Report of the
NSF CISE AC Midscale Infrastructure Committee:
A Working Group of the Computer and Information Science and Engineering Advisory Committee

Committee Members:

P. Barford, U. Wisconsin,
F. Berman*, RPI
S. Corbató, 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 Advisory Committee Member

May 2014

1. Introduction

The Mid-scale Infrastructure Committee (MIC) of the National Science Foundation (NSF) Computer and Information Science and Engineering (CISE) Advisory Committee (AC) was formed at the May 2012 CISE-AC meeting to represent the CISE community's interests in mid-scale research infrastructure investments, to help CISE assess how well its current mid-scale infrastructure investment meets the needs of the community, and advise CISE on future directions in similar investments. It was charged to address four issues:

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

The MIC's full charge can be found in Appendix A of this report.

The committee held one all-hands meeting in October 2012, and numerous smaller meetings and conference calls in 2012 and 2013. The MIC reported findings and recommendations regarding the four charge items at the CISE-AC meetings in November 2012 and May 2013, and again summarized those findings at the January 2014 CISE-AC conference call. Those presentations, as well as this report, can be found at http://gaia.cs.umass.edu/MIC.

NSF defines midscale infrastructure as "infrastructure investments that are larger than those supported by the Major Research Instrumentation (MRI) program (capped at $4M) and smaller than those supported by the Major Research Equipment and Facilities Construction (MREFC) account (typically $100M or more)." The overall landscape for mid-scale infrastructure (MI) is shown in Figure 1.

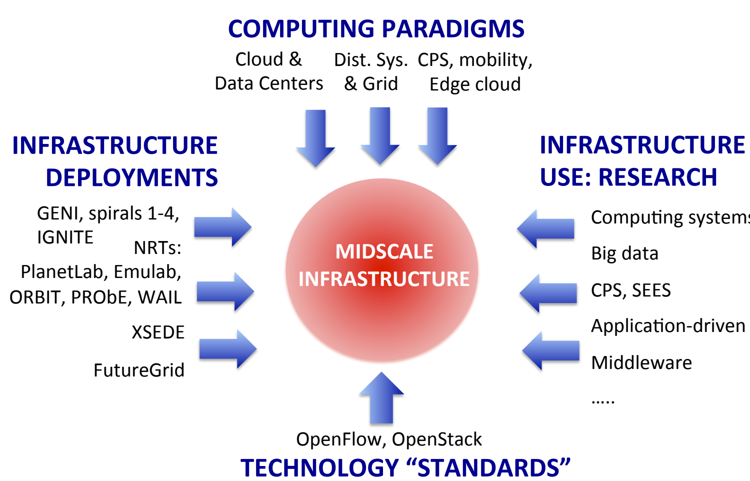

Figure 1: Landscape for mid-scale research infrastructure

Over the past decade, there have been a number of research infrastructure (RI) deployments in both the systems and grid computing communities that will inform future RI and MI deployments. Network research testbeds such as PlanetLab, Emulab, ORBIT, PROBE and WAIL – several of which have their roots in the CISE Network Research Testbeds Program (see NSF Solicitation 03-508 and [NRT 2002]) – and FutureGrid from the grid community are among the notable research infrastructure testbeds. GENI (Global Environment for Network Innovations) and XSEDE (Extreme Science and Engineering Discovery Environment) are larger-scale deployed research systems in the systems and grid communities, respectively. During this time, open standards for network (OpenFlow) and end-systems (OpenStack) virtualization together with experimental control frameworks (e.g., as developed in GENI), have made it possible for researchers to create end-end slices of virtualized infrastructure for their experimental systems research.

Computing paradigms have also evolved and matured over this period. Cloud computing has emerged as a central computing paradigm, with large-scale data centers as well as "edge clouds" supporting mobility and computing close to the network edge. Cyber-infrastructure has also played an increasingly important role as a research tool for discovery of new fundamental knowledge across broad areas of science, engineering, social sciences, and more [NSF Sustainable 2010].

Within CISE areas, experimental research infrastructure has played a central role in advancing research across all areas, from network and computing systems research to cyber-physical systems, to middleware, to application areas including Science, Engineering and Education for Sustainability (SEES). Indeed, the scope and importance of experimental research infrastructure is so broad and deep that the MIC often struggled with the definitional scope of midscale infrastructure noted in its charge. Ultimately, after much internal discussion, and discussion with NSF staff, the MIC scope was taken to be mid-scale research infrastructure aimed primarily at CISE research in the areas of cloud/network/grid systems as a first priority, but cognizant of the ultimate use of such systems at the application layer.

