

## Editor's Pick: NVIDIA Delivers the Personal Supercomputer

Tesla GPUs make desktop cluster-level computing performance accessible.

by Anthony J. Lockwood | Published January 14, 2009

Dear *Desktop Engineering*

Reader:



Before desktop computers, the only computer you had access to was guarded by guys wearing white shirts and pocket protectors. You'd beg them to run your job, but it was always "by the end of the week; we have other jobs before yours." Lately, we have taken the empowerment that desktop computers gave us back to the future with cluster computing. Only now it's job schedulers and available CPUs that force you to wait to run a job. Ah, but once again, the field tilts in your favor. NVIDIA just introduced its Tesla Personal Supercomputer, which brings cluster-computing performance to your desktop.

The first thing you have to understand about the NVIDIA Tesla Personal Supercomputer is that it is a technology that will be leveraged by computer manufacturers like Dell, Lenovo, Microway, or Penguin Computing. (Yes, you craft-your-own types will be able to purchase Tesla components.) It is also a set of system specifications — Quad-core AMD Opteron or Intel Xeon CPU, 12GB memory, and the like. So, you have to look for NVIDIA's certification logo on a unit to ensure you're getting the real McCoy.

That said, so what is a Tesla Personal Supercomputer? Basically, it is a parallel-core computer within a multi-core workstation. It starts with an NVIDIA Tesla card. The Tesla C1060 has 240 processor cores per GPU (graphics processing unit) with a GPU memory access of 102GBs peak bandwidth per GPU. Each Tesla C1060 GPU delivers 933 GFlop single-precision and 78 GFlop double-precision performance, according to NVIDIA's specs.

A key here is that all this is based on NVIDIA's CUDA massively parallel architecture, which, to make a long story short, is designed to make multi-core processing efficient and easy. But the thing here is that the Tesla card works in cahoots with a multi-core processor. It can take on and solve complex computational processing problems in less time than the CPU can do it and, by doing so, it frees the CPU to do other stuff. So, performance blasts through the roof: NVIDIA says that a Tesla Personal Supercomputer can give you up to 250 times the performance of a regular workstation in the same box and using essentially the same electronics found in standard workstations.

The bottom line for engineers and scientists is that you get on your desktop hundreds of cores to process your CFD and oil and gas simulations or crunch massive datasets. You also get the satisfaction of knowing that queuing up your job on the cluster, then hoping that load demands do not bump you off schedule for a few days, is a thing of the past. The speed and performance boost should be complemented by a boost in your productivity, not to mention energy saved bypassing the data center.

And, get this, NVIDIA estimates that workstations based on the Tesla Personal Supercomputer will come in around \$10,000. They'll plug into your regular old power strip, and the cherry on top is that the specification calls for quiet operation.

Check out NVIDIA's Tesla Personal Supercomputer in **today's Pick of the Week write-up**. You'll find lots of links to more information about manufacturers offering desktop supercomputers already.

Thanks, pal. — Lockwood

Anthony J. Lockwood

Editor at Large, *Desktop Engineering*

Magazine

Courtesy of: [DE Online](#)