

Windows Server 2008 SPECjbb2005 performance and power consumption on Dell, HP, and IBM blade servers

Executive summary

Dell Inc. (Dell) commissioned Principled Technologies (PT) to measure the SPECjbb2005 performance per watt of Windows Server 2008 Enterprise x64 Edition on the following dual-socket blade servers:

- Dell PowerEdge M600
- HP BladeSystem c-Class
- IBM BladeCenter H Type 8852

In our testing with Windows Server 2008, the Dell PowerEdge M600 produced better performance/watt than the HP BladeSystem c-Class and the IBM BladeCenter H Type 8852 at each blade count we tested.

Figure 1 presents the performance/watt for each of the blade server systems with Windows Server 2008 by configuration. Higher results show better system performance/watt. We calculate performance/watt by dividing the SPECjbb2005 result by the average power consumption in watts during the period the system achieved peak performance. For complete details of the performance of each Java Virtual Machine (JVM) by warehouse for each server, see the Test results section.

Figure 1 shows the blades in the Dell PowerEdge M600 achieved the best performance/watt at every configuration under Windows Server 2008. With 10 blades installed in all three systems, the Dell PowerEdge M600 achieved 23.64 percent higher performance/watt than the HP BladeSystem c-Class. The Dell PowerEdge M600 also achieved 29.92 percent higher performance/watt than the IBM BladeCenter H Type 8852.

KEY FINDINGS

- On Windows Server 2008, the Dell PowerEdge M600 achieved better performance/watt than the HP BladeSystem c-Class or the IBM BladeCenter H Type 8852 at every configuration we tested (see Figure 1).
- With 10 blades installed and running Windows Server 2008 in all three systems, the blades in the Dell PowerEdge M600 achieved 23.64 percent higher performance/watt than the HP BladeSystem c-Class and 29.92 percent higher than the IBM BladeCenter H Type 8852.
- With 16 blades in each chassis, the Dell PowerEdge M600 achieved 21.08 percent better performance/watt than the HP BladeSystem c-Class.
- At the maximum blades configuration, the Dell PowerEdge M600 achieved 27.66 percent better performance/watt with 16 blades than the IBM BladeCenter H Type 8852 did with 14 blades.
- The Dell PowerEdge M600 used 9.59 percent less power per blade with 16 blades than the HP BladeSystem c-Class did with 16 blades and 10.82 percent less power per blade than the IBM BladeCenter H Type 8852 did with 14 blades.

	Dell PowerEdge M600 Blade System	HP BladeSystem c-Class	IBM BladeCenter H Type 8852	Percentage performance/watt increase Dell over HP	Percentage performance/watt increase Dell over IBM
1 blade	456.5	335.9	249.2	35.92	83.18
2 blades	642.95	482.2	383.7	33.33	67.55
10 blades	891.04	720.66	685.82	23.64	29.92
Maximum blades	914.91 (16 blades)	755.6 (16 blades)	716.66 (14 blades)	21.08	27.66

Figure 1: Performance/watt results for each server by blade configuration. Higher numbers are better.

In the maximum blades configuration, we installed the maximum number of blade servers in each chassis: 16 for the Dell and HP and 14 for the IBM. In the maximum blades configuration, the Dell PowerEdge M600 achieved 21.08 percent better performance/watt with 16 blades installed than the HP BladeSystem c-Class did with 16 blades installed. The Dell PowerEdge M600 achieved 27.66 percent better performance/watt with 16 blades installed than the IBM BladeCenter H Type 8852 with 14 blades installed. In comparing the maximum blades configuration for the Dell PowerEdge M600 and the IBM BladeCenter H Type 8852, we are normalizing a 16-blade result to a 14-blade result.

Furthermore, the Dell PowerEdge M600 achieved 35.92 and 33.33 percent higher performance/watt than the HP BladeSystem c-Class in one- and two-blade configurations, respectively. The Dell PowerEdge M600 achieved 83.18 and 67.55 percent higher performance/watt than the IBM BladeCenter H Type 8852 in the one- and two-blade configurations, respectively.

Workload

SPECjbb2005 is an industry-standard benchmark created by the Standard Performance Evaluation Corp. (SPEC) to measure a server's Java performance. (Note: SPEC and the SPECjbb2005 are trademarks of the Standard Performance Evaluation Corporation.) SPEC modeled SPECjbb2005 on the three-tier client/server architecture, with the middle layer as the primary focus. According to SPEC, "Random input selection represents the first (user) tier. SPECjbb2005 fully implements the middle tier business logic. The third tier is represented by tables of objects, implemented by Java Collections, rather than a separate database."

(www.spec.org/jbb2005/docs/UserGuide.html).

SPECjbb2005 utilizes multiple special data groups and multiple threads as it runs. Each data unit is a "warehouse," which is a roughly 25MB collection of data objects. Each thread represents an active user posting transaction requests within a warehouse. The benchmark run begins with one warehouse and then increases the number of warehouses; its goal is to saturate the server's processor capacity. As the number of warehouses increases, so does the number of threads. The benchmark's results portray the server's throughput in business operations per second or SPECjbb2005 bops. A higher number of SPECjbb2005 bops is better. (For more information on SPECjbb2005, go to www.spec.org.)

