Performance and performance per watt of the Accelrys DMol^3 application on Intel- and AMD-processor-based servers

KEY FINDINGS

- The Dual-Core Intel Xeon processor 5160-based server outperformed the Dual-Core AMD Opteron 285 in a dual-processor configuration by 47 percent and 45 percent respectively on the AFY and TiN DMol^3 workloads.
- The Dual-Core Intel Xeon processor 5160-based server achieved between 81 percent and 108 percent better performance / watt than the Dual-Core AMD Opteron 285-based server while running the DMol^3 AFY and TiN workloads.
- In both the one- and two-processor configurations, the Intel Xeon processor 5160-based server outperformed the AMD Opteron 285-based server by significant margins on the Accelrys DMol^3 workloads.

Executive summary

Intel Corporation (Intel) commissioned Principled Technologies (PT) to measure the performance and power consumption of the DMol^3 workload on dual-processor servers using the following two processors:

- Dual-Core Intel Xeon processor 5160
- Dual-Core AMD Opteron 285

Accelrys is a company that provides software and service solutions to pharmaceutical, biotechnology, and industrial chemical research organizations. The DMol^3 software is a Density Functional Theory quantum mechanical code for use in chemistry and materials science. The AFY and TiN workloads perform differing complex calculations to predict the properties of chemicals and materials used by industrial processes. (See About DMol^3 and the two workloads for more details.)

Figure 1 illustrates the relative peak performance of each server in a two-processor configuration while running the DMol^3 AFY workload. The Dual-Core Intel Xeon processor 5160-based server finished the AFY workload about 47 percent faster than the Dual-Core AMD Opteron 285-based server. This difference translates into significant time savings for users. Figure 1 also illustrates the performance / watt for the servers, which we calculated as performance / watt = (3600 / (the benchmark's duration in seconds)) / (average power consumption in watts during the time period in which the benchmark was delivering peak performance). This formula converts the elapsed time the benchmark took to complete into a runs (or jobs) per hour metric, which we then use to compute the performance / watt. As Figure 1 shows, the Dual-Core Intel Xeon processor 5160-based server delivered about 108 percent more performance / watt than the Dual-Core AMD Opteron 285-based server.

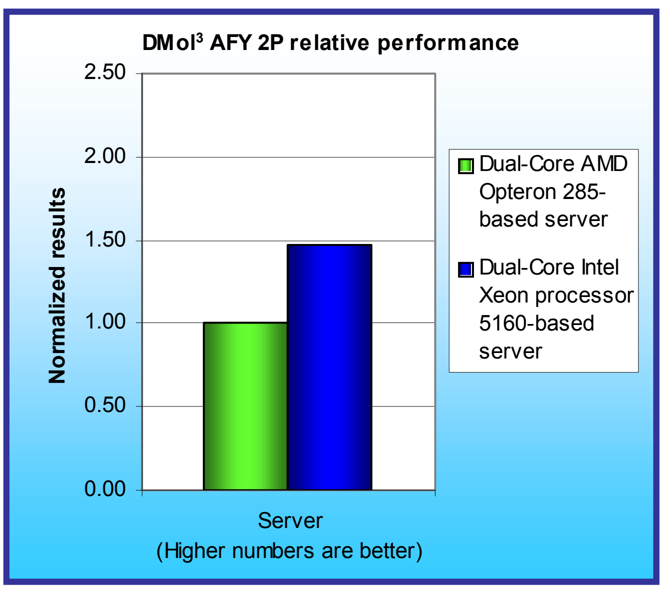

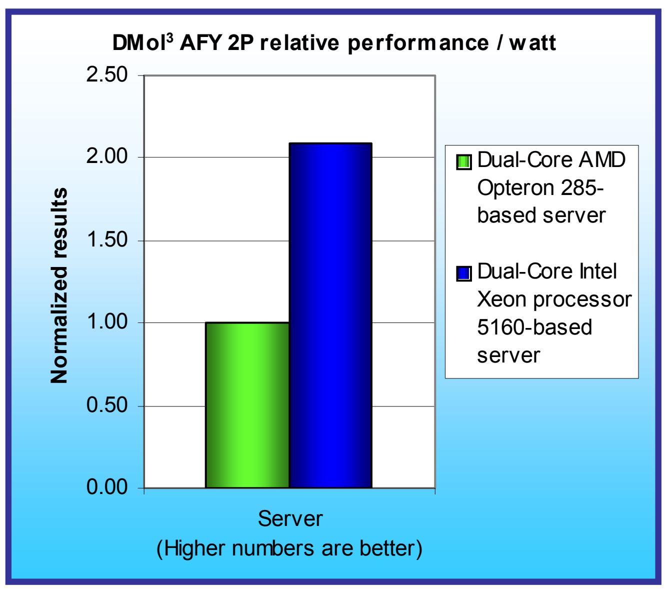

Figure 1: Normalized peak performance and performance per watt of the servers with two processors running the AFY DMol^3 workload. Higher numbers are better.

