Mid-Range Computing in Support of Science

at Office of Science Laboratories

Report of a Workshop October 21-22, 2008

Table of Contents

Executive Summary

Computation continues to play an increasingly important role in enabling the Office of Science to support innovative basic scientific research that improves people's lives. Leadership Class computers that perform at the highest possible capability are critical to the advancement of science is many areas, but midrange computing also plays a vital and growing role in advancing science in disciplines where capacity is as important as capability. Demand currently seems to be limited only by the availability of computational resources. Berkeley Lab does not seem to be alone in projecting the growth of midrange computational cores for FY09 at over 33 per cent. Human resources to develop new algorithms and improve the scaling and performance of existing applications is also seen as a bottleneck.

Provisioning sufficient resources is challenging. A single approach will not be satisfactory because the requirements are many, diverse, and often contradictory. Integrated disciplinary research efforts can achieve value in centrally managing and allocating hardware and software focused on a narrow set of applications. The most common approach is a centrally managed lab-wide cluster, which can offer a more capable resource with consistent usage, but requires mutually acceptable allocation policies and sustained multi-year funding. A subset of this approach is to centrally manage project or group purchased resources. In light of growing costs of refresh, space power and cooling, along with the growing maturity of third party providers, several labs are investigating outsourcing some types of computing.

Competitive pressures are a strong reason to keep the cost of midrange computing as low as possible. Externally mandated energy efficiency requirements, however, may soon force a reevaluation of the various business models used to provide midrange computing. A consistent set of cost analysis principles would help clarify the value of tailoring hardware and software to particular algorithms and the communities that use them versus more general systems that can support a broader set of applications. A range of contractual and technical issues need to be addressed before any large scale movement to third party "clouds" can take place. A small interlab effort to clarify relevant cost analysis principles seems worthwhile.

The set of tools used to allocate and manage resources is another area that would benefit from information sharing. The tools, like the resources they manage, are constantly evolving and sharing dynamic information requires focus and purpose.

Finally, there is the issue of data management. As the datasets get bigger, moving them becomes more of an issue. Even labs that have mature processes for managing computing resources have much less robust plans for managing data or computational workflow. This is another area where a set of best practices could contribute to the ensured viability of midrange computing.

Introduction

This report contains a summary of the status of midrange computing efforts at the ten Office of Science Laboratories. It represents a follow up to the discussions that occurred during the presentation of the laboratory business plans in the spring of 2008. The increasing importance of scientific computing at all scales to the success of science has been widely discussed and many efforts in the Office of Advanced Scientific Computing Research (ASCR) such as the Leadership Computing Facilities, NERSC, and SciDAC fill a portion of the computing ecosystem needed for DOE leadership in science. One description of this ecosystem is shown in Figure 1, taken from a BESAC report[1].

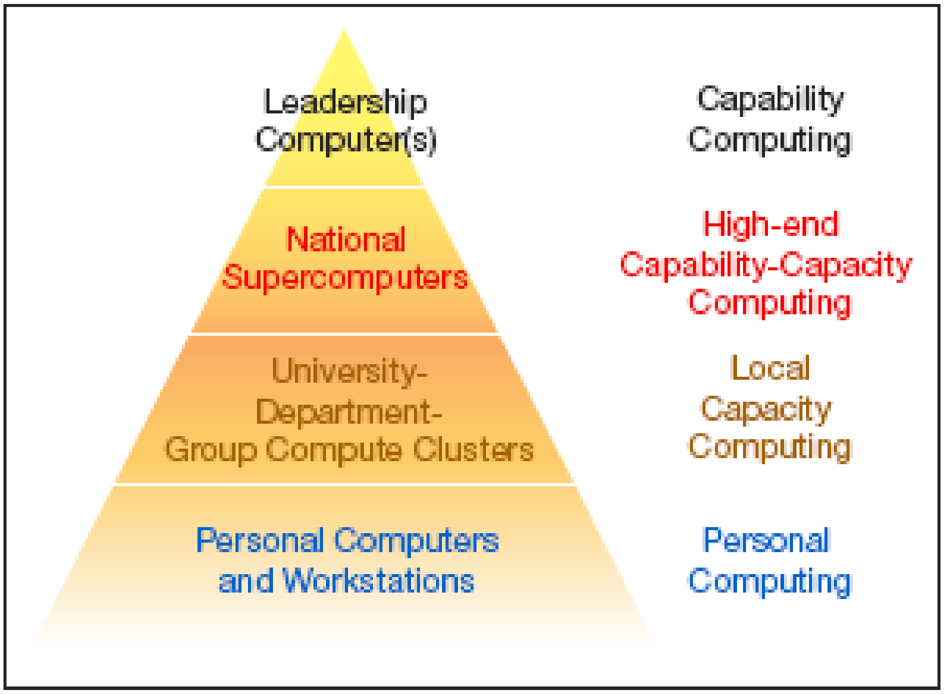

Figure 1 Updated version of the "Branscomb Pyramid"[2] of computing resources.

