



Hardware

Skip the Data-Center — High-Performance Computing and Simulation Under Your Control

At-a-Glance

Through a partnership with Exxact Corporation, we offer a powerful, integrated platform for *MedeA* atomistic simulations running VASP, LAMMPS and related *MedeA* modules on NVIDIA Graphical Processing Units (GPUs) and Intel or AMD CPUs. State-of-the-art hardware and design bring affordable, high-performance computing out of the data center and into your office. Optionally, the *MedeA* materials simulation environment can be pre-installed on the system prior to shipment and is ready to run.

You can also connect to your HPC hardware from a laptop or desktop running the *MedeA* user interface, where the HPC hardware only hosts the *MedeA* JobServer and TaskServer services.

Key Benefits

- Minimal setup time
- High-Performance GPU Computing without complex hardware setup and Linux system configuration
- *MedeA Environment* benefits, including job management and performance monitoring
- Remote access from anywhere via configured virtual private network (VPN)
- Workstations are quiet enough for the office, and need no special cooling or electrical requirements
- Fast, integrated tech support and scientific consulting on hardware, OS, and *MedeA*

Note

We focus here on state-of-the-art GPU-based workstations, but Exxact can also provide HPC data center hardware to run *MedeA* software, such as CPU and GPU clusters.

Our Partnership with Exxact Corporation

Materials Design is proud to partner with Exxact Corporation to supply cost-effective HPC hardware for materials simulations. We collaborate on validating and benchmarking *MedeA* software on Exxact's systems.

Exxact Corporation was founded in 1992, and provides technology solutions that enable many of the world's top national laboratories, research institutes, universities, startups, and Fortune 1000 companies to conduct cutting-edge artificial intelligence, drug discovery, and advanced engineering simulations, while also powering a host of other HPC, datacenter, and visualization applications.




Fig. 1: Examples of Exxact VASP Optimized GPU Advanced Simulation Workstation (a), 4U Rack-Mountable Workstation (b), and 4U 8 GPU Server (c)

System Integration

As computational challenges increase in complexity, Materials Design works with Exxact engineers to provide the right balance of computation, networking, storage, and memory in order to deliver cost effective and reliable simulation performance.

Quality Service

Exxact corporation's extensive HPC expertise and unparalleled customer support allows customers to increase and sustain their computational infrastructure through personalized service, life cycle

management, product allocation, and end-of-life product sourcing.

Designed for Performance and Ease of Use

With the MedeA materials simulation environment preinstalled, you can simply unpack your system, connect monitor, keyboard, and mouse, plug it in; and start modeling at a remarkable level of performance, all under your own control, with or without a dedicated data center.

Exxact Corporation MedeA configurations can be designed to be quiet and energy efficient, plug into a standard wall socket, and not require special cooling.

Key Features

Exxact Corporation systems provide high efficiency for all engines supported by the *MedeA Environment*:

- *MedeA VASP* — Vienna Ab-Initio Simulation Package
- *MedeA LAMMPS* — Large-scale Atomic/Molecular Massively Parallel Simulator
- *MedeA GIBBS* — Monte Carlo simulation for Fluid properties and Sorption
- *MedeA MOPAC* — Molecular Orbital PACKage for fast screening of molecular systems and solids
- *MedeA Gaussian GUI* — The standard in Computational Chemistry
- *MedeA PhaseField*

A typical configuration can provide a peak performance of well over 3 TFLOPS, is expandable through the addition of compute nodes or GPUs, is accessible via a VPN for secure remote access, and supports access to the *MedeA* simulation environment via laptop or desktop MedeA clients.

