Routing Security

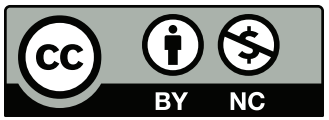

What is Routing Security?

- 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.

History of Routing Security

- Radia Perlman's dissertation: *Network Layer Protocols with Byzantine Robustness*, 1988.
- Bellovin's "Security Problems in the TCP/IP Protocol Suite" (1989)
- More work starting around 1996.
- Kent et al., 2000 (two papers).
- Many more since then, but no adoption

Why So Little Work Until Recently?

- 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?

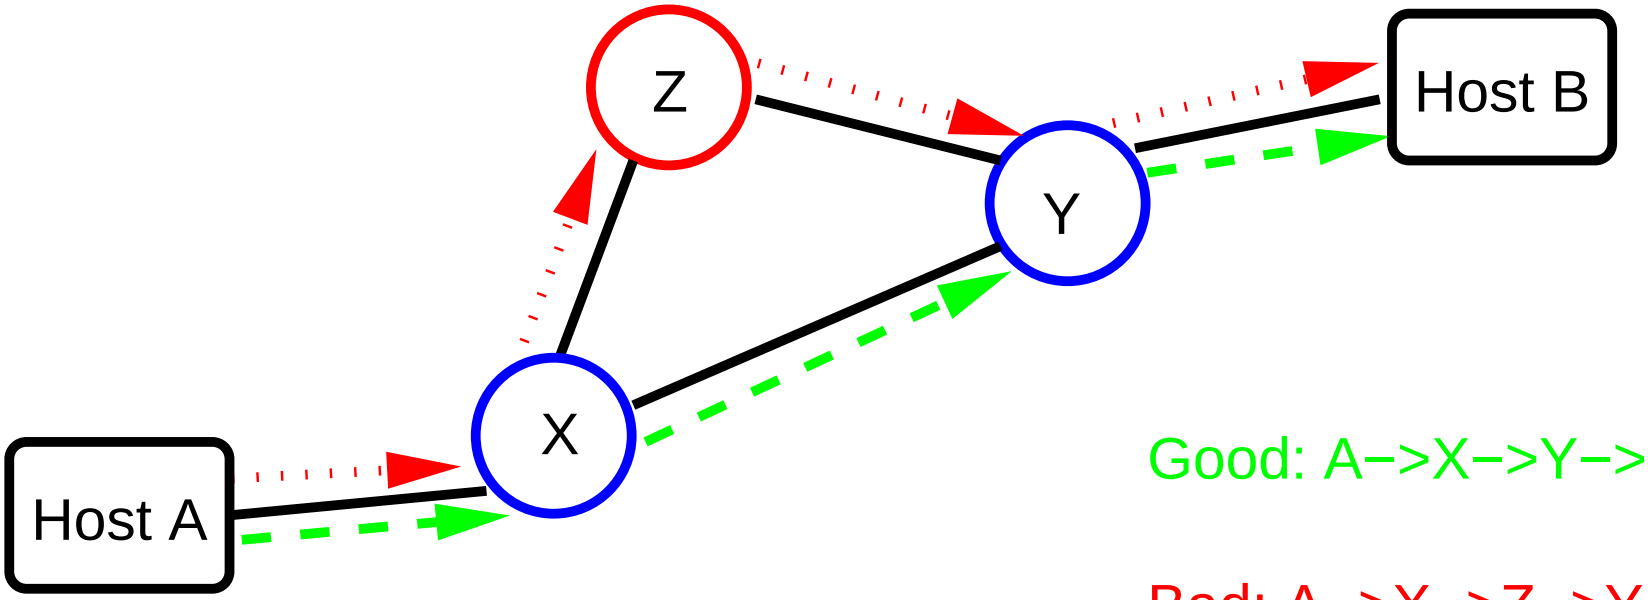

But how can this happen?

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.
- Attacker can lie to other routers.

