To: Ryan Fogle, EPA and John Clinger, ICF

From: Pierre Delforge, NRDC

Date: October 10, 2017

Re: Computer Server Memory and Power Supply Power Demand Analysis

---

Contents

This study focuses on providing the evidence to support a more representative idle power memory allowance and stronger power supply (PSU) efficiency requirements at low loading points in the development of version 3 of the ENERGY STAR® specification for computer servers. This memo supplement NRDC's comments on Draft 2 of Version 3.0 of the computer server specification, submitted on 9/1/2017.

I. Memory Allowance

Additional memory dataset mapped against the proposed ENERGY STAR approach

Draft 2 Version 3.0 Specification for Computer Servers proposes an additional memory adder of 0.125(watts)*AdditionalMem(GB) beyond 4 GB of installed memory. We calculated the additional idle power required by 92, Dell and HP server configurations using online power calculators[1,2] provided by these manufacturers and used regression analysis to identify memory adder equations that would significantly reduce the amount of energy "given away" for high-memory configuration. The models we studied spanned the last two generations of server technology with launch dates ranging from 2014 to 2017.[3] According to

---

[1] Online versions may be found here: http://dell-ui-eipt.azurewebsites.net/#/, and https://paonline56.itcs.hpe.com/?Page=Index. We were unable to identify working calculators from either IBM and Lenovo.

[2] "HP developed the HP Power Advisor using data collected from testing HP ProLiant servers. Each test starts with a system fully configured with the maximum number of processors, memory, hard drives, expansion cards, and power supplies. Proprietary software exercises the processors to the highest possible power level and operates all peripherals while taking voltage and current measurements. Testing continues for all levels of processor support at all speeds, with different memory amounts and hard drive sizes. During development cycles, HP retests revised or updated servers to ensure calculator integrity." Source: http://resources.idgenterprise.com/original/AST-0027751_TC090816TB.pdf.

"Dell EIPT is a model driven tool supporting a large number of products and configurations for infrastructure sizing purposes. EIPT models are based on hardware measurements with operating conditions representative of typical use cases." Source: http://www.dell.com/calc.

[3] In general, we analyzed low-end server configurations with no storage or I/O cards for example. As a sanity check, we ran the number is a fully configured HP Proliant DL380 system with 16GB DIMMs and found lower watts per additional

HP and Dell online calculator results are based on actual power measurements across a vast array of server configuration options.

The Draft 2 ENERGY STAR approach resulted in much higher allowances than justified by the calculators, in effect providing a giveaway for these configurations, particularly for high-end configurations as can be seen in the charts below, one of which is zoomed in to provide additional detail.

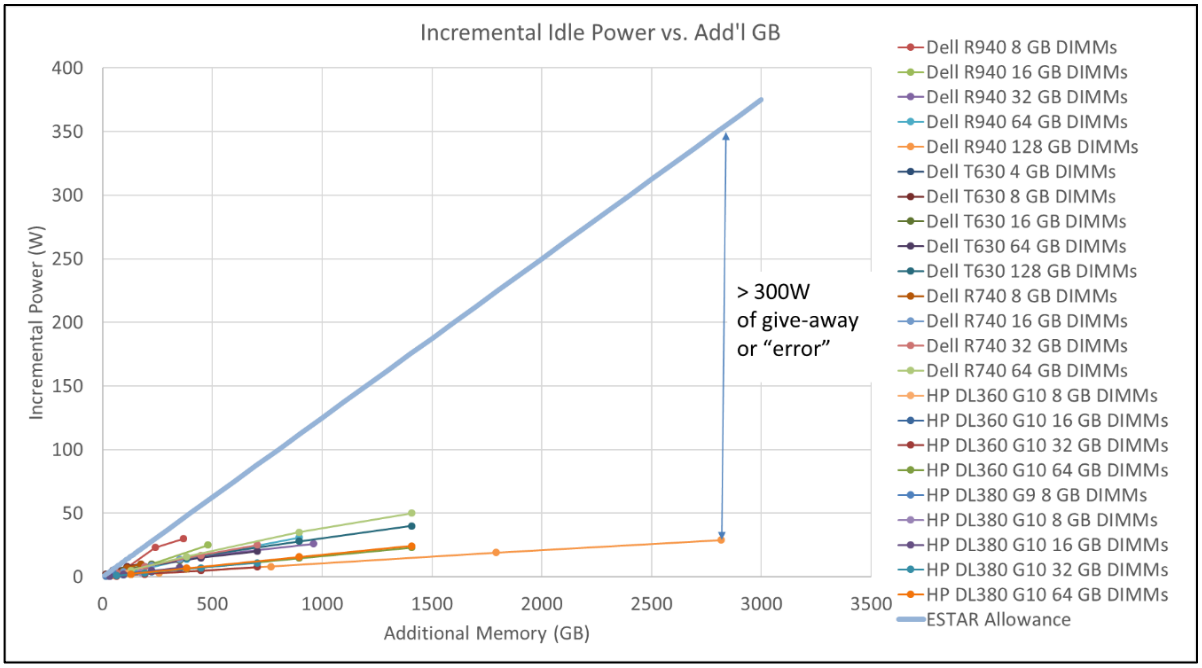

GB (0.017 vs 0.019 W/GB for the low-end system). See “Data (GB)” tab of the MemAnalysis workbook. The configurations analyzed are documented in far-right tabs of this workbook.

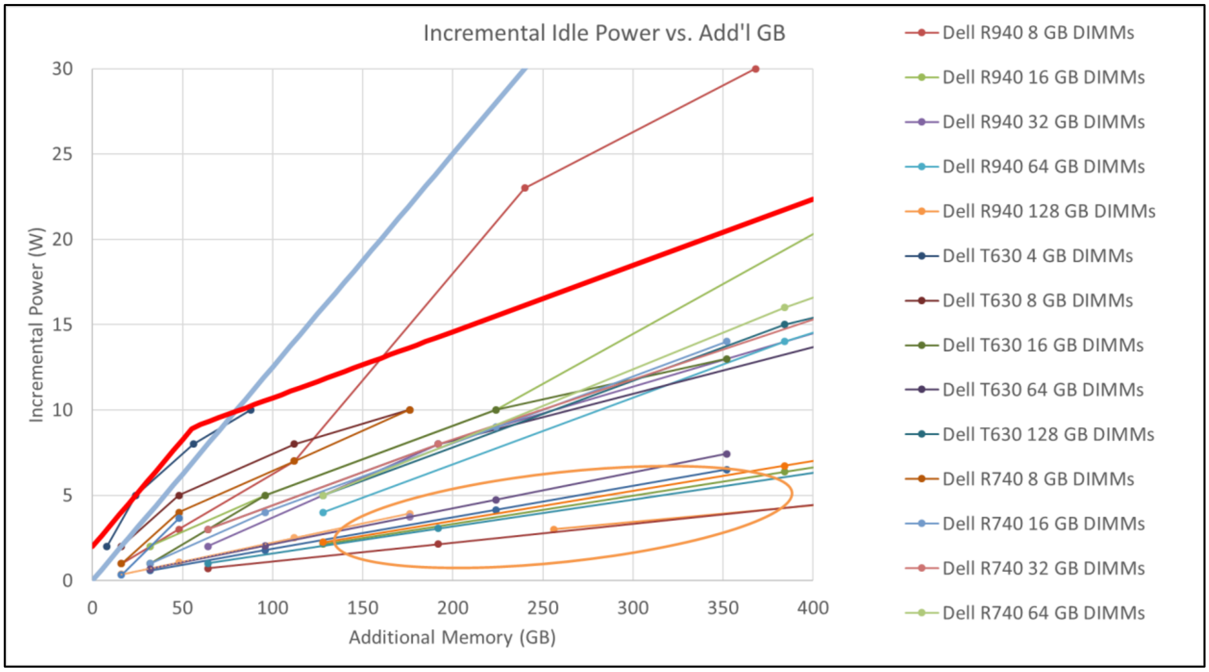

We have identified the most efficient cluster of configurations with the orange circle in the above chart. These lines have slopes of less than **0.033 W/GB**. These configurations have larger average DIMM module size then the less efficient configurations. And they tend to represent the latest server generation.[4]

The differences between the allowance and estimated power draw from the calculators can also be represented as absolute and percent error as shown in the charts below. The ENERGY STAR approach reflected an average percent error of 32% across on server configurations studied.

