X-Silicon

Real-Time Compute Built for the Physical World

One processing platform for persistent autonomy.

Raptor™ for drones and unmanned systems. Kinetic™ for robotics.

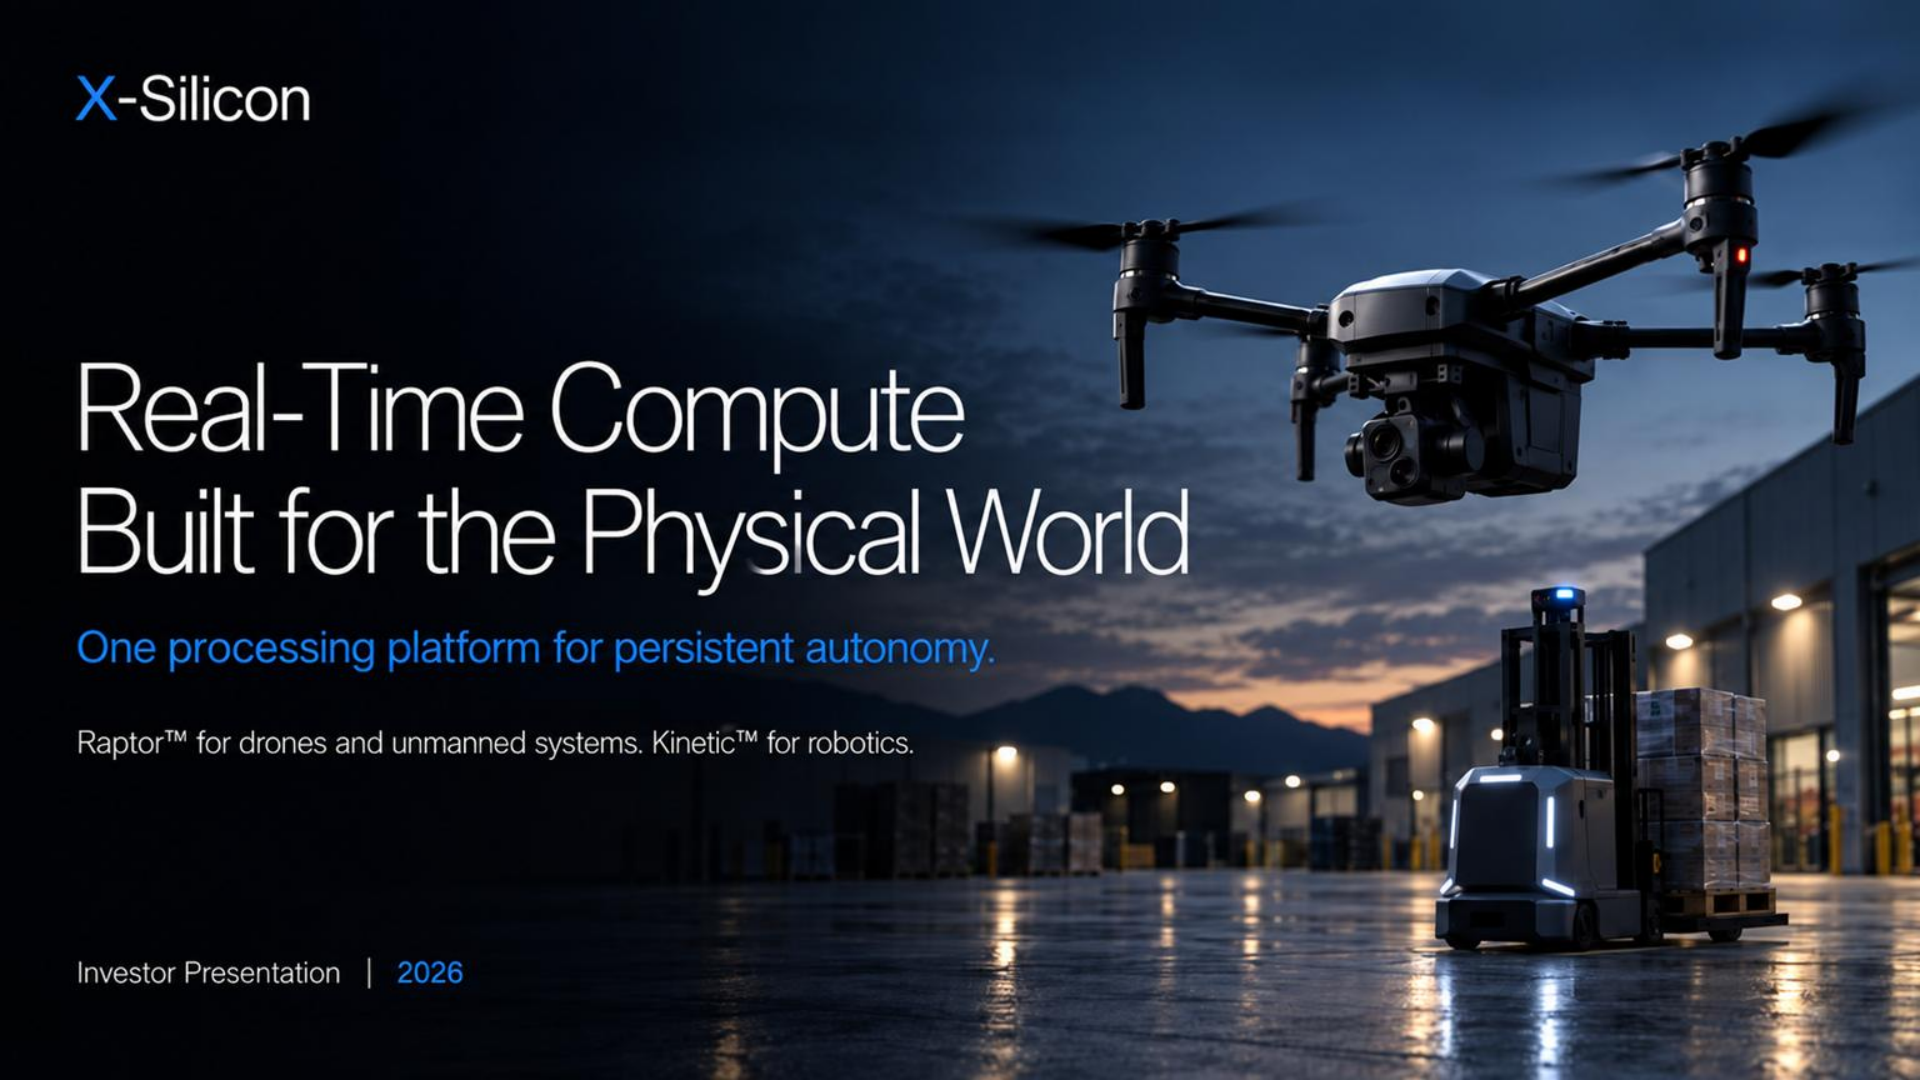

Investor Presentation | 2026

X-Silicon

Autonomy Is Becoming a Major New Compute Market

Drones and robots are moving from automated systems to AI-native autonomous machines — creating a new platform opportunity for onboard compute.

DRONES & UNMANNED SYSTEMS

Real-time autonomy in constrained environments

PHYSICAL AI

AI is moving from feature to system

ROBOTICS

From fixed automation to intelligent machines

$220B PHYSICAL AI REFERENCE TAM

2028–2032 cumulative semiconductor opportunity across autonomous machines

THE OEM PROBLEM

