LAB AUDIT REPORT

EMC IT

A Blueprint for Data Center Efficiency

By Brian Garrett
With Mark Peters

April, 2009

Table of Contents

ESG Lab Reports

The goal of ESG Lab reports is to educate IT professionals about emerging technologies and products in the storage, data management and information security industries. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab report was sponsored by EMC.

Executive Summary

In 2004, EMC faced explosive growth in applications, servers, and storage in its five worldwide data centers. With space, power, and cooling reaching the limits of the current infrastructure, EMC's IT organization was faced with the prospect of an expensive data center upgrade. This report explores the results of three continuing IT initiatives that have enabled EMC to deliver increased efficiency and reduced costs—without a data center upgrade.

1. **Data Center Efficiency**: Improve power, cooling, and space efficiency of existing data centers
2. **Storage Efficiency**: Move from two tiers of storage to five using software and automation to reduce costs
3. **Server Efficiency**: Use server virtualization to consolidate the number of servers in the data center

The bottom line benefits EMC achieved over the past five years can be used as a practical blueprint for organizations looking to improve the efficiency, effectiveness, and environmental sustainability of their existing IT infrastructure by leveraging existing technologies and best practices. The keys to success began with a strategic vision, sound business planning processes, executive buy-in at the highest level, and EMC's own professional services engagements. Continuous improvements in data center air flow and cooling played a vital role in increasing the density—and effectiveness—of existing data centers. And finally, hardware, software, and process improvements were used to increase the efficiency of storage and server infrastructure.

The balance of this report describes the processes and technologies used to increase the efficiency and effectiveness of EMC's IT infrastructure, which so far (the five years from 2004-2008) has resulted in:

- $74 million in data center, storage, and server savings
- $12.5 million in data center power, cooling, and space savings
- Extending the useful life of a legacy data center
- Increasing data center energy efficiency by 34%
- Increasing storage utilization by 19%
- Improving the productivity of storage administrators by 170%
- Reducing server footprint by 60% during the first phase of virtualization and consolidation

FIGURE 1. EMC IT EFFICIENCY SAVINGS SUMMARY (2004-2008)

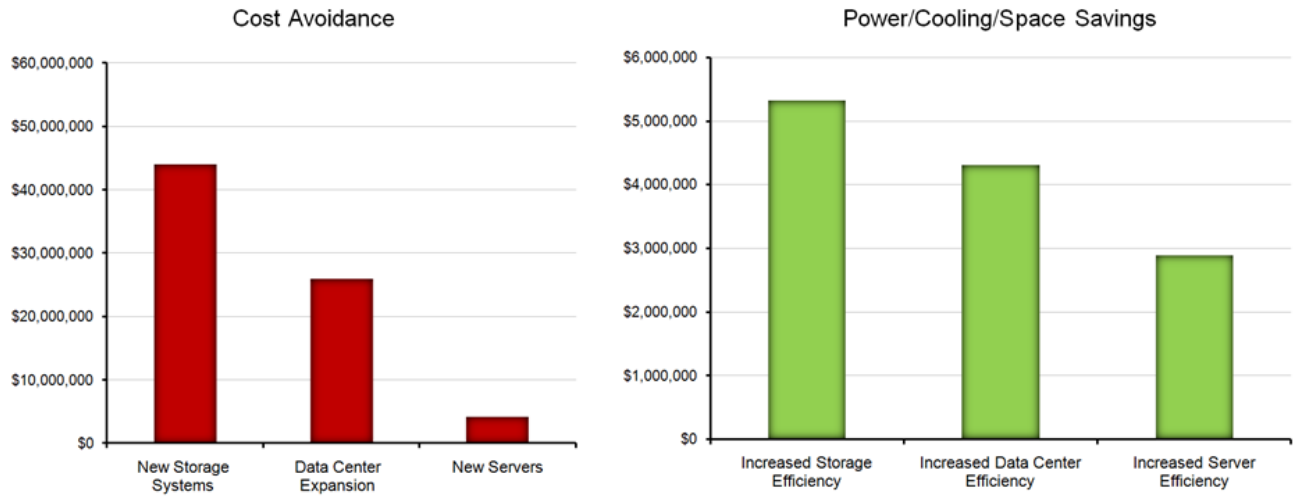

ESG Lab IT Audit

The EMC IT efficiency efforts examined in this report began in 2004. At that time, EMC had five data centers located around the world. A "one size fits most" philosophy was being used for the routine deployment of new applications, servers, and storage systems. Storage capacity was growing at an annual rate of 60%, driven by a combination of organic growth and acquisitions. A 30 year old core data center in Westborough, Massachusetts was running out of space, power, and cooling.

The cost of a new data center with state of the art energy and cooling efficiency was estimated at more than $120M. A $26M data center upgrade was being considered to accommodate future growth. A business planning process began and EMC professional services were engaged with a goal of determining the steps that could be taken to improve the efficiency of existing data centers and IT infrastructure.

Data Center Efficiency

As much as half—and sometimes more—of the electricity purchased from a utility company can be lost due to the inefficiency of power and cooling systems in the data center. As a result, optimizing the delivery of power and cool air in the data center can lead to significant savings.

