

Suggested Reference:

*“A Process-Oriented Approach to Engineering
Cyberinfrastructure”*

F. Berman, J. Bernard, C. Pancake, L. Wu
<http://director.sdsc.edu/pubs/ENG/>

A Process-Oriented Approach to Engineering Cyberinfrastructure

**Report from the Engineering Advisory Committee
Subcommittee on Cyberinfrastructure**

***Francine Berman**, [Chair] San Diego Supercomputer Center and U.C. San Diego*

***James Bernard**, Iowa State University*

***Cherri Pancake**, Oregon State University*

***Lillian Wu**, IBM*

February, 2006

Executive Summary

In the last two decades, computational and information infrastructures have become increasingly important and ubiquitous as enablers for Engineering (and more broadly, NSF) researchers, educators, and practitioners. In 2002, NSF convened a Blue Ribbon Advisory Committee to evaluate the opportunities and challenges in developing and deploying enabling “cyber-infrastructure” for its communities. The committee’s report, commonly referred to as the “*Atkins report*,” provides a compelling vision of the opportunities for NSF’s research and education communities through the development and provision of enabling cyberinfrastructure:

The next steps to achieving the vision for Engineering described in the Atkins Report is a plan of action for understanding the engineering communities’ requirements, resources, and needs, as well as for integrating cyberinfrastructure efforts within the broader portfolio of the Engineering Directorate’s research and education efforts. In this report, we describe a process that can assist Engineering to develop an integral and integrated set of cyberinfrastructure programs that effectively support Directorate efforts in research and education, and ensure that it is responsive to a changing landscape as engineering, science, and technology advance over the next decade.

Table of Contents

1. Introduction.....	1
1.1. Charge to the Committee	1
2. A Foundation for Evolving Needs and Priorities.....	2
3. Framework for ENG’s Cyberinfrastructure Efforts.....	4
3.1. The Innovation Loop.....	4
3.2. Organizational Relationships	5
3.3. Principles for Ensuring Cyberinfrastructure Success	7
4. Funding and Program Models for Cyberinfrastructure.....	9
4.1. Structuring ENG’s Cyberinfrastructure Portfolio.....	9
4.2. An “Infrastructure” Funding/Program Model.....	10
4.3. Partnership Opportunities	11
4.4. Creating “Links” to Power the Innovation Loop	12
5. Evaluating Cyberinfrastructure Programs.....	13
5.1. Framing Questions	13
5.2. Establishing Accountability for Cyberinfrastructure Projects	14
6. Summary of Recommendations.....	17

1. Introduction

In the last two decades, computational and information infrastructures have become increasingly important and ubiquitous as enablers for NSF communities and projects. In 2002, NSF convened a Blue Ribbon Advisory Committee to evaluate the opportunities and challenges in developing and deploying enabling “cyber-infrastructure” for its communities. The committee’s report, commonly referred to as the “*Atkins report*,”¹ provides a compelling description of the opportunities for NSF’s research and education communities through the development of an enabling cyberinfrastructure:

The term *infrastructure* has been used since the 1920s to refer collectively to the roads, power grids, telephone systems, bridges, rail lines, and similar public works that are required for an industrial economy to function. ... The newer term *cyberinfrastructure* refers to infrastructure based upon distributed computer, information, and communication technology. If infrastructure is required for an *industrial* economy, then we could say that cyberinfrastructure is required for a *knowledge* economy. [p. 5]

[Cyberinfrastructure] can serve individuals, teams and organizations in ways that revolutionize *what they can do, how they do it, and who participates*. [p. 17]

Just as good public infrastructure must be use-driven, good cyberinfrastructure must be science- and engineering-driven in order to further NSF’s mission. For the Engineering Directorate, cyberinfrastructure is a core focus area, not just to enable Engineering projects, programs, and communities, but also because the design, development, and deployment of cyberinfrastructure are themselves engineering challenges.²

1.1. Charge to the Committee

In the Spring of 2005, Engineering AD John Brighton appointed an Engineering Advisory Committee Subcommittee on Cyberinfrastructure (**EAC-CI**) to provide input to the Directorate and its Cyberinfrastructure Working Group. The EAC-CI committee has worked closely with the Engineering Directorate (**ENG**) to evolve an approach that addresses the needs of the Directorate’s communities and coordinates with NSF’s new Office of Cyberinfrastructure (**OCI**) and Cyberinfrastructure Council (**CIC**). EAC-CI includes the four EAC members who authored this document. Jo Culbertson, Suvrajeet Sen, and Abhijit Deshmukh served as ENG’s liaisons to the committee.

Dr. Brighton charged EAC-CI to identify the key issues and challenges in developing ENG’s strategy for developing cyberinfrastructure to promote discovery and innovation in engineering.

¹ *Revolutionizing Science and Engineering Through Cyberinfrastructure – Report of the National Science Foundation Blue Ribbon Advisory Committee on Cyberinfrastructure*. National Science Foundation, January, 2003. Available online at www.nsf.gov/cise/sci/reports/atkins.pdf.

² In this report, the capitalized “Engineering” and ENG refer to the Engineering Directorate at NSF and the lower-case “engineering” refers to the engineering community.

With the support of Dr. Brighton and his successor, Acting AD Richard Buckius, EAC-CI covered a broad range of topics. This report presents our findings and recommendations.

2. A Foundation for Evolving Needs and Priorities

The Engineering Directorate currently has several large-scale projects involving the development of targeted cyberinfrastructure, including NCN, NEES, and most recently, CLEANER. Each of these projects requires a broad spectrum of enabling technologies, from computational simulation, to data management and archiving, to the remote operation of shared laboratory equipment. Such projects provide substantial experience with the development of targeted cyberinfrastructure in service to research and/or community goals, however the development, funding, and use of more broad-based Engineering cyberinfrastructure remains a challenge.

