DOCKETED

Docket Number:	14-AAER-02
Project Title:	Computer, Computer Monitors, and Electronic Displays
TN #:	205736
Document Title:	California Investor Owned Utilities Comments: CA IOUs & NRDC August 2015 CASE Response to CEC proposal
Description:	N/A
Filer:	System
Organization:	California Investor Owned Utilities
Submitter Role:	Public
Submission Date:	8/14/2015 8:58:36 AM
Docketed Date:	8/14/2015

Comment Received From: California Investor Owned Utilities
Submitted On: 8/14/2015
Docket Number: 14-AAER-02

CA IOUs & NRDC August 2015 CASE Response to CEC proposal

Additional submitted attachment is included below.

Computers

Codes and Standards Enhancement (CASE) Initiative

For PY 2015: Title 20 Standards Development

Response to CEC Standards Proposal for

Computers

Docket #14-AAER-2A

August 7, 2015

Prepared for:

PACIFIC GAS & ELECTRIC COMPANY

SOUTHERN CALIFORNIA EDISON

SAN DIEGO GAS AND ELECTRIC

SOUTHERN CALIFORNIA GAS COMPANY

NATURAL RESOURCES DEFENSE COUNCIL

Prepared by:

Nate Dewart, ENERGY SOLUTIONS
Pierre Delforge, NATURAL RESOURCES DEFENSE COUNCIL
Vojin Zivojnovic, AGGIOS
Davorin Mista, AGGIOS
Peter May-Ostendorp, XERGY CONSULTING
Katherine Dayem, XERGY CONSULTING

This report was prepared by the California Statewide Utility Codes and Standards Program and funded by the California utility customers under the auspices of the California Public Utilities Commission.

Table *of* Contents

1 Executive Summary

The Pacific Gas and Electric Company (PG&E), Southern California Edison (SCE), Southern California Gas (SCG), San Diego Gas & Electric (SDG&E) Codes and Standards Enhancement (CASE) Initiative Project seek to address energy efficiency opportunities through development of new and updated Title 20 standards. Individual reports document information and data helpful to the California Energy Commission (CEC) and other stakeholders in the development of these new and updated standards. The objective of this project is to develop CASE Reports and subsequent documents that provide comprehensive technical, economic, market, and infrastructure information on each of the potential appliance standards.

This document provides recommendations and supporting analysis in response to the CEC's Computers Staff Report to ensure that California maximizes cost-effective and feasible energy savings. It also includes some research on panel self-refresh, as well as in Appendix A, the discrete graphics testing results which informed the proposal, and in Appendix B, testing results of business desktop computers with security and manageability features responding to stakeholder concerns regarding the applicability of the CEC's proposed standard to these products. Here is a summary of the CASE Team's recommendations by topic:

1. **Typical Energy Consumption (TEC) Base Allowance:**

- *Desktops:* The CASE Team supports CEC's proposed levels based on analysis of the current market, though we recommend the levels be more stringent to account for future market improvement by the effective date of 2018.
- *Notebooks:* Based on analysis of the current market, CEC's proposed levels overstate the allowance needed even for high-performing notebooks. The CASE Team proposes more stringent levels based on the current market and accounting for future improvement by the effective date of 2017.
- *Thin clients:* CEC's proposed levels overstate the allowance needed given that these products do not need an allowance for data storage. The CASE Team proposes more stringent levels based on analysis of the current market and to account for future market improvement by the effective date of 2018.

2. **Display adder:** CEC's proposed levels significantly overstate the adder needed. The CASE Team proposes more stringent levels based on current market analysis and includes a new display adder equation.
3. **Discrete graphics adder:** The CASE Team supports the proposed no-adder approach given existing technology and market trends of graphics switching between discrete and integrated graphics, although recognizes an adder is needed for products without integrated graphics and proposes a two-tier standard for this situation.
4. **Memory adder:** CEC's proposed levels overstate the adder needed for memory. The CASE Team proposes calculating this adder on a per-memory-module basis rather than the per-GB basis in the current proposal.
5. **Secondary storage adder**: CEC's proposed adder is not needed, as secondary internal storage drives need not be powered on during idle modes. The CASE Team proposes this

adder to be zero, with a clarification to the test procedure to ensure that the specifications for secondary storage are explicit and they are allowed power down during testing.

6. **Duty cycle:** The CASE Team reiterates the need to adjust cost-effectiveness calculations to account for real-world computer usage, by increasing both the estimated time in idle modes and the energy consumption when in idle. The result is an increase in both per-unit and statewide savings from the standards not currently calculated, and potentially justifies an even more stringent standard. The CASE Team also recommends use of only the conventional duty cycle for compliance.
7. **Power Supplies:** The CASE Team continues to recommend the cost-effective 80 PLUS Gold requirement for desktop internal power supplies, an efficiency requirement at 10% load, and power factor requirements at all loads.
8. **Power management:** The CASE Team recommends that computers transition to hibernate mode after 4 hours or less in sleep mode, and also a requirement for proximity sensors & auto brightness control.
9. **Definitions and Certification & Data Submittal Requirements:** The CASE Team recommended several modifications and additions to the definitions in its May 29th, 2015 submittal and is currently working on a possible joint proposal, where applicable, with Information Technology Industry Council. The CASE team also recommends certification requirements and reporting requirements for Title 20 Table X to ensure optimal compliance for the products covered by the proposal.

2 Typical Energy Consumption (TEC) Base Allowances

2.1 Desktops

The Aggios desktop optimization project, funded by Natural Resources Defense Council (NRDC) and building off the notebook research funded by the California Investor-Owned Utilities (IOUs) (2014), demonstrated at the April 15th workshop, exemplifies that mainstream desktops can be optimized to meet the proposed standard levels, with a combination of software configuration changes and cost-effective power supply replacement. The proposed desktop base allowances are appropriate for current technology, but **should be reduced** to account for expected technology improvements by effective date 2018. Computer technology is making progress on energy efficiency, with the current trend expected to continue and should be taken into account when setting standards for 2018. Given the potential software improvement opportunities and adjustments based on historical trends of annual ENERGY STAR level improvement of roughly 10% as highlighted by ITI/Technet in their April 15th, 2015 workshop presentation (2015) and an effective date of 2018, we propose a TEC base allowance of **36 kwh/yr** (50 kWh/yr with 10% improvement compounded over three years).

For further context, the proposed level is supported by a "bottom-up", quantitative approach that:

- Quantifies idle power consumption of individual components;
- Analyzes the dependencies between components (e.g. power supply unit, a.k.a. PSU, and I/O interfaces);

- Determines the sum of individual component idle power adjusted for component dependencies;

This approach is based on the principle that the power drawn by the electronic device and its components should be kept in proportion to the necessary and useful work delivered by the device. The main challenge when implementing and verifying this principle on real devices is how to determine and quantify the work delivered.

In the long and short idle states, as defined by the ESTAR test procedure, the computer is conducting the following tasks:

- Occasional background tasks such as maintenance and updates, however the processor is spending most of short and long idle time in C6/C7 low-power states. No ongoing active processing tasks;
- Other hardware keeps the internal physical connections (bus, memory) and external I/O interfaces (PCI, USB, eSATA, DVI, HDMI) adequately powered and alive;
- Software communicates with internal and external components to refresh their state, as drivers and components need regular refresh (including the display refresh);
- Devices wait for internal and external events (by interrupt or polling) to transition the computer to the active state with acceptable latencies;
- External interfaces wait to detect connections to newly added external peripherals;

Due to current imperfections in the hardware components (including the PSU), interconnections, operating system software, interface communication protocols and wake up procedures, components occasionally consume disproportionally more energy than necessary for the relatively small amount of work they typically provide in idle. Taking forthcoming technology improvements and diminishing imperfections into account, the base TEC level of 36 kWh/yr is a reasonable allowance for such computers with a typical set of (unpopulated) external interfaces.

We are currently testing high-performance configurations to determine if high-end computers with a large number of high-speed external I/O interfaces can also meet this level. Should this level prove to be limiting, we will explore the definition and computation of expandability adders. We believe that an adder based on the presence of expandability interfaces (e.g. internal expansion slots or USB ports) would be more appropriate than one based on technical specifications such as number of cores and processor frequency, for the following reasons:

- Expandability interfaces better represent the additional power needed by higher-capability computers in idle modes;
- Technical specifications evolve very rapidly due to natural technology evolution and would cause a large portion of the computers on the market to migrate to the highest category over time ("category creep").
- The number of expandability interfaces is less prone to rapid evolution due to the space constraints and cost.

In contrast, the categories in the ENERGY STAR framework are based on processor performance scores, but recent processor architectures, such as Intel's 4th generation Core architecture

(Haswell), can achieve very low idle power levels regardless of their maximum performance, due to deeper power management states, process improvements, and improved power scalability. Thus, processor p-score (the number of cores multiplied by the clock speed per core) is an insufficient proxy for a system's idle power draw and should be avoided as criteria for establishing categories or granting higher allowed TEC to systems.

2.2 Notebooks

CEC proposed limits for notebooks are far too generous and should be re-assessed. Notebook products now represent a majority of the mainstream computer market and should be scrutinized carefully for efficiency improvement potential.

As reported in its letter to Information Technology Industry Council (CEC 2015a), the CEC's analysis for determining standards was based on the highest performing products from the ENERGY STAR 5.2 Qualified Products List, which dates back to 2008-2012, and is now an outdated dataset. For example, the CLASP (2014) research of 2013 data showed that roughly 57% of notebooks on the market would meet ENERGY STAR Version 6 levels, and additional IOU research conducted in 2014 on available products determined that the market for ENERGY STAR Version 6 notebooks is likely already quite high, as purchasing a non-ENERGY STAR qualifying product from leading OEMs proved challenging. With further analysis of the May 2015 ENERGY STAR 6.1 QPL (2015) including subtraction of the estimated display and memory energy consumption from these products, the CASE Team suggests that CEC should reconsider its analysis and modify the proposed allowances. See Table 2.1 below for details. Roughly 75% of all notebooks made available in 2014 would meet the CEC proposed base allowance. **We instead recommend a base allowance of 19 kWh/yr.** This level corresponds to the median over the last 10 months (July 2014-April 2015) for high-performance (I3 category) units in the ENERGY STAR QPL (2015) discounted by 10 percent twice to account for the annual natural TEC reduction trend by 2017.

Table 2.1: ENERGY STAR 6.1 QPL Analysis - Typical Energy Consumption for Notebooks (kWh/yr)

	Category 0	**Category I1**	**Category I2**	**Category I3**	**Overall**
	# of Entries	**# of Entries**	**# of Entries**	**# of Entries**	**# of Entries**
(All)	111	1,679	727	714	3,280
(2014-1st Half)	74	1,480	643	647	2,844
(2014-2nd Half)	5	122	36	27	190
(2015-1st Half)	32	77	48	40	197
	Top 75% (TEC)	**Top 75% (TEC)**	**Top 75% (TEC)**	**Top 75% (TEC)**	**Top 75% (TEC)**
(All)	12.4	27.5	26.8	35.7	29.9
(2014-1st Half)	12.0	27.9	28.9	36.5	32.2
(2014-2nd Half)	12.4	7.6	24.6	27.5	25.6
(2015-1st Half)	14.0	N/A	18.3	35.9	24.5
	Top 50% (TEC)	**Top 50% (TEC)**	**Top 50% (TEC)**	**Top 50% (TEC)**	**Top 50% (TEC)**
(All)	10.9	11.0	19.2	27.7	22.7
(2014-1st Half)	8.7	21.0	21.4	30.6	25.4
(2014-2nd Half)	10.9	7.0	18.2	23.0	19.7
(2015-1st Half)	14.0	N/A	13.1	24.5	18.1
	Top 25% (TEC)	**Top 25% (TEC)**	**Top 25% (TEC)**	**Top 25% (TEC)**	**Top 25% (TEC)**
(All)	3.6	7.0	13.9	22.2	16.4
(2014-1st Half)	4.2	11.4	15.9	23.6	18.5
(2014-2nd Half)	3.6	6.5	13.8	19.3	14.4
(2015-1st Half)	14.0	N/A	10.3	20.2	11.5
	Top 10% (TEC)	**Top 10% (TEC)**	**Top 10% (TEC)**	**Top 10% (TEC)**	**Top 10% (TEC)**
(All)	2.2	5.5	8.9	17.8	10.8
(2014-1st Half)	2.2	6.8	9.7	20.3	12.2
(2014-2nd Half)	2.7	5.8	9.0	16.5	11.0
(2015-1st Half)	14.0	N/A	8.5	13.6	9.2

2.3 Thin clients

By definition thin clients are computers with lower capabilities than desktop computers. For example, they typically have no rotational storage media (hard disk, optical disk), as highlighted in more detail in the thin client definitions, Section 8.5.2. As such they should be able to meet lower limits than desktop computers. ENERGY STAR v6.1 sets different limits for thin clients and desktops. **We propose that a specific thin client limit be set at 20 kWh/yr**: the desktop level minus the roughly 16 kWh/yr of TEC required by 3.5" magnetic hard drives in desktop computers.[1] This reflects the fact that thin clients do not have a Hard Disk Drive (HDD) and therefore do not need to include disk power in the idle levels.

[1] Based on an analysis of hard drive idle power consumption data from Tom's Hardware, available at: http://www.tomshardware.com/charts/hdd-charts-2013/-26-Power-Requirement-at-Idle,2917.html

3 Adders

While we support the principle of functional allowances (a.k.a. adders) to cover features and performance capabilities that cannot be powered down in idle mode, like the display in short idle, adder amounts should be set to what is necessary using best-practice cost-effective technologies, and should only apply to features or performance capabilities that cannot be power down in idle mode. CEC's proposal aligns with ENERGY STAR v6, providing four adders —displays, discrete graphics cards, memory and storage—and we recommend adjustments to all four.

3.1 Display Adders

The display adders proposed by CEC based on ENERGY STAR v6 are far higher than required by current display technology and would result in ineffective standards for integrated desktop and notebook computers. We propose revised display adders based on the real power needs of current display technology per the ENERGY STAR v6.1 QPL.

Display adders are necessary to account for the energy used by the display of integrated desktops and notebooks in short idle mode. However, the ENERGY STAR v6 display adders used in CEC's proposal are far higher than the difference between short and long idle in the QPL. That difference is a conservative proxy for display power since it can include the power of other components than the display.

Figure 3.1 shows that CEC proposed adders, based on ENERGY STAR v6, are on average 23% higher than the difference between short and long idle in the QPL for integrated desktops, and 59% higher for notebooks.

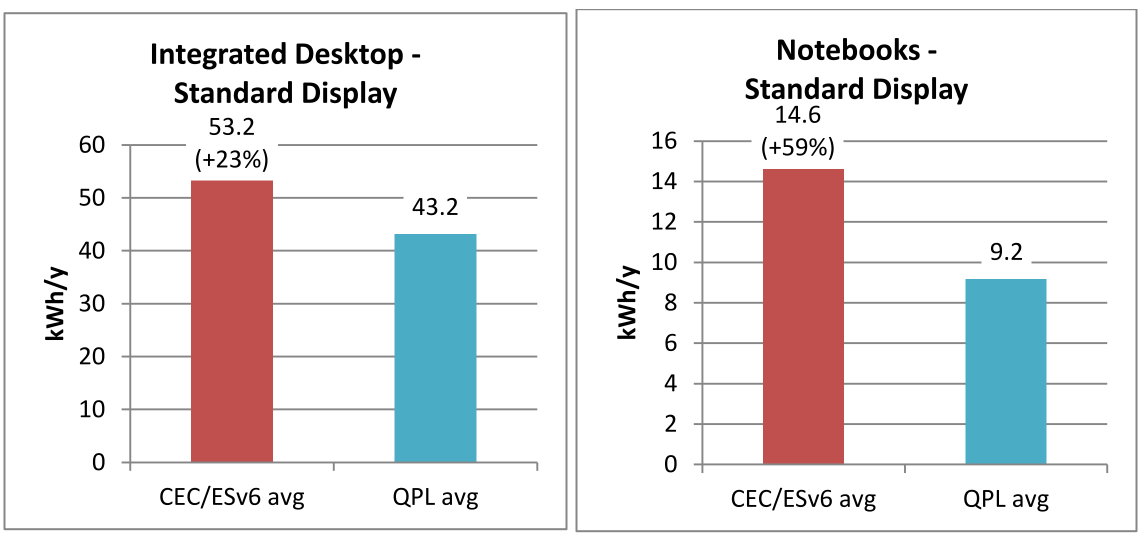

Figure 3.1 Average Display Allowances Across QPL for May 2014 – April 2015 Models – Standard Display

These differences have a huge impact on the effectiveness of the standards, because display allowances are of the same order of magnitude as the base allowance for the system: 53 vs. 50 kWh for integrated desktops, 15 vs. 30 for notebooks. In a TEC approach where excessive adders can give inefficient systems a free pass to comply, it is critical to pay attention to adders like displays so that the overall cost-effective and feasible energy savings are optimized.

3.1.1 Proposal for Display Adders

We propose revised display adders that are closer to the real needs of recent computers in the ENERGY STAR QPL. Our proposal is approximately 15% lower than the average difference between short and long idle power in the QPL for computers registered with ENERGY STAR from May 2014 to April 2015 (in order to represent the latest technology, which will be mainstream by 2017 / 2018). We have chosen a value that is lower than the average, because the difference between short and long idle can include more than just display power. For example, it could also include power savings from placing the disk and other components in lower power modes in long, and to some extent, short idle.

Our proposal uses a hyperbolic tangent equation, similar to that proposed in ENERGY STAR display spec Version 7 draft 2:

Computer type	Display allowance (kWh/year)
Integrated desktops	8.76 x 0.35 x (1+EP) x (0.5 x r + 16 x Tanh (0.004 x (A-55) + 0.25) + 0.3)
Notebooks	8.76 x 0.3 x (0.25 x r + 8.5 x Tanh (0.003 x (A-70) + 0.22) + 0.3)

Where:

- r = Screen resolution in megapixels
- A = Viewable screen area in in^2
- EP = 0.15 for enhanced performance displays of any size

Figure 3.2 shows how these proposed levels compare with ENERGY STAR v6.1 and with the QPL on average.

