To Be Presented at CMG 2008, December 7-12, 2008

Green Capacity Planning: Theory and Practice

Amy Spellmann
Optimal Innovations
amy@optimalinnovations.com

Richard Gimarc
HyPerformix, Inc.
rgimarc@hyperformix.com

Charles Gimarc, Ph.D.
LSI Corporation
Charles.Gimarc@lsi.com

Green Capacity Planning is a holistic planning approach for today's data centers that incorporates environmental considerations into traditional computer and network capacity planning. This paper presents the methodology, terminology, and practical application of the approach with a case study that evaluates the capacity, storage, and energy footprint of an existing eCommerce system. The case study quantifies the use of virtualization to reduce energy consumption as the system is scaled to meet the demands of a growing business.

1 Introduction

Managing today's data centers demands consideration of energy use in addition to the traditional goals of achieving Service Level Agreements (SLAs) for availability and performance. IT budgets are sky-rocketing due to increasing costs for electricity. At the same time, data center energy requirements are exceeding available power and there is growing concern about the impact on our limited natural resources. As a solution, Green Capacity Planning extends the established methods and best practices for mainframe, server and network capacity planning with an ***energy footprint projection***. This is a natural extension that views energy as a new resource for capacity planners to include in their measurement, analysis, predictions and recommendations. Green Capacity Planning does not alter the capacity planner's goal of determining a cost-effective way to deal with change; it simply adds our precious natural resources as a new ingredient.

Today's data centers have limited electrical and space capacity, therefore, it is essential to determine when these limits will be reached. Green Capacity Planning provides the means for:

- Predicting when the availability of these finite physical resources will be exceeded
- Evaluating cost-effective alternatives to postpone the need for data center expansion or construction

In this paper, we introduce Green Capacity Planning, a new approach for capacity management that includes optimization of costs for energy and computing equipment as well as server capacity and response times (performance). The goal is to provide the most cost-effective solutions that meet business SLAs. First, we describe the methodology and calculation of the energy footprint projection. Then, the methodology and value of this new approach are illustrated in a case study of a growing eCommerce system. The case study includes the use of virtualization as a mechanism for reducing the energy footprint of the system. While virtualization is certainly a valid choice for consolidating servers and reducing energy consumption, it is not the cure-all for every situation [MOLL2006] [NORT2007]. Virtualization is a good choice for our case study from a cost and energy use standpoint. However, the key, as always, is to ***model*** the virtualized environment, determine SLA achievability, scalability, and project energy footprints ***prior to*** acquisition, implementation and deployment. Since Green Capacity Planning includes more interacting factors than traditional capacity planning, evaluation and comparison of the larger number of optimization alternatives requires a more disciplined approach.

To set the stage, consider a few factoids from industry experts:

- 48% of IT budgets are spent on energy [APC2008].
- By 2010, IDC expects that for every dollar of new server spending, an additional $0.70 will be needed for power and cooling [IDC2008].
- 2% of the world's carbon dioxide emissions come from IT equipment--the same amount of pollution as the airline industry creates [WALL2007].
- "With the advent of high density computer equipment such as blade servers, many data centers have maxed out their power and cooling capacity" [GART2008].
- New data center construction projects cost $250M and up [TECH2006]. Google spent $600M each for four new data centers in 2007 [DCK2008].

These are staggering figures, indicating huge opportunities for businesses to reduce energy costs and be more responsible citizens. As capacity planners, we have the opportunity to prove our value: impact key business decisions through the delivery of actionable metrics, projections, and information that support Green IT and energy cost-cutting initiatives. At a global level, we help our planet, assisting in the effort to preserve our natural resources.

Most IT organizations have little or no responsibility for energy costs. Facilities Management is responsible for the power to the data center, but has little insight into the use of IT equipment. An InformationWeek survey confirms that almost no one in the IT organization is compensated based on saving energy, and only 22% of IT shops are responsible for managing power consumption [INFO2007]. Green Capacity Planning bridges the gap between IT and Facilities to enable better planning for data center efficiency. As stated in an InformationWeek review [INFO2008] of a recent McKinsey report [MCK2008]: "… there's too little good demand and capacity planning within and across IT, business and facilities management functions . . . Most organizations could double their efficiency by 2012 if they [implemented best practices]."

Clearly there is a compelling need for Green Capacity Planning. How do we do it? Section 2 describes the methodology and terminology of Green Capacity Planning: the required metrics, how they are interpreted, and how traditional capacity planning is extended to incorporate an energy footprint projection.

The case study presented in Section 3 demonstrates the application of Green Capacity Planning on an eCommerce system:

- Predicting system capacity
- Projecting the energy footprint
- Evaluating virtualization as a means for reducing the energy footprint
- Calculating cost-effectiveness of the solution
- Comparing energy footprints and costs

Green Capacity Planning is straightforward, based on processes that can be standardized and replicated across settings. Instead of overbuilding using lore-based rules of thumb; it is time to lower costs, reduce energy consumption and conserve natural resources with best practices and improved information for IT, Facilities and ultimately, the business.

2 Methodology & Terminology

Capacity planning practices exist in many businesses today. Additionally, there has been much written about the subject. A small sampling includes [MENA2002] [GUNT2007] [DOMA2002] [WAGO2001] [MART2001]. A good working definition of capacity planning was offered by Menasce [MENA2002]:

> *"Capacity planning is the process of predicting when future load levels will saturate the system and determining the most cost-effective way of delaying system saturation as much as possible."*

In general, the various capacity planning methodologies and processes are fairly consistent. Green Capacity Planning can be efficiently incorporated; we are simply adding environmental factors to the analysis and decision making.

We propose the following definition of Green Capacity Planning, building on Menasce's traditional one:

> *"Green Capacity Planning is the process of predicting when future business demand will exceed the availability of IT equipment, energy and space in the data center and then determining the most cost-effective way to meet SLAs and delay saturation."*

The next section describes the new combined methodology.

2.1 Methodology

Typical capacity planning methodology includes the steps shown in Figure 1:

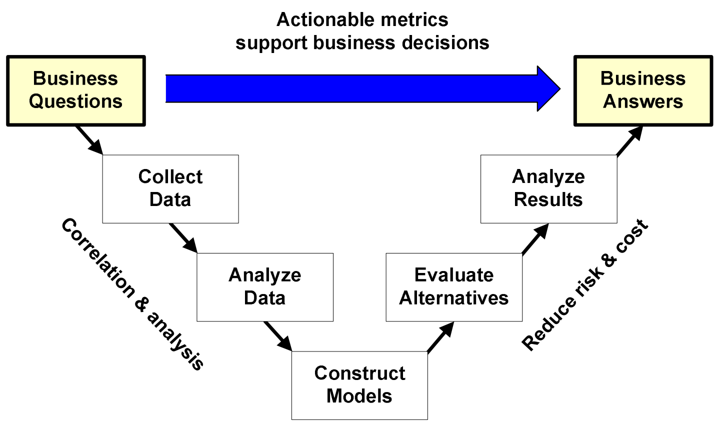

Figure 1. Typical Capacity Planning Methodology

The process starts and ends with the business. The initial business questions direct and focus the effort. The repeatable process shown in Figure 1 ultimately provides answers back to the business. Through the data collection and analysis steps, an application profile is developed that describes the performance characteristics of the system under study. Next, models are constructed (based on the application profile) that are used to predict when the system will reach capacity limits and to evaluate alternatives that satisfy the business requirements.

This same set of steps applies to Green Capacity Planning, the primary difference lies in the data collection, modeling and analysis steps: we need to include a new resource, namely power consumption. Historically, Facilities has had the primary responsibility for energy and electricity requirements in the data center. In today's environment, IT must partner with Facilities to effectively optimize data center capacity and minimize energy consumption. Facilities plays a role in Green Capacity Planning by assisting with data collection and analysis of power consumption metrics.

Data collection for traditional capacity planning includes the following metrics:

- Workload volumes: current and projected
- IT infrastructure equipment: current and proposed
- Resource consumption: CPU, memory, I/O, and network

Green Capacity Planning extends data collection to include power consumption of the system components (current and proposed) as well as the associated site energy usage to evaluate and project energy footprints.

There are a variety of methods for collecting power consumption metrics: benchmarks, vendor specifications, and real time measurements. ***Measurement of the energy footprint provides the best insight since power requirements vary with different applications and workloads.*** In general, energy consumption is highest during peak hours and lowest during idle times. Vendor specifications and benchmarks are useful for predicting power consumption, but energy footprint projections should always be validated with measurements. As with traditional capacity planning, monitoring system resource usage should continue on an ongoing basis. Facilities joins the process as the expert and primary stakeholder in power management for the data center.

Once the baseline application profile and energy footprint are developed, modeling continues along an iterative process. Predictive models are developed to determine infrastructure changes required to meet the increasing workloads. Once a configuration is identified that can support the projected load, the energy footprint can then be computed. If the resulting footprint meets business objectives (within energy and budget constraints), then the process is complete; otherwise, identify alternative configurations or components that may reduce the footprint, and repeat the modeling and analysis steps.