To achieve the vision described in the *Atkins Report*, NSF must establish ongoing processes that will make it possible to target and leverage cyberinfrastructure as it evolves over the next decade. Since the emerging cyberinfrastructure cannot be fully characterized in advance, we suggest putting in place a process that will position ENG to recognize and be responsive to a changing landscape over the next, critical decade. This, in turn, will enable ENG to focus its initial efforts and accelerate progress, both within the Directorate and Foundation-wide (through OCI and CIC).

We suggest that ENG begin by establishing a process for identifying and tracking what kinds of cyberinfrastructure tools, technologies, and components are most important for ENG research and education. The process should include regular updates as technologies evolve and cyberinfrastructure matures. This amounts to developing *user requirements* that:

- are general enough to adequately represent a broad spectrum of current engineering efforts (not just specific projects, no matter how large their scale)
- are specific enough to aid in the setting of priorities as to what components of cyberinfrastructure must be developed first
- are specific enough to suggest metrics of success for determining whether or when ENG cyberinfrastructure efforts have achieved the goal of enabling project efforts
- can evolve over time as the portfolio of ENG projects and programs evolves

The resulting **Cyberinfrastructure User Requirements** will provide a critical first step, helping to ensure that cyberinfrastructure efforts make engineering researchers and educators more competitive. That is the objective of our first recommendation.

Recommendation 1: A process for identifying an evolving portfolio of representative projects requiring Cyberinfrastructure should be developed by ENG. The process should be used to identify – and evolve – a set of specific *Cyberinfrastructure User Requirements* that can drive the design, development, and deployment of useful engineering-oriented cyberinfrastructure.

It is also essential that ENG have a clear understanding of the full range of current cyberinfra-

structure resources available for engineering research and education. We suggest developing an **Engineering Cyberinfrastructure Resources** catalog identifying

- cyberinfrastructure resources available to, and used by, the engineering community through efforts funded by ENG
- cyberinfrastructure resources funded through other Directorates and Offices at NSF that are being leveraged by the engineering community
- cyberinfrastructure resources funded by other government agencies or the industrial sector that are being leveraged by the engineering community

Once the catalog is developed, it can be used in conjunction with the Cyberinfrastructure User Requirements as the basis for identifying targets for ENG investments. The Directorate will have a comprehensive context in which to discuss priorities: what is urgent; what should be deferred; what can be left to another group; and what is most appropriate for joint efforts. This leads to our second recommendation.

Recommendation 2: A process for tracking the Engineering Cyberinfrastructure Resources (from all funding sources) currently being used or leveraged by Engineering research and education communities should be developed by ENG. This catalog should be used in conjunction with the user requirements to identify the most appropriate targets for ENG cyberinfrastructure efforts.

The process established through these initial recommendations will lay the groundwork for a targeted approach to cyberinfrastructure investments. The next sections focus on opportunities for ENG to expand on this foundation in the way it handles the organization, funding, and evaluation of its cyberinfrastructure efforts.

Cyberinfrastructure is a work-in-progress, not just because functional components have not been coordinated and made to interoperate within a useful “end-to-end” framework, but because many of the most challenging problems are still the focus of active research. Further, this situation is expected to continue for at least the next decade – as cyberinfrastructure evolves, new challenges will be identified. There is an opportunity for ENG to take a leadership role by identifying a handful of “**Engineering Cyberinfrastructure Research Challenges**,” or targeted research problems whose solution will increase the functionality, usability and stability of next-generation cyberinfrastructure. Our third recommendation reflects this concept of using cyberinfrastructure as a “real-world driver” for new ENG research.

Recommendation 3: A small set of Engineering Cyberinfrastructure Research Challenges should be identified by ENG that will be critical to the viability of next-generation cyberinfrastructure, and they should be supported through targeted research efforts.

A solid understanding of user requirements, current cyberinfrastructure resources, and high-priority research challenges will help ENG provide leadership for its communities in developing

and leveraging cyberinfrastructure. It will also provide the foundation for an integrated investment approach that can serve as a model for other groups.

3. Framework for ENG’s Cyberinfrastructure Efforts

The effectiveness and success of ENG’s cyberinfrastructure efforts will derive in large part from the organizational model. In effect, the model must provide a framework that integrates providers, users, and evaluators in a functional way that promotes accountability as well as resulting in useful, usable, stable, and evolutionary cyberinfrastructure.

3.1. The Innovation Loop

Engineering research and education *must* be the drivers of the infrastructure in order to maximize the cyberinfrastructure’s potential as an enabler of new discovery. Ideally, the relationships among programmatic research, development, and deployment form an ***Innovation Loop*** (Figure 1) which links complementary efforts and provides an evolutionary spiral resulting in ever more capable infrastructure and “reality-based” research.

The Innovation Loop focuses on three programmatic components that are key to cyberinfrastructure activities for the engineering and other NSF communities:

- *Research* efforts, which focus on discovery and innovation
- *Development* efforts, which adapt research methods and mechanisms to yield prototypes suitable for limited-scale use
- *Infrastructure deployment* efforts, which “harden” prototypes for robust, broad-scale use and add ongoing maintenance and support



Figure 1. The Innovation Loop

The three components are linked through the Innovation Loop as follows. Current infrastructure deployment supports research (deployment → infrastructure-enabled research). Promising research approaches can be further developed and engineered into new infrastructure prototypes (research → further development). Promising and proven prototypes under development can be further documented, integrated within existing systems, and engineered to become a part of the deployed infrastructure (prototype development → deployment). The challenges of current deployed infrastructure provide a driver for new research which can build the next generation of infrastructure (deployment challenges → research challenges).