Test results

For testing, we installed a given number of blade servers into the chassis and ran SPECjbb2005 on all servers on Windows Server 2008. Before starting the SPECjbb2005 benchmark, we logged into the system and allowed the servers to sit idle for 10 minutes. We then started recording power for 2 minutes. This process meant that all systems were idle for 12 minutes before we began the benchmark.

In each test configuration, we ran two JVM instances at the same time, a common practice on servers with many processors. To compute the overall score for the system, SPECjbb2005 sums the scores of all the JVMs. SPECjbb2005 computes the score of each JVM by taking the average of the results during mixes when the server is running at peak performance. In our testing, all servers achieved peak performance during mixes 4 through 8. (In SPEC's terms, these results are from "compliant" runs, meaning that we can disclose them publicly without posting them on the SPEC Web site with all the files SPEC usually requires. We do present here all the data necessary to reproduce these results.) In the tables below, we show the SPECjbb2005 results for each blade for a given configuration.

Figure 2 shows the idle power usage (in watts) for the Dell PowerEdge M600, HP BladeSystem c-Class, and IBM BladeCenter H Type 8852 at all blade configurations. Lower power is better.

	Dell PowerEdge M600	HP BladeSystem c-Class	IBM BladeCenter H Type 8852
1 blade	388.0	514.5	674.3
2 blades	505.2	668.8	817.0
10 blades	1,583.2	1,981.7	1,898.3
Maximum blades	2,427.6 (16 blades)	3,006.8 (16 blades)	2,447.0 (14 blades)

Figure 2: Idle power usage (in watts) of the test servers before the median peak runs for each blade configuration. Lower numbers are better.

Figure 3 shows the average power usage (in watts) for Dell PowerEdge M600, HP BladeSystem c-Class, and IBM BladeCenter H Type 8852 at all blade configurations. Lower power is better. To calculate the average power, we

recorded the power during the SPECjbb2005 benchmark and averaged the power during the period the system achieved peak performance. The power the Dell PowerEdge M600 required running with 16 blades is almost identical to the power the IBM BladeCenter H Type 8853 required running with 14 blades.

	Dell PowerEdge M600	HP BladeSystem c-Class	IBM BladeCenter H Type 8852
1 blade	467.0	580.2	751.3
2 blades	662.6	803.3	974.7
10 blades	2,383.5	2,695.2	2,727.9
Maximum blades	3,720.4 (16 blades)	4,115.4 (16 blades)	3,650.1 (14 blades)

Figure 3: Average total power usage (in watts) of the test servers during the median peak runs for each blade configuration. Lower numbers are better.

To calculate the performance/watt we used the following formula:

Performance/watt = the benchmark's score/average power consumption in watts during the period the system achieved peak performance.

For the 2-, 10-, and maximum-blade configurations, we divided the average power shown in Figure 3 by the number of blades. We then divided the benchmark's score for each blade by the average power. The formula for these blade configurations were as follows:

Performance/watt by blade = (benchmark score by blade/[total average power/number of blades])

We then averaged the performance/watt for all blades in the given configurations.

For each configuration, we performed three runs of SPECjbb2005 and recorded the power during these runs. The results below are the median of three test runs.

Figure 4 shows the SPECjbb2005 results, average power per blade, and performance/watt for the Dell PowerEdge M600 with one blade installed in the chassis.

	SPECjbb2005 bops	SPECjbb2005 bops/JVM	Average power per blade watts)	Performance/watt
System 1	213,180	106,590	467.0	456.5

Figure 4: SPECjbb2005 results, average power usage (in watts), and performance/watt for the Dell PowerEdge M600 during the median run for the one-blade configuration. Higher performance/watt is better.

Figure 5 shows the SPECjbb2005 results, average power per blade, and performance/watt for the Dell PowerEdge M600 with two blades installed in the chassis. We calculated the average performance/watt by averaging the performance/watt scores of the two systems.

	SPECjbb2005 bops	SPECjbb2005 bops/JVM	Average power per blade watts)	Performance/watt
System 1	212,776	106,388	331.31	642.23
System 2	213,249	106,625	331.31	643.66
Average performance/watt				642.95

Figure 5: SPECjbb2005 results, average power usage (in watts), and performance/watt for the Dell PowerEdge M600 during the median run for the two-blade configuration. Higher performance/watt is better.

Figure 6 shows the SPECjbb2005 results, average power per blade, and performance/watt for the Dell PowerEdge M600 with 10 blades installed in the chassis. We calculated the average performance/watt by averaging the performance/watt scores of the 10 systems.

	SPECjbb2005 bops	SPECjbb2005 bops/JVM	Average power per blade (watts)	Performance/watt
System 1	212,280	106,140	238.35	890.63
System 2	205,765	102,883	238.35	863.29
System 3	213,357	106,679	238.35	895.15
System 4	213,940	106,970	238.35	897.59
System 5	213,801	106,901	238.35	897.01
System 6	212,492	106,246	238.35	891.52
System 7	212,346	106,173	238.35	890.90
System 8	213,200	106,600	238.35	894.49
System 9	214,281	107,141	238.35	899.02
System 10	212,318	106,159	238.35	890.79
Average performance/watt				891.04

Figure 6: SPECjbb2005 results, average power usage (in watts), and performance/watt for the Dell PowerEdge M600 during the median run for the 10-blade configuration. Higher performance/watt is better.

