# ENERGY STAR* Version 5.0 System Implementation

**Whitepaper**

***Published by Intel with technical collaboration from the U.S. Environmental Protection Agency***

***February 2009***

Document Number: 321556-001

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

*"Energy Star* has been a key element in Intel IT's sustainability efforts for over 3 years as we continue to focus on using energy efficient technology. Energy Star compliance of our PC fleet results in lower operational costs and is a critical piece in delivering on our commitment to lower our overall IT carbon footprint."*

**Diane Bryant**
**Vice President and Chief Information Officer**
**Intel Corporation**

*"The latest update to the Energy Star specification is an important step toward aligning the industry on aggressive goals to reduce computer energy consumption. The Climate Savers Computing Initiative is committed to collaborating with the EPA to drive adoption of its metrics in the customer base very aggressively."*

**Lorie Wigle**
**President, Climate Savers Computing Initiative**
**Intel Corporation**

*"EPA thanks the IT industry for their 15 year partnership on the ENERGY STAR program and their commitment to energy efficiency. We look forward to more successes including with the new ENERGY STAR 5.0 specification for computers with its more holistic approach to computer energy efficiency across a broader range of products. We believe this paper from Intel will prove to be a valuable tool in broadly communicating the new ENERGY STAR requirements."*

**Kathleen B. Hogan**
**Director, Climate Protection Partnerships Division**
**U.S. Environmental Protection Agency**

§

# Contents

1 Preface ........................................................................................ 11

2 Introduction.................................................................................. 13

2.1 The ENERGY STAR* Program: An Overview ........................................... 14

2.2 How ENERGY STAR* Specifications are developed ................................. 16

2.2.1 Guiding Principles and the Specification Development Process ....... 16

2.2.2 Examples of ENERGY STAR* Product Marketing Campaigns: *Low Carbon IT* and *Change the World* ........................................................ 17

2.2.3 Verification Testing ............................................................ 18

2.3 ENERGY STAR* Version 5.0 ............................................................... 18

2.3.1 ENERGY STAR* Partnership and Submittal of Products ................ 20

2.3.2 ENERGY STAR* in the Channel Market ..................................... 20

2.4 Taking Advantage of Power Management Settings ................................. 21

2.5 Terminology ........................................................................... 22

2.6 ENERGY STAR* and Computer Platforms .............................................. 25

2.7 ENERGY STAR* Common Platform Requirements ................................... 27

2.8 Full Network Connectivity and Proxying ................................................ 27

2.9 Reference Documents ..................................................................... 28

3 Notebooks and ENERGY STAR* Requirements ................................................ 30

3.1 Notebooks, Battery Life and AC Energy ................................................ 30

3.2 ENERGY STAR* Basics for Notebooks .................................................. 30

3.3 The Notebook Platform ................................................................... 33

3.4 Power Delivery: A System's Power Amplifier ......................................... 35

3.5 ENERGY STAR* Requirements for the AC Brick ....................................... 37

3.6 Power Breakdown of the Notebook ..................................................... 40

3.7 Power sensitivity to Power Delivery Efficiencies ..................................... 41

3.8 Optimizing a Notebook for ENERGY STAR* ............................................ 44

3.9 Summary ................................................................................ 47

4 Desktop Computers and ENERGY STAR* Requirements ...................................... 48

4.1 Category Definitions ..................................................................... 48

4.2 Power Supply Considerations ........................................................... 49

4.2.1 ENERGY STAR* Requirements ................................................ 50

4.2.2 Overall Efficiency Measurements ............................................ 50

4.2.3 Power Supply Sizing ........................................................... 50

4.2.4 5 V Standby (5 VSB) Efficiency .............................................. 53

4.2.5 External Power Supplies for Desktop Systems ............................ 54

4.3 Operational Modes ....................................................................... 54

4.3.1 Idle State ........................................................................ 57

4.3.2 Sleep Mode ...................................................................... 58

4.3.3 Off Mode ......................................................................... 59

4.4 Summary ................................................................................ 60

5 Workstations and Small-Scale Servers ........................................................ 62

5.1 ENERGY STAR* Workstation and Small-Scale Server Specification Scope...... 62

5.2 Small-Scale Servers and Specification Application .................................. 63

5.3 Workstation Category and Definition .................... 64
5.4 Workstation Scalable Specification .................... 65
5.5 Measurements .................... 68
5.5.1 Max_Power .................... 68
5.5.2 Typical Energy Consumption .................... 69
5.5.3 Tools and Calibration .................... 70
5.5.4 Power Management and Settings .................... 70

6 Thin Clients and ENERGY STAR* Requirements .................... 72
6.1.1 Server-based Client Compute Architecture .................... 72
6.1.2 Category Definitions .................... 74

7 Future Specification Considerations .................... 76

8 Summary .................... 78

Appendix A Creating a Disk Image .................... 80

A.1 Modifications to the Default User/policy in Rollout Image (i.e.; "Ghosting") .................... 80

A.2 Microsoft Windows* XP .................... 81

A.3 Microsoft Windows Vista* .................... 83

Appendix B Sample TEC Calculations .................... 86

B.1 Desktop, Integrated Desktop, Notebook Computers .................... 86

B.2 Workstations .................... 87

Appendix C Energy Star* Version 5.0 and Version 4.0 Summarized .................... 88

About the Authors .................... 90


## Figures


Figure 1. Specification Development Cycle .................... 17
Figure 2: Microsoft Windows Vista* Computer and Monitor Sleep Settings .................... 22
Figure 3. Typical CAT A Notebook System (Integrated graphics, UMA Memory) .................... 34
Figure 4. Typical CAT B Notebook (discrete graphics) .................... 35
Figure 5. Notebook Power Delivery System .................... 36
Figure 6. Example AC Brick Efficiency over Different Loads .................... 38
Figure 7. Typical Notebook Idle Power Breakdown .................... 40
Figure 8. Typical Notebook Sleep and Off Power Breakdown .................... 41
Figure 9. Idle Power distribution across efficiencies .................... 42
Figure 10. Sleep Power distribution across efficiencies .................... 42
Figure 11. Idle Effects of Low Power States .................... 44
Figure 12. Percent Spundown Time Operating on Battery .................... 46
Figure 13. PSU Efficiency as Percentage of Output .................... 51
Figure 14. PSU Efficiency per Power for 450 W PSU .................... 52
Figure 15. PSU Efficiency per Power for 250 W PSU .................... 53
Figure 16. Idle State System Power State Example .................... 55
Figure 17. Sleep (S3) Mode System Power State Example .................... 56
Figure 18. Off (S5) Mode System Power State Example .................... 56
Figure 19. System Idle Power Factors .................... 58
Figure 20. Scaleable Typical Electricity Consumption .................... 66

Figure 21. Architecture - Terminal Services ........................................................ 73
Figure 22. Architecture – Virtual Hosted Desktop................................................. 74
Figure 23. Power Scheme ............................................................................ 84

## Tables

Table 1. $E_{TEC}$ Requirements – Notebooks ........................................................... 31
Table 2. Operational Mode Weighting for Conventional Notebooks.......................... 31
Table 3. Active Mode Requirements for Average Efficiency .................................... 37
Table 4. No Load Requirements ............................................................................ 38
Table 5. ENERGY STAR* Requirements for Desktops ............................................ 49
Table 6. Small-Scale Server Requirements ........................................................... 64

§

# Revision History

| Revision Number | Description | Revision Date |
|---|---|---|
| -001 | • Initial release. | February 2009 |

§

# 1 Preface

In July of 2009 a new version of the ENERGY STAR* specification for computers will take effect. This new version specifies many new system level requirements that must be met in order to carry the ENERGY STAR logo on computers. As such, Intel and the EPA have collaborated on creating this document that can provide system providers, from the largest to smallest, with key understandings on how to specify and test system configurations that can best meet the required power levels of the new specification.

§

# 2 Introduction

Over the last twenty-five years, computers have become pervasively used tools that have enhanced the productivity in the office and enhanced entertainment and utility within the home. Their remarkable growth has been fueled by amazing advancements in performance, capability and affordability. As the number of computers has grown, so has the need for delivery and deployment in increasingly energy conscious ways. More energy friendly computers can have an effect on both the available energy capacity as well as on the ecological impacts of generating additional electricity to meet growing demands.

Moving forward, there will continue to be a need for greater levels of computer performance and capability that will also be coupled with the need to manage energy consumption. Intuitively, it would seem that delivering greater performance/capability would be at odds with managing energy consumption. However, innovations by Intel and others in the industry have enabled delivery of technologies that can help offset and, in some cases, even reduce the energy consumed by the computer. These innovations have typically focused on optimizing the energy efficiency and performance when the computer is actively being used while minimizing the actual energy consumption when the computer is in a state of prolonged inactivity.

Today's computers, such as desktops and notebooks, have many power saving capabilities built into them. Examples are the "sleep" and "hibernate" modes that can significantly reduce the amount of energy consumed during inactive states. When these capabilities are turned on during periods of inactivity, it has been estimated to reduce the overall amount of energy consumed by computers by up to 50%[1].

In order to help encourage adoption and use of these energy saving technologies, in 1992 the US Environmental Protection Agency (EPA) established its voluntary program, called ENERGY STAR, to cover first computers and later other categories of office equipment and other products. The ENERGY STAR program for computers has the goal of generating awareness of energy saving capabilities, as well as differentiating the market for more energy-efficient computers and accelerating the market penetration of more energy-efficient technologies.

In the middle of 2009, the EPA will update the ENERGY STAR computer specification to Version 5.0. The new version is intended to expand on the tests and methodologies presented in Version 4.0 of the program, implemented in 2007, and expand the scope of the program to address some new product areas for the first time. EPA anticipates that new testing criteria and power/energy limits could reduce the amount of energy consumed by an average of 30%[2]. EPA routinely sets a target goal of about twenty-five percent compliance for each of the platform categories and this will also be the case for the Version 5.0 Specification.

[1] Assumes a baseline desktop configuration that consumes ~258 kW-hr/year without power management (Max = 118W (3%) + Sleep = 4W (5%) + idle = 65W (37%) + Off = 3W (55%)) and ~124 kW-hr with power management ((Max = 118W (3%) + Sleep = 4W (30%) + idle = 65W (12%) + Off = 3W (55%)).

[2] Savings Estimates for the ENERGY STAR* Voluntary Labeling Program (2007). Sanchez, Marla, Carrie Webber, Richard Brown and Gregory Homan. Climate Change Action Plan (CCAP) Model version 061121. Lawrence Berkeley National Laboratory.

The remainder of this document will discuss how the new ENERGY STAR specification applies to Desktop PCs, Notebook PCs, Workstation Computers, Small-Scale Servers, and a new product category in the program: Thin Clients. The final ENERGY STAR Version 5.0 specification will contain additional provisions for Game Consoles; these products are not within the scope of this document and are not discussed.

In particular, this document will describe the key system components that impact energy consumption in general as well as describe ways in which system designers can make choices that decrease a system's energy draw. Finally, this document will relate the impact of those choices to the requirements of the ENERGY STAR specification, as well as how to specify and test system configurations that can best meet the required power levels of the new specification.

## 2.1 The ENERGY STAR* Program: An Overview

ENERGY STAR is a voluntary program aimed at reducing green house gas emissions and is jointly administered by the US EPA and the US Department of Energy. Through the ENERGY STAR program in 2007, EPA helped Americans save 180 billion kilowatt-hours (kWh)—about 5 percent of U.S. electricity demand—prevent the emissions of 40 MMTCE of greenhouse gases, and save $16 billion on their energy bills. More than 70 percent of American households recognize the ENERGY STAR label. Other ENERGY STAR program highlights include:

**Offering More Qualified Products to More Consumers**

- More than 2,000 manufacturers are using the ENERGY STAR label on over 40,000 individual product models across 50 product categories.
- Americans purchased about 500 million ENERGY STAR qualified products in 2007, bringing the total to more than 2.5 billion since 1992.
- EPA expanded the suite of ENERGY STAR qualified products to include digital-to-analog converter boxes (DTAs), decorative light strings, and commercial dishwashers and ice machines, and completed important revisions to the specifications for residential light fixtures and roofing products.

**Raising the Bar for New Home Construction**

- Despite the downturn in the new housing market, more than 120,000 new homes were constructed to meet ENERGY STAR guidelines in 2007, which represents about 12 percent of the U.S. new housing starts in 2007 and brings the total to almost 840,000 qualified homes nationwide.

**Improving the Comfort and Efficiency of Existing Homes**

- More than 38,000 homeowners are enjoying greater savings and comfort in their homes thanks to state and locally sponsored Home Performance with ENERGY STAR programs. Seven sponsors launched new programs in 2007, bringing the total to more than 20 programs in 21 states.

**Saving Energy in the Commercial and Industrial Sectors**

- Almost 800 organizations and individuals nationwide have joined the ENERGY STAR Challenge to improve the energy efficiency of their commercial buildings by 10 percent or more.
- Use of EPA's energy performance rating system experienced tremendous growth; building owners and operators used the system to rate the performance of more than 62,000 buildings— doubling the number of buildings rated in just one year and representing more than 7.5 billion square feet of building space.
- More buildings than ever have qualified for the ENERGY STAR, over 4,000, representing more than 740 million square feet. These buildings use nearly 40 percent less energy than typical buildings.
- Almost 40 auto assembly, corn refining, cement—and for the first time, petroleum refining—facilities in the United States have earned the ENERGY STAR for superior energy performance.

The EPA strives to maximize international harmonization of test methods and efficiency requirements, seeing the benefit to the environment and our wide array of partners. Currently, EPA has agreements to harmonize regarding office equipment requirements with the European Union, Switzerland, Canada, Australia, New Zealand, Japan, and Taiwan. To account for differing voltage conditions in various markets, the program includes requirements in many specifications that products be tested and reported at the voltage/frequency combinations to which they will be marketed as ENERGY STAR qualified. The Version 5.0 computer specification includes guidance on the market applicability of three established voltage/frequency combinations for ENERGY STAR testing. Products must meet the same established requirements at all voltage/frequency combinations to which they will be sold and promoted as ENERGY STAR qualified.

In addition to international harmonization, EPA endeavors to work with other entities where program requirements come together. EPA does, where possible, harmonize with Federal Energy Management Program (FEMP) on standby requirements. Where the agencies do not align on requirements, manufacturing partners who wish to sell to the US federal government must meet both ENERGY STAR and FEMP requirements. EPA also joins forces with non government organizations who share the common goal of improving efficiency of products. For example, Climate Savers Computing Initiative and EPA's ENERGY STAR program are striving to harmonize on technical internal power supply requirements, as demonstrated by the power supply requirements described later in this document, as well as further each others' climate change marketing efforts.

EPA and the IT industry have a long history of partnering on driving down energy consumption of computers, monitors, and other products and thus also reducing greenhouse gas emissions. EPA continues to have in interest in IT product categories as they represent a sizable and growing segment of domestic and global energy use. For example:

- The US has more than 180 million computers in use; accounting for about 2% of nation's annual electricity consumption.
- There are an estimated 150 million monitors operated in the U.S. today, split roughly between homes and offices.
- The US has approx 260 million imaging products in use; accounts for 3% of nation's annual electricity consumption or more than $7.5 billion in energy.
- Servers and datacenters use about 61 billion kWh accounting for 1.5% of total U.S. consumed and costing $4.5 billion annually to power.

According to McKinsey & Company improving the efficiency of commercial and residential electronics is one of the lowest cost paths to reducing greenhouse gas emissions in a continuum of efficiency and clean energy tactics.

ENERGY STAR delivers these savings. If all computers sold in the United States meet the ENERGY STAR 4.0 requirements, the savings in energy costs will grow to about $2 billion each year and greenhouse gas emissions will be reduced by the equivalent of greenhouse gas emissions from more than 2 million vehicles. Given the additional energy savings expected from the Version 5.0 specification, the ENERGY STAR computer specification will continue to deliver on these expectations from consumers and enterprises seeking to reduce energy costs and reduce impact on the environment.

EPA also recognizes the potential for savings IT can offer through offering myriad tools to improve the efficiency of our systems, for example. As noted in the Climate Group's Smart 2020 Report, IT can be a significant part of the solution to climate change, offering a potential of 14% reducing in greenhouse gas emissions globally, or 22% in the US. EPA considers this potential as it looks for opportunities to reduce greenhouse gas emissions.

## 2.2 How ENERGY STAR* Specifications are developed

As the ENERGY STAR Program continues to grow, the ENERGY STAR specification development process had evolved with consideration of several key factors. This section discusses the guiding principles considered in identifying products for coverage in the ENERGY STAR program, the development process, and issues after finalization, including marketing efforts and compliance/verification.

### 2.2.1 Guiding Principles and the Specification Development Process

EPA considers a number of factors when identifying new product areas to address in the program or revising existing specifications: cost-effective efficiency, maintained (or enhanced) performance with respect to the overall market, significant energy savings potential for a product category, achievable efficiency improvements with non-proprietary technology, product differentiation opportunities in the market, and effectiveness of a labeling program in the market. The final decision involves a balance of these considerations, along with collaboration with contacts in industry and program stakeholders.

Below is a graphic describing the specification development process. Transparency and collaboration are integral to the process. EPA plans for multiple opportunities for stakeholder comment through release of public drafts; the agency strives to publish details on stakeholder meetings and comment opportunities publicly as well. EPA values the contribution of stakeholders involved in creation and revision of ENERGY STAR specifications and encourages interested organizations or individuals to get involved by going to the ENERGY STAR website (ENERGYSTAR.gov) and requesting program updates as they are available.

**Figure 1. Specification Development Cycle**

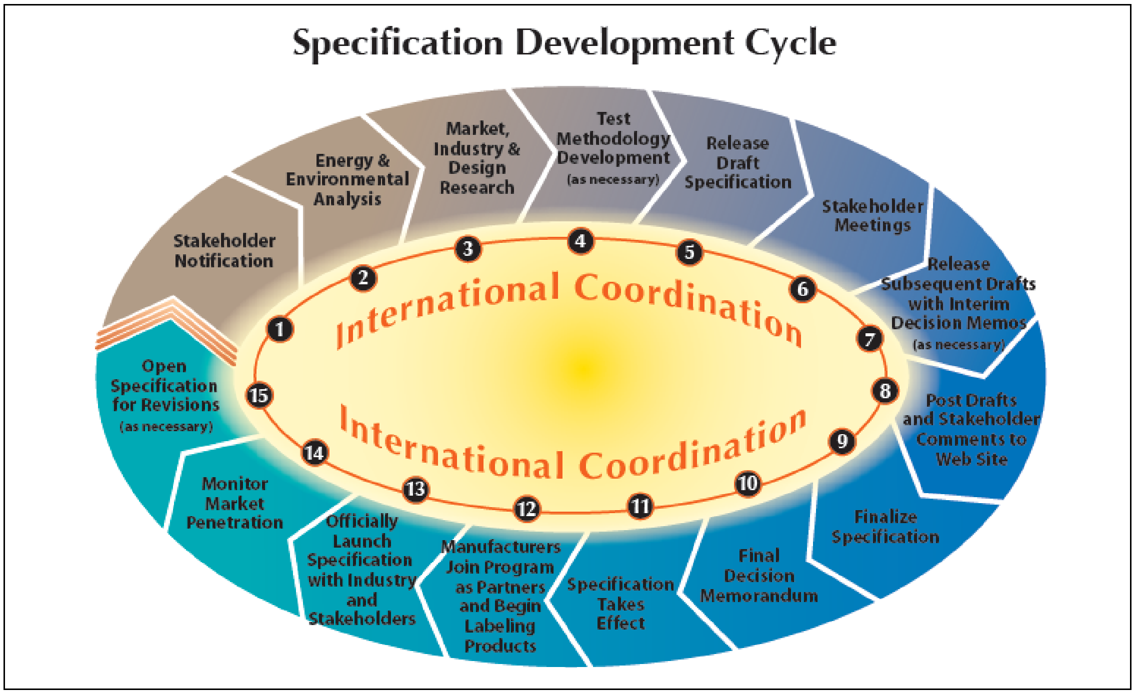


## 2.2.2 Examples of ENERGY STAR* Product Marketing Campaigns: *Low Carbon IT* and *Change the World*

The *Low Carbon IT* campaign was launched by the U.S. Environmental Protection Agency to improve enablement rates for computer power management. Given that Lawrence Berkeley National Labs estimates power management settings have been employed by only five to 10 percent of U.S. organizations, this program represents a significant energy (and energy cost) savings opportunity for interested organizations. If all office computers and monitors in the United States were set to sleep when not being used, the country could save more than 44 billion kWh or $4 billion worth of electricity and avoid the greenhouse gas emissions equivalent to those of about 5 million cars each year.

The ENERGY STAR Low Carbon IT Campaign is also partnering with the Climate Savers Computing Initiative, a non-profit group that includes technology firms, energy companies and nongovernmental organizations working to promote the use of more energy-efficient computers and increase the use of computer power management. Both efforts share similar goals and cooperate with each other on technical specifications and marketing through a strategic partnership.

Organizations that take the ENERGY STAR Low Carbon IT Campaign are asked to enable the power management, or sleep mode, on their computers and monitors. The easiest way to fulfill this agreement is to ensure that new computer purchases are ENERGY STAR qualified – all ENERGY STAR qualified computers are shipped with power management features enabled. For organizations that join this program, EPA provides free assistance to help implement power management, an estimate of the organization's energy and carbon savings, and official recognition from the agency. Organizations can join the campaign by visiting http://www.energystar.gov/lowcarbonit.

*Change the World, Start with ENERGY STAR* is an annual campaign each fall that focuses on individual consumers and provides an opportunity to take small, individual steps that make a big difference in the fight against global warming. The pledge has advice on ways to save at home and at work through ENERGY STAR, including purchasing ENERGY STAR qualified computer equipment. Organizations are encouraged to participate by becoming a pledge driver. Pledge drivers demonstrate environmental commitment by educating others about global warming and saving energy, and have access to communications materials to support efforts. Further information is available at http://energystar.gov/changetheworld.

### 2.2.3 Verification Testing

Through over 15 years of shared effort, EPA and ENERGY STAR Computer Partners have built a recognized and valued ENERGY STAR brand. The Verification Testing initiative, introduced in Version 5.0, maintains the value of this brand requires ensuring products labeled with the ENERGY STAR deliver on their promise to the consumer. Under this audit program, a limited number of qualified computers (no more than five) are selected from each manufacturer on an annual basis for testing at an independent test laboratory. This program allows EPA and its Partners to maintain the high quality and savings purchasers have come to expect when selecting ENERGY STAR. EPA maintains a procedure manual for this program on the ENERGY STAR website.

## 2.3 ENERGY STAR* Version 5.0

Version 5.0 of the ENERGY STAR specification for computers replaces Version 4.0 of the specification that has been in effect since mid-2007. Version 5.0 consists of a single tier, or phase of the program for the products covered by this document.

Version 5.0 will go into effect on July 1, 2009 and will require all systems manufactured on, or after, this date to meet the new requirements in order to ship with the ENERGY STAR logo. Similar to Version 4.0, there is no grandfathering for existing systems that meet pre-5.0 requirements to the new program. Systems will have to be retested and resubmitted, in their ”as-shipped” configuration, in order to continue to carry the logo.

The table below is intended to illustrate high-level changes from Version 4.0 to 5.0.

