Homeland Security
Science and Technology

2017 | Cyber Security Division
IMAM Mid Term Meeting

Science of Internet Security: Technology and Experimental Research

CAIDA/UCSD
PI k claffy

30 Jan 2017

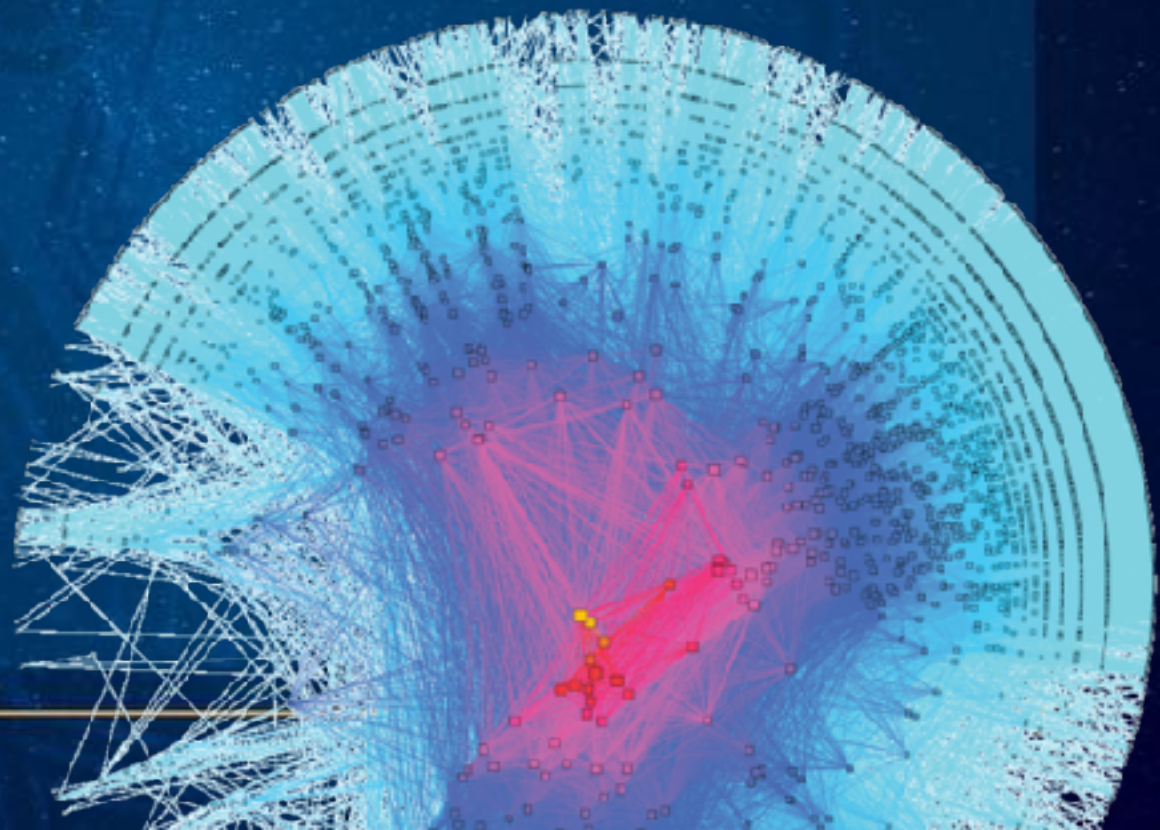

Team Profile

The Center for Applied Internet Data Analysis (CAIDA)

– Founded by PI and Director k claffy

– Independent analysis and research group

– 20+ years of data collection, curation, analysis, sharing

– Supporting measurement infrastructure, tool development, interactive access to measurement, data analysis capabilities

– Hundreds of data users

– Located at the UC San Diego Supercomputer Center

Key personnel: Dan Andersen, Alberto Dainotti, Marina Fomenkov, Alistair King, Bradley Huffaker, Young Hyun, Alistair King, Vasilieos Giotsas

Need: Macroscopic Models and Assessment of the global Internet

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

- Archipelago (Ark) platform (170 nodes and growing) supports active measurement studies of the Internet
- Ark gathers and shares the largest set of network topology data used for a broad spectrum of scientific research
- SISTER builds on this platform to involve a broader cross-section of the security research community
- Six targeted measurement needs to support assessments related to Internet security and stability