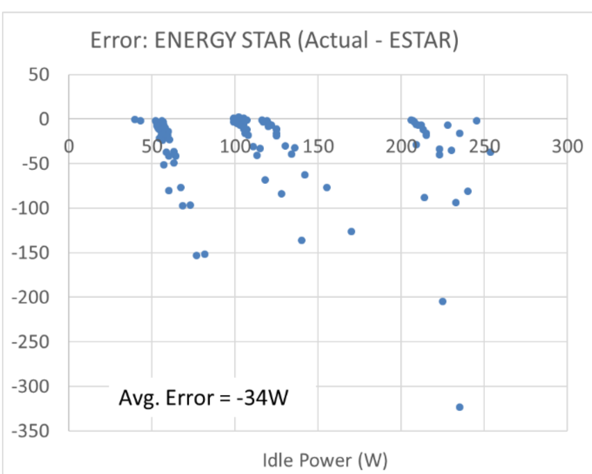

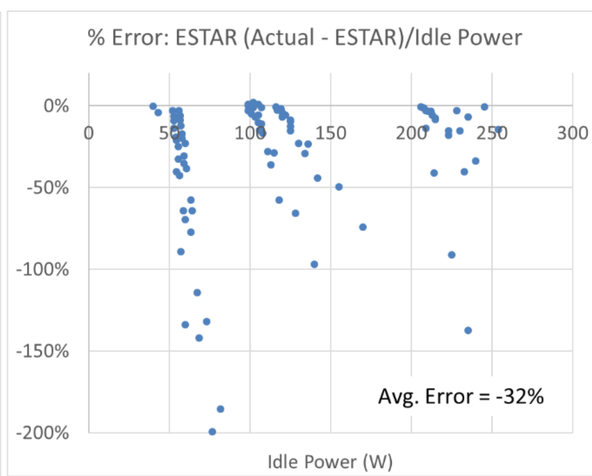

[4] See MemAnalysis workbook, "Data Sort WpGB" tab for data sorted by W/GB.

The most accurate, but complex, approach

In order to find an equation that would best fit the dataset without putting undue pressure on, or being overly generous to, systems of any size, we ran many regression analyses against variables that had the potential to impact server idle power. These include:

- GB of additional memory
- Number of added DIMMs
- DIMM size
- DIMM type (DDR generation, RDIMM vs. LRDIMM) – we looked briefly at NVDIMMs but did not perform detailed analysis[5]
- Memory speed (MT/s)
- Memory rank
- Idle power of baseline system

We achieved minimum average error of 10% in regression runs that included the following variables:

- GB of additional memory
- DIMM size
- Idle power of baseline system

The charts below compare these results to the above results using the Draft 2 ENERGY STAR approach. Note that we have limited the y axis range to -200 to accentuate the difference where data is clustered; as shown above, there are a few data points beyond -200 in the ENERGY STAR chart.

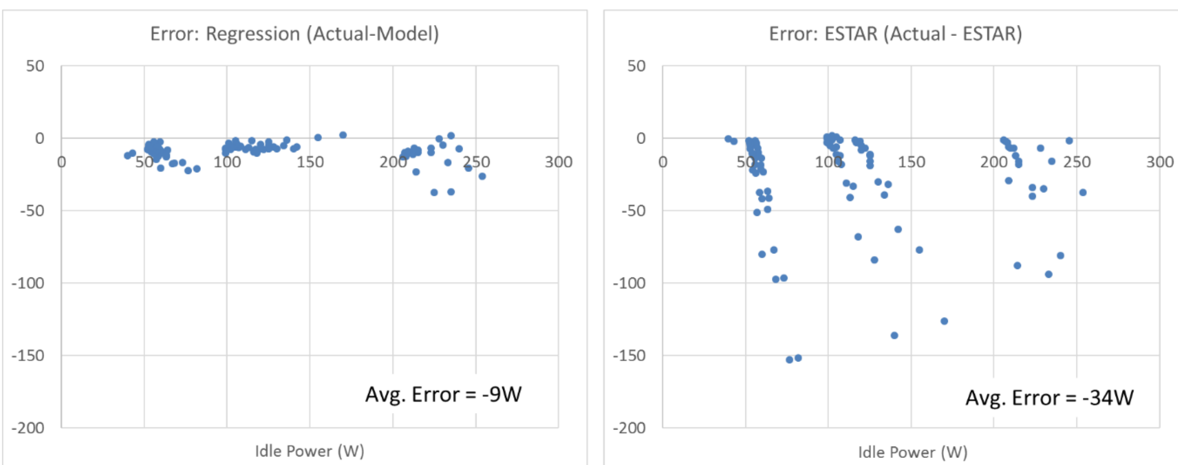

[5] Memory terms explained: http://www.dell.com/learn/us/en/04/campaigns/poweredge-memory

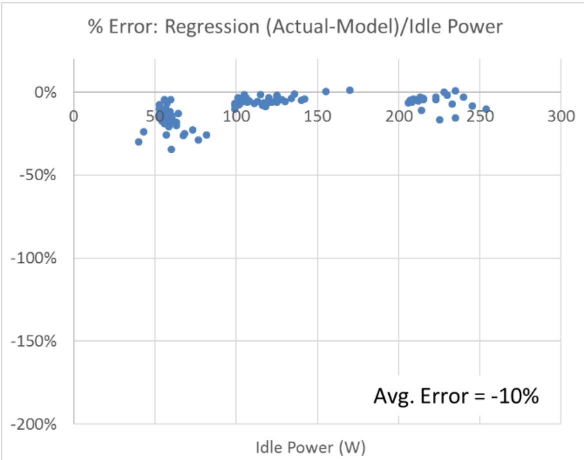

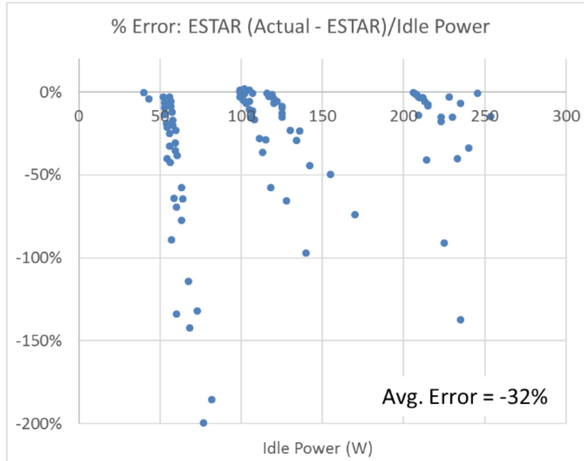

The equation for this approach is:

$$IncrementalPower(W) = -7.2 + (0.05 * Additional\,Mem(GB)) - ((5.25E-06) * Additional\,Mem(GB)^\wedge 2) + (422 * (-0.014 * Ln(DIMMSize) + 0.0747)) + (0.05 * BaselinePower(W))$$

The best balance between simplicity and accuracy

We found the best balance between simplicity and accuracy using the following equation, reflected in the plot below the equation:

$$IncrementalPower(W) = Min(2 + 0.125 * Additional\,Mem(GB), 6.82 + 0.039 * Additional\,Mem(GB))$$