**Table 1: Comparison between Version 4.0 and Version 5.0**

| Version 4.0 | Version 5.0 |
|---|---|
| Scope | |
| *Eligible Products: Desktops, Integrated [Desktops], Notebooks/Tablets, Desktop-Derived Servers, Workstations.* | Thin Clients are added to the list of eligible products. |
| Power Supplies | |
| *Internal Power Supplies: ≥80% efficient at 20%, 50%, and 100% of rated load, with 0.9 Power Factor at 100%.* | Efficiency has been raised to 85% at 50% of rated load, and 82% at both 20% and 100%; the test standard has been updated to the current appropriate version. Power Factor requirements remain at the same level. |

| Version 4.0 | Version 5.0 |
|---|---|
| *External Power Supplies: must meet ENERGY STAR no-load and average active efficiency requirements for External Power Supplies.* | The applicable ENERGY STAR requirements changed to reference the new Version 2.0 EPS specification, effective November 2008. |
| Desktop, Integrated Desktop, Notebook/Tablet | |
| *Categorized into A, B, and C for Desktops and A and B for Notebook/Tablets.* | New categories are implemented in the 5.0 specification and defined in section *3)B)1* of the final requirements. The new categories reflect products reviewed during development of ENERGY STAR levels. |
| *Operational Mode power requirements (Off, Sleep, Idle).* | Version 5.0 introduces a Typical Energy Consumption (TEC) approach that combines the three operational mode measurements with defined usage pattern to calculate an annual energy consumption requirement. |
| Workstations | |
| *$P_{TEC}$ determined using a 10% weighting for Off, 20% for sleep, and 70% for Idle. Requirement of less than or equal to 35% of max power, scaled by hard drive number.* | The usage weightings are revised in Table 4 of the Version 5.0 specification, with the max power percentage now 28%. |
| Desktop-Derived Servers | |
| *Desktop-Derived Servers are addressed under the same requirements as desktop computers, with the exception of sleep requirements.* | These products are referred to as Small-Scale Servers in the Version 5.0 specification and are addressed with product-specific requirements that maintain Version 4.0's categories A and B, and associated power requirements. |
| Power Management | |
| *Wake On LAN required to be enabled on shipment for all computers shipping through Enterprise Channels.* | As an alternative, products shipping to enterprise may shipped without WOL not enabled provided features are sufficiently accessible from both the client operating system user interface and over the network. |
| General | |
| *No guidance provided on customer software and management pre-provisioning.* | Language and resource references for manufacturers installing pre-provisioned software packages are included in section *3)C*. |

| Version 4.0 | Version 5.0 |
|---|---|
| *Qualifying families of products allowed with provisions in the specification.* | Version 5.0 maintains the language from Version 4.0, but added is a clarification that all units/configurations associated with a qualified model designation meet ENERGY STAR requirements. For systems where non-qualifying configurations exist, an option is included for manufacturers to identify qualifying configurations with a consistent "identifier" in the model name/number in all marketing and qualification references. |

### 2.3.1 ENERGY STAR* Partnership and Submittal of Products

To qualify products for the ENERGY STAR program, manufacturers partner with the program, after which product data is reviewed to determine whether or not products may be marketed as ENERGY STAR qualified. There is no cost to become an ENERGY STAR partner, though manufacturers must carefully consider program requirements and commitments when deciding to pursue partnership. Below is a general reference on the ENERGY STAR partnership process

- **Review** program requirements and commitments;
- **Partner**
    - Complete and submit a Manufacturer Partnership Agreement Packet and submit to join@energystar.gov;
    - Approved manufacturers receive a Welcome Aboard Packet via email that includes program information and usernames and passwords to access ENERGY STAR online submittal and contact update tools;
- **Test**
    - Product testing is completed by the manufacturer and data is submitted to EPA for review;
    - Manufacturers receive feedback on approval or rejection of products.
- **Market**
    - Approved products may be marketed as ENERGY STAR qualified upon receipt of feedback. Manufacturers are encouraged to make use of marketing support and tools accessible from the ENERGY STAR website.

Interested manufacturers of the computers referenced in this document (and in the scope of the ENERGY STAR computer program) can find further information on this process on the ENERGY STAR website at www.energystar.gov/join.

### 2.3.2 ENERGY STAR* in the Channel Market

The channel market is an important avenue for the ENERGY STAR program, both in terms of sales and implementation of energy efficient IT. Below are notes and resources for channel organizations; solution providers are also encouraged to review this document for further details on the new specification that will aid in educating customers.

- Become an ENERGY STAR Partner. PC makers can join the thousands of partners who have already teamed with ENERGY STAR to save energy in homes and businesses through energy efficient products. Please go to www.energystar.gov/join.

- Take Advantage of Federal Government Requiring ENERGY STAR qualified PCs and power management. The Environmental Policy Act of 2005 requires that the federal government purchase ENERGY STAR qualified computers. Executive Order 13423 requires that federal government computers are set to enter low power sleep mode when inactive. For more information, please go to www.energystar.gov/fedofficeenergy.

- Exercise Caution When Re-Configuring an ENERGY STAR Qualified Computer. Not only optimized to meet strict energy efficiency requirements, ENERGY STAR qualified computers are also configured to enter low-power sleep modes and use Wake-on-LAN functions. System integrators and value-added resellers are encouraged to test any ENERGY STAR qualified configurations they modify in order to ensure that they still meet ENERGY STAR specifications.

- Become an ENERGY STAR Low Carbon IT Ally. Low Carbon IT Allies help spread the word about the energy savings from activating sleep settings on computers and purchasing ENERGY STAR qualified computers. It's easy: we provide the informational materials and resources, and you simply share them as you see fit. VARs and system integrators can become Allies at www.energystar.gov/lowcarbonit.

## 2.4 Taking Advantage of Power Management Settings

When attempting to re/configure computers, System Integrators or OEMs should be made aware of specifications that go into an ENERGY STAR qualified computer. ENERGY STAR qualified computers are optimized to meet strict energy efficiency requirements and they accomplish this not only by hardware differences, but also in how they are configured to enter low-power sleep modes, use Wake-on-LAN functions and to use other hardware specific power management features. Please minimize the impact on these settings by:

- Ensuring that the power management settings adhere to the ENERGY STAR guidelines. (choosing the energy saving power scheme in the OS as the default should get one most of the way there)

- Ensure that all drivers for the hardware are the correct versions for the hardware. Also ensure that the drivers are up to date as most issues experienced by customers with power management stem from improper or outdated drivers for the hardware.

- Ensure that the PC transitions between high and low power states after modifications.

- Ensure that the processor spends most of it's time in the lowest power C state supported by the platform (Usually C2 or higher). If not, the system may not be configured properly and missing or has improperly installed hardware drivers.

The diagram below shows the power options property sheet for Windows* Vista that allows users to control the various power settings. There are three predefined setting schemes, “Normal”, “Maximum Performance” and “Power Saving”. Within those schemes, there are a multitude of settings that can be controlled including monitor idle timeouts, system idle timeouts and processor performance states. ENERGY STAR requirements require that the “Turn off monitor” idle timeout setting be set (by default) to 15 minutes or less for AC

("Plugged in") operation, and "Sleep" system idle timeout be set for 30 minutes or less, when on AC power[3].

**Figure 2: Microsoft Windows Vista* Computer and Monitor Sleep Settings[4]**

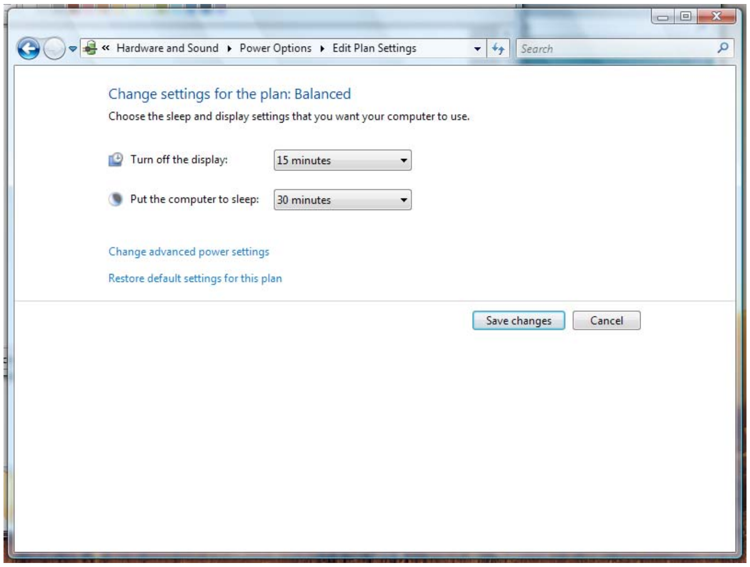


If rolling out Windows XP or Windows Vista* through the use of images (i.e.; Using Symantec Ghost* or similar) there is a method, albeit lacking centralized control (fire-and-forget option), a system administrator or OEM can set and ensure compliance to monitor power management policies. Please refer to Appendix A for more information.

## 2.5 Terminology

Some of the terminology used by the EPA in the ENERGY STAR specification is different than the language used in the computer industry. This section provides a definition of commonly used terms in order to prevent confusion of these terms. Other terms used in this document are listed here as well.

| Term | Description |
|---|---|
| AC | Alternating Current. |
| CAD | Computer Aided Drafting. |
| CAE | Computer Aided Engineering. |
| Capability | A set of features that enhance the usability and/or experience of a (compute) product; or, provide the ability to accomplish tasks or activities. |
| DC | Direct Current. |

[3] DC power settings are not specified by ENERGY STAR.
[4] Microsoft* product screen shot(s) reprinted with permission from Microsoft Corporation.

| Term | Description |
|---|---|
| DRAM | Dynamic Random Access Memory. This is the primary type of memory used in computer systems today. |
| ECC | Error Correcting Code. Error correcting code is a mechanism for improving the reliability of computer memory that allows the detection and correction of some types of memory errors. |
| Energy consumption | The amount of AC (wall plug) energy consumed by a system over a given period of time (hour, week, year) and is measured in kilowatt-hours (kW-hr). |
| Energy efficiency | The amount of AC (wall plug) energy consumed by a system to run a desired usage-based workload and is measured in kilowatt-hours (kW-hr). Usage-based workloads should include both active and non-active states that are reflective of the end user's use of the system. |
| Energy Efficient Performance | The intersection of great performance, expanded capabilities and energy efficiency. |
| Efficiency | Efficiency has two definitions. In the context of power delivery, efficiency is equal to (power out / power in) and can be represented by the symbol $\eta$ (*eta*). Another definition is that efficiency is a measure of the production of work versus the cost (time, energy, money, etc.). |
| Environmental impact | An assessment of any change to the environment whether adverse or beneficial, wholly or partially resulting from the activity of creating, using, or disposing of items (we) produce. |
| GPU | Graphics Processing Unit. This is typically a silicon component on the motherboard or on an add-in card that processes graphics information for display. |
| Discrete GPU | A graphics processor with a local memory controller interface and a local, graphics-specific memory. |
| Idle State | The state in which the operating system and other software have completed loading, a user profile has been created, the machine is not asleep, and activity is limited to those basic applications that the system starts by default. |
| MIPS | Millions of Instructions Per Second. MIPS are sometimes used as an indicator of computer or CPU performance. |
| Network Interface | The components (hardware and software) whose primary function is to make the computer capable of communicating over one or more network technologies. Examples of Network Interfaces are IEEE 802.3 (Ethernet) and IEEE 802.11 (Wi-Fi). |
| Performance | The compute throughput and responsiveness at a component level – or – Compute throughput and responsiveness at a system level. |

| Term | Description |
|---|---|
| OEM | Original Equipment Manufacturer. In the context of this paper these are computer manufacturers. |
| PF | Power Factor. This is the ratio of actual power, Voltage (V) * Current (I) at the same instant of time, to the apparent power level, Voltage * Current, uncorrected by time or phase. In an AC power system the cyclical waveforms on voltage with peak of 115V do not correspond to the current, eg. 3.5A, being consumed. A 90V corresponding to 3.5A, 315Watts, is different than 100 to 2.2A, or 220 Watts. |
| PFC | Power Factor Correction. This is used in computer power supplies to improve the Power Factor of the system. |
| Power | The measurement of energy consumption in watts at a specific point in time and under a fixed, static condition. Examples: maximum power, active and idle. |
| Sleep Mode | A low power states that the computer is capable of entering automatically after a period of inactivity or by manual selection. A computer with sleep capability can quickly "wake" in response to network connections or user interface devices with a latency of ≤ 5 seconds from initiation of wake event to system becoming fully usable including rendering of display. For systems where ACPI standards are applicable, Sleep mode correlates to ACPI System Level S3 (suspend to RAM) state. |
| Off Mode | The power consumption level in the lowest power mode which cannot be switched off (influenced) by the user and that may persist for an indefinite time when the appliance is connected to the main electricity supply and used in accordance with the manufacturer's instructions. For systems where ACPI standards are applicable, Off Mode correlates to ACPI System Level S4 or S5 states, where applicable. |
| UMA | Unified Memory Architecture. Computer memory architecture used for integrated graphics implementations where the system memory is used for graphics memory as well. |
| Wake On LAN (WOL) | Functionality which allows a computer to wake from Sleep or Off Mode when directed by a network request via Ethernet. |
| Typical Energy Consumption (TEC) | A method of testing and comparing the energy performance of computers, which focuses on the typical electricity consumed by a product while in normal operation during a representative period of time. For Desktops and Notebooks, the key criterion of the TEC approach is a value for typical annual electricity use, measured in kilowatt-hours (kWh), using measurements of average operational mode power levels scaled by an assumed typical usage model (duty cycle). For Workstations, requirements are based on a TEC power value calculated from operational mode power levels, maximum power, and an assumed duty cycle |

| Term | Description |
|---|---|
| Full Network Connectivity | The ability of the computer to maintain network presence while in sleep and intelligently wake when further processing is required (including occasional processing required to maintain network presence). Maintaining network presence may include obtaining and/or defending an assigned interface or network address, responding to requests from other nodes on the network, or maintaining existing network connections, all while in the sleep state. In this fashion, presence of the computer, its network services and applications, is maintained even though the computer is in sleep. From the vantage point of the network, a sleeping computer with full network connectivity is functionally equivalent to an idle computer with respect to common applications and usage models. Full network connectivity in sleep is not limited to a specific set of protocols but can cover applications installed after initial installation. |

## 2.6 ENERGY STAR* and Computer Platforms

The ENERGY STAR specification recognizes and specifies seven different types of platforms but then classifies these platforms into three different categories with specific requirements. The six platforms addressed by this document are defined below, followed by a summary of the primary requirements each must meet:

**Notebook Computer:** A computer designed specifically for portability and to be operated for extended periods of time either with or without a direct connection to an AC power source. Notebooks must utilize an integrated computer display and be capable of operation off an integrated battery or other portable power source. In addition, most notebooks use an external power supply and have an integrated keyboard and pointing device. Notebook computers are typically designed to provide similar functionality to desktops, including operation of software similar in functionality as that used in desktops. Tablet PCs, which may use touch-sensitive screens along with or instead of other input devices, are considered Notebook Computers.

**Desktop:** A computer where the main unit is intended to be located in a permanent location, often on a desk or on the floor. Desktops are not designed for portability and utilize an external computer display, keyboard, and mouse. Desktops are designed for a broad range of home and office applications

**Integrated Desktop Computer:** A desktop system in which, the computer and computer display function as a single unit which receives its AC power through a single cable. Integrated desktop computers come in one of two possible forms: (1) a system where the computer display and computer are physically combined into a single unit; or (2) a system packaged as a single system where the display is separate but is connected to the main chassis by a DC power cord and both the computer and computer display are powered from a single power supply. As a subset of desktop computers, integrated desktop computers are typically designed to provide similar functionality as desktop systems.

**Thin Client:** An independently-powered computer that relies on a connection to remote computing resources to obtain primary functionality. Main computing (e.g., program execution, data storage, interaction with other Internet resources, etc.) takes place using the remote computing resources. Thin Clients covered by this specification are limited to devices with no rotational storage media integral to the computer. The main unit of a Thin Client covered by this specification must be intended for location in a permanent location (e.g. on a desk) and not for portability.

**Workstation:** A high-performance, single-user computer typically used for graphics, CAD, software development, financial and scientific applications among other compute intensive tasks. To qualify as a workstation, a computer must:

- Be marketed as a workstation;
- Have a mean time between failures (MTBF) of at least 15,000 hours based on either Bellcore TR-NWT-000332, issue 6, 12/97 or field collected data; and
- Support error-correcting code (ECC) and/or buffered memory.

Additionally to be defined as a workstation the system must also have a number of other characteristics (picked from a list) that will be discussed in the workstation section of this paper.

**Small-Scale Server:** A computer that typically uses desktop components in a desktop form factor, but is designed primarily to be a storage host for other computers. explicitly to be a host for other computers or applications. A computer must have the following characteristics to be considered a Small-Scale Server:

- Designed in a pedestal, tower, or other form factor similar to those of desktop computers such that all data processing, storage, and network interfacing is contained within one box/product;
- Intended to be operational 24 hours/day and 7 days/week, and unscheduled downtime is extremely low (on the order of hours/year);
- Capable of operating in a simultaneous multi-user environment serving several users through networked client units; and
- Designed for an industry accepted operating system for home or low-end server applications (e.g., Windows Home Server, Mac OS X Server, Linux, UNIX, Solaris).

Small-Scale Servers are designed to perform functions such as providing network infrastructure services (e.g., archiving) and hosting data/media. These products are not designed to process information for other systems or run web servers as a primary function. Small-Scale Servers covered by the Version 5.0 specification for computers are limited to computers marketed for non-datacenter operation (e.g. homes, small offices).

| Product Category | Category Requirements |
|---|---|
| Notebooks | $E_{TEC}$ Requirements (category A, B, and C) |
| Desktop | $E_{TEC}$ Requirements (category A, B, C, and D) |
| Integrated Desktop Computers | $E_{TEC}$ Requirements (category A, B, C, and D) |
| Workstations | $P_{TEC}$ Requirement |
| Thin Clients | Off Mode Requirements<br>Sleep Requirements (if applicable)<br>Idle Power Requirements (category A and B) |
| Small-Scale Servers | Off Mode Requirements<br>Idle Power Requirements (category A and B) |

This paper will talk about each of the different product categories (e.g. treats the desktop and integrated desktop computers the same as they have the same requirements). Each section will describe the specific requirements for each unique category in more detail.

## 2.7 ENERGY STAR* Common Platform Requirements

The ENERGY STAR specification has a number of requirements which are common across all of the platform categories. This section will outline these common attributes.

All systems are tested "as shipped", unless otherwise specified. In general the tester is not allowed to enable or disable any power management settings specifically for testing purposes, unless the testing process explicitly calls out to do something differently.

Power management is emphasized in the ENERGY STAR specifications, and each system category must meet requirements as appropriate for the technology and usage case involved:

- Display's sleep mode (blank display after idle) should be enabled to activate within 15 minutes or less of idle.
- Platform's sleep mode (enter sleep mode when idle) should be enabled to activate within 30 minutes or less of idle; small-scale servers and thin clients, only, are exempt from this requirement.
- Platform's Gigabit Ethernet Link should switch to a lower rate mode (100 Mb or 10 Mb) when entering the sleep mode.

Additionally, systems which ship into an enterprise market are required to ship with Wake On LAN (WOL) capability[5] in the sleep state. For enterprise market shipments, version 5.0 of the specification includes a provision to allow WOL capability to be turned off at shipment provided WOL features are easily enabled from both the client OS and network. Systems targeted for consumer channels are not required to enable the higher power WOL in the sleep state. Thin Clients and Small-Scale Servers are exempt from the platform power management requirements above, but must meet the display power management requirement as applicable. Thin Clients not designed for off-hours scheduled updates are exempt from the WOL requirement.

## 2.8 Full Network Connectivity and Proxying

The ENERGY STAR Version 5.0 specification introduces eliminates the requirement of implementing full network connectivity and instead introduces incentives in its place. These incentives are designed to help adoption of any standards that arise from Ecma TC32-TG21 which is a task group working on a standard for proxying common network traffic.

Full network connectivity is defined as the ability of the computer to maintain network presence while in sleep and intelligently wake when further processing is required (including occasional processing required to maintain network presence). Maintaining network presence may include obtaining and/or defending an assigned interface or network address, responding to requests from other nodes on the network, or maintaining existing network connections, all while in the sleep state. In this fashion, presence of the computer, its

[5] Set to wake-up on traffic from management workstations is the preferred setting.

network services and applications, is maintained even though the computer is in sleep. From the vantage point of the network, a sleeping computer with full network connectivity is functionally equivalent to an idle computer with respect to common applications and usage models. Full network connectivity in sleep is not limited to a specific set of protocols but can cover applications installed after initial installation.

The ability for OEMs to take advantage of these incentives is contingent on EPA/ENERGY STAR confirmation that any and all standards being used do in fact conform to the ultimate goals and challenges laid out by ENERGY STAR. The most important aspect of proxying from an implementation standpoint is that it will not require modifications to legacy systems in order for proxying capable systems to benefit from the energy savings.

Other requirements that are specific to a product will be called out in those specific sections.

## 2.9 Reference Documents

| Document | Document No./Location |
|---|---|
| Advanced Configuration and Power Interface (ACPI) Specification | http://www.acpi.info/spec.htm |
| ENERGY STAR* Program Requirements for Computers | http://www.energystar.gov/ia/partners/prod_development/revisions/downloads/computer/Version5.0_Computer_Spec.pdf |
| ENERGY STAR Program Requirements for Single Voltage External Ac-Dc and Ac-Ac Power Supplies | http://www.energystar.gov/ia/partners/product_specs/program_reqs/EPS_Eligibility_Criteria.pdf |
| ENERGY STAR - Online Product Submittal System | http://www.energystar.gov/ops |
| Generalized Internal Power Supply Efficiency Test Protocol | http://www.efficientpowersupplies.org |
| Hybrid Hard Disk And Ready Drive* Technology: Improving Performance And Power For Windows Vista Mobile PCs | http://download.microsoft.com/download/5/b/9/5b97017b-e28a-4bae-ba48-174cf47d23cd/STO008_WH06.ppt |
| Power Supply Design Guide for Desktop Platform Form Factors | http://www.formfactors.org |

§

# 3 Notebooks and ENERGY STAR* Requirements

## 3.1 Notebooks, Battery Life and AC Energy

The notebook market has always had an appreciation for low power operation because of the reliance on batteries as a power source; the lower the power, the longer the battery life. AC power is a different story where notebooks traditionally optimize for performance, not power. Because many of the battery life optimizations can affect system performance or provide annoyances to the user (like the screen blanking when reading a document), power management features are typically disabled when the notebook is in AC mode.

AC power is becoming more important as energy conservation is seen more and more as a desirable trait and as energy costs increase. The key for a notebook designer is to balance which power management techniques to use in AC operation to provide the best performance and usability while meeting the ENERGY STAR energy metrics. The good news for notebook designers is that a host of power management tools are already available, and in most cases it is just a matter of enabling power management features in both battery and AC operating states. Due to the increased industry and end user awareness of energy conservation, the EPA has targeted challenging new energy consumption standards on computers where only 25% of today's notebooks are expected to pass the new ENERGY STAR logo requirements. These are not easy requirements to meet and you cannot expect to meet all requirements on all product models.

This section will review the requirements for ENERGY STAR notebooks, explain where much of the power is distributed throughout the notebook and review the key areas to investigate in order to meet these new requirements.

## 3.2 ENERGY STAR* Basics for Notebooks

Based on the latest Energy Star specification for computers Version 5.0, computers must meet the requirements below to qualify as ENERGY STAR.

The new specification defines Typical Energy Consumption (TEC) as a method of testing and comparing the energy consumption of computers, and focuses on the typical electricity consumed by a product while in normal operation over a year. The TEC approach calculates annualized energy consumption in KWhr by a formulae which weights the off, sleep and idle power of a system; where the weighting is based on the typical usage patterns over a year and is determined using the formula below:

$E_{TEC} = (8760/1000) * (P_{off} * T_{off} + P_{sleep} * T_{sleep} + P_{idle} * {}_{Tidle})$

where all Px are power values in watts, all Tx are Time values in % of year, and the TEC $E_{TEC}$ is in units of kWh and represents annual energy consumption based on mode weightings in Table 3. The formulae basically converts the "average weighted power" (Poff*Toff + Psleep*Tsleep + Pidle*Tidle) in watts to yearly Watt-hours by multiplying by the number of hours in a year (8760 hours/year) and converting from Watts to KiloWatts (by dividing by 1000). Finally, certain hardware-based capabilities are allowed defined capability adjustments in kWh ("adders") that are added to TEC limits when determining the TEC limit a computer must meet. An illustration of TEC calculation is provided in Appendix B of this document.

