Routing Security

Routing

What is Routing?

- A wants to send a packet to B
- The packet will be passed from router to router across the Internet's very complex topology
- What is the best route?
- What are the criteria for the "best" route? Fastest? Most reliable? Cheapest? Highest bandwidth?
- How is it determined?

Two Different Routes

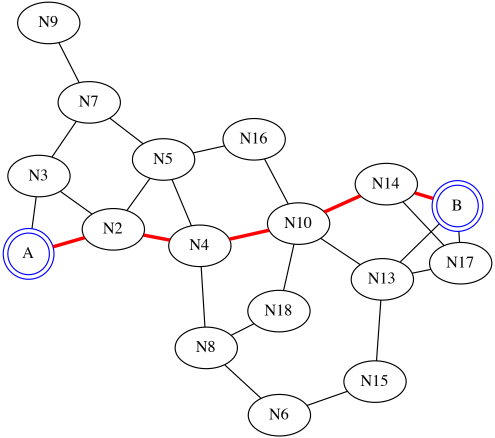

Two Different Routes

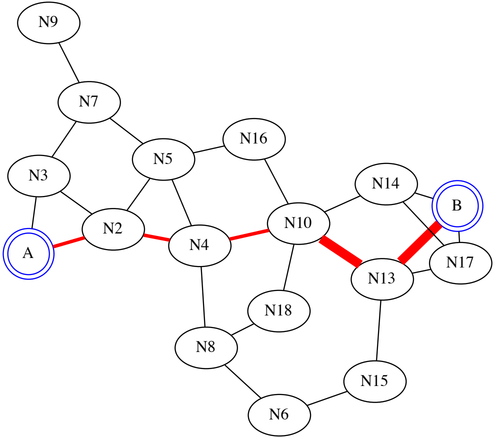

A Link Failure

Suppose the link from `N4` to `N10` fails—the Internet must route around the failure.

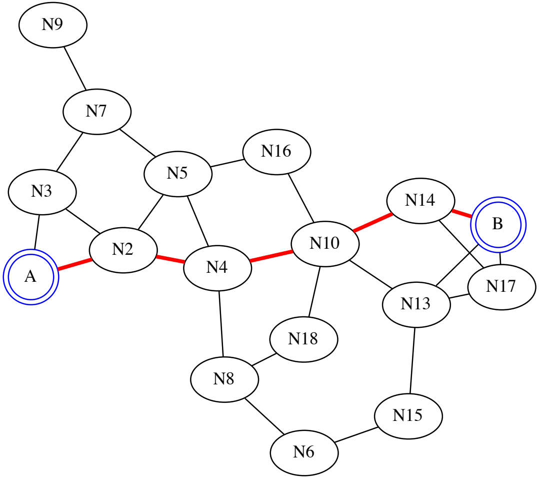

A Link Failure

Suppose the link from N4 to N10 fails—the Internet must route around the failure.

Note that the N2—N4 path is no longer used, because the N2—N5—N16—N10 route is "better" than N2—N4—N8—N18—N10

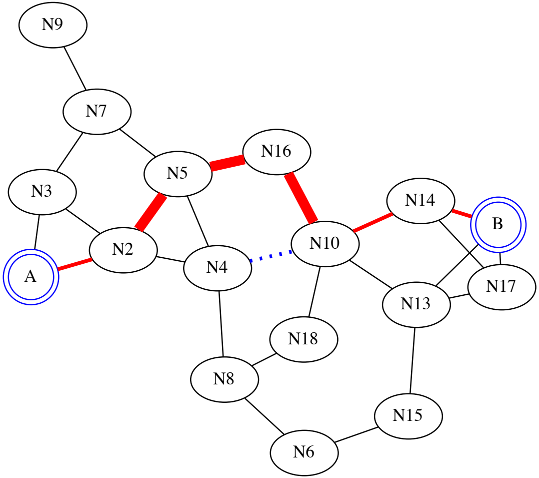

Routing Protocols

- Routers talk to each other to describe the network topology
- From this, each router *independently* calculates the best path for packets to follow
- Several different algorithms are in use