To defer a $26M data center upgrade, EMC created a multi-year strategy to increase space and energy efficiency, as well as operational effectiveness, in its existing data centers. Significant efficiency gains have been realized from data center cooling and airflow improvements, including:

- Hot and cold aisle design with a goal of eliminating hot spots
- Hot air return plenum for a more efficient removal of heated air
- More efficient Computer Room Air Conditioning (CRAC) units
- Filler panels over rack units with no equipment installed
- Floor "pillows" to reduce wasted air flow around cables
- Selective in-row cooling
- Ultrasonic humidification
- Monthly CFD analysis to identify and rectify hot spots with vented floor tiles

ESG Lab examined the Data Center Infrastructure Efficiency, or DCIE, at a primary EMC data center to quantify the gains realized to date. The results have been impressive.

The Green Grid, an industry association of IT professionals, defines DCIE as total IT equipment power (servers, storage, switches, etc.) divided by total facility power (as delivered by the utility company).[1]

A DCIE of 50% indicates that half of the power delivered to the data center is used to actually run IT equipment in the data center. A theoretically impossible DCIE of 100% would indicate that the data center is perfectly efficient with no loss in electricity. The Uptime Institute estimates that a typical data center has a DCIE of 40% and a DCIE of 62.5% is achievable with generally available technology.[2]

Looking Forward to Smarts

EMC's data center efficiency efforts are continuously evolving and improving. For example, EMC Smarts software is at the center of an effort that is planned for production deployment in 2009. EMC is integrating monitoring equipment from Johnson Controls with EMC Smarts software. The Smarts correlation engine will help improve the categorization and prioritization of building and infrastructure events.

As shown in Table 1, EMC's DCIE has improved from 50% to 67% over the past five years. Better cooling efficiency allowed more power and equipment to be deployed over each floor tile in the data center. As a matter of fact, EMC began at 60 watts per square foot in 2004 and is currently at 100 watts per square foot.

[1] DCiE can also be defined as 1 over Power Usage Efficiency (PUE)

[2] www.uptimeinstitue.org

Power and cooling efficiency improvements enabled the storage and server consolidation efforts described later in this report. The cost of power and cooling savings before and after EMC's data center efficiency efforts are itemized in Table 1 and depicted in Figure 2. A cost of commercial power of $0.10 per kilowatt was assumed for the cost of power and cooling calculations.[3]

TABLE 1. THE BOTTOM LINE FOR THE FIRST 5 YEARS

	BEFORE (Q1'04)	AFTER (Q4'08)	SAVINGS
Efficiency (DCIE)	50%	67%	34% more efficient
Watts per Square Foot	60 watts	100 watts	66% more power density
Power	$3.7M	$1.2M	**$2.6M**
Cooling	$2.5M	$0.8M	**$1.7M**

FIGURE 2. DATA CENTER INFRASTRUCTURE COSTS (POWER AND COOLING EFFICIENCY)

