

WHITE PAPER

RISC-V at Foundry Scale: MIPS and GlobalFoundries Industrialize Physical AI with the STAC Framework

Authors:

Stephen Sopko

Semiconductor & Deep Tech Analyst

Ron Westfall

VP and Practice Leader - Infrastructure and Networking

Steven Dickens

CEO and Principal Analyst

OCTOBER 2025



Executive Overview - RISC-V at Foundry Scale

A new phase of computing is emerging - where distributed intelligence beyond cloud data centers now must operate inside the machines that move, build, and protect the physical world. To function safely and predictably, this **Physical AI** requires silicon that can sense, think, act, and communicate in real time. It also requires an open, verifiable path from design intent to certified silicon.

That is the opportunity **MIPS** seeks to capture, and their 2025 acquisition by GlobalFoundries gives them the investment and foundry embedding to accelerate the curve.

From Open Instruction Set to Industrial Execution

RISC-V began as a symbol of architectural freedom - a clean, open instruction set unburdened by legacy licensing. But openness alone does not deliver the rigor, safety, and manufacturability demanded by automotive, industrial, and defense systems.

Because until now, RISC-V lacked a true and customizable **industrial perspective** - a design and development framework beyond academia and catalog providers. While RISC-V has a long legacy in industrial use, it has been an open story, not an integrated co-development story.

By integrating its proven multi-threaded processor architectures directly into a **tier-one foundry workflow**,

MIPS and GlobalFoundries have created **RISC-V at Foundry Scale** - the convergence of open standards with enterprise-grade production discipline. It is a leading RISC-V IP vendor operating inside a global foundry ecosystem, with transistor-level optimization, workload/application-grade validation, and **frequent silicon shuttles** that transform rapid design iterations into practical results.

The outcome is not just faster prototyping; it offers a new industrial model for open hardware.

Defining Physical AI

Artificial intelligence is migrating from **chatbots to robots** - from text prediction to torque control, from data centers to distributed systems that must interpret and respond to the physical world. MIPS implements **Physical AI**: through the class of workloads that must **Sense, Think, Act, and Communicate** within deterministic time boundaries.

These systems depend on **data movement and control**, a MIPS expertise for over 40 years, not just data-center scale inference from catalog silicon providers.

The success of Physical AI is measured in microseconds and efficient use of wafer space, not teraflops.

MIPS's mission is to deliver the silicon foundation for that world: real-time, safety-certified, and customizable. Rather than compete with CPU or GPU vendors chasing cloud AI, MIPS focuses where latency, safety, and lifecycle matter most.

The STAC Framework - A Closed Loop for Physical AI

At the heart of this transformation is the MIPS **STAC** framework - Sense, Think, Act, Communicate.

STAC is both a design language and a product map:

- **Sense:** P-series scalable embedded applications processor clusters that capture and preprocess data at the edge.
- **Think:** S-series edge AI inference engines to interpret data and make decisions.
- **Act:** M-series real-time high-performance microcontroller for field-oriented control and real-time low-latency control loops.
- **Communicate:** I-series data orchestration clusters for real-time event driven data movement tasks.

The addition of Communicate to STAC by MIPS completes the loop. It also defines the space where MIPS's multi-threaded architectures excel - orchestrating data movement between accelerators, memories, and sensors with maximum efficiency & predictable latency.

The Execution Model: Atlas Explorer + GlobalFoundries

Physical AI requires extremely tight collaboration between software, hardware, and manufacturing - not sequential hand-offs between multiple vendors. MIPS and GF are building exactly that continuum.

- **Atlas Explorer** serves as the digital twin of each workload, providing cycle-accurate modeling and architectural tuning without the cost or delay of FPGA prototypes.
- **GlobalFoundries' laboratory environment** converts those models into silicon, validating designs at transistor depth.

This **model-to-wafer workflow** transforms RISC-V development into a living process: not just 'open' but also measurable and repeatable.

