User Support and the Virtual Machine Room

Scott Koranda
skoranda@ncsa.uiuc.edu

Bruce Loftis
bruce@ncsa.uiuc.edu

John Towns
jtowns@ncsa.uiuc.edu

National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign
605 East Springfield Avenue
Champaign, IL 61820

ABSTRACT: NCSA and the National Computational Science Alliance are moving toward a Virtual Machine Room (VMR) that will seamlessly connect many distributed computational facilities, data archives, virtual reality facilities, and even large scientific instruments using high-speed networks and the new "grid" technologies. The grid and the Virtual Machine Room will radically transform the way computational scientists work. The VMR offers new challenges and requires new models and new technologies for supporting users.

1. THE ALLIANCE VIRTUAL MACHINE ROOM

The National Center for Supercomputing Applications[1] (NCSA) and its partners in the National Computational Science Alliance (Alliance) are currently deploying the infrastructure necessary to establish the Alliance Virtual Machine Room[2] (VMR). The goal is to make the many resources available at the distributed set of sites within the Alliance appear to be managed in a single coherent fashion as if they were co-located and under the direct management of a single center. This distributed infrastructure is a model that also maps well to a number of efforts along more traditional high-end computing efforts, such as NASA's Information Power Grid[3] and the US DOE's Advanced Strategic Computing Initiative[4] (ASCI). These efforts are providing insights into the solutions needed for distributed research environments that will emerge as the collaborative spaces in which science will be conducted in the future. These grid ecomputing environments will be likewise interconnected to support interdisciplinary and interagency research. The initial planning and work was the topic of an invited talk at an RCI meeting, which shows industry interest in developing and integrating their global sets of distributed resources.

The Alliance, one of two partnerships funded by the National Science Foundation's Partnerships for Advanced Computational Infrastructure (PACI) program, has as a mission to prototype an advanced computational infrastructure for the 21st century and includes more than 50 academic, government and industry research partners from across the United States. Our current near-term plans provide for establishing the VMR to support the PACI user community with compatibility and extensibility to interact with other distributed environments of various types.

In this paper we outline a variety of issues related to how users interact with such a computing environment and what strategies are being developed to support computational science and engineering research in the Alliance Virtual Machine Room.

2. THE USER'S PERSPECTIVE

The Alliance VMR dramatically transforms how computational scientists perceive and utilize their high performance computing (HPC) environment. Below we detail allocations, job submission and monitoring, authentication, and remote data access and storage in the VMR. We further discuss accessing the VMR through a web portal interface.

2.1. Allocations

Traditionally, for Alliance resources, allocations of computing time are site and architecture specific. The allocation currency is a service unit (SU). The number of SUs that a particular calculation might cost against a particular allocation is based on the CPU time used, along with other factors depending on the site and resource used.

In the first major phase of the VMR, allocation will continue to be specific to resources at particular sites. During this phase the infrastructure to support less specificity will be developed, but users will benefit from a single interface to submit jobs to any particular resource within the VMR.

During the second major phase of the VMR, allocations of computing time are still architecture specific, but not site specific. Scientists receive SUs for computing on a particular architecture, but the site of the resource is irrelevant. Users compute on any machine of the architecture type for which they have an allocation, and charge their cycles to the single allocation. Accounting happens centrally, and users monitor their architecture specific allocation rather than separate allocations at separate sites.

In the third major phase of the VMR, allocations of computing time will be both architecture and site independent. Scientists will receive SUs for computing on any Alliance resource, and the translation of CPU cycles along with other factors to SUs will be normalized for each particular resource, producing a usage amount in a unit defined for the entire VMR environment, a Grid Service Unit or GSU. The normalization depends on a variety of details including cost per cycle and demand for computing time. Codes that efficiently run on more than one architecture spend less time waiting in the queues since they are more readily dispatched to a host with available resources.

As more resources in the HPC community are connected to form additional grids, the Alliance expects to pursue exchange agreements with other grid environments. In this third major phase of the VMR, scientists might transparently run jobs and spend their allocations on resources external to the Alliance grid. With the rising and falling demand for cycles, exchange agreements between grid environments will allow resources to flow from grid to grid, enabling greater resource utilization and throughput. The reader might recognize the similarity of the described scenario with electrical power grids, and thus the choice of terminology.