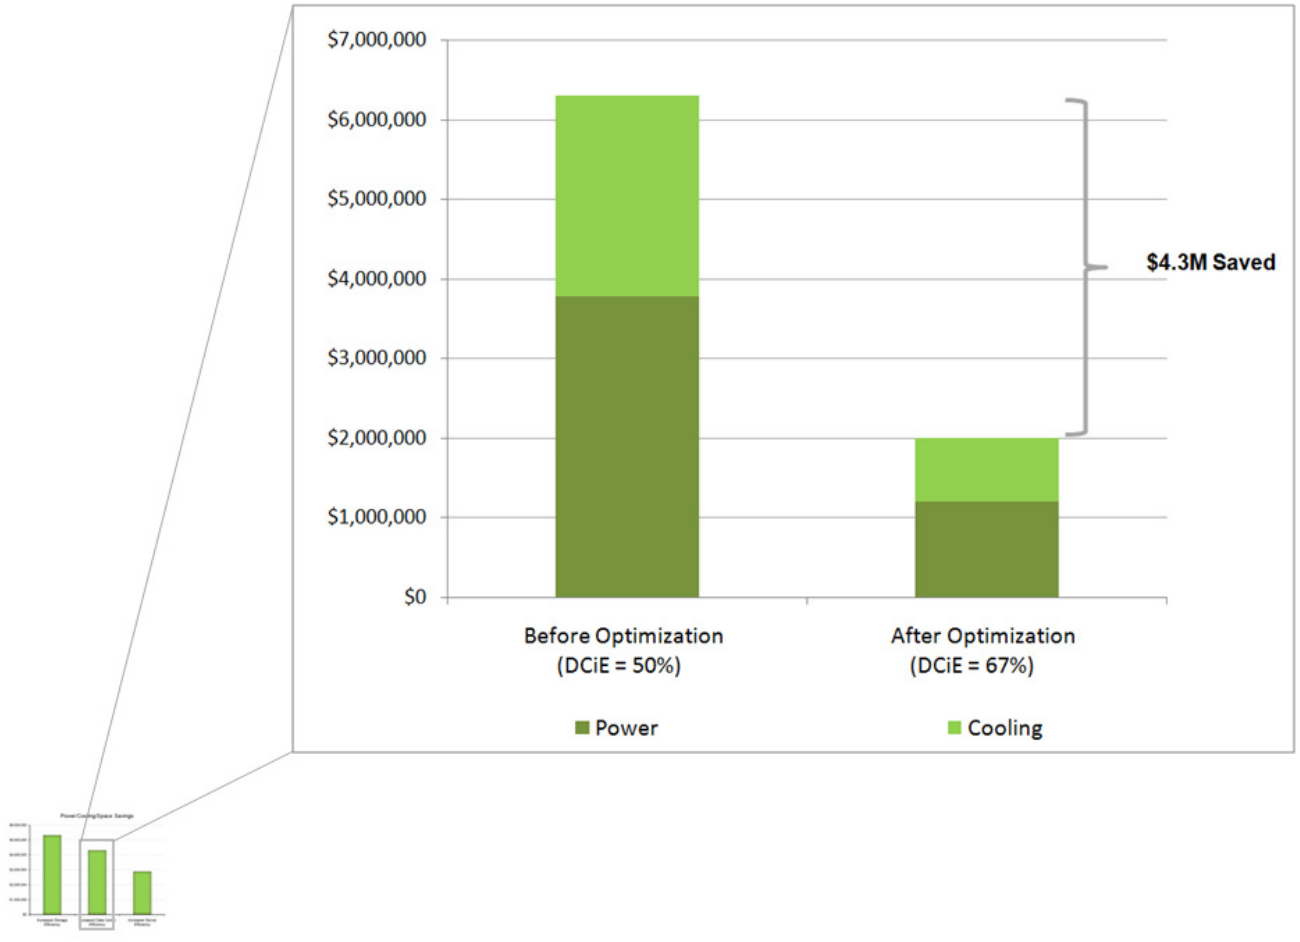

[3] According to the US department of energy, the average commercial retail price of power was $0.1049 in Oct 2008. *www.eia.doe.gov/cneaf/electricity/epm/epm_sum.html*

Why This Matters

Organizations of all sizes are running out of power, cooling, and space in the data center. As a matter of fact, 75% of respondents to a recent ESG survey report that electric power consumption and sourcing is a risk or constraint impacting the business. Sixty-five percent report a current or planned initiative with a goal of improving the efficiency of IT facilities, including data center power and cooling infrastructure.[4]

Increasing the efficiency of power and cooling in the data center can have a ripple effect throughout an IT organization. As EMC has proven, improving data center efficiency can be used to extend the useful life of an existing data center, avoid the cost of data center expansion, and consolidate more servers and storage per floor tile—all while reducing the ongoing costs of power and cooling by 50%.

Data center efficiency efforts can also have a significant impact on greenhouse emissions and global sustainability. Our calculations indicate that the combined effect of the EMC's IT efficiency efforts presented in this report have reduced emissions by 60 million pounds of carbon dioxide. That's equivalent to the carbon offset produced by 67 thousand trees over 27 years.

Storage Efficiency

When EMC's IT efficiency efforts began back in 2004, the majority of EMC's online IT assets resided on high-end EMC Symmetrix storage arrays (Tier 1). The balance resided on a second tier of modular EMC CLARiiON arrays (Tier 2). A tiered storage infrastructure initiative was designed with a goal of reducing costs and extending the life of existing data centers. A business case was built with a goal of expanding the existing two tier storage infrastructure. More cost effective, large capacity drives within Symmetrix and CLARiiON storage arrays were introduced over time to create third and fourth tiers of online storage. A fifth tier of CLARiiON disk library (CDL) capacity was added for disk-based backups.

EMC professional services and a wide variety of EMC software products have been used to increase the efficiency of EMC's tiered storage infrastructure. EMC EmailXtender is a notable example, automatically moving e-mail to more cost effective tiers of storage over time.[5] This also expanded the size of end-users' mailboxes, which improved end-user productivity and morale, centralized management, and made it easier to respond to legal and regulatory discovery requests. Additional EMC software products that have played a key role in increasing the efficiency of EMC's tiered storage infrastructure include EMC ControlCenter, Avamar, Smarts, and Documentum.

Moving to a tiered storage infrastructure has dramatically reduced the power consumption—and cost—of delivering information services at EMC. Consider, for example, the difference between a legacy Tier 1 FC drive and a high capacity Tier 4 SATA drive. The Tier 1 drive consumes more than seven times the amount of power than the Tier 4 drive ($774 per year vs. $100 per year), because it is designed to emphasize performance over capacity. Of course, for a given amount of data you'd also need more Tier 1 drives than Tier 4 drives which would actually exacerbate the situation. However, assuming that suitable non-performance-critical data is moved, these savings can add up quickly as data is migrated to more cost effective tiers of storage over time. Moving to dense, high capacity drives not only reduces the ongoing costs of power and cooling, but it also reduces space requirements in the data center.

Leveraging Flash

In 2009 EMC (as well as some users) will start to realize great performance and power saving benefits by using Enterprise Flash Drives (EFDs) – initially for its Microsoft Exchange environment, then extending to several others. For performance-oriented applications, the judicious use of EFDs can dramatically reduce power, cooling, and space requirements. A recent ESG Lab study found that 6 EFDs outperformed 150 FC drives for a performance sensitive online database application.[6] An analysis of the cost of acquisition, maintenance, power, and space showed the EFD configuration to be 88% more cost effective.

