2009

Computing laboratory sustainability & utilization: initiatives for a greener education

Kristian Stokes

COMPUTING LABORATORY SUSTAINABILITY & UTILIZATION:

INITIATIVES FOR A GREENER EDUCATION

by

Kristian Stokes

Thesis submitted in partial fulfillment of the requirements
for the degree of
Master of Science in
Computer Security and Information Assurance

Department of Networking, Security & Systems Administration

Rochester Institute of Technology

B. Thomas Golisano College
of
Computing and Information Sciences

August 2009

Rochester Institute of Technology

B. Thomas Golisano College
of
Computing and Information Sciences

Master of Science in
Computer Security and Information Assurance

Thesis Approval Form

Kristian Stokes

COMPUTING LABORATORY SUSTAINABILITY & UTILIZATION:

INITIATIVES FOR A GREENER EDUCATION

Thesis Committee

Name	Signature	Date
Dr. Sumita Mishra, Chair		
Dr. Yin Pan, Committee Member		
Dr. Bo Yuan, Committee Member		

Thesis Reproduction Permission Form

Rochester Institute of Technology

B. Thomas Golisano College
of
Computing and Information Sciences

Master of Science in
Computer Security and Information Assurance

COMPUTING LABORATORY SUSTAINABILITY & UTILIZATION:

INITIATIVES FOR A GREENER EDUCATION

I, Kristian Stokes, hereby grant permission to the Wallace Library of the Rochester Institute of Technology to reproduce my thesis in whole or in part. Any reproduction must not be for commercial use or profit.

________________ ________________________________

Date Signature of Author

ABSTRACT

COMPUTING LABORATORY SUSTAINABILITY & UTILIZATION: INITIATIVES FOR A GREENER EDUCATION

Kristian Stokes

Department of Networking, Security & Systems Administration

Master of Science

Environmental, social and economic sustainability has been a recent focus of academia and industry to further invest in the future and lower operating costs. In this thesis, various emerging techniques are applied in an academic laboratory environment to record, analyze and project green IT initiatives. Student utilization metrics are evaluated to observe typical computing operation during a normal academic quarter. CPU load, network usage, power consumption, and room temperature are analyzed. Sustainable power and virtualization practices are recommended based on projections and assessments of current and recommended configurations. The findings of this study can be used across an academic environment to further reduce energy usage and meet the objectives of the Institute's goals towards a green and sustainable environment.

ACKNOWLEDGMENTS

I would like to acknowledge Dr. Yin Pan, Dr. Sumita Mishra and Dr. Bo Yuan, whose support and constant checkups made this thesis complete.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Thesis Statement

The current state of the economy and environmental issues are quite widespread. Global dialogues are frequent to address the seriousness of the economic recession and global warming. "Since 2001, there has been a torrent of new scientific evidence on the magnitude, human origins and growing impacts of the climatic changes that are under way," said Mr. Holdren, who is the president of the American Association for the Advancement of Science. "In overwhelming proportions, this evidence has been in the direction of showing faster change, more danger and greater confidence about the dominant role of fossil-fuel burning and tropical deforestation in causing the changes that are being observed." [1].

Scientific findings are assisting the customer awareness of green initiatives, thus adding pressure on businesses to develop and utilize greener alternatives. Doubled with the current economic situation, businesses are tasked with reducing budgets while still maintaining profit margins. Educational institutions are no different from the current goal of reducing costs while still maintaining research and academic initiatives.

The goal of this study is to gather, analyze and develop findings for laboratory usage given a typical Networking, Security & System Administration (NSSA) lab. The study involves capturing weekly student usage across systems. Simultaneous power and cooling usage are captured and recorded during a typical spring quarter at Rochester Institute of Technology (RIT). Other green IT initiatives such as virtualization, power management and storage consolidation are explored to further complement utilization studies.

This initiative focuses assisting with RIT's sustainability effort to develop measurable ways of reducing the carbon footprint of the Institute. The ultimate goal of this effort is to develop a set of defined techniques using the NSSA Labs as a pilot environment. The Institute can further use these findings to deploy across all computing labs.

1.2 Understanding The Need

The US Department of Energy estimates an average of 600 kWh per year is used for a single PC and monitor. Up to two thirds of that energy is wasted because PCs are running at full power when not in use [2]. This energy estimate is further compounded by the energy costs in New York state. Compared to other Middle Atlantic states, Pennsylvania and New Jersey, New York state currently possess the highest cost of electricity in all sectors. In the commercial sector, New York currently operates 47% above the national average [3].

The Gartner Group identified green IT as the number one Strategic Technology for 2008 [4]. The Gartner Group is a fact-based consulting service that helps clients use and manage IT to enable business performance. Technologies that the group identifies are viewed as areas of significant impact on the enterprise over the next three years combined with a need for major dollar investment, and the risk of being late to adopt. Gartner further states that the global carbon dioxide (CO_2) emission footprint from IT and communication technologies is now 2 percent of global CO_2 emissions, placing it on par with the aviation industry [5].

The foundation of the computing industry centers on manufacturers developing smaller, faster, cheaper and better electronics than previous versions. This is essentially what keeps the technology train rolling. The impact of this industry can be viewed as over a billion computers have been built, with the sole intention of eventually decommissioning them [6].

Green IT can be broken into three areas [7]:

- The impact of building and acquiring technologies.
- The impact of operating those technologies in your enterprise.
- The impact of disposing of these technologies after they have served their useful life in your organization.

This thesis studies each of these areas to better understand the impact of technologies in an educational environment. RIT has the opportunity to charge ahead in the area of computing sustainability to further comply with the Institute's commitment to the American College & University Presidents' Climate Commitment [8].

1.3 American College & University Presidents' Climate Commitment

On Earth Day 2009, Rochester Institute of Technology President, William Destler, signed the American College and University Presidents' Climate Commitment [9].

This commitment provides support and a framework for America's colleges and universities to go climate neutral.

> As part of this commitment, we will need to come up with an institutional strategic plan for reducing energy consumption and greenhouse gas production, and provide institutional structures and support that can move the campus toward the goals set in the plan, explains Destler. This plan will include a green building policy, an energy-star procurement policy, encouragement of public transportation, green power production and purchasing, and waste minimization. [9]

Colleges and universities are laying the groundwork for developing large climate neutral academic institutions over a series of several years. This challenge is to result in zero net greenhouse gas (GHG) emissions. This can be achieved by minimizing GHG emissions and using carbon offsets or other measures to mitigate the remaining emissions. Electricity consumption plays a large role in the calculation of total emissions. Reduction in computing energy will result in a large gain for the Institute to meet this commitment. The number of signatories to date accounts for 643 colleges and university presidents supporting this commitment.

1.4 Implications of Research

The importance of this research not only results in cost savings for the Institute, but impacts the way of life for all people on this planet. Reducing the carbon footprint of large scale computing environments result in a big win in the area of green computing. Using a flexible lab environment as a case study for this research, real data and initiatives can be developed to be recommended at the Institute level to achieve even greener computing across campus.

All tests designed in this study develop ways to produce measurable cost savings. Computing sustainability techniques should not be limited to corporations and data centers. RIT serves as a large stakeholder due to the number of IT support services, computing labs and students living on campus.

1.5 Hypothesis

In an academic computing laboratory environment, increasing the efficiency of system utilization will result in economical and ecological advantages of energy saving and hence, reduce overall carbon footprint of the institution.

This thesis will validate this hypothesis via analyzing empirical data collected in a NSSA department computing lab, which can be considered as typical computing lab in institutes that offer similar curriculum in computing and information sciences.

Analysis is conducted on the collected data to correlate it with measurable cost savings. Recommendations are be provided to identify substantial areas of environmental advantages and cost savings.

1.6 Summary

Green initiatives can be viewed on a global scope. Understanding and implementing solutions that result in "big win" can be challenging, but rewarding. This study will closely align with RIT's recent focus on sustainability and associated green programs including the American College & University Presidents' Climate Commitment. A viable study in this area with empirical data collected from an academic computing lab identifies actual energy and cost savings for the Institute. Therefore, the results from this study adds visibility for pursuing green initiatives.

Chapter 2

Related Work

2.1 Analogous Studies

Green IT initiatives have been continually growing in corporations and academic institutions. IBM has an outstanding track record for comprehensive global environmental management. IBM saved 4.6 billion kWh of electricity consumption and avoided 3.1 million metric tons of CO_2 emissions from 1990 to 2007 [10]. Recently IBM has reduced data centers from nearly 200 to fewer than ten. IBM not only actively greens itself, but its customers by using the same techniques and technologies for business services.

Indiana University initiated a pilot project in the School of Education to place the department's computers in hibernation after two hours and 15 minutes [11]. The four week pilot program decreased energy usage by 48.3 percent for 11 desktop computers. If approved for full deployment, the university could save at least $500,000 in energy costs and 15,000 tons of CO_2 emissions annually. To assist in educating users involved in the pilot program, a GoGreen Gadget was developed to outline individual, building and year-to-date savings in CO_2 emissions.

Other universities are doing their part by pledging to practice smart computing. Hosted by the Climate Savers Computing Initiative, colleges and universities across the country joined together in a movement to reduce the energy consumption of computers [12]. The initiative, Power Down for the Planet, gathered college campuses across the nation to share a commitment to sustainable computing practices. The challenge centered on a month long competition to see which university could recruit the largest percentage of their campus community to pledge support to Climate Savers Computing.

In March 2009, the IT Power Management Summit was held and hosted by the Climate Savers Computing initiative and the EPA's Energy Star Program. During the summit, a number of CIOs presented their successful deployment of power management schemes. Forrester Research analyst, Doug Washburn, noted 11 myths about power management:

1. I'm likely to see larger gains in the data center.
2. The business case for PC Power Management isn't very compelling.
3. I have to buy new, energy efficient hardware to reduce energy consumption.
4. The power used turning my PC on negates any benefits of turning it off.
5. My screen saver is saving me energy.
6. Turning my PC on and off will reduce its performance and useful life.
7. I can't run updates, backups and patches for PC in low power states.
8. There's no clear entry point – where do I begin?
9. I have no way of tracking and reporting the benefits.

10. I don't own my power bill so there's little incentive for me to reduce it.

11. My PC users will not tolerate any downtime for power management.

The number one myth, "I'm likely to see larger gains in the data center." was proven false in Figure 2.1. This chart shows desktop PCs and monitors contribute to 39 percent of power consumption, while servers and cooling contribute to 23 percent.

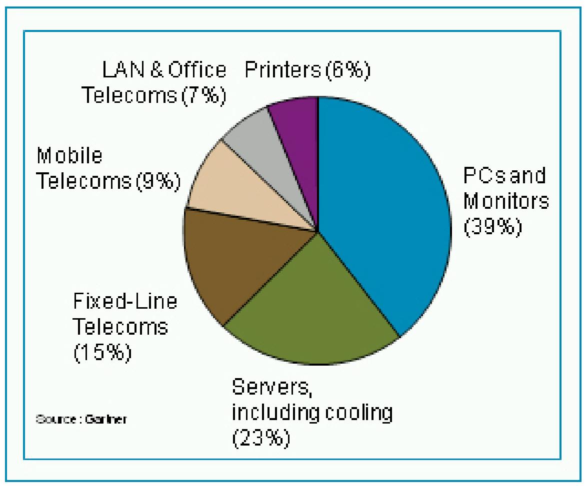

Figure 2.1 Typical IT power consumption.

This study further confirms that sustainable computing practices should not be limited to only the data center. A great deal of potential gain exists with typical workstations and computing labs.

2.2 Verdiem Surveyor

Commercial tools are available to track and reduce power usage across an enterprise. Educational institutions can also be considered an enterprise, the main focus being reducing unnecessary expenditures. Verdiem's Surveyor software centrally manages power policies over a network, which provides the IT staff full control over the current state of the enterprise. Power policies can be configured for the enterprise to trigger a shutdown procedure after an extended period of idle usage. If the IT staff requires a patch or configuration window, the workstations can be started remotely for this procedure.

Results of this software have been reported to significantly reduce utility costs and greenhouse gas emissions. Lake Washington School District cut PC power usage almost 2.5 kWh annually and saved almost 3 million pounds of greenhouse gas emissions produced in generating electricity [13]. To date, the use of Verdiem Surveyor has not resulted in a single help desk call for the school district.

2.3 Somniloquy

The primary reasons for leaving a computer switched on during times of limited usage stems from the ability to ensure remote access to local files, receiving e-mails, instant messaging and file sharing applications. With the current power saving schemes of sleep/suspend-to-RAM (ACPI state S3) and hibernate (ACPI state S4) [14], power is saved, but the PC will remain unresponsive to network commands.

Somniloquy attempts to solve this problem through the use of a hardware prototype which listens on a standard Ethernet interface for wake instances and transmits wake commands via USB to the host computer. Initial testing shows a system using Somniloquy consumes 11x to 24x less power than a PC in idle state [15]. The external secondary processor handles network communication and, if necessary, notifies the Somniloquy daemon residing on the host system for network changes such as new IP addressing. Specific application behaviors can be developed using the Somniloquy framework to handle unique environments. The prototype hardware and software has the potential to be incorporated into standard NIC cards and operating system power settings. Integrating the storage capabilities on standard NIC cards, or using memory blocks to store information to be processed upon wake-up is certainly plausible in the future. Operating system developers can further refine the Somniloquy framework to accept general user configuration. This can be similar to energy saving options currently used, such as placing the computer or display to sleep after a period of time. A user should have the ability to configure the network stack to remain active during a large file download while the computer turns into a low-power standby without download interruption.

2.4 Google PowerMeter

Google's mission is to organize the world's information and make it universally accessible and useful [16]. To further extend this information, Google is developing software to easily measure personal power consumption through the use of their PowerMeter [17] application. Google has paired up with electricity providers in the US and Canada to provide daily power usage data. Once the data is properly measured, informed decisions can be made on total power consumption of devices.

Energy usage is not itemized like our phone bills. We currently receive a billing statement from the local utility provider of the total usage for the month. We continue to pay the balance and unknowingly, we have no idea how much power the TV in the living room is using or how much it is costing for the family computer to remain idle for quick convenience. Worse yet, we as consumers are unable to project how much money and energy we could be saving by performing simple duties such as turning off printers when not in use. Google's initiative to itemize these items and provide a real-time viewpoint on residential power usage contributes to a large gain for user awareness.

2.5 Microsoft Hohm

Microsoft is also developing a home energy monitoring solution of their own. Microsoft Hohm [18] will utilize advanced analytics licensed from Lawrence Berkely National Laboratory and the U.S. Department of Energy to give customers highly personalized energy-saving recommendations. Customers will simply enter their postal code and a listing of basic information about their home such as occupants, appliances and systems. Based on this information, Microsoft Hohm will provide an energy report with recommendations.

2.6 Summary

The related work in this area is significantly growing in scope and subject matter. Most of the research has traditionally focused on data center efficiency. This study shifts the focus to workstation efficiency through the use of sustainable practices. Corporations and Universities such as IBM and Indiana University are actively pursuing these areas to further reduce energy consumption and operating costs. Software such as Verdiem Surveyor, Google PowerMeter and Microsoft Hohm assist with managing and identifying areas to save. Hardware such as the Somniloquy prototype assist with easing the concerns of users when implementing sustainable practices.

Chapter 3

Methodology & Approach

Many methodologies have been developed for full datacenter sustainability. This study is designed to leverage existing computing sustainability techniques and apply them to non-conventional aspects of workstations and computing labs.

3.1 Standards & Metrics

As with many technology and engineering characteristics, metrics and standards exist to allow for widely accepted products and points of comparison. Proprietary products often lock corporations into a single vendor essentially backing their procurement department in a corner. For this study, a number of standards and metrics are examined and included based on the relevance or application to workplaces and computing labs. Additional standards exist for environmental building design, regulation of hazardous substances and e-waste disposal.

Figure 3.1 Energy Star label.

3.1.1 Energy Star

Since 1992, the Energy Star program has made its mark on appliances, buildings and electronics. Energy Star is a joint program of the U.S. Environmental Protection Agency and the U.S. Department of Energy. The mission of the program is to identify and promote energy-efficient products to reduce greenhouse gas emissions. The program started with labeling computers and monitors, then expanded to major appliances, office equipment, lighting, home electronics, and more.

March 2009 marked the establishment of Energy Star 5.0 computer monitor specifications. With each revision, the standards become more stringent for manufacturers to fulfill. Power consumption is strictly outlined for on, off and sleep modes. Tiers are regulated with this specification stating that all tiers must comply with new version 5.0 requirements to retain the Energy Star label. This is determined by the manufacture date of the product. Older Energy Star qualifications are not automatically grandfathered to the latest version [19].

