

DOCKETED

Docket Number:	16-AAER-02
Project Title:	Appliance Efficiency Rulemaking for Computers, Computer Monitors, and Signage Displays
TN #:	214159
Document Title:	California Investor Owned Utilities Comments: California Investor Owned Utilities: CA IOU Response Computers
Description:	N/A
Filer:	System
Organization:	California Investor Owned Utilities
Submitter Role:	Public
Submission Date:	10/24/2016 3:17:32 PM
Docketed Date:	10/24/2016

Comment Received From: California Investor Owned Utilities

Submitted On: 10/24/2016

Docket Number: 16-AAER-02

California Investor Owned Utilities: CA IOU Response Computers

Additional submitted attachment is included below.

Computers

Codes and Standards Enhancement (CASE) Initiative
For PY 2016: Title 20 Standards Development

Response to CEC 45-day Express Terms for
Computers

Docket #16-AAER-2

October 24, 2016

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

Nate Dewart, ENERGY SOLUTIONS

This report was prepared by the California Statewide Investor-Owned Utilities Codes and Standards Program and funded by the California utility customers under the auspices of the California Public Utilities Commission.
Copyright 2016 Pacific Gas and Electric Company, Southern California Edison, Southern California Gas Company, and San Diego Gas & Electric.

All rights reserved, except that this document may be used, copied, and distributed without modification.
Neither PG&E, SCE, SoCalGas and SDG&E nor any of its employees makes any warranty, express or implied; or assumes any legal liability or responsibility for the accuracy, completeness or usefulness of any data, information, method, product, policy or process disclosed in this document; or represents that its use will not infringe any privately-owned rights including, but not limited to, patents, trademarks or copyrights.

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 seeks to address energy efficiency opportunities through development of new and updated Title 20 standards. This document provides recommendations and supporting analysis in response to the CEC’s Computers 45-day Express Terms.

The CASE team generally supports the energy efficiency standards for computers proposed by the CEC. We support CEC’s cost-effective and feasible base allowances and expandability adder levels for desktops and thin clients, discrete graphics adder levels, as well as the power supply and power management requirements for workstations and small-scale servers. The CASE team also recommends important modifications with supporting analysis in order to ensure the forecasted energy savings are achieved. The IOUs’ highest priority comments address language that could create loopholes or new elements of the standard that may not be sufficiently stringent:

Highest Priority Comments Summary

Issue	Recommended Changes
1) GPU definitions and exemption criteria need to be tightened to eliminate potential loopholes.	<ul style="list-style-type: none"> • Clarify that this exemption is for discrete graphics only and increase high-expandability computer discrete GPU thresholds from 400 and 600 GB/s to 800 and 1,200 GB/s, respectively. • Tighten definitions for discrete and integrated graphics as well as frame buffer bandwidth to ensure that next-generation integrated graphics cannot erroneously be counted as discrete graphics.
2) Expandability score calculation method and certain interface score require clarification to avoid exaggerated expandability scores and to accurately reflect power supply sizing requirements.	<ul style="list-style-type: none"> • Expandability score calculation requires more explicit wording to ensure correct summing of interfaces. • Since Thunderbolt 3.0 implements USB Power Delivery, its expandability score should be updated to reflect the multi-level approach already taken for USB Power Delivery. • PCI Express-related interfaces should be described clearly in terms of the number of <i>supported</i> PCIe lanes (wired lanes), not just the physical size of the PCIe slot.
3) Definitions related to several TEC adders in Table V-8 require clarification	<ul style="list-style-type: none"> • CEC should clarify that an individual piece of computer hardware (e.g. system memory) can only receive <i>one</i> TEC adder (e.g. system memory adder or high-bandwidth system memory adder, but not both). • The calculation method for system memory bandwidth should be clearly defined, and its scope should be limited to system memory (RAM) and not higher level CPU cache memory or graphics-specific memory. • The add-in card definition should explicitly exclude any add-in card hardware that already receives other TEC adders (e.g. discrete graphics cards, video surveillance cards, etc.).
4) Definitions for workstation to avoid loopholes.	<ul style="list-style-type: none"> • The workstation definition should be modified to ensure that only workstations receives the workstation exemption.
5) Mode weightings	<ul style="list-style-type: none"> • The mode weighting should not include “remoke wake,” as ITI has proposed on the October 10th public hearing.

6) Workstations and high-expandability computers testing and reporting	<ul style="list-style-type: none"> • Workstations and high-expandability computers should be required to be tested and report TEC to ensure CEC can assess a growing portion of statewide energy use, and to remove the additional incentive to receive the exemption.
--	---

2 Clarify and Strengthen GPU Language

2.1 High expandability computer exemption criteria

2.1.1 Summary of recommendation: the high expandability language should read as follows:

(2) If the computer is manufactured before January 1, 2020, a power supply of 600 watts or greater and a discrete ~~or integrated~~ graphics with a frame buffer bandwidth of ~~400~~ 800 gigabytes per second (GB/s) or greater; or

(3) If the computer is manufactured on or after January 1, 2020, a power supply of 600 watts or greater and a discrete ~~or integrated~~ graphics with frame buffer bandwidth of ~~600~~ 1,200 gigabytes per second (GB/s) or greater.

