


Security III: Availability, DDoS, and Routing Security

15-441 Fall 2019

Profs Peter Steenkiste & **Justine Sherry**

**Carnegie
Mellon
University**

Your Feedback



I like to think of Dr. Weiss as my teaching teacher

2. Prizes (candy, stickers, t-shirts) are motivating and help students stay involved during lecture. (66% agreed)

34% of you are unsatisfied with my candy. I have brought more variety.

2. Ensure consistency when answering questions on Piazza. (90% agreed)

I suspect consistency is due to changes in the course.

This year: going to ask TAs to leave answers to “lead” project TAs when they are unsure

Next year: hopefully fewer changes in the course mean all TAs will be on same page



Your Feedback



I like to think of Dr. Weiss as my teaching teacher

1. *Projects are very ambitious and need more specifics on what needs to be done. (<100% agreed)*

I'm going to ask for more feedback at the end of class today — did you feel the same way about Project 2?

Re: Project 1: What about offering a “highlighted” version of the RFC to draw your attention where to look?



Your Feedback




I like to think of Dr. Weiss as my teaching teacher

- Provide more tests to run locally (allowing for better debugging), or
- Increased submission limit (other classes allow 20+ submissions).
- CGI program doesn't work well. Tests in the handouts can be improved.
- "Currently, the feedback from Autolab is not sufficient for debugging."
- "A lot of time spent trying to get P2 testing working."

"A hill worth dying on": An issue to pursue with wholehearted conviction and/or single-minded focus, with little or no regard to the cost.





When staging doesn't match production.



Why are my Twitter Friends making these jokes?

- In developing a real product, the only people who “give you tests” to try against are your users.
- But “deploying” code to them is not free!
 - When things break, your users get angry
 - Send you mean emails
 - Quit your project and move to a competitor
 - **You might get fired if you make this happen with any frequency.**



Why are my Twitter Friends making these jokes?

- Best industry practices:
 - Try to figure out everything that could possibly go wrong before deploying.
 - Write tests to make sure those bad things don't happen.
 - *Then* “deploy”.



Why we have autolab limits

- This is a senior-level systems class. We are almost about to send you out to the big leagues!
- Think of “autolab” as “deploying” — you get fast and immediate feedback from your users. But it’s not free!
 - Then again, in industry, if you “test” on all your users with buggy code a dozen times in one week, you’ll probably get fired.
- We’re still giving you *lots* of submissions.
 - But we want you to slow down and think about fixing things before deploying.



What I see on my side

- Very very few of you are getting close to your autolab limits.
 - Some of you could even benefit from submitting a little more often :-)
- And yet... we're also seeing some of the highest project scores I've seen in the three times I've taught this course.
- The training wheels are working! You're becoming much stronger developers!



What were the four requirements for a secure communications channel?



What do we need for a secure comm channel?

- Authentication (Who am I talking to?)
- Confidentiality (Is my data hidden?)
- Integrity (Has my data been modified?)
- Availability (Can I reach the destination?)



A Chinese ISP momentarily hijacks the Internet (again)

By Robert McMillan

IDG News Service | Apr 8, 2010 5:59 PM PT

For the second time in two weeks, bad networking information spreading from China has disrupted the Internet.

On Thursday morning, bad routing data from a small Chinese ISP called IDC China Telecommunication was re-transmitted by China's state-owned China Telecommunications, and then spread around the Internet, affecting Internet service providers such as AT&T, Level3, Deutsche Telekom, Qwest Communications and Telefonica.

MORE LIKE THIS

China's Great Firewall spreads overseas

China telecom operator denies hijacking Internet traffic

Research experiment disrupts Internet, for some

[on IDG Answers](#) →
What is a BGP hijack?

<http://www.computerworld.com/article/2516953/enterprise-applications/a-chinese-isp-momentarily-hijacks-the-internet--again-.html>



Internet-Wide Catastrophe—Last Year



One year ago today TTNNet in Turkey (AS9121) pretended to be the entire Internet. And unfortunately for the rest of the Internet, many large network providers believed them (or at least believed them in part). As far as anyone knows, it was a mistake, not a malicious act. But the consequences were far from benign: for several hours a large number of Internet users were unable to reach a large number of Internet sites. Twelve months later we can take a look at what happened, and whether we've learned much in the intervening time.

Early Christmas Eve morning 2004, TTNNet (AS9121) started announcing what appeared to be a full table (well over 100,000 entries) of Internet routes to all of their transit providers. I was on call that Christmas (as I am this Christmas; I'm sensing a bad pattern here). So around 4:30 in the morning US Eastern Standard Time, I started getting paged.



DDoS Attack Hits 400 Gbit/s, Breaks Record

A distributed denial-of-service NTP reflection attack was reportedly 33% bigger than last year's attack against Spamhaus.



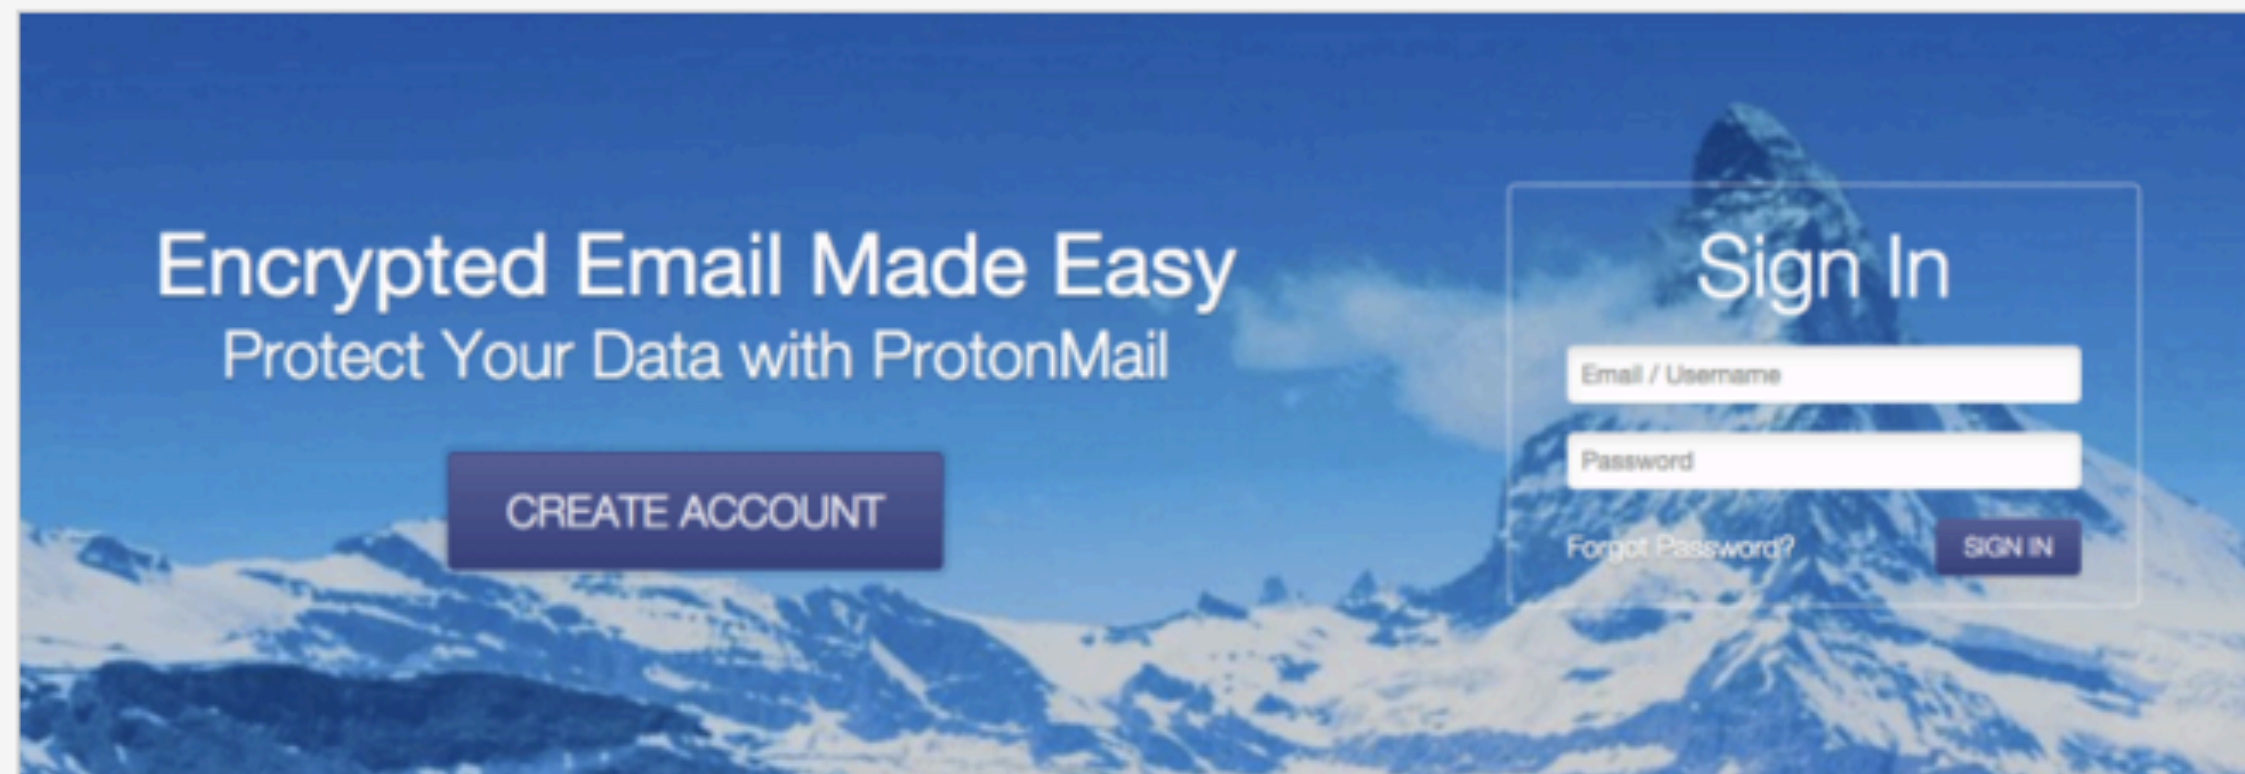
ProtonMail On Battling A Sustained DDoS Attack

Posted 23 hours ago by [Natasha Lomas \(@riptari\)](#)

897
SHARES



Next Story



Encrypted webmail provider, ProtonMail, has been fighting a wave of DDoS attacks since November 3 that, by last Friday, had taken its service offline for more than 24 hours. At the time of writing the attacks are still coming.

They have included what ProtonMail co-founder Andy Yen [described](#) as a “co-ordinated assault” on its ISP that exceeded 100Gbps and attacked not only the Swiss datacenter but routers in various locations where the ISP has nodes — taking multiple services offline, not just ProtonMail’s email.

CrunchBase

ProtonMail

FOUNDED
2013

OVERVIEW

End-to-end encrypted email, based in Switzerland. ProtonMail is a new service that provides easy to use secure email. ProtonMail's secure email system is designed around the principle of zero access. This means user data cannot be read by ProtonMail and turned over to third parties because ProtonMail servers do not store user encryption keys. The service is backwards compatible with insecure email ...

LOCATION

Geneva, 07

CATEGORIES

Messaging, Email, Data Security, Security

FOUNDERS



Goals of this lecture

- Understand attacks on availability *in the network*.
- Many attacks at the application layer — bugs in code — go take 18-487 to learn more about those.
- This class focuses on attacks on availability in the network.



Two classes of attacks on availability today

- **Resource Exhaustion**

- DDoS
- SYN Floods

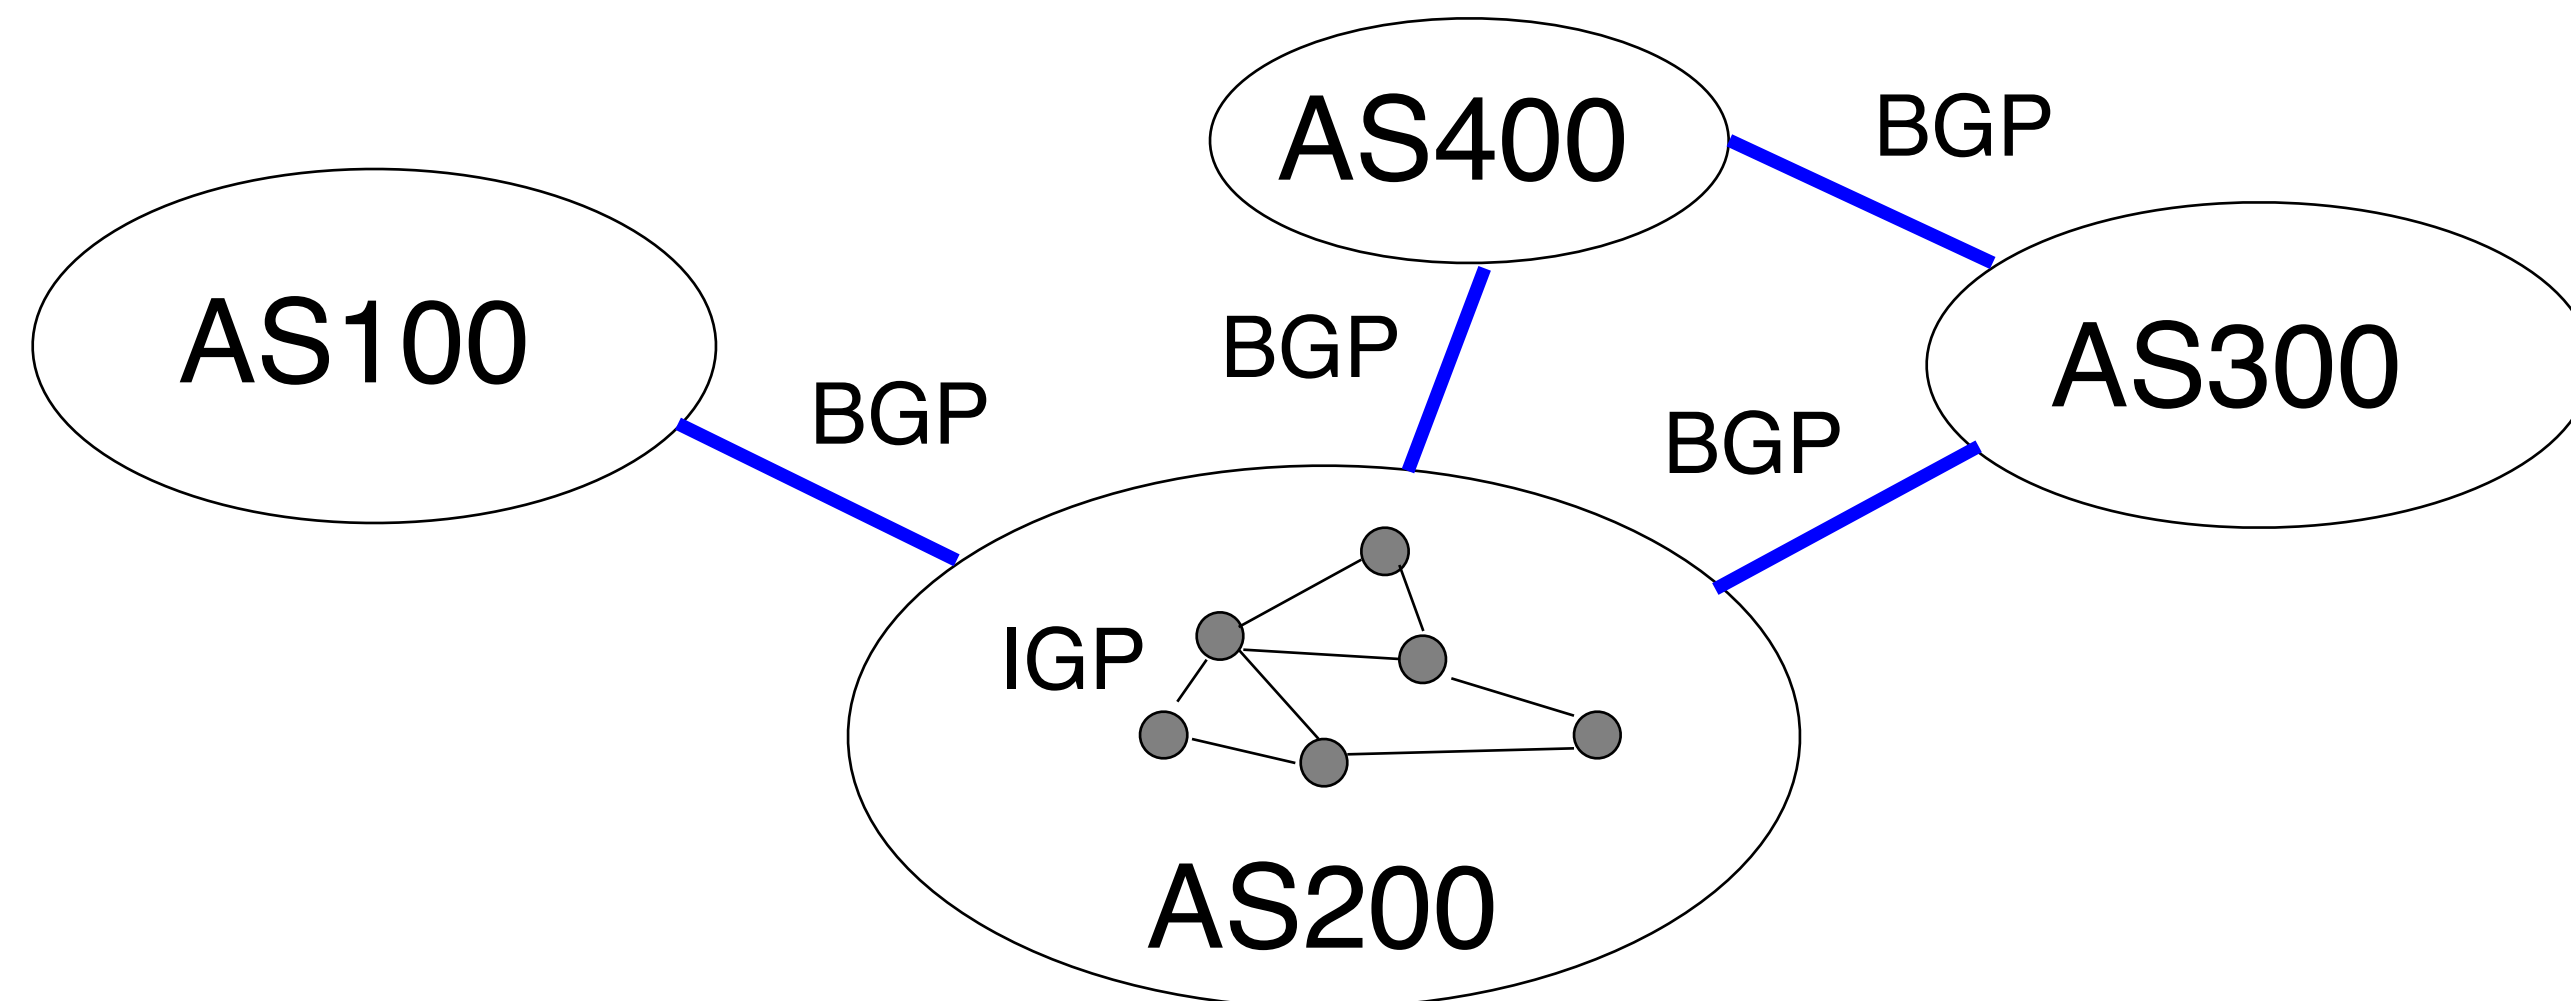
- **Routing Attacks**

- We'll talk about flaws in BGP
- There are so many kinds of attacks we're not discussing though!
- Take 18-487 with Prof. Sekar!