Three Pillars of Deployment Success

Every MIPS design engagement rests on three structural enablers that make Physical AI deployment faster, safer, and

more economical:

1. **Open Standards** - RISC-V ISA combined with industry protocols for transparent, portable design.
2. **Workload-focused Technology** - deterministic, event-driven architectures supporting accelerator attach, mixed OS environments, and built-in safety logic.
3. **Collaborative Co-Design** - virtual evaluation through Atlas Explorer linked to GlobalFoundries' shuttle fabrication for continuous optimization.

Together, these pillars define a reliable path from imagination to production.

Markets and Momentum

MIPS technology now scales across four markets where Physical AI is already transforming operations:

- **Automotive:** deterministic control and safety verification across ADAS and in-vehicle networks.
- **Industrial & Robotics:** precision automation and long-lifecycle reliability for Industry 4.0.
- **Networking & Data Center Infrastructure:** efficient data-plane processing for smart NICs and DPUs.
- **Aerospace & Defense:** secure, traceable compute for mission-critical autonomy.

Each domain benefits from the same foundry-embedded workflow, ensuring consistent quality from prototype to production.

A New Standard for Open Hardware

RISC-V at Foundry Scale is more than a slogan - it is the new baseline for credible open silicon. By uniting the freedom of an open instruction set with the reproducibility of foundry-grade manufacturing, MIPS and GlobalFoundries have proven that openness and reliability are not opposites but complements.

This paper outlines how that model works - from the **STAC** framework that defines Physical AI, through the **Atlas Explorer + GlobalFoundries** workflow that validates it, to the **market applications** where it creates measurable advantage.

In short, this is how MIPS is turning RISC-V from an architectural movement into an **industrial platform** - one designed to sense, think, act, and communicate at foundry speed.



Market Context - Why Physical AI Needs Foundry Scale

Artificial intelligence has moved out of the cloud and into the world.

Inference once confined to data-center clusters now occurs inside vehicles, factories, and defense systems - environments where milliseconds define success or failure. These are **Physical AI** systems: machines that sense, decide, act, and communicate in real time. They bring intelligence into the physical domain and expose the limits of the existing silicon supply chain.

From Chatbots to Robots

The first wave of modern AI was dominated by digital assistants and recommendation engines - chatbots that lived entirely in software. The next wave belongs to robots and embedded platforms that must interpret their surroundings and respond instantly. Each camera frame in an autonomous vehicle, each packet in an industrial network, and each control signal in a drone demands a closed feedback loop: **Sense** → **Think** → **Act** → **Communicate**.

The Communicate layer is the bridge between action loops - the point where bits become behaviors. It sees data-plane performance as a safety issue: if messages arrive late or out of order, the system fails. Latency, determinism, and security are now design constraints equal in importance to raw compute throughput.

Where Today's Tools Fall Short

Most of the hardware available to engineers was never designed for these requirements.

- **Latency & congestion:** Data-plane traffic crosses multiple general-purpose cores and interconnects, each adding unpredictable delay.
- **Generic catalog IP:** Commodity cores treat every workload alike, limiting differentiation and forcing OEMs into identical performance envelopes.
- **Fragmented development cycles:** Hardware and software evolve on separate timelines, creating costly re-spins and certification delays.
- **Prototype friction:** FPGAs and custom ASIC prototypes can take months, breaking the feedback loop between model and market.

The consequence is familiar: what designers simulate struggles to match what they manufacture. Physical AI diminishes and collapses under that lag.

Why Foundry Scale Matters

Solving these problems requires more than clever IP. It demands an execution model that joins **design, modeling, and fabrication** inside a single, continuously validated pipeline. That is the promise of **RISC-V at Foundry Scale**, delivered through the tie-up between **MIPS and GlobalFoundries**.

