

Cadence Brings Chip Verification to the Next Level with AI Agents

Karl Freund

Cambrian-AI Research

March 23, 2026

Sponsored by Cadence Design

Introduction

The chip design process has become a lynchpin in the global digital economy as cloud computing and artificial intelligence grow rapidly. The economic impact is huge. Forrester projects that, across 14 major countries, the digital economy will reach 16.5T USD and capture 17% of GDP by 2028. APEC and similar policy groups often quote that roughly 70% of new value created over the next decade will be based on digitally enabled platforms, even if the directly measured “digital sector” remains closer to the teens as a share of GDP.

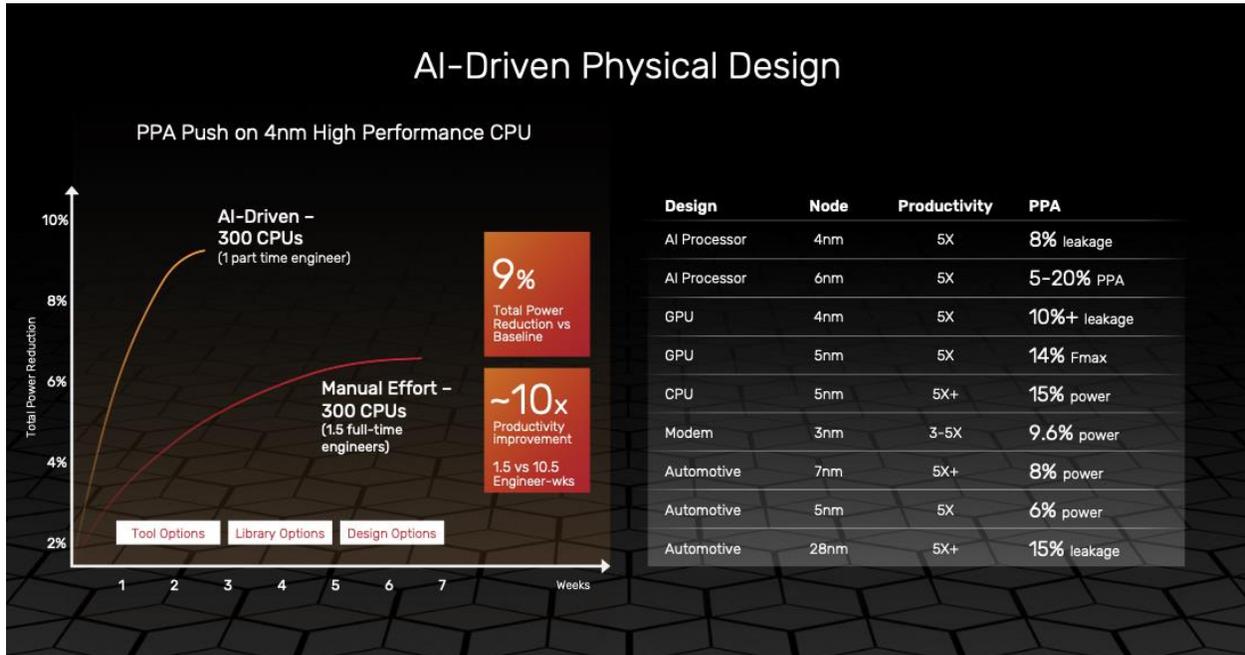
Unfortunately, this economic flywheel is burdened by the time and cost of chip development, which is typically around two-three years and two-three hundreds of millions of dollars. But what if that time and expenses were reduced dramatically to as little as two months and tens of millions of dollars? Then the flywheel would spin ten times faster, as would the pace of innovation across nearly every industry and economy.

This is not just a pipe dream. In fact, every EDA company is already hard at work to achieve these goals through the application of AI to accelerate the design process and lower its cost. This paper will examine Cadence’s latest innovation in this trend, the ChipStack AI Super Agent, and its emerging role at the cornerstone of automation of silicon and system design.

