Document: internet2-crcc-report-200607.html

Comments to Steve Olshansky
<steveo@internet2.edu>

Authors:
Ken Klingenstein, Internet2
Kevin Morooney, Penn State University
Steve Olshansky, Internet2

Rev. 17-July-2006

Final Report: A Workshop on Effective Approaches to Campus Research Computing Cyberinfrastructure

April 25-27, 2006 Arlington, VA

Sponsored by the National Science Foundation - Grant No. OCI-0627970,
Pennsylvania State University, and Internet2

Introduction

Cyberinfrastructure has become a key enabler for scholarly research. Faculty and researchers are becoming increasingly reliant on a mix of high-performance computing and communications (HPCC) hardware, software, networking, virtual organizations, and key research computing support professionals. To help develop a greater understanding of the key campus challenges in cyberinfrastructure, NSF sponsored a workshop developed by Penn State, with assistance from Internet2, in April, 2006. This workshop brought together a combination of CIOs and high level campus technical representatives – CTOs and others with similarly broad responsibilities – to share approaches and common problems, and to strategize about ways in which they would be able to improve their respective institutions' support for the demands of current and future research computing. Attended by almost 70 people, representing 40+ US research universities, NSF and Internet2, the workshop was well received and feedback to date indicates that it was highly valuable to the participants on several levels.

The key findings from the workshop include:

- Campus IT infrastructure roles are varied and vital to the support of overall national cyberinfrastructure, including data center support, campus-based system administrative

support; specialized local and regional networking, provision of campus computing resources, provisioning authentication, authorization and virtual organization services, etc. Scalable and consistent approaches to these needs and those related to them appear to all rely on such campus efforts.

- Major negative reinforcements exist in the current environment. For example, grant solicitations at several major funding agencies seem to favor "autonomous, small clusters in closets" over more sustainable and secure resources. Personal lambdas are sought by scientists less for their performance needs than to obviate the "friction" of campus network security. Lack of coordination of Institutional Research Boards (IRBs) creates major obstacles to inter-institutional data collaborations. Campuses may be legally liable to continue to provide services that were committed to by researchers who have since left these institutions; such situations can dissuade a campus from future commitments.

- There is increasing conflict between research computing systems and campus security procedures. Capabilities available to researchers are being limited by campus concerns about the vulnerability of research systems to attack, and the exposure of personal and sensitive information stored on them.

- The profile of researchers who use campus research computing cyberinfrastructure seems different from those who use national resources. Many campus research cyberinfrastructure users do not need massive computational resources. Most often, their critical need is in data lifecycle management, helping to mitigate the problems caused by transient graduate student data managers. Other needs are for readily available if not high performance cycles, and relatively primitive, but local, visualization systems over some high-end but inconvenient facility.

- There are several campuses which are very active in supporting local research cyberinfrastructures, but even they are challenged with chronically insufficient resources to support their researchers, and do not perceive that their problems and input are being adequately recognized and addressed by federal funding agencies.

- Among these leadership institutions, there is a considerable amount of undocumented but seemingly common approaches to provisioning research computing services and resources. Those range from which architectures are most appropriate for selected categories of computational algorithms (and even software license structures), to management options for shared storage infrastructures. Documenting and sharing this knowledge would clearly have value.

- The absence of planning guides that reflect the business cases developed by the leadership institutions is a problem that could be addressed by white papers, which could map issues and identify the opportunities and the challenges. Such analyses are needed for both partners in a campus cyberinfrastructure: the central administration and its IT organization, and the research communities that need to understand the benefits of partnering with a point of

"institutional drag."

- Among campuses that are less active or those that do not provide central IT cyberinfrastructure support, there is increased recognition of the need for cyberinfrastructure development. This introduces additional questions:
  - By what mechanism, and from what source(s), will this work be funded, particularly during the initial construction?
  - Who is concerned about this, both among campus and national research communities, and the funding agencies; are they concerned "enough?" Who will drive the national discussion and local instantiation of that national discourse?