2.2. Job Submission and Queuing

Today scientists queue jobs on Alliance hosts by logging in and submitting a job using LSF, PBS, or whatever batch system is installed locally. Queuing jobs at more than one Alliance site requires scientists to master the details for each machine's batch system, and also the local queue policies. Queued jobs run only when necessary resources, including memory and processors, are available on the host on which a job is queued.

During the first phase of the VMR, scientists need not master the details of each machine's batch system. Instead, when submitting jobs they specify job resource requirements using a single common interface, regardless of where the job is to run. This interface, built on the infrastructure of the Globus Metacomputing Toolkit[5], maps job resource requirements to local batch system commands when queuing the job. During this phase users still need, however, to specify a particular host on which to run, and each job is limited to the resources available on that specific host. In addition, scheduling of queued jobs is still completely determined by local policies.

Users will see more dramatic changes in job submission and queuing during the second major phase of the VMR when it truly becomes "virtual". Using the same Globus job submission tools, users will continue to specify job resource requirements, but will no longer specify particular hosts on which to run jobs. Instead users will simply specify the architecture on which their job can run, and the VMR scheduler will route the job to be queued on the host with that architecture that has or will have the necessary resources available the soonest. This basic global scheduling across homogeneous architectures promises to substantially increase resource utilization and throughput performance.

In addition to the above changes, users will also be able to exploit resources at more than one site across homogeneous architectures in the context of a single job. Co-scheduling will allow distributed applications that are capable of simultaneously using resources at more than one site to do so. Further, the coupling of global and local scheduling will allow for advanced reservations to be made for necessary resources.

The VMR becomes completely virtual during the third major phase when users are able to queue jobs for any resource in the Alliance. Codes which can exploit more than one architecture will simply be submitted with a list of job resource requirements, and the job will be dispatched to whichever host has the necessary resources available the soonest. This should dramatically increase resource utilization and computational throughput. Making this possible will require tightly coupling global and local scheduling, along with a sophisticated routing infrastructure. Exchange agreements with other grid environments will allow users to transparently run their jobs on machines located anywhere within an accessible grid, further increasing throughput, and allowing users to spend much less time waiting in the queues.

2.3. Authentication and Security

Strong security for the Alliance is a top priority, and thus from the initial deployment of the VMR users will be required to access the Alliance Grid using the Alliance Public Key Infrastructure (PKI)[6]. The Alliance PKI is built upon the Grid Security Infrastructure (GSI)[7], a component of the Globus Metacomputing Toolkit.

The GSI package uses X.509 v3 digital certificates for authentication, which is a common and established international standard for digital certificates. Alliance certificates are issued by a central certificate authority[6], and are accepted for authentication by all Alliance-allocated computational resources. In addition GSI supports "single sign-on", allowing users to authenticate once "to the grid" and then navigate the grid, transfer files, and submit jobs without having to again enter a pass phrase. GSI-enabled versions of common tools such as SSH and FTP are available and allow users to navigate the grid using tools they already understand.

Later, VMR users will be able to use their Alliance digital certificates to authenticate to other grids, for which agreements have been made. Likewise, Alliance computational resources will accept digital certificates issued and signed by other certificate authorities.

2.4. Remote Data Access and Mass Storage

A number of Alliance sites currently have mass storage resources available to users, including UniTree, HPSS and ADSM. During the first major phase of the VMR users, and most importantly their jobs, will have remote access to these mass storage systems. A job running at any Alliance site will be able to fetch necessary input or data files from the mass storage systems, and then later push files back to the mass storage. A common set of commands interfacing to each of the different systems will make it unnecessary for users to know the details of each system. More enhanced services will become available later with data staging, pre-fetching, and post-store for jobs.

The VMR again becomes more virtual during the second major phase with the implementation of a single virtual Alliance mass storage for all Alliance users. Users will access this virtual mass storage using an interface or set of commands with no reference to where the data is actually stored. The details of where data is to be stored and how it is to be moved will be based on access requirements and transparently handled by the infrastructure. With the third major phase of the VMR this virtual access will be extended to provide access to data for jobs that run outside the Alliance on other grid environments.

2.5. Web Portal Interface

In the preceding sections we have detailed how the Alliance VMR will dramatically change the way computational scientists interact with their HPC resources, and have considered the full spectrum of the user's perspective ranging from allocations to security to job submission and control. None of these, however, will affect and increase productivity for computational scientists as much as the VMR web portal interface.

The web portal will provide a graphical, seamless, and integrated interface to the VMR, enabling people to readily exploit VMR resources. Since it is a *web* portal, users can interact and monitor their jobs from almost anywhere, using their favorite web browser. The graphical interface allows users to point and click to monitor a job or query the state of a machine, rather than having to login and remember what options are necessary for the command line versions of the tools.

