Fermilab

Computing Division

Strategic Plans

FY2010

Table of Contents

Computing Division

Strategic Plan for the Computing Division (2010-2013)

Victoria White
10/30/09

Mission

The Computing Division's mission is to play a full part in the mission of the laboratory and in particular to proudly develop, innovate, and support excellent and forefront computing solutions and services, recognizing the essential role of cooperation and respect in all interactions between ourselves and with the people and organizations that we work with and serve.

Context and Assessment of Current State

The strategic plan for the Computing Division is guided by the strategic directions for the laboratory, which in turn are guided by the strategic plans for the field of high energy physics and related disciplines as presented through the HEPAP P5 report of May 2008. Existing memoranda of understanding and project commitments are also folded into CD strategic planning.

We are involved in a broad range of scientific programs of the lab as scientific collaborators on CDF, D0, MINOS, MiniBoone, Minerva, CMS, SDSS, DES, Pierre Auger, JDEM, US Lattice QCD, COUPP and CDMS. Our scientists may, in the coming years, participate in other scientific efforts at the lab – such as the Mu2E experiment, various neutrino experiments, and other experiments formed either to use existing facilities at Fermilab or to exploit the physics made possible by Project X, Under the guidance of the Center for Particle Astrophysics additional astrophysics experiments may be included in the portfolio of lab projects.

In prior years we had been involved in software infrastructure for the SiD ILC detector, in ILC accelerator modeling efforts and in ILC controls engineering and software work. With the decreased emphasis on ILC work we expect to have less involvement in this area, but rather to focus our limited resources on support for nearer term future detectors and more generic software tools.

It is within the mission of the Computing Division to provide DAQ and Trigger (Hardware and software) solutions and expert support to experiments and some testbeam users. We maintain expertise and services for a pool of detector electronics equipment known as PREP. We maintain legacy software DAQ applications and provide expertise and support for control system components (hardware and software) developed by CD. We assist experiments with their software infrastructure and with management and performance of their scientific codes. We provide tools and assistance to scientific programs to help support collaboration, to enable data analysis, to advise and assist with software and database coding and software infrastructure, to design new detectors and understand old detectors. Much of this is done through support of tools used widely in our field. Limited resources mean that we cannot provide all of the experiment and testbeam support we would like to, and that we are asked for, and so difficult choices have to be made about how best to deploy our limited resources in areas where it is either most needed, or can have the highest impact on the experimental program.

We are involved in collaborative efforts to provide the innovative and forefront computing solutions needed for the scientific programs through Advanced Computing research programs of work.

We provide services and computing solutions to all the scientific programs of the lab through common solutions and shared scientific computing facilities and services.

We are taking a strong leadership position in the Open Science Grid and we are working to ensure that the wide area networking infrastructure will support a globally distributed computing infrastructure.

We work on the remote operations center for the LHC and help facilitate access to and monitoring of the accelerator and the CMS detector.

We build and operate innovative and cost effective high performance computing facilities for Lattice QCD scientific research as well as leading the project to build and operate facilities at BNL, JLab and Fermilab. We leverage that expertise to provide smaller high performance computing facilities for Computational Cosmology and for Accelerator modeling and other parallel processing modeling applications related to accelerator design.

We provide the highly-reliable central management information systems (MIS) and services that support the Laboratory's core business and administrative systems, in areas such as finance, accounting, human resources, asset management, and cost, schedule, and performance management.

We provide the central IT infrastructure for the laboratory through site networking, site-wide Windows infrastructure, email and helpdesk systems, central web servers and many other IT services and database applications that support the entire laboratory. We have started on a process of adopting ITIL principles and have set a goal to have a critical subset of the core IT services of the lab ISO20000 certified within two years. In addition we want to more broadly adopt ITIL for scientific services.

We lead the lab in implementing and fully embracing a vigorous and continuous process for computer security and we continuously assess the threats and risks in our environment and evolve our program and protections to account for changes in the cyber environment and changes in DOE and federal regulations.

We plan for and execute projects to build and maintain (with FESS) adequate data center infrastructure to support the scientific computing and IT services of the lab.

We continuously develop and evolve our ES&H programs, our management systems to plan, track and monitor our resources and performance, our QA program, our document management systems, our formalization of work procedures and our internal communications.

All of the above efforts are effective and well serving their stakeholders. However technology is evolving rapidly in all aspects of computing and scientific services and solutions must scale with increased demands. Constant evolution of all of the services and revitalizing of expertise is essential for the success in the future.

Vision

We expect Fermilab Computing Facilities and Fermilab staff working on computing solutions (includes engineering) to continue to excel and be second to none in the high energy physics and astrophysics world.

We expect Run II computing to move into a steady stable state focused on analysis of condensed datasets with a major reprocessing of Run II data having been completed in 2010. However we expect Run II experiments to have an option to reprocess data should that be necessary for understanding a discovery or new physics. Grid computing resources at Fermilab and elsewhere would be used for such reprocessing. We expect MINOS computing to remain in steady stable state through 2012.

We envision fully provisioned CMS computing facilities and services and expect to handle vast amounts of data from the LHC at Fermilab and to support rapid access to and processing of that data for all of US-CMS and for the wider CMS experiment. We expect to be playing a strong role in making Fermilab a welcoming and efficient place to work by supporting the tools and facilities that users need to be productive. We expect use of the ROC for experiment monitoring and for accelerator monitoring to change and evolve.

We expect CD to continue to play a strong role in support of the experimental Astrophysics program and we expect CD people to be active in DES, JDEM, Auger and COUPP, and possibly super-CDMS and other experimental astrophysics programs, as resources permit. We expect the main CD technical involvement to be in planning for and building the JDEM Science Operations Center at Fermilab.

We expect to have a continuing important role in support of high performance computing (HPC) for Lattice QCD. With many large supercomputers (from both DOE and NSF) becoming available for allocation of computing time we anticipate developing a clear set of delineating features and goals for high performance computing at Fermilab – not only for Lattice QCD but for other computationally intensive science such as Accelerator modeling and Computational Cosmology.

We will offer services and computing support to Minerva, Nova, Testbeam efforts and other Intensity Frontier experiments, as they evolve. We expect this to be primarily in building up and helping experiments use our shared services, in system engineering, and in DAQ. We envisage a small amount of targeted help being provided, where it is most needed, with many activities competing for these scarce resources. We will continue to push hard on making operations of our scientific facilities and services and of our core IT services ever more efficient – to enable us to put as much effort as possible on assisting the future scientific program, even though our plate is over full with the many commitments we have.

We expect to see Grid Computing and globally distributed computing and services in full routine production mode, having evolved considerably between now and 2012 with likely both experiment computing resources, as now, and HPC computing resources made available on the Grid.

We expect to have gained some initial understanding of whether there are benefits that might be gained from cloud and/or virtualization in computing.

We expect to see fully centralized management of IT infrastructure and services at the lab, encompassing all of the currently disjoint and private IT, telephony and cyber security infrastructure of

AD and other areas of the lab. We expect to be running an increasing number of powerful systems and services for a larger number of people at the lab and worldwide. We envision receiving ISO20000 certification for the lab in core IT service delivery by 2011.

We envision leading the way in bringing in modern information management systems that can really help us all manage our work, understand our resources and performance at all levels and communicate more easily.

We envision a continuous process of hiring new staff, with an emphasis on hiring junior staff, of training and re-training of existing staff because technology will continue to evolve rapidly and everyone in CD must expect to participate in a continuous process of developing their skills.

We really hope to find a way to have considerably more scientists as members of CD by 2012, thus strengthening the successful way of working that teams up scientists, engineers, computing professionals and technical and support staff to work together on computing solutions.

We expect to continue to increase our participation in collaborative work with computer scientists, other sciences, university researchers and with NSF and DOE large facilities and to join with NCSA and ANL and others in working together to advance computing and networking support and innovation for our scientific programs.

Stakeholders

The stakeholders are the funding agencies (DOE and NSF), the lab management, divisions and sections, and the broad lab user community, the many scientific collaborations that we work with, the consortia we belong to (such as the Open Science Grid and the National Lattice QCD project), the multidisciplinary projects that we participate in (such as SciDAC projects), the other collaborative endeavors that we undertake with ANL, DESY, CERN, University of Chicago, University of Wisconsin and others, and finally the members of the CD.

Goals and Objectives

1. Provide excellent, secure, and, where necessary, innovative computing solutions to support the scientific program of the lab.
    a) Provide stability and support for the computing and core software systems to help maintain scientific productivity of the Run II, Experimental Astrophysics and Neutrino programs. Provide help to the smaller experiments and testbeam efforts of the lab, as resources permit.
    b) Ensure that the CMS Computing Facilities at Fermilab are 2^{nd} to none and that they fully support the LPC at Fermilab and all of the U.S. Tier 2 facilities in enabling rapid scientific results from CMS.
    c) Provide HPC facilities for Lattice QCD, Computational Cosmology and Accelerator modeling applications and ensure these facilities are used in conjunction with other DOE Office of Science and NSF computational facilities.

d) Provide leadership and support for the Open Science Grid and continue to support opportunistic use of Fermilab Facilities for certain science programs outside of the Fermilab scientific program.

e) Participate in the scientific and technical specifications of future scientific programs and LHC-related upgrades bringing CD experience and computing expertise.

f) Assist various scientific programs with their software infrastructure, their information system and database needs and their use of common services and supported tools (such as GEANT, ROOT) - as resources permit.

g) Carry out integrated system engineering projects in support of the scientific programs of the lab and maintain legacy electronics and DAQ systems as necessary.

2. Selectively carry out scientific computing research in areas where innovation and new ideas might lead to significant gains in functionality, cost, performance or efficiency in computing solutions.
3. Provide highly-reliable and cost-effective business and management information systems that meet laboratory needs and support laboratory operations at all levels.
4. Develop secure, efficient and cost-effective IT infrastructure and information applications to support all scientific programs, divisions, sections and centers of the lab.
5. Create a stimulating and rewarding working environment for talented scientists, computing professionals, engineers, technical and administrative staff. Support and encourage our scientists to participate in the scientific program of the lab.
6. Further develop and integrate our CD management, measurement, planning, analysis, and workflow processes and information systems to help us plan, execute, and measure our work and progress on all of the above goals.

Strategies

1. Provision the underlying facility and networking infrastructure to support the above goals.
2. Pay attention to “green” in our facility planning and usage in and all our procurements.
3. Continue to encourage the use of common and shared computing solutions.
4. Nimbly adjust over time the balance of effort on the current scientific program, the LHC programs and future scientific programs.
5. Improve operational efficiency by ensuring we have effective monitoring, incident tracking, problem resolution, release management, configuration management and customer service level agreements in place – i.e. adopt ITIL approach for all of our service delivery.
6. Ensure that all staff are given access to training, tools and incentives to develop professionally in a way that is aligned with the lab’s scientific program needs. Increase the amount of yearly job-specific and/or management training that each person in CD receives until it closely matches industry norms.
7. Improve communication and collaboration within and outside CD. Use tools such as strategic plans, structured meetings, status reports, quarterly reports, and technical seminars, communication forums and division newsletter.

8. Ensure each scientific program – current and future – has an assigned CD liaison.
9. Ensure every scientist has an active scientific mentor within or outside of CD.
10. Work closely with stakeholders and potential collaborators to identify target areas for innovation and target areas for competitive proposals for funding.
11. Increase visibility and acceptance of our work (both R&D work and production solutions) by attendance at conferences, workshops and other forums.
12. Increase usefulness and acceptance of our work by ensuring that training and education of users, as well as well-written and maintained documentation is an integral part of our services and deliverables to our stakeholders.
13. Identify next-generation technologies and tools and develop vision statements for how such technologies and tools might support our mission – in each of the strategic plans we write for major areas of CD work.

Resource Needs

The resource needs for the entire division have to come from the rollup of all of the strategic plans. We cannot plan for large growth in CD resources – even though we might like to. Instead we have to rely on judicious and flexible assignments of personnel as the programs evolve and look for efficiencies in operations and opportunities for consolidation labwide.

Progress Indicators

Each Activity area of the Computing Division's work will develop indicators that help us to assess each year whether progress is being made towards our strategic goals.

Additional Information

Our values will be, as stated in our mission statement, to put great emphasis in working collaboratively with others and in taking pride in our work. Everyone in CD will make it their personal concern to ensure that they and our entire organization work safely. No-one will take unnecessary risks or cut corners on processes and everyone will maintain awareness that an accident or incident can not only result in personal injury or damage to property, but also have far reaching consequences for the lab.

The strategic plan of the computing division is designed to accept that there may be unforeseen changes in schedule, funding or other resources and certainly unexpected changes in technology. The most important factors that will allow us to respond to such changes are a) the agility of a well trained staff; b) strong collaborations and ties with people outside of the lab; and c) a good planning, tracking and management process that allows us to fully understand how our resources are being deployed at all times.

Strategic Plan for the Scientific Research in the Computing Division (2009-2013)

Patricia McBride
10/30/2009

Mission

The Computing Division's mission is to participate fully in the mission of the laboratory. For many years, particle physics has progressed through the development and use of complex computing techniques. CD Scientists engage in all areas of the Scientific Research Program and facilitate communication and strategic planning in areas such as computing resources, detector development and algorithm development.

Context and Assessment of Current State

Scientists in the division are active in the physics research and provide a direct connection to the experimental and theoretical program at the laboratory. Over 30 CD Scientists are involved in many areas of research including the Intensity Frontier, the Cosmic Frontier, the Energy Frontier. Scientists in the division are often called upon to communicate with the science community regarding many computing issues.

CD scientists' enables a strong connection between the groups in the division and the primary mission of the laboratory. This connection enables the division to set strategic directions for data collection and processing, data archiving, simulation and analysis. The direct interaction between computing experts and the scientist has been key to the success of scientific computing at Fermilab since the technology and the science driven requirements are constantly evolving.

Scientists in the division have been involved in a broad range of scientific programs of the lab as scientific collaborators on CDF, D0, FOCUS, SELEX, KTeV, MINOS, MiniBooNE, CMS, SDSS, DES, Pierre Auger, JDEM, US Lattice QCD, COUPP and CDMS. CD scientists are participating in new scientific efforts at the lab – such as the Mu2e experiment and the LBNE neutrino experiments. In partnership with PPD and the Center for Particle Astrophysics, we plan to explore additional scientific opportunities.

Several CD scientists have been involved in the planning studies for the ILC detectors and accelerator. The division has significantly decreased the involvement in the ILC over the past year. Studies focused on the physics reach for the muon collider could be initiated if this becomes a high priority of the laboratory in the near-term.

Particle physics experiments rely on innovation in DAQ and Trigger hardware to collect and filter the large data volumes created by modern detectors. Scientists in CD work side by side with the engineering and computing staff to design and build these systems for the experiments in the lab program. They also direct computing operations activities for the experiments and take leadership roles in data production, calibration and simulation tasks.

CD scientists work on the development of new scientific software and on the validation of the scientific software code for experiments and for the theory and accelerator communities. They also participate in data production, calibration and analysis and help to design new detectors and understand existing detectors. They aim to have the highest impact on the experimental and theoretical program.

CD scientists have become leaders in detector simulation and modeling. CD scientists are leaders of COMPASS accelerator modeling project and are active in the GEANT4 project.

Scientists in the division participate in national and international scientific committees and panels often in leadership roles. They are often asked to serve on reviews and oversight committees in science and computing.

Vision

World class scientific research is a strategic goal for the laboratory and the division. The division expects CD scientists to remain active in the scientific program at the laboratory and to work with DOE and other laboratories and universities to set the direction for scientific computing in particle physics in the US.

CD scientists are expected to be in leadership roles within the scientific community, with their collaborations, and in scientific societies, committees, panels and reviews.

CD scientists are expected to lead the data operations effort and participate in CDF and D0 physics analysis through the end of Run II accelerator operations and for a number of years beyond in many cases. The scientists involved in Run II will migrate towards the Intensity Frontier program at that time.

CD scientists are expected to play a leading role in making Fermilab a welcoming place to work on CMS physics analysis. Division scientists will participate in the simulation, data production and planning for computing for the experiment. They will participate in data analysis efforts and are expected to take on leadership roles in physics.

CD scientists are expected to continue to play a strong role in the experimental Astrophysics program and to be active in DES, SNAP, Auger and COUPP, and possibly super-CDMS and other experimental astrophysics programs, as resources permit. They are also expected to be leaders in planning the JDEM Science Operations Center at Fermilab.

CD scientists are expected to have an important role in high performance computing (HPC) for Lattice QCD and in the scientific program. We anticipate developing a strategy not only for Lattice QCD but for other computationally intensive science such as Accelerator modeling and Computational Cosmology.

CD scientists are expected to continue to play leadership roles in accelerator modeling and detector simulations both at the laboratory and in the context of national and international collaborations.

The effort in the Intensity Frontier area is expected to grow: in the neutrino program, in Mu2e and in development of the kaon program. Scientists from the division will help to develop strategies for simulations, for offline software and computing and for trigger and DAQ for the experiments in the Intensity Frontier program. They will participate in the development of the physics program and in detector design studies. They will aide in the development of computing facilities for the Intensity Frontier and other future scientific programs at the laboratory.

R&D for future detectors and computing systems both offline and online is part of the scientific mission of the laboratory. Division scientists and staff will continue to be part of this effort. They will continue to work to develop algorithms and software for effective data collection and monitoring.

The division hopes to have more scientists at all levels as members of CD by 2012, thus strengthening the successful way of working that teams up scientists, engineers, computing professionals and technical and support staff to work together on computing solutions.

Mentoring with students and postdocs helps to drive innovation in scientific computing and to push the division to keep up with modern trends in computing as well as in the science. The division will pursue avenues to increase the connection to the postdocs and fellows at the laboratory.

CD Scientists will continue to play important roles in the management of the division.

Stakeholders

The stakeholders are the many scientific collaborations that we participate in, the consortia we belong to (such as the National Lattice QCD project), the multidisciplinary projects that we participate in (such as SciDAC projects), the other collaborative endeavors that we undertake with universities and other laboratories, and finally the funding agencies and the members of the CD.

Goals and Objectives

1. Enable the best science possible and communicate the scientific and strategic goals of the laboratory.
2. Provide leadership for the computing and software systems to maintain the scientific productivity of the Run II, Experimental Astrophysics and Neutrino programs. Provide help and guidance to all experiments and testbeam efforts of the lab, as resources permit. Participate as collaborators in data-taking and physics analysis.
3. Ensure the success of the CMS Scientific program. Help to manage the CMS computing and software efforts at Fermilab, support the LPC at Fermilab and enable rapid scientific results from CMS once data taking starts up. Participate in CMS commissioning, data-taking and physics analysis.
4. Provide leadership for the HPC facilities for Lattice QCD, Computational Cosmology and Accelerator modeling applications and ensure these facilities are used in conjunction with other DOE Office of Science and NSF computational facilities.
5. Provide leadership for the Fermilab Computing Facilities to support of the HEP science program and for science programs outside of the Fermilab scientific program.
6. Develop and support scientific software for detector simulations, accelerator simulations and physics generators. Provide consulting for the user community and lead simulation studies for the current and future laboratory program when resources permit.
7. Participate in the scientific and technical specifications of future scientific programs at the laboratory and LHC-related upgrades.

8. Advise the scientific programs in planning their software frameworks and use of common services and supported tools (such as GEANT, ROOT).
9. Plan integrated electronics engineering projects in support of the scientific program and plan for support of legacy electronics and DAQ systems as necessary.
10. Carry out scientific computing research in areas where innovation and new ideas might lead to significant gains in functionality, cost, performance or efficiency in computing solutions.
11. Create an open, secure, stimulating and rewarding working environment for science.

Strategies

1. Maintain a balance of effort on the current scientific program, the LHC programs and future scientific programs.
2. Encourage and support the direct participation of scientists in the division in the scientific program and in scientific conferences.
3. As more of the lab's scientific activities move off-site, enable scientists to travel to collaboration meetings and operations activities.
4. Ensure that the scientific staff is given access to training and incentives to develop professionally in accordance the lab's scientific program needs.
5. Support participation in national and international scientific organizations such as APS.
6. Ensure each scientific program – current and future – has adequate representation with the division and is assigned a CD liaison.
7. Ensure every scientist has an active scientific mentor within or outside of CD.

Resource Needs

The resources for the scientific program come from several budget and reporting (B&R) categories including Proton Research, Astrophysics, Accelerator Research, Electron Research, Detector R&D, Computational HEP and Theory. The needs and requirements for the scientific research program are determined through a dialogue among the division office, the directorate, the scientific centers, the collaborations and the individual scientists.

Progress Indicators

Scientific progress will be measured by scientific publications, by technical notes, by leadership positions within the scientific collaborations and by recognition by scientific societies.

Additional Information

Our values will be, as stated in our mission statement, to put great emphasis in working collaboratively with others and in taking pride in our work. The excellence and integrity of the scientific results are the marks of success for the laboratory. Every CD scientist will make it their personal concern to ensure that they and our entire organization work safely. No one will take unnecessary risks or cut corners on processes and everyone will maintain awareness that an accident or incident can not only result in personal injury or damage to property, but also have far reaching consequences for the lab.

Strategic Plan for Run II Computing

Qizhong Li and Rick Snider
9/14/2009

Mission

To provide support for the computing needed to meet the physics goals for the Run II experiments via innovative, cost effective and efficient computing solutions in designed to maximize the research productivity of the participating physicists.

Context and Assessment of Current State

The Computing Division currently provides support for all aspects of the major hardware computing systems, data handling and storage, analysis software infrastructure, database and network administration. The CD also supports raw data logging and archiving at D0, raw data archiving services at CDF, desktop OS support at CDF, production operations at both experiments, and online systems administration at both CDF and D0. Support for the grid computing software infrastructure, database applications, project disk and production operations are provided jointly by the experiments and the CD. The experiments alone provide support for all software pertaining directly to the detector, reconstruction or physics analyses. CDF also supports a stand-alone dCache pool used for transient data storage. A proprietary job submission and workflow management system supported by CDF during the past five years will be phased out by the start of FY2010 and replaced by an equivalent system supported by the CD. D0 holds full support responsibility for an on-site analysis cluster, clued0.

