Appendix C: More on Organizational Issues

This appendix provides some additional detail surrounding the organization of the ACP within NSF.

Organizational Alternatives

Alternative organizational structures can be made to work given the appropriate level of cooperation across organizational boundaries. Given sometimes competing objectives, no single organizational approach can capture seamlessly every aspect of what is to be accomplished. Having said this, the organizational approach can be a useful tool for emphasizing and promulgating the most fundamental, cherished, and/or difficult-to-achieve goals, and communicating to everyone involved (NSF and the research community it serves) the goals and their priorities. The greatest challenges were discussed in the body of the report, and our organizational recommendations emphasize successfully addressing these challenges with minimum disruption of the existing NSF organization.

There are many ways the ACP could be organized, and it is useful to list some alternatives considered and their perceived shortcomings:

Overlay the ACP on the current organization. The Panel believes that the INITIATIVE, to be fully successful, must be an agency mission with the highest priority and the highest visibility, and the organization of the ACP within NSF should reflect this. In our opinion, the ambitious goals of the INITIATIVE cannot be achieved by business as usual, but neither does it demand radical changes.

Making fine-tuning changes. For example, a fine tuning might consist of simply combining ACIR and ANIR into a single division within CISE with responsibility for all infrastructure. Again, the Panel believes this doesn't place adequate importance and visibility on ACP to be successful.

Centralize the ACP in a single organization. The Panel considered the option of creating a separate organization within the Office of the Director to centrally manage an infrastructure program on behalf of all directorates, similar to the Office of Polar Programs. This has the disadvantages that it does not involve either computer scientists or domain scientists integrally in the ACP. We believe that such an organization, in spite of best intentions, could evolve toward a typical 'information systems' organization that focused on procurement and operations, and might give inadequate attention to new technological opportunities and would not be sufficiently responsive to the needs of the end-user communities.

Distribute the ACP among all Directorates. In this approach, each directorate would take complete and exclusive responsibility for the infrastructure and applications supporting its respective community. The problems with this approach are very clear from the preceding discussion. Like the Office of the Director solution, it might focus too much on procurement and operation of current technologies, may result in excessive duplication of effort, and over time could create a serious balkanization of infrastructure becoming a serious obstacle to interdisciplinary collaboration and programs.

Technology Transfer

It is useful to specify in slightly more detail the horizontal technology-transfer dimension in Figure 4.2. The research actually includes two distinctly different flavors of research:

Fundamental, longer-term research in information technology and its applications. This type of research pursues revolutionary new ideas and fundamental understanding without being constrained by the current environment. This type of research, while extremely important, generally falls outside the scope of the ACP, with the exception of new long-horizon research on systems (social and technological) specifically supporting cyberinfrastructure and applications.

Applied, nearer-term research in information technology and its applications. This type of research seeks nearer-term outcomes that take strong account of and explicitly try to change and enhance the current environment. Its outcomes often include working prototypes fitting within an existing environment that can later be leveraged as a starting point for development, after they prove their mettle and after refinement through end-user experience.

Research outcomes that are deemed promising for the science and engineering research community are moved into the development and operational phases. A more detailed description of these phases includes:

Development of applications and infrastructure includes a set of activities resulting in a set of working, interoperable, and maintained implementations of working infrastructure and applications. The outcome is a set of software distributions that are interoperable and work within a prescribed environment of equipment and other software, including commercially available software.

Operations of applications and infrastructure include a set of activities resulting in a production environment where executing infrastructure and applications support end-users while they conduct science and engineering research. This includes provisioning, wherein the required facilities and equipment and various software modules are acquired and integrated and tested, and support of users in effectively using the capabilities.

A role of infrastructure is to reduce the time and effort and cost required for the development, provisioning, and operations of applications. Experience indicates that these phases do not follow sequentially, but rather it is most effective to repeat them in a process of successive refinement (in the context of software development, this is sometimes called the *spiral* model).

