the green grid™
get connected to efficient IT

USING VIRTUALIZATION TO IMPROVE DATA CENTER EFFICIENCY

EDITOR:

RICHARD TALABER, VMware

CONTRIBUTORS:

TOM BREY, IBM

LARRY LAMERS, VMware

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner. Version 1.0 Release Date: TBD

ABSTRACT

This paper outlines some of the advantages, considerations, processes, and implementation strategies needed to reduce server power consumption in a data center using virtualization techniques.

Virtualization has long held the promise of energy savings due to its server consolidation capabilities. This paper provides those responsible for reducing energy consumption with a description of the advantages and processes that should be considered when virtualizing servers in the data center.

the green grid™

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner. Version 1.0 Release Date: TBD

TABLE OF CONTENTS

1. Server Classification within a Typical Data Center....4
   - 1.1 Servers and Energy Consumption....7
   - 1.2 Accelerating Business Services....9
   - 1.3 Utilization Across Server Categories....9
   - 1.4 Server Consolidation and Virtualization....10
2. Preparing for Virtualization – Gathering Data....11
   - 2.1 Identifying the Server Population....11
   - 2.2 Determining Compute Resources Required....12
   - 2.3 Estimating Capacity Requirements....14
   - 2.4 Application Workload Characterization....16
3. Other Considerations....18
   - 3.1 Metrics and Measurements....18
   - 3.2 Establish Service Level Objectives....18
   - 3.3 Prioritization/Phases....19
   - 3.4 Technology Considerations....20
   - 3.5 Process Considerations....20
   - 3.6 Measurements of Success....21
4. Conclusion....21

the green grid™

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

INTRODUCTION

Data center IT equipment drives the amount of physical infrastructure (e.g., power distribution units, UPS, generators, cooling units) needed to maintain the availability of data center assets. Both the demand side (IT) and the supply side (physical infrastructure) of the data center consume energy.

On the IT side, combining workloads onto fewer servers can significantly improve the energy efficiency of a data center. Conversely, adding new servers into a data center without first considering the utilization of existing compute nodes will cause an unwanted and avoidable increase of data center energy consumption. This also holds true for data center physical infrastructure. "Over–sizing" of power and cooling equipment is endemic to existing data centers and represents a considerable waste of energy. Changes, either up or down, on the IT side will also have an impact on the efficiency of the physical infrastructure supporting IT.

the green grid™

The Green Grid recommends the use of Data Center Efficiency (DCE) metrics, which enable data center operators to quickly estimate the energy efficiency of their data centers, compare the results against other data centers, and determine if any energy efficiency improvements need to be made. A data center must deploy the necessary gauges, metering, and logging of data to enable the accurate and long-term measurement of data needed in order for PUE/DCiE trending and other efficiency metrics to be meaningful. Advances in defining productivity indicators[1], performance[2] and government programs[3] now are beginning to offer guidance to purchasers of computer servers. This guidance helps purchasers consider the productive work a computer server can exert for each watt of power that is consumed. Servers are the leading consumers of IT power and, as such, are the focus of this paper. Storage and networking equipment will be the subject of future works provided by The Green Grid.

Data centers often are plagued with underutilized servers. Huge energy savings and/or increases in productivity for little or no additional energy consumption can be realized using virtualization technologies. This paper will provide those interested in energy savings with a new avenue by which to immediately reduce their current energy consumption and/or substantially reduce the rate of increase of their data centers for the foreseeable future. While virtualization as a technology is interesting, this paper will only provide guidance on the steps that need to be put into place to exploit this technology for the express purpose of energy savings.

Consolidation using virtualization is one powerful tool that can be applied to many data centers to drive up server efficiency and drive down energy consumption. This paper will provide the background needed to understand both the business and technology aspects of virtualization before an administrator embarks on the energy savings journey of virtualization.

This White Paper is the first paper published by The Green Grid on the energy saving aspects of virtualization. Subsequent white papers from The Green Grid will further describe other aspects of virtualization both within the IT and non-IT domains.

1. SERVER CLASSIFICATION WITHIN A TYPICAL DATA CENTER

Compute, storage, and network equipment dictates the amount of physical infrastructure required in a data center, including power distribution units, generators, cooling units, and other IT support equipment. In addition, the expected growth of the data center also will determine the equipment in place (PDUs, pumps, chiller units, etc.). The typical data center is predominantly occupied by low-cost and underutilized

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

volume servers, otherwise known as x86 servers. Servers in combination with their workloads have varying requirements for performance and uptime, but to simplify the discussion, this paper will categorize them as Innovation,* Production, or Mission-Critical (see Figure 1). Servers are generally used to create, maintain, and execute solutions on behalf of business. In some cases, the solution is so important that it is deemed mission critical. These business classifications drive the types of architecture, processes, and infrastructure that are deployed alongside the solution. These classifications also help people understand what types of opportunities and effort would be involved in consolidation and virtualization. Although there are many ideas on how to categorize servers, categorization is not the intent of this paper. It will use the three broad categories only as an attempt to help readers visualize their opportunities.

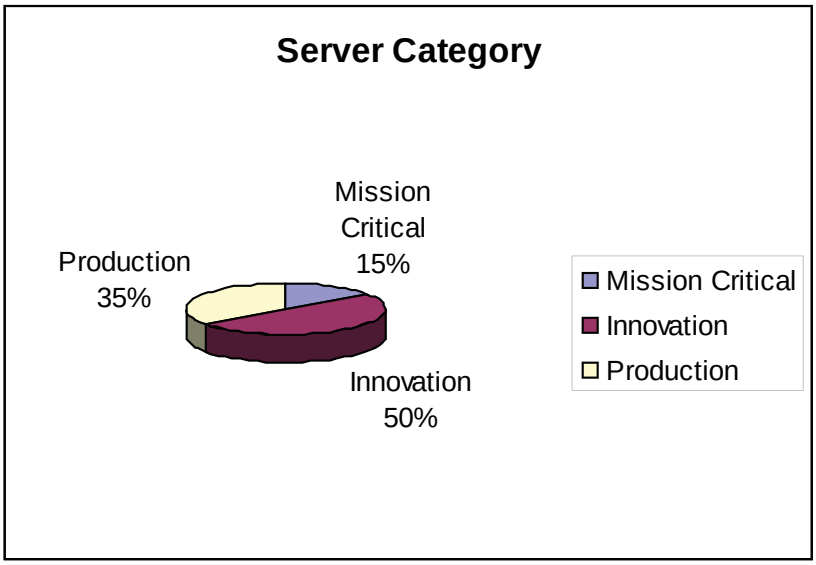

Figure 1. Typical server category percentage

Much of the proliferation of volume servers is due to the need to quickly innovate and deliver solutions to the marketplace. Instead of taking 12 months or longer to deploy applications in large scale, IT professionals working in shared-compute environments can provision new servers and deploy applications in weeks using volume servers. People creating solutions in these shared-compute environments typically have been focused on speed and flexibility.

