




ILANDS

Integrated Library for Advancing
Network Data Science

CNS-2120399



CIRC PI Meeting

7 Nov 23

kc claffy

CAIDA/UCSD



Major Goals

Enhance Internet measurement infrastructure that can support sustained longitudinal measurements and new experiments by upgrading and integrating:

- 1. 100GB two-way traffic capture**
- 2. BGP routing data collection**
- 3. Outreach and community engagement**



Task 1: Traffic Data Infrastructure Enhancements

- 1. Build 100GB traffic monitor**
- 2. Test and evaluate monitor**
- 3. Deploy monitor**
4. Manage and share traces
5. Establish data enclave at CAIDA
6. Augment with other (external) data sources
7. User training and support



Task 1: Traffic Data: Accomplishments


100GB traffic monitor built, tested, evaluated, deployed

- Redesigned software
- Specified/ordered/tested hardware
- Shipped system to (volunteer) host
- So many obstacles
- **Finally saw packets this month! (Y2Q4..)**
- Plan 2 more in 2024



Task 1: Traffic data: Accomplishments

Establishing data enclave at CAIDA (in process)



- Developed **storage specifications**
- Ordered and received hardware: 60x14TB drives (**840 TBs raw spinning disk**) in a 4U JBOD (Western Digital) chassis connected to 1U server
- **Established working group to integrate data/curricula** into UCSD DSMLP (datahub.ucsd.edu) and **National Research Platform (NRP)**



Task 2: BGP Data Infrastructure Enhancements

1. Enhance BGPStream service broker
2. **New interface to BGP data at RouteViews**
3. **Integrate active measurement with BGP collectors**
4. Enable OpenBMP capabilities (deferred)
5. **Enable data integrity controls**
6. **Update RouteViews infrastructure administration**

Task 2: New interface to RouteViews BGP data

<https://bgp2go.caida.org/> (authenticated)

BGP2GO

Searchbar

Search for IP prefix, ASN, or BGP Community


Selected Prefix

pre=2405:6e00:f800:fefe:0:bad:bad:bad/128

Some stats

OVERALL: 222,442 occurrences (111,182 withdrawals) in 8,257 files (6.26 GB) with 3 unique collectors.
 CURRENT: 69,391 occurrences (34,689 withdrawals) in 2,810 files (1.62 GB) with 2 unique collectors.
 31% of occurrences | 34% of files | 26% size

Further functions



Some plots

Set filters

CLEAR ALL FILTERS X YEAR:2017 - AUG -

Years	Months	Days	Collectors
2012	JAN (16) +	1 (4,200) +	1 ROUTE-VIEWS2
2013	FEB (69) +	2 (3,813) +	2 ROUTE-VIEWS3
2014	MAR (134) +	3 (3,932) +	3 ROUTE-VIEWS4 (54,177) +
2015	APR (107,950) +	4 (4,136) +	4 ROUTE-VIEWS6
2016 (7) +	MAY (1) +	5 (4,042) +	5 ROUTE-VIEWS.A
2017 (222,321) -	JUN (13) +	6 (4,435) +	6 ROUTE-VIEWS.F
2018 (104) +	JUL (27,929) +	7 (7,923) +	7 ROUTE-VIEWS.B
2019 (10) +	AUG (69,402) -	8 (6,119) +	8 ROUTE-VIEWS.CH
2020	SEP (16,865) +	9 (7,027) +	9 ROUTE-VIEWS.C
2021			10 ROUTE-VIEWS.EQIX

filters

OR
OR
OR
OR
OR
OR
OR
OR
OR
OR

AND

AND

AND

(YEAR1 OR YEAR2 OR YEAR...)
 AND
 (MONTH1 OR MONTH2 OR ...)
 AND
 (DAY1 OR DAY2 OR DAY10)
 AND
 (COLL1 OR COLL2 OR COLL...)



Task 2: BGP data infrastructure enhancements

Enhance BGPStream service broker

- **Incorporated BGPReader** (a BGPstream utility) **into BGP2GO** (CAIDA science gateway prototype) platform
 - **for users who want to leverage both capabilities**
 - **Lowers barrier** to use of BGPstream
- BGP Data integrity studies
 - **Investigating anomalies** in RouteViews updates