The first version of the VMR web portal will offer simple tools for job submission, job monitoring, resource and queue status queries, grid and network queries, usage monitoring and allocation management. Using familiar elements of a web form like text boxes, radio buttons, and pull-down menus users will specify what program to run, the location of any input or data files, the location of where output is to be stored, and the details of how the job is to be submitted such as what architecture and queue, as well as processor and memory requirements. Job monitoring will especially be streamlined. Users need only to click on a particular job from a table showing all their pending or running jobs to see all the details presented in tabular and graphical form. In a similar way scientists will easily be able to click and quickly review the queue status for a particular machine, or monitor the network performance connecting Alliance sites.

The portal also promises to make allocation and project management simpler and more efficient, allowing scientists to do more science with their allocated resources. Simply by pointing and clicking, a principal investigator (PI) for a project will be presented with graphics and tables illustrating how resources have been used, allowing the PI to track usage and plan how best to utilize his remaining allocation. PIs will be able to quickly add or remove people from projects, and manage their projects as a whole. In addition the web portal interface to the Alliance VMR will support a completely electronic process for submission and review of proposals to use Alliance resources for science.

In later phases of the VMR a more sophisticated portal with more sophisticated “tools” or components will offer scientists a completely customizable and extensible environment for job submission and monitoring. Each user will be able to customize the portal to streamline job submission for a particular code or codes. Streamlined and customized job monitoring tools will report only the details most desired and useful for each individual user. In addition the extensibility of the portal environment will offer hooks to users who want to *dynamically* interact with their applications via the portal, providing any level of desired interaction with a running code. Advanced portal tools might include a “job turnaround predictor” which would analyze the past performance of a code and, combined with current load and network performance data, predict when a particular job will finish.

Future versions of the VMR web portal will continue to reflect advancements of the VMR and will offer graphical interfaces to advanced services. As an example, agreements with other grid environments might allow users to “browse” amongst a catalog of available hosts, with predictions for cost, turnaround, and similar details. Scientists will be able, with a click of the mouse, to submit jobs to any resource in the catalog based on these types of “market variables”. As a second example, advanced portal tools might allow researchers to “plug” one portal tool into another and seamlessly connect codes producing data with tools for data analysis, data mining, and visualization.

3. NEW OPPORTUNITIES FOR USER SUPPORT

We face a variety of challenges as we consider supporting Alliance VMR users. The staff responsible for supporting Alliance users lives in 4 different U.S. time zones, and while there is some overlap of resources, in general each VMR site has its own distinct hardware vendors. Clearly, each configuration and local environment differs in many ways, and for example we do not expect support staff in Maui, Hawaii to provide detailed support for users of the SGI Origin 2000 at Boston University. Still, fundamental issues facing VMR users remain that each support staff person needs to understand. In the following sections we discuss some of the challenges, issues, and opportunities that we face in making the VMR a valuable resource for users of Alliance high-performance computing systems.

3.1. Distributed or Centralized User Support?

The Alliance Grid and Virtual Machine Room link together many computational resources. Each individual resource, whether it’s a batch host or a mass storage system, has its own unique set of tools, software, and system policies. The various VMR batch hosts represent different architectures, they run with different operating systems, and each offers specific compilers and other software. Likewise, a UniTree mass storage system offers a different interface than does a HPSS storage system. Although some common architectures and operating systems exist across the grid, because of local system administration policies, the user experience with each resource is truly unique at some level.

While the VMR will go a long way in making the uniqueness of the individual resources transparent to Alliance scientists, those staff or consultants who support Alliance users invariably need to look behind the virtual façade to investigate problems in detail. At some point consultants will be asked to investigate the details of jobs running on a specific Alliance host or the details of user data stored on a specific Alliance storage system. Because of the large number of different architectures, operating systems, compilers, tools, software, and the like it is unrealistic and inefficient to expect consultants and those who support Alliance scientists to master the details of every computing resource making up the VMR.

This suggests a model for VMR support that is highly distributed, with staff at each VMR site fielding and handling only those user problems specifically pertaining to their local systems. Not all sites, however, maintain full time dedicated consulting groups, and often those people providing support wear many hats and can only dedicate a small percentage of their time to user support. Since the VMR deployment at a site may significantly increase requests for support for resources available at the site, it is unrealistic to expect the staff at these sites to handle the entire workload.