Recall: Internet routing

- An Interior Gateway Protocol (IGP) is used to route packets within an AS: Intra-domain routing
- An Exterior Gateway Protocol (EGP) to maintain Internet connectivity among ASs: Inter-domain routing



What kind of routing algorithm is
BGP?



What are the other kinds of routing algorithms we discussed in this class (not BGP)?



How does BGP work?

Internet routers communicate using the Border Gateway Protocol (BGP):


- Destinations are **prefixes** (CIDR blocks)
 - Example: 128.2.0.0/16 (CMU)
- Routes through **Autonomous Systems** (ISPs)
- Each ISP is uniquely identified by a number
 - Example: 9 (Carnegie Mellon)
- Each route includes a list of traversed ISPs:
 - Example: 9 ← 5050 ← 11537 ← 2153

**Carnegie
Mellon**




Principles of operation

- Exchange routes
 - AS100 announces 128.1.1.0/24 prefix to AS200 and AS300, etc
- Incremental updates



UPDATE message example

Prefix: 128.1.1.0/24
Nexthop: 192.208.10.1
ASPath: 100



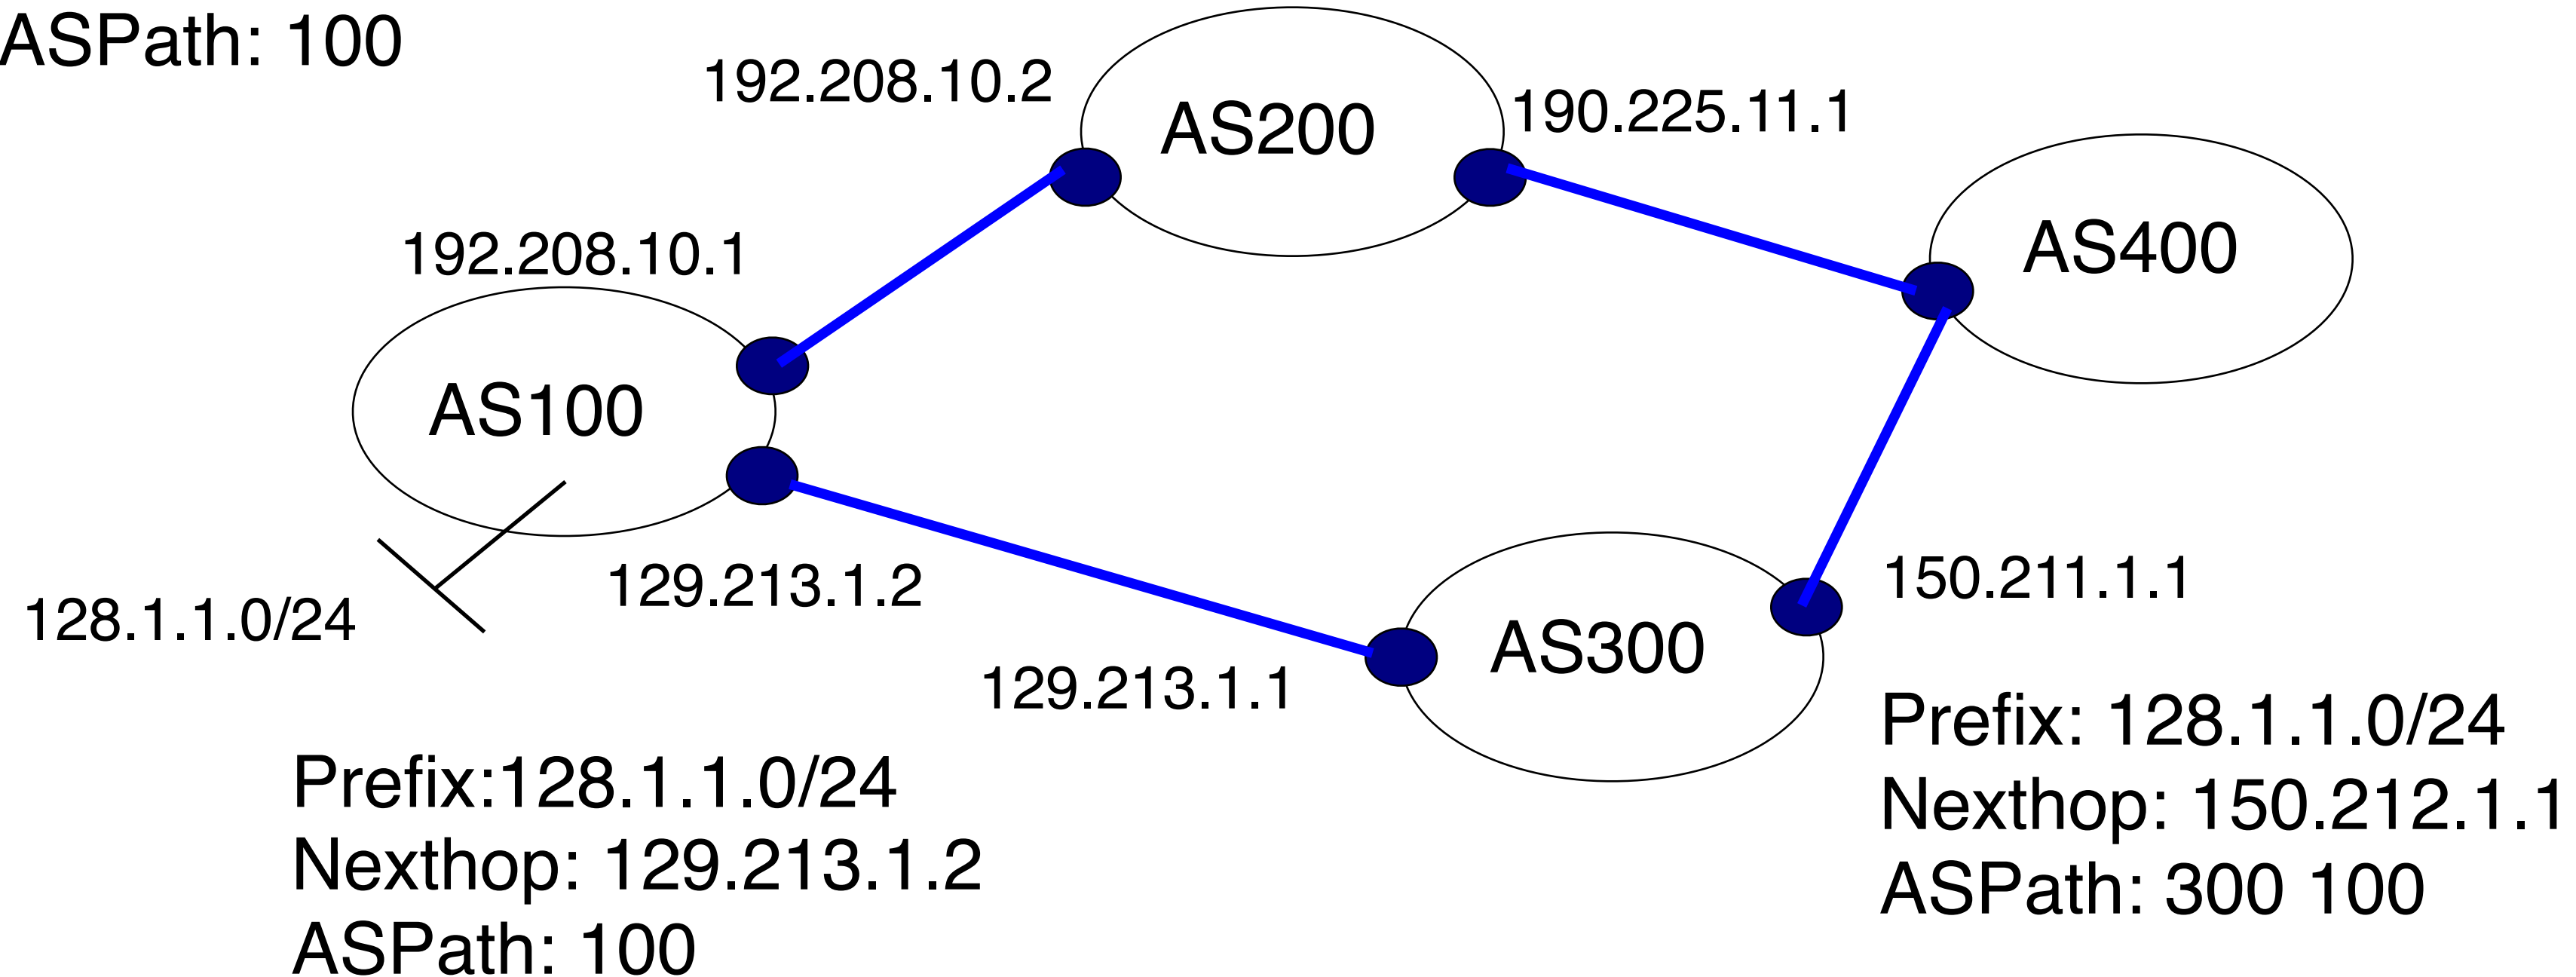
Prefix: 128.1.1.0/24
Nexthop: 129.213.1.2
ASPath: 100



Route propagation

Prefix: 128.1.1.0/24
Nexthop: 192.208.10.1
ASPath: 100

Prefix: 128.1.1.0/24
Nexthop: 190.225.11.1
ASPath: 200 100



All you need is one
compromised BGP speaker



Pakistan Telecom: Sub-prefix hijack



Corrigendum- Most Urgent

**GOVERNMENT OF PAKISTAN
PAKISTAN TELECOMMUNICATION AUTHORITY
ZONAL OFFICE PESHAWAR**

Plot-11, Sector A-3, Phase-V, Hayatabad, Peshawar.
Ph: 091-9217279- 5829177 Fax: 091-9217254
www.pta.gov.pk

NWFP-33-16 (BW)/06/PTA

February ,2008

Subject: Blocking of Offensive Website

Reference: *This office letter of even number dated 22.02.2008.*

I am directed to request all ISPs to immediately block access to the following website

URL: <http://www.youtube.com/watch?v=o3s8jtvvg00>

IPs: 208.65.153.238, 208.65.153.253, 208.65.153.251

Compliance report should reach this office through return fax or at email

peshawar@pta.gov.pk today please.

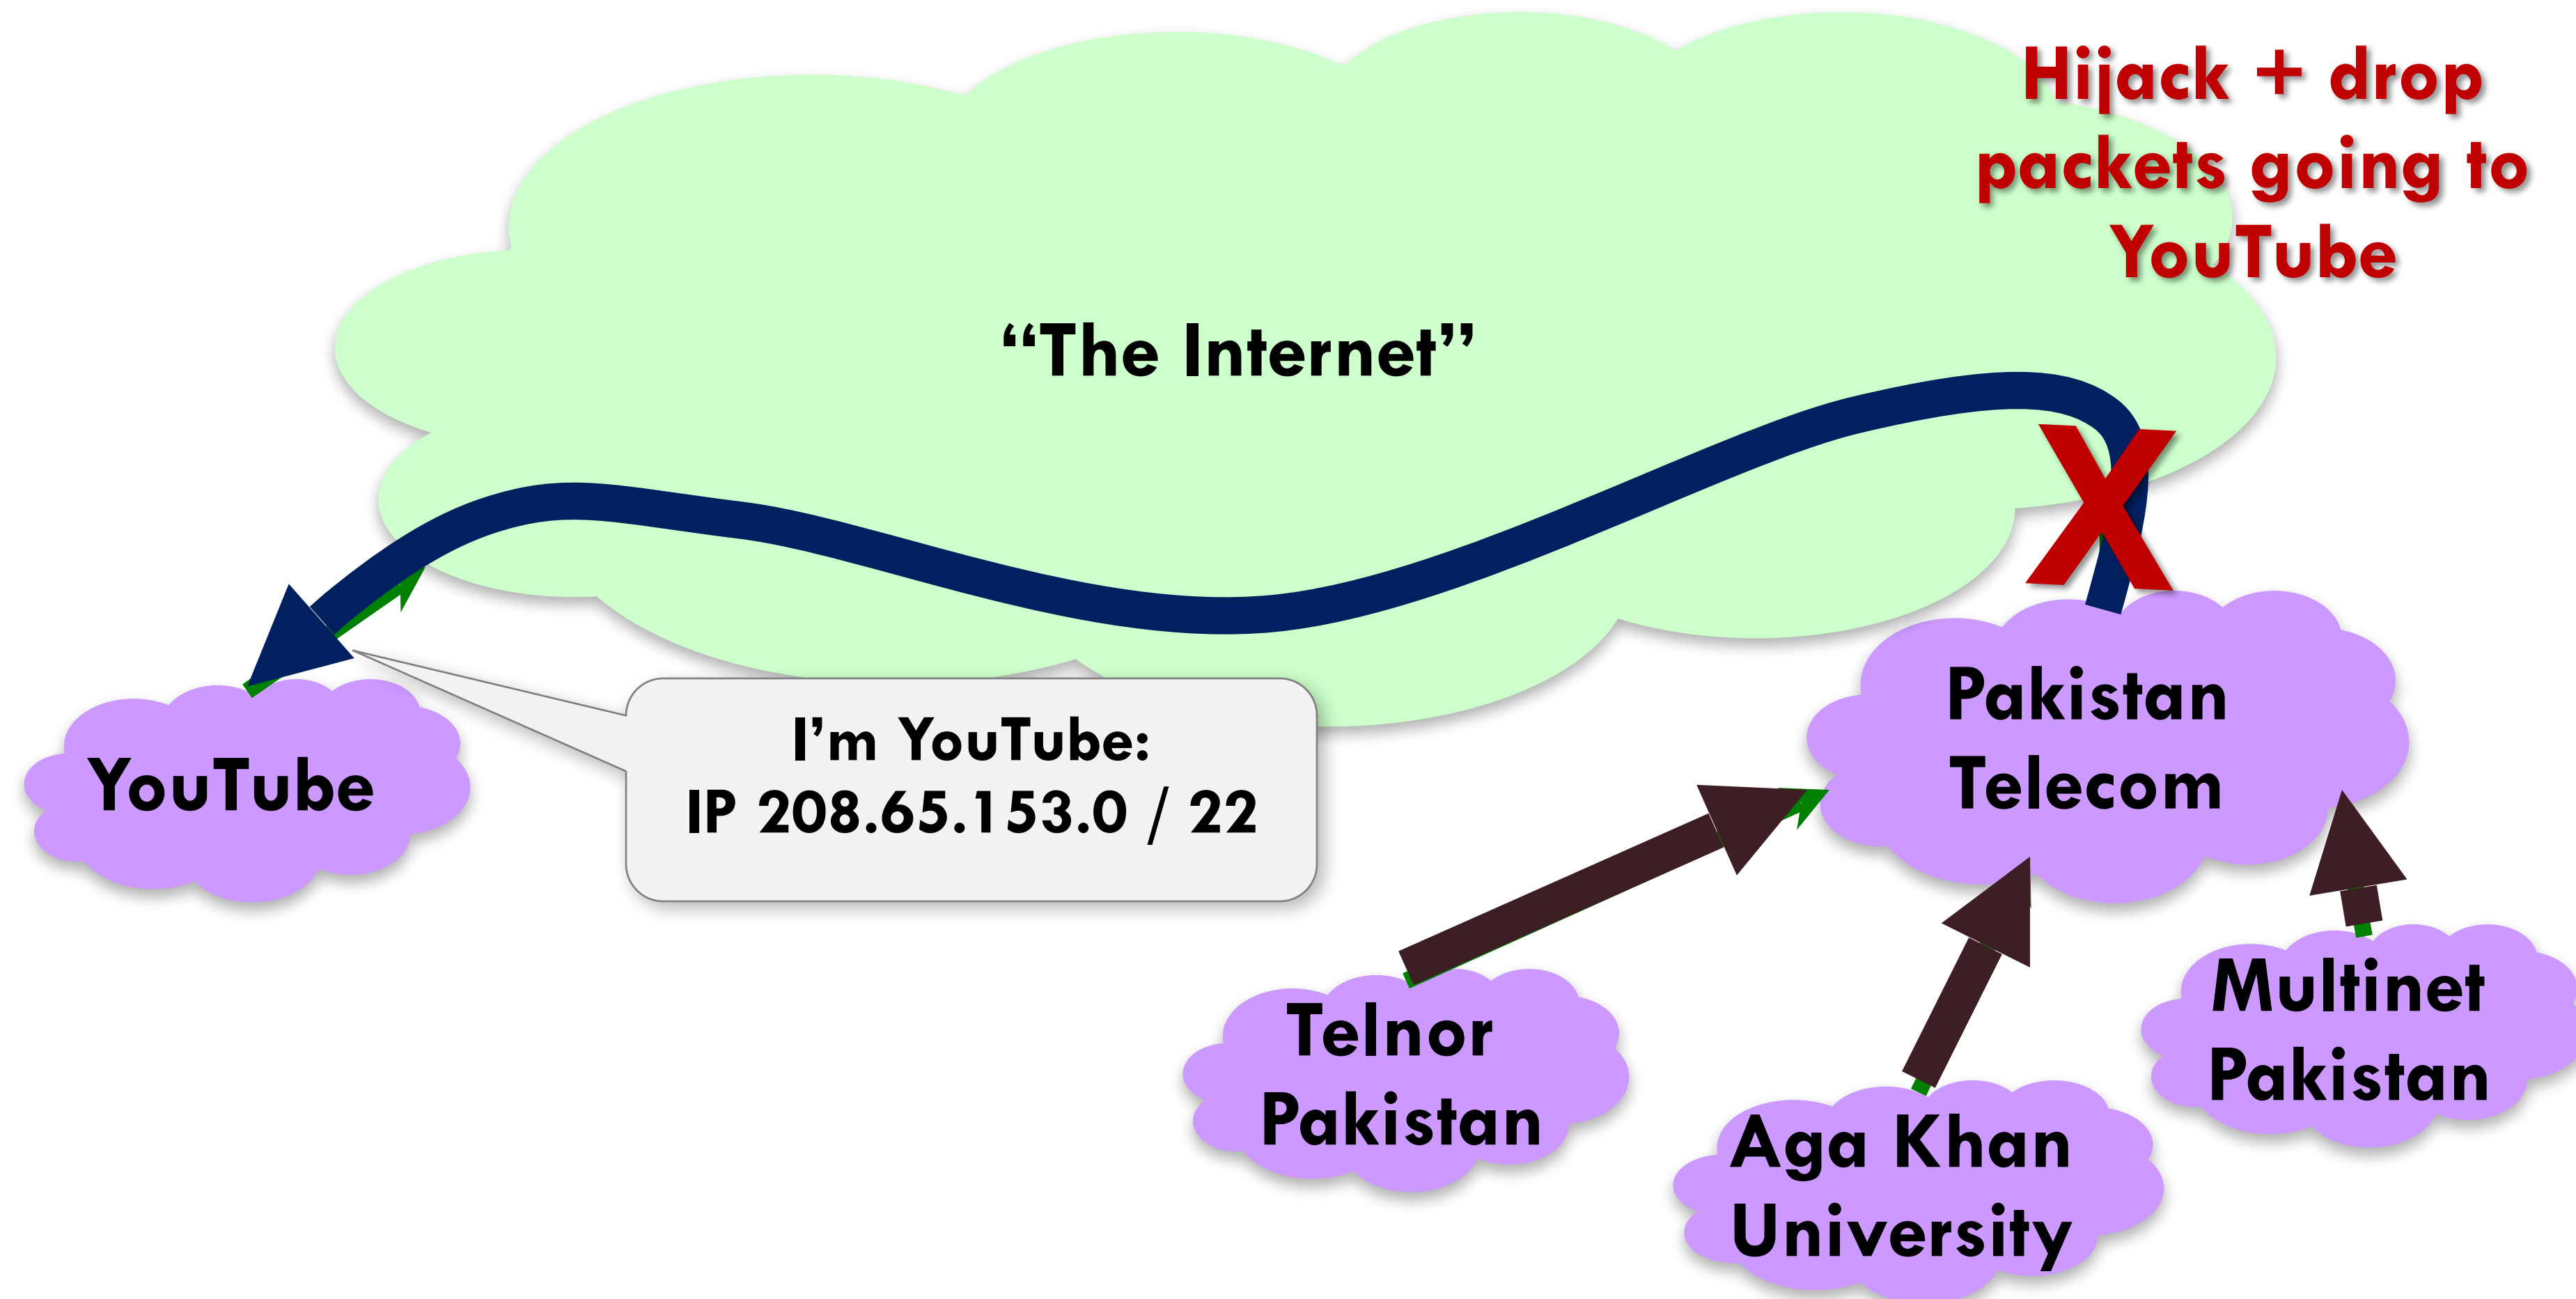
YouTube

Multinet
Pakistan



Pakistan wanted to send an iBGP announcement to blackhole traffic to YouTube...