A new version of Energy Star desktop computer specifications is set for a July 2009 release. Computer power supplies (PSU) are measured by percentage of efficiency. For example, if a 500 watt power supply is 50% efficient, it would actually draw 1000 watts under a full load. The rest of the energy is then converted into waste heat

Fraction of Rated Load	**20%**	**50%**	**100%**
80 PLUS	80%	80%	80%
80 PLUS Bronze	82%	85%	82%
80 PLUS Silver	85%	88%	85%
80 PLUS Gold	87%	90%	87%

Table 3.1 80 PLUS certification standards.

which adds to the temperature of the ambient air. The new Energy Star specifications require 80 PLUS bronze certification which is tested under 20%, 50% and 100% loads of the maximum rated power of the PSU [20].

To date, the cumulative amount of savings through Energy Star products accounts for $254.7 billion with 1,070 million metric tons of carbon emissions avoided [21].

3.1.2 The Green Grid

The Green Grid [22] represents a non-profit trade organization comprised of IT professionals pushing for energy efficiency in data centers. Data center managers around the world are running into limits related to power, cooling and space. These factors contribute to the primary reasons for focusing on improving energy efficiency. With the additional stress on the power grid and the rapid growth of information technology, corporations are looking for solutions of energy efficiency. To assist in evaluating energy efficiency, The Green Grid developed two metrics, power usage effectiveness (PUE) and Data Center Infrastructure Efficiency (DCIE).

Each metric is designed to improve operating efficiency of the computing environment. The result also provides a metric for comparison against other environments. Overall, effectively measuring current energy usage provides a point on the scale to make improvements and determine their overall effectiveness. After all, if you can't measure it, you can't control it.

3.1.3 Power Usage Effectiveness (PUE)

The most prominent energy efficiency metric is the power usage effectiveness (PUE). This simple metric identifies the ratio of total energy used by the building and the energy used by the computing equipment. The total energy variable includes lighting, cooling and general electricity distribution.

$$PUE = \frac{Total\,Facility\,Power}{IT\,Equipment\,Power}$$

A PUE of 1.0 indicates perfect 100% efficiency with all power dedicated to IT equipment only. In April 2009, Google announced their average datacenter PUE is down to 1.19, with their most efficient datacenter achieving 1.12. While these are outstanding metrics, with proper datacenter design, a PUE of 1.6 should be achievable. Many corporate datacenters with a PUE of 3.0 or higher are not uncommon, which proves to be plenty of room for improvement [23].

PUE Score	**Rating**
3.0+	Poor
2.0-2.9	Average
1.7-1.9	Fair
1.4-1.8	Good
1.3	Very Good
1.2	State of the art
1.0 or less	Recheck your calculations

Table 3.2 PUE scoring.

3.1.4 Datacenter Infrastructure Efficiency (DCiE)

The reciprocal of PUE defines DCiE:

$$DCiE = \frac{1}{PUE} = \frac{IT\,Equipment\,Power}{Total\,FacilityPower} \times 100\%$$

- IT Equipment Power: This variable encompasses all IT equipment including storage, network devices, KVM switches, monitors, workstations, servers, laptops and all equipment using to monitor and control the data center.
- Total Facility Power: This variable encompasses everything that supports the IT infrastructure. This includes power delivery equipment, cooling systems and lighting.

For example, if the PUE is determined to be 3.0, this represents three times the energy demand necessary to power IT equipment. This ratio can be further extended to calculate new power demands of a system. If a server requires 500 watts, and the PUE is 3.0, this equates to a demand to the server of 1500 watts. The reciprocal DCiE, represents the total percentage of power IT equipment consumes. For example, if the DCiE is 33% (equivalent to a PUE of 3.0), this means that IT equipment consumes 33% of the power in the computing environment [24].

3.1.5 Corporate Average Data Center Efficiency (CADE)

While PUE and DCiE metrics specifically examine power, actual computational efficiency is disregarded. This is clearly identified in a computing environment with excessive server sprawl. RackForce, a web hosting provider, found itself adding 20 servers a day with low utilization on every server [10]. Both PUE and DCiE metrics do not account for this utilization variable. This fundamental flaw has yet to be fully realized by datacenters standardizing on PUE metrics.

An additional metric, Corporate Average Data Center Efficiency (CADE) takes into account the energy efficiency of facilities, their utilization rates and the level of utilization of servers. With the extra data points of utilization factored into the metric, the ending result allows for greater visibility into the actual IT efficiency. Why continue to power a rack of servers with low utilization 24x7 when the same job can be performed in a virtual environment?

$$CADE = (\text{Facility Efficiency}) \times (\text{IT Asset Efficiency})$$

- Facility Efficiency = Energy delivered to IT / energy drawn from utilities
- IT Asset Efficiency = Average CPU utilization across all servers. [25]

McKinsey & Company, a management consulting firm, notes in their July 2008 Revolutionizing Data Center Energy Efficiency report that data centers should adopt CADE and use the metric to track and double energy efficiency by 2012 for the quickest and easiest way to improve an organization's return on assets and reduce GHG emissions [26]. Figure 3.2 illustrates the initial breakdown of the metric and their related CADE levels:

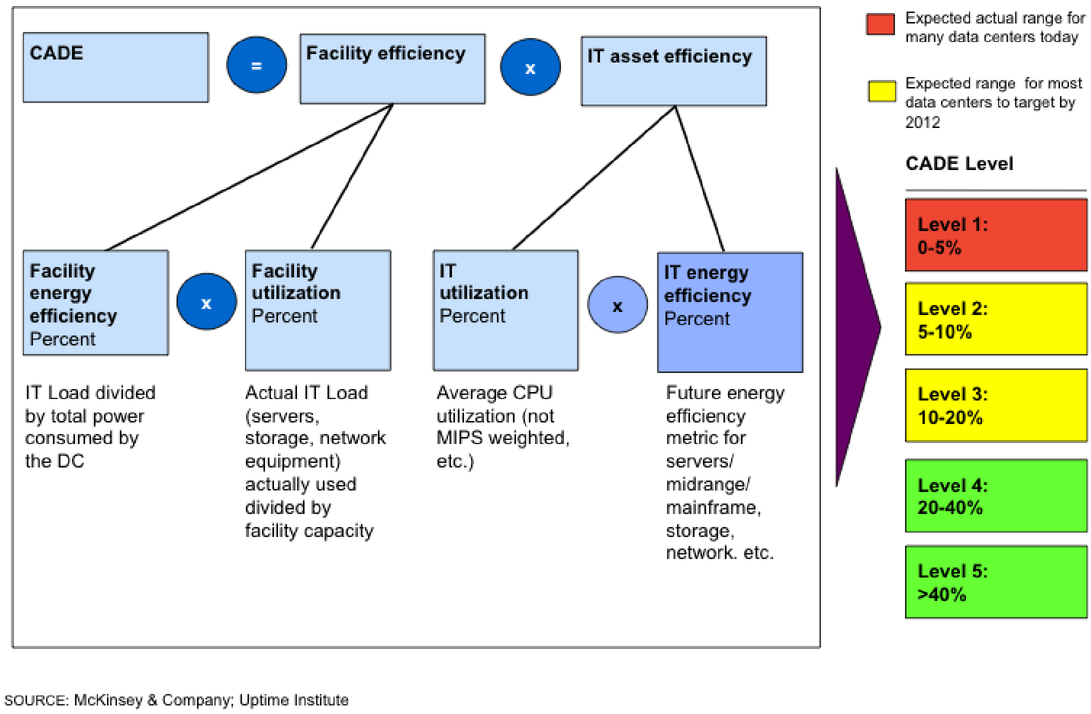

Figure 3.2 CADE breakdown.

3.1.6 Future Metrics

The Green Grid [22], is also in the process of developing longer term metrics that provide additional granularity for PUE and DCiE.

$$PUE = \frac{1}{DCiE} = CoolingLoadFactor(CLF) + PowerLoadFactor(PLF) + 1.0$$

Where all factors are ratios that are divided by the IT Load and:

- 1.0 represents the normalized IT Load. Effectively this is the IT Load Factor (ILF) but is always 1.0.
- Cooling Load Factor (CLF) is the total power consumed by chillers, cooling towers, computer room air conditioners (CRACs), pumps, etc. divided by the IT Load.

- Power Load Factor (PLF) is the total power dissipated by switch gear, uninterruptible power supplies (UPSs), power distribution units (PDUs), etc. divided by the IT Load.

These metrics will be designed to address the blurring of the lines between the IT equipment and facility infrastructure. The latest cooling technologies are integrated closely with IT equipment which further blurs the lines of these components. The Green Grid will look at these and other possible PUE and/or DCiE related metrics in the future [24].

3.2 Where and How to Measure

Many different choices exist when measuring a computing environment. When attempting to gain total facility power for the PUE metric, a reading should be taken from the power feed into the computing environment. This may also become difficult when datacenters blur with office buildings and are not a set, stand-alone structure. In this case, if available, sub metering should be used to isolate facility equipment.

- Measurement at the UPS is also viable; however these calculations result in only an approximation due to the inherent power inefficiencies.

- Measurement at a metered Power Distribution Unit (PDU) presents an active measuring technique for IT equipment, but usually fails to capture total facility power such as lighting and cooling. A PDU measurement can also provide granular per-outlet visibility into each device to identify extreme power inefficient devices.

Figure 3.3 Metered Power Distribution Unit (PDU).

- The final place to measure power usage is the actual CPU on the IT equipment. Again, this measurement will not provide total facility power, and is not recommended primarily due to the fact that IT personnel commonly replace entire devices, not individual CPUs.

Recommended approaches for measuring datacenter power usage are metered rack PDUs and intelligent rack PDUs that have the capability to measure individual outlets. This measurement provides granular, per outlet level information that can be clearly identified and quickly dealt with accordingly [25].

3.3 Data Gathering

Once a decision has been made on where to measure power usage, the data for each device should be gathered in regular intervals. The shorter the interval, more data points become available to reach an accurate analysis. Peak hours and seasons should also be monitored effectively to identify stress on the systems and associated power distribution. This data can also provide simultaneous power usage management on a per rack basis. If a rack is underutilized or a rack is under the risk of tripping the breaker, the active polling will provide additional visibility to these issues.

Each IT device commonly contains nameplate power ratings from the manufacturer. These ratings are often misconstrued which causes inefficient power distribution. Raritan, a leading supplier of rack power management, discovered the following on their own datacenter:

Through tests in its own data center, Raritan determined that rules-of-thumb percentages of nameplate ratings simply don't work [25]. Across 59 servers, 15 had average power consumption of 20 percent or less, 29 had 21 to 40 percent, 9 had 41 to 60 percent, 4 had 61 to 80 percent and 2 had 81 percent or more. Even at peak power consumption 49 of the servers were 60 percent or less of their nameplate rating. Many data center planners use 70 percent of nameplate which means there is a lot of stranded power in many data centers. On the other hand, at peak power consumption 5 of the 59 servers were at 81 percent or more of nameplate and therefore at risk of shutting down. The message is that in terms of power consumption, it is important to know what is going on at the individual device, not some aggregated average which may mask problems both on the high and low side.

3.4 Summary

Several metrics are currently used to measure the overall efficiency of a computing environment including PUE, DCiE and CADE. Each metric contains strengths and weaknesses associated with the data inputs. For example, PUE and DCiE metrics do not include actual server load efficiency when calculating. Due to this reason, this study will focus on CADE metrics for the NSSA computing environment. Future metrics will further break down cooling and power load factors to independent variables to extend insight into these areas. Gathering accurate and reliable data for equipment and facility power can be challenging without the proper equipment, however once the effort is started, the return on investment is likely significant.

Chapter 4

Experiment

An experiment was developed to gain quantifiable data on a traditional academic computing lab. Conventional sustainable data center practices and techniques are applied to a computing lab workstation environment and equally analyzed for operating efficiencies. The experiment is designed to scientifically test the hypothesis outlined in Section 1.5.

4.1 Case Study

A system administration lab was chosen as the basis of this study due to the number of workstations, lab infrastructure and usage during an academic quarter. The system administration lab is the largest NSSA lab containing black Antec workstations also used in the NSSA Projects lab. The lab contains 80 workstations, one core switch and two servers for infrastructure support. The topology and device breakdown can be found in Appendix A.1 All lab equipment was monitored over the course of 7 weeks, April 6 to May 22 2009. The lab is designed to support introduction and advanced system administration courses that leverage VMware Workstation virtualization technology used in hands-on labs.

SNMP agents were installed on all devices with a polling interval of 120 seconds. Data collection consisted of network response time for each node, CPU load and memory usage. Network response time was specifically monitored to properly gauge metrics on workstations that are online and not disconnected from the network or forced in a reboot loop. CPU load and memory usage was captured to extrapolate lab utilization during open and closed hours.

A wireless sensor network was deployed in the ceiling of the lab to capture temperature data over the course of the experiment. Four sensors were strategically placed within the lab ceiling to form a wireless mesh network for data collection and recording. Once captured, data is forwarded to a central gateway and stored in a database for later data correlation. Heating Ventilation and Cooling (HVAC) data is recorded from an AutomatedLogic controller system which provides cooling metrics for the lab.

This case study is designed to establish a baseline for current power usage and project the outcome of recommended green initiatives.

4.2 Data Collection Methods

The overall approach to this thesis is to collect useful, quantifiable data to conclude whether or not increasing the student utilization efficiency will be advantageous in making the environment green. The following areas are explored to gather sufficient data for a cost saving analysis.

4.2.1 Student Utilization

Utilization of computing labs is recorded to estimate student usage and potential computing power of the NSSA Labs. This initiative requires the use of each node reporting on a regular basis (120 seconds) the current processor and memory usage of the system. This is easily performed through the use of SNMP reporting to a centralized server.

From this data, captured over a period of seven weeks during spring quarter, provides average CPU utilization, network usage, uptime, etc. Overall computing utilization is estimated to provide an assessment of cost and energy savings pending application of green initiatives.

4.2.2 Power Utilization

Energy usage is the largest focus of this study. RIT spends approximately $11 million per year in essential utilities, including gas and electric costs [27]. The Institute is facing the nationwide escalation of utility costs and the continual growth of the campus to further impact the current financial standing.

Power readings are recorded over time and extrapolated based on student usage. Power readings are captured on the following devices:

- Typical bench equipment (desktop computers, CRT/LCD monitors)
- Additional devices (bench routers, switches, hubs, etc.)

Where applicable, each reading includes standby, idle, hibernate, normal and high power usage of devices. Student usage data is then correlated against these metrics to estimate power usage during a typical quarter.

4.2.3 Cooling

Cooling also contributes to a large percentage of cost in a computing environment. This experiment incorporates temperature recording captured during five minute intervals throughout the course of the study.

A wireless sensor network (WSN) is deployed in the ceiling of the lab to constantly record temperature readings over an extended period of time. This data can then be averaged and evaluated for cooling initiatives. Heating Ventilation and Cooling (HVAC) data is recorded from an AutomatedLogic controller system which provides airflow metrics for the lab.

4.2.4 Virtualization

The NSSA Labs currently employ the use of desktop virtualization. This initiative has many benefits in and of itself to allow for reduction of hardware and increase the number of possible operating systems for the student.

This study further explores the realm of virtualization to incorporate projects to further virtualize the NSSA curriculum. The increasing pressure to offer classes abroad via online collaboration creates interesting solutions without sacrificing lab quality.

The use of Remote Laboratory Emulation System (RLES) is incorporated into this study to expose current initiatives of virtualizing courses. Several courses such as advanced forensics and web/network auditing currently take advantage of this technology by simulating the lab environment to the student.

4.3 Technologies

The following technologies are used to observe and measure results:

4.3.1 Simple Network Management Protocol (SNMP)

SNMP version 2 is used to collect system data across all nodes in the lab environment. This data consists of CPU utilization, memory usage, disk access, uptime, etc. All

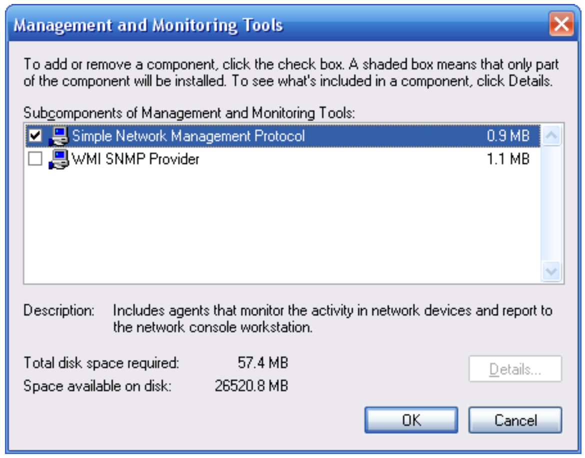

Figure 4.1 Windows SNMP component installation.