Six Targeted Needs

1. Macroscopic Security and Stability Monitoring and Analysis: tools for studying outages and BGP hijacking
2. Map of peering interconnections at the router level
3. Map of peering interconnections at the facility level
4. Measurements of TCP vulnerabilities
5. Software to infer grey-market IPv4 address transfers
6. On-demand router-level mapping

Detecting traffic interception using BGP anomaly detection:

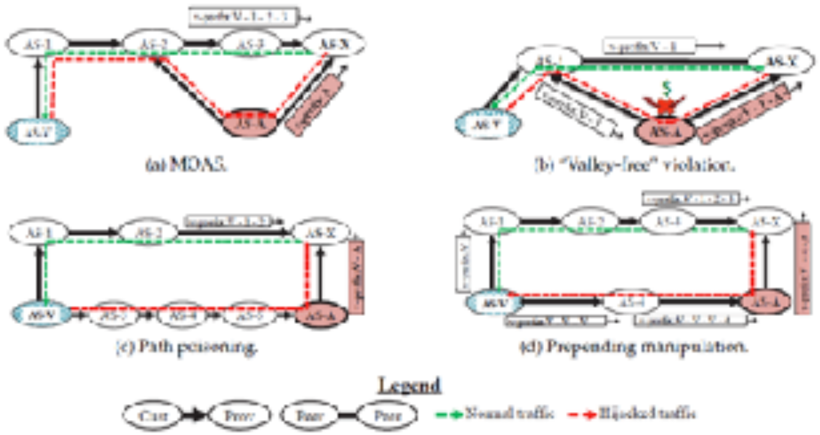

Need 1: Support for Macroscopic Security and Stability Monitoring and Analysis

1. Generate target list of prefixes to probe daily
   a. sanitize BGP data from sliding 1-week window of RV/RIS data
   b. dynamically identify which IPs to probe
   c. Result: s/w module that continuously queries BGP Watcher application

2. AS traceroute: overcome errors from third-party addresses, other traceroute artifacts.
   a. more accurate cross-validation of BGP vs traceroute incongruities
   b. use AS relationship data to estimate router ownership

Approach 1: Support for Macroscopic Security and Stability Monitoring and Analysis

3. Measurement and data processing pipeline.
   a. generate baseline reachability map that is fine-grained in spatial and temporal granularity
   b. allow comparison of triggered traceroutes with recent history
   c. Support: interactive monitoring for detection of MITM BGP hijacks
   d. Support: Internet outage detection and analysis (IODA) system
   e. Curated daily files of AS paths (BGP/tr) for reference, e.g., post-event analysis of security incidents (share via IMPACT)

