

Shared internet-scale measurement platforms

Berat Can Şenel^{*,**}, Maxime Mouchet^{*,**},

Justin Cappos[†], Olivier Fourmaux^{*}, Timur Friedman^{*,**}, Rick McGeer[‡]

^{*}LIP6-CNRS lab, Sorbonne Université, ^{**}LINCS lab, [†]NYU Tandon School of Engineering, [‡]US Ignite

Critical Research Questions

The internet research community benefits from being able to deploy experimental software and conduct measurements from vantage points broadly scattered across the network. The currently widest-spread platform, RIPE Atlas, present in 6% of IPv4 and 9% of IPv6 ASes, allows a few types of measurements from mostly resource-limited devices. PlanetLab offers general software deployment, but the US-based instance is gone, while PlanetLab Europe continues with just 36 nodes that offer a four-year-old Linux distribution on nodes that may be twice that age. The software-only Seattle testbed came closest to offering the community the ability to run both general code and measurements from a large number of vantage points. Can we draw lessons from experience going back over twenty years (see table) to break through and finally put into production an internet-scale general-purpose fully measurement-capable platform? This white paper addresses two of the critical questions raised by the WOMBIR workshop: (1) “What infrastructure is needed for such [internet measurement] data?” (2) “How can we develop best practices to facilitate cooperation or collaboration with commercial service providers as we collect data?”

Issues of Most Concern

The main issue of concern is whether experience can guide us in building a contemporary platform capable of running general code that, as compared to earlier platforms, will be: easier to maintain (a smaller code base), easier to run (using industry-standard tools), and easier to deploy to large numbers of vantage points (simple software-only installation).

We will suggest EdgeNet (<https://edge-net.org>)¹ as a platform that meets these criteria and describe its current features. We will solicit discussion on the issue of which features would render such a platform maximally useful to the internet research community, and in particular to the measurements community.

Moreover, the issue of how to engage a diverse set of actors worldwide to host platform nodes is of great concern. We take

¹EdgeNet got its start thanks to an NSF EAGER grant, and now benefits from a VMware Academic Program grant and a French Ministry of Armed Forces cybersecurity grant.

inspiration from systems that have enjoyed volunteer community support for hosting nodes: DIMES, the Ono plug-in, RIPE Atlas, and Seattle, as well as, outside the internet research domain, BOINC. We will encourage discussion of how to deploy beyond universities and research laboratories (including Science DMZs), to NREN backbones, consumer ISPs, mobile devices, backbone ISPs, datacenters, the emerging edge cloud, CDNs, IXPs, homes, and businesses.

NSF Support

We see a promising NSF role in facilitating access to the broadest possible set of vantage points. Since internet research is international, we would encourage the NSF to act in collaboration with its counterparts in other countries. It can broker access to some environments through its influence alone. For other environments, it could fund access for platform providers, much as it offers Google Cloud credits to researchers today.

We would also encourage the NSF to provide funding, in collaboration with its partners in other countries, for ongoing platform management. Experience has shown that PlanetLab-type infrastructures can be effectively maintained by small teams of just a few engineers per continent, but to do so reliably requires long-term funding with a time horizon on the order of a decade.

The NSF is also well positioned to encourage open standards and best practices.

Example Experiments

To demonstrate the potential of internet-scale general-purpose platforms, we consider two experiments currently running on the EdgeNet platform. First, the Iris system surveys the internet’s IPv4 IP layer route topology. It requires vantage points distributed around the world for good measurement coverage, and it requires nodes that are powerful enough to launch the large number of probes necessary for full discovery of the load-balanced paths. Second, CacheCash is a decentralized CDN backed by a cryptocurrency. A platform capable of running custom code and test it across geographically distributed nodes is a natural fit for developing and benchmarking new content distribution paradigms.

	currently active	number of vantage points	open to run code	open to run measurements	open measurement data
Ark (prev. Skitter)	Y	over two hundred	to researchers	to researchers	Y
BrsMark	N	over ten	to collaborators	to collaborators	Y
DIMES	N	hundreds	N	N	Y
ETOMIC	N	tens	to researchers	to researchers	Y
GENI racks	Y	tens	Y	Y	slice-dependent
iPlane	N	all PlanetLab nodes	N	N	traceroute-only
M-Lab	Y	hundreds	vetted only	Y	Y
NLANR AMP	N	one hundred	N	N	Y
NIMI	N	tens	Y	Y	slice-dependent
Ono plug-in	N	over one hundred thousand	N	N	N
perfSONAR	Y	thousands	N	to NREN operators	N
Pinger	Y	tens	N	N	Y
PlanetLab {Europe}	Y	tens, formerly hundreds	Y	Y	slice-dependent
RIPE Atlas (prev. TTM)	Y	ten thousand	N	Y	Y
SamKnows	Y	thousands	N	N	N
Seattle	N	tens of thousands	Y	TCP and UDP	slice-dependent