CSE508 Network Security

2/3/2016

Core Protocols: BGP and DNS

Michalis Polychronakis

Stony Brook University

IP Addressing and Forwarding

Packets are routed based on their dst. IP address

Router's task: for every possible IP address, forward packet to the next hop

Table lookup for each packet in a routing table

For 32-bit addresses, 2^{32} possibilities! ➔ impractical

Solution: hierarchical address scheme

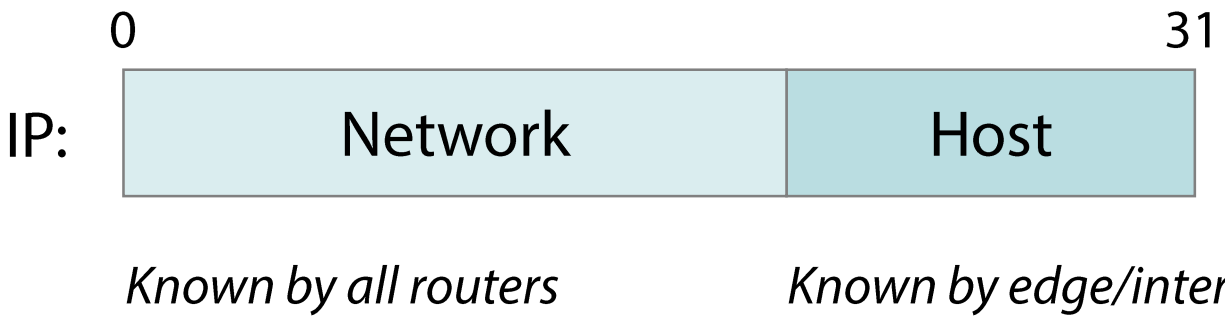

IPv4 Address Classes

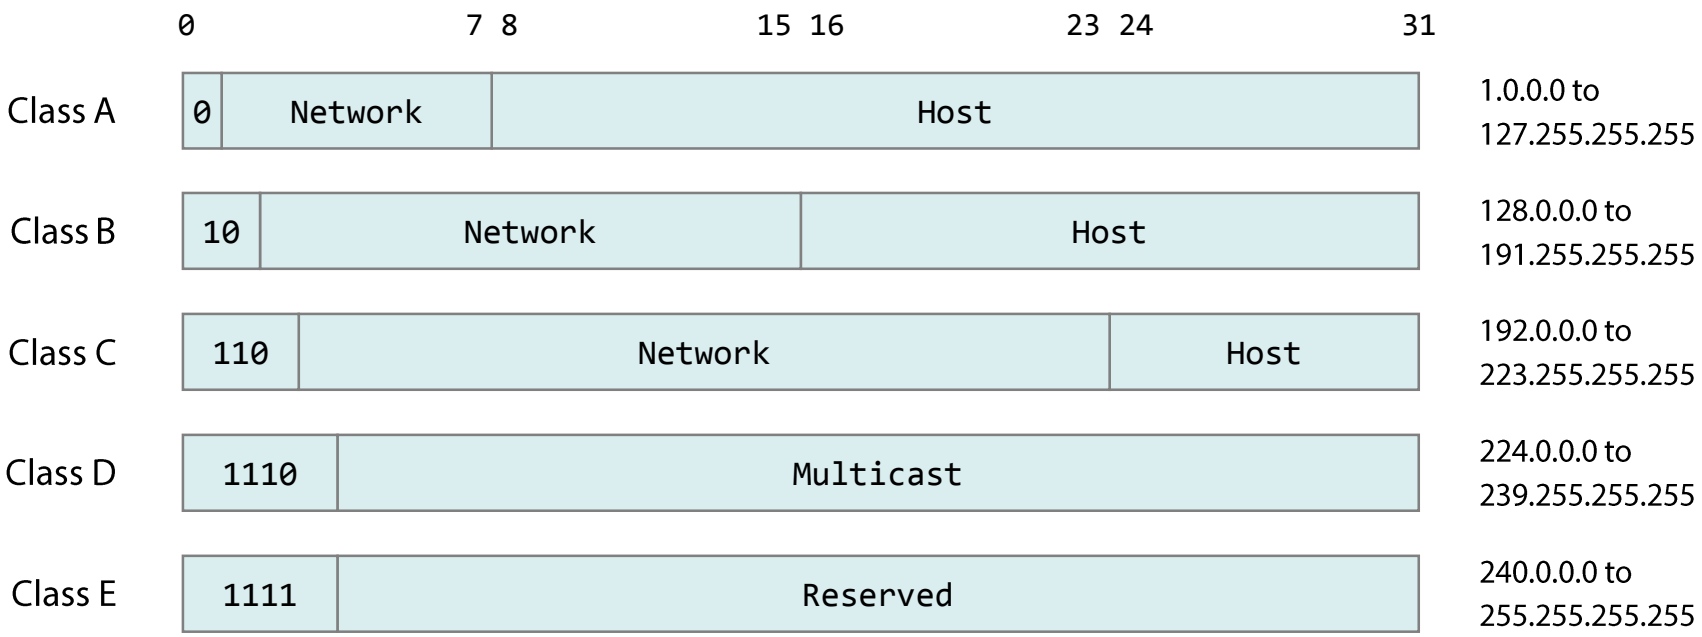

MAP OF THE INTERNET

THE IPv4 SPACE, 2006

THIS CHART SHOWS THE IP ADDRESS SPACE ON A PLANE USING A FRACTAL MAPPING WHICH PRESERVES GROUPING -- ANY CONSECUTIVE STRING OF IPs WILL TRANSLATE TO A SINGLE COMPACT, CONTIGUOUS REGION ON THE MAP. EACH OF THE 256 NUMBERED BLOCKS REPRESENTS ONE /8 SUBNET (CONTAINING ALL IPs THAT START WITH THAT NUMBER). THE UPPER LEFT SECTION SHOWS THE BLOCKS SOLD DIRECTLY TO CORPORATIONS AND GOVERNMENTS IN THE 1990's BEFORE THE RIRs TOOK OVER ALLOCATION.