These represent Innovation environments, which are used to invent new products, modify existing products, develop or enhance processes that make the company more competitive or productive, and much more. These environments include servers typically referred to as Test, Development, Quality Assurance, Volume, etc. In addition to requiring fast deployment, the Innovation environment alternates between periods of active use and periods of inactivity. During times of activity, variations in data, users, transaction rates, operating system patch levels, and other changes occur on a regular basis in order to properly develop and test solutions. The workloads typically are isolated to reduce the impact of those changes on other users who also are innovating. Once a solution is developed in the Innovation environment, the workload is typically moved into a Production environment.

The Production environment is one where change is much more controlled and scalability is extremely

* - By definition, anything that is not production-related or mission critical.

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

important. Deployment speed and flexibility take a back seat to service level requirements. Some of the solutions within an organization are vital to the business and move from the innovation category of systems to a Mission Critical Environment when placed into production. In other words, if the solution were to become unavailable for any period of time, it would have a severe impact on the business. That impact could be the immediate loss of dollars, the loss of customer loyalty, or some other significant effect on the business. These solutions tend to be placed in an infrastructure that will guarantee high performance and availability, often at a significant cost. While mission-critical servers get a good deal of attention and consume many more resources than the typical Production server, they normally account for the smallest number of servers in the data center.

the green grid™

Although this paper could spend a great deal of time debating how many categories of servers there are, the point to be made here is that different servers, workloads, and solutions require different types of infrastructure and processes to embrace virtualization for energy efficiency. As an example, it is unlikely that an IT administrator would run an Innovation workload on a server alongside Mission-Critical workloads. When determining where to start a consolidation and virtualization effort, it must be understood that the Innovation category is where to find the most idle servers and the servers requiring the fewest compute resources. This area represents the best opportunity for consolidation. It is here that consolidation and virtualization will provide the greatest energy reduction in the data center.

The remainder of this paper will refer to a sample environment that includes 500 servers, which are categorized in Table 1 below.

CATEGORIES	INNOVATION	PRODUCTION	MISSION CRITICAL	TOTAL
Server Count	250	175	75	500
Utilization	3%	6%	10%	5%
Watts (173 Per Server)	43,250	30,725	1,297.5	86,500
Consolidation Ratio	15:1	10:1	5:1	~10:1
Remaining Servers	17	18	15	50
Post-Consolidation Utilzation	50%	50%	50%	50%
Post-Consolidation Watts	3,910	4,140	3,450	11,500
Energy Savings	39,340	26,135	9,525	75,000

Table 1. Breakdown of servers in sample data center environment

One study from the VMware Information Warehouse has revealed that more than 20% of the servers assessed by the VMware CapacityPlanner tool in 2007 are running below 0.5% utilization (see Figure 2). This information is drawn from data collected from more than 300,000 servers running in thousands of companies around the world. Approximately 75% of the assessed servers were running below 5% utilization. This data demonstrates that there is a significant opportunity to reduce the number of physical servers and save a tremendous amount of energy simply by eliminating those servers that are no longer in use, sharing seldom-used servers, reducing farms of servers that support a common application, and employing virtualization.

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

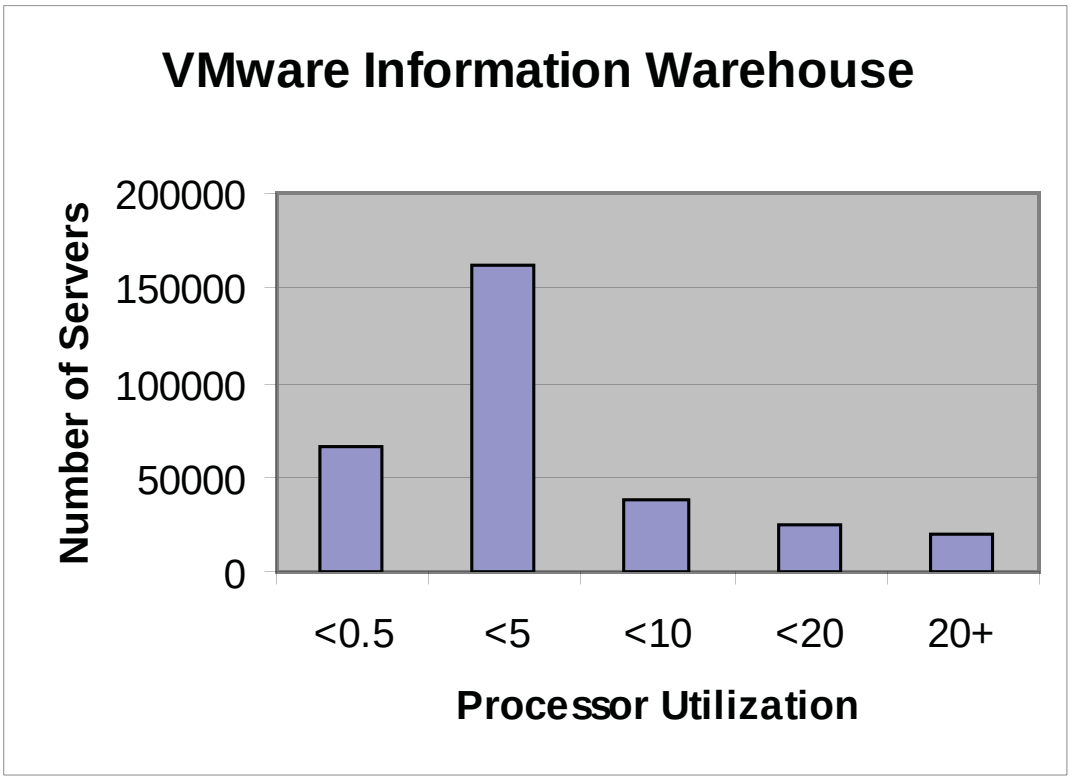

Figure 2. Processor utilization

Reducing the physical server count simplifies the IT infrastructure and does more than just reduce energy consumption. A simplified structure, given lower server counts, results in better service levels (i.e., the guarantee of compute service provided to an application), both from a physical server and business service perspective.

