CAIDA Ark

Archipelago Measurement Infrastructure

UC San Diego

JJ Jamison

jj@caida.org

March, 2025

About JJ Jamison - jj@caida.org

35 years working on/supporting US DoD and NSF funded Research projects at:

- A US Government funded Research & Development Center, an early ISP, the University of Illinois, and companies like Cisco Systems and Juniper Networks

Currently working for:

- CAIDA/UCSD from EAFIT University in Medellín, Colombia

Focused on:

- Network Measurement Research (CAIDA) and distributed computing (EAFIT/NRP)

UC San Diego

caida

APOLO
Centro de computación científica
Universidad EAFIT

NRP
NATIONAL RESEARCH PLATFORM

About CAIDA & Ark

The Center for Applied Internet Data Analysis (CAIDA):

- was founded in 1997
- conducts network research and builds research infrastructure
- supports large-scale data collection, curation, and data distribution
- based at the San Diego Supercomputer Center, located on the campus of the University of California San Diego

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

- provide insights 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
- inform science, technology, and communications public policies

Archipelago Project (Ark)

CAIDA's active measurement infrastructure serving the network research community **since 2007**.

Ark monitors provide raw data for most of CAIDA’s macroscopic Internet data sets.

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

Ark Monitor Deployment

https://www.caida.org/projects/ark/locations/

287 nodes in 187 autonomous systems across 213 cities in 66 countries

The benefits of hosting an ARK node

- **Ark monitor statistics provide for “a view” of the internet from the hosting institution’s perspective.**
- **Ark Nodes can be leveraged to manage and optimize host networks**
  - Develop custom applications using NiteOwl Python based primitives
  - Make use of community developed network management & optimization applications available in the CAIDA github repository:
    - https://github.com/CAIDA/ark-community
- **Ark Hosts Support Real Internet Research**
  - Broaden the view of the global Internet for the network research community
  - Enable researchers, including researchers at your own institution, to conceive, develop, and test their models and methods.

Archipelago (Ark) Monitor Statistics

Statistical information for the topology traces taken by each individual Ark monitor provide for “a view” of the internet from the hosting institution’s perspective.

Ark monitors collect data by sending scamper probes continuously to destination IP addresses. **Destinations are selected randomly from each routed IPv4 /24 prefix on the Internet** such that a random address in each prefix is probed approximately every 48 hours (one probing cycle). A single monitor won't probe all prefixes, but the prefixes it does probe will be randomly distributed, which gives a good sample cross section of the address space. As each probe travels from the monitor to its final destination, it passes through several IP addresses (ie, routers) which are owned by different autonomous systems (ASes).

Data in the following slides show statistical information for the topology traces taken by the Ark monitor at EAFIT.

https://www.caida.org/projects/ark/statistics/monitor/med-co.html

med-co

EAFIT
Medellin, CO (2)

Time range	2025-02-05 09:24 to 14:26 UTC (5 hours) (1 day ago)
Total traces	12000
Traces with responding destinations	1310 (10.917%)
ASes with responding destinations	414 (16.768% out of 2469)
Prefixes with responding destinations	1068 (12.800% out of 8344)
Protocol version	v4

ISPs, border routers, or firewalls might be dropping/blocking ICMP packets. Or it could be that the destination IP address does exist.

Note that this does not impact our goal of mapping the core of the internet.

Archipelago (Ark) Monitor Statistics include:

Median RTT per country and US state.

Median round-trip times (RTTs) from each Ark monitor to various countries (and US states).

Path Dispersion

The paths taken by different traces show where the majority of a monitor's traffic travels.

Path Length Distributions

Path length distributions show, on aggregate, how well-connected a monitor is to the rest of the Internet.

RTT Distribution

An RTT distribution shows the general latency between a monitor and the rest of the Internet.

RTTs vs. Distance

Comparing RTT with geographical distance helps show if latency is primarily due to distance or other factors.

Median RTT per country and US state

- These graphs show the median round-trip times (RTTs) from the probing monitor to various countries (and US states). The country/state is obtained via NetAcuity for each hop's IP address. All RTTs to that country/state are collected and the median is calculated and assigned a color (with a maximum displayed value of 1000ms). The location of the monitor is indicated by a black box on the maps.
- By looking at the country/state level differences in RTT, we can see how political boundaries can contribute to changes in packet transmission speed.
- The round trip time for a (IP level) hop is the time (in milliseconds) that it takes for a packet to be sent from an Ark monitor to that hop and for that hop's response to be received by the monitor. Non-responding hops are ignored, and hops in a routing loop are removed.
- In general, one will expect to have the lowest median RTTs in the country that hosts the Ark monitor, with RTTs increasing slowly with geographic distance. However, some countries will have significantly higher median RTTs, which indicates that the slower speed is due to infrastructural issues. If these only occur transiently, they point towards a misconfiguration or temporary failure instead of a systemically slower network.

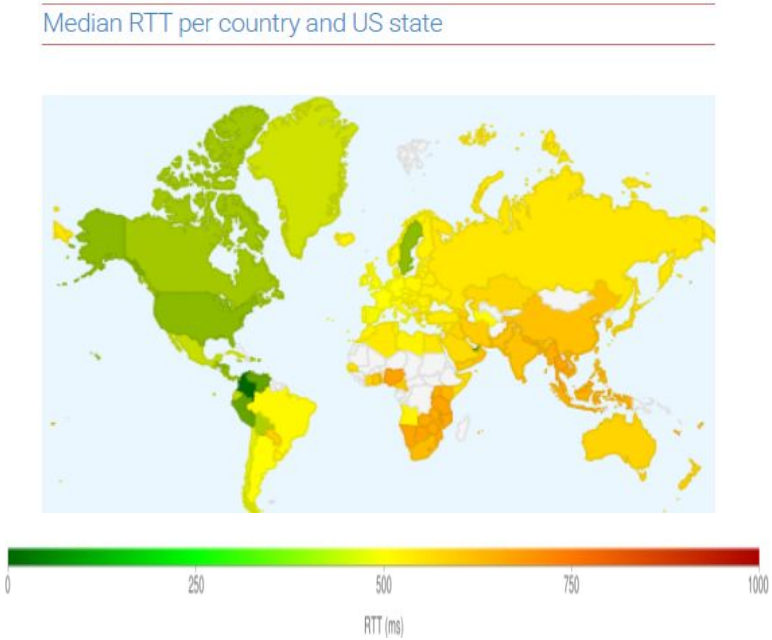

Path Dispersion

- The paths taken by different traces show where the majority of **a monitor's path measurement traffic** travels.
- **It is important to recognize that these graphs are meant to illuminate the routing from a monitor, and not to show the volume of traffic normally flowing on the links or their bandwidth.** Because of this, an AS/IP that might only be used for a small amount of actual traffic (but routes to a large section of the address space) can seem disproportionately large on the graph.