More Autonomy Is Creating a Compute Problem

As onboard intelligence grows, OEMs are being pushed toward more compute, more memory, and more power.

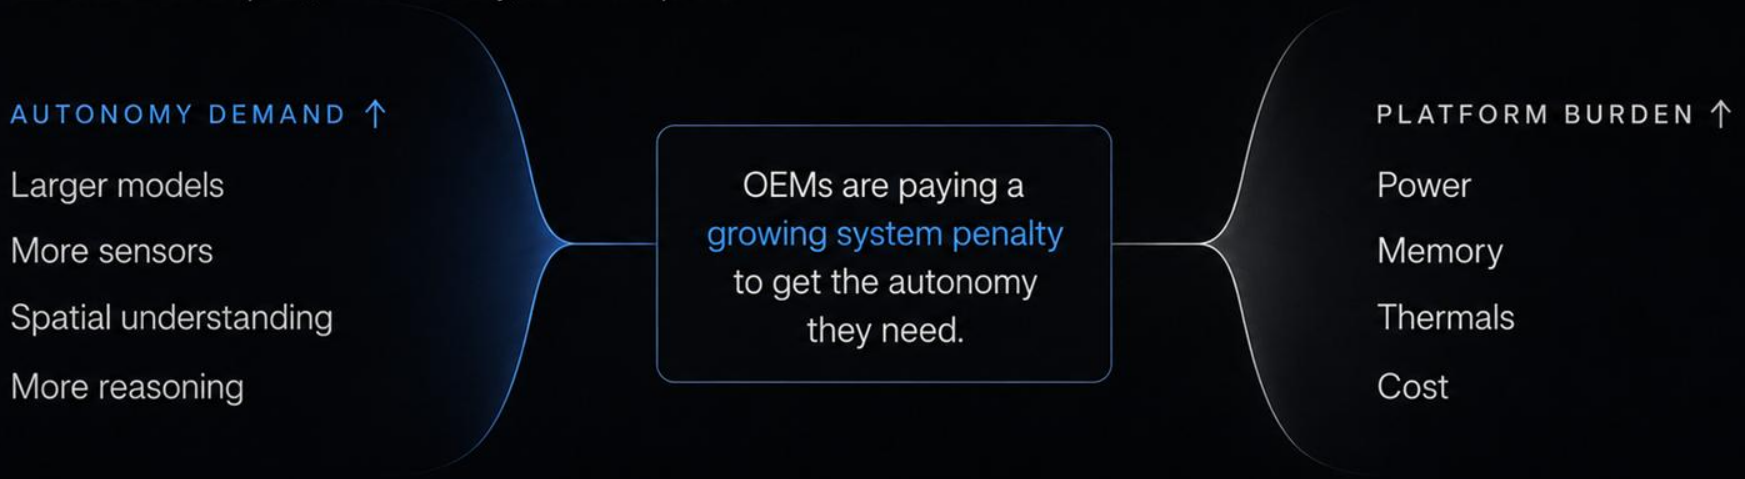

WHAT HAS TO CHANGE

Autonomous Machines Need a New Compute Platform

The next generation must increase performance per watt while overcoming memory-capacity and data-movement bottlenecks — without increasing system cost.

DO MORE ONBOARD
Increase AI performance per watt

UNDERSTAND THE WORLD
Continuous spatial and environmental awareness

MAINTAIN CONTEXT
Keep working state close without repeated memory movement

STAY INSIDE THE ENVELOPE
Power · Thermals · Size · Cost

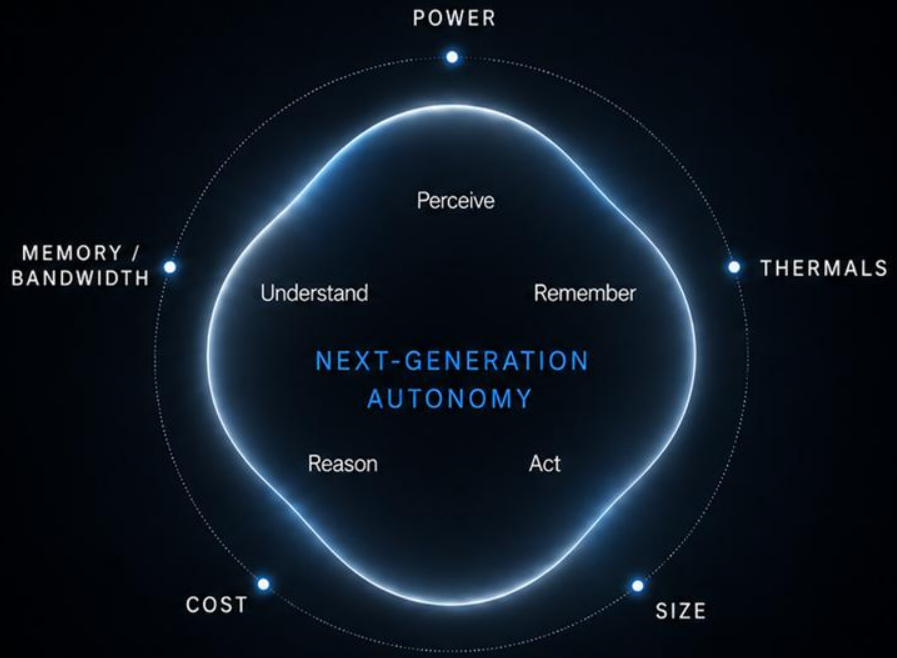

More autonomy per watt. Less memory movement.

