LBNL-6608E

ERNEST ORLANDO LAWRENCE
BERKELEY NATIONAL LABORATORY

Demonstration of Data Center Energy Use Prediction Software

Henry Coles, Steve Greenberg and William Tschudi
Environmental Energy Technologies Division

September 2013

DISCLAIMER

This report was prepared as the result of work sponsored by the California Energy Commission. It does not necessarily represent the views of the Energy Commission, its employees or the State of California. The Energy Commission, the State of California, its employees, contractors and subcontractors make no warrant, express or implied, and assume no legal liability for the information in this report; nor does any party represent that the uses of this information will not infringe upon privately owned rights. This report has not been approved or disapproved by the California Energy Commission nor has the California Energy Commission passed upon the accuracy or adequacy of the information in this report.

ACKNOWLEDGEMENTS

This work was supported by the California Energy Commission's Public Interest Energy Research program and by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Program, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Special thanks to the key partner, Lawrence Livermore National Laboratory, for providing the facility used for the demonstration and technical support during the project. Specific appreciation to Anna Maria Bailey, Marriann Silveira, Andrew Ashbaugh, and Curtis Tom.

Lawrence Berkeley National Laboratory contracted with Syska Hennessy Group to provide technical expertise and resources, including on-site data collection, software expertise, and reporting support. Thanks to Ian Kucma, the lead Syska Hennessy Group engineer for this project and to Vali Sorell for overall technical guidance.

Thanks to Romonet for providing software technical support and guidance; specifically Liam Newcombe, Jim Rodriquez, and Gavin Cummings.

ABSTRACT

This report documents a demonstration of a software modeling tool from Romonet that was used to predict energy use and forecast energy use improvements in an operating data center.

The demonstration was conducted in a conventional data center with a 15,500 square foot raised floor and an IT equipment load of 332 kilowatts. It was cooled using traditional computer room air handlers and a compressor-based chilled water system. The data center also utilized an uninterruptible power supply system for power conditioning and backup. Electrical energy monitoring was available at a number of locations within the data center.

The software modeling tool predicted the energy use of the data center's cooling and electrical power distribution systems, as well as electrical energy use and heat removal for the site. The actual energy used by the computer equipment was recorded from power distribution devices located at each computer equipment row.

The model simulated the total energy use in the data center and supporting infrastructure and predicted energy use at energy-consuming points throughout the power distribution system. The initial predicted power levels were compared to actual meter readings and were found to be within approximately 10 percent at a particular measurement point, resulting in a site overall variance of 4.7 percent. Some variances were investigated, and more accurate information was entered into the model. In this case the overall variance was reduced to approximately 1.2 percent.

The model was then used to predict energy use for various modification opportunities to the data center in successive iterations. These included increasing the IT equipment load, adding computer room air handler fan speed controls, and adding a water-side economizer.

The demonstration showed that the software can be used to simulate data center energy use and create a model that is useful for investigating energy efficiency design changes.

Keywords: data center simulation, simulation software, data center prediction, data center energy use, predictive modeling, data center modeling, prediction software

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

EXECUTIVE SUMMARY

Introduction

It is estimated that data centers in the United States currently consume approximately 2 percent of the nation's electrical energy. It is assumed this ratio is similar state-wide for California. In data centers, a large part of this electrical energy, often 50 percent or more, is consumed by the power distribution and cooling systems—major components of the infrastructure required to support the electronic equipment providing value to the end user.

A focus in recent years has been to understand and reduce infrastructure energy use by modifying existing data center systems or changing design and/or operational practices.

It would be helpful for data center designers and managers to be able to predict energy savings, as a basis for financial benefit estimates, before undertaking major design or operational changes. The process of estimating energy savings in new or retrofit projects could be streamlined by using a simplified method for predicting data center energy use (requiring less input data than other more involved methods yet still providing valuable results).

This demonstration evaluated the use of a software modeling tool that predicts energy use while requiring few resources associated with data entry or measurement.

At the time of the demonstration, Lawrence Livermore National Laboratory was evaluating design changes for one of its data centers located in Livermore, California. They partnered in the demonstration by offering the use of a 15,500 square foot production data center and by providing staffing support for this demonstration to investigate design changes.

Lawrence Berkeley National Laboratory contracted with Syska Hennessy Group to use the software model, and provide consulting and technical expertise based upon their experience with the modeling software.

Methods

The goals of this demonstration were to evaluate the ease of use of this data center modeling software, determine its accuracy in predicting energy use, and evaluate its usefulness in predicting energy use for mechanical retrofits or IT equipment load changes.

The initial accuracy of the model was determined by comparing electrical power levels predicted by the software to power readings from electrical power meters in place at the demonstration site. The model estimates system component power use by entering the component type into the model.

A more accurate final model (referred to as the "calibrated model") was created by making adjustments to the model inputs after better component information was obtained. It was then used as a baseline to predict energy use for a number of design options and loading scenarios.

The following design or operational changes were studied:

- An increase in the IT computing capability, resulting in a power increase
- Addition of fan speed control for computer room air handlers
- Addition of a water-side economizer

The following hypothesis focused this demonstration's efforts: "Data center predictive modeling software, using simplified data input, can produce easy to obtain results that are useful for energy-use prediction (existing and proposed)."

Results

The initial energy use predictions determined by the software were compared to readings from existing, installed power meters. Some significant variances were found. The model predicted values in conflict with the actual meter readings, and upon investigation it was determined that some meter readings were incorrect. In other instances equipment energy use information entered in the model was not accurate compared to the actual equipment in operation. Updates to the model were made to better reflect the facility operation. The model run with these corrections was called the "calibrated model."

The uncalibrated model provided an overall site input power variance of 4.7 percent when the software results were compared to the initial meter readings. The maximum variation at any single metered location was 10 percent before the model was calibrated. This software will often be used without applying the considerable resources necessary to investigate and resolve variances. Therefore the accuracy of the results are likely to be similar to those encountered during this demonstration for the uncalibrated ("initial") model.

The calibrated model provided an overall site input power variance of less than 1.2 percent when the software results were compared using the corrected meter readings. The maximum variation using that model at any single metered location was 7.5 percent.

After the calibration process, the model was used as the baseline for investigating the energy use performance for a number of modifications being considered.

As Lawrence Livermore National Laboratory was considering a number of energy use improvement modifications, it took site constraints into account and evaluated the following three scenarios:

- Increasing the IT equipment load
- Implementing computer room air handler fan speed controls
- Adding a water-side economizer system

Utilizing the software, the models for the three scenarios were constructed and energy use predictions were available after approximately four person-days of effort.

Conclusions and Recommendations

- Close agreement between predicted and measured energy use appears to be more associated with confirming the site infrastructure topology, component performance, and/or quality of on-site metering than it does with calculations originating in the model software.
- Correction activities to identify variance causes are straightforward but may require considerable resources from appropriate subject matter experts to resolve.
- The researchers did not quantify the software's ability to accurately predict energy use at conditions other than what was used to create the calibrated model.

- This study's hypothesis—that the evaluated modeling software could provide useful energy predictions using relatively simple data input using minimal resources—was borne out in the demonstration.
- To reduce the time needed to learn the proper use of this software, training from the software provider is recommended.
- If a more thorough understanding of this software's ability to predict data center energy use is desired, evaluations comparing actual energy use measurements and predicted energy use with an expanded set of inputs (e.g., weather, IT load) from the same site, as well as from data centers of other configurations, are recommended. This is an opportunity for further research.

CHAPTER 1: Introduction

The United States currently consumes approximately 2 percent of the electrical energy provided countrywide (Koomey 2012). There are no recent studies available, therefore it is assumed California has the same state-wide ratio. A large part of this electrical energy—fifty percent or more by many accounts—is consumed by the power distribution and cooling systems (commonly referred to as *infrastructure*) required to supply and support the information technology (IT) equipment providing value to the end user.

