

PERFORMANCE AND ENERGY ADVANTAGES OF DELL ENERGY SMART SERVERS AND LIEBERT COOLING SYSTEMS

David Moss
Data Center Thermals

Ramesh Radhakrishnan and Jenwei Hsieh
Enterprise Solutions Engineering

Greg Darnell
Server/Storage Performance Analysis Lab

Scott Hanson
Dell Enterprise Technology Center

June 2007



Contents

Executive Summary	3
Introduction	3
Generational Differences in Dell Servers	4
Test Configuration	4
Dell PowerEdge Energy Smart Servers	4
Power and Performance Comparisons of Dell PowerEdge Servers	5
Combining Energy Efficient Compute and Cooling Systems	5
Key Findings	6
Test Environment Overview	6
Scenario 1: Constrained on Power or Cooling	7
Chilled Water Facility	7
DX Facility	8
Scenario 2: Fill the Racks	8
Chilled Water Facility	8
DX Facility	8
Scenario 3: Direct Replacement with Same Number of Systems	9
Chilled Water Facility	9
DX Facility	9
Scenario 4: Replace with Less Systems to Achieve Similar Performance	9
Chilled Water Facility	10
DX Facility	10
Conclusion	10
Appendix A: Energy Efficiency for HPC workloads	11
Appendix B: Comparing Energy Smart and Blade Server Solutions	12
Appendix C: Best Practices and Cooling Overview	13

Executive Summary

The data center of today is a much different place than it was a decade ago. New high-density servers are accelerating the demand for increased power and cooling capacity in the data center while the cost of upgrading or building new facilities may not be in the budget. Traditionally, data centers were built to last up to 10 years but today's centers are running low on cooling and power capacity in half that time. In addition to capacity concerns, energy costs, government legislation and usage limits imposed by local utilities are increasing the pressure businesses face in managing the data center. As a result, IT organizations are looking for more energy-efficient products to lower their total cost of ownership.

Dell's new line of PowerEdge Energy Smart servers are the first industry-standard servers configured to optimize energy efficiency by not only improving performance but reducing power requirements as well. Studies described in the "Generational Differences in Dell Servers" section show that the Dell PowerEdge Energy Smart server increased performance by 80% over previous generation servers while requiring only 60% of the power, resulting in a 200% increase in performance per watt.

The Energy Smart servers offer significant improvements even in high performance computing environments where the PowerEdge Energy Smart 2950 was shown to provide an average of 17% to 23% better performance per watt over a similarly configured HP DL380 G5. (See Appendix A for more details.)

The Dell PowerEdge Energy Smart server can dramatically increase performance per watt as compared to the last generation server; and, when combined with Liebert™ cooling solutions can drive total energy savings in the data center even further.

In addition to engineering advances, Dell's commitment to energy optimization includes its partnership with Emerson Network Power™ for its best-in-class Liebert cooling systems. Dell worked with Liebert to devise four scenarios, outlined in the paper below, that address common customer problems by utilizing Dell Energy Smart servers and Liebert cooling. In these scenarios where customers were either constrained on power or needed to replace older systems with newer systems for either more power or space, Dell demon-

strates increased performance, delayed need for costly data center expansions or retrofitting, decreased power consumption, and/or the recapture of data center floor space. Key results showed:

- Up to a 250% performance gain with no increase in facility power requirements
- Up to an 80% performance increase with a reduction in facility power by 42%
- As much as a 65% reduction in facility power while maintaining the same level of performance

Introduction

Not long ago it was easy to add capacity to the data center. If there was space in a rack, a server could be placed into it. The amount of power that was available at the rack level was enough to power anything placed in the rack. Floor space to hold additional racks of servers was the key delimiting factor in capacity for the data center.

Server performance has increased significantly in recent years and the power requirements to run those servers have also increased. At the same time, servers are now available in form factors that make it possible to load more of them into a single rack. Many data centers were designed to only accommodate up to 5kW per rack. High density servers address the space issue, but bring new power and cooling challenges to the data center.

Many data centers were designed for a life expectancy of 10 years but are experiencing early obsolescence due to the new high density hardware. The issue is that many businesses do not have it in their budget to expand or to build new centers so they need ways to "right-size" their current facilities. This includes figuring out how to get the most computational power given a fixed set of kilowatts available in the data center.

Businesses that are able to meet their IT demands without increasing power or cooling demands can forego data center redesign or expansion for several more years and increase the return on capital investment for the original data center. As a result, there is an ever-increasing focus on maximizing data center resources.

In the past there was concern that energy efficiency also meant diminished performance. This is no longer

the case. Today's data center can benefit from technically advanced servers that consume much less power but deliver a noticeable improvement over the performance of past generations, such as Dell's PowerEdge Energy Smart servers. Coupled with advanced cooling technologies, companies are realizing decreased energy usage, increased performance and better space utilization to prolong the life of the data center.

Generational Differences in Dell Servers

Dell's ninth generation of PowerEdge servers incorporates the latest technology available in terms of processing, memory, power, and thermal design. The PowerEdge Energy Smart server line is based on the ninth generation platform and is designed to consume less power while still providing excellent performance.

This section describes the differences in performance and power between generations of Dell PowerEdge Servers. Some of these differences are due to industry changes and advances in technology and many are due to advances in thermal engineering by Dell in the PowerEdge Energy Smart server line.