WHY X-SILICON FOR PHYSICAL AI

Physical AI Needs More Than an AI Accelerator

X-SILICON
PHYSICAL INTELLIGENCE IN REAL TIME

Current platforms optimize isolated functions. X-Silicon is architected to coordinate inference, spatial context, control and working state as one real-time system.

TODAY: HETEROGENEOUS PHYSICAL AI PLATFORMS

Conventional heterogeneous architecture pattern

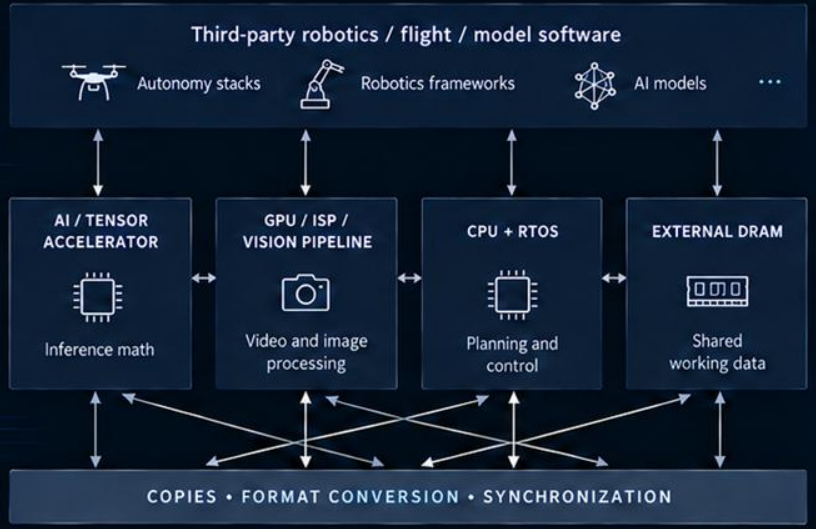

Multiple execution domains

State repeatedly moved

Integration burden shifts to the customer

ARCHITECTURAL SHIFT

X-SILICON: PHYSICAL AI-NATIVE PLATFORM

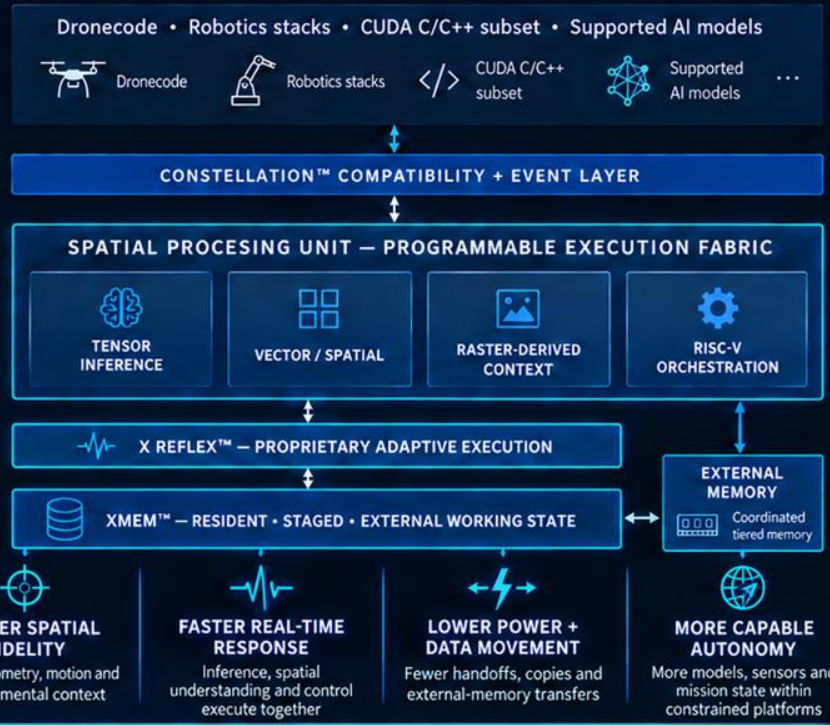

HIGHER-QUALITY PHYSICAL INTELLIGENCE—DELIVERED MORE EFFICIENTLY IN REAL TIME.

X-Silicon creates a richer understanding of the physical world while reducing latency, power and unnecessary data movement.

Target architecture; performance and implementation benefits remain subject to FPGA and silicon validation.

THE X-SILICON PLATFORM

Built for AI-Native Autonomous Machines

X-Silicon combines heterogeneous acceleration, persistent working context, and runtime orchestration to execute the autonomous workload as one coordinated platform.

XMEM — PERSISTENT CONTEXT

Keeps high-value working state close to execution.

SPATIAL PROCESSING UNIT (SPU)

PURPOSE-BUILT AUTONOMY ACCELERATOR

RISC-V CONTROL
System control & real-time coordination

TENSOR ENGINE
AI models & inference

VECTOR / SPATIAL
Perception, estimation & mathematics

GPU-DERIVED RASTER
Graphics, mapping & rendering

Runs AI models, perception, spatial understanding, estimation, graphics, and real-time autonomy workloads.

CONSTELLATION — ORCHESTRATION

Schedules and coordinates workloads across the processing fabric.

X-SILICON PHYSICAL-AI PLATFORM

Raptor™
Drones & Unmanned Systems
FIRST COMMERCIALIZATION PLATFORM

Kinetic™
Robotics
SECOND GROWTH PLATFORM

One heterogeneous processing fabric for the autonomous workload — not another isolated compute engine.

