



DELL DELIVERS PERFORMANCE AND ENERGY EFFICIENCY OVER HP AND IBM BLADE SERVERS

Dell™ PowerEdge™ blade servers are designed for high performance and efficient energy use. In Principled Technologies® blade server tests, the Dell PowerEdge M905 provided both higher peak performance and lower power consumption than the HP ProLiant BL685c G5 or IBM® BladeCenter LS42—delivering significantly higher performance per watt to help optimize data center efficiency.

When evaluating a server, performance and power consumption are two key factors that organizations should take into account. Perhaps most important, however, is the efficiency with which the server can translate energy into useful work performed in the data center—a factor captured by the key metric of *performance per watt*. Maximizing performance per watt can lead to significant improvements in IT productivity without increasing power consumption.¹

In January 2009, Principled Technologies performed tests commissioned by Dell to compare the performance and power consumption of three AMD Opteron™ processor-based blade servers: the Dell PowerEdge M905, HP ProLiant BL685c G5, and IBM BladeCenter LS42.² As the results show, the Dell server not only achieved the highest peak performance of the three tested servers, but also did so while using the least amount of power—providing significantly higher performance per watt than the HP and IBM servers to help optimize data center efficiency.

TEST ENVIRONMENT AND METHODOLOGY

The Dell PowerEdge M905 was housed in a Dell PowerEdge M1000e enclosure, the HP ProLiant BL685c G5 was housed in an HP BladeSystem c-Class enclosure, and the IBM BladeCenter LS42 was housed in an IBM BladeCenter H (8852) enclosure. Each blade server was configured with four quad-core AMD Opteron 8356 processors at 2.3 GHz and 32 GB of RAM, and ran the Sun Solaris 10 OS.

The tests were based on version 1.07 of the Standard Performance Evaluation Corporation (SPEC) SPECjbb2005 benchmark. SPECjbb2005 utilizes multiple special data groups and multiple threads as it runs, with each data unit referred to as a *warehouse*—a roughly 25 MB 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 until it has saturated the server's processor capacity. As the number of warehouses increases, so does the number of threads. The benchmark's results portray

Related Categories:

- AMD Opteron processors
- Benchmarks
- Blade servers
- Dell PowerEdge blade servers
- Energy efficiency
- Green IT
- Performance
- Performance characterization
- Power and cooling

Visit DELL.COM/PowerSolutions for the complete category index.

¹ For more information on how best practices for energy efficiency can help increase IT productivity, see "Compute More, Consume Less: Smart Policies Unleash Data Center Productivity," in *Dell Power Solutions*, March 2009, DELL.COM/Downloads/Global/Power/ps1q09-20090176-Esser.pdf.

² See "SPECjbb2005 Performance and Power Consumption on Sun Solaris 10 Multi-processor AMD-Based Blade Servers from Dell, HP, and IBM," by Principled Technologies, January 2009, DELL.COM/Downloads/Global/Solutions/DellM905SPECjbb0109.pdf.

	IBM BladeCenter LS42	HP ProLiant BL685c G5	Dell PowerEdge M905
JVM 1	90,957 bops	94,186 bops	101,453 bops
JVM 2	92,262 bops	94,081 bops	102,010 bops
JVM 3	91,671 bops	92,834 bops	92,936 bops
JVM 4	93,223 bops	91,250 bops	91,974 bops

Figure 1. Median SPECjbb2005 results for each Java Virtual Machine on the tested blade servers

the server's throughput in business operations per second (bops).³

SPECjbb2005 requires a Java Virtual Machine (JVM) on the test system. The test team used the Sun Java Platform, Standard Edition (SE) 6 Runtime Environment (build 1.6.0_06-p-b04) and Java HotSpot Server VM (build 13.0-b04, mixed mode) with the default installation settings. After installing these components, editing the SPECjbb_config.props file to include disclosure and license information, and editing the SPECjbb.props file to allow the servers to each run four JVM instances, the test team created a shell script to issue the Java run command to launch the benchmark.

The team used the Solaris `psrinfo -pv` command to retrieve system processor information before beginning the testing, using the processor assignments to set the processor affinity appropriately on each system. The team then set the `psrset` command to assign groupings of cores to JVMs. The core numbers differed among the systems, so the `psrset -c` commands differed slightly among the three blade servers; the test team configured these settings to provide maximum performance on all servers.

The shell script also included several Java options that controlled JVM performance:

- `-Xms3350m`: Set the minimum heap size to 3,350 MB
- `-Xmx3350m`: Set the maximum heap size to 3,350 MB (because the minimum and maximum heap sizes were the

same in these tests, the heap size stayed at a constant 3,350 MB)

- `-Xmn2600m`: Set the JVM nursery size to 2,600 MB
- `-XX:+UseParallelOldGC`: Set Java to use parallel garbage collection
- `-XX:+AggressiveOpts`: Enabled performance compiler optimizations
- `-XX:AllocatePrefetchStyle=2`: Set code-style prefetch instructions
- `-XX:ParallelGCThreads=2`: Set the number of GC threads to two

Server power consumption during each test was recorded using an Extech Instruments 380803 power analyzer and data logger. Before starting the benchmark, the test team logged into the servers and allowed them to sit idle for 8 minutes. To gauge idle power consumption, the

test team then recorded server power consumption for 2 minutes while the server was running only the Solaris 10 OS. This process meant that each system was idle for 10 minutes before the tests began. After the tests began, the team recorded server power consumption at one-second intervals. Average power consumption was calculated by averaging power consumption during the time the server was producing its peak performance results.