Normal Behavior

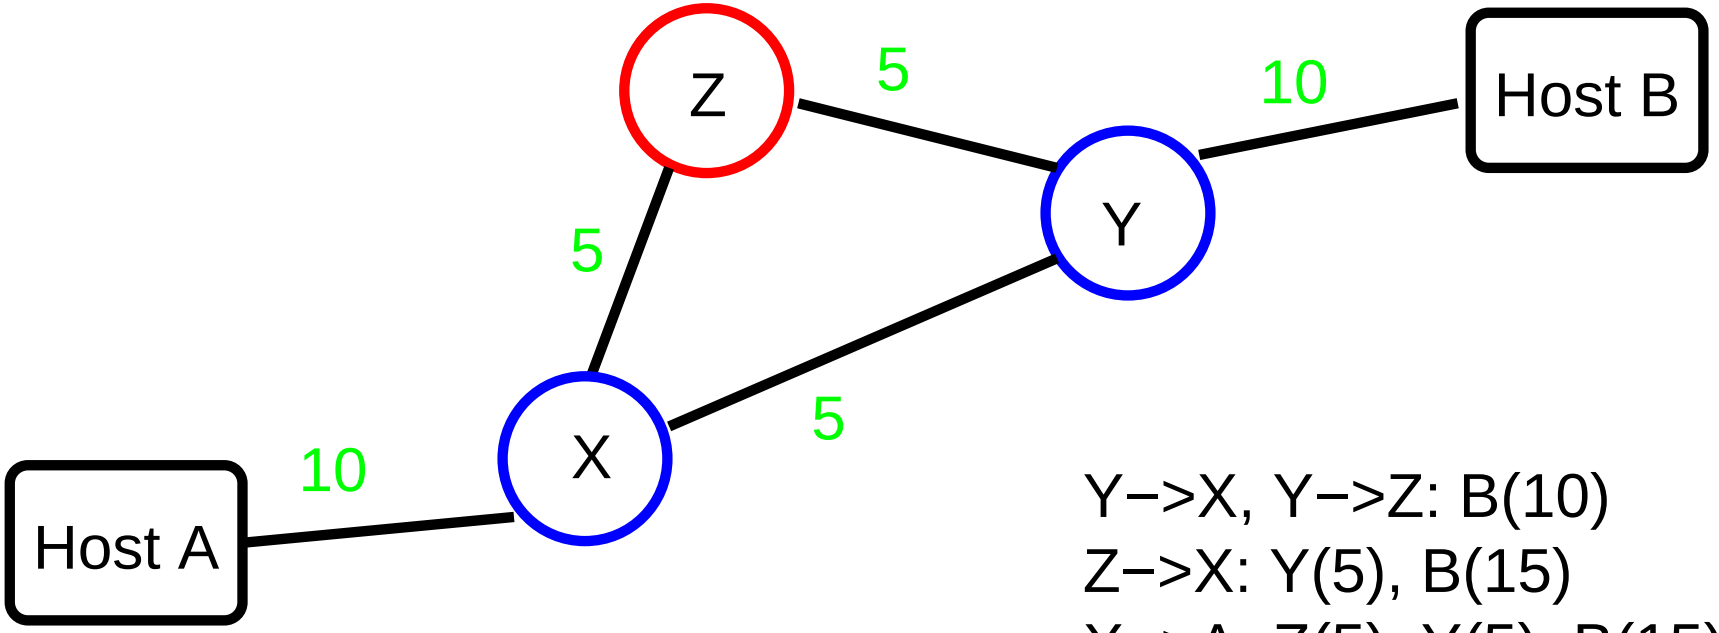

But Z Can Lie

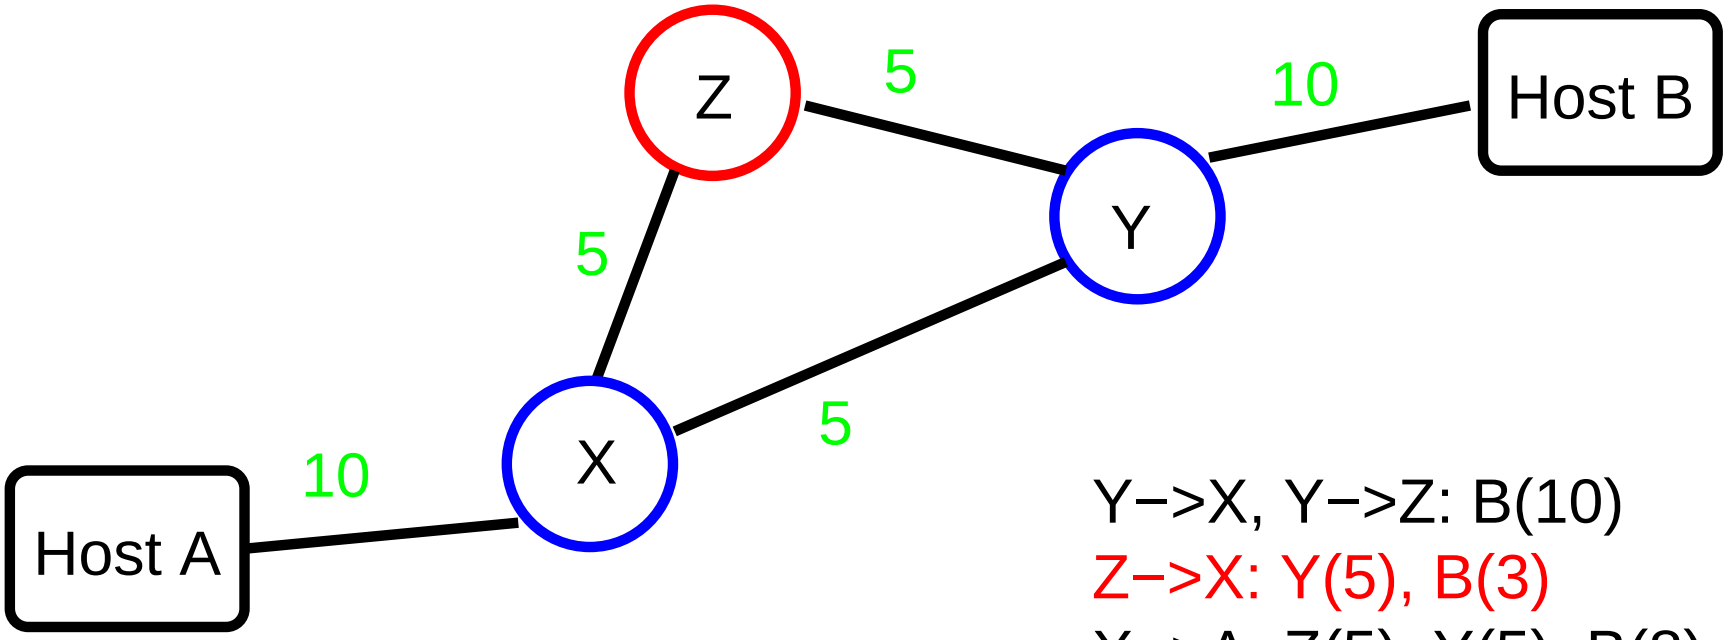

Note that X is telling the truth as it knows it.

Using a Tunnel for Packet Reinjection

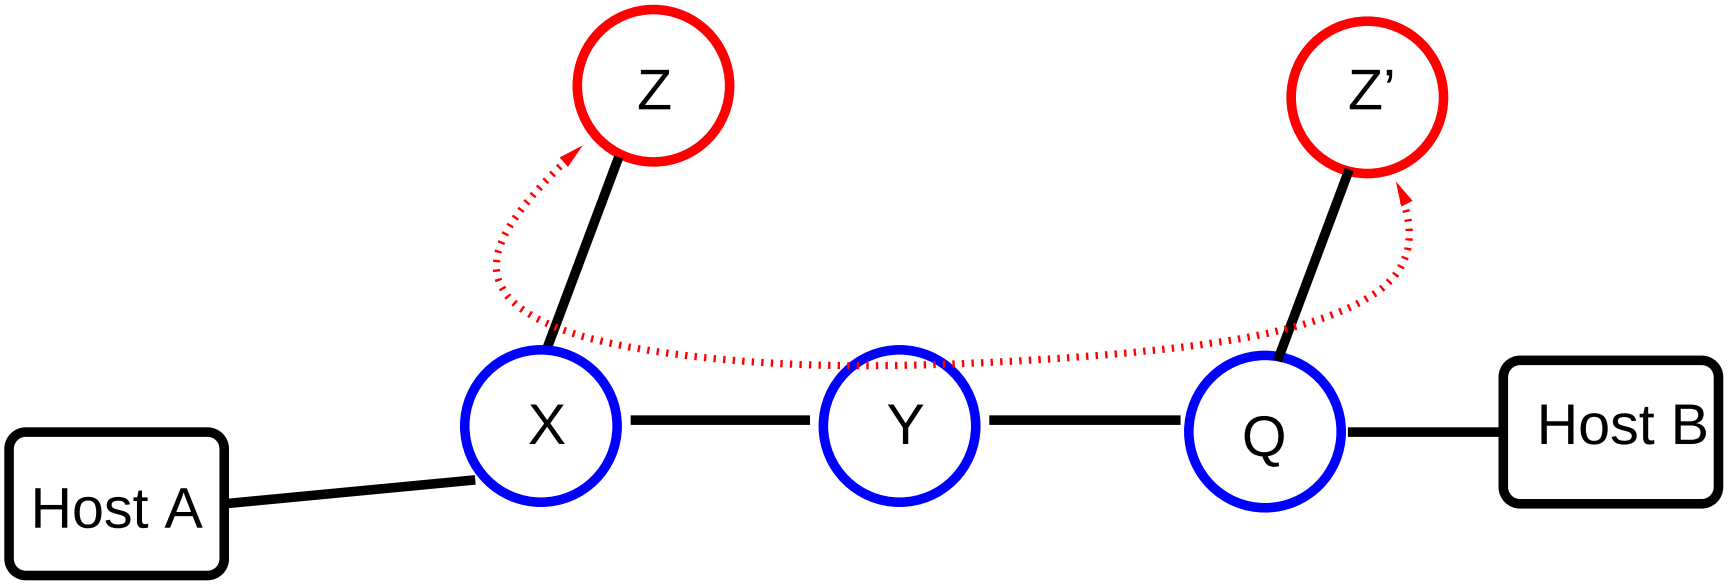

Why is the Problem Hard?

- 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.)

Routing in the Internet

- Two types, internal and external routing.
- Internal (within ISP, company): primarily OSPF.
- External (between ISPs, and some customers): BGP.
- Topology matters.

OSPF (Open Shortest Path First)

- Each node announces its own connectivity. Announcement includes link cost.
- Each node reannounces all information received from peers.
- Every node learns the full map of the network.
- Each node calculates the shortest path to all destinations.
- Note: limited to a few thousand nodes at most.