Figure 7 shows the SPECjbb2005 results, average power per blade, and performance/watt for the Dell PowerEdge M600 with 16 blades installed in the chassis. We calculated the average performance/watt by averaging the performance/watt scores of the 16 systems.

	SPECjbb2005 bops	SPECjbb2005 bops/JVM	Average power per blade (watts)	Performance/watt
System 1	214,683	107,342	232.53	923.26
System 2	212,560	106,280	232.53	914.13
System 3	213,299	106,650	232.53	917.31
System 4	212,896	106,448	232.53	915.58
System 5	212,538	106,269	232.53	914.04
System 6	213,004	106,502	232.53	916.04
System 7	212,856	106,428	232.53	915.41
System 8	213,971	106,986	232.53	920.20
System 9	212,781	106,391	232.53	915.08
System 10	214,006	107,003	232.53	920.35
System 11	213,322	106,661	232.53	917.41
System 12	211,849	105,925	232.53	911.08
System 13	213,416	106,708	232.53	917.81
System 14	205,347	102,674	232.53	883.11
System 15	214,192	107,096	232.53	921.15
System 16	213,116	106,558	232.53	916.52
Average performance/watt				914.91

Figure 7: SPECjbb2005 results, average power usage (in watts), and performance/watt for the Dell PowerEdge M600 during the median run for the 16-blade configuration. Higher performance/watt is better.

Figure 8 shows the SPECjbb2005 results, average power per blade, and performance/watt for the HP BladeSystem c-Class with one blade installed in the chassis.

	SPECjbb2005 bops	SPECjbb2005 bops/JVM	Average power per blade watts)	Performance/ watt
System 1	194,879	97,440	580.2	335.9

Figure 8: SPECjbb2005 results, average power usage (in watts), and performance/watt for the HP BladeSystem c-Class during the median run for the one-blade configuration. Higher performance/watt is better.

Figure 9 shows the SPECjbb2005 results, average power per blade, and performance/watt for the HP BladeSystem c-Class with two blades installed in the chassis. We calculated the average performance/watt by averaging the performance/watt scores of the two systems.

	SPECjbb2005 bops	SPECjbb2005 bops/JVM	Average power per blade watts)	Performance/ watt
System 1	192,835	96,418	401.63	480.13
System 2	194,517	97,259	401.63	484.31
Average performance/watt				482.22

Figure 9: SPECjbb2005 results, average power usage (in watts), and performance/watt for the HP BladeSystem c-Class during the median run for the two-blade configuration. Higher performance/watt is better.

Figure 10 shows the SPECjbb2005 results, average power per blade, and performance/watt for the HP BladeSystem c-Class with 10 blades installed in the chassis. We calculated the average performance/watt by averaging the performance/watt scores of the 10 systems.

	SPECjbb2005 bops	SPECjbb2005 bops/JVM	Average power per blade watts)	Performance/ watt
System 1	193,720	96,860	270.13	717.14
System 2	194,069	97,035	270.13	718.43
System 3	195,288	97,644	270.13	722.94
System 4	194,702	97,351	270.13	720.77
System 5	194,722	97,361	270.13	720.85
System 6	195,312	97,656	270.13	723.03
System 7	194,755	97,378	270.13	720.97
System 8	194,986	97,493	270.13	721.83
System 9	194,137	97,069	270.13	718.68
System 10	195,012	97,506	270.13	721.92
Average performance/watt				720.66

Figure 10: SPECjbb2005 results, average power usage (in watts), and performance/watt for the HP BladeSystem c-Class during the median run for the 10-blade configuration. Higher performance/watt is better.

Figure 11 shows the SPECjbb2005 results, average power per blade, and performance/watt for the HP BladeSystem c-Class with 16 blades installed in the chassis. We calculated the average performance/watt by averaging the performance/watt scores of the 16 systems.

	SPECjbb2005 bops	SPECjbb2005 bops/JVM	Average power per blade watts)	Performance/watt
System 1	194,624	97,312	257.19	756.73
System 2	193,347	96,674	257.19	751.76
System 3	194,245	97,123	257.19	755.25
System 4	194,567	97,284	257.19	756.50
System 5	194,650	97,325	257.19	756.83
System 6	195,350	97,675	257.19	759.55
System 7	194,641	97,321	257.19	756.79
System 8	194,543	97,272	257.19	756.41
System 9	195,365	97,683	257.19	759.61
System 10	194,144	97,072	257.19	754.86
System 11	194,916	97,458	257.19	757.86
System 12	193,660	96,830	257.19	752.98
System 13	194,268	97,134	257.19	755.34
System 14	194,111	97,056	257.19	754.73
System 15	193,873	96,937	257.19	753.81
System 16	193,152	96,576	257.19	751.00
Average performance/watt				755.63

Figure 11: SPECjbb2005 results, average power usage (in watts), and performance/watt for the HP BladeSystem c-Class during the median run for the 16-blade configuration. Higher performance/watt is better.

To try and better understand the HP's performance, we experimented with changing certain options in the system BIOS. The HP blades have an option in their BIOS called "Power Regulator for ProLiant." There are four options you can choose from:

- o HP Dynamic Power Savings Mode
- o HP Static Low Power Mode
- o HP Static High Performance Mode
- o OS Control Mode

HP Dynamic Power Savings Mode is the default setting. When we set the system to HP Static High Performance Mode, we see performance and power increase to roughly what they were under Windows Server 2003. Figure 12 shows the results of our experiment for a single blade comparing Windows Server 2003 versus Windows Server 2008 with different settings of the Power Regulator for ProLiant BIOS setting.

