Green Networking

Sudha K
UG Students,
VSB Engineering College,
Karur.

Sujitha K
UG Students,
VSB Engineering College,
Karur.

Selvanayaki P
UG Students,
VSB Engineering College,
Karur.

Sudar Vizhi S
UG Students,
VSB Engineering College,
Karur.

Abstract- Green networking refers to the practice of using networking resources more efficiently while maintaining or increasing overall performance. Sustainable IT services require the integration of green computing practices such as power management, virtualization, improving cooling technology, recycling, electronic waste disposal, and optimization of the IT infrastructure to meet sustainability requirements. Recent studies have shown that costs of power utilized by IT departments can approach 50% of the overall energy costs for an organization. While there is an expectation that green IT should lower costs and the firm's impact on the environment, there has been far less attention directed at understanding the strategic benefits of sustainable IT services in terms of the creation of customer value, business value and societal value. This paper provides a review of the literature on sustainable IT, key areas of focus, and identifies a core set of principles to guide sustainable IT service design.

Keywords: *IT service design, IT infrastructure, Energy efficiency .*

I. INTRODUCTION

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. This was followed by the ENERGY STAR program in 1992, which established Energy efficiency specifications for computers and monitors.

Green networking refers to the practice of using networking resources more efficiently while maintaining or increasing overall performance. Sustainable IT services require the integration of green computing practices such as power management, virtualization, improving cooling technology, recycling, electronic waste disposal, and optimization of the IT infrastructure to meet sustainability requirements. Recent studies have shown that costs of power utilized by IT departments can approach 50% of the overall energy costs for an organization. While there is an expectation that green IT should lower costs and the firm's impact on the environment, there has been far less attention directed at understanding the strategic benefits of sustainable IT services in terms of the creation of customer value, business value and societal value. This paper provides a review of the literature on sustainable IT, key areas of focus, and identifies a core set of principles to guide sustainable IT service design.

II. FACTORS DRIVING THE ADOPTION OF GREEN COMPUTING:

The following trends are impacting data centers, and to a lesser degree, desktop computers, and driving the adoption of green networking practices:

A. The Rapid Growth of Internet

The increasing reliance on electronic data is driving the rapid growth in the size and number of data centers. This growth results from the rapid adoption of Internet communications and media, the computerization of business processes and applications, legal requirements for retention of records, and disaster recovery. Internet usage is growing at more than 10 percent annually leading to an estimated 20% CAGR in data center demand. Video and music downloads, on-line gaming, social networks, e-commerce, and VoIP are key drivers. In addition, business use of the Internet has ramped up. Industries such as financial services (investment, banking, and insurance), real estate, healthcare, retailing, manufacturing, and transportation are using information technology for key business functions . The advent of the Sarbanes-Oxley Act with its requirement to retain electronic records has increased storage demand in some industries at 50 percent CAGR . Disaster recovery strategies that mandate duplicate records increases demand further. Finally, many federal, state, and local government agencies have adopted e-government strategies that utilize the Web for public information, reporting, transactions, homeland security, and scientific computing .

B. Increasing Equipment Power Density

Although advances in server CPUs have in some cases enabled higher performance with less power consumption per CPU, overall server power consumption has continued to increase as more servers are installed with higher performance power-hungry processors with more memory capacity . As more servers are installed they require more floor space. To pack more servers in the same footprint the form factor of servers has become much smaller, in some cases shrinking by more than 70% through the use of blade servers. This increase in packaging density has been matched by a major increase in the power density of data centers. Density has increased more than ten times from 300 watts per square foot in 1996 to over 4,000 watts per square foot in 2007, a trend that is expected to continue its upward spiral.

C. Increasing Cooling Requirements

The increase in server power density has led to a concomitant increase in data center heat density. Servers require approximately 1 to 1.5 watts of cooling for each watt of power used. The ratio of cooling power to server power requirements will continue to increase as data center server densities increase.

D. Increasing Energy Cost

Data center expenditures for power and cooling can exceed that for equipment over the useful life of a server. For a typical $4,000 server rated at 500 watts, one study estimated it would consume approximately $4,000 of electricity for power and cooling over three years, at $0.08 per kilowatt-hour, and double that in Japan. The ratio of power and cooling expense to equipment expenses has increased from approximately 0.1 to 1 in 2000 to 1 to 1 in 2007 . With the likely increase in the number of data centers and servers and the advent of a carbon cap-and-trade scheme, the cost of energy for data center power and cooling will continue to increase .