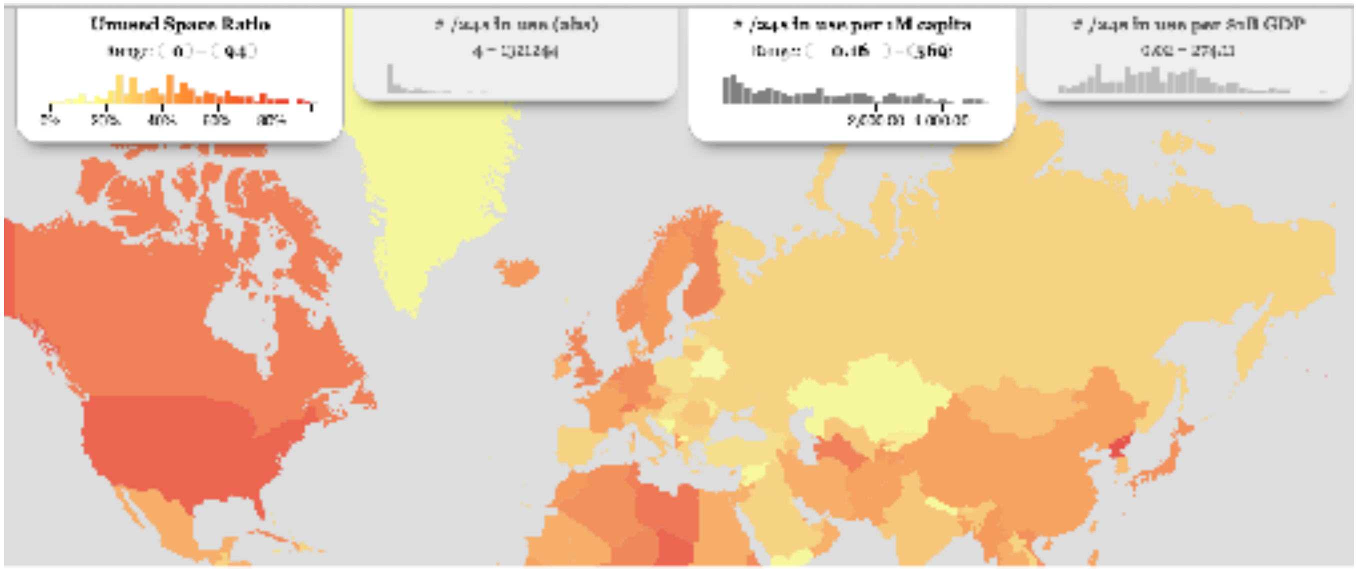

[Approach 1: Supporting IODA project]

Center for Applied Internet Data Analysis
University of California San Diego

Internet Outage Detection & Analysis

Challenge

- Real-time detection of Internet outages
- Global view
- Analysis capabilities

Screenshot of the IODA dashboard highlighting outages in Turkey

Scientific Impact

- Better understanding of network reliability and macroscopic events
- Methodologies for outage detection
- Reference datasets

Solution

- Process data from diverse and distributed sources
- Combine measurements: *control-plane, passive data-plane, active data-plane*
- Modular system
- Interactive visual interfaces

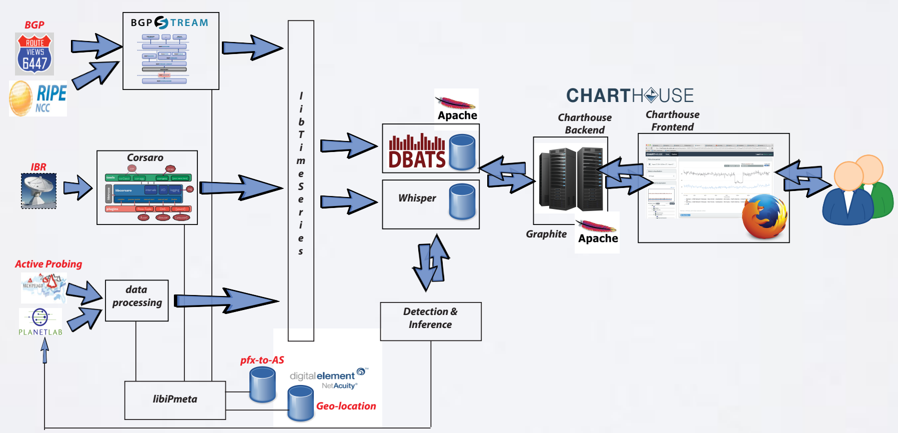

High-level architectural view of the IODA system

Broader Impact

- Data and service available to researchers
- Collaborations with industry and government
- Re-usable, open source code and frameworks
-10 papers,
20 presentations *(IETF, NANOG, IMC, RIPE, …)*

NSF CNS-1228994, UC San Diego
PI: Alberto Dainotti, CoPI: KC Claffy

[Approach 1: Supporting IODA project]

Center for Applied Internet Data Analysis
University of California San Diego

Detecting and Characterizing Internet Traffic Interception based on BGP Hijacking

Challenge

- Near-realtime detection of traffic interception attacks based on BGP prefix hijacking
- Global view

Solution

- Detect suspicious events on the control plane; trigger data-plane active measurements.
- Auxiliary datasets (e.g., AS relationships) to assist with classification of events
- Modular system