Many critical systems used by these experiments, such as data handling, are mature. Continued gains in stability, robustness and reductions in operational workload should be expected in these cases, although with diminishing returns on investment. The grid computing systems of the experiments are approaching maturity, but still require significant effort to maintain and operate. A few issues, such as movement of CDF data between Fermilab and remote grid computing sites, may require additional development in order to meet the physics goals of the experiment. The continued linear growth of the datasets and resulting computing demand of the experiments may continue to stress the scalability of all systems.

As the experiments reach the end of data taking, the available resources from the collaborations will decline. It is imperative that the CD be in a position to provide on-going support for all critical computing systems required by the experiments throughout this period and, if required, to take over responsibility for those critical systems at risk of losing support from the experiment. These goals must be achieved within budget constraints that are at best flat and probably declining.

Vision

Major developments by the experiments have been completed. System are transitioning into a maintenance phase that includes implementing changes directed at improved robustness and reduced operational load for both service providers and users. Service and infrastructure monitoring will be increasingly automated; typical service failures will be detected by the monitoring rather than by users, and will alert service providers accordingly. Most if not all simple and common operational tasks will be automated, with significant progress made toward automating more complicated tasks. Systems and procedures will be documented.

Most computing resources used by the experiments will be accessible via grid portals. Production activities and data analysis will use these portals when appropriate and cost effective. Data movement between off-site computing and on-site data handling systems will proceed via fully grid-compliant mechanisms.

Effort provided by the CD will become increasingly directed at simplifying the complexities and workload associated with utilizing the available computing facilities.

The operational load required to run the computing infrastructure for the Run II experiments will begin to decline before the end of this year, with further drops throughout the remainder of the experiments.

Stakeholders

The stakeholders for Run II computing are the CDF and D0 collaborations, the Computing Division, Fermilab management and the US and foreign funding agencies.

The customers of Run II computing are the CDF and D0 collaborations.

Goals and Objectives

1. Provide sufficient computing to the Run II experiments to meet the physics goals of each.
2. Support on-going computing operations at CDF and D0.
3. Migrate on-site computing farms into Fermigrid where possible and advisable. (Migrating clued0 or all CAB resources into Fermigrid are not goals of this Strategic Plan or the D0 Collaboration.)
4. Support grid-compliant solutions to data movement and job handling.
5. Expand the role of Grid computing within the experiments to the extent possible and advisable. (Execpt for selected CPU-intensive applications, expansion of D0 analysis computing to the Grid is not a goal of this Strategic Plan or the D0 Collaboration.)Transition as many systems as possible into a low-cost maintenance phase.
6. Improve the robustness, efficiency and usability of the computing systems.
7. Reduce the workload required to operate the combined computing systems at the experiments.
8. Improve and expand the level of automation available for system operations, monitoring, data processing and user services.

Strategies

Both experiments should adopt grid-based computing solutions as allowed within the available effort. The scope of computing included within this strategy should balance the cost of implementation and future support against possible gains in efficiency, synergies with existing projects and support workload reductions. Such solutions will benefit the experiments by enabling the use of a more diverse array of computing resources around the world, and better consolidating the support effort needed across the experiments. These solutions will benefit the lab by better enabling economies of scale to apply to the computing plant and associated grid infrastructure, and by providing a more cost effective solution to the peak loading problems presented by each individual stakeholder. The CD should provide leadership to the experiments within this area. We should note that the experiments may have already identified certain computing problems for which grid-based solutions are not cost effective.

For the remaining development needed to improve operations, the CD should take a leadership role in understanding requirements, defining supportable solutions and planning and managing projects. Clear milestones, performance metrics and documentation requirements should be an integral part of every project plan. Adopting common, community or lab supported solutions whenever feasible can help reduce the support load by leveraging effort external to the Run II experiments or REX department.

Synergies with existing and future projects should be exploited to the maximum extent possible. The Run II experiments share many common elements and strategies with MINOS, CMS and other active or proposed experiments at the Lab. The data handling and storage systems for these experiments, in particular, have benefited greatly from the use of common solutions provided within a central services model. The Run II experiments and the CD should continue to seek new ways to leverage effort across projects in all development efforts. The grid computing and data movement issues at the Run II experiments are prime examples of where such synergies might result in significant savings.

A high priority should be placed on increasing stability and robustness since these issues feed directly into operational workload.

Solutions should pay particular care to end-user workflow and usage in order to provide solutions that are effective.

The CD has the goal of bringing all IT management at the lab into ISO 20000 compliance. For the Run II experiments, this transition is expected to improve the quality of IT services and procedures provided by the CD. REX should assist in bringing these services and procedures into compliance.

Resource Needs

Through the course of FY2009, the department added sufficient effort to the Run II experiments to fully staff the computing projects. The demand for this effort was expected to remain high for most of the year. As major development work at CDF ended in the third quarter of the year, effort began to shift away to other tasks. Although further declines were anticipated into the beginning of FY2010, some effort has been retained to assist with the development needed to improve operations. D0 completed a major initiative to complete the transition to grid processing of raw data production. Again, some effort has been retained to assist with improving operations. After achieving operational improvements, the department intends to move operations into a maintenance mode in which a smaller number of people collectively share support responsibilities across the full range of services supported by the department. Flat funding for human resources through the end of data taking may be necessary in order to carry the increased load created by the migration of effort from the experiments, regardless of whether data taking ends in 2009 or 2010.
Additional effort will be required to address ISO 20000 compliance issues

A flat funding profile for equipment through the end of data taking is expected and will be adequate to provide for the computing needs through this period.

Once data taking ends, the support for production activities will become greatly reduced, as will the need for the associated hardware. Some moderate to large scale re-processing should be expected to occur during the first two years after data taking. The resources required to perform such activities should be retained in some form during this period.

Based upon the progress of existing analyses, we anticipate the need for significant analysis computing to continue three to four years after data-taking. It is possible that a small number of analyses will extend into a fifth year after data-taking. The scale of available analysis computing assets should remain level for at least three years, possibly ramping downward thereafter. This subject should be re-evaluated as the end of data-taking approaches and the scale of post-data taking analysis becomes evident.

All activities directly related to raw data taking and archiving will cease at the end of data taking.

Progress Indicators

Physics output: The primary objective for the computing project is to support the physics at the experiments. A sustained physics output with no noted or significant limitations imposed by the available computing systems is the most important metric.

Grid computing: All available slots on computers purchased by the experiments are utilized with high CPU efficiency. The scale of local operations is expected to be between 5k and 10k slots with on the order of 10k successful jobs completed per day.

Robustness and operational workload: A significant decline in the rate of problem reports from users will be an important indicator of progress. The number of manual interventions required per week on a particular system and the time expended in addressing problems will demonstrate the level of stability for that system, and provide a metric related to operational workload.

User efficiency enhancements: These are difficult to measure. The number of physics analyses per active researcher will provide some clue as to improvements in the usability of the systems. The time difference between data taking and first results, and the number of such first results also provides some estimate of researcher productivity. Timely publication of the “core analyses” as described in the Run II Task Force Report from March, 2006, is an essential goal and should be considered a baseline metric for the success of analysis computing.

Strategic Plan for CMS and LHC (2009-2012)

Patricia McBride, Ian Fisk, Lothar Bauerdick
8/12/2009

The CMS and LHC Strategic Plan covers the coming three years and is updated annually.
The LHC strategic plans have always covered a period of development followed by an operations period with data collection. It is not the first time that CMS has hoped that the period of stable data collection is experienced before the next revision of this document, but CMS is particularly confident that this time the transition to operations will occur within the time frame addressed in the strategic plan. The Tier-1 facility has completed the extended procurement ramp arriving at the full complexity Tier-1 facility; the detector has collected cosmic data during two extended runs; the distributed computing infrastructure has been exercised at the expected scale; and while continued development effort is needed in several key areas, functional systems for early running exist.

Mission

The mission is to develop, innovate, operate, and support excellent and forefront computing solutions and services for the CMS experiment at CERN, and the U.S. CMS Software and Computing Project. In particular, we have build and will operate a Tier-1 regional computing center as part of the Worldwide LHC Computing Grid and a CMS analysis computing facility for the LHC Physics Center at Fermilab, and to collaborate closely with similar efforts at other CMS institutions in the US and internationally. It is also our goal to facilitate analysis, remote collaboration, and remote operations at Fermilab for the CMS collaboration and for the LHC accelerator.

Context and Assessment of Current State

The LHC program at FNAL is comprised of several related activities. FNAL is the U.S. host lab for the U.S. CMS project. Activities include detector maintenance and operations, detector upgrade R&D, operation of a remote operations center for remote detector operations, and data operations for CMS and for LHC accelerator monitoring. FNAL hosts the CMS Tier-1 computing facility for the Americas. The core group of application framework software developers and a significant fraction of the distributed computing and data management developers are employed by FNAL. US-CMS is a leader in the development and deployment of the Open Science Grid that benefits both US-LHC experiments and other science communities.

The LHC and CMS monitoring activities have successfully demonstrated the ability to contribute to operations at a distance. The Remote Operations Center (ROC) final facility is complete and has successfully contributed to the Magnet Test and Cosmic Challenges (MTCC); 2008 and 2009 CMS Cosmic Tests in the P5 cavern (CRAFT, CruZet); and first beam and tracker integration activities, including shift work, real-time monitoring, and data processing and reprocessing. The LHC at FNAL facility is a model for remote operation centers, and technical challenges regarding secure access to CERN-based systems are being overcome. The successful utilization of the LHC at FNAL center has spurred the deployment of 2 other large remote centers in Hamburg and Beijing.

The Tier-1 computing facility has completed the final phase of the procurement and deployment cycle, which successfully achieved 100% scale of all important capacity metrics by the end of FY 2008. The

Tier-1 facility is utilizing a node replacement model where nodes are replaced with faster hardware to accommodate the experiments need for increased capacity, while keeping the center complexity roughly constant.

Software developers at FNAL are drivers in the development and support in 3 areas critical to CMS success: the CMS Software Framework, the CMS Workflow Management, and the CMS Data Management projects.

- The CMS Software Framework was redesigned using FNAL expertise in 2006. This massive undertaking, involving coordinating contributions from global CMS, was completed on schedule. The software was successfully validated during subsequent computing and software challenges and it now in the 3rd major release. It is expected that over the time period covered in this strategic plan the framework support effort will drop significantly, but if previous experimental experience is any indicator the effort will not drop during the first 18-24 months of data collection.
- The CMS workflow management system controls the organized processing for reconstruction, skimming, and simulation at the Tier-0 at CERN, the 7 national Tier-1 centers and the 40 Tier-2 computing centers. Additionally the CMS workflow project develops the analysis tools used by the thousands of the physicists on CMS for submission to Tier-2 and Tier-3 computing centers. CMS at FNAL is not the only contributor to the workflow project, with leadership and development for analysis tools coming through INFN, but FNAL developers have acted as the architects and lead implementers for the organized processing components and several of the shared components.
- The Data Management project in CMS includes the Dataset Bookkeeping System (DBS), the Data Transfer System (PhEDEx), and the Distributed Conditions Database (Frontier). CMS at FNAL provides developers critical to the success of DBS and Frontier.

The Open Science Grid receives funding from both the DOE and the NSF for running the Grid as a distributed facility and to enable it to contribute to grid based storage, higher level distributed computing services, packaging and deployment, and support. US-CMS continues to play a leading role in the OSG, and to benefit from advances provided by OSG efforts.

Vision

The US-CMS Tier-1 computing facility at FNAL will remain the largest and most capable remote computing center for CMS. US-CMS will meet its obligations to the international experiment with processing, storage, and network resources commensurate with the size of the US fraction of the collaboration. FNAL will continue to bring to bear the professional expertise at facility operations and deployment required to support the Tier-1 center.

By 2010, FNAL will be the best place in the world to perform CMS physics analysis. The combination of a critical mass of expertise at the LHC Physics Center (LPC), adequate computing resources to support the local analysis community, access to locally stored data samples, and a strong connection to experiment monitoring and operations will result in a world-class analysis center.

The Remote Operations Center at FNAL will provide a unique opportunity in the US for CMS collaborators to participate in detector operations and data monitoring. The operations center at FNAL will foster more inclusive participation in shift work in the US.

The software development teams at FNAL will deliver and support critical components for the experiment success in the areas of framework, workflow management, and data management. CMS will continue to benefit from the institutional expertise available at FNAL and FNAL will continue to be seen as an innovative force in HEP software development.

Stakeholders

The sponsors of the CMS and LHC work are the U.S. LHC Research Program funded by the DOE and the NSF, and the Fermilab core program.

Customers are the Fermilab and University scientists doing CMS research, and the CMS collaboration.

Effort and deliverables are provided from many groups both internal and external to the Laboratory.

Goals and Objectives

1. Operate a reliable Tier-1 computing facility that meets the requirements for custodial data storage, data serving capacity, reconstruction processing capacity, and local and wide area data serving capacity as outlined in the CMS computing technical design report (CTDR).
2. Provide a functional data operations team to peer with the CERN-based team. The CMS data operations model calls for two teams: one team operating at CERN during the European day and one operating at FNAL during the U.S. day. The teams have equivalent responsibility and authority during their respective shifts.
3. Provide sufficient local analysis computing, user storage, and data serving capacity to support a community of roughly 200 participating scientists.
4. Develop, support and operate distributed computing infrastructure to make efficiency use of the CMS dedicated and opportunistic resources through the OSG common grid infrastructure. This includes the use of processing resources for simulation, analysis, and event reconstruction as well as the delivery of data management components to make efficient use of distributed storage resources.
5. Facilitate the development, deployment, and support of the CMS core framework to give a competitive scientific advantage to CMS in data analysis.
6. Provide opportunities for U.S. physicists to actively participate in detector commissioning and operations from the U.S. through remote monitoring.

Strategies

For the facility, CMS utilizes a strategy of active monitoring of the centers to measure readiness and reliability. Augmenting this is a detailed set of site monitoring and alarms. The CMS Tier-1 has moved from a facility doubling procurement model, to a node replacement model with the goal of maintaining a similar level of complexity.   Additionally, for the last two years the Tier-1 has used a buy ahead model where approximately one third of the capacity needed in a fiscal year is procured at the end of the previous year to provide extra time for commissioning and ensure systems are in production early in the year. This will be continued as the machine schedule merits and the budget allows.

In workflow and data management as well as framework development, US-CMS has practiced active engagement with the international experiment, while maintaining a strong local core team.

CMS remains committed to the success of Open Science Grid and maintains close integration with OSG development and deployment projects.

Resource Needs

We expect the engineering resource needs profile to slowly decrease by 2-3 FTE total through these three years, in terms of the overall FTE count, due to flat funding. There is the possibility of compensating effort through the CMS upgrade project. The biggest change will be the focus of the supported effort from roughly equal amounts of operations and development effort currently to a more solid majority in operations. In the first 2 years of operations we expect the development needs will stay high, with more significant reductions in year 3.   The physics effort will increase and additional scientists, RAs and guest scientists will be required to start the physics program of the LPC.

The funding profile for equipment procurements will average to about flat with some structure year-to-year based on the machine schedule.

Progress Indicators

CMS will transition from a preparing to a running experiment beginning with the start of LHC operations in November 2009.   The goal for CMS is to achieve a level of service reliability and transparency in the distributed services comparable to what can be achieved with local access to computing resources.

By January 2010

- CMS expects be able to successfully execute 25,000 jobs per day on the FNAL batch computing resources through a combination of local submissions to the LPC and grid submissions to the Tier-1.
- The local community should be able to submit 20,000 jobs per day across LPC analysis and remote resources.
- CMS expects to be able to reliably accept 400MB/s of data from CERN to tape.
    - Time from deciding to transfer data to achieving the desired throughput should be less than 1 hour.
- The goal is to provide transfers from FNAL storage to Tier-2 centers at between 50MB/s, to the worst connected Tier-2 centers, and 500MB/s for the best connected Tier-2 centers.

By July of 2010

- CMS expects that the number of job submissions to FNAL and to U.S. distributed resources will increase slowly, but that the volume of data being accessed and the diversity of samples requested will increase.
- The Tier-1 center will execute at least 4 reprocessing passes of the custodial raw data.

- In 2010 the accelerator is likely to be collecting low luminosity data until April, but ramping toward higher luminosity. The physics groups at the LPC will have access to sufficiently large samples that analyses of rarer signals will be possible.

Strategic Plan for Intensity Frontier Experiments (2010-2014)

Lee Lueking
8/11/09

Mission

Our mission is to provide the agreed upon online and offline computing resources for the experiments that are part of Fermilab's Intensity Frontier program. This includes several neutrino related efforts, Mu to e, and g minus 2 experiments. These experiments have a broad range of time frames and physics goals, but similar needs in terms of the offline computing architectures to achieve their goals. Computing Division also provides some services and assistance to meet online needs by special agreement and contained in each experiment's MoU with CD.

Context and Assessment of Current State

The experiments included here have differing scales of need and a broad range of schedules. Today, two of the experiments (MINOS and MiniBooNE) have been operating for several years and their requirements are well understood. One experiment (MINERvA) is coming on line now, and another (NOvA) will begin commissioning an integration prototype early in 2010, and their full detectors in 2013. The Liquid Argon detector program has a significant prototype (ArgoNeuT) taking data through 2010 to be followed by MicroBooNE in 2011. The remaining experiments are in design phases but significant resources are needed to produce simulations and help to build and shape the Intensity Frontier program of tomorrow.

Vision

The Intensity Frontier program rivals in size and scope the Run II programs at Fermilab. Coordinating computing resources for a large number of small experiments is in many ways more challenging than the Run II efforts due to the many potential models of data processing and analysis. The vision of the Computing Division support for this part of the Fermilab program is to provide communication among, and common solutions for, the many collaborations with a broad range of needs.

Stakeholders

The following experiments are the stakeholders for the CD resources and services described here:

1. MINOS
2. MiniBooNE
3. MIPP
4. MINREvA
5. NOvA
6. ArgoNeut
7. MicroBooNE
8. DUSEL Water Cerenkov detector
9. DUSEL Liquid Argonne detector

10. Mu2e
11. G minus 2

In addition, other experiments and projects may be included as they are identified and/or require computing resources that can be provided under this umbrella.

Goals and Objectives

1. Identify, understand and document the computing needs of the stakeholders
2. Work with the computing division to identify existing solutions, or design new ones, that will meet the needs and timeframes of the stakeholders
3. Provide software development and support in the areas of physics applications
4. Provide software development and support in the areas of data and production management
5. Provide operational tools and support for using the computing resources provided by Fermilab CD.

Strategies

1. To achieve the goals with a minimum of manpower common solutions are needed that extend to all of the experiments needs
2. Communication among all the stake holders and the computing division in the form of meetings, mailing lists, and web based applications.
3. Since there is a large overlap of collaborators across many of the experiments (especially in the Neutrino area) having common tools and architectures for computing makes moving among the various experiments easier.
4. Computing hardware needs to be flexibly assignable to specific experiments as needs arise.
5. Using computing resources external to Fermilab must be easily integrated in the resource needs for all or most experiments.

Resource Needs

1. Manpower of the order of 6-7 FTE's
2. Central disk storage to grow to over 1.2 PB by 2014
3. Interactive login computing resources of 400 cores by 2014
4. GRID processing resources of over 2200 job slots by 2014
5. Central archive storage (Enstore tape) of 5.5 PB by 2014

Progress Indicators

1. Meeting the needs of the running experiments throughout 2010
2. Providing resources for experiments beginning in 2010
3. Monte Carlo production milestones for all experiments
4. Architecting and building interactive cluster
5. Providing scientific computing needs for all Intensity Frontier Programs
6. Scientific publications

Strategic Plan for The Cosmic Frontier (2009-2011)

Steve Kent
7/28/2009

Mission

The mission of the CD Experimental Astrophysics program is to explore the cosmic frontier.

The Experimental Astrophysics program enables Fermilab to collaborate with universities and other institutions internationally in projects that are aligned with the mission of the laboratory, that are at the forefront of research in their particular field, and for which the laboratory can contribute unique benefit.

Additionally, some members of the Computing Division participate in experimental astrophysics projects that are led by other parts of the Laboratory.

Context and Assessment of Current State

The program is primarily centered on the Experimental Astrophysics Group (EAG), but includes participants from several other departments throughout the division. The program operates under the umbrella of the Center for Particle Astrophysics (FCPA), which oversees the overall program within the Laboratory. The Dark Energy Survey (DES) and its two major components - the camera project (DECAM) and Data Management, are in the construction phase. The NASA/DOE Joint Dark Energy Mission, which supercedes the SuperNova Acceleration Probe (SNAP), is being reviewed as part of the Astro2010 Decadal Survey. The Auger project is nearing completion of construction and is in full operation. The Cold Dark Matter Search (CDMS) is in operations at Soudan. COUPP are just beginning underground operation of a rebuild of their current chamber which is now 4 kg, and are engaged in the system commissioning of a 60 kg chamber which will begin underground operation at Fermilab when commissioning is successfully completed. The Computational Cosmology Initiative (CCI), being led by the Theoretical Astropysics Group in PPD but will involvement of CD members, is producing its first results. The 21 cm BAO dark energy project is is in an intense phase of feasibility studies and R&D. The Holometer project, which seeks to detect Planck scale effects, is approaching the proposal stage. The Sloan Digital Sky Survey (SDSS) has completed construction of the long-term archive.