A focus in recent years has been to understand and reduce infrastructure energy use by making modifications to existing sites and to improve designs for new facilities. Before making those modifications, it is advantageous for data center designers and management to estimate the potential energy savings, to help justify the investment.

Without a simple modeling tool, the resources needed to evaluate the impact of a proposed design change can be large, which can be a barrier to accurately estimating the benefits. Therefore, it is desirable to develop a method for simulating data center energy use that simplifies the data needed and thus requires fewer resources compared to traditional methods.

This project, funded by the California Energy Commission's Public Interest Energy Research (PIER) program, demonstrated the use of a software modeling tool that provided energy use predictions while using a simple and effective approach to data entry. The software was developed by Romonet, a United Kingdom (UK)-based software company, with offices in London, San Francisco, and New York.

Lawrence Livermore National Laboratory (LLNL) offered to host the demonstration since it was starting the process of evaluating design changes for one of its data centers in Livermore, California. This involved access to a 15,500 square foot production data center, as well as the lab's personnel resources.

Syska Hennessy Group, a firm experienced with the use of the Romonet modeling software, was contracted to build the model and provide technical expertise for the data center technologies being evaluated. Syska Hennessy and Lawrence Berkeley National Laboratory (LBNL) conducted site visits to gather the required design and operational information and to obtain power measurements to input into the Romonet modeling software.

1.1 Demonstration Goals

This demonstration's goals were to evaluate the ease of use of this data center modeling software, determine its accuracy in predicting energy use, and evaluate its usefulness in predicting energy use in "what if" scenarios. The accuracy of the model was determined by comparing software-predicted electrical power levels to power readings from electrical power meters in place at the demonstration site. The question that leads to our hypothesis is: Does the Romonet software provide valuable energy use prediction while using simplified data input?

Hypothesis: Using simplified input the Romonet software provides modeling results that are useful for datacenter energy use prediction.

Once the initial variance of the model was determined, refinements were made to the model (referred to hereinafter as the "calibrated model"). The major refinement was an updated value for the chiller power consumption. This more accurate value was obtained from the chiller manufacturer after a serial number and years of service were provided. The calibrated model was then used to compare energy use for various design and loading scenarios. In this way, predictions of energy savings were made between the baseline and the design modifications or IT equipment load changes. The goal was to provide data center operators with a model that they could use with confidence to develop a basis for justifying financial benefits of proposed modifications.

1.2 Organizations Involved

A number of organizations were involved in this demonstration. Lawrence Livermore National Laboratory hosted the demonstration in one of their data center facilities and provided facility-related technical support during data gathering activities. Romonet provided technical support associated with use of their software. Syska Hennessy Group provided assistance in site data collection, modeling using the Romonet software, and technical support. Lawrence Berkeley National Laboratory provided data collection, management oversight, and reporting.

CHAPTER 2: Methods

The demonstration consisted of two basic parts. First, investigate the modeling software ease of use and prediction accuracy by developing an initial model and comparing the model energy use predictions to measurements taken from electrical power meters. Second, use the modeling software to predict the energy impact of implementing three different load or design change scenarios.

A Romonet software model was created using documentation provided by the site. This model is referred to as the initial model. Results from IT equipment and PDU readings were compared to the initial model.. The variances found between the initial model results and on-site metering were investigated. The variances found were reduced by adjusting the model inputs using improved equipment performance information and by correcting for faulty on-site power measurements as appropriate. The adjustments from this process created the calibrated model.

The calibrated model was used to predict the energy impact of implementing three different load or design change scenarios. These involved two cooling infrastructure design changes and an increase in IT equipment load. The attributes of the demonstration site are presented below, along with a description of the demonstration details:

- Simulation Model Development
- Model Prediction Variance Investigation and Adjustment
- Simulation of Design Changes

2.1 Demonstration Data Center Description

2.1.1 Site General Description

Lawrence Livermore National Laboratory operates an enterprise data center on the LLNL campus to support the data processing requirements of the laboratory. The data center, shown in Figure 2-1, provides common space and infrastructure to support the computing needs of multiple user groups. It is housed within a 50-year-old building which was retrofitted in 2007 to its current configuration. The single-story building houses 15,500 square feet of raised-floor data center space and 1,500 square feet of administration space. The basement of the building houses the infrastructure to serve the data center and some adjacent office building space.

Figure 2-1: Demonstration Data Center IT Equipment

Source: LLNL Infrastructure and Operations. August 2012

2.1.2 Critical Power Schematic

The LLNL critical power distribution system (Figure 2-2) consists of two (2) 1,000 kilovolt-ampere (kVA) Uninterruptible Power Supply (UPS) modules and twenty (20) 150 kVA Power Distribution Units (PDUs). *Critical power* refers to the power being supplied to the IT equipment. A majority of the IT equipment is supplied with two power feeds from different UPS modules, to provide power supply redundancy in case of a UPS module failure. The remaining IT equipment is supplied with a single power source.

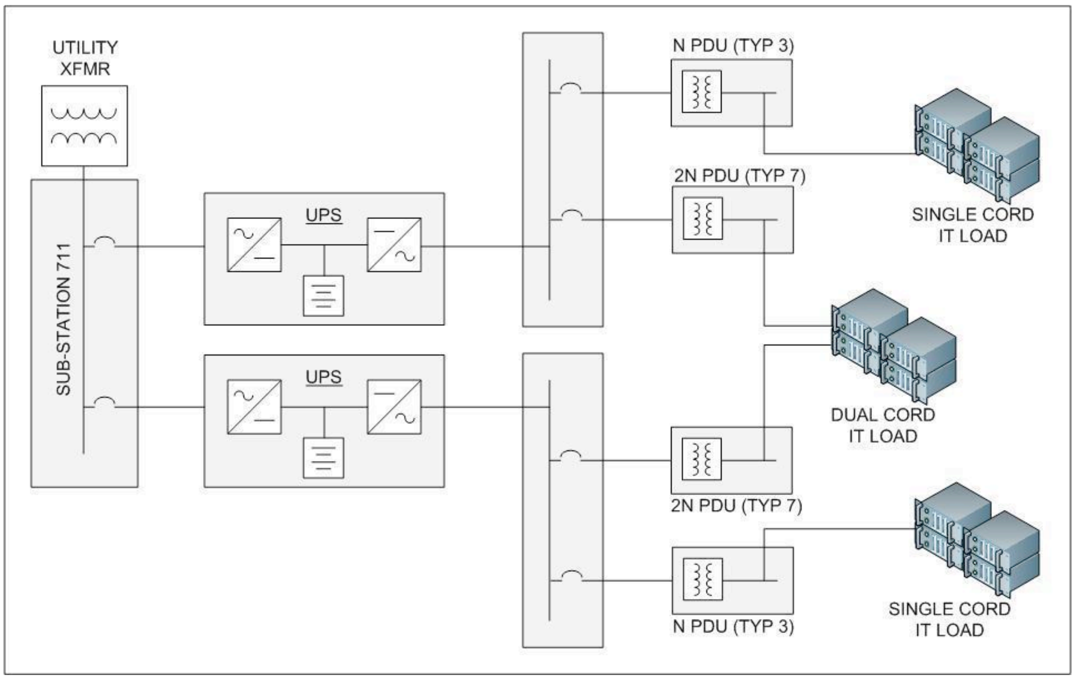

Figure 2-2: LLNL Demonstration Data Center Critical Power Distribution Schematic

2.1.3 Critical Power Capacity and Redundancy

The data center is configured to support 1,110 kilowatts (kW) of IT load with a mix of dual-feed (2N) and single-feed (N) IT equipment power supply configurations. Table 2-1 lists the amount of IT equipment power data center capacity supported by 2N and N power feeds and the mix found during the demonstration period.

