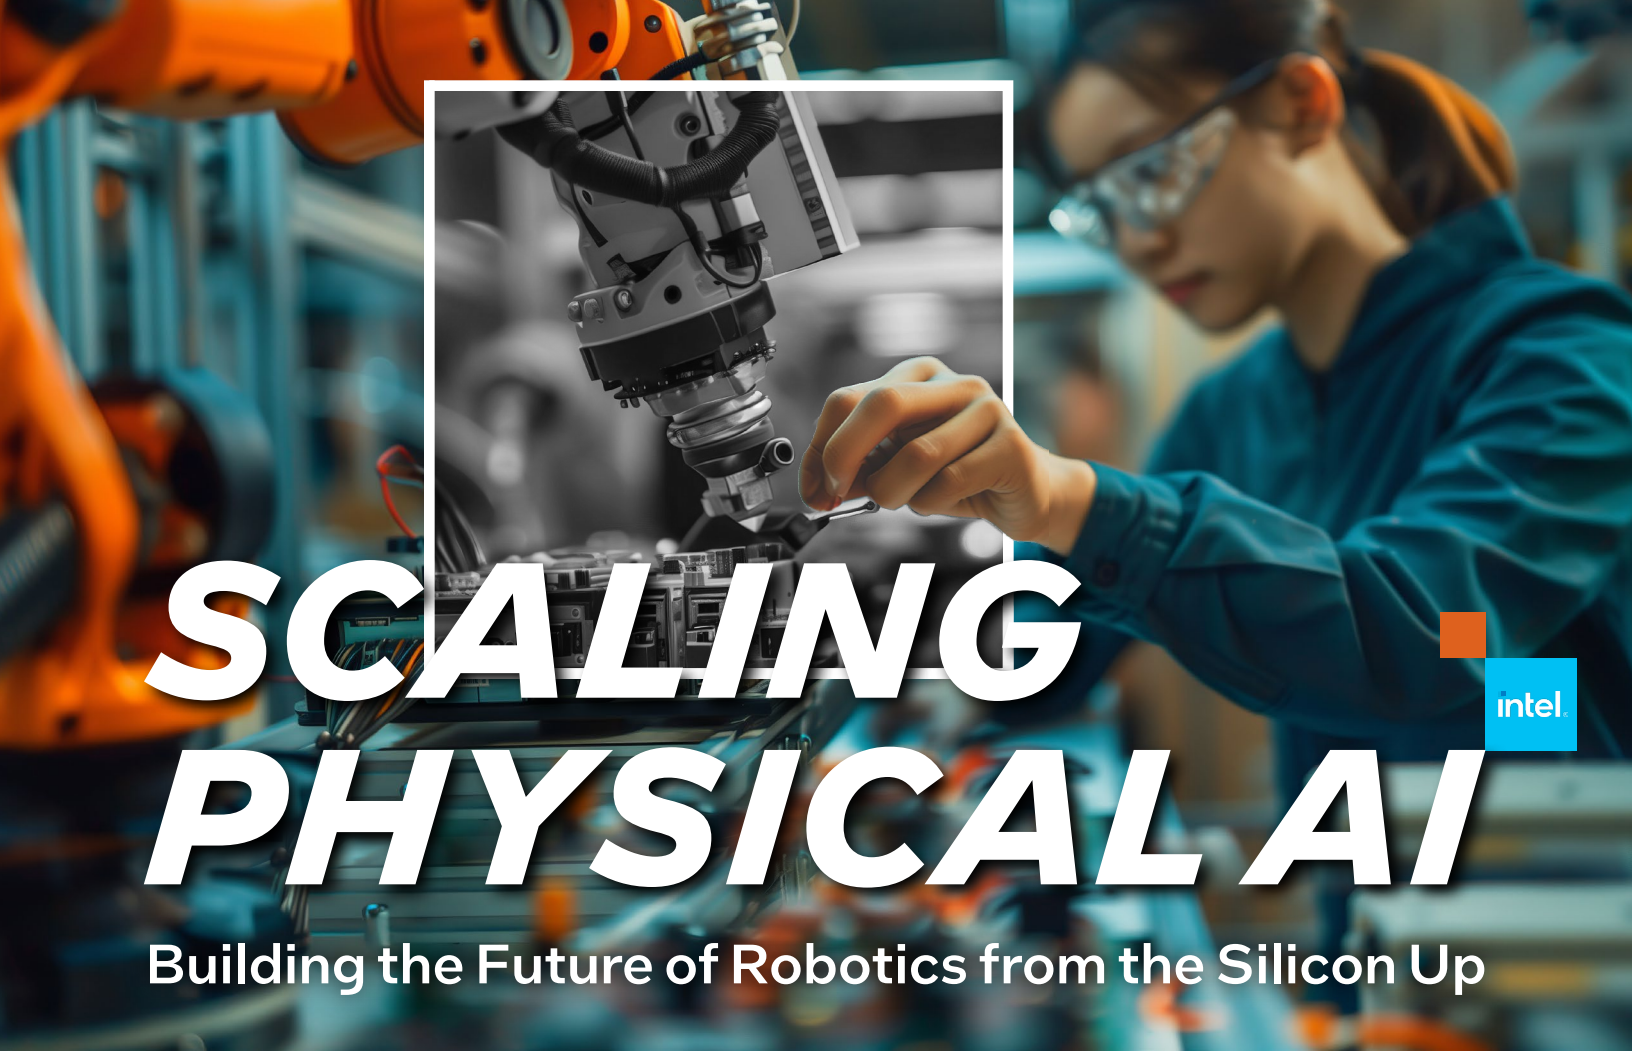

SCALING PHYSICAL AI

Building the Future of Robotics from the Silicon Up

Ricky Watts
Senior Director, Intel Industrial & Robotics Division

Rajesh Kurusetty
Director of Product, Robotics & Edge AI, Intel Industrial & Robotics Division

Roger Tang
APJ Industrial Segment Manager, Intel Industrial & Robotics Division

I. MARKET FORCES AND STRATEGIC SHIFTS:

The New Era of Robotics

We are living through a transformational era in robotics—a shift from traditional automation to intelligent systems that understand, adapt, and collaborate. Robotics is no longer just about repetitive tasks; it's about autonomy, perception, and real-time decision-making. At the center of this shift is Physical AI—the convergence of artificial intelligence with robotics in the physical world.

Reshoring, labor shortages, and an urgent need for flexibility are reshaping global manufacturing. The robotics market is on a steep trajectory, driven by demand for intelligent machines that can navigate complex environments, operate reliably around people, and continuously learn from data.

From autonomous mobile robots (AMRs) in warehouses to collaborative arms on factory floors, the industry is evolving quickly. But this evolution brings new requirements—performance, flexibility, reliability, and sustainability—and it demands more than just software innovation. It requires foundational rethinking of the hardware, architecture, and ecosystem that enable robotics at scale.

II. INTEL'S PLATFORM APPROACH:

Built from the Silicon Up

Intel is laying the groundwork for the next wave of robotics by doing what it does best - building from the silicon up. Our approach is rooted in the belief that intelligent machines need intelligent infrastructure: validated hardware, deterministic control, and scalable AI performance working in harmony.

This isn't just about launching products. It's about building a modular,

Intel's full stack portfolio is transforming and enabling new Intel + partner-led offerings to better solve customer problems at the edge.

future-ready platform customers can trust. From low-power edge compute to high-performance CPUs and integrated GPUs, Intel's silicon portfolio is optimized for robotics workloads—from vision and SLAM to real-time control and AI inferencing.

Openness and modularity are at the core. We support industry standards like ROS 2, OPC UA, and real-time Linux and provide validated building blocks through platforms like OpenVINO™, Geti, and Edge Controls for Industrial. This modular architecture simplifies integration, consolidates workloads, and accelerates time-to-value. In a hybrid world where cloud training and edge inference coexist, interoperability—not lock-in—wins.

