IODA - Internet Outages: Detection & Analysis

Alberto Dainotti

alberto@caida.org

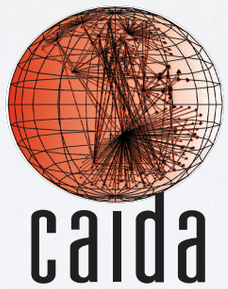

Center for Applied Internet Data Analysis
University of California, San Diego

CAIDA

intro

The Center for Applied Internet Data Analysis (CAIDA) is an independent analysis and research group based at the University of California's San Diego Supercomputer Center. CAIDA investigates both practical and theoretical aspects of the Internet.

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

CAIDA

research highlights

- topology analysis
  - Internet-scale router alias resolution
  - comparing IPv6 & IPv4 topology
  - Internet topology data sharing
- security & stability
  - large-scale Internet outages
  - botnet activity
  - BGP hijacks
- Internet peering analysis
  - inferring AS relationships
  - AS ranking
- interconnection economics
  - modeling peering strategies
  - transit pricing
- modeling complex networks
  - using hidden metric spaces
- geolocation analysis
  - comparing geolocation services
  - IP reputation vs. governance
- future Internet
  - IPv6
  - Named Data Networking
- visualization

Center for Applied Internet Data Analysis
University of California San Diego

http://www.caida.org/publications/

CHRONOLOGY

CAIDA and Internet Outages

- ***Jan/Feb 2011*** - Internet Kill Switch in Egypt and Libya

- ***Nov 2011*** - We present a novel approach to study Internet Outages by **combining different types of Internet measurements**

- ***Jan 2012*** - We present a study on the impact of natural disasters on the network infrastructure

- ***Sep 2012*** - NSF funds CAIDA to further develop our methodology and build an **experimental operational deployment** to monitor the public IPv4 Internet (**IODA**)

- ***2012 - 2015*** — more science and a lot of engineering

- ***Today*** a prototype that starts to be quite usable

BEFORE IODA

post-event manual analysis

- Country-level Internet Blackouts during the Arab Spring

Dainotti et al. "Analysis of Country-wide Internet Outages Caused by Censorship" ACM SIGCOMM IMC 2011

Egypt, Jan 2011
Government orders to shut down the Internet

- Natural disasters affecting the infrastructure

Dainotti et al. "Extracting Benefit from Harm: Using Malware Pollution to Analyze the Impact of Political and Geophysical Events on the Internet" ACM SIGCOMM CCR 2012

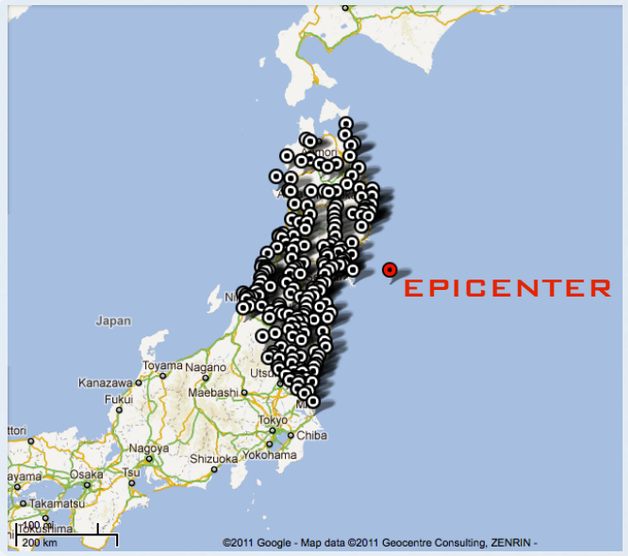

Japan, Mar 2011
Earthquake of Magnitude 9.0

BEFORE IODA

post-event manual analysis

Egypt, Jan 2011
Government orders to shut down the Internet

4 months of work

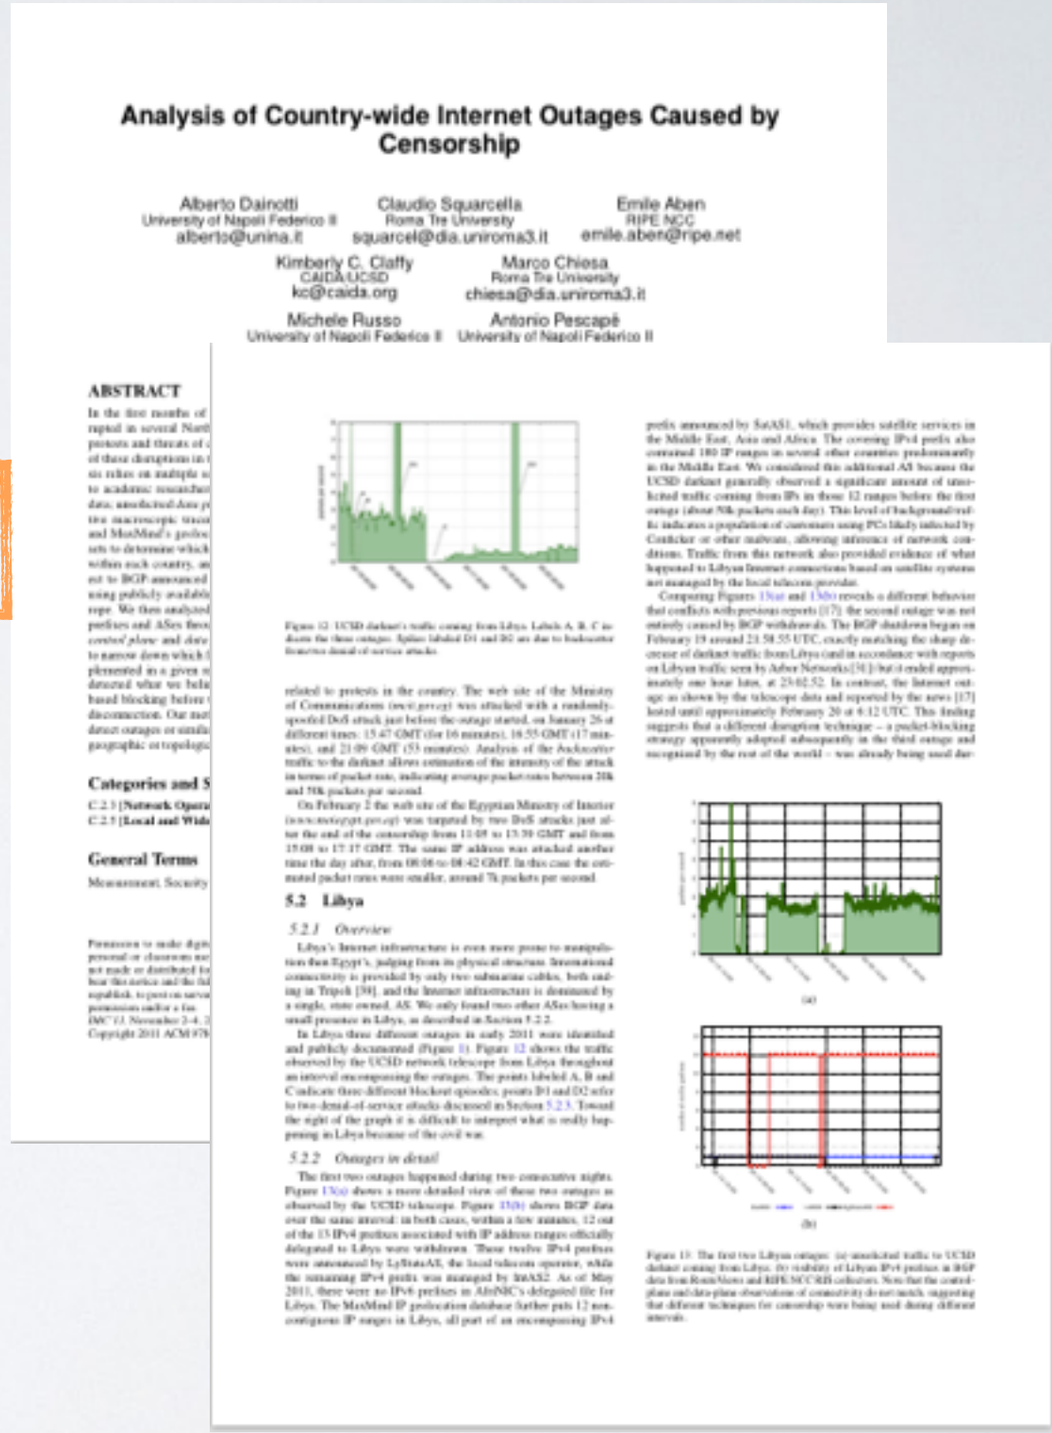

Analysis of Country-wide Internet Outages Caused by Censorship

Alberto Dainotti
University of Napoli Federico II
alberto@unina.it

Claudio Squarcella
Roma Tre University
squarcel@dia.uniroma3.it

Emile Aben
RIPE NCC
emile.aben@ripe.net

Kimberly C. Claffy
CAIDA/UCSD
kc@caida.org

Marco Chiesa
Roma Tre University
chiesa@dia.uniroma3.it

Michele Russo
University of Napoli Federico II

Antonio Pescapé
University of Napoli Federico II

ABSTRACT

In the first months of ...

Categories and S

General Terms

Measurement, Security

5.2 Libya

5.2.1 Overview

5.2.2 Outages in detail

IODA TODAY

live Internet monitoring

Last Christmas we made it possible for anybody to follow the North Korean disconnection almost live