SPATIAL PROCESSING

X-SILICON

What Is a Spatial Processing Unit?

A Spatial Processing Unit (SPU) is a purpose-built heterogeneous accelerator that consolidates the compute required to perceive, understand, reason, and act in the physical world.

WHERE AM I?
Determine position, orientation and movement.

WHAT SURROUNDS ME?
Interpret objects, terrain, people and free space.

HOW IS IT CHANGING?
Track motion, conditions and evolving context.

WHAT SHOULD I DO NEXT?
Apply AI and real-time control to decide and act.

Spatial awareness establishes the machine's understanding of the physical world. Autonomous intelligence turns that understanding into action.

WHAT IT CONSOLIDATES

AI acceleration • Vector / spatial processing • GPU-derived raster • Real-time compute control

Replaces a stitched autonomy-compute stack — not every processor in the system.

FIRST COMMERCIAL PLATFORM

Raptor™

Built for Drones & Unmanned Systems

Raptor brings the X-Silicon Physical AI platform into a drone-optimized implementation—designed for systems where power, endurance, real-time response, and onboard autonomy are mission-critical.

MORE AUTONOMY ONBOARD

Perception, spatial understanding, navigation, and AI decision-making closer to the mission.

MORE MISSION PER WATT

Increase autonomous capability within constrained power and thermal envelopes.

GREATER MISSION INDEPENDENCE

Support continued operation through degraded connectivity, limited external positioning, and changing environments.

More intelligence. **More** endurance. **More** mission capability.

Autonomy Must Continue When External Guidance Does Not

When positioning, communications, or remote guidance are degraded, the machine must use onboard sensors and retained context to determine where it is—and what to do next.

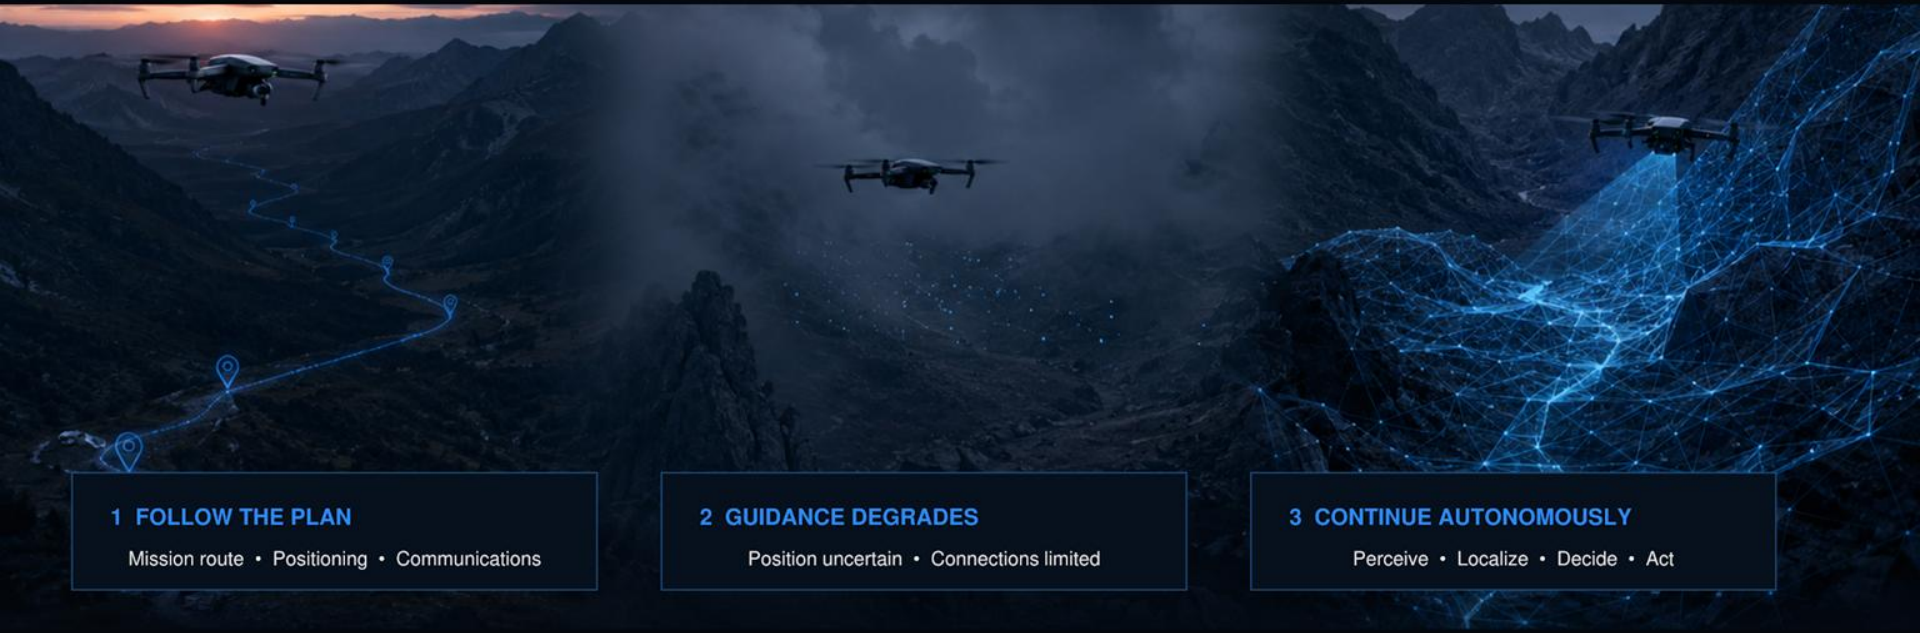

The machine must understand where it is, what surrounds it, and what to do next.

RAPTOR TARGET MARKETS

Focused on SWaP-constrained drones that require more onboard autonomy.