The top two layers of the computing ecosystem, within the Office of Science, are primarily the responsibility of ASCR. All of the laboratories have well established methodologies for dealing with the bottom layer of the pyramid. The focus of this report is on the strategies and plans of the laboratories for the second tier of computing. It is important to note that this pyramid is not just hardware. At each layer there are requirements for support, system management, cybersecurity, and physical infrastructure. In addition, at each layer of the ecosystem there are corresponding requirements for storage and management of data, both generated by the computing as well as experimental data that is related to the computing.

There are a number of reasons why it is important to evaluate the planning in this area at this time.

- The demand in the scientific community for this class of computing is rapidly increasing as the scale of the largest systems increases;
- The complexity of managing and securing this class of system is increasing;
- The power and infrastructure requirements are increasing at a time when energy efficiency is becoming more critical;
- The number of alternative ways for providing these capabilities to scientists are increasing due to both research community initiatives such as the "Grid" and commercial alternatives such as "Clouds" operated by companies like Amazon and Google; and
- Future evolution of microprocessor design will present significant challenges for this class of computing.

In order to develop a better understanding of the impact of midrange computing on the SC Laboratories, The Deputy Director for Programs of the Office of Science (SC-2) charged the ASCR to convene a workshop to address the following issues:

1. assess the current computing and networking capacities of the laboratories;
2. summarize their current and projected needs in these areas based on existing or approved projects and activities; and
3. summarize their needs in these areas based on proposed but not yet approved activities.

To address these questions ASCR held a workshop on October 21-22, 2008. Each of the SC Laboratory Directors was asked to send between 1-3 people to the workshop who could discuss the requirements for midrange computing at their laboratories and their plans for providing these services. All ten laboratories were represented at the workshop. The next section will discuss the structure of the workshop.

It was clear from the workshop that the laboratories had significant and growing requirements for midrange computing. It was also clear that managing data was as important for a number of the laboratories as computing. The laboratories had a variety of plans for providing these capabilities ranging from shared central resources to individual hardware owned by individual projects. All the laboratories are moving to more centralized management of these resources driven partially by cost but primarily by cybersecurity. The funding models also range from direct funding by projects to overhead funded resources. It was also clear that the individual Offices of the Office of Science think about the integration of computing into their projects in different ways with HEP supporting the highest level of integration.

Workshop structure

The workshop was held October 21-22, 2008 at the Hilton in Gaithersburg, Maryland. In addition to representatives from the laboratories representatives from each of the SC Programs were invited and five of the Programs were represented.

The first day of the meeting was taken up with reports by the laboratories on the requirements for midrange computing. These presentations used a standard template (Appendix A). A detailed analysis of the requirements will be presented in the next section. The second day was devoted to a detailed discussion of the laboratories plans for providing midrange computing for science. A number of possible joint activities resulted from the discussions on the second day.

Role of Midrange Scientific Computing within the DOE Office of Science

What is Midrange

Roughly speaking, midrange computing is represented in Figure 1 as Local Capacity Computing, but it can be either capability or capacity. While not the top of the pyramid, midrange computing may supply some unique capability aspect such as very large memory per processor, unique I/O or other capabilities not found on the general purpose supercomputers at the top of the pyramid. The following guideposts were presented at the Workshop:

- Geometric Mean of NERSC Capability and Desktop ~ 20 teraflops
- Geometric Mean of cores on leadership computer and desktop ~ 700
- Geometric Mean of NERSC Capacity and desktop ~ 13,000 CPU hours
- Geometric Mean of High End System Cost and desktop Cost ~ $500,000

In other words, a midrange system can offer significant compute power and may command a significant investment. Midrange computing systems are generally not found in individual offices, but they may not be in managed datacenters either, instead being located in “closets”.

Midrange computing serves a variety of needs that can be divided into roughly two categories:

1. serve as an onramp to effective use of Leadership Class machines; and
2. support for science not served by Leadership Class machines.

Onramp to Leadership Class Computing

Effective use of the Leadership Class (LC) systems would not be possible without extensive access to midrange systems for training, application development and staging results for pre- and post-production analysis. Midrange computing also acts as a supplement to LC computing. Berkeley Lab estimates that 25% of its midrange computing resources supplement projects with inadequate LC allocations.

Training

The LC systems are extremely valuable resources used by world class scientists to solve critical problems of national and global importance. You don’t give the keys to a Formula I racer to someone who has never driven and you don’t allocate LC resources to inexperienced users. So where do inexperienced uses become experienced? On mid-range systems. Demonstrating proficiency on midrange systems is a prerequisite for LC use.

Staging

Although an LC system may be essential for a particular simulation, frequently pre- and post- production runs on smaller midrange systems is necessary for preparation and analysis.

Development

Although there will always be a need to test the performance and correctness of new or improved LC applications, the majority of the effort can be done on midrange systems. Debugging on a midrange system is a more appropriate use of resources. LC applications need to demonstrate correctness and scalability on smaller midrange systems before being run on the full LC machine.

