



POWER MANAGEMENT SYSTEM DESIGN GUIDE

TABLE OF CONTENTS

| | |
|--|----|
| Section 1: Introduction and Benefits of S3 | 3 |
| Section 2: Guide Purpose and How to Use Guide | 5 |
| Section 3: S3 User Experience | 6 |
| Section 4: S3 Wake Mode..... | 8 |
| Section 5: System Recommendations for S3 Reliability..... | 10 |
| Section 6: S3 Latency | 11 |
| Section 7: Windows Logo Program Overview..... | 13 |
| Section 8: Network Proxy Ecma TC38-TG4 Standard..... | 15 |
| Section 9: ENERGY STAR v5.0 | 17 |
| Section 10: General Design Principles for S3 Associated Hardware | 18 |
| Section 11: System Manufacturer Recommendations | 19 |
| Section 12: Implementing Power Management for USB Devices | 22 |
| Section 13: Reliability Testing and Tools for S3..... | 24 |
| Section 14: Glossary of Terms | 27 |

SECTION 1: INTRODUCTION AND BENEFITS OF S3

The problem of disruption of global climate by human-produced greenhouse gases in the atmosphere is coming to be understood as the most dangerous and intractable of all the environmental problems caused by human activity. The problem is highly intractable because the dominant cause of the disruption, emission of carbon dioxide from fossil fuel combustion, arises from the process that currently supplies nearly 80 percent of civilization's energy. (John Holdren (Harvard, Kennedy School of Government) UN Investor Summit on Climate Risk, May 10, 2005). The ubiquitous PC and the global information and communications technology industry that support it account for approximately [2 percent of all global CO₂ emissions](#)¹. As such, computer energy consumption is an appropriate target for energy conservation strategies. Power management is one, and arguably the most promising, of the available PC energy conservation strategies. By simply lowering an idle, unused computer's energy consumption, significant energy savings can be realized without sacrificing system productivity or performance.

Today [90 percent](#)² of desktops do not use power management, even though today's computers support this feature in their operating systems, power management software and platform components and design. Adopting power management can generate computing energy savings up to [60 percent](#)³ without impacting productivity or performance. Implementing power management allows a desktop or laptop to quickly transition into and out of a lower energy state, consuming as little as three to five W of power versus more than 50 W when left in the Idle state. Further overall energy savings can be realized by a decrease in computer system cooling and a reduction in total energy to cool the room environment. Other benefits include longer laptop battery life, reduced computer system fan noise and a reduction in total CO₂ emissions. Climate Savers Computing Initiative's goal is to reduce computer power consumption by 50 percent and thereby reduce global CO₂ emissions by 54 million tons per year. CSCI plans to achieve this goal in two ways: first through enterprise and consumer adoption of energy-efficient systems; second through the implementation of computer power management.

The question can be asked, "What is computer power management (PM)?" PM is the feature that lowers the computer system energy state to a lower-power state when inactive or idle. Advanced Configuration and Power Interface (ACPI) is the general power management standard adopted by the computer industry. A well-managed system using PM will place the system into a low-energy state when not in use for an extended period of time. S3 (sleep) is the preferred ACPI state over S4 (hibernate). In the S3 state, a PC sleeps when it is idle, wakes quickly to the task generated locally or remotely, then puts itself back to sleep. S3 is also referred to as "Suspend to RAM" as RAM is still powered. In the S4 state, memory from RAM is saved to non-volatile memory and is powered down.

S3 offers the best trade-off between power savings and the ability to wake quickly and begin doing real work. In S3, the major power usage devices such as CPUs, graphics controllers, hard drives and optical devices are completely turned off. Memory is placed in a low-power, self-refreshed state. Certain other devices such as the keyboard, mouse and network controllers receive only enough power to enable them to respond to a wake-up signal which will in turn cause the PC to fully awaken. Since the tasks that were being worked on before S3 are usually still resident in system memory or RAM, the PC can almost immediately begin doing useful work. Today, vast numbers of computers are left fully on 24 hours a day, seven days a week without adopting PM. Computer users are either unaware of the energy-saving benefits of PM or are fearful that systems will not recover reliably. In fact, S3 is a mature technology and the corner case wake issues are constantly being addressed. Many businesses, government offices and consumers utilize S3 technology successfully today.

What are the energy-saving benefits of using S3 technology? Below is an example highlighting the energy differences between enabling S3 versus simply allowing a PC to idle. Today's relatively efficient dual core CPU desktops are capable of drawing five W or fewer in S3 and around 60 W in idle. Assuming a 15-hour idle period (5 p.m. to 8 a.m.), an average work year of 250 days, an average energy cost of 15 cents per kilowatt-hour and a 1,000 PC installation. A savings of 55 W in S3 over idle would yield more than \$30,000 due to reduction of

206 kilowatt-hours. Further energy savings would be realized from a decrease in environmental cooling. Now consider the vast majority of systems that are always on, sitting idle for most of the day. The math demonstrates significant potential worldwide energy and CO₂ emissions savings of \$55 billion and 54 million tons each year respectively.

References:

- 1) 2 percent of all global CO₂ emissions
(<http://www.gartner.com/it/page.jsp?id=503867>)
- 2) 90 percent
(<http://www.climatesaverscomputing.org/about/faq/#11>)
- 3) 60 percent
(<http://www.climatesaverscomputing.org/about/faq/#11>)

SECTION 2: GUIDE PURPOSE AND HOW TO USE GUIDE

The CSCI Power Management System Design Guide is a reference document on how to build energy-efficient, power-managed client platforms. Specifically, this guide presents guidelines and recommendations to achieve Standby (S3) state reliability, while reducing Standby, Resume and Sleep latencies resulting in positive user experiences. PM implementation depends on successful interactions of the various computer platform elements. Consistent and dependable interactions are necessary between the operating system (OS), microprocessor, chipsets, BIOS, devices, PC cards, monitors and software (drivers and applications) in order for PM to operate as anticipated. This guide will focus on aspects of the computer components, as well as platform design, operating systems and software to achieve consistent, reliable and efficient Standby operations across platforms. It contains material for computer OEMs, component manufacturers, system integrators and manufacturers, software engineers and end users. This guide can be used as a reference to build power-efficient client systems, including the anticipated regulatory ENERGY STAR compliant proxy (Ecma TC38-TG4) devices and is divided into 11 additional sections and a glossary of terms beyond this section, with each addressing a barrier to reliable and/or efficient PM implementation at the system or user level. The sections do not need to be read in numerical order and the sections are not ordered in any priority beyond the first two sections. Each section independently addresses specific issues pertaining to client PM system or component design or implementation principles.