- There are apparent benefits to the coordination of cyberinfrastructure development at national and institutional levels. Computing jobs frequently move from campus to national centers and vice versa. Data provisioning for large data sets, as exemplified in the Large Hadron Collider (LHC) activities, should be engineered carefully using multiple tiers in their design. Authentication and authorization for national resources need to leverage and integrate with campus identity management. Improved communication mechanisms need to be established and nurtured in all of these areas.

The workshop website, with links to the proposal, agenda/presentations, and roster, is at: http://middleware.internet2.edu/crcc/.

The CIO and Campus Cyberinfrastructure

The first segment of the workshop focused on perspectives, both technical and policy, for CIOs. The CIO participants were particularly interested in the following:

- Campus cyberinfrastructure is not just about the technology. We need to understand and engage the research community, bridge the cultures, enhance the collaborative relationships on campuses and between campuses, and learn from each other. What is the process by which the workshop participants can best continue sharing and collaboration among this community? How can we best interoperate and integrate among campus and national cyberinfrastructure efforts?

- We need to address the primary drivers and value propositions for campus involvement in research support – focusing especially on funding, balancing priorities, and sustainable support. Value propositions need to be established for both campus administration and for researchers. What are useful campus techniques or incentives to encourage faculty to think institutionally as they craft their proposals?

- On campus, an increasingly complex mix of central and distributed components is a given – how can we best balance and manage in this environment, and provide the support needed in a sustainable manner? Safe, secure, cost-effective operations – these are primary institutional goals that must be approached to place the fewest limits on researcher capabilities.

- How can we most effectively move beyond campus boundaries in inter-organizational

collaborations, while minimizing any possible friction introduced by campus boundaries? How can we trust remote users as much as trust those on our own campuses? Can we use emerging federations of trust for our purposes, or are new federation structures needed?

- At several funding agencies, review panels do not consider the security and sustainability of research computing resources that may be awarded, nor do they factor in the development of shared, leveragable resources. The funding model we often see – small, investigator-initiated grants – feeds the problem of autonomous clusters hidden in closets, which in turn makes support issues more challenging. It is worth considering a holistic view that would promote larger sharable, campus systems. The campus is in fact a logical nexus for the development of cyberinfrastructure. Along with other NSF programs that span disciplinary areas, how can NSF most effectively build cross-cutting campus cyberinfrastructure programmatically?

- Improved investment in federated Identity Management (IdM) would significantly benefit the research and education community at large, encouraging national development and implementation of enterprise IdM in support of research computing. Also, as Virtual Organizations (VOs) become a more useful construct beyond big team science, there is a need for dynamic management, e.g., the ability to dynamically create VOs without unnecessary friction, and provide the appropriate access controls and authentication, as needed. One step in providing more effective support to VOs could be a clearinghouse that can contain, and distribute VO metadata for campuses who want to provide institutional-level support services.

- In some ways the federal funding model appears to be going into a set of high-end initiatives, such as building petascale machines and grids, without a similar focus on important, if more pedestrian (and hence more tractable) needs. The most urgent needs for researchers on campus, especially considered across the major science funding agencies, appear to be data management, improved security including identity management, and convenient resources such as computing and visualization platforms more than massive high-performance resources.

- Long term, we plan to create a sustainable campus funding model for research computing. Individual grant funding tends to be episodic and unpredictable, researchers switch institutions but institutional commitments remain, creating serious support issues for the university while increasing the possibility of researchers deprioritizing investments that are not science-critical. Campuses could establish some baseline service offerings, particularly in those areas (such as data management training, and authentication and authorization) where there is little campus expense and considerable researcher benefits.
- We want and need to be able to leverage each others' ongoing efforts, and learn from each other. All of us must address cost-recovery components, and integration of our campus identity management systems into virtual organization support services. Common models and shared wisdom are essential to an effective national cyberinfrastructure.

