GREEN COMPUTING - DESKTOP COMPUTER POWER MANAGEMENT AT THE CITY OF TULSA

By

CHARLES LETCHER

Bachelor of Music in Piano Performance

University of Tulsa

Tulsa, Oklahoma

1991

Submitted to the Faculty of the
Graduate College of
Oklahoma State University
in partial fulfillment of
the requirements for
the Degree of
MASTER OF SCIENCE

GREEN COMPUTING - DESKTOP COMPUTER POWER MANAGEMENT AT THE CITY OF TULSA

Thesis Approved:

N. Park Thesis Advisor

K.M. George

Subash Kak

Name: CHARLES LETCHER

Date of Degree: MAY 2013

Title of Study: GREEN COMPUTING - DESKTOP COMPUTER POWER MANAGEMENT AT THE CITY OF TULSA

Abstract: One type of Green Computing focuses on reducing power consumption of computers. Specialized software like 1E/Nightwatchman aids in reducing the power consumption of desktop computers by placing them in a low power state when not in use. This thesis describes the implementation of 1E/Nightwatchman power management software on two thousand desktop computers at the City of Tulsa. It shows the method used to predict power savings of $100,000.00 per year and compares the prediction to the actual savings after one year of operation.

TABLE OF CONTENTS

Chapter Page

1 INTRODUCTION **1**

1.1 What is Green Computing? 1

1.1.1 Green Computing Origins 1

1.1.2 Sustainability and the United Nations 2

1.1.3 Significance of Green Computing 2

1.1.4 Definitions of Green Computing 3

1.1.5 Benefits of Green Computing 6

1.1.6 Carbon Footprint 7

1.1.7 Is there any Computer Science in Green Computing? 8

1.2 Green Computing Efforts 9

1.2.1 Green Computing's Goal of Reducing Power Consumption . . 9

1.2.2 Data Centers 10

1.2.3 Energy Aware Scheduling 11

1.2.4 Server Virtualization 13

1.2.5 Computer Replacement to Save Energy 15

1.2.6 Router Power Efficiency 16

1.2.7 EOC - Container Based Data Center 16

1.2.8 Software to Manage Power in Buildings 17

1.2.9 Telecommuting 17

2 POWER MEASUREMENT AND SAVINGS **19**

2.1 Power Measurement 19

2.1.1 CMOS Single Transistor Energy Cost 19
2.1.2 Static Power and Leakage 20
2.1.3 Dynamic Power 21
2.2 Power Reduction 23
2.2.1 Dynamic Voltage and Frequency Scaling 23
2.2.2 Cooling, Waste Heat and PUE 23
2.2.3 Power/Performance Metrics 25
2.2.4 Data Centers Power Measurement Metrics 27
2.2.5 High Performance Computing 28
2.2.6 High Performance Computing Power Measurement 31
2.2.7 Optimal Power to Performance Ratio 32
2.3 Desktop Power Management Concepts 33
2.3.1 Power Savings Potential 34

3 CASE STUDY - PRELIMINARIES 37

3.1 Scope of Case Study 37
3.2 1E/Nightwatchman Software Features 37
3.3 1E/Nightwatchman Software Assumptions 38
3.4 Mayoral Policy on Energy Efficiency 39
3.5 1E/Nightwatchman Software Purchasing Decision 41
3.6 Resistance to Implementation 41
3.7 Power Savings Estimate 44

4 CASE STUDY - Implementation 48

4.1 Implementation Planning 48
4.1.1 Power Management Naive Solution for Desktop Computers . . 48
4.1.2 Management Directive - Manager Approval 50
4.1.3 Power Down Options and Behavior 51

4.1.4 Deployment Plan - Green Computing Table 52
4.1.5 Policy Exception Handling 55
4.1.6 Power Management Settings 58
4.2 Implementation 60
4.2.1 Pilot - Phase I 60
4.2.2 Pilot - Phase II 63
4.2.3 Staggered Software Deployment 66
4.3 Wakeup Issues - Last Man Standing Algorithm 66
4.3.1 Wakeup Test Results 71

5 CASE STUDY - RESULTS 74

5.1 One Year Later 74
5.2 Wake Up Results 74
5.3 Actual Costs 75
5.4 Time In State - 2011 75
5.5 Time In State - Day of the Week 77
5.6 Impact of not power managing 78
5.7 Actual Savings 79

6 CONCLUSION 82

BIBLIOGRAPHY 85

LIST OF FIGURES

Figure Page

2.1 Single Transistor Energy Equation [1] *(t=1)* 19

2.2 Dynamic Energy Equation [2] 22

2.3 Dynamic Energy Equation 2 [2] 22

2.4 PUE efficiency interpretation for data centers. [3] 25

3.1 City of Tulsa Desktop Energy Policy [4] 39

3.2 5% Time in Off State from Baseline 40

3.3 Cost to Purchase 1E Nightwatchman Software 41

3.4 Cost to Purchase 1E Web Wakeup Software 42

3.5 Yearly Cost per PC at the OTC for City of Tulsa at $.05/KWh . . . 45

3.6 Yearly Cost per PC outside the OTC for the City of Tulsa at $.10/KWh 45

3.7 Lighting Cost at OTC, December 2008 46

3.8 Electricing Cost at OTC from PSO, January 2010 47

3.9 Electricing Cost at OTC paid to Williams, January 2010 47

4.1 Deployment Plan Diagram 50

4.2 Power Consumption Values Supplied with 1E software 52

4.3 Options for Delaying Shutdown 53

4.4 15 Minute Poweroff Countdown Message 53

4.5 Force Poweroff Message 54

4.6 Green Computing Table 54

4.7 Exception Form (1) 56

4.8 Exception Form (2) 57

4.9 Exception Form (3) 57
4.10 Exception Form (4) 58
4.11 Nightwatchman Stealth Mode 59
4.12 Location Groups 59
4.13 Pilot Phase I Installation List 60
4.14 Altiris Deployment Console 61
4.15 Altiris 1E/Nightwatchman Installation Job 62
4.16 Pilot Phase II Installation Approval List 64
4.17 NOTICE: Deployment of NightWatchman to IT Computers 65
4.18 1E Web WakeUp Architecture 67
4.19 1E WakeUp Multi-Agent Installation 68
4.20 Nightwatchman server sends wakeup messages to target computer. These broadcast messages are not relayed by the router. 69
4.21 The primary agent sends wakeup packets to the powered off computer. 69
4.22 The computer that is waking up sends this JUSTWOKE packet. . . . 70
4.23 1E/Nightwatchman Web Wakeup 72

5.1 Sample Wakeup Results 75
5.2 Actual Costs by Month - 2011 76
5.3 COT Nightwatchman Power States 76
5.4 Yearly Time in State - 2011 77
5.5 Weekly Time in State 78
5.6 Actual Costs of Unassigned Desktops - 2011 79
5.7 Total kWh Saved 2011 80
5.8 Total Savings Report 2011 80
5.9 Lessons Learned 81

CHAPTER 1

INTRODUCTION

1.1 What is Green Computing?

1.1.1 Green Computing Origins

"Although the term 'green computing' and its alternative 'green IT' have recently become widely popular and taken on increased importance, their conceptual origin is almost two decades old. In 1991 the Environmental Protection Agency (EPA) introduced the Green Lights program to promote energy-efficient lighting." [5] "The Green Lights program encourages U.S. corporations to install energy-efficient lighting technologies." [6] "Green Lights is a non-regulatory program. The motivation for this program was energy-efficiency whenever profitable while maintaining or improving lighting quality." [6] Green Lights was originally a government regulatory effort to promote energy efficiency without sacrficing performance.

Another source states, "The 'green computing' idea started in 1992 when the US Environmental Protection Agency (EPA) launched Energy Star," a voluntary labeling program.[7, p. 74] The EPA's Energy Star program is still active.[8] Energy Star labels promote consumer awareness of the power efficiency of products. This regulatory program makes energy efficiency information publically available to allow informed evaluation of builders and buildings as well as products.

The earliest use of the keyword "Green Computing" in an IEEE journal is from 2005 when the term was used to describe differences between green and red computing.[9, p. 3] In this case, green computing was meant as safe, secure processing, while red computing was meant as less safe web browsing and downloading. This security term does not apply to the power efficiency meaning of "Green Computing" in common usage today. This thesis reviews other uses of the term Green Computing in the literature through the introduction chapter.

1.1.2 Sustainability and the United Nations

An important early milestone in the definition of terms for Green Computing comes from the United Nations. United Nation's resolution A/RES/42/187 summarizes the importance of sustainablity as "Believing that sustainable development, which implies meeting the needs of the present without compromising the ability of future generations to meet their own needs, should become a central guiding principle of the United Nations, Governments and private institutions, organizations and enterprises,..." [10]

"The term (sustainability) is often used in reference to the potential longevity of the environmental and human ecological systems, such as climate, agriculture, forestry and human communities in general." [11] "Computing plays a critical role in our society, thus it has a critical role in our society, thus it has a special responsibility for sustainability and green movement." [11]

Social responsibility includes proper stewardship of the environment. Energy efficiency, energy aware design, minimizing the energy used in manufacturing, and designing for recycling are all aspects of Green Computing that contribute to fulfilling the computer scientist's role in creating sustainable systems.

1.1.3 Significance of Green Computing

It is important to understand the amount of energy used by computing compared to all energy used. "Information and communications technology (ICT) infrastructure accounts for roughly 3 percent of global electricity usage and the same percentage of greenhouse gasses (GHGs)". [7, p. 74]

"Because green IT is a very small segment of the aggregate electrical energy equation, its ecological impact might not be great overall, but it's an integral part of the wider green movement, and many of its manifestations are significant." [7, p. 77]

1.1.4 Definitions of Green Computing

The terms Green Computing, Green IT and sustainable computing seem to be used without distinction in meaning between them. The terms relate computers to the environment, whether through power consumption, or the materials and processes involved in the construction/destruction of computers. The focus is on minimizing the negative impact of these factors. For example, when the term Green Computing is applied to computer power consumption, it means minimizing the power consumption which lessens the negative impact of consuming power. When the term Green Computing is applied to construction/destruction of computer equipment, it means designing, producing and handling equipment in the least harmful way. "Thus, green IT includes the dimensions of environmental sustainability, the economics of energy efficiency, and the total cost of ownership, which includes disposal and recycling." [12]

"Green computing also includes the goals of controlling and reducing the environmental footprint of computing by minimizing the use and discharge of hazardous materials, conserving water and other scarce resources, and reducing waste throughout the value chain [1]. Green computing encompasses IT product use over its lifecycle, and the recycling, reuse, and biodegradability of obsolete products. These goals seek to minimize the ecological footprint of IT products and services for companies and their customers." [13] The Green computing term applies not only to power efficiency, but also the choice and management of material used in the construction and disposal of computers.

"Green computing is the study and practice of using computing resources efficiently, by reducing the use of hazardous materials, maximizing energy efficiency during the products lifetime, and promoting recyclability or biodegradability of defunct products and factory waste." [14] This focus on efficiency in Green Computing suggests that the negative environmental impact will be lessened by Green Computing efforts.

Sustainable Computing refers to all aspects of the computer life-cycle. "Supporting business critical computing needs with least possible amount of power and sustainable computing." [15] In this case, sustainable computing refers to reducing the negative environmental impact. "Each stage of a computer's life, from its production, throughout its use, and into its disposal, presents environmental problems." [12]

Sustainable IT refers to service strategies. While "we define green computing as the practice of maximizing the efficient use of computing resources to minimize environmental impact." [5] "We define sustainable IT services in broader terms to include the impact of IT service strategies on the firms and customers societal bottom line to include economic, environmental, and social responsibility criteria for defining organizational success." [5] "The primary driver of sustainable IT is corporate social responsibility (CSR), especially as it applies to firms impact on the economy, environment, and society at large [56]." [13] As previously stated, the United Nations defined sustainability as meeting the needs of the present without compromising the ability of future generations to meet their own needs. [10]

While the terms Green IT and Green Computing have similar meaning, the term Green IT could be more specifically applied to an organization IT department's implementation of Green Computing practices. "Green IT, also known as Green Computing, refers to the study and practice of designing, manufacturing, and using computer hardware, software, and communication systems efficiently and effectively with no or minimal impact on the environment. Green IT is also about using IT to support, assist, and leverage other environmental initiatives and to help in creating green awareness." [16, p. 4]

In Europe, where there is more government regulation, compliance is an addtional motivation for Green IT. [17] This provides an opportunity to study company's responses to Green IT regulation.

A survey of 14 Danish companies led to the identification of 3 different Green IT strategies in response to regulation. They are Storefront (Creating an Image); Tuning (Improving Current Operations) and Redesign (Reinventing the Company). [17] The strategies correspond to tiers of increasing commitment to Green IT. The StoreFront level consists of "buying new servers that are more energy efficient, or an investment in IT for tracking emissions at production facilities". [17] and conclude that "behind the facade was very little substance." [17] "Tuning was the most common strategy employed. Its convenient and leads to improvements without requiring drastic changes." [17] Adopting redesign strategy may involve Green IT as part of a broader Corporate Social Responsibility (CSR) program. Redesign suggests a response to regulation that exceeds what the regulations require.

Profit driven efficiency, or first wave Green Computing, has be applied to many aspects of computing. "The first wave of green computing strategies focused on increasing data center efficiency. Infrastructure costs, power and workload management, thermal management, product design, virtualization, and cloud computing strategies remain the central focus of business value oriented strategy and tactics that have had the desired effect of lowering costs and environmental impact in terms of energy use." [13]

David Wang, the data center architect for Teradata, has specialized in thermal management solutions for the Miamisburg, OH-based data warehousing company since 1996. ""I've raised the issue [of green computing] because, for me, its both a business question and an ethical question,' [18] He describes "the green computing movement - a multifaceted, global effort to reduce energy consumption and promote sustainability." [18] This definition highlights the ethical responsibility for Green Computing.

The social responsibility motivation of Green Computing is a less powerful motivator than both profit through energy efficiency and regulatory compliance. "While energy efficiency is the principal driver of green computing, environmental reporting standards will be the primary driver of sustainable IT services. Voluntary and global guidelines for sustainability reporting have existed for some time and are continuing to evolve." [5]

"'Green' isn't all that new, of course. Energy management has always been an important topic for laptop manufacturers, who have been working on this for years as they struggle to manage weight versus battery life issues" [19] Battery effficiency in battery powered devices should not necessarily be considered Green Computing. The term seems to be applied recently to reviewing existing systems for power efficiency (and environmental impact) when those factors had not been seriously considered before. The idea is that when power efficiency and waste reduction are an integral part of the design process, it won't be Green Computing, it will be good design.

1.1.5 Benefits of Green Computing

Benefits of Green Computing "Green IT benefits the environment by improving energy efficiency, lowering greenhouse gas emissions, using less harmful materials, and encouraging reuse and recycling". [12]

"Green IT also strives to achieve economic viability and improved system performance and use, while abiding by our social and ethical responsibilities." [12]

"The US Environmental Protection Agency (EPA) estimated that providing computers with a sleep mode reduces their energy use by 60-70 percent" [12]

"The energy use of the nations servers and data centers in 2006 is estimated to be more than double the electricity that was consumed for this purpose in 2000." [20]

1.1.6 Carbon Footprint

Green Computing has also received attention because of the computer's contribution to carbon dioxide emmisions. "Over the last several years the link between energy use and carbon generation and the desire to lessen both has given rise to the green computing label." [5] "Reducing electric power consumption is a key to reducing carbon dioxide emissions and their impact on our environment and global warming." [12]

The 1E/Nightwatchman software calculates the carbon footprint based on a single carbon conversion factor. This carbon emissions factor is based on the efficiency of power production and transmission. Because most commercial installations do not generate CO2 directly, the CO2 is emitted when power is produced initially. The carbon conversion kg/KW factor could obtained from the power company, but would not reflect the transmission loss to your specific location. The problem is made more complex by multiple generation sources on the power grid. The carbon footprint of energy used is a factor of how efficient the power generation facility is.

You could try to precisely calculate a carbon footprint by multiplying the total energy used by some factor that indicates the amount of carbon produced per unit at a specific generation facility. But, there is transmission loss through the grid, so the amount of power generated is greater than the amount recieved. Because the grid has multiple generation sources, calcuating the exact transmission loss is complex.

Unless you can choose the efficiency of your power generation, the carbon footprint minimization problem reduces to a power consumption minimization problem. Most organizations can not choose the efficiency of their energy providers, and have limited ability to retask computing to more efficient environments. Presumably if purchasing greener power is an option, we could calculate the cost/benefit based on comparing the carbon emissions of the two power production methods. Unless you can choose the efficiency of your power generation, the carbon footprint minimization problem reduces to a power consumption minimization problem.

Environmentally responsible choices, like Google's placement of a datacenter next to a hydro-electric dam, can dramatically reduce the carbon footprint of computing. Most organizations can not choose the efficiency of their energy providers, and have limited ability to retask computing to more efficient environments. But, for dynamic cloud scheduling solutions, if you can do distributed computing at a datacenter with more efficient energy production, you can lower the carbon footprint. In this case, power cost may be lower, but have higher carbon factor output such as using power from coal.

In general, minimizing the energy used also minimizes the carbon footprint. "A key green objective in using computer systems and operating data centers is to reduce their energy consumption, thereby minimizing the greenhouse gas emissions." [12] Because the carbon factor is not precise, we apply a general rule that reducing power consumption reduces carbon emissions.

"Each PC in use generates about a ton of carbon dioxide per year." [12]

1.1.7 Is there any Computer Science in Green Computing?

"Although currently there's a lot more hype than science or engineering behind green computing, there's also great promise behind achieving the goals that are being articulated. [19]

Green Computing focuses on aspects of minimizing the negative environmental impact. One category of papers are high level overviews that try to put Green Computing in the context of global warming, peak oil and global economic strategy. These papers are therefore inter-disciplinary, and rely on expertise in economics, climate science and public policy.

Another type of Green Computing papers are those that present an idea to address an existing computer science problem. Many topics can be re-analyzed with a fresh perspective of "What are the power implications of this solution?" or "Can we re-write this algorithm/architecture so that it is power aware and power optimized?" In these cases, the computer science comes from the nature of the existing problem, not from making it power aware.

1.2 Green Computing Efforts

1.2.1 Green Computing's Goal of Reducing Power Consumption

The primary goal of Green Computing efforts seems to be aimed at reducing power consumption. This section gives an overview of types of efforts that have been put in place.

"The most obvious cost savings associated with using green computing methods is in energy costs. Simply put, a computer (or other device) that requires less power to operate, costs less to operate." [21]

"Power management is one of the major approaches in green computing to minimize the power consumption while the processing performance can be maintained at the target level." [22] Batch scheduler on the server farm level allows switching servers off. [22]

"The energy used by the nations servers and data centers is significant. It is estimated that this sector consumed about 61 billion kilowatt-hours (kWh) in 2006 (1.5 percent of total U.S.electricity consumption) for a total electricity cost of about $4.5 billion." [20] "Among the different types of data centers, more than one-third (38 percent) of electricity use is attributable to the nations largest (i.e., enterprise-class) and most rapidly growing data centers." [20]

1.2.2 Data Centers

"Energy management has now become a key issue for servers and data center operations, focusing on the reduction of all energy-related costs, including capital, operating expenses, and environmental impacts." [23] "Creating an energy-efficient data center is paramount to curbing runaway power consumption and accommodating greater data center capacity." [3]

"Data center energy and emissions costs are a major concern in green IT analysis because more than half of all IT-related electrical costs are generated there." [7] "In 2005, the power and cooling cost for servers worldwide was US$26 billion, and that cost is forecasted to top $40 billion by next year". [7]

"Information and communications technology (ICT) infrastructure accounts for roughly 3 percent of global electricity usage and the same percentage of greenhouse gasses (GHGs)1-3 but it seems to have a far greater role in the green debate than that." [7]

"In recent years, Microsoft and other companies have built data centers in central Washington to take advantage of the hydroelectric power produced by two dams in the region. The Microsoft facility, which consumes up to 27 megawatts of energy at any given time, is powered by hydroelectricity." [18]

"Another Microsoft data center, located in Dublin, Ireland, is expected to become operational in 2009 and, thanks to Irelands moderate climate, the 51,000-square-meter facility will be air cooled, making it 50% more energy efficient than other comparably sized data centers." [18]

Environmentally responsible choices, like Google's placement of a datacenter next to a hydro-electric dam, can dramatically reduce the carbon footprint of computing. Most organizations can not choose the efficiency of their energy providers, and have limited ability to retask computing to more efficient environments.

1.2.3 Energy Aware Scheduling

Energy Aware Scheduling uses the idea that you can balance the speed of a number of processors using digital voltage scaling DVS, against a certain number of tasks. Energy aware scheduling algorithms often give a tradeoff between performance and energy savings. One example is energy aware Dynamic Level Scheduling (DLS). "The scheme utilizes both time and energy to schedule tasks. The algorithm attains a higher energy saving by rewarding task processor pairs which are more energy efficient." [24] This energy-aware DLS (EDLS) immediately sacrifices a slight amount system response (0%-20%) in exchange for a large power savings (50%).

The test cases used for predicting power savings for EDLS included turning off a processor/core, or reducing the power of a processor/core. The power savings comes from the reduced energy consumption of the processors when some of them are shut down. The slow down comes from the schedulers inability to predict what the future processor requirements will be.

Another model for Energy Aware scheduling identified that memory contention would be a problem for memory bound processes. "We define a memory contention when several memory-bound tasks are concurrently executed on the same machine." [25] In this case, the model applies a penalty to the Estimate Time of Completion (ETC) because it assumes a higher cache miss rate. The power savings projected for this model are 17%. [25]

Task scheduling to meet performance parameters such as time and power is an NP-Complete problem.[26] "Therefore, many heuristics have been developed for real-time scheduling algorithms." [27] One approach is to compare new power aware scheduling heuristics to analyze power savings potential. For example, comparing Shortest Task First for Computer with Minimal Energy First with Speed Adjustment (STFCMEF-SA) algorithm, and Longest Task First forComputer with Minimal Energy First with Speed Adjustment (LTFCMEF-SA) algorithm, the researchs shows energy savings of a 1% to 3% while potentially sacrificing some performance.[27] "STFCMEF-SA algorithm and LTFCMEF-SA algorithm may not guarantee optimal speeds to minimize energy consumption." [27]

Another power aware task scheduling model compared shortest task first on the most power efficient computer with longest task first on the most power efficient computer. These results show that shortest task first provides a slight power savings of less than one percent compared to longest task first. [28] The percentage would increase in an environment with more variety in the power efficency of the servers available. The authors argue that because of the large amounts of power involved, this is still a significant savings.