AS Path Dispersion (by AS Hop)

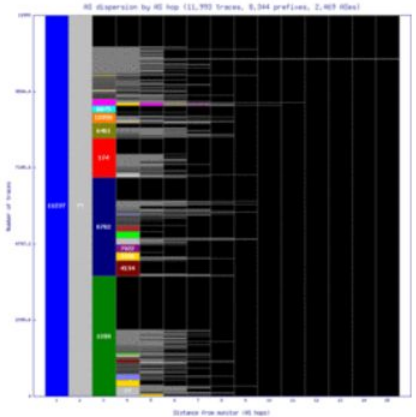

AS connectivity near monitor

AS Path Dispersion (by IP Hop)

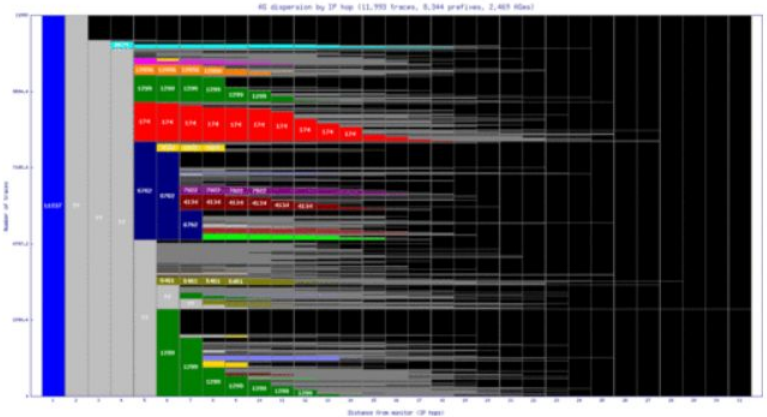

AS connectivity near monitor (with IP hop distances preserved)

IP Path Dispersion (by IP Hop)

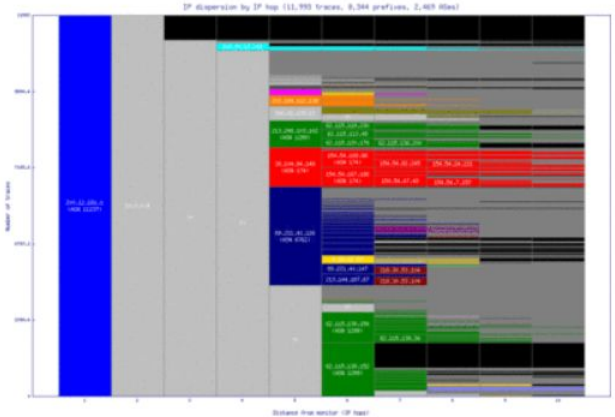

IP connectivity near monitor

AS Path Dispersion (by AS Hop)

- This graph will give you an immediate view of the AS peering relationships near the monitor. It is primarily useful to compare with the hosting organization's own information about what ASes are providing transit for your data, to see whether it matches what the Ark probes have discovered.
- In this graph, we show the AS-level path dispersion, where all adjacent hops within the same AS are collapsed into a single hop. Typically, the first hop will be all one AS (ie, the local network that the monitor is in), with the second or third hop starting to split into different ASes as probes go to their destinations. This is seen as one solid contiguous column with several smaller column segments to its right.

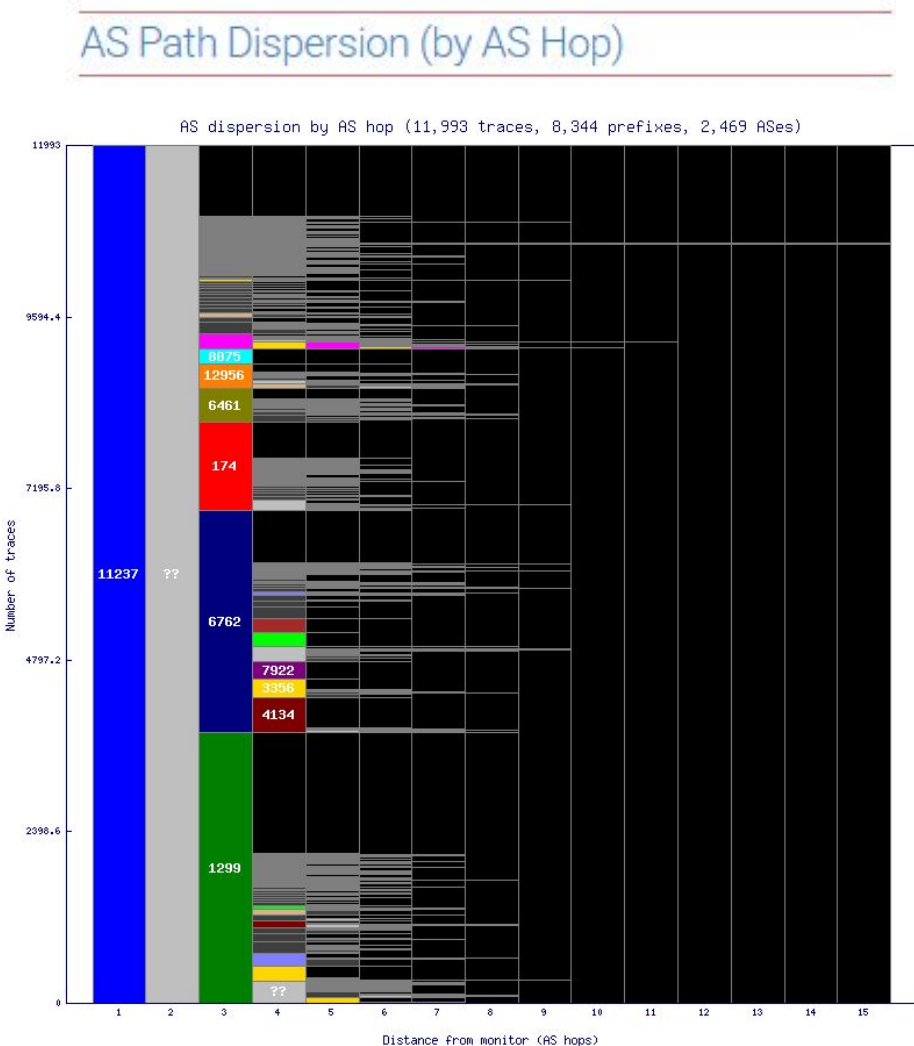

AS Path Dispersion (by IP Hop)

- This graph tells you more (compared to the AS Dispersion by AS Hop graph) about where different ASes transit data to their peers.
- For instance, some routers within an AS will pass off their packets to many other ASes at a single IP hop, whereas others will send some packets to different ASes but continue to move data within their own AS.
- It is useful to compare these results with the hosting organization's information about what ASes are providing transit for your data.

