


BOOST THROUGHPUT AND DISCOVERIES WITH NVIDIA TESLA K80.

Three Reasons to Deploy Tesla K80 in Your Data Center



The path forward is parallel. Traditional CPUs no longer deliver generational performance gains like they used to, and HPC experts agree that only powerful accelerators like GPUs will deliver the performance required to drive future scientific discoveries.

The NVIDIA® Tesla® K80 is the world's most powerful accelerator built for Machine Learning and HPC applications, offering three powerful reasons to add them to your data centre.

Reason 1: Top Applications are GPU Accelerated

Many of the compute-intensive applications running in your Machine Learning and HPC system are GPU-accelerated. In fact, an independent study by Intersect360 Research shows that 70% of the most popular HPC applications have built-in support for GPUs.


For example, chemistry applications consume a lot of compute cycles in a typical data center. According to the study, most computational chemistry applications offer support for GPUs today. The NVIDIA Tesla K80 GPU accelerator can deliver a significant performance boost for these applications.

Intersect360's Top Chemistry Applications

Gaussian	GAMESS	CP2K
NAMD	Quantum Espresso	Q-Chem
VASP	LAMMPS	ADF
AMBER	Schrodinger	Accelrys
NWChem	CHARMM	

Green means GPU-accelerated

TOP HPC APPLICATIONS 5-10X FASTER WITH TESLA K80




Reason 2: Replace up to Eight CPU Servers with One Tesla K80-Powered Server

While time to discoveries is critical, compute resources are often scarce. This limits the scale of science that researchers can achieve.


The Tesla K80 redefines what's possible for scientists and researchers. For Many GPU-optimised applications, a single K80 server can deliver higher throughput than eight dual-socket CPU servers.

For example, running a popular computational chemistry application - HOOMD-Blue a single server with two Tesla K80 accelerators delivers 1.7x faster performance compared to 8 servers with dual-socket Haswell CPUs and 3.4x faster performance than four servers with dual socket Haswell CPUs.

Application Performance Benchmark
4 CPU Servers vs 1 Tesla K80 Server



Application Performance Benchmark
8 CPU Servers vs 1 Tesla K80 Server



Reason #3: Higher Throughput Leads to More Discoveries

Data centre managers all face the same challenge: how to meet the demand for computer resources that often exceed cycles available in the system.

Tesla K80 can dramatically boost the throughput of your data centre, completing more compute jobs every day. When your key applications are accelerated by a third of the systems nodes having Tesla K80 accelerators, the data centre would experience more than two-fold increase in overall throughput.


Accelerated computing enables researchers to tackle problems that simply would not be possible with CPU-based approaches. For example, TU Dresden expanded their supercomputer with Tesla K80 accelerators when traditional systems could no longer meet the growing demand by scientists and researchers. For their current mix of applications, TU Dresden estimates that 64 servers with Tesla K80 deliver similar performance of 1,400 CPU servers.

“The demand for GPU was so high that the IT department decided to replace their planned purchase of a future CPU system with 128 nodes with Tesla K80 GPUs, and utilisation jumped to 80% on the first day.”

Guido Juckeland, TU Dresden

System Throughput Increases when Nodes are Accelerated

CPU-only System



100 Jobs per Day

1/3 of System Nodes with Tesla K80 Accelerators



220 Jobs per Day



OCF is a high performance compute, storage and data analytics provider. Our flexible, scalable and unobtrusive high performance server clusters are powered by NVIDIA GPU accelerators to deliver application performance gains for customers.

Contact OCF

www.ocf.co.uk | +44 (0)114 2572200
info@ocf.co.uk