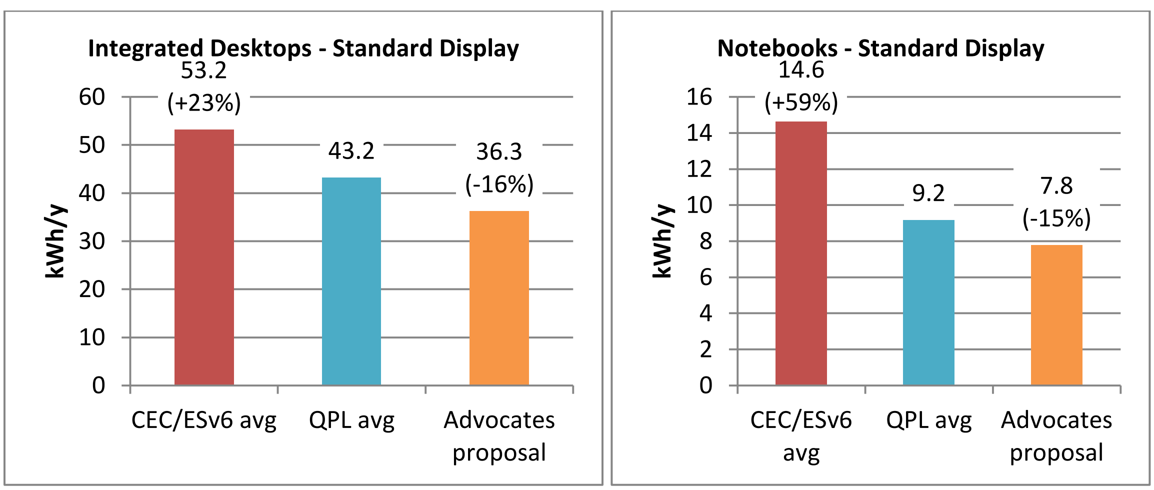

Figure 3.2: Comparison of Advocate Display Adder Proposal vs. CEC/ENERGY STAR v6.1 and QPL (May 14-April 15) – Standard Display

Our proposed levels yield "pass rates" of 21% of the QPL for integrated desktops (54% for Enhanced Performance Displays, or EPDs), and 29% for notebooks (44% for EPDs) as shown in Table 3.1. The substantial number of units achieving these levels demonstrates their technical feasibility and broad availability in the market today. Note that these are not real pass-rates; we're using conservative estimates of display power.

Table 3.1: Number of QPL models with Short-Long idle delta lower than proposed display adder, percentage of products of total that meet the proposed levels.

	Standard Display	EPD
Integrated Desktops	40 (21%)	11 (54%)
Notebooks	163 (29%)	9 (44%)

The proposed equations have been developed to allow units to meet the proposed levels across the spectrum of screen size and resolutions as shown by Figures 3.3 through 3.6. These figures compare the ENERGY STAR Version 6.1 display allowances and CASE Team proposed levels to the measured values of the difference between short and long idle, across screen area and resolution.

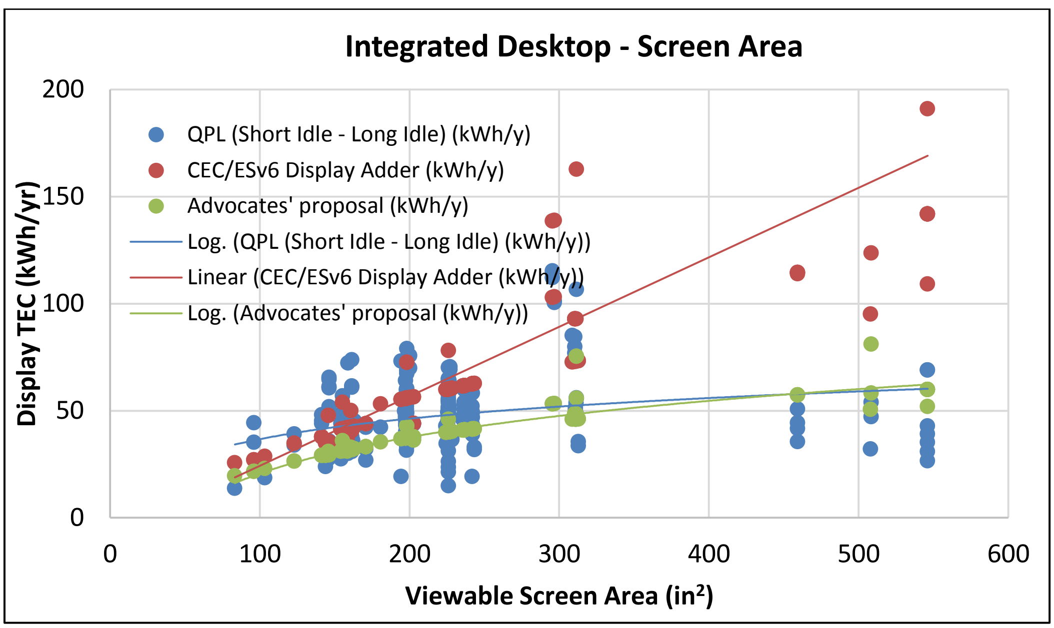

Figure 3.3: Comparison of ENERGY STAR Version 6.1 Display Allowances and QPL-Reported Short-Long Idle Difference across Screen Area.

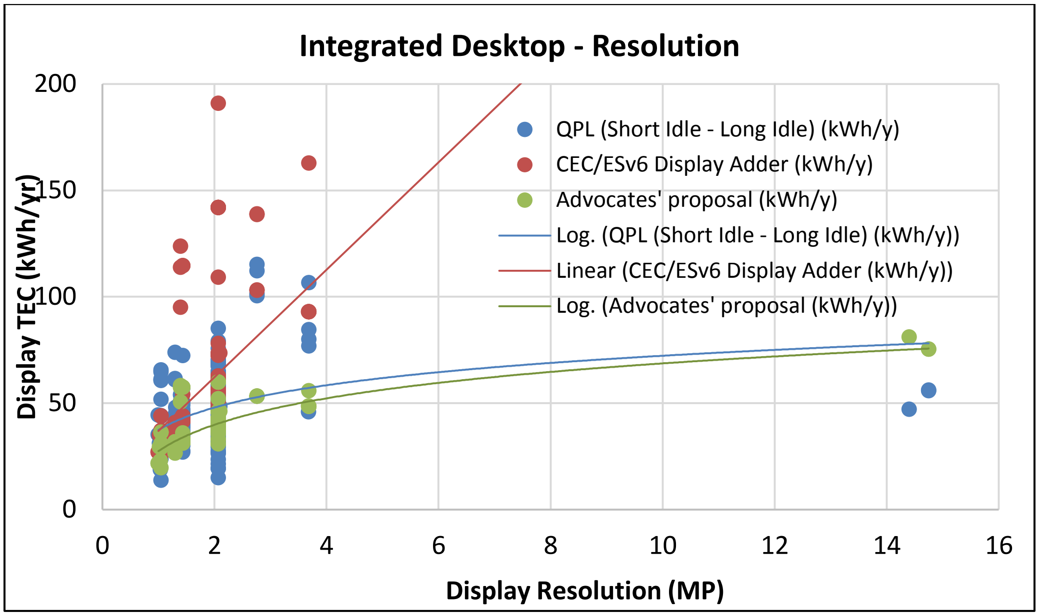

Figure 3.4: Comparison of ENERGY STAR Version 6.1 Display Allowances and QPL-Reported Short-Long Idle Difference across Resolution.

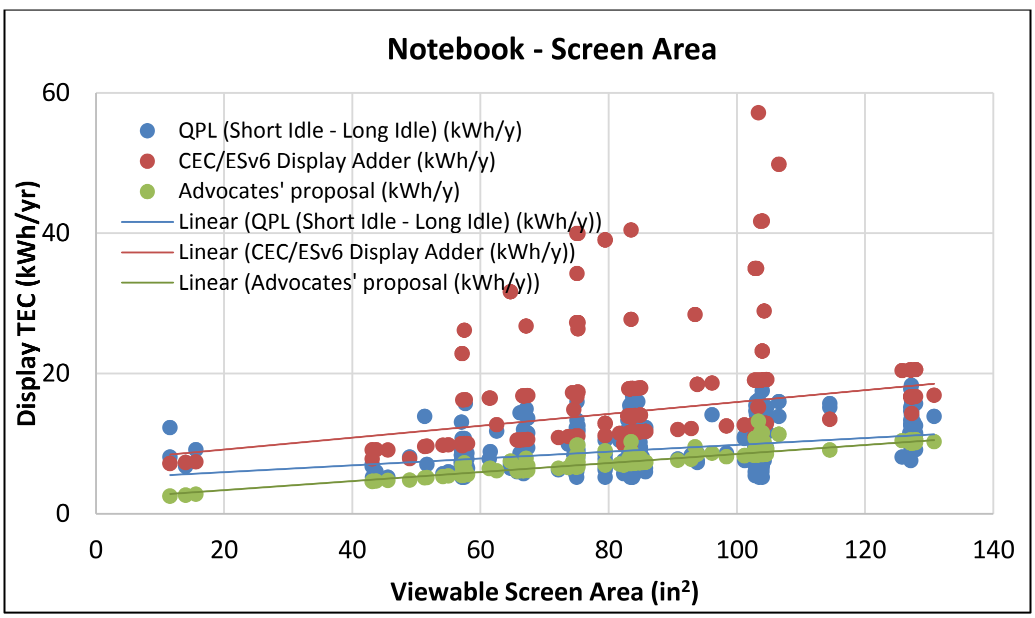

Figure 3.5: Comparison of ENERGY STAR Version 6.1 Display Allowances and QPL-Reported Short-Long Idle Difference across Screen Area.

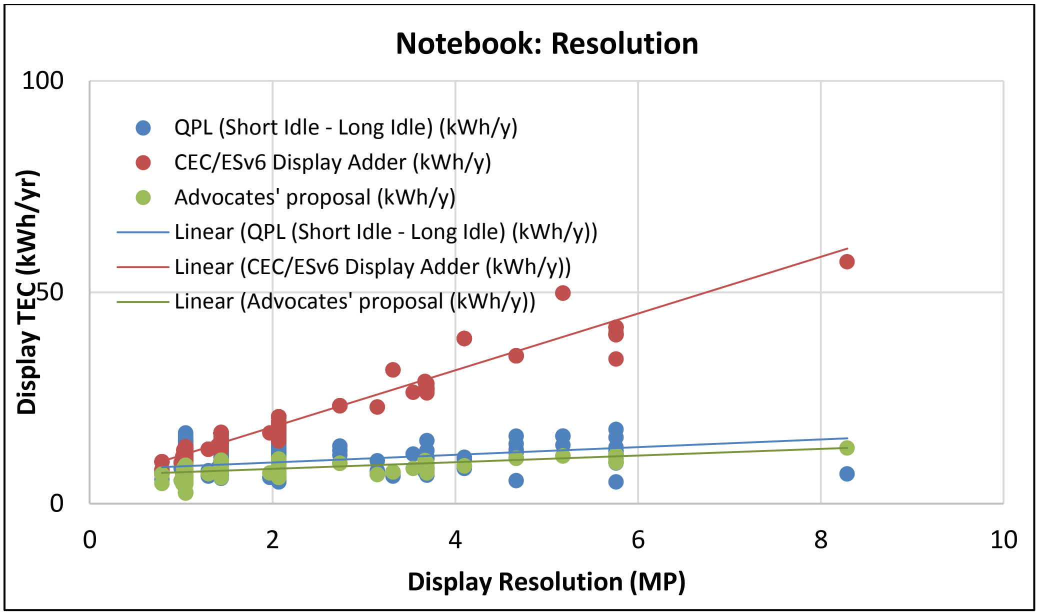

Figure 3.6: Comparison of ENERGY STAR Version 6.1 Display Allowances and QPL-Reported Short-Long Idle Difference across Resolution.

Enhanced Performance Displays (EPDs): Our proposal does not include a scaling factor for EPDs for notebooks because the QPL shows that the resolution factor in our proposed equations gives EPDs a sufficient allowance as it is.

3.2 Discrete Graphics

The CASE team generally supports the March 2015 staff proposal in establishing no functional adders for computer systems containing discrete graphics processing units (dGfx). Discrete graphics continue to make efficiency gains, integrated graphics serve mainstream graphics needs very efficiently, and for users requiring higher performance, switchable graphics are able to eliminate the energy impacts of discrete graphics cards altogether during short and long idle, all of which obviates the need for an adder in most systems. Only a small portion of desktop systems may require discrete graphics adders in cases where platforms do not support integrated graphics.

Below, we provide relevant trends and data that support CEC's proposal as well as considerations for modifications to the proposal to more comprehensively treat the diverse graphics configurations available on the market. Our comments here are specifically directed at desktop computers, although discrete graphics adders for notebooks and integrated desktops are also unnecessary for similar reasons.

3.2.1 Continued Discrete Graphics Processing Unit (GPU) Efficiency Improvements

GPU manufacturers have dramatically improved the efficiency of their products in the past half-decade. Until 2011, many discrete GPU products consumed as much energy as a mainstream desktop computer itself. Today, thanks to process improvements and enhanced power management, even the highest specification cards have power demands less than 10 W (DC) during idle modes (Tom's Hardware, 2015). Since 2011, the CASE team, along with CLASP and NRDC, has been developing a dataset of discrete GPU power measurements (37 dGfx in all) and have observed a stepped decrease in energy consumption. Figure 3.7 illustrates the incremental TEC impacts of a broad sampling of dGfx, displayed as a function of Frame Buffer Bandwidth (FBB) in GB/s (a proxy for overall card performance and graphics throughput). The chart also displays the FBB limits of each ECMA performance category. In 2011, a typical G4 dGfx contributed an additional 102 kWh/yr to a desktop's TEC, but now only adds 35 kWh/yr, a 66% decline.

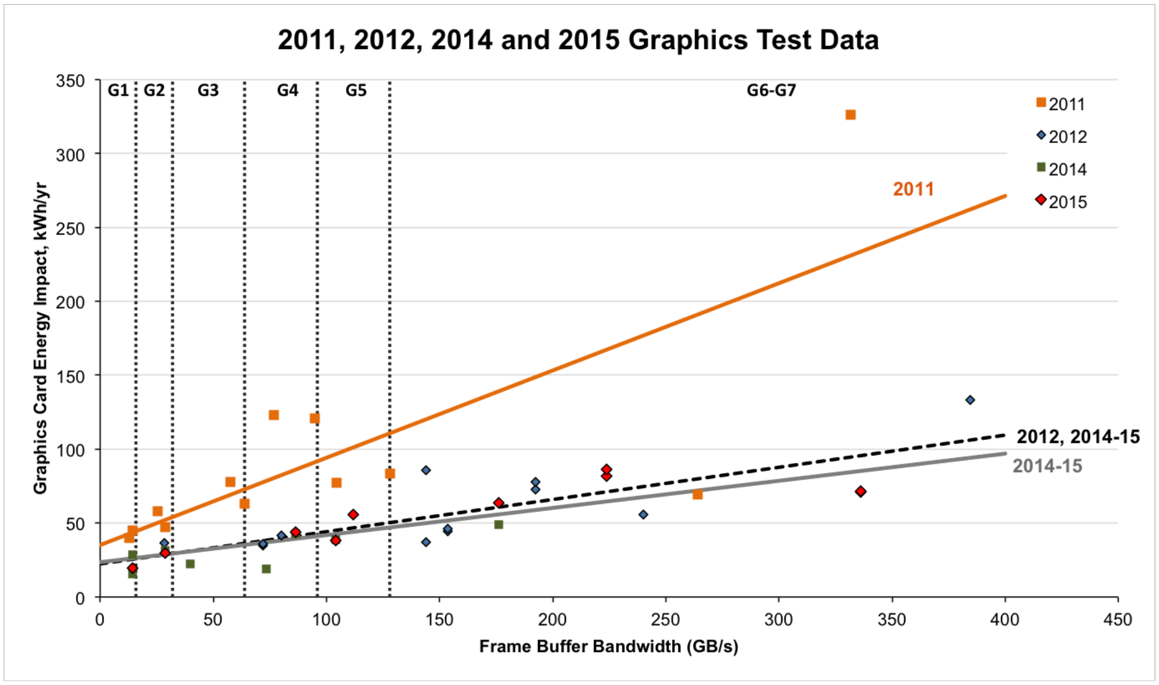

Figure 3.7: IOU, NRDC and CLASP measurements of dGfx, 2011 – 2015, with linear trends. Data collected in 2014 and 2015 were combined for the trend shown in gray.