TEST RESULTS: PERFORMANCE AND EFFICIENCY

During the tests, each server ran four JVMs simultaneously. To compute the overall score for each server, SPECjbb2005 sums the scores of its JVMs; the scores of each JVM are calculated by taking the average of the results when the server is running at peak performance. Figure 1 shows the median SPECjbb2005 results for each JVM across three runs of the SPECjbb2005 benchmark, while Figure 2 illustrates the total median SPECjbb2005 results for each blade server. A higher SPECjbb2005 score indicates that the server handled more Java requests and thus delivered greater throughput.

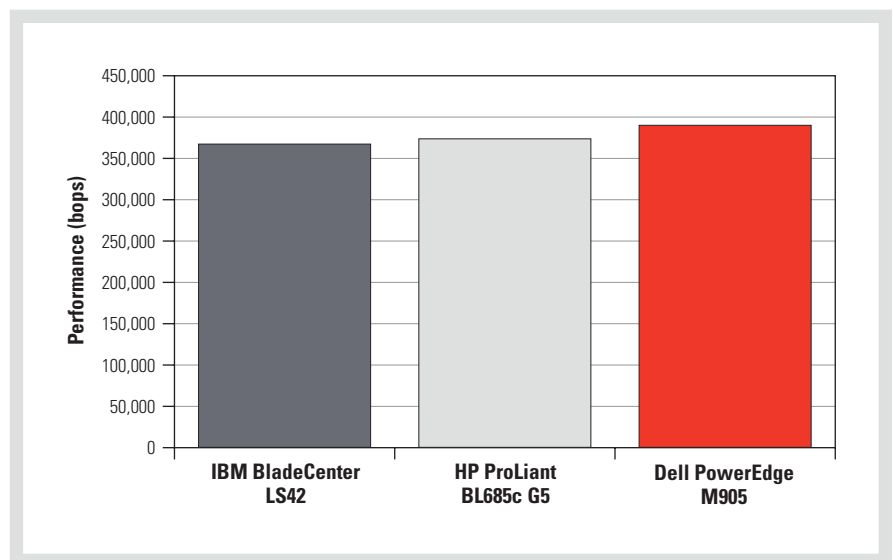


Figure 2. Total median SPECjbb2005 results for the tested blade servers

³ For more information on SPECjbb2005, visit www.spec.org/jbb2005.

As these figures show, the Dell PowerEdge M905 blade server produced the highest peak performance of the three systems tested, handling a total median of 388,373 bops. This represented 4.3 percent higher performance than the HP ProLiant BL685c G5 (which achieved a total median performance of 372,351 bops) and 5.5 percent higher performance than the IBM BladeCenter LS42 (which achieved a total median performance of 368,113 bops).

Figure 3 shows the average power consumption for each server when idle and during peak performance. Figure 4 shows the performance-per-watt results, which have been normalized to show relative increase over the server with the lowest result, the IBM BladeCenter LS42. Performance per watt was calculated by dividing the total median peak SPECjbb2005 performance by the average power consumption in watts when

“As the results show, the Dell server not only achieved the highest peak performance of the three tested servers, but also did so while using the least amount of power—providing significantly higher performance per watt than the HP and IBM servers to help optimize data center efficiency.”

the server was delivering peak performance during the median run. As these figures show, the Dell PowerEdge M905 blade server provided not only the highest overall performance of the three tested servers, but also the lowest power consumption—delivering 37.6 percent higher performance per watt than the HP

ProLiant BL685c G5 and 46.7 percent higher performance per watt than the IBM BladeCenter LS42.

POWERFUL, ENERGY-EFFICIENT BLADE SERVERS

Optimizing energy use can be critical to controlling operational costs in enterprise data centers. As these Principled Technologies tests demonstrate, Dell PowerEdge blade servers are designed to offer high performance while minimizing power consumption—providing the foundation of an effective and comprehensive strategy for data center efficiency. [u](#)

	IBM BladeCenter LS42	HP ProLiant BL685c G5	Dell PowerEdge M905
Idle	856.2 W	707.1 W	484.1 W
During median benchmark run	995.9 W	945.0 W	716.2 W

Figure 3. Average power consumption for the tested blade servers




Figure 4. Normalized performance per watt for the tested blade servers

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Full Principled Technologies test report:
DELL.COM/Downloads/Global/Solutions/DellM905SPECjbb0109.pdf

Dell PowerEdge blade servers:
DELL.COM/Blades

Dell Solaris solutions:
DELL.COM/Solaris