CAIDA update

PI k claffy, CAIDA
ISI/USC
Marina del Rey, CA
10 February 2017

CAIDA Update

- Data collection activities
  - Ongoing measurements
  - Data storage status
  - Data dissemination statistics
  - Recent publications
- • Related other activities
  - New data infrastructure
  - Related research activities
- Open issues
  - Portal, New Data Types

Data Collection Infrastructures

- Ark Platform (as of Sept 2016)
  - 170 monitors in 59 countries
  - 74 IPv6-enabled
  - 124 Raspberry PIs
- UCSD Network Telescope
  - As of January 2017, captures more than 1TB of compressed traffic trace data per day.
  - 28 TB: last full month (Aug 2016)
  - 182 TB: 2015
  - 211 TB: YTD 2016 (as of 9/13/16)
  - 288 TB: last 12 months at NERSC (as of 9/13/16)
  - 703 TB: total archived at NERSC

CAIDA Datasets/Requests

Datasets	Requests
Active Toplogy Measurements w/ Skitter	0
OC48 Peering Point Traces	3
Backscatter	10 (4 rejected)
DDoS 2007 Attack Dataset	3 (1 rejected)
IPv4 2013 Census Dataset	3
IPv4 Routed /24 Topology	0
IPv4 Routed /24 DNS Names	0
IPv6 Topology	0
Internet Topology Data Kits (ITDK)	2 (1 withdrawn)
Patch Tuesday Dataset	3 (1 rejected)
Three Days of Conficker Dataset	4
Two-Days-in-2008 Telescope Dataset	3 (1 rejected)
UCSD Real-time Network Telescope Dataset	5 (1 rejected, 3
UCSD Telescope Darknet Scanners Dataset	7 (1 rejected)
Witty Worm	2

New and Upcoming Data Sets

- (2) Macroscopic Internet Topology Data Kit (ITDK)
  http://www.caida.org/data/internet-topology-data-kit/
- IPv4 2013 Census Dataset
  http://www.caida.org/data/active/ipv4_2013_census_dataset.xml
  (available from IMPACT only)
- UCSD Network Telescope -- Darknet Scanners Dataset
  http://www.caida.org/data/passive/telescope-darknet-scanners_dataset.xml
  (available from IMPACT only)
- AS Border Mapping Dataset (coming soon)
  http://www.caida.org/publications/papers/2016/bdrmap/
- AS to Facilities Dataset (coming soon)
- Spoofer data

External Publications Using IMPACT Data

Statistics for publications that make use of the UCSD Network Telescope Dataset 2005-2016.

http://www.caida.org/data/publications/bydataset/index.xml#UCSD Network Telescope

UCSD Network Telescope	102
backscatter-2004-2005	8
backscatter-2006	4
backscatter-2007	8
backscatter-2008	17
backscatter-generic	4
backscatter-tics	3
code-red worm	7
code-red-generic	2
telescope-2days-2008	12
telescope-3days-conficker	14
telescope-educational	3
telescope-generic	7
telescope-patch-tuesday	2
telescope-real-time	5
witty worm (public)	1
witty worm (restricted)	17
witty-generic	6

Tools under consideration

- Vela: On-Demand Topology Measurement Service of CAIDA’s Ark infrastructure
  - Web interface https://vela.caida.org/
  - Command-Line interface

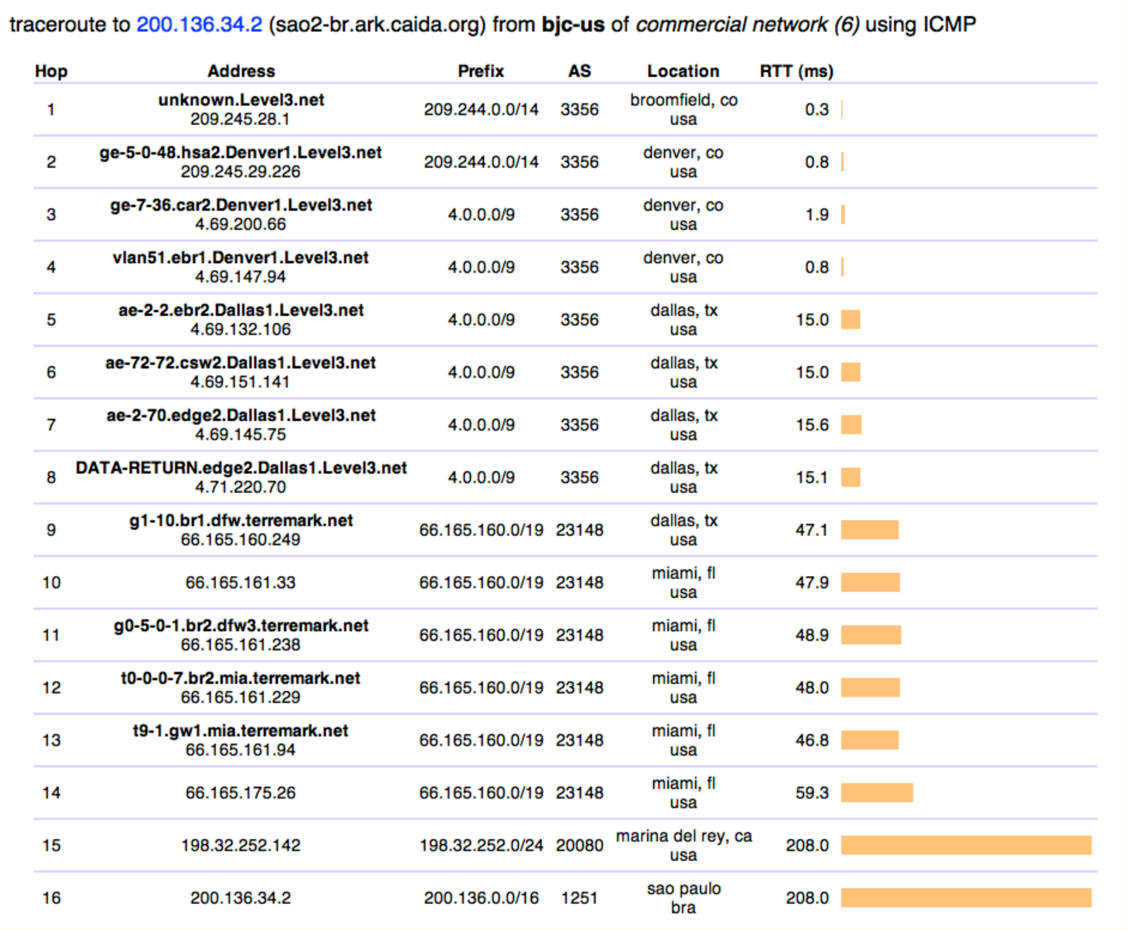

traceroute to 200.136.34.2 (sao2-br.ark.caida.org) from **bjc-us** of *commercial network (6)* using ICMP

Hop	Address	Prefix	AS	Location	RTT (ms)
1	**unknown.Level3.net** 209.245.28.1	209.244.0.0/14	3356	broomfield, co usa	0.3
2	**ge-5-0-48.hsa2.Denver1.Level3.net** 209.245.29.226	209.244.0.0/14	3356	denver, co usa	0.8
3	**ge-7-36.car2.Denver1.Level3.net** 4.69.200.66	4.0.0.0/9	3356	denver, co usa	1.9
4	**vlan51.ebr1.Denver1.Level3.net** 4.69.147.94	4.0.0.0/9	3356	denver, co usa	0.8
5	**ae-2-2.ebr2.Dallas1.Level3.net** 4.69.132.106	4.0.0.0/9	3356	dallas, tx usa	15.0
6	**ae-72-72.csw2.Dallas1.Level3.net** 4.69.151.141	4.0.0.0/9	3356	dallas, tx usa	15.0
7	**ae-2-70.edge2.Dallas1.Level3.net** 4.69.145.75	4.0.0.0/9	3356	dallas, tx usa	15.6
8	**DATA-RETURN.edge2.Dallas1.Level3.net** 4.71.220.70	4.0.0.0/9	3356	dallas, tx usa	15.1
9	**g1-10.br1.dfw.terremark.net** 66.165.160.249	66.165.160.0/19	23148	dallas, tx usa	47.1
10	66.165.161.33	66.165.160.0/19	23148	miami, fl usa	47.9
11	**g0-5-0-1.br2.dfw3.terremark.net** 66.165.161.238	66.165.160.0/19	23148	miami, fl usa	48.9
12	**t0-0-0-7.br2.mia.terremark.net** 66.165.161.229	66.165.160.0/19	23148	miami, fl usa	48.0
13	**t9-1.gw1.mia.terremark.net** 66.165.161.94	66.165.160.0/19	23148	miami, fl usa	46.8
14	66.165.175.26	66.165.160.0/19	23148	miami, fl usa	59.3
15	198.32.252.142	198.32.252.0/24	20080	marina del rey, ca usa	208.0
16	200.136.34.2	200.136.0.0/16	1251	sao paulo bra	208.0

