Routing Security: an Oversimplification

EOF / Istanbul

2006.04.25

Randy Bush <randy@psg.com>
Steve Bellovin <smb@cs.columbia.edu>

<http://rip.psg.com/~randy/060425.eof-routesec.pdf>

Router Security

- Go to any Routing Ops Security Tutorial
- TCP/MD5 session protection
- ACLs on everything
- ssh, not telnet. no http, …
- DDoS on links
- Attacks on routers as hosts

Routing Security?

- The unique threat is attackers **using routing protocols**
  - To divert traffic
  - To alter traffic
- We have some ability to lessen the danger, but not enough!

History of Routing Security

- Radia Perlman dissertation: *Network Layer Protocols with Byzantine Robustness*, 1988
- Bellovin: *Security Problems in the TCP/IP Protocol Suite*, 1989
- Work accelerates 1996
- Kent et alia two papers in 2000
- Endless jawing in the IVTF

Why So Little Progress?

- The problems are technically very difficult
- Simple routing is already a very complex operational issue
- It is not traditional communications security
- Installed base & transition problem
- Under-motivated vendor$

What is so Different?

- Well-studied communication and host security issues are buggy code and/or bad protocol design
- Routing is vulnerable with good code and good protocols
- **The problem is a dishonest peer**
- Hop-by-hop authentication is not sufficient

Diversion Attack

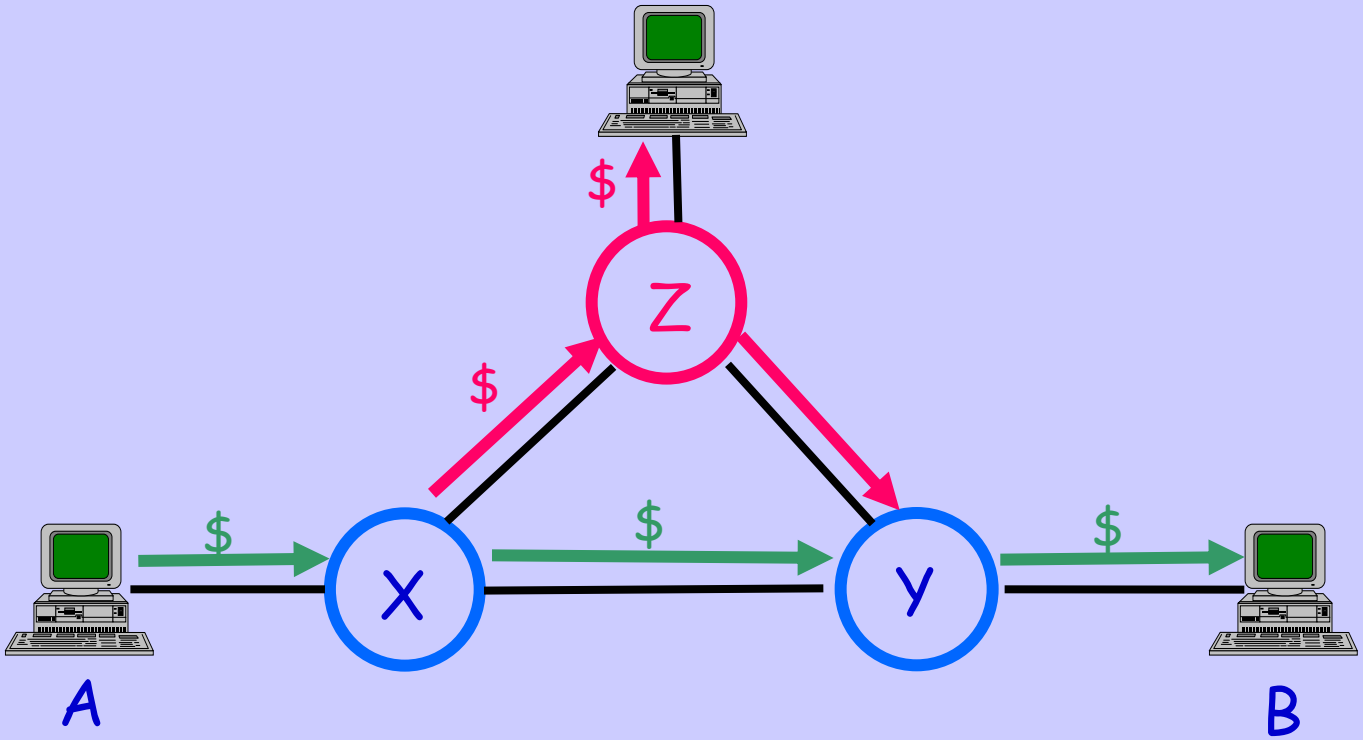

Expected Path - A->X->Y->B

Diverted Path - A->X->Z->Y->B

How does Attacker Do It?

- Routers select lowest cost path toward destination on a hop by hop basis
- Attacker 'owned' router lies about cost
- And we must assume that random routers can be owned

How Does Z Do It?