The MIC formed four subcommittees – one to deliberate on each of the four charges above and then report back to the larger committee for discussion. The following sections describes the activities, findings and recommendations for each of these four charges:

- Section 2.1 identifies different ways in which community infrastructure requirements might be derived, and ***describes the outcomes of an extremely successful white paper process conducted to elicit requirements and recommendations from the broader community.***
- Section 2.2 recommends that ***a paper/report documenting research advances resulting from the use of experimental infrastructure*** could play a valuable role in informing the larger research community of the importance and impact of research infrastructure. ***Qualitative criteria and quantitative metrics for assessing the value and impact of research infrastructure*** are also identified and discussed.
- Section 2.3 discusses funding and management models for mid-scale infrastructure, ***addressing issues of testbed timescale, the development of physical and human capital, and funding models that incent campus co-investment and create/leverage public-private partnerships.***
- Section 2.4 broadly ***discusses virtualized research infrastructure and the importance of leveraging existing GENI resources and experience, synthesizing***

the valuable lessons learned from this on-going, important and unique mid-scale infrastructure.

2.1 Community Infrastructure Requirements: CCC White Paper Process (S. Corbató, E. Lazowska, B. Maggs, D. Raychaudhuri)

The MIC sub-committee on Community Infrastructure Requirements discussed the question of how mid-scale RI requirements should be derived within the CISE research community and, in the spring of 2013, worked with the CCC to conduct a white paper process aimed at eliciting such requirements.

Various mechanisms for obtaining community input were discussed:

- a decadal study (similar to those done for Astronomy), which was considered too "heavyweight" and too long-time scale;
- relying on CISE-AC deliberations alone (led by the MIC), which was considered to not draw enough community input;
- commissioning a National Academy of Sciences study, which was considered to be a bit heavyweight, given that similar input might be drawn from an NSF-sponsored workshop;
- an NSF-sponsored workshop, perhaps preceded by white paper submissions, which has been successfully used many times within CISE to elicit community input, including for network research testbeds [NRT 2002];
- a white paper process, which could more quickly elicit community requirements.

The last two approaches were ultimately judged to be the best options. The subcommittee then moved to work with the CCC to undertake such a white paper process.

An Open Call for White Papers on CISE midscale infrastructure needs was issued on March 20, 2013 through the CCC (Computing Community Consortium) website and blog. The Call is shown in Appendix B.

Ten white papers were received in response to the Call:

- Cappos *(NYU Poly),* "Three Computing Infrastructure Needs"
- Chase, Baldine, *(Duke, RENCI / UNC)*, "CCC Mid-Scale Response"
- Feamster *(GaTech),* Banerjee *(Wisconsin)*, "An Open Observatory for the Internet's Last Mile"
- Fox, von Laszewski (Indiana), Fortes (Florida), "Mid-Scale Infrastructure Investments for Computing Research: A FutureGrid Perspective"
- Katz-Bassett *(USC),* Levin *(Maryland),* Zariffs *(USC)*, Feamster *(GaTech)*, "The Transit Portal: A Testbed for Internet-scale Routing"
- Krieger, Bestavros, Appavoo *(Boston U.),* "Enabling an Open Cloud"

- Landweber *(Wisconsin),* Elliott *(BBN),* "Mid-Scale Infrastructure Investments for Computing Research"
- McKeown, Parulkar, Peterson, Shenker, *(Open Network Lab),* "NSF OpenCloud"
- Ricci *(Utah)*, "The Need for Flexible Mid-scale Computing Infrastructure"
- Weiss *(Evergreen State),* Mache *(Lewis & Clark),* "Mid-scale infrastructure for experimentation in networking and security"

The white papers submitted span a range of research infrastructure areas including an Internet last mile observatory, a future smart grid testbed, an Internet-scale routing testbed, open cloud computing laboratory, network security testbed, mobile edge networks. The submitters included the designers, developers, and operators of many of the most important research infrastructure testbeds in our community, including GENI, PlanetLab, Orbit, FutureGrid, and Emulab. The white papers were read by MIC members and NSF staff, who were deeply impressed by the depth of thought, clarity of vision, and insightful "lessons learned" through collective decades of experience in research infrastructure development and operation. The readers were also struck by the *maturity* of the white papers. Our community has clearly learned much from previous testbed efforts and is much wiser and pragmatic, while still bold in vision. The community has been enriched by this experience, and now has research leaders with a critical mass of expertise and experience in this area. Indeed, the white papers were collectively so impressive, and deemed of such value to the community, that we hope to publish a summary of the experiences, vision, and recommendations outlined in these papers at a future date (see next steps, below).