Figure 2 illustrates the relative peak performance and performance / watt of each server in a two-processor configuration while running the DMol^3 TiN workload. The Dual-Core Intel Xeon processor 5160-based server finished the TiN workload about 45 percent faster than the Dual-Core AMD Opteron 285-based server and delivered about 103 percent more performance / watt.

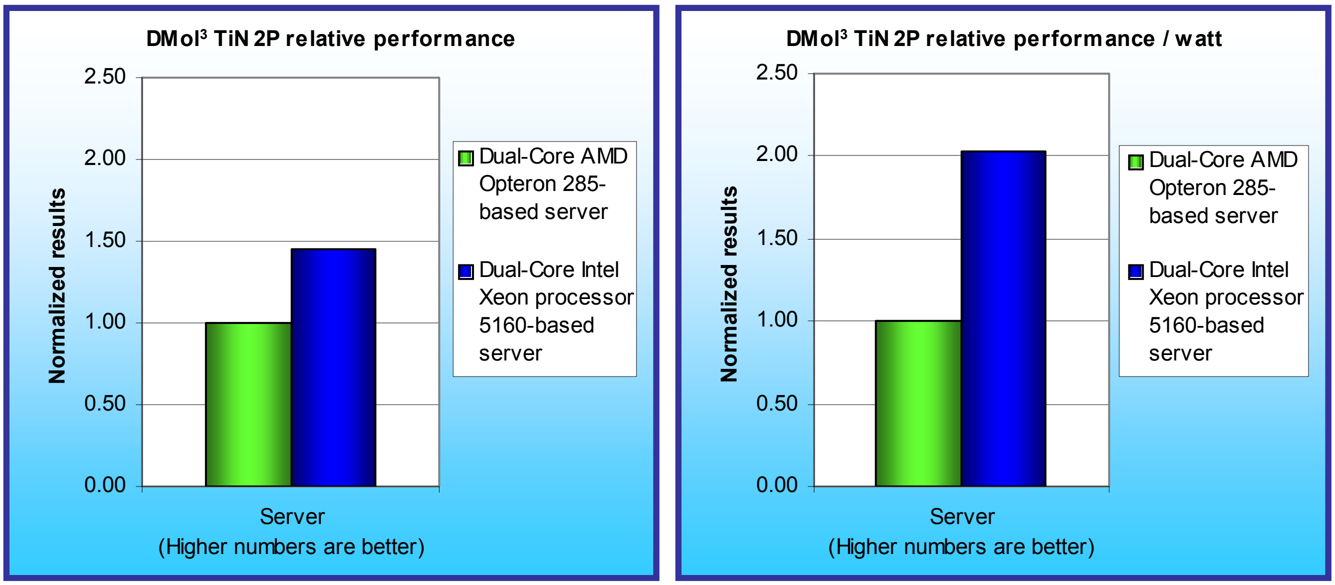

Figure 2: Normalized peak performance and performance per watt of the servers with two processors running the TiN DMol^3 workload. Higher numbers are better.

Figure 3 shows the relative peak performance of each server in a one-processor configuration while running the DMol^3 AFY workload. The Dual-Core Intel Xeon processor 5160-based server finished the AFY workload about 42 percent faster than the Dual-Core AMD Opteron 285-based server. The figure also shows that the Dual-Core Intel Xeon processor 5160-based server delivered about 81 percent more performance / watt.

