




WHITE PAPER

MIPS I8500 and GlobalFoundries: Boosting AI Data Movement Efficiency

Integrating Data and Control Planes for Physical AI Workloads

Authors:

Stephen Sopko

Semiconductor & Deep Tech Analyst

OCTOBER 2025



Executive Summary

Artificial intelligence is moving from data centers into the physical world - from cloud-hosted chatbots to in-the-field robots: autonomous machines, industrial controllers, and connected infrastructure. These Physical AI systems must sense, think, act, and communicate (STAC) in real time. Yet most current architectures remain divided between control-plane CPUs and data-plane accelerators, introducing latency, power inefficiency, and software complexity.

The **MIPS I8500** addresses this gap as the next generation of MIPS's data-movement microarchitecture - a design lineage already proven in millions of deployed units at customers such as Mobileye - now evolved into a native RISC-V implementation. The I8500 is a soft IP release, independent of process technology, and can be fabricated in multiple nodes, integrating control- and data-plane processing within a unified, highly scalable, multithreaded architecture. The result: a programmable, efficient, and standards-based foundation for customers building real-time, event-driven systems in the communications, storage, automotive, and industrial domains.

Already sampling to lead customers, the I8500 translates decades of MIPS data-movement expertise into the RISC-V era by providing freedom with guardrails for designers who demand determinism, scalability, and ecosystem credibility.


Market Context: Physical AI and the Integration Challenge

AI workloads have evolved. Training and inference remain data-center staples, but the next frontier is AI that interacts with the real world in real time - machines that move material, steer vehicles, manage power, or route data. These workloads demand ultra-low latency and predictable response, yet are often constrained by fragmented system architectures.

In traditional designs:

- **Data-plane tasks** (packet forwarding, sensor I/O, accelerator control) run on specialized processors or firmware.
- **Control-plane logic** (management, scheduling, OS) runs on a general-purpose CPU.

The boundary between the two becomes a bottleneck, especially as edge systems absorb more AI capability. Every hand-off between cores or subsystems adds latency, burns power, and complicates software development.



Customer Pain Points

Networking & Communications: Packet processing and flow-control engines often rely on discrete match processors or proprietary ASICs, limiting flexibility and programmability.

- **Storage Controllers:** Increasing concurrency between NVMe, SSD, and AI-driven caching requires deterministic multi-thread handling.
- **Automotive Gateways & ADAS:** Real-time safety logic must coexist with Linux-class software, forcing redundant compute domains.
- **Industrial and A&D Systems:** Long-lifecycle designs need open standards and configurability, yet certification and toolchain fragmentation slow adoption.

Across these markets, customers need data-movement performance with software freedom - and a migration path from proprietary IP to open, verifiable architectures.

Technology Discussion: The MIPS I8500

The MIPS I8500 is a **fully featured, RISC-V-compliant processor IP core** engineered for intelligent data movement. It delivers the throughput and flexibility once reserved for proprietary network processors within a standard ISA and toolchain.

As a fully-featured RISC-V processor, the I8500 runs Linux, RTOS, or bare-metal workloads on standard toolchains, and gives designers full control over memory topology and system integration to optimize programmable dataflow for their own SoC.

Architecture Strengths

- **Four threads per core:** Each hardware thread (or “hart”) operates independently, allowing one thread to manage a hardware accelerator while others process control logic or I/O tasks. This minimizes idle cycles and ensures consistent latency.
- **Cluster scalability:** Up to six cores per cluster and 64 clusters per system deliver exceptional scale (up to a maximum of 384 cores / 1,536 threads) without abandoning deterministic behavior. Built for high-throughput, low-latency match processing, the I8500 scales efficiently across cores and clusters to handle modern data-movement workloads.
- **Programmable pipeline:** Supports Linux, RTOS, or bare-metal deployment; developers can mix real-time control with user-space applications on the same device.
- **Flexible memory topology:** Cache plus directly addressable instruction / data RAMs allow designers to tune latency and cost for their workload.
- **Hardware integration:** Open interfaces enable easy coupling with NPUs, FPGAs, or custom ASIC blocks; customers can tailor their SoC architecture for programmable dataflow.