E. Restriction on Energy Supply AD Access

Companies such as Google, Microsoft, and Yahoo with the need for large data centers may not be able to find power at any price in major American cities. Therefore, they have built new data centers in the Pacific Northwest near the Columbia River where they have direct access to low-cost hydroelectric power and don’t need to depend on the overtaxed electrical grid. In states such as, California, Illinois, and New York, the aging electrical infrastructure and high costs of power can stall or stop the construction of new data centers and limit the operations of existing centers. In some crowded urban areas utility power feeds are at capacity and electricity is not available for new data centers at any price .

F. Low Server Utilization Rates

Data center efficiency is a major problem in terms of energy use. The server utilization rates average 5-10 percent for large data centers. Low server utilization means that companies are overpaying for energy, maintenance, operations support, while only using a small percentage of computing capacity.

G. Growing Awareness of It Impacts on the Environment

Carbon emissions are proportional to energy usage. In 2007 there were approximately 44 million servers worldwide consuming 0.5% of all electricity. Data centers in the server-dense U.S. use more than 1% of all electricity. Their collective annual carbon emissions of 80 metric megatons of CO_2 are approaching the carbon footprint of the Netherlands and Argentina . Carbon emissions from operations are expected to grow at more than 11% per year to 340 metric megatons by 2020. In addition, the carbon footprint of manufacturing the IT product is largely unaccounted for by IT organizations.

III. IMPLEMENTING GREEN COMPUTING STATERGIES

Transitioning to green computing has involved a number of strategies to optimize the efficiency of data center operations in order to lower costs and to lessen the impact of computing on the environment. The transitioning to a green data center involves a mix of integrating new approaches for power and cooling with energy efficient hardware, virtualization, software, and power and workload management.

A. Data Center Infrastructure

Infrastructure equipment includes chillers, power supplies, storage devices, switches, pumps, fans, and network equipment. Many data centers are over ten years old. Their infrastructure equipment is reaching the end of its useful life. It is power hungry and inefficient. Such data centers typically use 2 or 3 times the amount of power overall as used for the IT equipment, mostly for cooling . The obvious strategy here has been to invest in new data centers that are designed to be energy efficient or to retrofit existing centers.

B. Power and Workload Management

Power and workload management software could save $25-75 per desktop per month and more for servers. Power management software adjusts the processor power states (P-states) to match workload requirements. It makes full use of the processor power when needed and conserves power when workloads are lighter. Some companies are shifting from desktops to laptops for their power-management capabilities.

C. Thermal Load Management

Technology compaction in data centers has increased power density and the need for efficient heat dissipation. Power use by ventilation and cooling systems is on par with that of servers. Typical strategies for thermal management are variable cooling delivery, airflow management, and raised-floor data center designs to ensure good air flow, more efficient air conditioning equipment, ambient air, liquid heat removal systems, heat recovery systems, and smart thermostats .

D. Product Desing

For example, microprocessor performance increased at approximately 50% CAGR from 1982 to 2002. However, performance increases per watt over the same period were modest. Energy use by servers continued to rise relatively proportionally with the increase in installed base . The 1709PICMET 2009 Proceedings, August 2-6, Portland, Oregon USA ©

2009 PICMETAuthorized licensed use limited to: University of Pittsburgh. Downloaded on January 26, 2010 at 20:16 from IEEE Explore. Restrictions apply. Shift to multiple cores and the development of dynamic frequency and voltage scaling technologies hold great promise for reducing energy use by servers. Multiple-core microprocessors run at slower clock speeds and lower voltages than single-core processors and can better leverage memory and other architectural components to run faster while consuming less energy. Dynamic frequency and voltage scaling features enable microprocessor performance to ramp up or down to match workloads. Moving beyond microprocessors, the energy proportional computing concept takes advantage of the observation that servers consume relatively more energy at low levels of efficiency than at peak levels.

Therefore, the goal is to design servers that consume energy in proportion to the work performed. Since microprocessors have more quickly acquired energy-saving capabilities, it is expected that CPUs will consume relatively less energy than other components. Therefore, it will be necessary for major improvements in memory, disk drives, and other components to reduce their power usage at higher levels of utilization. Energy proportionality, which promises to double server efficiency with the potential for large energy savings for data centers, should become a primary goal for equipment designers.

E. Virtualization

Virtualization has become a primary strategy for addressing growing business computing needs. It is fundamentally about IT optimization in terms energy efficiency and cost reduction. It improves the utilization of existing IT resources while reducing energy use, capital spending and human resource costs . Data center virtualization affects four areas: server hardware and operating systems, storage, networks, and application infrastructure. For instance, virtualization enables increased server utilization by pooling applications on fewer servers. Through virtualization, data centers can support new applications while using less power, physical space, and labor. This method is especially useful for extending the life of older data centers with no space for expansion.