1.1 SERVERS AND ENERGY CONSUMPTION

The energy consumption of a server is not a linear relationship with work or the productivity output of that server. An idle server may consume more than 40% of the power that a fully utilized server does.

If the typical server is running at 5% utilization and the target utilization rate is only 50%, then energy consumption could be greatly reduced by improving server utilization and turning off unused servers. Across the various categories of servers, utilization could be seen at approximately 3% for Innovation servers, 6% for Production servers, and 10% for Mission-Critical servers. These are business hour averages.

Decreasing the number of physical servers does more than reduce the energy consumption of the servers; it also has a positive ripple effect through the entire data center. Server consolidation can result in reduced cooling load (the energy used to remove heat from the data center), increased UPS backup time, and longer generator backup times. In addition, lowering server counts reduces the number of interconnects needed and thereby reduces IT configuration and maintenance in the data center. The result is improved service levels and greater energy efficiency across the data center.

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

Improving energy efficiency using virtualization is an important activity. This study found that at 10% utilization, the server utilized 173 watts of power. Adding an additional 10% of load, which took the machine to 20% utilization, only increased the power utilization from 173 watts to 189 watts. In Figure 3 below, the blue line represents the amount of incremental power load placed on a machine for each additional 10% of processor utilization. At 100% processor utilization, the power utilization peaked at 276 watts. As processor utilization increases, power utilization increases at a much lower rate. (See Figure 4.)

Figure 3. Performance to power ratio

SPEC ® and the benchmark name SPECpower_ssj TM are trademarks of the Standard Performance Evaluation Corporation. For the latest SPECpower_ssj2008 benchmark results, visit http://www.spec.org/power_ssj2008/results/power_ssj2008.html. See http://www.spec.org for more information about SPEC.

In Figure 4, the first bar chart shows the effect of combining two typical workloads into a virtual environment using data points from Figure 3. The result is predicted to be approximately 10% processor utilization and a 173 watt load. Combining enough workloads together to achieve a 50% processor utilization would result in a 230 watt load. If it took 10 workloads to achieve this result (10 physical servers consolidated into 1), the data

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

center would reduce its energy requirement from 1,730 watts (10 servers @ 173 watts each) to only 230 watts (1 server at 50% utilization = 230 watts). Consolidation yields a savings of almost 1,500 watts in this scenario or a savings of 87%.

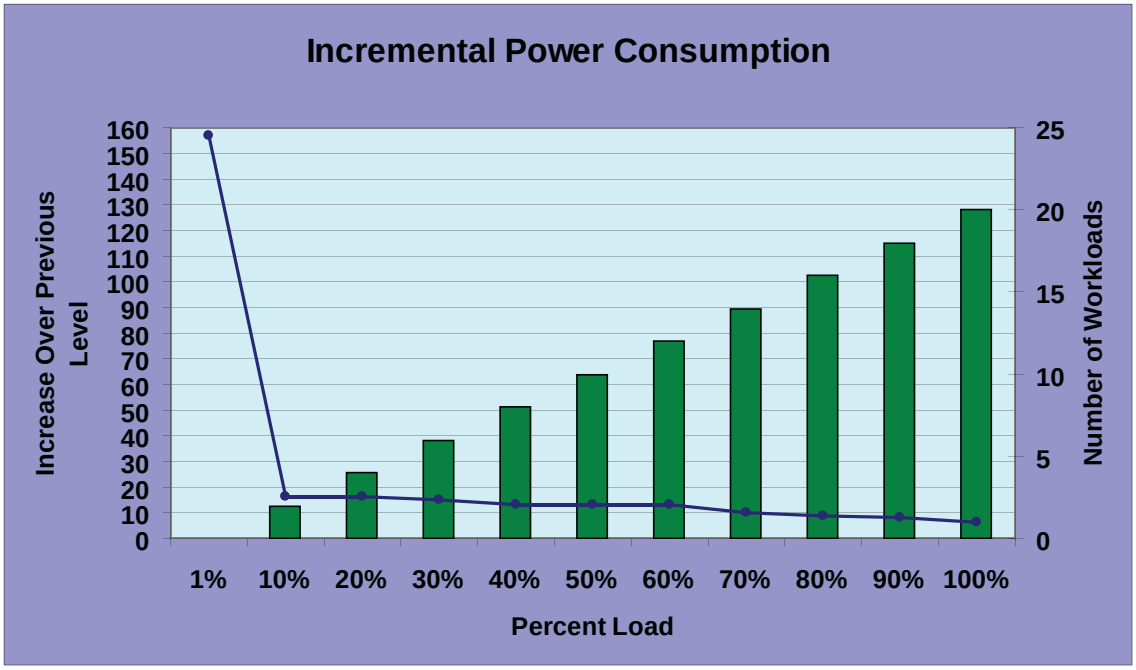

Figure 4. Incremental power consumption

1.2 ACCELERATING BUSINESS SERVICES

In addition to producing energy savings, virtualization technology promotes the acceleration of business services. By setting up virtual machine templates for common server configurations, deploying a “server” is significantly faster than installing an operating system and application on a bare metal server platform. It is even faster than installing an application on an existing server platform that already has an operating system and base configuration loaded.

The ability to deploy virtual servers on demand allows an organization to accelerate its business services. Typically, obtaining a new server requires budgeting and financial approvals, platform selection and sizing, vendor selection, provisioning, installation, patching, securing, and networking. Deploying virtual servers off of a template means dynamic allocation of new capacity within minutes of a business service request. Therefore, both energy efficiency and improved business services can be achieved with virtualization technology.

1.3 UTILIZATION ACROSS SERVER CATEGORIES

The typical Innovation environment contains a large number of servers that have low utilization characteristics and undefined schedules of usage. Some Innovation servers may go days or weeks without being used. However, when in use, an Innovation server may consume every resource available on the server. These servers have a highly variable workload.