It is important to note that the *model for success* differs for each stage:

- The goal of research programs is discovery, disciplinary advances, and innovative new ideas.
- The goal of development programs is to produce prototypes which expand on proof-of-concept research and allow evaluation of usefulness, usability, and potential community impact.
- The goal of deployed infrastructure programs is to provide a useful, usable, and functional basis for enabling research and education. Usability, usefulness, performance, robustness, and other characteristics are more important than innovation for these efforts.

In the same way, funding and execution mechanisms must be adjusted according to the particular stage in the Innovation Loop. Research programs, for example, are inherently high-risk, and the decision to continue funding may consider peripheral benefits as much as progress along the proposed path. For a deployment effort, on the other hand, the ability to meet milestones and deadlines will be critical for continued funding.

The real opportunity for ENG – and more broadly, NSF – is to develop organizational models which not only encourage the research, development, and deployment efforts, but also link them synergistically in a way that promotes new discovery and better infrastructure over time. Section 4.4 offers specific recommendations for how to do this.

3.2. Organizational Relationships

The Innovation Loop provides a focus for programs and inter-program relationships targeted at creating and leveraging cyberinfrastructure. ENG's organizational model for cyberinfrastructure should facilitate these synergistic programs. At the same time, ENG's cyberinfrastructure model should complement the NSF-wide model being developed by the OCI and the CIC.

There are many players in cyberinfrastructure, and an organizational infrastructure which coordinates and defines the roles for each is critical for success. At this point in time, NSF's cyberinfrastructure constituency includes the following groups:

- *ENG and other Directorates.* ENG is responsible for promoting the progress and success of engineering. Its cyberinfrastructure efforts focus on developing and promoting the infrastructure needed to support innovation, education, advances, and new discovery in the engineering community.
- *Engineering Cyberinfrastructure Users.* That portion of the engineering community whose efforts are enabled by cyberinfrastructure in all its forms – computational resources, data and storage resources, software tools and technologies, networking resources, visualization resources, scientific instruments, human resources, etc.
- *Engineering Cyberinfrastructure Providers.* ENG-funded projects and providers whose efforts provide cyberinfrastructure (compute, data, software, human, etc.) resources for the broader engineering community. This includes project-oriented cyberinfrastructure efforts such as NEES, NNIN, and CLEANER, as well as more broad-based cyberinfrastructure providers.
- *OCI Cyberinfrastructure Providers.* NSF funds a number of centers and projects to create

and deploy cyberinfrastructure on behalf of the national community. They serve the Engineering user community (providing computational, data, software, or human infrastructure), but are funded through OCI rather than ENG. Such providers include the NSF supercomputer centers, TeraGrid partners, etc.

- *Cyberinfrastructure Council (CIC)*. Chaired by NSF Director Arden Bement, this group includes the Directorate ADs and Office Directors. CIC engages in strategic planning, development and review of budget recommendations and oversight of cyberinfrastructure efforts for the Foundation.

Figure 2 illustrates the interrelationships among these groups.

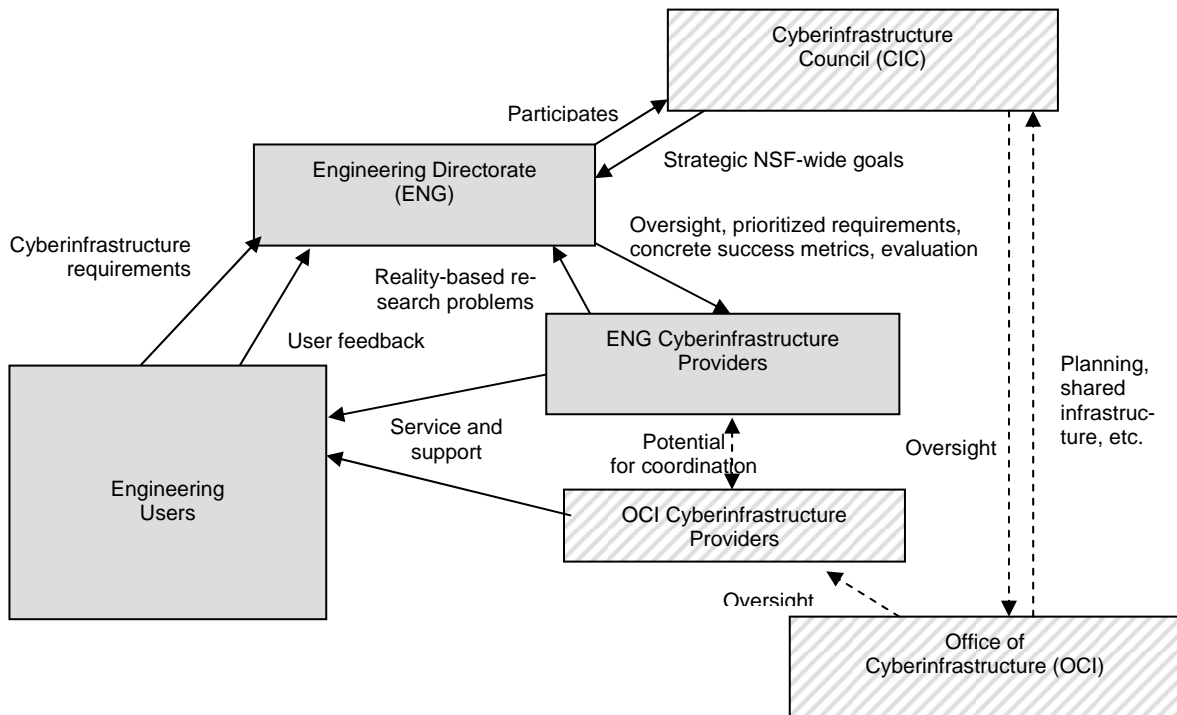


Figure 2. *Organizational Relationships for Engineering Cyberinfrastructure*
(patterned boxes are NSF components not part of/funded by ENG)