The sections are organized as follows:

- Section 1 – Introduction and Benefits of S3
- Section 2 – Guide Purpose and How to Use Guide
- Section 3 – S3 User Experience
- Section 4 – S3 Wake Mode
- Section 5 – System Recommendations for S3 Reliability
- Section 6 – S3 Latency
- Section 7 – Windows Logo Program Overview
- Section 8 – Network Proxy Ecma TC38-TG4 Standard (formerly TC32-TG21)
- Section 9 – ENERGY STAR v5.0
- Section 10 – General Design Principles for S3 Associated Hardware
- Section 11 – System Manufacturer Recommendations
- Section 12 – Implementing Power Management for USB Devices
- Section 13 – Reliability Testing and Tools for S3
- Section 14 – Glossary of Terms

These design guidelines are recommended by CSCI to improve clients' Sleep reliability and reduce latencies during Sleep and Resume cycles. The guidelines are not CSCI member-company design requirements. For specific member-company design requirements, parties should contact that company directly. This guide does not provide specific PM software product recommendations. Rather, it addresses general design principles for making PM faster and more reliable.

SECTION 3: S3 USER EXPERIENCE

The energy savings and productivity potential of S3 has not been realized in the computing community. Laptop owners often opt to use hibernation (S4), which writes system context to non-volatile storage, usually a hard drive. Hibernation, which predates S3, has long been employed on laptops to conserve battery life. However both single-user and IT-administrated network productivity suffers with Hibernation due to the long latency required to transition back to the operating system.

End-user desktop usage for Standby is largely to conserve AC power. Standby is used by a significant percentage of business desktop customers, which are largely managed by their IT departments. Business desktop customers that have a well-controlled network environment and use OS images as shipped from the top-tier OEM vendors generally have successful S3 experiences. Because of low and stable power grids, business desktop users generally prefer S3 instead of Hibernation.

The advent of hybrid Sleep offers the opportunity for both notebook and desktop users to enjoy the low latency of S3 and the data protection of S4. Upon entry into S3, the Hibernation system context is written to non-volatile storage. If there is no loss in AC power, the system will resume quickly from S3. If AC power is lost, system context can be retrieved from non-volatile storage.

The end-user community would need to be acquainted with an explicit recognition and associated functionality of Standby, in addition to more intuitive interfaces for managing operations within this state. Current user experience for Standby is inconsistent not only across OS and system vendors, but even within different models from the same vendor. Intuitive interfaces are needed for any state transition, including S3, S4 or hybrid S3.

Specifically, the new usages being defined by the network proxy Ecma TC38-TG4, allowing devices to maintain network presence in Standby, assist in defining requirements around user experience in Standby state.

1. The Basics

- S3, in conjunction with hybrid Sleep, should never cause a data loss. S3 should be at least as reliable as the appliance state transitions users are most familiar with: turning on and turning off.
- Systems in S3 must evidence a compelling decrease in power consumption with respect to the computer's active or idle state.
- Systems returning from S3 must evidence a compelling reduction in resumption latency with respect to the computer's resume time from Hibernation or Off.

2. Consistent Power Settings

- Users should experience consistent default power settings on the platform, including the configuration of the power buttons and their defined usage. For instance, users should understand the default function of the Power button as well as what happens when the Power button is half-pressed.
- The methodology of activating or exiting Standby needs to be consistent.

3. Lid Closing/Opening Operations

- Users should enjoy a consistent experience on various lid operations (e.g., close, open) and timing of these operations.

4. Display

- The platform should offer a consistent Display experience to the users, especially in S3 transitions. This applies to externally attached projectors and monitors.

5. Manage S3 operations

- Owing to limited current use of Standby, users need to be presented consistent and an easy-to-use interface for managing functionality within Standby. These requirements apply to use of the Ecma network proxy Standby usages:
 - Users should be presented with an intuitive user interface to enable and disable such functionality, on demand.
 - Design of clear user-settable timers for when the platform will enter Standby, and triggers that will transition platform from Standby to Idle or Active.
 - Present users with only those configuration options that are supported by the platform. For instance, there have been user reports on enabling "Wake On LAN" in the OS, only to discover that the LAN/WLAN components either were not powered on in the Standby rail, or were disabled in the BIOS.

6. Generate meaningful events and remedial information

- Users should be able to get events and remedial suggested actions regarding Standby (e.g., when any Standby-related operations are modified).

7. Managing software behavior

- Users may install software that may inadvertently prevent a platform from entering Standby or wake up the platform for events that the user did not "register."
 - The users should be able to get a list of non-compliant applications and related activity and have interface options to manage those.
 - At the time of SW application installation, the user should be presented with Standby-Wake compliance or non-compliance of those applications. Currently, users can install software that inadvertently breaks their configured Standby-Wake functionality. Such applications should be flagged at install and run times, with meaningful events and recommended remediation.

SECTION 4: S3 WAKE MODE

For ENERGY STAR 5.0 Client computer compliance, OEMs must ship computer systems with power management or S3 enabled. Sleep mode lowers a system's energy consumption while maintaining rapid response to full functionality, in line with ENERGY STAR's goals of promoting energy-efficient products. In addition, ENERGY STAR-compliant platforms must also be shipped with Wake on LAN (WoL) enabled from Sleep when operating on AC power. ENERGY STAR client systems with Ethernet capability only need the ability to enable and disable WoL for Sleep. WoL is an Ethernet computer networking standard that allows a computer to be turned on or woken up remotely by a network message sent by another computer on the network. Those systems with full network connectivity via a proxy are granted greater operational mode weighting in Sleep for the ENERGY STAR TEC (typical energy consumption) calculation. Proxy agents, which are in the process of being defined by Ecma, handle certain kinds of network activity on behalf of their host systems. This allows computers to spend more time asleep.

Keeping a computer in Standby with WoL can reduce power consumption when compared to keeping the same computer in Active (S0) state. However, enabling WoL can increase power consumption while in S3 compared to systems without WoL enabled in S3. WoL requires extending Standby power rails to additional power wells, e.g. the wake-up circuitry for LAN or WLAN controllers, which boosts a system's power consumption in S3. In the past, the EPA has allowed up to 700 milliwatts for WoL capability. Note that network proxy, depending on its capabilities, may require more Standby power than WoL.

To address this seeming contradictory ENERGY STAR requirement of enabling a more power-consuming feature, CSCI makes the following recommendations with respect to WoL and network proxies. For ENERGY STAR compliance, platforms must be shipped with WoL enabled by default. However, users should determine best implementation models once systems arrive. For example, Standby and Standby with network proxy enablement should be dependent on the specific desired usages such as 24x7 remote access, remote manageability, etc. CSCI supports network proxy over WoL for ease of use and simplicity of implementation. See Section 8: Network Proxy Ecma TC38-TG4 Standard for more usages. In such usages, Standby with network proxy will provide substantial energy savings as compared to keeping the PC in continual active state 24x7. If no such usages are desired, then the PC should be placed in Sleep (S3) without WoL enabled or Off (S5), thus lowering the user's energy consumption relative to S0. CSCI further recommends that network proxy and WoL be disabled in DC mode (i.e., on battery) by turning off the power supply to network devices. Disabling network proxy and WoL should be done by the system firmware or hardware platform.

To facilitate enabling and disabling network proxy and WoL, vendors should provide an easy network proxy configuration in the network setup application. A software mechanism to enable and disable the network proxy and WoL should be provided. CSCI recommends that network proxy and WoL settings in network profile be configured by the native operating system or an OEM's network setup application. These configurations would allow most users to enable and disable the proxy or WoL depending on network profile. Computer users would then set up their computers to the appropriate network proxy (see Section 8: Network Proxy Ecma TC38-TG4 Standard) or WoL configuration for their usage models.

