Configuring Ninth-Generation Dell PowerEdge Servers

for High-Performance Computing Environments

Dell engineers tested performance settings for ninth-generation Dell™ PowerEdge™ servers using representative computation- and memory-intensive applications. Benchmark performance under various BIOS settings and memory configurations was also measured.

BY RIZWAN ALI; BARIS GULER; RAMESH RADHAKRISHNAN, PH.D.; AND VISHVESH SAHASRABUDHE

Related Categories:

Benchmarks

Dell ninth-generation servers

Dell PowerEdge servers

High-performance computing (HPC)

Parallel systems

Visit www.dell.com/powersolutions for the complete category index.

Ninth-generation Dell PowerEdge servers feature the Intel® 5000X or 5000P chipsets, each of which has a fast dual frontside bus (FSB) to enable increased CPU and memory bandwidth, quad-channel double data rate 2 (DDR2) memory in fully buffered dual in-line memory modules (DIMMs), and PCI Express I/O architecture. A team of Dell engineers in June 2006 used high-performance computing (HPC) benchmarks—Linpack; LS-DYNA; STREAM; NAS (NASA Advanced Supercomputing) Parallel Benchmarks (NPB); MIMD (multiple instruction, multiple data) Lattice Computation (MILC); FLUENT; and OOCORE—comparing various settings to help determine optimal configurations for HPC environments.

System performance is sensitive to several factors that depend on an application's behavior and how it interacts with the underlying hardware. Other than compiler and software optimizations, the hardware configuration (including memory, CPU speed, and disk speed) and BIOS-modifiable performance flags have an impact on application performance. This article focuses on two of these aspects: BIOS settings and memory configuration. All the results described in this article were achieved on a single server; the effect of these settings may vary in an HPC cluster. Furthermore, the results may also vary based on application data sets and parameters. Figure 1 describes the hardware and software environment that the Dell team used for testing performed on the Intel 5000X–based ninth-generation Dell PowerEdge 1950 server.

Benchmarks for evaluating HPC performance

The following synthetic and application-centric benchmarks were used in the Dell tests. These benchmarks and applications represent a broad spectrum of HPC workloads.

Linpack. This is a popular benchmark for HPC environments. The High-Performance Linpack (HPL)[1]

[1] For more information about HPL, visit www.netlib.org/benchmark/hpl.

Chipset	Intel 5000X
CPU	*BIOS tests:* Intel Xeon 5060 processor at 3.2 GHz with 2×2 MB L2 cache and 1,066 MHz FSB *Memory tests:* Intel Xeon 5080 processor at 3.73 GHz with 2×2 MB L2 cache and 1,066 MHz FSB
Memory	PC2-5300 fully buffered DIMMs from various vendors
Disk	Serial Attached SCSI (SAS)
OS	Red Hat® Enterprise Linux 4 Update 2 with support for Intel Extended Memory 64 Technology (EM64T)
Compilers	Intel C, C++, and Fortran Compilers version 9.0, build 20050914 (used where source code available)
Message Passing Interface (MPI)	MPICH version 1.2.6
FLUENT	Version 6.2.16
LS-DYNA	Version mpp970s and revision 6763.205

Figure 1. System configuration for BIOS and memory tests on a ninth-generation Dell PowerEdge 1950 server

implementation is commonly used to rank supercomputers on the TOP500 Supercomputer Sites list.

LS-DYNA. This general-purpose, transient dynamic finite-element program can simulate complex real-world problems.[2] LS-DYNA is optimized for shared and distributed memory on UNIX®, Linux®, and Microsoft® Windows® platforms. The specific LS-DYNA workload used in the Dell study was neon_refined.[3]

STREAM. The STREAM benchmark is a synthetic benchmark that measures sustainable memory bandwidth.[4] The Dell test team used the average of the results from the COPY, SCALE, ADD, and TRIAD kernels.

NAS Parallel Benchmarks. NPB is an application-centric suite of benchmarks that has been widely used to measure and compare the performance of parallel-processing computers.[5] The NPB suite consists of a set of eight programs, which are derived from computational fluid dynamics (CFD) code. The Dell team used the IS (Integer Sort) and LU (Lower-Upper Diagonal) B Class programs, and then calculated the sum of the IS and LU results to evaluate performance.