2.1.2 Background

First and most importantly, we recommend that CEC clarify that its GPU criterion apply only to discrete GPUs. The current exemption language applies to integrated GPUs, but this graphic type does warrant a description of “high expandability”; in discussion with industry during the pre-rulemaking, it was agreed upon that this exemption would only apply to discrete graphics. The intent was to recognize the need for additional power draw for computers with the expanded functionality of a discrete graphics card beyond integrated graphics. Moreover, the IOUs do not support ITI’s proposal, presented at the October 10, 2016 hearing,¹ to include separate system memory bandwidth limits for integrated graphics products as part of this exemption. There is currently no data to support the inclusion of integrated GPUs in this definition or the levels that are being proposed for exemption.

To the contrary, industry’s current efforts and projections support no exemption, such as the Gen-Z consortium, which will provide the means to include high bandwidth memory solutions into computers where “memory-centric architectures minimize data movement, reduce power consumption, reduce latency, and increase data access parallelism.”² This translates into, as an example, low-cost HBM at 200 GB/s per stack³ with a likely minimum of 2 stacks, being

¹ http://docketpublic.energy.ca.gov/PublicDocuments/16-AAER-02/TN213959_20161011T090753_ITITechNet_Oral_Comments_on_the_Rulemaking.pdf

²

³ <http://wccftech.com/sk-hynix-samsung-micron-hbm-hmc-ddr5-hot-chips/>

mainstream products that will easily meet the thresholds for exemption of 400 GB/s by, if not before 2019. Additionally, given HBM's inclusion in discrete graphics (e.g., the Radeon R9 Fury X and Titan X Maxwell), resulting in increases in FBB and decreases in power draw (an inverse relationship between idle power draw and performance), the integrated graphics trends will follow suit, at a minimum. The additional current technological advantages compared with discrete graphics (e.g., multiplane overlay (MPO),⁴ duty cycle control⁵ and panel self-refresh⁶) additionally justify no exemption for integrated graphics.

Second, and also as important, the current definition of "High expandability computer" establishes discrete GPU thresholds of 400GB/s and 600GB/s by 2019 and 2020, respectively, as criteria for system exemption, but based on the available data and the industry reported trends, the GPU thresholds should be revised upward for several reasons. First, recent third-party measurements⁷ indicate that today's highest bandwidth GPUs can already easily be accommodated by CEC's proposed GPU adders. Figure 1 illustrates a series of very recent third-party discrete GPU idle measurements, plotted as incremental TEC in kWh/yr (assuming 75% power supply efficiency and the conventional idle duty cycle of 4,380 hours per year). All but one card, a dual-GPU model, fall well below the Tier 1 adder, and one card could already achieve Tier 2 levels. The latest generation of high-bandwidth discrete GPUs are extremely efficient on a TEC-per-bandwidth basis, with 512 GB/s cards idling at the same power levels as previous generation G1 and G2 cards.

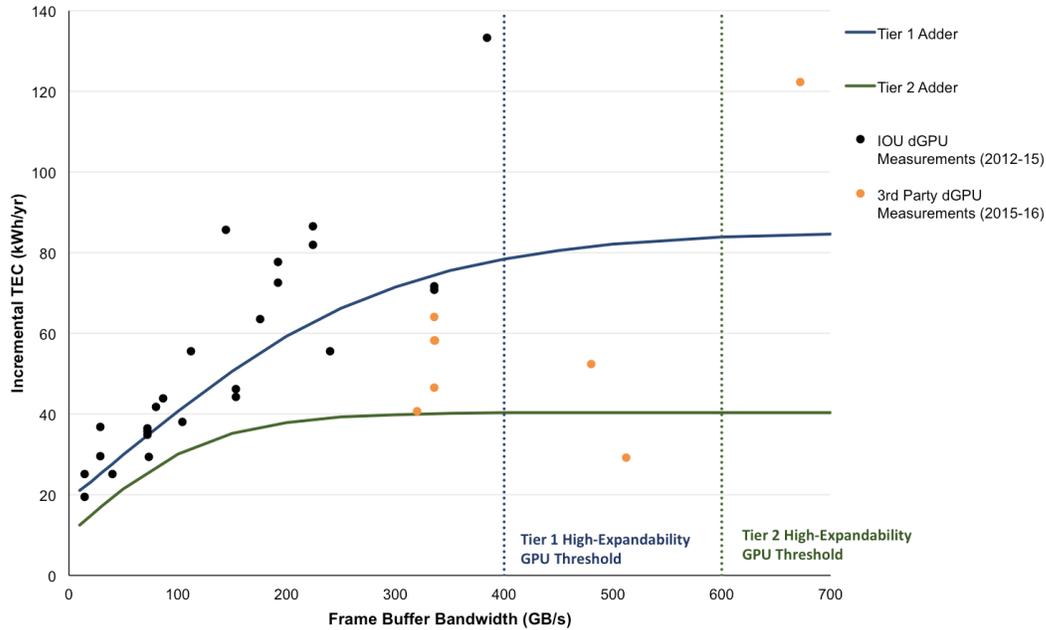
Figure 1: Desktop dGPU Dataset, Adders and FBB Exemption Thresholds

