



ROBERT FRANCES GROUP
Business Advisors to IT Executives

120 Post Road West, Suite 201, Westport, CT 06880 Phone: 203-429-8950 Fax: 203-429-8930

The Rise to Power of...Power: Dealing with the new Data Center Constraint

Jerald Murphy
Senior Vice President and Service Director
Robert Frances Group



Table of Contents

Executive Overview	3
Power Costs and Challenges today	4
Moore's Law: a two-edged sword.....	4
Increased density = more computing power.....	5
Increased density = more power/ft ² consumed	5
Power cycle doubled	6
Increase power lead to increased cooling	6
Increased cooling requirement itself increased power consume.....	6
Increase AC needs increases acoustics in the data center, also difficult to deliver the volume of air needed	7
Shift moving from performance as a gating factor, to power as a gating factor	7
Graphs on power growth and % utilization of CPUs	8
Graph on increased power consumption	9
Critical trend shift	10
Direct Benefits of Power Efficient Blades	10
Reduced power consumption – lower energy costs	10
Reduced Cooling requirement – lower cooling costs, which itself reduces power consumption.....	11
Additional Benefits of Power Efficient Blades	11
Additional space available in data center	11
Reduced facilities costs	12
Increased server density	12
Better power management techniques.....	12
More effective space utilization.....	13
Reduced system failure rate and better uptime	13
Even without benefits, you cannot afford to ignore power problem	14
Case study: Gains that can be achieved by Smart power design.....	15
Conclusion.....	16



Executive Overview

For the last several years, power issues in the enterprise have been rapidly escalating. While industry has seen the direct benefit of Moore's Law with respect to the price/performance curve, it exists with a dark underbelly companies have not noticed until this past year: power consumption has increased at the same exponential rate that performance has increased.

IT executives have been relatively unaware of this phenomenon until recently because the direct cost of this problem has not been borne by the IT shop. Typically, facilities managers are responsible for paying the electric bill. Until 2-3 years ago, electrical usage has been considered a commodity. Electrical costs have been relatively cheap, total consumption was relatively low, and the emphasis in the business has been increasing performance.

The last year has seen a fundamental shift in the performance metrics for IT in the enterprise. Performance of integrated circuits has gotten to the point that raw computation is now the commodity. Across industry, the average utilization of most processors in the data center is 15-20%. At 20% utilization, there is not always a compelling need for the fastest processor, but for the processor that can efficiently meet business objectives.

At the same time, power has moved from being a commoditized utility to being an operational constraint. The cost of electricity has started to skyrocket, along with the rest of the energy industry. Utilities often do not have enough capacity to satiate the ever-increasing data center demands. Finally, data centers do not have the physical capability to handle the power and cooling requirements of a data center full of blade servers, even if the company can afford it and the utility industry can provide it.

RFG believes power efficiency will move to become the number one IT issue in the data center for the next 3-4 years. Already, RFG has interviewed many IT executives in 2005 and found power already to be the number one IT issue for many large financial institutions.

In order to get out from under this enormous problem, IT executives will have to tackle it from every angle. Not only will IT and facilities have to work closer for power planning, but IT executives must shift their computing criteria priorities from price performance to power performance. While each area of the data center must be examined for power efficiency, it is clear that the increasing density of servers in the data center both provide the greatest current limitations and greatest potential for efficiency gains. IT executives should look closely at server manufacturers who focus on power efficiency to get the greatest ROI.



Additionally, IT executives should look not just at the cost reduction of these systems, but understand the domino ROI effects that grow out of this power efficiency. Among the other areas to benefit are reductions in data center space, reduced data center / facilities management costs, reduced system failure rate, and more efficient space utilization.

Power Costs and Challenges today

IT executives will always be presented with challenges, and those challenges are constantly shifting. The most recent shift that has been happening the last three years, almost without notice, is the consumption of power in the data center. This paper will review the history of how this has come about, why IT executives must focus on this today, and how the benefits of a power-efficient approach to data center design, especially with blade servers, yields benefits far beyond the reduction in power consumption.