### Table 1. $E_{TEC}$ Requirements – Notebooks

| | Notebook Computers (KWh) | |
|---|---|---|
| **Power Supply** | Avg eff, No Load eff (EnergyStar EPS 2.0 Specification) | |
| **TEC (kWh) Limits** | Category A: ≤ 40.0 | |
| | Category B: ≤ 53.0 | |
| | Category C: ≤ 88.5 | |
| **Capability Limit Adjustments** | | |
| **TEC (kWh)** | Memory | 0.4 kWh (per GB over 4) |
| | Premium Graphics (for Discrete GPUs with specified Frame Buffer Widths) | Cat. B: 3 kWh (FB Width > 64-bit) |
| | Additional Internal Storage | 3 kWh |

### Table 2. Operational Mode Weighting for Conventional Notebooks

| Operational Mode | Conventional Operational Mode Weighting | Proxying Operational Mode Weighting[6] |
|---|---|---|
| Toff | 60% | 45% |
| Tsleep | 10% | 30% |
| Tidle | 30% | 25% |

$(8760/1000) * (P_{off} * 60\% + P_{sleep} * 10\% + P_{idle} * 30\%)$

As mentioned in the introduction, the ENERGY STAR specification requires that notebooks are tested as shipped, and further they are required to be shipped with the following power management features enabled: System should be set to enter a sleep state after 30 minutes or less of idleness and the display should be set to blank after 15 minutes or less of idleness.

---

[6] Upon approval by EPA and EC of a non-proprietary proxying standard

EnergyStar has three notebooks categories which provide a different TEC limit (for each category): A, B, and C. Each category represents a different type of notebook configuration, where the more highly configured notebooks have a higher TEC limit. Category C is more highly configured than Category B, which is more highly configured than Category A. Category C has the highest TEC limit, while Category A has the lowest TEC limit. A notebook must meet the configuration criteria as specified in the Version 5.0 specification, which is summarized below:

**Category A:** All notebook computers that do not meet the definition of Category B or Category C below will be considered under Category A for ENERGY STAR qualification.

**Category B:** To qualify under Category B, notebooks must have:

- A Discrete GPU.

**Category C:** To qualify under Category C, notebooks must have:

- Greater than or equal to 2 Physical Cores;
- Greater than or equal to 2 gigabytes (GB) of System Memory; and
- A Discrete GPU with a Frame Buffer Width greater than 128-bit.

The power of the notebook is measured at the AC wall socket and is measured with an approved meter as described in Appendix A of the ENERGY STAR specification. The notebook test conditions for power mode testing are:

- Configure the notebook to blank its screen after 1 minute,
- Configure networking to appropriate conditions by doing one of the following:
  - *Notebooks with Ethernet capability*: connect the Ethernet port of the notebook to an active switch which supports the highest supported network throughput (typically a gigabit switch with today's systems) and turn off all wireless (radio) devices; or
  - *Notebooks without Ethernet capability*: power to a wireless LAN radio (e.g. IEEE 802.11) should remain on during testing and must maintain a live wireless connection to a wireless router or network access point, which supports the highest and lowest data speeds of the client radio, for the duration of testing.
- Remove the battery from the system (while leaving the AC cord attached). This is to ensure that the battery is not charging during the test (which would make power measurements inaccurate). If the system has an integrated battery (can't be removed) then use other means to ensure the battery is not charging during power measurements. Note, a charging battery during the test can add significant power to the measurement (50Whr charging battery can add 50W to a power measurement!).
- Do not plug anything else into the system (as shipped)
- Primary hard drives may not be power managed ("spun-down") during Idle testing unless containing non-volatile cache integral to the drive (e.g. "hybrid" hard drives). 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 if it is enabled when the product is shipped

Other than what is mentioned above, all other configurable parameters of the notebook must remain in the "as-shipped" state.

The system is then shut down, and the tester begins to measure power consumed in off during a 5 minute period, recording the average over this period. The system is then rebooted, and after reaching a stable state (finished booting) the tester will wait between 5 and 15 additional minutes and then will start measuring and averaging the power for the next 5 minutes. Finally, the system is placed into the sleep state where the power is again measured and averaged over 5 minutes.

This is the extent of the testing involved for notebooks. It is important to remember the testing configuration: screen blanked, wireless devices and hard drives configured appropriately and battery removed. It is also important to remember the testing environment: connected to an active Ethernet switch (or wireless access point), battery removed, and waiting 5-15 minutes from boot before taking the idle measurement. This idle time is important as it allows the system to reach an idle condition, allowing one-time processes to complete before measurement. Full testing provisions are detailed in Appendix A of the ENERGY STAR computer specification.

## 3.3 The Notebook Platform

For the purpose of understanding where power is consumed in the system, let's define a generic notebook (see diagram below, Figure 3) such that it will be easier to understand the power breakdowns we illustrate later. The platform consists of the following main components:

- CPU – a chip which contains computer's execution engine or processor, Memory controller and Graphics Controller in the case of an internal graphics system
- PCH – chip which contains much of the platform's display and I/O bus controllers
- DRAM – platform memory
- Audio – in this case the block refers to the audio CODEC (controller is in PCH)
- HDD – Hard Disk Drive, the controller is in the PCH chip
- ODD – Optical Disk Drive (CD or DVD drive), the controller is in the PCH chip
- LAN – This refers to the physical layer of the Ethernet chip, the Ethernet MAC controller is integrated within the Intel PCH, but requires an external PHY chip. In the power blocks it will also be referred to as COMM Other – This refers to special subsystems on the notebook for controlling the integrated keyboard, mouse and for controlling much of the power delivery system
- CPU VR – This is the component which delivers power to the CPU
- Plat VR – This is symbolically the component which delivers power to the rest of the platform. It should be noted that there are multiple voltages, VRs and FETs delivering power to the rest of the platform, but for simplicity we represent it as a single block with a single efficiency.
- Display – This is typically the LCD panel for notebooks. Because the panel is turned off during all ENERGY STAR testing, it is listed here only to prevent the inevitable "hey you forgot the panel!" statement. There will be no panel breakout in the ENERGY STAR power numbers.
- AC Brick – This refers to the external power adapter used to provide the notebook with power. For EPA purposes, its power is part of the total power of the platform.

**Figure 3. Typical CAT A Notebook System (Integrated graphics, UMA Memory)**

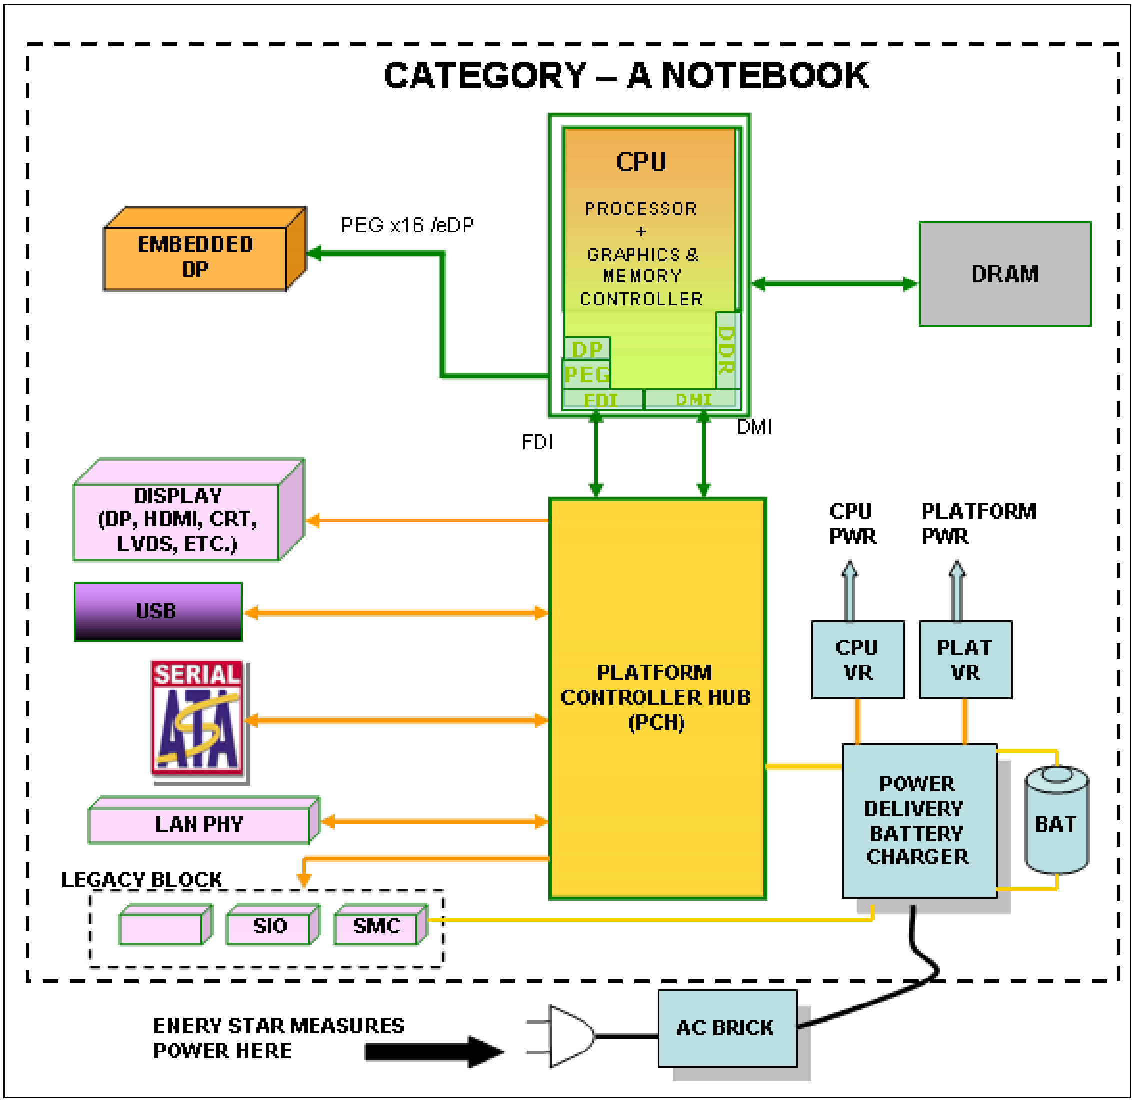


Figure 3 represents a type A platform with an integrated graphics solution using UMA (frame buffer is part of system DRAM memory). While we will not be doing a power breakout for a Type B or C notebook, one is illustrated in Figure 4. Note that the graphics chip is connected to a special graphics bus for Type B/C notebooks, and has a separate frame buffer. One of the distinguishing features of a Category C vesus Category B notebook is the width of this discrete graphics memory bus; to qualify as Category C this bus must be wider than 128-bits (in addition to the other criteria of 2 or more CPU cores and 2Gbytes or more of system memory)ENERGY STAR has provided an idle TEC requirement of 53 kWh for the Type B notebook and 88.5KWh for the Type C notebook versus the 40 kWh idle TEC requirements for the Type A notebook.

**Figure 4. Typical CAT B Notebook (discrete graphics)**

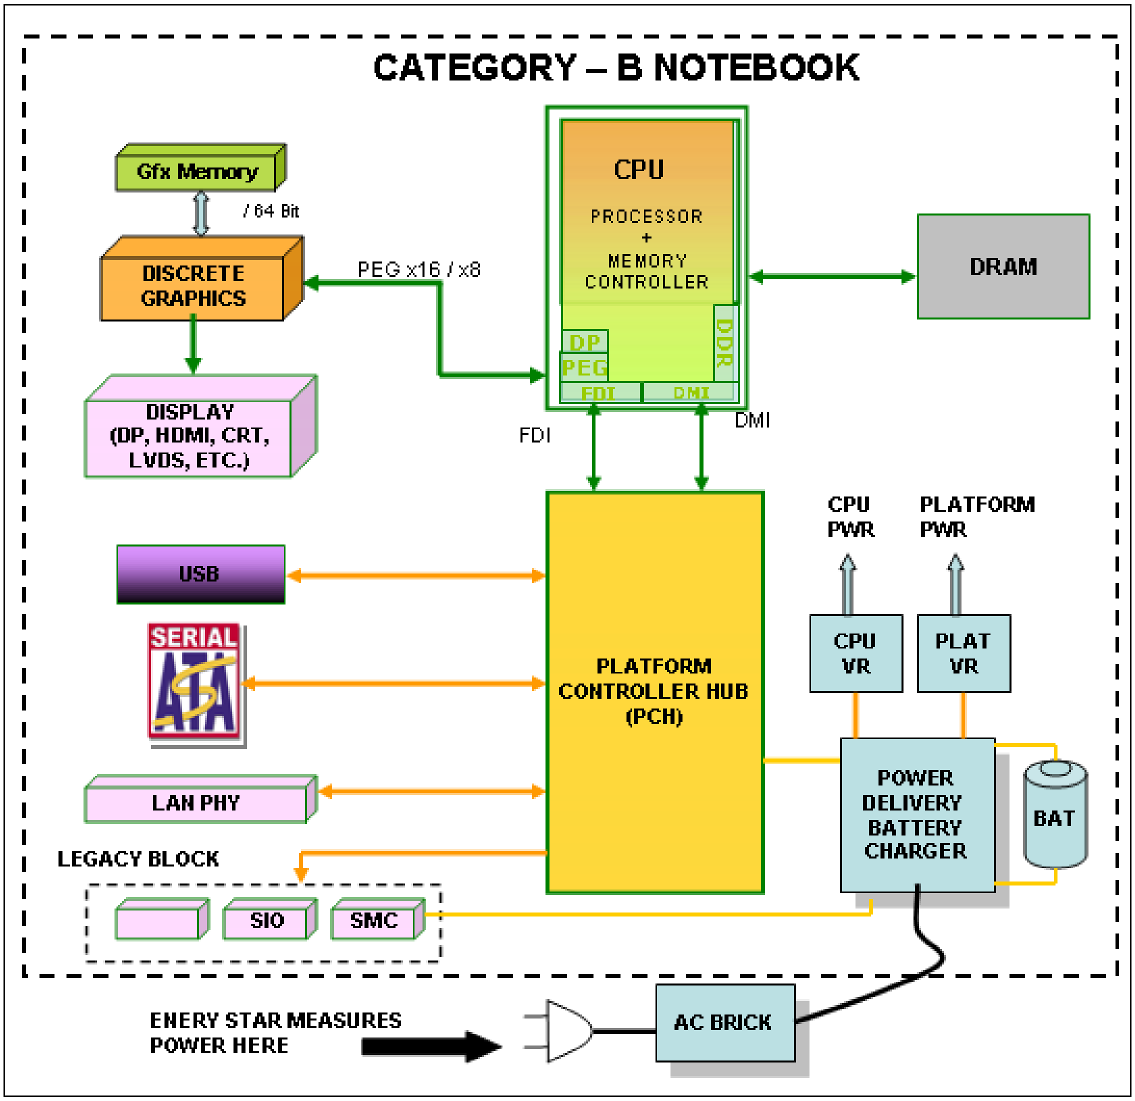


## 3.4 Power Delivery: A System's Power Amplifier

To meet the ENERGY STAR specification, power delivery is the most critical aspect of the design. While notebooks already optimize power delivery for idle, off and sleep to maximize battery life, the AC configuration adds one additional layer of power delivery (AC Brick) making power delivery key to an efficient AC design. For notebook designers, it is important to understand the difference between the AC and DC power delivery.

**Figure 5. Notebook Power Delivery System**

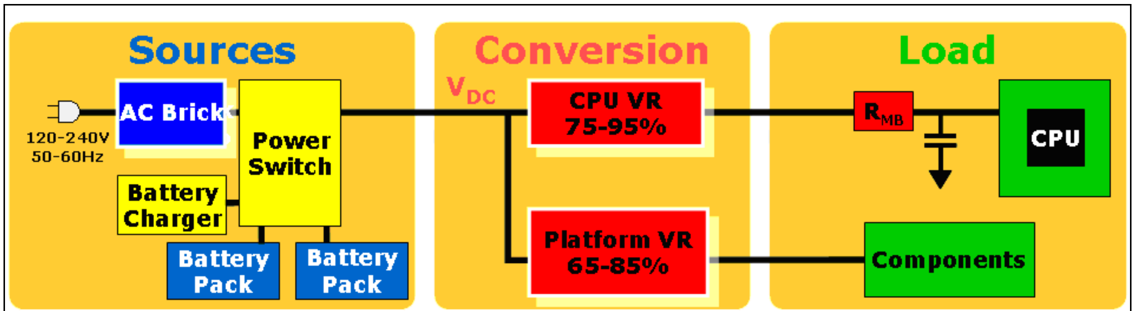


AC brick,efficiency is dictated by the Energy Star EPS spec

The notebook's power delivery can be modeled by three major categories (see Figure 5):

1. Load
2. Power Conversion
3. Power Source

For our purposes we will simplify the load into two areas: the CPU and the rest of the platform. The CPU load changes dramatically over time (micro-seconds) and so it is modeled separately. The rest of the platform consists of many different voltage rails (loads), but will be represented as a single load represented by the "Component" block.

The second aspect is power conversion. Here power is converted from the voltage delivered by the power source, to the voltages needed by the loads. The efficiency of this conversion depends on a number of things:

- The input voltage (VDC) to the regulator
- The quality of the regulator
- The load on the output of the regulator

For purposes of this paper, the power conversion is represented by to two voltage regulators to match how the loads are modeled: the first for the CPU load (CPU Voltage Regulator or VR) and the second representing many voltage regulators that supply the rest of the platform (Platform VR). Associated with each power conversion VR is an efficiency represented by the amount of power delivered to the load divided by the amount of power input into the VR in order to deliver that load (Pout/Pin). As shown in the diagram, the efficiency can vary. This efficiency can vary by manufacturing distributions (changes due to manufacturing variations), but more importantly each VR's efficiency will vary depending on load presented to it.

The input voltage to the power conversion also affects the VR's efficiency. A higher input voltage typically results in lower conversion efficiency compared to a lower input voltage (VDC). This is important because batteries will have a much lower output voltage than what the AC brick delivers (a fact dependent on the charging circuits), so inherently the power delivery (CPU VR and PLT VR) will be less efficient in AC mode than in DC mode.

In the case of ENERGY STAR, the power source will be the AC brick. Again the AC brick can be thought of providing as a power conversion from AC power to some lower AC or DC voltage and therefore it has a power conversion efficiency associated with it also.

For example a power conversion component which is 80% efficient will require 1 W to deliver 0.8 W to a load. For purposes of this paper, we will represent the efficiency of the CPU VR by CPUeff, the efficiency of the Platform VR by PLTeff and the efficiency of the AC Brick by ACeff.

Right away you will note that the power of the CPU load is amplified by the efficiency of the CPU VR (1/CPUeff), and then further amplified by the AC Brick (1/ACeff). The total power of the system can then be modeled by the equation:

**Equation 1**

$$AC(W) = \frac{1}{ACeff}\left(\frac{1}{CPUeff}\cdot CPULoad + \frac{1}{PLTeff}\cdot PLTLoad\right)$$

For the CPU, if you assume 80% efficient regulation, and the AC brick is 80% efficient, then a 1 W CPU load requires 1.3 W of AC power. If you assume 75% efficient load for the platform, then a 1 W platform load would require 1.7 W of AC power. As can be seen the power delivery amplifies the power needed at the load.

## 3.5 ENERGY STAR* Requirements for the AC Brick

For AC power operation, the AC brick is one of the critical factors affecting power in the platform and has an additional set of requirements outlined in a separate ENERGY STAR specification: *ENERGY STAR Program Requirements for External Power Supplies (EPS) version 2.0*. In this section we will review these requirements and analyze how they affect notebook designs.

This document outlines three major requirements:

1. Minimum Average Efficiency in Active Mode
2. Maximum Power at no load
3. Power Factor requirements

The first requirement outlines the average efficiency required by the AC brick and this is dictated by Table 3 which is copied from the EPS 2.0 specification. Notebooks typically fall into the category of EPS that have nameplate power of greater than 49W and therefore must have an average efficiency of 87% or higher.

**Table 3. Active Mode Requirements for Average Efficiency**

| Nameplate Output Power (Pno) | Minimum Average Efficiency in Active Mode (expressed as a decimal)* |
|---|---|
| 0 to ≤ 1 W | ≥ 0.480 * Pno + 0.140 |
| > 1 to ≤ 49 W | ≥ [0.0626 * Ln (Pno)] + 0.622 |
| > 49 W | ≥ 0.870 |

The second requirement comes from another table in the EPS 2.0 specification which outlines the no-load requirements (see table 4) which show that the no load power (power the AC brick draws when plugged into the wall but not into the notebook) must be less than

0.3W if nameplate power is less then 50W or less than 0.5W if the nameplate power is 250W or less and 50W or greater.

### Table 4. No Load Requirements

| Nameplate Output Power ($P_{no}$) | Maximum Power in No-Load (Ac-Dc) |
|---|---|
| 0 to < 50 watts | ≤ 0.3 watts |
| ≥ 50 to ≤ 250 watts | ≤ 0.5 watts |

Average efficiency is defined by measuring and averaging efficiencies at four different loads (100%*Po, 75%*Po, 50%*Po and 25%*Po); where Po is the maximum rated output power of the AC brick. This is illustrated in Figure 6 (note the blue circles).

### Figure 6. Example AC Brick Efficiency over Different Loads

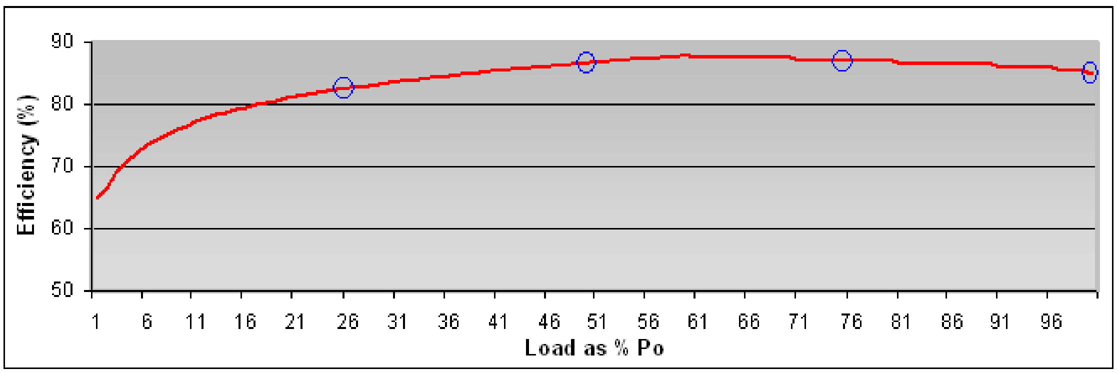


This diagram shows the efficiency of a brick on the vertical axis, and the load (as a percent of the maximum rated load, Po) on the horizontal axis; the red line representing the efficiency of the AC brick over various loads. As noted you can see the efficiency is lower for low loads, and also trails off for very high loads. The four blue circles represent the efficiency/load points to measure and average in order to meet the ENERGY STAR requirement Average Efficiency. For simplicity assume that the Po for this particular AC brick is 100 W (so the percentage and load values are the same).

Do not rely on the ENERGY STAR External Power Supply specification for the AC brick to create an acceptable power brick. This specification is generic and covers any AC brick a consumer device may use and does not require high efficiencies at loads that matter to the *notebook* ENERGY STAR specification. More specifically, the specification requires high efficiency from 25% to 100% of nameplate load. Because the brick is designed to meet maximum power requirements for simultaneous operation and charging, the laptop will operate at a low fractional load during the ENERGY STAR idle test.

The ENERGY STAR specification for external power supplies (EPS) determines the load fraction by the ratio of current (amps) to nameplate current, the load fractions in the specification (i.e. 25%, 50%, 75%, 100%) do not always match the ratio of power (W) to nameplate power. Because the computer specification is based on power (W) it is important to consider the actual load at idle and the AC brick nameplate power rating.

Therefore, when purchasing or designing AC bricks for notebooks, implement low load efficiency requirements that will correspond to the loads of interest at Idle, Sleep and off states. Maximizing efficiency at these loads will decrease the overall power draw of your notebook computer (less loss through the AC brick).

As we shall see next, the majority of notebook power (as it concerns ENERGY STAR) is consumed by the power delivery.

The last requirement has to do with power factor correction, which is now required for any AC brick which will draw more than 100W for the mains (wall outlet). The EPS 2.0 specification requires a power factor of 0.8 or greater at 100% rated load when tested at 115volts@60Hz.