III. WHAT THE OPEN EDGE PLATFORM ADDS:

From Building Blocks to Blueprints

While the principles above set the direction, Intel's Open Edge Platform turns strategy into practice with a layered, end-to-end approach designed specifically for edge AI and robotics. Think of it as a complete stack with clear hand-offs—from silicon through software, tools, and reference systems—so teams can design, validate, and scale with confidence.

1. A Layered Architecture That Meets You Where You Build

- Intel® Silicon → CPUs with integrated GPUs and NPUs for balanced AI, control, and media.
- Intel's Open Edge Platform → The integration plane: real-time OS support, orchestration hooks, SDKs, and security scaffolding.
- Intel® AI Edge Systems & Verified Reference Blueprints → Sized, verified, and benchmarked system recipes with recommended workloads and expected performance, spanning multiple form factors.
- Edge AI Suites & Libraries → Curated AI pipelines and building blocks (e.g., benchmarking tools, sample applications, SDKs) that map to vertical solutions in manufacturing, retail, media & entertainment, and metro/smart city.

2. Open Source Flexibility A permissible open source license with source code available on GitHub for customers to pick and choose the code they would like to experiment and/or take it to production.

3. Base Platform Support for Robotics & Real-Time Control Applications
For robotics, the platform emphasizes end-to-end determinism and performance: Linux LTS BSP with PREEMPT-RT, optimized EtherCAT drivers, camera drivers, oneAPI compilers/libraries, Intel® Real-Time Libraries, RT tuning tools, BIOS RT BKC, optional ACRN hypervisor, and ROS 2 middleware support. You can assemble a stack that couples sub-10 us control loops with modern AI inference and media pipelines on the same platform.

4. Benchmark-Ready, Right Out of the Box
The platform includes micro-level and end-to-end benchmarking for customer platform evaluations, so teams can quantify gains across media pipelines, AI inference, and control workloads before committing to deployment. Intel's materials highlight competitive, workload-level advantages for video analytics and performance-per-watt, with reference boards, expected performance targets, and recommended configurations documented to speed sizing and procurement.

5. Vertical Fit: Solutions Framed by Real Use Cases
Rather than treating AI as an abstract capability, the platform anchors solutions by vertical use cases—Manufacturing (quality/AOI, pick-and-place, predictive), Retail/Media (IVA, kiosk, analytics), and Metro (traffic, safety). The Edge AI Suites tie use-case templates to accelerators, models, and sample apps, shortening the path from pilot to production.

6. A Roadmap for Humanoids & Complex Systems
Looking ahead to humanoids and higher-DoF platforms, the stack accommodates whole-body control, multi-sensor fusion, and multi-modal AI (vision-language models for planning and perception), while preserving safety and determinism. The objective: merge compute, sensing, control, and learning into a serviceable, upgradeable system with lifecycle support aligned to industrial realities.

IV. ENGINEERING FOR REALITY:

Reliability, Openness, and Lifecycle

Industrial customers don't need proof-of-concepts; they need reliable, maintainable systems with long lifecycles. Intel's approach bakes in serviceability and upgradeability, aligns to 10–15-year component lifetimes, and keeps the ecosystem open so integrators can mix and match best-of-breed components across silicon, OS, middleware, and application frameworks. The result is a path to scale that avoids lock-in and reduces total cost of ownership.

V. PROOF THROUGH PARTNERSHIP:

Physical AI in Action

The fastest way from idea to outcome is ecosystem execution. Below is a curated set of partner implementations—spanning inspection, logistics, safety, humanoids, and machine vision—that illustrate how Open Edge Platform + Intel® silicon + Edge AI Suites translate to real-world ROI.