	Windows Server 2003 HP Dynamic Power Savings Mode (default)	Windows Server 2008 HP Dynamic Power Savings Mode (default)	Windows Server 2008 HP Static High Performance Mode
Idle power (watts)	511.0	512.5	509.9
Average power (watts)	609.6	582.3	609.2
SPECjbb2005 bops	206,993	194,362	207,430
Performance/watt	339.6	333.8	340.5

Figure 12: Comparison of Windows Server 2003 and Windows Server 2008 results with different BIOS settings showing idle power (in watts), average power usage (in watts), SPECjbb2005 results, and performance/watt for the HP BladeSystem c-Class. Higher performance/watt is better.

From these results, Windows Server 2008 and Windows Server 2003 appear to utilize this setting differently. While the results did change with the new BIOS option, the change in power and performance are about equal, so there is not much affect on the performance per watt (about 2 percent).

Figure 13 shows the SPECjbb2005 results, average power per blade, and performance/watt for the IBM BladeCenter H Type 8852 with one blade installed in the chassis.

	SPECjbb2005 bops	SPECjbb2005 bops/JVM	Average power per blade watts)	Performance/watt
System 1	187,225	93,613	751.3	249.2

Figure 13: SPECjbb2005 results, average power usage (in watts), and performance/watt for the IBM BladeCenter H Type 8852 during the median run for the one-blade configuration. Higher performance/watt is better.

Figure 14 shows the SPECjbb2005 results, average power per blade, and performance/watt for the IBM BladeCenter H Type 8852 with two blades installed in the chassis. We calculated the average performance/watt by averaging the performance/watt scores of the two systems.

	SPECjbb2005 bops	SPECjbb2005 bops/JVM	Average power per blade watts)	Performance/watt
System 1	186,989	93,495	487.36	383.68
System 2	187,052	93,526	487.36	383.81
Average performance/watt				383.74

Figure 14: SPECjbb2005 results, average power usage (in watts), and performance/watt for the IBM BladeCenter H Type 8852 during the median run for the two-blade configuration. Higher performance/watt is better.

Figure 15 shows the SPECjbb2005 results, average power per blade, and performance/watt for the IBM BladeCenter Type H 8852 with 10 blades installed in the chassis. We calculated the average performance/watt by averaging the performance/watt scores of the 10 systems.

	SPECjbb2005 bops	SPECjbb2005 bops/JVM	Average power per blade watts)	Performance/watt
System 1	187,614	93,807	272.79	687.75
System 2	187,150	93,575	272.79	686.05
System 3	186,934	93,467	272.79	685.26
System 4	187,158	93,579	272.79	686.08
System 5	186,814	93,407	272.79	684.82
System 6	186,946	93,473	272.79	685.30
System 7	187,726	93,863	272.79	688.16
System 8	186,866	93,433	272.79	685.01
System 9	186,638	93,319	272.79	684.17
System 10	187,033	93,517	272.79	685.62
Average performance/watt				685.82

Figure 15: SPECjbb2005 results, average power usage (in watts), and performance/watt for the IBM BladeCenter H Type 8852 during the median run for the 10-blade configuration. Higher performance/watt is better.

Figure 16 shows the SPECjbb2005 results, average power per blade, and performance/watt for the IBM BladeCenter H Type 8852 with 14 blades installed in the chassis. We calculated the average performance/watt by averaging the performance/watt scores of the 14 systems.

	SPECjbb2005 bops	SPECjbb2005 bops/JVM	Average power per blade watts)	Performance/watt
System 1	188,509	94,255	260.72	723.02
System 2	186,951	93,476	260.72	717.04
System 3	187,189	93,595	260.72	717.96
System 4	187,007	93,504	260.72	717.26
System 5	187,388	93,694	260.72	718.72
System 6	187,502	93,751	260.72	719.16
System 7	186,821	93,411	260.72	716.54
System 8	186,690	93,345	260.72	716.04
System 9	186,168	93,084	260.72	714.04
System 10	187,140	93,570	260.72	717.77
System 11	186,177	93,089	260.72	714.07
System 12	186,125	93,063	260.72	713.87
System 13	186,577	93,289	260.72	715.61
System 14	185,655	92,828	260.72	712.07
Average performance/watt				716.66

Figure 16: SPECjbb2005 results, average power usage (in watts), and performance/watt for the IBM BladeCenter H Type 8852 during the median run for the 14-blade configuration. Higher performance/watt is better.

Concerned about IBM's relatively low SPECjbb scores, we investigated further. We verified that the system's processors were running at full speed and there were no heat issues.

We tested all systems using 4GB of RAM. The IBM BladeCenter H Type 8852 had only four RAM slots, all of which were full. The Dell PowerEdge and HP BladeSystem each had eight RAM slots, of which we used 4. RAM was running at 667 MHz for all three test systems.

As a sanity check, we ran SiSoftware Sandra's memory test, which reported the IBM server's memory bandwidth as half that of the Dell server's: 10,656 MB/s vs. 21,312 MB/s. SiSoftware Sandra also reported the IBM server had two memory channels, while the Dell server had four memory channels.