MORE ONBOARD INTELLIGENCE WITHIN THE MISSION ENVELOPE

X-SILICON

1 TACTICAL & DEFENSE

ISR, FPV, loitering and autonomous mission systems

MISSION AUTONOMY

2 PUBLIC SAFETY & DFR

Police, fire, emergency response and search-and-rescue

PERSISTENT RESPONSE

3 INDUSTRIAL INSPECTION

Utilities, energy, mining, rail and infrastructure

LONGER MISSIONS

4 AGRICULTURE

Crop intelligence, spraying, mapping and multispectral sensing

SENSING ENDURANCE

5 COMMERCIAL & ENTERPRISE

Surveying, mapping, logistics and specialized aerial systems

SCALABLE DEPLOYMENT

RAPTOR SWEET SPOT

Small-to-mid-size autonomous drones where every watt and gram competes with flight time, payload, and mission endurance.

1 – 25 KG
Typical Platform Class

5W – 25W
Compute Power Envelope

30+ MIN
Mission Endurance Target

MORE AUTONOMY PER WATT, GRAM, AND SYSTEM DOLLAR.

WHY THE ARCHITECTURE MATTERS

Built for Real-World Autonomy Results

Raptor's SPU, XMEM and Constellation architecture reduces the fundamental costs of autonomy—data movement, memory latency and orchestration overhead—so more work gets done onboard with less power.

ARCHITECTURAL ADVANTAGE

Less Data Movement

Keep working context close in XMEM.
Minimize off-chip traffic.

Lower Energy per Task

Higher compute density and efficient memory architecture.

Lower & More Predictable Latency

On-chip state + distributed scheduling for real-time responsiveness.

More Compute

The SPU integrates tensor, vector/spatial, RISC-V and modified GPU/raster processing on one SoC.

Simpler Integration

Constellation coordinates workloads and abstracts the underlying hardware complexity.

MORE AUTONOMY ONBOARD. BETTER SYSTEM OUTCOMES.

MEASURED RESULTS ON RAPTOR GEN-1 PROTOTYPE*

Energy per Inference	External Memory Bandwidth	End-to-End Latency	Sustained AI Throughput/Watt	Compute Consolidation
3.2x	**5.6x**	**2.8x**	**4.1x**	**~3–5x**
Lower	Lower	Lower	Higher	Fewer chips
vs. discrete GPU solution	vs. GPU solution	at iso-accuracy	vs. GPU solution	vs. typical autonomy compute stack

* Internal lab measurements on representative autonomy workloads. Results will vary by use case and system.

IMPACT WHERE IT MATTERS

More Missions Completed
Longer endurance and higher mission success.

Lower Power
More capability within tight power envelopes.

Lower System Cost
Fewer chips, less memory, simpler integration.

Thermal Efficiency
Less heat for the same compute and workload.

Higher Reliability
Fewer components and simpler systems.

THE ARCHITECTURAL DIFFERENTIATOR

Raptor is not just more compute.
It's a different way to architect autonomy.

SECOND GROWTH PLATFORM

Kinetic™

Physical-AI Compute for Robotics

The same X-Silicon foundation for robots that must perceive, understand, adapt, and work continuously in the physical world.

MORE INTELLIGENCE AT THE EDGE
Richer perception and autonomous workloads

KEEP CONTEXT ALIVE
Maintain environment, task, and state as the robot works

OFFLOAD → CONSOLIDATE
Start beside the host. Move more compute onto X-Silicon over time.

Keep the host. Move the Physical AI.

Material Handling → Field / Inspection → Mobile Manipulation → Advanced Robotics

KINETIC TARGET MARKETS

Focused on High-Value Autonomous Robotics

X-SILICON

Kinetic targets robotics systems where spatial intelligence, persistent context and compute consolidation create the greatest platform value.

1 WAREHOUSE / MATERIAL HANDLING

AMRs | Autonomous Forklifts
Pallet | Yard | Dock Automation

PRIORITY

Navigation, fleet coordination and persistent task state

2 INDUSTRIAL / FIELD ROBOTICS

Inspection | Mining | Energy
Harsh-Environment Robots

PRIORITY

Spatial awareness, sensor fusion
and continuous autonomy

3 MOBILE MANIPULATION

Picking | Machine Tending
Material Transfer | Mobile Manipulators

PRIORITY

Perception, geometry, task state and motion coordination

4 SERVICE / COMMERCIAL ROBOTICS

Facilities | Delivery
Cleaning | Security

PRIORITY

Lower power and cost, persistent
context and scalable deployment

5 ADVANCED AUTONOMOUS ROBOTICS

Humanoids | Assistive Systems
Specialized Field Platforms

PRIORITY

Higher workload density, richer spatial state and consolidation

KINETIC SWEET SPOT

Autonomous robots where every watt, compute module, and memory movement adds cost, heat, integration complexity, or battery burden.

PHYSICAL-AI OFFLOAD
Move perception + spatial workloads off the host

PERSISTENT STATE
Keep pose, object, environment, and task context close to the compute

COMPUTE CONSOLIDATION
Companion platform today → primary robotics SPU over time

INITIAL FOCUS: Warehouse, industrial, field, and mobile-manipulation robotics.

These segments best match Kinetic's spatial-compute, persistent-state, and host-offload advantages.

KINETIC PLATFORM

Kinetic in the Autonomous Robotics Stack

X-SILICON

A Physical AI platform that integrates compute, persistent context, and runtime orchestration with the broader robotics ecosystem.