Moore's Law: a two-edged sword

Moore's Law states that the complexity of integrated circuits will double every eighteen months. From an historic perspective, this was more an observation by Gordon Moore than a law. However, it has remained incredibly consistent over the last 40 years, and will likely hold true for at least another decade.

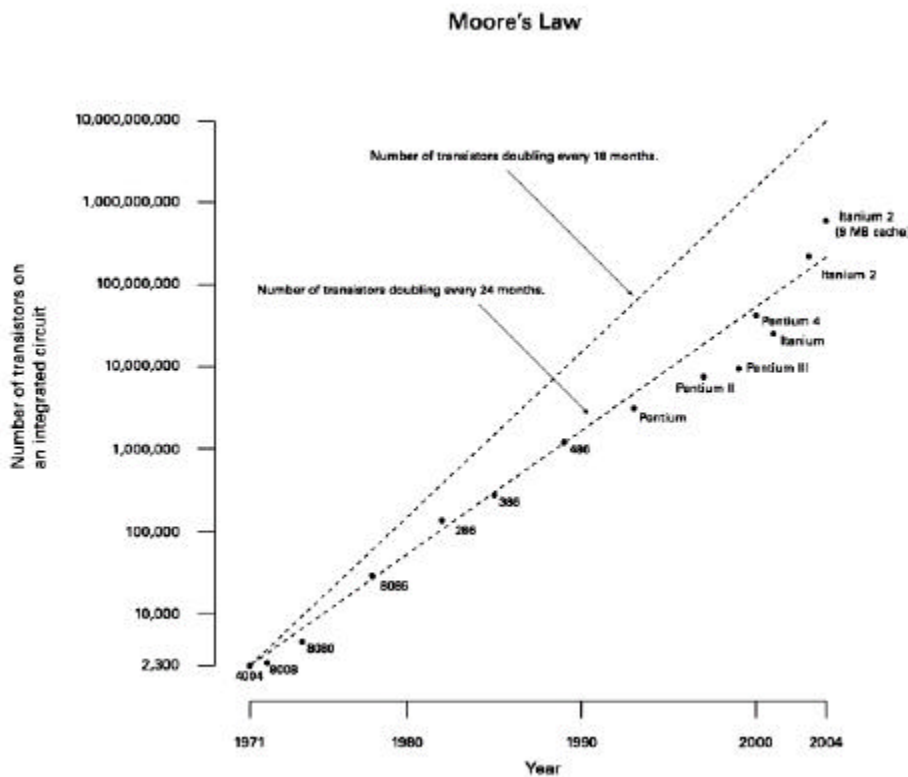


Figure 1 Moore's Law Growth Chart – Source Wikipedia, 2006



Increased density = more computing power

The good news about Moore's Law is that it has allowed more computing power to be applied for the same price. In practice, this has meant, over time, more computing power for less money. From 1958 to 1992, inflation increased at an average 3.66%, while computer costs decreased on average of 19.13%.¹ Buying computing cycles has been a procrastinator's paradise. Vendors have continued to leapfrog over each other which better benchmark performance. At the same time, competition has forced prices to steadily decline over time.

Increased density = more power/ft² consumed

The bad news about Moore's Law is that power consumption in chips has been growing at the same geometric rate as computing power. As geometries shrink, more computing is done in less area. Additionally, as clock cycles increase, the amount of power consumption per processor goes up. In reality, this has actually made the power problem even worse than the performance growth.

In early computers, transistors that drove computation were made with bipolar designs. These circuits had low voltage changes, and could switch quickly because the bias current was always on. Unfortunately, the constant current draw was very inefficient. Early mainframe computers had to be water cooled, and supercomputers built with these basic circuits could consume 1500 amps of current.

With the advent of complimentary, metal oxide semiconductors (CMOS), the circuits only consumed power when they changed state. While these circuits were switched more slowly (since they initially switched at 5 volts as opposed to the 60 millivolts of bipolar circuits), they consumed dramatically less power. The advent of these circuits during the start of the personal computer age allowed the rapid, low cost growth of the personal computer industry.

Of course, Moore's Law was hard at work over the years. The reduction of lithography sizes allowed voltages to decrease, more sophisticated processing to occur in a smaller area, which then followed by high switching speeds.