Science not targeted at Leadership Class Computing

Although LC systems are necessary for solving many of our most pressing problems, there are other problems where LC systems are inappropriate. Not all science needs LC level resources to make progress. Given the level of compute power now available in a mid range system, we should not be surprised to see world class science being done with them. There are plenty of difficult, if smaller computationally, problems waiting to be addressed and getting on the cover of Science (Figure 2), as did the work of Gutowski[3] *et al.*

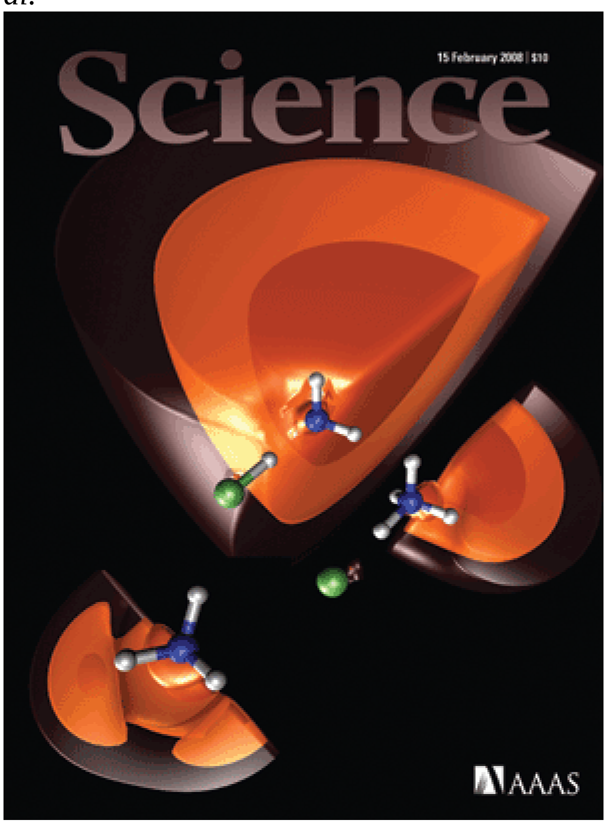

Figure 2 Science cover from 15 February 2008

Figure 2 is a depiction of the interaction of an excess electron with the hydrogen-bonded complex NH_3···HCl, which induces formation of the ionic pair $NH_4^+Cl^-$ solvated by the excess electron in contrast to widely familiar acid-base behavior in solution. The image shows the structures of three possible systems and highlights the areas that correspond to 10%, 30%, and 50% of the excess electron.

A BES chemistry user was quoted as saying “We need to probe a relatively large range of conditions (temperature, pressure, chemical environments, mechanical perturbations, etc) to understand a physical or chemical phenomena or even an experimental observation and that typically requires a large number of calculations of similar size (but not consuming huge numbers of processors). Also, we often have novel ideas on a particular subject and need to run a quick calculation to confirm a pseudo-hypothesis before proceeding forward.”

Serial/scalability issues

The Science Labs presented several instances when applications were not candidates for LC allocations because they could not make effective use of such a powerful resource. A common problem for these applications was the serial nature of the algorithm or its inability to scale to large numbers of processors. How can this be? One reason has been the performance safety net Moore’s law used to provide. Applications could expect to see steady improvement in performance with each new generation of processor that was faster than the previous one. While Moore’s law still holds for the number of transistors on a chip, the transistor speed has not improved significantly for several years. The result is the steady increase in the number of cores per chip to keep pace with increased performance expectations. This is putting pressure on application developers to create parallel programs instead of maintaining serial codes, but parallel programs are more difficult to write and many scientist/developers lack the skills necessary.

Although the number of serial applications can be expected to decrease over time as new generations of researchers are more familiar with parallel programming methods, the issue of scalability is not so easily solved. Naïve, intuitive algorithms, the kind most often conceived by scientist/developers, often don’t scale to large numbers of processors without the sophisticated use of advanced data location and thread management tools. In many cases completely new algorithmic approaches are necessary to avoid the bottlenecks imposed by high latency and low bandwidth. These bandwidth/latency issues occur at multiple points in any high performance system including cache, memory, interprocessor communication, and I/O.

While it would be wonderful if scalable, parallel applications were easy to develop; this is not the case now not likely to be so for the foreseeable future. With effort, most applications can be made more scalable. There are, however, cases where the effort is not worth the investment. An example is rapid prototyping, where new ideas need to be tested quickly first on the correctness of the approach and optimized for performance and scalability only after validation. Midrange computing will continue to serve an important function in supporting applications that haven’t yet or will never achieve the kind of performance that an LC machine can supply.

Trivially Parallel Applications

There is another class of problems that are capable of excellent performance per core and can scale well but are not suitable candidates for an LC allocation because they don't make use of the expensive high performance interconnect that typifies an LC resource. Frequently derided as "trivially parallel", their principal characteristic is the small amount of interprocessor communication they do, i.e., they have very large compute to communicate ratios and are capable of running on large clusters with cheaper commodity networks. Having an application with these characteristics can be very desirable as they can potentially command resources much larger than any LC machine, i.e., clouds as large as the Internet itself, but the application can also be used to absorb unused midrange cycles. These applications fall into a couple of broad categories: parameter sweeps and statistical sampling, sometimes collectively referred to as Monte Carlo methods. Although there is a quantum of work that requires little communication, within that work unit there may indeed be a requirement for a high performance interconnect as may be found on a shared memory node of a cluster.

