InspireSemi™

Disruptive Next Generation HPC-AI Accelerated Computing Platform

Thunderbird™

"Supercomputer cluster-on-a-chip"

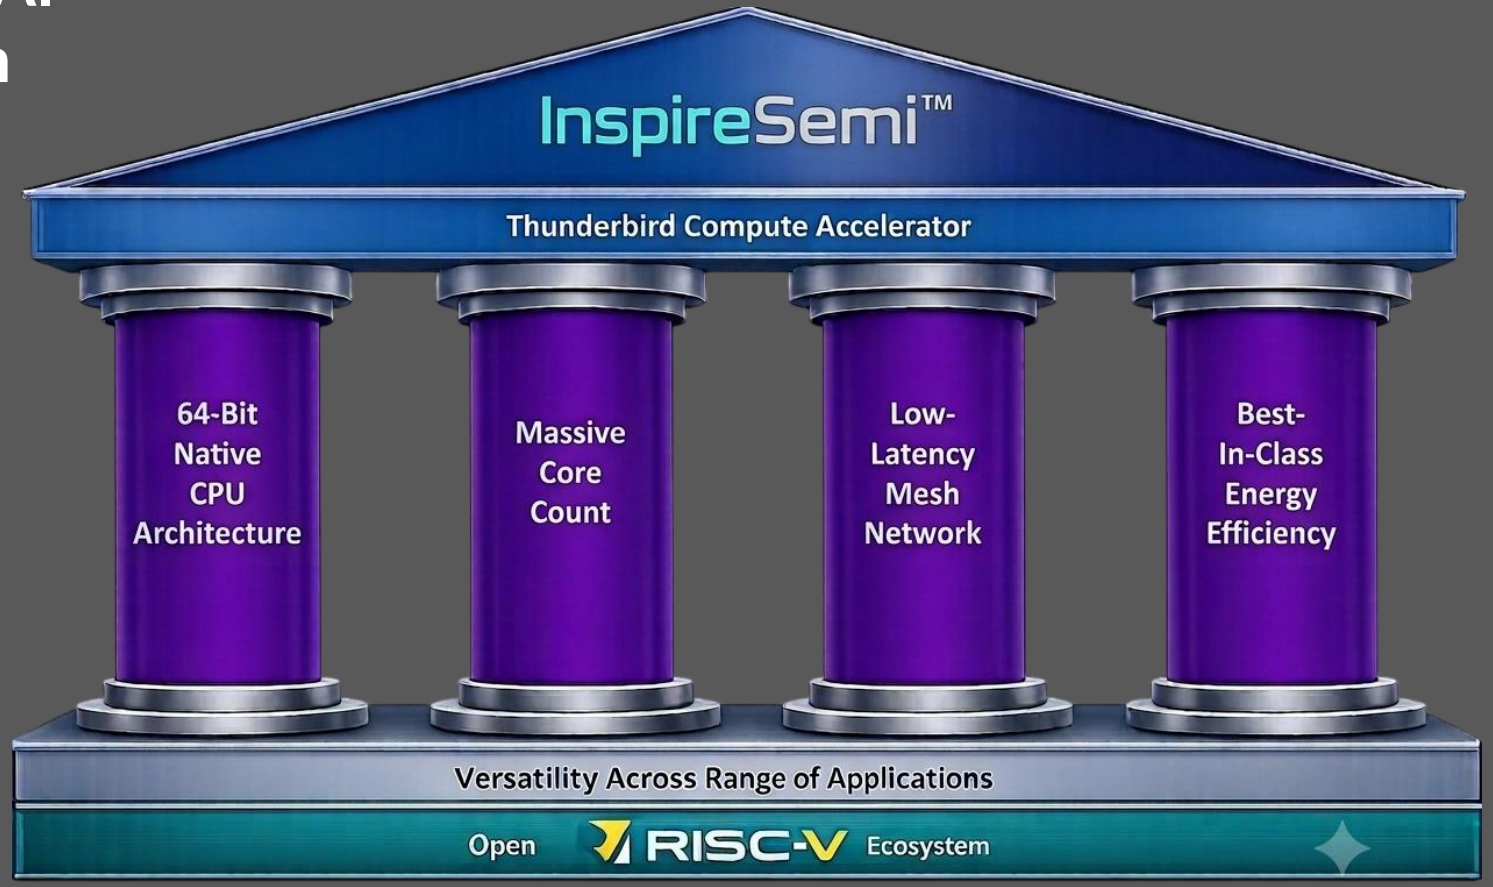

Company Introduction

July 2026

InspireSemi at a glance

- Fabless semiconductor startup based in Austin, TX
- Experienced, accomplished leadership team
- World class partners: design, supply chain, go to market
- “Bubble-proof” AI & HPC investment, targets unmet needs of multiple attractive high growth markets
  - Ideal for large scale deployment and orchestration of Agentic AI workloads
  - Initial beachhead in HPC: purpose-built platform for the $30B HPC accelerated computing market
- Best in class energy efficiency
- Thunderbird “Supercomputer cluster-on-a-chip” achieved functional first silicon
- Recognized global partners to deliver turnkey solutions across multiple markets and geographies

tsmc

RISC-V®

Lenovo

PENGUIN COMPUTING
A subsidiary of SMART Global Holdings, Inc.

World Wide Technology

POWERED BY DELL Technologies

SUPERMICRO

E4 COMPUTER ENGINEERING

Revolutionizing HPC-AI datacenters

One Thunderbird 6K board has more CPU cores than an ***entire rack of standard Intel, AMD, or ARM servers!***

- Less cost: datacenter footprint, servers, networking, cables, power, cooling