- **Open standards ensure flexibility:** RISC-V and industry protocols free customers from proprietary lock-ins.
- **Transistor-level co-optimization** within GF enables safety, power, and yield tuning unreachable to IP-only vendors.
- **Monthly silicon shuttles** compress validation from quarters to weeks, closing the time-to-prototype gap.
- **Atlas Explorer** forms the digital front end of that pipeline - a virtual environment where workloads are modeled, profiled, and stress-tested before silicon exists. Its outputs feed directly into GF's lab for rapid proof in real wafers.
- **Manufacturing freedom** enables an exit from this iterative development model to any manufacturer, anywhere in the world - not just GlobalFoundries.

This is not merely faster simulation; it is a design-model-foundry loop that collapses the boundaries between idea, model, and device. The result is reproducible silicon that performs as predicted - a requirement for safety-critical markets that can't afford uncertainty.

MIPS's Focus and Position

Unlike CPU or GPU suppliers chasing the next large-language-model benchmark, MIPS targets the connective infrastructure of Physical AI: the **data-movement and event-driven subsystems**


that coordinate sensors, accelerators, and actuators. Storage controllers, DPUs, automotive gateways, and industrial robots all rely on this orchestration fabric. By specializing in deterministic throughput and real-time responsiveness, MIPS fills the space between high-level AI compute and low-level mechanical control-the very gap where latency turns into liability.

The Physical AI Imperative

Autonomous mobility, Industry 4.0, aerospace, and defense modernization all hinge on the same foundation: dependable, time-bounded data flows between sensing, thinking, and acting elements. Delivering that requires silicon that is **safe, customizable, and immediately manufacturable**.

The MIPS + GlobalFoundries combination industrializes that vision. Open RISC-V design freedom meets foundry-grade rigor, producing Physical AI platforms that scale from virtual model to certified hardware without breaking the loop.

Physical AI has evolved beyond research labs and cloud APIs. It now lives on factory floors, inside vehicles, and aboard satellites-and it can only thrive when connections are manufactured at **foundry scale**.





STAC Framework - The MIPS Design Language for Physical AI

As artificial intelligence moves from digital assistants to autonomous machines, the way we design computing systems must evolve. Cloud AI was about training and inference. Physical AI is about interaction - perception, reasoning, response, and feedback happening continuously in the real world.

Earlier frameworks for autonomous systems stopped at **Sense** → **Think** → **Act**.

MIPS extends this lineage by adding Communicate, creating a closed feedback system tuned for real-time, event-driven computing.

Together, these four stages form **STAC** and constitute the organizing principle for MIPS's Physical AI strategy.

STAC is both a conceptual model and an engineering discipline. It defines how information flows through intelligent machines and how silicon, software, and safety interlock to make that flow deterministic. Each stage corresponds to a distinct compute requirement -and to a specific part of the **MIPS RISC-V portfolio**.

Sense - Turning the World into Data

Every action begins with perceiving the world around us.

Physical AI depends on an uninterrupted stream of input from cameras, radars, lidars, and network interfaces that translate the physical world into digital form. These signals must be gathered and pre-processed close to their source, often under strict thermal and power limits.

The **MIPS P8700** architecture addresses this first step with embedded application processors optimized for **data-movement efficiency** rather than sheer clock speed. They handle concurrent sensor feeds, manage local memory hierarchies, and execute lightweight AI models for noise reduction and signal fusion - all in real time.

For automotive and industrial customers, this means sensor data arrives coherent and time-stamped, ready for the decision layer.

For networked systems, it means packets are prioritized before congestion builds.

Key design priorities include:

- **Multithreading and low-latency I/O** for high-rate sensor traffic.
- **Integrated safety** preventing a single fault from propagating.
- **Long-lifecycle support** suited to embedded markets measured in decades.

Sensing is where Physical (digital) AI meets the analog world; tying it to deterministic throughput to make it trustworthy.

Think - Turning Data into Decisions

Once information is captured, it must be interpreted.

The Think phase is where reasoning occurs -inferences drawn, correlations made, trajectories predicted. Unlike cloud inference, embedded thinking happens under **strict temporal budgets**. A braking system or robotic arm cannot wait for a delayed frame of data.