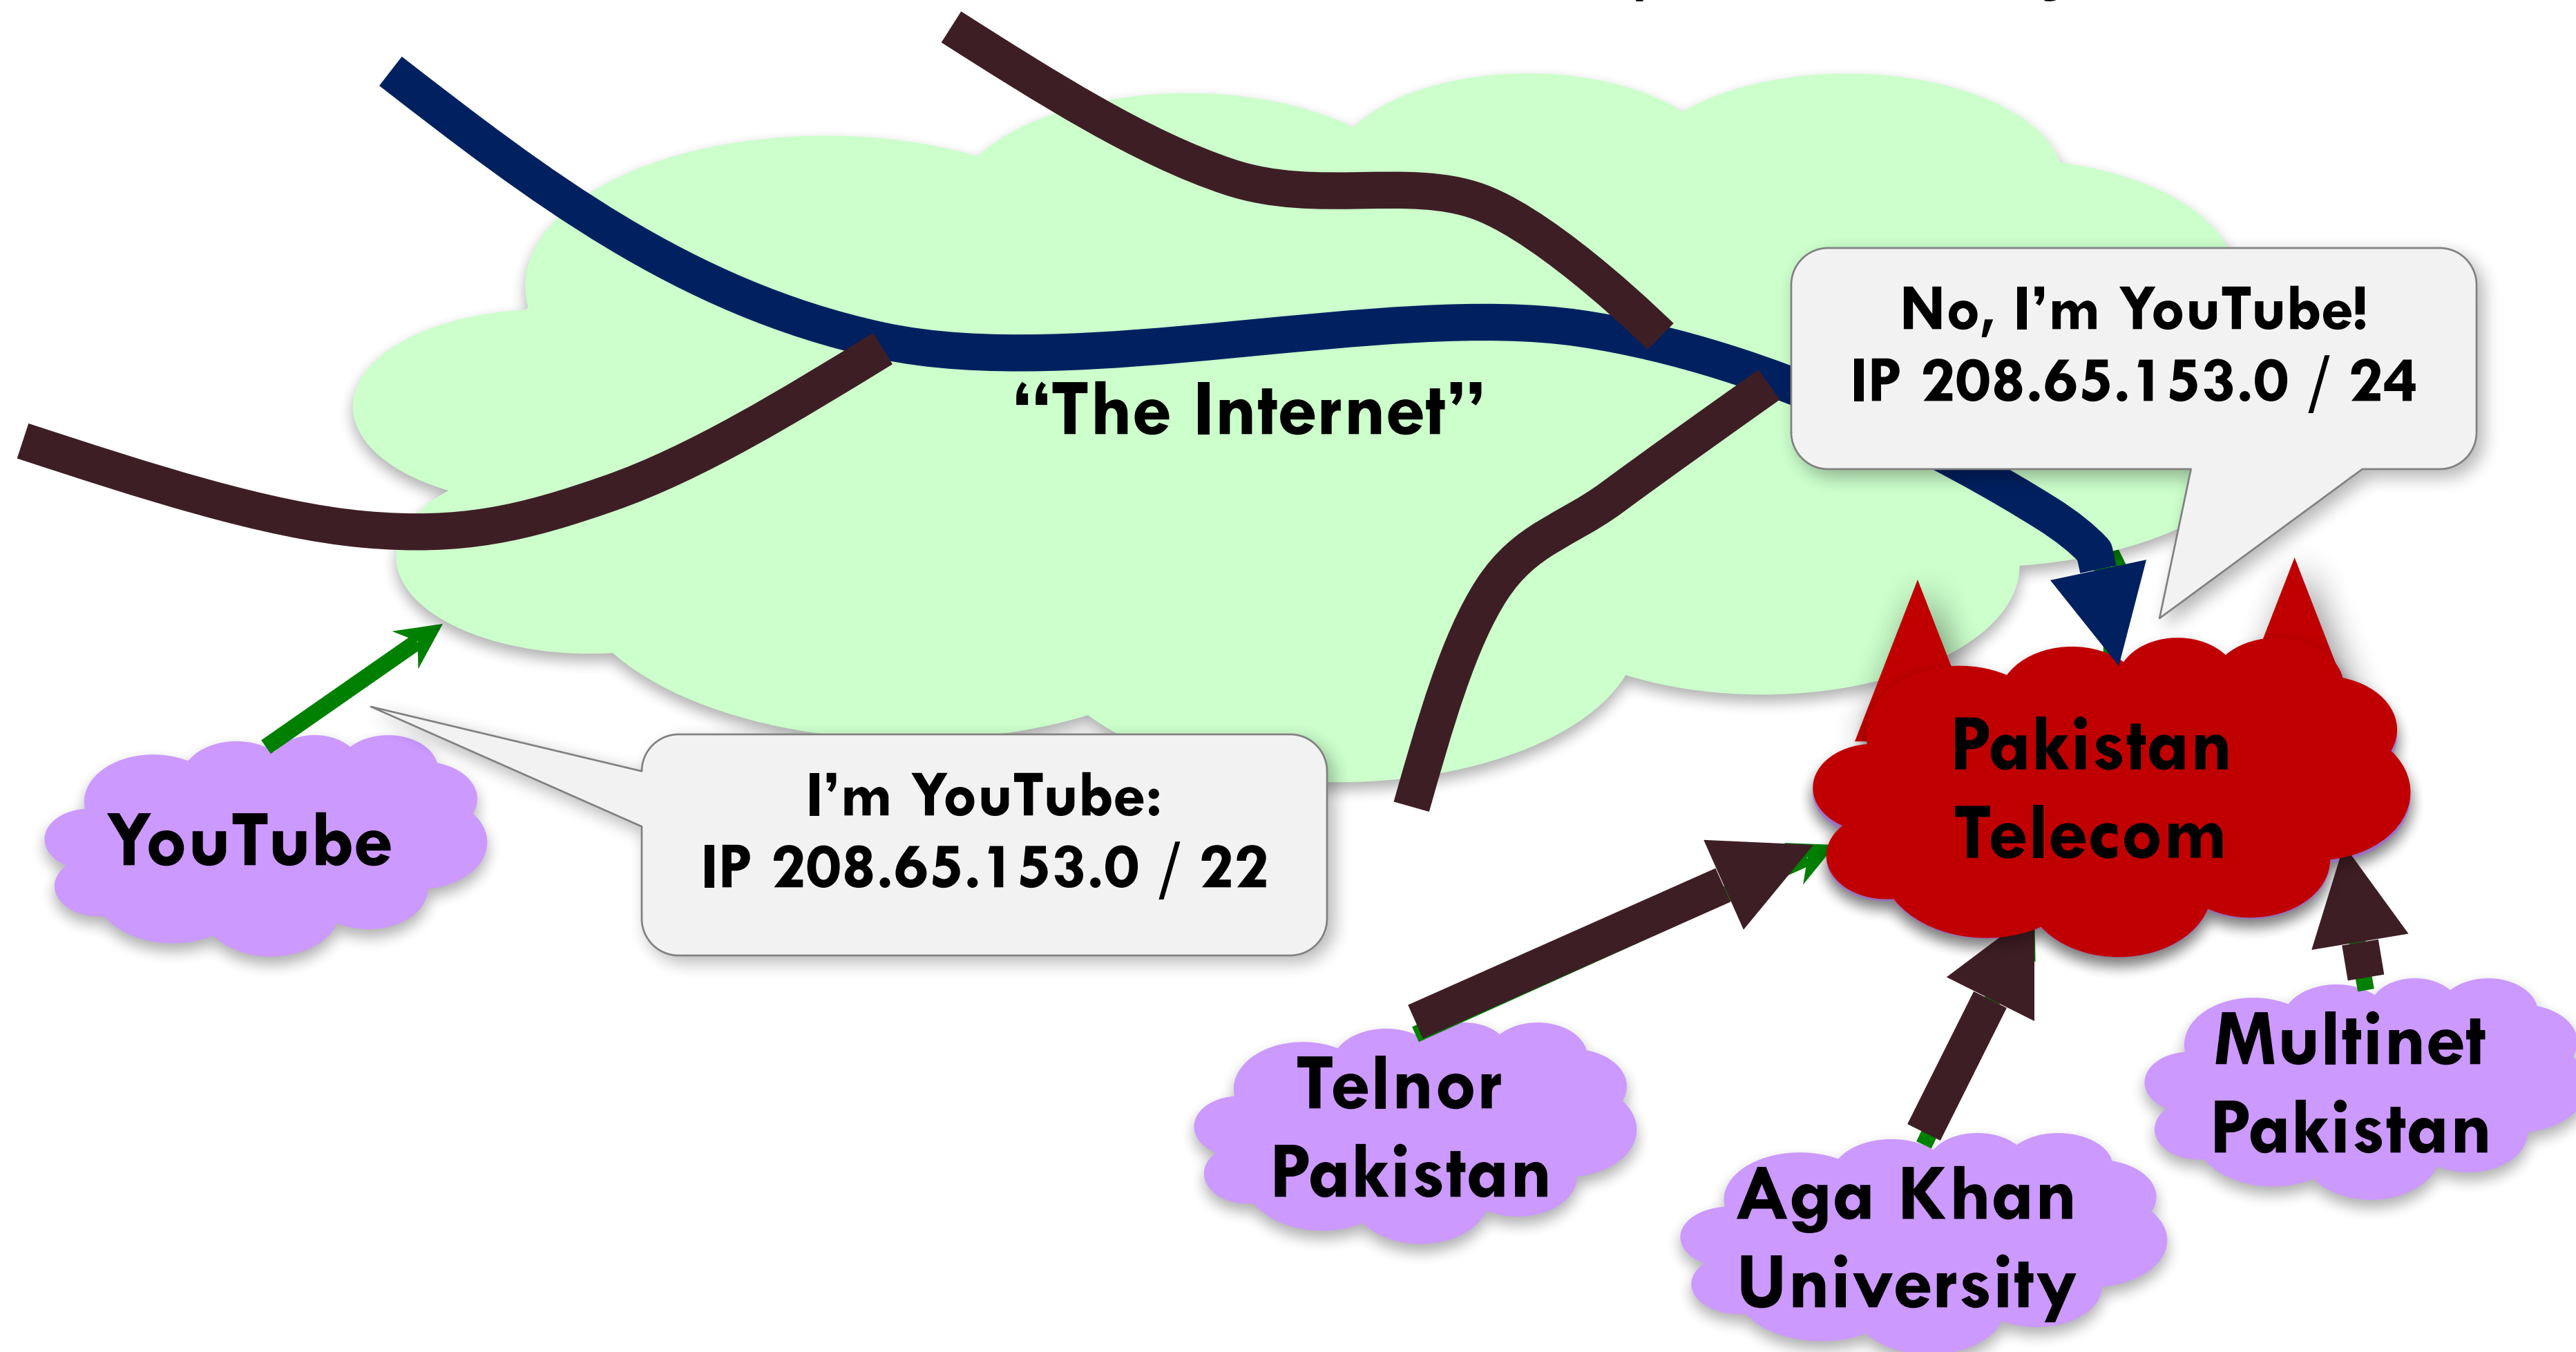
Pakistan Telecom: Sub-prefix hijack



Block your own customers.

But they accidentally sent an eBGP announcement to blackhole YouTube!