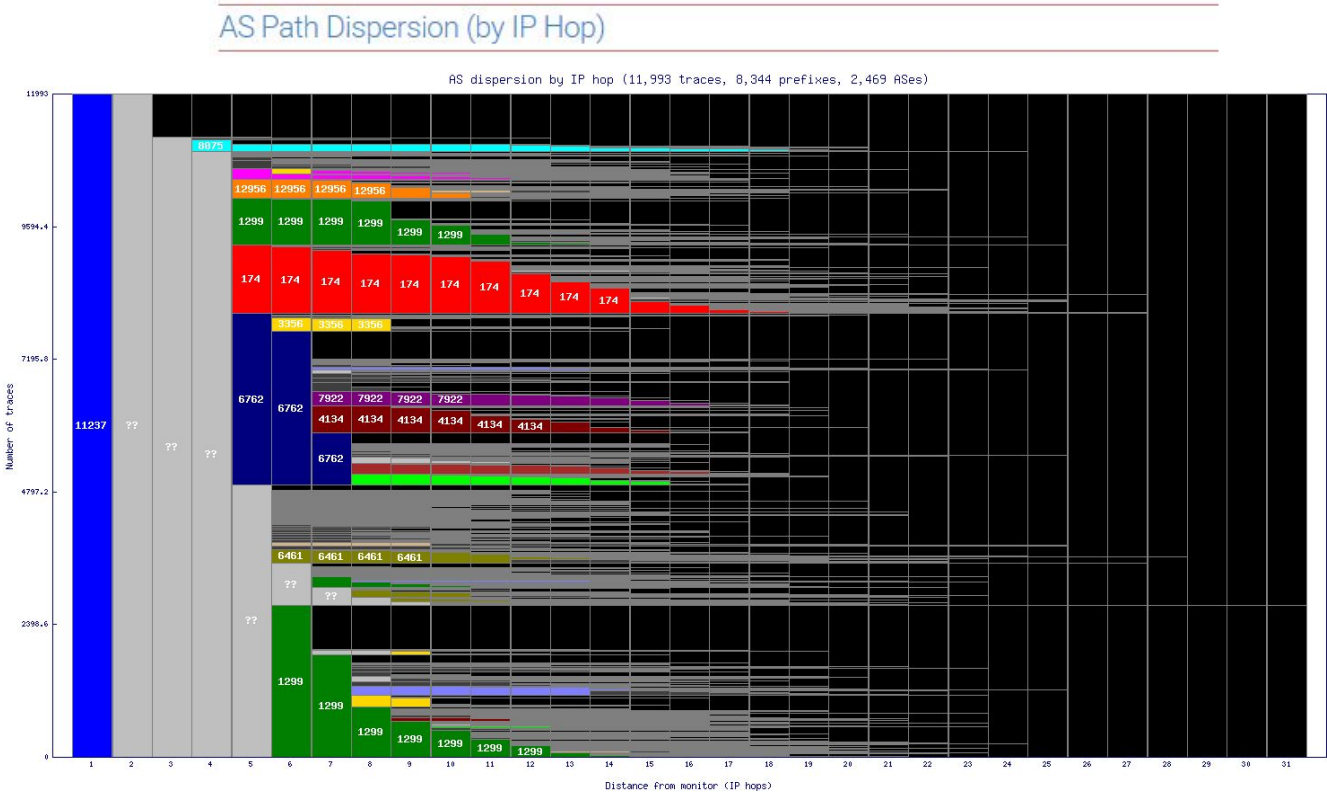

IP Path Dispersion (by IP Hop)

- In this graph, we show the IP-level path dispersion, colored and annotated by corresponding AS number. As with the AS dispersion graphs, you might encounter several column segments of roughly the same size next to each other, which indicates a string of IP addresses that are common for many paths.
- One unique quality that only shows up in the IP-level graph, however, is the symmetry caused by **load-balancing**. This is seen as a block splitting into two or more equal segments (within the same AS) that have the same pattern of column segments to their right, causing a repetition of IP addresses within the same column.

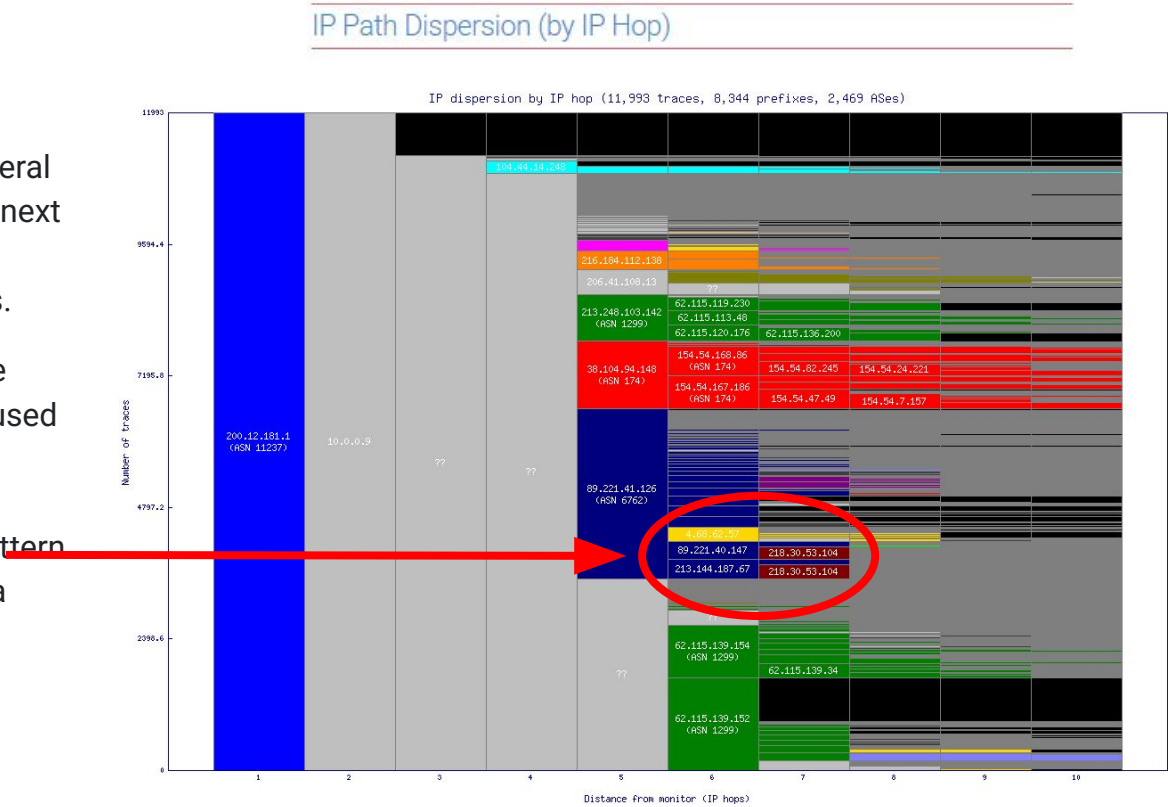

Path Length Distributions

Path length distributions show, on aggregate, how well-connected a monitor is to the rest of the Internet.

CCDF of IP path lengths for responding destinations

percentile	10th	25th	50th	75th	90th	Max
IP path length	9	11	14	15	17	26

CCDF of AS path lengths for responding destinations

percentile	10th	25th	50th	75th	90th	Max
AS path length	4	4	4	5	6	12

CCDF of IP path lengths for non-responding destinations

percentile	10th	25th	50th	75th	90th	Max
IP path length	6	6	9	13	16	33

CCDF of AS path lengths for non-responding destinations

percentile	10th	25th	50th	75th	90th	Max
AS path length	3	3	3	4	5	11

RTT Distribution

An RTT distribution shows the general latency between a monitor and the rest of the Internet. By showing the distribution of RTT values to all responding destinations, we can get a sense of how varied the speeds are for connecting to different points in the Internet.

When the CCDF graph has a nearly vertical dropoff point, that indicates that RTT values fall within a narrow range. This tends to mean that a bottleneck exists within the monitor's connectivity that dominates over individual destination path variation. A more gradual curve, on the other hand, indicates greater variability in the response times of destinations, which tends to scale directly with the path length distribution.

