White Paper on Power-Saving Modes of Microsoft® Windows Server® 2008 SP2 on BladeSymphony 320

September 2009

Rev.1.0.0

Hitachi Ltd.

Enterprise Server Division

Content

1 Overview

In recent years, the growing use of IT and information processing has led to an increase in server numbers and performance. This has resulted in a sizeable increase in power consumption, and the cost of electricity has become a significant factor. Reducing power consumption is now a major cost-saving issue.

Both Hitachi and Microsoft have been addressing this problem, and have been actively developing new technologies to improve the power efficiency of servers. Hitachi was amongst the first manufacturers to adopt the Intel® Xeon® processor 5500 series, with its improved power saving features, in the BladeSymphony product line. Microsoft over the years has introduced many improvements to the power efficiency of its Windows Server® operating system, and added further power saving features to Microsoft® Windows Server® 2008.

In this paper we evaluate the power saving features of Microsoft® Windows Server® 2003 R2 SP2 (Windows Server 2003) and Microsoft® Windows Server® 2008 SP2[1] (Windows Server 2008) on Hitachi BladeSymphony 320 A4 server blades equipped with the latest Intel® Xeon® 5500 series processors. The evaluation shows that Windows Server 2008 requires approximately 10% less power[2] than Windows Server 2003. An evaluation was also made of how the power options of Windows Server 2008 and changes in the BIOS settings of server blades affect the power consumption. As a result, it was demonstrated that optimal settings can be selected based on usage profiles and that no external power measuring devices are required for the optimization, since BladeSymphony servers are equipped with a power monitoring function[3].

2 Improved Power-Saving Modes in Windows Server 2008

In general, it is the processor that consumes the most power in a server, so the greatest server power savings are gained by reducing the power consumption of the processors. A standard called the Advanced Configuration and Power Interface (ACPI) provides an interface that allows the OS to control the power consumption of servers by means of Processor Performance State (P-State) and Power State (C-State). The P-State defines the relationship between performance and power consumption of the processor, while C-State

[1] Windows Server 2008 SP2 was not officially supported at the time of writing this document. The evaluation is based on SP2 Beta edition.

[2] Approximately 10% less power was consumed to complete the same processing.

[3] Function for monitoring the power consumption (average, maximum, and minimum values) of individual server blades and server chassis. See BladeSymphony User's Guide for the models that support this function and its detailed usage.

defines the relationship between the idle state and power consumption of processor. An OS complying with the ACPI standards can optimize the relationship between processing performance and power consumption by controlling P-State and C-State.

Both Windows Server 2003 and Windows Server 2008 support the ACPI standards and can control P-State and C-State. However, Windows Server 2008 can save more power than Windows Server 2003 based on the following enhancements:.

(1) Multiprocessor support

The P-State and C-State controls of Windows Server 2003 support only single processors, whereas Windows Server 2008 supports multiple processors and can individually control processors and cores.

(2) Improved P-State and C-State Controls

The algorithms of P-State and C-State control have been refined, and power management in Windows Server 2008 improved over that of Windows Server 2003.

(3) Effective P-State Control with Power Options (Default)

The user can select a method for P-State control from the Power Options window of Windows Server 2003 and Windows Server 2008, but the default values are different: Windows Server 2003 does not control P-State by default, whereas Windows Server 2008 does. Due to this difference, Windows Server 2008 is able to save power by P-State control immediately after OS installation. Table 1 shows the available Power Options.

Table 1: Power Options of Windows

Power options		Settings
Windows Server 2003	Windows Server 2008	
Always On (Default)	High Performance	Always sets P-State to “P0” and demands high performance.
Server Balanced Processor Power and Performance	Balanced (Default)	Sets P-State appropriately and balances performance and power consumption.
-	Power Saver	Always sets P-State to “Pn”. Power consumption is reduced but performance decreases.