Sandra showed only that the IBM BladeCenter H Type 8852 used an Intel 5000P chipset, with no code names or other details. It did report that the memory controller supports eight memory slots, the same as the Dell and the HP. It is possible that using the optional IBM BladeCenter Memory and I/O Expansion Blade, which provides four additional DIMM slots, is necessary to access the other two memory channels. The IBM documentation says that when the using the Memory expansion blade, the optimal memory configuration is to use two pairs of matching DIMMs; one pair on the board in slots 1 and 3 and the other pair on the expansion blade in slots 5 and 7. Because the system we tested did not ship with this part, we were unable to test this hypothesis.

In conclusion, we believe that the IBM BladeCenter H Type 8852's relatively low SPECjbb scores are probably a function of its lower memory bandwidth.

Test methodology

Perform the following BIOS-level operations on every system before installing the OS:

1. Make sure to configure systems with RAID 1. Use the disk controller utility for this, not the OS.
2. Set the partition to be the whole disk.
3. Disable HW prefetcher and Adjacent line prefetcher in BIOS. Leave all other values at their defaults.

Create the base image:

1. For each installation, begin by installing a fresh copy of Microsoft Windows Server 2003 R2, Enterprise x64 Edition Service Pack 2 on each blade server. (If you are using a support pack with an

express install option, do *not* use it. Select Custom Install and only install drivers. Otherwise, you may be installing unnecessary software, which may affect the results of the test.)

- a. For the licensing mode, use the default setting of five concurrent connections.
 - b. Enter a password for the administrator logon.
 - c. Select Eastern Time Zone, and check date and time.
 - d. Use typical settings for the Network installation.
 - e. Assign a computer name. We used the format “<Brand>Server<#>”, where Brand is Dell, HP, or IBM and X is the blade number in the chassis (1–16 for Dell and HP and 1–14 for IBM) (e.g., IBMServer1).
 - f. Leave the default “WORKGROUP” for the workgroup.
 - g. Finish installation.
 - h. Install SP 2.
 - i. Run Live update and install the following updates. Our update date was November 29, 2007.
 - Windows Server 2003 Security Update for Windows Server 2003 x64 Edition (KB943460)
 - Windows Server 2003 Windows Malicious Software Removal Tool x64 - November 2007 (KB890830)
 - Windows Server 2003 Cumulative Security Update for Internet Explorer 6 for Windows Server 2003 x64 Edition (KB939653)
 - Windows Server 2003 Security Update for Outlook Express for Windows Server 2003 x64 Edition (KB941202)
 - Windows Server 2003 Security Update for Windows Server 2003 x64 Edition (KB933729)
 - Windows Server 2003 Security Update for Windows Server 2003 x64 Edition (KB936021)
 - Windows Server 2003 Update for Windows Server 2003 x64 Edition (KB933360)
 - Windows Server 2003 Security Update for Windows Server 2003 x64 Edition (KB938127)
 - Windows Server 2003 Security Update for Windows Server 2003 x64 Edition (KB921503)
 - Windows Server 2003 Security Update for Windows Server 2003 x64 Edition (KB936782)
 - Windows Server 2003 Update for Windows Server 2003 x64 Edition (KB932596)
 - Windows Server 2003 Security Update for Windows Server 2003 x64 Edition (KB926122)
 - Windows Server 2003 Security Update for Windows Media Player 6.4 (KB925398)
 - Windows Server 2003 Update for Windows Server 2003 x64 Edition (KB936357)
 - Windows Server 2003 Cumulative Security Update for Outlook Express for Windows Server 2003 x64 Edition (KB929123)
 - Windows Server 2003 Security Update for Windows Server 2003 x64 Edition (KB935839)
 - Windows Server 2003 Security Update for Windows Server 2003 x64 Edition (KB935840)
 - Windows Server 2003 Security Update for Windows Server 2003 x64 Edition (KB924667)
 - Windows Server 2003 Update for Windows Server 2003 x64 Edition (KB927891)
 - Windows Server 2003 Security Update for Windows Server 2003 x64 Edition (KB932168)
 - Windows Server 2003 Security Update for Windows Server 2003 x64 Edition (KB930178)
 - Windows Server 2003 Security Update for Windows Server 2003 x64 Edition (KB925902)
2. Upgrade from Microsoft Windows Server 2003 Enterprise x64 Edition to Microsoft Windows Server 2008 Enterprise x64 Edition.
- a. Insert the Microsoft Windows Server 2008 Enterprise x64 Edition DVD into the DVD-ROM drive.
 - b. At the Install Windows Server 2008 screen, click Install now.
 - c. In the Get Important Updates for Installation screen click on the Go online to get the latest updates for installation option.
 - d. At Type your product key for activation screen, enter your activation key, uncheck Automatically activate Windows when I'm online, and click Next.
 - e. Click on Windows Server 2008 Enterprise x64 (Full Installation), check the box next to I have selected the edition of Windows that I purchased, and click next
 - f. Check I accept the license terms checkbox, and click Next.
 - g. Click Upgrade for type of install.
 - h. At the Compatibility Report screen, click Next.
 - i. When the installation is finished, press Ctrl + Alt + Del to log in.