CCDF of destination RTTs

percentile	10th	25th	50th	75th	90th
RTT (ms)	91.733	119.396	180.437	254.854	296.181

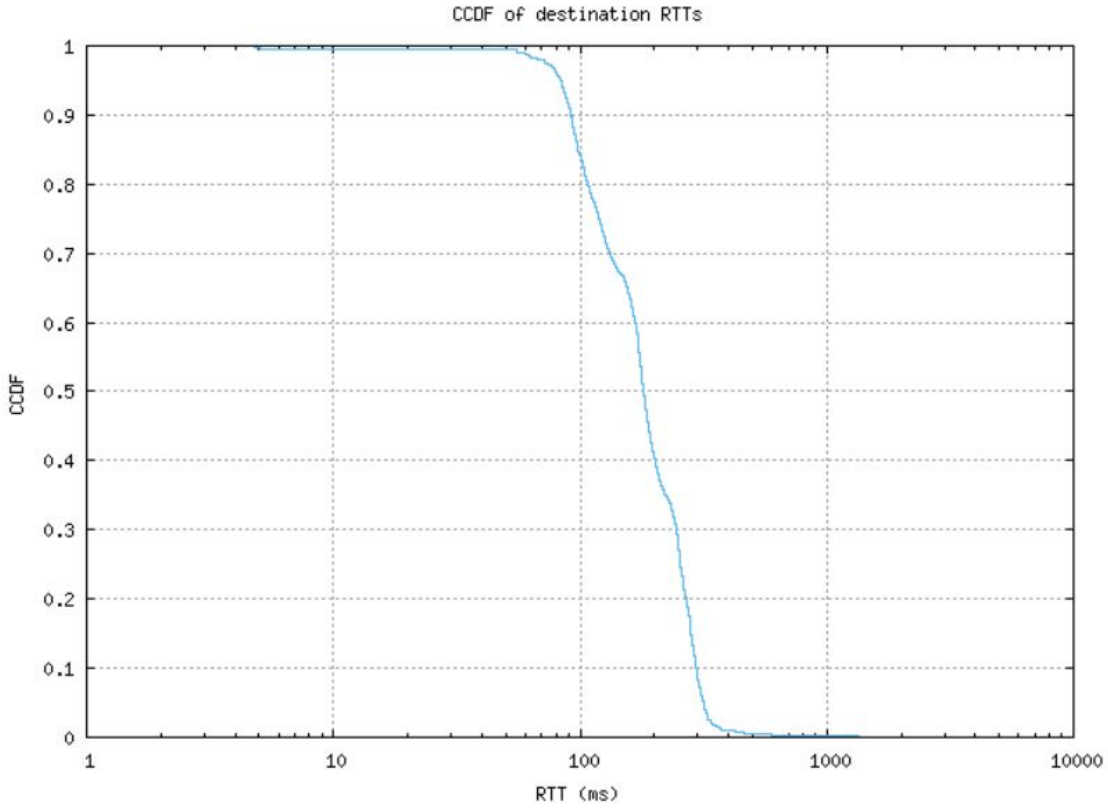

RTTs vs. Distance

med-co

Medellín, Colombia

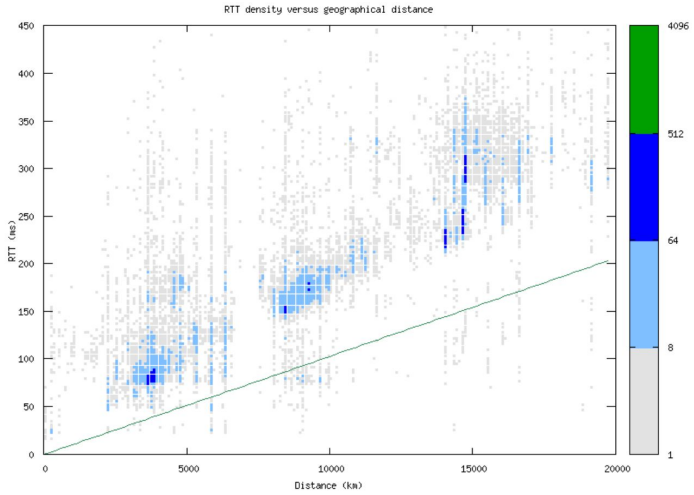

iad3-us

Leesburg, VA, USA

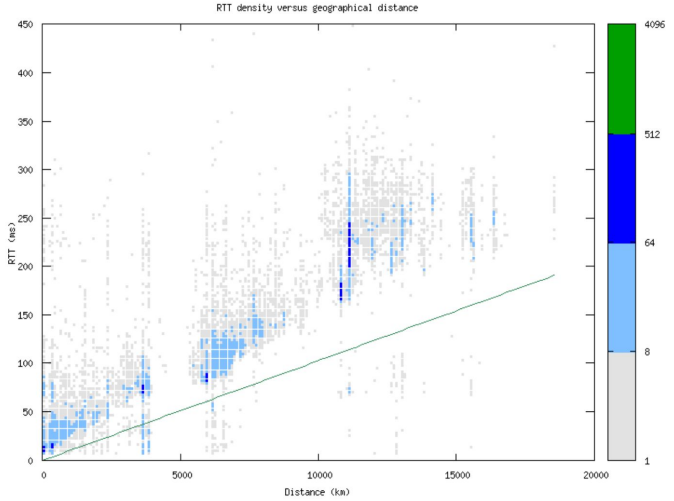

Comparing RTT with geographical distance helps show if latency is primarily due to distance or other factors. Well-connected monitors show RTT distributions that track speed of light, forming a linear baseline.

Is latency primary due to distance or other factor?

med-co

Medellín, Colombia

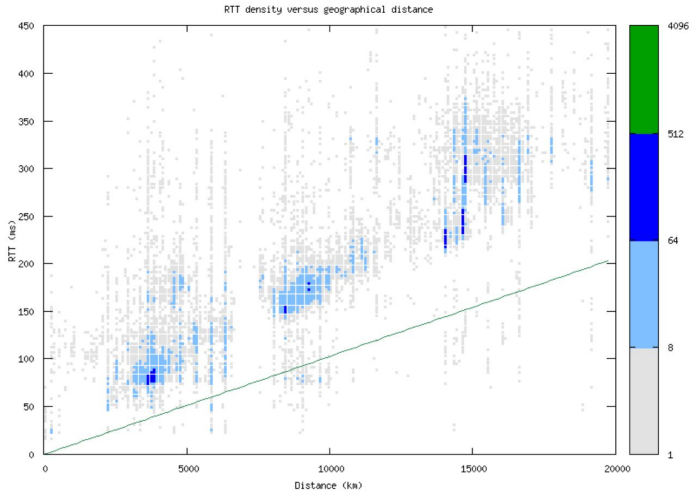

hkg-cn

Hong Kong, China

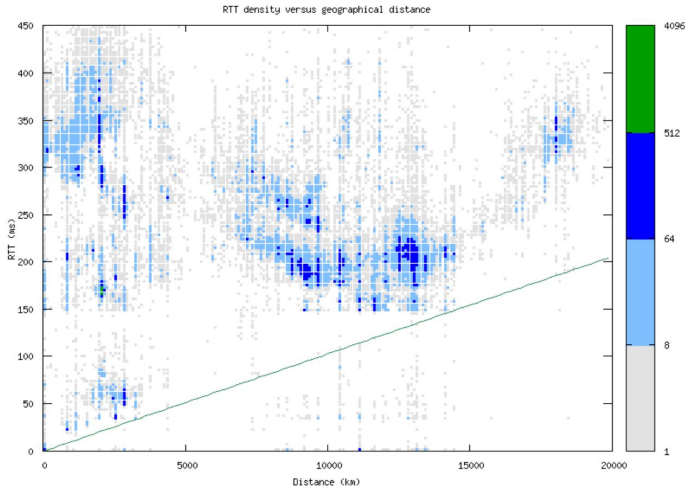

Less well-connected monitors (on right) show more scattered RTT distributions.