Parameter sweeps

Parameter sweeps consist of running the same application, which could run on an LC machine but more typically use midrange systems, over and over with different input parameters. This can be a brute force optimization method (perhaps used in the absence of an objective function) or the characterization of a parameter space, such as mapping the potential energy surface of a chemical reaction. Searches can frequently be performed with this technique.

One researcher was quoted as saying "one really needs to run tens of thousands or possibly hundreds of thousands of jobs that sample different regions of conformation space. It is at this point that the major supercomputer facilities around the country fail abysmally. Most of them are structured to avoid running many jobs independently on many CPU's. (For instance, at NERSC in Berkeley one is only allowed to have something like four jobs in the queue at one time.)"

Monte Carlo methods

Another approach repeats an application (or piece of an application) over and over to build up a statistical representation of a parameter space. In many instances this is the only way to measure the accuracy of a simulation. Global climate models use ensembles of runs to quantify the inherent uncertainty of nonlinear models.

Monte Carlo methods are often the method of last resort, however, because the amount of computation is generally greater than other approaches. This may not seem to matter if computation is cheap and a useful answer can be arrived at faster than a method that uses less computation but maybe doesn't scale as well to large numbers of processors so it takes longer. This advantage may disappear, however, if energy costs are considered rather than, or in addition to, time to solution.

Unique requirements

As previously mentioned, LC resources are focused on high performance computation using a fast, low latency interconnect. Trivially parallel applications are not LC candidates because they don't make effective use of an expensive component: the interconnect, but there are other classes of problems that simply won't run efficiently or at all on LC machines because of some unique hardware requirement such as large amounts or fast access to secondary, tertiary or archival storage; very large memory; or real time requirements.

Data-Intensive

Another class of problem is not suitable because it makes minimal use of the computation units because most of the time is spent moving data rather than processing it, so called data-intensive applications.[4] The Office of Science hosts many data-intensive applications, usually associated with a unique data generator such as the ATLAS[5] experiment at the Large Hadron Collider (LHC), the Linac Coherent Light Source (LCLS)[6], and the Large Synoptic Survey Telescope (LSST)[7] to name just a few of the large physics data generators coming on line soon or under development. Other disciplines, from biology and climate to materials, are rapidly coming to grips with the need to analyze large amounts of data.

There are many applications that would like to access large amounts of memory in random patterns. Available solutions include shared memory machines, which do not scale to very large systems because of cost issues, and hardware assisted gather-scatter functions, which requires expensive special purpose hardware.

Real time

Making the most efficient use of an LC resource normally requires a batch queue. Interactive use of the entire LC resource is only considered under special conditions, if at all, so applications that must deliver results in real or near-real time are also not LC resource candidates. This is not to say that an LC resource is not involved, it may be the generator of the data, but the real-time analysis is typically done using a more interactive environment using midrange computing.

Visualization

Visualizing the results is a task best suited to an interactive environment. With midrange computing it makes more sense to dedicate the entire system to producing the visual information.

Tight coupling with experiment

Real-time control of an experiment typically requires a dedicated midrange system. If very tightly coupled, however, it may be considered part of the instrument and not a separate system, as is the case for detectors in many physics experiments, but this is not always the case. The LSST is an example. The data analysis system is not part of the instrument, but it should be able to respond to cosmic events in a timely manner so that other telescopes can be focused on rare events such as supernovae.

Status and Barriers at the Labs

The following summarizes how individual labs manage their midrange computing.

Ames Laboratory

The strategy is to

- Limit intergroup allocation hassles (groups control users and jobs)
- Avoid the dispersion of assets
- Avoid lockin
- Build clusters designed for the algorithms they run
- Provide a single manager with consistent software management for all clusters
- Seek out and test novel solutions before deploying them

A single staff person supports 50 clusters, but datacenter space has filled up more rapidly than expected. Taking into consideration power and space requirements, Ames is investigation scalable storage container units that could be placed near wind generators. There are areas in northern Iowa and southern Minnesota that have nearly steady 15-20 mph winds. The current system from Rackable can deliver 40 TeraFLOPS for under $5 million and a next generation product could deliver 140-160 TeraFLOPS (at 400-460 MF/Watt)

Argonne National Laboratory

Argonne's Laboratory Computing Resource Center (LCRC) was established in 2003. Currently, its primary resource is Jazz, a terascale Linux cluster that will be replaced in 2009. A staff of 3.5 FTEs support Argonne applications. It is available at no charge to ANL employees and projects and meets many midrange needs. Its charter is to help as many Argonne groups as possible to use the HPC facility. Allocations are granted by a committee and is under the overall guidance of the Computational Science Advisory Committee. Startup accounts of 1000 hours require no justification and are granted immediately through an online project creation and management facility. Streamlined project proposals are reviewed quarterly. Tutorials and Hands-on Training include:

- Introduction to Jazz and MPI
- Parallel Programming
- Performance Tuning
- Advanced MPI
- Advanced Parallel Programming
- Introduction to PETSc

Expert consulting and user support is available for many software development tools, including 6 brands of compilers. Support is also provided for application installation services, licenses for common applications and problem and performance analysis.