All contributors emphasized the importance of increased infrastructure investment for computer science research. Although the white papers received spanned a wide range of research interests and technology domains, a common vision was articulated in many papers (excerpted here from one submission):

> *"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."*

Interestingly, this common vision was articulated by white paper authors with roots in the networking community (who have worked over time towards this vision "bottom up") as well as authors with roots in the grid community (who have worked over time towards this vision "top down"). Examining this vision a bit more deeply, we find the following characteristics:

- a multi-tiered system, with national/regional R&E backbones and core/edge networks, interconnecting data centers and campus-based clouds. Several papers suggested a relatively small number of larger-scale data centers connected to a larger number of smaller campus-based cloud infrastructure;
- a sliced, and virtualized system, allowing one (logically shared) physical infrastructure;
- a programmable system, providing a platform for innovation;

- a federated system, allowing for organic growth, and a skin-in-the-game business model (see section 2.3);
- a system accessible to different researcher communities at different levels in architecture (see also discussion in section 2.3 regarding different categories of infrastructure use):
  - IaaS: infrastructure as a service, down to bare machine
  - PaaS: experimental platforms (e.g., end-end networked experimental cloud platform) as a service
  - SaaS: application software as a service. Here, the goal would be to support application developers willing to work closely with PaaS and IaaS developers, as opposed to providing "cycles as a service".

Some notable observations in the white papers, borne from years of experience, were:

- The importance of a clear, consistent architecture of testbed design, control, and management.
- The need for the testbed to be based on open community-supported software standards such as OpenFlow and OpenStack.
- The need for realistic edge networks, particularly for use in mobile experiments, where ubiquity of access (here, e.g., via LTE) is important.
- The importance of sustainable business models for the experimental facility was pointed out by several respondents, calling for co-investment by the NSF, campuses and industry (see section 2.3). Several authors also noted the need for different types of investments over the different phases of a testbed's lifetime.
- The value of interaction with industry, in order to avoid obsolescence, to track technology trends in a timely manner, and to provide a conduit for research results into practice.
- The educational value for hands-on courses and research experience was noted in several white papers.
- Somewhat surprisingly, limited input was received on the topics of cyber-physical systems, security, and optical networks, which the MIC judged to be quite important.

In summary, the white papers provided thoughtful input reflecting deep experience and articulating the value of midscale infrastructure to the CISE community. Multiple respondents articulated the need for a nationwide infrastructure for large-scale networking and cloud computing experiments that would enable the academic research community to contribute towards advancements in data centers, cloud computing technologies, mobile services and the future Internet – areas of national importance from both competitiveness and human capital points of view. MI represents experimental infrastructure than no small collection of researchers can do alone, and thus requires national-level leadership and scope.

The common infrastructure vision argues for a multi-tiered system that is sliced, deeply programmable, virtualized and federated. There are many common views on how to get there, with emerging software-defined network technology, open stack standards, and virtual networks playing a central role. With this common vision and common technology base, there were differences as well in terms of the details of testbed architecture, control and management.

Next Steps. The MIC hopes that the white papers have informed the development of NSF Solicitation 13-602 and will continue to inform any future resulting activity. The papers themselves were so impressive that the MIC hopes to assemble authors to synthesize a paper based on selected white papers, and published in the open literature for the benefit of the technical community.

2.2 A framework for understanding the value of novel infrastructure to transformational research for the CISE community (P. Barford, J. Fortes, J. Kurose, B. Lyles, K. Marzullo, J. Mogul, J. Rexford)

This MIC sub-committee addressed the second charge item: "How can CISE determine a framework for understanding the value of novel infrastructure to transformational research for the CISE community?" The first subsection below addresses the challenge of articulating the value of infrastructure to the community; the next two subsections discuss qualitative and quantitative criteria that might be applied by NSF in evaluating any specific proposals for midscale infrastructure, and in on-going assessments.

2.2.1 Articulating the value and importance of infrastructure

Some members of the CISE research community may have a healthy skepticism about the value of investing in experimental infrastructure. CISE could provide valuable historical perspective, using concrete examples, to illustrate the benefits of government investment in infrastructure. Some examples could include recent testbeds within the systems research community such as PlanetLab, Emulab, Orbit, Probe, Internet2, and GENI, as well as prominent earlier examples such as Berkeley UNIX; examples from the grid community could include high-speed optical infrastructure, CCNIE, and FutureGrid. Such a perspective, which could consider both MI and MRI-scale infrastructure, would be broadly read if published in a visible venue such as *Communications on the ACM*.

This historical perspective could capture what surprises came out of creating and using the experimental facilities, as well as how long these platforms took to produce valuable results and how much the platforms cost to build, maintain, and operate. These past projects can also offer "lessons learned" about what could have been done differently, including whether the impact might have been larger with additional funding. The whitepapers discussed in Section 2.1 might provide a valuable starting point for such an effort.