3.2.2 Integrated DGfx Adequately Address Mainstream Graphics

As currently implemented, even today's most efficient dGfxs require some additional energy consumption in idle. New technologies and trends, however, suggest that in the vast majority of cases, future discrete GPUs will not warrant a TEC allowance. For example, both Intel and AMD have moved to integrate relatively high-performance GPUs into their processors. Intel's Core processors have included integrated GPUs since 2010, and AMD's "Accelerated Processing Units" (APUs) have existed since 2012. Both organizations tout the performance and capability of integrated graphics products to address the graphics needs of mainstream customers, even for certain gaming applications.[2]

An analysis of the advertised performance characteristics of current integrated GPUs suggests that these products are currently capable of addressing graphics needs in ECMA categories up to G4. Most products are currently equivalent to G2 discrete GPUs based on their supported memory

[2] See descriptions from Intel (http://www.intel.com/content/www/us/en/architecture-and-technology/hd-graphics/hd-graphics-video.html) and AMD (http://www.amd.com/en-us/innovations/software-technologies/technologies-gaming/apu).

bandwidth. Several products in AMD's A-series APU line are equivalent to G3. The memory bandwidth of current integrated GPUs is limited by system memory, which is shared between the processor and GPU; however, recent product announcements from Intel state that the company's next generation of integrated GPUs—namely Iris Pro 5200 and 6200—will include an additional 128 MB of dedicated onboard graphics memory, bringing FBB values well into the G4 range.[3] Benchmark tests conducted by Tom's Hardware show that the Iris Pro 6200 is comparable in performance to an AMD Radeon R7 250X (FBB 72 GB/s, G4) or an NVIDIA GeForce GTX 560 (FBB 128 Gb/s, G5).[4] Intel's HD 4000 integrated GPU is comparable to the NVIDIA 6800 Ultra (FBB 35 GB/s, G3) and the NVIDIA 9500 GT DDR2 (FBB = 16 GB/s, G1/G2).[5]

Integrated GPUs already appear in the vast majority of desktops and can perform at levels equivalent to low-end G1 through G3 discrete cards. G1-G3 discrete graphics provide no performance benefit compared to modern integrated graphics, they therefore warrant no additional energy allowance.

3.2.3 Switchable/Hybrid Graphics Key to Mitigating Discrete GPU Impact

Manufacturers may need to incorporate high-end discrete graphics for enhanced performance in advanced business PCs or higher end home PCs focused on gaming and media. In such systems, graphics adders are still unnecessary as long as the system contains an integrated GPU, as the vast majority of systems do. Graphics switching—sometimes referred to as hybrid or dual graphics—can then be used to place the discrete graphics card in a low-power mode in short and long idle mode when there are no applications loaded and no windows open, and the system is therefore not handling advanced rendering or video processing workloads.

Graphics switching is already widely available for mobile dGfx products and some desktop products as well. NVIDIA's Optimus technology allows mobile dGfxs to be powered down under certain scenarios (i.e. during idle modes or while performing less graphics-intensive tasks such as browsing the Web). Similarly, AMD's "Radeon Dual Graphics" technology provides switchable graphics capability for several existing pairings of APUs and discrete GPUs on both mobile and desktop platforms. IOU measurements confirm that AMD's technology can dramatically lower the impact of discrete GPUs when enabled (see Figure 3.8), although the configuration tested was not able to completely power down the discrete GPU.

[3] Approximately 75 GB/s using system memory and dedicated graphics memory.

[4] Tom's Hardware (http://www.tomshardware.com/reviews/intel-core-i7-5775c-i5-5675c-broadwell,4169-6.html)

[5] Tom's Hardware (http://www.tomshardware.com/reviews/gaming-graphics-card-review,3107-7.html)

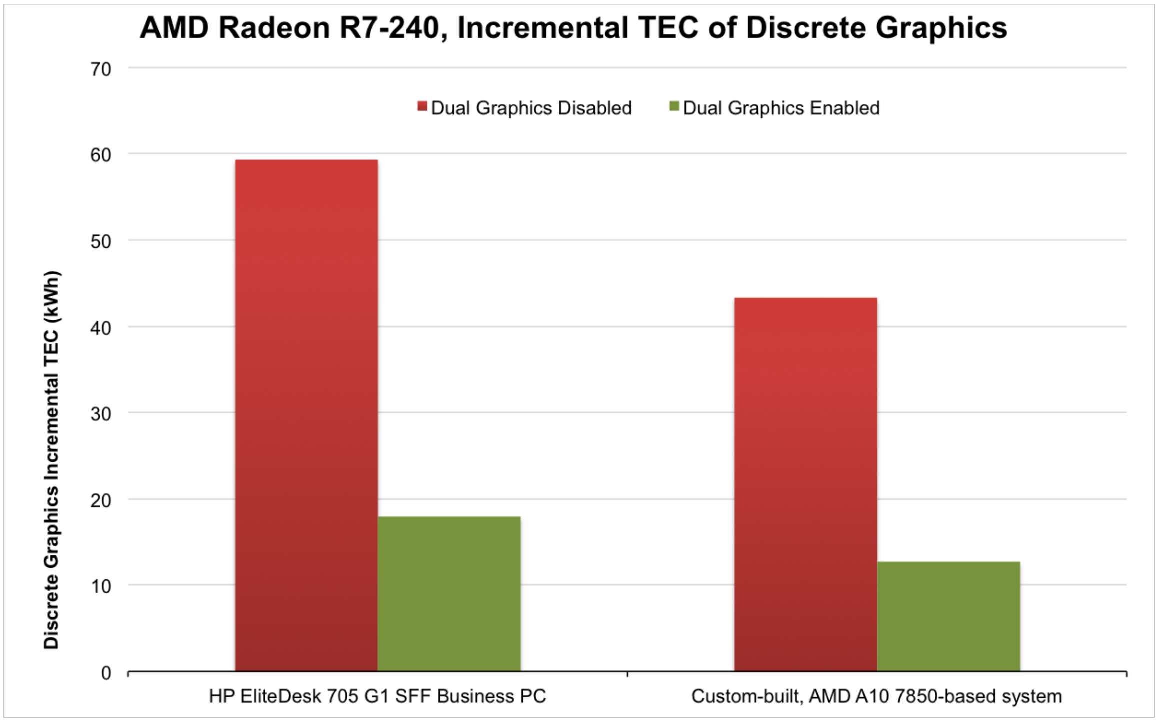

Figure 3.8: Measurements of Radeon Dual Graphics incremental TEC impacts compared to baseline systems without discrete graphics

Third parties have also begun developing vendor-agnostic desktop graphics switching solutions that can be used on a wide range of GPUs. Both Fujitsu and Asus have recently partnered with software manufacturer LucidLogix to incorporate Lucid's "VirtuWatt" graphics switching technology into certain desktop products. Fujitsu's implementation is branded as "Virtu Green" and is available in select CELSIUS workstations and ESPRIMO desktops.[6] Asus has more recently incorporated VirtuWatt on gaming PCs to enable an "Eco Energy Mode" with low idle power.[7] VirtuWatt is a combination software/hardware solution. Software is used to "virtualize" GPUs, route requests to the appropriate graphics component, and channel output signals to a single port (usually the motherboard's display port). A power microcontroller is incorporated into the motherboard or the card itself to power manage the discrete GPU (see Figure 3.9 below).

[6] A description of Fujitsu's collaboration with LucidLogix:http://www.businesswire.com/news/home/20120809005092/en/Lucid-Virtu-Enables-'0-Watt'-Graphics-Functionality-Selected#.VWTimWBJZTY.

[7] A description of Asus' implementation of VirtuWatt technology: http://www.techpowerup.com/201666/lucid-and-asus-announce-desktop-pc-power-optimization.html.

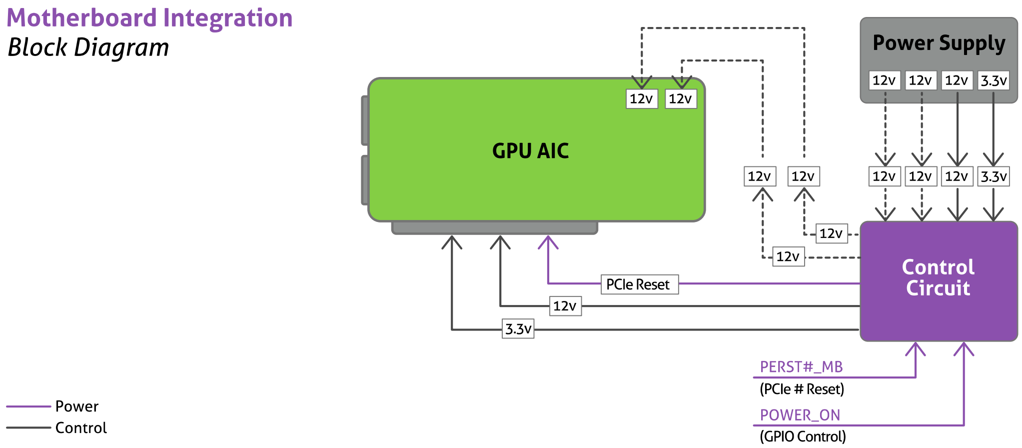

Figure 3.9: VirtuWatt motherboard integration. Discrete graphics card (green) is power managed by control circuit (purple). Source: LucidLogix, 2015

Regarding the costs of graphics switching, we know that many current NVIDIA and AMD products already support this technology (particularly in the mobile space) and that desktop OEMs are voluntarily implementing third-party solutions from vendors like LucidLogix that involve software changes and the addition of a simple and inexpensive microcontroller to the motherboard.[8] At the April 15, 2015 staff workshop on CEC's draft staff proposal, Dell Inc. indicated that the addition of graphics switching capabilities would incur a $5 incremental cost at retail. This figure establishes a conservative basis for cost effectiveness. At a $5 incremental cost, a $0.16/kWh cost of electricity, and a 5-year product lifetime, graphics switching would need to save only 6.25 kWh/yr for consumers to incur no net costs over the product life. As Figure 3.7 shows, these savings are readily achievable and economically justifiable for any discrete GPU in our dataset. Even the lowest power discrete GPUs in our dataset would generate highly cost-effective savings, with benefit-cost ratios in excess of 2.5. The benefit-cost ratio for certain higher powered G7 GPUs could exceed 13.

Although graphics switching is prevalent in notebooks and integrated desktops, challenges related to implementation in desktops may be preventing broad penetration. In particular, users may expect to connect multiple displays to a desktop, and often these displays are connected to the discrete GPU (although many integrated GPUs can support multiple physical monitor connections as long as the ports are integrated into the motherboard). Even in a single-display situation, users may opt to connect their monitor to the discrete GPU's adapters. In a graphics switching solution, however, information for the display(s) is routed through the integrated adapter on the motherboard. If the user does not connect display(s) to the motherboard, power management of the discrete GPU and associated power savings will not be achieved. The CASE team recognizes this challenge, and suggests that clarifications in the test procedure (specifically, stating that if graphics switching is enabled as-shipped, the display shall be connected to the motherboard during testing) can allow manufacturers to take credit for the energy savings achievable through graphics

[8] Power microcontrollers, such as those used in battery charge management, can be purchased at volume for less than $1, according to data available on electronics parts supplier DigiKey.com. Even at large retail markups (e.g. 300%), this translates into incremental retail costs of less than $5.

switching. Analogous to the situation of power management, the consumer may choose to disable graphics switching by connecting displays through the discrete GPU, but energy savings tips and smart default settings can guide the majority of users to realize savings by connecting displays using the motherboard adapter.

CEC should maintain its current no-adders position for dGfx and encourage further adoption of this sensible and highly cost-effective technology.

3.2.4 Considerations for Discrete Graphics Only Systems

The CASE Team supports the staff proposal of no additional graphics adders for systems that contain an integrated GPU. We recommend that the CEC clarify that integrated graphics are either part of the system's main processor or another motherboard-mounted component.

A limited number of desktop platforms today rely solely on discrete graphics. Intel's Haswell Extreme processors are one such example. Although uncommon in mainstream desktops, such processors can be found in OEM high-end gaming PCs like the Alienware Area 51. As these systems rely solely on high-performance discrete GPUs to drive any connected displays, IOU research suggests that some additional TEC allowance is warranted.

The CASE team measured current generation discrete GPUs and present test methods and results in Appendix A. The team used the results of this testing and testing carried out in 2014 to develop a two-tier adder for discrete-only GPU systems. The adder has the mathematical form:

$$TEC_{dGPU} = X*\tanh(Y*(FBB-Z)+B)+C$$

where FBB is the frame buffer bandwidth in GB/s and X, Y, Z, B, and C are parameters used to adjust curve characteristics. The slope of the curve progressively decreases with increasing FBB, and asymptotically approaches a slope of zero at high FBB. The CASE team selected this form for the adder to place an upper bound on its impacts, while still allowing continued innovation and development of higher performance cards. The proposed curve has the same mathematical form as that used by ENERGY STAR in their v6 television and draft v7 displays specifications.

The proposed Tier I adder includes many products available on the market today across every performance category:

$$TEC_{dGPU} = 65*\tanh(0.004*FBB+0.05)+15$$

[Note that here Z = 0] About 44% of the discrete GPUs tested by the IOUs in 2012, 2014, and 2015 pass the proposed Tier I adder (Figure 3.10).

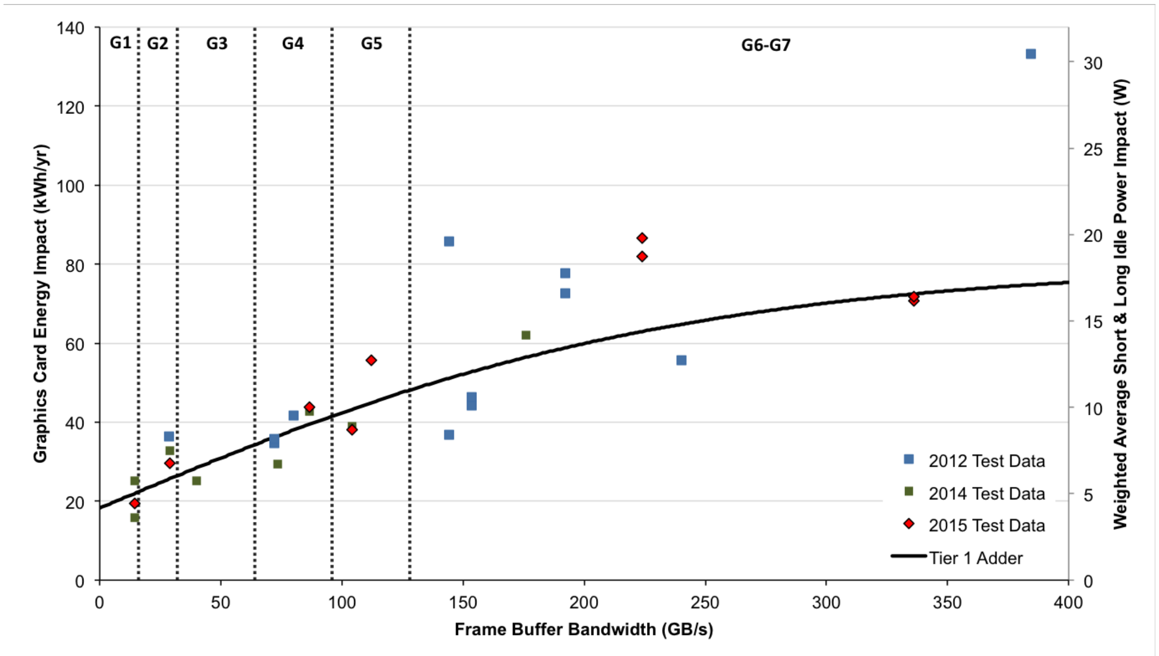

Figure 3.10: Proposed Tier I adder for discrete graphics-only systems.

The Tier II proposal is based on two technologies that have already been implemented in a limited number of discrete GPUs and systems, but will likely be mainstream by the effective date of Tier II in 2019 (assumed to be one year after the effective date of Tier I). The first is reducing discrete GPU power in long idle mode when the monitor is blank, as implemented in AMD's ZeroCore Technology. AMD claims that this type of power management can reduce long idle power of the GPU by at least 95%[9]. We generated the proposed Tier II curve by applying a 95% drop in long idle power on the 2014 and 2015 data, and constructing a curve that allows GPUs to pass across performance categories (Figure 3.11):

$$TEC_{dGPU} = 35*\tanh(0.005*(FBB\text{-}64)+0.05)+25$$

[9] A graphical description of AMD ZeroCore Technology: http://www.amd.com/PublishingImages/graphics/tables/hi-res/power-efficient-gpu-amd-zerocore-large.png, and AMD whitepaper including a section on ZeroCore: https://www.amd.com/Documents/amd_powertune_whitepaper.pdf

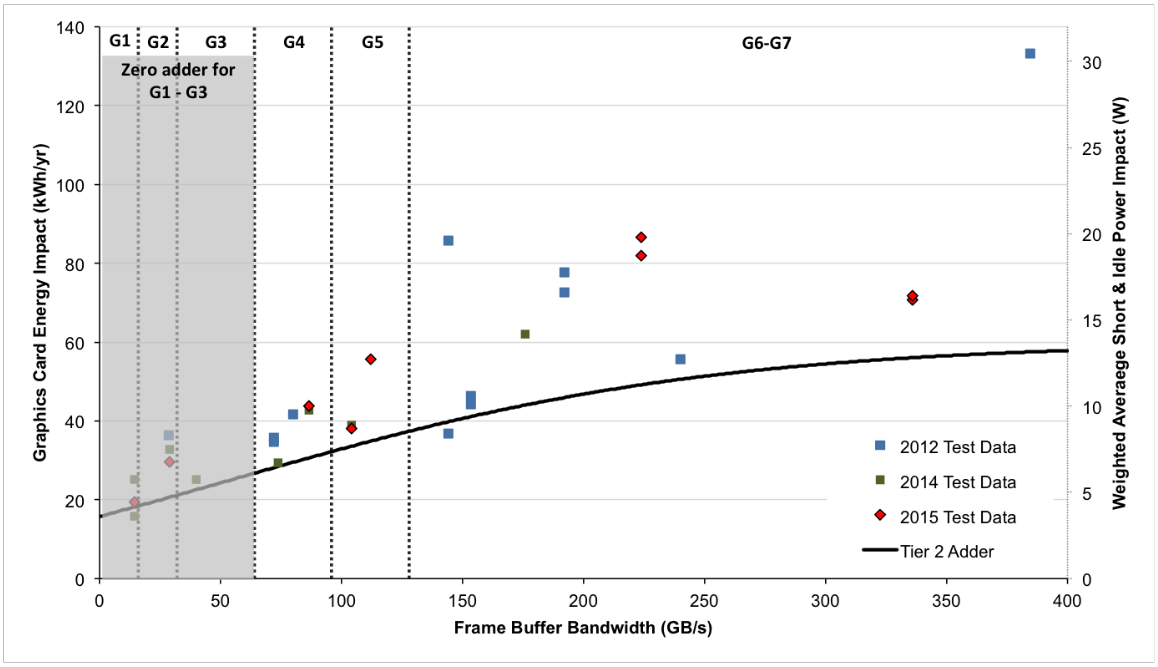

Figure 3.11: Proposed Tier II adder for discrete graphics-only systems.