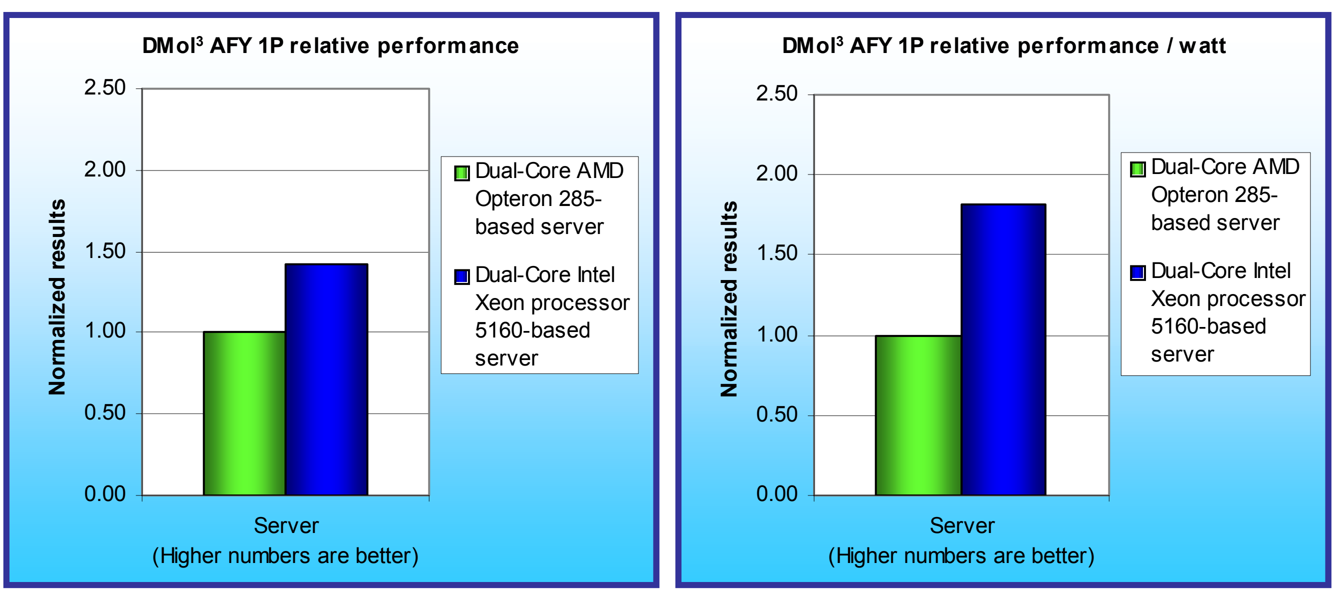

Figure 3: Normalized peak performance and performance per watt of the servers with one processor running the AFY DMol^3 workload. Higher numbers are better.

Figure 4 illustrates the relative peak performance and performance / watt of each server in a one-processor configuration while running the DMol[3] TiN workload. The Dual-Core Intel Xeon processor 5160-based server finished about 47 percent faster than the Dual-Core AMD Opteron 285-based server and delivered about 99 percent more performance / watt.

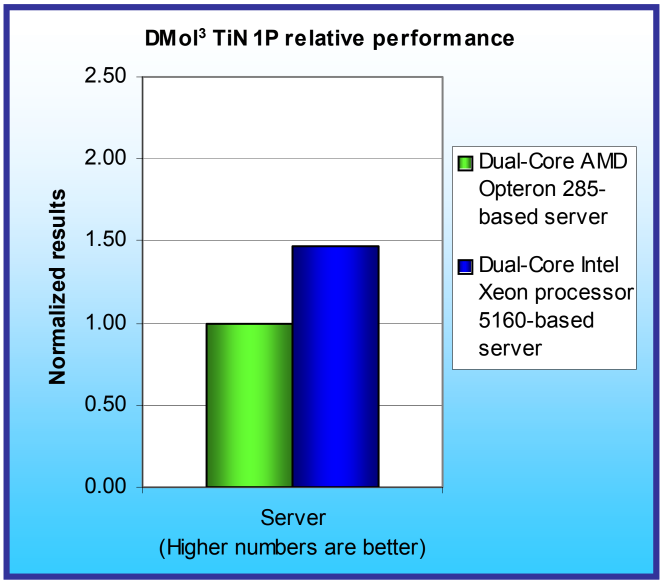

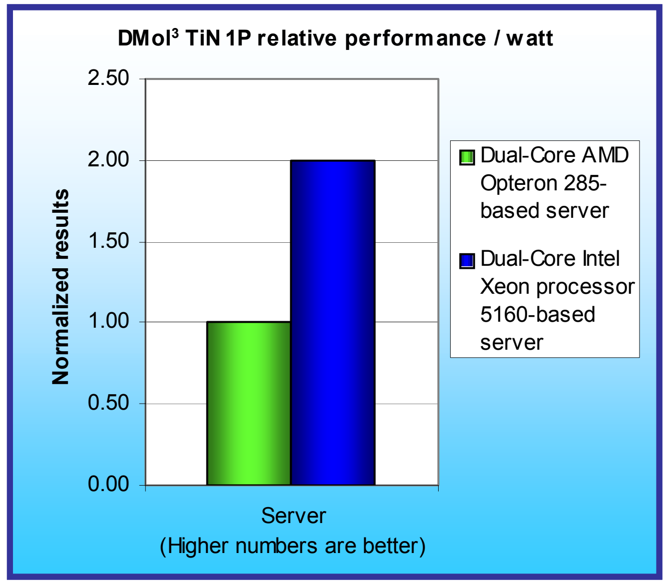

Figure 4: Normalized peak performance and performance per watt of the servers with one processor running the TiN DMol[3] workload. Higher numbers are better.