On the other hand, VMR resources are quite geographically distributed, as are the scientists who compute on those resources. It is not altogether unlikely that a user in the eastern U.S. might awaken early in the morning to check the details of a computation running on a host in the western U.S. (or even a host located on the Hawaiian island of Maui), only to find that the computation has failed for some unknown reason. If the only consultants with expertise and knowledge about a specific resource are those “on site”, then, in this instance, he might have to wait a few—possibly many—hours to consult support staff.

Additionally, we expect that a significant fraction of user problems will not be directly related to resources located at a specific VMR site. For example, a user might have trouble using his Alliance digital certificate and the PKI software to access the grid from his local workstation. Questions about using the common job submission interface will most likely also be more general in nature and not related to the details of any single VMR site.

Taken together, the above arguments suggest a model for VMR support that is more centralized, with 24 by 7 service. Providing this level of centralized support with available human resources, however, presents a substantial challenge. The National Computational Science Alliance is very truly an alliance; VMR sites share common goals and strategies for achieving the goals, but progress requires consensus. Any solution for Alliance-wide VMR user support must recognize and champion the unique makeup and policies of each individual site.

Thus, determining how best to support VMR users requires examining the balance between available human resources, and the goal of providing targeted support during the hours when scientists really work.

3.2. VMR Support Goals

Before determining whether the VMR user support structure should be highly distributed or highly centralized, it is useful to list two basic goals for VMR user support.

The first important support goal is providing a single point of access to consulting support. No user should be confused where to start looking for support. This is especially important for later major phases of the VMR when a user may not know on which physical machine his job is running, or where his data is physically stored.

The second important goal is that the VMR support network and infrastructure be completely transparent to the user. A scientist experiencing a problem with a VMR resource and prevented from making progress with his science, should only know that he has requested assistance and that a knowledgeable and helpful consultant has responded. He should not have to assist his own support by routing his concern from site to site until finding the most appropriate person to help him.

3.3. The Alliance Virtual Consulting Office

Taking into account the desire to provide user support "on demand" during the hours when people actually work, while at the same time recognizing the limitations imposed by the nature of the Alliance and its being geographically distributed, we propose a two-tier model for providing user support for the VMR, and the establishment of the Alliance Virtual Consulting Office (VCO).

The first tier of the VCO would be a 24 by 7, immediate response team trained to provide a "useful" level of support for *all* VMR resources. The NCSA Consulting Group and the NCSA Technology Management Group (TMG) would together make up this immediate response team, with the consultants providing support for high-performance issues and TMG staff providing operations support. During business hours consultants would answer user questions about compilers, parallel programming, the Globus software infrastructure, and the like. TMG staff would handle questions about system availability, network performance, and related operations issues on a 24 by 7 basis.

As we noted previously, it is unrealistic and inefficient to expect support staff to master the details of all VMR resources, and so defining the "useful" level of support provided by the first tier of the VCO is a challenge. We propose that NCSA consultants be trained to answer basic questions about compilers, parallel programming libraries, debugging tools, performance tools, and the like for all VMR resources. We further propose that NCSA consultants be fully trained to provide detailed support for those issues not directly related to a specific site, such as the common interfaces for job submission and access to mass storage. NCSA TMG staff will require training on monitoring the systems and networks comprising the VMR. Implementing this model presents a challenge to NCSA, and will require additional staff for both the consulting and TMG groups.

The second tier of the VCO would be made up of the support staff from the other VMR sites. Issues that cannot easily be resolved by the NCSA consultants or TMG staff, and which involve particular resources at a site, would be routed to the support staff at that site.

While this model for VMR user support does not provide for expert support available 24 hours a day—and so might not immediately help the user in the eastern U.S. having trouble with a system in Maui—it does provide for some level of support at any time of day. We assume that a quick response from a perhaps "non-expert" support person, followed by a later response from an expert, is more beneficial to a user than no response until much later.

We also note that some Alliance partners not providing computational resources to the VMR might still offer their support staff to contribute to some degree to the VMR support effort. We expect that this type of distributed support would be extremely helpful to the NCSA support groups providing the top tier of VMR support, especially with general issues not relating to the details of a specific VMR resource, such as topics in parallel programming like MPI or OpenMP.

Like any consulting office or "helpdesk", users would be able to contact the VCO in a number of ways: email sent to *consult@alliance.edu* would route directly to the VCO, a single phone number would be published for those wanting to speak directly with a VCO consultant, and VMR documentation would be available at a single comprehensive web site[2]. Additionally, a link on the web portal interface would point all users directly to the VCO.