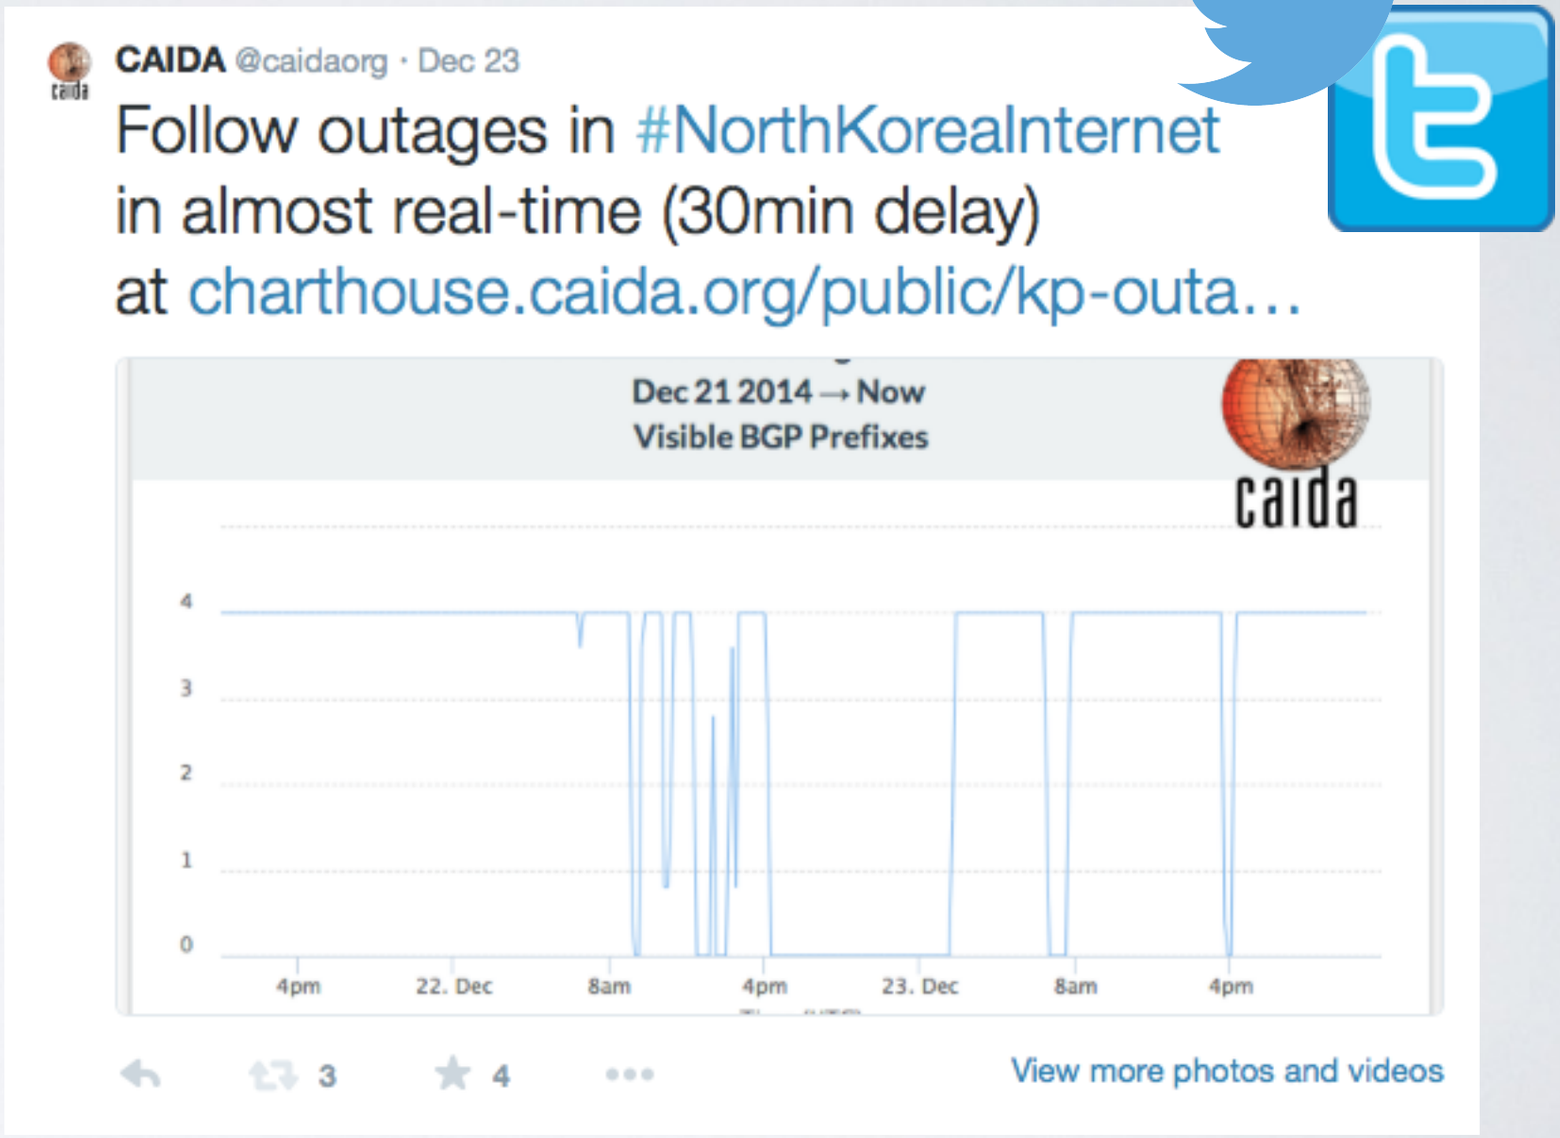

IODA TODAY

let's see how Internet providers are doing in the US

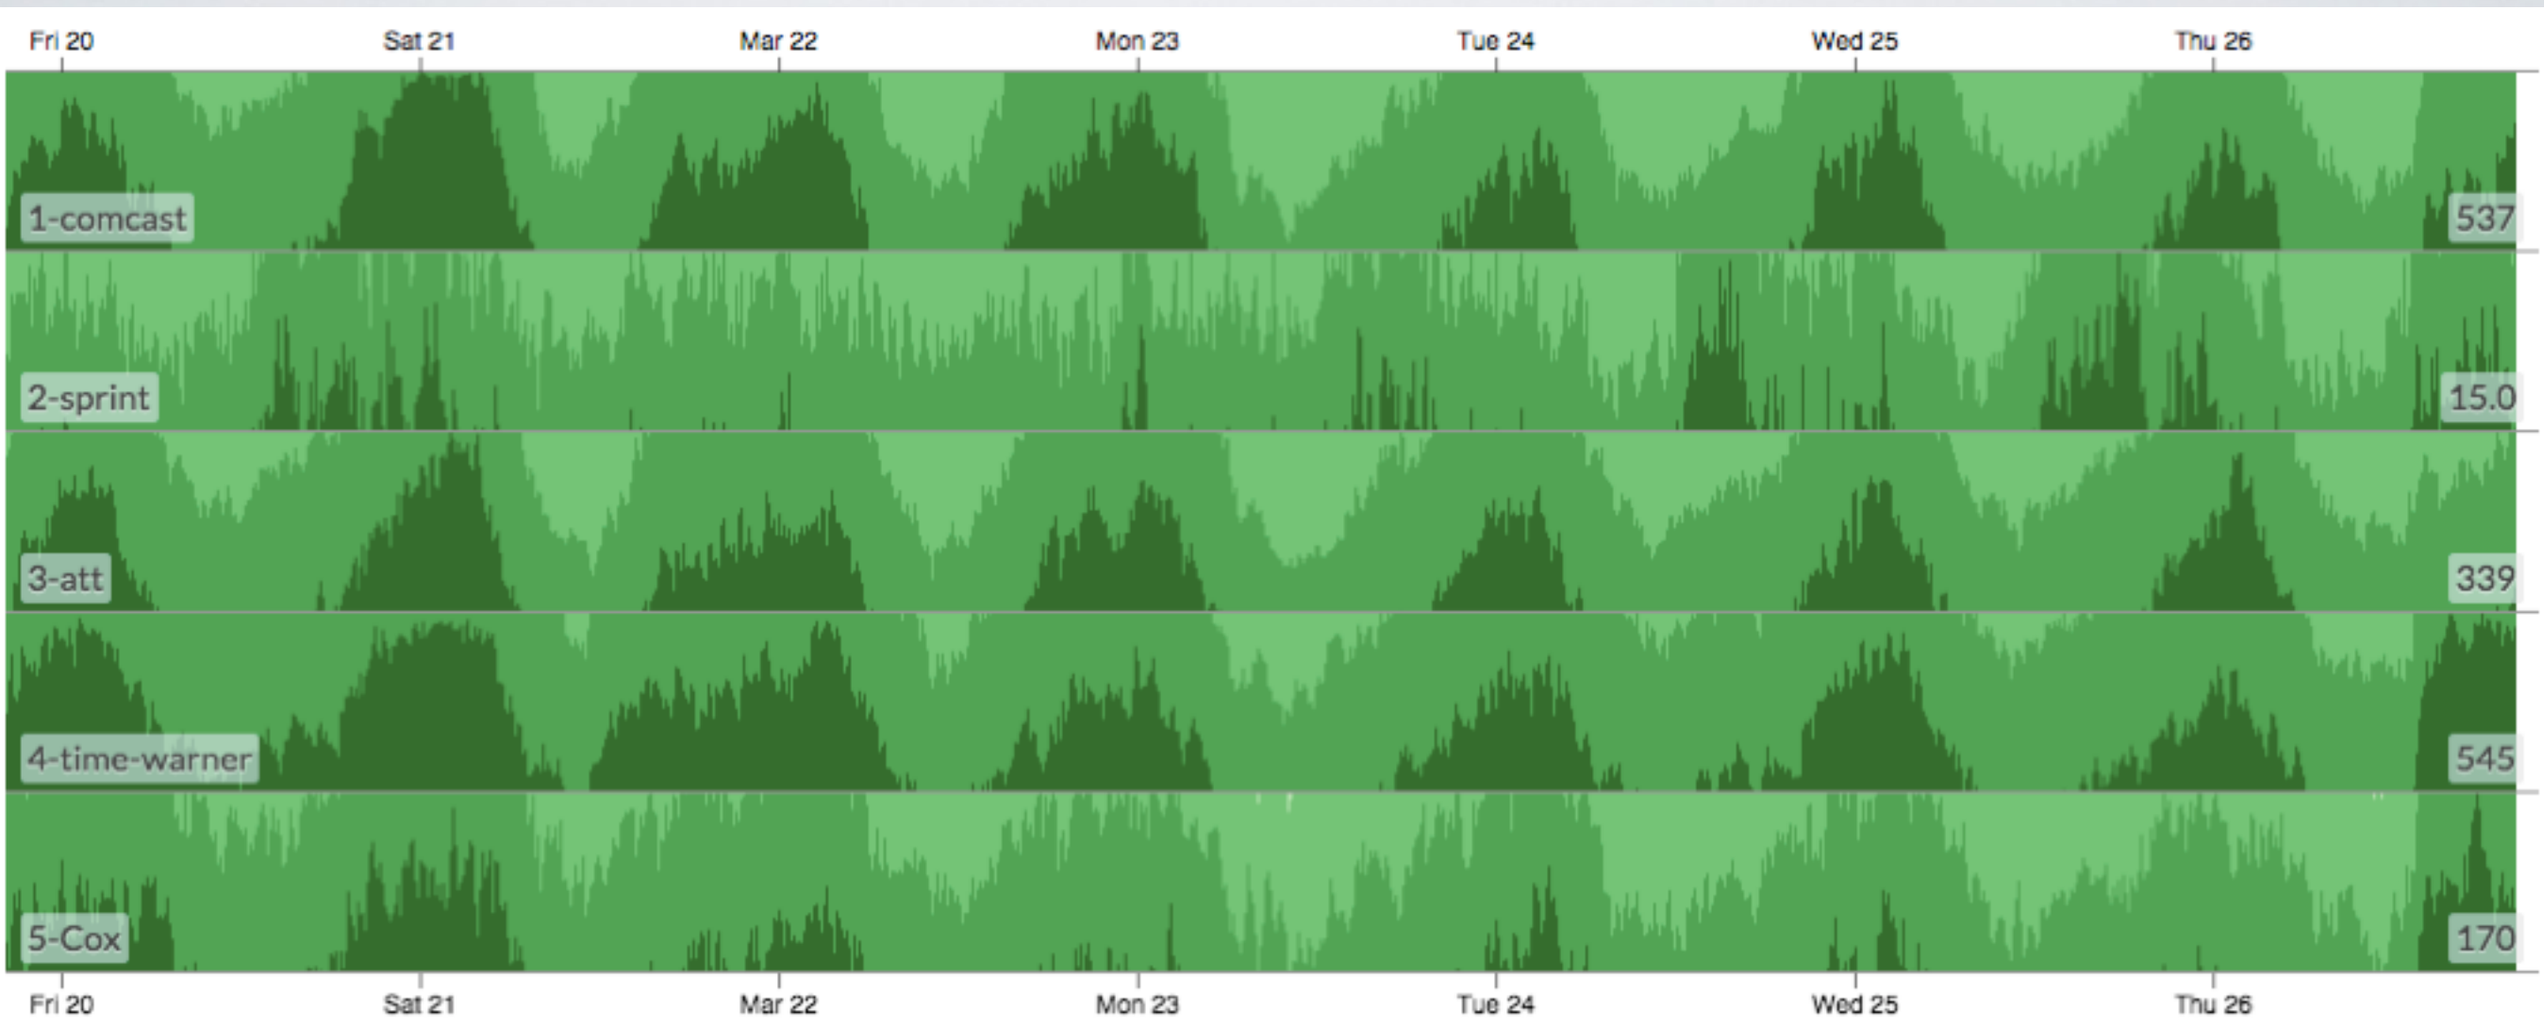

EVERYTHING LOOKS FINE…

IODA TODAY

Same type of graphs at the end of August 2014

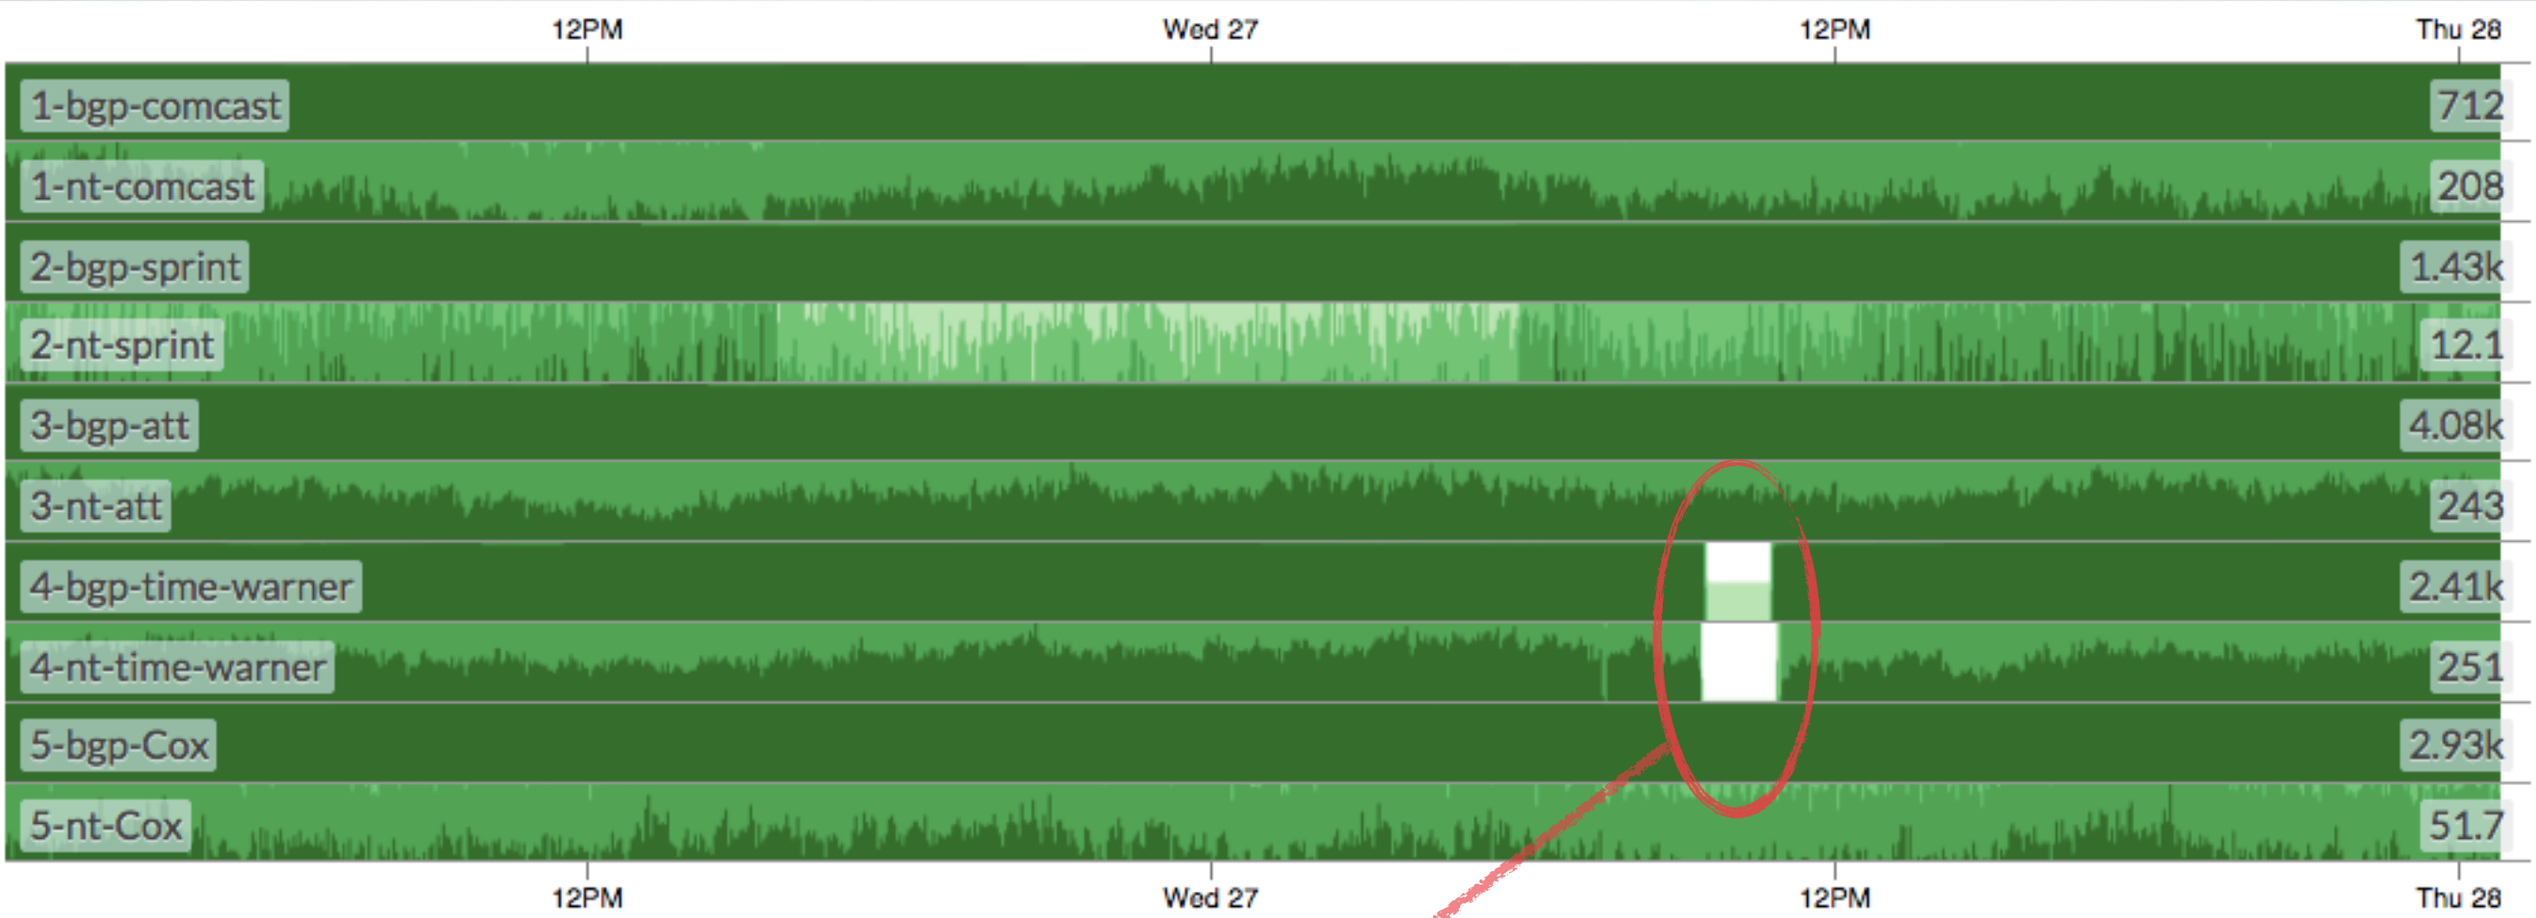

Series: *10* | # Points: *28800* | Data resolution: *minute*

TWO INDICATORS SHOW SOMETHING WENT WRONG WITH TIME WARNER

IODA

the system at a glance

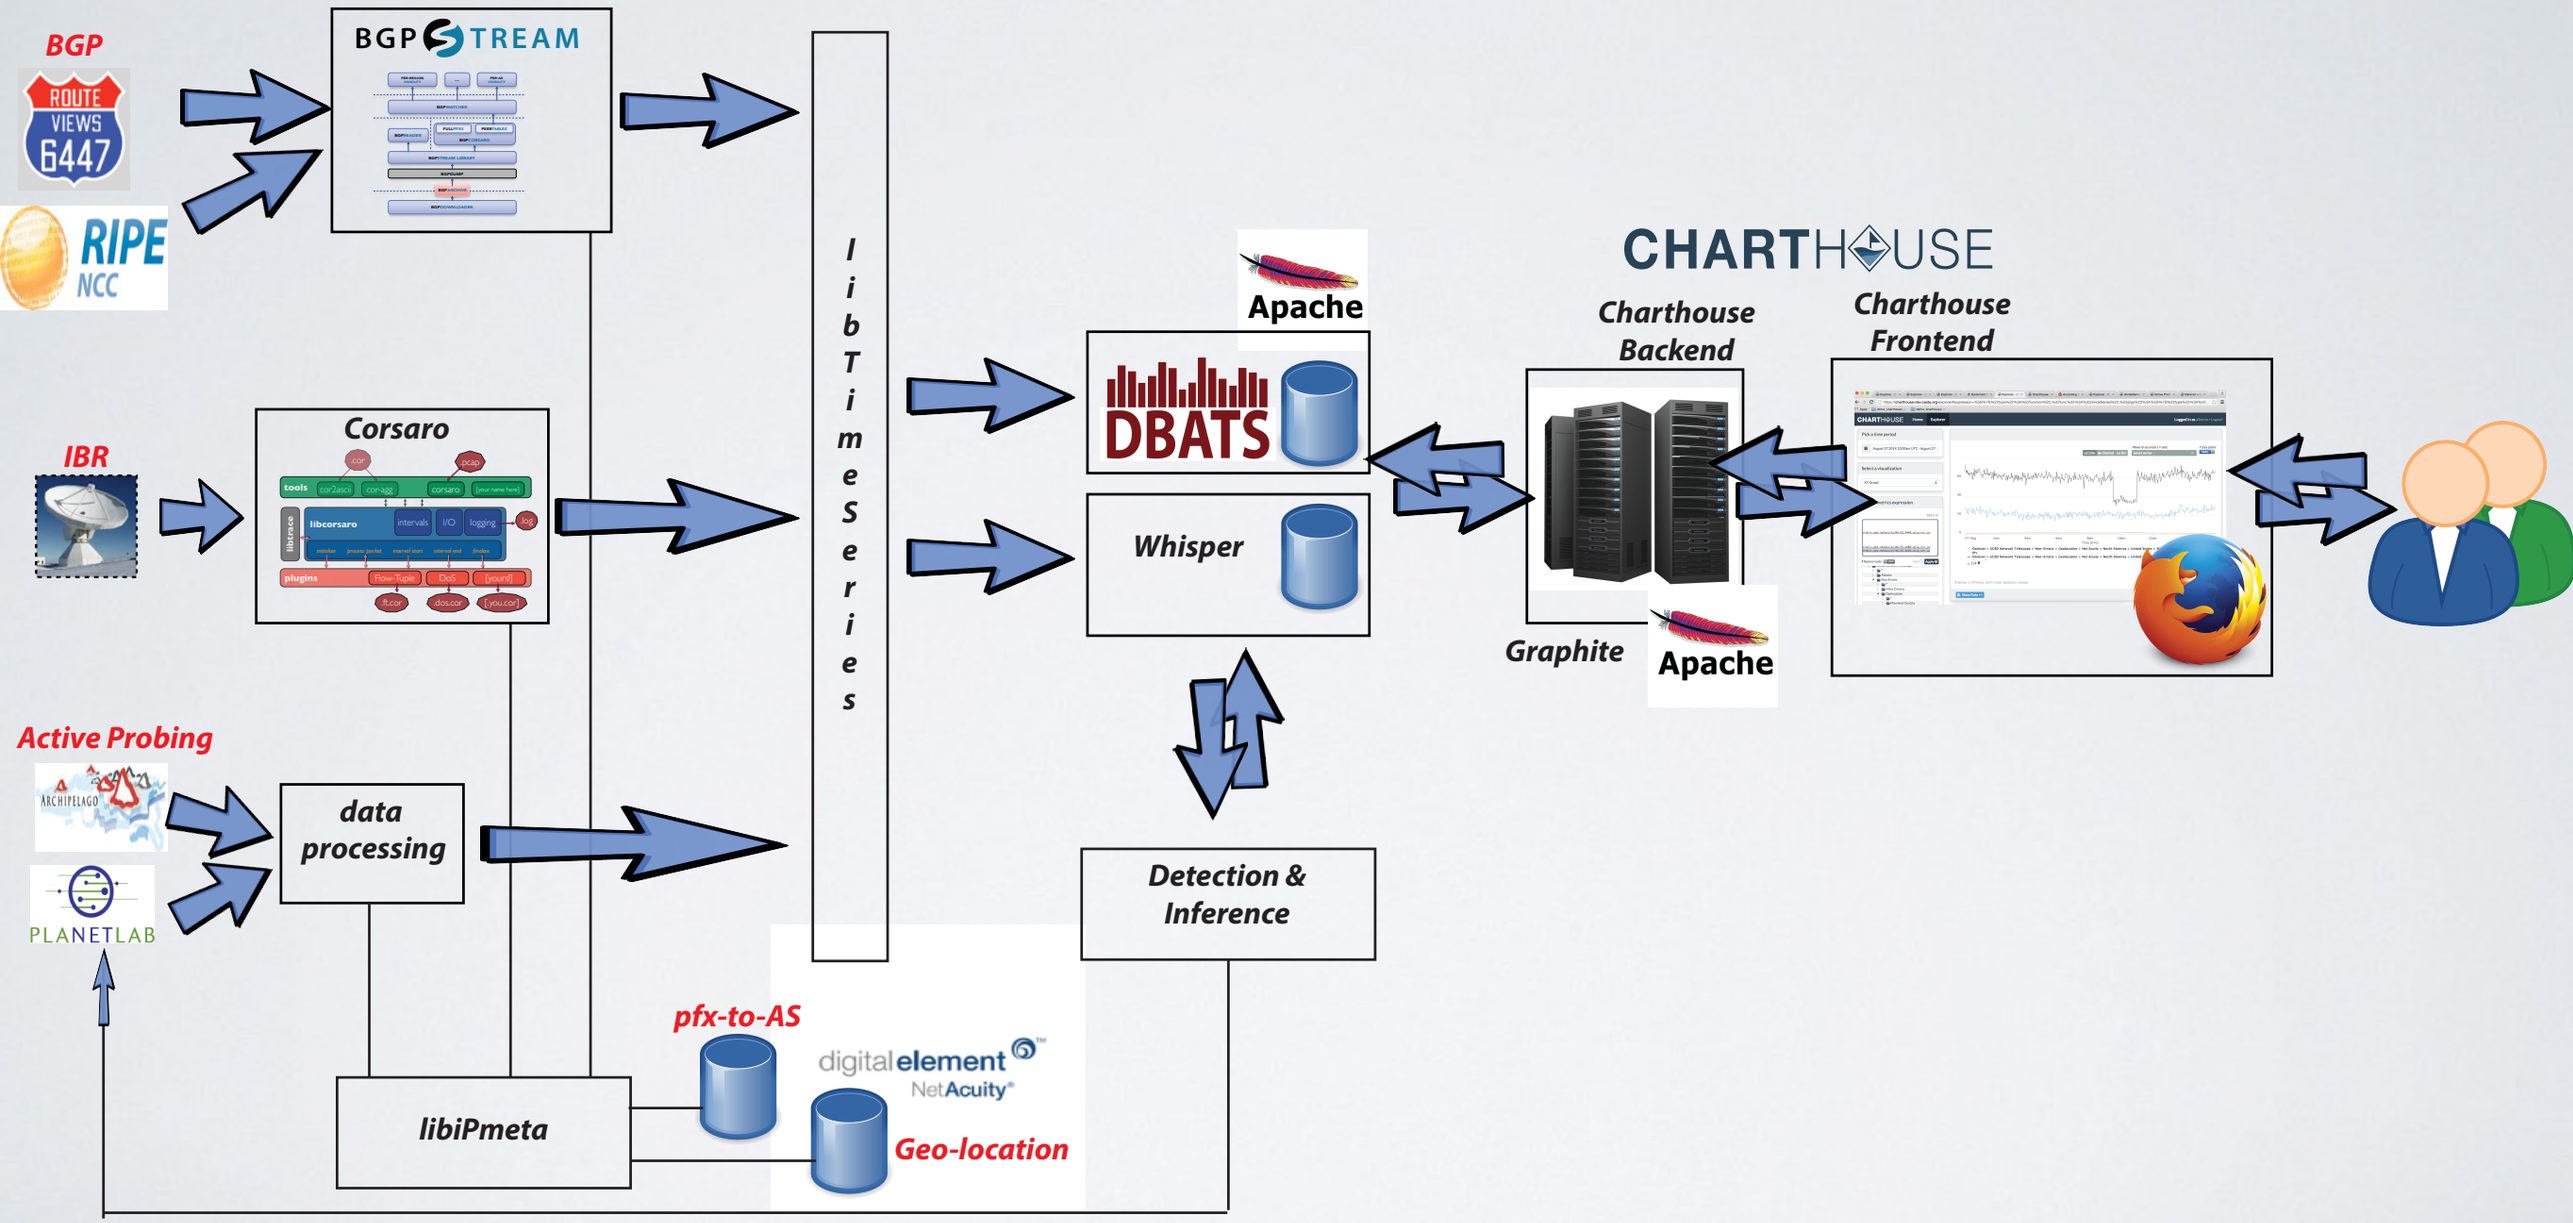

IODA

Internet Outages: Detection & Analysis

- **multiple** types of **sources and methodologies**
  - Routing Plane [BGP]
  - Data Plane
    - Active probing [pinging + traceroutes]
    - Passive [IBR]
  - *easy to plug new sources*

- **meta-data** to extract *liveness* signals for various aggregations *(countries, counties, cities, /24 address blocks, prefixes, ASNs)*
- combining signals to **detect & monitor**
- **trigger ad hoc active measurements** when an event is detected
- **visual interface** for analysis and dashboards CHARTHOUSE

BGP

Border Gateway Protocol

- BGP
  - The protocol that establishes routes between ISP networks (***autonomous systems***) all over the world
- Autonomous Systems (AS)
  - *"a set of routers under a single technical administration, using an interior gateway protocol and common metrics to determine how to route packets within the AS, and using an inter-AS routing protocol to determine how to route packets to other ASs." - RFC 4271*

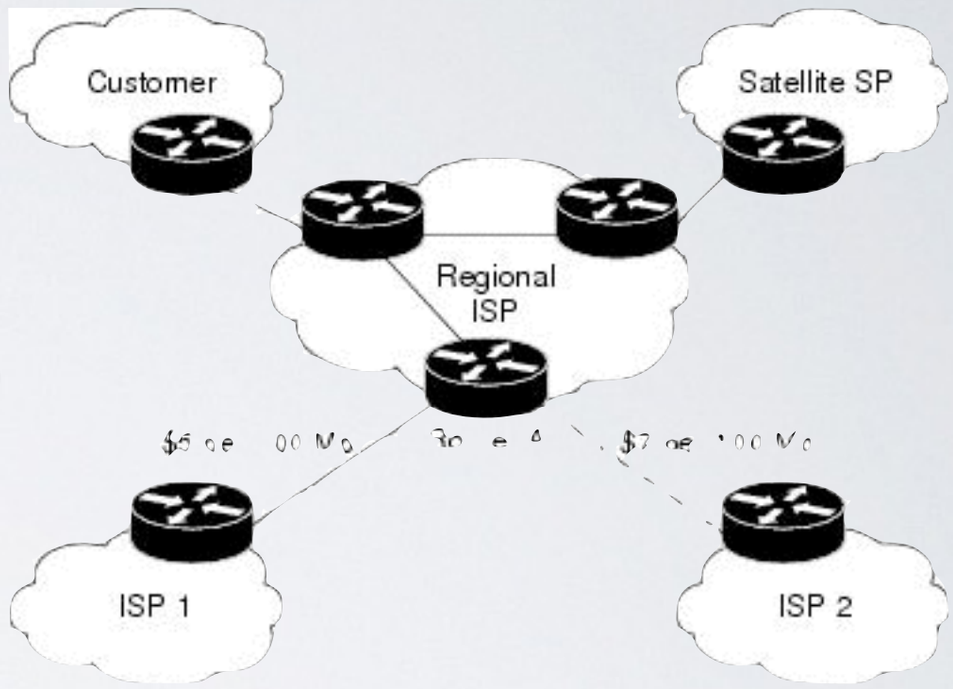

- Network Prefixes
  - Smaller networks are identified by a network address and a network mask
  - e.g. *prefix 192.172.0.0/16 is assigned to AS99 and is reachable by AS67 through the* ***AS path:*** *AS67➔AS44➔AS15➔AS99 ➔ 192.172.0.0/16*
  - AS paths are computed by exchanging ***BGP update messages***: *"Hey, I'm AS44 and can reach 192.172.0.0/16 through AS15➔AS99"*

BGP DATA COLLECTION

Route Collectors and Route Monitors

- BGP measurement projects establish peering sessions with ASes to receive their routing tables (no exchange of other traffic)

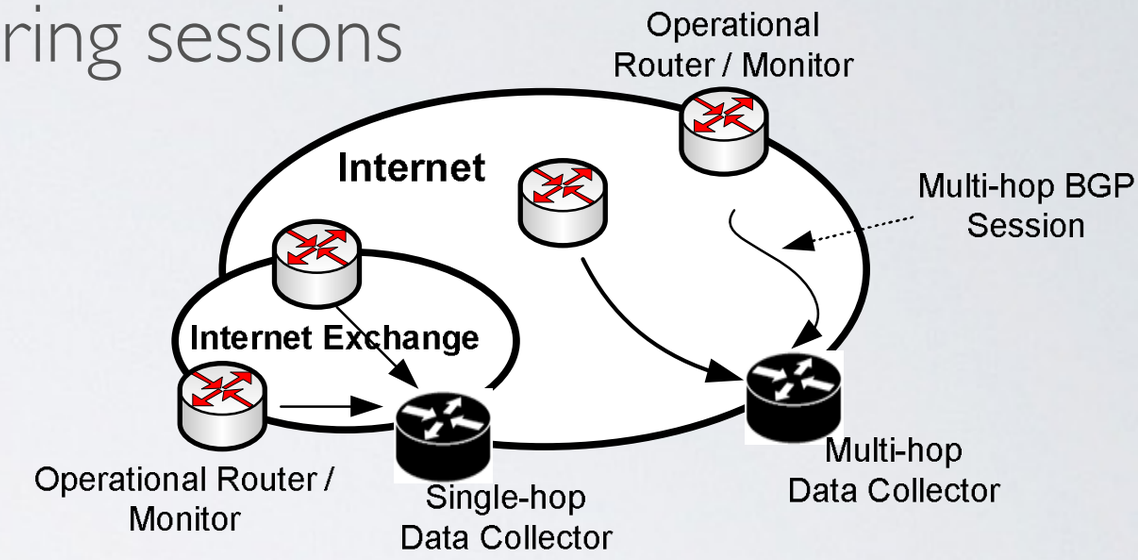

- RouteViews (Univ. Oregon): *371 peers*
- RIPE RIS (RIPE NCC): *508 peers*
- BGPmon (Colorado State Univ.): *330 peers*