Cyberinfrastructure spans all the disciplinary communities served by NSF. In addressing the organization of Foundation-wide cyberinfrastructure efforts, the *Atkins Report* recommended a separate office, reporting to the NSF Director. This recommendation has now been implemented through the OCI. At the time of this writing, the OCI is in an early stage of development and Dr. Dan Atkins has been announced as its Founding Director. To avoid duplication of efforts or critical gaps in services, it is important that there be close coordination between ENG-funded and OCI-funded cyberinfrastructure efforts. Similarly, the groups’ oversight functions for their respective cyberinfrastructure providers should be coordinated and consistent. This is the focus of our fourth recommendation:

Recommendation 4: ENG should coordinate with OCI and CIC in the provision of Cyberinfrastructure. In particular, mechanisms should be put in place to ensure integrated infrastructure and consistent evaluation processes.

The decision of whether a separate organizational group is also required within ENG is beyond the scope of this report. It is important, however, that ENG explicitly consider the importance and priority of its cyberinfrastructure efforts in the broader contexts of its full portfolio and its relations with other Directorates, and establish an organizational structure that supports and promotes cyberinfrastructure at the appropriate level.

Recommendation 5: ENG should determine where cross-cutting cyberinfrastructure activities fit within the Directorate and ensure that the organizational structure promotes and supports cyberinfrastructure at a level appropriate with ENG’s strategic goals.

3.3. Guidelines for Ensuring Cyberinfrastructure Success

The broad goal of cyberinfrastructure is to foster innovation in engineering research and education. To achieve this goal, it is important to distinguish cyberinfrastructure from other types of ENG activities. The following guidelines provide a framework for developing new funding models and programs which promote the development of successful cyberinfrastructure.

1. ***Cyberinfrastructure should be treated as infrastructure.*** Like other forms of infrastructure that support and enable new discovery, cyberinfrastructure efforts should focus on providing a useful, usable, and stable foundation for research and education efforts. The addition of new infrastructure and the elimination of poorly performing infrastructure must be managed so as not to destabilize or retard the efforts of users.
2. ***Like other infrastructure, cyberinfrastructure will need to be maintained (not just built).*** Cyberinfrastructure resources – be they machines, networks, sensors, or data – will need to be maintained and upgraded/evolved on a regular basis. Even more importantly, the trained and seasoned professional workforce who provide key cyberinfrastructure services need to be considered an integral part of the infrastructure. Both resources and people must be sustained in order to continue achieving success in cyberinfrastructure.
3. ***Cyberinfrastructure efforts should emphasize development and deployment – which yield short and medium-term impacts – more than research.*** In particular, cyberinfrastructure research, development, and deployment should facilitate the Innovation Loop. Cyberinfrastructure research efforts should be targeted to critical missing elements without which cyberinfrastructure cannot succeed.
4. ***New cyberinfrastructure efforts should “bootstrap” infrastructure efforts in emerging cyber-communities.*** Cyberinfrastructure should leverage successful infrastructure efforts as they emerge. This means exploiting the current foundation of understanding, tools, and technologies when acceptable and available, rather than starting from scratch.

5. *It will be necessary to clearly define the notions of “success” and “usefulness” for Cyberinfrastructure resources and services.* NSF traditionally defines and measures success in terms of the research model, which is an exploratory or open-ended process whose primary outcomes are publications, demonstrations, and/or testimonials. For cyberinfrastructure, it will be necessary to consider how success differs in the context of development and deployment efforts. An overriding criterion for success should be “did we enable something that would not have been as easy or that would not have otherwise been possible?”
6. *Cyberinfrastructure resources and services that help “the few” achieve research prizes are not the same ones that will attract “the many” to daily use.* Different approaches will be needed to address these very different needs. Cyberinfrastructure efforts simply cannot succeed if they are expected to be “everything for everybody.”
7. *It will be necessary to incentivize quality cyberinfrastructure efforts.* Development and deployment activities are not part of the traditional academic reward structure (i.e., do not normally lead to tenure and promotion). NSF and its Directorates will have an important and proactive role to play in challenging academic perceptions and helping to generate awareness of the importance of shared infrastructure to future research and education.

These principles are not well understood, and the challenge of putting them into practice accounts for much of the current confusion about how cyberinfrastructure functions and how it fits within existing research and development efforts. ENG has a unique opportunity to bring these principles to the forefront and to teach, by example, how they can be interwoven into NSF programs.

The key is to be explicit about what is needed and wanted, and to promote accountability in evaluation and assessment. It should be clear from the outset which portion of the Innovation Loop a particular program is addressing, and evaluation metrics and mechanisms, funding timeframes, deliverables, goals, etc., should be appropriate.

For example, cyberinfrastructure deployment projects should include the long-term maintenance (and continued enhancement) of the resources and sustainability for the people who provide cyberinfrastructure services. There also should be general agreement between ENG and the infrastructure provider as to the degree of risk that is acceptable. While research is generally accepted to be high-risk, development activities should at least include proactive risk management and mitigation, and deployment must generally be low-risk in order to justify its return on investment. It is critical at the outset for NSF and the infrastructure provider to be in agreement about what is expected and how it will be measured.