The idea behind actively managing voltage by processor is the same idea as is used by Intel SpeedStep and AMD PowerNOW!. With this software, a desktop PC's processor will be reduced when the CPU is not under load. If Energy Aware Scheduling like EDLS were used on a computer that was also being actively power managed by Intel SpeedStep and AMD PowerNOW!, the additional savings would be very little, if any. Because you get a 50% power savings from Intel SpeedStep or AMD PowerNOW! already, that power savings from addition DVS would have marginal benefit. Essentially you would have 2 methods trying to modify the power consumption in response to load, so contention between the methods could lead to additional inefficiency.

The possible contention between different DVS algorithms (i.e. EDLS and PowerNOW!) raises the question of what level of the system should implement power management. What if the processor would automatically lower the voltage or shut off processors when the pipeline was empty? In that case, the length of the pipeline plus the amount of time to change voltage would dictate the amount of power saved. If that time were reduced to zero, the processor would automatically reduce its power consumption when not under load. This would bring the processor much closer to the goal of proportional processing in which energy used is proportional to work done.

"There are three key components for implementing dynamic voltage scaling (DVS) in a general-purpose microprocessor: an operating system that can intelligently determine the processor speed, a regulation loop that can generate the minimum voltage required for the desired speed, and a microprocessor that can operate over a wide voltage range." [29]

"Control of the processor speed must be under software control, as the hardware alone may not distinguish whether the currently executing instruction is part of a compute-intensive task or a nonspeed-critical task." [29]

Energy aware caching for disk arrays, like PRE-BUD, allow more disks to be kept in a low power state for a longer period of time. "Energy dissipation in parallel disks can be reduced by traditional power management strategies that turn idle disks into low-power modes or by directly shutting down idle disks. The traditional power management schemes can suffer great time and energy overheads that are induced by waking a disk up many times." [30] "Empirical results show that PRE-BUD is able to reduce energy dissipation in parallel disk systems by up to 50 percent when compared against a non-energy aware approach. Similarly, our strategy is capable of conserving up to 30 percent energy when compared to the dynamic power management technique." [30]

"It is known that disk systems can account for nearly 27% of the total energy consumption in a data center [Maximum Throughput, Inc., 2002]. Even worse, the push for disk I/O systems to have larger capacities and speedier response times have driven energy consumption rates upward." [30]

1.2.4 Server Virtualization

"Virtualization is at or near the top of every 'green computing' list. Virtualization may be considered a no-brainer, but there are a surprisingly limited number of applications for virtualization." [21]

”Virtualization software manages large clusters of separate, individual servers as if they were a single, virtual computer. In a data center with thousands of servers, there can be considerable idle time as equipment waits for sequentially assigned tasks using electricity and cooling power while doing no computing. Virtualization harnesses these resources by increasing server sharing and utilization, thereby reducing electrical cost.” [7, p. 76]

”When you have multiple underutilized physical servers consider virtualization.” [21] ”Physical servers take up physical space, use up energy resources, and require cooling. Consider that a physical server using, at maximum, 20% of that server's CPU, memory, storage and network bandwidth is still using 100% of that server's other components (fans, LEDs, optical devices, network interface card(s), and power supply). This Physical overhead can be reduced by consolidating these underperforming servers into a set of virtual servers hosted by one large physical server.” [21]

”Data centers and computational clusters use virtualization because of the many benefits it offers over the use of traditional stand-alone servers such as ease of management, enhanced security, and reduced costs.” [31]

”Optimal power efficiency is obtained when virtual machines are consolidated so that all the resources for a given host are fully utilized.” [31]

”A recent study of six corporate data centers found that most of the 1,000 servers were using just 10-35% of their processing power [12]. IBM estimated that the average capacity utilization of desktop computers was just 5% [13]. All of this wasted capacity still requires power, space, and cooling, as well as administration.” [32] ”Proper use of powerful modern server class hardware can eliminate the wasted capacity. First, by using products such as those developed by VMware [14], a physical server can be carved into a number of 'virtual machines' that share the disk space, memory, and processing capacity of the physical server while remaining entirely isolated from each other.” [32]

”A virtual machine can represent almost any operating system, and virtual machines representing different operating systems can reside on the same physical server hardware. Natural consolidation occurs when multiple physical servers are turned into virtual servers and aggregated on one physical server. This reduces the space, cooling, and power requirements of the data center.” [32]

1.2.5 Computer Replacement to Save Energy

Because of the high energy cost of manufacture, it is generally not energy efficient to replace less efficient computers with more efficient computers.

One manufacturer (Teradata) noticed that successive generations of their servers consumed significantly less power to perform the same amount of work, stating that ”today’s computers or servers can perform at a much higher level compared to those only a year or two old while consuming similar or even less energy during operation.” [33]

This allowed them to consider the question of whether it would be energy efficient to replace computers with newer models. The basic calculation is to compare the energy cost of manufacturing a computer to the energy saved by replacing it with a more energy efficient computer.

The key number in this analysis is the very high amount of energy required to manfacture a new computer. An estimate of 1700 kWh was used in the Teradata study. [33]

Consider the power savings in the case with a one to one replacement of a computer for a more power efficient model. For example, replacing a server that consumes 100 watts of power with one that consumers 90 watts would yield about a 10 watt savings. When compared with the high estimated manufacturing cost of 1700 kWh, it would take 170,000 hours of operation to justify replacement based on energy cost alone.

This is far beyond the useful life of computers. [33]

"Green computing must take the product life cycle into consideration, from production to operation to recycling." [33] To complete the energy cost replacement calculation, we would need to consider the energy required to recycle the computer. This additional factor would make computer replacement for energy savings even less justifiable.

1.2.6 Router Power Efficiency

In 2008, a rough estimate by Lawrence Berkeley Laboratory researcher Bruce Nordman was "that although data center network equipment accounted for 2 terawatt hours per year in the US, residential network equipment accounted for 7.3 terawatt hours and commercial office equipment accounted for 8.8 terawatt hours, not including cooling costs."[34] This estimate indicates that data center routers consumes far less than office and home routers. This observation led to the suggestion that "the resurgence of network-centric applications on which servers, proxies, and switches closer to the network core will be key to cutting overall consumption, and the continued deployment of peer-to-peer and grid computing networks in which desktops assume the role of servers will be crucial to cutting energy use and costs." [34]

Energy Efficient Ethernet is an IEEE standard "which can maintain rudimentary connectivity for devices in low-power states, and an IEEE standards track effort to speed up or slow down Ethernet link rates". [34] The IEEE ratified standards for energy efficient ethernet in September 2010. The estimate is that 50% power will be saved by enabling power saving features in routers. [35]

1.2.7 EOC - Container Based Data Center

One example is creation of a virtualized data center termed "Environmentally Opportunistic Computing (EOC) , which engages sustainable computing at the macro scale" [36] at Notre Dame. This solution is a containerized data center with no cooling unit. "From the outset the system was designed for operation utilizing only direct free cooling via outdoor or greenhouse air." [36] All output of air goes into the greenhouse. "During cold weather, the heat generated by the data center is vented into the greenhouse, saving both cooling costs for the data center and heating costs for the greenhouse. During hot weather, heat production and delivery is balanced by services migration." [36] The 100 servers in this "Sustainable Distributed Data Center (SDDC) container" [36] are dynmically scheduled, using a temperature aware algorithm, in the Notre Dame Condor cloud, which is called "Green Cloud." [36] This solution demonstrates that modern data centers can be dstirbuted in containers, possibly eliminating the need for active cooling.

1.2.8 Software to Manage Power in Buildings

Computer software has proven the ability to reduce power consumption in the management of office buildings. "the Energy Star power-management tool Portfolio Manager showed potential energy savings of 6 to 15 percent for more than 20 buildings." [7, p. 76] This example demonstrates that software programs can be effective in reducing power consumption.

1.2.9 Telecommuting

Telecommuting offers the potential to save huge amounts of energy. "An ITIF report found that, if only 14 percent of existing American office jobs were converted to work-from-home jobs, the savings would be dramatic: estimated at 136 billion vehicle travel miles annually in the US by 2020 and 171 billion miles by 2030.[10]" [7, p. 77] "The federal government has been relatively unsuccessful in achieving even modest telecommuting goals." [7, p. 77] Some unknown change in how managers handle off-site employees will be required for the government to meet even modest telecommuting goals. [7, p. 78]

CHAPTER 2

POWER MEASUREMENT AND SAVINGS

2.1 Power Measurement

2.1.1 CMOS Single Transistor Energy Cost

"Most microprocessors today are fabricated using the CMOS (Complementary Metal Oxide Semiconductor) technology. All digital circuits, regardless of their complexity, work by switching the state (on-to-off, off-to-on) of the millions/thousands/hundreds of transistors in the device. The reason that CMOS is so popular is that significant power is only drawn while the transistors in the CMOS device are switching between on and off states. In other words: no processing/calculating implies no (minimal) electrical current drawn." [1]

"For any electrical device, the power it dissipates is equal to P = V x I where V is the voltage and I is the current. Furthermore, the average amount of energy consumed equals E = P x t where P is the average power dissipated in the device and t is the time period of interest. If the speed of a CPU clock is S MHz, then a transistor could switch S million times in one second. Thus, the amount of energy used by a single transistor in one second is" [1]

$$E_{per\_\mathrm{sec}} = V \times I \times t \times S$$

Figure 2.1: Single Transistor Energy Equation [1] *(t=1)*

2.1.2 Static Power and Leakage

"The static power of a CMOS circuit is determined by the leakage current through each transistor." [29] Static Power is the amount of power consumed by a CMOS circuit that is not related to the switching of the transistors. It is wasted power that is dissipated as heat.

"With the continuous scaling of CMOS devices, leakage current is becoming a major contributor to the total power consumption."[29] "Reducing the supply voltage for transistors requires that the threshold voltage be decreased as well. Decreasing the threshold voltage increases the subthreshold leakage current." [29]

"The ratio of leakage power to total power is of particular interest in a modern high-performance microprocessor." [37] The leakage power ratio is the leakage power divided by the total power consumed. Leakage power for an early Itanium processor varies from .2 to .4 during various tests. [37] This means that 20% to 40% of power consumed by this CPU under load is wasted when compared to a design with no leakage power.

Models have been developed to help designers understand device leakage in CMOS circuits. Total leakage includes gate-to-channel, edge direct tunneling (EDT) and subthreshold leakage. [38] "The major leakage mechanisms in DG devices are the subthreshold leakage and the gate (gate-to-channel + EDT) leakage." [38] In this case, DG means double gate transistors.

Reducing the leakage current is a complex design problem. "The goal is to optimize the channel profile to minimize the off-state leakage while maximizing the linear and saturated drive currents." [29] "With reduction of threshold voltage (to achieve high performance), leakage power becomes a significant component of the total power consumption in both active and standby modes of operation" [29]

2.1.3 Dynamic Power

In order to understand the energy consumed by a computer, it is important to understand dynamic power measurement. A Watt is a point in time measurement. The sum of the fluctuating point in time measurements is the energy used measured in Joules.

"Today, the limits are based on fundamental physics, as exhibited by the dynamic power equation: $P = CV^2f$ where P is power, C is capacitance, V is voltage, and f is frequency." [39]

"However, the most popular measurement of machine energy for benchmarking is based on measuring the power consumption in Watts [35], [36]. Although power (Watt) is a good benchmark on a time interval, it does not reflect the overall energy of the processes. This paper takes execution time into account and continuously monitors the power of the entire processes. Thus, the energy is (Joule, Watt / Second) more important than power (Watt)." [2]

"The energy is obtained by integrating the power over execution time as shown on equations 4 and 5. We also had the result shown in Joule in addition to Watts." [2]

"Our measurement comes along with the current signal we intercept between the machine and the wall AC source. As we mentioned before, in reality, the system power (Watts) usage is not a constant. Thus, continuously monitoring the system power will bring us a more accurate picture. As a result of a continuous random variable, the dynamic energy E_{dyn} is " [2]

$$E_{dyn} = \tau \times P_{dyn} = \tau \times \frac{1}{\tau} \int_{t=0}^{\tau} p(t)dt = \int_{t=0}^{\tau} p(t)dt \qquad (4)$$

Figure 2.2: Dynamic Energy Equation [2]

where

$$P_{dyn} = \frac{1}{\tau} \int_{t=0}^{\tau} p(t)dt \qquad (5)$$

Figure 2.3: Dynamic Energy Equation 2 [2]

"Note that the energy E_{dyn} is in Joules when time t is measured in seconds and power P in Watts. For example, equipment consumes 10 Watts, and it takes one second to complete a task, then we need 10 Joules for each task. To obtain the total energy usage of a process, we simply record the signal through its entire execution. Actually, to make the record more precise, we have the process repeat itself for at least 120 seconds for measurement purpose. So, even if the process is so fast that all we get are aliasing signals, we can still get a rough average power." [2]

"Dynamic power dissipation consists of two components. One is the switching power due to charging and discharging of load capacitance. The other is short circuit power due to the nonzero rise and fall time of input waveforms.” [29]

2.2 Power Reduction

2.2.1 Dynamic Voltage and Frequency Scaling

”Dynamic Voltage and Frequency Scaling (DVFS) is used in order to obtain lower energy consumption on a dynamic CPU load; vendors refer to DVFS as P-states (P=performance). This allows for energy saving in processors tuned for laptops or marketed as EE (energy effcient) [13, 4].” [40]

”Per Core Power Gating (PCPG) (or dynamic core gating [DCG]) [8, 5, 11] is another technology, that is a hardware feature allowing the cores in a multi-core CPU to shut themselves off in absence of any load [15]. DVFS is explicitly controlled by the OS-scheduler, but on the other hand PCPG is controlled, in hardware, automatically. The PCPG states are also called C-states. The i7 processors has a hardware unit called Power Control Unit (PCU) that controls the C-states of a core depending on the load.” [40]

2.2.2 Cooling, Waste Heat and PUE

Each watt of power consumed is dissipated as heat. When the computer is in a cooled environment, the power cost calculations should include the cost of cooling. Likewise, in a heated environment, the power cost should be discounted because of the useful work the computer does in heating the environment.

”A Gartner study found that data centers, with their associated servers, air conditioning, fans, pumps, uninterruptible power supply (UPS), and so on, use 100 times the energy per square foot of an office building.6” [7] Data centers therefore require more concentrated cooling than office space.

”Cooling is the removal of waste heat generated by power conversion, racks, and any other equipment in the data center. Cooling can be improved by more efficiently removing the waste heat, or by generating less heat in the first place." [32]

"To a first order approximation, both cooling and provisioning costs are proportional to the average energy that servers consume, therefore energy efficiency improvements should benefit all energy-dependent TCO components." [23] TCO stands for total cost of owernship.

"Supercomputers consume a huge amount of electrical power and generate a tremendous amount of heat. Consequently, keeping a large-scale supercomputer reliably functioning requires continuous cooling in a large machine room."" [41]

"high voltage AC power can reduce waste heat by eliminating PDUs and lowering current and circuit requirements." [32]

"Almost all power consumed by IT equipment turns into waste heat that must be removed, or IT equipment can't function properly and reliably." [42] "Every watt of electricity consumed by IT equipment, an extra one and one half (1.5) watts of electricity is needed" [42] The figure is 1.5 instead of 1 because of ineffieciencies in the cooling process. "As a current rule of thumb, 1 megawatt (MW) of power consumed by a supercomputer today typically requires another 0.7 MW or cooling to offset the heat generated" [41]

Managing heat effectively is important because running computers at higher temperatures reduces reliability. "Hot spots in data centers are detrimental to hardware reliability as, for every 10 °C (18 °F) above 21 °C (70 °F), reliability declines by 50%."[42] "According to Arrhenius' equation as applied to computer hardware, every 10 °C increase in temperature results in a doubling of the system failure rate." [41] "The lesson learned is that by keeping the power draw lower, we can lower a supercomputer's system temperature, thus improving system reliability, which, in turn, contributes to better availability and productivity." [41]

The PUE is the ratio of the total power used in the data center divided by the power used by the computer equipment in the data center. The lower the PUE, the more efficient the datacenter. Using the PUE as measure of data center efficiency, we can speculate on what the values of an efficient data center would be. For a naive interpretation, a PUE of 1.0 would mean that all power consumed by the data center is used computers, and none is used for lighting or cooling. A PUE of 2.0 would mean that each watt of power used would need to be actively cooled excluding lighting. Figure 3 shows industry accepted standards for the interpretation of the PUE as a measure of efficiency.

PUE	Level of Efficiency
3.0	Very Inefficient
2.5	Inefficient
2.0	Average
1.5	Efficient
1.2	Very Efficient

Figure 2.4: PUE efficiency interpretation for data centers. [3]

Because the PUE is defined with respect to total energy consumed, it is possible to have a PUE of less than 1.0 if energy generation is included as part of the data-center design. One example is using the heat generated by the computers to power a heat driven absorption chiller. "The main challenge addressed is generating enough high temperature heat from the blade components inside the data center, and then capturing and transporting that high quality heat effectively and efficiently to the Li-Br absorption chiller." [3] When a solar array is added, a datacenter can achieve a PUE of 0.94 with current technology. [3]

2.2.3 Power/Performance Metrics

"There are two popular energy-efficient metrics currently in use: the energy-delay product (i.e., ED^2) and performance per watt (e.g., floating-point operations per second per watt or FLOPS/watt)." [43]

"The ED^2 metric represents the energy consumed by an application (E) multiplied by its execution time (i.e., delay) of that application (D) to the power n, where n = 1, 2, ... The ED^2 captures the translation of energy into useful work. For example, a small ED^2 means that energy is more efficiently translated into performance, i.e., smaller execution delay." [43] "Using ED^2 with n $\geq$ 2 as an efficiency metric to compare two supercomputer designs has a biased effect towards a massively parallel HPC architecture. A large n value not only emphasizes the performance aspect of a HPC system, but it also exaggerates the performance gained from the massive parallelism.More specifically, the ED^2 metric with n $\geq$ 2 increases exponentially with respect to the number of processors in a supercomputer.[44]

"The Green500 List uses performance per watt (PPW) as its metric to rank the energy efficiency of supercomputers." [45] "The 'performance per watt' metric is defined as: PPW = Performance / Power. The Green500 List explicitly specifies that performance is defined as the achieved maximal performance by the Linpack benchmark on the entire system." "Power is defined as the average system power consumption during the execution of Linpack with a problem size that delivers" the maximum performance. [45] The units for this metric are GFLOPS (Giga FLoating-point OPerations per Second) per watt. [45]

"the performance-power ratio is a better metric for efficient supercomputing, at least relative to power consumption." [44] "The FLOPS^n/W metric is really a derivative of EDn" with the value of n equal to one. [44]

For a comprehensive power/performance metric, reliablity must be included, but it currently is not. We know that computers become less reliable when operating at higher temperature. Additionally, a computer that fails to operate over sustained periods of time wastes energy during those failure periods, which is not reflected in the metric. "Despite the 'elusiveness' of a new efficiency metric, we still advocated the need for one in supercomputing and discussed the issues that a new metric should consider: reliability, productivity, power efficiency (performance-power ratio), and sustained performance." [44]

2.2.4 Data Centers Power Measurement Metrics

"Data center energy and emissions costs are a major concern in green IT analysis because more than half of all IT-related electrical costs are generated there, from small installations to massive facilities with thousands of servers and tens of thousands of associated workstations." [p. 75][7]

"Total power consumption. In a recent study, this metric was the most popular with 68% of IT managers specifying its use. The cost of power and the kilowatts used are typically included in the baseline assessment [9]. This metric can be useful in tracking power usage by facility, function, application, and employee. Accountability for electricity usage by IT organizations has been highlighted since it is a cost that can easily be tracked and it is a large part of the IT budget. Making power cost a discrete line item in the IT budget invites action to become more efficient and generate cost savings." [5]

"Power usage effectiveness (PUE). PUE is equal to Total Facility Power/IT Equipment Power. IT equipment power is defined as the load associated with computers, storage, network equipment and peripherals [33, 44]. Total facility power is the total power measured at the utility meter. A PUE of 2.0 indicates that data center demand is twice as high as the power necessary to power the IT equipment. A PUE value of 1.0 would indicate 100% efficiency with all power consumed by IT equipment." [5]

"Data center infrastructure efficiency (DCiE). DCiE = 1/PUE. This ratio is equivalent to the PUE. In the above example IT equipment uses 50% of the power in the data center. The other 50% is of power demand is typically required for cooling. As IT equipment uses less energy per unit of performance, then less energy is needed for cooling and DCiE will move higher [33].” [5]

”Data center performance efficiency (DCPE). DCPE = Useful Work/Total Facility Power. This ratio is informed by PUE and DCiE. However, it is much more complex to define and measure ’useful work’ performance as a standard metric [44].” [5]

2.2.5 High Performance Computing

The performance to power design problem leads to two questions: how to measure performance, and how to measure power. For measuring power, the Green Destiny team used a Yokogawa digital power meter that was plugged into the same power strip as the system.[46] The author used the Watts Up Pro power meter for verifying the power usage parameters supplied in the 1E software. The Watts Up Pro meter was also recommended and used by Tabeli in his Villanova thesis.

Measuring performance is more complex. The Green Destiny team used MIPS and Mflops at a clock speed to describe overall system performance. They ran benchmark programs like Linpak for CPU performance, STREAM for memory performance, Bonnie for I/O performance. The result of measuring power performance is that you can generate power efficiency measurements when you do the ratio of performance to power consumption.

