Case Study:
Accelerating Data Center Efficiency

Today's tech-hungry and tech-reliant planet elevates data centers to a status of ultra-critical infrastructures. This is particularly true for government agencies, e-commerce providers, educational institutions and other settings where real-time data analysis and constant uptime are a must. Data center operational costs are rising due to increased computing demand and fluctuating electricity prices. Add to this increased government pressure to improve efficiency and reduce environmental impact [see *Federal Mandates* sidebar, page 4], and data centers are increasingly ripe for state-of-the art operational management strategies.

The OSIsoft PI System™ gathers data from multiple sensors and serves as a universal translator that synchronizes data center operations to boost efficiency, improve planning, and reduce IT and facility costs. With high-performance applications in both the public and private sector, the PI System integrates into a wide variety of systems to track energy use, air handling and cooling, water use, humidity, and even integrated renewables and carbon reporting. Coordinating data in one universal platform allows facility managers to prioritize and monetize every vital aspect of their data center operations. The PI System delivers "value now and value over time," says Steve Sarnecki, OSIsoft's Vice President of U.S. Public Sector Sales. "What we can do is not only identify and capture savings across a system of systems – we can baseline that."

This case study addresses continuous improvement in data center management. Following are three distinct examples of projects where the PI System has provided elegant, customized solutions for critical infrastructure at eBay Inc., Lawrence Livermore National Laboratory and Weill Cornell Medical College.

eBay Inc.

E-commerce powerhouse eBay.com reached a major milestone in December 2012 with more than 400 million listings. According to year-end data, eBay.com sellers transact an auto part every second; a vehicle every minute; a pair of retro sunglasses every two minutes. With its critical technological infrastructure, Dean Nelson, eBay Inc.'s Vice President of Global Foundation Services, recognizes "The foundation of our revenue is our data centers."

In 2010, the company brought online a new flagship data center in South Jordan, Utah, that handles more than one-third of the infrastructure for its business units – Marketplaces, PayPal and GSI. The site was built to earn the U.S. Green Building Council's LEED (Leadership in Energy & Environmental Design) Gold certification [see *Skanska's Mission Critical* sidebar, below] and includes sustainable features such as recycled building materials, natural cooling and rainwater collection for the facility's water chillers.

The PI System provided the only scalable software capable of managing the more than one-half million data streams the company sought to capture at its data centers, according to eBay Inc. Distinguished Engineer Jeremy Rodriguez. From a business perspective, real-time data became a critical enabler in eBay Inc.'s successful efforts to reduce electricity usage and keep energy costs-per-transaction

Skanska's Mission Critical Center of Excellence

Skanska's Mission Critical Center of Excellence, which shepherded development of eBay's green flagship data center to LEED Gold standards, is comprised of mechanical and electrical engineers, estimators and personnel versed in constructing highly technical and sensitive facilities to support critical 24/7 operations.

In 2009 Skanska announced a strategic partnership with OSIsoft to deploy its PI System in data centers. The partnership focuses on three areas: energy, water and infrastructure efficiency; creating feedback loops to drive efficient design and lower development costs; and boosting data center reliability. The eBay project was where Skanska and OSIsoft "proved a concept to come to market with," says John Coster, Vice President at Skanska. The partners took both real-time and historical data and trended it out to predict performance. A system was put in place to maximize efficiency. "Even if you build it perfectly, if you don't manage the data appropriately you're going to miss opportunities," Coster adds. "The PI System is the tool that fills the gap."

Data from the PI System allowed Skanska to develop innovative modular data center solutions that deliver proven savings. Each year U.S. data centers use about 388 billion liters of water, enough to fill 155 million Olympic-sized pools, and 272 gigawatt hours of energy, enough to power 44 million houses, according to Skanska. If each of these data centers used its Modular Solution, U.S. data centers would fill only 31 million Olympic pools and power 27 million houses. It would also reduce CO2 emissions by 67 trillion tons annually.

Clients including eBay, Royal Bank of Canada, TELUS Communications, IBM, Goldman Sachs, AT&T, and Verizon have worked with Skanska to construct new data centers and renovate, upgrade and expand existing facilities.

contained while growing technical infrastructure 32 percent. In addition to collecting environmental, IT and facility data from its electrical and mechanical systems, the PI System collects electrical utility data from renewable energy generation sources eBay Inc. continues to install at its data centers.