ASUS IoT — Industrial-Grade Edge AI, From AOI to Autonomous Machines

What they do: ASUS IoT delivers ruggedized edge systems (e.g., PE-series) and a software stack (AISVision 365) designed to take computer vision and robotics from lab to line—fast.
Why it matters:

- Time-to-value: AISVision 365 lets teams upload, label, train, and deploy in a day, with prebuilt algorithms for classification, segmentation, anomaly detection, and object detection.
- Hardware fit: Scalable Intel® Core™ i7/i9 platforms support dual-GPU options for high-throughput vision + real-time motion control (e.g., palletizer robots), and wide temperature, anti-shock and vibration designs for harsh environments.
- Use cases: AOI/defect inspection, pick-and-place, object character recognition, and autonomous machines (e.g., fruit-picking robots combining Intel® Core™ i7 + Intel® Arc™ MXM for perception and motion planning).

Takeaway: ASUS IoT operationalizes the "3V" continuum—Vision, Vehicles, and Video Analytics—showing how standardized Intel edge platforms and a packaged vision workflow compress pilot timelines and make factory-scale rollout practical.

MVTec — HALCON AI² + OpenVINO®: Speed Where It Counts

What they do: MVTec's HALCON AI² integrates with OpenVINO®, unlocking dramatic inference speedups on Intel CPUs, and now default Intel GPU and NPU support on Intel® Core™ Ultra.
Why it matters:

- Performance: Representative tests show order-of-magnitude latency gains when moving from CPU-only to OpenVINO-optimized execution (float32/float16) across classification, segmentation, OCR, and detection tasks.
- Portability: Developers can target CPU, integrated GPU, and NPU as needed—without rewriting model pipelines—supporting flexible BOMs and field upgrades.
- Breadth: HALCON covers virtually every classic and deep-learning vision primitive used in industrial automation.

Takeaway: With HALCON AI² + OpenVINO®, teams get predictable acceleration on Intel platforms, making it easier to hit real-time budgets in AOI, OCR/OCV, pick-and-place, and more—without bespoke optimization.

NexCOBOT — Functional Safety, Controllers, and Kits for Humanoids

What they do: NexCOBOT provides SCB100 functional safety control boards, all-in-one controllers (motion + AI + safety), EtherCAT masters, and humanoid/polyped design kits, built on Intel® x86 and integrated with Intel's software stack (oneAPI, RT Linux).
Why it matters:

- Safety as a feature: TÜV-aligned safety modules (e.g., SIL3/PLe) and real-time Linux give integrators a tested path to collaborative and humanoid systems that must operate near people.
- Consolidation: With workload consolidation on next-gen Intel® Core™ Ultra ("Panther Lake") boards, teams can co-locate HMI, AI inference, robotics control, and safety under a hypervisor.
- Interoperability: Proven EtherCAT compatibility with a wide range of servos and I/O simplifies brownfield integration.

Takeaway: NexCOBOT reduces risk and time-to-certification for safety-critical robotics, while keeping the stack open and serviceable—vital for production deployments.

Synapticon — Pushing Functional Safety for AI-Controlled Robots

What they do: Synapticon's POSITRON platform targets the next frontier of robotic safety: safe motion for high-DoF systems, safe human detection, and behavioral safety (e.g., safe guided falling, safe situational motion) as autonomy increases.
Why it matters:

- Beyond basic safety: Moves from STO/SSx to dynamic, behavior-aware safety aligned with emerging standards (e.g., ISO 25785-1 for dynamic stability; ISO/IEC TR 5469 for AI).
- Scalability: POSITRON advances from safe motion monitoring <10 DoF to >50 DoF, enabling humanoids with complex kinematics.
- Path to certification: Combines safety-rated hardware/software with camera-based human detection, aiming at SIL2/SIL3 paths where applicable.

Takeaway: As robots gain capability, safety must evolve. Synapticon's work provides a credible blueprint for AI-aware safety in humanoids and next-gen cobots.

