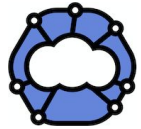

Shared Internet-Scale Measurement Platforms

Berat Senel[*,**], Maxime Mouchet[*,**], Justin Cappos[†], Olivier Fourmaux[*], Timur Friedman[*,**], **Rick McGeer[‡]**

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

EdgeNet 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.

usignite NYU TANDON SORBONNE UNIVERSITÉ Berkeley UNIVERSITY OF CALIFORNIA University of Victoria

General-Purpose Edge Clouds for Measurement Platforms

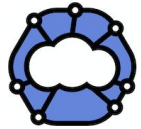

- Ex: PlanetLab hosted iPlane, others
- Many Vantage Points for Measurement
- Cost of platform amortized over services

Large-Scale Edge Clouds

Name	Dates Active	Size (Sites)	Footprint	Environment		
				Hardware	Software	Control
PlanetLab	2002-20	750	World	Linux Server	Containers	Custom
GENI	2009-	65	US	5-10 node rack Programmable Network	VMs, Containers	Custom (2)
SAVI	2011-	8	Canada	5-10 node rack Configurable Network	VMs, Containers	OpenStack
VNODE	2010-15	8	Japan	5-10 node rack Programmable Network	VMs, Containers	Custom
Glab	2009-14	17	Germany	25 node rack Programmable Network	VMs, Containers	Custom

usignite

NYU TANDON

SORBONNE UNIVERSITÉ

Berkeley UNIVERSITY OF CALIFORNIA

University of Victoria

usigni
TANDON
UNIVERSITE
UNIVERSITY OF CALIFORNIA
University of Victoria

Platform Status

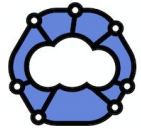

Name	Status	Cause of Death
PlanetLab	Dead	Discontinued Funding
GENI	Semi-supported	Discontinued Funding
SAVI	Dead	Discontinued Funding
VNODE	Dead	Discontinued Funding
GLab	Dead	Discontinued Funding

usignite

NYU TANDON

SORBONNE UNIVERSITÉ

Berkeley UNIVERSITY OF CALIFORNIA

University of Victoria

What Killed Them?

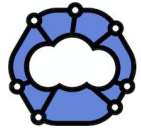

- Hardware failure
  - Hardware should be renewed on a 3-4 year cycle
  - Means: $500-$750/year/node in hardware refresh costs *alone*
  - PlanetLab: $650K-$1M/year
- Lack of central support
  - Maintaining distributed hardware infrastructures is expensive
  - PlanetLab: 6 FTEs, 770 sites
  - GENI: 5 FTEs, 65 Sites
  - $5K-$15K year/site in central administration costs
  - Rob Ricci: “Each site is weird in its own way”

Bottom Line: $10,000 site/year (PlanetLab); $50,000 site/year (GENI)

usignite NYU TANDON SORBONNE UNIVERSITÉ Berkeley UNIVERSITY OF CALIFORNIA University of Victoria

What Killed Them? - II

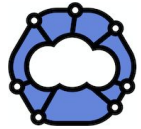

- Lack of onsite support
  - Research boxes aren't a local admin's day job
  - Hardware replacement/installation are slow, networks reassigned, etc
- Infrastructure-specific software stack
  - Maintenance & training burden without value-add
  - (Note: SAVI avoided this trap!)

usignite NYU TANDON SORBONNE UNIVERSITÉ Berkeley UNIVERSITY OF CALIFORNIA University of Victoria

Four Big Takeaways

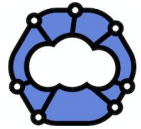

- CS Funding Agencies don't do maintenance
  - Set up to fund 3-5 year projects, grad students
- Hardware is heinous
  - Breaks and needs fixing
  - Bulk of the cost of both onsite and centralized support
  - PlanetLab reached its maintenance limit
- Don't write software unless you're actually adding value
  - If industry-standard stuff works, use it
    - Free maintenance
    - Free tutorials/answers/educational material
- Be as cheap and easy as possible for local sites

usignite NYU TANDON SORBONNE UNIVERSITÉ Berkeley UNIVERSITY OF CALIFORNIA University of Victoria

Live-off-the-Land

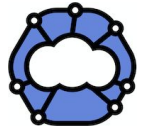

- Successful infrastructures use *local*, *existing* resources
  - Internet used *existing* computers, phone lines
  - Web used *existing* Internet, computers
    - First Web was a thin overlay on FTP
    - *Nine-line* shell script was the first web server!
- Spread is simple: software download and voluntary federation