AI has Already Impacted Design Productivity

Up to this point, AI has been broadly applied to physical design automation, or “place and route” work once a logical design has been completed, as well as providing conversational access to documentation and applying reasoning for problem solving. The impact of AI on this work is significant, improving productivity, Power, Performance and Area (PPA). Reinforcement learning tools like Cadence Cerebrus and Synopsys DSO.ai have now accelerated tape-out of thousands of chips with users reporting a typical 3x productivity uplift in PPA exploration and

notable gains such as power reductions of up to 15%, die size reductions, and typically boosting performance by 10%.



But while these results are impactful, we are now just at the beginning of AI' impact on the overall design workflow as Agentic AI emerges to solve more difficult challenges.

Agentic AI in Chip Design

Agentic AI refers to autonomous, goal-oriented AI systems that can independently reason, plan, and execute multi-step tasks with minimal human supervision. Unlike chatbots that only generate text, agentic AI interacts with software (in this case EDA tools), and data to achieve specific objectives. It represents a shift from passive, reactive AI to active, proactive agents capable of handling complex, dynamic, and cross-functional workflows.

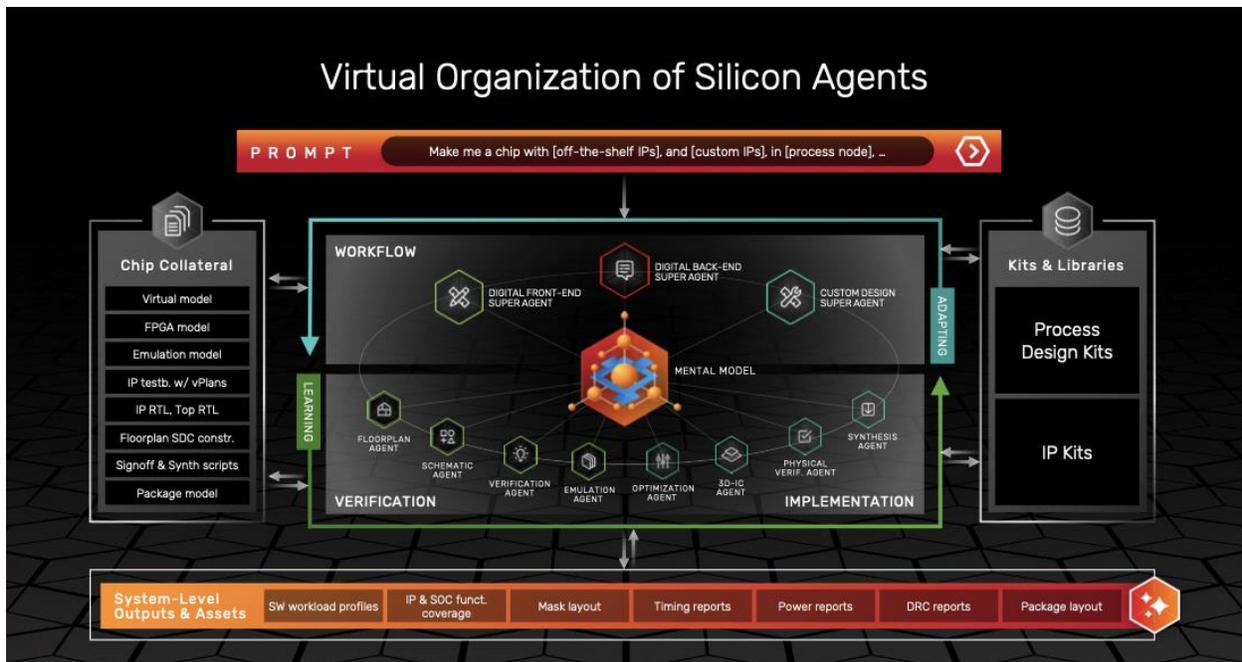
If this sounds ideal for EDA design workflow, it is. Cadence has built its first silicon agents to concentrate on the labor-intensive design verification phase first and then apply agents more broadly to other areas of chip and system design as they progress.