0 1 14 15 16 19 →
3 2 13 12 17 18
4 7 8 11
5 6 9 10

= UNALLOCATED BLOCK

No green patches after 2011…

Autonomous Systems

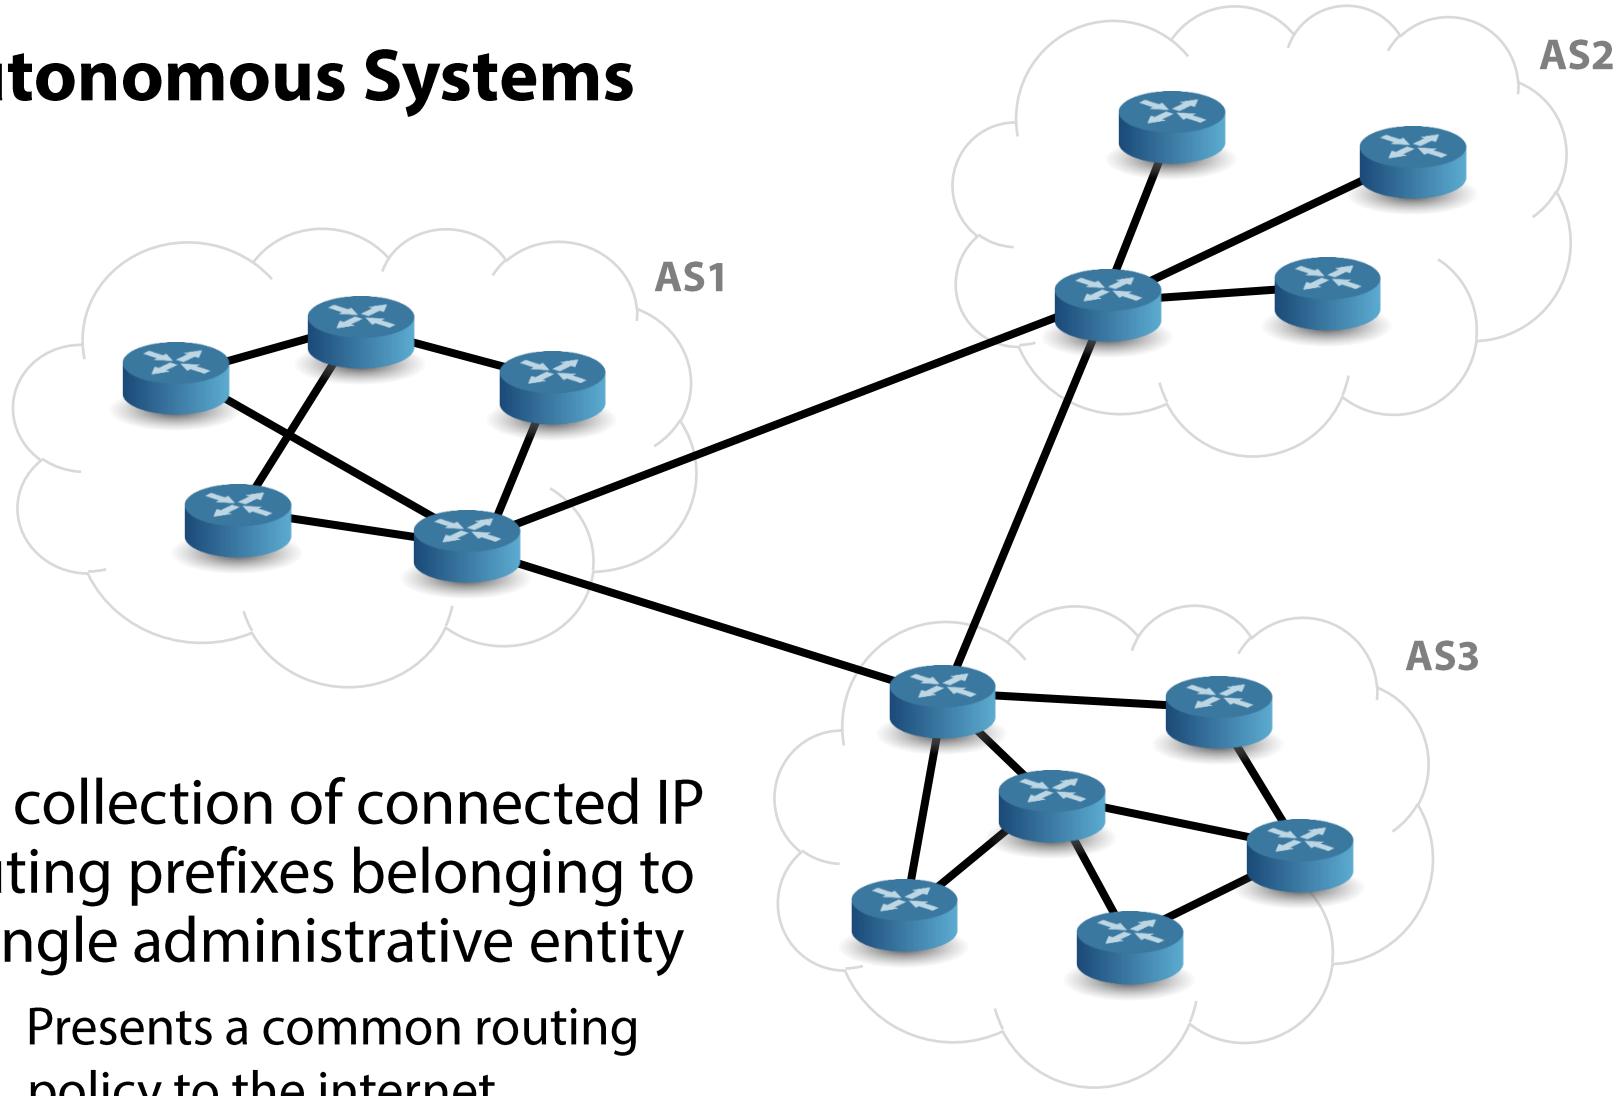

AS: collection of connected IP routing prefixes belonging to a single administrative entity