The little-noticed effect on the power side of the equation is the enormous growth of power consumed by these devices. Like the story of the sage who asked for one grain of rice on a chess board, doubling to two grains on the second square and four on the fourth, the total power consumed over the past ten years has been not large enough to notice. However, this geometric expansion of power has finally caught up with the data center.

One of the reasons it has taken so long to notice is that the person who owned the power budget has not been the same person as the one who owned the data center. Typically,

¹ "Computers and Growth", Jorgenson, D and Stiroh, K, Econ. Innov. And new Tech. , 1995, Vol3, p 295.



facilities managers control the physical aspects of a building: plumbing, electrical, and cooling. While IT managers need power and cooling, up until the last 2-3 years the power consumed was not dramatic enough to warrant detailed planning. Facilities managers know data centers needed cooling of stacked systems, but the cooling requirement was easily addressed with raised floors with cooling channels underneath perforated tile floors. The typically power requirement per rack did not exceed 2 kilowatts.

In the last three years, Moore's Law has not been that sympathetic.

Power cycle doubled

While Moore's Law has not been good to the power problem, it is actually worse than it even looks on the surface. A power catch-22 is at play here that works to compound the problem.

Moore's Law shows that circuits will reduce in size every 18 months. At the same time, clock speeds are increasing. With smaller circuits, electrons don't have to travel as far to switch, allowing the clock rate to increase. These factors have a multiplicative effect on performance. In addition, with more circuitry running on the same size wafer, increased performance occurs at a lower price. However, faster circuit switch speed increases the amount of energy released. Faster moving electrons increase friction, which increase heat dissipation. So, with faster circuits, and more circuits on a chip, heat dissipation grows geometrically as well.

Increase power lead to increased cooling

As denser chips with faster clock cycles have led to a dramatic increase in power consumption, another problem that needs to be addressed is how to cool these chips. As lithographies shrink and clock speeds increase, the heat that is generated on each chip has gone up. Even as companies move to dual-core processors with reduced clock cycles, the problem does not go away. Multiple cores on a chip means that a greater percentage of the chip is transitioning at the same time. So, while multi-core chips at lower clock speeds may reduce the power per CPU, the total power and heat dissipated by each chip will likely continue.

Increased cooling requirement itself increased power consume

This increased requirement for cooling directly results in larger chiller units that are needed to cool these systems. These larger cooling systems themselves require more electricity, which further exacerbates the problem. The large density of heat generated requires a much larger volume of air flow to get rid of the heat generated. This directly leads to the need for more air conditioning for data centers, which itself consumes up to 50% of the power requirement for the data center



Increase AC needs increases acoustics in the data center, also difficult to deliver the volume of air needed

Not only do the larger chilling systems require more power, but they also take up more space, which becomes yet another problem that needs to be addressed. Many companies that RFG has talked with find that, even if they have the money for the power and can engineer a way to get the cooling to the systems, they do not have the physical space to put these systems in. The problems with cooling here are threefold. First, the chilling systems take up floor space that was calculated to be used for servers. Facilities managers typically plan building construction and space in 7 – 10 year plans, with facility use often planned for 20 years. If space reserved for data centers ends up not being enough space, there usually isn't spare space in the facility left to handle extra equipment. If more space is needed, that can entail additional building construction, which falls out of time requirements required by IT.

An alternative to putting chiller units closer to compute resources is to increase the airflow in the raised floor columns to be able to generate the airflow necessary to cool everything in a rack. The problem with this approach is that to get special modifications to the data center environment usually need to be made in order to handle larger airflow densities. Some of these actions include more raised floor space, special-purpose airflow grids, and additional, special-purpose airflow conduits. So, even if the facilities were planned and designed to handle the power and electricity needed, current system power densities would still make it difficult to handle the cooling requirements.