Vision

The vision for the Dark Energy Survey (DES) is to deliver a working instrument and operations systems so it can deliver fundamental new science in the first year of operations. Since the character of the NASA/DOE Joint Dark Energy Mission (JDEM), (which has superceded SNAP) has changed in the past year, the vision is to maintain a minimal but still critical mass of effort in anticipation of an Announcement of Opportunity in order to ensure that Fermilab will be in a position to lead a science proposal and to have an opportunity to host the mission's Science Operations Center. The vision for the Auger project is to confirm and understand the apparent anisotropy signal, which will have a seminal impact on the field of cosmic ray physics. The vision for CDMS is either to detect or provide the world's best upper limits to the detection of Dark Matter. COUPP is advancing the sensitivity frontier in the search for the direct detection of Dark Matter using heavy liquid, room temperature, bubble chambers as nuclear recoil detectors for the observation of dark matter scattering. They hope to either detect or else achieve the best cross-section limits for both spin dependent and spin independent dark matter-nucleus scattering. The vision for the Holometer project is to develop the scientific case and technical capabilities needed for the experiment so it can be executed in three years. The vision for the 21 cm BAO project is tentative, since it involves an international collaboration and will be large funded by external sources. The vision for the SDSS is to maintain the data archive and mine it for new science results for the next 5 years. The vision for the Cosmological Computing Initiative is that it deliver simulations of use to DES in three years. The vision for any future projects, such as the Large Synoptic Survey Telescope (LSST) is to maintain contact with the projects and people leading them.

While the core size of EAG is expected to remain stable, it is expected that the Experimental Astrophysics program will continue to attract collaborators from other experiments within the laboratory as the makeup of the laboratory's program changes in the next several years. The Fermilab Long Range Plan envisions growth in the overall astrophysics program of a factor of 3 over its current size in the next 5 to 10 years.

Stakeholders

The stakeholders for each project consist of: 1) the funding agencies; 2) collaborators both within and external to Fermilab; and 3) the public scientific community at large. Unlike traditional particle physics experiments, data from particle astrophysics experiments often have legacy value and utility to researchers outside the experiment.

Goals and Objectives

1. Dark Energy Survey:
    a. DECAM: Work with DES collaboration to deliver DECAM to CTIO and to conduct survey planning and mock data challenges to prepare for start of survey operations by 2011. Delivery of DECam includes participating in installation and commissioning and ongoing technical support of Fermilab deliverables.
    b. DESDM: Provide a secondary DES archive node containing a full copy of the raw and processed data.
    c. Develop a plan for DES survey operations and science research support.
    d. Participate in DES operations and science research.
2. JDEM: Work with the members of the original SNAP collaboration to develop a winning dark energy proposal for the NASA/DOE Joint Dark Energy Mission in response to an anticipated Announcement of Opportunity expected sometime in calendar year 2010. Develop a prototype of the Science Operations Center, leading to preparation of a compelling proposal.
3. CDMS: Work with the CDMS group in PPD to ensure that data are acquired and processed in a timely fashion.
4. COUPP: Commission and operate the 60 kg chamber underground at Fermilab until acceptable operation and background rates are achieved. Once these goals are achieved the plan is to move this chamber deep underground at SNOLAB (~late 2010). The potential improvement of the dark matter limits is a factor of >1000 beyond those of the COUPP 2008 Science publication.

5. AUGER: Work with the Auger group in PPD to ensure that data are processed in a timely fashion while also developing the science case for Auger North. Analyze data to monitor stability of instrument subsystems, and propose for new hardware and/or software tests as necessary.
6. Holometer: Work with member of PPD and external collaborators to design and construct the experiment and operate.
7. 21 cm BAO: Work with the collaboration to prepare a design report, build the collaboration, and develop a proposal for a10% prototype.
8. Computational Cosmology Initiative: Ensure that CCI simulations meet the requirements for analysis of DES and JDEM data and are competitive with the best simulations in the world.
9. Ensure that SDSS archive continues to operate for the benefit of Fermilab scientists, the SDSS collaboration, and the scientific community and general public.
10. FUTURE: Identify and develop new opportunities in Experimental Particle Astrophysics. [See http://astro.fnal.gov/Retreat/Retreat0409/index.html and http://astro.fnal.gov/Retreat/Retreat0609/index.html.]

Strategies

The broad strategies to achieve the strategic goals are:

1) provide a productive working environment for astronomical data collection and analysis;
2) utilize shared resources within the Computing Division, especially those developed for larger physics experiments, where possible, to maximize efficiency;
3) leverage the Computing Division and EAG's expertise in conducting large, data-intensive surveys, by applying common knowledge and solutions to similar areas in different projects;
4) collaborate with colleagues from affiliated project institutions.

Resource Needs

At present EAG has12 members (including postdocs) - 10 scientists and 2 computational physicists. Three scientists in other departments in CD participate at least part-time. Additionally, the particle astrophysics projects draw on computing professionals, engineers, and technician in other departments in CD and with scientists and personnel in other divisions. Through the FCPA, the particle astrophysics projects work closely with scientists in other divisions and numerous visitors. To maximize the effectiveness of these people, space within the FCPA is required to accommodate them.

Progress Indicators

For DES, the major progress indicators are passing DOE reviews and meeting the various project milestones, particularly those related to data challenges.
For JDEM, progress is defined by achieving internal project milestones, working towards the delivery of science and SOC proposals.
For CDMS, progress is defined by data release publications, with analysis.
For COUPP, fundamental progress in this science is controlled by the rate of un-controlled backgrounds measured in events per kilogram per day. This translates into published cross-section upper limits every year or two. The migration of chambers to lower depths: e.g. the 60 kg from commissioning to underground at Fermilab to SNOLAB are also milestones.
For AUGER, progress is defined by data release publications, with analysis.

For CCI, progress is defined by the size (volume of universe), detail (number of particles) and accuracy (physics) of the simulations.
For the Holometer, progress will be measured by submission, approval and successful funding of the proposal.

Additional Information

From time to time, opportunities arise to broaden the Particle Astrophysics and the Computing Division's participation in these projects in areas where the Computing Division has significant strengths.

Strategic Plan for Network and Virtual Services (FY2010)

David Coder, Ray Pasetes, Mike Rosier, Phil DeMar
8/14/2009

Mission

- Provide all aspects of networking support: architecture, design, acquisition, installation, operation, monitoring, maintenance, and documentation of cabling plant, device infrastructure, and network services necessary to support the Laboratory's network needs.
- Provide centrally managed virtualization services to support the Laboratory’s scientific and business needs.
- Provide centrally managed storage and data protection in support of the Laboratory programs. Work closely with our customers to deliver storage solutions that meet or exceed our customer’s expectations while maintaining the manageability, flexibility, functionality and security of the services we provide.
- Work with collaborations having special distributed system needs to achieve their performance requirements.
- Engage in network research and development efforts that directly benefit the Laboratory’s pursuit of its scientific mission.

Context and Assessment of Current State

The strategic plan for networking is guided by the strategic plan for Central Services, the strategic plan for Scientific Computing Facilities, the Computing Division Strategic Plan and the strategic directions of the Laboratory; MOU’s, SLA’s and OLA’s established with internal and external organizations; and by Computer Security guidance and directives.

Responsibility for network infrastructure and services at the Laboratory, except for the Accelerator Division’s local-area networks, resides with the Computing Division (CD):

- The Network Services (CD/LSCS/NVS/NS) Group is responsible for data forwarding, the physical data cable plant, equipment installation services, network support and services such as DNS, DHCP, NTP.
- The Storage Network Services (CD/LSCS/NVS/SNS) Group is responsible for storage area networking (SAN) and network attached storage (NAS) servers, storage arrays, storage switching and access controls.
- The Virtual Services (CD/LSCS/NVS/VS) Group is responsible for promoting and implementing enterprise-class virtualization services based on a virtual server infrastructure, including host servers, virtual network switching, virtual storage and virtual infrastructure security. The Virtual Services group will build upon the basic virtual infrastructures of the former MIS Department, TD and CD.
- The Wide Area Networking and Network Research (CD/SCF/DMS/WAN) Group is responsible for providing off-site network services and supporting wide-area network infrastructure, including shared management of the ESnet Chicago-area MAN, with neighboring Argonne National Laboratory. The Wide-Area Networking group is also responsible for network infrastructure support of large scale, high performance scientific computing facilities dedicated to high impact distributed movement on a wide-area scale, such as the US-CMS Tier-1 Center. Network

research & development associated with scientific computing facilities is also a group responsibility.

Network Services

The network consists of a mixture of current and obsolete technologies and is reaching capacity limits in several areas:

- Approximately 25% of the network equipment is classified as End-of-Life (EOL) or End-of-Sale (EOS) by the manufacturer;
- Network hubs and 10Base2 Thin-Net coaxial cable is still in production use, though such technology has been largely extinct in networks for a decade;
- Many buildings in the fixed-target areas and Site 37/38/39 areas are connected by aging multimode fiber optic cable. This cable may be capable of supporting 1Gb connections in some areas depending on the condition of the fiber and the distance between switches but is not capable of 1Gb in all areas, nor capable of 10GB in any area;
- Chassis-based switches are not capable of operating at maximum performance due to the mixture of "Classic", CEF256 and CEF720 modules in the same chassis, plus most 6500-series switches are operating in Hybrid IOS/CatOS mode rather than in Native IOS mode;
- The Wilson Hall network consists of managed and unmanaged switches, and also unmanaged hubs. Fiber optic cables from offices and cubes to the 8th floor Fiber Central terminate in obsolete MT-RJ connectors, requiring new network interface modules in Fiber Central switches and new terminations to support gigabit uplinks;
- The DNS system frequently fails and needs to be replaced with an IP Address Management (IPAM) solution to accommodate IPv6 and mandated DNSSEC requirements.
- The DHCP system has reached its capacity and is often overwhelmed by address requests when large numbers of client systems request IP addresses, such as when there are conferences on site.
- Large numbers of high performance computers and storage servers are being added with corresponding saturation of network interfaces. Many switches and interface modules need to be upgraded from 1Gb to 10Gb or from 10Gb to 20Gb or 40Gb with corresponding cable upgrades from multimode to single mode within buildings.

In FY2009, one external assessment of networking was completed. This assessment evaluated the following areas: Network Infrastructure, Critical Network Services, Network Backups, Strategic Plan, Network Management, VoIP, WAN Circuits and Internet Connectivity. This assessment resulted in a series of action items to improve the performance and reliability of network infrastructure and services.

Wide Area Networking and Network Research

The current state of the Fermilab wide-area network infrastructure is a highly reliable, high performance, capacious facility operated at the forward edge of network technology. The Chicago metropolitan-area network (MAN) provides the Laboratory with a resilient, flexible, and scalable optical network infrastructure that is capable of satisfying it current and projected off-site bandwidth needs. The Laboratory network perimeter infrastructure is redundant and capacious. The US-CMS Tier-1 Center's network infrastructure has been upgraded to very high performance, high density Cisco Nexus 7000 core switching devices, and is well positioned to accommodate the extremely high local and wide-area network demands of that facility. Network research efforts that align with and directly support the

Laboratory's wide-area data movement activities are under way. The major focus area for future enhancements is upgrading the MAN and WAN perimeter infrastructures to be capable of supporting 100Gb/s network technologies.

Virtual Services

Currently, there are 2 virtual host servers running in a development environment (Wilson Hall 5 West Server Room) and 1 host server running in a production environment (FCC2). All storage is hosted on HP EVA 6100's, which are considered enterprise-class storage arrays. Virtual machines and storage can be moved from one datacenter to another datacenter if configured properly. HP SIM is currently being used to monitor virtual host server hardware. The current virtual infrastructure contains virtual machines, which can reside behind a network firewall or on a DMZ network. This is accomplished by creating separate physical VLANs using interfaces on a network firewall, which in turn connects to multiple switches.

Though there are some virtual desktops running in the environment described above, there should be more consideration to the idea of expanding the use of virtual desktops to areas where sensitive information is accessed on a daily basis or where certain configurations are required.

The potential benefits of running a virtual infrastructure in our current environment are not being fully realized due to lack of resources and higher priority projects. Moving forward with a dedicated group will help the lab use current resources more efficiently, thus minimizing future growth in datacenter space and energy consumption.

Storage Network Services

Formed in January 2009 as part of a re-organization which was partly influenced by the ITIL/IS020K project, the dedicated central storage group provides 4 main customer services:

1. SAN – Direct connection to centrally managed disk for block access.
2. NAS – Indirect connection to the centrally managed disk via an IP network for file access.
3. AFS – A global file system accessible from anywhere in the world.
4. Backups – Data protection of lab servers and non-experiment data.

The centrally managed shared resources are used lab-wide, with customer use spanning both scientific and business needs. Currently over 5000 systems use the SAN and NAS services to access ~1PB of storage. The AFS service – a major component underpinning the central web servers and several neutrino experiments – is accessed by collaborators around the world. Demand for the backup service continues to increase ~30-40% annually.

Assessment of current state of each service:

Overall

- Monitoring must be improved throughout all services to provide a better understanding of the relationships between the various components. Currently, there is a lack of an end-to-end monitoring system to fully see the performance of all components. Current monitoring is reactionary. Most reports are manual.

SAN

- A major upgrade of the SAN switches has placed the fabric (network components only, not storage) in good position. There is ample capacity and the new equipment is well positioned to adapt to upcoming CEE (DCE) technologies.
- There is NOT enough storage capacity on the SAN in terms of performance and space in order to react to changes in production which may adversely affect a shared resource. In 2009, several actions taken to solve performance problems involving FermiGrid, Minos and D0 proved that a higher amount of reserve capacity must be maintained in order to react to (unpredicted) increased workloads from the various experiments. Failure to provide this reserve could result in the shared resource negatively affecting many areas simultaneously.
- Further isolation between experiment areas will need to be achieved to reduce impact. This can be achieved by reducing the sharing of storage arrays between experiments.

NAS

- A major upgrade of the NAS servers in 2009 has significantly increased the performance capacity of the NAS service. Additional servers may be needed in the short term depending on the additional load from the upcoming neutrino experiments or increased activity from existing experiments (CDF, D0, FermiGrid, ILC, Minos, Minerva)
- Due to the highly shared nature of the NAS system and the lack of QoS control available from the vendor, access controls must be addressed at the end-user systems.
- Increase education and involvement in customer processes to better understand and develop solutions that meet customer expectations, or at the minimum, set expectations based on funded solution.

AFS

- Sufficient storage and performance exists to maintain normal growth for the next 2-3 years.
- Automated tools should be developed to reduce the management overhead required when evacuating/re-populating the AFS servers for maintenance tasks.

Backups

- Though the current backup system is adequate today, it is becoming apparent that improvements are required to:
    - Meet the exploding data growth (adapt to 100-200% annual data change)
    - Protect different types of structured (DB) and unstructured file types
    - Provide encryption
    - Provide archiving
    - The operations effort to maintain the existing backup service is high due to the highly user-interactive nature of resolving backup issues. Many issues are the result of a client system error (misconfiguration, upgrade, etc) that require active involvement from the service maintainers.

Vision

By 2014, the Laboratory will require local and wide-area network facilities capable of forwarding at hundreds of gigabits per second, storing and accessing over a hundred petabytes of data, and virtualizing hundreds of servers and workstations. Distributed computing collaborations will place significant demand on the networking and storage infrastructures, requiring the ability of network, storage and virtual server systems to handle very high peak network and data processing loads.

Network infrastructure will be instrumented to interact with applications & middleware, and adaptable to meet their data movement requirements. End-to-end data circuits will provide wide-area network path bandwidth and latency guarantees for distributed applications.

By 2014, the network will be architected in a modular topology to meet the service level requirements of experiments and users of the Lab's network, storage and virtualization services. Based on MOU's and SLA's, various modules will require 99.999% reliability (no more than approximately 5 minutes of service disruption per year), other modules will require different levels of reliability that can be engineered into each module as appropriate.

By 2014, IPv6 will be fully deployed and operational in local and wide area networks.

By 2014, IP telephony (VoIP) will be fully deployed and operational at the Lab.

By 2014, enterprise-class server and switch virtualization services will be fully deployed and operational at the Lab. The use of virtual infrastructures will be just as common as stand-alone servers for most general-purpose computing applications.

By 2014, enterprise storage will consist of a back-end storage infrastructure with a user-facing service of data availability and protection. The exact back-end technology will be based on MOU, SLA or OLA criteria of storage space, access speed, security (encryption), access frequency (for hierarchical storage capabilities), and long-term archival requirements. SLA's will include Disaster Recovery or Business Continuity criteria (Recovery Time Objective and Recovery State Objective).

Stakeholders

The list of stakeholders includes everyone at the Lab who uses networking, shared storage or virtual servers, including:

- Experiments such as CDF, DZero, MINOS, MiniBooNE, MINERvA, and Grid collaborators;
- Off-site experiments such as NOVA, DUSEL, DSM, MINOS, CDMS, and Auger;
- Providers of wide-area and Internet services to the Lab such as ESnet;
- Experiments with specific service requirements for data movement, storage, performance or reliability such as CMS and the LHC Remote Operations Center (ROC).

Strategic Goals

The strategic goals for networking are based on a set of high level, architectural principles. These principles define a core philosophy that helps ensure decisions on the design, implementation, and upgrade of the Laboratory network are made consistent with a common strategic direction. They should provide the basis for all levels of networking decisions within the organization, from the design of major projects to the implementation of small project tasks. The strategic goals are not drafted to be specifically applicable to any one component of this plan.

<u>Shared Goals of Network Services & Wide-Area Networking Groups</u>

- Network designs or configurations will be kept as simple as requirements allow in order to minimize implementation and support effort, improve reliability, and simplify troubleshooting by following established industry Best Practices such as the Cisco Enterprise Campus

Architecture, the Cisco Enterprise Composite Network Model and the Cisco Enterprise Data Center Architecture;
    - Modularity can help make each design element simple and more easily understood; can establish failure domains, security policy domains, availability/performance domains, etc.
    - Modularity can simplify network scalability by scaling within and between modules.
    - Modular designs allow the addition or removal of modules as new services and functions are needed without changing the underlying network design.
- Network infrastructure capacity needs to be significantly increased in order to meet present MOU's and SLA's;
- High capacity, high density switch fabric will be used to the greatest extent practicable, in order to minimize management and maximize performance;
- Network infrastructure will be restored to the forward edge of established network technology, while following industry Best Practices of architecture and design;
- Network reliability needs to be improved through the use of redundancy where appropriate, or by alternative routing and switching paths where appropriate.
- Users and associated computing resources affiliated with a specific organization, experiment, project or location should be aggregated into network modules in order to provide structure to the facility network. A modular topology will allow the implementation of technology and reliability appropriate to meet the SLA's of each module;
- By default, data forwarding path determination will be via the use of standard routing protocols, specifically OSPF and BGP, with use of policy or other types of routing limited to alternate path forwarding needs for specific applications or facilities.
- Network support staff requires training to keep up with current network technologies and methodologies.
- Commercial service offerings and supported open-source tools & software will be preferred over in-house solutions.
- Network monitoring infrastructure will be deployed consistent with a vision of providing network information to users, applications, and middleware.

<u>Network Services Group Specific Goals:</u>

- Wireless network media will be deployed as an integral component of general network access;
- Remote access from off-site locations via Virtual Private Networking (VPN) will be deployed as an integral component of general network access;
- Problematic and obsolete technology will be replaced by current, more flexible and more reliable products, for example:
    - The QIP DNS/DHCP system will be replaced with an appliance-based system that is capable of scalability and DOE-mandated security (DNSSEC);
    - The Cisco ACE and Alteon load balancers will be replaced by an F5 application delivery appliance;
    - The GPS receivers for network time (NTP) will be replaced with current, supported GPS and CDMA receivers;
    - End-of-Life/End-of-Sale (EOL/EOS) equipment as designated by the manufacturer, and obsolete equipment such as hubs.

Wide Area Networking & Network Research Group Specific Goals

- Redundancy designed into the Laboratory's off-site network access will be based on complete physical and geographical diversity, to the maximum extent practical.
- Aggregate off-site network capacity will continuously be maintained at a level to meet peak traffic load requirements for Laboratory facilities, experiments, and applications.
- Performance analysis services for Laboratory wide-area, distributed applications will be automated to the extent practicable.
- Research efforts in network monitoring (perfSONAR) and data circuit services will be integrated into and directly support Laboratory facilities, experiments, & applications.

Virtual Services Goals

- Use virtualization as a driver toward more centralized IT services. Virtualization requires a centralized environment to reach the economies of scale required for maximum cost savings.
- Provide virtualization services that have broad customer appeal by aggressively communicating the many benefits that can be achieved through the use of virtualization technologies.
- Provide a virtual infrastructure which provides rapid provisioning of virtual machines and applications.
- Provide cost-savings analysis data to support increased use of virtualization technologies at Fermilab.

Storage Network Services Goals

- Flexibility -- the service should stay nimble enough to adapt the services to changing workloads while minimizing service outages.
- Maintain enough capacity in terms of storage space and performance in order to support the above objective.
- Keep staff trained in the relevant areas for the specific technologies placed into production.
- Investigate new technologies which can add to or enhance current services.
- Implement monitoring to aid in diagnostics, provide metrics and trending information as well as show the value of the various services to our customers.
- Architectural designs should take into account virtual server and virtual networking technologies. Where appropriate, couple solutions around these technologies to provide flexibility, simplicity, manageability and security.
- Automate repetitive tasks as much as possible; provide documentation and FAQs to reduce operation effort as much as possible.
- Engage users to better understand requirements, use cases and expectations. Educate customers in the capabilities of the services, document known limitations and design considerations to aid customers in the use of the services.

Strategic Objectives:

Strategic objectives are tangible targets for efforts or activity areas that are intended to be the means of achieving strategic goals. They may be specific enough to be applicable to only one area of activity, or may span multiple areas. There are normally timeframes associated with strategic objectives.