3 Power-Saving Effect of Windows Server 2008

We measured and compared the power consumptions of Windows Server 2003 and Windows Server 2008 installed on BladeSymphony 320 A4 server blades.

3.1 System Configuration

Table 2 shows the configuration of the BladeSymphony 320 server used for the evaluation.

Table 2: System configuration of BladeSymphony 320

Server blade	One BladeSymphony 320 A4 server blade [Configuration of server blade] ・Processor 2 Intel® Xeon® X5570 (2.93GHz/8M) processors ・Memory 24GB (6 DDR3 Registered DIMM 4GB) ・Internal disk 2 SAS disks (147GB, 10krpm) ・Network interface 4 ports of 1 Gbit-Ethernet (SERDES)
OS	Microsoft® Windows Server® 2003 R2 Enterprise x64 Edition SP2 And Microsoft® Windows Server® 2008 Enterprise Edition Server SP2 (Beta)
Load generator software	Transaction processing simulator

The BladeSymphony 320 A4 server blade was equipped with two Xeon X5570 processors, six 4GB memory RDIMMs, and two internal HDDs. The latest BladeSymphony servers can monitor the power and measure the power consumption of server blade without special devices.

We generated load on the server blade with application software and simulated transaction processing. This software can adjust the amount of load on servers, so the power consumption for a range of system loads could be evaluated. We measured the power consumption and the transaction processing performance of ten different loads in a range from lowest to highest.

3.2 Power Consumption of BladeSymphony 320 A4 Model

Figure 1 shows the measured power consumption of the BladeSymphony 320 A4 server blade. The default Power Option “Balanced” was selected in Windows Server 2008, and the equivalent setting of “Server Balanced Processor Power and Performance” was selected in Windows Server 2003. For the remaining settings the default values were kept.

