CSE508 Network Security

Stony Brook University

2024-02-13

Core Protocols: BGP

Michalis Polychronakis

Stony Brook University

IPv4 Addressing and Forwarding

Packets are routed based on their destination IP address

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

Table lookup for each packet in a routing table

32-bit addresses, 2^{32} possibilities ➔ impractical to maintain 2^{32} entries

Solution: hierarchical address scheme

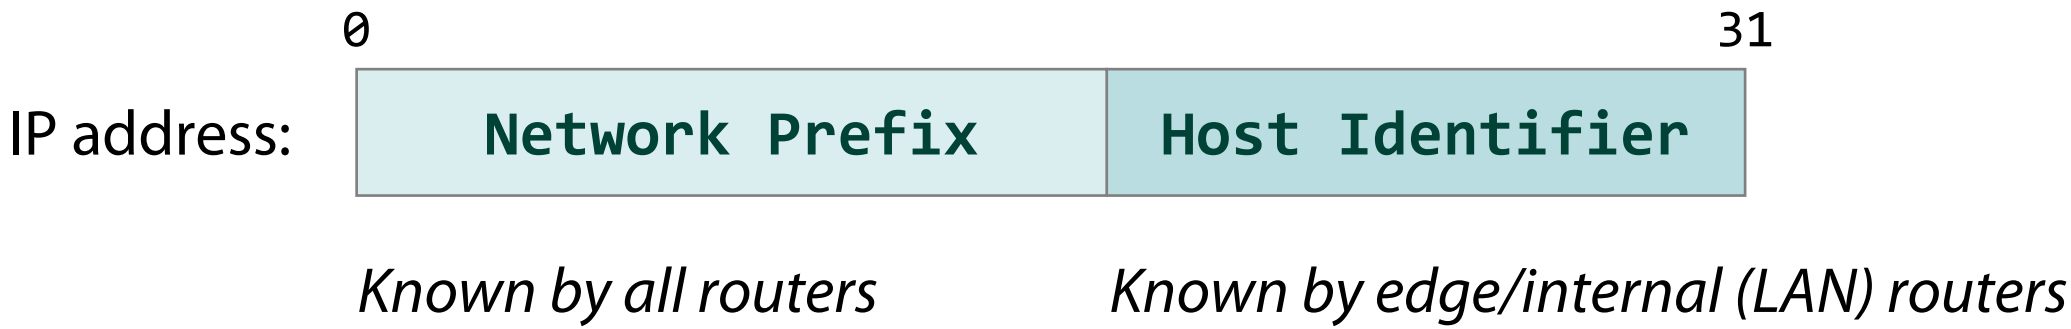

IPv4 Address Classes

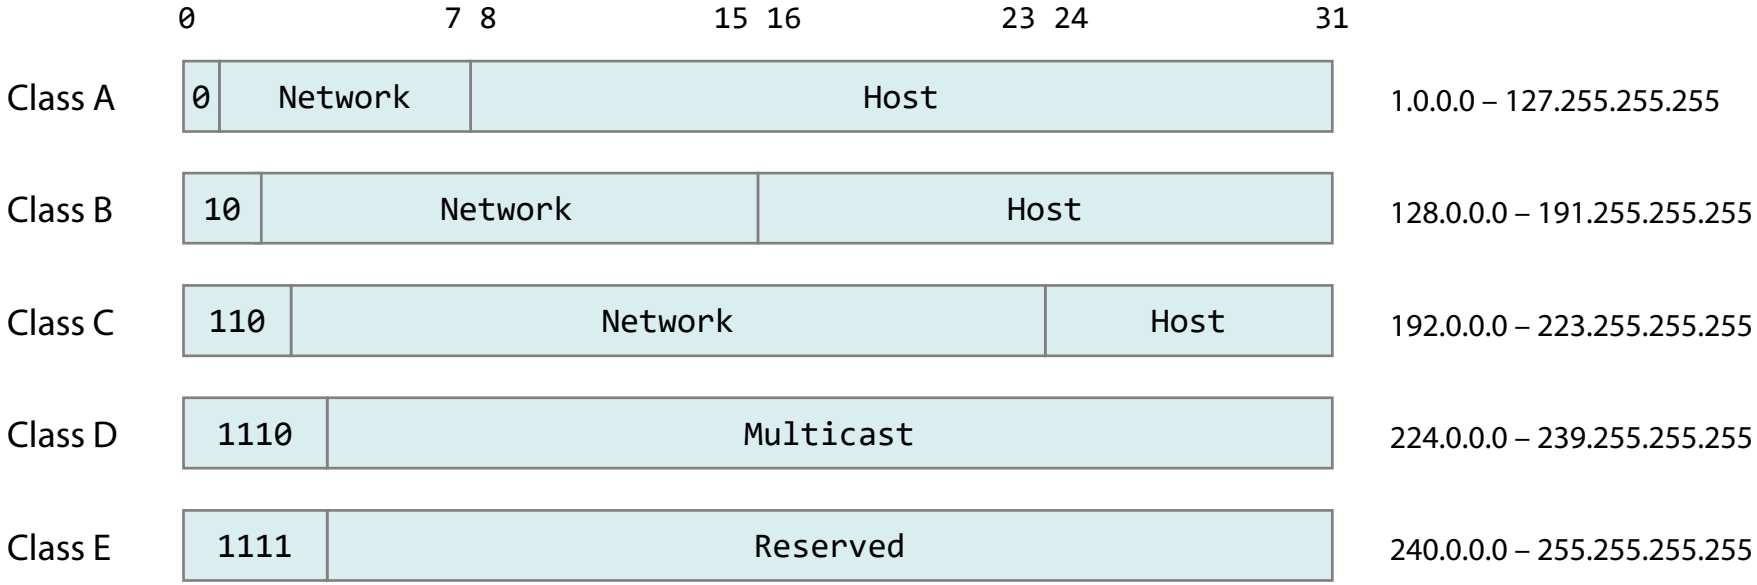

Classless Inter-Domain Routing (CIDR) was introduced in 1993

- Replaced the *classful* A/B/C network addressing architecture
- IP addresses are now associated with a *subnet mask*
- Allocations to ISPs and end users can be made on any address-bit boundary

