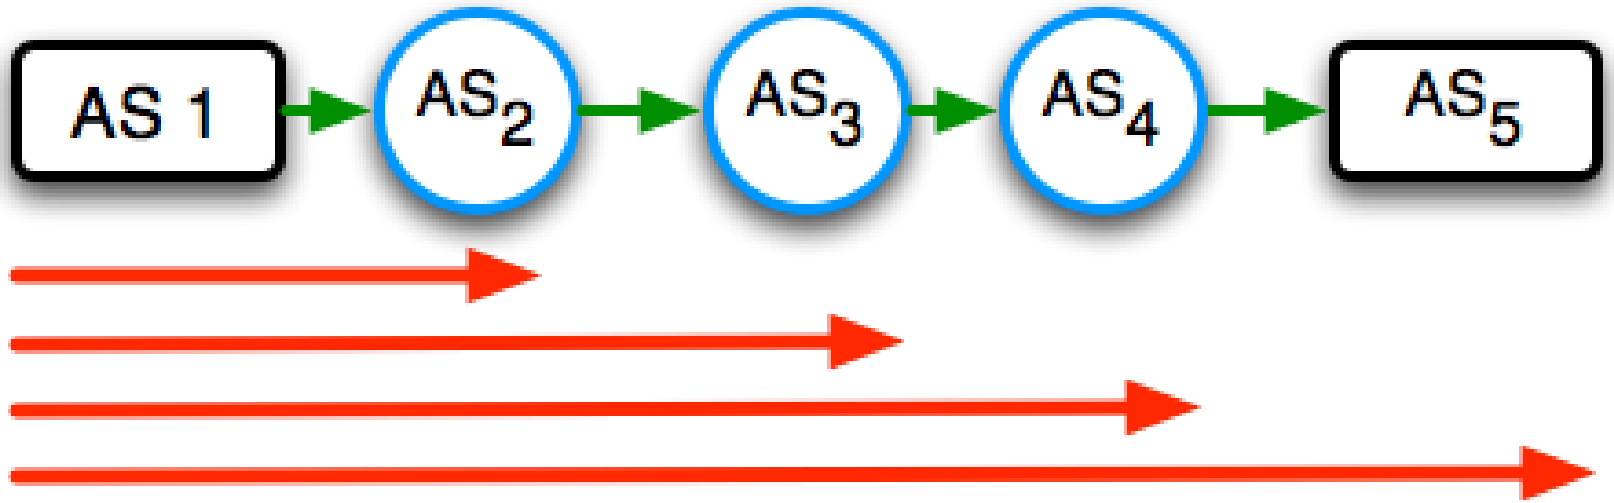

- Signing recursively: each advertisement signs everything it receives, plus the last hop.

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

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 BGP 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.

Wireless Security

Wireless makes network security much more difficult

- Wired:
  - If Alice and Bob are connected via a wire, Eve can only eavesdrop if she has physical access to that wire* (exceptions?)
- Wireless:
  - Everybody shout (broadcast) as loud as you can
  - Friendly to eavesdropping

Evil Toaster

(with wireless card)

Access Point

Internet

Infrastructure mode

Evil Toaster

(with wireless card)

Internet

Ad hoc mode

Finding wireless networks is easy

- wardriving
- warbiking
- warwalking
- warrailing
- warkitteh

(http://thehackernews.com/2014/08/how-to-weaponize-your-cat-to-hack-your_9.html)

