Department of Energy National Laboratories and Plants

Leadership in Cloud Computing

U.S. DEPARTMENT OF
ENERGY

Prepared by the National Renewable Energy Laboratory
(NREL), a national laboratory of the U.S. Department of Energy,
Office of Energy Efficiency and Renewable Energy; NREL is
operated by the Alliance for Sustainable Energy, LLC.

TABLE OF CONTENTS

U.S. DEPARTMENT OF ENERGY LABORATORIES AND PLANTS

The 22 U.S. Department of Energy (DOE) national laboratories and plants that comprise the nation's federal scientific research and defense systems provide strategic scientific and technological capabilities. Their collective goal is to meet the nation's challenges and priorities in these areas, which often reach beyond the scope of academia and private industry; and also to ensure that our government has access to these crosscutting discoveries and innovations.[1]

The following are brief overviews of the missions of each entity and the Information Technology (IT) professionals who contributed to this report.

Special thanks to the National Renewable Energy Laboratory, including Chuck Powers, Matt Fish, Joelynn Schroeder, and Kakie Walker.

AMES LABORATORY

The Ames Laboratory is a U.S. Department of Energy Office of Science national laboratory operated by Iowa State University. Established in the 1940s with the successful development of the most efficient process to produce high-quality uranium metal for atomic energy, the lab now pursues a broad range of scientific priorities. The Ames Laboratory creates innovative materials, technologies, and energy solutions. Using expertise, unique capabilities, and interdisciplinary collaborations to solve global problems. Building on that strength in the development and use of new materials, the Ames Laboratory scientists have expanded their work into seven main research areas. Their goals are to expand scientific knowledge and turn their discoveries into technology that will aid people throughout the world.

Diane Den Adel, Ames Laboratory, USDOE • 111 TASF • Ames, IA 50011 • Phone: 515-294-1061 • Email: ddenadel@ameslab.gov
Web: http://www.ameslab.gov/

ARGONNE NATIONAL LABORATORY

Argonne National Laboratory (Argonne) seeks solutions to pressing national problems in science and technology. The nation's first national laboratory, Argonne is one of the U.S. Department of Energy's largest national laboratories for scientific and engineering research. Argonne's mission is to apply a unique mix of world-class science, engineering, and user facilities to deliver innovative research and technologies. Argonne's programmatic activities cover all aspects of the innovation ecology: basic research, technology development, and prototype development and testing. Argonne regularly works with industry to transfer their innovative work to the marketplace.

Paul Domagala, Argonne National Laboratory • 9700 S. Cass Avenue • Argonne, IL 60439 • Phone: 630-252-5197
Email: domagala@anl.gov • Web: http://www.anl.gov/

BROOKHAVEN NATIONAL LABORATORY

Brookhaven National Laboratory (BNL) was established in 1947. It is a multiprogram lab conducting research in physical, biomedical, and environmental sciences; energy technologies; and national security. BNL has received seven Nobel Prizes for discoveries made at the lab.

Jim Allegue, Brookhaven National Laboratory • P.O. Box 5000, NY 11973 • Email: allegue@bnl.gov • Web: http://www.bnl.gov/

FERMI NATIONAL ACCELERATOR LABORATORY

Fermi National Accelerator Laboratory (Fermilab) advances the understanding of the fundamental nature of matter and energy by providing leadership and resources for qualified researchers to conduct basic research at the frontiers of high energy physics and related disciplines. Fermilab's broad scientific program pushes forward on three interrelated frontiers of particle physics. Each uses a unique approach to making discoveries, and all three are essential to answering key questions about the laws of nature and the cosmos.

Mark Kaletka, Fermi National Accelerator Laboratory •
P.O. Box 500 • Batavia, IL 60510-5011 • Email: kaletka@fnal.gov
Web: http://www.fnal.gov/

[1] U.S. Department of Energy, Office of Science. "Laboratories." Office of Science Online, http://science.energy.gov/ laboratories/ Accessed April 19, 2011.

IDAHO NATIONAL LABORATORY

The Idaho National Laboratory (INL) is the U.S. Department of Energy's national nuclear laboratory. INL serves a distinctive and unique role in civilian nuclear research, while operating and maintaining the core of DOE's essential capabilities and infrastructure needed for nuclear energy research, development, demonstration, and deployment. Its size, remote location, and safeguards and security provide an environment where the laboratory can test nuclear, chemical, electrical transmission, and other energetic systems under postulate normal and abnormal conditions.

Troy Hiltbrand, Idaho National Laboratory • P.O. Box 1625
Idaho Falls, ID 83415 • Phone: 208-526-1092
Email: Troy.Hiltbrand@inl.gov • Web: http://www.inl.gov/

LAWRENCE BERKELEY NATIONAL LABORATORY

The Lawrence Berkeley National Laboratory (LBNL) conducts unclassified research across a range of scientific disciplines. Its key efforts are in fundamental studies of the universe, quantitative biology, nanoscience, new energy systems, and environmental solutions; and the use of integrated computing as a tool for discovery. Founded in 1931, the laboratory boasts 11 scientists who have won the Nobel Prize, and has many other distinguished awards to its credit.

Adam Stone, Lawrence Berkeley National Laboratory
1 Cyclotron Road Mail Stop 65-0113 • Berkeley, CA 94720-8105
Phone: 510-486-4650 • Email: adstone@lbl.gov
Web: http://www.lbl.gov/

LAWRENCE LIVERMORE NATIONAL LABORATORY

The Lawrence Livermore National Laboratory (LLNL) was founded in 1952. LLNL is dedicated to ensuring the safety and security of the nation through applied science and technology in nuclear security, international and domestic security, and energy and environmental security.

Mark Dietrich, Lawrence Livermore National Laboratory
P.O. Box 808 • Livermore, CA 94551-0808
Phone: 925-423-4628 • Email: dietrich9@llnl.gov
Web: https://www.llnl.gov/

LOS ALAMOS NATIONAL LABORATORY

Los Alamos National Laboratory (LANL) is a premier national security research institution. Since 1943, the lab has delivered scientific and engineering solutions for the nation's most crucial and complex problems. LANL's primary responsibility is ensuring the safety, security, and reliability of the nation's nuclear deterrent. The lab also advances bioscience, chemistry, computer science, earth and environmental sciences, materials science, and physics disciplines.

James Franzen, Los Alamos National Laboratory • P.O. Box 1663
Los Alamos, NM 87545 • Phone: 505-665-6341
Email: jfranzen@lanl.gov • Web: http://www.lanl.gov/

NATIONAL NUCLEAR SECURITY SITE

For more than sixty years, the Nevada National Security Site (NNSS) has played a vital role in ensuring the security of the U.S. and its allies. Today, the site continues to provide a unique and indispensable extension of the national laboratories' experimental capabilities in support of the Stockpile Stewardship Program. The site also has become the nation's leader in Homeland Security with respect to nuclear/radiological testing, training, and emergency response. In addition to ongoing environmental cleanup of historic nuclear research and testing areas on NNSS, non-defense research, development, and training activities are conducted in cooperations with universities, industries, and other federal agencies.

Bob Hillier • National Nuclear Security Administration, Nevada Site Office • P.O. Box 98518 • Las Vegas, NV 89193-8518
Phone: 702-295-0411 • Email: hillierm@nv.doe.gov
Web: http://www.nv.doe.gov/main.aspx

NATIONAL RENEWABLE ENERGY LABORATORY

In operation since 1977, The National Renewable Energy Laboratory (NREL) is the nation's only laboratory dedicated solely to renewable energy and energy efficiency research and development. NREL develops renewable energy and energy efficiency technologies and practices, advances related science and engineering, and transfers knowledge and innovations to address the nation's energy and environmental goals. These areas span from understanding renewable resources for energy, to the conversion of these resources to renewable electricity and fuels, and ultimately to the use of renewable electricity and fuels in homes, commercial buildings, and vehicles.

Matt Fish, National Renewable Energy Laboratory
15013 Denver West Parkway • Golden, CO 80401
Phone: 303-275-3641 • Email: matt.fish@nrel.gov
Web: http://www.nrel.gov/

OAK RIDGE NATIONAL LABORATORY

The Oak Ridge National Laboratory (ORNL) is a multi-program science and technology laboratory established in 1943. The lab conducts basic and applied research and development to create scientific knowledge and technological solutions that strengthen the nation's leadership in key areas of science. These include increasing the availability of clean and abundant energy, restoring and protecting the environment, and contributing to national security.

Bruce Wilson, Oak Ridge National Laboratory • P.O. Box 2008
Oak Ridge, TN 37831 • Phone: 865-574-6651
Email: wilsonbe@ornl.gov • Web: http://www.ornl.gov/

PACIFIC NORTHWEST NATIONAL LABORATORY

The Pacific Northwest National Laboratory (PNNL) has delivered leadership and advancements in science, energy, national security, and the environment since 1965. The lab conducts applied research in information analysis, cyber security, and the nonproliferation of weapons of mass destruction; research in hydrogen and biomass-based fuels to reduce U.S. dependence on oil; and works to reduce the effects of energy generation and use on the environment.

Clay Hagler, Pacific Northwest National Laboratory
P.O. Box 999 • Richland, WA 99352 • Phone: 509-372-4487
Email: clay.hagler@pnl.gov • Web: http://www.pnl.gov/

PANTEX

Pantex is the National Nuclear Security Site's production integrator and provider of the nation's nuclear deterrent to the U.S. Department of Defense. Pantex was originally constructed by the U.S. Army in 1942 to load and pack conventional artillery shells and bombs in support of the World War II effort. After evolving through several iterations to support the nation's war efforts, Pantex's mission is now to safely and securely maintain the nation's nuclear weapons stockpile and dismantle weapons retired by the military. The plant's future includes life extension programs designed to increase the longevity of weapons in the stockpile.

Sean Dougherty, Pantex • P.O. Box 30020 • Amarillo, TX 79120
Phone: 806-477-6925 • Email: sdougher@pantex.doe.gov
Web: http://www.pantex.com/index.htm

PRINCETON PLASMA PHYSICS LABORATORY

The Princeton Plasma Physics Laboratory (PPPL) is a national center dedicated to plasma and fusion science with a leading international role in developing the theoretical, experimental, and technology innovations needed to make fusion practical and affordable. Since 1951, PPPL has worked with collaborators across the globe to develop fusion as an energy source for the world, and conduct research along the broad frontier of plasma science and technology.

Steve Baumgartner, Princeton Plasma Physics Laboratory
P.O. Box 451 • Princeton, NJ 08543-0451 • Phone: 609-243-2820
Email: sbaumgar@pppl.gov • Web: http://www.pppl.gov/

SANDIA NATIONAL LABORATORIES

Sandia National Laboratories (Sandia Labs) has developed science-based technologies that support national security since 1949. The lab develops technologies to sustain, modernize, and protect the U.S. nuclear arsenal; prevent the spread of weapons of mass destruction; defend against terrorism; protect national infrastructures; ensure stable energy and water supplies; and provide new capabilities to the U.S. armed forces.

Kelly Rogers • Sandia National Laboratories
PO Box 5800 • Albuquerque, NM 87185-0165
Phone: 505-844-5391 • Email: gkroger@sandia.gov
Web: http://www.sandia.gov/

SAVANNAH RIVER NATIONAL LABORATORY

The Savannah River National Laboratory (Savannah River) was founded in 2004 and is the applied research and development laboratory at the Savannah River Site (SRS). The lab is dedicated to solving complex national defense, homeland security, and nuclear material problems. They also provide applied research in environmental management, energy security, and technologies.

John Longo, Savannah River National Laboratory
Savannah River Site • Aiken, SC 29808 • Phone: 803-557-9911
Email: john.longo@srnl.doe.gov • Web: http://srnl.doe.gov/

SAVANNAH RIVER SITE

The Savannah River Site (SRS) is a long-term national asset dedicated to protecting public health and the environment while supporting the nation's nuclear deterrent and the transformation of the site for future use. Constructed in the 1950s, SRS's original mission was to produce the basic materials used in nuclear weaponry to support the nation's defense programs. In 1981, SRS began the shift into environmental stewardship and other areas of clean energy research. The site's current transformation objectives target impact in three business segments—national security, clean energy, and environmental stewardship.

Bruce Wilson• Aiken, SC 29808 • Phone: 865-574-6651
Web: http://www.srs.gov/

SLAC NATIONAL ACCELERATOR LABORATORY

The SLAC National Accelerator Laboratory (SLAC) is dedicated to the design, construction, and operation of state-of-the-art electron accelerators and related experimental facilities for use in high-energy physics and synchrotron radiation research. Founded in 1962, SLAC is a multipurpose laboratory for astrophysics, photon science, accelerator, and particle physics research. The lab boasts six Nobel Prize winning scientists.

Imre Kabai, SLAC National Accelerator Laboratory
2575 Sand Hill Road, Mail Stop 58 • Menlo Park, CA 94025-7015
Phone: 408-218-9604 • Email: imre@slac.stanford.edu
Web: http://www.slac.stanford.edu/

THOMAS JEFFERSON NATIONAL ACCELERATOR FACILITY

The Thomas Jefferson National Accelerator Facility (JLab) began operation in 1995. The lab provides forefront scientific facilities, opportunities, and leadership essential for discovering the fundamental nature of nuclear matter, to partner with industry to apply its advanced technology, and to serve the nation and its communities through education and public outreach. Scientists from around the world use the laboratory's facilities to conduct their research.

Andy Kowalski, Jefferson Laboratory • 12000 Jefferson Avenue
Newport News, VA 23606 • Phone: 757-269-6224
Email: kowalski@jlab.org • Web: http://www.jlab.org/

Y-12 NATIONAL SECURITY COMPLEX

The Y-12 National Security Complex (Y-12) maintains the safety, security, and effectiveness of the U.S. nuclear weapons stockpile. Y-12 also reduces the global threat posed by nuclear proliferation and terrorism, and provides safe and effective nuclear propulsion systems for the U.S. Navy. Built in 1943 in support of World War II, the plant's unique emphasis is the processing and storage of uranium and development of technologies associated with those activities.

Jeffrey Jones, Y-12 National Security Complex • 602 Scarboro Road • Oak Ridge, TN 37830 • Phone: 865-576-2335
Email: jonesja@y12.doe.gov • Web: http://www.y12.doe.gov/

ABOUT THIS REPORT

In this report, the 22 U.S. Department of Energy (DOE) laboratories and plants (research organizations) share thoughts, vision, and direction for cloud computing within the DOE complex. The laboratories' leadership in rapid adoption and innovation of cloud computing within DOE is showcased, as is their stewardship of taxpayer dollars. The intent of this report is to serve as a status report and a vehicle to share implementations and best practices for cloud computing across the nation.

The last several years have seen significant advances in efficiencies in how Information Technology (IT) delivers products and services to their enterprises. Perhaps the most promising technology is cloud computing. Almost unheard of before 2006, this architecture is revolutionizing how previously costly and resource-intensive applications and services are delivered. Cloud computing is also aligning IT functions better with the business side, providing needed software and services just in time and at the exact level the business actually needs.

In the past, IT drove the capabilities it pushed out to the business. Due to pressure from shrinking budgets and increased sophistication of the user base, competition is forcing alignment of IT objectives to the business. Because IT organizations have traditionally been slow in responding to business needs, users are increasingly finding their own computing solutions. Today, IT organizations must find ways to deliver tangible value to their clients with the "time to value" significantly reduced. Now, unless a critical value proposition exists that requires the software, platform, or infrastructure be built in-house, services provisioned from a public cloud provider may be the best way to quickly get services to the business without adding trained staff and infrastructure. IT products and services sourced in the cloud are expected to reduce costs, increase service quality, and improve responsiveness (time to value) in serving business needs.[2]

In addition to the increased efficiency and effectiveness of delivering IT commodity services through the cloud, a sustainability component also exists. U.S. companies that move to a cloud computing infrastructure can save upwards of $12.3 billion in energy costs by 2020.[3] Additionally, Gartner, the leading IT research and advisory company, projects that revenue from cloud computing will near $152.1 billion in 2014, an astounding 39% growth in revenue in the five years from 2009, when cloud computing took IT by storm.[4] It is estimated that global IT spending will top $3.7 trillion in 2012, a 2.5% increase over 2011.