Vela/Henya Web Interface to Topology Measurements and Data

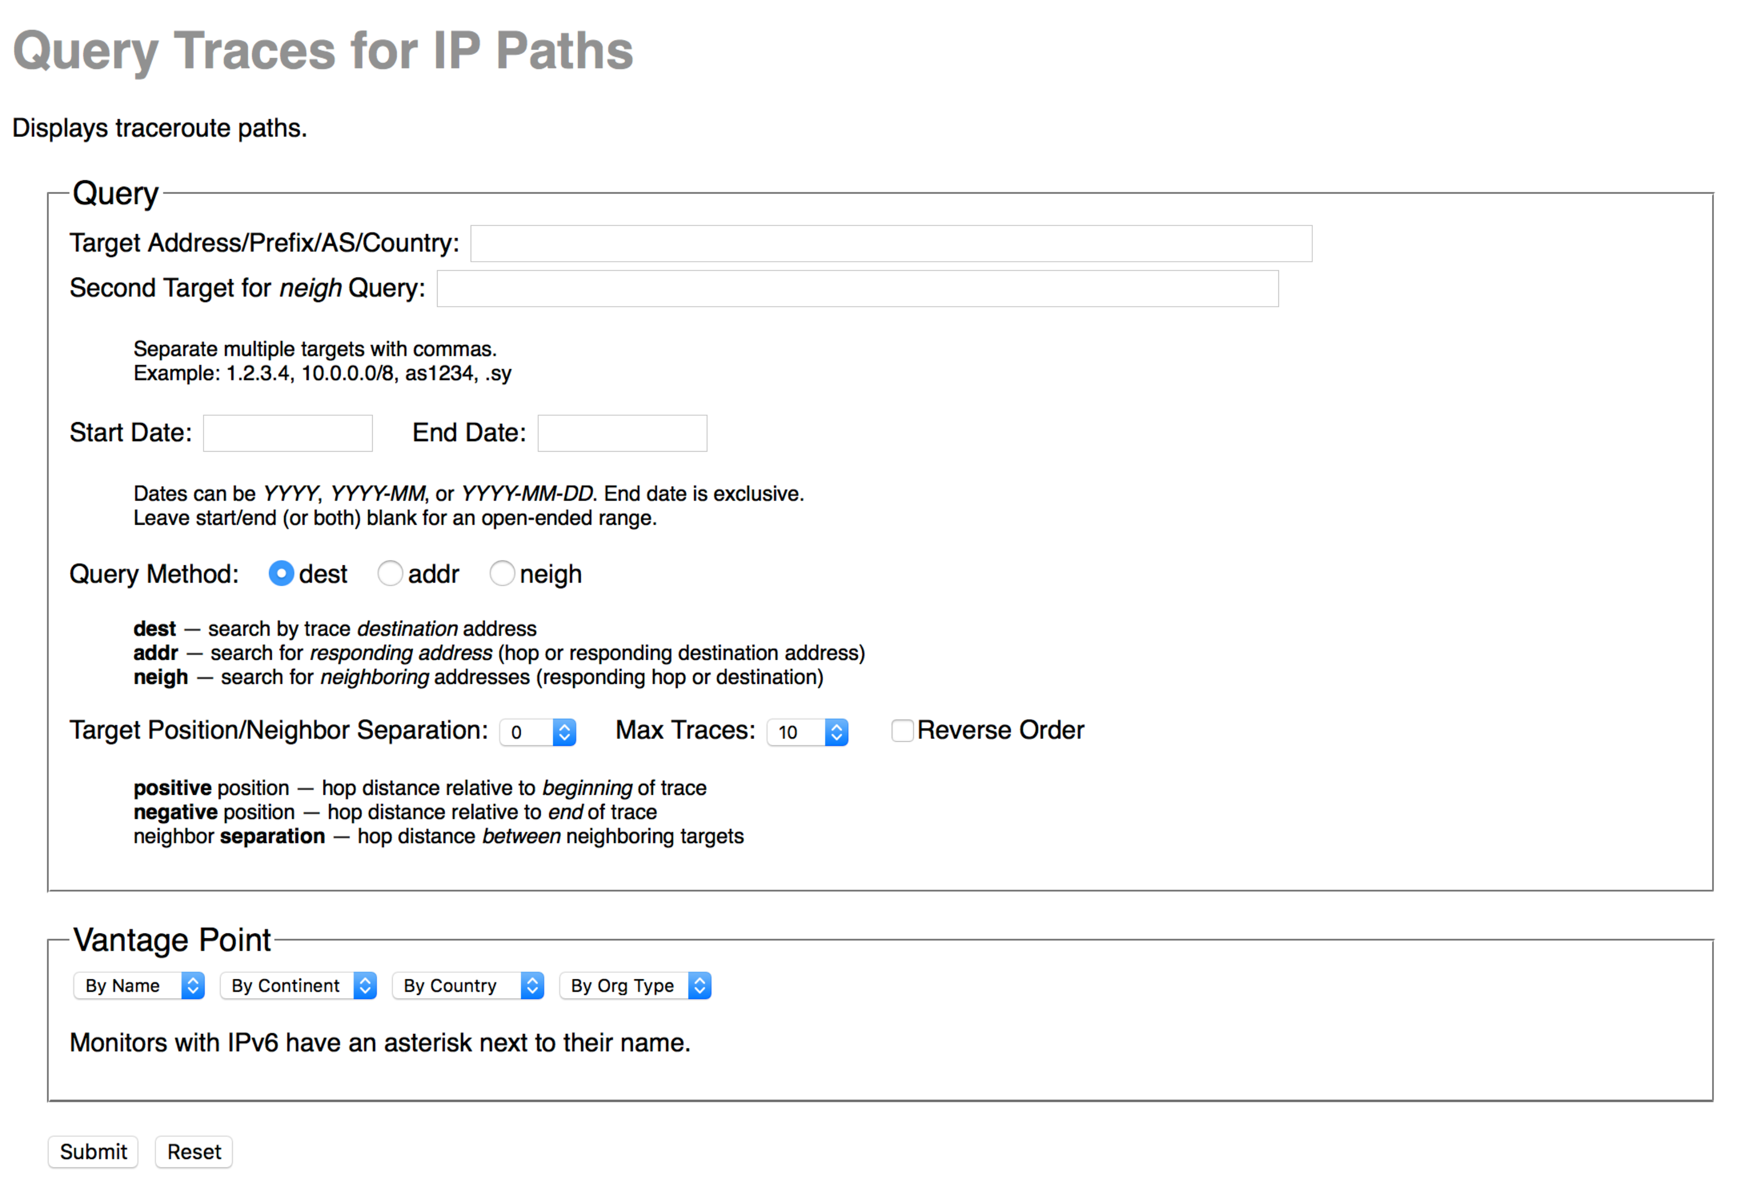

Query Traces for IP Paths

Displays traceroute paths.

Query

Target Address/Prefix/AS/Country:

Second Target for *neigh* Query:

Separate multiple targets with commas.
Example: 1.2.3.4, 10.0.0.0/8, as1234, .sy

Start Date: End Date:

Dates can be *YYYY*, *YYYY-MM*, or *YYYY-MM-DD*. End date is exclusive.
Leave start/end (or both) blank for an open-ended range.

Query Method: ● dest ○ addr ○ neigh

dest — search by trace *destination* address
addr — search for *responding address* (hop or responding destination address)
neigh — search for *neighboring* addresses (responding hop or destination)

Target Position/Neighbor Separation: 0 Max Traces: 10 □ Reverse Order

positive position — hop distance relative to *beginning* of trace
negative position — hop distance relative to *end* of trace
neighbor **separation** — hop distance *between* neighboring targets

Vantage Point

Monitors with IPv6 have an asterisk next to their name.

Tools: Henya

- Henya: Large-Scale Internet Topology Query System
  - Access via the Vela web interface https://vela.caida.org/
  - 9 years of “Routed /24” trace routes
    - 47 billion traces in 20TB of files
    - growing yearly by 10 billion traces
  - 1 year of “Prefix Probing” trace routes
    - growing yearly by 9 billion traces

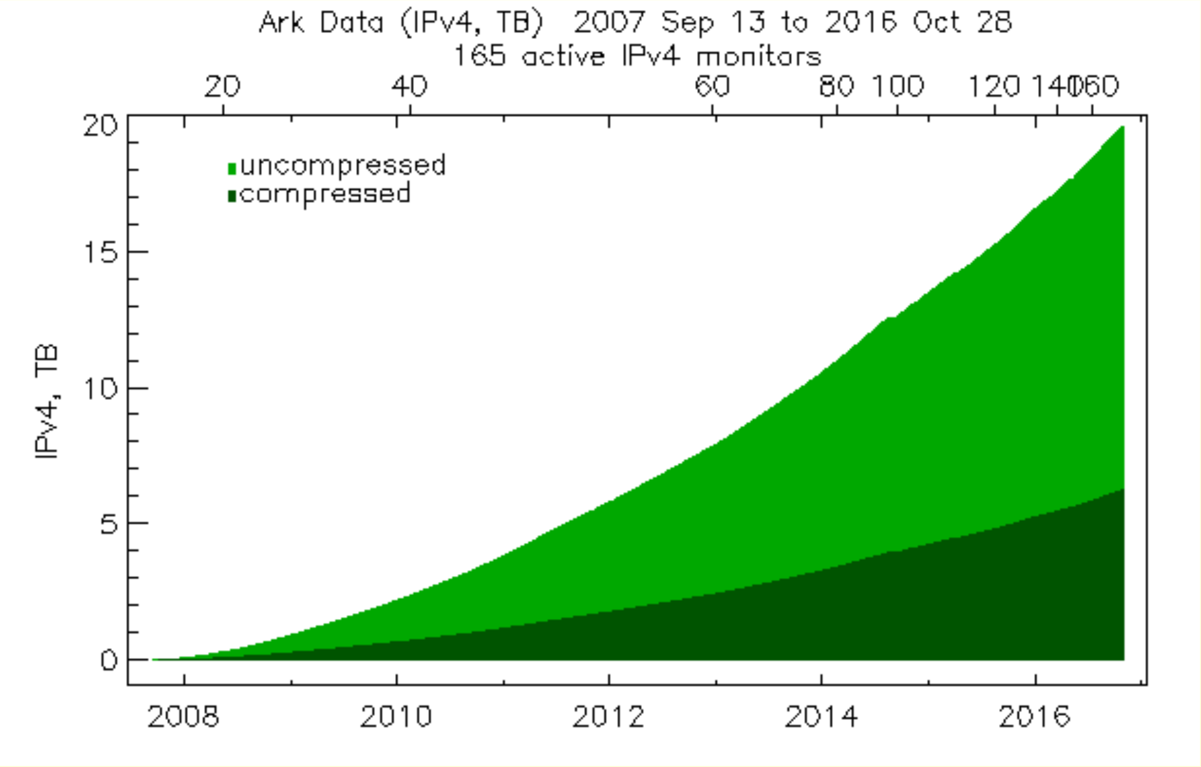

Henya Topology Queries

- find occurrences of traceroute path elements
- 《targets》 = IP addreses, prefixes, ASes, or countries
- Queries:
  - traceroutes toward 《targets》
  - traceroutes containing one or more 《targets》
- Parameters:
  - measurement vantage points
  - data collection time periods
  - position of 《targets》 in path
  - hop distance between sets of 《targets》

Henya Query Complexity

- the most complex case:
  - traceroutes containing two or more 《targets》
    - precisely: traceroutes containing some hop h1 ∊ 《targets1》, h2 ∊ 《targets2》, ···
  - example: traceroutes containing hops in both 《Germany》 and 《Japan》

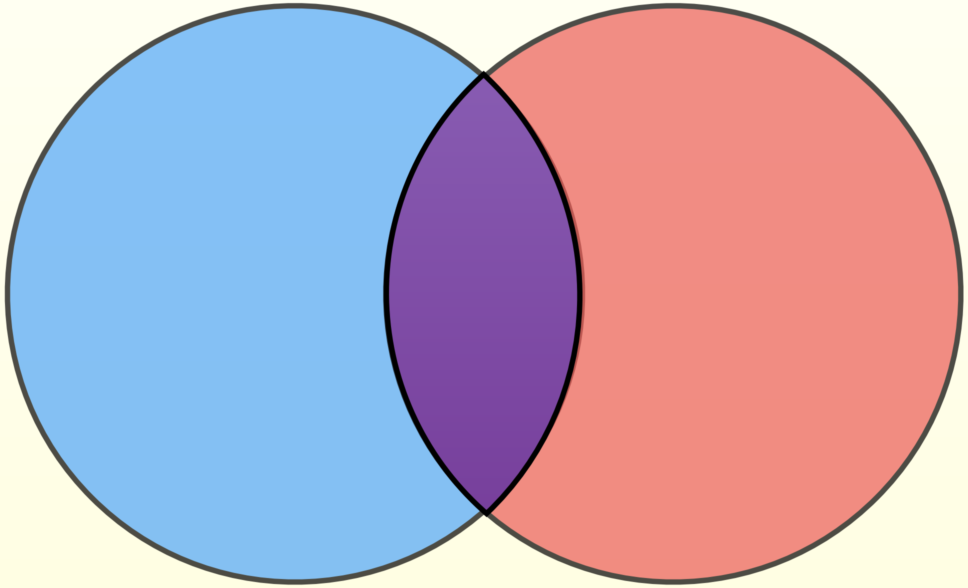

traceroutes with 《Germany》

traceroutes with 《Japan》

find intersection of two sets

- harder:
  - traceroutes with hops in 《Germany or UK or France》 and hops in 《ATT or Level3 network》 and hops in 《Amsterdam Internet Exchange》

Vela and Henya Access Policies

- Currently accepting requests for accounts on Vela
- Currently accepting requests for early access to Henya and a subset of total topology dataset.

Restricted Dataset Requests

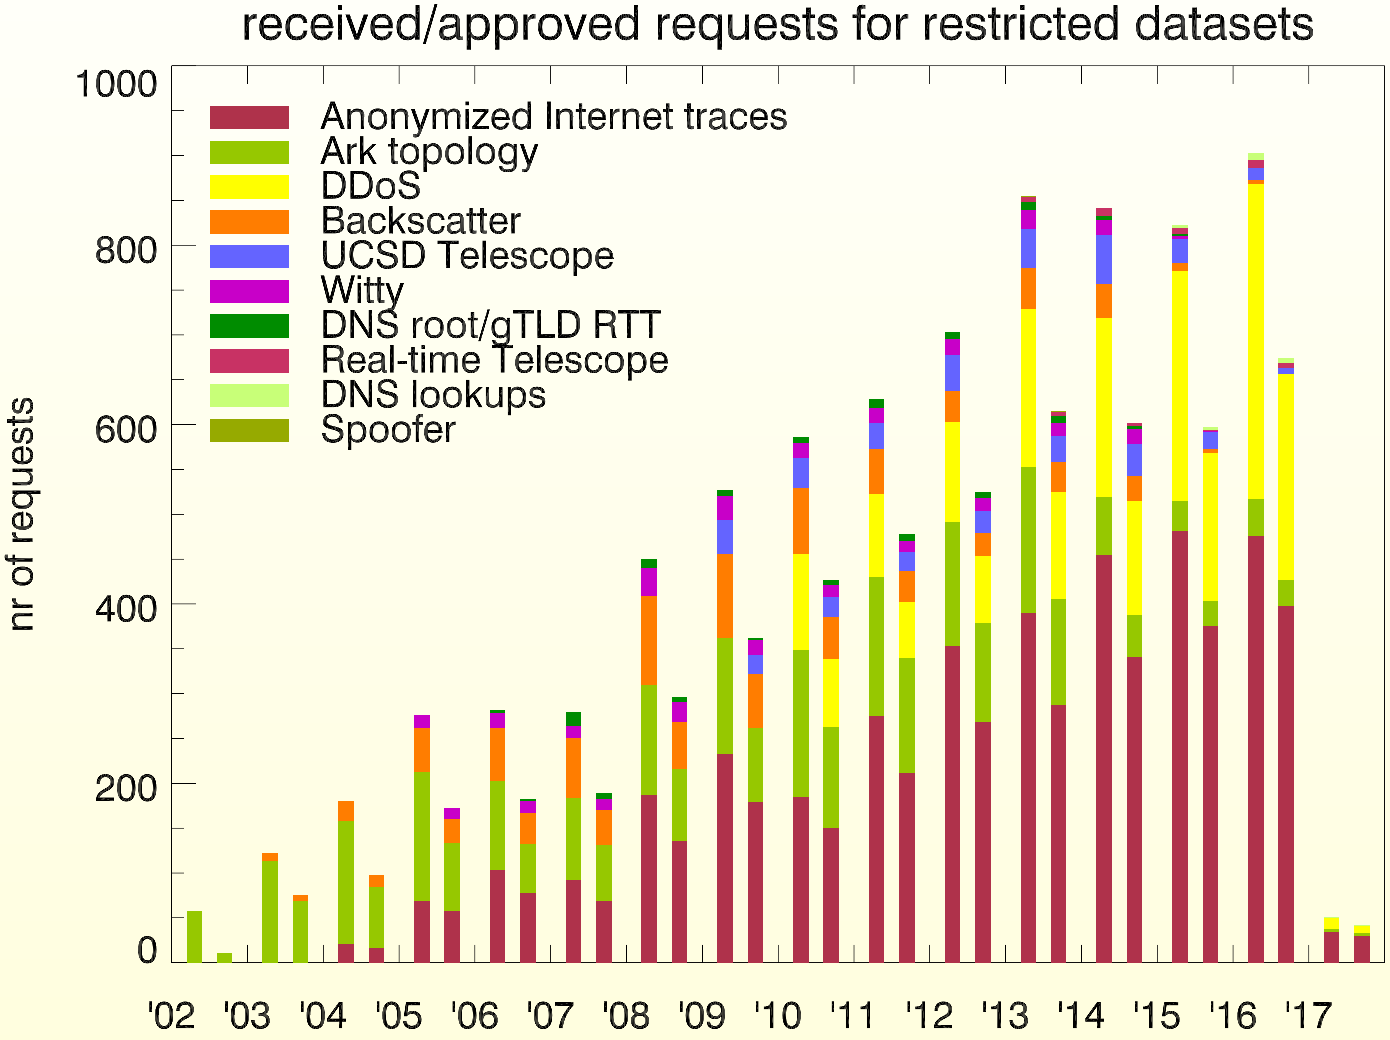

* This graph now includes all passive traces (including OC192). Previous graphs included only OC48 requests.