Archipelago data is available to researchers

- Raw traceroute data 2007-Present (IPv4-IPv6)
  - > 7 TB of trace data
- Curated topology snapshots: Internet Topology Data Kit (ITDK), two per year
  - Router-level topology
  - Router-AS assignment
  - DNS Names
  - Geolocation
- Traceroute-derived IPv4 and IPv6 AS links

Search

Displaying **12** of **126** Results for Query: **"ark"**

Type Clear Filters

- [x] **26** Datasets
- [ ] **14** Presentation
- [ ] **64** Papers
- [ ] **7** Recipes
- [ ] **3** Software
- [ ] **11** Media
- [ ] **1** Collection

Date

1998 2024

Access (22)

filter access

Tags (117)

filter tags

Status

- [ ] **14** complete
- [ ] **12** ongoing
- [ ] **2** supported

Ark IPv4 Routed /24 Topology (tlv-il) This dataset contains the Ark IPv4 team-probing data captured by the tlv-il and tlv3-il Ark monitors. A subset of the IPv4 Routed /24 Topology dataset. It contains information useful for...	2023-12-13	topology ark ...More (2)	
Ark IPv4 Routed /24 Topology (iev-ua) This dataset contains the Ark IPv4 team-probing data captured by the iev-ua Ark monitor. A subset of the IPv4 Routed /24 Topology dataset. It contains information useful for studying the...	2023-01-23	topology ark ...More (2)	
Ark IPv4 Routed /24 AS Links Data from the IPv4 Routed /24 Topology Dataset are processed by using RouteViews BGP data to identify the Autonomous System (AS) associated with each responding IP address and...	2020-05-17	topology ark ...More (4)	Ark IPv4 Routed /24 Topology RouteViews Prefix To AS Mappings ...More (83)
Ark IPv4 prefix-probing data This dataset results from traceroute-based measurements running on the Archipelago (Ark) measurement infrastructure.	2015-12-08	topology ark ...More (3)	Radian: Visual Exploration Of Traceroutes Cross-AS (X-AS) Internet Topology Mapping ...More (9)
Ark DNS traffic measurements Contains quarterly samples of the DNS query and response traffic resulting from the DNS lookups. No longer ongoing; last data are from April 2014.	2015-07-24	topology DNS ...More (3)	
Ark IPv6 Routed /48 Topology This dataset contains information useful for studying the IP- and AS-topology of the IPv6 Internet. The goal of these measurements is to discover how much subnetting exists in IPv6...	2015-02-27	topology ark ...More (2)	Follow The Scent: Defeating IPv6 Prefix... How To Parse Ark IPv6 Warts File?
Ark IPv6 Topology DNS Names The IPv6 DNS Names Dataset provides fully-qualified domain names for IP addresses seen in the traces of the IPv6 Topology Dataset	2014-05-30	DNS ark ...More (4)	Ark IPv4 Routed /24 DNS Names Ark IPv6 Topology ...More (20)
Ark IPv6 Topology AS Links Data from the IPv6 Topology Dataset are processed by using RouteViews BGP data to identify the Autonomous System (AS) associated with each responding IP address and collapsing the...	2008-12-12	topology ark ...More (4)	Ark IPv4 Routed /24 AS Links Source Identification Of Spoofed DDoS Attac... ...More (18)
Ark IPv6 Topology These are all the Ark IPv6 probing data, collected by a globally distributed set of IPv6-enabled Archipelago (Ark) monitors. These data contain information useful for studying the IP and AS...	2008-12-12	topology ark ...More (3)	Ark IPv4 Routed /24 Topology Ark IPv6 Topology DNS Names ...More (39)
Ark IPv4 Routed /24 DNS Names The IPv4 Routed /24 DNS Names Dataset provides fully-qualified domain names for IP addresses seen in the traces of the IPv4 Routed /24 Topology Dataset	2008-03-01	topology DNS ...More (5)	Ark IPv4 Routed /24 Topology Ark IPv6 Topology DNS Names ...More (19)
Ark IPv4 Routed /24 Topology These are all the Ark IPv4 team-probing data, collected by a globally distributed set of Archipelago (Ark) monitors. IPv4 Routed /24 Topology dataset. It contains information useful f...	2007-09-13	topology ark ...More (3)	Ark IPv4 Routed /24 AS Links Ark IPv4 Routed /24 DNS Names ...More (190)
Ark IPv6 Topology Dataset The IPv6 Topology Dataset contains information useful for studying the IP- and AS-topology of the IPv6 Internet. The focus of this measurement is on discovering topology and not on finding...	2008-12-12	active topology ...More (4)	Tracking IPv6 Evolution: Data We Have And... IPv6 Alias Resolution Via Induced...

CAIDA's IPv4 and IPv6 AS Core Graph

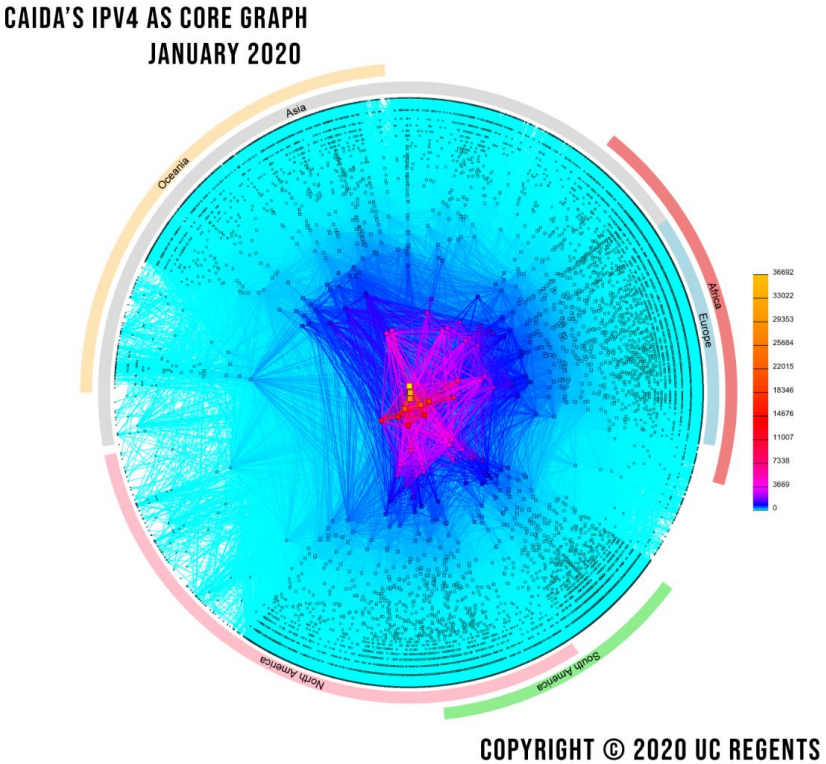

IPv4 AS Core 2020.