The Green Capacity Planning iterative approach is shown in Figure 2:

1. Infrastructure capacity and workload growth analysis – What infrastructure is required to support the projected workload?
2. Energy footprint projection – What is the infrastructure's energy footprint?
3. Cost calculation; infrastructure and energy – What is the cost for IT equipment and infrastructure energy?
4. Optimization (iterate over steps 1-3) – Does the infrastructure cost and energy footprint satisfy our business requirements?
5. Communication (business solutions) – What solution do we recommend to the business?

Communication is critical to the process as results must be presented to Facilities, IT and the business. Good business decisions depend on standardized metrics and actionable information to support the collaboration between these groups. Viable system configuration alternatives and corresponding energy footprints provide insight into cost, efficiency and scalability. As with traditional capacity planning, it is essential to continuously report the cost-effectiveness and value of Green Capacity Planning using quantifiable results.

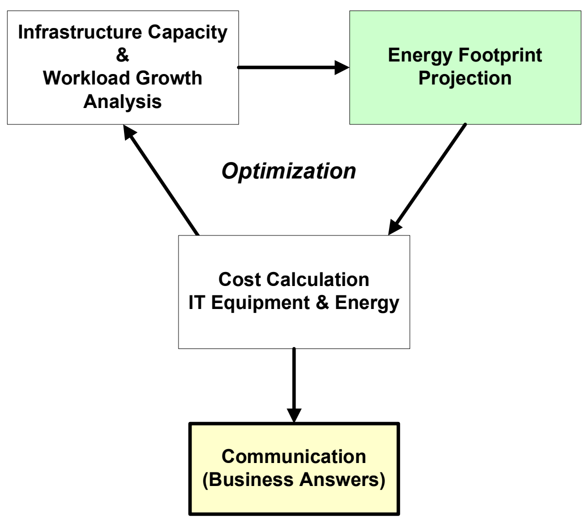

Figure 2. Iterative Approach to Green Capacity Planning

2.2 Energy Footprint Projection

The key to implementing Green Capacity Planning is the energy footprint projection. The energy footprint is a data center-level metric that describes the amount of energy used to support a system. The metric includes both the IT equipment (e.g., servers) and associated site infrastructure (e.g., cooling, power distribution and lighting). The question this section will answer is: How do we determine the amount of energy consumed (present) and required (future) by a system?

Before describing an energy footprint, we need to review the relationship between power, time and energy. Power and time are the two primary factors that determine the energy requirements of an electrical apparatus (for example, a server). Power, expressed in terms of watts (W), describes the quantity of energy used at a point in time. Energy, expressed in kilowatt hours (kWh), relates to the power used by the apparatus over a specified period of time. As the wattage increases, so does its energy consumption. You can decrease your energy consumption by reducing the amount of time you are using the apparatus. Alternatively, you can use a more energy efficient apparatus with a smaller power specification (W).

The two major consumers of energy in a data center are illustrated in Figure 3 [ACC2008]:

- *IT equipment* (servers, storage, networks)
- *Site infrastructure* (cooling, fans, lighting, power distribution units, etc.)

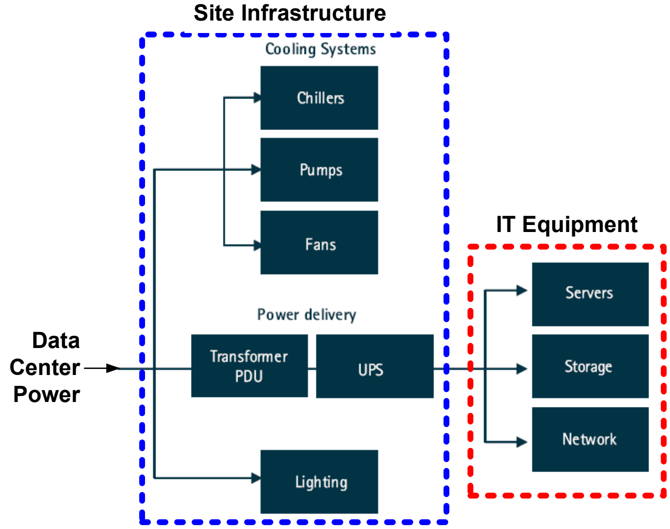

Figure 3. Data Center Power Consumers

The energy footprint describes the amount of energy consumed by IT equipment and its supporting site infrastructure for a given period of time (e.g., one month). Developing an energy footprint is a 3-step process:

1. Determine the power (W) required by each IT component.
2. Determine the power (W) used by the supporting site infrastructure.
3. Estimate time-based energy usage (kWh) by scaling the total power used by hours.

This 3-step process is summarized as follows:

$$\text{efp(hours)} = \left[\frac{\left(\sum \text{IT Power}\right) + (\text{Site Power})}{1{,}000}\right] \times \left(\text{hours}\right) \text{kWh}$$

Our approach to developing an energy footprint will be to:

- *Compute* power requirements for the IT equipment.
- *Estimate* site infrastructure power required to support the IT equipment.

Based on our approach, we refine our energy footprint expression as follows:

$$\text{efp(hours)} = \text{IT\_kWh}_{\text{computed}}(\text{hours}) + \text{Site\_kWh}_{\text{estimated}}(\text{hours})$$

Estimation of site infrastructure power will be based on the Power Usage Effectiveness metric [TGG2007]. PUE relates the two major components of data center power. Note that these are exactly the same two power consumers that we include in our energy footprint.

$$\text{PUE} = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}} = \frac{\text{Site Power} + \text{IT Power}}{\text{IT Power}} \geq 1.0$$

As an example, a PUE of 2.0 implies that for every watt of power used by servers, storage and networks, there is an additional watt used by the site infrastructure to support the IT equipment. A PUE of 1.0 implies that all power consumed by the data center is used exclusively for IT equipment, no additional site infrastructure (e.g., cooling) is required. We will use PUE and IT Power to estimate the power consumed by the site infrastructure.

Using PUE to estimate site infrastructure power simplifies our energy footprint expression to:

$$\text{efp(hours)} = \text{PUE} \times \text{IT\_kWh}_{\text{computed}}\text{(hours)}$$

For this paper, we will assume an average PUE of 2.0 (supported by [KOOM2007] [EPA2007]).

IT Equipment Power Calculation - Servers: Server power usage will be based on three factors:

- Idle power consumption (active idle)
- Maximum power consumption (processor intensive)
- Average CPU utilization

In [TGG2008] the author states that server power usage scales linearly with CPU utilization. Additional verification of this statement can be made by examining SPECpower benchmark results [SPEC2008]. By assuming a linear relationship, we have the following expression to estimate server power usage:

$$\text{Power@Util\%} = (\text{MaxW} - \text{IdleW}) \times \text{Util\%} + \text{IdleW}$$

As an example, suppose we have a server that consumes 177 watts when idle and 262 when driven at maximum load. The power consumed at 70% utilization can be estimated as follows:

$$\text{Power@70\%} = (262 - 177) \times 70\% + 177 = 236.5 \text{ W}$$

The energy footprint for this server for a month is computed as:

$$\text{efp(month)} = 2.0 \times \left(\frac{236.5}{1{,}000}\right) \times \left(24 \times \frac{365}{12}\right) = 345.3 \text{ kWh}$$

This example described two levels of power utilization: processor intensive and active idle. These metrics characterize IT equipment available today. However, systems are emerging that dynamically manage power usage in response to workload, with finer granularity than "idle" and "peak" levels. This variation will occur within the server through changing fan speeds, CPU clock speeds, powering up and down cores of a multi-core processor, and turning on and off bus adapters as use demands. This dynamic power management will occur within the storage subsystem through spinning down or up individual disks or racks of disks, again as use demands. Once dynamic power management is fully utilized, the power envelope will be more complex than that used in the examples of this paper. As a result, total energy use over long time periods will be reduced. See [IBM2002] and [INTL2008] for additional information.

2.3 Energy Cost Calculation

The cost of energy for a system is based on two factors:

- Energy footprint (in terms of kWh)
- Cost of energy (dollars per kWh)

The monthly cost for energy can be computed as follows:

$$\text{EnergyCost(month)} = \text{efp(month)} \times (\text{dollars}/\text{kWh})$$

The residential cost of a kilowatt-hour in the United States varies from 5 to 16 cents, as shown in Figure 4 [EIA2008]. Commercial rates are similar, typically negotiated based on term of service (e.g., annually). For this paper, we will use the U.S. estimate for mid-tier and enterprise data centers of 6.2 cents per kWh [EPA2007].

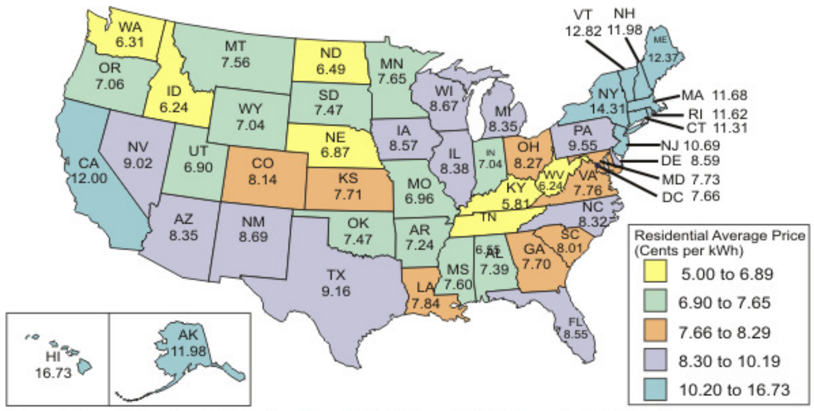

Source: Energy Information Administration, Form EIA-861, "Annual Electric Power Industry Report."