Figure 5 illustrates the power usage of the two servers in two-processor configurations as they were executing the DMol[3] AFY workload. The red lines indicate the power measurement interval, the time during which the server was delivering peak performance on the DMol[3] AFY workload and during which we captured power measurements. Lower power consumption is better. The Dual-Core Intel Xeon Processor 5160-based server achieved its peak performance while drawing less power—9 percent less—than the Dual-Core AMD Opteron 285-based server. (The drop in power consumption back to idle state for both the Dual-Core Intel Xeon Processor 5160-based server and the Dual-Core AMD Opteron 285-based server occurred when each of those servers finished the workload.)

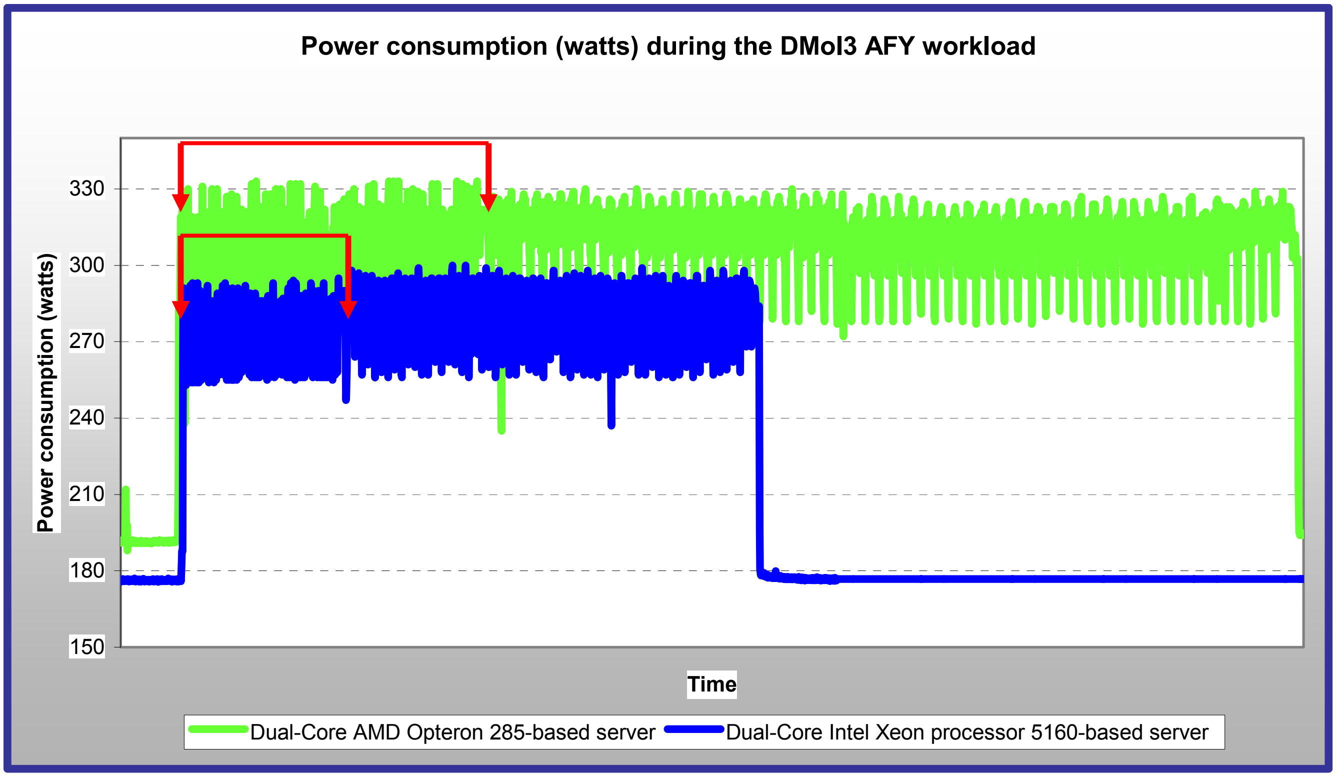

Figure 5: Power consumption (in watts) of each of the servers in dual-processor configurations as they executed the DMol[3] AFY workload. Lower power consumption is better.

Figure 6 illustrates the power usage of the two servers in dual-processor configurations as they were executing the DMol[3] TiN workload. The red lines again indicate the power measurement interval, the time during which the server was delivering peak performance on the DMol[3] TiN workload and during which we captured power measurements. Lower power consumption is better. This graph shows no idle power measurement for the TiN workload because that workload executed immediately after the AFY workload. The graph also does not include the AFY power consumption. The Dual-Core Intel Xeon Processor 5160-based server achieved its peak performance while drawing less power—10 percent less—than the Dual-Core AMD Opteron 285-based server. (The drop in power consumption back to idle state for both the Dual-Core Intel Xeon Processor 5160-based server and the Dual-Core AMD Opteron 285-based server occurred when each of those servers finished the workload.)