Example Specifications

GPU Advanced Simulation Workstation Tower

A single workstation in a standard tower case with up to 4 GPUs, up to 56 CPU cores.

| Property/Device | Value |
|-----------------|------------------------------|
| CPU processor | 1x Intel Xeon W-3400 Series |
| GPU | Up to 4x NVIDIA RTX 6000 ADA |
| Memory | Up to 1TB DDR5 ECC Memory |
| MedeA | Version 3.x (most recent) |

Rackmountable GPU Workstation

A single workstation in a 2U rack mount chassis with 64 to 128 cores.

| Property/Device | Value |
|-----------------|---|
| CPU processor | 2x 4th/5th Gen Intel Xeon Scalable Processors |
| GPU | Up to 4x Double-Wide GPUs: RTX 6000 ADA, A800, etc. |
| Memory | 16 DDR5 ECC DIMMs (up to 2TB) |
| Storage | Multiple 3.5" Hard Drive and 2.5" SSD |
| Interconnect | 25GB Ethernet or 100GB InfiniBand |
| MedeA | Version 3.x (most recent) |

HPC 8 GPU Server

An HPC server in a 4U rack mount chassis with 64 to 128 cores.

| Property/Device | Value |
|-----------------|---|
| CPU processor | 2x AMD EPYC 7003/7003X CPUs |
| GPU | Up to 8x Double Wide GPU: H100, RTX 6000/5000 Ada, and more |
| Memory | 16 DDR4 ECC DIMM Slots (Up to 4TB) |
| Storage | 8x 3.5" Hot-swap (4x SATA, 4x NVMe) |
| MedeA | Version 3.x (most recent) |

HPC Cluster

Please contact support@materialsdesign.com or sales@exxactcorp.com for additional specs and recommendations.

Performance and Scaling on CPUs and GPUs.

When deciding on the appropriate hardware for your simulation needs, it is helpful to understand how speed and memory requirements scale with the size of the system (usually, number of atoms) for a chosen accuracy of the computational method. In order to make optimal use of your hardware and to obtain results in the shortest time possible, the *MedeA* software supports two main parallelization modes for your simulations:

- High-throughput mode, where you run many small calculations in parallel, using *MedeA*'s out-of-the-box support for high-throughput computing. In this mode, no communication is required between individual calculations, therefore this is often the most scalable mode. With this mode, you can expect near-perfect “strong scaling”, limited only by the number of cores and the available memory of your hardware.
- Algorithmic parallelization mode, where individual algorithmic operations of the simulation codes (such as VASP and LAMMPS) are spread out across multiple CPU and GPU cores. To run a single large calculation on hardware with many CPU or GPU cores, the speed of the calculation will depend on many factors, for example system size, chosen computational accuracy, number of CPU/GPU cores and their clock speed, memory bandwidth and latency, and the interconnect latency and bandwidth between processors. It may also be limited by disk I/O speed and total available memory. Performance scaling with the number of processor cores typically tapers off when the number of parallel threads exceeds a certain limit (“weak scaling”).

You can also combine these two parallelization strategies.

With the availability of modern GPUs, a key purchasing decision is whether and how much to invest in GPUs. Both the VASP and LAMMPS software are optimized for NVIDIA GPUs but also perform well on multi-core CPUs. Our benchmark calculations indicate that an investment in GPUs makes most sense for computations on large systems (several hundred atoms), where GPUs substantially outperform CPUs. These speedups are

less dramatic for smaller systems.