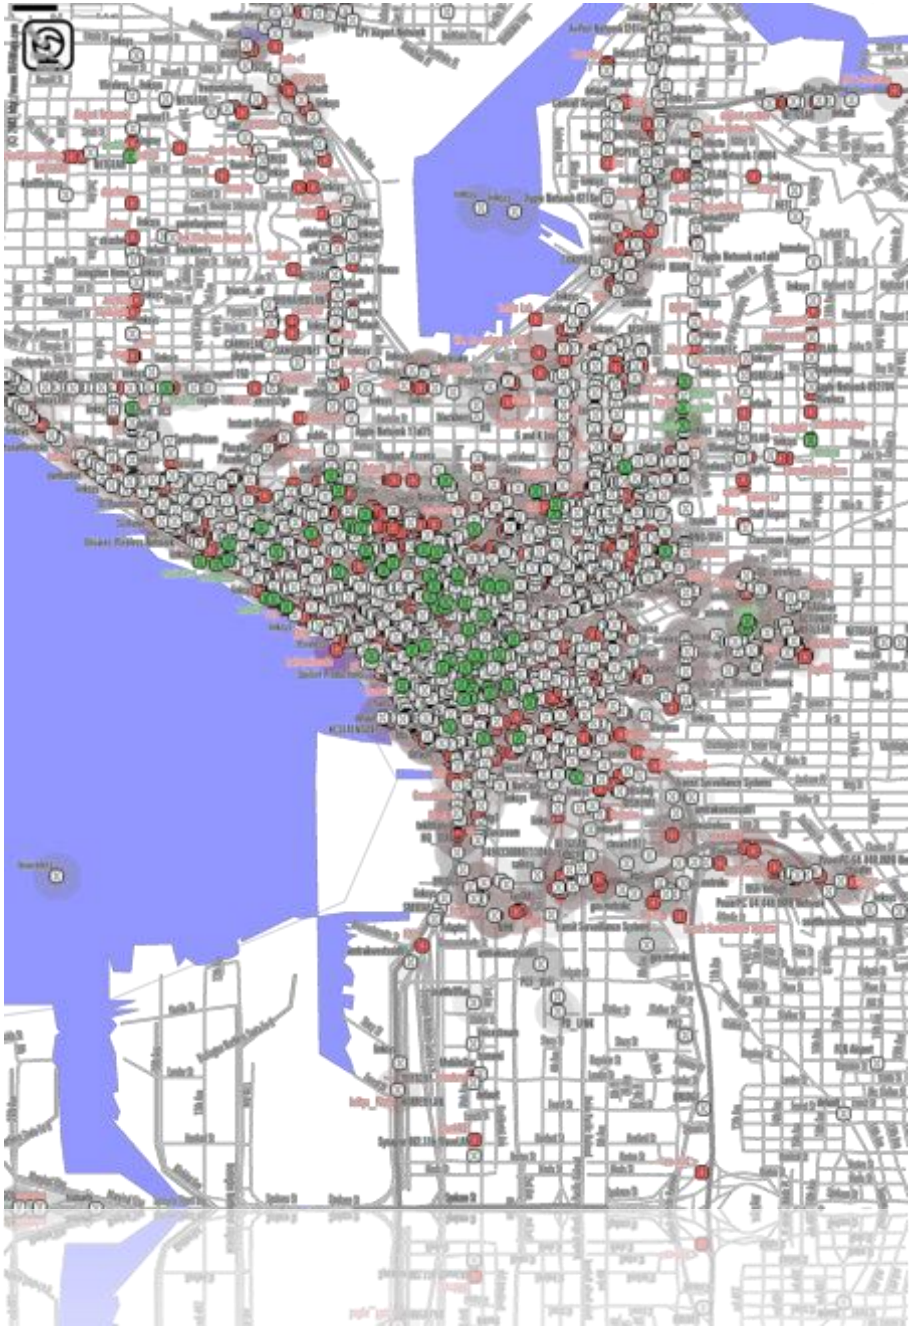

Several online repositories
Check out https://wigle.net/

Wireless Networking: 50,000 ft view

- Protocols defined in IEEE 802.11 standards
- Access points (APs) may periodically broadcast *beacon frames* to advertise its presence (and some configuration parameters)
- Authentication:
  - client sends *authentication frame* to AP
  - if successful, client sends *association request frame* to AP, requesting allocation of resources
    - if successful, AP responds with *association response frame*
- Data sent via *data frames*
- Session Termination:
  - AP sends *disassociation frame* and *deauthentication frame*