For each of these performance metrics, it is possible to show the power efficiency by dividing performance by the power consumption. This the simple idea behind the Green 500 supercomputer list. The Green 500 shows supercomputers that perform with the highest power efficiency and provides recognition for teams that optimize for this area of design. ”It is a common practice to use average power for reporting FLOPS/watt metric for the Green500 list.” [43] Even though average power is used for the Green 500 FLOPS/watt metric, it can be argued that maximum instantaneious power would be a better measure because it can indicate when to best measure power due to fluctuations. [43]

The Supercomputing in Small Spaces project was started in 2001 by members of the Los Alamos National Lab. They identified lower power building blocks that could be used to create a low power supercomputer. One result of the project was the creation of Green Destiny, a power efficient supercomputer. Green Destiny consumed only 3.2 kilowatts of power while running diskless, "but with performance slightly better than a 172-processor Cray T3E 900 (circa 11/2001 TOP500 List) when running Linpack." [46]

The Supercomputing in Small Spaces team recognized the idea "that the failure rate of a compute node doubles with every 10.C (18.F) increase in temperature above 75.F, and temperature is proportional to power density." [46] They designed their supercomputer to run reliably at high temperatures. "Green Destiny provided reliable compute cycles with no unscheduled downtime from April 2002 to April 2004, all while sitting in an 85.F dusty warehouse." [46] By using low power building blocks, they were able to provide a reliable super-computer without active cooling (refrigeration). Green Destiny ran at higher temperatures.

The Green Destiny team provided an example that if nodes in the computer were spaced out to allow adequate air flow, active cooling could be avoided. There will likely continue to be massive super-computers with huge actively cooled facilities. In the future, this same technology can be used to create reliable computers that do not need active cooling simply by spacing the nodes and providing adequate air flow.

"A major issue in contemporary data centers, hosting such high computation facilities, is the high energy consumption in their operations."[47]

"However, as supercomputers transition from petascale to exascale systems, well have to dramatically rethink the underlying way in which computation is performed." [39]

The K computer is a new super computer. Preliminary test showed 828 mega-flops/Watt while performing the LINPACK benchmark on 4 racks of nodes and

48tera-flops. [48]

”As a result, supercomputers must be designed to provide significantly better power performance. This contrasts sharply with previous decades, when reliance on commodity parts drove supercomputer evolution at the cost of power and device efficiency.” [39]

”Historically, CMOS scaled by decreasing the capacitance per device, while increasing the total number of devices with Moores law. Successive generations also experienced increases in clock frequency. Despite the exponential growth derived by increasing capacitance and clock frequency, decreasing supply voltages held the entire equation in balance.” [39]

”However, clock-rate increases have resulted in increases in per-processor power consumption, which have now reached the upper limits of practical acceptance, even with the mitigating contributions of reduced logic voltage levels. Clock rate shouldn’t increase over the coming decade as much as it has over the previous two, which will largely eliminate an important contributing factor to performance increase. Also, voltage levels will diminish only slightly because we’re now reaching the limits of reliable circuit switching. Thus, we can expect little future power advantage from this other source of past improvements.” [39]

”Indeed, the use of hardware-supported speculative computing was one factor in increasing power consumption. The current era in processor core size has largely eliminated processor complexity as the second important source of performance gain.” [39]

”Finally, active power management and fault tolerance mechanisms will be included within the new architectures. Where core logic is used only lightly, clock rates will be lowered to further reduce power consumption. Fault detection, isolation, and reconfiguration mechanisms will be incorporated for fault tolerance. The mean time between failures of future systems comprising hundreds of millions of cores will be too short to support conventional check-pointing and restart methods.” [39]

”Projecting todays systems into exascale performance, even adjusting for technology improvements, will result in supercomputers that exceed 100-MW power budgets.12” [39] [49]

2.2.6 High Performance Computing Power Measurement

”Power utilization is a primary issue for high performance computing (HPC). Understanding it through physical measurements provides the critical first step to developing effective control techniques. For example, knowing the power consumption of each job in an HPC system will enable us to design a power-aware job scheduler and enhance the systems power efficiency.” [50]

Power measurement for the Jaguar supercomputer at Oak Ridge National Laboratory (ORNL) ” remains an ad hoc process.” [50] It relies on a combination of hardware based methods for measuring wattage and software-based methods to estimate wattage. [50]

Hardware based methods for measuring power can be distinguished as a standalone meter or an integrated sensor. ”The use of meter requires a measurement platform be physically built first whereas the use of sensor can skip this step. As we will see, today we have more options to use integrated sensors for each domain. In other words, the measurement process is slowly being standardized.” [50]

”The use of an inline meter is a common power measurement method for a server. Watts up? PRO, for example, is a popular inline meter used by many because of its cheap price (around $150), reasonable accuracy (+/- 1.5%), and okay sampling rate (1 Hz) via USB interface. Energy benchmarking such as Green500 and SPECpower ssj2008 use this type of measurement method.” [50]

”To measure the power usage of server components such as processors, memory and peripherals, a common method used is to intercept the power rails coming from the nodes power supply unit. In practice, current is measured instead since the voltages on these rails remain constant.” [50]

”The new generations of CPU, such as AMD K10, use multiple power rails for a CPU. The cores, memory controller, and I/O links are all powered up with separate rails. Yet there is no easy way to measure the power usage of these CPU components.” [50] ”New high-end processors add on-chip power sensors and an integrated power control unit to allow each processor to perform real-time monitoring of each cores current, power, and voltage states.” [50]

”The U.S. Environmental Protection Agency has recently determined that source energy is the most equitable unit for energy efficiency evaluation. Source energy represents the total amount of raw fuel required to operate the site. In other words, it accounts for all production, transmission, and delivery losses. The evaluation of the power loss will require simultaneous power measurements.” [50]

2.2.7 Optimal Power to Performance Ratio

”The list of fastest supercomputers is maintained by the Top500 [2], based on the high-performance LINPACK (HPL) benchmark [1]. The list ranks the supercomputers based on performance using the metric of floatingpoint operations per second (FLOPS).” [51] ”Many analytical and statistical models have been developed to automate the process of tuning the parameter space of HPL for achieving maximum performance.” [51]

”The most popular metric for energy efficiency is the performance-per-watt metric (i.e., FLOPS/watt). It is used by the Green500 to rank the most energy-efficient supercomputers in the world.” [51]

Some HPC researchers took an existing 64 node HPC test cluster and attched a Watts Up! Pro energy meter to attempt to understand the relationship between the maximum performance settings for the cluster in FLOPS, compared to the most energy efficient settings in FLOPS/watt.[51] The model used multi-variable regression analysis select the optimal settings. We could reasonably expect that the most energy efficient settings to be the same as the maximum performance settings because of the static power problem. Instead, they found that the most energy efficient settings were slightly below the maximum performance settings (3.65 TFLOPS to 3.62 TFLOPS). Though the researchers did not identify the cause, it is plausible that high performance settings produce more heat and therefore the components are slightly hotter and therefore less efficient. To verify this, we would need to monitor the temperature for the nodes. Additionally, changing the model to include the amount energy required for active cooling would increase the gap between the most energy efficient settings and the performance optimized settings. Because of other costs like personnel and facility, it is not clear that the fastest supercomputers would be set to run a few percent slower to save a few percent in energy costs.

"In our experiments, we use a Watts UP? Pro ES power meter" [51]

2.3 Desktop Power Management Concepts

"PCs, which are left powered on during night or weekends irrespective of user activities, waste 1.5 TWh of electricity per year." [52]

Desktop computer usage can be divided into different modes of operation: Active, Hibernation, Standby and Shutdown. [15] In active mode, the computer is not power managed. In hibernation mode, memory is stored in persistant storage. In standby mode, the processor is stopped, but memory remains active. From shutdown mode, the computer must restart. [15] These modes of operation can be used by an algorithm to reduce power consumption. [15]

"The most attractive energy-savings modes tend to be those with the highest wake-up penalities, such as disk spin-up time, and thus their use complicates application deployment and greatly reduces their practicality." [23]

2.3.1 Power Savings Potential

The power savings potential from Green Computing efforts is really the amount of power wasted. Many computer components have a static power consumption even when the computer is doing no useful work. One design ideal for proportional computing would be the reduction of static power consumption to a minimum level.

"For example, the EPA recently estimated annual savings from using enterprise power-management software for desktops could be US $25 to $75 per computer - you do the math for your company (Computerworld, 'Green from the Roots,' 4 Aug. 2008)." [14]

The idea behind powering down a desktop when not in use is to eliminate the power that is used while the computer is doing no useful work. "Essentially, even an energy-efficient server still consumes about half its full power when doing virtually no work." [23] "We see that peak energy efficiency occurs at peak utilization and quickly drops as utilization decreases. Notably, energy efficiency in the 20 to 30 percent utilization range - the point at which servers spend most of their time - has dropped to less than half the energy effeciency at peak performance." [23]

One author proposes that hardware should be designed so that energy use be proportional to use. "energy-proportional hardware could obviate the need for power management software, or at least simplify it substantially, reducing power management to managing utilization." [23] If computer desktop's power consumption were proportional to the work done, it would be unnecessary to place the computer in a low power state because no energy would be saved. This would render specialized power management software, like Nightwatchman, obsolete.

The power wasted by current server hardware power overhead is about 50%. "Energy-proportional computers would enable large additional energy savings, potentially doubling the efficiency of a typical server." [23]

A study of power savings potential on database servers observed that implementing low-power "resource management algorithms within the DBMS to exploit the power-saving modes of energy-aware hardware systems (e.g., CPU,storage devices, and memory). " [53] could save "up to 47% of energy". [53] Implementing DBMS power aware query optimizer would provide a 3% - 19% energy savings with minimum sacrifce in performance. This savings was calculated from a model that assigned power savings to execution plans. The model accepted more power effiicient plans while sacrificing performance. [53] The model "also systematically studied the energy/performance patterns of our workloads, and found that queries with power efficient plans that are missed by traditional query optimizers are very common." [53]

Although low-power supercomputing has resulted in impressive efficiency and competitive performance, many system researchers argue that it still has two disadvantages. First, most low-power supercomputing solutions require architectural modifications, for example ... Second, the tremendous growth in supercomputing performance has largely been spurred by commodity parts designed for PCs and servers." [41]

"We are building a system with multiple power and performance modes and then dynamically adapting the system's power and performance mode to match the current workload - the aim is to reduce power while maintaining performance. For instance, a processor-bounded workload requires the highest clock frequency (and voltage) to maintain performance, whereas a memory-bound workload can use a lower frequency and voltage for better energy effeciency while still maintaining performance." [41]

"Ideally , the appropriate time to scale down the processor's voltage and frequency is whenever the processor is waiting for data from memory accesses or I/O operations because there's no reason for it to set and spin its wheels at the maximum voltage and frequency while doing nothing." [41]

A similar approach can be taken with network loads. Power savings can be achieved on switches by adding dynamic transimission capacity control. The idea is to have multiple ports connecting two locations, powering off ports when not in use. A test environment was set up with traffic varying up to 8Gb/s balanced over 9 Giga bit ethernet ports. A model was used to find an optimal control interval to adjust the number of active ports. Acceptable performance was defined as no packet loss. Proportional power savings have been demonstrated in testing using this method. [54]

"The challenge for power-aware algorithms is thus to place processors in low-power mode only when doing so won't reduce performance." [41]

"A comprehensive plan for achieving green computing really does require an architectural approach." [19]

CHAPTER 3

CASE STUDY - PRELIMINARIES

3.1 Scope of Case Study

This case study documents a 1E/Nightwatchman software deployment on 2000 desktop computers at more than 10 offices for the City of Tulsa government. 1E/Nightwatchman software allows centralized management of computer power states with centralized reporting. The value of this case study is as a reference for future deployments.

Review of a case study can mitigate risk. Increasingly complex software is deployed as vendors supply more sophisticated products to solve real world problems. Often the people who evaluate, purchase, deploy and maintain these products have limited technical understanding of how these products work. Inexperience and lack of technical understanding can have serious negative impact on success. Review of a case study can help to inform decision makers about potential issues to mitigate risk.

3.2 1E/Nightwatchman Software Features

- Desktop computers can be power managed in groups with a shared policy.
- Groups of power managed desktop computers can be formed by physical location or organization role.
- Power management policy is applied by a physical location group, or by organization role group, or both.
- Power policies include a day and time for shut down.
- Computer users may be allowed to manually cancel a shutdown in progress.

- A shutdown policy has the option to detect if a computer has been idle, and skip the shutdown if not idle.

- There may be multiple shutdowns times in a power policy for a single day.

- Customizable scripts are run as part of shut down in order to save documents or handle errors. One example of how the City of Tulsa used the custom scripts was to cancel the shutdown if the computer had the Extra! mainframe terminal emulation program running.

- Centralized reporting stores time in state (on/off/standby/active/idle) for each desktop computer in one minute intervals in a database for reporting.

- Centralized reporting captures shutdown/wakeup errors for proactive troubleshooting.

- Comprehesive management reports are provided.

- The web wakeup feature allows individuals to power on their computers from a remote location.

3.3 1E/Nightwatchman Software Assumptions

- Central power management of desktop computers is more effective than relying on individuals.

- Desktop computers can be shutdown remotely with minimal negative impact.

- Assignment of a computer to a group (location or organization) is done manually by an administrator.

- For newly added computers, an administrator will be able to tell location and organization information by computer name.

3.4 Mayoral Policy on Energy Efficiency

Power management policy for the City of Tulsa's desktop computers was first established in 2007 by Mayor Kathy Taylor in "City of Tulsa Energy Conservation and Efficiency Plan". The policy read, "In an effort to save electricity, all City computers are to be turned off at night unless it is absolutely necessary not to do so." [4]

Electricity and Lighting

Electricity is among the most difficult commodities to control. From the large scale level with complicated energy arrangements with electric suppliers, to use, a compilation of thousands of individual consumers with various electrical needs, steps may be taken to turn the curve on electrical consumption by all departments.

TEAM meetings, held in the Mayor's Office, will continue to explore working with suppliers for municipal rate reduction. These meetings will also assess the increased feasibility of energy independence through: Hydroelectric, Solar, Wind, Biomass, Landfill Methane Recovery, and other technologies which could lead City government to energy independence. This could also be used to leverage existing energy rates. Experts from around the country will continue to be consulted, as they have recently been, on cleaner, more cost-efficient, and renewable solutions. It is not a question of if, but when.

Public Works

Light emitting diode (LED) street lights, according to the 2007 U.S. Conference of Mayors "Best Practices Guide," cost less than half what incandescent lights do, yet use less electricity and last five times longer.

City-Wide

In an effort to save electricity, all City computers are to be turned off at night unless it is absolutely necessary not to do so.

All lights are to be turned off when leaving a room.

There will be increased use of motion detection switches as opposed to manual on/off switches.

Turn coffee warmers off when the pot of coffee will no longer be served. Use a coffee mug, not a disposable cup.

Lights are to be turned off in areas where work is not done or where natural light is sufficient.

City of Tulsa employees shall not use plug in fans or heaters.

City employees are to avoid using lamps that are mostly aesthetic in purpose.

Energy saving light bulbs shall be used, where feasible, in replacement of inefficient and hot incandescent bulbs.

58

Figure 3.1: City of Tulsa Desktop Energy Policy [4]

Publishing a desktop computer power policy, and then relying upon the individual to take action to enforce that power policy, was ineffective. Initial results from the 1E/Nightwatchman pilot showed that computers were in the off state less than 5% of the time. [see Figure 3.2] We determined this by installing the 1E/Nightwatchman software in stealth mode. In stealth mode, the software monitored the time in state including whether the computers were powered off. This low voluntary compliance with policy would have been a significant factor in the decision to purchase centralized power management software, if the data had been available.

Daily Time In State

Shows the percentage of time machines have spent in each reported power state.

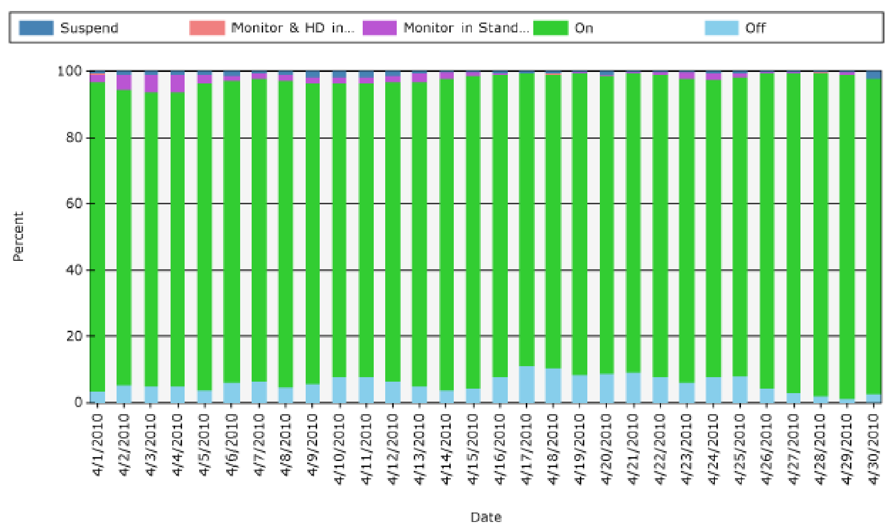

Date	# PCs Reporting State	% Time In State
4/1/2010	101	
	Off	3.41
	On	93.14

Generated: 4/30/2010 10:09:09 AM

Figure 3.2: 5% Time in Off State from Baseline

3.5 1E/Nightwatchman Software Purchasing Decision

In July 2009 Information Technology management at the City of Tulsa purchased licenses of 1E/Nightwatchman (NW) software for 2000 of the citys 2750 desktop computers for $24,790.00. NW allows centralized power management and data collection for desktop computers running the Windows OS. The city also purchased licenses for 500 computers of 1E/Wakeup (WU) software for $2,440.00.

Establish cost to purchase:

CITY OF TULSA

PURCHASE REQUISITION

Req. No.				DATE	06/01/09	PAGE	01 OF 01
Ship To Address:							
Department	Information Technology Dept			Ship To Code	500		
Address	175 E 2nd St Suite 660						
City, State	TULSA, OKLAHOMA			Postal Code	74103		
Requester	GEORGE SMITH			Requester ID	291		
Req.Line No.	0001	Item No.	204 - - 00 - 000 - 0	Capital Approval No.		Bid No.	
Description	1E NightWatchman including maintenance and support for one year						
Qty.	2,000	SKU	EA	Unit Price	$12.38	Total	$24,760.00
Required Date	ASAP						
Company (Fund)	1080	Account	53	Center (Element)	034030	Vendor No.	
Vendor Name	1E			Contact Name	Mark Kassel		
Address	5 Penn Plaza			Phone	972-378-5650	Fax	917-339-7333
City, State	New York, NY 10001			Zip Code	10001	P.O.#	
Req.Line No.		Item No.	- - 00 - 000 - 0	Capital Approval No.		Bid No.	
Description							
Qty.		SKU	EA	Unit Price		Total	$0.00
Required Date							
Company (Fund)		Account		Center (Element)		Vendor No.	
Vendor Name				Contact Name			
Address				Phone		Fax	
City, State				Zip Code		P.O.#	
Refer Questions Concerning Req. To	George Smith 918-596-9325 gsmith@cityoftulsa.org			Total Est.Cost	$24,760.00		

Figure 3.3: Cost to Purchase 1E Nightwatchman Software

3.6 Resistance to Implementation

The City of Tulsa's first effort in Desktop Computer Power Management was initiated in July 2009 with the purchase of software to manage power. Though the software was purchased in July 2009 , the software implementation did not begin in January 2010. The delay in implementation from June 2009 to January 2010 was due to many factors.

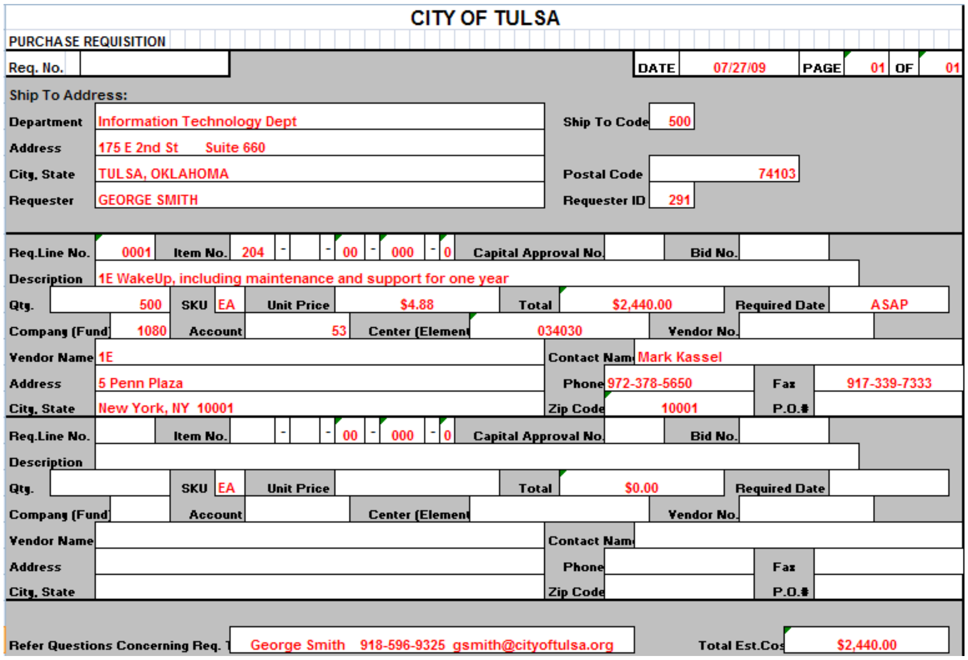

CITY OF TULSA

PURCHASE REQUISITION

Req. No. | DATE 07/27/09 | PAGE 01 OF 01

Ship To Address:

Department: Information Technology Dept | Ship To Code: 500
Address: 175 E 2nd St Suite 660
City, State: TULSA, OKLAHOMA | Postal Code: 74103
Requester: GEORGE SMITH | Requester ID: 291

Req.Line No.: 0001 | Item No.: 204 - - 00 - 000 - 0 | Capital Approval No. | Bid No.
Description: 1E WakeUp, including maintenance and support for one year
Qty.: 500 | SKU: EA | Unit Price: $4.88 | Total: $2,440.00 | Required Date: ASAP
Company (Fund): 1080 | Account: 53 | Center (Element: 034030 | Vendor No.
Vendor Name: 1E | Contact Nam: Mark Kassel
Address: 5 Penn Plaza | Phone: 972-378-5650 | Fax: 917-339-7333
City, State: New York, NY 10001 | Zip Code: 10001 | P.O.#

Req.Line No. | Item No. - - 00 - 000 - 0 | Capital Approval No. | Bid No.
Description
Qty. | SKU: EA | Unit Price | Total: $0.00 | Required Date
Company (Fund) | Account | Center (Element | Vendor No.
Vendor Name | Contact Nam
Address | Phone | Fax
City, State | Zip Code | P.O.#

Refer Questions Concerning Req. T: George Smith 918-596-9325 gsmith@cityoftulsa.org | Total Est.Cos: $2,440.00

Figure 3.4: Cost to Purchase 1E Web Wakeup Software

The reason given by IT managament for placing the 1E/Nightwatchman implementation on hold was related to the recent deployment of Altiris centralized desktop management and control suite. This deployment caused some widespread interruption of service, dramatic slow-down of desktop computers, and hostile response toward IT. IT management was unwilling to commit to another city-wide control software deployment because of the risk of similar problems. City IT workers were skeptical of whether the software would work properly in the city's environment, and recommended caution.