Visualizing IPv4 and IPv6 Internet Topology at a Macroscopic Scale

The CAIDA AS Core visualization depicts the Internet's Autonomous Systems' (ASes) geographic locations, number of customers, and interconnections.

For the IPv4 visualization we used CAIDA's Jan 2020 IPv4 Internet Topology Data Kit (ITDK) and AS Relationship data. We obtained the raw IPv4 topology data for the ITDK by performing traceroutes to randomly-chosen destinations in each routed BGP prefix using **159 Ark monitors located in 50 countries.**

https://www.caida.org/projects/as-core/2020/

CAIDA Archipelago Data Usage: Outcomes

Inferring Persistent Interdomain congestion

Dhamdhere et al — SIGCOM '18

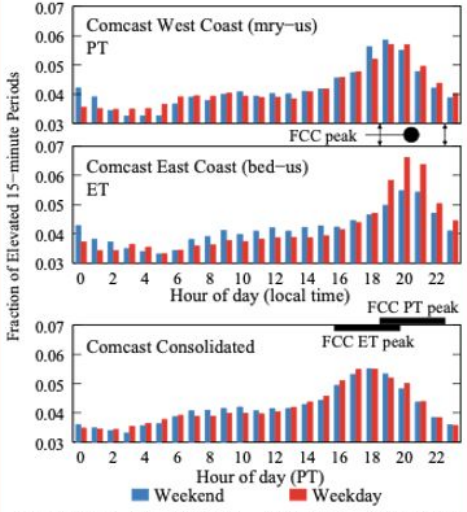

Distribution of recurring 15-minute congestion preiods from 2017 as seen from VPs in Comcast

CAIDA/UCSD, MIT, U. Waikato, TU Munich

Illuminating Router Vendor Diversity Within Providers and Along Network Paths

Albakour et al — IMC '23

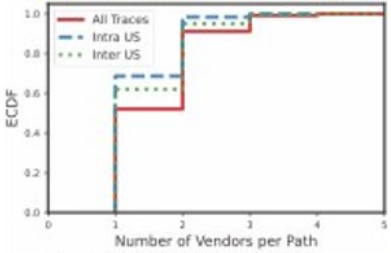

Router vendor diversity of a path

MPII, Delft U. Tech

Zeroing in on Port 0 Traffic in the Wild

Maghsoudlou et al — PAM '21

Traffic between top 10 (source AS, destination AS) pairs involved in port 0 traffic in the MAWI-short dataset.

Max Planck

(sample papers)

Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet

Luckie et al — CCS '19

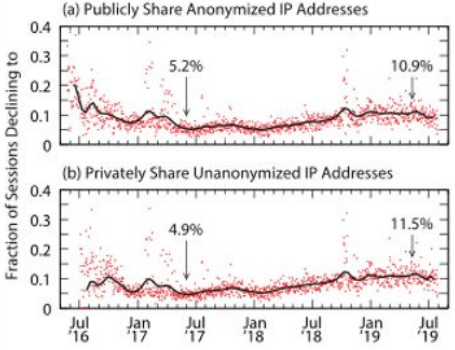

Fraction of Spoofer project tests with the daemonized client per day with sharing restrictions, overtime. The daemonized client was released inMay2016. The Bézier curves hows the trend that the percentage of private tests grew from 5.2% of tests in June 2017, to 10.9% in May 2019.

Waikato, Naval Postgraduate, CAIDA

Solar superstorms: planning for an internet apocalypse

Jyothi et al — IMC '23

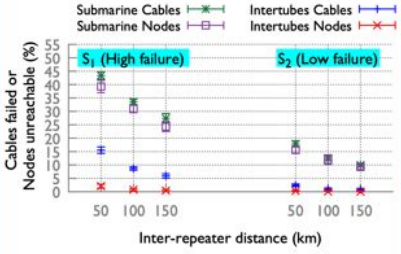

Cable and nodes failures under two states of non-uniform repeater failures ((1, (2). In each state, repeaters in a cable are assigned a failure probability based on the highest latitude (!) endpoint of the cable. Three levels of failure are ! > 60, 40 < ! < 60, and ! < 40. Assigned failure probability per repeater in (1 is [1, 0.1, 0.01] and in (2 is [0.1, 0.01, 0.001] across the three levels respectively.

UC Irvine, VMware Research

Typical uses of Ark nodes

- Topology measurements such as ICMP, TCP, and UDP traceroute and ping measurements to a broad cross section of the routed address space
- Performance measurements such as one-way delay, loss, and jitter
- Measurements of or leveraging DNS infrastructure
- **Traffic from Ark nodes running the default suite of tests is 95% traceroute and 5% ping**
- **Current Ark tests run at approximately 100 pps, using small packets, so typically generate less than 1 Mbps of traffic**

Next Generation Active Measurement Programming (NiteOwl)

- Python-based programming environment deployed on Ark nodes
- Customized for a next-generation Internet Active Measurement Infrastructure.
- **Researchers can use available primitives to build their own measurement tools.**
- More powerful approach to building distributed measurement systems.
- Extensive Public Documentation: https://www.caida.org/catalog/software/scamper/python/

NiteOwl Measurement Primitives

Ping	***MDA traceroute***	***Pair-wise IPID Alias Resolution***	***MIDAR estimation***
Traceroute	***ICMP pkt capture***	***Common src A.R.***	***MIDAR discovery***
DNS query	***TCP behavior inference***	***Find alias in subnet***	
HTTPS query	***UDP probes***	***Bulk IPID Alias Resolution***	

Container-based Ark node software

- Ark container images are Debian-based, with Ark-specific packages.
- CAIDA publishes images to the Docker Hub Container Image Library
- Hosts of containerized Ark monitors are required to follow the standard onboarding process.