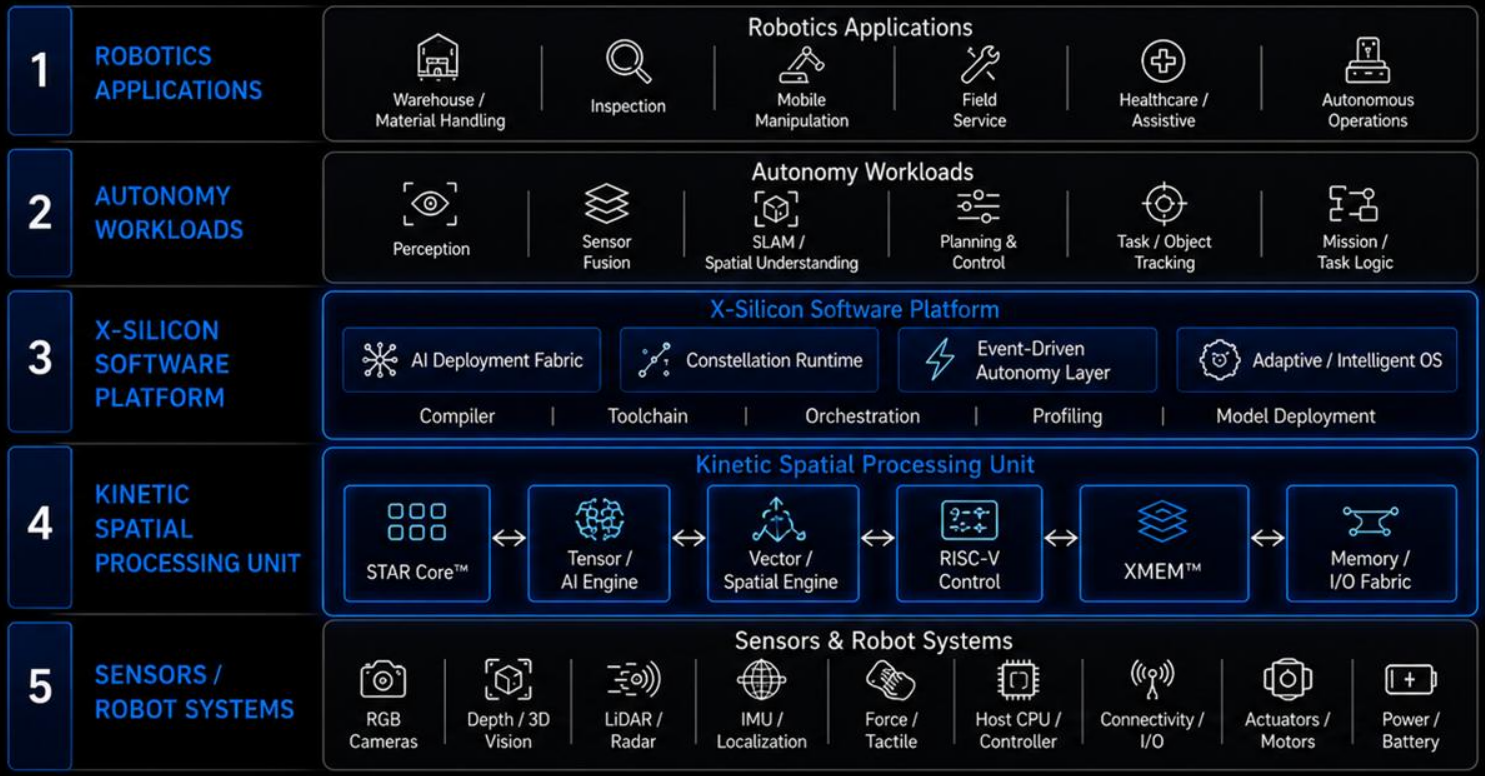

KEEP PHYSICAL AI AT THE EDGE

Execute critical autonomy workloads onboard.

MAINTAIN WORKING CONTEXT

Preserve task and environmental state with XMEM.

REDUCE DATA MOVEMENT

Keep compute and high-value state closer together.

INTEGRATE WITH ROBOTICS ECOSYSTEMS

Support established applications, autonomy software, sensors, and control systems.

One platform connecting robotics software to physical-world execution.

PRODUCT MOAT

The Moat Is the Platform—Not Just the Silicon.

X-Silicon converts protected silicon IP into complete, vertically focused platforms designed around how autonomous-machine OEMs build and deploy products.

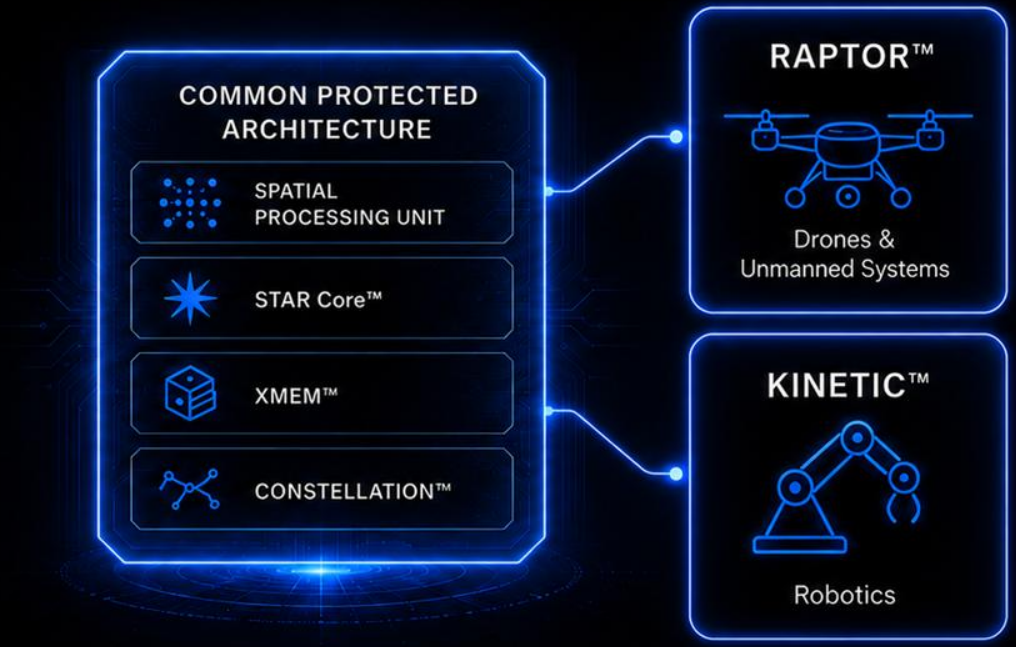

01 **PROTECTED COMPUTE FOUNDATION**
Spatial Processing Unit, STAR Core™ and XMEM™ provide the differentiated architectural foundation.

02 **PURPOSE-BUILT FORM FACTORS**
Modules and reference platforms are engineered around each market's power, thermal, size, sensor and deployment requirements.