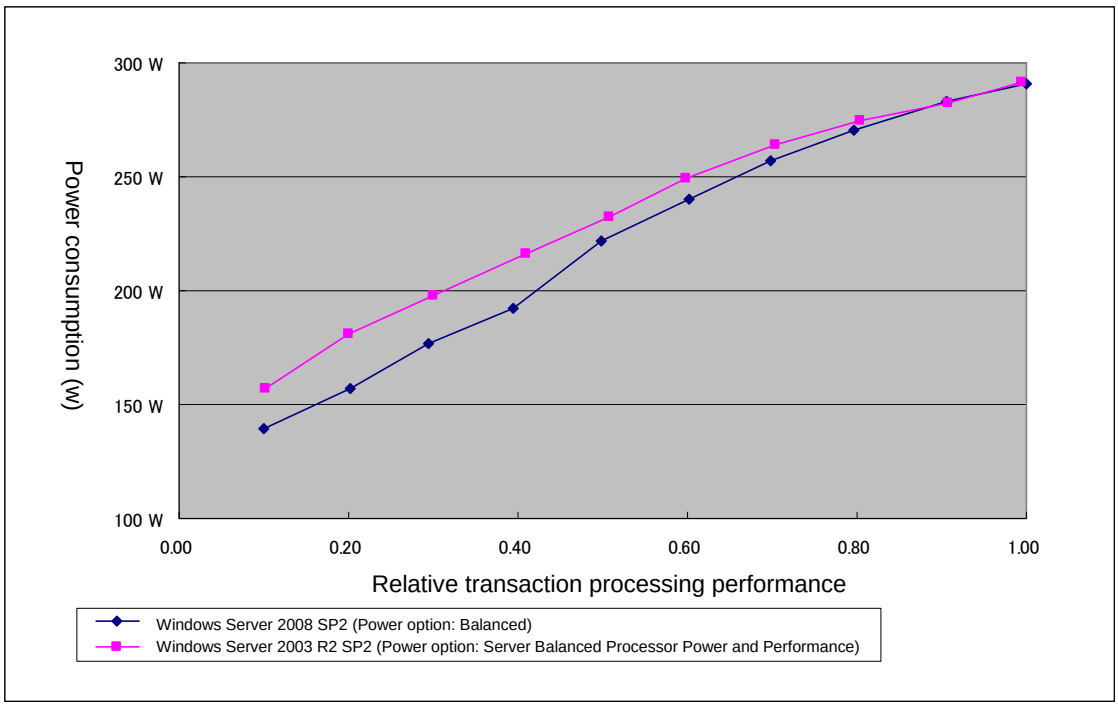

Figure 1 Power Consumption of BladeSymphony 320 A4 Server Blade

The horizontal axis of Figure 1 represents the relative transaction processing performance[4] for each measurement point when the maximum performance obtained with this system configuration is normalized to 1. The vertical axis shows the server power consumption. With this figure, the power consumption per processing performance can be compared. The power consumption of Windows Server 2008 is significantly reduced when the processing performance is below 0.8. With a load of 0.4 or below, the power consumption decreases by approximately 10%.

[4] In Figure 1, the maximum transaction processing was achieved when the highest load was generated by load generator software in Windows Server 2008.

A server chassis of BladeSymphony 320 can hold up to ten server blades, as well as other electric components such as management modules, switches and fans . Table 3 below presents the total power consumption for a chassis populated with ten server blades, using the results presented in Figure 1.

Table 3 Total Power Consumption of Server Chassis Installed with Ten Server Blades of BladeSymphony 320 A4 Model

Relative Transaction Processing Performance	0.2	0.4	0.8
Microsoft® Windows Server® 2003 R2 SP2	2.5kW	2.9kW	3.6kW
Microsoft® Windows Server® 2008 SP2	2.3kW	2.7kW	3.5kW

(Note 1) The above is based on a theoretical configuration of a server chassis with one management module, one 1GbpsLAN pass-through module, five cooling fan modules, and four power supply modules installed.

(Note 2) Internal data (evaluating the relation between the power consumption of blade and server chassis) was used for estimating the power consumption. A power supply of 200V and internal temperature of 25 degrees Celsius are assumed.

Overall power consumption of a populated server chassis was reduced by approximately 7% (relative performance = 0.4) when upgrading Windows Server 2003 to Windows Server 2008.

4 Effects of Power Options and BIOS Settings on Power Consumption

Windows Server 2008 offers three Power Options as shown in Table 1. Some BIOS settings of the blade server affect the operating frequency of the processors and influence the processing performance and power consumption of the server. These settings can be optimized according to the usage profile if the power consumption of server during operation is known. The latest BladeSymphony servers can monitor the power and measure the server power consumption without using external devices (Figure 2), therefore the settings are easily optimized.

