CAIDA update

PI k claffy, CAIDA
University of Wisconsin
Madison, WI
29 September 2016

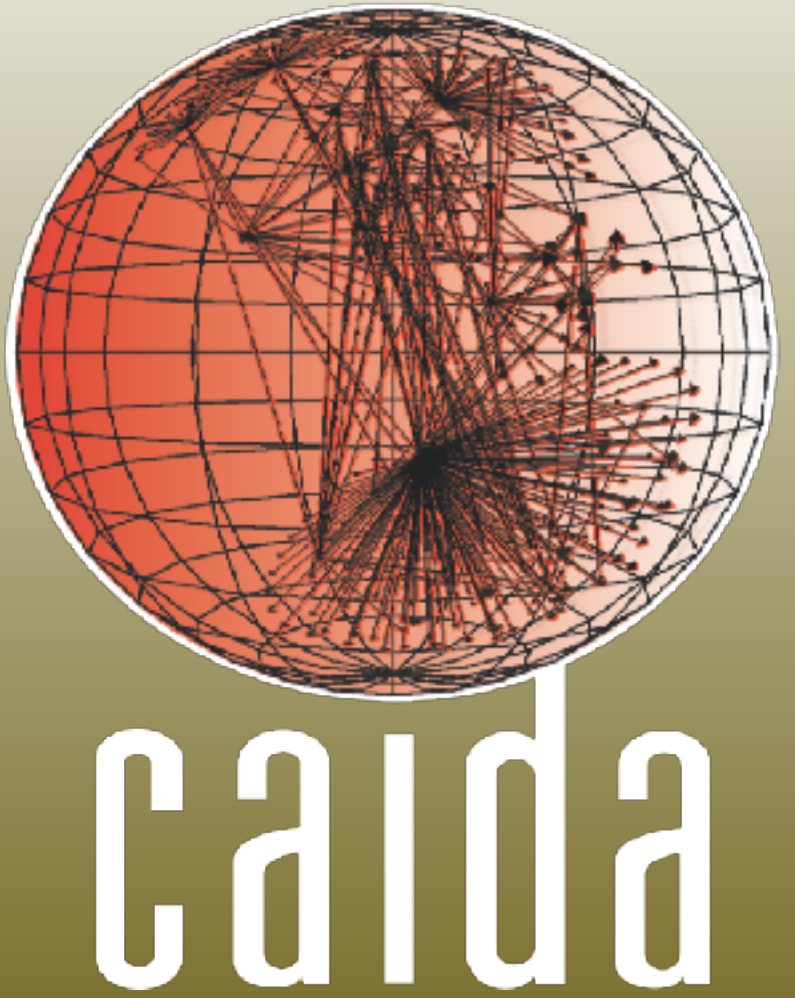

DHS IMPACT Project: CAIDA update

- 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, MOAs, Data Sets, Users

Ark Measurements

- Ongoing data collection
  - IPv4 and IPv6 topology
  - spoofer
  - (evolving) congestion measurements
- Ark Platform (as of Sept 2016)
  - 163 monitors in 57 countries
  - 71 IPv6-enabled
  - 116 Raspberry PIs

Ark Derived Data Sets

- Derived data sets
  1. IPv4 Routed /24 Topology Dataset (September 2007 - ongoing)
  2. IPv4 Routed /24 AS Links (September 2007 - ongoing)
  3. IPv6 AS Links (December 2008 - ongoing)
  4. Pv4 Routed /24 DNS Names Dataset (ongoing)
  5. IPv4 Prefix-probing dataset (ongoing)
  6. Internet Topology Data Kits (ITDK) (regular)
  7. IPv6 Topology Dataset (ongoing)
  8. IPv6 DNS Names Dataset (ongoing)
  9. AS Rank & AS Relationships
  10. **AS Relationships -- with geographic annotations** (New)

- 2016 Topology Requests
  - 29 received
  - 21 approved

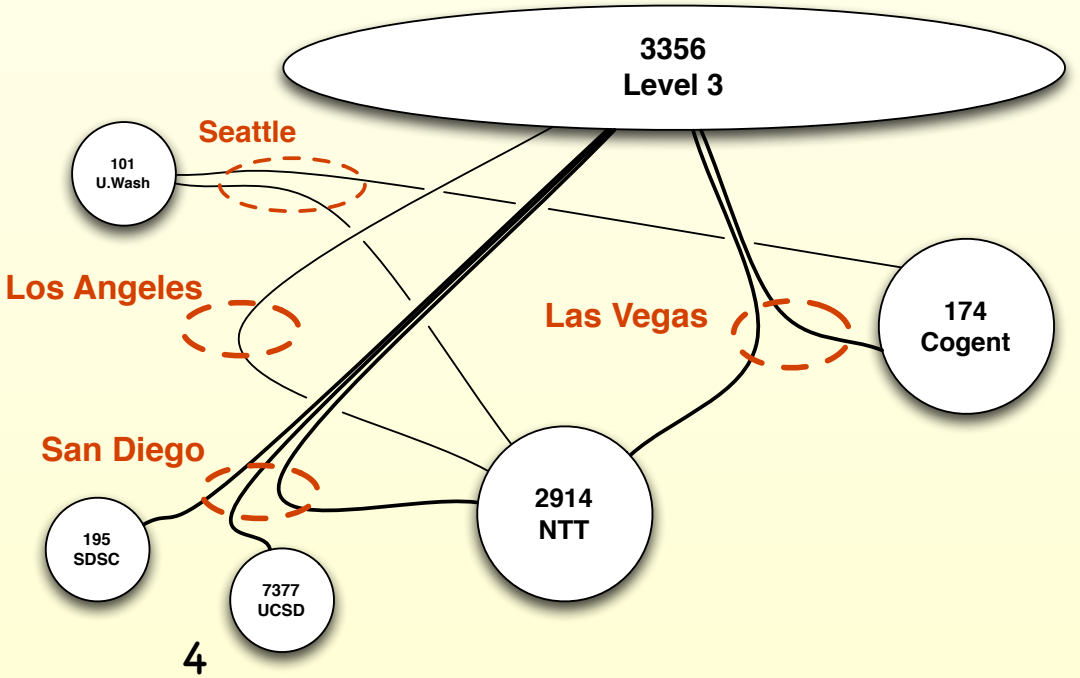

Ark Data Coverage - Sept 2007 to Present

- Duty cycle of currently active Archipelago monitors.
- IPv4 Routed /24 Topology Dataset