Brookhaven National Laboratory

Brookhaven's key research activities rely on the ability to capture, store, and process large (hundreds of terabytes to multi-petabyte) datasets created through experiments and high performance simulation. Data can be generated at rates over 100 MS/second. Brookhaven is one of just ten LHC Tier 1 sites worldwide (Brookhaven is the US ATLAS Tier 1 site and Fermi is the CMS Tier 1 site). One of Brookhaven's strengths is the RHIC-ATLAS Computing Facility, a network-centric Linux farm (thousands of heterogeneous compute nodes) with 8 PB of automated storage (growing to 30 PB in 2010) supporting thousands of simultaneous jobs. In contrast, midrange needs are modest, mostly cluster augmentation and replacement, involving a number of small dedicated clusters with about 100 to 500 cores (CFN, BNL Cluster, Chemistry, NSLS, etc.). Brookhaven is interested in studying the tradeoffs between commercial and lab owned and operated systems.

Fermi National Accelerator Laboratory

Fermilab considers *all* of its scientific computing facilities as midrange and all are programmatically funded. All scientific computing is designed, procured and managed by the Computing Division and is housed in one of three well-interconnected computing facilities:

1. Feynman Computing Center (FCC)
2. Grid Computing Center (GCC)
3. Lattice Computing Center (LCC)

The GCC and LCC have made adaptive re-use of experimental buildings and were built incrementally for high density and efficiency according to a multi-year plan beginning in 2002. Fermilab serves as the interface to the Open Science Grid[8], jointly funded by NSF and DOE.

Thomas Jefferson National Accelerator Facility

Computing efforts at Jefferson Lab are focused on support of its mission: experimental physics data acquisition, storage, and analysis (farm computing); and Lattice Quantum Chromodynamics (LQCD) theory calculations of fundamental quantities. The USQCD Collaboration consists of nearly all high energy and nuclear physicists in the US involved in LQCD. The Jefferson Lab portion of LQCD I cluster is allocated to collaboration members on a peer-reviewed basis. This midrange machine performs smaller analysis jobs than the leadership class machines and is made up of a 256 node machine (6n) and a 396 node machine (7n). Since these clusters can be configured for a single application (LQCD), they can be better optimized than general purpose clusters, i.e., they can be memory and disk lean, use a pruned fat tree network due to the highly local communications pattern, and provide lower aggregate bandwidth to disk with the overall impact that they can provide 50% more computing capacity per hardware dollar spent. They area two to five times more cost effective for analysis jobs than leadership class machines.

Lawrence Berkeley National Laboratory

Berkeley Lab supports about 4,000 scientists and staff as well as about 4,000 guests working in a diverse scientific portfolio. In a recent poll, 38% of the scientists said they depend on cluster computing for their research and 69% said they are interested in using a Lab owned cluster, with early-career scientists twice as likely to be ‘very interested’ than their later-career peers. The Information Technology Division manages 35 clusters containing 5126 processor cores. These clusters serve over 500 scientific users and are managed by 5.15 FTEs plus one recently added user support FTE. The projected growth in cores for FY09 is over 33%.

The Scientific Cluster Support (SCS) Program began in 2003 with free support for a group of pilot clusters. This has evolved to where the cluster owners now pay for administration, but the service was 'too successful,' resulting in scalability issues; no shared storage; and data and code that was not easily portable between clusters. The current MetaCluster system has the clusters more closely interconnected so storage and other resources can be shared, a master job scheduler that submits to all clusters, and a 'super master' node that manages the others. The software toolkit is Perceus[9].

Most of the PI-owned clusters have been optimized for tightly-coupled parallel computation, which will be provided by a new institutional cluster (16 TF, 1500 cores), but researchers also have serial computing needs. LBNL is investigating 'cloud' services to meet these needs in a cost-effective manner and is having productive conversations with ANL.

LBNL has a cost-effective service model for midrange computing, but the business model is a challenge. Barriers include stable funding and datacenter capacity. To improve datacenter efficiency and capacity, the IT Division has been collaborating with the Environmental Energy Technologies Division, which has led to the establishment of an IT datacenter testbed and an environmental monitoring system using a wireless sensor network.

Oak Ridge National Laboratory

ORNL is DOE's largest science and energy laboratory with 4,250 employees and 3,900 research guests annually featuring the world's most powerful open scientific computing facility, the nation's largest concentration of open source materials research, the $1.4 billion Spallation Neutron Source and managing the billion dollar U.S. ITER project.

Program	Units
CNMS	3
SNS	1
NCCS	2
CCSD	4
BES-materials	1
BES-chemistry	1
BER-climate	1
NE	2
NSSD	1

Table 1 OIC Support

The ORNL Institutional Cluster (OIC) business model, which supports midrange computing, is based on the model adopted by NNSA/ASC and implemented at LLNL. It features a large vendor contract negotiated centrally and managed centrally. Programmatic costs pay a one time fee of $250K for one hardware unit and a monthly service fee of $1500 per unit. Lab overhead funding covers approximately 1 FTE for support, Red Hat Enterprise Linux licenses and network switch maintenance while the service fee covers additional system administration time and other costs approved by the OIC Steering Committee. OIC currently has two clusters delivering 16 units with over 3,300 compute processors. The programmatic distribution of the units is shown in Table 1The advantages of the OIC include a way of managing IT risks and costs while providing high-quality computing resources. A users group creates a peer-to-peer community that facilitates technical information exchange and representative governance for the OIC steering committee.