In addition to supporting WoL, an energy-efficient computer with full network connectivity needs to respond intelligently to commands while sleeping. Platforms with WoL should wake the system only when user scenarios require the system to wake up. To support selective wake-up, platforms should be configured to expose Wake programmability through WMI in the device driver. System manufacturers should validate that every Wake device can properly wake the system and that the Wake source is properly identified.

Platforms and devices should support reporting the source of the Wake event to the operating system. This reporting allows the operating system to take the proper response action. For example, if the Wake source is due to the RTC, the operating system can assume that no user is present and therefore not turn on the display. In

addition, platforms and devices should enable the OS to log diagnostic information for troubleshooting purposes. For example, an administrator can determine the reason why a particular PC is waking spuriously.

Beyond the platform, device-specific recommendations should be implemented to aid with Wake source identification and subsequent intelligent Wake response. If the Wake source cannot be properly identified the device requesting attention will not be properly serviced. Below, CSCI makes recommendations for each of the common devices: fixed-feature devices, ACPI control-method devices, PCI express devices and legacy PCI devices.

- For fixed feature devices, integrators should follow the operating system implementation recommendations for device-drivers, and wake devices should be wired with the corresponding ACPI PM1 register bits, e.g., PWRBTN_STS, SLPBTN_STS, RTC_STS, PCIEXP_WAKE_STS.
- For ACPI control-method devices, Wake devices should be wired with corresponding GPE register bits. In addition, these ACPI devices should align the GPE usage with hardware setups used in the ACPI firmware and execute a Notify (0x02) in the GPE handler. Lastly these drivers should follow operating system implementation recommendations for device drivers.
- For systems with PCI express devices, Wake devices should be wired with PM1 PCIEXP_WAKE_STS register bit and should not assert legacy GPE events. These devices should grant Windows native control via _OSC. These drivers should also follow operating system implementation recommendations for device drivers.
- For systems with legacy PCI devices, the Wake devices should be wired with corresponding GPE register bits. The drivers should not share any signals with PME# or from different root buses or PCI hierarchies. To find the wake source for PME, for instance, an operating system may have to walk the entire PCI bus, which is time consuming. Adding other Wake signals to PME may create so much overhead that essentially the OS is presented with an “interrupt storm.” Some Wake devices may never be serviced. System designs should align GPE usage with a hardware setup in ACPI and should execute a Notify (0x02) in the GPE handler. Lastly these drivers should follow operating system implementation recommendations for device drivers.

Consult these Windows OS references for wake-up methodologies that OEMs and integrators should implement:

- Kernel-Mode Driver Architecture: Implementing Power Management
(<http://msdn.microsoft.com/en-us/library/ms798306.aspx>)
- Supporting Devices That Have Wake-Up Capabilities
(<http://msdn.microsoft.com/en-us/library/aa490140.aspx>)

SECTION 5: SYSTEM RECOMMENDATIONS FOR S3 RELIABILITY

Power management has been available on computer systems for many generations and upgrades. However, most users are either not aware of the feature or purposefully choose not to enable it due to historical reliability and latency concerns. Previous software and hardware models and configuration have not always allowed for a reliable return to normal user conditions. Many users returning from S3 experienced computers that would simply hang or terminate in a blue screen. This activity could result in data loss as systems would require a total reboot. Improvements in PM from both hardware and software vendors have greatly enhanced the reliability and expediency of systems returning from low-power Sleep to normal user conditions. Below is a summary of recommendations OEMs can implement during development and assembly phases to ensure platforms support reliable and swift transitions into and out of PM.

OEMs are recommended to conduct stress tests during system development to correct S3 instability at the system level. IHVs and ISVs are also recommended to conduct stress tests using a reference platform for reliability confirmation – 300 cycles of S3 entry and exit, for example. End users are recommended to update device drivers, OEM applications or OS patches because OEMs, ISVs or IHVs might have updated their software after conducting stress tests. Updating software will result in reliable and quick Suspend and Resume for end users. In addition, client users are recommended to periodically update device drivers and applications when available from OEM Web sites.

Certification processes that confirm a system's hardware and firmware configuration with respect to S3 reliability are strongly recommended. For example, as part of its Windows logo-certification program, Microsoft requires Windows logo tests. These tests include software tools that run repeated S3 loops on the hardware and firmware configurations. Products that pass these operating system tests are permitted to use a "Certified for Windows" logotype, which certifies that the hardware or software has had some share of testing by Microsoft to ensure S3 compatibility, among other requirements. Although Linux and other operating systems may not have official certification programs, OEMs should confirm a system's S3 reliability using S3 testing loops.

SECTION 6: S3 LATENCY

A system in a Suspended or Sleep mode (S3) consumes less power than when in the Idle state (S0). Waking a system from a sleeping state is not instantaneous, but requires a brief waiting period before a system returns to its Active state. Some device drivers or applications cause prolonged latency, resulting in extended or frustrating user experiences. OEMs, ISVs and IHVs should design systems and components to minimize the S3 recovery period.

1. OS Analysis Tool Utilization

Microsoft provides tools to minimize S3 latency for those systems that have installed Windows OSs such as Vista and Windows 7. One such tool is Microsoft xPerf. OEMs and ISVs should utilize this and similar analysis tools to minimize S3 latency during platform and software development. In future revisions, Microsoft may want to consider simplifying the analysis tool to allow IT administrators as well as non-technical users to determine and fix sources of Suspend and Resume latencies.

- The Microsoft suggested analysis tool (<http://msdn.microsoft.com/en-us/performance/cc825801.aspx>)
- Resume performance (<http://www.microsoft.com/whdc/system/sysperf/resumeperf.mspix>)

2. Windows Logo Certification and Latest Drivers

Windows logo requirements include a reduced S3 Sleep recovery time. Users should select PCs that are Windows logo certified. These systems must meet certain requirements including S3 latency checks. In addition, Microsoft's Velocity program is targeted to ensure a satisfactory Windows OS experience, which includes short S3 latencies. To ensure reduced S3 latency, OEMs, ISVs and IHVs should consider complying with Microsoft's Velocity program. Employing Windows logo- and Velocity-certified systems will greatly improve reliable and quick Suspend and Resume. Users are encouraged to regularly verify and update device drivers and applications from OEM Web sites.

3. Unnecessary Applications Closure Before S3

Before suspending or hibernating the system, users should close any documents that may be open and sharing over the network as well as any unnecessary applications.

4. Resume Time Considerations on Wake via Network Proxy

Certain low-power network proxy usages such as those being defined by the network proxy Ecma TC38-TG4 rely on a relatively rapid resume time of the host platform. On a typical network Wake scenario performed using a new TCP connection request (TCP SYN), the Ecma proxy will trigger a host Wake. In these scenarios, it is critical that once the host resumes, the network device (NIC) host network stack and all necessary services – like Host Intrusion Prevention systems, firewalls, etc. – are operational prior to the expiration of the TCP connection request retry connection timeouts. Failure to return the host to a fully functional network will result in a futile Wake, and may potentially diminish user experience with network Wake.