The Production category servers usually has more users per application workload and higher transaction

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

volumes than the Innovation servers category, but fewer servers categorized as Mission- Critical. Many Production workloads have a pattern of utilization that increases at the start of the business day and decreases at the end of the business day. Some of these servers may have high utilization rates, but utilization is typically in the 5% range during normal business hours. There is a low probability that a significant number of these servers has processor utilization of more than 30%. This discussion does not refer to memory allocations, which typically only vary 4% during the time a server is powered on.

Mission-Critical servers also have typical processor utilization rates below 30%. If these servers are running at a low utilization, it may be due to over-engineering of the solution in order to ensure high service levels—without concern for energy consumption. In other words, they may only require two cores of processor capacity for the normal workload, but they have eight installed to ensure enough capacity at peak hours or in the event of unforeseen growth in demand. For most users of Mission-Critical servers, utilization is not a concern; meeting the required service level is. Until recently, more servers and lower utilization has been an accepted practice to guarantee performance. High availability and failover capability was more important than utilization percentages.

Across all the categories of servers, there are some that run at more than 30% utilization, but that is not common. The typical server runs at less than 5% utilization and more than 90% of all x86 servers run at less than 10% utilization. This means that a large opportunity exists to consolidate servers in a typical data center. Remember, utilization does not tell the whole story. A 5% utilization rate on a 700 MHz single-core server is not the same as 5% utilization on a 4-core, 3.6 GHz machine. The second example requires much more processing capacity than the first.

Of course, processor consumption is not the only concern. The typical x86 server utilizes approximately 1.5 GB of memory, but that accounts for about half of the memory that is normally installed on these servers for a utilization of approximately 50%. Each memory DIMM consumes additional energy even when idle.

Although local storage has higher rates of consumption, the typical server's throughput infrastructure for both storage and network utilizes less than 1% of the server's capacity, as the average number of bytes per second passed through the infrastructure is around 300,000 bytes per second through the disk subsystem and 150,000 bytes per second through the network.

1.4 SERVER CONSOLIDATION AND VIRTUALIZATION

Server consolidation can be achieved in several ways. The two typical approaches are: 1) combining applications onto a single server and a single operating system instance and 2) virtualizing application workloads in virtual machines and hosting them on a hypervisor*.

As the economy expanded in the late 1980s and 1990s, companies focused on speed-to-market. No longer was it acceptable to wait 18 months for applications to be developed, and a centralized development approach was no longer viable. Many companies chose a decentralized model, buying more and more x86 servers as the prices went lower and lower. When the economy faltered in 2001, CIOs were asked to look for ways to reduce operating costs.

* A hypervisor is software that is capable of hosting multiple virtual machines. Hypervisor software allows operating systems and applications that otherwise would run alone on a single physical server to run on a server shared with numerous other operating systems and applications.

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

Combining applications onto a single server was one way to reduce the number of servers in a data center and therefore lower costs. In 2001, the bulk of volume servers found in a data center typically supported 4 GB or less of RAM. Because the typical application workload required a little more than 1 GB of RAM, the best consolidation possible was a 3:1 ratio.

Although awareness of low server utilization existed, the process and methodology for combining applications onto a single server was not well understood. It would take a significant amount of work by business and IT experts within an organization to verify that multiple workloads (i.e., application workloads) could be run on a single operating system in a Production environment. Even if tests proved that the workloads could coexist in a production environment, there was no assurance that modifications to the operating system, application software, or other changes would not cause application failures.

the green grid™

The related maintenance and issues associated with testing updates to the operating system and applications served as a barrier to the consolidation approach of combining multiple operating systems and applications onto one server. Basic cost, reliability, and service level issues kept most data centers from proceeding with physical consolidation. This type of consolidation was most effective in reducing the number of servers to support multiple instances of like applications. Database, Mail and Collaboration, and Thin Client Application Servers are all examples of applications that may benefit from collapsing multiple instances or workloads into less physical servers.

Virtualization is a form of consolidation. The process of virtualizing an application workload encapsulates both the operating system and the application(s) into a virtual machine. A hypervisor runs on a server platform that hosts one or more virtual machines. From a file system standpoint, a virtual machine is a set of files that resides on non-volatile storage typically within a server or inside a SAN. One file defines the resources that the virtual machine gets to consume based upon parameters defined by the administrator. Since each virtual machine has its own operating system, the maintenance issues associated with the physical consolidation approach do not arise. Placing limits on each virtual machine as to how much of any given resource it can consume alleviates concerns about one workload overtaking all the resource on a machine. This enables different types of workloads with varying characteristics to be combined effectively to improve server utilization.

Server consolidation is not always the right answer. Mission-Critical application workloads may be virtualized to provide better business continuity by meeting service levels, but it is not advisable to combine this virtual machine with another virtual machine. In the Mission-Critical environment, protection of service levels is much more important than improving utilization, achieving energy savings, and minimizing risk of downtime for the applications identified as mission critical.

In addition, workloads in the Innovation environment can be unpredictable, so it may make sense not to combine those with production applications that require more stability and fault tolerance.

Servers that have privacy, security, or regulatory restrictions or ultra-high service level requirements tend not to be the primary or first candidates for virtualization or consolidation. It's likely that other servers also are not good candidates due to technical or business considerations.

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

2. PREPARING FOR VIRTUALIZATION – GATHERING DATA

2.1 IDENTIFYING THE SERVER POPULATION

The first step toward virtualization is discovering all the servers in the organization and taking an inventory of their compute resources along with their associated application workloads. This information is needed to construct a plan for consolidation using virtualization. It is easy to lose track of servers as time goes by, as mergers and acquisitions take place, as people procure servers without going through central IT, or as servers are retired or repurposed. The inventory process may reveal servers that are well-managed and have clear ownership established as well as servers that are not being used at all and possibly were abandoned. If an organization has multiple networks, domains, or directories, every one of them needs to be inventoried. There are tools and services available from numerous vendors that can automate the inventory and discovery process.