Characteristics of Internal Networks

- Common management.
- Common agreement on cost metrics.
- Companies have less rich topologies, but less controlled networks.
- ISPs have very rich—but very specialized—topologies, but well-controlled networks.
- Often based on Ethernet and its descendants.

How Do You Secure OSPF?

- Simple link security is hard: multiple-access net.
- 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.
- 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.)

Address Authorization Certificate

- Each router has certain interfaces and hence direct network reachability
- Each router therefore has a certificate binding its public key to its valid addresses
- Note well: the CA has to *know* the proper addresses for each router
- But that's the norm in OSPF environments

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.

POP Topology

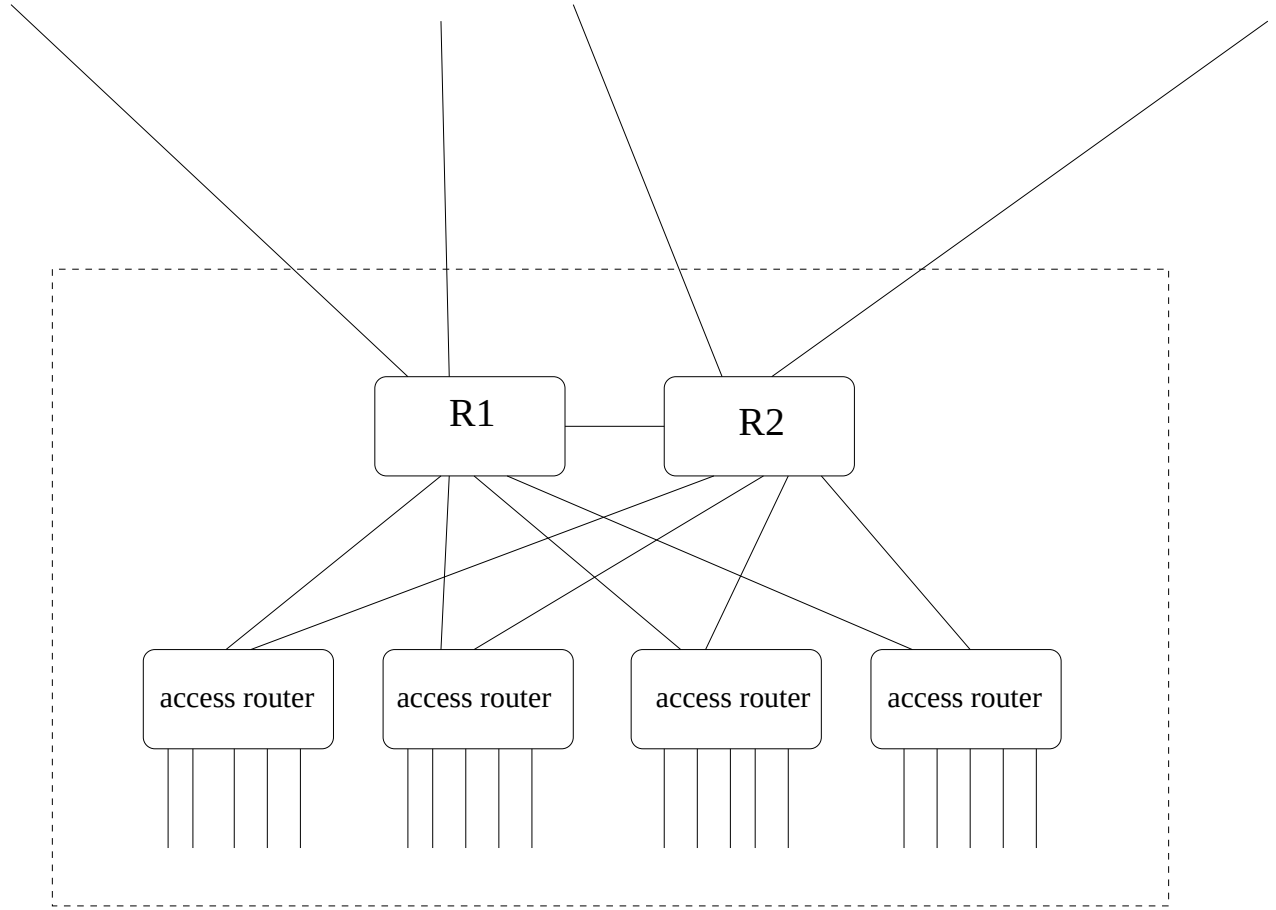

Noteworthy Points

- A lot of attention to redundancy.
- Rarely-used links (i.e., R1→R2)
  Link cost must be carefully chosen to avoid external hops.
- May have intermediate level of routers to handle fan-out.