⁴ <http://www.anandtech.com/show/9582/intel-skylake-mobile-desktop-launch-architecture-analysis/6>

⁵ <http://www.anandtech.com/show/8355/intel-broadwell-architecture-preview/3>

⁶ http://www.phoronix.com/scan.php?page=news_item&px=Intel-DRM-Next-PSR-FBC

⁷ <http://media.bestofmicro.com/J/R/506151/original/31-Overview-Idle.png>,
<http://www.tomshardware.com/reviews/amd-radeon-r9-fury-x,4196-7.html>



Secondly, a step change in GPU frame buffer bandwidth (FBB) is rapidly approaching due to the advent of GDDR5x and high bandwidth memory (HBM) as well as growing interest in GPU-intensive applications like virtual reality (VR). CEC’s GPU thresholds may be outdated by the time of the standard as a result. As VR undergoes dramatic growth prior to the standard’s adoption — with annual VR headset shipments expected to eclipse 100 million units by 2020⁸ — GPU vendors and memory suppliers will prepare to push increasingly high-bandwidth products Samsung is already producing second-generation HBM products (HBM2), enabling discrete GPUs with 1,024 GB/s of FBB. The company plans to double achievable memory densities in the next year, increasing potential FBBs to 2,048 GB/s.⁹ **Recent memory forecasts released by Samsung estimate that the mainstream GPU market will support FBBs in the range of 725 GB/s by the time CEC’s standard would go into effect in 2019.**¹⁰ This is 80% above the proposed 400 GB/s threshold. The high-end GPU market would support approximately 1,000 GB/s, which is 150% above the proposed threshold of the Tier 1 exemption.

Regarding costs, although the first generation of GPUs integrating higher bandwidth memory will come at a price premium, they will quickly drop in price and become tomorrow’s mainstream cards. As a recent example of such price declines, NVIDIA’s Titan X discrete GPU at 336 GB/s was released in 2015 at prices above \$1,000 and was the highest bandwidth and performance card measured by the CASE team. However, only a year later, NVIDIA released its GTX 1080 GPU and \$599 and with dramatically better performance *and* efficiency than the Titan X. The related GTX

⁸ http://thefarm51.com/ripress/VR_market_report_2015_The_Farm51.pdf

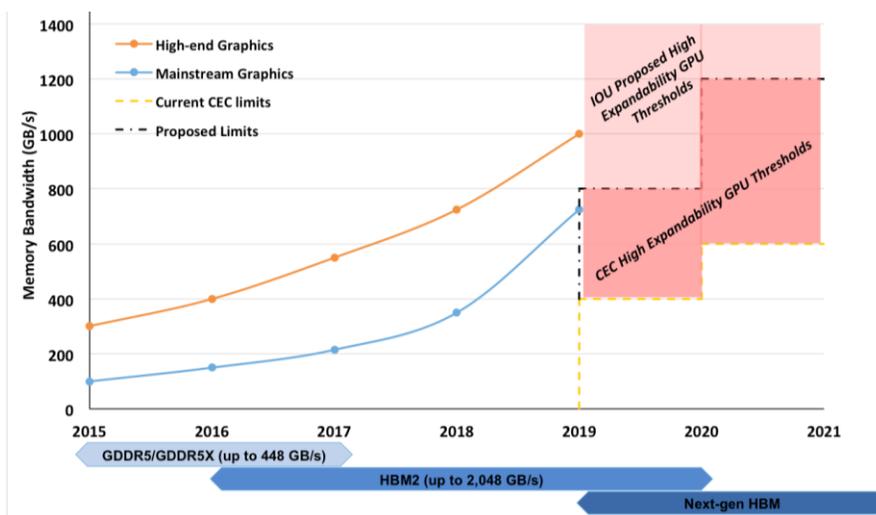
⁹ <http://arstechnica.co.uk/gadgets/2016/01/graphics-cards-with-1024gbs-bandwidth-samsung-begins-hbm2-production/>