http://www.caida.org/data/about/

Users downloading public data

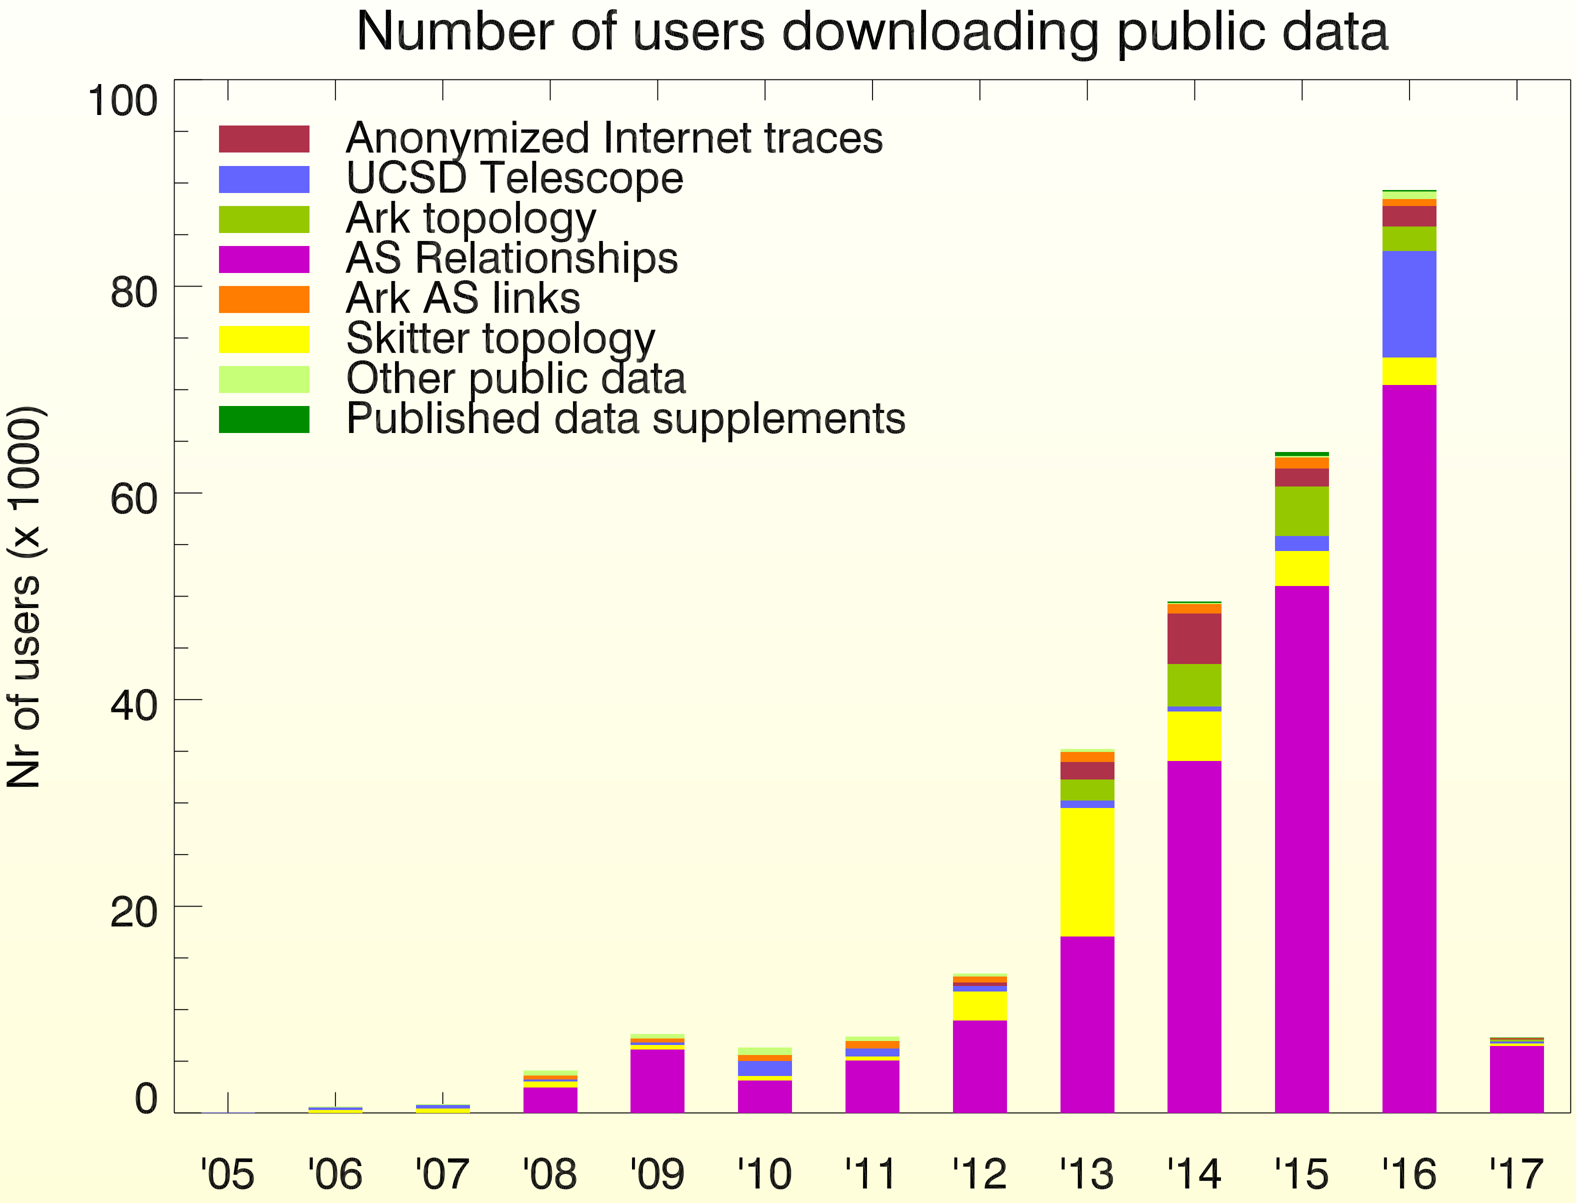

http://www.caida.org/data/about/

Users downloading restricted data

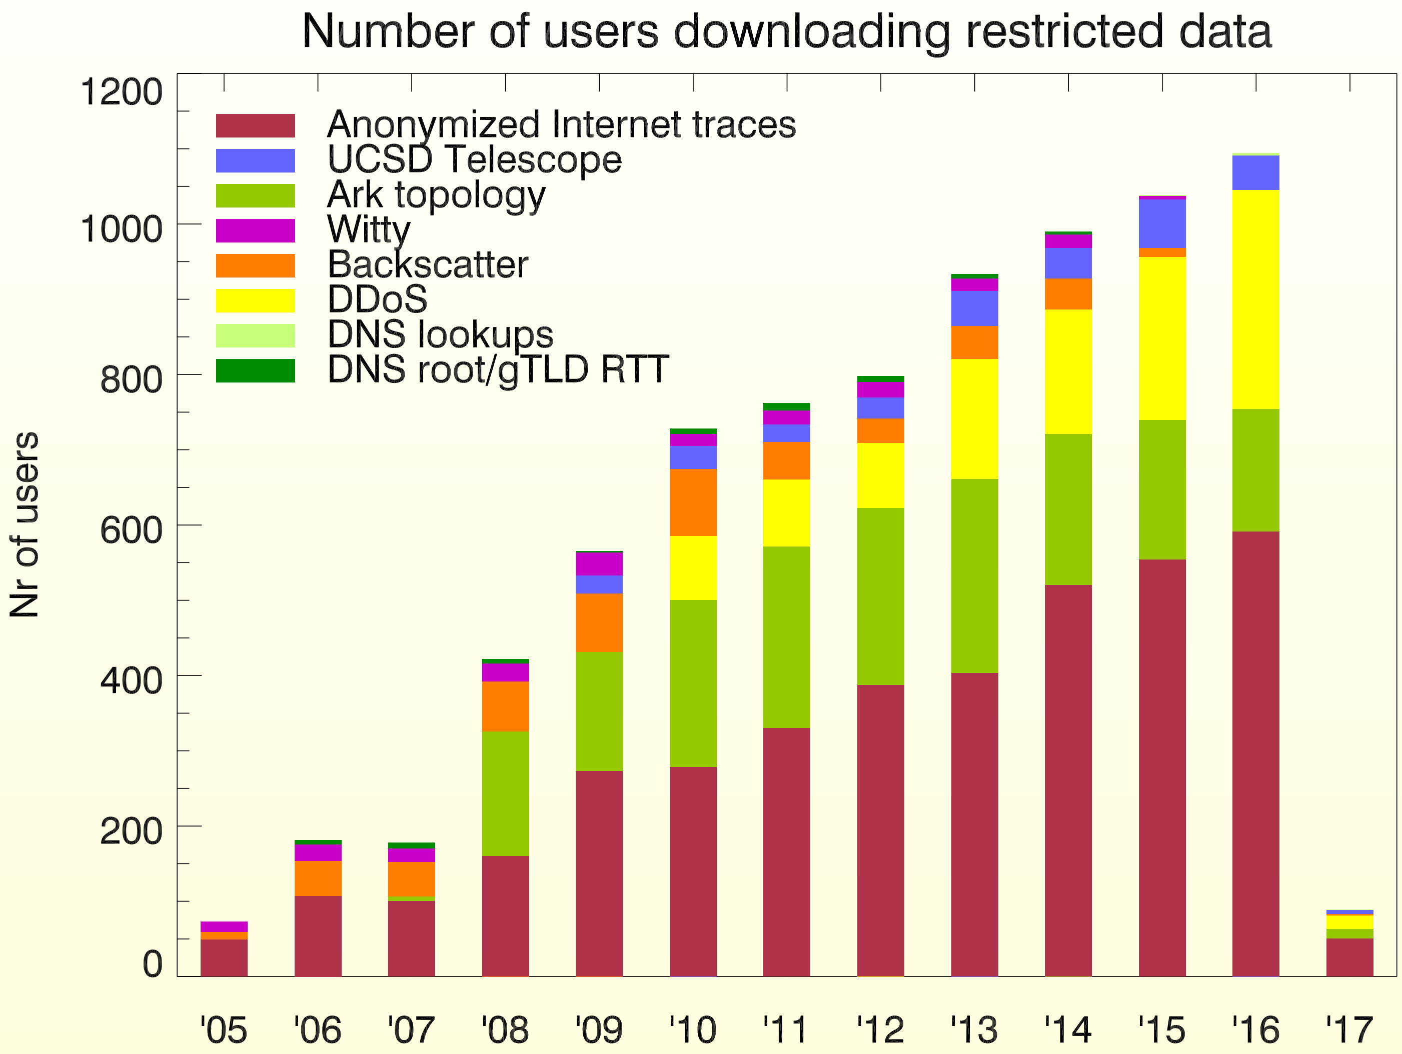

* This graph now includes all passive traces (including OC192). Previous graphs included only OC48 downloads.