- Presents a common routing policy to the internet

AS number defined as 16-bit integer

- ~47,000 ASNs as of 2014, assigned by IANA

Map of the internet, 1982

Ovals: sites/networks
Rectangles: routers

Created by Jon Postel

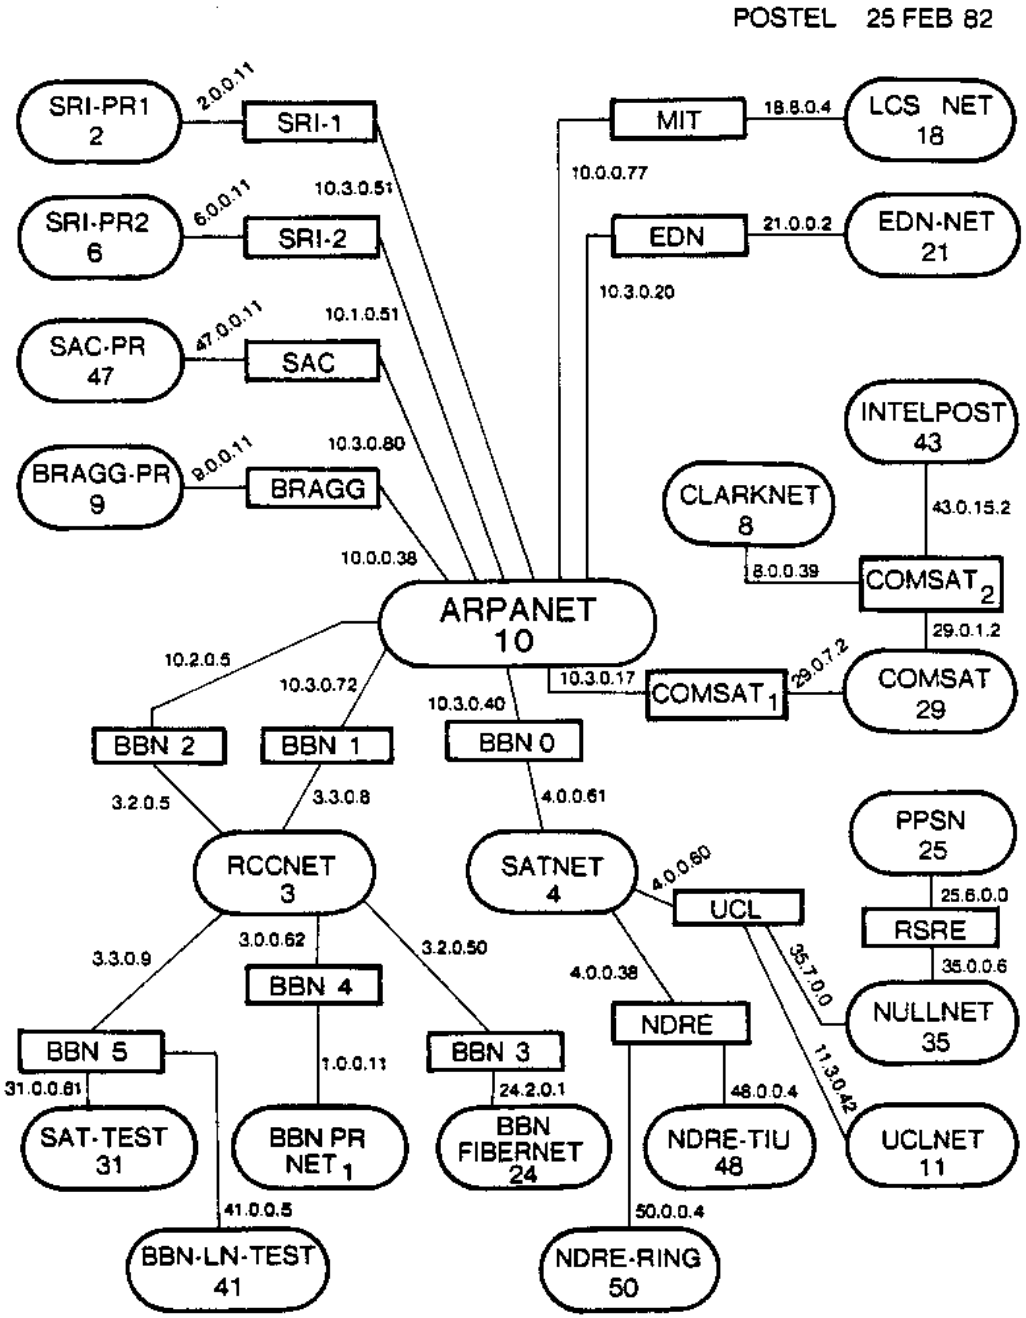

CAIDA's IPv4 AS Core
AS-level INTERNET GRAPH

Skitter January 2000

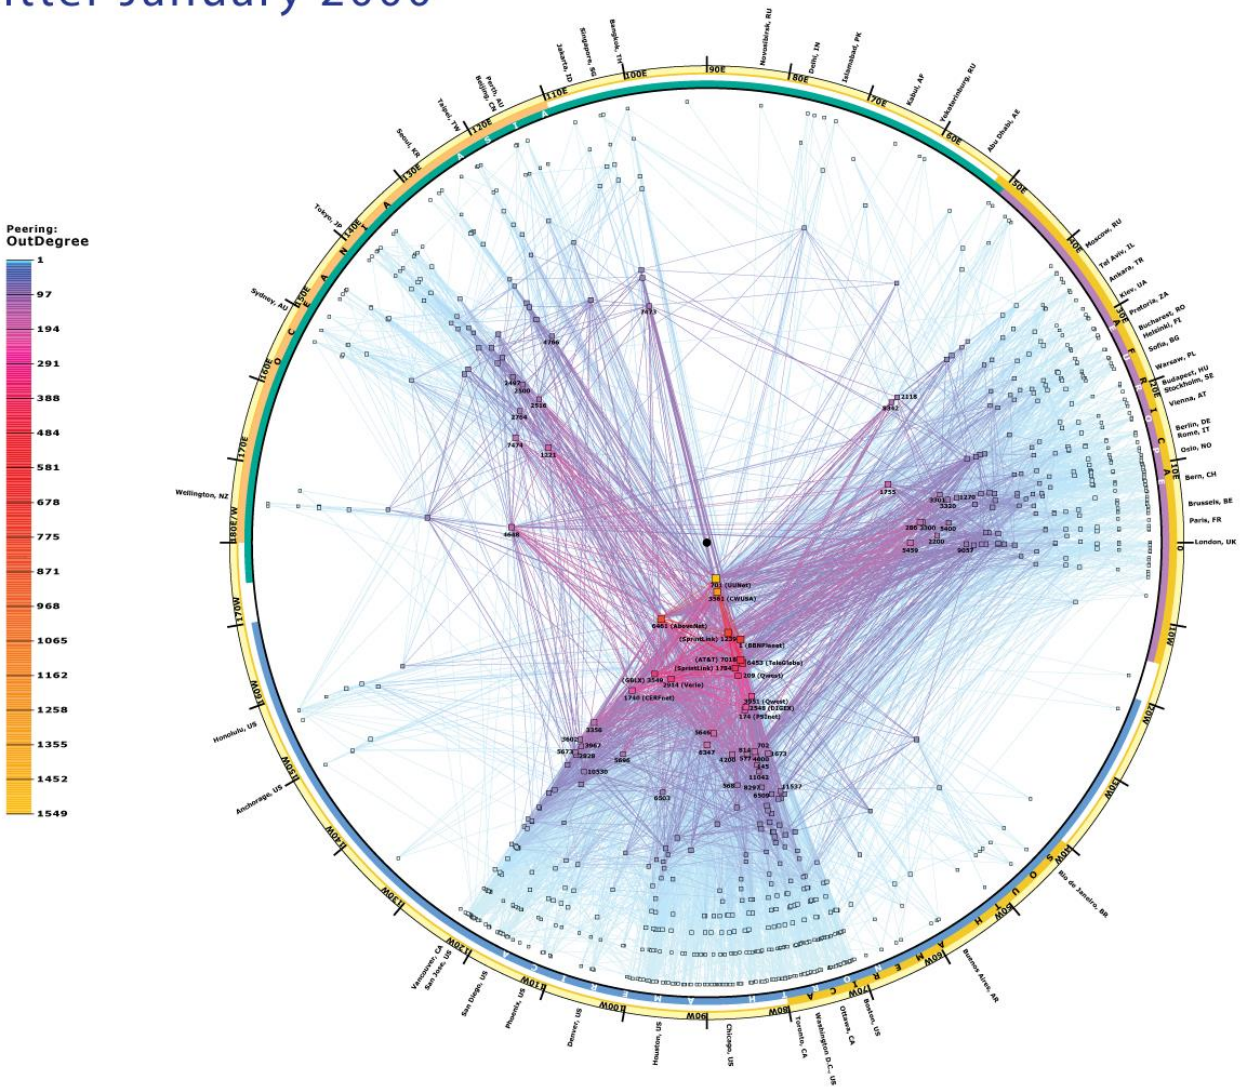

CAIDA's IPv4 AS Core
AS-level INTERNET GRAPH

Archipelago August 2007

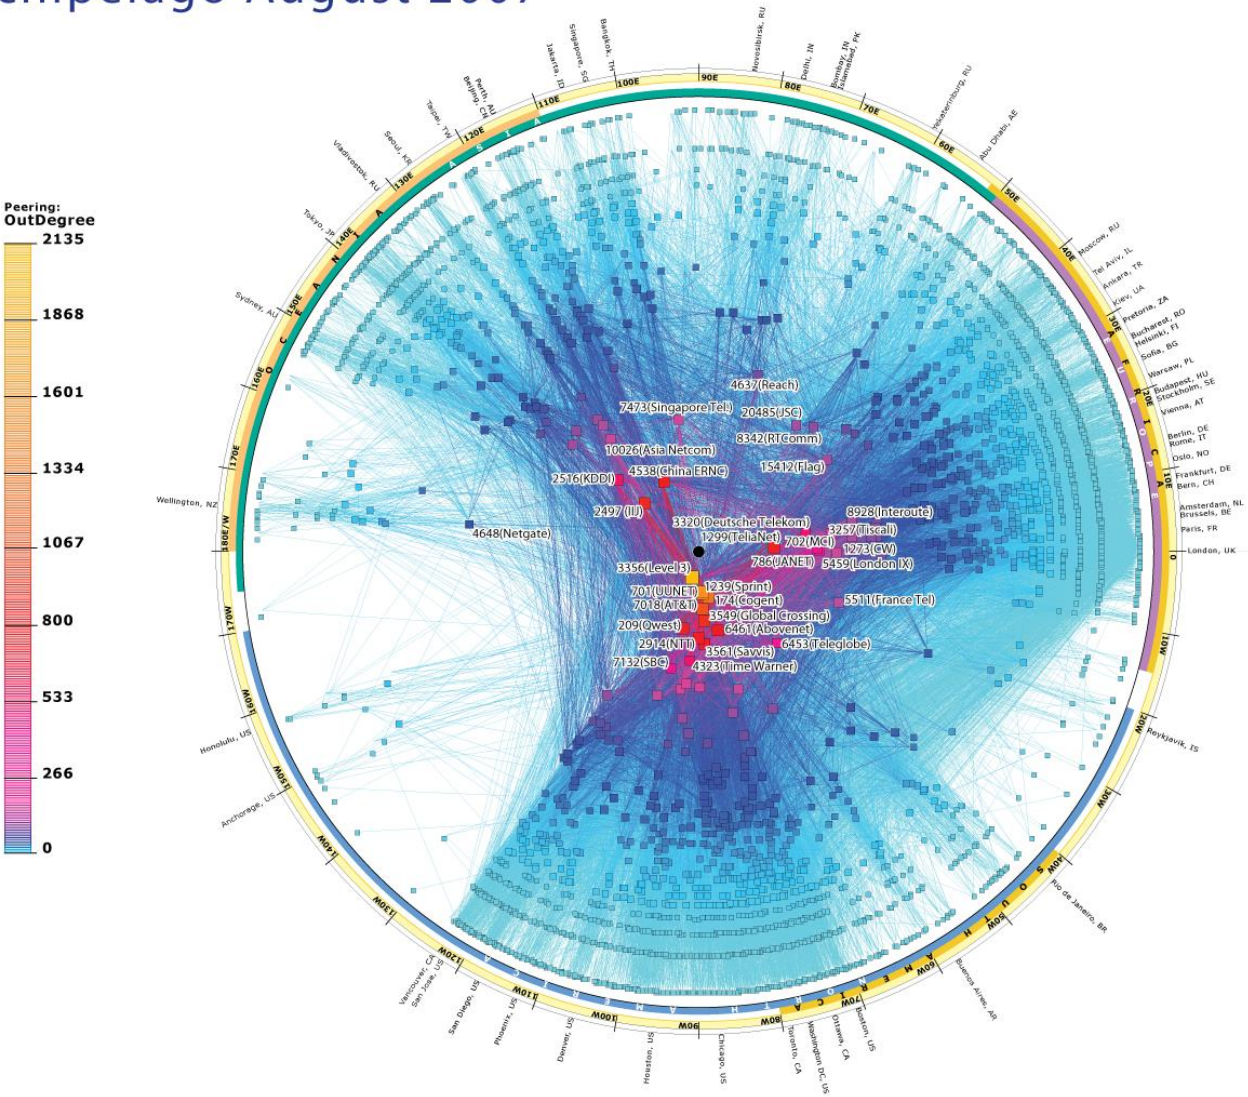

CAIDA's IPv4 AS Core
AS-level INTERNET GRAPH

Archipelago January 2015

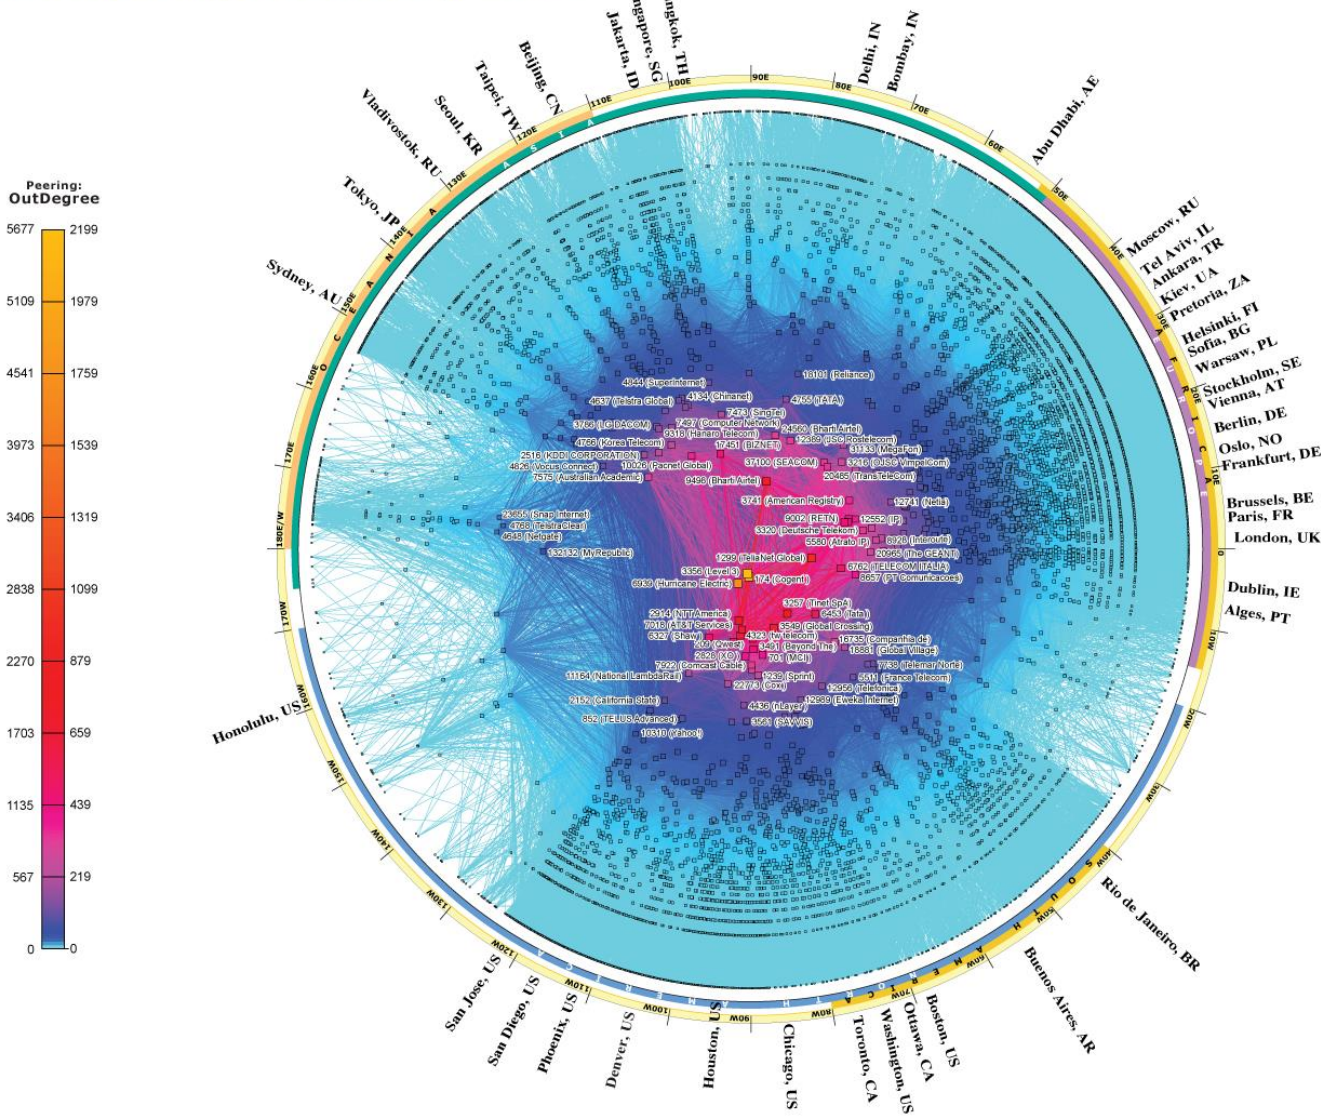

Internet Routing

Routers speak to each other to establish internet paths

- Exchange topology and cost information
- Calculate the best path to each destination