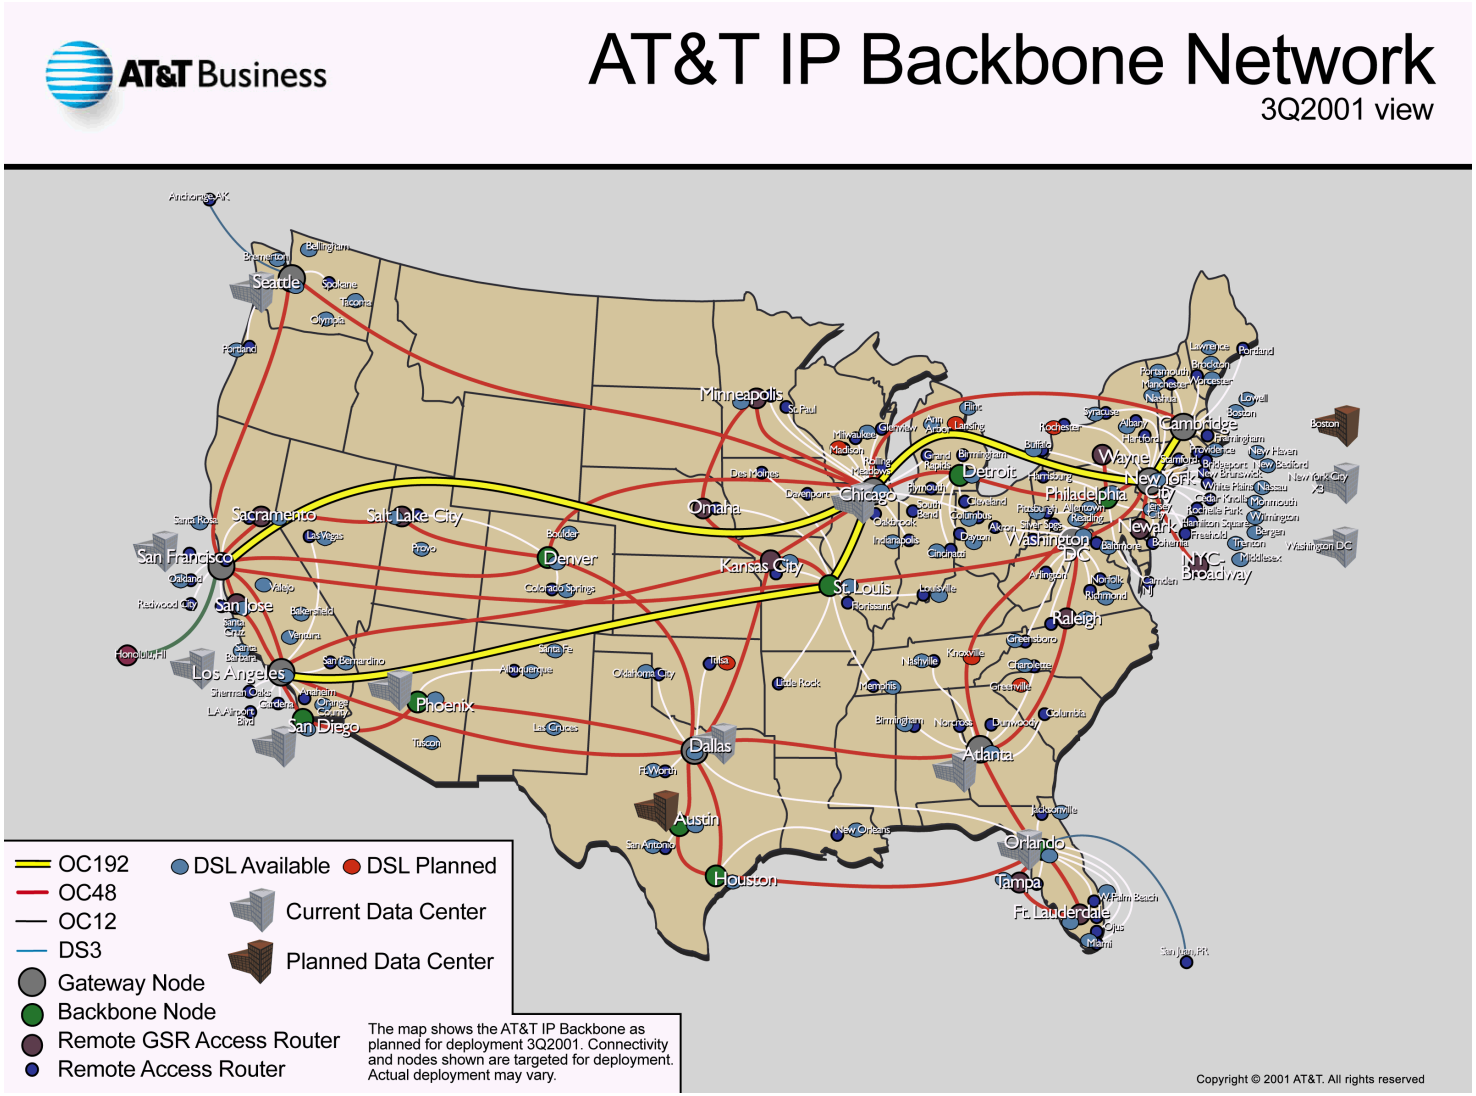
AT&T Business
AT&T IP Backbone Network
3Q2001 view
Seattle
Portland
Sacramento
San Francisco
San Jose
Los Angeles
San Diego
Phoenix
Salt Lake City
Denver
Omaha
Minneapolis
Chicago
Detroit
Wayne
Cambridge
New York City
Philadelphia
Newark
Washington DC
NYC Broadway
Kansas City
St. Louis
Dallas
Houston
Austin
Atlanta
Raleigh
Orlando
Tampa
Ft. Lauderdale
OC192
OC48
OC12
DS3
Gateway Node
Backbone Node
Remote GSR Access Router
Remote Access Router
DSL Available
DSL Planned
Current Data Center
Planned Data Center
The map shows the AT&T IP Backbone as planned for deployment 3Q2001. Connectivity and nodes shown are targeted for deployment. Actual deployment may vary.
Copyright © 2001 AT&T. All rights reserved