After obtaining the server inventory information, group the resources into pools of workloads that can coexist together. It's already been indicated that Innovation servers and Production servers would most likely not coexist. As mentioned earlier, some of those servers may have technical, political, security, privacy, or regulatory concerns that prevent them from sharing resources. To work through this list of concerns, your IT department may need to ask a few more questions about the workloads. Five items of information are needed:

- Who owns the server?
- What is the purpose of the server?
- What is the service level requirement?
- What is the location of the server?
- Is there a requirement for this workload to remain on a dedicated physical resource?

Answering these questions allows you to pool workloads into groups. You then can determine the overall group requirement for compute resources. For example, the exercise may yield a list of 100 Innovation workloads that all reside in the same location. You then can determine the resource requirements for the entire group. In this case, you may determine that you need 5 to 10 servers to support the resource requirements of these 100 workloads.

As a part of the inventory process, the following minimum information should be gathered:

- Make
- Model
- Processor(s) (sockets, cores, threads, speed, cache, front side bus)
- Memory size and speed
- Network (number of ports, speed of each port)
- Local storage (number of disk drives, capacity, RAID configuration)
- Operating System and patch level
- Applications installed
- Services running

Just because software is installed on a server does not mean that it is being used. A check of the services that are running may better indicate which software is in use. However, not all software requires a service to

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

be able to execute, so in some cases, the installed software list may be all that is available. Many vendors ask for and record the name of the installer and the company name during the installation. This information should be collected as part of the software inventory in order to help identify server ownership.

2.2 DETERMINING COMPUTE RESOURCES REQUIRED

What is needed next is a method that estimates the processor resource required for an application workload in a normalized form. Given that the average reported utilization for servers around the world is near 5%, an extremely scientific method for normalizing this information is not needed until average utilization rates are increased above 50%.

the green grid™

The easiest method, and one that is used by the vendors providing consolidation and analysis tools, is to simply multiply the reported utilization (% Processor utilized) by the total processor capacity (maximum processor frequency * number of processors * number of cores). This produces a MHz requirement that can be applied to target servers. You can refine this calculation by normalizing for multi-core overhead as well. Remember that the average processor is underutilized and typically is not the resource that limits consolidation.

In the information survey of machines assessed in the 2007 CapacityPlanner mentioned earlier, the average reported utilization was 5.365% during prime-time hours and 12.39% during periods of highest sustained load. Using the normalization techniques described above to adjust the demand to a MHz rating, the actual demand levels (including overhead for virtualization) turn out to be 266 MHz during prime-time hours and 595 MHz during peak loads.

The typical server today has four cores running at 3.6 GHz for a total capacity of 14.4 GHz. If the number of workloads combined together would not exceed 50% utilization in the event that all the loads peaked at the same time, the combined workloads should demand no more than 7.2 GHz of processing power. Using the average peak sustained load numbers, this means that 12 typical workloads could be combined into a single server.

The reality is that servers do not all peak at the same time. The actual utilization of a 12:1 consolidation ratio would most likely run well below the 50% target. It must be noted that this is an average among all servers assessed, and the survey does not make it possible to identify each server as whether it is an Innovation, Production, or Mission-Critical server. It is logical to assume that capacity utilization rates are lowest in the Innovation environment and highest in the Mission-Critical server environment. It also is logical to assume that consolidation rates can be much higher than 12:1 in the Innovation environment and lower in the Mission-Critical environment.

The average data center server has 2.6 GB of RAM installed and consumes approximately 43.5% of that amount of memory, or 1.145 GB. A consolidation ratio of 12:1 dictates that 16 GB of RAM will need to be installed in a virtual host to support this ratio. Once a server starts up and opens its primary applications, the memory allocation typically varies only slightly, about 4%. Therefore, having available, unallocated memory in a server is not as necessary as it is for processor capacity, where demand can fluctuate by more than 100%. The typical server performs 35 physical disk I/O operations per second during business hours. Peak sustained loads average 125 physical disk I/O operations per second. Today's servers are purported to support tens of thousands of I/O operations per second, thus making a consolidation ratio of 12:1 extremely doable. The average number of bytes per second through the disk infrastructure (including reads and writes)

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

is approximately 420,000 with a peak sustained load of almost 2.8 million. Again, servers today have capacity that far exceeds the estimate of combining 12 typical servers together.

While most servers today have two network interface cards installed with the capability of transferring 1 GB of data per second, the average number of bits per second sent through typical server interfaces is approximately 1 Mbps, or .1% of the total capacity. Even at peak sustained loads, the demand only climbs to 5.8 Mbps, or slightly under 1% utilization of a single network interface card.

the green grid™

As a result, consolidation ratios in the 10:1 to 15:1 range can be achieved without placing stress on the server resources. However, every organization contains servers that have high demands placed on them. Also, there are servers with utilization patterns that fluctuate; these servers occasionally have extremely high demands placed on them. These should be avoided in initial consolidation efforts, but fortunately they tend to make up a very small percentage of servers in the data center. They typically represent only 3% to 5% of the server infrastructure, leaving the vast majority of the infrastructure as candidates for virtualization.
Also, placement is not permanent in a virtualized environment. If a combination of workloads causes a critical processor load condition, then workloads can be moved to a different physical server that is not constrained.

2.3 ESTIMATING CAPACITY REQUIREMENTS

Servers have different application workload characteristics based upon the type of application and the population of users who rely upon that application. Applications that are used primarily during business hours will peak during business hours. Some applications peak during off hours. Even business applications may peak during off hours due to backup processing and maintenance windows.

It is important to note that the application workload can be characterized at various times of the day, week, and month or longer if the business cycle dictates (e.g., changed to handle year-end activities). This technique ensures that peak and sustained-load times are captured in addition to normal usage. Servers should be sampled at least once per hour over a 30-day period to develop a workload profile. Using this technique, a load profile with a minimum of 720 samples can be developed.

By gathering information this way, an organization can combine workloads that have complementary characteristics. For example, combining workload A, which consumes 10% processor capacity during business hours and very little at night, with workload B, which consumes 10% processor capacity during the night and very little during business hours, delivers the best possible results.

Data centers also need an estimate of the capacity requirements for those workloads that are being consolidated and/or virtualized. There is a lot of discussion around sampling. It has been suggested that real-time data collection is more valuable than sampling over time. However, the primary objective in the collection process is to gather statistically relevant data while doing no harm to the service levels of the applications that rely on the server that is being evaluated. The more data that is collected and the more often it is collected, the more harm done because it adds overhead to the server.