http://www.caida.org/data/about/

Public data downloaded

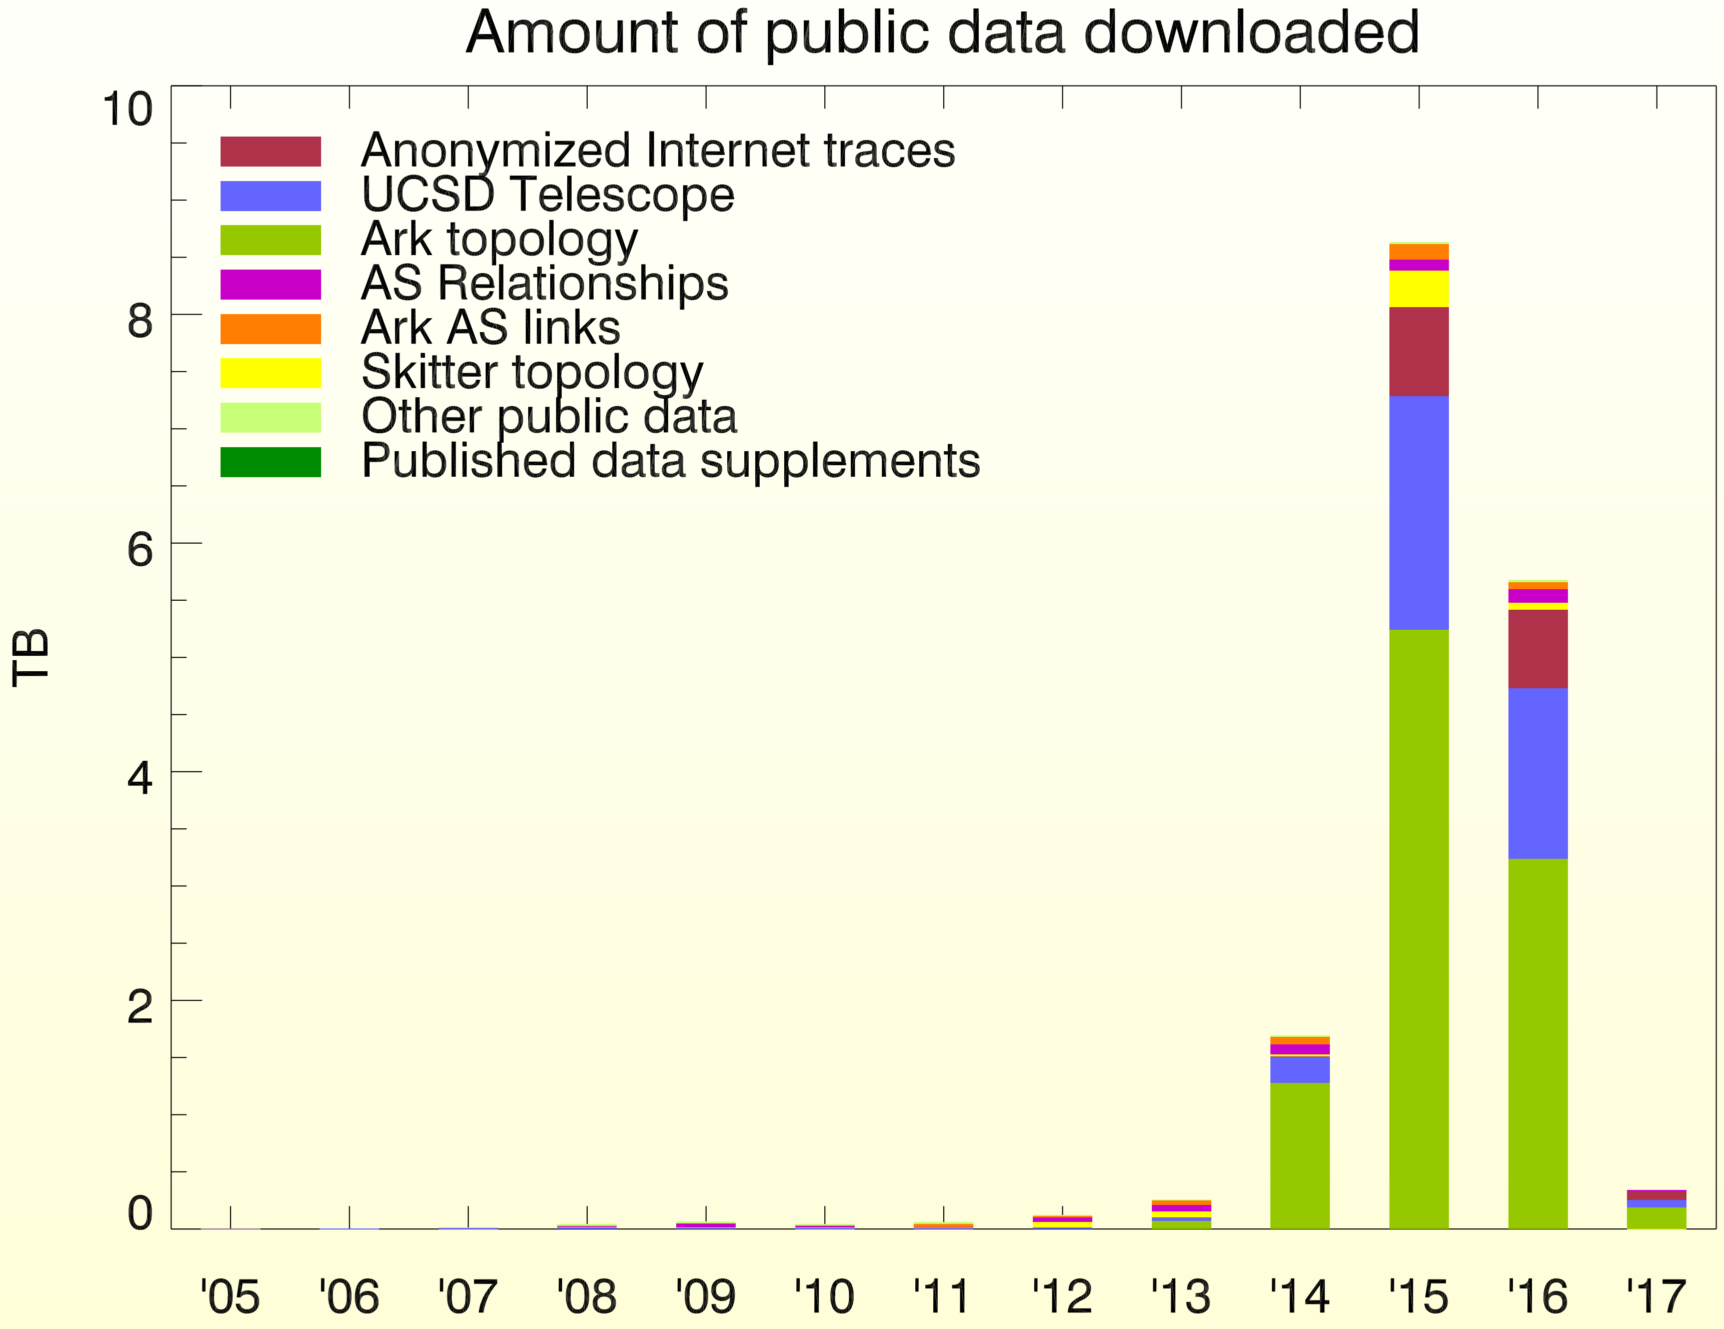

http://www.caida.org/data/about/

Restricted data downloaded

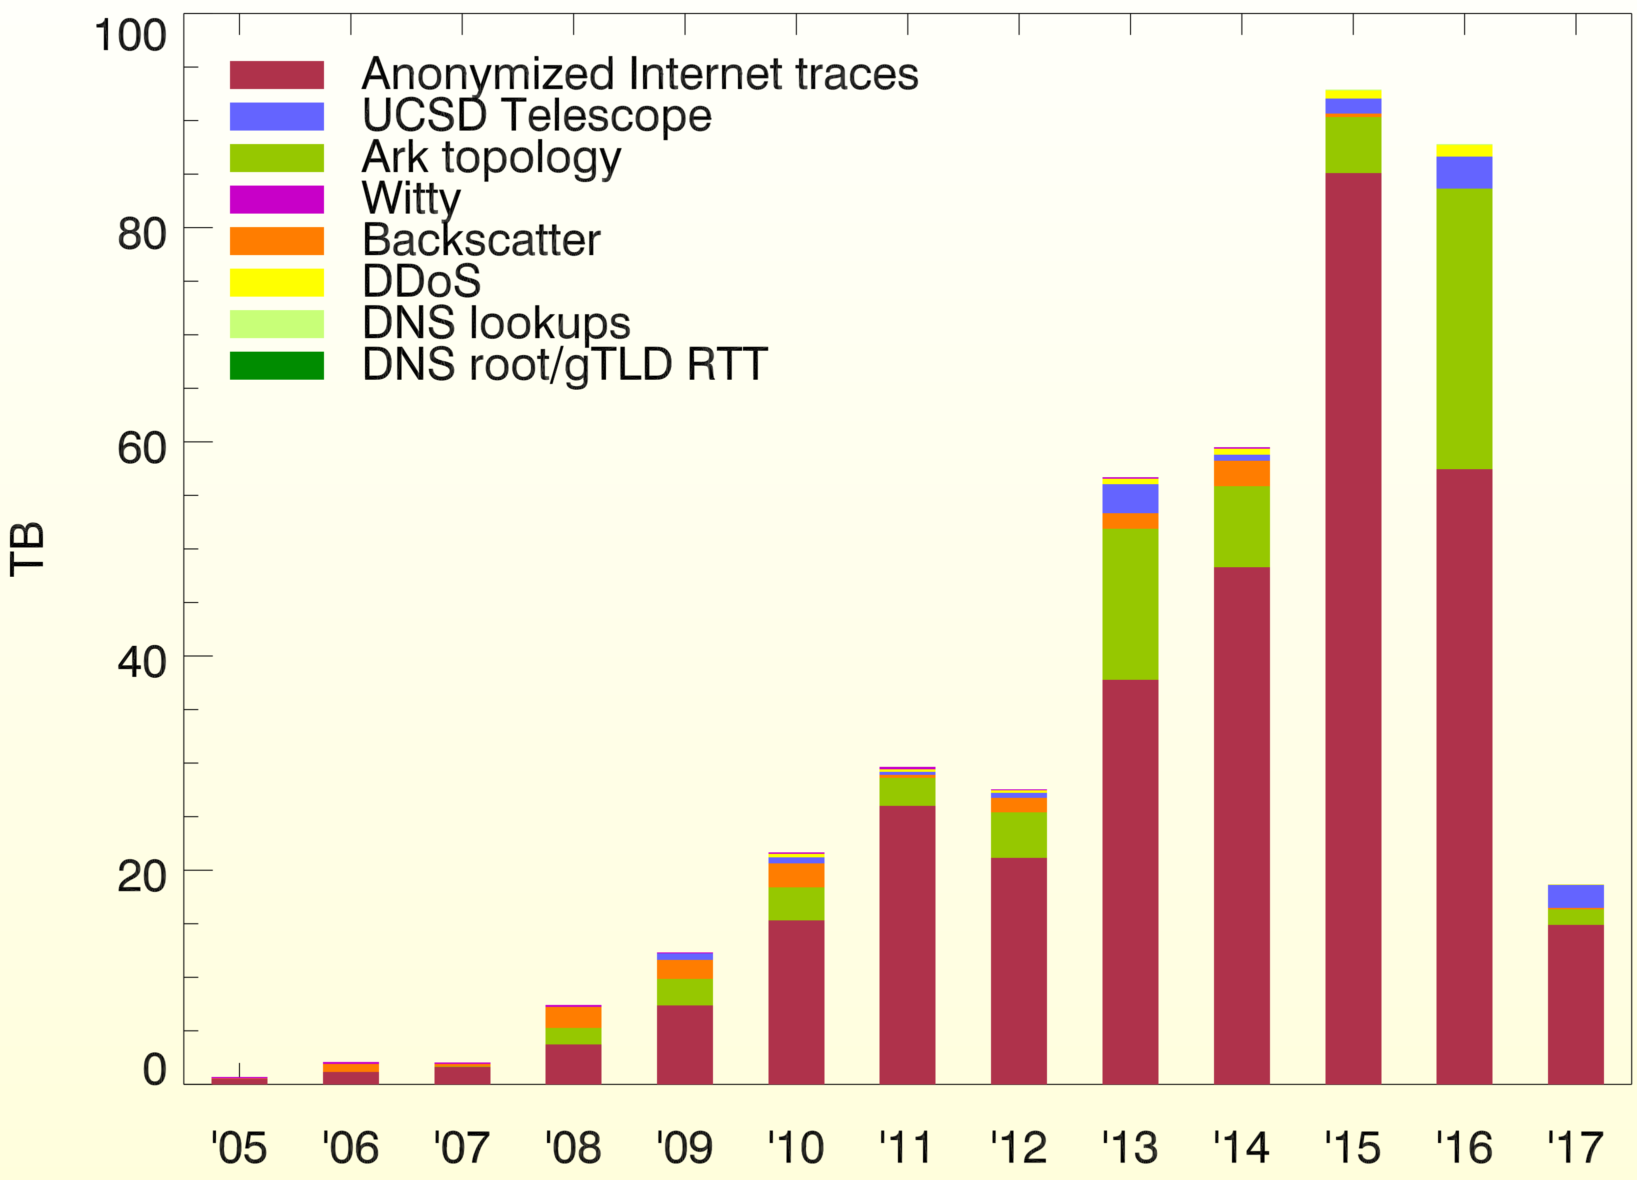

- drop in topology data in 2016 due to making topology data public

http://www.caida.org/data/about/

Recent Related R&D Activities

- DHS: Spoofing measurement (spoofer.caida.org)
- New DHS project: Science of Internet Security: Technology and Experimental Research (SISTER)
- NSF: Internet Outage Detection and Analysis (IODA) (ioda.caida.org)
- NSF: Internet congestion mapping system (beamer.caida.org)

Software Systems for Surveying Spoofing Susceptibility

- DHS S&T funded project that seeks to minimize Internet's susceptibility to spoofed DDoS attacks

- Goal: develop, build, and operate multiple open-source software tools to assess and report on the deployment of source address validation (SAV) best anti-spoofing practices.