Pakistan Telecom: Sub-prefix hijack

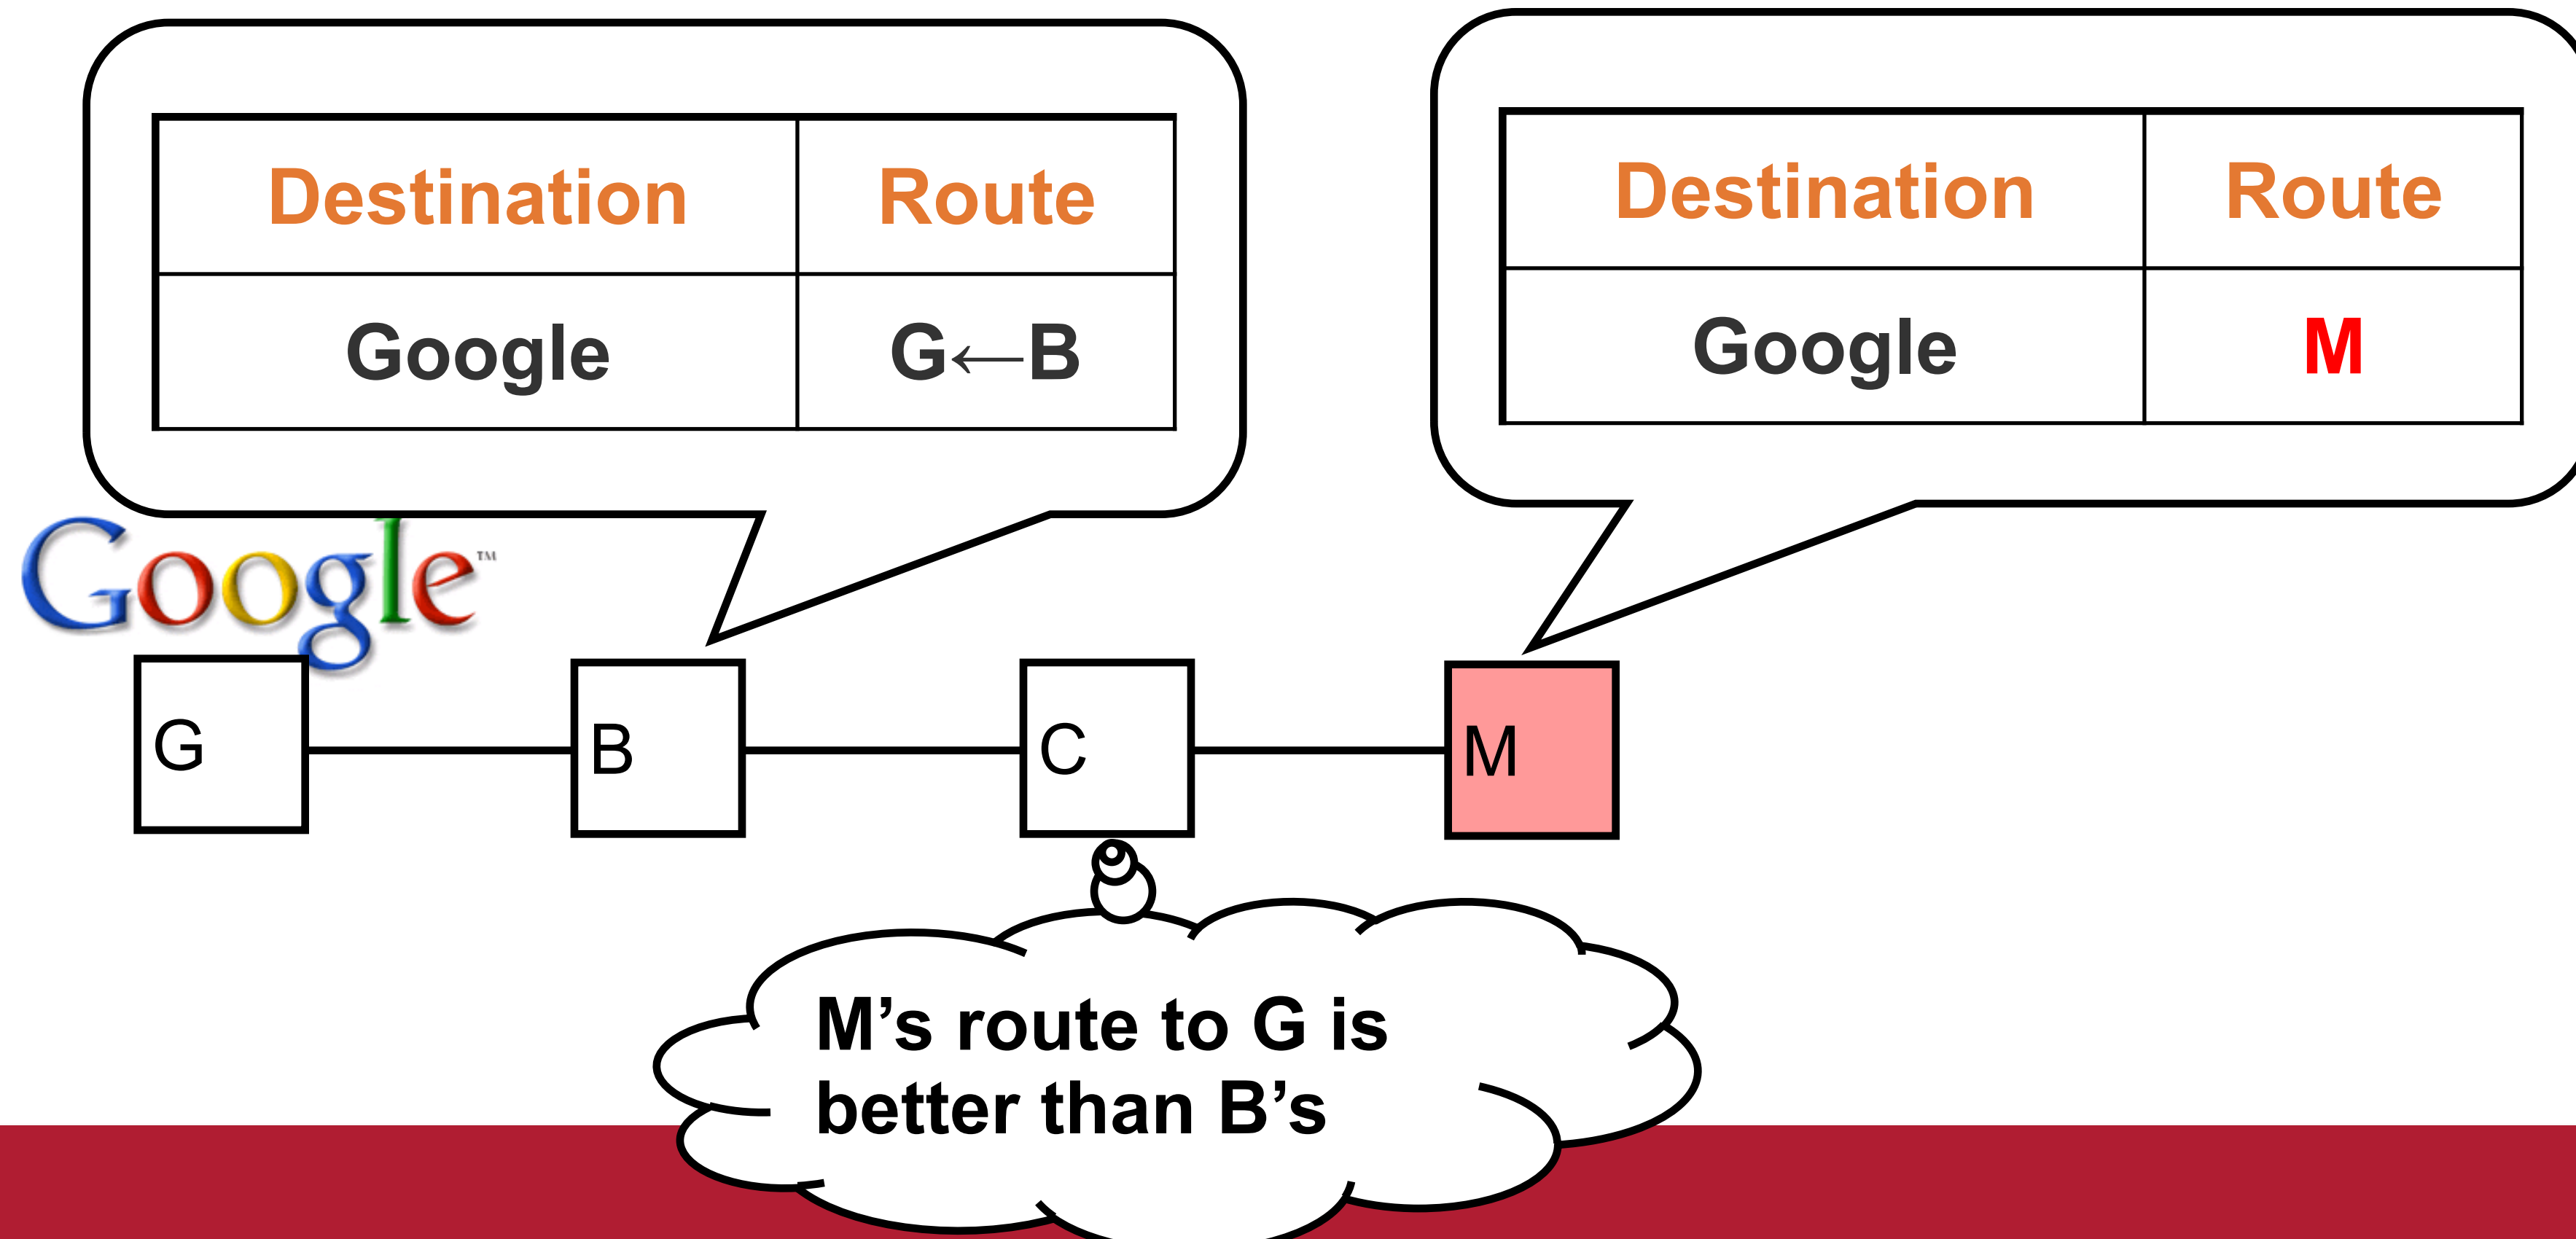


Potential attack objectives

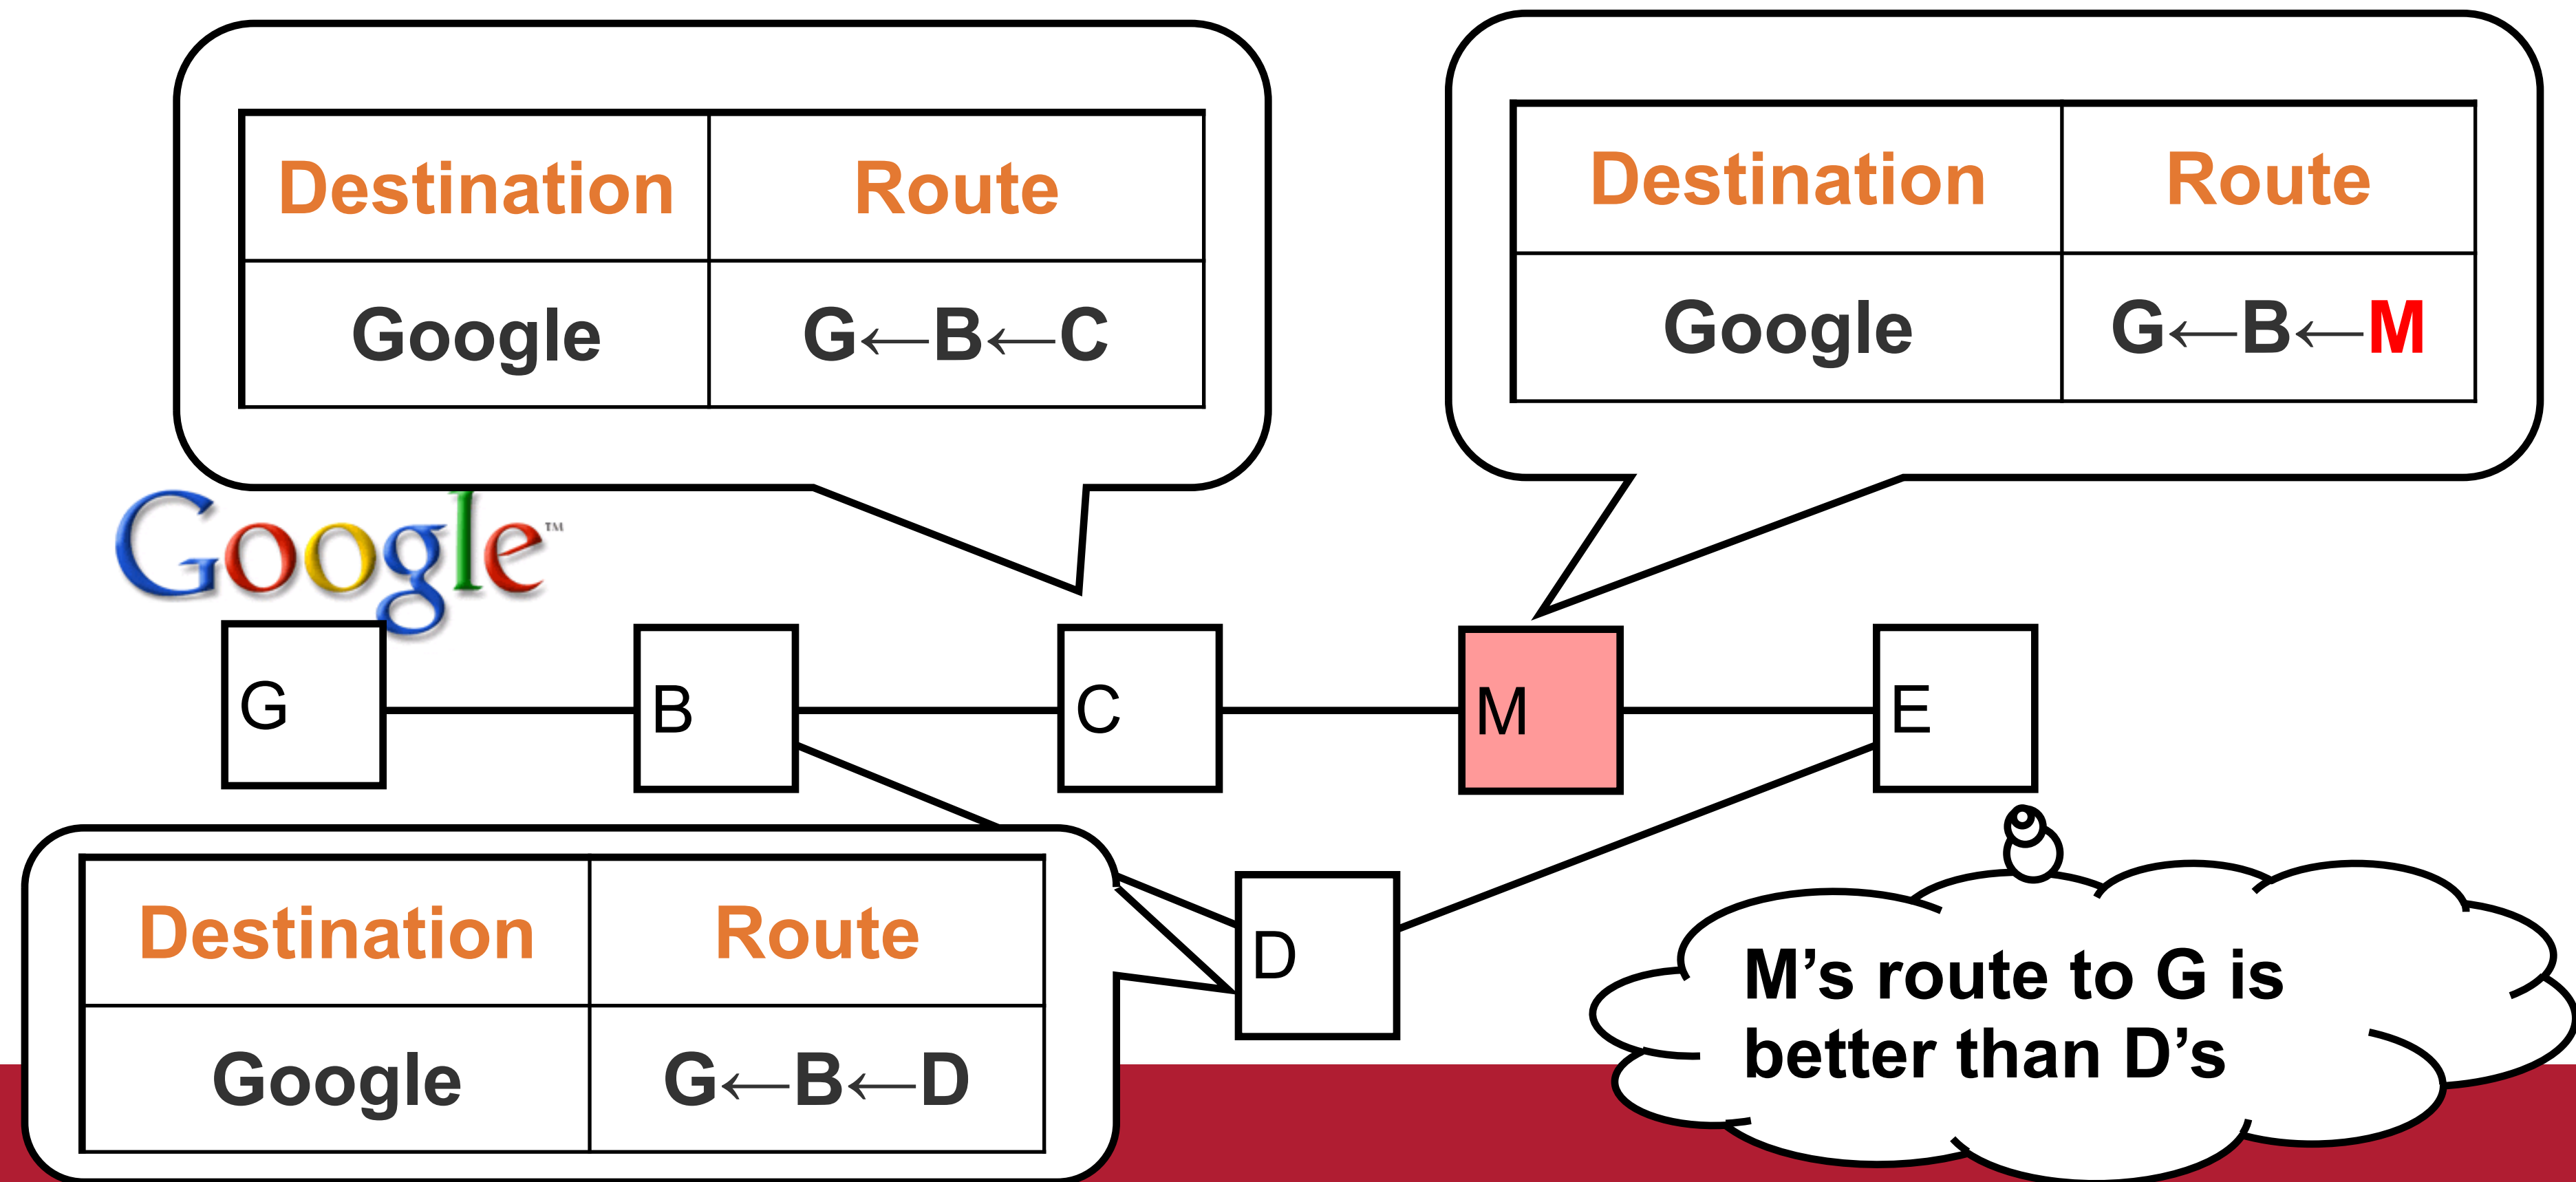
- Blackholing – make something unreachable
- Redirection – e.g., congestion, eavesdropping
- Instability
- But more often than not, just a mistake!



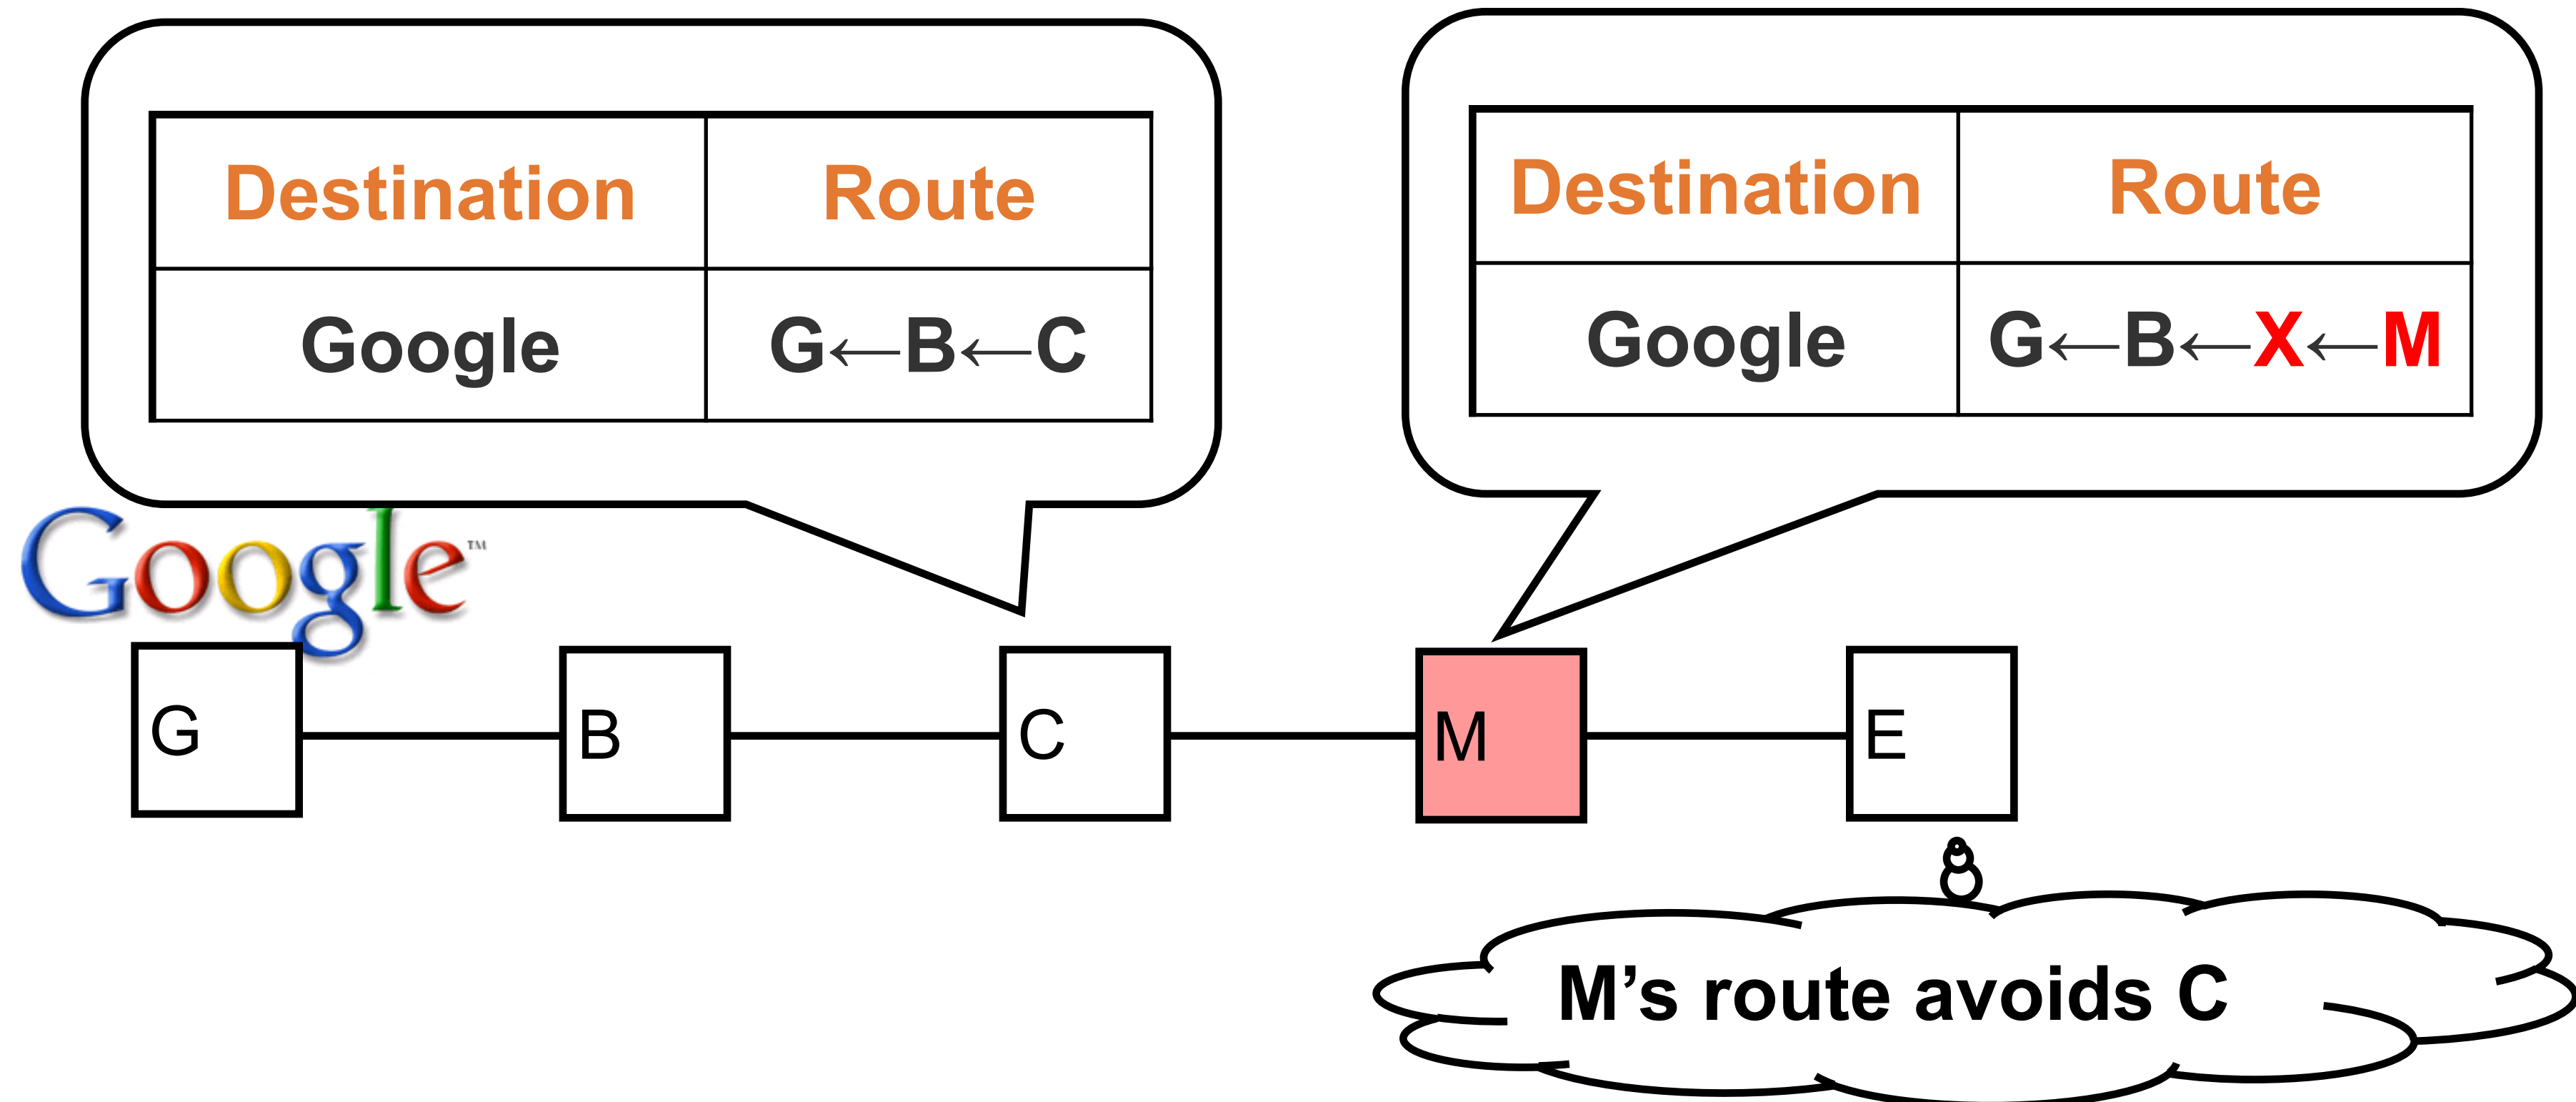
Unauthorized origin ISP (prefix theft)



AS-path truncation



AS path alteration



How can we fix this problem?