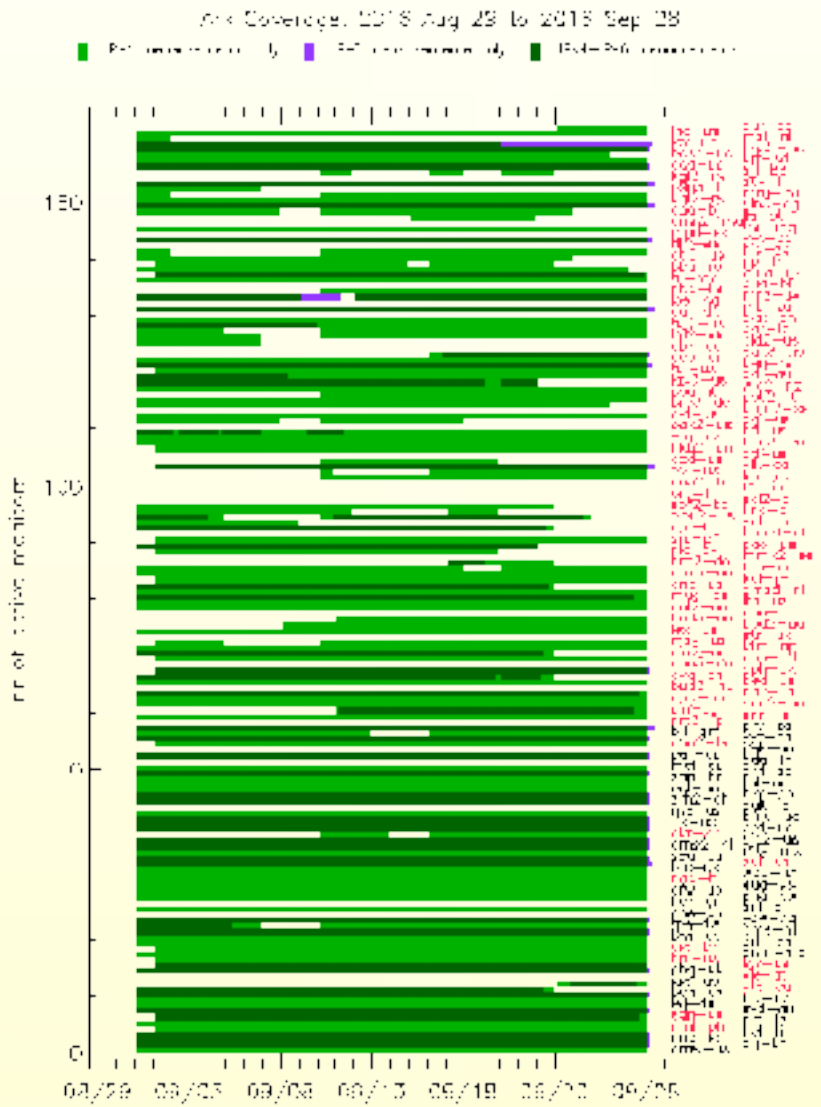

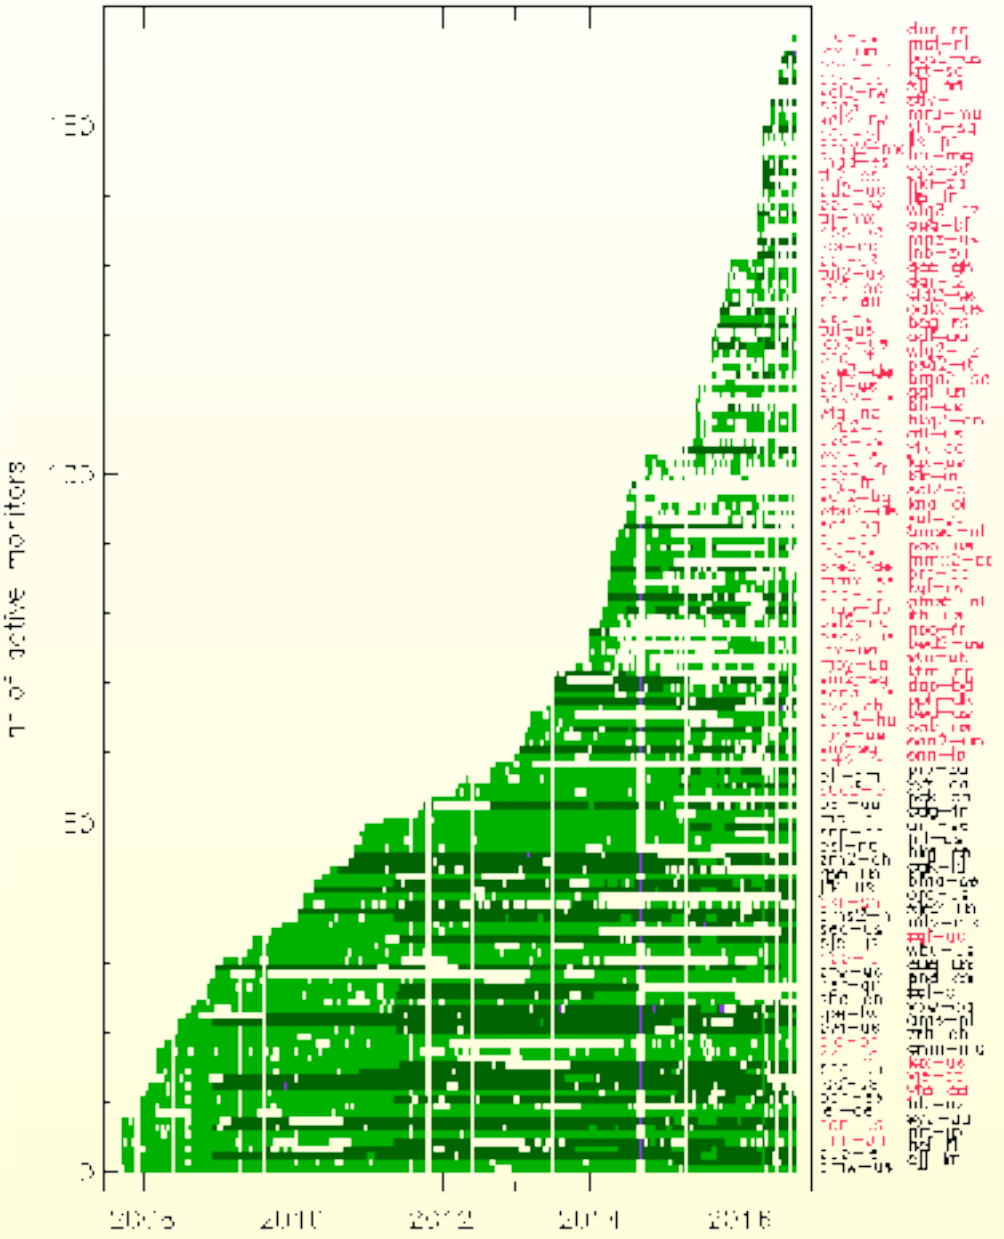

http://www.caida.org/projects/ark/coverage/

Storage Size of Ark Data Sept 2007 to Present

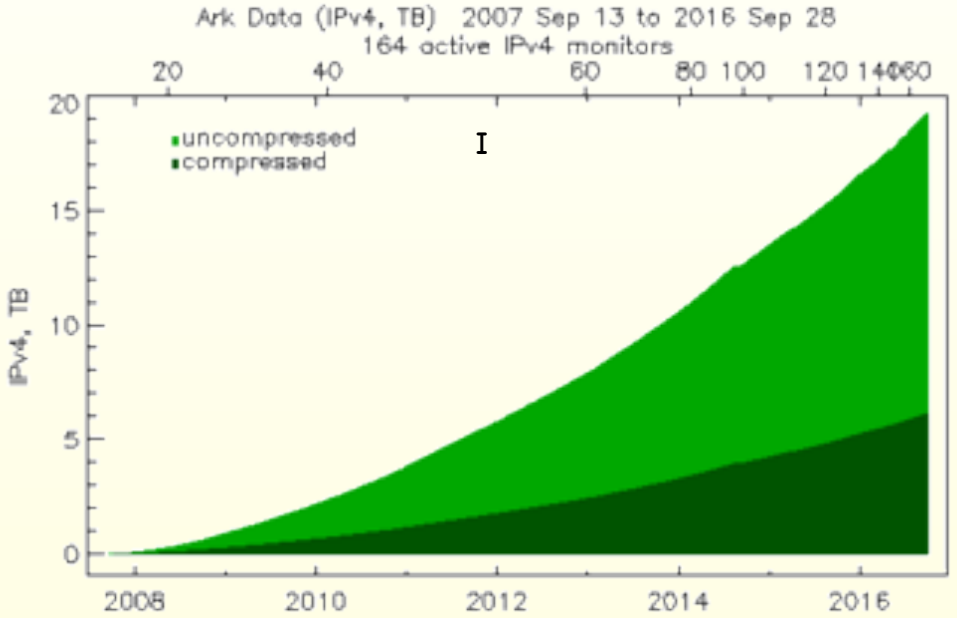

IPv4

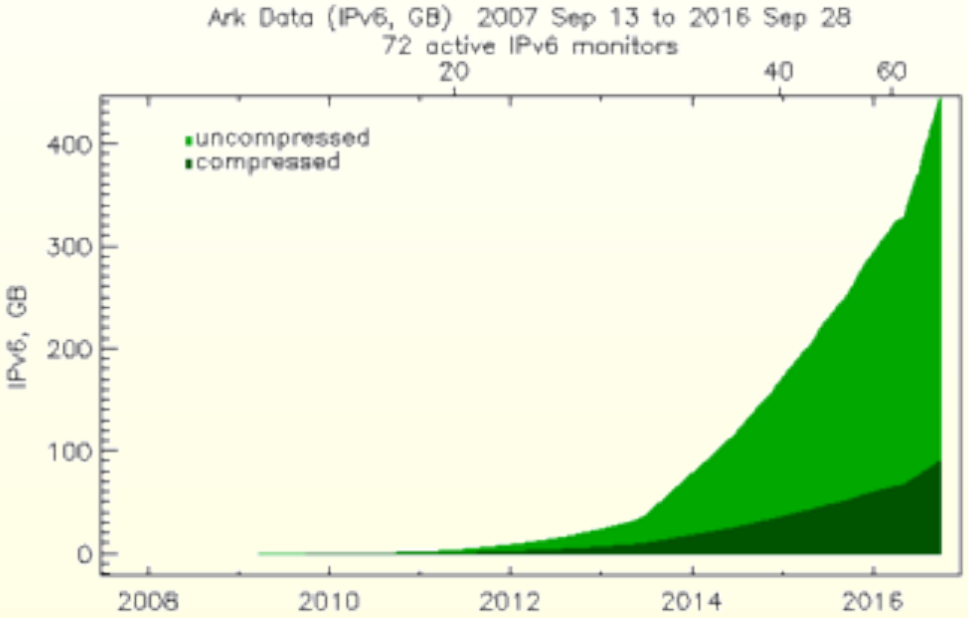

IPv6

http://www.caida.org/projects/ark/coverage/

Ongoing Measurements: Unsolicited Traffic (aka IBR)

- UCSD Network Telescope
  - 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
- 2016 Archival Telescope Requests
  - 13 received
  - 7 granted (directly from CAIDA)
  - 9 received and approved via IMPACT
- 2016 Real-time Telescope Requests
  - 8 received
  - 5 granted (directly from CAIDA)

Passive Trace Collection

- Passive infrastructure
  - two monitors with taps and Endace 10GE capture cards on 2 links (Equinix): San Jose and Chicago.
  - links upgraded to 100G: **lost San Jose** (Sep 2014), **lost Chicago** (May 2016). Current h/w can't do 100G
  - Trying to move Chicago monitor to 10G link in NYC
  - Still considering options

- 2016 Passive Internet Data
  - Monthly traces Jan-April 2016 from Chicago monitor
- 2016 Passive Requests
  - 348 received
  - 296 granted

Data Storage

- DOE's NERSC supporting archive of telescope data (no re-charge)
- Expanded CAIDA storage
  - added 125 TB of space to loki.caida.org
    - 95TB Feb 2015. 30TB Oct 2015
  - added 74TB space to junior.caida.org
    - 37TB April 2016, 37TB May2016

New Data Sets in 2016

- Topology: IPv4 Prefix-Probing
  http://www.caida.org/data/active/ipv4_prefix_probing_dataset.xml
- Anonymized Internet Traces 2016
  http://www.caida.org/data/passive/passive_2016_dataset.xml
- AS Relationships — with geographic annotations - **public**
  https://www.caida.org/data/as-relationships-geo/
- 2016-03 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)

