

InspireSemi™

Disruptive Next Generation HPC-AI Accelerated Computing Platform

InspireSemi provides revolutionary high-performance, energy-efficient accelerated computing platforms for High-Performance Computing (HPC), AI, graph analytics, and other compute-intensive workloads. The Thunderbird “supercomputer cluster-on-a-chip” is designed to address multiple underserved and diversified industries.

OVERVIEW

New standards of speed, efficiency, and versatility

Existing HPC, AI, & Graph Analytics accelerated compute solutions are no longer “good enough”. They are expensive, hard to program, power hungry, and only work for limited applications. InspireSemi is poised to set new standards of speed, versatility, and efficiency for computationally-intensive applications. Led by an accomplished team with a proven track record, we are leveraging an efficient operating model with world class outsourced manufacturing partners (e.g.- TSMC, ASE). Leading industrial companies, national labs, and computer OEMs value our technology, innovative North America design team, and scalable supply chain.

Technology

Our versatile Thunderbird “supercomputer cluster-on-a-chip” accelerated computing architecture contains thousands of modern, powerful, and efficient 64-bit CPU cores, all connected via our high-speed mesh network fabric. These and other architectural advantages result in real world application benefits, and it is compatible with an established, thriving open software ecosystem to support a broader range of applications than GPUs.

MARKETS

Engineers and scientists can open new frontiers with fast, energy efficient, easy to program, compute accelerators, with an open standard CPU software programming model that eliminates vendor lock-in.

AI developers can iterate their AI surrogate models on GPUs with the high precision, physics-based “ground truth” models on Thunderbird to ensure accuracy and accelerate innovation, insights, and market impact.

Graph analytics applications now can scale and support extremely large models by leveraging the thousands of high-performance interconnected CPU cores and cache memory on each chip.



Alex Gray, Founder, CTO & President

Alex is the founder of InspireSemi and inventor of its processor architecture. He is a versatile technical leader skilled in many areas of electronic design and other disciplines. He previously founded CryptoCore Technologies and several small businesses. He has over a decade of design experience with companies including SunPower & SolarBridge Technologies.



James Hickman, Chairman & CEO

James is a West Point graduate and former US Army intelligence officer who has had an extensive business and investment career spanning more than 25 years. He has been CEO of multiple startups and was investor #1 in InspireSemi. He has served as Chairman of the Board since 2022 and became CEO upon his predecessor's retirement on January 1, 2026. He earned a BS in Mathematics and Computer Science from the US Military Academy at West Point and a MS in Accounting and Information Management from the University of Texas at Dallas.



Doug Norton, Chief Marketing Officer

Doug has over 35 years of experience in business development, marketing & sales of HPC, AI and enterprise IT solutions in both public and startup companies. He brings a customer-focused approach to market development, go to market strategy, creating OEM, channel, & technology partnerships, and building worldwide sales teams. He began his career at IBM, going on to various leadership positions at Cadence Design Systems, CoWare, Newisys, Virtana, & Nimbit. He is also President of the non-profit HPC-AI Society.



Thomas Fedorko, Chief Operating Officer

Thomas is an entrepreneurial executive bringing more than 35 years of hands-on technical and business leadership in semiconductor Operations. He began his career as a product engineer at Motorola Semiconductor who delivered the first microprocessor for the Apple Macintosh and went on to lead Operations teams for several successful startup companies including Eta Compute, Uhnder, Bluetechnix & Luminary Micro.



Jack Cartwright, Chief Financial Officer

Jack has over 20 years of experience in finance, M&A and operations, with a strong focus in high growth venture & PE backed startups. He most recently served as CFO of Datum Tech, a leading data and advertising technology company. Prior to Datum, he was CFO of Spruce Services, Workrise (RigUp), and finance roles at prestigious organizations such as BT Global, Atkins Global, and Cisco Systems.