- Less complexity, interconnects, points of failure
- No need to build new datacenters

vs.

Ideal for large scale deployment and orchestration of Agentic AI workloads

Addresses the need to accelerate All HPC-AI software

What customers always wanted…not "yet another GPU"

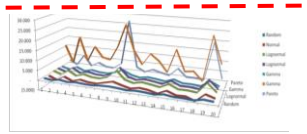
Financial simulations

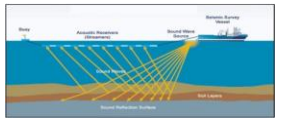
Geology: Seismic

Datacenter GPUs primarily focused on AI

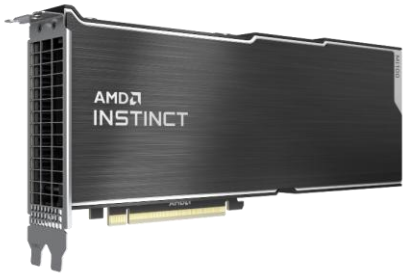

Agentic AI Orchestration


Financial Trading & Graph Analytics

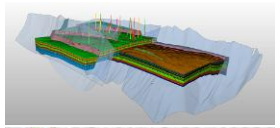
Energy: Reservoir Modeling & Sim

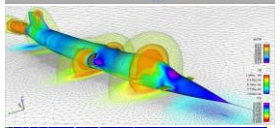
CAE/Computational Fluid Dynamics

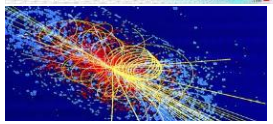
Government Lab Simulations

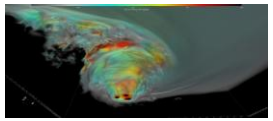
Climate & Weather Modeling


Cybersecurity & Defense

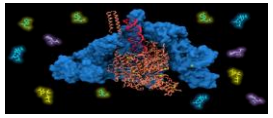
Genomics, Pharma, Life Sciences

InspireSemi Thunderbird

Versatile 4 chip platform delivers unprecedented capability

- ***Massive CPU core count: 6,144/board (FP64)***
- ***High performance: 24 TFLOPS (FP64)***
- ***Low power: 75 GFLOPS/watt (FP64)***
- ***Low latency: 10x faster***
- ***Determistic + Predictable Performance***
- ***Value: exceptional performance/$***