Cloud computing has enabled a fundamental shift in the staff productivity paradigm that held that work can only be accomplished effectively when staff are physically in the office together. Work can now be accomplished from home or from anywhere around the globe, cutting back on the need to travel in order to collaborate.

The Federal Chief Information Officer's (CIO) 25 Point Implementation Plan to Reform Federal Information Technology Management mandates a shift to a "cloud first" policy.[5] This mandate recognizes that the private sector is already on board with cloud computing and is already reaping the benefits of the flexibility and scalability of cloud computing technologies. While the private sector increased its capabilities while lowering its costs with cloud computing, government entities were seeing failures in programs with traditional data center environments.

[2] HP Software Professional Services. "Enable cloud service strategies by running IT like a business." HP Software Cloud Consulting Service online, http://h20195.www2.hp.com/V2/GetPDF.aspx/4AA3-3784ENW.pdf Accessed May 20, 2012.

[3] Canu, A. "The history and future of cloud computing." Forbes online, http://www.forbes.com/sites/dell/2011/12/20/the-history-and-future-of-cloud-computing/ Accessed 6/16/2012.

[4] Petty, C and van der Meulen, R. "Gartner Says Worldwide IT Spending Figures Show Mixed Results for 2012." Gartner Newsroom online, http://www.gartner.com/it/page.jsp?id=1975815 Accessed 6/16/2012.

[5] Kundra, V. "25 point implementation plan to reform federal information technology management." Chief Information Officers Council online, http://www.cio.gov/documents/25-Point-Implementation-Plan-to-Reform-Federal%20IT.pdf Accessed May 20, 2012.

History of Computing

Innovations in computing have skyrocketed since the birth of Hewlett-Packard (HP) in a California garage in 1939.[6] Bell Laboratories followed HP's innovations with the first true computer in 1940—the Complex Number Calculator. The following years saw an explosion in computing technologies, with room-sized computing devices that continued to improve in speed, storage, and computing capabilities. Just 10 years after the Complex Number Calculator, the first commercially-produced computer was developed, as were the first standards governing computing. These early computers operated at 90% utility—an achievement that today's computing systems envy. Computing came to the attention of the public in 1951, when the U.S. Census Bureau UNIVAC computer came online.

In 1962, the Laboratory Instrumentation Computer was commercialized, bringing computing to the scientific world and offering real time laboratory data processes. The year 1964 introduced networking with computers that could communicate with each other and with peripheral devices, as well as the first supercomputer. Computing was finally a viable tool for sharing information across entities. Since then, computers have revolutionized the way people collaborate, advance, and discover things.

Cloud computing was officially coined in 1997 by Emory professor Ramnath Chellappa, who likely based it on the cloud symbol used to diagram IT infrastructure and the internet.[7, 8]

The growing demand for increased capabilities at reduced costs in the 1980s began the push toward a computing model that shifted away from large, expensive supercomputers, without sacrificing capability. The explosion of internet-based innovations in the 1990s led to the advent of cloud computing. During this decade, grid and utility computing paved the way for organizations to collaborate and to rent computing capabilities, opening the market for smaller businesses. Application Service Providers, then gave the world internet-enabled applications. Application Service Provider companies licensed a single application to multiple users, enabling the outsourcing of services. Software as a Service (SaaS) arrived.

Evolution of Computing Models

Cloud computing is the next evolutionary step in enterprise IT. Just as the mainframe computer gave way to the more sophisticated client/server model that is prevalent in today's computing world, cloud computing is well on its way to sending the client server model into history.[9]

What is Cloud Computing?

Most people make regular use of SaaS solutions in their day to day lives. Booking trips through services like Expedia.com, checking bank accounts through software banks provide, tracking UPS packages, and much more, are all hosted in the cloud.

These same capabilities have rapidly expanded and are providing a wide variety of internet-based services to organizations, public and private. Cloud computing is the delivery of compute and storage capacity as a service that allows IT organizations to provide hosted applications and services to their users on demand. Its goal is to increase the value of delivered products and services, with value defined as a function of cost and utility. Cloud computing extends current IT capabilities without increasing capital expenditures and allows organizations to pay for only the applications and services they need.

Because multiple organizations share the provider's product and all costs associated with it, the overall benefit is magnified by the economies of scale. Organizations can afford far more than would be possible individually, meaning that overall costs are significantly reduced. And because cloud-based solutions are isolated, the security risks are significantly reduced.

The National Institute of Standards and Technology (NIST) defines cloud computing as "a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."[10]

[6] Computer History Museum. "Timeline of Computer History." Computer History Museum online, http://www.computerhistory.org/timeline/ Accessed 6/16/2012.

[7] Canu, A. "The history and future of cloud computing." Forbes online, http://www.forbes.com/sites/dell/2011/12/20/the-history-and-future-of-cloud-computing/ Accessed 6/16/2012.

[8] Stark, C. 2012. "The history of cloud computing." CETROM online, http://www.cetrom.net/blog/the-history-of-cloud-computing/ Accessed 6/16/2012.

[9] Bias, R. "The evolution of IT towards cloud computing." Cloudscaling online, http://www.cloudscaling.com/blog/cloud-computing/the-evolution-of-it-towards-cloud-computing-vmworld/ Accessed May 20, 2012.

[10] Mell, O. and Grance, T. 2011. "The NIST definition of cloud computing." National Institute of Standards and Technology (NIST) Computer Security Division online, http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf Accessed May 20, 2012.

Cloud computing has seven essential characteristics:

- **Agility –** the ability to change rapidly, efficiently, and effectively
- **Device and location independence –** users can connect from anywhere
- **Virtualization –** allows servers and storage devices to be shared and utilization be increased
- **Multi-tenancy –** sharing of resources and costs across a large pool of users
- **Reliability –** multiple redundant sites
- **Scalability –** dynamic deployment of resources on a fine-grained, self-service basis near real-time
- **Application programming interface –** accessibility to software that enables machines to interact with cloud software.

Three service models and four deployment models are defined by NIST and will be covered in this report.

SERVICE MODELS

Service models include Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS).

IaaS: This is the most basic cloud service model where a Cloud Service Provider (CSP) provides "computers" generally as virtual machines, as well as storage and networks (they may provide physical computing devices, as well). These resources are supplied on demand from large pools of physical resources resident on the provider's servers.

PaaS: PaaS builds on IaaS by providing the necessary tools and software stacks used by the cloud user to assemble solutions. While IaaS may or may not provide the base operating system, PaaS would typically include all of the necessary Web servers, data manipulation tools, and development languages needed to implement a software system.

SaaS: In this model, the cloud provider delivers a completed software solution, such as a package shipping service, customer relationship management, or travel booking service to the cloud user. The user is typically not responsible for any of the software installation or management, but may be involved with some level of configuration to meet specific needs.

DEPLOYMENT MODELS

Deployment models include private cloud, community cloud, public cloud, and a hybrid of these.

Private cloud: Provisioned for exclusive use of a single organization.

Community cloud: Provisioned for use by a community of users with shared purpose or common requirements. An example includes providing cloud services to meet the specific needs of U.S. Government agencies.

Public cloud: Provisioned for use by the general public. In some cases, public cloud can be single tenant, where specific equipment is dedicated to specific customers or groups of users.

Hybrid cloud: Deployment of one or more of these models connected by technologies that enable portability.

Cloud Security

In December 2010, the Office of Management and Budget published the 25 Point Implementation Plan to Reform Federal Information Technology Management. The plan outlined a shift to cloud services requiring agencies to use cloud-based solutions where feasible.

While many promises of cloud computing are compelling, the DOE complex cannot fully move forward without focusing on cyber security issues. Gartner raises seven security issues when moving to the cloud: privileged user access, regulatory compliance, data location, data segregation, recovery, investigative support, and long-term viability.

ADOPTION OF THE CLOUD AND RISK PROFILE

The amount of effort required to deploy, build, or use cloud services will vary from case to case, and service to service. Adoption of the cloud comes with many benefits and also security considerations. Risks can be limited by maintaining islands of risk and isolating the effect of security events.

Cloud solutions have varying level of flexibility, control, and risk management. The three cloud service models show how their benefits align with security considerations.

FEDRAMP

In order to address such issues of risk and security, the 25 Point Implementation plan also established a strategy

SaaS	PaaS	IaaS
Benefits: lowest total control of ownership, lowest time to production	**Benefits:** lower total control of ownership, lower time to production, increased flexibility and management control	**Benefits:** Reduced total control of ownership, reduced time to production, complete flexibility and management control
Cyber Security Considerations: Least amount of organizational control and flexibility for controls and configurations, most dependent on service provider for cyber protections, incident response, and disaster recovery	**Cyber Security Considerations:** Security of underlying architecture, access to information systems from outside entities or service providers, moderate level of dependence on service provider for cyber protections, incident response, and disaster recovery	**Cyber Security Considerations:** Location of primary and backup data centers, physical protection of data and information systems, location of support personnel

to “approve once and use often,” a government-wide risk program aimed to add consistency to the application of security controls and eliminate redundant efforts.

To facilitate this approach the Federal Risk and Authorization Management Program (FedRAMP) was established by the Office of Management and Budget in December 2011, for the Security Authorization of Information Systems in Cloud Computing Environments, and is expected to result in a cost-effective, risk-based approach that will enable the rapid adoption and usage of cloud services. The program provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services.[11]

The goals of FedRAMP include:

- Accelerating the adoption of secure cloud solutions through reuse of assessments and authorizations
- Increasing confidence in security of cloud solutions
- Achieving consistent security authorizations using a baseline set of agreed upon standards
- Ensuring consistent application of existing security procedures
- Increasing confidence in security assessments
- Increasing automation and near real-time data for continuous monitoring.

The program seeks to accelerate the adoption of cloud solutions by increasing confidence in the technology’s security. This centralized assessment and authorization agency allows for reuse of findings across government entities, speeding the implementation of cloud services to government programs. A CSP must apply for accreditation through FedRAMP and once in the database, government entities can pull from the authorized CSPs pool to quickly implement cloud-based technologies.

Energy Risk and Authorization Management Program (E-RAMP) is being implemented to establish security control baselines that align with DOE policies and NIST guidance. E-RAMP is expected to be the DOE organization’s front-end to FedRAMP to ensure the consistent implementation of controls across the department, assist with the Office of Management and Budget reporting requirements, and facilitate the efficient and effective provisioning of cloud services for department organizations.

CONCLUSION

Risks associated with cloud service providers do not fall only on CSPs. Organizations must take responsibility, reviewing, and understanding the risks associated with each service provider.

DOE laboratories are showing leadership by moving toward contracting with FedRAMP-approved CSPs exclusively to meet the mandates of the Federal CIO’s “cloud first” policy.

RightPath – DOE/NNSA Cloud Strategy

Navigating the increasingly complex maze of government policies, evolving technologies, and conflicting information related to cloud computing and a virtual workforce can be a daunting and even overwhelming task. Going down the wrong path can lead to costly project overruns, millions of dollars wasted, and government initiatives that fail catastrophically. A typical government deployment path tends to involve multiple technology point solutions or a single solution provider in an attempt to accelerate their journey to the cloud. This fragmented approach and the resulting lack of meaningful impact has resulted in

[11] http://www.gsa.gov/portal/category/102371 (Accessed September 6, 2012)

widespread hesitancy to adopt cloud technologies on a broad scale with many government agencies “putting their toes in the water” versus aggressively leveraging technology to drive dramatic cost efficiencies and new capabilities.

Lack of a solid framework unifying these initiatives will result in a government-wide failure to deliver a cost effective and scalable solution that meets the needs of a 21st century virtual workforce. To effectively speed adoption of emerging technologies, each agency must adopt a proven, unifying framework that coordinates people, processes, and technology to deliver a leaner government.

DOE/NNSA have partnered to deliver a new set of capabilities, supported by innovative cyber security and policy reforms that will provide increased flexibility and agility, lower costs, and improved communications/collaboration for our employees through the RightPath program. In addition, the methodology/framework used to deliver these capabilities will be corporately captured and provided to other government agencies as a blueprint for delivering virtual workforce, mobility, and cloud-based solutions in a rapid, agile, and effective manner.

CURRENT STATE

DOE/NNSA currently operates in “silos of excellence” whether internal or external. Internally, IT services and capabilities are delivered independently across a variety of program lines and there is great redundancy in the overall federal IT portfolio. Externally, DOE/NNSA provide strategic direction and oversight for a distributed IT architecture that is managed largely through indirect budgets at each M&O site. Due to mission diversity, contract clauses, and technology limitations, this architecture has developed in a decentralized manner with federation of capabilities available in a small subset of use cases.

While this architecture and approach has served DOE/NNSA well over the years, advancements in technology present compelling opportunities to re-architect service delivery to realize new capabilities to support our mission lines and to reduce the costs of existing capabilities throughout the enterprise.

CLOUD VISION

DOE/NNSA is pursuing a “cloud of clouds” approach as one of the strategic elements of the IT modernization strategy. Realizing that a one-size-fits-all strategy is not appropriate for our diverse mission lines, the RightPath program is investing in the delivery of a secure cloud services brokerage technology, YOURcloud, that will connect a diverse customer set spanning federal and M&O constituencies to a federated marketplace of cloud service providers. This approach should simplify acquisition, provide an easier on-ramp for cloud service providers, deliver hardware cost saving based on economies of scale, reduce licensing costs, transition from a CAPEX to an OPEX model, and offer improved business agility (taking the process of procuring, configuring, and deploying a server from months to minutes).

Innovation around cloud services brokerage is critical to cloud computing success within DOE. It allows sites to maintain full autonomy to manage their workloads, provides for a federated capability, and provides a common security baseline for sites to leverage (normalizing risk and reducing site specific effort). This effort will build a base for Shared First initiatives (through the Shared Services enclave in YOURcloud and the App Store) while fully avoiding a “centralize and consolidate” strategy that would violate site autonomy, reduce mission capabilities, and drive down innovation. Effectively, delivery of the cloud broker provides that capability to have best of both worlds; fully fusing cost savings and innovation into an agnostic orchestration platform.

KEY INITIATIVES

In the coming 12 months, RightPath is focused on delivering three signature technologies:

ONEvoice – a comprehensive collaboration solution connecting scientist to scientist, Headquarters (HQ) to field, and fed to contractor in a rich, immersive technology platform. Initial deployment will federate Microsoft Lync between sites and HQ allowing desktop video, voice, instant messaging, web conferencing, desktop sharing, and presence capabilities across geographic boundaries. This effort will provide immediate and meaningful improvements in cross-site collaboration. In addition, this collaboration technology stack will be federated to VMware Socialcast, a business focused, internal social network capability. The social network will connect people with similar interests, allow for the creation of communities of interest, facilitate cross-site project teams, and bring new capabilities such as town halls, crowd sourcing, and social analytics to the enterprise.

OneNNSA Network – delivers a secure overlay network that provides FIPS 140-2 encrypted communication paths between sites and HQ. This approach provides a high bandwidth and secure transport for consuming services from YOURcloud and utilization of the ONEvoice stack. In addition, the network will provide federated identity management using SAML 2 standards and a new enterprise virtual directory that will enable single sign on for cloud services and provide a basis for future implementations of HSPD-12 both logically and physically.

YOURcloud – a secure cloud services brokerage capability based on IOD, developed by LANL and re-platformed in version 3 to address the DOE/NNSA enterprise requirements. YOURcloud will provide a self-service portal for IaaS offerings across multiple cloud services providers on premise, corporately provided, and public (i.e. Amazon EC2). YOURcloud will provide a diversity of choice to sites for IaaS providers while allowing sites to maintain full autonomy of their workloads. In addition, YOURcloud will also provide a Shared Services Enclave that will provide the foundation for the Enterprise App Store in DOE/NNSA.

CONCLUSION

RightPath will serve as a key element in the overall DOE IT modernization plan and provide an architectural foundation that will serve as an innovation catalyst for both federal and contractor IT programs as they seek to retool their IT investments to deliver key technologies of the future such as cloud computing, mobility, social computing, and big data/analytics. In particular, the YOURcloud offering will be an enabler for cloud computing capabilities throughout the enterprise. The cloud broker technology will minimize cyber security efforts in cloud deployment, provide a robust marketplace, allow for common use applications between sites, and provide a low cost destination point for data center consolidation.

THE AMES LABORATORY

The Ames Laboratory Information Systems (IS) office includes cloud service considerations as a part of each IT initiative. The goal is identifying opportunities to provide excellent research support with minimal overhead.