Test Configuration

Four generations of Dell PowerEdge servers were tested with the SPECjbb2005® benchmark. The servers were configured as similar as possible, but due to changes in technology between different generations, the exact same components could not be used. Table 1 provides details on the configuration of the servers used in these tests.

	6th Generation	7th Generation	8th Generation	9th Generation
Server	PowerEdge 1650	PowerEdge 1750	PowerEdge 1850	PowerEdge Energy Smart 1950
CPU	2 x 1.4 Ghz Pentium® III	2 x 3.2 Ghz Xeon®	2 x 2.8 Ghz Dual-Core Xeon® 7030	2 x 2.33 Ghz Dual-Core Xeon 5148LV
RAM	4 GB	4 GB	4 GB	4 GB
Disks	2 x 18GB 10K RPM	2 x 18GB 10K RPM	2 x 36GB 15K RPM	2 x 73GB 10K RPM SAS

Table 1 – Configuration of Dell PowerEdge Servers for SPECjbb2005 tests

The SPECjbb2005 benchmark is the SPEC® (Standard Performance Evaluation Committee) benchmark for evaluating the performance of server-side Java. SPEC is a non-profit group of computer hardware and software vendors, system integrators, universities, research organizations, publishers and consultants. This benchmark evaluates the performance of server side Java by emulating a three-tier client/server system, and emphasizes the middle tier. It also measures the performance of processors, cache, and memory hierarchy. SPECjbb2005 models more closely how today's Java business applications found commonly across the business world are designed and implemented. For the latest SPECjbb2005 benchmark results, visit <http://www.spec.org>.

Dell PowerEdge Energy Smart Servers

The results in Figure 1 show a steady increase in power consumption from 6th generation (PE1650) to 8th generation (PE 1850) servers. In that time frame, memory and processor technology were constantly improving with increased performance but at the expense of increased power demands.

Dell and the industry realized the problem with this trend and responded with newer, more energy efficient technologies that help increase the performance curve without a steep increase in power usage. Dell's Energy Smart design, coupled with industry changes, resulted in a significant decrease in power draw in its ninth generation servers (Energy Smart 1950), as shown in Figure 1.

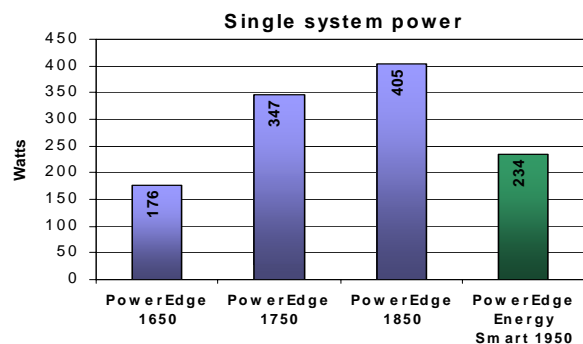


Figure 1 - Dell PowerEdge Energy Smart Power Savings

Power and Performance Comparisons of Dell PowerEdge Servers

Generational performance increases have been dramatic within the last decade. Comparing 6th generation servers to the PowerEdge Energy Smart servers, it can be seen that performance has increased about 10 fold. Figure 2 shows the performance increases measured using the SPECjbb2005 benchmark.




Figure 2 - SPECjbb2005 Score on Dell PowerEdge Servers

If we combine the power savings and performance measurements to derive performance per watt, Dell Energy Smart engineering and industry technology advances are even more evident. Figure 3 shows an 80% performance increase in the Dell PowerEdge Energy Smart server while only using 60% of the power as compared to the 8th generation server - a performance per watt increase of over 200%.

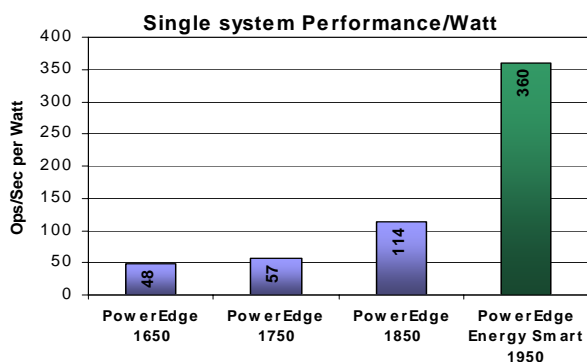


Figure 3 - Performance per watt of SPEC jbb2005 on Dell PowerEdge Servers

The performance per watt improvements are not just limited to applications commensurate with business type workloads; similar improvements are seen in several other compute intensive workloads that put significant stress on the CPU and memory subsystems.

Figure 4 shows several compute intensive workloads and the relative performance per watt increase comparing 8th generation and 9th generation Energy Smart servers.