Currently, only 2 of our measured GPUs would pass the proposed Tier 2 adder. However, if the discrete GPUs tested in the 2012, 2014, and 2015 projects achieved the 95% lower long idle power, 24% of them would pass the proposed level.

3.3 Memory

In its March 2015 proposal, CEC included a functional adder for memory of 0.8 kWh/yr per GB of installed physical memory based on the ENERGY STAR v6.1 computer specification. This adder is outdated, is inconsistent with the driving factors behind memory energy use, and most importantly, grants overly conservative and unwarranted allowances for memory.

ENERGY STAR developed its original adder using data gathered during the 2010-2012 timeframe, but the memory landscape has since evolved. Today's computers use DDR3 DRAM, whereas ENERGY STAR's dataset was gathered at a time when less efficient DDR2 technology was still available. Furthermore, computer memory is expected to begin transitioning to newer, more efficient DDR4 technology, which draws less power when in an active mode and has improved support for power management when in lower power modes.

Even compared to today's designs using DDR3 memory, the staff proposal is too conservative for establishing an energy efficiency standard. The CASE Team has identified computer models on the ENERGY STAR QPL with identical model numbers and equivalent (or occasionally higher performance) processors, but differing amounts of physical memory, allowing for a comparison of real memory TEC impacts. As Figure 3.12 shows, the total amount of installed physical memory is a relatively poor indicator of energy use and, therefore, should not be used as a basis for granting functional allowances.[10] Furthermore, the proposed 0.8 kWh/yr per GB value is overly conservative and can grant over 50 kWh/yr of additional TEC to high-end systems (twice that amount as 16 GB DIMMs become available, enabling 128 GB memory installs in high-performance systems).

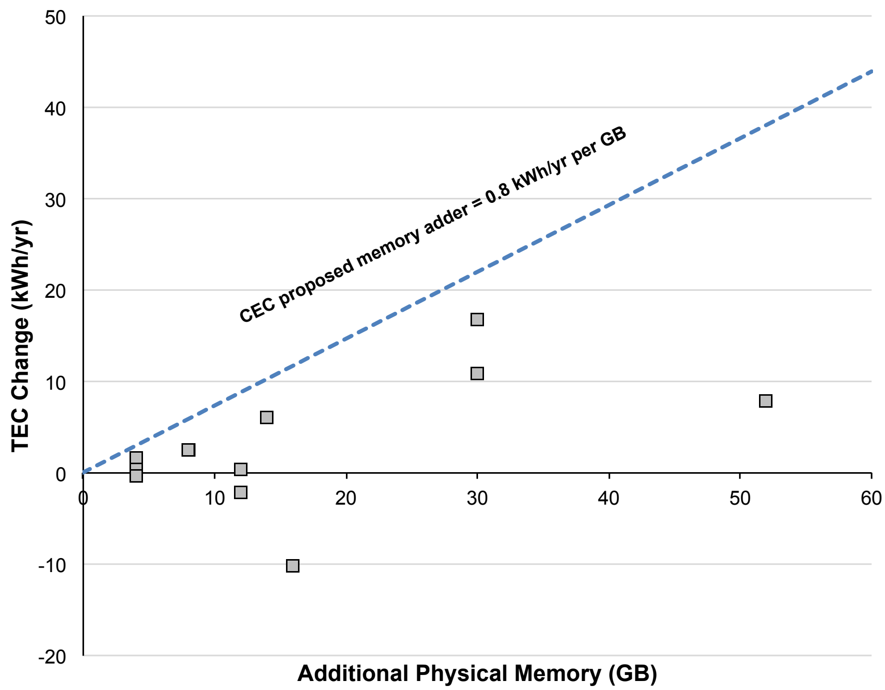

Figure 3.12: Incremental TEC consumed by desktop computers with additional installed memory. Source: ENERGY STAR qualifying products list, 2015

Recent testing conducted by Aggios suggests that the incremental power consumed by memory tends to scale more closely with the number of populated dual inline memory module (DIMM) slots on the motherboard than with overall capacity. Aggios' testing also shows that memory typically consumes about 1.3 to 2.2 kWh/yr per installed DIMM, based on today's DDR3 technology.

The CASE Team proposes an alternative memory adder informed by Aggios' testing and our current understanding of memory energy usage. Our proposal grants no adder for the first DIMM installed in the system (its energy use is covered by the base TEC allowance), but provides a 2 kWh/yr adder for each additional installed DIMM. Mathematically expressed, this is

$$\mathrm{TEC}_{\mathrm{memory}} = 2 * (n_{\mathrm{DIMM}} - 1)\ \ \mathrm{kWh/yr},$$

where n_{DIMM} is the total number of installed DIMMs. In systems where physical memory is, for some reason, completely integrated onto the motherboard or otherwise not user-serviceable, this adder would not apply.

[10] In two isolated instances, manufacturers reported *lower* measured TEC despite a higher memory configuration. We acknowledge that such results are counterintuitive but include the data to illustrate our point that installed memory capacity does not necessarily correlate to energy use.

3.4 Secondary Storage

In its March 2015 Staff Report, CEC included a 26 kWh/yr allowance for secondary storage drives. The CASE Team sees this allowance unnecessary. Primary storage drives need to remain active during idle mode to provide the operating system with quick access to critical files; however, secondary hard drives—which mainly exist to provide storage for extremely large files, media, or backups—can generally be spun down under short and long idle conditions through power management. Short and long idle modes, as defined by the ENERGY STAR v6.1 test procedure and IEC 62323 standard, have no applications loaded and no windows open, and therefore do not require access to secondary storage.

Operating systems used in some of today's computers already support separate power management of secondary drives. For over a decade, Mac OS-based systems have come equipped with a low-level, command-line utility to manage secondary drives.[11] A similar tool is available for Linux systems as a third-party add-on.[12] The various versions of the Windows operating system currently do not natively support separate power management for secondary drives; however, the basic power management infrastructure is already in place and has enabled the development of a variety of third-party utilities for secondary drive power management on Windows.[13] Implementing power management of secondary drives is, thus, a matter of upgrading existing software capabilities to more easily expose power management settings for individual drives and ensure that drives are appropriately power-managed by default.

The CASE Team recommends that the CEC grant no additional power allowances for secondary drives, encouraging manufacturers to make power management for these drives standard on all systems instead. This will involve implementing software changes, either by the computer manufacturers themselves or by operating system developers. In addition, we recommend that CEC clearly define secondary storage to align with ENERGY STAR's definition of "additional internal storage," specifically: "any and all internal hard disk drives (HDD) or solid state drives (SSD) shipping with a computer beyond the first. This definition does not include external drives."

In order to ensure that manufacturers are allowed to count the impacts of secondary drive power management when qualifying systems, we also recommend that CEC clarify language in the referenced ENERGY STAR v6.1 test procedure and IEC 62323 standard. The ENERGY STAR v6.1 test procedure refers to IEC 62623, Ed.1.0, 2012-10 for details concerning testing of short and long idle power demands. The IEC standard states that during short idle power mode testing, "long idle power management features should not have engaged (for example, HDD (if available) is spinning and the EUT is prevented from entering sleep mode)". The IEC standard does not describe how secondary storage devices should be configured during short idle power mode testing under ENERGY STAR v6.1. It does, however, address secondary storage devices in ENERGY STAR v5.2-compliant *long* idle measurement procedures, in which it allows as-shipped power management on secondary drives to be enabled during long idle testing[14]. The CASE team

[11] "pmset", short for "power management settings", has been a part of Mac OS X since 2002.

[12] "hdparm" is an open-source utility for Linux hard drive power management.

[13] Third-party Windows utilities include revoSleep, Drive Power Manager, and Hard Disk Sentinel.

[14] The IEC standard states for "ENERGY STAR v5.2-compliant long idle" that "If more than one internal hard drive is installed as shipped, the non-primary, internal hard drive(s) may be tested with hard drive power management enabled as shipped".

recommends that similar language be used to encourage power management of short idle modes enabled as shipped.

Thus, CEC should clarify test procedure language so that additional drives may be spun down in short idle, using similar language to the IEC 62623 language for ENERGY STAR v5.2-compliant long idle.

Finally, in the event that CEC maintains this adder, it needs to tighten potentially ambiguous language that currently does not specify whether to apply the adder only *once* (per ENERGY STAR v6.1) or for *each* secondary storage drive. The CASE team recommends that this adder need only be applied once, as per ENERGY STAR, and because there is no reason to have multiple drives active in short and idle modes.

4 Panel Self-Refresh

The CASE team continues to investigate the technical feasibility, energy savings potential and cost effectiveness of panel self-refresh (PSR) technology, cited in our original 2013 CASE proposal as a promising energy-saving measure in desktops and notebooks.

The image on a computer's display must be refreshed 60 or more times per second. Displays and the computer's GPU normally must continue this refresh process even when the on-screen image is *static*. PSR is a technology pioneered by Intel and display manufacturer partners that allows a computer's GPU to rapidly power down when the frame is not changing. The frame displayed on the screen is buffered using a small amount of DRAM (less than 10MB) integrated into the display panel. In order for the technology to work, the computer's GPU must be able to negotiate with the timing controller on the display panel to hand off control from the GPU to the frame buffer. This "hand-off" occurs rapidly—on the order of milliseconds—making PSR highly deployable even during brief periods of static screen time. This offers the potential for significant computer energy savings in both idle and active modes.

PSR relies on several enabling technologies and standards. First, the GPU must support PSR. Intel's 4th generation Core architecture (Haswell) and associated integrated graphics support the feature through the Embedded DisplayPort (eDP) interface (the eDP standard has supported PSR since version 1.3, released in 2011). The connected display panel must also contain additional DRAM for the frame buffer, and its timing controller must support PSR. LG's Shuriken panel, intended for notebooks, is one example. PSR technology is currently implemented in a number of notebook products (mainly ultrabook style). Prominent product examples include the HP Spectre x360 and Lenovo Yoga series.

Although the desktop ecosystem is still not fully prepared for PSR technology, the energy savings opportunity is promising. Intel claims that PSR is able to save energy during the vast majority of a computer's on time. During system idle modes, Intel estimates that the video frame is static over 95% of the time (Figure 4.1). During a variety of other real-world tasks, such as reading PDF documents, browsing the web, or editing word processing documents, the frame is still estimated to be static over 80% of the time. This means that any PSR savings should be able to apply to all short idle periods (we assume that during long idle, the integrated GPU is already in a power-saving state

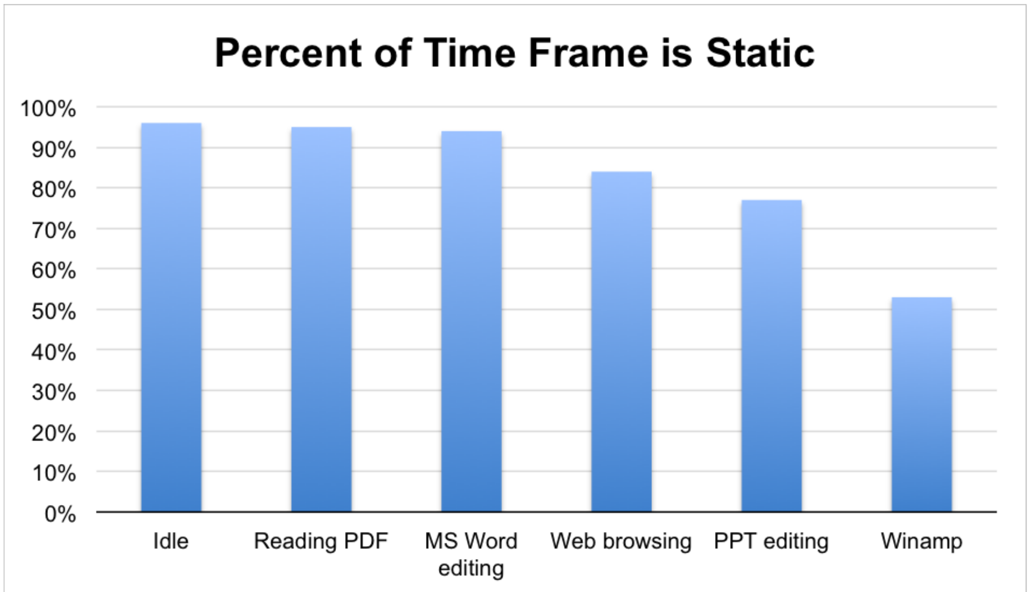

Figure 4.1: Idle frame time. Source: Linder, 2012.

As to the actual energy savings achievable and the associated costs, little data is currently available. Upon introducing the technology for notebook applications, Intel claimed 0.5W savings (likely DC power) and 45 minutes to 1 hour of battery life extension. Given the slightly larger power footprint of desktop CPU-GPU packages, this value may be greater on desktop platforms.

Our research indicates that PSR is a viable near-term energy savings mechanism for notebooks and potentially integrated desktops. The CASE team continues to research both the energy savings potential and associated costs of this technology for desktop platforms and will provide updates in future submittals.

5 Adjustment to Cost-effectiveness Analysis

5.1 Duty Cycle

5.1.1 Network Connectivity

The CASE Team recommends that CEC should require a single conventional duty cycle for determining TEC. ENERGY STAR's network connectivity mode weightings are meant to provide an incentive for manufacturers to implement network connectivity in sleep modes. However there is no evidence that network connectivity in sleep mode actually reduces computer on time. These numbers are therefore arbitrary, and may under-represent active modes, weakening the standards. This could significantly reduce energy savings over the life of the standard, particularly if network connectivity in sleep mode becomes widely available. In the absence of evidence of the benefits of network connectivity, it would be prudent, and simpler, to treat all computers equally and use a single duty cycle.

5.1.1 Estimates for Cost-effectiveness

The CASE Team continues to recommend that CEC utilizes the revised duty cycle for desktops proposed by the IOUs in determining energy savings estimates and cost-effectiveness rather than adopt the ENERGY STAR Version 6.1 duty cycle, which is based on just two, outdated studies; it does not consider several additional studies of computer duty cycle. The IOUs compiled these numerous studies and proposed a more representative duty cycle in their most recent submission at the end of 2014 (CA IOUs 2014). Moreover, inclusion of the most recent and only California-focused study conducted by the California Plug Load Research Center in 2013 further supports the recommendation for desktops, which in the non-residential sector are in non-sleep or off modes 77% of the time. Based on a residential and non-residential weighted average by sample size (each study assigned a value of 1-10, rather than absolute sample size) we recommend the revised mode weighting for conventional desktops below in Table 5.1.

Table 5.1: Estimated Duty Cycle for Each Form Factor

Mode	Conventional Desktops, Integrated Desktops and Thin Clients	Notebooks	Workstations	Small-scale Servers
Off	30%	25%	35%	0%
Sleep	10%	35%	10%	0%
Long Idle	20%	10%	15%	100%
Short Idle	40%	30%	40%	0%

5.2 Real-World Adjustment Factor

The CASE Team proposes that CEC utilize our latest published $TEC_{adjustment}$ factors (Figure) with a 50% weighting when applying them to energy savings and cost effectiveness estimates.

The CASE Team has been examining the impacts of real-world usage on computer energy use since mid-2014 and hase generally found that ENERGY STAR-reported TEC values need to be adjusted to determine actual energy use in the field. We provided initial results of this research as part of our October 2014 CASE report addendum (CA IOUs 2014). Our most recent analysis shows that computers can consume anywhere from 15 to 40% more energy under real-world conditions compared to ENERGY STAR-reported TEC (see Figure), not including the duty cycle adjustment discussed above. As discussed in earlier comments, such adjustments for field usage could increase energy savings for the same cost, and therefore increase cost effectiveness for the proposed standard. We recommended CEC use a TEC adjustment factor in its analyses. Specifically, we recommend adopting the adjustments presented in Figure with a 50% weighting. Accordingly, this would mean an adjustment of 7.25% for desktops, 12.5% for integrated desktops, and 20.5% for notebooks.

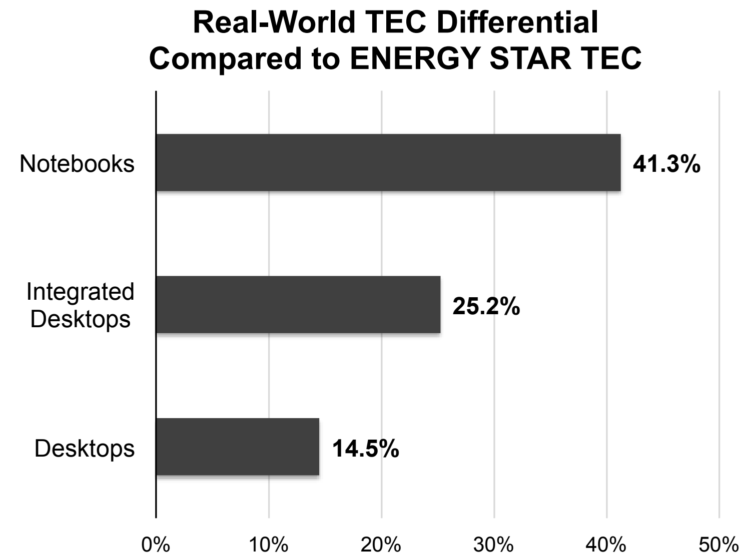

Figure 5.1: Percent increase in TEC under real-world test conditions compared to ENERGY STAR measured TEC

While some of the energy-using components in a system (e.g. hard drives, optical drives, network interfaces) demand about the same amount of power whether the system is under short idle conditions or intense workloads, other system components—CPUs, GPUs, power supplies and memory, for example—will be driven harder and consume more power under real-world loading.