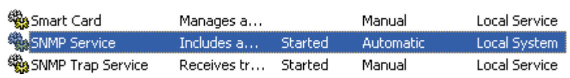

Figure 4.2 SNMP service installed and started.

client data is collected and stored on a centralized SNMP server.

The standard Windows XP Professional SNMP agent was installed on all lab PCs. The agent runs as a background service that reports data to the server at all times, during normal operation, logout, and under high load. Third-party software exists to provide this functionality, however the intent was to provide a seamless transition into the study without raising student interest or interrupting normal lab operation. The lab uses Deep Freeze, a system integrity product to prevent changes on the system. This required the assistance of lab staff to "thaw" the lab machines, install and configure the SNMP agent, modify firewall settings, reboot and freeze the image for the changes to have a lasting effect. This change was not mentioned to students. This is due to the potential for students to skew results based on their direct access to the nodes being monitored.

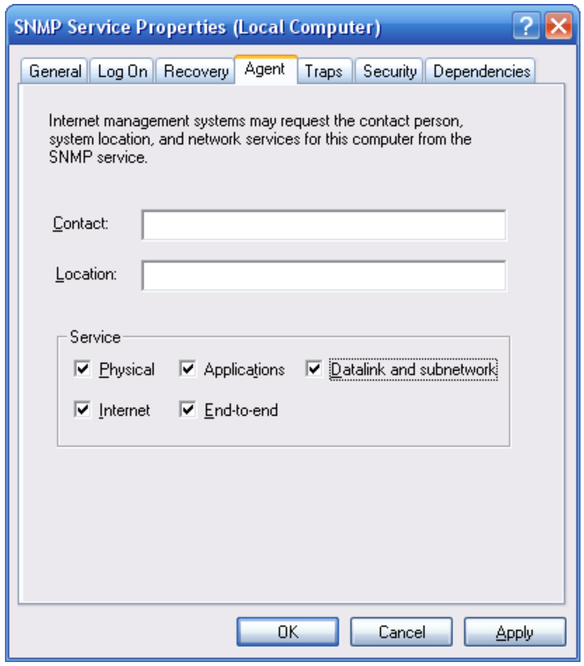

Figure 4.3 Windows SNMP agent configuration.

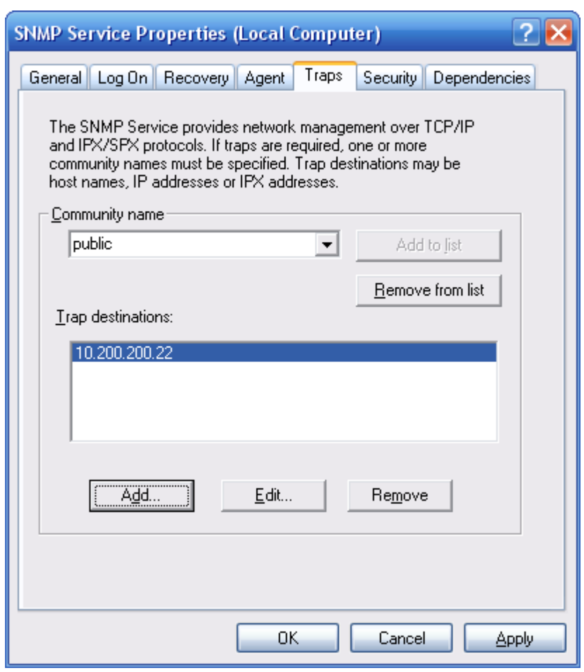

Figure 4.4 Windows SNMP agent configuration.

4.3.2 SNMP Data Aggregation

SolarWinds Orion Network Performance Monitor serves as the SNMP collection engine to poll clients and efficiently store data in a MS SQL database. This data is then graphically analyzed and exported using a number of display formats.

With the assistance of the sales staff at SolarWinds, a temporary academic license was granted specifically for this study. SolarWinds Orion Network Performance Monitor was chosen to provide a commercial toolset for gathering and analyzing data from SNMP agents. The toolset is proven to be extremely versatile with scalability to thousands of nodes. The provided SNMP management information bases (MIBs), proved essential for vendor specific SNMP reporting (Cisco, Dell, etc.). All SNMP nodes reported utilization data every 120 seconds to the SolarWinds Orion Windows 2003 virtual machine. Based on this stored data, utilization metrics can be calculated such as percent of uptime, CPU usage, disk volume activity, etc.

4.3.3 Wireless Sensor Networks (WSN)

Temperate data is captured to correlate cooling metrics and determine hot spots in the lab where warm and cool air mix. To accomplish this task, four Crossbow MICA2 motes were programmed and stored in the ceiling to record temperature data. Reporting periods consisted of sampling every five minutes. Collected data is then transmitted over a 433MHz wireless radio spectrum to a Crossbow MoteView server which stores the data for correlation. The data is formatted in a PostgreSQL database which then can be used by MoteView's graphing features or third party programs. The motes create an active mesh network that change over time if nodes fail or route health decreases. During testing, a single mote would survive approximately two weeks on a set of AA batteries.

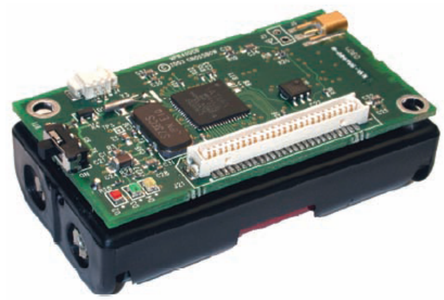

Figure 4.5 Crossbow MICA2 mote.

4.3.4 Electricity Usage Monitoring

A P4320 Kill A Watt PS [28] surge protector was used to measure power usage of selected electrical devices in the lab. Voltage, line frequency, watts, amperage, KWH/Leakage is captured on various devices to record average and state changing power usage. This data is used to establish a baseline and develop a power usage projection.

For example, when measuring a lab PC, the Kill A Watt device provided an in-line measurement tool to provide accurate metrics on wattage and amperage. A PC was tested under load and with several varying state changes, including off, on, standby and hibernate. Additional scenarios were developed such as initial power on of a bench which recorded minimum and maximum values during power up.

Figure 4.6 P4320 Kill A Watt PS.

4.4 Summary

The experiment to test the working hypothesis involves the use of a NSSA department computing lab containing 80 workstations that regularly report their utilization. CPU and memory usage are captured via SNMPv2 and stored using SolarWinds Orion Network Performance Monitor. Additional metrics are taken such as room temperature from a deployed wireless sensor network and power consumption under varying operating loads. This experiment is designed to mimic other academic computing labs. The results gained from this study should closely relate to other labs with similar hardware.

Chapter 5

Results

This section examines the results captured from the experiment outlined in Chapter 4.

5.1 S3 Sleep Configuration

During initial power testing, the configured standby power measurements appeared extremely high for a sleep state. Upon further investigation, the configured sleep state of the system was set to S1. Four sleep states exist ranging from S1-S4, with S0 being on and S5 being off (see Appendix B.2). Power consumption can be dramatically improved if the proper sleep levels are configured on the operating system level and BIOS.

S3 sleep must be properly configured in the BIOS settings before the operating system can detect the capable sleep states. Most modern systems are compatible with all sleep levels. Output 1 displays correct power settings to ensure proper BIOS configuration.

Output 1 Recommended ACPI BIOS Configuration

```
Suspend Mode [Auto]
Repost Video on S3 Resume [No]
ACPI 2.0 Support
ACPI APIC support [Enabled]
```

The settings were tested and captured on a standard syslab PC. Options and capabilities may exist under different motherboards and BIOS versions.

Upon discovery of the inefficient S1 sleep state, further research was conducted to force the system to an S3 (suspend to RAM) state. The research revealed a Microsoft utility called "dumppo.exe", which continues to remain, for the most part, undocumented. The command line utility provides insight to the supported sleep states of the system, CPU throttling, etc. The utility is freely available: ftp://ftp.microsoft.com/products/Oemtest/v1.1/WOSTest/Tools/Acpi/dumppo.exe.

Output 2 dumppo.exe power capabilities output

```
C:\>dumppo.exe cap
power capabilities
System power capabilities
Power Button Present....: TRUE
Sleep Button Present....: FALSE
Lid Present.............: FALSE
System states supported.: S1 S3 S4 S5
Hiber file reserved.....: TRUE
Thermal control.........: FALSE
CPU Throttle control....: FALSE
Processor min throttle..: 100
Processor throttle scale.: 100 (1%)
Some disk will spindown.: TRUE
System batteries present: FALSE
System batteries scale..: (G:0 C:0) (G:0 C:0) (G:0 C:0)
Ac on line wake ability.: Unspecified
Lid wake ability........: Unspecified
RTC wake ability........: S4 - hibernate
Min device wake.........: Unspecified
Default low latency wake: Unspecified
```

The syslab systems are capable of S1, S3, S4 and S5 sleep states. Knowing this information, the dumppo.exe utility provides capabilities to force the operating system to a user configured sleep state. The following command was issued to force Windows XP to S3:

Output 3 dumppo.exe force S3 sleep

```
C:\>dumppo.exe admin minsleep=s3 maxsleep=s3
Admin policy overrides
Min sleep state......: S3
Max sleep state......: S3
Min video timeout....: 0
Max video timeout....: -1
Min spindown timeout.: 0
Max spindown timeout.: -1
```

The command modifies a binary registry setting to force a S3 sleep. This misconfiguration stems from Windows incorrectly identifying the capable power states upon initial installation. This could be attributed to insufficient drivers at the time of installation or misconfiguration of BIOS power settings.

5.2 Hibernate & Deep Freeze

Hibernate (S4) functionality on a modern computer typically resembles the shutdown (S5) state due to the majority of devices being powered off. The remaining power consumption results from trickle current delivered to the power button. The system state is saved to a hibernation file and upon reboot the file is loaded to restore the session [29]. Hibernation presents the lowest power state available, thus making the power mode attractive for green initiatives.

The system administration lab, and other labs administered by NSSA, use a system integrity software package named Faronics Deep Freeze [30]. Deep Freeze prevents permanent changes from affecting the operating system. Any changes in configurations or files during the user session will be reverted to a known good, "frozen" state upon reboot. During each reboot, the frozen system data is read and loaded discarding all previous data the system contained. This prevents the system from recording accidental or malicious changes to the operating system.

The addition of Deep Freeze to the system administration lab environment, complicates the hibernate state. The hibernation state essentially simulates the shutdown state, except the current system state is written to the hard drive. Restoring from hibernation re-initializes the computer, walking through BIOS prompts on startup. In the case of Deep Freeze, the re-initialization process acts like a normal system startup which causes an inherent incompatibility with the software. With Deep Freeze installed, the protected operating system will load instead of the hibernated state. In fact, during testing, enabling hibernation on a frozen system was impossible. The fundamental incompatibility is how Deep Freeze prevents any changes to the system state, especially the creation of hibernation files.

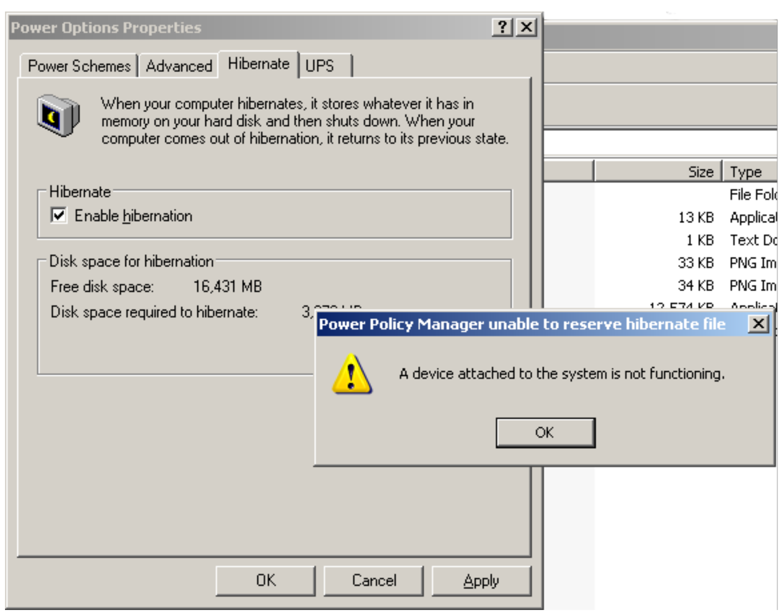

Figure 5.1 Error presented during hibernation under Deep Freeze.

5.3 Power Usage

Power usage metrics were measured using a Kill A Watt [28] surge protector outlined in section 4.3.4. Each scenario was tested under differing operating modes for each device. Monitors were tested powered off (no LED), on standby (amber LED) and on (green LED). PC towers were tested as fully populated with common peripherals such as mouse, keyboard and a network connection. PC towers were tested under varying operating loads and power states. For this study, watts will serve as the primary metric of focus. Identifying the operating watts further computes into kilowatt hours (kWh) and other comparable metrics. For a list of key terms, please see Section B.1.

5.3.1 Syslab PC Tower

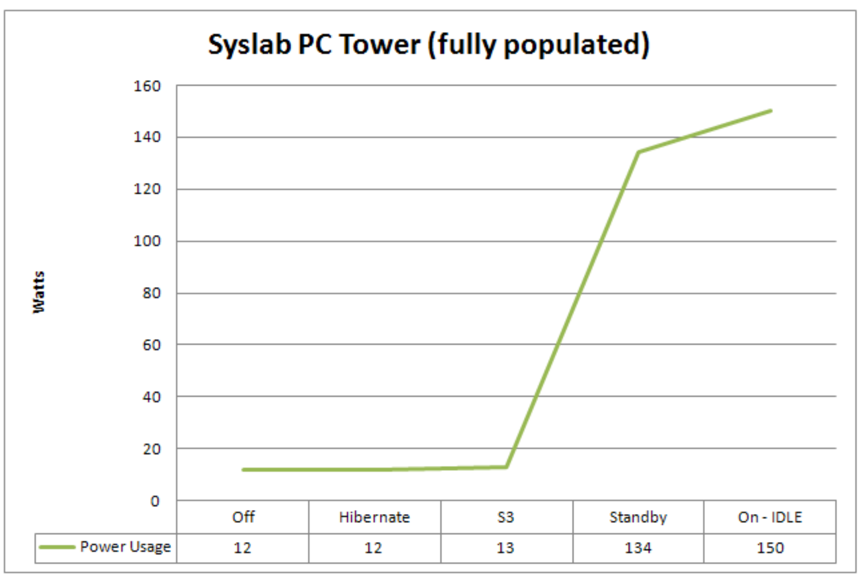

Figure 5.2 1 Syslab PC Tower.

Figure 5.2 and Table 5.1 display the power usage metrics for a single syslab PC tower. The varying power states of the system are noted. The state labeled “Standby” corresponds to the S1 sleep state, mentioned in Section 5.1. The corrected S3 sleep

Mode	Volt	Amp	Watt	Power Factor %
Off	121.5	0.1	12	74
Hibernate	119.9	0.1	12	65
S3	120.2	0.1	13	65
Standby	121	1.1	134	100
On - IDLE	121.2	1.2	150	100

Table 5.1 1 Syslab Tower (fully populated).

state is labeled as "S3". With proper configuration, the wattage of S1 - Standby 134, is reduced to 13 watts for S3 sleep. This is a 90% decrease in power consumption just on sleep states.

Another important note is that at total system shutdown, labeled "off", the PC tower still consumes 12 watts of power. This is very typical with consumer electronics and is often referred to as "vampire power", "phantom load", or "leaking electricity". The only way to eliminate vampire power is to cut power to the device at the source. In 2001, US President George W. Bush issued Executive Order 13221, which states that every government agency, when it purchases commercially available, off-the-shelf products that use external standby power devices, or that contain an internal standby power function, shall purchase products that use no more than one watt in their standby power consuming mode. [31].

5.3.2 Syslab CRT/LCD Monitors

Figure 5.3 and Tables 5.2, 5.3 display power consumption metrics for CRT and LCD monitors used in the NSSA system administration lab. Both monitors have a great characteristic of zero watts used during off and standby modes. Figure 5.3 displays a comparison of CRT and LCD technology, the LCD monitor uses 34% less wattage.