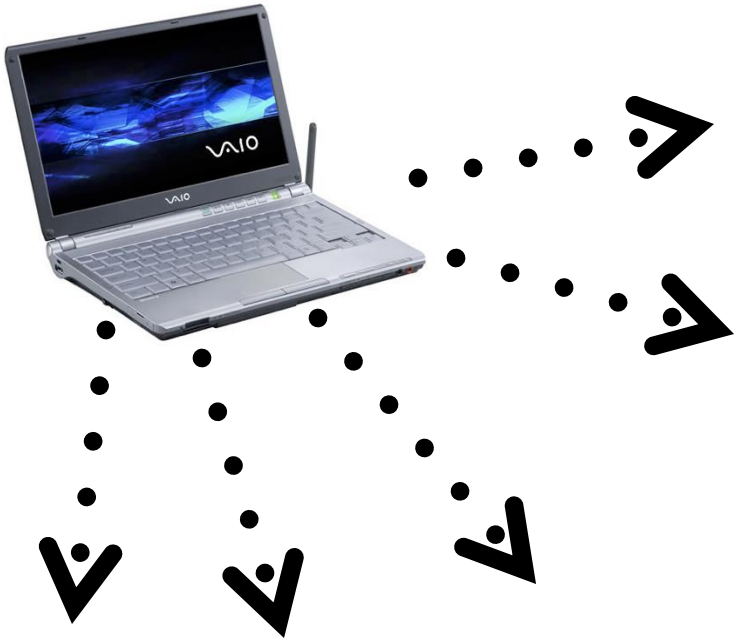

Unsecured wireless:

Problem #1:

Everybody is the receiver.

Unsecured wireless:

Problem #2:

Any one can join.