Pacific Northwest National Laboratory

PNNL midrange computing consists of 40 dedicated clusters comprising about 5,000 processor cores in addition to an older general access shared memory system with 128 cores.

Table 2shows the approaches to midrange computing tried or considered at PNNL. PNNL's strategy is to continue support for dedicated clusters; pursue co-investments involving the Lab, Program and Project (providing infrastructure, including storage resources, to allow general access to dedicated clusters on a non-interference basis); pursue collaborations to allow midrange access to leadership platforms; and invest in data management technology and infrastructure.

Approach	Pros	Cons	Experience
Individual Project Buys & Maintains	Dedicated cycles for project.	Unused cycles; not all projects need dedicated clusters.	Successful at PNNL; informal arrangements allow some other projects access.
Laboratory Overhead Buys; Projects charged for access	Cycles available to all projects; costs associated with benefitting projects.	Expensive for projects without critical subscription rate.	Unsuccessful at PNNL. Projects that can afford want dedicated cycles. Other projects not willing to pay.
Laboratory Overhead Buys & Maintains	Cycles available to all projects.	Overhead charging not equitable for projects not benefitting.	Successful at PNNL, but Systems tend to be over-subscribed; limited effectiveness without ongoing commitment to support
Co-investment between Lab, Sponsors, and Projects	Dedicated level of cycle availability for investors. General cycles available. Benefits of scale.	Differences in timing of funding between projects challenging.	Have not tried at PNNL.
Cloud (external provider)	No up-front investment.	Cyber Security issues. Performance relative to clusters.	Have not tried at PNNL.

Table 2 Midrange Computing Approaches at PNNL

Princeton Plasma Physics Laboratory

PPPL provides a centralized facility for capacity computing, standardized on no more than two architectures, with centralized storage and backup services. The goal is to mimic leadership class systems to enable applications to move from midrange to LC. It also provides a resource for applications that do not scale or do not have allocations elsewhere. Primarily overhead supported ($76K/year for hardware), with occasional OFES support, about 150 researchers use the Linux clusters with about 1300 cores. Of these, about 150 support serial applications, 32 support applications with large memory requirements (8 GB/cpu) and 160 are in dedicated clusters. The rest are in Infinband (576) or Ethernet (384) connected clusters. The systems average about 80% utilization but the wait times are steadily increasing. Over 200,000 jobs will be run in 2008 with 30% being single CPU jobs, 50% use 4-16 CPUs and 20% use 16-32 CPUs and the average run time being 200 hours of wall time.

SLAC National Accelerator Laboratory

All SLAC Scientific Computing hardware remains program funded while Lab overhead pays for support for (optional, through intelligent matchmaking) pooling/sharing resources as well as support for agreed standard software, file systems, etc. Over 7000 processor cores are used in data analysis “farms.” Another 492 cores make up Infiniband and Myrinet clusters, there is a 73-core shared memory multiprocessor, application-specific clusters, a research prototype cluster and an Apple cluster. Disk servers provide over one petabyte of network attached storage (supporting xrootd, NFS, AFS, and Lustre) backed up by six STK Powderhorn Silos with a capacity of up to six petabytes (being upgraded to an initial capacity of 13 TB. The SLAC Scientific Computing Sub Council gives the programs the opportunity to set common priorities and (likely multiple) standards and advises computing management on scientific priorities. The Lab assures career continuity for key computing expertise likely to be needed by and funded by the programs. The Business model for provisioning of power, cooling and space is yet to be determined. SLAC is evaluating the performance of two Sun Black Box systems while planning a new, up to 24 MW, facility.

Program-focused Grids can be a good solution to some needs, e.g. OSG and BaBar poor-man’s grid. About half the computing cycles for the SLAC HEP program are provided by collaborators in other countries. (SLAC can claim it is not a computer-center hugger). Collaborators is a key word: distribution of data-intensive tasks requires commitment to the long-term provision of storage; and matching the computing architecture to the evolving science task requires involvement in the science. There are also many tasks (e.g. compute intensive, trivially parallel) that are totally cloud-ready if the price is right! People and their interactions are key elements: easy access, even serendipitous access, to people who can help apply computing hardware and software to science is even more important than access to cycles – this is a major benefit that lab computing can bring to the programs.

Thoughts on the future

It became clear during the workshop that the laboratories have developed a number of approaches to providing midrange computing from multi location partnerships such as

LQCD through centrally managed and/or allocated resources at one laboratory to individual group clusters. It was also clear that most laboratories were moving to some form of central management, if only to secure the systems adequately. It also appears to be the case that, with the exception of large HEP experimental collaborations, the cost of computing is not included in the cost of the work proposed to the SC Programs. Finally, even laboratories that have relatively mature processes for managing the midrange resources have much less robust plans for managing the data or the computational workflow.