A significant factor in normal city worker's resistance to desktop power management was attributed to the "big brother" factor. Normal city workers are resistant to big brother software [55] that remotely monitors usage and controls computer usage. Employees were concerned that the software was monitoring their computer usage to measure their job performance and gather evidence against them for possible disciplinary action and termination. Nightwatchman does gather information about whether the computer is idle, and for how long, but does not have a method to determine whether the worker is being productive. It is impossible to use the 1E/Nightwatchman software to determine if the worker is being productive while they are active on the computer.

Other obstacles to implementation were:

- It would be difficult to test the installation properly to ensure that the software did not destabilize the target PC. The team decided to have multiple pilots and a staggered roll-out schedule to overcome this objection.

- It was well known that Wake On LAN broadcast packets were not relayed by routers on the city network, so there were concerns about whether the wake up feature would work at all city locations. As it turned out, the 1E/Nightwatchman software was successful in using active agents by subnet to broadcast these wakeup packets.

- Many city employees needed to VPN (virtual private network) and remote control their desktops after hours, which would be impossible if the personal computers (PCs) were off. This required activation of the 1E/Nightwatchman web wakeup features. Other work done to support remote wakeup included custom modification of 1E/Nightwatchman's web form for waking up computer, integration of the web form into the Juniper web VPN product, proper designation of network subnets in the Agent control section of the software, and testing. This allows worker to wakeup up their work computers from home using the VPN web page.

- There was no list of PCs which could not be turned off, so called production boxes that performed 24 hour productive work. Mission critical programs ran on desktop computers (instead of a server) that would have serious consequences if shutdown. The implementation team would have to create a list of these computers that could not be power managed. The list exceeded 400 computers, but also included individuals who strongly objected to power management, as well as computers that needed to stay on.

In January 2010, the author initiated the implementation of 1E/Nightwatchman software at the City of Tulsa because of Green Computing ideas from Dr. Park's OSU computer architecture class. This paper details the results of work done to implement 1E/Nightwatchman software for that class project. The author had to reassure IT management that the list concerns would be addressed before permission was given to proceed. Without Dr. Park's vigorous encouragement, this effort would not have started.

3.7 Power Savings Estimate

To justify the cost of purchasing 1E/Nightwatchman software, IT management guessed that $100,000.00 per year could be saved if 2000 Desktop computers were power managed by turning off the computers when they were not in use. The actual estimate, as well as the results after 1 year of operation, confirmed that initial guess. 1E software provides a saving calculator based on number of computers, time of power off, and cost of electricity per kWh which also matches the estimate and results.

Estimate of savings was $100,000.00 for the first year. OTC stands for One Technology Center, and it the main city hall. Outside OTC refers to all locations except city hall.

1000 computers at the OTC, 1000 computer outside the OTC: Total savings: $42,300.00 + $84,600.00 = $126,900.00 per year

Estimated savings for each Desktop PC at the OTC at $65.70 - $23.40 = $42.30 per year. Total savings estimate for 1000 PCs at the OTC is 1000 * $42.30 = $42,300.00 per year.

Estimated savings per Desktop PC outside the OTC at $131.40 - $46.80 = $84.60 year. Total savings estimate for 1000 PCs outside the OTC is 1000 * $84.60 = $84,600.00 per year.

Lighting estimates for the OTC used a $0.06021 cost per kilowatt hour. Based on review the rate paid to PSO through Williams, our estimate used the lower figure of

Table I. Yearly Cost per PC at the OTC for the City of Tulsa at $.05/KWh

Scenario	Desktop PC 100 Watts, 19" LCD 50 Watts	Laptop PC 35 Watts
Worst Case 24 hours per day usage, 365 days per year	150Watts/Hour * 8760Hours * 1Kilowatt / 1000 Watts * $0.05/ Kilowatt = **$65.70 total cost for 1 year of use**	35Watts/Hour * 8760Hours * 1Kilowatt / 1000 Watts * $0.05/ Kilowatt = **$15.33 total cost for 1 year of use**
9 hours per day usage, power on Monday night for patch updates	150Watts/Hour * 3120 Hours * 1Kilowatt / 1000 Watts * $0.05/ Kilowatt = **$23.40 total cost for 1 year of use**	35Watts/Hour * 3120 Hours * 1Kilowatt / 1000 Watts * $0.05/ Kilowatt = **$5.46 total cost for 1 year of use**
8 hours per day usage, turned off when not in use	150Watts/Hour * 2080Hours * 1Kilowatt / 1000 Watts * $0.05/ Kilowatt = **$15.60 total cost for 1 year of use**	35Watts/Hour * 2080Hours * 1Kilowatt / 1000 Watts * $0.05/ Kilowatt = **$3.64 total cost for 1 year of use**

Figure 3.5: Yearly Cost per PC at the OTC for City of Tulsa at $.05/KWh

Table II. 1 Yearly Cost per PC outside the OTC for the City of Tulsa at $.10/KWh

Scenario	Desktop PC 100 Watts, 19" LCD 50 Watts	Laptop PC 35 Watts
Worst Case 24 hours per day usage, 365 days per year	150Watts/Hour * 8760Hours * 1Kilowatt / 1000 Watts * $0.10/ Kilowatt = **$131.40 total cost for 1 year of use**	35Watts/Hour * 8760Hours * 1Kilowatt / 1000 Watts * $0.10/ Kilowatt = **$30.66 total cost for 1 year of use**
9 hours per day usage, power on Monday night for patch updates	150Watts/Hour * 3120 Hours * 1Kilowatt / 1000 Watts * $0.10/ Kilowatt = **$46.80 total cost for 1 year of use**	35Watts/Hour * 3120 Hours * 1Kilowatt / 1000 Watts * $0.10/ Kilowatt = **$10.92 total cost for 1 year of use**
8 hours per day usage, turned off when not in use	150Watts/Hour * 2080Hours * 1Kilowatt / 1000 Watts * $0.10/ Kilowatt = **$31.20 total cost for 1 year of use**	35Watts/Hour * 2080Hours * 1Kilowatt / 1000 Watts * $0.10/ Kilowatt = **$7.28 total cost for 1 year of use**

Figure 3.6: Yearly Cost per PC outside the OTC for the City of Tulsa at $.10/KWh

$.05 per KWh at the main city hall. Using the higher rate would have increased the estimated savings.

		# OF BULBS	WATTS/BULB	TTL WATTS	DIVIDER	KW	5:30-MIDNIGHT HRS OF OPERATION	TOTAL KWH	COST/KWH	CHG/DAY	DAYS/MO	TTL MO CHG	
Dec-08													
	8 LG 1000 W DIRECTED DOWN	8	1000	8000	1000	8	6.5	52	0.0621	3.2292	31	100.1052	
	28 SM 250 W - DIRECTED UP	28	250	7000	1000	7	6.5	45.5	0.0621	2.82555	31	87.59205	
	37 LIGHTS - 2ND FLOOR	37	32	1184	1000	1.184	6.5	7.696	0.0621	0.4779216	31	14.8155696	
	31 LIGHTS- 3RD FLOOR	31	32	992	1000	0.992	6.5	6.448	0.0621	0.4004208	31	12.4130448	
	197 LIGHTS/FL - 4TH -9TH FLOORS (6 FLS)	1182	32	37824	1000	37.824	6.5	245.856	0.0621	15.267658	31	473.2973856	
	207 LIGHTS/FL - 11TH - 15TH FLOORS (5 FLS)	1035	32	33120	1000	33.12	6.5	215.28	0.0621	13.368888	31	414.435528	
								572.78					
												$1,102.66	TTL DEC PER & SW
												69216.03	TTL DEC ELE COST
												1.59%	% OF TOTAL
												$1,102.66	TTL DEC PER & SW
												12	MONTHS
												$13,231.91	ANN PER & SW

Figure 3.7: Lighting Cost at OTC, December 2008

PSO charges a higher rate outside city hall, averaging more than $.10 per KWh. The actual rate paid varies by location from $.08 per kWh to $.14 per kWh with rates varying by meter. No suporting documention is provided in the thesis for this. The author's opinion is that $.10 per KWh is a conservative guess.

The reduced cooling cost was not included in the estimate of savings. Almost all power consumed by computers is dissipated as heat. Desktop computer power savings have a double impact on Green Computing during warm weather. Including the reduced cooling costs would have increased the savings estimate.

OTC UTILITY BILL FOR :	Jan-10
Actual Cooling Consumption For Month Of :	Dec-09
2008 Actual Cooling Cost	$ 28,100.72
2009 Billable Cooling Cost	$ 28,710.79

*Per U.S.A. agreement CHW Bill is for Previous Months Usage

Actual Heating Consumption For Month Of :	Dec-09
2008 Actual Heating Cost	$ 34,829.00
2009 Billable Heating Cost	$ 38,311.90

*Per U.S.A. agreement HHW Bill is for Previous Months Usage

Total Heating and Cooling	$ 67,022.69

OTC - Values - PSO 630 / AEP 246 Rate. Level 3 Fuel Rates

ELECTRICITY CONSUMPTION FOR:		Jan-10
Number of Days per billing	[Metasys Variable I.D.]	31
OTC Kilowatt hours	[wtc-pt]	987,972
OTC On Peak Demand	[wtc-pd7]	-
OTC On Peak-D Date	[wtc-px7]	N/A
OTC On Peak-D Time	[wtc-pz7]	N/A
OTC Off Peak Demand	[wtc-pd24]	2,017
OTC Off Peak-D Date	[wtc-px24]	01/20/10
OTC Off Peak-D Time	[wtc-pz24]	13:45
OTC On Peak Reactance	[wtc-pr7]	0.00
OTC Off Peak Reactance	[wtc-pr24]	(816.0)
Energy Charge (.0366*KWH)		$ 36,159.78
Fuel Cost Adjustment Mult.		$ (0.020036)
Fuel Cost Adjustment / KWH		$ (19,795.01)
Billing Demand, On-Peak 2-9, M-F, June, July, August and September		1,779.4
Charges: On-Peak Period Demand (@ $4.78/KW)		$ 9,875.61
OFF-PEAK		
Off-Peak Period Demand (KW) ($.87/KW)		2,017.0
Charges: Off-Peak Period Demand ($.87/KW)		$ 1,754.81
Peak Reactance (maximum of pr7 and pr24)		(816.0)
Base KW Allowance (30% of Monthly Max. Demand)		(605.1)
Applicable Peak Demand to KVAR Calc		2,017.0
Excess KVAR (.3 x Demand KW - KVAR)		210.9
KVAR Charge		$ 65.39
Reliability Rider (RVU) Factor		$ 0.000890
Reliability Rider (RVU) Charge		$ 879.30
Primary purchase Fee Rider (PFR)		$ 0.000606
Primary purchase Fee Rider $		$ 598.71
[redacted]		$
Demand Side Mgmt. Recovery Rider (DSM)		$ 0.000304
Demand Side Mgmt. Recovery Rider (DSM) $		$ 300.34
Regulatory Asset Recovery Rider (RAR)		$ 0.000311
Regulatory Asset Recovery Rider (RAR) $		$ 307.26
Purchased Power Capacity Rider (PPC)		$ 0.000579
Purchased Power Capacity Rider (PPC)		$ 572.04
Base Service Charge		$ 200.00
Sub-Total		$ 32,250.00
AEP/PSO Misc. Charges (+)		$ -
Merger Credit Apportioned per Consumption		$ -
AEP/PSO Misc. Credits (-)		$ -
OTC Apportioned Franchise Fee		$ 642.89
County/State/City Tax		$ 2,746.73
Total Of Taxes & Miscellaneous Charges		$ 3,389.62
TOTAL BILL		$ 36,025.13
Cost Per KWH		$ 0.03607

Billed Utilities Total	$ 103,047.82

Figure 3.8: Electricing Cost at OTC from PSO, January 2010

INVOICE

Invoice No: 567000008670
Invoice Date: 08-MAR-10
Terms: IMMEDIATE
Due Date: 08-MAR-10
Invoice Currency: US Dollar

REMITTANCE INSTRUCTIONS
Please remit payment by check to:
WILLIAMS HEADQUARTERS BUILDING CO.
Dept 650
TULSA OK 74182

Williams

CITY OF TULSA (OK)
TULSA PUBLIC FACILITIES AUTHOR
175 E 2ND ST
OTC - 8TH FLOOR
TULSA OK 74101

FOR BILLING ENQUIRES
Please Contact:
Kim Helm (918) 573-5108
M-F 8:00 am - 5:00 Central Time
P.O. Box 2400
Tulsa, OK 74101-9567

Item Description	Quantity	Uom Desc	Unit Selling Price	Extended Amount
METERED UTILITY CHARGE - Chilled Water 3450-5315401-040768	1 757	MONTH	28710.79	28,710.79
METERED UTILITY CHARGE - Heating Water	1	MONTH	38311.9	38,311.90
METERED UTILITY CHARGE - Electricity 3450-5315201-040768	1 756	MONTH	36025.13	36,025.13
		INVOICE TOTAL US Dollar		103,047.82

Figure 3.9: Electricing Cost at OTC paid to Williams, January 2010

CHAPTER 4

CASE STUDY - Implementation

4.1 Implementation Planning

4.1.1 Power Management Naive Solution for Desktop Computers

A naive solution for computer power management of a desktop computer is to power it off when not in use. "Turning off the system when not in use. This is the most basic energy conservation strategy for most systems." [12] This is the approach behind the practical section of this thesis, implementing 1E software. This is also the approach use by other modern research such as the Somniloquy solution. The naive solution is valid, if you provide a method for remotely powering on a desktop computer. Allowing the computer to be remotely powered on allows software updates to be applied, as well as the ability to remotely access it. With this remote wakeup ability, it is feasible to power off desktop computers when not in use for large power savings, even for complex organizations like a city government.

The 1E power management software allows a central administrator to set rules for powering off computers when not in use. The implementation of rules like "turn off all computers at 6:00PM except those in use," allow the 1E software dramatically save energy. These power policies can be applied to groups of computers by location or organzation.

A computer that is powered off may need to be remotely powered-on. Many modern network interface cards (NICs) support wake on LAN. Wake on LAN is a broadcast packet that allows a remote computer to send a signal to wake a powered off computer. In order to recieve the wake on LAN packet, the network interface card will stay powered while the rest of the computer is powered off. Our direct observation is that this takes about half a watt of power. Unfortunately, many routers are set to not relay wake on LAN beyond the originating subnet for security reasons.

A sophisticated solution that uses hardware to solve the problem of waking a computer is Somniloquy. The Somniloquy solution allows a computer to be powered off when not in use, but allows the ability to power on the computer by having a low power USB device simulate the computer. In fact, the Somniloquy solution allows certain PC functions like BitTorrent to active even when the PC is powered-off. [56] 1E software uses agent computers on the same subnet as the target computer. The agent will receive a message from the Nightwatchman server and will broadcast the wakeup message(s) originating from the same subnet.

The authors of the Somniloquy solution did an informal survey to better understand the motivation for not powering off a PC. For home machines that were left powered on, "29% did so for remote access, 45% for quick availability and 57% for applications running in the background, of which file sharing/downloading (40%) and IM/e-mail (37%) were most popular. In the office environment, 52% of respondents left their machines on for remote access, and 35% did so to support applications running in the background, of which e-mail and IM were most popular (47%)." [56]

Because some computers are relied upon to do valuable work when a person is not present, identifying computers that are doing useful work when a person is not there is critical. 1E software identifies a computer as idle if a person has not interacted with it for some period of time (i.e. 15 minutes). Desktop computers can be setup to run data processing jobs that should be run on a properly monitored and supported server computer. Turning off these computer can cause serious data processing problems for the computer owners. Identification of computers that could not be powered off was a key factor in success in the project.

The 1E/Nightwatchman solution implemented in the thesis allows the function of remote wakeup to address the remote access need. Our experience was that computer power-on took 3-10 minutes, and functioned properly more than 90% of the time. The 1E server did properly log computers that did not wake. The troubleshooting for this was an on-going minor issue. In general, the support group waited until a person complained about a computer not waking instead of pro-actively checking the report.

4.1.2 Management Directive - Manager Approval

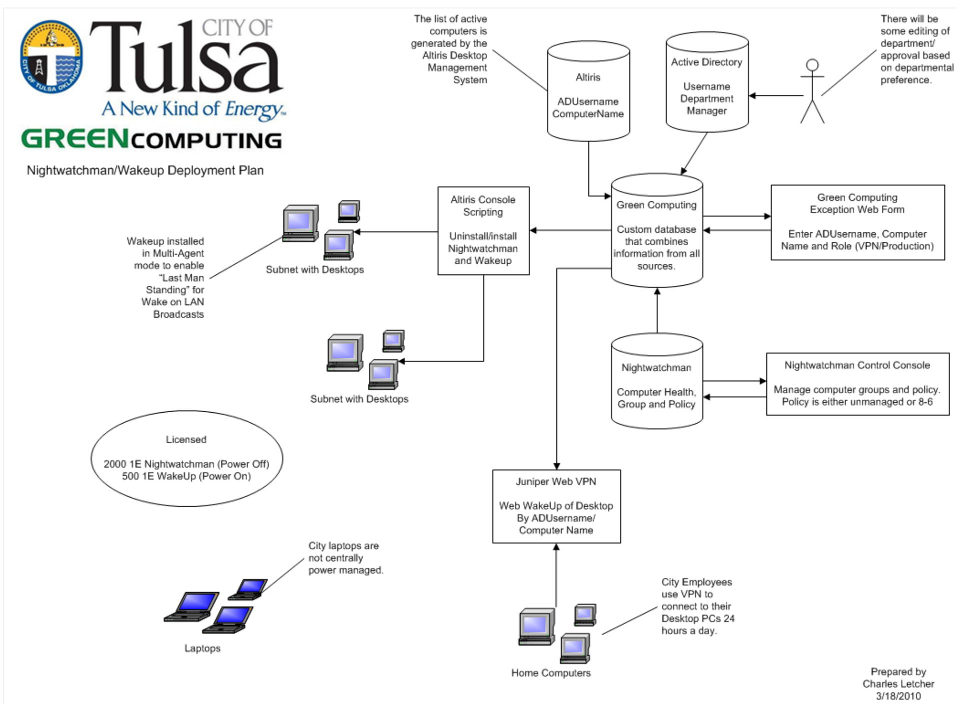

Figure 4.1: Deployment Plan Diagram

To facilitate rapid deployment, the city IT management establishing the directive that all desktop PCs would be turned off at 6:00PM, and left off until manually powered. They encouraged aggressive deployment of the power management software in order to make the savings goal for the year in spite of obstacles. They recognized that provisions would need to be made to accommodate VPN users and 24 hour desktops. IT management placed an additional restriction that managers must approve before the software was deployed in their area. This additional manager approval step was put in place to help ease user acceptance.

A problem with requiring manager approval before implementation was that the city did not maintain a complete list of PCs and owners. The city also did not maintain an accurate operational organizational chart, just a cost center chart. A city employee may be paid by a different department, which makes it difficult to programmatically identify the approving manager for an employee. Fortunately, the recent Altiris deployment provided a fairly complete list of AD (Active Directory) user names and corresponding computer name. Also, spreadsheets were used in the recent Exchange implementation that roughly showed the actual organizational unit for each AD user. From this information, we constructed a new Green Computing database table that listed each computer, the owner, its power management policy, the approving manager, and whether approval was given.

This new Green Computing table was web enabled with a grid showing all computers assigned to a manager, and a function added to allow web based approval. After the mangers approved, a WSH script was used to install the power management software. Because installations were automated, with lists of installations coming from this table, multiple workers could update the Green Computing table through a web form and thus control software installation.

4.1.3 Power Down Options and Behavior

1E/Nightwatchman identify the model of computer. The model of the computer is stored to the central database. That model is used to pull the power usage of the computer. Tests of specific computer using the Watts Up! meter verified the software's estimatated power usage. Figure 4.2 shows a sample of the power consumption estimates by model.

1E/Nightwatchman software allows the option of the computer's user to delay

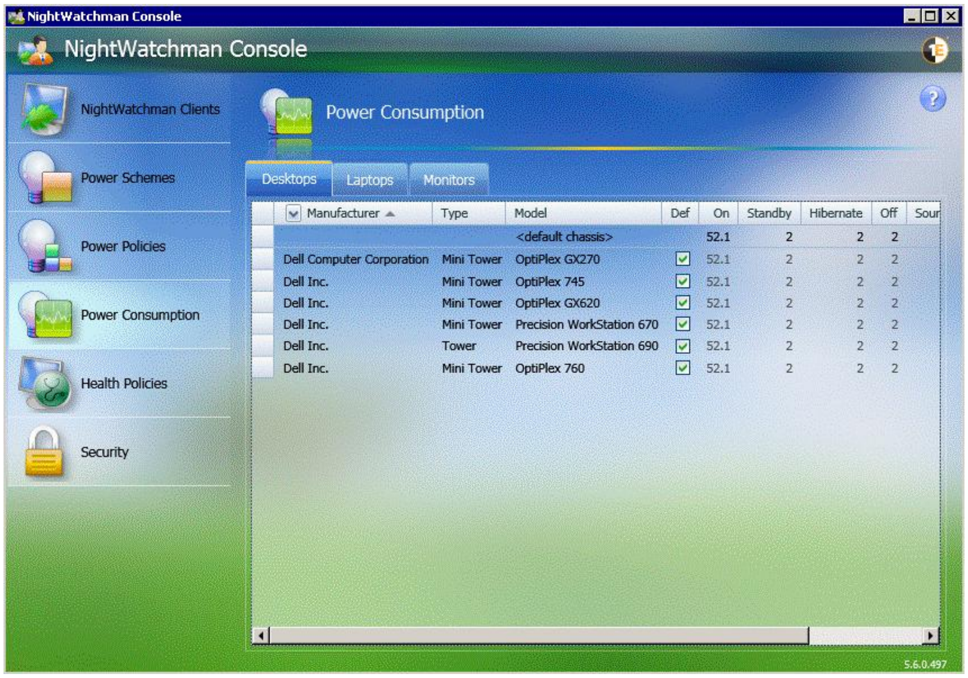

Figure 4.2: Power Consumption Values Supplied with 1E software

shutdown using a pop-up window. Whether to force the shutdown, or allow the user to cancel, is configured in the power policy. Figure 4.3 shows the pop-up screen that allows you to delay shutdown.