The Impact on EDA Professionals

While AI’s impact on design is dramatic, it will not replace design engineers even in the long-term. Instead, it will help design teams be more productive and accelerate innovation. Engineers will use these tools to shorten the time for each iteration of verification now and additional workflow elements over time.

The agents will become both advisors and implementation tools accessing EDA workflow tools. Entry-level design engineers will ramp up more rapidly and will spend less time asking questions, making mistakes and seeking advice of more senior engineers. Senior-level engineers can then spend more of their time on building the “mental models” for more advance designs. Overall, design teams will operate more efficiently and quickly, helping the business accelerate its designs and time to market. And keeping engineers-in-the loop- remains essential to avoid missteps and design flaws.



The Super Agent for Chip Verification and the “Mental Model”

The Chip Stack AI Super Agent suite from Cadence includes agents for key skills which access Cadence EDA tools, from verification now to implementation phases in the future. This is the next step of a major retooling of EDA for design engineers.

The challenge with applying AI to the full workflow is that one needs to build a database and software that adequately captures the design intent and priorities of the chip. What is its function? What optimizations like pruning and compression are desired? How should the design trade-off performance, power efficiency and speed? What numeric precisions should be exploited? What is the optimal zone to maximize value and profitability? In ChipStack, a company which Cadence acquired last November, these factors are embodied in the Super Agent’s “Mental Model”.

The Mental Model provides the ground truth for the chip’s behavior intention and goals while the “Super Agent” orchestrates the entire workflow. Cadence plans to extend capabilities beyond verification as the design progresses.



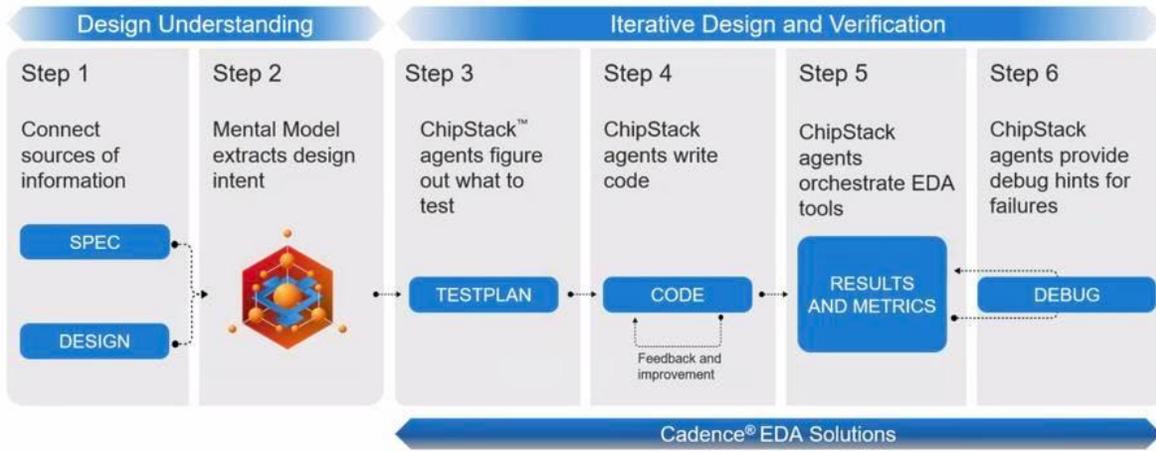
The impact of this approach is remarkable, with test correctness rising from a baseline of 42% improvement using RAG to some 83% using the Mental Model. The design inputs include RTL, design specs, behavioral models and a block diagram. The Mental Model Agent then elaborates on the hierarchy, connectivity and various parameters, which it then enriches. Finally, it reasons the organization, relationships and abstractions to produce the mental model, defining the various relationships and hierarchy of the functions the chip is designed to perform.