In order to integrate real-time data into corporate decision making, as well as compare the data center's performance to original building design requirements, eBay Inc. presents its facilities and IT data as a single, holistic system with an enterprise-wide interface. The PI System monitors 1,000 data center devices with roughly 500,000 data points. Using the PI SDK™, PI Tag Configurator and Microsoft® Excel®, eBay Inc. configured 60,000 tags in 20 minutes. It also built a 200-page web portal using PI WebParts and PI ProcessBook, and integrated with Microsoft's SharePoint® as well as custom Silverlight® applications.

In one pane of glass, stakeholders can see the entire data center – all mechanical systems and electrical usage from the substation down to the server or branch circuit. A select few engineers and executives can remotely review data center operational summaries such as IT load and Power Usage Effectiveness [see *What is PUE?* sidebar, page 6] and perform condition based maintenance. eBay Inc.'s flagship data center in Salt Lake City was designed for an annualized PUE no greater than 1.15, which amounts to "huge savings," says John Coster, Vice President at Skanska Mission Critical. Its Phoenix Modular Data Center solutions are routinely achieving partial PUEs lower than 1.05, even at temperatures as extreme as 119 degrees Farenheit.

With its 60,000 square feet of raised floor, the data center is expandable by four times its current size. Continuous monitoring helps eBay Inc. understand when to add server racks. This sidesteps upfront development of an over-compensating mega-data center and helps align IT costs with business goals as they evolve. "IT can then support the business rather than the other way around," Coster says.

The PI System provides a solution for total cost of ownership that is akin to third-party validation, Coster says, and eBay Inc. would not be able to reach the savings it has achieved without it. "Everybody talks about operational costs to build a project – what about the costs it takes to run it?" he asks.

Lawrence Livermore National Laboratory

Due east of the San Francisco Bay Area in Livermore, California, Lawrence Livermore National Laboratory (LLNL) is focused on national security and the application of world-class science and technology to the important issues of our time. One of 16 national labs in the country, the facility is overseen by the U.S. Department of Energy's National Nuclear Security Administration.

Built ahead of time and under budget in 2004, LLNL's $91 million, 253,000-square-foot Terra Scale Simulation Facility serves as a large data center encompassing numerous classified and unclassified high-performance computing systems that allow scientists to simulate nuclear weapons use, weather patterns, energy-efficiency projects, electrical grid management and myriad other phenomena. In 2009, the Terra Scale Simulation Facility earned LEED Gold certification from the U.S. Green Building Council. Its sustainability plan is largely focused on achieving ambitious gains in power management.

"The LEED process challenged us to find ways to benchmark," said Noah Goldstein, LLNL Spatial Scientist and also a LEED-Accredited Professional. "Because ours is a massive high-performance computing center there was nothing to benchmark against."

In addition to providing space for roughly 250 workers, the Terra Scale Simulation Facility is home to supercomputers Dawn, award-winning IBM Blue Gene/P, and now Sequoia,

a Blue Gene/Q. Sequoia was recently ranked as the single most powerful computing system in the world. Clocking in at more than 16.3 sustained petaflops (quadrillion floating point operations per second), the supercomputer earned the No. 1 ranking on the industry standard Top 500 list of the world's fastest supercomputers released in June 2012 at the International Supercomputing Conference in Hamburg, Germany.

"We bring all of these different data streams and computations into the system and look at them in a common dashboard and figure out what data is significant. All of the individual data streams have a different look to them, but when we put it into the PI System we have a common source of information we can manipulatein different ways."

— Marriann Silveira, LLNL

Sequoia enables simulations that explore phenomena at a level of detail never before possible. Sequoia is mostly water cooled and consists of 96 racks, 98,304 compute nodes, 1.6 million cores and 1.6 petabytes of memory.

According to Marriann Silveira, the Terra Scale Simulation Facility's Lead System Engineer, the PI System has been critical in meeting LLNL's constantly evolving high-performance computing needs. "The PI System helps us understand how our buildings and how our systems operate," she said, citing high-performance computing racks, electrical metering equipment, building management systems and energy utility data.