Another way to reach the community is to stress the "virtuous cycle" between experimental facilities and systems research [Peterson 2007], in which the process of building and deploying a real system helps reveal the next set of research challenges, and forces researchers to grapple with a richer, multi-dimensional space of competing design goals. In addition, creating and managing an experimental infrastructure often raises new research questions. For example, designing and building PlanetLab and GENI led to new ideas for server and network virtualization, techniques for managing multiple types of resources, distributed trust models, federated management, and automated configuration and measurement – topics that arguably are just as important as the experiments that run on these platforms.

Mid-scale IT systems are the "new normal" for today's IT industry. It is urgent for research and education to create the science and engineering practices for creating and managing these systems, and training a new generation of computer scientists and engineers. Having researchers build and manage these kinds of facilities – and publish what they learn – is an important complement to industry efforts driven by immediate (and often proprietary) business concerns.

The community should also be part of a larger discussion regarding the appropriate size and scope of experimental facilities. A facility that is "too small" cannot support larger experiments that evaluate new protocols at scale, or strive to attract real users by offering good performance and reliability. That said, a facility that is "too large" is expensive to build, maintain, and manage, and draws resources away that could go to multiple smaller projects. Larger mid-scale infrastructure, however, has tended to be shared among the community, which both amortizes cost and creates a larger, more stable community of expertise. Similarly, a facility mixing multiple kinds of technology (e.g., wireless and wired networks, optical networking and routers) enables novel research that revisits the traditional boundaries between protocol layers. Yet, building separate facilities for different technologies or protocols layers can help reduce risk. These questions have no easy answers, but engaging the community in meaningful discussion and debate on these topics can help build confidence that CISE is identifying good ways to balance these competing trade-offs, and help guide making the hard decisions about what kinds of infrastructure to support.

2.2.2 Making the case for infrastructure: non-quantitative criteria

The primary reason to fund midscale infrastructure is to enable research experiments that are otherwise impossible to perform. The community of researchers and practitioners clearly understand that scale matters, and that certain scale-related challenges only appear when running in a real (rather than simulated/emulated) environment. Much CISE-funded research may be conducted using small-scale infrastructure or pooled resources (such as PlanetLab or public clouds). Midscale infrastructure research requires larger-scale experimental infrastructure, often with access to the underpinnings of the infrastructure itself. For example, an MI research project might address the structure or efficiency of a computational infrastructure operating at a scale beyond that which could fit into a small number of racks, but at smaller scale than that of the largest commercial systems (e.g., Google, Amazon, Facebook, Microsoft). There is much valuable research that can be done at the mid-scale level. MI is also important from a national competitiveness point of view. Without MI, academic research and training would be constrained to relatively small-scale systems, while much of practice moves to large-scale system.

Beyond the research itself, there are numerous other benefits to be gained from MI research:

- **Human capital.** Computer systems researchers will be trained by exposing them to the design, operational, and measurement issues needed for careers in academic research or industry R&D.
- **Pathways to practice.** The results of CISE-funded systems research are more easily ingestible by industry, startups, and government practitioners when demonstrated in systems working at realistic scales. MI activities should articulate possible pathways to practice. For examples, is the infrastructure design compatible with some aspects of standard practice, or does it assume a clean-slate approach?

In evaluating MI activities, the following considerations are also important:

- **Sharing and multi-investigator use.** Given the level of investment needed, MI activities should support shared and/or multipurpose use through software definition, partitioning, dynamic provisioning or other means. This will contribute to sustainability and flexibility as research needs evolve over time.
- **Scaling.** How might the MI be scaled up in the future, or replicated by other researchers? That is, CISE may prefer proposals that can be used as building blocks for future, larger infrastructures, or for parallel deployments.
- **Usability.** Can the MI be efficiently and effectively used by experimenters? How easy is it to learn how to use a research infrastructure and run experiments at much larger scale than many researchers have previously experienced? Proposals could include a discussion of how to evaluate usability.
- **Uniqueness.** In what way does a specific MI activity offer a unique facility, rather than duplicating what is already available? Why would the goals of an MI activity not be met by using a shared "public cloud"?
- **Technology refresh and sustainability.** How can an MI activity cope with technology change and turnover, and avoid obsolescence? What is the sustainability model for an MI activity (see section 2.3)?
- **Data Sets.** Can the MI generate and/or house useful data sets? Can it produce valuable data on service and user behavior?

2.2.3 Making the case for infrastructure: quantitative criteria

One component for assessing the value of experimental infrastructure is to identify metrics that capture a variety of meaningful characteristics. Metrics are standards of measure that have a number of benefits. First, metrics are less vulnerable to subjective interpretation than qualitative assessments (e.g., users "like" the infrastructure). Metrics also enable relative comparisons between different infrastructures. Finally, some metrics lend themselves to automated collection, which simplifies on-going assessment of infrastructures.