Restricted Dataset Requests, 2016

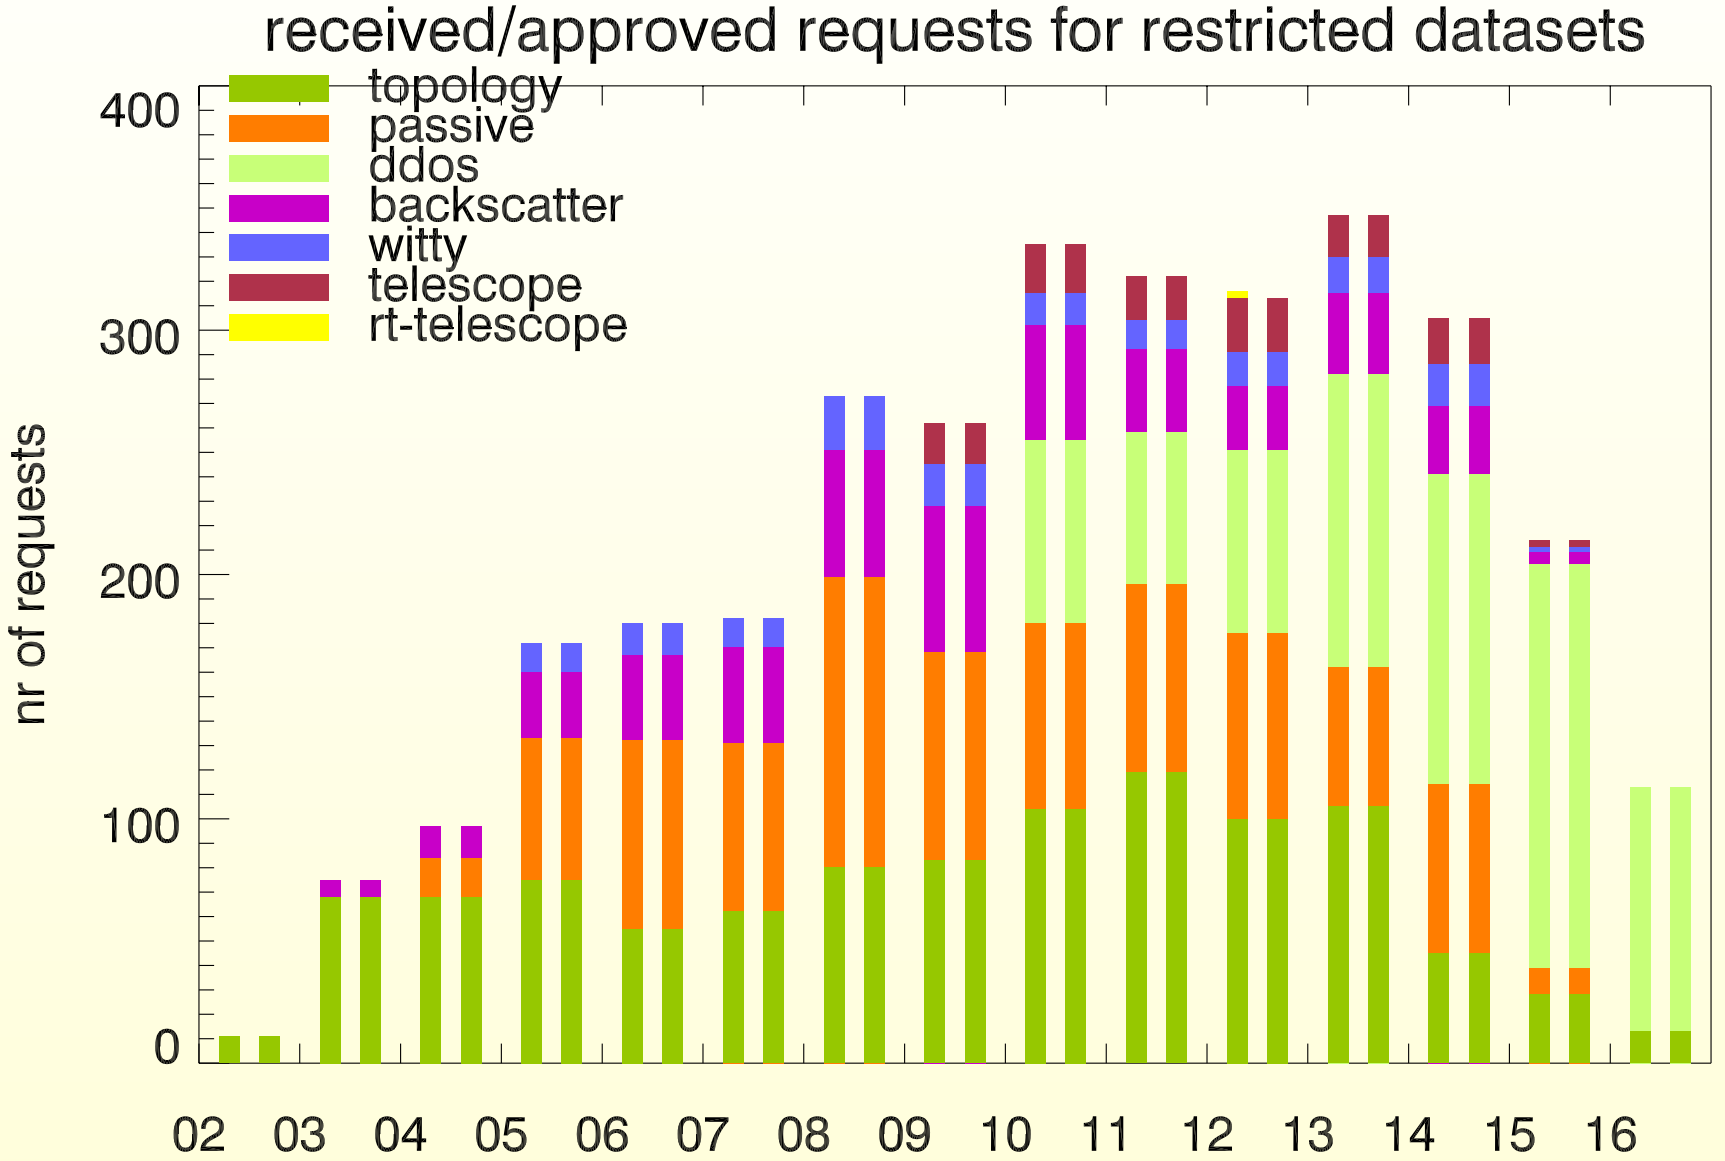

- drop in requests continues 2015-2016 due to making topology data public

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

Users downloading public data

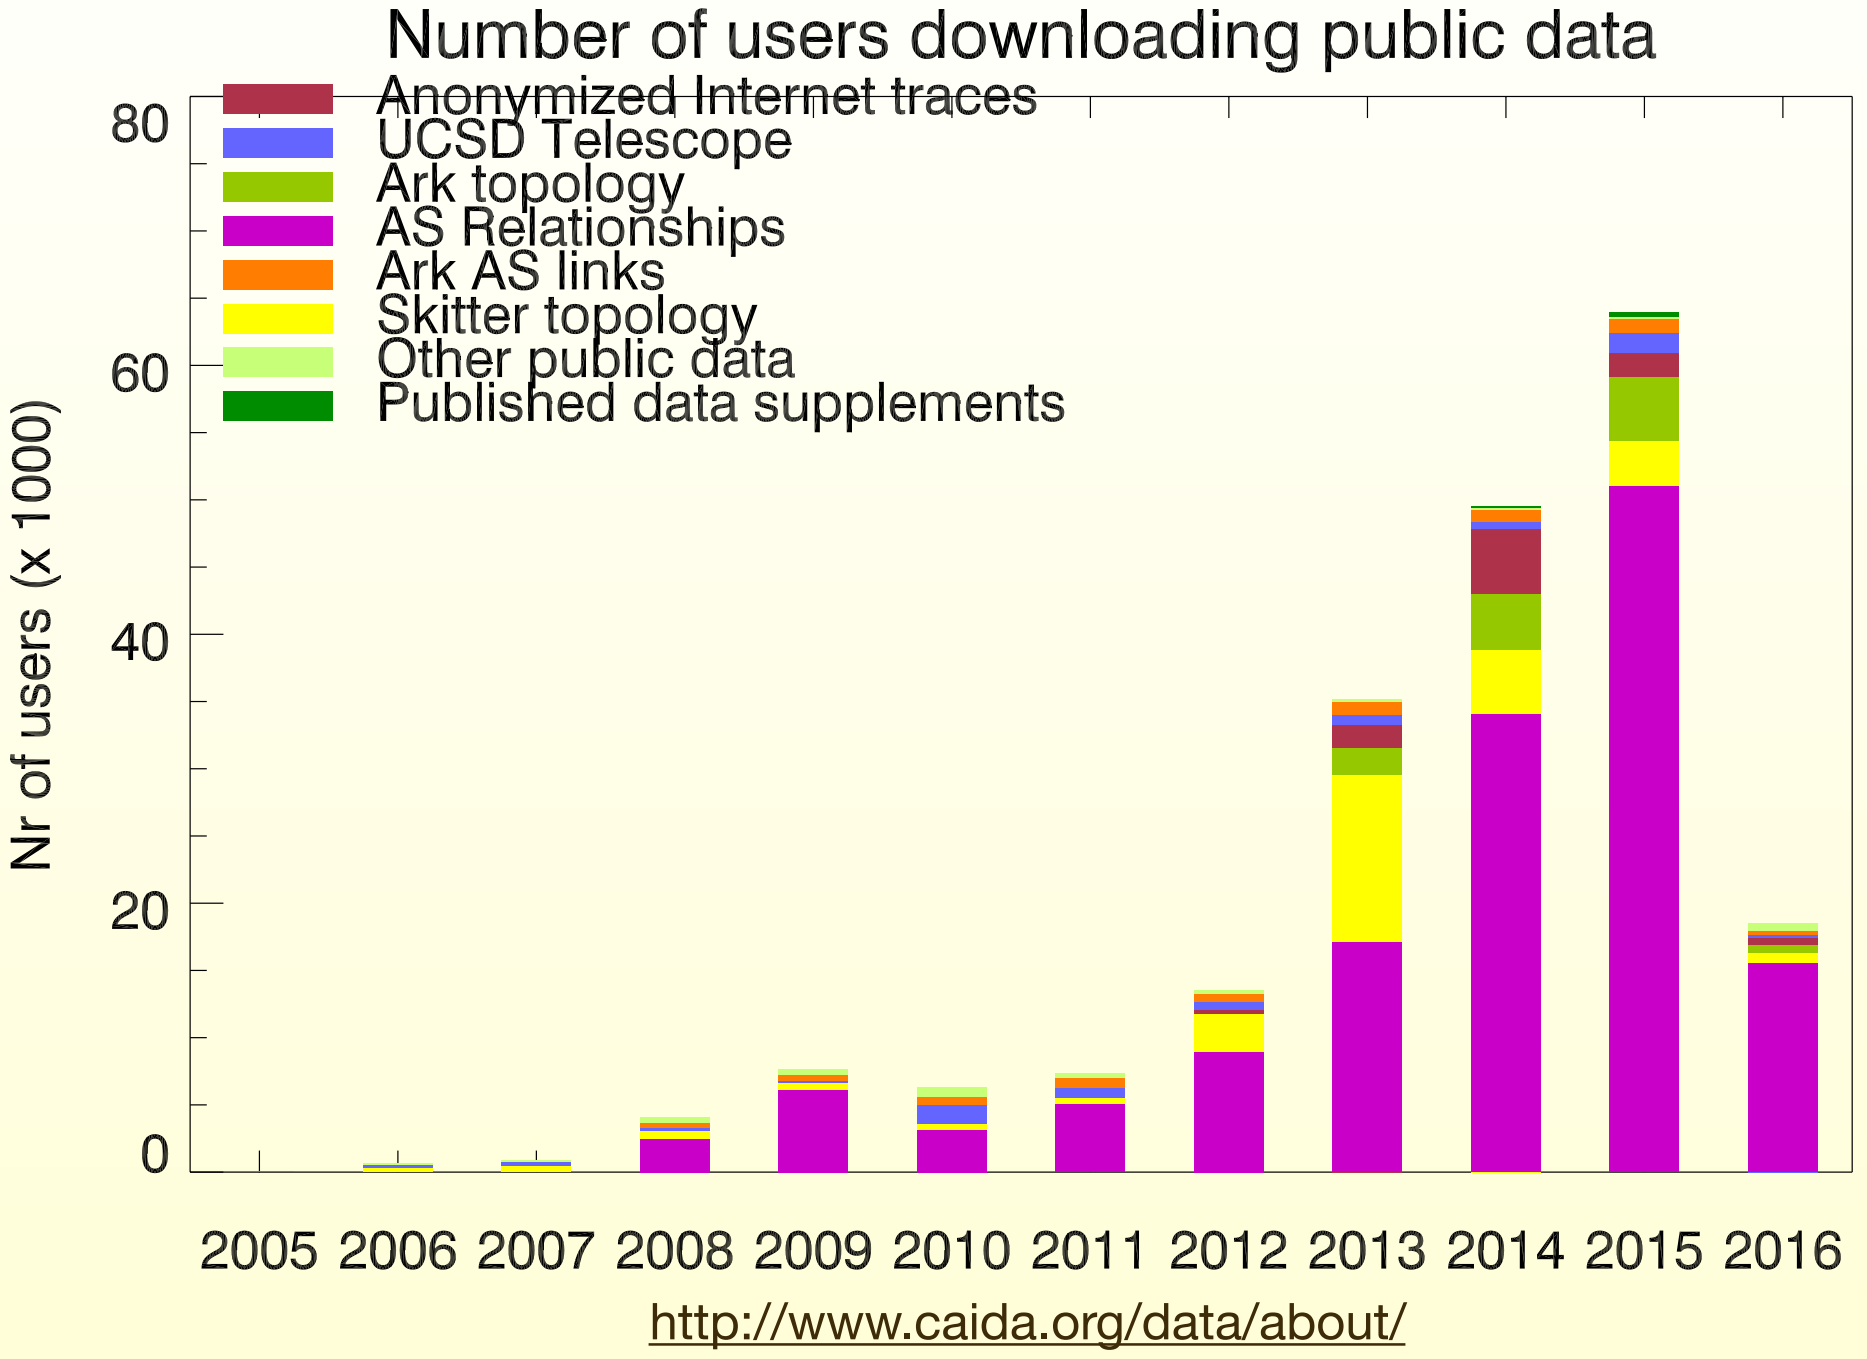

http://www.caida.org/data/about/

Users downloading restricted data

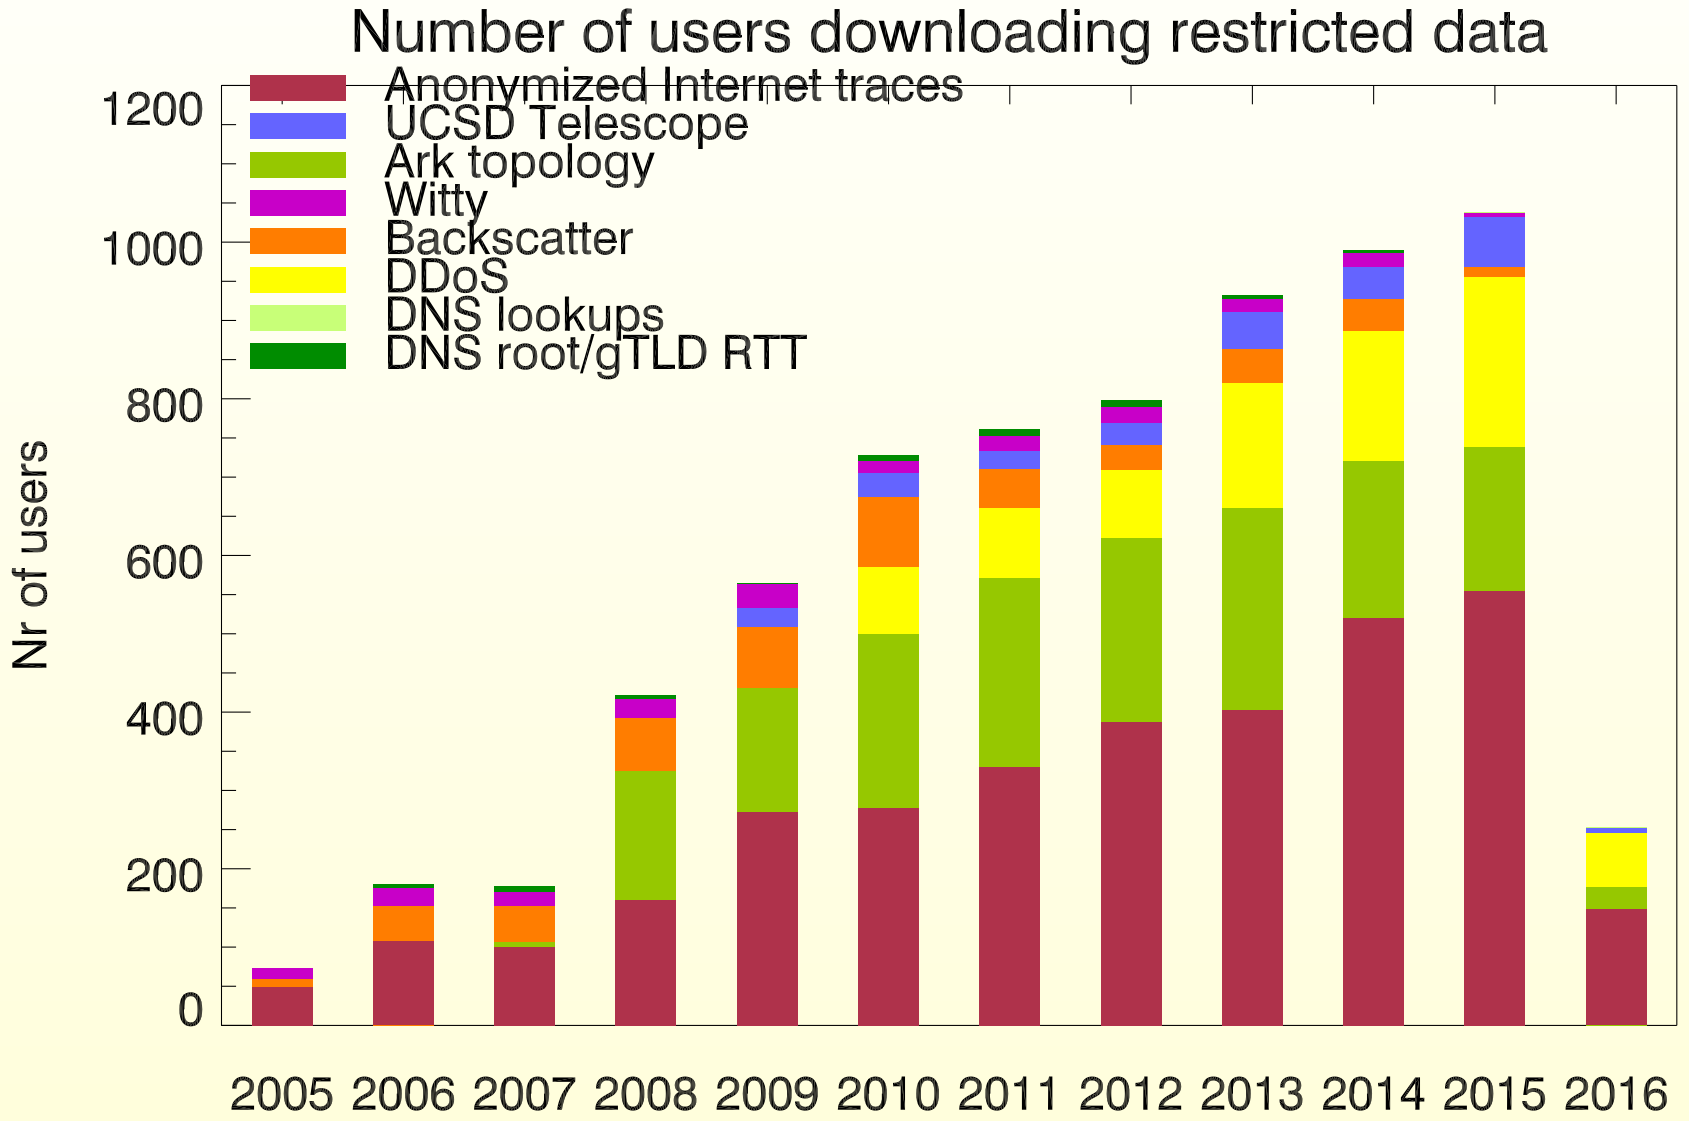

http://www.caida.org/data/about/

Public data downloaded

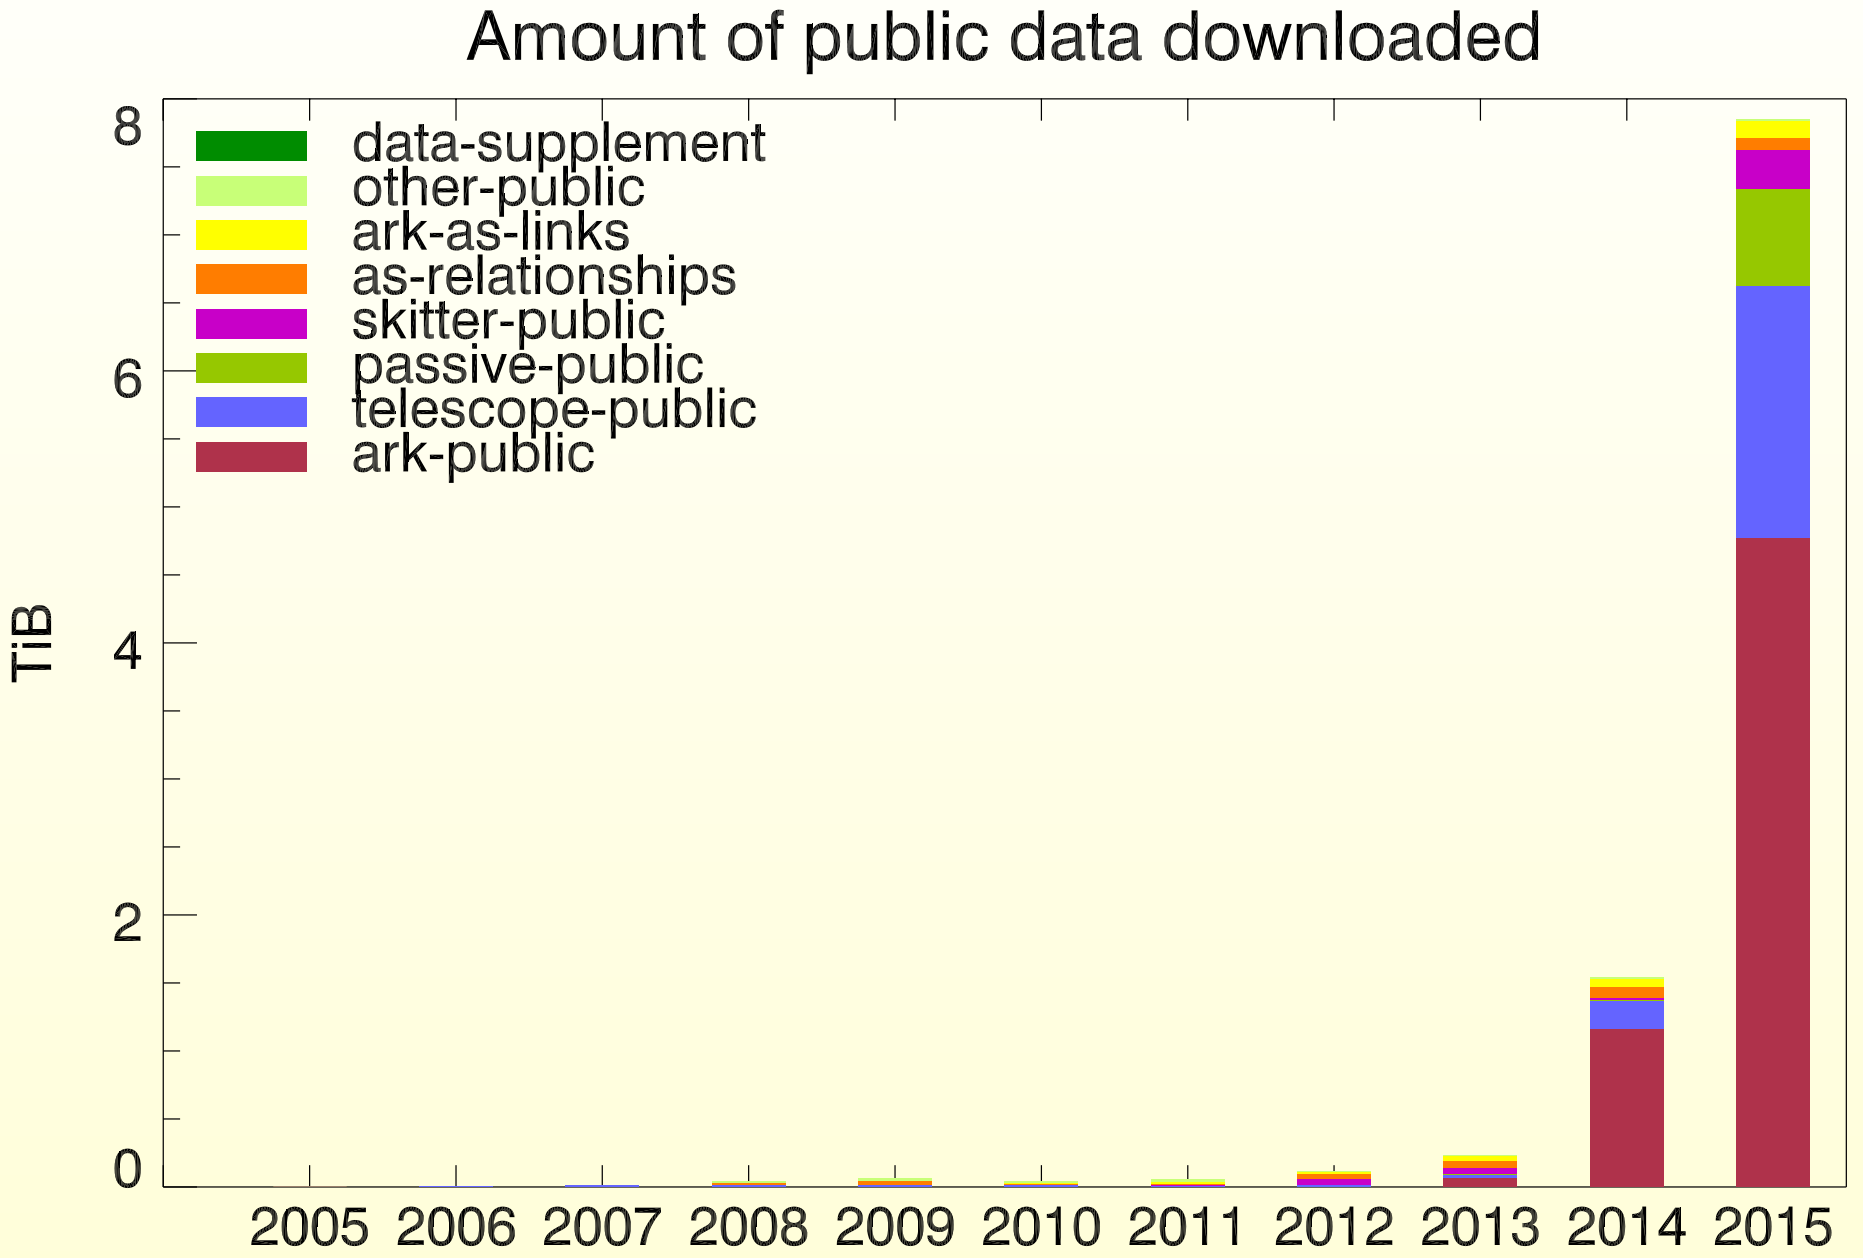

http://www.caida.org/data/about/

Restricted data downloaded

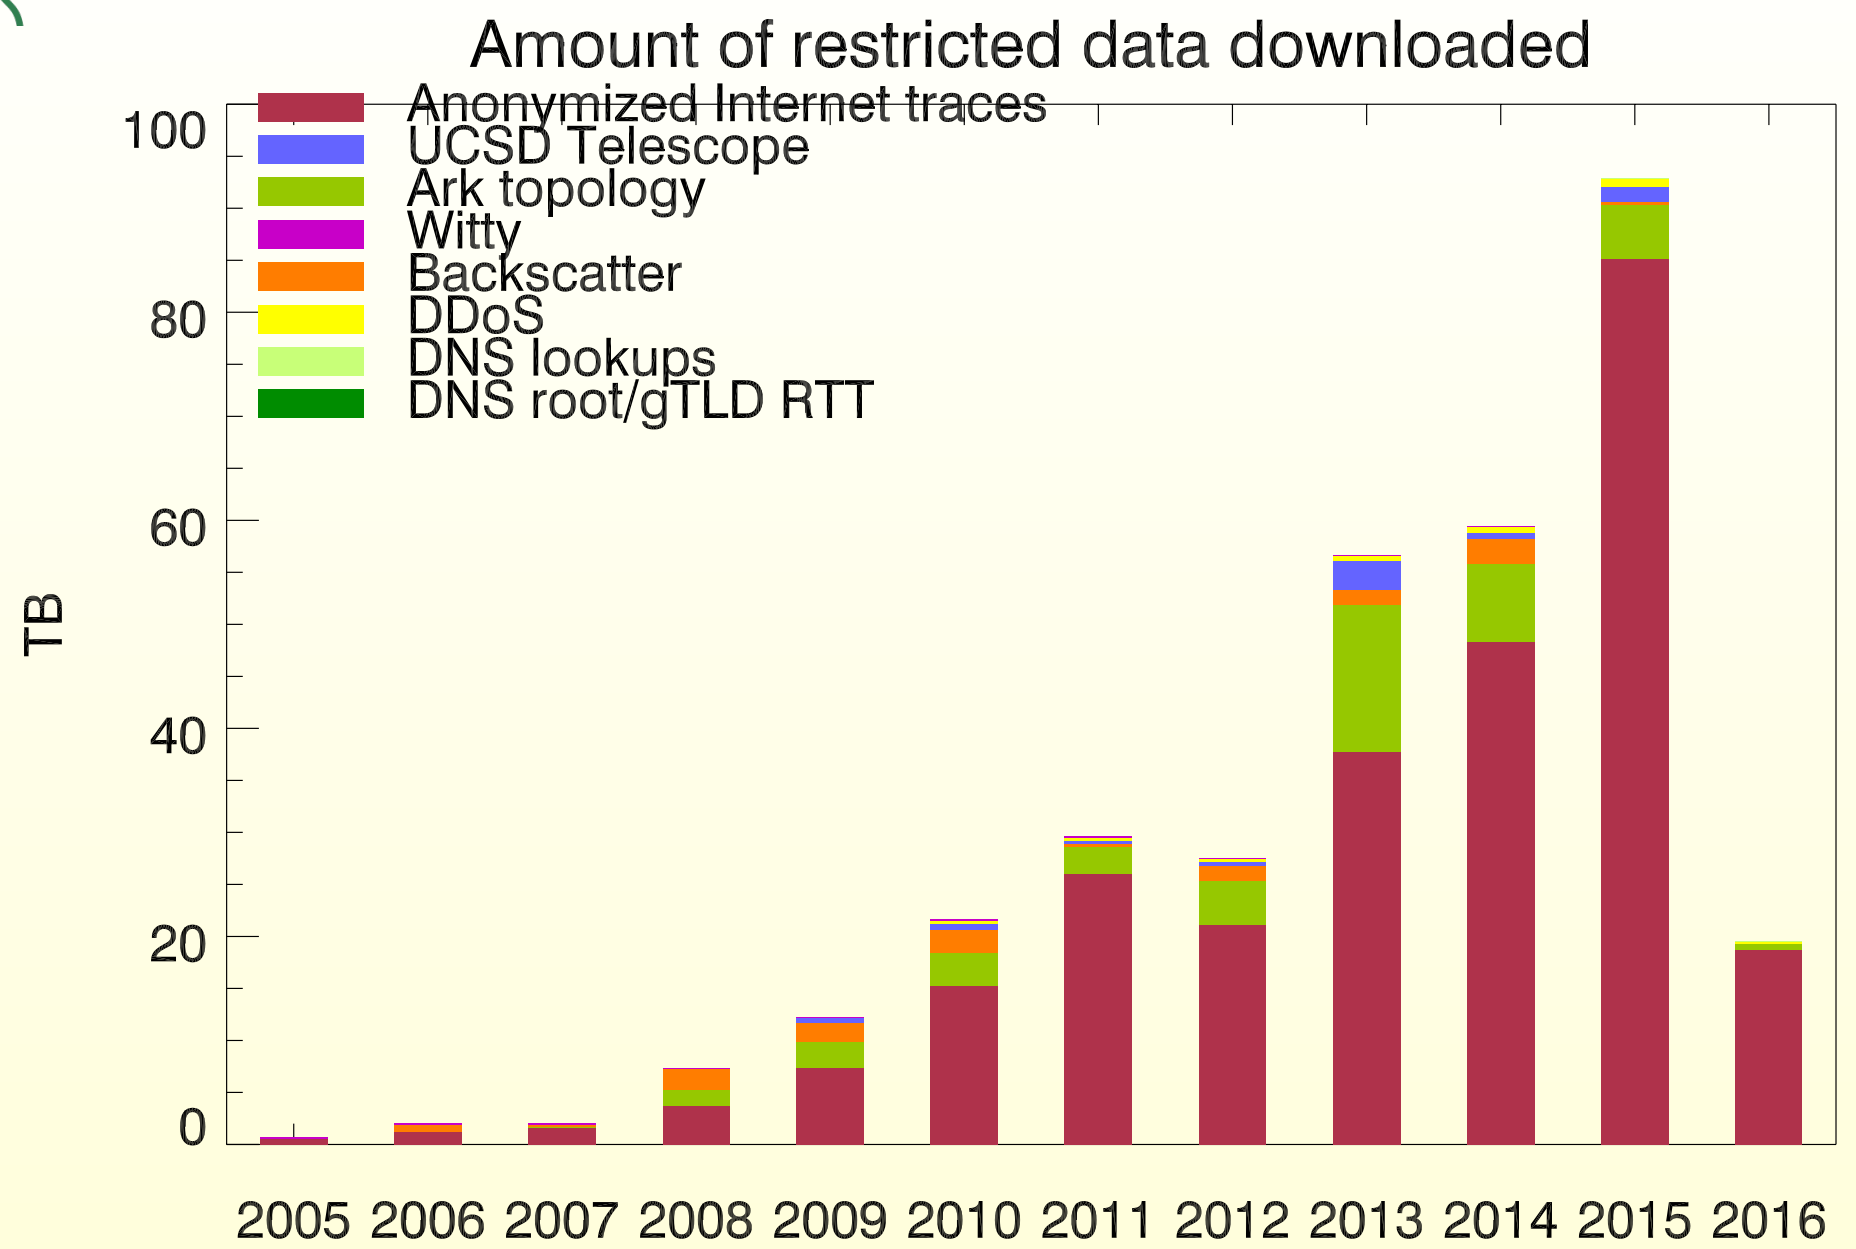

- drop in topology data in 2016 due to making topology data public
- drop in passive data accompanying loss of monitor access

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

Recent publications

- A. Dainotti, E. Katz-Bassett, and X. Dimitropoulos, “**The BGP Hackathon Report**", ACM SIGCOMM Computer Communication Review (CCR), Jul 2016.
- k. claffy, "**The 8th Workshop on Active Internet Measurements (AIMS8) Report**", ACM SIGCOMM Computer Communication Review (CCR), Jul 2016.