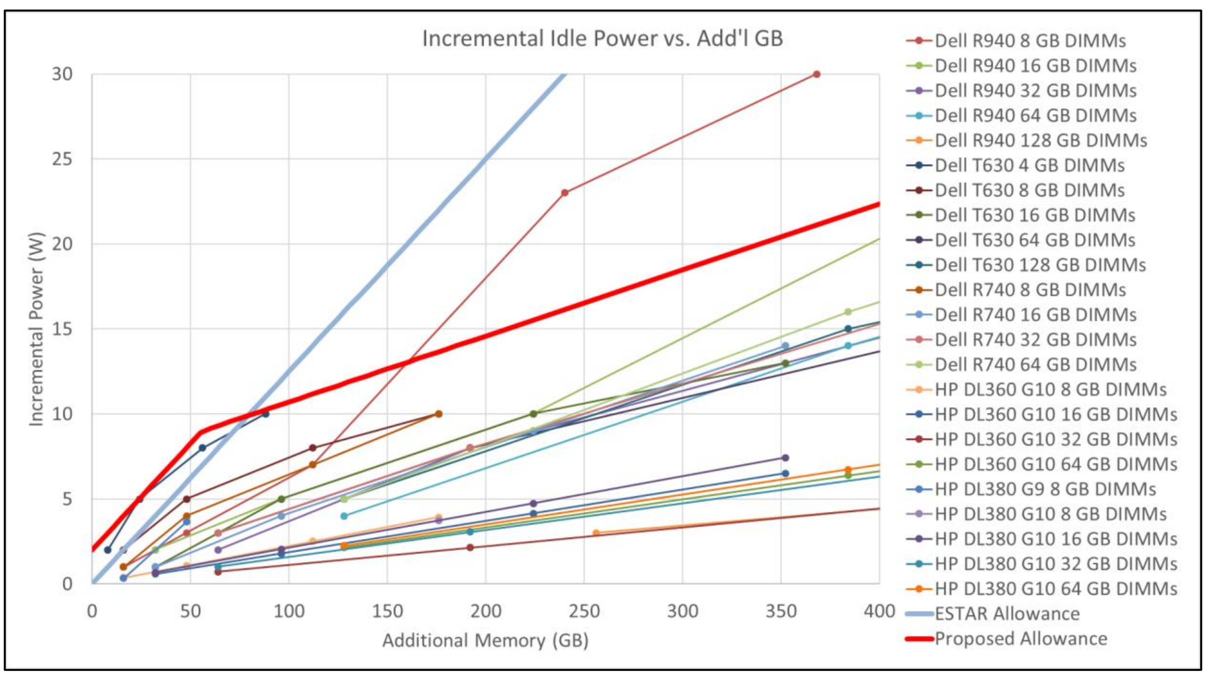

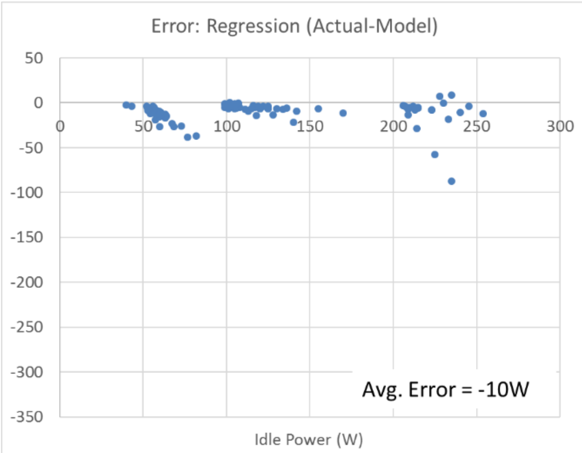
Error: Regression (Actual-Model)
50
0
-50
-100
-150
-200
-250
-300
-350
0
50
100
150
200
250
300
Avg. Error = -10W
Idle Power (W)

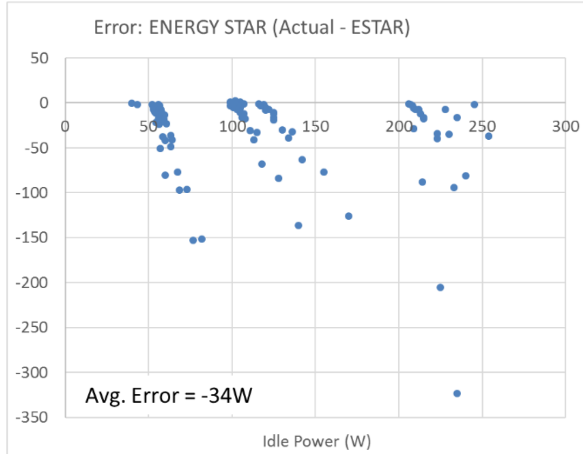
Error: ENERGY STAR (Actual - ESTAR)
50
0
-50
-100
-150
-200
-250
-300
-350
0
50
100
150
200
250
300
Avg. Error = -34W
Idle Power (W)

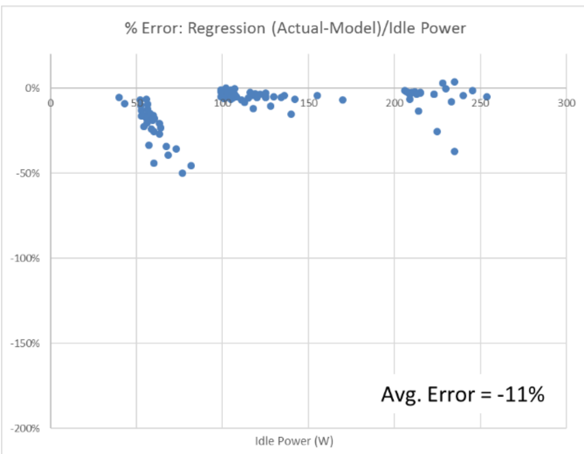
% Error: Regression (Actual-Model)/Idle Power
0%
-50%
-100%
-150%
-200%
0
50
100
150
200
250
300
Avg. Error = -11%
Idle Power (W)

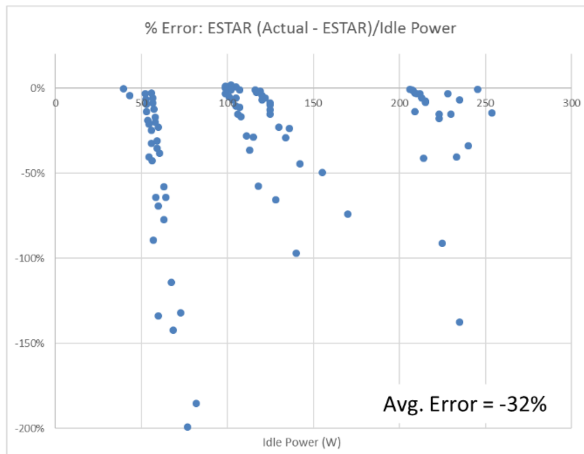
% Error: ESTAR (Actual - ESTAR)/Idle Power
0%
-50%
-100%
-150%
-200%
0
50
100
150
200
250
300
Avg. Error = -32%
Idle Power (W)

Using DIMM count as a metric

Finally, we performed one last analysis that explored the possibility of using DIMM count as the basis for an additional memory adder. The best average error we could get is -19% (-18W). It is clear from the plot below that there is a large range of incremental power values associated with each additional DIMM count.

