

Fermilab Computing
Strategic Plan
FY2016 - FY2020

A Technology Vision for 2020

 **Fermilab**

Table of Contents

1. EXECUTIVE SUMMARY	3
2. MISSION AND VISION STATEMENT	5
COMPUTING MISSION	5
COMPUTING VISION	5
3. PREFACE	6
4. INTRODUCTION	7
5. GUIDING PRINCIPLES	8
6. WHERE ARE WE TODAY?	9
CORE COMPUTING DIVISION TODAY	9
SCIENTIFIC COMPUTING DIVISION TODAY	10
GOVERNANCE AND INFORMATION SECURITY TODAY	12
7. WHAT IS DRIVING THE NEED TO ADVANCE FERMILAB'S COMPUTING TECHNOLOGY?	13
8. WHAT IS OUR VISION FOR 2020?	14
CORE COMPUTING DIVISION VISION	18
SCIENTIFIC COMPUTING DIVISION VISION	22
GOVERNANCE AND INFORMATION SECURITY VISION	27
APPENDIX A: DEVELOPING THE COMPUTING STRATEGY	33

1. Executive Summary

Nearly two centuries ago, the English chemist Humphrey Davy wrote, "Nothing tends so much to the advancement of knowledge as the application of a new instrument. The native intellectual powers of men in different times are not so much the causes of the different success of their labors, as the peculiar nature of the means and artificial resources in their possession." Davy's observation that advantage accrues to those who have the most powerful scientific tools is no less true today.

"Whether describing the advantages of high-energy particle accelerators (such as the Large Hadron Collider) or powerful astronomy instruments (which yield insights into the universe's expansion and dark energy), ever-more powerful scientific instruments continually advance knowledge. Each such scientific instrument, as well as a host of others, is critically dependent on computing for sensor control, data processing, international collaboration, and access."*

One of the cornerstones of Fermilab's world-renowned scientific research is its Computing leadership and resources, which enable the broad particle physics community to deliver scientific results in a timely fashion. Computing's staff are internationally recognized experts in programming languages, high-performance computing and networking, distributed-computing infrastructure, peta-scale scientific data management and physics simulation. Computing's services support large-scale computing, data-management and data-analysis facilities for the CMS experiment and the LHC Physics Center; computational cosmology; the Digital Sky Survey; the Dark Energy Survey; neutrino and rare-process experiments; and computational cosmology. Furthermore, Computing is a leader in the Open Science Grid multi-disciplinary distributed computing infrastructure and hosts major computing projects such as the CMS Tier-1 center and Lattice QCD computing.

In addition to continuing to enable Fermilab's scientific research and world-renowned projects, Computing must also help drive the laboratory to better leverage emerging digital technologies. In particular, with 1,800 full-time employees and an additional 2,000 facility users representing more than 40 different countries and 120 institutions, our end users demand equally innovative world-class business solutions, architecture and security that are agile, flexible, easy to use and cost effective. The strategy document put forth here outlines our multi-year plan to refocus our computing investments to better meet the needs of the laboratory and significantly improve the end-user experience.

This document will, in non-technical terms, identify the main themes, technologies and philosophies we will employ in order to deliver on mission outcomes. This plan is not an overhaul of Computing; instead, it is a change in emphasis as we shift from locally developed computing solutions to modern cloud-based applications, where applicable, and continue to drive innovative, agile scientific computing solutions with broad applicability to our program.

At the core of this strategy are eight underlying themes.

- We will continue to innovate our scientific solutions.
- We will leverage industry best practices to deliver and support our solutions in order to ensure the overall consistency, efficiency and speed of our performance.
- We will significantly increase our cybersecurity program to ensure we are excellent custodians of the digital data we are entrusted with.
- We will leverage industry-leading partners and vendors and adopt mature pre-packaged computing solutions that have been proven to be successful to these other organizations.
- We will strive to simplify the end-user experience by providing tools, interfaces and mobility applications so that our employees' work experiences are similar to what they are used to at home.
- We will continue to modernize our systems and infrastructure to better support the laboratory's mission.
- We will provide a work environment that makes it easy to attract top talent and retain the talent that we employ.
- We will build stronger partnerships with industry as well as the other areas within the Office of Science with an emphasis on ASCR.

In summary, this strategy outlines the roadmap well into the next decade and beyond and establishes a holistic digital environment that can provide maximum responsiveness and flexibility with secure and reliable access to the data when and where it's needed, on any device.

* Daniel A. Reed, Jack Dongarra, Communications of the ACM, Vol. 58 No. 7, Pages 56-68

2. Mission and Vision Statement



Computing Mission

Deliver timely, innovative computing solutions and services that enable Fermilab to achieve its scientific mission, efficiently execute the business of the laboratory and provides a modern user experience

Computing Vision

Provide secure and reliable access to information and services, any time, any place, on any device, in full collaboration with our scientific community

3. Preface

April 10, 2016

Dear Nigel,

In the last decade, “flat budgets” have taken their toll on our ability to sustain our technological advantage over time, resulting in a computing portfolio that is complex, aging and costly to maintain. In the business world, significant changes have occurred in the way successful enterprises deliver enabling information technology, morphing from a model of highly custom applications development to commercial package applications that are integrated to the needs of the enterprise. In addition, industry is now moving towards cloud-based services, further changing the model for IT delivery.

In the scientific world, compute engines are moving more and more into accelerator-based processors with increasing parallelization. HEP’s software stack must adapt to these emerging technologies in order to conduct science in an efficient and cost-effective way. As these systems get more complex, we as a laboratory must provide “portals” to enable our researchers to efficiently interact with the variety of compute engines (ASCR HPC, commercial clouds, U.S. Grid systems, etc.) while hiding the complexity from them.

Today, Fermilab Computing is spending most of its effort just keeping things operating well. We must continue to take steps to modernize our systems and approaches and consolidate applications to a manageable number to gain efficiencies. Streamlining operations will free up the necessary resources that will allow us to keep pace with technology and provide both the modern systems and the type of interfaces to those systems that people have come to expect outside of work.

I commissioned this report to provide to both the people in Computing and the laboratory’s senior leadership our vision as to how we want to evolve our enterprise to better serve the laboratory mission and to tackle the issues above head-on. By following this plan, Computing will be well positioned to handle the needs of the high luminosity LHC and the DUNE experiment and provide a modern IT environment and infrastructure that will carry Fermilab well into its second 50 years.

On behalf of the Computing organization,



Rob Roser
CIO

4. Introduction

The overall objective of the computing strategy is to increase the laboratory's effectiveness and efficiency through the proper application of technology.