Providing the “ground truth” for AI in science & engineering

Accelerating the datacenter HPC-AI market

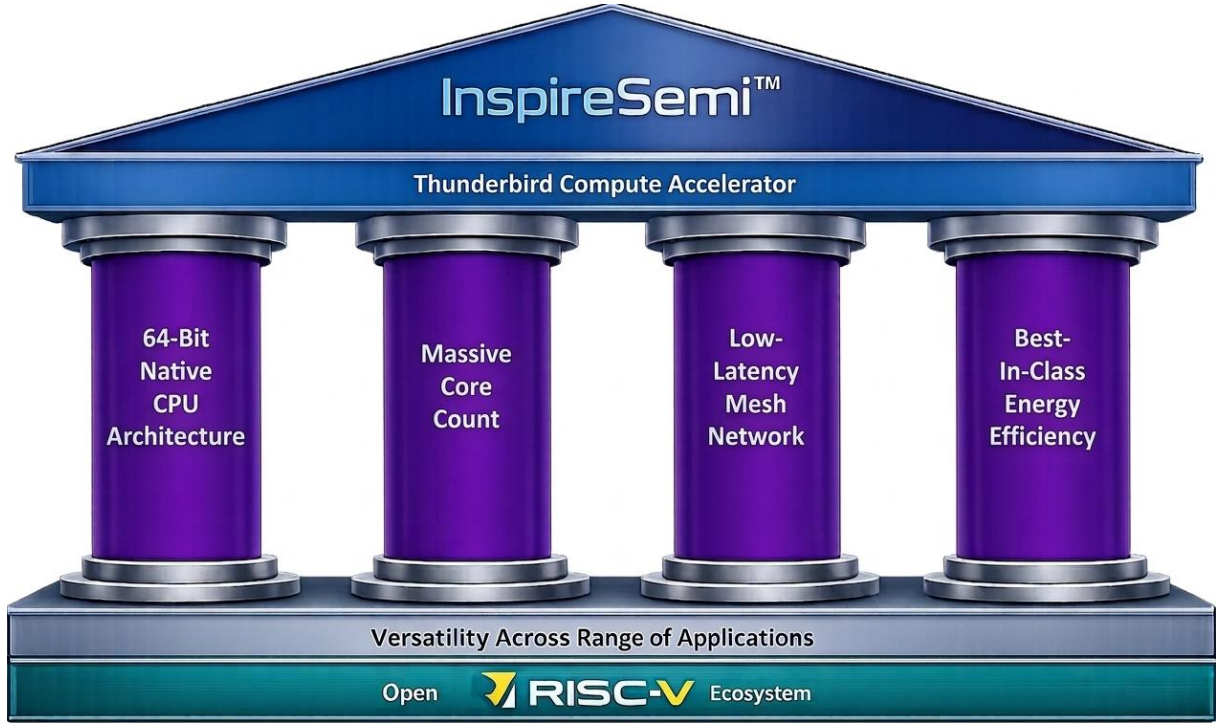

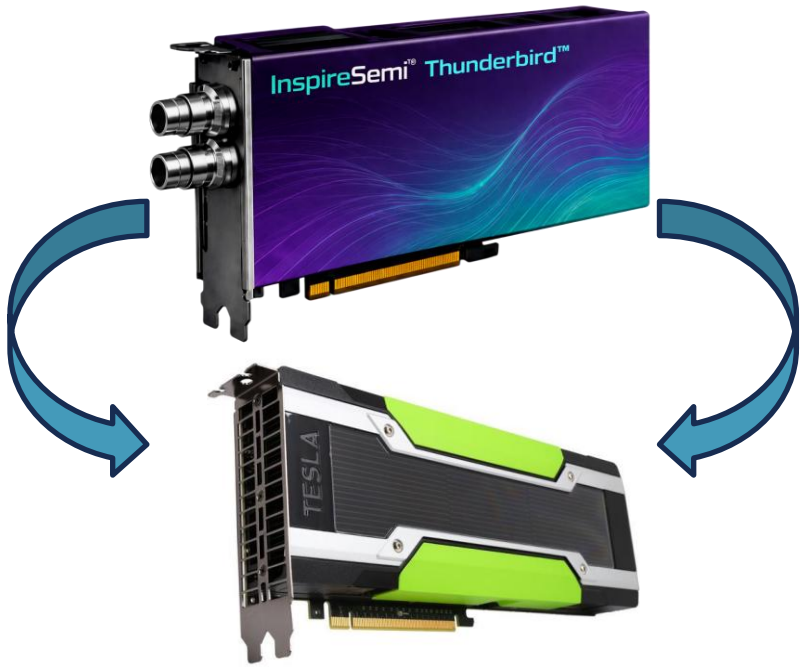

- Versatile all-CPU architecture applicable to all HPC-AI software, much of which does not benefit from GPUs
- High precision, high performance, low power 64-bit native processors to solve “big math” problems required by most HPC software