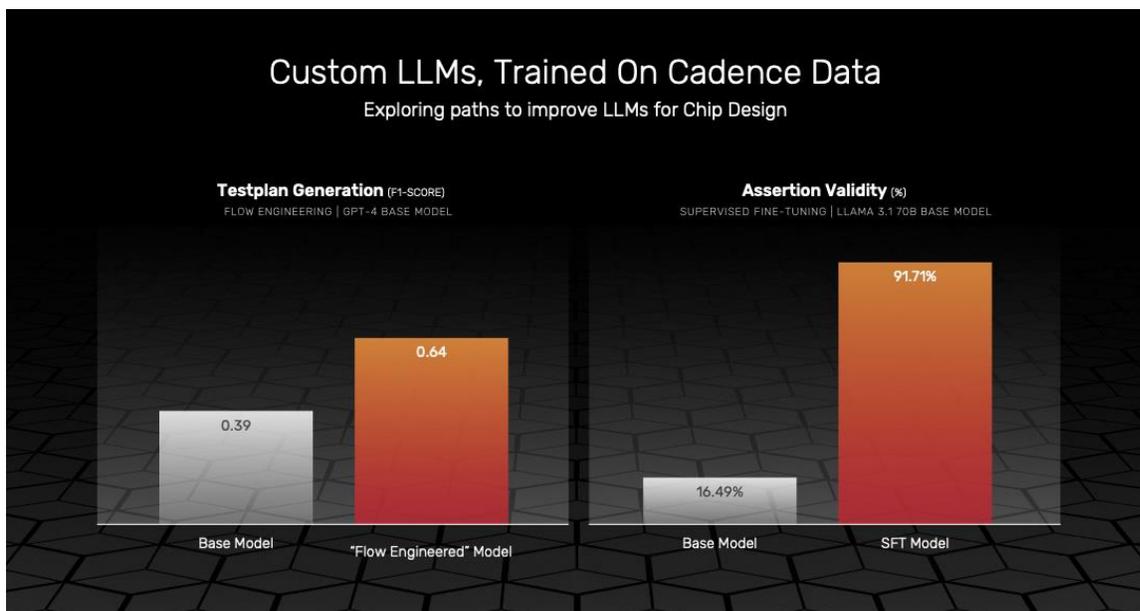
The end-to-end verification is an iterative process of determining what parameters to test, writing the code to conduct those tests (also an iterative process), measuring the results against

goals and providing debug hints for designers to explore. Again, the need for a human-in-the-loop is apparent in the slide below.