What tools from the last two lectures might we use?



BGP Security Requirements

- Verification of address space “ownership”
- Authentication of Autonomous Systems (AS)
- Router authentication and authorization (relative to an AS)
- Route and address advertisement authorization
- Route withdrawal authorization
- Integrity and authenticity of all BGP traffic on the wire
- Timeliness of BGP traffic




Idea #1: RPKI & Origin Authentication

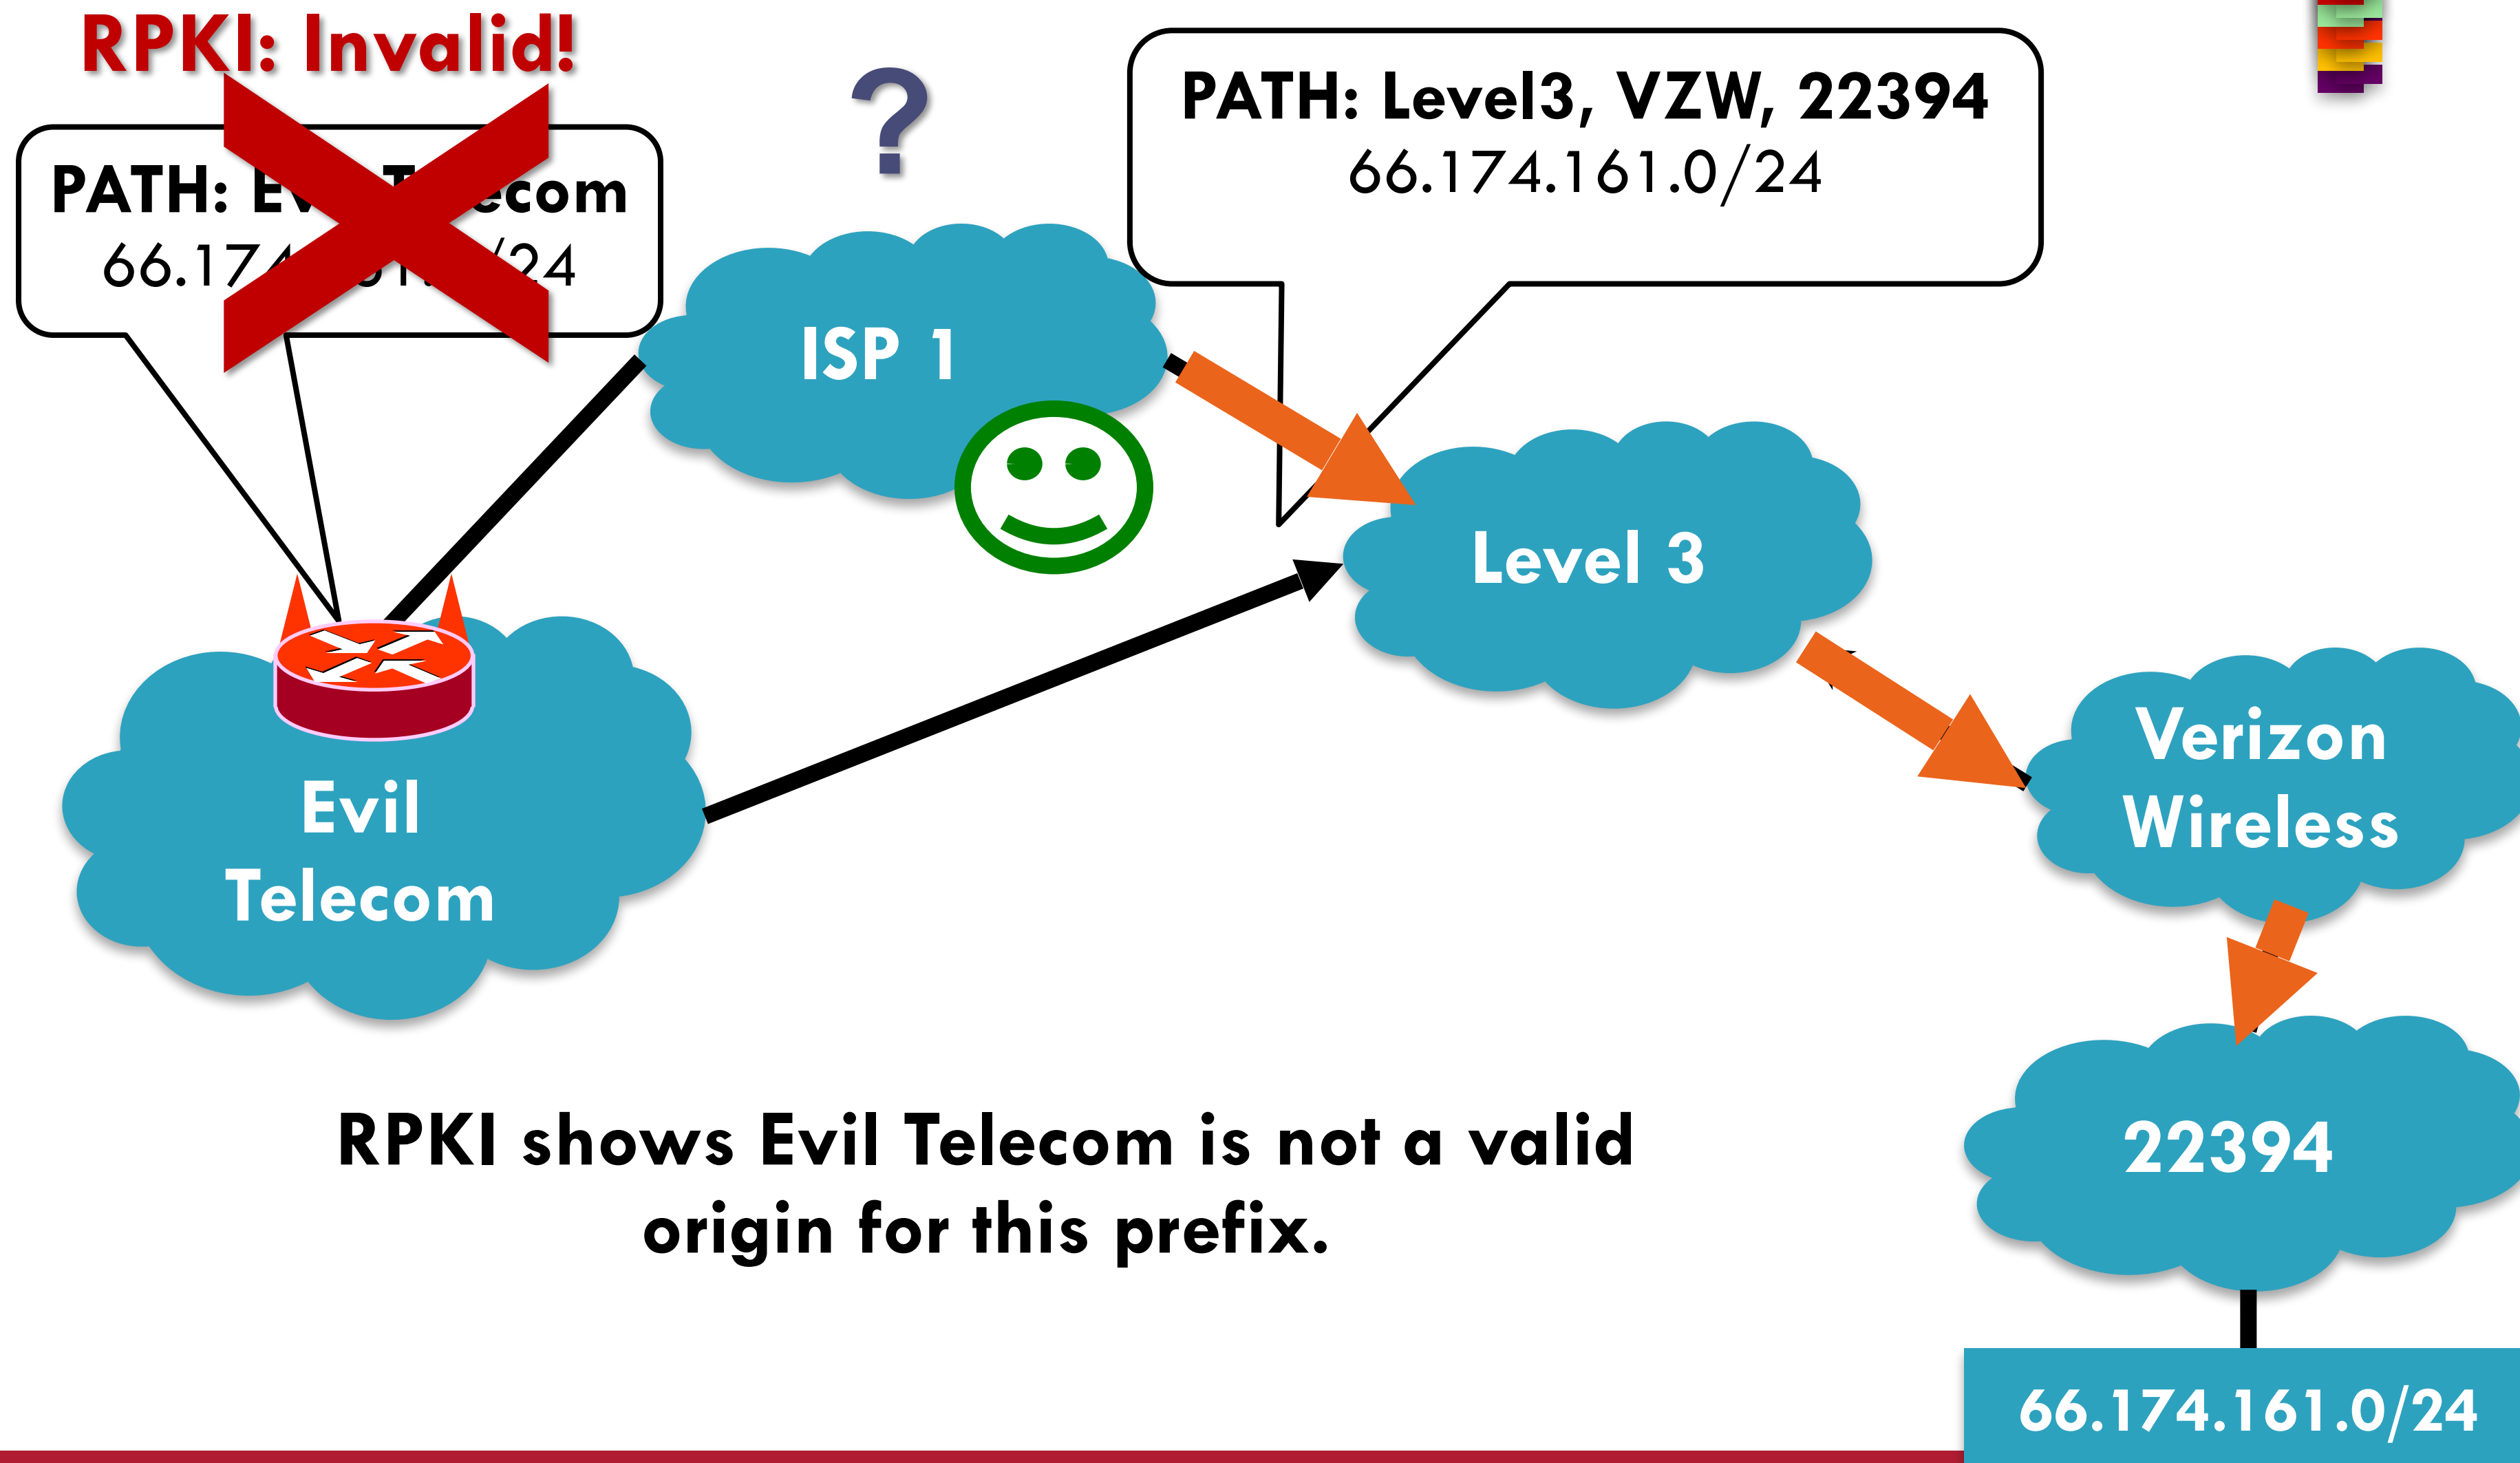

- Have all legitimate network operators *register* their prefixes along with a public key with a central authority.
 - Called: “RPKI” for *Routing Public Key Infrastructure*
- Whenever I announce my prefix, I sign my announcement.
- Anyone can verify that I am indeed allowed to originate this prefix.
-



Resource Public Key Infrastructure (RPKI): Certified mapping from ASes to public keys and IP prefixes.



Securing the Internet: RPKI

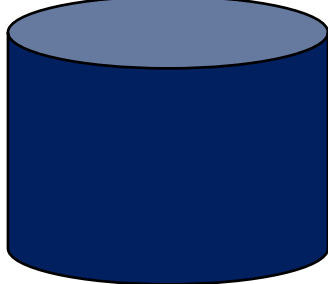


RPKI shows Evil Telecom is not a valid origin for this prefix.