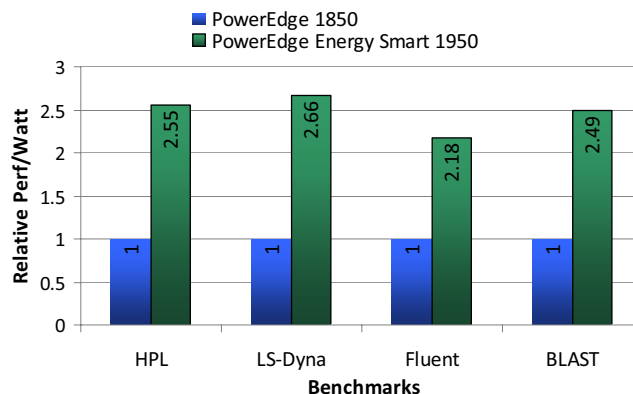


Figure 4 - Relative Performance per Watt for HPC Benchmarks

In summary, recent technology advances and Dell engineering have resulted in a line of 9th generation Energy Smart servers that offer significant improvement in performance and power usage, even in intensive computing environments. To further illustrate the performance of Dell Energy Smart servers in high-performance computing (HPC) environments, a recent study compared the Dell PowerEdge Energy Smart 2950 to a similarly configured HP DL380 G5 system. The Dell PowerEdge ES 2950 system was measured to provide on the average of 17% to 23% better performance per watt over the three HP DL380 G5 configurations across four HPC applications. See Appendix A for a complete overview of the study.

Combining Energy Efficient Compute and Cooling Systems

With the demand for computing power skyrocketing and energy costs rising, the challenge of data center cooling becomes more complex - and more critical. Emerson Network Power's Liebert cooling solutions can help data centers cut costs and increase data center efficiency. The Liebert X-treme Density (XD) family represents a new class of cooling system designed specifically for high density systems.

In this section, the advantages of combining several different Dell PowerEdge Energy Smart servers with cooling systems that enhance overall energy savings are

outlined. In partnership with Liebert, Dell has analyzed these hardware combinations to show significant benefits to the data center.

Because industry best practices for cooling play a crucial role in the energy efficiency of the data center, these guidelines along with an overview of the different types of cooling systems are included in Appendix C. Newer technologies such as VFD (Variable Frequency Drives) and Digital Scroll™ are also discussed.

Four scenarios are analyzed below that show significant improvements using Dell Energy Smart servers and efficient Liebert cooling systems in facilities with either Chilled Water (CW) or DX environments. These scenarios are representative of challenges a customer may face and include: 1) Constraining power or cooling, 2) Filling the racks, 3) Direct replacement with same number of systems, and 4) Replace with less systems to achieve similar performance. The savings are shown relative to the performance and power signature associated with Dell's 8th generation equipment. Similar comparisons could also be made with earlier servers.

Key Findings

Table 2 shows a summary of findings for the Chilled Water and DX facilities. Details for each scenario for both the Chilled Water and DX facilities are covered in this section.

Scenario	Server QTY Increase/Decrease		Facility Power Inc / Dec		Relative Performance	
	CW	DX	CW	DX	CW	DX
Facility Power Constrained	+91%	+75%	-0-	-0-	3.5X	3.2X
Fill all Racks	+250%	+250%	+80%	+94%	6.4X	6.4X
Replace Each Server 1:1	-0-	-0-	-42%	-46%	1.8X	1.8X
Reduce Energy; Same Performance	-44%	-44%	-65%	-67%	-0-	-0-

Table 2 - Key Findings for Multiple Scenarios

Utilizing Dell PowerEdge Energy Smart servers and Liebert cooling produced impressive results in terms of efficiency and performance improvements:

- Up to a 250% performance gain with no increase in facility power requirements
- Up to an 80% performance increase with a reduction in facility power by 42%
- As much as a 65% reduction in facility power while maintaining the same level of performance

Test Environment Overview

As a baseline, consider 50 racks of last generation servers, specifically the Dell PE1850. Based on Dell testing, the power consumed by a single server of this generation is approximately 400 watts when running the SPECjbb2005 benchmark. Configured with two dual-core 2.8 GHz Intel® Xeon® Processor 7030 processors; 4 x 1 GB DDR2 memory; 2 x 36 GB, 15K RPM drives; and redundant power; the PE1850 has a measured SPECjbb2005 rating of 46,360. In contrast, a similarly configured Dell PowerEdge Energy Smart 1950 consumes 234 watts and has a SPECjbb2005 rating of 84,312. The Energy Smart PE1950 is configured with the same number of hard disks and memory, and uses Intel's Low Voltage Xeon® processors.

For purposes of comparison, this paper assumes a raised floor cooling limit of approximately 5 kW per rack. For the PE1850, that 5kW limit amounts to about 12 servers per rack. Assuming 50 racks, the baseline is then 600 servers. Each scenario will include a comparison to this baseline 600 servers and their complement of raised floor cooling without energy saving features. CRAC (Computer Room Air Conditioner) and CRAH (Computer Room Air Handler) performance for all analyses assumes typical sensible limits associated with nominal return temperature and humidity conditions. For the baseline chilled water configuration, three Liebert CS125W chilled water systems were chosen which represents a 20-30% overage in capacity. No VFD (Variable Frequency Drive) was assumed. For the baseline DX case, five CS105W units were assumed representing an overage of about 35-45% based on the nominal capacity limits assumed. The facility is assumed to be fairly well optimized with adherence to at least 4 of the 6 best practices stated in Appendix C.

Figures 5 and 6 show the power breakout for both chilled water and DX environments used as the baseline in all 4 scenarios.

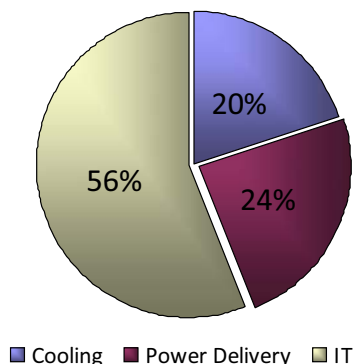


Figure 5 - Power Breakout for Chilled Water Facility




Figure 6 - Power Breakout for DX Facility

The scenarios discussed are a study of IT and cooling equipment and do not touch on power delivery. For the purposes of comparing power at the facility level, an assumption was made to use a composite efficiency of 70% to account for the electrical delivery to both the IT equipment and to the cooling system components.