The first step is to identify a set of metrics that best captures meaningful measures of costs, scale, use, and impact. While many metrics are possible, an initial list of potential metrics include those that have been used by existing infrastructures such as Condor, Emulab, and PlanetLab:

Quantitative Impact Metrics: number of

- papers based on work done in the infrastructure
- software artifacts based on use of the infrastructure
- patents filed/granted based on work done in the infrastructure
- proposals (both funded and unfunded) that include the use of the infrastructure
- companies that are started based on work done in the infrastructure
- classes that use the infrastructure

Quantitative Use/Subscription Metrics:

- Number of users and their geographic diversity (both for researcher-users and users of services supported by the infrastructure), including return users who use the infrastructure for more than one project
- Number of experiments/projects (e.g., as in Emulab)
- Number of user-hours for research or teaching (e.g., as in PlanetLab)
- Number of CPU-hours (e.g., as in Condor)
- Number of bytes transmitted and/or stored
- Number of institutions and their geographic diversity
- Resource utilization
- Distribution of usage durations (how long does a user/project run on the infrastructure)

Scale Metrics:

- Number of CPU cores
- Storage capacity (RAM, disk)
- Bandwidth and/or switching capacity
- Number of general (e.g., rack-mount servers) and specialized (e.g., NetFPGAs) devices

Cost Metrics:

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

Important to the effective use of metrics for assessing the value of infrastructures is to understand their relative importance, particularly within the context of a project's lifetime. Cost and scale are critical metrics for assessing feasibility prior to funding. Use metrics are critical throughout the lifetime of a project, but particularly in the early days, when impact is less clear. Finally, impact is perhaps the most important metric after the infrastructure has been established and in use for some time.

In addition to evaluating proposed infrastructures, these metrics are valuable for educating the community about the costs, impacts and importance of existing experimental facilities.

2.3 Funding and management models (S. Corbató, F. Berman, J. Fortes, J. Kurose)

The MIC subcommittee on funding and management models focused in particular on the economic and inter-sector aspects of information infrastructure, described by the questions below.

- **What are the best models of funding community mid-sized ($4M-$100M) infrastructure?**
- **How can CISE involve industry in mid-scale infrastructure?**

MIC interpreted *infrastructure* broadly, including computation, networking, digital data, software, and human support within its discussions. The key theme of our three main recommendations is to *increase the effectiveness of mid-size infrastructure investments*, making the most out of NSF efforts. The recommendations are briefly discussed in the sub-sections below.

Note that the MIC recognized that infrastructure is most effective when funded differently than research. In particular, infrastructure should be assessed based on its broad impact (e.g., via qualitative and quantitative metrics, as discussed in Section 2.2), and should be funded so that it can support the user community without interruption and over a substantial period of time. Although MIC focused on the economic and inter-sector aspects of information infrastructure, it was clear that there are larger issues that must be addressed over the long term to maximize the benefit and effectiveness of NSF infrastructure. These issues include:

- *Roughly what proportion of the NSF budget should be spent on a) production infrastructure that broadly enables research, b) research enabled by production infrastructure, and c) infrastructure research -- prototypes of infrastructure that expand the functionality and capability of production infrastructure?* It is important to note that each of these three categories represents a distinct use of "infrastructure," with very different needs and operating modes. It is unlikely that a single facility would simultaneously serve all three categories, unless very coarsely partitioned.

- *What is the right mix of compute infrastructure, networking infrastructure, data infrastructure, software systems, and human support within the infrastructure budget?* This mix will vary with the type (e.g., among the three categories above) and use of midscale infrastructure. Our subcommittee noted, however, that human support is often the most critical, yet under-budgeted and underappreciated (at least initially) – a finding echoed in earlier reports [NRT 2002]. Our full MIC noted that CISE's infrastructure investment of roughly 4% of its (pre-ACI) budget (see page 6 of May 2013 CISE-AC briefing) was lower than (or in the case of Engineering, comparable to) other NSF directorates. No one felt this level of investment was too high.

- *What parts of the research infrastructure landscape should NSF be responsible for and what parts should/can NSF rely on from campuses, states, other agencies, other sectors*? As evidenced in our discussion below, we recommend broad, long-term partnerships rather than a strict partitioning of responsibility.

Continuous attention to these broad, long-term issues will be critical in optimizing NSF investments. While definitive answers to these issues were beyond the scope of the MIC, but the issues themselves were always present in our discussions.

Recommendation 2.3.1: Build longer-term sustainable facilities at the mid-scale that more strongly link NSF infrastructure investments with science outcomes, impact.

The focus of this recommendation is to increase the potential of NSF infrastructure to enable new research and discovery.