¹⁰ <http://hexus.net/tech/news/ram/85751-samsung-start-hbm-mass-production-early-next-year/>
<http://allnewsgood.com/hi-tech-innovation-news/sk-hynix-samsung-micron-and-multi-layered-memory-the-plans-and-specifications/>

1070 GPU even outperforms the Titan X, but at a price point of \$379.¹¹ As a result, other recent best-in-class GPUs have dropped dramatically in price. The GTX 980, a benchmark for VR-capable GPUs originally purchased by the CASE team at over \$550 in 2015, can now be purchased at under \$325 only a year later. Based on this information, even in the most conservative scenario of GB/s remaining flat (far different than the industry forecasts described above) there is evidence that prices will have dropped significantly and close to today’s “mainstream” prices. For example, the AMD Radeon Fury X has a FBB of 512 GB/s which is very close to the Tier 2 threshold today and is priced at \$389. Based on IOUs’ tracking of products tested in 2010-2012, where prices dropped on average 15% per year, the price would be ~\$230 by 2020, slightly above the upper range of “mid-range, mainstream” cards and lower than the range of the “enthusiast” cards, and significantly below the range for the “high-end” cards, according to Priceonomics.com.¹²

For the reasons above, we recommend that CEC double its current GPU thresholds from 400 to 800 GB/s (until January 2020) and from 600 to 1,200 GB/s (after January 2020). We have illustrated our proposal in Figure 2 with Samsung’s memory market forecast plotted alongside. The proposed thresholds are still expected to exclude the high-end GPU market segment, allowing the latest, highest bandwidth cards to clear the GPU requirements for the high expandability computer exemption.

Figure 2: CEC and IOU High Expandability GPU Thresholds and Frame Buffer Bandwidth Forecast



2.2 Discrete and integrated graphics definitions

The “discrete graphics” or “discrete graphics GPU” definition is extremely important to the standard, as it defines a major hardware component that can garner significant TEC adders and an exemption to TEC requirements. CEC currently defines a discrete GPU based on the presence of

¹¹ <http://www.theverge.com/circuitbreaker/2016/5/7/11615806/nvidia-gtx-1080-1070-pascal-specs-price-release-date>

¹² <https://priceonomics.com/graphics-cards/>

two key components: a local memory controller and dedicated graphics memory. This definition is sufficiently vague that it could allow future integrated GPUs (for example, AMD's Zen processors¹³) with these attributes to be treated as discrete GPUs and earn the associated adder, weakening the standard for mainstream computers.

Until the release of the current express terms, discrete GPUs were understood to be an add-in card that could be completely separated from the motherboard (in desktops, particularly). Indeed, all of the IOU discrete graphics data submitted to CEC to inform discrete GPU adder levels were derived from graphics add-in cards. As a result, we suggest the following modifications to the discrete GPU definition:

“Discrete Graphics” or “Discrete Graphics GPU” means a graphics processing unit (GPU) a discrete hardware component containing one or more graphics processing units (GPUs) with a local memory controller interface and local graphics-specific memory. Discrete GPUs are not packaged on the same die or substrate as the CPU.

A small and related adjustment to the integrated graphics definition is required for clarity:

“Integrated graphics” or “Integrated Graphics Processing Unit (GPU)” means a graphics solution GPU that does not contain is not a discrete graphics GPU.

Additionally, the current language for GPU is confusing, because it states that a GPU is separate from the CPU, then later suggests that it can be integrated with the CPU. Our simplified definition removes this confusion:

“Graphics processing unit (GPU)” means an integrated circuit, separate from the CPU, a computer component that is designed to accelerate the rendering of two-dimensional or three-dimensional content to displays. A GPU may be either integrated with the CPU or discrete.

2.3 Frame buffer bandwidth definition and calculation

The frame buffer bandwidth concept exists within the standard to drive several important adders or exemptions related to discrete graphics; however, current express terms need to be more clear on how frame buffer bandwidth is defined and calculated, for clarity and to prevent double-counting of integrated and discrete GPU bandwidth. We recommend the following definition changes:

“Frame buffer bandwidth” means the rate at which data can be read from or stored exchanged between a GPU’s local memory controller and graphics-specific memory within discrete, integrated, or hybrid graphics, expressed in gigabytes per second (GB/s). It is calculated by: (graphics memory data width [bits] x graphics memory data rate [MHz]) / (8 x 1000).

¹³ <http://wccfttech.com/xbox-one-may-be-getting-a-new-apu-based-on-amds-polaris-architecture/>

3 Clarify Language for Expandability Score and Calculations

3.1 Expandability Score Calculation Instructions

CEC's current expandability score calculation methodology could be misconstrued in a way that would dramatically inflate expandability score values and result in large numbers of systems receiving higher base TEC allowances, or worse, being exempted altogether from TEC requirements. The wording currently reads as follows: "Sum the product of each interface score as determined by Table V-1 multiplied by the number of such interfaces present in the system as sold or offered for sale." As currently worded this could be interpreted as: "sum all scores and then multiply by the total number of interfaces."