- 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

Sample Thunderbird customer & analyst feedback

BERKELEY LAB

"High-precision calculations are vital for generating reliable scientific data, which serves as the foundation for large language models. These include biomedical research, drug design, medical devices, climate change research, and applications that require deep simulation and modeling."
- Kathy Yelick, Director of Lawrence Berkeley National Laboratory (LBNL)

Argonne NATIONAL LABORATORY

"Double precision calculations are crucial for AI in science and engineering to ensure AI surrogate models are accurate and not hallucinating. The need for physics-based models is not going away soon."
- Charles Edward Catlett, Senior Computer Scientist

World Wide Technology

"With the momentum of AI and the convergence of AI and HPC, it is time to look outside the status quo and leverage a new technology base, like InspireSemi's Thunderbird product line. Thunderbird is ideal for workflows that require the highest performance and lowest power. It is easy to integrate, making it a valuable addition to the HPC and AI industry."
- Earl J. Dodd, Global HPC Business Practice Leader

Intersect360 RESEARCH

"The pursuit of AI has been a tremendous boon for high-performance architectures across the board. For pure AI investments, most of the attention is on GPUs, but many organizations are seeking a more versatile solution, built on processing elements that are suited to a variety of HPC, AI, and analytics workloads. This is where we see a market opportunity for companies like InspireSemi with its Thunderbird platform."
- Addison Snell, CEO

InspireSemi Thunderbird in the News

techradar.

Supercomputer-on-a-chip goes live: single PCIe card packs more than 6,000 RISC-V cores, with the ability to scale to more than 360,000 cores

- InspireSemi has announced the successful tapeout of the Thunderbird I Accelerated Computing chip for fabrication at TSMC

tom's HARDWARE

Thunderbird packs up to 6,144 CPU cores into a single AI accelerator and scales up to 360,000 cores — InspireSemi's RISC-V 'supercomputer-cluster-on-a-chip' touts higher performance than Nvidia GPUs

- The Holy Grail of supercomputing chip design is an architecture that combines the versatility and programmability of CPUs with the explicit parallelism of GPUs, and InspireSemi strives to achieve just that

TECHSPOT

Move over GPUs, with 1,536 cores the Thunderbird RISC-V CPU is ready to eat your lunch

- Open source enables small industries to participate in the accelerator boom

ALL ABOUT CIRCUITS

InspireSemi Announces Tapeout of Thunderbird Accelerated Computing Chip

- InspireSemi's new compute chip couples the parallel processing of GPUs with the versatility of CPUs

JPR Jon Peddie Research

InspireSemi announces tape-out of RISC-V HPC chip

- InspireSemi's Thunderbird I RISC-V chip offers high-performance computing for underserved applications, with an emphasis on energy efficiency and competitive pricing

Open software ecosystem solves customer porting challenges

- Supports open standard software, eliminates vendor lock-in (unlike competing compute accelerators)
  - **Linux!** provides access to existing HPC-AI software ecosystem
  - Plus Zephyr and FreeRTOS (Real-Time Operating Systems)
  - Eliminates need to learn and use proprietary software stacks
- Also supports OpenMP offload model
- Uses standard CPU-style programming models
  - No need for CUDA, ROCM, etc. that GPUs require
  - No need for disruptive software algorithm rewrites
  - Standard compiler, OpenMP, MPI, etc. approaches
- Leverages key RISC-V software ecosystem
  - Address-accurate QEMU model
  - Standard compilers, e.g.- GCC, Gfortran, GDB toolchains
  - Standard HPC libraries, e.g. – BLAS, LAPACK, FFTW

World class supply chain partners

- TSMC - Wafer fab
  - World's largest semiconductor foundry
  - Developing the most advanced process nodes
  - Secured 12nm wafer capacity for 2026+
  - Secured 3nm wafer capacity for 2027+

- ASE – Chip package & test
  - Worlds largest and highest quality OSAT (Outsourced Semiconductor Assembly & Test)
  - Leading edge package design

- Imec: Value Chain Aggregator (VCA)
  - Enable early access to tier-1 supply chain
  - Support engineering and early-prod volumes