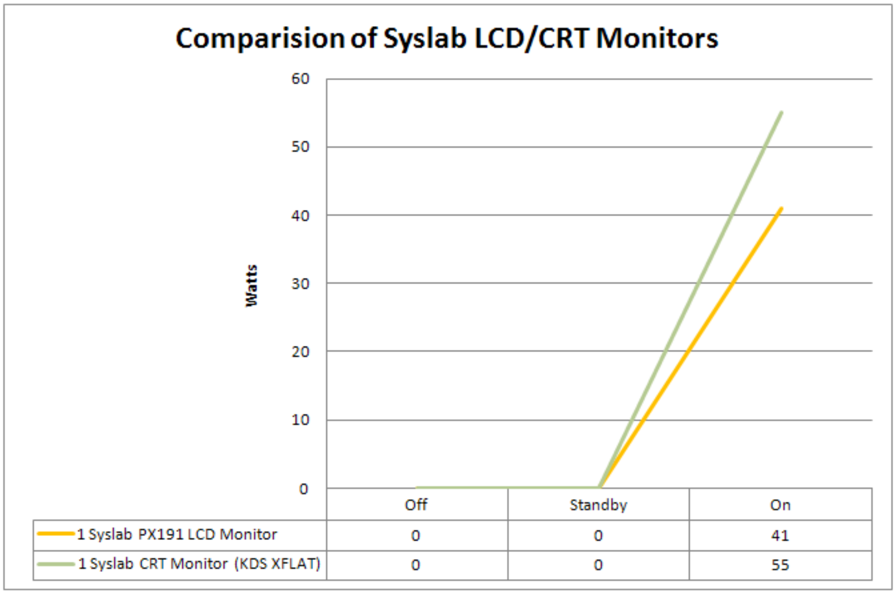

Figure 5.3 Syslab CRT/LCD Monitor.

Mode	Volt	Amp	Watt	Power Factor %
Off	121.5	0	0	5
Standby	121.3	0	0	7
On	121.2	0.4	41	73

Table 5.2 1 Syslab PX191 LCD Monitor.

Mode	Volt	Amp	Watt	Power Factor %
Off	121.2	0	0	14
Standby	121.3	0	0	24
On	121	0.6	55	73

Table 5.3 1 Syslab CRT Monitor (KDS XFLAT).

5.3.3 Syslab Bench Comparison

Figure 5.4 and Table 5.4 display the power metrics for a single workstation, PC tower and LCD monitor, in varying power states. With the workstation in standby mode

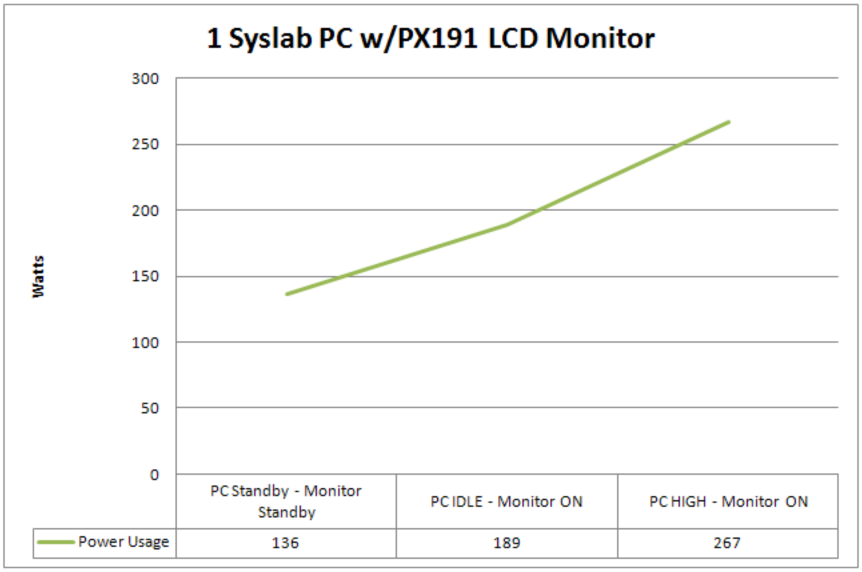

Figure 5.4 1 Syslab PC w/PX191 Monitor.

Mode	Volt	Amp	Watt	Power Factor %
PC Standby - Monitor Standby	120.7	1.1	136	100
PC IDLE - Monitor ON	121	1.6	189	96
PC HIGH - Monitor ON	120.7	2.1	267	98

Table 5.4 1 Syslab PC w/PX191 LCD Monitor.

(S1), power consumption was measured at 136 watts. During idling conditions, the workstation was measured with the monitor on and PC CPU utilization at 0%. The resulting idling condition over standby is a 39% increase in wattage. A stress test of the system was performed using SiSoftware Sandra Lite [32] v15.99 "Processor Arithmetic" benchmark. The benchmark results in a CPU utilization of 100% for several minutes allowing sufficient time for power readings. High utilization of the system equates to an increase of 96% over standby (S1) and a 41% increase over PC idle power consumption. High power consumption is not uncommon for the lab during classes due to the heavy use of VMware virtualization.

Figure 5.5 and Table 5.5, 5.6 display a comparison of 3 workstations per bench to

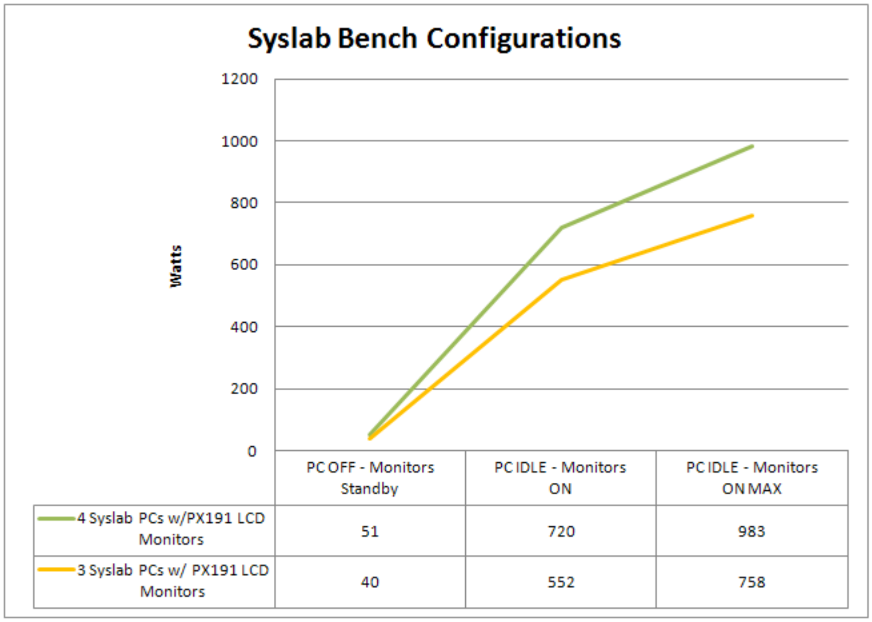

Figure 5.5 Syslab Bench Configurations.

Mode	Volt	Amp	Watt	Power Factor %
PC OFF - Monitors Standby	121.8	0.5	40	59
PC IDLE - Monitors ON	119	5.9	552	78
PC IDLE - Monitors ON MAX	121.9	8.2	758	

Table 5.5 3 Syslab PCs w/PX191 LCD Monitors.

Mode	Volt	Amp	Watt	Power Factor %
PC OFF - Monitor Standby	121.5	0.7	51	60
PC IDLE - Monitor ON	117.7	7.8	720	79
PC IDLE - Monitor ON MAX	121.9	10.6	983	

Table 5.6 4 Syslab PCs w/PX191 LCD Monitors.

4 workstations per bench. This comparison is investigated due to the migration to VMware Workstation, students are no longer using all four workstations. Students are able to run multiple operating system environments on a single physical workstation

and achieve the same results. The four workstation per bench configuration was standard during Spring 2009. The movement to three workstations per bench is now in effect to further reduce power consumption. The testing conditions attempted to emulate typical lab conditions of PCs off or idle and monitors on or standby. During Spring 2009, a four workstation bench in idling conditions consumed 720 watts of power. Reducing the workstation number to three per bench lowers the wattage to 552, a 23% decrease.

Another scenario was tested to record maximum wattage during a full bench startup. All four workstations started the test in the off position with monitors in standby. All workstations were powered on rapidly to create a bench "power on" effect. The maximum wattage was recorded as 758 and 983 for three and four workstation setups respectively. During startup, the workstations use 37% more wattage.

5.4 System Utilization

To identify the average state of workstation usage, system utilization data is captured during a typical academic quarter. The system administration lab floor plan is divided in the center, which is referred to as left and right halves in this study. SNMP agents are installed on all workstations that report CPU, memory and network utilization every 120 seconds. Normally student utilization metrics are performed via headcounts during a set interval to identify lab usage. For this study workstation activity is the primary focus, this essentially equates to CPU usage. It is known that as CPU usage increases, power usage increases. This data attempts to identify the average CPU utilization during all recorded times.

5.4.1 System Availability

Several caveats are present with this data. First, due to the fundamental purpose of the system administration lab, students are free to disconnect network cables, modify configuration settings, etc. Each workstation contains two network interface cards, one that connects to a student network and one that connects to an imaging network. Most students will only manipulate the student network and not actively alter the imaging network. The primary reporting mechanism for the workstations is to communicate to the SNMP server via the imaging network. Due to the unique learning environment, 100% uptime and data capture cannot be expected. However, through the use of Deep Freeze (see Section 5.2), the integrity of the configuration and operating system can be guaranteed upon reboot. Uptime metrics have been included to provide an idea on the average time on reporting.

Overall, during the period of 4/06/09 to 5/22/09, academic weeks 5 to 11, the SNMP server was running full-time capturing data. The total availability average for the workstations in the lab results to 47.72%. This means that out of all the data captured, 47.72% of the time workstations were reachable on average, which resulted in proper data collection. The remaining 52.28% of the time the workstations were not reachable due to network disconnection.

LEFT AVERAGE	53.20%
RIGHT AVERAGE	41.50%
TOTAL AVERAGE	47.72%

Table 5.7 Syslab availability average.

Detailed availability metrics can be viewed in Table B.1

5.4.2 CPU Utilization

CPU utilization metrics are captured to identify the overall average operating state of the lab. For instance, if the average CPU utilization is 50%, then the average operating state is between idling and high conditions. The measurement period coincided with the availability metrics, from academic week 5 to 11.

LEFT AVERAGE	0.43%
RIGHT AVERAGE	0.20%
TOTAL AVERAGE	0.32%

Table 5.8 Syslab CPU load average.

Overall, the total CPU load percentage is extremely low. With a total average of 0.32%, the lab can be considered in an idle state all of the time. Based on this data, idling conditions with monitors in standby will be considered the "Normal Operation".

Detailed CPU load metrics can be viewed in Table B.2.

5.5 Corporate Average Data Center Efficiency

When evaluating efficiency metrics, CADE can be considered one of the best, due to the inclusion of IT utilization metrics (see Section 3.1.5). To compute the CADE on the system administration lab, total facility power metrics such as lighting and cooling are needed. At this point, these metrics are unattainable due to the lack of individual metering per room.

$$CADE = (\text{Facility Efficiency}) \times (\text{IT Asset Efficiency})$$

- Facility Efficiency = Energy delivered to IT / energy drawn from utilities
- IT Asset Efficiency = Average CPU utilization across all servers. [25]

With this stated, CADE can still be investigated based on the data collected thus far. With the IT utilization level of 0%, we can establish a CADE level 1 0-5% efficiency rating. Unfortunately, many datacenters and labs are in this red area that needs drastic improvement for the future. Sleeping or powering down underutilized systems can increase this efficiency rating.

5.6 Cost-benefit Analysis

The primary motivating factor for pursing Green IT sustainable initiatives is cost savings. This section focuses on the potential for reduction in operating costs of the system administration lab. Based on the power usage data gathered in Section 5.3, projections can be made to calculate total energy usage of the lab. When assessing these projections, lab configurations can be easily compared to estimate savings and further recommend operating adjustments to further comply with sustainable practices.

Based on the single workstation and bench metrics captured in section 5.3 the total power usage of the lab can be extrapolated given the known number of workstations. Raw calculations and results can be found in Appendix B.5. The raw data is based on the following calculations:

5.6.1 Total Wattage

The total wattage for all devices in syslab are computed via the following formula:

$$Total\,Wattage = Quantity\,of\,Devices \times Watts\,Per\,Device$$

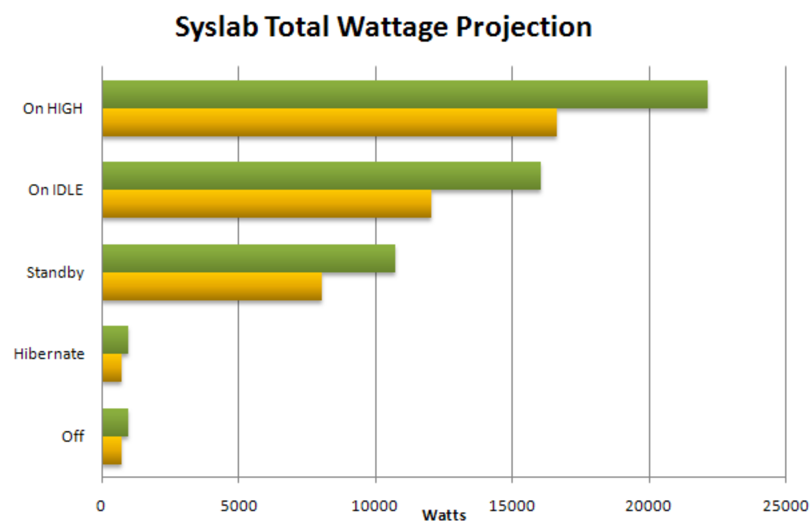

	Off	Hibernate	Standby	On IDLE	On HIGH
4 PCs per Bench	960	960	10720	16048	22128
3 PCs per Bench	720	720	8040	12046	16606

Figure 5.6 Syslab Total Wattage Projection.

Figure 5.6 displays a bar chart representing the total wattage consumed in syslab. The devices include PCs, LCD/CRT monitors and network hubs. These projections assume the specific scenario is played out through the entirety of the test. This means that the "On IDLE" test is computed to reflect no user interaction, enabling the idle (0% CPU utilization) condition.

Figure 5.7 displays a bar chart focusing on the current normal operation compared to the recommended operation. Table 5.9 displays the operating states between the two. The primary difference resides in the PC operating state. In the past, the normal

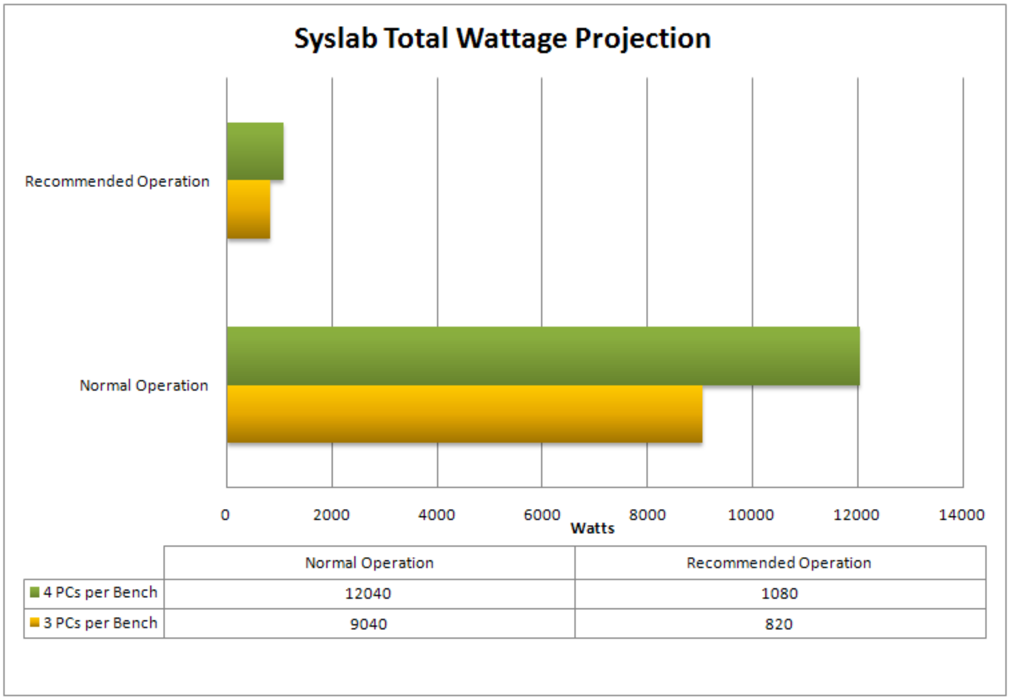

Figure 5.7 Syslab Normal vs Recommended Projection.