MAC Filtering

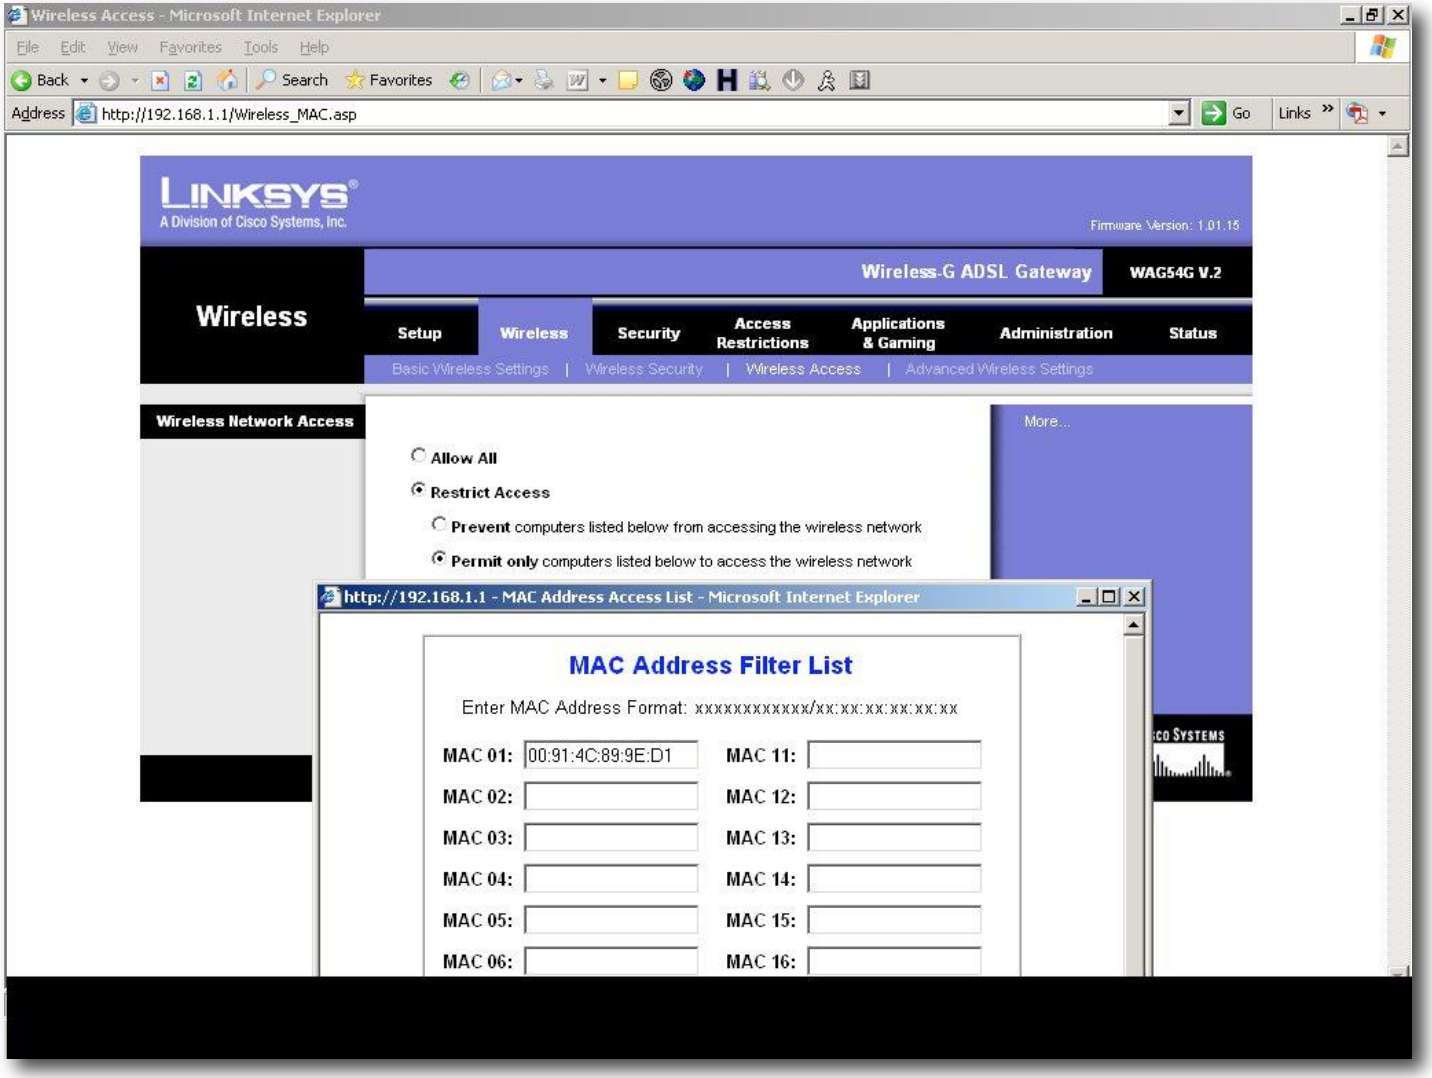

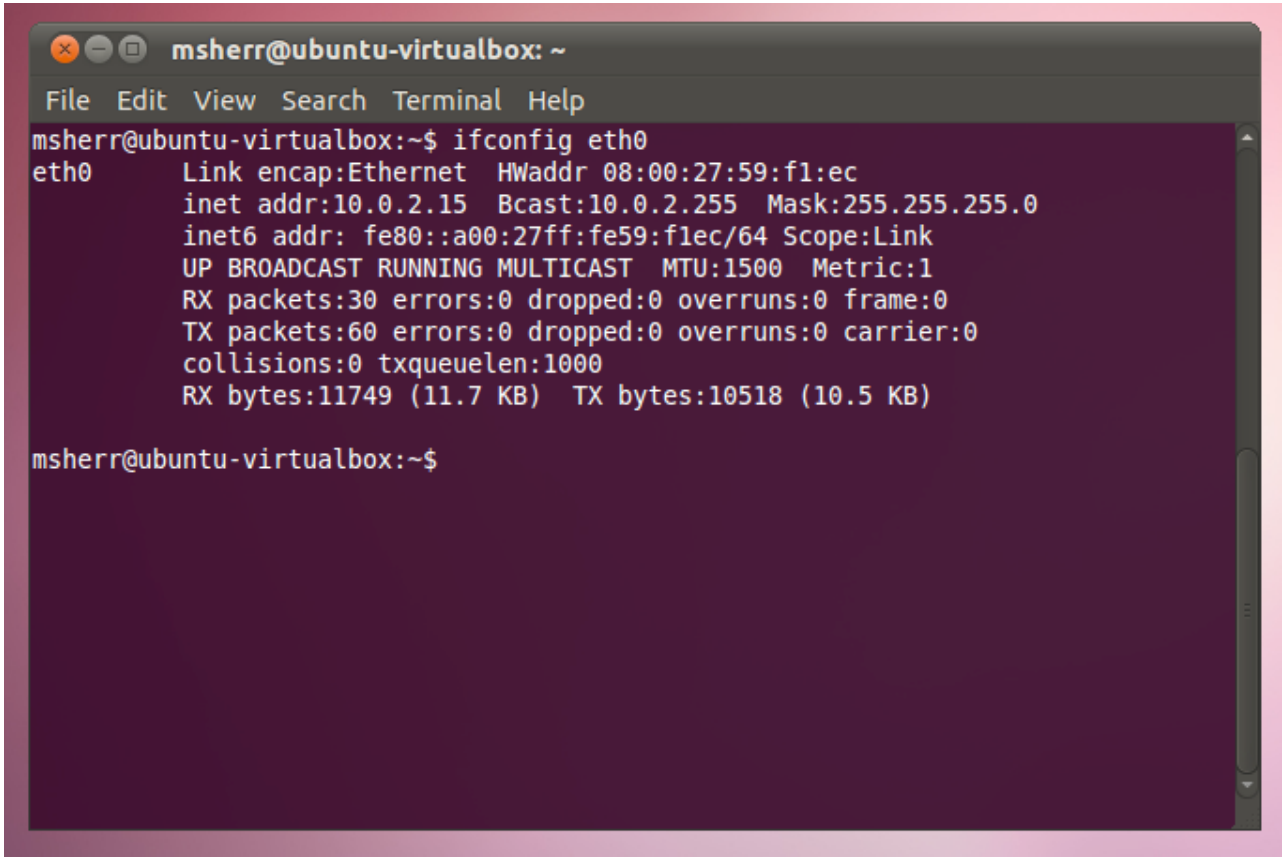
msherr@ubuntu-virtualbox: ~
File Edit View Search Terminal Help
msherr@ubuntu-virtualbox:~$ ifconfig eth0
eth0 Link encap:Ethernet HWaddr 08:00:27:59:f1:ec
inet addr:10.0.2.15 Bcast:10.0.2.255 Mask:255.255.255.0
inet6 addr: fe80::a00:27ff:fe59:f1ec/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:30 errors:0 dropped:0 overruns:0 frame:0
TX packets:60 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:11749 (11.7 KB) TX bytes:10518 (10.5 KB)
msherr@ubuntu-virtualbox:~$