- ***Build longer-term sustainable facilities at the mid-scale.*** There is a considerable ramp-up period for infrastructure development and adoption/use by the research community; infrastructure lifetime should be carefully considered to fully realize infrastructure investment and impact. As discussed in section 2.2, measure impact through outcomes and science. Prioritize usefulness of infrastructure for the community.
- ***Establish periodic discipline-wide assessment within research areas to drive new facility design and funding priorities.*** Draw on broad community participation and input. Include and address successes and lessons learned from previously supported facilities; see also our recommendations in Section 2.1.

Recommendation 2.3.2: Create shared business models between NSF and the community

The focus of this recommendation is to strengthen partnerships between NSF and other entities to support enabling infrastructure that can advance the research community. Shared funding responsibilities are key for productive partnerships and deep engagement. Although the business models will vary by sector, community, etc., each provides an opportunity for the NSF to extend its infrastructure reach and expand its research impact.

- ***Develop funding models to incent campus co-investment in shared-use research facilities***. Expect longer (2X?) campus commitment (e.g., through campus matching funds) beyond typical MRI/CRI award duration with funding transition as part of the project plan. All plans should incorporate realistic operational expenses. Key stakeholders should be expected to participate – researchers, CIO, VPR, service providers, and more.

- ***Explore the development of community-led non-profits as vehicles to sustain large mid-sized projects and focus institutional and partner co-investments.*** Explore the possibility of using a large science project management approach with appropriate checks and balances, stakeholder advisory groups. Examples of this include UCAR / Atmospheric science, AURA / astronomy, IRIS / seismology, NSF MREFC facilities.
- ***Create and leverage public-private partnerships.*** Partner with the for-profit and non-profit private sector to utilize private sector facilities (e.g. data repositories, clouds, computers) as platforms for NSF-funded research. Develop joint programs with the private sector to create and provide infrastructure for the NSF research community at community venues (e.g. campuses, libraries, centers).
- ***Co-sponsor enabling infrastructure across NSF directorates and with other R&D agencies.*** Create cross-NSF facilities for computation, data stewardship, etc. Partner with DOE, NIH, NIST, NASA, etc. to develop, maintain and provide compute, data, software, networking and other kinds of information infrastructure for the research community.

Recommendation 2.3.3: Optimize mid-sized infrastructure investments

The focus of this recommendation is to ensure that the NSF community is getting the most out of NSF infrastructure investments. This means that NSF investments should strategically leverage existing infrastructure, maximize the lifetime of useful infrastructure, and provide a spectrum of infrastructure options.

- ***Ensure that facilities support a broad spectrum and scale of meritorious research projects.*** Assess facilities based on quantifiable outcomes (see section 2.2). Use metrics of success that link assessment, innovation, and sustainability.
- ***Explore the use of commercial cloud services, when appropriate.*** Consider a variety of platforms including those developed by NSF grantees and those provided by commercial organizations. The MIC noted that current research overhead rules (exempting infrastructure purchases but not infrastructure service purchases) dis-incents the use of commercial large-scale cloud services in favor of traditional equipment and software purchases. The incentivization implicit in current indirect cost rules should be addressed.
- ***Require viable sustainability plans beyond the duration of NSF funding.*** Provide at least a 3-5 year horizon for "ramp in" funding and expect cost-sharing / partner investment / in-kind contributions after grants expire.

2.4 Leveraging GENI and Beyond (P. Barford, S. Corbató, J. Kurose, K. Marzullo, B. Lyles, D. Raychaudhuri)

This subcommittee began its discussions with briefings and reviews of current experimental research testbeds, including GENI, making the following observations:

- **Virtualization.** Over the past five years, virtualization – both in the network (e.g., OpenFlow) and in the end systems (e.g., OpenStack) – has become a key technology-enabler for experimental research infrastructure (RI) testbeds. Critical to exploiting these technologies is the creation of a control framework for resource allocation and access, identity, and authentication.
- **Converging interests in the cloud.** The committee noted three separate constituencies with cloud-based experimental research interests: *(i)* systems researchers interested in performing experimental research in the underlying operating system, networking, and control technologies; *(ii)* researchers interested in the development of middleware making the underlying infrastructure more easily and efficiently usable in myriad application domains, *(iii)* application users interested in "cycles". *Despite these convergent interests in cloud infrastructure, these communities have been historically distinct (indeed, with too little cross-community discussion),* with the former set of researchers coming more from the CISE/CNS community and the latter set coming from the CISE/ACI community. As noted in Section 1, the MIC's scope was primarily aimed at mid-scale research infrastructure for first of these constituencies, but cognizant of the ultimate use of such systems at the application layer.