http://www.routeviews.org

https://www.ripe.net/data-tools/stats/ris

http://bgpmon.netsec.colostate.edu

BGP DATA COLLECTION

RouteViews Collectors

http://www.routeviews.org

BGP DATA COLLECTION

IPv4/IPv6 BGPmon peers around the world

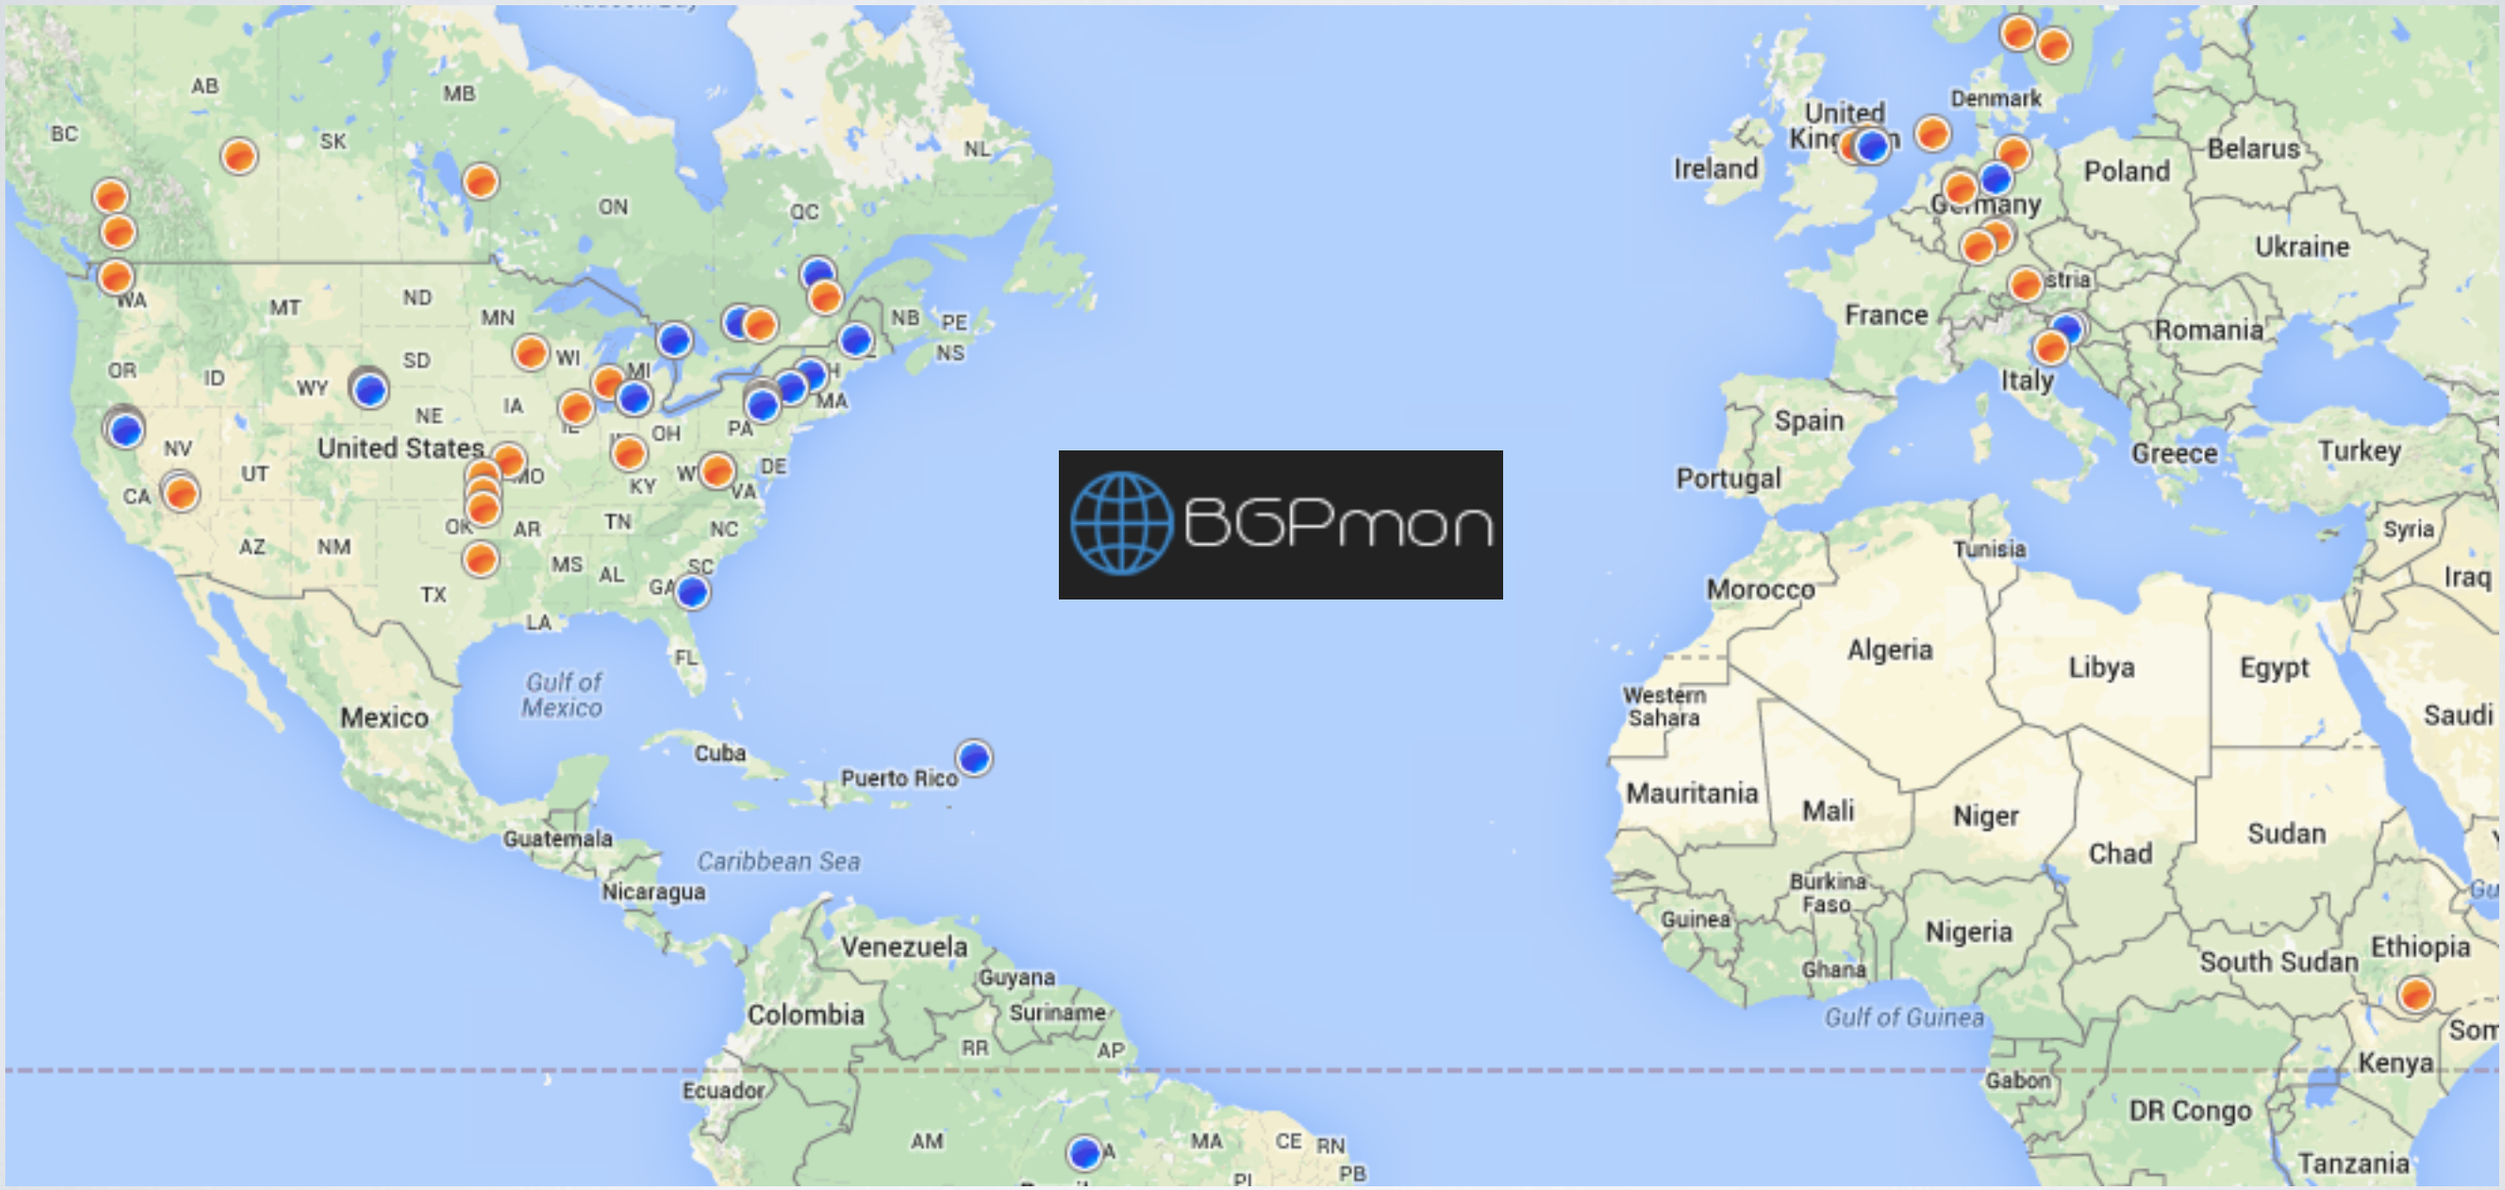

BGP: OUTAGES

North Korea (AS131279) - Dec 2014

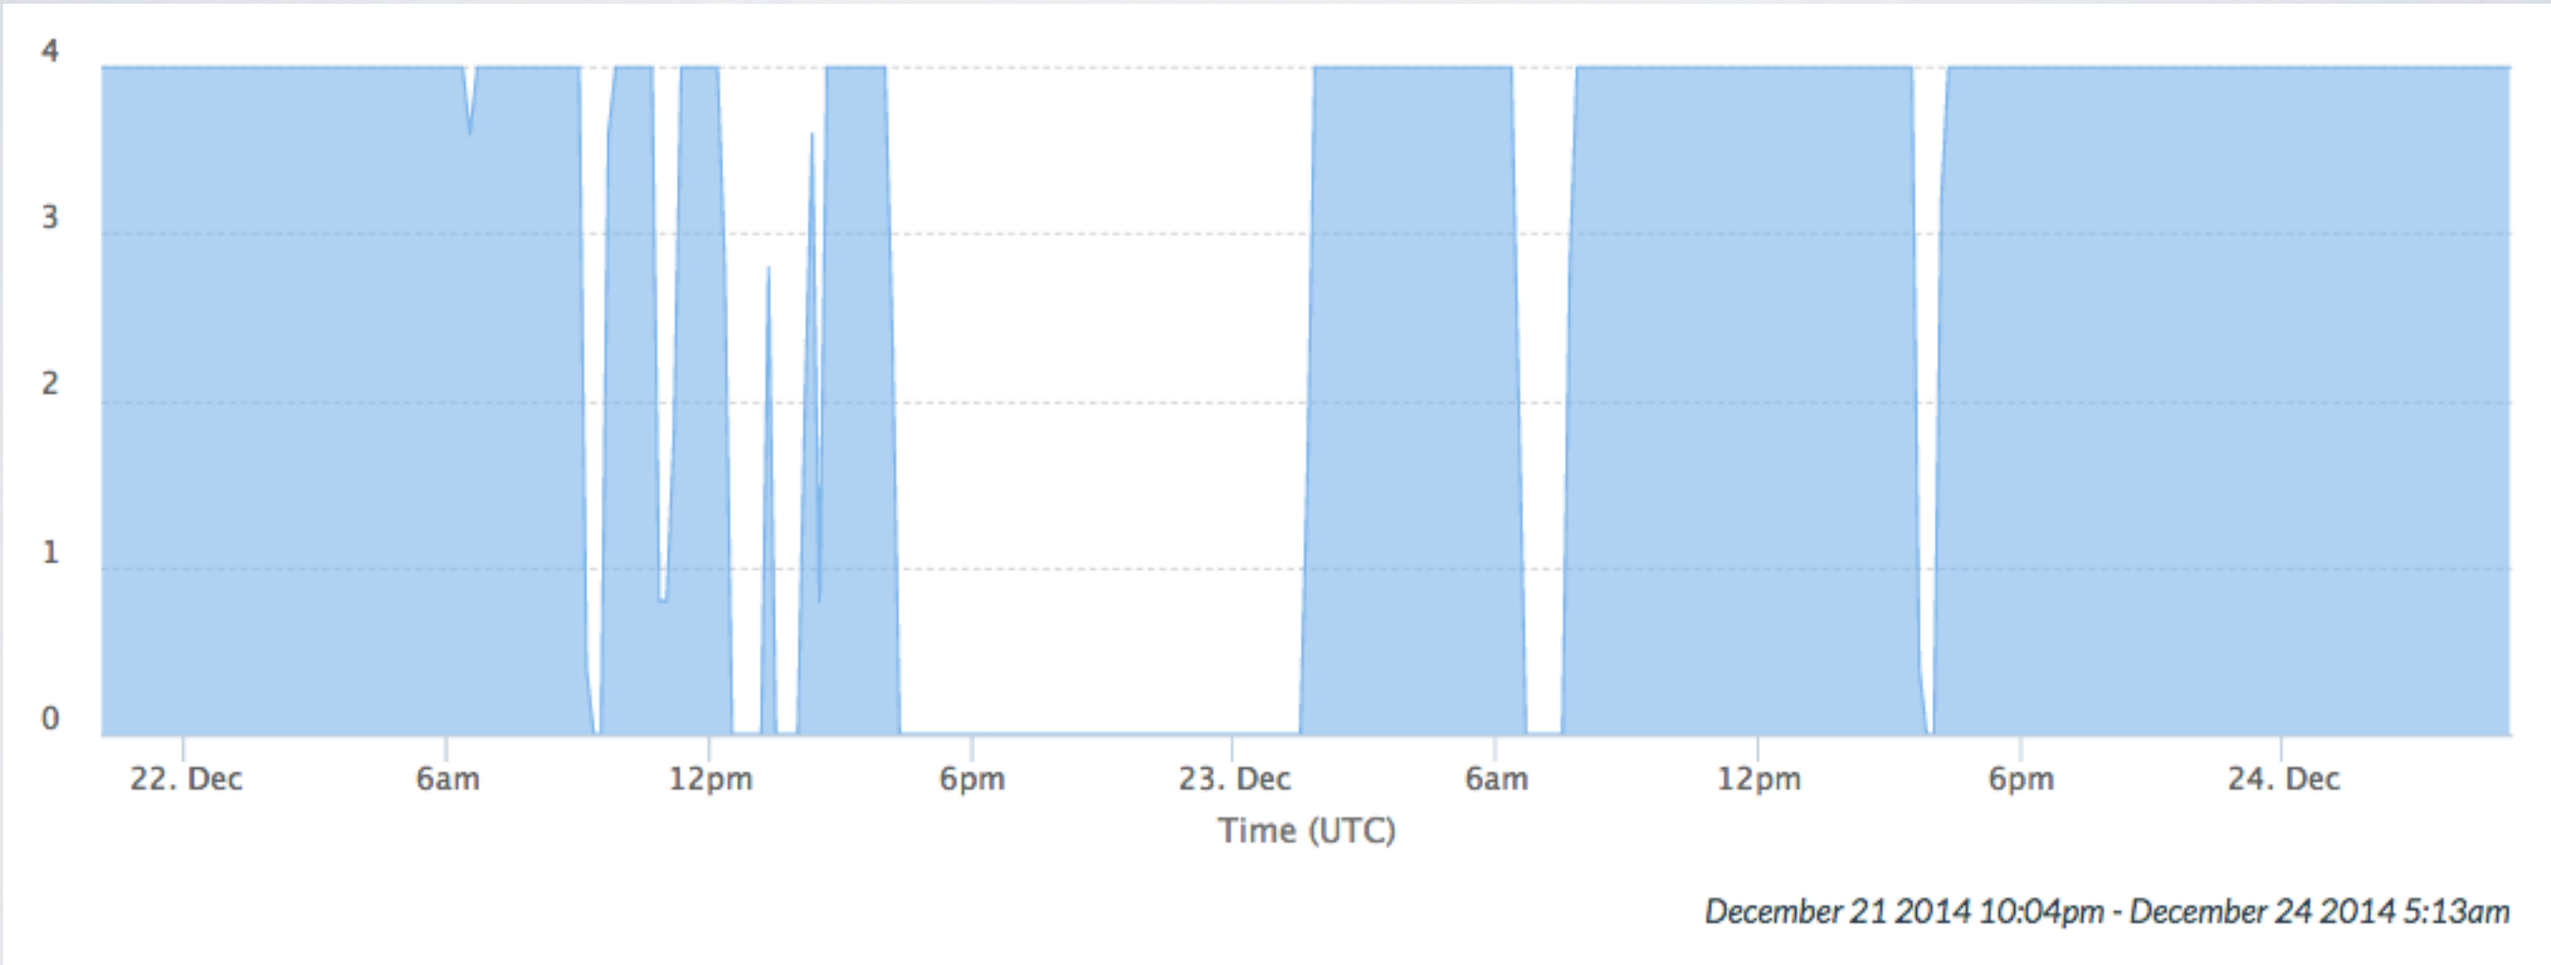

BGP: OUTAGES

Time Warner Cable - 27th Aug 2014

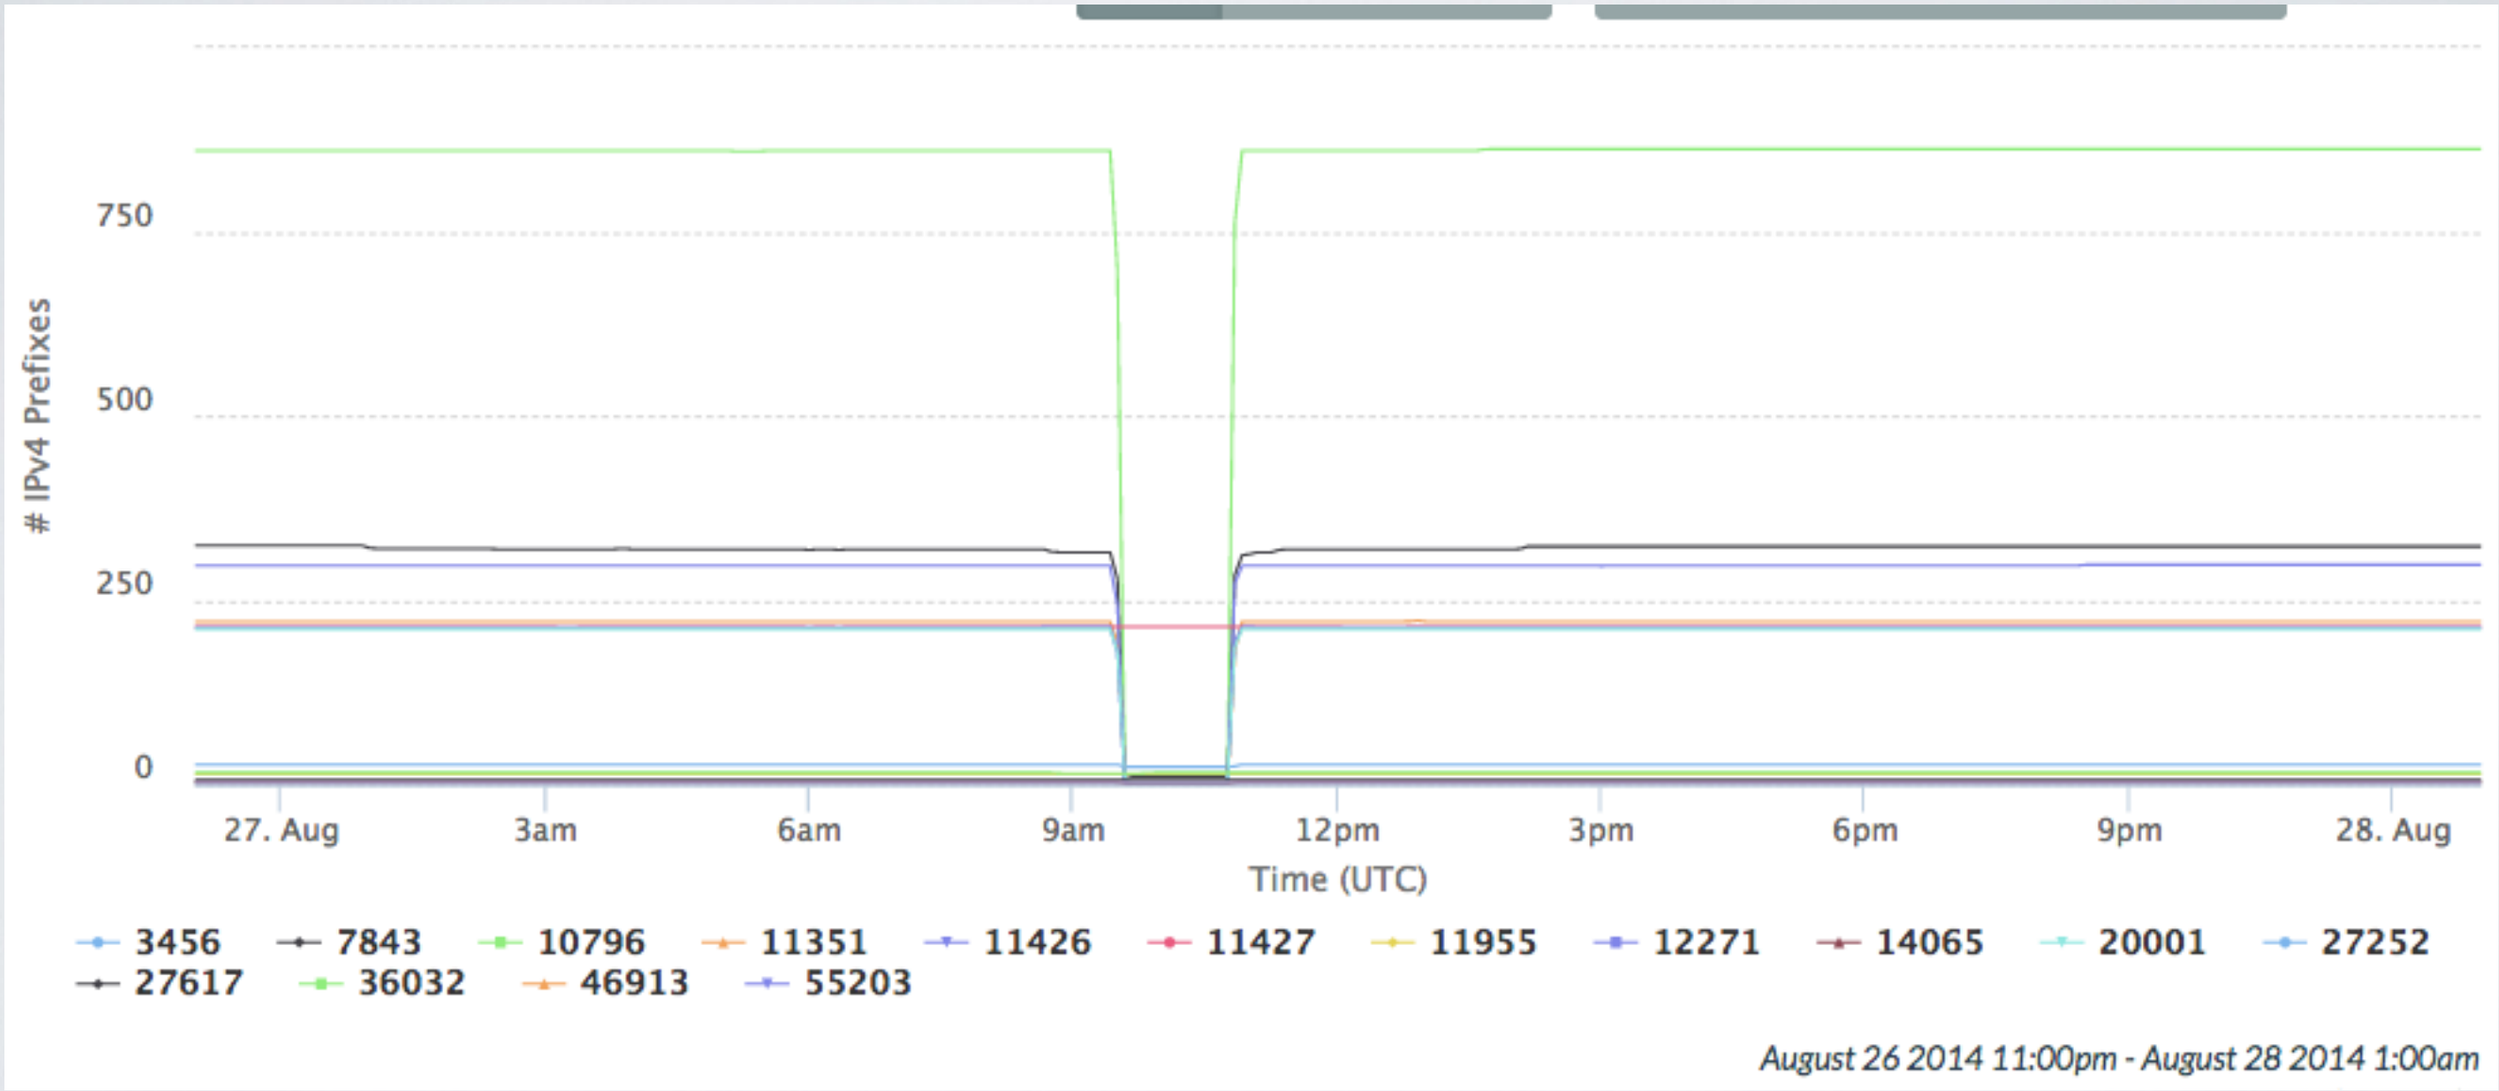

BGP: OUTAGES

Time Warner Cable - 27th Aug 2014

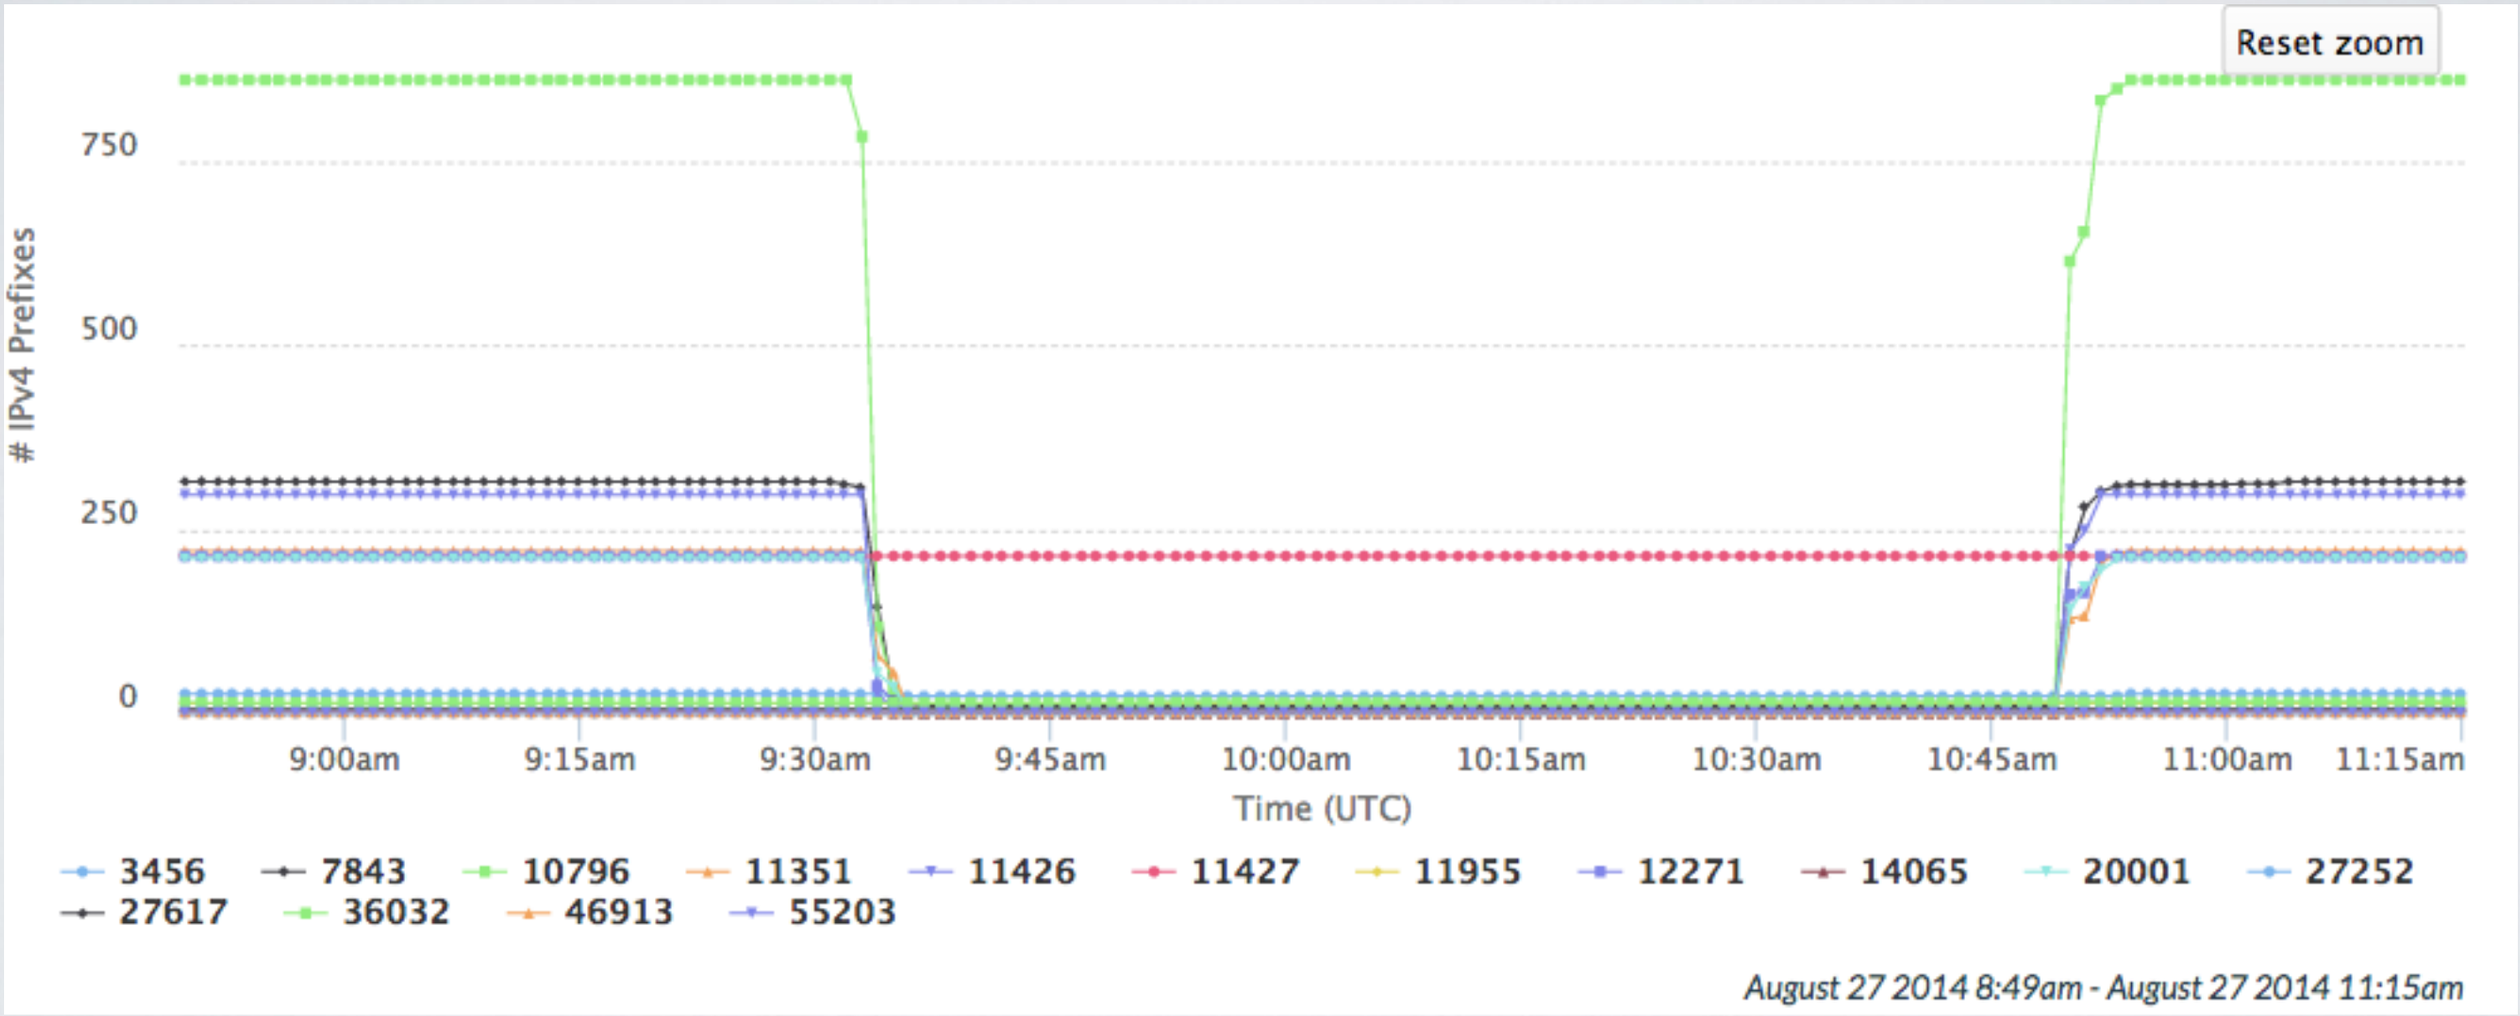

BGPSTREAM

our BIG DATA framework for BGP

- Enables **large-scale realtime BGP analysis**
