Dell PowerEdge 12th Generation Server System Profiles

A Dell technical white paper that shows the performance and energy efficiency differences between the user-selectable BIOS System Profiles

Bruce Wagner

Solutions Performance Analysis

Contents

Tables

Figures

January 2014 | Rev 1.10

Executive summary

Customers choosing one of the performance-per-watt optimized system profiles instead of the performance-optimized one can experience up to 60% reduced system idle power consumption and an overall power efficiency increase of 19% when the system if fully utilized. Given industry research that an Enterprise server's annual fully burdened cost per watt is $2.69, a savings of $134 per server, per year can be realized.

Introduction

The new Dell™ PowerEdge™ 12th generation servers support Dell Active Power Controller (DAPC) system profiles that serve as "easy buttons" for applying default BIOS and firmware values targeted for performance, performance per watt, or RAS for dense configurations. Table 1 shows the BIOS settings for each System Profile.

Table 1. "Easy Button" BIOS settings correspondence

		System Profile				
Subsystem	**Setting**	**Performance**	**Performance per watt (OS)**	**Performance per watt (DAPC)** *(default)*	**Dense configuration**	**Custom**
CPU	CPU Power Management	Maximum Performance	OS DBPM	System DBPM (DAPC)	System DBPM (DAPC)	System DBPM (DAPC)
	Turbo Boost	Enabled	Enabled	Enabled	Disabled	Enabled
	C1E	Disabled	Enabled	Enabled	Enabled	Enabled
	C-States	Disabled	Enabled	Enabled	Enabled	Enabled
Memory	Frequency	Maximum Performance	Maximum Performance	Maximum Performance	Dense Configuration Optimized	Maximum Performance
	Voltage	Auto	Auto	Auto	Maximum	Auto
	Patrol Scrub	Standard	Standard	Standard	Extended	Standard
	BIST Duration	Standard	Standard	Standard	Extended	Standard
	Refresh Rate	1x	1x	1x	2x	1x
Thermal	Thermal Algorithm	Maximum Performance	Minimum Power	Minimum Power	Dense Configuration	Maximum Performance

The system profiles can be selected through the following BIOS menus in Figure 1, Figure 2, and Figure 3.