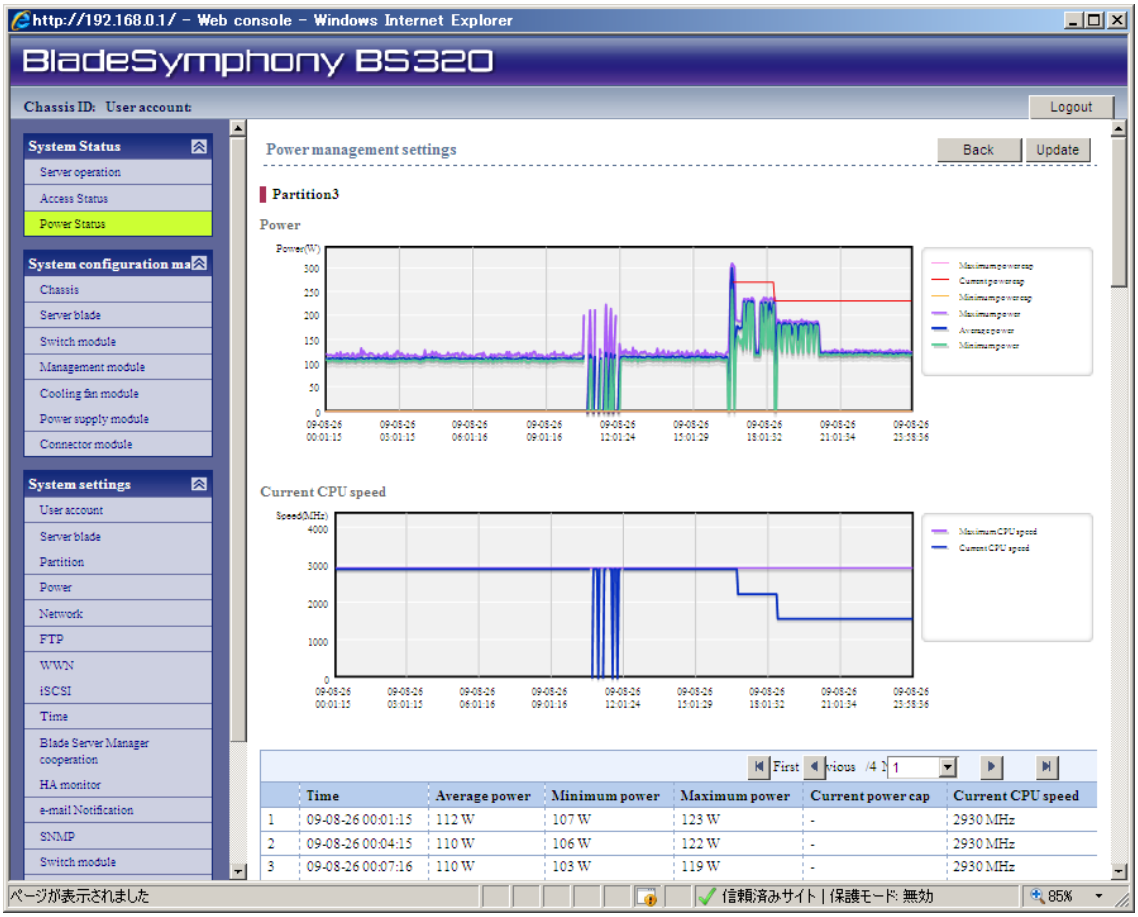

Figure 2 Sample Screen of Power Monitoring Function
(Server Blade Power Consumption History)

We evaluated how the Power Options of Windows Server 2008 and BIOS settings of the server blade affect the performance and power consumption of a server blade that was configured as described in Section 3.1.

Figure 5 shows the measured power consumption based on combinations of the Power Options of “High Performance” / “Balanced (default)” in Windows Server 2008 (Figure 3) and “Enable (default)” / ”Disable” of Turbo Boost technology[5] in the BIOS settings (Figure 4).

[5] Intel Xeon processor 5500 series of BladeSymphony 320 A4 model offers Intel Turbo Boost technology. Processing speed may increase when the processor frequency is increased under certain conditions. Changes in operating frequency affect the power consumption, and the effects of Turbo Boost technology vary significantly depending on the application.