Network Services Objectives

- Upgrade core network capacity and reliability to meet or exceed MOU, SLA and OLA requirements:
    - 2010 – Upgrade core switch fabric to 2Tb capacity with 80Gb backplane capacity;
    - 2010 – Start to deploy MPLS in the network core;
    - 2011/2012 – Deploy next generation (40/100Gb) backbone links as needed and as available;
- Upgrade building/experiment distribution switch capacity and reliability to exceed MOU, SLA and OLA requirements:
    - 2010 – Upgrade distribution switch fabric to 2Tb capacity as necessary;
- Upgrade or replace obsolete network devices as necessary to assure network performance and reliability. This includes devices for which software or firmware upgrades are no longer available, which use obsolete data link technologies, or are identified as EOL/EOS by the manufacturer, or which don't support SSH:
    - 2010 – Upgrade all Hybrid CatOS/IOS switches to Native IOS software;
    - 2010 – Replace all CatOS-only switches with IOS equivalents;
    - 2010 – Replace all shared-media devices (e.g. hubs) with dedicated media devices (switches).
    - 2010/2011 – Replace remaining Category 3 cable with Category 5e or 6 as appropriate.
    - 2011 – Improve chassis switch performance by replacing "classic" bus modules with CEF720 or better modules;
    - 2011 – Upgrade access switches to support 1000BT;
    - 2011 – Upgrade/replace approximately 50% of the remaining EOL/EOS devices;
    - 2012 – Upgrade/replace remaining EOL/EOS devices;
- Upgrade DNS/DHCP application
    - 2010 – Implement appliance-based IP address management system, replacing the error-prone QIP system;
    - 2010 – Implement DNSSEC as mandated by the DOE;
- Upgrade Wireless networks
    - 2010 – Upgrade wireless access points in FCC and Wilson Hall
    - 2010 – Install wireless mesh network in Village and selected Fixed Target areas
    - 2011 – Expand wireless coverage to all areas of the Lab as appropriate
    - 2012 – Upgrade WiSM modules to next-generation Wireless LAN Controllers
- Upgrade physical cable plant
    - 2010 – Install new fiber-optic cable between Wilson Hall and LCC (for GCC) to support new communication requirements and to provide path diversity between FCC and GCC.
    - 2010 – Upgrade fiber and copper cable infrastructure in FCC
    - 2010/2011 – Replace multimode fiber between buildings in the Site 38 area with single mode fiber to support higher throughput.
- Implement Network Access Control system
    - 2010 – Deploy Cisco NAC solution to authenticate users, enforce minimum security criteria to connect to the network, and to help identify the types of computers and other devices connected to the network.
- Upgrade network security systems
    - 2010 – Implement network firewall systems in Wilson Hall and FCC distribution switches.
    - 2010 – Implement network firewall systems for Visitor wireless connections.

- Implement IP Telephony (VoIP)
    - 2010 – Upgrade the VoIP pilot infrastructure to support Exchange integration for Unified Communication, and to support scalability to be able to expand VoIP usage across all areas of the Lab.

Wide Area Networking and Network Research Objectives

- Off-site network device infrastructure and operations
    - Facilitate, support, and upgrade as necessary, a fully redundant perimeter network infrastructure that provides the Laboratory with the high bandwidth data channels necessary for its offsite data movement requirements. Timelines:
        - 2010/2012 - Increase off-site 10GE data channels to meet data movement requirements
        - 2010/2011 - Deploy next generation end-to-end circuit router & border routers
        - 2011/2014 - Implement 100Gb/s off-site links on perimeter infrastructure
    - Deploy capacious end-to-end data paths to remote sites involved in high impact data movement with the Laboratory, including support infrastructure to facilitate use, monitoring, and troubleshooting of those paths. Timelines:
        - 2010/2014 - Facilitate optimal use of network paths between CMS T0/T1/T2 sites, including increasing bandwidth capacity to meet evolving requirements, and deployment of dynamic circuit capabilities
- Network support & analysis services for distributed, high-impact scientific computing:
    - Deploy higher density, higher capacity network infrastructure in support of high impact scientific computing facilities. At the current time, this is limited in scope to the CMS Tier-1 facility. Timelines:
        - 2010 - Complete upgrade of CMS Tier-1 facility LAN to be redundant core 10GE-based data center switches, with very high capacity switch fabric
        - 2010/2011 - Migrate Tier-1 servers to direct connections on CMS data center switches; support initial implementation of 10GE-connected CMS host systems
        - 2010/2012 - Upgrade interconnections to 100Gb/s or more between core switches, and to 80Gb/s for access switches
        - 2012/2014 - Expand data center switches to high density chassis, higher capacity switch fabric; replace access switches with data center-class switches
- ESnet Chicago Metropolitan Area Network (MAN) design, engineering, and operations
    - Operate the Chicago MAN, which supports all of the Laboratory's off-site access, at a very high level of reliability, increase MAN channel deployment to meet bandwidth needs, introduce next generation optical network technology as it emerges, and develop a MAN test-bed capability. Timelines:
        - 2010 - Establish 2nd Ciena DWDM hub on site, geographically separate from FCC
        - 2010 - Implement a 10GE channel to ANL to serve as an initial MAN testbed
        - 2010 – Upgrade the west leg (FNAL-ANL-Level3 PoP) of the MAN to 20GBS SDN capacity, for closer parity with the east leg
        - 2011/2014 - Upgrade the MAN to 100Gb/s capability & deploy 100GE channels
        - 2013/2014 - Replace or renew the ComEd fiber component of the MAN, as appropriate
- Network research & development efforts

- Develop an R&D support infrastructure to participate in international & national, advanced, optical wide area network research initiatives. Timelines:
    - 2010 - Deploy a full service IPv6 test-bed network facility
    - 2010 – Develop content management system for performance analysis
    - 2010/2012 - Develop & deploy End Site Control Plane systems
    - 2010/2012 - Develop Network Weather & Performance Service E-Center
    - 2011/2012 – Automate performance analysis methodology
    - 2010/2014 - Support Laboratory participation in additional advanced network and distributed systems R&D projects and collaborations

<u>Virtual Services Objectives</u>

- Develop a long-range plan to deal with future architecture, budget, scope, staffing, and technology selection concerns:
    - 2010 – Meet with CD management and stakeholders to gather ideas and expectations, which will help set the priorities and direction of the virtual services group.
    - 2010/2011 – Develop documentation to describe the current and future virtual architecture, including networking, security, failover, storage, dependencies, and other key elements to aid in long-range planning.
- Upgrade current MIS virtual infrastructure to support additional production virtual machines:
    - 2010 – Build out MIS virtual infrastructure to by adding 6 additional CPU's, memory, and licenses across 2 VMware ESX servers.
    - 2010 – Expand SAN storage capacity with new drives to support additional virtual machines.
    - 2010 – Install a second VMware vCenter Server to provide redundancy of management and monitoring tools.
    - 2010/2011 – Upgrade MIS virtual infrastructure to vSphere 4 Enterprise Plus version.
- Implement consolidated backups in virtual environment:
    - 2010/2011 – Investigate, design, and implement a backup strategy for virtual machines utilizing VMware's consolidated backup feature.
- Implement additional capabilities in analysis, monitoring, and reporting of metrics in a virtual environment:
    - 2010 – Investigate and Implement 3rd-party or built-in tools used to monitor production environments.
- Demonstrate viability of new application delivery mechanisms for production use:
    - 2010/2011 – Investigate/Install/Configure VMware View ThinApp , Microsoft App-V, or other application virtualization products.
    - 2012 – Investigate and install load-balanced Terminal Servers, TS Remote App, or other applications used for running applications using terminal services.
- Demonstrate viability of thin client/virtual lab technologies for production deployment:
    - 2011 – Install and configure thin clients for use as kiosks, training machines, and contractor machines.
    - 2012 – Incorporate technologies that allow fast provisioning and reimaging of virtual desktops.
- Implement a tiered services architecture utilizing various virtualization technologies working closely with storage and networking groups:
    - 2011 – Compare virtualization offerings by multiple providers to identify the most cost-

effective products for each type of virtualization technology in use.
    - 2012 – Identify each tier (i.e. desktops running on iSCSI storage, servers running on Fiber Channel storage, highly available servers running on virtualized storage) and select the software, hardware, protocols, and other components which will be used to implement each one.
    - 2012 – Develop processes and tools to aid customers in selecting a tier that meets their performance needs and budget constraints.

<u>Storage Network Services Objectives</u>

- Service Consolidation
    - SAN
        - 2010 Merge MIS/CD/(TD) SAN Networks – includes incorporating existing SAN fabrics in WH/CD via FC and possible iSCSI SAN at TD
        - 2011 Replace aging equipment in MIS/TD with standardized SAN equipment (HDS)
    - Backups
        - 2010 Merge MIS/CD/ (TD) Backups – investigate existing requirements and match to existing infrastructure. Where required, upgrade existing infrastructure to meet/exceed requirements from MIS (and TD).
        - 2010 Implement hardware based encryption services to protect PII data
        - 2010 Implement offsite storage practices w/ 10yr retention. Includes the planning of keeping proper equipment to actually recover data that is 10yrs old and developing a plan to test periodic restoration and verification of data.
- Service Upgrades
    - AFS
        - 2010 Upgrade AFS servers to v1.6 of software using Solaris ZFS filesystem.
        - 2011 Replace AFS servers with Linux servers.
        - 2012 Upgrade AFS service with new features from YFS SBIR project.
        - 2013 Integration of mobile devices (Windows Mobile, Palm, etc) with AFS/YFS file system for data access.
    - Backups
        - 2010 Implementation of disk based library for TiBS, further increasing per server backup capacity.
        - 2010 Expansion of Commvault backup system to fully support Exchange and SharePoint services for both backups and archiving.
        - 2010 Increase LTO-4 capacity as appropriate for service demand.
        - 2010 Continue investigation into Single-instance-store (data de-duplication) and Virtual tape libraries for feasibility and price point that is appropriate for the applications at the lab.
    - NAS
        - 2010 Investigate/implement the use of BlueArc Mercury platforms. These new platforms are a low-cost alternative to existing Titan infrastructure that also underpins new NFSv4.1 server implementation. This new implementation allows for multiple servers to serve data in parallel providing better scaling for the shared NAS service.
        - 2011 Implementation of NFSv4.1 pNFS with integration into Kerberos infrastructure.

        - 2011 Investigation into Tier 0 storage (SSD) to be used in conjunction with existing storage devices to provide better performance and scalability.
        - 2012 Investigation of incorporating an HSM system into the NAS service
        - 2013 Implementation of HSM system into NAS service
    - SAN Fabric
        - 2010 Complete upgrade of SAN Fabric to 4Gb.
        - 2011 Implement CEE (DCE) and upgrade to 8Gb fibre w/ 10Gb IP uplink.
        - 2012 Implement full CEE with existing 40Gb/100Gb IP infrastructure
        - 2013 Implement DR/BC
    - SAN Storage
        - 2010 Replace aging HDS storage with 2nd generation storage
        - 2010 Replace aging 3PAR storage (reaching EOS) with 2nd generation storage
        - 2010 Expand existing storage capacity and performance to fully support Exchange, SharePoint, Virtualization Service, migration of DB servers to SAN and expansion of AFS, Backup and NAS services.
        - 2010 Research/implement DDN S2A6620 storage devices. If proven to work, these devices will provide a lower cost storage alternative for experimenters that will work with the NAS service at a price point that is sustainable and ensures experiment isolation.
        - 2010 Investigate the use of hardware virtual controllers (USP-V) or s/w based virtual controllers (such as Falconstor) to add flexibility and reduce downtimes required for storage reconfigurations/migrations.
        - 2011 Implementation of Virtual storage controllers
        - 2011 Research/Investigate costs and requirements for disaster recovery and business continuity (DR/BC).
        - 2012-2013 Implement DR/BC for storage.
- Service Operations
    - 2010 Establishment of user meetings to promote services, better understand user needs and issues and raise awareness of the limits of the various services and use cases.
    - 2010 Implement SAN monitoring that is heterogeneous, aids in diagnostics, provides trending as well as modeling capabilities.
    - 2010 Provide adequate training to personnel in new technologies/equipment used in production.
    - 2010 Increase documentation using established document repository (docdb or SharePoint) for user training, document processes, etc.
    - 2011 Establish processes and procedures that will comply with the upcoming ITIL/ISO20K requirements for Change Management and Continuous Improvement Plans

Strategies

Network Services

The strategy for network architecture and networking support is the Fermilab Modular Architecture, a site-specific implementation of the Cisco Enterprise Campus Architecture and the Cisco Enterprise Composite Network Model in order to improve scalability, flexibility and reliability.

- Modularity can help make each design element simple and more easily understood and can establish failure domains, security policy domains, availability/performance domains, etc.
- Modularity can simplify network scalability by scaling within and between modules.

- Modular designs allow the addition or removal of modules as new services and functions are needed without changing the underlying network design.

Fermilab computer rooms are being re-engineered as a site-specific implementation of the Cisco Enterprise Data Center Architecture.

While our model is based on Cisco's best practices architecture, there is no requirement to use Cisco products so we are not "locked in" to any vendor.

Wide Area Networking and Network Research

High-level wide-area networking strategies employed to achieve our strategic objectives in a manner consistent with our strategic goals include:

- Maintaining and upgrading wide area network-supported infrastructure to keep it as capacious as practical. Capacious network infrastructure helps avoid application-level performance problems, and provides the necessary agility to accommodate changing needs.
- Aligning wide-area networking infrastructure and support with DOE science research efforts, as well as technology trends in the general research & education community. Technology decisions and deployments need to be made in context of the global collaborations the Laboratory supports.
- Integrating network research efforts with the Laboratory's advanced scientific computing activities. Applied network research directly applicable to Laboratory distributing computing efforts benefits both the results of the research, and the Laboratory's overall scientific mission.

Virtual Services

- Use automation to achieve goals in providing provisioning (i.e. Windows PowerShell or visual basic scripts that interact with VMware API's).
- Implement application virtualization technologies (i.e. VMware ThinApp or Microsoft App-V).
- Expand current virtual infrastructure by adding additional storage and processing.
- Investigate 3rd-party applications which enhance the advanced features of our current platform.
- Actively encourage the use of virtualization services, which requires having an environment that is stable, flexible, secure, and easy to use.
- Engage in long-range planning exercises to help set the future direction of virtualization technologies at Fermilab.

Storage Network Services

- Implement high-availability/redundancy in all areas. Service outages should eventually become an anomaly and not a regularly scheduled maintenance window.
- Implement standards of operation – document processes and procedures for daily operations and changes. Use this documentation to review processes to aid continuous improvement plans (CIP).
- Engage customers early and often to understand needs, issues and better anticipate future requirements.
- Standardize monitoring. Along with alerts, the monitoring should include end-to-end quality of service (QoS) information, trending information and, if possible, modeling information. The latter two can be used for Capacity Management. The monitoring tools should be able to

monitor heterogeneous systems. From the tools a set of requirements can be placed on the types of equipment we will choose to purchase.

- Continuous pursuit of new technologies and automation that can reduce operational costs, increase efficiencies and/or change the "way we do things" in a positive manner.
- Maintain a skilled and cross-trained staff to ensure that data protection is not just about the different technologies put into place, but also the personnel required to operate and maintain the systems. The loss of a single person should not place the service in jeopardy.
- Delivered solutions should be modular, flexible, simple, and secure whenever possible. Manageability and maintainability of the solution should be stressed at all levels. Designs should always be cognizant of virtual server and network technologies and deployments. When possible, designs should encompass all three aspects providing s single packaged solution.

Resource Needs

Network Services

Historically, the level of effort, in terms of both personnel and M&S costs, has remained relatively constant for network support over the years, however significant investment will be required in the next several years:

- To replace the aging fiber infrastructure in several areas of the Lab;
- To replace obsolete equipment presently in service;
- To increase capacity and reliability to meet the increasing needs of experiments and general computing.

Wide Area Networking and Network Research

Wide-area networking research & development efforts will be increasing as we initiate two new projects funded by the Department of Energy. Part of the grant funding will be used to bring in additional software development resources for those projects. The rest of the effort for those projects will come out of the existing staff. In most cases, the new R&D effort aligns with current staff job efforts & responsibilities. But some redirection of effort away from current ongoing wide-area network support will be unavoidable. We will require some increases in efficiency and productivity to sustain the Laboratory's wide-area network service at current performance levels. New management tools will need to be procured if that expanding scope of work is to be met within current staffing levels.

Additional effort and skill sets will likely be needed to provide guidance in optimizing use of emerging distributed applications across our increasingly capacious network infrastructure. We anticipate this to become an increasingly important aspect of Laboratory network support for our collaborations that are involved with wide-area distributed computing systems.

Over the longer term, the Laboratory's networking efforts need to be consistent with the DOE Office of Science strategic directions to develop Petascale science facilities supported by terabit/sec network infrastructure. Additional resources, both personnel and capital expenditure, may be needed to keep the Laboratory at the forward edge of the Office of Science's strategic direction.

Virtual Services

In order to build up our virtual infrastructure in the near future, resources from other groups will likely be required to assist in migrating physical machines to virtual machines. Cooperation with networking, storage, operating system, and application groups will be essential in achieving a successful virtual environment.

Depending on the goals set by management, additional staff or contractors may be necessary to achieve these goals.

Storage

With respect to the scientific needs, often times the exact requirements are not well known due to the very nature of their business – one of a kind, never been developed. As such workloads can vary significantly from initial prototyping/modeling to actual production and may not have been accurately predicted when initially providing a solution. Sufficient additional capacity in terms of both storage space and storage performance must be maintained in order to allow the services to adapt to changing workloads and patterns of the various customers.

As demand for the various services increases, limited staff resources will negatively impact the following areas in order:

1. Effort to investigate new technologies
2. Effort to investigate improving current services
3. Effort to maintain standard operations

Progress Indicators

Network Services

The level of progress in attaining strategic objectives for Network Services will be determined through a combination of four factors:

1. Meeting externally-imposed milestones, such as DOE-mandated policy compliance deadlines;
2. Measurement and publication of performance and reliability data to demonstrate that MOU, OLA and SLA requirements have been met;
3. Service desk ticket statistics showing failure recurrence, time to close, and other statistics to show that the network is becoming more reliable and that the Network Services support staff is responding and resolving issues as required;
4. Feedback from stakeholders.

Wide Area Networking and Network Research

The level of progress in attaining strategic objectives for this plan will be determined through a combination of three factors:

Comparison between the timeline expectations for strategic objectives listed in this plan, and what is actually achieved within those time frames. This comparison is not intended to be absolute. It is expected that there will be some time shifting in implementation of identified objectives, given the dependencies on technological evolution, personnel resources, and changing requirements. Rather, the progress is better gauged by how closely implementation compares to the general trend outlined for the objective.

Measurement and observation of how on the capacity and capabilities of the network infrastructure and services compare to the utilization and performance at any particular time. Insufficient capacity or capabilities to meet current requirements is a potential indicator that progress needs to be greater.

Feedback from stakeholders. In the end, network infrastructure and support services exist to satisfy the needs of the stakeholders, and they should be the ones to determine how well their needs are being met.

Virtual Services

The level of progress in attaining strategic objectives for Virtual Services will be determined through a combination of four factors:

1. The number of active virtual servers in production use;
2. Measurement and publication of performance and reliability data to demonstrate that MOU, OLA and SLA requirements have been met;
3. Service desk ticket statistics showing failure recurrence, time to close, and other statistics to demonstrate the reliability of the virtual infrastructure and that the Virtual Services support staff is responding and resolving issues as required;
4. Feedback from stakeholders.

Storage Network Services

The level of progress in attaining strategic objectives for Storage Network Services will be determined through a combination of four factors:

1. The amount of networked storage in use;
2. Measurement and publication of performance and reliability data to demonstrate that MOU, OLA and SLA requirements have been met;
3. Service desk ticket statistics showing failure recurrence, time to close, and other statistics to demonstrate the reliability of the storage infrastructure and that the Storage Network Services support staff is responding and resolving issues as required;
4. Feedback from stakeholders.

FY10 Strategic Plan for Central Services (2010-2012)

Mark Kaletka, Jack Schmidt
10/30/2009

Mission

Provide common central computing services which broadly underpin the Laboratory's Open Science and general technical and business missions.

Context and Assessment of Current State

The strategic plan for central services is guided by the Computing Division Strategic Plan and the strategic directions of the Laboratory program, and MOU's and SLA's established with other organizations, including internal Laboratory organizations (other Div's/Sec's). The strategic plan is also guided by cybersecurity guidance and directives imposed on the Laboratory by DOE.

The Computing Division provides a comprehensive set of central services which support the daily scientific and business functions of the lab (web, email, print, storage, database and application development, backup, etc.), the computing support infrastructure (patching, inventory, configuration mgmt, antivirus, authentication, metrics and accounting, etc.), and the computer security infrastructure (inventory, scanning, automated controls, etc.). Several of the services have grown to be very complex, with limited depth of knowledge in the support groups (one or two "experts"), leading to concerns for future reliability.

In response to the findings of the 2009 DOE Safeguards and Security audit of cybersecurity, the Laboratory has launched a "Tune IT Up" campaign to improve the management of IT resources across the Laboratory. Although the campaign has specific deliverables required to complete the corrective action plans for the S&S audit, the higher level goal is to improve the management of IT at the Laboratory in a sustainable way, by making changes to IT management policies and procedures to ensure a more consistent application of best practices. Part of this campaign includes the consolidation of IT support from existing silos of the Laboratory into the Computing Division. IT support from PPD and FIN are being consolidated, with TD to follow during FY2010.

The requirements of Open Science are understood and incorporated into the existing environment and work model in areas like planning, risk assessment, and mitigation. This work is done in coordination with Open Science Grid.

Central services continue to follow the strategy and vision of IT Service Management in the ITIL framework. The Service Desk is the focus for ITSM. Central services are approximately half way through the roadmap for adoption of ITIL and ISO20000 certification, with maturing processes for incident and problem management, and service level management, and developing processes for change and configuration management.

Vision

Be the provider of innovative, high-quality, secure common computing services. By providing these common services, contribute to the Laboratory's successful execution of the current scientific program, while positioning the Laboratory to successfully compete for future scientific projects.