- Published Container: https://hub.docker.com/r/caida/ark
- Ark monitor onboarding info: https://www.caida.org/projects/ark/moc/

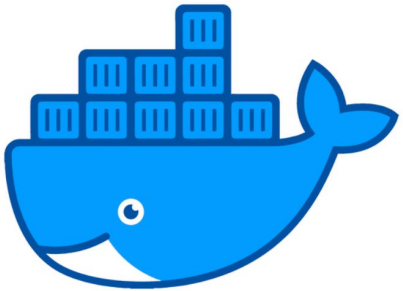

Raspberry Pi 4B based Ark Network Monitor

- Broadcom BCM2711, quad-core Cortex-A72 (ARM v8) 64-bit SoC 1.5GHz
- 8GB LPDDR4-3200 SDRAM
- Gigabit Ethernet; 2 USB 3.0 ports; 2 USB 2.0 ports
- 2 × micro-HDMI ports
- 64GB Class-10 high endurance SD card
- 5V/2A power supply
- **Used for deployments where containers are not practical**

Container based Ark vs Raspberry Pi

- Containerized Ark deployment is preferred by CAIDA because:
  - Much Faster Deployment
  - No need to purchase, configure, and then ship a Raspberry Pi
  - No issues with clearing customs or being charged import duties
  - No risk of hardware being misplaced
- Containerized Ark deployment is often preferred by hosting institutions because:
  - Other containerized tools such as RIPE Atlas and perfSONAR can run alongside Ark
  - Engineers and Researchers can use NiteOwl primitives to develop homegrown network measurement/optimization tools
  - Host institutions can leverage a library of community developed tools

NRP NATIONAL RESEARCH PLATFORM & INTERNET2

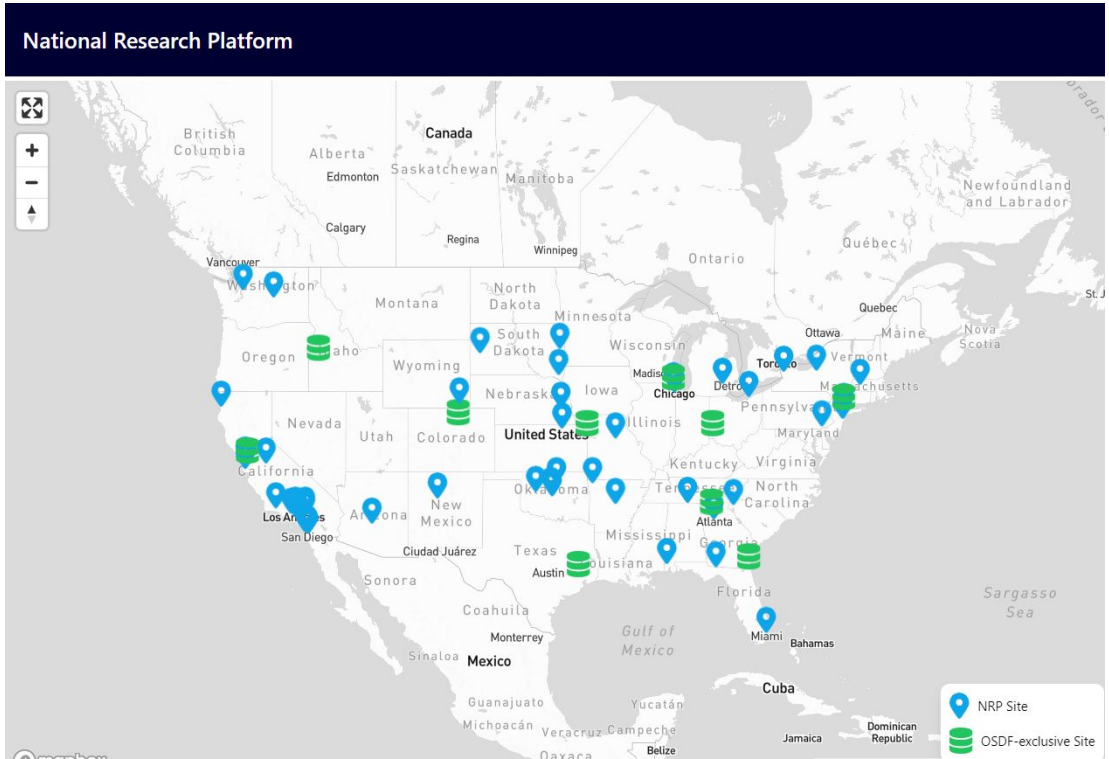

- The National Research Platform is a partnership of more than 50 institutions, led by researchers at UC San Diego, University of Nebraska-Lincoln, and UC Berkeley.
- 77 Sites are hosting 426 NRP Kubernetes Nodes
- **41 Ark monitors are running on NRP nodes**
- **Some of these nodes are also running RIPE Atlas & perfSONAR**

	ARK	RIPE Atlas	perfSONAR
Description	A globally distributed measurement platform that operates measurement campaigns, curates results into an Internet topology data kit, and supports vetted access to programmable environment	A globally distributed measurement platform that operates several measurement campaigns and supports broad execution of limited set of measurements	Network measurement toolkit that provides federated coverage of network paths, interface to schedule measurements, store, visualize data
Type of Measurements	ping, traceroute, DNS, HTTP, UDP probes, packet capture, IP address alias resolution (tightly coordinated)	Built-in: RTT to first 2 hops, ping/traceroute/SSL queries to routing table, DNS queries to root servers. User-defined: ping, traceroute, DNS, TLS and NTP query to any destination	Throughput (TCP and UDP), round-trip time, one-way delay, one-way packet loss, network path
User-Defined Scheduled Measurements	Yes	System of credits based on participation	Yes
On-demand Measurements	Yes	System of credits based on participation	Yes
Incoming Measurements	Yes; site-programmable	No	Yes; site-programmable
Target Control	Yes; site-programmable	Any non-local destination	Yes; site-programmable
Type of Distribution	Hardware or software	Hardware device or software package	Software
Data Storage	Central	Central	Local or central; many archive types
Storage Architecture	Data stored in various documented topology related datasets	Data stored in service provider infrastructure + Google BigQuery	Data stored in user infrastructure
Access to Archive Measurements	Central web interface	Central web interface	Local or central web interface
Source of Funding	US Government Research Funding	RIPE Members (ISPs)	US Government Research Funding