The **MIPS S8200** inference engine blends AI and classical compute at the **edge**, emphasizing:

- **Heterogeneous workloads** - control logic, neural inference, and signal processing side by side.
- **Tight integration with Sense** so pre-processed data flows directly into decision loops.
- **Hardware–software co-design** allowing AI libraries and real-time operating systems to share one safety domain.

Leveraging the extensibility of RISC-V, the S8200 supports custom vector and ML extensions without abandoning standard toolchains. It delivers localized, power-efficient reasoning where correctness and timing outrank raw scale.

In STAC, Thinking is measured not by model size but by how reliably a system reaches the right decision before physics moves on.

Act - Turning Decisions into Motion

Decisions matter only when they move something.

The Act stage converts intent into physical behavior - spinning a motor, adjusting a control surface, regulating a valve, or switching network paths. Timing is absolute: a late signal can create instability or danger.

The **MIPS M8500** real-time controller specializes in **deterministic control and functional safety**.

It combines precise interrupt handling, safety-certified microcontroller cores, and redundant lock-step operation where required by ISO 26262. Core capabilities include:

- **Cycle-accurate response** measured in microseconds.
- **Built-in diagnostic logic** for continuous self-testing.
- **Configurable safety levels** (ASIL-B through ASIL-D) engineered into the RTL itself.

For industrial robotics, this yields smoother motion and less jitter.

For EVs, it ensures torque stability and battery-management safety.

For aerospace and defense, it keeps actuators and sensors synchronized across distributed nodes.

Where Sense and Think handle decisions, Act guarantees **results**.

Communicate - Closing the Loop

Adding **Communicate** transforms the model from a one-way pipeline into a living system.

After an action, information must flow back - confirming execution, updating models, and coordinating across nodes. In distributed AI, this communication dominates total workload. Data must move continuously among processors, accelerators, and storage with minimal latency.

The **MIPS I8500** architecture embodies this stage. It is a multi-threaded RISC-V engine purpose-built for **data movement and packet orchestration**.

Rather than processing payloads, it directs them - managing queues, routing traffic, and enforcing quality of service.

Each hardware thread functions as an independent context, enabling one core to manage several streams concurrently without enlarging silicon area.

Distinct advantages:

- **High throughput per mm²** for dense performance.
- **Power-proportional efficiency** that scales with workload.
- **Programmable pipelines** customers can tailor to networking or control protocols.
- **Open-standard programmability** under Linux, RTOS, or bare metal.

This layer also showcases the **foundry-scale feedback loop**. Through **Atlas Explorer**, engineers model multi-threaded workloads; GlobalFoundries' monthly shuttle runs translate those models into verified silicon.

Communicate is therefore both architectural and procedural - the bridge between **compute and control** and between **virtual design and physical fabrication**.

STAC as a Closed Feedback System

Together, **Sense–Think–Act–Communicate** form a continuous loop:

1. Sense captures.
2. Think interprets.
3. Act executes.
4. Communicate synchronizes and informs the next Sense cycle.

Within **Atlas Explorer**, this framework lets hardware and software teams **co-simulate every stage** before tape-out, guaranteeing that timing, safety, and performance targets align across disciplines.

This closed-loop design makes Physical AI systems **adaptive, safe, and verifiable** - and it is the same structure MIPS uses internally to design and test its RISC-V IP.

Why STAC Matters


The value of STAC extends beyond taxonomy.


It gives OEMs a vocabulary for mapping complex systems to concrete silicon choices, provides software teams predictable boundaries between sensing, reasoning, and control code, and embeds safety analysis directly into design.

Where prior architectures chased universality, STAC embraces **specialization and coordination**. In Physical AI, success is not one chip doing everything - it is every layer working together in real time.

STAC is the grammar of intelligent machines - and MIPS is shipping it to lead customers today

By completing the feedback path between compute and control, Communicate transforms STAC from a framework into a living, self-verifying system. With it, MIPS and GlobalFoundries give RISC-V not just a new architecture, but a new language for the physical world.