	Normal Operation	Recommended Operation
PCs	On - IDLE	S3 Sleep
CRTs	Standby	Standby
LCDs	Standby	Standby
Hubs	On	On

Table 5.9 Normal vs recommended operating comparison.

operation has been online and idle, waiting for user input. This can be confirmed by the current power configuration in Figure A.2 and A.3. The recommended operation is to set the system to S3 standby for greater power savings. With this setting enabled in the BIOS and operating system, a 91% reduction in power usage is possible.

5.6.2 Total Kilowatts

A kilowatt hour is computed for all devices via the following formula:

$$Kilowatt\, hour = \frac{Total\, Watts}{1000}$$

This formula utilizes the total watts of the devices measured in Section 5.3. Normally kilowatt hours (kWh) are a separate measurement taken on the device during the time of measuring. The device is normally measured continuously for an hour creating a kWh reading for the device. This is performed to capture normal usage and changes in the device such as a motor starting in a refrigerator. With electronic devices, the irregularities of power consumption over time are less frequent and drastic.

To compute a measurement of watts in time, the above formula was used to provide kWh. During measurement, each device seldom drifted out of the recorded wattage, 3% max, which created a steady kWh baseline.

$$Kilowatts\, per\, Day = Total\, Watts \times \frac{24\,(hours)}{1000}$$

$$Kilowatts\, per\, Year = \; Kilowatts\, perDay \times 7\,(days) \times 52\,(weeks)$$

Each kilowatt bar chart represents a projection of kilowatt hour usage during an identified time period. Normal static baselines are displayed and compared to normal and recommended operating conditions. Again, the standby scenario was tested using the default power settings in syslab which is S1 sleep. If the sleep level is changed to S3, dramatic improvements can be made to the total kilowatts used.

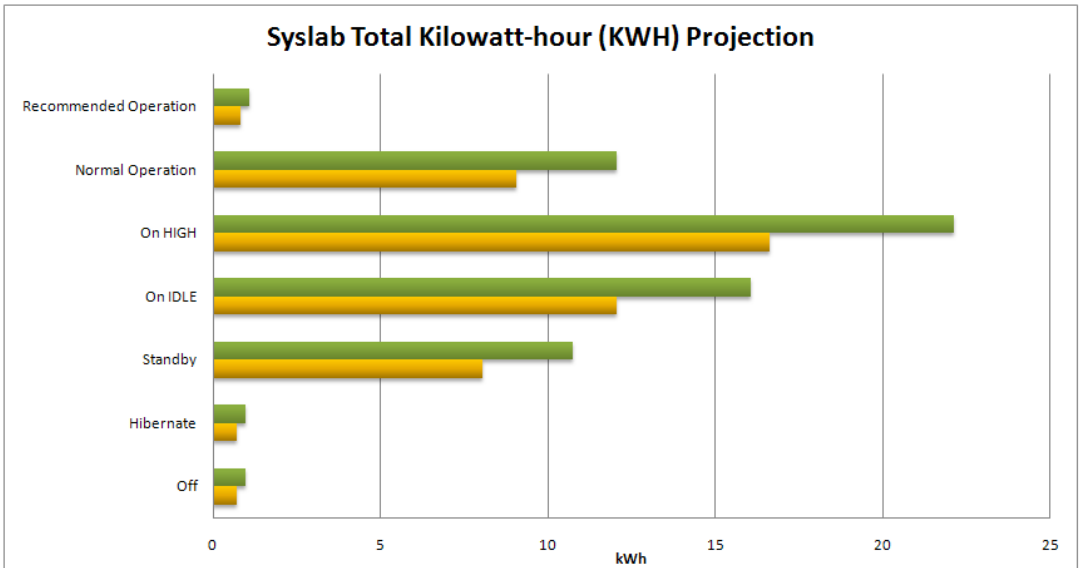

	Off	Hibernate	Standby	On IDLE	On HIGH	Normal Operation	Recommended Operation
4 PCs per Bench	0.96	0.96	10.72	16.048	22.128	12.04	1.08
3 PCs per Bench	0.72	0.72	8.04	12.046	16.606	9.04	0.82

Figure 5.8 Syslab Total Kilowatt-hour (kWh) Projection.

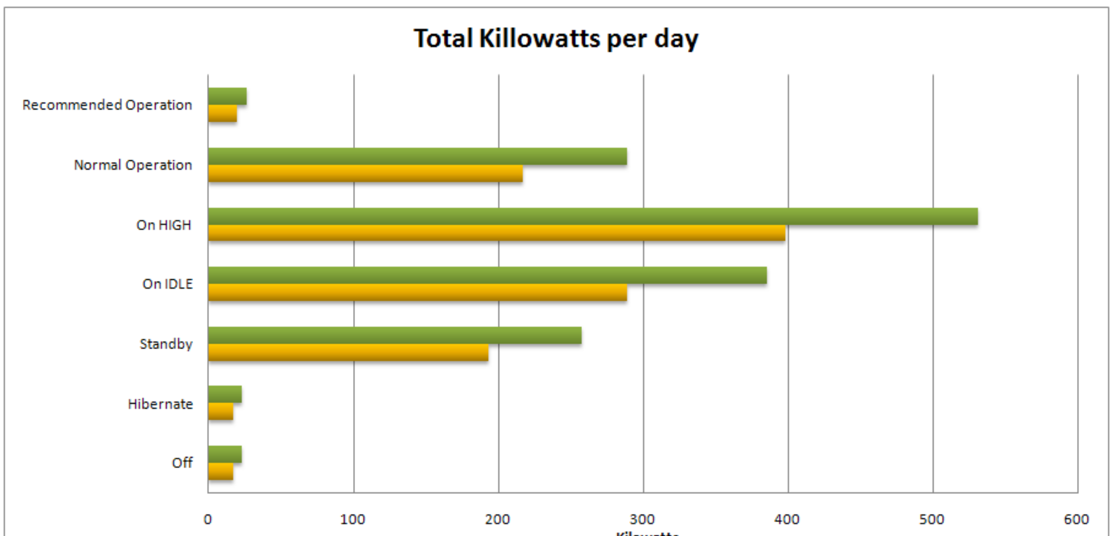

	Off	Hibernate	Standby	On IDLE	On HIGH	Normal Operation	Recommended Operation
4 PCs per Bench	23.04	23.04	257.28	385.152	531.072	288.96	25.92
3 PCs per Bench	17.28	17.28	192.96	289.104	398.544	216.96	19.68

Figure 5.9 Syslab Total Kilowatts per day.

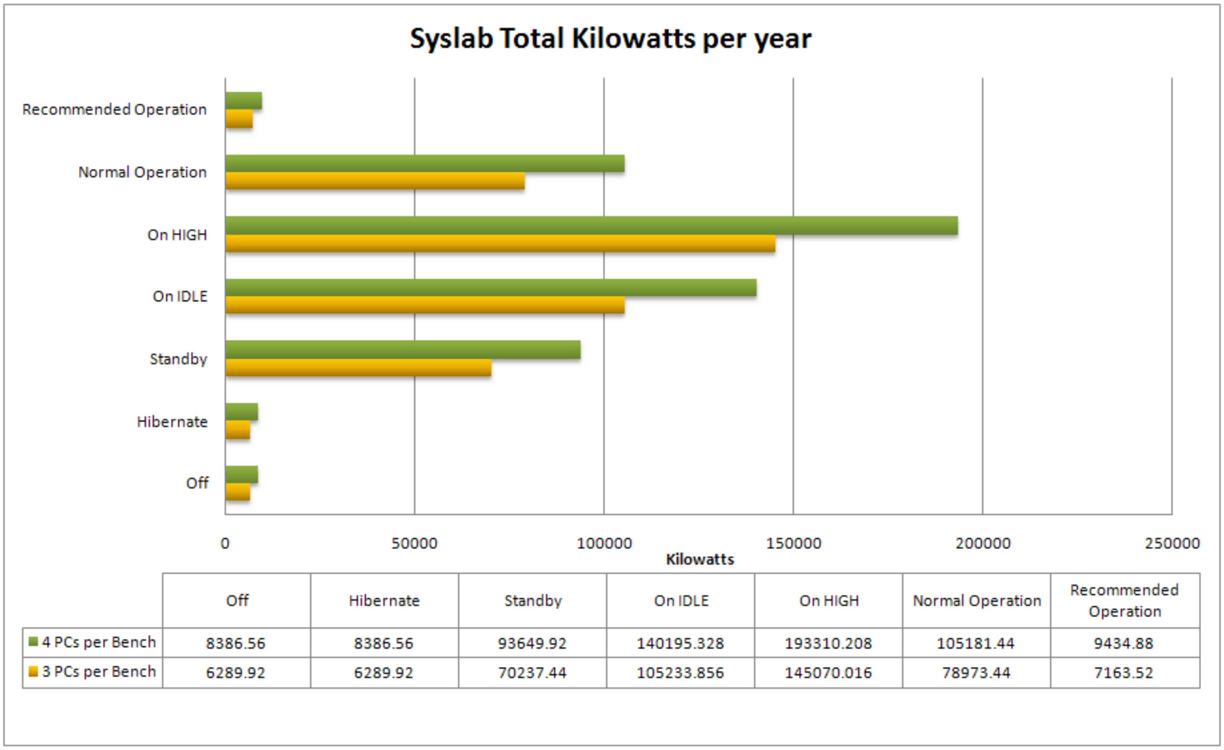

	Off	Hibernate	Standby	On IDLE	On HIGH	Normal Operation	Recommended Operation
4 PCs per Bench	8386.56	8386.56	93649.92	140195.328	193310.208	105181.44	9434.88
3 PCs per Bench	6289.92	6289.92	70237.44	105233.856	145070.016	78973.44	7163.52

Figure 5.10 Syslab Total Kilowatts per year.

5.6.3 Cost Savings

One of the most prominent benefits with Green IT initiatives is the cost savings on daily operation. This benefit catches the C-level staff, administrators and board members alike. Unfortunately, power metrics are often ignored simply because the operating costs are considered as "necessary overhead" and are paid out of a separate financial account. Power costs also have a habit of steadily rising due to the addition of new workstations, servers, etc. This steady rise is extremely difficult to track over time and get a handle on the usage. Reports similar to this section should be created quarterly and reviewed to demonstrate which departments can improve. Holding people accountable for this process is essential for a successful deployment in power reduction.

Cost calculation differs between each location and associated electric provider. Peak and off-peak hours may be identified and charged differently depending on the time of day. Minimum kilowatt charges may apply for the first designated number of kilowatts, then costs change for extra usage. Different tiers exist for residential, commercial and sub-tiers depending on the size of the customer. In general, power cost calculation can be challenging. For this study, the average retail price of electricity to commercial New York customers was used. During April 2009 the average price was 14.24 cents per kilowatthour [3]. Based on this identified cost, the total cost was calculated via the following:

$$Total\,Cost = Total\,Kilowatts\,(hour/day/year) \times Cost\,per\,kWh\,(0.1482)$$

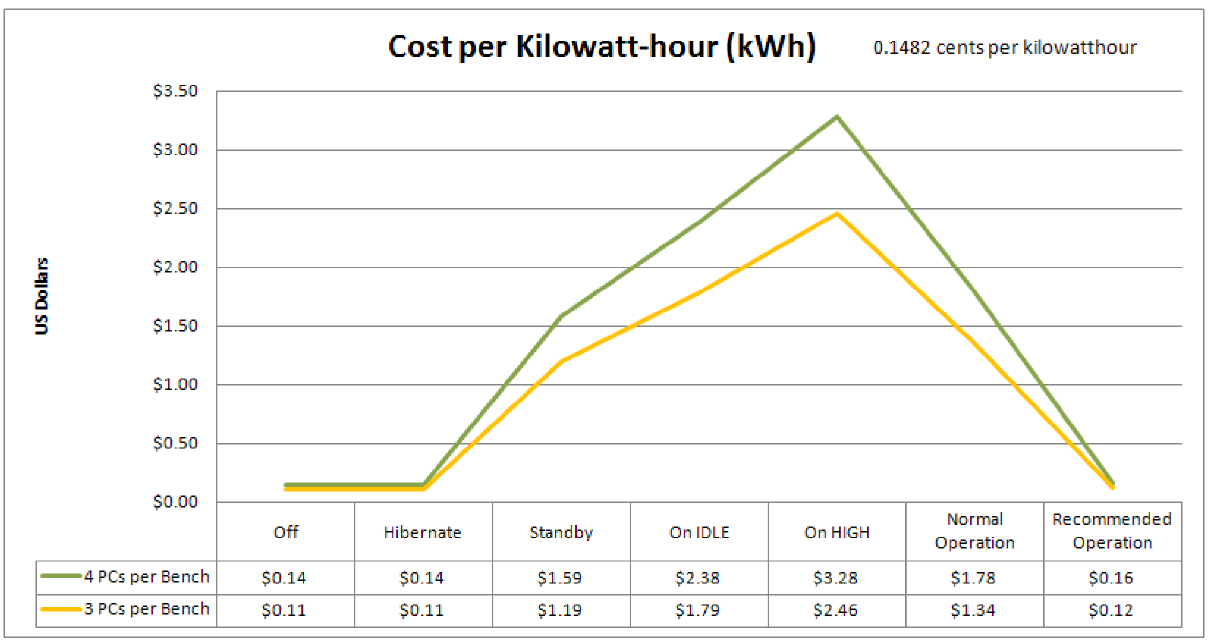

Figure 5.11 Cost per Kilowatt-hour (kWh).

Figure 5.11, 5.12 and 5.13 display the cost per kWh, day and year respectively. Each line graph follows the same overall trend with varying scales. Focusing on

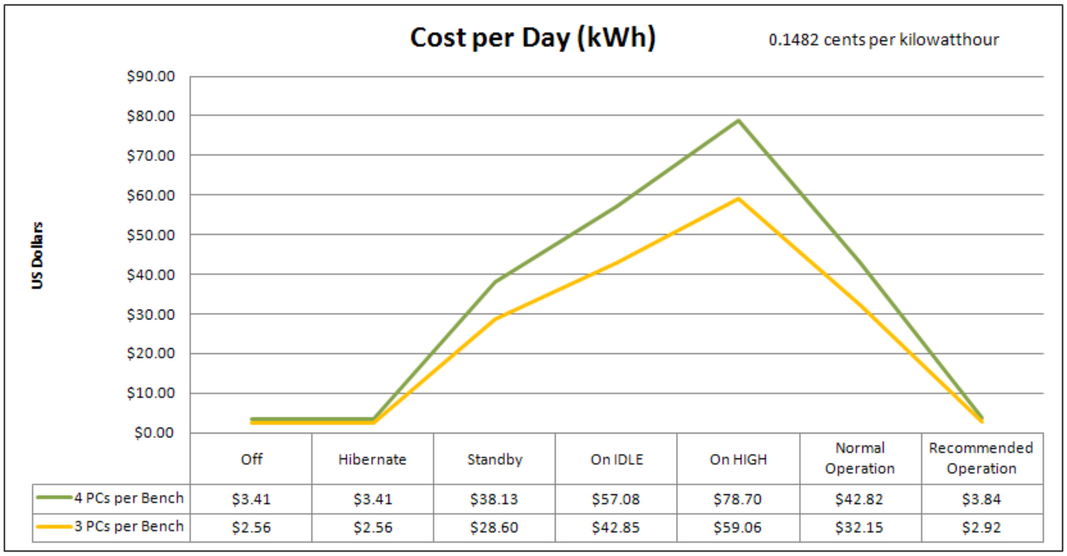

Figure 5.12 Cost per Day (kWh).