Virtual servers use less power and have higher levels of efficiency than standalone servers . Virtualization technology was originally developed by IBM (as CP/CMS in the 1960’s) to increase the utilization efficiency of mainframes. More recently the concept has been applied to x86 servers in data centers. With the use of a hardware platform virtualization program called a hypervisor, or virtual machine monitor (VMM), multiple operating systems can run concurrently on a host computer. The hypervisor controls access to the server’s processor and memory and enables a server to be segmented into several “virtual machines”, each with its own operating system and application. For large data centers, server usage ranges from 5-10 percent of capacity on average. With virtualization, server workloads can be increased to 50-85 percent where they can operate more energy efficiently. Less servers are needed which means smaller server footprints, lower cooling costs, less headcount, and improved manageability.

Green Networking Metrics

Power-related metrics currently dominate green computing. Several energy-efficiency related metrics have been proposed to help IT organizations understand and improve the efficiency of data centers.

A. Energy Efficiency

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 volume of kilowatts used are typically included in the baseline assessment. 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.

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 . 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.

Data Center Infrastructure Efficency (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.

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 .

Other energy efficiency benchmarks. An alternate approach to energy efficiency monitory at the data center level is to build energy efficiency into the initial design of components and systems and to adaptively manage system power consumption in response to changes in workload and environment. These benchmarks include Analysis tool, Energy Bench, Swap, Energy Star, SPECPower, and Joule Sort.

B. Environmental Impact

Carbon Footprints

Regulations to reduce green house gas emissions worldwide will likely be forthcoming soon as a carbon tax or cap and trade scheme is being considered by the U.S. government and the Intergovernmental Panel on Climate Change (IPCC). Already some businesses are requesting that their partners provide information on carbon dioxide production. One emerging strategy is to purchase electricity from renewable energy sources such as wind, solar, or hydro. Google has adopted this strategy, although the low-cost hydro energy it has tapped into has significant environmental drawbacks that offset its attractiveness long term. The key metric here is the volume of carbon dioxide that is produced by various business processes and products—the carbon footprint.

The Elements of a Green Network

There is clear evidence that working towards a green network is not only beneficial in terms of the cost savings associated with power consumption but also delivers a future proof resilient network that can sustain the risks associated with climate change.

These risks, such as flooding and limited access to power, need to be included in current business continuity and disaster recovery planning to ensure ongoing business success. The benefits of green networking can also lead to reduced operational costs, reduced power usage, increased network & power efficiency.

To obtain a full understanding of the benefits of moving towards green networking an analysis of the total cost of ownership (TCO) of the existing network needs to be completed. TCO refers to the cost of the network lifecycle, including all costs paid by operators in purchasing and deploying equipment, implementing, and maintaining until the end of life.

Area's involved in the analysis are:

- Building infrastructure
- Server, computing and storage
- LAN
- WAN

Green competitiveness can be interpreted in terms of lifecycle, dividing resources consumed in various stages of the period into four categories: materials, manpower, energy, and land, and thus deriving a sound framework for the green evaluation system.

- Materials: physical resources purchased by operators for network operation, such as equipment and supplementary engineering materials.
- Manpower: labour cost spent by operators in the process of network development and operation.
- Floorspace: floorspace resources used by operators to place network equipment.

- Energy: energy consumed by operators' networks, mainly referring to power and cooling resources.

All this information already exists within the customers financial and management systems. we will complement this with technical and structural analyzes:

- Number and class of elements
- Systems and network structure (centralized, distributed, ...)
- Efficiency and pollution value of equipment used
- Operational and design process

Review device hardware configuration, where we will analyze power consumed versus resource (KPBP,Mbyte,MIPS), this eventually supported by online or offline measurements of our analyzers

IV. DESIGN PRINCIPLES FOR EFFICIENT NETWORK ARCHITECTURES

Green networking means efficient networking. Efficient network design combines improvements in consumption and consolidation for increased manageability and lower lifecycle cost. In this tip, learn more about the specifics of efficient network design you can implement to keep your network green, as well as avoiding wastes of bandwidth, power and budget.

Efficiency is a broad term, especially in network architectures, but there are several key elements:

- Consumption
- Consolidation
- Manageability
- Lifecycle cost

Each of these elements is related, and their synergies create the semblance of a "total system."