"We bring all of these different data streams and computations into the system and look at them in a common dashboard and figure out what data is significant," Silveira says. "All of the individual data streams have a different look to them, but when we put it into the PI System we have a common source of information we can manipulate in different ways."

Looking to the future, LLNL is working to achieve the U.S. Department of Energy's goal of exascale computing with several of its high-performance computing systems. If achieved, exascale will represent a thousandfold increase in computing capabilities over petascale.

Federal Mandates

Energy conservation legislation has gained significant momentum in the past several years and placed increased focus on data center efficiency. In 2005, President Bush signed an Energy Policy Act (EPAct) that mandated all federal buildings be metered for energy use by October 2012. Two pieces of legislation passed in 2007 pushed energy-efficiency goals much further:

Executive Order 13423 sets a goal for federal agencies to reduce energy use by 3% each year, and reduce water use by 2% each year, through 2015. It also requires that at least 50 percent of all renewable energy required under EPAct 2005 come from new resources.

The Energy Independence and Security Act of 2007 defined further energy management requirements and amended parts of the National Energy Conservation Policy Act:

- Each agency must improve construction design for federal buildings to reduce energy consumption from 2006 through 2015 (up to 30%).
- Energy managers identified for each facility are responsible for:
  - Completing energy and water evaluations of facilities once every four years.
  - Putting operation and maintenance plans in place, and measuring and verifying energy and water savings.
  - Entering energy use data into a benchmarking system.
- Future projects must include an estimate of the projected energy performance of the building or space and details of what energy-efficient systems or tools will be used.

Executive Order 13514 was signed in October 2009, making reducing greenhouse gas emissions a priority, and requiring agencies to develop comprehensive sustainability plans.

Goals include:

- Measuring and reducing Scope 1, 2 and 3 greenhouse emissions, starting in 2010.
- Minimizing waste and pollution creation, achieving at least a 50% diversion of solid waste by 2015.
- Reducing use and disposal of toxic materials and toxic chemicals.
- Reducing paper use and purchasing at least 30% post-consumer recycled paper.
- Ensuring all new buildings planned after 2020 are designed to achieve zero-net-energy by 2030.
- Developing integrated sustainability plans.

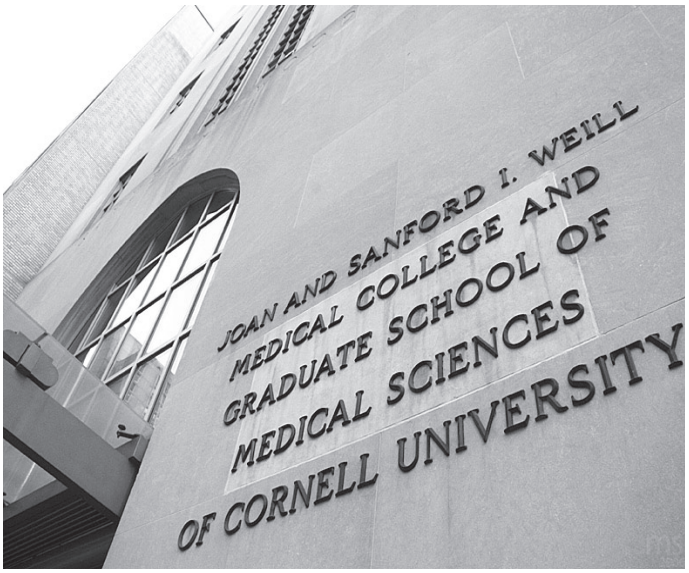

Weill Cornell Medical College

Professors, researchers and students at Weill Cornell Medical College in New York are employing applied mathematics and computer-based technologies that enable physician-scientists to attack complex medical problems formerly beyond their reach. Behind the scenes, state-of-the-art supercomputers are used for advanced modeling, simulations of biological systems, analysis of large data sets, literature mining and imaging.

In developing and integrating technologies to support its Department of Physiology and Biophysics, Institute for Computational Biomedicine and Computational Genomics Core Facility, Weill Cornell Medical College laid out a vision to create a first class computational resource center for one of the top-ranked clinical and medical research institutions in the country.