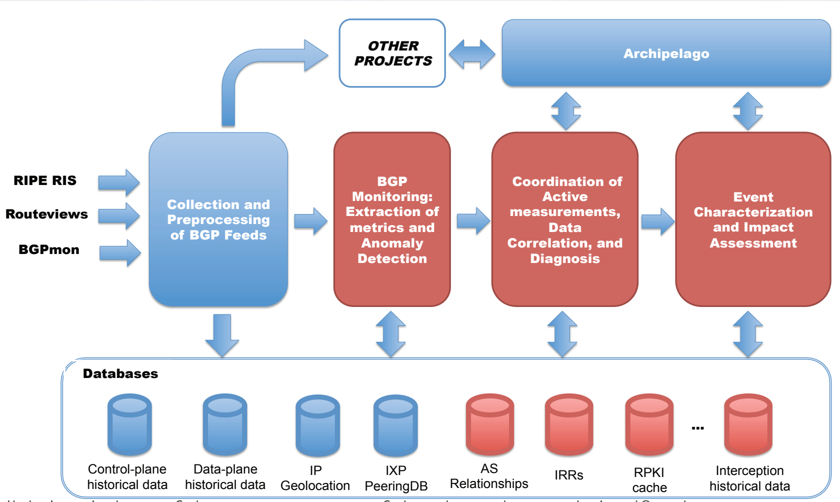

High-level view of the components of the detection and classification system

Scientific Impact

- Methodology for near-realtime detection of BGP-based traffic interception attacks
- Reference datasets of anomalies

Broader Impact

- Publications and presentations at conferences
- Open source, re-usable software (BGPStream)
- BGP Hackathon 2016

NSF CNS-1423659, University of California San Diego
PI: Alberto Dainotti

Need 2: Mapping Peering Interconnections at the Router Level

- (collaborating PI: Matthew Luckie, now at U. Waikato)
- TCP/IP has no notion of interdomain boundaries
- Refine algorithms to infer router ownership and peering interconnections
- Deploy s/w on Ark to conduct AS border mapping
- Validation of borders
- Versions of the s/w for resource-constrained platforms

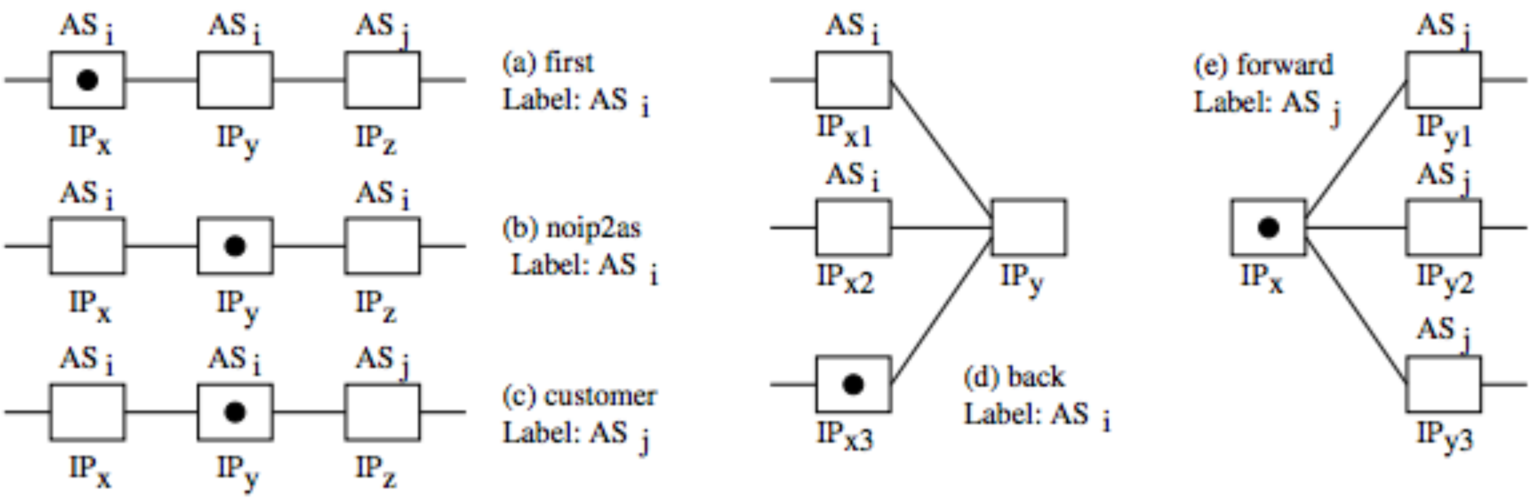

Router ownership inference heuristics used on traceroutes to annotate interfaces with owners.