Key questions considered for each project include:

- Are there existing offerings in the SaaS, IaaS, or PaaS spaces which can potentially provide this service?
- What are the potential risks?
- What are the potential impacts to user experience by using a cloud-based service?
- Is there a business case where additional cyber security risks are involved?
- What is the opportunity cost of providing staff with the potential to learn new technologies or increase their expertise?

The Ames Laboratory is utilizing cloud computing to reduce the cost of services while saving on administrative effort and implementation time. This allows IS to focus on aligning services with the mission of the laboratory. At the Ames Laboratory the objective is to optimize IT; implementing cloud computing services is one way to accomplish this goal.

The Ames Laboratory is focused primarily on utilizing SaaS deployments. The laboratory's initial SaaS application was a small system used for tracking compressed gas cylinders which streamlines effort to track cylinders and calculates the end user cost for the cylinder.

Vision

The Ames Laboratory views cloud computing services as an important strategic component in a successful information infrastructure that is capable of supporting world-class research. It is imperative to pay attention to the changing offerings of cloud-based services to take full advantage of opportunities for improved growth and efficiency. The Ames Laboratory continues to evaluate cloud technologies to deliver mission-enabling products and services.

It is anticipated that cloud computing services will reduce the risk of IS losing its control over decision-making by adjusting technology, processes, and roles, and by providing an easily accepted and transparent set of services and costs. Cloud computing services allow IS to evaluate long-term problem solving and decision-making processes with the opportunity to focus on the protection of data by limiting what goes to the cloud and encrypting stored data.

Key Initiatives

As research demands on information services become more critical and dynamic, the Ames Laboratory anticipates the need for a scalable, reliable, secure, and flexible infrastructure, extending beyond the traditional boundary of site-controlled infrastructure. The Ames Laboratory will evaluate cloud alternatives first and select an option that provides a secure, reliable, and cost-effective solution. If cloud services are selected, they may take the form of commercial services, private cloud implementations, or a combination of the two. In preparation, the following key initiatives are underway:

COLLABORATION AND CONTENT MANAGEMENT TOOLS

Iowa State University is the Ames Laboratory's contractor, and they are positioning themselves as a private cloud provider. The Ames Laboratory is in the process of employing two services from Iowa State University, including:

- Collaboration software tools simplifying business workflow processes.
- Content management providing the resources for long-term administration and storage of documents.

VIRTUAL DESKTOP INFRASTRUCTURE

The Ames Laboratory is developing a private infrastructure cloud using VmWare's Virtual Desktop Infrastructure (VDI) for administrative infrastructure. This effort leverages the existing physical infrastructure investment to centralize and provide managed, high-performance desktop environments to administrative staff. The expected result is reduced hardware costs and improved energy efficiency, plus streamlined desktop support and cyber incident response. In addition, user data is redirected to a backed-up central file storage server, increasing data availability and improving overall continuity of operations.

PUBLIC FACING WEBSITE

The Ames Laboratory plans to pursue the use of cloud services for hosting its public facing website. The goal is providing a quality website with high availability and reducing administrative costs with managed hosted services.

PAYROLL COMPENSATION

The Ames Laboratory is evaluating a Human Resources (HR) compensation and analysis tool to streamline laboratory-wide benchmarking to assist with employee retention and recruitment. The current payroll compensation process requires manual effort to perform this task. Utilizing a cloud-based application saves numerous hours of effort with no additional administrative infrastructure burden.

FEDERATED IDENTITY MANAGEMENT

A key challenge for all DOE laboratories is managing identities of external collaborators who participate in projects. A national lab federated identity project linking all interested DOE laboratories is being led by LBNL. The Ames Laboratory is participating in this project.

The DOE labs are leveraging the work of Internet2 to support federated identity management through InCommon and Shibboleth which allows researchers to use their lab identity to authenticate to resources at other institutions.

ARGONNE NATIONAL LABORATORY

Argonne National Laboratory has taken a leadership role within the DOE community on advanced computational and IT service delivery technologies. Many of the precursor technologies on which cloud computing is built have their roots in Argonne programs and tools such as the Globus Toolkit, GridFTP, TeraGrid, and Open Science Grid. As cloud computing came of age in the 21st century, Argonne translated this expertise into ground-breaking research and development work through the Nimbus project, an integrated set of tools that deliver the power and versatility of infrastructure clouds to scientific users and Magellan, a DOE-funded nationwide scientific cloud computing testbed. Argonne operates the Argonne Leadership Computing Facility, home to three of the world's fastest and most energy-efficient computers.

While computational science research is generally not directly applicable to the mission support side of the laboratory, operational IT maintains a close relationship with and draws heavily on Argonne's computational legacy and in-house expertise. Argonne was an early advocate of cloud technology, executing some of the first in-depth evaluations of cloud and SaaS offerings and sponsoring awareness and educational opportunities within Argonne and the DOE community at large.

Approach

Argonne has a simple but effective strategic approach to cloud computing that is aimed at optimizing overall business value of IT services. A portfolio management approach is applied to all IT projects and services. Sourcing options, whether cloud, SaaS, or in-house application, are evaluated and chosen based on merit in the project execution phase. The portfolio management process uses the following criteria when selecting and sourcing IT services: strategic value, health/safety/security, efficiency and process improvement, compliance, total lifecycle cost and benefit, enterprise risk, and sustainability.

Key Initiatives

Beginning in 2008, Argonne's Computing and Information Systems division took its first steps into cloud-based services with the adoption of Adobe Connect, a hosted Web conferencing service. In the 2009-2010 timeframe, Argonne performed in-depth evaluations of messaging and collaboration, frameworks (PaaS), portals, and computation and storage cloud services. Providers included Google, Force.com, Appian, Amazon, Appirio, Right Scale, Cloud Scaling, SGI, Penguin Computing, and Parabon.

Most notably, Argonne executed a sizable Google Enterprise Apps proof of concept in 2009. The conclusion was that there was not, at that point in time, sufficient benefit to merit outsourcing messaging and collaboration services to the cloud.

Argonne has adopted cloud-based offerings where the business analysis has been favorable. The laboratory currently uses Adobe Connect, SalesForce, YouTube, Ustream, and Flickr enterprise offerings. Procurements are also under way for Google Maps Engine, which will be a foundational component of the Argonne Enterprise GIS services and Oracle cloud services.

Despite the hype surrounding cloud services, the fact remains that the service model is still maturing, albeit rapidly. Inhibiting factors such as ill-suited procurement procedures, inability to agree on terms and conditions, and risk acceptance remain however. Fortunately, these barriers are well known and industry is working with federal entities toward solutions. Adoption of cloud-based services will undoubtedly grow, and become a larger portion of Argonne's IT service portfolio. Sourcing of IT service will always remain, as stated in the approach, a matter of sound business decisions and mission alignment and not a predisposition for any particular technology.

BROOKHAVEN NATIONAL LABORATORY

Brookhaven National Laboratory (BNL) is currently using both SaaS and PaaS cloud offerings in production. In addition to the production use, there are also separate evaluations for additional SaaS products, and a comprehensive evaluation of Amazon Web Services (AWS) for large scale IaaS deployment.

LawLogix is an SaaS package that assists BNL personnel with electronic visa processing. Although LawLogix was a learning experience for the lab, it's now properly integrated with the on-site PeopleSoft deployment. This allows data to move between the cloud and on-site Human Resources Management System without the extra time and errors associated with duplicate data entry.

AWS is already being used to provide services that are isolated from BNL's network. The first service is a monitoring service that checks on the availability of BNL's public facing services to ensure that they are online and available to the outside world. Having an isolated system is invaluable, because it can truly emulate the end-user experience of accessing content from a different network.

The second service is a set of computing resources that can be used in an emergency to provide vital communications within Brookhaven's cyber security group and to and from DOE. Additional systems and services can be provisioned rapidly, allowing for a flexible and secure area to facilitate communications in the event that the laboratory's network is unavailable.

In addition to the existing, small scale use of AWS, there is a larger scale proof of concept project that is ongoing, where all aspects of Amazon's IaaS offering are being evaluated in many dimensions. This project is already yielding much valuable information.

Moving forward, the laboratory is evaluating cloud-based solutions as an alternative to traditional off-site tape vaulting and disaster recovery. There are some areas where BNL may have no choice but to move into the cloud space. Of the five Applicant Tracking software vendors being considered to supplement the hiring and onboarding process, all of them are cloud-based SaaS packages, even though that was not a requirement.

Unique Challenges of SaaS

The integration of LawLogix with existing BNL systems highlighted many challenges that are specifically related to SaaS offerings. In addition to the normal considerations regarding the quality and documentation of 3rd party application programming interfaces, there were also many factors that were not anticipated. Because the data is not on-premise, the only available option for integration was the provided Web services application programming interface. The release schedule for software updates is entirely at the discretion of the cloud service provider. This significantly impairs the lab's ability to manage the risk associated with software updates. Finally, the availability of a development environment was at the whim of the cloud service provider, making it impossible to work on the integration between Brookhaven's existing systems and LawLogix while the development environment was offline.

There are also many benefits that are directly related to these challenges. Brookhaven is no longer responsible for applying security patches or maintaining hardware related to the service. New features and bug fixes are delivered without any effort on the part of Brookhaven's development staff. Finally, the responsibility to back up the data that is housed by LawLogix is no longer Brookhaven's responsibility.

Even taking all these additional considerations into account, LawLogix was still the best option for BNL. Now, equipped with a better understanding of the potential pitfalls that are unique to SaaS offerings, Brookhaven is better equipped to evaluate SaaS packages in the future.

Infrastructure Proof of Concept

The goal of the proof of concept is to determine if AWS can be used to work toward one of three major goals of the proof of concept: reducing costs, decreasing cyber security risk, and enabling external collaboration, which has recently been a pain point for BNL.

Six scenarios were devised to be evaluated against the stated goals:

- An Internal SharePoint Server, to test performance and cost effectiveness of running a complex software stack on an Amazon Instance, which is joined to the BNL domain and complies with all BNL policies

- A Scientific Computing Cluster, to measure the cost and performance of running computing jobs in Amazon EC2
- A Physical-to-Cloud migration, to test the feasibility of migrating applications from legacy hardware or on-site virtualization products to Amazon Instances, and to measure the costs associated with such a migration
- An External SharePoint Farm, to test external collaboration scenarios, including ADFS for authentication
- An "Emergency" Exchange server, to test the possibility of running their own, familiar email platform for emergency communications via Amazon EC2 in case of an event that would cause BNL's on-site email servers to go offline
- A load-balanced external website, using Amazon's Elastic Load Balancers to test system redundancy, and scaling in the event of heavy load.

In addition to the stated goals, a cyber security risk register is being compiled, as well as a list of governance recommendations to develop strategies for effectively controlling the use of IaaS offerings in general, and AWS specifically.

As the lab completes its various cloud-related research projects, it's becoming clearer that cloud services are not always more cost effective, available, or secure than traditional software, platforms, or infrastructure. Nor can they instantly decrease the skill set required to provide effective IT services. They are simply a shift in the way that computing services are provided—away from specialized hardware, platforms and software, away from high upfront costs and difficult to predict recurring costs and toward commodity hardware, platforms, and software, and flexible usage and billing options.

With this shift it must be acknowledged that many traditional policies and processes simply may not apply. This is where the adherence to programs like FedRAMP will allow organizations including BNL to adapt their policies to accommodate cloud services faster and with less effort that it was previously possible.

With a thorough understanding of the unique challenges that cloud services present, as well as the benefits they provide, BNL is poised to be able to make the best use of the available, FedRAMP-approved cloud service providers to enable faster and more flexible service delivery.

FERMI NATIONAL ACCELERATOR LABORATORY

Strategy

Fermilab's Computing Sector supports the scientific mission of the laboratory through developing and supporting innovative and cutting edge computing solutions and services for Fermilab. The services provided that are relevant to cloud computing fall into three main areas. First, there are the massively compute- and data-handling-intensive activities associated with simulation and analysis of detector and accelerator data. These are mainly handled by Grid facilities at Fermilab. Second, there are the commercial but specialized types of enterprise applications associated with running any business, such as HR and finance systems. Third, there is the IT infrastructure of servers, networks, and applications such as email and document management which support broadly and equally the scientific and business activities of the laboratory.

No single approach to cloud computing satisfies requirements for all these areas, so Fermilab's cloud computing strategy allows approaches which are tailored to the needs of each.

The strategy for data-intensive scientific computing relies heavily on Grid computing, which can be considered a form of PaaS cloud computing. Grid computing is the only approach suitable for the massive amounts of data generated by the Tevatron and LHC experiments. As an active member and contributor of the Open Science Grid (OSG), Fermilab led the early development and adoption of Grid computing within HEP. Commercial public IaaS cloud offerings, such as Amazon EC2, have been evaluated and some interoperability has been demonstrated with Grid services, but these offerings remain cost-prohibitive for data-intensive applications. At Fermilab, non-data-intensive scientific computing is also supported on IaaS and PaaS private clouds.

For commercial enterprise business applications, Fermilab has adopted a strategy of first considering SaaS cloud solutions when implementing new systems or making major upgrades to existing ones. The primary goal of this strategy is to reduce the overall cost of operation and maintenance of these systems by adapting business processes to standard commercial SaaS solutions, rather than allowing business processes to drive customization of expensive in-house deployments. This reduces internal support costs, frees IT resources for other projects, and avoids very expensive test cycles when customizations must be brought forward in new releases of products. Additional goals include reduced infrastructure costs, high availability, and disaster recovery.

Finally, for infrastructure computing, the strategy is to use private IaaS cloud offerings to provision servers for a wide variety of applications used by scientific and staff users. The chief goals are more efficient utilization of computing facility infrastructure (power, cooling, floor space), lower cost and reduced procurement cycles for deploying hardware for new services, high availability, redundancy, and disaster recovery. A key outcome is that services provisioned to use this private IaaS cloud infrastructure are considered compliant with ITIL management processes for Capacity, Availability, and Continuity, with little or no additional effort for the service provider.

Implementation

Fermilab's cloud computing implementation currently includes in-house IaaS and PaaS private clouds, and externally hosted commercial SaaS. In-house IaaS and PaaS service offerings are further tailored to the requirements of scientific users and developers, and to the requirements for enterprise and infrastructure applications and databases. Commercial SaaS solutions are used primarily for specific business applications.

At Fermilab, the Grid infrastructure is supported by a private IaaS cloud which provides redundancy and high reliability for these important services. Private IaaS and PaaS clouds also augment the Grid for scientific computing by providing hosting services for scientific computing users and applications. This private cloud is also interoperable with public clouds, such as Amazon EC2.

Virtualization and cloud computing are central to Fermilab's computing strategy. Most major applications—scientific as well as business—are now developed and/or operated in virtualized IaaS environments. Commercial SaaS solutions are always considered when planning new business applications, particularly when Fermilab expects that COTS solutions will meet our requirements without customizations.

- FermiCloud IaaS, whose primary mission is to provide a platform for scientific research that integrates with scientific applications, concentrating on tailoring to the various scientific research needs
- Enterprise SaaS, whose primary mission is to move applications which are not the lab's core competency to SaaS, freeing IT resources for internal projects and reducing operational costs
- Virtualization Services IaaS, whose primary mission is to support a broad range of commercial and business applications in a production environment, with emphasis on stability and wide commercial acceptability and support.

FermiCloud

FermiCloud is an IaaS private cloud service which supports development and integration for primarily scientific applications and users, as well as production operation of certain grid and other services. It is also used as a testbed for open source cloud computing frameworks.

FermiCloud users get access to virtual machines without intervention of a system administrator, i.e. the end user is responsible for administering their own virtual machines. Virtual machines are created by the users and destroyed when no longer needed or used. Storage is provided from network-attached storage systems and high-performance channel-attached systems.

The FermiCloud hardware platforms have been configured to offer a flexible IaaS hosting environment and include high performance Infiniband networking that supports High Performance Computing (HPC)/MPI applications. Benchmarks using the FermiCloud infrastructure have shown little to no performance penalty using the "virtual HPC" capabilities when compared to equivalent "bare metal" HPC provisioning.

A goal of FermiCloud in the coming year is to provide a model for how an IaaS facility can be integrated with distributed computing operations of international scientific collaborations, with special attention to unified authorization, authentication, and accounting standards, and to understand interoperability requirements with other virtualized distributed cloud infrastructures and to further demonstrate interoperability.