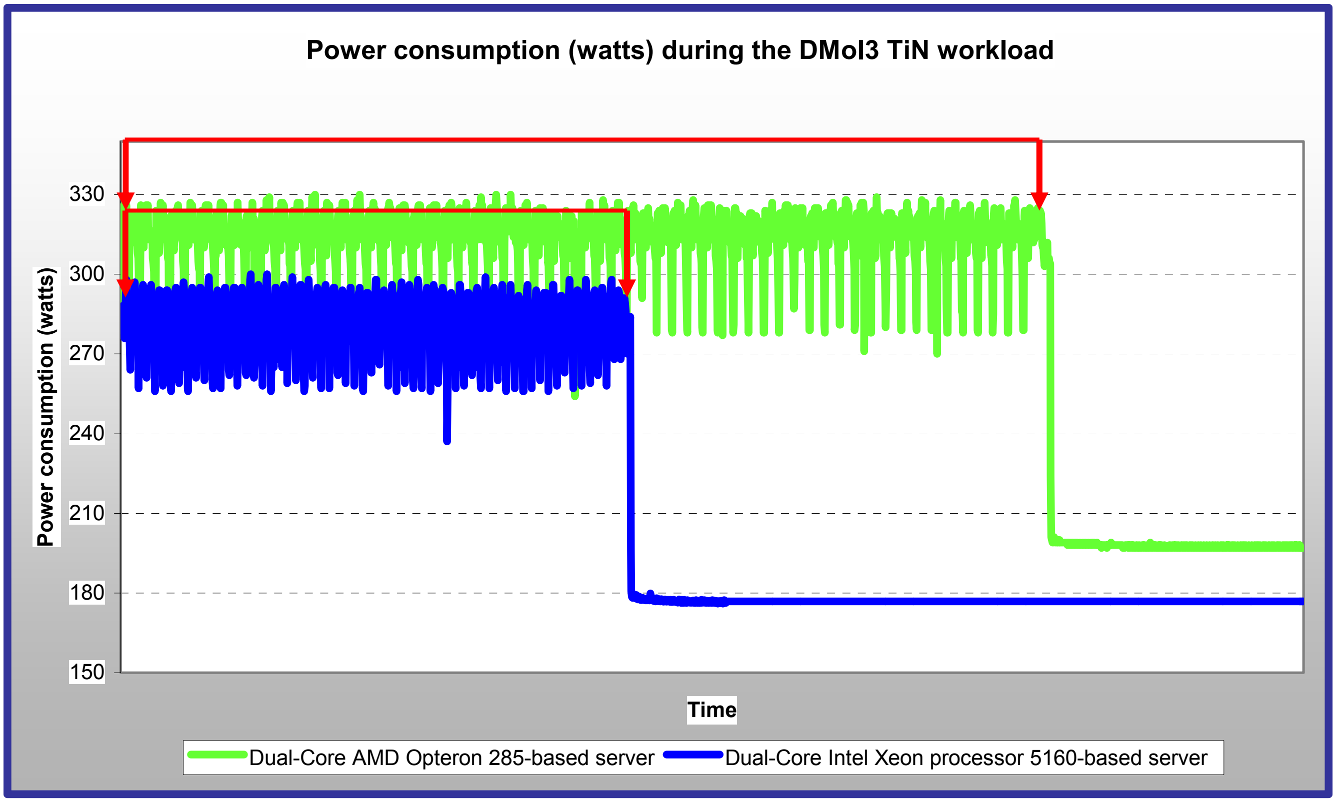

Figure 6: Power consumption (in watts) of each of the servers in dual-processor configurations as they executed the DMol[3] TiN workload. Lower power consumption is better.

About DMol[3] and the two workloads

Intel provided Principled Technologies (PT) with the scripts, compilers, and workloads to build and run the Accelrys DMol[3] workloads. Per Accelrys, "DMol[3] is a unique Density Functional Theory (DFT) quantum mechanical code that allows users to study problems in chemistry and materials. DMol[3] predicts the properties of materials at the atomic level using the principles of quantum mechanics. The methods employed yield accurate results that can be used to design new catalysts, improve the efficiency of chemical plants, or to discover stronger, lighter structural materials, to list just a few examples. The approach used in these simulations is quite computationally intensive; users require high-performance computers in order to perform realistic simulations in a reasonable timeframe." (http://www.accelrys.com/products/mstudio/modeling/quantumandcatalysis/dmol3.html)

AFY refers to a type of zeolite, minerals that companies can use for catalysis and separations. Companies can manipulate zeolites to perform specific industrial processes by changing their shape and composition. DMol[3] can predict the effects that these changes would have on the material's behavior.