- j. Click the Administrator icon.
 - k. Enter a password, and press Enter.
3. To improve Java performance, enable large pages in memory on all servers. To enable this service, the administrator must first assign additional privileges to the user who will be running the application. (We assigned this privilege only to the administrator, because we used that account for our tests.) To enable large pages, select the following:
 - Control Panel→Administrative Tools→Local Security Policy→Local Policies→User Rights Assignment: Add Administrator
 - “Lock pages in memory,” add users and/or groups
4. Turn the screen saver off.
5. Customize desktop to put Computer on it.
6. Set screen to 1024x768, 32-bit color.
7. In the Manage your server dialog, select Don't display this page at logon.
8. Turn off Automatic updates.
9. Set the server to login automatically.
10. Install SPECjbb and the JVM on the server.
11. Empty Recycle Bin.
12. Create the image on server.

For all other blades, use Ghostcast to install the image. Use the system we just completed configuring to create a Ghost image on the server.

During initial testing, the HP server would not load large pages with the amount of Java heap size we were using due to memory constraints. To free more memory on the HP, we uninstalled HP System Management Homepage through the Server 2008 programs and features option on all HP blade systems. This allowed the HP blades to use the same heap size as the Dell and HP chassis. Removing this program on the HP did not hinder performance in any way.

After capturing images to servers

1. Each server needs a unique host name. Change the host name from the one you used for the base image to one of the form “<Brand>Server<#>”, where Brand is Dell, HP, or IBM and X is the blade number in the chassis (1–16 for Dell and HP and 1–14 for IBM) (e.g., IBMServer2).

We tested with the following settings for each enclosure:

Dell PowerEdge M600:

- Server Power Throttling Enabled: Checked
- Redundancy Policy: AC redundant
- Enable Dynamic Power Supply Engagement: Checked

HP BladeSystem c-Class:

- Enclosure Power Mode: AC redundant
- Enable Dynamic Power Savings Mode: Checked

IBM BladeCenter H Type 8852:

- Domain 1: Redundant without performance impact
- Domain 2: Redundant without performance impact
- Acoustic mode: Disabled

Power measurement procedure

To record each blade system's power consumption during testing, we used an Extech Instruments (www.extech.com) 380803 Power Analyzer/Datalogger. We connected the power cord from the system's power supply to the Power Analyzer's output load power outlet. We then plugged the power cord from the Power Analyzer's input voltage connection into a power outlet. We used this setup for each power supply in the chassis.

Because each of the three servers has six power supplies, we used six Extech Power Analyzers for testing. We connected all Extech Power Analyzers to one monitoring system to record the power draw of the systems.

We used the Power Analyzer's Data Acquisition Software (version 2.11) to capture all recordings. We installed the software on a separate PC, to which we connected all Power Analyzers via a separated RS-232 cable for each Extech. We captured power consumption at 1-second intervals.

To gauge the idle power usage, we recorded the power usage for 2 minutes while each server was running the operating system but otherwise idle.

To compute the total power, we took the wattage sum from each of the meters. We averaged the power usage during the period the server was running the benchmark. We call this time the power measurement interval. See Figures 2 (idle power consumption) and 3 (average peak power) for the results of these measurements.

SPECjbb2005 configuration

We used SPECjbb2005 version 1.07, dated March 15, 2006. We followed SPEC's run rules. (For more information about SPECjbb2005 and its run rules, see www.spec.org/jbb2005/docs/RunRules.html.) We installed SPECjbb2005 by copying the contents of the SPECjbb2005 CD to the directory C:\SPECjbb2005v1.07 on the server's hard disk.

SPECjbb2005 requires a Java Virtual Machine on the system under test. We used the BEA JRockit(R) (build P27.4.0-10-90053-1.6.0_02-20071009-1827-windows-x86_64, compiled mode) JVM for this testing and left the default installation settings.

After installation, as per the run rules, we edited the SPECjbb_config.props file in the root SPECjbb2005 directory to include disclosure information about the server and our license information. SPECjbb2005 uses this file when generating the results output for each run. We also modified the SPECjbb.props file to change the number of JVM instances to two. This change allows a server to run two JVM instances during testing.

We created a batch file, which we placed in the root SPECjbb2005 directory, to issue the Java run command to launch the benchmark. During testing, we used the command prompt window within Microsoft Windows Server 2008 x64 Edition to run this batch file, the text of which is as follows:

```
@echo off
set path="C:\jrockit-jdk1.6.0_02\bin";%path%

set JVM=2
:: Set JAVA_HOME to Java.exe path.
set JAVA_HOME="C:\jrockit-jdk1.6.0_02\bin"

:stage1
set PROPFIELD=SPECjbb.props
set JAVAOPTIONS= -Xms256m -Xmx256m
rem set JBBJARS=.\jbb.jar;.\check.jar
set JBBJARS=.\jbb.jar;.\jbb_no_precompile.jar;.\check.jar;.\reporter.jar

set CLASSPATH=%JBBJARS%;%CLASSPATH%

:stage2

echo Using CLASSPATH entries:
```

```

for %%c in ( %CLASSPATH% ) do echo %%c
@echo on
start /b C:\jrockit-jdk1.6.0_02\bin\java.exe %JAVAOPTIONS% spec.jbb.Controller -
propfile %PROFFILE%
@echo off
set I=0
set J=F
:LOOP
set /a I=%I + 1
echo.
echo Starting JVM Number %I% with Affinity to CPU %J%
echo.

@echo on
start /AFFINITY %J% /B C:\jrockit-jdk1.6.0_02\bin\java.exe -Xms1600m -Xns1300m -
Xmx1600m -XXaggressive -XXlargepages -XXcallprofiling -Xgc:genpar -
XXthroughputCompaction -XXlazyUnlocking -XXtlasize:min=4k,preferred=256k
spec.jbb.JBBmain -propfile %PROFFILE% -id %I% > multi.%I%
@echo off
set J=%J%0
IF %I% == %JVM% GOTO END
GOTO LOOP
:END

:egress