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

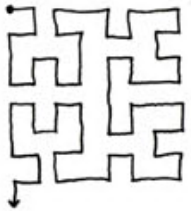

= UNALLOCATED BLOCK

No green patches after 2011…

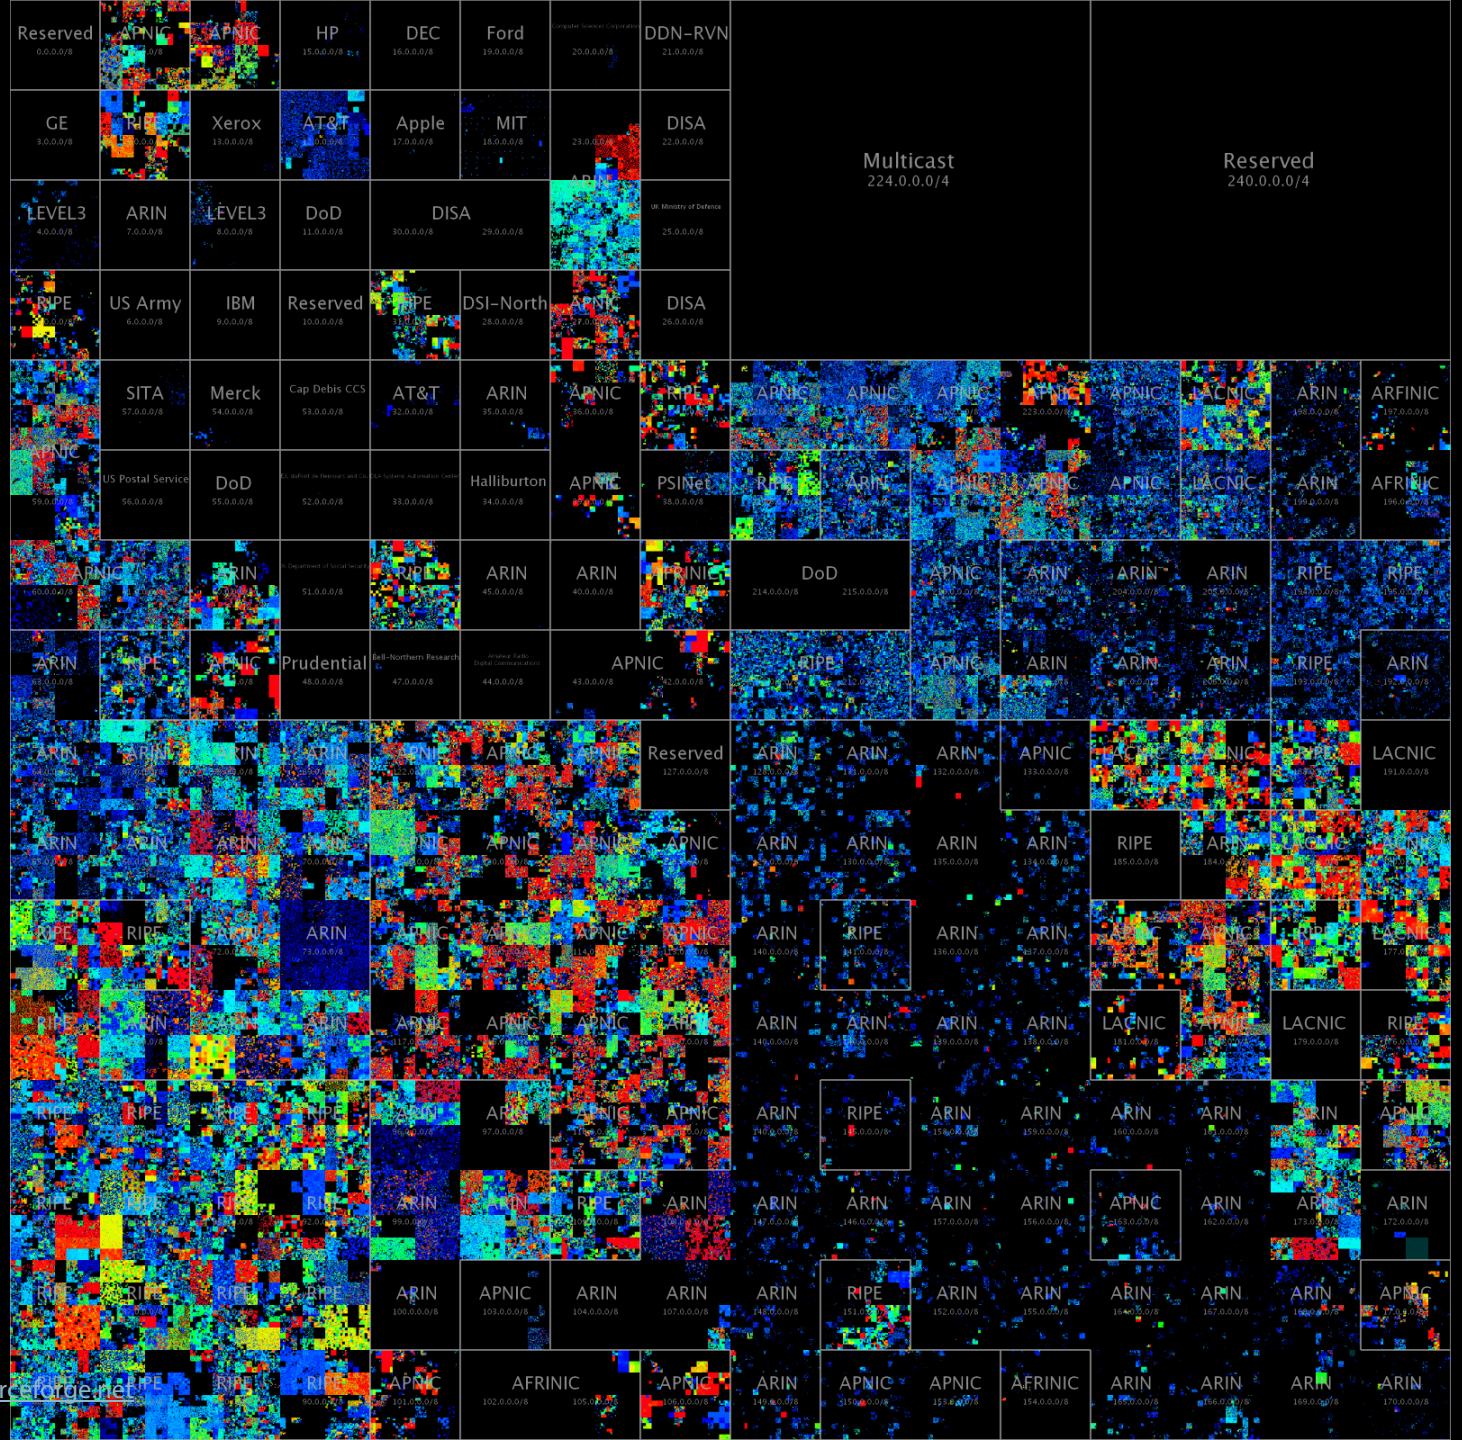

IPv4 Census Map

June – October 2012

Utilization

100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%

Prefix Sizes

= /8

= /12

= /16

= /20

= /24

420 Million hosts that responded to ICMP Ping at least 2 times between June and October 2012
Source: Carna Botnet

http://census2012.sourceforge.net

Autonomous Systems

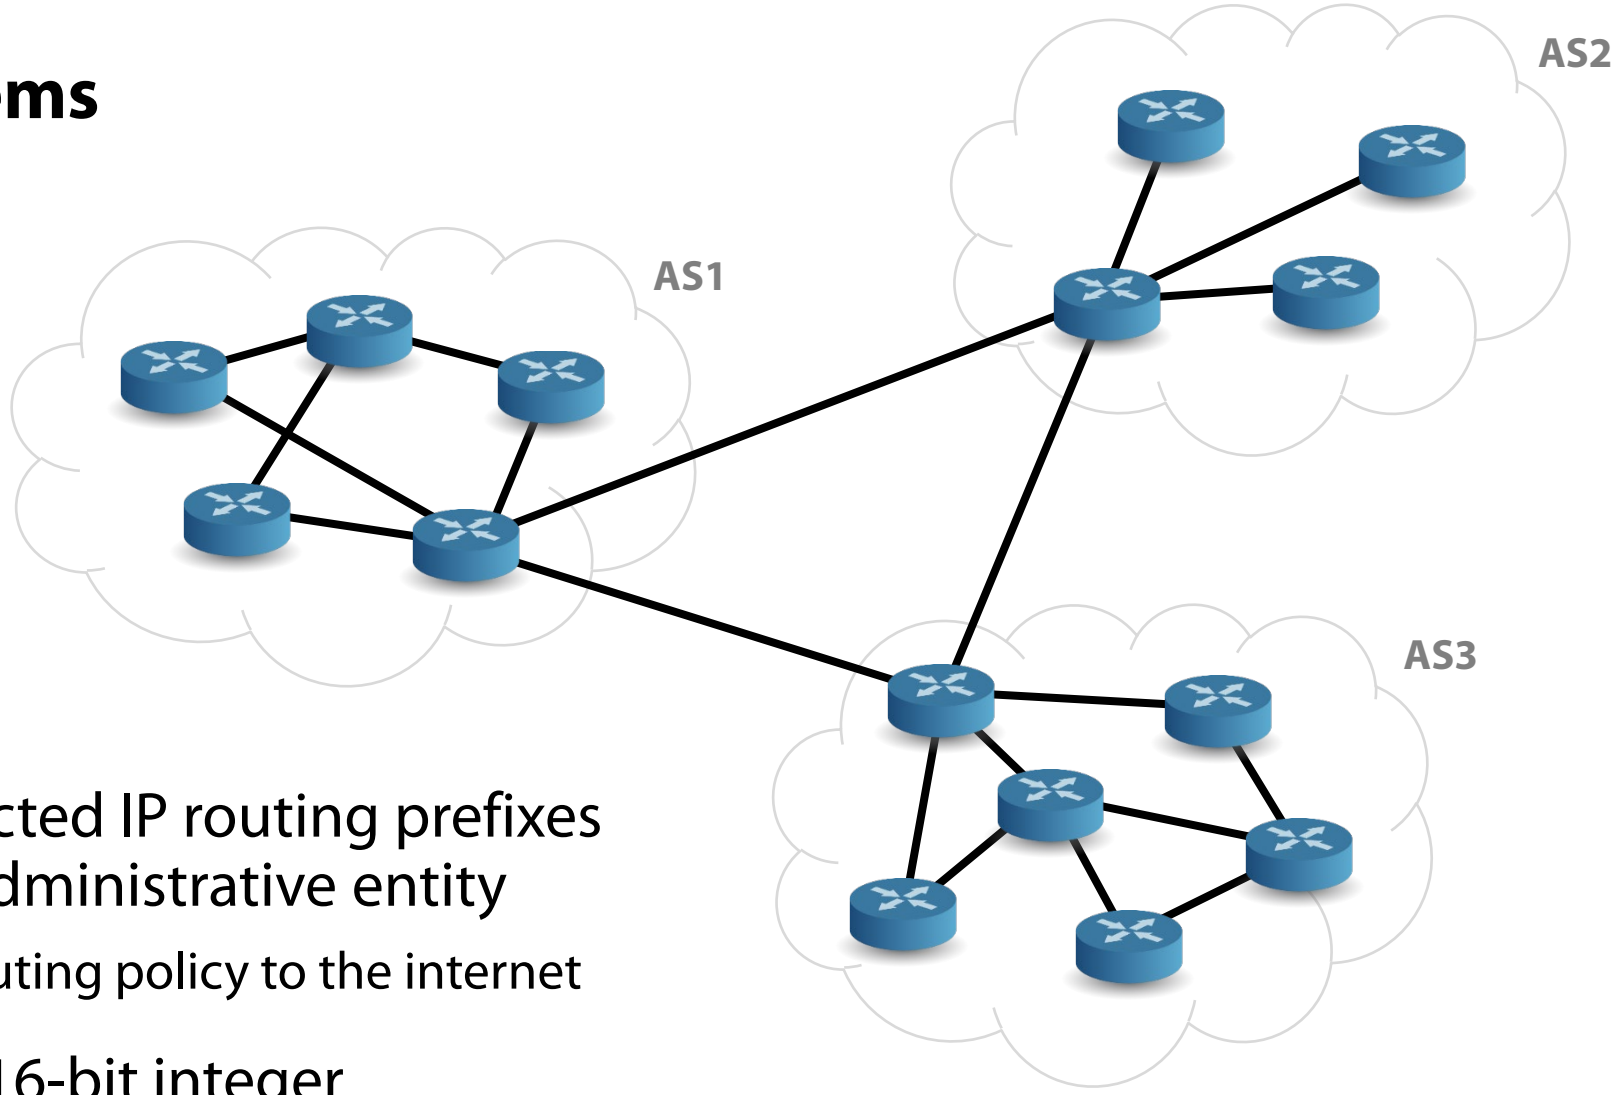

AS: collection of connected IP routing prefixes belonging to a single administrative entity

- Presents a common routing policy to the internet

AS number defined as 16-bit integer

- 99,857 ASNs as of February 2021, assigned by IANA

WHOIS

Query–response protocol [RFC 3912] used for querying public databases that store an Internet resource's registered users/assignees

- Domain names
- IP addresses
- Autonomous systems

Additional information

- Registrant name and contact details (not always available due to privacy concerns)
- Nameservers
- Dates related to registration

```
mikepo@konami:~> nslookup hexlab.cs.stonybrook.edu

Name:   hexlab.cs.stonybrook.edu
Address: 130.245.42.42

mikepo@konami:~> whois 130.245.42.42

NetRange:       130.245.0.0 - 130.245.255.255
CIDR:           130.245.0.0/16
NetName:        SBU-130-245-0-0-16
NetHandle:      NET-130-245-0-0-1
Parent:         NET130 (NET-130-0-0-0-0)
NetType:        Direct Allocation
OriginAS:
Organization:   State University of New York at Stony Brook (SUNYASB-Z)
RegDate:        1988-10-25
Updated:        2023-10-16
Ref:            https://rdap.arin.net/registry/ip/130.245.0.0

mikepo@konami:~> whois -h whois.cymru.com 130.245.42.42

AS      | IP               | AS Name
5719    | 130.245.42.42    | SUNYSB, US
```

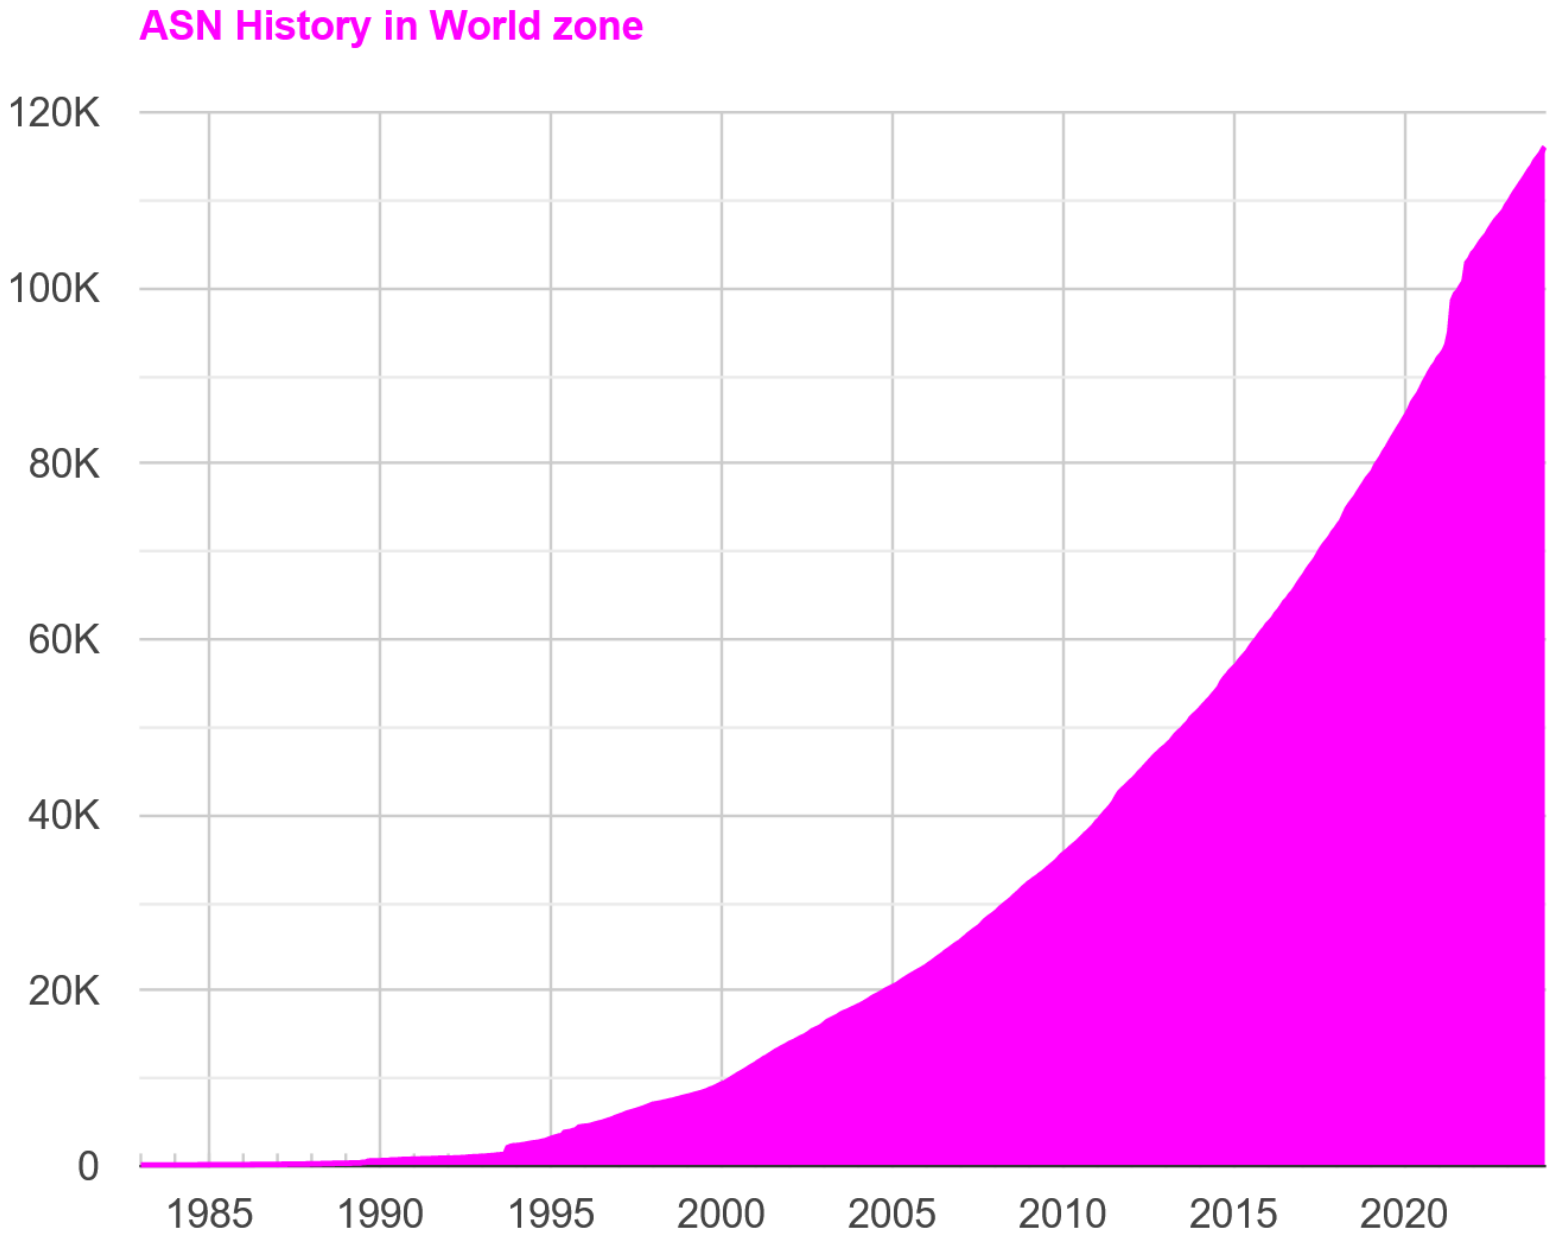

https://www-public.imtbs-tsp.eu/~maigron/rir-stats/rir-delegations/world/world-asn-by-number.html

