NSF CISE AC Midscale Infrastructure Committee

May 2013 CISE AC meeting

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 AC member

Midscale Infrastructure Committee(MIC): charge

today 1. How should community infrastructure *requirements* be derived?

Nov. 2012 2. How can CISE articulate a framework for *understanding the value* of novel infrastructure to transformational research?

today 3. What are the best models of *funding* community mid-scale infrastructure?

Nov. 2012 4. *Future research infrastructure:* leveraging GENI and beyond

“Midscale (research) infrastructure”:
larger than MRI, CRI, smaller than MREFC
$4M - $100M

NSF CISE Mid-scale Infrastructure Committee (MIC)

Charge 3:

Funding and management models

S. Corbató, Chair (U Utah)
F. Berman (RPI)
J. Fortes (U Florida)
J. Kurose (U Massachusetts)

May 17, 2013

Charge to the Committee (#3)

- **What are the best models of funding community mid-sized ($4M-$100M) infrastructure?**
  - *Complicated because infrastructure is often best built through community cooperation and federation with a focus on interoperability, rather than as done with MRI (Major Research Instrumentation) and CRI (Computing Research Infrastructure) -- most highly stand-alone projects funded*

- **How can CISE involve industry in mid-scale infrastructure?**

Current Status – What NSF spends on Infrastructure

Directorate/ Office	FY2010 Enacted	RI Funding/ Percentage	FY2012 Request	RI Funding/ Percentage
OPP	451.16	321.43/71%	477.41	338.02/71%
OCI (pre-CISE)	214.28	150.38/70%	236.02	148.06/63%
GEO	889.64	367.79/41%	979.16	364.96/37%
OIA	275.04	93.04/34%	336.25	93.14/28%
MPS	1,351.84	353.73/26%	1,432.73	305.51/21%
BIO	$714.54	135.45/19%	794.49	132.93/17%
SBE	255.25	43.56/17%	301.13	58.04/19%
CISE (pre-ACI)	618.83	30.60/5%	728.42	30.60/4%
ENG	618.16	32.83/5%	761.42	31.33/4%
OISE	47.83	.10/.2%	58.03	.10/.1%
CISE + ACI			964.44	178.66/19%

Current Status – What CISE – ACI spends on Infrastructure

Project	Total funding (anticipated)	FY11 Funding
PRObE - NSF Parallel Reconfigurable Observational Environment for Data Intensive Super-Computing and High End Computing	$10.77M	$1.72M
The Global Environment For Network Innovation (GENI)	$102.5M	$11M
Cooperative Association for Internet Data Analysis (CAIDA)	$14.6M	$0.69M
Network Emulation Lab (Emulab)	$16.3M	$1M
ORBIT: Open-Access Research Testbed for Next-Generation Wireless Networks	$8.39M	$0.2M

MIC Recommendations: Increase the effectiveness of mid-size infrastructure investments

❖ ***Key themes:***

1. More strongly link infrastructure investments with science outcomes, impact

2. Create shared business models with other sectors

3. Optimize mid-size infrastructure investments

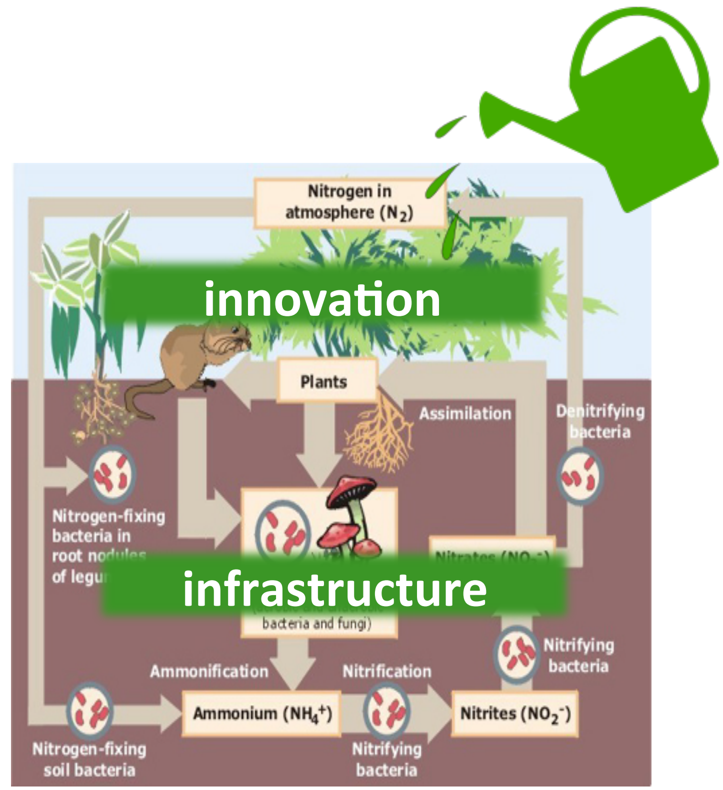

Credit: Wikipedia

Strongly link infrastructure investments with impact, outcomes

- ***Recommendation:*** **Build longer-term sustainable facilities at the mid-scale**
  - *Measure impact through outcomes, science; prioritize users*