MILC. The code developed by the MILC Collaboration is used in high-energy physics for simulations of 4-D special unitary (SU) lattice gauge theory on MIMD parallel-processing systems. The workload used for this application was partially based on the MILC Collaboration's public lattice gauge theory code.[6]

FLUENT. A popular CFD application suite, FLUENT is used commonly in HPC environments. The FLUENT applications allow users to perform CFD analysis around their particular models. Several benchmark data sets (workloads) available from Fluent Inc. were used in the Dell tests.[7] To represent FLUENT benchmark performance, the Dell team took the geometric mean of the results.

OOCORE. An out-of-core matrix solver, OOCORE handles matrix equations that are too large for the cache.[8] This benchmark writes large amounts of data to the disk and thus also tests the disk I/O performance of the server.

Performance tests using various BIOS settings

The first part of the Dell study examined the effects that various BIOS settings have on HPC application performance. As part of this study, the test team ran benchmark tests on a Dell PowerEdge 1950 server equipped with the Intel 5000X chipset, two dual-core Intel Xeon® 5060 processors at 3.2 GHz, and eight 1 GB PC2-5300 fully buffered DIMMs. The BIOS settings in this study controlled Intel Hyper-Threading Technology (HT Technology), snoop filter, and processor prefetch capabilities.

Intel Hyper-Threading Technology. Available in Intel CPUs, HT Technology provides thread-level parallelism on each processor to enable efficient use of processor resources, high processing throughput, and enhanced performance for multi-threaded software.

Snoop filter. The Intel 5000X chipset uses a snoop filter to reduce the number of snoop cycles (which are generated to maintain cache coherency) on the processor FSB. Also known as a cache coherency filter, the snoop filter is a cache structure that exists on the chipset to help reduce unnecessary snoop traffic.

Prefetching. Fetching instructions or data from the memory well before the processor needs it is defined as prefetch. The prefetched instructions or data may simply be the next items in the program fetched while the current instruction is being executed. Prefetching may also be part of a complex speculative prediction algorithm, in which the processor tries to anticipate and fetch the appropriate instructions or data in advance.

In this study, the Dell team used four permutations of these three BIOS settings. Configuration A represents the baseline setting; the other three configurations are derived by starting from configuration A

[2] For more information about LS-DYNA, visit www.lstc.com.

[3] For more information about the neon_refined workload, visit www.topcrunch.org/benchmark_details.sfe?query=3&id=60.

[4] For more information about STREAM, visit www.cs.virginia.edu/stream/ref.html.

[5] For more information about NAS Parallel Benchmarks, visit www.nas.nasa.gov/Software/NPB.

[6] For more information about the MILC code, visit www.physics.utah.edu/~detar/milc/milcv6.html.

[7] For more information about these benchmark data sets, visit www.fluent.com/software/fluent/fl5bench/flbench_6.2.

[8] For more information about OOCORE, visit www.nsf.gov/pubs/2006/nsf0605/nsf0605.jsp.

and then changing one setting at a time. The configurations used in the study were as follows:

- **Configuration A:** Baseline setting—HT Technology off, prefetch on, and snoop filter off
- **Configuration B:** HT Technology enabled (prefetch on and snoop filter off)
- **Configuration C:** Snoop filter enabled (HT Technology off and prefetch on)
- **Configuration D:** Prefetch disabled (HT Technology off and snoop filter off)

The results of comparing configurations B through D against the baseline setting (configuration A) are shown in Figure 2. For example, the first set of bar charts shows the relative performance of each of the seven benchmarks when run under configuration B and compared against configuration A.

Evaluating the impact of Intel Hyper-Threading Technology

The first set of bars in Figure 2 show the effect of enabling HT Technology. NPB and OOCORE experienced a slight degradation in performance when HT Technology was enabled, whereas the other benchmarks showed little or no difference. Disabling HT Technology is recommended unless a significant performance benefit can be obtained from it.

Determining the effect of the snoop filter

The second set of bars in Figure 2 show the effect of enabling the snoop filter. The STREAM benchmark experienced more than 30 percent improvement with the snoop filter enabled. Similarly, the performance of NPB and MILC also improved. However, Linpack, LS-DYNA, and OOCORE experienced a slight degradation in performance. The snoop filter can cause some degradation in performance depending on the workload, so testing an application using both snoop filter settings is recommended to determine the optimal setting.