The following figure illustrates these relationships:

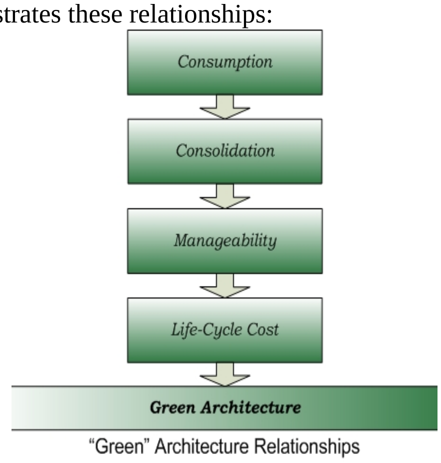

Figure 1: Green Elements

Efficiency Factor	Industry Description	Additional Factors	Design Principle
Consolidation	Server virtualization	FW virtualization VSANs Chassis-based elements w/ modular services WDM technology	If tested, stable and viable, consolidate network infrastructure into a single element. Keep the design simple and allow the technology the chance to meet your requirements.
Consumption	Power (KW) usage and space	Port density Media type (Fiber vs. Copper) Wireless capability Redundancy overkill GBICs vs. SFPs vs. WDM	Design around technologies which use the latest in energy-efficient components. Focus on simplified yet scalable designs which utilize multi-service network elements that don't need additional equipment to take advantage of new features

Figure 2: Consolidation Design Principles

There is a reason each of these elements is portrayed with a shading effect. The purpose is to show that there are in fact different shades of green, and though it may be possible to create a design that encompasses all of these factors, benefits can result from focusing on just one.

From a design perspective, there are really two elements that can be thought of as inputs to network design:

- Consumption
- Consolidation

"Consumption" is the broadest of terms used most often to describe the power and space usage of network elements such as servers, routers, switches, firewalls and SANs. There are, however, other points that can be related to this term, but they aren't as easy to differentiate. This table outlines the common and alternative thinking points in terms of consumption.

Consolidation is a distinct design option that can mitigate your consumption issues and provide an avenue for increased manageability -- and subsequently decrease your cost of support. Here are a couple of technologies that consolidate infrastructure:

- Virtualization (includes server, firewall, SAN, routers, switches, desktops)
- Chassis-based installation (FWSM, WSM, RSM, VPNSM, etc.) We will note that there are caveats to this approach that shouldn't be taken lightly Notably:
- Security
- Scalability
- Implementation time

The true trick to "getting green" is applying the principles without sacrificing these factors, or you risk losing the gains forged within the design itself.

V. THE FACTS

1. Most desktop PCs now ship with a 1 Gbps interface, an increase from the 100 Mbps interface that was standard a few years ago.
2. A typical 1 Gbps interface consumes roughly 2W more than a 100 Mbps interface.
3. Ethernet interfaces continue to consume power at this level even when the network is idle.
4. The 1 Gbps interface in the switch port at the other end of the link also consumes 2W more than a 100 Mbps interface.
5. The upgrade from 100 Mbps to 1 Gbps thus consumes an extra 4W for each connected PC.
6. The IEEE project task force has proposed that interfaces reduce speed while idle or lightly loaded. Proposed is a method by which interfaces negotiate a speed reduction when load drops and an increase when load increases.

VI. CONCLUSION

This paper offered a review of current thinking and suggested factors that should be considered for a sustainable IT strategy. Future research should address the relationship between customer value, business value, and societal value and how sustainable IT strategies will impact each. It would seem that these concepts should be mutually supportive.

However, many business professionals view them to be at odds with each other, or at least to involve tradeoffs that may not always be beneficial for the company. More research is needed to fully understand the market impact of a sustainable IT services strategy. Beyond cost savings are there benefits from sustainability oriented business strategies that customers are willing to pay for? Does sustainability for IT services create competitive advantage? Finally, a model for the development and implementation of sustainable IT services needs to be developed. This model will likely involve the integration of the IT organization's sustainability initiatives with the enterprise-level model and throughout the corporate ecosystem.

REFERENCES

1. Baldwin, E. and M. Curley, Managing IT Innovation for Business Value, Intel Press, 2007.
2. Bailey, M., M. Eastwood, T. Grieser, L. Borovick, V. Turner, and R.C.Gray, Special Study: Data Center of the Future, New York: IDC#06C4799, 2007.
3. Barroso, L. A. and U. Hölze, "The Case for Energy-Proportional Computing," IEEE Computer Society, December, pp. 33-37, 2007.
4. Barreto, L., H. Anderson, A. Anglin, and C. Tomovic, "Product Lifecycle Management in Support of Green Manufacturing: Assessing the Challenges of Global Climate Change," Proceedings of the International Conference on Comprehensive Product Realization 2007(ICCPR2007), June 18-20, Beijing, 2007.