Figure 4. Residential Energy Cost per Kilowatt-Hour

As an example, consider the server used in the previous section. The monthly energy footprint for this server running at 70% utilization was computed to be 345.3 kWh. If our data center contained 200 of these servers, then the estimated monthly cost of energy would be:

$$\text{EnergyCost(month)} = (200 \times 345.3\text{ kWh}) \times (\$0.062/\text{kWh}) = \$4{,}281.72$$

We will utilize the energy footprint projection and cost calculations to compare systems in our case study.

3 Case Study – Applying Green Capacity Planning

The baseline for our case study is a published TPC-W benchmark result [TPC2002]. TPC-W specifies an eCommerce workload that simulates a Web-based retail site where emulated users can browse and buy books. The TPC-W result serving as our baseline is:

- Dell result submitted for review in May 2002 [DELL2002]
- Scale factor of 100,000 items
- 76,000 emulated browsers
- Throughput of 9,708 WIPS

The choice of this particular benchmark and result provides the following:

- A publicly available description of a non-trivial application and accompanying implementation.
- Sufficient information to create a high-level baseline model for our predictive work.
- Servers with readily available power metrics from the Dell Datacenter Capacity Planner [DELL2008].

3.1 Business Goals

Assume we have a business that has an implementation of the benchmarked system. Our marketing group is predicting a 5% monthly growth rate over the next two years as the Web site is rolled out to additional regions worldwide. The primary business goals we need to meet are:

- Determine the upgrade path to support the 2-year plan
- Ensure that we meet our SLAs:
    - Response times (equal or better than current)
    - CPU utilization (less than 70%)
- Reduce power consumption and energy costs as much as possible

To satisfy these business goals, using the Green Capacity Planning approach of Figure 2, we:

- Predict system performance, capacity and energy consumption
- Determine configuration changes required to support growth while achieving SLAs
- Compute the energy footprint
- Calculate energy and IT equipment costs
- Evaluate the use of virtualization to reduce energy consumption
- Communicate cost-effectiveness.

Cost-effectiveness will be evaluated based on the price of the IT equipment (servers and storage) required to support the growth and the total monthly energy expense (IT equipment and site).

3.2 Current System Description (TPC-W)

The baseline configuration is illustrated in Figure 5 [DELL2002]. For this study we will focus on servers and storage, no network components are included. However, our methodology applies equally to network equipment.

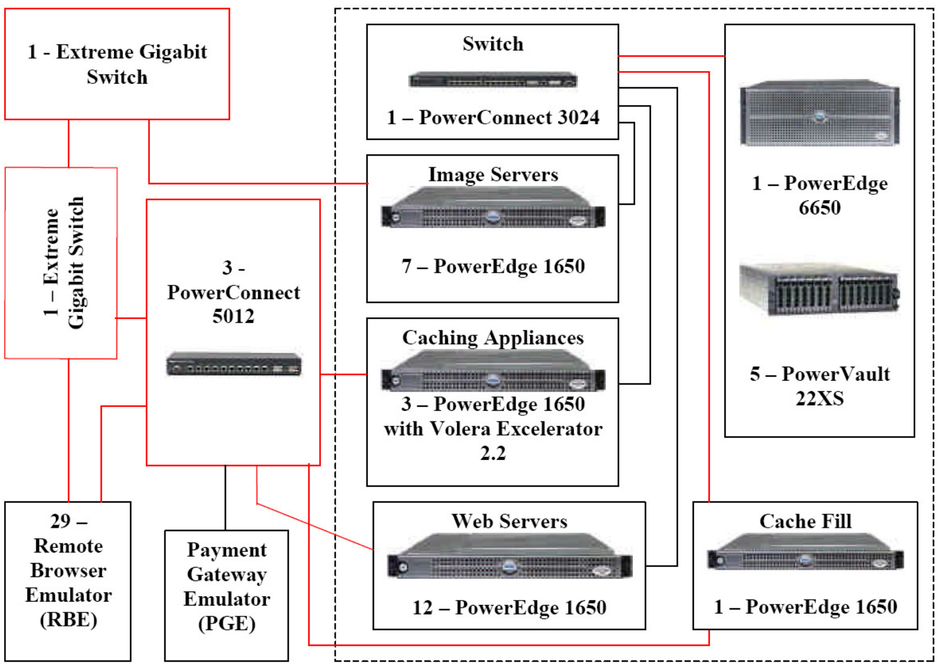

Figure 5. TPC-W Benchmark Configuration

Figure 6 lists the baseline servers and their utilization levels. Note that the table lists two utilization values:

- *Reported CPU Utilization* is the metric that was included in the benchmark's Full Disclosure Report [DELL2002]. This metric was collected and reported by Microsoft's Performance Monitor and describes the percent of the number of processors in use.
- *Capacity Utilization* is a companion metric that describes the percent of the available processing power that is being used on the server. Capacity Utilization is an output of the modeling tool that we used for scenario evaluation in this case study [HYPR2008].

Capacity Utilization is a better measure of the remaining processing capacity in a server. This metric takes into account the SMP scalability factors of chips, cores and threads to better report the actual amount of processing capacity in use. As a quick example, consider a 2-processor hyper-threaded server (four logical CPUs). If you had enough work to keep two of the four logical CPUs busy, Reported CPU Utilization would be 50%, but Capacity Utilization would be in the 80-85% range. Capacity Utilization is a better estimate of the actual usage of a server's available computational horsepower. See [COCK2006] for a discussion of the utility of different utilization metrics.

Server	Reported CPU Utilization	Capacity Utilization	Quantity	Server Type
Database Server	78.2%	88.5%	1	Dell PowerEdge 6650
Web Server(12)	75.5%	77.1%	12	Dell PowerEdge 1650
Image Server(7)	52.9%	55.3%	7	Dell PowerEdge 1650
Cache Appliance – AD	51.1%	53.6%	1	Dell PowerEdge 1650
Cache Appliance – SEARCH	47.8%	50.3%	1	Dell PowerEdge 1650
Cache Fill Server	33.8%	36.4%	1	Dell PowerEdge 1650
Cache Appliance – SEARCH2	28.2%	30.6%	1	Dell PowerEdge 1650

Figure 6. Baseline Servers and Utilization

There are two different kinds of storage in this system. First, each server has its own boot, operating system, and local storage drive pair. These drives provide temporary storage of application data and scale as servers are added. Cost and energy requirements of this storage are included in the server metrics. Second, database storage, a significant portion of the system,, is assumed to also scale at a rate of 5% per month, both in terms of capacity and aggregate throughput. Database storage is separate from any server and is tracked and upgraded as necessary to maintain the system SLA. The cost and energy requirements of storage host adapters is contained in the database server component. Figure 7 shows the initial database storage configuration.

Purpose	Enclosure #	# of Disk	RAID Level	Capacity (GB)
Log	1	10	0 / 1	135
Backup	2	7	5	393
Database	3, 4, 5	24	0	816

Figure 7. Baseline Database Storage Configuration

The baseline system [DELL2002] is already optimized for the initial workload (as this is an unstated consequence of TPC benchmarks), so we are starting with very little headroom. This means that there are some system components (e.g., the Database and Web servers) that will need to be upgraded in the first month. Now that we have collected the data for the application profile in terms of capacity and performance, we need the new ingredient, the energy footprint.

3.3 Energy Footprint

We used Dell’s Datacenter Capacity Planner (DCCP) [DELL2008] as our source of power metrics for servers and storage. Figure 8 shows a sample screen from DCCP. The Dell disclaimer, which is appropriate to note here, states that the results of the tool are “approximate and conservative”. Another key point in the disclaimer relates to workloads: they vary between older servers and the newer ones:

> *For Dell PowerEdge Servers shipping prior to June 6, 2006, the test applications used were extreme, chosen to push the system components to their maximum possible power levels. … for PowerEdge servers shipping after June 6, 2006 (PE1950, PE1955, PE2950, PE2900, PE6950, PE1435, PE860) and going forward we have adopted the industry standard SPECjbb2005 benchmark for baseline power measurements to provide more useful power predictions across the entire spectrum of configurations.*

Therefore, it is important to note that our calculations and comparisons are for early planning purposes only and must be verified at each step in the process. In some cases, error may be introduced into our energy estimates since we do not have an exact one-to-one match.

Ideally the energy footprint of our baseline system comes from our production measurement tools, but for this study, we used DCCP for the baseline as well as the proposed new configurations.