[4] Source: ESG Research Report, *Global Green IT Priorities*, November 2008.

[5] Throughout Q109 EMC has been migrating mailboxes to its latest generation email archiving solution called SourceOne. EMC expects the entire organization to be on SourceOne by the end of the second quarter.

The progress that EMC has made since the tiered storage infrastructure effort began in 2004 is summarized in Figure 3. The savings began to take effect in 2005 as Tier 3 and 4 capacity was deployed in earnest. Archiving efforts, which enabled better use of Tiers 3 and 4, began to take effect in 2006. Forty-three percent of EMC IT's worldwide storage capacity resided on Tier 3 and 4 at the end of 2008.

FIGURE 3. INCREASING EFFICIENCY WITH TIERED STORAGE AND ARCHIVING

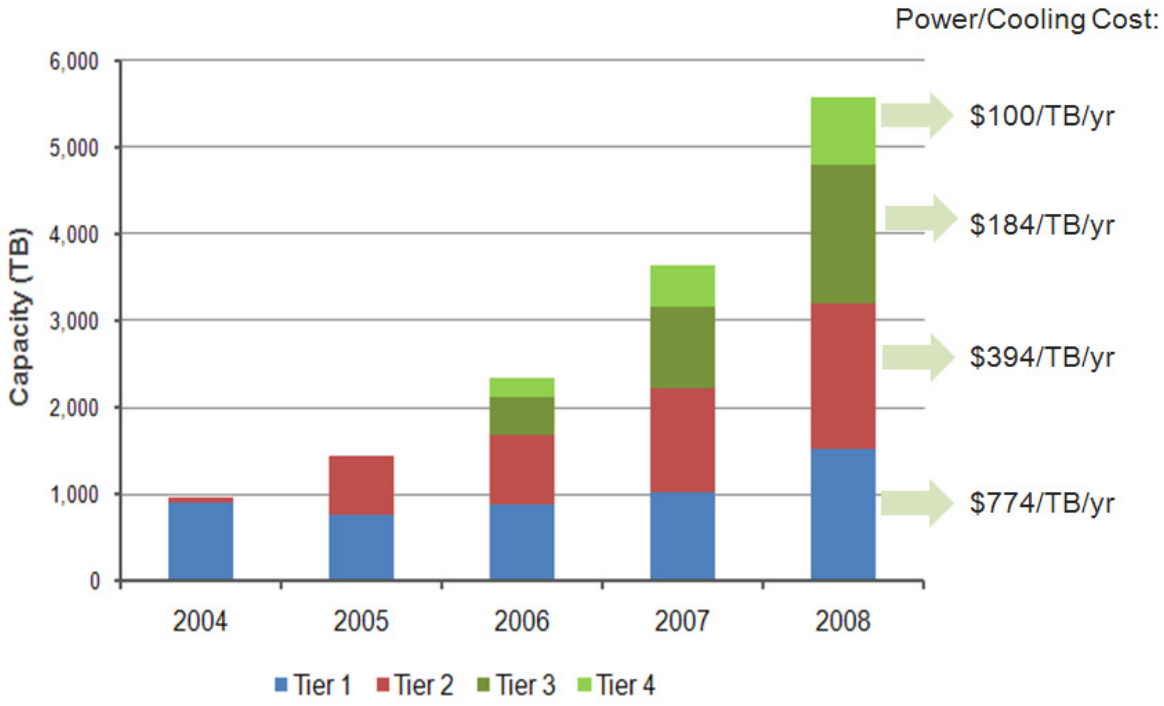

The benefits of moving from two static tiers of FC drives to a tiered storage infrastructure with a mix of FC and cost effective SATA drives are summarized in Table 2.

TABLE 2. THE BOTTOM LINE FOR THE FIRST 5 YEARS

	BEFORE (Q1'04)	AFTER (Q4'08)	Savings
Capacity	962 TB	5,613 TB	
Tiers	2	5	
Utilization	50%	69%	19%
Power	$5.8M	$2.6M	**$3.2M**
Cooling	$3.2M	$1.4M	**$1.7M**
Space	$0.6M	$0.3M	**$0.3M**

With all the detail contained in Table 2, it is worth restating the key result achieved here: *EMC experienced a greater than 500% capacity increase, and yet still saved over $5M (or 54%) of its original power, cooling and space running costs.*

In addition to the power, cooling, and space savings that have translated to millions of dollars of savings over the past five years, EMC has increased storage utilization by 19%. Storage utilization is the ratio of used vs. deployed storage. Utilization rates of 50% or less are common across the industry for a number of reasons,

[6] Source: ESG Lab Validation Report, *EMC CLARiiON CX4*, December 2008

including capacity set aside for future growth and capacity assigned to retired applications. While increasing utilization is a great goal for any organization because it reduces capital equipment and ongoing costs, it can be extremely difficult in practice. EMC was able to achieve these dramatic utilization improvements with the help of management buy-in at the very top of the organization. As part of the information efficiency effort, each business unit owner met with IT to set performance, availability, and utilization goals on an application-by-application basis.

As shown in Figure 4, an ESG Lab audit of EMC IT's storage infrastructure has confirmed that reducing space, power, and cooling requirements by 56% has saved $5.3M in ongoing infrastructure costs.

