


BOOST YOUR CODE

NVIDIA® TESLA® K80 FOR THE ACCELERATED DATA CENTER




ENGINEERED TO MAXIMIZE APPLICATION PERFORMANCE

DUAL GPU ACCELERATOR

Dual GPU design allows for higher overall application throughput.

GPU BOOST

Dynamic GPU boost automatically maximizes application performance by taking advantage of any available power headroom.



24 GB GPU MEMORY


Double memory enables the K80 to run bigger data applications.

2X SHARED MEMORY


Enables more concurrent threads to deliver significant speedup without changes to GPU-accelerated code.

A GIANT LEAP IN PERFORMANCE


PEAK DOUBLE PRECISION FLOPS



PEAK MEMORY BANDWIDTH




PERFORMS 2-4X FASTER THAN ITS PREDECESSORS



*CPU comparison not available.
CPU Server: Dual-socket E5-2697 v2 @ 2.7GHz, GPU Server: Dual-socket E5-2697 v2 @ 2.7GHz, Dual Tesla M2090/K20/K80; K80 GPU Boost enabled.

MOLECULAR DYNAMICS

AMBER: DAYS TO SIMULATE 1 MICROSECOND



CPU Server: Dual-socket E5-2697 v2 @ 2.7GHz, GPU Server: Dual-socket E5-2697 v2 @ 2.7GHz, Single Tesla M2090/K20/K80; K80 GPU Boost enabled, Dataset: 1 microsecond PME-JAC_NVE simulation on AMBER 14


Eight of these Tesla K80s in one system combine 16 GPUs in a node—that's almost 1.4 microseconds aggregate of MD per day for a 25K atom system!

For the same workload, a dual-socket CPU server would take over two and a half months to complete.

Prof. Ross Walker, San Diego Supercomputer Center

QUANTUM CHEMISTRY

TERACHEM: WEEKS TO SIMULATE 25ps OF TRPCAGE AB INITIO PROTEIN DYNAMICS



CPU Server: Dual-socket E5-2697 v2 @ 2.7GHz, GPU Server: Dual-socket E5-2697 v2 @ 2.7GHz, Dual Tesla M2090/K20/K80; K80 GPU Boost enabled, Dataset: TeraChem TRPCage. Extrapolated from a single point energy calculation


TeraChem is used by researchers worldwide to understand electronic structures for discoveries ranging from efficient photovoltaic material to drug development.

Tesla K80 delivers 2.8x performance over Fermi GPUs, enabling researchers to accelerate their computation from months to weeks.

Prof. Todd Martinez, Stanford University

OIL AND GAS

ECHELON: WEEKS TO COMPUTE RESERVOIR SIMULATION



Dual 2x Xeon E5-2630v3 @ 2.4GHz and 1x K80 node; Dual Xeon E5-2640 @ 2.5GHz and 1x K10 node; Dual Xeon E5-2620 @ 2.0GHz and 1x M2090 node. Dataset: Dual-porosity simulation of a hydraulically fractured reservoir with 2M active cells for over forty years of production.

Oil and gas companies can accelerate reservoir simulations with Tesla K80 by 2-3x compared to Tesla K10 and M2090 GPUs.

In addition, Tesla K80 handles 3-4x larger models, so engineers can run more detailed simulations.

Ken Esler, Senior Physicist, Stoneridge Technology

TEST DRIVE TESLA K80 FOR FREE

Take your GPU-accelerated application for a spin on a remote server.

Sign up at www.nvidia.com/gputestdrive