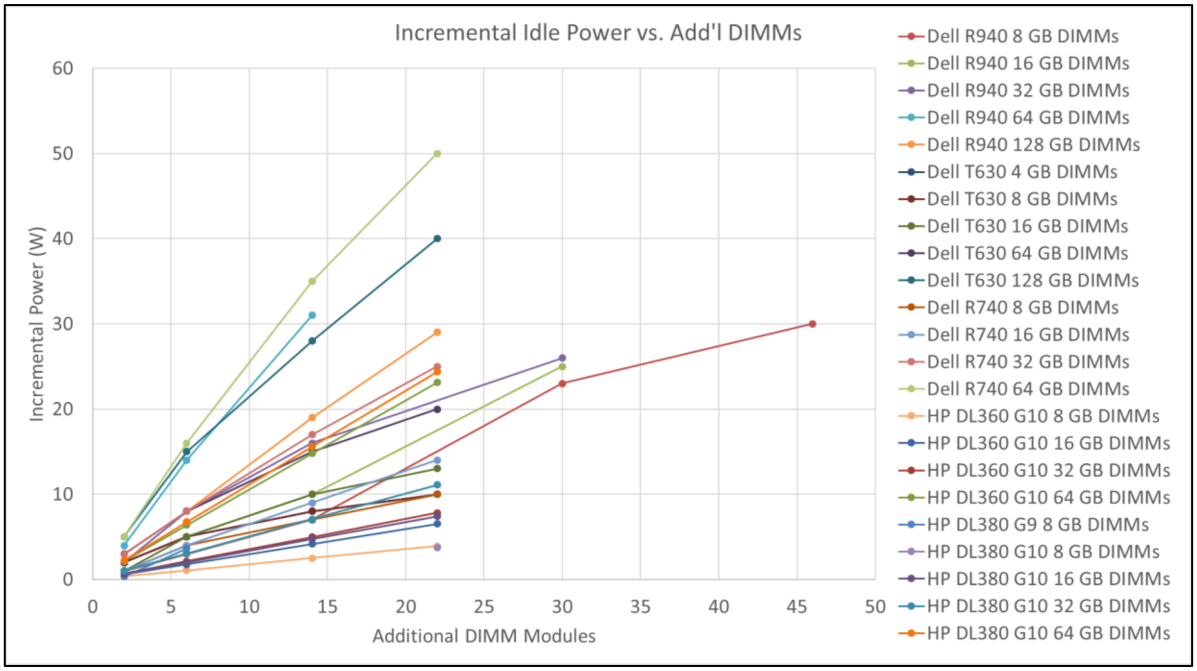

II. Power Supply Efficiency

As noted in NRDC's September 1, 2017 comments to EPA regarding Draft 2 Version 3.0 Specification for Computer Servers, the average load of typical servers is in the single digits or low double digits. Power supplies are typically oversized to ensure they can support periods of maximum load reliably without burning out, and many servers use redundant power supplies that share the load between both power supplies, leading to the load on each power supply being half of what it would be on a single power supply or in a failover configuration.

Key Finding: The average server power supply spends most of its time in the 3 to 20 percent load range.

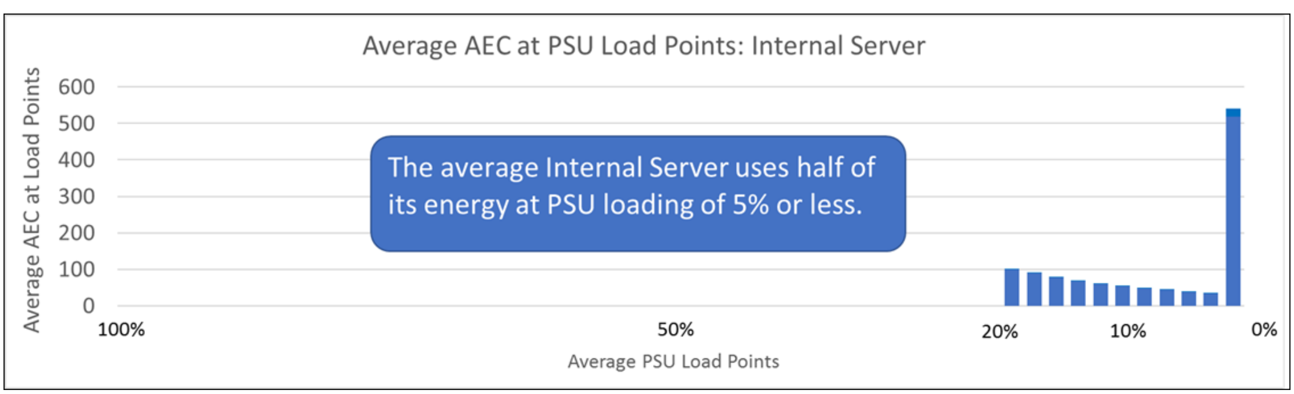

Recommendation: Revise power supply efficiency requirements to better represent real-world loading conditions, and to maximize the energy savings effect of PSU efficiency requirements.

Analysis

Below we document how we derived these numbers, we include the other server classes defined by LBNL, and we estimate how much energy could be saved by bringing a less efficient Platinum power supply up to the levels proposed by NRDC in their September 1st comments.

Our analysis of the SpecPower database[6] reflects the following energy distribution by PSU load level for what LBNL calls Internal Servers in their 2016 U.S. Datacenter Energy Efficiency Study.[7]

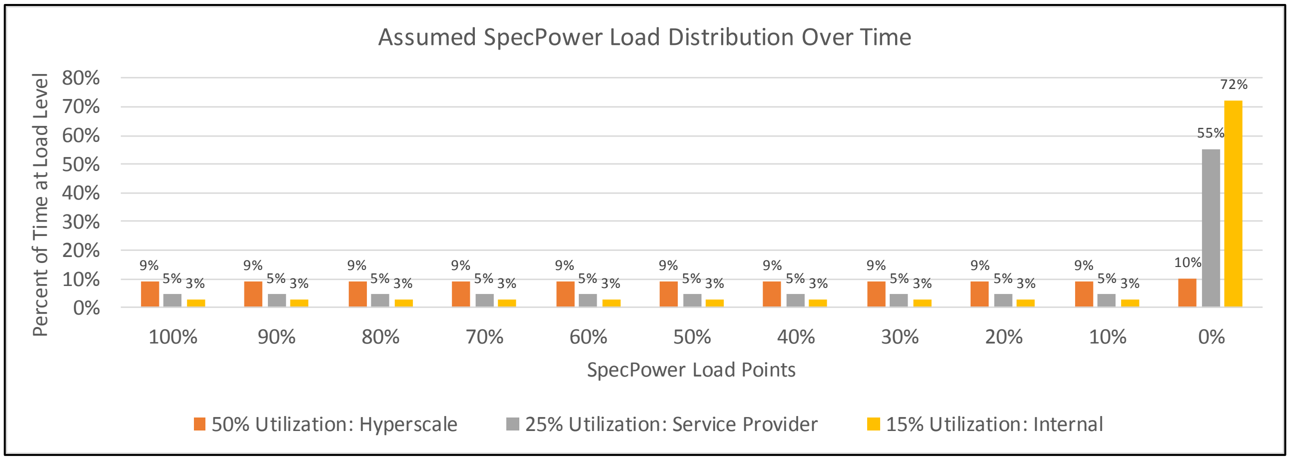

These time distributions result in average power distributions (shown in the legend above) that match those reported in the LBNL study for the year 2020.

We then analyzed the SpecPower dataset[8] to find the *PSU loading* distributions at each *SpecPower system load level* for models launched in the 2014-2017 timeframe. We assumed that all listed servers have redundant power supplies that share the load in normal operations without "Hot Spare" or equivalent feature enabled.[9] In other words, we added a second PSU to all servers reported to have a single PSU, thereby doubling the reported PSU capacity.

---

[6] File name: "power_ssj2008-results-20170920-182840.csv" downloaded from https://www.spec.org/cgi-bin/osgresults?conf=power_ssj2008

[7] https://eta.lbl.gov/publications/united-states-data-center-energy

[8] https://www.spec.org/power_ssj2008/results/

[9] For background on "Hot Spare" technology, see http://i.dell.com/sites/doccontent/business/solutions/whitepapers/en/Documents/hot-spare-whitepaper.pdf or http://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/power-and-cooling-innovations_030216.pdf. Our understanding is that this or equivalent technology by other manufacturers is not commonly used.

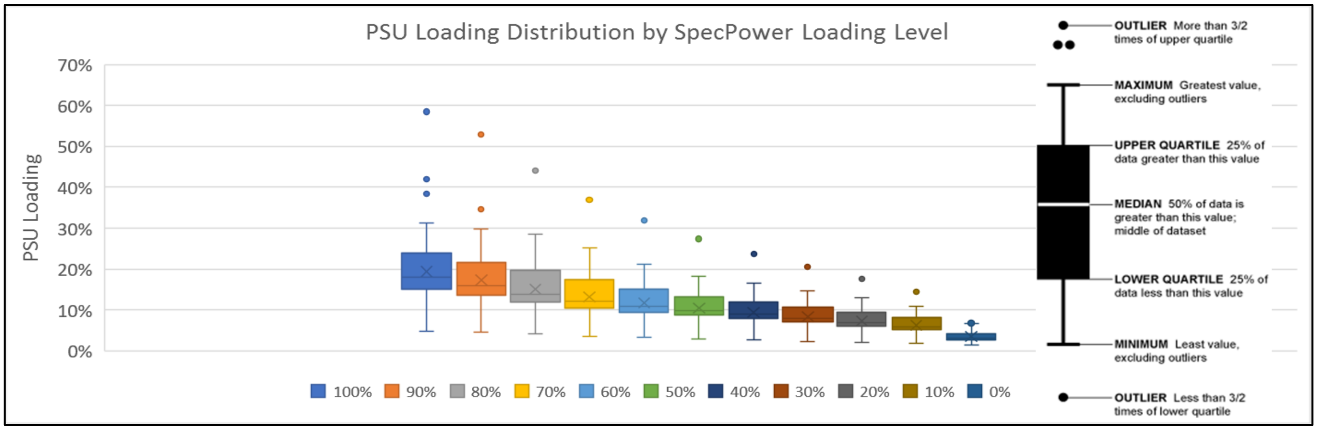

We then used the average PSU loading levels shown above to plot the average watts over time for these server classes:

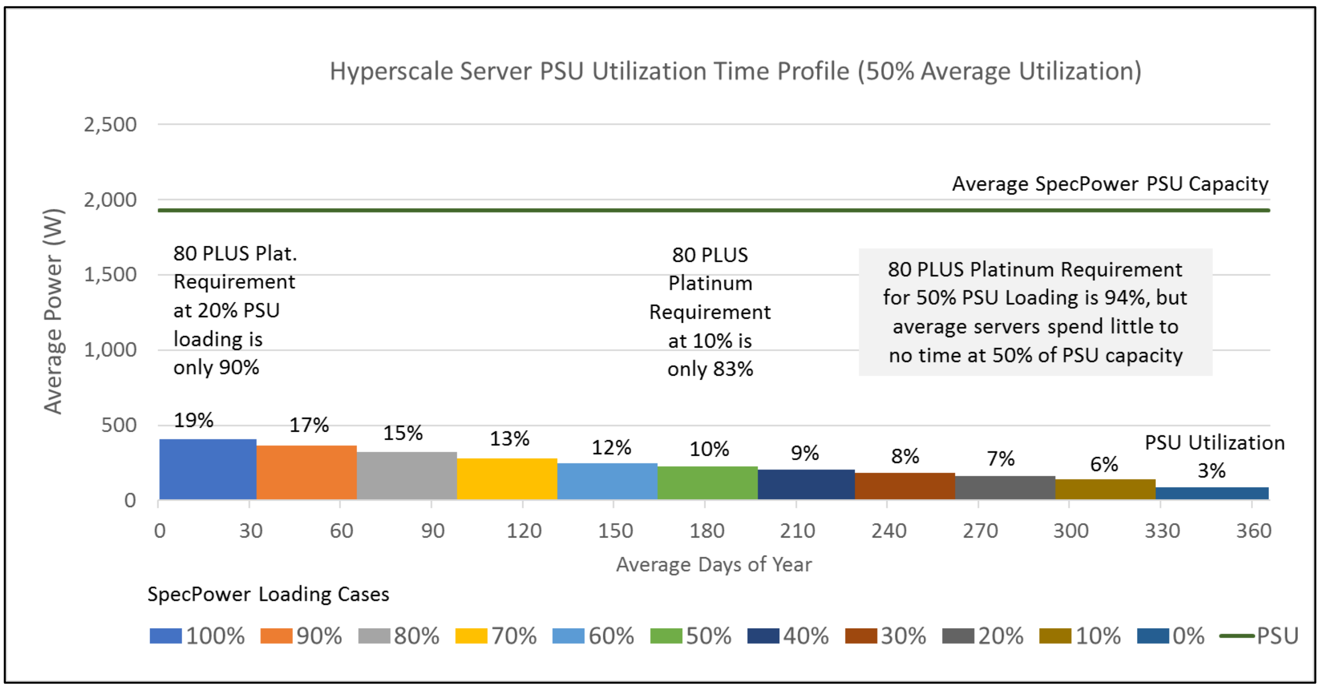

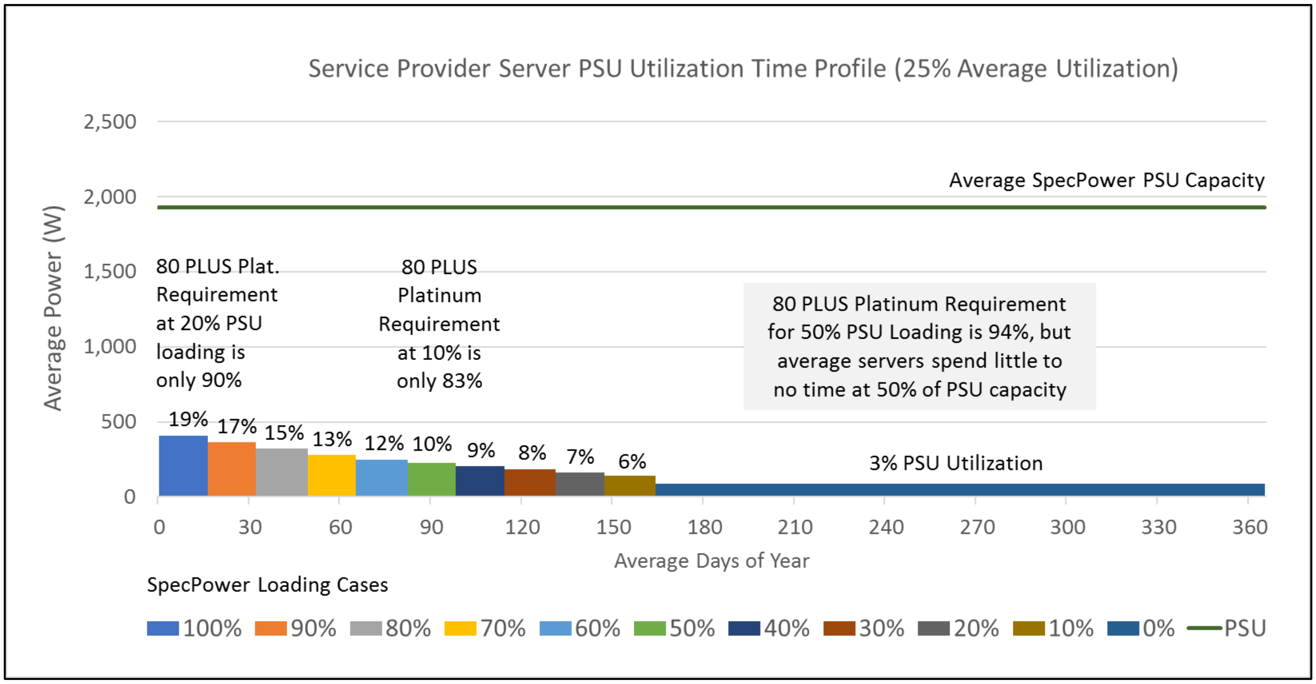

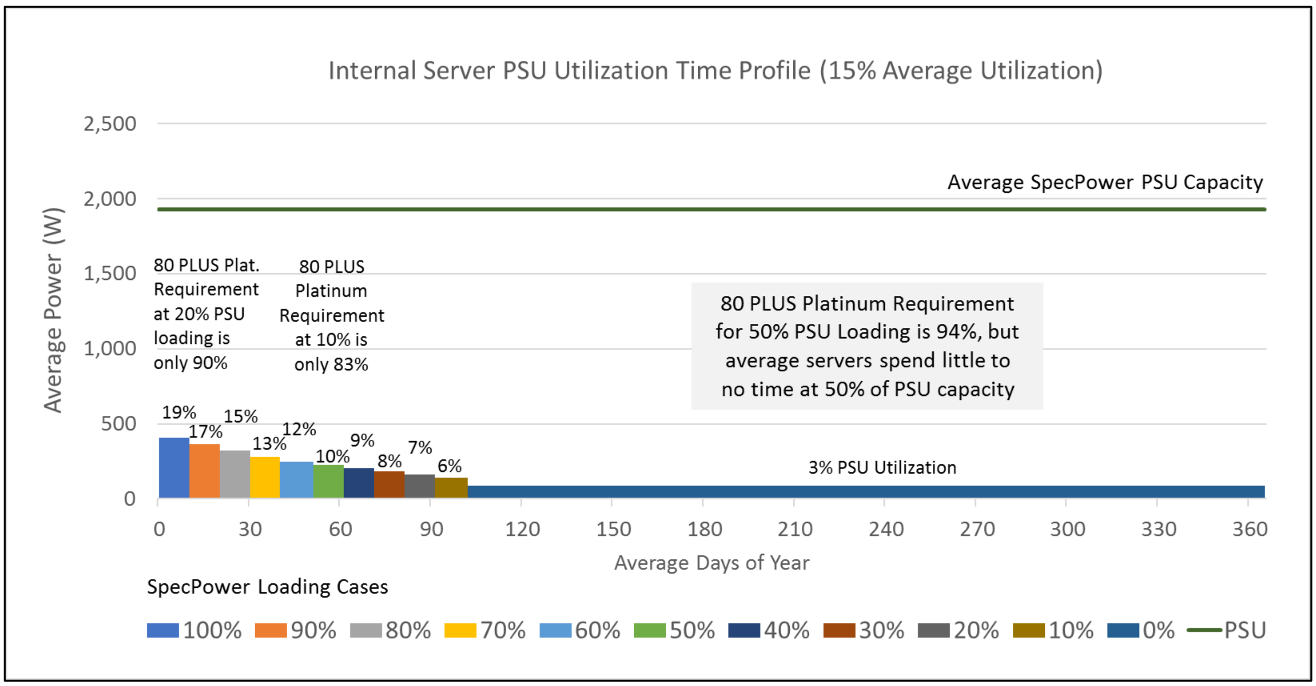

We then plotted the average energy (area of the boxes above) by PSU load point first as a percent:

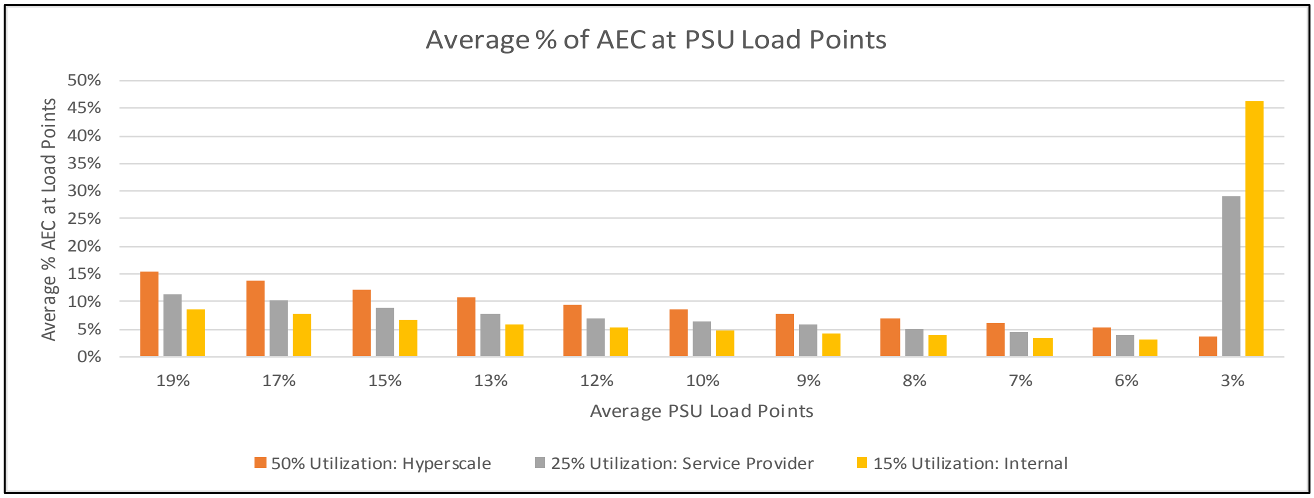

then as kWh/yr:

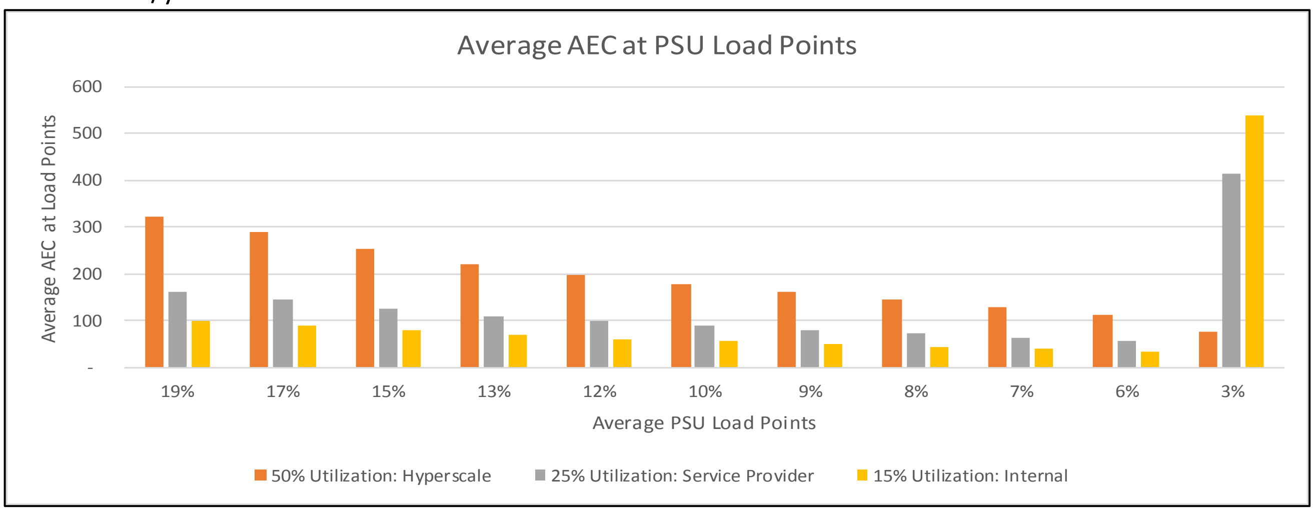

To better visualize how this distribution of energy relates to 80 PLUS test points, and to emphasize the importance of PSU efficiency