Bellman-Ford Routing

- Each node tells all of its neighbors: "Here are the destinations I'm connected to, with the following costs"
- Each node also tells its neighbors: "Here are the destinations I've heard about from my neighbors, with the link cost added to the cost I heard about"
- In other words: a node knows $\langle$destination, cost$\rangle$ pairs; it always chooses the cheapest path
- Worst-case complexity: $O(|E| \cdot |V|)$

Dijkstra's Algorithm

- Each node tells all of its neighbors: "Here are the links I have, with the following costs"
- Each node also tells its neighbors: "Here are the links and costs I've heard about from my neighbors"
- In other words, each node eventually learns the full topology of the graph
- Each node can then calculate the cheapest path to each destination
- Worst-case complexity: $O((|E| + |V|) \cdot \log |V|)$
- (Both algorithms are used on the Internet)

What is the “Cost” of a Link?

- There’s no one answer!
- Different ISPs will have different views
- Some will want to give high bandwidth to customers; others will want to optimize for latency
- Sometimes, it’s necessary to tinker with link costs to better balance traffic
- We need different solutions for different parts of the Internet

Two Kinds of Routing

Interior Routing

- Routing within an organization
- One cost metric, agreed upon internally
- Generally uses Dijkstra's Algorithm, as OSPF or IS-IS

Exterior Routing

- Used between ISPs and other organizations, called *autonomous systems* (AS)
- More-or-less Bellman-Ford using hop-count as BGP (Border Gateway Protocol)
- (Technically, it's a "distance vector" protocol, but it more closely resembles Bellman-Ford. However, many business and policy constraints are possible with BGP.)
- The actual mechanisms are *extremely* complex, to implement all sorts of complex policies
- The details are *way* beyond the scope of this course

InterISP Routing

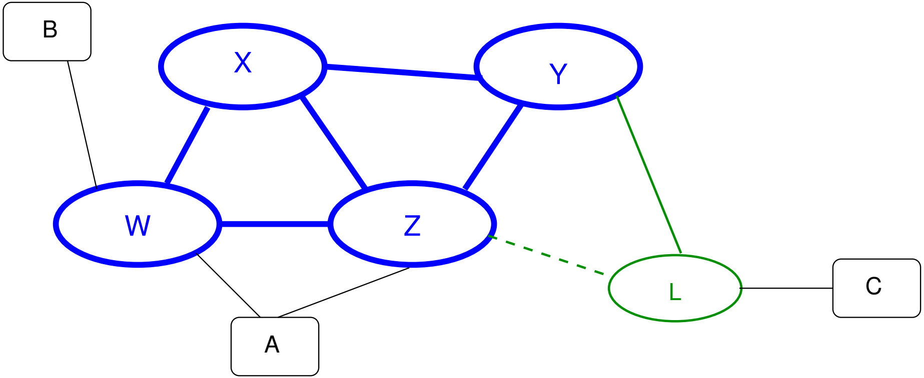

InterISP Routing

- “Tier 1” ISPs are peers, and freely exchange traffic.
- Small ISPs buy service from big ISPs.
- Different grades of service: link L-Z is for customer access, not transit. C→B goes via L-Y-X-W, not L-Z-W.
- A is multi-homed, but W-A-Z is not a legal path, even for backup.
- BGP is distance vector, based on ISP hops. Announcement is full path to origin, not just metric.
- An announcement for a network is [*net*,{AS-path}]

External Routing via BGP

- No common management (hence no metrics beyond hop count)
- No shared trust
- Policy considerations: by intent, not all paths are actually usable
- Columbia University has connections to multiple ISPs—but outsiders are not permitted to route traffic through Columbia; the links are for Columbia's use only

Exchange Points?