InterISP Routing

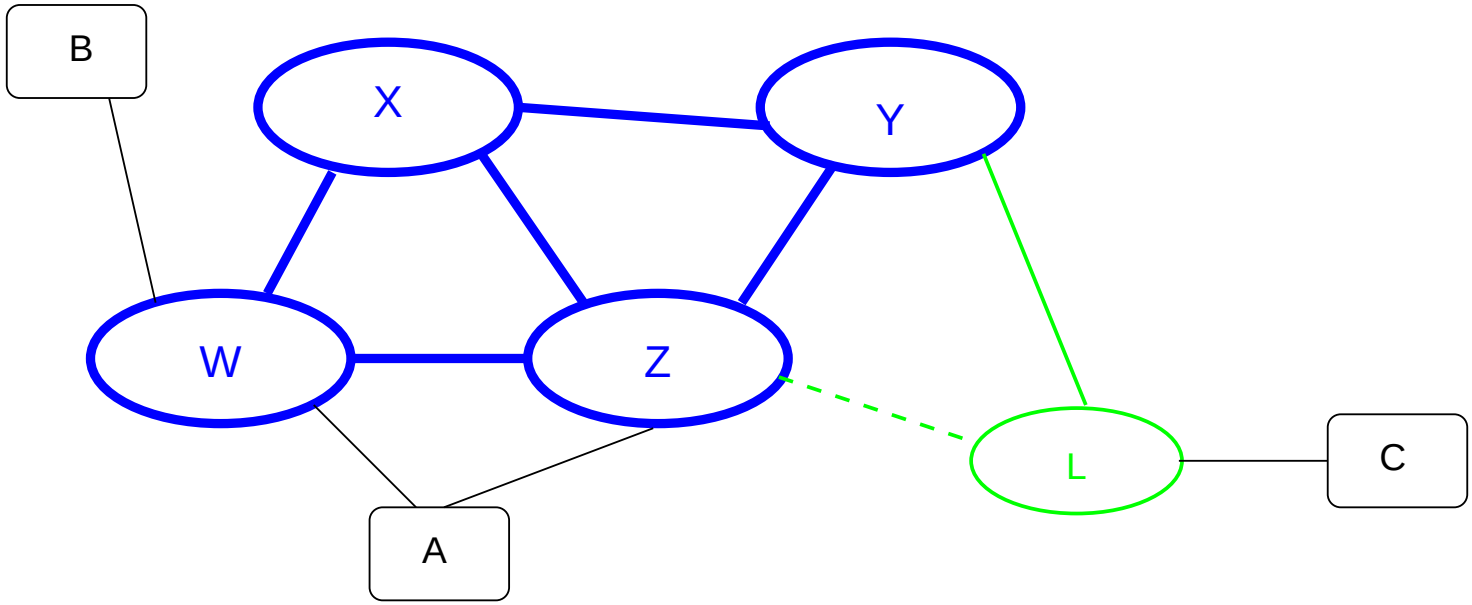

InterISP Routing

- "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.

Path Vectors

- Route advertisements contain a prefix and a list of ASs to traverse to reach that prefix
- Example: if B owns address block 10.0/16, L would see ⟨10.0/16, {Y,X,W,B}⟩
- ASs do not see paths filtered by upstream nodes. Y sees ⟨10.0/16, {X,W,B}⟩ and ⟨10.0/16, {Z,W,B}⟩; if Y only forwards the former to L, L will know nothing of the path via Z

Policies

- ISPs have a great deal of freedom when choosing the "best" path
- While hop count is one metric, local policies (i.e., for traffic engineering) count more
- These policies — in general, not disclosed publicly — affect with path neighbors will see

Long Prefixes and Loop-Free Routing

- Routers ignore advertisements with their own AS number in the path
- This is essential to provide loop-free paths
- Routers use longest match on prefixes when calculating a path
- These two facts can be combined to form an attack

Longer Prefix Attack

- Suppose B owns 10.0/16. Z sees ⟨10.0/16, {W,B}⟩
- A advertises ⟨10.0.0/17, {A,W}⟩
- Z will route packets for 10.0.0/17 to A — it has a longer prefix
- W will never see that path, and hence won't pass it to B — the path (falsely) contains W, so it will be rejected by W

Filtering

- ISPs can filter route advertisements from their customers.
- Doesn't always happen: AS7007 incident, spammers, etc.
- Not feasible at peering links.