## 3.6 Power Breakdown of the Notebook

For the notebook we described previously, measured in an ENERGY STAR idle configuration, a typical power breakdown for a category A notebook is illustrated in Figure 7.

**Figure 7. Typical Notebook Idle Power Breakdown**

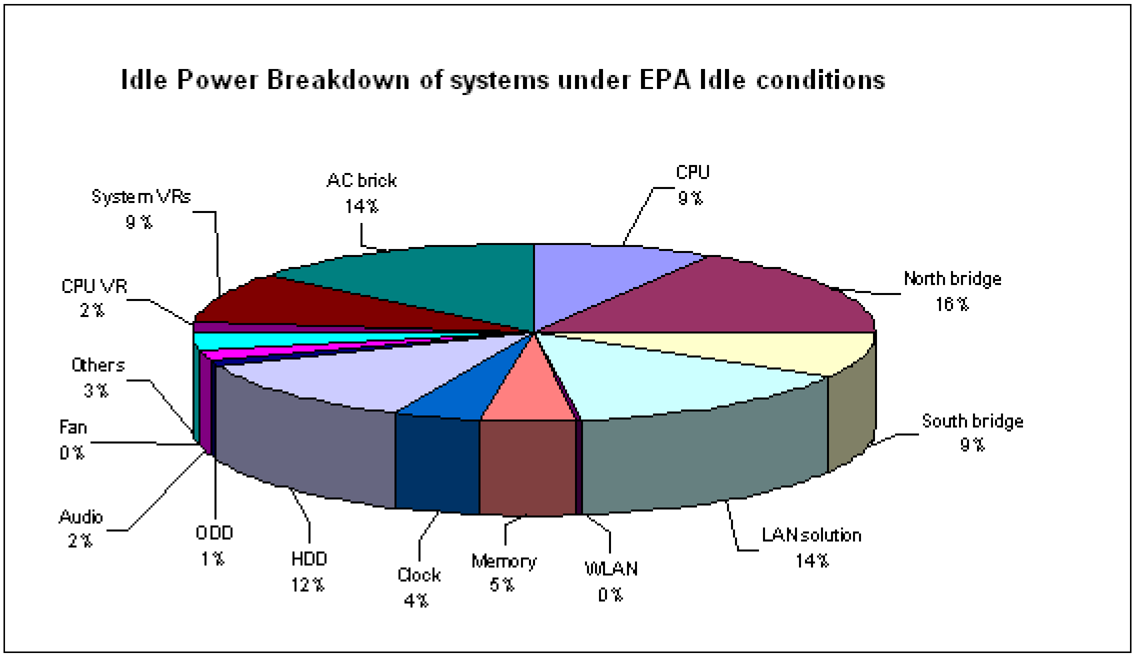


Power delivery represents 25% of the total power (AC Adapter, system VRs and CPU VR) assuming a CPU, Platform and AC Adapter efficiencies of 80%, 88% and 86% respectively. The LCD Panel is not shown as the backlight is blanked and no power is drawn. The power breakdown of the other platform components is also illustrated.

Another major component of power here is represented by the LAN solution (Ethernet PHY and small VRs for PHY).

Note that if a 52W-hr battery was inserted into the AC powered system, then the power draw could go up substantially as the system would add another 52W load and require an additional 69W of AC power to charge the battery. This is why we remove the battery when measuring power (if it is not fully charged, it could easily disrupt idle power measurements).

For the notebook we described previously, measured in an ENERGY STAR sleep (S3) and off (S5) configurations (with WOL disabled), a typical power breakdown for a notebook is illustrated in Figure 7. In this case the power delivery consists of 42% in sleep and 63% in off. The reason why the portion for power delivery is bigger in sleep and off than in idle is due to lower conversion efficiency in AC Adapter and VRs with lighter output current.

**Figure 8. Typical Notebook Sleep and Off Power Breakdown**

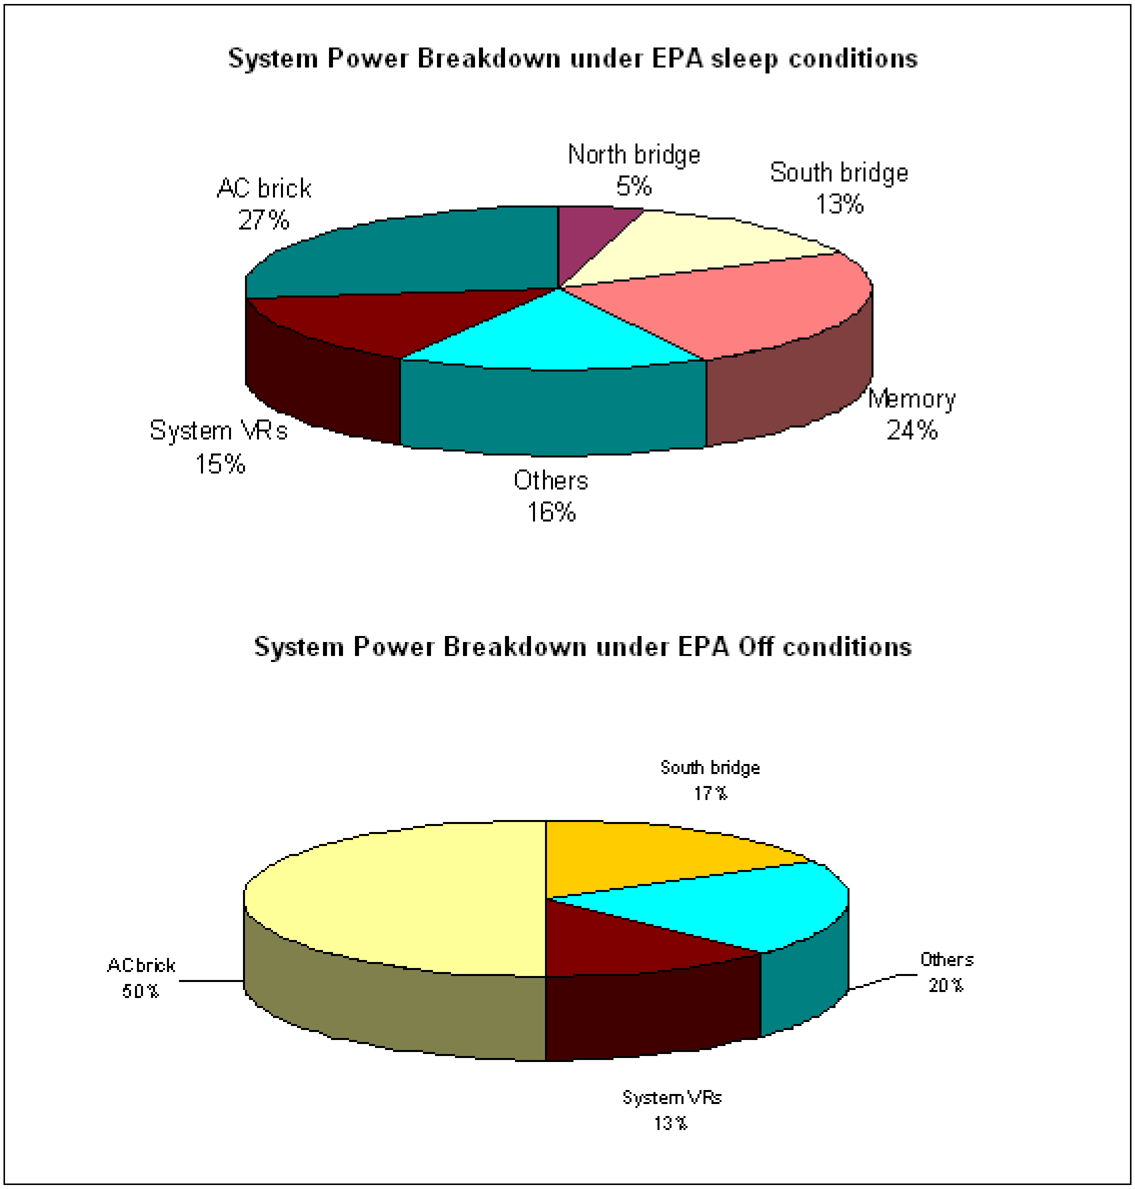


We have shown that the power distribution is the major component in the idle and sleep powers, let us next look at the sensitivity of this power to the efficiencies of the power delivery system.

## 3.7 Power sensitivity to Power Delivery Efficiencies

As we noted already, VR efficiencies change over various loads; they tend to get less efficient at very low loads, or very high loads. Let us look at how the platform power distribution changes based on the power delivery efficiencies changing.

### Figure 9. Idle Power distribution across efficiencies

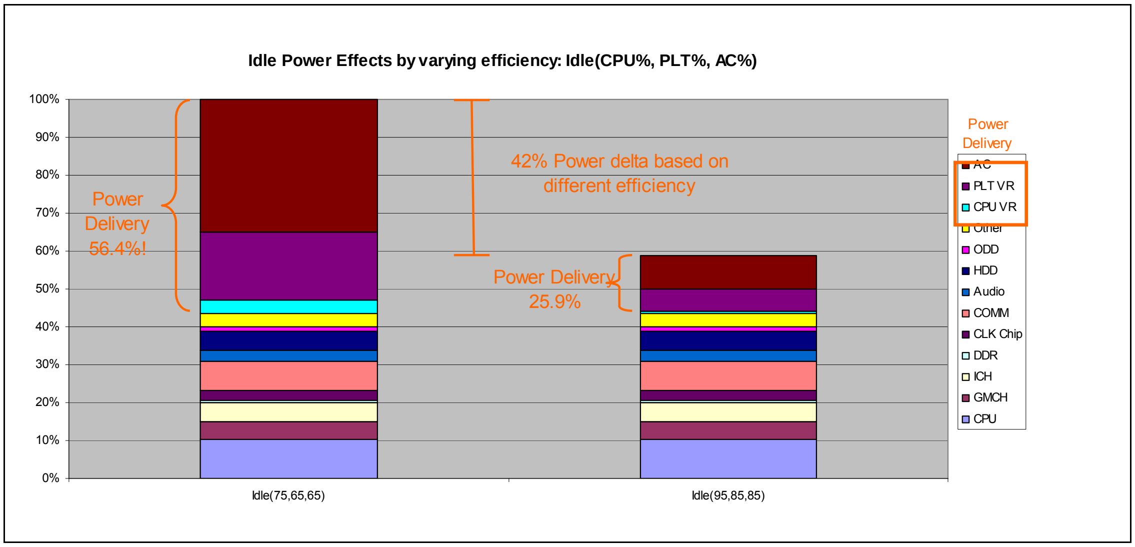


In Figure 9 we show the effects of changing the efficiency of the power delivery (CPU VR, PLT VR and AC) efficiencies. The first bar graph shows the platform power breakdown for efficiencies of 75%, 65%, 65% respectively and the second bar graph of 95%, 85%, 85% respectively. As can be seen the total power of the system dropped by 42%. In the first case the power delivery contributed to 56.4% of the platform, while in the more efficient case it only contributed to 26% of the platform power.

In Figure 10 we have modified the sleep power delivery efficiency (PLT VR and AC brick) from 50%, 50% in the first bar graph to 75%, 75% respectively for the second bar graph.

### Figure 10. Sleep Power distribution across efficiencies

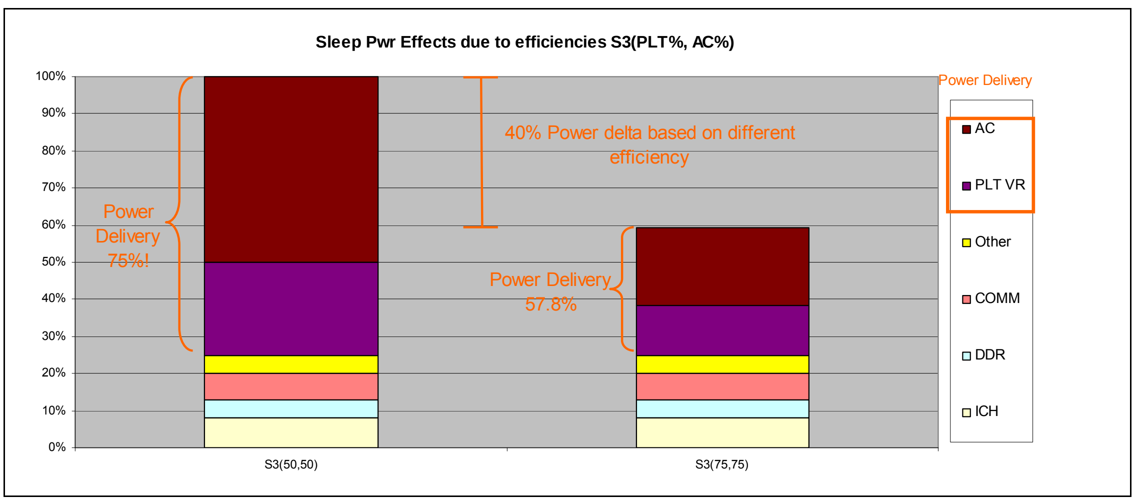

In this case we see that the power varies by 40% between the two configurations where in the first configuration power delivery contributed to 75% of the total power, while in the second case it dropped to 58% of the total power.

As can be seen, power delivery and the efficiencies of the components have a major impact on platform power (both idle and sleep) and will greatly impact the platform's ability to meet the ENERGY STAR requirements.

Another area that will greatly affect platform power is altering the load through various power management features. As an example to illustrate this point we show the same system in an idle ENERGY STAR configuration where we change the default CPU idle state.

The computer industry uses C-state terms to describe the activity of the CPU. When it is executing it is said to be in a C0 state, and while it is idle it can be in one of the many Cx states (C1, C2, …) where the higher the number, the lower the power and the higher the exit latency (lower performance). These states are entered and exited hundreds of times each second.

Why have multiple C-states? Each C-state trades off lower power for higher exit latency; so the system is tuned to have the best balance of power and performance (C-state exit latency) for a given configuration.

For this example we are going to illustrate the differences in platform power by varying the default C-state: in the first case we will show the higher power C2 idle CPU state, while in the second case we will show the lower power C3 idle CPU state. As can be seen, changing the default C-state from C2 to C3 resulted in a 29% platform power delta on the same platform (see Figure 11). These low power states not only affect the CPU, but you can notice how the power in the GMCH, DRAM and power delivery also changes.

**Figure 11. Idle Effects of Low Power States**

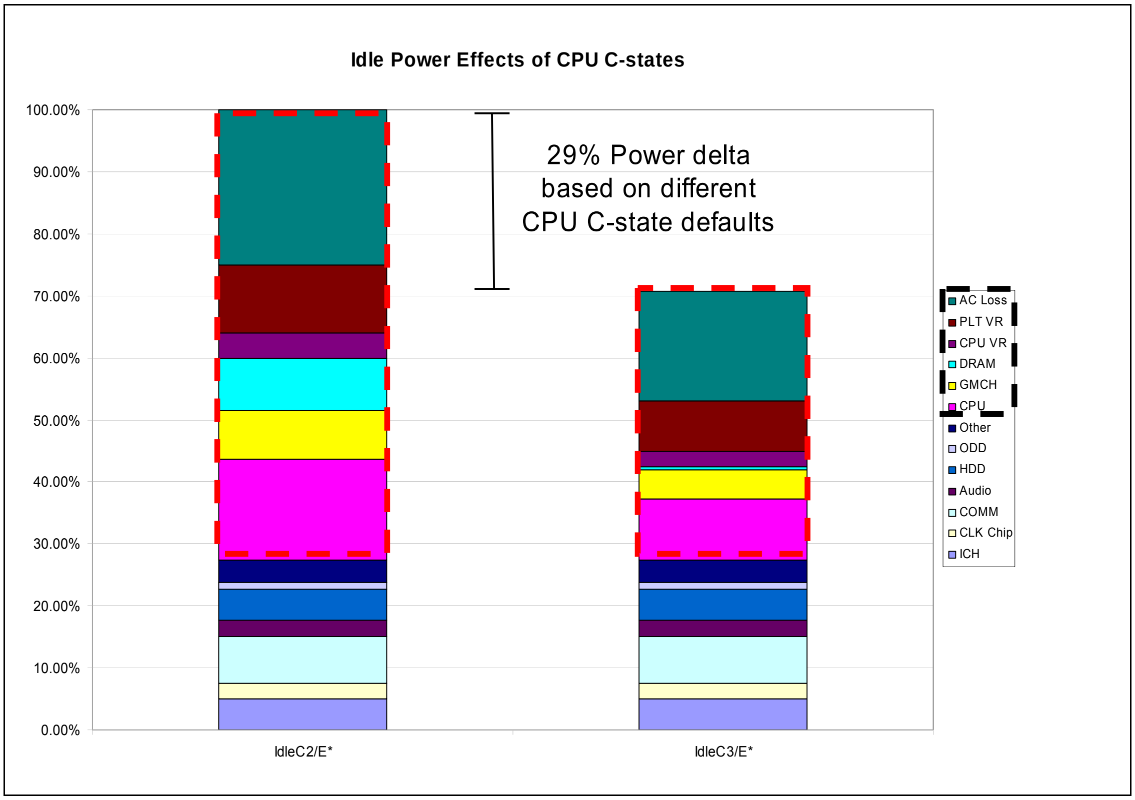


## 3.8 Optimizing a Notebook for ENERGY STAR*

We have illustrated how AC power is affected by a number of different platform factors. The most important factor is power delivery; and given that most platform and CPU power delivery designs are already optimized for battery operation (and are efficient at operating loads) the AC brick should be closely scrutinized for efficiency over the ENERGY STAR loads of interest. Some key recommendations for OEMs are:

- Use AC bricks with high efficiency in off mode. (@0.0~1.0W output power)
- Make sure the power rails for WOL is turned off in off mode.
- Minimize the leakage in rest-of-platform.
- Optimize the 3.3V voltage regulator for high efficiency in off mode.

The next area to examine is battery power management features which are traditionally disabled when running on AC and configuring the platform to enable these under the AC mode of operation. As an example, there is a large power reduction by using C3 as the default C-state under AC mode versus a C2 state.

If the power is still high in idle, then the next step is to look for crying babies. This is a term used to describe devices which do not rest when idle, and in the process wake everyone else up (like a baby crying at night, if the baby sleeps through the night everybody sleeps, but if the baby cries every 15 minutes nobody sleeps). A crying baby device which generates activity when idle will keep all of the other subsystems out of their low power states (memory, busses, clocking etc). And as we discussed, this additional power load then gets amplified by the power delivery system.

USB* devices are classic crying babies which generate lots of bus traffic when the device is sitting idle (keeping the rest of the system busy). In general, integrating USB devices into the platform can result in a non-optimal design for power management. If you cannot avoid integrating a USB device, then insure the USB device can support a "selective suspend" mode such that when the device is idle its function driver can suspend the device and then the USB driver can shutdown the USB controller so it will not generate idle activity. Further, if you have having trouble meeting ENERGY STAR idle power requirements, check to see if you have an integrated USB device.

The type of Operating System (OS) can also affect idle power. The way the OS schedules work can influence the amount of time the CPU spends in idle and will affect the idle power of the platform. Additionally an operating system that schedules daemons to pop-up and do work when the system is idle will also affect idle power. Having one of these pop-up just prior to your ENERGY STAR idle test could be problematic and is the reason why the system is allowed to sit idle for up to 15 minutes prior to actually testing the system and the power is averaged over 5 minutes. Also, a provision has been made for the tester to go through the first boot process and bring the system up to its normal use state, further reducing the chance of this impacting the idle test.

Not all hardware is the same either. Different CPU's, chipsets and devices will have different power attributes and support different power management features, which can dramatically reduce the load when the system is idle. Again as illustrated earlier, a power management feature that can save 1 W of power will translate to 1.7 W (or so) of power at the AC input.

For example, notebooks with Intel Celeron (high power) processors can still meet E-Star version 5.0 unless OEMs use the bad combination of devices for LAN, Memory, Clock, HDD, Audio, Fan, VRs, and AC brick.
Recommendation for OEMs (to meet ENERGY STAR Version 5.0 with Intel® Celeron processor)

- Use Intel Celeron processors based on Intel's 45nm process technology (Vs Celeron processors based on Intel's 65nm process technology)
- Enable the power saving features provided by Intel
- Do not use high-power components for LAN, Memory, Clock, HDD, and Audio
- Optimize the power delivery efficiency for VRs and AC brick

With Sleep and Off, again the majority of the power budget is dedicated to power delivery (typically over half). The AC brick should be optimized for these targeted loads (1 W – 2 W). Again you cannot rely solely on the ENERGY STAR External Power Supply specification for external bricks as they do not account for these low loads (idle is marginal), so the first action is to optimize the AC brick for these low loads, and then minimize the load for sleep and off. In many cases an external power supply that has better efficiency than is required for the ENERGY STAR specification is needed.

Low power technologies are always being created and any paper on low power system design should cover some of these new technologies. One such technology is the field of large non-volatile hard disk drive (HDD) caches. The concept is to place a fast (compared to the HDD speed) non-volatile memory, i.e. NAND FLASH, between the CPU and the HDD, and then store the most often accessed HDD into the cache. Because the memory is faster, the overall platform performance is increased, and because the HDD is utilized less (the majority of accesses now occur to this FLASH), the power of the system becomes lower as the HDD can remain much longer in its low power state. There are few technologies that both speed up the system and lower its power; so this is a great technology: mobile systems will want it to decrease power consumption, while higher-end systems will want it just to get the performance boost. Please note that ENERGY STAR requires HDD spin down to be disabled during the idle test unless the HDD is a hybrid drive as described above.

How much does this technology boost performance and lower system power? Performance will vary with application, but the HDD is considered one of the I/O bottlenecks of platform performance. Access times to HDD memory can be measured in milliseconds (say ~8ms), while accesses to FLASH memory can be measured in hundreds of nano-seconds (say ~120ns); hence FLASH is roughly 10,000 times faster than the HDD. Users won't see this speed-up, as overall performance is much more complicated (certain percentage of HDD accesses, cache misses, etc…), but it certainly reduces the impact of HDD latency to overall system performance.

Additional performance can be gained by "pinning" commonly used applications within the FLASH memory. Basically think of the application remaining in the FLASH memory such when using/loading these applications the HDD is not used at all.

What is the power impact of these non-volatile caches? The HDD contributes about 7-10% of system power (our breakdown shows about 7%), as such the maximum upside is removing this power component. The majority of the HDD power is that needed to spin the disk; hence one area to look at is how does such a cache impact the ability for the system to spin-down the HDD? Ruston Panabaker of Microsoft published such a study in a WinHec '06 presentation[7] and a graph from this presentation is reproduced below:

**Figure 12. Percent Spundown Time Operating on Battery**

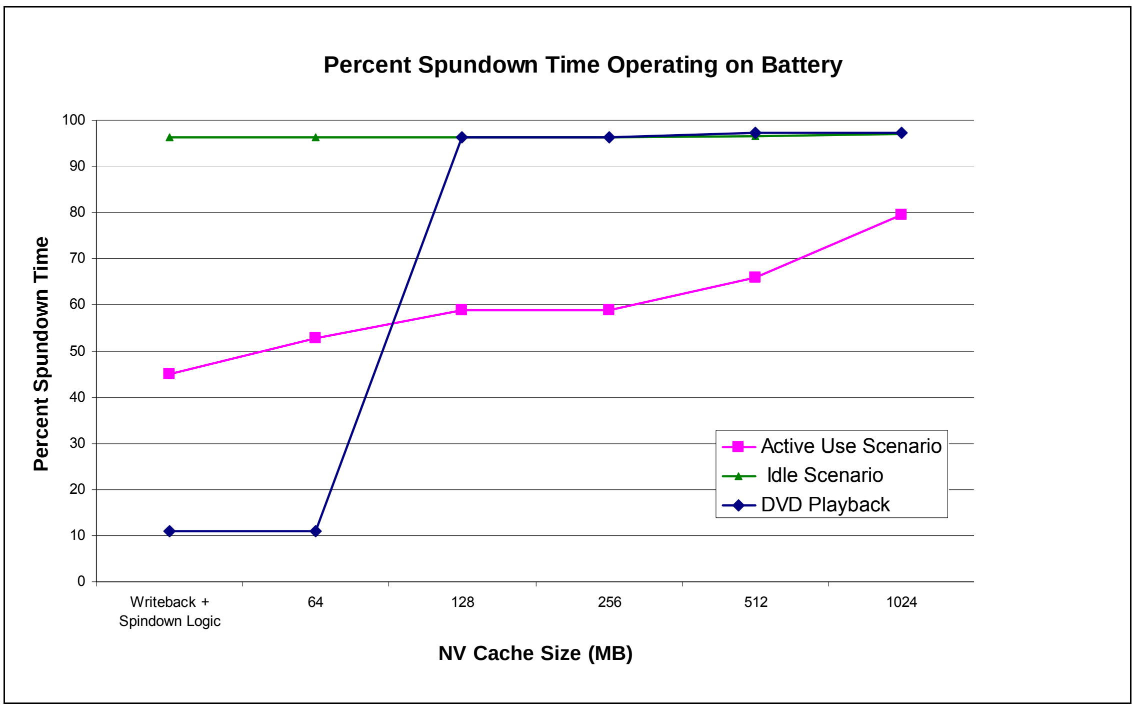


This graph[8] shows the percent of time the HDD spins-down based on the size of the NV Cache. Assuming that these caches will be 512Mbytes or larger, you'll see that an idle system would be spun down over 95% of the time. Further it shows that even running an

---

[7] **Hybrid Hard Disk And ReadyDrive™ Technology: Improving Performance And Power For Windows Vista* Mobile PCs** (download.microsoft.com/download/5/b/9/5b97017b-e28a-4bae-ba48-174cf47d23cd/STO008_WH06.**ppt)**

[8] **Hybrid Hard Disk And ReadyDrive™ Technology: Improving Performance And Power For Windows Vista Mobile PCs** (download.microsoft.com/download/5/b/9/5b97017b-e28a-4bae-ba48-174cf47d23cd/STO008_WH06.**ppt)**

active workload such as DVDE playback the HDD can be spun-down 95% of the time and during a very active usage the HDD could be spun down over 65% of the time.