- You will sometimes read that ISPs interconnect at “exchange points”
- That used to be broadly true. Now, it’s much more for small ISPs
- Major—“Tier 1”—ISPs all serve the same major cities and have interconnections in most of them
- They use *hot potato* routing: get rid of packets as soon as possible

Hot Potato Routing

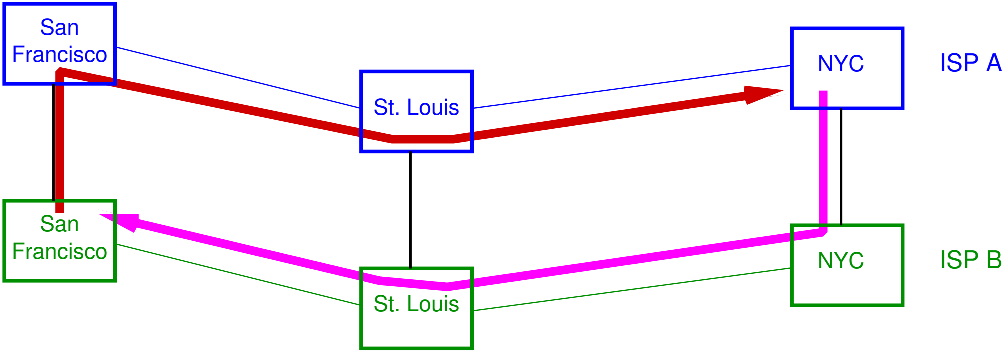

Each ISP hands off packets as early as it can to a neighbor—let them bear the expense of transporting it

Routing to What?

- We can't possibly calculate routes to 2^{32} different hosts—it would be far too expensive, in time and space
- Instead, we calculate routes to *networks*
- But what is a network?

Networks

- Think of an American phone number: +1 212-854-1754
- From outside the US, routing is done on the +1: the United States
- Within the US, routing is done on the area code, 212, Manhattan
- In Manhattan, the phone company uses 854; the university worries about the 1754 line
- The Internet does the same thing: it routes on the *network number*, but the boundary shifts
- Thus: AT&T owns network 12/8 and advertises that to the world
- Other ISPs look only at those 8 bits—but within AT&T, much more of the 32-bit number is used for routing
- This limits the scope of routing changes—and can hide some attacks

Longest Prefix Routing

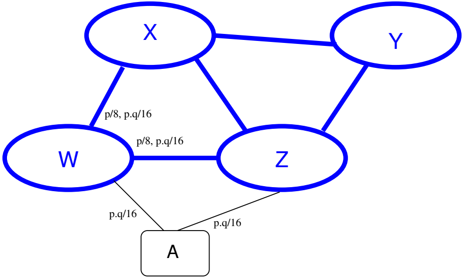

- ISP W owns network p/8; it delegates p.q/16 to customer A
- A wants a second link to ISP Z
- How does routing work? All traffic to p/8 should go to W.
- A announces p.q/16 to W and Z
- ISPs X and Y use that route to reach A instead of p/8—p.q/16 has a *longer prefix*

Routing Insecurity

What is Routing Security?

- Bad guys play games with routing protocols.
- Traffic is diverted.
    - Enemy can see the traffic.
    - Enemy can easily modify the traffic.
    - Enemy can drop the traffic.
- Cryptography can mitigate the effects, but not stop them.

Why So Little Work?

- It's a really hard problem.
- Actually, getting routing to work well is hard enough.
- It's outside the scope of traditional communications security.

How is it Different?

- Most communications security failures happen because of buggy code or broken protocols.
- Routing security failures happen despite good code and functioning protocols. The problem is a dishonest participant.
- Hop-by-hop authentication isn't sufficient.

The Enemy's Goal: Divert Traffic

The enemy wants traffic to pass through Z, for monitoring, modification, etc.