IT Loads	IT Equipment Supported by 2N Power (kW)	IT Equipment Supported by N Power (kW)	Total (kW)
Data Center Capacity	630	480	1,110
Demonstration Load	325	7	332

Table 2-1: Demonstration IT Equipment Power Infrastructure Support Type
Source: LLNL Infrastructure and Operations. August 2012.

2.1.4 Infrastructure Power Consumption Schematic

As shown in Figure 2-3, the mechanical systems are powered by two substations: 709 and 710. These substations also feed some non-data center loads such as a portion of the chiller plant that is used to cool a nearby office building. The cooling towers and pumps from the LLNL central condenser water loop (supplying condenser water to the data center's chillers) are not connected to substations 709 or 710.

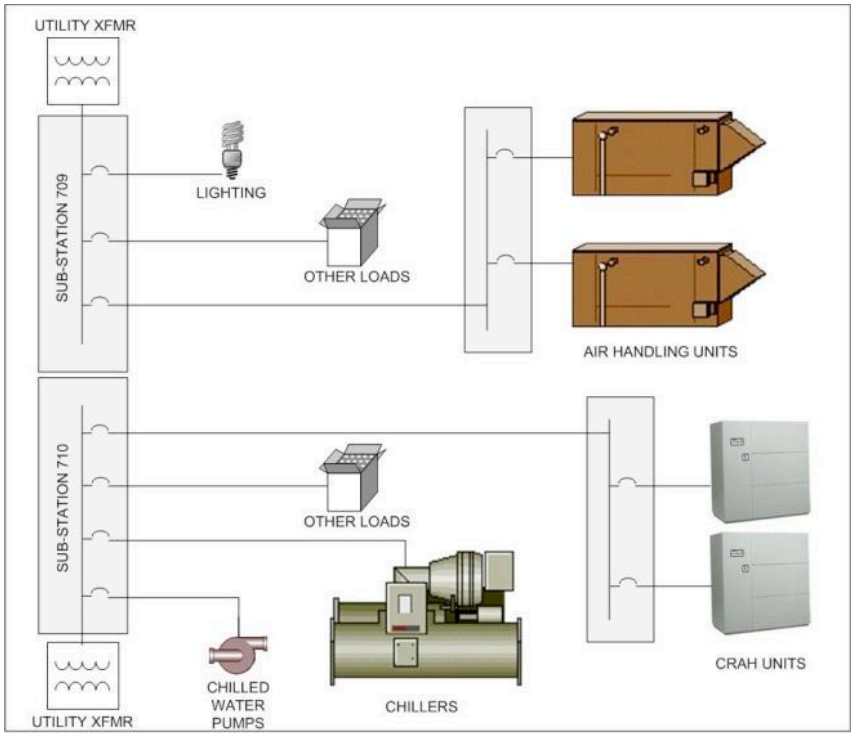

Figure 2-3: Demonstration Data Center Mechanical Power Schematic

2.1.5 Data Center Cooling Schematic

The data center is cooled by thirteen (13) chilled water Computer Room Air Handlers (CRAHs) located on the raised floor and eight (8) chilled water Air Handling Units (AHUs) located in the basement supplying air to the raised floor. Two (2) additional CRAH units are located in the UPS room to condition the UPS and battery systems.

The CRAHs and AHUs are supplied with chilled water from two (2) 400-ton chillers located in the basement of the data center building, shown as one chiller in Figure 2-4. The chillers also supply chilled water to support an adjacent five-story office building. The chillers are connected to the LLNL campus-wide condenser water loop supplied from a central cooling tower system that provides remote heat rejection and pumping capacity.

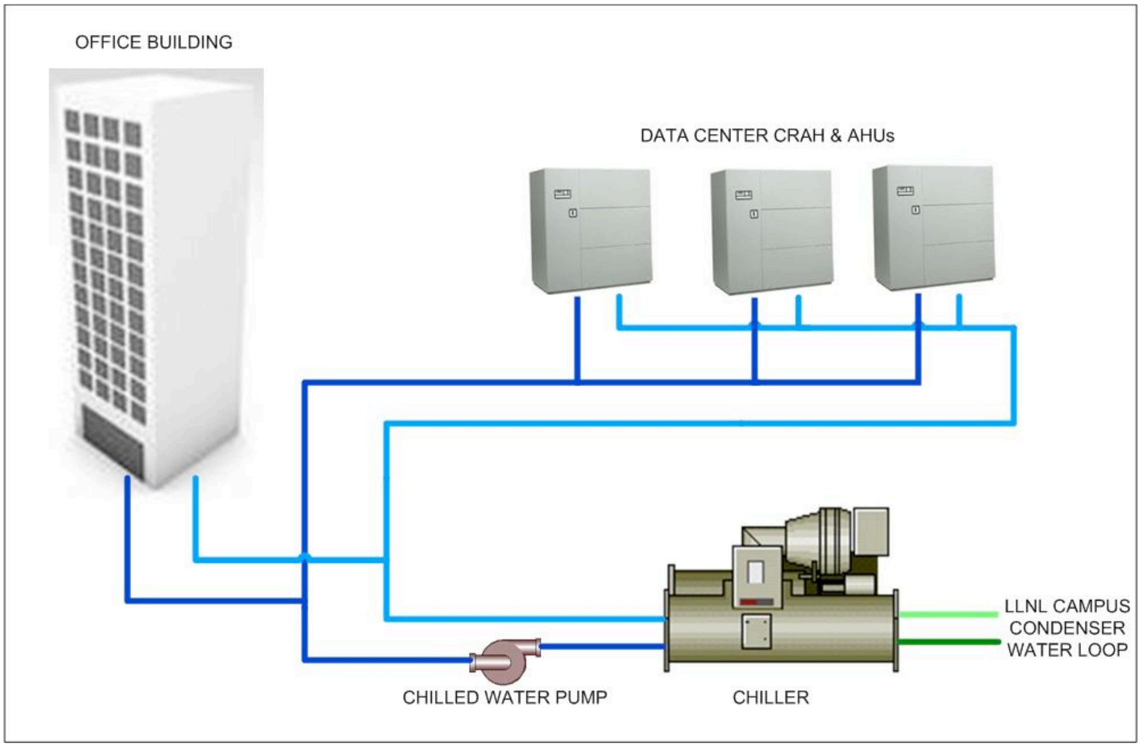

Figure 2-4: Demonstration Data Center Chilled Water Schematic Flow Diagram Representation

2.2 Simulation Model Development

The model for this demonstration was developed by gathering the data center design information, along with a detailed list of electrical power distribution and cooling components. This information was entered into the software to form the basis for a model that provided predictions for energy use performance as a function of IT equipment power, operational settings, and outside environmental conditions. The software provides a large library of generic common data center infrastructure devices and includes a number of IT equipment makes and models along with the energy use performance associated for each. Using the library significantly speeds the process of entering the information into the software because efficiency performance attributes of these devices are contained in the library and do not need to be researched and entered in detail.

This section covers the following subjects:

- Software Suite Description
- Information Survey
- Model Creation
- Variable Inputs and Limitations
- Initial Prediction, Analysis, and Model Adjustments
- Predicting Energy Use from Design or Operational Changes
    - Increase IT equipment power consumption

- o Add computer room air handler fan speed control
- o Install a water-side economizer capability.

2.2.1 Software Suite

The complete Romonet Software Suite consists of five modules:

Portal: a software as a service (SAAS) designed for business and financial management

Enterprise, **Economics**: capital prioritization, return on investment (ROI) and total cost of ownership (TCO) analysis

Energy, **Engineer**: energy analysis modeling

The analysis for this demonstration used two of the five modules: Romonet Energy 2.1.0 and Romonet Engineer 2.1.0.

2.2.2 Romonet Survey Form

To assist with entering technical information needed by the software to describe the data center and supporting infrastructure, a survey form is provided. The survey asks for make and model information for all devices that are a part of the power distribution and cooling infrastructure. Details on the IT equipment are also included in order that the software can calculate the total electrical power consumption and heat generated as a function of selected utilization. In addition, the power and cooling infrastructure connections or schematic designs are recorded.