It is important to note that the performance and power measurements for the Dell systems are measured values. The assumptions made for cooling in each of the scenarios and the associated energy consumption are estimates provided by Liebert.

Scenario 1: Constrained on Power or Cooling

Many businesses are faced with constraints on their ability to increase power or cooling in the data center. In this scenario, the addition of equipment exceeding the current facility power or cooling requirements might require the building of a new facility or the upgrade of

an expensive chiller to the facility. In this case, the customer could consider the replacement of the baseline 600 last generation servers with a number of Power-Edge Energy Smart 1950 servers that fit within the current facility power and cooling envelope.

Chilled Water Facility

The original cooling assumption included three CS125W units operating with excess capacity but with no VFD running at about 3.5 watts to cool every 10 watts of IT equipment.

The replacement with new servers in the chilled water facility results in 548 new servers within the same facility power and cooling envelope. The relative performance associated with this replacement is 3.5 times the previous performance (as shown in the example calculation below) at no greater operating costs.

Example Calculation:

8G Server- 600 servers

- 600 x 405 watts = 243 kW
- 3 x Liebert CS125W = 86.6 kW
- Server plus Cooling: 243 + 86.6 = 329.6 kW
- Cooling as percent of server load: 86.6 / 243 = 35.6%
- Composite SPECjbb score: 600 x 46,360 = 27,816,000

9G Energy Smart - 1148 servers

- 1048 x 234 watts = 269 kW
- 2 x Liebert CS125W plus 1 XD system = 56 kW
- Server plus Cooling: 269 + 56 = 325 kW
- Cooling as percent of server load: 56 / 269 = 20.8%
- Composite SPECjbb score: 1148 x 84,312 = 96,790,176

Relative comparisons

- Similar total power: 329.6 kW (8G); 325 kW (9G)
- Performance: 96,790,176 / 27,816,000 = ~ 3.5 times more performance
- Cooling, as a percent of server load went from 35.6% to 20.8%, a reduction of just over 40%

With a cooling refresh that includes a single XD system and two rather than three CS125W CRAH units equipped with VFD, the power to cool relative to IT load drops to about 2.1 watts to cool 10 watts of IT. That essentially results in a 40% decrease in the relative cooling energy.

The efficiency improvements described by these changes in cooling solutions would be even more dramatic if the change from the initial baseline also

included correction of some of the best practices described in Appendix C.

Before	After	Difference
600 previous generation servers	1148 Energy Smart PE1950	
243 kW IT load	269 kW IT load	
50 racks at 12 servers per rack	28 racks at 41 servers per rack	Less space consumed
27,816,000 composite SPEC score	96,790,176 composite SPEC score	3.5X performance
3 (qty) CS125W CRAHs (No VFD)	2 (qty) CS 125W CRAHs w/VFD	Less space consumed
No XD System	1 XD System	
471 kW total facility power for these servers	464 kW total facility power for these servers	
3.5 watts driving cooling equipment for 10 watts IT	2.1 watts driving cooling equipment for 10 watts IT	~40% reduction

Table 3 - Performance Increase Within the Same Power and Cooling Envelope

DX Facility

In the DX facility, the CRACs were replaced with three CS70W; and Digital Scroll was assumed for all. A part of the cooling burden was replaced with a single XD system. As a percentage of IT power, the power required by the cooling systems dropped by 8%. In this refresh, the implementation of the XD solution allows the density of the racks to be increased past the previous limit of 5 kW. This refresh assumed 25 racks with all 42U filled with servers.

Before	After	Differences
600 previous generation servers	1050 Energy Smart PE1950	
243 kW IT load	246 kW IT load	
50 racks at 12 servers per rack	25 racks at 42 servers per rack	Less space consumed
27,816,000 composite SPEC score	88,527,600 composite SPEC score	3.2X performance
5 (qty) CS105W CRACs (No digital cross)	3-4 (qty) CS105W CRACs with digital cross	Less space consumed
No XD System	1 XD System	
611 kW total facility power for these servers	606 kW total facility power for these servers	
6 watts driving cooling equipment for 10 watts IT	5.5 watts driving cooling equipment for 10 watts IT	8% reduction

Table 4 - Performance Increase within same Power and Cooling Envelope

Scenario 2: Fill the Racks

In scenario 2, the customer goal is to maximize compute power and performance by completely filling the available rack space. With the 5 kW raised floor cooling limit assumed for the 8G servers, only 12 servers per rack were possible. With the addition of XD systems, the racks can be fully populated. This scenario assumes there is no power limit on the facility and the racks are completely filled with PowerEdge Energy Smart 1950 servers.

Chilled Water Facility

The replacement with new servers in the chilled water facility results in 2100 new servers in the same physical space. Facility power nearly doubled and there was a 6.4X improvement in performance. Full racks are enabled by the addition of 2 Liebert XD systems. One CRAC is removed. The reduction in power to cool relative to IT load is about 25%.