Secure BGP (Kent et al.)

- Each node signs its announcements.
- That is, X will send $\{W\}_X, \{Y\}_X, \{Z\}_X$.
- W will send $\{B\}_W, \{A\}_W, \{X\}_W, \{X : \{Z\}_X\}_W$.
- Chain of accountability.

Problems with SBGP

- 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. (About 275K routes. . .)
- How to deal with route aggregation?
- What about secure route withdrawals when link or node fails?
- Dirty data on address ownership.

Certificate Issuance

- Who issues prefix ownership certificates?
- Address space comes from upstream ISP or RIRs
- RIRs really are authoritative — hence they're a monopoly
- If an RIR makes a mistake, the prefix is off the air
- Is this a risk worth taking?

Certificate Tree

- The *RIR*s (Regional Internet Registries) give addresses to big ISPs and big end users
- Accordingly, the RIRs should issue certificates
- (Really, it should be ICANN, but the politics of that are too painful)
- Small ISPs and small customers get address space from their own ISPs
- Every ISP is thus a certificate holder and a certificate issuer
- These are *authorization certificates*, not *identity certificates*

Authorization Certificates

- The identity of the certificate holder is irrelevant
- What matters is the authorization: the certificate contains IP address ranges
- The signing party has its own certificate listing larger ranges of IP addresses, and hence the right to delegate them

Signed Origin BGP

- Suppose only the origin was digitally signed: $\langle$10.0/16, B$\rangle$
- In addition, all polices are (securely) published in some database
- Receiving node verifies origin, then compares received path against all policies
- Query: is the received path *consistent* with policies?
- Advantage: many fewer signatures

Problems with SOBGP

- Sill have monopoly RIRs
- ISPs don't like to publish policies
- Clever attackers can play games in the middle of the path

Happy Packets

- Philosophy: don't worry too much about routing security
- Crucial metric: do packets reach their destination?
- What about confidentiality? If it matters, encrypt end-to-end
- But what about traffic analysis?

Link-Cutting Attack (Bellovin and Gansner)

- Suppose that we have SBGP and SOSPF.
- Suppose the enemy controls a few links or nodes. Can he or she force traffic to traverse those paths?
- Yes. . .

Is Link-Cutting Feasible?

- Attacker must have network map.
  Easy for OSPF; probably doable for BGP—see "Rocketfuel" paper.
- Can attacker determine peering policy? Unclear.
- How can links be cut?
  Backhoes? "Ping of death"? DDoS attack on link bandwidth?

