GlobalLogic®
A Hitachi Group Company

Beyond the Chip: Why the Next Semiconductor Advantage Won't Be Won on the Silicon Alone



A CTO at a fabless company signs off on a design she has spent the better part of a year developing and has spent several hundred million dollars to complete. Her primary worry isn't that the chip meets the engineering specifications. It's whether it will meet its market window.

One missed corner case in verification means a re-spin: four to six months to get silicon back, a market window gone, and the lead handed to whoever shipped first. The cost of building a chip has climbed from tens of millions of dollars a decade ago to a few hundred million dollars today, and the cost of being wrong has climbed with it. The chip is no longer the hardest part. The confidence to commit it is.

The competitive boundary in semiconductors is moving beyond silicon. It now runs through the intelligence and adaptability of the company that designs, manufactures, and services the chip.

The shift the boardroom is missing

Most executive teams are asking the wrong question about AI. The visible race is to put AI inside the chip: faster accelerators and smarter edge processors. That race matters. The larger and more defensible lever is to put agentic AI around the chip, across the full lifecycle, from design and verification through manufacturing, yield, field service, and customer enablement.

> "Customers stopped buying just the chips a while ago. They buy complete solutions, the silicon plus the software and the ecosystem around it, and they expect it to be ready to use. The companies that accept they are now solutions driven businesses will pull away from the ones still selling components."
>
> Ravi Thummarukudy, Senior VP & GM Semiconductor Business Unit at GlobalLogic & CEO, Mobiveil, Inc (a GlobalLogic Company)

The semiconductor industry is on track to more than double in size, from $775 billion in 2024 to a projected $1.6 trillion by 2030, as AI demand and a new generation of computing architectures reshape the market (McKinsey, 2026). Growth on that scale resets where the advantage sits. Semiconductor companies are becoming software companies, whether they intend to or not.

Customers no longer buy a die in a package. They buy a working solution: the silicon, plus the drivers, SDKs, reference designs, and the ecosystem that makes it usable in a real product. The decisive lead in AI silicon was built not on transistors alone but on the software ecosystem wrapped around them. Every serious player now faces the same choice: become a solutions and ecosystem company, or supply parts to one that did.

Treated as a chip feature, AI is a line on a datasheet. Treated as an operating-model change, it becomes the thing that lets a company design faster, validate with more confidence, lift yield, and turn its installed base into a continuing relationship. The companies that grasp that difference will pull away within a single product cycle.

Where the pressure is showing up

Three pressures surface first. They are what semiconductor leaders describe when you ask what slows them down.

The first is getting it right the first time. Verification and validation are the highest stakes work in the building. Manufacturing is slow and costly, and a design error caught in the lab can cost four to six months and the market window with it. Validation is harder still, because no company controls the full stack: hardware, firmware, third-party silicon, software, and the customer's real application all must work together, and integration risk grows at every layer. Leaders feel this one most sharply, because the downside is counted in quarters and market position, not engineering hours.

The second is the distance between AI ambition and AI in production. Most semiconductor organizations are interested in AI and are stuck in pilots. The reason is rarely the model. It is the data and the processes beneath it. This is a precise, long-established industry, and the data it has collected for decades was never shaped for AI. Without that foundation, initiatives stall at proof-of-concept and never reach the floor where they would pay off.

The third is the gap between the silicon roadmap and software readiness. Hardware and software can no longer be built in sequence. Waiting for a chip before starting its software means arriving late to markets where last year's product already looks dated. Keeping pace takes parallel development: shift-left engineering, digital twins, and pre-silicon simulation, so the software, ecosystem, and validation are ready when the silicon is.

From tasks to value loops

Most AI narratives get one thing wrong about this industry: AI will not give engineers their time back, and it will not take over the hard judgment. Three decades of design automation have already shown that automating routine work does not lighten the load; it raises the ceiling on complexity and creates more to do. AI will follow the same pattern. Its value isn't fewer engineers. It's engineers working at a scale and with a confidence that was previously out of reach.

Agentic AI makes that scale usable. A traditional workflow runs in a line: inputs in, decisions made, work done. An agentic loop runs as a cycle. Agents sense signals across systems and telemetry, decide the next best action against goals and constraints, then act and learn, improving with each pass.

A few loops stand out in semiconductor work. Verification agents that learn from prior tape-outs discover risk before silicon is committed. Yield agents that close the loop between fab telemetry and design decisions. Field-service agents that turn installed equipment into a stream of operating data and an ongoing customer relationship, so a product sold once becomes a service that keeps earning.

Why reliability is the whole game

In most industries, an over-eager AI produces an awkward answer. In semiconductors, it results in a respin or a stalled fab. The cost of an autonomous wrong decision is too high for the hands-off automation being marketed elsewhere. A second constraint is just as firm: core chip design is a company's most guarded IP, and few will expose it to external, cloud-based models.



> "The scariest outcome in semiconductors isn't a slow chip; it's a wrong one. A re-spin costs four to six months, and missing the market window by a quarter can take a company from leading to trailing. That is why getting it right the first time matters more than raw speed"
>
> Ravi Thummarukudy, Senior VP & GM Semiconductor Business Unit at GlobalLogic & CEO, Mobiveil, Inc (a GlobalLogic Company)

Reliability is the precondition here, not an add-on. Agentic AI becomes deployable only when it is built with hybrid reasoning and humans on-the-loop, with deterministic guardrails around the decisions that carry real cost, and with a data-sovereignty model that keeps proprietary designs and models under the company's own control. Trust isn't a starting assumption. It builds release by release as the system earns a wider mandate. That discipline separates an AI that demos well from one that runs in production.

Proof it runs in production

This is not a theory for us. As part of Hitachi, we build and harden these systems within one of the world's most demanding industrial ecosystems, including Hitachi's own semiconductor and equipment businesses. That makes Hitachi our Customer Zero: real products and real operational risk, where an agentic system must work before it is offered to anyone else. The shift it creates is the prize. An equipment business stops seeing itself as a maker of boxes and starts operating intelligent systems, where monitoring and service run as an autonomous, learning loop instead of a cost center. The same shift is open to any semiconductor company ready to take it.

Why this combination, and why now

For a CXO who has heard "chip to cloud" from three vendors this quarter, the sharper question is who can do all of it. Most partners own one or two parts of the picture. This combination owns three: deep silicon engineering through Mobiveil, system-level integration and software through GlobalLogic, and the experience of running semiconductor equipment in production through Hitachi.

Customer Zero ties them together. Problems get solved inside Hitachi's own businesses, on real silicon and real operations, before they reach anyone else. It shows in the work itself: long-term transformation built on proven, industrial-grade IP, the kind of durable engineering capability we scale for semiconductor clients, rather than project work that ends at handoff.

The cost of standing still

A leader who reads this and waits eighteen months doesn't stand still while the wait plays out. The ground moves. An entrant unburdened by legacy methods reframes the problem and moves faster. The roadmap slips a generation behind. The product on its way to market is already aging, because here a window is a quarter, not a year. Standing still is a decision to fall behind.

Related reading: The New Physics of Software executive brief

Where to start

The path forward doesn't begin with a large bet. It begins with one workflow. Get in touch, and we'll start by identifying the highest-friction workflow in your semiconductor lifecycle, where intelligence and reliability would be most visible to the business, and what it would take to move it from pressure point to advantage.

© 2026 GlobalLogic — A Hitachi Group Company. All rights reserved. GlobalLogic, the GlobalLogic logo, and other GlobalLogic marks are trademarks and/or registered trademarks of GlobalLogic in the United States and/or other countries.