


**Managing Energy Efficiency
in the Data Center**

A tactical approach can deliver
immediate problem resolution plus
long-term savings and efficiencies



Managing Energy Efficiency in the Data Center

Executive summary

Power, cooling, and floor space are top-of-mind issues these days for IT and facilities executives and staff struggling with dwindling resources, increasing costs, growing demands from the business, and concern for the environment. Adding new hardware and growing the data center to accommodate individual business needs and requests is no longer a viable response. Organizations must find ways to manage today's energy demands against ever-growing constraints—as well as have a plan in place for the future.

However, the answer isn't as simple as upgrading to newer technology or consolidating storage. Today's disk drives and servers may be smaller, but they also draw higher power—and even with an upgrade to new platforms, some pieces of legacy equipment typically must remain in operation. While virtualization and consolidation reduce multiple storage footprints into a fewer frames and servers—freeing up available physical space—power and cooling still remain a problem.

EMC, like many of our customers, was concerned about the ability of our data centers to operate efficiently and effectively—from both a cost and environmental standpoint—as well as to support continued, rapid growth. We, too, had historically expanded the footprint of our main data center to accommodate growth and customer demands—a practice made even easier because we could supply our own product.

We had already realized dramatic savings from extensive storage tiering and consolidation (conducted as part of a comprehensive, five-year information lifecycle management (ILM) initiative), which had significantly reduced the number of physical storage frames in our main data centers. However, by early 2007, our main data center was once again running out of power and was experiencing cooling distribution issues as well as hot spots. Coincidentally, we had approximately 800 servers in that environment approaching end of service life (EOSL) that were going to need to be replaced or virtualized.

We knew that we lacked the knowledge, tools, and resources within the data center to address this urgent problem quickly and effectively, and made the decision to engage EMC® Consulting Services to help us. We believed that the investment in energy efficiency specialists and services would speed time-to-results and enable us to successfully address our immediate power and cooling issues as well as assist us in achieving our longer-term energy efficiency objectives by:

- Identifying and implementing thermal improvements in our power distribution chain, IT consumption, and cooling infrastructure
- Implementing storage management policies, procedures, and technologies within our production IT environment to reduce costs and enhance scalability and flexibility
- Assessing opportunities for additional virtualization and server consolidation and developing an implementation roadmap

We not only achieved our business/energy efficiency objectives, but we also exceeded our expected return on investment. This paper highlights the results realized and lessons learned from this effort, and provides considerations and guidelines for organizations facing similar challenges.

What's in your data center?

Establishing a baseline

Before you can resolve existing energy efficiency issues—much less plan for the future—you need to thoroughly understand the current state of your data center. This includes knowing what hardware you have, where/how that hardware is physically located, how these systems are powered and cooled, how much power and cooling they demand, how efficiently they are operating, and other physical and operational details.

To establish this baseline, the EMC Consulting team conducted an audit and physical assessment of our main data center (a week and a half in duration) that began with gathering and/or validating known inventory data about our servers, storage, and energy components (UPS, power, and cooling distribution components) as well as facility space. Floor plans, online drawings, cooling system diagrams or mechanical prints, and current listings of IT devices in the data center were instrumental to this discovery phase.

EMC Consulting also worked with us to document our current and future data center and IT asset requirements related to servers, storage, floor space, power, and cooling, based on our existing “business as usual” (BAU) infrastructures and policies. They also validated the assumptions we used to build our storage and server consolidation strategies and our energy efficiency strategies, including—but not limited to—our energy component, server, and storage inventory; data center, power, server, and storage utilization; three-year growth related to the data center, power, and infrastructure components; and cost of our IT infrastructure, power, and cooling requirements.

A comprehensive power, thermal, and cooling efficiency analysis followed. Power energy efficiency specialists performed a physical inspection of the data center infrastructure (e.g., hot and cold aisles, security of panels); measured electrical infrastructure loads; and conducted a power usage effectiveness (PUE) analysis¹; cooling and space utilization analysis; thermal imaging; and a computational fluid dynamic (CFD) analysis².

Analyzing the results

The measurements and findings obtained during the data center energy efficiency assessment were analyzed in order to formulate recommendations to achieve increased operational efficiencies and effectiveness and to generate a detailed plan for improving power distribution and cooling efficiencies. The assessment team used a “power calculator”—specialized software—to determine actual power consumption for a possible technical refresh, and employed a differential financial model to validate power costs, replacement cost data for IT assets, and data center square footage cost. The team also compared the BAU costs to the optimized costs over a three-year timeframe and prioritized recommendations based on business and financial criteria.

Weighing tradeoffs with solution options

The results of the analysis, delivered via a score card, confirmed that although our main data center power distribution systems had available capacity to support additional load, there was minimal bulk power capacity to support future growth. Our cooling distribution system added further limitations.