In addition to these recommendations, this report highlights a roadmap of campus issues that identifies key themes emerging from the workshop, a list of potential leverage points to identify

possible opportunities for community action, and a map or inter-institutional issues which describes major challenges and approaches to interrealm sharing of resources.

Roadmap of Campus Issues

1. Planning

There are a number of major issues to be considered in planning for campus efforts. While an exhaustive list was out of scope for the workshop, some key issues were discussed.

- Making a business case and negotiating commitments, factoring in the needs and wishes of campus administration, central IT, and key users/communities
- Moving from strategic commitment to tactical plans: lining up required seed funding, implementing infrastructure improvements, and encouraging faculty/researcher involvement.

Some specific issues to consider:

- **Shared computing facilities and data centers**
  Almost all institutions report that shared computing facilities are perceived as mutually beneficial once engineered. The level of those facilities varies greatly, from simple data center space to shared condominiums where researchers install blades or disks. The researchers realize greater computing access than stand-alone systems (the whole is greater than the sum of its parts due to load balancing); departments cite the absence of machine room costs and the chance to transform that space into offices or labs as benefits; campus facilities management appreciates the reduced energy, cooling and remodeling costs; security is improved; and the campus enterprise is more robust.

  The leadership institutions reported the return of motor pool computing – providing a variety of computing vehicles in order to match needs to resources. Some jobs need large memory; some need fast I/O, and some need the ability to run multiple jobs simultaneously. In some cases, software licenses for scientific packages are priced per CPU, effectively eliminating some architectures.

- **Connectivity**
  Fiber access is often a costly challenge outside major metropolitan areas, if not in procuring it then in lighting it. The growth of Regional Optical Networks (RONs) is certainly a positive step, but situations vary greatly depending upon the specific circumstances or location of research facilities.

  In many cases, campus networking costs are among the hardest expenses to specifically allocate to or recover from a project, and also among the most important and costly. At the same time, connectivity is a key component of cyberinfrastructure, and one which needs to be upgraded to meet the current and anticipated demands of researchers, and of the campus as a whole.

- **Evolving funding and cost recovery models**
  Given the nature of HPCC and the rate at which it is expanding, it is important to identify cost-effective funding models that serve to bootstrap or bridge support for central resources to be utilized by multiple projects, but which can be difficult to fund from individual project grants. In addition, taking advantage of broadly leveragable resources and economies of scale can and should be a factor in central planning, to the extent that this is practical. There is growing desire to be able to amortize these longer-term investments across multiple projects, if this can legitimately be factored into project budgets.

  The workshop resulted in recognition of the need for a better cost recovery model for attendant support and infrastructure not directly specified in project budgets (e.g., how are robotic storage and backup tape systems funded, since they span across multiple projects?).

  Two examples of storage cost recovery strategies were presented:

  - Some campuses rent storage capacity along with support, upgrades, etc. (typically on a per/TB basis), to PIs/projects rather than have them use their funding to purchase actual disk drives which decline in cost as they increase in capacity and regularly become obsolete. This approach allows service levels to be maintained, and provides ongoing revenue streams to central IT for expansion, support, and upgrades. Ultimately this model has some benefits as allocated funding is more effectively utilized, central IT can take advantage of economies of scale, and storage capacity can be more quickly adjusted to project requirements – whether up or down – over the life of the project.
  - Central IT pays for core storage infrastructure – centrally managed, high transfer rate, etc. – but researchers pay for the actual disks they acquire that plug into the disk farm.

- **Campus IT Policy development**
  Policy and related enforcement issues on campuses can be particularly challenging to some scholarly resource sharing, and research requirements should be addressed early in the campus planning process and continually tweaked to meet evolving needs. In many cases, policies are much more effective when they are not developed in isolation from the research community and imposed upon them by central IT, but rather created collaboratively and facilitated where appropriate by education and outreach mechanisms. There is a clear benefit to building trust and partnerships among the community, which bear long-term benefits to all concerned. There is significant value in facilitating effective communication in both directions, and developing strong relationships among researchers and campus IT professional staff.