TiN refers to Titanium Nitride, an extremely hard, ceramic material, that the semiconductor industry uses as a barrier material in thin films. Hydrogen plays a key role when vendors grow TiN thin films by chemical vapor deposition. DMol^3 can predict the effect of surface hydrogen atoms, which can act as contaminants.

Accelrys provided the two DMol^3 workloads, AFY and TiN, to Intel, and Intel supplied those workloads to PT.

The DMol^3 program performs complex calculations in its AFY and TiN workloads and reports the completion time for each. We used this completion time as the raw result for each workload.

Test results

Table 1 shows the median completion times of the two servers in both one- and two-processor configurations while executing the DMol[3] AFY workload. Per agreement with Accelrys, we present only normalized results. In the two-processor configuration, the Dual-Core Intel Xeon processor 5160-based server outperformed the Dual-Core AMD Opteron 285-based server by 47.2 percent. In the one-processor configurations, the Dual-Core Intel Xeon processor 5160-based server showed a 42.3 percent performance increase over the Dual-Core AMD Opteron 285-based server.

Server / # of processors	1P	2P
Dual-Core Intel Xeon processor 5160-based server	0.58	0.31
Dual-Core AMD Opteron 285-based server	1.00	0.60

Table 1: Median relative completion times of the servers with varying processor numbers running the DMol[3] AFY workload. We calculated each result by dividing the completion time by the completion time of the single-processor Opteron 285-based server. That score is 1.0. Lower numbers indicate faster times.

Table 2 shows the average power usage of the test servers with one processor and two processors while idle and during the median peak runs of the AFY workload. With one processor, the Dual-Core Intel Xeon processor 5160-based server consumed 207.1 watts, 4.4 percent less power than the Dual-Core AMD Opteron 285-based server, which consumed 216.6 watts. With two processors, the Dual-Core Intel Xeon processor 5160-based server was 9.1 percent more power efficient, consuming 274.9 watts, than the Dual-Core AMD Opteron 285-based server, which consumed 302.9 watts.

Server / # of processors	1P		2P	
	Idle power (watts)	Average power (watts)	Idle power (watts)	Average power (watts)
Dual-Core Intel Xeon processor 5160-based server	156.3	207.1	191.8	302.4
Dual-Core AMD Opteron 285-based server	160.2	216.6	176.3	274.9

Table 2: Average power usage (in watts) of the test servers with varying processor numbers while idle and during the median peak runs of the DMol[3] AFY workload. Lower numbers are better.

Table 3 shows the median completion times of the two servers in both one- and two-processor configurations while executing the DMol[3] TiN workload. In the two-processor configuration, the Dual-Core Intel Xeon processor 5160-based server outperformed the Dual-Core AMD Opteron 285-based server by 45.5 percent. In the one-processor configurations, the Dual-Core Intel Xeon processor 5160-based server showed a performance win for the Dual-Core Intel Xeon processor 5160-based server of 47.2 percent.

Server / # of processors	1P	2P
Dual-Core Intel Xeon processor 5160-based server	0.53	0.31
Dual-Core AMD Opteron 285-based server	1.00	0.56

Table 3: Median relative completion times of the servers with varying processor numbers running the DMol[3] TiN workload. We calculated each result by dividing the completion time by the completion time of the single-processor Opteron 285-based server. That score is 1.0. Lower numbers indicate faster times.

Table 4 shows the average power usage of the test servers with one processor and two processors while idle and during the median peak runs of the TiN workload. With one processor, the Dual-Core Intel Xeon processor 5160-based server consumed 209.2 watts and thus was 5.1 percent more power efficient than the Dual-Core AMD Opteron 285-based server, which consumed 220.5 watts. With two processors, the Dual-Core Intel Xeon processor 5160-based server consumed 282.6 watts, a 9.9 percent lower power usage than the Dual-Core AMD Opteron 285-based server, which consumed 313.6 watts.

Server / # of processors	1P		2P	
	Idle power (watts)	Average power (watts)	Idle power (watts)	Average power (watts)
Dual-Core Intel Xeon processor 5160-based server	160.2	220.5	191.4	313.6
Dual-Core AMD Opteron 285-based server	155.6	209.2	176.3	282.6

Table 4: Average power usage (in watts) of the test servers with varying processors numbers while idle and during the median peak runs of the DMol^3 TiN workload. Lower numbers are better.

Test methodology

Intel configured and provided the test servers. Figure 7 summarizes some key aspects of their configurations; Appendix A provides detailed configuration information.