- Stacked modular framework
- To monitor for **outages** and **hijacks** we run ~30 instances of BGPCorsaro in parallel
- data is filtered and aggregated
- creates a **global view of BGP** every minute

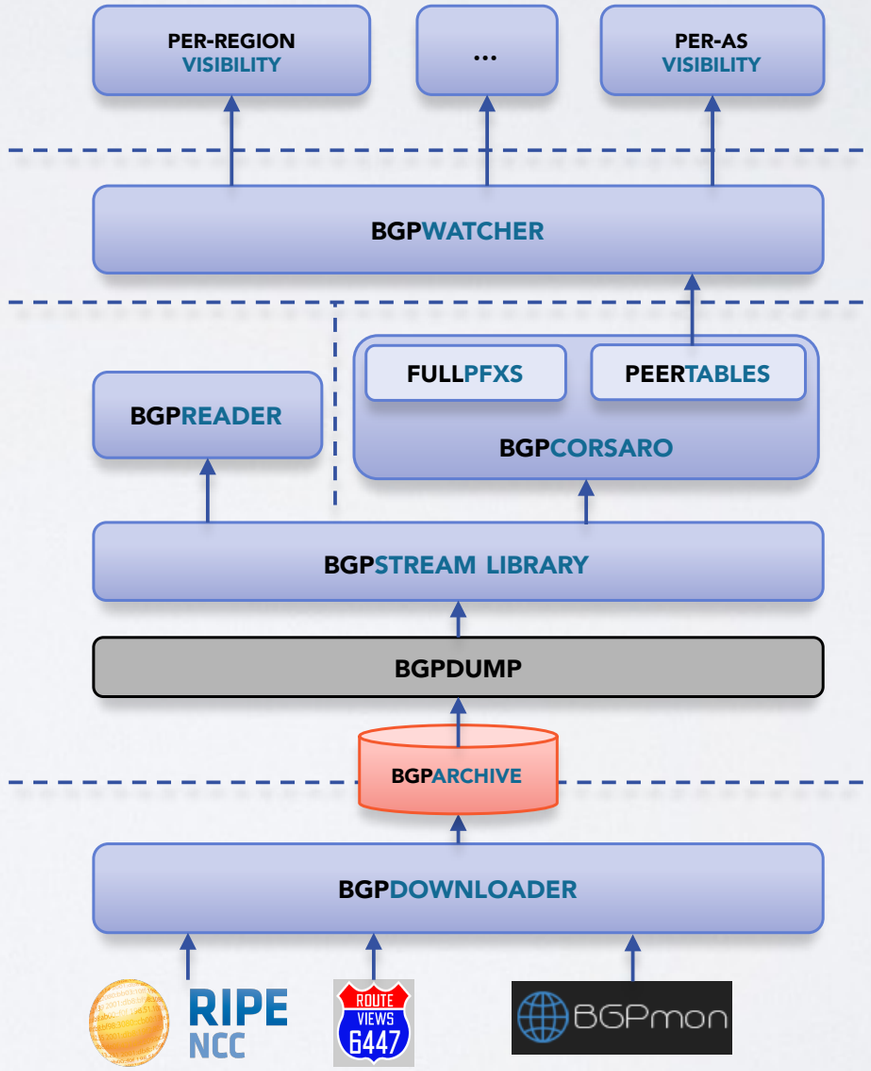

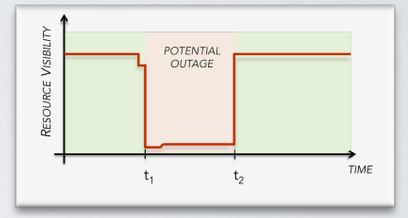

CHARTHOUSE

Center for Applied Internet Data Analysis
University of California San Diego

MONITORING..

..the measurement & monitoring infrastructures

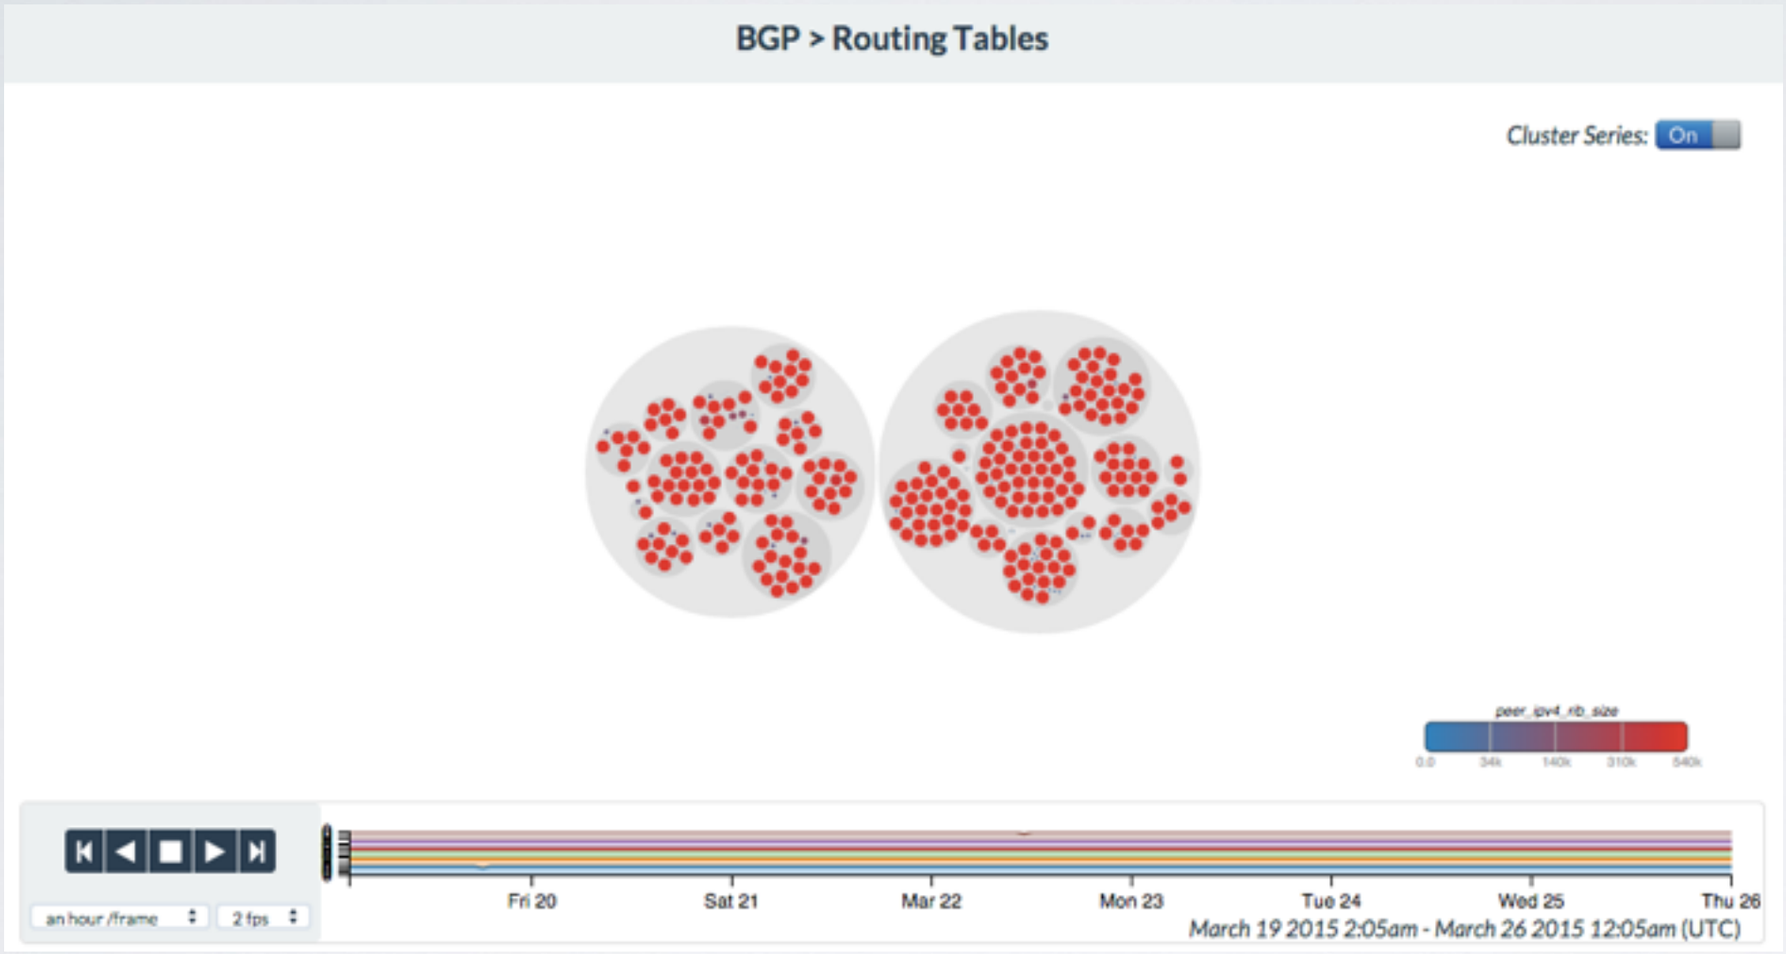

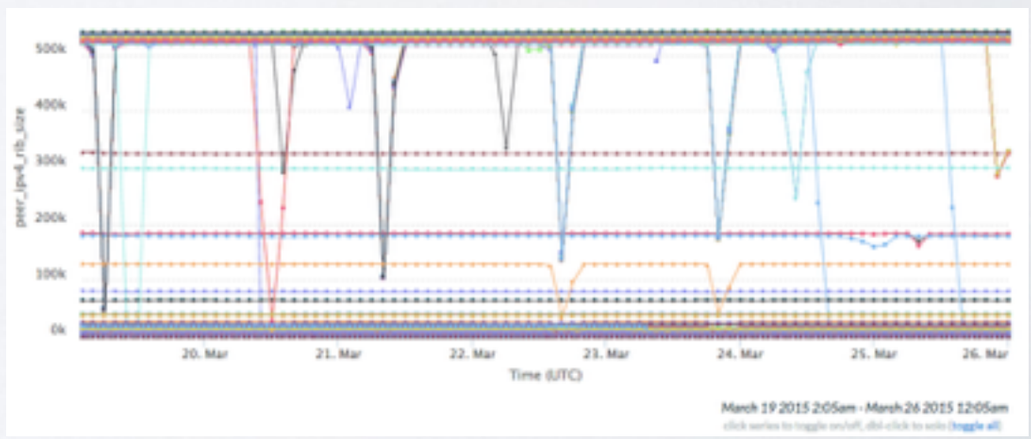

IBR

"Extracting benefit from harm.."

- Use *Internet Background Radiation (IBR), mostly generated by malware-infected hosts as a "signal"*