In some cases, when resources are not available for an on-site survey and follow-up confirmation, the data from the survey form can be used as the only input to the software. In this demonstration, however, we were able to compare the software prediction to actual metered values and then make model input refinements.

The survey asks for information regarding the infrastructure devices, including:

- Quantity of Units
- Manufacturer
- Model Number
- Operational Set Points
- Control Sequence Descriptions
- Power Infrastructure Design Topology (N, N+1, 2N, etc.)

In addition, the survey also asks for information on the quantity and type of IT equipment. Manufacturer and model information can be entered but are not needed if the power consumption of the IT equipment for the data center is known (as was the case in this demonstration).

2.2.3 Model Creation

2.2.3.1 Entering Mechanical and Electrical Power Schematics and Device Attributes

With the help from LLNL site facility personnel, the schematics of the critical power distribution and mechanical (cooling) systems were determined. The heat energy (transferred using air or liquid) and the electrical power energy flow information for the data center and supporting infrastructure were entered into the software. The resulting diagram displayed by the software is shown in Figure 2-5. The red lines represent heat flow from one object to another. In some cases the heat flow is accumulated from one or more objects and split to be cooled by one or more objects as shown in the data center area in Figure 2-5. If the heat from a device is lost to the outside environment or generates no heat this is indicated by a missing red line connected to another device. The heat is removed from each room (indicated by the blue rectangles) via CRAH objects (indicated by propeller icons). The blue lines represent electricity flow starting from the medium voltage switch and ending at the final load. The net result is a schematic that describes the electrical and heat energy flows so that an energy balance is obtained for groups of devices contained in a single room or enclosure.

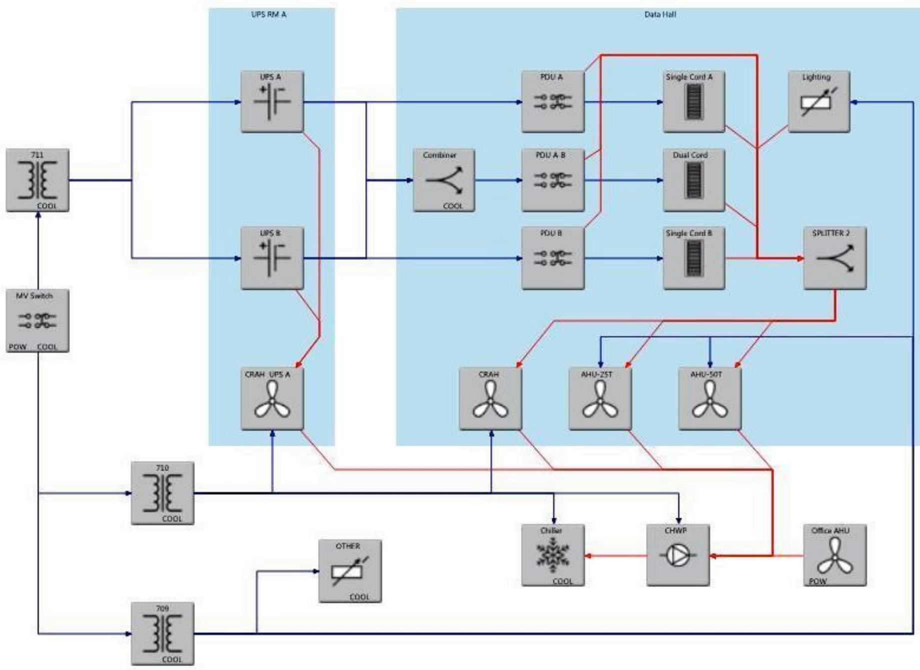

Figure 2-5: (Romonet Screen Shot) Heat Energy and Electrical Power Energy Flow Schematic

The information from the schematic input shown in Figure 2-5 was converted to an energy model pictured and displayed by the software, as shown in Figure 2-6.

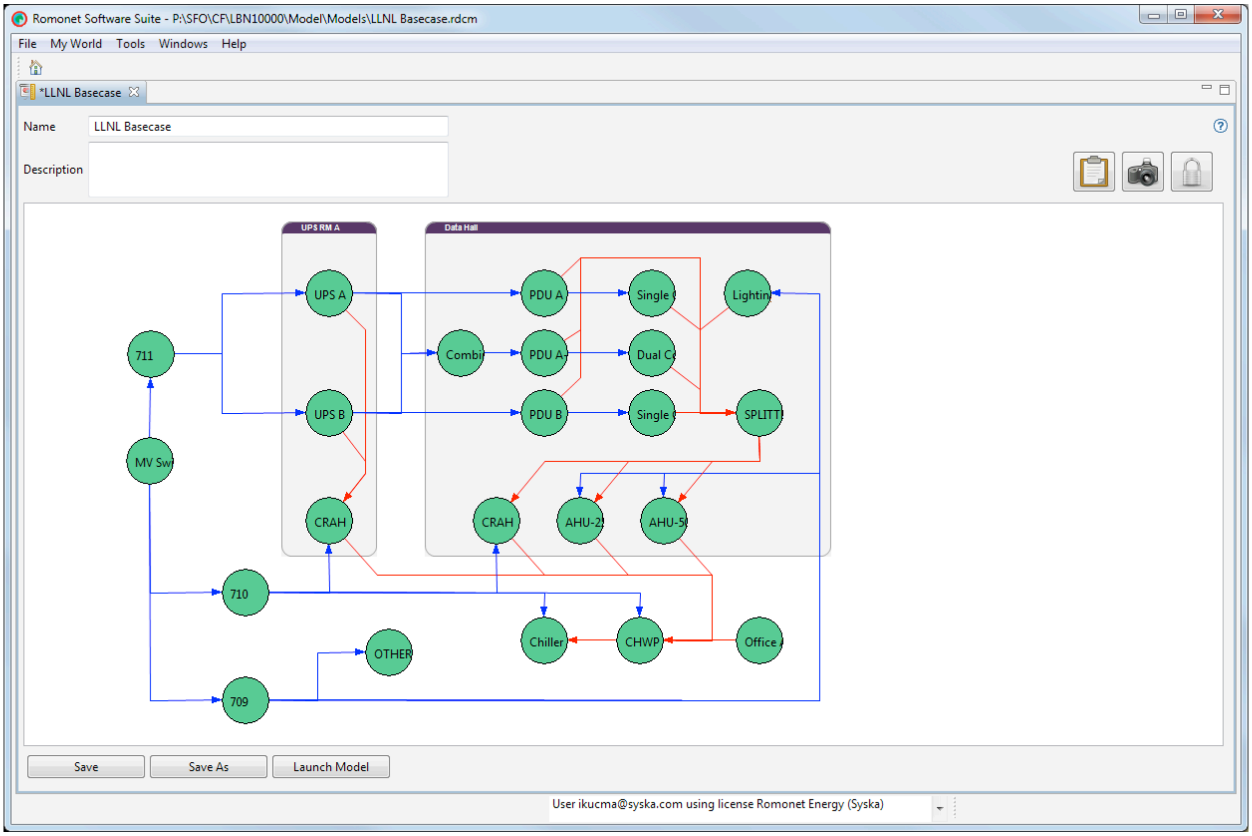

Figure 2-6: (Romonet Screen Shot) Energy Model Image of Existing LLNL Data Center

The performance for each device type in the schematic needs to be identified by selecting from a library of common devices included with the software. However, the library may not contain all the devices required. In that case, the user can create a device and describe the performance using attribute input points provided in the software. Figure 2-7 shows a list, as presented by the included library feature, of a number of chillers with operating conditions that can be selected.

