



Center for Applied Internet Data Analysis
San Diego Supercomputer Center

9500 Gilman Drive #0505
La Jolla, California 92093-0505
Tel: (858) 534-5000
Fax: (858) 534-5113

March 26, 2024

The Archipelago (Ark) Project¹ is an active measurement infrastructure for Internet topology discovery operated by CAIDA. The primary goals are to

- reduce the effort needed to develop and deploy sophisticated large-scale measurements,
- provide a step toward a community-oriented measurement infrastructure by allowing collaborators to run their vetted measurement tasks on a security-hardened distributed platform, and
- collect and distribute datasets that support a variety of research, engineering, and operational interests that promote a robust, scalable global Internet infrastructure.

Ark is tailored specifically for active network measurement. This allows Ark to be simpler than other general-purpose distributed experimental platforms and to concentrate on providing facilities that directly address the needs of networking research. For example, we provide a facility for communication and coordination that makes it easier to write distributed measurements that must work together to achieve a goal. Our aim is to lower the barrier to bringing novel and interesting measurement techniques to life.

Organizations that wish to host an Ark monitor agree to our Memorandum of Cooperation (MoC)² and donate rack space, power, and cooling for a server dedicated to Ark measurement. Each Ark measurement node broadens the view of the global Internet for the network research community and increases the completeness and accuracy of data representing the topological structure of the Internet.

Our main goal in deploying Ark monitors is to achieve a geographically and topologically diverse view of the Internet. As of March 2024, the Ark infrastructure consists of 159 active monitors, distributed over 57 countries, 66 of which are IPv6-capable. Geographically³, monitor distribution is concentrated in North America (56 nodes) and Europe (55 nodes), and the remaining monitors are in Asian-Pacific Rim (30 nodes), South America (8 nodes), and Africa (10 nodes).

An important focus of Ark is large-scale traceroute-based topology measurements from multiple vantage points to the global IPv4 and IPv6 Internet. The results of these measurements, which have been ongoing since September 2007, are made available to the research community in several datasets. The primary dataset, the “IPv4 Routed /24 Topology Dataset”⁴, contains 16 TB of data as of March 2024. In addition, we provide the results of DNS name lookups for all IP addresses observed in traceroute paths⁵; Autonomous System (AS) relationships⁶; and “Macroscopic Internet Topology Data Kits (ITDKs)”⁷ that provide router-level topology derived from IP paths using multiple alias-resolution techniques.

Network researchers use CAIDA topology data to conceive, develop, and test their models and methods. The range of scientific experiments has successfully demonstrated our vision of a metaphorical distributed measurement “operating system” to support empirical Internet science.

1 <https://www.caida.org/project/ark/>

2 <https://www.caida.org/projects/ark/moc/>

3 <https://www.caida.org/projects/ark/locations/>

4 https://www.caida.org/catalog/datasets/ipv4_routed_24_topology_dataset/

5 https://www.caida.org/catalog/datasets/ipv4_dnsnames_dataset/

6 <https://www.caida.org/catalog/datasets/as-relationships/>

7 <https://www.caida.org/catalog/datasets/internet-topology-data-kit/>

Sincerely,

Kimberly Claffy
Director, Center for Applied Internet Data Analysis
Senior Research Scientist, San Diego Supercomputer Center
Adjunct Professor, Computer Science and Engineering, UC San Diego

For questions, email monitor-info@caida.org

About CAIDA

The Center for Applied Internet Data Analysis (CAIDA), based at the University of California's San Diego Supercomputer Center in San Diego, USA, is a collaborative undertaking among organizations in the commercial, government, and research sectors aimed at promoting greater cooperation in the engineering and maintenance of a robust, scalable global Internet infrastructure.

CAIDA investigates practical and theoretical aspects of the Internet in order to:

- provide macroscopic insights into 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.

<https://www.caida.org/>