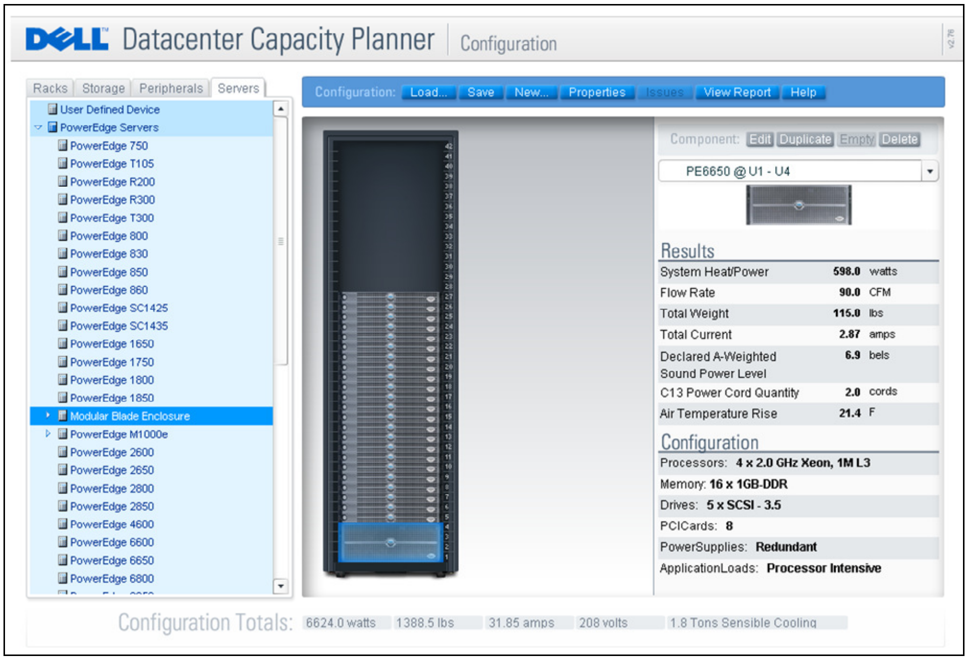

Figure 8. Dell Datacenter Capacity Planner

For the baseline system, the total watts required (sum of the individual components) range from 8,841 watts (active idle) to 11,324 watts for maximum usage. The details are shown in Figure 9.

System Component (Quantity)	Idle Power	Max Power	Total Idle Power	Total Max Power
PowerEdge 1650 (23)	177	262	4,071	6,026
PowerEdge 6650 (1)	405	598	405	598
PowerVault 220S (5)	873	940	4,365	4,700
Configuration Total			**8,841**	**11,324**

Figure 9. Baseline Power Requirements

Energy Footprint Calculation: For simplicity, we assume a utilization of 70% for computing server power consumption and the maximum power values for database storage. With these assumptions, the energy footprint for each component type in the first month in our plan (baseline) is computed as follows:

$$\text{efp(PE1650@month)} = \frac{\left((6{,}026 - 4{,}071) \times 70\%\right) + 4{,}071}{1{,}000} \times \left(24 \times \frac{365}{12}\right) = 3{,}971\,\text{kWh}$$

$$\text{efp(PE6650@month)} = \frac{\left((598 - 405) \times 70\%\right) + 405}{1{,}000} \times \left(24 \times \frac{365}{12}\right) = 394\,\text{kWh}$$

$$\text{efp(PV220S@month)} = \frac{4{,}700}{1{,}000} \times \left(24 \times \frac{365}{12}\right) = 3{,}431\,\text{kWh}$$

The total IT equipment energy footprint for the entire baseline configuration is:

$$\text{efp(IT Equipment@month)} = 7{,}796 \text{ kWh}$$

The energy footprint of the entire configuration (IT equipment plus site infrastructure) can be computed as follows:

$$\text{efp(month)} = \text{PUE} \times \text{efp(IT Equipment@month)} = 2.0 \times 7{,}796 \text{ kWh} = 15{,}592 \text{ kWh}$$

From the energy footprint, we can estimate the total monthly cost to be:

$$\text{EnergyCost(month)} = 15{,}592 \text{ kWh} \times \left(\$0.062 / \text{kWh} \right) = \$966.71$$

The total baseline monthly cost for energy (equipment and infrastructure) is $966.71.

Now that the baseline application profile is complete, the next step is to determine the upgrade path for our 2-year planning horizon.

3.4 Capacity Modeling Scenarios

To evaluate the upgrade scenarios, we used a capacity modeling tool from HyPerformix [HYPR2008]. Our goal is to determine the hardware requirements as the workload increases (5% per month) and to upgrade or add servers just in time. Upgrades are staged so that server capacity utilization never exceeds our SLA of 70%.

Although we are focused on server utilization, we also need to ensure that response time does not increase relative to the baseline system. In our case study, by constraining server utilization and always upgrading to faster devices, we assume that response times do not exceed baseline levels. Detailed analysis beyond the scope of this paper is required to determine response times of individual workloads and transactions.

We examined two capacity planning scenarios for our eCommerce system:

- *Server Upgrade* – Whenever a server reaches its utilization threshold (70%), we upgrade the server.
- *Virtualization* – Instead of upgrading individual servers, we virtualize the saturated server(s). Virtualization is used for all server upgrades except for the database.

For each scenario, we used the following guidelines:

- Minimize the number of server upgrade points (less disruptive to the business) by limiting the number of server changes to no more than one per 3-month period.
- Retire old resources; that is, replace old servers with faster and more efficient servers.
- Fully utilize new resources (only in the Virtualization scenario). For example, when we introduce a new virtual host, we use as much of that new server as soon as possible.

Figure 10 shows an overview of each scenario. Note that there are a number of configuration changes that are common to both scenarios.

Server Upgrade Scenario	Virtualization Scenario
5% monthly growth for 2 years	5% monthly growth for 2 years
70% capacity threshold	70% capacity threshold
Optimize the number of upgrades	Optimize the number of upgrades
Dell PE 2950 for the database	Dell PE 2950 for the database
Database storage upgrade to PV MD1120	Database storage upgrade to PV MD1120
Dell PowerEdge 1850 for all other servers	***Virtualize all other servers onto a pool of Dell PowerEdge 1950s*** ***Optimize use of virtual hosts***

Figure 10. Scenario Overview

3.5 Server Upgrade Scenario

The timeline of infrastructure changes is shown and described in Figures 11 and 12. We used our modeling tool to determine the infrastructure required to successfully grow the workload by 5% per month over the 2-year planning horizon. Recall that servers are upgraded when they reach 70% capacity utilization.

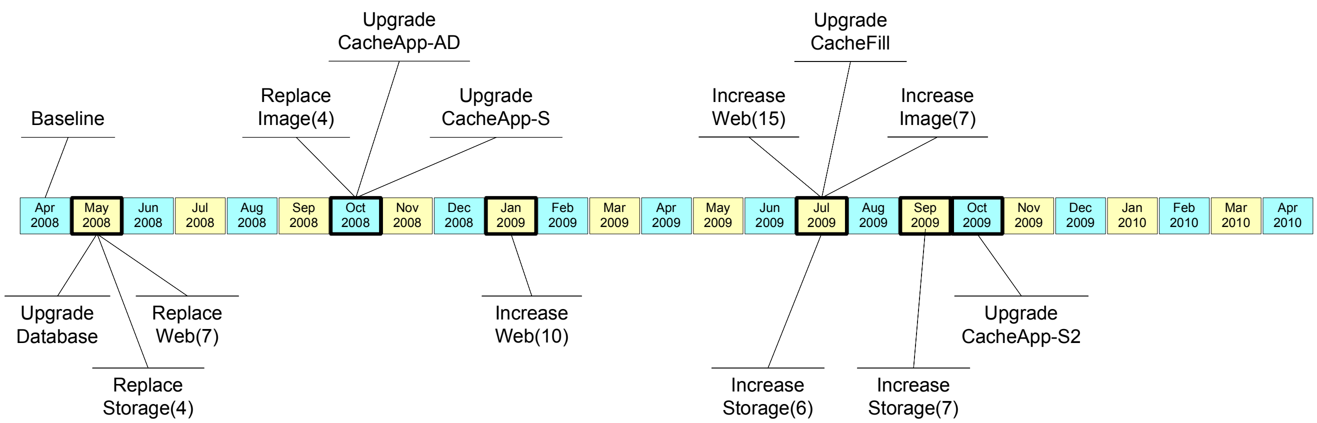

Figure 11. Server Upgrade Scenario - Timeline

Date	Configuration Change
Apr 2008	Baseline system
May 2008	Upgrade Database server to a PowerEdge 2950 1.8GHz Replace the 12 Web servers with 7 Dell PowerEdge 1850 servers Replace database storage with 4 Dell PowerVault MD1120
Oct 2008	Replace 7 Image servers with 4 Dell PowerEdge 1850 servers Upgrade Cache Appliance-AD to a Dell PowerEdge 1850 Upgrade Cache Appliance-S to a Dell PowerEdge 1850
Jan 2009	Increase Web servers from 7 to 10 PowerEdge 1850 servers
Jul 2009	Increase Web servers from 10 to 15 PowerEdge 1850 servers Upgrade Cache Fill server to a Dell PowerEdge 1850 Increase Image servers from 4 to 7 PowerEdge 1850 servers Increase database storage from 4 to 6 PowerVaults
Sep 2009	Increase database storage from 6 to 7 PowerVaults
Oct 2009	Upgrade Cache Appliance-S2 to a Dell PowerEdge 1850

Figure 12. Server Upgrade Scenario - Configurations

The first upgrade point occurs in the second month (May 2008). Since the database is saturated, we chose to replace the database server with a newer PowerEdge 2950. We also replaced the original Ultra 320 SCSI storage subsystem with newer SAS (Serial Attached SCSI) technology. This complete replacement of storage was done to:

- Reduce power by changing from 3½" to 2½" SFF (small form factor) disks
- Move to current SAS technology that will allow scaling for 2 years
- More than double the storage density by moving from 12 disks in a 3U space to 24 disks in a 2U space
- Achieve higher performance to meet the demands of the system after 2 years of workload growth

Figure 13 graphically shows the modeling results. Each line represents the capacity utilization of a server (tier). The chart shows a gradual increase in server utilization up to the upgrade threshold (approximately 70%) and then the utilization decreases after the upgrade. Note that there are five server upgrade points over the 2-year planning period. Single server tiers show a marked decrease in utilization, dropping from around 70% to the 20-30% range. We sized multiple server tiers to have less of a drop in utilization at the upgrade points; the Web and Image tiers show utilizations ranging from 50-70%.