While much of the benefit for power was originally targeted at notebook battery operation, the same technology will also help the AC power characteristics of any system using this technology. As discussed previously, this 7-12% of HDD power savings will then get amplified by the platform VR efficiency and further by the AC brick efficiency. Of course some of the best ways to save power in personal computers is to enable the monitor and system sleep states when the system is idle. Use of NVM to cache HDD represents just one of the background technologies the industry and system developers should be encouraged to develop to improve energy efficiency. These technologies are also not limited to notebooks or even desktop PCs. As described in subsequent chapters, workstations and enterprise class machines rely on increased amounts of bulk memory (HDDs). The HDD's power percentage of the system and hence, the energy savings along with performance gains can exceed 20% in these configurations.

The new ENERGY STAR specification requires OEMs to ship systems with power management features enabled. As mentioned previously, "a system is required to be shipped with the monitor enabled to sleep after 15 minutes of idle, while the entire system to is required to enter sleep after 30 minutes or less of idle". The monitor sleep mode will enable ENERGY STAR platforms to enter power states that bring the total system power to below required TEC levels.

## 3.9 Summary

Several choices in components and system configuration can greatly influence the power use of the computers you supply and the energy bills of your customers. This chapter outlined examples of these practices including:

- Tune efficiency of the AC brick so that it reaches its higher efficiency levels when the computer is idling. Avoid having idle fall at less than 25% of the AC bricks nameplate load.
- Ensure that system power management features (e.g. C-state selection) are tuned for both battery and AC operation, and are aggressive enough to meet these new ENERGY STAR AC power requirements.
- Ensure that the platform is enabled to meet ENERGY STAR shipping requirements. In particular that the monitor will enter sleep after 15 minutes or less of idleness and that the system will enter a sleep state after 30 minutes or less of idleness while on AC power.
- Ensure that devices are generating no activity when idle (avoid crying babies).

§

# 4 Desktop Computers and ENERGY STAR* Requirements

## 4.1 Category Definitions

In order for a desktop system to be eligible to meet ENERGY STAR it must first meet the definition of a desktop computer. A desktop computer is defined as a computer system intended to be located in a permanent location, i.e., desk or floor, and utilizes an external monitor, keyboard, and mouse. Within the ENERGY STAR requirements there are four categories of desktops systems:

Category A: All desktop computers that do not meet the definition of Category B, Category C, or Category D below will be considered under Category A for ENERGY STAR qualification.

Category B: To qualify under Category B, desktops must have:

- Equal to 2 Physical Cores; and
- Greater than or equal to 2 gigabytes (GB) of System Memory.

Category C: To qualify under Category C, desktops must have:

- Greater than 2 Physical Cores.

  In addition to the requirement above, models qualifying under Category C must be configured with a minimum of 1 of the following 2 characteristics:

- Greater than or equal to 2 gigabytes (GB) of System Memory; and/or
- A Discrete GPU.

Category D: To qualify under Category D, desktops must have:

- Greater than or equal to 4 Physical Cores.

  In addition to the requirement above, models qualifying under Category D must be configured with a minimum of 1 of the following 2 characteristics:

- Greater than or equal to 4 gigabytes (GB) of System Memory; and/or
- A Discrete GPU with a Frame Buffer Width greater than 128-bit.

The new specification defines Typical Energy Consumption (TEC) as a method of testing and comparing the energy consumption of computers, and focuses on the typical electricity consumed by a product while in normal operation over a year. The TEC approach calculates annualized energy consumption in KWhr by a formulae which weights the off, sleep and idle power of a system; where the weighting is based on the typical usage patterns over a year and is determined using the formula below:

$$E_{TEC} = (8760/1000) * (P_{off} * T_{off} + P_{sleep} * T_{sleep} + P_{idle} * {}_{Tidle})$$

Where all Px are power values in watts, all Tx are Time values in % of year, and the TEC $E_{TEC}$ is in units of kWh and represents annual energy consumption based on mode weightings. The formulae basically converts the "average weighted power" (Poff*Toff + Psleep*Tsleep + Pidle*Tidle) in watts to yearly Watt-hours by multiplying by the number of hours in a year (8760 hours/year) and converting from Watts to KiloWatts (by dividing by 1000).

Each of the ENERGY STAR categories contains TEC requirements, calculated based on AC power measurements at the wall plug for idle, sleep, and off power. Power supply efficiency and power supply power factor requirements need to be met as well. Finally, certain hardware-based capabilities are allowed defined capability adjustments in kWh ("adders") that are added to TEC limits when determining the TEC limit a computer must meet. An illustration of TEC calculation is provided in Appendix B of this document and a summary of these requirements are listed in table 6. For further details see the *ENERGY STAR Program Requirements for Computers: Version 5.0* specification.

**Table 5. ENERGY STAR* Requirements for Desktops**

| **Requirement** | |
|---|---|
| **Power Supply** | ≥ 85% eff at 50% of rated output and ≥ 82% eff at 20% and 100% of rated output<br>Power Factor ≥ 0.9 at 100% of rated output |
| **TEC (kWh)** | **Category A**: ≤ 148.0<br>**Category B**: ≤ 175.0<br>**Category C**: ≤ 209.0<br>**Category D**: ≤ 234.0 |
| **Capability Adjustments** | |
| **Memory** | 1 kWh (per GB over base)<br><br>*Base Memory*:<br>Categories A, B and C: 2 GB<br>Category D: 4 GB |
| **Premium Graphics** | Cat. A, B: 35 kWh (FB Width ≤ 128-bit)<br>50 kWh (FB Width > 128-bit)<br><br>Cat. C, D: 50 kWh (FB Width > 128-bit) |
| **Additional Internal Storage** | 25 kWh |

In addition to the above requirements, when a desktop system is shipped it must be delivered with the display's sleep mode set to activate after 15 minutes or less of user inactivity and with the desktop system sleep mode set to activate after 30 minutes or less of inactivity.

A number of requirements exist, regardless of distribution channel, for Wake On LAN configurations. For the specifics of the requirements, refer to Table 8 in the *ENERGY STAR Program Requirements for Computers: Version 5.0* specification.

## 4.2 Power Supply Considerations

ENERGY STAR has specific requirements for the computer system's internal power supply. In addition to these specific requirements, there are a number of aspects of the computer power supply that must be considered when building an ENERGY STAR compliant desktop computer system. These aspects of selecting a proper power supply will be discussed in this section.

### 4.2.1 ENERGY STAR* Requirements

The ENERGY STAR computer specification as documented in *ENERGY STAR Program Requirements for Computers: Version 5.0* requires that the internal power supplies for compliant computers be at least 85% efficient at 50% of rated output and 82% efficient at 20% and 100% of rated output. In addition, the power supply needs to have a power factor of at least 0.9 measured at 100% of the rated output capacity. In order to meet the power factor requirements, internal power supplies will need to incorporate active power factor correction.

### 4.2.2 Overall Efficiency Measurements

Desktop computer systems typically use internal power supplies with multiple outputs. These outputs are 12 V, 5 V, 3.3 V, 12 V, -12 V and 5 VSB (or 5 V standby). Power supply efficiency for the multiple output power supplies is defined in terms of a sum of the individual outputs. Equation 2 defines the efficiency for a single output power supply.

**Equation 2**

$$\eta_{PSU} = \frac{Pout}{Pin}$$

However for a multiple output power supply the efficiency definition is slightly different and shown in Equation 3.

**Equation 3**

$$\eta_{PSU} = \frac{\sum_{i} Pout,i}{Pin}$$

In many cases, the sum of the output power capability for the individual outputs exceeds the total output capacity of the power supply. For these cases, a method of proportional loading has to be applied to measure or calculate the efficiency of the power supply. This proportional loading method is explained in detail in the *Generalized Internal Power Supply Efficiency Test Protocol* which is available from www.efficientpowersupplies.org.

### 4.2.3 Power Supply Sizing

The ENERGY STAR computer specification's internal power supply requirements effectively provide a window of high efficiency that extends from 20% to 100% of the rated capacity of the power supply. Internal power supplies for desktop computers are available in a number of capacities that vary from approximately 150 W up to 1000 W or more. Generally the efficiency of the power supply drops off significantly as the load falls below 20%. Because the power supply is one of the largest contributors to power loss in the system, it is important to maximize the efficiency in order to comply with the power targets for the various system categories. Below are two case examples to illustrate this concept. For these two cases, assume the efficiency curve in terms of percentage of the output is equivalent. This curve is shown in Figure 13.

### Figure 13. PSU Efficiency as Percentage of Output

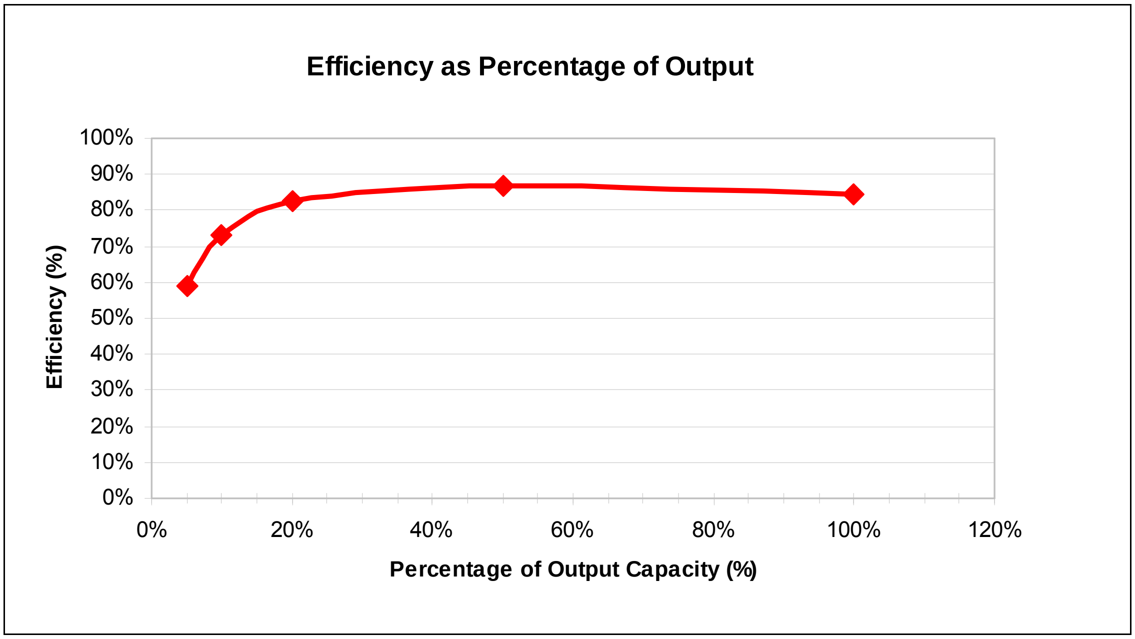


**NOTE:** Graph is an estimated representation for illustrative purposes only.

**CASE 1**

For Case 1, suppose the system we are trying to configure to meet the desktop category D has a 450 W internal power supply. The category D TEC specification is 234kWh (Say corresponding idle is 65W AC wall power). The AC wall power is calculated or derived as shown in Equation 4.

### Equation 4

$$AC(W) = \frac{DC(W)}{\eta_{PSU}}$$

In order to minimize the AC wall power, we need to maximize the power supply efficiency. As can be seen from Figure 14, for a 450 W power supply, the efficiency is low when the system is at 65 W AC wall power. Because the category D system at idle is on the part of the power supply efficiency curve where the efficiency is low, the remaining power budget for the other system components is much less than if the power supply efficiency was higher.

**Figure 14. PSU Efficiency per Power for 450 W PSU**

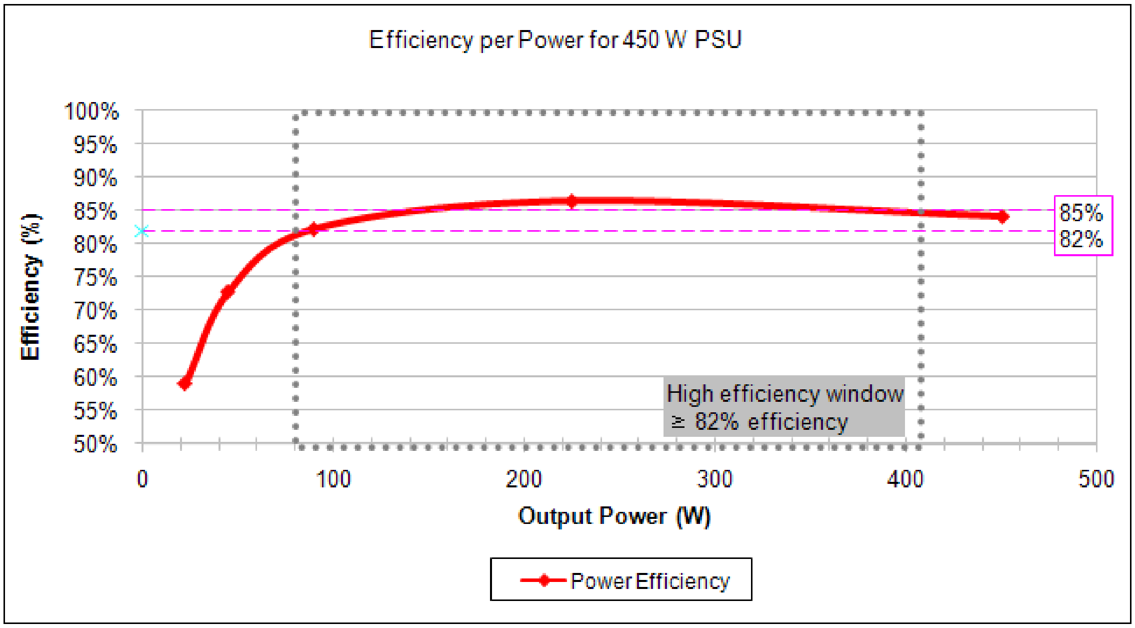


**NOTE:** Graph is an estimated representation for illustrative purposes only.

**CASE 2**

For Case 2, assume the system we are trying to configure for category D has a 350 W power supply. Assuming the platform power during sleep $P_{sleep}$ to be 3.5W and during off $P_{off}$ to be 1W, the idle power is 65W (using TEC equation). In this case, Figure 15 shows that for an idle state of 65 W AC wall power (Corresponding TEC of say 234 kWh), the system is within the high efficiency window of the power supply.

Since the system is operating within the high efficiency window of the power supply when in the idle state, the power supply losses are minimized which allows additional budget for other system components.

Figure 15. PSU Efficiency per Power for 250 W PSU

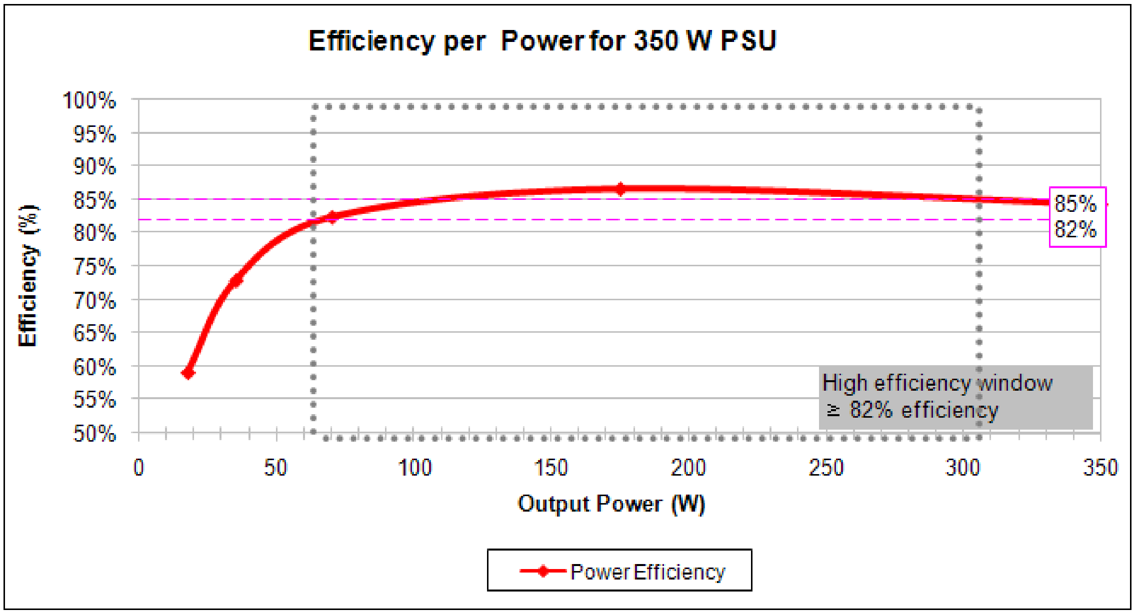


**NOTE:** Graph is an estimated representation for illustrative purposes only.

The examples in Case 1 and Case 2 above show that building the system with a power supply that is the proper size will maximize the power supply efficiency and allow the most flexibility for selecting other components in the system.

### 4.2.4 5 V Standby (5 VSB) Efficiency

As discussed previously, desktop systems typically have a multiple output power supplies. Of these multiple outputs, the 5 VSB (or 5 V standby) output is unique. This output is always present when the power supply is plugged into the AC wall outlet. The other outputs are off until the system powers on the entire power supply unit. Generally, the 5 VSB output is generated inside the power supply using its own circuit. It is this 5 VSB output that powers circuitry in the system when the system is in the off and sleep states.

Because the 5 VSB output is the only output present for the sleep and off states it is possible to measure the efficiency of this output independent of the other outputs. In order to meet the sleep and off targets for ENERGY STAR, the power supply will need to have good efficiency performance for the 5 VSB output. For the overall power supply efficiency measurements, the 5 VSB efficiency only has a small affect because of the relatively small current capability. Because of this it is possible to have a power supply that is ≥ 80% efficient overall and yet has poor 5 VSB efficiency. The *Power Supply Design Guide for Desktop Platform Form Factors* available at www.formfactors.org has guidelines for 5 VSB efficiency that can be used to guide the purchase decisions to ensure good efficiency for this portion of the power supply[9].

[9] http://www.formfactors.org/developer/specs/PSU_DG_rev_1_1.pdf

**Recommended 5 VSB Efficiency**

| Load | Efficiency |
|---|---|
| 100 mA | ≥ 50% |
| 250 mA | ≥ 60% |
| ≥ 1 A | ≥ 70% |

### 4.2.5 External Power Supplies for Desktop Systems

Desktop systems that use external power supplies should use ENERGY STAR qualified external power supplies or use power supplies that meet the external power supply requirements defined at www.energystar.gov/powersupplies.

## 4.3 Operational Modes

TEC requirements for desktops are calculated using the three operational mode power measurements present in Version 4.0: Idle, Sleep, and Off. As with the other platform testing, the power of the desktop is measured at the AC wall socket and is measured with an approved meter as described in Appendix A of the ENERGY STAR specification. The desktop test conditions for power mode testing are:

- Configure the desktop to prevent external monitors from powering down to ensure it stays on for the full length of the Idle test
- Configure networking to appropriate conditions by connecting the Ethernet port of the desktop to an active switch which supports the highest supported network throughput (typically a gigabit switch with today's systems) and turn off all wireless (radio) devices that might be installed
- Ensure that a standard mouse, keyboard, and external display are installed for all tests
- Primary hard drives may not be power managed ("spun-down") during Idle testing unless containing non-volatile cache integral to the drive (e.g. "hybrid" hard drives). 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 if it is enabled when the product is shipped

Other than what is mentioned above, all other configurable parameters of the desktop must remain in the "as-shipped" state. Appendix A of the *ENERGY STAR Program Requirements for Computers: Version 5.0*, contains detailed requirements for the power measurement equipment. In addition, a power meter that is capable of logging power data over time at a rate of at least one reading per second is useful for creating a record of power over time which can be used for more detailed analysis of the power consumption of the system.

Each of the system operational modes has different power states for the system components. Figure 16 through Figure 18 show examples of which elements of the system are powered in the various operational modes.