With a small staff overseeing its research facilities, automation was key. The college also needed to be flexible enough to monitor its IT infrastructure and facility information simultaneously. To keep all of its operations optimized, the PI System infrastructure was installed.

The PI System helps Weill Cornell Medical College synthesize critical data from what might otherwise be seen as disparate sources. For instance, some 2,000 processors conduct more than 40 Tflop/sec (40 trillion floating point operations per second) of computing across the college. Due to rigorous scientific research requirements, standardizing desktop computers is not possible – the college runs Windows, Macs and Linux.

Through custom-built PI Interfaces, data is collected, integrated and disseminated among desktops, servers, printers, switches, building management systems – and any other relevant systems the college wants to monitor.

"The question of, 'Is this affected by that?' can only be answered when you have all data in the same place," says Jason Banfelder, Technology Engineer at the Weill Cornell Medical College.

What Weill Cornell possesses now is a quick, holistic view of key resources. Case in point is its physical plant, which consists of three server rooms between two separate locations in Manhattan and Ithaca that at full load can draw nearly 1 megawatt of energy to run and cool. With the density of computing ever increasing, energy efficiency is critical.

The PI System instructs what nodes can be shut down when not in use to conserve power. Yet while saving energy, it's also critical that computational performance is in no way impacted, and latency is not perceived by end users. A target number of stand-by nodes are therefore ready to run new jobs immediately.

By using the PI System, Weill Cornell Medical College was able to immediately realize an 8% reduction in energy use – and likely double that value when considering related HVAC (heating, ventilation and air conditioning) cost savings – all while maintaining a high service level.

Another benefit comes in the form of the full operational history the facility retains indefinitely. Having this historical information available on call helps drive optimal IT management for the college.

A demonstration of the PI System's interfaces can be found in the halls of Weill Cornell Medical College outside the Department of Physiology and Biophysics, on a large touchscreen called the "Big Board." The dashboard allows stakeholders to see the vital data that Banfelder describes as "real-time, all the time."

Enduring value

Only 15 years ago, data centers had a relatively small footprint. Today they have truly become a nerve center of global enterprise. Costs, regulations and resource constraints are colliding with the rapid growth of new and existing data center developments that are powering today's dynamic technological, civic and business needs, such as those profiled in this case study.

OSIsoft's portfolio of PI System customers operate critical assets ranging from manufacturing facilities to nuclear power plants. The secure data acquisition layers provided by the PI System are especially suited for data centers, where reliability is paramount. And through its membership in The Green Grid, an industry consortium of end-users, policy-makers, technology providers, facility architects and utility companies collaborating to improve the resource efficiency of data centers, OSIsoft remains on the front line of data center innovation.

While some costly vendor-specific data center management solutions can be isolated and proprietary in their approach, the PI System gathers data from multiple sensors and serves as a universal translator that synthesizes data center operations to maximize efficiency, improve planning and reduce IT costs throughout the facility. In realizing these otherwise unattainable costs savings, the PI System reaps expansive data center management benefits for both public and private sector clients as resource and space constraints intensify.

The key, according to OSIsoft's customers, is the extent which the PI System allows critical data centers to baseline for continuous improvement, delivering value now and value over time.

What is PUE?

Power Usage Effectiveness (PUE) is a metric created by members of the Green Grid and embraced as an industry standard to determine the energy efficiency of a data center. PUE is calculated by dividing the amount of power entering a data center by the power used to run the computer infrastructure within it. PUE is therefore expressed as a ratio, with overall efficiency improving as the quotient decreases toward 1.

Driving a data center's PUE as close to 1 as possible is an increasingly popular goal in an industry where every kilowatt of energy recouped amounts to annual savings of about $1,000 as well as a reduction of 5 tons of carbon dioxide emissions.

Energy Reuse Effectiveness (ERE) is now being broken out of the PUE metric, as it skews the PUE reading. ERE refers to re-use of a data center's waste heat – a particularly effective strategy in colder climates.

Data center infrastructure efficiency (DCIE), meanwhile, is the reciprocal of PUE, expressed as a percentage that improves as it reaches 100%.