In the past decade, the face of Fermilab Computing has changed dramatically. In the information technology (IT) world, the Core Computing Division (CCD) was established, consolidating most of the IT organizations from the various support organizations and providing a single leadership and focus. At that time, what had traditionally been called "Computing Division" was split into 3 organizations: CCD, Scientific Computing Division (SCD) and the Office of the CIO (OCIO). This new structure is better equipped to manage the enterprise of computing efficiently.

Furthermore, Computing as an enterprise has moved away from a model based on managing applications and computers to a services model based on the Information Technology Infrastructure Library (ITIL). ITIL is a widely accepted approach to IT service management, which has been adopted by individuals and organizations. ITIL provides a cohesive set of best practices, drawn from the public and private sectors across the world. A services model enables resources to be leveraged across an entire organization, resulting in lower costs with agreed-upon levels of customer service. Over the past 5 years we have committed to this approach and became ISO20k-certified in 2012 (and again in 2015), acknowledging our success in deploying a services model to manage our computing enterprise.

The consolidation of IT across the laboratory resulted in considerable savings and now provides a consistent way of managing IT. However, there is still significant work remaining to capitalize on this consolidation. There remain in excess of 150 business applications that require significant support as well as a wide range of technologies to execute these applications. With a proper strategy and execution, it is possible to reduce this number dramatically and standardize our technologies to simplify operations, all while improving the customer experience.

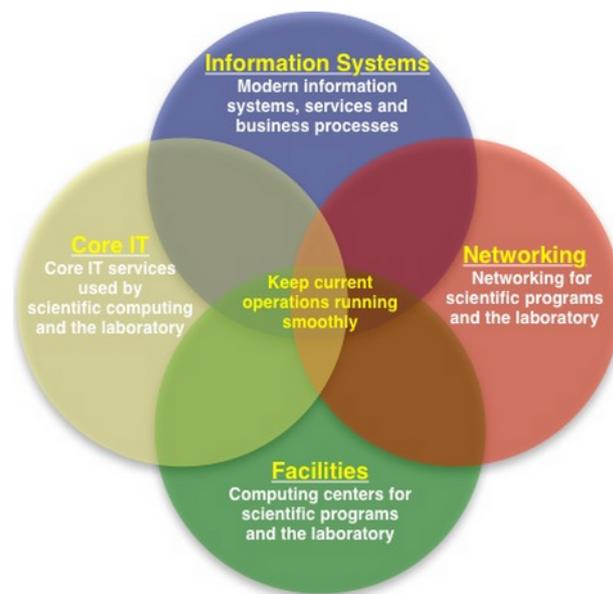
5. Guiding Principles

The following principles clearly and distinctly provide direction for Fermilab Computing's strategy going forward.

- Enable mission and resiliency while managing risk
- Make doing the right thing the easy choice
- Focus on the end-user experience and leverage, where possible, mobile devices to empower people to be productive wherever they are
- Use technology to enhance our competitive advantage for researchers
- Manage information as a critical laboratory asset
- Instill the right level of cyber security to protect our information and reputation while maintaining the laboratory's open science mission
- Recognize science is collaborative and never stops
- Simplify and reduce the overall legacy application portfolio
- Maintain an applications architecture strategy of preferring commercial and/or cloud solutions over custom developments
- Focus on our greatest assets—our employees

6. Where are we today?

Core Computing Division Today



Core Computing Division, established in 2009, is the IT arm of the laboratory. The overarching goal of this approximately 120-person organization is to deploy and support excellent and innovative cost-effective computing solutions and services that support Fermilab's mission. As such, the primary goals of the division are to:

- Keep current operations running smoothly
- Provide reliable and efficient computing centers for scientific programs and the laboratory
- Provide the network infrastructure and services for scientific programs and the laboratory
- Sustain and improve core IT services used by scientific computing and the laboratory

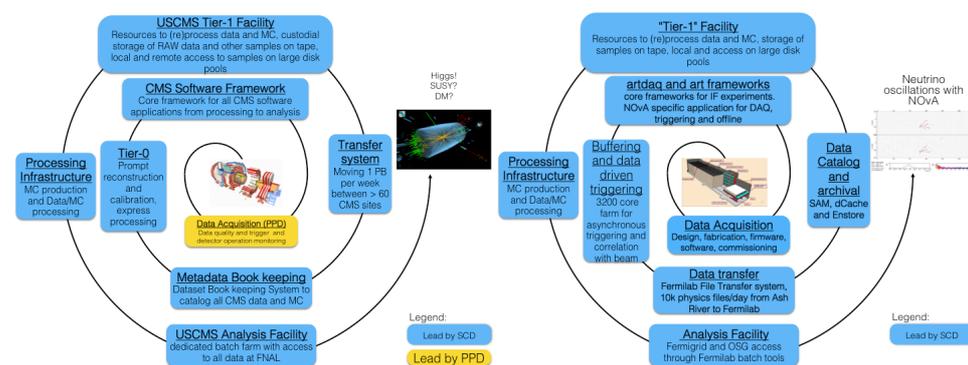
Modernize the laboratory's information systems and services and work to improve business processes

The division is comprised of four main areas:

- **Facilities:** responsible for the laboratory's data centers and act as landlord for our buildings
- **Networking and Communications Services:** provides network connectivity to everyone in the laboratory, provide and administer email services and conduct network research
- **Core IT:** encompasses cybersecurity, authorization/authentication services; enterprise services such as virtual servers; global disk storage; web and SharePoint; general operating system support for the laboratory's business systems; Windows and Mac desktop services including updates and antivirus support; the laboratory's Service Desk, which supports employees and users; and licensing/contracts for all the laboratory's software (Microsoft, Oracle, Workday, etc.)
- **Information Systems:** provides all the laboratory's business financial, property, ESH&Q applications and Records Management, Technical Publications, History and Archives, Printing/Duplicating Services and Library Services

The main challenge of Core Computing Division is to maintain our diverse portfolio of hardware and applications while finding the resources to modernize both the hardware and application stacks. Since CCD's budget is almost entirely made up of fixed maintenance and license costs and more than half of its staff are responsible for ongoing application and user support, it is difficult to make progress in this area. Extra people and extra dollars go a long way in helping.

Scientific Computing Division Today



The Scientific Computing Division, currently about 140 people strong, provides scientific software solutions and deploys and operates scientific computing facilities in support of the Fermilab program. Its staff works closely with the other laboratory organizations to deliver world-class computing services, operations and software engineering support to Fermilab-based experiments, CMS and the high-energy physics community at large. It also interacts with the experimental community to determine capacity, availability and capability requirements and to procure resources as necessary. Furthermore, it engages in R&D activities required to maintain or advance capabilities necessary for the success of Fermilab's future physics program.