Enterprise SaaS

In the last two years, Fermilab has moved several important applications to SaaS hosted cloud solutions. Fermilab Time & Labor uses Kronos for time and labor reporting, with successful integration into our backend Oracle financials and HR systems. Fermilab also successfully moved from an in-house help desk ticketing system to the hosted ServiceNow cloud solution and will soon replace their HR job applicant tracking with a hosted service.

Commercial SaaS hosted solutions will be considered as alternatives to deploying or upgrading in-house enterprise systems. Fermilab will be strongly considering SaaS hosted solutions as alternatives for future upgrades of financials, HR, and email systems, in particular.

Virtualization Services

Virtualization Services operates high quality IaaS private cloud services for development, integration, and production operation of enterprise applications and databases. Services are not typically available directly to end users, rather, they are offered to administrators (system, application, or database) within the IT organization who are deploying enterprise-level systems. This ensures a high level of stability in the deployed virtual systems.

Since Virtualization Services supports many important laboratory business functions, its private cloud is designed for reliability, availability, and performance. In order to support commercial applications, Virtualization Services' private cloud is based on VMware rather than open-source technologies. As with FermiCloud, storage is provided from high-availability network-attached storage systems and high-performance channel-attached systems.

Virtualization Services utilizes modern backup and replication tools capable of providing data de-duplication, instant image-level and file-level restorations, backup to disk, distributed job engines, and self-service to allow system administrators to view and restore files for the systems they manage. This technology can save time and money over traditional agent-based backup systems and gives us the ability to perform backups over the SAN, LAN, or a combination of both.

VDI

Fermilab is investigating VDI but has not yet made a strong move in this direction. This year, a pilot project provided standardized virtual desktops with a "bring your own device" arrangement for students participating in summer programs at Fermilab. The pilot was considered successful and Fermilab plans to expand the deployment in the coming year. Targets for VDI deployment include: kiosks; training PCs; payroll clerks; stock room / shipping / receiving / property clerks; developer systems; "loaner" desktops; etc.

IDAHO NATIONAL LABORATORY

With the complex-wide focus on driving costs down, while at the same time increasing the strategic relevance of work activities, it is imperative that IT rethinks how to provide services that enable research and business engagement. In order to accomplish this, it is important to revisit the service portfolio and make a clear distinction between those services that provide strategic value to the organization and those that are necessary, and can be fulfilled as a commodity through a cloud service provider. When services are classified as commodities, an analysis must be performed to determine the relative cost of supporting service execution internally against what can be found externally in the market. The summary of this analysis forms the foundation for defining which commodity services make good candidates for moving to the cloud. Idaho National Laboratory (INL) has developed a service framework which identifies current commodity services and includes the future state of how they should be delivered in order to optimize strategic value to the laboratory.

History

In 2010, INL was faced with a vexing challenge. The lab's existing Lotus Notes infrastructure was quickly becoming unmaintainable and was failing to meet the needs of the laboratory as it pushed toward a future characterized by national and international collaboration and partnerships. The organization was faced with a dilemma: it could modernize the existing infrastructure which would require a multi-year capital investment commitment, or it could choose to leap-frog the current paradigm of an internally-hosted email solution and move directly into the cloud, leveraging the rapidly growing market of cloud communication and collaboration. With the vision for cloud utilization, which had been established at the highest levels of government, INL decided to bypass the infrastructure-heavy alternative in favor of the cloud to provide communication and collaboration services to its personnel.

After engaging with the lab's workforce to determine the factors that were most important in a cloud solution, a request for proposal was released delineating INL's requirements, both from a technical perspective and from a user experience perspective. As a result, a contract was awarded to Unisys, a reseller of Google Apps for Government services. This service promised many capabilities that were above and beyond of the scope of a comparable, on-premise solution including:

- Dynamic supply scaling of resources both up and down to meet the demands of the laboratory
- Significantly greater storage capacity for email
- Ability to carry much of the cost as operating expense instead of relying upon capital investment.

INL's engagement of Unisys for the implementation and deployment of Google Apps for Government has not been without its challenges, especially as INL has pioneered new ground in the area of cloud computing within the DOE complex.

Information security has been of particular significance in this endeavor. As INL moves forward, the technical security around the encryption mechanisms of information both in-transit and at-rest have been reviewed and assessed. In this assessment, it was determined that the overall risk profile associated with this information decreased due to the modernization of infrastructure in the cloud and the advanced protection mechanisms in place by Google.

Issues relating to the handling of International Traffic in Arms Regulation 2011 and export control have been at the forefront of much of the activity surrounding the implementation of Google Apps for Government. INL worked jointly with DOE HQ and the U.S. Department of State in creating enhancements to existing regulations and definitions associated with International Traffic in Arms Regulation 2011 data, establishing a standard that will facilitate other labs in moving communication and collaboration to the cloud. Internally, the policies and procedures surrounding the use and management of export control data required review and revision. Resulting updates to policies and procedures at INL ensure that work is conducted with at least an equal level of effectiveness in the cloud as when using an on-premise infrastructure solution. By using a team approach in collaborating with the rest of the laboratory and other agencies, INL successfully established a path for the management of data in the cloud. This approach can be leveraged by other laboratories as they move forward with cloud initiatives.

As a result of INL's experience with the Google Apps for Government contract, the laboratory has developed a risk-based, data-centric approach to cloud procurements. Standard requirements language has been developed for future cloud contracting activities to ensure that protection of the laboratory's data is dependent on the level of associated risk. Trying to protect all data exactly the same using the same level of control quickly exaggerates operating cost. By utilizing a graded approach, the lab will ensure adherence to the right level of control for the level of risk inherent to the data moving to the cloud.

Key Initiatives

INL has set forth a vision with respect to the future use of cloud solutions. With each new service requirement, cloud solutions will be evaluated and given preference over on-premise COTS and custom-developed applications when solution requirements can be met. This will ensure that INL focuses its internal resources on those services which are of highest strategic relevance. Those solutions which are commodity in nature will be managed by service providers possessing the expertise in delivering cloud services in an optimal manner.

INL has realized a significant demand for hosting infrastructure and is currently working to establish a cloud procurement framework with defined providers. At an enterprise level, the lab will establish a contractual vehicle to facilitate the acquisition of IaaS by lab programs. The framework will ensure central oversight over usage and data management. Key to this cloud framework will be information security, scalability of service, communication and coordination with the service provider, and capabilities from an execution and reporting perspective. With the framework in place, duplication of effort in procuring IaaS will be eliminated or reduced and centralized technical points of contact will ensure that the lab manages proliferation of cloud infrastructure. The end results include programmatic flexibility and scalability while ensuring operational sustainability.

Over the past two years, INL has been heavily involved in forging new territory in the arena of cloud utilization. Several important lessons have been learned that will enable the lab to optimize its processes moving forward. As a result, the future success of cloud computing at INL is assured. Cloud computing is a fundamental aspect of INL Information Management's vision of strategic partnership with the laboratory and ensures that internal resources are committed in areas of highest relevance to the mission. This can only be accomplished by leveraging cloud services to fulfill essential, but non-strategic functions. Cloud utilization will allow INL to take advantage of centers of expertise throughout the industry, while using economies of scale to drive costs down. This makes the implementation of cloud services to meet commodity needs an essential element of enabling the mission of the laboratory.

LAWRENCE BERKELEY NATIONAL LABORATORY

Lawrence Berkeley National Laboratory (LBNL) was an early adopter of cloud technologies and continues to pursue numerous cloud-based solutions for scientific and business problems. With its diverse portfolio of open, fundamental research, LBNL has been in a unique position to test and adopt a wide range of cloud solutions. From scientific computing to collaboration, to business systems, LBNL has taken a leadership role in the technical, operational, and policy aspects of cloud computing.

Cloud Foundations

From the earliest days of the modern cloud, LBNL's strategy has been to build the human, technical, and policy foundations for the intelligent selection, deployment, and management of cloud computing services. The human foundations began with aggressive efforts to educate both IT professionals and scientific users on the expanding toolbox provided by cloud computing. Outreach included classes on programming for PaaS offerings, a seminar series on the implications and architectures of cloud computing, and numerous talks, seminars, and demonstrations.

On the technical side, LBNL put in place multiple "cloud enabling" technologies including approaches to providing intrusion detection and central logging at cloud providers. The most important of these was Shibboleth. As a standards-based approach to federated authentication, Shibboleth makes it easy to provide new cloud services to LBNL staff in a secure manner, without exposing login credentials to third party providers.

Finally, LBNL has been a strong participant in the national-level dialogue around cloud policy and risk management, including giving numerous talks and consultations to agencies and research organizations about its risk management strategy for cloud-based systems.

As part of the initial foray into cloud computing, LBNL conducted a business-wide application portfolio assessment to determine the readiness of various applications for movement to cloud models. From that initial portfolio assessment, the lab has piloted numerous applications on various cloud models and moved many services to production on cloud systems.

Cloud Portfolio

LBNL's largest cloud service rollout to date has been Google Apps. Google Docs and Sites was rolled out over 3 years ago, calendar migration was completed in 2010 and, as of November 2010, the migration of mail from an IMAP system is essentially complete.

Together, these applications represent a robust suite of tools for collaboration and productivity. LBNL is on track to recognize savings of $2 million dollars over five years from the switch. But even more importantly, the lab's scientists continue to benefit from the ongoing improvements and extensions to the suite. Scientific collaborations small and large throughout the lab make use of Google Docs and Sites, and use the productivity tools like calendar and email pervasively to enable scientific work. Tens of thousands of Google Docs have been created and shared, and some of the most visited scientific websites at the laboratory are served by Google Sites.

The lab has also used the cloud to extend the cloud—by using features like Google Apps Marketplace, which enables one click deployments of third party applications integrated into existing security and authentication models. LBNL has widely deployed Smartsheet and Gqueues, two collaborative applications for task and project management which have been widely adopted at the laboratory.

Desktop and laptop backups were also moved to a cloud provider, and the lab is currently piloting the move of additional backup systems to cloud systems.

On the business systems side, three SaaS applications have been rolled out including, Point and Ship (for managing shipping), Daptiv (for operations project management), and Taleo, a SaaS Talent Management Application which was the lab's first major business application in the cloud.

On the scientific side, LBNL has been an early adopter of using and evaluating IaaS platforms for scientific computing. IT Division in collaboration with Computing Sciences, conducted numerous tests on Amazon's EC2 services for various scientific computing workloads, and NERSC's report on these tests has been widely cited in the scientific community. ALS physicists Changchun Sun and Hiroshi Nishimura along with LBNL IT staff Kai Song, Susan James, Krishna Muriki, Gary Jung, Bernard Li, and Yong Qin recently explored the use of Amazon's VPC service to transparently extend the ALS compute cluster and software environment, into the public Cloud to provide on-demand compute resources for particle tracking and NGLS APEX development. Their work was presented during the poster session at the International Particle Accelerator Conference earlier this year.

While commercial cloud computing offerings are not yet a good fit for classic high performance scientific computing workloads, the laboratory continues to use and explore a "cloud model" for these services internally as well as explore where commercial offerings can best extend scientific computation services.

Next Steps

As the tools of cloud computing continue to develop, LBNL will continue as an adopter and evaluate these new technologies as they enter the marketplace.

This year, through a partnership with Amazon, IT has been making small grants of AWS credit available to researchers to pilot the use of AWS for specific scientific workloads.

LBNL cyber security has also been testing Google's new Big Query online data processing service to speed up queries on network logs to more quickly locate evidence of malicious behavior. Early testing suggests that the service is capable of providing results for queries that used to take 15 minutes in less than 5 seconds.

Finally, as LBNL enters a phase of refreshing its major enterprise business systems, the lab continues to deeply evaluate cloud products like Workday to understand how these could be part of the future of enterprise resource planning at the laboratory.

Leadership in Cloud

As a uniquely positioned early adopter, LBNL continues to work hard to share its findings with the broader research and education community. Through papers, invited talks, and participation in scientific collaborations, the laboratory has helped to spread knowledge of this emerging trend, in terms of both its promises and limitations, to the wider community. For example, CIO Rosio Alvarez and High Performance Computing Services group lead Gary Jung shared experiences with scientific computing at GovCloud2011. In addition to numerous other talks at meetings and private discussions with other research organizations, Dr. Alvarez also serves as chair of Google's Government Advisory Board and has led panel discussions at numerous government and education focused events.

LBNL has been an early adopter of cloud technologies and remains convinced that, when implemented judiciously, these technologies can be "good for science and good for the planet." Whether enabling scientific collaboration or extending the resources available for scientific computing, LBNL is firmly committed to making use of cloud services and sharing lessons learned with the government, education, and research communities.

LAWRENCE LIVERMORE NATIONAL LABORATORY

Lawrence Livermore National Laboratory (LLNL) is one of DOE's key national security laboratories, administered by the National Nuclear Security Administration (NNSA).

To meet mission goals, LLNL Lab Director Dr. Parney Albright has implemented a "One Lab" strategic philosophy to unite mission objectives and align all lab resources to fulfilling those goals. IT remains a key enabler of these objectives, particularly in driving excellence, cost effectiveness, agility, and ensuring a secure, transformative, and innovative future for the laboratory. Development and production of shared IT services and resources are key enablers of the "One Lab" vision, and cloud technologies will, in time, play an increasingly important role in this space.

Current State

LLNL's approach to adoption of cloud services and technologies has been measured, considering the extraordinary data security requirements mandated for NNSA weapons laboratories. As cloud service providers mature in their security models and execution, particularly in the protected federal cloud space, LLNL is developing initiatives to begin leveraging these emerging technologies to solve business problems.

FISMA-based risk analysis in the federal protected cloud space is maturing as key providers achieve certification. This has enabled critical conversations to begin on what unique business needs may be resolved in properly-protected and risk-mitigated cloud spaces. Those conversations will include migration plans for a very gradual, measured move of certain business functions into protected federal spaces, in particular the DOE and NNSA shared community spaces planned to be offered through RightPath and NNSA Network Vision (2NV) initiatives.

Beginning in Fiscal Year (FY) 10, LLNL invested in commissioning an enterprise-wide private cloud, hosted in the Enterprise Data Center. This capability has since grown to host nearly 400 virtual servers across a variety of business and programmatic functions, with capacity to near 800. The cluster features high availability using vMotion technologies from VMWare's vCenter suite, and is built on the Cisco Unified Computing System platform. Expansion plans include splitting the cluster across multiple facilities to enable geographic separation and stability should an event impact the Enterprise Data Center.

Private clouds have also been stood up in LLNL's mission-oriented principle directorates for specific programmatic purposes, including the National Ignition Facility, the Weapons and Complex Integration principle directorate, and the Global Security principle directorate.

Cloud Computing Vision

LLNL plans to leverage cloud technologies to deliver secure, efficient business value to support the lab's critical national security missions. Services will be sourced primarily from LLNL's private cloud but will leverage external resources when it makes sense—from both a business and a cyber security perspective—to do so. All potential cloud services will be carefully evaluated and selected using key federal standards (e.g. FISMA); any services selected will be managed as an extension of existing IT resources within LLNL.

Key Initiatives

Private cloud capacity expansion. LLNL's enterprise private cloud is architected such that it is extremely responsive to incremental investments. A key LLNL CIO Program strategy is to gradually invest in capacity and delivery expansion on an annual basis to meet increasing programmatic and business needs for virtual infrastructure and platforms.

Infrastructure on Demand (IOD). Launched as a new LLNL CIO Program service offering in FY12, the IOD service features a Web-enabled self-service portal through which virtual servers may be requested. The service leverages HP's Operations Orchestrator suite for workflow automation, and the IOD system executes these automated workflows to deploy virtual servers within about an hour once approvals are obtained.

VDI. LLNL piloted a VDI initiative in FY12 leveraging the Citrix XenDesktop and XenApp suites. FY13 plans include a production service launch of approximately 500 seats, targeting approximately ten commonly-used business applications as part of the hosted and streaming service, in addition to those contained in LLNL's core operating environment image.

Shared enterprise services. Strategic roadmaps include leveraging LLNL's private cloud to expand shared enterprise service offerings such as shared tools (software development, quality assurance, bug tracking, etc.), shared storage, shared data protection services, shared databases, and shared mid-tier technologies, particularly in the Oracle-based layered product space.