at load levels, we then plotted the average energy distribution for Internal Servers (15% average system loading) with a PSU load level scale that spans 0% to 100%:

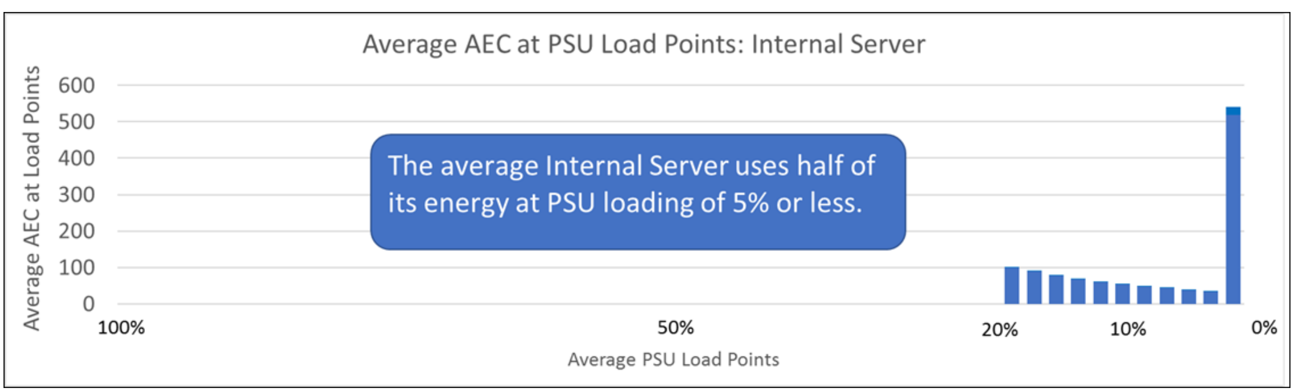

We analyzed the 80 PLUS database to understand the distribution of efficiency levels among Platinum servers in the charts below. And we included as open circles the proposed changes to efficiency allowance levels noted in NRDC's September 1st comments.