Before	After	Differences
600 previous generation servers	2100 Energy Smart PE1950	
243 kW IT load	491 kW IT load	
50 racks at 12 servers per rack	50 racks at 42 servers per rack	
27,816,000 composite SPEC score	177,055,200 composite SPEC score	6.4X performance
3 (qty) CS125W CRAHs (No VFD)	2 (qty) CS 125W CRAHs w/VFD	Less CRAC space
No XD System	2 XD SystemS	
471 kW total facility power for these servers	852 kW total facility power for these servers	
3.5 watts driving cooling equipment for 10 watts IT	2.6 watts driving cooling equipment for 10 watts IT	~25% reduction

Table 5 - Performance Increases with Filled Racks

DX Facility

In the DX facility, the CRACs were replaced with six CS70W; and Digital Scroll was assumed for all. The replacement with new servers in the DX facility results in 2100 new servers in the same physical space. Additional space is needed for one extra CRAC. Facility power nearly doubled but there was a 6.4X improvement in performance. Full racks were enabled by the addition of two Liebert XD systems. The reduction in power to cool relative to IT load is less than 5%.

Before	After	Differences
600 previous generation servers	2100 Energy Smart PE1950	
243 kW IT load	491 kW IT load	
50 racks at 12 servers per rack	50 racks at 42 servers per rack	
27,816,000 composite SPEC score	177,055,200 composite SPEC score	6.4X performance
5 (qty) CS105W CRACs (No Digital Scroll)	6 (qty) CS70W CRACs with Digital Scroll	More CRAC space required
No XD System	2 XD SystemS	
611 kW total facility power for these servers	1185 kW total facility power for these servers	
6 watts driving cooling equipment for 10 watts IT	5.8 watts driving cooling equipment for 10 watts IT	3% reduction

Table 6 - Performance Increases with Filled Racks with DX

Scenario 3: Direct Replacement with Same Number of Systems

In this scenario, the server count is held steady but all are upgraded to ninth generation, Energy Smart servers.

Chilled Water Facility

The facility power dropped by 42% with an 80% improvement in performance when replaced with the same number of Energy Smart servers.

Before	After	Differences
600 previous generation servers	600 Energy Smart PE1950	
243 kW IT load	140 kW IT load	
50 racks at 12 servers per rack	50 racks at 12 servers per rack	
27,816,000 composite SPEC score	50,587,200 composite SPEC score	1.8X performance
3 (qty) CS125W CRAHs (No VFD)	2 (qty) CS 125W CRAHs w/VFD	Less space consumed
No XD System	0 XD System	
471 kW total facility power for these servers	273 kW total facility power for these servers	42% reduction
3.5 watts driving cooling equipment for 10 watts IT	3.5 watts driving cooling equipment for 10 watts IT	No appreciable reduction in cooling efficiency with this choice

Table 7 - Replacement of Same Quantity Servers

The power requirements for this quantity of servers are too small to warrant the deployment of Liebert XD, so CRAC cooling was continued. The heat reduction allows a CRAC to be eliminated. Based on the relation of CRACs to server load there doesn't appear to be a real reduction in cooling as a percentage of load. The absolute cooling power has obviously dropped. The number of racks could have been reduced and still fit within the 5 kW assumption for raised floor capacity.

DX Facility

In the DX facility, the CRACs were replaced with four smaller CS70W units; and Digital Scroll was assumed for all. Facility power dropped by 46% with an 80% improvement in performance. The power requirements for this quantity of servers are too small to warrant the deployment of Liebert XD, so CRAC cooling is continued. The energy used for cooling in relation to the IT load has dropped about 15%. The number of racks could have been reduced and still fit within the 5 kW assumption for raised floor capacity.

Before	After	Differences
600 previous generation servers	600 Energy Smart PE1950	
243 kW IT load	140 kW IT load	
50 racks at 12 servers per rack	51 racks at 12 servers per rack	
27,816,000 composite SPEC score	50,587,200 composite SPEC score	1.8X performance
5 (qty) CS105W CRAHs (No Digital Scroll)	4 (qty) CS70W CRAHs w/VFD	Less space consumed
No XD System	0 XD System	
611 kW total facility power for these servers	332 kW total facility power for these servers	46% reduction
6 watts driving cooling equipment for 10 watts IT	5.2 watts driving cooling equipment for 10 watts IT	~15% reduction

Table 8 - Replacement of Same Quantity Servers with DX

Scenario 4: Replace with Less Systems to Achieve Similar Performance

In this scenario, the goal is to achieve the same level of performance while replacing older generation systems with fewer Energy Smart servers.

Chilled Water Facility

When 600 PowerEdge 1850's were replaced by 336 PowerEdge Energy Smart 1950's to provide the same performance, the facility power dropped by 65%. The power requirements for this quantity of servers are too small to warrant the deployment of Liebert XD, so CRAC cooling was continued. The heat reduction allows a CRAC to be eliminated and the remaining CRACs down-sized. Based on the relation of CRACs to server load, this analysis actually shows an increase in cooling power as a percentage of load. The absolute cooling power has obviously dropped. Both the rack and CRAC footprint were reduced.