Intra-domain routing: set up routes within a single network/AS

- **RIP** (Routing Information Protocol): distance vector
- **OSPF** (Open Shortest Path First): link state

Inter-domain routing: set up routes between networks

- **BGP** (Border Gateway Protocol)
- Advertisements contain a prefix and a list of ASes to traverse to reach that prefix

Internet Routing

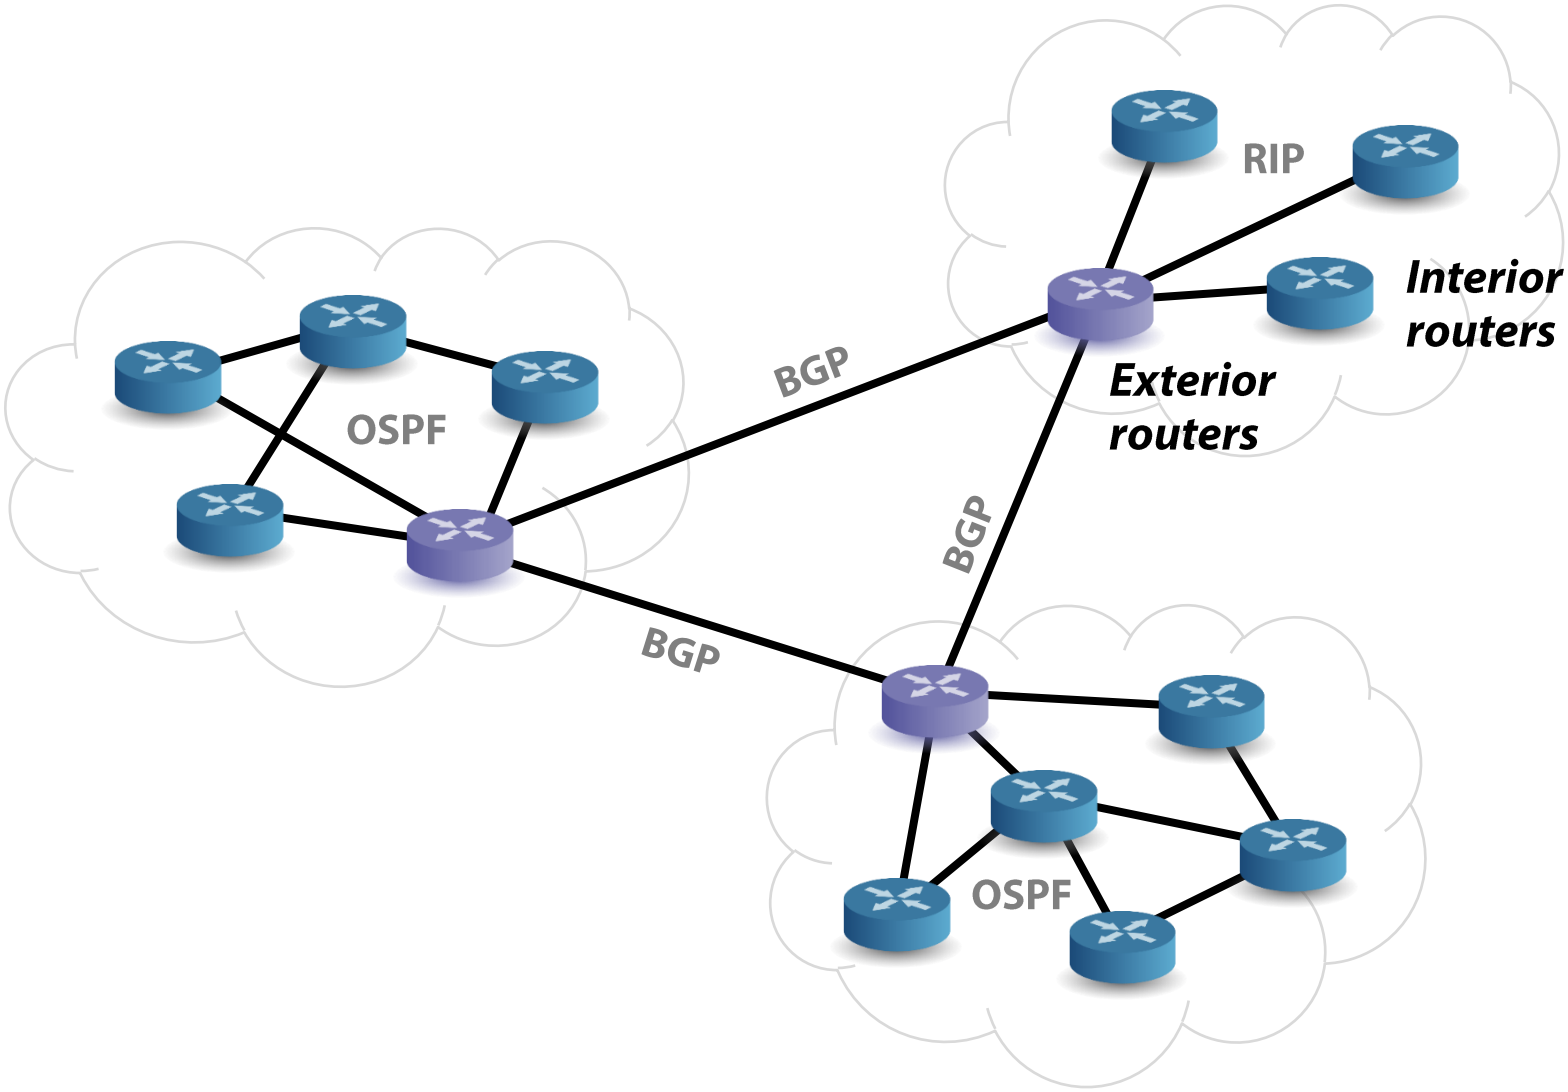

BGP Security Issues

No authentication of path announcements

- Neighbor adjacencies can be "secured" using MD5 digests

BGP messages are sent over TCP connections

- All the usual problems: eavesdropping, content manipulation, …

Misconfigurations are easy

- Complex interactions

Attackers can lie to other routers

Routing Attacks

Blackholing: false route advertisements to attract and drop traffic

Redirection: force traffic to take a different path, either for eavesdropping/manipulation (MitM) or for causing congestion

Instability: frequent advertisements and withdrawals and/or increased BGP traffic to cause connectivity outages

How?

- Configuration mistakes
- Insider attacks
- Compromised routers (vulnerability exploitation, default credentials, …)
- Traffic manipulation