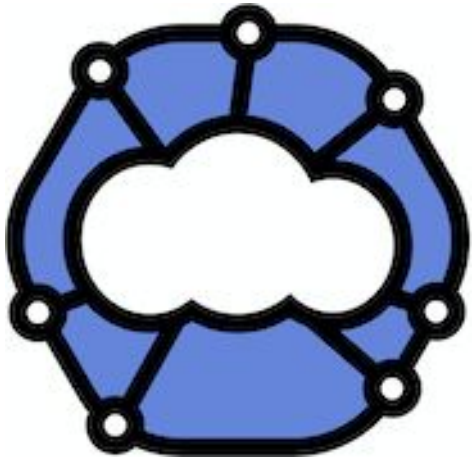

=

“Hello, World” around the world in 5 minutes

PLANETLAB

+

zero-touch install

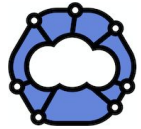

~$0 per-site maintenance

Third-party tools
Copious documentation

+

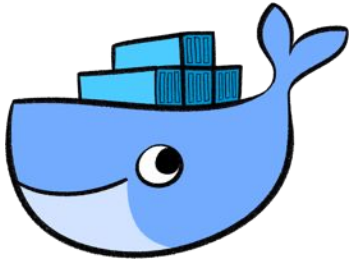

usignite

NYU TANDON

SORBONNE UNIVERSITÉ

Berkeley UNIVERSITY OF CALIFORNIA

University of Victoria

The Bottom-Up KubeEdge

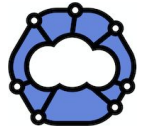

- Previous Edge Clouds Ran on Dedicated Hardware
  - Made hosting a rack/machine expensive
- But every site has a local Cloud…
  - Downloadable image onto *existing* local Cloud
  - Boots, “phones home”, to connect
  - Node adds to Edge-Net in 5 minutes!
- Site can host Edge-Net node part- or full-time

Status

- EU/Japan/Brazil/US joint project
  - Rio Workshop, November 2017
  - Canada has joined
- 36 Sites, 28 Active
  - 33 US, 1 South America, 2 Europe
- Ready to accept alpha users/systems
- Soon: backbone nodes added using Google Cloud

Experiments on edgeNet

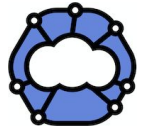

- CashCache: Decentralized CDN backed by cryptocurrency
- Paris Traceroute: robust network measurement
- Data Capsules: secure data transport across untrusted nodes

usignite NYU TANDON SORBONNE UNIVERSITÉ Berkeley UNIVERSITY OF CALIFORNIA University of Victoria

