CSSI Element: DIBBS EI: Platform for Applied Network Data Analysis (PANDA)

PI: kc claffy, Alberto Dainotti, Alistair King
Institution: University of California San Diego (UCSD)

Award #: 1724853
Office of Advanced Cyberinfrastructure (OAC)

challenges

- domain-specific knowledge required to understand data and format
- limited programming ability by many researchers
- ad-hoc code difficult to reuse

data platform integration

Platform for Applied Network Data Analysis

A Platform for Applied Network Data Analysis (PANDA) will offer researchers more accessible, user-friendly tools for collecting, analyzing, querying, and interpreting measurements of Internet infrastructure.

- Archipelago Active Internet Measurement Platform, Supporting Components, and Derivative Data (AS Relationships/Links/Types; ITDK)
- ASRank: Comparison of routing and economic relationships among ISPs
- BGPStream: Efficient framework for routing (BGP) data analysis
- Periscope: Extend topology measurement coverage using public infrastructure
- MANIC: Measurement and Analysis of Interdomain Congestion
- Spoofer: Assessment of IP source address validation best practices

existing systems

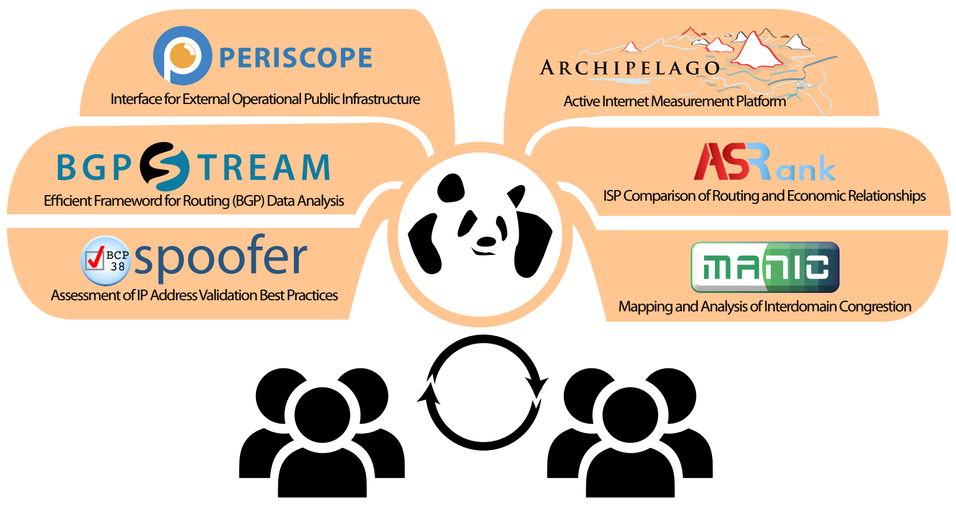

research enabled

The Impact of the General Data Protection Regulation on Internet Interconnection

Ran Zhuo, Bradley Huffaker, kc claffy, Shane Greenstein

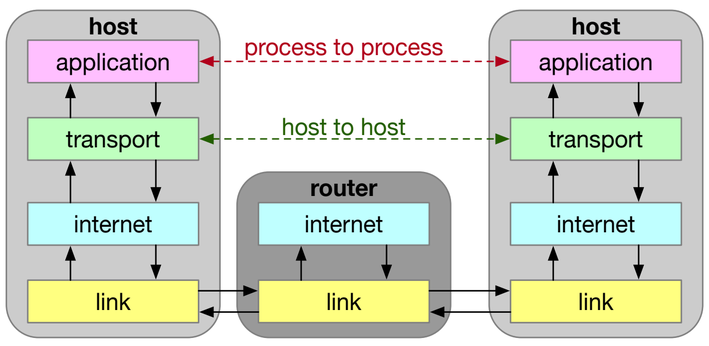

The European Union's General Data Protection Regulation (GDPR) imposes restrictions on processing and sharing of personal data of EU residents. Some economists predicted that a reduction in data usage at the application layer would negatively impact incentives for negotiating interconnection. Using topology data from Archipelago and BGPStream, we confirmed the lack of any measurable change in the number of network-level interconnections at the Internet layer before vs. after GDPR went into effect.