**Figure 16. Idle State System Power State Example**

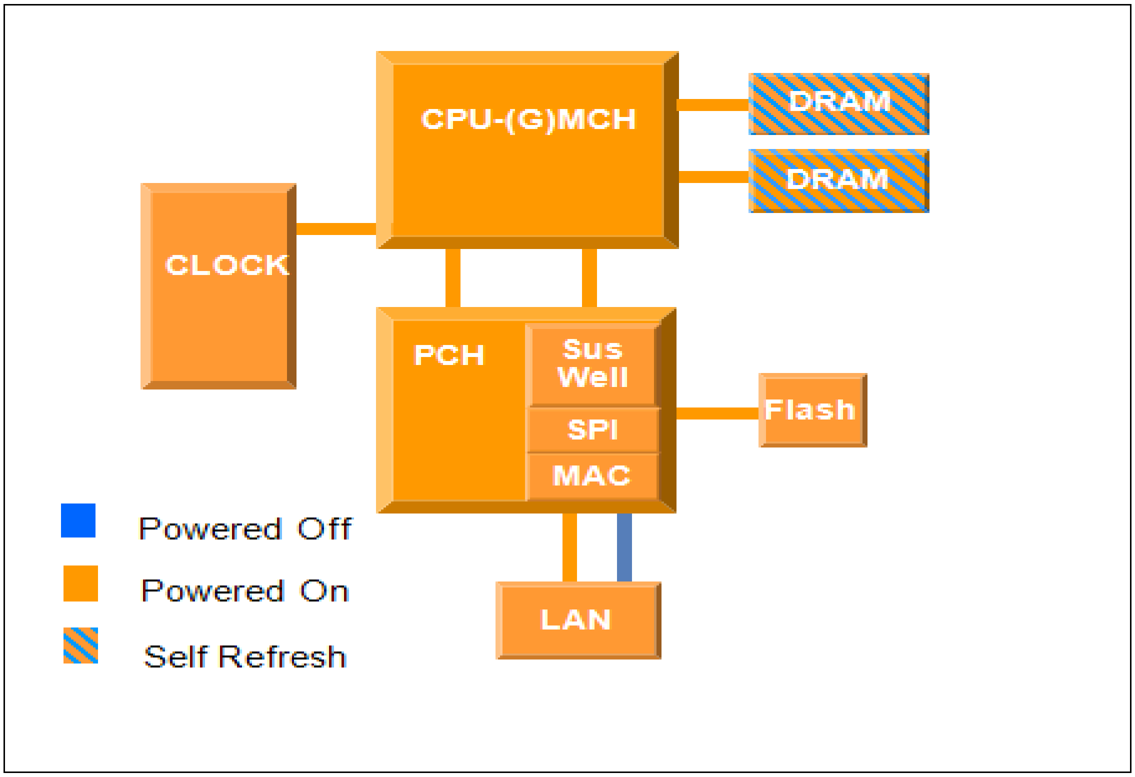

### Figure 17. Sleep (S3) Mode System Power State Example

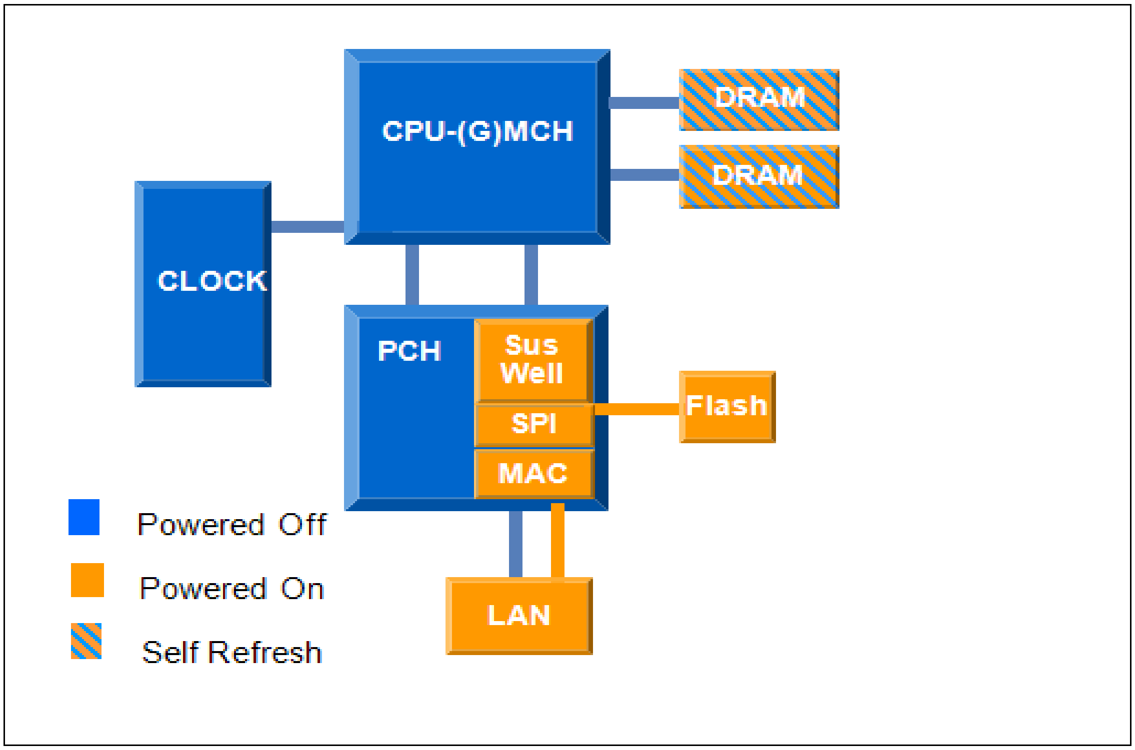


**NOTE:** Wake On LAN enabled.

### Figure 18. Off (S5) Mode System Power State Example

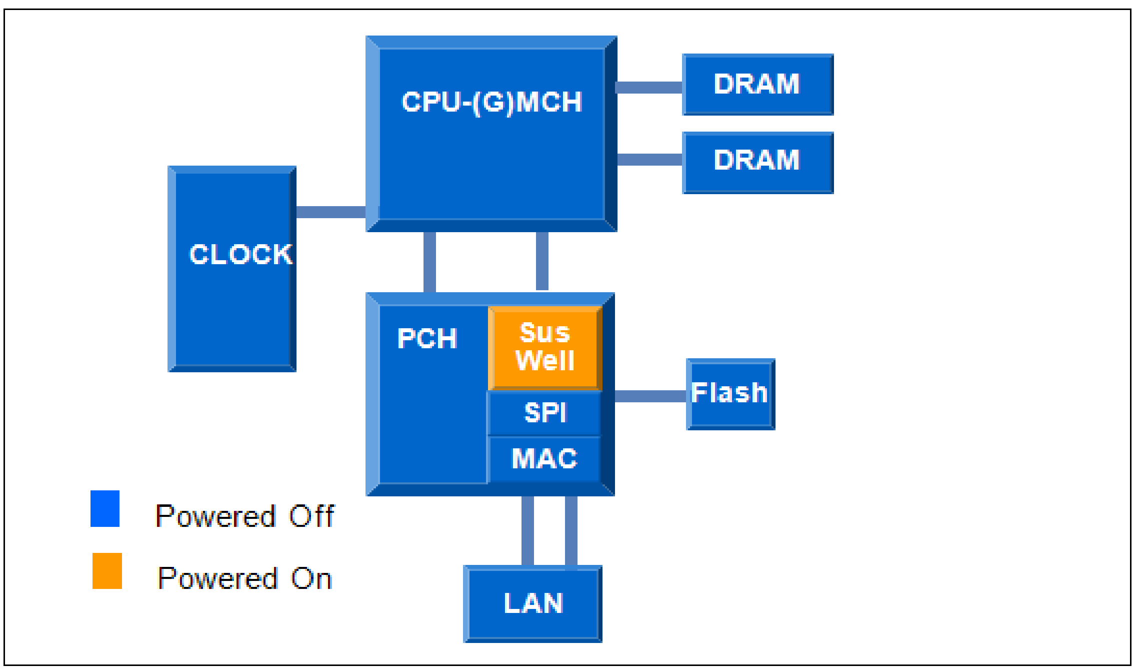


**NOTE:** No Wake On LAN

### 4.3.1 Idle State

Idle State is the state the computer system is in when it is fully powered on but not active and "the system responds to external events in real time"[10]. This mode is equivalent to the *Advanced Configuration and Power Interface (ACPI) Specification* G0 system state.

#### 4.3.1.1 Idle State Test Procedure

The Idle State power should be determined by performing the following steps:

1. Turn on the system and begin recording elapsed time either from the time of power-on or from the time of login
2. Once the operating system is loaded, close any open windows
3. After 5-15 minutes have elapsed from time of power on or time of login, either manually or through automation of the test equipment true power measurements should be made at a rate of 1 per second
4. Continue to collect power values for 5 minutes
5. The average of the 300 true power measurements is reported as the Idle State power level

#### 4.3.1.2 Factors that Affect Idle State Power

Version 5.0 of the ENERGY STAR requirements for computers sets an aggressive target for computer system energy consumption. Therefore, in order to configure a compliant system, the system components must be selected carefully. Motherboards, power supplies, processors, hard drives, voltage regulators, etc. should all be considered when integrating or designing a computer system.

Figure 19 is an example of how some of the components in the system vary in terms of their idle power consumption for Category A or B systems. In this example, if you configured a system using all components at the high power end of the range, the resulting system would not be compliant with the Category A or Category B requirements. On the other hand if you configured a system using all components from the low power range of the distribution, the system would be compliant with the Category A and B requirements. Use of components from the average range in this example would allow compliance with the desktop Category B requirements.

If the power is still high in idle, then the next step is to look for crying babies. This is a term used to describe devices which do not rest when idle, and in the process wake everyone else up (like a baby crying at night, if the baby sleeps through the night everybody sleeps, but if the baby cries every 15 minutes nobody sleeps). A crying baby device which generates activity when idle will keep all of the other subsystems out of their low power states (memory, busses, clocking …).

USB devices are classic crying babies which generate lots of bus traffic when the device is sitting idle (keeping the rest of the system busy). In general, do not integrate USB devices into the platform. If you cannot avoid integrating a USB device, then ensure the USB device can support a "selective suspend" mode such that when the device is idle its function driver can suspend the device and then the USB driver can shutdown the USB controller so it will not generate idle activity. Further, if you have having trouble meeting

---

[10] *Advanced Configuration and Power Interface Specification, Revision 3.0b*. Retrieved December 4, 2006 from http://www.acpi.info/spec.htm

ENERGY STAR requirements, check to see if you have an integrated USB device that might be affecting idle.

In addition to the tradeoffs that can be made between components for power, system designers need to consider performance, cost and feature tradeoffs as well when designing or building ENERGY STAR systems.

**Figure 19. System Idle Power Factors**

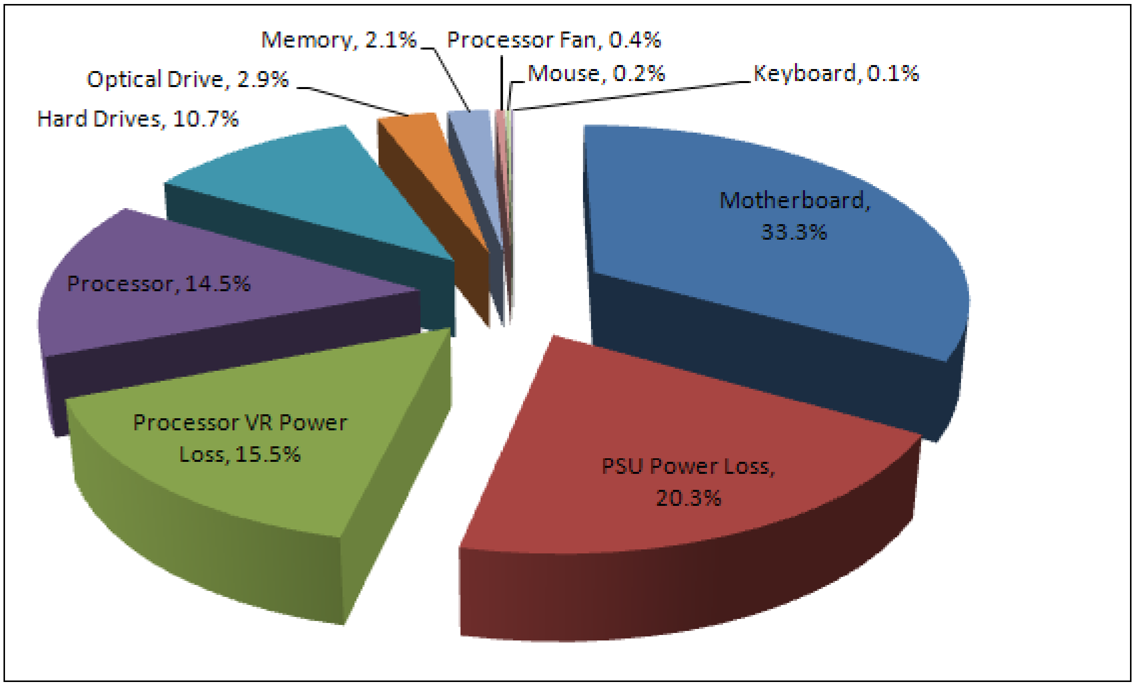


**NOTE:** Graph is an estimated representation (Average Power) for illustrative purposes only.

## 4.3.2 Sleep Mode

Sleep mode is the state the computer system is in when it is placed into the standby (Standby is the Windows XP nomenclature) state. This mode is equivalent to the *Advanced Configuration and Power Interface (ACPI) Specification* G1/S3 sleeping state. The operating system context is preserved in system memory for this state. The power requirements for Sleep Mode are the same for all three desktop system categories.

### 4.3.2.1 Sleep Mode Test Procedure

The Sleep Mode power should be determined by performing the following steps:

1. Place system in Sleep state
2. Either manually or through automation of the test equipment true power measurements should be made at a rate of 1 per second
3. Collect 5 minutes of values
4. The average of the 300 true power measurements is reported as the Sleep Mode power level

5. If testing both Wake On LAN (WOL) enabled and WOL disabled, steps 1-4 above should be performed for both WOL settings

#### 4.3.2.2 Factors that Affect Sleep Mode Power

As can be seen in Figure 17, when the system is in the sleep mode or S3 state, portions of the system are still powered. In this mode the system is powered only by the 5 VSB portion of the power supply so the efficiency of the 5 VSB portion of the power supply has a significant impact on the power measured at the AC wall plug. Refer to Section 4.2.4 for additional information.

For systems with Wake On LAN (WOL) enabled in the sleep state, powering the LAN to look for wake traffic consumes some additional power. Also, the link speed of the LAN in this state affects the magnitude of this additional power. For example, the difference between 100 Mb WOL and 10 Mb WOL is 485 mW for Intel® 82566 Gigabit Platform LAN Connect Networking Silicon.

Figure 17 indicates that the system memory is the self-refresh state in the sleep mode. Because of this the number of DRAM devices present and their associated self-refresh power impacts the sleep power of the system.

Other aspects of the system to consider when performing an analysis of sleep mode power or optimizing the system to reduce sleep mode power are the efficiency of motherboard regulators that are powered in this mode as well as the number and type of devices that are enabled and allowed to wake the system.

### 4.3.3 Off Mode

The off mode is the state the computer system is in when it is shutoff but still plugged into the AC power outlet. This mode is equivalent to the *Advanced Configuration and Power Interface (ACPI) Specification* G2/S5 or soft-off sleeping state. The operating system context is not preserved in this state.

#### 4.3.3.1 Off Mode Test Procedure

The Off Mode power should be determined by performing the following steps:

1. Place system in Off state
2. Either manually or through automation of the test equipment true power measurements should be made at a rate of 1 per second
3. Collect 5 minutes of values
4. The average of the 300 true power measurements is reported as the Off Mode power level

#### 4.3.3.2 Factors that Affect Off Mode Power

As can be seen in Figure 17, when the system is in the Off mode or S5 state, portions of the system are still powered. In this mode the system is powered only by the 5 VSB portion of the power supply so the efficiency of the 5 VSB portion of the power supply has a significant impact on the power measured at the AC wall plug. Refer to Section 4.2.4 for additional information.

For systems with Wake On LAN (WOL) enabled in the Off state, powering the LAN to look for wake traffic consumes some additional power. Also, the link speed of the LAN in this state affects the magnitude of this additional power. For example, the difference between 100 Mb WOL and 10 Mb WOL is 485 mW for Intel® 82566 Gigabit Platform LAN Connect Networking Silicon.

If Wake On LAN is not enabled or used in the Off state, implementing a hardware switch or some other mechanism to remove power from the LAN in the off state can save a significant amount of power.

Other aspects of the system to consider when performing an analysis of Off mode power or optimizing the system to reduce off mode power are the efficiency of motherboard regulators that are powered in this mode as well as the number and type of devices that are enabled and allowed to wake the system from the Off mode.

## 4.4 Summary

Several choices in components and system configuration can greatly influence the power use of the computers you supply and in the energy bills of your customers. This chapter outlined examples of these practices including:

- Match the capacity of the power supply and the power draw of the computer in “idle” mode so that it reaches its higher efficiency levels when the computer is idling (typically 40-75W). Avoid having idle fall at less than 20% of the power supply nameplate load.
- Ensure that CPU (C-states) and other component power management features are set such that the components run in less active modes
- Ensure that monitor power management and computer sleep settings (S-states) are enabled.
- Ensure that devices are generating no activity when idle (avoid crying babies).
- Consider tradeoffs between power consumption, performance, cost and feature when designing or building ENERGY STAR systems.

§

# 5 Workstations and Small-Scale Servers

ENERGY STAR criteria and metrics are meant to coexist and co-motivate the market to achieve the computing demanded in the most energy efficient manner. Workstations and servers can be classified as business critical computing devices; where throughput and capabilities are critical to the operation of that business. The business criticality is easily observed with the computer systems specifications and configurations commonly employed in this market. Indeed customers of such systems place a premium on these systems that is well beyond the value placed on personal computer requirements for these reasons. For workstations, the workload demands can be in either or both attended (user) and unattended (networked and little user interaction) environments. For servers, the uptime and reliability in an unattended operation serves as the computing backbone for those businesses. Specifically for workstations, operational responsiveness for applications such as CAD/CAE, financial simulations, engineering modeling, graphical or multimedia content creation, and scientific analysis can imply business operation success or failure. In the attended mode, business operation decisions based on these computing responses dictate the capability and capacity of the system. In the unattended or batch operation modes, these very same machines are in many use conditions relied upon to drive completion of project critical milestones, such as scientific, engineering or business modeling. Therefore, any efficiency criteria should comprehend the use condition to co-exist with the market incentives.

## 5.1 ENERGY STAR* Workstation and Small-Scale Server Specification Scope

The intent of the ENERGY STAR workstation and small-scale server specification remains to consider an energy metric which scales to the configuration and use demands of such systems under Version 5.0 while avoiding impact to the critical aspects for which the system was configured and purchased in the first place. By doing so, one develops a premium classification that both encourages energy savings and supports the varied, increasing capacity and capability requirements in these markets.

The Computer specification only addresses workstations and small-scale servers for this category of machines. Enterprise class servers are explicitly not part of ENERGY STAR for Computers, but in a separate specification development effort, ENERGY STAR for Computer Servers; accordingly, these products are outside the scope of this document. In addition to the classification difference for workstations and small-scale servers in the Version 5.0 specification, the approach and criteria applied for energy efficiency are de-coupled from those for desktops and notebooks. Requirements for small-scale servers remain stable with those under the previous specification (with a change to idle categorization). Requirements for workstations maintain the general structure of those in the previous specification, continuing to balance the scalable business needs and the energy consumed.

One key tenet on computing energy efficiency is the system’s ability to automatically migrate between power levels of operation. Indeed, this power management capability and the seamless transition in various modes are already market critical in both mobile and server environments. The unattended business criticality for server or enterprise environments, uptime requirements, and the fact that this environment is not self contained, drives increased challenges beyond a mobile computer system. From an energy efficiency standpoint, many technologies take advantage of the time-scale of operation that is to determine how long each component can be at a lower power condition thus enabling the technology's seamless operation. An attended activity is evaluated in 100's of milliseconds of operational time slice, to save energy, yet appear to be fully responsive; whereas unattended activity is sliced at a scale of 10's of microseconds. Response or the latency to wake-up from these lower power conditions without impacting the use condition is the key technical energy savings challenge. Conversely, simply offering a lower power condition without taking the wake-up response into account will be unlikely to "motivate" energy savings. Solely limiting the power level without accounting for power management transitions either constrains the efficient use of the system or increases the use of older non-energy efficient systems (e.g. purchasing more systems that do less work, but, have lower inactive power levels). Encouraging or providing incentives for the development of power management (and reduced wake-up latencies) was a significant consideration in the development of the ENERGY STAR specifications on workstations.

While ENERGY STAR for Computer specifications for workstations does comprehend some of these concepts, the workstation metric primarily targets the attended operational characteristics and the power levels of non-active states. Small-Scale servers are also included in the specifications and focused on its similarity to a standard desktop personal computer.

## 5.2 Small-Scale Servers and Specification Application

Small-Scale servers are so aptly named, driven by use of a desktop personal computer chassis to possibly provide both attended operation but, most definitely provide a secondary level of networked and business services (e.g.; print and file services, simplified mail services). The secondary level of service increases the use of the personal computer and minimizes use conditions such as sleep. Generally such devices, despite some being used 24x7, can tolerate rebooting or significant standby conditions where the response time is not business critical. Once the use conditions and business requirements drive a critical need for continuous- availability, enterprise class machines may become a more suitable option than small-scall servers, with advantages including redundancy and fail-over features at various levels to support such expectations. Small-scale servers may remain a viable option for non business critical or periodic applications, with the potential advantage of minimization of the facilities or infrastructure investment required for enterprise-class implementations.

The requirements for small-scale servers remain aligned with those present in Version 4.0 for all desktops. As noted, the user expectations in this category of product reflect its use as a standalone attended computer, as well as a standalone system using the network to distribute services. For the ENERGY STAR limits, testing methods for Small-Scale servers are the same as for the desktop category, to determine Idle and Soft Off power levels. Since the expected operation to provide 24 x7 localized storage and communication service precludes long term sleep in this usage model, small-scale servers are exempt from the sleep requirements. Technologies however, at sub-application visibility levels, that can exploit sleep-modes of the subsystems to enable lower idle power or faster standby recovery, are highly encouraged.

Small-Scale Servers are defined in Section 2 of this document (*ENERGY STAR and Computer Platforms*). The definition includes:

+ Pedestal, Tower form factor

+ Capable of operating in a simultaneous multi-user environment serving several users through networked client units; and

+ Designed for an industry accepted operating system for home or low-end server applications (e.g., Windows Home Server, Mac OS X Server, Linux, UNIX, Solaris).

+ High reliability, high availability, 24/7, 7d/wk, low unscheduled down time[11]

The Version 5.0 Specification evaluates small-scale servers using the same operational mode requirements from the previous specification. Category C no longer exists for this product area, and the two categories for small-scale servers are defined below, with requirements summarized in table 7:

Category A: All Small-Scale Servers that do not meet the definition of Category B will be considered under Category A for ENERGY STAR qualification.

Category B: To qualify under Category B Small-Scale Servers must have:

- Processor(s) with greater than 1 physical core or greater than 1 discrete processor; and
- Minimum of 1 gigabyte of system memory
-

**Table 6. Small-Scale Server Requirements**

| **Small Scale Server Operational Mode Power Requirements** | |
|---|---|
| **Off Mode**: | ≤ 2.0 W |
| **Idle State**:<br>Category A:<br>Category B: | <br>≤ 50.0 W<br>≤ 65.0 W |
| **Capability** | **Additional Power Allowance** |
| Wake On LAN (WOL)<br><br>(*Applies only if computer is shipped with WOL enabled*) | + 0.7 W for Off |

## 5.3 Workstation Category and Definition

Workstations are task critical computing systems that rely heavily on operation response (speed), reliability, large dataset manipulation, and rich data creation or manipulation. Depending on the application employed, the configuration and feature demand changes. Some businesses require high end graphics and rendering, while others require financial,

[11] The description on Small-Scale Servers is considered aggressive with regards to the high reliability and high availability comments. As noted in the general category descriptions in this paper, though used continuously in unattended tasks and at times 24x7, Small-Scale Servers tolerate more down-time than enterprise class servers, are not employed as business critical solutions, and are not designed to be part of multi-system integrated compute structure ensuring high reliability or high availability or application scalability. These areas are generally why full-featured enterprise class solutions are employed. The energy significance is that to support business critical needs on high reliability, availability and scalability, configurations and system architectures provide such as items such as redundant power supplies, resource fail-over, and intersystem interconnect and controls.

business, or scientific modeling. A workstation must scale to the capacity and performance demanded. The use condition also varies from critical compute response in attended applications, to being part of a net_batch pool of systems tasked with large scale unattended computing. As a result of these use conditions and incremental features and cost; workstations are driven to maximum configurations and maximum effective use of the system's features.

The workstation definition highlights both the minimum and optional list of features to constitute a workstation. The definition highlights the features which support the compute intensive, high capacity usage model of these systems. If the system does not lend itself to be classified as a workstation, one can apply as an ENERGY STAR compliant desktop system under the various categories A, B, or C. The criteria may also not address some mobile workstations. Some of these computing devices have inherently compromised the performance or features for the sake of mobility, cost or application criticality. The natural tradeoff will determine which category definition is best suited based on its configuration (use condition prioritization).

The workstation definition focuses on several areas: ECC or secure extensible memory, reliability, high end computing, high end I/O or graphics. These are broken into a mandatory set and optional list, reflecting the various configuration differences within a workstation class machine.