Below, we provide an example for a system with the following attributes:

- 4 of - USB 3.0 or 3.1 Gen 1 (score each: 10)
- 2 of USB 3.1 Gen 2 (score each: 15)
- 4 of Unconnected USB 2.0 motherboard header (score each: 10)
- 1 of PCIe x16 or higher (only count mechanical slots) (score each: 75)
- 1 of Thunderbolt 3.0 or greater (score each: 100)

If misinterpreted, the result could be:

$$(10+15+10+75+100) \times (4+2+4+1+1) = 210 \times 12 = \mathbf{2,520}$$

We believe CEC's intended calculation would result in the following score, almost seven times smaller:

$$(4 \times 10) + (2 \times 15) + (4 \times 10) + (1 \times 75) + (1 \times 100) = 40 + 30 + 40 + 75 + 100 = \mathbf{385}$$

To clarify the calculation and prevent such confusion, we suggest the following wording:

1. ~~Sum the product of~~ Identify the score for each individual interface type score as determined by Table V-1 and then multiply by the total number of occurrences of that particular interface type multiplied by the number of such interfaces present in the system as sold or offered for sale. Finally, sum the subtotals for all interface types.

Additionally, current language instructs that "Each individual interface may only receive one score." This suggests that a product could only receive a single 10 W if it had 5 USB 3.1 Gen 2 ports, whereas we believe CEC's intent would be to allot 10 W to *each* USB 3.1 Gen 2 port, for a total score of 50 W. We suggest the following modification for clarity:

Each instance of an individual interface may only receive one score.

3.2 Thunderbolt 3.0 Score

The current Thunderbolt 3.0 expandability score of 100 W may overestimate power supply sizing requirements, which would inflate expandability scores and cause category creep. Thunderbolt 3.0 implements USB Power Delivery to negotiate power delivery to peripherals, and there is no guarantee that all Thunderbolt controllers will provide power to peripherals up to the maximum

100W (it depends on the USB-PD power profile supported by the controller).¹⁴ A controller may be able to *receive* 100 W of power — for example, in a notebook for charging purposes — but only *deliver* 15 or 30 W of power to Thunderbolt peripherals like external hard drives.

The IOUs recommend striking the Thunderbolt 3.0 row of the table entirely, allowing manufacturers to elect the USB-PD level that best fits their implementation of Thunderbolt. Thunderbolt 2 and below would still retain their own 20W expandability score.

3.3 PCIe Lane Notation

CEC has adopted the “x16” shorthand to refer to PCIe slots supporting 16 PCIe lanes. However, a PCIe x16 port may only be *wired* for 8 PCIe lanes, even though it could physically be connected to 16 lanes. This provides far less expandability. CEC should clarify language in Table V-8 as follows:

PCIe port with 16 or more PCIe lanes fully supported ~~x16 or higher (only count mechanical slots)~~

4 Improve Other Key Definitions Related to Adders and Exemptions

4.1 “Full capability” mode weighting and sleep mode

The IOUs are opposed to ITI’s proposal from the October 10, 2016 45-day language hearing to alter the requirements of section 1604(v)(5)(B)(2) from “full capability” to “remote wake”. CEC’s current language requires that systems provide several key network services while in sleep in order to use the more forgiving full capability duty cycle weighting. ITI has requested that those requirements be limited to “remote wake” capability, because remote wake functionality is supported by all OS and hardware suppliers.¹⁵ ITI’s proposal would, in effect, allow all computers to use a less stringent duty cycle. To our knowledge, this is not in the spirit of CEC’s mode weighting proposal and would weaken TEC requirements for the entire standard by about 16% (the proposed remote wake duty cycle contains 16% fewer idle hours than the conventional duty cycle, effectively giving systems a 16% credit).

We are not opposed to CEC’s current allowance of full capability mode weighting for systems that meet the provided criteria. However, we caution against allowing systems with “alternative sleep mode” to participate in this mode weighting option, which could significantly weaken Tier 1 of the standard.

The full capability mode weighting concept evolved out of the ENERGY STAR v6.0 specification development process in the 2010 timeframe at a time when sleep mode in computers was

¹⁴ <http://www.intel.com/content/www/us/en/io/thunderbolt/thunderbolt-overview-brief.html>

¹⁵ http://docketpublic.energy.ca.gov/PublicDocuments/16-AAER-02/TN213959_20161011T090753_ITITechNet_Oral_Comments_on_the_Rulemaking.pdf

synonymous with the ACPI S3 state. The alternative mode weighting was used as an incentive to provide robust network connectivity while still maintaining ACPI S3's low power levels.