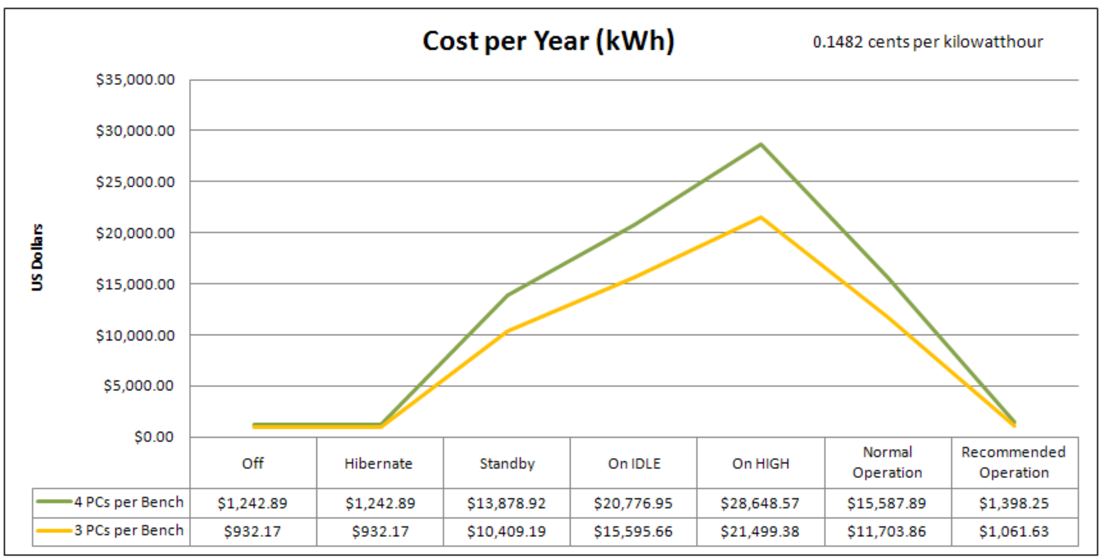

Figure 5.13 Cost per Year (kWh).

Experiment	Total Kilowatts per year	Cents per kWh	Cost
Current Cost - Normal Operation (4PC)	105181.44	0.1482	$15,587.89
Projected Cost - Recommended Operation (3PC)	7163.52	0.1482	$1,061.63
Difference	98017.92		$14,526.26
Reduction in energy	93.19%		

Table 5.10 Normal vs recommended operating cost and energy comparison.

yearly savings, transitioning from the current “normal operation” to the recommended operation will save 98017.92 kWh. This equates to a 93% reduction in energy. When using the New York average cost of electricity, $14,526 can be saved by using a sleep operation and reducing the number of PCs per bench (see Table 5.10).

5.7 Carbon Footprint Analysis

Under normal operating conditions 75.5 metric tons of carbon dioxide are produced by the system administration lab. This can be reduced to 5.1 metric tons by instating the recommended operation.

Table 5.10 identifies 98017.92 kilowatt-hours of electricity saved if recommended operations are followed. The U.S. Environmental Protection Agency provides a Greenhouse Gas Equivalencies Calculator [33]. When using kilowatt-hours of electricity saved from normal to recommended operation, environmental equivalency data is provided. The metrics in Table 5.11 are based on one year of savings. For example, after one year of reducing power consumption to the recommended levels, the NSSA system administration lab can save 70.4 metric tons of CO_2. Specific calculations and statistics are provided by the EPA [34].

70.4 metric tons of avoided CO_2
Annual greenhouse gas emissions from 12.9 passenger vehicles
CO_2 emissions from 7,990 gallons of gasoline consumed
CO_2 emissions from 164 barrels of oil consumed
CO_2 emissions from 0.94 tanker trucks' worth of gasoline
CO_2 emissions from the electricity use of 9.8 homes for one year
CO_2 emissions from the energy use of 6.4 homes for one year
Carbon sequestered by 1,805 tree seedlings grown for 10 years
Carbon sequestered annually by 16 acres of pine or fir forests
Carbon sequestered annually by 0.49 acres of forest preserved from deforestation
CO_2 emissions from 2,933 propane cylinders used for home barbeques
CO_2 emissions from burning 0.37 railcars' worth of coal
Greenhouse gas emissions avoided by recycling 24.3 tons of waste instead of sending it to the landfill

Table 5.11 Equivalency results.

5.8 Cooling

Temperature was measured during the designated test period from academic weeks 5 to 11. Four wireless sensors were placed in the ceiling of the lab to take readings every five minutes and forward the data to a central storage server. The sensors were dispersed throughout the room in strategic locations to capture a targeted quarter of the lab (Figure 5.14). The primary purpose of this data is to identify "hot spots", a situation where not enough cold air is being delivered to the air intake of the computer equipment. Hot spots can lead to equipment failures and unneeded stress on cooling systems.

Sensor:	Temp Average:
SENSOR 1	71.65937669
SENSOR 2	73.32965933
SENSOR 3	70.27346475
SENSOR 4	71.93037572
Total Average	71.74591044

Table 5.12 Temperature averages.

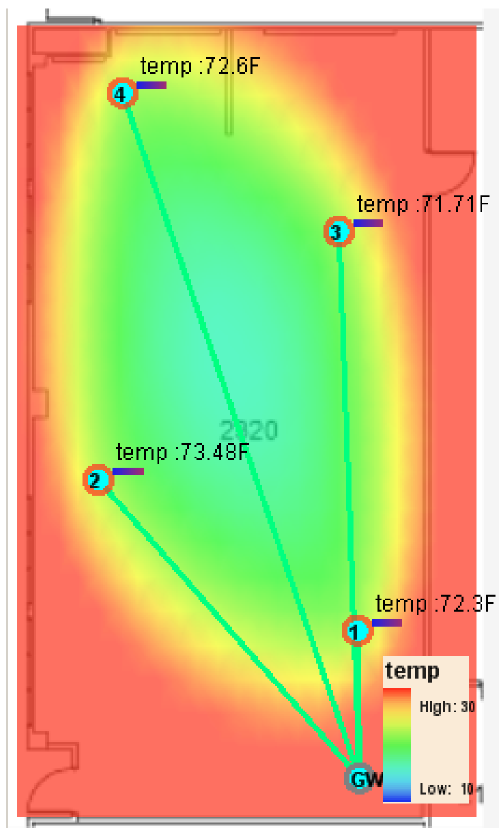

Figure 5.14 Sensor location overlay with floorplan.

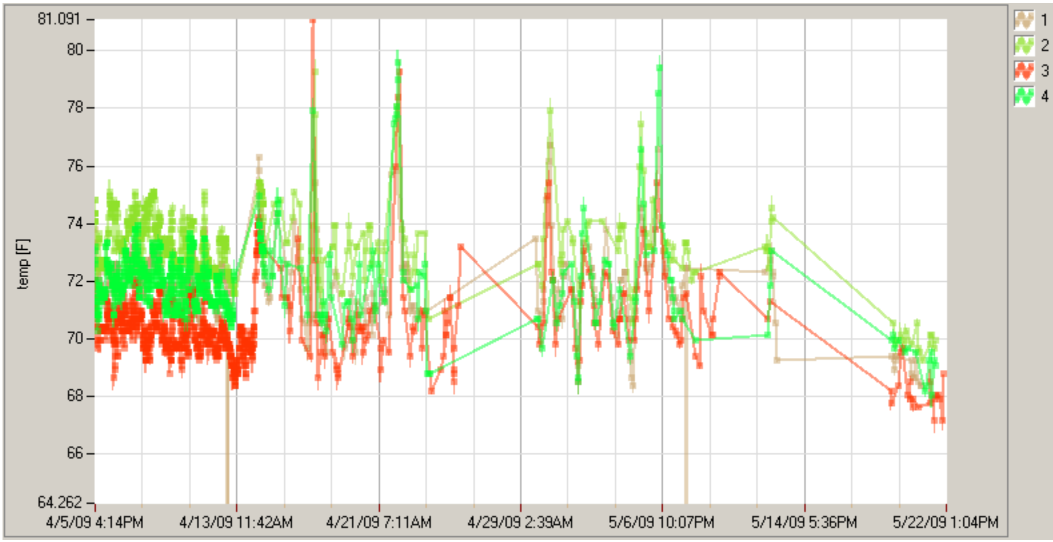

Figure 5.15 Full temperature graph.

Overall, a total of 164,753 data points were collected from the four sensors. Based on this data, an average of 71.74 degrees Fahrenheit was recorded. This is extremely close to the standard targeted room temperature of 72 degrees Fahrenheit. Table 5.12 displays averages per sensor to indicate slight temperature differences throughout the room. Sensors 2 and 3 are slight outliers, averaging several degrees above or below the targeted temperature. Given the dataset, the sensors did not uncover extreme hotspots within the system administration lab. This is primarily due to the advanced AutomatedLogic HVAC air handlers compensating for the varying temperatures.

5.9 Summary

When evaluating the utilization of systems, a very low average of 0.32% CPU load was concluded over course the of the experiment. Based on this metric, a recommendation is made to sleep the inactive workstations to save energy and reduce component stress. The captured utilization further decreases the CADE metric of the lab due to the on-idle state of the systems 24x7. During power consumption testing, configuration inefficiencies were found that resulted in identifying the default sleep state to S1. When adjusting this configuration parameter on Deep Freeze systems, S4 hibernate sleep was found inoperative due to the limitation of the system integrity software. Therefore, this leaves a recommendation of enabling S3 sleep on all capable systems. Significant results are captured when identifying the power consumption differences of S1 and S3 sleep states. When factoring for all lab equipment, a 93% reduction in energy is possible by enabling S3 sleep and reducing the number of PCs per bench to three. This equates to 98017.92 kilowatts saved per year attributing to a projected cost savings of $14,526.26 at 0.14 cents per kilowatt-hour, the New York state average for commercial customers. This savings is equivalent to 70.4 metric tons of CO_2 saved or 7,990 gallons of gasoline consumed. Cooling metrics resulted in limited hotspot identification with a general temperature variance of 2-3 degrees Fahrenheit.

Chapter 6

Recommendations

6.1 S3 Sleep

The primary and most important recommendation of this study is utilizing S3 sleep capabilities of the systems. This will create a 91% reduction in energy usage which translates into significant environmental and financial savings for the Institute.

- For NSSA labs using Deep Freeze, dumppo.exe must be used to configure minimum and maximum sleep for S3 state (see Section 5.1).
  - S4 sleep is not recommended due to the fundamental operation of Deep Freeze (see Section 5.2).
- For other Institute labs, dumppo.exe must be used to configure minimum sleep for S3 and maximum sleep for S4 states.

S3 standby (suspend to RAM) configurations are recommended to enter after 20-30 minutes of inactivity. As tested, exiting standby mode to a fully functional desktop resulted in 3-5 seconds of restore time. Exiting hibernation mode to a fully functional desktop resulted in 30-45 seconds of restore time. The added benefit of hibernation mode is the ability to restore the system state after a power failure. In the NSSA labs, this benefit is debatable due to the use of Deep Freeze, which would destroy the system state after a power failure anyway.

Another added benefit of enabling sleep states is the reduced operating time on the system components. Power supply failure has been a constant problem for quite some time. By reducing the time of operation, components are not stressed for extensive periods of time, thus attributing savings for reduced replacement part cost.

Energy consumption between S3 sleep and S4 hibernate states should also be considered. During the system restoration process, which is initiated upon exiting each state, a wattage spike occurs which is attributed to the processor turning on, hard drives spinning up, etc. The wattage spike is significant for the duration of system restoration, which can range from several seconds for S3 to 30-45 seconds for S4. If the systems are configured for a short hibernation time, students may find themselves powering on the systems too often which causes user frustration and increases the average power consumption of the lab.

S3 standby state is highly recommended. The short restore time and added energy conservation results in a win-win situation. If the hibernation state is pursued in non-Deep Freeze labs, students may voice their dissatisfaction of the added 30-45 seconds for the restore time.

6.1.1 Energy Star EZ GPO

For Windows Active Directory domains, Energy Star provides a free utility to allow centralized control of user power management settings [35]. This tool is targeted for Windows 2000 and XP deployments due to the lack of Group Policy Objects (GPOs) in the operating system. Starting with Windows Vista, GPOs are built-in for this functionality. EZ-GPO, combined with Windows Task Scheduler can provide significant cost savings while still maintaining patching windows and providing centralized control. http://www.energystar.gov/index.cfm?c=power_mgt.pr_power_mgt_ez_gpo

6.2 Low-Hanging Fruit

- Screensavers are not recommended, they do not conserve energy and may prevent the system from entering standby states. It is recommended that PCs and monitors continue to enter standby after 20-30 minutes of inactivity.

- During holiday breaks, or extended periods of minimal usage, labs should be powered down.

- Adopt Energy Star procurement strategies. Recommending the purchase of Energy Star products to organization procurement administrators will result in the support of green products and extended usage of the product until end of life.

- Teach user awareness to employees and students to proactively turn off devices. User awareness is essential for developing successful green programs. Most focus will be brought if administrators or departments are held accountable for power usage. Awards and recognition can be given to departments that meet or exceed green goals.

- Mandate system administrators to enable energy efficient settings on operating systems.

Other universities have performed similar sustainable acts, RIT has the opportunity to reap great benefits.

6.3 Long-Term Sustainability Strategies

- Develop a full shutdown procedure for each lab that is simple to implement and execute. For example, a shutdown feature in Deep Freeze can be used to power down the lab from the administration console. Power should be then cut to the lab to prevent power usage from vampire devices. To wake the systems up the following morning, a wake-on-LAN script can be configured and scheduled for minimal user interaction. The tools and technology are readily available to meet these goals.

- Granular energy costs. Clearly identifying the energy cost of each device will greatly assist in retiring inefficient devices, or creating a reduced operating schedule. Google and Microsoft are working to provide this functionality to the average user with PowerMeter and Hohm.

- E-waste. Electronic components are quickly deemed obsolete and moved to the dumpster to make way for new technology. A retirement plan must be developed to include responsible e-waste disposal.

- Sustainable practices must be included from the beginning to enable a foundation of green IT strategies. Planning for power attributes and efficient cooling is essential to reduce long-term operating costs and the environmental footprint.

- Manufacturers:
  - Develop devices that do not use vampire power when “off”.
  - Develop power strips that have the capabilities to cut power automatically to vampire devices when not in use.

6.4 Virtualization

The virtualization of operating systems is the number one green initiative undertaken by corporations. This is due to the enormous benefits of operating system portability and server consolidation, which consequently reduces strain on cooling and power distribution equipment.

During the start of school year 2008-2009, several NSSA courses were redesigned to integrate closely with VMware Workstation. Although not specifically mentioned during the transition, the migration to VMware assisted with sustainable techniques. After one academic year of virtualizing the labs, students primarily used two out of the four workstations available per bench. Allowing students to virtualize multiple operating systems on one physical host eliminates the need for additional clients.

This unforeseen side effect of virtualizing labs helps in a number of ways:

- Reducing the required power by 23% (reduction from 4 to 3 PCs see Figure 5.13)
- Reducing the required cooling capacity by 25% (reduction from 4 to 3 PCs)
- Allowing spare systems to serve as extra parts or hot swaps.
- Adding additional workspace to an already crowded workbench, allowing better airflow across PCs.
- Reducing the PC requirement allows further distribution of LCD monitors and proper retirement of outdated CRT monitors.
- By using linked clones as diff images, additional storage is no longer needed due to the drastic reduction in file size.

Moving to virtual labs should be a priority for similar system administration courses due to the above benefits and the wide adoption across industry. If deemed viable in the future, PCs should be further reduced to two per bench.

6.4.1 Remote Laboratory Emulation System (RLES)

The NSSA department also employs the use of a Remote Laboratory Emulation System (RLES) to further virtualize lab work for remote access. This system is primarily used for graduate level courses, but does possess sufficient computing resources for undergraduate work. This system in itself has several green qualities.

First, the system is clustered using VMware VirtualCenter software to provide a segregated environment for students and faculty while hosting a wide assortment of courses. The use of virtualization software enables many users to use a single system with performance quotas given to each user. A similar non-virtualized system would require a physical host system for each operating system needed.

Second, the physical location of the cluster is co-located with RIT's Information Technology Services (ITS) datacenter. Co-location refers to sharing space, in this case the NSSA department is sharing space with RIT's ITS datacenter. This option provides several advantages such as sufficient power, cooling, security and facilities management. ITS currently plans to virtualize most of their operating environment to further reduce power and cooling consumption for the Institute.

These points are made to show recommended sustainable strategies that the NSSA department has made.

Chapter 7

Conclusion

The research and findings of this study conclude that enacting green initiatives in non-conventional, workstation environments not only benefit cost savings, but the people and the planet.

As stated in the initial hypothesis, the goal of this thesis is to validate increasing the efficiency of system utilization and measuring the resulting economical and ecological advantages through the use of empirical data. This study found system utilization efficiency to be currently extremely low, a total average of 0.32%. Based on this assessment, green initiatives such as enabling energy saving states can be enacted to take advantage of inefficient utilization. Increasing student utilization through the use of additional lab courses or assignments will help to offset the current low utilization. However, the important aspect of this study resides in the power management strategy to be followed when the systems are not intended to be used such as overnight and academic break weeks.