Approach 2: Mapping Peering Interconnections at the Router Level

1. Explore methods that yield most accurate inferences of inter domain boundaries
2. Apply heuristics to traceroute data to annotate maps w/ device ownership critical for study of interconnects
3. Supports MANIC platform (Measurement and Analysis of Interdomain Congestion) [see demo]

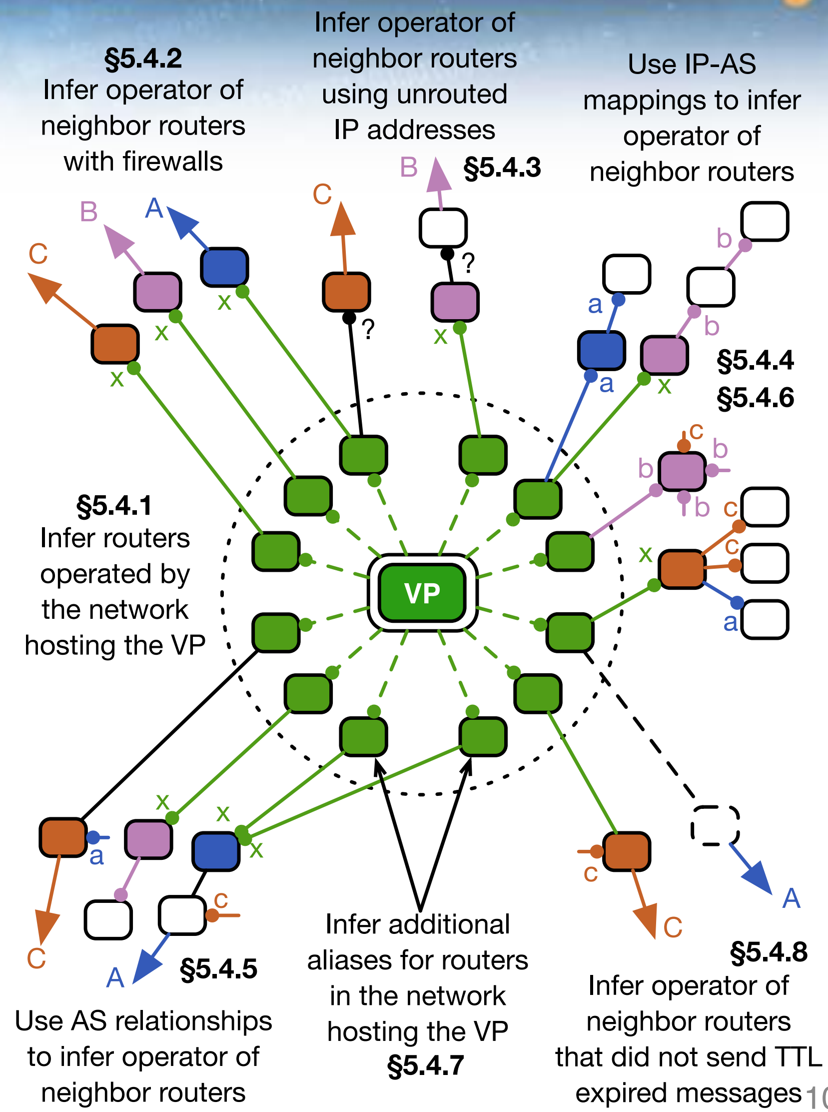

Conceptual mapping of heuristics to infer border routers (Diagram taken from Luckie, et al. IMC2016 paper.)

Need 3: Mapping Peering Interconnections at the Facility Level

1. Manually assemble/maintain IXP database from publicly available data sources (PCH, PeeringDB, etc.)
2. Develop new techniques to infer engineering approach to interconnection
3. Alias resolution of interconnection IP addresses
4. Merge above techniques to generate facility-aware map