Prefix Hijacking

Announce someone else's prefix

Victim prefers shortest path

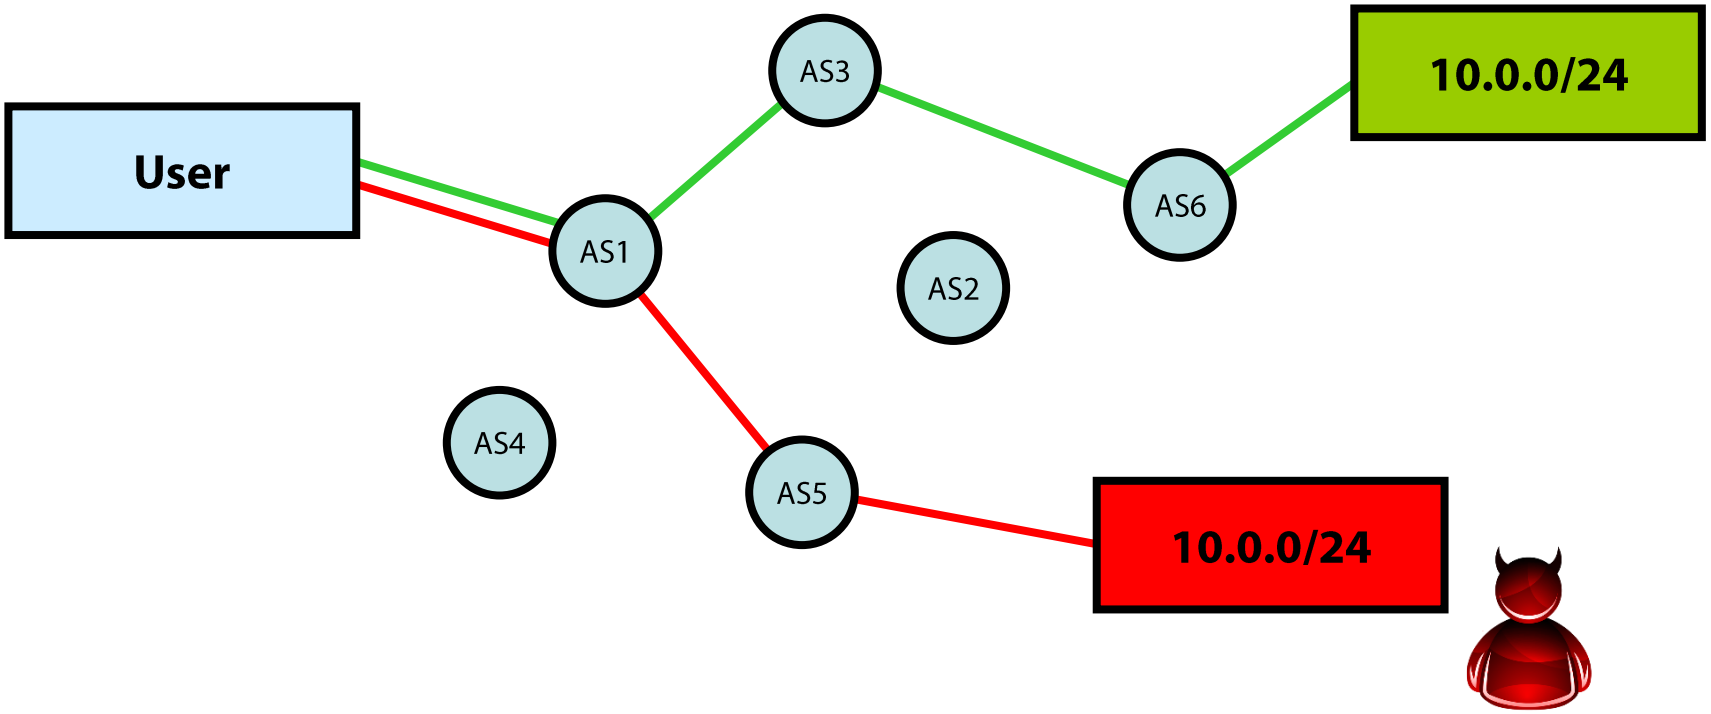

Prefix Hijacking

Announce a more specific prefix than someone else

- Victim prefers more specific path

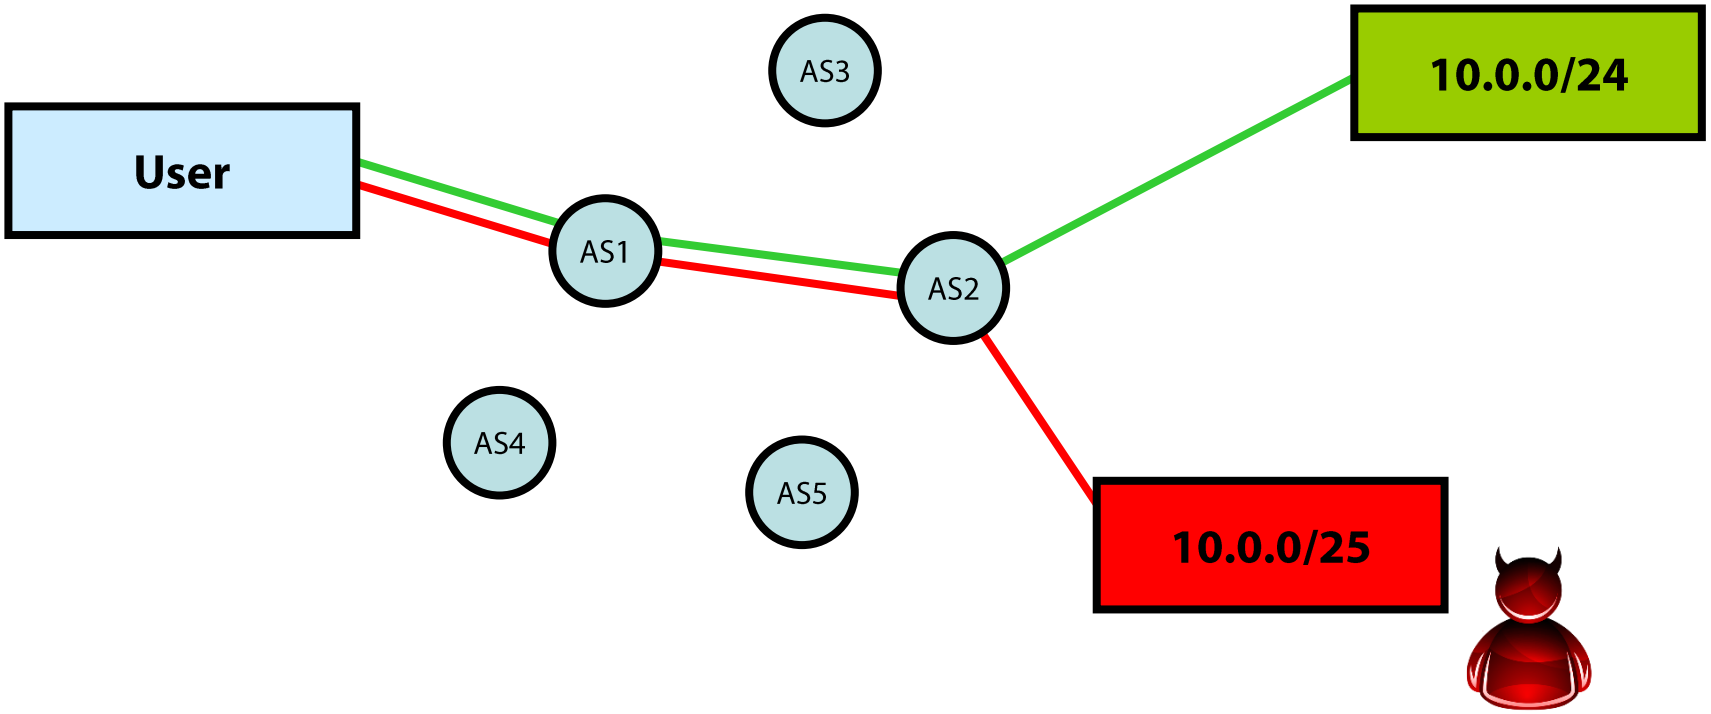

Pakistan's Accidental YouTube Re-Routing Exposes Trust Flaw in Net

BY RYAN SINGEL 02.25.08 | 10:37 AM | PERMALINK

Government: you have to block this YouTube video

Pakistan Telecom: OK

Use URL filtering? No

Change the DNS record? No

Use IP blocking? No

Let's just blackhole 208.65.153.0/24

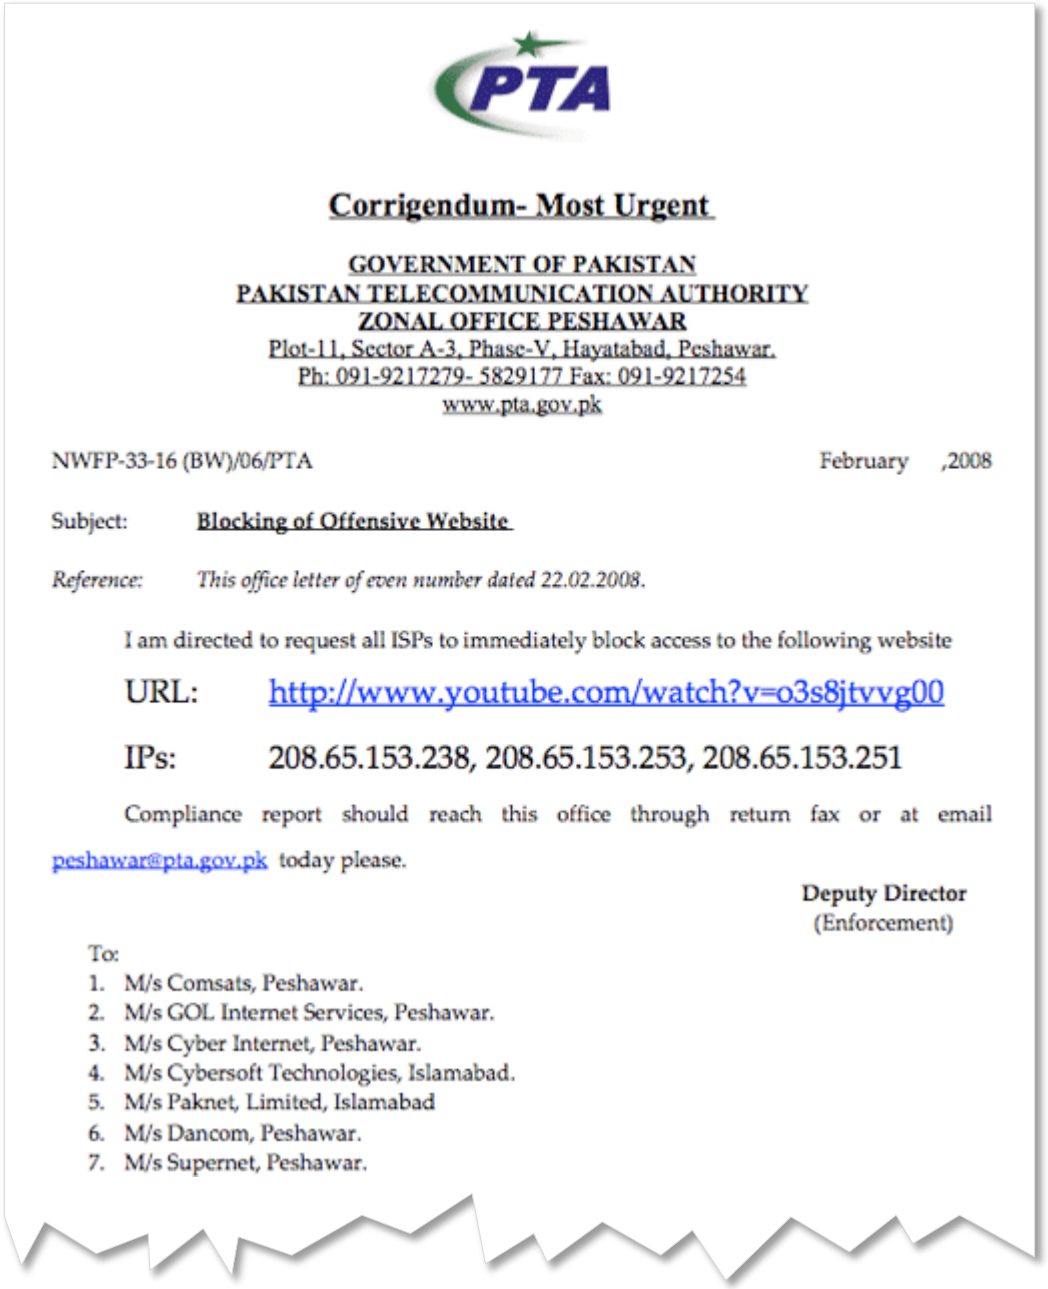

Corrigendum- Most Urgent

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.