Recommendation 6: ENG should adopt a set of guidelines distinguishing cyberinfrastructure from other types of research and education efforts, and work to ensure that ENG cyberinfrastructure program is consistent with these guidelines.

4. Funding and Program Models for Cyberinfrastructure

The development and deployment of successful cyberinfrastructure does not happen by chance. It is critical to plan and structure funding efforts so that they will have significant and lasting impact on engineering research and education. In this section, we explore funding and programmatic models pertinent to ENG's cyberinfrastructure activities.

4.1. Structuring ENG's Cyberinfrastructure Portfolio

To establish a structure for ENG's investments in cyberinfrastructure, it is first essential to arrive at some key decisions about the balance with which different types of efforts should be included in the portfolio and about their priorities. We formulated the major issues as a series of *framing questions*, as follows.

- What **proportion of ENG efforts and funding** should be allocated to cyberinfrastructure-related research, education, development, and deployed infrastructure, respectively?
- What **kinds of groupings (single PI, project groups, centers, etc.)**, and in what proportion, are appropriate for each? What funding model is required to ensure that these efforts are successful?
- What are the **realistic costs (startup, continuing, long-term, growth)** associated with each kind of effort? How much of ENG's budget can be "mortgaged" for continuing or long-term efforts?
- Under what circumstances should **multiple projects addressing the same technical issues be sought** (e.g., distinct research efforts targeting the same technical problem) and under what circumstances should economies of scale be sought (e.g., centralized hosting of community data collections)?
- Which efforts should be linked to form the Innovation Loop? What are the **programmatic and budget issues associated with linking and coordination**?
- **Which efforts should be fostered from within ENG**, and which efforts should the Directorate encourage be funded elsewhere (other directorates, OCI, non-NSF academic efforts, the private sector, etc.)?

By discussing and deciding these issues, ENG will be able to identify a set of "structuring priorities" for its portfolio. This is particularly important for cyberinfrastructure, where the kind of *de facto* priorities established by the engineering community – through its submission of successful proposals – are not likely to represent a coherent, long-term strategy.

Recommendation 7: ENG should develop a framework of priorities and apply it to purposefully structure the cyberinfrastructure portfolio (driving both funding and programmatic decisions).

4.2. An “Infrastructure” Funding/Program Model

Engineering cyberinfrastructure should be useful to and usable by a substantial constituency, be stable and evolutionary so that it continues to meet community needs over time without major disruption, and should provide effective functionality, capacity, and capability over and above what Engineering researchers, educators, and practitioners can develop on their own.

Although many Directorates at NSF have their own versions of infrastructure (e.g., GEO procures planes and ships for its research projects, CISE facilitates community software development, and OPP provides residential, experimental, and other support at its South Pole station), the scope and complexity of cyberinfrastructure provides unprecedented challenges. The development of a viable and appropriate **Infrastructure Model** for supporting cyberinfrastructure will be fundamental to its success.

NSF’s traditional funding/program model enables *research*, which focuses on innovation and the discovery of new ideas. In contrast, the priorities for infrastructure efforts are usefulness and usability, along with stability, reliability, and capability. Table 1 summarizes key differences between research and infrastructure models.

	(Current) Research Model	Infrastructure Model
Goals	New discovery, innovation	Provision of useful and usable tools/technologies to support research and education
“Deliverables”	Innovative results, as represented by published results in peer-reviewed conferences and journals, presentations, recognition, etc.	Explicit lists of tasks and tools, with explicit specifications, timeframes, quality measures, and expectations.
Timeframe	Grant awards are typically 3-5 years	Timeframes should incorporate realistic startup, continuation, growth, and transition times. Typically requires 5-10 years to be cost-effective
Selection Criteria	Peer-reviewed evaluation, focusing on intellectual merit, broad impact, and other programmatic and/or Directorate goals	Evaluation by infrastructure experts as well as users. Selection criteria should include compliance of proposed infrastructure to user needs, soundness of approach, track record of provider, cost-effectiveness of plan, etc.
Evaluation	Mid- or post-award evaluation of research against plan as described in annual reports	Evaluation of infrastructure at regular intervals, using concrete metrics of usefulness, impact, usability, accountability of providers, etc.
Costs	Typically initial outlay for equipment, followed by relatively steady personnel costs	Significant initial outlay with major outlays at intervals. In addition, recurring costs for maintenance (upgrades of hardware, new/improved versions of software), operation, and labor for infrastructure development and user services. “Last mile” costs of documentation, training, access, targeting and retargeting of infrastructure to current platforms can be substantial.
Transition Approach	PIs submit new research proposals or engage in discussions with NSF about other possible continuation mechanisms for research	When one provider replaces another, products of infrastructure must be explicitly transitioned between providers in a way that minimizes disruption for users.

Table 1: Comparison of Research Model and Infrastructure Model

As indicated in Table 1, successful infrastructure awardees should demonstrate that they can provide support, service, and added value to ENG researchers, educators and practitioners. It is critical that infrastructure be grounded in a sound business model - a model which considers the needs of the “market,” the appropriate “products” and “services,” the logistics of coordination with other infrastructure awardees, and the strategic impact of the infrastructure to current and evolving Engineering communities. Moreover, commitment from the host institution to support infrastructure providers should also be considered.

Because of this, it is important that the review of proposals from infrastructure providers should be performed by individuals who understand how infrastructure is built and operated, so proposals can be assessed appropriately. Some of the issues of infrastructure provision are better understood in the private sector, and such expertise should be included in selection and review panels. The track records of candidate infrastructure providers must also be considered. In some ways, the selection of cyberinfrastructure providers more closely resembles the choice of candidate startups by venture capitalists than the prioritization of peer-reviewed research projects by academic panels.