Recall the sample data center introduced on page 9 with 1,000 volume servers and a server population categorized as 50% Innovation, 35% Normal Production, and 15% Mission-Critical Production. If these 1,000 servers were spread across two data centers, then six pools of capacity (each data center containing three categorized types of servers) with similar service level requirements exist. To simplify this example, let's assume there is just one of these data centers and 500 Innovation servers. Capacity and utilization

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

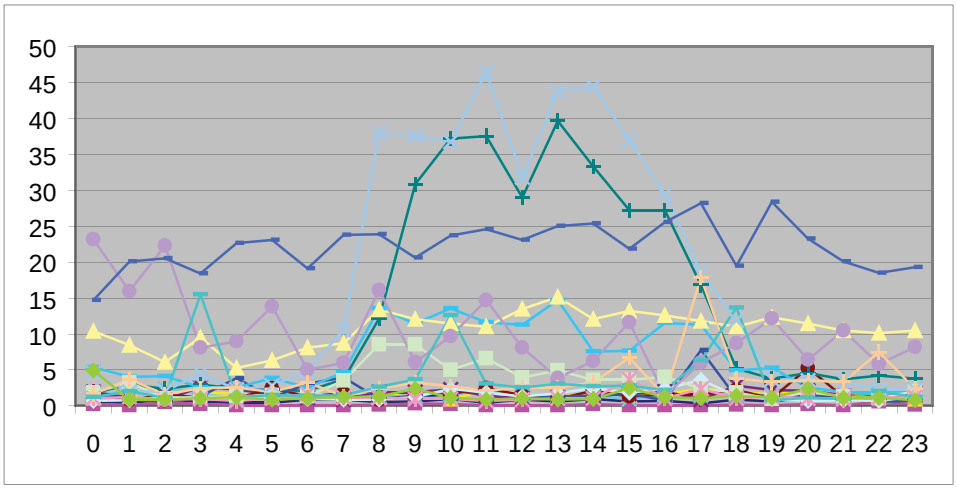

Figure 5. Non-normalized demand variance

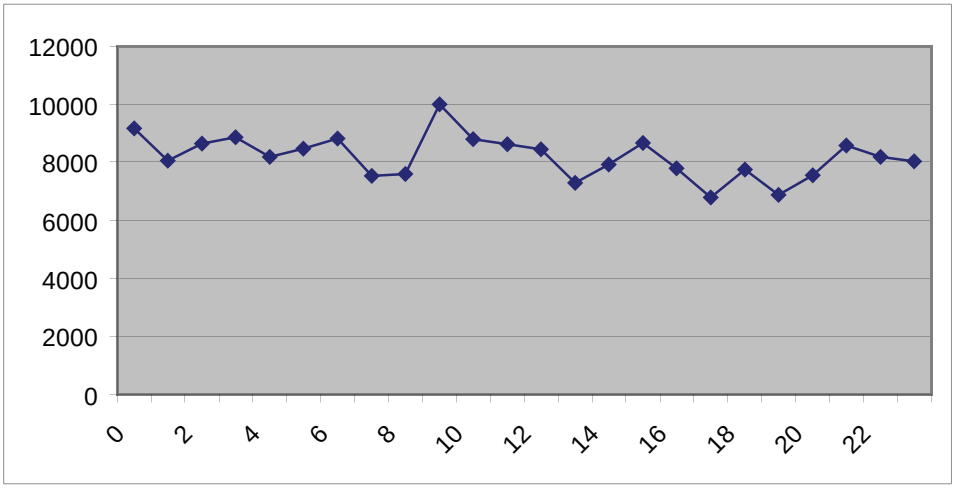

Figure 6. Normalized processor load

Figure 6 demonstrates the smoothing effect. This chart is created by normalizing the processor load by using the formula Processor Count * Processor Speed * Processor Utilization and then adding together this value for all servers for each hour of the day. The total demand variance is minimal. This smoothing effect becomes more pronounced as the number of servers in the group increases. On the opposite end of the spectrum, combining only two loads together would have almost no smoothing effect at all. As consolidation ratios increase, the smoothing effect is more pronounced as individual variations in workload demand are averaged out across a larger number of overall workloads.

As mentioned earlier, the estimated average server places a 266 MHz load on a shared environment during business hours, with peak loads reaching an average of 595 MHz. Adopting the strategy of assuming the worst—that all workloads could potentially peak at the same time—it would make sense to use 595 MHz as the planning number. If the strategy assumes that workloads peak at various times, then the 265 MHz load number may be used. Taking the most conservative approach possible, the 595 MHz number, the total

* Innovation, Production, and Mission-Critical servers.

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

requirement for the Innovation environment is 297.5 GHz of processing power. The average allocation of memory is approximately 1.15 GB. The total memory requirement is calculated by multiplying 500 servers by 1.15 GB of RAM, which equals 575 GB.

These numbers should be adjusted to accommodate growth estimates and also desired utilization level. A target objective of 50% processor utilization requires almost 600 GHz of processing power. An objective to run at 90% utilization of memory requires approximately 640 GB (575 GB/.9) of RAM. The memory and processing capacity requirements for the sample 500-server pool could actually be accommodated on 42 servers, each containing four cores at 3.6 GHz with 16 GB of installed RAM. Innovation servers run below the average utilization rates and consume a below-average amount of RAM.

the green grid™

Energy Savings: From an energy consumption standpoint, the data center has moved from having 500 servers utilizing 161 watts each (80,500 watts in the case above) to 42 servers utilizing 214 watts each (8,988 watts at 50%). That represents an 89% decrease in server energy consumption, which does not include associated reductions in cooling and other infrastructure requirements.

This same exercise can be carried out for each of the three server categories, although the level of consolidation may be less with the other two.

2.4 APPLICATION WORKLOAD CHARACTERIZATION

So far, this paper has only characterized workloads by their service level objectives and locations in order to determine capacity requirements for each pool of resources. The idea was to identify workloads that can be combined in order to determine compute resource requirements. You also must consider that there will be a percentage of servers—regardless of their assigned category (Innovation, Production, or Mission-Critical)—that will require further study and may not be part of your initial consolidation effort. There are various reasons for this, but some include:

Technical

A server may have specific hardware or peripherals that do not allow it to be virtualized. Or the software vendor may not have certified its application to run in a virtual or shared environment.

