



OVERVIEW

**Country or Region:**

United States

Industry:

Computational Biomedicine

Customer Profile:

With the creation of the Institute for Computational Biomedicine (ICB), Weill Medical College of Cornell University will realize the full potential of mathematics and computation to enhance the study of medicine. Employing the tools of applied mathematics and computer-based technologies will enable physician-scientists to attack complex medical problems formerly beyond their reach.

Business Situation:

Executive management set a vision for creating a first class computational resource for the research staff.

Solution:

PI IT Monitor infrastructure was installed to enable applications to enhance services levels, optimize operations, and save energy.

Benefits:

- Real-time dashboard now available to research staff to optimize job scheduling
- Historical information available to drive optimal IT management
- An 8% reduction in IT server energy consumption while maintaining a high service level

High Performance Computing

Background

The Weill Medical College of Cornell University has founded a new institute for computational biology and biomedicine. The aim is to bring together a critical mass of theorists and experimentalists to tackle the most challenging questions in systems biology.

In the last decade, there have been major technological advances in this field: the development of large scale methods for recording cellular activity, high throughput tools for monitoring gene expression, rapid screens for measuring behavioral performance, and powerful mathematic tools (statistical and dynamical systems-based) for analyzing and interpreting complex datasets.

These advances open the door to major scientific discoveries and conceptual breakthroughs. The Institute's mission is to provide the resources and nurture the collaborative, multi-disciplinary environment that will make such breakthroughs possible.

Initial Situation and Issues

A group of a 100 researchers run computer models to describe and explore the properties of biological systems at Cornell University's Weill Medical College, Institute for Computational Biomedicine. To service the faculty, students, and staff, there is a computer cluster composed of 100 compute nodes with 2 processors per node. Some models are compute intensive, such as those that explore the way proteins in cell membranes control the signaling that goes on between the complex world inside the cell and the entities outside. Other models are data intensive: for example, post-Genome Project identification of the the functions of the proteins expressed by DNA.

The compute cluster is an essential tool that allows the staff to complete their work. How they use the compute cluster can affect the performance and efficiency of the compute nodes, yet the users had no data on the results of their actions.

The Director, Dr. Harel Weinstein, set the vision to build a first-class computational resource for his staff, enabling Jason Banfelder, System Administration Manager and Associate Technology Engineer, and his team to think about new ways that they can create a compute resource, providing a high level of service. Ideally, what the Institute needed was a method to track service levels in relation to computer operations, connecting IT

The initial challenge was how to provide a holistic view of the key Institute resources. Resources included servers and their utilization as well as server room temperature.

After data center resources were understood and base-lined, future projects could be prioritized.

The Solution: PI IT Monitor Infrastructure

Jason needed end-to-end management tools, but many of these tools are developed for large IT departments, assuming a large IT budget for tools and staff. Ideally, there would be a tool which could be used in modular fashion to phase in service management. The tool Jason eventually found to meet his needs is OSIsoft's IT Monitor, a tool proven to work in data intensive operations like process control data in manufacturing. A lesson learned from Edward Deming, "The key is to practice continual improvement and think of manufacturing as a system, not its bits and pieces" can be applied to IT thinking of IT as a system, and not the components.

Once the infrastructure was identified the first project was the "Big Board" to collect and communicate a holistic view of key system resources.



Big Board on Display at ICB Entrance

The "Big Board" was a proof of concept project that increased and improved visibility of data center system performance, helping IT maintain high services level. After six months, the "Big Board" also became the first point of reference for the users using the ICB to help plan their work more efficiently.

PI IT Monitor has both real time data and historical trending analysis capabilities. The real time data is useful if you know how to interpret the numbers and baselines are set. The historical data provides a comprehensive way to measure change, comparing current operations to the past, and establish new baselines. Both of these data types enable the IT team to close the loop in measuring operations, creating feedback, taking action, and measuring effects of change. The closed loop systems are ideal for automation; enabling quicker decisions to be made with less man power.

As part of continuous improvement, the next project implemented was Cluster Power Management. The project is save energy without compromising compute service levels. The way this was achieved was to power down the compute nodes when they are not needed, saving electricity for the nodes and associated cooling. This functionality has been implemented in closed loop feedback system where job requests can trigger additional resources coming on line, as nodes are idle, commands are sent to power down. PI IT Monitor data enabled this application.

The efforts in power cycling servers have built a closer connection to the facilities department in improving the power and cooling services to the server room. Jason's group is planning future efforts to more tightly manage the server room temperature to avoid overcooling, and detect as early as possible thermal changes increasing the operational risk of the compute cluster.