```

In the batch file we set the Java options that control the performance of the JVM as follows:

<i>Xms1600m</i>	This option sets the minimum heap size. We set the minimum and maximum heap sizes to be the same, so the heap size would stay a constant 1,600 MB.
<i>Xns1300m</i>	This option sets the nursery size to 1,300 MB.
<i>Xmx1600m</i>	This option sets the maximum heap size.
<i>XXaggressive</i>	This option essentially tells the JVM to perform at maximum speed.
<i>Xgc:genpar</i>	This option tells Java to use generational parallel garbage collection.
<i>XXthroughputCompaction</i>	This option adjusts the compaction ratio dynamically based on live data in the heap.
<i>XXlazyUnlocking</i>	This option determines when the JVM releases locks.
<i>XXtlasize:min=4k,preferred=256k</i>	This option sets the thread-local area size the JVM uses. We specified a minimum and preferred setting for testing.
<i>-XXlargepages</i>	This option tells the JVM to use large pages, if they are available, for the Java heap and other areas in the JVM.
<i>-XXcallprofiling</i>	This option enables the use of call profiling for code optimizations.

Test execution

To calculate the average power during peak performance, we needed all systems to be running at maximum performance at the same time. To achieve this, we needed all blade servers to start SPECjbb2005 at the same time, which we accomplished by using batch files to start SPECjbb2005 on all systems under test.

On each system under test, we created a batch file in the startup folder that would start as soon as the operating system loaded and then sleep, or sit idle, for 720 seconds. After 720 seconds, the batch file would search, once per second, for a run.txt file in the SPECjbb2005 directory. To begin the test, we used a batch file on a controller system that copied the run.txt file to all systems. Once that batch file had copied the run.txt file to the systems under test, the running batch files would start the SPECjbb2005 benchmark. By starting SPECjbb2005 this way, we ensured that all clients started within 1 second of each other.

Appendix A – Enclosure configuration information

Figure 17 provides detailed configuration information about the enclosures, which we present in alphabetical order.

Enclosure	Dell PowerEdge M600	HP BladeSystem c-Class	IBM BladeCenter H Type 8852
General dimension information			
Height (inches)	17.5	17.5	15.75
Width (inches)	19	19	19
Depth (inches)	30.5	31.0	28.0
U size in server rack	10	10	9
Number of blades	16	16	14
Power supplies			
Total number	6	6	2
Wattage of each	2,360	2,250	2,900
Cooling fans			
Total number	9	10	2 blowers
Dimensions (H x W) of each	3.5 x 3	3.5 x 3	4.5x11.5
Voltage	12 volts	12 volts	200-240 volts
Amps	7 amps	16.5 amps	5.5 amps

Figure 17: Detailed configuration information about the enclosures.

Appendix B – Blade system configuration information

Figure 18 provides detailed configuration information about the blade server systems, which we present in alphabetical order.

Servers	Dell PowerEdge M600	HP BladeSystem c-Class	IBM BladeCenter H Type 8852
General processor setup			
Number of processor packages	2	2	2
Number of cores per processor package	4	4	4
Number of hardware threads per core	1	1	1
System Power Management Policy	Balanced	Balanced	Balanced
CPU			
Vendor	Intel	Intel	Intel
Name	Quad-Core Intel Xeon processor E5345	Quad-Core Intel Xeon processor E5345	Quad-Core Intel Xeon processor E5345
Stepping	B	7	7
Socket type	LGA 771	LGA 771	LGA 771
Core frequency (GHz)	2.33 GHz	2.33 GHz	2.33 GHz
Front-side bus frequency (MHz)	1,333 MHz	1,333 MHz	1,333 MHz
L1 cache	32 KB + 32 KB (per core)	32 KB + 32 KB (per core)	32 KB + 32 KB (per core)
L2 cache	2 x 4 MB (each 4 MB shared by two cores)	2 x 4 MB (each 4 MB shared by two cores)	2 x 4 MB (each 4 MB shared by two cores)
Platform			
Vendor and model number	Dell PowerEdge M600	HP ProLiant BL460c	BladeCenter HS21
Motherboard model number	Dell 0MY736	HP 435458-B21	IBM 8853C2U
Motherboard chipset	Intel 5000P	Intel 5000P	Intel 5000P
Motherboard revision number	X31	91	B1
BIOS name and version	Dell 1.1.0	HP I15 1/24/2008	IBM 1.10 1/31/2008
BIOS settings	Disabled Hardware Prefetcher and Adjacent Cache Line Prefetcher	Disabled Hardware Prefetcher and Adjacent Cache Line Prefetcher	Disabled Hardware Prefetcher and Adjacent Cache Line Prefetcher
Chipset INF driver	Microsoft 6.0.6001.18000	Microsoft 6.0.6001.18000	Microsoft 6.0.6001.18000
Memory module(s)			
Vendor and model number	Samsung M395T2953EZ4-CE65	Micron MT18HTF12872FDY	Hynix HYMP512F72CP8D2-Y5
Type	PC2-5300	PC2-5300	PC2-5300
Speed (MHz)	667 MHz	667 MHz	667 MHz
Speed in the system currently running @ (MHz)	667 MHz	667 MHz	667 MHz
Timing/Latency (tCL-tRCD-tRP-tRASmin)	5-5-5-15	5-5-5-15	5-5-5-15
Size	4 GB (4 x 1 GB)	4 GB (4 x 1 GB)	4 GB (4 x 1 GB)
Number of RAM modules	4	4	4
Chip organization	Dual side	Dual side	Dual side