EdgeNet Architecture

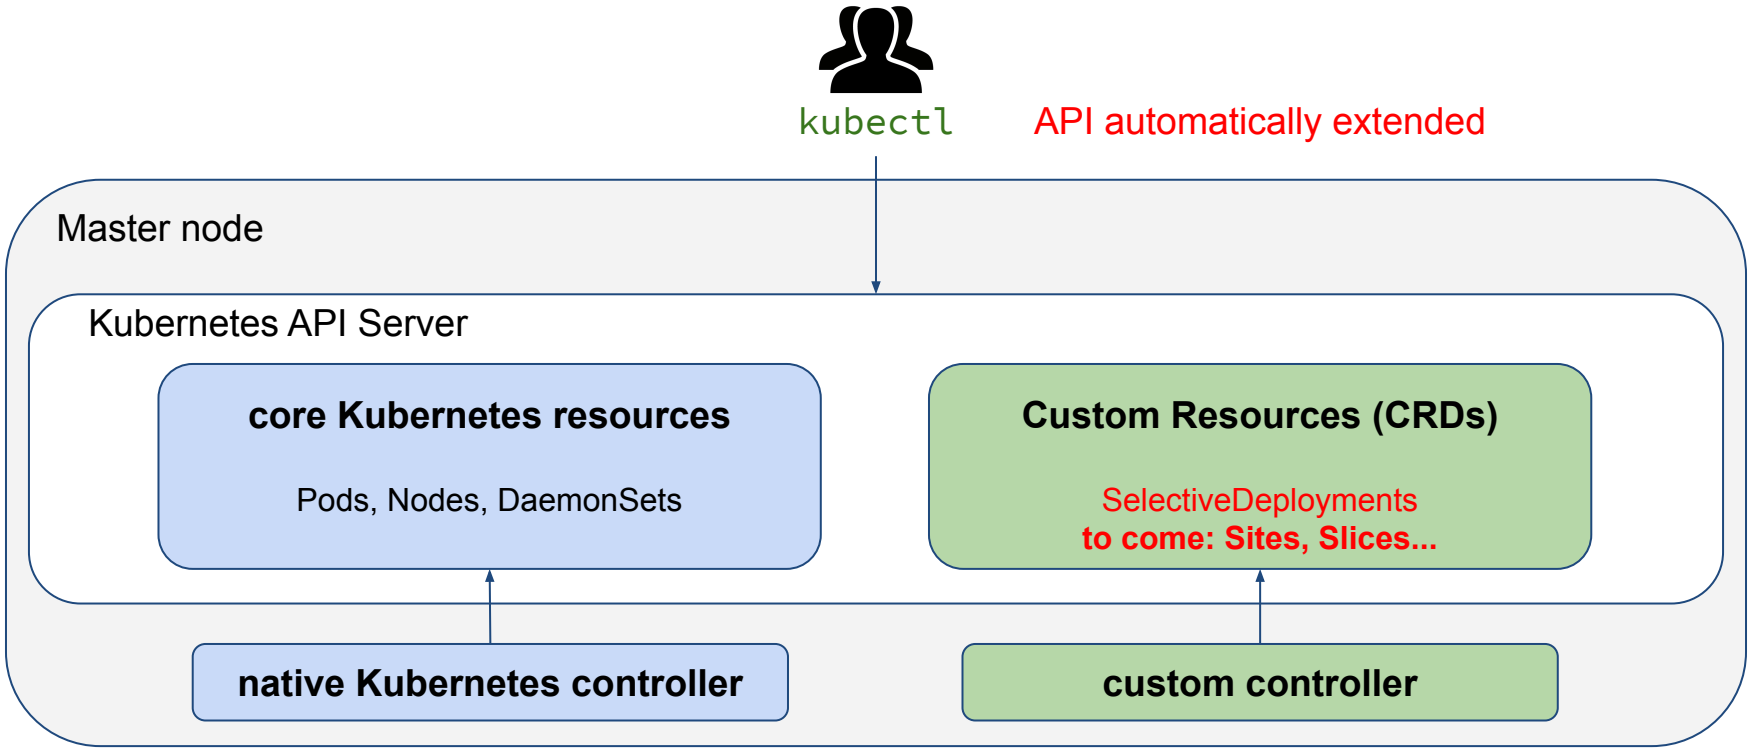

usignite NYU TANDON SORBONNE UNIVERSITÉ Berkeley UNIVERSITY OF CALIFORNIA University of Victoria

Thanks to Collaborators

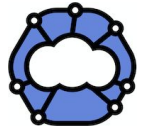

Eric Allman, Hadi Bannazadeh, Ilya Baldin, Andy Bavier, Andi Bergen, Mark Berman, Peter Boothe, Ketan Bhardwaj, Sushil Bhojwani, Marshall Brinn, Jack Brassil, Chip Elliott, Sudhakar Ganti, Ada Gavrilovska, Jim Griffioen, John Kubiatowicz, Al Leon-Garcia, Thomas Lin, Ken Lutz, Sean McGeer, Nitesh Mor, Hausi Muller, Aki Nakao, Pat O'Connell, Riz Panjwani, Tom Mitchell, Larry Peterson, Albert Rafetseder, Glenn Ricart, Rob Ricci, Ciro Scognamiglio, Robert Taylor, Ulrike Stege, Vic Thomas

https://edge-net.org/

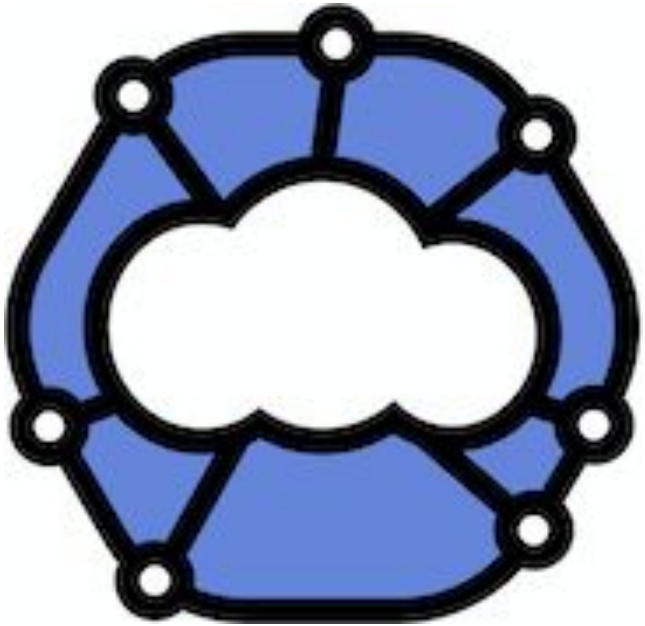

usignite

NYU TANDON

SORBONNE UNIVERSITÉ

Berkeley UNIVERSITY OF CALIFORNIA

University of Victoria