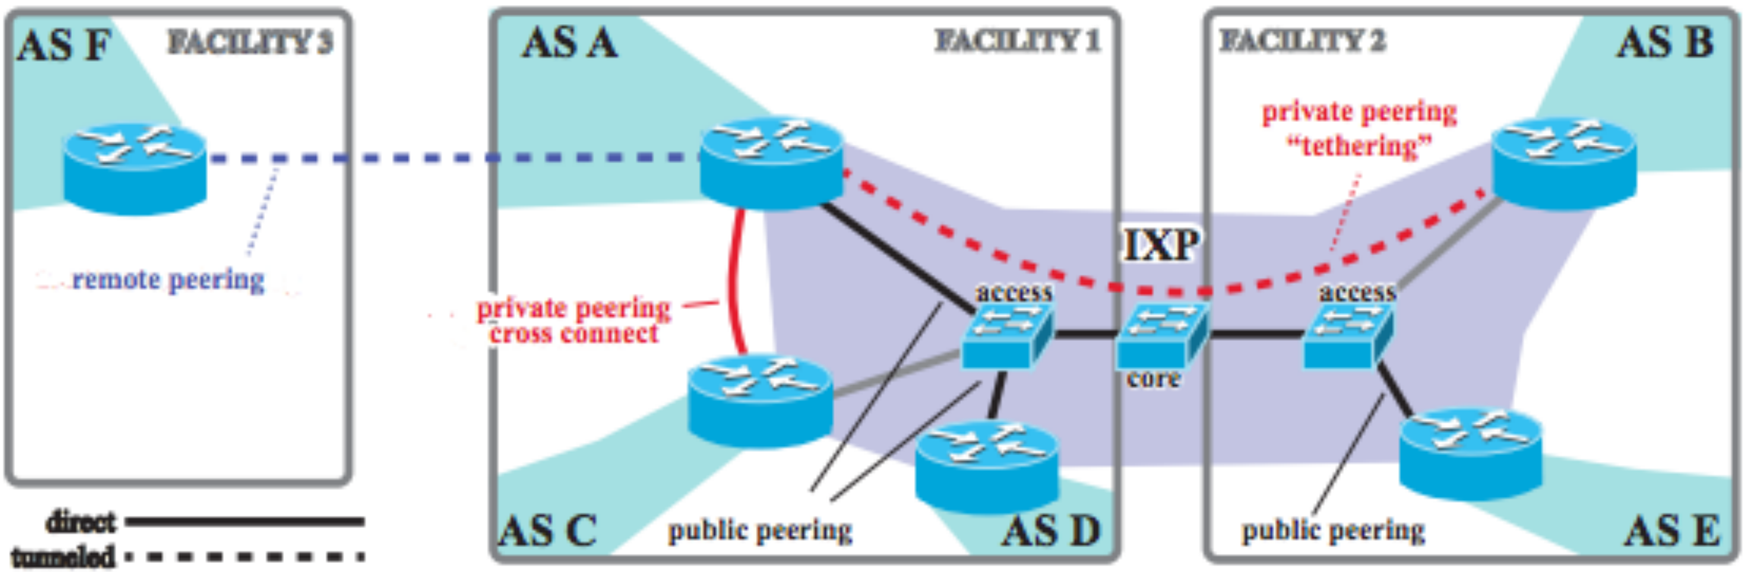

Approach 3: Mapping Peering Interconnections at the Facility Level

- Prototyped and ran this technique with NSF support on a few facilities in Europe (CoNEXT 2015 paper)
- Exploring feasibility of scaling methods to hundreds of facilities