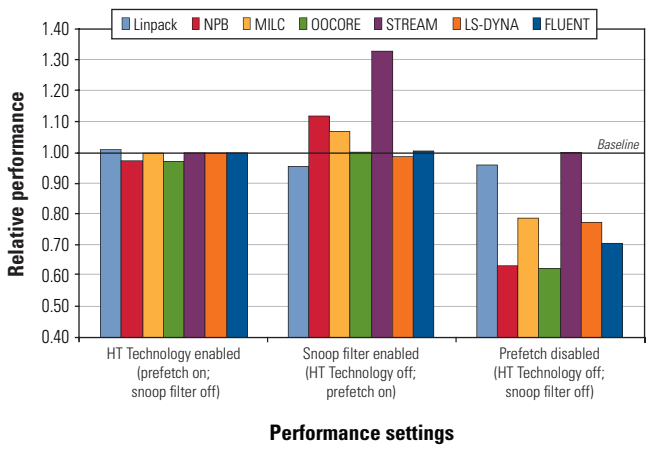

Figure 2. Relative performance of HPC applications for various BIOS configurations

Examining the benefits of prefetching

The last set of bars in Figure 2 show the effect of disabling prefetch settings in the BIOS. Except for STREAM, which was not affected, all the benchmarks showed significant degradation in performance when prefetch was disabled. Thus, enabling the processor prefetch settings is recommended for optimal performance of compute-intensive applications.

Performance tests using different memory configurations

The second part of the Dell study focused on comparing various memory configurations. Just as in the BIOS tests, the Dell test team used the seven benchmarks—Linpack, NPB, MILC, OOCORE, STREAM, LS-DYNA, and FLUENT. As part of this study, the Dell team ran benchmark tests on a PowerEdge 1950 server equipped with the Intel 5000X chipset and two dual-core Intel Xeon 5080 processors at 3.73 GHz. The team ran the tests on three memory sizes, each of which were configured with varying types and numbers of DIMMs:

- **2 GB configuration:** Two 1 GB DIMMs and four 512 MB DIMMs
- **4 GB configuration:** Two 2 GB DIMMs, four 1 GB DIMMs, and eight 512 MB DIMMs
- **8 GB configuration:** Four 2 GB DIMMs and eight 1 GB DIMMs

In addition to running each benchmark on the different memory configurations, the Dell team measured the average performance for all benchmarks on each configuration, as shown in Figures 3 and 4. The results represented in Figures 3 and 4 were obtained with the snoop filter enabled and disabled, respectively. For both figures, the baseline for the results is the system performance when using dual 1 GB DIMMs.

For the 2 GB configuration, performance improved on average by 30 percent when using four 512 MB DIMMs compared to two 1 GB DIMMs. Similarly for the 4 GB configuration, the four 1 GB DIMMs outperformed the two 2 GB DIMMs. This demonstrates the performance advantage of populating at least four DIMMs in the system so that each memory channel has at least one memory DIMM.

The results for four 1 GB DIMMs, eight 512 MB DIMMs, four 2 GB DIMMs, and eight 1 GB DIMMs show that configurations with one DIMM per channel performed similarly to configurations with two DIMMs per channel, especially when the snoop filter was disabled. Another interesting observation is that, with the snoop filter enabled, a higher delta in performance occurred between the baseline configuration—two 1 GB DIMMs—and the