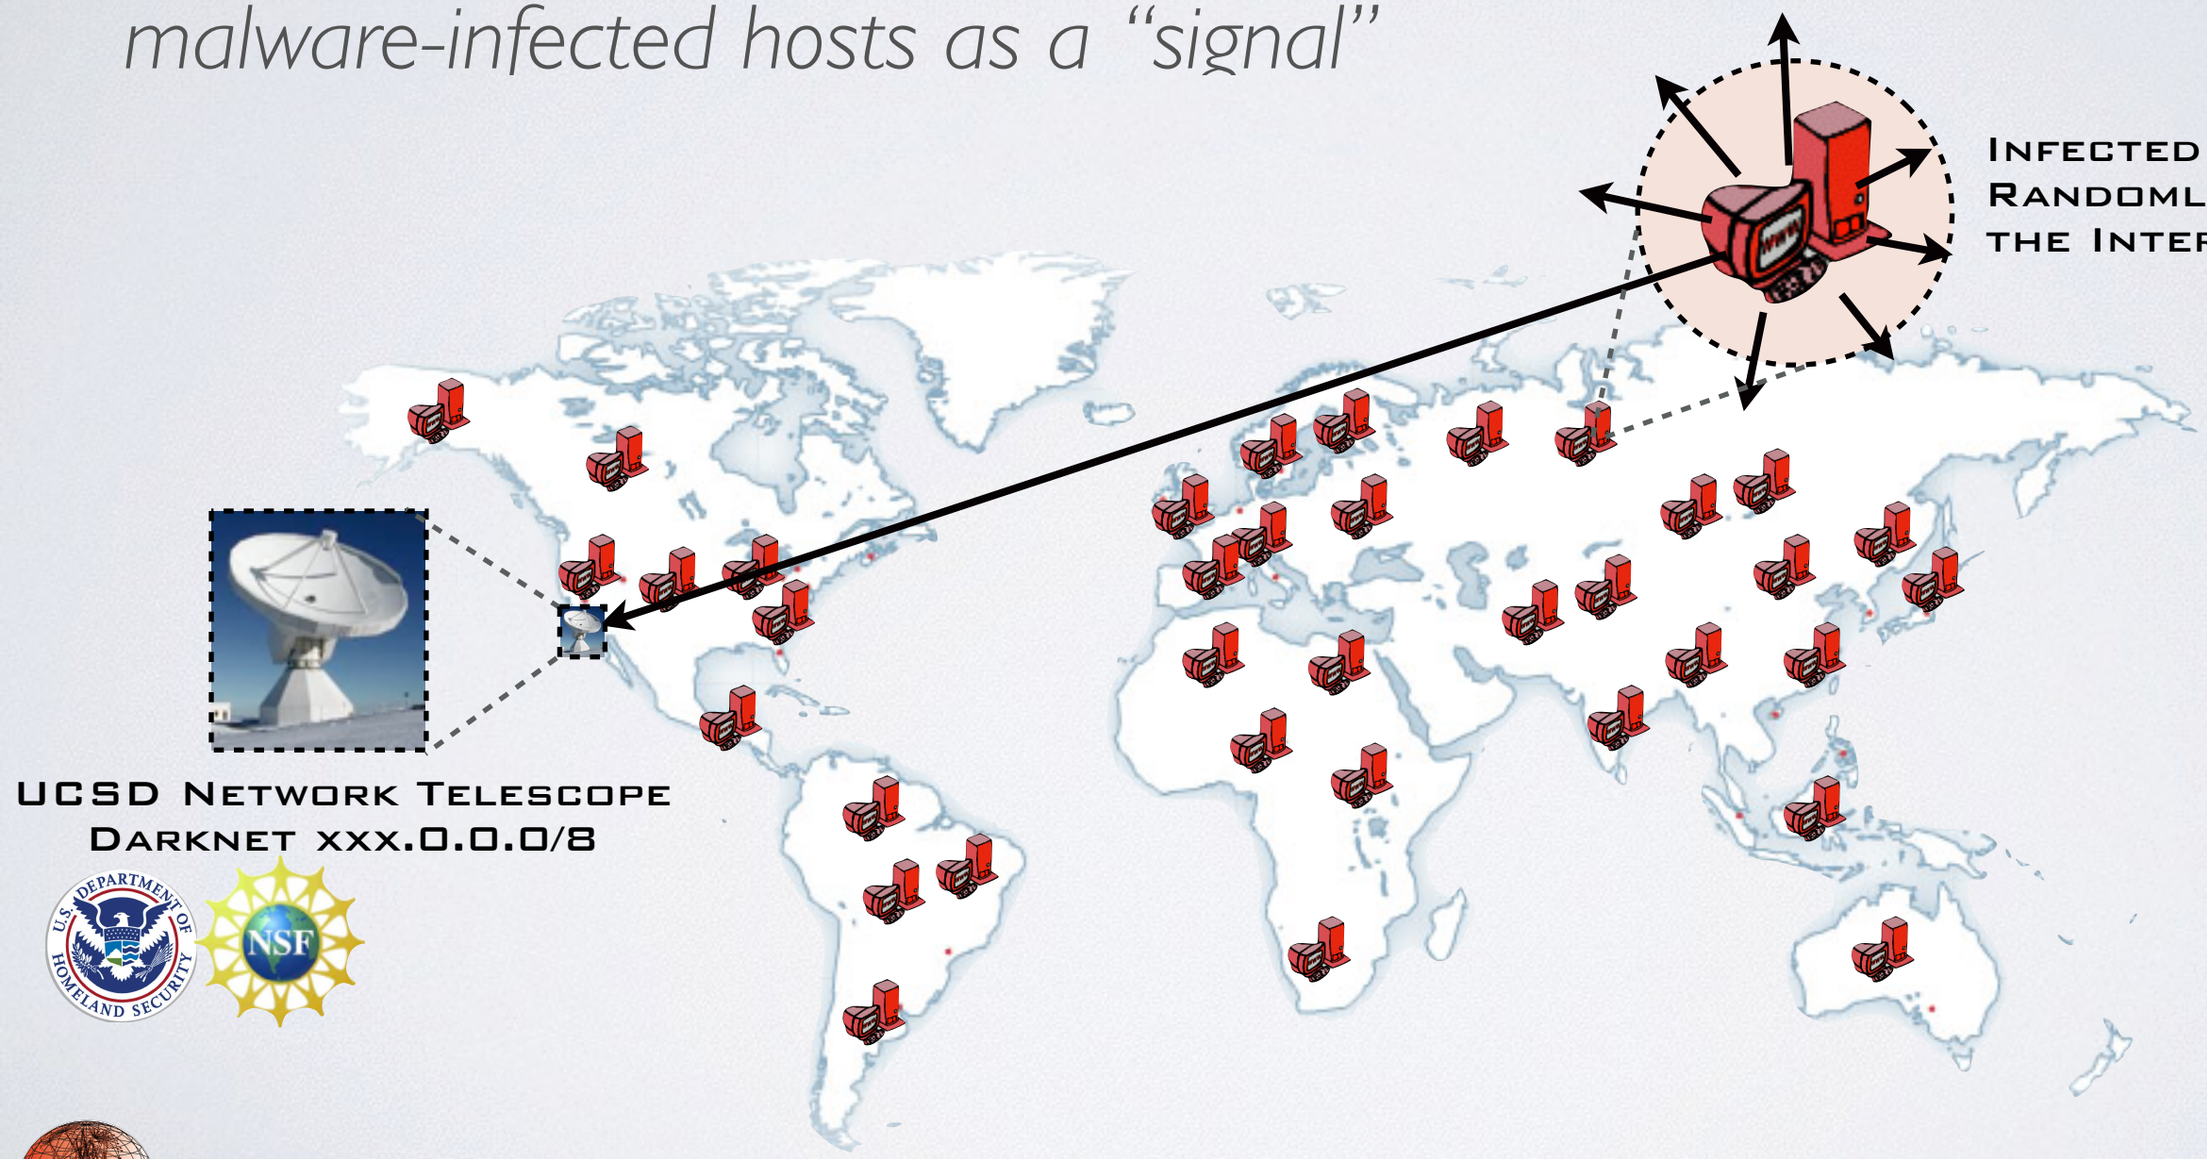

IBR

Time Warner Cable - 27th Aug 2014

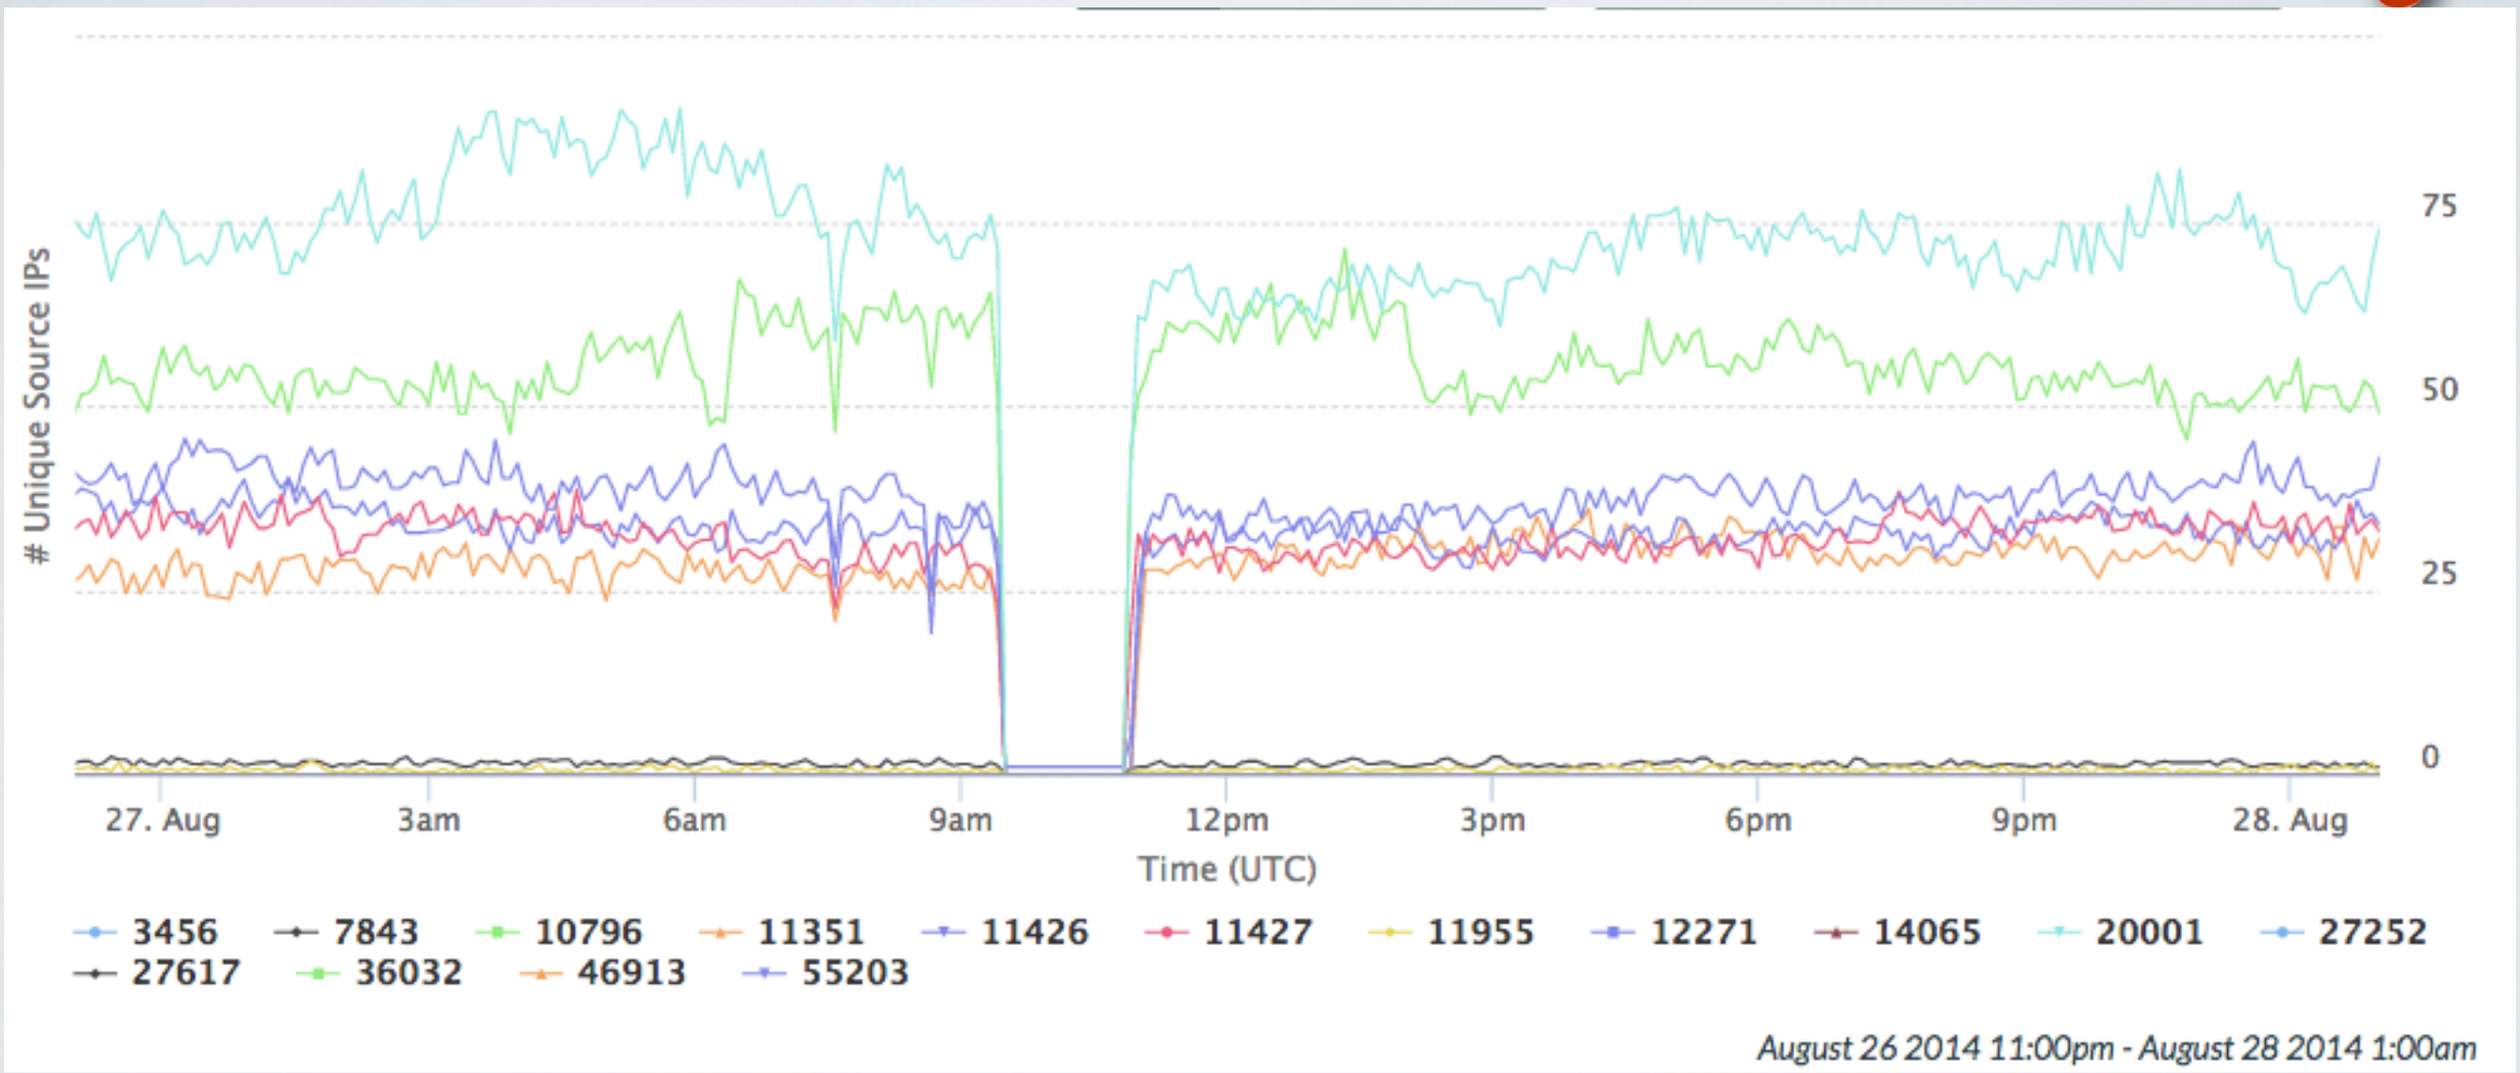

IBR

TWC outage: a look at few ISPs

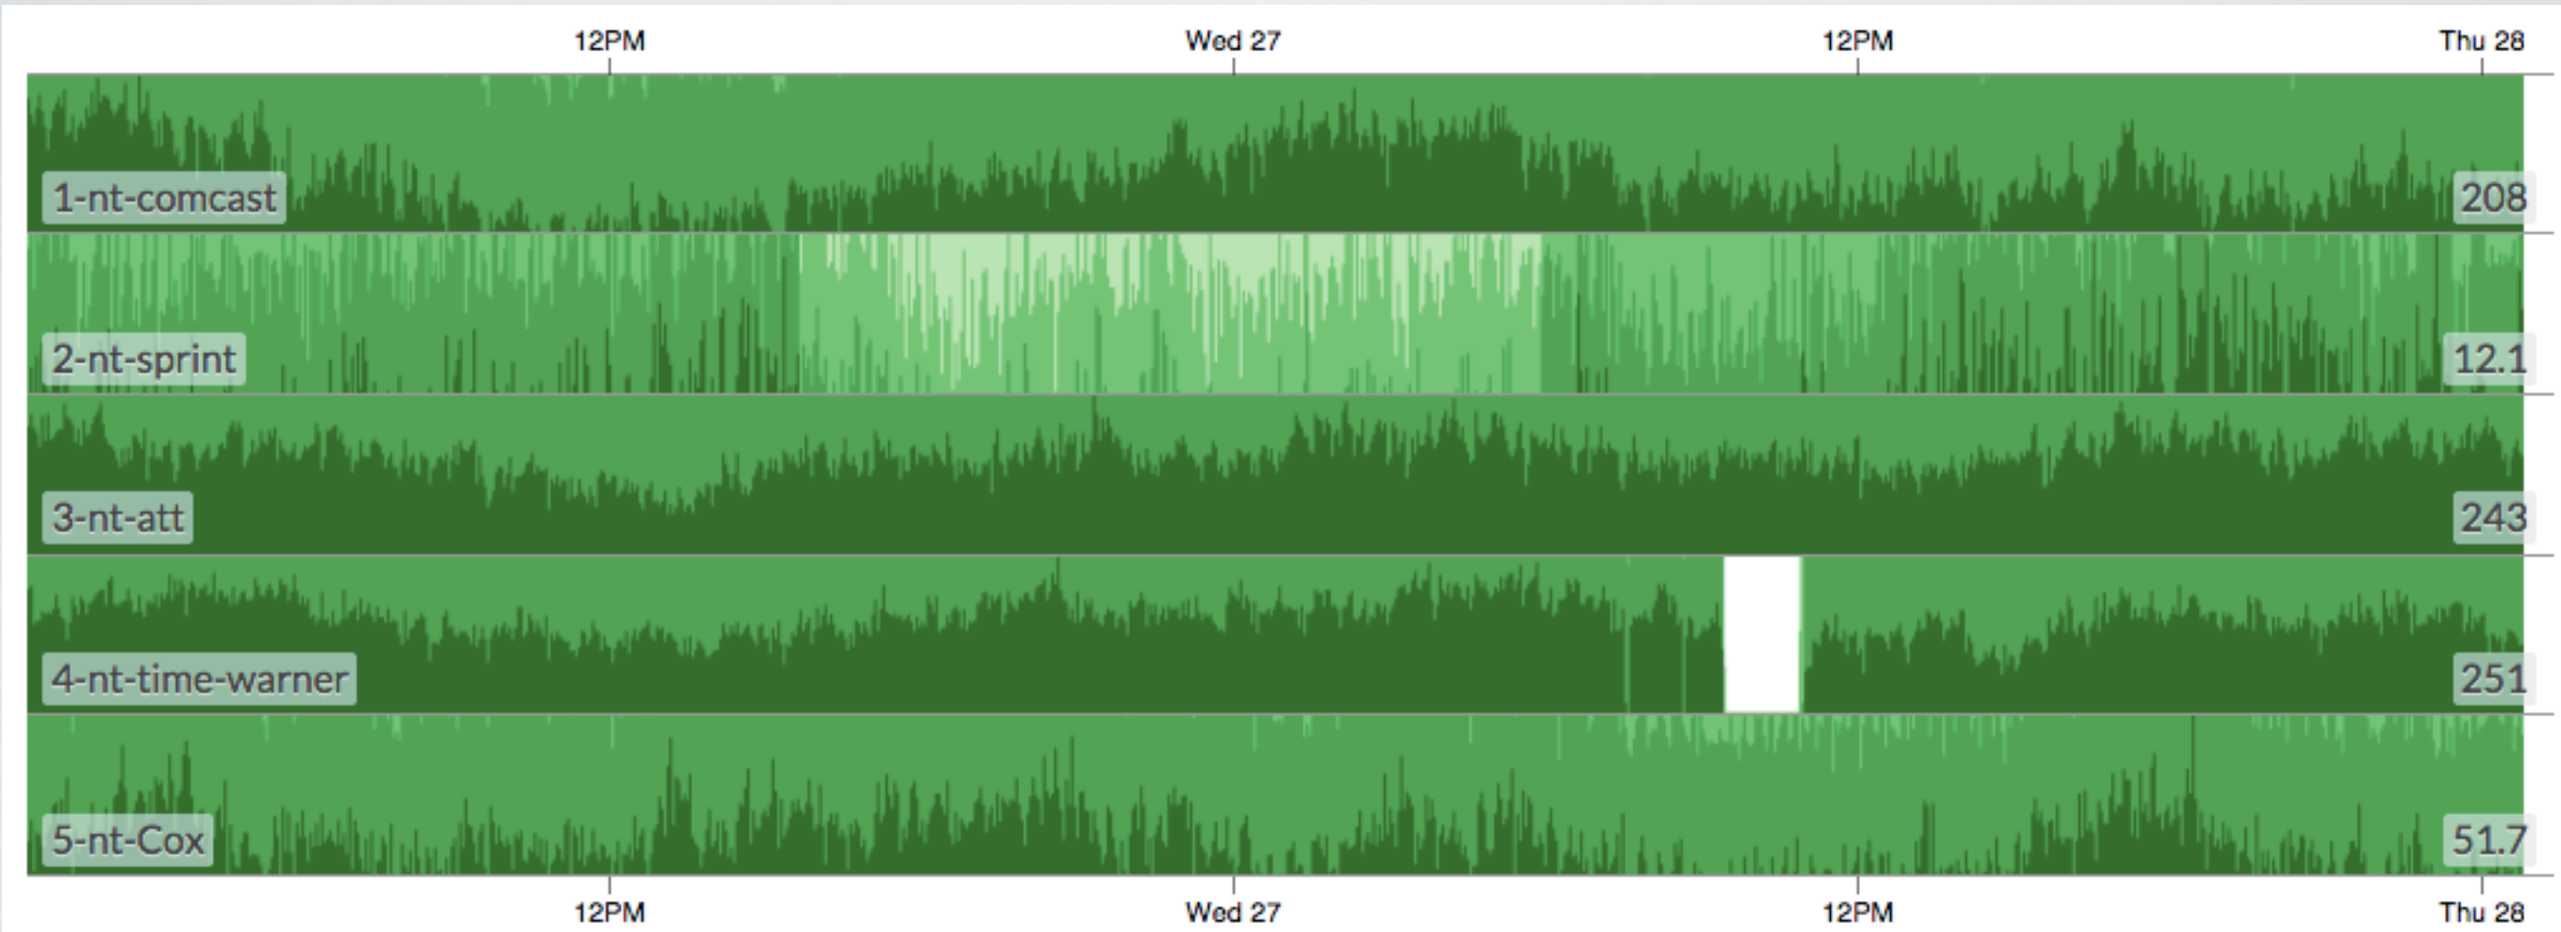

IBR + BGP

TWC outage: a look at few ISPs

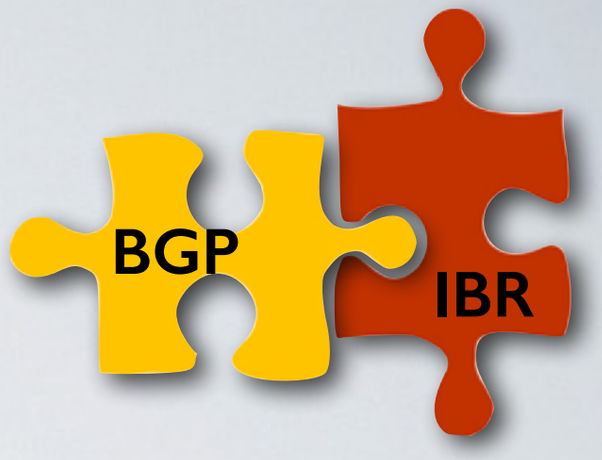

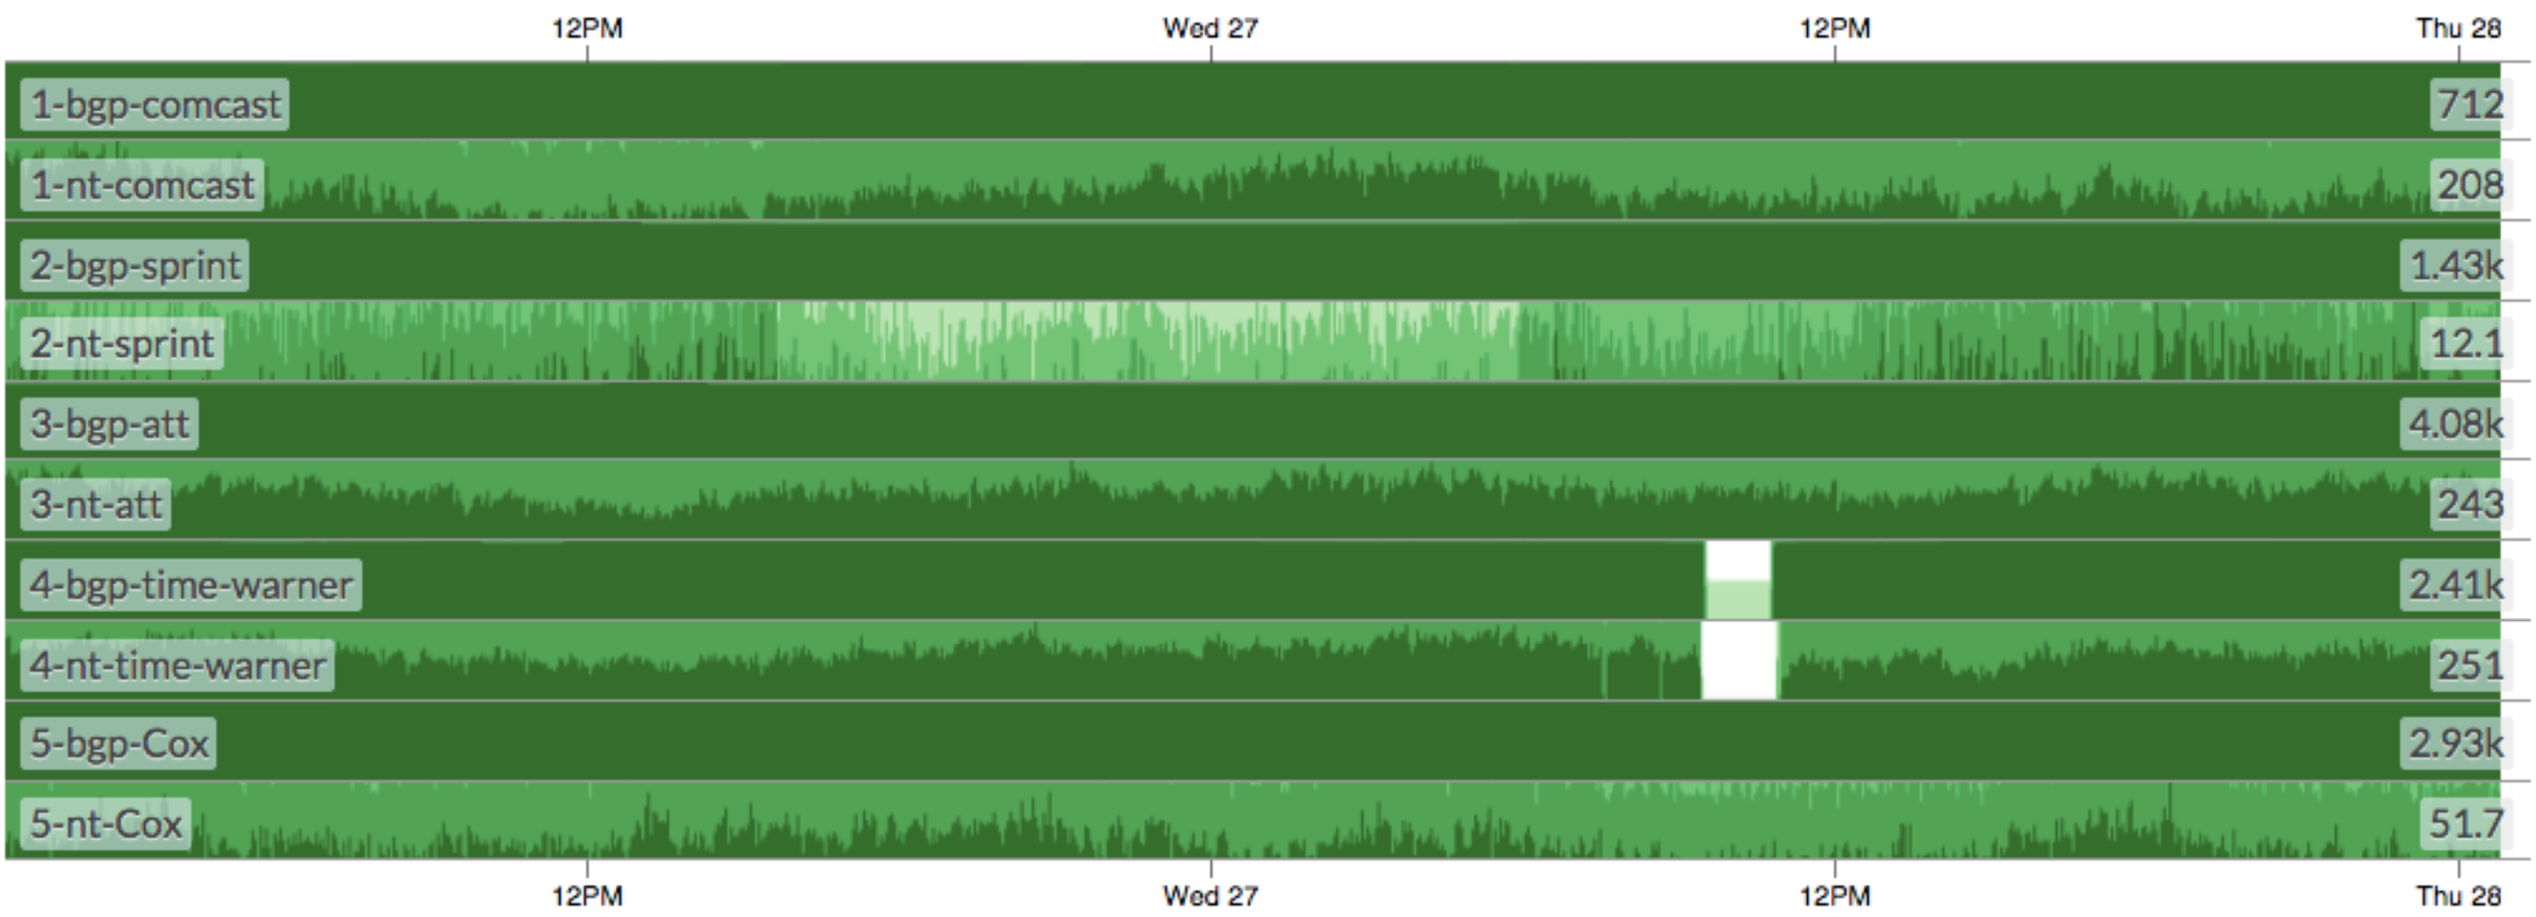

Series: *10* | # Points: *28800* | Data resolution: *minute*

IBR + BGP

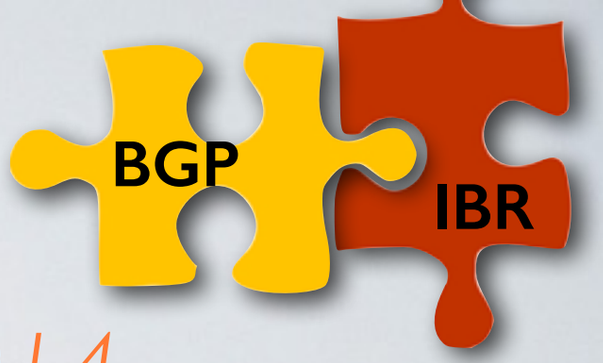