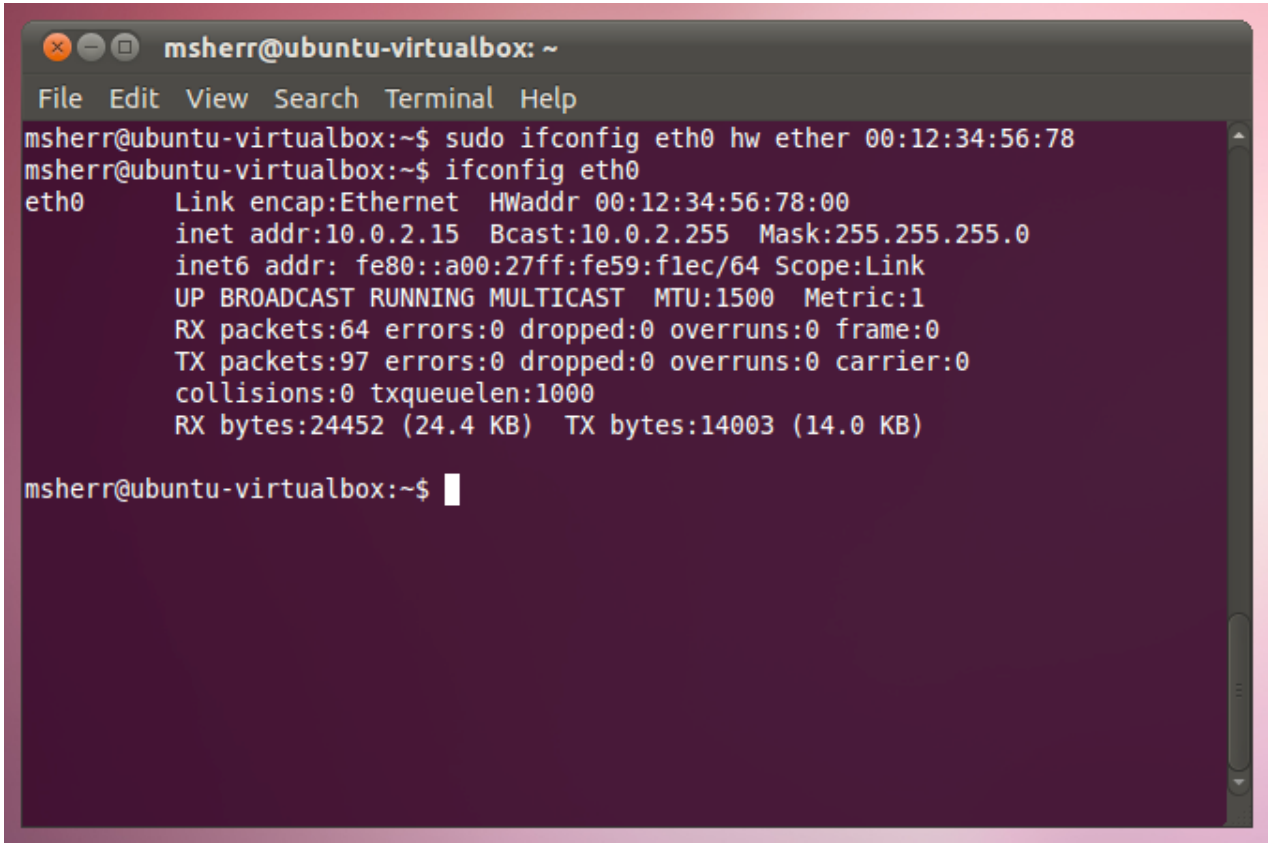
msherr@ubuntu-virtualbox: ~
File Edit View Search Terminal Help
msherr@ubuntu-virtualbox:~$ sudo ifconfig eth0 hw ether 00:12:34:56:78
msherr@ubuntu-virtualbox:~$ ifconfig eth0
eth0 Link encap:Ethernet HWaddr 00:12:34:56:78:00
inet addr:10.0.2.15 Bcast:10.0.2.255 Mask:255.255.255.0
inet6 addr: fe80::a00:27ff:fe59:f1ec/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:64 errors:0 dropped:0 overruns:0 frame:0
TX packets:97 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:24452 (24.4 KB) TX bytes:14003 (14.0 KB)
msherr@ubuntu-virtualbox:~$

SSID hiding

- APs broadcast **Service Set Identifiers (SSIDs)** to announce their presence
- In theory, these should identify a particular wireless LAN
- In practice, SSID can be anything that's 2-32 octets long
- To join network, client must present SSID
- Crappy security mechanism for preventing interlopers:
  - Don't advertise SSID
  - Problem:
    - To join network, client must present SSID
    - This is not encrypted, even if network supports WEP or WPA

Wireless Security

Let’s sprinkle on some of that crypto magic sauce