Today, implementation of sleep in computers is shifting through the use of alternative operating systems (Chrome, Android, etc.) and Microsoft's adoption of "Modern Standby" in Windows 10. These sleep implementations offer a level of functionality that is a blend between today's ACPI S3 sleep and long idle states and were *expressly designed to preserve* network connectivity in a low power state. These sleep alternatives *do not require additional incentives* to maintain network connectivity, especially since CEC has already accommodated them through a broader sleep mode definition and alternative sleep mode power limits.

The IOUs encourage CEC to remove all references to alternative sleep mode in this particular provision:

In order to use the "full capability" mode weighting a computer shall have the following features enabled as shipped:

- i. Maintain Ethernet (IEEE 802.3-2015) or wireless (IEEE 802.11-2012) network addresses and network connection capability while in ACPI System Level S3 Sleep Mode ~~or an alternative to ACPI S3 sleep mode~~;
 - ii. Resume from ACPI System Level S3 Sleep Mode ~~or an alternative to ACPI S3 sleep mode~~ upon request from outside the local network; and
- 4.2 iii. Support advertising host services and network name while in ACPI System Level S3 Sleep Mode ~~or an alternative to ACPI S3 sleep mode~~. One TEC adder per hardware element

As currently written, the express terms may allow a single hardware element to receive multiple TEC adders in Table V-8. Table V-8 should state that any individual hardware device may receive *only one* TEC adder.

4.3 System memory bandwidth

The current system memory bandwidth definition does not provide sufficient clarity on how system memory bandwidth should be calculated for the purposes of the associated adder. The term is relatively new in computer efficiency policy and has no precedent (e.g. an existing ENERGY STAR definition), so specificity is important. We suggest a definition that clarifies the technical underpinnings of system memory bandwidth and provides clear instructions on how it is to be calculated:

"System memory bandwidth" means the rate at which data can be read from or stored into the of data transfer between an integrated GPU's memory controller and any on-package, graphics-specific memory, expressed in gigabytes per second (GB/s). It is calculated by: (memory data width [bits] x memory data rate [MHz]) / (8 x 1000).

4.4 Add-in card

CEC's current add-in card definition makes several exclusions for devices that either are not connected through an internal expansion slot or that reside outside the chassis of the computer (hard drives, external peripherals, etc.). However, it does not exclude several add-in card products that already receive separate TEC adders. These include discrete GPU add-in cards, video surveillance cards, and wired Ethernet/fiber cards with a transmit rate of 10 Gb/s or greater. In addition to clarifying that hardware devices may only qualify for one TEC adder (see further comments on the stringency of new standard elements), we recommend revising the add-in card definition as follows:

“Add-in card” means a removable, internal hardware device that can be installed in a computer peripheral component interconnect (PCI) or other ~~slot~~ physical port. Add-in card does not include hard disks, system memory, ~~or~~ removable devices that are intended to operate outside of a computer chassis, or components that are explicitly listed elsewhere in Table V-8. It also does not include cards that split, physically extend, or convert a slot type.

5 Other Comments and Clarifications – Existing Definitions

5.1 Workstation

The IOUs request that the CEC strengthen its workstation definition to ensure that only workstations receives the workstation exemption. The CA IOUs, NRDC, and ITI docketed a joint proposal on the workstation definition in October 2015¹⁶. This definition was better aligned with ENERGY STAR and required that workstation products meet at least 3 of the 5 listed “additional criteria.” CEC staff maintained this basic definition in the *Revised Analysis of Computers, Computer Monitors, and Signage Displays*, but dropped the criteria involving “five or more logical expansion ports” and reduced the requirement to meeting 2 of the 4 remaining additional criteria.

The IOUs recommend that CEC tighten its definition by either: 1) reinstating language equivalent to the agreed-upon definition from October 2015 or 2) revising the current Criteria 3B as suggested below. If this is not adjusted, most products (i.e. all desktops and even notebooks) would meet the PCIe requirement.¹⁷

~~Supports~~ Includes four or more lanes of PCI-Express (PCIe) slots, ~~other than discrete graphics, connected to accessory expansion slots or ports~~ one of which must support at

¹⁶ http://docketpublic.energy.ca.gov/PublicDocuments/14-AAER-02/TN206287_20151006T100251_California_Investor_Owned_Utilities_Comments_CEC_Title_20_Compu.pdf

¹⁷ Example: <http://www.gigabyte.com/products/product-page.aspx?pid=5806#sp> - 1 x PCI Express x16 slot, running at x16 (PCIEX16), 1 x PCI Express x16 slot, running at x8 (PCIEX8) (shared bandwidth with the x16 slot so reduces the x16 slot to 8 lanes when populated), 1 x PCI Express x16 slot, running at x4 (PCIEX4) (disabled if SSD connected to M2H_32G connector,) 3 x PCI Express x1 slots and 2 x M.2 connectors (PCIe x4/x2/x1 SSD support) – for a total of 23 PCIe lanes (Gen 3))

least 16 PCIe lanes, where each lane has a bandwidth of 8 ~~gigabytes~~ gigabytes per second (Gb/s) or more.