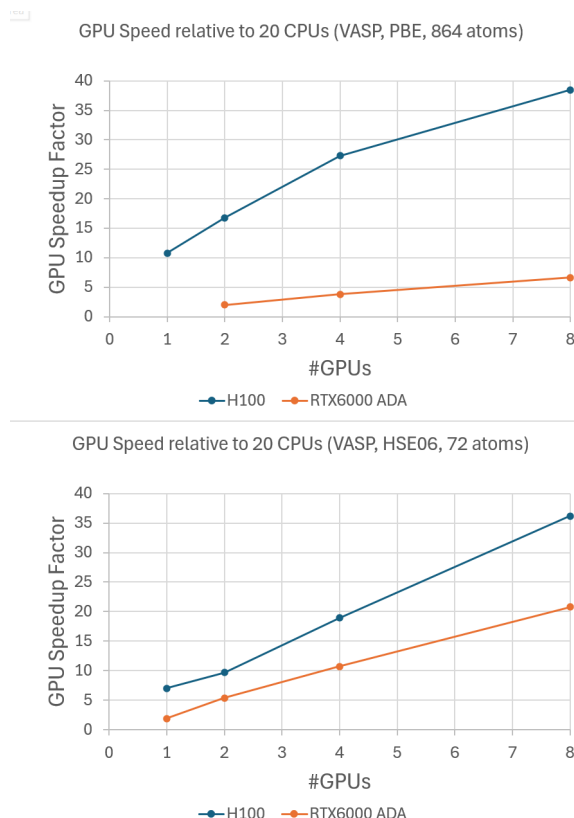


Fig. 2: Performance of two types of NVIDIA GPUs relative to calculations on an Intel Xeon Silver 4416+ 20 core CPU. The computational speedup factor is shown for two typical VASP calculations: 864 atoms with PBE density functional and blocked Davidson Optimizer (top), and 72 atoms with the HSE06 hybrid functional and Damped Molecular Dynamics Optimizer (bottom). The speedups are relative to the CPU-only system.

For illustration, the Figures 2 and 3 show speedup factors with higher-end NVIDIA GPUs compared to a traditional 20-core CPU, for typical VASP calculations with different density functionals and optimizers.

For additional benchmarks and discussions of your needs, please contact customer support.

Required Modules

- *MedeA Environment*

Recommended Modules

- *MedeA VASP*
- *MedeA MOPAC*
- *MedeA Gaussian GUI*
- *MedeA GIBBS*

- *MedeA HT-Launchpad*
- *MedeA PhaseField*

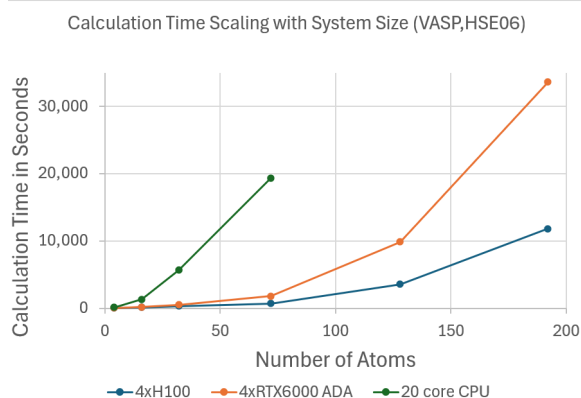
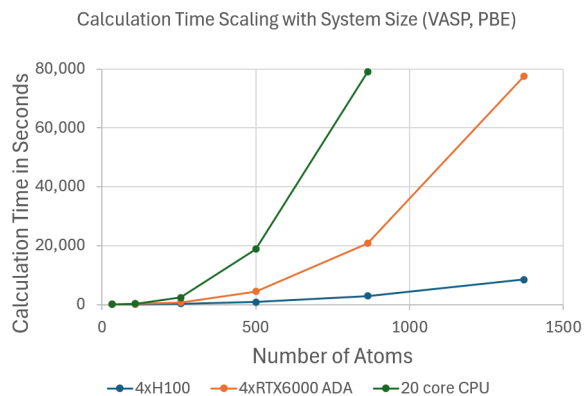


Fig. 3: Comparison of calculation times as a function of the number of atoms, for 4xNVIDIA RTX6000 ADA GPUs and 4xNVIDIA H100 GPUs, compared to computations on an Intel Xeon Silver 4416+ 20 core CPU, for two types of VASP calculations: PBE density functional and blocked Davidson Optimizer (top) and HSE06 hybrid functional and Damped Molecular Dynamics Optimizer (bottom)

Supported Modules

- All *MedeA* components

Find Out More

Learn more about how *MedeA* can support your work through capabilities such as [Databases](#), [Builders](#), [Compute Engines](#), [Forcefields](#), [Property Modules](#), [Analysis Tools](#), and [High-Throughput](#).