When a computer receives a shutdown command from the server, or reaches a scheduled shutdown time, a pop-up with a countdown timer is provided. Whether the user can cancel the shutdown is part of the power policy. The time for the countdown is also set in the power policy.

4.1.4 Deployment Plan - Green Computing Table

In order to track the status of installation a new Green Computing table was created to combine information from various sources. Inputs to this new table came from 4 sources: The list of computers, and usernames came from Altiris; the management structure came from Active Directory; the group and power policy came from the Nightwatchman database; finally user objections were captured in a custom web form.

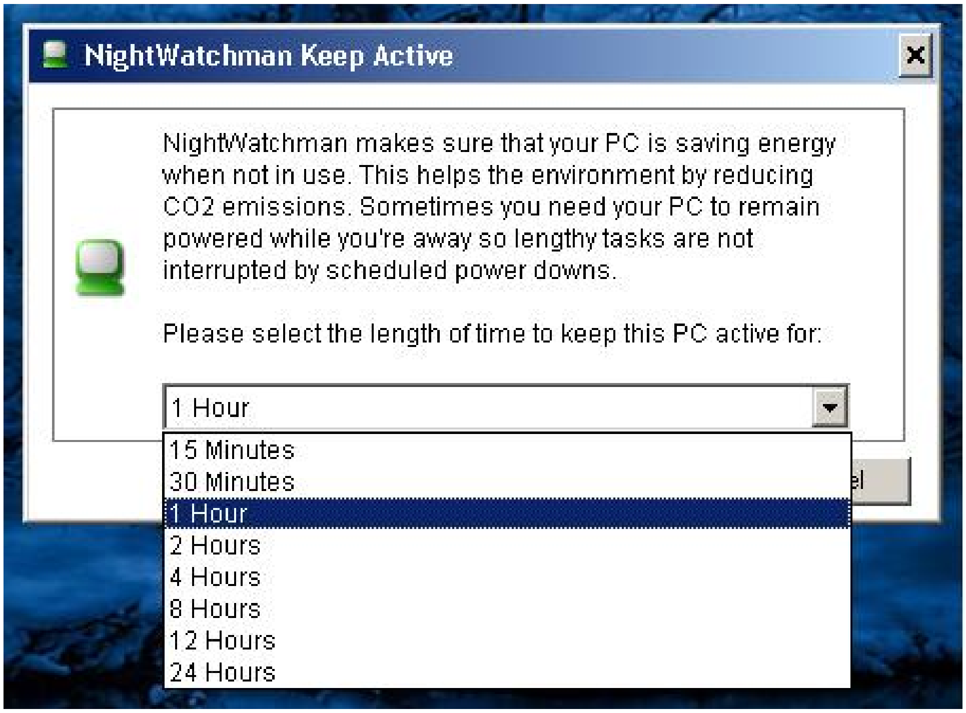

Figure 4.3: Options for Delaying Shutdown

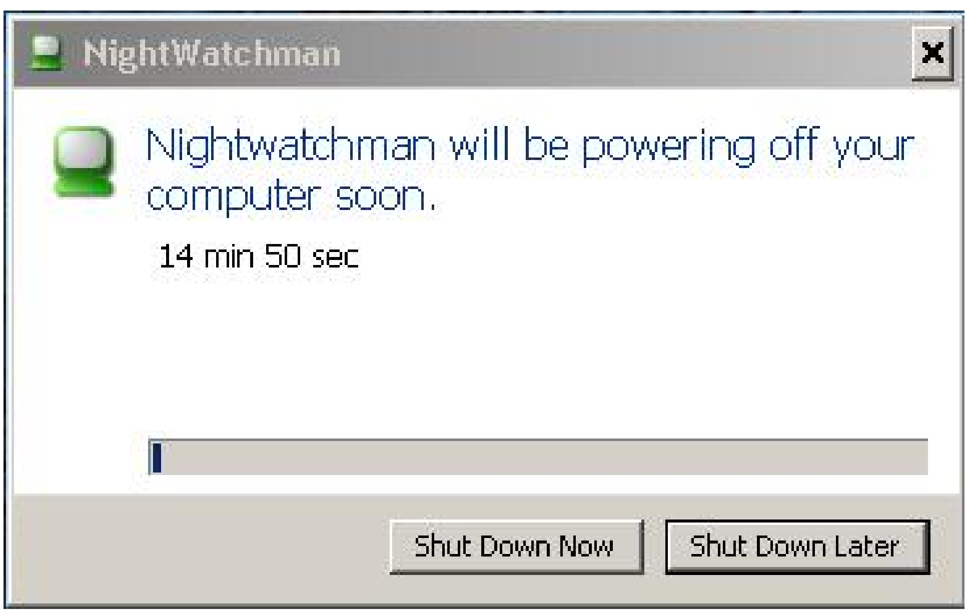

Figure 4.4: 15 Minute Poweroff Countdown Message

Figure 4.5: Force Poweroff Message

C:\Documents and Settings\cletcher\Desktop\test.accdb — Thursday, March 18, 2010
Table: dbo_GreenComputing — Page: 1

Properties

Attributes:	Password Saved, Linked ODBC	Connect:	ODBC;DSN=PWSqlSvr1;UID=texter;APP=2007 Microsoft Office system;DATABASE=Cityweb
DateCreated:	3/11/2010 1:09:59 PM	DefaultView:	2
DisplayViewsOnSharePointSit	1	FilterOnLoad:	False
GUID:	{guid {8905EA2A-E696-4CF8-A2A0-694DDD710F23}}	HideNewField:	False
LastUpdated:	3/17/2010 11:00:48 AM	NameMap:	Long binary data
OrderBy:	[dbo_GreenComputing].[ADUser], [dbo_GreenComputing].[ComputerName]	OrderByOn:	True
OrderByOnLoad:	True	Orientation:	Left-to-Right
RecordCount:	-1	SourceTableName:	dbo.GreenComputing
TotalsRow:	False	Updatable:	False

Columns

Name	Type	Size
pk_id	Long Integer	4
ADUser	Text	50
ComputerName	Text	50
Category	Text	50
Expanation	Memo	-
dtUpdated	Date/Time	8
AltirisComputerName	Text	50
ImplementationNote	Text	50
NightwatchmanInstalled	Text	50

Figure 4.6: Green Computing Table

The green computing table is used to generate 2 sources of output: 1) automated generation of install/un-install scripts for Altiris and 2) a list of power managed computers so that VPN users can wake-up their powered-off desktop.

The Green Computing table's unique information is the users Category (Prod, VPN, OK), and an explanation why they are in that category. Prod means a 24/7 computer that should never be power managed. VPN means an active VPN user who has logged in to the VPN system in the last 90 days. OK means OK to power manage.

The green computing table alsos tracks manager approval and installations. Additional information is stored in the source system. For example, if you want to know the name of the user, you use the ADUser to look up the name in active directory. Additional information about the computer or user is gained by joining to the Altiris database by Computername, or Active Directory by ADUsername. If you want to know the Power Policy, you look it up in Nightwatchman using the computer name.

The Green Computing table is used by the Green Computing Web Form to store User preferences. This custom table was created for the project.

4.1.5 Policy Exception Handling

We created a new web form to allow users to register their exceptions to power management. The form contained detailed explanation about the purpose of the project. The web form also contained instructions about how to find your computer name. Finally, the web form allowed the user to see the computers they had logged on to, and register the computer as an exception. We learned that inviting objections was the wrong strategy. Only authorized IT employees had access to this web form after strong negative responses from the organization.

The custom web form was initially deployed to achieve approval from managers to power manage computers. The development of the form took several iterations.

The initial version of the web form included an explanation of the intention of the initiative, an area to enter an exception for power management, and a grid showing previous entered exceptions. After the initial version was written and deployed, IT management reworded the form, which was reviewed by a technical writer. The final version of the web form included a web grid of all computers assigned to a manager, and the power management settings applied. This custom web form was created

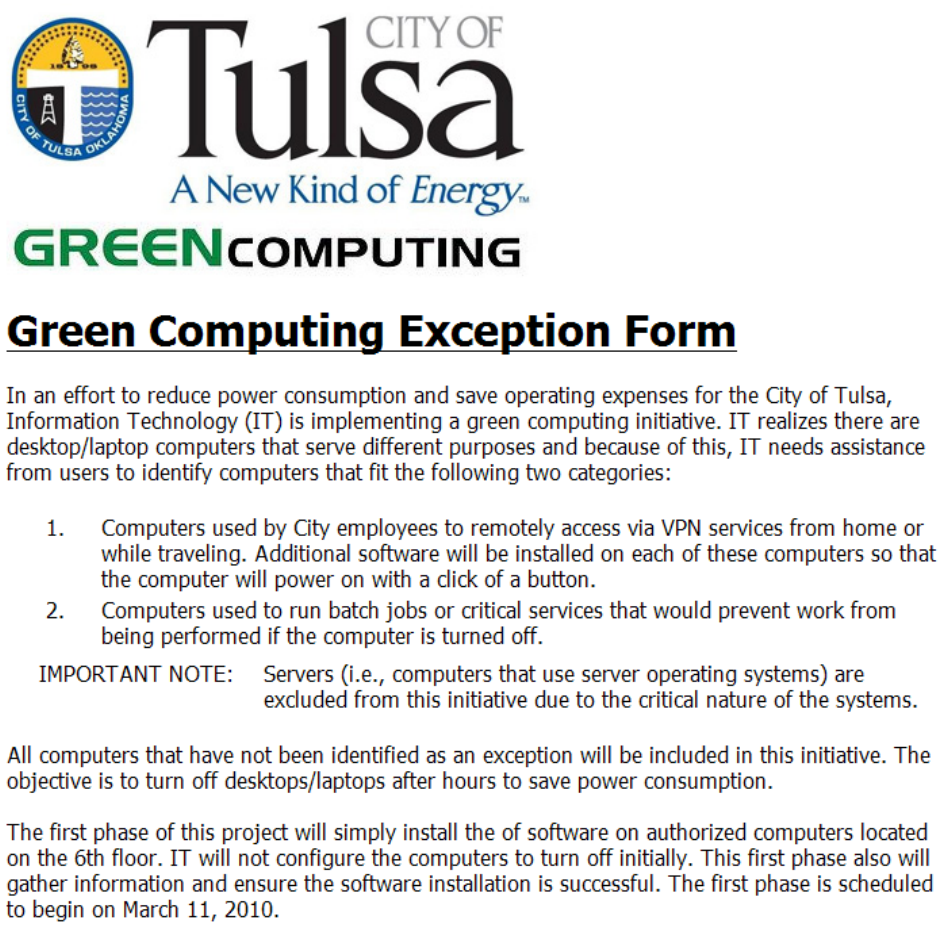

CITY OF Tulsa

A New Kind of Energy™

GREEN COMPUTING

Green Computing Exception Form

In an effort to reduce power consumption and save operating expenses for the City of Tulsa, Information Technology (IT) is implementing a green computing initiative. IT realizes there are desktop/laptop computers that serve different purposes and because of this, IT needs assistance from users to identify computers that fit the following two categories:

1. Computers used by City employees to remotely access via VPN services from home or while traveling. Additional software will be installed on each of these computers so that the computer will power on with a click of a button.
2. Computers used to run batch jobs or critical services that would prevent work from being performed if the computer is turned off.

IMPORTANT NOTE: Servers (i.e., computers that use server operating systems) are excluded from this initiative due to the critical nature of the systems.

All computers that have not been identified as an exception will be included in this initiative. The objective is to turn off desktops/laptops after hours to save power consumption.

The first phase of this project will simply install the of software on authorized computers located on the 6th floor. IT will not configure the computers to turn off initially. This first phase also will gather information and ensure the software installation is successful. The first phase is scheduled to begin on March 11, 2010.

Figure 4.7: Exception Form (1)

for this project to allow Users to enter exceptions to the power management policy. The web form uses IIS integrated security to capture the ADUsername.

The second phase of this project includes configuring the software to turn off authorized computers starting at 6:00 p.m. each evening, Monday through Friday.

Excluding a Computer from the Green Computing Initiative

To exclude a computer from this cost savings initiative based on the criteria mentioned above, use the following steps to determine the computer name and then enter it in the Computer Name field in the form below. Be sure to select a Category describing the computer and give an Explanation as to why the computer should be an exception.

Figure 4.8: Exception Form (2)

1. Press the *Start* button and the *Break* button simultaneously. *Figure 1* should appear on the screen as shown below.
2. Click on the *Computer Name* tab and locate the *Full Computer Name* as indicated in *Figure 2*.

Figure 1

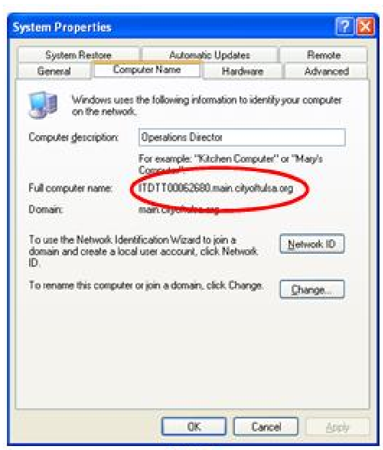

Figure 2

Figure 4.9: Exception Form (3)

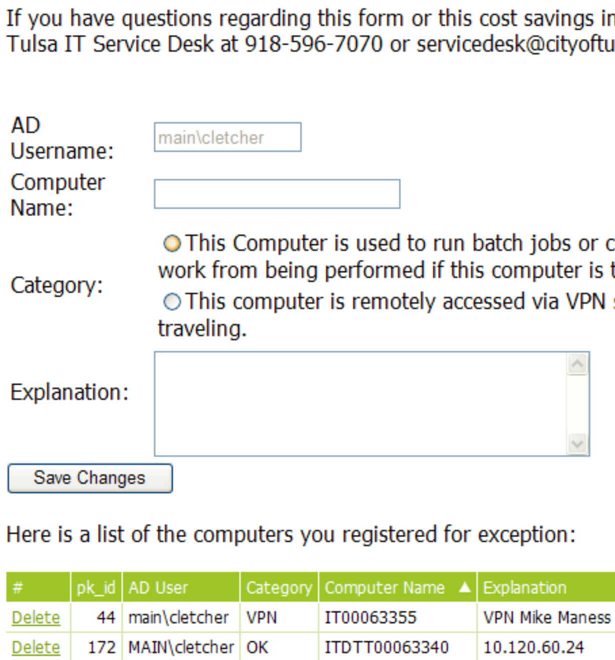

If you have questions regarding this form or this cost savings initiative, please contact City of Tulsa IT Service Desk at 918-596-7070 or servicedesk@cityoftulsa.org.

AD Username: main\cletcher

Computer Name:

Category:
- This Computer is used to run batch jobs or critical services that would prevent work from being performed if this computer is turned off.
- This computer is remotely accessed via VPN services from home or while traveling.

Explanation:

Save Changes

Here is a list of the computers you registered for exception:

#	pk_id	AD User	Category	Computer Name ▲	Explanation
Delete	44	main\cletcher	VPN	IT00063355	VPN Mike Maness
Delete	172	MAIN\cletcher	OK	ITDTT00063340	10.120.60.24

Figure 4.10: Exception Form (4)

4.1.6 Power Management Settings

The NW software allows rules based enforcement of power management settings. This is accomplished by applying NW Power Policies to NW Groups of computers. All new computers are added to an unassigned group. In the case of the City of Tulsa, the unassigned group had the stealth mode power policy assigned. In stealth mode no system tray icon is displayed, and no power off is scheduled, but reporting statistics are gathered. Once the computer is moved to its corresponding group, then a different policy may apply, because a different power policy may be assigned to that group.

At the City of Tulsa, the default unassigned group applies the Stealth Mode policy:

Individual computers are assigned a power policy by group. Computers are assigned to a group by drag-drop, or by updating the Nightwatchman database:

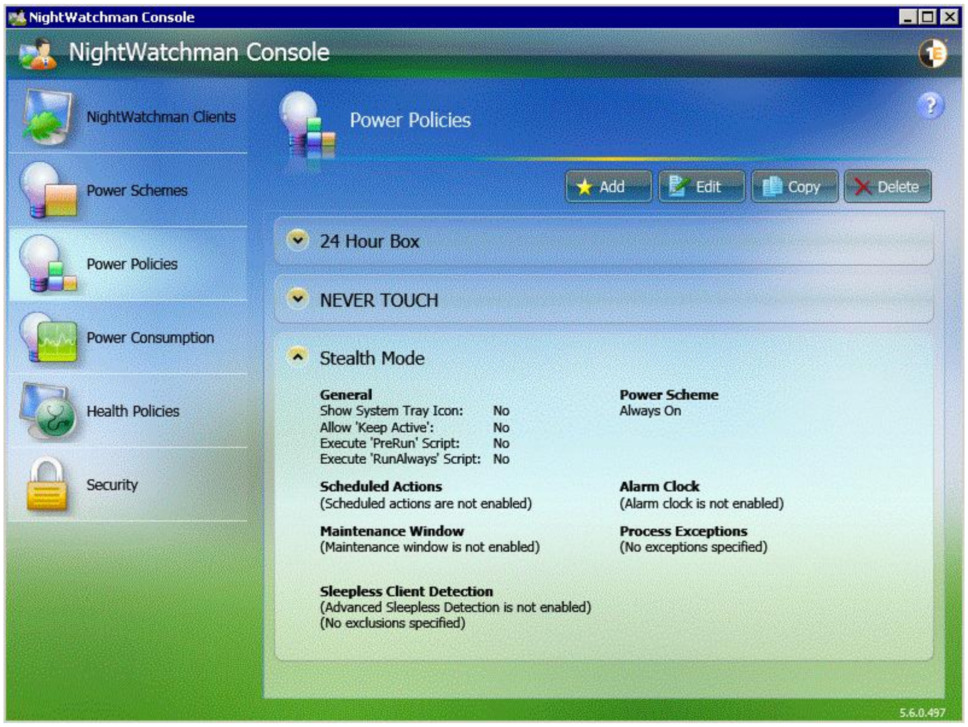

Figure 4.11: Nightwatchman Stealth Mode

Figure 4.12: Location Groups

4.2 Implementation

4.2.1 Pilot - Phase I

10 computers were identified by the manager of desktop support as the test computers for a first pilot. In the Nightwatchman console power policies are applied to groups of computers, not individual computers. By default, the computers show up in the Unassigned Building group, and must be moved to their correct organizational group. In this case, the computers were moved to the Pilot Phase 1 group.

Figure 4.13: Pilot Phase I Installation List

The power policy used for initial deployment (Pilot Phase 1) was called Stealth Mode. This mode has zero visible impact: no system tray icon, no program shows as installed, and there is no performance slow-down. However, reporting statistics are sent back to the Nightwatchman server by each individual computer.

A pilot project of installation on 10 test computers was initiated in February. Letcher installed the software using a new custom Altiris job that executes a MSI silent install using these command options:

INSTALLDIR="C:\\1E\\NightWatchman" SYSTEMTRAY=OFF REPORTING=ON REPORTINGSERVER=T1ITPAPP06 PIDKEY=COTNWM5-2TXU-248X-9051-XXXX

Remote installation and de-installation was accomplished for the first pilot using the Altiris deployment console. Installation is accomplished by dragging and dropping a job onto a computer.

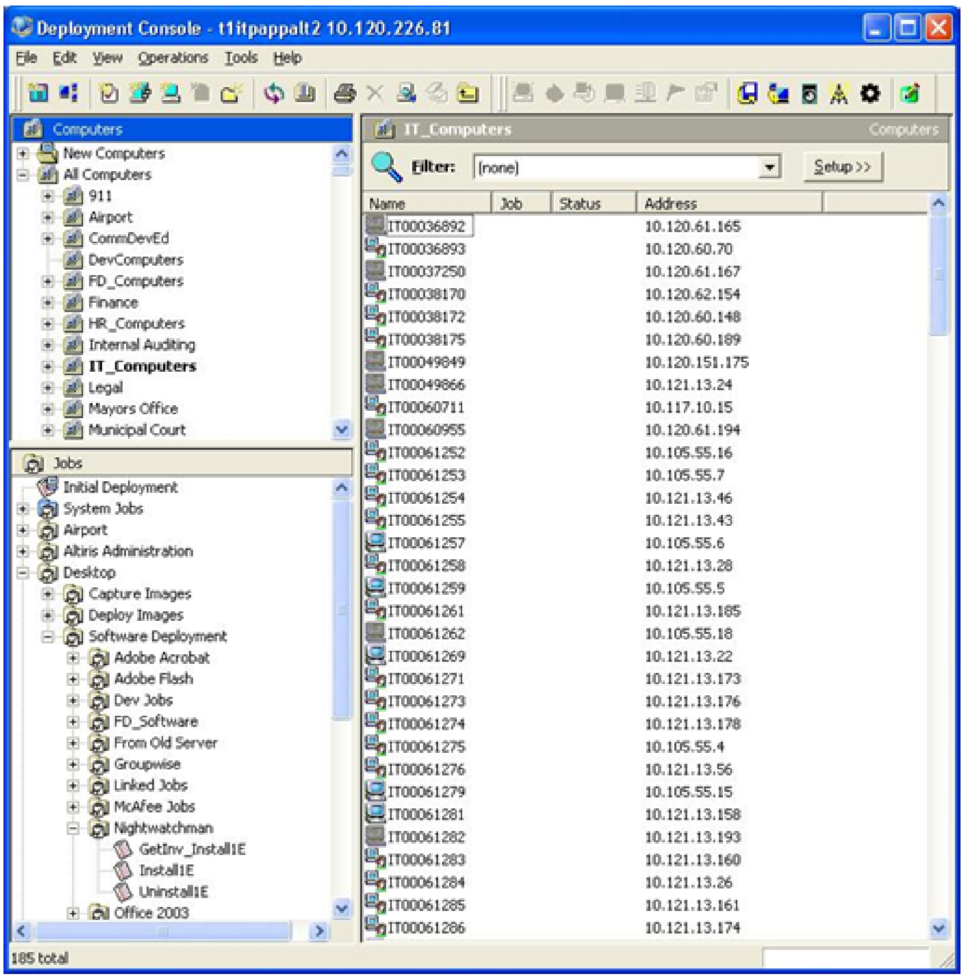

Figure 4.14: Altiris Deployment Console

This Altiris job silently installed the nightwatchman client by pushing an .MSI to the desktop, and then running MSIEXEC with the specified command line parameters.