Workstation criteria:

- Marketed as a workstation
- Has MTBF of at least 15 k-hrs
- Support ECC and / or buffered memory

Meet three of the following six optional characteristics:

- Support power for high end graphics (i.e., PCI Express* 6-pin 12V power feed)
- System is wired for greater than x4 PCI Express on motherboard in addition to the graphics slot(s) and/or PCI-X support
- Does not support UMA graphics
- Includes 5 or more PCI, PCI Express or PCI-X slots
- Capable of MP support for two or more processors (must support 2 or more packages/sockets)
- Be qualified by at least ISV product certification, these certifications must be completed within 3 months of qualification

## 5.4 Workstation Scalable Specification

The workstation criteria were redesigned to scale with the capabilities of the system, thus supporting the premise of co-existing or co-motivation of the market demands of the system with the energy efficiency targets. The metric is known as a scalable version of the Typical Electricity Consumption (TEC)[12]. Although accurate energy efficiency would reflect energy consumption to accomplish a set of typical tasks, the active mode task suites vary, making the active-mode aspect difficult to define in the current ENERGY STAR specifications. The assumption in focusing on the inactive portions to represent TEC is that the market forces

[12] The EPA document refers to this as Typical Electricity Consumption, and should be considered synonymous in this context

would motivate the performance side while the ENERGY STAR specification or use of TEC could focus on energy use in inactive states: Idle, Sleep and Off. With an indeterminate workload (tasks) or work cycle defined, weighting factors used in the current TEC calculations currently reflect a balance of power management features the existing architectures and platforms contain. Both the active workload and the weightings to accurately represent these "Typical Electricity Consumption" modes are expected to be considerations in identifying an efficiency benchmark for the next revision of the ENERGY STAR specification.

**Figure 20. Scaleable Typical Electricity Consumption**

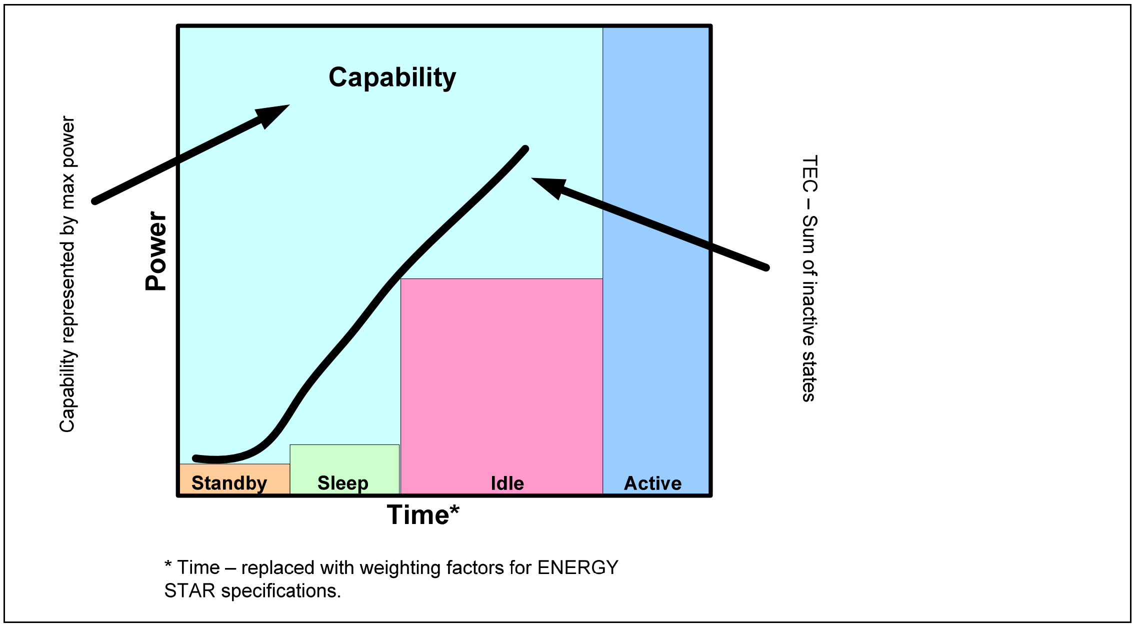


**NOTE:** Simplified diagram of scalable Typical Electricity Consumption index based on system power levels

The additional challenge is to scale the metric or specification based on the capabilities configured on the platform. Given the variety of features and innovations the industry continues to produce a simplified proxy of power to reflect capability is employed. The resulting metric becomes an efficiency index that compares the "typical energy consumed" against the "maximum capability configured in the system". As the compute demand and capability of the computing devices increases, the definition of an energy efficient system is to keep the typical energy assumption low relative to that peak capability. Scaling in this fashion promotes energy efficient computing without trying to dictate the features, purpose or work accomplished. Maximum power and the scaling provisions that enable power to serve as a proxy for capability is discussed later in this document. The current ENERGY STAR version of this metric, though limited in several areas, such as active mode consumption and work-cycle representation, does effectively scale to the various market demanded system configurations and encourages both the implementation and seamless operation of low power states. The simplified illustration of the scalable version of TEC demonstrates the accumulation of inactive power states as a ratio to "capability" of the system defined by the peak active power state. This efficiency index method is also referred to as Ratio Typical Electricity Consumption, or RTEC.

Specifically, the current ENERGY STAR limit for workstations is

- $P_{TEC} \leq 0.28*[P_{max} + (\# HDD * 5)]$
- $P_{TEC}$ is the sum of: 55% idle power + 10% sleep power + 35% off power.

The weights used on the power levels are not a direct reflection of the use condition of these systems. The weights are first to balance the inactive power consumption employed in the various architectures and system configurations. The adjustment in the formula from version 4.0 was to more closely reflect or prioritize the inactive modes that are more often employed on a workstation in operation. The cumulative value not only focuses the effort on reducing power in all inactive modes, but, encourages power management technology development to allow tradeoff in power consumption across various system states. For example, in developing higher I/O capability, the technology is encouraged to strive for lower power consumption waiting (lower idle power) and yet sufficient energy to wake-up (increase sleep power) and be ready for operation. Previous methods of pre-allocating budgets assume one already understands both the system's performance demand and sleeping capabilities. Attempting to predict technology growth, and task critical demands with fixed limits in each category limits technology growth and does not motivate industry investments in energy efficient performance innovations (e.g. getting more work done for the energy consumed).

Maximum power is determined or demonstrated by applying activity benchmarks which reflect the maximum capability of the system in terms of power. Even with well established benchmarks such as MIPS per watt, a key issue was whether MIPS reflected the "capability" demanded by the system/market. One could argue that availability, reliability, database size, graphical response, or modeling regressions can eventually be synthesized to MIPS and hence MIPS per watt; but, unfortunately, there is no translation of this market's value set to centralize to a single MIPS number. In other words, one could have a very good MIPS but, fail to be efficient in the application. In many cases use of this metric would result in burning more energy to accomplish the task even if the MIPS per watt were very efficient. Such a metric is, however, applicable for known fixed computational workload such as in scientific numerical calculations. Therefore, even though MIPS per watt could be an applicable metric in some cases, the variety of valued non-MIPS activities dictate the need for a different set of activities to generate a power proxy for a workstation's configured capabilities.

The activities or performance benchmarks that scale power levels well with the platform, CPU, memory, and graphics capability are LINPACK and SpecViewPerf (in concurrent operation). Both benchmarks are publicly available and supported across multiple platforms and architectures, key aspects for objectivity. Both benchmarks may have been tuned to demonstrate performance, which actually helps also track use of peak power consumption as a proxy for system capability. Do note that the use of these performance benchmarks is not to determine or measure performance, nor does it judge whether these are appropriate performance indices. However, the benchmarks do meet the scaling, power proxy, and public-availability objectives. From initial testing from various companies, these routines, run in parallel, exercise each of the major platform components to about a 70-80% level (except for HDD and network). These levels are purely a figure of merit as to how well they exercise the features and components in the system and their ability to scale with increasing capability in these areas. The percentage should not be used in place of actual measurements. Indeed, the main premise in using these routines is to establish the power proxy of the capability of the system. Use of key performance benchmarks is also intended to limit arbitrary increases in power, just to meet ENERGY STAR criteria. Market demands of lower power during peak performance will drive that end of the power consumption, thus producing the co-motivation aspect of energy efficient performance. To be sure, a power proxy for performance is merely a compromised solution until a performance or energy efficient performance metric can be developed.

*LINPACK* is a performance benchmark which solves a mathematical matrix array. The mathematical matrix-array is configurable to reflect the memory capacity, operational threads and CPU cores in the system. By adjusting these parameters in the routine (available as prompts in the routine published by Intel), one can appropriately exercise and hence represent in terms of power, the capability scaling of these platform components.

*SpecViewPerf* is a performance benchmark which targets both the GPU and its associated memory, while also loading other parts of the system. This graphics benchmark has been shown to be a reasonable exerciser to scale platform power levels with the graphics capability. One must set sufficient memory allocation to allow SpecViewPerf to run at maximum performance and power consumption in parallel with LINPACK. With insufficient memory to run the routine, excessive paging will occur, and the system will demonstrate lower power consumption than the routines separately. Experiments have found that in many cases 1-1.5MB is sufficient to maximize the SpecViewPerf run.

There has been several graphics developments that have evolved, one of which is use of dual graphics engines, either in combined or individually controlling displays. Since SpecViewPerf does not scale to multiple GPU engines, testing with a single engine enabled is sufficient for compliance testing.

A subcomponent of the system not heavily exercised by the routines above is the mass storage, i.e. hard disk drives (HDD). Though there are routines and benchmarks, such benchmarks for the HDDs cannot be conducted in parallel with LINPACK and SpecViewPerf. Therefore, an additive value per HDD is allocated in the power proxy for maximum capability.

## 5.5 Measurements

### 5.5.1 Max_Power

Maximum power on the platform is determined by running SpecViewPerf and LINPACK in parallel. Configuring each benchmark to ensure that they maximize the execution of various aspects of the subsystems is listed below. It is recommended that initial setups review a power profile of each routine to ensure proper testing durations and cycles are employed. The assessment will also confirm that the two benchmarks are indeed operating in parallel and register the power level caused by exercising these aspects of the system. The maximum power value is captured by the power meter sampling at a rate of 1 sample per second or greater.

#### 5.5.1.1 SpecViewPerf

SpecViewPerf revisions and application details can be obtained via the SPEC web sites. The specific instructions identify setting options which should reflect the configuration under test or evaluation.

#### 5.5.1.2 LINPACK-configurations

LINPACK binaries offered by Intel:

http://www.intel.com/cd/software/products/asmo-na/eng/perflib/mkl/266857.htm

General configuration instructions for max power testing are:

5. Determine the max memory installed in the system

6. Subtract the amount needed to run SpecViewPerf. and other system operations that would allow concurrent benchmarking operation
    a. This amount to subtract is highly dependent on the system architecture and configuration.
    b. Underestimating what is needed for SpecViewPerf and other operations, may cause LINPACK to run excessive paging of memory, resulting in a lower maximum power
7. Determine maximum matrix size (reminder: adjust the memory size per step 2) via this example script in Linux[13]:

This AWK script will calculate it on a Linux machine:

```
awk '
BEGIN {
printf "Maximum matrix dimension that will fit in RAM on this machine: "
}
/^MemTotal:/ {
print int(sqrt(($2*1000)/8)/1000) "K"
}
' /proc/meminfo
```

For those not running Linux the equation is:

$$\sqrt{\frac{(MemTotal)*1000}{8} \Big/ 1000}$$

- Identify the number of cores and threads feasible, based on the processor architecture
- Enter these values in the configuration prompts or header for the LINPACK run.

## 5.5.2 Typical Energy Consumption

The platform under evaluation must be configured as shipped. The power management settings when the system is shipped should be configured as documented in the ENERGY STAR specification. Some power management settings may conflict with the power measurement testing required for calculating of TEC and can be disabled during the TEC testing; however, they need to be enabled for shipping systems.

The descriptions of the various tests below solely highlight key considerations and changes to previous descriptions in this document. Users should reference the previous testing methods and considerations for the operational modes as shown in Section 4.3.

### 5.5.2.1 Idle

The power level should be evaluated at the respective operating system idle conditions. Of particular note would be to turn off prepackaged monitoring features, daemons, and other such activities not absolutely needed to have the O/S running. This should be a relatively steady state value. Any significant variations during initial test development may identify background activities that may need to be turned off in the outgoing configuration and while

[13] One needs to adjust the memory available to LINPACK per step 2 above, or the maximum power demonstrated by result in a lower value than running LINPACK on its own.

running this test. Such items should be end user accessible to be re-enabled, and documented prior to data submission.

#### 5.5.2.2 Sleep

Turn off wake-up timers or events. Management utilities that enable detection and migration through various power management states can also be turned off.

For those systems, whose architectures do not support a specific sleep state, the mode to use as a replacement is Idle. The formulae and weights for the TEC limits and calculations were determined accounting for this type of configuration.

#### 5.5.2.3 Off Mode (S5)

There are no workstation specific changes to the testing methods and considerations for the off mode as shown in Section 4.3.3.

### 5.5.3 Tools and Calibration

For simplified pass/fail criteria and for the data submission, the minimum meter accuracies are listed in Appendix A of the *ENERGY STAR Program Requirements for Computer* for testing equipment requirements. Many system manufacturers will however, want to better characterize and profile the power consumption across the various operating conditions and variety of sub-system components. Some tools will provide the ability to synchronize and register power and energy levels at time intervals across the test suites. Given the operational nature of the tests, such as determining maximum power, the difference between power and energy is likely a topic of investigation for a system developer. Providing documentation of anomalies as part of the dataset for ENERGY STAR compliance is recommended to ensure consistency to future compliance data gathering.

### 5.5.4 Power Management and Settings

As noted previously, sophisticated power management schemes are often employed to support peak performance, and run efficiently under both low loads and periods of inactivity. As the current testing addresses solely static conditions, i.e. maximum capability, Idle, Sleep, and Off Mode; some power management features could disrupt the determination of maximum capability, or waking up the system components during Idle, Sleep, or off. Therefore, the power management features should be enabled by default as shipped, identified in the data submission to the EPA, but may be turned off or disabled during the testing required for calculating the TEC.

§

# 6 Thin Clients and ENERGY STAR* Requirements

### 6.1.1 Server-based Client Compute Architecture

Thin clients are part of a client compute architecture commonly classified as a ***server-based computing model***. Thin Clients are primarily the remote display and human interface in this computing model. Most or all of the computing activity is supported by a central server. The types of services commonly deployed under a server-based compute model are Terminal Services and Virtual Hosted Desktops also called Virtual Desktop Infrastructure or VDI. All server-based computing models fully rely on both the network and servers to manage and run the thin-clients on the system. Thin clients contain limited compute capability and cannot operate in a stand-alone mode and are increased in capability to reduce the burden on the network, network bandwidth, and servers. As implied by this computing model, the operational and energy efficiency is determined by the complete structure, not solely on the thin client.

The efficiencies gained by a server-based computing model are the overall resource savings in the data center equipment (servers, network, storage, redundancy, etc…) and data center infrastructure (e.g. cooling, power delivery, and management services), in addition to the resource savings on the thin clients. A common metric for applicability is the ratio of number of clients supported by a given network and server capacity installed. This ratio of clients to server_network varies based on the local acceleration capability of the client and the compute burden to support each of the clients. The compute burden of Terminal Services (fixed application to each user) is much smaller than the compute burden of Virtual Hosted Desktops (VHD) [replicates a full O/S + PC environment] thereby supporting many more Terminal Clients than VHD's. Energy efficient options for applications that can utilize a server-based computing model should include the server and network resources, the compute burden per client, and ratio of client to server_network, in addition to the (thin) client efficiency.

Terminal Services is the "classic" server-based model. Here, the client is merely a display and input device. All computation is done centrally on the server and all data is stored in a datacenter. Nothing is persistent on the client. It is the most proven, reliable server-side model, harkening back to the days of mainframe computing. Remote Display Protocol (RDP) or Independent Computing Architecture (ICA) are used to deliver an image of the server-based application to a terminal viewer on the client, and return keystrokes and mouse clicks to the server.

## Figure 21. Architecture - Terminal Services

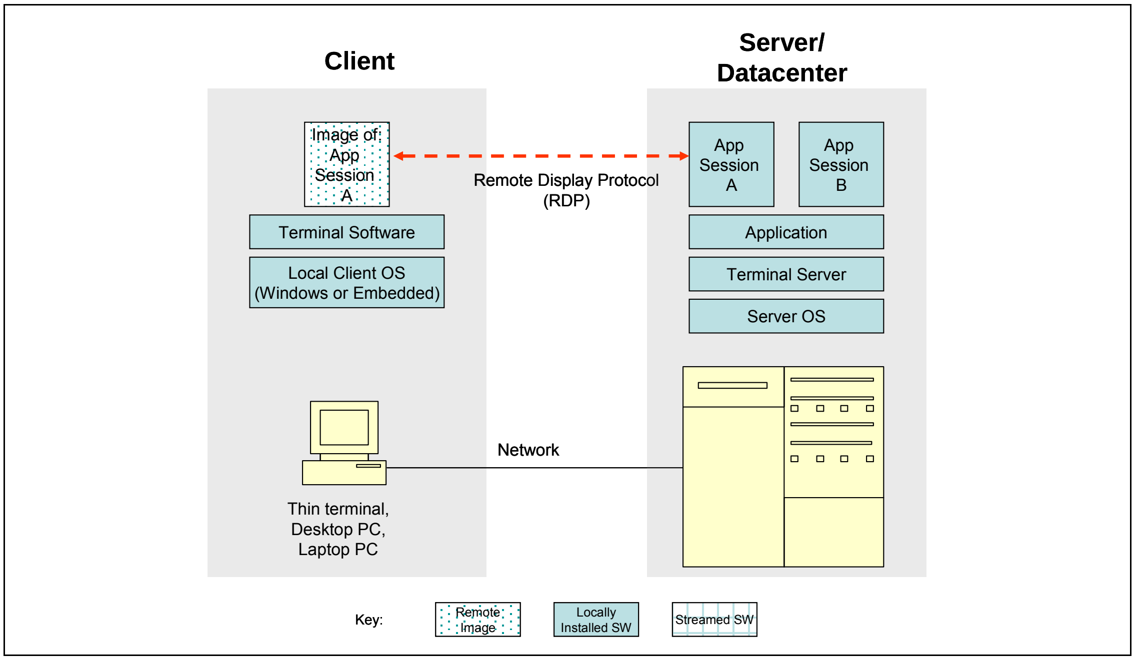


Terrminal services are commonly deployed where:

+ users are in a fixed location,
+ constant high speed, high bandwidth network access is available, and,
+ Workloads are highly deterministic, compute-lite, and graphics-lite.

A ratio of a large number of clients with limited workloads to servers make this compute model efficient where compute requirements are low, mobility is not required, and independent operation is not needed. Typical applications include: bank tellers accessing the transaction system, call center workers entering orders in database and healthcare professionals working with text-based patient records.

A recent entry to server-side computing is Virtual Hosted Desktops (VHD), a generic term encompassing commercial products such as VMware's Virtual Desktop Initiative (VDI), or Citrix's XenDesktop Server. For VHD, all computation and storage are centralized, with images of the user's desktop pushed over the network to the client via RDP or other protocol. The major difference is that VHD offers each user their own complete virtual machine, including the OS, applications and settings. VHD is designed to replicate the user experience of a rich PC with all the management and security of server-side models.

**Figure 22. Architecture – Virtual Hosted Desktop**

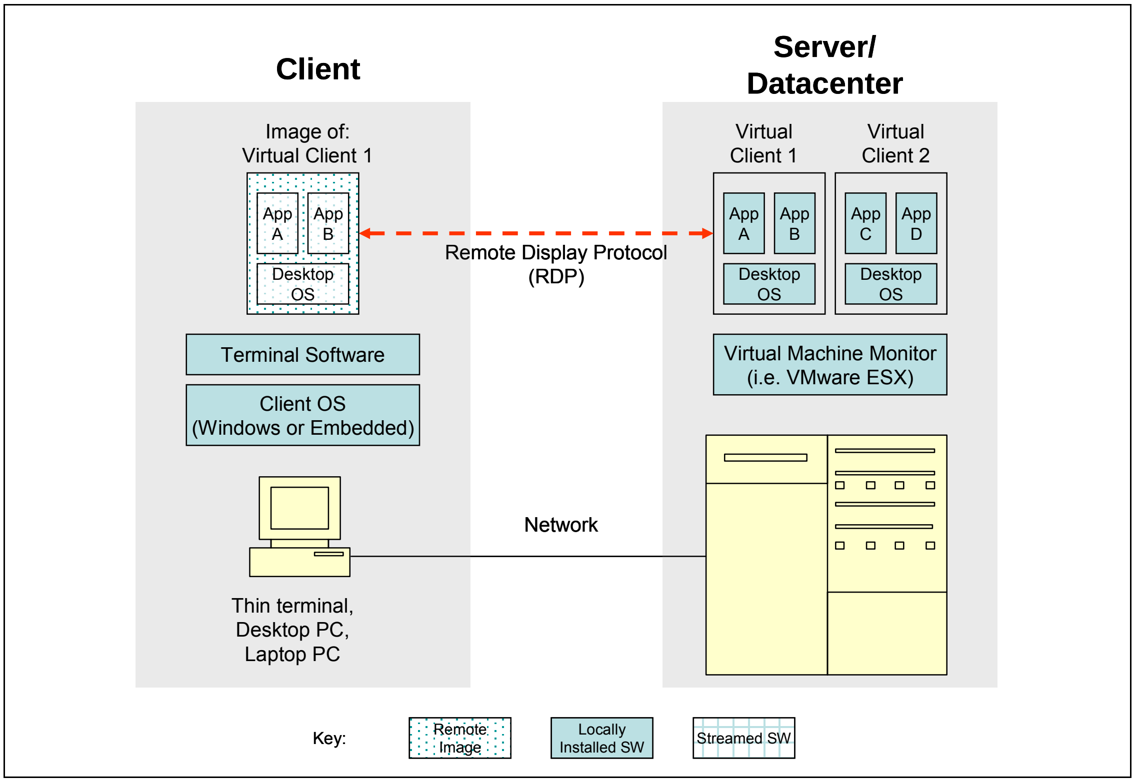


VHD is a relatively new technology, and suffers from the same limitations inherent in server-side models, specifically performance challenges as the number of users grows, graphics limitations, lacks stand alone operation and does not offer a mobile solution. Hosting multiple full virtual PC operation and interface places an increased burden on a high bandwidth, persistent network, and resilient compute services. Over time, these limitations may be addressed through innovation on both the server and network side; along with increasing capabilities built into the (thicker) thin clients. Indeed a number of higher capability thin client devices have been introduced with localized graphics rendering, rich media acceleration, and other capabilities. Accommodating higher capability thin clients may become challenging given the limits chosen for Energy Star. As noted in the category definitions, the EPA has defined a category B, recognizing the increased localized capability afforded by these machines.

### 6.1.2 Category Definitions

In order for a thin client to be eligible to meet ENERGY STAR it must first meet the definition of a thin client. A thin client is defined in the Version 5.0 ENERGY STAR specification as follows:

*Thin Client: An independently-powered computer that relies on a connection to remote computing resources to obtain primary functionality. Main computing (e.g., program execution, data storage, interaction with other Internet resources, etc.) takes place using the remote computing resources. Thin Clients covered by this specification are limited to devices with no rotational storage media integral to the computer. The main unit of a Thin*

*Client covered by this specification must be intended for location in a permanent location (e.g. on a desk) and not for portability.*

This definition and product category represents the first time these products can be qualified under the ENERGY STAR program. The definition above is intended to limit the scope of the products covered to "desktop replacements" without rotational storage media. "Mobile thin clients" and hybrid clients with integral hard drives are not covered under Version 5.0. Similar to the other product categories, requirements are applied to thin clients under a category structure.

Category A: All Thin Clients that do not meet the definition of Category B, below, will be considered under Category A for ENERGY STAR qualification.

Category B: To qualify under Category B, Thin Clients must:

- Support local multimedia encode/decode.

These categories determine appropriate levels for Idle State. All ENERGY STAR thin clients must meet requirements for Idle State and Off Mode; systems supporting a sleep mode must meet also meet a specified Sleep Mode requirement. Systems ship with Wake On LAN enabled qualify for a small power allowance in low power modes. These requirements are listed in 5. For further details see the *ENERGY STAR Program Requirements for Computers: Version 5.0* specification.