- ***Recommendation:*** **Invest in MRI and CRI projects that provide a test bed for infrastructure innovation**
  - *Focus on both innovation in infrastructure (CISE - ACI) and infrastructure-enabled innovation (ACI)*

- ***Recommendation:*** **Establish periodic discipline-wide assessment within research areas to drive new facility design and funding priorities**
  - *Draw on broad community participation and input*
  - *Include successes and lessons learned from previously supported facilities (Analog: Astronomy & Astrophysics decadal survey)*

Create shared business models between NSF and the community

- ***Recommendation:*** **Develop funding models to incent campus co-investment in shared-use research facilities**
  - *Expect longer commitment beyond typical MRI/CRI awards with funding transition as part of the project plan*
  - *Incorporate realistic operational expenses*
  - *Expect key stakeholders to participate -- researchers, CIO, VPR, service providers*

- ***Recommendation:*** **Explore development of community-led non-profits as vehicles to sustain large mid-sized projects and focus institutional and partner co-investments**
  - *Large science project management approach with appropriate checks and balances, stakeholder advisory groups, e.g.*
    - *UCAR / Atmospheric science, AURA / astronomy, IRIS / seismology*
    - *NSF MREFC, DOE Office of Science*

Create shared business models between NSF and other sponsors

- **Public-private partnerships**
  - *Corporate support*
  - *Foundations*
- **Public-academic partnerships**
  - *Campus co-investment*
  - *Academic consortia*
- **Public-public partnerships**
  - *State co-investment*
  - *Co-sponsorship with other agencies*
  - *Co-sponsorship with other NSF Directorates*

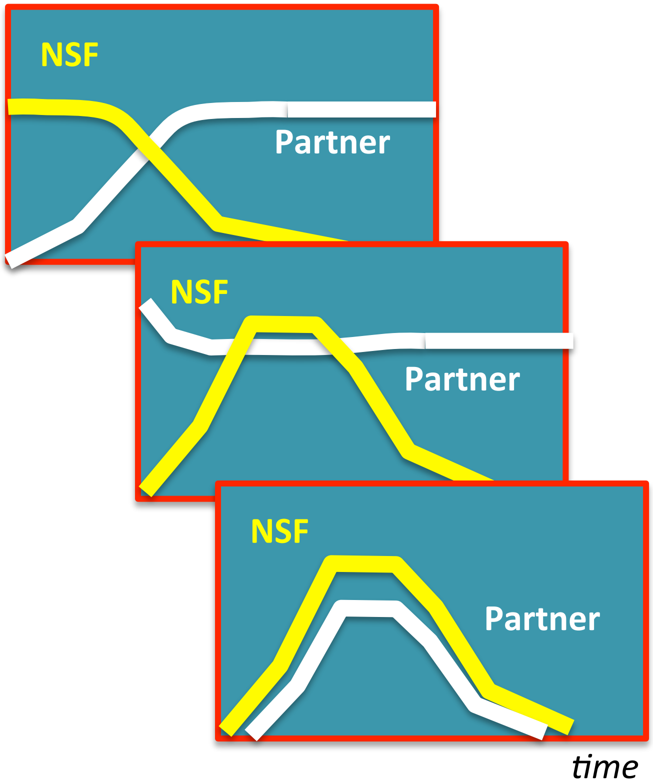

different co-investment models

Optimize MI Investments

- ***Recommendation:*** **Ensure that facilities support a broad spectrum and scale of meritorious research projects**
  - *Assess based on quantified outcomes*
  - *Link assessment, innovation, sustainability*

- ***Recommendation:*** **Explore research-focused cloud services**
  - *Developed by NSF grantees*
  - *Provided by commercial organizations*
  - *FINDING: Current research overhead rules strongly dis-incent large-scale cloud utilization in favor of traditional equipment purchases*

- ***Recommendation:*** **Require viable sustainability plans beyond the duration of NSF funding**
  - *3-5 year horizon for "ramped in" funding / cost-sharing / partner investment / in-kind contributions after the expiration of grants*