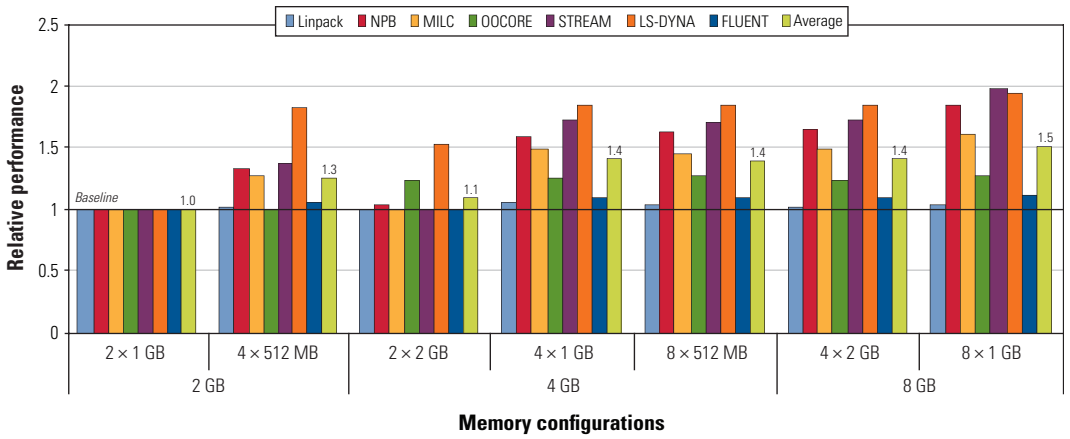

Figure 3. Memory configuration results with snoop filter enabled

rest of the configurations. With the snoop filter disabled, the average performance delta was around 30 percent; however, with the snoop filter enabled, the delta in performance was 40 percent. Furthermore, with the snoop filter enabled, there was approximately a 10 percent difference in performance between the two 8 GB configurations: the eight 1 GB DIMMs outperformed the four 2 GB DIMMs. Therefore, if an application benefits from using the snoop filter, it may show further improvements with fully populated DIMMs.

Optimal settings for ninth-generation Dell servers in HPC environments

The Dell study evaluated the performance effects of multiple BIOS settings and memory configurations for HPC benchmark applications running on a ninth-generation Dell PowerEdge 1950 server. As demonstrated in this study, using the appropriate BIOS settings and memory configurations can help achieve optimal server performance in HPC environments. Similarly, at least four DIMM slots should be populated on ninth-generation Dell servers to take advantage of the four memory channels offered on the Intel 5000X and 5000P chipsets. Heeding these guidelines can enable enterprise IT organizations to take advantage of the enhanced performance capabilities of the new-generation Dell PowerEdge servers.

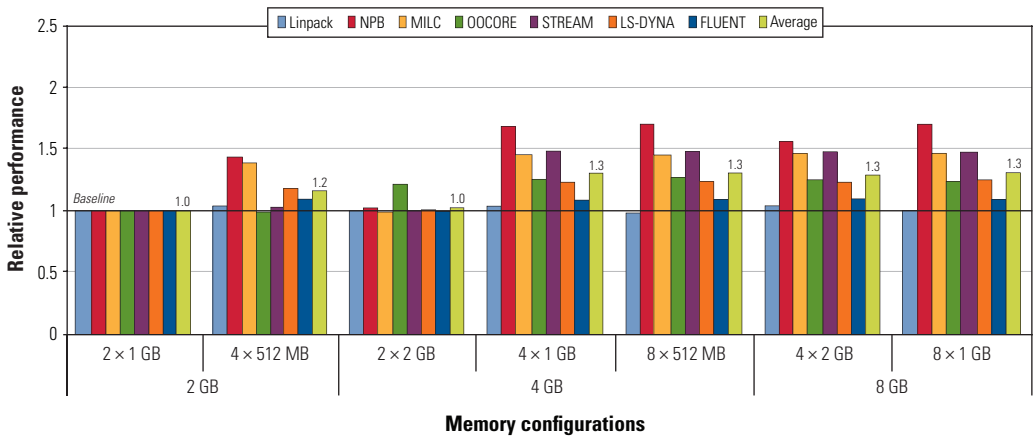

Figure 4. Memory configuration results with snoop filter disabled

Rizwan Ali is a systems engineer in the Scalable Systems Group at Dell. His current research interests include performance benchmarking, cluster architecture, parallel applications, and high-speed interconnects. He has a B.S. in Electrical Engineering from the University of Minnesota.

Baris Guler is an HPC application specialist in the Scalable Systems Group at Dell. His current research interests are parallel processing, diskless HPC clusters, performance benchmarking, reservoir engineering and simulation, and numerical methods. Baris has a B.S. in Petroleum and Natural Gas Engineering (PNGE) from the Middle East Technical University in Turkey, and an M.S. in PNGE from Pennsylvania State University.

Ramesh Radhakrishnan, Ph.D., is a member of the Scalable Systems Group at Dell. His interests include performance analysis and characterization of enterprise-level benchmarks. Ramesh has a Ph.D. in Computer Engineering from The University of Texas at Austin.

Vishvesh Sahasrabudhe is a systems engineer and adviser in the Scalable Systems Group at Dell. He has a bachelor of technology in Electrical Engineering from the Indian Institute of Technology in Bombay and an M.S. in Computer Science and Engineering from The Ohio State University.