When leveraging the use of energy saving sleep states, a 91% reduction in power usage is possible per PC. Removing one PC per bench further increases the power savings to 93%. This equates to 98017.92 kilowatts saved per year attributing to a projected cost savings of $14,526.26 at 0.14 cents per kilowatt-hour, the New York state average for commercial customers. These findings confirm the first portion of the hypothesis relating to economical advantages.

Reduction in energy usage favorably affects ecological aspects due to avoided CO_2 emissions. After one year of reducing power consumption to the recommended levels, the NSSA system administration lab can save 70.4 metric tons of CO_2. These findings confirm the second and final portion of the hypothesis relating to ecological advantages.

The data presented in this study supports the decision to reduce the number of PCs and properly manage power consumption for the remaining systems. Running a clean and energy efficient lab should be the goal. Although this study focused on a single large lab, the savings can be multiplied through remaining NSSA labs and eventually conveyed across all Institute labs.

The recommendations described in Chapter 6 do not require significant financial investments. The changes can be performed using existing technologies and trusted, free tools. Administrators can integrate these recommendations into a procurement refresh or an annual operating system update.

7.1 Future Work

RIT's Information Technology Services are following their own sustainable techniques through the use of virtualization and tracking energy usage. Currently an active kW graph exists to monitor live energy usage http://eve.rit.edu/~slpits/. This is a prominent initiative, but granular views and control should be added to make this tool more effective.

While data centers are the primary focus for corporations, general workstations should be a greater concern on our campus due to the high technological aspect of the Institute. Administrators, support staff and a vast number of students live on this campus of technological innovation.

Individual energy monitoring should be pursued to increase user awareness for administrators. This will provide departments with traceable data points to focus on energy usage for their particular environment.

Bibliography

[1] E. Rosenthal and A. C. Revkin, “Science Panel Calls Global Warming ’Unequivocal’,”.

[2] A. D. Little, “Energy Consumption by Office and Telecommunications Equipment,” Technical report, U.S. Department of Energy (2002) .

[3] U. D. of Energy, “Average Retail Price of Electricity to Ultimate Customers by End-Use Sector, by State,”, 2009.

[4] G. Consulting, “Gartner Identifies the Top 10 Strategic Technologies for 2008,” Press Release (2007).

[5] G. Consulting, “Gartner Says Data Centres Account for 23 Per Cent of Global ITC CO2 Emissions,” Press Release (2007).

[6] M. James, “Is green IT an illusion?,” ComputerWeekly.com (2007).

[7] V. Corporation, “IT Unplugged: Slash PC Power Use at the Network Level,” Technical report (2009) .

[8] T. S. of the American College University Presidents’ Climate Commitment, “American College University Presidents’ Climate Commitment,”, 2009.

[9] “RIT Joins in Climate Commitment,”, 2009.

[10] C. Baroudi, J. Hill, A. Reinhold, and J. Senxian, *Green IT For Dummies* (Wiley Publishing, 2009).

[11] G. Staff, “Power Management Project May Save Indiana University $500K a Year,” Greener Computing (2009).

[12] C. S. C. Initiative, “About Power Down for the Planet,” (2009).

[13] V. Corporation, “Lake Washington School District Lowers Greenhouse Gas Emissions and Saves over $200,000 Annually,” (2008).

[14] ACPI, “Advanced Configuration Power Interface,”.

[15] Y. Agarwal, S. Hodges, R. Chandra, J. Scott, P. Bahl, and R. Gupta, “Somniloquy: Augmenting Network Interfaces to Reduce PC Energy Usage,” (2009).

[16] G. Corp., “Company Overview,”.

[17] G. Corp., “Google PowerMeter,”.

[18] M. Corp, “Microsoft Hohm,”.

[19] E. STAR, “ENERGY STAR Program Requirements for Displays,”.

[20] J. Ng, “New Energy Star 5.0 Specs for Computers Become Effective Today,” DailyTech (2009).

[21] E. Star, “Energy Star and Other Climate Protection Partnerships 2007 Annual Report,” Technical report (2007) .

[22] T. G. Grid, “About The Green Grid,”.

[23] G. Staff, “Keysource Cuts Energy Use 45 Percent at U.K. Data Center,”.

[24] M. Christian Belady, A. Andy Rawson, D. John Pfleuger, and S. Tahir Cader, "Green Grid Data Center Power Efficiency Metrics: PUE and DCIE," Technical report, The Green Grid (2008) .

[25] H. Chan and G. More, "The Whys and Hows of Measuring Power in your Data Center,".

[26] J. Kaplan, W. Forrest, and N. Kindler, "Revolutionizing Data Center Energy Efficiency," Technical report, McKinsey Company (2008) .

[27] R. G. R. I. of Technology, "Dept. home page,".

[28] P. International, "Kill A Watt PS,".

[29] M. D. Network, "Windows Driver Kit: Kernel-Mode Driver Architecture System Sleeping States,".

[30] F. Corporation, "Faronics Deep Freeze Windows Editions,".

[31] T. Economist, "Pulling the plug on standby power,".

[32] S. Zone, "SiSoftware Sandra,".

[33] U. E. P. Agency, "Greenhouse Gas Equivalencies Calculator,".

[34] U. E. P. Agency, "Calculations and References,".

[35] E. Star, "EZ GPO,".

[36] *The New Oxford American Dictionary, Second Edition*, E. McKean, ed., (Oxford University Press, New York, 2005).

Appendix A

Case Study

A.1 System Administration Lab

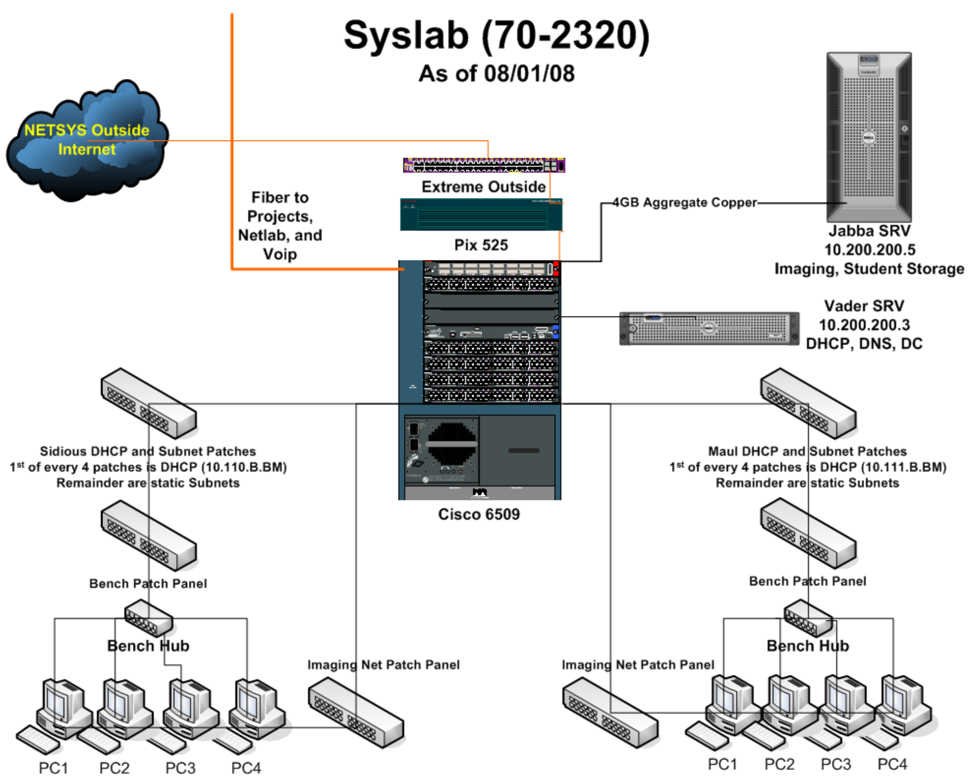

Figure A.1 System Administration lab topology.

Syslab Devices	**Quantity**
Benches	20
PCs	80
CRT Monitors (KDS XFLAT)	52
LCD Monitors (PX191)	28

Table A.1 Syslab devices.

CPU	Pentium 4 3.40GHz
Memory	3.0 GB RAM
Motherboard	ASUS P5AD2-E PREMIUM
Hard Drive	120GB

Table A.2 Syslab Antec PC System Specifications.

Antec Model TPII-380 TRUEPOWER 2.0 - 380 WATT PSU						
AC INPUT 115V/10A; 60Hz/50Hz						
DC OUTPUT	+5V	+12V1	+12V2	+3.3V	-12V	+5V SB
MAX	35A	16A	16A	28A	1.0A	2.0A
MIN	0.5A	0.4A	0.4A	0.5A	0A	0A
MAX LOAD 360W / 28A						

Table A.3 Syslab Antec PSU Specifications.

Scheme	**Setting**
Turn off monitor	20 mins
turn off hard disks	Never
System standby	Never
Hibernation	Never

Table A.4 Configured power options.

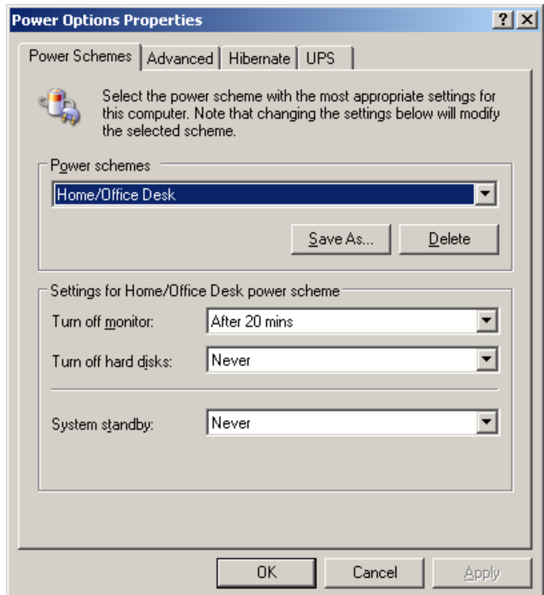

Figure A.2 Current power configuration.

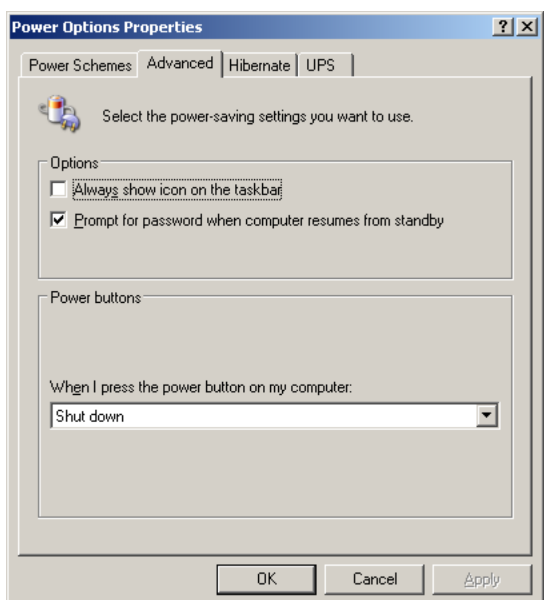

Figure A.3 Current power configuration.

Output 4 Current Power Configuration.

```
C:\>dumppo.exe cap
power capabilties
System power capabilties
Power Button Present....: TRUE
Sleep Button Present....: FALSE
Lid Present.............: FALSE
System states supported.: S1 S4 S5
Hiber file reserved.....: FALSE
Thermal control.........: FALSE
CPU Throttle control....: FALSE
Processor min throttle..: 100
Processor trottle scale.: 100 (1%)
Some disk will spindown.: TRUE
System batteries present: FALSE
System batteries scale..: (G:0 C:0) (G:0 C:0) (G:0 C:0)
Ac on line wake ability.: Unspecified
Lid wake ability........: Unspecified
RTC wake ability........: S4 - hibernate
Min device wake.........: Unspecified
Default low latency wake: Unspecified

C:\>dumppo.exe admin
Admin policy overrides
Min sleep state......: S1
Max sleep state......: S4 - hibernate
Min video timeout....: 0
Max video timeout....: -1
Min spindown timeout.: 0
Max spindown timeout.: -1
```

Appendix B

Sample Data

B.1 Key Terms

Volt - the SI unit of electromotive force, the difference of potential that would drive one ampere of current against one ohm resistance [36].

Ampere - a unit of electric current equal to a flow of one coulomb per second. The SI base unit of electric current, 1 ampere is precisely defined as that constant current which, if maintained in two straight parallel conductors of infinite length, of negligible circular cross-section, and placed 1 meter apart in a vacuum, would produce between these conductors a force of 2 10 7 newton per meter [36].

Watt - the SI unit of power, equivalent to one joule per second, corresponding to the power in an electric circuit in which the potential difference is one volt and the current one ampere [36].

Power Factor - the ratio of the actual electrical power dissipated by an AC circuit to the product of the r.m.s. values of current and voltage. The difference between the two is caused by reactance in the circuit and represents power that does no useful work [36].

1 kilowatt = 1000 watts

1 kilowatt-hour - represents the usage of one kilowatt of energy for an hour. This is the basic unit by which electricity is charged. If you run a 100-watt electrical bulb for 10 hours, you have used 1kWh of energy.

B.2 System Power States

The following section outlines the detail in regards to the difference in each sleep level [29].

System Power State S1

System power state S1 is a sleeping state with the following characteristics:

Power consumption - Less consumption than in S0 and greater than in the other sleep states. Processor clock is off and bus clocks are stopped.

Software resumption - Control restarts where it left off.

Hardware latency - Typically no more than two seconds.

System hardware context - All context retained and maintained by hardware.

System Power State S2

System power state S2 is similar to S1 except that the CPU context and contents of the system cache are lost because the processor loses power. State S2 has the following characteristics:

Power consumption - Less consumption than in state S1 and greater than in S3. Processor is off. Bus clocks are stopped; some buses might lose power.

Software resumption - After wake-up, control starts from the processor's reset vector.

Hardware latency - Two seconds or more; greater than or equal to the latency for S1.

System hardware context - CPU context and system cache contents are lost.

System Power State S3

System power state S3 is a sleeping state with the following characteristics:

Power consumption - Less consumption than in state S2. Processor is off and some chips on the motherboard also might be off.

Software resumption - After the wake-up event, control starts from the processor's reset vector.

Hardware latency - Almost indistinguishable from S2.

System hardware context - Only system memory is retained. CPU context, cache contents, and chipset context are lost.

System Power State S4

System power state S4, the hibernate state, is the lowest-powered sleeping state and has the longest wake-up latency. To reduce power consumption to a minimum, the hardware powers off all devices. Operating system context, however, is maintained in a hibernate file (an image of memory) that the system writes to disk before entering the S4 state. Upon restart, the loader reads this file and jumps to the system's previous, prehibernation location.

If a computer in state S1, S2, or S3 loses all AC or battery power, it loses system hardware context and therefore must reboot to return to S0. A computer in state S4, however, can restart from its previous location even after it loses battery or AC power because operating system context is retained in the hibernate file. A computer in the hibernate state uses no power (with the possible exception of trickle current).

Power consumption - Off, except for trickle current to the power button and similar devices.

Software resumption - System restarts from the saved hibernate file. If the hibernate file cannot be loaded, rebooting is required. Reconfiguring the hardware while the system is in the S4 state might result in changes that prevent the hibernate file from loading correctly.

Hardware latency - Long and undefined. Only physical interaction returns the system to the working state. Such interaction might include the user pressing the ON switch or, if the appropriate hardware is present and wake-up is enabled, an incoming ring for the modem or activity on a LAN. The machine can also awaken from a resume timer if the hardware supports it.

System hardware context - None retained in hardware. The system writes an image of memory in the hibernate file before powering down. When the operating system is loaded, it reads this file and jumps to its previous location.