ASN Statistics by country in World zone

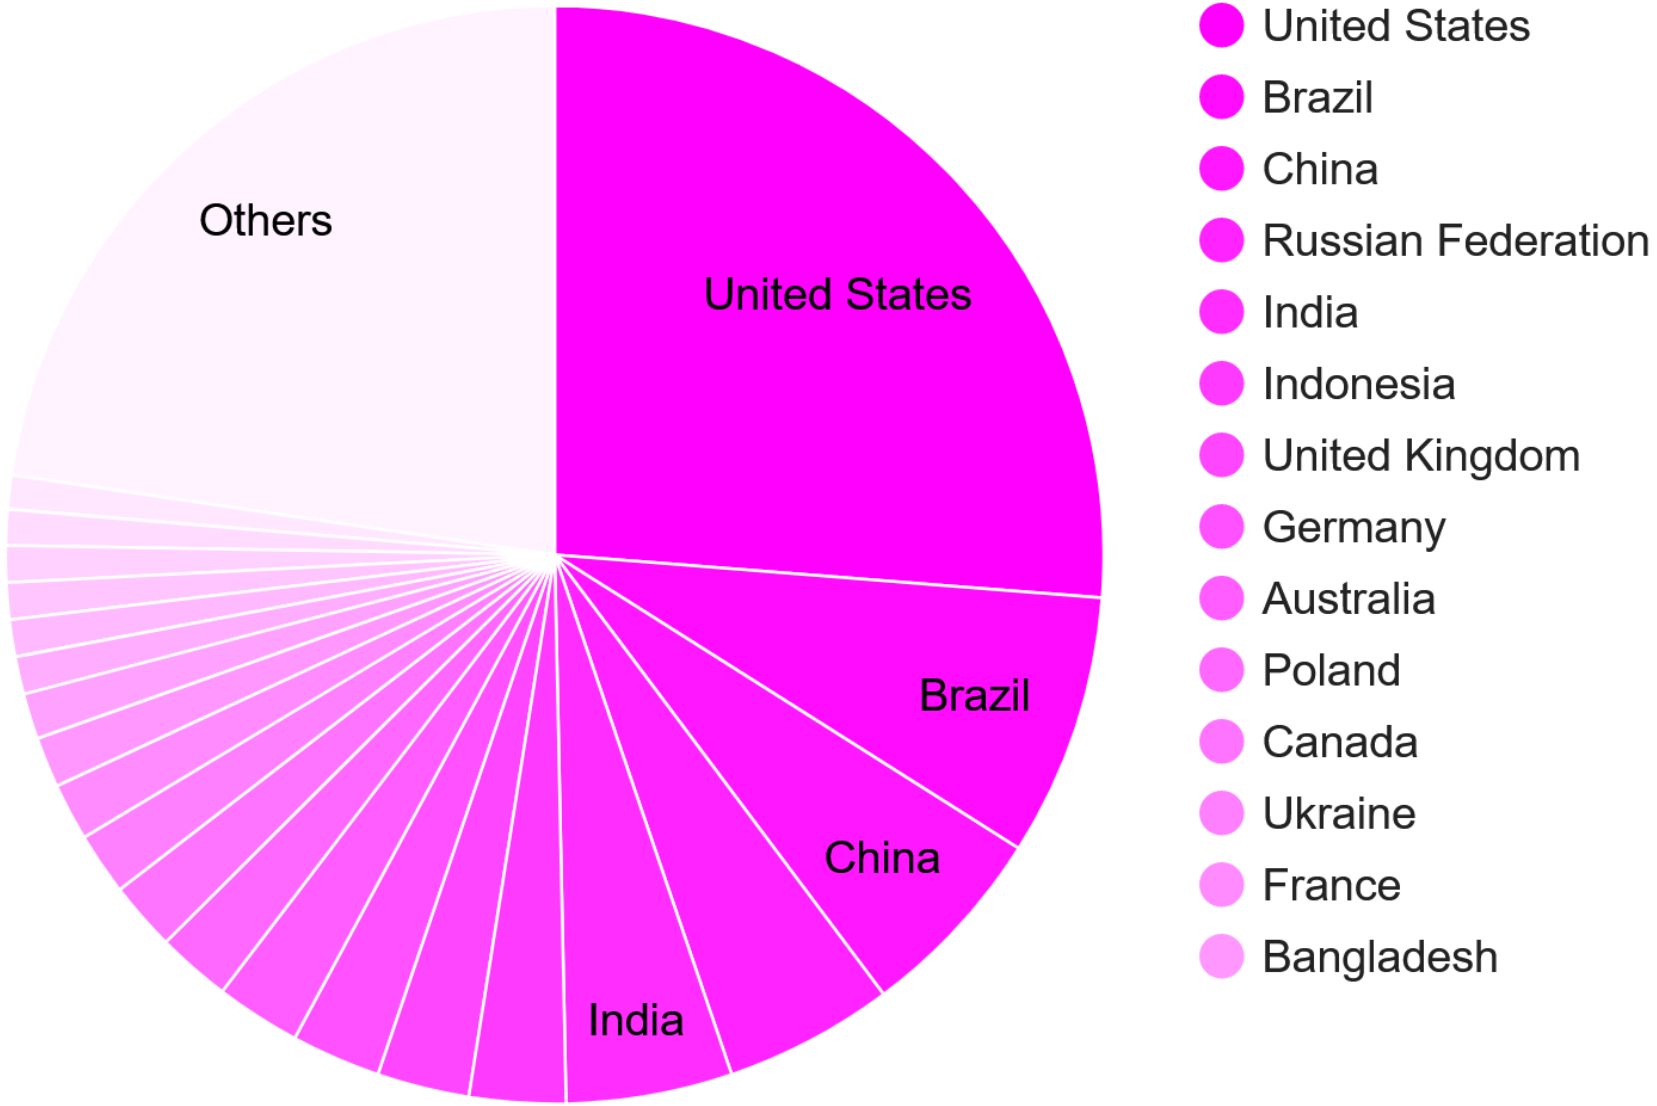

https://www-public.imtbs-tsp.eu/~maigron/rir-stats/rir-delegations/world/world-asn-by-number.html