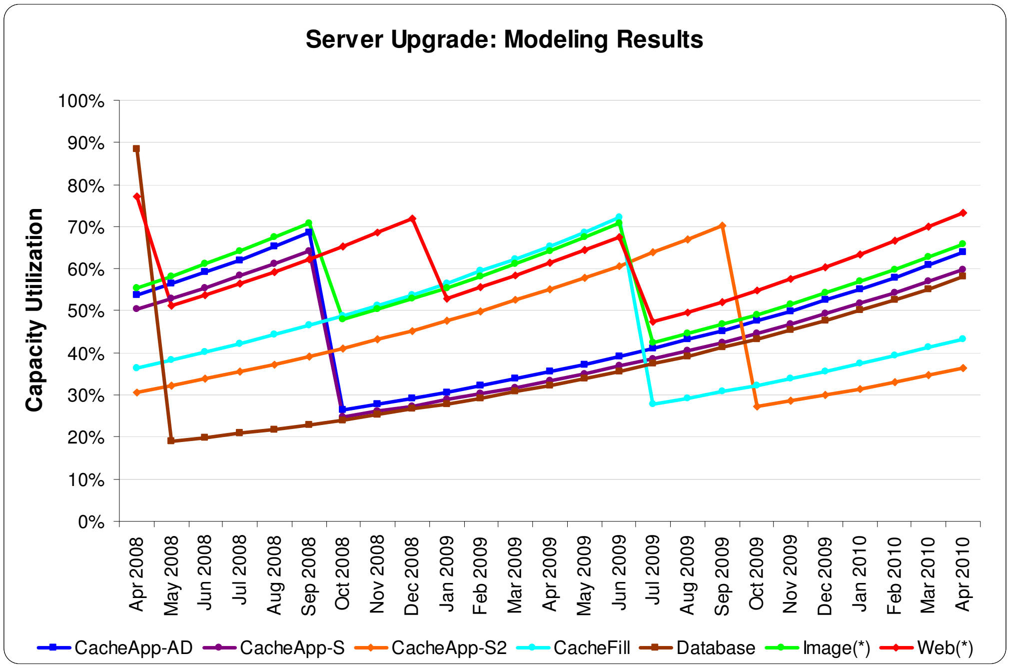

Figure 13. Server Upgrade Scenario - Modeling Results

The monthly energy footprint for the Server Upgrade scenario is shown in Figure 14. The only significant decrease in the energy footprint is in the second planning point (May 2008); this is where we did an entire replacement of the database storage subsystem. We do not see much of an energy savings with the servers. As expected, the Web and Image tiers have the largest footprints since there are multiple servers at each of those tiers. Most importantly, this chart shows that we are not achieving the goal to reduce energy consumption.

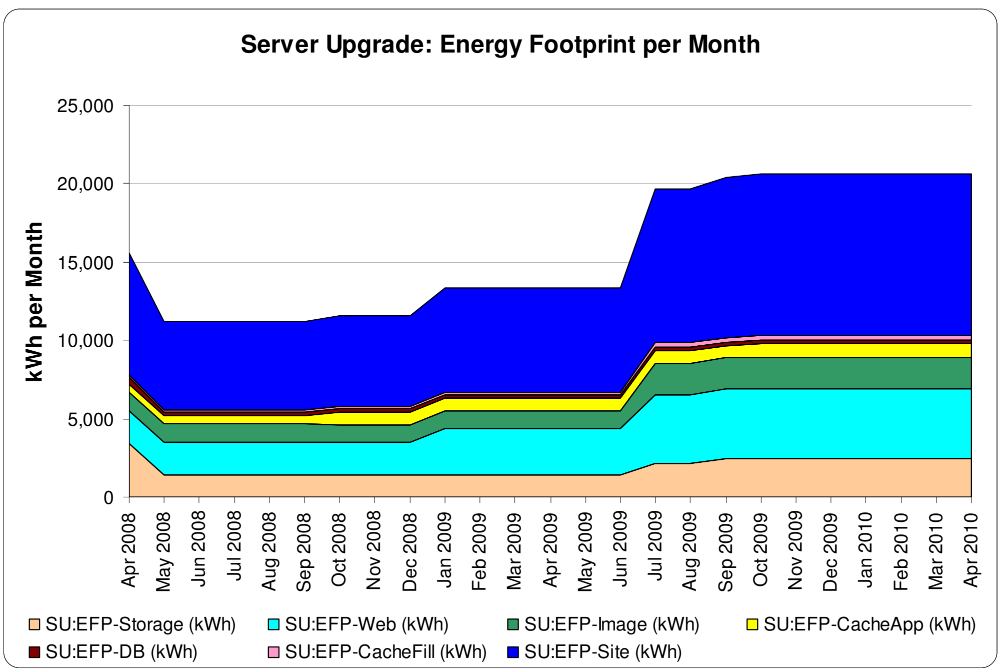

Figure 14. Server Upgrade Scenario - Monthly Energy Footprint

In an effort to reduce the energy footprint, we now evaluate the use of virtualization to decrease the number of servers and exploit new technology available in power efficient servers.

3.6 Virtualization Scenario

In this scenario, we used virtualization to reduce the count and energy footprint of the Web and Image server tiers. Recall that the Web tier initially has 12 servers, and there are 7 on the Image tier. These two tiers will be moved onto a pool of balanced virtual hosts. Our model represents VMware as the virtualization environment [VMWA2008].

The timeline of infrastructure changes is shown and described in Figures 15 and 16. In this scenario, virtualization is introduced whenever we have to upgrade the Web and/or Image tiers.

Recall that one of our guidelines was to fully utilize a new resource when available. This can be seen in May 2008. In addition to virtualizing the 12 Web servers, we also virtualize the four “cache” servers onto the new virtual hosts. It makes sense to utilize this new resource as much and as soon as possible.

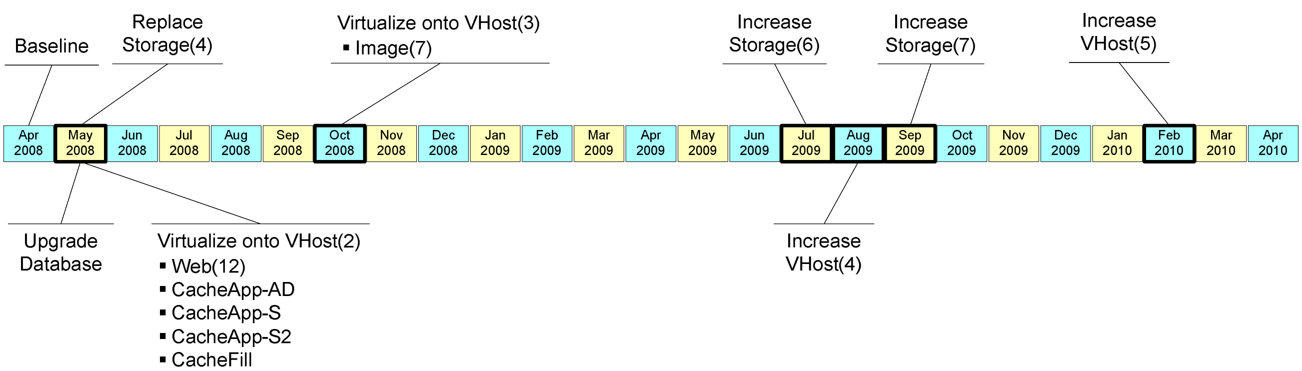

Figure 15. Virtualization Scenario - Timeline

Date	Configuration Change
Apr 2008	Baseline system
May 2008	Upgrade Database server to a PowerEdge 2950 1.8GHz Virtualize all 12 Web servers and all 4 Cache servers across 2 PowerEdge 1950 III servers (VHosts) Replace database storage with 4 Dell PowerVault MD1120
Oct 2008	Increase VHosts from 2 to 3 PowerEdge 1950 III, and virtualize all 7 Image servers across the pool of VHosts
Jul 2009	Increase database storage from 4 to 6 PowerVaults
Aug 2009	Increase VHosts from 3 to 4 PowerEdge 1950 III and rebalance the Web, Cache and Image servers across the VHosts
Sep 2009	Increase database storage from 6 to 7 PowerVaults
Feb 2010	Increase VHosts from 4 to 5 PowerEdge 1950 III and rebalance across the VHosts

Figure 16. Virtualization Scenario - Configurations

Figure 17 shows the modeling results from the Virtualization scenario. The green vertical bars indicate the number of virtual hosts (VHosts). The overlaid capacity utilization lines for the servers show that as servers are virtualized, their utilization lines merge into the VHost line. The VHost line represents the overall average utilization of the pool of VHost servers. As we introduce a new VHost and rebalance the load, you see a corresponding drop in capacity utilization.