  Recent GENI activities in the CISE/CNS community have evolved towards virtualized cloud infrastructure (e.g., GENI racks with OpenFlow andL2 VLANs, InstaGENI, ExoGENI) – an investment to be leveraged and built upon in the future. In the CISE/ACI community, FutureGrid has a significant cloud component, and an XSEDE cloud survey was performed in 2013.
- **Potential missing pieces.** While some applications of virtualization in research infrastructure have been underway for some time and are maturing, others areas are still evolving. Data-center-scale and edge-cloud computing (e.g., supporting cyber-physical systems, mobility, wireless) are two such evolving areas.

Mid-scale Infrastructure: how "deep" and how "big"

The subcommittee also discussed the question of "how deep" (i.e., how close to the underlying computing and networking hardware) the provided research infrastructure should go. Should the base unit available to researchers be the bare machine, a virtual machine, or a choice of either? With choice, the research infrastructure can be more things to more people, but then also risks the many perils that accompany increased complexity. In the networking domain, wired L2 services is a common baseline in GENI, with some activities in wireless (WiMaX, to date) and less in optical. Mixing network technologies enables novel research across boundaries but again introduces risks due to increased complexity. *Understanding and managing the risk-versus-complexity (i.e., testbed scope) tradeoff will be critical for any MI activity going forward.*

A question closely related to that of testbed scope is that of testbed size. MI activities must also consider how "big" the testbed should be before the activity sees diminishing returns with increasing size for the target use cases. Past experience has taught the community that it is preferable to start small (but always working with users from early on), iterate and grow, rather than aim for a single large-scale launch. Federation provides a mechanism to grow incrementally, as well as providing partial resource "buy in" (see section 2.3) that leverages central investment.

The above questions of RI depth and size both speak to the need for a clear vision/statement of the research enabled, as well as research precluded, by specific RI. Specific success criteria – both qualitative and quantitative as discussed in Section 2.2 – would naturally be considered here.

Mid-scale Infrastructure: more than just iron, wires, and code

Our subcommittee noted that when discussing "research infrastructure," the hardware (and to some extent the software) often tend to dominate the discussion. Experience (e.g., as reflected in the white papers discussed in Section 2.1 and our own experience), however, has shown that RI success ultimately hinges on many other factors as well.

- **Architecture.** It is critical that the overall architecture of the research infrastructure fit its use. Decades of systems research has shown that the right overall architecture obviates or eases numerous challenges in the future; conversely, an ill-fitting architecture can complicate and confound even the most talented engineers building research infrastructure that instantiates an ill-fitting architecture. This suggests when designing a testbed, one should invest a great deal of time up front, getting the best architect and resulting architecture possible.
- **Importance of operations/administration/management (OAM), end-user interaction.** While testbed developers/operators as well as funding agencies all acknowledge the importance of OAM and end-user interaction, the committee's sense is that this critical activity is often underfunded. Indeed, OPEX costs can dominate, increasing as infrastructure becomes more "production quality." This observation echoes that of a network research testbed report a decade ago [NRT 2002], suggesting this is a lesson hard-learned. Equally important is the human infrastructure – the interaction between RI providers and users. Our sense is that GENI has done an outstanding job of building and engaging a broad community of users, through events such as the GENI Engineering Conferences and the GENI education workshop.

GENI

GENI has been a unique and remarkable project for the research community, of unprecedented size and scope. Countless researchers have benefitted from its realized vision of large-scale, instrumented and deeply programmable experimental research infrastructure; it is playing a key role in the US IGNITE program.

There are certainly physical GENI resources that can be leveraged moving forward. Indeed, a number of on-going GENI activities are built around the vision discussed in Section 2.2 of an end-end sliced, virtualized, and deeply programmable cloud-like computing and networking research infrastructure.

More important than the physical assets, however, are the hard-won experiences, insights, and lessons learned by the GENI leadership and management team. This invaluable resource must be maximally leveraged going forward.

References

[NRT 2002] J. Kurose *et al.*, Report of NSF *Workshop* on *Network Research Testbeds*, November 2002, http://*gaia.cs.umass.edu*/testbed_workshop.

[NSF Sustainable 2010] NSF-Sponsored Workshop Report: Sustainable Funding and Business Models for Academic Cyberinfrastructure Facilities, November 2010. http://www.cac.cornell.edu/SRCC/

[Peterson 2007] Larry Peterson and Vivek S. Pai. 2007. Experience-driven experimental systems research. *Commun. ACM* 50, 11 (November 2007), 38-44. DOI=10.1145/1297797.1297820 http://doi.acm.org/10.1145/1297797.1297820

APPENDIX A: MIC CHARGE

The CISE **M**idscale **I**nfrastructure **C**ommittee (MIC), which is a subcommittee of the CISE Advisory Committee, will represent the community's interests in mid-scale research infrastructure investments.