High Utilization

Servers that exhibit utilization above the norm could dominate the resources in a shared environment. These servers' workloads should be evaluated to determine if the high utilization can be reduced through application optimization prior to adding them to a shared environment. In the early stages of consolidation and virtualization, it is important to give your project the highest probability of success. This means mitigating any potential risk. Including high-utilization workloads certainly may be feasible, but it is not worth the risk early on in your project.

Volatile Workload

Some workloads are very unpredictable. When these types of workloads become active, they drive up resource demands and could affect the rest of the workloads in the shared environment. Volatile workloads represent a small percentage of the typical data center environment and also should be avoided in the early stages of a consolidation/virtualization project.

Privacy and Security

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

Some workloads rely on extremely sensitive data. Although loads in virtual servers are protected from access by other loads on the same physical machine, the end-user community may not initially grasp that concept. As your consolidation project demonstrates success, it will become less difficult to convince end users that their workloads and their sensitive data will be protected.

Regulatory

There are servers that may be regulated by a government agency due to the types of applications that the server supports. For example, if the server is used to manage a production line for food or industrial gases, you may need permission from a regulatory agency before being allowed to consolidate or virtualize this workload.

Business Sensitivity

If you are in the early stages of a consolidation/virtualization project, your internal users will not be familiar with the technology. At this point, it is unlikely that your users will allow you to virtualize mission-critical applications. Even if they do, you should wait to move those workloads until you have gained experience with the new architecture. Once these criteria are applied, you may be left with a set of servers that will not be part of the consolidation/virtualization categorization plan as described above.

Organizations usually try to keep the total cost of ownership (TCO) for the Innovation environment as low as possible, including the need to incorporate power consumption costs in the TCO. Utilization objectives and consolidation ratios for this server pool typically will be much higher than for the Production or Mission-Critical pools.

Organizations may place restrictions on which workloads can coexist with which with regard to specific servers that have privacy, security, or regulatory concerns. These restrictions may limit the consolidation ratios for groups containing those servers.

Once you have grouped all of your servers into the appropriate categories and sub-categories, you will need to assign utilization objectives. While you may select utilization targets at or above 50% for the Innovation pool, you may only feel comfortable with a range of 25% to 50% for the Production pool. Your consolidation ratio will be driven by the utilization percentages for not just processors but for the other components as well. At the end of the exercise, you will be able to identify your best opportunity for savings, which will most likely be in the Innovation area. The Innovation environment normally has the most servers, the lowest utilization, and the highest tolerance for error. As data centers' costs rise regarding energy consumption and availability of energy, the numbers of virtualized servers may increase to allow more processing at lower levels of energy consumption.

In summary, you not only need to characterize the workloads that you are planning on consolidating, but you also must characterize the target environment into which you are sending those workloads. You need to determine how busy it should be, how much it will cost, what type of technology it should have in order to achieve the service levels required or meet security/privacy objectives, and so on. Also, can your data center infrastructure handle the increased power and cooling densities that virtualization places upon them due to the increases in each server's utilization and accompanying power consumption? You also need to consider that if a single server fails now, many more applications will be affected.

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

3. OTHER CONSIDERATIONS

3.1 METRICS AND MEASUREMENTS

Organizations commonly use processor utilization as a measurement for success, with 5% utilization considered to be inefficient (poor), 25% to be better, 50% to be good, 75% to be very good, and 100% to be best. However, this can be misleading because these are corporate-wide averages. Some percentage of servers will not be candidates for virtualization. These non-virtualized servers are still part of the overall computation for corporate-wide server utilization.

Non-virtualized servers may have represented 15% of the physical server population prior to the virtualization project. Post-virtualization, these same servers may represent 85% of the server population.
For example, consider the data center with 1,000 servers. With typical consolidation ratios of 10:1, it is reasonable to assume that 850 physical servers will be virtualized and consolidated to 85 physical servers running at 50% utilization, with the other 150 non-virtualized servers running at 5% utilization. If your definition of success is to achieve 50% processor utilization, it would not have met your objective, even though you have achieved significant cost reductions. You would be left with 235 physical servers with an average utilization of 21% (85 servers @ 50% utilization and 150 @ 5% utilization).
Additional metrics and measurements to consider include:

- Reduction in energy consumption
- Reduction in server count
- Reduction in peripheral components
- Reduction in floor space
- Reduction in cooling requirements
- Reduction in business service delivery

3.2 ESTABLISH SERVICE LEVEL OBJECTIVES

Service level objectives are a key consideration for consolidation. These objectives are different within each of the Innovation, Production, and Mission-Critical workload environments.

Mission-Critical application workloads demand stable environments and a high service level, while Innovation environments change constantly. Arguably, the number one cause of application outages is change. Combining Mission-Critical and Innovation workloads places additional risks on Mission-critical tasks and is not recommended. This is why it is so important that you find and label all the Mission-Critical, Production, and Innovation servers in your data center and then establish service level objectives for each and every one. The Innovation environment should have the lowest TCO, the highest rate of change, and the most volatility. The Production environment should have a steady utilization rate. Fault tolerance is important here because these applications support the day-to-day operations of the company. The loss of a Mission-Critical application workload has an impact on an organization's financial health and/or customers every second that the application is unavailable or not performing at expected levels.

Service level objectives drive hardware purchasing decisions, redundancy in configuration, throughput requirements, disaster recovery plan (DRP) and business continuity program (BCP) objectives, and much more. Once the service level objectives have been established and the dedicated hardware's memory, CPUs, etc. are known, it is possible to predict how much to spend on each of the three environments. To do so, summarize the total capacity requirements for each environment and establish a baseline cost for each type of capacity.

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

3.3 PRIORITIZATION/PHASES

A consolidation/virtualization project plan must be constructed in a way that limits risk in the early phases yet realizes significant progress toward the objectives, metrics, and measurements as the organization gains consolidation experience.

The initial stages of the plan should include servers with workloads that have low service level objectives and high potential consolidation ratios. As described earlier, the following steps should be taken:

- Rank servers by type of workload
- Rank servers within each type of workload by processor utilization
- Examine the service levels established for each application workload

Although it is possible to combine workloads that have various service level objectives, avoid doing so during the early stages.