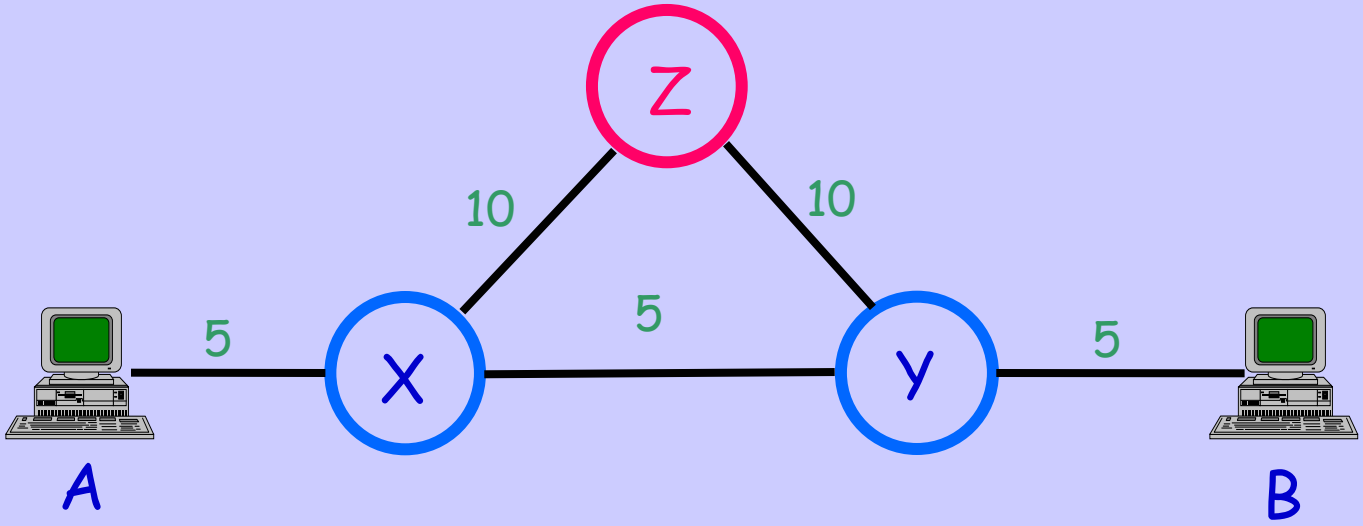

Y tells X and Z that costs are B:5
X tells A and Z that costs are Y:5 B:10
Z tells X that costs are Y:10 B:15

Z tells X that costs are Y:10 **B:5**

X now sends B's traffic to Z!!!

Why is this a Hard Problem?

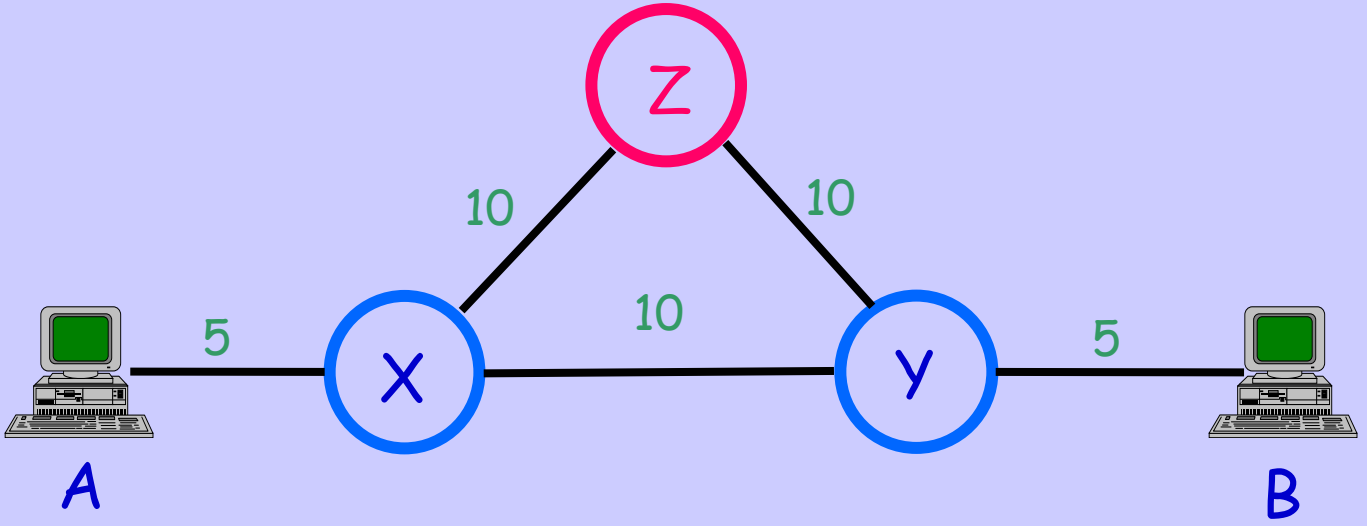

- X does not really know Z's links
- X does not really know Y's links
- They trust each other re costs!