NSF defines *midscale infrastructure* as infrastructure investments that are larger than those supported by the Major Research Instrumentation (MRI) program (capped at $4M) and smaller than those supported by the Major Research Equipment and Facilities Construction (MREFC) account (typically $100M or more). Programs like the NSF-wide Major Research Instrumentation Program (MRI) and the CISE Computing Research Infrastructure Program (CRI) support the creation, enhancement, and operation of world-class computing research infrastructure, but the size of these projects are limited (MRI projects are capped at $4M and CRI at $3M).

CISE has its own Computing Research Infrastructure program (CRI), but these are capped at $3M. CISE has funded a few midscale infrastructure projects, including projects such as PlanetLab, Emulab, ORBIT and most recently, GENI and PRoBE.

MIC's charge is to help CISE assess how well its current mid-scale infrastructure investment meets the needs of the community and advise CISE on future directions in similar investments.

MIC's first activity will be to use data provided by CISE to develop a baseline understanding of CISE's investments in mid-scale infrastructure over the years, to include, but not limited to:

- Investment areas and amounts over the years
- Comparisons across other directorates and /or relevant research areas today
- Relevant outcomes and value of those investments to the research community

In the coming year, MIC shall provide CISE with feedback and recommendations related to the following:

1. How should community infrastructure requirements be derived?
    a. What are the best mechanisms for surfacing requirements? Some mechanisms include decadal studies, academy studies, workshops, recommendations from the CISE AC.
    b. How should investments be prioritized when several areas have real needs but there is insufficient funding?
2. How can CISE determine a framework for understanding the value of novel infrastructure to transformational research for the CISE community?
    a. What is the appropriate balance between infrastructure investment to research

investment?

b. What metrics should be applied to evaluate infrastructure development and the associated research outcomes?

3. What are the best models of funding community mid-scale infrastructure? This is complicated because infrastructure is often best built through community cooperation and federation, and with a focus on interoperability, rather than as done with MRI and CRI, in which the most highly-rated stand-alone projects are funded. How can CISE involve industry in mid-scale infrastructure?

4. GENI is now in its fifth year of development and prototyping. A meso-scale infrastructure is now being deployed across the country.

    a. How should NSF decide which campuses come on board? How should these be paid for?

    b. What are the models of sustainability that NSF should consider for GENI?

    c. How can the GENI resources be leveraged as CISE moves forward?

Appendix B.

The following call for white papers was issued through the CCC (Computing Community Consortium) website and blog http://www.cccblog.org/2013/03/20/call-for-white-papers-on-mid-scale-infrastructure-investments-for-computing-research/ on March 20, 2013:

"Mid-Scale Infrastructure Investments for Computing Research
The Computing Community Consortium (CCC) is seeking community input to better understand the potential needs and payoff for additional investments in mid-scale infrastructure for computing research.

The National Science Foundation spends significantly less on shared research infrastructure for computing research than it does for many other fields. By "shared research infrastructure" we mean experimental hardware and/or software and associated instrumentation that serves a significant portion of the research community (versus a small set of investigators). In other fields, such shared research infrastructure includes equipment such as telescopes, ocean observatories, supercomputers, and field stations.

We specifically are interested in "mid-scale" infrastructure investments, defined as investments of over $4 million but under $100 million. Infrastructure investments in the $100,000 - $4 million range are accommodated by NSF's Major and Computing Research Instrumentation (MRI and CRI) programs. Infrastructure investments of $100 million or more fall under NSF's Major Research Equipment and Facilities Construction (MREFC) Program. GENI, PlanetLab, Orbit, FutureGrid, and Emulab are examples of mid-scale infrastructure investments with significant impacts on our field.

While it is possible that the current level of investment is appropriate and our field does not require additional investment in shared research infrastructure, it seems equally likely that we are underinvesting and that there are some potential investments that would significantly enhance the research capabilities of our field. We would like to solicit your opinions and ideas, in the form of short white papers. What sort of investment in mid-scale infrastructure can you envision that would drive computing research forward? What research infrastructure do you think is most appropriate (and for which broad class of research questions) and why? To what extent can such infrastructure be shared (and at what level)? How should such infrastructure be administered? How would technical standards be developed by the community? Is federation of resources possible? How should the research community organize to utilize such infrastructure most effectively and to provide input on its operation and technical evolution? How do you see the infrastructure being funded? Will the infrastructure be used for relatively short-lived activities, long-term projects, or a mix? How can the design of the infrastructure enable effective transfer of research results to the broader community, including industry and students?

White papers should not exceed 10 pages in length (shorter is better!) and are due by April 15, 2013. Depending on the level of interest generated by this call for white papers, a follow-on workshop may be held at which papers may be presented. In either case, it is our intention to approach the NSF and other funding agencies with the results of this effort."