Stakeholders

Sponsors: Computing Division management; Laboratory management; DOE management;

Customers: All users of Fermilab computing resources, world-wide, regardless of affiliation;

Providers:

Goals and Objectives

Overarching Goals

- Achieve "operational excellence" by following best practices for service delivery, quality and change control, customer service and satisfaction, etc., through implementation of the ITIL framework leading to ISO20000 certification by 2011. This is a transformational cultural change which will have significant impacts on all aspects of delivery of central services.
- Fully centralized management of IT infrastructure and services at the lab, encompassing all of the currently disjoint and private IT infrastructure of other areas of the lab (a Computing Division strategic goal).
- Stable and secure operating environment which is flexible and responsiveness to user needs.
- Services "marketed" and used throughout the Laboratory (scientific & business applications).
- Plans developed one year ahead and resources in place at least 6 months ahead of need.

Central Services

- Provide robust, stable and secure production central IT services, including email, web, printing, global file services, etc.;
    - Increase the level of support and the level of security for all centrally-provided services;
- Provide robust, stable and secure infrastructure to support the major operating systems, e.g. sw distribution, patching, inventory & config control, licensing, etc.
- Provide innovative new services which anticipate customer demand;

Customer Services

- Increase customer support staff and end-user effectiveness through training and documentation.
- Improve end-user experience through more intuitive and user-friendly interface for customer support.
- Better automation and integration of account processing.
- Provide support for major desktop operating systems – Windows, Linux, MacOS.
- Collaborate and integrate efforts across the whole Laboratory, including cross-pollinating best practices. Successfully market services to other Div/Sec.

Strategies

- "Sell" management support of the high-level goal to centralize management of IT infrastructure and services at the Laboratory. Achieving this goal will require a set of business management strategies which demonstrate the cost-effectiveness and other benefits to the Laboratory.

- Apply formal ITIL-derived best practices to the delivery and support of services. The ITIL framework applies across the whole range of activities and implements a number of key processes, among them continuous service improvement.
- Demonstrate the effectiveness of these first two strategies by achieving ISO20000 certification by 2011.
- Vigorously encourage use of central services; This means central services have to be useful, easy to use, flexible, responsive, performant, etc., and aggressively marketed to the end-users.
- Continue to encourage the involvement of the user community in the decision process, through e.g. Windows Policy, Unix Users, GCSC meeting, sysadmin meeting, web users group, database users group, network policy group, etc.
- Increase efficiency and integration across the Laboratory, collaborating better with other Div/Sec/Exp's on common projects, and by aggressively looking for opportunities to reduce redundancies and consolidate efforts.
- Reduce effort and increase efficiency through effective collaboration with other labs and institutions (particularly taking advantage of FRA relationships with ANL and UC).
- Outsource or use consultants where appropriate, to fill gaps in effort or expertise. This includes considering outsourcing major efforts where cost/benefit analysis supports the decision.
- Think long-range for enterprise solutions – not "labware", use supported (commercial or open source) solutions, pay attention to TCO in build-vs-buy decisions.
- Use common methodologies, tools and frameworks for application development to achieve consistency and efficiency. Applications should share common support data and methods (not duplicate them).
- Use automation wherever practical and cost-effective.
- Investigate and adopt rigorous development and test methodologies which provide rapid turnaround for projects without sacrificing production quality.
- Maintain life cycle processes which anticipate user needs, changes in technology, growth (or decline) in demand, etc. and allow for tactical plans to be developed twelve to eighteen months ahead, with implementation six months ahead of need.

Resource Needs

Staffing for central services will necessarily remain approximately static. Expertise in several critical areas – web services, email, shared storage, backup, for example – is very thin with perhaps one real "expert" with little backup and limited prospects for developing other staff as backup. In addition there is the need to free enough effort to continue to evolve current services and investigate new ones.

The transformation of our IT processes into the ITIL framework will require additional resources (effort) in the short term, before beginning to produce larger long term gains as more efficient processes are established and IT services are consolidated across the Laboratory. Consolidation will make some of these resources available in the long term, but we will have to rely on use of consultants for much of this effort.

Limited staff resources will increase the need for effective collaboration and aggressive consolidation of efforts across the Laboratory, as well as for use of contractors or consultants and outsourcing. The alternative is to reduce levels of service and slow (or halt) the introduction of new services, which will have an adverse effect on the Laboratory's ability to support the current and future program.

Progress Indicators

Indicators of effective progress include:

- Progress towards service excellence indicators derived directly from the ITIL framework;
- Use of central services by a substantial portion of the Laboratory's scientific and business program, as measured by the number of organizations and users supported, rather than equivalent local solutions;
- Improvement of the levels of central services, as measured through SLA and MOU agreements and actual service delivery, and through internal, peer and external reviews, etc;
- Maintaining high "scores" for internal, peer and external reviews of computer security process and documentation, combined with acceptable low actual rates of vulnerabilities, incidents and similar technical indicators.

Strategic Plan for Scientific Computing Facilities (2010-2014)

Jason Allen, Jon Bakken, Eileen Berman, Keith Chadwick, Glenn Cooper, Matt Crawford, Phil DeMar, Catalin Dumitrescu, David Fagan, Stu Fuess, Gabriele Garzoglio, Burt Holzman, Alexander Moibenko, Stanley Naymola, Gene Oleynik, James Simone, Amitoj Singh, Stephen Wolbers (editor)
10/28/2009

Mission

The Scientific Computing Facilities at Fermilab (SCF) provide scalable management, storage and movement of data, and scalable computational resources required for the scientific program.

Context and Assessment of Current State

In 2009 Fermilab's Scientific Computing Facilities are well aligned with the scientific program's needs. Indeed, they have often helped define the model for world-class scientific computing, as there is a long history of technical leadership with respect to the design, provisioning and management of scientific facilities, scalable high-performance computing, large-scale data management, and data movement over wide area networking. The facilities are involved in implementation and alignment with the Open Science Grid. The facilities are especially important for the computing needs of the Run 2 experiments CDF and D0, CMS, MINOS and MiniBooNE, particle astrophysics experiments, other small experiments including test beams. The facilities design and operates compute clusters with high-performance network interconnects and scalable disk storage systems for high-performance parallel computing as part of the LQCD.OMB300 project to provide a national computing infrastructure for lattice QCD. Other initiatives and programs - especially those directed at the laboratory's future, require excellent scientific computing facilities. These include future experiments at the Intensity and Cosmic Frontiers and, specifically, computational programs in: beam and RF structure modeling of accelerators and cosmological simulations.

The scientific facilities are continuously improving in areas such as availability, compute systems monitoring and management, performance benchmarking, operating system installation procedures, operating system patching, computing performance, data caching, data storage, data movement and network systems monitoring.

Elements of strategy that have led to the current state include:

- Appropriate and close relationships with scientific organizations (experiments, projects, divisions/centers, outside partners).
- Awareness of the technical evolution of facility components.
- Willingness to determine a good balance between self-integration and procured systems.
- Use of common tools and processes whenever possible, including alignment within the organization and with lab and industry directions.
- Reduction of effort needed for support by careful planning and the use of appropriate metrics.
- Proper attention to aspects of the proper operation of large facilities, for example, backups and security.
- A willingness to outsource work within and outside Fermilab in order to achieve efficiencies.

Organizationally, Fermilab scientific computing facilities are managed by five departments in the SCF quadrant: CMS Computing Facility, High Performance and Parallel Computing Facilities, Fermilab Experiments Facilities, Grid Facilities, and Data Movement and Storage. All departments use similar operational concepts and they differ only in the specific services provided and stakeholders served.

The CMS Computing Facilities are responsible for the US CMS Tier 1 (T1) computing facility. This includes the acquisition, installation, and operation of large computing, storage and data movement facilities that form part of the CMS worldwide computing system. Important connections and coordination with CMS, the Open Science Grid (OSG), and other Grid and computing organizations are required for effective management and operation of these systems. A second facility, serving the LHC Physics Center (LPC) is also acquired and managed by the CMS Computing Facilities department. The computing systems extend beyond Fermilab and demand excellent connectivity to CERN and to every CMS T1 and T2 center worldwide.

High Performance and Parallel Computing Facilities (HPC) supply a high-performance computing environment for the Lattice QCD computing community as part of a national hardware infrastructure program. The LQCD facility includes several cluster-based parallel supercomputers, scalable disk-based storage systems (dCache and Lustre), a login service and a four node prototype production cluster having GPGPU (General Purpose Graphical Processing Unit) co-processors. In addition, the HPC department has available a GPGPU system to support investigations of scientific software optimizations on many-core (cores > 8) processors. HPC supports other growing high-performance computing efforts, in particular the Cosmological Computing Initiative, accelerator beam modeling and superconducting RF cavity design simulations. Software development common to all HPC computing is an important component of this activity.

The Fermilab Experiments Facilities provide well architected ensembles of networked computers, possessing well-administered operating systems, management of and connectivity to storage systems, basic infrastructure middleware including batch systems, and support for experiment infrastructure. The Run 2 experiments in particular benefit from the excellent system administration, maintenance, and migration from old to new hardware of these systems. Other user communities, including a growing intensity frontier program, are also beneficiaries of this service.

The Grid Facilities provide software development support for Grid activities including development related to authorization, storage, resource allocation, security, metrics analysis and accounting. There is a transition to grid operations as current development completes and the grid moves into full production for the LHC and others. The Fermilab grid facility (FermiGrid) enables the experiments at Fermilab to have priority access to their own computing resources and supports sharing of these resources in an opportunistic fashion across the various experiments and projects that participate in FermiGrid. In addition, FermiGrid resources are available to members of the Open Science Grid on an opportunistic basis. Interoperability with other external grid infrastructures (e.g., EGEE, TeraGrid) is increasing.

Data Movement and Storage provides large storage systems, both cache (disk) and permanent (robotics and tape). Data on older media has been and continues to be migrated to higher density media allowing for better management of the data. These systems underlie all of the scientific computing and are required for local, campus-wide, national and international data movement. Development of dCache as part of a worldwide collaboration continues. Wide area networking and network research provides leading-edge high performance data movement for the scientific programs' needs. Responsibility for long-term data storage and access is provided. Interconnectivity and coordination with other facilities is especially important. R&D efforts including work on dynamic paths, monitoring, and high-throughput transfers are all part of this activity.

All facilities, where appropriate and not in conflict with mission needs, do or will provide opportunistic, reciprocal access of their facilities using grid methods, via the Fermilab campus integrated grid facility, FermiGrid.

Vision

Fermilab provides and will continue to provide a center of excellence in computational and data management facilities to the Fermilab scientific program. This includes but is not limited to:

1. To provide steady and stable Run II computing facilities.
2. To provide a fully operational CMS Tier-1 computing facility, including both the CMS production center committed to international CMS needs and the LPC center committed to USCMS analysis needs.
3. To have a leading role in the World Lattice QCD Facilities
4. To provide HPC facilities where appropriate to the Fermilab scientific community.
5. To be an important contributor to HEP, Astroparticle and other general computing needs.
6. To be a full partner in national and international scientific computing efforts.
7. To provide computational facilities for any Fermilab scientific activity, now and in the lab's future.
8. To design and build general scientific computing facilities needed for the Intensity Frontier experiments.
9. To have a leading role in the Open Science Grid and provide a premier grid facility on the OSG.
10. To be aligned with other scientific and industry facility practices and to lead in defining those practices whenever appropriate.
11. To provide performant, leading edge, and reliable WAN capability.
12. To participate in research in grid, cloud, and WAN computing.

Stakeholders

The primary stakeholders are the Fermilab scientific community (including but not limited to Fermilab accelerator-based experiments, CMS, Lattice QCD Computing, accelerator simulations, and particle astrophysics). Other stakeholders include the Computing Division, all other Divisions, Centers and Sections at Fermilab, other laboratories and universities involved in the Fermilab scientific program, the Department of Energy, the National Science Foundation, the consortia that we belong to (such as the Open Science Grid and the National Lattice QCD project), and other partners in the Fermilab scientific program.

Goals and Objectives

1. Provide reliable and supportable computing facilities including large computational and storage capabilities and wide area data movement.

2. Provide a leading world-class center for scientific data. Provide facilities that are responsive to experiments' and projects' needs by working with their scientists and engineers and computer scientists and taking computational models into account.
3. Provide and expand US QCD, cosmological computing and other general HPC computing.
4. Engage in capacity planning, including the potential for provisioning capacity on contingency, taking into account the costs (hardware and effort) for a sustainable facility.
5. Architect, design, acquire, install, commission and manage general scientific computing facilities for intensity frontier and other scientific users.
6. Develop, deploy and operate the grid service model.
7. Continue development of a usable security model.
8. Further develop the Computing Division's internal service model.
9. Integrate with experiment, national, international and HEP cyber-infrastructures.
10. Be aware of opportunities coming from OHEP funding agencies and peers.
11. Acquire and operate systems consistent with the laboratory energy efficiency and Environmental Management Plan.

Strategies

1. Interact and coordinate with experiments' and projects' computing experts and computing models.
2. Use metrics and costs to clearly state business methods that include the ability to trade off cost, schedule, and performance.
3. Operate the facilities at very high efficiency and reliability, reduce the operating costs whenever possible, and use commonly developed or acquired tools whenever possible.
4. Develop, improve and maintain metrics and use them for feedback to software and service providers.
5. Maintain awareness of the linkage between the computing and the physical facilities. Contribute to energy efficiencies by keeping resources utilized, keeping systems optimally coupled to the Facilities HVAC, and linked as appropriate to central power related metrics.
6. Engage in research appropriate to achieve goals and objectives.
7. Represent Fermilab facilities at appropriate technical and other conferences and meetings.
8. Work jointly with the FNAL Computer Security Team.
9. Migrate tape storage to higher density newer technologies as quickly as possible.

Resource Needs

A well trained and organized staff is critical for the successful operation of any and all facilities. The Fermilab scientific computing facilities have such a staff and should strive to keep it trained and up to date in modern programming and support tools and methodology. Sufficient staff will have to be provided in key areas, matching the strengths of the individuals with the needs of the operation. Skills to document costs, provide metrics and state business and policy models will continue to be needed.

The facilities themselves will grow and evolve as the scientific program changes. Older equipment will have to be retired as necessary. The facilities will have to be sited carefully as the computer room space, power and cooling availability will continue to be a critical resource at the laboratory.

Progress Indicators

1. Increased understanding of facilities utilization and a linkage of facility use to scientific utility, including latencies in the procurement process.
2. Increased commonality of approach in facility underpinnings.
3. Movement to service-based interface to experiments.
4. Increased robustness of service offered, for a constant level of support effort.
5. Success in meeting CMS metrics and performance levels.
6. Success in meeting milestones for US LQCD project.
7. Progress goals for all externally funded programs.

Strategic Plan for Computational Physics Tools and Applications (2010)

Panagiotis Spentzouris, 8/27/2009

Mission

The mission of the Computational Physics Tools and Applications (CPTA) area of the Computing Division (CD) is to develop and apply scientific software tools for accelerator and detector design, physics analysis and theoretical studies, to help users deploy and utilize these tools, and to provide expertise in computational physics at the laboratory.

Context and Assessment of Current State

The CD CPTA effort is recognized as a valuable resource for supporting HEP scientific computing at the laboratory. We develop and support computational physics tools and their applications for the major activities in the strategic plan of CD, such as Run II and LHC analysis support, accelerator modeling, future experiment and detector design simulation support and development, and computational physics R&D in areas useful to the above activities.

Most scientific tools and libraries used by the Run II experiments are in and will remain in maintenance mode. Some maintenance will be needed to accommodate changes to operating systems. Development of ROOT continues driven primarily by the needs of the LHC experiments, as is the development of other HEP and mathematics libraries. Planning for this development is done in collaboration with the CMS offline software and LCG applications area groups. Legacy Fortran-based tools such as GEANT3, CERNlib and most standard physics generators are mature products.

Event generators and tools that work with them are critical for the analysis of Run II and LHC data. We develop and support the necessary tools and collaborate with the experiments to help maximize their utilization, in order to enhance physics productivity and further advance the understanding of the data.

We are members of the GEANT4 collaboration and participate to the development of the GEANT4 toolkit. This toolkit is a key software package for CMS and other scientific programs at the laboratory, and we expect that it will be used in the design of Project X and other high-intensity frontier experiments, and future lepton collider experiments.

Our intense involvement in software for the SiD ILC detector has been de-scoped in light of the change of priorities of the US High Energy Physics efforts. The SiD effort is redirected to research and support for current or near term future detectors, and generic software tools. We do maintain a lower level of support for the design of a generic future lepton collider calorimeter, with emphasis on precision jet calorimetry.

We support tools used to simplify software development at a low level and develop and support tools for high level applications, such as offline frameworks and DAQ systems. Some commercial tools are made available to the users, but the support is now minimal. We also provide consulting on software design, both for ongoing experiments and potential new efforts. The level of support on software design is determined by the overall CD priorities, with highest priority assigned to RUN II and CMS, followed by other approved Fermilab projects (NoVa, Mu2e).

We provide simulation support and expertise in the major accelerator activities of the laboratory's mission: RUNII operations (Tevatron, Main Injector, Booster), future accelerator R&D (Project-X, future lepton collider), and computational accelerator physics tools development (SciDAC project). We collaborate with Fermilab AD and TD personnel on accelerator physics applications, and are members of the Accelerator Physics Center (APC), where we participate in accelerator physics R&D projects.

We are leading the SciDAC accelerator modeling project, ComPASS. In the context of ComPASS, we collaborate with math and computer science teams in R&D for high performance solvers, framework technology, and other computational infrastructure. Also, through SciDAC we have access to computational and accelerator physics resources beyond those available at Fermilab. We leverage those resources to help our accelerator physics R&D and application activities. We have access to the NERSC and ALCF supercomputer facilities, and to a parallel PC cluster that we operate using lQCD surplus hardware. We port and deploy our accelerator codes on these facilities.

We make available tools required for modern global scientific collaboration such as video conferencing, electronic logbook, document database, and agenda tools; we also provide consulting on the use of these tools to users.

Vision

We envision the CPTA activities to continue enabling the Fermilab CD to remain a center of excellence in scientific tools development, and to help make Fermilab a welcoming and effective place to work, through the development and support of applications for present and proposed Fermilab experiments and accelerators. For FY09-10, Run II and LHC will continue to be high priority activities, with mid-term future accelerator R&D (Project-X) ramping up, and significant effort going to the design of future neutrino experiments, astrophysics experiments, possible rare decay mode experiments, future lepton collider physics and theoretical studies.

We expect to continue supporting and contributing to the development of detector simulation frameworks, DAQ systems, software strategies and algorithmic approaches, both for specific experiments and for generic detector advances, and to begin participating in R&D for Run II data curation.

We anticipate strengthening our position in the GEANT4 (G4) collaboration by playing a major role in the forming of a US G4 development team, and continuing to contribute in the areas of hadronic physics, model validation, performance enhancement and code quality. We also expect to continue providing support and G4 application guidance to CMS and we are planning to provide similar support to other Fermilab experiments and projects.

We expect to continue our participation in event generator code development and its application, and increase our effort to enable integration of event generator software into useful tools for the analysis of experimental data.

We expect to continue language support for scientific programming by providing expert design consultation and courses for improvement of physics developer's skills. We also expect to continue working toward improvement of language attributes and associated tools relevant to HEP, toward increase of compiler efficiency, and toward optimal utilization of multi-processor and multi-core resources.

We anticipate continuing installation and limited user support for a selected set of debugging and correctness-checking tools which are of prime importance to CMS and other experiments.

We expect to continue to provide support for the utilization of Collaboration Tools.

We expect to support multi-particle dynamics simulation activities relevant to RunII accelerator performance and to the design of a brighter proton source at the laboratory. We expect to continue developing and deploying high-performance accelerator modeling tools in order to enable the above activities and also support future accelerator R&D. By the end of FY10 we expect to have developed capabilities and applications for multi-physics, multi-particle modeling applications.

We expect to continue providing leadership for the SciDAC accelerator modeling collaboration, the COMPASS project.

We expect to continuously search for opportunities to enhance the safety and environmental soundness of our activities, and to develop and monitor metrics and other ways to measure our resources, our performance, and the quality and effectiveness of the tools we provide to the community.

Stakeholders

Our primary stakeholders are the Fermilab experiment user community (including RunII, neutrino, and lattice QCD users), US-CMS, the LPC user community, the Mu2e experiment, the NoVa experiment, the G4 collaboration, the Fermilab astrophysics and accelerator communities, the ComPASS SciDAC collaboration and its sponsors (the ASCR, and HEP offices directly, and NP and BES indirectly), and physics collaborations working on the design of future experiments and accelerators. Within the Fermilab organization, our stakeholders include physicists and other researchers from AD, APC, CD, CPA, PPD, and TD.

Goals and Objectives

1. Develop and support modern computational physics tools that enable the development of HEP applications:
    a) Educate the user community in the use of modern software techniques, scientific tools and math libraries.
    b) Develop and make available to the community tools for studying and enhancing code performance.
        1) Profiling tools, along with database, visualization, and query components to facilitate investigations.
        2) Studies of compiler efficiencies and coding techniques.
        3) Investigations into techniques for optimal use of muti-core architectures.

c) Investigate and support collaboration tools to foster scientific collaboration through effective global communication and meetings.
d) Develop parallel computing algorithms and solvers, and expertise on available (through SciDAC or other sources) parallel computing algorithms and solvers.
e) Continue to develop and support tools for offline computing (frameworks) and DAQ
f) Investigate and support commercial tools for scientific computing.

2. Develop and support applications in support of the Fermilab experimental program:
   g) Support HEP specific tools and libraries that can be used by several experiments.
      1) Develop scientific software libraries at the request of the experiments.
      2) Collect/distribute software developed by the HEP community.
      3) Provide minimal support for legacy products.
   h) Provide support to the Fermilab user community for simulation tools, physics generators, physics analysis tools visualization tools, and applications of our offline framework and DAQ tools and libraries.
   i) Develop code and contribute to testing as part of the GEANT4 collaboration. Align the work as closely as possible to the needs of the Fermilab scientific program.
   j) Develop the software infrastructure needed to perform effective physics simulations and studies, for research into detectors for future experiments (at Fermilab and elsewhere).