- **https://spoofer.caida.org/** **<— plz download now!**

- **Will share data through IMPACT**

Software Systems for Surveying Spoofing Susceptibility

Recent Tests

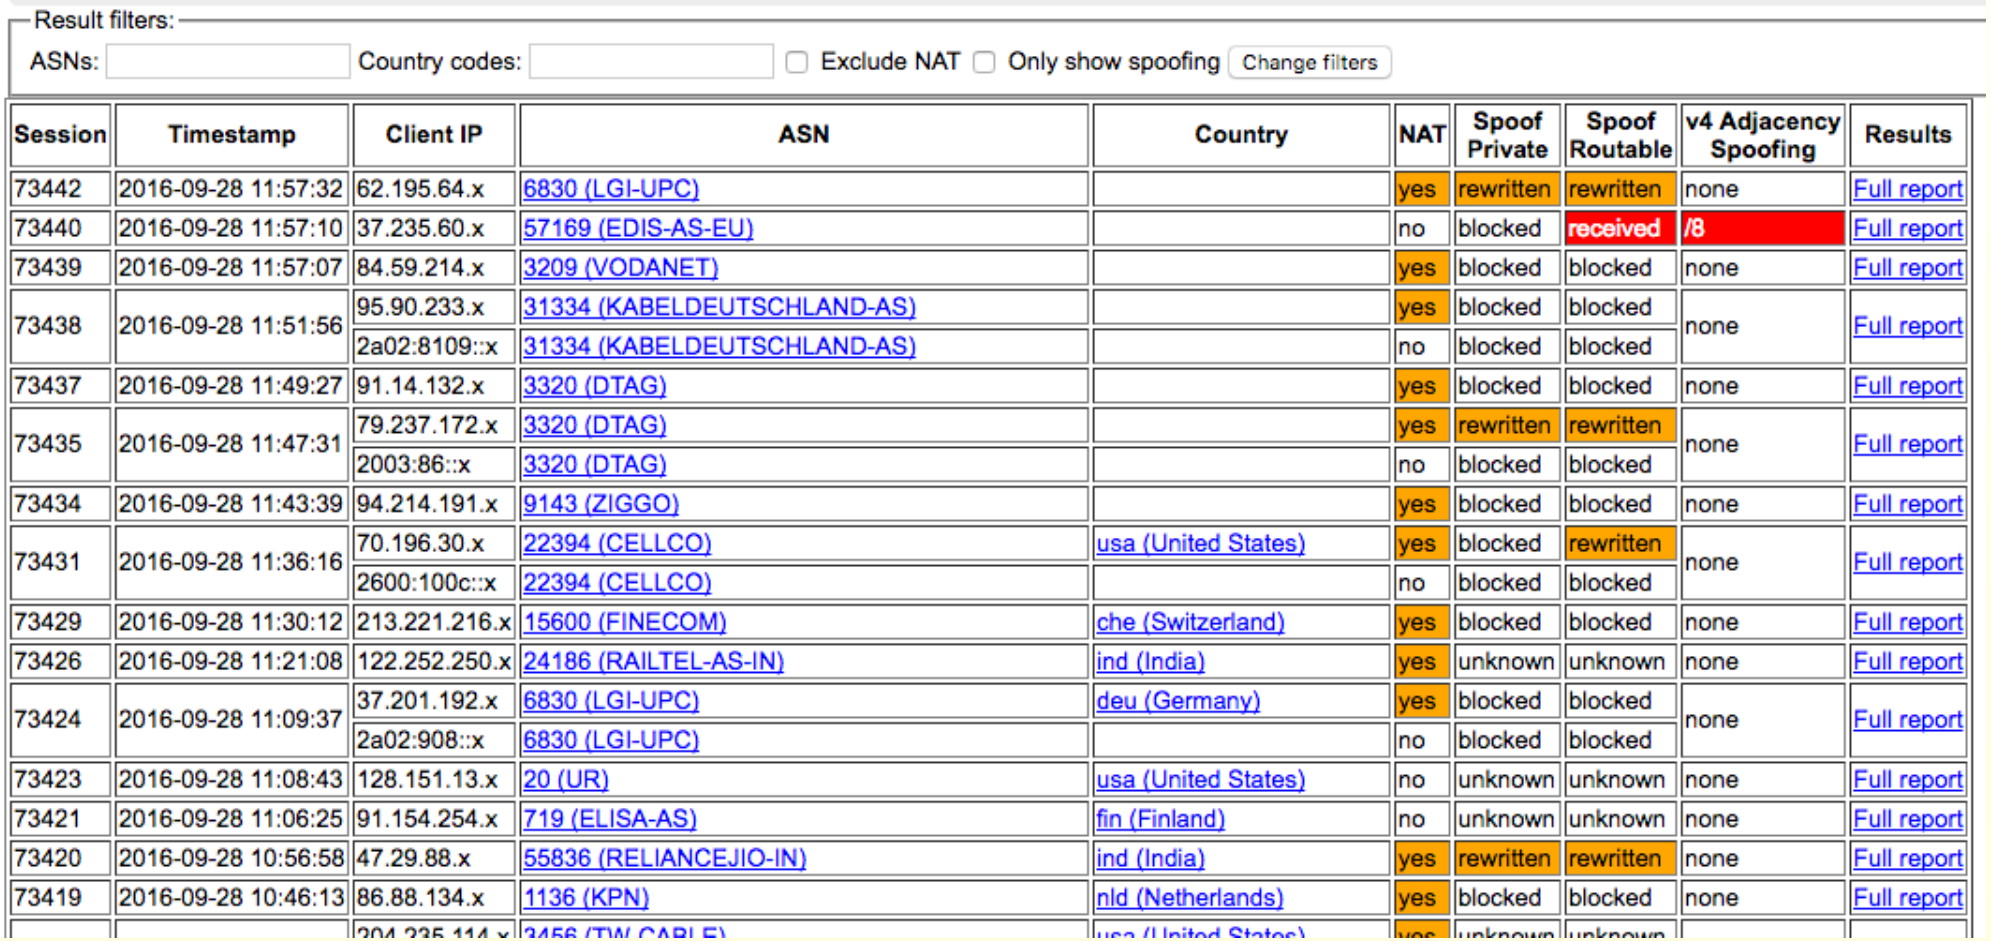

Result filters:

ASNs: Country codes: ☐ Exclude NAT ☐ Only show spoofing

Session	Timestamp	Client IP	ASN	Country	NAT	Spoof Private	Spoof Routable	v4 Adjacency Spoofing	Results
73442	2016-09-28 11:57:32	62.195.64.x	6830 (LGI-UPC)		yes	rewritten	rewritten	none	Full report
73440	2016-09-28 11:57:10	37.235.60.x	57169 (EDIS-AS-EU)		no	blocked	received	/8	Full report
73439	2016-09-28 11:57:07	84.59.214.x	3209 (VODANET)		yes	blocked	blocked	none	Full report
73438	2016-09-28 11:51:56	95.90.233.x	31334 (KABELDEUTSCHLAND-AS)		yes	blocked	blocked	none	Full report
		2a02:8109::x	31334 (KABELDEUTSCHLAND-AS)		no	blocked	blocked		
73437	2016-09-28 11:49:27	91.14.132.x	3320 (DTAG)		yes	blocked	blocked	none	Full report
73435	2016-09-28 11:47:31	79.237.172.x	3320 (DTAG)		yes	rewritten	rewritten	none	Full report
		2003:86::x	3320 (DTAG)		no	blocked	blocked		
73434	2016-09-28 11:43:39	94.214.191.x	9143 (ZIGGO)		yes	blocked	blocked	none	Full report
73431	2016-09-28 11:36:16	70.196.30.x	22394 (CELLCO)	usa (United States)	yes	blocked	rewritten	none	Full report
		2600:100c::x	22394 (CELLCO)		no	blocked	blocked		
73429	2016-09-28 11:30:12	213.221.216.x	15600 (FINECOM)	che (Switzerland)	yes	blocked	blocked	none	Full report
73426	2016-09-28 11:21:08	122.252.250.x	24186 (RAILTEL-AS-IN)	ind (India)	yes	unknown	unknown	none	Full report
73424	2016-09-28 11:09:37	37.201.192.x	6830 (LGI-UPC)	deu (Germany)	yes	blocked	blocked	none	Full report
		2a02:908::x	6830 (LGI-UPC)		no	blocked	blocked		
73423	2016-09-28 11:08:43	128.151.13.x	20 (UR)	usa (United States)	no	unknown	unknown	none	Full report
73421	2016-09-28 11:06:25	91.154.254.x	719 (ELISA-AS)	fin (Finland)	no	unknown	unknown	none	Full report
73420	2016-09-28 10:56:58	47.29.88.x	55836 (RELIANCEJIO-IN)	ind (India)	yes	rewritten	rewritten	none	Full report
73419	2016-09-28 10:46:13	86.88.134.x	1136 (KPN)	nld (Netherlands)	yes	blocked	blocked	none	Full report
		204.235.114.x	3456 (TW-CABLE)	usa (United States)	yes	unknown	unknown		

http://spoofer.caida.org/recent_tests.php

Spoofer request for help

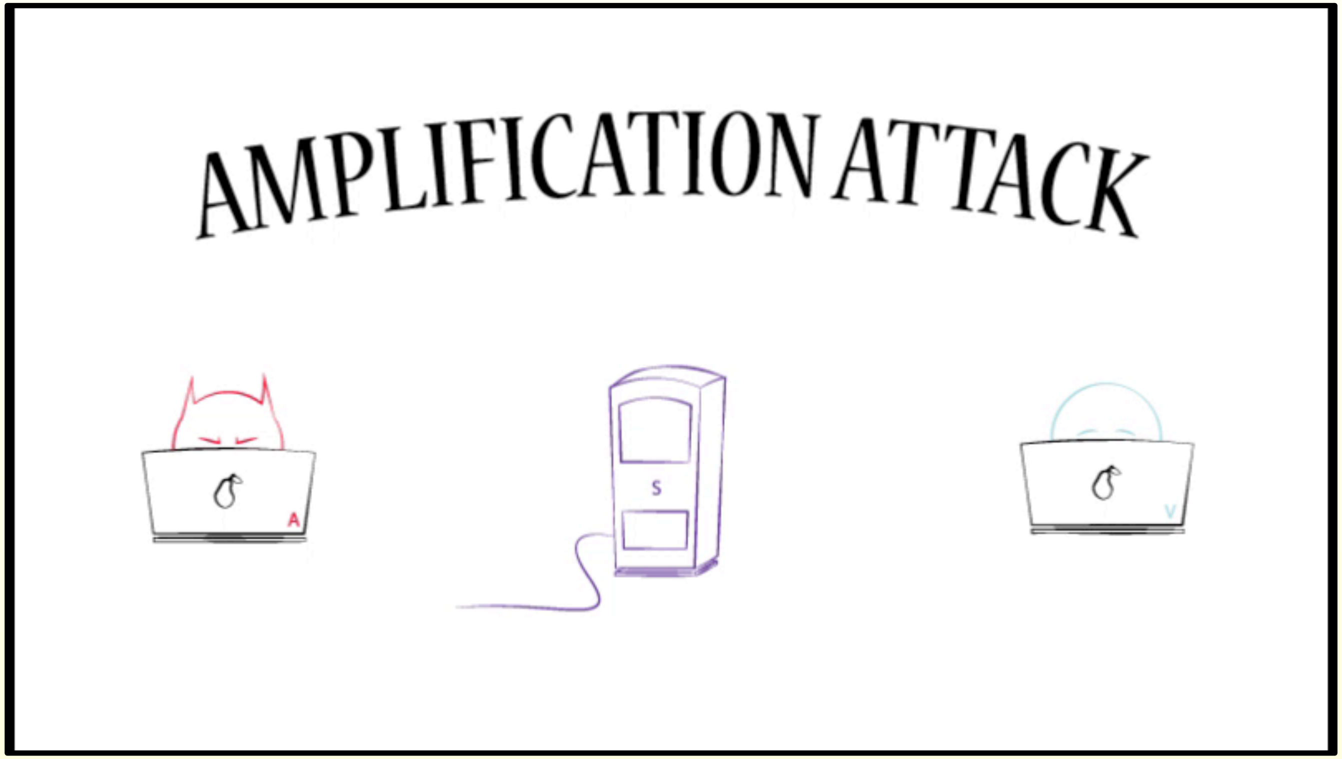

The video will explain to a general audience the dangers of IP spoofing.