FIGURE 4. DATA CENTER INFRASTRUCTURE COSTS (STORAGE SYSTEMS)

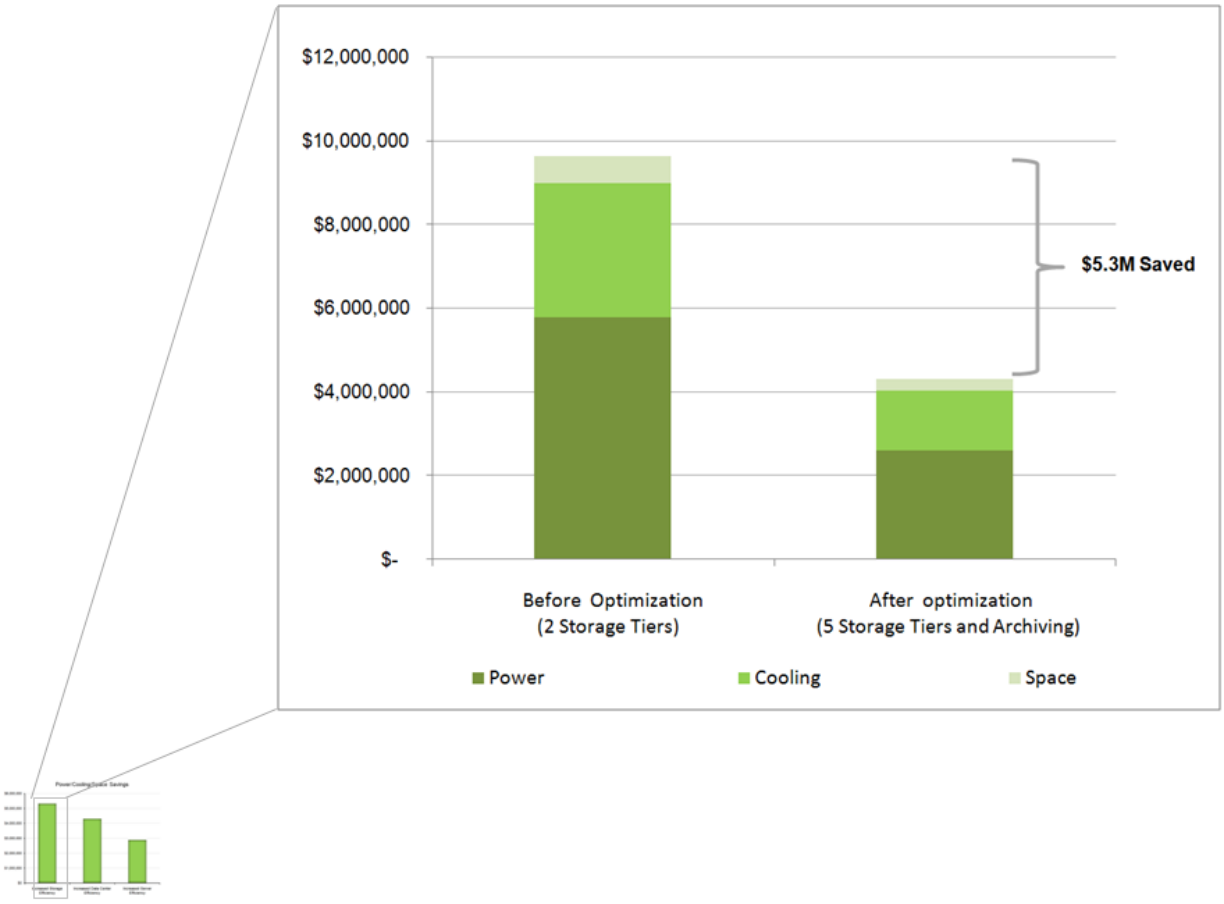

EMC's five year storage tiering initiative has demonstrated a number of benefits beyond the quantifiable power, cooling, and space savings presented above. E-mail archiving eliminated mailbox quotas across the company, which increased employee productivity. Moving from an all tape backup strategy to a mix of disk and tape reduced the cost of tape media from $1M to $20,000 annually. Last, but not least, a consolidated storage environment increased the operational efficiency of the IT staff. As a matter of fact, the amount of storage that each full time engineer manages increased from 85 TB in 2004 to 230 TB by the end of 2008 – an increase of 170%.

Why This Matters

Mechanically spinning disk drives consume a lot of space, power, and cooling in the data center. Moving infrequently used data to bigger, slower drives over time can dramatically reduce the ongoing costs of storing information.

With online storage capacity growing at 60% and a two tier strategy relying predominantly on power hungry Tier 1 storage, EMC was running out of data center space, power, and cooling. A tiered storage strategy with archiving software and five tiers of storage has been embraced over the past five years. An estimated $44M of capital equipment costs for new Tier 1 storage systems has been avoided. A space savings of 56% and a reduction in the ongoing cost of power, cooling, and space of $5.3M have been confirmed.

Server Efficiency

Consolidating onto fewer, more efficient servers can be used to dramatically reduce power, cooling, and space requirements in the data center. In this section of the report, we examine the savings realized by EMC's ongoing efforts to improve server efficiency with VMware-enabled server consolidation.