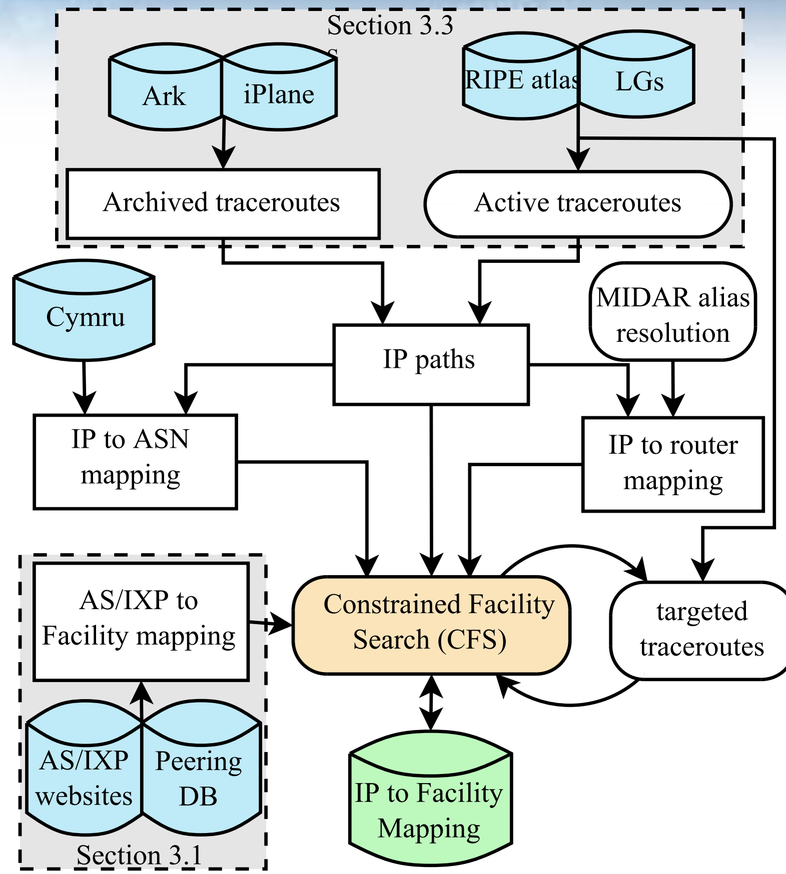

Interconnection-to-facility mapping

Need 4: Measurements of TCP Behavior to Understand Security Vulnerabilities

- Measure vulnerability of TCP implementations
- Develop measurement techniques to test for vulnerabilities
- “Resilience of Deployed TCP to Blind Attacks”, Luckie, et al. IMC2015 (Lead author Luckie now at U. Waikato)
- We ran this experiment once and need to rerun

Device	OS date	Blind reset in	Blind reset out	Blind SYN in	Blind SYN out	Blind data behind	Blind data ahead	Port range
Cisco 2610 12.1(13)	2002-01	× (A)	✓ (I)	× (R)	✓ (C)	× (A)	✓ (C)	seq.
Cisco 2610 12.2(7)	2002-01	× (A)	✓ (I)	× (R)	✓ (C)	× (A)	✓ (C)	seq.
Cisco 2650 12.3(15b)	2005-08	✓ (C)	✓ (I)	✓ (C)	✓ (C)	× (A)	✓ (C)	40785
Cisco 7206 12.4(20)	2008-07	✓ (C)	✓ (I)	✓ (C)	✓ (C)	× (A)	✓ (C)	54167
Cisco 2811 15.0(1)	2010-10	✓ (C)	✓ (I)	✓ (C)	✓ (C)	× (A)	✓ (C)	46166
Cisco 2911 15.1(4)	2012-03	✓ (C)	✓ (I)	✓ (C)	✓ (C)	× (A)	✓ (C)	39422
Juniper M7i 8.2R1.7	2007-01	× (A)	✓ (I)	× (R)	✓ (I)	× (A)	✓ (C)	181
Juniper EX9208 14.1R1.10	2014-06	✓ (C)	✓ (I)	✓ (C)	✓ (I)	× (A)	✓ (C)	13769
Juniper MX960 13.3	2015-05	✓ (I)	✓ (I)	✓ (C)	✓ (I)	× (A)	✓ (C)	13033
Juniper J2350 12.1X46-D35.1	2015-05	✓ (I)	✓ (I)	✓ (C)	✓ (I)	× (A)	✓ (C)	12481
HP 2920 WB.15.16.0006	2015-01	✓ (C)	✓ (C)	✓ (C)	✓ (C)	✓ (I)	✓ (I)	14273
HP e3500 K.15.16.0007	2015-06	× (A)	✓ (I)	× (R)	✓ (C)	✓ (I)	✓ (I)	15611
Brocade MLX-4 5.7.0bT177	2014-10	✓ (I)	✓ (I)	✓ (C)	✓ (C)	✓ (C)	✓ (C)	const.
Pica8 Pronto3290 v2.6	2015-05	× (A)	✓ (I)	× (R)	✓ (C)	× (A)	× (A)	HBPS

Laboratory testing of blind TCP attacks against BGP-speaking router and OpenFlow-speaking switches.