The Scientific Computing Division activities are aligned along three functional areas, crosscutting all of Fermilab's scientific programs, which are briefly described below.

- Development, Integration and Research:** enhances or improves aspects of computing for experiments to positively impact their ability to arrive at physics results. The division has many such activities encompassing the full stack of software for HEP. Most notably, it provides frameworks and tools for physics analysis and reconstruction, real-time software and engineering for DAQ, simulation packages and applications and workflow and production tools. In addition, the division collectively houses an enormous store of knowledge about and expertise in scientific computing, and experiments depend on our consultation for developing their specific computing solutions. To achieve our goals efficiently and cost effectively, we leverage expertise and tools from all scientific programs to create a coherent computing

“ecosystem” for all our users. We also have strong ties with other institutions and participate or lead collaborative R&D efforts on new architectures in order to transform our major toolkits. We especially seek collaborations with the DOE Office of Advanced Scientific Computing Research (ASCR) to leverage expertise and maximize opportunities of participation to relevant R&D programs.

- **Scientific Facility:** provides resources for data processing, simulation, storage, and analysis; connections to the data acquisition systems of Fermilab based experiments; and the networks to interconnect these both locally and globally. In addition, the facility provides the services that allow users to execute workflows on facility resources and manage the archived data. (Facility utilities—power, cooling, etc. and infrastructure, such as networking, are supported by the Core Computing Division.)

The core of the data storage facility consists of seven 10,000-slot robotic tape libraries. With the current tape technology (T10000-D), the tape library complex has a potential capacity of nearly an exabyte. Access to the data on tape relies on a complex of over 100 enterprise-class tape drives. Average data transfer to and from tape is currently over 1.5 petabytes per month (50 terabytes per day) and reached 6 petabytes per month during the previous LHC running period. A robust disk system serves as the front-end to the tape system and is supported by the data management services and tools necessary for efficient local and global distribution of content.

The staff that operates the facility possess a wide range of skills including detailed knowledge of processor, tape and disk storage hardware architectures, network architectures, the Linux operating system (Fermilab is home to Scientific Linux) and computer security. The skills of the facility personnel are as valuable as the physical resources.

- **Science Operations and Workflows:** provides the essential services needed for experiments to access very large data sets across distributed computing infrastructures, both local and remote, in a robust and scalable way. Staff operates services needed for distributed computing, such as batch, grid, authentication, authorization and virtualization, as well as services for cloud and high performance computing. In addition, staff develop and operate the database applications required by the scientific collaborations, provide and operate services that enable collaboration such as code repositories and wikis and build and test services for applications.

The Scientific Computing Division participates in, supports and enables every aspect of the Fermilab scientific program. It constantly faces the challenge of optimizing deployment of resources, both to support operations of the current experiments and projects and to prepare the scientific software, facility and services for the future program of the laboratory. This task becomes even more challenging by the evolution of computing technologies driven by industry requirements, and, ultimately, limited by declining budgets.

Governance and Information Security today

Computing is a service-oriented organization that demonstrates excellence through effective and efficient management and information security practices supported by governance processes including enterprise architecture, service management, project management, financial management, quality assurance and organizational change. These processes, policies, roadmaps and plans help ensure that computing investments are aligned with Fermilab's strategic plan, goals and objectives. We strive to ensure we are able to answer the following questions:

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

Governance and information security reside in the Office of the CIO so that adherence and management is consistent for all of Computing. As a key component of the Information Management System, which is part of the Fermilab Contractor Assurance System (CAS), there are now many bodies that meet regularly and many processes that have labwide purview and are continuously operating under the authority of but below the level of the IT Executive Council:

- Computer Security Board
- IT Policy Board
- Enterprise Architecture Board (internal to Computing)
- Service Management bodies and meetings (e.g., Change Authorization Board)
- Information Systems Portfolio Management Team
- Scientific Computing Portfolio Management Process

We define success as ensuring that investments are made in a manner that maximizes the use of limited resources to deliver computing capabilities that are aligned with the strategic goals of the laboratory.

7. What is driving the need to advance Fermilab's computing technology?

The lines between the physical and digital business are blurring. Successful businesses are leveraging technology in new ways to increase market share and provide a competitive advantage. Fermilab is no different—Computing has become one of the core enablers of science. Modern experiments demand more innovative computing solutions than ever before. Within the IT world, technology is also changing rapidly. Custom solutions have become the dinosaur of the computing ecosystem, and companies are turning to cloud-based applications that not only streamline the development process but also reduce the total cost of ownership. Thus, Fermilab needs to have a strategy in place that will leverage digital solutions to its advantage or risks being left behind.

At the same time, users' expectations are changing. With the increased commercialization of computing services (Google, Amazon, etc.) readily available, we need to rethink our digital position. People want and expect the same type of experience in their work environment that they already have as part of their home consumer experience. This means we must not only provide systems that are more intuitive to use, but on a variety of platforms. We must enable workers to be productive, not only at their console, but anywhere, by optimizing use via mobile devices.

As we have seen recently in the news, personal identifiable information (PII) has been compromised in the commercial world (Target, Home Depot, etc.) and in the U.S. government (OPM Breach). The emerging cyber security threats mandate significant changes in our cyber security controls and processes. We can no longer simply run an open network for all Fermilab business as we have in the past. Architecting a cyber security program that protects our digital assets and reputation while maintaining an open scientific network will be a challenge that will require substantial effort and resources.

All of these changes demand that we create a strategic plan that presents a vision of where we want to be as an organization over the next 5 years with an initial 3-year focus. This plan will not only help prioritize the Computing spending of our current budgets, but will also provide clarity into how incremental additional investments would provide tangible outcomes.

8. What is our vision for 2020?



Once again, at the core of this strategy are eight underlying themes that comprise our overall Computing strategy:

- We will continue to innovate our scientific solutions.
- We will leverage industry best practices to deliver and support our solutions in order to ensure the overall consistency, efficiency and speed of our performance.
- We will significantly increase our cybersecurity program to ensure we are excellent custodians of the digital data we are entrusted with.
- We will leverage industry-leading partners and vendors and adopt mature pre-packaged computing solutions that have been proven to be successful to these other organizations.
- We will strive to simplify the end-user experience by providing tools, interfaces and mobility applications so that our employees' work experiences are similar to what they are used to at home.
- We will continue to modernize our systems and infrastructure to better support the laboratory's mission.
- We will provide a work environment that makes it easy to attract top talent and retain the talent that we employ.
- We will build stronger partnerships with industry as well as the other areas within the Office of Science with an emphasis on ASCR.

Following these 8 themes, we will be able to improve on the delivery of business and scientific solutions in a timely and cost-effective manner.

There are a few initiatives going forward that are larger than any division and thus listed below. Following these, each of the divisions presents their detailed strategic plans.