Figure 1. POST splash screen with F2 = System Setup

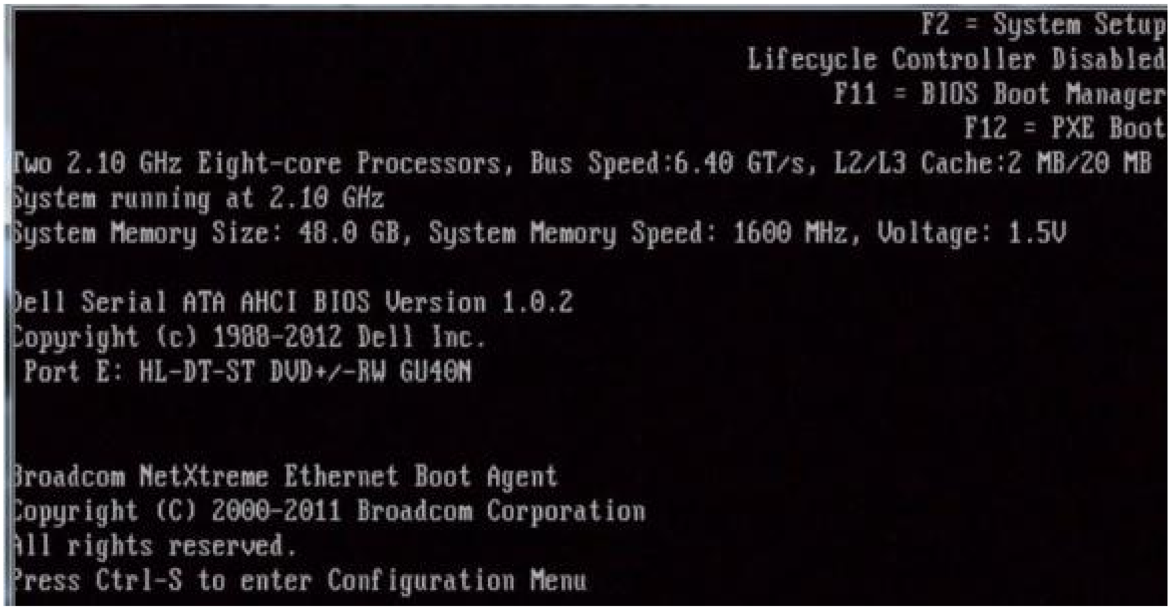

Figure 2. BIOS Settings menu

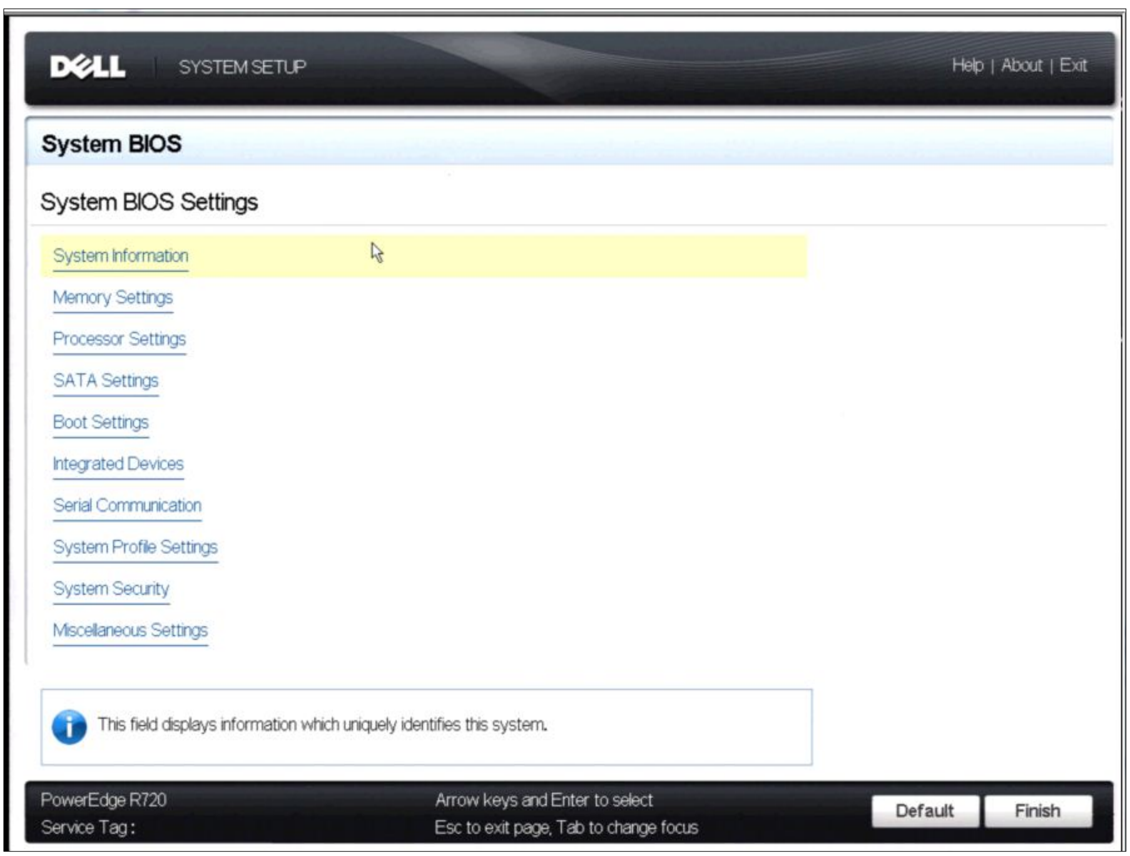