RISC-V at Foundry Scale - How MIPS + GlobalFoundries Turn Openness into Industrial Silicon

RISC-V began as a language of freedom. To power **Physical AI** (machines that must sense, think, act, and communicate under tight, deterministic budgets) that language needs an execution model that is **manufacturable, certifiable, and repeatable**.

MIPS and GlobalFoundries (GF) provide that model by uniting **fit-for-purpose RISC-V microarchitectures** with a **foundry-embedded, feedback-driven workflow**. The result is not a single chip but a way to go from **idea** → **verified model** → **proven silicon** with the rigor enterprises expect.

Design Ethos: Determinism over Generic Peak Scores

Instead of chasing general-purpose benchmarks, MIPS aims to build **workload-custom, event-driven subsystems** that map to the STAC loop:

- **Sense / Think / Act / Communicate** are treated as distinct compute roles, with clear timing and safety boundaries.
- **Multithreaded pipelines** keep execution units busy across concurrent data flows; when one thread stalls

on I/O, another advances-sustaining throughput and power efficiency.

- **Programmable data paths** let teams tailor packet orchestration, sensor fusion, or control loops without discarding standard toolchains.
- **Safety-ready logic** (with paths to ASIL levels depending on configuration) and **virtualization/mixed-OS** options support real systems where consolidation and isolation both matter.

This approach reflects a simple truth from decades of deployed systems: **data movement is the new compute**. Optimizing how information moves among sensors, accelerators, memories, and actuators yields more deterministic systems-without exotic process nodes or runaway power budgets.

Proven Foundation, Modernized on RISC-V

MIPS's current data-movement engines build on a heritage of high-volume deployments in demanding, real-time environments. The step forward is the migration of that proven design DNA to **RISC-V**, pairing open-ISA flexibility and modern toolchains with a microarchitecture tuned for **low-latency orchestration**. Customers gain an **open, extensible** platform backed by silicon that has already informed software stacks and safety practices in the field.

The Three Execution Pillars in a Single Flow

Open Standards → Workload-Custom Technology →

Collaborative Co-Design used to be separate phases. At foundry scale, they become a single, continuous method:

1. Open Standards - Freedom That Scales

RISC-V plus widely adopted industrial/automotive protocols give teams a transparent base to extend instructions, attach accelerators, and mix Linux/RTOS/hypervisors- without vendor lock-in or toolchain fragmentation.

Openness, here, is disciplined: shared interfaces, visible boundaries, auditable behavior.

2. Workload-Custom Technology - Engineered for Reality

Multithreading, programmable micro-engines, native accelerator hooks, and safety-ready RTL deliver **predictable latency** and **throughput per watt** for the actual job: moving data deterministically between sensing, decision, and action. Because it's RISC-V, domain-specific instructions can be added **without breaking** software continuity.

3. Collaborative Co-Design - From Model to Wafer

Atlas Explorer serves as the digital twin for STAC systems: teams model workloads, profile cycle-accurate behavior, and tune micro-architecture **before** tape-out. Those models feed a **regular MPW/shuttle cadence at GF** (often monthly by node), where prototypes validate timing, power, and safety assumptions in real silicon. Measured results flow **back** into Atlas to refine RTL and software. Instead of "design → fabricate → debug," hardware and software evolve **in lockstep**.

Foundry as Laboratory: A Verified Continuum

Embedding the co-design loop **inside a tier-one foundry ecosystem** turns prototyping from a sporadic event into a **repeatable lab practice**:

- **Transistor-level co-optimization** aligns layout and libraries with GF process capabilities (e.g., automotive-grade platforms), improving correlation between models and wafers.
- **Regular shuttle runs** shorten the distance from architectural hypothesis to measured silicon behavior, so timing, leakage, thermals, and safety mechanisms get verified early and often.
- **Traceable, long-lifecycle flows** support industries where availability and compliance outlast product launch windows.