3.4. NCSA Ticketing System

A necessary piece of the infrastructure needed to support the VCO is the NCSA Ticketing System (NTS)[8]. NTS is a web-based helpdesk utility for issuing and tracking support issue or "tickets." Support staff access NTS using a standard web browser, though a Sybase database engine powers the system. The NCSA Information Resources Group (IRG) designed, coded, and deployed NTS, which has been in production at NCSA for over a year. While commercial helpdesk software is readily available, the NCSA IRG group determined through experimentation and testing that none is sufficiently flexible to provide the necessary infrastructure for support at NCSA and within the Alliance. In addition, the IRG group required that the NCSA ticketing infrastructure be scalable and provide inter-site and inter-helpdesk support capabilities. Indeed, any trouble ticket system for the Alliance VCO must smoothly and seamlessly integrate with any existing system at any of the VMR sites.

Consultants create a NTS ticket after receiving email or a phone call from a user. Tickets are then dispatched or assigned to various groups at NCSA or the Alliance, such as the Systems Group or the High Performance Data Management Group. Tickets may, however, be assigned directly to the *dispatching* group. NCSA currently has two dispatching groups: the NCSA Consulting Group and the Technology Management Group. That consultants or TMG staff can assign tickets directly to themselves is an important flexibility provided by NTS. Consultants typically solve user problems without having to route a ticket to a specialized technical group, and so assigning tickets directly to consultants removes an unnecessary layer from the support infrastructure; a separate staff for routing and assigning tickets is unnecessary.

While this distinction may not seem immediately important, it does allow NTS to be scaled up as the ticketing system for the Alliance VCO. A distributed Alliance support team, comprised of staff at the various VMR sites, can use NTS for both dispatching tickets to specialized technical staff, and for directly supporting users. A separate group for the sole purpose of routing VMR trouble tickets is unnecessary.

An additional NTS flexibility, making it particularly scalable to the VCO, is that tickets can be created, assigned, and routed by any consultant. No central authority is necessary. Each consultant may, upon seeing a new email from a user, create a ticket and begin working on the problem. With VMR support staff being geographically distributed, this decentralized workflow model is ideal.

3.5. Specialized Support for Common Software

The VMR provides the opportunity to develop new ways to support scientific software. The first step is an online database of all scientific software available throughout the Alliance. This software repository will maintain important information about each package or library: included are version level, vendor contact, local coordinator for the package, and a pointer to instructions for using the software at the local site.

We also have the opportunity to provide effective Alliance-wide distributed support for some of these scientific packages and libraries. For example, the chemistry community requires help from Alliance support staff. Often this support is less a question of interacting with the local computing environment, and is more a question of helping the user interact with Gaussian98 (or some other chemistry application) to solve a particular science problem. There are Ph.D.-level computational chemists at most of the Alliance resource partner sites. These computational chemists often have different expertise and differing familiarity with the chemistry packages. Using the NCSA Ticketing System, we are establishing a VMR chemistry support group. Relevant tickets will be dispatched to the VMR chemistry group, and the most appropriate chemistry support staff throughout the Alliance will investigate the tickets. We see this distributed support mechanism as an opportunity to provide even better scientific support to the chemistry community. We anticipate expanding and scaling the discipline-specific distributed support to include math libraries and tools and eventually support of structural engineering and CFD codes.

3.6. Supporting The New Technologies

The Alliance VMR and the web portal interface expose users to new technologies like the Globus Metacomputing Toolkit and web portal technologies such as XML. Some fraction of scientists will desire to exploit these technologies directly to further enhance their HPC environment, and will naturally turn to support staff for assistance. While developers might provide some level of support to those scientists extending the VMR using tools the developers have produced, it is unlikely they will have the resources to support not only Alliance VMR users but other grid users as well.

Should then VMR consultants be trained to some level to support VMR infrastructure such as Globus or XML? Most likely yes, all consultants should have some familiarity and be able to provide some level of support, if nothing more than pointing users to the appropriate documentation. A more efficient strategy is to train a group of specialists able to provide specific support for infrastructure pieces of the VMR, in much the same way that professional chemists support chemistry applications across the Alliance, as detailed above. Such specialists might work closely with developers and provide support for users of other grid environments as well. Should "grid computing" become ubiquitous, we expect that commercial companies will be formed to meet the demand for specialized "grid consulting".