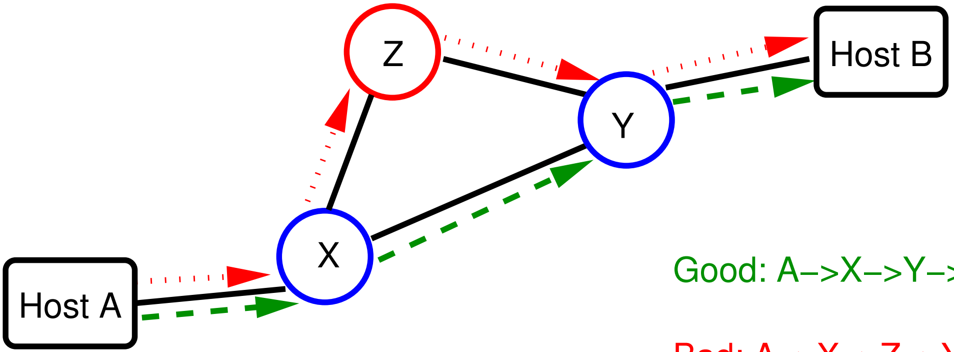

Good: A–>X–>Y–>B

Bad: A–>X–>Z–>Y–>B

But how can this happen?

Routing Protocols

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

Normal Behavior

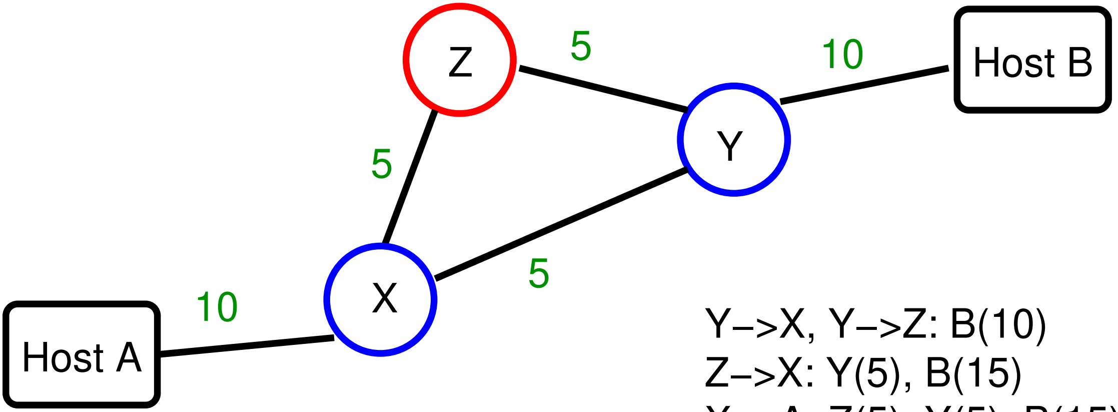

Y–>X, Y–>Z: B(10)
Z–>X: Y(5), B(15)
X–>A: Z(5), Y(5), B(15)