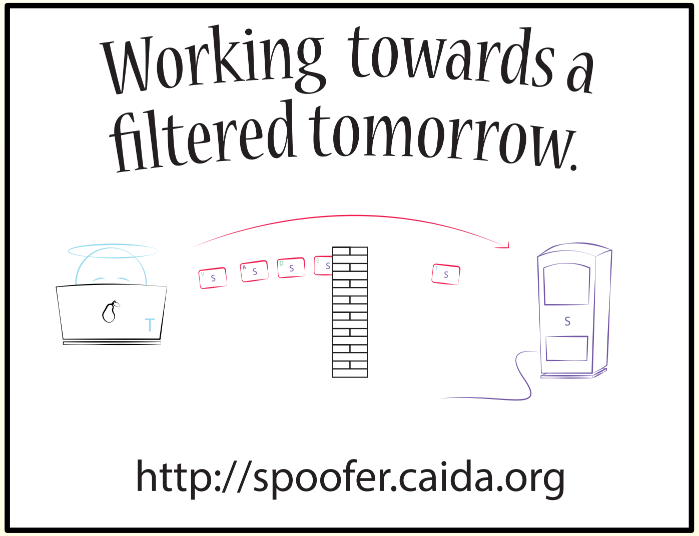

We will end the video with a requester help.

Science of Internet Security: Technology

- Using the versatile Ark measurement platform, we will conduct measurements and analysis for documented explanations of structural and dynamic aspects of the Internet infrastructure relevant to cybersecurity vulnerabilities
  - Task 1: Support for Macroscopic Security and Stability Monitoring and Analysis
  - Task 2: Mapping Peering Interconnections at the Router Level
  - Task 3: Mapping Peering Interconnections at the Facility Level
  - Task 4: Measurements of TCP Behavior to Understand Security Vulnerabilities
  - Task 5: Identifying Grey Market IPv4 Address Transfers
  - Task 6: Internet Router-Level Topology Mapping on Demand

- Task 1:
  - IPv4 Prefix-Probing Dataset
    http://www.caida.org/data/active/ipv4_prefix_probing_dataset.xml
- Task 2:
  - AS Border Mapping Dataset (February 2017)
- Task 3:
  - AS to Facilities Mapping Dataset (February 2017)
  - AS to Facilities Mapping Dataset annotated w/ approach to interconnection (private peering with cross-connect, public peering, private interconnects over the public switch fabric, and remote peering) (February/March 2017)
  - Alias resolved Interconnection (router-level map) (April 2017)
  - Global facility-aware map of interconnection (May 2017)

Detection and analysis of large-scale Internet infrastructure outages (IODA)

- Developing methods to infer location and extent of outages
- **Goals: (1) investigate and define strategies and methodologies to fuse diverse data sources to detect & characterize outages, (2) define and refine *system* requirements for continuous monitoring & (near) real-time analysis (3) testing & experimental deployment**
- Part of a 3 year NSF-funded SATC project

Detection and analysis of large-scale Internet infrastructure outages (IODA)

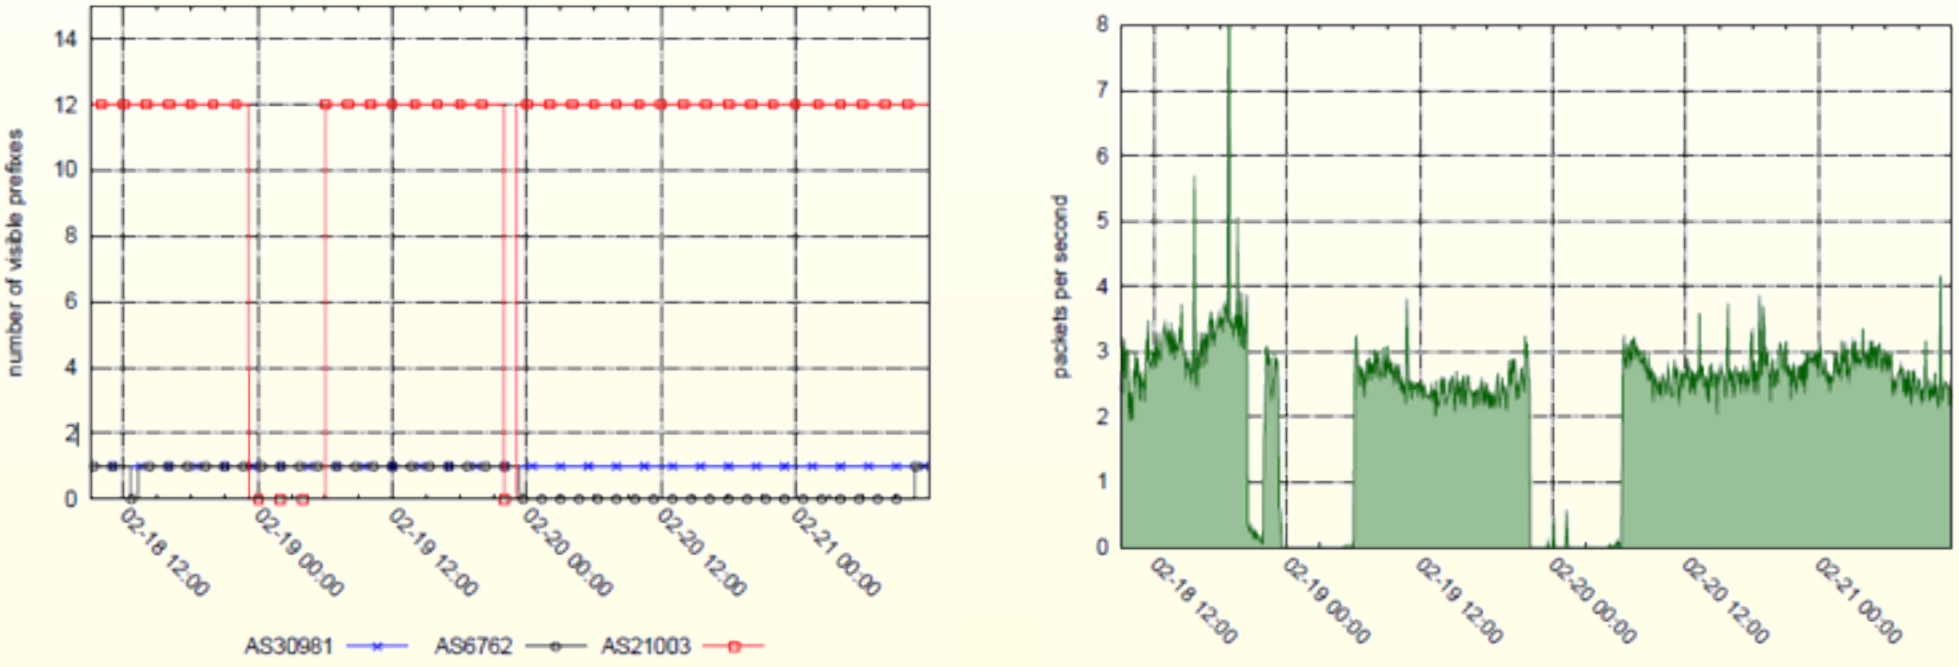

(a)

(b)

Libyan outages: (a) visibility of Libyan IPv4 prefixes in BGP (RouteViews,RIPE NCC RIS);

(b) unsolicited traffic to UCSD telescope from Libya.

IODA After Four Years (Today)