3.7. Desktop Data Sharing

Providing support for users accessing the VMR through the web portal interface is especially challenging. Traditionally, users connect to a host using standard tools like telnet and a simple line terminal, and when they encounter a problem they can simply copy the plain text session output and email it to the consultants. With a web portal interface, however, users cannot simply email a copy of what they are seeing. Clearly, providing support for VMR web portal users requires a different approach.

Desktop sharing, also called data sharing or data conferencing, promises to provide the new approach necessary for supporting VMR web portal users. Data conferencing systems allow consultants to interactively work with a user and directly see what the user is seeing. Desktop sharing technology is fast maturing and on some platforms is already ubiquitous, being provided as part of the operating system. The ITU T.120 data conferencing standard is considered robust and is "incorporated" into the H.323 standard for audio, video, and data communications across IP-based networks, better known as internet desktop videoconferencing. Many different videoconferencing products are available from well-established vendors such as PictureTel, Intel, and Microsoft. Each of these solutions provide data conferencing based on the T.120 standard, allowing people to collaborate and share desktops even if they are using systems from different vendors.

Other web based data conferencing solutions exist separate and apart from videoconferencing solutions. Services such as WebEx[9] allow people to data conference "on demand" using a Java enabled web browser and without having to have previously installed or configured dedicated videoconferencing software. Currently WebEx runs on Microsoft Windows, Apple Macintosh, Linux, and Solaris platforms, and user and consultant need not be running on the same platform in order to data conference.

The benefits of data conferencing for providing user support are many. During a data conferencing session a consultant can directly see what the user is seeing, and can easily pickup details and clues that the user might have missed or disregarded as not important. Most data conferencing solutions allow for some level of dynamic interaction so that the parties connected not only see what the other person sees but can also, with appropriate permission, take control and manipulate the remote desktop. In this way it is much easier for a consultant to "become" the user and directly investigate the problem "inside" the user's environment. This approach should significantly reduce the time spent exchanging email, providing access to files, and checking environment variables.

The ability to multicast and have more than two parties data conference allows others like system administrators to join the discussion directly, further enhancing the level of support. Other applications of data conferencing for consultants include direct demonstration of visually enabled software and tools like graphical debuggers. Direct demonstration of such tools is much more efficient and powerful than simply typing instructions out and emailing them. In the future, data conferencing sessions might be recorded and then presented to a user for later reference.

Although data conferencing technologies promise exciting new ways for consultants to support users and in particular VMR users, security and privacy issues inherent when sharing a desktop need to continually be addressed and monitored as the technologies mature.

3.8. VMR Tools for Consultants

A necessary part of the VMR software infrastructure will be a set of tools consultants can use to investigate the details of any job running on any particular host. Included should be tools for querying the state of batch queues and batch hosts, gathering detailed process information, inquiring about pending jobs, and other common utilities usually found as part of a job batch system. A single common interface should allow consultants to investigate the details of a job regardless of what system a job is running on or what batch manager runs on a particular host. In addition tools should be available allowing consultants to query and investigate the details of the global queuing and job routing infrastructure built on top of Globus.

While these types of tools will be provided at some level for all users, it is also helpful for support staff to have “hooks” not necessarily available to general users. These specialized tools might allow consultants access to things such as job submission transcripts, detailed batch manager queries, individual system and network logs, and the like.

ACKNOWLEDGEMENTS

We are grateful to Steve Quinn for details about the NCSA ticketing system. This work has been partially supported by the NSF Partnerships for Advanced Computational Infrastructure (PACI) program.

REFERENCES

[1] http://www.ncsa.uiuc.edu/
[2] http://www.ncsa.uiuc.edu/SCD/Alliance/VMR/
[3] http://www.nas.nasa.gov/About/IPG/ipg.html
[4] http://www.sandia.gov/ASCI/
[5] I. Foster and C. Kesselman, *Globus: A Metacomputing Infrastructure Toolkit*, Intl J. Supercomputer Applications, 11(2):115-128, 1997. See http://www.globus.org.
[6] R. Butler and V. Welch, *Alliance Public Key Infrastructure: History and Future*, NCSA white paper.
[7] R. Butler, D. Engert, I. Foster, C. Kesselman, S. Tuecke, J. Volmer, V. Welch, *Design and Deployment of a National-Scale Authentication Infrastructure*, Globus preprint.
[8] Personal communication with Steve Quinn.
[9] http://www.webex.com