- **Portability remains:** the methodology is open-standard and repeatable, preserving options for multi-foundry strategies.

Each wafer is more than a prototype - it's a **feedback** signal that tightens the model-to-manufacturing loop.

Why This Matters for Physical AI

Physical AI collapses if simulation diverges from hardware. The MIPS + GF model narrows that gap:

- **Speed:** Shuttle-verified silicon keeps hardware schedules aligned with evolving software workloads.
- **Certainty:** Correlated models reduce surprises at bring-up and de-risk certification.
- **Safety:** Safety-ready IP, verified within automotive-grade flows, supports deterministic behavior where it's non-negotiable.

Across automotive gateways, industrial robotics, networking/ dataplane infrastructure, and mission-critical systems, the pattern is identical: **deterministic communication between sensing, decision, and control**, made practical by a workflow that is **open where it matters, customized where it counts, and verified where it's critical**.

The Bottom Line

RISC-V at Foundry Scale is not a slogan; it's a **working operating model** for building trustworthy silicon:

- A **deterministic microarchitecture** optimized for data movement.
- An **open, extensible ISA** that preserves software continuity while enabling domain-specific gains.
- A **co-design loop** that binds Atlas Explorer to a foundry-embedded prototyping pipeline, converting virtual wins into physical certainty.

MIPS offers the freedom of customer choice in manufacturing partner. Part of RISC-V at Foundry scale and openness is to go where customers are, which means geographically, in their workload, and in their choice of manufacturing partner. This is how openness becomes infrastructure-and how Physical AI gets engineered for the real world.



Target Market Use Cases

Physical AI is not a single market-it is a pattern appearing across many. Wherever real-time intelligence must interact with the physical world, the same needs emerge: safety, determinism, and efficient data movement.

Because MIPS designs, validates, and manufactures at foundry scale, its technology applies consistently across industries that once required entirely different architectures.

Four sectors illustrate this breadth.

Automotive - Safety, Latency, and Lifecycle

Modern vehicles are distributed computing platforms on wheels. They contain hundreds of processors performing thousands of real-time tasks-from motor control and battery management to in-vehicle networking and driver assistance. Autonomy demands not only AI inference but **deterministic communication** among all of these systems.

The **MIPS STAC** framework maps perfectly onto this environment:

- **Sense** → P8700 processors capture and pre-filter camera, radar, and LIDAR data.
- **Think** → S8200 inference engines perform object classification at the edge.
- **Act** → M8500 controllers handle power-train and chassis control with ASIL-D-ready logic.
- **Communicate** → I8500 cores orchestrate data flow between domains, ensuring every decision reaches its actuator on time.

GlobalFoundries' **automotive-grade processes** make this silicon inherently compliant with ISO 26262 and AEC-Q100. Because Atlas Explorer and GF's monthly shuttles validate each generation of the architecture under these standards, customers receive silicon that is both **safety-ready and production-ready**. Long-lifecycle process availability-ten years or more-meets OEM supply commitments and reduces recertification overhead.

For automakers, this means fewer ECUs, lower power draw, and the ability to consolidate workloads on an open-standard platform verified from model to wafer.

Industrial & Robotics - Precision and Longevity

Factories and robotic systems operate under similar constraints but harsher conditions. Vibration, heat, and electrical noise demand processors that are **predictable, durable, and verifiable**. Industrial automation increasingly depends on AI to adjust to variable materials or unexpected inputs-yet timing remains paramount.

MIPS delivers deterministic control through its **Act (M8500)** layer, while **Sense (P8700)** and **Think (S8200)** cores manage sensor fusion and on-device analytics. The **Communicate (I8500)** layer links these systems enabling machine-to-machine collaboration without latency spikes.

Because the entire workflow is validated during the Atlas Explorer to GF pipeline, developers can simulate factory conditions inside Atlas Explorer, prototype at foundry speed, and deploy hardware with confidence it will behave identically in the field.