Map of the internet, 1982

Ovals: sites/networks
Rectangles: routers

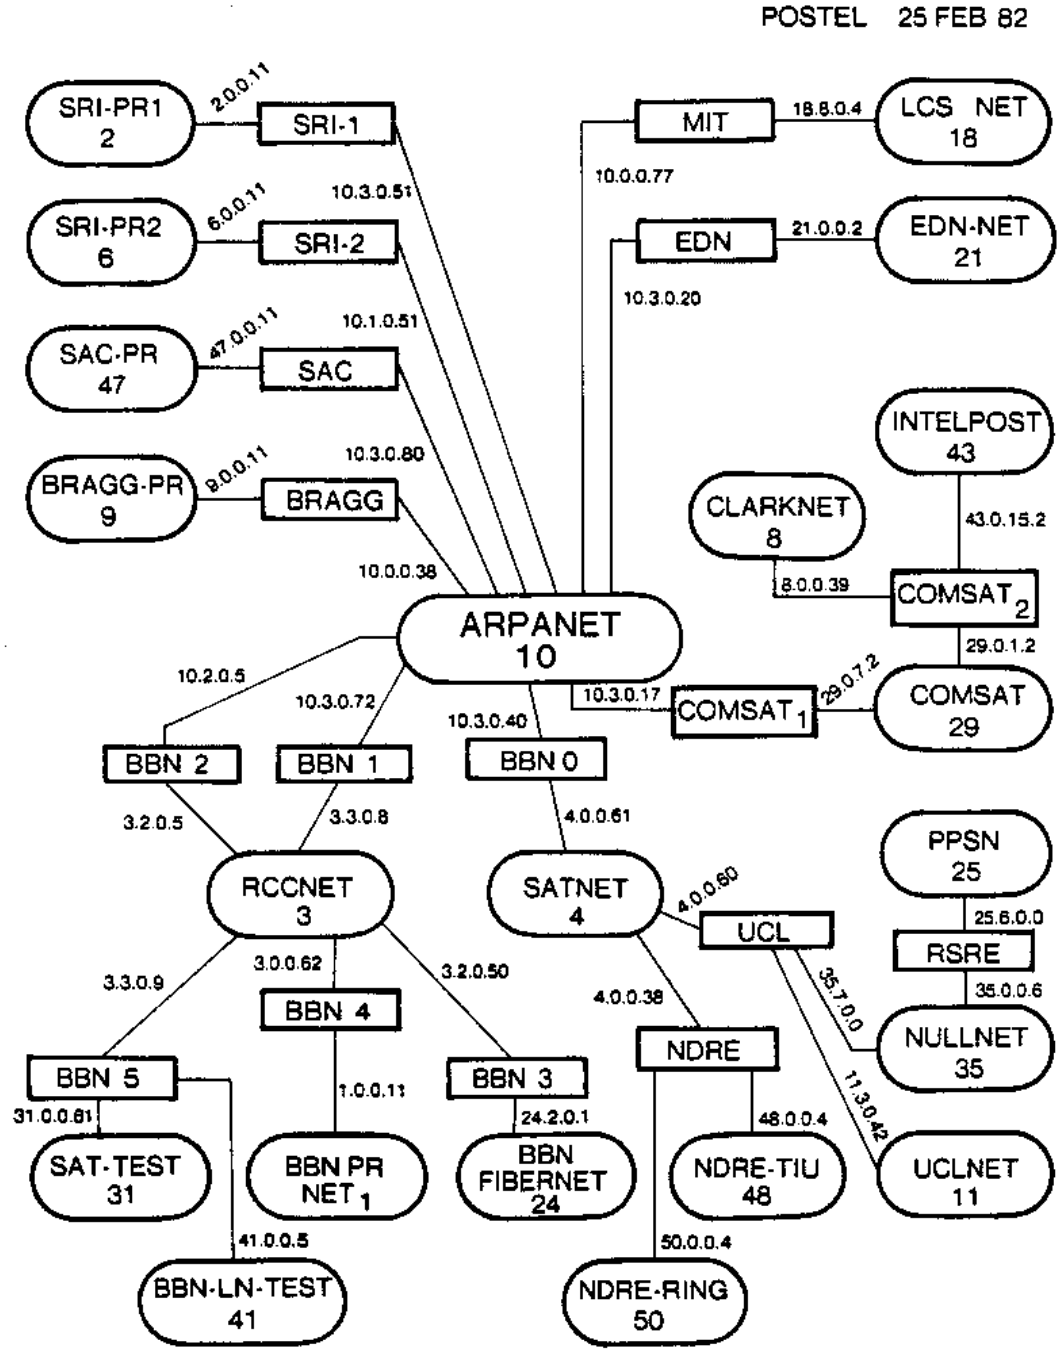

CAIDA's IPv4 AS Core
AS-level Internet Graph

Skitter
January 2000

220,533 IP addresses

5,107 ASes

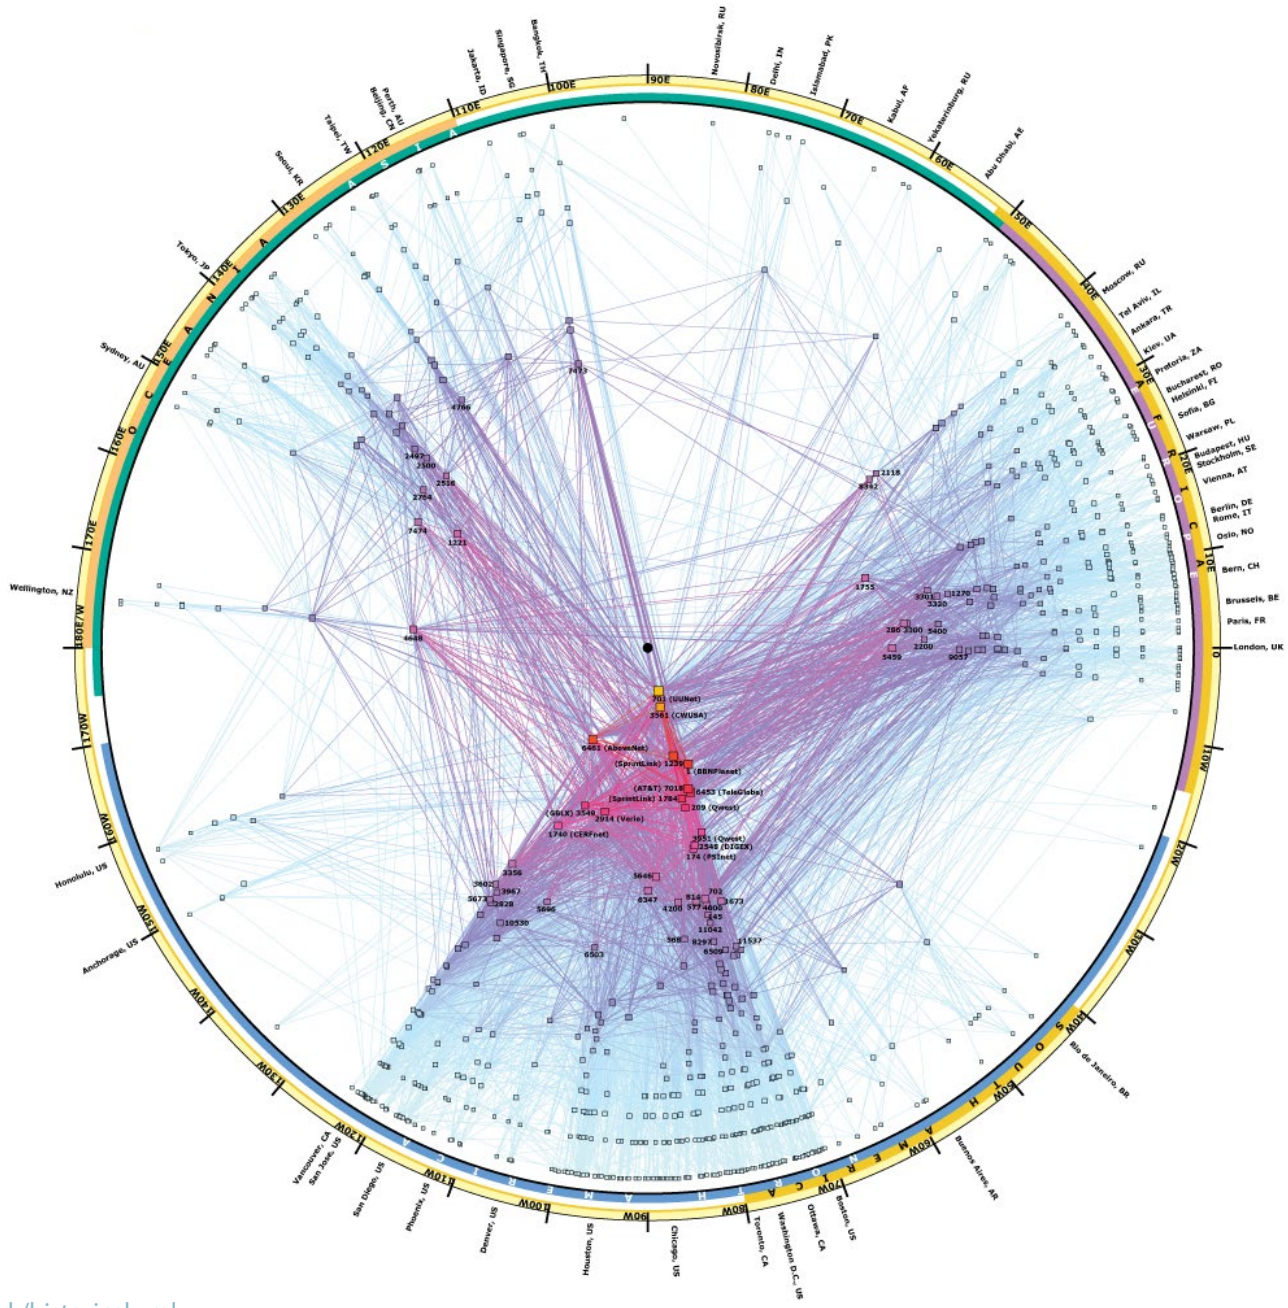

CAIDA's IPv4 AS Core
AS-level Internet Graph

Archipelago
August 2010

16,802,061 IP addresses

26,702 ASes

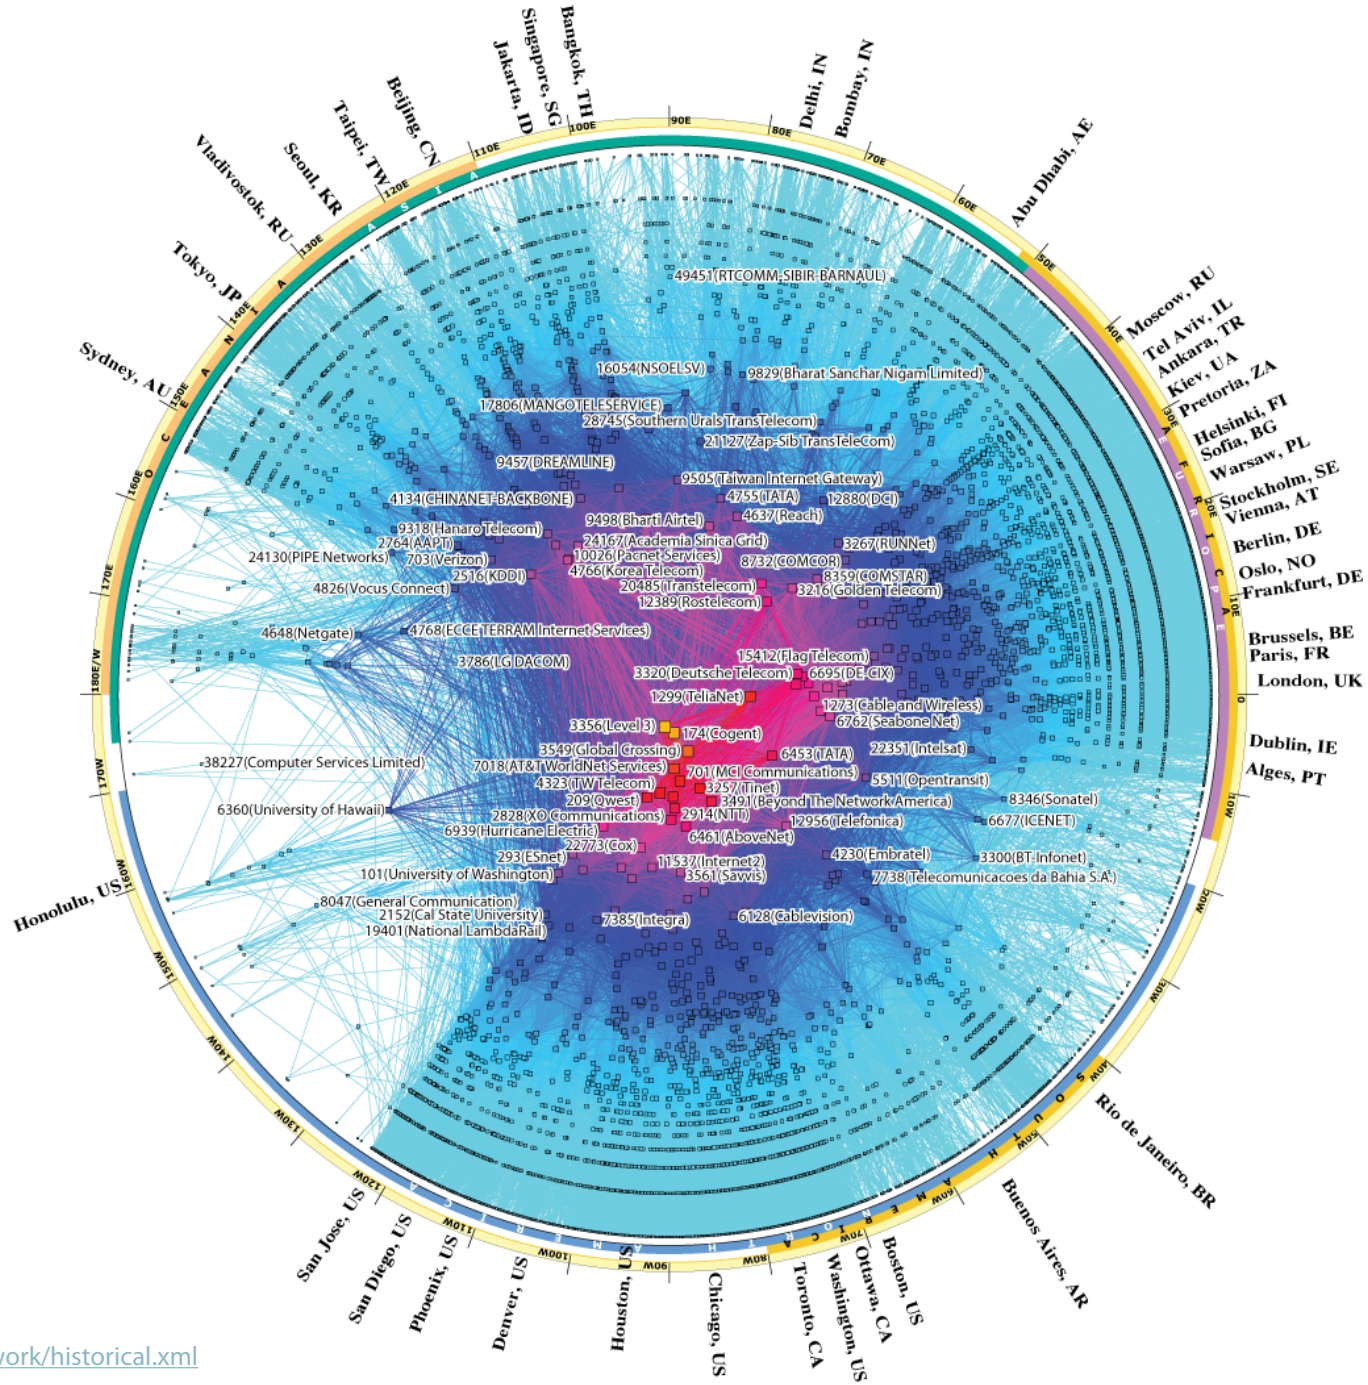

CAIDA's IPv4 AS Core
AS-level Internet Graph

Archipelago
February 2017

50 million IP addresses

47,610 ASes

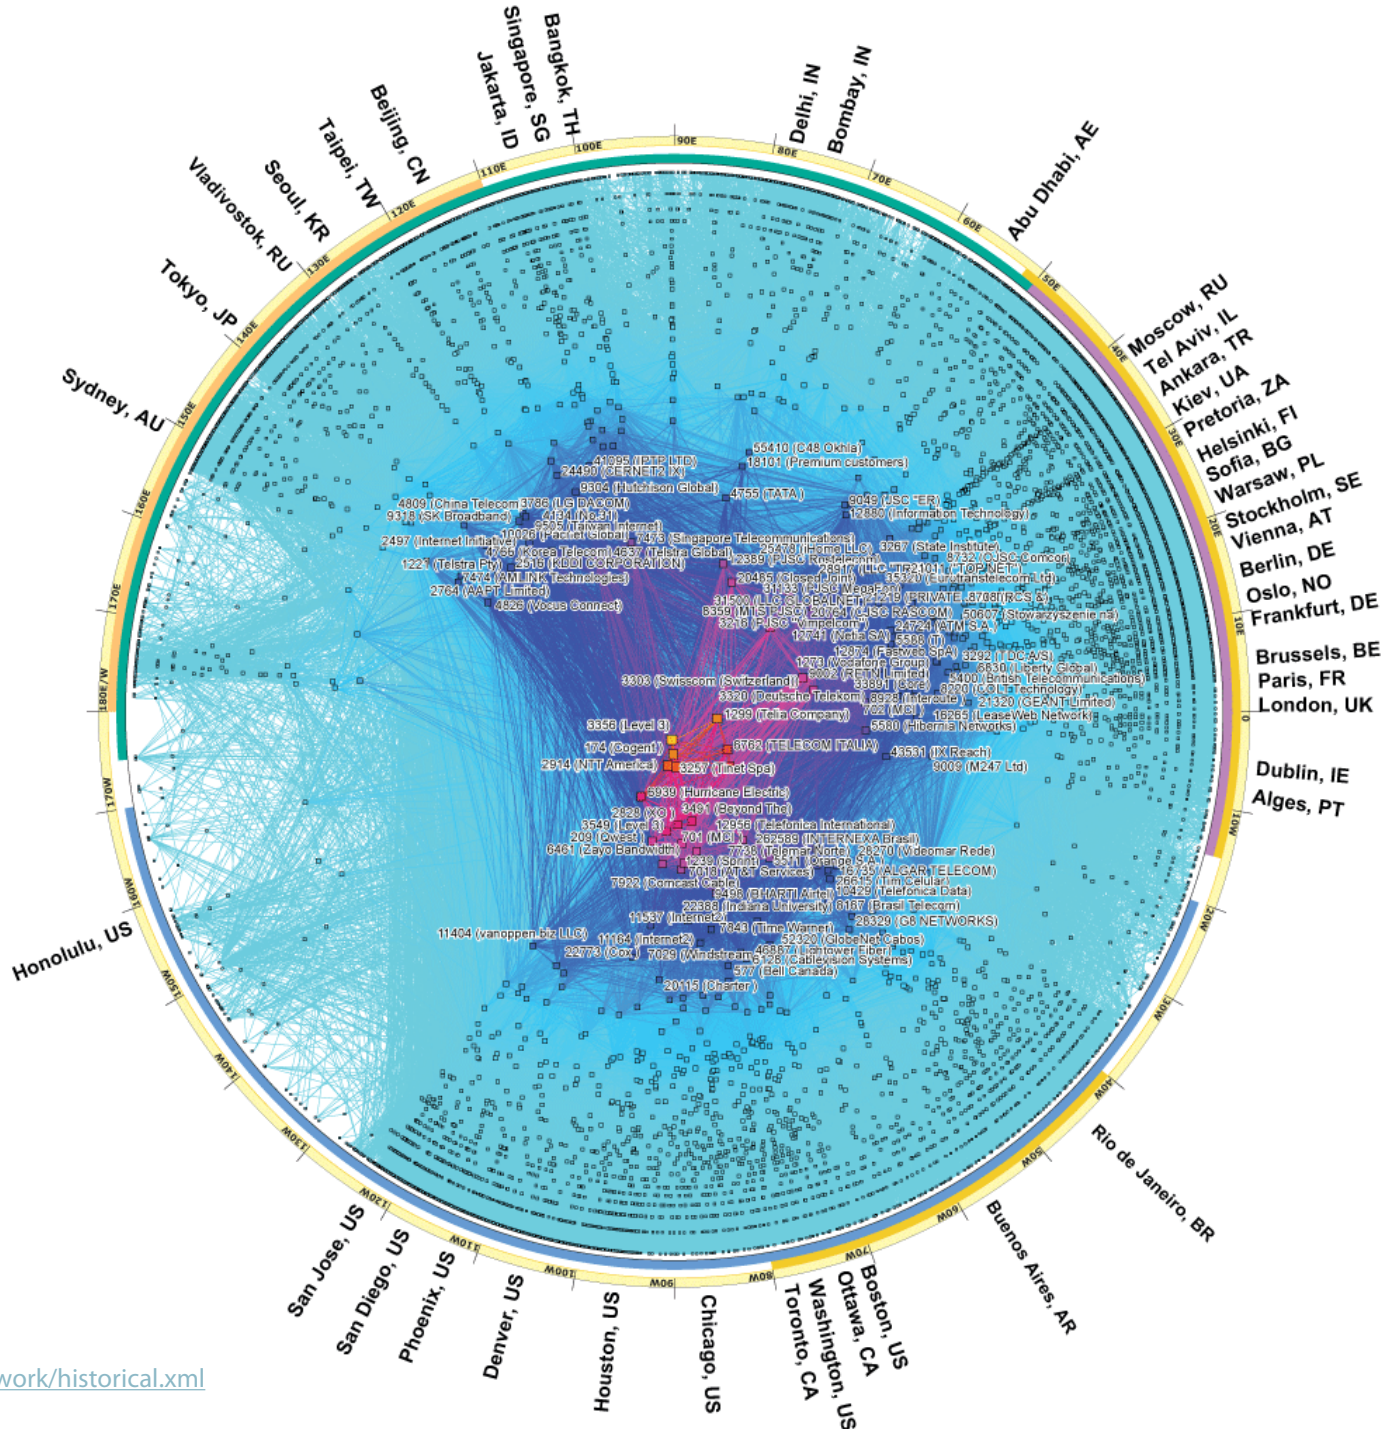

CAIDA's IPv4 AS Core
AS-level Internet Graph

Archipelago
January 2020

64 million IP addresses

61,290 ASes

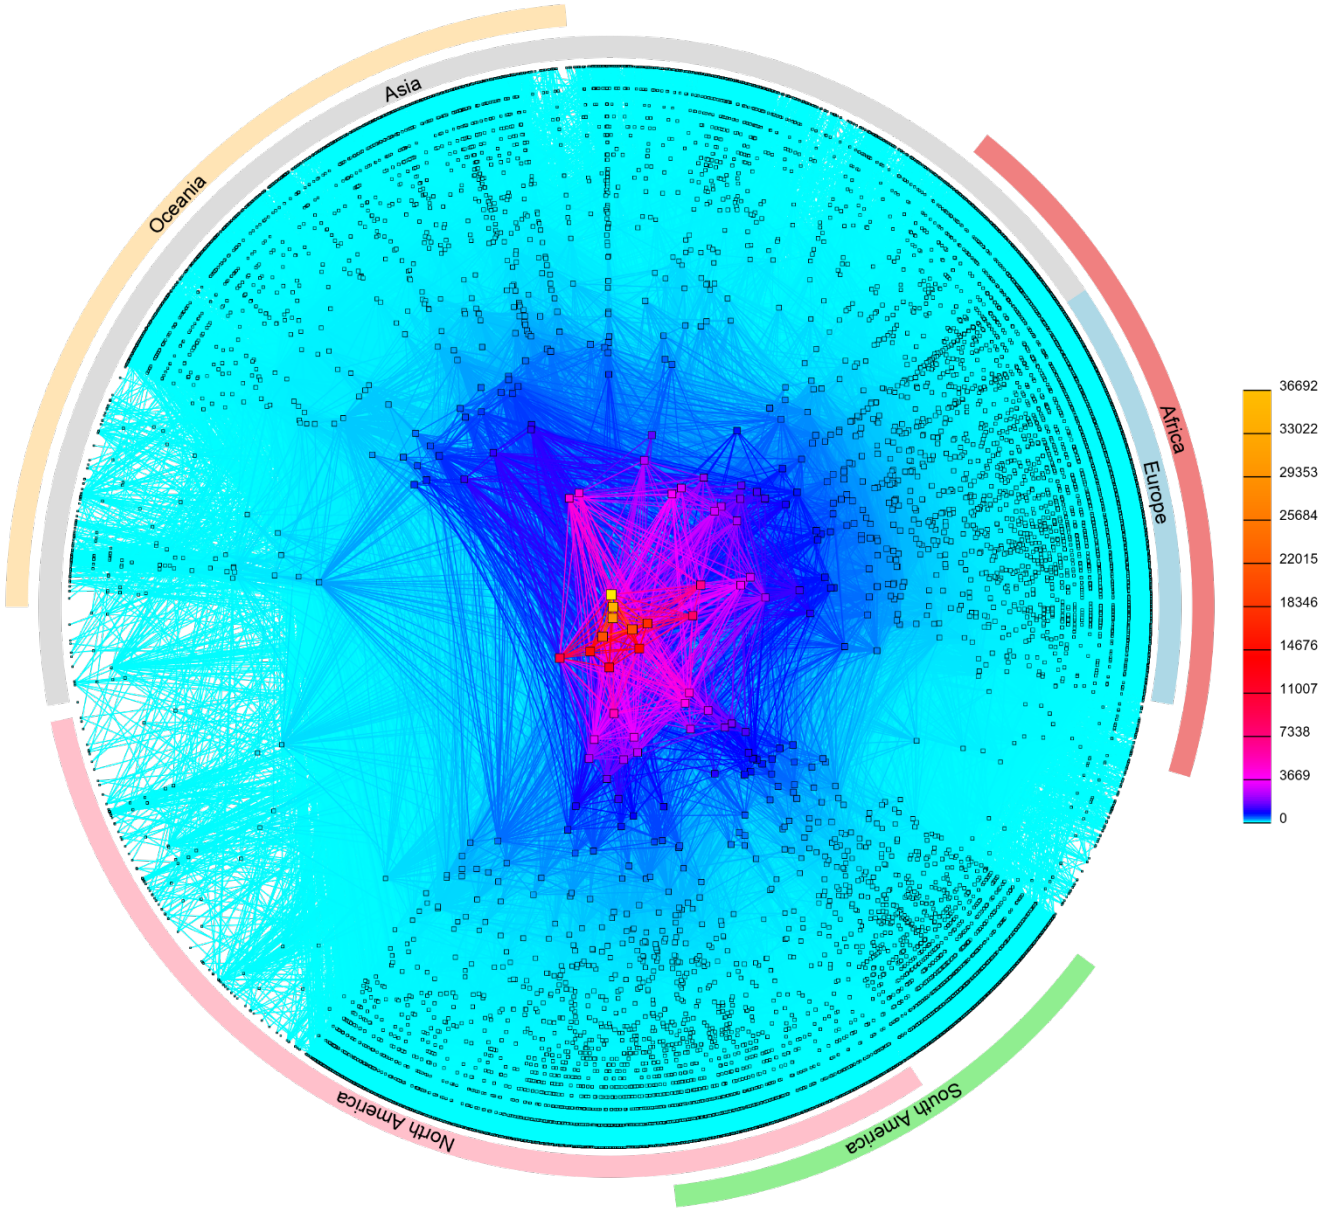

Internet Routing

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

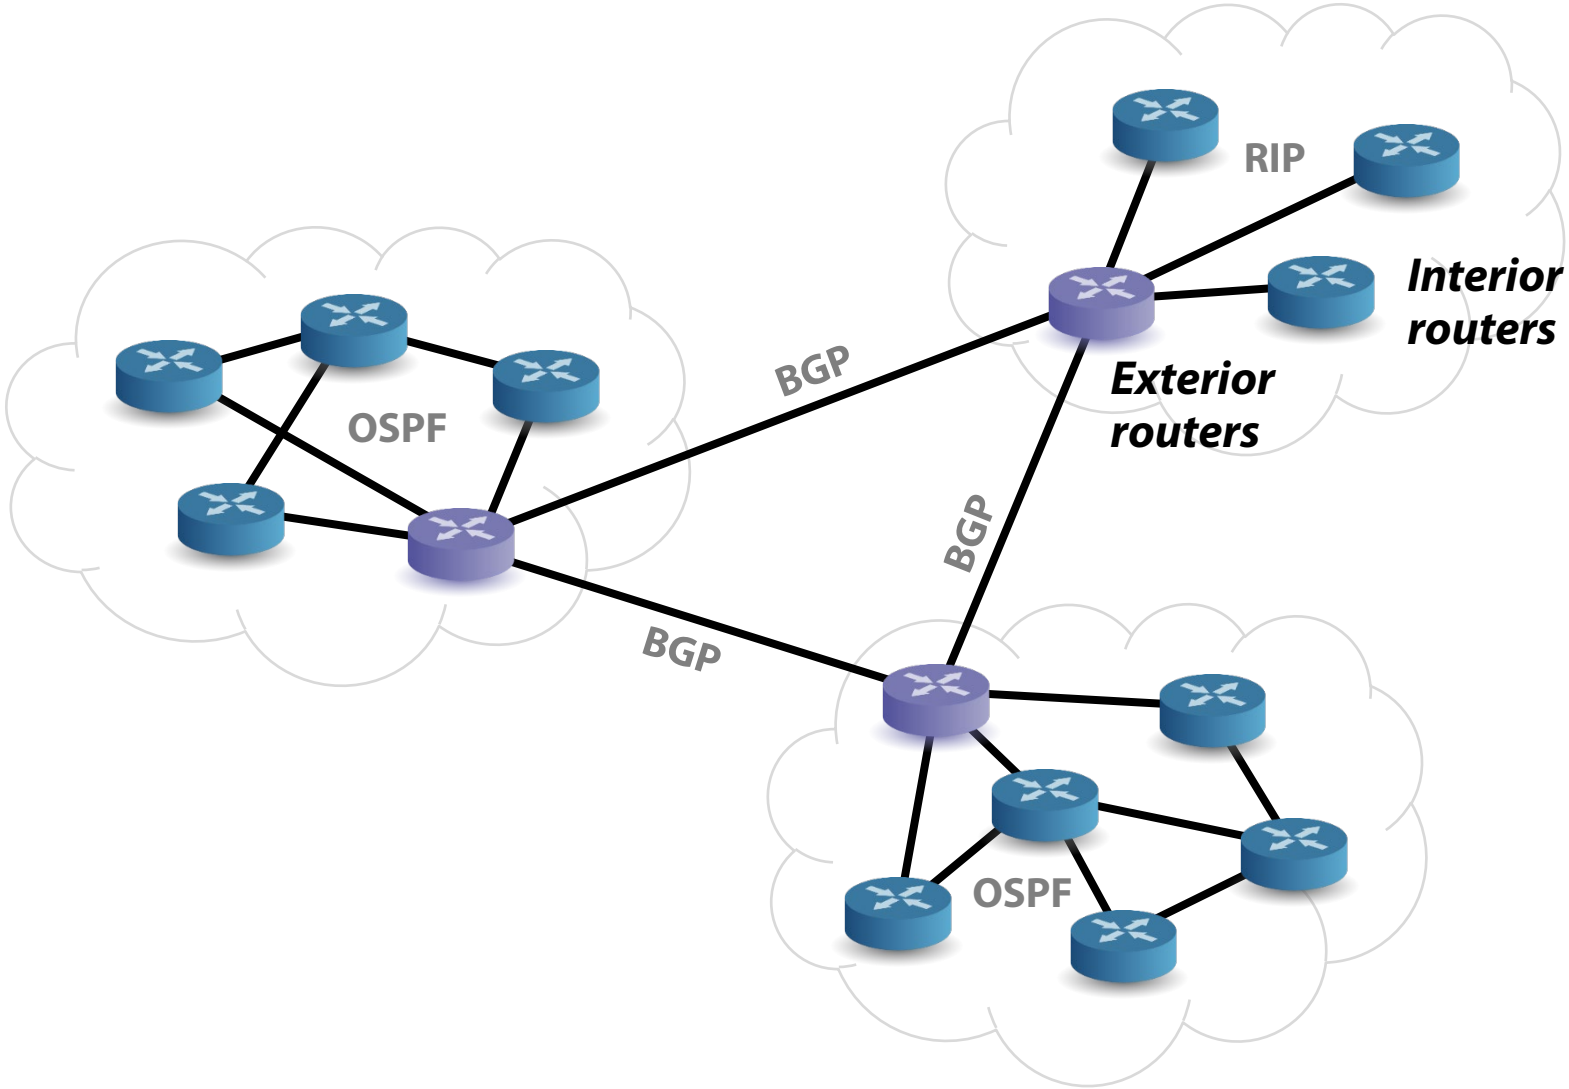

BGP (Border Gateway Protocol)

The de facto standard inter-AS routing protocol in today's Internet

- BGP is what enables subnets to advertise their existence to the rest of the Internet

Main goals:

- Obtain subnet reachability information from neighboring ASs
- Propagate the reachability information to all internal routers
- Determine "good" routes to subnets based on the obtained reachability information and the policies of the involved ASes

Path-vector routing protocol

- Maintains path information that is updated dynamically
- Makes routing decisions based on paths, network policies, or rules configured by network administrators

Root Causes of 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

- BGP is a complex protocol, with complex interactions

Attackers can lie to other routers

Routing Attacks

Blackholing

False route advertisements to attract and drop traffic

Redirection

Force some or all traffic to take a different network path ➔ sniffing, interception (MitM), flooding/congestion

Instability

Frequent advertisements and withdrawals, or increased BGP traffic to cause connectivity disruption

How?

Misconfigurations, insider attacks, compromised routers, BGP traffic manipulation, …