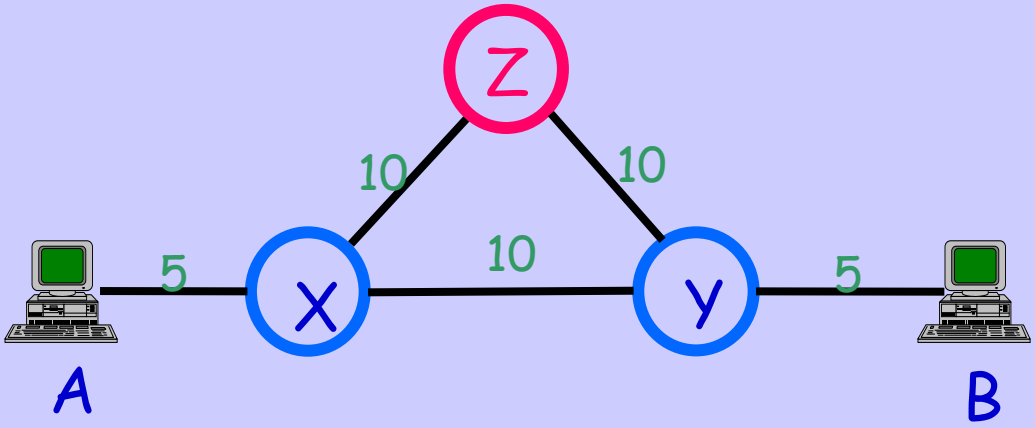

- Validating IP prefix ownership does not help, as Z is not lying about B's owning it
- Using IRR-like peering map does not help, as Z is not lying about who connects to whom

One Approach

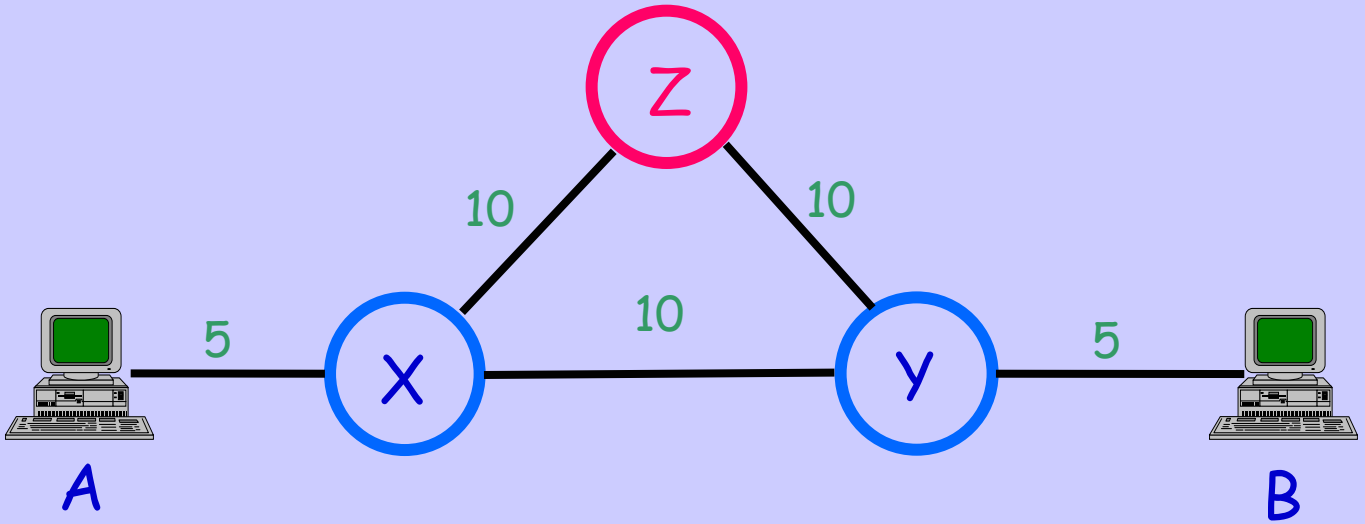

- B cryptographically signs the message to Y Sb(Y->B=5)
- Y signs messages to X and Z encapsulating B's message Sy(X->Y=10 Sb(Y->B=5)) and Sy(Z->Y=10 Sb(Y->B=5))
- Z can only sign Sz(X->Z=10 Sy(Z->Y=10 Sb(Y->B=5)))
- Now X can verify paths and costs
- **Forward path signing** solves the 'simple' case

Costs

- Very crypto-CPU-intensive
  - Use caching
  - Use pre or delayed validation
  - Moore's 'Law' is your friend
  - Most announcements are boring
- Expense is highest when routing is changing, just when we need validation the most ☹

Trust Issues

- How to know ISP X 'owns' address space A?
- How to know AS 42 has the right to originate X's address space A
- How to know AS 666 legitimately received the announcement from X?
- So there are three classes of trust,
  - IP address ownership
  - Right to Originate
  - Right to Propagate

Address Space Ownership

- Luckily, IP space delegation is a natural hierarchy
- RIR signs address allocations to ISPs/LIRs using RIR certificate
- ISP/LIR signs allocations to downstreams using its ISP/LIR certificate
- ISP/LIR signs allocations to End Sites

Thanks

- Steve Bellovin, whose ideas and work I liberally stole
- NSF via award ANI-0221435
- Internet Initiative Japan