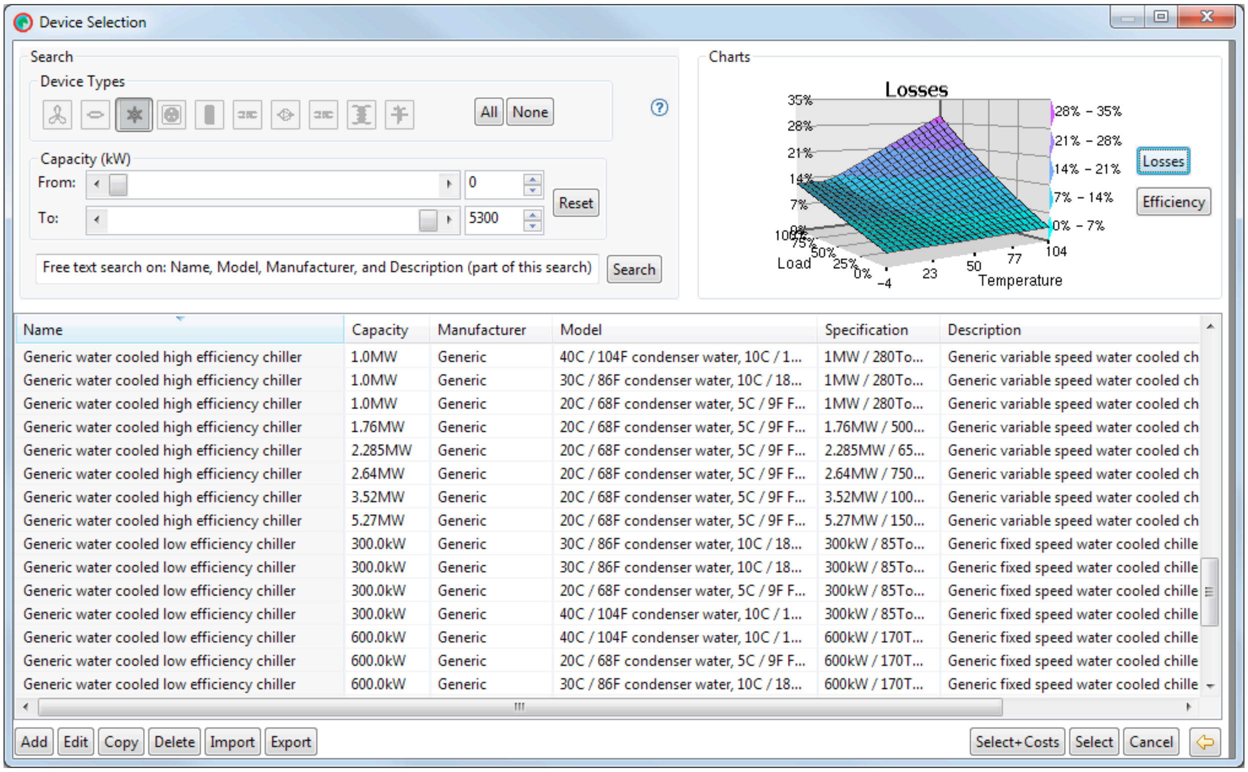

Figure 2-7: Data Center Cooling Component Performance Attributes Are Entered by Selecting from the Romonet Library. Screen Shot of Partial List of Chiller Models Contained in the Software Database Are Shown as an Example

2.2.3.2 Model Variable Inputs and Limitations

The model was used to predict energy use for a single short demonstration period. Three readings from the on-site power meters, flow meters, and temperature sensors associated with the demonstration were collected during August 20, 2012, at 11 am, 2:00 pm and 3:45 pm.

The IT equipment electrical power consumption is a key input variable. The data center power distribution architecture had power distribution units at the end of each IT equipment row. These PDUs were equipped with power meters displaying the power consumed, and meter readings were recorded during the demonstration period. The sum of the IT equipment power indicated at the output of the 20 PDUs was 332 kW; this value was used in all modeling runs.

The outside environmental conditions were recorded at the start and end of the demonstration period. Average readings were used to select a fixed environmental condition for all modeling runs. These conditions were easily selectable using the software.

Not all assignable electrical or other energy uses were included in the model. For example, not included in this analysis were the natural gas energy supplied to the boilers providing steam to the basement air handling units for humidification, the data center share of electrical energy needed for the site cooling tower, and supplied condenser water. In addition, the fuel needed for periodic testing of the diesel generator systems was not an output from the software. The software module used for the demonstration did not have an output predicting the water use of the supporting cooling towers.

2.3 Model Prediction Variance Investigation and Adjustment

The model output contained a number of predicted electrical power flows at a number of locations. For a subset of these points there were power meters at the facility, including those indicated by a star in Figure 2-8. These common locations provided a means to compare the model predictions to corresponding actual meter readings. The common points were:

- Outputs from main transformers supplying the mechanical systems, and IT equipment
- Inputs and outputs from the UPSs
- Power supplied to the chillers

One of the key inputs to the model is the IT equipment energy use. These values were obtained by reading the displays indicating output power from 20 PDUs supplying power to the IT equipment. The IT equipment power measurement points are indicated (represented by three PDUs: A, B, and C) in Figure 2-8, with the star icons to the right of the PDUs indicating that the output power was recorded. For this demonstration, there are no predicted values for IT equipment power, as the actual IT power from PDUs was entered into the software.