Deputy Director
(Enforcement)

To:
1. M/s Comsats, Peshawar.
2. M/s GOL Internet Services, Peshawar.
3. M/s Cyber Internet, Peshawar.
4. M/s Cybersoft Technologies, Islamabad.
5. M/s Paknet, Limited, Islamabad
6. M/s Dancom, Peshawar.
7. M/s Supernet, Peshawar.

AS36561 (YouTube) announces 208.65.152.0/22

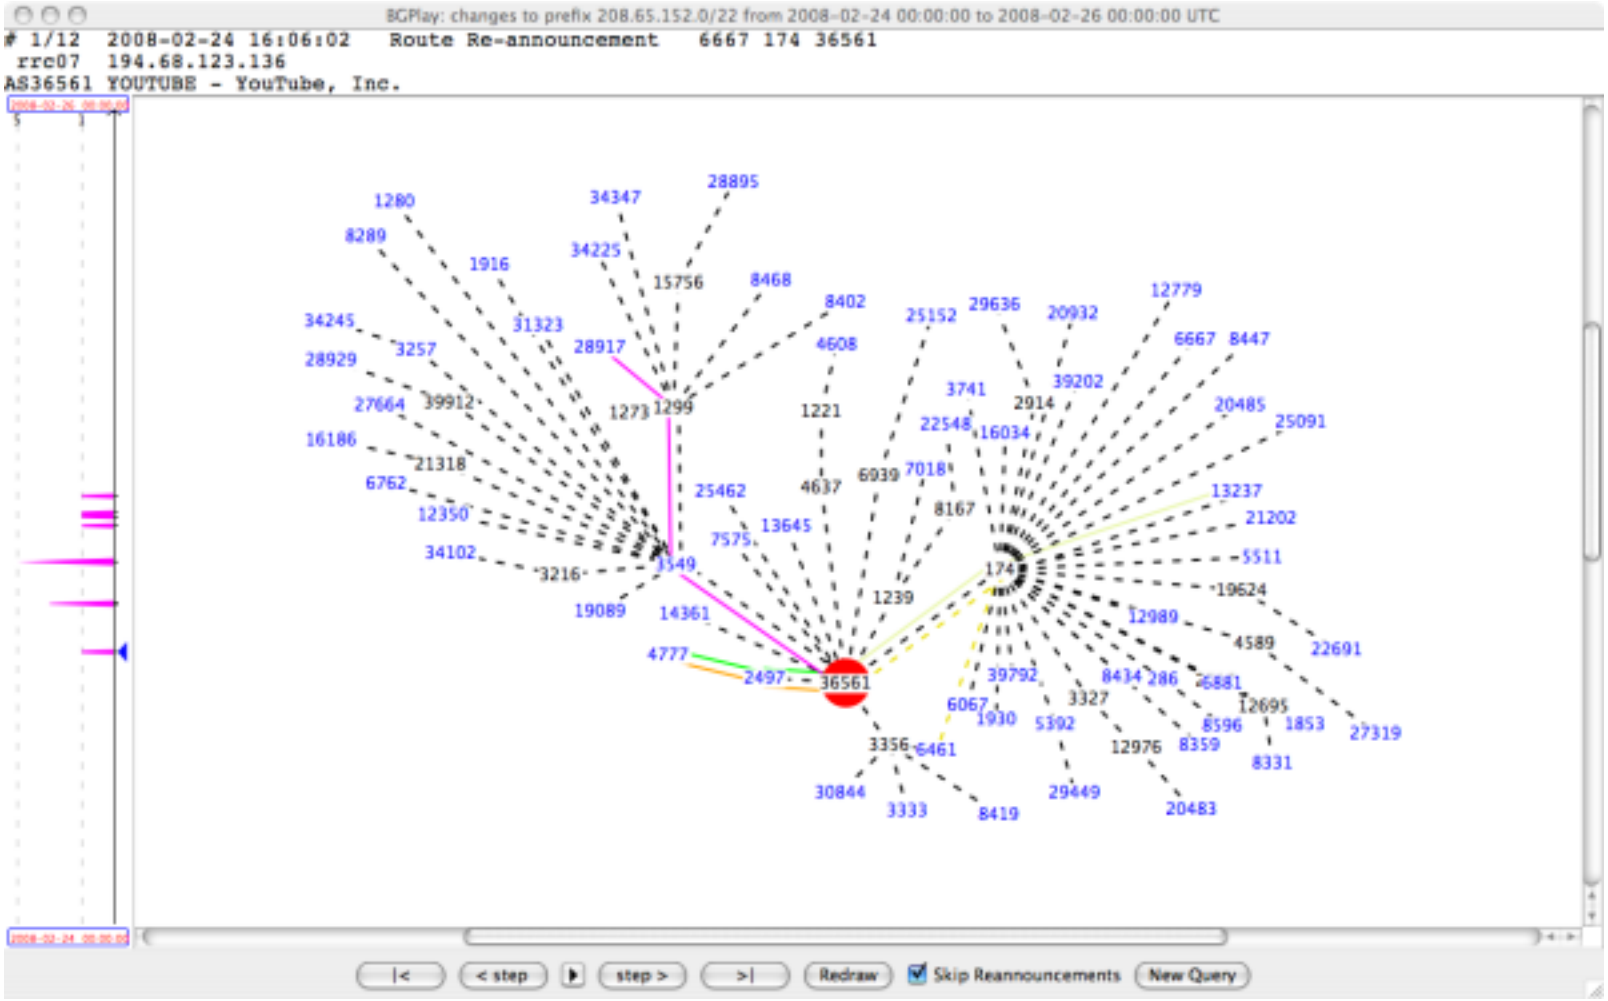

The prefix 208.65.153.0/24 is not announced on the Internet before the event

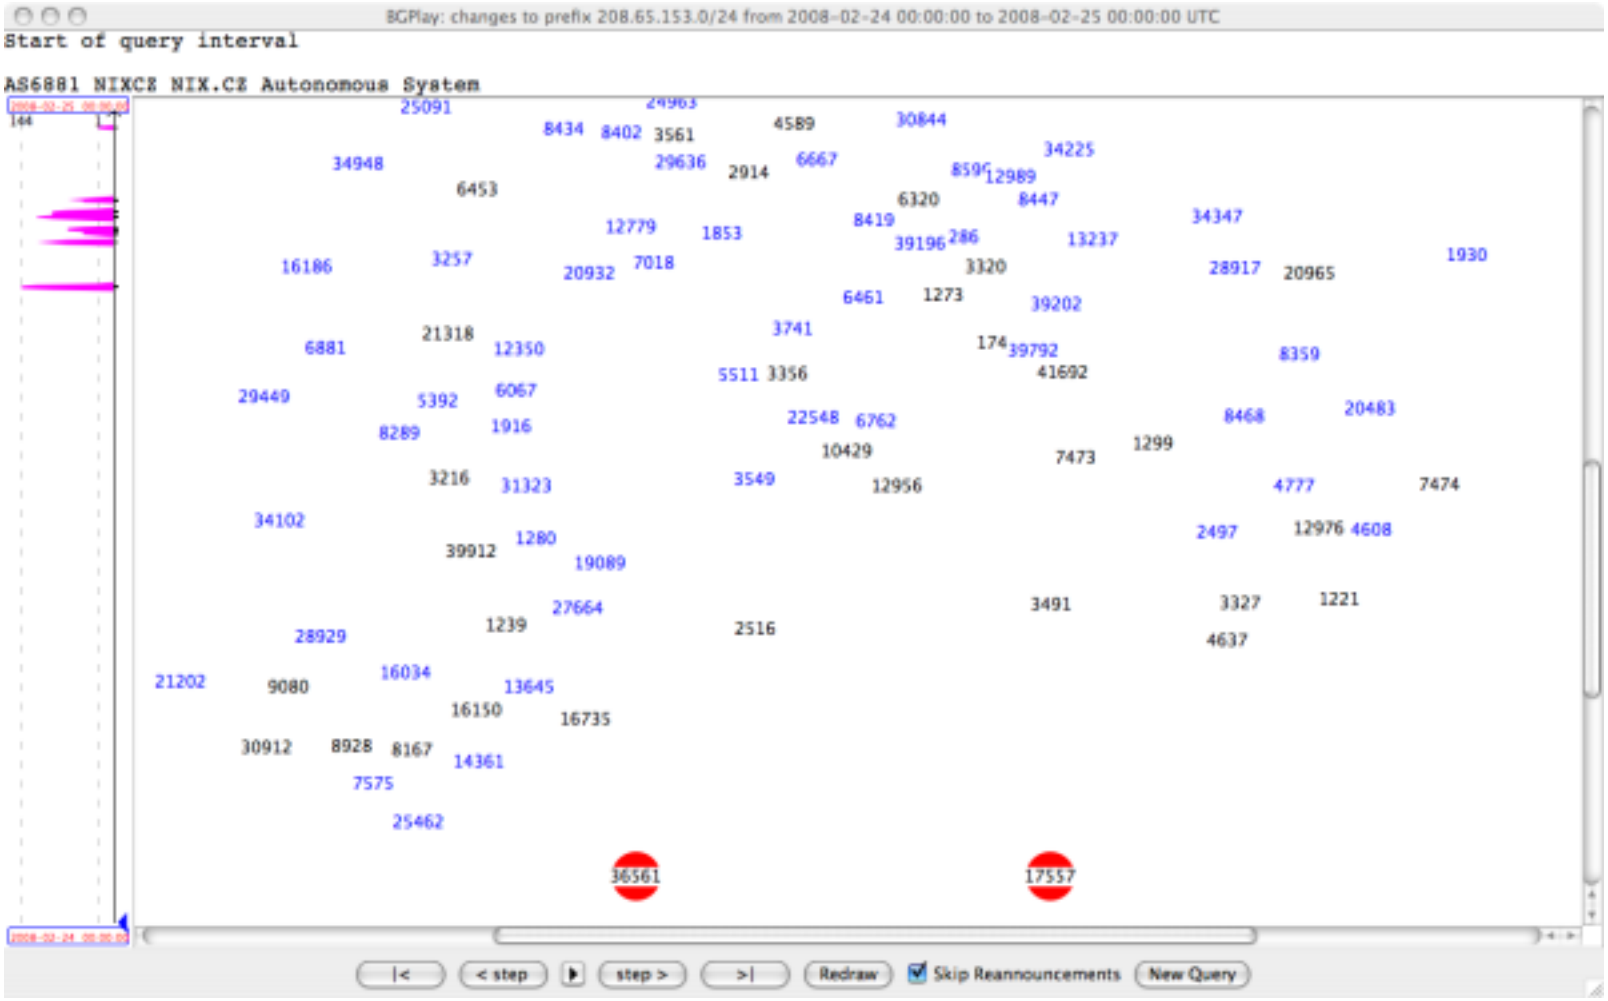

AS17557 (Pakistan Telecom) announces 208.65.153.0/24

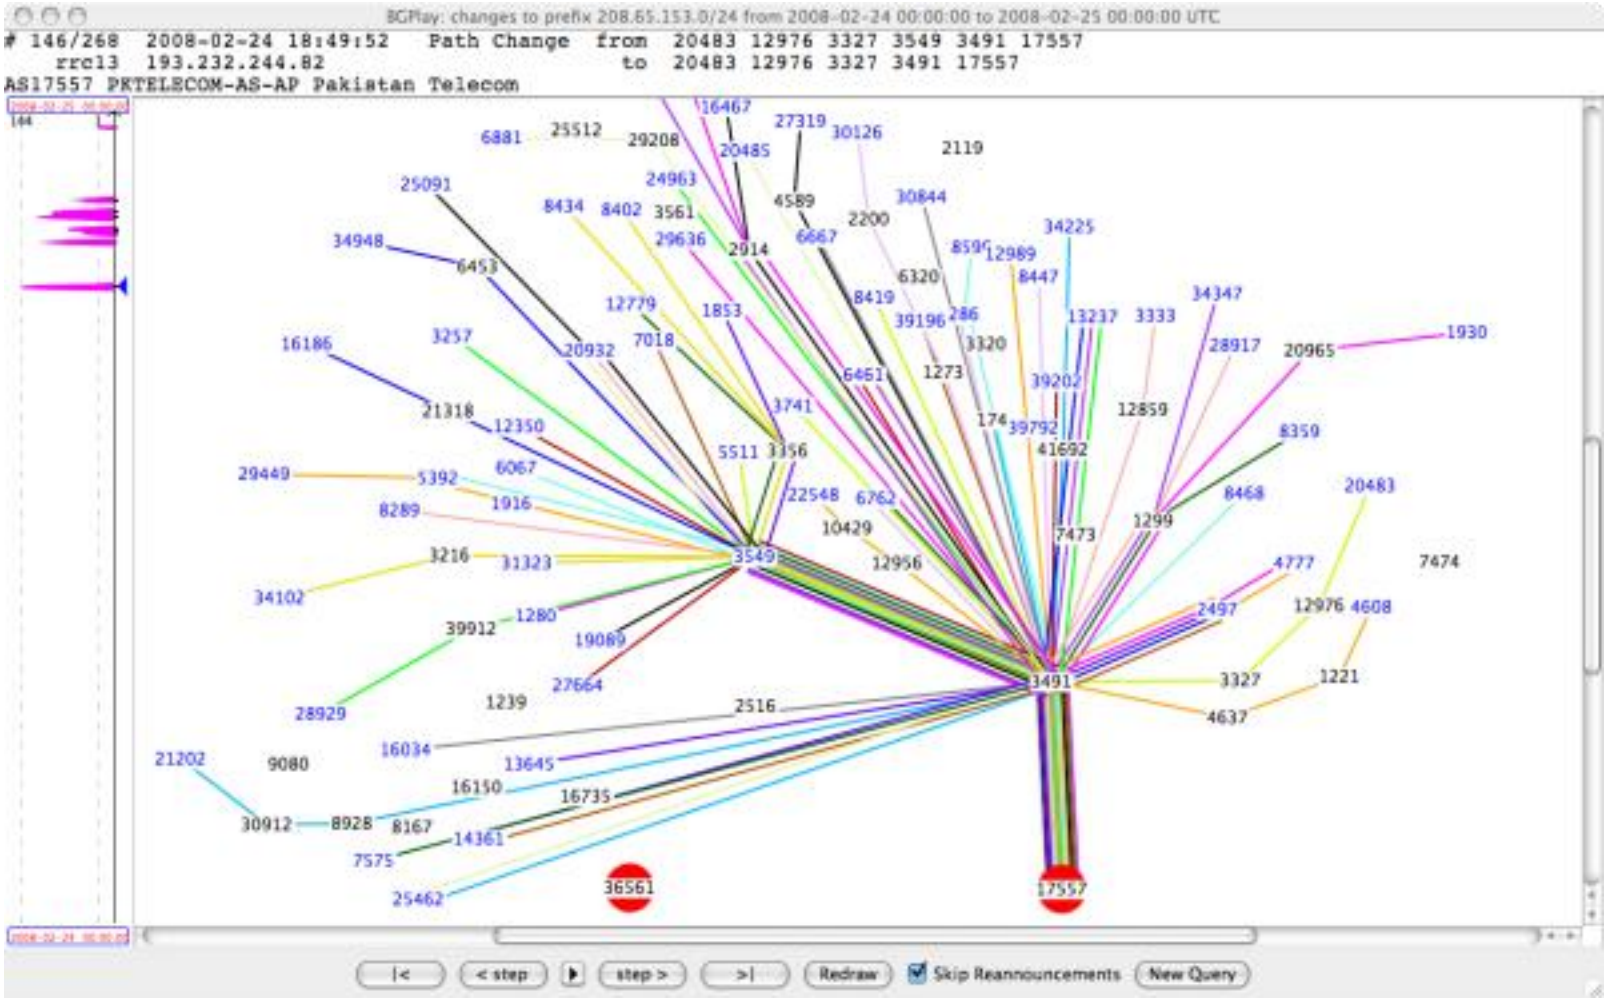

Other Notable Incidents

April 2010: China Telecom announced bogus paths to 50,000 IP prefixes

- Enabled traffic interception

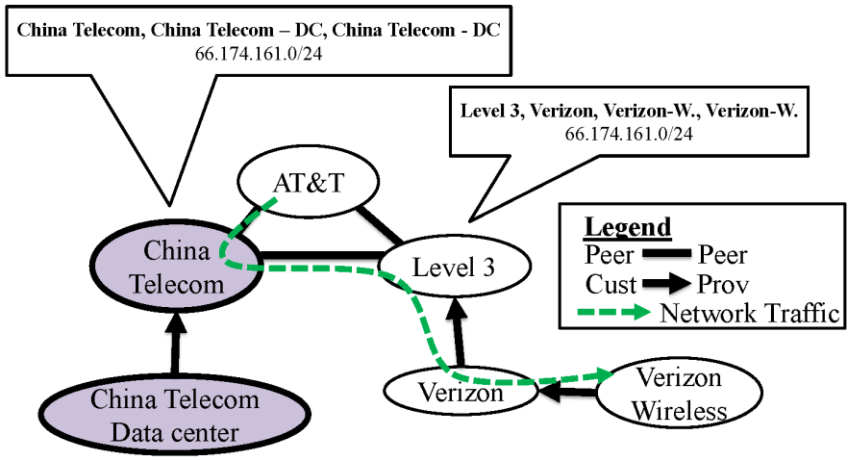