Server design issues often do not help either. When power was a commodity, manufacturers focused on putting in cheap parts, not efficient parts. Many servers today have inefficient power supplies, which wastes as much as 35% of the power delivered to the server. Cheap, noisy, inefficient fans can lead to hot spots, when these fans fail they cannot be easily replaced which is an availability issue. Air flow in server racks has not been paid attention to closely either. While data center best practices today dictate front to back movement of air in the system, earlier designs that have bottom to top cooling lead to uneven heat dissipation leading to heat problems for servers, the rack and the room. In today's rack-based environments, this can lead to hot spots in the computer room as well, higher system failure rates in the top half of fully loaded server racks.

Shift moving from performance as a gating factor, to power as a gating factor

Based on the above discussion, it is clear to see that a fundamental shift has occurred in the server industry. This shift is fundamental, profound, with huge impacts in IT executives decision making criteria and the vendors requirements for delivering effective solutions.

The key changes that have occurred are threefold. First, integrated circuit performance has become a commodity. Second, power and cooling have moved from being



commodities to critical system constraints. Third, space is an equally important tertiary constraint that springs out of the power and cooling limitations. RFG recommends IT executives shift their planning considerations from price/performance to power efficiency to address current (for some) and future (for most) data center designs.

Graphs on power growth and % utilization of CPUs

The graphs below depict the changes in power growth and percentage of CPU utilization for servers in the data center.

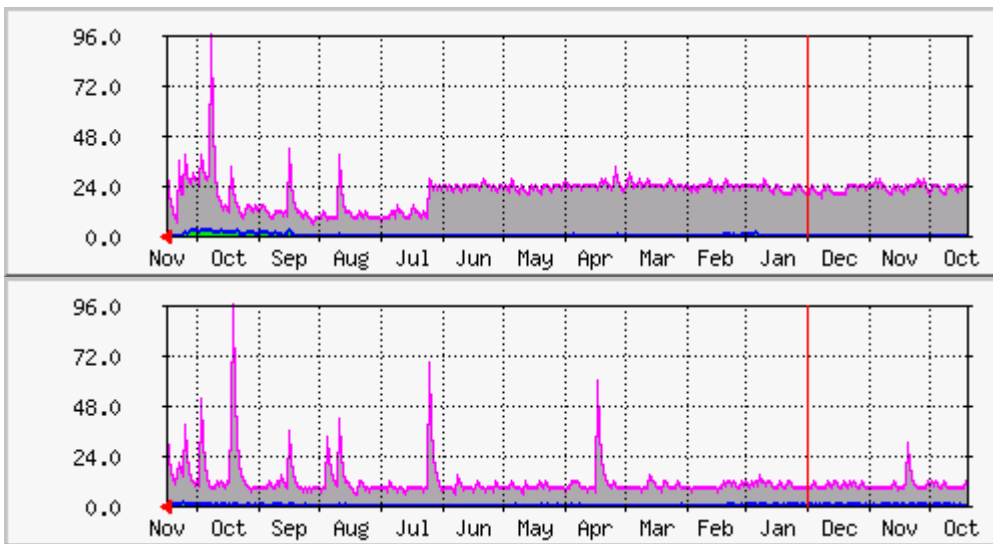


Figure 2 Typical CPU Utilization Figures

The figures above are typical of CPU utilization for typical data center servers. The significant thing to notice is that typical CPU utilization for servers generally doesn't exceed 25%.

A more sobering statistic is the power consumption of processors over time. Since the first Intel 8088 processor was manufactured in 1979, power consumption per processor has increased from 4.7 Watts to over 130 Watts for a Pentium IV processor running at 3.4 Ghz (see Figure 3). While much of the power consumption on servers is directly attributed to these CPU increases, other server factors are not helping. RFG finds that most servers have internal components that are inexpensive, but not efficient. For example, power supplies often only provide 50% deliver of utility power to the machine-