Green network data. LLNL is considering outsourcing green network data hosting to secure, federal providers associated with RightPath/2NV. Since green (external) network data has been thoroughly reviewed and released, LLNL's risk assessment of potential outsourcing shows an opportunity to realize efficiency and financial gains with this initiative.

Disaster Recovery in the Cloud. Disaster recovery is often an enterprise's afterthought, and many are caught unprepared when an event impacts the data center. The cloud is expanding disaster recovery service vectors and forcing enterprises to rethink their data protection posture. As cloud service offerings mature, LLNL intends to carefully analyze these for potential use in fortifying the lab's disaster recovery procedures and investments.

Specific application ventures into the cloud. As requirements for time-to-market shrink and business drivers demand increasingly more responsive and agile IT organizations, use of cloud-based application services will dramatically increase. LLNL plans to very carefully analyze such application opportunities from risk-based perspective and employ such services as necessary to fit business functions.

It is no secret that government agencies have not rushed to become bleeding- or leading-edge early adopters of cloud technologies. Caution, with careful emphasis on cyber security, data protection, and efficiencies, all support a measured approach. LLNL's cloud computing vision and direction balances this required caution with future optimism of cloud services becoming a value-adding functional extension of on-site IT resources and capabilities.

LOS ALAMOS NATIONAL LABORATORY

Los Alamos National Laboratory (LANL) embarked on a journey to the cloud three years ago with the launch of IOD, a cloud service broker written as a collaboration between industry and government. The ultimate goal for this project was to develop a cloud to consolidate data centers, speed provisioning, and enhance the laboratory's Green IT posture. The first step was to consolidate where it could and virtualize a sizable portion of the environment.

In 2006, LANL operated with the following servers and applications with its business systems directorate:

- Approximately 300 Intel-based HP Proliant and Dell servers
- Over 32 Web applications which received 10,000 hits daily
- Fifty Citrix servers with 70 applications
- An application portfolio that includes Lotus Notes/Domino, WebSphere, SharePoint, Project, SQL Server, Exchange, and more.

To assist with its consolidation planning, LANL used Novell PlateSpin Recon to gather workload profiles for its Windows and Linux hosts. These physical server utilization metrics were instrumental in determining which servers to target for decommissioning.

To continue closing the gap, LANL used Novell's PlateSpin Portability Suite to migrate physical machines, operating on disparate hardware, to the new virtual platform. An attractive attribute of Novell's product is the ability to not just migrate to a virtual infrastructure, but also the ability to move the virtual machines (VW) back to a physical environment. The enterprise, using Vizioncore vRanger with VMware Consolidated Backup, also backed up the entire VM environment for disaster recovery.

Current State

FROM VIRTUALIZATION TO CLOUD COMPUTING

Only after the virtual environment was running optimally and to best practices, did LANL implement the second phase of the project, which was to create an IaaS private cloud.

This cloud illustrates the following characteristics as defined by NIST:

- On-demand self-service
- Broad network access
- Resource sharing
- Scalable metered services.

ON-DEMAND SELF-SERVICE WEB PORTAL

LANL created a self-service Web portal where system administrators request a virtual server with the push of a button and provide relevant details such as operating system, CPU, memory, and disk space. The monthly cost is dynamically calculated and displayed based on the requirements the system administrator (customer) inputs into the online form. The requirements can be adjusted (for instance, less CPU) at any time and the rate is updated accordingly.

Lifecycle Management is an important component of IOD. Users may request a virtual server for a maximum period of one year, at which time they must renew the system or it will be shut down and decommissioned automatically. At various stages of the notification process, the server moves through an expiration workflow over a 30-day period, generating six notification emails as the system moves from an active state, to shut down, to archived and deleted.

Another feature of the Web portal is the industry's first Green IT Smart Meter, a dynamic calculator that displays the up-to-date green IT savings by computing the amount of energy saved employing virtualization as opposed to deploying physical systems.

BROAD NETWORK ACCESS

Since the portal and servers are Web-based, users may access any of these resources wherever they can connect—on campus or afield.

RESOURCE SHARING

The infrastructure team monitors growth in the infrastructure using VMware's Capacity IQ and Operations Manager to measure growth trends and ensure that there is enough capacity available for its customers. In the event that a large spike in demand constrains resources, prioritization is given per the level of service requested at time of provision: gold tiered systems receive highest priority; silver systems, next priority; and bronze systems, the lowest priority. The cost of these tiers of service range accordingly. SLA selection is one of the on-demand options from which customers choose when requesting the system and can be modified as the needs for the system change.

SCALABLE AND METERED SERVICES

System administrators may change the attributes of their virtual server at any time by logging into the portal. An email follows that sums up the change as well as the new cost chargeback. Cost is based on requested level of service and amount of system resources, not actual used, but requesters can adjust as needed.

Storage

LANL also made an investment to virtualize its storage as part of the new infrastructure by adding NetApp V-series and 2PB of tier 2 (SATA) storage. This investment enables the enterprise to virtualize and manage its existing storage arrays that allows for native integration with vCenter and NFS presentation of storage to vSphere.

Cloud Computing Attributes

To deliver a secure private cloud, LANL deployed VMware vCloud Director, which provides a web-based user interface to consume cloud resources. VMware vShield App and VMware vShield Edge are the tools LANL uses to secure the private cloud as well as other tools that LANL developed itself.

Security

LANL applied the same security policies it employed in its physical environment to the virtual infrastructure as the rules are based on roles. Role-based access to administration and reporting interfaces simplifies the security administration. Provisioning security services is now a lot quicker than in a physical environment as REST-based scripts are employed to place systems in the appropriate enclave. When the request for a server is submitted, a line manager must provide approval before the server is provisioned.

Various servers, desktops and applications operate within LANL's cloud. Desktops are not persistent images and a user accesses a new desktop each time they log in. Roaming profiles ensure that the desktop retains the same look and feel for each user. When a VM moves from one server to another, the security follows accordingly.

LANL extended the traditional government framework for hardware appliances and VLAN's and built the same enclaves in the virtual environment, but also created sub-enclaves for the virtual desktops that are managed by vShield. With vShield, virtual machines are grouped by security policy, with no need for dedicated resources or clusters or with a need to associate VMs to hosts or clusters. Also, not only do policies follow VMs as each move to another host, but other system attributes such as a scaling, load balancing, high availability, DRS, and user-driven changes can be managed dynamically along with the security in the virtualized or cloud-based infrastructure. Any VM that is not acting as expected is automatically shuttled to a remediation enclave where it's fixed and moved back into the system automatically.

Results

The results have been remarkable, with the enterprise meeting its goals more effectively than expected. Namely, LANL achieved the following:

Consolidated its infrastructure

- Physical servers decommissioned: 105
- Data centers retired: 3

Earned return on investment (ROI)

- Estimated time frame for ROI: 2 years
- Actual time frame for ROI: 9 months

Reduced costs

- To date (April 2011), LANL has calculated these savings: $1.4M in cost avoidance and $1.4M in cost savings.

Saved energy and became more green

- Won NNSA Best in Class Pollution Prevention Award for server Virtualization in 2009
- Won InformationWeek 500 Top Government IT Innovators and SANS National CyberSecurity Award for Cloud Security in 2011
- Measured the following direct and indirect energy savings.

Gained added value from the storage investment

- By virtualizing its existing storage, LANL found value in being able to de-duplicate its data in storage, create snapshots of data for backup and restore procedures, and rapidly provision desktops.

Critical success factors

- Collecting physical server utilization metrics was critical in identifying the servers and data centers to decommission. "You don't know what you need if you don't know what you have."
- Spending the up-front time in the virtualization stage, because "you have to have best practices around virtualization before you can take the next step [private cloud]."
- Including lifecycle management as part of the IOD process. Decommissioning is automatic unless the owner takes action.

Lessons learned

- LANL estimated an ROI for the virtualization project in two years, but it realized ROI within nine months.
- Publishing the energy savings and updating in real time contributed to a positive perception of the program.
- Investing in a NetApp storage area network for the virtual environment was deemed a "wise decision" as doing so provided value-added opportunities.

Cloud Computing Vision

The YOURcloud vision based on IOD is to deliver a secure cloud broker solution that will allow multiple organizations to securely consume cloud services across multiple private and public cloud providers.

The NIST Reference Architecture, SP 500-292, defines a cloud broker as an entity that manages the use, performance, and delivery of cloud services, and negotiates relationships between cloud providers and cloud consumers. As cloud computing evolves, the integration of cloud services can be too complex for cloud consumers to manage. In such cases, a cloud consumer may request cloud services from a cloud broker, instead of contacting a cloud provider directly. Cloud brokers provide a single point-of-entry to manage multiple cloud services. The key defining feature of a broker, distinct from a provider, becomes the ability to provide a single consistent interface to multiple differing providers, whether the interface is for business or technical purposes.

A major portion of the project scope is related to standing up a geographically desperate private cloud infrastructure that is owned and operated by NNSA. Cloud infrastructure is essentially the underlying hardware/software resources that provide network, storage, and compute resources.

Key Initiatives

A key aspect of the YOURcloud project is security for both the cloud broker solution and the private cloud infrastructure. The certification and accreditation of the system will be based on NNSA and FedRAMP guidelines. Several key milestones and deliverables have been defined in the project schedule that keeps security personnel closely involved. This allows for potential risks and issues to be detected and addressed early on. The YOURcloud project will leverage centralized security services as dictated by the RightPath IPT.

Integration with the OneNNSA network will allow for a secure transport between an organization's campus networks and the YOURcloud service.

Integration with the planned Federated Identity management solution being developed under the 2NV RightPath umbrella is a critical aspect of improved

The service broker component of the YOURcloud solution requires a custom software development effort using the agile methodology. Several industry solutions were evaluated during the early stages of the project but none had the full feature set and security required by the NNSA.

The NNSA Private Cloud Infrastructure Standup will leverage commercial cloud service providers to accelerate deployment and reduce costs. The commercial cloud service provider will share the operational management responsibilities but NNSA will remain in control of the infrastructure.

Once YOURcloud powered by IOD has been moved into production, the project team will shift its direction towards Migration Assistance. Several labs and plants have already been selected for this process which will include training and technical assistance related to using the YOURcloud service as well as a documented approach for migrating workloads into the cloud.

YOURcloud will also be leveraged by some of the new collaboration services being developed under the 2NV RightPath umbrella.

NATIONAL RENEWABLE ENERGY LABORATORY

Almost all work that the National Renewable Energy Laboratory (NREL) performs is information intensive. Today, NREL opportunities are more dynamic than ever and must be responded to quickly or they will be lost. To meet these needs, it is imperative that NREL IT infrastructure and service delivery be responsive and adaptable to the needs of the lab's user base.

NREL's legacy systems carry a high capital investment cost and those investments require approximately 85% of the lab's IT staffing resources to maintain the current state. To address these issues, NREL's IS support service leadership were tasked with creating five-year strategic plans aimed at innovating and advancing the services and performance of laboratory operations' key functions. Adoption of cloud-based technologies emerged as a key element of the IT strategy, with an eye toward transforming the capabilities and performance of the IT organization. Called NREL Cloud, this strategy represented a dramatic shift in the way IT services were delivered—a shift away from legacy applications and infrastructure developed in-house to a more efficient, secure, and cost effective solution.

Current State

Cloud computing technologies are expected to:

- Improve responsiveness to customer requirements
- Reduce capital and operating costs
- Better manage the IT service lifecycle.

NREL's IT organization incorporated these technologies into the lab's core strategy before cloud computing was recognized with a long history of using SaaS to filter email and within the lab's library system. The goal was to improve the "time-to-value" for IT service delivery, making IT a mission-enabling partner by doing more with less while increasing value. The gains in operational efficiency positions the organization to redirect trained staff to focus on delivering more value to clients.

NREL has a large number of customized on-premise software solutions that take a significant amount of time to show value given their development and maintenance requirements. This leaves the lab's user base discouraged and thereby disengaged with the IT organization. Because of this, NREL is primarily focused on SaaS applications in the cloud computing strategy. SaaS providers remotely host and manage the software and associated data, providing access to the service over the internet from any location and device. This frees NREL's technicians to focus on mission-critical services, rather than the installation and maintenance of software applications.

NREL's first ventures into cloud computing are focused on Human Resources applications for payroll and applicant tracking. NREL currently partners with Ceridian and Kenexa, significantly reducing the costs and mitigating the risks associated with these sensitive functions. By utilizing CSPs who are experts in these areas, NREL's data may be more secure than it would be relying on its own security resources, which isolates risk. NREL also currently uses the Service Desk service management application Service Now and learning management system Success Factors, implemented in 2011 and early 2012.

NREL began moving public-facing websites to the cloud by leveraging the IaaS capabilities of Terremark and Amazon. These CSPs help keep public sites separate from the lab's computing environments, mitigating the security risks inherent in Web-based infrastructures.

SaaS applications also have a huge perceived value to staff. Applications are always up-to-date on the latest versions, eliminating NREL's need to purchase costly application updates and increasing staff productivity as new functionality is released. NREL IT can exceed expectations by providing updates not only more quickly, but automatically.

Cloud Computing Vision

NREL will extensively use cloud technologies to efficiently and effectively deliver both commodity and mission-enabling IT products and services. All services will be sourced on NREL's private cloud, DOE's community cloud, or in the public cloud based on value (where value is a function of utility and cost). To achieve this vision, all commodity IT products and services will be evaluated to determine delivery methods based on value. Cloud service providers will be considered as an extension of IT's capabilities. It is anticipated that many of the commodity IT services will initially be sourced in the public cloud. Over time, it is expected that cloud computing will be NREL's primary source for the delivery of mission-enabling IT capabilities.

Key Initiatives

NREL is moving away from custom solutions for every client requirement—CSPs offer proven solutions, providing a better product at reduced costs. Looking to the future, as more entities move to cloud computing solutions, costs will continue to decrease while efficiencies continue to rise.

NREL plans to use CSPs that adhere to FedRAMP guidelines whenever possible and to leverage the Authority to Operate (ATO) issued by FedRAMP. Because the lab needed to leverage the Terremark IaaS before FedRAMP became available, NREL sponsored the ATO for the needed service from this CSP at significant time and cost. Once FedRAMP and E-RAMP are fully underway, the lab will leverage ATOs across all government agencies for all cloud-based services, reducing costs, and exhibiting financial stewardship of lab and taxpayer resources.

Industry recognizes cyber security as one of the largest challenges to fully realize the benefits of cloud computing. In reality, the security of a service provided by a cloud provider is a chief consideration in the delivery of the service. Therefore, cloud service providers have a vested interest in managing the risks associated with the delivery of the service and protecting its customer's data. By sourcing IT service delivery from multiple distributed and secure infrastructures, risk islands are created that isolate the impact of potential security breaches, lowering NREL's overall risk profile.

To further support NREL's cyber security initiatives, the lab's virtual private cloud will continue to expand, leveraging SaaS, IaaS, and PaaS cloud capabilities. In addition to shifting the IT service management function and learning environments into the cloud, NREL is piloting cloud-based business applications. This move is beginning with the implementation of Microsoft Office 365 user licenses to evaluate the potential of moving desktop productivity tools into the shared services model.

NREL knows that computing is not a one-size-fits-all proposition. To provide a complete cloud computing platform that does not compromise security or quality of service, NREL is developing its Private Cloud 2.0 (PC 2.0). PC 2.0 supports the value of existing applications, while driving transformative innovation across the lab. It will be architected to produce flexibility in service delivery, and improved efficiency and availability, while providing clients with a superior computing experience.

The next generation of NREL's private cloud will not be a monumental shift from PC 2.0. Future iterations will extend existing investments, enabling IT to achieve unprecedented results in an evolutionary manner. NREL will also utilize PC 2.0 to implement a VDI that will provide secure access to applications and data from any device, wherever and whenever the client needs it.

NREL Cloud is a comprehensive, integrated strategy for infrastructure, application platform, and client computing that spans the private and public cloud environments. This initiative is based on standards that can be supported on multiple cloud environments and CSPs, ensuring that NREL's cloud computing environment is portable and scalable, and that CSPs are strategically sourced based on cost, availability, and services provided. NREL has developed a virtual private cloud that is accessible from inside the NREL firewall only, supporting security of data and the network itself.