Prefix Hijacking

Announce someone else's prefix

Victim prefers the *shortest* path

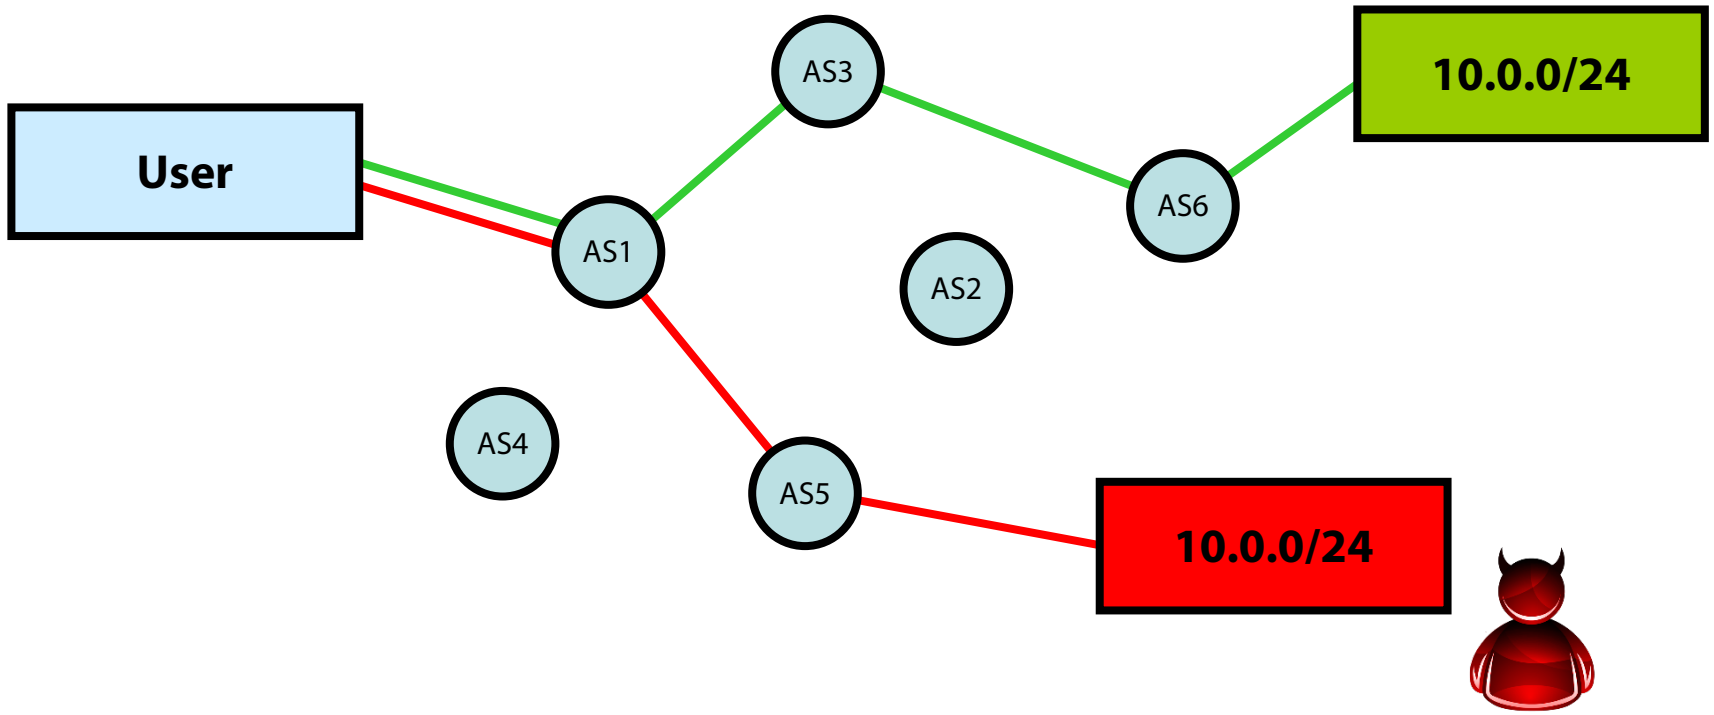

Prefix Hijacking

Announce a more specific prefix than someone else

Victim prefers the *more specific* path

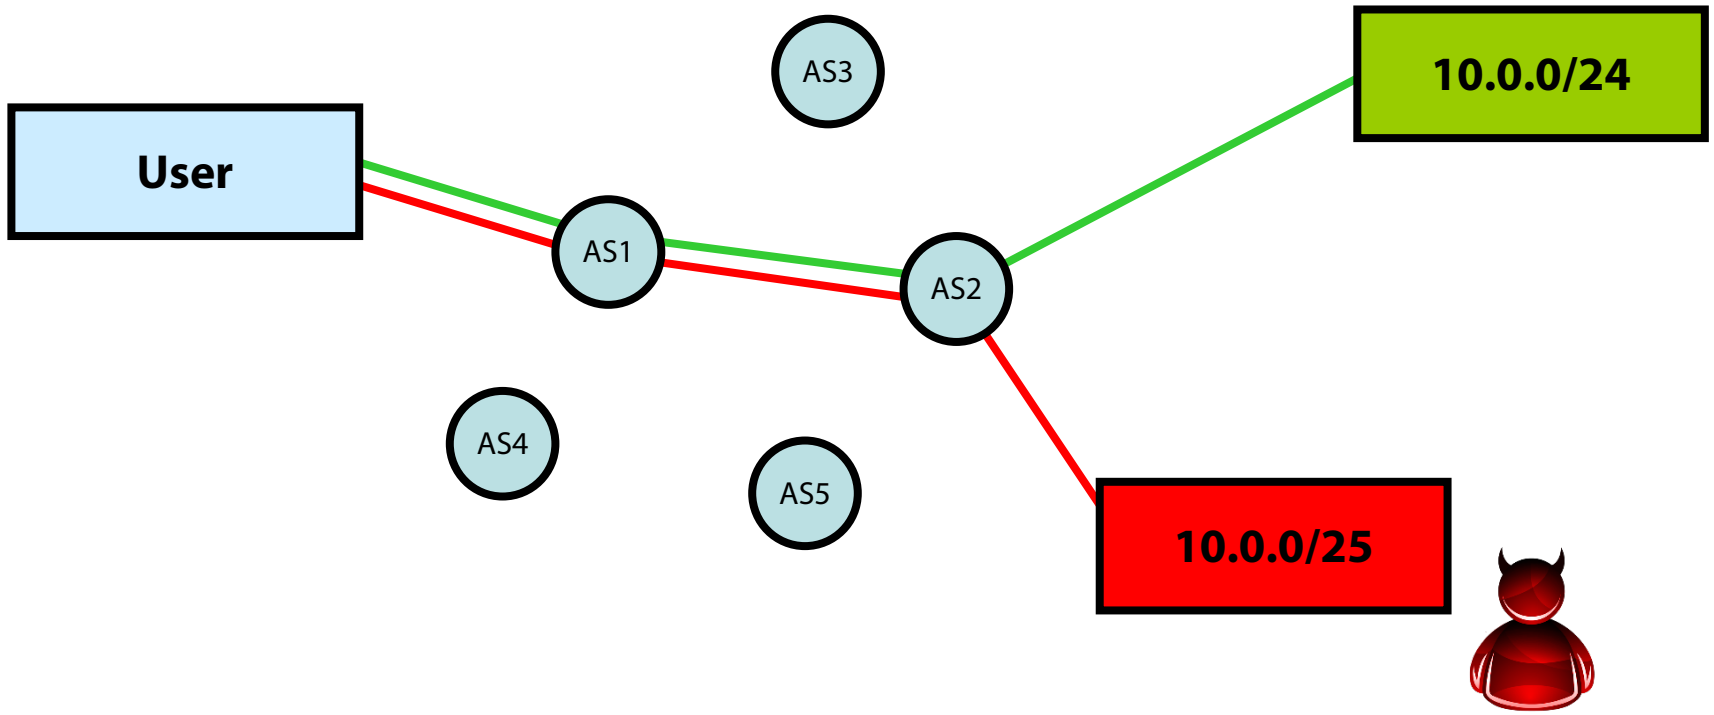

BGP leak causing Interne

https://bgpmon.net/bgp-leak-causing-internet-outages-in-japan-and-beyond/

BGPMON Now part of OpenDNS

HOME BLOG ABOUT US PRODUCTS AND SERVICES CLIENT PORTAL

BGP leak causing Internet outages in Japan and beyond.

Posted by Andree Toonk – August 26, 2017 – BGP instability – No Comments

Yesterday some Internet users would have seen issues with their Internet connectivity, experiencing slowness or parts of the Internet as unreachable. This incident hit users in Japan particularly hard and it caused the Internal Affairs and Communications Ministry of Japan to start an investigation into what caused the large-scale internet disruption that slowed or blocked access to websites and online services for dozens of Japanese companies.

In this blog post we will take a look at the root cause of these outages, who was affected and what networks were involved.

Starting at 03:22 UTC yesterday (aug 25) followers of @BGPstream would have seen an increase in alerts involving Google. The BGPstream alerts were informing us that Google was announcing the peering lan prefixes of a few well known Internet exchanges. This in itself is actually a fairly common type of incident and typically indicates something isn't quite right within the networks hijacking those prefixes and so these alerts were the first clues that something wasn't quite right with Google's BGP advertisements.

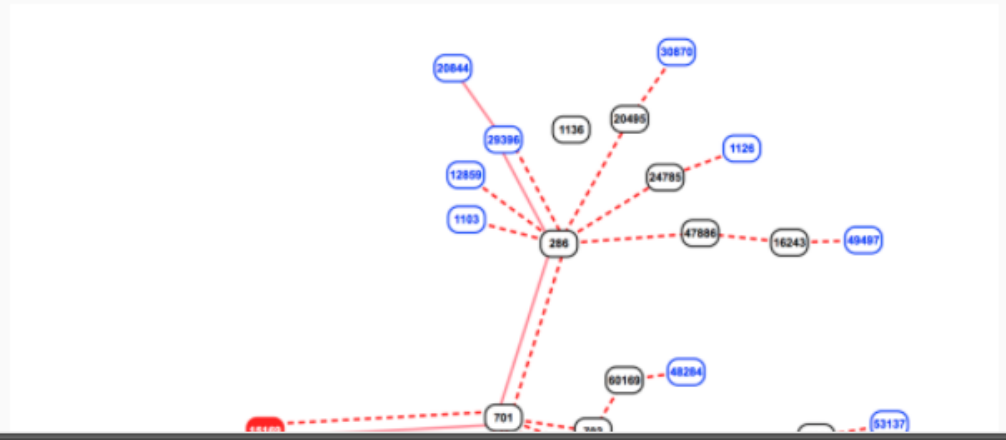

Latest Tweets

Tweets by @bgpmon

BGPmon.net
@bgpmon
New blog: Route leak via Google and Verizon causing internet outages in japan and beyond bgpmon.net/bgp-leak-causi…
Aug 26, 2017

BGPmon.net Retweeted

bgpstream
@bgpstream
BGP,OT,KP,Korea, Democratic People's Republic of,-,Outage affected 4 prefixes, bgpstream.com/event/95347
Aug 14, 2017

BGPmon.net
@bgpmon
Another Internet outage in Syria. Interestingly one remaining Syrian telecom network still BGP visible 91.144.0.0/20 twitter.com/bgpstream/stat…

Embed View on Twitter

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: sure

Use URL filtering?

Nope

Change the DNS record?

Nope

Use IP blocking?

Nope

Blackhole 208.65.153.0/24?

Yeah!