Graph on increased power consumption

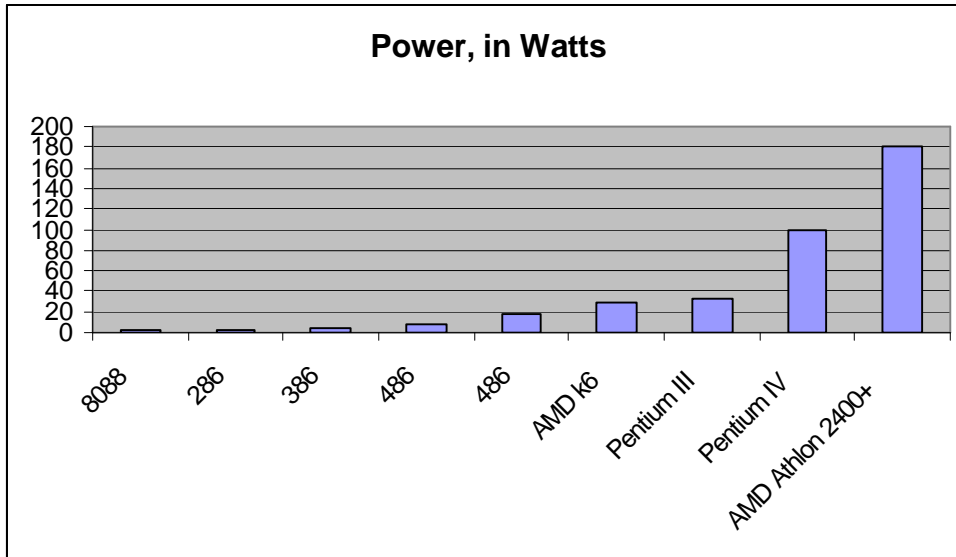


Figure 3 Evolution of CPU power consumption Source Robert Frances Group

This geometric growth in power consumption has led power and cooling to be the biggest impediments today to data center growth. The growth in power has grown so acute, that in many cases the utility company doesn't have enough power to provide the facility, even if they can pay for it.

Additionally, increased power consumption can lead to higher failure rates in servers. A common airflow for cooling comes from the bottom of the data center floor, rising through the servers until the heated air comes out the top of the.

The problem with large density systems such as racks is that, by the time the air reaches the servers in the top of the rack, the air is no longer cool. This directly causes significantly higher failure rates to systems in the top half of a server rack. Best practice is to have computing systems vent air from the front to the back of the computing system delivering fresh cool air to each server, regardless of location in the rack. Even with front to back cooling proper air flow practices such as employing blanking plates and hot aisle/cold aisle should be used.

The alternatives to smartly designed data center where air delivery is uniform is to place special-purpose air cooling systems immediately adjacent to these racks or to decrease the density of the racks to allow the cooling capacity not to be exceeded. In either case, the net result is a decrease in available data center space for servers. RFG has found many clients, especially those with high density server utilization, that are not able to use more than 50% of the available floor space. An alternative to adding floor space is to employ directly-attached special purpose cooling systems, such as the IBM Rear Door Heat eXchanger, which can remove up to 55 percent of the server heat. While products like IBM's rack-based heat exchanger do not do away with heat problems they can act as



important tools for customers looking to create denser solutions than their computer room might allow.

Critical trend shift

The combination of dramatic computation ability and the massive increase in power consumption has led to the following dramatic key shifts:

- *The CPU has moved from a critical differentiator to a fungible commodity*
- *Power consumption have shifted from being a commodity to a critical resource constraint*

IT executives need to shift their priorities from picking CPUs and servers that have the best processing performance to those that can process work with the least amount of energy. Current energy consumption in large data centers can approach 1/3 of the total IT budget. The good news for IT executives is that selecting more efficient power designs not only has direct energy reduction benefits, but that benefits can extend in other, equally compelling areas throughout the data center.

Direct Benefits of Power Efficient Blades

Since RFG suggests that IT executives need to focus on power efficiency as a critical design criteria, it makes sense to review what benefits can be gained from looking at systems that focus on this efficiency. RFG believes there are both direct and indirect benefits. The direct benefits can be shown as a result of the decrease in power consumption of the chips. Indirect benefits are just as significant, including increased system capacity, reduced management costs, reduced system failure rates, and reduced facilities costs.

We will examine each of these benefits in turn.

Reduced power consumption – lower energy costs

The most directly measurable benefit of focusing on power efficient servers is lower energy costs. As mentioned previously, a fully – populated large data center can generate large costs. RFG has found that power consumption for a large data center can exceed two million dollars a year. Systems that can increase power efficiency by 20% will yield a \$400,000 annual cost savings. Current designs shown by some systems vendors show that these savings are achievable.