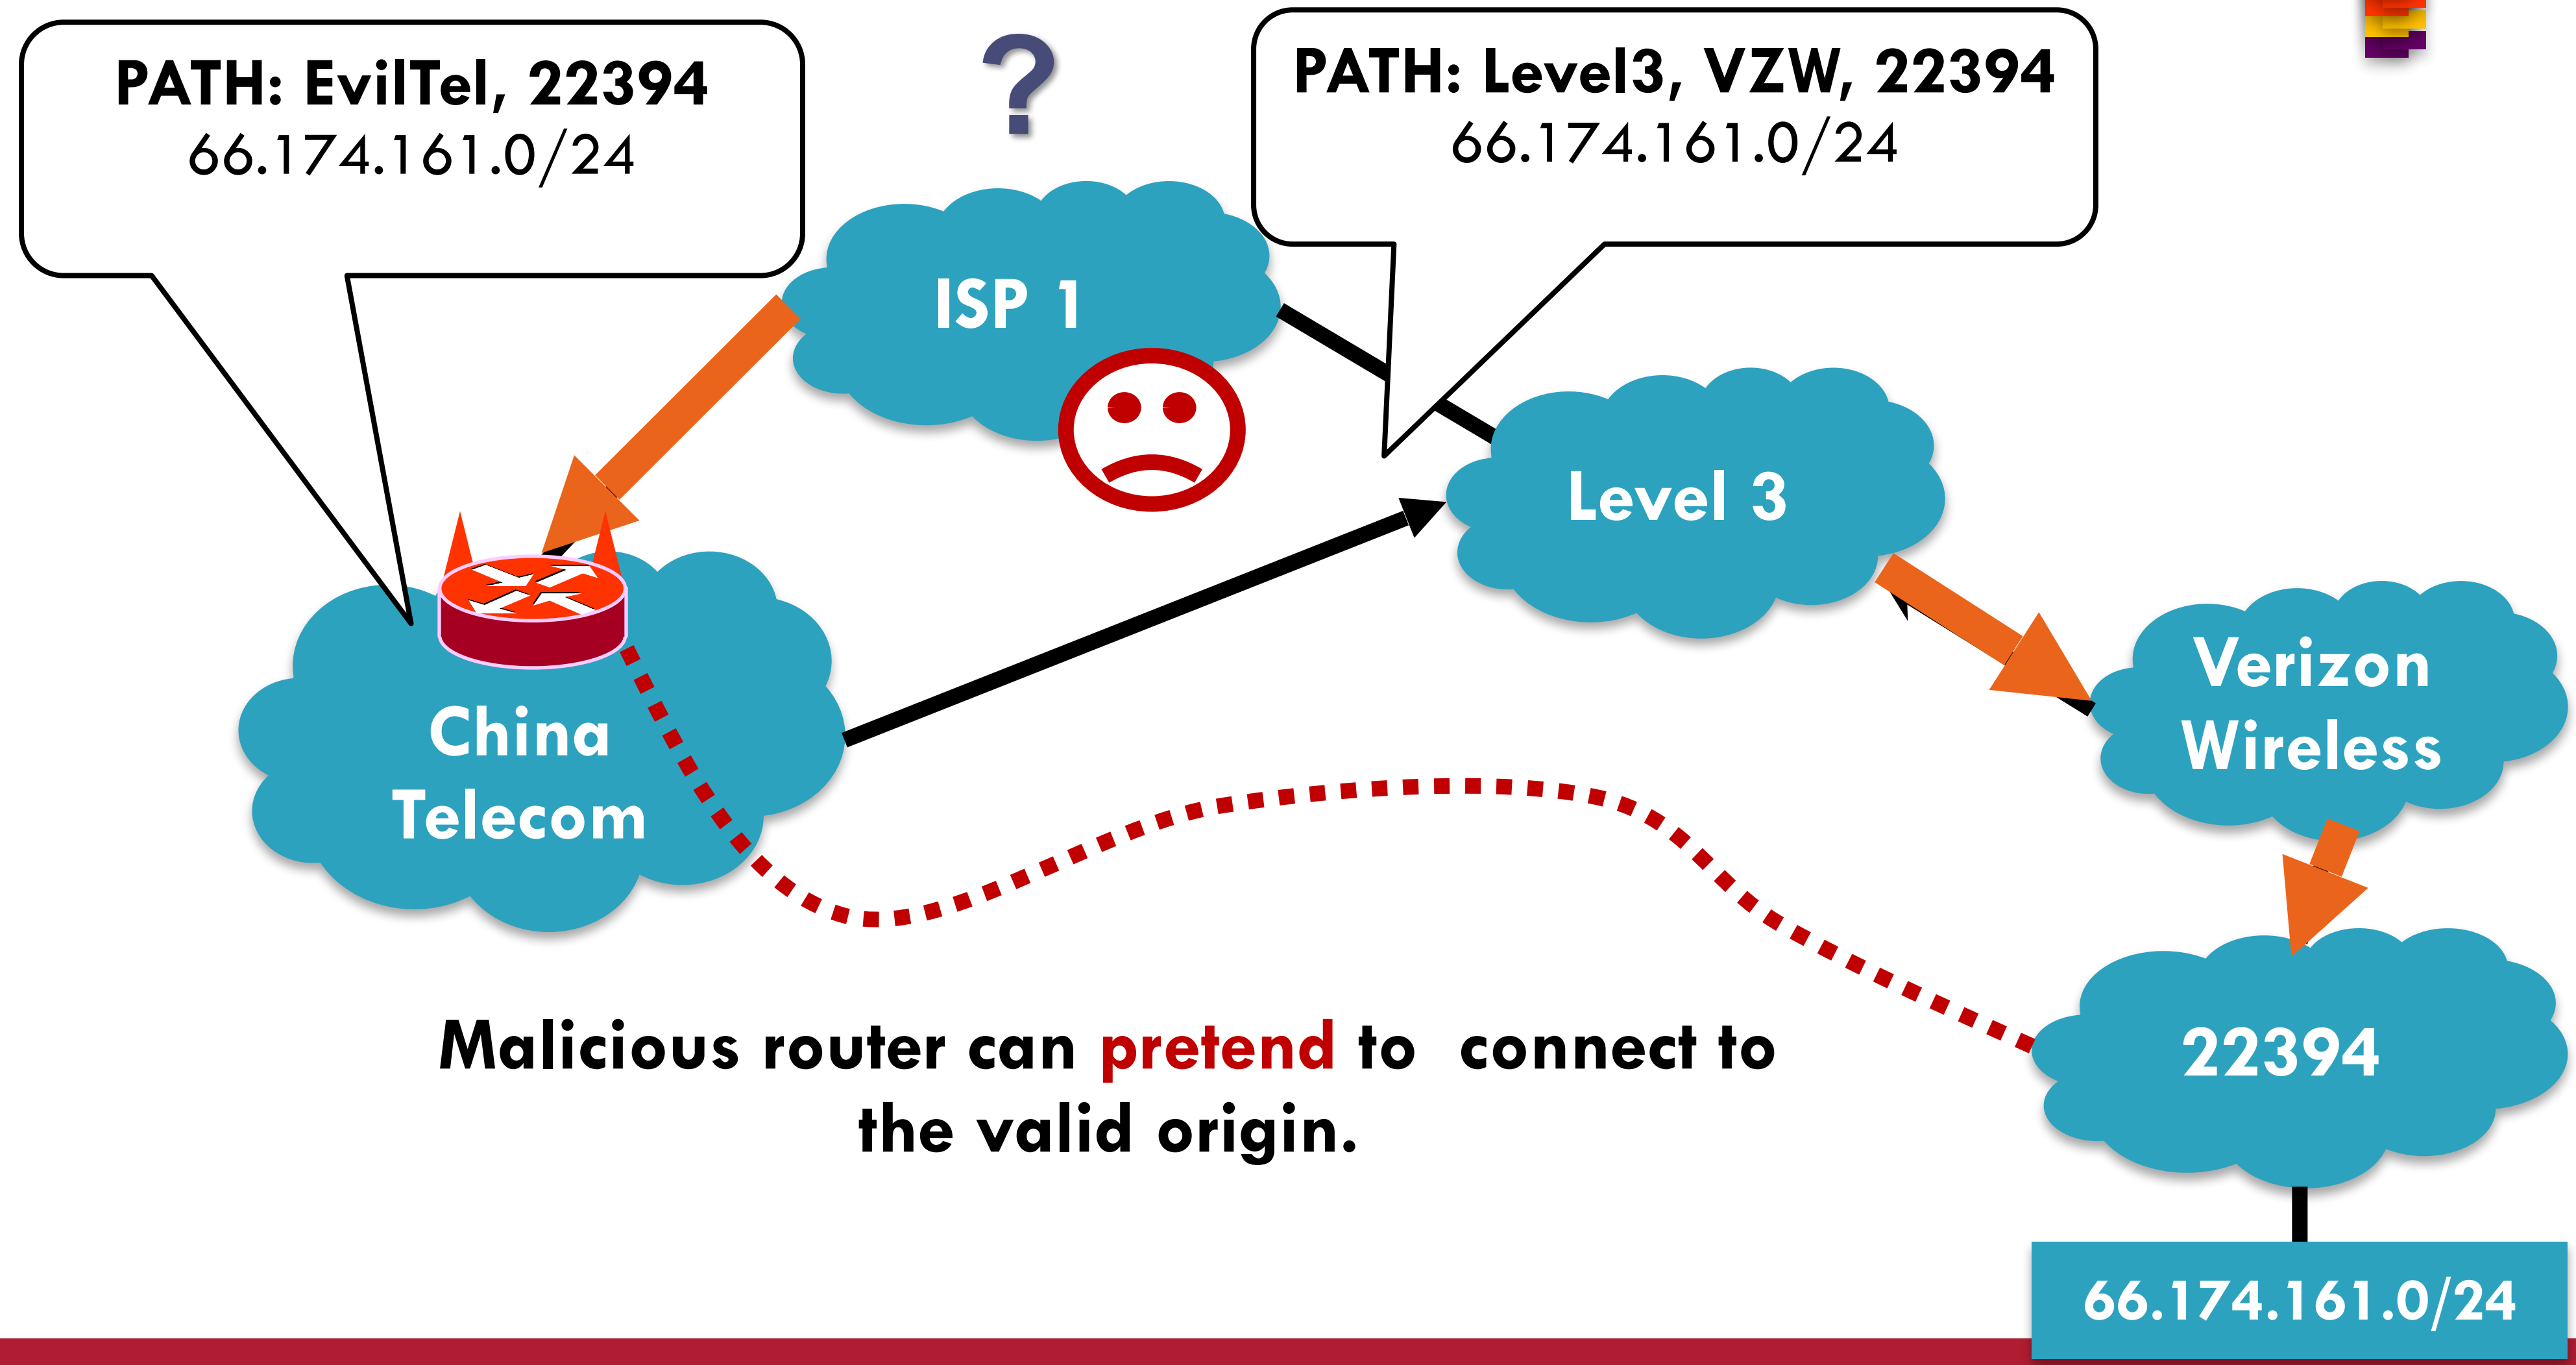


Why is this solution insufficient?



Resource Public Key Infrastructure (RPKI): Certified mapping from ASes to public keys and IP prefixes. 

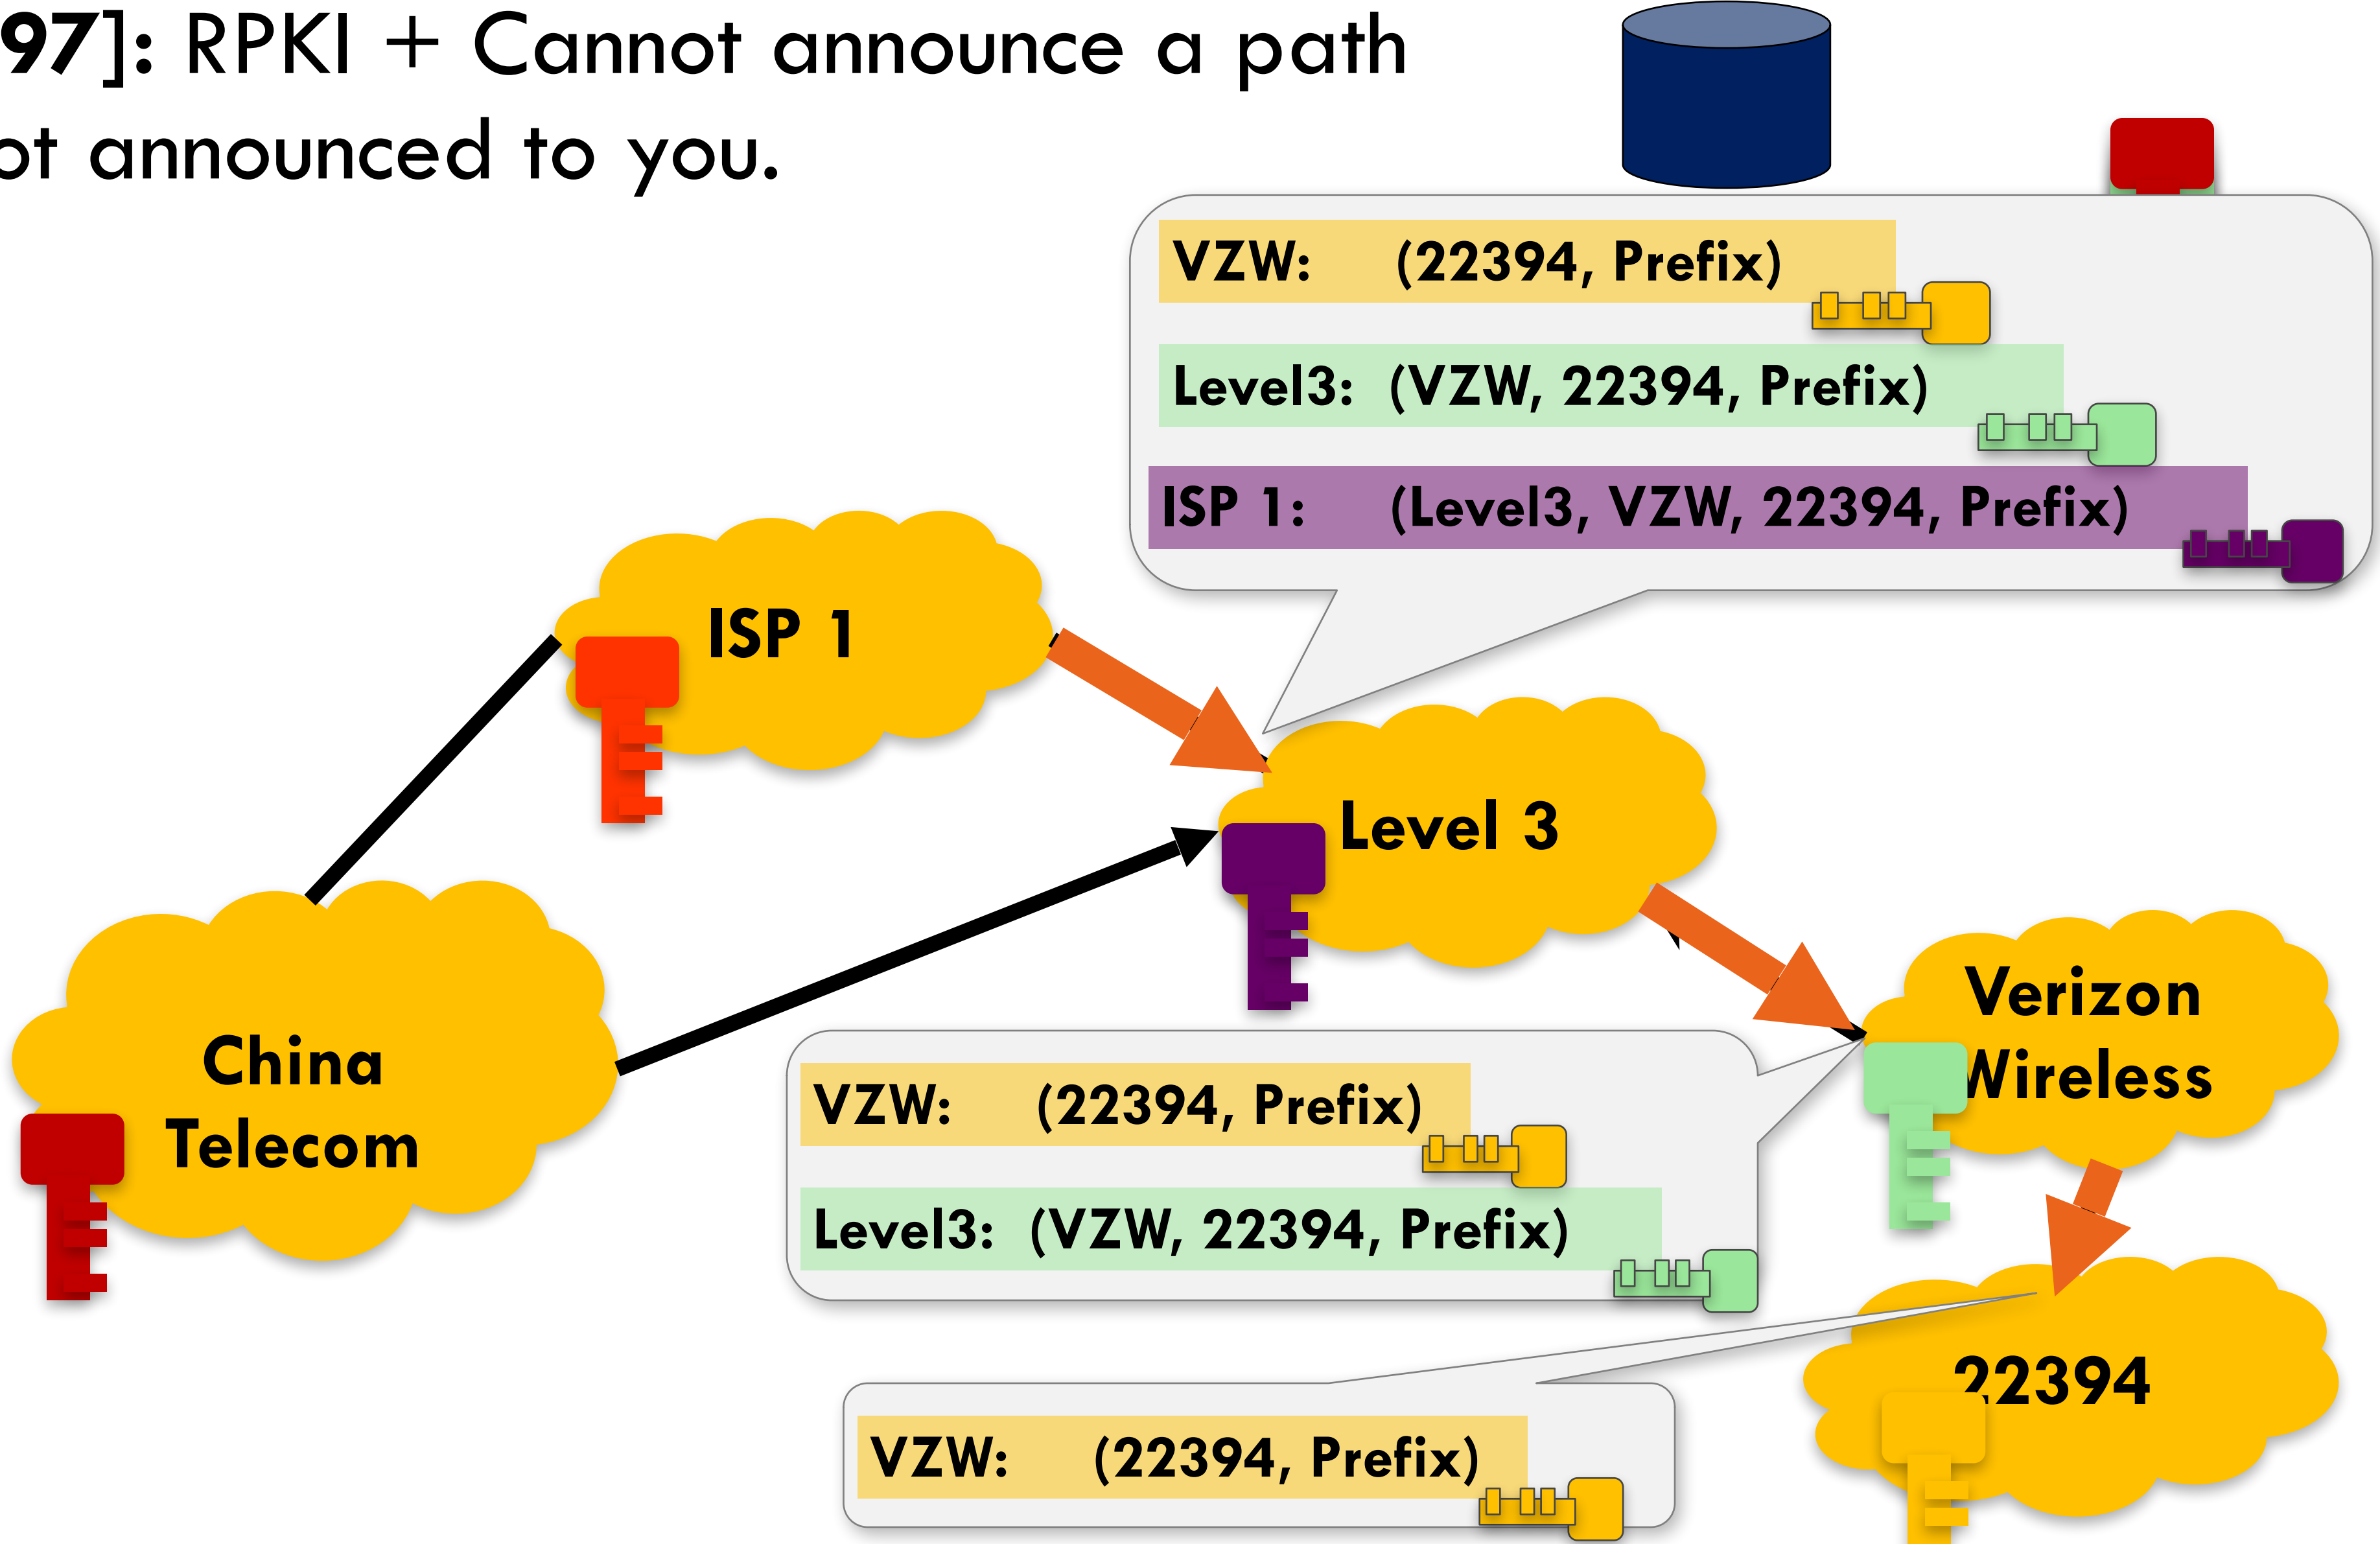
But **RPKI** alone is not enough!



Malicious router can **pretend** to connect to the valid origin.