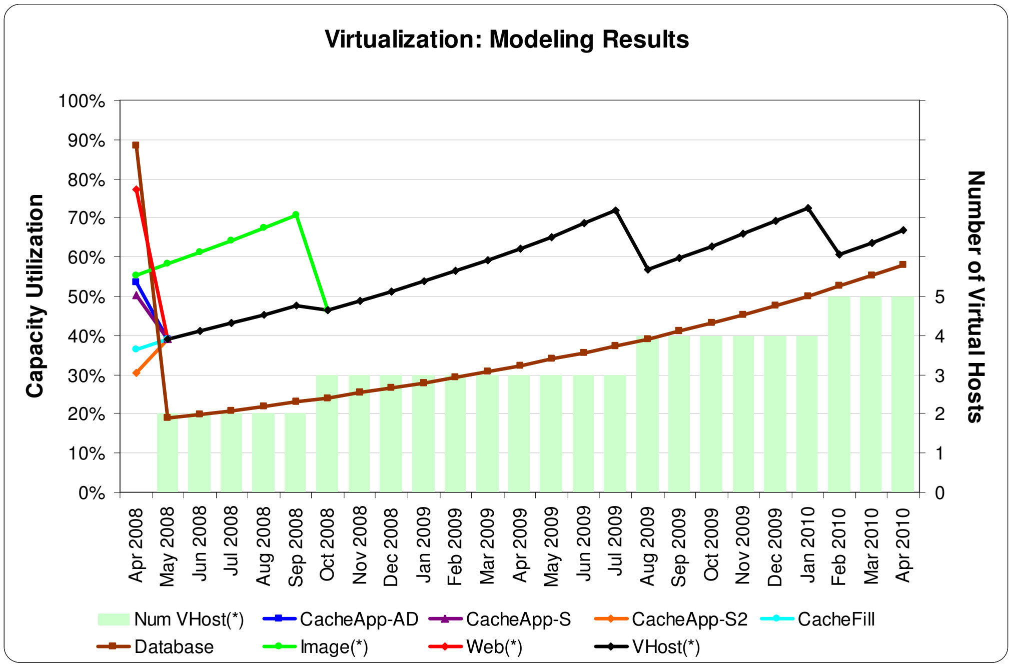

Figure 17. Virtualization Scenario - Modeling Results

The monthly energy footprint for the Virtualization scenario is shown in Figure 18. These results show a significant reduction in power consumption (compared to the Server Upgrade scenario in Figure 14). Even after the required server and storage upgrades for the 2-year growth we are using approximately half the energy compared to the baseline. We are increasing the workload volume by 5% per month; at the end of 2 years the expanded system has 3.2 times the capacity of the baseline, at approximately half the energy utilization.

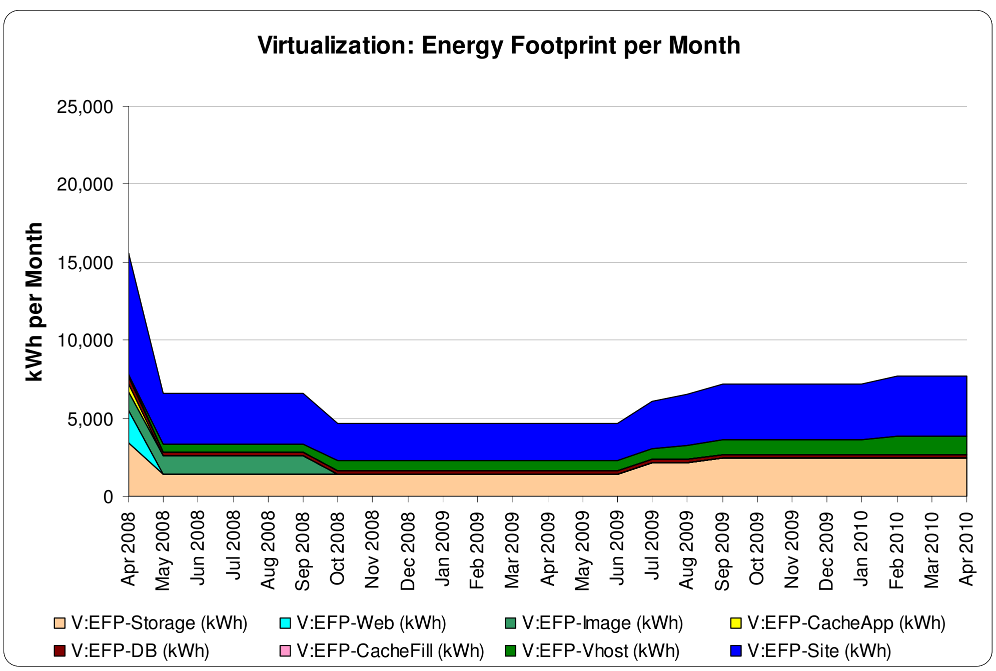

Figure 18. Virtualization Scenario - Monthly Energy Footprint

Next, we compare the two sets of scenario results, side-by-side, including the energy footprints, energy and IT infrastructure costs and space.

4 Comparison: Server Upgrade vs. Virtualization

The results of the two scenarios show that the virtualization approach has a significantly smaller energy footprint, and hence, reduced energy cost. This section will compare and contrast the trenchant observations made when comparing the results from the two scenarios.

4.1 Monthly Energy Footprint

The energy footprint of the virtualized environment is about 40% less than the non-virtualized case (see Figure 19). But more importantly, the energy footprint of the virtualized configuration at the end of two years is approximately half of our original baseline configuration. Figure 19 is a simplified view of the stacked component charts from Figure 14 and 18.

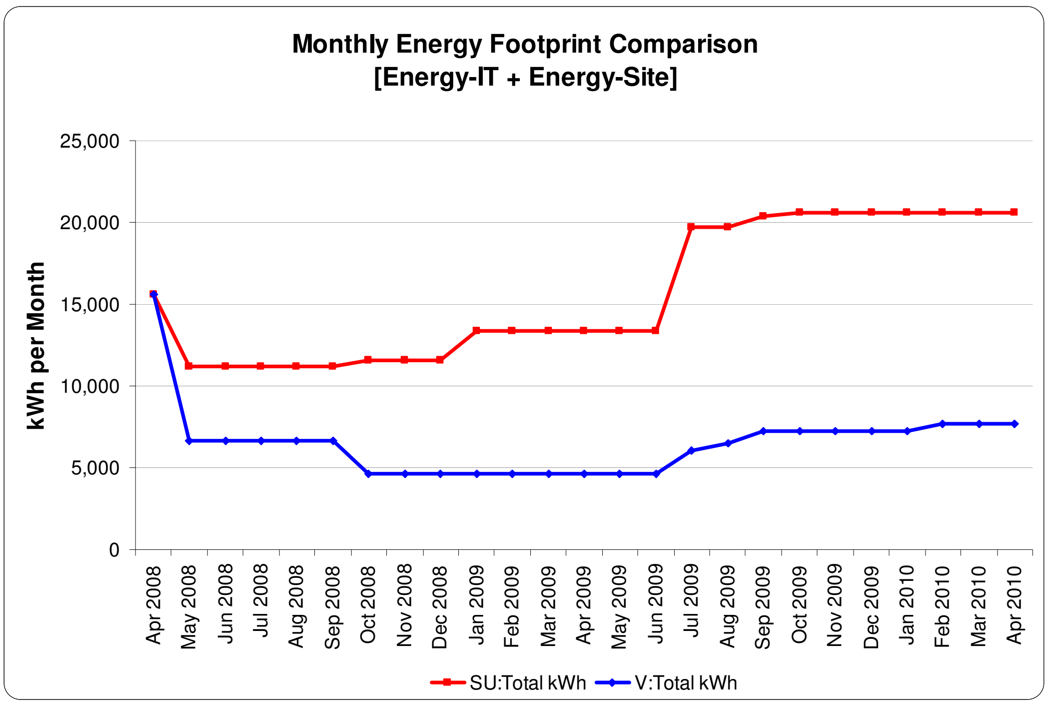

Figure 19. Comparison - Monthly Energy Footprint

4.2 Cumulative Monthly Cost

Figure 20 shows the cumulative monthly cost comparison. In this chart, cost includes both energy and new IT equipment. The two cost curves are almost identical up through September 2008, where the Server Upgrade (SU) begins to rise faster than in the Virtualization scenario (V). The reason for this increase is the six new servers purchased for the Server Upgrade scenario.