Two options for addressing these problems were proposed. The first included increasing UPS capacity to support future growth, instituting a systematic approach to efficiently distributing power throughout the data center; increasing the number of PDUs in the data center; replacing existing rack-mounted power strips with metered power strips to monitor power usage at the rack level; removing non-critical loads from the UPS distribution infrastructure; and installing lighting controls to modulate lights depending on occupancy. For cooling, it was recommended that we position racks in a hot aisle/cold aisle arrangement according to industry standards to improve the cooling efficiency of the CACUs, minimizing gaps between racks, installing blanking panels and air restriction devices where appropriate, adding floor tiles to utilize the full volume of air being delivered, and calibrating the temperature and humidity probes on the CACUs to ensure proper operation.

The second option recommended adding separate dedicated power metering for the data center to determine power usage, and utilizing a ceiling plenum as a return plenum for direct hot air return to the CACU units to improve the effectiveness and to minimize the mixing of hot and cold air streams within the data center. Strategically placing return air grills would increase unit efficiency and reduce mixing in the data center as well as contribute to lower equipment inlet temperatures. This option would provide additional bulk capacity for future load requirements, enable us to support a higher rack heat load density, and provide a method for determining actual power usage for the data and associated appurtenances. However, it would require significant capital investment and cause interruption to data center operation.

The analysis also identified additional opportunities to improve our data center according to industry best practices—for example, rerouting cables, installing occupancy sensors to minimize lighting costs during minimal occupancy times—and included practical strategies for reducing electrical power consumption in the data center, with detailed estimated savings, guidelines, and limitations outlined.

The power system analysis produces recommendations to reduce current operating expense and defer or lower future capital expense, looking at input and output power, IT load and input power factors, harmonics measurements, UPS load factor, efficiency, and energy calculations.

The cooling infrastructure analysis looks at the bulk cooling and distributions of the as-built data center and analyzes the ability to support the existing data center loads and available capacity for future deployment.

The thermal analysis produces recommendations to reduce current operating expense and defer or lower future capital expense by optimizing placement of equipment and evaluating efficiencies in the air flow, based on pressure analysis, tile flow measurements, and heat removal calculations.

The CFD analysis, thermal imaging, tile air flow measurements, heat removal calculations, and field measurements are gathered to optimize temperature management and heat dissipation. The analysis involves a model that simulates the data center cooling infrastructure and loads, enabling consultants to run different scenarios and compare results.

¹Power usage effectiveness analysis is the corresponding load on the data center in power, cooling, lights, etc. for every watt of power your IT equipment uses.

²Computational fluid dynamics generates a “weather map” of hot spots in the data center, where they are, and what can be done to address them.

Results: immediate and long-term operational and business benefits

We chose to go with option 1 and the ceiling plenum from option 2, which we determined to be the most efficient and cost-effective short-term solution. Implementing the recommended thermal improvements and improving efficiencies in the data center's power distribution chain, IT consumption, and cooling not only resolved our existing problems but **extended the life of the data center by more than 15 months—a significant cost avoidance**—through follow-on application alignment, server virtualization, and data migration.

Using VMware® virtualization technology to reduce and consolidate the environment, we were able to virtualize 1,357 servers onto 231 physical machines, thereby eliminating the infrastructure needed to support 1,126 servers. This resulted in an infrastructure (space, power, and cooling) savings of over \$4.6 million over five years—a 67 percent reduction. Furthermore, we avoided the significant costs associated with the purchase and maintenance of these 1,126 servers as well as the network equipment needed (over 2,000 production and network ports). Additionally, nearly 3,159,726 pounds of CO₂ emissions will be eliminated over five years.

Operationally, data center administrative personnel are now able to manage nearly three times the storage capacity thanks to consolidation efficiencies, and power and infrastructure costs are being managed more efficiently. In fact, the \$700,000+ infrastructure savings do not include additional savings from cancelled maintenance contracts and reduced administrative costs from the elimination of 165 EOSL servers.

Virtualization has enabled EMC to address immediate server and life support issues, avoid a tech refresh, and has extended the life of the main data center.

After adding a ceiling plenum as recommended in option 2 in late 2007, we were able to reduce the temperature in our main data center by a measurable average of 3.4 degrees F, further improving our cooling efficiency, extending the life of our HVAC capacity, and reducing CO₂ emissions that otherwise would have been required to cool that additional 3.4 degrees.

While we are again looking at expanding our main data center to sustain present and future organic growth, all of the efficiency measures we deployed—including consolidation and virtualization—allowed us to put off an expensive data center upgrade for three years—extending the life of our main data center from 2005 to 2008 and resulting in a \$13 million cost avoidance.

Lessons learned: planning to get the most value from your data center assessment

If you are considering undertaking an assessment of your data center, it's important to understand the associated resource, time, and skill requirements. The following should be taken into consideration when preparing to undertake an assessment of your data center.