Additionally, the IOUs recommend that for Criteria 2, CEC change the language from “supports” to “has installed.” Error Correcting Code (ECC) is the primary criteria that determines whether a system is considered workstation. There is now evidence that some desktops can already support ECC¹⁸ and that HBM2 supports ECC.¹⁹

Finally, an additional requirement could be for the products to be marketed as “workstations” as well. This should not preclude the other changes, but it could provide some additional differentiation from traditional desktops.

5.2 Limited Capability Operating System

We believe CEC’s intent is that a limited capability operating system cannot provide *any* of the four features listed in its definition. This is an important distinction, since under the current definition, one might interpret an operating system without, for example, support for multiple user profiles as a limited capability operating system. We suggest:

“Limited capability operating system” means an operating system that performs basic operations and that does not meet any of the following:

- (1) Have automatic power management features;
- (2) Support USB devices;
- (3) Have Graphical User Interface (GUI); or
- (4) Support multiple user profiles or distinguish between users.

5.3 Switchable graphics

We support CEC’s decision to recognize switchable graphics capability but have two suggested corrections to the definition. First, hybrid graphics may be a confusing term, as manufacturers often market graphics solutions that simply share memory between discrete and integrated GPUs as hybrid graphics. Switchable graphics is a clearer term that more directly speaks to the technology’s power management capabilities. Second, we would like to ensure that the definition recognizes that the control, power management, and coordination of the discrete and integrated GPUs is an automated process (i.e. it cannot be accomplished manually through configuration changes or swapping monitor connections). Our revised definition is provided below:

~~“Hybrid Switchable graphics” means a functionality that allows automatically places the system’s primary or master Discrete Graphics to enter in a low-power state when not required in favor of Integrated Graphics. The functionality also supports automatic waking of a primary or master discrete GPU from a low-power state. This functionality allows graphics rendering by lower power and lower capability~~

¹⁸ <http://www.amd.com/Documents/merlin-falcon-product-brief.pdf> and <http://www.amd.com/en-gb/products/embedded/processors/g-series>).

¹⁹ <http://monitorinsider.com/HBM.html>

~~integrated GPUs while on battery or when the output graphics are not overly complex while then allowing the more power consumptive but more capable discrete GPU to provide rendering capability when the system requires it.~~

5.4 Off mode

CEC's current definition does not clarify that the off mode referenced on page 6 of the express terms applies only to computers. As the regulatory language is intermingled with computer monitors, televisions, and consumer audio and video equipment. We recommend that this be titled "Computer Off Mode."

6 Other Comments and Clarifications – New Definitions

We recommend several new definitions in order to ensure that exemptions and adders are properly allocated.

6.1 PCI Express Lane

PCI Express (PCIe) lanes are used to describe interfaces eligible for the expandability score in Table V-8 as well as in other key definitions in the standard. For clarity, we recommend adding the following definition:

"PCI Express Lane" refers to a set of wire pairs, one pair for transmission and one for reception, that are used to send and receive data between a computer system and a peripheral hardware device according to the PCI Express standard.

6.2 PCI Express Port

PCI Express (PCIe) ports are used to describe interfaces eligible for the expandability score in Table V-8 as well as in other key definitions in the standard. For clarity, we recommend adding the following definition of PCIe ports:

"PCI Express Port" means a connector on the motherboard that allows for the installation of PCI Express hardware devices and communication with the computer system over PCI Express Lanes. PCI Express Ports must conform to the form factor specifications of the PCI Express standard.

6.3 Integrated Liquid Cooling

As integrated liquid cooling may count toward a system's expandability adder score, it needs to be defined. We suggest:

"Integrated liquid cooling" means a closed system that cools electronic components, such as a CPU or GPU, using heat transfer liquid that is mechanically pumped.

6.4 USB Motherboard Header

USB motherboard headers currently may count toward the expandability score, but have not been clearly defined. We suggest the following definition for clarity:

“USB Motherboard header” means a physical connector located on a computer’s motherboard that provides data and power connections for 2 or more downstream USB ports.

6.5 Integrated Displays Effectiveness

Currently, the effective date for integrated display EPD adders is listed as July 1, 2019. This date should be corrected to align with the effectiveness date of the computer standard (January 1, 2019) rather than the computer monitors standard. If this is not corrected, integrated EPDs will not receive any adder for the first 6 months of the computer standard.