Time Warner Cable - 27th Aug 2014

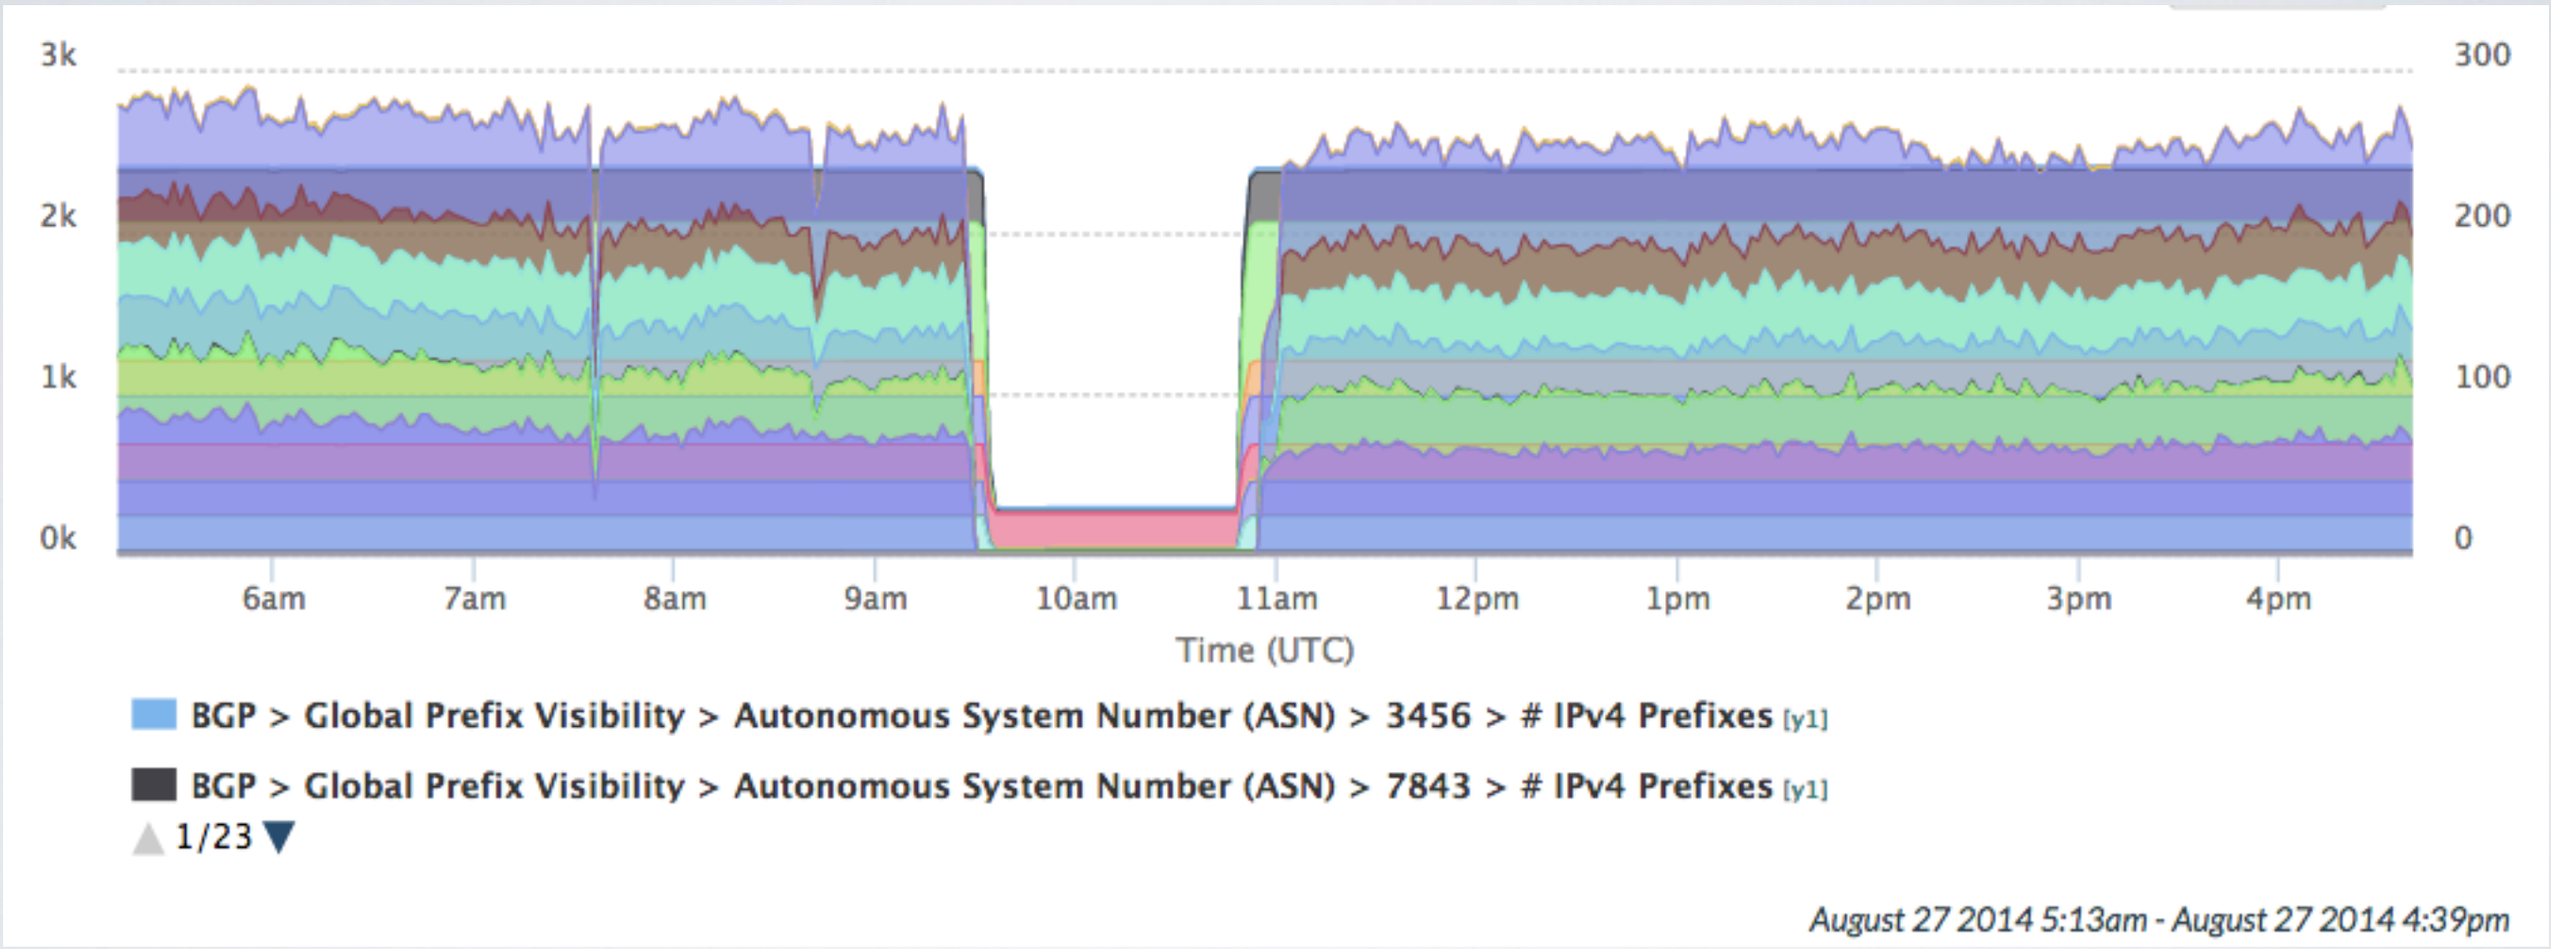

IBR + BGP

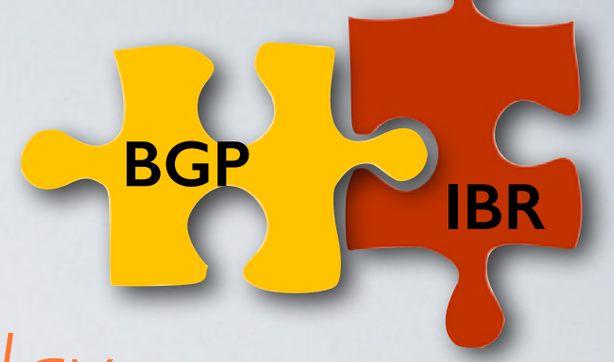

BGP convergence/propagation delay

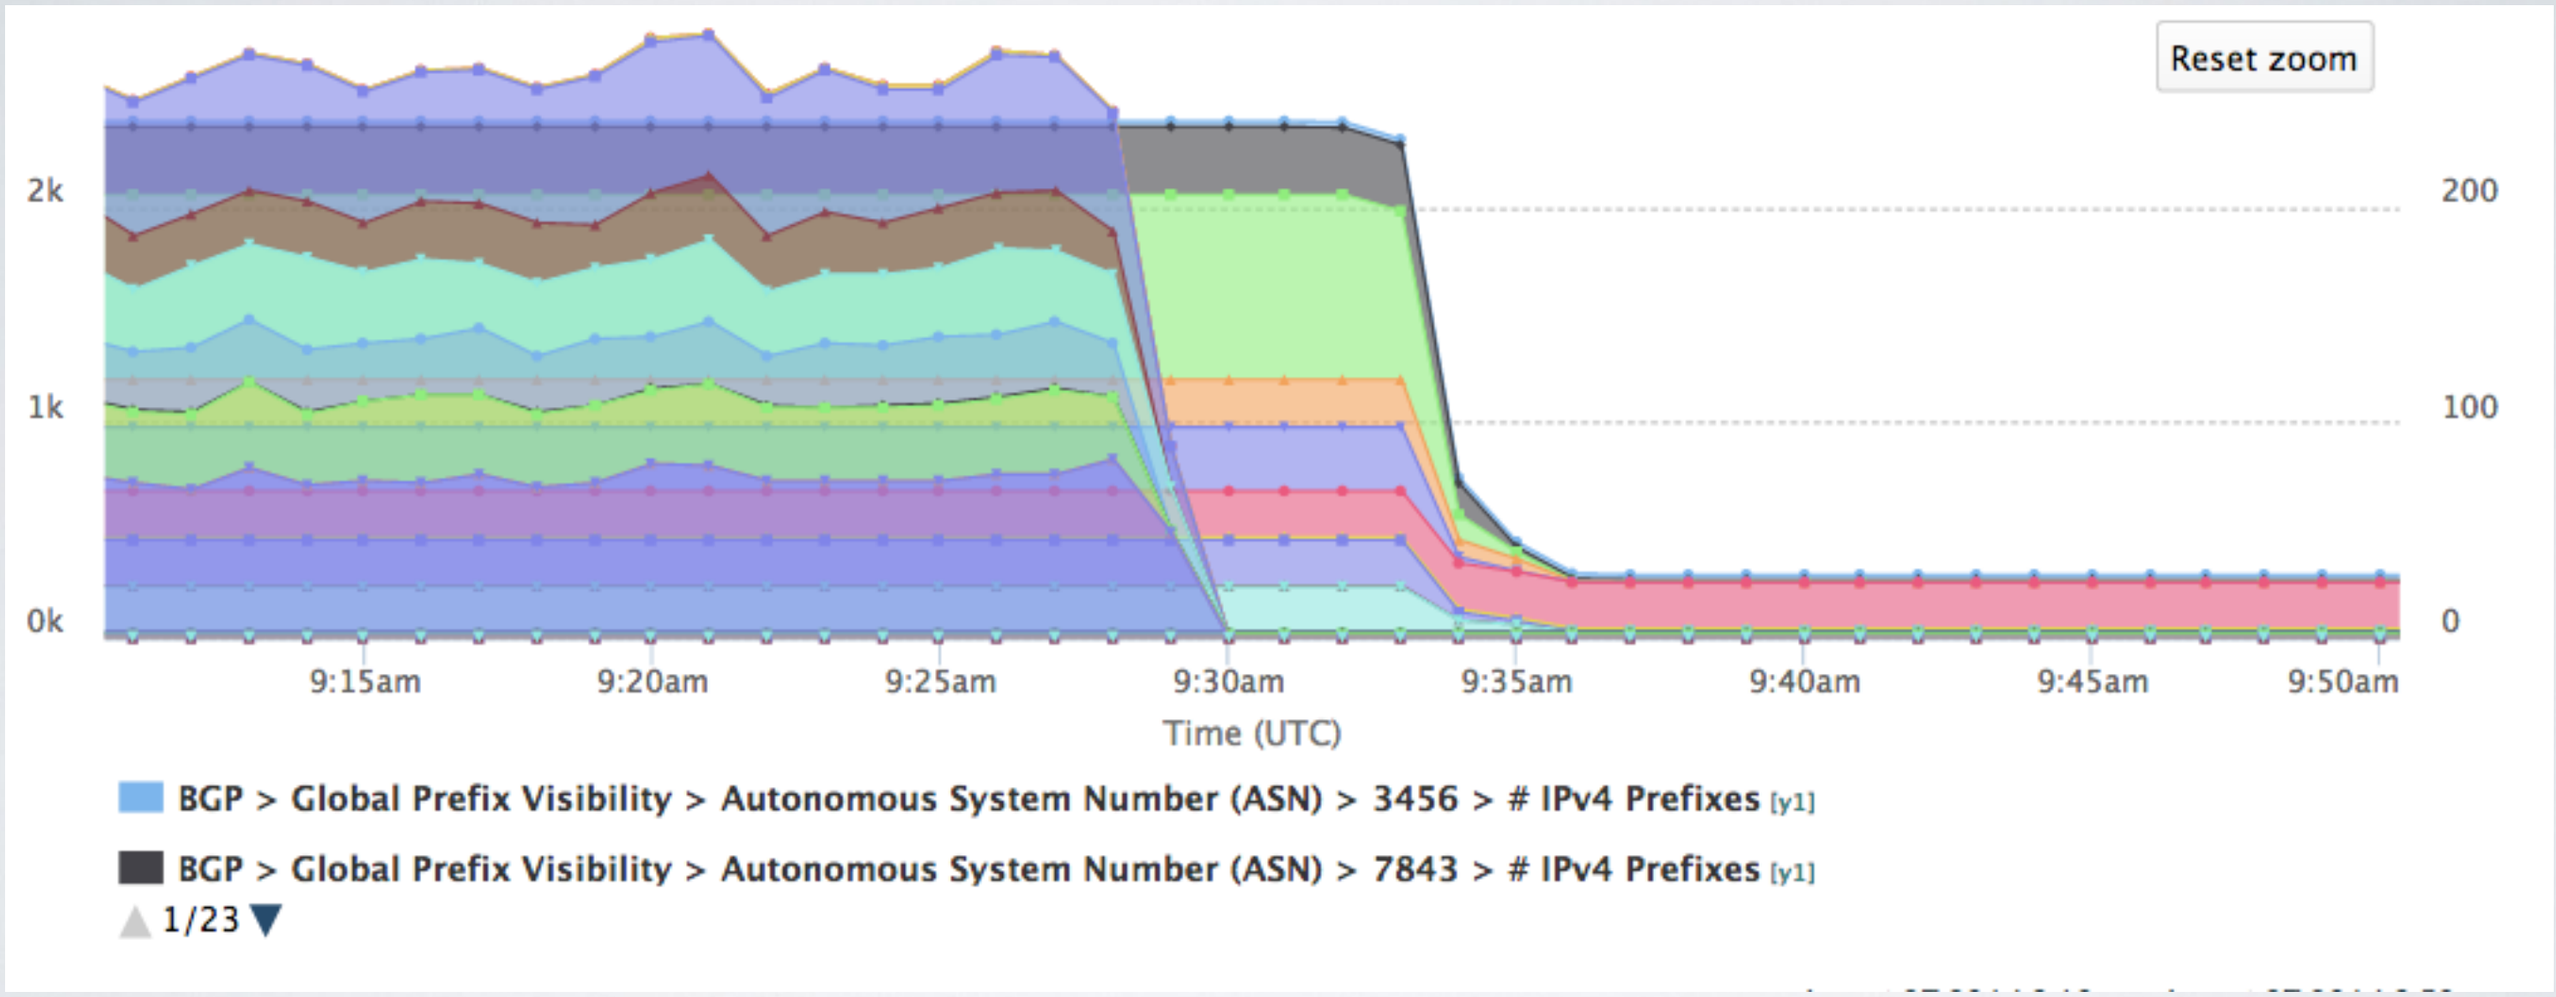

IBR

TWC outage: AS11427 seen by 2 Darknets + BGP

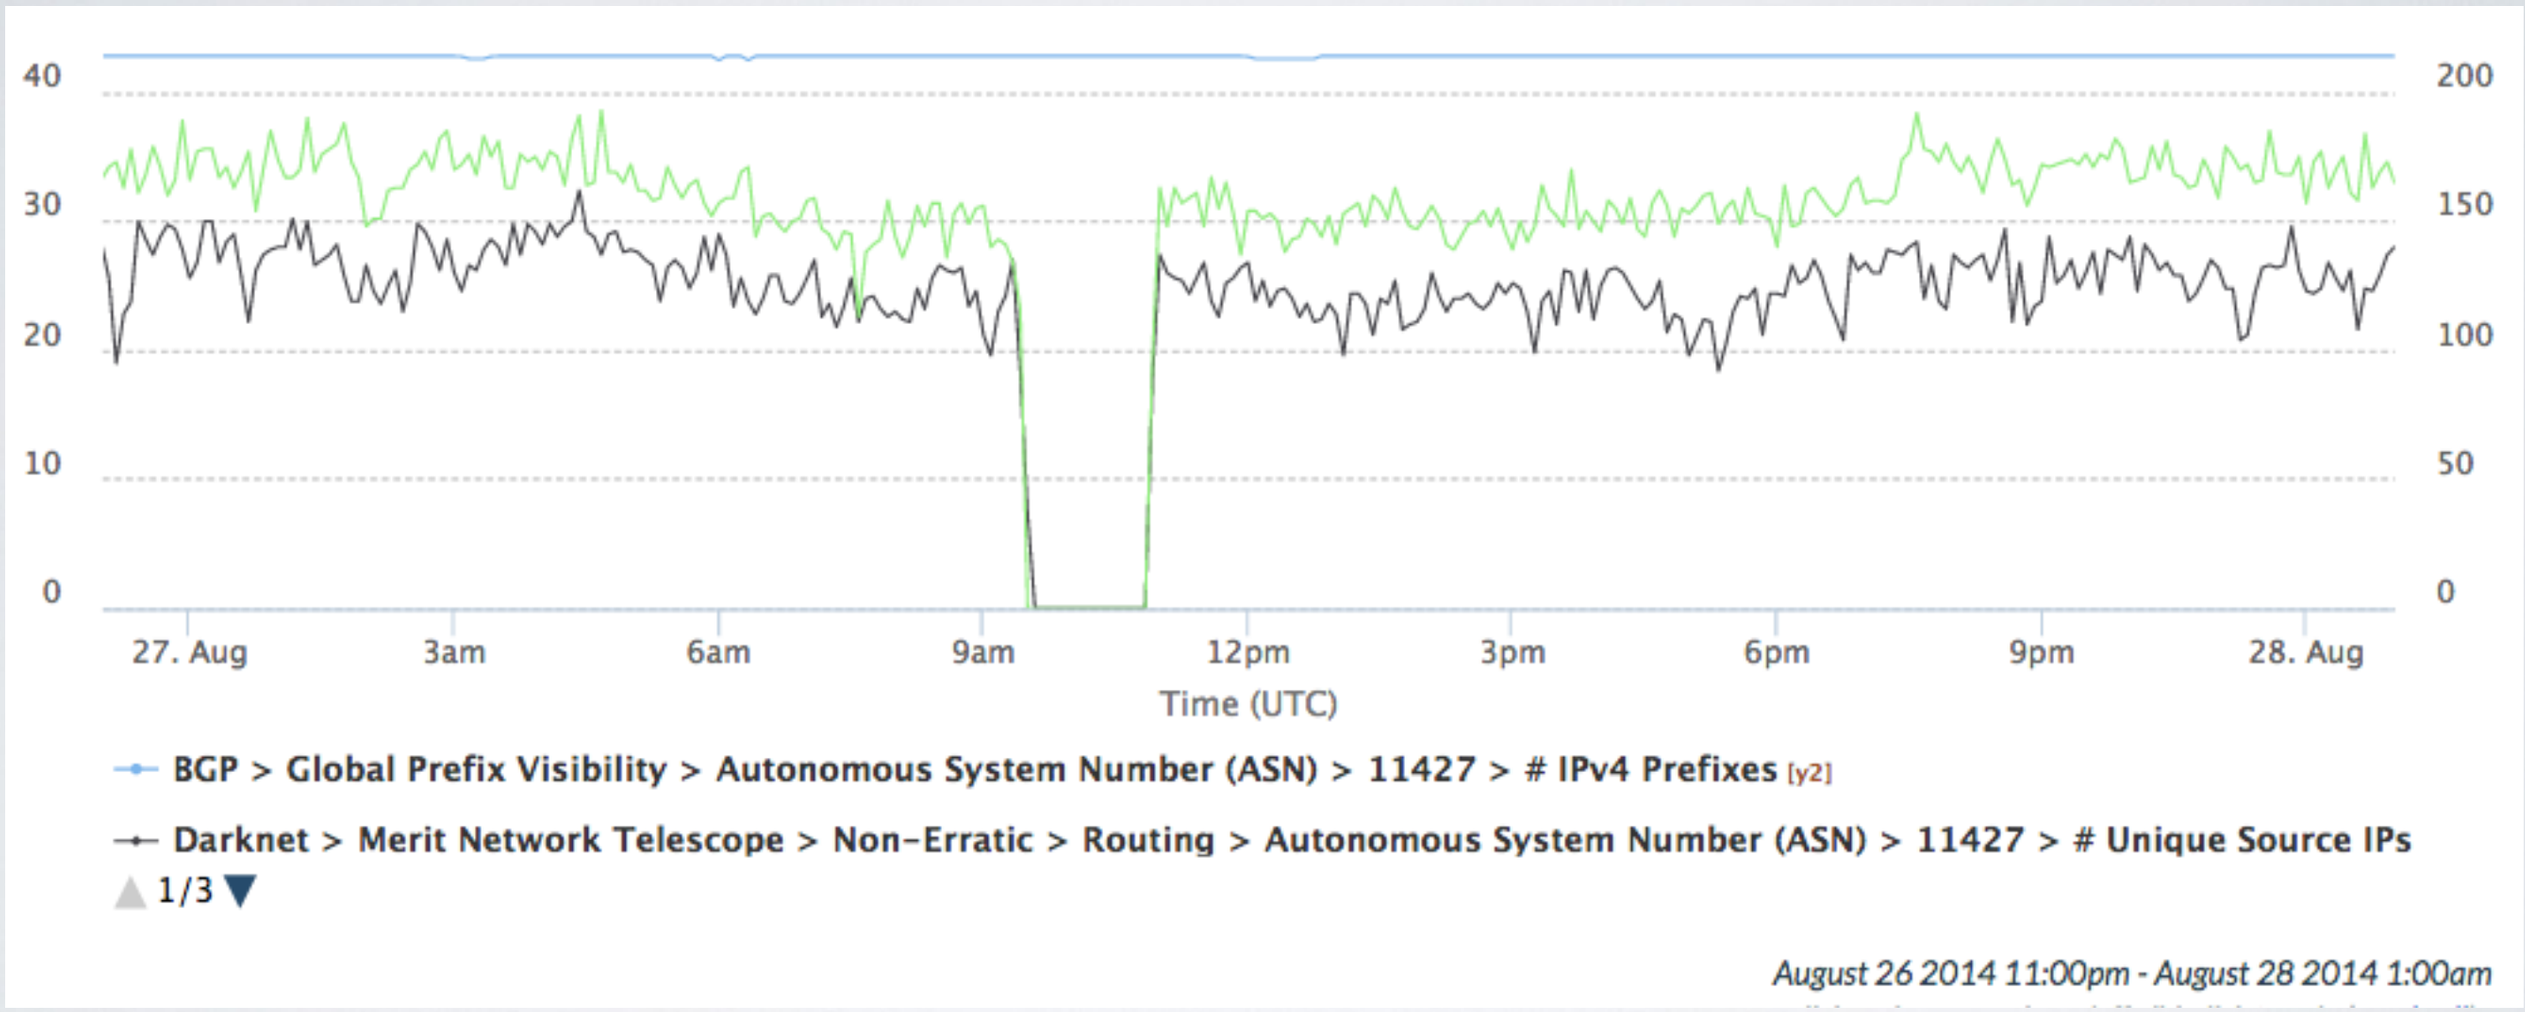

2nd Darknet provided by Merit Networks Inc.
Collaboration with Michalis Kallitsis
http://www-personal.umich.edu/~mgkallit/