Charge 1

Community infrastructure requirements:

CCC white paper process

S. Corbato
E. Lazowska
B. Maggs
D. Raychaudhuri

Community Infrastructure Requirements: background

- *goal*: solicit community input on midscale infrastructure requirements
  - options discussed, not taken: CISE-AC only, decadal studies (e.g., Astronomy), NAS study, NSF-sponsored workshop
- white paper process:
  - solicited through CCC, advertised via CCC blog, mailing lists

    "The Computing Community Consortium is seeking community input to better understand the potential needs and payoff for additional investments in mid-scale infrastructure for computing research … "
  - 10 white papers, including inputs from many impactful MI activities (Emulab, FutureGrid, GENI, Openflow, Planetlab), experimental systems researchers

White papers:

[Cappos] Cappos *(NYU Poly),* "Three Computing Infrastructure Needs"

[Chase] Chase, Baldine, *(Duke, RENCI / UNC)*, "CCC Mid-Scale Response"

[Feamster] Feamster *(GaTech),* Banerjee *(Wisconsin)*, "An Open Observatory for the Internet's Last Mile"

[Fox] Fox, von Laszewski (Indiana), Fortes (Florida), "Mid-Scale Infrastructure Investments for Computing Research: A FutureGrid Perspective"

[Katz] Katz-Bassett *(USC),* Levin *(Maryland),* Zariffs *(USC)*, Feamster *(GaTech)*, "The Transit Portal: A Testbed for Internet-scale Routing"

[Kreiger] Krieger, Bestavros, Appavoo *(Boston U.),* "Enabling an Open Cloud"

[Landweber] Landweber *(Wisconsin),* Elliott *(BBN),* "Mid-Scale Infrastructure Investments for Computing Research"

[McKeown] McKeown, Parulkar, Peterson,Shenker, *(Open Network Lab),* "NSF OpenCloud"

[Ricci] Ricci *(Utah)*, "The Need for Flexible Mid-scale Computing Infrastructure"

[Weiss] Weiss *(Evergreen State),* Mache *(Lewis & Clark),* "Mid-scale infrastructure for experimentation in networking and security"

A common vision:

Is there a need for midscale infrastructure?   *Yes!!*

> *"A nationwide, multi-tiered system (national/regional R&E backbones, data centers, campuses) that is sliced, deeply programmable, virtualized, and federated so that research experiments can run `end to end' across the full suite of infrastructure."*

- *multi-tiered system (national/regional R&E backbones, data centers, campuses:* core/edge networking, computation, clouds
- *sliced, virtualized:* one (logically shared) physical infrastructure
- *programmable*: platform for innovation
- *federated:* organic growth, skin-in-the-game business model

Observations (1):

- accessible to different researcher communities at different levels in architecture
  - *IaaS:* infrastructure as a service, down to bare machine
  - *PaaS:* experimental platforms (e.g., end-end networked cloud platform) as a service
  - *SaaS:* application software (SaaS)
- [Chase, Kreiger, Landweber, McKeown, Ricci] building bottom up, [Fox] top-down: *converging* to similar place
  - architectural, control differences
- importance of clear, consistent architecture of testbed design, control, management
- open software: OpenFlow, OpenStack

Observations (2):

- edge networks:
  - WiMax, mostly via existing GENI sites (wireless ubiquity a challenge)
  - measurement of wireless, cable access nets [Feamster]
- limited input from:
  - cyberphysical systems: [Landweber] only
  - security
  - optical (some)
- sustainable business models often addressed:
  - NSF, campus co-investment, working with industry
  - investment timescales: [Landweber, McKeown, Ricci]
- interaction with industry

Observations (3): other visions

- education value noted in several white papers
- a couple of other, more tightly focused whitepapers:
  - edge network: measurement observatory
  - BGP routing

MIC white papers: summary

- valuable, thoughtful input reflecting deep experience, articulating midscale infrastructure value
- multiply-articulated MI vision: *nationwide, multi-tiered system .. sliced, deeply programmable, virtualized, and federated*
  - many common views on how to get there, but some differences as well (architecture, control, management)
- next steps: what's valuable to NSF?
  - broadening community input (CPS, security)
  - sustainability, review & evaluation processes
  - control/architecture/management approaches
  - whitepapers ideas out to broad audience?