3. Develop and support high-performance accelerator modeling applications:
   a. Continue to maintain, support and further develop (interfaces, build system, testing, and documentation) the Synergia framework.
   b. Continue to enhance the capabilities of the Synergia framework with new physics modules
      1) incorporate beam-beam and impedance physics modules
      2) interface to SciDAC e-cloud generation libraries
   c. Continue to maintain and support CHEF.
   d. Continue to develop, in the context of ComPASS, multi-physics simulation capabilities and application of these capabilities.
   e. Utilize SciDAC tools for e-cloud, space-charge, and impedance Project-X (Main Injector, Recycler) and Mu2e applications.
   f. Participate, in the context of the APC, on the e-cloud modeling validation effort based on MI measurements.

Strategies

We will develop new or utilize existing accelerator modeling codes, with modern software architecture, flexible user interface, and well supported build systems. In order to obtain good performance and software design we will seek collaboration with other computational accelerator physicist and computer scientists. For successful accelerator applications and code validation we will collaborate with machine physicists and operators.

For High Performance Computing infrastructure for accelerator modeling we will rely on the support of the Fermilab LQCD group, and the NERSC and ALCF centers, while for the development of software collective services we will seek collaboration with FermiGrid researchers.

We will continue to build expertise in the programming languages, compilers, build systems, and development environments needed for scientific computing for the LHC, anticipated new experiments, and potential future lepton collider experiments.

We will continue support for HEP products and math libraries by developing expertise on existing and/or developing new capabilities.

We will promote the use of common and shared solutions. For example, monitoring and data collection tools motivated partially by the GEANT4 effort can find use locally in the context of grid reliability projects for the lattice effort, and (potentially) globally in the stabilization of RunII operations and general grid applications.

We will increase visibility and acceptance of our work by attendance and presentations at conferences, workshops, and other forums.

We will identify widely used tools and technologies which may be under-appreciated in the physics community, along with next-generation technologies, and develop visions of how these may be introduced to the community to better support our mission.

Recognizing that many of the deliverables appropriate for our activities are inherently challenging to quantify and measure, we will develop better metrics and quantitative measures of the impact of our advances.

We will work with potential collaborators to identify target areas for innovation. In particular, programs of collaboration with the universities which are part of the FNAL managing team, on projects involving student participation and professorial mentorship, can be a valuable resource and an advantage in creating competitive funding proposals.

Resource Needs

The execution of the planned activities for accelerator simulations requires the same resource allocation planned (but not yet executed) for the 2008-09 period: the addition of two FTE for DAQ development and studies of new computing techniques and technologies for future HEP applications.

The resources focused on detector simulations and generators will grow from current levels by 1 FTE, to accommodate the needs of operation and data curation for Run II, CMS infrastructure support and development in the crucial early data-taking period, and medium-range detector simulation activities (including increased practical user support for GEANT4). The resources for physics analysis tools will stay at a fixed level throughout this period. Product support resources (in FTEs and maintenance costs) will be decreased slightly, assuming decisions are reached on which products no longer warrant central support.

Progress Indicators

Successful applications of existing code capabilities, development of applications of the new capabilities, addition of users and collaborators, code releases, code and application porting on supercomputers.

Successful applications lead to talks at workshops and conferences, technical publications, and physics publications.

Maintaining or expanding our non-core funding sources is an overall measure of the success of our goals and our strategic planning.

Feedback from the LPC, and the CMS experiment, will help evaluate the efficacy of our support for that LHC experiment.

Feedback from the Mu2e experiment and the Mu2e accelerator design team will be used to evaluate our success in supporting the detector and accelerator chain design for the experiment.

The NoVa DAQ team and experiment management will provide feedback and evaluations on our NoVa DAQ activities.

Run II physics groups and the Run II department in the CD will provide feedback on our activities for that purpose.

GEANT4 collaboration meeting reports and (roughly) monthly meetings of the performance and validation group (which include the Geant4 spokesperson) will provide feedback for GEANT4 work. There are also quantifiable metrics, including net performance speedup, in this area.

ComPASS collaboration meeting reports and progress reports to the funding agencies.

Instances of "serendipitous" adaptation of tools developed for one focus, by physics groups benefitting in similar ways on different projects, can validate the strategy of promoting common and shared solutions.

Progress reports from experiments and other Fermilab projects receiving services and support.

Additional Information

We expect most major development of software tools for the LHC to be completed before the end of 2009. Once LHC data taking starts, the performance of these tools will be challenged by the scale and problems of real data analysis. This could result in increased demands for support effort by the tools developers in order to assure success.

FY10 Strategic Plan for Enterprise & Collaborative Systems (2010-2014)

R. Karuhn, S. Nolan, J. Trumbo
10/30/2009

Mission

Provide highly reliable, extensible, secure and cost effective enterprise & collaborative systems (ECS) and services in support of day-to-day laboratory operations, and scientific projects and experiments.

Context and Assessment of Current State

ECS supports the Laboratory's core business and administrative systems. These systems are comprised primarily of the Oracle e-Business Suite, PeopleSoft HRMS, Sunflower Asset Management and Deltek Cost and Schedule Control systems. These systems are upgraded regularly to enhance existing systems functionality and service and to maintain currency with vendor support requirements. Upgrades are scheduled and implemented in accordance with business needs and vendor support requirements. The Laboratory's core business information systems are up-to-date and reside at version levels that are supportable by commercial software vendors. These systems are highly reliable and have excellent support capability.

ECS also implements and supports web and collaborative IT solutions. These solutions include a wide variety of QA and workflow applications, web-based reporting, and audit and project management systems solutions. In addition, ECS provides support for numerous homespun information systems and services that support internal division objectives, and scientific operations.

Significant emphasis is placed on minimizing customizations to vendor provided software. This industry "best practice" is a key aspect of the Laboratory's Systems Development Life Cycle (SDLC) methodology for its enterprise systems. It is also a key tenet of the ECS strategic plan. Software development is managed using structured SDLC methods and tools. These tools allow for collaboration on development activities and provide an institutional memory of procedural and problem solving information. Controls such as automated approval hierarchies provide assurance that Laboratory personnel and contractors carry out the systems development procedures management has prescribed for them. Electronic documentation is in place for management, auditor and operational review, and a clear separation of responsibility exists between individuals performing software development (Application Developers) and individuals migrating changes to production systems environments (Production Control).

In FY2008, two external assessments of the Central Services organization were completed. The first concentrated rather narrowly on the current state of the Computing Division's Help Desk and it included recommendations to align the Help Desk with the Service Desk components of the Information Technology Infrastructure Library (ITIL) framework, an internationally recognized and accepted framework for IT service management (ITSM). The second assessment looked more deeply at IT services across the Laboratory and made recommendations for a roadmap leading to ISO20000 certification on a two year time scale (FY2011), as well as organizational changes consolidating the "silos" of IT support across the Laboratory. These assessments, and in particular the ISO20000 certification roadmap, will be used to guide the strategy for ECS. The ECS strategic plan will also be influenced by cyber security guidance as well as directives provided to the Laboratory by DOE.

Vision

Become the preferred provider of highly reliable and cost effective shared IT services in support of day-to-day laboratory operations, scientific projects and experiments. In doing so, promote common methods and tools, and centralized IT services that are appropriately aligned with industry best practices, and that provide for robust IT governance and business continuity assurance.

Stakeholders

- Laboratory user community and scientific collaborations
- Funding agencies (DOE and NSF)
- FRA, Board of Directors, and governance bodies
- Sr. Laboratory management
- Division & Sections
- CD personnel

Goals and Objectives

- Provide highly reliable, extensible, secure, and cost effective enterprise systems and services in support of day-to-day laboratory operations, and scientific projects and experiments.

  - Exceed 99% uptime performance for all production systems
  - Maintain version currency for vendor provided hardware and software
  - Provide timely and accurate information for decision making
  - Implement computer security policies that appropriately safeguard information systems.

- Centralize management and operation of all shared IT systems and services at the lab in order to optimize resources and lessen operating costs.

  - Eliminate redundant systems
  - Consolidate duplicate, competing services
  - Focus priorities and efforts
  - Increase accountability

- Achieve “operational excellence” and ISO20000 certification by 2011 by adopting industry best ITIL practices for service delivery, quality, change control, and customer service and satisfaction.

- Establish and maintain a central IT Service Catalog that is advertised and available for use throughout the laboratory.

- Establish accurate, achievable tactical (1-2 year) plans and budgets, and execute them according to plan. Attain a cost objective of 98-100% of plan.

- Provide quality ECS management practices in terms of planning, leadership, organization, control and communication to ensure results are achieved on time and within budget. Ensure products and services meet customer requirements as demonstrated by 88% positive customer feedback. Measure success through feedback received by customers through an annual Customer Satisfaction Surveys.

- Maintain staff turnover at less than 10% annually.

- Perform all activities in a safe manner. No recordable injuries.

Strategies

1. Establish strong IT governance to appropriately prioritize IT investments

2. Consolidate redundant IT systems and competing services

3. Centralize management of common site IT functions and run as “shared services”

4. Invest in tools and methodology:

    a. Define frameworks and effective approaches for managing software development
    b. Provide direction and guidance in the use of structured analysis and design techniques (SADT)
    c. Provide project management review and governance

5. Adopt industry best practices:

    a. Establish strong relationships with industry and other DOE contractor organizations
    b. Become active participants in technology conferences and user group meetings
    c. Apply formal ITIL-derived best practices to the delivery and support of services.

6. Implement commercially supported, “off-the-shelf” software whenever possible:

    a. Don’t reinvent the wheel
    b. Deliver highly reliable, extensible, secure and cost effective enterprise systems
    c. Minimize customization, require business justification
    d. Adapt business processes wherever possible
    e. Reduce dependency on “key” personnel

7. Establish core teams for major IT initiatives

    a. Define clear project goals
    b. Give project manager complete responsibility and authority
    c. Encourage widespread participation

   d. Utilize outside expertise where business value is added

8. Increase customer autonomy and ownership

   a. Develop functional analysts within the business units
   b. Strengthen customer training and systems knowledge
   c. Document business processes
   d. Assist customers in becoming responsible for business IT functions

9. Establish mutually agreed upon Service Level Agreements (SLA's) with customers

   a. Identify measures of importance and agree upon metrics
   b. Establish and commit to plans
   c. Prioritize service requests jointly with customers

10. Build strategic vendor relationships

    a. Maintain a strong presence with vendors and a keen awareness of their technical direction
    b. Participate in customer sponsored user groups
    c. Maintain currency with vendor product releases
    d. Implement technologies that demonstrate cost effectiveness

Resource Needs

ECS is presently understaffed due to a recent RIF, staff attrition, a reorganization, and an increase in the number of systems and services provided. ECS presently has three open positions, which when filled, will help reduce the impact of this shortfall.

Cross-training is underway to help mitigate reduced staffing levels and associated risk, but it is anticipated that staffing shortfalls will remain in key service areas such as database administration and application development.

Over time, consolidation of redundant systems and services will help free up resources, but in the short-term, we need to rely on the use of consultants for staff augmentation.

Progress Indicators

- Business and administrative systems projects in excess of $500,000 achieve 90% of the schedule, budget and technical milestones specified in the Approved Project Plan.

- Business and administrative systems Information Management (IM) plans are in alignment with the Laboratory's Strategic Plan and are in place by the end of each fiscal year.

- Business and administrative systems IM products and services meet customer requirements as demonstrated by 88% positive customer feedback.

- The business and administrative IM projects are completed as identified in the IM plans and demonstrate measurable improvement and cost effective services and products.

Strategic Plan for Grids (2010-2012)

Eileen Berman, Keith Chadwick, Stu Fuess, Gabriele Garzoglio, Burt Holzman, Don Petravick, Ruth Pordes, Stephen Wolbers
08/13/09

Fermilab's scientific program involves collaborations that are national or international in scale. Computing for these activities is spread among many facilities[1], and a single facility needs to serve many experiments, projects, and collaborations.

Grid (distributed) computing seeks to organize computing along the roles of these experiments, projects, and collaborations, ("virtual organizations"), and resource providers ("sites"[2]) One of the main goals of grid computing is to reduce the coupling between sites and virtual organizations to a minimal number of transparent, standard interfaces and business processes.

Virtual organizations gain from this approach in that they are able to add sites at low marginal cost, and sustain operations. Fermilab gains because it is able to support computing for new virtual organizations at low marginal cost. Both parties gain because the coupling between sites and virtual organizations is reduced to a thin interface, allowing each party to focus on optimizing its processes with relative independence.

Grid techniques are required by the Large Hadron Collider (LHC) software community[3], therefore embracing them is a necessity. Improving upon the current set of techniques and identifying tractable additions and expansions would provide benefit for Fermilab virtual organizations and our external grid collaborators. Primary reference points for locating such mechanisms are the Fermilab Computing Division and Fermilab-hosted experiments/projects, the Open Science Grid (OSG) and the Worldwide LHC Computing Grid (WLCG.) Constraints on acceptable techniques may come from cyber security.

The grid approach creates maximal value when it is broadly accepted. It is necessary to both collaborate with others and to build on our existing strengths. This will both help the virtual organizations and strengthen our ties into a national cyberinfrastructure, enabling us to keep pace with this infrastructure and participate in its evolution.

Fermilab has a history of working closely with experiments/projects, thinking carefully about operations, interoperable (and especially default-allow) security, data movement and storage. These skills will contribute positively to a grid effort which must apply the knowledge gained from the acquisition of these skills to creation and operation of the grid infrastructure itself.

The Fermilab Computing Division faces the challenge of supporting the Run II experiments in a period of increasing luminosity and decreasing available effort, assisting CMS during the critical first years at the new energy frontier, facilitating the use of core grid services of the future programs of Fermilab and maintaining flexibility to understand and participate in the evolution of distributed computing in support of the Fermilab scientific program.

---

1 A set of resources that provide a particular service, e.g. grid facility, storage facility.
2 An infrastructure of heterogeneous resources under a single administrative domain. A single site may contain multiple facilities.
3 CMS, ATLAS, ALICE, LHCb

Mission

To enhance and expand the body of grid software, business methods, and deployment community that is broadly accepted by the Fermilab site and Fermilab based virtual organizations.

To expand and improve upon the relationship that exists between Fermilab based virtual organizations, distributed computing infrastructures (grids) and sites.

To assess and measure the acceptance and effectiveness of these techniques.

To perform targeted research and development to enhance the methods and technologies used in Grid production systems, specifically including computer security, in our distributed open science environment.[4]

Assessment of Current State (2009)

Fermilab offers a production grid facility (FermiGrid), which provides a set of common services and a portal to the OSG. Through this facility, virtual organizations access both FermiGrid and OSG resources. CDF, CMS and DØ generate large MC samples, CDF runs analysis jobs on FermiGrid and OSG resources and DØ has reconstructed more than 1B events and runs primary production on these resources. FermiGrid and OSG are utilized by many experiments/projects at Fermilab. Stakeholder interoperability and opportunistic use of resources commonly occur. Support of OSG virtual organizations is ongoing. FermiGrid is one of the preeminent computing facilities in the OSG. OSG communities have access to significant opportunistic computing resources at Fermilab (~150,000 CPU hours/month). Operational security practices and procedures governing security at the grid-site boundary are being investigated and realized. The continued deployment of high availability techniques is contributing to the strengthening of the distributed infrastructure and increased availability of the offered services. Extensions to the deployed automated monitoring have increased timely response to operational incidents. The SAZ[5] and GUMS[6] supported service levels were increased to 24x7 support. Application of ITIL (Information Technology Infrastructure Library) techniques is continuing and will contribute to the understanding and operation of the infrastructure as an enterprise level activity.

CD is actively involved in grid middleware development, integration, deployment, and support for authorization and identity management, accounting, resource and workload management, data handling, and storage and data movement. Many of these activities involve collaboration with external Grid projects and organizations. The VO Services Project[7] and SAMGrid Project[8] have wound down development and are focused mainly on operations. CD expertise in the above areas is an often-requested capability. Acceptance and use of the glideinWMS[9] deliverables has expanded to additional

---

4 From Wikipedia: Open research is research, conducted in the spirit of free and open source software. If the research is scientific in nature it is frequently referred to as open science; open research can also include social sciences and the humanities.

5 Site AuthoriZation Service

6 Grid User Management Service

7 Provides software solutions for VO user registration and fine-grained authorization for access to grid-enabled resources

8 A general data handling system designed for experiments/projects with large (petabyte-sized) datasets and widely distributed production and analysis facilities

9 Condor glidein workload management system that provides a simple mechanism for accessing grid resources

Fermilab VOs. Awareness and application of good security practices and policies is increasing and contributing to worldwide efforts. Strengthening of the grid infrastructure to production quality has progressed and many components are being heavily used in production. Investigation of enterprise architecture techniques and their application to distributed computing middleware (including generation of metrics-based reports) has begun. Application of these techniques in the area of service monitoring and troubleshooting is being investigated. User support is ongoing and dedicated effort is available for problem triaging and user assistance.

Vision

To integrate HEP computing at US universities to the common grid distributed computing infrastructure. This allows coherent HEP computing to emerge in the US, and ensures that Fermilab is anchored in this infrastructure, allowing it to evolve its computing facilities and expertise in a way that gives maximal advantage to the US HEP community.

To influence the labs non-HEP program (QCD, Astrophysics) to evolve along similar or identical lines.

To be seamlessly integrated into the distributed computing system for CMS as a Tier-1 center and as an analysis facility including the production deployment of a workload management system. Processing resources for event reconstruction, data selection, and regional analysis will be transparently accessible through the OSG infrastructure and will have been scaled to support the additional resource needs. Custodial data storage from the experiment and event simulation centers, as well as data serving to remote analysis facilities will be accessible through reliable and secure common grid interfaces.

In the coming few years, the Run II experiments will have completed their raw data taking and will be focusing on Monte Carlo generation and data analysis. Migration to stable operations will have occurred. This includes a set of core computing services that perform at the scale required to support these activities. Computing services will have been migrated to common grid solutions where this is seen as the most efficient program of work, weighing the development cost against the potential savings in operations and support.

For all experiments/projects where legacy solutions are operated, any transition to maintenance will include the identification of sufficient effort to provide adequate support for the program of the virtual organization.

Fermilab will provide, in addition to the current advanced set of core services, a production set of distributed computing workload management tools in use by multiple virtual organizations, an enterprise level grid aware metrics aggregation service, and monitoring and alarming infrastructure, and the deployment of High Performance Computing (HPC) capabilities for distributed computing. The work with the storage implementation in the Virtual Data Toolkit[10] will enable Fermilab to continue as the primary expert in grid enabled storage solutions for the OSG. Fermilab will have sufficient effort and expertise to facilitate and guide the transition of smaller communities to common grid services and components where desired.

Grid operation will have developed into a long-term sustainable activity with minimal operational effort. Available operational tools will provide information to operations teams, security personnel, and grid

10 VDT – The software stack containing the OSG grid software (services, clients, tools.)

users. Virtualization technologies will be used to improve the user experience and capabilities and ease operational load. There will be seamless interoperation with alternate authentication mechanisms (such as Shibboleth, used in university campus grids). Commercial cloud technologies will be interfaced to, as appropriate, to enhance and extend users access to available resources.

Contributions will be made to research and development to enhance the capabilities and effectiveness of Grid computing in areas of computer security, production distributed infrastructures, and large-scale data, high throughput grids.

Stakeholders

The sponsors of the grid work are the Fermilab Computing Division base program and the DOE and NSF in the SciDAC-2 and NSF sponsored projects (OSG, CEDPS). Contributions are made from the US CMS software and computing project. Effort and deliverables are provided from many groups both internal and external to the Laboratory.

The stakeholders are:

1. The running experiments based at Fermilab (CDF, D0, Minos, MiniBooNE), the neutrino program, and the astrophysics community.
2. CMS.
3. Simulation and theory including accelerator modeling, LQCD.
4. The Computing Division.
5. The members of the Open Science Grid Consortium.
6. The Worldwide LHC Computing Grid (WLCG).
7. The greater US HEP community, including QCD and astrophysics communities.
8. The SBIR community.

Goals and Objectives

1. Perform all work activities with awareness and understanding of appropriate safety practices and procedures.
2. Operate, support, and evolve a robust, effective, secure, local production grid facility, which supports the scientific program of Fermilab (FermiGrid).
3. Increase FermiGrid's operational efficiency by the application of High Availability and ITIL techniques, enhanced monitoring and alarming, and the creation of additional operational tools.
4. Provide a grid services platform including services that are contributions to the OSG and help Fermilab achieve its goals. Special attention should be paid to implementations compatible with the infrastructure as outlined in the CMS Computing Technical Design Report (CTDR). Collaborate with external projects in the development and deployment of these grid services and grid tools.
5. Understand and develop grid security policies, processes, and procedures in order to ensure the secure operation of the production grid facility. Assist the experiments/projects in developing operational security expertise.
6. Understand and work to help provide grid accessible storage solutions that provide for the needs of Fermilab based VO's.

7. Develop security processes and procedures surrounding the software development life cycle. Apply these processes and procedures to existing and future projects.
8. Work towards seamless interoperation between the national and regional grids that Run II and CMS worldwide collaborations depend on– in particular the EGEE (Enabling Grids for E-sciencE Project) and OSG.
9. Work towards interoperation with additional national and regional grids such as TeraGrid and campus grids.
10. Maintain visibility as a leading member of the OSG, used by all Fermilab scientific groups and take leadership roles within this consortium.
11. Maintain and develop relationships and collaboration with sites, software contributors, committees and other bodies sufficient to increase the usability and administration of grid infrastructure integral to the Fermilab mission. Special attention should be paid to the needs of the DOE environment.
12. Understand/Improve the relationship between VOs, resource providers, and Grid providers and the user experience working with grid tools.
13. Offer excellent support to the Fermilab experiments/projects to enhance and enable their efficient use of FermiGrid, OSG, and worldwide grid facilities. Provide a support level that is appropriate to enable the virtual organization to operate in an efficient manner.
14. Promote use of distributed computing facilities by additional experiments/projects including the neutrino and astrophysics programs as needed.
15. Provide outreach and communication activities to further the movement of Fermilab experiments/projects to using the distributed grid infrastructure.
16. Investigate the application of MPI technologies in the grid environment. Deploy these technologies in the grid environment to support the MPI applications.
17. Investigate the utility of scientific cloud computing to the grid community and apply to distributed computing at Fermilab as appropriate.
18. Investigate the utility of commercial cloud offerings to the grid community and utilize these offerings at Fermilab as appropriate.
19. Investigate the utility of virtualization to the grid community and apply to distributed computing at Fermilab as appropriate.
20. Identify a grid/distributed computing based path forward for HEP and astrophysics communities.
21. Contribute to research and development collaborations for advancing computer security for open science.