5. Checklist of S3 Latency Issues by OEMs, IBVs, IHVs and CPU/Chipset Vendors

- Design CPUs to restore context without the need for multiple CPU resets.
- Design CPUs to restore context without the need for multiple CPU code patch loads.
- Multi-core CPU designs should allow for code patch loads to be done in parallel.
- Optimize code execution efficiency in the BIOS S3 Resume path. As the system BIOS transitions from assembly code to C, even greater care must be taken in avoiding unnecessary code paths. S3 Resume times have increased by five to six times in the transition from assembly to Tiano-based BIOS paradigms.
- Non-host processors (NHP) that aid the primary CPU(s) to facilitate such things as sideband manageability are becoming more popular. NHPs must not greatly impact S3 Resume time. NHPs should do the minimum

amount of initialization and then let the host processor gain control. If there is handshaking between the host CPU and the NHP, the host CPU should never have to wait on the NHP.

- The boot CPU must be allowed to turn on its cache as quickly as possible in the S3 Resume path. X86 processors are extremely slow, especially on code fetches from non-volatile storage, when their caches are disabled.
- Card and chipset vendors should ensure that PCI Express devices train quickly and reliably (<http://www.pcisig.com/specifications/pciexpress/specifications/>).
- Card and chipset vendors should ensure that USB devices get ready quickly and reliably.

SECTION 7: WINDOWS LOGO PROGRAM OVERVIEW

The [Windows Logo Program](#)¹ is designed to address the current and future market needs of customers using the Windows platform. The Windows logo signifies the compatibility and reliability of systems and devices with Windows operating system. It gives customers confidence that a product has been thoroughly tested with Microsoft-provided tools and ensures a good user experience.

The Windows Logo Program helps partners to innovate and bring a premium experience to market, thereby improving their ability to increase market share. The program strives to continuously improve its processes, responsiveness and partner satisfaction.

Getting Started

To help ensure that compatibility with Microsoft Windows is built in during hardware and driver design, follow these [steps](#)²:

1. Design for compatibility with industry standards.

From a user's perspective, new hardware should install correctly on the first try, without interfering with other devices and applications on the PC. To enable this basic experience, Windows must be able to configure and manage devices. Industry standards are the foundation for hardware compatibility with Windows management of system and device configuration. For information about standards for Windows compatibility, see:

- [Design PC Fundamentals: Overview](#)³
 - Industry standards such as USB and PCI that define the implementation mechanisms for [Plug and Play](#)⁴ and [power management](#)⁵.
- [Device Fundamentals: Overview](#)⁶
 - Windows compatibility guidelines, device class specifications and protocol standards, such as networking and communication protocols.

2. Design for the Windows Logo Program.

Designing to meet the criteria for the Windows Logo Program ensures that your product is compatible with industry standards, is easy to install and is reliable and secure. For details, see:

- [Microsoft Windows Logo Program System and Device Requirements](#)⁷
 - Criteria for servers, desktop PCs, mobile PCs and hardware components to qualify for the Windows Logo Program.

3. Design for reliability, stability and security.

For details about designing for system and device reliability, see:

- [Security and Reliability Strategies](#)⁸
 - White papers, checklists and access to resources for creating reliable, secure drivers.
- [Development Tools and Testing](#)⁹
 - Information about development and testing tools for verifying reliability and stability.

4. Design and test for system-wide performance.

Each component in a PC or server system must be tested in common configurations, with consideration for issues such as startup and resume behavior, performance, and battery life. For details, see:

- [System Performance](#)¹⁰
 - Design guidelines, system information and tools for performance issues.
- [Driver Performance: Best Practices](#)¹¹
 - Design and testing guidelines for optimizing driver performance.

5. Test with current versions of Windows, including beta testing.

Rigorously test all your devices and systems under the current versions and beta release of Windows, including service packs. For information about Windows beta, see Beta Testing Frequently Asked Questions:

<http://www.microsoft.com/whdc/resources/respec/default.aspx>

6. Attend WinHEC, the Windows Hardware Engineering Conference.

WinHEC provides the year's best opportunity for technical professionals across all segments of the PC industry to meet and investigate the future directions for quality and innovation with the Windows platform. For details, see the WinHEC Web site: <http://www.microsoft.com/whdc/winhec/default.aspx>

Windows Logo Requirements

- System and Hardware Requirements
(<http://www.microsoft.com/whdc/winlogo/hwrequirements.aspx>)
- Specifications: Hardware and Firmware Standards
(<http://www.microsoft.com/whdc/resources/respec/specs/default.aspx#a>)

References:

- 1) Windows Logo Program
(<http://www.microsoft.com/whdc/winlogo/default.aspx>)
- 2) steps
(<http://www.microsoft.com/whdc/winlogo/getstart/design.aspx>)
- 3) Design PC Fundamentals: Overview
(<http://www.microsoft.com/whdc/system/default.aspx>)
- 4) Plug and Play
(<http://www.microsoft.com/whdc/system/pnppwr/pnp/default.aspx>)
- 5) power management
(<http://www.microsoft.com/whdc/resources/respec/specs/pmref/download.aspx>)
- 6) Device Fundamentals: Overview
(<http://www.microsoft.com/whdc/device/default.aspx>)
- 7) Microsoft Windows Logo Program System and Device Requirements
(<http://www.microsoft.com/whdc/winlogo/hwrequirements.aspx>)
- 8) Security and Reliability Strategies
(<http://www.microsoft.com/whdc/driver/security/default.aspx>)
- 9) Development Tools and Testing
(<http://www.microsoft.com/whdc/driver/security/default.aspx>)
- 10) System Performance
(<http://www.microsoft.com/whdc/system/sysperf/default.aspx>)
- 11) Driver Performance: Best Practices
(<http://www.microsoft.com/whdc/driver/perform/default.aspx>)

SECTION 8: NETWORK PROXY ECMA TC38-TG4 STANDARD

Conventional computer devices maintain networking only in higher power Active and Idle modes. As such, when these devices go into a much lower power Sleep mode, the network connectivity is lost. An inadvertent effect of Sleep is that it denies remote applications access to these devices until such time as the device is powered back on. 24x7 access to a device would require it to be fully powered in the "ON" mode, at all times, consuming energy even when not doing any useful work. Applications like hosting remote media sharing would require this device to be fully powered even when the remote access is sporadic. Users, therefore, have a choice of either saving energy or getting 24x7 application access. The objective of network proxy standardization is to provide application access and energy savings at the same time, giving the end user the best of both worlds while enhancing experience.

The ENERGY STAR v5.0 Computer Specification introduced the concept of "Full Network Connectivity" as the ability of a sleeping device to maintain network presence using a low-power "Proxy" that would wake up the higher power computer device when further processing is required. The specification provided for increased allowances for any device implementing such a system and therefore incentivized its creation and deployment. This proxy is a lower power component that allows a higher power device to remain in Sleep, thereby reducing overall energy consumption for that device considerably. The typical desktop PC uses anywhere from 30 to 60 W of power at Idle and typically less than two W in sleep. Even if the proxy doubles the sleep power to four W, the energy savings potential is huge if the Proxy allows end users to enable sleep more often.