Power efficient systems designs should not only focus on power-efficient components like processors, but also look at issues such as efficient power supply conversion,



improved thermal characteristics, and better system power management. IT executives should press chip and systems manufacturers to focus on these areas of efficiency, and show development plans that incorporate such thinking.

Reduced Cooling requirement – lower cooling costs, which itself reduces power consumption

As mentioned earlier, power and cooling can run in a vicious cycle: more power leads to more cooling, which itself requires more power. The converse is also true: more efficient systems design require less power, which requires less cooling. Less cooling itself then requires less power. Most cooling systems require approximately 50% of the power the IT systems consume. Therefore a reduction of 20% of the energy costs for data center would be:

	System power costs	Cooling Power Costs	Total Power Costs
Standard	\$2,000,000	\$1,000,000	\$3,000,000
Power efficient	\$1,600,000	\$800,000	\$2,400,000
Savings	\$400,000	\$200,000	\$600,000

Table 1 Annual Power cost savings for large data center

Source, Robert Frances Group

Additional Benefits of Power Efficient Blades

Beyond the direct (significant) cost savings for power efficient design, several significant benefits accrue that may not be immediately noticeable.

Additional space available in data center

Because of current power limitations, dense servers (from a power consumption perspective) lead IT managers to leave racks empty. RFG has talked with several IT executives that run large data centers that have reached power capacity when the center only had 50% of the available rack space utilized. In 2000, when an IT company bought 100 servers, the data center manager would ask how many U each server was to calculate the number of racks needed to house them. Today, that question is far more likely to be how many watts does each server draw? Today, power and cooling limit density far more often than real estate space in the rack.

If IT managers can deploy more efficient blade servers, better power utilization will then allow additional servers to be installed in the rack, dramatically increasing the number of servers per rack. More compute power for every kilowatt of energy is really the end game. IT executives can then have the freedom to either add more compute capacity into



the data center, or use this additional space for other elements, such as storage capacity or networking.

Depending on the nature of data flow in the data center, power consumption from networking devices can be significant as well. The additional space savings from more efficient servers can help by giving more airflow headroom to networking devices. In some cases, data centers that were at 50% physical capacity can now be given the ability to add needed compute capacity without expensive infrastructure upgrades.

Reduced facilities costs

As power and cooling consumption goes down with efficient design, facilities management costs also become lower. Reduced cooling equipment reduces the total number of devices for facilities to manage. This can help in allowing further data center consolidation, which can reduce both operational and facilities staff. Additional facilities savings include a reduction in the number of uninterruptible power supplies (UPS) needed, fewer breaker panels, and few power distribution units (PDUs). As large companies are looking to consolidate data centers, especially through merger and acquisition, a 20% gain in efficiency could lead to a reduction by 30% in the overall data center consolidation costs, depending on the existing utilization rates of the centers that are candidates for consolidation. Server consolidation will allow an additional reduction in the amortized hardware maintenance costs (software costs should be similar), which will lead to reduced FTE costs as well as a multiplicative power benefit, since there will be fewer servers to run, which each run at lower power.

Increased server density

At current power consumption of most of the existing blade servers, fully populated racks would need to be able to support over 25 kilowatts per rack. Most customer data centers cannot handle this level of power consumption and the associated heat dispersion requirement. Often, they are limited to putting 15, 10, 7, or as low as 4KW/rack. The only way to get greater power densities with these chips to day is to put the racks in special-purpose "hot-spot" rooms that are designed to handle more concentrated cooling requirements. The problem with this option is that IT ends up trading space for cooling. The problem with trading space for cooling is that many data centers have run out of space, so special purpose cooling facilities are always an expensive, but not always viable, option.

Better power management techniques

Some vendors today are delivering smarter tools for predicting and managing power consumption. These tools are primarily aimed at giving customers better information about their actual power draw so they can make smarter decisions about density, redundancy, cost, etc. Some vendors have taken this to even higher levels with strong power management that can also use this software to reduce the cost of infrastructure that users bring to the rack.