The development phase can be further subdivided into some constituent functions:

Conceptualization and analysis. Identification of an opportunity, assessment and analysis of needs, and development of a detailed set of requirements. It is particularly important in this ACP that this incorporate the outcomes and prototypes from NSF-sponsored information technology research.

Design. Choose an architecture (how to divide and conquer the implementation) and develop a plan (including identification of what designs can be reused, what can be purchased off the shelf, and what needs to be developed).

Implementation and testing. Programming new software or adapting existing software (e.g. from a research prototype), testing and refinement in the intended operational environment. In many cases a starting point may be a prototype arising from research. The outcome is a single software distribution to be used everywhere.

Maintenance and upgrade. Repair defects identified during provisioning and operations, and add new capabilities and features based on user needs identified during operational experience.

The Role of CISE

To be sure that the ACP leverages the most advanced technologies, utilizes state-of-the art software development processes and methodologies, remains fixated on advancing the information technologies themselves utilizing the opportunity to conceptualize and experiment with new applications in research, captures commonalities of need, and enables rather than hinders cross-disciplinary collaboration and programs, we recommend a continuing strong involvement of CISE in the management and implementation of ACP. In addition, like the other directorates, CISE's own research should also be a target of new applications and an opportunity to utilize the expanded infrastructure.

Activities within CISE comprise the three layers shown in Figure 4.3 of Section 4: core technologies, social and technological systems, and applications. The core technology layer includes a diverse set of individual technologies, available commercial products, processes, and best practices. The two higher layers comprise a set of integrated and coordinated activities, each activity dealing with three related activities: research, development, and provisioning and operations. The systems layer focuses predominantly on the infrastructure supporting applications, and the applications layer on new ways of conducting science and engineering research that are built upon this infrastructure. This architecture thus preserves both the vertical (core technology through infrastructure through applications) and horizontal (technology transfer from research through development and use) structure described earlier.

This layered architecture suggests that the horizontal grouping take precedence over the vertical, primarily because it will be more effective at capturing commonalities and coordinating activities across end-user communities. These are the most difficult goals to achieve, particularly so within NSF because of the separation of scientific and engineering disciplines in NSF and in the research community served by NSF. This organizational structure places these goals of commonality and coordination as the most explicit and visible within the organizational structure.

There are some additional points of evaluation that emphasize technology transfer:

Applied research seeks to influence the direction of development. Thus, a point of proposal evaluation should be the appropriateness of the research in light of end-users needs and a roadmap for serving those needs over the coming years, and the post-evaluation should be based in part on success in moving the ideas and prototypes into development.

A point of evaluation for development proposals should be ongoing processes for the technology transfer from research outcomes. Since specific research outcomes cannot be anticipated, there is an element of uncertainty, suggesting annual adjustments in direction. (The cooperative agreement mechanism for funding serves this end well.) Post-evaluation should focus on whether the development activity resulted in a stable and supported software distribution, whether and how proactively it has been provisioned, and (allowing a reasonable time for diffusion) and user satisfaction as well as (secondarily) how many users have been attracted.

A further point of evaluation for provisioning and operations is how effectively and proactively the organization has worked with developers to make the distributions available to users, and how effectively user experience and problems have been fed back to developers for maintenance, upgrade, and new capabilities.

There also needs to be an application program within CISE, collaborating closely with the other Directorates. This has several important purposes:

Like the other Directorates, the CISE research community can itself be served by the innovative application of information technology.

CISE involvement ensures the proximate and ongoing engagement of technology experts in identifying, formulating, and implementing new applications.

CISE takes primary responsibility to identify commonalities of need among different scientific and engineering disciplines that can be served by shared applications and infrastructure, to avoid unnecessary duplication of effort, and to empower future scientific collaboration across disciplines.

CISE takes responsibility for identifying and promulgating generic applications of wide interest across the NSF community. Examples include collaboration, data storage and archiving, digital libraries, numerical tools, and similar capabilities, all realized keeping in mind customization and extension to meet discipline-specific needs.