A New Core Capability at Fermilab

This year, for the first time, Fermilab added to its list of core capability *Advanced Computer Science, Visualization, and Data*. Fermilab's expertise in Advanced Computer Science, Visualization, and Data enables discovery in particle physics. This core capability complements theory and experiment as a means to increase scientific knowledge through data collection, storage, reconstruction and analysis, as well as through scientific simulations. Fermilab has a remarkable history of developing, delivering and deploying computing technologies for the scientific community. This capability should now give us a "full seat at the table" when we are looking to partner with ASCR (Office of Scientific Computing Research) on R+D issues of interest to us. Thus, as we go forward, we as a laboratory will need to continue to invest in this capability so that we maintain it and use it to our fullest advantage.

Next Generation Computing Facility





To move forward as a modern lab, we need a modern data center. Thus, we have embarked on the design of just such a building. This Next Generation Computing Center (NGCC) is envisioned as supporting the next generation of computing at Fermilab and transforming the laboratory's core capability in advanced computer science, visualization and data. The NGCC will support the major initiatives outlined below, including HEP Cloud and the Active Archive Facility. Supporting these major initiatives requires ongoing improvements in laboratory infrastructure to keep pace with rapidly developing technologies and scientific demands.

Three critical components have been identified for the NGCC. The first is big data storage. This component will provide large disk storage, tape storage, and networking and will be located underground to protect systems and data from natural disasters. The second component will provide space needed for computing systems. This will include space for scientific computing and business computing systems that will be deployed using open air, state-of-the-art, green cooling systems. The third component of the NGCC will provide functionality that is currently unavailable at Fermilab's Feynman Computing Center by providing collaboration space and additional meeting space. Computer scientists and their technical teams, as well as collaborating scientists from around the world, will benefit from the availability of small group workspaces for collaboration, medium-sized meeting rooms and a 300-seat auditorium that is large enough to accommodate division-wide meetings.

The building is planned to be sited right across from Feynman Computing Center, where CHL resides now. An architectural sketch of the proposed building is at the top of this section. We are currently seeking private funding for this building in order to expedite a construction start.

Increased Collaboration

One of our goals moving forward will be to leverage our computing abilities and capabilities in order to impact the broader scientific community. Through HEP Cloud, for instance, scientists from across the Office of Science will be able to more effectively use computational resources by making it simpler for them to compute. Fermilab will leverage its data storage facility to store and serve data for the broader scientific community, and we will establish a revolutionary terabit network between Fermilab and Argonne National Laboratory to improve scientific collaboration. This network will better leverage each laboratory's computational core competencies. Finally, through the National Computing Strategic Initiative (NCSI), we will partner with ASCR on the Exascale initiative to ensure that this next generation class of high performance computing offers maximal value to the HEP community. Furthermore, it's time to leverage some of our intellectual property and partner with industry to provide new revenue streams for the lab.

Core Computing Division Vision

The thrust of Core Computing Division's efforts can be summed up by three simple goals:

- Improve our users' experience both at the laboratory and when they use our services remotely
- Modernize and automate the laboratory's business processes
- Enable mobility—provide access to information and services any time, any place, on any device

Our first goal is **to improve the user experience**, making it easy for users, not only when they start work at the laboratory (onboarding), but also as they do their day-to-day work and need services and support, so they can accomplish what they come to the laboratory to do. We also want to improve our employees' experience, and much of this happens naturally as a by-product of improving the users' experience. We can do this by taking the following actions:

- Provide a simple, user-oriented interface to the laboratory's service management system, ServiceNow. The current user interface, particularly for non-IT-experts, is overly complicated and could be simplified. Argonne National Laboratory has done something similar and has had good results.
- Improve the user onboarding process. Our current procedures require users and staff to perform actions in multiple applications (e.g., FermiWorks and ServiceNow), and there are multiple places where manual intervention is required. This process is error-prone and confusing to users, and it does not allow us to easily accommodate user preferences for account names. Onboarding can be streamlined and automated with the implementation of Identity Management. In addition, we need to provide better tools for new users to help them find laboratory resources.
- Move ahead with enterprise search. Finding the information you need is straightforward when you know what "vault" the information is kept in. DocDB, SharePoint, the central Fermilab Web, TeamCenter and other "vaults" all contain information that users and employees need to find. One solution, a portfolio project called Document Management, was halted about 5 years ago by the COO due to insufficient effort to start. The first step in document management is to enable searching through all possible vaults at the laboratory and deliver context-aware, relevant results regardless of where the documents reside—this is called enterprise search. Preliminary work on enterprise search has begun, but it is stalled due to lack of effort and appropriate hardware.
- Networking in the Fermilab Village. Today, the Village uses rather basic network capabilities provided by the laboratory's network department. There is a proposal on the table from Comcast to provide Village residents with internet and cable TV service that would separate them from the

laboratory's business network and improve residents' quality of life.

- Provide unified communications.
 - Our science users are primarily driven by Google services. We therefore need to develop better file sharing integration here.
 - Video service is still being requested. Some experiments are considering purchasing it on their own, but would like us to support it for them. A coherent labwide plan would be better. Virtual conference rooms and desk meetings are increasingly popular and should be considered as part of this plan.
 - Presence awareness, knowing who is online, should be available for email, messaging and video.

Our second goal is **to modernize the business processes of the laboratory.**

Automating items wherever possible, eliminating paper and removing redundancy are all hallmarks of modernizing.

- There are several financial projects that are under study by Computing and Finance that require funding to start:
 - Our payroll system, PeopleSoft, will enter end-of-life status at the end of 2017. We must either replace it with a system more compatible with FermiWorks or deploy Payroll in the FermiWorks system.
 - Everyone is familiar with buying items on Amazon. Procurement wants to deploy a similar system at the laboratory. The system would allow employees to select predefined common items from predetermined vendors via a web interface. There would be electronic approval, a purchase order would be automatically cut to the vendor and automatic payment would be made after the item was received. This would significantly cut down on staff effort and would also streamline the procurement process. Our planning indicates this project should be done in phases, with the first phases consisting of about 500 employees and 30 vendors.
 - Budget and planning: Fermilab currently completes its budget and planning activities using multiple, disconnected processes, tools and systems. This creates a lack of consistency that makes aggregating the data for analysis and planning complicated and time consuming, and it hinders required interactions between organizational units and consistent reporting. The laboratory desires a single, modern system and common processes that will result in budget and planning activities that are streamlined, more robust and easier to use. This project is underway.
 - Our prime contract is being changed and will require us to collect more information from small businesses. The DOE calls this project "M&O Subcontract Reporting Capability (MOSRC)." We have understood the requirements as explained to us and believe it will

cost between \$100-200K to implement. The DOE is requiring a fully automated system by November 2016, and we estimate it will take 7 to 8 months to deploy once started.