But Z Can Lie

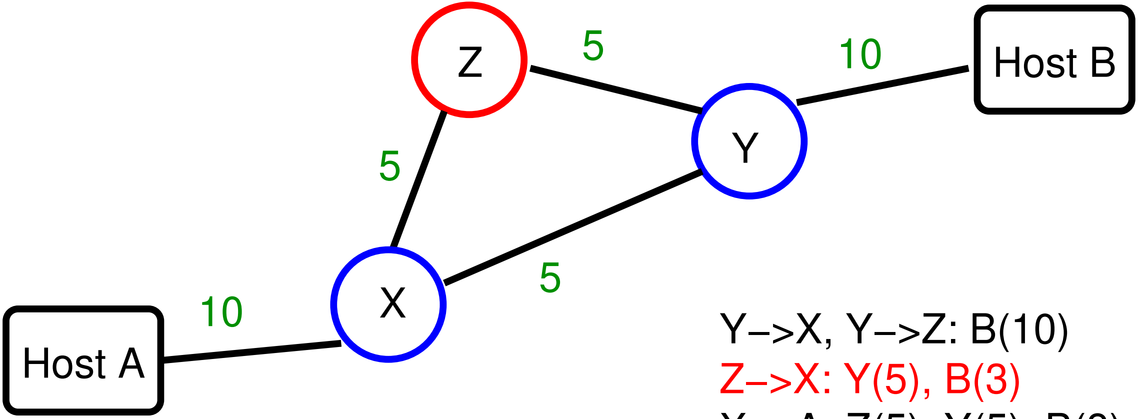

Y–>X, Y–>Z: B(10)
Z–>X: Y(5), B(3)
X–>A: Z(5), Y(5), B(8)

Note that X is telling the truth as it knows it

Why is the Problem Hard?

- X has no knowledge of Z's real connectivity.
- Even Y has no such knowledge.
- The problem isn't the link from X to Z; the problem is the information being sent. (Note that Z might be deceived by some other neighbor Q.)

Monitoring Traffic

Suppose that Z wants to monitor traffic from A to B. It uses a routing attack to divert traffic from intermediate router X.

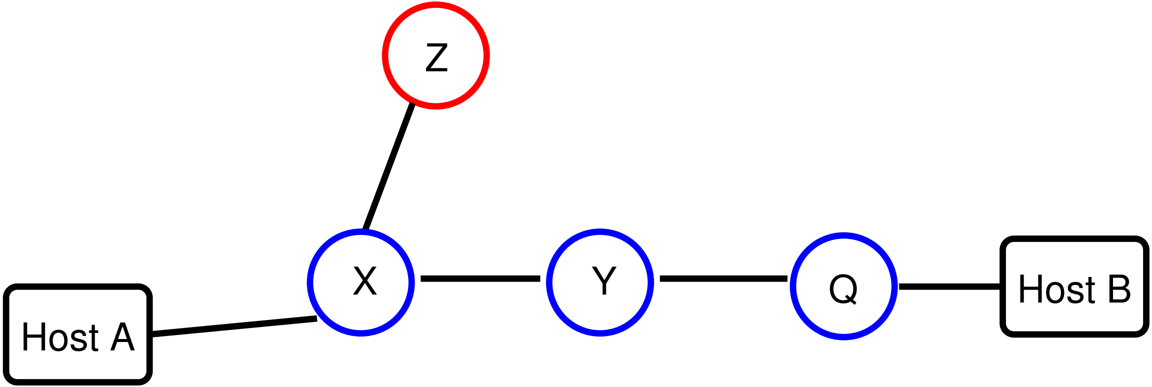

How does the traffic get to B? If Z sends it to X, it will just loop back.

Using a Tunnel for Packet Reinjection

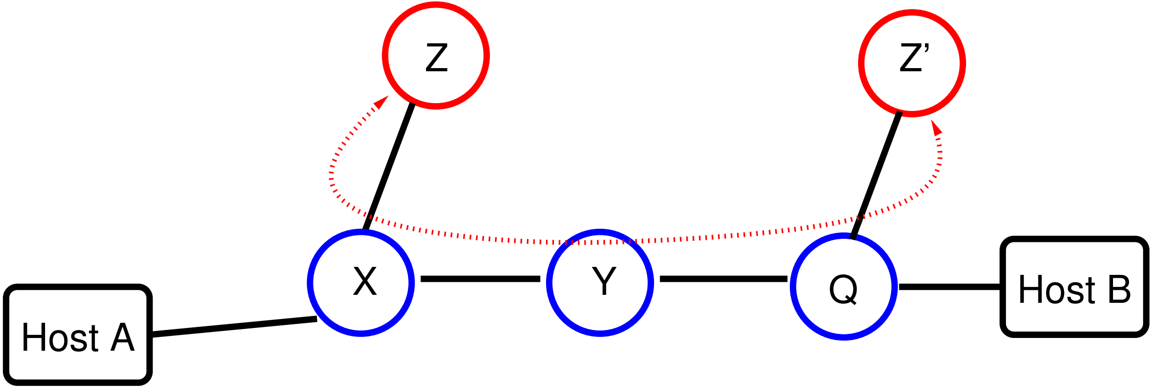

Set up another router near B, and *tunnel* traffic using a virtual link from Z to Z′.