With the Proxy operational, any incoming valid connection will fire off a PME or Wake event to bring the host back into Active. ENERGY STAR v5.0 has established different weightings for proxy-enabled and conventional (with no proxy) devices. ENERGY STAR, however, does not define what standard to utilize as none was available at the time the specification was written. A cooperative process was created in the European standards body Ecma International (TC38-TG4) between industry and energy-efficiency advocates to set forth that standard. ENERGY STAR mandates that the proxy standard, developed by the industry group, be approved by the US Environmental Protection Agency (EPA) and the European Union (EU), as meeting the goals of ENERGY STAR. In addition, ENERGY STAR has formulated a measure for "Typical Energy Consumption" (TEC) for networked devices which measures the overall annual energy consumption using power consumed in Sleep, Idle and Off with corresponding times in each of these states. This translates to the following equation, as defined in ENERGY STAR v5.0:

Where P_x is power value in W and T_x is time value in percentage of the year based on the product category and weighting. Those systems implementing full network connectivity during sleep receive greater sleep ratings relative to those without it.

Ecma Technical Committee 32 Task Group 21 (TC38-TG4) has been chartered to establish the industry standard for a low-power network proxy. TC38-TG4 has broad industry representation and participation, ranging from computer, processor, chipset and computer-based device manufacturers, software operating system vendors, networked devices manufacturers, and independent energy efficiency consultants. TC38-TG4's scope is to define network proxying of ICT devices to reduce energy consumption. Based on the adopted timeline and process, TC38-TG4 standard will gain Ecma General Assembly approval in January 2010. A successful standard rests on consensus building, and as such, this timeline may be changed by TC38-TG4 membership.

TC38-TG4 defines network proxy as "an entity that maintains network presence for a sleeping higher-power ICT device." The charter's scope of work includes the specification of "the protocols that network proxies must handle to maintain connectivity while hosts are asleep...the proxy behavior including ignoring packets, generating packets and waking up host systems...[and] the information exchanged between hosts and proxies." It does not include the information syntax and exchange methods. The scope also includes maintenance of published work as well as liaison and cooperation with other standards organizations.



For further details, please check:

<http://www.ecma-international.org/memento/TC38-TG4.htm>

The Ecma TC38-TG4 defined Proxy system is one of the key usage models for this guide. It is expected that the guidelines published herein for improving S3 reliability and reducing S3 Resume and Sleep latencies will be a crucial building block for TC38-TG4-compliant systems. These guidelines could be used in their entirety or as parts by system manufacturers to build reliable, energy efficient and user responsive systems based on network proxying. If, as expected, TC38-TG4 gains approval from the US EPA and the EU, then those systems could be qualified under ENERGY STAR.

SECTION 9: ENERGY STAR V5.0

The US EPA's ENERGY STAR program provides a large incentive for OEMs and VARs to build, label and ship ENERGY STAR-labeled desktops, notebooks and workstations. The federal government has been mandated by Congress to purchase only ENERGY STAR products when available. As there is a labeling program for computers, the federal government will only buy labeled PCs. Many state, local and university buyers use the same purchasing guidelines as well. As much as 25 percent of all buyers – broken out as 15 percent federal, 10 percent state and local – are mandated to purchase ENERGY STAR. As a result, it is advantageous to build ENERGY STAR-compliant systems.

The ENERGY STAR program places considerable weight on the amount of energy a system uses when Idle, with secondary focus on the amount of energy a system uses when in Sleep and Off. To pass ENERGY STAR, it is best to focus on the amount of power consumed in Idle and optimize for that.

In addition to optimizing for Idle power, an implementer may also want to implement a Sleep proxy as defined by the Ecma standard for Sleep proxies. By doing so, the OEM receives a more favorable duty cycle to use for calculating the TEC. This duty cycle can provide up to 25 percent more energy for their platform to be used in any state by their platform.

SECTION 10: GENERAL DESIGN PRINCIPLES FOR S3 ASSOCIATED HARDWARE

Overview

Computer architecture is increasing in complexity with each succeeding generation. Processor packages are now split into multiple cores to allow parallel execution of tasks to increase throughput. Auxiliary processors to offload such tasks as manageability from the system CPU are becoming more prevalent. Circuit logic is being moved from gates to firmware to allow for more design flexibility and the ability to update what was once circuit board logic in the field. All of these changes, while good for the PC ecosystem, present challenges to a favorable S3 user experience. Consideration for S3 Resume time must begin at the start of system design, not as an afterthought.

Design Principles

- Processors: The complexity of x86 technology processors is increasing from product cycle to product cycle. In a few years, processor packages could contain 32 cores or more. Since the processor packages lose power in S3, these cores must be reinitialized. This can and does strain S3 Resume times, which continue to be pressured downward by OS vendors such as Microsoft. A few basic guidelines could help minimize the impact to S3 Resume times.
 - CPU architectures should allow for cache enablement as quickly as possible in the S3 execution stream.
 - Processor core patch loads must be allowed to occur in parallel.
 - Multiple CPU resets should be avoided, or reset mechanisms must be designed that do not destroy the entire CPU state.
- Non-host processors (NHP): Due to the increasing complexity of chip set features, more and more logic is being moved away from gates and state design to auxiliary CPUs that execute firmware. This aids in design flexibility and upgradeability, but S3 Resume time can suffer.
 - NHPs should do the minimum amount of initialization and turn over control to the host processor.
 - Designs should allow the host processor and NHP execution flows to be as independent of each other as possible.
 - If there must be coordination of a host processor and an NHP, ideally the host processor – which is responsible for S3 Resume execution – should never have to wait on the NHP.
- Bus initialization: To achieve the maximum bandwidth, busses are often fine tuned for peak performance. This fine tuning is time consuming and slows down S3 resume.
 - Bus tuning should yield a few parameters that can easily be saved in non-volatile storage during system POST. In this manner, in S3 Resume, only the parameters need be reapplied. This avoids lengthy tuning algorithms in the S3 execution flow.
- Memory initialization: As system RAM speeds are driven faster and faster, the tolerances for reliable performance are getting tighter. Memory is now essentially fine tuned before it can be used. Fine tuning algorithms such as memory training should be done in hardware if at all possible.
- Some of the issues related to Suspend and Resume and Hibernate and Resume are tied to dual and multi-core CPUs, issues typically addressed via QFE, driver or BIOS updates. Check if the issue is related to dual and multi-core CPU by selecting single-core option typically exposed in BIOS.
- The hybrid Sleep feature in Windows Vista and Windows 7 saves Hibernate file and then enters Standby. Many devices, and in particular Wake-capable devices, need special handling in their devices and drivers for handling hybrid Sleep transitions and need to be designed taking into consideration the applicable requirements.