IBR

TWC outage: a couple of hours before

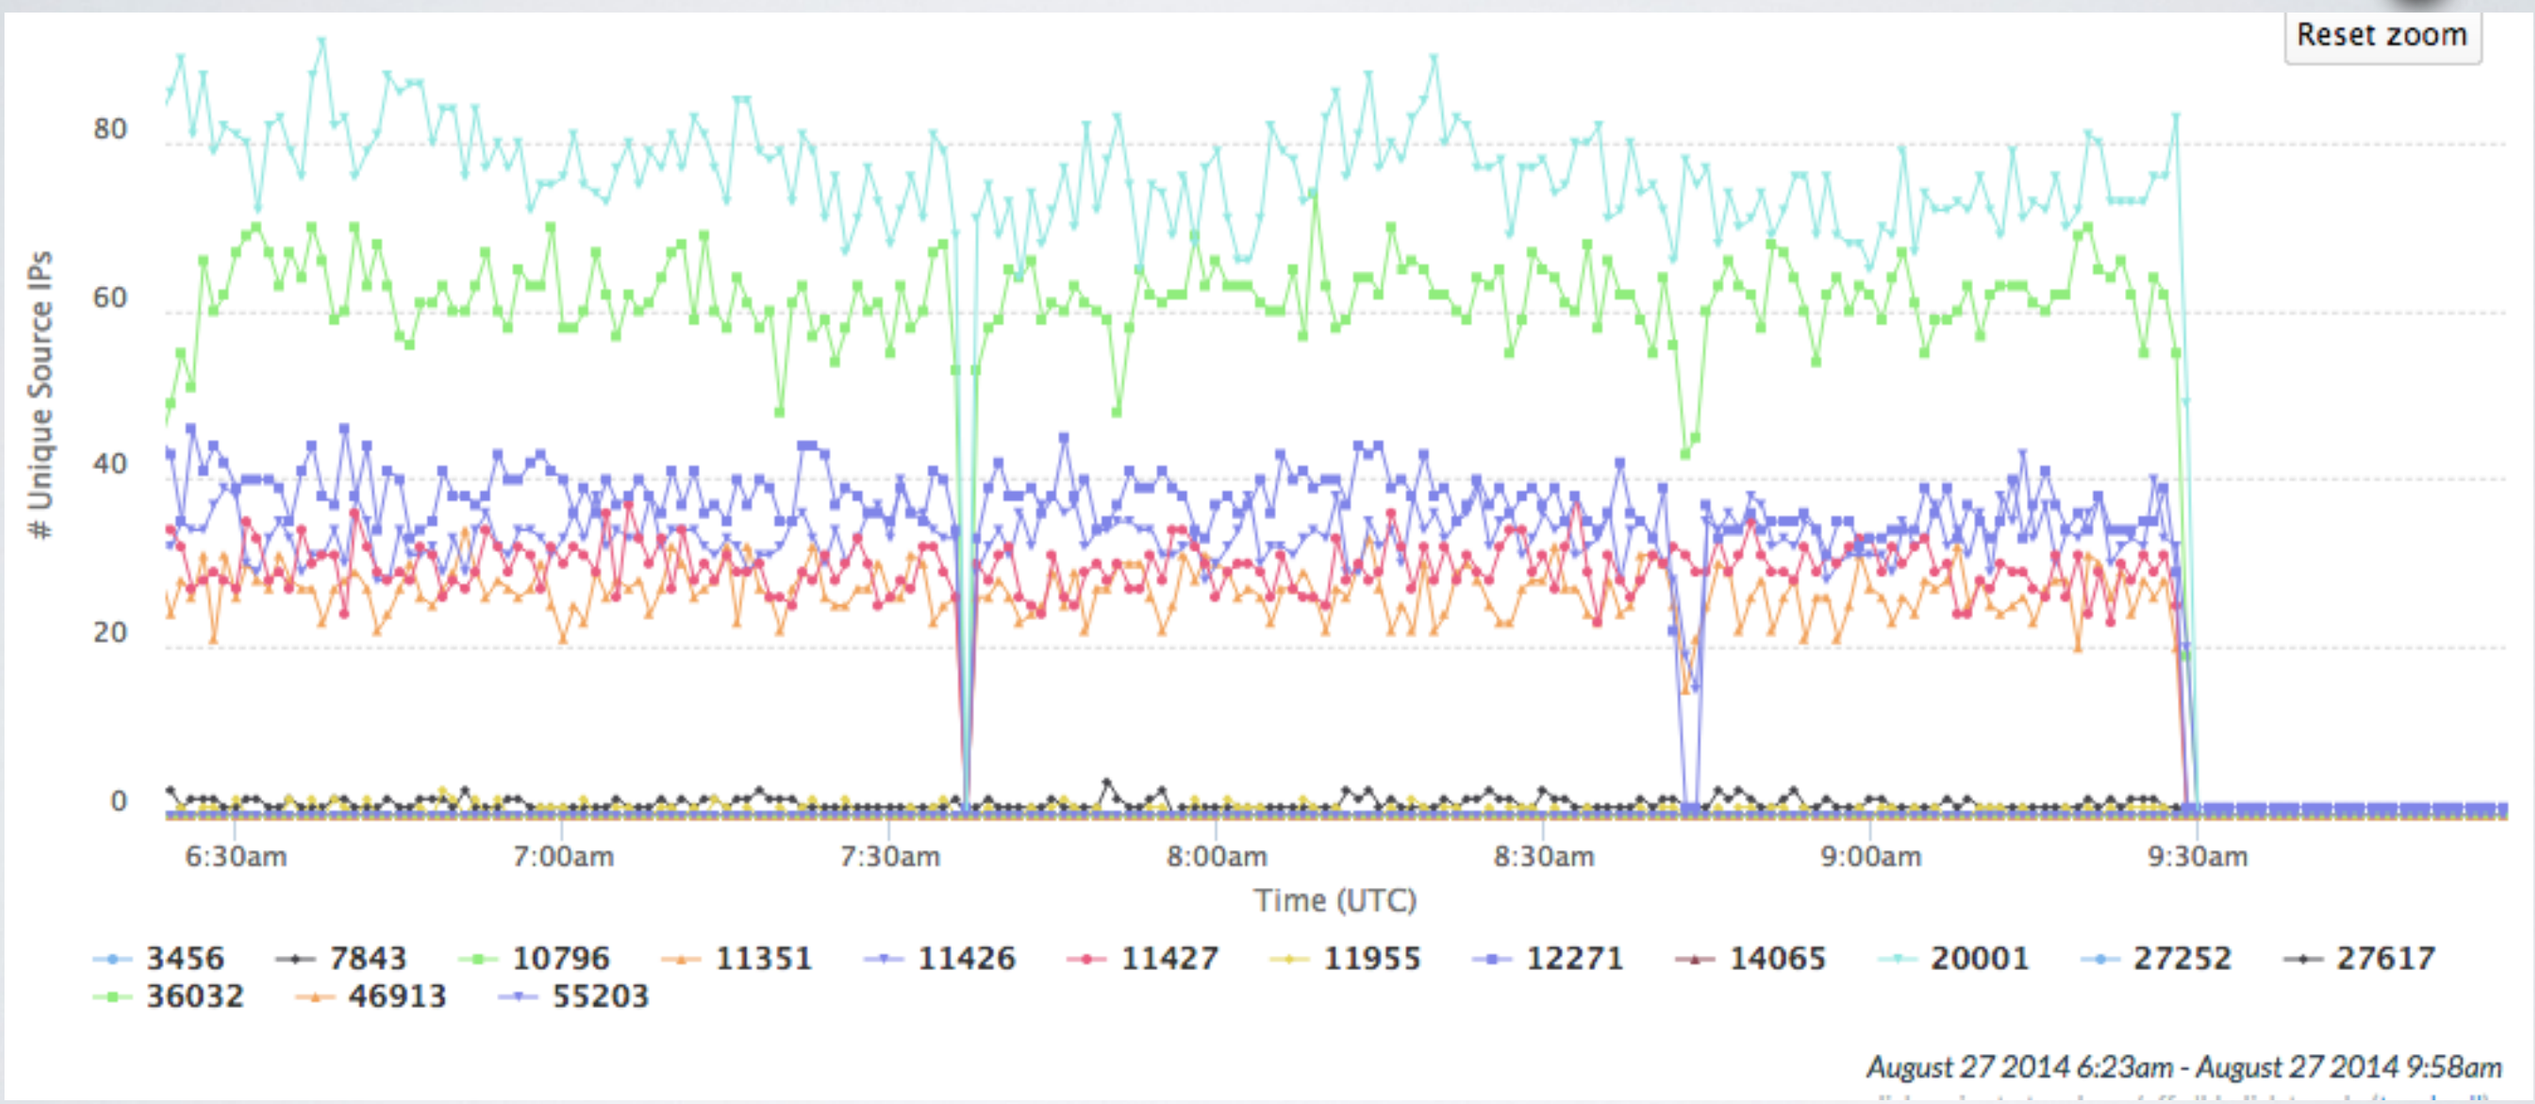

BACKSCATTER

e.g., SYN+ACK replies to spoofed SYNs

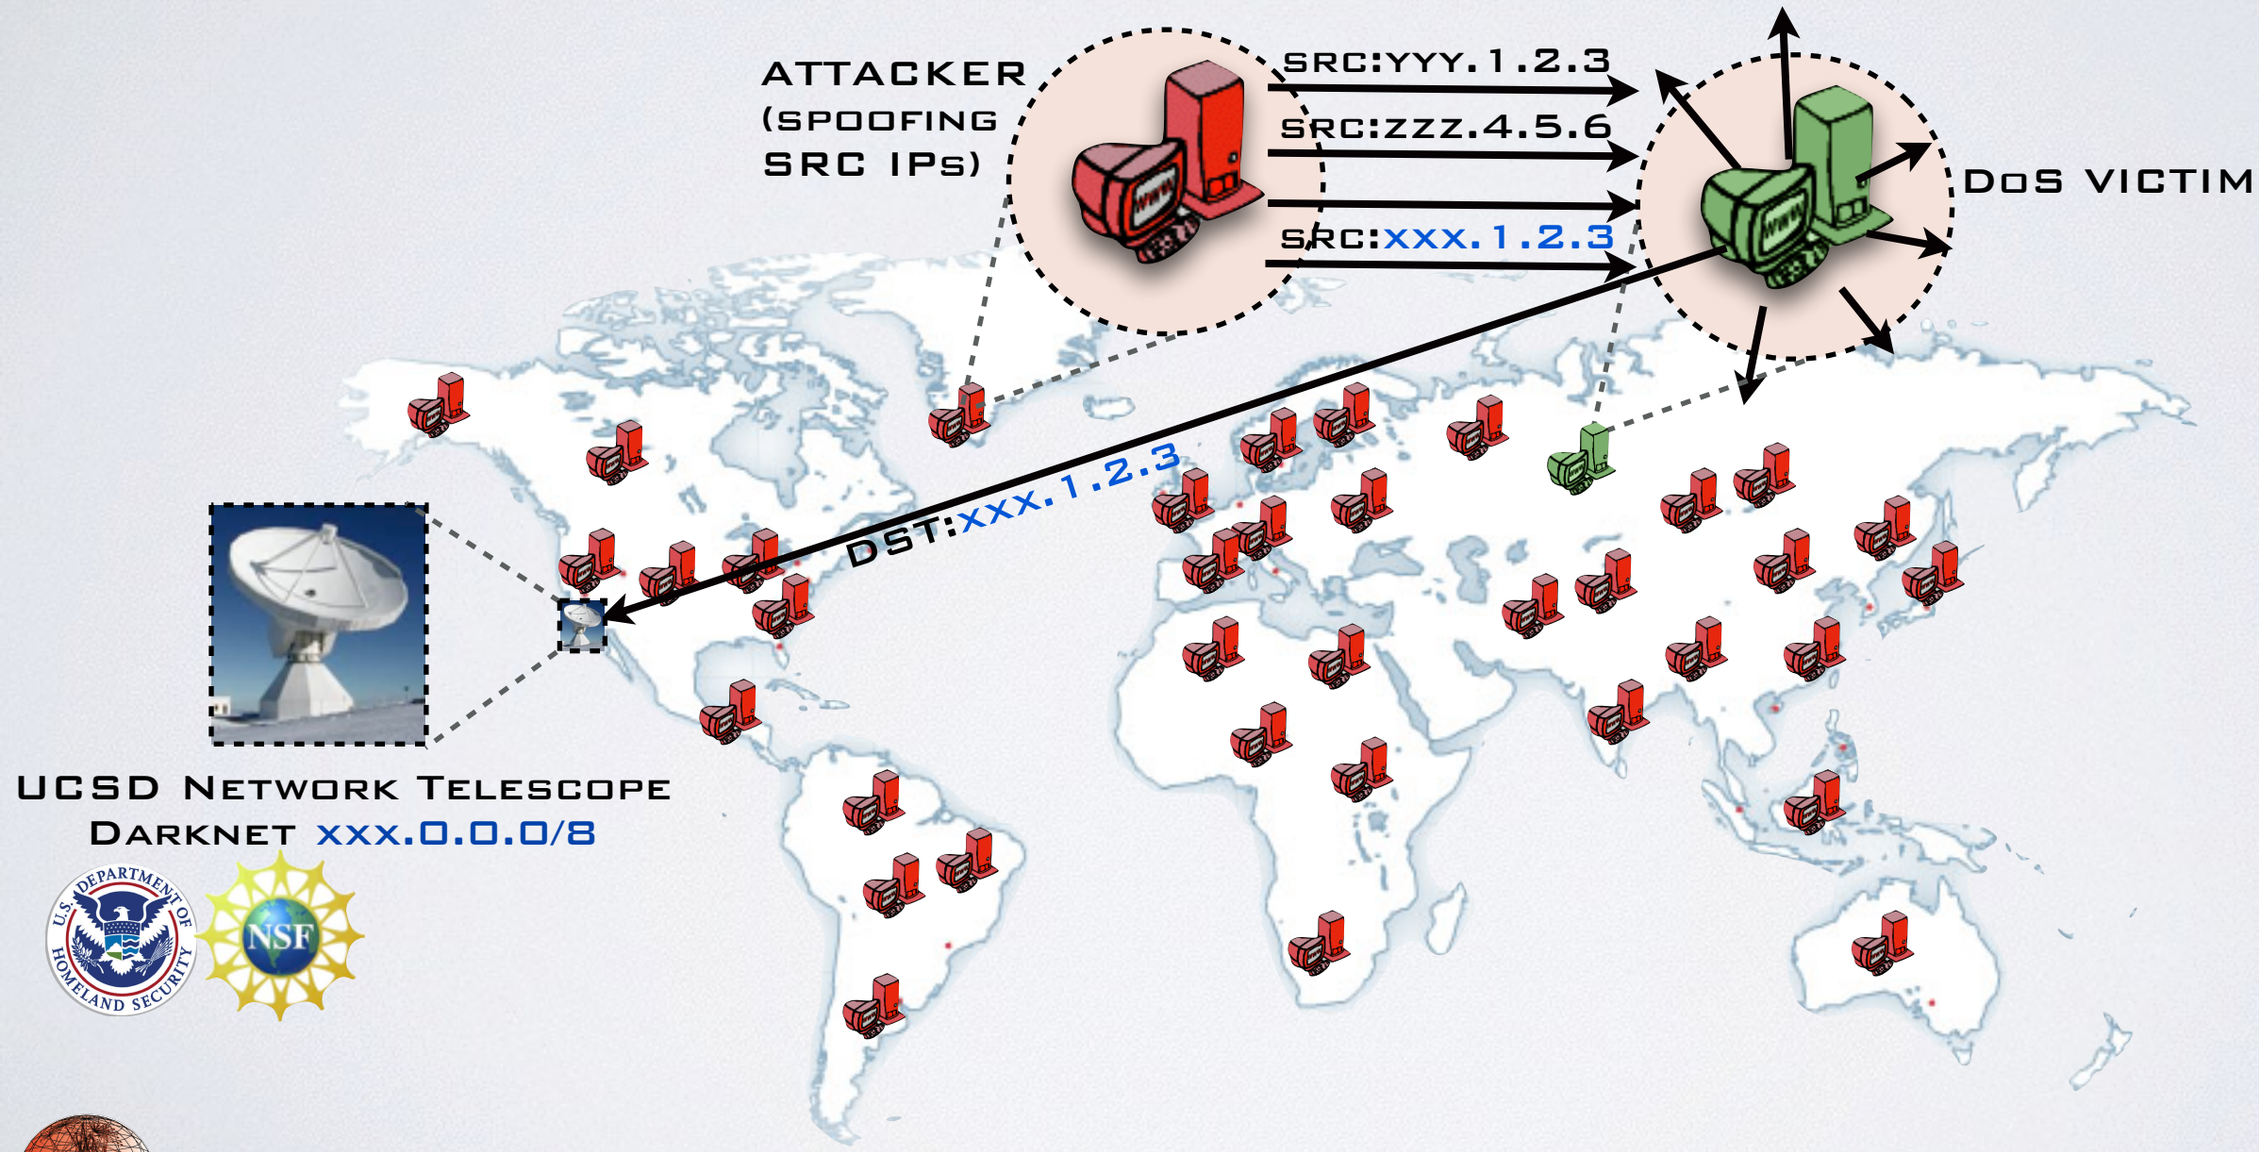

SPOOFED IBR

forged/corrupted packets

- IBR contains also packets with a spoofed source address.

If not properly filtered, they can alter our inferences

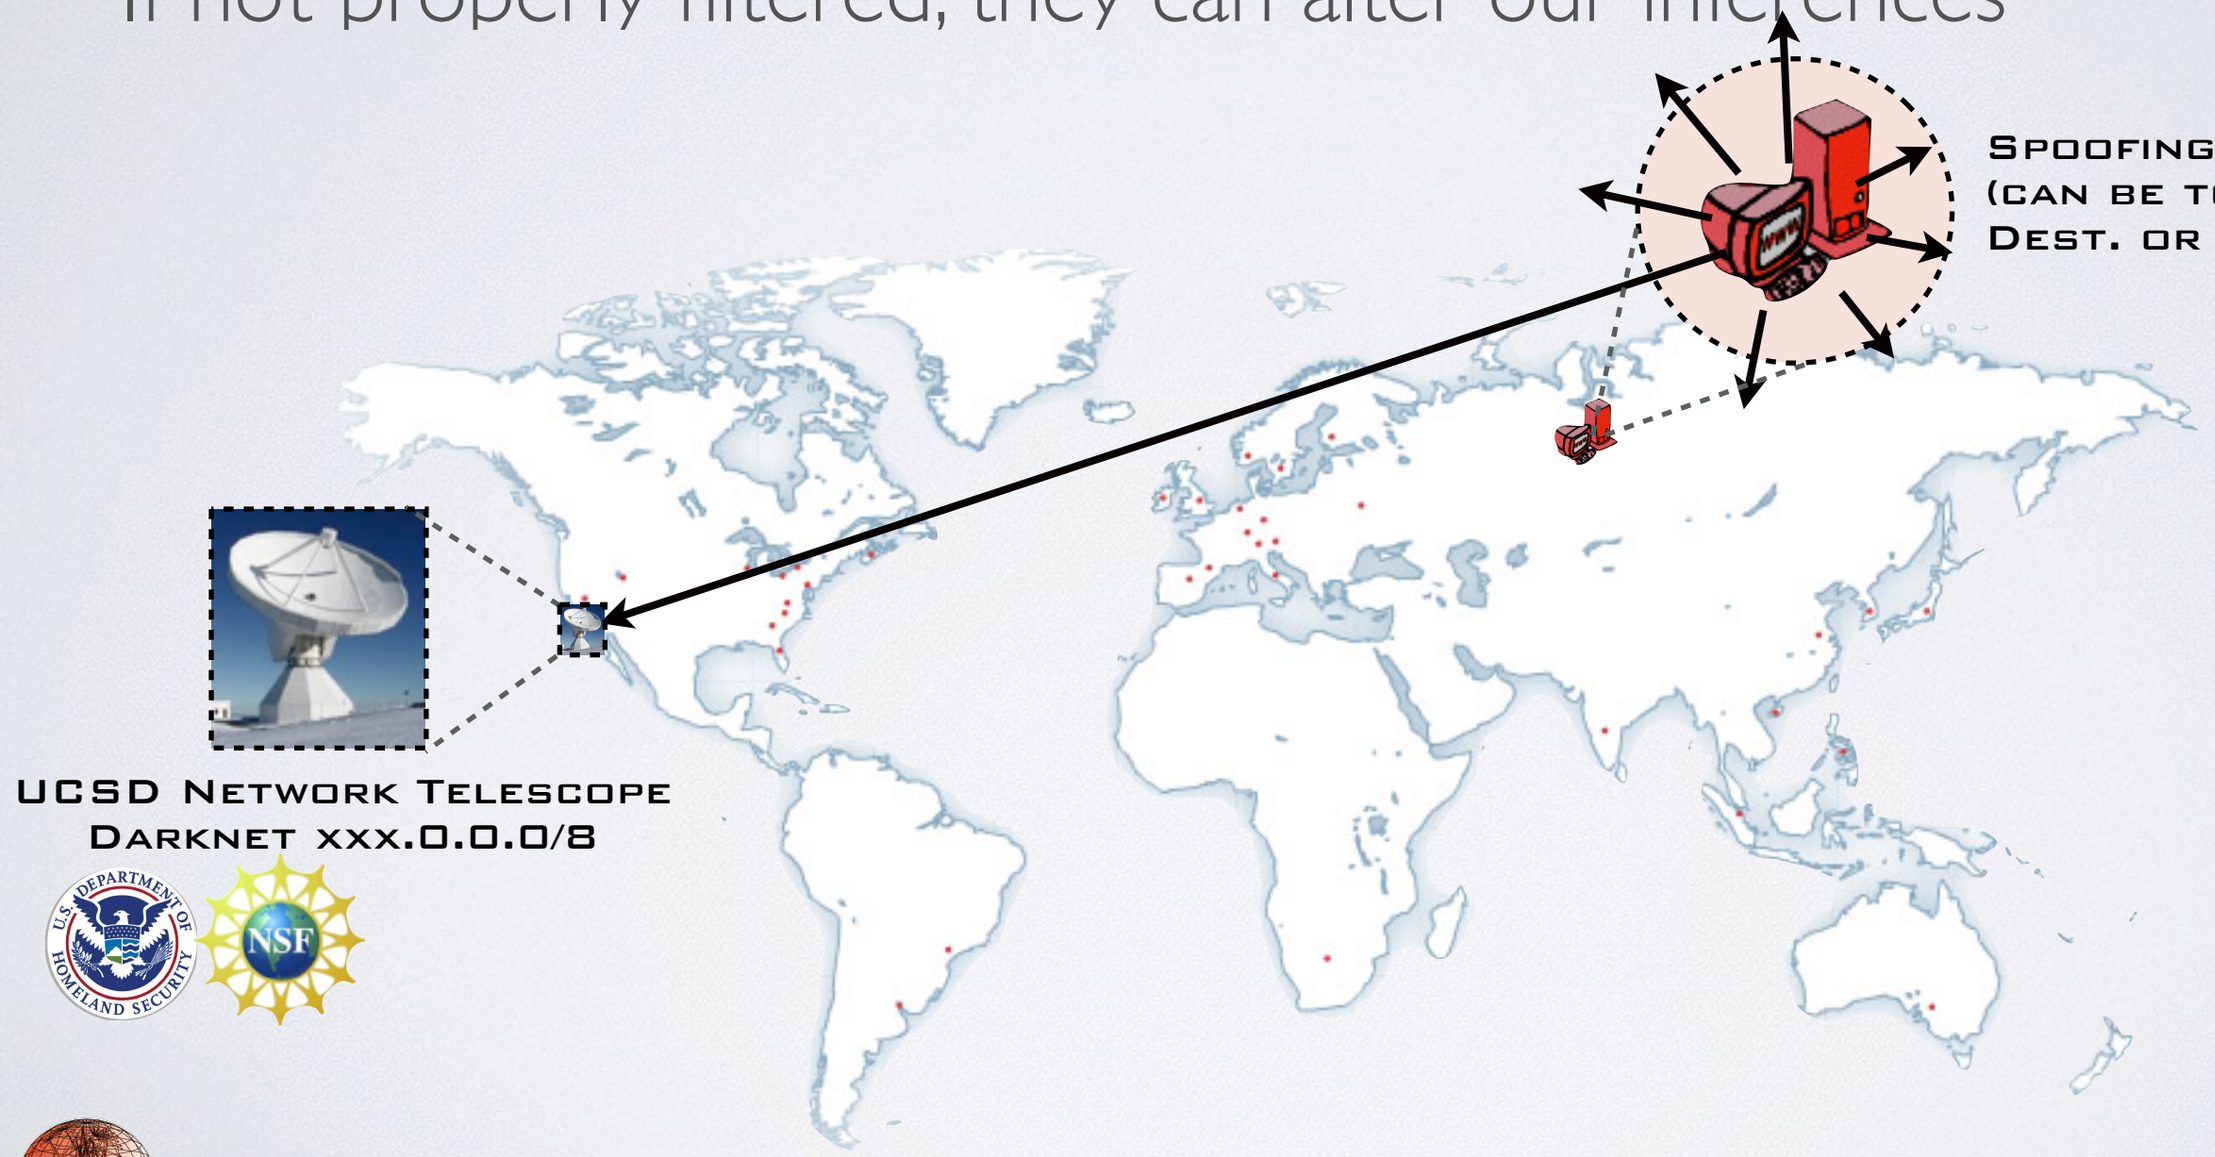

IBR

cleaning the signal

- We developed a methodology to monitor IBR and generate **filters** that automatically remove large-scale and bursty spoofing
- Small-scale spoofing does not constitute a problem
- Also, targeted spoofing attacks are difficult to implement:
  - each network has a "fingerprint" that the attacker would need to know
  - absence of traffic cannot be reproduced
  - we monitor more than one darknet
  - darknet address blocks are not widely known

	Filter	Characterization
General	TTL> 200 and not ICMP	Large-scale/Bursty
	Least signif. byte src addr 0	Large-scale/Bursty
	Least signif. byte src addr 255	Large-scale/Consistent
	Protocol 0	Large-scale/Bursty
	Protocol 150	Large-scale/Consistent
	Same Src. and Dst. Addr.	Small-scale

- We also monitor IBR to identify large-scale coordinated activities that create noise in the signal and exclude such traffic
  - e.g., backscatter from spoofed scanning for open DNS resolvers

Dainotti et al. "Estimating Internet Address Space Usage through Passive Measurements", ACM SIGCOMM CCR 2014

CORSARO

our packet processing framework

- Enables **large-scale near-realtime traffic analysis**

- Modular framework.