Before	After	Differences
600 previous generation servers	336 Energy Smart PE1950	Large reduction
243 kW IT load	79 kW IT load	67% reduction
50 racks at 12 servers per rack	16 racks at 21 servers per rack	Less rack space
27,816,000 composite SPEC score	28,328,832 composite SPEC score	
3 (qty) CS125W CRAHs (No VFD)	2 (qty) CS90W CRAHs w/VFD	Less CRAC space
No XD System	0 XD System	
471 kW total facility power for these servers	163 kW total facility power for these servers	65% reduction
3.5 watts driving cooling equipment for 10 watts IT	4.5 watts driving cooling equipment for 10 watts IT	Increase in power to cool relative to IT power

Table 9 - Similar Processing Power with Reduction of Servers

DX Facility

When replacing older generation systems with fewer Energy Smart servers, the facility power dropped by 67%. The power requirements for this quantity of servers are too small to warrant the deployment of Liebert XD, so CRAC cooling was continued. The heat reduction allows two CRACs to be eliminated and the remaining CRACs were down-sized. Based on the relation of CRACs to server load, this analysis showed a similar relationship in cooling power as a percentage of load. The absolute cooling power has obviously dropped. Both the rack and CRAC footprint were reduced.

Before	After	Differences
600 previous generation servers	336 Energy Smart PE1950	Large reduction
243 kW IT load	79 kW IT load	67% reduction

Before	After	Differences
50 racks at 12 servers per rack	16 racks at 21 servers per rack	Less rack space
27,816,000 composite SPEC score	28,328,832 composite SPEC score	
5 (qty) CS105W CRACs (No Digital Scroll)	3 (qty) CS970W CRACs with Digital Scroll	Less CRAC space required
No XD System	0 XD System	
611 kW total facility power for these servers	199 kW total facility power for these servers	67% reduction
6 watts driving cooling equipment for 10 watts IT	6 watts driving cooling equipment for 10 watts IT	Similar cooling power relative to IT power

Table 10 - Similar Processing Power with Reduction of Servers in DX

Conclusion

Ever-increasing demands for computing power have customers struggling to find ways to optimize the performance and energy efficiency of the data center. Dell's PowerEdge Energy Smart servers provide customers with an immediate solution to:

- Maximize performance per square foot, potentially delaying costly expansions or new construction
- Reduce energy consumption resulting in significant cost savings
- Increase performance in the existing power envelope

By increasing performance, decreasing power consumption, and enabling the recapture of data center floor space, Dell's Energy Smart servers drive down data center costs and total cost of ownership. And, when combined with Liebert's state-of-the-art cooling solutions, Energy Smart servers can drive total energy savings in the data center even further.

Other server hardware vendors are taking an alternative approach by endorsing blades and proprietary rack cooling solutions for tackling data center power and cooling challenges. Appendix B compares Dell's Energy Smart servers to existing blade solutions.

Dell is leading the charge in helping customers solve power and cooling challenges with its industry-leading servers designed specifically to optimize energy efficiency and its partnerships with world-class companies such as Liebert. To find out more, go to www.dell.com.

Appendix A: Energy Efficiency for HPC workloads

Energy efficiency is becoming a top priority for high-performance computing users as well. To illustrate the advantages of the Energy Smart server for high-performance computing users, the performance and power consumption of the Dell PowerEdge Energy Smart 2950 were measured when running high-performance computing workloads and applications. The performance, power and energy-efficiency of the Dell server were also compared to the HP ProLiant DL380 G5. The Dell Energy Smart (ES) PE2950 demonstrated superior energy-efficient performance over the ProLiant DL380 G5 across all the workloads tested.

The table below summarizes the system configurations that were used for this comparison study.

CPU	Dell Energy Smart PE2950	HP ProLiant DL380 G5
CPU	Xeon® 5148LV	Xeon® 5148LV, Xeon® 5150, Xeon® 5160
Chipset	Intel® 5000X	Intel® 5000P
Memory	4 x 1GB PC-25300 FB-DIMM	4 x 1GB PC-25300 FB-DIMM
Disk(s)	3 x 36GB 2.5 10K SAS (RAID0)	3 x 36GB 2.5 10K SAS (RAID0)
OS	RHEL4 U4	RHEL4 U4

Table 11 - Server Configurations Tested

To measure the energy-efficiency or performance per watt, power was measured when running the benchmark (during steady state). The performance per watt represents the energy efficiency of the system and is computed by dividing the performance metric measured during the run by the measured power. Applications used in this study are:

High-Performance Linpack (HPL): The High-Performance Linpack (HPL) implementation is commonly used to rank supercomputers on the TOP500 Supercomputer Sites list. For more information about HPL, visit www.netlib.org/benchmark/hpl.

LS-DYNA: This general-purpose, transient dynamic finite-element program simulates complex real-world problems. The specific LS-DYNA workloads used in this study were neon refined and 3 cars. For more information about LS-DYNA, visit www.lstc.com.

FLUENT: The FLUENT applications allow users to perform CFD analysis around their particular models. Several benchmark data sets (workloads) are available from Fluent Inc. and the medium dataset was used in these performance tests. For more information on Fluent, visit www.fluent.com/software/fluent.