Figure 3. System Profile Settings menu

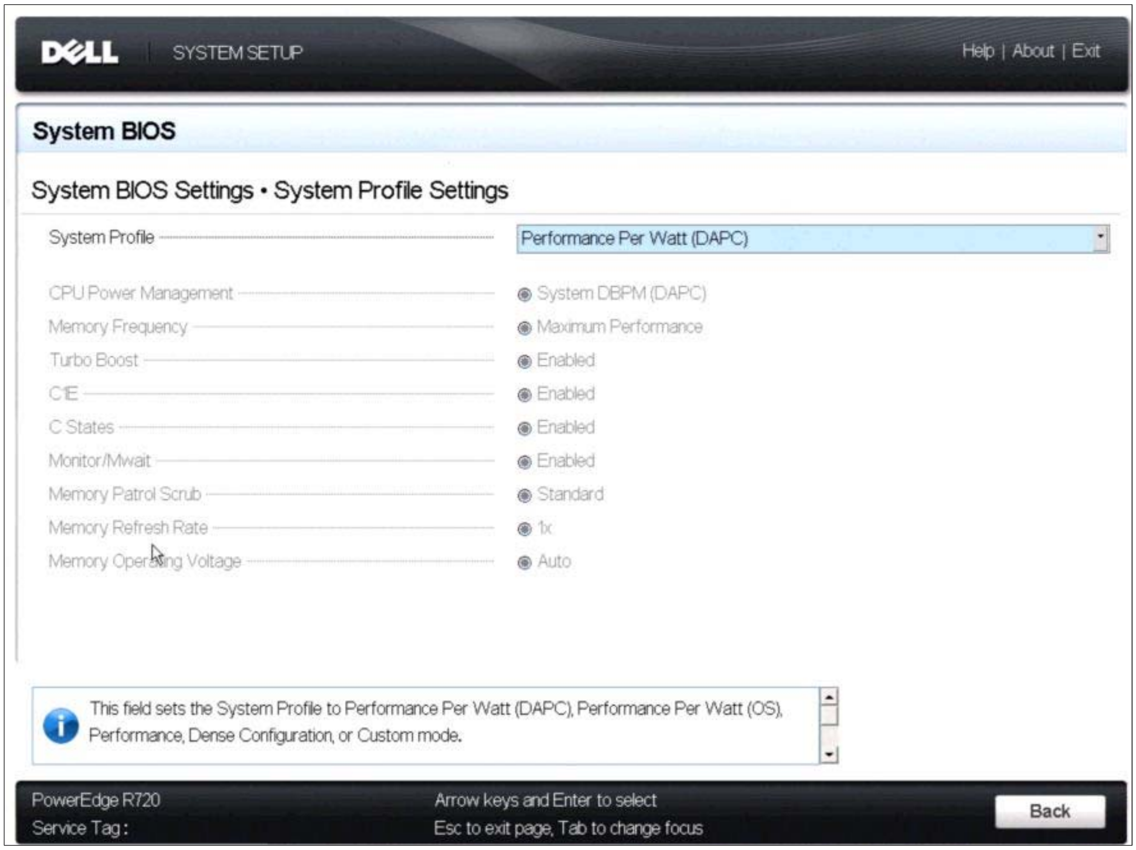

Microsoft® Windows Server® power plans can be adjusted through Control Panel menus, shown in Figure 4 and Figure 5:

Figure 4. Microsoft Windows Server 2008 power plan menus

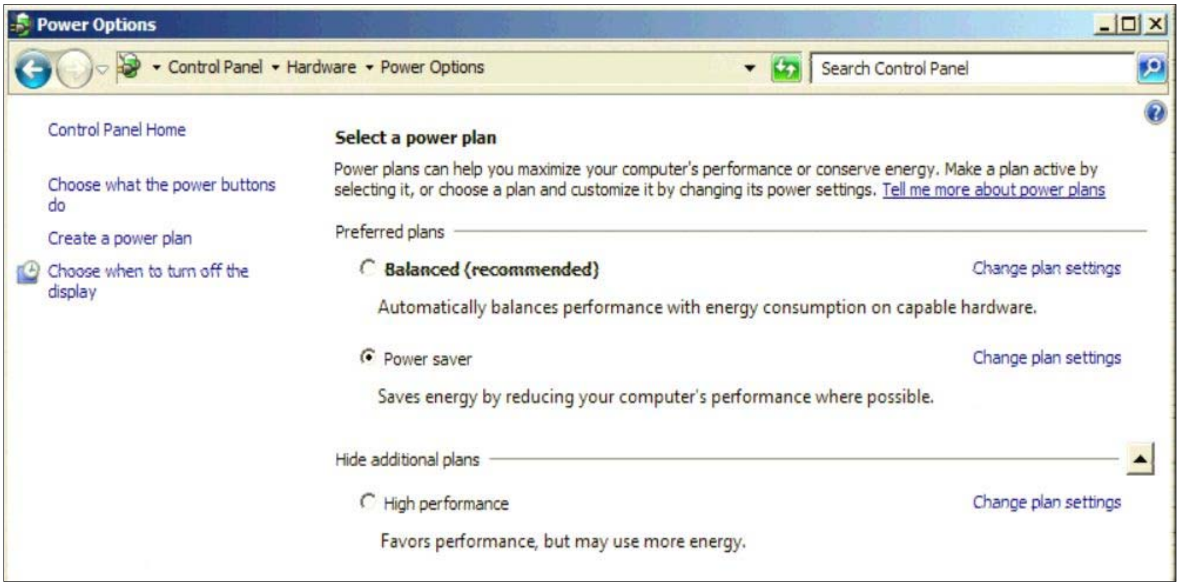

Figure 5. Microsoft Windows Server 2008 Power Saver plan settings

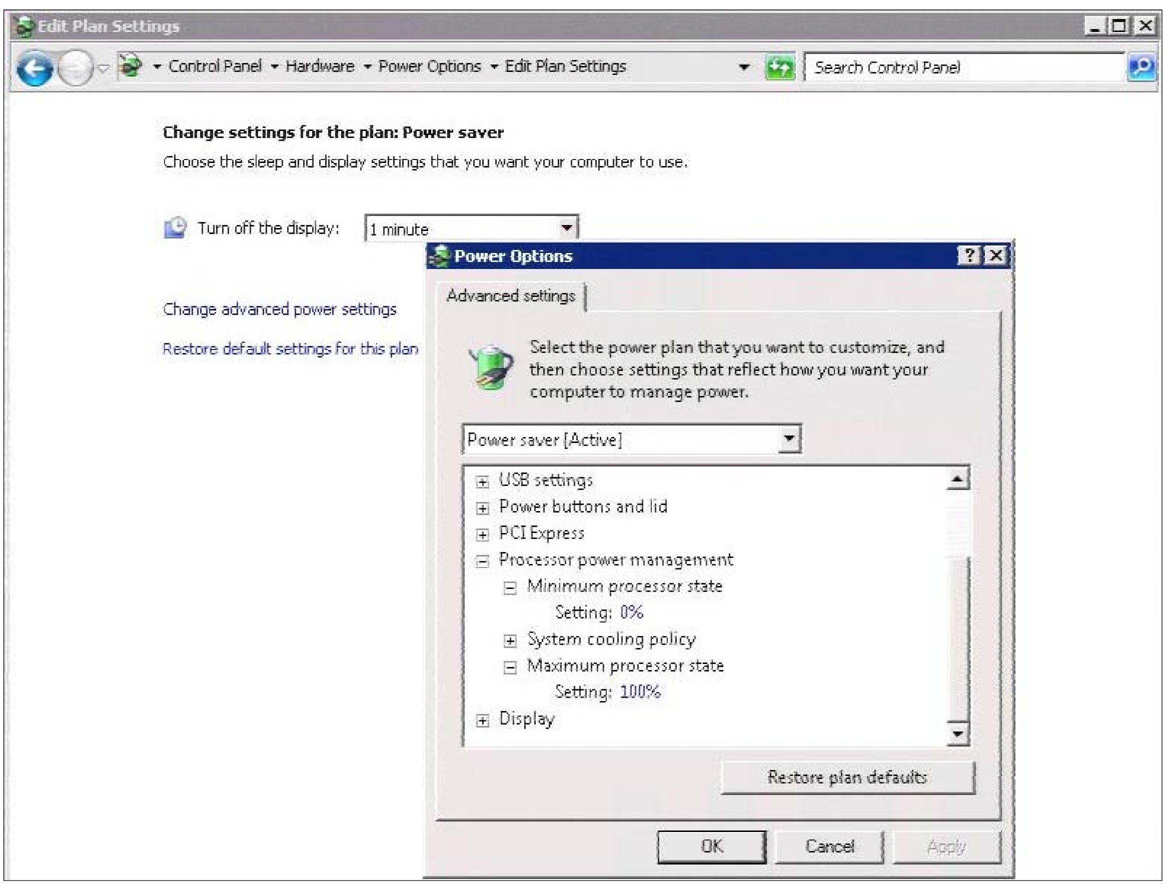

Test Methodology

We used the SPECpower_ssj2008 industry-standard benchmark[1] to compare server power efficiency results for each of the “easy button” user-selectable system profiles.

The SPECpower_ssj2008 metric is a calculation of the performance-to-power ratio across the full range of system workload levels from active idle to 100% utilization. The workload itself models transactions commonly associated with managing a business warehouse including new order, payment, order status, delivery, stock level, and customer report generation.

Each measurement point is of the number of ssj_ops (server side Java operations per second) processed per watt of power consumed. The more ssj_ops the system under test (SUT) can output for a given number of watts of power input; the better its efficiency.