How Do You Secure Interior Routing?

- Shared secrets guard against new machines being plugged in, but not against an authorized party being dishonest.
- Solution: digitally sign each routing update (expensive!). List authorizations in certificate.
- ☞ The authorizations describe what networks a given router is *allowed* to announce
- (Experimental RFC by Murphy et al., 1997.)
- Note: everyone sees the whole map; monitoring station can note discrepancies from reality. (But bad guys can send out different announcements in different directions.)

It’s Never Done

- To my knowledge, no one has ever used this
- Most internal security problems are attacks on hosts
- Maybe some routers have been compromised to do this—but we haven’t heard about it. . .

It's Never Done

- To my knowledge, no one has ever used this
- Most internal security problems are attacks on hosts
- Maybe some routers have been compromised to do this—but we haven't heard about it. . .

☞ But it was inadvertent routing errors that started me doing network security: "This accident—what if someone did this on purpose?"

Attacks on BGP

- They’re real
- They’ve been going on for more than 20 years
- There are a variety of techniques available to stop them—but not all are used

The AS7007 Incident

- A router in Florida “deaggregated” many routes
- The cause—malice, ignorance, or a bug—is not clear
- That is, it broke up, say, a /8 into 2^{16} /24s
- A /24 is a longer prefix than a /8, so it’s used preferentially
- Most of the traffic on the Internet tried to head to one small ISP. . .

The Pakistan YouTube Incident

- The government of Pakistan wanted to ban YouTube because of what they regarded as objectionable videos
- The mechanism: change the interior OSPF routing tables to discard packets intended for YouTube
- Somehow—just how isn't clear—these routes leaked into the global BGP system
- Result: most traffic intended for YouTube went to Pakistan and was discarded

Spammers and BGP

- Spammers don't want to get blocked or traced
- Some use BGP hijacking: announce some prefixes, send spam, withdraw the announcement
- You can't tell where the spam came from!

Securing Routing

Filtering

- ISPs can filter route advertisements from their customers.
- Filtering can happen on prefix lengths, too
- Doesn't always happen: AS7007 incident, spammers, etc.
- Not feasible at some links—how does Z know that A is entitled to p.q/16?
- Complex

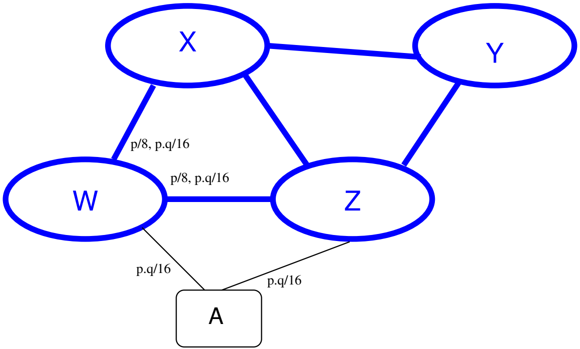

RPKI: Router Public Key Infrastructure

- Every AS has a certificate that lists its legal IP address blocks
- These certificates are used to sign route origin announcements
- And—they can be used to issue sub-certificates for subsets of their ranges
- Thus: `W` has a certificate for p/8, i.e., p.0.0.0–p.255.255.255.
- They can issue a certificate to A for p.q/16
- These certificates are both route-signing *and* CA certificates!