- k. claffy, A. Dhamdhere, D. Clark, and S. Bauer, "**Report of AT&T Independent Measurement Expert Background and supporting arguments for measurement and reporting requirements**", Aug 2016
- k. claffy, D. Clark, S. Bauer, and A. Dhamdhere, "**Policy challenges in mapping Internet interdomain congestion**", in Telecommunications Policy Research Conference (TPRC), Aug 2016.
- C. Orsini, A. King, D. Giordano, V. Giotsas, and A. Dainotti, "**BGPStream: a software framework for live and historical BGP data analysis**", accepted to Internet Measurement Conference (IMC), Nov 2016
- M. Luckie, A. Dhamdhere, B. Huffaker, D. Clark, kc, “**bdrmap: Inference of Borders Between IP Networks**”, IMC, Nov 2016
- R. Padmanabhan, A. Dhamdhere, Emile Aben, kc claffy, Neil Spring, “**Reasons Dynamic Addresses Change**”, IMC2016**.**

Recent Related R&D Activities

- Ark Topology Query System (vela.caida.org)
- Spoofing measurement (spoofer.caida.org)
- new DHS project: Science of Internet Security: Technology and Experimental Research (SISTER)
- ATT/FCC Independent Measurement Expertise
- Internet congestion mapping system (beamer.caida.org)
- Workshops: BGP Hackathon, AIMS8 (Feb 2016)
- BGPStream real-time BGP data processing system
- Periscope: Unified Interface to Administratively Distributed Measurement Infrastructure (www.caida.org/tools)
- FCC MBA: Detecting CGN in the ISP

Periscope: Unifying Looking Glass Querying

- publicly-accessible overlay that unifies LGs into single platform, automates their discovery and use
- crowd-sourced and cloud-hosted querying mechanisms to scale querying resources
- handles bursts of requests, intelligent controller
- enforces pre-existing rate limit LG policies
- Automatically extracted 1691 LGs in 297 ASs (Dec 15)
- V. Giotsas, A. Dhamdhere, and k. claffy, "Periscope: Unifying Looking Glass Querying", PAM, Mar 2016. http://www.caida.org/publications/papers/2016/periscope/

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)

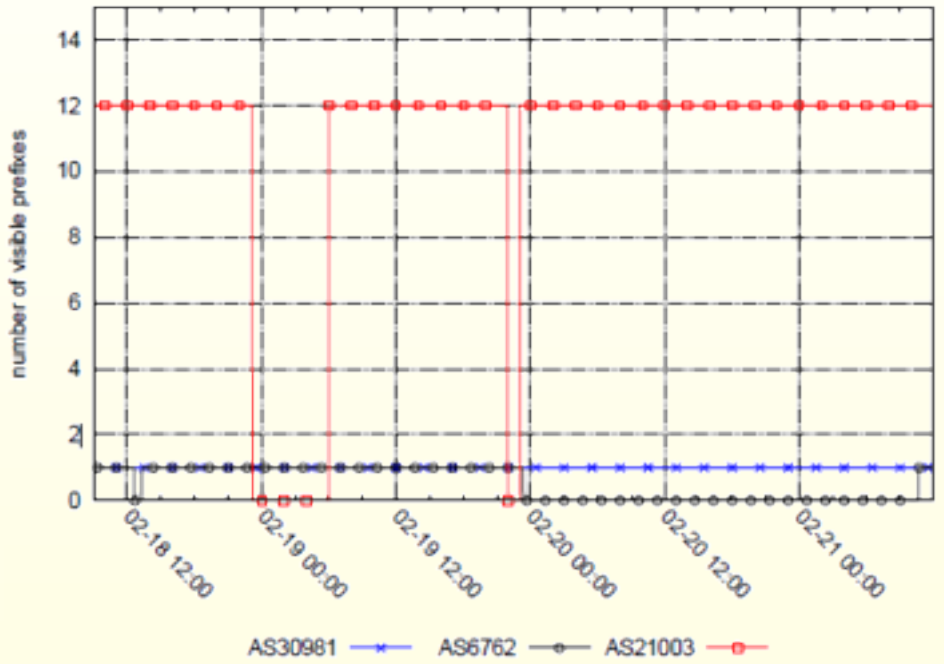

(a)

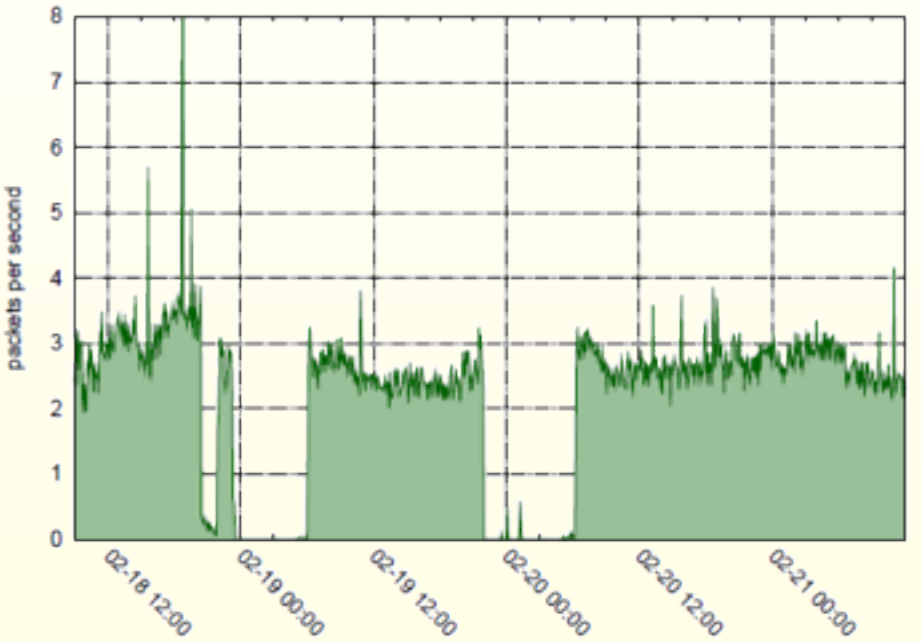

(b)

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

(b) unsolicited traffic to UCSD telescope from Libya.