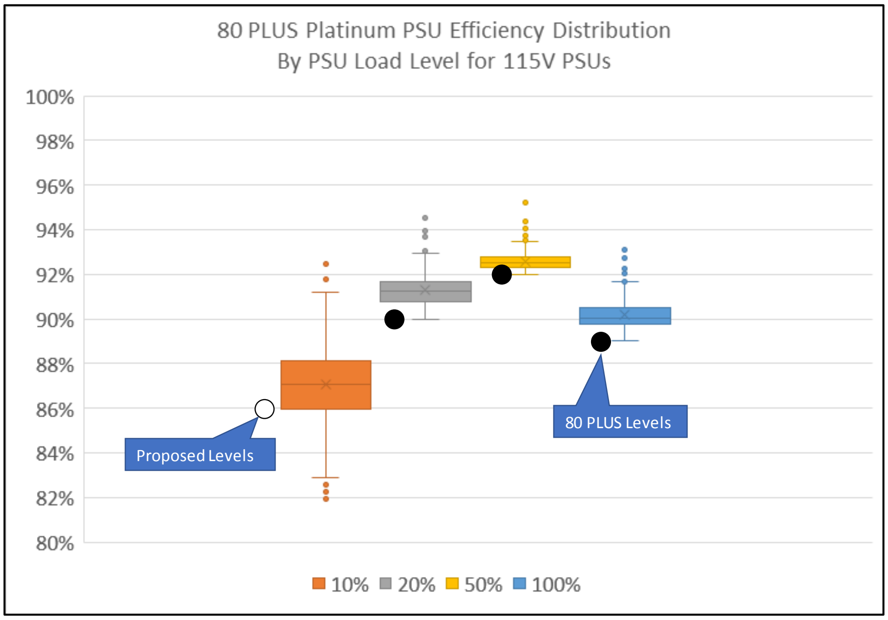

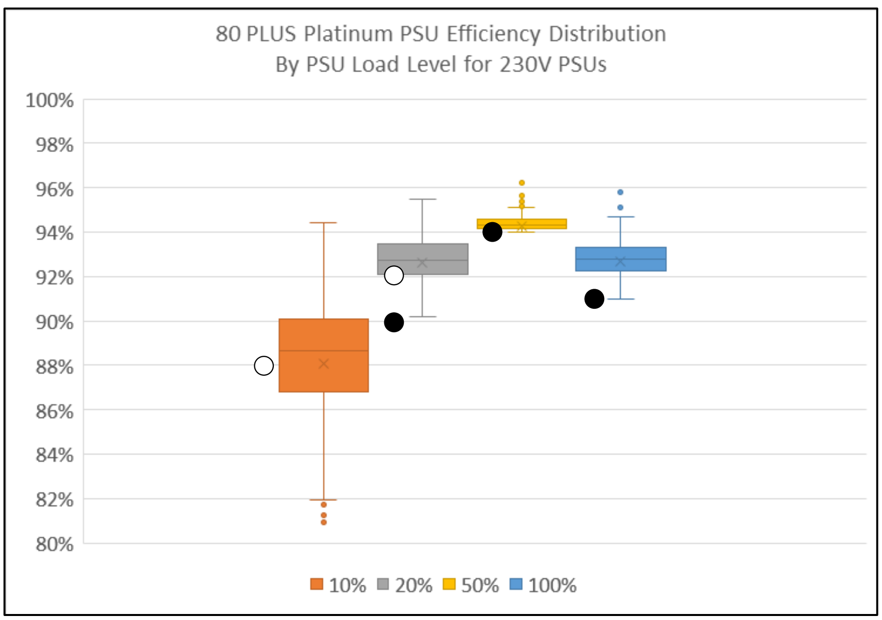

Based on this data and lacking clarity about whether the servers in the SpecPower database had 115v or 230v power supplies, we calculated savings for the average Internal Server as follows. The second row of the table shows our PSU savings assumptions by PSU load level relative to an 80 PLUS Platinum power supply at the lower end of the efficiency spectrum.

Savings estimates for the average Internal Server

PSU Loading	19%	17%	15%	13%	12%	10%	9%	8%	7%	6%	3%	Total
Potential PSU Savings	1%	1%	2%	2%	2%	4%	4%	4%	4%	4%	4%	3%
Average AEC (kWh/yr)	100	90	79	69	61	55	50	45	40	35	541	1,165
Potential PSU Savings (kWh/yr)	1	1	2	1	1	2	2	2	2	1	22	37
Average AEC w/Savings (kWh/yr)	99	89	77	67	60	53	48	43	38	34	519	1,128

A 4% improvement in PSU efficiency at low load levels results in a 3% system-level efficiency gain because of the high proportion of energy used at low PSU load levels.

These savings are reflected as the blue bars in the following stacked bar chart.

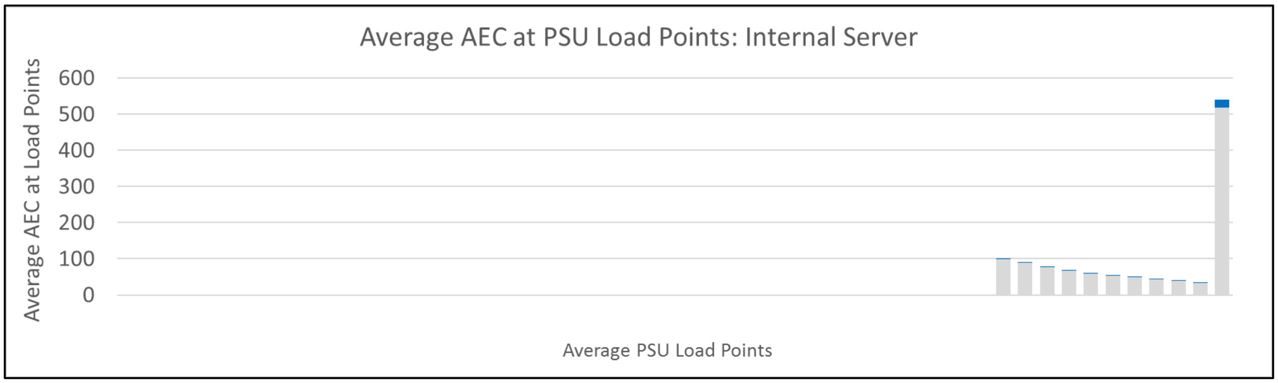

Appendix 1: Dataset Details

The goal of this analysis is to develop recommendations for ENERGY STAR program requirements in the 2019 timeframe. To do this, we mined data from both the ENERGY STAR QPL and the SpecPower database. SpecPower data, in particular, goes back many years / server generations. In order to understand how far back (in generations or years) to look in these databases, we evaluated the distribution of server “Year Available” in the QPL. Approximately 80% of the servers on the QPL were launched in 2014 and later.

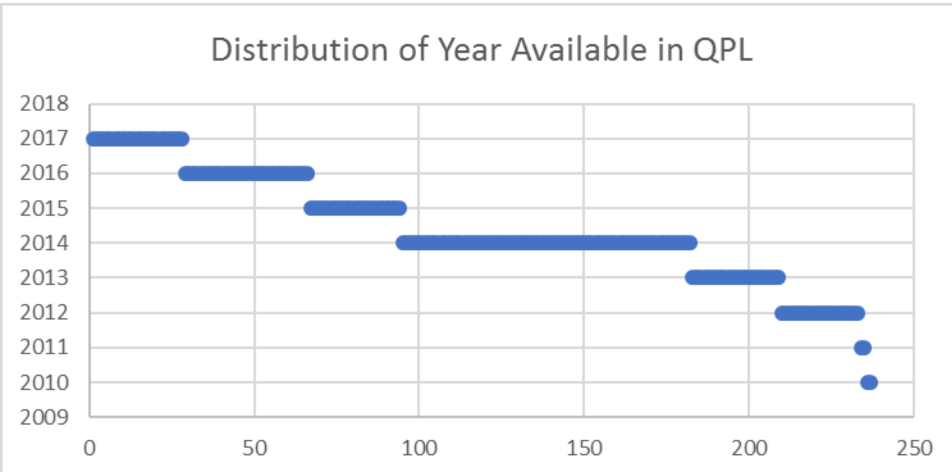

To understand how launch year maps to server generation, we mapped these characteristics for HP servers listed in both the QPL and the SpecPower database. Server generations are not indicated by other manufacturers.