Strategies

Work with specific experiments/projects to both gain domain knowledge and foster acceptance of the distributed computing approach.

Use CMS and the Run II experiments as exemplars and critical leverage to evolve and expand distributed computing techniques and technologies.

Undertake leadership and supporting roles in various organizations, including the Open Science Grid, and evolving the organization to build acumen in applicable techniques.

Partner and collaborate with other significant contributors in developing and deploying distributed computing middleware and production facility capabilities.

Provide common grid services at Fermilab, including the operation of the FermiGrid campus grid.

Evolve site operational strengths to a distributed computing context.

Work with other departments in CD to facilitate the deployment of grid software and its operation and support.

Collaborate with computer science, mathematics, and domain science groups in the DOE Office of Science and the NSF to advance the state of the art, improve the effectiveness and mitigate the risks associated with open science distributed computing

Resource Needs

All resources purchased for the LHC will be entirely consumed with LHC activities, so there should not be any planned reliance on securing large numbers of opportunistic computing resources from these resources.

FermiGrid resources can be expected to require growth over the next three years, both in the area of service machines (i.e. Gatekeepers) and worker nodes.

CMS effort will contribute to both the experiment specific services and to the common grid services and infrastructures, increasingly shifting to operations from development over the next 2 years.

With the increase in luminosity for Run II, and the increasing interest in workload management systems by more experiments/projects, support of these users and the resources they require will occupy a larger fraction of effort.

Increase in the types and variety of storage services necessary for efficient and effective grid use by virtual organizations and the associated use cases will require increased effort.

Progress Indicators

Development project milestones will be tracked and reported. For virtual organization layers, compliance with standards will be tracked. The following will be reported monthly, taking account of the virtual organization needs:
-- The % of jobs through grid interfaces per virtual organization.
-- The number of different types of jobs using grid interfaces.
-- Number of sites offering turnkey grid access per virtual organization.
-- Efficiency of use of grid resources.
-- Numbers of incidents a week and the successful resolution of said incidents.
-- Number of problems involving associated services.
-- Decreasing time spent in administration of grid services and resources.
-- Decreasing time spent by the users to get up to speed and sustain use of the services and resources.

Risks

The risks of this strategic plan include:

1. The continued support and attention of the underlying Grid technology groups – Condor (batch system support for highly distributed computing) and Globus (toolkit for building grid systems) – to deliver the standard middleware used by all grids.
2. Failure at the annual review for funding of the Open Science Grid.
3. Failures in communication with the virtual organizations to ensure their ongoing requirements and plans are met and/or insufficient resources made available to meet their goals.
4. Failure of internal communication within a virtual organization such that the Fermilab grid activities have no reliable communication mechanism with the VO.
5. Failure to have adequate staffing levels to meet needs and/or goals.
6. CD funding levels may be inadequate for the needed ongoing hardware replacement/upgrades to manage hardware lifecycles.

Strategic Plan for Engineering (2009-2012)

Simon Kwan, Gustavo Cancelo, Vince Pavlicek, Alan Prosser
8/24/2009

Mission

The mission of CD Electronic Systems Engineering (ESE) Department is to provide high quality design, implementation, maintenance and project management support for hardware systems for HEP and Particle-Astrophysics experimentalists. Hardware systems include board and system level hardware and firmware, infrastructure electronics and DAQ systems. In addition, the department will consolidate its position as a leader in integrated test stand development through the development of application and networking software for these systems.

Context and Assessment of Current State

CD Electronic Systems Engineering Department currently provides high quality electronic engineering services to the scientific community at Fermilab and other high-energy physics laboratories and university users. This strategic plan focuses on keeping ESE's leadership position in the field. This strategy relies on the strategic direction set by the laboratory and the division. To fulfill our mission, we also have to coordinate and cooperate with the following activities:

1. FNAL operation – running experiments and test-beam users, construction of new experiments, support of PREP and the Accelerator facilities,
2. LHC – CMS and upgrades for the SLHC,
3. ASTROPHYSICS – JDEM and it's Science Operations Center,
4. Data Acquisition system
5. R&D for future experiments/projects, and advanced accelerator instrumentation and detector R&D.

These activities are either clients, collaborators, or both.

Successful operation of the FNAL accelerators and the collider and neutrino experiments is the first priority for the lab and ESE continues to provide hardware support to these groups as defined in our MOUs.

ESE is currently responsible for maintenance and repair of the HEP equipment available from PREP. Despite the diminishing value of the equipment pool due to aging of its assets, activity continues at a moderate level in support of test beam activities and R&D for new experiments and projects. ESE will continue to retrain the technicians employed in these activities so that they can be utilized in other areas and projects. In tandem, there is the need to strategically reduce the diversity of the equipment pool to minimize the amount of labor needed and yet to keep it a viable and pertinent resource to the Fermilab user community.

The ESE strategic vision recognizes the importance of supporting the DAQ and control systems where appropriate and feasible by enabling the synergies of software and hardware teams working together to support the mission of the division and laboratory. The department plays an important role in helping to set the standards for the HEP application of the Micro Telecommunications Computing Architecture (uTCA).

This plan supports the operation of the forward pixel detector for CMS and this project has been fundamental in positioning ESE for future CMS Tracker upgrades in the scope of the SLHC. ESE engineers have started to play a significant role in ongoing activities for two important areas for the proposed SLHC upgrades. These areas are:

1. Power distribution system studies for the CMS Tracker
2. Optical communications component testing for the Versatile Link Common Project for ATLAS and CMS

We expect our role will grow further within the next couple of years as CMS gears up for the Phase 1 Tracker upgrade.

ESE will continue to leverage its experience with pixel and strip detectors to pursue a leadership role in the detector and systems engineering R&D for the international lepton collider. This effort is expected to build upon recent collaborations with ASIC designers in the Particle Physics Division. The development of powerful data acquisition systems (for example, CAPTAN) is being pursued to also provide new opportunities for collaborations with HEP researchers and the university community worldwide.

This plan also supports the Astrophysics strategy of the division by providing engineering expertise to build a strong proposal for JDEM including a Science Operations Center proposal. ESE is interested in continuing the collaboration with the Particle Astrophysics Center on pursuing various interesting projects.

While there is a decreased emphasis on ILC work, ESE will continue to support the LLRF systems developed for facilities at A0, the Horizontal Test Facility at the Meson Lab, and the cryomodules at NML. If funding could be secured, there is also plan to conduct R&D towards LLRF system for the High Intensity Neutrino Source project (HINS). ESE also collaborates with the ILC 9mA experiment at DESY-FLASH which will answer important questions about the technology proposed in the RDR.

There are quite a few new experimental initiatives including DUSEL, muon-to-electron conversion, and Liquid Argon (LAr) detectors. The department will pursue engineering opportunities with these experiments.

Vision

By 2012 the CD Electronic System Engineering Department will be the premier provider of engineering system services for the scientific community at Fermilab. With the end of the Run II in 2011, there will be no need to continue D0 and CDF hardware support. This effort has been near a minimum for some time and those individuals involved are principally engaged in other projects. By the end of 2011, ESE should be heavily involved in the upgrade projects for CMS (pixels, strip tracker, and trigger). This effort shall ramp up to about a third of the department (~6 FTE). The second third of the department's effort shall be working on HEP projects such as NOVA, JDEM, g-2, COUPP, low-noise CCD readout, and maintaining the PREP pool of equipment for test beam and experimental R&D. Lastly, a third of the group should be working on R&D for new projects including DAQ for DUSEL, Muon-to-electron conversion, LAr, and new initiatives in particle astrophysics, and in the advanced accelerator and detector R&D areas for Project X. The correct balance between the different experiments and projects shall be determined each year based on the laboratory's strategy and availability of funds. By 2012, the future of the lepton collider project should begin to become clearer as the LHC data produces results that may or may not support an ILC type program.

Stakeholders

CMS:

USCMS collaboration at Fermilab and other US institutes, CMS and LHC management teams

Upgrades for SLHC

CMS collaboration, CMS and LHC management teams, collaborators at CERN and institutes working on joint ATLAS/CMS project on Versatile Link., various working groups (e.g. power distribution, trigger, readout) on the CMS upgrade

NOvA:

Project collaboration and management team

CDF:

CDF collaboration, users, and management team

D0:

D0 collaboration, users, and management team

Astrophysics:

CD's Experimental Astrophysics Dept., Fermilab Particle-Astrophysics Center, JDEM management team, collaborators at LBL and SLAC, future astrophysics experiment collaborations).

PREP logistics

CD's CSS Dept.

Advanced accelerator R&D:

AD RF department, LLRF collaboration, ILC, Porject X, HINS project.

Support of BPM/IPM systems:

Fermilab management and MI-BPM, TeV-BPM and TeV-IPM users

PPD:

Physicists and ASIC designers working on 3D and new trigger proposal for SLHC upgrade, strip telescope for IHEP(Beijing)

MTEST and test beam users

Test beam users who would use the pixel telescope (.e.g T992)

Mu2e:

Mu2e collaboration and management team

Goals and Objectives

- Perform all activities with good engineering practices and a high level of professionalism and excellence
- Adhere strictly to safety guidelines established by the Laboratory ES&H section and OSHA
- Adhere strictly to documents and procedures as documented and approved by the Division and Division's Safety officer
- Provide users with high quality and well documented engineering services

- Provide users with high quality support for engineering projects
- Provide expert advice and help in designing engineering systems for experiments
- Perform engineering related design and software reviews for projects
- Provide members of the department with challenging and rewarding assignments
- Provide members of the department with opportunities to enhance their skills both through work assignments and training opportunities and to advance their professional career

Strategies

- Position ESE to take on new projects for the department (e.g., SLHC, DUSEL, LAr, mu2e, HINS, project X etc.). We'll achieve that by leveraging our experience and continue to establish a reputation for excellence through successful delivery of systems. The group has the core competencies and the overall domain expertise to take on these projects. Nevertheless, the group working on these projects requires training to make a considerable contribution to the effort.
- Maintain a strong sense of self esteem among team members by keeping them involved, encouraging peer cooperation and positive interactions, maintaining an open door policy, reminding members they are encouraged to speak freely, and demonstrating that the opinions of everyone on the team are important.
- Develop engineering systems with a forward looking perspective to see how reuse and leveraging of effort could provide solutions in multiple areas. The expertise gained on the PHENIX Pixel Plane project and instrumentation for the CMS Forward Pixel Detector is being used to create new, powerful data acquisition and control systems for future devices (VIP) and systems (SLHC upgrades, high performance test beam pixel and strip detector telescopes).
- Continue encouraging our technical staff to pursue research opportunities in data acquisition systems, detector instrumentation, and timing and control systems. We'll achieve that by continuing collaborating with colleagues at Fermilab and around the world, participating at conferences/workshops and writing technical papers and reports about our work.
- Remain professional at all times with clients or other members of the team, and seek out management to resolve conflicts when necessary to avoid appearances of unprofessional behavior. Professionalism is a key component to a successful business model whether working for internal or external clients.
- Reduce the equipment pool supported at PREP so that users can be provided with a core set of pertinent assets. All other redundant and unnecessary items will be directed to surplus.

Resource Needs

The leadership role the ESE has kept during its many years of existence leverages its continuous investment in the improvement of our engineers and technicians. Thus, we will require the resources necessary to support our training needs. Furthermore, we utilize summer and coop students to fulfill a two-fold goal: they provide the section with fresh skills to help us in several different projects while we provide them with real life projects that augment their education.

The 19 members of the engineering department will be fully occupied with significant work at the beginning of FY10. If SLHC and other new experimental efforts (e.g. DUSEL and the muon-to-electron conversion experiment) ramp up, we will need more engineers possibly in FY10 and for sure in FY11. We will need more software support for FY10 for the new DAQ projects. Due to the reorganization of

CD, we need to formalize a work plan together with the CET group within FPE/ASD on new DAQ projects.

Progress Indicators

The progress of the various projects will be measured based on the work breakdown structure (WBS) of each project. The performance of the group cannot be measured solely in the compliance with the WBS of each project but the WBS provides an objective measurement to assess the status of our activities. Thus, subjective metrics must be implemented to better gauge the group's performance. Subject metrics will be determined in case by case basis. Examples of such metrics include:

- Best practices analysis: how much time should we need to perform the task and how much effort did we employ
- Cost analysis: is the project's cost within budget?
- ROI analysis: how is the return on our investment (both labor and M&S)? Should we relocate funds to another part of the project? Are we spending resources wisely?
- Are the major milestones of a project being met on time? Are obstacles and risks identified and addressed before they jeopardize a project's schedule?
- Safety compliance indicators.

Strategic Plan for Computer Security (FY10-13)

Joe Klemencic, Ron Cudzewicz, Mark Leininger
10/30/2009

Mission

The mission for Computer Security is to:

- Implement and oversee the security life cycle process, including required documentation, processes and certifications.
- Manage all computer security related audits.
- Contribute resources to development of security strategy for Open Science community
- Perform R&D in the area of Open Science security
- Decrease the response time to threats through process automation;
- Deliver robust computer security tools that are user-friendly
- Encourage a participatory culture of Integrated Computer Security, which gets people "on board";
- Remain proactive in understanding and guarding against new threats while maintaining an appropriately open computing environment;
- Continuously improve computer security training programs to sharpen customer and sysadmin skills in response to new threats and technologies;
- Facilitate the day-to-day execution of the laboratory computer security program.

Context and Assessment of Current State

The Computing Division Computer Security Team provides computer security and safeguards guidance to the general laboratory community on cyber best practice methods, sensible and safe computing operations, a balance between government cyber objectives and the mission of the laboratory and facilitates the day-to-day execution of the computer security program. The threat and regulatory landscape rapidly change, and computer security must be nimble to quickly assess new threat impacts to the laboratory's mission and implement solutions that fit a risk/cost model, while maintaining compliance with government regulations

Future Challenges

- Keeping up with the rapidly increasing bandwidth requirements of the Physics program while capturing enough data for adequate trending, inspection and analysis.
- Developing sufficient capabilities to positively impact the developing Grid security model.
- Maintaining an accurate vulnerability assessment and services inventory of network connected nodes with minimal impact to the computing resources.
- Create modular strategies that are amendable to impending governmental requirements such as HSPD-12.
- Ensuring that the openness of the laboratory Internet presence is preserved through expedited threat analysis and remediation.

Vision

We will provide guidance for cyber best practices and policy enforcement to the laboratory. We will continue to evaluate emerging threats within our risk model. We will continue to be flexible technology leaders relying on state of the art alternatives over prescriptive controls, which often produce negative impacts on the Laboratory's mission. We will continue to strengthen our outside relationships to ensure the laboratory is best in class and leaders in computer security practices without the reliance on over hyped technologies. We will lead the Grid security efforts to ensure sensible but enforceable policies are created. We will collaborate with our peers both inside and outside of the laboratory.

Stakeholders

Stakeholders for Computer Security are

1. Funding agency (DOE) , lab senior management, and everyone with a stake in the lab's reputation
2. Domestic and International FNAL supported scientific research
3. Non-scientific support computing

Goals and Objectives

Major objectives and goals are:

1. Ensure current laboratory computing security policies are enforced
2. Emerging threats are properly assessed against the laboratory operating environment
3. Keep the cyber security life cycle process current
4. Assess effectiveness of operational, management and technical controls
5. Become integrated in the Grid security practices and policy discussions

Strategies

1. Computing policy enforcement

   Continue regular meetings with the GCSC's and system administrators to keep them aware of current computing policies. Evolve the scanning and vulnerability detection and management process to include detection of unaccepted risks.

- Threat assessment strategies

  Assess vulnerability reports from various sources such as peer contacts, mailing lists, news groups and web sites. Monitor our existing data metrics and expand our data gathering scope.

We must also work diligently to keep pace with the rapidly expanding bandwidth offerings.

- Cyber security life cycle strategies

  Encourage the usage of the ST&E tool for recording of the various stages of controls. Roll out the CSA application to feed data into the risk assessment stage.

- Controls assessment strategies

  Expand our scanning and vulnerability detection methods through less intrusive mechanisms. Ensure that documented controls are correctly assessed and implemented through the cyber security life cycle process.

- Authentication strategies

  Promote the use of Central Authentication

- Grid security strategies

  Become a more active participant in the Grid policy and best practice discussion through the recently hired FTE. We intend to be a leader of Scientific Grid Computing while ensuring governing computer security policies are sensible, not prescriptive and enforceable while accepted by the worldwide community.

Resource Needs

The years 2010-2013 are transition years for Fermilab. The LHC will be ramping up in while the Tevatron experiments will ramp down. Both activities will bring in new expectations of resource access and data movement. CST will need to remain sensitive to the scientific communities' needs while complying with regulatory objectives. Grid computing is one area where there is a direct conflict between the scientific community and the computer security expectations by the Government. We will require that major Grid participants conduct Grid computing in a way that does not sacrifice computer security objectives. Over the next year or two, the network bandwidth will far exceed our collection capabilities, forcing us to revert to sampling of datasets for anomaly detection until we can develop higher speed collection processes. With the growing number of regulatory obligations, Counter Intelligence data requests, rapidly emerging threats and current and future support of specialized applications, additional effort especially in the area of software development must be allocated to ensure the software processes remain nimble and extensible to meet the ever changing needs.

Progress Indicators

1. Audit results
2. Lab PEMP goals
3. Semi-annual reports and a taking stock meeting with stakeholders.

4. Metrics on trending analysis
5. Metrics on vulnerability detection and repeat offenders
6. Metrics on incident reports
7. Metrics on planned and unplanned downtimes.
8. Published project progress and milestones.

Additional Information

The total amount of core work exceeds current levels of staffing. Failure to produce adequate additional funding and resources will severely impact the ability of Computer Security to meet its goals and responsibilities.

Strategic Plan for Facility Operations (2010)

Adam Walters
8/11/09

Mission

Operation, maintenance, improvement, design & construction of Computing Division computer rooms, buildings & grounds, office and public areas in the Feynman Computing Center with offices and high availability computing; the Grid Computing Center with high density computing and specific high availability computing rooms (Networking & Tape Robot Room); LCC with LQCD and AMR computing; WH8-Fiber Central and specific office areas in Wilson Hall. The support and services for the areas listed above are provided by Facility Operations and are required for Computing Division to carry out its mission.

Current State

- The Data Center availability metrics have exceeded the documented goals. The Division has kept up with new requirements and procurements by constructing, upgrading & optimizing computer rooms. In the last year, much focus was placed on the improvement of FCC2.

- Buildings have been operating reliably, but have fallen behind in the replacement of aging infrastructure. Maintenance of critical equipment and infrastructure is performed by FESS or outside contractors.

- Offices have kept up with personnel needs, although barely, and need options for growth. There is need for a better office planning process.

Vision

- Continue to achieve excellent availability metrics. Provide space, power and cooling commensurate with stakeholder computing needs using modern computer room capacity planning practices and tools. Improve power efficiency in operation of computer rooms. Address the need for Test & Development computer room capacity and develop a plan for computer rooms/server rooms across the Fermi site which are managed by other Division/Sections.

- Provide reliable operation, maintainability and needed improvements to buildings. Adopt industry good practices to decrease energy usage within FCC where possible.

- Provide sufficient office space to meet personnel needs. Upgrade selected office space to achieve modern esthetics and efficiency standards. Enhance the planning process for office space.

Stakeholders
Computing Division, running experiments, other D/S (BSS, FESS, PPD), GRID users and others

Goals and Objectives
Provide necessary computer room infrastructure maintenance, improvement and operational services for existing and new computer equipment.

Provide necessary building utilities maintenance, improvement and operational services.

Provide suitable & safe working conditions for all CD personnel including adequate work spaces; environment: ventilation, air quality, heating, cooling, humidity; ergonomics: workstations; furniture & storage; renovation & construction of offices; personnel relocation

Strategies
1. Optimize the maintainability of the Data Centers through comprehensive maintenance contracts
2. Proactively correct failures and address infrastructure risks and gaps in monitoring to increase the availability of the Data Centers
3. Provide Data Center capacity to meet computer acquisition timetables and work with stakeholders to plan for space and installations
4. Complete provisioning of remote monitoring for incident alerts and collection and analysis of energy effectiveness of computer rooms and building infrastructure
5. Follow architected model for computing which includes the addition of FCC3 High Availability computer room and plans to fallow FCC1
6. Develop & execute a plan for better capacity forecasting & planning
7. Execute replacement plans for aging/failing equipment
8. Modernize building and execute improvements
9. Formal maintenance program & documentation
10. Provide options for additional office spaces

Resource Needs
Capacity planning and forecasting is an area where better processes and tools will be necessary to predict the future needs for data center capacity. This is coupled with development and management of monitoring applications & tools targeted at understanding capacity usage and energy effectiveness. The tools must assist in the planning and change management of computing assets throughout the life cycle (procurement, modification, retirement).

The new applications & tools associated with Capacity Planning, Facility Access & Security, Office Planning, Asset Tracking and other areas require a significant effort by database knowledgeable personnel.

Given the forecast & schedule for new construction of computer rooms, we need to conclude that some existing computing equipment may need to be relocated and incorporate this fact into our planning.

Progress Indicators

Achieve Data Center Availability Goals as described in the Facility Operations Outages and Availability Report (Docdb# 3123). Formalization of the process to forecast data center capacity 1-5 years.

Execution of general building improvements with emphasis on the replacement of FCC heat pumps (office units).