As of April 20th, the project is on track for successful deployment to 2000 city

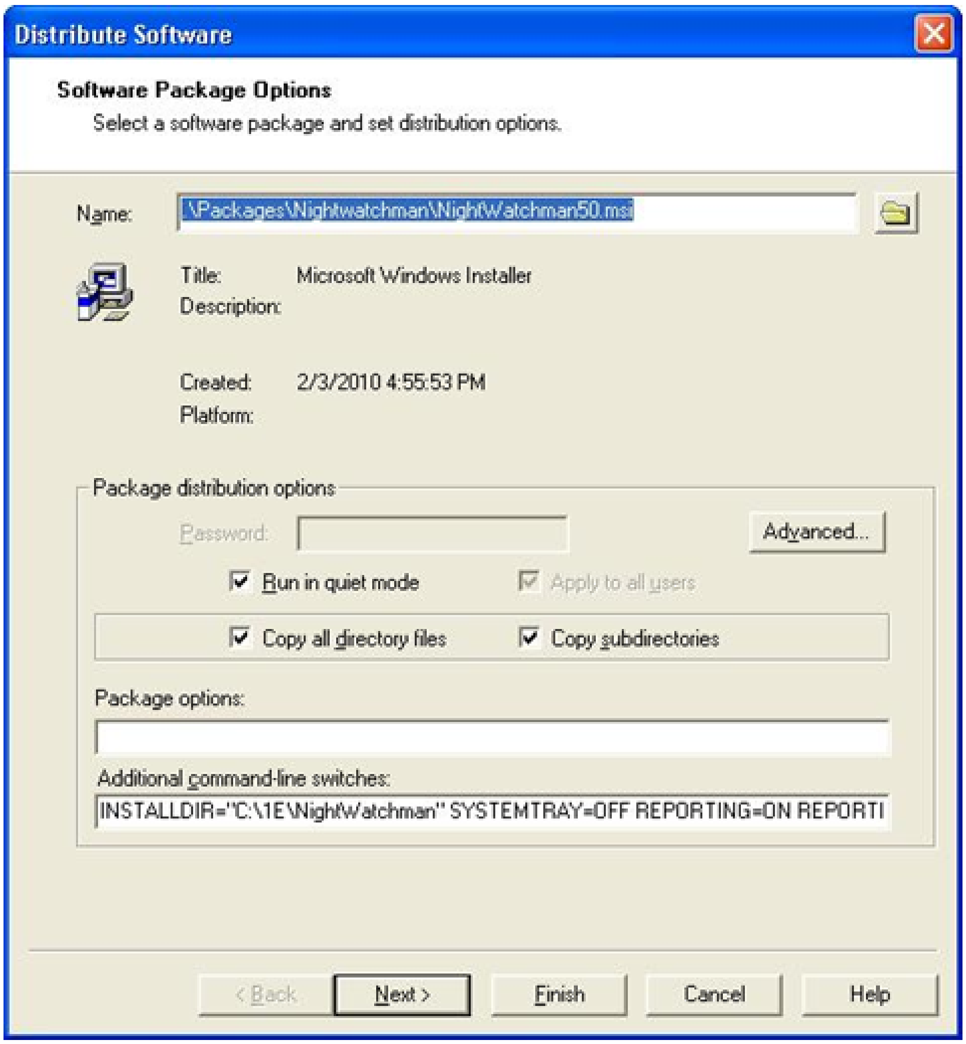

Figure 4.15: Altiris 1E/Nightwatchman Installation Job

computers by May 10th. By July 1st, we will be powering off these computers automatically.

4.2.2 Pilot - Phase II

Pilot Phase II involved silent installation on 151 additional computer on the 6th floor of city hall. This was the first attempt at a reproducible process, concrete steps that could be re-used for future deployments.

- Load computers from Altiris into the Green Computing table by appending.
- Identify managers based on AD information and usernames (update Green Computing table).
- Notify managers, and include a spreadsheet (add all computers assigned to manager by web-form)
- Pull the currently installed workstations from Nightwatchman and update the green computing table.
- Generate an altiris script from the green computing list, everyone who is approved but not installed to install nighwatchman.

IT management dictated installing Nightwatchman software on all computers on the 6th floor ASAP in March.

Change control request for the initial deployment to the 6th floor. The change management committee consisting of 12 IT managers deliberated for 30 minutes on the request. Besides general resistance to change, the main issue seemed to be whether exceptions would be enforced by computer name, or by IP address. We could not store exceptions by IP address because we use DHCP, which assigns dynamic IP addresses.

An aborted trial installation occurred in March for 193 computers using a voluntary web form. Managers responded for only 40 of the computers. All but 2 computers were registered as exceptions to be VPN or production boxes. By policy on the web form, VPN boxes were not eligible for the NW installation without the WU software as well. Pilot Phase 2 ended in failure without implementation. We could not proceed without manager approval, and could not get approvals using a web form. We needed human interaction to overcome objections.

Manager Approved	Group	ComputerName	Category	ADUser	Installed
No	Auditing	IADTT00062963	OK	MAIN\rmaxwell	
No	Auditing	IADTT00062956	OK	MAIN\tmcintosh	
No	Auditing	IADTT00062961	OK	MAIN\sjackson	
No	Auditing	IADTT00062962	OK	MAIN\pdoerflinger	
No	Auditing	IADTT00062964	OK	MAIN\spotter	
No	Auditing	IADTT00062959	OK	MAIN\swagner	
No	Auditing	IADTT00062955	OK	MAIN\kbrader	
No	Auditing	IADTT00062958	OK	MAIN\rrose	
No	Auditing	IADTT00062965	OK	MAIN\cackley	
No	Auditor	D49646	OK	MAIN\npickard	
No	Auditor	IADTT00062957	OK	MAIN\npickard	
No	Axsom - PW	PW36784	OK	MAIN\mnguyen	
No	Axsom - PW	PW37799	VPN	MAIN\daxsom	
No	Axsom - PW	PW37077	VPN	MAIN\AKOMAREK	
No	Axsom - PW	PW00064243	VPN	MAIN\DLISTER	
No	Axsom - PW	PW36377	OK	MAIN\mnguyen	
No	Axsom - PW	PW38037	VPN	MAIN\UBLACK	
No	Axsom - PW	IT00064550	VPN	MAIN\DWIGHTH	
No	Axsom - PW	PW36785	OK	MAIN\hpresident	
No	Axsom - PW	PW36723	OK	MAIN\shill	
No	Axsom - PW	PW35740	VPN	MAIN\DWIGHTH	
No	Axsom - PW	PW36869	OK	MAIN\jpankey	
No	Axsom - PW	PW38061	VPN	MAIN\JWoodward	
No	Axsom - PW	PW37770	VPN	MAIN\JWoodward	
No	Axsom - PW	IT00064528	VPN	MAIN\JGRAY	
No	Axsom - PW	PW36386	Production	MAIN\PALSAPPS	
No	Axsom - PW	IT00064553	Production	MAIN\DPOWELL	
No	Axsom - PW	PW38062	OK	MAIN\tluu	
No	Axsom - PW	prD4G7DB71	VPN	MAIN\DEWYBROYLES	
No	Axsom - PW	IT00063537LT	OK	MAIN\mnguyen	
No	Axsom - PW	PW36725	OK	MAIN\kcrone	
No	Barnes - DS	IT00036893	OK	MAIN\dajcue	
No	Barnes - DS	IT00062131	OK	MAIN\jcue	
No	Barnes - DS	PW36766	OK	MAIN\rleonard	
No	Barnes - DS	POLICE9999	OK	MAIN\ddavid	
No	Barnes - DS	IT00064517	OK	MAIN\jlbarnes	
No	Berg - Service De	IT00064531	OK	MAIN\dbest	

Figure 4.16: Pilot Phase II Installation Approval List

For the Pilot Phase II, a notification email was sent to all IT Users Fri 3/26/2010 5:09 PM

Pilot Phase 3 began when IT Management assigned a Project Manager to the project to overcome obstacles. The project manager obtained approval to silently

User Impact: All Information Technology (IT) Computers at the One Technology Center.

Description of the Change: Information Technology (IT) will begin deployment of the NightWatchman power monitoring software to Information Technology computers for the purpose of measuring and reporting on power consumption data. Critical systems previously identified by users will not be included in this deployment. The automatic shut-down feature of NightWatchman will not be activated during this phase of the software implementation. However, certain software configurations on some individual computers may require a reboot for the initial software installation to complete. Otherwise, this phase of the implementation should have no impact to users.

Estimated Begin Time of Service Change: Monday March 29th at 9:00 a.m. (CDT)

Estimated Time of Completion: Tuesday March 30th at 6:00 p.m. (CDT)

Who to Call If You Have Questions: Service Desk - (918) 596-7070

COT Ticket #: 81614

Change Management Ticket #: 1858

E-mail Approved by Department Head: Jim Langster

Figure 4.17: NOTICE: Deployment of NightWatchman to IT Computers

install Nightwatchman on all of the Phase II computers, in reporting mode.

4.2.3 Staggered Software Deployment

The deployment schedule was managed by the project manager. The schedule was set by location, with one location per week. Exceptions were processed by the service desk, with a ticket being created for each exception. Exceptions were entered into the Green Computing web form to control the installations. Installations were done using the script below.

The final approach to implementation involved these steps: 1) Load computers from Altiris into the Green Computing table by appending 2) Identify managers based on AD information and usernames (update Green Computing table) 3) Notify managers, and include a spreadsheet (add all computers assigned to manager to web-form) 4) Pull the currently installed workstations from Nightwatchman and update the green computing list 5) Generate an altiris script from the green computing list, everyone who is approved but not installed to install nighwatchman.

This is the WSH script used to schedule the Altiris installations of 1E software:

```
for /f %%f in (1ECL.txt)
do "\\t1itpappalt2\eXpress\axSched.exe" %%f "Install1E" /t "2010-02-10 12:00"
/dsn "Altiris eXpress Database"
pause
//computer names are put in 1ECL.TXT by the query, one per line.
```

4.3 Wakeup Issues - Last Man Standing Algorithm

Almost all of the IT staff used RDP (Remote Desktop Protocol) to remotely control their computers at work from home through a VPN (Virtual Private Network). This meant that the power off desktop computers was on hold until we implemented an adequate VPN/web wakeup solution. Fortunately, IT management had already purchased 500 licenses of 1E Wakeup agent. The 1E wakeup software is an additional software license that allows web wakeup of powered off computers using Wake On LAN Magic-Cookie broadcasts from an agent.

"One disadvantage of powering down, or hibernating, a computer is that it is often convenient for administrators to apply software updates during a time when a user will not be using the computer, often late at night. Generally this is an automated update. If the computer is powered down or hibernated, however, it will miss the scheduled update time. To solve this, you can enable a feature of most Ethernet cards called Wake On LAN (WOL). Generally, WOL is enabled with a BIOS setting. Once WOL is enabled, a computer that receives a special network packet will power itself on. This works because even when a computer is turned off (or hibernating), the Ethernet card is still receiving power and can thus receive network packets." [21]

WakeUp Architecture from page 11 of The 1E Web WakeUp Guide:

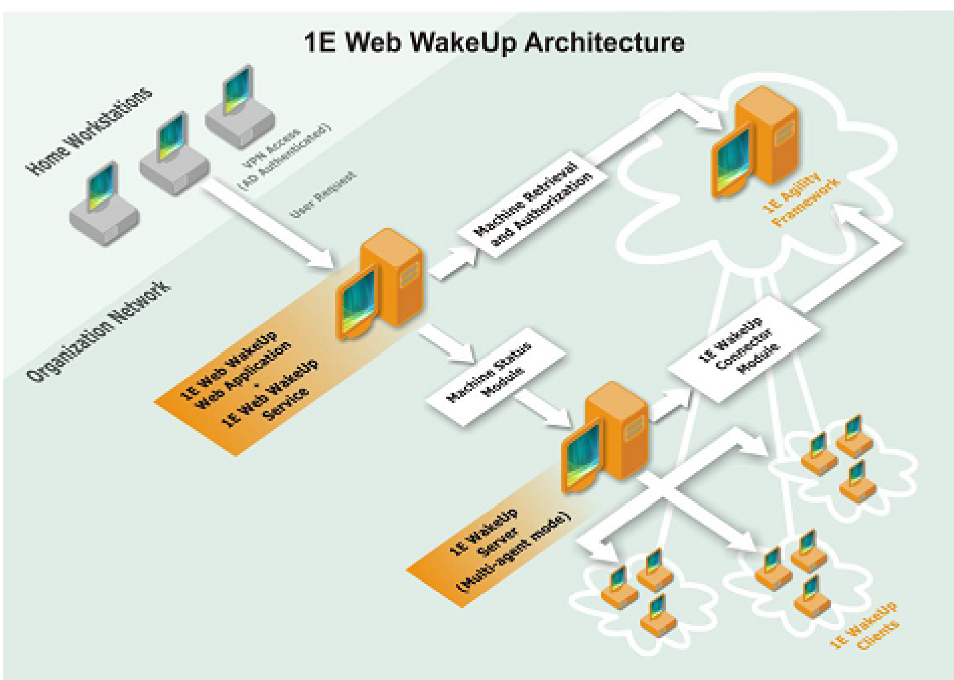

Figure 4.18: 1E Web WakeUp Architecture

Multi-Agent Installation from The 1E WakeUp Installation Guide

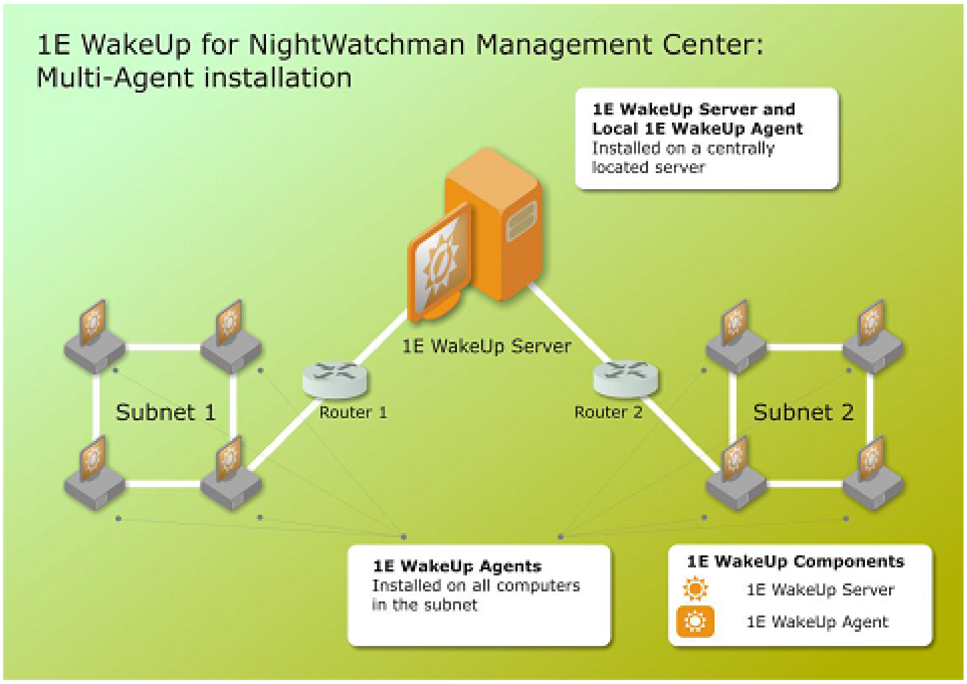

Figure 4.19: 1E WakeUp Multi-Agent Installation

The 1E Wakeup software is closely integrated with Nightwatchman. The Wakeup Server is installed on the same web server as the Nightwatchman Server, and it uses the same database. The agent is installed with a REPORTINGSERVER command line option which tells the agent where the reporting server is located. The 1E Wakeup software takes advantage of Wake On LAN capabilities built into almost all modern NICs. The agent uses a default Magicpacket port of 1776 to receive UDP wakeup broadcasts. The server uses a default of port 1777 to receive TCP communication from the agent. See Appendix J for Wireshark analysis of packets sent.

The city's network is not capable of passing directed subnet broadcasts. The routers are configured so that a broadcast must originate from the subnet. The desktop support manager was concerned that would mean that the wakeup software would not work. According to 1E technical support, this network limitation is common, and their software completely handles it by using agents to broadcast the wakeup packets, and using a last man standing algorithm. Testing verified this as true.

A last man standing algorithm so that one computer with 1E wakeup installed on each subnet is always left powered on. If there is a need to wakeup computers, that primary computer receives a list using TCP/IP from the server which tells which computers to wake-up (in that subnet). If that primary computer starts to shutdown, it sends a wakeup packet to the alternate in the subnet, which then handles the wakeups from that point on.

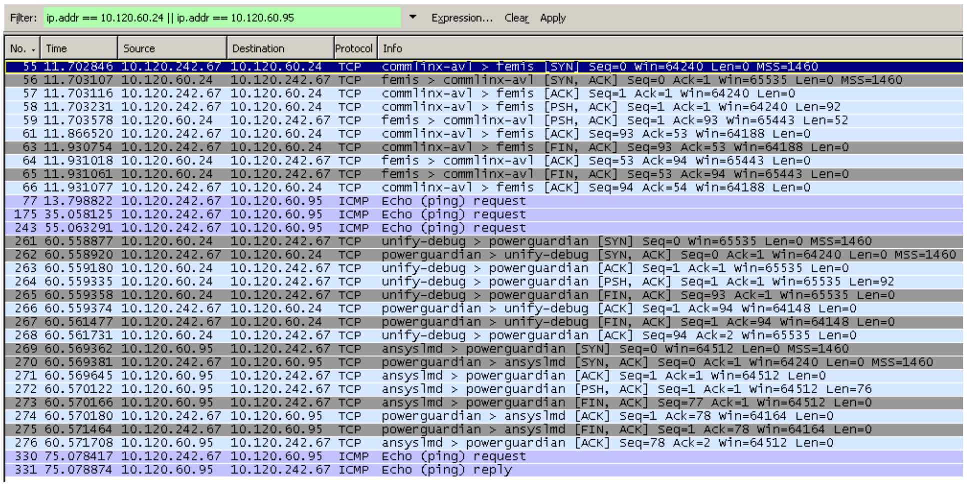

Filter: ip.addr == 10.120.60.24 || ip.addr == 10.120.60.95 Expression... Clear Apply

No.	Time	Source	Destination	Protocol	Info
55	11.702846	10.120.242.67	10.120.60.24	TCP	commlinx-avl > femis [SYN] Seq=0 Win=64240 Len=0 MSS=1460
56	11.703107	10.120.60.24	10.120.242.67	TCP	femis > commlinx-avl [SYN, ACK] Seq=0 Ack=1 Win=65535 Len=0 MSS=1460
57	11.703116	10.120.242.67	10.120.60.24	TCP	commlinx-avl > femis [ACK] Seq=1 Ack=1 Win=64240 Len=0
58	11.703231	10.120.242.67	10.120.60.24	TCP	commlinx-avl > femis [PSH, ACK] Seq=1 Ack=1 Win=64240 Len=92
59	11.703578	10.120.60.24	10.120.242.67	TCP	femis > commlinx-avl [PSH, ACK] Seq=1 Ack=93 Win=65443 Len=52
61	11.866520	10.120.242.67	10.120.60.24	TCP	commlinx-avl > femis [ACK] Seq=93 Ack=53 Win=64188 Len=0
63	11.930754	10.120.242.67	10.120.60.24	TCP	commlinx-avl > femis [FIN, ACK] Seq=93 Ack=53 Win=64188 Len=0
64	11.931018	10.120.60.24	10.120.242.67	TCP	femis > commlinx-avl [ACK] Seq=53 Ack=94 Win=65443 Len=0
65	11.931061	10.120.60.24	10.120.242.67	TCP	femis > commlinx-avl [FIN, ACK] Seq=53 Ack=94 Win=65443 Len=0
66	11.931077	10.120.242.67	10.120.60.24	TCP	commlinx-avl > femis [ACK] Seq=94 Ack=54 Win=64188 Len=0
77	13.798822	10.120.242.67	10.120.60.95	ICMP	Echo (ping) request
175	35.058125	10.120.242.67	10.120.60.95	ICMP	Echo (ping) request
243	55.063291	10.120.242.67	10.120.60.95	ICMP	Echo (ping) request
261	60.558877	10.120.60.24	10.120.242.67	TCP	unify-debug > powerguardian [SYN] Seq=0 Win=65535 Len=0 MSS=1460
262	60.558920	10.120.242.67	10.120.60.24	TCP	powerguardian > unify-debug [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
263	60.559180	10.120.60.24	10.120.242.67	TCP	unify-debug > powerguardian [ACK] Seq=1 Ack=1 Win=65535 Len=0
264	60.559335	10.120.60.24	10.120.242.67	TCP	unify-debug > powerguardian [PSH, ACK] Seq=1 Ack=1 Win=65535 Len=92
265	60.559358	10.120.60.24	10.120.242.67	TCP	unify-debug > powerguardian [FIN, ACK] Seq=93 Ack=1 Win=65535 Len=0
266	60.559374	10.120.242.67	10.120.60.24	TCP	powerguardian > unify-debug [ACK] Seq=1 Ack=94 Win=64148 Len=0
267	60.561477	10.120.242.67	10.120.60.24	TCP	powerguardian > unify-debug [FIN, ACK] Seq=1 Ack=94 Win=64148 Len=0
268	60.561731	10.120.60.24	10.120.242.67	TCP	unify-debug > powerguardian [ACK] Seq=94 Ack=2 Win=65535 Len=0
269	60.569362	10.120.60.95	10.120.242.67	TCP	ansyslmd > powerguardian [SYN] Seq=0 Win=64512 Len=0 MSS=1460
270	60.569381	10.120.242.67	10.120.60.95	TCP	powerguardian > ansyslmd [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
271	60.569645	10.120.60.95	10.120.242.67	TCP	ansyslmd > powerguardian [ACK] Seq=1 Ack=1 Win=64512 Len=0
272	60.570122	10.120.60.95	10.120.242.67	TCP	ansyslmd > powerguardian [PSH, ACK] Seq=1 Ack=1 Win=64512 Len=76
273	60.570166	10.120.60.95	10.120.242.67	TCP	ansyslmd > powerguardian [FIN, ACK] Seq=77 Ack=1 Win=64512 Len=0
274	60.570180	10.120.242.67	10.120.60.95	TCP	powerguardian > ansyslmd [ACK] Seq=1 Ack=78 Win=64164 Len=0
275	60.571464	10.120.242.67	10.120.60.95	TCP	powerguardian > ansyslmd [FIN, ACK] Seq=1 Ack=78 Win=64164 Len=0
276	60.571708	10.120.60.95	10.120.242.67	TCP	ansyslmd > powerguardian [ACK] Seq=78 Ack=2 Win=64512 Len=0
330	75.078417	10.120.242.67	10.120.60.95	ICMP	Echo (ping) request
331	75.078874	10.120.60.95	10.120.242.67	ICMP	Echo (ping) reply

Figure 4.20: Nightwatchman server sends wakeup messages to target computer. These broadcast messages are not relayed by the router.