For a customer that requires redundancy, generally they run twice the amount of power to the rack. This is required in case one of the two power sources fails. With power management software like IBM's Power Executive, users can decrease the size of each of the redundant power feeds. During crisis, the software keeps the system up and running without any down time.

More effective space utilization

Customers using blade servers designed for efficient power may not be able to fill the rack completely given current data center limitations. However, they can get more servers, thus processing power, into each rack. This allows existing standard raised-floor cooling topologies to meet the cooling requirements without eating away at often-scarce floor space. If IT executives implement more efficient blade processors, they will be able to get more processing capability per kilowatt of power consumed. IT executives should look for system performance metrics that can demonstrate improved capacity/kW.

Reduced system failure rate and better uptime

A further result of the lower power utilization, reduced and more efficient cooling, is a reduction in system failure rates.

While most systems employ front-to-back cooling, overall system design should be looked at to make all aspects of systems more efficient. For example, many systems in use today still employ a large number of small fans for direct system cooling. Each fan is a point of failure, and a potential source of local hotspots, which further increases the probability of failure. Improved shared thermal solution systems focus on efficiency, with shared thermal management employ fewer, more efficient fans and heat sinks. This both reduces costs, reduces the number of components to fail, while reducing overall power consumption. These systems are also more manageable, since the larger, shared fans can be designed to be hot swappable in the event of failure. Traditional 1U/2U servers and even some poorly designed blades are built using 8-12 non hot swap fans in each server. Each of these fan can consume between 5-10W each. In medium to large solutions this adds up to 100s or even 1000s of power consuming fans, which can draw many kilowatts of power, further exacerbating the power problem.

Additional thought for design with an emphasis on power and cooling approach of some blade servers can lead to less points of failure as well. For example IBM's BladeCenter can use 4 power supplies to redundantly power up to 14 servers and associated switching. To do this in a typical 1U stacked solution would require well over 28 power supplies. The BladeCenter is also a great example of this rethinking in a shared thermal solution-built with two hot swap blowers that consume only about 100W. These two parts replace the 112 fans typical in other servers that consume well over 600W. Thus, the similar cooling function can be performed with both increased reliability and dramatically reduced power.



Even without benefits, you cannot afford to ignore power problem

Even if the IT executives didn't buy all the benefits described, the reality is that they often do not have a choice when it comes to data center growth and cost options. Facility management, not IT, usually owns and controls business real estate. Regardless of what IT executives claim they need to process business requirements, they are often told "they've got what they've got", and that they have to live with the data center real estate they have been given.

Often times, this reality is the IT executives own doing. RFG has talked with several CIOs recently about their data center problems. Many of them had planned on building their new data center based on power requirements they saw at the start of the design process. In 2002, the typical per-rack power consumption was two kilowatts. In 2005, just three years later, the per-rack power consumption for blade servers jumped to 20 kilowatts.

The order-of-magnitude increase caught most every IT executive by surprise, which in turn caught facilities managers and even the utility industry, incapable of responding to the mushroomed growth in power. As a result of 9/11 and the Northeast power outage of 2003, many companies have moved backup data center plans from metropolitan to rural areas. This recent move, in conjunction with the geometric power requirement growth, has left many of the areas with utilities unable to supply "metropolitan" power requirements to these "rural", historically lower power areas.

Additionally, the increased density of servers has led to weight issues. That is, the floor loading is such that it cannot support the weight of a fully populated rack of servers at today's weights, 1 rack of 1U servers can approach 1800 lbs; many older data centers were designed for much lower floor ratings. Not only to blades need to use less power, they also need to be lighter.

RFG has interviewed some CIOs who have even had to generate their own primary power, as the local utility was incapable of guaranteeing the 20 megawatts the data center needed. Even when IT manager work with facility management to generate power, legal then needs to get involved, since most communities have both zoning, sound and environmental considerations that need to be met if the company generates its own power on site.



Case study: Gains that can be achieved by Smart power design

To put some real numbers to the power statements, let's look at a sample study of comparing legacy blade servers with systems that focus on optimizing power consumption.