Server	Dual-Core AMD Opteron 285-based server	Dual-Core Intel Xeon processor 5160-based server
Processor frequency (GHz)	2.6 GHz	3.0 GHz
Single/Dual-Core processors	Dual	Dual
Motherboard	HP Proliant DL385 G1 399684-001	Intel S5000PSL
Chipset	AMD 8111/8131 chipset	S5000P chipset
RAM (4GB in each)	4 x 1GB PC-3200	4 x 1GB PC2-5300
Hard Drive	1x Seagate Cheetah ST336754LC 15K RPM 37 GB drive attached through Adaptec 39320A-R PCI-X 133MHz SCSI controller	1x Seagate Cheetah ST336754LC 15K RPM 37 GB drive attached through Adaptec 39320A-R PCI-X 133MHz SCSI controller

Figure 7: Summary of some key aspects of the server configurations.

The difference in RAM types reflects the capabilities of the two motherboards: the Intel S5000PSL motherboard offered two independent front-side busses at a speed of 1333 MHz and contained Fully-Buffered DIMM (FBDIMM) modules that used commodity DDR2 PC2-4400 533MHz memory components. The Dual-Core AMD Opteron 285 motherboard supported 184-pin DDR memory. The highest memory speed available for the Dual-Core AMD Opteron 285-based server was DDR PC3200 400MHz RAM.

We used the default BIOS settings on each server.

We began our testing by installing a fresh copy of SUSE Linux Enterprise Server 9 (64 bit) on each server. We followed this process for each installation:

1. Assign the computer name Server.
2. Enter a password for the root account.
3. Select Eastern Time Zone.
4. Use typical settings for the Network installation.
5. Use the default name for the workgroup.

Installing and configuring the DMol^3 workloads

Intel also provided the following test software, which we used to build, install, and configure the test workloads:

- Fortran compiler version 9.1.033
- Intel C++ compiler version 9.1.039
- Intel MKL compiler version 8.1
- Intel MPI compiler version 2.0.1
- the DMol^3 benchmark, which contains two workloads: AFY and TiN
- scripts for executing the workloads

We explain below how we installed each of these tools and the benchmark.

Installing the Intel 32-bit compilers

1. Copy the Fortran compiler (l_fc_c_9[1].1.033.tar.gz) and the C compiler (l_cc_c_9[1].1.039.tar.gz) to the root directory: cp l_fc_c_9[1].1.033.tar.gz /.
2. Untar the Fortran and C compilers: tar –xzvf l_fc_c_9[1].1.033.tar.gz.
3. Install the Fortran compiler by changing directory to l_fc_c_9[1].1.033 and entering "./install.sh".
4. Install the C compiler by changing directory to l_cc_c_9[1].1.039 and entering “./install.sh”.
5. Source the iccvars and the ifortvars scripts by typing the following:
   source /opt/intel/cc/9.1.039/bin/iccvars.sh

source /opt/intel/fc/9.1.033/bin/ifortvars.sh

Installing the Intel MPI Library

1. Copy the MPI (l_mpi_p_2.0.1.012.tgz) to the root directory: cp <filename>.
2. Untar the library: tar –xzvf l_mpi_p_2.0.1.012.tgz.
3. Install it: ./install.sh.

Installing the Intel Math Kernel Library (MKL)

1. Copy the MPI (l_mkl_p_8.1.014.tgz) to the root directory: cp <filename>.
2. Untar the library: tar –xzvf l_mkl_p_8.1.014.tgz.
3. Install it: ./install.sh.

Installing the DMol^3 benchmark

1. Copy the file (dmol3bm.tar) to the root directory and untar. It will extract into /DMol3.
2. Change directory to /DMol3/src and edit the "Makefile" file as follows:
   a. Find the macro "MPIDIR" and ensure the path for the Intel MPI installation is correct. In our case it was "/opt/intel/mpi/2.0.1/".
   b. Find the macro "MKL_PATH" and ensure the path points to the 32 bit MKL libraries. In our case it was "/opt/intel/mkl/8.1/lib/32".
3. Type "make" to build the benchmark installation.
4. If the build is successful, the file "dmol3_mpi.exe" will appear in /DMol3/ia32.

Running the DMol^3 workloads

A single script runs both workloads. We executed that script as follows:

1. Source the iccvars, ifortvars, and the mpivars scripts by typing the following:
   source /opt/intel/cc/9.1.039/bin/iccvars.sh
   source /opt/intel/fc/9.1.033/bin/ifortvars.sh
   source /opt/intel/mpi/2.0.1/bin/mpivars.sh
2. Run the mpdboot command: cd to /root and type mpdboot -v -d -n 1.
   a. If you get a message stating that the system cannot open the mpd.hosts file, type "touch mpd.hosts.
   b. You will not get a command prompt after this, so open anther console and perform step 1 again, then proceed to step 3.
3. Get CPU utilization information by typing the following and saving the output to a results file: "vmstat 1".
4. Change directory to DMol3/benchamrks and run the benchmark by typing: "./RunBM".
5. When the benchmark completes, record the "time all done" result for both the AFY and TiN workloads.