Servers	Dell PowerEdge M600	HP BladeSystem c-Class	IBM BladeCenter H Type 8852
Hard disk			
Vendor and model number	Fujitsu may2073rc	Seagate St973402SS	IBM 26K5777
Number of disks in system	2	2	2
Size	73 GB	72 GB	73.4 GB
Buffer size	16 MB	16 MB	8 MB
RPM	10,000	10,000	10,000
Type	SAS	SAS	SAS
Controller	Dell SAS 6/iR Integrated Blades Controller	Smart Array E200I controller	LSI Adapter, SAS 3000 series
Controller driver	Microsoft 1.25.6.22	HP 6.6.0.64	LSI 1.21.28.0
Operating system			
Name	Microsoft Windows Server 2008, Enterprise x64 Edition Service Pack 1	Microsoft Windows Server 2008, Enterprise x64 Edition Service Pack 1	Microsoft Windows Server 2008, Enterprise x64 Edition Service Pack 1
Build number	6001	6001	6001
Microsoft Windows update date	None	None	None
File system	NTFS	NTFS	NTFS
Kernel	ACPI x64-based PC	ACPI x64-based PC	ACPI x64-based PC
Language	English	English	English
Microsoft DirectX version	10	10	10
Graphics			
Vendor and model number	ATI ES1000	ATI ES1000	ATI ES1000
Chipset	ATI ES1000	ATI ES1000	ATI ES1000
BIOS version	BK-ATI VER008.005.031.000	BK-ATI VER008.005.013.000	BK-ATI VER008.005.031.000
Type	Integrated	Integrated	Integrated
Memory size	32MB	32MB	16MB
Resolution	1,024 x 768	1,024 x 768	1,024 x 768
Driver	Microsoft 6.0.6001.18000	Microsoft 6.0.6001.18000	Microsoft 6.0.6001.18000
Network card/subsystem			
Vendor and model number	Broadcom BCM5708S NetXtreme II GigE	Broadcom BCM5708S NetXtreme II GigE	Broadcom BCM5708S NetXtreme II GigE
Type	Integrated	Integrated	Integrated
Driver	Microsoft 3.7.19.0	Microsoft 3.7.19.0	Microsoft 3.7.19.0
Optical drive			
Vendor and model number	None installed	None installed	None installed
USB ports			
Number	2	2 (with adapter attached)	2
Type	USB 2.0	USB 2.0	USB 2.0

Figure 18: Detailed configuration information about the blade server systems.



Principled Technologies, Inc.
1007 Slater Road, Suite 250
Durham, NC 27703
www.principledtechnologies.com
info@principledtechnologies.com

Principled Technologies is a registered trademark of Principled Technologies, Inc.
All other product names are the trademarks of their respective owners

Disclaimer of Warranties; Limitation of Liability:

PRINCIPLED TECHNOLOGIES, INC. HAS MADE REASONABLE EFFORTS TO ENSURE THE ACCURACY AND VALIDITY OF ITS TESTING, HOWEVER, PRINCIPLED TECHNOLOGIES, INC. SPECIFICALLY DISCLAIMS ANY WARRANTY, EXPRESSED OR IMPLIED, RELATING TO THE TEST RESULTS AND ANALYSIS, THEIR ACCURACY, COMPLETENESS OR QUALITY, INCLUDING ANY IMPLIED WARRANTY OF FITNESS FOR ANY PARTICULAR PURPOSE. ALL PERSONS OR ENTITIES RELYING ON THE RESULTS OF ANY TESTING DO SO AT THEIR OWN RISK, AND AGREE THAT PRINCIPLED TECHNOLOGIES, INC., ITS EMPLOYEES AND ITS SUBCONTRACTORS SHALL HAVE NO LIABILITY WHATSOEVER FROM ANY CLAIM OF LOSS OR DAMAGE ON ACCOUNT OF ANY ALLEGED ERROR OR DEFECT IN ANY TESTING PROCEDURE OR RESULT.

IN NO EVENT SHALL PRINCIPLED TECHNOLOGIES, INC. BE LIABLE FOR INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH ITS TESTING, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL PRINCIPLED TECHNOLOGIES, INC.'S LIABILITY, INCLUDING FOR DIRECT DAMAGES, EXCEED THE AMOUNTS PAID IN CONNECTION WITH PRINCIPLED TECHNOLOGIES, INC.'S TESTING. CUSTOMER'S SOLE AND EXCLUSIVE REMEDIES ARE AS SET FORTH HEREIN.