03 **INTEGRATED SOFTWARE PLATFORM**
Constellation™, developer tools, production software and drivers make the architecture programmable and deployable.

04 **OEM SYSTEM INTEGRATION**
The platform integrates with sensors, control systems, autonomy software and customer applications.

05 **COMPOUNDING VERTICAL EXPERTISE**
OEM requirements, workload tuning and deployment learning strengthen every subsequent product generation.

The silicon creates the advantage. The complete platform creates the moat.

X-SILICON

MEET THE TEAM

Experienced Team Who Has Designed and Shipped Billions of Chips

Billions of units shipped into Tier-1 OEM production platforms.

Industry pioneers who have built, scaled, and commercialized the platforms that defined modern graphics, mobile, and AI compute.

Dan Nilsson
Founder & CEO
ex-Qualcomm • AMD/ATI • Thundercomm
Qualcomm AMD

Steven Turner
CTO
ex-Qualcomm • AMD • ATI
Qualcomm AMD

Dr. Atif Zafar
Chief Scientist
Co-founder — IRANE
Co-founder — Pixlia
Founding Member — RISC-V Graphics Extension Group

Larry Polizzotto
CFO/COO
ex-Dell • IBM • Honeywell
DELL IBM Honeywell

Colin McKellar
Technical Advisor
ex-Imagination • Sony
Imagination SONY

Daniel G. Nilsso
CTO, RealX
VP Engineering, X-Silic
Tieto • Ericsson
tieto ERICSSON

Advisory relationships support architecture validation, OEM engagement, ecosystem development, and go-to-market execution.

X-SILICON

COMPLETE STAR CORE MICROARCHITECTURE FPGA PROGRAM

SEED ROUND A | FPGA PROGRAM WINDOW: OCTOBER 2026–DECEMBER 2027

$3.8M total program • **$1.9M** Estonia/EU RUP • **$230K** ZPA/partner funding • **$1.67M** investor delta

PHASE 1

$800K — XMEM PROOF

FUNDING TARGET: OCT 2026 | ~5 MONTHS

Memory residency • XMEM-on/off traffic • lock/isolation • counters

Investors $320K

→

PHASE 2

$2.43M — XMEM + SPU PROOF

ADVANCED INTEGRATION | ~9 MONTHS

Tensor/vector/spatial execution • RISC-V orchestration • multi-context workloads

Investors $729K

→

PHASE 3

$633K — FULL SYSTEM-LEVEL PROOF

PROGRAM TARGET: DEC 2027

Integrated FPGA platform • compiler/runtime integration • Raptor/Kinetic workload demos

Investors $633K

Investor capital closes the funded-program gap while X-Silicon retains core architecture control and IP ownership.

PRODUCT DEFINITION

What Customers Actually Buy

X-SILICON

Customers buy deployable Raptor and Kinetic platforms—each built on the same SPU + XMEM + Constellation architecture.

RAPTOR™

Drones & Unmanned Systems

- SPU Silicon / Module
- Production Software
- Constellation Runtime
- SDK / Developer Tools
- Reference Platform
- Integration & Support

INITIAL DEPLOYMENT

Integrates with flight-control, autonomy and sensor ecosystems.

PLATFORM EXPANSION

Consolidates more onboard Physical AI and autonomy workloads over time.

KINETIC™

Robotics & Autonomous Machines

- SPU Silicon / Module
- Production Software
- Constellation Runtime
- SDK / Developer Tools
- Reference Platform
- Integration & Support

INITIAL DEPLOYMENT

Starts alongside the existing host and robotics software stack.

PLATFORM EXPANSION

Moves more spatial, perception and autonomy workloads onto Kinetic over time.

COMMON X-SILICON PLATFORM

SPU
programmable tensor, vector/spatial, RISC-V and GPU-derived execution

XMEM™
resident working-state and reduced data movement

CONSTELLATION™
runtime orchestration and workload portability

DEVELOPER TOOLS

REFERENCE SOFTWARE

One platform. Two products. Multiple revenue layers.

Silicon, modules, software, developer tools, integration and support across the customer lifecycle.

INVESTMENT THESIS

Why X-Silicon. Why Now.

X-SILICON

The convergence of autonomy, AI and onboard compute creates a generational Physical AI opportunity.

01

LARGE, GROWING PHYSICAL AI MARKET

Autonomous drones, unmanned systems and robots are scaling rapidly—creating demand for more capable intelligence inside the machine.

02

AI-NATIVE AUTONOMY SHIFT

Machines are moving from rule-based operation toward continuous perception, spatial understanding, reasoning and action.

03

ONBOARD AUTONOMY IS A REQUIREMENT

Connectivity, cloud access and positioning cannot always be assumed. Critical autonomy must operate locally and in real time.

04

DIFFERENTIATED PHYSICAL AI PLATFORM

The Spatial Processing Unit, XMEM™ and Constellation™ coordinate compute, working context and software for continuous autonomy.

05

RAPTOR FIRST, PROVEN NEED

Drones and unmanned systems provide the strongest first fit: severe SWaP constraints, onboard autonomy and immediate market pull.

06

KINETIC FOLLOWS, PLATFORM UPSIDE

Robotics extends the same architecture into a second high-value market without funding a separate clean-sheet platform.

07

VALIDATION MOVING FORWARD

Customer requirements, technical partners and staged FPGA proof progressively reduce engineering and market risk.

08

CAPITAL LEVERAGE ACCELERATES PROOF

European, sovereign and non-dilutive programs can extend investor capital and accelerate platform validation.

One Physical AI architecture. Raptor first. Kinetic follows. Proof-driven commercialization.

CLEAR FOCUS
Raptor first. Kinetic next.

ARCHITECTURAL MOAT
SPU. XMEM. Constellation.

MARKET VALIDATION
Requirements, partners and proof.

CAPITAL LEVERAGE
Strategic and non-dilutive programs.

THANK YOU

X-SILICON

Real-Time Compute Built for the Physical World

One processing platform for persistent autonomy.