The Panel envisions an applications program in SBE as well. As in CISE, this has more than one purpose. One goal is to identify new applications to serve social scientists. A second goal is to involve social scientists in studying the application of information technology to groups and organizations, both how this can be done effectively and its impact.

The systems layer of the proposed architecture of ACP will fall predominantly within CISE (as well as the social sciences, see below). It groups a set of activities relating to heterogeneous compositions of diverse technologies with social constructs (groups, organizations, and communities). The idea is to serve several interests:

Infrastructure and applications emphasize systems composed of many technologies, including processing, storage, communications, and software. While there is considerable relevant activity in systems in many (and perhaps even all) of the current programs in CISE, the Panel believes it is time (and the challenges of the ACP focus attention on this) to view systems as a first-class target for research. By providing a focused effort on systems, more research attention can be focused here[1].

The infrastructure portion of ACP prominently involves systems issues. Thus, the systems layer is where activities surrounding infrastructure development, provisioning, and operations could reside.

Social systems figure prominently in the context of both the applications and infrastructure portions of the ACP. There is a deep integration of social systems with technological in both applications and infrastructure (the latter less obvious, but relating to the human organizations involved in development, provisioning, and operations). Thus, a systems activity should include social system issues, and thus collaborate with SBE.

Both applications and infrastructure suffer from a serious disconnect in fundamental objectives between technological researchers and end-user communities. The typical attitude of users is “we need it right away”, while technologists appropriately assert that “we don’t know the right way to do this until we do the research, and developing this right now will set in stone premature and suboptimum assumptions”. The Internet is an inspirational example of how these legitimate competing interests can reach a compromise by coupling both deployed infrastructure and its applications to *both* research outcomes and to end-user experience. The NSF middleware ACP is a recent example of how a program can be designed to base development of deployable infrastructure on a coupled and coordinated program of applied research, prototyping, and productizing of research outcomes.

The Role of Other Directorates

The post-evaluation of programs within the non-CISE directorates should focus on how substantially and beneficially the actual conduct of research has been changed, and how widely and effectively the new applications are actually used. Over time, resources should flow to directorates most successful in effectively using information technology to beneficially change the conduct of research.

This aspect of the ACP is a substantial change from the current Partnerships portion of the PACI program. In particular, as described in Section 5, the direction and funding of some portion of ‘enabling technologies’ and ‘applications’ would shift to the Directorates, where they would be subject to ordinary peer review. Of course, we expect the present PACI’s to participate in these competitions, often involving collaborations with domain scientists and engineers. Rather than placing the burden on Centers to find partners to participate in applications, this would shift the direct control to Directorates and base awards on a competitive peer review process. If proposals are initiated by the domain experts, and if successful proposals demonstrate directly the enthusiasm and commitment to revolutionizing the conduct of research in the discipline and not simply serving their own narrow requirements, much more will be accomplished. The direct involvement of domain-expert program managers within the directorates will stimulate interest and involvement among more domain scientists and engineers, and they will also serve as coordinators to make sure that the aggregate activity funded out of the Directorate forms a coherent and complete ACP serving an entire domain of science or engineering.

The future direction of the PACI program is discussed in Section 5.

These programs within the directorates are also expected to work closely with each other and with CISE. All proposals to a Directorate should be evaluated in part on the credibility of its plan to execute its vision by working with NSF-funded centers or others, and also in its coherence to the overall ACP. For example, does the proposed activity make appropriate and maximum use of centrally developed infrastructure, does it anticipate opportunities to serve a larger community, and does it avoid duplication of effort with related activities in other Directorates? Does it have a credible plan to develop and support production technology and applications for the benefit of the entire discipline? For this reason, all such programs should be considered joint programs with the relevant portions of CISE, and proposal evaluation should be a joint responsibility involving both domain and technology experts.