The objective of phase 1 of the project is to get high consolidation ratios, meet service level objectives, and implement a process that promotes manageability of the environment.

High consolidation ratios demonstrate reduced TCO as well as reductions in floor space, cooling requirements, and energy consumption. Another important aspect of the Innovation environment is the rate of change. The server configurations, software, and data for Innovation workloads change on a regular basis. Business services can be accelerated by allowing the user community to allocate and manage capacity on demand as well as to create backup strategies and versions of their application environments on demand. Production servers have more predictable utilization patterns. Production environment users generally are not allowed to make changes without going through a change management process, and service level objectives typically are well understood. This means that the TCO for this type of environment is higher. Fault tolerance, backup plans, and disaster recovery plans are key factors in designing the environment to house Production workloads.

The Production environment's typical consolidation ratio may be lower than that of the Innovation environment in order to accommodate necessary headroom for Production service levels. Therefore, TCO may need to be adjusted for the Production environment. You also may be choosing more advanced technology for the host components to ensure higher throughput and better service levels.

Some workloads are difficult to identify because you are trying to delineate between Mission-Critical Production and Normal Production servers. Both types of categories may have the same naming convention, so the initial focus will be on resource utilization patterns. The characteristics of Mission-Critical servers include ultra-high transaction rates, large databases, high I/O counts, high network packet counts, top processor percentiles, high memory consumption, and so on.

When trying to determine which servers to consolidate in phase 2 and which in phase 3, it might be easier to put server workloads with the following characteristics into phase 3:

- Servers with highly volatile utilization patterns
- Servers requiring more than 4 GB of RAM
- Servers with utilization of more than 20% of the target processor capability
- Servers with I/O rates greater than 10,000 I/O per second

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

- Servers with network bytes per second of greater than 5,000,000

If your environment is like that of most other organizations, you will have put approximately 10% to 15% of your total infrastructure into phase 3.

3.4 TECHNOLOGY CONSIDERATIONS

Consolidation ratio, service levels, and projected growth are all key factors in determining what technology to use to host a virtual environment. You now know that the typical server runs at or below 5% utilization of processor resources, but you also should know that the typical volume server utilizes less than 1% of the throughput capability of both the disk and network components of the server. Memory is the one resource that is consumed at a more reasonable 60% utilization. When making technology decisions for target platforms, consider that memory is often the factor that limits consolidation ratios.

Consolidation ratios vary depending upon the environment that you are attempting to consolidate. The Innovation environment often will enjoy high consolidation ratios, frequently exceeding 20:1, while the Production and Mission-Critical environments may achieve consolidation ratios that are closer to 10:1. The following chart depicts the resource requirements for consolidating typical servers at various ratios. While combining 10 servers requires a large amount of board memory, the requirements for the other resources are relatively low.

Processor Core Count	Processor Speed MHz	Consolidation Ratio								
		5	10	15	20	25	30	35	40	60
2	2400	24%	**48%**	72%	96%	*	*	*	*	*
4	2400	12%	24%	36%	**48%**	60%	72%	84%	96%	*
8	2400	6%	12%	18%	24%	30%	36%	42%	**48%**	72%
2	3600	16%	32%	**48%**	64%	80%	96%	*	*	*
4	3600	8%	16%	24%	32%	40%	**48%**	56%	64%	96%
8	3600	4%	8%	12%	16%	20%	24%	28%	32%	**48%**
	RAM GB	8	16	23	31	39	47	54	62	93
	Disk GB	500	1000	1500	2000	2500	3000	3500	4000	6000

Figure 7: Resource Requirements vary with Consolidation Ratio

3.5 PROCESS CONSIDERATIONS

- An IT process to accommodate self-service provisioning versus physical provisioning
- Change management to extend Mission-Critical practices
- Capacity management to extend Mission-Critical practices and the planning horizon
- Performance tuning of existing server platforms
- Consolidation increases the criticality of the server platform, requiring service level objectives to be established for all shared environments.
- Reporting that monitors health and measures resource utilization

3.6 MEASUREMENTS OF SUCCESS

- CPU utilization

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

- Server platform count reduction
- Business service improvements (due to the reduced time to allocate capacity with virtual versus physical servers)
- Deployment benchmarking to gain support within the organization
- Reduction in energy consumption and utility bill
- More room to increase capacity
- Floor space reduction
- Cooling reduction
- UPS that covers IT for longer periods of time
- Backup generator(s) that cover IT for longer periods of time
- Service level improvement

the green grid™

4. CONCLUSION

Virtualization is a mature technology that works well and is here to stay. It saves energy and money in your utility bill, increases computing throughput, frees up floor space, and facilitates load migration and disaster recovery. The virtualization savings entitlement can be multiplied by upgrading the power and cooling infrastructure to run efficiently when IT loading is substantially reduced due to consolidation and virtualization.

Looking forward, the three major challenges that virtualization poses to physical infrastructure are dynamic high density, under-loading of power/cooling systems, and the need for real-time, rack-level management. These challenges can be met by row-based cooling, scalable power and cooling, and predictive management tools. All three of these solutions are based on design principles that simultaneously resolve functional challenges and increase efficiency.

It is essential in a virtualized environment to manage the data center as a holistic system comprised of both IT and physical infrastructure components. Because virtualization generally results in power savings, it can affect how data centers are powered and cooled. Understanding the impact of virtualization on physical infrastructure is the key to mapping out a sound strategy.

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

REFERENCES

1 Belady, C. and Patterson, M. (2008) The Green Grid Productivity Indicator;
http://www.thegreengrid.org/gg_content/White_Paper_15_-_TGG_Productivity_Indicator_063008.pdf

2 Anderson, D. et al. (2008) A Framework for Data Center Energy Productivity
http://www.thegreengrid.org/gg_content/White_Paper_13_-_Framework_for_Data_Center_Energy_Productivity5.9.08.pdf
and http://www.spec.org/power_ssj2008/

3 US Department of Energy/US Environmental Protection Agency ENERGY STAR Program. ENERGY STAR Version 1.0 Program Requirements for Computer Servers: Draft 3.
http://www.energystar.gov/ia/partners/prod_development/new_specs/downloads/servers/Draft3_Server_Spec_110408.pdf

the green grid™

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.