InspireSemi™ Thunderbird™

Disruptive Next Generation HPC-AI Accelerated Computing Platform

Providing the "ground truth" for AI in science & engineering

Accelerating the datacenter HPC-AI market

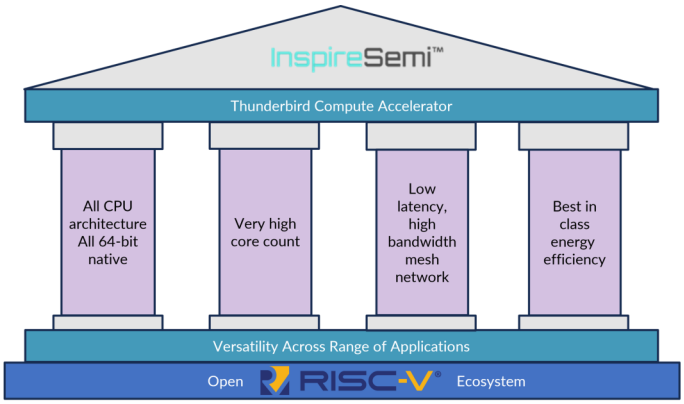

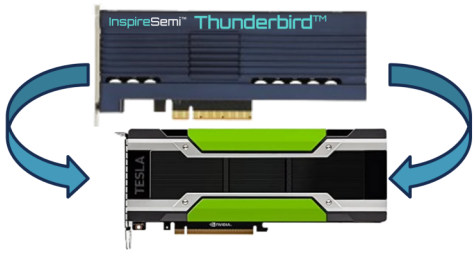

- Versatile all-CPU architecture applicable to all HPC-AI software, much of which does not benefit from GPUs
- High precision, high performance, low power 64-bit native processors to solve "big math" problems required by most HPC software

- Thunderbird runs the highly accurate physics-based HPC simulations needed to train AI surrogate models for science and engineering
- GPUs run fast AI surrogate models to approximate long-running physics-based HPC simulations

Thunderbird accelerates <u>All</u> HPC-AI software


Financial simulations

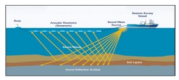
Geology: Seismic


Financial Trading & Graph Analytics

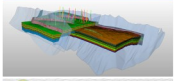
Energy: Reservoir Modeling & Sim

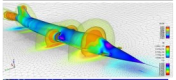
CAE/Computational Fluid Dynamics

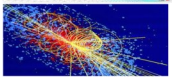
Government Lab Simulations

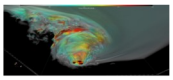
Climate & Weather Modeling


Cybersecurity & Cryptography

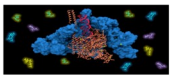
Genomics, Pharma, Life Sciences

Datacenter GPUs primarily focused on AI

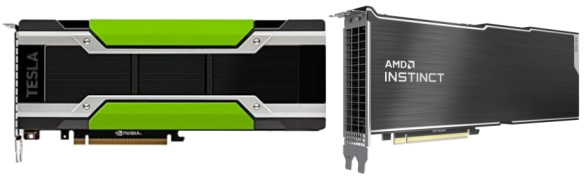

InspireSemi Thunderbird

Highly differentiated "supercomputer-cluster-on-a-chip"

- Versatile platform delivers unprecedented capability
- 4 chip PCIe card delivers >6,000 64-bit CPU cores (FP64)
- Innovative high-bandwidth, low-latency on-chip network
- Best-in-class for both Performance/$ and Watt
- Large scale computing power with much lower TCO - can replace many racks of servers and expensive high-speed networking