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

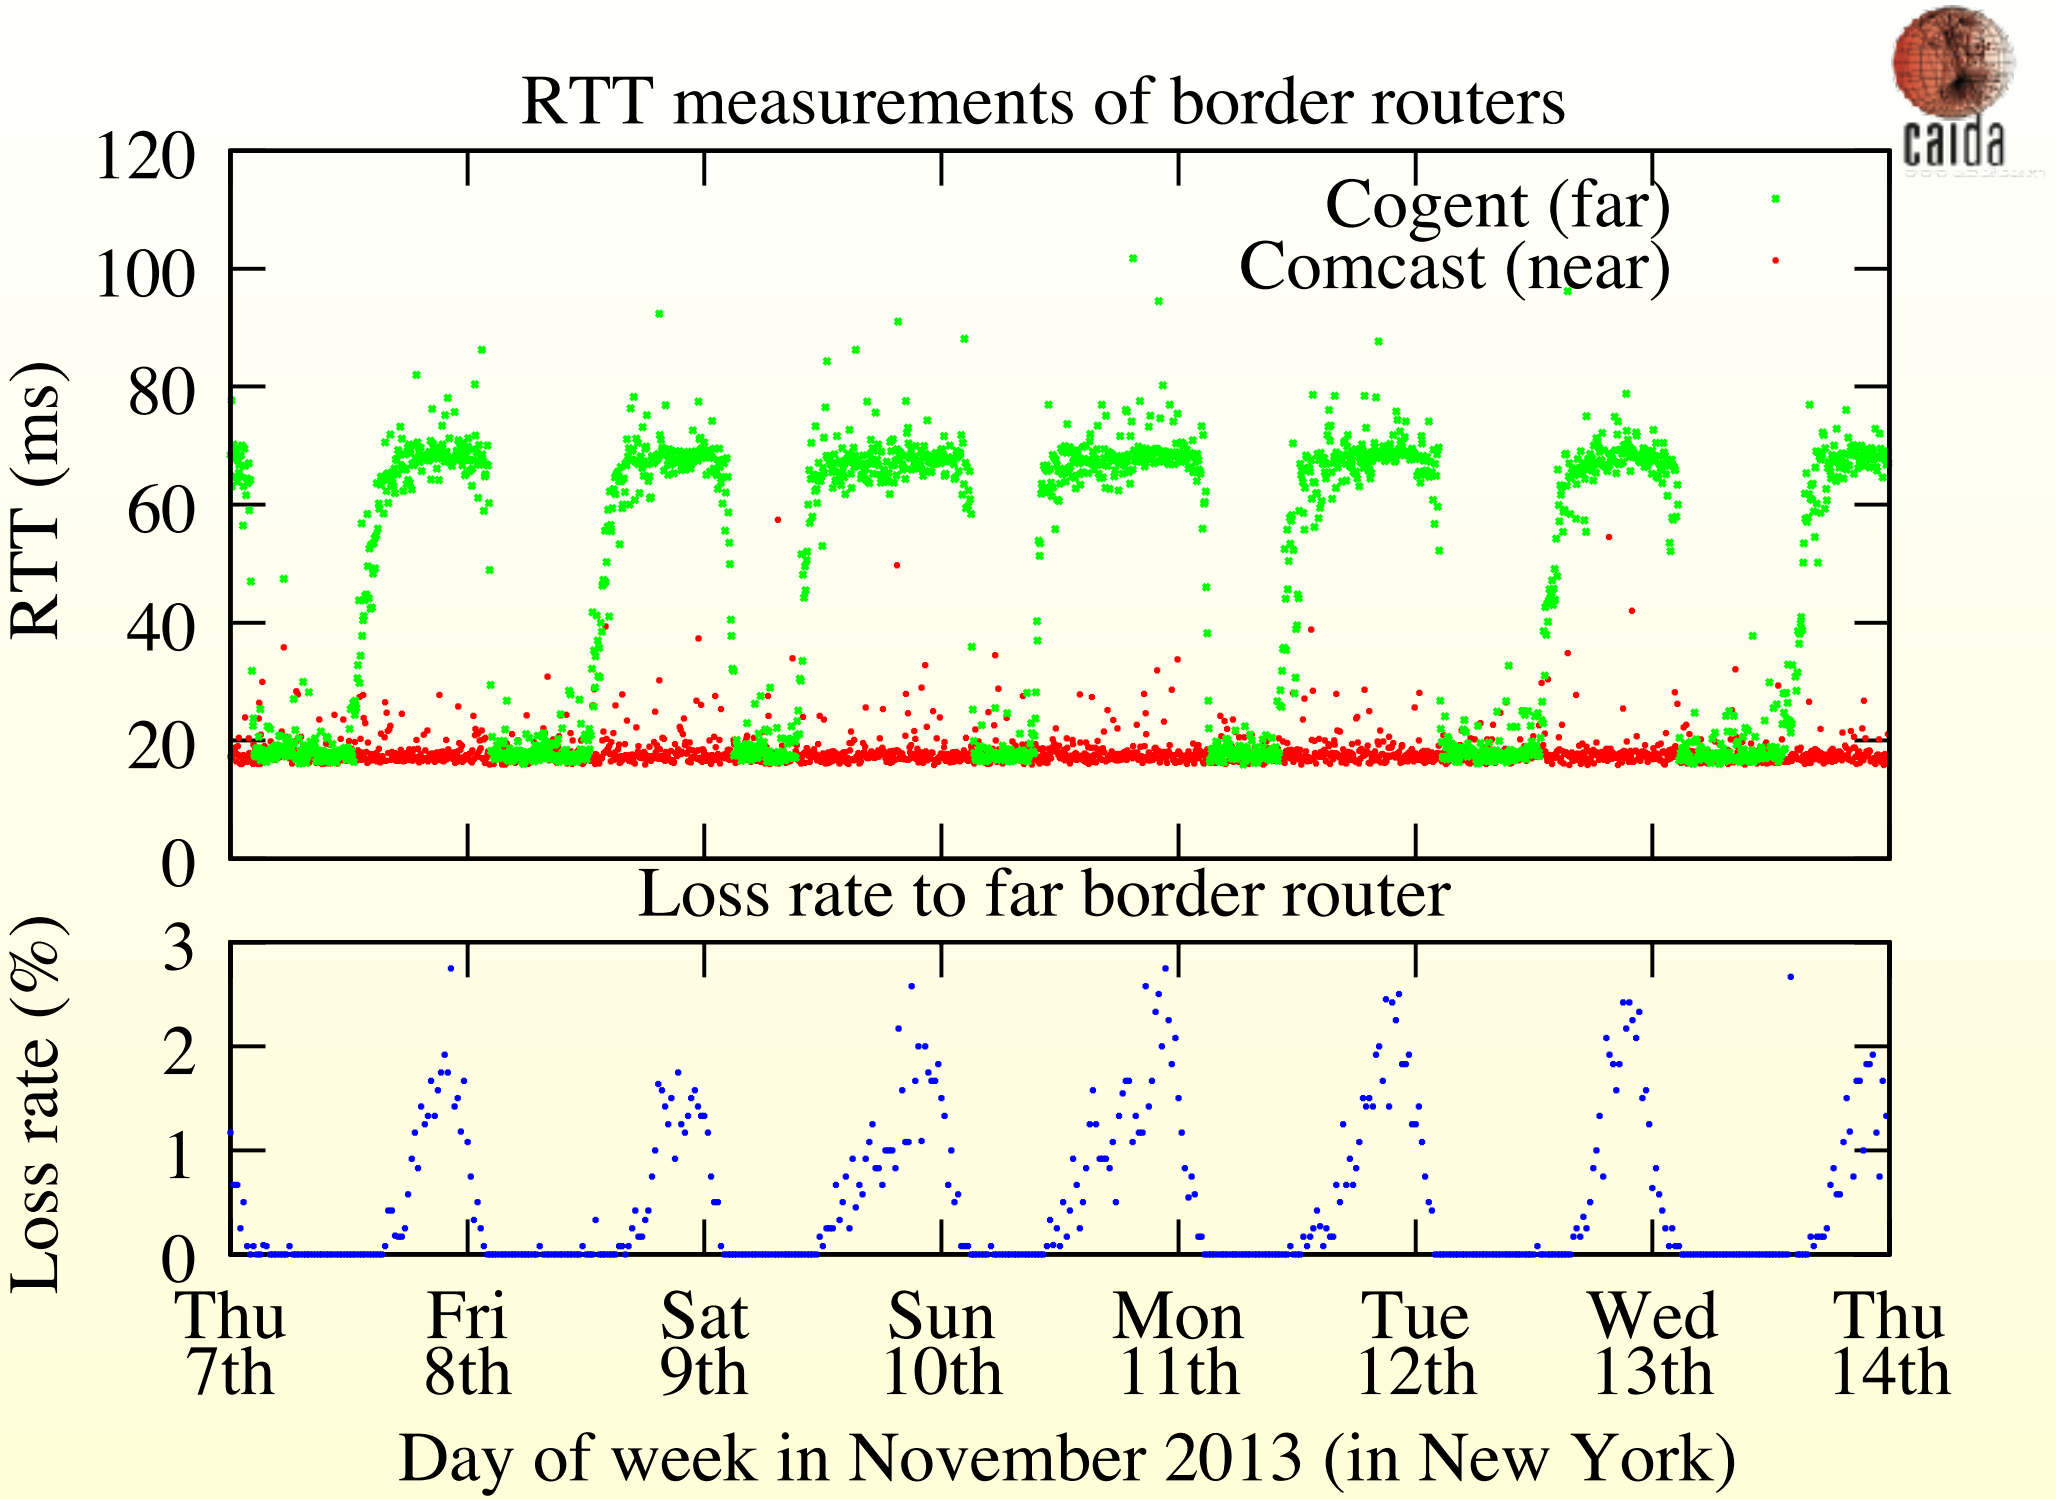
RTT measurements of border routers
Cogent (far)
Comcast (near)
RTT (ms)
120
100
80
60
40
20
0
Loss rate to far border router
Loss rate (%)
3
2
1
0
Thu 7th
Fri 8th
Sat 9th
Sun 10th
Mon 11th
Tue 12th
Wed 13th
Thu 14th
Day of week in November 2013 (in New York)