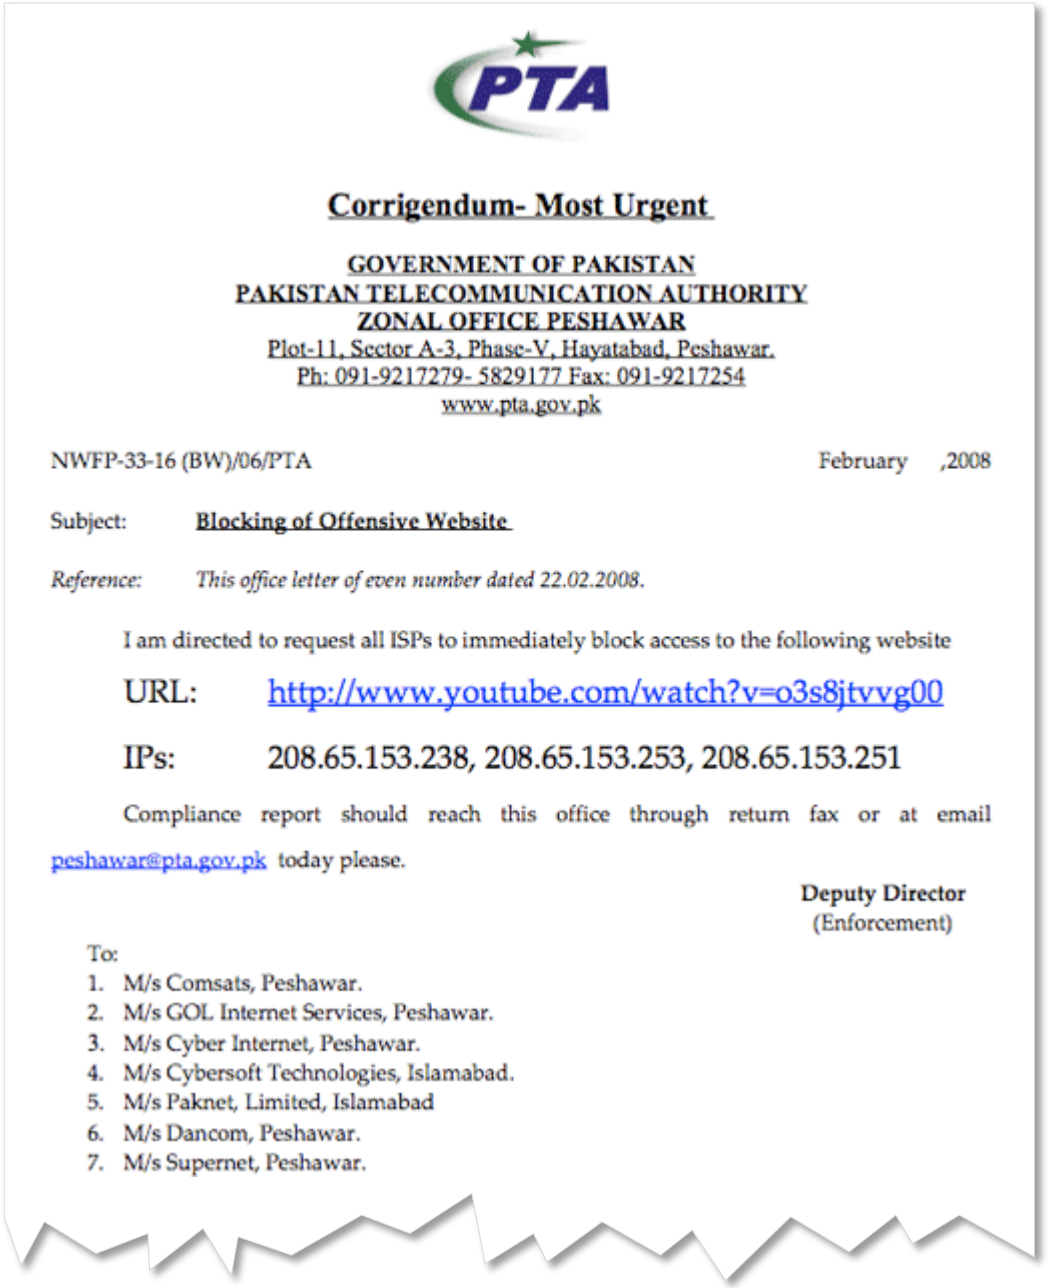

Corrigendum- Most Urgent

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

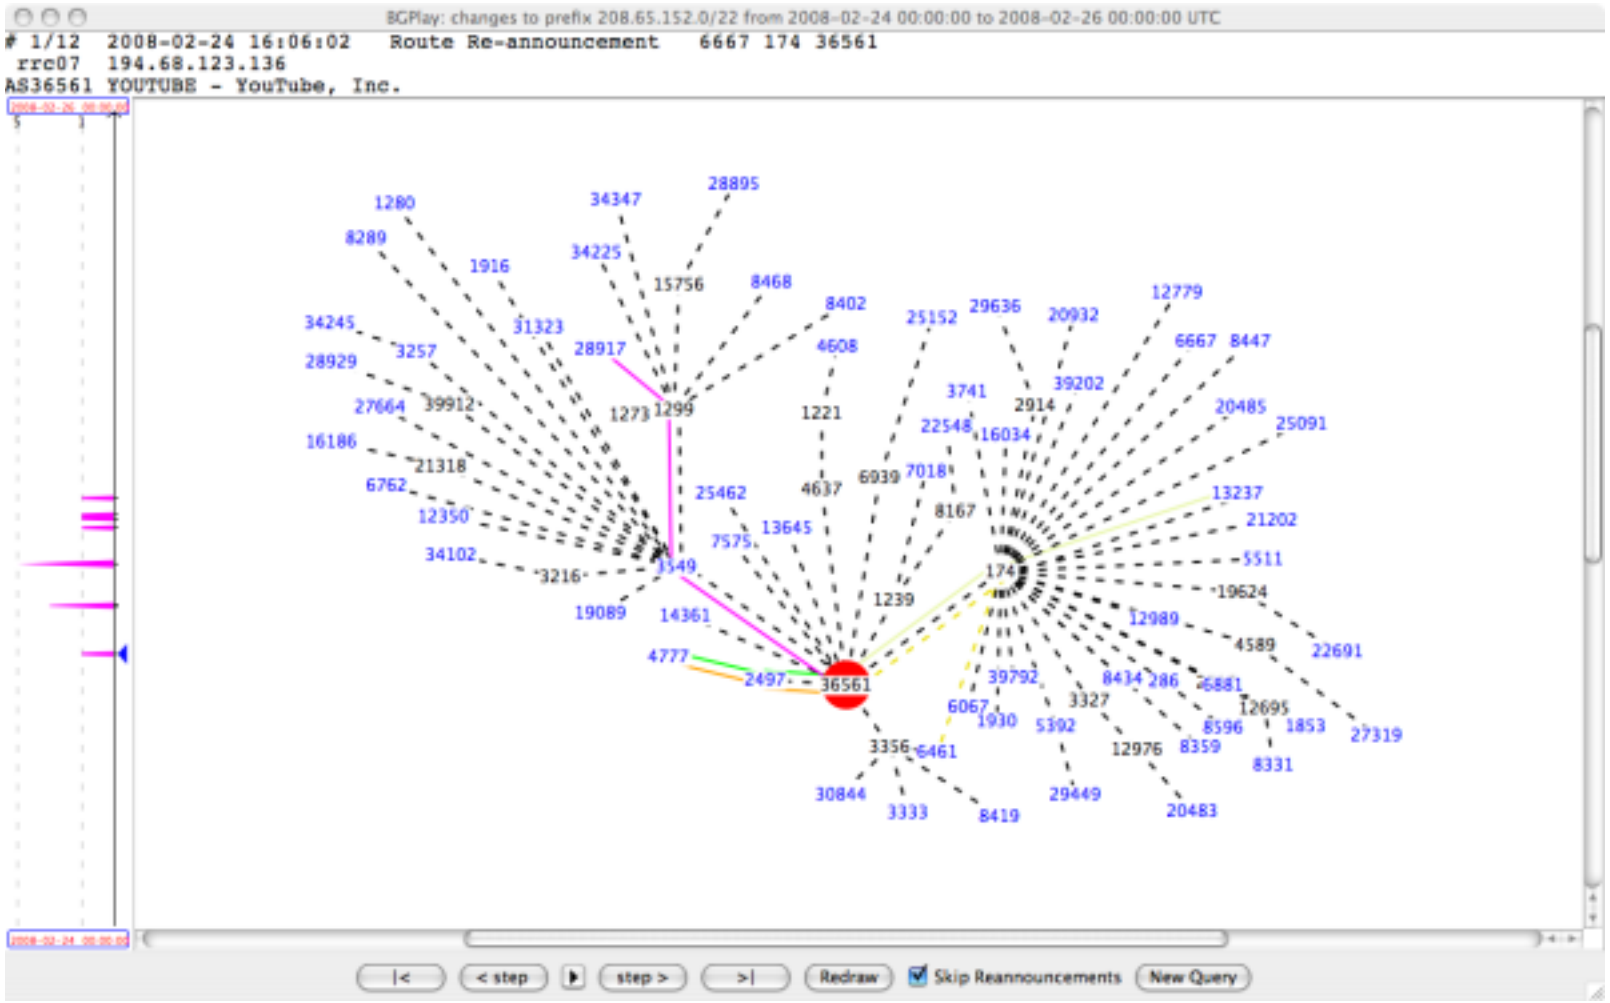

The prefix 208.65.153.0/24 is not announced on the Internet before the event

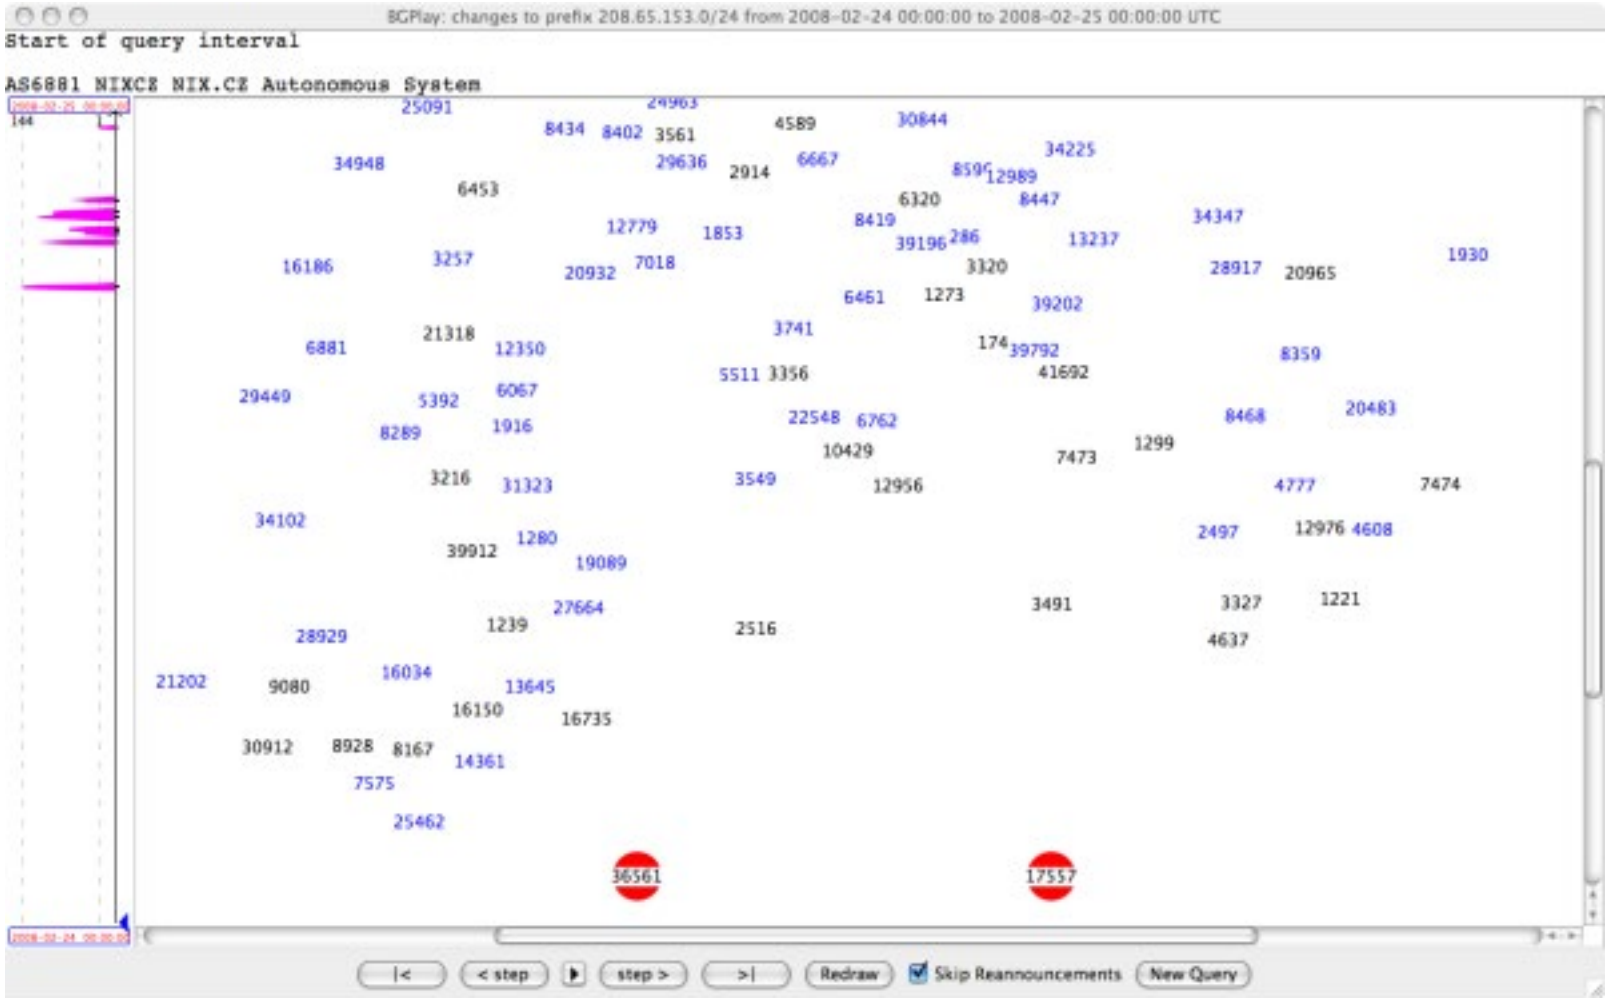

AS17557 (Pakistan Telecom) announces 208.65.153.0/24

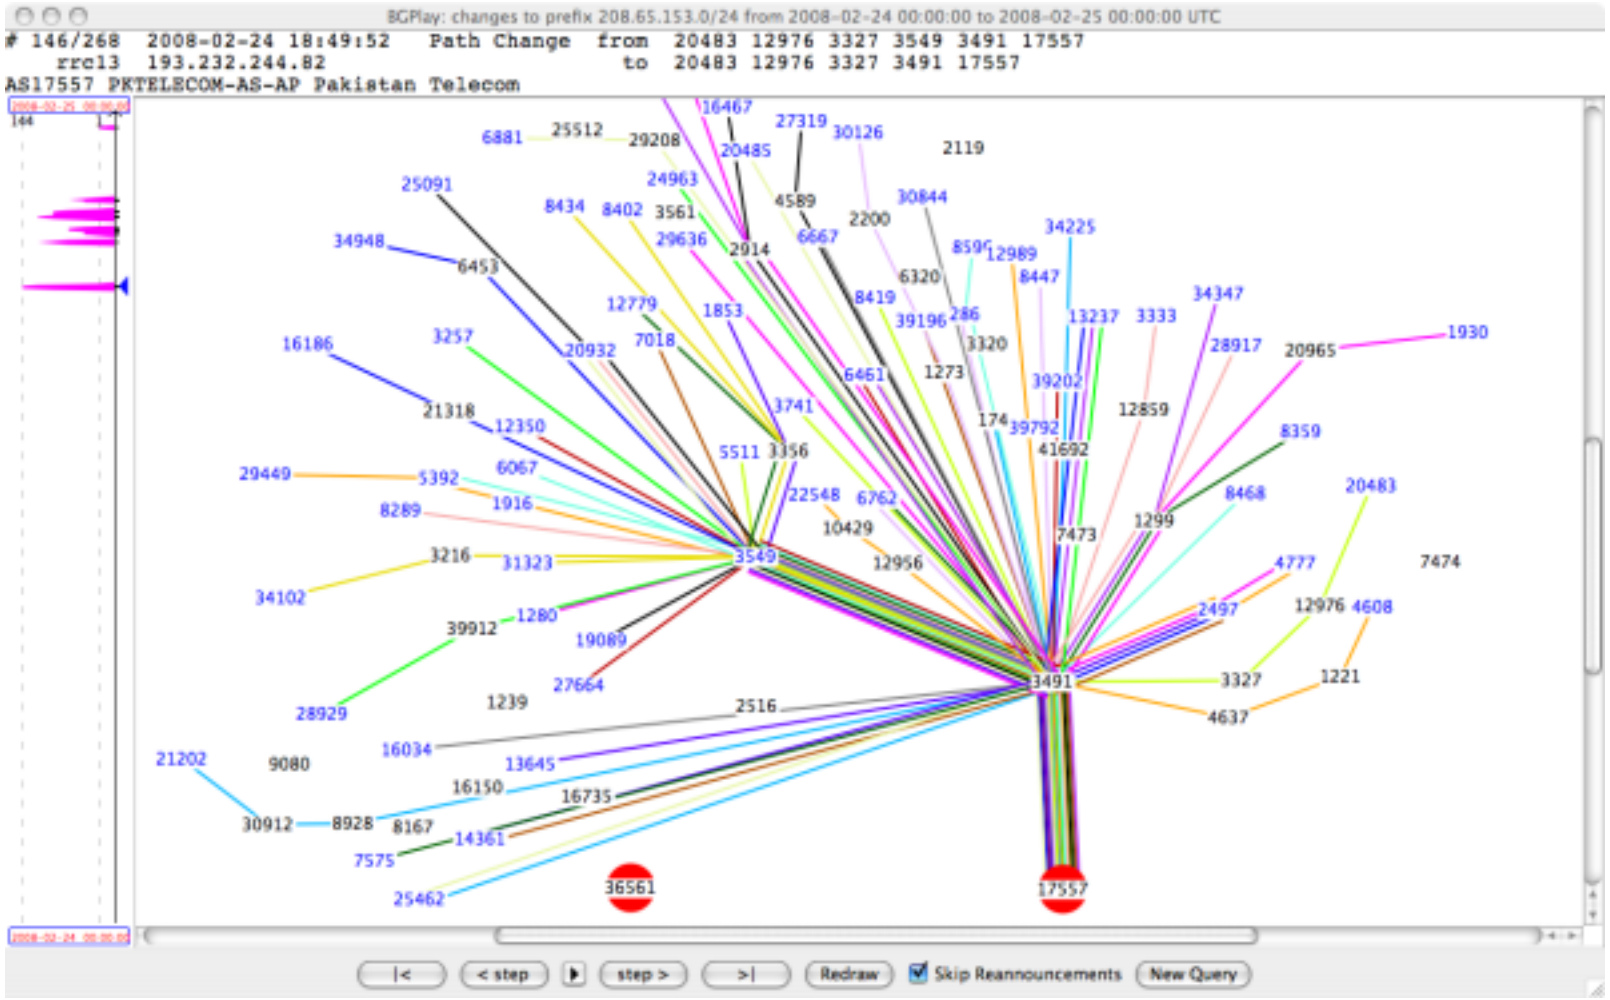

Other Notable Incidents

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

Enabled traffic interception

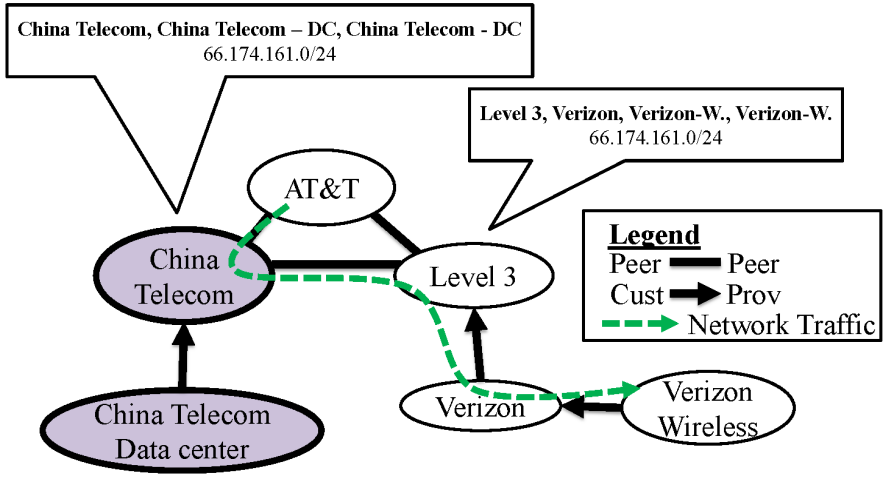

February 2014: hijacking of 51 networks (incl. Amazon, Digital Ocean, OVH)