Unintended consequences: Effects of submarine cable deployment on Internet routing

Roderick Fanou, Bradley Huffaker, Ricky Mok, KC Claffy

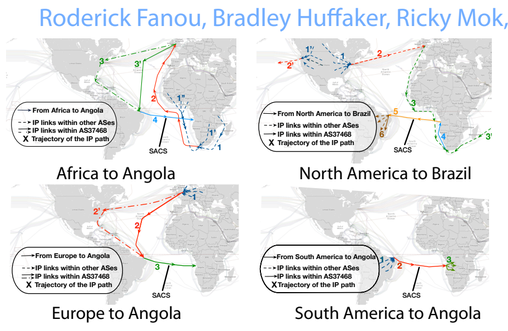

Using topology data from Archipelago and BGPStream, and AS information from ASRank, we evaluated the effects of a new transatlantic cable on the performance of paths that used it. Most source-destination pairs the crossed the cable benefited from its deploying relative to their previous path, but for a surprising number of source-destination pairs, suboptimal routing after the cable deployment actually led to worse performance

Inferring persistent interdomain congestion.

A. Dhamdhere, D. Clark, A. Gamero-Garrido, M. Luckie, R. Mok, G. Akiwate, K. Gogia, V. Bajpai, A. Snoeren, k. claffy

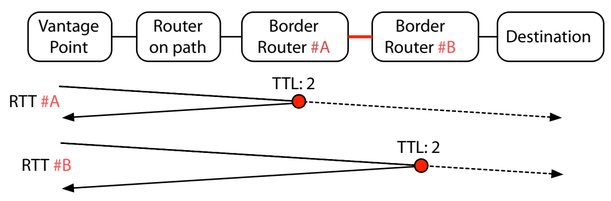

The Time-series latency probing (TSLP) method sends TTL-limited packets expiring at border routers #A and #B to measure link delay patterns.

We implemented a system based on the Time Series Latency Probes (TSLP) technique that identifies links with evidence of recurring congestion. During our window of study, we did not find evidence of widespread endemic congestion on links directly connecting access ISPs and content providers, although some such links exhibited recurring congestion patterns.

outreach

Active Internet Measurements (AIMS 11th)

CAIDA hosts annual Active Internet Measurement Systems (AIMS) workshops to promote discussion and collaborations between academics, industry, policymakers, and funding agencies, to better understand challenges and opportunities of active Internet measurement infrastructure.

Workshop on Internet Economics: Knowledge of Internet Structure: Measurement, Epistemology, and Technology (WIE 10: KISMET)

This workshop series provides a unique forum for researchers, commercial service providers, economists, theorists, policy makers, and other stakeholders to empirically inform emerging Internet regulatory and policy debates.

Regulation When Platforms Are Layered

William Lehr, David Clark, Steven Bauer, kc claffy

This study proposed the use of a model of the layered platform nature of the Internet ecosystem to improve the rigor of analysis of several ongoing regulatory debates.

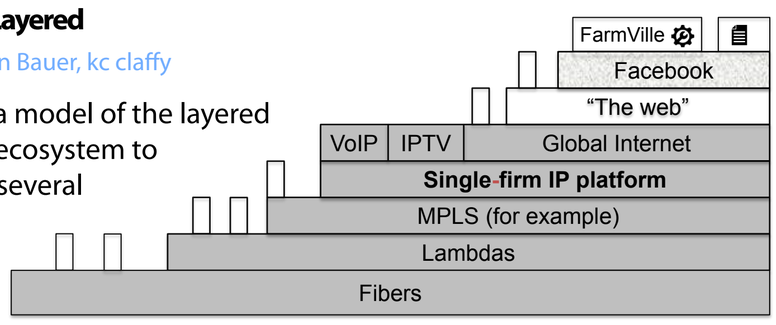

SDSC SAN DIEGO SUPERCOMPUTER CENTER

UC San Diego

The Center for Applied Internet Data Analysis (CAIDA) conducts network research and builds research infrastructure to support large-scale data collection, curation, distribution and scientific analysis. Located at the San Diego Supercomputer Center at UC San Diego, CAIDA designs, deploys and maintains computational, data analysis and visualization services that illuminate the most pressing problems of today's Internet infrastructure.

caida