This combination yields **match-processor performance**, including rapid evaluation and routing of data streams, without resorting to proprietary accelerators. Multi-threaded efficiency scales nearly linearly with workload parallelism, maximizing performance per mm² of silicon.

STAC System-Level Integration

Within the broader MIPS **Sense-Think-Act-Communicate (STAC)** framework, the I8500 occupies the “Communicate” domain: orchestrating data between sensors, accelerators, and networks. Its low-latency context switching and deterministic threading let physical-AI systems maintain real-time feedback loops - closing the gap between decision and action.

Example Applications

- **Smart NIC / DPU:** Integrates control and data planes for packet orchestration, replacing multi-chip designs with a single programmable fabric.
- **Storage Controller:** Handles concurrent I/O paths and AI-based caching with consistent response time.
- **Automotive Gateway:** Supports mixed-criticality workloads, such as Linux-based infotainment alongside RTOS-driven safety logic, on one platform.
- **Industrial Robotics and A&D:** Provides deterministic control, safety monitoring, and data exchange within Industry 4.0 and defense-grade networks.

Across all cases, customers can **simulate, customize, and validate designs in Atlas Explorer**—MIPS’s virtual platform enabling software/hardware co-design before silicon. Atlas Explorer functions as a **digital twin** environment, allowing cycle-accurate profiling and early optimization of cache, memory, and thread utilization.

Atlas Explorer and GlobalFoundries Co-design

MIPS’s integration within **GlobalFoundries (GF)** transforms its execution model. The company now combines open-standard IP development with **foundry-scale validation** - a capability no other RISC-V provider currently offers.

Atlas Explorer is a software development and optimization platform from MIPS, designed as a Visual Studio Code extension for their portfolio of RISC-V compute subsystems, enabling pre-silicon workload analysis and digital-twin modeling throughout product lifecycles.

Combined with GF’s monthly shuttle program, MIPS empowers customers to prototype and prove silicon rapidly, shortening the feedback loop between design intent and physical realization. This connection delivers customers:

- **Fabrication credibility:** Access to GF’s proven automotive- and industrial-grade processes.
- **Acceleration of time-to-silicon:** Iterative validation using real wafer prototypes rather than purely virtual models.
- **Road-map alignment:** Foundry-verified power, performance, and thermal models that match production behavior.

Together, MIPS + GF extend open innovation to industrial scale, i.e., **RISC-V at Foundry Scale**, bridging the gap between concept and qualified silicon.



Strategic Implications

For System Designers: I8500 clusters provide a ready-to-deploy foundation for event-driven systems where every microsecond counts. Designers can tailor coherence, cache, and accelerator links while staying within standard RISC-V toolchains.

For OEMs and Hyperscalers: Programmability across Linux, RTOS, and bare metal enables unification of control and data planes - simplifying boards, reducing BOM costs, and lowering power per transaction.

For Investors and Ecosystem Partners: The I8500 demonstrates that RISC-V is no longer experimental. MIPS delivers a mature, revenue-class architecture with decades of field-proven data-movement lineage and industrial-grade foundry backing.


MIPS's mission is not to replicate the general-purpose CPU market but to lead the space where **data movement, real-time control, and AI converge**. This focus positions the company as the **reference IP provider for Physical AI** platforms across automotive, industrial, communications, and storage.

Conclusion + HyperFRAME Research POV

HyperFRAME Research believes the emergence of Physical AI marks a structural shift in semiconductor demand. Success will hinge on architectures that combine deterministic performance, software flexibility, and ecosystem credibility.

The MIPS I8500 embodies these traits. By merging proven data-movement heritage with a modern RISC-V foundation and GlobalFoundries' industrial infrastructure, MIPS has created an execution model that others will struggle to replicate.

For customers seeking to move AI from data centers into machines that sense, think, act, and communicate in the real world, the I8500 represents not just, a product but a pathway to build with confidence and scale.





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

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.