The SPECpower measurement test bed shown in Figure 6 consists of four hardware components working together to collect a server's power consumption and performance data while the SUT is exercised with a workload exercising the SUT with a predefined workload :

- **Server Under Test (SUT)** is the system driven by the SSJ-based workload. The SUT’s performance and power consumption characteristics are captured and measured by the benchmark.
- **Power Analyzer** is used to measure and record the power consumed by the SUT.
- **Temperature Sensor** is used to capture the temperature of the environment where the SUT is being benchmarked.
- **Controller System** is a separate system that applies the SUT workload while capturing power and temperature measurements.

Figure 6. SPECpower test bed

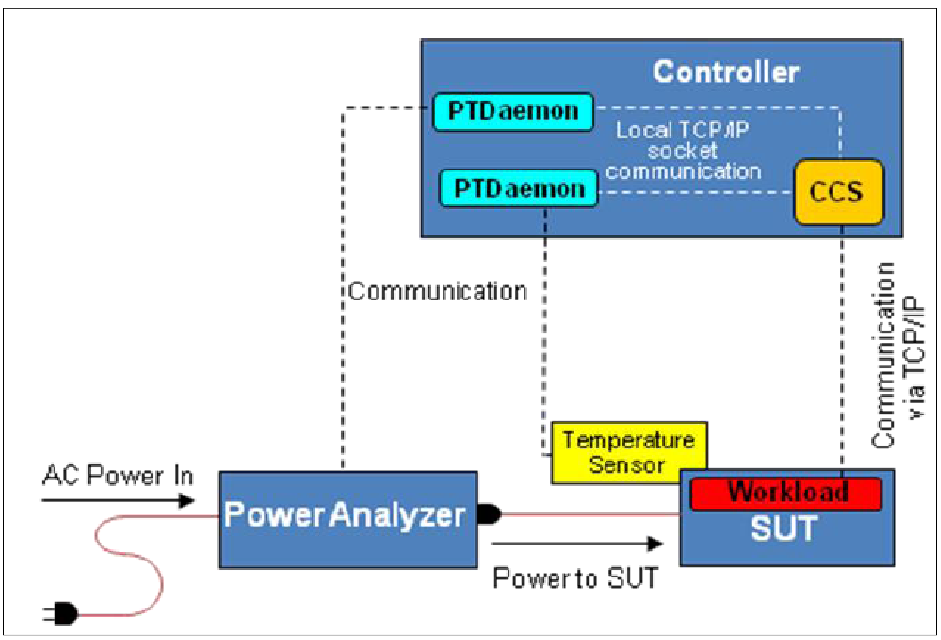

[1] See Appendix C: References for a link to the Standard Performance Evaluation Corporation’s SPECpower_ssj2008 website for more information about the SPECpower_ssj2008 benchmark.