SECTION 11: SYSTEM MANUFACTURER RECOMMENDATIONS

System manufacturers should validate that the firmware on their systems have the following abilities:

1. Support processor performance and idle states.

- Take advantage of chipset and processor components that can dynamically scale their power consumption with workload demand.
 - Select a processor that supports multiple ACPI P-states.
 - Select a processor that supports ACPI C-states.
- Enable Windows control of processor performance and Idle states. Validate use of processor power-management technologies by using PwrTest.
- ACPI firmware must accurately describe the processor performance and idle state capabilities. Windows performs stringent validation of processor PM firmware descriptions and will disable processor performance and Idle states if errors are found in the firmware data.
- OS power management policy is recommended to be "balanced" to allow CPU to be "adaptive," utilizing multiple P-states.

2. Support processor performance and Idle states.

- Take advantage of ASPM (Active State Power Management) for PCI Express devices.
 - Select a chipset that supports ASPM L0s and L1.
 - Select a processor that supports ACPI C-states.
 - Select a PCI Express device that supports ASPM L0s and L1.
 - Enable Windows or BIOS control of PCI Express power management.
 - Device example: Discrete Graphics, GbE, Wireless LAN, etc.
- Take advantage of CLKREQ for PCI Express devices and CLKRUN for PCI or LPC.
 - Select a chipset or a clock generator that supports CLKREQ and CLKRUN.
 - Select a PCI Express device or a PCI/LPC device that supports CLKREQ/CLKRUN.
- Take advantage of Selective Suspend for USB devices.
 - Select a chipset that supports Selective Suspend.
 - Select a USB device that supports USB Selective Suspend.
 - Enable Windows or BIOS control of USB Selective Suspend power management.
 - Device example: Camera, FingerPrint, Bluetooth, Wireless WAN, etc.
- Take advantage of DIPM/HIPM for SATA devices.
 - Take advantage of Device Initiated Power Management and Host Initiated Power Management for SATA devices.
 - Select a chipset that supports HIPM.
 - Select a SATA device that supports DIPM.
 - Enable Windows or BIOS control of HIPM and DIPM.
 - Device example: SATA HDD, SATA ODD.

3. Provide Aux power to the 802.11 Wireless and Ethernet Network Interface Cards (NICs) when the platform transitions into S3 in AC mode.

4. Validate that other core chipset PM technologies are enabled in the platform.

- If a platform has a chipset low-power processor, then implement capability to power that processor when the platform transitions into S3 in AC mode.

5. Take advantage of the latest graphics device display PM technologies.

- Integrate ALS devices and provide firmware and software support for integrating ALS control with Windows.

- Enable Windows to control display brightness through ACPI methods. This allows the display brightness to automatically decrease when the system goes on battery power.
- Incorporate display power conservation technologies that do not noticeably impact display quality.
- Use LED-backlit displays to reduce overall system power consumption when the display is on, such as in mobile PCs.

6. Enable very fast Resume from Sleep (S3).

- Design firmware for Resume from Sleep performance. BIOS initialization times should be less than 500 ms as observed by the PwrTest utility.

7. Wake the system automatically when a mobile PC's battery capacity is low.

- The platform should automatically wake when the remaining capacity is just enough to put the system into Hibernate (S4) state. Windows will automatically detect the low battery capacity and hibernate the system.
- Implement hardware and firmware support for Wake on critical battery capacity.

8. System firmware

- Design system firmware to prevent the system from waking on timer alarm when the system is on battery power, except for low battery conditions or if configured accordingly. This helps prevent a mobile PC system from waking up accidentally in an enclosed environment such as a briefcase or a laptop bag.

9. Implement Windows 7 improvements to Wake on LAN (WoL).

- Update hardware and drivers to NDIS 6.20 to take advantage of new Windows 7 functionality.
- Provide support for at least 8 Wake patterns.
- Implement ARP and NS power management offloads on NIC.
- Support new packet type patterns.
- Support low power on media disconnect.
- Implement Windows 7 improvements to WoL.
- Support the 802.11 Wireless offload capabilities.

10. Adjust display brightness.

- Set display brightness to required max in user environment.
- Display brightness is a major power-consuming device. Sixty to 100 nit is recommended for normal usage, balancing power and brightness.

11. Implement network proxy (Ecma TG32-TG21).

- Update hardware and drivers to take advantage of extending Sleep time.
- This is applicable to the Ethernet (802.3) and wireless LAN (802.11) link protocols.

12. Implement the wireless Power Save feature.

- It is recommended that the 802.11 Wireless adapter implements the (Legacy) 802.11 Power Save mechanisms. This Power Save mode is defined in IEEE 802.11-2007 standard, and included in the Wi-Fi Alliance basic interoperability device certification and testing.
- Wireless LAN may implement the Quality-of-Service U-APSD based power save mechanisms, as defined in the standard 802.11-2007.
- Wireless LAN may implement the WMM Power Save mechanism defined by Wi-Fi Alliance WMM Specification.

References:

- Mobile Battery Life Solutions Guide
(http://www.microsoft.com/whdc/system/pnppwr/mobile_bat.msp)

- Mobile Battery Life Solutions Guide (Windows 7)
(http://www.microsoft.com/whdc/system/pnppwr/mobile_bat_Win7.mspix)
- Brightness Control in WDDM
(<http://www.microsoft.com/whdc/device/display/aero/brightness.mspix>)
- ENERGY STAR Computer Specification version 5.0
(http://www.energystar.gov/ia/partners/prod_development/revisions/downloads/computer/Version5.0_Computer_Spec.pdf)
- ACPI 3.0b specification
(<http://acpi.info/DOWNLOADS/ACPIspec30b.pdf>)
- Brightness Control in WDDM
(<http://www.microsoft.com/whdc/device/display/aero/brightness.mspix>)
- NDIS 6.20
(<http://msdn.microsoft.com/en-us/library/dd568055.aspx>)
- IEEE 802.11 Standard
(<http://standards.ieee.org/getieee802/download/802.11-2007.pdf>)
- Designing Power Friendly Devices
(http://download.intel.com/technology/EEP/designing_power_friendly_devices.pdf)

SECTION 12: IMPLEMENTING POWER MANAGEMENT FOR USB DEVICES

Universal Serial Bus (USB) specification was introduced to the computer industry in 1996 as a portable choice for connecting multiple peripheral devices to a computer. Since its inception, USB has been a tremendous success and today is a ubiquitous interconnect on almost all electronic devices and computers. USB 2.0 specification, released in 2002, bumps up the transfer data rates to 1.5 Mbps ("low"), 12 Mbps ("full") and 480 Mbps ("high"). USB 3.0 specification was released in mid-2009, which adds significant power and performance enhancements.

However, USB is a "poll" based protocol, which means that the host continually polls USB devices for data, even when there is no data from the device. USB devices only deliver data when polled by the host; therefore, the host has to continually poll all USB devices. This simplistic architecture has serious power issues on the host. Specifically, USB's poll architecture results in significant power expense at the CPU, memory and chipsets, as every poll results in a series of memory fetches and cache snoops, preventing host complex from entering low power state.