  Campuses may also want to consider reward systems to encourage researchers to think institutionally. Space, budget, recognition, and other incentives may help.

2. Development and Implementation

Workshop participants recommended the following areas of focus for development and

implementation:

- Facilitating research through provisioning of resources, services, identity management and virtual organization services, etc.

- Creating effective partnerships with researchers, Institutional Research Boards, network managers and security professionals, etc.

Again, the full breadth of development activities were beyond the workshop scope, but some specific issues emerging from the workshop included:

- **Inter-organizational collaboration**
  Many projects are using Virtual Organizations (VOs) as tools to bring together researchers spanning multiple institutions, bringing several problems to the forefront: how to enable the required Authentication, Authorization, and Accounting (AAA) without introducing more friction than necessary? Interrealm/federated trust mechanisms would seem to be an obvious solution, but not enough institutions are currently equipped (in terms of having enterprise directories in place, the MACE/Internet2 Shibboleth System, or participation in a trusted federation structure), to currently provide these services and progress is difficult. Yet the practice of VOs issuing individual credentials to remote users is becoming too burdensome, both for the resource holders issuing them and the users trying to manage them. Furthermore, campus compliance and audit requirements are not supported by these external and ad hoc approaches, yet another source of friction.

  Example issues include:
  - The portability of the data
  - Ways to address regulatory compliance (e.g. HIPAA/FERPA) when users of protected data are outside your institution and determination of who is responsible
  - How to manage overlapping VO memberships, including campus and PI level
  - Resolving or mapping relevant policies between institutions.
  - Deployment of enabling tools such as the Shibboleth system, myVocs, and GridShib

- **Architecture design**
  Computational needs are often subject to change, emphasizing the need for architectural approaches which are as extensible, scalable, and as cost-effective as possible in order to permit more efficient aggregation and utilization of resources. Scarcity is driving ingenuity on some campuses – e.g., pooling HPCC resources in a condo model, harvesting short-term cycles when available to supplement guaranteed resources, and building grids to avoid expensive systems sitting idle when they could be doing productive work on other projects. Running counter to this desire among central IT is the need by some/many projects to utilize an architecture optimized to their particular requirements. This inherent tension may require creative compromises in order to best meet the needs of both sides.

  Since different systems architectures often require code to be developed in such a way as to take best advantage of the resources being utilized, it is critical for both sides to establish

effective communication as early in the development process as possible. Some architectures have relatively steep learning curves, requiring particular attention to support and education components.

Many campuses are working to implement diverse architectures, trying to proactively determine what their particular user community will want, and provide it to the extent this is practical. Solutions that emphasize sharable architectures, rather than individual designs with little likely application beyond the project at hand, are more likely to be effective in that they are serving as catalysts for interdisciplinary projects and serve the needs of their users now and in the future. Being able to take advantage of economies of scale and to build systems that can be more easily repurposed for other users is particularly important in this arena, since the systems and components are often so costly.

In some ways there are conflicting goals to be resolved: making big users more productive versus. making high-end computing facilities easier to use, creating more demand for those scarce resources.

- **Physical and Logistical support represent more significant problems than cycles**
  The relative useful lifetimes of HPCC resources is short, compared to telescopes or other large scientific instrumentation. Cost recovery is a difficult problem, along with the related issue of how to fund infrastructure and support. In many ways, there seems to be an inherent disincentive to move toward implementing economically feasible and scalable data centers, even as the need for quality physical environments becomes more apparent.