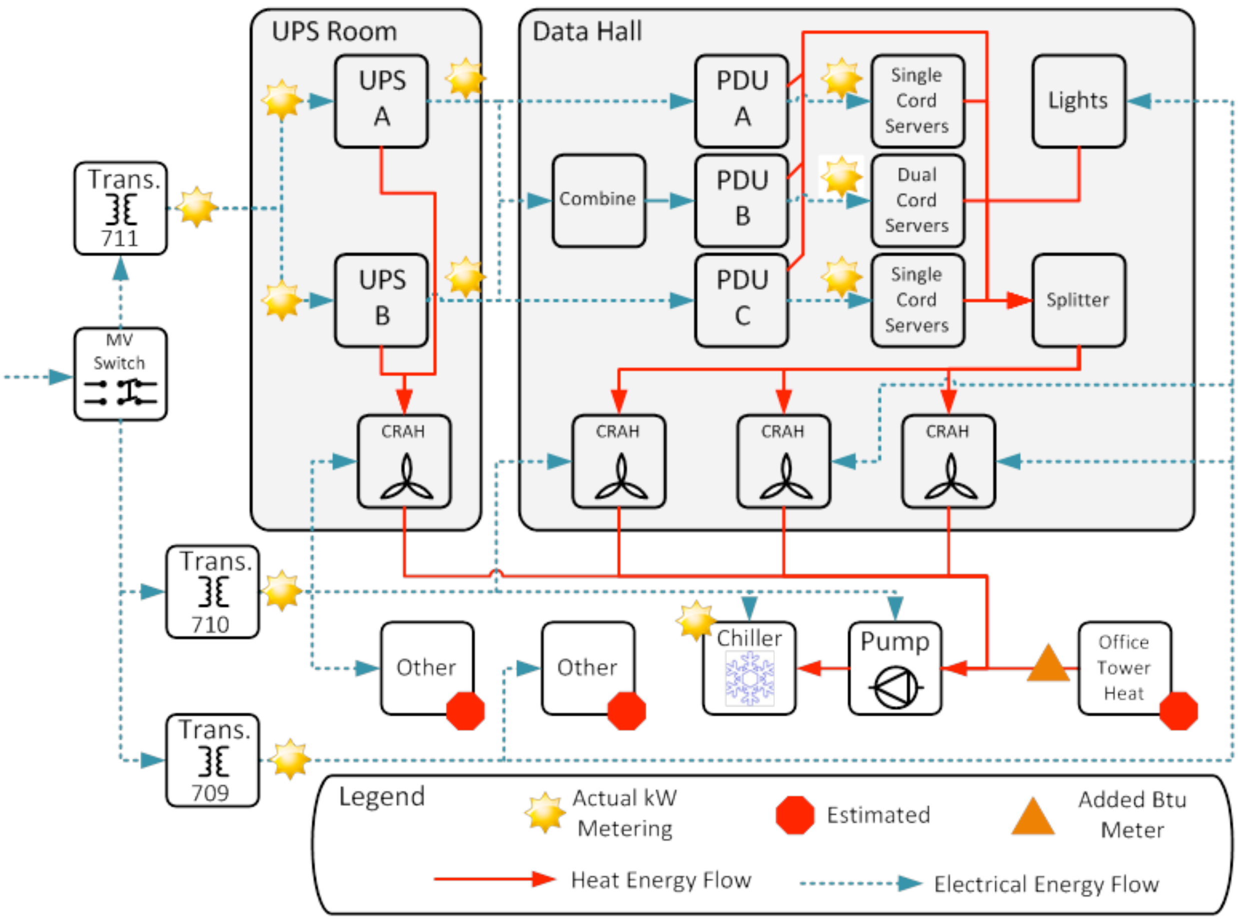

Figure 2-8: Actual Power Meter Locations, Estimated Loads, and Additional Btu Metering

The resources available to investigate the variances found between predicted and measured electrical power values were limited. Details on selected variances and how they were addressed are discussed in Section 3. Reasons for variances included:

- Power meters displayed incorrect values.
- Devices consuming electrical power were not initially accounted for.
- Device performance selected from the software database did not match the actual performance.
- A large cooling load supported by the data center chiller plant (but not supporting the data center) was not easily accounted for.

A second site visit was completed in an attempt to understand the variances and enter updated device performance information or support estimation of more accurate meter readings. This updated information was used to create the calibrated model. The process (predict electrical power values, compare to actual values, and update the software or meter readings) was repeated until the variances were acceptably small.

Details of the results and adjustments to create the calibrated model are described in Chapter 3 (Results).

2.4 Simulation of Design Changes

A goal of this demonstration was to investigate the ease of use of the Romonet software to quantify three proposed efficiency improvements. The calibrated model was used as a baseline. The baseline model was updated with information related to the proposed modifications. Predictions for the three scenarios were then compared to the baseline. The scenarios were evaluated by combining the proposed modifications in the order below. These were the combinations proposed by the LLNL site management.

Three modifications were investigated:

1. Increasing the IT Equipment Load

At the time of the demonstration, the data center had a much lower IT load than the rack space, electrical power, and cooling capacity would allow. There were plans to increase the IT equipment, and the team wanted to investigate how the energy efficiency would change with an increase in load from 332 kW to 1,110 kW. The results from this investigation could be used to justify accelerating the consolidation of IT equipment at LLNL.

2. Adding Air Handling Unit Fan Speed Controls

Another demonstration (Coles et. al. 2012) showed that a significant amount of energy (8 percent of the total site electrical power consumption) could be saved by controlling the fan speeds of the computer room air handlers (CRAHs). The LLNL Data Center Master Plan Team (LDCMPT) was interested to see what energy savings would be expected as part of a payback period analysis. The addition of CRAH fan speed controls would involve adding variable frequency drives to supply power to the existing CRAH fan motors or require that replacement electronically commutated (EC) motor-driven fans be installed. In addition, air temperature monitoring equipment is typically needed as part of the CRAH fan speed control system. Included in the model inputs was the constraint that all CRAH units be operating. The model used for this proposed modification also included the increased IT load, as described above.

3. Adding a Water-Side Economizer

In the baseline design, the cooling tower water (chiller condenser supply) fed to the chillers is supplied by the LLNL campus-wide cooling tower loop. The temperature of this campus-wide loop cannot be controlled to optimize the efficiency of the one data center because it supports many buildings and other services that require a constant temperature.

The most cost-effective method for adding a water-side economizer (WSE)-type feature to support the data center, given the current design, was to add a dedicated cooling tower. With this approach the temperature of the cooling water supplied by the cooling tower can be controlled to save electrical energy and minimize water consumption while keeping the use of the chiller to a minimum. This approach also reduces the electrical power consumption by the chillers because there are more hours during the year where the added cooling tower can supply all the required cooling. The baseline design is shown in Figure 2-9 and the baseline design with the water-side economizer added is shown in Figure 2-10. The model used for this proposed modification also included the increased IT load and CRAH fan speed controls, as described above.

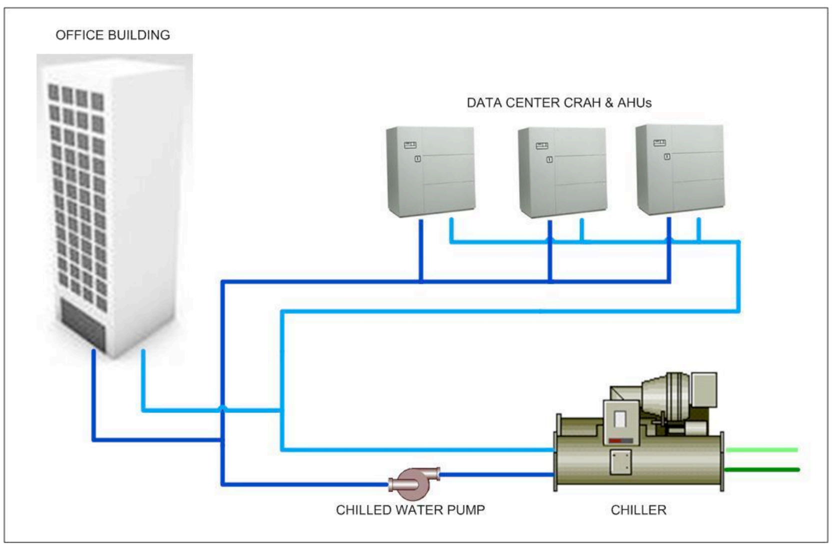

Figure 2-9: LLNL Current Unmodified Cooling Water Flow Schematic (Baseline)

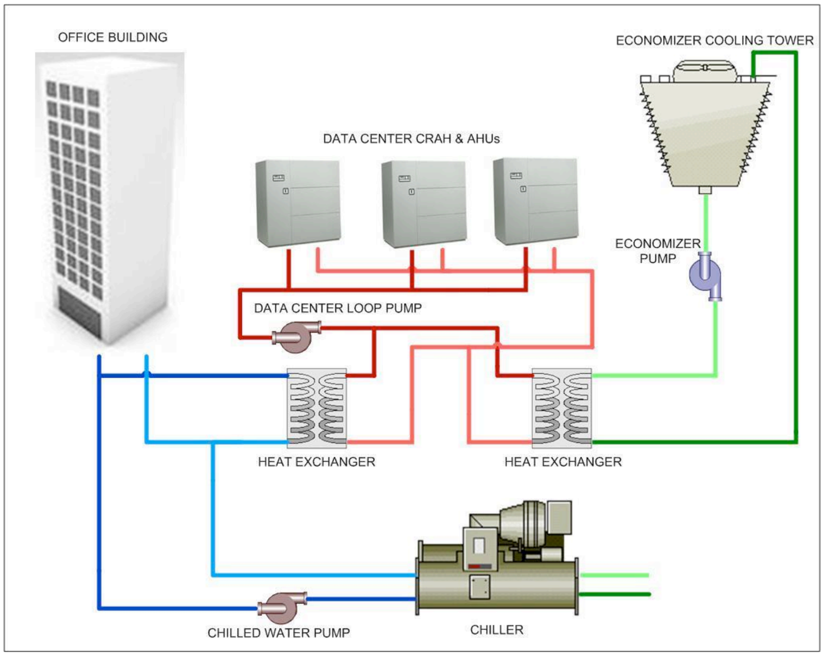

Figure 2-10: LLNL Proposed Water-side Economizer Added Flow Schematic

CHAPTER 3: Results

The following results are discussed in this chapter:

- Quality of initial and calibrated model predictions
- Predicted results for proposed design changes
    - IT equipment load increase
    - Adding air handler fan speed control
    - Adding water-side economizer
- Ease of use and resources required

3.1 Quality of Initial and Calibrated Model Predictions

After the first site visit, the electrical and mechanical systems information was entered into the software, along with the IT equipment power level and outside environmental conditions. At this point the software had enough data to run a simulation. This initial simulation was completed, and electrical power predictions from the software were compared to readings obtained from on-site power meters.