Staffing—Certain data center staff may have knowledge of your operations and your business that is helpful and perhaps essential to the success of your project. These individuals—who are dedicated to the daily operations of your facilities—must be prepared and available to work with consultants to provide needed information upon request.

A project sponsor should be assigned as a single point of contact for issue resolution, activity scheduling, interview scheduling, and information collection and dissemination. In addition, a project manager/contact should be assigned to work with the consulting team's project manager. This individual should have a concrete knowledge of the current data center and its power, storage, and server requirements, as well as be able to identify appropriate support personnel whose services or knowledge may be required.

A complete energy efficiency initiative—including the planning for and execution of the recommendations delivered by the assessment team—is a major cross-functional team effort involving network and systems architects and run teams, consolidation and virtualization specialists, and other professional services and customer support experts to architect the recommended changes, perform installation, integration consolidation, migrations, and decommission old environments. Our data center initiative ultimately involved several hundred individuals from across the organization.

Knowledge/skill/time—Power, cooling, and mechanical engineers are required to perform a physical assessment of a data center. Most data centers do not have the knowledgebase and the specialized tools—as well as the time—to determine computational fluid dynamics, take individual power readings, and do a CRD analysis. Following the physical assessment, virtualization specialists can help you plan the virtualization and consolidation that will be driven by the assessment, and information infrastructure specialists/strategic consultants can help you build and implement an information infrastructure management strategy.

Documentation—A good understanding of your physical/storage environment and system performance is required before beginning a physical assessment—e.g., rack power readings, performance of individual servers, and tiering/SRDF[®] requirements for data. The success of the physical assessment, and the speed with which it can be delivered, is dependent on the amount of detailed/documented information—e.g., floor plans, a current list of the IT devices in the data center, online drawings, cooling system diagrams or mechanical prints, power readings of the racks, and performance data on individual servers—that can be made available to the assessment team. If existing documentation is not sufficient, your staff or assessment team resources will need to gather this information before the assessment can begin.

Duration—A data center assessment can span weeks or months, depending on the available baseline data and the size and complexity of the data center. (Because our assessment was being driven by looming EOSL deadlines, we performed our physical assessment over a week and a half.) Other contributing factors included availability of system admin staff to provide assets and configurations of systems in the data center as well as utilization stats, availability of data center staff to work with electricians and power specialists to measure electrical consumption in the data center, and availability of data center staff to participate in interviews on an as-needed basis to help the assessment team understand usage patterns, service-level agreements, and other required information. Remember, too, that continuous improvement follow-on activities such as server virtualization and application migration can extend over several years depending on availability of staff and available windows for downtime.

Obtaining business buy-in

Any major data center initiative requires executive-level buy-in, especially from the CIO. However, this can be particularly challenging when requesting support for an energy efficiency initiative; many IT executives simply cannot justify making changes if the data center is “working for them.” They are less concerned about power, space, and cooling, and about gaining efficiencies or reducing costs than they are about their primary goal—keeping servers up and running to meet application uptime and ensure that users can do what they need to do as quickly as possible. Infrastructure is simply not a priority.

Executive buy-in can best be gained by demonstrating the potential cost savings and partnering with the facilities organization is a particularly powerful way to drive this message home. Because facilities is responsible for the expense of running a data center (and, thus, to drive costs out while supporting the requirements of the business), they will be highly motivated to champion an energy efficiency initiative because of the millions of dollars of savings in power costs—as well as the cost avoidance of expanding data centers—that can be realized. Compelling arguments can also be made for how this initiative can help the organization address known pain points such as performance, turnaround time, and functionality and efficiency in the data center.

Can you benefit from a data center assessment?

If your data center is out of space, power, cooling, or any combination of the three—or if you are in the predicament of needing to do a significant capital expenditure to upgrade your facility or build a new data center—an energy efficiency assessment of your data center can help you:

- Reduce demand for power and cooling and free up floor space
- Lengthen the current usable lifetime of your facility
- Reduce operating expenses and avoid capital expenditures
- Increase operational efficiency and utilization
- Reduce downtime
- Deploy new applications
- Upgrade/replace older, power-hungry hardware
- Adopt newer technology paradigms
- Mitigate risk
- Go green

A data center assessment can take many paths, from resolving immediate pain points to a new infrastructure strategy for how your company uses its IT assets to tiering of storage and to consolidation and virtualization (with the biggest benefits ultimately realized from consolidation). A successful data center assessment—while targeting energy efficiency issues—can help you build and maintain a modern, energy- and cost-efficient and effective data center infrastructure that can continue to support your business in the future.

Even if you do not consider energy efficiency in the data center to be a challenge today, it **will** be at some point in the future. Our experience shows that it is easier to be proactive by understanding your current data center now—and learning what you can do to prepare for and influence the future—than wait until you need to take immediate action to address an urgent problem.



EMC Corporation
Hopkinton
Massachusetts
01748-9103
1-508-435-1000
In North America 1-866-464-7381
www.EMC.com