Thunderbird addresses *all* AI-HPC customer needs

Disrupting on everything that matters

	InspireSemi Thunderbird	CPU - Servers	GPU Accelerators	AI Accelerators	FPGA
Applicability / Versatility	Applicable to all HPC-AI applications since CPU based	Applicable to all HPC-AI applications since CPU based	Applicable to most AI applications but only some HPC apps	Applicable only to AI applications	Very limited, for prototyping and special cases
Performance	High for broad range of HPC apps	Slow, need h/w accelerators	High for AI and some HPC apps	High for AI only	Medium
Energy consumption	Low ~80W/chip (150W max)	High 650W+/chip (+ many servers)	Very High ~1450W+	High ~600W to Very High 20kW	Med ~300W
Scalability	256 chips	1-4 chips/server	2-8 chips	Varies	1 chip
Cost & Complexity	Low: one PCIe card vs. full rack of connected servers	High: many servers, InfiniBand, switches, cables, cooling …	Med to High: depends on vendor, form factor	Med to $Millions: depends on vendor, form factor	Med $8K-$10K
Programming model	Standard CPU-like, Any language, Full instruction set	Standard CPU, Any language, Full instruction set	Specialized C variant (CUDA, ROCM, SYCL)	Proprietary, obscure	Hardware description language (Verilog)
Software ecosystem	Open-source, Linux, compilers, libraries, AI frameworks, existing applications	Robust	Limited, proprietary	AI frameworks and proprietary software stacks	None

Accomplished Leadership Team

Alexander Gray, Founder, President & CTO

- 20 years experience in tech startups, entrepreneurship
- CryptoCore, SolarBridge, SunPower
- Holds 9 patents
- BSEE, University of Illinois at Urbana-Champaign

James Hickman, Chairman & CEO

- 25+ year extensive business and investment career
- Has been CEO of multiple startups
- Was investor #1 in InspireSemi
- Former US Army intelligence officer
- BS Mathematics & Computer Science, US Military Academy at West Point
- MS Accounting & Information Management, University of Texas at Dallas

Jack Cartwright, CFO

- 20+ years experience in tech startups and public companies
- Datum, Spruce Services, BT Global, Atkins Global, Cisco
- Former US Army Logistics & Operations officer
- MBA, University of Texas; MS Accounting, University of Miami

Thomas Fedorko, COO

- 35+ years hands-on technical and business leadership in semiconductor operations in both large IDM and startups
- Eta Compute, Uhnder, Bluetechnix, Black Sand (Qualcomm), Luminary Micro (TI), Oak Technology, Motorola SPS
- Technical degree from DeVry University and graduate of the Motorola Management Institute

Doug Norton, CMO

- 35+ years experience; enterprise, startups, Federal
- Viviota, Nimbix, TMT, Virtana, Newisys, CoWare (Synopsys), Cadence, IBM
- President of The HPC-AI Society, Mentor at UT ATI (Austin Technology Incubator at University of Texas)
- RISC-V International: member HPC-SIG & Marketing team
- BSEE, Missouri University of Science & Technology

World class team driving breakthrough solutions

- Delivering breakthrough computational power and radically enhanced efficiency to solve real world problems
- Proven Thunderbird “supercomputer cluster-on-a-chip” architecture well positioned to attack multiple high growth market verticals
- Ideal for large scale deployment and orchestration of Agentic AI workloads
- Versatile, high-precision, low-power accelerated computing platform provides the “ground truth” the AI industry needs to solve real problems
- Recognized global partners to deliver turnkey solutions across multiple markets and geographies