Task 2: BGP data infrastructure enhancements

Integrate active measurement with BGP collectors

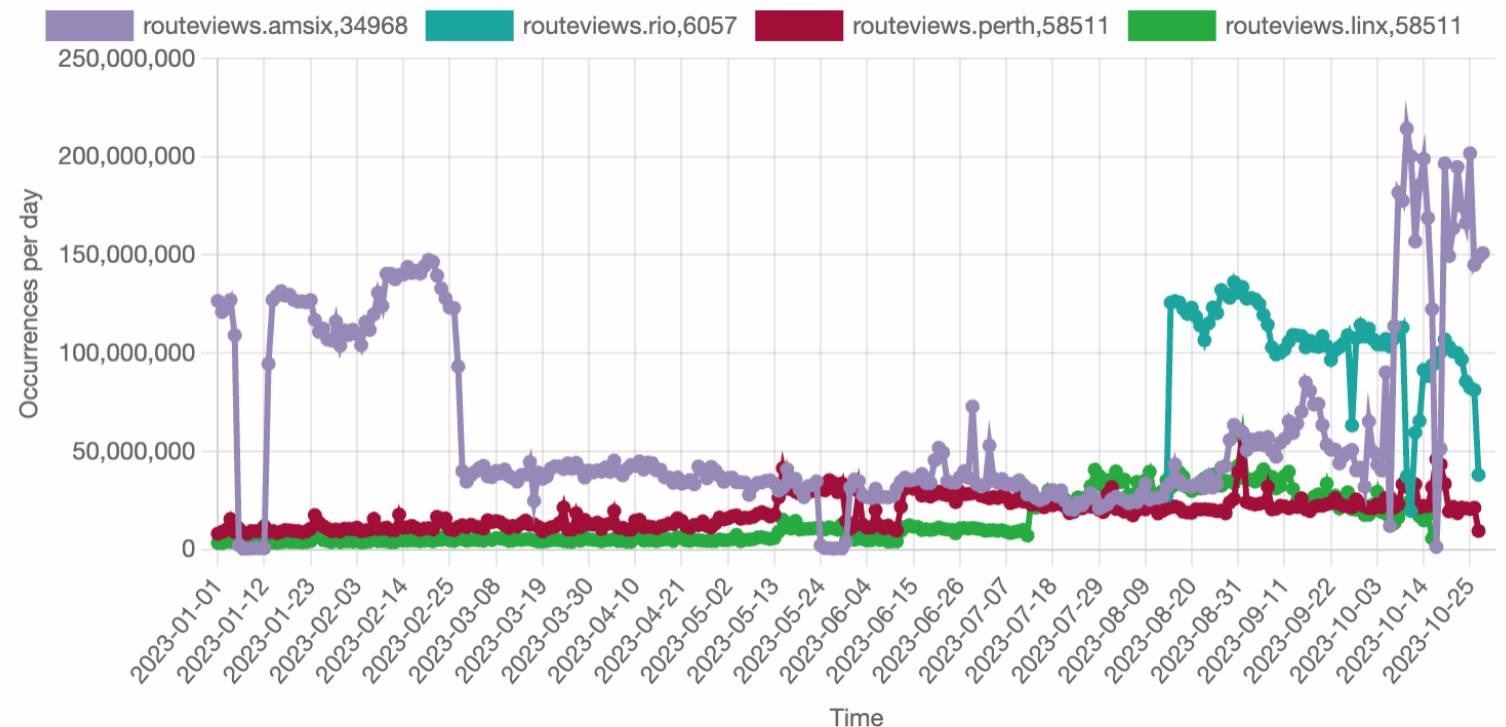
- **Extended scamper to read and write compressed files**, reducing writes to disk by orders of magnitude
- **Deployed scamper to prototype active measurements on existing BGP collector infrastructure**
 - Leveraged Periscope extensions as middleware



Task 2: BGP data infrastructure enhancements

Create BGP data integrity controls

- RV data quality checks (per-peer counts, file sizes)
- Created **peerstats** tool to detect **noisy peers**



<http://nids.caida.org:45000/cgi-bin/peerstats.sh>



Task 2: BGP data infrastructure enhancements

RouteViews Infrastructure Administration Updates

- **Implemented security-related updates**
- **Deployed 3 new collectors** (Pacific Wave Exchange in Los Angeles (CENIC), Chile (PIT), and Malaysia (DE-CIX))
- **Updated hardware/VM collectors** in 6 locations
- **Updated collectors to use FRR**
- **Added over 100 new peers**
- **Improved reliability monitoring**
- **Updated RV website, improved code documentation**
- **Added 300 more external publications using RV data (>1,000 total) and made all publications searchable**



Task 3: Outreach and Community Engagement

- 1. Catalog Management (catalog.caida.org)**
- 2. Ongoing user support (mattermost server)**
- 3. Biannual newsletters (data-announce@caida.org)**
- 4. Annual community workshops (hybrid)**
5. Annual community surveys
6. Sustainability plan



Task 3: Outreach and Community Engagement

Outreach and Community Engagement

- **Catalog updates: new (traffic and topology) data, tools**
- **Meetings with potential monitor hosts: CENIC, DREN**
- **RV engagement: Peering Forums, NANOG 88 & mnNOG 5**
- **8 IMC presentations**
- **Two in-person AIMS workshops this year**
- **FCC BGP security workshop presentations**
- **FCC public comment**

https://catalog.caida.org/paper/2023_a_path_forward

CAIDA ILANDS Team



kc claffy
Principal Investigator



David Clark
Co-Principal Investigator



Bradley Huffaker
Technical Project Manager



Dan Andersen
Systems Administrator



Elena Yulaeva
Data Administrator



William Herrin
Consulting
Software Architect and Engineer




Brendon Jones
Consulting
Research Programmer



Thomas Krenc
Postdoc from Naval
Postgraduate School



Matthew Luckie
Consulting
Research Scientist



Alex Ma
Manager Scientific Projects
Website Administrator



Victoria Nguyen
Systems Administrator



Philip (Leo) Pascual
Data Systems Analyst



Hammas Tanveer
PhD. student from U Iowa