  Long-term sustainability of required support infrastructure is quickly becoming one of the most challenging issues, particularly in the unpredictable sponsored funding environment in which most campuses find themselves. This includes, but is not limited to, storage, networking, data/network security, support staff (including training), performance tuning, troubleshooting, problem solving, backup, disaster recovery, data-center infrastructure (power, A/C, physical security). PIs are pushing central IT to provide support for their needs, often beyond the funding available in their respective grants. In some cases central IT may not even be aware of all of the clusters on their campuses, given how inexpensive and relatively easy they are becoming to do the initial setup (especially when the vendor is still in the room) – if not to optimize or support effectively over time. On occasion, central IT is finding itself competing with faculty for grant funding. A growing challenge faced by many campuses is how to best avoid this friction, or at least the perception of it.

  Bootstrap funding is often overlooked and difficult to obtain, in that many of these support structures need to be established prior to deploying specific computational resources specified and funded in a particular project. Creating the infrastructure required to effectively support HPCC needs to be planned and implemented in advance of large projects, so that central IT will not be caught in continuous catch-up mode, or that the PIs and their projects suffer, or at are not as efficient and effective as they could be. As more campuses address this issue, a planning guide will be essential.

- **Education and outreach**
  Education and outreach are ongoing challenges, and much like security, must be addressed as a process rather than an end goal. Researchers don't need or want to become computer scientists or security specialists, but outreach efforts to them can be very valuable in facilitating effective communication. Specific examples derived from workshop participants include:
  - Researchers tend to be autonomous and instrumental in their focus, in contrast to central IT which by its nature takes the broader systems view, which in turn brings conflicting culture issues into play.
  - In an environment in which particular hardware architectures have specific requirements in order to most effectively utilize their capabilities, educating users/researchers /support staff about how to structure new code to maximize use of particular resources is becoming ever more important.
  - Security awareness training for research IT staff is critical, especially since the security field is constantly shifting and staff and researchers are transient.
  - VPs for Research (VPRs), Deans, Department Chairs, and Provosts are often not equally informed when it comes to these sorts of issues, which is not surprising given their particular roles and priorities. There needs to be a coordinated effort to engage them and enable ongoing communication channels, for the benefit of all concerned. This would include the identification of current loci of policy and decision authority affecting this space and assessment of how this is working, and determination if subunits, such as departments and colleges, represent the best institutional level for such decisions.
  - It is desirable to see an education component, training undergraduates and post-doctoral fellows. Discussion focused on integrating integrated into the curricula and the role (if any) that central IT play in this.

  There needs to be enhanced understanding of the appropriateness and benefits of using supercomputing and national scale centers, as opposed to local HPCC resources. Depending on the particular project, in many cases it can be more productive to use local resources, because of better accountability, availability, and the ability to control and customize the resources to the particular needs of the project. Some wider awareness of the options for research should be encouraged, especially for new researchers.

  In planning outreach and education in this context, it is important to recognize and emphasize the public good to the research community and to the campus, and the often indirect ways in which researchers will benefit. "All gain from raising the water level..."

- **Assessment**
  Ongoing evaluation and assessment are essential to enable long-term goals. Much of that assessment must come from the users and stakeholders in order to ensure they have a strong voice in the process. In particular, as researchers come to appreciate the impact of their decisions and processes, both on the central infrastructure as well as on fellow researchers, they will be are more able to participate as active partners and ensure the success of their

projects.

Given the rapid rate at which technology is evolving, and as new projects come on board and old ones retire, building flexibility and a strong feedback mechanism into the planning and deployment process will be increasingly crucial.