Power measurement procedure

To record each server's power consumption during each workload, we used an Extech Instruments (www.extech.com) 380803 Power Analyzer / Datalogger. We connected the power cord from the server under test 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 the Power Analyzer's Data Acquisition Software (version 2.11) to capture all recordings. We installed the software on a separate Intel–processor-based PC, which we connected to the Power Analyzer via an RS-232 cable. We captured power consumption at one-second intervals.

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

We then recorded the power usage (in watts) for each server during the testing at one-second intervals. To compute the average power usage, we averaged the power usage during the time the server was producing its peak performance results. We call this time the power measurement interval. See Figures 5, 6, 8, and 9 for the results of these measurements.

Appendix A – Test server configuration information

This appendix provides detailed configuration information about each of the test servers, which we list in alphabetical order.

Processors	Dual-Core AMD Opteron 285	Dual-Core Intel Xeon processor 5160
System configuration information		
General		
Processor and OS kernel: (physical, core, logical) / (UP, MP)	2P4C4L / MP	2P4C4L / MP
Number of physical processors	2	2
Single/Dual-core processors	Dual	Dual
System Power Management Policy	Always On	Always On
CPU		
Vendor	AMD	Intel
Name	Dual-Core AMD Opteron 285	Dual-Core Intel Xeon processor 5160
Stepping	2	4
Socket type	940	LGA 771
Core frequency (GHz)	2.6 GHz	3.0 GHz
Front-side bus frequency (MHz)	2000 MHz HyperTransport	1333 MHz Dual Independent Busses (DIB)
L1 Cache	64 KB + 64 KB	32KB + 32KB
L2 Cache	2 MB (1 MB per core)	4MB (Shared)
Platform		
Vendor	Dual-Core AMD Opteron 285	Dual-Core Intel Xeon processor 5160 server
Motherboard model number	HP Proliant DL385 G1, P98330CRHSQ06H	Intel S5000PSL
Motherboard chipset	AMD 8111/8131 chipset	Intel 5000P Chipset
Motherboard revision number	C03	92
Motherboard serial number	USE606N50M	QSSL61900260
BIOS name and version	Hewlett Packard, version A05, 12/15/2005	American Megatrends, version S5000.86B.02.00.0054
BIOS settings	Default	Default
Memory module(s)		
Vendor and model number	Viking VI4CR287228ETPA2	Micron MT18HTF12872FDY
Type	PC-3200 Registered DIMM	FB-DIMM using PC2-5300 components
Speed (MHz)	400	667MHz
Speed in the system currently running @ (MHz)	400	667MHz
Timing/Latency (tCL-tRCD-iRP-tRASmin)	3-3-3-8	5-5-5-12
Size	4096 MB	4096 MB
Number of RAM modules	4	4
Chip organization	Double-sided	Double-sided
Channel	Dual	Dual
Hard disk		
Vendor and model number	Seagate ST336754LC	Seagate ST336754LC
Number of disks in system	1	1
Size	37 GB	37 GB
Buffer Size	8 MB	8 MB
RPM	15,000	15,000
Type	SCSI	SCSI

Controller	Adaptec SCSI 39320A-R PCI-X 133MHz	Adaptec SCSI 39320A-R PCI-X 133MHz
Controller driver	modprobe aic79xx	modprobe aic79xx
Operating system		
Name	SUSE Linux Enterprise Server 9	SUSE Linux Enterprise Server 9
File system	/ (Linux) 33.1GB / reiserfs /swap 1GB	/ (Linux) 33.1GB / reiserfs /swap 1GB
Kernel	Version 2.6.5-7.267-smp	Version 2.6.5-7.267-smp
Language	English	English
Graphics		
Vendor and model number	ATI Rage XL	ATI ES1000
Chipset	ATI Rage XL PCI (B41)	ATI ES1000 PCI
BIOS version	GR-xlcpq-5.882-4.333	GR-xlints3y.09a-4.332
Type	Integrated	Integrated
Memory size	8 MB	8 MB
Resolution	1024 x 768	1024 x 768
Driver	VESA	VESA
Network card/subsystem		
Vendor and model number	Broadcom BCM5704 dual NetXtreme Gigabit	Intel PRO/1000 EB Network Connection
Type	Integrated	Integrated
Driver	tg3	e1000
Optical drive		
Vendor and model number	Samsung SN-124	LITE-ON LTN-529SV
Type	CD-ROM	DVD/CD-ROM
Interface	Internal	Internal
USB ports		
# of ports	3	6
Type of ports (USB 1.1, USB 2.0)	USB 1.1	USB 2.0

Figure 8: Detailed system configuration information for the two test servers.

Principled Technologies, Inc.
4813 Emperor Blvd., Suite 100
Durham, NC 27703
www.principledtechnologies.com
info@principledtechnologies.com