As shown in Figure 5, a first phase of VMware server virtualization software has been used to consolidate 1,250 servers down to just 250. Older servers consuming two units of rack space were replaced with larger, more efficient servers consuming four units of rack space. An ESG Lab audit of the space, power, and cooling consumed before and after consolidation indicates that the new consolidated server environment consumes 60% less space, 70% less power, and 70% less cooling.

FIGURE 5. INCREASING EFFICIENCY WITH SERVER VIRTUALIZATION

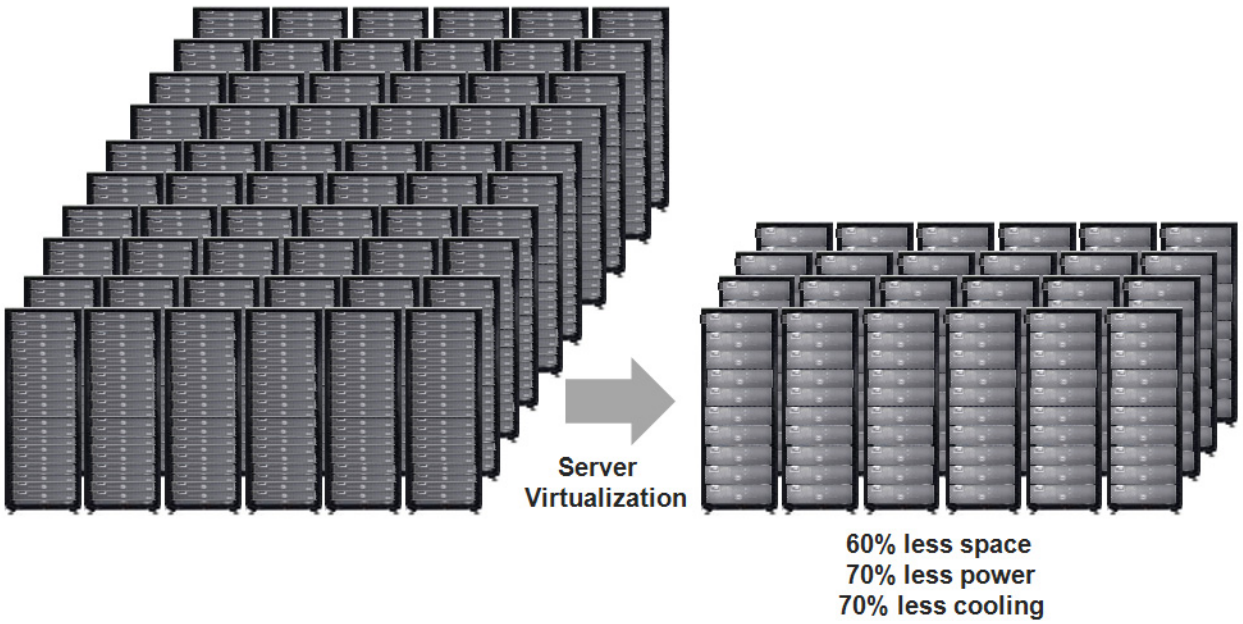

The initial phase of EMC's server consolidation began in 2006. EMC and VMware professional services were engaged to help design, test, and document best practices.

Two tiers of virtual server infrastructure were defined. The first tier was defined with high availability and disaster recovery for critical applications. This tier of approximately 20 mission critical applications, which bumps up to around 27 at the end of each fiscal quarter, represents a minority of the approximately 500 applications supported by the worldwide EMC IT organization. A second tier with high availability only was designed for the majority of applications that can tolerate the delay of a restore after a disaster.

Production implementation began with new applications and now includes the replacement of older servers as they are decommissioned. Approximately 33% of eligible applications have been migrated to date.

The results of this server consolidation are summarized in Table 3 and Figure 6 (see next page).

Sweep the Floor

The first wave of server consolidation, which focused on new application deployments and Tier 2 applications, has been used to justify a "sweep the floor" consolidation initiative that is currently underway at EMC. ESG Lab expects that the ongoing savings realized by this initiative will increase dramatically in coming years as the balance of servers are consolidated (approximately 33% today with a goal of 100%). A planned increase in the ratio of virtual to physical servers (approximately 5:1 today with a goal of 40:1) will magnify savings even more.

TABLE 3. THE BOTTOM LINE FOR THE FIRST 3 YEARS

	BEFORE (Q1'06)	AFTER (Q4'08)	Savings
Servers	1,252	250	
Racks	60	24	60%
Power	$2.5M	$0.7M	**$1.8M**
Cooling	$1.5M	$0.5M	**$1.0M**
Space	$0.11M	$0.05M	**$0.06M**

The Four Server Efficiency Phases in Detail

For a crisp summary of what EMC did here are its approaches and achievements/aims at each phase; it is of course not prescriptive for all other organizations, but nonetheless it gives clarity on the step-by-step value.

Phase 1: Virtualize new dedicated application environments

Instead of buying 15-20 physical servers to support a given application project, EMC was able to buy just 3-4 and to load multiple Virtual Machines on each to support all production and non-production environments. Overall this enabled EMC to avoid buying 640 physical servers through December 2008. However, because servers were still dedicated to specific applications, the consolidation ratio suffered (5:1 P2V ratio) and CPU utilization was not optimized.