Raptor™ | Kinetic™

Daniel Nilsson | Founder & CEO
dan@x-silicon.ai

Larry Polizzotto | COO & CFO
larryp@x-silicon.ai

x-silicon.com

PHYSICAL
REAL WORKLOADS
REAL WORLDS

THE PHYSICAL AI WORKLOAD GAP

AI Inference Is Only One Part of the Autonomy Loop

Physical AI must transform continuous sensor input into spatial context, decisions and control—within real-time power and latency constraints.

X-SILICON
COMPUTE THE PHYSICAL WORLD

SILICON FOR A MORE AUTONOMOUS WORLD

TODAY — CONVENTIONAL HETEROGENEOUS PHYSICAL AI PATTERN

Each engine is efficient at one task; the autonomy loop crosses multiple boundaries.

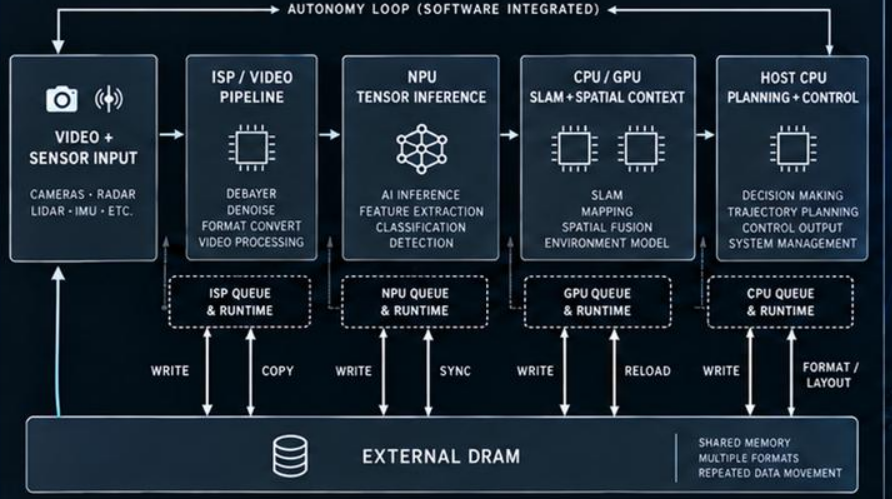

Separate engines • host coordination • repeated state movement • latency, power and integration tax

VS

X-SILICON — PHYSICAL AI EXECUTION PATH

The complete perception-to-action loop is coordinated across one programmable platform.

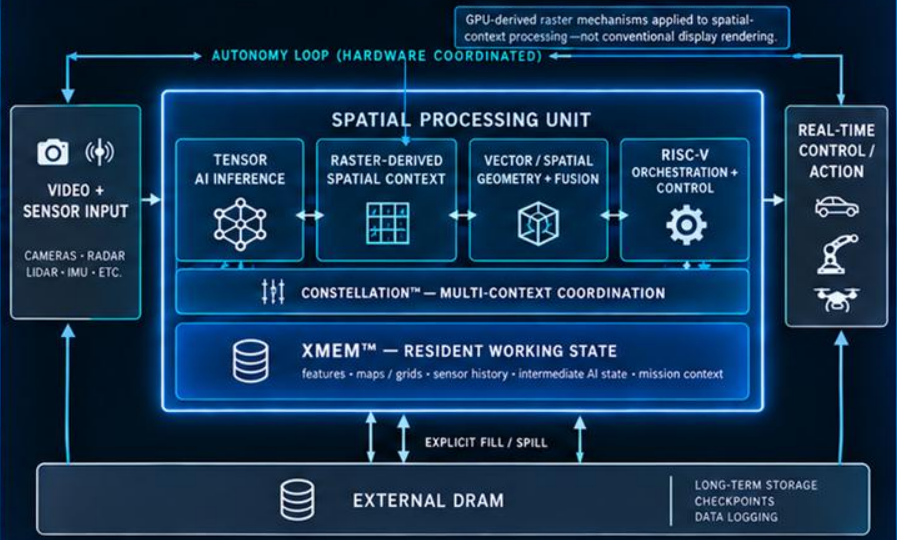

RUP PHASE 2 TARGET PROOF: inference + spatial context + multi-context execution + RISC-V orchestration

HIGHER SPATIAL FIDELITY
Richer geometry, motion, position and environmental context

FASTER CLOSED-LOOP RESPONSE
Perception, spatial understanding, planning and control are coordinated in real time

LOWER POWER + DATA MOVEMENT
Working state is retained and shared instead of repeatedly copied and reloaded

MORE CAPABLE AUTONOMY
More sensor streams, richer models and longer mission context remain available

THE PHYSICAL AI RESULT | A richer understanding of the real world, faster perception-to-action response and greater autonomy capability within constrained power and memory.

Target architecture and RUP Phase 2 proof objectives; benefits remain subject to FPGA and silicon validation.

BUSINESS MODEL

Multiple Revenue Layers. One Platform.

X-Silicon captures value across the customer lifecycle—from evaluation and design integration through production volume and installed-base support.

One design win can create multiple revenue streams over the life of the platform.

X-Silicon is a **product semiconductor and software** company—**not an IP-licensing business.**

Direct product model:
NRE → silicon/modules per **unit** → software and support.

The Team Behind The Team.

Industry pioneers who have built, scaled, and commercialized the platforms that defined modern graphics, mobile, and AI compute.

Dave Orton

Former CEO & SVP — ATI / AMD
Former CEO — Aprina Imaging

Graphics and compute industry leader with direct experience scaling GPU architectures to global OEM platforms.

Jon Peddie

Founder — Jon Peddie Research

Graphics industry pioneer and analyst. Forty years tracking every major compute architecture transition. Advisor and industry validator.

Eric Powers

Lead AI — DreamWorks Animation
Former Technical Director — Disney Feature Animation

Production AI and advanced graphics leader deploying AI-native rendering and content pipelines at global scale across film and animation.

Jim Edson

Former VP — ARM
Former VP — Qualcomm
Former VP — Motorola

Decades of silicon commercialization, business development, and ecosystem partnership experience across the industry's most consequential platforms.

X-SILICON

Advisory relationships support architecture validation, OEM engagement, ecosystem development, and go-to-market execution.