Many of the laboratories cited as one reason for their way of providing midrange computing the need to keep the cost of proposals low for competitive reasons. However, all the laboratories that keep statistics are experiencing significant contention for these resources. This is reflected either in ratios of request to resources in the range of 3-4 or increasing wait times for jobs to run. In addition, most laboratories are experiencing pressure on their physical infrastructure (space, power and cooling) to support this class of computing.

From the workshop it is clear that no one model will fit all the requirements and that both cost and non cost considerations must be weighed. The underlying principal of the Office of Science to support the best science within the constraints of the budget must be the first guide. It is critical to note that the purchase cost of the hardware is not the major part of the cost in most cases. The next sections will discuss the pros and cons of varying models for provisioning and their underlying business models.

Provisioning

The first provisioning model is the resource at one or multiple sites supporting an integrated disciplinary research effort. Examples of this sort of situation are the collaboration around the LHC and LQCD. They are characterized by development of community software and long range sharing of data. In these cases having a centrally managed and allocated resource has significant value. The technical issues of doing this are relatively straightforward; however, the community must be sociologically ready. The Climate research community is moving in this direction. The largest advantage is that the hardware and the support infrastructure can be specialized for the science lowering the overall cost and helping to build the community.

The second provisioning model is the centrally allocated lab-wide cluster. This has the advantage that the varying use by scientists can be smoothed out with all scientists having access to a more capable resource. The disadvantage is that allocation policies at the laboratory have to be worked out. In addition, since charge back has been shown to lead to very bad outcomes in this space this sort of resource needs continual support as a strategic resource for the laboratory. The multi year funding that is required to refresh these systems is an issue for many laboratories. The underlying business model here is that computing is a strategic resource for the laboratory and providing it enables the laboratory to lead in science.

The third provisioning model is project or group purchased but centrally operated resources. This maintains the independence of the individual groups, but achieves some economies of scale by central management. Sharing of unused cycles is more difficult in this model. However, development of Grid tools could make it easier to provide general access on a non-interference basis to these resources.

Finally, in light of growing costs of refresh, space and power and the growing maturity of third party providers, outsourcing the computing is being investigated by several labs.

Argonne is investigating Infrastructure as a Service (IaaS), sometimes known colloquially as “cloud computing”. Security is an obvious concern, but a preliminary analysis of Amazon Web Services (AWS) was promising. Amazon recently published its security procedures[10]. The SOX and SAS 70 certification appear comparable to NIST 800.53. Argonne believe is will likely be able to rely on this for secure infrastructure with the following observations:

- Risks are augment by additional personnel to be trusted
- Risks are reduced by additional separation of duties
- It is possible to configure machines with restricted network tunnels to site machines only
- Most threats and mitigations are identical to current security analysis and plans
- The datacenter has State of the Art physical protections
- There are internal and external network protections
    - AWS uses a default deny architecture at multiple tiers
    - The firewall rules are under user/VM control
    - Global IDS is run by AWS staff
- VM protections include
    - Images are encrypted using a user key that is unavailable to AWS
    - There is custom VM isolation within Xen
- Insider thread protection via two person rules and business process oversight
- There is not much in the way of penalties or guarantees
- Most threats and mitigations are identical to current security analysis and plans. Of 18 identified risk factors
    - 12 were the same risk
    - 3 had mixed risk
    - 2 had increased risk
    - 1 had reduced risk
- The user effectively assumes all the risk
- Low risk of AWS staff exploiting access to data/services (Amazon corporate procedures seem reasonable and accredited by SOX and (soon) SAS-70 audit procedures
- Low risk of other AWS customers attacking from “inside” AWS (internal design isolates instances and service access is under Argonne control)
- Current techniques and tools apply to attacks from the general Internet (access under Argonne control)
- Insider attacks from ANL staff is the same as for all other ANL operated services.

- Other mitigations
    - AWS internal OpSec measures as good or better than ANL
    - Data stored in AWS can be encrypted (file system or file level)
    - Communications with AWS services may require SSL
    - ANL AWS credentials must be protected by ANL (rekeying possible under ANL control)
    - Image maintenance and security management is ANL responsibility (leverage current patching knowledge and infrastructure)

A comparison of the features of Clouds vs. Grids is presented in Table 3 and a comparison of Cloud Challenges and Advantages is presented in Table 4.

Clouds	**Grids**
X as a service	Shared Resources via API
Driven by Industry/Business	Driven by Science/Academia
Virtualization based	Hardware based
Pay as you go (dynamic)	Quid pro Quo
Pros • Common currency ($) • Economies of scale • Meme of the week • Direct business relationship	Pros • Science focus and collaborative development • Functioning today
Cons • Expensive (at least now) • Data hard to deal with • Single vendor lockin	Cons • Labor intensive and custom • Mesh trust/business arrangements • Economy not worked out

Table 3 Clouds vs. Grids

Challenges	**Advantages**
Security Analysis and (perception) management	Reduced local machine room, power and admin demands
Data transfer and access	Rapid response growth and deployment of services
Virtual Machine image creation/management	Disaster recovery
Usage capping and subaccount billing	Execution environment archive and exchange
User/Group identification, authentication, authorization and management	Economies of scale and competition drive $ down

Table 4 IaaS Cloud Challenges and Advantages

Argonne concludes that the bottom line for IaaS looks very promising. They are starting an internal pilot study of business use. An initial, small contract with Amazon is almost in place and they are looking at security, return on investment and application matching.