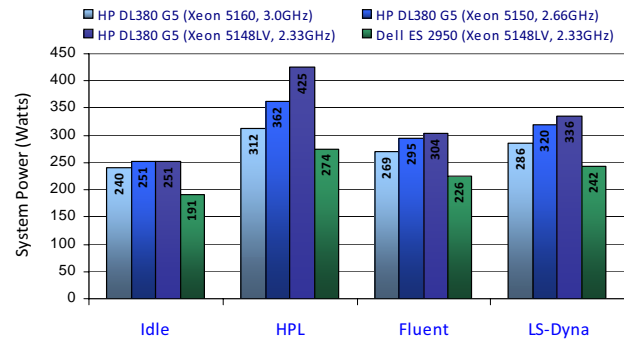


Figure 7: Measured System Power Consumption

Figure 7 shows the measured power consumption for the servers at idle and heavy utilization conditions when running the HPC applications. Power measured at idle represents the power drawn by the server when the OS is up and running but there are no applications running that stress the CPU. The CPU utilization at idle is close to 0%. The Dell PowerEdge ES 2950 system is observed to draw 25%-31% lower power compared to the HP system when the system is idle.

High-Performance Linpack (HPL) is a CPU and memory intensive benchmark and is observed to cause the highest power draw amongst all the applications. The Dell PowerEdge ES 2950 is observed to draw significantly less power compared to all three HP DL380 G5 configurations. The average power draw (including both idle and heavy utilization scenarios) for the Dell PowerEdge ES 2950 system was measured to be 18%-40% lower compared to the HP system configurations.

Figure 8 shows the relative performance/watt across the three HP configurations and the Dell PowerEdge ES 2950 system. The HP DL380 G5 system with the Xeon® 5160 (3.0GHz) CPUs is being used as the baseline for the performance/watt comparison. Due to the significantly lower power consumption of the Dell PowerEdge ES 2950 system, it exhibits better performance/watt and therefore superior energy-efficiency over the three HP DL380 G5 configurations. The Dell PowerEdge ES 2950 system was measured to provide on the average

(across the 4 HPC applications) 17% to 23% better performance per watt results over the three HP DL380 G5 configurations.




Figure 8: Normalized performance/watt when running HPC applications

Appendix B: Comparing Energy Smart and Blade Server Solutions

In June 2007, the Dell Solutions engineering team measured the performance and power consumption of the Dell PowerEdge Energy Smart 1950 server when running the SPECjbb2005 benchmark. The performance, power and energy-efficiency of that Dell server was compared to that of the separately tested Dell PE1955 and HP BL460c Blade system servers. The energy-efficiency of the blade systems was measured by a third party testing organization commissioned by Dell, Principled Technologies. A detailed report of their testing methodology and results are available in their Test Report titled, "SPECjbb2005 performance and power consumption on Dell and HP blade servers," Principled Technologies, June 2007. Table 12 summarizes the system configurations that were used in this comparison.

Table 12 - Energy Smart and Blade Server Configuration

The Energy Smart PE1950 server was configured similarly to the individual HP and Dell blade servers and tested at the Dell solutions engineering lab. To measure the energy-efficiency or performance per watt, power was measured when running the benchmark and readings from peak stable warehouses were used to calculate energy efficiency. The performance/watt is computed by dividing the performance metric measured during the run by the measured power and represents the energy efficiency of the system.

Figure 9 shows the average measured power consumption for a single Dell PE1955 blade, HP BL460c blade and a single ES PE1950 server when running SPECjbb2005 benchmarks. Both blade systems were fully populated when the measurements for the average power consumption were conducted. Based on Principled Technologies' test results, the single Dell and HP blades when configured with two Intel Xeon E5345 (2.33GHz, 1333FSB) CPUs consume 7.7% and 10.9% more power, respectively, compared to the Energy Smart PE1950 server configured with two Xeon L5320 (1.86GHz, 1066MHz) CPUs.




Figure 9 - Average System Power Draw

Figure 10 shows the average energy efficiency for the Dell and HP blade servers and compares it with the Dell PowerEdge ES 1950 system. The Dell and HP blade servers when configured with Xeon E5345 (2.33GHz, 1333FSB) CPUs are 19% and 14.4% more energy efficient compared to the Energy Smart PE1950 server configured with the Xeon L5320 (1.86GHz, 1066MHz) CPUs due to the faster FSB and CPU clock speeds.

	Dell Energy Smart PE1950	Dell PE1955 Blade Server	HP BL460c Blade Server
CPU	Quad-Core Intel® Xeon® L5320 (1.86GHz, 1066MHz FSB, 2x4MB L2)	Quad-Core Intel® Xeon® E5345 (2.33GHz, 1333MHz FSB, 2x4MB L2)	Quad-Core Intel® Xeon® E5345 (2.33GHz, 1333MHz FSB, 2x4MB L2)
Chipset	Intel 5000X	Intel 5000P	Intel 5000P
Memory	4 x 1 GB PC2-5300 FB-DIMM	4 x 1 GB PC2-5300 FB-DIMM	4 x 1 GB PC2-5300 FB-DIMM
Disk(s)	2 x 36GB 2.5 10K SAS	2 x 73GB 2.5 10K SAS	2 x 73GB 2.5 10K SAS
OS	MS Windows 2003 Server	MS Windows 2003 Server	MS Windows 2003 Server




Figure 10 - Energy Efficiency Comparison

The comparison between the Energy Smart PE1950 and Blade Systems showed that in terms of energy-efficiency, blade systems offer up to a 20% improvement over the Dell Energy Smart server. However customers that are not constrained by data center floor space should consider Energy Smart servers due to the high costs associated with blade systems and proprietary cooling solutions needed to support their high density. The price of the cooling required to achieve maximum density was not considered in the testing results. Cooling methods endorsed and supported by different Blade providers vary in price and will further increase the total cost of the computing solution making the Energy Smart servers an attractive solution for the customers computing needs.