To illustrate the concept, consider the energy savings calculations below between a base system (B) and efficient system (E). Each system's energy budget consists of "scalable" and "fixed" portions. Real-world workloads only exercise the power-scalable pieces of hardware, so in the energy savings calculation, the $TEC_{adjustment}$ only multiplies the scalable part of the energy budget:

$$TEC_{B,real\ world} = TEC_{B,scalable} * TEC_{adjustment} + TEC_{B,fixed}$$

$$TEC_{E,real\ world} = TEC_{E,scalable} * TEC_{adjustment} + TEC_{E,fixed}$$

$$TEC_{savings,real\ world} = TEC_{B,real\ world} - TEC_{E,real\ world}$$

$$= (TEC_{B,scalable} - TEC_{E,scalable}) * TEC_{adjustment} + TEC_{B,fixed} - TEC_{E,fixed}$$

The $TEC_{adjustment}$ should only apply to a portion of the energy savings, but which portion and how much? Our technical team has established two basic criteria for applying this factor. Energy savings measures that 1) provide meaningful savings in idle and active modes and 2) are power-proportional (generate larger savings as the system's workload and power increase) are applicable. Figure 1 illustrates a number of promising energy-savings strategies for computers and maps them by mode and power proportionality. Measures highlighted in green are considered applicable. Measures highlighted in red either do not apply to idle/active modes or provide fixed savings.

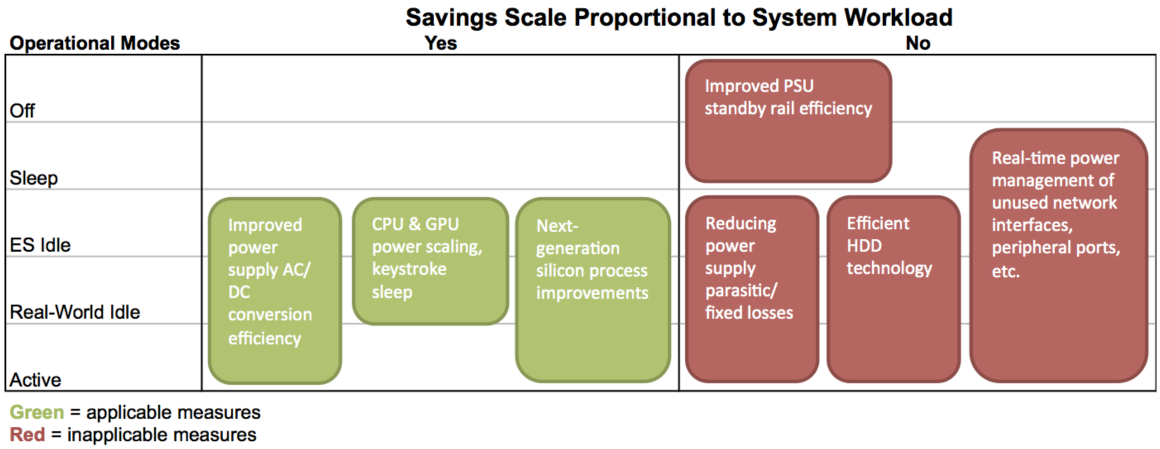

Figure 1.2: Matrix of computer energy efficiency measures and applicability to $TEC_{adjustment}$ multiplier

The energy savings from several compliance strategies are likely to be increased when taking real-world usage into account. Improving AC/DC power supply efficiency remains a promising, cost-effective compliance path, and associated savings will increase if systems are more heavily loaded. Similarly, die shrinkage and continued refinements to silicon components will continue to reduce power demand at a variety of load points. Implementation of low-latency, device-level power management—sometimes referred to as "keystroke sleep"—will also help to lower power demand across a range of CPU and GPU load conditions.

We cannot anticipate the exact paths that manufacturers will pursue to comply with the standard, but we assume that at least 50% of the savings generated by the standard will be derived from a combination of the strategies highlighted in green. The CASE Team therefore proposes that CEC utilize our latest published $TEC_{adjustment}$ factors (Figure) with a 50% weighting when applying them to energy savings and cost effectiveness estimates.

6 Power supply requirements

6.1 Desktops, integrated desktops and thin clients

We strongly recommend that CEC include efficiency requirements for internal power supplies, in addition to TEC requirements. We propose 80-PLUS Gold levels with additional 10 percent load efficiency requirements of 84 percent, and a power factor requirement at all four load points.

6.1.1 80-PLUS Gold levels

CEC's proposal does not include power supply efficiency requirements, contrary to ENERGY STAR and EU Ecodesign. External power supplies are already subject to federal standards, resulting in a transformation of the market, whereas many internal power supplies in today's computers are still very inefficient.

Opponents of this requirement argue that power supply efficiency is only one of the pathways for meeting TEC requirements, and manufacturers should be given the flexibility to meet TEC requirements in whichever way they want. We agree with this generally, however not having a power supply efficiency requirement would still leave opportunity for additional cost-effective savings. Moreover, having the requirement guarantees energy savings in active mode; CEC draft standards are appropriately focused on idle mode, and some of the potential compliance techniques such as graphics switching may not save as much or any energy in active mode.

We propose 80-PLUS Gold based on break-even cost-effectiveness analysis utilizing data provided by Information Technology Industry Council (ITI) in response to the Invitation to Participate (2013) and the April 15th, 2015 computers workshop (2015). The assumptions and results are highlighted below in Table 6.1.

In summary, ITI reports a 2013 incremental cost of $12.15 for improving PSU efficiency between the market baseline of 68% to 80 PLUS gold, for a 300 Watt PSU, which we estimate to be typical size. Assuming $.16 per kWh and a 5-year design life, the amount of energy savings needed for the customers to incur no additional cost is 15 kWh/yr. The conversion efficiency of the baseline and efficient computer systems is defined as:

$$\eta_{base} = \frac{E_{DC}}{E_{AC,base}} \text{ and } \eta_{eff} = \frac{E_{DC}}{E_{AC,eff}}.$$

Since only the power supply is being replaced, the DC power budget on the systems, E_{DC}, is the same in both cases. The AC energy consumption of the more efficient system can then be written as:

$$E_{AC,eff} = E_{AC,base} \frac{\eta_{base}}{\eta_{new}}.$$

The energy saved by replacing the power supply is then:

$$E_{saved} = E_{base}(1 - \frac{\eta_{base}}{\eta_{new}}).$$

Using this conversion, the energy used by the base system needs to be greater than 57 kWh/yr. A typical desktop computer today uses about 180 kWh/yr accounting for real-world energy use (both real-world adjustment factor and duty cycle modifications), with about 95% of this energy used in idle/active modes that would be impacted by improved power supply efficiency. We recommend using this value for the baseline since the savings would be achieved as one compliance pathway, however if assuming the power supply savings is accounted for after the TEC requirement is in place, the energy consumption would be 72 kWh/yr using the real-world adjustment factor and duty cycle modifications. Regardless of the scenario used, the CASE Team proposal would be cost-effective, with **benefit-cost ratios ranging from 1.28-3.17** when using experience curves and anticipating the costs by the effective date of 2018. Benefit-cost ratios for 80 PLUS Silver and Bronze using the same methodology are also provided for reference, and are even greater.

Table 6.1: Incremental Cost of Efficiency Analysis - PSUs

PSU Efficiency Improvement	**2013 Values**	**Source**
68% to 80 PLUS Bronze	$ 5.18	ITI Response to CASE proposal, July 29, 2013
68% to 80 PLUS Silver	$ 10.35	
68% to 80 PLUS Gold	$ 12.15	
68% to 80 PLUS Platinum	$ 16.88	
Assumptions		
$/kwh	$ 0.16	CEC Staff Report 2015
Design Life	5	CEC Staff Report 2015
Real World Energy Use Factor	0.35	Calculated based on CASE Team proposed Real-World Adjustment Factor and revised duty cycle
Idle mode savings only adjustment	0.95	Based on ENERGY STAR QPL 2015 % of Sleep and Off Mode relative to TEC
CAGR	4%	Communications w/ Power Integrations October 2013.
Years until effective date	5	CEC Staff Report 2015
Breakeven TEC savings per year		
68% to 80 PLUS Gold	12.38	Calculated using above incremental cost, $/kWh and design life
68% to 80 PLUS Silver	10.55	
68% to 80 PLUS Bronze	5.28	
Breakeven TEC		
68% to 80 PLUS Gold	57	Calculated, using efficiency at 20% Load
68% to 80 PLUS Silver	53	
68% to 80 PLUS Bronze	31	
"Baseline" Avg. TEC		
Business Desktop	180	Calculated using 140 kwh/yr (ITI presentation, April 15, 2015), Real World Energy Use Factor, and idle mode only adjustment
CEC Proposal	72	Calculated w/ Real World Energy Use Factor
Proposed Base Allowance	50	CEC Staff Report 2015
Proposed Memory Adder	6	CEC Staff Report 2015, assumed 8 GB.
Annual TEC Savings		
68% to 80 PLUS Gold (for Business Desktop)	39.21	Calculated.
68% to 80 PLUS Gold (for CEC Proposal)	15.80	Calculated.
Life Cycle Cost Savings		
68% to 80 PLUS Gold (for Business Desktop)	$ 31.37	Calculated.
68% to 80 PLUS Gold (for CEC Proposal)	$ 12.64	Calculated.
Benefit Cost Ratios (2018)		
68% to 80 PLUS Gold (for Business Desktop)	**3.17**	Calculated.
68% to 80 PLUS Gold (for CEC Proposal)	**1.28**	Calculated.
68% to 80 PLUS Silver (for Business Desktop)	3.40	Calculated.
68% to 80 PLUS Silver (for CEC Proposal)	1.37	Calculated.
68% to 80 PLUS Bronze (for Business Desktop)	5.81	Calculated.
68% to 80 PLUS Bronze (for CEC Proposal)	2.34	Calculated.

6.1.2 10 Percent Load Efficiency Requirements

In addition to 80 PLUS, we strongly recommend that CEC includes efficiency requirements of 84 percent at 10 percent load for desktops, workstations and small-scale servers.

10 percent load is a proxy for the idle load range. As computers, workstations and servers are becoming better able to scale power between idle and active mode, their idle load point has fallen below 20 percent and can be found anywhere between 5 and 15 percent for most computers. And even in active mode, computers are better able to dynamically ramp down their power demand when not performing resource-intensive tasks. As a result, typical computers spend an increasing share of their time in the 5-15% load range.

Unfortunately, the 80 PLUS standard test points of 20, 50 and 100 load focus on the active load range, and do not guarantee a decent efficiency below 20 percent. An 80 PLUS power supply with poor efficiency at the idle load point, would significantly impact overall system efficiency.

In fact, the 80 PLUS program has been testing all power supplies at the 10 percent load point since January 2012, despite this load point not being part of the 80 PLUS standard. The test data is available on the 80 PLUS website[15]. An analysis of this data in Figure 6.1 shows that the range of efficiencies is twice as large at 10 percent as at 20 percent load, confirming that the 10 percent load range is not consistently optimized.

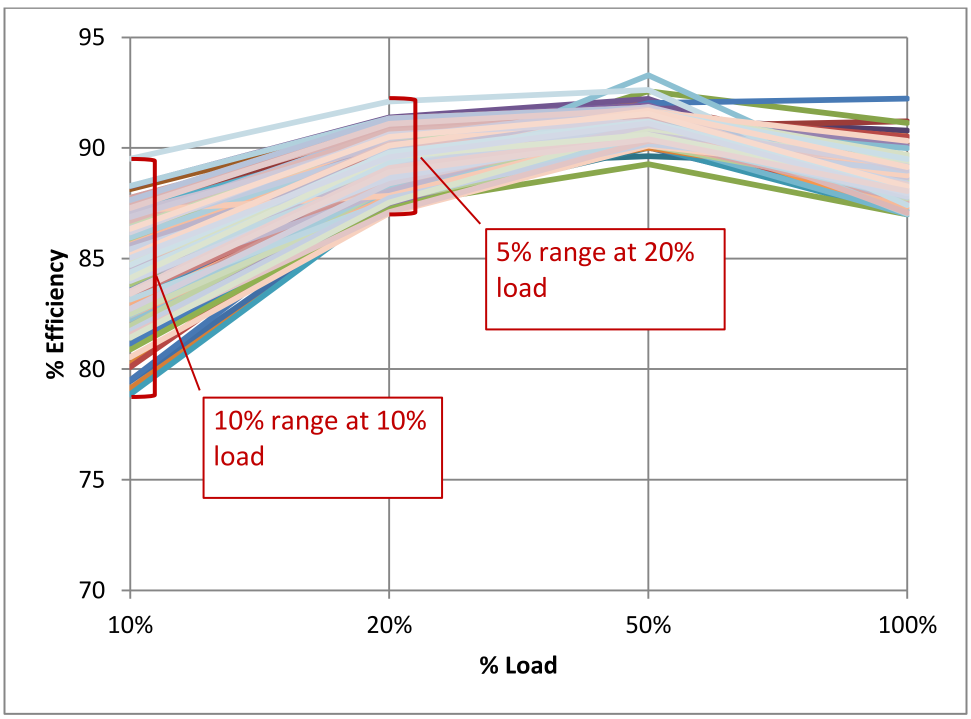

Figure 6.1: Efficiency Profiles of 80 PLUS GOLD Power Supplies

[15] http://www.plugloadsolutions.com/80pluspowersupplies.aspx

We recommend the CEC adopts efficiency requirements of 84 percent for 80 PLUS Gold, in alignment with ENERGY STAR v6.1's power supply efficiency incentive allowance.

This requirement does not significantly add to the test burden for manufacturers, since the 10 percent load test can be performed using the same test setup and equipment as other load points, adding negligible time and cost to the 80 PLUS testing.

6.1.3 Power Supply Power Factor Requirements

The CASE report provided recommendations both for internal power supply efficiency and power factor, mirroring requirements of the ENERGY STAR Version 6.1 computer specification and the 80 PLUS program. We would like to clarify and expand upon our recommendations in this regard. After further analysis, we recommend that CEC requires a more comprehensive set of power factor requirements for computer internal power supplies that covers all major load conditions and harmonizes with ENERGY STAR Version 6.1 and 80 PLUS program requirements. This is important in order to prevent a loophole that would mostly eliminate the benefits of the current power factor requirements

Our proposed requirements, seen in Table 6.2 below, would cover 10, 20, 50, and 100% load points as with proposed efficiency levels. They would be identical to ENERGY STAR Version 6.1 and all levels of 80 PLUS PSU certification (Standard, Bronze, Silver, Gold, and Platinum). They would also extend modest power factor requirements down to lower load levels where computers are anticipated to spend most of their operational hours.

Table 6.2: Proposed Power Factor Requirements

	Load Condition			
	10%	**20%**	**50%**	**100%**
80 PLUS Standard requirements	-	-	-	0.9
80 PLUS Bronze/Silver/Gold/Platinum requirements	-	-	0.9	-
ENERGY STAR Version 6.0 requirements	-	-	-	0.9
CASE report proposed requirements	-	-	-	0.9
New IOU proposed requirements	**0.8**	**0.8**	**0.9**	**0.9**

A random sampling of 80 PLUS certification reports from 2013 indicates that a wide variety of power supplies — varying in rated output wattage, efficiency level, form factor, and manufacturer — will be able to meet these requirements. Despite this widespread compliance with the 80 PLUS program requirements, a more comprehensive power factor requirement is in the best interests of California's rate payers, utilities, and grid operators to encourage improved power quality on the grid. It would also prevent a loophole whereby power supplies could be designed to meet power factor requirements at 50% and 100% load, but switch off power factor correction at lower load levels in order to gain in efficiency. This behavior was already observed in external power supplies by a 2012 European study.[16]

7 Power management requirements: Hibernate

In addition to ENERGY STAR's basic power management requirements (display off after 15 minutes and power down to low-power mode after 30 minutes), CEC should require that computers transition to hibernate mode after 4 hours or less in sleep mode.

In sleep mode, computers continue to draw between 1 watt (notebooks) to 2 to 3 watts (desktops) for as long as the computer is unused. This could last days, or weeks. While sleep mode with a latency of 5 seconds or less is justified when the computer is used frequently (such as for a lunch break), it is not justified when computers are unused for long periods of time, such as over the weekend or when people are away on vacation. Many notebooks are already configured by default to transition to hibernate automatically after several hours in sleep mode, because of battery life considerations. Desktops should do the same. The capability already exists in all computers today, it just needs to be implemented by default.

In the April 15th, 2015 workshop, stakeholders conveyed that display off and auto-power down requirements may not be appropriate for some particular computer uses. The CASE Team is supportive of limited exemptions of power management requirements if these uses can be clearly and narrowly defined.

8 Proximity sensors and auto-brightness control (ABC)

In addition to time-based power management, there is an opportunity for CEC to require power management based on the presence of the user in proximity of the computer and ambient lighting levels:

Proximity-based power management: require occupancy sensors on notebooks and integrated desktops so that when no one is in the room, there is no need for the display to be on and other computer features ready to respond within a millisecond. This is an opportunity to transition the computer into long idle mode, including switching off the display, and engaging other long idle power management strategies such as powering down the disk and other components.

Auto-brightness control: this capability is already available in most notebooks because of battery use reasons. It should be implemented in integrated desktops and enabled by default in both notebooks and desktops.

9 Product Definitions and Information Provision

9.1 Summary

CEC's definitions of the types of computers covered by the standards need to be revised to avoid any misinterpretations and ensure that they do not unintentionally open up loopholes in the standards. The CASE Team included proposed definitions in its May 29th, 2015 submittal and is

[16] ITU-GeSI 2012: An energy-aware survey on ICT device power supplies

currently working on revisions and a possible joint proposal, where applicable, with the Information Technology Industry Council.