A pipeline of plugins:

- filter traffic *(e.g., spoofed)*
- classify it *(e.g., backscatter)*
- tag packets *(e.g., by geolocation)*
- aggregate tagged packets and extract time series data points at each interval *(e.g., 1 min)*
- writes data in our high-performance time-series DB

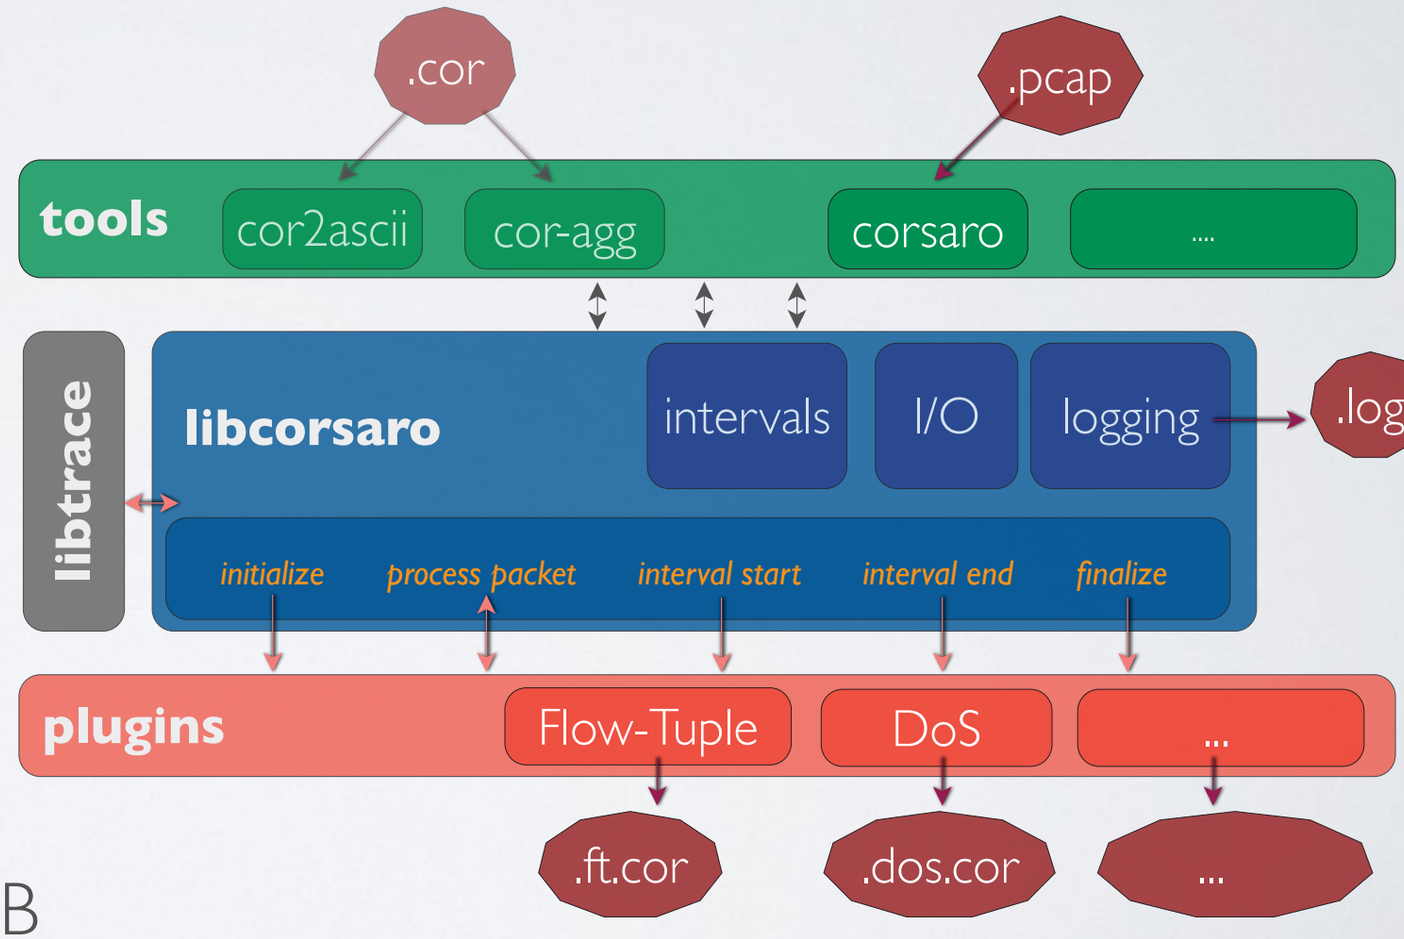

ACTIVE PROBING

background info

- Collect measurements by injecting packets into the network
  - Ping
    - Sends "ECHO requests" packets to a destination host and receives back "ECHO replies"
      - ***if:*** the host is reachable, is up, is configured to reply to echo requests, etc...
      - measures reachability and Round Trip Time
  - Traceroute
    - Similar concept but tries to "*ping*" each hop on the path to the destination host through a careful use of the TTL IP header field
      - ***if:*** the hop is reachable, is up, is configured to reply, etc...
      - measures reachability, Round Trip Time for all (replying) hops, enables inference of IP-level path and AS-level path

ACTIVE PROBING

collaboration with ISI/USC

- John Heidemann's methodology: "*Trinocular*"
  - probing based on *pings* (ICMP echo requests)
  - includes an outage inference methodology based on Bayesian principles
  - /24 IPv4 blocks granularity
  - currently 3 vantage points

- Inferences and raw data shared through the DHS PREDICT project
  - we started working on importing historical data into Charthouse for analysis

- Planning to collaborate to integrate realtime feed with our system

Quan et al. "Trinocular: Understanding Internet Reliability Through Adaptive Probing", ACM SIGCOMM 2013

ARK

Archipelago

- CAIDA active measurement infrastructure
  - supports ongoing 24/7 Internet-wide topology measurement as well as customized experiments
  - IPv4, IPv6, TCP, UDP, ICMP
  - active since *2007*
- *107* vantage points (and growing)

- Planning to implement a variation of ISI's *Trinocular* exploiting the availability of more vantage points

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

PINGING IN THE RAIN

probing from PlanetLab

- Collaboration with Aaron Schulman (Stanford) and Neil Spring (Univ. Maryland)
  - *Schulman, Spring, "Pingin' in the Rain", ACM Internet Measurement Conference 2011*
- Originally focused on how weather affects residential Internet connections in the US
- Probing from PlanetLab
  - *1342 nodes at 666 sites.*
- We observe drops in % of hosts replying to pings

http://www.cs.umd.edu/~schulman/thunderping.html

https://www.planet-lab.org

PLANETLAB

PINGING IN THE RAIN

ACTIVE PROBING

probing from PlanetLab

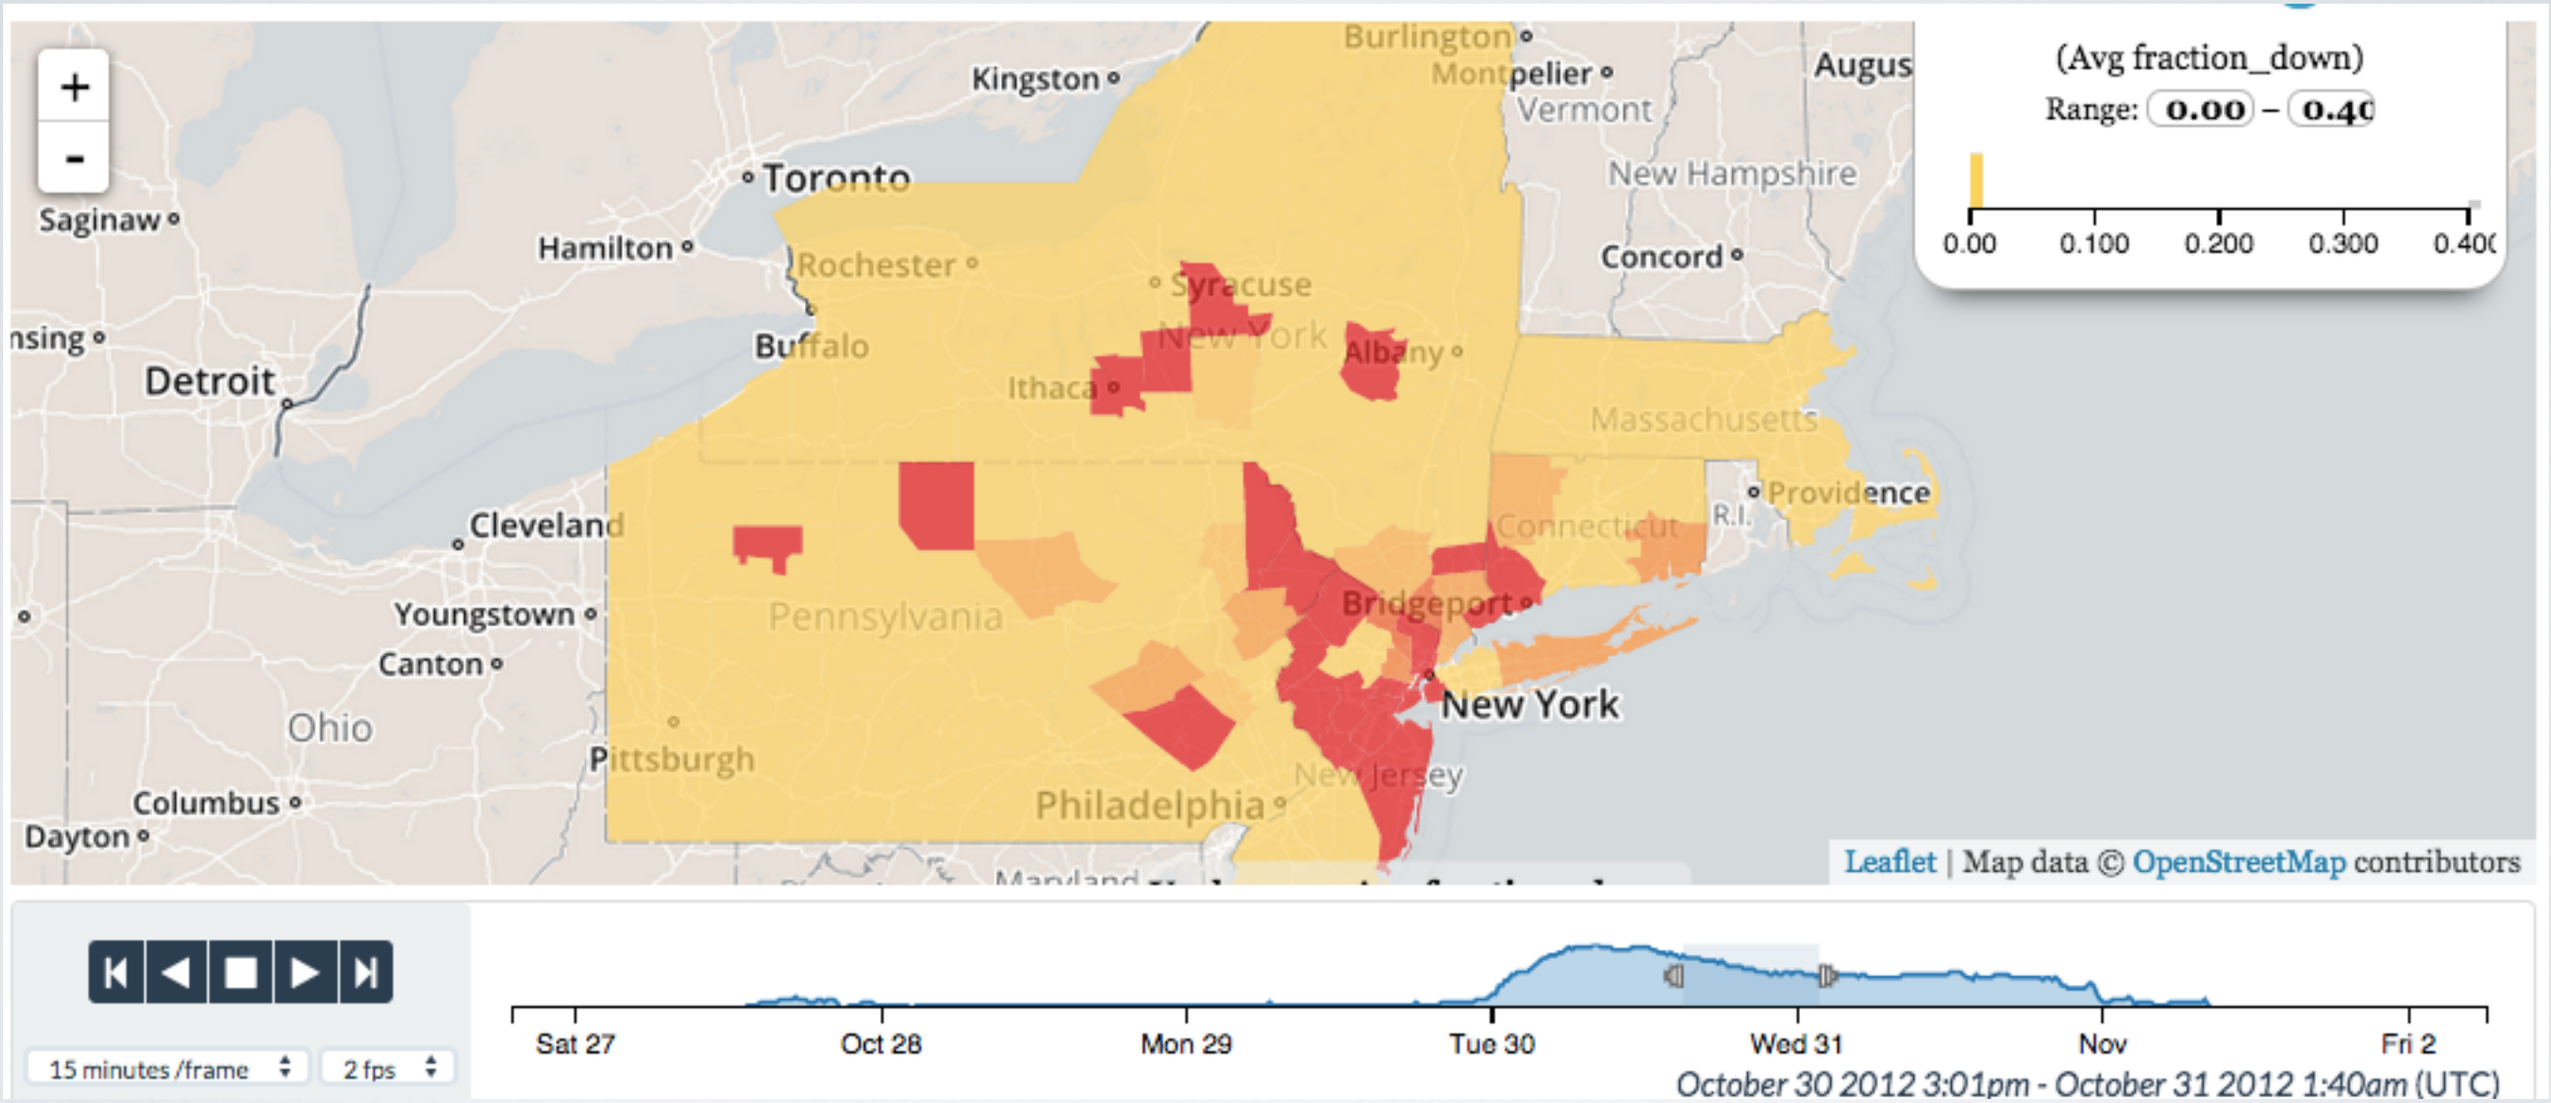

ON DEMAND PROBING

Ark measurements on demand

- Ark offers *(to authorized users only)* an API for on demand probing
- helps with validation, allows finer time granularity, multiple vantage points, diagnostics with traceroutes, not only ICMP (e.g., TCP, UDP), …

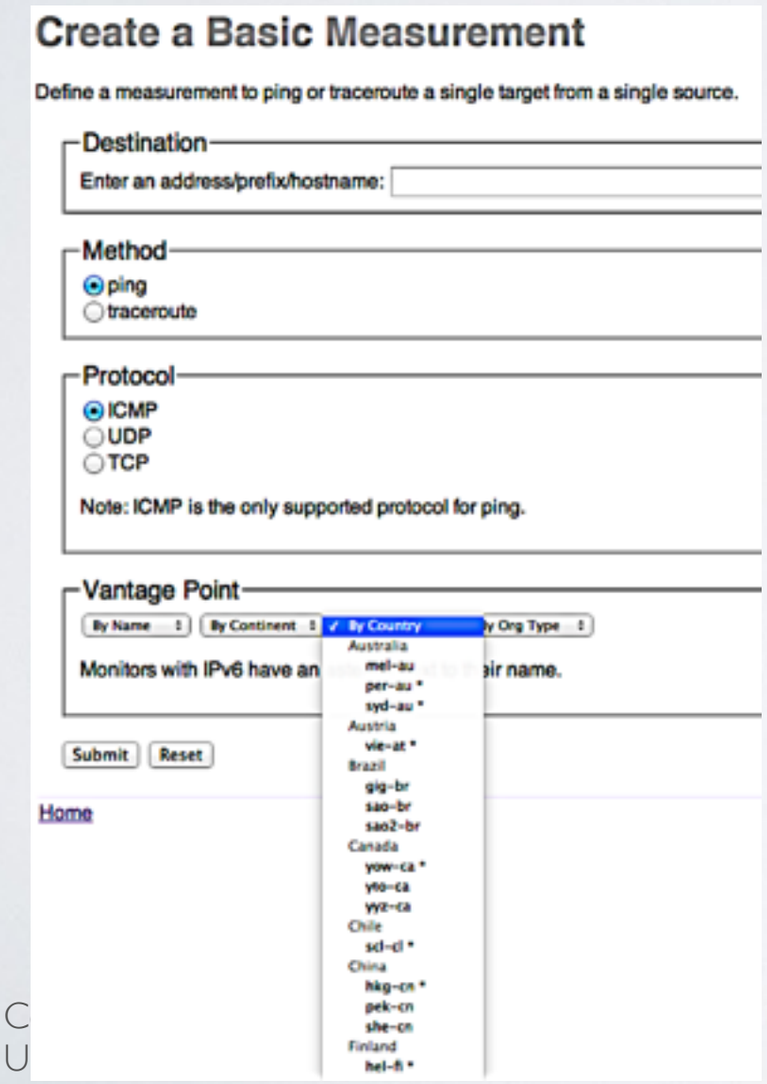

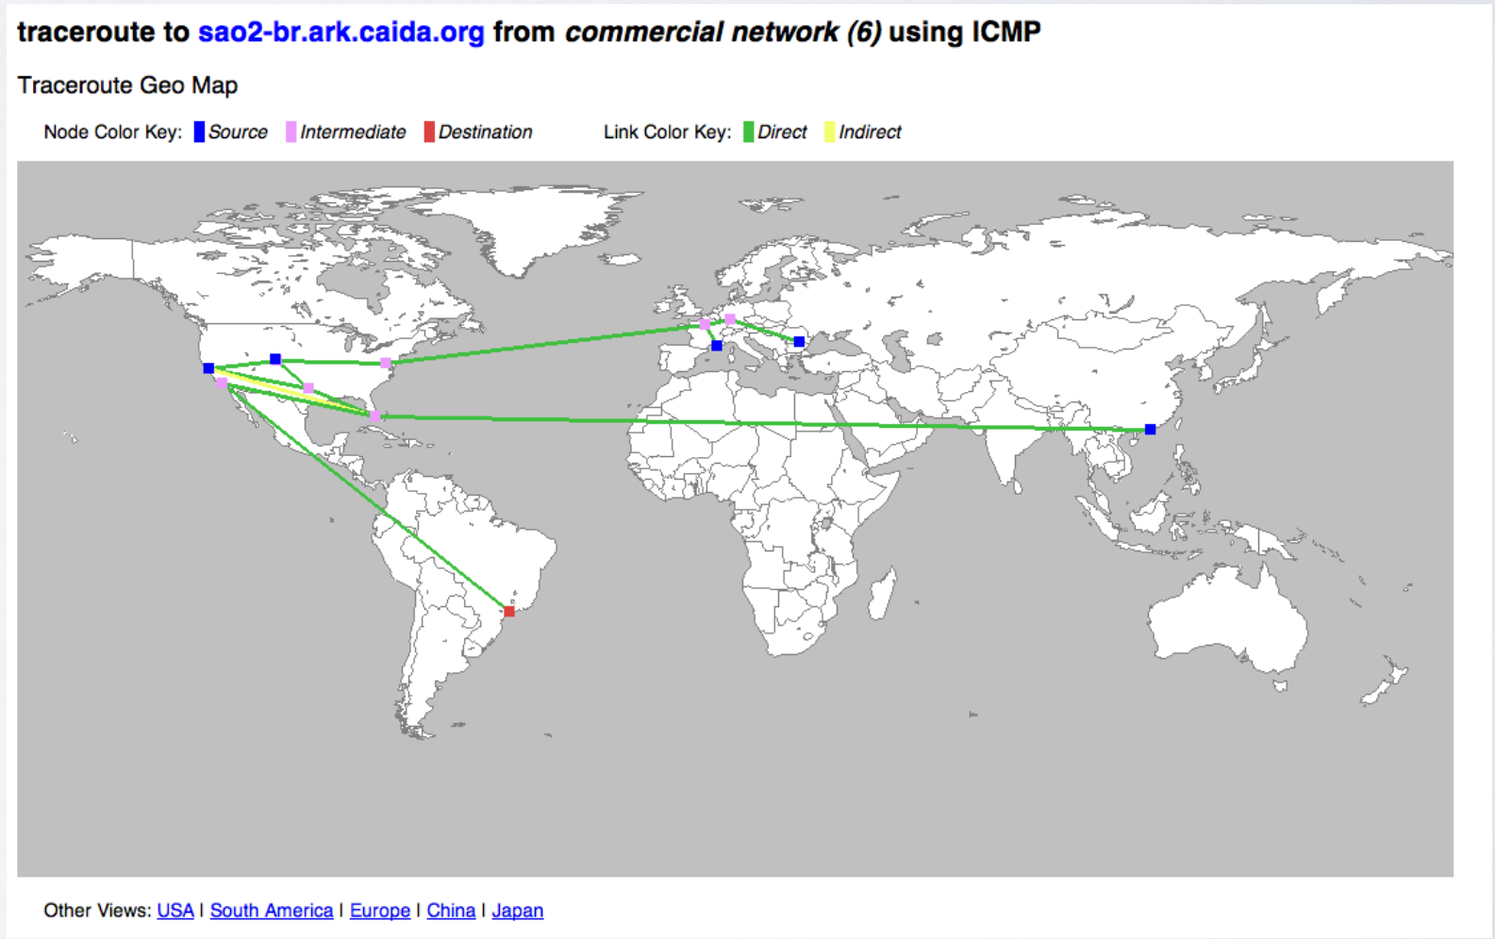

TEAM

credits

Vasco Asturiano
Karyn Benson
KC Claffy
Alberto Dainotti
Marina Fomenkov
Young Hyun
Bradley Huffaker
Ken Keys
Alistair King
Ryan Koga
Alex Ma
Chiara Orsini
Josh Polterock
Cindy Wong

THANKS

questions?

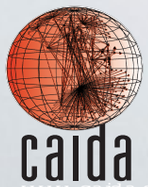

Center for Applied Internet Data Analysis
University of California San Diego