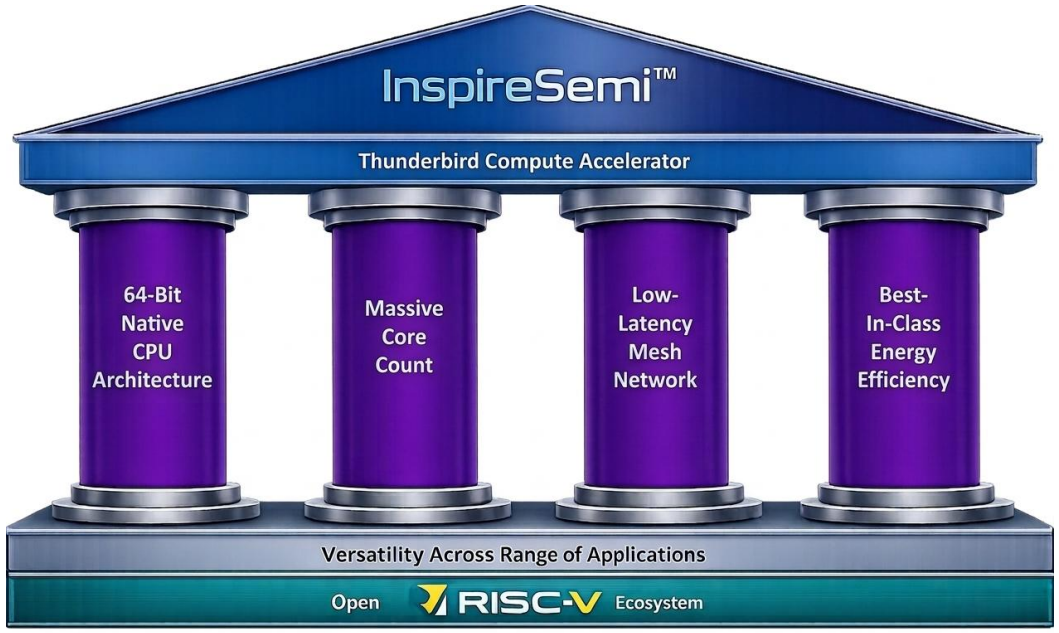

Lenovo
PENGUIN COMPUTING
A subsidiary of SMART Global Holdings, Inc.
World Wide Technology
POWERED BY DELL Technologies
SUPERMICRO
E4 COMPUTER ENGINEERING

Disclaimer

This presentation contains statements which constitute “forward-looking information” within the meaning of applicable securities laws, including statements regarding the plans, intentions, beliefs and current expectations of InspireSemi with respect to future business activities and operating performance.

Often, but not always, forward-looking information can be identified by the use of words such as “plans”, “expects”, “is expected”, “budget”, “scheduled”, “estimates”, “forecasts”, “intends”, “anticipates”, or “believes” or variations (including negative variations) of such words and phrases, or statements formed in the future tense or indicating that certain actions, events or results “may”, “could”, “would”, “might” or “will” (or other variations of the forgoing) be taken, occur, be achieved, or come to pass. Forward-looking information includes, but is not limited to, information regarding: (i) the business plans and expectations of the Company including expectations with respect to production and development; and (ii) expectations for other economic, business, and/or competitive factors. Forward-looking information is based on currently available competitive, financial and economic data and operating plans, strategies or beliefs as of the date of this presentation, but involve known and unknown risks, uncertainties, assumptions and other factors that may cause the actual results, performance or achievements of InspireSemi, to be materially different from any future results, performance or achievements expressed or implied by the forward-looking information. Such factors may be based on information currently available to InspireSemi, including information obtained from third-party industry analysts and other third-party sources, and are based on management’s current expectations or beliefs. Any and all forward-looking information contained in this presentation is expressly qualified by this cautionary statement.

Investors are cautioned that forward-looking information is not based on historical facts but instead reflect InspireSemi’s management’s expectations, estimates or projections concerning future results or events based on the opinions, assumptions and estimates of management considered reasonable at the date the statements are made. Forward-looking information reflects InspireSemi’s current beliefs and is based on information currently available to it and on assumptions it believes to be not unreasonable in light of all of the circumstances. In some instances, material factors or assumptions are discussed in this presentation in connection with statements containing forward-looking information. Such material factors and assumptions include, but are not limited to: the impact of the COVID-19 pandemic on the Transaction or the company; the ongoing conflict between Russia and Ukraine and any actions taken by other countries in response thereto, such as sanctions or export controls; and anticipated and unanticipated costs and other factors referenced in this presentation and the Filing Statement, including, but not limited to, those set forth in the Filing Statement under the caption “Risk Factors”. Although the Company has attempted to identify important factors that could cause actual actions, events or results to differ materially from those described in forward-looking information, there may be other factors that cause actions, events or results to differ from those anticipated, estimated or intended. Forward-looking information contained herein is made as of the date of this presentation and, other than as required by law, the Company disclaims any obligation to update any forward-looking information, whether as a result of new information, future events or results or otherwise. There can be no assurance that forward-looking information will prove to be accurate, as actual results and future events could differ materially from those anticipated in such statements. Accordingly, readers should not place undue reliance on forward-looking information. Should one or more of these risks or uncertainties materialize, or should assumptions underlying the forward-looking information prove incorrect, actual results may vary materially from those described herein as intended, planned, anticipated, believed, estimated or expected.