- We need to modernize our travel system. Specifically, we need an enterprise-level travel system with end-to-end electronic integrations for users and for our back-end financial systems.
- We need a talent and recruitment enhancement for our FermiWorks implementation that makes it easier to recruit and retain scientific talent.
- We are operating our phone system on an already end-of-life-status 5ESS switch. We have a contract with AT&T until 2017 and option years for 2 more years. We need this replaced before the end of the end of the contract, since there is a reasonable chance AT&T will not renew it again. A year-long study has occurred, and many plans were considered. The result is to continue using our copper infrastructure via digital gateways and move to VoIP for new construction. After implementation, the laboratory will have a modern phone infrastructure, but unfortunately, there will essentially be no overall cost savings to Fermilab. Estimated cost is \$2M over 2 years.
- We have been working on modernizing our public Web presence with the Office of Communication. Known as the Web Modernization project, this effort's goal is to make it easy for both new and experienced users to find and access laboratory resources via the Web.
- Outside of the laboratory, the work experience is becoming more and more mobile. But many of our resources are reachable only via archaic systems that do not function well for people using mobile devices or when accessed from outside the laboratory network. We need to evolve and become more modern as well. "Any device, anytime, anywhere" is an accurate representation of where we want to be. This is a very large project. First steps would be to improve the mobile experience for timecards, entering purchase requisitions and viewing laboratory Web pages effectively on tablets and phones.
- We must automate electronic approval for phone bills. We send out two boxes of paper bills each month for bills starting at around \$30 per phone.
- We must untangle the intertwining of applications and eliminate old applications. Workday was supposed to be the catalyst for this, but to be successful, we need to take the next step and construct a **data service hub** that will provide commonly used business data to all consumers independent of the source. This will permit us to upgrade or replace individual legacy applications without forcing other applications to change where and how they obtain the data.

- Many of our business processes span multiple applications. To automate these, we need a **modern workflow engine** that can interface to multiple applications.
- We must help our business and scientific users obtain timely data analytics, KPIs and dashboards. To do this, we need a **data warehouse** to store historical information for reporting purposes. This will also let us dismantle some really old systems we keep around just to provide historical data for reporting.
- We must accelerate the work begun on identity management. The current work has concentrated on account lifecycle and will soon go into production. The real work of identity management is managing roles. When one becomes a group leader at the laboratory, for example, all responsibilities associated with that role should be enabled for that person. Equally as important, the roles should be taken away when one no longer is a group leader. This project has significant connections to many business processes at the laboratory and will need to evolve as applications are introduced or upgraded.

Scientific Computing Division Vision

As we execute the plan for U.S. High Energy Physics, our scientific computing program must face some major challenges:

- As computing architectures evolve, we must change the paradigms for how we construct our algorithms, write our codes and organize our analysis workflows.
- We must continue making effective use of high-performance computing (HPC) for large computational problems and increase the number of applications that are able to do so. This will be essential as computing in the U.S. enters the exascale era and more powerful DOE/ASCR supercomputers become available.
- New technologies, such as cloud computing, provide new, potentially cost-effective options for deploying computing resources in the future, requiring development of new services for on-demand reliable resource allocation.
- We must ensure that our globally distributed community of users and collaborators have ubiquitous, scalable, simple access to and use of up-to-date software, data and computing for data acquisition, simulation, processing and analysis.

The Scientific Computing Division strategic initiatives for the next few years are focused to address these challenges. They are motivated by our mission to provide, lead and innovate computing, software tools and expertise for the Fermilab scientific program. They are driven by our principles to ensure coherence and effectiveness by encouraging sharing of facilities and common approaches to service and software solutions. The major initiatives for the next few years are:

- **Delivering a new paradigm for HEP computing facilities** through a single managed portal to a dynamic, heterogeneous set of computing and storage resources. We are developing and will deploy a system that provides cost-effective and efficient “elastic” deployment of resources while allowing us to enable, with cost recovery, utilization of Fermilab computing assets by external programs:
 - The HEP Cloud Facility
 - The Active Archival Facility
- **Conducting R&D for advanced computing and computing fabric solutions** for the facilities of the future focusing on the ever increasing speed of the network links and the management needed for their use by our globally, distributed, data-intensive science:
 - Networking R&D for Software (SDN) and Name Defined Networks (NDN) in partnership with ASCR institutions and the community

- The Terabit network link between Fermilab and Argonne National Laboratory facilities
- **Performing common experiment analysis and DAQ solutions** that provide new levels of support to ongoing and planned Fermilab experiments and projects. Challenges to be met include accommodating the different timelines and priorities and ensuring full partnerships with and leveraging of collaborative efforts:
 - The LArSoft Phase II analysis ecosystem for LArTPC detectors
 - Extending the artdaq DAQ solution for Fermilab experiments
- **Enabling effective use of emerging new computing architectures** that demand vectorization and multithreading. This includes evolving our scientific software toolkits and developing the necessary new tools and algorithms to utilize and take advantage of the exascale era of U.S. computing including:
 - Leveraging partnerships and collaborations with ASCR institutions
 - Seeding and coordinating cross-cut development efforts across HEP experimental programs
 - Creating a Center for Numerical Computation in HEP
- **Improvement in the effectiveness, timeliness and productivity of HEP analyses**—experiment and collaboration—via the adoption and deployment of new paradigms, techniques and technologies through:
 - Improving processing operations through automation and intelligence
 - Adopting Big Data technologies

More detail is given for selected initiatives below:

- **HEP Cloud Facility:** As an example of the challenges facing us, during the next decade the increased event rates and complexity of the HL-LHC will push computing needs to approximately 100 times more than current HEP capabilities can handle, generating exabytes of data! HEP must plan now on how to efficiently and cost effectively process and analyze these vast amounts of new data. The industry trend is to use cloud services to reduce costs of provisioning and operating, provide redundancy and fault-tolerance, rapidly expand and contract resources (elasticity) and pay only for the resources used. The HEP Cloud Facility concept is a path to this evolution, providing “complete solutions” transparently to all users with agreed-upon levels of service and routing user workflows to local (“owned”) or remote (“rental”) resources based on efficiency, cost, workflow requirements and target compute engine policies. Achieving this goal requires significant development in several new key areas of work:
 - Understanding, procurement and use of commercial cloud resources from Amazon, Google and Microsoft (reduce risk and increase cost effectiveness through competition)
 - Developing decision-making algorithms to determine what the most cost-effective resources to use at any given time are for particular user workflows, processing and analysis