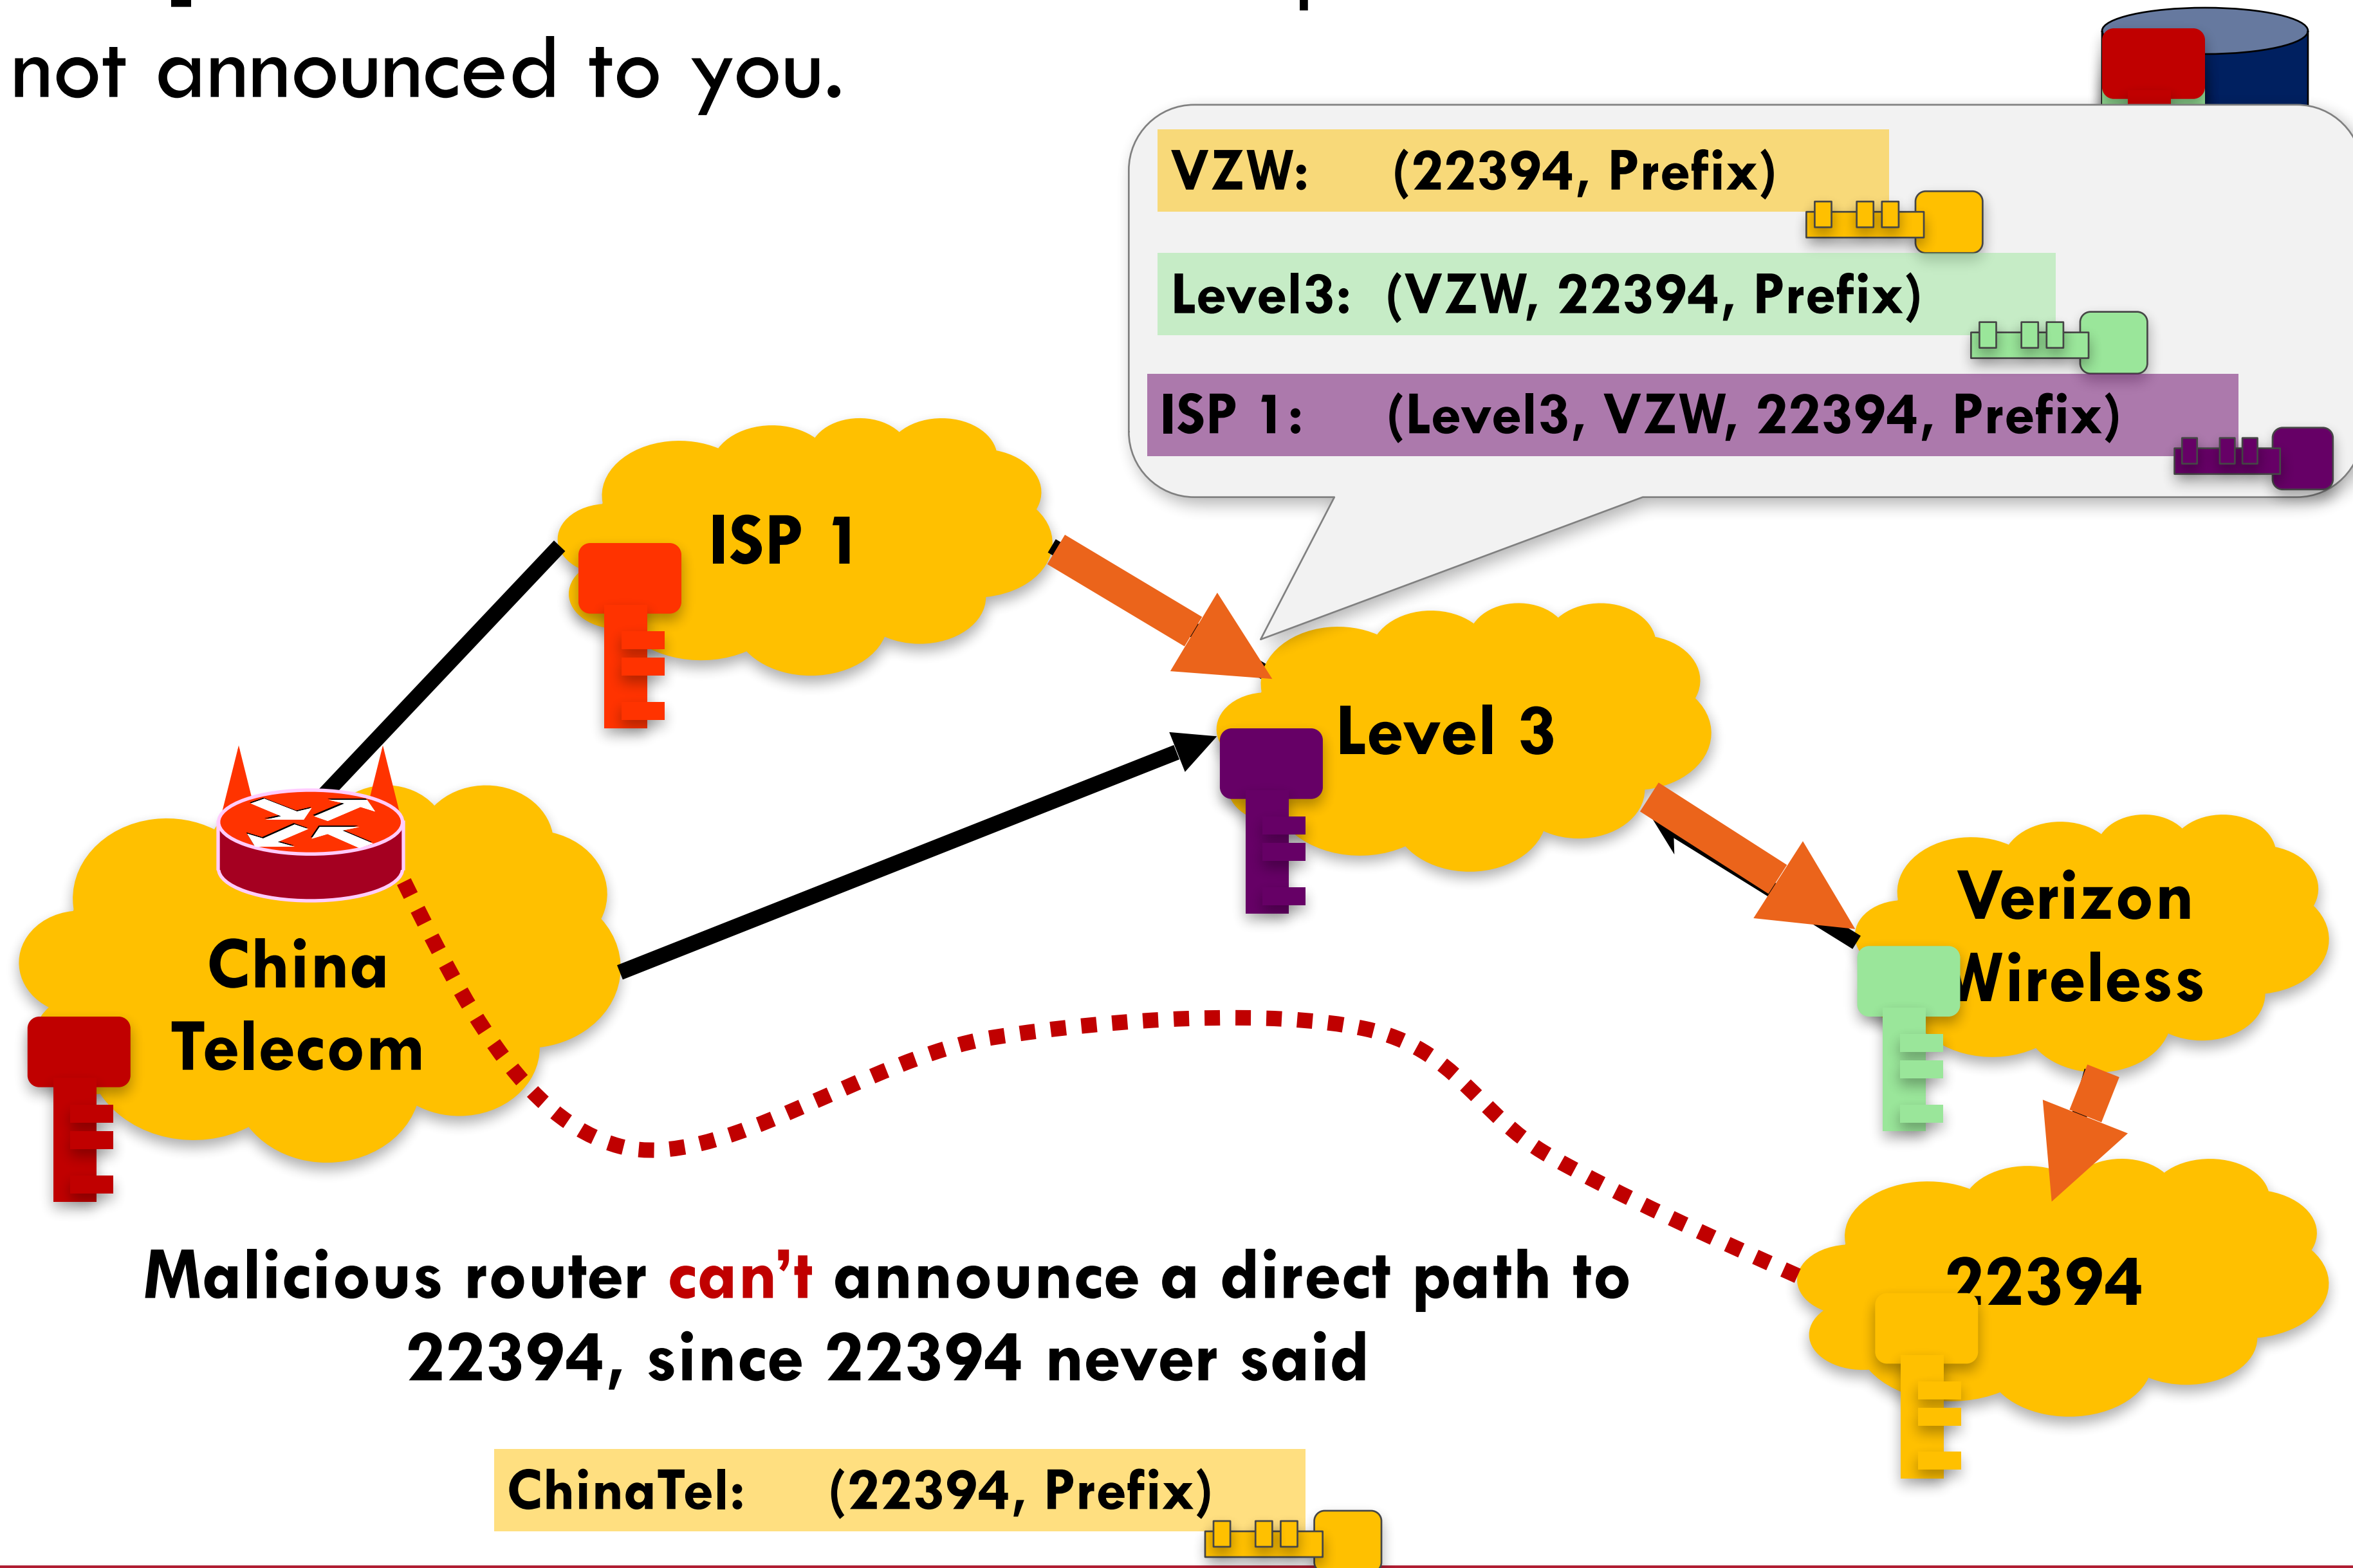
S-BGP [1997]: RPKI + Cannot announce a path that was not announced to you.



Public Key Signature: Anyone with 22394's public key can validate that the message was sent by 22394.



S-BGP [1997]: RPKI + Cannot announce a path that was not announced to you.



S-BGP Secure Version of BGP

- Address attestations
 - Claim the right to originate a prefix
 - Signed and distributed out-of-band
 - Checked through delegation chain from ICANN
- Route attestations
 - Distributed as an attribute in BGP update message
 - Signed by each AS as route traverses the network
 - Signature signs previously attached signatures
- S-BGP can validate
 - AS path indicates the order ASes were traversed
 - No intermediate ASes were added or removed



What might be hard about upgrading BGP to S-BGP?



S-BGP Deployment Challenges

- Complete, accurate registries
 - E.g., of prefix ownership
- Public Key Infrastructure
 - To know the public key for any given AS
- Cryptographic operations
 - E.g., digital signatures on BGP messages
- Need to perform operations quickly
 - To avoid delaying response to routing changes
- Difficulty of incremental deployment
 - Hard to have a “flag day” to deploy S-BGP



S-BGP Deployment Challenges

- Need ISPs to **agree on** and **deploy** a new protocol!
 - These are competing organizations!
- Economic incentives?
 - Doesn't improve performance
 - Hard to convince customers to pay more for security
- No benefit to unilateral deployment
 - Need entire path to deploy SBGP/soBGP before you get any benefit!
 - Like IPv6.... But worse 😞



Has S-BGP been adopted?

- Sadly, no
- If you solve this or want to solve this you can go to grad school
 - Or join a big company's networking team
 - Lots of people will thank you
 - You will be very popular at Internet parties



Summary

- BGP was built on the assumption of cooperation
- Assumption fails due to attacks... and just to errors.
- Proposed fixes are many, but all have some limitations
 - S-BGP
 - Relies on a PKI
 - Potentially significant overhead
- Very hard to retrofit security in an existing model!




DoS: General definition

- DoS is **not** access or theft of information or services
- Instead, goal is to stop the service from operating
- Deny service to legitimate users

- Why?
 - Economic, political, personal etc ..



Smurf amplification DoS attack




- Send ping request to broadcast addr (ICMP Echo Req)
- Lots of responses:
 - Every host on target network generates a ping reply (ICMP Echo Reply) to victim

Prevention: reject external packets to broadcast address



Modern day example (May '06)

DNS Amplification attack: (×50 amplification)



580,000 open resolvers on Internet (Kaminsky-Shiffman'06)



“Resource Asymmetry”

- One attacker with one server generating traffic probably cannot completely overwhelm the victim.
- Smurf and DNS attacks:
 - Attacker can harness arbitrary machines (lots of them!)
 - Receiver is just one server.
 - “Resource Asymmetry” is the problem.



How much traffic do I need to
overwhelm a receiver?



Look up: Victim, Year, Bandwidth of Attack



ddos bandwidth




Evolution of (D)DoS in history

- Point-to-point DoS attacks
- TCP SYN floods, Ping of death, etc..
- Smurf (reflection) attacks
- Coordinated DoS
- Multitarget DDoS
- P2P botnets

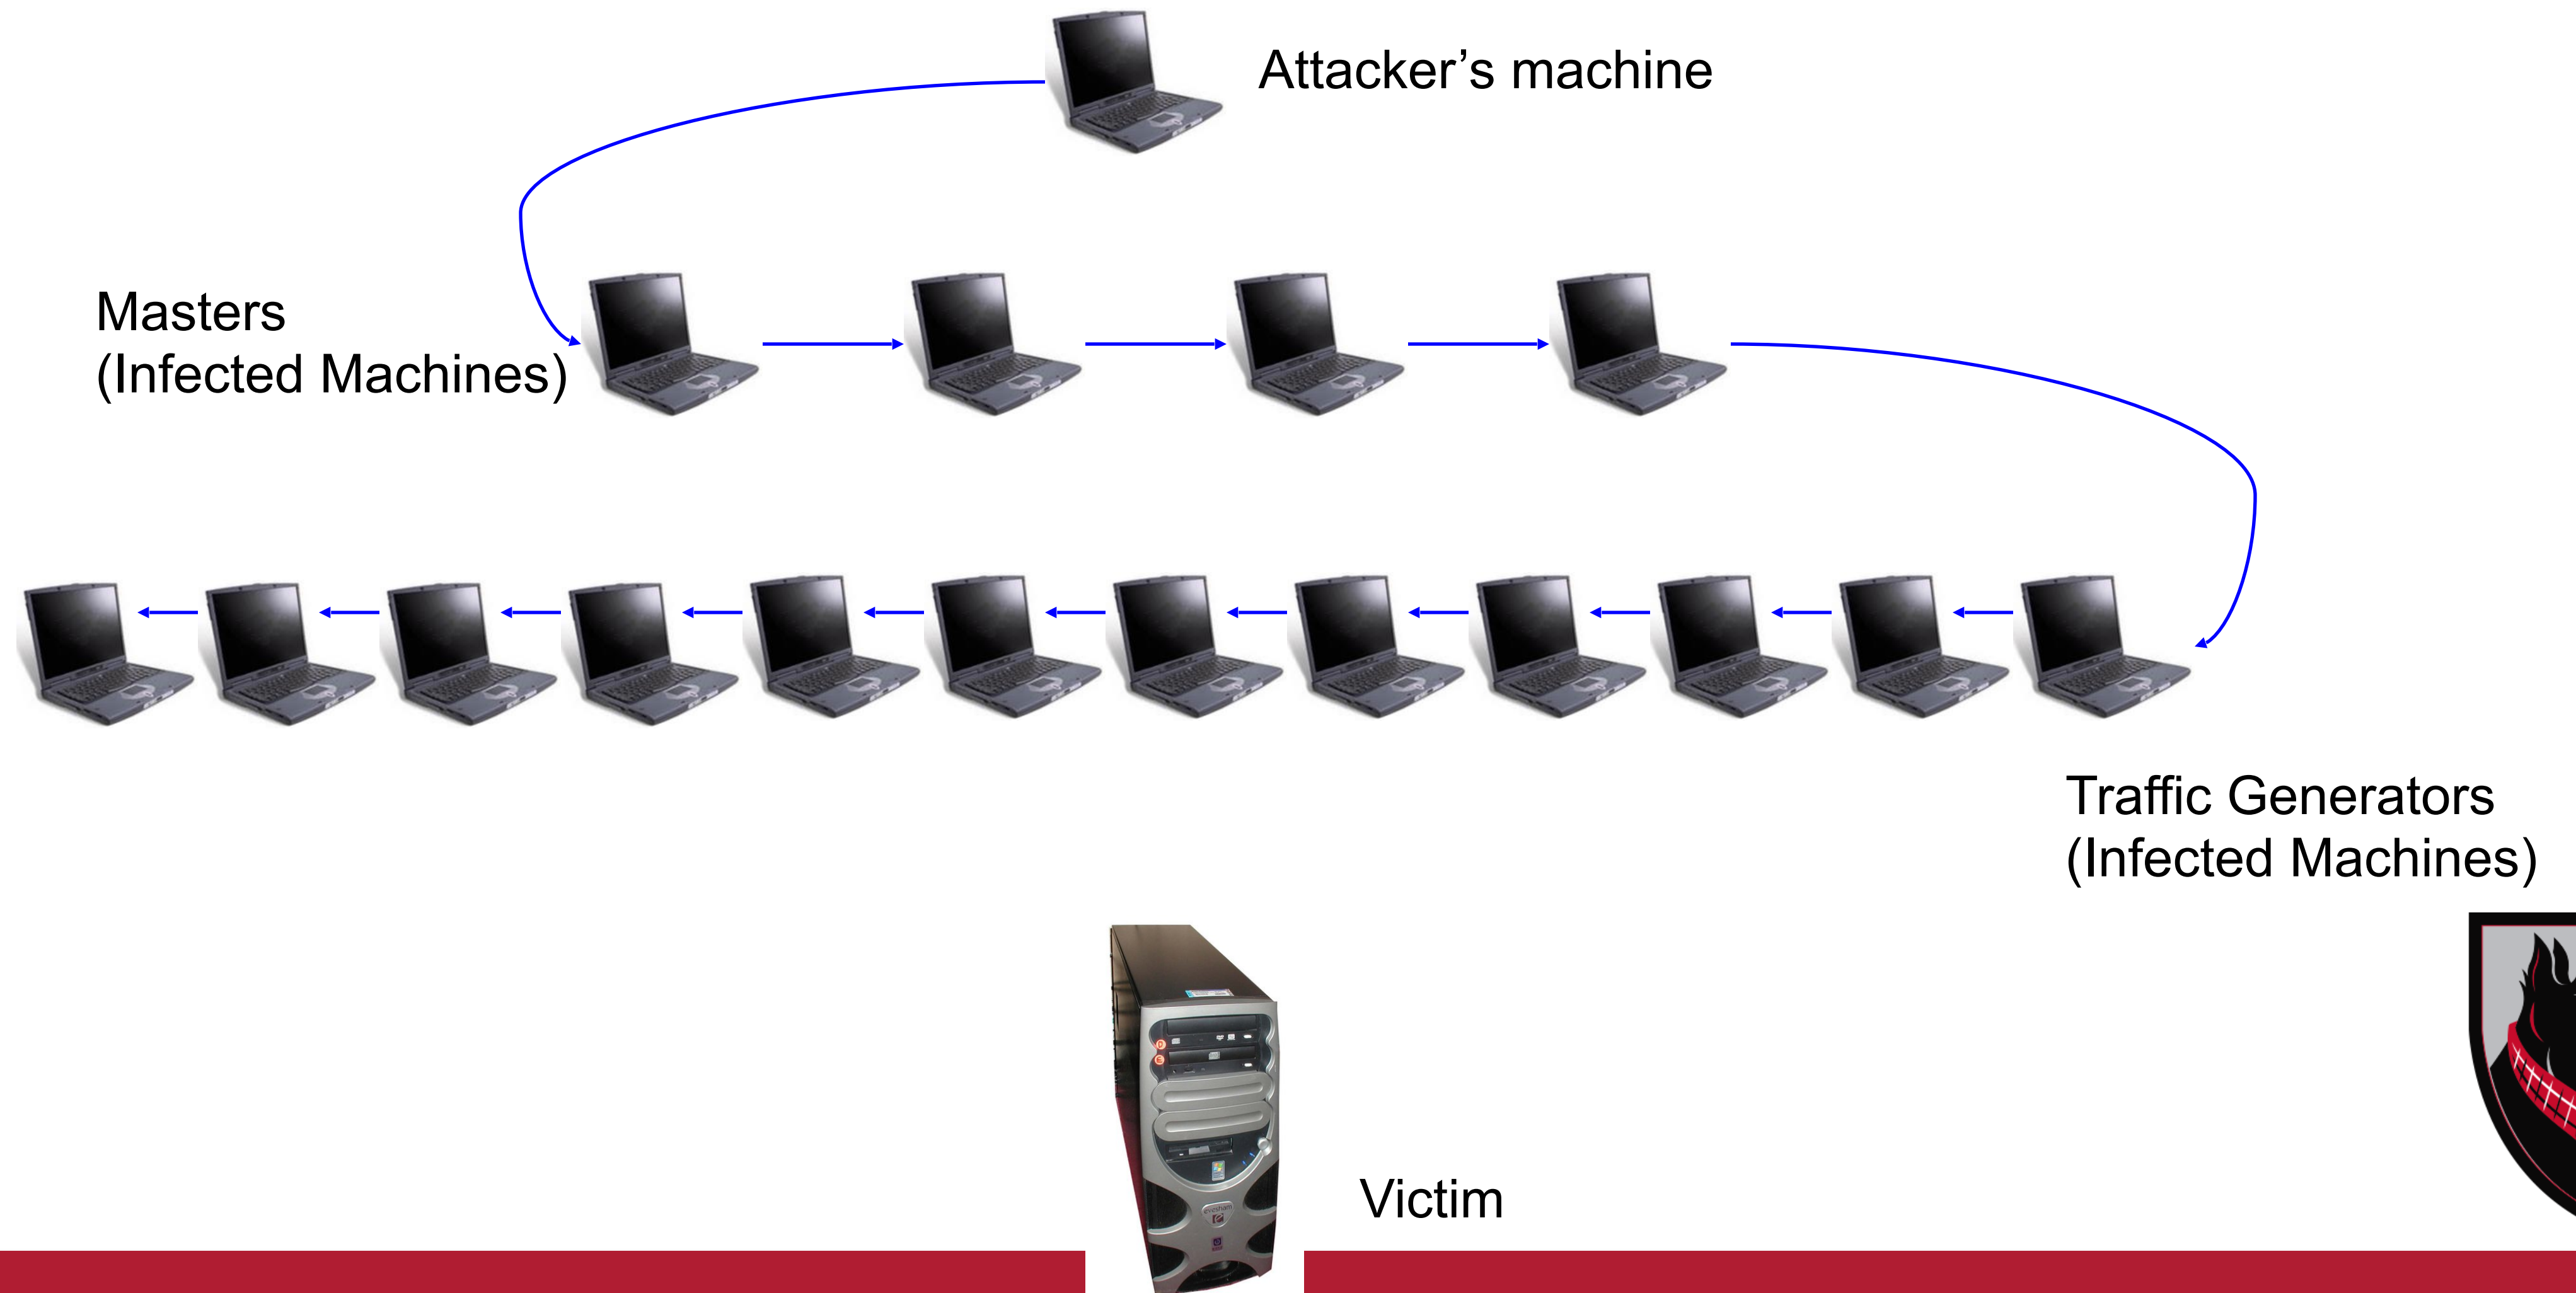


Coordinated DoS


- Simple extension of DoS
- Coordination between multiple parties
 - Can be done off-band
 - IRC channels, email...