Internet2/NRP Ark Measurement Nodes

41 nodes in 35 autonomous systems across 37 cities

6 of these nodes are in Internet2 PoPs

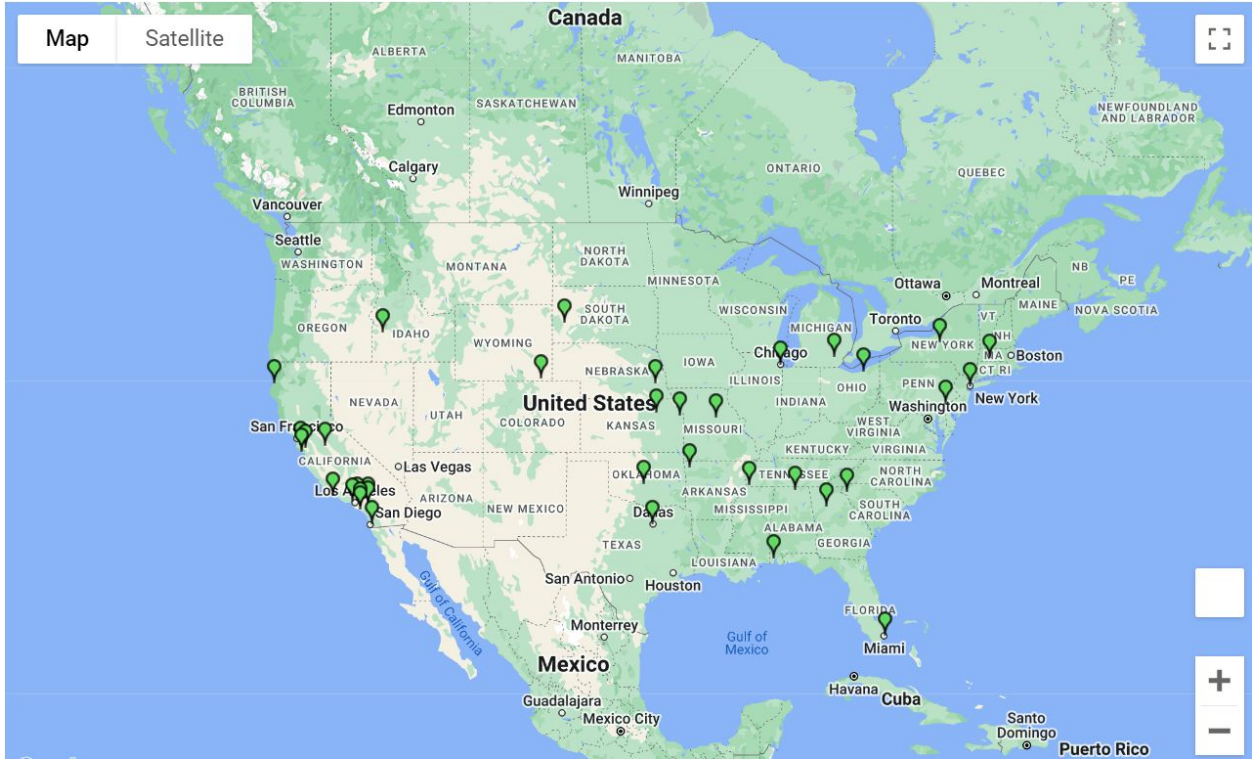

Internet2's 40+ Performance Assurance Service Nodes are being added to NRP and will be running Ark, RIPE Atlas and perfSONAR

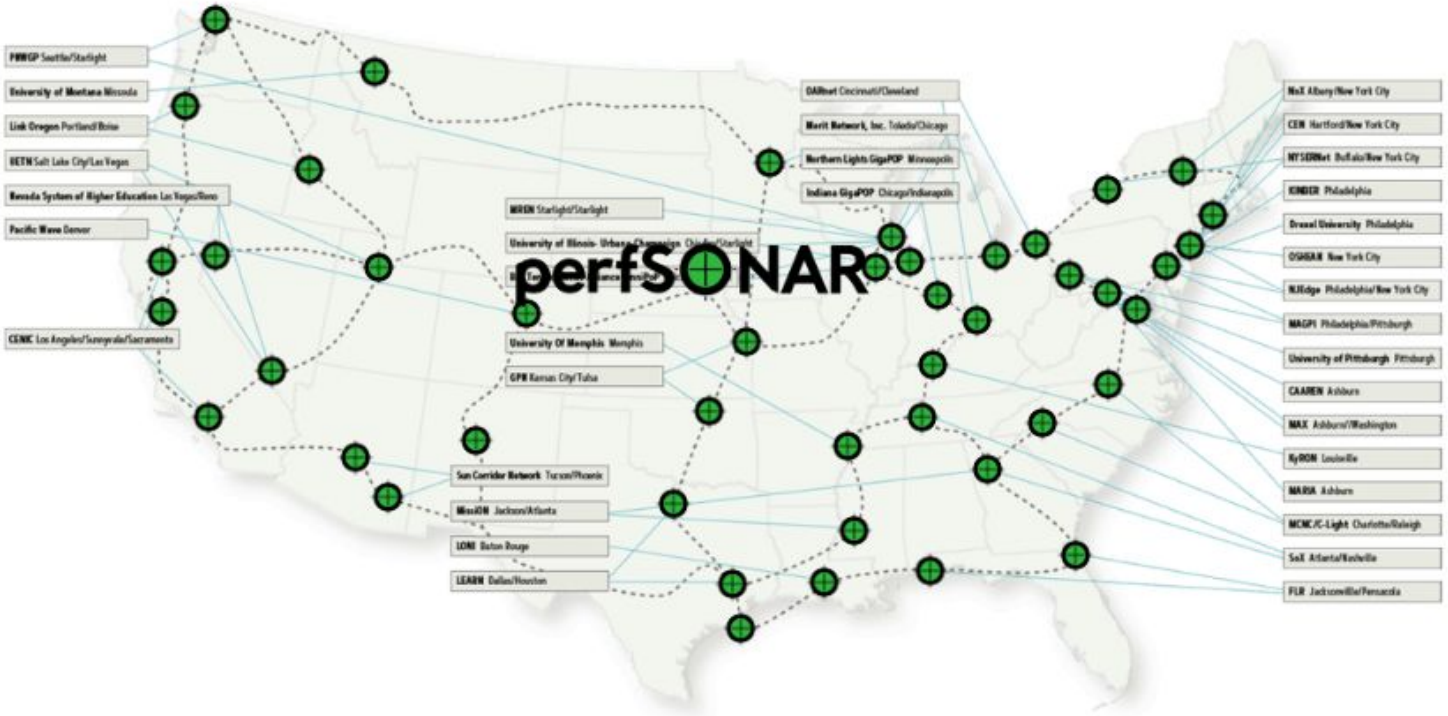

Currently these nodes are only running perfSONAR

Internet2's Local Pref Probe

- Internet2 (USA) used NiteOwl python based primitives to build a Local Pref Probe that they deploy on existing Ark nodes.

- They designed and implemented a method to infer an aspect of routing policy of individual ASes: whether they reach Internet2 through an R&E route.

- Of those that they could classify, targets in 83% of prefixes would reach Internet2 through an R&E route. The asymmetric routing for the other 17% can be attributed to the local configuration at origin.

- The code for this application will be made available in the CAIDA github repository:
  - https://github.com/CAIDA/ark-community

INTERNET2®

RNP, CUDI, & RedCLARA are all deploying Ark nodes

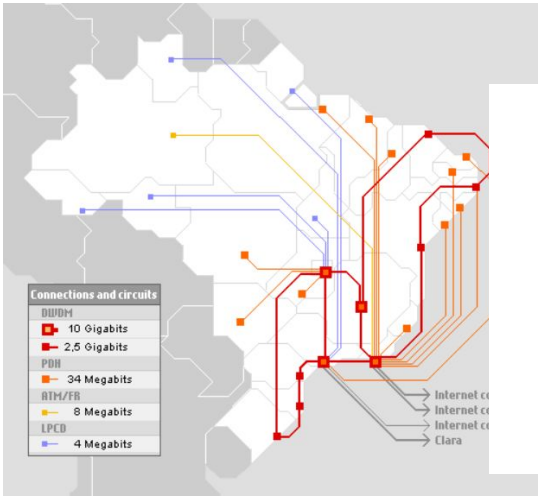

RNP - Brazil (6)

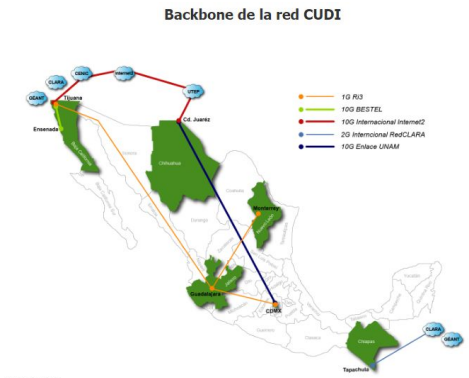

CUDI - México

RedCLARA - Latin America (3)

Thanks!

www.caida.org/projects/ark

For questions, or to offer hosting: **jj@caida.org**