- **Security and data management**
  More capabilities, and significantly larger quantities of data are being generated (much of this sensitive and regulated), which brings attendant security problems to the fore. Security and appropriate management of this data must be addressed, and supported, as a process rather than an end goal. Example issues include:
  - Responsibility for the management and security of the data (The PI, university, or sponsoring/funding body?)
  - Individual servers run by individual faculty or small groups – which are generally not centrally managed and in some cases not even known to central IT and which can create one of the largest security threats to a number of campuses. Moving sensitive data from central IT supported systems to these local systems often creates huge security risks
  - What restrictions do Institutional Review Boards (IRBs) place on the data?
  - Network isolation, which can help in some cases, also can discourage broader access and availability while ignoring physical security issues in the data center.
  - Encryption – including appropriateness , and management of keys
  - Foreign students – managing access to restricted data or systems
  - Local or remote data generation and the need for access by remote users collaborating with PI
  - Developing tools for managing groups and privileges, e.g. Internet2 Grouper™ and Signet™ projects
  - Visualization – important to many researchers, but costly to implement well for broad use
  - Data protection and the challenges for central IT being familiar with the kind of data it may be held responsible for protecting (not having a medical center does not mean that there are no HIPAA issues to be aware of – patient-identifiable data may exist without central IT being aware of it). Processes need to be developed with these sorts of challenges in mind, e.g. asking about HIPAA or other regulatory issues when researchers provide data to be managed.
  - The natural tension between security and the needs of an open research community. Campuses must balance between the somewhat orthogonal goals of flexibility, security and performance.

Security reviews are becoming more common, but at the same time more challenging to perform, especially as they span multiple departments and multiple institutions. The environment is becoming increasingly complex, with more devices to protect, and more potential threat vectors enabled by this complexity. Similarly, in these distributed computational environments, more users means more security awareness training and more complex policies and procedures required.

Potential leverage points

Potential opportunities for community action emerged from discussions:

- **Coordination of input to NSF on topics of shared interest**
  There was agreement among participants that they want to see an organized means by which the campus IT cyberinfrastructure professionals could provide input and proposed alternatives to NSF, on a range of topics affecting their community and for which NSF policy can or does play a central role. Examples could include, but are not limited to:
  - Observations about consequences of the NSF funding model, including the individual grants that feed the problem of autonomous clusters hidden in closets
  - NSF's proposed layered model – workgroups, campuses, petascale machines , and suggestions regarding where does central IT fit into this model
  - The cyberinfrastructure needs of typical campus researcher, and recognition that while the petascale machine and large scale grids may meet some large modeling/simulation projects, this represents a relatively small percentage of the campus researcher base, especially when viewed cross-agency. Often the most pressing needs are in data management and simple visualization tools.
  - Encouraging NSF to support pure storage acquisitions in addition to computational resources NSF storage guidelines
  - Working with NSF and campus leaders to promote campus development and facilitate culture change in how researchers view campus infrastructure.

- **Shared libraries of key domain applications ported to various architectures**
  Since various systems architectures require specific software optimization in order to perform to their fullest potential, there would be broad benefit in maintaining a library accessible to researchers, to facilitate leveraging previous work, learning from others, and eliminating duplication of effort.

- **Ongoing collaboration among campuses**
  Workshop attendees were enthusiastic about this opportunity to network with each other; to share stories; and to share model policies, SLA and MoU templates, suggested approaches (e.g. for education, outreach, support, etc.), and common problems/solutions. Consensus was that there is a real need for continuation of this collaboration, which could take many forms. Early, relatively simple steps in this direction could include wikis and mailing lists dedicated to various topics of broad interest. Further opportunities to meet in person, while more difficult to coordinate, seem to be of interest as well. These could include meetings arranged in conjunction with meetings likely to be attended by at least subsets of this group, such as SC, Internet2 member meetings, Committee on Institutional Cooperation (CIC) meetings, or other venues attracting a similar audience.

  Given the extensive presence of medical research, there was also interest in closer interaction and collaboration with the Association of American Medical Colleges (AAMC).

- **Shared vocabularies/semantics**
  Supporting the broad adoption of relevant models and vocabularies, both technical and policy, can be a significant step toward enabling inter-realm collaboration and related access control systems.

- **Facilitating inter-organizational collaboration**
  As inter-organizational collaboration becomes more common, there are a number of ways in which campuses can work locally toward a common goal of facilitating federated IdM. Examples include the sharing of business cases, approaches, policy frameworks and templates, and working to overcome local barriers to adoption of supporting infrastructure – e.g., the Shibboleth system.

  A community effort to encourage the adoption of technologies and policies to enable collaboration would be a benefit to the community as a whole.