USB power efficiency can be addressed in close coordination with various components, as discussed in [1, 2]:

1. OS and Driver Model

A power-friendly USB stack is critical to the overall USB energy efficiency. The USB drivers control the USB Host Controller behavior. Hence, any enhancements can result in significant power advantages for USB.

2. USB Host Controller

The host controller can cache the USB Schedule Information to prevent multiple schedule reads from memory. This can prevent the chipsets from entering deeper Power Save states.

3. USB New Fast Suspend Power State

USB has a "Suspend" power state that is limited by large entry and exit latencies, which are magnified when OS behavior is factored in. This largely indicates an introduction of a more efficient low-power state, which can assist in making USB more power efficient. This new L1 state [1] greatly reduces the entry and exit latencies (10s of usecs) while maintaining the functional capabilities of the USB Standby state. The USB controller and device can negotiate this state that allows Idle USB devices to quickly enter or exit standby low-power state.

4. USB Selective Suspend

USB has the selective suspend feature to reduce system power conservation. The selective suspend allows the USB driver to suspend an individual port selectively while the device is idle. When any USB device is active, CPU is prevented from entering low power state like C3 or less due to a "poll" based protocol, so that the selective suspend can reduce CPU power as well as USB device power itself. This technology requires USB driver support.

5. USB Devices

Recommendation:file:///Volumes/Profile/dscheste/Desktop/Untitled%20Attachmentns for USB devices [1] includes:

- USB devices could use Interrupt endpoint to indicate if a device requires service and use the Bulk endpoint for transferring data to or from the device. This translates to more power efficiencies on the entire platform.
- Using Isochronous transfers for streaming devices: USB is a shared bus and therefore, multiple streams share the same bandwidth. This introduces inefficiencies when bulk endpoints are used for data transfers. Using Isochronous, time-scheduled data transfers can be more efficient when properly allocated by the host drivers.
- Minimize the polling rate of the device as much as possible while maximizing device-buffering capacity.

- Devices having "Instant ON" integrated buttons require a continually running Interrupt endpoint to poll the button. This is power inefficient, so it recommended that USB devices work with platform designers to explore alternate options like using platform out of band mechanisms or ACPI modifications.

6. Application Software

Application software should be designed to maximize USB power efficiencies. For instance, USB applications can ensure that they clean up the drivers on exit or inactivity period. This can avoid dangling transactions pending on the device, which if not cleaned up, can result in continual retries.

References:

- High-Speed USB Platform Design Guidelines
(http://www.usb.org/developers/docs/hs_usb_pdg_r1_0.pdf)
- Making USB a More Energy Efficient Interconnect
(http://download.intel.com/technology/itj/2008/v12i1/2-usb/2-Making_USB_a_More_Energy_Efficient_Interconnect.pdf)

SECTION 13: RELIABILITY TESTING AND TOOLS FOR S3

Introduction

It is important that Sleep and Resume transitions be reliable and deterministic. In addition to potential data loss, failure to provide a reliable experience will discourage the use of low-power states. This section details a number of resources and tools to test and troubleshoot reliability issues.

1. Determining Platform Power Leakage

In general, good design practice and enabling power management defaults on chipsets can go a long ways toward reducing platform energy consumption. Platform leakage should be identified and removed. To identify leakage sources, the platform should be placed into S3, S4 and/or S5, and the voltage measured on the power rails to ensure appropriate rails are near zero. If a design mistake was made, then power leaks through resistors and transistors and bad design choices for pull-ups will register at higher than expected power. Sources of leaks can be negated by using intelligent power-gating or ensuring appropriate devices or ports are turned off in S3, S4 or S5. Identifying and removing these voltage leakages will ensure good power savings during these low-power states. In addition, validate that core chipset power-management states are enabled and set as default on chipsets. For example, on USB, optimize power management to ensure the CPU C3 states can be entered. Additional details on USB power management are discussed in USB section of this guide.

2. Power State Transitions

Reliable, deterministic and consistent Power State transitions are essential to the overall user experience. A Windows computer will transition between Power States in the following scenarios:

- Boot – initial startup or boot of Windows
- Shutdown (S5) – shutdown and transition to the off state
- Sleep (S3) – suspend to memory
- Hybrid Sleep – suspend to memory and disk
- Hibernate (S4) – suspend to disk
- Resume – resume operation from Sleep or Hibernate

It is essential that each of these transitions occur in a reliable and consistent fashion. Therefore each scenario needs to be thoroughly tested. Microsoft provides a number of tools to measure, analyze and troubleshoot power state transitions. There are three main tools recommended by Microsoft for testing power transitions. They are as follows:

- B powercfg.exe: Power configuration tool, provided with Windows Vista and later.
- pwrtest.exe: Power management test tool, ships as part of the Windows Driver Kit (WDK). For more information on how to obtain the WDK, please refer here:
<http://www.microsoft.com/whdc/DevTools/WDK/WDKpkg.msp>
- Windows Performance Tools Kit: Windows performance analysis tools, includes xPerf and related utilities. For more information on how the Windows Performance Tools kit refer here:
<http://www.microsoft.com/whdc/system/sysperf/perftools.msp>

3. Sleep, Hibernate and Resume

Microsoft recommends using PwrTest to test Sleep, Hibernate and Resume. The whitepaper, [On/Off Transition Performance Analysis of Windows Vista](#)¹, outlines recommended test procedures and practices for measuring and analyzing performance related to Power State transitions.

Additional information on how to use PwrTest can be found in the WDK documentation sections:

- PwrTest Documentation
(<http://msdn2.microsoft.com/en-us/library/aa906552.aspx>)
- PwrTest Sleep Scenario
(<http://msdn2.microsoft.com/en-us/library/aa906550.aspx>)

Additional information on how to use the Windows Performance Analysis Tools can be found here:

- Windows Performance Analysis
(<http://go.microsoft.com/fwlink/?LinkId=103276>)

In addition to Microsoft tools there are a number of third-party tools for testing Power State transitions. The following tools are recommended:

- Sleeper, developed by PassMark Software, is a utility to test the ability of PC systems to enter and resume from Sleep and Hibernate:
(<http://www.passmark.com/products/sleeper.htm>)