Miner connections were redirected to an attacker-controlled mining pool

Attacker collected the miners' profit (estimated $83,000 in 4 months)

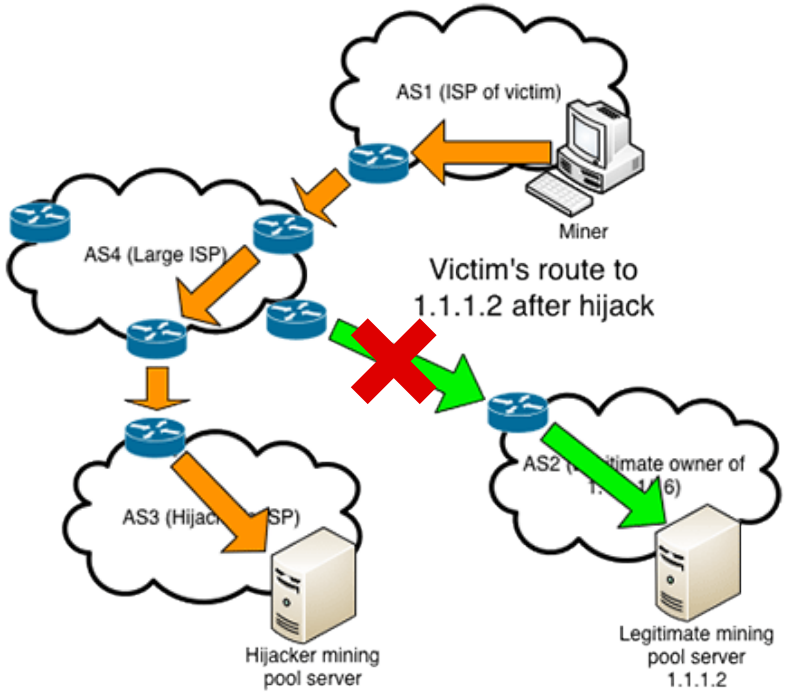

Popular Destinations rerouted to Russia

Posted by Andree Toonk – December 12, 2017 – Hijack – No Comments

Early this morning (UTC) our systems detected a suspicious event where many prefixes for high profile destinations were being announced by an unused Russian Autonomous System.

Starting at 04:43 (UTC) 80 prefixes normally announced by organizations such Google, Apple, Facebook, Microsoft, Twitch, NTT Communications and Riot Games were now detected in the global BGP routing tables with an Origin AS of 39523 (DV-LINK-AS), out of Russia.

Looking at timeline we can see two event windows of about three minutes each. The first one started at 04:43 UTC and ended at around 04:46 UTC. The second event started 07:07 UTC and finished at 07:10 UTC.

Even though these events were relatively short lived, they were significant because it was picked up by a large number of peers and because of several new more specific prefixes that are not normally seen on the Internet. So let's dig a little deeper.

One of the interesting things about this incident is the prefixes that were affected are all network prefixes for well known and high traffic internet organizations. The other odd thing is that the Origin AS 39523 (DV-LINK-AS) hasn't been seen announcing any prefixes for many years (with one exception below), so why does it all of sudden appear and announce prefixes for networks such as Google?

Latest Tweets

Tweets by @bgpmon

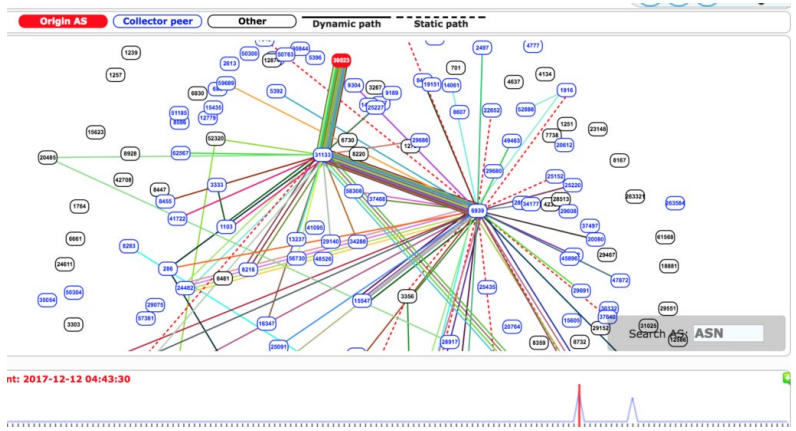

BORDER GATEWAY PROTOCOL ATTACK —

Suspicious event hijacks Amazon traffic for 2 hours, steals cryptocurrency

Almost 1,300 addresses for Amazon Route 53 rerouted for two hours.

DAN GOODIN - 4/24/2018, 3:00 PM

125

Amazon lost control of a small number of its cloud services IP addresses for two hours on Tuesday morning when hackers exploited a known Internet-protocol weakness that let them to redirect traffic to rogue destinations. By subverting Amazon's domain-resolution service, the attackers masqueraded as cryptocurrency website MyEtherWallet.com and stole about $150,000 in digital coins from unwitting end users. They may have targeted other Amazon customers as well.

The incident, which started around 6 AM California time, hijacked roughly 1,300 IP addresses, Oracle-owned Internet Intelligence said on Twitter. The malicious redirection was caused by fraudulent routes that were announced by Columbus, Ohio-based eNet, a large Internet service provider that is referred to as autonomous system 10297. Once in place, the eNet announcement caused Hurricane Electric and possibly Hurricane Electric customers and other eNet peers to send traffic over the same unauthorized routes. The 1,300 addresses belonged to Route 53, Amazon's domain name system service

In a statement, Amazon officials wrote: "Neither AWS nor Amazon Route 53 were hacked or compromised. An upstream Internet Service Provider (ISP) was compromised by a malicious actor who then used that provider to announce a subset of Route 53 IP addresses to other networks with whom this ISP was peered. These peered networks, unaware of this issue, accepted these announcements and incorrectly directed a small percentage of traffic for a single customer's domain to the malicious copy of that domain."

China Telecom has been using poisoned internet routes to suck up massive amounts of US and Canadian internet traffic

CORY DOCTOROW / 6:15 AM FRI OCT 26, 2018

In a new paper published in the journal *Military Cyber Affairs* researchers from the US Naval War College and Tel Aviv University document the use of BGP spoofing by China Telecom to redirect massive swathes of internet traffic through the company's routers as part of state military and commercial espionage efforts.

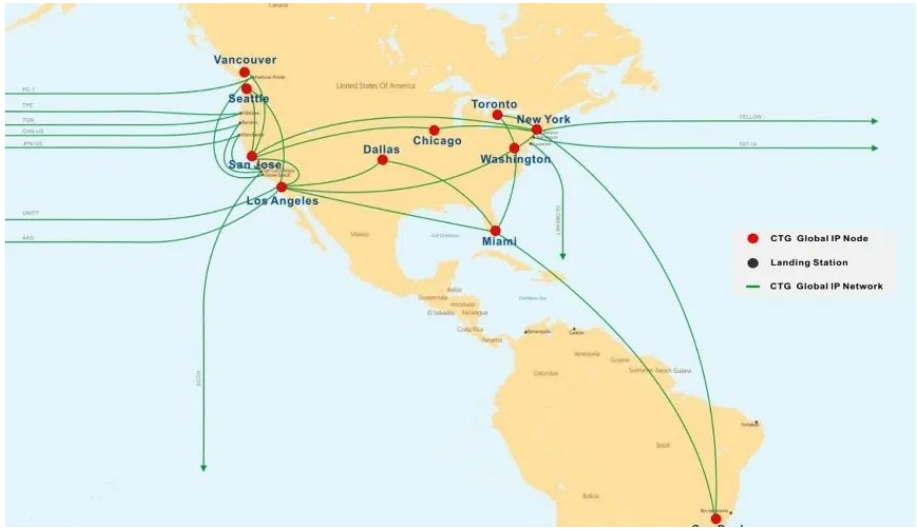

BGP is a notoriously insecure protocol used to route internet traffic; by design it is dynamic and responsive, moving traffic away from congested routes and onto those with more capacity: this flexibility can be exploited to force traffic to route through surveillance chokepoints, as well as for censorship (publishing BGP routes to censored services that dead-end in nonexistent addresses are a common technique in repressive regimes).

The researchers logged global BGP route announcements and discovered China Telecom publishing bogus routes that sucked up massive amounts of Canadian and US traffic and pushed it through Chinese listening posts. Much of today's internet traffic is still unencrypted, meaning that the entities monitoring these listening posts would have been able to read massive amounts of emails, instant messages and web-sessions.

Truth Behind the Celer Network

https://medium.com/coinmonks/truth-behind-the-celer-network-

Truth Behind the Celer Network cBridge cross-chain bridge incident: BGP hijacking

SlowMist ·
Published in Coinmonks · 8 min read · Aug 20, 2022

Background

Celer Network officials stated on August 18 that between 3:45 and 6:00 Beijing time, certain cBridge users were directed to malicious smart contracts. Initially, the cBridge front-end interface was suspected of being compromised by DNS hijacking.

Completely different from the previous cross-chain bridge hacking incidents such as Nomad, Wormhole, Ronin, Harmony, etc., this attack was not caused by bugs in smart contracts and cross-chain protocols or the intrusion of related servers, and the cross-chain assets locked in cBridge have also been kept safe. **In this attack, the hackers directly targeted the underlying infrastructure in the Internet architecture outside the Celer system, and allowed cross-chain users to access a "phishing" front-end user interface within a period of time by deceiving the Internet's underlying routing protocol (BGP).**The Celer network was able to limit damages due to their prompt response.This is because the Celer Network team has a 24-hour

nanog: Yet another BGP hijackin

https://seclists.org/nanog/2022/Aug/236

Nmap.org Npcap.com Sectools.org Insecure.org

SECLISTS.ORG

Site Search

nanog mailing list archives

NANOG

By Date By Thread

List Archive Search

Yet another BGP hijacking towards AS16509

From: Siyuan Miao <siyuan () misaka io>
Date: Tue, 23 Aug 2022 01:54:50 +0200

```
Hi folks,

Recently I read a post regarding the recent incident of Celer Network and
noticed a very interesting and successful BGP hijacking towards AS16509.

The attacker AS209243 added AS16509 to their AS-SET and a more specific
route object for the /24 where the victim's website is in ALTDB:
(Below is our IRRd4 server NRTM logging, UTC timezone)

irrd.log-20220817.gz:31106270-ADD 96126

irrd.log-20220817.gz:31106280-

irrd.log-20220817.gz:31106281-as-set:     AS-SET209243

irrd.log-20220817.gz:31106306-descr:      quickhost set

irrd.log-20220817.gz:31106332-members:    AS209243, AS16509

irrd.log-20220817.gz:31106362:mnt-by:     MAINT-QUICKHOSTUK

irrd.log-20220817.gz:31106392-changed:    crussell () quickhostuk net 20220816

irrd.log-20220817.gz:31106438-source:     ALTDB

irrd.log-20220817.gz:31147549-ADD 96127

irrd.log-20220817.gz:31147559-
```

Mitigating BGP Threats

Neighbor authentication

Only authorized peers can establish a given BGP neighbor relationship

TTL check

Most external peering sessions are established between adjacent routers

Good idea: set TTL=1 ➔ an attacker X hops away can still set TTL=1+X

Better idea: set TTL=255 and accept only packets with TTL=255 ➔ an attacker further away cannot spoof such a packet

BGP prefix restrictions, sanity checks, and filtering

Accept only a certain number of prefixes, ignore unwanted/illegal prefixes, limit the number of accepted AS path segments, …

ACLs to explicitly permit only authorized BGP traffic

According to existing security policies and configurations

Securing BGP

Secure BGP (S-BGP)

Each node signs its announcements

Secure origin BGP (soBGP)

Origin authentication + trusted database that guarantees that a path exists

BGPPSec

Allow recipients to validate the AS path included in update messages

Many deployment challenges

No complete, accurate registry of prefix ownership

Cannot react rapidly to changes in connectivity

Cost of cryptographic operations

Incremental deployment not always possible

Need for a public key infrastructure

Resource Public Key Infrastructure (RPKI)

Certified mapping from ASes to public keys and IP prefixes [RFC6480]

- Signed records that associate a BGP route announcement with the correct originating AS number

Deploying RPKI has two distinct stages

Route Origin Authorizations (ROAs)

- Signed route information advertised through BGP

Route Origin Validation (ROV)

- Validation of the cryptographic signatures of other networks' route information

https://manrs.org/2023/10/coordination-key-to-largest-rpki-deployment/