Measurement System

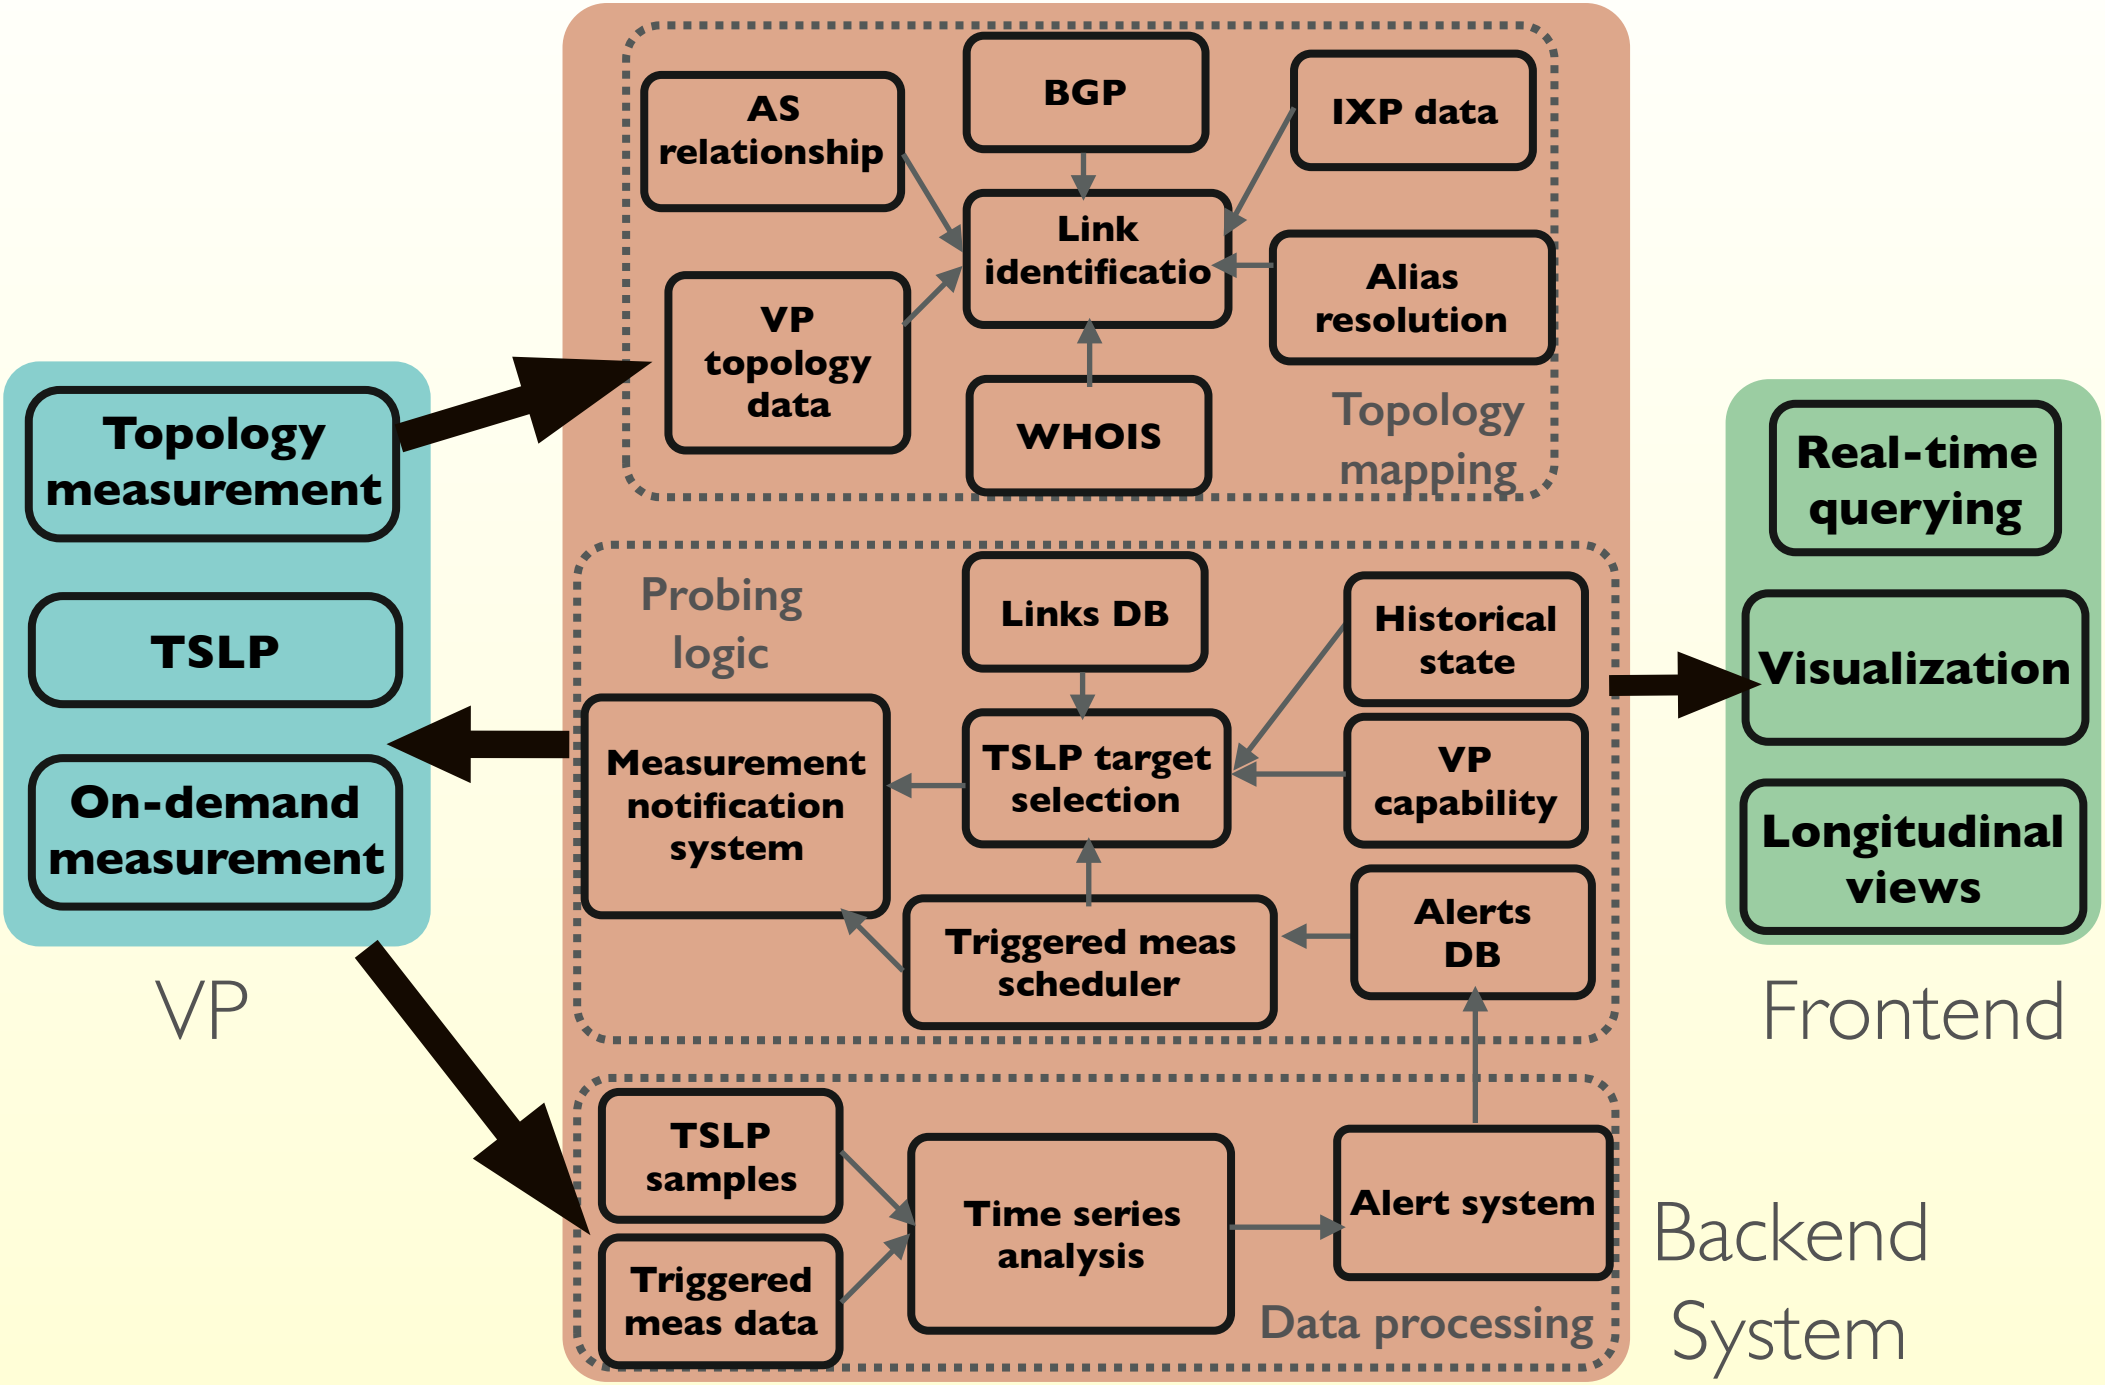

Software Systems for Surveying Spoofing Susceptibility

- 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 Change filters

Session	Timestamp	Client IP	ASN	Country	NAT	Spoof Private	Spoof Routable	v4 Adjacency Spoofing	Results
73442	2016-09-28 11:57:32	82.195.64.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
		2003:66::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 11:35:16	70.196.30.x	22394 (CELLCO)	usa (United States)	yes	blocked	rewritten	none	Full report
		2600:100e::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
73428	2016-09-28 11:21:08	122.252.250.x	24560 (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 11:08:25	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
		204.235.144.x	2452 (TW-CABLE)	usa (United States)	yes	unknown	unknown		

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

Vela/Henya Web Interface to Topology Measurements and Data

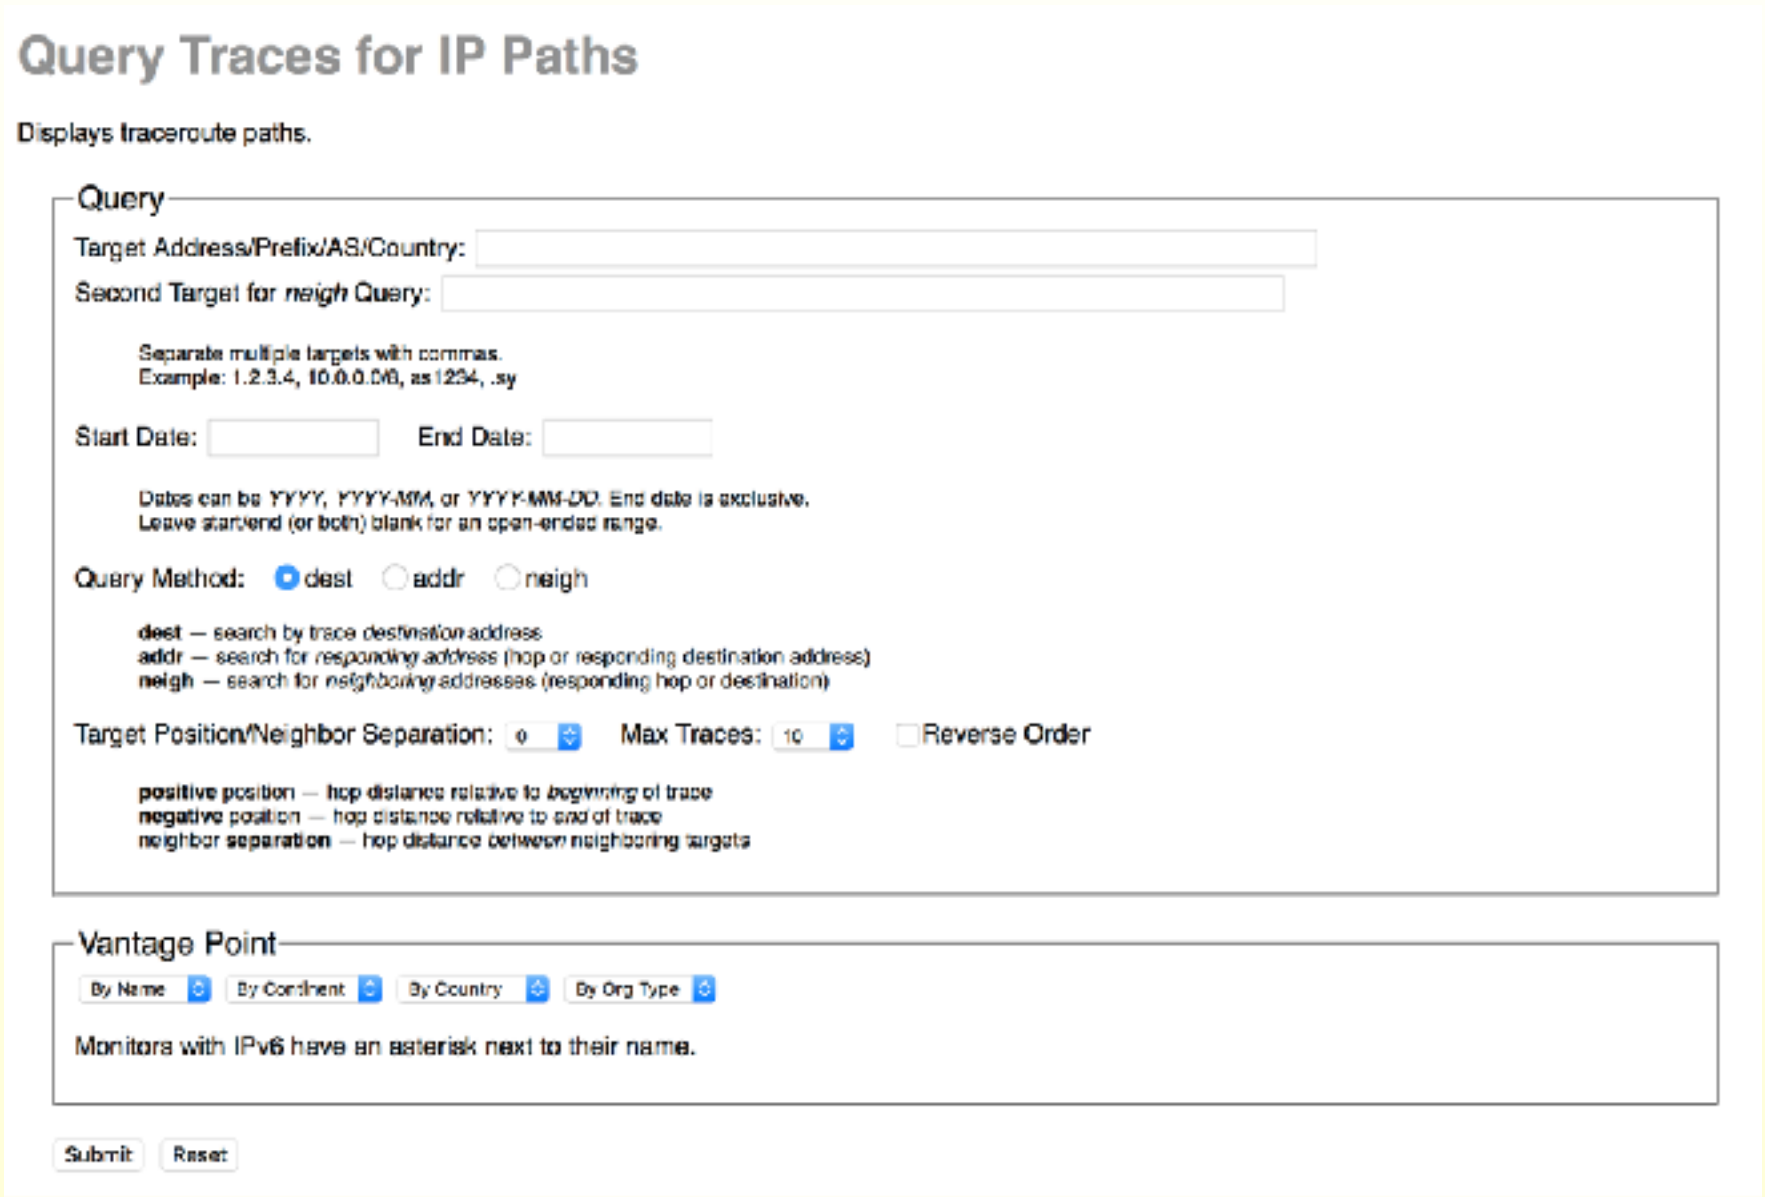

Recent Marketing/Outreach Efforts

- CAIDA Information Sheet
  - http://www.caida.org/publications/posters/eps/caida-infosheet-2016.pdf
- DHS S&T C4 Information Sheet
  - http://www.caida.org/publications/posters/eps/c4-infosheet-2016.pdf
- CRA Congressional Visit
  - http://blog.caida.org/best_available_data/2016/09/27/cra-congressional-visit-to-washington-d-c/
- Videos
  - Spoofer request for help
  - 17 years of Internet Topology

Center for Applied Internet Data Analysis (CAIDA)

http://www.caida.org/
info@caida.org

UC San Diego

caida

Overview

The Center for Applied Internet Data Analysis (CAIDA) conducts network research and builds research infrastructure to support large-scale data collection, curation, and data distribution to the scientific research community. The group, located at the federally funded San Diego Supercomputer Center located at the University of California, San Diego, designs, deploys and maintains a growing number of computational, data analysis and visualization services. The group also ships and maintains small form factor measurement instrumentation to networks around the world, extending its Archipelago (Ark) Internet measurement platform for use by the network and cybersecurity research community. CAIDA researchers develop novel techniques to collect, analyze, query and visualize the resulting data.

Mission Statement: CAIDA investigates practical and theoretical aspects of the Internet, focusing on activities that:

⇒ provide insight into the macroscopic function of Internet infrastructure, behavior, usage, and evolution,

⇒ foster a collaborative environment in which data can be acquired, analyzed, and (as appropriate) shared,

⇒ improve the integrity of the field of Internet science, and

⇒ inform science, technology, and communications public policies.

Research and Analysis

CAIDA's research spans Internet topology, routing, security, economics, future Internet architectures, and public policy. Our infrastructure, software development, and data sharing activities support measurement-based Internet research, both at CAIDA and around the world, with focus on the health and integrity of the global Internet ecosystem.