Additional Resources

- On/Off Transition Performance Analysis of Windows Vista
(http://www.microsoft.com/whdc/system/sysperf/On-Off_Transition.msp)
- Windows Driver Kit
(<http://www.microsoft.com/whdc/DevTools/WDK/WDKpkg.msp>)
- Windows Performance Tools Kit
(<http://www.microsoft.com/whdc/system/sysperf/perftools.msp>)
 - Tools, testing and reliability
 - 1) Windows Performance Analysis Tools
(<http://msdn.microsoft.com/en-us/performance/cc825801.aspx>)
 - 2) Performance Explorer tools and information (MSDN home page)
(<http://go.microsoft.com/fwlink/?LinkId=103276>)
 - 3) Power Event Monitoring Tool Sample Application (WHDC Web site)
(http://www.microsoft.com/whdc/system/pnppwr/powermgmt/PM-apps_samp.msp)
 - 4) PwrTest (WDK documentation)
 - a. PwrTest Battery Scenario (WDK documentation)
(<http://msdn2.microsoft.com/en-us/library/aa906541.aspx>)
 - b. Pwrtest /DISK scenario information in “Windows Disk Idle Detection” (WHDC Web site)
(http://www.microsoft.com/whdc/system/pnppwr/powermgmt/Disk Idle_Detection.msp)
 - c. PwrTest Execution State Scenario (WDK documentation)
(<http://msdn2.microsoft.com/en-us/library/aa906553.aspx>)
 - d. PwrTest Info Scenario (WDK documentation)
 - e. (<http://msdn2.microsoft.com/en-us/library/aa906543.aspx>)
 - f. PwrTest PPM Scenario (WDK documentation)
(<http://msdn2.microsoft.com/en-us/library/aa906546.aspx>)
 - g. PwrTest Sleep Scenario (WDK documentation)
(<http://msdn2.microsoft.com/en-us/library/aa906550.aspx>)
 - 5) Using PowerCfg to evaluate System Energy Efficiency
(<http://www.microsoft.com/whdc/system/pnppwr/powermgmt/powercfg.msp>)
 - 6) On/Off Transition Performance Analysis of Windows Vista
(http://www.microsoft.com/whdc/system/sysperf/On-Off_Transition.msp)
 - 7) Free Tool for S3/S4 stress testing
(<http://www.passmark.com/products/sleeper.htm>)
 - a. OEM or device vendors can test S3 reliability or latency using the free tools
- Various system components that can be tested, including:
 - Operating System - for overall performance and consistent behavior. OS should support basic S3 operations, including capabilities to override some ongoing sleeping tasks to take platform into Sleep.
 - Device drivers - for blocking the system from going into Sleep, not lagging system resume latency.

- Applications - capable of resuming when the system resumes from Sleep, not blocking system from Sleeping.
- Proxy - For systems implementing Ecma Network Proxy, the OS should be able to configure, initiate and terminate operations of the Proxy.
- User experience in terms of system response time should be maintained across multiple Sleep-Resume cycles.

References:

- 1) On/Off Transition Performance Analysis of Windows Vista
([http://www.microsoft.com/whdc/system/sysperf/on-off transition.msp](http://www.microsoft.com/whdc/system/sysperf/on-off_transition.msp))

SECTION 14: GLOSSARY OF TERMS

ACPI - Advanced Configuration and Power Interface specification is an open standard for unified operating system-centric device configuration and power management.

- S0: Working.
- S1: All processor caches are flushed, and the CPU(s) stop executing instructions. Power to the CPU(s) and RAM is maintained; devices that do not indicate they must remain on may be powered down.
- S2: The CPU is powered off.
- S3: Commonly referred to as Standby, Sleep or Suspend to RAM. RAM is still powered.
- S4: Commonly referred to as Hibernation. All content of main memory is saved to non-volatile memory such as a hard drive, and is powered down.

APM - An application programming interface which enables an operating system running an IBM-compatible personal computer to work with the BIOS (part of the computer's firmware) to achieve power management.

BIOS - Basic Input/Output System (BIOS), is a de facto standard defining a firmware interface. The BIOS is boot firmware, designed to be the first code run by a PC when powered on. The initial function of the BIOS is to identify, test and initialize system devices such as the video display card, hard disk and floppy disk and other hardware. BIOS programs are stored on a chip and are built to work with various devices that make up the complementary chipset of the system. They provide a small library of basic input/output functions that can be called to operate and control the peripherals.

BIOS POST - BIOS power on self-test is the common term for a computer's, router's or printer's pre-boot sequence. The same basic sequence is present on all computer architectures. It is the first step of the more general process called initial program load (IPL), booting or bootstrapping. The term POST has become popular in association with and as a result of the proliferation of the PC. It can be used as a noun when referring to the code that controls the pre-boot phase or when referring to the phase itself. It can also be used as a verb when referring to the code or the system as it progresses through the pre-boot phase. Alternatively this may be called "POSTing."

CPU - A central processing unit is an electronic circuit that can execute computer programs.

Ecma TC38-TG4 - Ecma is an international, private (membership-based) nonprofit standards organization for information and communication systems. ECMA TC38-TG4 has been chartered to establish the industry standard for proxy.

IHV - An independent hardware vendor (IHV) is a business term for companies specializing in making or selling computer hardware, usually for niche markets.

ISV - Independent software vendor is a business term for companies specializing in making or selling software, designed for mass marketing or for niche markets.

LAN - Local Area Network is a computer network covering a small physical area, like a home, office or small group of buildings, such as a school, or an airport. The defining characteristics of LANs, in contrast to wide-area networks (WANs), include their usually higher data-transfer rates, smaller geographic place and lack of a need for leased telecommunication lines.

MSDN - The Microsoft Developer Network is the portion of Microsoft responsible for managing the firm's relationship with developers: hardware developers interested in the operating system (OS), developers standing on the various OS platforms, developers leveraging the API and scripting languages of Microsoft's many applications.

NIC - A network card, network adapter, network interface controller, network interface card or LAN adapter is a computer hardware component designed to allow computers to communicate over a computer network.

NDIS - Network Driver Interface Specification is an application programming interface (API) for network interface cards (NICs).

OEM - Original Equipment Manufacturer is typically a company that uses a component made by a second company in its own product, or sells the product of the second company under its own brand.

PM - Power management is a feature of some electrical appliances, especially copiers, computers and computer peripherals such as monitors and printers, that turns off the power or switches the system to a low-power state when inactive.

QFE - A networking hardware acronym standing for Quad Fast Ethernet, a type of vendor-specific network card with four network interfaces.

SW - Software is a general term used to describe a collection of computer programs, procedures and documentation that perform some task on a computer system.

SYN - Synchronize packet in transmission control protocol (TCP).

TCP - Transmission Control Protocol is one of the core protocols of the Internet Protocol Suite. TCP operates at a higher level, concerned only with the two end systems, for example, a Web browser and a Web server. In particular, TCP provides reliable, ordered delivery of a stream of bytes from a program on one computer to another program on another computer.

USB - Universal Serial Bus is a serial bus standard to connect devices to a host computer. USB was designed to allow many peripherals to be connected using a single standardized interface socket and to improve plug and play capabilities by allowing hot swapping; that is, by allowing devices to be connected and disconnected without rebooting the computer or turning off the device.

WDK - The Windows Driver Kit is a software toolset from Microsoft that enables the development of device drivers for the Microsoft Windows platform. It includes documentation, samples, build environments and tools for driver developers.

WHQL - Windows Hardware Quality Labs Testing or WHQL Testing is a process which involves running a series of tests on third-party (i.e., non-Microsoft) hardware or software, and then submitting the log files from these tests to Microsoft for review.

WLAN - A wireless LAN is a wireless local area network that links two or more computers or devices using spread-spectrum or OFDM modulation technology based to enable communication between devices in a limited area. This gives users the mobility to move around within a broad coverage area and still be connected to the network.

WoL - Wake on LAN is an Ethernet computer-networking standard that allows a computer to be turned on or woken up remotely by a network message sent usually by a simple program executed on another computer on the network.