B.3 Workstation Availability

Node (RIGHT)	IP Address	Average Availability	Node (LEFT)	IP Address	Average Availability
MAULS011	10.200.251.11	46.01%	SIDIO011	10.200.250.11	5.36%
MAULS012	10.200.251.12	2.52%	SIDIO012	10.200.250.12	81.21%
MAULS013	10.200.251.13	0.03%	SIDIO013	10.200.250.13	75.24%
MAULS014	10.200.251.14	57.17%	SIDIO014	10.200.250.14	30.67%
MAULS016	10.200.250.17	42.66%	SIDIO021	10.200.250.21	46.99%
MAULS021	10.200.251.21	1.17%	SIDIO022	10.200.250.22	30.03%
MAULS022	10.200.251.22	57.94%	SIDIO023	10.200.250.23	61.91%
MAULS023	10.200.251.23	55.56%	SIDIO024	10.200.250.24	51.27%
MAULS024	10.200.251.24	55.79%	SIDIO031	10.200.250.31	37.40%
MAULS031	10.200.251.31	49.68%	SIDIO032	10.200.250.32	59.28%
MAULS032	10.200.251.32	44.17%	SIDIO033	10.200.250.33	54.75%
MAULS033	10.200.251.33	49.19%	SIDIO034	10.200.250.34	42.20%
MAULS034	10.200.251.34	31.70%	SIDIO041	10.200.250.41	50.08%
MAULS041	10.200.251.41	47.69%	SIDIO042	10.200.250.42	44.63%
MAULS042	10.200.251.42	50.78%	SIDIO043	10.200.250.43	40.51%
MAULS043	10.200.251.43	42.94%	SIDIO044	10.200.250.44	56.24%
MAULS044	10.200.251.44	35.47%	SIDIO051	10.200.250.52	60.18%
MAULS061	10.200.251.61	44.61%	SIDIO052	10.200.250.51	28.12%
MAULS062	10.200.251.62	47.18%	SIDIO053	10.200.250.53	39.70%
MAULS063	10.200.251.63	48.07%	SIDIO054	10.200.250.54	38.59%
MAULS064	10.200.251.64	5.11%	SIDIO061	10.200.250.61	69.18%
MAULS071	10.200.251.71	51.72%	SIDIO062	10.200.250.62	77.04%
MAULS072	10.200.251.72	50.32%	SIDIO063	10.200.250.63	73.00%
MAULS073	10.200.251.73	0.23%	SIDIO064	10.200.250.64	63.46%
MAULS074	10.200.251.74	49.25%	SIDIO071	10.200.250.71	58.56%
MAULS081	10.200.251.81	45.77%	SIDIO072	10.200.250.72	84.99%
MAULS082	10.200.251.82	51.32%	SIDIO073	10.200.250.73	69.33%
MAULS083	10.200.251.83	50.93%	SIDIO074	10.200.250.74	70.58%
MAULS084	10.200.251.84	48.70%	SIDIO081	10.200.250.81	78.38%
MAULS091	10.200.251.91	27.40%	SIDIO082	10.200.250.82	76.06%
MAULS092	10.200.251.92	56.96%	SIDIO083	10.200.250.83	78.03%
MAULS093	10.200.251.93	49.09%	SIDIO084	10.200.250.84	72.53%
MAULS094	10.200.251.94	60.70%	SIDIO092	10.200.250.92	73.79%
MAULS101	10.200.251.101	49.96%	SIDIO093	10.200.250.93	28.31%
MAULS102	10.200.251.102	47.82%	SIDIO094	10.200.250.94	65.21%
MAULS103	10.200.251.103	47.95%	SIDIO101	10.200.250.101	50.15%
MAULS104	10.200.251.104	31.82%	SIDIO102	10.200.250.102	65.09%
MAUL AVG		41.50%	SIDIO103	10.200.250.103	66.33%
			SIDIO104	10.200.250.104	21.77%
TOTAL AVERAGE		47.72%	SIDIO111	10.200.250.111	22.08%
			SIDIO113	10.200.250.113	11.62%
10.200.250.254	10.200.250.254	94.55%	SIDIO114	10.200.250.114	24.50%
10.200.251.254	10.200.251.254	77.24%	SIDIOUS AVG		53.20%
esx3.vm.nssa.labs	10.200.201.30	100.00%			
MONSTER	10.200.202.18	100.00%			
REDBULL	10.200.202.14	99.99%			
SPRITE	10.200.202.17	100.00%			
SURGE	10.200.202.13	100.00%			
SYSLABCORE.netsys.labs	172.30.1.253	100.00%			
INFRASTRUCTURE AVG		96.47%			

Table B.1 Syslab Workstation Availability - Sample Data.

B.4 Workstation Utilization Sample Data

Node (RIGHT)	Average CPU Load	Peak CPU Load	Node (LEFT)	Average CPU Load	Peak CPU Load
MAULS011	1.00%	95.00%	SIDIO011	0.00%	58.00%
MAULS012	0.00%	100.00%	SIDIO012	1.00%	99.00%
MAULS014	0.00%	92.00%	SIDIO013	0.00%	62.00%
MAULS016	0.00%	53.00%	SIDIO014	0.00%	94.00%
MAULS022	1.00%	97.00%	SIDIO021	0.00%	53.00%
MAULS023	0.00%	87.00%	SIDIO022	0.00%	60.00%
MAULS024	0.00%	90.00%	SIDIO023	0.00%	100.00%
MAULS031	0.00%	93.00%	SIDIO024	0.00%	94.00%
MAULS032	1.00%	91.00%	SIDIO031	0.00%	52.00%
MAULS033	0.00%	85.00%	SIDIO032	0.00%	94.00%
MAULS034	0.00%	91.00%	SIDIO033	1.00%	71.00%
MAULS041	0.00%	40.00%	SIDIO034	0.00%	39.00%
MAULS042	0.00%	100.00%	SIDIO041	1.00%	100.00%
MAULS043	0.00%	87.00%	SIDIO042	1.00%	92.00%
MAULS044	0.00%	64.00%	SIDIO043	2.00%	98.00%
MAULS061	0.00%	49.00%	SIDIO044	0.00%	52.00%
MAULS062	0.00%	50.00%	SIDIO051	0.00%	76.00%
MAULS063	0.00%	100.00%	SIDIO052	0.00%	59.00%
MAULS064	0.00%	76.00%	SIDIO053	0.00%	72.00%
MAULS071	0.00%	90.00%	SIDIO054	0.00%	55.00%
MAULS072	0.00%	100.00%	SIDIO061	1.00%	97.00%
MAULS073	2.00%	35.00%	SIDIO062	0.00%	92.00%
MAULS074	0.00%	88.00%	SIDIO063	1.00%	55.00%
MAULS081	1.00%	91.00%	SIDIO064	0.00%	93.00%
MAULS082	0.00%	87.00%	SIDIO071	0.00%	100.00%
MAULS083	0.00%	95.00%	SIDIO072	0.00%	96.00%
MAULS084	0.00%	88.00%	SIDIO073	1.00%	91.00%
MAULS091	0.00%	100.00%	SIDIO074	0.00%	50.00%
MAULS092	0.00%	100.00%	SIDIO081	0.00%	100.00%
MAULS093	0.00%	96.00%	SIDIO082	1.00%	98.00%
MAULS094	0.00%	88.00%	SIDIO083	2.00%	95.00%
MAULS101	0.00%	88.00%	SIDIO084	0.00%	65.00%
MAULS102	1.00%	61.00%	SIDIO092	2.00%	83.00%
MAULS103	0.00%	60.00%	SIDIO093	1.00%	68.00%
MAULS104	0.00%	65.00%	SIDIO094	0.00%	90.00%
MAUL AVG	0.20%	82.06%	SIDIO101	0.00%	91.00%
			SIDIO102	0.00%	91.00%
TOTAL AVERAGE	0.32%	80.65%	SIDIO103	1.00%	84.00%
			SIDIO104	0.00%	99.00%
SYSLABCORE.netsys.labs	2.00%	9.00%	SIDIO111	0.00%	36.00%
esx3.vm.nssa.labs	16.00%	14589.00%	SIDIO113	2.00%	93.00%
INFRASTRUCTURE AVG	9.00%		SIDIO114	0.00%	91.00%
			SIDIOUS AVG	0.43%	79.48%

Table B.2 Syslab Workstation CPU Utilization - Sample Data.

B.5 Power Projections Sample Data

Experiment	Device	Mode of Op.	Quantity	Watt	Total Watt	kWh	Kilowatts per day	Killowatts per year
Off	PCs	Off	80	12	960	0.96	23.04	8386.56
	CRTs	Off	52	0	67	0.067	1.608	585.312
	LCDs	Off	28	0	0	0	0	0
	Hubs	Off	20	0	0	0	0	0
Total kWh					1027	1.027	24.648	8971.872
Total Cost						$0.15	$3.65	$1,329.63
Hibernate	PCs	Hibernate	80	12	960	0.96	23.04	8386.56
	CRTs	Standby	52	0	0	0	0	0
	LCDs	Standby	28	0	0	0	0	0
	Hubs	Off	20	0	0	0	0	0
					960	0.96	23.04	8386.56
Total Cost						$0.14	$3.41	$1,242.89
Standby	PCs	Standby	80	134	10720	10.72	257.28	93649.92
	CRTs	Standby	52	0	0	0	0	0
	LCDs	Standby	28	0	0	0	0	0
	Hubs	Off	20	0	0	0	0	0
Total kWh					10720	10.72	257.28	93649.92
Total Cost						$1.59	$38.13	$13,878.92
On IDLE	PCs	On - IDLE	80	150	12000	12	288	104832
	CRTs	On	52	55	2860	2.86	68.64	24984.96
	LCDs	On	28	41	1148	1.148	27.552	10028.928
	Hubs	On	20	2	40	0.04	0.96	349.44
Total kWh					16048	16.048	385.152	140195.328
Total Cost						$2.38	$57.08	$20,776.95

Table B.3 Syslab Total Power Projections - Sample Data.

Experiment	Device	Mode of Op.	Quantity	Watt	Total Watt	kWh	Kilowatts per day	Killowatts per year
On HIGH	PCs	On - HIGH	80	226	18080	18.08	433.92	157946.88
	CRTs	On	52	55	2860	2.86	68.64	24984.96
	LCDs	On	28	41	1148	1.148	27.552	10028.928
	Hubs	On	20	2	40	0.04	0.96	349.44
Total kWh					22128	22.128	531.072	193310.208
Total Cost						$3.28	$78.70	$28,648.57
Normal Operation	PCs	On - IDLE	80	150	12000	12	288	104832
	CRTs	Standby	52	0	0	0	0	0
	LCDs	Standby	28	0	0	0	0	0
	Hubs	On	20	2	40	0.04	0.96	349.44
Total kWh					12040	12.04	288.96	105181.44
Total Cost						$1.78	$42.82	$15,587.89
Recommended Operation	PCs	S3 Sleep	80	13	1040	1.04	24.96	9085.44
	CRTs	Standby	52	0	0	0	0	0
	LCDs	Standby	28	0	0	0	0	0
	Hubs	On	20	2	40	0.04	0.96	349.44
Total kWh					1080	1.08	25.92	9434.88
Total Cost						$0.16	$3.84	$1,398.25
Other Devices	3COM OfficeConnect Dual Speed Hub 8	On	19	2	38	0.038	0.912	331.968
	NETGEAR 4 PORT DS104	On	19	3	57	0.057	1.368	497.952
	2950 Switch	On	19	22	418	0.418	10.032	3651.648
	2514	On	19	26	494	0.494	11.856	4315.584
	2621XM	On	19	17	323	0.323	7.752	2821.728
Total kWh					1330	1.33	31.92	11618.88
Total Cost						$0.20	$4.73	$1,721.92

Table B.4 Syslab Total Power Projections - Sample Data.

B.6 Wireless Sensor Network Raw Data

Id	Time	parent	voltage [V]	temp [F]	light	accel_x [g]	accel_y [g]	mag_x [mgauss]	mag_y [mgauss]	mic
2	4/11/2009 19:34	0	2.4556	73.335	0	-0.26	0.12	32.276	32.546	409
3	4/11/2009 19:34	0	2.413	70.538	0	-0.06	0.2	32.681	32.681	414
1	4/11/2009 19:34	0	2.4177	72.156	0	-0.3	-0.06	32.411	32.276	417
4	4/11/2009 19:34	0	2.3764	72.008	0	0.22	-0.02	32.816	32.681	423
2	4/11/2009 19:34	0	2.4556	73.335	0	-0.26	0.12	32.141	32.546	414
3	4/11/2009 19:34	0	2.4177	70.538	0	-0.06	0.2	32.681	32.681	419
1	4/11/2009 19:34	0	2.4177	72.156	0	-0.3	-0.06	32.411	32.276	414
2	4/11/2009 19:34	0	2.4556	73.335	0	-0.26	0.1	32.141	32.546	409
4	4/11/2009 19:34	0	2.3719	72.008	0	0.22	-0.02	32.816	32.681	420
3	4/11/2009 19:34	0	2.413	70.538	0	-0.06	0.2	32.681	32.681	419
1	4/11/2009 19:35	0	2.4223	72.156	0	-0.3	-0.06	32.411	32.276	409
2	4/11/2009 19:35	0	2.4556	73.335	0	-0.26	0.12	32.141	32.546	411
4	4/11/2009 19:35	0	2.3719	72.008	0	0.22	-0.02	32.816	32.681	421
3	4/11/2009 19:35	0	2.4177	70.538	0	-0.08	0.2	32.681	32.546	415
1	4/11/2009 19:35	0	2.4177	72.156	0	-0.3	-0.06	32.411	32.276	409
2	4/11/2009 19:35	0	2.4556	73.335	0	-0.26	0.12	32.141	32.546	413
4	4/11/2009 19:35	0	2.3719	72.008	0	0.22	-0.02	32.816	32.816	416
3	4/11/2009 19:35	0	2.4177	70.538	0	-0.08	0.2	32.681	32.681	416
1	4/11/2009 19:35	0	2.4177	72.156	0	-0.3	-0.06	32.411	32.276	409
2	4/11/2009 19:35	0	2.4556	73.335	0	-0.26	0.1	32.141	32.546	413
4	4/11/2009 19:35	0	2.3719	72.008	0	0.22	-0.02	32.816	32.681	418
3	4/11/2009 19:35	0	2.4177	70.538	0	-0.06	0.2	32.681	32.681	413
1	4/11/2009 19:35	0	2.4177	72.303	0	-0.3	-0.06	32.411	32.276	410
2	4/11/2009 19:35	0	2.4556	73.335	0	-0.26	0.12	32.141	32.546	413
4	4/11/2009 19:36	0	2.3719	72.008	0	0.22	-0.02	32.816	32.681	422
3	4/11/2009 19:36	0	2.4177	70.538	0	-0.08	0.2	32.681	32.681	415
1	4/11/2009 19:36	0	2.4177	72.303	0	-0.3	-0.06	32.411	32.276	409
2	4/11/2009 19:36	0	2.4604	73.335	0	-0.26	0.12	32.141	32.546	414
4	4/11/2009 19:36	0	2.3719	72.008	0	0.2	-0.02	32.816	32.681	417
3	4/11/2009 19:36	0	2.413	70.538	0	-0.08	0.2	32.681	32.681	420
1	4/11/2009 19:36	0	2.4223	72.303	0	-0.3	-0.06	32.411	32.276	412
2	4/11/2009 19:36	0	2.4556	73.335	0	-0.28	0.12	32.141	32.546	419
4	4/11/2009 19:36	0	2.3719	72.008	0	0.22	-0.02	32.816	32.681	422
3	4/11/2009 19:36	0	2.413	70.538	0	-0.06	0.2	32.681	32.681	417
1	4/11/2009 19:36	0	2.4177	72.303	0	-0.32	-0.06	32.276	32.276	413
2	4/11/2009 19:36	0	2.4556	73.482	0	-0.26	0.12	32.141	32.546	412
4	4/11/2009 19:36	0	2.3719	72.008	0	0.22	-0.02	32.816	32.816	418
3	4/11/2009 19:36	0	2.4177	70.538	0	-0.06	0.2	32.681	32.546	414
1	4/11/2009 19:37	0	2.4177	72.303	0	-0.3	-0.06	32.411	32.276	407
2	4/11/2009 19:37	0	2.4556	73.482	0	-0.26	0.12	32.141	32.546	414
4	4/11/2009 19:37	0	2.3719	72.156	0	0.22	-0.02	32.816	32.681	429
3	4/11/2009 19:37	0	2.4177	70.685	0	-0.06	0.22	32.681	32.681	414
1	4/11/2009 19:37	0	2.4177	72.303	0	-0.3	-0.06	32.411	32.276	412
2	4/11/2009 19:37	0	2.4556	73.482	0	-0.26	0.1	32.141	32.546	416
4	4/11/2009 19:37	0	2.3719	72.156	0	0.22	-0.02	32.816	32.681	416

Table B.5 Wireless Sensor Network - Raw Data. (Sample Set)