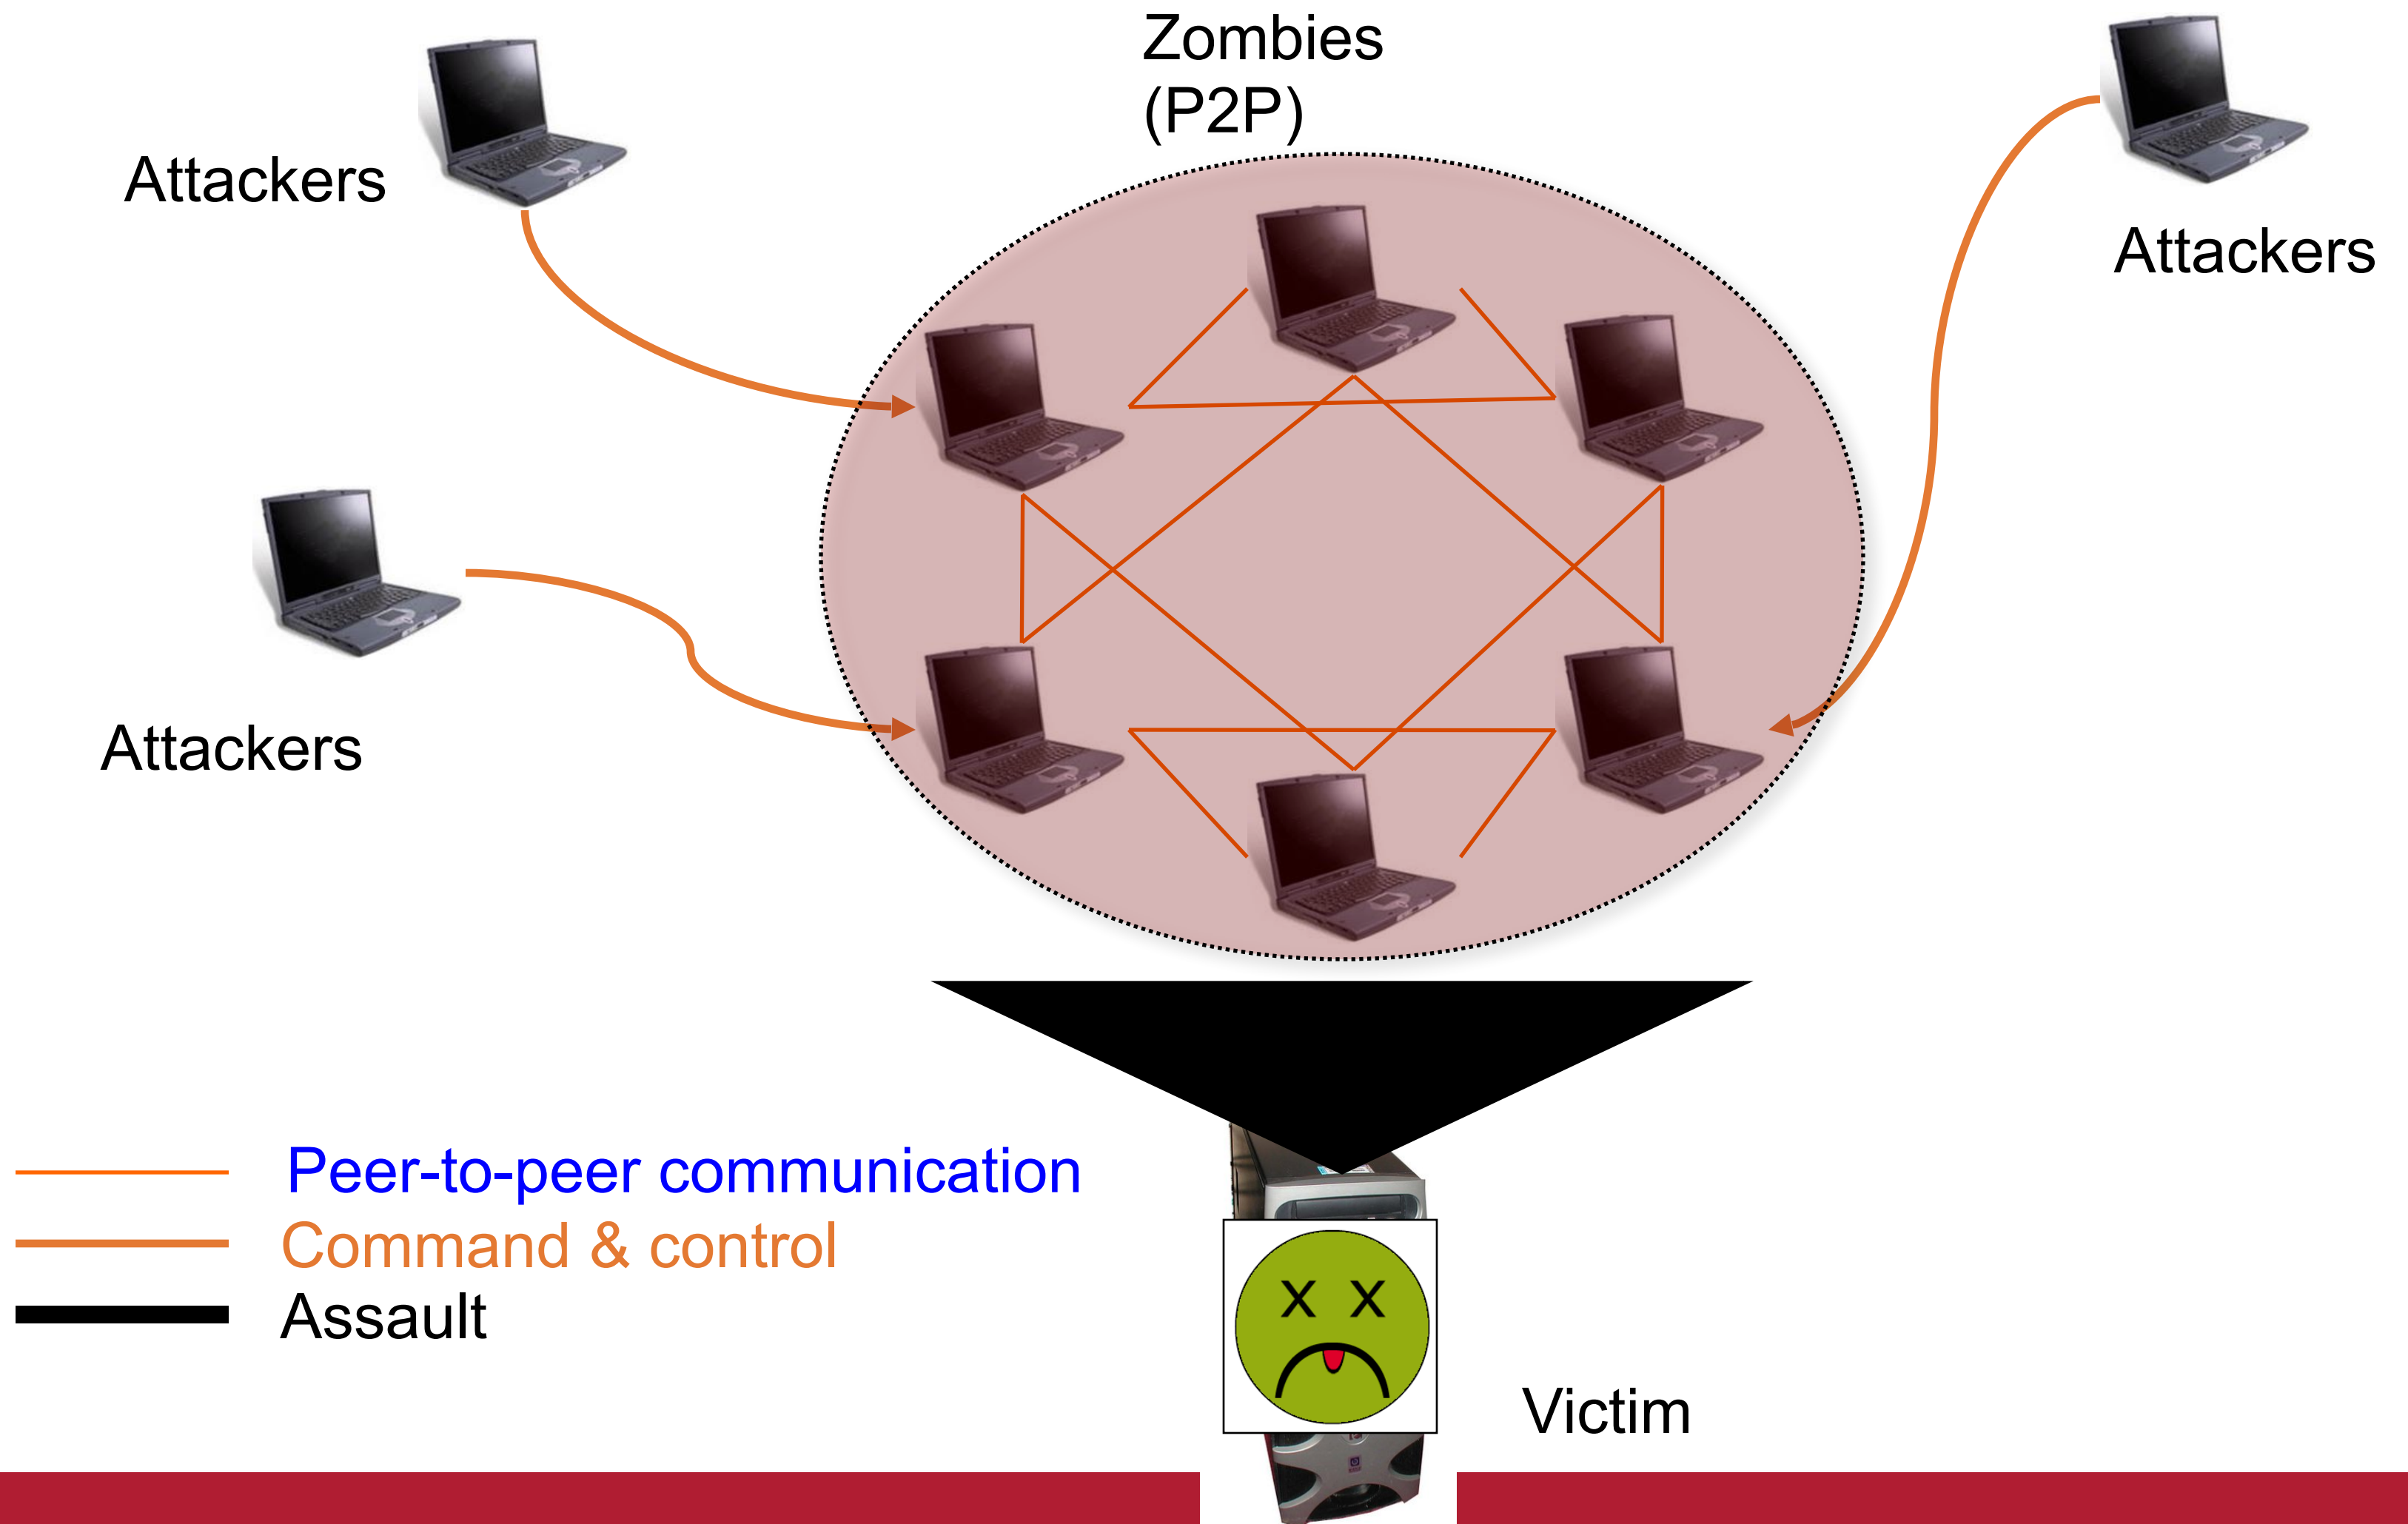
Typical DDoS setup circa 2005



Typical DDoS setup circa 2005



Modern Botnet setup



Goal: Overload the Host and Disable their Availability

- Multiple ways to achieve overload!
 - Smurf and DNS amplification attacks overload the network link.
 - Botnets can do that too.




DoS Attacks Characteristics

- Link flooding causes high loss rates for incoming traffic
- TCP throughput


$$\downarrow BW = \frac{MSS \cdot C}{RTT \cdot \sqrt{q}} \uparrow$$

- During DoS few legitimate clients served



Content Distribution Networks (CDNs)

- CDN company installs hundreds of CDN servers throughout Internet
- Replicated customers' content
- How can this help DDoS?
- Legitimate requests can still go through
- Attack scale must be higher




Some CDNs even specialize in DDoS Defense!

Cloudflare now offers unmetered DDoS attack mitigation

Posted Sep 25, 2017 by [Ron Miller \(@ron_miller\)](#)



Next Story



Crunchbase

Cloudflare

FOUNDED
2009

OVERVIEW

Cloudflare is a web performance and security company that provides online services to protect and accelerate websites online. The company's online platforms include Cloudflare CDN, which distributes content around the world to speed up websites; Cloudflare optimizer that enables web pages with ad servers and third-party widgets to download Snappy software on mobiles and computers; CloudFlare ...

LOCATION

San Francisco, CA

CATEGORIES

Security, Web Hosting, Advertising, Analytics, Ad Server, Enterprise Software

FOUNDERS


Michelle Zatlyn



Finding the Zombies and Killing Them

Bot Detection and Removal

Detection, notification, and prevention against malicious software. Have you noticed any suspicious email account activity, unusual error messages, or unfamiliar browsers? Your computer may be infected by a "bot," malicious software that secretly uses your computer to send spam, host phishing sites, and steal your personal information.



How our proactive bot notification works


The XFINITY Internet Security bot notification tool looks for patterns coming from your home network that match our infection libraries. If we suspect that a device on your home network is

Goal: Overload the Host and Disable their Availability

- Multiple ways to achieve overload!
 - Smurf and DNS amplification attacks overload the network link.
 - Botnets can do that too.
 - May also try to overload at the application or transport layer, e.g.:
 - Send a database a lot of very large queries
 - Open lots of TCP connections — “SYN attack”



TCP SYN Flood I: low rate (DoS bug)



Single machine:

- SYN Packets with **random source IP addresses**
- Fills up backlog queue on server
- No further connections possible



SYN Floods

(phrack 48, no 13, 1996)

OS	Backlog queue size
Linux 1.2.x	10
FreeBSD 2.1.5	128
WinNT 4.0	6

Backlog timeout: 3 minutes

- ⇒ Attacker need only send 128 SYN packets every 3 minutes.
- ⇒ Low rate SYN flood



How to prevent SYN flood attacks

- Non-solution:
 - Increase backlog queue size or decrease timeout
- Correct solution (when under attack) :
 - **Syncookies**: remove state from server
 - Small performance overhead



Syncookies [Bernstein, Schenk]

- Idea: use secret key and data in packet to gen. server SN
- Server responds to Client with SYN-ACK cookie:
 - $T = 5\text{-bit counter incremented every 64 secs.}$
 - $L = \text{MAC}_{\text{key}}(\text{SAddr}, \text{SPort}, \text{DAddr}, \text{DPort}, \text{SN}_C, T)$ [24 bits]
 - key: picked at random during boot
 - $\text{SN}_S = (T \cdot \text{mss} \cdot L)$ ($|L| = 24 \text{ bits}$)
 - **Server does not save state** (other TCP options are lost)
- Honest client responds with ACK ($\text{AN}=\text{SN}_S$, $\text{SN}=\text{SN}_C+1$)
 - Server allocates space for socket only if valid SN_S .



What about attacks on applications
— like RPC calls and database
queries?



Client puzzles

- Idea: slow down attacker
- Moderately hard problem:
 - Given challenge C find X such that
$$\text{LSB}_n (\text{SHA-1}(C \parallel X)) = 0^n$$
 - Assumption: takes expected 2^n time to solve
 - For $n=16$ takes about .3sec on 1GHz machine
 - Main point: checking puzzle solution is easy. Pushes resource requirements to attacker!
- During DoS attack:
 - Everyone must submit puzzle solution with requests
 - When no attack: do not require puzzle solution



What about a DDoS attack on a web server?
(There is a simple mechanism, invented at
Carnegie Mellon, that you have all used)



CAPTCHAs

- Idea: verify that connection is from a human



- Applies to application layer DDoS [Killbots '05]
 - During attack: generate CAPTCHAs and process request only if valid solution
 - Present one CAPTCHA per source IP address.



What do net operators do?

- Best common operational practices:
- <http://nabcop.org/index.php/DDoS-DoS-attack-BCOP>
- Often, blackholing malicious looking IPs and rerouting to custom “Scrubbers” / Firewalls



THIS IS A SAD STORY



I HAVE JUST LISTED A TON OF
PROBLEMS WITH THE INTERNET
NONE OF WHICH ARE FULLY SOLVED



What needs to happen to fix BGP?
Why is solving the BGP security
problem challenging?



What needs to happen to fix BGP?
Why is solving the DDoS security
problem challenging?



Summary...

- Today: two classes of attacks on Internet availability.
 - Routing attacks on BGP to prevent traffic from reaching victim
 - Need to validate routes... but getting all 50k+ networks to upgrade is challenging.
 - DoS and DDoS to overwhelm resources of victim
 - Modern bonnets mean attackers can amass large amounts of resources to overrun victims
 - No “off button” on the Internet — all traffic is allowed through by the network, even if it is unwanted :(



PROJECT FEEDBACK

Side #1: Thoughts and recommendations for Project 2
NEXT YEAR (besides, “get the tests working early”, sigh)

Side #2: Thoughts, recommendations, and requests for
Project 3 or anything in this class THIS YEAR