NREL Cloud will enable IT to fundamentally redefine its relationship with one client base by supporting a focus on client requirements, rather than on providing the technologies needed to support applications. NREL Cloud infrastructure will enable IT to produce services in a self-service model and position itself as a PaaS/SaaS provider to the lab and the IT organization. The strategy will accomplish this goal by providing logical pools of resources and by combining enterprise PC resources with those provided by public cloud providers. When implemented, the result will appear to be near-infinite resources, on-demand, with cost structures and performance levels tailored to deliver value to clients.

The lab has already moved several DOE public-facing Web applications into the cloud and will continue to explore public hosting. These include the OpenEI, Smartgrid.gov and the SmartGrid Data Hub, Solar Decathalon, EnergySavers, Building Technologies, and the Solar Media Gallery.

While it may seem like the lab is taking risks shifting to a model where laboratory data is stored outside of NREL's firewall, the reality is that the lab has been doing this for quite a while. NREL has personally identifiable information that resides in the cloud starting from the moment candidates apply for a job via Kenexa. That reality continues as employees get paid via Ceridian and take courses that are tracked using Success Factors. Additionally, much of the lab's financial data exists in organizations outside of the NREL walls as well—the lab regularly sends data to Washington, D.C., where it moves beyond NREL's control. The lab uses SaaS applications for banking, external transaction processors to administer P-Card transactions, and many vendors have details regarding purchases.

What this adds up to, is that although the lab did not recognize SaaS as a strategy, much of NREL's "private" information has resided outside of the lab's walls for years.

The reality is that external services can and will be much stronger than NREL's own capabilities. The lab does not currently have failover capabilities on critical business systems—a significant risk that will be mitigated by the Cloud Now strategy. Additionally, most SaaS applications store data in the cloud rather than locally on laptops that can be left in taxis. PC 2.0 and virtualization will help IT improve the value provided to the lab by using its resources more efficiently. All of these benefits lead to stronger partnerships with clients across the lab and a level of risk mitigation the lab would not have achieved without the move to cloud computing technologies.

NEVADA NATIONAL SECURITY SITE

In 2003, the Information Services Division at the Nevada National Security Site (NNSS) implemented its first Storage Area Network, separating the disk storage from a physical host. Later in 2005, Information Services Division started the virtualization of over 300 physical servers, separating the operating systems from a physical host. In 2009, Information Services Division started a pilot project for implementation of a VDI, separating an end-user's entire desktop from a physical host. The NNSS has been building on the components that make up cloud computing and plan to keep separating traditional IT services in way that makes them accessible when and where they are needed.

Current State

SERVER AND STORAGE VIRTUALIZATION

Today, the NNSS has a very mature server and storage virtualization infrastructure. Leading with a virtualize first strategy, the remaining NNSS physical footprint remains only to support virtualization and Oracle.

VDI

There are approximately 1200 persistent virtual desktops in use today. This replaced half the physical end-user desktops with zero clients. Enabling the end user's desktop and tools to be available from any work location, travel, or home, enabling a more mobile workforce. Savings for implementing VDI is estimated to be around $800,000 annually. The NNSS has shared its lessons learned and best practices for its VDI implementation with other NNSA labs through conferences and other contacts.

APPLICATION VIRTUALIZATION

Supporting a modular approach to the services utilized by the end user, application virtualization allows us to separate the applications from both the physical and virtual hosts.

Cloud Vision

With virtualization and cloud technologies maturing and the operating system, applications, storage, memory, and processors all becoming independent of each other, a trend has developed in business where the calculation, movement, and storage of data does not have to occur at the desk of an end user or at their company's data center. Instead of distributing a specific amount of processing, memory, or storage power to desks or servers across the company, the NNSS is taking the sum of all that power and centralizing it into a private cloud to make it accessible in a dynamic way to those who need it. The NNSS plans to leverage the RightPath offerings in the future, allowing DOE to also enjoy benefits that come from sharing commodity based IT services. In addition, the NNSS will take a "Cloud-first" approach when evaluating applicable PaaS and SaaS solutions that provide low cost and high value to the business.

Key Initiatives

To provide immense value at lost cost to our customers, certain key initiatives are being focused on this FY and beyond:

VIRTUALIZE THE END-USER EXPERIENCE

Regardless if an employee is technical or administrative, the NNSS wants to provide each employee with the ability to access their data and applications from wherever and whenever they need them. This will be accomplished by continuing the implementation of our current VDI deployment, completing the virtualization of all applications, and introducing a VDI infrastructure that can be used for complex computations and 3D modeling. The end-user experience will be in a "cloud of clouds" and accessible on demand.

NETWORK VIRTUALIZATION

The NNSS is currently working on a project to refresh its entire network infrastructure. As part of this refresh, the NNSS will be taking advantage of one physical network for all types of communication. Similar to server and desktop virtualization, this will allow NNSS networks to be separate from the physical routers and switches, allowing reduced cost and flexibility in a secure manner.

HIGH SPEED BACKBONE

To virtualize the end user experience, virtualize the network, and access cloud resources, a decent amount of network throughput is necessary. The NNSS has focused on installing a 40Gbps backbone between the NNSS and its North Las Vegas location. All user facilities at the NNSS will be connected via high speed fiber connections.

SAAS IMPLEMENTATIONS

The NNSS is looking to evaluate the replacement of its Service Desk, Firehouse, HR, and Financial management systems with SaaS solutions. These solutions provide a lower cost than is currently paid to manage these systems, and would free up valuable resources to focus on NNSS mission specific opportunities.

OAK RIDGE NATIONAL LABORATORY

Oak Ridge National Laboratory (ORNL) is a multi-functional laboratory with science, particularly data- and computationally-intensive science, that spans the full range of confidentiality needs, including data centers that make research data available to the general public, user facilities where researchers from around the world collaborate, fundamental research, export controlled and confidential technology collaborations, and many levels of classified information. Within that environment, cloud computing offers ways to increase organizational agility, adapt to rapidly changing needs and demands, bring in new capabilities, and allow both IT professionals and researchers to focus their time and energy on the aspects of research and business needs which are most critical. Achieving those benefits requires collaboration across a broad range of skills and groups, including the owners of key processes, hardware and software engineers, cyber security personnel, and policy experts.

Current State and Future Work

The initial focus on implementation of cloud computing technologies has been SaaS implementations where the provider brings unique capabilities, has a clear cost advantage, or enables ORNL staff to focus on more core activities and competencies. In some cases, such as external hosting of Web meetings through LiveMeeting and WebEx, the services are largely commodity in nature, and the array of competing offerings drives down the offering costs and ORNL does not need to develop extensive expertise in the underlying technologies to support the services. In other cases, such as SaaS and service partnerships in cyber security intrusion prevention and monitoring, the service providers are both able to bring specific skills and data to bear on the problem and enable cyber security staff to focus on tasks and monitoring more specific to ORNL.

Moving forward, ORNL is evaluating all new IT efforts and many aspects of existing computing services with a “Cloud First” approach. While cost is a factor often considered first in evaluating cloud initiatives, going to the cloud is not necessarily a direct cost savings, particularly when internal services are delivered with a high level of efficiency, are integrated with other internal systems, or are otherwise highly adapted to business processes. Email in the cloud is a particularly interesting example of the trade-offs and analysis. ORNL’s Exchange-based infrastructure is extremely efficient, with a cost per mailbox which is very close to the external price per mailbox of cloud providers. Given that moving to email in the cloud will still have some residual internal costs plus the full external mailbox costs, it is likely that moving to email in the cloud would not result in an immediate or easily quantifiable cost savings. This move is also complicated by the layered email security measures put in place to address the multiple email attack methods, the tools for handling encrypted emails, and the tools for email delivery to mobile devices.

However, a move to the cloud does not have to provide immediate or easily quantifiable cost savings. Moving to the cloud is one means of removing the funding challenges associated with periodic email infrastructure upgrades. It can also eliminate a significant amount of labor associated with server and application maintenance and patching, potentially allowing staff time previously associated with those tasks to more creative and business-value creating activities. Perhaps most importantly, cloud-based services can provide for capabilities and access mechanisms that aren’t practical for internally-hosted solutions, particularly in terms of the Bring Your Own Device trends. On the cost side of the equation, the full integration costs of moving to an external provider must be considered, including the integration to existing business processes, potential changes in those processes, the integration of the external provider with enterprise identity management, changes in risks surrounding the provider, as well as the more traditional migration project costs.

To evaluate email in the cloud, specifically, ORNL has completed an analysis of its email infrastructure, including costs and touch-points with other business processes. ORNL has also worked with Microsoft to encourage their efforts to secure a FISMA moderate ATO for Microsoft Office 365 and has completed a solutions alignment workshop with them to evaluate cost advantages and pain points for differing tiers of usage for Microsoft Office 365. This workshop provided valuable information both for the specific contexts associated with Microsoft Office 365, as well as evaluation questions to use in addressing other cloud migration projects. The evaluation is a particularly good template for evaluation of hybrid approaches, where some services and/or some users may remain internally hosted, while others are moved into external cloud providers.

To better support the research enterprise, ORNL has completed the design and started procurement for a Phase 1 of a research hybrid cloud, which will form the core of a stack of IaaS, PaaS, and SaaS tools to support research and development. This project will improve the speed with which projects can stand up new capability, and it aims to provide a set of tools that will be useful for many aspects of most projects. Recognizing that the needs of research projects can be complex and diverse, this effort is targeting a set of key tools in use so that projects have the choice to stand up their own infrastructure, use cost-competitive common infrastructure, or make use of external cloud services.

As part of the move of the enterprise computing capability to a new data center, ORNL has also procured the initial hardware for a business cloud. This infrastructure will enable extension of ORNL's existing aggressive use of virtualization for enterprise computing into a true cloud model. Additional work is planned both on the hardware/software end as well as from the policy perspective to achieve the full vision, with a hybrid cloud capability using external hosting to enable load shifting, operations independent of ORNL's internal networking, and additional disaster recovery capabilities.

In some cases, such as external hosting of Web meetings or transitioning an internal Office Communicator Server to an external Lync service, the transition to a cloud provider is relatively straightforward. This is particularly the case for services that are relatively independent of other business processes and services, and these transitions can have very clear and immediate cost savings benefits. In other cases, the transition to the cloud may have little initial cost benefit, but provide needed capabilities or agility. By pursuing a Cloud First strategy, ORNL is looking at all major upgrades and new initiatives to determine what types of cloud implementations will provide the greatest value for achieving the lab's mission. Carefully considering cloud options can force a re-evaluation of existing assumptions about the way things have always been done, which in turn can provide value far above a straightforward replacement of an internal service with an external cloud provider. Tools like the solutions workshop can also identify key needs across cloud projects, such as a service-oriented identity management infrastructure, which can then enable a much broader array of business effectiveness improvement projects. Cloud First is not an end, in and of itself, but rather a means by which ORNL will continue to improve its effectiveness in meeting the mission to deliver world-class science.

PACIFIC NORTHWEST NATIONAL LABORATORY

Pacific Northwest National Laboratory (PNNL) has a highly-diversified mission with key initiatives needing rapid deployment of an ever increasing amount of computational resources. PNNL is looking to commercial cloud service providers and internal cloud computing models to effectively distribute the overhead of computing to the providers that can best fulfill the need with the least amount of precious PNNL staff time required to operate it. The PNNL cloud strategy is underpinned by at an objective to reduce the overall staff hours required to operate IT infrastructure that is non-differentiating for the lab and redirecting that human capitol toward the development of innovative infrastructures to further consolidate systems and freeing researcher time for additional research.

Vision

PNNL leadership endeavors to honor PNNL staff as being leaders in their fields of research and to free them from as much of the burden of IT systems administration as possible while providing a diverse portfolio of cost-effective information technology capabilities from which to assemble a system appropriate for the success of any project.

PNNL has identified the following as the key principles to be considered in the pursuit of revolutionizing computing at PNNL with the use of cloud computing:

- Improve the laboratory's capability to deliver differentiating research
- Reduce labor and expense of deploying project IT resources
- Facilitate continued growth while limiting expense of space and operations
- Reduce the ratio of time scientists spend doing IT administration versus research
- Reduce the laboratory's carbon footprint.

In support of these principles the following categories of cloud computing technology were identified as being particularly valuable.

Cloud-based Business Systems—Infrastructure supporting specialized business functions and non-differentiating core applications such as email and the PNNL home page.

Project Infrastructure for Hosted Solutions—Infrastructure for Information Systems developed and hosted on behalf of partners or customers.

Collaboration Zones—Projects that require the collaboration with a disparate group of partners across a wide geographic area stand to benefit from cloud-based solutions that support account federation and geographically distributed content delivery.

Onsite Commercial Cloud Technology—Beyond getting a server, some projects need the power of cloud scalability but their data is not appropriate for commercial clouds.

Commodity Compute Cycles—While PNNL has a tradition of on premise HPC, there is more that can be done to accelerate results by using the capacity of cloud computing providers to support projects that could benefit from scale but do not effectively utilize HPC specific infrastructure.

Commercialization of PNNL Services—The exposure of PNNL-developed systems and services as commercialized SaaS applications. Whether monetized or not, the simplification of proliferation of signature systems would be an asset to the lab and scientific community as a whole.

Current State

PNNL business systems have been consolidating onto a virtual infrastructure for over three years. The systems operating the lab are now over 80% virtualized and housed in 1 of 2 high efficiency data centers, the newest of which is rated at a PUE of 1.18. This infrastructure forms the core of the PNNL on-premise cloud solution. In the last year PNNL Information Management Services has deployed a Self-Service server provision portal that provides a true IaaS capability to any staff member.

This service not only automates the building of a server in a true "cloud" fashion, it also automates the provisioning of an IP address, subscription to online backups, registration with PNNL property tracking, registration with system management and registration with security services.
This automation reduced the time to deliver a server to a researcher from over a week to less than an hour and also improved compliance with asset tracking, security scanning, and system management in general. An onsite cloud

has proven to be an essential component in reducing IT management overhead and getting "research-ready" systems in the hands of our researchers in record time.

PNNL's approach to commercial CSPs is also forward looking. Most institutions currently contract with CSPs on a project by project basis. PNNL identified early on that negotiation of appropriate contracts with CSPs was a lengthy, technical and time consuming task which would be extremely inefficient to replicate for individual projects. So, PNNL IM Services has taken the lead to develop a portfolio of CSPs with pre-negotiated contracts, terms, conditions, security controls, and centralized billing. This portfolio is then offered to all research projects and proposals; allowing for the rapid acquisition of cloud services, requiring only a security review and a charge code. The CSP portfolio comprises services that the research community has requested and will adapt as their needs change; it currently consists of Amazon Web Services and Microsoft Azure. PNNL's contracts with both of these vendors are first-of-a-kind as neither had previously executed an enterprise cloud agreement with a federal entity. As such, these agreements have taken nearly a year to negotiate. And now other laboratories and federal agencies are benefiting from our work and lessons learned. In taking this enterprise approach, PNNL IM services has in essence taken on the role of cloud services broker by partnering closely with cloud service providers to present the simplest, quickest path possible for researcher to utilize cloud computing. This enterprise broker approach has already saved numerous hours of staff time and inspired a new wave of cloud based proposals that can now include cloud services with a high degree of confidence that they will be able to execute their project without the risk of contract issues or provisioning delays.

PNNL is also actively engaged with the identification of point solutions for specific business processes that may be fulfilled efficiently by cloud service providers. Currently, the PNNL Library index is hosted by SirsiDynix and our software developers organize themselves using the Agile Development tools hosted by RallyDev, both offered as SaaS cloud applications. PNNL encourages business groups to seek out cloud delivered solutions whenever possible and has an established protocol for the evaluation of each CSP as well as the information that would be hosted. Additional specialized services are currently in proposal/evaluation stage and we foresee the number of these solutions to increase as they fit well within the strategy as being generally costly to develop and maintain onsite and are non-differentiating for the lab.

In 2011 PNNL started an Institutional Computing program that consolidated much of the HPC capacity that had been distributed across multiple directorates and projects. This program has been highly successful and has delivered to the entire research community an HPC PaaS. The service allows for the purchase of HPC compute time in increments that are acceptable to project budgets and delivers elasticity for projects to run jobs much larger than they could afford individually.