The Role of the PACI’s

The impetus for this ‘matrix’ form of organization surrounding applications shifts responsibility for initiating application research and development away from the PACI’s and toward the scientific and engineering directorates. This is motivated by some perceived shortcomings in the current organization surrounding application initiatives in CISE. While domain scientists and engineers are encouraged to participate through ‘partnerships’, these are largely ad hoc collaborations driven by individual initiative rather than any common vision or direction. If the conduct of science and engineering research is to be revolutionized, this will be based on leadership for creating and executing a vision emanating from the non-CISE directorates within programs focused on this objective led by committed and visionary managers who are themselves domain experts. If the non-CISE directorates and the communities they serve have a major stake in the outcome, they can provide the necessary leadership to rally researchers around creating and executing a vision. We envision ACP program directors within each of the directorates to be domain scientists or engineers who possess a deep and abiding interest in revolutionizing the conduct of research in their respective fields through information technology. We expect them to motivate and lead their respective communities, as well as define coherent programs that systematically approach this challenge.

All this applies to the funding side of the equation, but when it comes to delivery the role of the PACI's and other centers serving this ACP may not be greatly changed. This separation of funding should not preclude the grouping of activities within centers where this makes sense, such as software development and operations. Investigators within the scientific and engineering research communities will likely seek the involvement of centers in prototyping and productizing their application ideas, to bring in needed software engineering expertise and to lend credibility to technology transfer.

Industry Involvement

Not only should commercial technologies be acquired in preference to development of similar technologies, but the goal of the ACP should be commercialization of both cyberinfrastructure and applications that prove to be widely used and beneficial. This will not be practical for some more esoteric and specialized applications. However, most of the cyberinfrastructure technologies and many applications should attract industrial interest, and longer term government support for ongoing development and operations should prove unnecessary.

To the extent suitable off-the shelf technologies are available, they should be acquired and used; researchers and NSF program managers need to be well connected to current and emerging commercial activities and seek alliances with them as appropriate. One centralized activity should systematically choose and license commercial solutions to avoid multiple (and incompatible) choices and to obtain favorable licensing terms. For example, prototypes and experimental results may originate from self-supported activities in industry as well as from NSF-supported researchers.

The information technology researchers participating directly or indirectly with this ACP have limited ability to integrate, maintain, and support their own research prototypes. This is the primary role of development organizations that start with these prototypes and end with an integrated and supported software distribution. It is generally healthy to consider alternative approaches and encourage competition among prototype solutions before choosing one to develop and deploy.

The intermediate- to long-term goal should be to commercialize all infrastructure and many application solutions developed in this ACP that are successful and gain a significant following, and withdraw from those that don't. It should not be necessary for NSF to support the development activities indefinitely in any particular area; rather, the goal should be to migrate those development dollars to new areas. Thus, a growing portion of the supported infrastructure and applications are expected to be off-the-shelf commercial technologies licensed with financial support from NSF, with NSF funding for prototyping and development continually redirected to the moving frontier of new (non-commercially supported) capabilities.

Infrastructure suffers from a 'chicken and egg' conundrum in the commercial world: Which comes first, the infrastructure or the applications? It is difficult to invest in new infrastructure with no applications available to provide value to users, and application investment usually follows existing infrastructure. Following the inspirational example of the Internet, this ACP seeks to use NSF investment coordinated across both infrastructure and complementary applications to 'jump start' new commercial markets, and later move those applications and supporting infrastructure together into commercial practice. The ultimate beneficiary will be not only the science and engineering research communities, but the U.S. economy and industry as a whole.

Participation by Other Agencies and Governments

The nature of the ACP is that it should provide value to all science and engineering research, regardless of whether it is funded by NSF and regardless of whether it is conducted in the U.S. or abroad. In fact, the more universally these applications and supporting infrastructure are adopted, the more value they offer to each participant. Thus, this should be viewed as a government-wide initiative and include strong international cooperation. NSF will be the leader, but it should seek broad participation by others.