- Providing the needed data handling, transfer and access protocols and technologies to ensure seamless provisioning of data for the computations to be executed
 - Extending the current distributed computing “administrative services” (e.g., accounting, monitoring, error reporting) to work seamlessly in the new environment
 - Development of user-friendly interfaces for all services provided
- **The Active Archival Facility (AAF)** is developed as a component of the HEP Cloud Facility to enable, with cost recovery, external scientific collaborations access to the Fermilab tape data storage resources. The AAF is seen as a valuable service to scientific efforts that are otherwise without access to a facility that provides massive capacity, yet ease of access, through well-exercised and -optimized data transfer tools. The key areas necessary to fully accomplish this goal are the development of:
 - Monitoring and accounting tools needed to provide necessary services to the storage facility operational support of many distinct internal customers
 - Enhanced data management tools and catalogs for easy data access
 - Access strategies for non-HEP customers that perform within their corresponding cybersecurity paradigms
- **R&D into advanced networks** utilizing Software Defined Networks (SDN) and Named Data Networking (NDN) technologies. Research projects with ASCR and partnerships with ESnet will allow us to deploy and operate these networks (including those needed by the Fermilab HEP Cloud Facility). We will:
 - Deliver seamless provisioning, usage and support for networking fabrics that deliver the most value to our users’ needs, wherever the users reside, whatever their data transfer and usage patterns
 - Advance the systems to use more modern technologies as they become available
 - Continue the cross-computing (SCD, CCD) research and development efforts and collaborations with ASCR and ESnet and pursue new opportunities through projects that are funded by ASCR
- **Common analysis ecosystem for LArTPC detectors:** The Snowmass 2013 community report and the P5 panel recommend increasing collaborative and shared approaches to HEP experiment software and computing to address resource constraints, cross-training and sustainability of the current and future workforce and agility in response to new hardware and software technologies. The neutrino experiments based on the Liquid Argon Time Projection Chamber (LArTPC) detector technology have common analysis needs and requirements:
 - We will establish a sustainable data simulation, reconstruction and analysis ecosystem for each experiment to integrate, deploy and operate as necessary to extract physics results from the LArTPC detectors they use.

- We will facilitate a sustained, productive collaboration across all experiments benefiting from the ecosystem, with ongoing development to address new requirements, enhanced algorithms, new technologies and new computing hardware.
- We will cover all aspects of the system needed to deliver validated physics results including: human interactions; collaboration and responsibilities; computing hardware and software; scientific workflows, processes, interfaces and communications.
- **Extending the artdaq DAQ solution** for Fermilab experiments. These currently include MicroBooNe, NOvA, DUNE 35-ton, Mu2e, ProtoDUNE, Icarus and SBN.
 - Increase the use of commercially available off-the-shelf components (IoTDAQ)
 - Make the common DAQ an open-source available toolkit
 - Integrate new intelligent read-out controllers, hardware managers and networking to improve performance and cost effectiveness
- **Establish common toolkits** through R&D projects to increase the availability, adoption and use of common software frameworks, promising technologies and methods and new hardware architectures for scientific computing solutions including for the challenging needs of the US CMS and DUNE experiments:
 - Short-term deliverables for common modules and approaches (computing patterns and methods) across the art and CMSSW frameworks to demonstrate the benefit of such an approach
 - Continuing to integrate development and support teams in cross-cutting development and support projects for U.S. CMS, neutrino experiments, astrophysics and the g-2 and Mu2e experiments
- **Integrating Big Data technologies** and improving the effectiveness, timeliness and productivity of experiment and collaboration analyses. Initial work includes NoVA, U.S. CMS, g-2, Mu2e and DES. Investigate, adopt, integrate and then deploy solutions using No SQL databases such as CouchDB data management solutions such as Hadoop and public domain analysis toolkits such as R to the benefit of the end users. Short-term goals include:
 - NoVA Library Event Matching (LEM) algorithms depending on No SQL database and advanced event-selection techniques in production use
 - Documented analysis use cases that can benefit from Big Data approaches
 - Execute prototypes in collaboration with U.S. CMS, g-2, and MINVERvA to demonstrate value in a well-defined timeline
- **A Center for Numerical Computation** evolving the current, well-regarded Lattice QCD, Pythia and other theoretical and simulation activities to include applied mathematics and computer science colleagues from ANL, UCLA and other groups through:

- Continuing to build collaborations with ANL and LBNL computer science and applied mathematics communities, as well as university groups such as at UCLA
- Participating as part of the HEPCFE and HSF collaborations
- Ensuring HEP and Fermilab are a full partner in the DOE ASCR high-performance and applications computing roadmaps

Governance and Information Security Vision

Computing will continue to emphasize continuous improvement to ensure we are following best practices, providing the right services at the right time and building out capabilities to meet strategic objectives. There four main areas in which these improvements will be focused:

- Strategic alignment
- Service delivery
- Project delivery
- Information security

In addition, it is crucial that we recognize the value of our people and provide a competitive work environment that is stimulating and rewarding.

Strategic Alignment: An effective Computing organization consists of many elements that work in concert to ensure that actions and activities are performed in a manner that protects the laboratory and optimizes the use of resources.

Given current budget and resource constraints, it is important to make informed decisions regarding the laboratory's investments to ensure that investments and decisions are aligned with the laboratory's strategic vision. In addition to rationalizing the current software-application portfolio to eliminate redundancy, streamline operations and reduce costs, we will focus on the development of strategic business and technology roadmaps. These roadmaps will be developed with laboratory functional area leaders and will serve as strategic blueprints that define the initiatives and actions necessary to meet current and future computing needs. In addition to communicating the planned future direction for Computing, these roadmaps will serve as a basis for improved investment decision-making and future resource skills planning.

While the roadmaps will chart the path forward for developing business capabilities across numerous computing-service domains, sound portfolio-management processes will be required to manage, update and occasionally re-prioritize the numerous initiatives and projects defined on the roadmaps. An effective process for managing the scientific computing hardware and software portfolio is in place and operating well, with refinements made occasionally to improve effectiveness. A portfolio process for business system applications was established and operated for several years, but needs to be transformed to function more effectively and interactively with key stakeholders.

Responsible financial management is key to achieving strategic goals in a cost-effective manner. An effective financial management system is critical to help plan, track, monitor and assess financial performance. In addition, as the technology landscape continues to evolve at an unprecedented pace, we must fully understand the cost of the computing services we deliver. As budget pressures mount, concrete methods of determining the costs of services delivered are necessary to benchmark our delivery costs against commercial offerings. Financial systems will need to be

put into place to demonstrate that services provided by Computing are cost-competitive and of high value or to identify areas in which additional action is needed to improve cost effectiveness.