The initial model provided an estimated 697 kW for the overall site power, compared to the measured 731 kW, as listed in Table 3-1. This variance was a negative 4.6 percent comparing the software results to the meter readings. Substation 710 (supplying power to the chillers) had the maximum variation contribution compared to the other substations. The initial model result for Substation 710 was 206 kW, compared to a measured 230 kW, as listed in Table 3-1; this variation was a negative 10.4 percent.

An investigation into the Substation 710 variance indicated that the chiller power estimated by the model was low (113 kW) compared to the measured value of 142 kW. More accurate information was obtained by the chiller manufacturer and entered in the model producing the results for the calibrated model. The calibrated model provided an overall site power variance of negative 1.2 percent comparing the software results (722 kW) to the meter readings (731 kW).

	Measured		Initial Model Results		Calibrated Model Results (change only chiller)	
	Metered	Sub-Metering	Totals	Components	Totals	Components
Substation 709	120		111		111	
Misc.		NM		111		111
Totals	120		111	111	111	111
Substation 710	230		206		231	
Chiller		142		113		138
Pumping/CRAH		NM		93		93
Totals	230		206	206	231	231
Substation 711	381		380		380	
UPS A Input		170		189		189
UPS B Input		191		192		192
Totals	381		380	381	380	381
Site Totals	731		697	698	722	723

NM = not measured

Table 3-1: Measured and Modeled Results
(See Figure 3-1 for a graphical representation)

Examples of the calibration process are described below.

Substation 709

For Substation 709, consumption of 84 kW was initially predicted using the software; the actual meter measured 120 kW. Additional investigation uncovered that an estimated 27 kW was being consumed by devices not listed in the schematics. This error resulted in a low value (84 kW) prediction from the software. When the estimate of 27 kW was added to the software for the missing devices, the variation between the actual reading and predicted value improved to 111 kW vs. 120 kW. A smaller variation could likely have been achieved by completing individual spot metering of some or all devices, but additional resources to attempt this were not within the project scope.

Substation 710

For Substation 710, consumption of 198 kW was predicted using the initial model, while the actual meter measured 230 kW. This large variance was associated with the power consumed by the chiller. The chiller manufacturer was contacted to obtain more accurate performance specifications for the particular chiller and to get information for estimating a fouling factor based on the time in service. After these updates were entered, the variance between the model prediction and actual meter reading for the chiller plant was about 4 kW (138 kW vs. 142 kW), as shown in Table 3-1.

Substation 711: UPS A (Input vs. Output)

The UPS devices display the input and output power levels, where indicated in Figure 2-8. It was observed that the output (180 kW) of UPS A was higher than the input (170 kW). Since this is technically impossible (during normal operation) one or both of the displayed power values must have been in error. Using the UPS vendor efficiency data and efficiency calculated from the other UPS device, a correction was applied to estimate the correct UPS A input value of 188 kW. The model predicted 189 kW, and that was assumed to be correct.

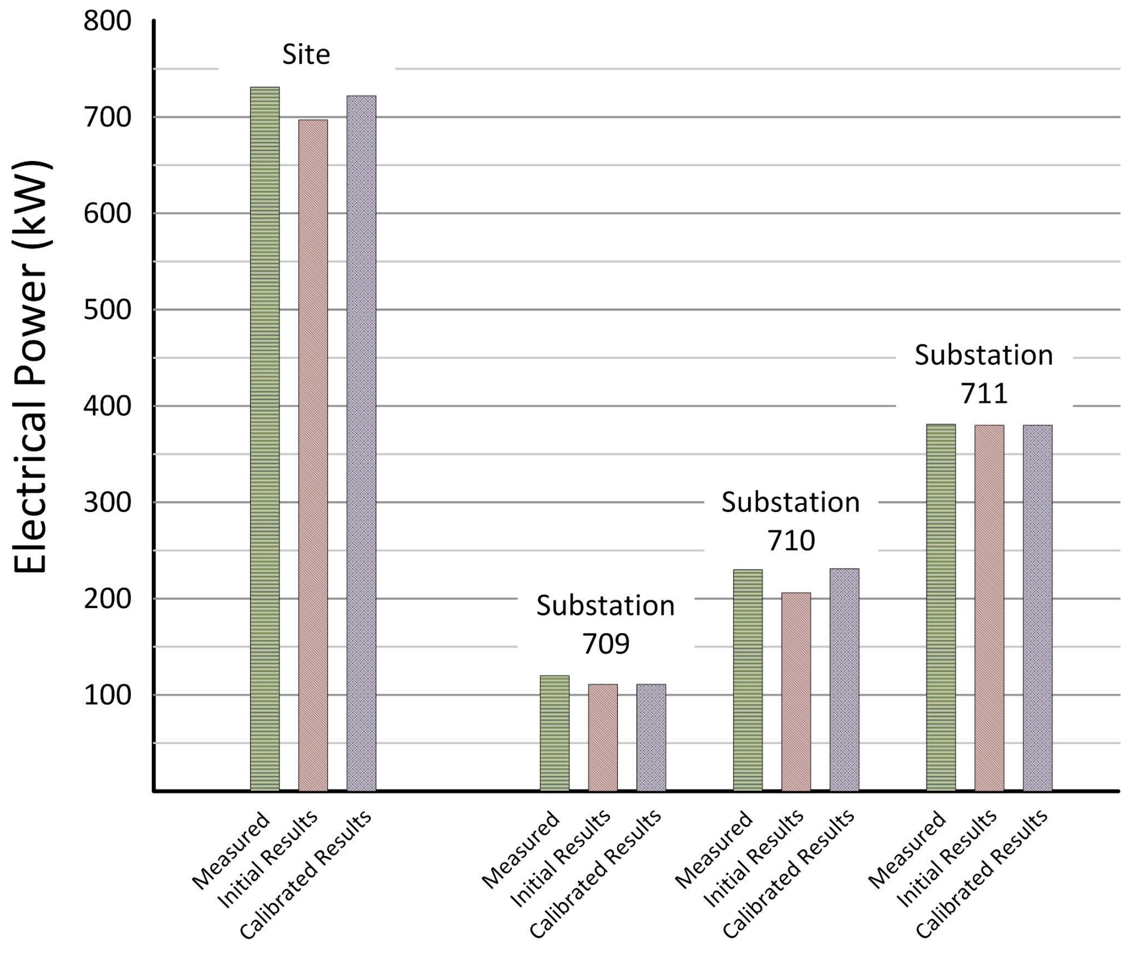

Figure 3-1: Measured and Modeled Results
(Initial and Calibrated Results Are from the Model)

3.2 Predicted Results for Proposed Design Changes

To demonstrate the flexibility of the software, three proposed design or operational changes were compared to the calibrated (baseline) results shown in Figure 3-2. The metric used for the comparison, between the baseline and each proposed change, was Power Usage Effectiveness (PUE). Power Usage Effectiveness is a common data center industry energy-use efficiency metric. It is formally defined as the ratio of total annual site energy to the IT input energy. In this study PUE was the ratio of total data center electrical power to the IT electrical power; lower is better. The PUE used here omits some energy use such as standby generator fuel and supporting office space loads.