Key Initiatives

PNNL has virtualized nearly every M&O application that can be and is now redirecting energy toward evaluating the needs of research projects in more detail. Some key needs have surfaced and are driving the key initiatives going forward. Future initiatives will give increased emphasis on the value to the research sector.

This successful Institutional Computing project has brought to the fore a need for low cost virtual machines that can access the same datasets and networks that are available to the HPC cluster. PNNL is actively exploring means for driving down the cost of virtualization by exploring alternative hardware and software to support research that needs scalability, and elasticity without the high availability of the current infrastructure.

In the effort to consolidate general purpose servers and HPC clusters, the need to consolidate storage systems was also identified as necessary in order to reduce researcher IT Admin time and to get systems into high-efficiency data centers. PNNL is in the process of deploying an institutional storage service to compliment the Institutional Computing system. This lab-wide service will provide the infrastructure for projects to buy-into and leverage the collective capacity and performance of the overall investment. This multi-tenant storage cloud will have a higher utilization rate, be hosted in PNNL's energy efficient data center, and only require 1-2 FTE to manage multiple Petabytes of data as opposed to the unknown number of staff hours spent managing independent file-systems in low-efficiency data centers or labs.

PNNL is a key stakeholder in the RightPath IPT – Cloud Policy subcommittee. This committee is tasked with the development of a reusable framework from which all DOE and other federal agencies can draft a local cloud usage policy. PNNL's experience in contract negotiations has been invaluable to the effort.

PNNL has partnered with ORNL, NREL, and Savannah River to evaluate the feasibility of moving email services to be hosted by Microsoft's Office 365 service. Together the four partners are working to fully investigate this service for its technical capabilities, its security implications and its financial impact. If successful, moving email to the cloud would improve PNNL's ability to support a global workforce, reduce our overhead of managing an email system, and allow for the redirecting of staff toward differentiating capabilities.

PNNL is tracking the progress of the FedRAMP program and plans to pursue FedRAMP approved CSPs whenever possible. PNNL sees the FedRAMP program as a natural analog to the enterprise broker model currently in place as it will continue to reduce the cost of onboarding new CSPs via the shared ATO model.

PNNL views cloud technology as a means to take a more granular approach to the way individual types of information are handled and hosted and allow for the laboratory to focus more of the talent of its world-class workforce on technologies and research that support its chartered missions. PNNL intends to continue to explore the best use of cloud technology for every project and deploy it whenever feasible.

PANTEX

Pantex has several missions—National Security, Nuclear Material Operations, Nuclear Explosive Operations, and High Explosive Operations. The diversity of the mission is also reflected in the wide range of computing solutions provided to support these missions. In order to more quickly adapt to business needs, Pantex became an early adopter of virtualization, Pantex has shown significant process improvement gains in provisioning, deploying, adapting, and maintaining computer resources. A semi-automated IaaS Cloud was the natural evolution of the Virtualization Strategy. Virtualization and Cloud strategies are not just for process improvement, they are also an integral part of the Pantex Energy Strategy. Pantex recently won the Best in Class award from the NNSA in the category of Comp. Energy &/or Fleet Management. While the award was for several different projects Pantex had in reducing energy use, a key piece of the energy reduction strategy is the virtualization effort.

The Pantex foundational private cloud infrastructure is proving to be a key tool in development of business solutions that will support the combination of the Pantex and Y-12 contracts.

Current State

The Pantex cloud environment is a semi-automated IaaS. All strategies, be they the Cloud or Virtualization are based on common themes:

- Address Business needs efficiently
- Manageable and measurable
- Reduce complexity and cost.

At Pantex, Security and Safety is so ingrained in everything they do, that they are not called out as separate considerations. Private Cloud development has been the focus thereby allowing full control over the security of data.

Cloud Computing Vision

Pantex will leverage the infrastructure in place to fully automate the IaaS cloud. VDI will be the next cloud service Pantex will be focusing on. Specific applications will be streamed to mobile devices and desk tops and will mark Pantex's start of developing an SaaS Cloud. As Pantex needs to support more and more personnel between two locations, Pantex and Y-12 Desktop as a Service will begin to come into play. Over time, Pantex fully expects to integrate the Private cloud model into a hybrid model with Nuclear Security Enterprise cloud offerings like YourCloud.

Key Initiatives

Pantex is either in the process of planning the portfolio/project or actively working projects in the following areas:

- Identity Management providing IdP for both the Pantex Private Cloud along with integration with RightPATH Identity Management initiatives
- Support of NNSA Production Office
- VDI and mobile computing projects to provide SaaS.

Virtualization and cloud strategies have already proved effective in reducing cost and increasing IT responsiveness to business needs, all while also reducing energy usage. For Pantex, cloud strategy is a natural evolution of the strategies already employed over the years.

PRINCETON PLASMA PHYSICS LABORATORY

Cloud computing is a significant trend in the Information Technology world with potential to increase agility, bring value to the organization and lower costs. Princeton Plasma Physics Laboratory (PPPL) is working to include cloud offerings and capabilities where cloud might provide mission and business value for the organization into our overall IT strategy.

In reality, there are numerous tradeoffs between cloud options and traditional computing options. The problem with creating a “cloud strategy” is that, by placing strategy focus on the technology rather than the mission, it’s easy to lose focus and assume that adopting cloud-based solutions is a sure path to mission benefits. Nearly every cloud solution has a functionally equivalent non-cloud alternative, so to maintain focus on the mission and business requirements, PPPL feels it is best to build the strategy around the business decisions to which each type of cloud offering is directed. This approach fosters more level-headed consideration and comparison of cloud and non-cloud options, and it establishes a stronger foundation for a long-term evolution toward cloud and cloud-like options as they mature.

Current State

PPPL began its foray into the cloud+ in 2002, long before the term became popular. Using Enviance for environmental compliance and reporting gave PPPL access to an enterprise application via Web browsers without the need for specialized servers, databases, and IT technical support.

HR services are hosted by Princeton University in a private cloud environment providing PPPL with a tier one Enterprise Resource Planning (ERP) solution at a fraction of the cost.

In 2011, after two years of planning, testing, and pilot analysis, PPPL shifted from its in-house email system to Google Apps Premier. In addition to avoiding the time and costs of a very expensive upgrade project and equipment replenishment, the migration provided capabilities in calendaring and collaboration which PPPL physicists and engineers did not have in the prior system.

To date, the benefits of PPPL’s cloud initiatives has been to provide value for the organization with better service and functionality at less cost in implementation and maintenance over the lifecycle of the application.

Moving Forward

PPPL plans to integrate cloud computing solutions into existing IT strategy where the cloud makes sense from the standpoint of mission, service to employees and collaborators, and is cost effective. As existing applications and equipment flow through their IT lifecycle, replacements and or additions will be considered based on the mission and operational requirements and a comparison of cloud and non-cloud options.

Although PPPL is an ‘open science’ and educational environment, security and data privacy are concerns in the cloud. Cloud Service Providers who are authorized through FedRAMP will be used where possible to facilitate the procurement and implementation process.

To date, PPPL’s major focus in cloud computing has been in the category of SaaS. As cloud and cloud-like options mature, PPPL will correspond the three major categories of cloud computing—IaaS, PaaS, and SaaS—to the three major business decisions of the organizational architecture:

- On what computing resources will PPPL run our operations? (IaaS)
- With what tools will PPPL build and run custom solutions? (PaaS)
- How should PPPL mix custom and off-the-shelf solutions? (SaaS)

At PPPL cloud computing is seen as an enabler. In addition to the usual benefits touted in the cloud such as flexibility, speed of deployment, agility (scalable), and cost effectiveness, the cloud gives a smaller laboratory like PPPL access to technologies that previously were out of our reach financially and allows technicians to focus on providing value added services versus maintenance. Today’s applications will naturally move toward a cloud model as they become more pervasively available through the Web, require more data processing, and span the boundaries of multiple devices.

SANDIA NATIONAL LABORATORIES

Sandia is a multi-program national security laboratory that plays a vital role in ensuring that the United States maintains science and engineering superiority. To continue to meet this role, the office of the CIO, has embarked on an aggressive cloud computing implementation which is guided by a comprehensive strategic plan and roadmap. The strategic plan focuses on supporting mission goals and priorities with effective and responsive information technology solutions. The plan is grounded in practical infrastructure and service delivery projects that will establish the foundation for sustainable, cost effective cloud computing capabilities at Sandia and across the complex. The planning horizon for this framework is three years with elements of it being updated regularly. This timeframe reflects the need for Sandia to keep current with the ever-changing cloud computing landscape.

Current State: Why Cloud?

Sandia IT provides support for traditional core IT capabilities covering a wide range of mission needs (from computing clients to basic infrastructure to high performance computing). However, the high cost and lack of agility in providing these capabilities is hampering IT's ability to assist in accomplishing the mission efficiently and effectively, leading to a do-it-yourself mentality across the organization, even within IT. A transformation to cloud can reduce costs and provide higher-value services agile enough to meet customer needs. Sandia's current infrastructure state includes a significant percentage of virtualized servers and is moving towards an integrated cloud infrastructure. A planned and coordinated effort with significant investment and backing is required to achieve an optimized cloud environment.

Drivers for this change include increasing complexity, rapid changes in the IT industry and in mission programs, heightened security, rising costs, internal and external collaboration needs, recently federal legislation and direction, and the need for technology innovation to bring Sandia IT to the forefront of advanced computing capability. Cloud computing solutions address these needs by taking a services-first, automated, virtualized resources approach, allowing IT to better scale and configure the infrastructure while giving customers greater flexibility, lower costs, and increased access to computing resources.

Vision: Cloud-of-Clouds

Sandia's vision is to establish a "Cloud-of-Clouds" solution to deliver the optimal mix of cloud-based shared service offerings to enable customer success. This vision, whose aim is to modernize and "right-size" IT for the laboratory, will be guided by the following strategic principles:

- Rapid, automated self-service provisioning
- Elastic, usage-based delivery of pooled computing resources
- Usage of commodity resources, open standards and automated processes
- Seamless integration of services, regardless of provider or location
- Reduced footprint and environmental impact
- Secure, ubiquitous Web-based access to services
- Maintain Security and Privacy of data throughout its lifecycle.

Responsive, manageable governance policies will enable our enterprise cloud architecture to reduce internal IT stove pipes and enterprise risk. The result will enable enterprise IT to have greater flexibility without compromising accountability.

Leveraging cheaper processors, faster networks, mobile devices and cloud aware applications will enable Sandia to become an innovator in cloud technology and position it as a service provider of choice to the greater NWC community.

The table below presents a condensed view of the strategic goals and outcomes. Key performance indicators are established to enable management to monitor success and effectiveness.

SNL Cloud Computing Goals and Outcomes

Goal 1:	Enable Sustainable, Cost-Effective Cloud Computing
Outcome:	The future infrastructure will be more agile and delivery greater value. Dynamic scalability and self-healing will support performance, business continuity, and disaster recovery. This will reduce risk, lower costs, and increase operational effectiveness while reducing the IT footprint by supporting data center consolidation.
Goal 2:	Establish and Manage Governance
Outcome:	The future infrastructure will ensure efficient and effective governance by integrating policies and procedures into the service lifecycle. The result will be a standardized, integrated and secure infrastructure enabling greater flexibility without compromising accountability.
Goal 3:	Drive Cloud Technology Innovation
Outcome:	The future infrastructure will place Sandia at the forefront of cloud technology by leveraging advances in networking, virtualization, storage, server and processing platforms, applications, and mobile computing to enable the delivery of required capacity and services when and where needed.
Goal 4:	Operate as a Service Provider
Outcome:	The future infrastructure will enable a fundamental shift in how we serve enterprise IT and mission customers, also positioning Sandia as a complex-wide provider of choice. Our features and service offerings will be delivered in an agile, reliable and secure manner to meet customer requirements.

Key Initiatives: Cloud Roadmap

The cloud roadmap is broken into three phases. The first phase, Core Design and Initial Implementation, establishes the core hardware/software infrastructure design, cost model, provisioning model and governance infrastructure and begins to deliver IaaS capabilities. This lays the groundwork for data center consolidation across the laboratory which is identified as an executive strategic project beginning in FY13.

The second phase, Cloud Service Delivery and Operation, establishes the Cloud-of-Clouds processes required for multi-cloud brokering, orchestration, data protection/security, and enables significant consolidation of data center resources. In addition, it allows for the entire service portfolio to be offered up on cloud resources (e.g., XaaS capability).

The last phase, Cloud Optimization and Integration, establishes a cloud-centric way of thinking, where we utilize cloud-centric tools to design cloud-aware applications, running on cloud infrastructure within a cloud-optimized data center.

The benefits of cloud computing will be realized by establishing goals and outcomes that maximize efficiencies and reduce the cost of providing computing services to IT customers, while reducing the overall IT footprint. Consolidating common services and virtualizing where possible will reduce maintenance efforts and enable sustainable, cost effective cloud computing.

Sandia is committed to working in partnership with DOE, NNSA, and other related agencies and contractors to better leverage the services identified herein and to execute the identified goals. In addition, Sandia will work collaboratively with all stakeholders to determine actions required for mission success and to take positive steps to achieve IT innovation and leadership in the Cloud Computing frontier.

SAVANNAH RIVER SITE/SAVANNAH RIVER NATIONAL LABORATORY

Enterprise•SRS defines a new business direction for the Savannah River Site (SRS). The Savannah River National Laboratory plays a key role in achieving the objectives of Enterprise•SRS. Innovative computing solutions contribute significantly to its continued and future success.

Current State

SRS continues to expand its private cloud for Windows, UNIX and Linux services and systems. Reduced hardware costs and carbon footprint are key drivers, in addition to achieving operational flexibility, improved uptime performance, and enhanced system and application availability.

SRS utilizes services based in the public cloud such as Taleo, BrassRing, Cvent, and MindLeaders.

SRNS is currently implementing and evaluating new operating system technologies such as Windows Server 2012 to enhance its private cloud infrastructure. On-demand rapid provisioning of new virtual servers will ensure SRS has an agile and cost effective computing infrastructure, enabling the innovation and agility necessary to meet the goals of Enterprise•SRS.

Moving Forward

Cloud computing, whether private, public, or hybrid, has the potential to enable Enterprise•SRS activities. Innovative cloud-based computing solutions can provide the capabilities to help build strong business and inter-agency support networks and to broaden stakeholder collaboration securely. For example, hosted virtual desktops (also known as VDI) can be provided which can extend SRS team flexibility to reduce greenhouse gases; enable non-traditional resource provisioning to support collaboration with new and potential partners engaged in the Small Modular Reactor program; and to expand the reach of the National Center of Radioecology.

SRS is transitioning from the traditional IT "build it" concept toward an approach that gives consideration of cloud alternatives priority, aligning with the "Cloud First" policy outlined in the 25-Point Implementation Plan.

The approach balances the policy objectives of "Cloud First" with the cyber security requirements to which we must adhere, all the while acknowledging the challenges imposed by limited financial resources.

SRS continues to:

- Work with DOE and NNSA to architect and evaluate cloud alternatives that have the potential to deliver cost-effective and secure alternatives for commodity IT services such as email, instant messaging, and calendaring
- Monitor implementation of FedRAMP to identify authorized CSPs
- Collaborate with current and potential technology providers to identify cloud technology solutions that will help meet the objectives of Enterprise•SRS
- Engage with current, future, and potential SRS partners and stakeholders to find ways to leverage existing infrastructure and capabilities to promote collaboration and to achieve full value of the national resources at SRS's disposal.

THOMAS JEFFERSON NATIONAL ACCELERATOR FACILITY

When deploying new services or upgrading older ones, Thomas Jefferson National Accelerator Facility (Jefferson Lab) seeks to find the most cost efficient solution that meets the functional requirements of its customers. In recent years, cloud services have offered new options for consideration. Jefferson Lab uses cloud services in a variety of applications, where solutions are available and cost effective for the lab. When analyzing for cost savings, Jefferson Lab looks at licensing costs, maintenance labor costs, internal and infrastructure requirements. Cyber security requirements are also considered. Today, Jefferson Lab has made use of SaaS and IaaS to provide reliable and secure services to their customers while reducing the labor and acquisition costs associated with maintaining an on-site service.