The following strategic initiatives will improve the efficiency and effectiveness of Computing:

- Develop strategic roadmaps to align the development of capabilities with strategic objectives
- Rationalize the current software-application portfolio to eliminate redundancy, streamline operations and reduce costs
- Improve the project portfolio management (PPM) process to provide governance and prioritization management to the suite of potential and in-flight strategic projects, to ensure alignment with strategic roadmaps and progress toward achieving strategic objectives
- Establish demand management processes to effectively capture computing needs and establish a higher level of involvement with customers on setting priorities and selecting work to be performed
- Develop financial methods to demonstrate that our services are cost effective and of high value, and identify areas requiring corrective action
- Mature the Fermilab Software Quality Assurance program to ensure that software developed or used at Fermilab is reliable and meets or exceeds established requirements and expectations
- Mature the Enterprise Architecture program to ensure we make good technological choices
- Establish expertise in organizational change management to improve the introduction and adoption of new computing capabilities throughout the Fermilab community

Service Delivery: Effective service delivery involves providing the right services that the laboratory needs to deliver science and agility to meet the demands of the customers and to operate efficiently in a cost-effective way. We will continue to promote and enable a centralized service-delivery model. An established service management system is in place and integrated with other Computing processes to manage and oversee the continuous improvement of the Computing service-delivery model and is based on industry best practices and standards: Information Technology Infrastructure Library (ITIL), Information Technology Service Management (ITSM) and the ISO20k standard. Computing strives to get to a state where service providers have maximum visibility into their service operations, customers have maximum visibility into the services provided and Computing has a modern and sustainable method for service delivery to the laboratory.

Numerous activities are underway to continuously improve our service-delivery model. Computing will continue to work on the following initiatives:

- Develop key performance indicators (KPIs) and performance metrics to demonstrate effective delivery of Computing services
- Continue to evolve service operation processes through regular review of captured metrics around event, incident, problem and request metrics and KPIs. Improve execution and visibility to the customers and drive communication channels, both for resolution and fulfillment activities
- Improve customer interfaces to ensure efficient access to services
- Continue to evolve service-design processes through assessment and regular review of established and new services. Fully integrate service-design processes into enterprise architecture, information security and Computing project management methodologies for strategic and tactical/operational project delivery
- Continue to provide visibility and quality-control measures on services deployed into production through service-transition processes. Continue to expand the transparency and communication channels around service projects being built and deployed through tactical/operational projects
- Provide overarching support for service providers through improved interfaces with the service-management tools

Project Delivery: Effective project delivery helps ensure that we make the best use of limited resources to complete projects that support and enhance the scientific program; transform business operations and the working environment; and improve the online experience for employees, users and visitors.

The Computing Project Management Office (PMO) has established and is using a tailored approach to manage enterprise-level computing projects. Following Project Management Body of Knowledge (PMBOK) and other industry best practices, a formal project management methodology is in place to help ensure that projects are completed according to plan and that expected outcomes and value are realized. The goal is to put just enough structure and process in place to help complete projects more effectively without being overly burdensome. A “PM-Lite” methodology has also been developed to aid in the successful delivery of small- to mid-sized projects that benefit from lightweight planning, tracking and reporting processes.

Strategically, we aspire to establish a world-class project management center of excellence recognized for regularly completing projects according to plan and delivering expected results. Effective project management is about more than finishing projects on time and within budget; it is about delivering expected outcomes and value. We will accomplish this by:

- Establishing improved processes for developing realistic and achievable plans
- Executing projects according to plan
- Taking prompt corrective action when necessary to meet target completion dates
- Ensuring that expected outcomes and value are realized

We will continue to leverage existing platforms, such as Microsoft Project, ServiceNow, SharePoint, and Redmine, to support project management activities. We will finalize the Computing Project Planning Reference Guide and develop a process for socializing the guide and its associated procedures to gain broader adoption. We will continue to refine our performance-tracking dashboard and identify ways to improve the effectiveness of progress reporting processes.

More specifically, we will execute the following initiatives:

- Develop a set of key performance indicators (KPIs) and performance metrics that will effectively and regularly assess the effectiveness of our project management approach
- Establish a project governance board that will periodically review the performance of key projects and initiatives and make adjustments as necessary to achieve desired outcomes
- Implement an improved approach for delivering on our commitments that will establish realistic and achievable project plans as well as accountability for completing those plans
- Incorporate an agile software-development approach into the Computing project management suite of tools based on techniques that are already being used successfully within the Scientific Computing Division
- Develop tailored training materials that will be delivered through various methods to socialize and integrate proven project management techniques throughout the Computing organization

Information Security: Information Security is necessary to safeguard enterprise information against unauthorized access, modification, disclosure and use. Due to recent cyber security incidents, the federal government accelerated its efforts to secure the enterprise. The Office of Management and Budget (OMB) defined a set of cyber security sprint action items that all departments and agencies must address. The following outlines Fermilab initiatives:

- Migrate independently managed web servers to a centrally managed system to increase Web server security
- Install upgraded firewalls to improve border protection
- Implement multi-factor authentication for all privileged account access to prevent data loss
- Implement multi-factor authentication for all access to sensitive data, including cloud-based solutions, per DOE mandate, to prevent data loss

The following table summarizes our governance and information security plans for the next three fiscal years. We explain our plans in detail further below.

Domain	FY16	FY17	FY18
Governance	<ul style="list-style-type: none"> Define Computing KPIs <ul style="list-style-type: none"> Performance metrics Service delivery metrics 	<ul style="list-style-type: none"> Compare KPIs and performance metrics against benchmarks 	
Enterprise Architecture	<ul style="list-style-type: none"> Develop strategic roadmaps <ul style="list-style-type: none"> FIN-IT Strategic Plan Refine PMT process 	<ul style="list-style-type: none"> Application Rationalization Develop strategic roadmaps <ul style="list-style-type: none"> TBD 	<ul style="list-style-type: none"> Infrastructure Rationalization Review/refine strategic roadmaps
Service Delivery	<ul style="list-style-type: none"> Improve ServiceNow customer usability Define Service Costs Implement Availability and Continuity processes Continue onboarding of SCD services 	<ul style="list-style-type: none"> Improve ServiceNow internal usability Analyze Service Costs Mature Service Management processes <ul style="list-style-type: none"> TBD TBD Continue service onboarding <ul style="list-style-type: none"> TBD TBD 	<ul style="list-style-type: none"> Optimize Service Costs
Project Delivery	<ul style="list-style-type: none"> Implement Demand Management methodology and process Implement PM (details) methodology (PMLite) 	<ul style="list-style-type: none"> Implement Resource Management & Planning methodology and process 	
Information Security	<ul style="list-style-type: none"> Migrate independently managed web servers to a centrally managed system to increase web server security. Install upgraded firewalls to improve border protection Implement 2-factor authentication for all privileged account access Implement 2-factor authentication for all access to sensitive data including cloud based solutions per DOE mandate 	<ul style="list-style-type: none"> Replace end-of-life intrusion detection systems Acquire and implement a commercial risk management tool to automate and standardize the preparation of risk assessments Implement tools to monitor and prevent data exfiltration (data loss prevention). 	

