



TESLA C1060
DATASHEET

NVIDIA TESLA COMPUTING PROCESSOR SOLVE TOMORROW'S PROBLEMS TODAY

PETASCALE COMPUTING WITH TERAFLOP PROCESSORS

The NVIDIA® Tesla™ C1060 computing processor enables the transition to energy efficient parallel computing power by bringing the performance of a small cluster to a workstation. With 240 processor cores and a standard C compiler that simplifies application development, Tesla scales to solve the world's most important computing challenges—more quickly and accurately.

FEEDING THE RELENTLESS DEMAND FOR HPC PERFORMANCE

The NVIDIA® Tesla™ C1060 closes the gap between the demands placed by application performance and performance delivered by the computing processor. With the massively parallel architecture of the GPU, scientists and engineers can get a quantum jump in performance and continue to advance the pace of their work, guiding us to faster discovery in drug research, weather modeling, oil and gas exploration, computational finance, and more.

MANY-CORE ARCHITECTURE DELIVERS OPTIMUM SCALING ACROSS HPC APPLICATIONS

Scientists and engineers have made the transition to 'many core' computing because their problems have reached a complexity that required them to look for new ways to boost their work. With parallel performance from 240 cores, capable of concurrent execution of thousands of computing threads, these professionals can get accurate answers, faster. The number of processor cores has replaced MHz as the new metric for performance.


Revolutionary scalar, single, and double-precision floating-point performance enables solving a wide range of high-performance computing applications, whose complexity has outstripped the CPU's ability to solve them.

HIGH-EFFICIENCY COMPUTING PLATFORM FOR ENERGY- CONSCIOUS ORGANIZATIONS


Unlike any other solution available in the HPC space, the Tesla C1060 delivers cluster-level performance in a power-efficient workstation. This "high density computing" allows every technical professional to have a dedicated high-performance supercomputer at their desk-side instead of building up a small cluster in the data center. This gives green-conscious organizations an excellent choice to achieve their desired level of computation without the build-up, management, and energy cost of a data center.

NVIDIA CUDA™ TECHNOLOGY UNLOCKS THE POWER OF TESLA'S MANY CORES


The CUDA C programming environment simplifies many-core programming and enhances performance by offloading computationally-intensive activities from the CPU to the GPU. It enables developers to utilize NVIDIA GPUs to solve the most complex computation-intensive challenges such as protein docking, molecular dynamics, financial analysis, fluid dynamics, structural analysis and many others.



MOLECULAR DYNAMICS



LIFE SCIENCES



SEISMIC EXPLORATION

FEATURES AND BENEFITS

MASSIVELY-PARALLEL MANY CORE ARCHITECTURE WITH 240 PROCESSOR CORES	Solve compute problems on your workstation that traditionally required a cluster installation.
WIDELY ACCEPTED, EASY TO LEARN CUDA C PROGRAMMING ENVIRONMENT	Easily express application parallelism to take advantage of the GPU's many core architecture.
SCALE TO MULTIPLE GPUS AND HARNESS THE PERFORMANCE OF THOUSANDS OF PROCESSOR CORES	Solve large-scale problems by scaling to 1000s of cores across multiple GPUs.
IEEE 754 SINGLE & DOUBLE PRECISION FLOATING POINT UNITS	Achieve the highest floating point performance from a single chip, while meeting the precision requirements of your application.
ASYNCHRONOUS TRANSFER CAPABILITY	Turbo charges system performance by overlapping data transfers with computation.
4 GB GLOBAL MEMORY	Compute on larger data sets than before.
SHARED DATA MEMORY	Groups of processor cores can collaborate using low latency memory.
HIGH SPEED, PCI-EXPRESS GEN 2.0 DATA TRANSFER	Fast and high-bandwidth communication between CPU and GPU.
TESLA PROCESSORS AVAILABLE IN FLEXIBLE FORM FACTOR	Tesla workstation products and 1U systems enable deployment in a wide range of environments.

TECHNICAL SPECIFICATIONS

SUPPORTED OPERATING SYSTEMS

- > Microsoft® Windows® XP (64-bit and 32-bit)
- > Linux® (64-bit and 32-bit)
 - > Red Hat Enterprise Linux 4 and 5
 - > SUSE 10.2 and 10.3

TESLA ARCHITECTURE

- > Massively-parallel many-core architecture
- > 240 scalar processor cores per GPU
- > Integer, single-precision and double-precision floating point operations
- > Hardware Thread Execution Manager enables thousands of concurrent threads per GPU
- > Shared data memory enables processor cores to collaborate on shared information at local cache performance

- > Ultra-fast memory access with 102 GB/s peak bandwidth per GPU
- > IEEE 754 single-precision and double-precision floating point

SCALABLE SOLUTIONS

- > Scalable from one to thousands of GPUs
- > Available as computing processor and 1U rack-mount GPU computing system.

SOFTWARE DEVELOPMENT TOOLS

- > C language compiler, debugger, profiler, and emulation mode for debugging
- > Standard numerical libraries for FFT (Fast Fourier Transform), BLAS (Basic Linear Algebra Subroutines), and CuDPP (CUDA Data Parallel Primitives)

PRODUCT DETAILS

- > Tesla C1060 Computing Processor
 - > 1Tesla GPU (240 processor cores)
 - > 4 GB dedicated memory
 - > Standard length, dual slot card that fits in one PCI Express x16 (Gen 1 or 2) slot (Gen 2 recommended for best results)

To learn more about NVIDIA Tesla, go to www.nvidia.com