February 2014: hijacking of 51 networks (incl. Amazon, Digital Ocean, OVH)

- Miner connections were redirected to a attacker-controlled mining pool
- Attacker collected the miners' profit (est. $83,000 in 4 months)

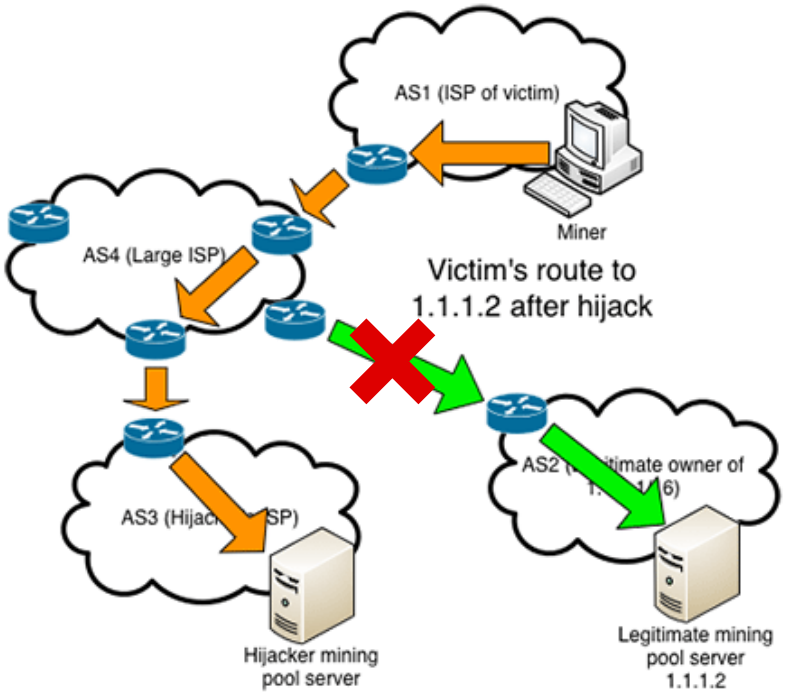

Securing BGP

Secure BGP (S-BGP)

- Each node signs its announcements

Resource Public Key Infrastructure (RPKI)

- Certified mapping from ASes to public keys and IP prefixes

Secure origin BGP (soBGP)

- Origin authentication + trusted database that guarantees that a path exists

Several other proposals… all facing many challenges

- No complete, accurate registry of prefix ownership
- Need a public-key infrastructure
- Cannot react rapidly to changes in connectivity
- Cost of cryptographic operations
- Not deployable incrementally