A Map of Inter-institutional issues

Federated trust, and in particular the facilitation of virtual organizations spanning organizational boundaries, is rapidly becoming a top goal of a broad spectrum of campuses and related service organizations.

Challenges include:

- PKI – key management, certificate authorities, policies – ongoing obstacles to wide adoption.
- Credential management – consistent Identification and Authentication (I&A) policies among collaborating organizations. Credential revocation and retirement policies also come into play.
- Deployment of required infrastructure at the campus level, e.g. enterprise directories and the Shibboleth system. Increased investment in federated IdM by funding agencies could facilitate adoption considerably
- Controlled vocabulary or common schema (e.g. eduPerson) particularly in the definition of roles used for role-based access control, and for attributes required to make informed access control decisions
- Determination of the role of federations, including agreement on required levels of assurance (LoAs) for specific application classes
- The best ways to convey VO membership consistently across all resource owners, while managing overlapping VOs and determining authorities for VO attributes
- Enhancing federal agency cooperation – there would be value in bringing together NASA, DoE, NIH, and NSF to the extent practical to agree on a common set of policies that would in turn facilitate the sorts of large collaborative projects that seem to be the wave of the future.

Resources, funding, and support assume different dimensions in this context, including ways to deal with users or resources beyond your borders

- Each organization is set up to support its own user community. In collaborative environments, challenges include the possibility or practicality of dedicating a portion of local resources for

inter-institutional projects
- The political issues raised by external users using local resources when seeking support from central administration

Regulatory issues
Boundary conditions tend to influence approaches. Frequently, technology precedes policy, leading to some potentially awkward issues but making a case for a new paradigm for regulatory efforts. Examples include:

- Individuals and/or enterprises as creators of digital signatures, which can encourage different approaches depending upon the context
- The inconsistency of interpreting regulations (e.g. HIPAA, FERPA) and technical drivers (e.g., data may fall under IRB regulation, and yet researchers may need to share it on a data grid.) Issues include data management in this sort of open environment and the value of more detailed and consistent guidance from relevant federal agencies about how best to implement the regulations
- Determination of who is responsible for data security in inter-realm collaborative environments and identification of the data "home"
- People are trying to collaborate, yet technology and policy can interfere (e.g., legal staff tend to be worried about university indemnification, especially for outside users. It is clearly better to have lawyers who are IT-savvy, if possible.) Issues include ways to educate the legal and regulatory communities and encourage effective communication and whether this is the appropriate leverage point, given the significant funding at stake
- The need for IRB offices to work together to facilitate collaboration, if the project/funding is there to drive it, and corresponding strategies to facilitate this and make it an efficient process

Acknowledgements

This report was authored by Ken Klingenstein, Internet2; Kevin Morooney, Pennsylvania State University (Program Committee Chair), and Steve Olshansky, Internet2. Special thanks for contributions by Jim Davis, Iowa State University, and program committee member Patrick Dreher, MIT.

References

- Association of American Medical Colleges (AAMC) http://www.aamc.org/
- Committee on Institutional Cooperation (CIC) http://www.cic.uiuc.edu/
- eduPerson schema http://www.educause.edu/eduperson/
- GridShib http://gridshib.globus.org/
- Grouper™ – groups management http://grouper.internet2.edu/

- InCommon Federation
  http://www.incommonfederation.org/
- Internet2
  - Middleware
    http://middleware.internet2.edu/
  - Networks
    http://networks.internet2.edu/
  - Security
    http://security.internet2.edu/
- myVocs – Virtual Organization Collaboration System
  http://www.myvocs.org/
- SC
  http://sc06.supercomputing.org/
- Shibboleth®
  http://shibboleth.internet2.edu/
- Signet™ – privilege management
  http://signet.internet2.edu/
- Workshop on Effective Approaches to Campus Research Computing Cyberinfrastructure
  http://middleware.internet2.edu/crcc/