Sample Link-Cutting Attack

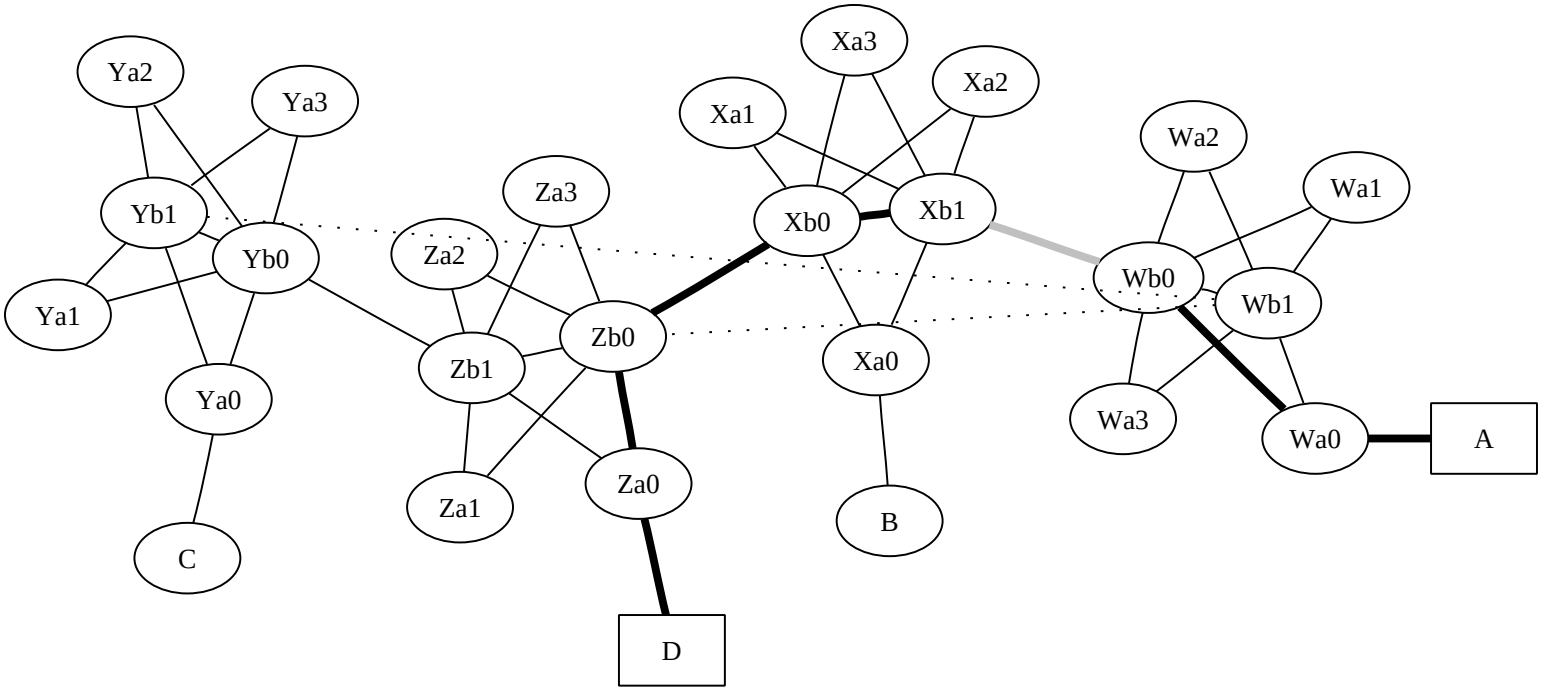

Cost of Link-Cutting Attacks on the Backbone

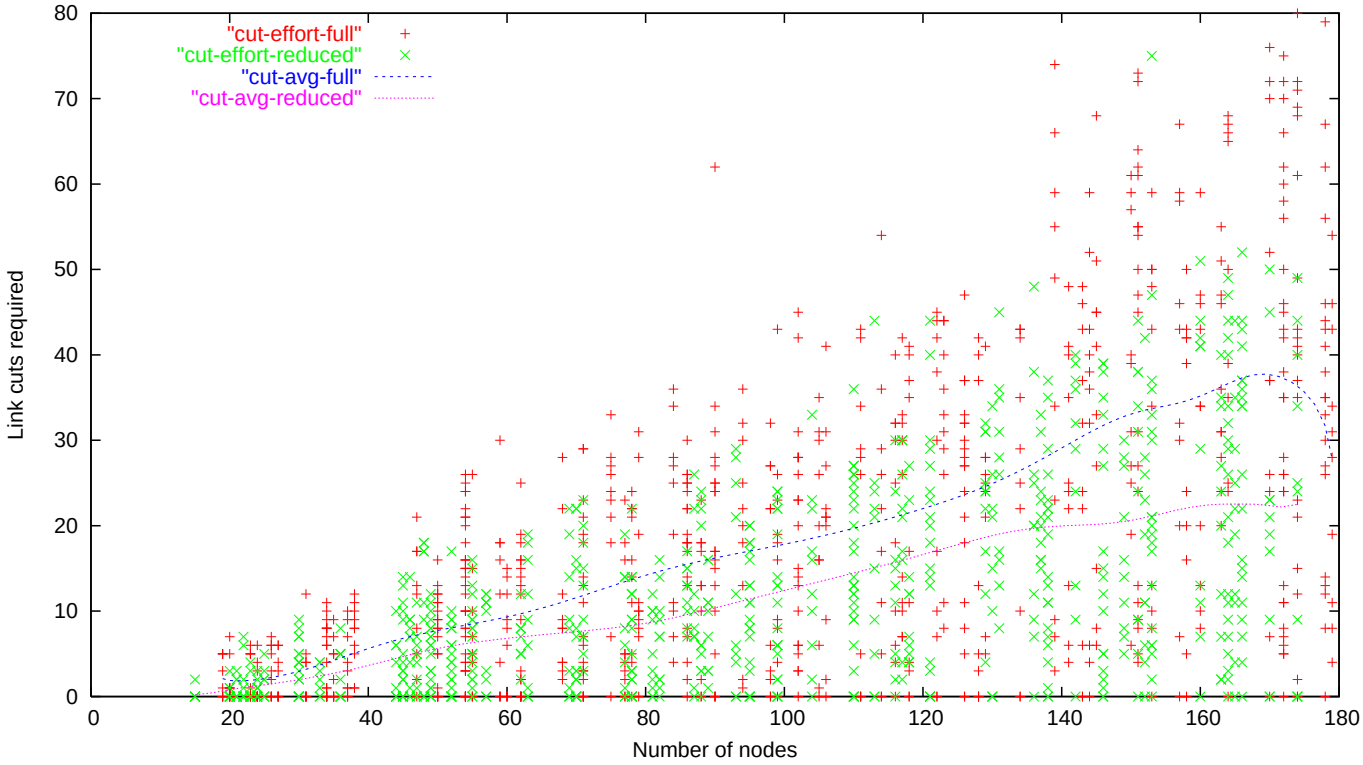

Defenses

- Hard to defend against—routing protocols are doing what they're supposed to!
- Keeping attacker from learning the map is probably infeasible.
- Feed routing data into IDS?
- Link-level restoration is a good choice, but can be expensive.
- Others?

Conclusions

- Routing security is a major challenge.
- Mentioned specifically in White House Cybersecurity document.
- Lots of room for new ideas.