Domain Name Service

DNS maps domain names to IP addresses

- "Phonebook" for the internet

Hierarchically divided name space

- .edu ➔ stonybrook.edu ➔ cs.stonybrook.edu ➔ www.cs.stonybrook.edu

Not a one-to-one mapping

- Virtual hosting: many names to a single address
- Load balancing/fault tolerance: single name to many addresses

Generic "directory" for other host-related information

- PTR records: map IP addresses to names (reverse lookup)
- MX records: find the mail server of domain
- CNAME records: aliases for other names (not IP addresses)
- TXT records: associate arbitrary data with a host or other name

Primarily uses UDP for queries/responses (port 53)

- TCP sometimes used for long responses and zone transfers

Recursive Name Resolution

Hosts know at least one local DNS server

- Uses the hierarchy of zones and delegations to respond to queries for which it is not authoritative.
- Caches responses for future queries (TTL specified by owner)

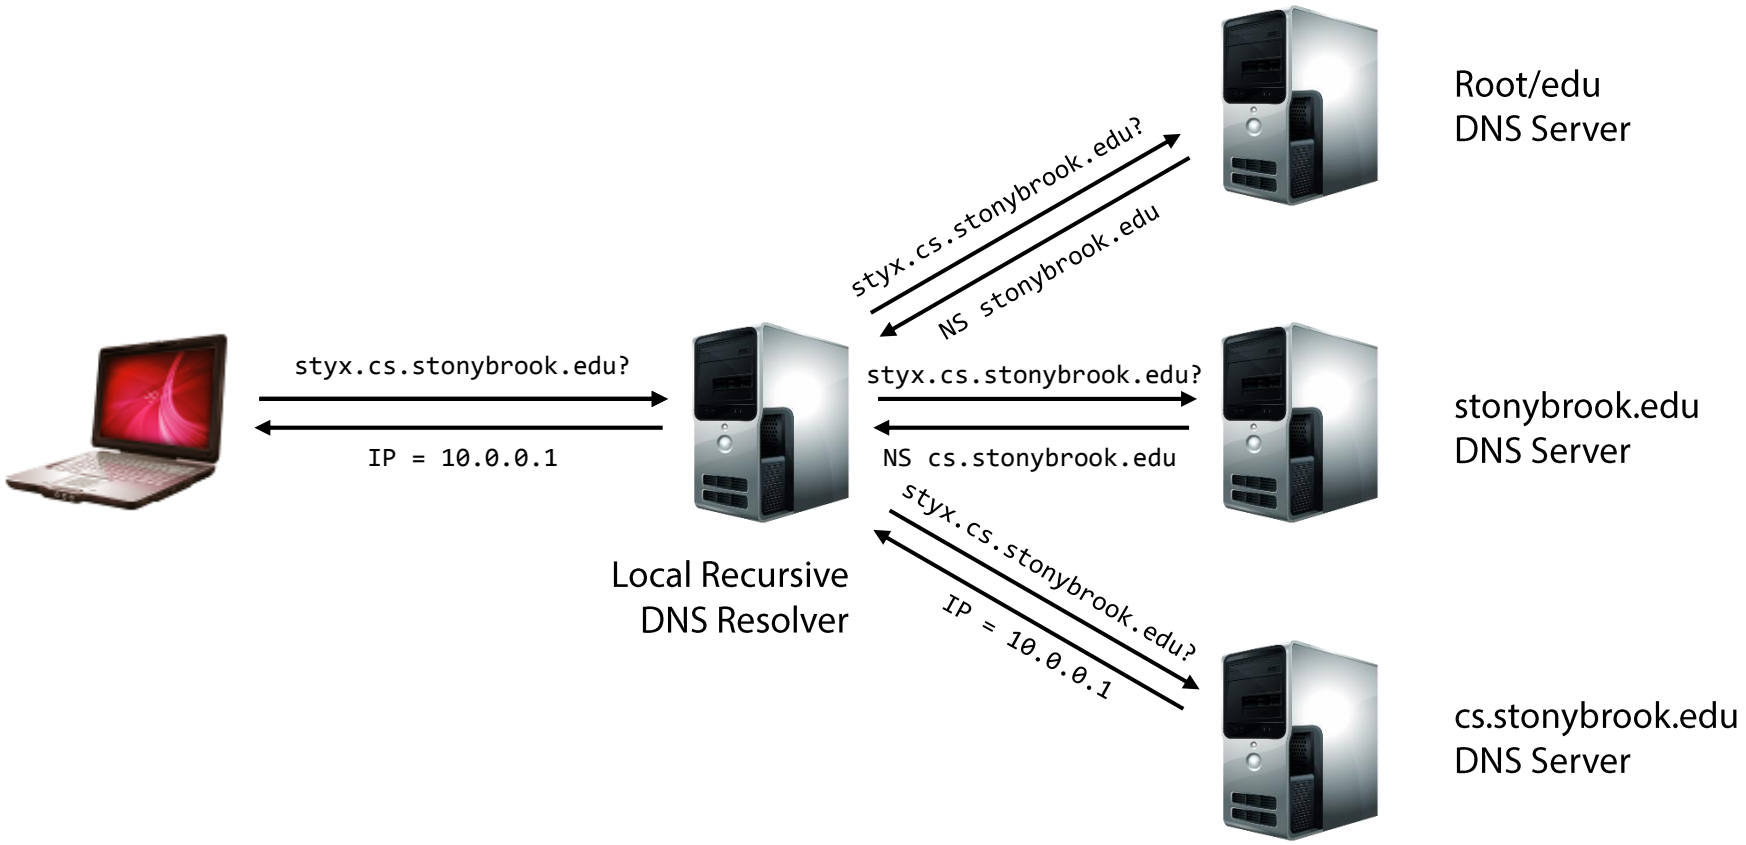

DNS Spoofing/Cache Poisoning

No authentication (reminds something?)

Responses can be spoofed!

- Point to a different address of the attacker's choosing
- Phishing, malware infection, …

Bellovin's cache contamination attacks (1990)

- Example: fake a PTR record for an attacker-controlled IP address to return a trusted hostname
- r-utilities perform name-based authentication (e.g., permit all hosts in .rhosts to rsh/rlogin in)
- The reverse lookup for the attacker's originating IP address when rsh/rlogin receives the connection will return a trusted name…
- Fix: cross-check the returned hostname by looking it up again

DNS Poisoning: Different Vantage Points

Off-path: attackers cannot observe any DNS queries and responses

- Blind packet injection: must guess the proper values in the response fields according to the query
- Race condition: forged response must arrive before the real one

On-path: attackers can passively observe the traffic (queries) and inject properly forged responses (MotS)

- Easy to mount in WiFi networks, by ISPs, …
- Race condition: forged response must arrive before the real one

In-path: attackers can block responses from reaching the victim, and inject forged ones instead (MitM)

- But then the attacker can do so much more…

DNS TXID

Synchronization mechanism between clients and servers

16-bit transaction identifier

- Randomly chosen for each query
- Response accepted only if TXIDs match
- Cached according to TTL (e.g., one day)

Attacker has to win a race

- Guess the correct TXID
- Response *src IP* and *dst port* should match query *dst IP* and *src port*

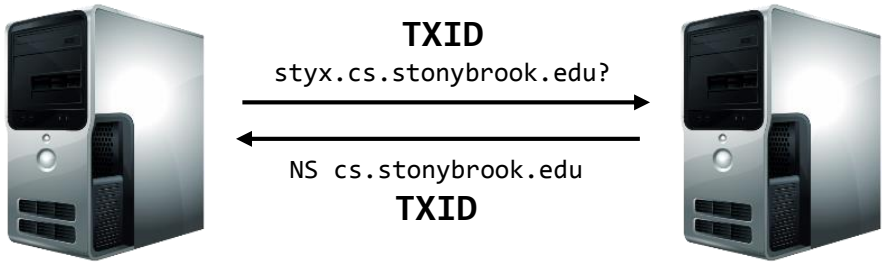

It's possible!

Kaminsky Attack

Attacker wants to take over example.com

Query the target resolver with any subdomain not in the cache

- Non-existent subdomains are fine: `foo1.example.com`
- Not affected by TTL (e.g., as would be the case for `www.example.com`)
- Causes the target resolver to query the authoritative server(s) for this domain

The attacker then floods the resolver with a large number of forged responses

- Each containing a different guess of the query's TXID
- **Fake referral**

```
;; ANSWER SECTION:
foo1.example.com.      120   IN A   10.0.0.10
;; AUTHORITY SECTION:
example.com.          86400  IN NS
ns1.example.com.
;; ADDITIONAL SECTION:
ns1.example.com.     604800  IN A   10.6.6.6
```

If the race is lost, repeat with a different subdomain!

Pharming

Mostly traffic redirection attacks at the client side

Malware can alter local DNS settings

- Change the system's (or the local router's) DNS server
- Add entries in /etc/hosts
- Example: DNSChanger: est. 4M infected computers, US$14M profit (FBI's "Operation Ghost Click")

Drive-by pharming

- A malicious web page contains JavaScript code that alters the local router's DNS server *from the inside LAN*

Dynamic pharming

- Switch mapping of bank.com between a malicious and a real IP
- First serve malicious script, then switch to the real site ➔ same origin policy is bypassed

Other DNS Attacks

Attacking registrars

- Social engineering, stolen credentials, …

DoS on root/critical servers or other targets

- DNS amplification attacks

Typosquatting/Registering expired domains

- Phishing – www.paypa1.com
- Hijack scripts hosted on expired domains still in use by other web pages

Zone transfers

- Reconnaissance

Server bugs

- System compromise

Censorship

DNSSEC

Goal: enable authentication and ensure the integrity of DNS requests and responses

- Non-goals: availability, confidentiality

Cryptographically signed resource records

- Resolvers can verify the signature

Two new resource types:

DNSKEY: creates a hierarchy of trust within each zone

- Name = Zone domain name
- Value = Public key for the zone

RRSIG: Prevents hijacking and spoofing

- Name = (type, name) tuple, i.e. the query itself
- Value = Cryptographic signature of the query results

Not a complete solution

- Enables DoS amplification/CPU exhaustion attacks
- Forgery of delegation records still possible
- No "last mile" protection