Charge 2: Articulate framework for understanding infrastructure value

- 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

Breaking down the charge

- explaining to community why CISE should fund research infrastructure (RI)
- non-quantitative criteria
- quantitative criteria

} applied in evaluating specific proposals or ongoing assessment of RI

Explaining why CISE should fund midscale RI

- *educate:* use concrete examples of past/current RI to illustrate benefits on investment in RI: Planetlab, Emulab, Orbit, Proble, Internet2, GENI and earlier examples (Berkely UNIX)
  - capture surprises, lessons learned, as well as timescale, funding scale
  - most unaware of CISE RI spending relative to other directorates
  - OCI examples less visible to CISE: optical infrastructure, FutureGrid
- stress "virtuous" cycle"* between facilities and experimental systems research
  - building/deploying real system reveals next set of research challenges
  - researchers wrestle with richer space of cometing design goals
  - example: PlanetLab, GENI – advances in virtualization trust models, federated management, auto config, measurement

*L. Peterson, V. Pai, "Experience-driven experimental systems research," Nov. 2007 *CACM.*

Non-quantitative criteria

- *midscale infrastructure:* enables research not possible at smaller scale, needing access to underlying infrastructure
  - e.g., structure or efficiency of RI operating at scale beyond a few racks
- *train CS systems researchers* (academic or industrial R&D careers): exposure to design/operational/measurement challenges
- *better paths to practice:* results of CS systems research more "ingestible" by industry
- proposals should address how research done on midscale RI
  - could transition to practice
  - could scale up
  - is efficiently, effectively performed by experimenter (usability), particularly for experiments at large scale

Non-quantitative criteria (more)

- effects on national competitiveness: training systems-facile researchers
- Uniqueness: what's new and unique? How is facility different from existing facilities?
- technology transfer vector: software, hardware, best practices?
- technology evolution/change: how will RI cope with technology change?
- Security: RI suitable for security research?
  - do no harm: can RI be used to amplify security threat?
- midscale RI shareable through virtualization/SD*, partitioning, dynamic provisioning?
- broad use: can midscale RI be linked to CPS and/or XSEDE (supercomputing) facilities?

Quantitative criteria

IMPACT METRICS

- # papers published using RI
- # software artifacts based on use of RI
- # patents filed/granted based on work done in RI
- # proposals (funded and unfunded) that include RI use
- # companies started based on work done in RI
- #classes that use the RI

USE/SUBSCRIPTON METRICS

- # users, geographic diversity
- # experiments/projects (e.g., as in Emulab)
- # user- hours for research or teaching (e.g., as in PlanetLab)
- # CPU-hours (e.g., as in Condor)
- # bytes transmitted and/or stored
- # institutions, geographic diversity
- resource utilization

Quantitative criteria (more)

SCALE METRICS

- # CPU cores
- storage capacity (RAM, disk)
- bandwidth and/or switching capacity
- # general (e.g., rack-mount servers) and specialized (e.g., NetFPGAs) devices
- geographic diversity of deployment

COST METRICS

- Initial equipment costs (including deployment, configuration)
- Initial software development associated with tools required to use, operate the RI
- On-going costs associated with equipment upkeep, maintenance
- Day-to-day infrastructure operation costs
- User support costs (including maintaining RI software); user training.
- University overhead and/or support

Quantitative criteria: commentary

- relative importance of criteria will change over time
  - cost, scale metrics important at proposal time
  - impact not significant at start, but critical at renewal or latter assessment times
  - use metric critical throughput, but particularly at start
- metrics valuable for educating community about impacts and costs of existing experimental facilities

Charge 4: Midscale Research Infrastructure

- *white papers mirror many earlier MIC RI discussion themes*
- *virtualization*: key RI technology (computation, networking)
  - Critical: control framework for allocation, access, identity, authentication
- *converging interests in "cloud"*
  - *data centers:* interesting, missing piece of RI story
- *industry: key partner* interested in "readily ingestible" research contributions
- importance of OAM, end user involvement
  - quality architecture, human infrastructure
- how "deep," how far does RI go (bare machine, wireless, optical, CPS)?
- how to "rightsize" MI?

MIC next steps

- all charged MIC subcommittees have all reported out
- new/next tasks: what's valuable to NSF?
  - MI project management/review?
  - broadening reach (e.g., CPS)?
  - report documenting findings, recommendations?

Questions and Discussion?

[backup slide]

Target Community?

- target community: system researchers, but …
  - "…support the larger scientific community, especially those forward-looking scientists interested in transitioning their codes from local clusters and GRID-based resources to the commodity cloud, and/or working with Computer Scientists to create new cloud-inspired models for scientific computing."
  - "…create, deploy and evaluate hardware, middleware, software and systems to implement, test and validate new ideas and methods, possibly in the context of new scientific experiments and at-scale important applications."

[backup slide]