10 Certification Requirements Proposal

Computers are highly configurable, with thousands of possible configurations per product family. Registering all possible configurations is not practical. We propose the following approach that should reasonably assure that all configurations comply, while representing a reasonable reporting requirement for manufacturers.

ENERGY STAR requires that the highest energy configuration of a family or series in each applicable ENERGY STAR category meets the specification. This is based on the assumption that if the highest energy configuration meets the specification, it is likely that all configurations of this family will. With CEC's single category approach, this would mean only one configuration registered for each computer family. This highest energy configuration would not be representative of computers actually sold in California, and would not guarantee that other models in the family meet the standards.

We propose the following approach:

For each form factor in a computer family, defined as computers having the same external physical shape (including screen size, but not color, for example: Tower, Desktop, Small Form Factor), register the following configurations:

Highest energy configuration: Vendor-defined configuration that represents the form factor's TEC worst-case.

Rationale: Energy worst case; one of the configurations the most at risk of non-compliance. However does not necessarily represent a commonly sold configuration.

Most typical: Vendor-defined configuration that is expected to be sold the most.

Rationale: Provides the best assurance that models sold the most comply. While manufacturers cannot predict the exact highest selling model at time of registration, this provides a directional requirement that the tested configuration should not be atypical, which can be enforced by analysis of the database.

Lowest energy configuration: Vendor-defined configuration that represents the form factor's TEC best-case.

Rationale: Provides a lower bound for the family and ensure that it can comply with no or minimal adders.

11 Data Submittal Requirements

The data submittal requirements could be enhanced by adding a few extra items (as proposed in red in the Title 20 Table X below). These would assist in any technical reviews or enforcement activities of computer energy efficiency for products sold on the California market.

The CEC should also provide definitions for some of the items in the data submittal table where these are not provided elsewhere in the document. For example, definitions for the different classes of GPUs would be required to inform both users and manufacturers.

Section 1606 of Title 20 - Table X

	Appliance	*Required Information*	*Permissible Answers*
V	Computers	Computer Type	Desktop, Integrated Desktop, Notebook, Small-Scale Server, Workstation, Thin-Client, Portable All-in-Ones, Mobile Thin Clients.
		Manufacturer	
		Model Name	
		Model Number	
		Operating System	
		CPU Name	
		Base Core Frequency (gigahertz)	
		Number of Physical ~~Execution~~ Cores	
		Amount of RAM (gigabytes)	
		Number of RAM modules	
		Discrete Graphics	None, G1, G2, G3, G4, G5, G6, G7, G8+
		Switchable graphics functionality	Yes, No
		Switchable graphics enabled during testing	Yes, No
		Does the computer have an integrated display?	Yes, No
		Diagonal screen size (inches)	
		Viewable screen area (square inches)	
		Resolution (megapixels)	
		Enhanced Performance Display	Yes, No
		Length of time of user inactivity before entering sleep (minutes)	
		Length of time of user inactivity before placing display into sleep (minutes)	
		Energy Efficient Ethernet Capability	Yes, No
		Internal PSU efficiency at 10 %, 20 %, 50 % and 100 % of rated output power	
		Internal PSU Power Factor at 10 %, 20 %, 50 % and 100 % of rated output power	
		External PSU efficiency	
		Off mode power (watts)	
		Sleep mode power (watts)	
		Long-idle power (watts)	
		Short-idle power (watts)	
		Total Annual Energy Consumption (kilowatt hours per year)	

References

Tom's Hardware, 2015. "2015 VGA Charts, 22. Power Consumption Idle." Accessed May 22, 2015. http://www.tomshardware.com/charts/2015-vga-charts/22-Power-Consumption-Idle,3685.html.

[CEC] California Energy Commission. 2015. Analysis of Computers, Computer Monitors, and Signage Displays. Staff Report. 2014 Appliance Efficiency Pre-Rulemaking. Docket Number 14-AAER-2.

[CEC] California Energy Commission. 2015a. Letter to Information Technology Industry Council. Computers. Computer Monitors, and Electronic Displays.

[CA IOUs] California Investor-Owned Utilities. 2014. Response to CEC Invitation to Submit Proposals. Codes and Standards Enhancement (CASE) Initiative: Title 20 Standards Development, Analysis of Standards Proposal for Computers, Addendum to submittal on August 6, 2013. CEC Docket #12-AAER-2A. Prepared by Energy Solutions, NRDC, Ecova and Aggios.

[CA IOUs] California Investor Owned Utilities. 2013. Codes and Standards Enhancement (CASE), Analysis of Standards Proposal for Computers. Prepared by Energy Solutions, NRDC, Xergy and Ecova. Submitted to CEC. August 6th, 2013

[CA IOUs & NEEA] California Investor Owned Utilities and Northwest Energy Efficiency Alliance. 2013. Updated Discrete Graphics Processing Unit Testing Results. Stakeholder Comments on the Draft 3 Version 6.0 ENERGY STAR Computer Specification. http://www.energystar.gov/sites/default/files/specs//California%20IOU%20%20NEEA%20Comments_Public.pdf

[CLASP] Collaborative Labeling and Appliance Standards Program and Natural Resources Defense Council. 2012. The Impact of Graphics Cards on Desktop Computer Energy Consumption.

[EPA] U.S. Environmental Protection Agency. 2015. ENERGY STAR; Computers for Consumers. Qualified Products List. http://www.energystar.gov/index.cfm?fuseaction=find_a_product.showProductGroup&pgw_code=CO. Accessed May,1 2015.

[EPA] U.S. Environmental Protection Agency. 2013a. ENERGY STAR Program Requirements. Product Specification for Computers. Final Test Method. Rev. Oct-2013. http://energystar.gov/products/specs/node/143.

[EPA] U.S. Environmental Protection Agency. 2013b. ENERGY STAR Program Requirements. Product Specification for Computers. Eligibility Criteria. Version 6.0. Rev. Oct-203 http://energystar.gov/products/specs/node/143.

[ITI/TechNet] Information Technology Industry Council. 2015. "PC's Cost Effectiveness/Technical Barriers," April 15, 2015. CEC Staff Workshop.

[ITI/TechNet] Information Technology Industry Council. 2015a. "PCs - Methodology/Framework." April 15, 2015. CEC Staff Workshop.

Linder, Brad. 2012. "Intel: Future could use less power Panel Self Refresh tech." Liliputing.com. Accessed August 5, 2015. http://liliputing.com/2012/04/intel-future-could-use-less-power-panel-self-refresh-tech.html

Appendix A:

Discrete Graphics Cards:

Testing and Energy Impacts

August 7, 2015

Prepared for:

PACIFIC GAS & ELECTRIC COMPANY

SOUTHERN CALIFORNIA EDISON

SAN DIEGO GAS AND ELECTRIC

SOUTHERN CALIFORNIA GAS COMPANY

Prepared by:

Peter May-Ostendorp, XERGY CONSULTING

Katherine Dayem, XERGY CONSULTING

Catherine Mercier, XERGY CONSULTING

Nate Dewart, ENERGY SOLUTIONS

This report was prepared by the California Statewide Utility Codes and Standards Program and funded by the California utility customers under the auspices of the California Public Utilities Commission.

1 Introduction

Discrete Graphics Processing Units (GPUs) provide enhanced graphics performance in computers. Discrete GPU energy use has been measured and documented previously first by efforts in 2011 by the Collaborative Labeling and Appliance Standards Program (CLASP) and the Natural Resources Defense Council (NRDC) and (CLASP 2012) then in 2012 and 2014 by the California IOUs (CA IOUs & NEEA 2013) (Figure 1.1).[17] These earlier measurements have documented a decrease in GPU energy use since 2011 as more power-efficient GPU architectures have been released.

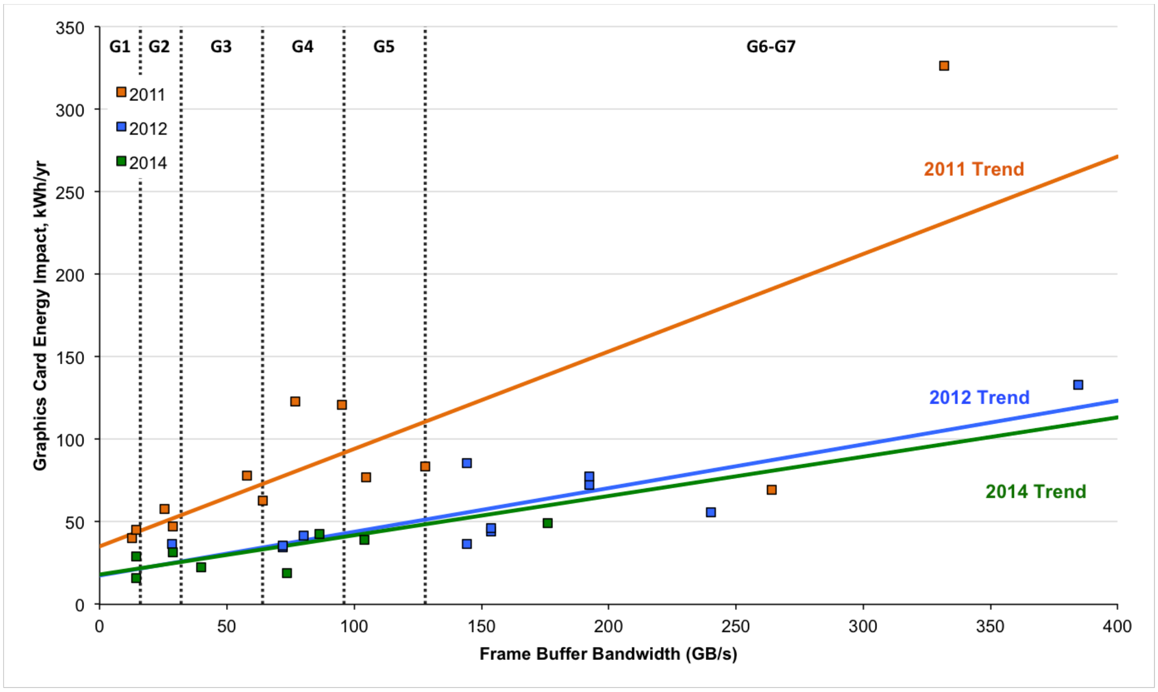

Figure 1.1: Previous results from discrete GPU testing carried out in 2011 (orange markers), 2012 (blue markers), and 2014 (green markers). Note: Graphics card energy impact (kWh/yr) shown on the y-axis is calculated by subtracting the measured system baseline (no discrete GPU) power from the power of that system with the GPU under test installed, and using the ENERGY STAR TEC equation. Each point represents the average of several tests. Linear trends drawn through each data set indicate a decrease in graphics card energy use through time.

Building on the previous work, we determined whether this trend is continuing with the latest GPU products, tested on updated systems. We focused on recently released, higher performance (as measured by frame buffer bandwidth (FBB)) cards.

[17] Note that the 2014 data, previously not available publicly, will be presented in this data report.

2 Test Plan and Procedure

In 2014, the CASE team procured and tested 8 discrete GPUs (GPU 25 through GPU 32, Table 2.1) on six systems previously used in the 2011 and 2012 tests (systems PC 1 through PC 6.2, Table 2.2).[18] In 2015, the CASE team tested 10 discrete GPUs on four systems (Table 2.1). Five of the GPUs were tested previously in 2014. Four GPUs were newly purchased, and one was provided by its manufacturer. The cards range across the performance spectrum, with focus in 2015 on recently released, high-performance cards.

Table 2.1: GPUs tested in 2014 and 2015

GPU Card ID	Test Year	GPU Mfr	GPU	Memory (MB)	ECMA Category (v6)	FBB (GB/sec)	Max GPU Power (W)
GPU 25	2014	AMD	Radeon R7-250	2048	4	73.6	75
GPU 26	Both	AMD	Radeon R9-285	2048	7	176	190
GPU 27	Both	AMD	Radeon R7-240	2048	2	28.8	30
GPU 28	Both	AMD	Radeon R7-260X	2048	5	104	115
GPU 29	Both	NVIDIA	GeForce GTX 750ti	2048	4	86.4	60
GPU 30	2014	NVIDIA	GeForce GT 720	1024	1	14.4	19
GPU 31	2014	NVIDIA	GeForce GT 730	1024	3	40	25
GPU 32	Both	NVIDIA	GeForce GT 730	1024	1	14.4	23
GPU 33	2015	NVIDIA	GeForce GTX 960	2048	5	112	120
GPU 34	2015	NVIDIA	GeForce GTX 970	4096	7	224	145
GPU 35	2015	NVIDIA	GeForce GTX 980	4096	7	224	165
GPU 37	2015	NVIDIA	GeForce GTX Titan X	12000	7	336	250
GPU 38	2015	NVIDIA	GeForce GTX Titan X	12000	7	336	300

Three of the four systems used for the 2015 testing were OEM models. We procured systems across a range of performance, using more recent CPU architectures than in earlier 2011-2014 tests. The fourth system (PC 10) was a custom-built, high-end gaming system. Two of the systems (PC 9 and PC 10) had AMD integrated graphics, which, when paired with a compatible AMD graphics card, should allow graphics switching – the ability to switch between the integrated and

[18] The 2014 testing was conducted by Ecova, Inc. This data was previously unpublished, thus we present it here. We limit the test method discussion to 2015 testing activities, which were conducted by two of the authors of this report at Xergy Consulting, LLC.

discrete GPUs according to activity and to power down the discrete card when not in use. PC 8 was selected for its purported graphics switching capabilities, but we were unable to observe proper switching operation in our testing. PC 7 was chosen as an entry-level system. With a regular tower form factor, it had space for expansion. PC 8 and PC 9 had more compact form factors.

Table 2.2: Hardware attributes of 2015 GPU Test Systems

PC ID	Test Year	Market Class	CPU Cores	CPU Speed (GHz)	Integrated Graphics	Base Memory (GB)	Number and Type of Memory DIMMs	Storage	Operating System
PC1	2014	Entry	2	3.06	Intel GMA X4500	2	1 x 2GB DDR3	500 GB HDD	Windows 7 Home Premium
PC2	2014	Slim	4	2.9	AMD Radeon HD 6550D	2	1 x 2GB DDR3	500 GB HDD	Windows 7 Professional
PC3.2	2014	Basic commercial	4	2.8	Intel HD 2000	2	1 x 2GB DDR3	100 GB HDD	Windows 7 Professional
PC4	2014	Budget gaming	4	3	AMD Radeon HD 4250	4	2 x 2GB DDR3	500 GB HDD	Windows 7 Home Premium
PC5	2014	Performance	4	3.3	Intel HD 3000	4	2 x 2GB DDR3	500 GB HDD	Windows 7 Home Premium
PC6.2	2014	Enthusiast gaming	4	3.3	Intel HD 3000	8	1 x 4GB DDR3	500 GB HDD	Windows 7 Home Premium
PC7	2015	Entry-level	2	3.4	Intel HD Graphics 4400	4	1 x 4GB DDR3	500 GB HDD	Windows 8.1
PC8	2015	Mid-range gaming	4	3.2	Intel HD Graphics 4600	8	2 x 4GB DDR3	1 TB HDD + 8 GB SSHD	Windows 8.1
PC9	2015	High-end business desktop	4	3.7	AMD Radeon R7	8	2 x 4GB DDR3	1 TB HDD	Windows 7 Professional
PC10	2015	High-end custom built gaming	4	4	AMD Radeon R7	16	2 x 8GB DDR3	1 TB HDD	Windows 8.1

An inherent challenge in power testing computer systems is acquiring stable and representative intervals over which measurements can be averaged. To help produce stable results, before testing we "seasoned" each system by performing all system updates and removing or manually exiting applications/processes that may start and run intermittently during system idle modes, such as network connection managers, system update schedulers, or pre-installed anti-virus software. In some cases, we had to run the test several times before we acquired an acceptably stable power log. We averaged power in stable 5-minute windows. Sample output is shown in Figure 2.1.

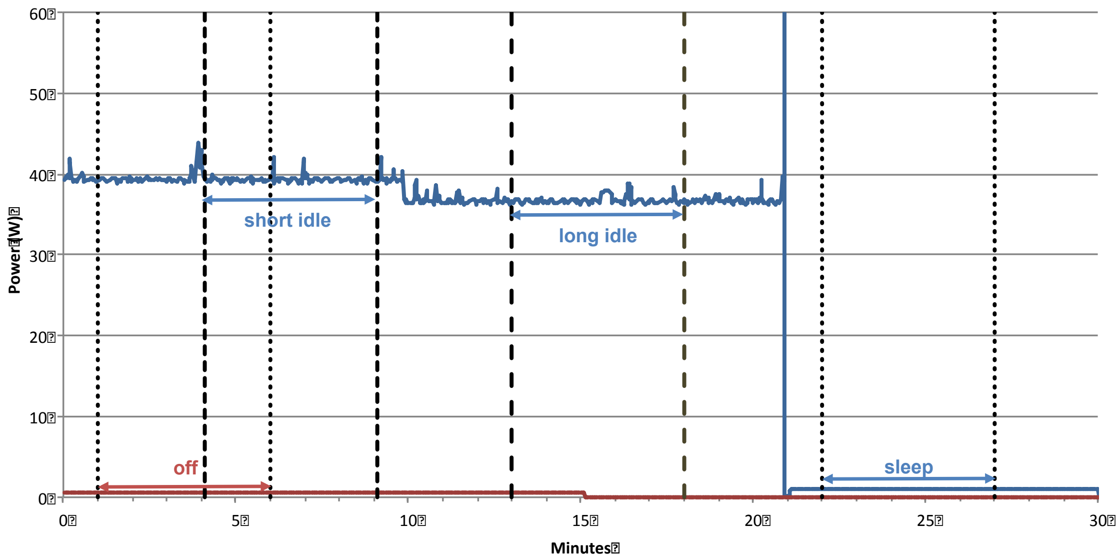

Figure 2.1: Example data log of test, with averaging intervals for each mode

Testing was carried out using calibrated, high-accuracy equipment: a Chroma 61502 AC power source and a Yokogawa WT1803 power analyzer. For each system, we first measured the baseline power state (i.e., tested the system without a discrete GPU). We measured short idle, long idle, sleep, and off power according to the ENERGY STAR test method (EPA, 2013a). Each GPU was installed in each system, and short idle, long idle, sleep, and off power were measured.