- Live detection and monitoring

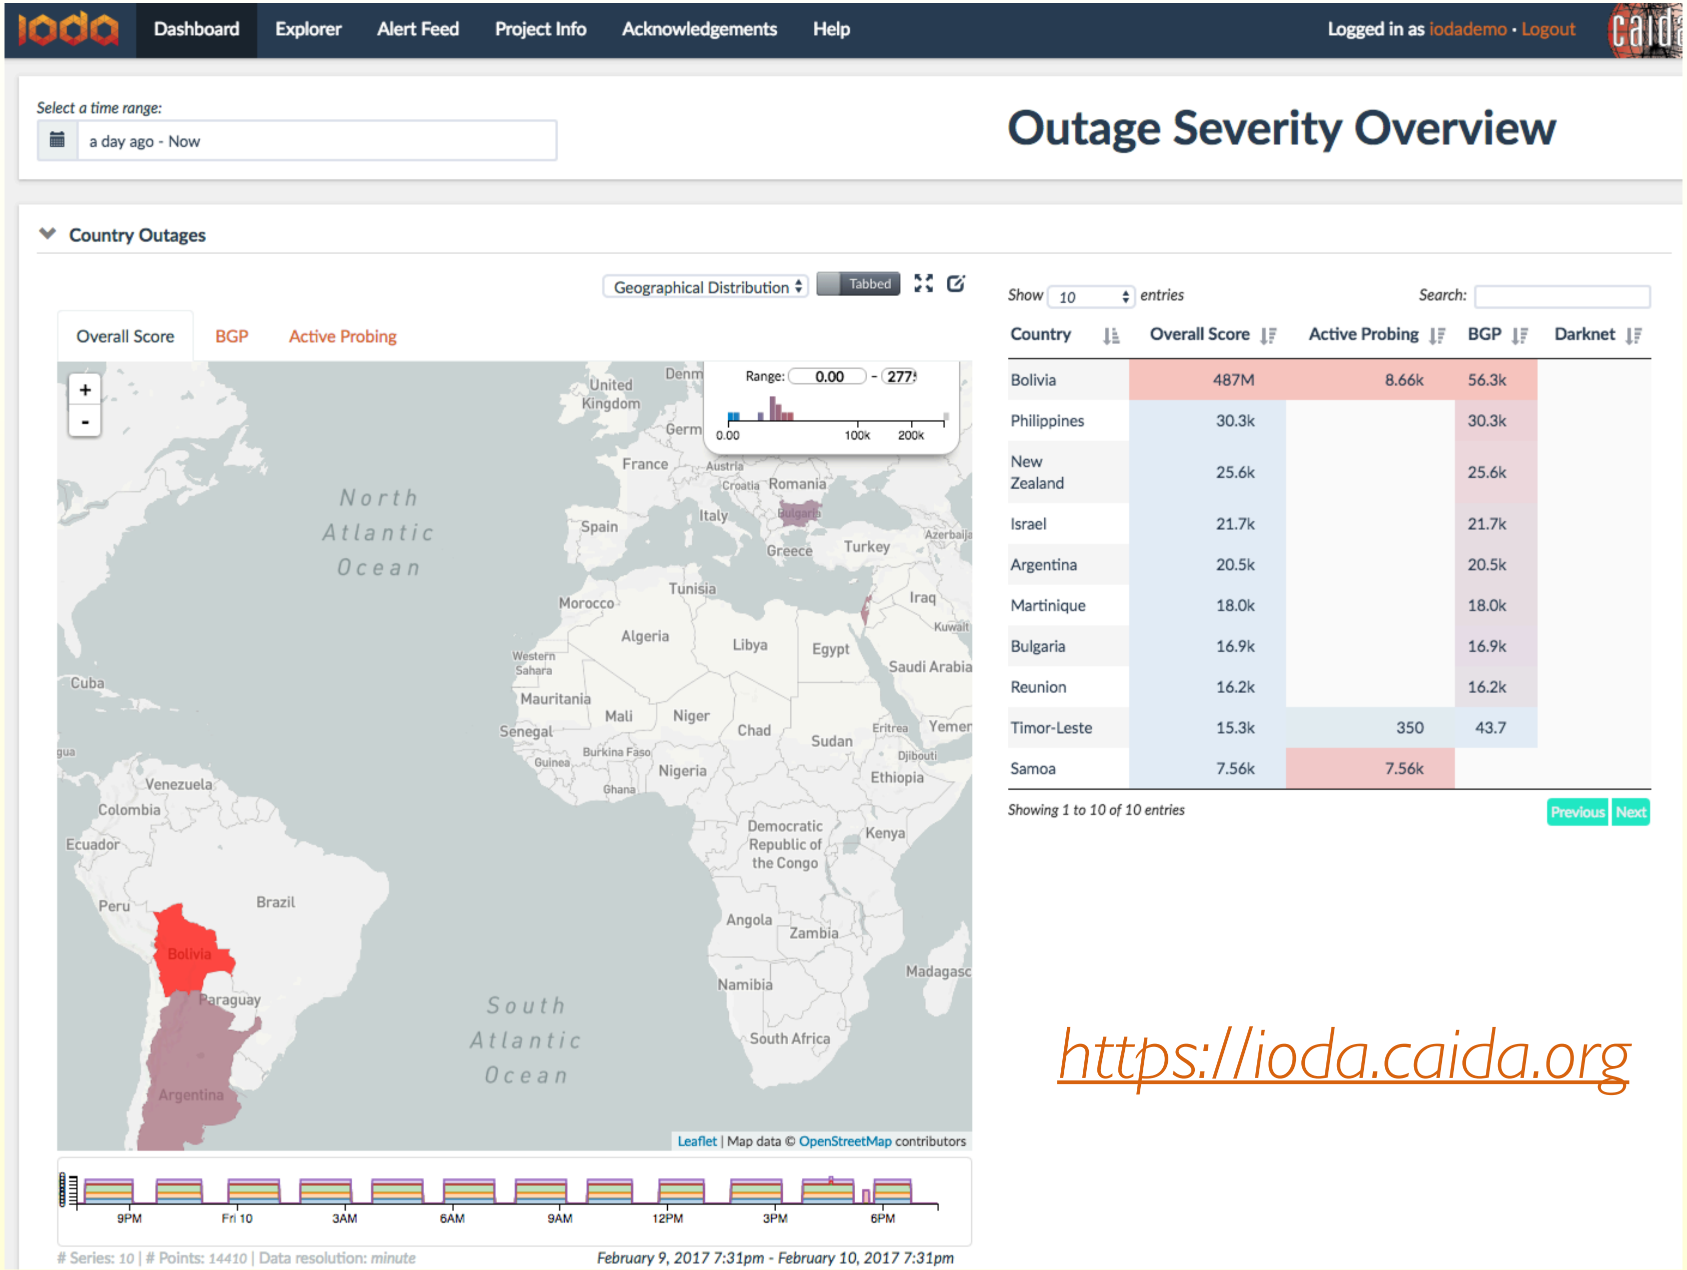

Select a time range: a day ago - Now

Outage Severity Overview

Country Outages

Series: 10 | # Points: 14410 | Data resolution: minute

February 9, 2017 7:31pm - February 10, 2017 7:31pm

Show 10 entries Search:

Country	Overall Score	Active Probing	BGP	Darknet
Bolivia	487M	8.66k	56.3k	
Philippines	30.3k		30.3k	
New Zealand	25.6k		25.6k	
Israel	21.7k		21.7k	
Argentina	20.5k		20.5k	
Martinique	18.0k		18.0k	
Bulgaria	16.9k		16.9k	
Reunion	16.2k		16.2k	
Timor-Leste	15.3k	350	43.7	
Samoa	7.56k	7.56k		

Showing 1 to 10 of 10 entries

https://ioda.caida.org

IODA City Map

- High-level system view

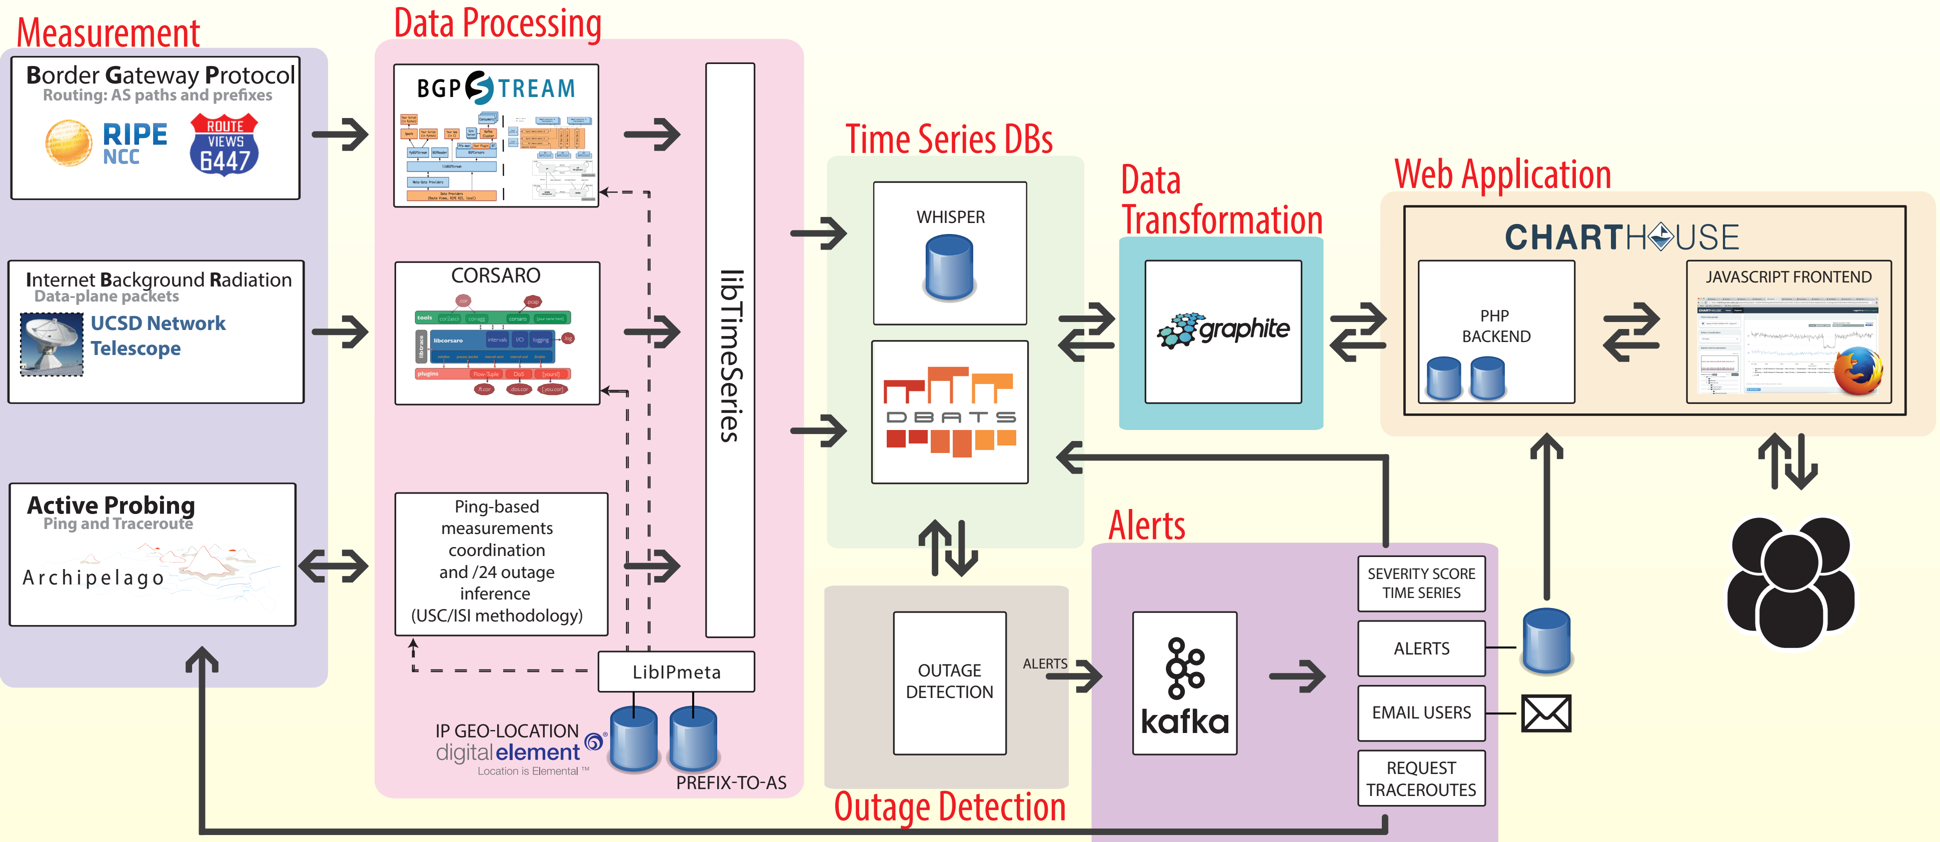

Mapping Interdomain Internet Congestion

- Developing methods to measure the location and extent of interdomain congestion

- **Goals (1) system to monitor interdomain links and their congestion state, (2) near real-time "congestion heat map" of the Internet, (3) increase transparency, empirical grounding of debate**