For this case study, we review power comparison run a one of EMEA's largest business partners for one of their customers weighing two different blade platforms: IBM BladeCenter and HP BladeSystem BL35p:

The general configurations were:

- IBM BladeCenter LS20 with dual processor 2.2Ghz AMD Opteron ; 4 GB memory, two Ultra-320 SCSI HDDs; and dual Ethernet ports,
- HP BladeSystem BL35p with dual processor 2.2Ghz AMD, 4GB memory, two IDE 40GB HDD, and dual Ethernet ports

The company was looking to purchase 4800 blades in 2006.

As a result of choosing the IBM BladeCenter servers, the company would realize the following benefits through a power-optimized configuration:

- A 20% direct power savings for running the servers
- A 20% reduction in costs to cool the computer room
- Reduction in number of racks needed to house the solution from 73 racks with HP to 59 racks with IBM
- Reduction in overall floor space required in the data center from over 1300 ft² with HP down to 1080ft² with IBM.
- Reduction in the total number of fans needed. The solution went from 19,200 for HP down to less than 700 – a reduction of 96% in potentially failing components.
 - Reduced heat from 46.1K BTUs to 31.9K BTUs

Racks needed	BTU	Annual Energy Cost
73	46.1kBTU	\$ 1,098,433.92
59	31.9kBTU	\$ 760,087.68

- With typical energy costs, this leads to an annual energy savings of 31%
- Assuming a 10% increase in energy costs, and adding the 4800 servers planned



	New Servers	Racks needed	BTU	Total BTU	Projected Energy cost 2006
Old	4800	18	11.4kBTU	57.5kBTU	\$1,507,070.40
New	4800	15	8.1kBTU	40kBTU	\$1,048,396.80
Projected Cost Savings 2006					\$458,673.60

Shows a predicted energy cost savings with power efficient design of \$458, 673 for 2006.

- o Efficient power supplies allowed the company to better use the power they had. RFG finds that this is becoming more important, especially as companies are increasingly gated by the power utility's ability to supply power.

Conclusion

RFG believes power will be the number one issue for most large company IT executives to address in the next 2-4 years. Ignoring this issue will NOT be an option. Power consideration must be incorporated into data center planning. This includes close cooperation with facilities management, intelligent reduction of unneeded applications, systems, and infrastructure, and reevaluating system design with an emphasis on power efficiency as a fundamental design criterion.

Power considerations must be incorporated into the details of data center planning, realistic; conservative planning assumptions must be made in predicting future requirements.

Power supply efficiency allows a vendor to take any set of commodity parts and run them at a lower cost. No matter the amount of memory, type of processor, number of hard disks, a vendor that has rethought how power is used can make a huge difference in the total power required in the data center. In keeping with this, vendors need to provide greater levels of choice for lower power processor support. Not all customers require the fastest system; some just require new capacity at as little power cost as possible.

Use of less power means less heat introduced to the data center. This, in turn, means less air movement inside the server, which allows the systems to run on less power, which in turn can allow the system to run more quietly. This reduced cooling can provide increased efficiencies for fans as well. A smart blade design could reduce the number of fans or blowers to just a handful, replacing the 100s needed for typical stand alone servers- again leading to lower power for moving air.

For older data centers, efficiency in power and cooling allow you to extract the most from each and every kilowatt. For those fortunate enough to be designing new data centers,



the cost savings for choosing efficient components can be massive: smaller UPS, fewer power distribution units, fewer breaker boxes. Additionally, smaller air handlers and AC units are needed, as well as less floor space, since each rack can house more processing power.

Finally, efficient chip and server design will become the key criteria for selecting data center blade servers. IT executives should look for these factors as key selection criteria, and systems designs should push aggressively for chip and server manufacturers to offer more efficient designs, including power management and planning into the sales lifecycle.

All rights reserved. 120 Post Road West, Suite 201, Westport, CT 06880. Telephone 203/429-8950. Facsimile 203/429-8930 www.rfgonline.com. This publication and all publications may not be reproduced in any form or by any electronic or mechanical means without prior written permission. The information and materials presented herein represent to the best of our knowledge true and accurate information as of date of publication. It nevertheless is being provided on an "as is" basis. Reprints are available for a nominal fee.