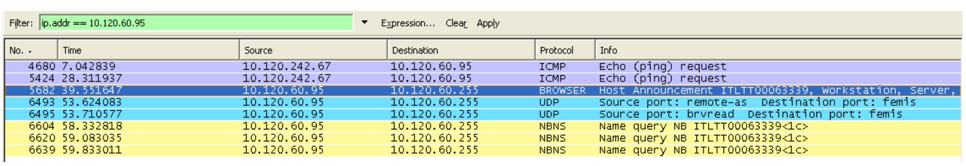

Filter: ip.addr == 10.120.60.95 Expression... Clear Apply

No.	Time	Source	Destination	Protocol	Info
4680	7.042839	10.120.242.67	10.120.60.95	ICMP	Echo (ping) request
5424	28.311937	10.120.242.67	10.120.60.95	ICMP	Echo (ping) request
5682	39.551647	10.120.60.95	10.120.60.255	BROWSER	Host Announcement ITLTT00063339, Workstation, Server,
6493	53.624083	10.120.60.95	10.120.60.255	UDP	Source port: remote-as Destination port: femis
6495	53.710577	10.120.60.95	10.120.60.255	UDP	Source port: brvread Destination port: femis
6604	58.332818	10.120.60.95	10.120.60.255	NBNS	Name query NB ITLTT00063339<1c>
6620	59.083035	10.120.60.95	10.120.60.255	NBNS	Name query NB ITLTT00063339<1c>
6639	59.833011	10.120.60.95	10.120.60.255	NBNS	Name query NB ITLTT00063339<1c>

Figure 4.21: The primary agent sends wakeup packets to the powered off computer.

In interesting behavior of the last man standing algorithm is that if we only have one computer with wakeup installed on a subnet, and Nightwatchman (or the user) tries to shut that computer down, it will instead reboot. It also means that it would be impossible to shut down all of the computers in a sub-net without unplugging a cord. We verified this behavior in the lab.

The city's network, like many networks, does not relay broadcast messages. In

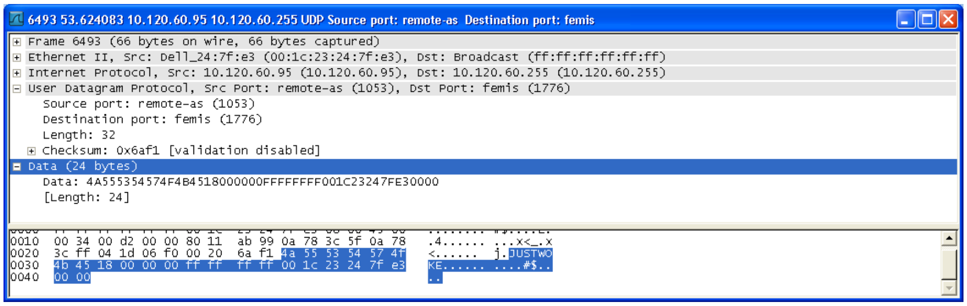

Figure 4.22: The computer that is waking up sends this JUSTWOKE packet.

this case, the software must be implemented in multi-agent mode using a last man standing algorithm.

To wake up a powered off desktop, 1E wakeup server sends a TCP message to the active agent on the subnet of the target computer, which sends a UDP wakeup on LAN Magiccookie to the target computer. Almost all modern NICs accept Wake On LAN magic cookies.

In order to evaluate whether our desktop computers accepted Wake On LAN magic cookies, we did a brief internet search. Since all modern network cards accept Wake On LAN magic cookies, we assumed this would not be a problem. Our plan was to quickly replace any bad network cards on a case by case basis. The replacement was not necessary.

To implement the Last Man Standing algorithm, the server designates a backup agent. If the primary is powered off, or fails to respond to health checks, the backup agent is designated as the primary agent, and a new backup agent is designated.

The Nightwatchman service starts when a computer boots. The Nightwatchman service is a program that runs in the background while the computer is on. The service starts after Windows loads, but before a person logs in. The Nightwatchman service running on each computer reports changes in state to the Nightwatchman server which logs that information to a database. It is the Nightwatchman service that shuts down a computer, but a built in feature of the network card that powers on the computer.

The purpose of broadcast wakeups is to power on large groups of computers to apply patches during the evening. The Altiris deployment console will push maintenance patches and software installs during the evening, but the computer must be powered on. Instead of doing mass wakeups, we established the policy that computers would not be powered down on Monday nights. This allowed a maintenance windows of Monday night when any software patch, installation or inventory could be done.

4.3.1 Wakeup Test Results

Management required that the wake on LAN feature be tested because of skepticism about whether the 1E/Nightwatchman solution would work in the city's environment. The software worked well in testing and through the first year of active use.

The VPN system needed to be able to power on a specific desktop. 1E Wakeup server provides a web interface that allows the entry of a computer name. Upon submission of this web form, the 1E wakeup server sends a TCP message to the active agent on the subnet of the target computer, which sends a UDP wakeup on LAN Magiccookie to the target computer. Because there were so many computers with VPN requirements (including all IT), we needed to develop an easy method for waking up computers from the VPN applicance.

The Juniper VPN device allows connection using RDP/remote desktop through a standard web browser. We configured the VPN appliance to have a permanent link to the Web Wakeup server, so that individuals can wake up their computers after they connect to the VPN system. After waiting 3 to 10 minutes for the computer to fully start, they could connect through the normal VPN RDP utility. The long 3 to 10 minute startup period was caused by old machines, Altiris startup processing, McAfee startup processing, Adobe updates, Java updates and confilcts between different startup routines.

Testing the wake on LAN required 2 different subnets of computers. Because our lab was on one subnet, we waited to test the wake on LAN until Pilot Phase 2. We tested the wake up on working desktops. We powered off the computers both manually, and using Nightwatchman power policy. After computers were powered down, we woke the computers up using the 1E web wakeup form.

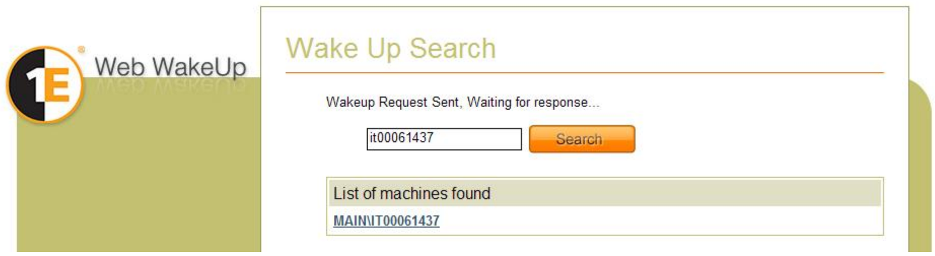

Figure 4.23: 1E/Nightwatchman Web Wakeup

The wakeup function worked very well, but required that each subnet be set up on the Nightwatchman server using the agent setup program. The web wakeup function did fail when you pulled the power cord from the wall instead of shutting down the computer manually. This failure was because the power state was still set to "ON" in the Nightwatchman database. One other shutdown failure for Dell computers was that the BIOS on one machine needed the "Wake On LAN" setting changed from "ON" to "ON BOOT FROM NIC". This counter-intuitive action was repeated on other similar computers.

For troubleshooting of the 1E web wakeup solution, we first confirmed that wakeup packets are being sent and recieved. We use WireShark to ensure that the wake on LAN packets were being sent to the target computer. The computer running Wireshark needs to be on the same subnet as the machine being awakened. You can see the signal being sent from the Nightwatchman server (Figure 4.20 t1itpapp06 10.120.242.67) to the current primary agent, and then the agent sending the wakeup Magic Packets to the target machine. The easiest method is to filter Wireshark results by the target machines IP address. If you see both of these signals, and the IP of the target computer matches, then there is a problem with that machines config (change the BIOS settings).

CHAPTER 5

CASE STUDY - RESULTS

5.1 One Year Later

The Nightwatchman system ran continuously for the year 2011. The savings reports generated by the Nightwatchman system met the expected savings.

5.2 Wake Up Results

A major resistance to powering off desktops is the perception from customers that they will not be able to remotely power on there computers when needed. The Nightwatchman system allowed the remote powering on the desktops through a web page available through the VPN. All web wakeup issues reported to the service desk were successfully corrected.

Figure 1 shows that, in general, it was possible to power-on computers using the remote wake-up function. In one week, of 62 attempts to wake up the low-power desktops, there were 6 failures. This is a failure rate of less than 10% without active management. The Nightwatchman software provides clear capturing of the wake-up errors. If a person were actively managing the error report, they would be proactively troubleshooting the errors. The city instead elected to passively react to service desk tickets of individuals reporting wake up errors.

With active management, it should be possible to reduce the number of errors, but may not be with the additional costs. If the failure rate of less than 10% is not acceptable, a desktop support person could proactively review the errors each week.

Weekly WakeUp Success Report

This report shows the WakeUp successes for the period 08 May 2012 to 14 May 2012

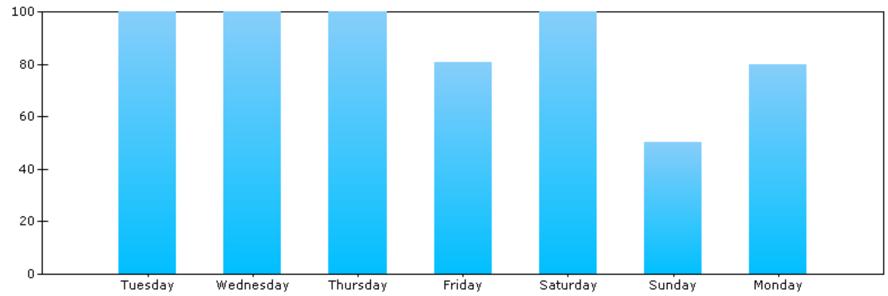

The total height of the column shows the percentage of WakeUp success, which includes machines that were woken up by WakeUp as well as machines that were already awake. The success rate for a day is calculated by the sum of woken up and awake over the total WakeUps for that day.

Day	Total WakeUps	Wokeup	Awake	Failed	Success (%)
Tuesday, 08 May 2012	9	6	3	0	100.00
Wednesday, 09 May 2012	8	7	1	0	100.00
Thursday, 10 May 2012	13	10	3	0	100.00
Friday, 11 May 2012	21	14	3	4	80.95
Saturday, 12 May 2012	4	4	0	0	100.00
Sunday, 13 May 2012	2	1	0	1	50.00
Monday, 14 May 2012	5	2	2	1	80.00

Figure 5.1: Sample Wakeup Results

Determining the cause of the failures, and taking steps to prevent future failures, would lower the failure rate. However, the personnel costs of active management may not be worth the savings.

5.3 Actual Costs

After one year of normal operation, the total usage in for the desktops with Nightwatchman installed was 516,312 kWh.

5.4 Time In State - 2011

The Nightwatchman software records the changes in power state for each desktop. The results from 2011 show that the computers were in use between 20% and 25% of the time. The Nightwatchman software policies resulted in the computers being powered off over 60% of the time. The desktop computers were idle for less than 15% of the time.

The results from 2011 show that the computers were in use between 20% and 25%

This chart shows cost for all organizations in all locations for the period January 2011 to December 2011.

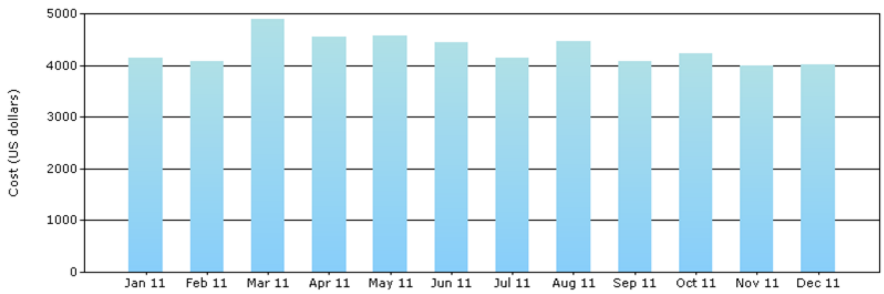

The height of each column shows the total daily cost in US dollars for all organizations in all locations.

	Average per machine			Totals			
Month	Energy	Cost ($)	Emissions (lbs	No. of PCs	Energy (kWh)	Cost ($)	Emissions (lbs
January 2011	24.06	2.41	31.98	1724	41,483	4,148	55,131
February 2011	22.39	2.24	29.75	1819	40,719	4,072	54,116
March 2011	27.54	2.75	36.61	1780	49,028	4,903	65,158
April 2011	26.44	2.64	35.14	1721	45,503	4,550	60,474
May 2011	26.99	2.70	35.87	1695	45,753	4,575	60,806
June 2011	27.00	2.70	35.88	1646	44,437	4,444	59,057
July 2011	25.87	2.59	34.37	1601	41,410	4,141	55,034
August 2011	28.04	2.80	37.26	1596	44,751	4,475	59,474
September 2011	26.55	2.65	35.28	1540	40,884	4,088	54,334
October 2011	27.93	2.79	37.11	1513	42,252	4,225	56,153
November 2011	26.83	2.68	35.65	1491	40,000	4,000	53,160
December 2011	27.24	2.72	36.20	1472	40,090	4,009	53,280
Avg per machine per month/totals	**26.41**	**$2.64**	**35.09**		**516,312**	**$51,631**	**686,178**

This table shows the number of computers in all organizations in all locations reporting into NightWatchman each month for the period January

Figure 5.2: Actual Costs by Month - 2011

Power State	Description
Off	This is a low-power state in which the desktop still responds to wake-on-LAN requests.
Hibernate/Standby	These are low-power states that allow quicker restoration of service than the Off state. The city rarely used these power states.
Monitor Idle	The computer had been idle for 15 minutes, and the monitor was turned off by the power management scheme.
On	The computer had been idle for less than 15 minutes because it was actively being used.

Figure 5.3: COT Nightwatchman Power States

of the time. The Nightwatchman software policies resulted in the computers being powered off over 60% of the time. The desktop computers were idle for less than 15% of the time.

These results indicate that the efficiency of the simple power management scheme was highly effective.

Yearly time in state for laptops and desktops

This report shows the percentage of time laptops and desktops have spent in each reported power state for all organizations in all locations for the period 01 Jan 2011 to 01 Dec 2011

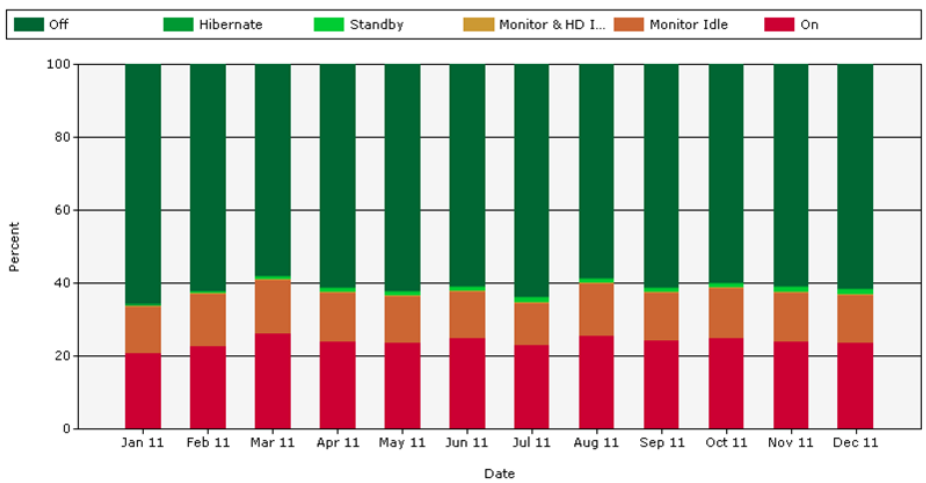

		Percentage Time in State (%)					
Month	Number of PCs	On	Monitor Idle	Monitor & Disk Idle	Standby	Hibernate	Off
Jan 2011	1724	21	13	0	-	1	66
Feb 2011	1819	22	15	0	0	1	62
Mar 2011	1780	26	15	0	1	0	58
Apr 2011	1721	24	14	0	1	0	61
May 2011	1695	24	13	0	1	0	62
Jun 2011	1646	25	13	0	1	0	61
Jul 2011	1601	23	12	0	1	0	64
Aug 2011	1596	25	15	0	1	0	59
Sep 2011	1540	24	14	0	1	0	61
Oct 2011	1513	25	14	0	1	0	60
Nov 2011	1491	24	14	0	1	0	61
Dec 2011	1472	24	13	0	2	0	62

Figure 5.4: Yearly Time in State - 2011

5.5 Time In State - Day of the Week

The power policy specified that desktops would not be powered off on Monday nights. This would allow applying Microsoft Software Updates/Patches as well as allow a maintenance for the automated installation of software. This caused the percentage of time in the On state to be higher on Monday and Tuesday. More computers were

On on Monday night and Tuesday morning to allow for software updates.

The results of one sample week show that on Monday and Tuesday, the time in the Off state was less than 40%. In contrast, on Saturday and Sunday the time in the Off state was more than 60%.

Weekly time in state for desktops

This report shows the percentage of time desktops have spent in each reported power state for all organizations in all locations for the period 07 May 2012 to 13 May 2012

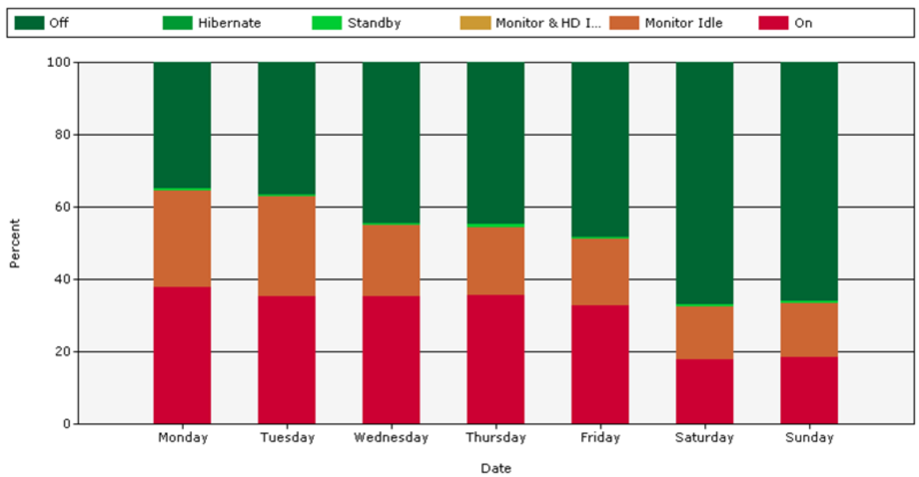

		Percentage Time in State (%)					
Day	Number of PCs	On	Monitor Idle	Monitor & Disk Idle	Standby	Hibernate	Off
Mon, 07 May 2012	1346	38	27	-	1	-	35
Tue, 08 May 2012	1344	35	28	-	1	0	36
Wed, 09 May 2012	1329	35	19	-	1	0	45
Thu, 10 May 2012	1317	35	19	-	1	-	45
Fri, 11 May 2012	1294	33	18	-	1	0	48
Sat, 12 May 2012	1258	18	14	-	1	-	67
Sun, 13 May 2012	1254	18	15	-	1	-	66

Figure 5.5: Weekly Time in State

5.6 Impact of not power managing

One report from 1E/Nightwatchman shows the effect of not power managing desktop computers. Average energy consumption for power managed computers was $30.04 while the average power consumption from the not power managed computers was $59.16. The power consumption of the UNASSIGNED desktops is roughly twice the cost of all other desktops. These other desktops include desktops that are designated to never power off. This result demonstrates that a 50% savings can be achieved through desktop power management.

The procedure for installing Nightwatchman on a new computer included a manual step of assigning the computer to a power management group by location. If the computer was not assigned to a group, no power management was done. Often this step was skipped during the Nightwatchman installation. By default, Nightwatchman places a desktop into the "UNASSIGNED" location. The city opted to not apply a power management policy to the UNASSIGNED group, so those desktops were not power managed until assigned to a group.

For the year of 2011, an average more than 100 desktops were not power managed because of a breakdown in this installation process. IT support people did not check the default "UNASSIGNED" location in the Nightwatchman management tool and properly assign computers to power management groups. An additional $3,200.00 could have been saved if this had been done on a regular schedule.

Yearly energy, cost and emissions by sub-location

This table shows yearly energy, costs and emissions for all organizations in all locations broken down by direct sub-location and includes yearly

	Average per machine			Totals			
Location tier name	Energy	Cost ($)	Emissions (lbs	No. of PCs	Energy (kWh)	Cost ($)	Emissions (lbs
UNASSIGNED	591.57	59.16	786.20	113	66,848	6,685	88,841
City of Tulsa	300.44	30.04	399.29	1496	449,464	44,946	597,337
Avg per machine per year/totals	**446.01**	**$44.60**	**592.75**	**1609**	**516,312**	**$51,631**	**686,178**

Figure 5.6: Actual Costs of Unassigned Desktops - 2011

5.7 Actual Savings

These results indicate that the efficiency of the power management scheme was effective. Total savings for the computers that had Nightwatchman software installed was $98,972.00.