- Part of a 3 year NSF-funded project on topology+congestion

Measurement System

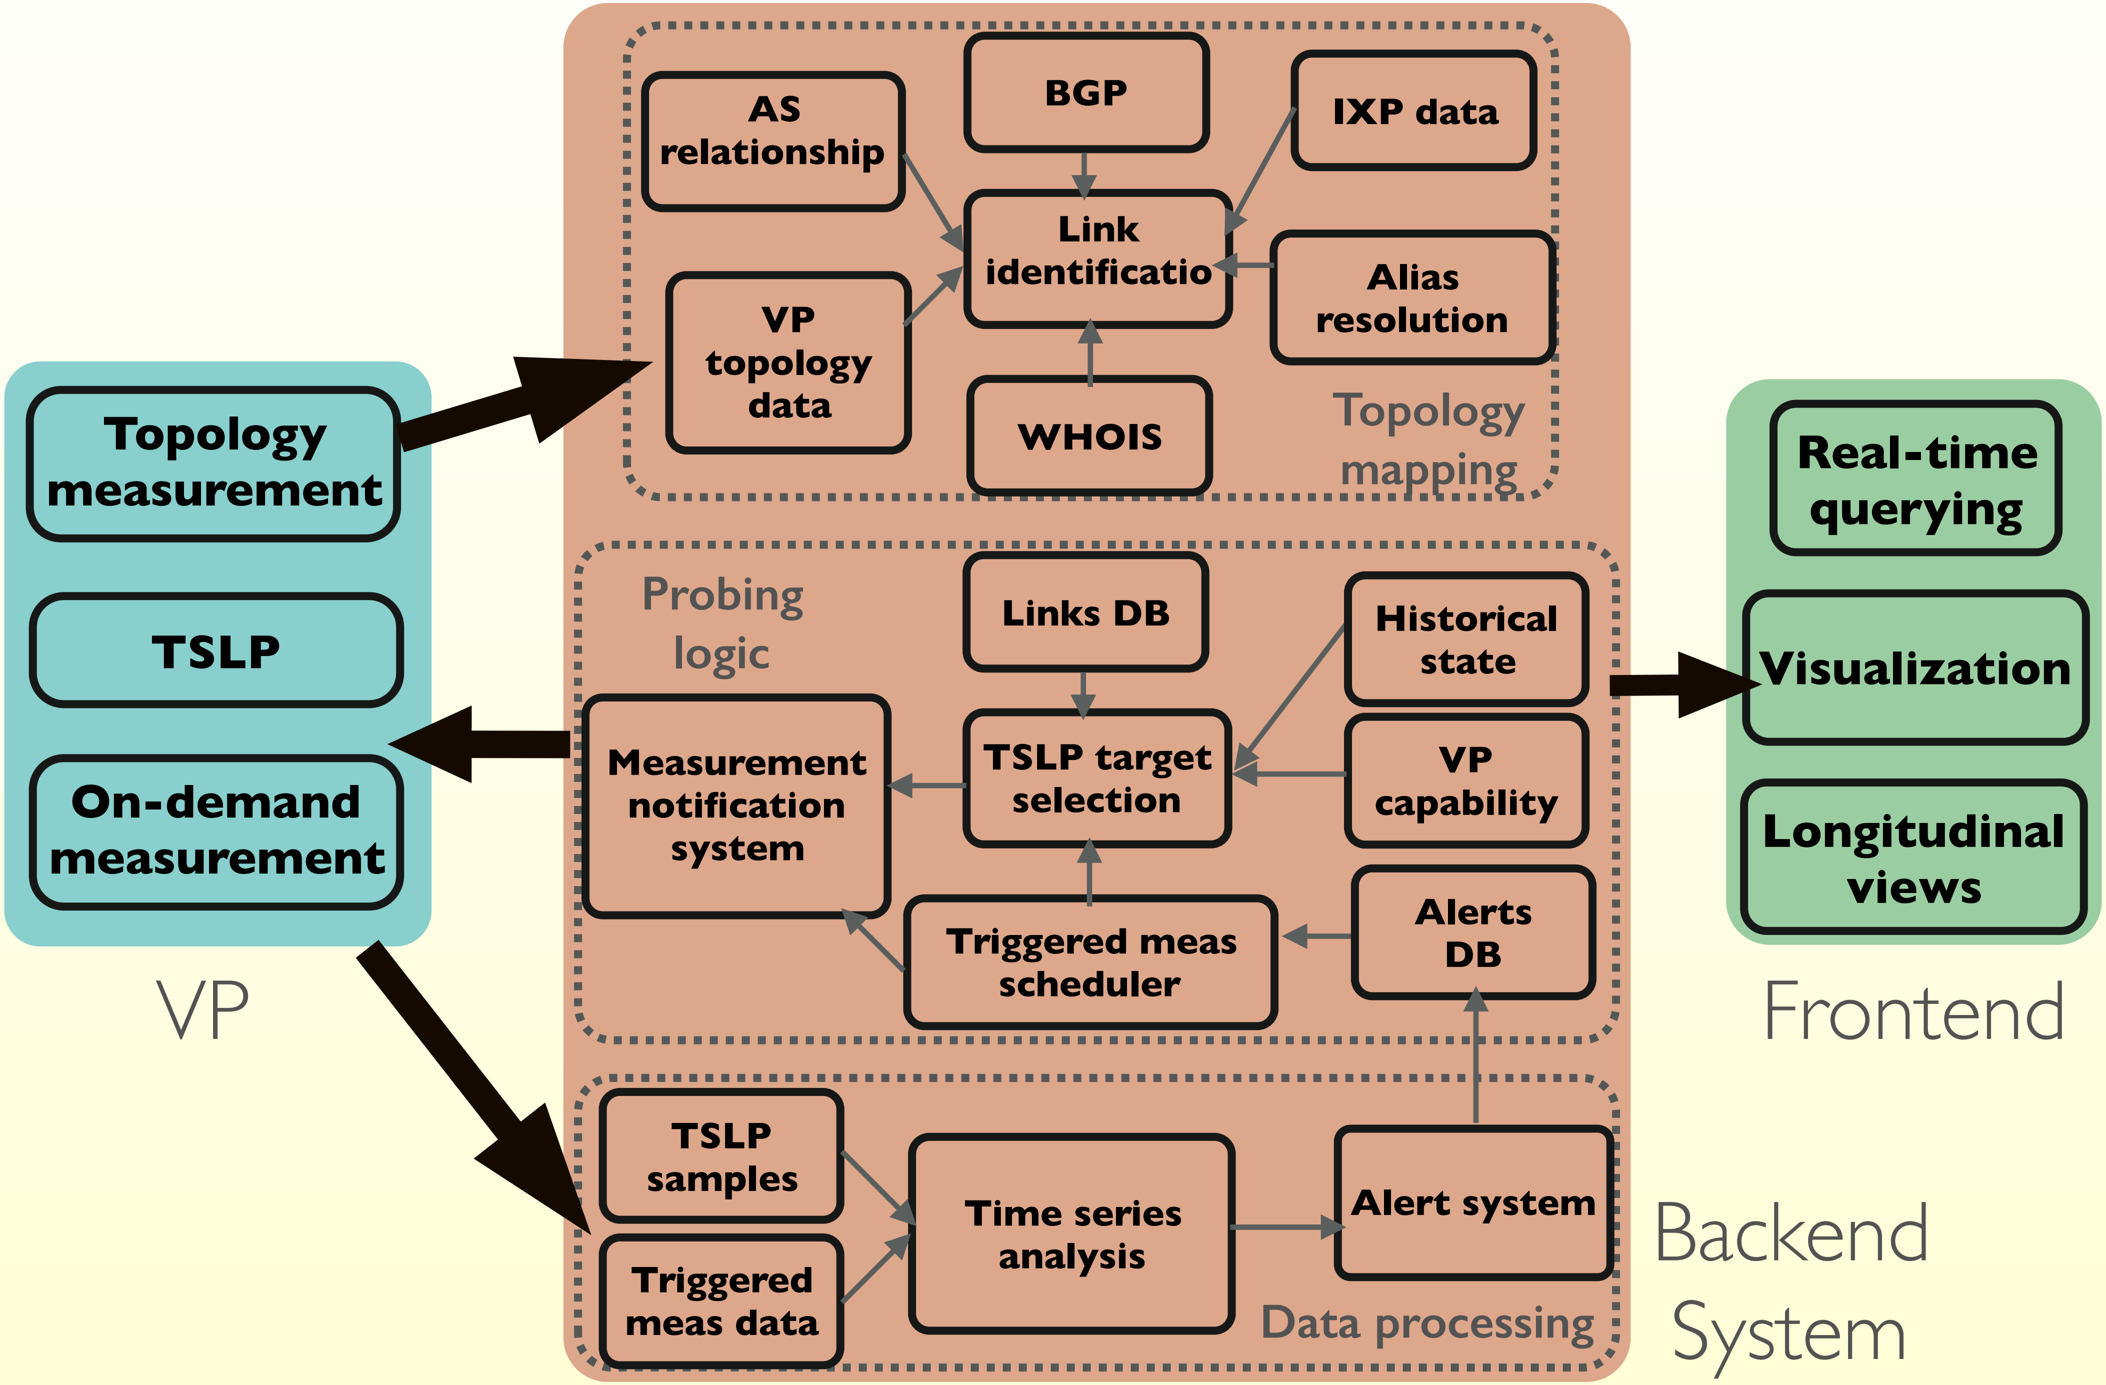

Mapping Interdomain Internet Congestion

Congestion seen between Comcast

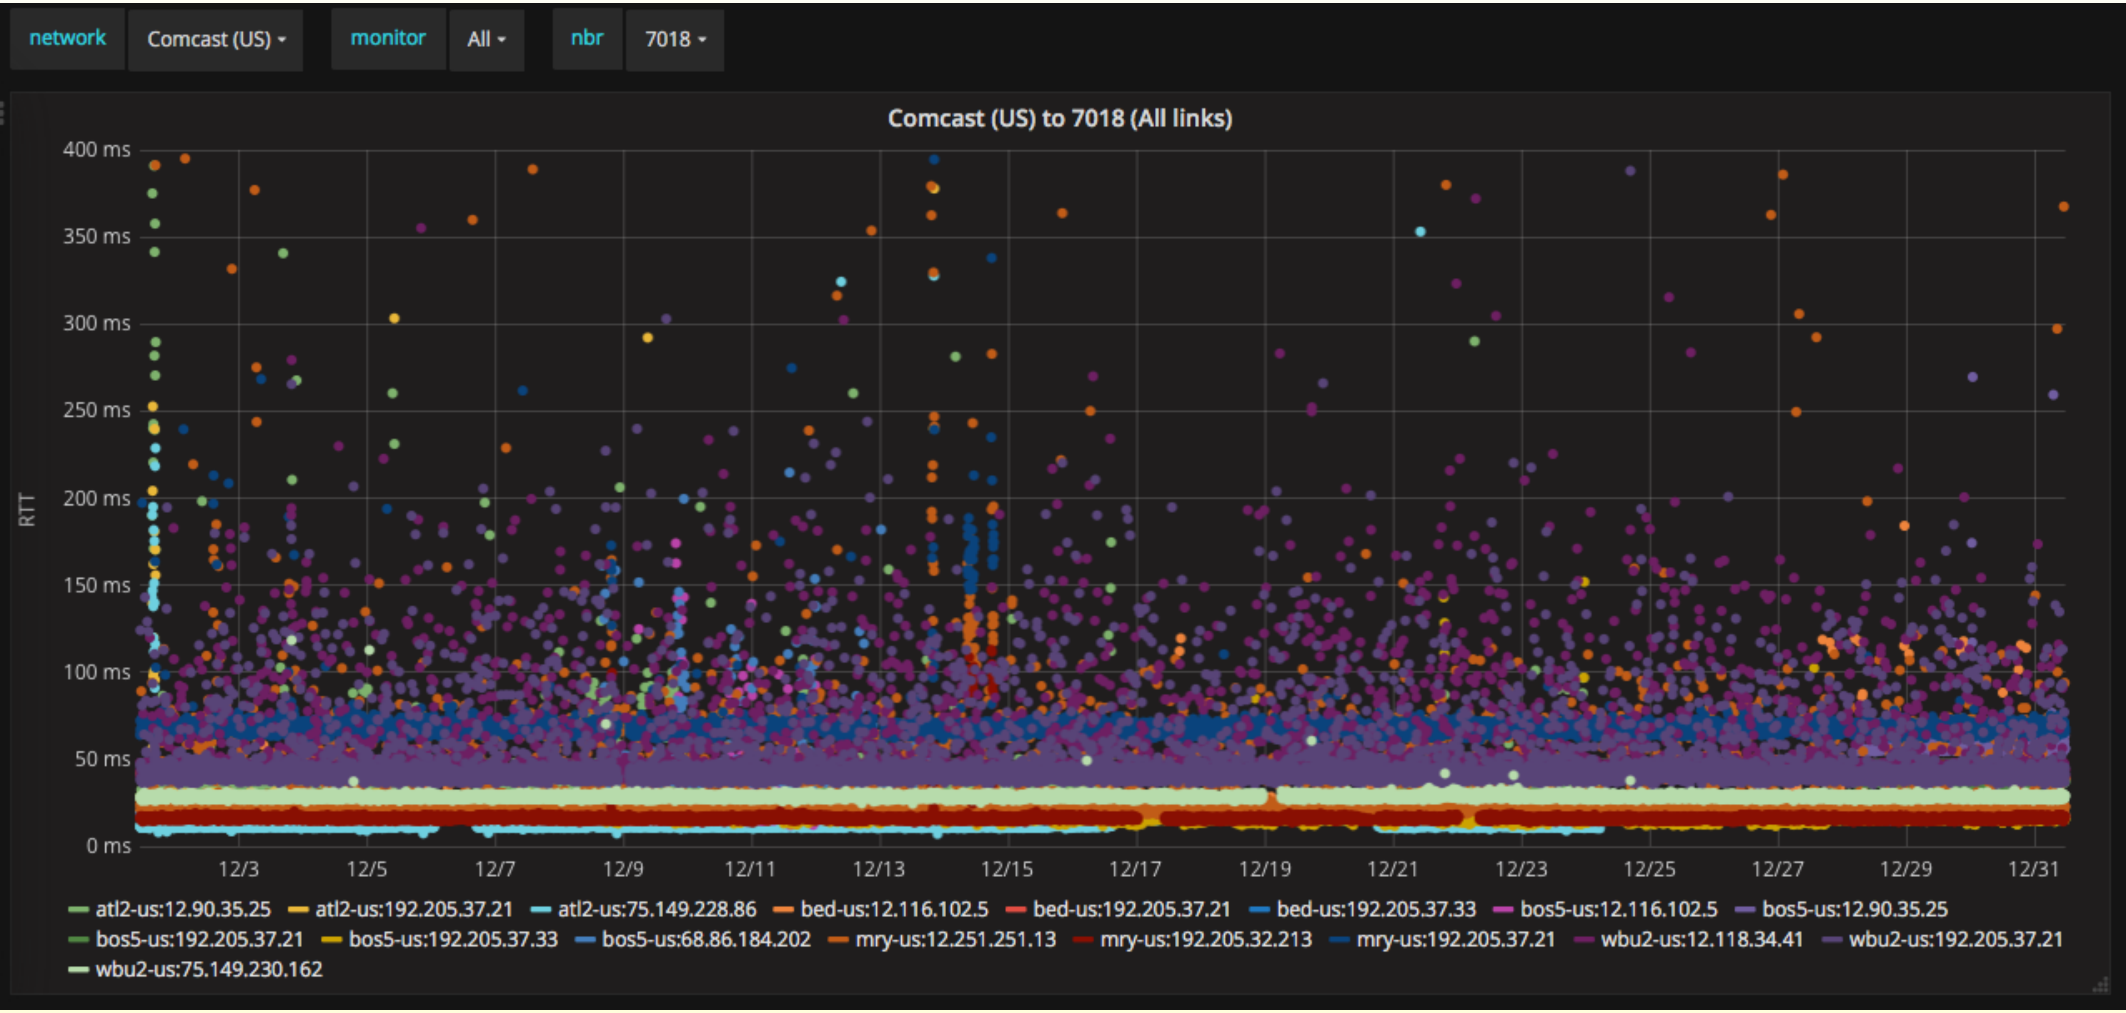

network Comcast (US)

monitor All

nbr 7018

Contact Information

PI: k claffy, CAIDA
kc@caida.org
http://www.caida.org/