Today, most of Jefferson Lab's recruiting software uses SaaS cloud services, including the marketing tools used to promote and cross-post Jefferson Lab positions (Jobs2Web) and management/tracking of open positions and applicants (Resumeware). Combined, this costs the lab about $70,000 per year, saving the costs to develop and/or maintain a solution hosted on site.

Occupational Health Management software is also a SaaS cloud service. The service manages the lab's medical appointments and related medical information for staff. Jefferson Lab spent $28,000 for this software, a significant savings when compared to what it would have cost to host the software on-site. In addition to licensing costs, Jefferson Lab saves a ¼ full-time employee by not having to provide maintenance on an on-site hosted solution.

The lab's I-9's are tracked using SaaS cloud services from LawLogix. At a cost of about $2,000 per year, this service reduces costs and labor.

Jefferson Lab also uses cloud services such as SurveyMonkey to perform surveys for on-site staff and Open House events and social media services such as YouTube, Facebook, and Flickr to promote the lab's scientific mission and to publish rich media content for public consumption, at little or no cost to the lab.

To improve the lab's cyber security posture, a SaaS cloud service by MxLogic McAfee is used to perform email filtering and virus detection. By utilizing a cloud service, signature updates for viruses and phishing emails are automatically deployed within hours instead of days. This cloud service lowers the lab's overall risk profile.

For backups and contingency planning, Jefferson Lab utilizes an IaaS cloud service from IronMountain to back up critical business data. The backups are secured with encryption and stored at mirrored data centers on disk storage for fast access. Since the data is stored at mirrored data centers, it is available when needed and can be recovered from wherever when needed.

These examples demonstrate how cloud services can and are being utilized at Jefferson Lab to reduce the IT operating costs associated with licensing, maintenance labor, and infrastructure while providing reliable services that do not increase cyber security risks.

SLAC NATIONAL ACCELERATOR LABORATORY

SLAC National Accelerator Laboratory pursues a world-class program of accelerator-based research and fundamental physics. SLAC's LCLS is the world's first high-energy free electron x-ray laser, opening up new frontiers in ultra-fast science (how chemistry really works), materials science and biology. SLAC leads U.S. research into electron accelerator technologies and plays a major role in the world's leading accelerator-based, satellite-based, telescope-based and computational studies of physics, astrophysics, and cosmology.

These programs are both data and compute-intensive and offer opportunities for utilizing internal, external or hybrid clouds or other forms of datacenter, storage, compute, platform, and application virtualization.

For the last 15 years, SLAC has advanced what is now the cloud concept in the service of data intensive science. In 1997, SLAC founded the Particle Physics Data Grid that has now evolved to become the OSG—a community cloud service for U.S. science. The major part of SLAC's scientific computing resources are offered to SLAC science as an OSG-accessible private cloud. Since the advent of commercial cloud services such as Amazon's EC2/S3, SLAC has regularly examined the economics of its in-house and OSG cloud facilities in relation to commercial services. Commercial services are currently uncompetitive in providing the sustained computing used by SLAC science. This is to be expected since much of the computing is data-intensive, and SLAC's in-house computational cloud resources are used from installation to decommissioning at above 80% of their capacity. There is a stronger case for using commercial cloud services to meet sudden peaks in computing need.

SLAC also maintains a full suite of business and enterprise systems to provide back-office capabilities for SLAC employees and users. SaaS is used in several key areas and server-consolidation through virtualization is increasingly used for the system functions that remain in-house.

Mission and Vision

The SLAC Computing Division's role in supporting the SLAC Mission is the following:

"To be the most efficient, customer focused, service oriented ,and capable IT organization to optimally support the laboratory."

Both "most efficient" and "optimally support" requires the SLAC Computing Division to investigate and incorporate cloud solutions and virtualization in the Computing Division's strategic plans.

SLAC is fortunate to have a new CIO and Computing Division Director on board, who was leading the development of the NASA Nebula Cloud Computing Platform. The NASA Nebula Project turned into OpenStack, the world's largest open source cloud platform.

Current State of Virtualization and Cloud Solutions at SLAC

Scientific:

- SLAC pooled and shared compute clusters, in place for over a decade and now a part of OSG
- LDRD proposal to burst the Fermi Space Telescope Pipeline into the Amazon Elastic Cloud
- BaBar physics experiment, operating world's largest database in 2000, now running an in-house, fully virtualized, private cloud.
- Distributed Filesystems - reaching massive throughput (6GB/sec/PB)
- Research and Development into robust, high-performance, multi-site scientific data-access systems:
    - Underpinning of the LSST object catalog database that must scale to 150PB.
    - Xrootd: fault tolerant high performance worldwide access to worldwide data, serving the Large Hadron Collider program.
- Virtualization pilots focused on the needs of particular science activities and on more general testing, login, UNIX services, and build servers.

Business and Enterprise:

- Server: Hyper-V, about 200 virtual servers already, 100 more this year
- Windows Application: Citrix
- Storage: EMC Clariions, Netapp filers
- Network: VLANs, Virtual Switches, VPN, Big5 Load Balancers
- SaaS, PaaS: Taleo Recruit, Drupal, External Forum, WorkSoft's Time and Effort
- Cyber: VMware virtual servers for vulnerability monitoring
- Software Revisions: GIT

Cloud Computing Vision

The SLAC environment is changing:

- Less and less of the IT budget is controlled by the CIO's office
- Almost single-program laboratory years ago to multi-program laboratory today
- IT customers can go out and buy services directly from the cloud
- "Generation Y" has different expectations: work anywhere/any time/on any device
- Commodity activity is not economical for in-house anymore

The traditional SLAC computing principles have to change as a result too:

- IT as a Utility → IT as the Driver of Innovation
- IT the Partner of Choice → IT the Service Broker of Choice
- In-house First → Cloud First
- Do everything → Outsource Commodity, Run Differentiating

SLAC Computing Division key initiatives:

- Research/plan for the use of cloud and virtualization technologies within the various IT roadmaps
- $500,000 IGPE money requested for new technologies, part of it for cloud/virtualization
- Data center virtualization task force established
- More (100-200) virtual servers planned on to the windows side
- More (25) virtual servers going to be installed on the Scientific Computing side
- New virtualized ERP environment design ongoing
- SLAC is looking for cloud collaboration opportunities with other labs

Y-12 NATIONAL SECURITY COMPLEX

The Y-12 National Security Complex is one of four production facilities in the National Nuclear Security Administration's (NNSA) Nuclear Security Enterprise. Y-12's unique emphasis is the processing and storage of uranium and development of technologies associated with those activities. IT within Y-12 not only supports the Y-12 plant mission of production activity, but also supports a general computing environment for email and WEB along with an ERP implementation for business support.

With the constant state of change of technology within the computing industry, Y-12 IT must remain agile to implement new technologies in a secure fashion that meets the requirements of the NNSA that seeks to lower costs, improve security, and enable enhanced communication. Effective alignment of IT strategies and priorities with Y-12's mission objectives remains a primary goal.

Current State

The Y-12 computing environment is composed of central infrastructure designed to deliver the underlying services and functions for business and production computing. The infrastructure is deployed with an appropriate level of redundancy and disaster recovery. Built on this infrastructure are business and production applications deployed to support either the small or large target set of users.

Success has been realized in the management of the computing environment in the areas of effective enterprise integration, work flow simplification, business and process improvements, automation of cyber security requirements, collaboration facilitation, regulatory reporting, and responsive overall infrastructure. Y-12 IT will continue to deliver computing solutions to build on these successes.

Cloud Computing Vision

Y-12 will utilize cloud technology to the extent that it provides efficiencies and capabilities necessary for the mission of the site and its communication with NNSA. It is expected as the technology matures and develops; additional cloud services will be deployed. Y-12 IT will evaluate these new services as they are established to determine their applicability and use.

Y-12 is moving toward utilizing the capabilities of the NNSA cloud implementation, YOURcloud. YOURcloud is one part the 2NV Information Technology transformation strategy. 2NV delivers a plan for enhancing communication capabilities to all sites while maintaining the requisite level of cyber and physical security.

Y-12 IT plans to initially utilize the IaaS service model to implement a private cloud to house servers for enhanced collaboration and information exchange with other sites and site offices.

YOURcloud also has capabilities for SaaS by defining organization wide services. Examples for these services may be the DOE PKI implementation and the ICAM implementation for HSPD-12. Y-12 IT plans to utilize these services as they become available.

Key Initiatives

For the last few years, Y-12 IT has modified its strategy on server purchase and deployment by emphasizing the procurement of blade servers instead of tower and rack-mounted servers in order to obtain the benefits of smaller footprint, centralized management, and reduced overall power and cooling consumption. Tied to this change of emphasis on the hardware, there has also been a change in emphasis the bare metal operating system that the blade servers run. The blade servers are typically configured with a virtualization hypervisor that permits the running of one or more virtual servers on the same physical blade server. Exceptions are only permitted if interfacing with specialized hardware is required or high resource consumption by the system is a prerequisite.

Implementation of internal services, such as email, WEB, and business ERP by Y-12 IT has most of the characteristics and benefits of a private cloud implementation. One notable exception is the ability for resource consumption and chargeback. As Y-12 ventures into cloud implementation, this is the one major change and benefit afforded for the long term.

In addition to server virtualization, another technology that is directly usable within the cloud computing paradigm is VDI. VDI is currently in the pilot phase and is being evaluated for applicability and usability within the current Y-12 IT environment and for functions with high risk and/or high maintenance requirements. A follow-on evaluation will be performed to ascertain its use and viability when running in the cloud environment.

Y-12 IT is well positioned to take advantage of the improved technologies provided under NNSA's OneNNSA strategy. The YOURcloud implementation provides sufficient flexibility to meet changing technology needs along with the required security controls to provide a cloud environment that can be protected at the level that NNSA requires. Finally, using YOURcloud along with the other OneNNSA technologies, efficiencies can be obtained for not only Y-12 but all sites while maintaining an acceptable level of risk.

CLOUD COMPUTING KEY TAKEAWAYS

Before moving toward the cloud, it is important to understand perceived benefits, both short-term and long-term, and to have a cloud strategy in place relating to overall organizational strategy that addresses issues like cyber security and long-term viability. Most organizations will look to integrate cloud computing into their current IT infrastructure, instead of discarding current practices and taking the costly road of starting from scratch.

The following items should be discussed with organization leadership, IT leadership, and help serve as a roadmap to achieving success with the cloud.

1. **Have a plan.** Work across the organization to develop a long-term, viable strategic plan focusing on delivering IT services through the cloud. Create measureable goals; establish priorities; scope; budget; and resources available. As with any large-scale project, know the risks associated and the projected ROI.
2. **Address security concerns.** With any IT computing service, there are typically cyber security concerns. Address these concerns from the beginning. Work with your cyber security experts to understand the risks and develop a plan for mitigating those risks. Test and validate wherever possible. Understand and utilize resources already in place. Though FedRAMP, E-RAMP, and RightPath are continuing to evolve, these frameworks are in place to help accelerate the process of moving to the cloud.
3. **Share successes and missteps.** Utilize the DOE lab and plant community to share successes and missteps when moving to the cloud. Build off successes and collaborate to make the move easier for all parties involved.
4. **Remember cloud services are evolving.** Understand the cloud and its risks and benefits. As cloud computing continues to evolve, know that risks and benefits may change.

Conclusion

Cloud computing is here and a readily available resource for organizations, inside and outside of the DOE Laboratory and Plant system. The examples in this report show the varying degree of movement to the cloud; there is not one solution, but many. All organizations highlighted are working toward incorporating the cloud into their IT strategic plans and DOE is making this transition more efficient by enabling the rapid adoption and usage of cloud services.

ACRONYMS

2NV – NNSA Network Vision

ATO – Authority to Operate

AWS – Amazon Web Services

BNL – Brookhaven National Laboratory

CIO – Chief Information Officer

DOE – U.S. Department of Energy

E-RAMP – Energy Risk and Authorization Management Program

ERP – Enterprise Resource Planning

FedRAMP – Federal Risk and Authorization Management Program

FY – Fiscal Year

HP – Hewlett-Packard

HPC – High Performance Computing

HQ – Headquarters

HR – Human Resources

IaaS – Infrastructure as a Service

INL – Idaho National Laboratory

IoD – Infrastructure on Demand

IS – Information Systems

IT – Information Technology

JLab – Thomas Jefferson National Accelerator Facility

LANL – Los Alamos National Laboratory

LBNL – Lawrence Berkeley National Laboratory

LLNL – Lawrence Livermore National Laboratory

NIST – National Institute of Standards and Technology

NNSS – National Nuclear Security Site

NREL – National Renewable Energy Laboratory

ORNL – Oak Ridge National Laboratory

OSG – Open Science Grid

PaaS – Platform as a Service

PC 2.0 – Private Cloud 2.0

PNNL – Pacific Northwest National Laboratory

PPPL – Princeton Plasma Physics Laboratory

ROI – Return on Investment

SaaS – Software as a Service

VDI – Virtual Desktop Infrastructure

REFERENCES

HP Software Professional Services. 2010. “Enable cloud service strategies by running IT like a business.” HP Software Cloud Consulting Service online, http://h20195.www2.hp.com/V2/GetPDF.aspx/4AA3-3784ENW.pdf Accessed May 20, 2012.

Baker, M. 2009. “An introduction and overview of cloud computing.” Mark Baker’s SSE pages online, acet.rdg.ac.uk/~mab/Talks/Clouds-La-Coruna09/Talk.ppt Accessed May 20, 2012.

Bias, R. 2011. “The evolution of IT towards cloud computing.” Cloudscaling online, http://www.cloudscaling.com/blog/cloud-computing/the-evolution-of-it-towards-cloud-computing-vmworld/ Accessed May 20, 2012.

Bitman, T.J. 2011.“Private cloud computing: emerging from the mist.” Gartner online, http://www.gartner.com/id=1709714 Accessed May 20, 2012.

Bitman, T.J. 2012. “Top five trends for private cloud computing.” Gartner online, http://blogs.gartner.com/thomas_bittman/2012/03/22/top-five-private-cloud-computing-trends-2012/ Accessed May 20, 2012.

Canu, A. 2011. “The history and future of cloud computing.” Forbes online, http://www.forbes.com/sites/dell/2011/12/20/the-history-and-future-of-cloud-computing/ Accessed 6/16/2012.

Computer History Museum. “Timeline of Computer History.” Computer History Museum online, http://www.computerhistory.org/timeline/ Accessed 6/16/2012.

Petty, C and van der Meulen, R. 2012. “Gartner Says Worldwide IT Spending Figures Show Mixed Results for 2012.” Gartner Newsroom online, http://www.gartner.com/it/page.jsp?id=1975815 Accessed 6/16/2012.

Ingthorsson, O. 2011. “5 cloud computing statistics you may find surprising.” Cloud Computing Topics online, http://cloudcomputingtopics.com/2011/11/5-cloud-computing-statistics-you-may-find-surprising/ Accessed May 20, 2012.

Jalona, S. and Chandrakar, A. 2008. “Evolution of IT Services Delivery Model.” Infosys online, http://www.infosys.com/global-sourcing/white-papers/documents/evolution-it-services.pdf Accessed May 20, 2012.

Kundra, V. 2010. “25 point implementation plan to reform federal information technology management.” Chief Information Officers Council online, http://www.cio.gov/documents/25-Point-Implementation-Plan-to-Reform-Federal%20IT.pdf Accessed May 20, 2012.

Mell, O. and Grance, T. 2011. “The NIST definition of cloud computing.” National Institute of Standards and Technology (NIST) Computer Security Division online, http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf Accessed May 20, 2012.

Netmetrix. 2011. “Top 10 cloud computing statistics.” Netmetrix online, http://netmetix.wordpress.com/2011/11/09/top-10-cloud-computing-statistics/ Accessed May 20, 2012.

Stark, C. 2012. “The history of cloud computing.” CETROM online, http://www.cetrom.net/blog/the-history-of-cloud-computing/ Accessed 6/16/2012.

U.S. General Services Administration. The Federal Risk and Authorization Management Program (FedRamp). U.S. General Services Administration online, http://www.gsa.gov/portal/category/102371, Accessed 6/16/2012.

*Images from: iStock/4615257 (cover),
20494874 (page 4), 13836310 (page 8)*

Printed with a renewable-source ink on paper containing at
least 50% wastepaper, including 10% post consumer waste.

U.S. DEPARTMENT OF
ENERGY

DOE/GO 102012-3750 • November 2012