Note that unlike NSF’s traditional Research Model, the Infrastructure Model must incorporate the recognition that infrastructure maintenance and transition both incur significant costs. Infrastructure transition, evolutionary or otherwise, must be done in a way that minimizes negative impact on infrastructure users. In particular, terminating infrastructure without adequate transition plans can have disastrous consequences on users, who often come to depend on the tools/technologies and may not be able to work effectively without them. (Imagine “sunsetting” your local highway or airport without providing working alternatives).

Recommendation 8: In coordination with other Directorates and OCI, ENG should develop a new funding/program model that explicitly addresses the nature of *infrastructure*. It should incorporate realistic models for evaluation, selection, funding, and assessment.

4.3. Partnership Opportunities

ENG has been an innovator Foundation-wide with respect to substantive partnerships between the academic and private sectors. The Directorate’s *Industry/University Cooperative Research Program*, for example, focuses on the development of centers for research that is of interest to both groups. Catalyzed by a small ENG investment, I/URCs are primarily supported by the industry members at each center. The *Engineering Research Centers* provide another example. In this case, the focus is definition, development, and validation of engineering technologies with the potential to spawn whole new industries or radically transform the methodologies of current industries. ERCs combine research and education with real-world technology and involve substantive participation by industry.

Cyberinfrastructure provides new opportunities for partnering with industry as well. In the U.K., the successful *e-Science Programme* supports cyberinfrastructure-related projects that are jointly proposed by industry and universities; industry provides half the funding and government the other half, and there are firm deliverables and milestones for all parties. We believe that partnerships with industry can greatly benefit cyberinfrastructure – providing an evolutionary path as the hardware, software, and data technologies developed and deployed for the academic envi-

ronment transition to broader user constituencies (and more substantive “markets”).

For such partnerships to be successful, NSF will need to assist in establishing policies that address the IP challenges faced by the academic and private sectors. We suggest that NSF require that software, data, and documents developed as part of NSF-funded cyberinfrastructure be open source. This type of requirement will be critical to “level the playing field”; without it, university-specific policies and procedures regarding IP and technology transfer handicap many researchers and force multi-institutional collaborations to adopt the lowest-common-denominator (which is almost invariably unacceptable to industry). Requiring cyberinfrastructure products to be in the open arena will significantly accelerate the creation of economic value and the impact of cyberinfrastructure on the nation.

Recommendation 9: ENG, in coordination with OCI, should develop guidelines requiring that the products of funded cyberinfrastructure projects be available in the open domain. This will both decrease the IP-related burden on individual researchers and accelerate technology transfer.

4.4. Creating “Links” to Power the Innovation Loop

The Innovation Loop offers a programmatic framework for creating, developing, deploying, and using cyberinfrastructure. It cannot function as a loop, however, unless new programs explicitly target the links among the stages and enable the synergies among them. Three types of programs are needed.

- **Infrastructure → Research Program.** In Section 2, we discussed the initiation of an “Engineering Cyberinfrastructure Research Challenges” program which could be used to galvanize interest in the engineering community in addressing some of the most difficult challenges for deployed cyberinfrastructure. This linking program would likely be suitable for funding using the traditional, peer-reviewed Research Model.
- **Cyberinfrastructure Research → Development Program.** This would focus on further funding for research projects which have yielded early prototypes that warrant further development and would be suitable for integration within existing cyberinfrastructure. One possible scenario is to enable research PIs to apply for additional funding to support prototype extension and hardening; selection panels would focus on the prototypes’ potential for broader impact and their suitability for more robust engineering. Such projects would have accelerated timeframes (1-2 years) and would be expected to yield concrete deliverables that are integration-ready.
- **Cyberinfrastructure Development → Deployment Program.** This program would provide additional funding for integrating robust, proven prototypes (the outcomes of the previous linking program, or of other development-oriented efforts) into existing cyberinfrastructure. Funding would support documentation, targeting to specific cyberinfrastructure resources, interoperability, maintenance, etc. These projects would be assessed using the Infrastructure Model. Funding would include an initial “integration/ documentation” phase and an “evolution/maintenance” phase, both consistent with the timeframes expected for infrastructure.

The development and funding of these linking programs is summarized in our tenth recommendation.

Recommendation 10: “Linking programs” should be developed by ENG to ensure the synergistic coordination of cyberinfrastructure-related research, development, and deployment programs.

5. Evaluating Cyberinfrastructure Programs

Cyberinfrastructure resources and services will need to be evaluated in a different context than traditional research awards. In particular, it will be critical to consider the effectiveness, usability, and usefulness of cyberinfrastructure, in the context of its ability to enable new research and education. In this section, we focus on development of metrics for assessing whether deployed and developing cyberinfrastructure is achieving its goals, and for determining whether the efforts associated with cyberinfrastructure have provided an adequate return on investment. We also suggest how ENG solicitations and awards can promote accountability and achievement of those metrics.

5.1. Framing Questions

In Section 4.1, we discussed the framing questions that must be resolved before funding priorities can be set. Another set of framing questions is pertinent to defining the notion of “success” for cyberinfrastructure before evaluation metrics and processes can be established. These include:

- Which **type of impact** is desired: deep impact on “the few” (i.e., “new science”) or broad impact on “the many” (daily use)?
- What is the **balancing point between target constituencies**, e.g. between supporting engineering research and (university) education, versus broader outreach to K-12, practitioners, and/or the general public?
- What are the **priorities for the types of resources/services** (compute servers, data servers, models, software tools, instruments, etc.) to be supported?
- What is the **balancing point between specialized resources/services for one domain, and generalized support** that helps all (but to a much more limited extent)?
- What **level of usability** is needed?
- How important is **interoperability across infrastructure**?
- What are the **acceptable modes of failure** for infrastructure? What modes of failure are not acceptable?
- What level of **QC/QA, software engineering, and risk management** are expected?
- What **level of accountability** is important for research, education, cyberinfrastructure-related

development and deployment efforts?