At the end of the two year plan, the cost of the Server Upgrade scenario exceeded the Virtualization scenario by approximately $20,000.

It is also interesting to note the slope of the cost lines. There were no configuration changes between February and June 2009. However, you can see that the Server Upgrade scenario curve increases at a faster rate than the Virtualization scenario curve. This is due to the higher rate of power consumption by the servers in the Server Upgrade scenario.

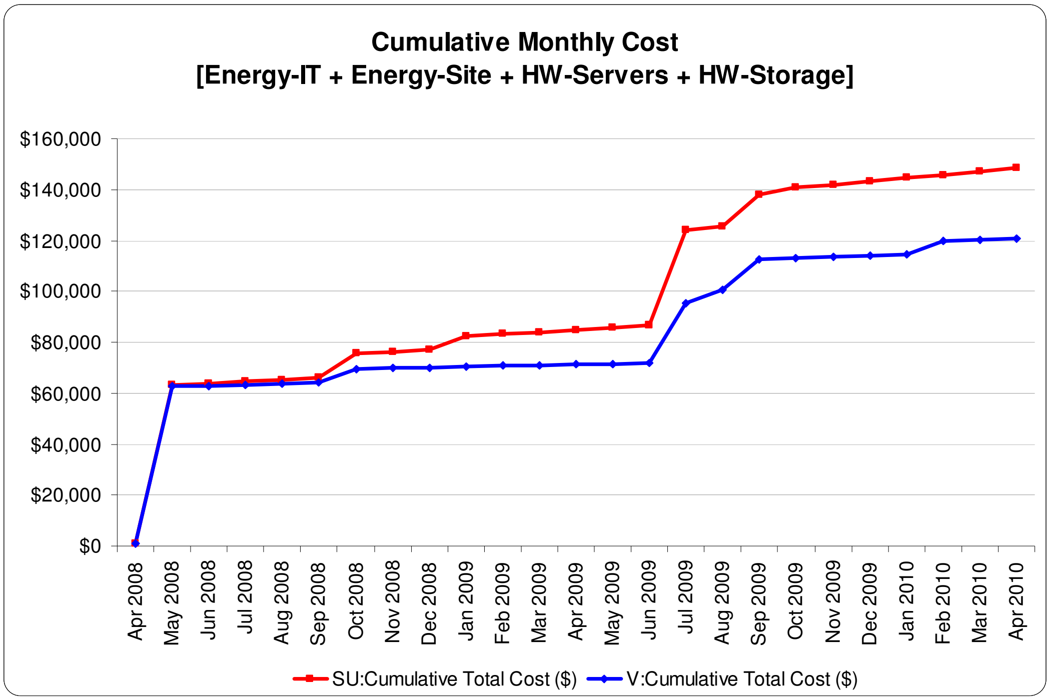

Figure 20. Comparison - Cumulative Monthly

Figures 21 and 22 show a more detailed view of the cumulative monthly cost. By examining the shapes in these two charts you can see that the IT equipment (servers and storage) is the largest cost component in both scenarios. In fact, the storage costs are identical (since the storage upgrade is the same in both scenarios). It is also clear that the cost of the pool of virtual host servers is less than the individual server upgrades used in the first scenario.

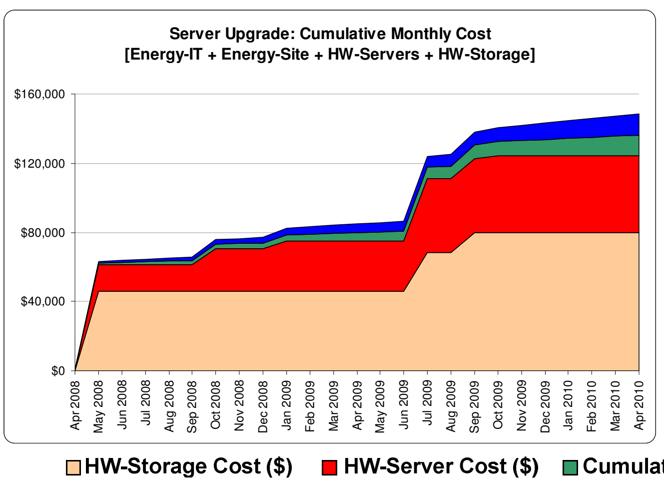

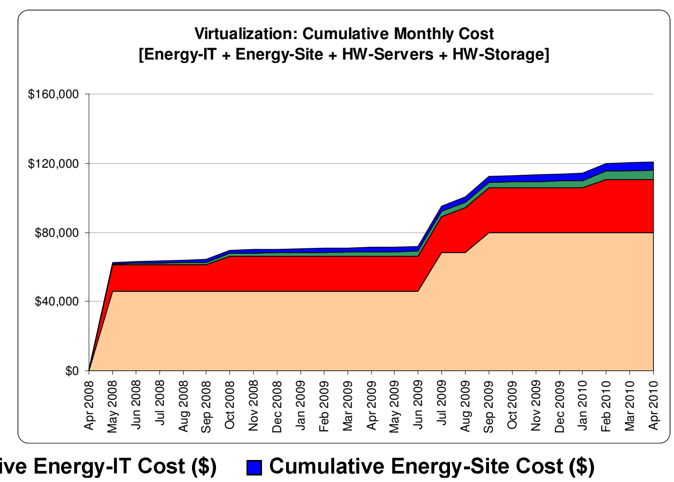

Figure 21. Server Upgrade Scenario - Cumulative Monthly Cost

Figure 22. Virtualization Scenario - Cumulative Monthly Cost

4.3 Server Count

Figure 23 shows a comparison of the number of physical servers for the baseline (April 2008) and at the end of the 2-year plan for both scenarios.

- The baseline configuration has a total of 24 servers; one database server (PowerEdge 6650) and 23 PowerEdge 1650s.
- The Server Upgrade scenario ended with 27 servers (3 more than the baseline). Again, there was one database server (PowerEdge 2950), and the original 1650s were replaced with 26 PowerEdge 1850s.
- The Virtualization scenario resulted in a significantly smaller configuration: 6 physical servers. There was a final pool of 5 virtual hosts (PowerEdge 1950s) in addition to the database server.

Both of these systems achieved the same SLA at the end of the study period, but with vastly different IT equipment requirements.

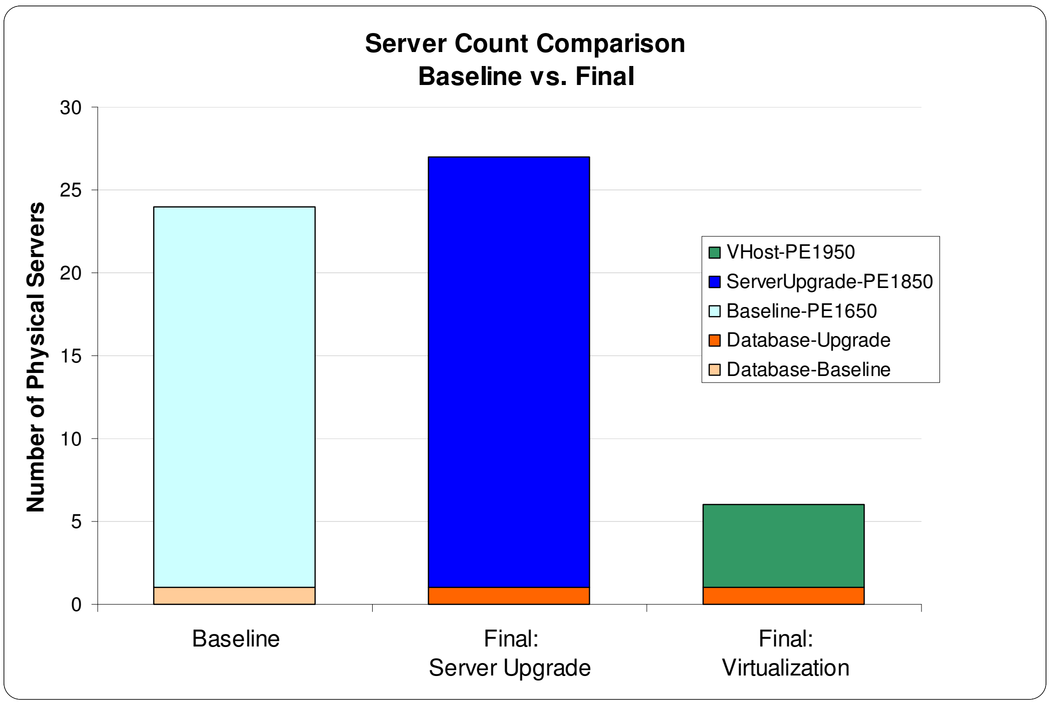

Figure 23. Comparison – Physical Server Count

4.4 Total Rack Space

Finally, we come to our space analysis results: total rack space. Figure 24 shows the total rack space required per month. As expected, the Virtualization scenario requires less rack space than the Server Upgrade scenario (or non-virtualized solution) by the end of the 2-year planning period. In fact, the Virtualization Scenario requires half the rack space of the Server Upgrade scenario.