Phase 2: Replace EOSL (end of service life) hardware with virtualized, shared servers

Instead of buying replacement servers for every existing server that reached "end-of -service-life", EMC built out its ESX servers and virtualized the application environments that had been running on the physical servers that were to be being decommissioned. Therefore, although 424 physical servers were decommissioned, only 62 virtualized servers were required to replace them, which produced an additional net reduction of 362 servers through December 2008. The combined results from the first two phases are shown in Table 3 and Figure 6. While progress was good and had yielded significant financial benefit, the combination of Phase 1 and Phase 2 had also resulted in many "islands of virtualization" with hardware still "owned" by the application(s) running on it. The consolidation ratio was better than Phase 1 alone, but - at 7:1 – still resulted in lower than optimal CPU utilization, which lead EMC to enter into its next two phases, which it describes as 'hyper-consolidation'.

Phase 3: SVC Project to build out highly consolidated, tiered service level hosting platform

EMC is architecting and deploying a consolidated, tiered and shared application hosting platform to achieve 40:1 consolidation ratios and optimal CPU utilization. It plans to us this platform to host all future application projects, breaking its prior paradigm of dedicated hardware for new applications. At an annual project demand for 150 physical servers at the old 5:1 P2V ratio, this would avoid the purchase of 750 servers over the next 5 years.

Phase 4: Sweep-The-Floor

Migrate the applications running on 1600 legacy servers over to 40 physical servers in the new SVC infrastructure (at the expected 40:1 P2V ratio) for an additional net server reduction of 1560. The savings (maintenance elimination and reduced energy consumption) are expected to be $8M over the next 5 years.

FIGURE 6. DATA CENTER INFRASTRUCTURE COST ANALYSIS (SERVERS)