Work environment: While Computing has literally many tens of millions of dollars of hardware on the floor, its greatest asset is its people. They are the ones that have the creative inspirations, technical knowledge, passion and inspiration to take technology and make it do something. Thus, to be successful over the next five years, we must provide a stimulating and rewarding environment, not only to attract new people, but also to keep the talented group that we do have. To this end, we will execute the following initiatives:

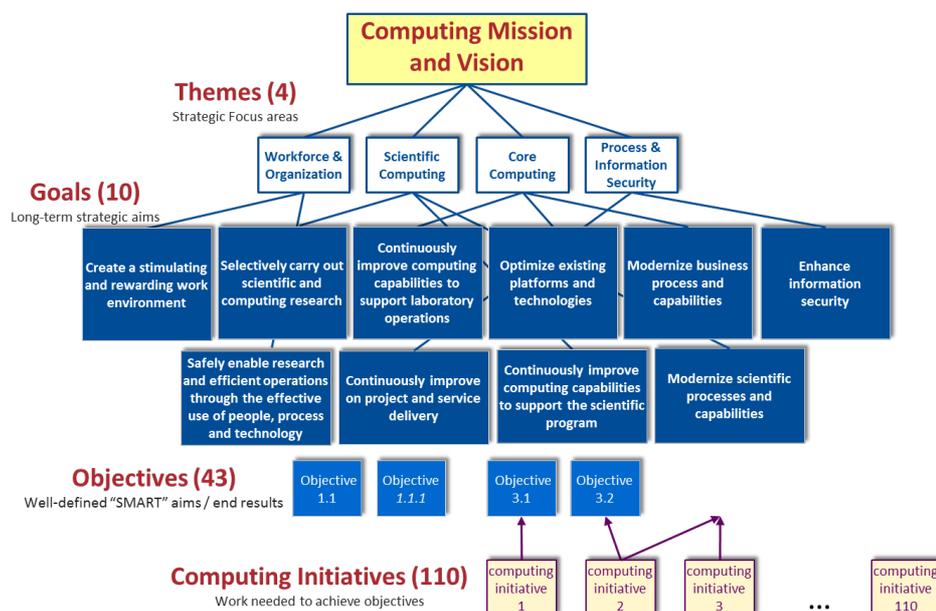
- Compete with industry.** First, we must be competitive with industry in the way they treat their computing professionals. In today's world, that means offering a more virtual environment that empowers people to contribute and be successful independent of their physical location. While we have started that process in recent years with telecommuting, we can and will do much better moving forward. Being competitive also means providing an attractive work environment. Over the next few years, we will improve the aesthetics of our work environment in both Feynman Computing Center and Wilson Hall and build more productive public spaces to foster more cross collaboration.
- Create a greater sense of identity.** Given that Computing folks are distributed at various locations throughout the laboratory, this is not easy to do. Establishing a computing "campus" as part of the laboratory's master plan is a good first step in this process. Toward that end, we have begun to design a showcase data center and meeting center that would reside close to Feynman Computing Center and be integrated into this new campus footprint.

- **Hire and retain minority candidates.** Computing has done well at hiring and retaining minority candidates. However, we can do much better. A diverse workforce in terms of knowledge and cultural background is essential to getting the best solutions proposed and built. In coming years, we will make diversity hiring a priority and find new ways to attract minority candidates to apply for our jobs.
- **Modernize Web presence.** A company's website says a lot about the company and its culture. Fermilab's website is woefully outdated and needs significant attention, not only to update content, but also to present it in a way that makes it simple to use and digest. Not all of Fermilab's business needs to be public. Separating employee information from that intended for the public should streamline both employees' and the public's experience.
- **Focus on HPI.** Finally, safety of our workforce is paramount. We have spent a lot of effort over the past 2 years training our staff in human performance improvement. (HPI). The ultimate goal of HPI is to make a direct and positive impact on business outcomes and help our organization become better and faster at achieving our business goals. We, in Computing, need to start putting this philosophy into practice. We will start with a few areas and establish success. We will then use these newly created "experts" to help leverage these techniques in other areas.

Appendix A: Developing the Computing Strategy

The Computing Strategic Plan is aligned to support the laboratory's agenda and evolving business capabilities by providing robust computing services that are implemented through industry best practice.

In developing the strategy, we executed the following process. With the senior management team, we established 4 main strategic focus areas referred to as themes and then established 10 high-level long-term strategic goals.



The 4 focus areas are:

1. Workforce and Organization
2. Scientific Computing
3. IT
4. Process and Information Security

The 10 long-term goals identified were:

- Create a stimulating and rewarding work environment
- Selectively carry-out scientific computing research
- Continuously improve computing capabilities to support laboratory operations
- Optimize existing platforms and technologies
- Modernize business process and capabilities
- Enhance information security
- Safely enable research and efficient operations through the effective use of

people, process and technology

- Continuously improve computing capabilities to support the scientific program
- Modernize scientific processes and capabilities

We then engaged a broader group within Computing—about 30 managers in all—to assess the current situation, identify initiatives, establish priorities and identify gap areas. These discussions took place over a series of half-day workshops. From those discussions, we identified 43 objectives and 110 initiatives needed to achieve those 43 objectives.

Based on this vast collection of data, a series of themes emerged. Those themes are summarized below.

Focus Area	Theme
Scientific Computing	<ul style="list-style-type: none"> • Lead and innovate • Explore state-of-the-art technologies • Collaborate effectively with experiments, scientific programs and international partners
Core Computing	<ul style="list-style-type: none"> • Follow mature technology • Adopt proven platforms and applications • Partner with Key vendors and leverage evolving market efficiency
Workforce Management	<ul style="list-style-type: none"> • Plan for talent and skill • Encourage/seek women and minority participation
Organization	<ul style="list-style-type: none"> • Attract and retain top talent • Strive for this to be the best place to work
Cyber Security	<ul style="list-style-type: none"> • Support mandate and mitigate risk • Strike balance between control and flexibility
Process (Governance)	<ul style="list-style-type: none"> • Adopt industry best practices • Follow continuous improvement and efficiency
Technology Transfer	<ul style="list-style-type: none"> • Apply technologies that enable science to industry transfer

Out of this series of themes, we determined 8 underlying themes that we have included in our vision.