Results 1: Performance per Watt (DAPC) versus Performance

Table 2. Compute throughput versus input power across the full workload range

Target load	(Perf/Watt DAPC) ssj_ops	(Perf/Watt DAPC) Power	(Performance) ssj_ops	(Performance) Power
100%	866,247	212	875,480	214
90%	787,737	194	791,399	207
80%	703,517	182	705,736	199
70%	613,679	167	618,043	192
60%	524,531	152	528,858	183
50%	437,506	140	441,469	176
40%	350,215	132	352,899	168
30%	262,010	121	264,335	160
20%	174,960	114	175,169	151
10%	86,747	105	88,520	143
Active Idle	0	83.6	0	134
Σssj_ops / Σpower =		**3,000**		**2,513**

Figure 7. Compute throughput versus system power consumption

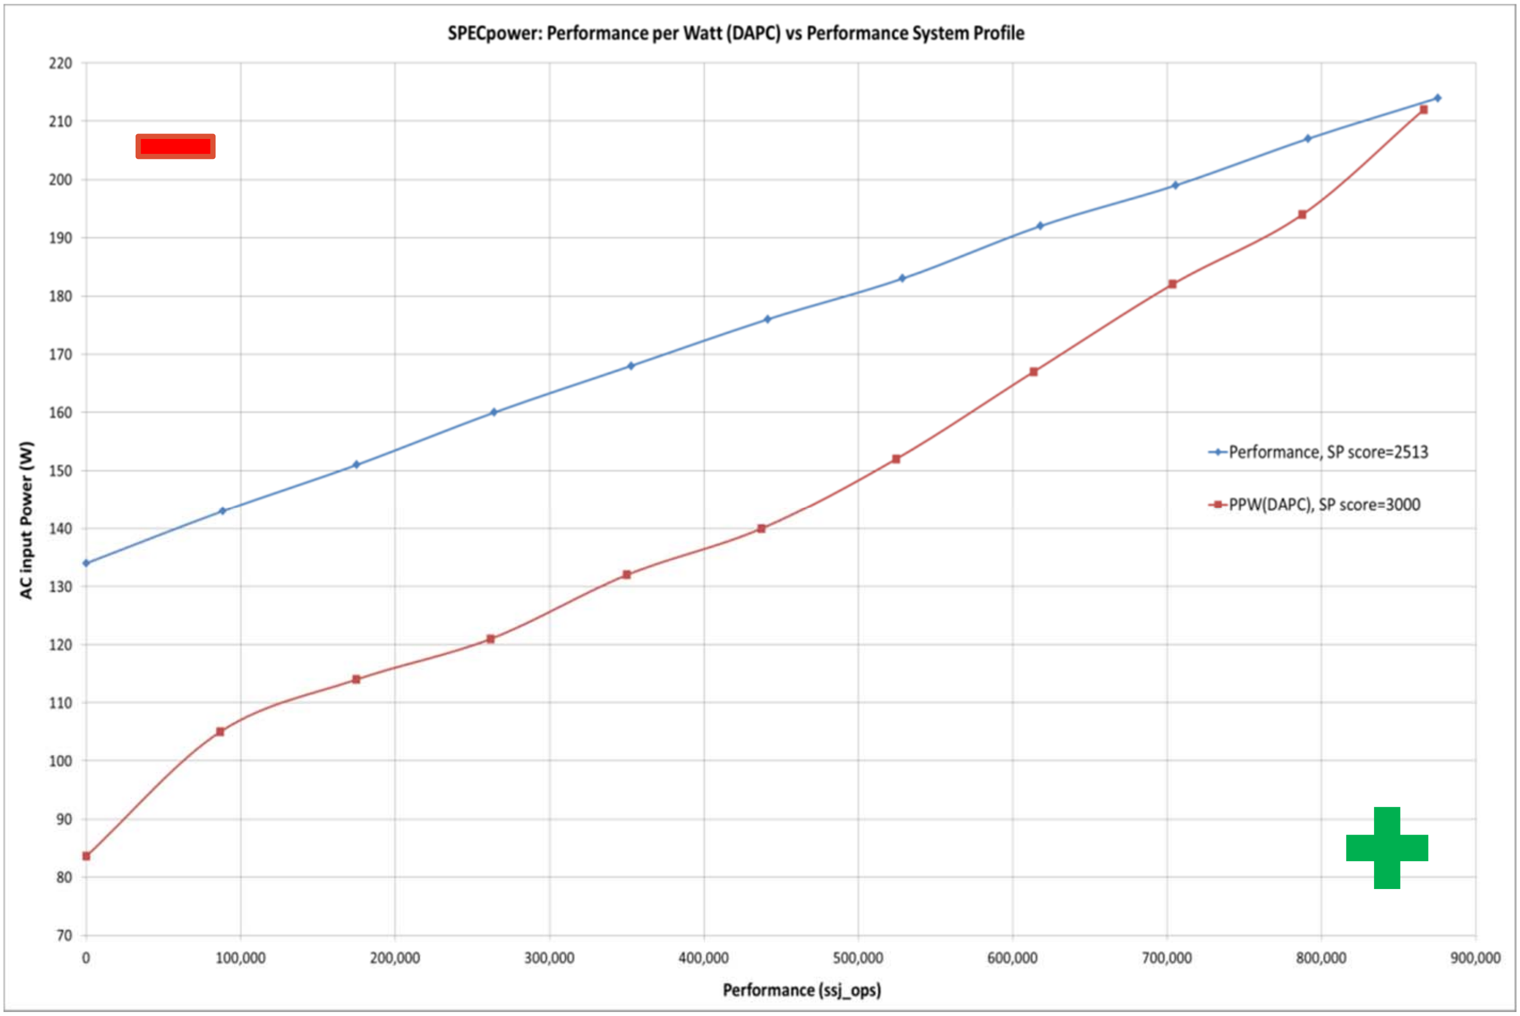

Summarizing the system behavior differences between Performance per Watt (**D**ell **A**ctive **P**ower **C**ontroller) and Performance-optimized profiles:

- Performance per watt (DAPC) system profile idle workload power advantage over Performance (W) = 50.4
- Performance per watt (DAPC) system profile idle workload power advantage over Performance (%) = 60%
- Performance per watt (DAPC) system profile overall efficiency advantage over Performance (%) = 19%

Results 2: Performance per Watt (DAPC) versus Performance per Watt (OS)

Table 3. Compute throughput versus input power across the full workload range

Target Load	(Perf/Watt DAPC) ssj_ops	(Perf/Watt DAPC) Power	(Perf/Watt OS-power saver default 5/75%) ssj_ops	(Perf/Watt OS-power saver default 5/75%) Power	(Perf/Watt OS-power saver optimum 0/100%) ssj_ops	(Perf/Watt OS- power saver optimum 0/100%) Power	(Perf/Watt OS-balanced) ssj_ops	(Perf/Watt OS-balanced) Power	(Perf/Watt OS-highperf) ssj_ops	(Perf/Watt OS-high perf) Power
100%	866,247	212	589,965	153	875241	215	866,247	215	879,249	215
90%	787,737	194	534,918	149	795212	203	787,737	206	795,177	207
80%	703,517	182	476,384	144	706992	182	703,517	185	704,295	199
70%	613,679	167	416,501	138	614191	172	613,679	174	620,136	191
60%	524,531	152	354,536	131	528268	160	524,531	164	529,418	182
50%	437,506	140	296,630	125	441732	145	437,506	152	443,465	173
40%	350,215	132	236,329	118	349904	132	350,215	139	354,411	163
30%	262,010	121	177,567	113	264109	121	262,010	127	266,751	151
20%	174,960	114	119,302	108	176166	113	174,960	116	176,732	137
10%	86,747	105	59,788	101	87413	105	86,747	106	88,166	118
Active Idle	0	83.6	0	83.2	0	83	0	83.6	0	84.7
Σssj_ops / Σpower =		**3,000**		**2,393**		**2,966**		**2,899**		**2,668**