Find out more: www.inspiresemi.com

InspireSemi™

SOLUTIONS

Setting new standards of performance, energy efficiency and versatility for a range of applications

InspireSemi's powerful Thunderbird "supercomputer cluster-on-a-chip" accelerated computing architecture is based on an array of thousands of modern, efficient, and powerful 64-bit CPU cores, tightly integrated with high-speed memory and networking. It is an ideal platform for high-performance computing (HPC), AI/machine learning, graph analytics, and other compute-intensive applications.

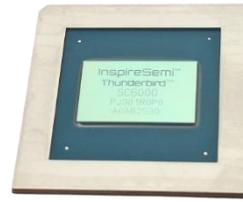
Leveraging the modern, efficient, open-standard RISC-V CPU ISA (instruction set architecture) provides access to an established, thriving ecosystem of software, libraries, and development tools. Our developer-friendly all-CPU programming model eliminates vendor lock-in and greatly simplifies software development, QA, and maintenance, since there is no need to support multiple and often proprietary software stacks (e.g.- CUDA, ROCM, FPGAs). This also enables the combination of HPC & AI on one platform for deeper insights and bigger advancements.

Our highly optimized, high speed mesh network fabric provides high bandwidth and low latency communications between cores. Tightly integrated with the high-performance, low-power CPU cores and distributed memory, it removes crucial bottlenecks from applications that depend on close cooperation between many independent threads. Arrays of up to 256 Thunderbird chips connect seamlessly via high-speed serial transceivers.

Compared to a modern datacenter GPU, Thunderbird provides all the raw computational power with far broader application to real world codes. Our advanced network fabric enables far better cooperation between cores. And our roots in extremely power-sensitive blockchain computing enable us to provide this performance with considerably lower energy consumption and cost than leading GPU's.

Raw performance: 24 TFLOPS (FP64)
Energy efficiency: 75 GFLOPS/Watt (FP64)
Utilization: network reduces latency 10x
Deterministic + Predictable Performance
Value: exceptional Performance/\$

Innovate with InspireSemi
Contact us to learn more



Thousands of powerful RISC-V CPU cores & established, thriving software ecosystem to support broader range of applications than GPUs



Developer-friendly CPU programming model greatly simplifies software & algorithm development, eliminates vendor lock-in, and provides maximum versatility for HPC, AI, Graph Analytics, and other applications



Ample high-speed SRAM memory tightly integrated with CPU processor cores



Dedicated acceleration for several cryptography algorithms



Extreme-bandwidth, low-latency, on-chip mesh network fabric for both inter- and intra-chip connectivity, up to 256 Thunderbird chips



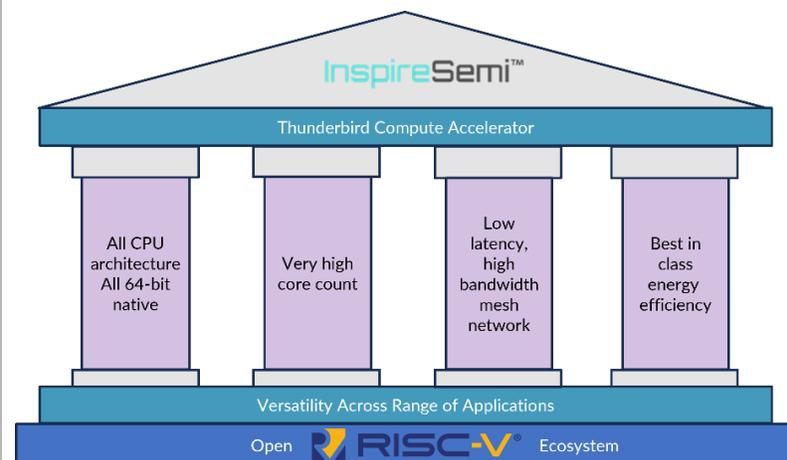
Deterministic + predictable performance addresses applications GPUs cannot, including real-time and real-safe computing



Support for large amounts of onboard LP-DDR memory, NVMe storage, PCIe, and 10Gb ethernet



Best in class energy efficiency



5301 Southwest Parkway
Building 1, Suite 100
Austin, TX 78735
(737) 471-3230
Info@inspiresemi.com