Some of the higher performance discrete GPUs required auxiliary power from the internal power supply. The stock power supplies for PC 7 and PC 9 did not include the required 8-pin and 6-pin auxiliary power connectors, so we adapted existing 12V connectors to power these cards. PC 8 uses a 19V external power supply coupled to internal DC-DC converters to achieve required system voltages. When auxiliary power for a discrete GPU is necessary, a second external power supply is required. Unfortunately, we were unable to obtain this secondary power unit during the course of the research, so were only able to measure discrete GPUs on this system that could be fully powered through the PCI express slot. We were also unable to test the high-performance/high-power GPU 38 in PC 9 due to power supply capacity constraints.

In total, we completed 4 baseline and 34 discrete GPU tests (Table 2.3).

Table 2.3: Tests completed on each system. Combinations with an asterisk (*) were tested twice: once with switchable graphics enabled, once with switchable graphics disabled

	PC 7	PC 8	PC 9	PC 10
Baseline	✓	✓	✓	✓
GPU 26	✓		✓	✓
GPU 27	✓	✓	✓*	✓*
GPU 28	✓		✓	✓
GPU 29	✓	✓	✓	✓
GPU 32	✓	✓	✓	✓
GPU 33	✓		✓	✓
GPU 34	✓		✓	✓
GPU 35	✓		✓	✓
GPU 37	✓		✓	✓
GPU 38	✓			✓

3 Results and Discussion

3.1 Graphics card energy impact

Discrete GPU power, TEC, and TEC deltas measured and calculated in the current study are presented in Table 3.1. Data from the 2014 study is included in Table 3.2. TEC was calculated according to the ENERGY STAR TEC formula:

$$TEC = 8.76(P_{off}T_{off} + P_{sleep}T_{sleep} + P_{long\_idle}T_{long\_idle} + P_{short\_idle}T_{short\_idle})$$

Where P_{off}, P_{sleep}, P_{long_idle}, and P_{short_idle} are average power measured in off, sleep, long idle, and short idle modes, respectively, and the ENERGY STAR conventional duty cycle values are T_{off} = 0.45, T_{sleep} = 0.05, T_{long_idle} = 0.15, and T_{short_idle} = 0.35.

The discrete GPU power and energy impacts (or "deltas") were calculated by subtracting the system baseline power or TEC from the system-plus-GPU power or TEC.

We calculated average graphics card TEC deltas by averaging data for each card across the systems on which it was tested (Table 3.3) For the five cards tested in 2014, we averaged across the 2014 and 2015 tests, omitting the maximum and minimum data values. For the five cards tested in 2015, we averaged all tests.

Table 3.1: 2015 Discrete Graphics Test Data. Graphics switching (AMD's "dual graphics") enabled for the tests shaded in gray.

GPU under test	GPU FBB (GB/s)	PC under test	Power, Short Idle (W)	Power, Long Idle (W)	Delta Power, Short Idle (W)	Delta Power, Long Idle (W)	Sleep Power (W)	Standby / Off Power (W)	TEC (kWh/yr)	Delta TEC (kWh/yr)
Baseline	-	PC7	24.3	23.0	-	-	0.82	0.56	107	-
Baseline	-	PC8	25.9	25.0	-	-	1.14	0.84	116	-
Baseline	-	PC9	25.2	24.4	-	-	2.37	0.13	111	-
Baseline	-	PC10	25.8	24.9	-	-	2.12	1.03	117	-
GPU 26	176	PC7	39.9	39.1	15.6	16.1	0.99	0.56	176	69.1
GPU 26	176	PC8	33.6	32.5	7.74	7.52	1.20	0.83	150	33.6
GPU 26	176	PC9	44.6	43.7	19.4	18.5	2.35	0.13	196	84.8
GPU 26	176	PC10	50.5	49.2	24.7	24.3	2.12	1.03	224	108
GPU 27	28.8	PC7	31.8	30.7	7.53	7.66	1.00	0.56	141	33.2
GPU 27	28.8	PC9	29.4	28.2	4.23	3.75	2.37	0.13	129	17.9
GPU 27	28.8	PC9	38.8	37.7	13.7	13.3	2.36	0.13	170	59.3
GPU 27	28.8	PC10	28.6	27.9	2.86	2.98	2.13	1.03	129	12.7
GPU 27	28.8	PC10	35.7	34.7	9.90	9.82	2.14	1.04	160	43.3
GPU 28	104	PC7	30.2	28.6	5.91	5.58	1.93	0.56	133	25.5
GPU 28	104	PC9	34.2	33.4	9.03	8.94	2.35	1.30	155	44.0
GPU 28	104	PC10	35.2	34.5	9.47	9.55	2.14	1.03	158	41.6
GPU 29	86.4	PC7	33.5	31.8	9.22	8.80	1.03	0.56	147	39.9

GPU under test	GPU FBB (GB/s)	PC under test	Power, Short Idle (W)	Power, Long Idle (W)	Delta Power, Short Idle (W)	Delta Power, Long Idle (W)	Sleep Power (W)	Standby / Off Power (W)	TEC (kWh/yr)	Delta TEC (kWh/yr)
GPU 29	86.4	PC8	33.5	33.2	7.55	8.17	1.20	0.82	150	33.8
GPU 29	86.4	PC9	40.3	39.2	15.1	14.0	2.36	0.13	176	65.6
GPU 29	86.4	PC10	36.6	35.3	10.9	10.4	2.14	1.05	164	47.2
GPU 32	14.4	PC7	27.6	27.0	3.34	3.91	0.99	0.56	123	15.5
GPU 32	14.4	PC8	28.8	28.2	2.92	3.18	1.24	0.84	129	13.2
GPU 32	14.4	PC9	36.6	35.6	11.4	10.4	2.35	0.13	161	49.6
GPU 32	14.4	PC10	30.6	29.8	4.85	4.86	2.14	1.04	138	21.3
GPU 33	112	PC7	34.8	33.1	10.5	10.1	1.01	0.58	153	45.5
GPU 33	112	PC9	40.6	39.2	15.4	14.0	2.36	0.13	177	66.5
GPU 33	112	PC10	38.6	36.6	12.9	11.7	2.14	1.03	172	54.8
GPU 34	224	PC7	40.3	38.8	16.0	15.7	1.03	0.60	177	69.9
GPU 34	224	PC9	47.5	46.0	22.3	20.8	2.38	0.13	208	96.6
GPU 34	224	PC10	44.1	42.6	18.4	17.7	2.13	1.04	196	79.6
GPU 35	224	PC7	40.5	38.9	16.2	15.8	0.99	0.64	178	70.8
GPU 35	224	PC9	50.5	48.8	25.3	23.7	2.38	0.13	221	110
GPU 35	224	PC10	44.0	42.5	18.3	17.6	2.15	1.05	196	79.2
GPU 37	336	PC7	39.3	36.7	15.0	13.7	0.99	0.57	172	64.2
GPU 37	336	PC9	42.5	39.6	17.3	15.2	2.40	1.10	188	77.0
GPU 37	336	PC10	42.9	39.0	17.1	14.1	2.14	1.04	188	71.1

GPU under test	GPU FBB (GB/s)	PC under test	Power, Short Idle (W)	Power, Long Idle (W)	Delta Power, Short Idle (W)	Delta Power, Long Idle (W)	Sleep Power (W)	Standby / Off Power (W)	TEC (kWh/yr)	Delta TEC (kWh/yr)
GPU 38	336	PC7	40.8	37.9	16.5	13.6	0.99	0.57	177	70.1
GPU 38	336	PC10	43.2	40.0	17.5	15.1	2.13	1.04	190	73.4

Table 3.2: 2014 Discrete Graphics Test Data

GPU under test	GPU FBB (GB/s)	PC under test	Power, Short Idle (W)	Power, Long Idle (W)	Delta Power, Short Idle (W)	Delta Power, Long Idle (W)	Sleep Power (W)	Standby / Off Power (W)	TEC (kWh/yr)	Delta TEC (kWh/yr)
Baseline	-	PC1	50.7	49.9	-	-	4.91	4.27	240	-
Baseline	-	PC2	33.2	32.9	-	-	1.22	0.19	146	-
Baseline	-	PC3.2	18.7	17.8	-	-	1.93	0.77	84.7	-
Baseline	-	PC4	67.9	67.9	-	-	5.46	3.28	313	-
Baseline	-	PC5	50.7	49.8	-	-	1.93	1.1	226	-
Baseline	-	PC6.2	56.9	55.6	-	-	2.54	1.37	254	-
GPU 25	73.6	PC1	59.2	53.5	8.51	3.61	n/a	n/a	271	30.8
GPU 25	73.6	PC2	40.2	35.2	6.96	2.33	n/a	n/a	171	24.4
GPU 25	73.6	PC4	76.9	71.8	9.02	3.90	n/a	n/a	345	32.8
GPU 25	73.6	PC5	60.2	56.2	9.50	6.38	n/a	n/a	263	37.5
GPU 25	73.6	PC6.2	55.6	49.6	-1.32	-6.04	n/a	n/a	242	-12.0
GPU 26	176	PC1	72.0	58.1	21.4	8.26	n/a	n/a	316	76.3

GPU under test	GPU FBB (GB/s)	PC under test	Power, Short Idle (W)	Power, Long Idle (W)	Delta Power, Short Idle (W)	Delta Power, Long Idle (W)	Sleep Power (W)	Standby / Off Power (W)	TEC (kWh/yr)	Delta TEC (kWh/yr)
GPU 26	176	PC2	48.1	37.8	14.9	4.85	n/a	n/a	198	51.9
GPU 26	176	PC4	85.1	71.7	17.3	3.83	n/a	n/a	370	58.0
GPU 26	176	PC5	70.4	57.6	19.7	7.84	n/a	n/a	297	70.6
GPU 26	176	PC6.2	61.8	51.5	4.94	-4.10	n/a	n/a	264	9.76
GPU 27	28.8	PC1	60.4	54.2	9.71	4.30	n/a	n/a	275	35.4
GPU 27	28.8	PC2	42.0	36.1	8.79	3.19	n/a	n/a	178	31.1
GPU 27	28.8	PC4	77.0	70.9	9.09	3.02	n/a	n/a	344	31.8
GPU 27	28.8	PC5	62.7	57.6	12.0	7.81	n/a	n/a	273	47.1
GPU 27	28.8	PC6.2	65.2	57.2	8.29	1.58	n/a	n/a	281	27.5
GPU 28	104	PC1	60.5	56.0	9.83	6.17	n/a	n/a	278	38.2
GPU 28	104	PC2	44.9	41.1	11.7	8.23	n/a	n/a	193	46.7
GPU 28	104	PC3.2	24.2	20.2	5.43	2.40	n/a	n/a	104	19.8
GPU 28	104	PC4	75.5	70.5	7.64	2.61	n/a	n/a	339	26.9
GPU 28	104	PC5	62.5	58.4	11.8	8.58	n/a	n/a	274	47.6
GPU 28	104	PC6.2	69.2	60.4	12.3	4.77	n/a	n/a	298	44.1
GPU 29	86.4	PC1	63.6	62.3	12.9	12.4	n/a	n/a	296	55.9
GPU 29	86.4	PC2	43.7	42.9	10.5	9.96	n/a	n/a	192	45.2
GPU 29	86.4	PC3.2	27.1	25.9	8.36	8.13	n/a	n/a	121	36.3
GPU 29	86.4	PC4	76.1	75.1	8.25	7.27	n/a	n/a	347	34.8

GPU under test	GPU FBB (GB/s)	PC under test	Power, Short Idle (W)	Power, Long Idle (W)	Delta Power, Short Idle (W)	Delta Power, Long Idle (W)	Sleep Power (W)	Standby / Off Power (W)	TEC (kWh/yr)	Delta TEC (kWh/yr)
GPU 29	86.4	PC5	63.6	62.5	12.9	12.7	n/a	n/a	282	56.3
GPU 29	86.4	PC6.2	64.9	63.5	8.04	7.90	n/a	n/a	289	35.0
GPU 30	14.4	PC1	56.6	54.9	5.92	4.98	n/a	n/a	265	24.7
GPU 30	14.4	PC2	36.7	35.7	3.41	2.81	n/a	n/a	161	14.1
GPU 30	14.4	PC3.2	22.6	21.8	3.88	4.00	n/a	n/a	102	17.2
GPU 30	14.4	PC4	73.7	72.8	5.85	4.90	n/a	n/a	337	24.4
GPU 30	14.4	PC5	58.6	57.7	7.90	7.94	n/a	n/a	261	34.7
GPU 30	14.4	PC6.2	72.4	59.1	15.5	3.49	n/a	n/a	306	52.1
GPU 31	40	PC1	54.7	53.9	4.00	3.98	n/a	n/a	257	17.5
GPU 31	40	PC2	40.2	39.2	6.96	6.25	n/a	n/a	176	29.6
GPU 31	40	PC3.2	24.0	22.9	5.25	5.09	n/a	n/a	107	22.8
GPU 31	40	PC4	75.2	74.5	7.29	6.61	n/a	n/a	344	31.0
GPU 31	40	PC5	59.8	58.9	9.16	9.10	n/a	n/a	266	40.0
GPU 31	40	PC6.2	59.7	58.4	2.81	2.76	n/a	n/a	266	12.2
GPU 32	14.4	PC1	51.6	50.8	0.92	0.89	n/a	n/a	244	3.99
GPU 32	14.4	PC2	38.3	37.4	5.07	4.48	n/a	n/a	168	21.4
GPU 32	14.4	PC3.2	21.8	21.0	3.04	3.24	n/a	n/a	98	13.6
GPU 32	14.4	PC4	72.3	71.7	4.44	3.82	n/a	n/a	331	18.6
GPU 32	14.4	PC5	58.1	57.3	7.44	7.51	n/a	n/a	259	32.7

GPU under test	GPU FBB (GB/s)	PC under test	Power, Short Idle (W)	Power, Long Idle (W)	Delta Power, Short Idle (W)	Delta Power, Long Idle (W)	Sleep Power (W)	Standby / Off Power (W)	TEC (kWh/yr)	Delta TEC (kWh/yr)
GPU 32	14.4	PC6.2	61.2	60.1	4.33	4.51	n/a	n/a	273	19.2

Table 3.3: Average graphics card energy impact for discrete GPUs. Note: Cards highlighted in orange were tested in 2014. Cards highlighted in blue were tested in both the 2014 and the present study, and the ΔTEC shown is averaged across both data sets. The rest of the cards were procured and tested in the present study.

GPU	FBB (GB/s)	Average ΔTEC (kWh/yr)
GPU 25	73.6	29.3
GPU 26	176.0	63.5
GPU 27	28.8	29.5
GPU 28	104.0	38.0
GPU 29	86.4	43.8
GPU 30	14.4	25.2
GPU 31	40.0	25.2
GPU 32	14.4	19.4
GPU 33	112.0	55.6
GPU 34	224.0	82.0
GPU 35	224.0	86.5
GPU 37	336.0	70.8
GPU 38	336.0	71.8

Graphics card energy impact as a function of frame buffer bandwidth is shown in Figure 3.1. Linear trends show that the energy impacts of discrete GPUs have not changed dramatically between the previous round of measurements in late 2014 and the current study. When assessing the need for adders for discrete graphics-only systems, therefore, the IOU team used the combined 2014 and 2015 data.

The 2015 data included data for two extremely high performance, enthusiast GPUs (37 and 38) with idle power and TEC well below the trends from 2011-2014 testing. We also see several cards in lower performance categories (G1-G3) that manage incremental TEC levels less than the trends.

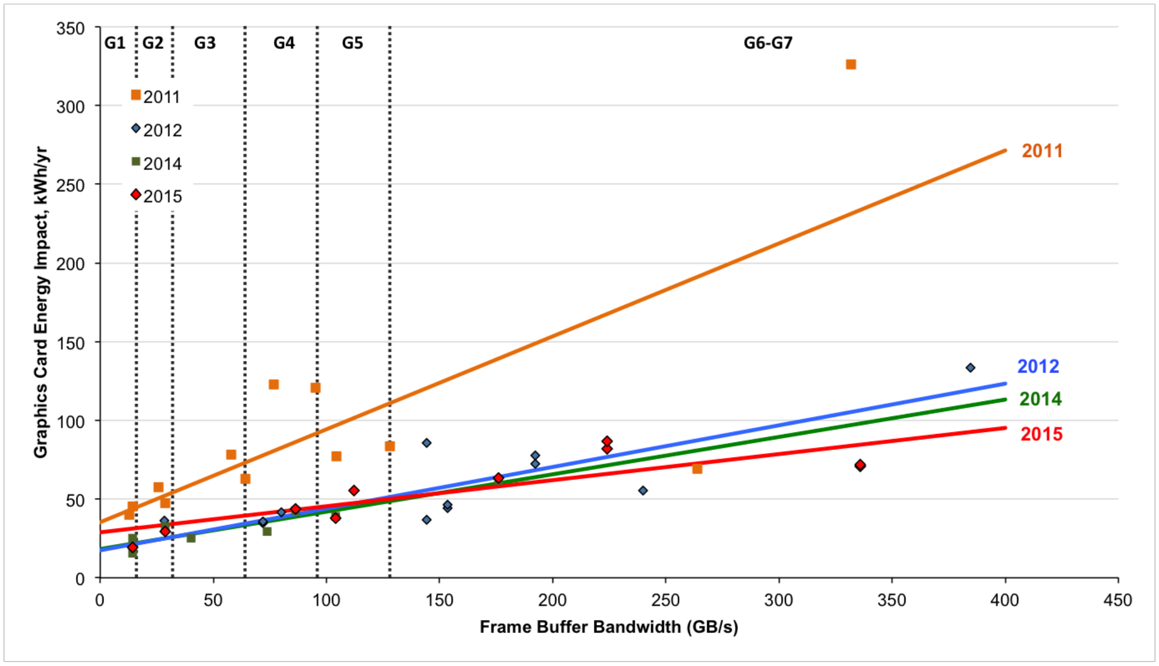

Figure 3.1: Figure 1.1, with average results from the current study in red.