Appendix C: Best Practices and Cooling Overview

The topics below summarize the best practices in Data Center design for improved reliability of the IT equipment and enhanced energy efficiency. Of primary importance is the arrangement of the racks into a "Hot Aisle / Cold Aisle" arrangement. This arrangement creates the best separation of the hot and cold air flows and minimizes the amount of hot air which "leaks" back into the servers. Once this is in place, the other five best practices supplement the separation of hot and cold air flow and are not mentioned in any particular order of effectiveness. With attention to these 6 practices together, it is common to see a significant cooling systems energy consumption savings.

Hot Aisle / Cold Aisle Arrangement - Are the racks set up with racks discharging air to a common aisle instead of into the input of the next row of racks?

Placement of the CRAC Units - CRAC units should be placed at the end of the hot aisle so that returning hot air doesn't mix with the cool air going to the rack inlets

Perforated Tile Placement - Perforated tile should only be in the cold aisles so to maximize the amount of cold air flowing into the racks

Optimal Rack Airflow - Blanking panels should be used on the front of the rack to push air through the servers and the obstructions removed from the rear of the rack to ensure movement of hot air out of the rack

Minimal Obstructions under the Raised Floor - the raised floor should be clear of clutter, waste and unnecessary equipment to reduce the horsepower required to move air to the racks

Sealed Cable Cut-Outs - All cable outlets feeding the racks should be sealed to reduce the loss of cold air to the hot aisle

When one considers the typical data center, there are various reasons the cooling equipment may be under utilized and operating inefficiently. Not following the aforementioned best practices is one reason. Under utilization can simply be the result of an over-deployment of cooling. It might be over-deployed inadvertently due to uncertainty of the load, or it might be over-deployed on purpose to handle a greater peak load, redundancy goals, or quick future server adds. Any or all of these reasons can contribute to cooling equipment operating at low partial load, but there are savings to be had by implementing various alternative cooling technologies.

Data Centers have two main types of modular raised floor cooling units that differ in the way they are equipped to chill the air. The more common modular unit, a Computer Room Air Conditioner (CRAC), uses compressorized refrigerant based cooling technologies. In contrast, other facilities choose to use the more efficient combination of a centralized chiller plant and Computer Room Air Handlers (CRAH). These modular CRAH units look just like CRACs. In fact, many people

often refer to the modular unit as a CRAC whether it truly is a CRAC or a chilled water CRAH.

The compressorized CRAC essentially does the real cooling within the unit itself, meaning energy is used in a vapor/compression cycle with a refrigerant. It still requires a condenser on the roof to eventually expel the heat. This technology is referred to as Direct Expansion or DX. Similarly the CRAC or facility is referred to as a DX unit or DX facility in contrast to a chilled water unit or facility.

The CRAH merely transfers heat from the data center air drawn through it to the chilled water that also circulates through it. That chilled water is pumped back to a centralized chiller where the water is chilled to low temperatures (commonly 45-50°F).

The reasons for choosing a CRAC unit are varied. Often, even if the facility has a chilled water system, it may not be dedicated to the data center; there could be a risk of planned or unplanned outages that may or may not be communicated with the IT staff. In tenant situations, it is often difficult for data center operators to secure Chilled Water systems, but a condenser on the roof is a relatively easy addition. Incremental adds of a CRAC unit may be immediately more appealing since the incremental add of a CRAH that pushes past the capacity of the chiller would trigger the addition of an expensive chiller system.

For Data Centers using compressorized refrigerant CRACs, the move to a Digital Scroll compressor can reduce the total Data Center cooling energy by minimizing the amount of compressor cycling and minimizing excessive humidity removal.

For Data Centers with Chilled Water CRAHs, the installation of air handlers with Variable Frequency Drives (VFD) enables a significant reduction in air handler energy consumption by providing the right amount of power

to blower fans based on the cooling load of the room. For simplicity, CRAC may be used interchangeably with CRAH throughout the remainder of the paper except where specifically note.

Besides air handlers that are most commonly used to pressurize a raised floor, there are other cooling technologies that operate above the raised floor. One such system, exclusive to Liebert, is the Liebert XD System. It requires a minimum of 6 kW per rack of load to operate efficiently, but once this level is reached, Liebert analysis shows a savings of as much as 30% in energy use as compared to a chilled water CRAH system. The XD system uses a non-compressorized refrigerant approach. A two phase heat transfer at the rack allows a very efficient means of absorbing the heat and moving it away from the rack to be transferred to the facilities cooling system.

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

Dell and *PowerEdge* are trademarks of Dell Inc.; *Intel* and *Xeon* are registered trademark of Intel Corp.; *Digital Scroll* and *Emerson Network Power* are trademarks of Emerson Electric Company; *Liebert* is a registered trademark of Liebert Corporation; *HP* and *Hewlett-Packard* are registered trademarks of Hewlett-Packard Company; *Red Hat* is a registered trademark of Red Hat Inc. *Linux* is a registered trademark of Linus Torvalds; *SPEC* and the benchmark name *SPECjbb* are registered trademarks of the Standard Performance Evaluation Corporation.

Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

©Copyright 2007 Dell Inc. All rights reserved. Reproduction in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell. Information in this document is subject to change without notice.