The number of power managed computers dropped from 1750 to 1500 by the end of 2011. A more aggressive installation procedure would have allowed the city to reach the goal of power managing 2000 desktop computers, and achieve additional savings.

kWh Savings, Monthly

Estimate monthly kWh savings that have been made by using NightWatchman

Date	Number of PCs	Total Actual (kWh)	Total Estimated (kWh)	Total Saving (kWh)	Average Actual (kWh)	Average Estimated (kWh)	Average Saving (kWh)
May 2011	1652	45641.22	94865.57	49224.35	27.63	57.42	29.80
Jun 2011	1610	44306.20	92186.43	47880.23	27.52	57.26	29.74
Jul 2011	1551	41303.70	85310.40	44006.69	26.63	55.00	28.37
Aug 2011	1560	44578.89	89951.22	45372.33	28.58	57.66	29.08
Sep 2011	1509	40819.08	85818.82	44999.74	27.05	56.87	29.82
Oct 2011	1478	42164.44	79595.65	37431.21	28.53	53.85	25.33
Nov 2011	1445	39823.55	83230.54	43406.99	27.56	57.60	30.04
Dec 2011	1423	39958.05	82190.06	42232.02	28.08	57.76	29.68
Jan 2012	1443	41903.67	82266.58	40362.91	29.04	57.01	27.97
Feb 2012	1498	45306.72	81891.26	36584.54	30.24	54.67	24.42
Mar 2012	1477	46897.80	85362.82	38465.02	31.75	57.79	26.04
Apr 2012	1474	45942.58	80370.60	34428.02	31.17	54.53	23.36
May 2012	1419	21544.31	34331.76	12787.44	15.18	24.19	9.01
Total		**540190.22**	**1057371.71**	**517181.49**	**358.96**	**701.62**	**342.66**

Figure 5.7: Total kWh Saved 2011

Yearly NightWatchman Savings Report

This report shows the savings made by NightWatchman power management in terms of cost and carbon emissions for all departments in all buildings for the period Jan 2011 to Dec 2011

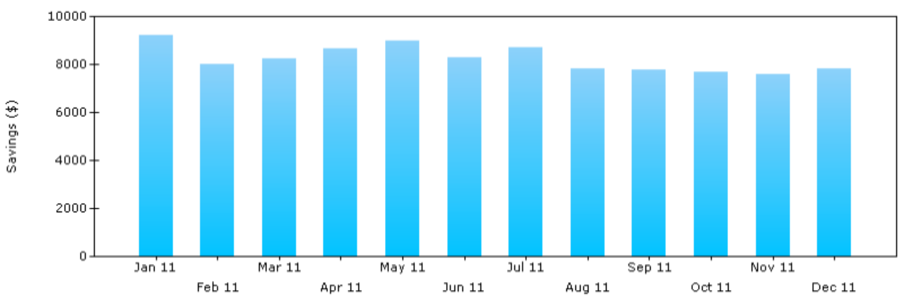

The total height of the column shows the savings for all of the PCs for all departments in all buildings for that month. The savings are calculated from the cost based on the "Always on" power behavior and the actual cost reported by NightWatchman.

	Average Per PC		Totals for selected group			
Month	Actual ($)	Savings ($)	Number of PCs	Actual ($)	Savings ($)	Savings Co2 (lbs)
Jan 2011	2.41	5.36	1724	4,148	9,238	122,776
Feb 2011	2.24	4.41	1819	4,072	8,015	106,519
Mar 2011	2.75	4.63	1780	4,903	8,242	109,542
Apr 2011	2.64	5.03	1721	4,550	8,653	114,995
May 2011	2.70	5.30	1695	4,575	8,979	119,337
Jun 2011	2.70	5.03	1646	4,444	8,275	109,972
Jul 2011	2.59	5.44	1601	4,141	8,715	115,823
Aug 2011	2.80	4.89	1596	4,475	7,812	103,815
Sep 2011	2.65	5.04	1540	4,088	7,760	103,136
Oct 2011	2.79	5.07	1513	4,225	7,667	101,898
Nov 2011	2.68	5.10	1491	4,000	7,599	100,989
Dec 2011	2.72	5.32	1472	4,009	7,836	104,143
Total	**$31.67**	**$60.62**		**$51,631**	**$98,792**	**1,312,946**

There are four main factors which affect costs and emissions.

- Behavior: The amount of time the PCs spend in each power state will affect how much electricity they use and therefore the cost and carbon emissions. This is the primary mechanism for NightWatchman to realize savings.
- Number of PCs under power management: The number of PCs under power management directly affects the overall reported electricity use and therefore the cost and emissions. This value can change as more NightWatchman clients are provisioned or new PCs are installed within the organization.
- Power consumption per PC: The power consumption of the PCs have a direct effect on the amount of electricity used and therefore also the cost and carbon emissions. The average power consumption will change as the mix of hardware in the organization changes over time - new models tend to consume less power than old models.
- Tariffs: The tariff in use may change over time and will affect the cost and/or carbon emissions.

Figure 5.8: Total Savings Report 2011

Green Computing Customers Lessons Learned Results

10/21/10

What went well?

Computer Shutdowns

Good Planning

Good Prior Communications

Smooth Installations

Identification of critical computers

Email notifications

Walkthroughs on day of installation

Different configuration assistance – Apple used to Remote in City PC

What didn't go well?

Message is not prompted at 5:45pm – message delay

Loss in productivity – boot up – 16 minutes

Shift changes affected by NightWatchman

Power Saving Mode takes 2 – 3 minutes – Dual Monitors

What can we do to improve?

Wakeup machines prior to employee coming in

Vendor request – fix prompt delay

Shift changes addressed with Computer Exception Form

Accurate/Updated user information – City Issue

Figure 5.9: Lessons Learned

CHAPTER 6

CONCLUSION

Green Computing is the effort to reduce the power consumption of computers through efficiency. It follows the ideal of sustainable development, progress without compromising the future. Though computers use just 3% of the power in the U.S., computer based solutions make Green Computing's influence have more affect.

The terms Green Computing, Green IT and Sustainable Computing can be used interchangeably. When they are actively pursued, they address the ethical responsiblity to protect the environment by promoting power efficiency. For corporations, Green Computing is an integral part of a Corporate Social Responsibility program. Three motivations for Green Computing are profit, regulation and ethics.

Power efficiency of use and manufacture are design issues. To recycle computers, they must be manufactured to facilitate recycling. To fully enable reduced power consumption, the components must be designed to consume power proportional to the amount of useful work done. Special Green Computing sustainability efforts will not be necessary when these principles are fully integrated into computer computer design.

The amount of carbon produced is proportional to the power consumed. If the source of power can be chosen, then sophistcated task provisioning can be used to run tasks at more efficient centers. Without the ability to run a task at a more efficient location, the carbon footprint minimization problem reduces to the power usage minimization problem.

Because of wasted power from idle components or leakage, there is at least a 50%

power savings to be gained through a power saving measures for servers and computers. Because many data center computers must also be actively cooled, the potential savings are increased. Solutions like virtualization and energy efficient scheduling provide viable power savings potential for servers, while power state management software like 1E/Nightwatchman provide immediate dramatic savings for desktop computers.

Replacing an inefficient computer with a more efficient model does not make sense because of the very high power cost for manufacturing computers.

Computer technology can be used to facility energy efficiency in other areas. For example, telecommuting could save huge amounts of fuel. Also, using software to power manage buildings has proven significant savings.

CMOS CPU power usage can be calulated by the number of transistors that switch. Because of scaling, static power consumption is a significant problem which represent wasted power.

Power measurement in Watts is a point in time measurement. Power measurement in Joules is the sum of the point in time measurements. Because of the high clock rate, and rapid changes in Watts, CPU measurements in Joules are estimates based on the sampling rate. Because of this, average Watts are used as the power part of the current leading power/performance metrics.

Reducing power in data centers must address the cooling issue. Running systems at higher temperatures increases energy costs and reduces reliability. Data centers which can avoid active cooling or use more efficient sources of power are able to reduce the total power consumed and have a lower PUE.

Power utilization has become a primary issue for high performance computing. The Green 500 shows supercomputers that perform with the highest power efficiency. New supercomputers like the K computer boast power efficient design.

Power measurement can be accurately done with inexpensive equipment. Standard tools include the Watts Up! Pro meter, which is less than $200, and is the standard for the Green 500. More advanced meters allow the monitoring to be done at higher rates or be permanently installed, but do not provide more accurate average power consumption.

The optimal power to performance ratio is predicted to be slightly below the optimal performance settings. This is because components consume more power as they heat up.

Idle desktop computers waste a huge amount of electricity. Approximatedly 1.5 TWh of electricity are wasted each year by idle computers. A 2008 EPA estimate for power savings potential is 25 to 75 dollars per year per computer. This energy could be saved by implementing proper desktop computer power management. Proper desktop computer power management means turning off computer components and whole computers when they are not in use.

The naive power management solution for desktop computers is to turn the computer off when not in use. This solution is viable will save 50% or more power when compared to a computer that is not power managed. Remote wakeup and access, and 24/7 production boxes must be addressed for large scale deployment inside organizations like a city government.

BIBLIOGRAPHY

[1] B. Zhong, M. Feng, and C.-H. Lung, "A green computing based architecture comparison and analysis," in *Green Computing and Communications (GreenCom), 2010 IEEE/ACM Int'l Conference on Int'l Conference on Cyber, Physical and Social Computing (CPSCom)*, pp. 386 –391, dec. 2010.

[2] C. Tseng and S. Figueira, "An analysis of the energy efficiency of multi-threading on multi-core machines," in *Green Computing Conference, 2010 International*, pp. 283 –290, aug. 2010.

[3] A. Haywood, J. Sherbeck, P. Phelan, G. Varsamopoulos, and S. Gupta, "A sustainable data center with heat-activated cooling," in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2010 12th IEEE Intersociety Conference on*, pp. 1 –7, june 2010.

[4] C. of Tulsa Communications, "City of tulsa energy conservation and efficiency plan 2007." `http://www.cityoftulsa.org/COTlegacy/documents/CityofTulsaEnergyConservationandEfficiencyPlan.pdf`.

[5] R. Harmon and N. Auseklis, "Sustainable it services: Assessing the impact of green computing practices," in *Management of Engineering Technology, 2009. PICMET 2009. Portland International Conference on*, pp. 1707 –1717, aug. 2009.

[6] Dazor, "Green lights: A bright investment in the environment." web link: `http://www.dazor.com/green-lights.html`.

[7] S. Ruth, "Green it more than a three percent solution?," *Internet Computing, IEEE*, vol. 13, pp. 74 –78, july-aug. 2009.

[8] EPA, "Energy star." web link: `http://www.energystar.gov/`.

[9] C. Landwehr, "Green computing," *Security Privacy, IEEE*, vol. 3, p. 3, nov.-dec. 2005.

[10] U. Nations, "42/187. report of the world commission on environment and development." Web link: `http://www.un.org/documents/ga/res/42/ares42-187.htm`, 1987.

[11] Y. Cai, "Integrating sustainability into undergraduate computing education," in *Proceedings of the 41st ACM technical symposium on Computer science education*, SIGCSE '10, (New York, NY, USA), pp. 524–528, ACM, 2010.

[12] S. Murugesan, "Harnessing green it: Principles and practices," *IT Professional*, vol. 10, pp. 24 –33, jan.-feb. 2008.

[13] R. Harmon, H. Demirkan, N. Auseklis, and M. Reinoso, "From green computing to sustainable it: Developing a sustainable service orientation," in *System Sciences (HICSS), 2010 43rd Hawaii International Conference on*, pp. 1 –10, jan. 2010.

[14] L. Wilbanks, "Green: My favorite color," *IT Professional*, vol. 10, p. 64, nov.-dec. 2008.

[15] I. Siddavatam, E. Johri, and D. Patole, "Optimization of load balancing algorithm for green it," in *Proceedings of the International Conference & Workshop on Emerging Trends in Technology*, ICWET '11, (New York, NY, USA), pp. 1344–1346, ACM, 2011.

[16] S. Murugesan, "Making it green," *IT Professional*, vol. 12, pp. 4 –5, march-april 2010.

[17] J. Hedman and S. Henningsson, "Three strategies for green it," *IT Professional*, vol. 13, pp. 54 –57, jan.-feb. 2011.

[18] P. Kurp, "Green computing," *Commun. ACM*, vol. 51, pp. 11–13, October 2008.

[19] J. Williams and L. Curtis, "Green: The new computing coat of arms?," *IT Professional*, vol. 10, pp. 12 –16, jan.-feb. 2008.

[20] EPA, "Epa report to congress on server and data center energy efficiency." Web link:
`http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Report_Exec_Summary_Final.pdf`, 2007.

[21] D. R. Herrick and M. R. Ritschard, "Greening your computing technology, the near and far perspectives," in *Proceedings of the 37th annual ACM SIGUCCS fall conference*, SIGUCCS '09, (New York, NY, USA), pp. 297–304, ACM, 2009.

[22] D. Niyato, S. Chaisiri, and L. B. Sung, "Optimal power management for server farm to support green computing," in *Cluster Computing and the Grid, 2009. CCGRID '09. 9th IEEE/ACM International Symposium on*, pp. 84 –91, may 2009.

[23] L. Barroso and U. Holzle, "The case for energy-proportional computing," *Computer*, vol. 40, pp. 33 –37, dec. 2007.

[24] V. Shekar and B. Izadi, "Energy aware scheduling for dag structured applications on heterogeneous and dvs enabled processors," in *Green Computing Conference, 2010 International*, pp. 495 –502, aug. 2010.

[25] F. Pinel, J. Pecero, P. Bouvry, and S. Khan, "Memory-aware green scheduling on multi-core processors," in *Parallel Processing Workshops (ICPPW), 2010 39th International Conference on*, pp. 485 –488, sept. 2010.

[26] H. Topcuoglu, S. Hariri, and M.-Y. Wu, "Performance-effective and low-complexity task scheduling for heterogeneous computing," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 13, pp. 260 –274, mar 2002.

[27] L. Zhang, K. Li, and Y.-Q. Zhang, "Green task scheduling algorithms with energy reduction on heterogeneous computers," in *Progress in Informatics and Computing (PIC), 2010 IEEE International Conference on*, vol. 1, pp. 560 –563, dec. 2010.

[28] L. M. Zhang, K. Li, and Y.-Q. Zhang, "Green task scheduling algorithms with speeds optimization on heterogeneous cloud servers," in *Green Computing and Communications (GreenCom), 2010 IEEE/ACM Int'l Conference on Int'l Conference on Cyber, Physical and Social Computing (CPSCom)*, pp. 76 –80, dec. 2010.

[29] K. Roy, S. Mukhopadhyay, and H. Mahmoodi-Meimand, "Leakage current mechanisms and leakage reduction techniques in deep- submicrometer cmos circuits," *Proceedings of the IEEE*, vol. 91, pp. 305 – 327, feb 2003.

[30] A. Manzanares, X. Qin, X. Ruan, and S. Yin, "Pre-bud: Prefetching for energy-efficient parallel i/o systems with buffer disks," *Trans. Storage*, vol. 7, pp. 3:1–3:29, June 2011.

[31] C. Humphries and P. Ruth, "Towards power efficient consolidation and distribution of virtual machines," in *Proceedings of the 48th Annual Southeast Regional Conference*, ACM SE '10, (New York, NY, USA), pp. 75:1–75:6, ACM, 2010.

[32] D. Przybyla and M. Pegah, "Dealing with the veiled devil: eco-responsible computing strategy," in *Proceedings of the 35th annual ACM SIGUCCS fall conference*, SIGUCCS '07, (New York, NY, USA), pp. 296–301, ACM, 2007.

[33] D. Wang, “Meeting green computing challenges,” in *Electronics Packaging Technology Conference, 2008. EPTC 2008. 10th*, pp. 121 –126, dec. 2008.

[34] G. Goth, “The netś going green: Multipronged approach might save costs, energy $ and the climate,” *Internet Computing, IEEE*, vol. 12, pp. 7 –9, jan.-feb. 2008.

[35] Lightwave.com, “Ieee ratifies new 802.3az standard to reduce network energy footprint.” web link: `http://www.lightwaveonline.com/articles/2010/10/ieeel`.

[36] P. Brenner, R. Jansen, D. Go, and D. Thain, “Environmentally opportunistic computing transforming the data center for economic and environmental sustainability,” in *Green Computing Conference, 2010 International*, pp. 383 –388, aug. 2010.

[37] S. Naffziger, B. Stackhouse, T. Grutkowski, D. Josephson, J. Desai, E. Alon, and M. Horowitz, “The implementation of a 2-core, multi-threaded itanium family processor,” *Solid-State Circuits, IEEE Journal of*, vol. 41, pp. 197 – 209, jan. 2006.

[38] S. Mukhopadhyay, K. Kim, C. T. Chuang, and K. Roy, “Modeling and analysis of leakage currents in double-gate technologies,” *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 25, pp. 2052 –2061, oct. 2006.

[39] R. Murphy, T. Sterling, and C. Dekate, “Advanced architectures and execution models to support green computing,” *Computing in Science Engineering*, vol. 12, pp. 38 –47, nov.-dec. 2010.

[40] A. Jaiantilal, Y. Jiang, and S. Mishra, “Modeling cpu energy consumption for energy efficient scheduling,” in *Proceedings of the 1st Workshop on Green Computing*, GCM ’10, (New York, NY, USA), pp. 10–15, ACM, 2010.

[41] W. chun Feng, X. Feng, and R. Ce, "Green supercomputing comes of age," *IT Professional*, vol. 10, pp. 17 –23, jan.-feb. 2008.

[42] D. Wang, "Meeting green computing challenges," in *High Density packaging and Microsystem Integration, 2007. HDP '07. International Symposium on*, pp. 1 –4, june 2007.

[43] B. Subramaniam and W. chun Feng, "Understanding power measurement implications in the green500 list," in *Green Computing and Communications (GreenCom), 2010 IEEE/ACM Int'l Conference on Int'l Conference on Cyber, Physical and Social Computing (CPSCom)*, pp. 245 –251, dec. 2010.

[44] C.-H. Hsu, W.-C. Feng, and J. Archuleta, "Towards efficient supercomputing: a quest for the right metric," in *Parallel and Distributed Processing Symposium, 2005. Proceedings. 19th IEEE International*, p. 8 pp., april 2005.

[45] H. P. K. C. W. F. R. Ge, X. Feng, "Power measurement tutorial for the green500 list." web link: `http://www.green500.org/sites/default/files/tutorial.pdf`.

[46] W.-C. Feng, A. Ching, and C.-H. Hsu, "Green supercomputing in a desktop box," in *Parallel and Distributed Processing Symposium, 2007. IPDPS 2007. IEEE International*, pp. 1 –8, march 2007.

[47] A. Banerjee, T. Mukherjee, G. Varsamopoulos, and S. Gupta, "Cooling-aware and thermal-aware workload placement for green hpc data centers," in *Green Computing Conference, 2010 International*, pp. 245 –256, aug. 2010.

[48] M. Yokokawa, F. Shoji, A. Uno, M. Kurokawa, and T. Watanabe, "The k computer: Japanese next-generation supercomputer development project," in *Low Power Electronics and Design (ISLPED) 2011 International Symposium on*, pp. 371 –372, aug. 2011.

[49] e. P.M. Kogge, "Exascale computing study: Technology challenges in achieving exascale systems, tech. report tr-2008-13." Computer Science and Eng. Dept., Univ. Notre Dame, 28 Sept. 2008 http://www.cse.nd.edu/Reports/2008/TR-2008-13.pdf.

[50] C.-H. Hsu and S. W. Poole, "Power measurement for high performance computing: State of the art," in *Green Computing Conference and Workshops (IGCC), 2011 International*, pp. 1 –6, july 2011.

[51] B. Subramaniam and W. chun Feng, "Statistical power and performance modeling for optimizing the energy efficiency of scientific computing," in *Green Computing and Communications (GreenCom), 2010 IEEE/ACM Int'l Conference on Int'l Conference on Cyber, Physical and Social Computing (CPSCom)*, pp. 139 –146, dec. 2010.

[52] P. Zhang and B. Helvik, "Towards green p2p: Understanding the energy consumption in p2p under content pollution," in *Green Computing and Communications (GreenCom), 2010 IEEE/ACM Int'l Conference on Int'l Conference on Cyber, Physical and Social Computing (CPSCom)*, pp. 332 –337, dec. 2010.

[53] Z. Xu, "Building a power-aware database management system," in *Proceedings of the Fourth SIGMOD PhD Workshop on Innovative Database Research*, IDAR '10, (New York, NY, USA), pp. 1–6, ACM, 2010.

[54] Y. Fukuda, T. Ikenaga, H. Tamura, M. Uchida, K. Kawahara, and Y. Oie, "Performance evaluation of power saving scheme with dynamic transmission capacity control," in *GLOBECOM Workshops, 2009 IEEE*, pp. 1 –5, 30 2009-dec. 4 2009.

[55] P. Attewell, "Computerization and controversy," ch. Big brother and the sweatshop: computer surveillance in the automated office, pp. 236–256, San Diego, CA, USA: Academic Press Professional, Inc., 1991.

[56] Y. Agarwal, S. Hodges, R. Chandra, J. Scott, P. Bahl, and R. Gupta, "Somniloquy: augmenting network interfaces to reduce pc energy usage," in *Proceedings of the 6th USENIX symposium on Networked systems design and implementation*, NSDI'09, (Berkeley, CA, USA), pp. 365–380, USENIX Association, 2009.

VITA

Charles Letcher

Candidate for the Degree of

Master of Science

Thesis: GREEN COMPUTING - DESKTOP COMPUTER POWER MANAGEMENT AT THE CITY OF TULSA

Major Field: Computer Science

Biographical:

Personal Data: Born in Miami, Oklahoma, USA on June 27, 1962.

Education:
Received the B.M. degree from University of Tulsa, Tulsa, OK, USA, 1991, in Piano Performance
Completed the requirements for the degree of Master of Science with a major in Computer Science at Oklahoma State University in May, 2013.

Experience:
Worked as Enterprise Webmaster at the City of Tulsa from March 2007 to July 2012