For industrial OEMs, this reduces costly line downtime and enables **predictive-maintenance AI** at the edge-without sacrificing real-time control.

Networking & Data Center Infrastructure - The Data-Plane Advantage

Data centers and network operators face a different kind of real-time problem: **moving enormous volumes of information efficiently**.

Here, the **Communicate** layer becomes the centerpiece.

The **I8500** provides multi-threaded packet processing and data-plane management, integrating seamlessly with Smart NICs, DPUs, and storage controllers. Its open RISC-V architecture allows customized extensions for encryption, compression, or traffic shaping-tasks previously handled by proprietary ASICs or cumbersome collections of catalog hardware.

Atlas Explorer enables architects to profile these workloads at cycle-accurate detail before hardware exists, while GF's monthly shuttles convert those models into test silicon for verification.

The result is a class of **event-driven, low-latency RISC-V subsystems** that free CPUs to focus on orchestration and application logic.

For hyperscalers and telecom providers, MIPS offers an escape from the power and cost ceilings of traditional network processors, bringing **data-plane efficiency once exclusive to Broadcom-class designs** into an open, foundry-verified ecosystem.

Aerospace & Defense - Security and Reliability at the Edge

Aerospace and defense applications stretch the boundaries of Physical AI. Satellites, UAVs, and tactical systems require intelligence that operates securely and autonomously, often beyond communication range.

Failure is not an option; every component must be **radiation-tolerant, tamper-resistant, and traceable** from design through production.

The MIPS-GF model meets these needs through its **closed-chain-of-custody fabrication** and **safety-verified IP library**. Atlas Explorer simulations feed directly into controlled GF laboratories where each design is fabricated, tested, and audited. This traceability-combined with RISC-V's inspectable ISA-supports security reviews that proprietary cores cannot match.

Because MIPS retains cross-foundry flexibility, aerospace and defense integrators can produce hardened derivatives at multiple facilities while maintaining a consistent verification baseline. The result is **mission-critical autonomy** powered by open standards and manufactured with enterprise-grade assurance.

A Single Architecture, Many Domains


Across these markets, the pattern repeats:

- **Safety and determinism**
- **Precision and longevity**
- **Throughput and efficiency**
- **Security and reliability**

All draw from the same STAC foundation and the same foundry-embedded design process. Customers gain differentiated performance without vendor lock-in, and an upgrade path that spans industries and decades.

This is the practical expression of **RISC-V at Foundry Scale**; a single, open architecture engineered to sense, think, act, and communicate wherever intelligence meets the physical world.





Analyst POV & Strategic View

RISC-V at Foundry Scale is an execution model, not just a marketing tagline. By binding open-ISA design, workload-custom microarchitectures, and a foundry-embedded feedback loop, MIPS + GlobalFoundries turn “open” into operational certainty for Physical AI at an industrial scale.

What Changes (and Why It Matters)

- **From IP to Infrastructure:** Openness is paired with a repeatable path from digital twin to proven silicon, tightening correlation between what teams model and what ships.
- **From Peak Scores to Determinism:** Data-movement efficiency and bounded latency become first-class metrics, not afterthoughts - exactly what STAC workloads require.
- **From One-off Tape-outs to a Lab Rhythm:** Regular MPW/shuttle access and transistor-level co-optimization shift prototyping from sporadic to systematic.
- **Silicon IP as Design Partner:** Instead of building a static catalog of IP, the co-development process increases end-product differentiation and meets specific OEM requirements.

Customer Value by Stakeholder

OEMs & Tier-1s (Auto/Industrial/Embedded)

- **Freedom with guardrails:** Co-design in Atlas, verify in GF flows, and carry forward on an open-standard toolchain.
- **Faster, safer ramps:** Model-to-silicon convergence shortens

cycles and reduces re-spins; safety-ready IP supports certification paths.

- **Differentiation without lock-in:** Domain-specific extensions and accelerator attach on RISC-V—without abandoning ecosystem software.