- What is the **balancing point between attracting new users for cyberinfrastructure and ensuring productivity for recurring users?**

ENG can provide important leadership to the Foundation by developing a process for formalizing priorities and deriving evaluation metrics that promote accountability as well as robust enabling infrastructure – and sharing that process with other Directorates and OCI.

Recommendation 11: ENG should discuss and resolve an appropriate set of framing questions for evaluating cyberinfrastructure programs and projects, and work to ensure that ENG programs respond to the resulting priorities.

5.2. Establishing Accountability for Cyberinfrastructure Projects

As with any infrastructure, cyberinfrastructure must be reviewed and evaluated to ensure success. ENG will need to define a process for reviewing and evaluating infrastructure on a recurring basis. Clarifications about what infrastructure will be reviewed, who will be doing the reviewing, what happens if expectations are not met, etc., will be key to promoting and ensuring accountability within the system. This will establish the kind of accountability that is crucial in developing useful and robust infrastructure.

Once the framing questions have been resolved, appropriate metrics of success can be incorporated into solicitations, review processes, and portfolio evaluations. We provide example evaluation criteria below for cyberinfrastructure in Table 2.

Many more criteria are possible and appropriate for different kinds of Cyberinfrastructure. Note that some metrics will require new types of instrumentation to provide evaluation data.

It will be critical for the review and evaluation process to actually *apply* the metrics established in program descriptions to assess cyberinfrastructure success. Infrastructure development and deployment projects, in particular, will require regular reviews by NSF to ensure accountability. Feedback from external review committees should be complemented by independent metrics of success to reduce subjectivity in the review process. By way of example, the textbox on Page 16 demonstrates how the tradeoffs might be balanced in defining and evaluating a program, *Program X*, to develop engineering cyberinfrastructure.

Recommendation 12: ENG should develop explicit metrics of success, evaluation processes, and expectations for accountability and incorporate them in solicitations for infrastructure programs and projects.

It is important that both ENG staff and review committees recognize that new infrastructure is hard to build and integrate. One approach would be to evaluate infrastructure efforts using a *red light/yellow light/green light* grading scale. Infrastructure which meets expectations receives a “green light” of approval and is allowed to continue/evolve over an additional review cycle (similar to “rolling tenure”). Infrastructure which receives a “yellow light” must demonstrate that the problems are addressed at the next review cycle or earlier. Infrastructure which receives

a “red light” should cease to be supported as infrastructure. Development of aggregate measures using the red light/yellow light/green light scheme could be applied to larger portfolios of infrastructure efforts, and/or to cyberinfrastructure providers.

Recommendation 13: Selection, review, and evaluation processes for cyberinfrastructure programs should reward usability and usefulness, and disincentivize poorly engineered infrastructure, or infrastructure without a sufficient user community.

Metric Type	What It Assesses	Example Metrics
Usage	Amount of use of resource by user community	Number of users of resource; Utilization, throughput (computation); Number of collections (data); Number of hits (web); Number of downloads (software), etc.
Usability	“Ease of use” of resource by user community	Turnaround time (computation); User satisfaction as assessed by surveys; Informal feedback from users; Software productivity measures
Deep impact	Importance of science and engineering enabled by resource	Publication in peer-reviewed journals and conferences; Community recognitions and awards; “Landmark” publications
Broad impact	Extensiveness of user community, accessibility of resources	Number of disciplines, communities served; Number of publications enabled, Number of courses, dissertations, and other educational vehicles enabled
Expanding Use of Cyberinfrastructure	“Growth” of cyberinfrastructure as an enabling technology	Number of new users (greater than some threshold of times) of cyberinfrastructure components and resources
Coordination of Cyberinfrastructure	Integration and interoperability of Cyberinfrastructure components	Number or percentage of times that resources or software is used together
Technology Transfer promoted by Cyberinfrastructure	Movement of academic cyberinfrastructure efforts to the private sector or “productization”	Number of deployed cyberinfrastructure tools and technologies initiated with the academic community and productized within the private sector
Workforce impact	Individuals involved in the provision of cyberinfrastructure	Number (gender, race, creed, level) of individuals involved in cyberinfrastructure-related professions; Number (gender, race, creed, level) of individuals with cyberinfrastructure-oriented education or training and their increase/decrease over time

Table 2. Example success metrics for cyberinfrastructure

An Example Strategy

Defining Success: The description for *Program X* and all solicitations would explicitly state that:

- *Program X* addresses major infrastructure development, specifically the development of cyberinfrastructure that will make it possible for interdisciplinary researchers to share, link, and exchange computational models for <insert programmatic goal here – e.g. “the transport of water, elements, contaminants, etc. through fluids, air, and soils”>.
- The infrastructure developed for *Program X* should satisfy the following criteria <insert appropriate ENG criteria here> and should be in “production mode” and utilized within the broader community within <insert appropriate timeframe here>.
- Justification must be provided for why the development of “new” infrastructure proposed for *Program X* will enhance cyberinfrastructure and provide improved functionality or capability within the proposed timeframe.
- Infrastructure developed and deployed for *Program X* must be capable of attracting new users from multiple disciplines and at varying levels (undergrads to professional researchers). The developers/providers must identify and assimilate user requirements from Engineering’s *Cyberinfrastructure User Requirements* document; both the communities and the procedures for doing this must be specified in the proposal.
- Infrastructure developed and deployed for *Program X* must coordinate with <insert appropriate resources here > from the *Engineering Cyberinfrastructure Resources* document; the procedures for doing this must be specified in the proposal.
- Proposed infrastructure should interoperate with <insert appropriate existing software here>. The proposal must identify how this will occur and what obstacles must be overcome in order to do so.
- The prototype infrastructure must be demonstrably robust and scale to simultaneous use by at least <insert number here> individuals. A beta version containing 50% or more of the proposed functionality must be available no later than <insert number here> months before the end of the project. The proposal should describe what user groups have committed to serving as early users and evaluators of that system.
- The proposed infrastructure must include plans for complete documentation and for maintenance and evolution over a 10 year period. The proposal must address what kind of risk management strategy will be used to ensure the infrastructure is successful. It must also lay out the QC/QA and formal software engineering procedures to be followed, include specifics about the project team’s previous experiences in these areas, and provide a realistic assessment of the human infrastructure required to develop and maintain the infrastructure over the long-term.

Evaluating success: The description/solicitations for *Program X* will also explicitly address the mechanisms whereby project success will be assessed:

- The PI is required to submit <insert regular timeframe of up to 12 months> self-assessment reports that detail the QC/QA and software engineering documents prepared during the evaluation period, evaluate the current risk situation, and propose adjustments to ensure that the project remains on schedule and on budget.
- The proposal must address how the early user program will be structured to ensure that initial usability, usefulness, and overall user satisfaction can be measured and fed back into the development cycle at least 12 months before the end of the project.
- At the end of the first year, the PI will propose the quality metrics (based on software engineering practice) and usability metrics (based on the user requirements) to be used in evaluating project success. The metrics will be applied by the project staff and reported in a document to be reviewed by a panel of CI experts.

6. Summary of Recommendations

The EAC-CI members appreciate the opportunity to provide input on cyberinfrastructure to the Engineering Directorate. The ideas and recommendations within this report are intended to help define concrete steps in which ENG can play an important and leadership role, both within the Foundation and within the broader science and engineering community.

For convenience, we reprise our recommendations below, ordered thematically.

Assessment Recommendations:

The following recommendations are designed to enable the Engineering Directorate to assess its needs, resources, and progress in a way that can be incorporated into strategic planning and programmatic initiatives.

Recommendation 1: A process for identifying an evolving portfolio of representative projects requiring Cyberinfrastructure should be developed by ENG. The process should be used to identify – and evolve – a set of specific *Cyberinfrastructure User Requirements* that can drive the design, development, and deployment of useful engineering-oriented cyberinfrastructure.

Recommendation 2: A process for tracking the *Engineering Cyberinfrastructure Resources* (from all funding sources) currently being used or leveraged by Engineering research and education communities should be developed by ENG. This catalog should be used in conjunction with the user requirements to identify the most appropriate targets for ENG cyberinfrastructure efforts.

Recommendation 12: ENG should develop explicit metrics of success, evaluation processes, and expectations for accountability and incorporate them in solicitations for infrastructure programs and projects.

Recommendation 13: Selection, review, and evaluation processes for cyberinfrastructure programs should reward usability and usefulness, and disincentivize poorly engineered infrastructure, or infrastructure without a sufficient user community.

Coordination Recommendations:

The following recommendations are designed to enable the Engineering Directorate to coordinate with other Directorates and OCI within NSF in the provision and use of Cyberinfrastructure.

Recommendation 4: ENG should coordinate with OCI and CIC in the provision of Cyberinfrastructure. In particular, mechanisms should be put in place to ensure integrated infrastructure and consistent evaluation processes.

Recommendation 8: In coordination with other Directorates and OCI, ENG should develop a new funding/program model that explicitly addresses the nature of *infrastructure*. It should incorporate realistic models for evaluation, selection, funding, and assessment.

Recommendation 9: ENG, in coordination with OCI, should develop guidelines requiring that the products of funded cyberinfrastructure projects be available in the open domain. This will both decrease the IP-related burden on individual researchers and accelerate technology transfer.

Planning Recommendations:

The following recommendations are designed to enable the Engineering Directorate develop a plan for incorporating Cyberinfrastructure programs and Cyberinfrastructure-enabled efforts into the broader ENG portfolio.

Recommendation 5: ENG should determine where cross-cutting cyberinfrastructure activities fit within the Directorate and ensure that the organizational structure promotes and supports cyberinfrastructure at a level appropriate with ENG's strategic goals.

Recommendation 7: ENG should develop a framework of priorities and apply it to purposefully structure the cyberinfrastructure portfolio (driving both funding and programmatic decisions).

Recommendation 11: ENG should discuss and resolve an appropriate set of framing questions for evaluating cyberinfrastructure programs and projects, and work to ensure that ENG programs respond to the resulting priorities.

Recommendations for Building the "Innovation Loop":

The following recommendations are designed to enable the Engineering Directorate build the "Innovation Loop" through programs and initiatives.

Recommendation 3: A small set of *Engineering Cyberinfrastructure Research Challenges* should be identified by ENG that will be critical to the viability of next-generation cyberinfrastructure, and they should be supported through targeted research efforts.

Recommendation 6: ENG should adopt a set of guidelines distinguishing cyberinfrastructure from other types of research and education efforts, and work to ensure that ENG cyberinfrastructure program is consistent with these guidelines.

Recommendation 10: "Linking programs" should be developed by ENG to ensure the synergistic coordination of cyberinfrastructure-related research, development, and deployment programs.