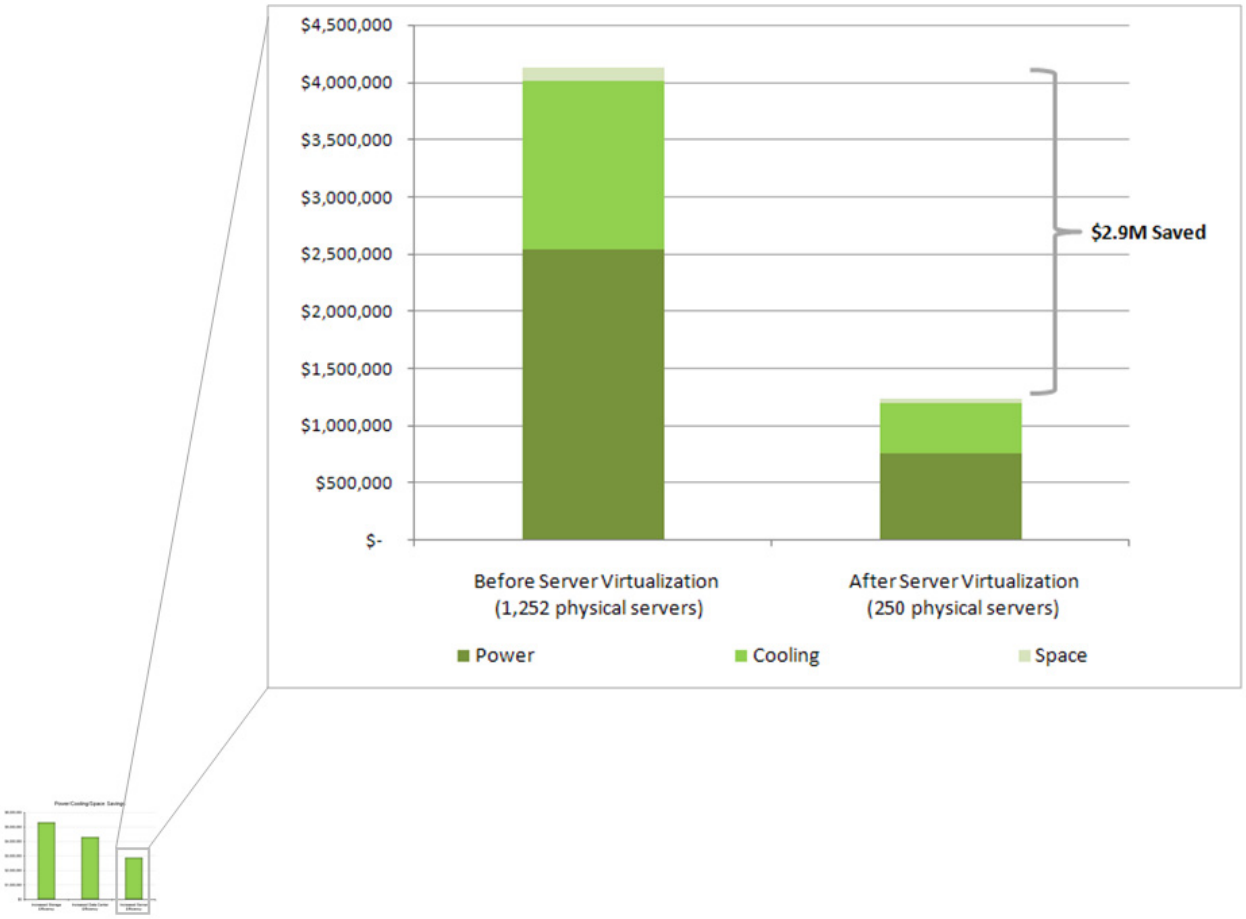

Why This Matters

Power hungry servers—and the never ending growth of applications running on those servers—were stressing the limits of EMC's IT infrastructure. A core data center nearing capacity needed a $26M upgrade to accommodate more power and cooling. ESG Lab has verified that EMC has realized significant savings during the first wave of VMware-enabled server consolidation. Consolidating from 1,252 to 250 servers helped extend the useful life of the data center by eight years as the cost of $4.2M worth of new servers was avoided. Space requirements were reduced by 60% as the ongoing costs of power and cooling were reduced by 70%.

ESG Lab's View

A growing number of IT organizations are struggling with the same challenges that EMC's IT organization was facing back in 2004. The growth of applications, servers, and storage arrays in the data center was straining the capacity of existing data centers. A new state of the art energy efficient data center was projected to cost more than $120M. Senior management challenged EMC IT to think differently—and more efficiently—and to follow the advice EMC gives its own customers. The challenge was not only to delay spending $26M on a data center upgrade, it was to take a step back and take a systems approach to create a flexible strategic IT plan that responded to and anticipated the needs of EMC's business. Working with EMC's professional services organization, a business planning process revealed that a data center expansion could be avoided and the useful life of a thirty year old data center extended. Investment would still be required, but it was appropriate for creating the flexible foundation that both IT and the business required.

ESG research indicates that EMC's IT organization is not alone when it comes to power and cooling issues.[7] With power consumption moving to the forefront as a strategic concern, 66% of IT decision makers whose organizations have active green projects say data center power and cooling initiatives are among the top IT-related programs important to the overall success of the organization's green strategy. Among organizations planning green programs, 59% intend to launch efforts to reduce data center power and cooling.

Among survey participants looking to reduce data center power and cooling requirements, 79% have deployed or are planning to deploy more power-efficient servers and 73% are investing in server virtualization as a way to reduce power consumption. However, few IT decision makers believe that simply improving their server power consumption profile is sufficient to meet data center power reduction targets. Most organizations are also implementing or planning to undertake a number of additional initiatives, including efforts to improve storage system power efficiency, create more efficient data center facilities infrastructure, and utilize data reduction technologies to decrease storage capacity requirements.

It should be noted that you don't need to buy all EMC (or all any other vendor's) products or services in order to get some of the benefits documented in this report. However, with new thinking and approaches, it's clear that new technologies (such as larger capacity drives and an e-mail archiving platform) played a key role in allowing EMC to achieve much of the savings. Such products, as well as professional services (without which it's unlikely the IT team would have achieved so much so quickly) are, of course, available from EMC.

A large part of any IT infrastructure budget can easily be absorbed by power and cooling. And much of it is essentially not used well (read: wasted). In other words, very significant cost and CO2 reductions—across most business units and applications—can be gained by focusing on power and cooling efficiency. As such, organizations would be well advised to follow EMC's lead in improving the air flow, cooling, and power efficiency of their existing data center infrastructure.

EMC is not done with its IT efficiency efforts; and, in reality, it never will be, as technological advances open up new efficiency opportunities all the time. Examples include the deployment of Tier 0 flash drives, the ongoing sweep the floor server virtualization effort, and continuous improvements in cooling and power distribution in the data center. As a matter of fact, EMC's IT management suggests that *the organization is only about 50% 'done.'* So, there's equally as much saving still to be realized as it has already enjoyed. This fact alone should motivate many other IT users to get started down the efficiency path as it speaks volumes to the potential.

One of the areas where EMC is not done yet is a return on investment analysis for each of the efficiency efforts presented in this report. While ESG Lab has confirmed that a three year return on investment was used to justify the tiered storage initiative back in 2004, the actual costs—and savings—incurred have not yet been calculated. When that effort is completed, ESG Lab is confident that the $12.5M in measurable savings in power and cooling costs highlighted in this report will represent a fraction of the savings that EMC has realized over the past five years. Besides the $74M in cost avoidance for a data center expansion, storage systems and servers, we're confident that savings due to the increased efficiency of the IT organization—and the increased productivity of employees—will prove that EMC has realized an excellent return on investment.

[7] Source: ESG Research Report: *Global Green IT Priorities*, November 2008

EMC has derived *huge* benefits for itself and provided a blueprint for others to follow suit. Even if some of the elements in EMC's overall savings are not directly replicable in every data center, and even if some reading this report wish to apply a skeptical level of 'discount' to some of the figures, what remains is nonetheless impressive. IT managers that are motivated to start—or simply to improve—their efforts in a drive to efficiency are encouraged to contact EMC and get a conversation started. As EMC's results demonstrate, this is a conversation that could be extremely valuable on many fronts.

20 Asylum Street
Milford, MA 01757
Tel: 508-482-0188
Fax: 508-482-0218

www.enterprisestrategygroup.com