It is also interesting to note that the rack space for the Server Upgrade scenario at the end of the 2-year period is approximately the same as that required for the baseline. Over this 2-year period, computational and storage capacity of the system increased by over 3.2 times, yet we did not see a corresponding increase in the rack space. The increase in computational density over this time period is due to advances in server and storage technologies. As a server or storage tray is replaced and upgraded, we bring in new technology with higher performance in the same or smaller space.

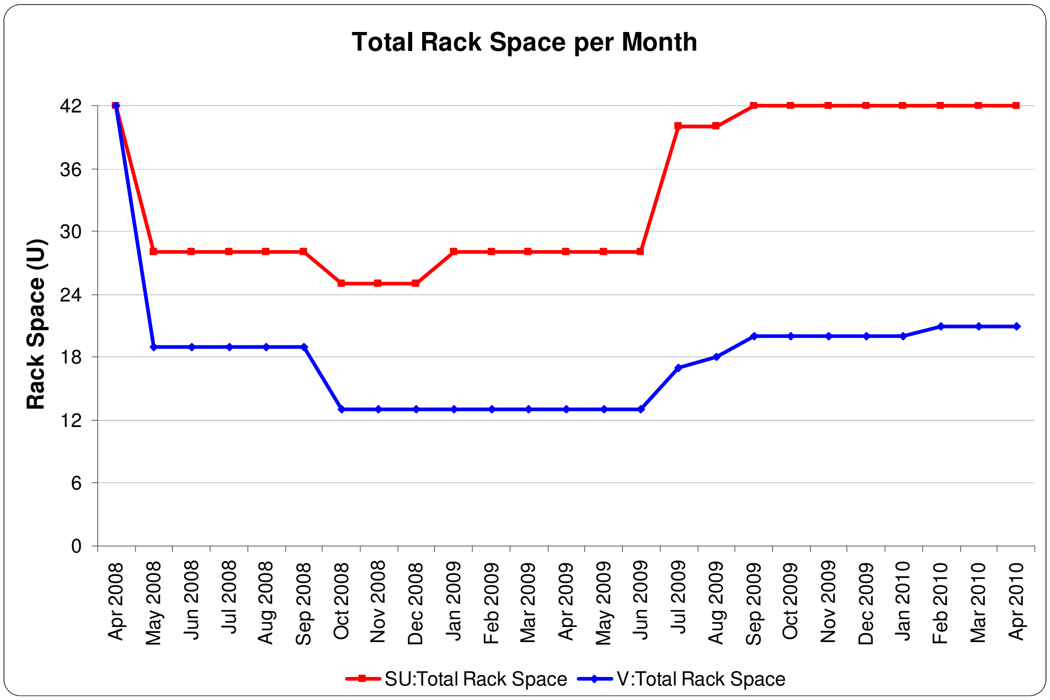

Figure 24. Comparison - Total System Rack Space Requirement

4.5 Scenario Comparison Conclusion

IT managers are concerned with adding computational capability into their data centers to satisfy the requirements of a growing business. We know that space is a limited resource in a data center; so to grow, we must process more transactions and store more data per cubic foot of space (that is, increase the computational density). Electricity is also a finite resource, since data centers have an upper limit on the amount of power they can utilize from their electric utility. In general, to effectively scale a data center of fixed size, the IT manager must find a way to make the “system” more efficient. In our case study, the virtualization path is the preferred choice for achieving the required efficiency.

5 Summary

Green Capacity Planning expands the scope of traditional capacity planning to include environmental considerations. Specifically, Green Capacity Planning adds power consumption as a new resource for consideration and analysis. The concept of an energy footprint is used to quantify the demand for power (our new resource). This enables us to develop estimates for energy consumption and cost.

The case study demonstrated how Green Capacity Planning can efficiently be integrated into the traditional capacity planning process. With the addition of power resource metrics, capacity planners can increase their value to the business by supporting and addressing today’s Green IT and cost-cutting initiatives.

Furthermore, the benefits of a standardized energy footprint metric are:

- Different proposed solutions, arrived at with different strategies, can be directly compared.
- Capacity plans and models for optimization would be expected to include an energy footprint projection.
- Estimated and actual impacts of infrastructure changes can be evaluated with respect to power consumption.

Today’s data centers have a physical facility of fixed and finite size. In addition, the electricity available to run data center equipment can only increase to the ceiling that is set by physical constraints of the building and attachment to the municipal electrical grid. Green Capacity Planning helps IT managers plan, specify, and deploy systems of increasing computational density, power and storage capacity that will meet the business' SLA and fit within the finite physical resource constraints.

> *"Sure, the Googles and Amazon.coms of the world can build state-of-the-art facilities next to dammed rivers or geothermal vents, but for the rest of us, the environmentally responsible thing to do is to squeeze every last ounce of potential out of the data centers we have."* [INFO2007]

We must strive to be environmentally responsible, cost-conscious capacity planners by incorporating ***energy footprint projections*** into our practices. Green Capacity Planning provides us with the thought process and approach for adapting to today's growing concerns about the scarcity, cost and use of our limited natural resources.

6 Acknowledgements

The authors would like to thank John Pflueger, Dell Technology Strategist and a Director of The Green Grid, for his support and guidance. His editorial and technical input contributed to the quality of our work.

7 References

[ACC2008] Accenture, "Data Center Energy Forecast Report", July 2008.

[APC2008] APC-MGE, www.apc-mge.com.

[COCK2006] Adrian Cockcroft, " Utilization is Virtually Useless as a Metric!", CMG 2006 International Conference.

[DCK2008] Data Center Knowledge, "Google Data Center FAQ, Part 2", August 26, 2008, www.datacenterknowledge.com/google-data-center-faq-part-2.

[DELL2002] TPC Benchmark W Full Disclosure Report, DELL PowerEdge 6650/1.6GHz with PowerEdge 1650/1.4GHz, May 31, 2002, www.tpc.org.

[DELL2008] DELL Datacenter Capacity Planner, www.dell.com/calc.

[DOMA2002] Bernard Domanski, "System Performance Management and Capacity Planning Tutorial", CMG 2002 International Conference.

[EIA2008] "Annual Electric Power Industry Report", Form EIA-861, Energy Information Administration, 2008, www.eia.doe.gov.

[EPA2007] "Report to Congress on Server and Data Center Energy Efficiency", U.S. Environmental Protection Agency, August 2008.

[GART2008] Gartner, Inc., www.gartner.com.

[GUNT2007] Neil J. Gunther, "Guerrilla Capacity Planning", Springer, 2007.

[HYPR2008] HyPerformix Capacity Manager, www.hyperformix.com, 2008.

[IBM2002] "Dynamic Power Management for Embedded Systems", 2002, www.research.ibm.com/arl/publications/papers/DPM_V1.1.pdf.

[IDC2008] IDC, www.idc.com.

[INFO2007] Information Week, "The Cold, Green Facts", September 3, 2007.

[INFO2008] Information Week, "McKinsey: Measure Data Center Efficiency Like Car Fuel Efficiency", April 30, 2008.

[INTL2008] " Dynamic Power Optimization for Higher Server Density Racks - A Baidu Case Study with Intel Dynamic Power Technology", 2008, http://communities.intel.com/openport/docs/DOC-1492.

[KOOM2007] Jonathan G. Koomey, "Estimating Total Power Consumption by Servers in the World", 2007, http://enterprise.amd.com/Downloads/Technology/Koomey_Estimating.pdf.

[MART2001] Fernando Martinez, "Your First Capacity Planning Process, a Survival Guide" CMG 2001 International Conference.

[MCKN2008] McKinsey and Uptime Institute Report, “Revolutionizing Data Center Efficiency”, 2008.

[MENA2002] Daniel A. Menasce and Virgilio A. F. Almeida, "Capacity Planning for Web Services", Prentice Hall PTR, 2002.

[MOLL2006] Chris Molloy, "Virtualization - Inhibitors to Server and Storage Virtualization, and How to Mitigate Them", CMG 2006 International Conference.

[NORT2007] Tim R. Norton, "The Myth of Precision Planning: Understanding Capacity in an Age of Virtual Parallelism", CMG 2007 International Conference.

[PRIN2008] Principled Technologies, Inc., www.principledtechnologies.com.

[SPEC2008] SPECpower_ssj2008 Benchmark, Standard Performance Evaluation Corporation (SPEC), 2008, www.spec.org.

[TECH2006] TechTarget, Bernard Golden, “Golden's Rules: Sun's Linux-friendly, virtual, portable data center”, October 23, 2006, searchenterpriselinux.techtarget.com/tip/0,289483,sid39_gci1225895,00.html].

[TGG2007] The Green Grid, “The Green Grid Data Center Power Efficiency Metrics: PUE and DCiE”, 2007, www.thegreengrid.org.

[TGG2008] The Green Grid, Mark Blackburn, “Five Ways to Reduce Data Center Server Power Consumption”, 2008, www.thegreengrid.org.

[TPC2002] TPC BENCHMARK W (Web Commerce) Specification, Version 1.8, Feb 2002, www.tpc.org.

[VMWA2008] VMware, www.vmware.com, 2008.

[WAGO2001] Dave Wagoner, "Practical Capacity Planning", CMG 2001 International Conference.

[WALL2007] Wall Street Journal, Monday, August 27, 2007.