3.2 Graphics switching

Despite claims of graphics switching capability on several of our systems, we had only mild success in observing and measuring it. We enabled AMD's "Radeon Dual Graphics" on PCs 9 and 10 with GPU 27 (AMD Radeon R7-240) installed, but were not successful in enabling switching on another card in the Radeon R7 series, GPU 28 (AMD Radeon R7-260X).

PC 8 was equipped with Lucid VirtuWatt, a graphics switching hardware-software feature that allows the user to specify which GPU individual applications should use. For example, the user can set games to use discrete graphics, while using integrated graphics for everyday tasks like word processing. Unfortunately, we were unable to successfully implement VirtuWatt graphics switching on this particular unit.

Despite difficulties encountered on our particular test systems, graphics switching still appears to be a promising technology for reducing idle power and energy use in systems with both integrated and discrete GPUs. The AMD Radeon R7-240 lowered energy use by about 70% when the Radeon Dual Graphics feature was enabled (Figure 3.2). The CASE team continues to evaluate switchable graphics solutions as they become available in new desktop systems.

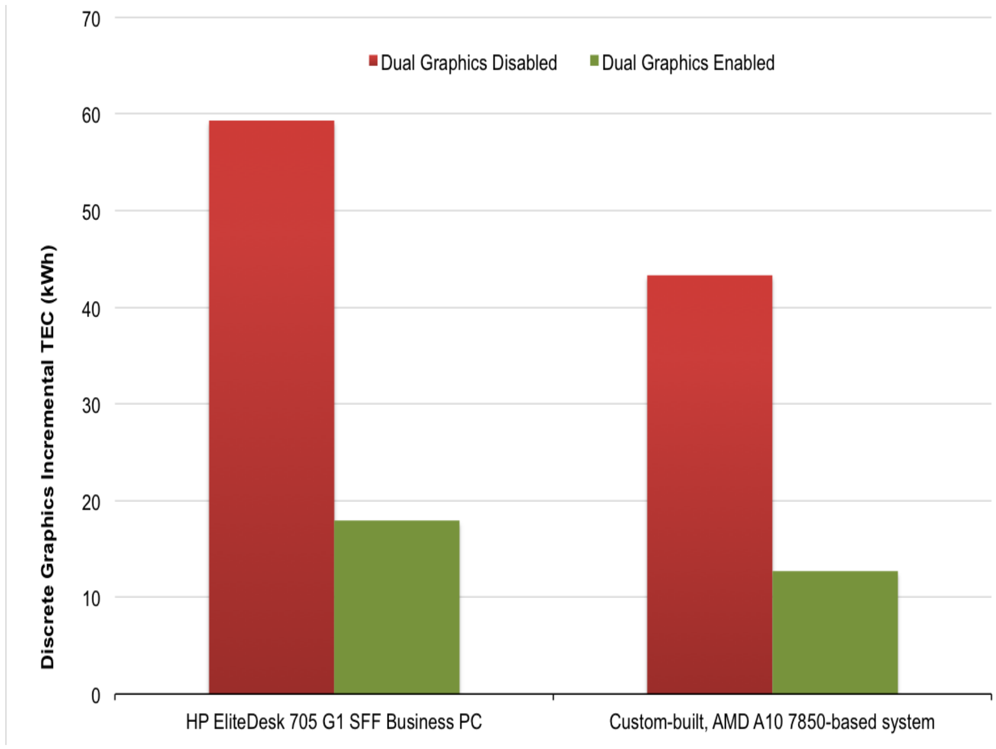

Figure 3.2: TEC impact of the AMD Radeon R7-240 GPU on the two compatible AMD systems, PC 9 (left) and PC 10 (right). Red and green bars show TEC with Radeon Dual Graphics turned off and on, respectively.

Appendix B:

Power Requirements of Security Features in Business Desktop Computers

August 3, 2015

Prepared for:

PACIFIC GAS & ELECTRIC COMPANY

SOUTHERN CALIFORNIA EDISON

SAN DIEGO GAS AND ELECTRIC

SOUTHERN CALIFORNIA GAS COMPANY

Prepared by:

Vojin Zivojnovic, AGGIOS, INC.

Davorin Mista, AGGIOS, INC.

This report was prepared by the California Statewide Utility Codes and Standards Program and funded by the California utility customers under the auspices of the California Public Utilities Commission.

1 Executive Summary

In this appendix, we present our initial findings on the impact of the newest security and manageability features on the idle power consumption of business desktop computers. The tests conducted on 3 representative computers (assembled, desktop tower and small form factor version) indicate no or minimal impact on the power consumption in idle and only a limited impact of around 0.5W while in the sleep state. Our findings appear to be in contrast with ITI's response to Commission's draft report, so additional information about the testing approach taken by ITI is requested.

2 Problem Statement

In their response to the CEC draft report, ITI (2015) stated that no provisions were made for end user performance and feature needs, listing a series of security-related technologies:

> ***Industry Response to CEC Staff Report for Computer Standards*** ***Page 19***
>
> - Absolutely no provision was made for end user performance and feature needs
>   - Security (Example: TPM, VPRO, TrustZone, DMTF DASH, etc.)

In order to investigate the impacts of such technologies on PC power consumption, we researched the available security technologies and analyzed the power draw of representative business PCs incorporating such security features.

3 Security Technologies in Personal Computers

The most robust way to make computers secure is by including security features directly in the hardware instead of relying on the software alone. Intel's vPro® is one such technology offered in computers targeting business customers, as well as security-minded consumers. In fact, vPro is Intel's marketing term referring to a vast collection of security-related features and technologies. In order to enable the vPro features, both the CPU as well as the motherboard must include vPro features.

While the complete feature list of vPro is quite long, some of the most noticeable features of a vPro enabled PC which a user will find include:

- Secure and OS independent remote management with AMT (Active Management Technology)
- Secure and transparent hard drive encryption using TPM (Trusted Platform Module)
- Secure wireless display (WiDi)

Historically vPro was mostly focused on secure remote management, hence vPro and AMT are often used synonymously in articles and discussions found on the internet. The AMT technology in its most recent versions (current 5^{th} generation Intel cores include AMT v9.0) is largely compliant with the DMTF DASH (Desktop and Mobile Architecture for System Hardware) standard. DMTF (Distributed Management Task Force) is an industry standards organization.

AMT allows remote management of PCs without any dependence on an operating system. In fact, as long as the PC is connected to a power source and connected to the network, the computer does not even need to be turned on in order to allow remote management, as AMT uses an out-of-band connection method to communicate with the PC. In other words, it does not rely on an agent running on the PC's operating system, instead it communicates with the PC on a different level, such as through firmware running on a dedicated management processor.

Another technology that all vPro-enabled PCs offer is TPM (Trusted Platform Module), enabling hardware-accelerated encryption and decryption as well as secure key storage.

TPM is a tamper-resistant integrated circuit included on vPro-compliant motherboards that can perform cryptographic operations (including key generation) and protect small amounts of sensitive information, such as passwords and cryptographic keys. The TPM standard is managed by the Trusted Computing Group (TCG), the current standard is TPM 1.2, while TPM 2.0 is in the process of being reviewed. Both AMD and Intel currently support TPM 1.2.

Applications using TPM include:

- File and folder encryption
- Local password management
- S-MIME e-mail
- VPN and PKI authentication
- Wireless authentication

The most widely used TPM-based feature is Microsoft's Bitlocker, used to securely encrypt entire hard drives, including the drive which the OS is installed on (typically referred to as the C-Drive). Bitlocker is a built-in feature of the professional edition of recent Windows versions, available since Windows Vista.

If the computer is equipped with a compatible TPM, BitLocker uses the TPM to lock the encryption keys that protect the data. As a result, the keys cannot be accessed until the TPM has verified the state of the computer. Encrypting the entire volume protects all of the data, including the operating system itself, the Windows registry, temporary files, and the hibernation file. Because the keys needed to decrypt data remain locked by the TPM, an attacker cannot read the data just by removing your hard disk and installing it in another computer.

4 Selection and Configuration of Business Desktop PCs for Analysis

We procured the following three different business-oriented PCs, which incorporate full vPro support, i.e. containing the features mentioned by the industry in their responses to the CEC and AGGIOS reports:

1. Assembled Desktop PC
    - ASUS Q87M-E Motherboard with Intel i7-4770 Core (also tested with the i7-4785T)

- 400W PSU - 80PLUS Platinum
- 8GB RAM, 500GB HDD (WD Blue)
- Intel XTU benchmark: 620 (584 marks with the i7-4785T)

2. Full Size Business Desktop
    - Tower PC with i7-4790 CPU
    - 280W PSU - 80PLUS Bronze
    - Intel XTU benchmark: 623 marks
3. Small Form Factor Business Desktop
    - Small form factor mini-PC with i7-4785T CPU
    - External 65W PSU – 87% efficient according to manufacturer's web site
    - USB-based dock with DVD drive and additional ports
    - Intel XTU benchmark: 581 marks

All 3 PCs are running Windows 8.1 Professional. Each PC was given at least 24 hours to run without any user interactions and with sleep disabled to allow for proper " aging. We also made sure to update all device drivers to their latest versions prior to conducting our power measurements.

The PCs we examined all had the same default configuration regarding vPro functionality:

- AMT is enabled by default, but not configured
- TPM is disabled by default

After performing power measurements in the default (as-shipped) state, we then configured AMT using the Intel System Configuration Software, SCS version 9.1.2.74 downloaded from Intel's website. We verified the correct configuration of AMT by using the AMT WebUI to remotely access and control the PCs in question, verifying the ability to access the PC even when the Windows is shut down, or while booted into the BIOS, as shown in Figure 4.1 and 4.2.

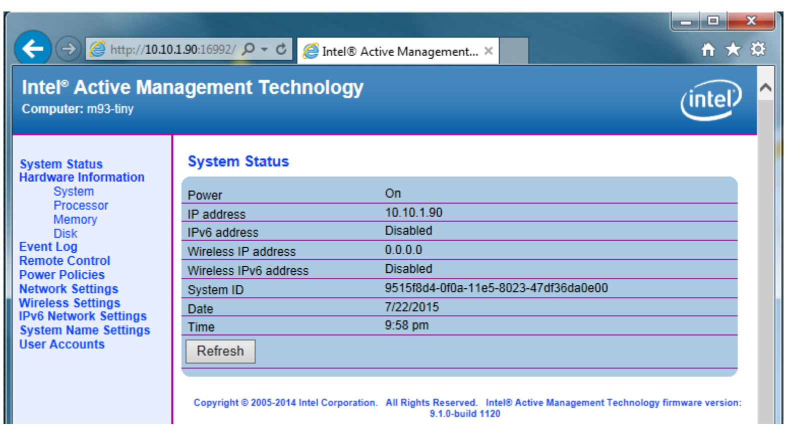

Figure 4.1: AMT System Status

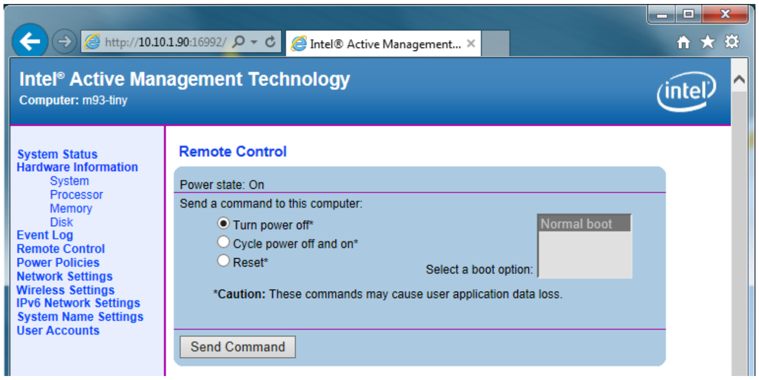

Figure 4.2: AMT Remote Control

We also enabled the TPM module in the BIOS and activated BitLocker in Windows for the C-Drive, waiting for the drive encryption to be completed before conducting power measurements, as shown in Figure 4.3.

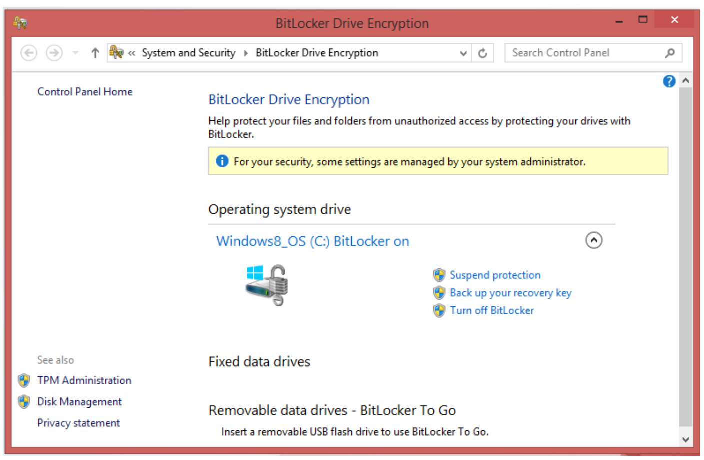

Figure 4.3: BitLocker Drive Encryption

We also verified the correct configuration of TPM using the Windows TPM Management utility, as shown in Figure 4.

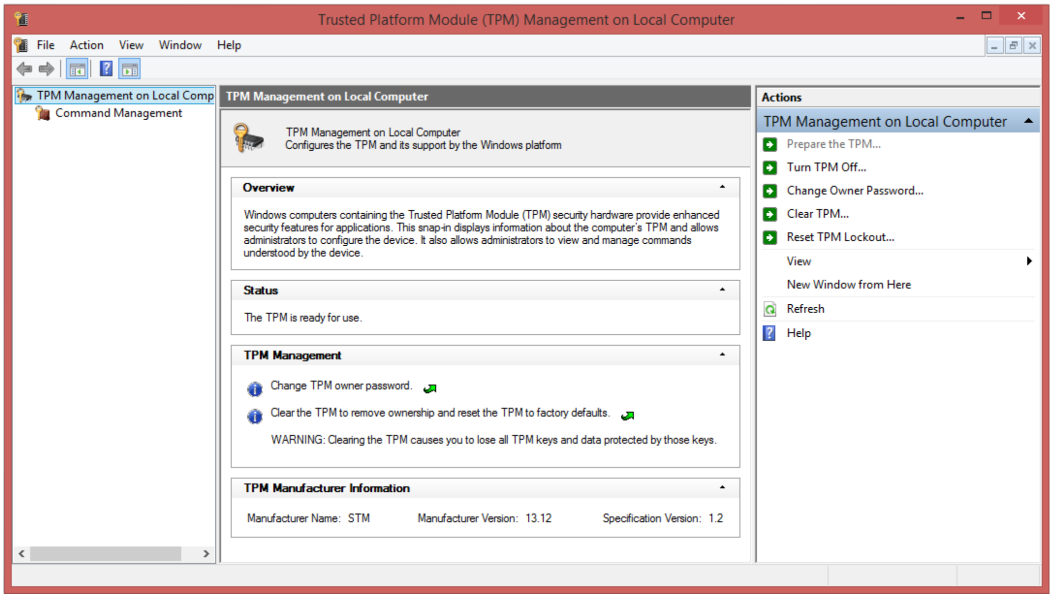

Figure 4.4: TPM Management Configuration

5 Measurement Results

The following are the initial results we have obtained when measuring the power draw of the business PCs tested.

Baseline Power Consumption, without AMT/TPM enabled and configured (no additional idle power optimizations implemented):

- ASUS Q87M with i7 4770/4785T:
  - 21.3W in Short-Idle
  - 20.9W in Long-Idle (15.5W when using a 1TB WD Green Drive)
- Full Size Business Desktop:
  - 27.2W in Short-Idle
  - 26.2W in Long-Idle
- Small Form Factor Business Desktop (without USB-based DVD dock):
  - 10.8W in Short-Idle
  - 9.8W in Long-Idle

When conducting the same power measurements after configuring AMT remote management we obtained the same measurement results in the idle states. We did however observe a slight increase in power consumption (1.05W vs 0.65W) in the OFF-state, when AMT was enabled and configured, compared to when AMT was disabled in the BIOS.

Also, when enabling the TPM module as well as activating BitLocker hard drive encryption, we could not observe any increase in idle state power consumptions of any of the 3 computers tested.

We are continuing the security/manageability project presented above by analyzing additional vPro features and further optimizing the idle power of the complete devices.

6 Conclusions and Questions for the Industry

Our observations on the selected desktops so far do not indicate that the security features mentioned by ITI increase the power requirements of desktop PCs in the short-idle and long-idle modes.

Based on our understanding of the TPM technology, we would not expect any noticeable increase in power draw due to encryption features offered by a TPM module, considering that these encryption features would only be triggered during active encryption tasks. In addition, given its relatively small size, the additional circuitry required by TPM should not require significant levels of idle power.

Similarly, the presence of remote management functionality should only have minimal impacts on power draw in the idle modes, considering that even without AMT an active network connection is already accounted for in the short-idle and long-idle states. Even if AMT does require a dedicated management and control processor to be running at all times, such small processors typically have very low power requirements (below 100mW).

Our preliminary observations presented above seem to be in contrast with the position taken by ITI in their response to the CEC Staff Report (2015). It would be helpful if ITI would additionally quantify the impact of security/manageability features on desktop idle power draw and provide information on the devices tested and the measurement setup and procedures followed.