Figure 8. Compute throughput versus system power consumption

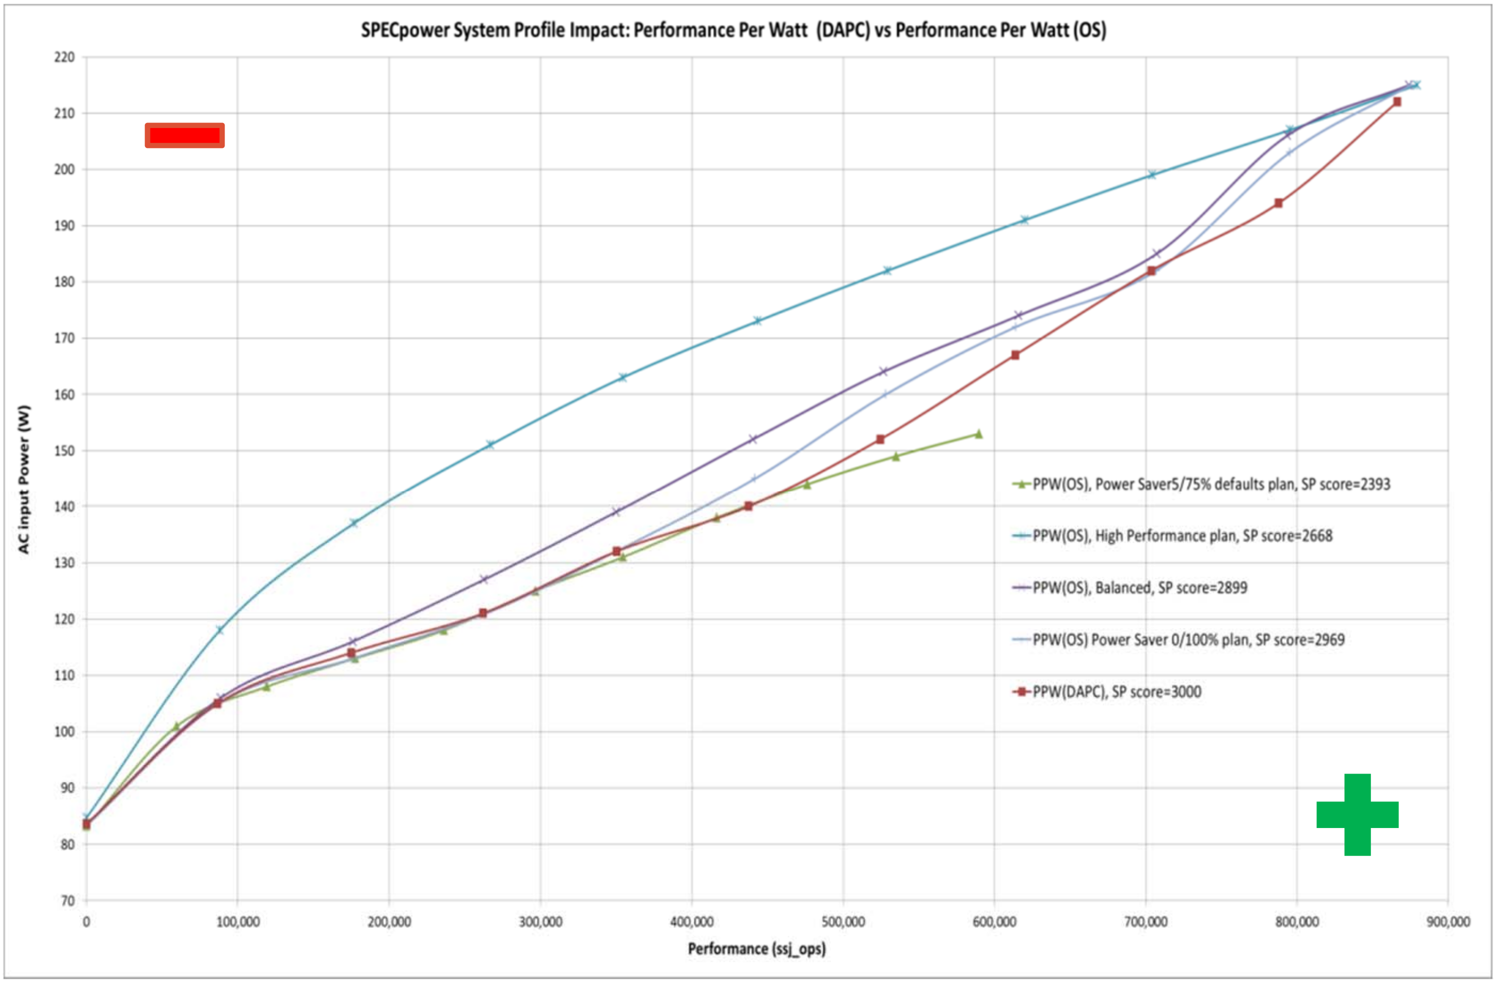

Key findings from a comparison of system responses under Performance Per Watt (DAPC) and Performance Per Watt (OS Windows Server 2008 R2 Enterprise) profile:

- Performance per Watt (DAPC) vs Performance per Watt (OS) – power saver 0/100% plan overall efficiency advantage (%) = 1%
- Performance per Watt (DAPC) vs Performance per Watt (OS) – balanced plan overall efficiency advantage (%) = 3%
- Performance per Watt (DAPC) vs Performance per Watt (OS) – high performance plan overall efficiency advantage (%) = 12%
- Performance per Watt (DAPC) vs Performance per Watt (OS) – power saver default 5/75% plan overall efficiency advantage (%) = 25%

Conclusions

Choosing the performance-per-watt optimized system profile instead of performance-optimized option reduces system idle power consumption by 60% while improving overall energy efficiency by 19% and impacting only peak compute performance by 1%.

The Dell Active Power Controller (DAPC) profile built into every PowerEdge server achieves even better overall energy efficiency than an OS-based one such as in Windows Server 2008 R2 Enterprise.