7 Test Method Corrections and Clarifications

7.1 Including of TEC requirements for workstations and high expandability computers in Section 1604(v)(5)

The IOUs highly recommend that the language be clarified to ensure that workstation and high-expandability computer TEC measurements are tested and reported. While these products are exempted from the TEC requirement, it is important that CEC have the data to assess a growing portion of statewide energy use. The proposed language is as follows:

(5) The test method for computers is the ENERGY STAR Program Requirements for Computers, Final Test Method (Rev. March-2016), with the following modifications:

...

(B) The total power consumption of a computer shall be calculated using Equation 1 in Section 3 of the ENERGY STAR Program Requirements for Computers, Eligibility Criteria Version 6.1 (Rev. March-2016).

1. Computers manufactured before July 1, 2021 shall use the “conventional” mode weighting of Table 3 for a desktop computer, a high-expandability computer, a workstation, a mobile gaming system, or a thin client, or Table 4 for a notebook computer or portable all-in-one computer, contained within Section 3 of the ENERGY STAR Program Requirements for Computers, Final Test Method (Rev. March-2016), unless they meet the criteria to use “full capability” mode weighting, below.

7.2 Monitor Connections for Computer Testing, Section 1604(v)(5)(D)

CEC’s test guidance for desktops requires testers to connect a computer monitor with 1920 x 1080 (Full HD) native resolution; however, it is possible that monitors with this native resolution will no longer be sold during the useful life of the standard. To future-proof the standard, we recommend adapting the language as follows:

(D) A computer monitor used in the testing of desktop computers shall have a native resolution of at least 1920x1080 pixels and use progressive scanning. The computer operating system shall be set to operate at a minimum of 1920x1080 pixels and progressive scanning.

CEC also provides useful guidance on the display port to which testers must connect the computer monitor. Given our recommendations on switchable graphics, we suggest altering the wording in 1604(v)(5)(D)(1) to:

1. If ~~hybrid switchable~~ graphics are available, choose the port that ~~enables hybrid graphics~~ supports this functionality.

Finally, CEC has provided a prioritized list of display ports for situations in which the system supports multiple varieties of display connectors. To future-proof the language, we suggest adding Thunderbolt and USB to the list, both of which already or soon will support display connectivity. We assume that CEC's intent is to require testers to connect the display to the connector that supports the highest resolution and bandwidth:

4. If there are multiple connector ports to choose from pursuant to subdivisions (D)(1) through (D)(3) of this section, connect the display to the port supporting the highest display resolution, or port using the first available from the port types listed below:
- i. Thunderbolt
 - ii. Display Port
 - iii. HDMI
 - iv. DVI
 - v. USB
 - vi. VGA
 - vii. Other

7.3 Automatic Brightness Control for Integrated Monitors, Section 1604(v)(5)(F)

CEC's express terms in this section instruct testers to ignore section 5.2(A)(1) of the ENERGY STAR test method, which CEC says requires testers to "not disable" automatic brightness control for integrated displays. However, section 5.2(A)(1) in the ENERGY STAR test procedure requires testers to "disable display dimming," which reads contrary to CEC's interpretation.

We request that CEC clarify its intent and suggest the following correction, assuming that CEC intends for testers not to follow ENERGY STAR's requirements in section 5.2(A)(1):

During testing, a notebook computer, mobile gaming system, portable all-in-one, or integrated desktop shall proceed using Section 5.2(A)(1) and ignore the direction ~~not~~ to disable automatic brightness control as described in Section 5.2(A) of the ENERGY STAR Program Requirements for Computers, Final Test Method (Rev. March-2016). If

automatic brightness control is supported, position a light source such that 300 lux directly enters the ABC sensor.

7.4 Sleep Mode Power Measurement, Section 1604(v)(5)(H)

CEC's requirements for sleep mode testing appear to reference the *long idle mode* test provisions from IEC 62623:2012 rather than the sleep mode provisions. We suggest the following modified language to clarify:

(H) The sleep mode power measurement shall be tested in a modified manner from the sleep mode test procedure described in IEC 62623:2012, section 5.5.3. Instead of measuring power after manually entering sleep mode, the power measurement shall begin no sooner than 30 minutes and no later than 31 minutes of user inactivity on the unit under test. This measurement shall follow the ~~long idle~~ sleep mode test without altering the unit under test.

8 Reporting Requirements

The IOUs recommend that power supply size be added to the reporting requirements in Table X, given that this is a criteria for the high expandability computer TEC exemption. While power supply model number is a good start, it does not guarantee that the power supply size will be included, and even if it is, there are often other numbers included in the model number as well which can make the sizing difficult to recognize. Requiring both would be helpful for compliance.

In addition, we suggest the following wording changes for clarity and accuracy:

- Number of CPU Cores
- ~~AC Adapter Size~~ External Power Supply (EPS) Rated Output (watts) (notebook computers)
- Total Battery Capacity (watt-hours) (notebook computers and Portable All-In-Ones only)