

Thoughts on Experimental Midscale Infrastructure and NSFCloud

Jim Kurose
School of Computer Science
University of Massachusetts
Amherst MA USA



Overview

- ❖ Background: CISE-AC midscale infrastructure committee
- ❖ High-level observations
 - importance of midscale research infrastructure
 - sustainability
 - community
- ❖ “In the weeds” observations

CISE-AC Midscale Infrastructure Committee

Subcommittee of CISE-AC, charged by F. Jahanian (2012-2014):

- ❖ How should community infrastructure *requirements* be derived?
- ❖ How can CISE articulate a framework for *understanding the value* of novel infrastructure to transformational research?
- ❖ What are the best models of *funding* community mid-scale infrastructure?
- ❖ *Future research infrastructure*: leveraging GENI and beyond

CISE-AC Midscale Infrastructure Committee

Midscale Infrastructure Committee

P. Barford, U. Wisconsin,
 F. Berman*, RPI
 S. Corbato, U. Utah
 J. Fortes*, U. Florida
 J. Kurose*, U. Massachusetts (co-chair)
 K. Marzullo, NSF
 E. Lazowska, U. Washington
 B. Maggs, Duke, Akamai (co-chair)
 B. Lyles, NSF
 J. Mogul, Google
 D. Raychaudhuri, Rutgers
 J. Rexford*, Princeton

**CISE Research Infrastructure: Mid-Scale Infrastructure -
 NSFCloud (CRI: NSFCloud)**

**PROGRAM SOLICITATION
 NSF 13-602**



National Science Foundation
 Directorate for Computer & Information Science & Engineering
 Division of Computer and Network Systems
 Division of Advanced Cyberinfrastructure

Full Proposal Deadline(s) (due by 5 p.m. proposer's local time):

December 17, 2013

Charge 2: Framework for understanding infrastructure value: findings

- ❖ Explaining to community why CISE should fund research infrastructure
 - concrete examples of past success
 - “virtuous cycle” between facilities and experimental systems research
- ❖ Qualitative value metrics: enabling research, training systems researchers, better paths to practice
- ❖ Quantitative metrics:
 - impact metrics
 - use/subscription metrics
 - scale metrics
 - cost metrics

Source: CISE-AC
MIC presentation

Charge 3: Funding models: findings

- ❖ More strongly link infrastructure investments with science outcomes, impact
- ❖ Create shared business models with other sectors
 - campus co-investments, industry, other agencies
- ❖ Optimize midscale infrastructure investments
 - enable research platforms, as well as (separate) cloud services

Source: CISE-AC
MIC presentation

Overview

- ❖ Background: CISE-AC midscale infrastructure committee
- ❖ High-level observations
 - importance of midscale research infrastructure
 - sustainability
 - community
- ❖ “In the weeds” observations

What you’re doing is *incredibly* important


- ❖ Midscale Infrastructure for experimental system research - meeting a critical *national* need
 - enabling, performing experimental research at (mid)scale
 - creating the next generation of experimental systems researchers
 - national competitiveness

“I love XXX (university) and will miss it a lot. The computer science faculty are absolutely top-notch, and the students are the best a professor could ever hope to work with. ... There is one simple reason that I'm leaving academia: I simply love work I'm doing at XXXX (company) working on problems that are orders of magnitude larger and more interesting than I can work on at any university. That is really hard to beat”

blog post, 11/2010

Infrastructure sustainability (AM discussion: "what will it cost me?")

- ❖ Leveraging partnerships: acknowledging “value proposition” to all involved:
 - institutional commitments (“skin in the game”)
 - industry collaborations



RESTORING THE FOUNDATION
The Vital Role of Research in Preserving the American
Dream, AAAS, 2014.

It's *much* more than \$:

- enabling long-term fundamental science advances floats all boats
- translational research
- educational opportunities, workforce training
- national competitiveness

Sustainability and (testbed) architecture

- ❖ Architecture embodies implicit or explicit sustainability model:
 - much like Internet's model of organizational autonomy
- ❖ Resource ownership/management:
 - center model (resources centrally owned/managed)?
 - federated model (locally owned “resources”)?
 - two-sided marketplace?
- ❖ Testbed use (AM discussion):
 - by (central) committee
 - prioritized algorithmically, on basis of contributed (federated) resources?

Sustainability and (testbed) architecture

- ❖ The “right” architecture and the “right” abstractions triumph over enhanced-functionality (architecture != wiring diagram)
 - Internet versus ATM
 - Openflow
- ❖ A danger in being all things to all people

Building a research community

As successful builders in research infrastructure you know this, but

- ❖ Researchers, and science it enables are “ends”
- ❖ Support, community forums, community-building critical
- ❖ Community building shout-out to GENI : GECs, education workshops, educational materials (M. Berman, S. Edwards, C. Elliott, N. Riga, V. Thomas)

A few “down in the weeds” comments:

Importance of:

- ❖ data (from application and infrastructure standpoint)
- ❖ measurement – well-instrumented infrastructure
- ❖ mobility
 - tie to recent NSF workshop on Future Research Infrastructure for the Wireless Edge?
- ❖ application use (e.g., cloud infrastructure in support of CPS)

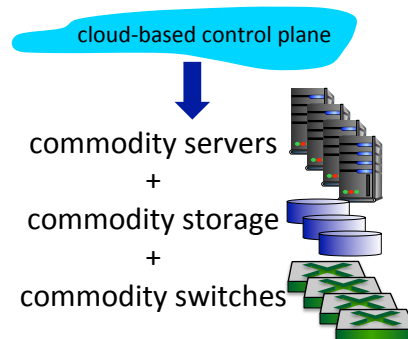
A few “down in the weeds” comments:

Importance of: SDN, NFV: “control plane”

- ❖ networking, distributed systems, PL, formal methods,

Proliferation of proprietary boxes:

Routers
Switches
CDN nodes
Firewalls
DPI
Carrier-grade NAT
Radio Access Net nodes
QoE monitors
WAN Accelerators



Summary

- ❖ Background: CISE-AC midscale infrastructure committee
- ❖ High-level observations
 - importance of midscale research infrastructure
 - sustainability
 - community
- ❖ “In the weeds” observations

Shout-out to NSF leadership



... for making this all happen