Hyperscalers & Network Operators

- **Efficient data planes:** Offload packet/storage plumbing to multi-threaded, event-driven engines; reserve CPU for orchestration and services.
- **Iteration at hardware speed:** Profile workloads pre-silicon, validate on shuttle silicon, and evolve designs without committing to proprietary ASICs.
- **Lower lifetime cost:** Reusable, open cores and a consistent verification path support scale and supply diversity.

Industrial & Automotive Integrators

- **Predictable compliance:** Safety-ready designs and traceable foundry flows align innovation pace with regulatory cadence.
- **Lifecycle assurance:** Long-lived processes and verified behavior reduce recertification risk and platform churn.

Investors & Analysts

- **Commercial maturity signal:** Evidence that open-ISA platforms can meet enterprise reliability expectations with repeatable, foundry-grade methods.
- **New value capture:** Co-design, digital twins, and continuous verification enable service- and platform-oriented revenue beyond static IP licensing.

The Pragmatic Center of RISC-V

The RISC-V ecosystem tends to split between “generic breadth” and “niche accelerators.” MIPS anchors the **pragmatic center** - open enough to flex, structured enough to scale. Integration with a tier-one foundry converts open hardware from ideal to **auditable process**. Decision-makers can expect:

- **Faster deployment** of verified designs (because models and wafers stay in sync).
- **Lower NRE and compliance drag** (because safety and determinism are designed-in, not bolted-on).
- **Higher manufacturability confidence** (because layout, libraries, and workloads co-evolve inside a repeatable lab environment).

The Takeaway

The MIPS + GF approach - **open where it matters, customized where it counts, verified where it's critical** - is how Physical AI becomes dependable infrastructure. That means RISC-V should no longer be evaluated on ideology; it competes on **execution quality**.

What's Next

- **For builders:** Map your STAC workload into Atlas, identify the latency-critical paths, and prototype on the next available shuttle.
- **For platform owners:** Standardize on open interfaces, reserve proprietary effort for domain-specific gains, and institutionalize the model-to-silicon loop.
- **For the ecosystem:** Treat foundry-embedded co-design as the baseline, not the exception.

Bottom line: The age of experimentation or catalog vendors is giving way to **RISC-V at Foundry Scale** - a disciplined, repeatable, optimized way to manufacture intelligence on open standards for the physical world.





ABOUT HYPERFRAME RESEARCH:

HyperFRAME Research delivers in-depth research and insights across the global technology landscape, spanning everything from hyperscale public cloud to the mainframe and everything in between. We offer strategic advisory services, custom research reports, tailored consulting engagements, digital events, go to market planning, message testing, and lead generation programs.

Our industry analysts specialize in rigorous qualitative and quantitative assessments of technology solutions, business challenges, market forces, and end user demands across industry sectors. HyperFRAME Research collaborates closely with your Analyst Relations, Product, and Marketing teams to build and amplify your thought leadership, positioning your expertise to enhance brand and product recognition. Through content that engages readers, viewers, and listeners alike, we ensure your voice resonates across channels.

CONTACT HYPERFRAME RESEARCH:

Steven Dickens

CEO & Principal Analyst | HyperFRAME Research

Email Address:

steven.dickens@hyperframeresearch.com

Telephone Number:

+1 845 505 1678

X: @StevenDickens3

LinkedIn: Steven Dickens

BlueSky: Steven Dickens

CONTRIBUTORS

Steven Dickens

CEO & Principal Analyst

Stephen Sopko

Semiconductor & Deep Tech Analyst

Ron Westfall

Analyst In Residence

INQUIRIES

Contact us if you would like to discuss this report and HyperFRAME Research will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "HyperFRAME Research." Non-press and non-analysts must receive prior written permission by HyperFRAME Research for any citations.

LICENSING

This document, including any supporting materials, is owned by HyperFRAME Research. This publication may not be reproduced, distributed, or shared in any form without the prior written permission of HyperFRAME Research.

DISCLOSURES

HyperFRAME Research provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