Degree to which we meet office space needs of personnel and the Division.

Project Plan / Schedule

FY09	FY10	FY11	FY12	FY13	FY14	$M
FCC3 HA CR Phase I						4.5
FCC2 CR Cooling Upgrade						4.5
	Re-Provision GCC CRA[1]					0.6
	FCC2 UPS Replacement[2]					0.3
	GCC HA Gen. - NR & TRR[3]					2.5
	LCC Cooling Upgrade					0.6
	LCC FlyWheel UPS					2.0
		Computer Room D[4]				8.0
		FCC CR Fire Det. Upgrd[5]				0.4
		FCC2 Structural Upgrade[6]				0.4
		FCC Office Improvements				0.4
			FCC3 HA CR Phase II			6.0
			Computer Room E			8.0
			WH Office Improvements			0.8
			FCC Office Improvements			1.8
				FCC Concrete Col. Repair		0.6
				FCC UPS/Electrical Upgrd[7]		5.0
				Re-Provision GCC CRB		2.2
					GCC UPS Upgrade[8]	4.5
					GCC CRF (ultra HD/HA)[9]	6.0
9.0	6.0	9.2	16.6	7.8	10.5	59.1

[1] Provision 208V electrical distribution, 15kW/rack density, 42" racks, remove ceiling
[2] Replace end of life Best 60kVA UPS with 100kVA UPS & upgrade electrical distribution
[3] Standby generator for GCC Network Rooms & Tape Robot Rooms and mechanical/cooling
[4] Provision high density Computer Room D and shell for Computer Room E and support areas
[5] VESDA fire detection systems in FCC computer rooms
[6] Modifications to FCC2 structural systems to accommodate racks up to 2000 lbs (908kg)
[7] Replace end of life UPS systems for FCC computer rooms
[8] Replace end of life UPS systems for GCC computer rooms
[9] Provision a computer room for high density / high availability computing
Next Generation Computing Facility – FY10 $25M; FY11 $25M

Strategic Plan for ES&H (2010)

Amy Pavnica
8/12/09

Mission

Our mission is to provide a safe and healthful workplace for all employees while having minimal impact on the environment.

Current State

Computer rooms:

- CD employees have worked over 3 million hours without a lost workday case.
- CD is heavily involved in the Federal Electronics Challenge, which is an executive order to go "green" with desktop/laptop purchases, use and disposal
- Our ES&H procedures have been formalized and controlled.
- We now have an internal auditing program and procedure.

Vision

Through continual improvement of our ES&H performance, and compliance to legal and regulatory requirements, our goal is to reduce ES&H impacts associated with our mission and the mission of the laboratory.

Stakeholders

Computing Division

Goals and Objectives

- Implement practices to reduce risks and to prevent injury to our employees.
- Maintain plans to minimize and respond to emergencies.
- Examine and guide ES&H practices of our contractors.
- Manage and reduce consumption of natural resources and generation of waste through design, reuse and recycling.
- Measure performance by monitoring and improving compliance with ES&H Policies.
- Provide information regarding ES&H practices and policy through dialogue, training, and openness with employees.

Strategies

CD Safety Plan (revised annually)

Resource Needs

- Budget for materials to address maintenance and improvements in ES&H

Progress Indicators

- Leading/Lagging Indicators (kept by ES&H)
- Accomplishments from annual CD Safety Plan

FY10 Strategic Plan for Division Infrastructure

(2010-2012)

Ruth Pordes, David J. Ritchie, Chander Sehgal
8/31/2009

Mission

The mission of the Computing Division's Infrastructure is to enable a more capable and effective CD staff by providing excellent administrative, logistical, communication, programmatic, and education/training support across all aspects of the Division's activities and organization.

Context and Assessment of Current State

The strategic plan for Division Infrastructure is guided by the strategic directions across the Computing Division. Division Infrastructure covers activities in support of the Laboratory's administration and infrastructure as well as those that support the growth of the of the Computing Division staff in their ability to execute their work activities.

Stakeholders

Stakeholders are the Laboratory's CIO, all Computing Division organizations, all members of CD, as well as the lab divisions and sections that leverage CD infrastructure activities.

Goals and Objectives

- Encourage continuous improvement in staff skill and performance to effectively accomplish the CD work program
- Support Division conference, committee, travel, and other such activities in response to Laboratory or external organization requests.
- Enable process improvement, training, and management activities that have broader scope and coverage than those required for the immediate work program services or project activities.
- Provide and improve internal and external administrative assistance, HR and infrastructure support for the division office, departments and project activities in the Computing Division
- Implement and facilitate communications internally to the Computing Division and externally to the Laboratory and our peers.
- Increase the coherence in the use of tools and technologies

Strategies

1. Work with all CD organizations to ensure that all staff members have access to training, tools and incentives to develop professionally in a way that is aligned with the lab's scientific program needs and their own professional goals.

2. Increase support in all administrative areas through ongoing improvements and specific projects targeted to advance specific technologies and processes.
3. Work with all CD organizations to enable the program and process improvements to enable them to increase the effectiveness and quality of their program of work.
4. Improve communication and collaboration within and outside CD.
5. Increase visibility and acceptance of our work (including scientific, production solutions and R&D) by responsibilities and attendance at conferences, committees, workshops and other forums.

Resource Needs

We have to rely on the communication and collaboration of personnel, as well as the contributions of staff in specific activities to improve the Division Infrastructure.

Progress Indicators

1. Advances in skills of Computing Division staff.
2. Measurable improvements in the efficiency of administrative processes.
3. Assignments and invitations to participate in conferences, committees, workshops, reviews and other such activities.

Additional Information

Provide a Division Infrastructure that meets the needs of the organizations and members, as well as makes recognized contributions to the Laboratory's, collaborations and external organizations infrastructure and administrative activities.

Strategic Plan for the
Office of Performance Management (2010-2013)

Bill Boroski
10/27/2009

Mission

The mission of the Office of Performance Management is to support work planning and execution within the Computing Division by providing financial management services and the frameworks, systems, and support necessary to improve project execution, operational effectiveness, and quality assurance in a quantifiable manner.

Context and Assessment of Current State

The Office of Performance Management comprises two sub-groups: the Financial Management Group and Project Management Group.

Financial Management

The Financial Management Group (FMG) is responsible for providing financial management services to the Division. This includes:

1) Participating in financial planning and tracking within CD;
2) Responding to data calls and preparing financial reports internal and external to CD;
3) Providing accounting support and other finance-related services to CD personnel (e.g., effort reporting and requisition processing); and
4) Managing software contracts and licensing agreements.

FMG works closely with Division line management to plan and prepare annual budgets; assess cost performance; and prepare cost forecasts. Budgeting and reporting processes are complicated due to the variety of internal and external funding sources and the corresponding set of associated rules and guidance. To accommodate this, a complex set of tools and systems exist and are used to support CD budget planning and tracking activities. While effective, opportunities exist to improve functionality and ease of use for line managers, particularly in the areas of budgeting, forecasting, and performance tracking.

Financial reporting to external stakeholders is completed in a timely and thorough manner. Financial reporting to internal stakeholders is provided through a set of online reports. Monthly cost and effort reports are prepared and made available to line managers to track performance against plan and identify concerns. Over time, additional tailoring of these reports to the needs of line managers will improve financial transparency and the flow of information throughout the organization.

Contract and software license management processes are in place and well-managed, but additional process improvements would reduce the time necessary to demonstrate compliance. Online effort-reporting and purchase requisition systems are in place and meeting internal needs, but there are opportunities for process improvements to increase efficiency (e.g., to reduce the time required to obtain requisition approval).

As CD continues to implement IT Service Management (ITSM), FMG will be called upon to help quantify the cost benefits and value of IT services provided by CD. FMG will work with the ISO20K implementation project managers to develop and implement financial management processes and procedures in accordance with ITIL frameworks.

At present, the financial monitoring and analysis capabilities of the FMG are limited by resource availability. Improving efficiency in other areas should help free resources for more trending and analysis. Longer-term goals include the development of appropriate performance and value metrics, the monthly analysis of cost and performance data, and production of meaningful scorecards and reports that document and trend performance, and identify areas for improvement.

Project Management

The Project Management Group (PMG) is responsible for providing project and operations management services to the Division and external stakeholders. Specific responsibilities include:

1) Executing project management responsibilities for the LQCD, Open Science Grid (OSG), and US CMS projects;
2) Establishing a Project Management Office (PMO) to help CD personnel more effectively plan, track, monitor, and assess progress; and to monitor and track the performance of the CD project portfolio;
3) Assisting in the planning and execution of projects internal to CD;
4) Coordinating the annual work and budget planning and review cycle through the use of strategic and tactical plans, and budgeting tools; and
5) Coordinating the implementation of the Fermilab Integrated Quality Assurance (IQA) program within CD;

Several members of the PMG have direct project management responsibilities, as designated project managers on the LQCD, OSG, and US CMS projects. All three projects are managed using established project management methodologies. Project management responsibilities are defined in the respective project management plans.

As the wide array of computing systems and services delivered by the Computing Division becomes more complex, and as demands on available resources continue to increase, it is critical that we implement a business process and supporting system to better manage the portfolio of CD service operations and new projects. In particular, we need to implement better methods of monitoring the delivery and measuring the performance of services provided. We also need a better process for planning, prioritizing, and tracking new projects required to meet current and future demands. Although some work has been done to develop frameworks, guidance documents and templates, all are in draft form and no tools or systems are in place to monitor, track, and report on the performance of service delivery, or the progress of CD projects. More work is required to implement effective portfolio management, which includes developing effective processes, meaningful scorecards, dashboards, and metrics to plan, prioritize, track, and convey performance information.

The PMG provides project management support on internal CD projects by working closely with CD staff and external stakeholders, and by deploying standard project management methodologies on internal initiatives. We will build on the success achieved to date by further improving the effectiveness of our planning and reporting tools and methodologies.

The PMG actively participates in the annual work planning and budgeting process by coordinating and overseeing the preparation of strategic and tactical plans; and by coordinating periodic reviews of progress against tactical plans and cost performance against budgets. Guidance documents and templates are in place and updated annually to reflect process improvements. We will further refine the business process, work flow document, and annual schedule for the planning process to spread out the work more evenly during latter part of the current fiscal year.

The PMG is responsible for implementing the Fermilab Integrated Quality Assurance (IQA) Program within CD and will continue to work with CD personnel at all levels to effectively implement IQA.

Vision

The Office of Performance Management fully supports the belief that CD is a world-class provider of high-performance computing and IT services, not only to Fermilab, but to the world-wide high energy physics community and beyond. We place great emphasis on working collaboratively with others and performing our work safely. We take great pride in our work.

We play an important role in CD by providing financial management services and systems, and by assisting with project and operations performance management. In particular, we support line management by providing services, tools, and frameworks that support planning, operations, compliance, and quality assurance.

Our goal is to provide services and management systems that will enable CD to better plan, track, monitor, and assess work performance. Many systems, services and tools exist and are used throughout the Division, but require further evolution in the spirit of continuous improvement. We will build on the existing capabilities of the FMG and PMG groups to increase the level of support and services provided to the Division. The skills and capabilities of individuals in these groups are broad and deep; they overlap in some areas and are complementary in others. We will leverage the existing skill set to assist the Division in meeting strategic goals and performance objectives. Where necessary, we will strive to expand our skill set through training, professional development, and an increase in staff size.

As the complexity and interaction of funding sources increase, and as budget pressures mount, CD will need to work ever more efficiently and effectively. We will need concrete methods of demonstrating that our services are cost-effective and that our commitments are met on time and within budget. OPM will play a leading role in implementing systems and processes that demonstrate that IT services provided by CD are cost-competitive and of high quality, and therefore of significant value to Fermilab and the high-energy physics community.

We will support the Division's efforts to implement IT Service Management and the goal of achieving ISO20000 certification in specific parts of the organization by the end of calendar year 2010. We will be an active participant in the execution of the ITSM implementation roadmap and will strive to improve our processes in accordance with ITIL best-practices as appropriate.

Stakeholders

Stakeholders include the funding agencies (DOE and NSF); Fermilab management; the Fermilab Finance Section, as well as other Divisions and Sections; the LQCD, OSG, and US CMS projects; scientific collaborations supported by CD; and all members of CD.

Goals and Objectives

1. Improve the annual work and budget planning process. Simplify and streamline the budget input process. Continue to improve our ability to plan, monitor, control, and predict needs and resources.
2. Implement a performance-monitoring system that tracks and reports performance against critical success factors, key performance indicators, and metrics.
3. Provide a clear, concise set of financial reports for CD line managers that improve financial accounting, transparency, and forecasting.
4. Respond to external and internal data calls in an accurate and timely manner.
5. Maintain responsibility for the annual Fermilab self-assessment of CD activities.
6. Participate in the transformation to IT Service Management and the achievement of ISO20000 certification. Support the planning and execution of transformation projects and initiatives. Develop the management systems necessary to support IT Service Management and the tools to quantify and measure the cost and value of services provided.
7. Satisfy the project management performance requirements for the LQCD, OSG, and US CMS projects.
8. Implement a functional Project Management Office that enables improved project planning and performance tracking.
9. Implement effective project portfolio management, which includes developing processes and systems to plan, prioritize, and approve projects; and meaningful metrics, scorecards, and dashboards to monitor, track, and convey project performance information.
10. Implement the Fermilab Integrated Quality Assurance Program at an appropriate level and in accordance with the laboratory implementation schedule.

Strategies

1. Maintain and as necessary improve the business process and schedule for the annual work planning and budgeting process that takes into account annual funding cycles and conflicting resource demands (e.g., performance evaluations, summer vacations, etc.). Improve the budget input process by simplifying methods for rolling up budget data. Better utilize resources to increase our ability to review, analyze, and trend performance data; to track actual performance against plan; and to better forecast future needs.
2. Develop and implement an online performance monitoring system that includes a high-level scorecard summarizing performance against key performance indicators, and lower-level reports that allow drill-down to understand performance achievements and trends.
3. Develop a set of online financial reports geared to the needs of line managers. The reports will contain the information managers need to track and report cost performance against budget.

4. Develop a schedule showing when external data calls are anticipated, based on past experience, to reduce last-minute report preparation. Delegate report preparation to teams of individuals within the OPM organization to increase spread and depth of knowledge. Complete all reports accurately and according to data call schedules.
5. Maintain a schedule for the annual Fermilab self-assessment based on past experience, gather input from CD personnel, prepare the CD self-assessment report, and submit it to the Fermilab Office of Quality and Best Practices.
6. Provide ITIL training to the OPM staff. Work closely with external consultants and CD personnel as appropriate to develop and execute work plans associated with ITSM practices. Lead efforts to develop systems and tools that will help quantify and measure the cost and value of CD-provided services.
7. Dedicate sufficient resources to the LQCD, OSG, and US CMS project management efforts to meet the performance requirements for these externally-funded, high-visibility projects.
8. Provide a set of guidance documents and templates designed to facilitate improved project planning and execution within CD. Provide in-house training and guidance for individuals charged with executing internal projects.
9. Develop a process for summarizing performance for the projects undertaken by CD. Commercially-available tools and methods will be researched to identify a cost-effective solution appropriate for CD. Implementation schedules and work plans will be developed and implemented based on resource availability.
10. Continue to coordinate and oversee the implementation of the Fermilab IQA program within CD. One member of the OPM will serve as the CD Quality Assurance Representative (QAR) on the laboratory implementation committee. This person will be responsible for working with CD line management and others to develop and execute a plan for implementing the Fermilab Integrated Quality Program in a measured and meaningful way.

Resource Needs

The current level of effort in OPM is barely sufficient to properly support the implementation and execution of a Project Management Office; the monitoring, trending, and analysis of financial and operational performance data; and the CD implementation of the IQA Program. Ideally, one additional person is needed to assist with CD project execution and IQA implementation, and one additional person is needed to help with financial trending and analysis, the creation of performance scorecards, and the monitoring and analysis of performance data.

Progress Indicators

1. A work flow and schedule for the FY2011 work planning and budgeting process will be vetted and published by May 2010.
2. An online performance monitoring system containing key performance indicators will be developed.
3. A set of financial reports tailored to the needs of CD line managers will be developed.
4. All requests for financial data calls will be met. Reports will be provided on time and with accurate information.
5. The OPM organization will manage the preparation and submission of the Fermilab FY2010 Self-Assessment of Scientific Computing Activities.

6. OPM will actively participate in the planning and execution of ITSM-related projects as requested. A management system to support ITSM, along with the tools necessary to quantify and measure the cost and value of services provided, will be developed.
7. All project management requirements for the LQCD, OSG, and US CMS projects, as defined in the respective project management plans, will be successfully met.
8. Project management guidance documents and templates will be developed and made available to CD staff. In-house training will be provided as requested. PMG staff will assist CD personnel in the planning and execution of projects, and will in some cases serve as project managers on specific initiatives.
9. A portfolio project management system will be implemented that facilitates improved planning and summarizes performance against goals for all CD projects that exceed pre-established thresholds.
10. Fermilab IQA implementation within CD will be accomplished in accordance with the approved laboratory implementation schedule.

Strategic Plan

For IT Governance and Oversight (FY10-13)

Don Petravick
10/30/09

Mission

To provide computing at Fermilab that is responsive to Lab's current and future needs, by establishing a system of comprehensive decision rights that yields agreement on goals, portfolios of coherent projects to meet the goals, and the ability to maintain IT operations.

Current State

The computing division is evolving into a more nimble and flexible organization. The division continues to increase the level at which it works – increasingly using knowledge from the outside world; buying tools and expertise. The division continues to develop core competencies and is increasingly able to maintain a sharper focus on the business needs of the laboratory, while maintaining a keen eye on IT and computing technology.

Because of these developments, the work of

- deciding what to do,
- choosing the correct set of tools, methods, and context to use,
- executing projects and operations well and deterministically,
- maintaining the right set of core competencies, and
- maintaining an organization that can adapt and respond to needs for change

require overall governance and oversight.

Governance

The governance function of the division decides what the division will do, and how it will be done, by assigning decision rights and accountabilities to roles in defined processes. The governance process itself produces processes, policies, roadmaps and plans which ensure that IT investments are aligned with the lab's strategic plans, business plans, goals and priorities.

While these governance processes are being coherently defined, several building blocks have been identified.

1. A Service-Oriented ISO2000/ITIL operational framework is used to provide a framework for services that include the decision rights needed to meet defined levels of reliability. The project has implemented a service desk, incident management, and problem management.

2. An enterprise architectural process that holds the decision rights for the common requirements, uniform principles and architectural models that constrain projects to assure they are consistent

with the envisioned future state of the laboratory. A nascent EA effort has been established, with the goals of establishing a first level of architectural maturity in year one.

Oversight

The oversight function of the division assesses the overall work of the division and synthesizes the needs for internal change. A essential input to oversight is key progress indicators, (KPI) whose formulation are overseen, and which are analyzed by the Office of Performance Management, and consequently made available to the oversight processes. The oversight processes are under development. It is envisioned that there will be at least two levels of oversight.

The first level of oversight, the IT executive Council, will represent the point of view of the whole laboratory, and will draw effort from throughout the laboratory, and focus on how well the division meets the current needs of the laboratory, and what the future needs of the laboratory are. The IT Governance Institute lists 4 key questions that governance must address:

- Are we doing the right things?
- Are we getting the benefits?
- Are we getting them done?
- Are we doing them the right way?

A second level of oversight which is to be designed, but is likely to be the IT council, will look at the division from the interior point of view, and being familiar with its internal processes, will guide the internal evolution of the laboratory's IT at a greater level of detail.

Vision

Enterprise Architecture:

- Translate the Labs vision and strategy into effective enterprise change by creating, communicating, improving and enforcing the key requirements, principles and models that describe the enterprise's future state and enable its evolution.

ITIL Process Management:

- The production work at Fermilab is conceived and run as services, organized about the ITIL Framework, and certified by ISO 20000.

Stakeholders

Fermilab as a whole, including divisions, centers and sections, as well as experiments and other scientific activities.

Goals and Objectives

Enterprise Architecture:

1. Organize
2. Staff

3. Facilitate the development of an information architecture capability, using master data as a key project
4. Architectural involvement in new projects with significant architectural impact.

ITIL Process Management:

1. All ITIL process will be implemented by the end of FY2010.
2. An identified set of services will be operating in an ISO20000 compliant fashion.

Strategies

Enterprise Architecture:

1. Coherently analyze the things the lab needs to construct, providing an enterprise level breakdown of the IT or computing scientific components, aligned with the overall architectural direction of the CD. Accrete architecture in response to the needs of the projects.
2. Use consensus methods to propose architecture; Architecture certified by IT governance.
3. Will work with the CD to establish competencies which are not really EA, but foundational to having effective EA. (Technical Architecture, Information Architecture, Business Process Architecture)

ITIL Process Management:

1. Keep the CIO involved, and use the CIO's guidance and authority to establish ITIL processes and expand them to their full scope.
2. Provide training for ITIL managers and others involved in managing the ITIL processes, guided by the ISO20000 training plan.
3. Budget for and provide effort for ITIL Managers and staff for ITIL processes in the stabilization (e.g. "commissioning) and operational phases.
4. Rely on the Lab Scientific Core Services Quadrant for IT management process infrastructure, and ongoing changes to support evolution of Fermilab ITIL processes, which is currently instantiated in the BMC Remedy tool.

Resource Needs

Enterprise Architecture: Staff, at least 1 FTE, EA tools, and a fitting of EA into the overall governance process. Continue toward organizational maturity and development in the areas of strategy, project/portfolio management, information architecture, business process architecture, and technology architecture.

ITIL Process Management: The overall need is to support the 11 roles established by the IS02000 Construction Project, which have reached the commissioning or operations phase in the ISO 20000 project. The primary resources needs are effort and training.

Progress Indicators

Enterprise architecture: EA charter and documented processes, EA work products, progress in EA maturity (as measured by outside EA maturity metrics).

ITIL Process Management: Indicated in the metrics for each individual management activity.