The Limits of RPKI

- A signs its announcement of p.q/16
- `W` adds its AS to the path, as does X
- V sees `[p.q/16,{X,W,A}]` from X
- Y sees `[p.q/16,{V,Z,A}]`—but it's evil, so it announces `[p.q/16,{Y,A}]` to V
- This is a shorter path, so V believes it
- RPKI protects the *origin*, not the *path*

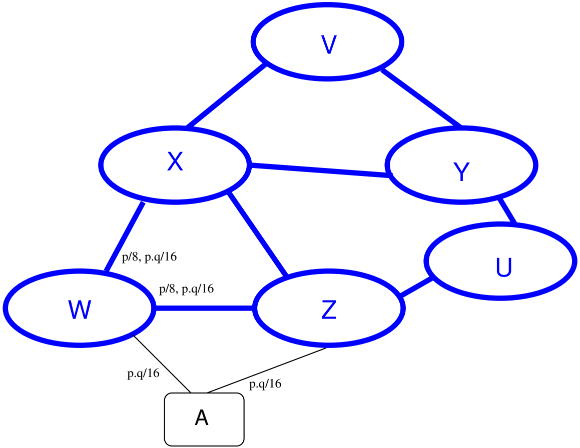

Path Protection: BGPsec

- Each node signs its *full* announcements, including the next hop
- A sends $\{p.q/16, \{W, A\}\}_A$ to W and $\{p.q/16, \{Z, A\}\}_A$ to Z
- W sends $\{p.q/16, X, \{p.q/16, \{W, A\}\}_A\}_W$ to X and W sends $\{p.q/16, Z, \{p.q/16, \{W, A\}\}_A\}_W$ to Z, etc
- Everything is signed—evil nodes can't cut out pieces of a route

Problems with BGPsec

- Lots of digital signatures to calculate and verify.
  - Can use cache
  - Verification can be delayed
- Calculation expense is greatest when topology is changing—i.e., just when you want rapid recovery.
- Lots of router RAM is needed
- What about secure route withdrawals when link or node fails?
- BGP is already horribly complex (there's much more announced than just the AS path); this makes it worse

No Uptake!

- Most security people think that BGPsec is necessary
- There are routing security incidents every day, and we’ve already seen attackers cutting out and reusing RPKI announcements
- But—it’s complex to operate, and there are new failure modes
- Lack of an economic model for deployment

New Failure Modes

- If your certificate expires, your site is off the air
- (Can you get adequate customer care?)
- If a government forces revocation of your certificate, you're off the air

No Economic Model

- Why should an ISP install BGPsec?
- A customer could pay them—but that only protects the customer if all other ISPs are using BGPsec
- BGP signatures are stripped off at the boundary between BGPsec speakers and those who don't speak it
- A, W, X, Z, and V speak BGPsec; Y and U don't. V hears routes to A from X and Y—which should it prefer?

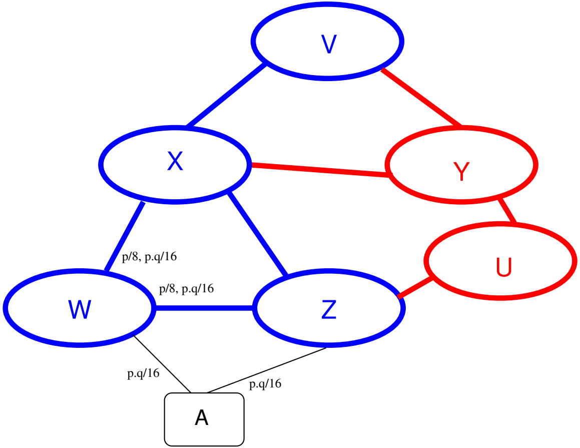

An Unsolved Problem

- We have the science to protect routing
- The protocol engineering has been done, with the cooperation of cryptographers, ISPs, and router engineers
- But most ISPs don’t want it because there’s no customer demand, and there’s no strong security story for limited-scale deployment
- And routing security problems keep happening...

Questions?

(Northern flicker, Riverside Park, January 18, 2020)