Cost analysis principles

In all of these cases the answer depends strongly on the cost analysis principles that are used. For example the ability to tailor hardware to a particular algorithm has some value as does the ability to make use of the computer by the wider community. It is also important to decide exactly what costs and benefits should be included in such an analysis. A number of the laboratories are currently evaluating their strategies and there was agreement at the workshop that a small interlab effort to clarify the principles would be useful. Especially as laboratories consider outsourcing some of this type of service to "cloud" providers like Amazon and Google it is clearly necessary to think carefully about the contractual and technical issues involved.

Resource Management

Once the resource has been procured and installed it must be operated and here again there are opportunities for sharing resources, software and knowledge.

Tools

The laboratories use a variety of tools to allocate and manage these resources. Some laboratories use more than one tool. It is clear that discussion between the labs would lead to better sharing and improved cost effectiveness. In addition, the sharing of these tools could lead to improved and broader understanding by the management of the laboratories of the impact of this class of computing on their science.

Staff

The laboratories provide different levels of support for these resources with some laboratories providing only the support to keep the systems running and some other labs providing more significant user support and even consulting. These services will be increasingly important as mid range computing evolves in the coming years.

Conclusion

Midrange computing plays a vital and important role in enabling the Office of Science to accomplish its missions. Its role is multifaceted, serving as a training, development and staging agent for the largest Leadership Class simulations and for advancing science in areas where the LC resources are inappropriate, as where large numbers of midrange scale runs are required or hardware with features not available at the LC level are required, e.g., data intensive or near real time computing. It is also the resource of last resort for applications that have not been moved to scalable parallel computing. There were numerous instances of users that desired very long run times, when perhaps an investment should be made to make the application more scalable and parallel to shorten the runs. Investments in research that would result in the improved productivity of application developers, allowing them to quickly turn algorithms into scalable parallel programs, could have a significant impact.

All the labs have moved to some form of centralized support. The growing size and complexity of clusters and associated data management has put an end to individual groups being able to support and operate these resources, although they are not always centrally located. Growing space and power requirements are becoming critical facility issues. The business models used range from complete reliance on programmatic funds for hardware, software and support to the nearly full use of indirect funds. The later case results in heavy pressure to minimize scale of the computational resources to make overhead rates more competitive. The usual result is that demand outstrips the available resources. The labs are constantly experimenting with a variety of approaches, in part driven by the evolution of the technology, including outsourcing the computing to clouds, in efforts to optimize the cost/benefit ratio. It is clear that there is no single solution that would satisfy everyone. At the same time, coordination could reduce redundant efforts at piloting new approaches.

Large, multi-investigator projects seem to better manage computing as one aspect of managing a project. Smaller, individual investigator projects seem to be the most likely not to properly externalize the necessary computing costs, which are often listed only in terms of hardware, if they are mentioned at all. Hardware, of course, is the tip of the iceberg. As with other facility infrastructure, however, there will always be portions of the IT infrastructure that will be difficult to charge directly for. The vitality of the labs comes in part from their ability to try new ideas via the LDRD process, so one should not expect overhead funded resources to completely disappear either.

Appendix A: Approved Mission Tables

Appendix B: Proposed Mission Tables

References

[1] Subcommittee on Theory and Computation of the Basic Energy Sciences Advisory Committee U.S. Department of Energy (Co-chairs: Bruce Harmon, Kate Kirby, and C. William McCurdy), "Opportunities for Discovery: Theory and Computation in Basic Energy Sciences", January 2005. Available online at http://www.sc.doe.gov/bes/reports/files/OD_rpt.pdf.

[2] NSF Blue Ribbon Panel on High Performance Computing (Lewis Branscomb, Chairman), "From Desktop to Teraflop: Exploiting the US Lead in High Performance Computing" (Washington DC: National Science Board NSB 93-205, August 1993).

[3] Soren N. Eustis, Dunja Radisic, Kit H. Bowen, Rafa A. Bachorz, Maciej Haranczyk, Gregory K. Schenter, Maciej Gutowski, " Electron-Driven Acid-Base Chemistry: Proton Transfer from Hydrogen Chloride to Ammonia", *Science* **319**: 936-939, 2008. [DOI: 10.1126/science.1151614]

[4] Data-Intensive Computing, IEEE Computer special issue, April, 2008.

[5] http://atlas.ch/

[6] http://lcls.slac.stanford.edu/

[7] http://www.lsst.org/

[8] http://www.opensciencegrid.org/

[9] http://www.perceus.org/

[10] http://s3.amazonaws.com/aws_blog/AWS_Security_Whitepaper_2008_09.pdf