Three generations of HP systems are represented in the QPL. The chart below maps them by year available:

QPL

Gen	2012	2013	2014	2015	2016	2017	Tot
Gen8	10		1				11
Gen9			8	4	6		18
Gen10						1	1

A similar analysis of the SpecPower dataset reveals the following:

SpecPower

Gen	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	Tot
G3		1	1												2
G4	1		3												4
G5			1		15	4									20
G6						8	2								10
G7							3	3							6
G8									1						1
G9												2	1		3
G10														2	2

Given that server generations are not generally indicated by manufacturers in either the QPL or the SpecPower spreadsheet, we’ll limit our analysis in general to server systems available in 2014 or later.

Appendix 2: Memory Configurations

Servers analyzed represent the most recent two generations from HP and Dell, with launch dates spanning from 2014 to present. The table below provides a summary of tested configurations. Additional details are captured in the accompanying workbook.

Additional Memory Calculations Table

Brand	Model	Memory Size (GB)	No. of DIMMs	DIMM Size	Baseline Memory Size (GB)	Server Gen	Server Available
Dell	PowerEdge R940 G14	32	4	8	16	Dell G14	2017
Dell	PowerEdge R940 G14	64	8	8	16	Dell G14	2017
Dell	PowerEdge R940 G14	128	16	8	16	Dell G14	2017
Dell	PowerEdge R940 G14	256	32	8	16	Dell G14	2017
Dell	PowerEdge R940 G14	384	48	8	16	Dell G14	2017
Dell	PowerEdge R940 G14	64	4	16	32	Dell G14	2017
Dell	PowerEdge R940 G14	128	8	16	32	Dell G14	2017
Dell	PowerEdge R940 G14	256	16	16	32	Dell G14	2017
Dell	PowerEdge R940 G14	512	32	16	32	Dell G14	2017
Dell	PowerEdge R940 G14	128	4	32	64	Dell G14	2017
Dell	PowerEdge R940 G14	256	8	32	64	Dell G14	2017
Dell	PowerEdge R940 G14	512	16	32	64	Dell G14	2017
Dell	PowerEdge R940 G14	1024	32	32	64	Dell G14	2017
Dell	PowerEdge R940 G14	256	4	64	128	Dell G14	2017
Dell	PowerEdge R940 G14	512	8	64	128	Dell G14	2017
Dell	PowerEdge R940 G14	1024	16	64	128	Dell G14	2017
Dell	PowerEdge R940 G14	512	4	128	256	Dell G14	2017
Dell	PowerEdge R940 G14	1024	8	128	256	Dell G14	2017
Dell	PowerEdge R940 G14	2048	16	128	256	Dell G14	2017
Dell	PowerEdge R940 G14	3072	24	128	256	Dell G14	2017
Dell	PowerEdge R740	32	4	8	16	Dell G14	2017
Dell	PowerEdge R740	64	8	8	16	Dell G14	2017
Dell	PowerEdge R740	128	16	8	16	Dell G14	2017
Dell	PowerEdge R740	192	24	8	16	Dell G14	2017
Dell	PowerEdge R740	64	4	16	32	Dell G14	2017
Dell	PowerEdge R740	128	8	16	32	Dell G14	2017
Dell	PowerEdge R740	256	16	16	32	Dell G14	2017
Dell	PowerEdge R740	384	24	16	32	Dell G14	2017
Dell	PowerEdge R740	128	4	32	64	Dell G14	2017
Dell	PowerEdge R740	256	8	32	64	Dell G14	2017
Dell	PowerEdge R740	512	16	32	64	Dell G14	2017

Dell	PowerEdge R740	768	24	32	64	Dell G14	2017
Dell	PowerEdge R740	256	4	64	128	Dell G14	2017
Dell	PowerEdge R740	512	8	64	128	Dell G14	2017
Dell	PowerEdge R740	1024	16	64	128	Dell G14	2017
Dell	PowerEdge R740	1536	24	64	128	Dell G14	2017
HP	ProLiant DL360 G10	32	4	8	16	HP G10	2017
HP	ProLiant DL360 G10	64	8	8	16	HP G10	2017
HP	ProLiant DL360 G10	128	16	8	16	HP G10	2017
HP	ProLiant DL360 G10	192	24	8	16	HP G10	2017
HP	ProLiant DL360 G10	64	4	16	32	HP G10	2017
HP	ProLiant DL360 G10	128	8	16	32	HP G10	2017
HP	ProLiant DL360 G10	256	16	16	32	HP G10	2017
HP	ProLiant DL360 G10	384	24	16	32	HP G10	2017
HP	ProLiant DL360 G10	128	4	32	64	HP G10	2017
HP	ProLiant DL360 G10	256	8	32	64	HP G10	2017
HP	ProLiant DL360 G10	512	16	32	64	HP G10	2017
HP	ProLiant DL360 G10	768	24	32	64	HP G10	2017
HP	ProLiant DL360 G10	256	4	64	128	HP G10	2017
HP	ProLiant DL360 G10	512	8	64	128	HP G10	2017
HP	ProLiant DL360 G10	1024	16	64	128	HP G10	2017
HP	ProLiant DL360 G10	1536	24	64	128	HP G10	2017
HP	ProLiant DL380 G10	32	4	8	16	HP G10	2017
HP	ProLiant DL380 G10	64	8	8	16	HP G10	2017
HP	ProLiant DL380 G10	128	16	8	16	HP G10	2017
HP	ProLiant DL380 G10	192	24	8	16	HP G10	2017
HP	ProLiant DL380 G10	64	4	16	32	HP G10	2017
HP	ProLiant DL380 G10	128	8	16	32	HP G10	2017
HP	ProLiant DL380 G10	256	16	16	32	HP G10	2017
HP	ProLiant DL380 G10	384	24	16	32	HP G10	2017
HP	ProLiant DL380 G10	128	4	32	64	HP G10	2017
HP	ProLiant DL380 G10	256	8	32	64	HP G10	2017
HP	ProLiant DL380 G10	512	16	32	64	HP G10	2017
HP	ProLiant DL380 G10	768	24	32	64	HP G10	2017
HP	ProLiant DL380 G10	256	4	64	128	HP G10	2017
HP	ProLiant DL380 G10	512	8	64	128	HP G10	2017
HP	ProLiant DL380 G10	1024	16	64	128	HP G10	2017
HP	ProLiant DL380 G10	1536	24	64	128	HP G10	2017
Dell	PowerEdge T630 G13	16	4	4	8	Dell G13	2015
Dell	PowerEdge T630 G13	32	8	4	8	Dell G13	2015
Dell	PowerEdge T630 G13	64	16	4	8	Dell G13	2015
Dell	PowerEdge T630 G13	96	24	4	8	Dell G13	2015

Dell	PowerEdge T630 G13	32	4	8	16	Dell G13	2015
Dell	PowerEdge T630 G13	64	8	8	16	Dell G13	2015
Dell	PowerEdge T630 G13	128	16	8	16	Dell G13	2015
Dell	PowerEdge T630 G13	192	24	8	16	Dell G13	2015
Dell	PowerEdge T630 G13	64	4	16	32	Dell G13	2015
Dell	PowerEdge T630 G13	128	8	16	32	Dell G13	2015
Dell	PowerEdge T630 G13	256	16	16	32	Dell G13	2015
Dell	PowerEdge T630 G13	384	24	16	32	Dell G13	2015
Dell	PowerEdge T630 G13	128	4	32	64	Dell G13	2015
Dell	PowerEdge T630 G13	256	8	32	64	Dell G13	2015
Dell	PowerEdge T630 G13	512	16	32	64	Dell G13	2015
Dell	PowerEdge T630 G13	768	24	32	64	Dell G13	2015
Dell	PowerEdge T630 G13	256	4	64	128	Dell G13	2015
Dell	PowerEdge T630 G13	512	8	64	128	Dell G13	2015
Dell	PowerEdge T630 G13	1024	16	64	128	Dell G13	2015
Dell	PowerEdge T630 G13	1536	24	64	128	Dell G13	2015
HP	ProLiant DL580 G9 R	32	8	4	16	HP G9	2016
HP	ProLiant DL580 G9 R	384	96	4	16	HP G9	2016
HP	ProLiant DL380 G9	8	2	4	4	HP G9	2014
HP	ProLiant DL380 G9	48	12	4	4	HP G9	2014

It's evident from the QPL composition charts below that our analysis covers most of the memory range of QPL models.

Spread of installed memory (GB) in QPL by model count