PlayRobot — Quadruped Inspection, Now with Real-World ROI

What they do: PlayRobot's AI quadruped inspection solution targets labor shortage and safety by combining remote and autonomous inspection modes, with on-device analytics using OpenVINO® and Intel® Core™ processors.
Why it matters:

- Productivity & Safety: 24/7 mobility over difficult terrain, 360° vision, optional 6-axis arms, and digital-twin integration reduce human exposure and increase coverage.
- Turn-key: A path from tele-op to autonomous routes ("walk, record, then repeat with stop points"), plus prebuilt safety/analytics workflows for factories and utilities.

Takeaway: Quadrupeds move inspection from scheduled sampling to continuous intelligence, with on-device AI handling detection and triage in the field.

Circulus — On-Device Physical AI for Humanoids

What they do: Circulus demonstrates fully on-device Physical AI for humanoids—showing CPU–GPU–NPU collaborative processing on Intel® Core™ Ultra platforms and retrofits of standard robots to Intel-based compute.
Why it matters:

- Cloud independence: Local perception, language, and control reduce latency and protect privacy.
- Modularity: Swapping in Intel compute (without Jetson) is feasible for education and development kits, scaling to production designs later.
- Developer-first: A path for labs and startups to iterate on voice/vision/actuation loops directly on the robot.

Takeaway: By proving rich multi-modal pipelines on compact Intel edge systems, Circulus points to cost-effective, offline humanoids that still feel responsive and capable.

Geek+ & Unitech — Logistics at Scale, Powered by Intel

What they do: Geek+ (with Unitech in Taiwan) provides an AMR/ASRS portfolio (RoboShuttle, SkyCube, smart sorting, moving systems, smart forklifts) plus a software suite—RMS/WES/IOP—for large-scale scheduling, mixed fleets, and digital operations.
Why it matters:

- System-of-systems: Thousands of robots, cluster path planning, mixed navigation types, and multi-floor/area management demand edge compute + orchestration that scales.
- Proven KPIs: Case studies show 24-hour shipping with 99.99% accuracy, real-time inventory, and dramatic space efficiency gains.
- Intel inside: Configurations range from Intel® Xeon® edge servers to Intel® Core™ i5 embedded controllers ("InstaMove"), matching workload needs to hardware profiles.

Takeaway: When the pain points are variability and volume, standardized Intel platforms and Open Edge tooling reduce the time and cost to expand robotic flows while maintaining SLA/quality.

VI. From Pilot to Production:
How to Move Faster

Start with the blueprint, not a blank page. Use Intel® AI Edge Systems & Verified Reference Blueprints to size correctly and set performance expectations early.

1. Anchor on a real-time-capable stack. Pair PREEMPT-RT Linux with optimized drivers (EtherCAT, cameras) and oneAPI/RT libraries to guarantee determinism where it matters.
2. Choose AI paths for your BOM. With OpenVINO® support across CPU, iGPU, and NPU—and HALCON AI² integration—you can hit latency targets without bespoke ports.
3. Design for reliability from day one. Leverage NexCOBOT/Synapticon solutions to embed safety as a first-class feature, especially for collaborative and humanoid use cases.
4. Use partners to accelerate outcomes. ASUS IoT for packaged vision, PlayRobot for mobile inspection, Geek+ & Unitech for logistics scale, Circulus for on-device humanoid R&D—don't reinvent the wheel.

VII. Conclusion:
Modular, Intelligent, and Open

The robotics industry is no longer in its infancy. It's scaling—and the demands on infrastructure, compute, and collaboration are growing with it. Intel's approach is grounded in reality: build from the silicon up, design with modularity, and scale through openness.

By empowering developers, integrators, and partners with validated hardware, optimized AI tools, and a cooperative ecosystem, we're helping shape a robotics future that is inclusive, scalable, and real.

Notices & Disclaimers

Performance varies by use, configuration and other factors.
No product or component can be absolutely secure.

Results based on specific use case; individual results may vary.
Standards compliance claims provided by partner, not independently verified by Intel.
Partner results are provided by the respective companies.

Intel technologies may require enabled hardware, software or service activation.
Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.