**Table 7: ENERGY STAR* Category Wall Power Requirements**

| **Requirement** | |
|---|---|
| **Power Supply** | ≥ 85% eff at 50% of rated output and ≥ 82% eff at 20% and 100% of rated output<br>Power Factor ≥ 0.9 at 100% of rated output |
| **Off Mode** | ≤ 2 W |
| **Sleep Mode** *(if applicable)* | ≤ 2 W |
| **Idle State**<br>**Category A:**<br>**Category B:** | <br>≤ 12.0 W<br>≤ 15.0 W |
| **Capability Adjustment** | |
| **Wake On LAN (WOL)**<br>*(Applies only if computer is shipped with WOL enabled)* | + 0.7 W for Sleep<br>+ 0.7 W for Off |

In addition to the above requirements, when a thin client is shipped it must be delivered with the display's sleep mode set to activate after 15 minutes or less of user inactivity.

§

# 7 Future Specification Considerations

The ENERGY STAR specification highlights the need to revisit the criteria in preparation of revisions. Highlighted in several areas in this document are concepts for criteria improvements to denote energy efficiency. The future revision serves as the opportunity to do so, and various industry organizations are current engaged by the EPA to study and propose evolving energy efficiency criteria. Pending the acceptability of these new metrics, the time frame also allows the EPA an opportunity to re-assess the ENERGY STAR limits themselves.

While advances in technology and energy saving techniques will likely educate future revisions to the ENERGY STAR requirements, potential areas that may be further pursued in future versions of the program include:

- Full benchmark approach integrated into TEC calculations for standard clients and workstations;
- Integrated display power consumption as part of ENERGY STAR metrics (and implemented equitably vis-à-vis external displays);
- Full sleep mode capability in thin clients.

The next revision of ENERGY STAR will incorporate the learning obtained via the compliance data submissions after the effective date of the current specifications. The current industry-wide organizations consulting with the EPA on an energy efficiency metric include ECMA*, BAPCO*, LBNL*, and SPEC.org*.

§

# 8 Summary

Both Intel and the EPA are each focused on growing adoption of energy saving technologies and products that can continue to improve usability and energy efficiency while also reducing energy consumption on the computer platform.

With the ENERGY STAR program for computers, the EPA has set the goal of generating awareness of these energy saving capabilities, as well as differentiating the market for more energy-efficient computers and accelerating the market penetration of more energy-efficient technologies.

This document provides some key insights into the delivery of desktop, notebook and workstation systems that are in a good position to meet the requirements in the new ENERGY STAR specification for computers.

§

# Appendix A Creating a Disk Image[14]

## A.1 Modifications to the Default User/policy in Rollout Image (i.e.; "Ghosting")

If rolling out Windows XP or Windows Vista through the use of images (i.e.; Using Symantec Ghost or similar) there is a method, lacking centralized control (fire-and-forget option), a system administrator or OEM can set and ensure compliance to monitor power management policies. When a user logs into a machine for the first time, assuming no roaming profiles in use, the default settings for the user comes from the default user account in the "Documents and Settings" (C:\Documents and Settings\Default User\ntuser.dat) folder in use by the system. Under most circumstances this is the case for power management settings[15]. By using regedt32.exe[16] and making the necessary edits all new accounts will now pick up these settings. To do this:

- Open regedt32.exe
- Highlight the HKEY_USERS branch.
- Under the Registry menu, click load hive and navigate to C:\Documents and Settings\Default User\ntuser.dat, or where ever the default user's profile is stored.
- Load that file into the hive naming the branch PMDefault.
- From here, manually change the power management settings for the currently logged in user, following the directions in A.2
- If regedit is already open switch to it and press F5 to refresh the current view. (Will export old settings otherwise)
- Highlight the key 'HKEY_CURRENT_USER\Control Panel\PowerCfg' and select "Export Registry file…" from the Registry menu in the menu bar.
- Export the file making sure that the selected branch option is selected.
- Name the file anything you would like and save it.
- Edit the reg file using a text editor changing the key prefixes (in every key entry) to 'HKEY_USERS\PMDefault\Control Panel\PowerCfg'.
- Go ahead and merge this reg file with the system by double clicking on it.
- From here, unload the PMDefault hive by highlighting 'HKEY_USERS\PMDefault ' and under the Registry menu, click unload hive.

---

14 Microsoft product screen shot(s) reprinted with permission from Microsoft Corporation.
15 There is a case where, if the last user to login changed their power management settings (with or without a reboot), and then a new user logs in next, this new user will pick up those settings.
16 Regedit.exe in XP or higher

## A.2 Microsoft Windows* XP

- Select Start > Control Panel from the Start Menu.

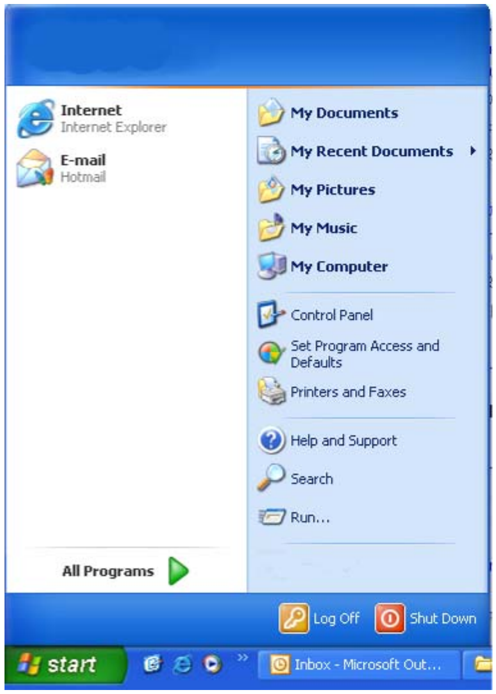


- Click the "Performance and Maintenance" icon in the Control Panel window then choose "Power Options".

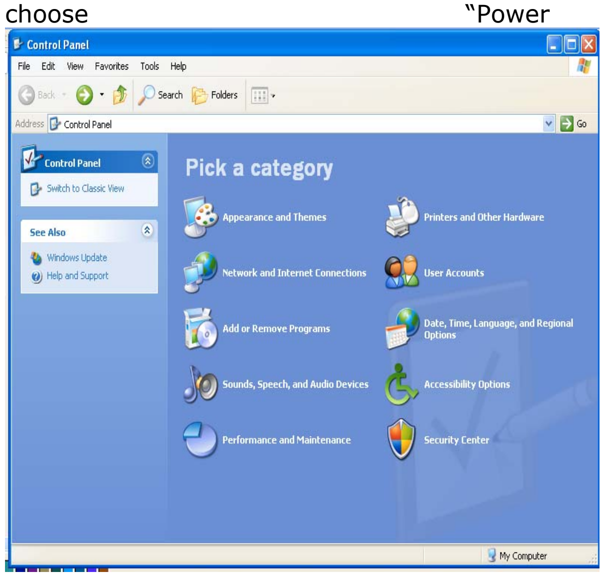

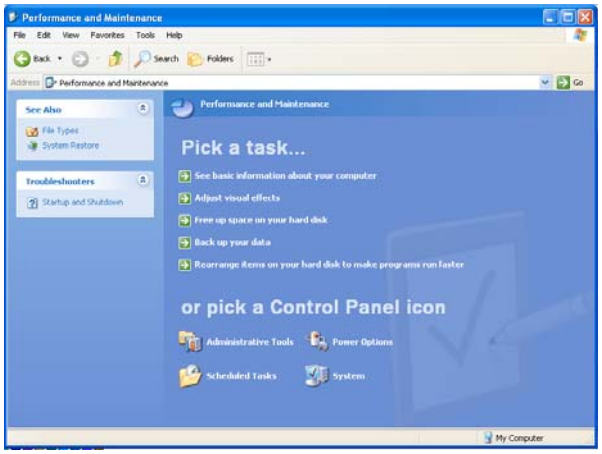


- Select the power scheme you would like to use as the default. This is normally "Home/Office Desk".


- Adjust the settings to the desired time out for monitor and standby. For ENERGY STAR this is 15 for the monitor and 30 for standby.
- Click Apply and OK to commit those changes.

## A.3 Microsoft Windows Vista*

- Select Start > Control Panel from the Start Menu.

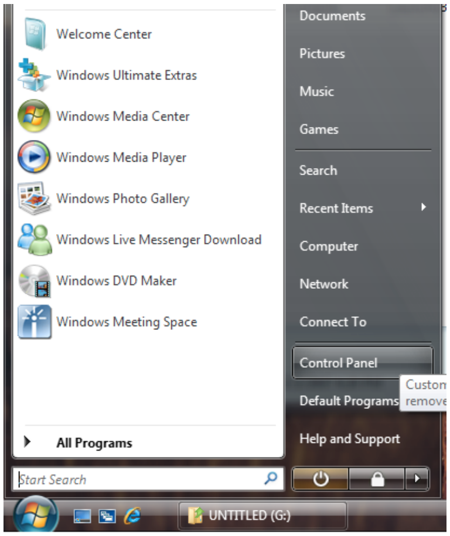


- Click the "System and Maintenance" icon in the Control Panel window then choose "Power Options".


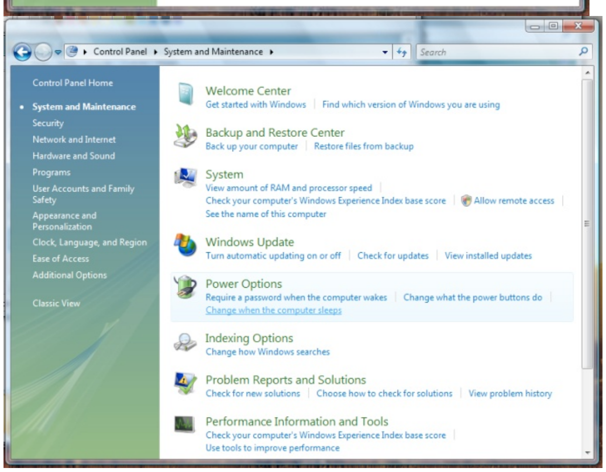

- Select the power scheme you would like to use as the default. This is normally "Home/Office Desk".

**Figure 23. Power Scheme**

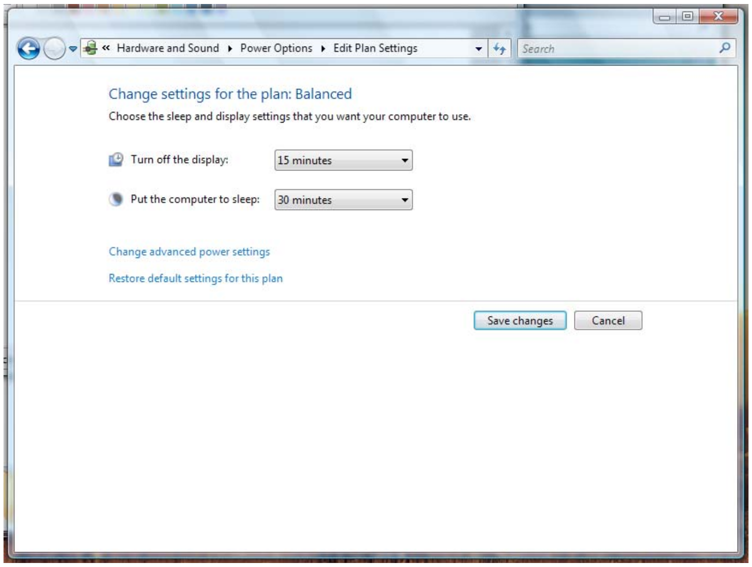


- Adjust the settings to the desired time out for the display and computer sleep. For ENERGY STAR this is 15 for the monitor and 30 for standby.
- Click Apply and OK to commit those changes.

§

# Appendix B Sample TEC Calculations

*This appendix is reproduced from the ENERGY STAR Version 5.0 Computer Specification.*

## B.1 Desktop, Integrated Desktop, Notebook Computers

Below is a sample TEC calculation intended to show how levels for compliance are determined based on functional adders and operational mode measurements.

**Example:** Below is a sample $E_{TEC}$ evaluation for a Category A Notebook Computer (integrated GPU, 8 GB Memory Installed, 1 HDD)

1. Measure values using the Appendix A test procedure.

   *Off =* *1W*
   *Sleep =* *1.7W*
   *Idle =* *10W*

2. Determine which Capability Adjustments apply.
   *Integrated Graphics? Does not apply for Premium Graphics.*
   *8GB Memory installed. Does meet memory adjustment level: 8 yields a 1.6kWh adjustment (4 * 0.4kWh).*

3. Apply Weightings based on Table 2 to calculate TEC:

   Table 2 (for conventional notebook):

| | |
|---|---|
| **Toff** | 60% |
| **Tsleep** | 10% |
| **Tidle** | 30% |

$$E_{TEC} = (8760/1000) * (P_{off} * T_{off} + P_{sleep} * T_{sleep} + P_{idle} * T_{idle})$$
$$= (8760/1000) * (P_{off} * .60 + P_{sleep} * .10 + P_{idle} * .30)$$
$$= (8760/1000) * (1 * .60 + 1.7 * .10 + 10 * .30)$$
$$= \mathbf{33.03\ kWh}$$

4. Determine TEC Requirement for the computer by adding any capability adjustments (step 2) to the Base TEC requirement (Table1).

   Table 1 (for notebooks):

| Notebook Computers (kWh) | |
|---|---|
| **Category A** | 40 |
| **Category B** | 53 |
| **Category C** | 88.5 |

ENERGY STAR TEC Requirement = *40 kWh + 1.6kW = 41.6 kWh*

5. Compare $E_{TEC}$ to the ENERGY STAR TEC Requirement (step 4) to determine if the model qualifies.

   *Category A TEC requirement: 41.6 kWh*
   $E_{TEC}$*: 33.03 kWh*
   *33.03 kWh< 41.6 kWh*
   ***Notebook meets ENERGY STAR requirements.***

## B.2 Workstations

Below is a sample $P_{TEC}$ calculation for a Workstation with 2 hard drives.

1. Measure values using the Appendix A test procedure.

   | | |
   |---|---|
   | *Off =* | *2W* |
   | *Sleep =* | *4W* |
   | *Idle =* | *80W* |
   | *Max Power =* | *180W* |

2. Note number of Hard Drives installed.
   *Two hard drives installed during test.*

3. Apply Weightings based on Table 4 to calculate $P_{TEC}$:

   Table 4:

   | | |
   |---|---|
   | **Toff** | 35% |
   | **Tsleep** | 10% |
   | **Tidle** | 55% |

$$P_{TEC} = (.35 * P_{off} + .10 * P_{sleep} + .55 * P_{idle})$$
$$= (.35 * 2 + .10 * 4 + .55 * 80)$$
$$\mathbf{= 45.10\ W}$$

4. Calculate the $P_{TEC}$ requirement using the formula in Table 3.

   $P_{TEC}$ = 0.28*[Pmax + (# HDD * 5)]
   $P_{TEC}$ = 0.28*[180 + 2 * 5)]
   $P_{TEC}$ = 53.2

5. Compare the adjusted $P_{TEC}$ to the ENERGY STAR levels to determine if the model qualifies.

   *45.10< 53.2*
   ***Workstation meets ENERGY STAR requirements.***

§

# Appendix C Energy Star* Version 5.0 and Version 4.0 Summarized

## Energy Star Version 5.0 Computers Criteria

| Platform | Power Supply | Requirements | Capability Adjustments |
|---|---|---|---|
| Notebooks & Tablets | No change | TEC (kWh)<br>Cat A: ≤40.0<br>Cat B: ≤53.0<br>Cat C: ≤88.5 | Memory: 0.4 kWh (per GB over 4 GB)<br>Discrete GPU (FB Width): CAT B 3kWh (>64-bit)<br>Additional storage: 3 kWh |
| Desktop, Integrated DT | 82-85-82% eff *(20-50-100% load)* + ≥0.9 PFC @100% rated output | TEC (kWh)<br>Cat A: ≤148<br>Cat B: ≤175<br>Cat C: ≤209<br>Cat D: ≤234 | Memory: 1 kWh (per GB over base)<br>Base Memory: CAT A, B, and C 2GB; CAT D 4GB<br>Discrete GPU (FB Width):<br>CAT A, B: 35 kWh (≤ 128-bit); 50 kWh (>128-bit); CAT C, D: 50 kWh (>128 bit)<br>Additional storage: 25 kWh |
| Thin Client | | Idle State: Cat A: ≤12.0 W; Cat B: ≤15.0 W<br>Off Mode: ≤2.0 W; ≤2.0 W | Sleep Mode: (if applicable) ≤2.0 W; ≤2.0 W<br>WOL: (if shipped with WOL enabled) 0.7W; 0.7W |
| Workstations | | $P_{TEC} \leq 0.28*[P_{MAX}{}^{1} + (\text{\# HDD's} *5)]W$ | |

[1]Notes: $P_{MAX}$ = Max Power based on Linpack & SPECviewperf BMs

| Notebooks | Cat C | Cat B | Cat A |
|---|---|---|---|
| CPU | ≥ 2 Cores | N/A | All systems not in CAT B or C |
| Gfx | Discrete GPU FBW> 128-bit | Discrete GPU | |
| Memory | ≥ 2 GB memory | N/A | |

| DESKTOP | Cat D | Cat C | Cat B | Cat A |
|---|---|---|---|---|
| CPU | ≥ 4 Cores | >2 Cores | = 2 Cores | All systems not in CAT B. C or D |
| Gfx (One or Both) | Discrete GPU FBW >128-bit | Discrete GPU | N/A | |
| Memory | ≥ 4 GB memory | ≥ 2 GB memory | ≥ 2 GB memory | |

| Thin Client | Cat B | Cat A |
|---|---|---|
| | Support Local multimedia (Encode/decode) | All systems not in CAT B |

TEC Weighting

| | Desktop | | Notebook | | Workstation |
|---|---|---|---|---|---|
| | Conventional | Proxying | Conventional | Proxying | Conventional |
| Toff | 55% | 40% | 60% | 45% | 35% |
| Tsleep | 5% | 30% | 10% | 30% | 10% |
| Tidle | 40% | 30% | 30% | 25% | 55% |

# Energy Star Version 4.0 Computers Criteria

| Platform | Power Supply | Idle (S0, Cx) | Sleep (S3) | Standby (S5) |
|---|---|---|---|---|
| Notebooks & Tablets | No change | Cat A: ≤14W<br>Cat B: ≤20W | ≤1.7W/ ≤2.4W[1] | ≤1.0W/ ≤1.7W[1] |
| Desktop, Integrated (AIO) DT, uP-DDS, Game Consoles | 80% efficient *(20%-100% load)* + >0.9 PFC | Cat A: ≤50W<br>Cat B: ≤65W<br>Cat C: ≤95W | ≤4.0W/ ≤4.7W[1,2] | ≤2.0W/ ≤2.7W[1] |
| uP/DP Workstations | | $P_{TEC} \le 0.28*[P_{MAX}^3 + (\# \text{ HDD's } *5)]W$ | | |

| Notebooks | Cat B | Cat A |
|---|---|---|
| Gfx | GPU w/ >128MB dedicated memory | All other systems |

| Workstations | |
|---|---|
| Mandatory | RDDR/FBD; "Workstation"; MTBF |
| Options 3 of 6 | 2P-skts; +Pwr_GFX; no-UMA; 5+PCI/E, 2+ISV; PCI-X/4xPCIE |

| DESKTOP | Cat C | Cat B | Cat A |
|---|---|---|---|
| CPU | Multi-Core or >1 discrete CPU | | All other systems |
| Gfx | GPU w/ >128MB dedicated memory | N/A | |
| Memory | Min 2GB memory | Min 1GB memory | |
| HDD | Min 2 HDD's | N/A | |
| Other | TV tuner or video capture w/ HiDef | N/A | |

*Choose 2 of 3*

***Notes:***

*[1]0.7W adder for WOL support (Wake on ME); latest spec draft also indicates systems with 'full network connectivity' i.e.. AMT, may be exempted from Sleep/Standby requirements.*

*[2]Desktop Derived Server is exempt from Sleep requirements*

*[3]PMAX = Max Power based on Linpack & SPECviewperf BMs*

# About the Authors

## Technical Contributors

**Rama Kolappan** – Intel Corporation
Rama Kolappan is a Senior Platform Application Engineer working on power delivery enabling for Intel's desktop platforms. He is responsible for risk analysis on Intel's desktop platforms for Energy Star and other regulatory programs. He is also responsible for customer platform enabling for OEMs and ODMs to meet Energy Star specifications. Rama started as a design engineer at Intel and moved to product development engineering where he led the yield improvement working group for 65nm processors. He delivered product health indicators used for planning manufacturing capacity in the factories and improved the validation processes, resolved production related issues and optimized test program development. Rama holds a bachelor's and Master's degree, both in Electrical Engineering.

**Henry Wong** – Intel Corporation
Senior Staff Platform Technologist, enabling Intel's platform power and thermal technologies. Mr. Wong is a 20yr Intel veteran, with 16yrs industry experience in digital and mix signal processor development, including the first Intel mobile chipset, 360SL, and first mobile Pentium (P54LM/P55C). Henry led the development and enablement of high efficiency and high reliability power conversion techniques, component thermal solutions, and system clocking networks for the Intel® Itanium® and Intel® Xeon® - based processor platforms. Mr. Wong authored and enabled key technologies such as Adaptive Voltage Positioning, Modular Direct Power Connect, Server Component High Impingement Mode Cooling, and Programmable Geared Differential Clocking for Multi-Time Domain Architectures. He is currently leading Intel's support of the enterprise industry power and computing efficiency initiatives, with organizations such as the EPA. Mr. Wong is a 1984 graduate from Yale University in both Semiconductor Physics and Econometric Modeling.

**Mohan Duggirala** – Intel Corporation
Mohan Duggirala is a Technical Marketing Engineer in Mobile Platforms group enabling Intel's mobile customers for extended battery life, power optimization & performance optimization on their latest platforms. He also supports electrical analysis and platform validation activities on customer mobile designs. In the 3 years he has been with Intel, Mohan has handled customer platform enabling activities for key APAC OEMs & ODMs, providing end to end hardware support to customers on their designs, enabling new features and technologies, coordinating sample delivery and issue debug and resolution. Mohan holds a Master's degree in Electrical Engineering from Western Michigan University, a Bachelors degree in Electronics and Communication Engineering from University of Madras and has 6 years of Industry experience.

**Katharine Kaplan** – US Environmental Protection Agency
Katharine Kaplan leads the ENERGY STAR Product Development Team, responsible for developing energy efficiency specifications for a full suite of IT and other electronics products, calculating program benefits, managing verification efforts, and creating all web based IT tools for the program. She managed the development of Version 4.0 and Version 5.0 ENERGY STAR specifications for computers and specifications for Imaging Equipment, Televisions, Set-Top Boxes, and Digital to Analog Converter Boxes. Also during her 14 years

at EPA, Katharine led EPA's Plug-In To eCycling program and negotiated settlements with companies that disposed of waste at the largest hazardous waste landfill in California. Before joining EPA Katharine worked for Booz, Allen & Hamilton, an international management consulting firm. She holds a Bachelor of Science in Journalism and a Master of Public Administration.

**Evan Haines** – ICF International, Inc. (technical consultant for EPA on ENERGY STAR)
Evan is an Associate at ICF International and has worked for the past 3 years as a consultant to the EPA's ENERGY STAR Labeled Products programs. He provided development support for the Version 5.0 ENERGY STAR computer specification revision and currently leads the administrative process under the Version 4.0 program requirements. Evan's concentrations are in energy-efficient electronics and lighting, with specification development and implementation projects covering battery chargers, external power supplies, and various residential lighting programs. Evan is a graduate of the University of Rochester with a Bachelor of Science degree in Electrical and Computer Engineering.

**Thomas Bolioli** – Terra Novum, LLC (technical consultant for EPA on ENERGY STAR)
Thomas Bolioli has been working in the IT field since 1993 in such roles as developer (web & application), small business IT consultant as well as a systems administrator for large user populations. Since 2001 Thomas Bolioli has been consulting to the Environmental Protection Agency in support of ENERGY STAR. Projects Thomas Bolioli has participated in were the Million Monitor Drive (an effort to get users to enable monitor power management) and a follow up which focused on having users and companies enable the system standby/sleep function on their computers. Other work for ENERGY STAR performed by Thomas Bolioli has been to support ENERGY STAR on the development of various energy efficiency and labeling specifications, including those for computers, Digital Television Adapters and set top boxes. Thomas Bolioli has also worked on energy efficiency, renewable energy and other miscellaneous IT projects for the state of New York and Massachusetts as well as numerous private clients.