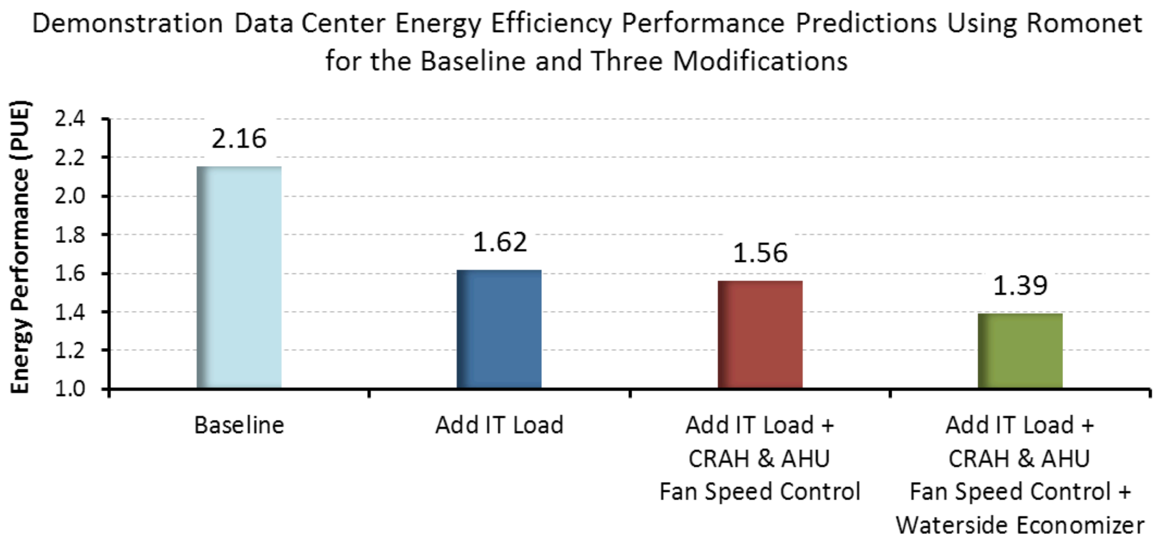

Figure 3-2: LLNL Data Center Energy Use Efficiency Design Change Prediction Comparisons

Adding IT load reduced the PUE from 2.16 to 1.62; although overall electrical load increased, the percentage of infrastructure load was reduced, resulting in a lower PUE value.

The model was then changed by adding CRAH and AHU fan speed control in addition to the IT equipment power increase. The results showed an additional PUE improvement, from 1.62 to 1.56.

The third modification included the two changes already included in the model (added IT load and CRAH and AHU fan speed control) and incorporated the addition of a water-side economizer feature. This change was predicted to make an additional PUE improvement from 1.56 to 1.39.

3.3 Software Ease of Use and Resources Required

A subject matter expert (SME) had used the Romonet modeling software on three previous projects. In addition, the SME completed the recommended three one-half days of training from Romonet.

The effort to gather the initial information and to construct and calibrate the model was 11 person days, as indicated in the top portion of Table 3-2. Data center operations personnel assisted with multiple tours of the facility and filled out the survey form described above.

The changes to the model supporting the proposed design changes required 8 person days, but those changes were relatively straightforward after the design was approved by LLNL. Changing the model for the IT equipment load addition was considerably easier, compared to the other two modifications.

		LLNL (person days)	**Syska Hennessy (person days)**
Model Development	Paper Survey	1	
	Site Visit # 1	1	1
	Develop Base Model		3
Model Calibration	Site Visit # 2	1	1
	Analysis and Model Edits		3
	Sub-Total	**3**	**8**
Investigate Proposed Design Changes	Modify Model: CRAH and AHU Fan Speed Control		1
	Modify Model: Add Water-side Economizer		3
	Analysis		4
	Sub-Total	**0**	**8**
	Total	***3***	***16***

Table 3-2: Resources Required for this Demonstration: Model Development and Calibration and Using the Model to Investigate Proposed Design Changes

CHAPTER 4: Conclusions and Recommendations

4.1 Model Prediction Quality

In this demonstration the variances between the calibrated model's energy use predictions and actual measurements were easily investigated because the load and weather conditions were very similar for both the initial and subsequent calibration readings.

Because most variances found in this demonstration were substantially resolved by correcting inputs to the model or by accounting for incorrect on-site meter readings, the underlying model appeared to be sound.

This modeling software will often be used to build a model for a given data center without applying resources to investigate and resolve differences between the first or initial model predictions and meter readings. Therefore the accuracy of the results using the Romonet modeling software will vary. The largest resulting variation at a particular point for this demonstration using the initial model was approximately 10 percent, resulting in a variance of 4.7 percent overall. This prediction quality is more than adequate for data center energy use analysis. After calibration, the overall variation was even better, at 1.2 percent, with a worst point variance of 7.5 percent.

In all three pre-calibration, high-variance comparisons, most of the variation was traced to incorrect meter display values and the chiller needing specific performance information.

The activities to determine the cause of the variances are straightforward but can require a considerable amount of time from appropriate subject matter experts to resolve.

This demonstration quantified the software's energy use prediction accuracy for the conditions used to create the calibrated model. However, considering that the maximum variance was 7.5 percent among the calibrated model comparison points, there is reason to believe the variance between predicted and actual values for other conditions not tested should be low, and therefore useful.

The hypothesis that data center predictive modeling software, using simplified data input, can produce results that are useful for energy-use prediction (existing and proposed), is true.

4.1.1 Required Resources

Obtaining low variances between the energy use predictions and site measurements is conceptually simple: find where the measured values do not match the predicted values, then either correct problems with the measured values or change the model inputs.

Investigating, resolving, and documenting the variances required five person-days in this demonstration.

In some cases, meter readings were found to be in error, and corrected meter reading values were estimated. In those cases, no model adjustments were made.

Using this software at other data center sites, in the same manner as this demonstration, may require more or less resources, depending on the quality of the inputs. These inputs include: correctness and quality of the supplied electrical and mechanical systems information and accuracy of the electrical power readings from components or power meters.

It is recommended that training be obtained before attempting to use this data center modeling tool. Considering the number of possible technical issues, a person attempting to use the energy-related modules should have considerable knowledge and experience with the infrastructure components that typically support the data center being modeled.

4.1.2 Recommendations

A suggested area for future study is to evaluate the software with an expanded set of inputs (for example, IT equipment load and environmental conditions) at the same site, relative to what was used for this demonstration.

In addition, it is suggested that the software be demonstrated at a number of data center sites of varied infrastructure design, to obtain additional empirical data to validate the prediction accuracy and ease of use over a wider application.

An additional opportunity for future research would be to compare the actual IT server power to the values predicted based on the IT equipment library provided with the software. This would provide confidence in predicting data center energy use when the actual server power is not easily measureable or when the server models are known in the case of a future deployment.

References

Coles, Henry, Steve Greenberg (Lawrence Berkeley National Laboratory), and Corinne Vita (Vigilent). 2012. *Demonstration of Intelligent Control and Fan Improvements in Computer Room Air Handlers*. Publication number LBNL-6007E.

Koomey, J. 2012. Data Center Energy Use is Moderating. Data Center Knowledge. Retrieved January 30, 2013. www.datacenterknowledge.com/archives/2012/10/15/koomey-data-center-energy-use-moderating/(last accessed 1/30/2013).

Glossary

AHU	air handling unit
Btu	British thermal unit
CRAH	computer room air handler
gpm	gallons per minute
IT	information technology
kW	kilowatt
LBNL	Lawrence Berkeley National Laboratory
LLNL	Lawrence Livermore National Laboratory
MV	medium voltage
PDU	power distribution unit
PUE	power usage effectiveness
ROI	return on investment
SAAS	software as a service
SME	subject matter expert
TCO	total cost of ownership
UPS	uninterruptible power supply
VFD	variable frequency drive
W	watt
WSE	water-side economizer