Need 5: Identifying Grey Market IPv4 Address Transfers (Year 2)

1. Detect (evidence of) prefix ownership change using BGP data, DNS and targeted traceroute data
2. Need history of router-level paths and RTT information from Ark monitors toward routed prefixes
3. Develop signatures for prefix transfers due to non-BGP speakers changing upstream providers

Need 6: Router-Level Topology Mapping on Demand (Year 2)

- Researchers have asked for real-time measurement analysis and extraction of router topologies on specific sets of addresses
- Implement a system with an interface for researchers to specify target prefixes/addresses to run our **midar** alias resolution tool from a set of Ark probes.

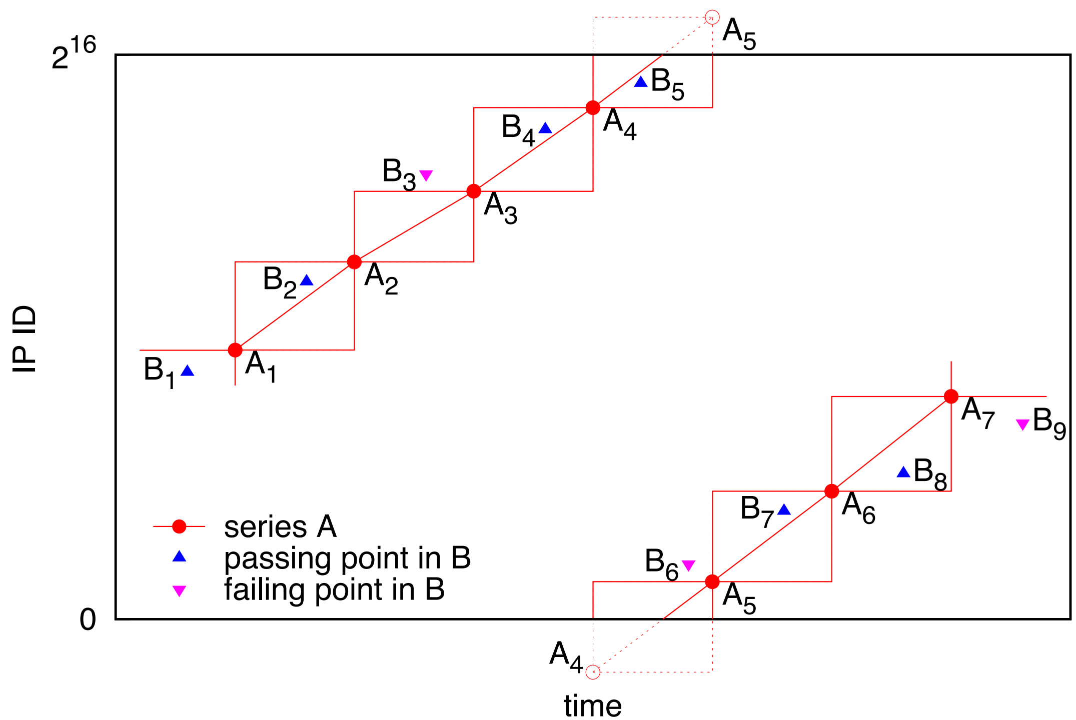

Illustration of the sample-wise execution of the Monotonic Bounds Test (MBT).

Benefits

- Enhanced scientific understanding and technical capabilities for empirically grounded macroscopic assessment of the global internet
- Leverage DHS supported, stable, trusted infrastructure for gathering comprehensive, trustworthy measurements of security-relevant properties and behavior of the global Internet
- Could run some of these measurements elsewhere but with less trust/confidence/precision

9

- To some degree, each of our approaches competes with and complements RIPE Atlas (http://atlas.ripe.net/)
- Others infrastructures might be applied but stakeholders have different focus or lack community interest and access
  - Internet Atlas (http://internetatlas.org/)
  - iPlane datasets (http://iplane.cs.washington.edu/data/data.html
  - zMap (https://zmap.io/), with results (https://censys.io)
  - ISI Census (http://isi.edu/ant/address)
  - Dyn (http://www.dyn.com/)

Contact Information

k claffy
CAIDA/UCSD
kc@caida.org
858-534-8333

SDSC
SAN DIEGO SUPERCOMPUTER CENTER

caida

UC San Diego