Mapping the Internet. We pursue Internet cartography to improve our Internet topology mapping capabilities using our expanding and extensible Ark measurement infrastructure. We work to improve the accuracy and sophistication of our topology annotation capabilities, including classification of ISPs and their business relationships. Using our evolving IP address alias resolution measurement system, we collect, curate, and release our flagship data product, the Internet Topology Data Kit (ITDK).

Mapping Interconnection Connectivity and Congestion. We use the Ark infrastructure to support an ambitious collaboration with researchers at Massachusetts Institute of Technology (MIT) to map the rich mesh of interconnection in the Internet, with a focus on congestion induced by evolving peering and traffic management practices of Content Distribution Networks (CDN) and access ISPs, including methods to detect and localize the congestion to specific points in networks. We undertake studies to pursue different dimensions of this challenge: identification of interconnection borders from comprehensive measurements of the global Internet topology; identification of the actual physical location (facility) of an interconnection in specific circumstances; and mapping observed evidence of congestion at points of interconnection. We produce related data collection and analysis to enable evaluation of these measurements in the larger context of the evolving ecosystem: quantifying a given ISP's global routing footprint; classification of autonomous systems (ASes) according to business type; and mapping ASes to

CAIDA's projects made possible by support from:

Department of Homeland Security Science & Technology Division

National Science Foundation

SDSC SAN DIEGO SUPERCOMPUTER CENTER

Spoofer request for help

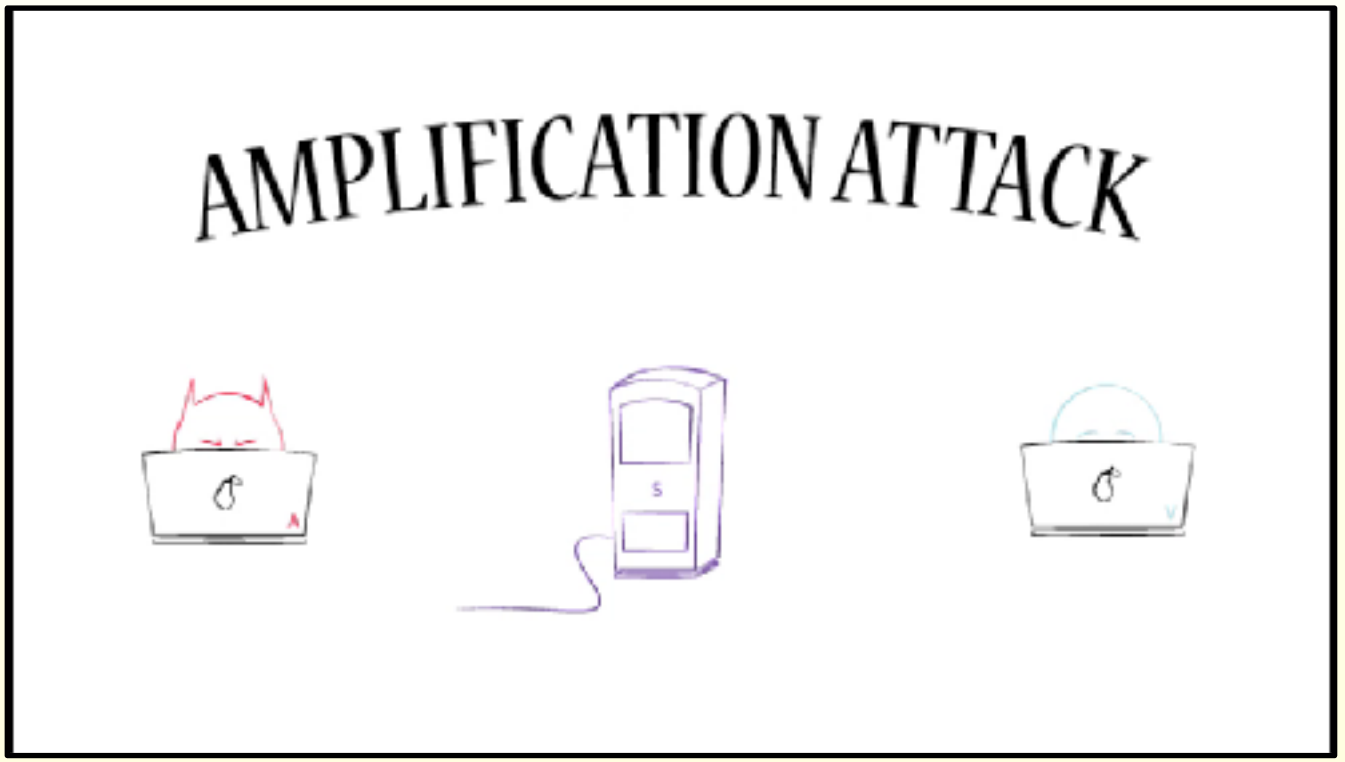

The video will explain to a general audience the dangers of IP spoofing.

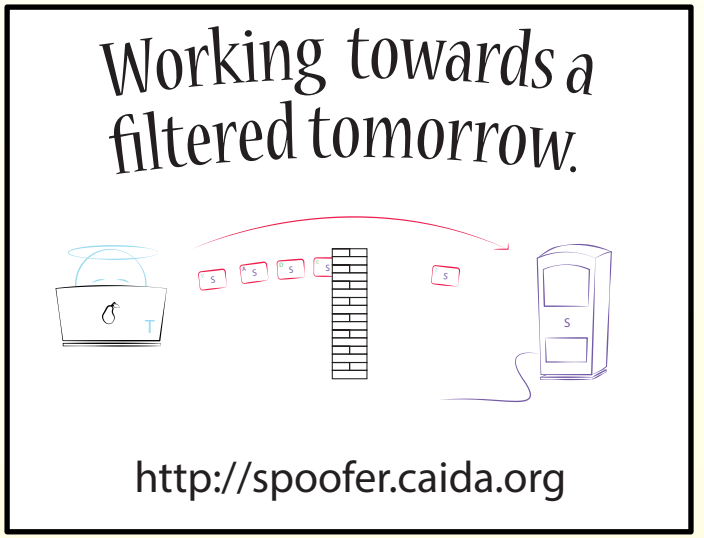

We will end the video with a requester help.

17 years of Internet AS topology

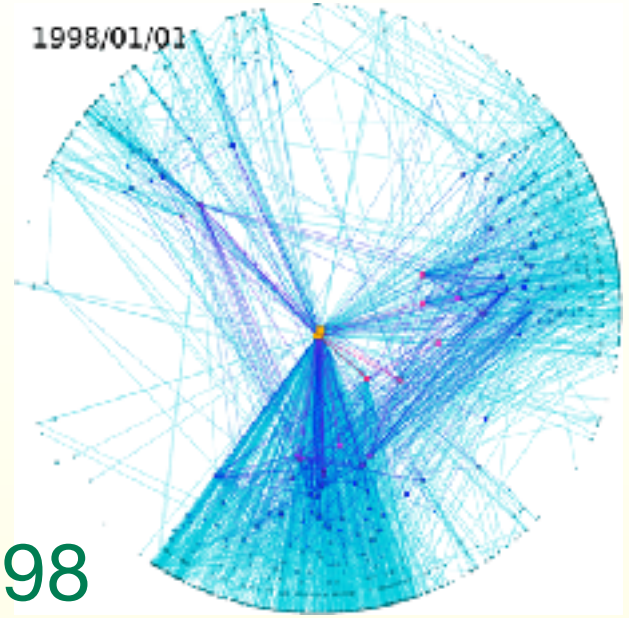

1998

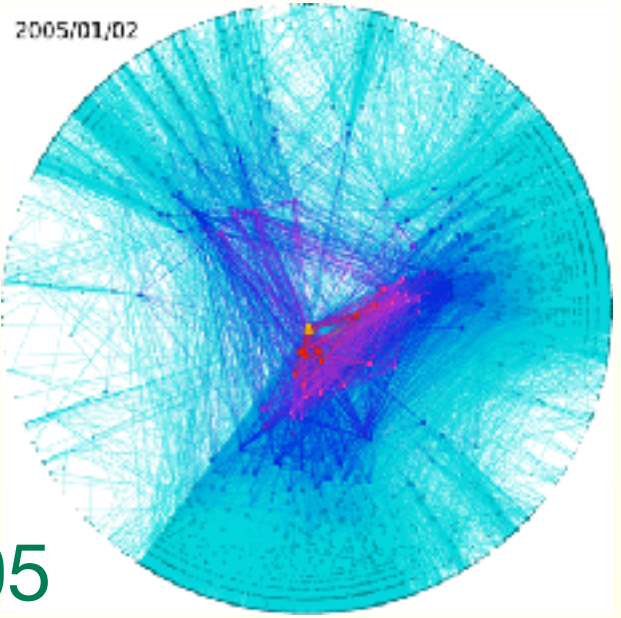

2005

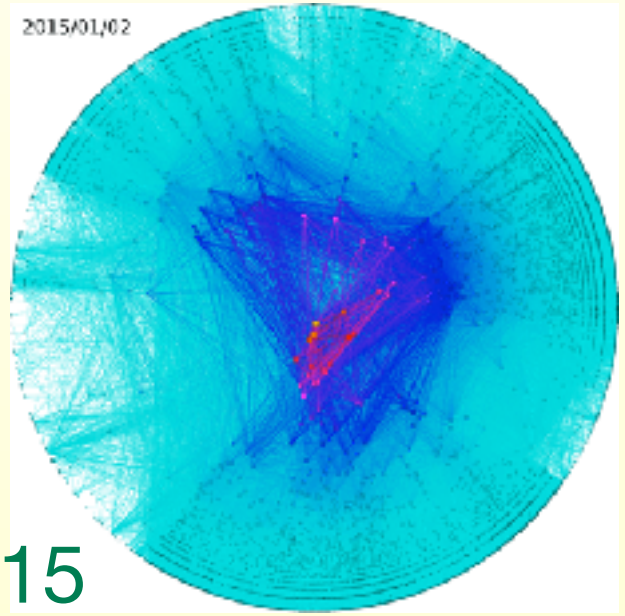

2015

goal: animations capturing growth & diversity of the Internet AS topology

Contact Information

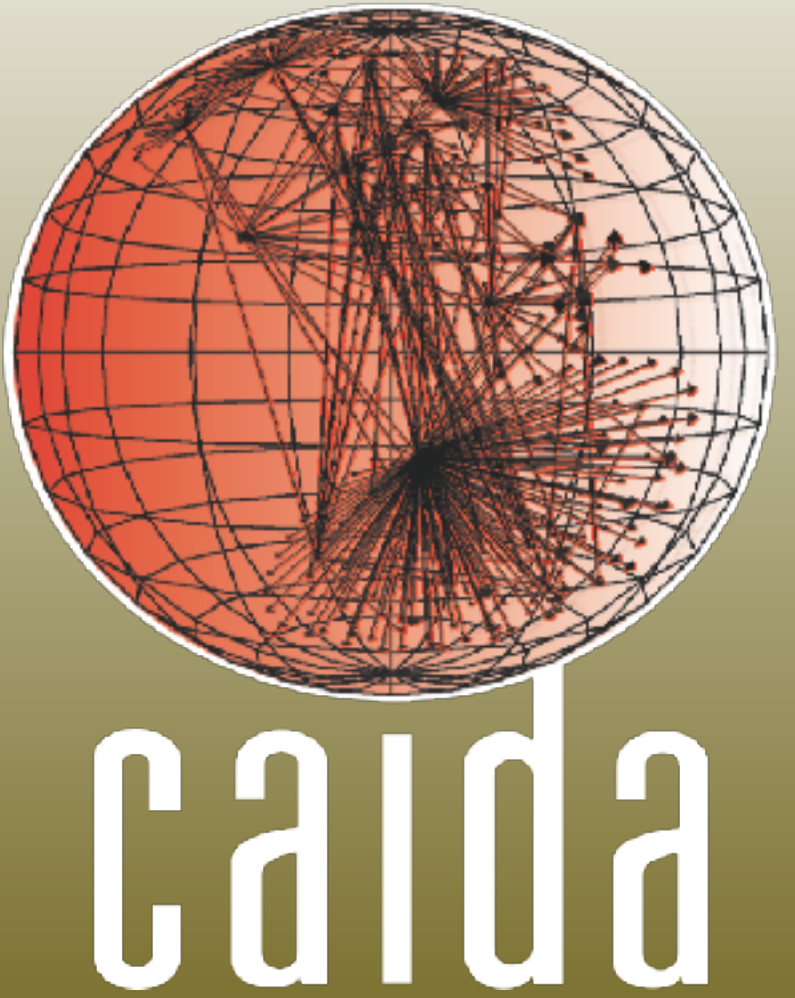

PI: k claffy, CAIDA
kc@caida.org
http://www.caida.org/