Version 1.10 Addendum

It is important to note that the performance-per-watt power management control algorithms evaluated here, base their adaptive system resource throttling decisions upon the CPU's reported total utilization at any given moment. Light workloads (e.g. ones that recruit fewer process threads than available CPU cores) may not exhibit sufficient total CPU utilization to trigger an instantaneous return to the system's highest throughput/lowest latency levels. For more detail, see the "12G PowerEdge Power Efficiency 'How To' v1.0" white paper listed in the appendix.

The SPECpower_ssj2008 metric used in this study is fully threaded and thereby models a server at its maximum potential just as it would be in a typical consolidated datacenter environment. Lightly-threaded, latency-sensitive application performance with less regard for system power consumption, may be better served by selecting alternate system profile settings.

Appendix A: System Under Test (SUT) configuration

The server configuration chosen for this study is one that was considered “typical” in that it closely resembles what customers were currently specifying for datacenter deployment. We used this same configuration for Dell’s 12th generation performance and power claims, and total cost of ownership (TCO) studies.

Table 4. Detailed system test configuration

Platform	Dell PowerEdge R620
Processor model	2 x Intel® Xeon® E5-2620
Processor frequency	2000 MHz
Processor L3 cache	15 MB
Physical cores	12
Logical processors	24
Memory frequency	1333 MHz
Memory details	8 x 4 GB 2Rx4 LV RDIMMs
Internal Storage	2 x 146GB 15K 2.5in SAS HDD
HBA	PERC H710p mini
BIOS version	1.2.1
iDRAC version	iDRAC7 Enterprise 1.06.06 (Build 15)
NDC	4 x 1Gb Base-T
Power supply quantity and rating	2 x 495W
Line Voltage	208VAC
Turbo Boost	Enabled
Processor prefetchers	Disabled
Logical processors	Enabled
Node interleaving	Disabled
C-states	Set by system profile
C1E	Set by system profile
Power management	Set by system profile

Appendix B: Power and Cooling Calculation Assumptions

To calculate a server's annual fully burdened cost per watt, we used an industry expert's formula[2] updated with a $0.11/kWh cost of power to come up with the figure of $2.69 per watt per year.

Figure 9. Annual operating cost of servers per watt

	Fully Burdened Annual Cost of Power	
Assumptions		
	Cost of power ($/kwh):	$0.07
	Cost of Facility ($):	$200,000,000
	Facility Amortization Years	15
	Size of Facility (Watts Critical Load):	15,000,000
	Cost of Money (%/year):	5%
	Average Power Consumption (%):	80%
	Power Usage Effectiveness	1.70
Fully Burdened Cost of Power:		
	Infrastructure Cost/Watt:	$1.28
	[=-PMT(AnnualCostMoney, FacilityAmortizationYears, FacilityCost, 0)/FacilityWattsCriticalLoad]	
	Power Cost/Watt:	$0.83
	[=FacilityWattsCriticalLoad/1000*AveragePowerConsumption*PUE*PowerCost*24*365/FacilityWattsCriticalLoad]	
	Fully Burdened Cost of Power (annualized):	**$2.12**
	[=+Infrastucture_CostPerWatt+PowerCostPerWatt]	

[2] See Appendix C: References for a link to the article "Annual Fully Burdened Cost of Power" by James Hamilton.

Appendix C: References

- The Standard Performance Evaluation Corporation's SPECpower_ssj2008 metric for volume, multi-node class computer efficiency. http://www.spec.org/power_ssj2008/
- Dell white paper: "12G PowerEdge Power Efficiency 'How To' v1.0" by John Jenne. http://en.community.dell.com/techcenter/extras/m/white_papers/20109038.aspx
- Dell white paper: "Dell Energy Smart Architecture and Power Management Adoption" by Brad Lawrence, John Jenne. http://en.community.dell.com/techcenter/extras/m/white_papers/20109094.aspx
- Dell white paper: "Right-Sized Power Systems: A Means to Improved Energy Efficiency" by Eric Wilcox, Mark Muccini http://en.community.dell.com/techcenter/extras/m/white_papers/20106404.aspx
- U.S. EPA "Energy Star Program Requirements for Computer Servers", v2.0. http://www.energystar.gov
- 80 PLUS computer power supply initiative. http://www.plugloadsolutions.com/80PlusPowerSupplies.aspx
- Annual Fully Burdened Cost of Power; James Hamilton; 2008: http://perspectives.mvdirona.com/2008/12/06/AnnualFullyBurdenedCostOfPower.aspx
- Dell Server BIOS LLD specification: "12G Dell Active Power Controller (DAPC) v1.1" by Wuxian Wu
- Dell's Power and Cooling Technologies website
- Microsoft OS Power Plan primer: http://windows.microsoft.com/en-us/windows7/Power-plans-frequently-asked-questions
- Dell PowerEdge R620 12th Generation Rack Server: http://www.dell.com/us/business/p/poweredge-r620/pd