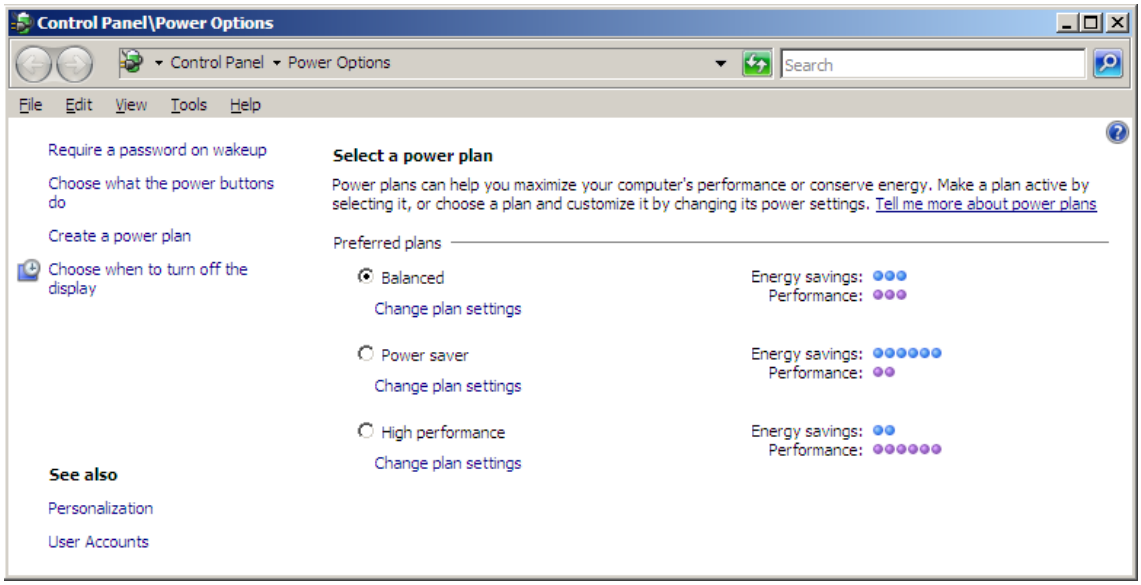

Figure 3　Power Options Setting of Windows Server 2008

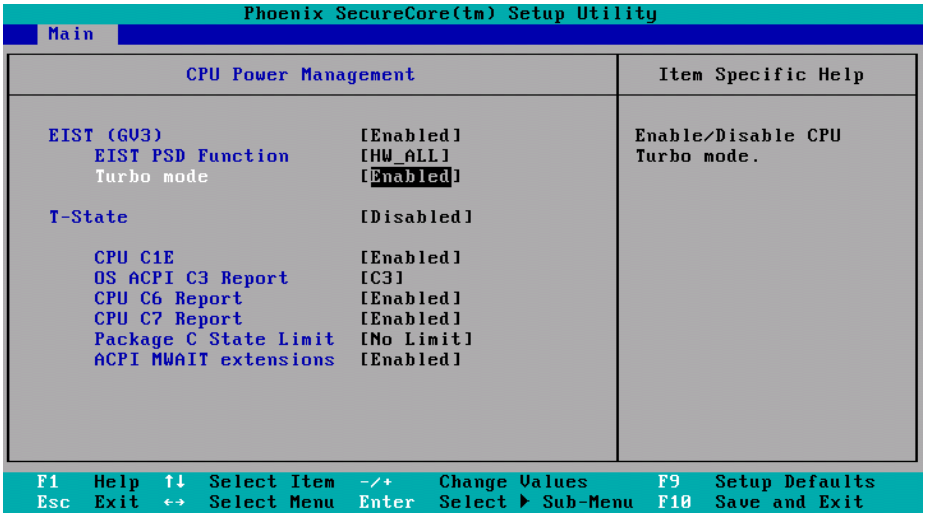

Figure 4　BIOS Settings of BladeSymphony 320 A4 Model

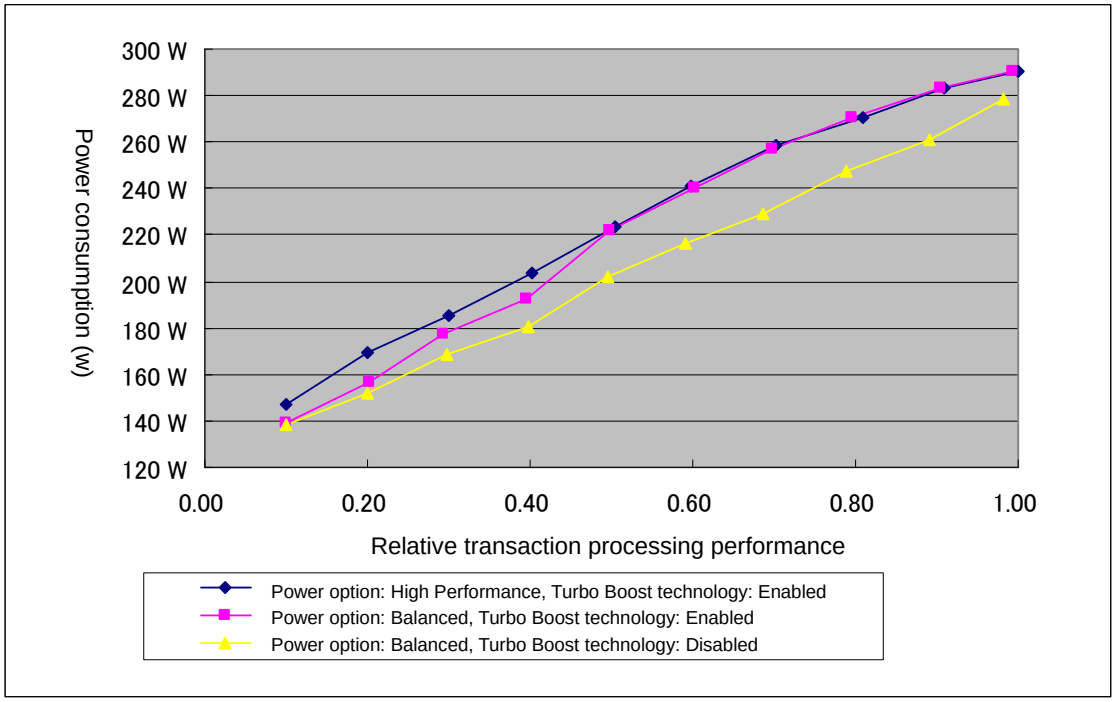

Figure 5 Power Consumption of BladeSymphony 320 A4 Model (Windows Server 2008)

When changing the power option from “Balanced” (default) to “High Performance”, the power consumption increased when the relative performance was below approximately 0.5. The total power consumption was reduced and the maximum ratio of transaction processing performance decreased by approximately 2% when the Turbo Boost technology settings in the BIOS were changed from “Enable” (default) to “Disable”. The Turbo Boost technology was not effective with the load generation used in this evaluation.

Based on the above results, the following settings for Power Options and BIOS are considered to be optimal:.

- Set the Power Options of Windows Server 2008 to “Balanced” and Turbo Boost technology of BIOS to “Enabled” to achieve high performance while reducing the power consumption when the processor is not operating at full workload. These values are set by default, so the user does not have to change OS or BIOS settings.
- Change the Turbo Boost technology of BIOS settings to “Disabled” to lower the overall power consumption while maintaining acceptable performance.

Changes in the settings of Windows Server 2008 or the BIOS will affect power consumption differently depending on the application, so the actual target application should be used for this evaluation. The BladeSymphony power monitoring function significantly eases this evaluation.

5 Summary

The following points were clarified when assessing the power saving effectiveness of Windows Server 2008 on the BladeSymphony 320 A4 model.

- When operating a server blade under the given operating conditions[6], Windows Server 2008 can use less power than Windows Server 2003.. In the given evaluation environment, Windows Server 2008 is able to reduce the power consumption by approximately 10% compared with Windows Server 2003. This indicates that users of Windows Server 2003 seeking to save server power could achieve improvements simply by upgrading to Windows Server 2008.
- The processing performance and power consumption of server blades changes depending on the Power Options of Windows Server 2008 and BIOS settings. The relationship between processing performance and power consumption described in section 4 can be used as a reference for the optimization of these settings. The latest generation of BladeSymphony servers can monitor the power consumption as required for this optimization without the need for additional power measuring equipment.

[6] Evaluation using the same server blade with the same BIOS settings (default) and the same Windows Power Option settings (with P-State control).