End-to-End Agentic AI Verification Powered by Cadence's ChipStack AI Super Agent

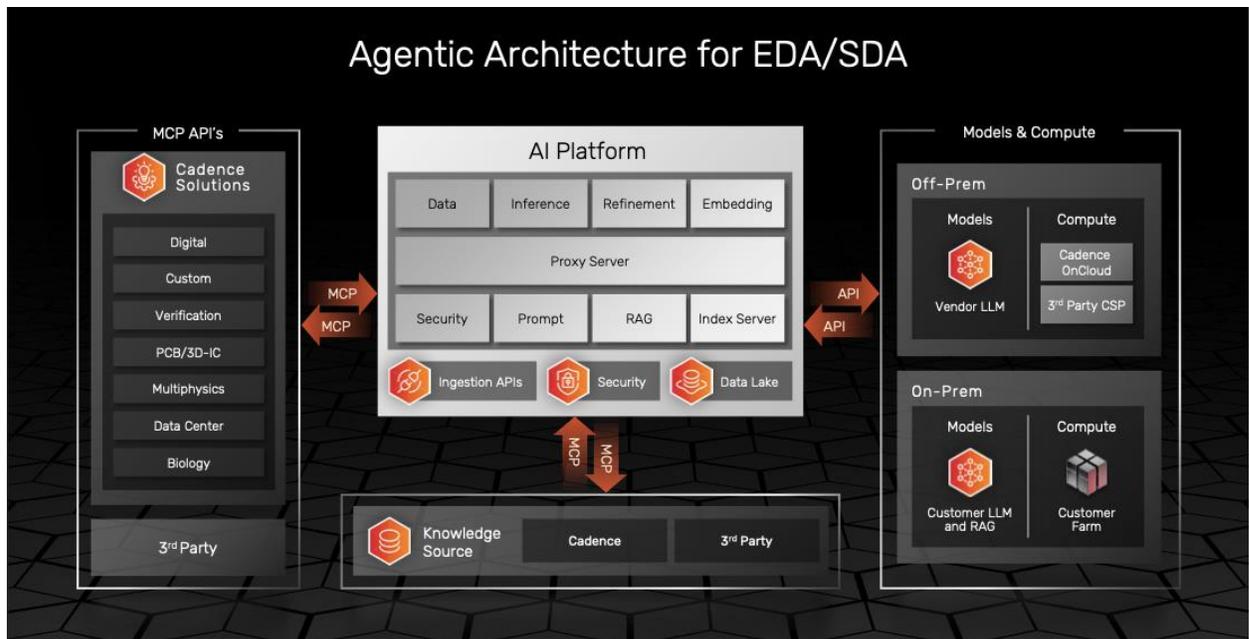


Cadence utilizes custom LLMs trained on historical design experiences. These LLMs have demonstrated a 20% improvement on test plan generation and an over five-fold increase in assertion validity using supervised fine tuning. Once again, we see that solid AI complemented by a human-in-the-loop produces the optimal results.



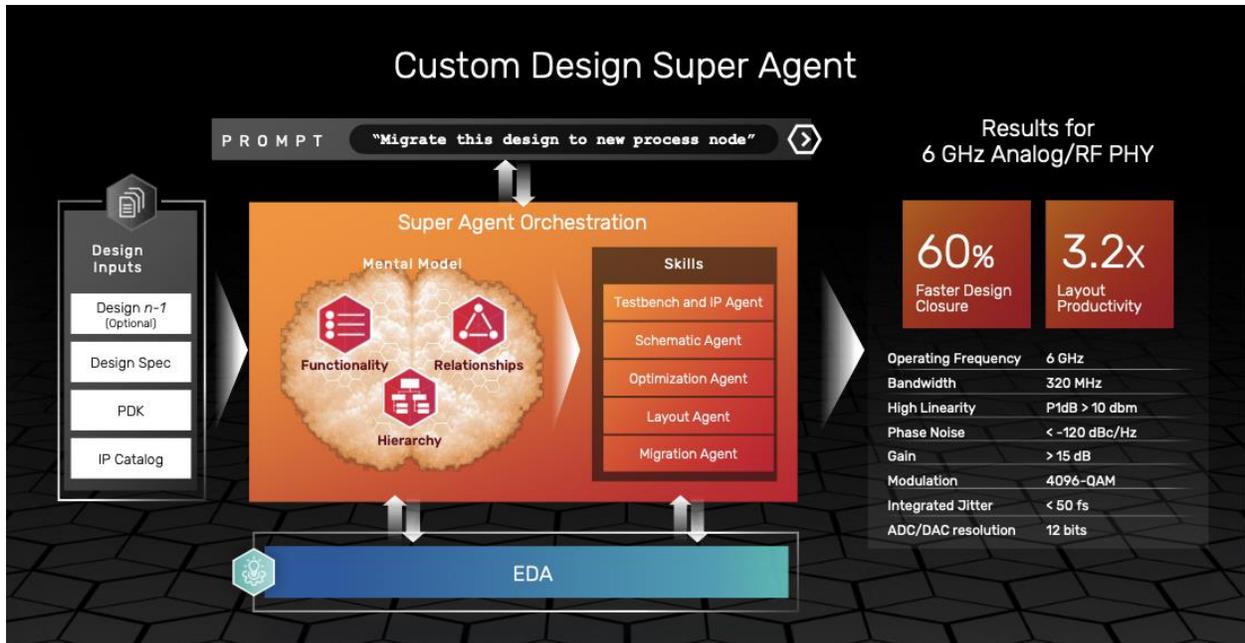
The Road Ahead: Custom EDA and SDA Agentic Architecture

The Cadence vision for AI does not just stop at the chip level. Management sees this same agentic architecture to be applicable for custom chip design and system-level design, or SDA.



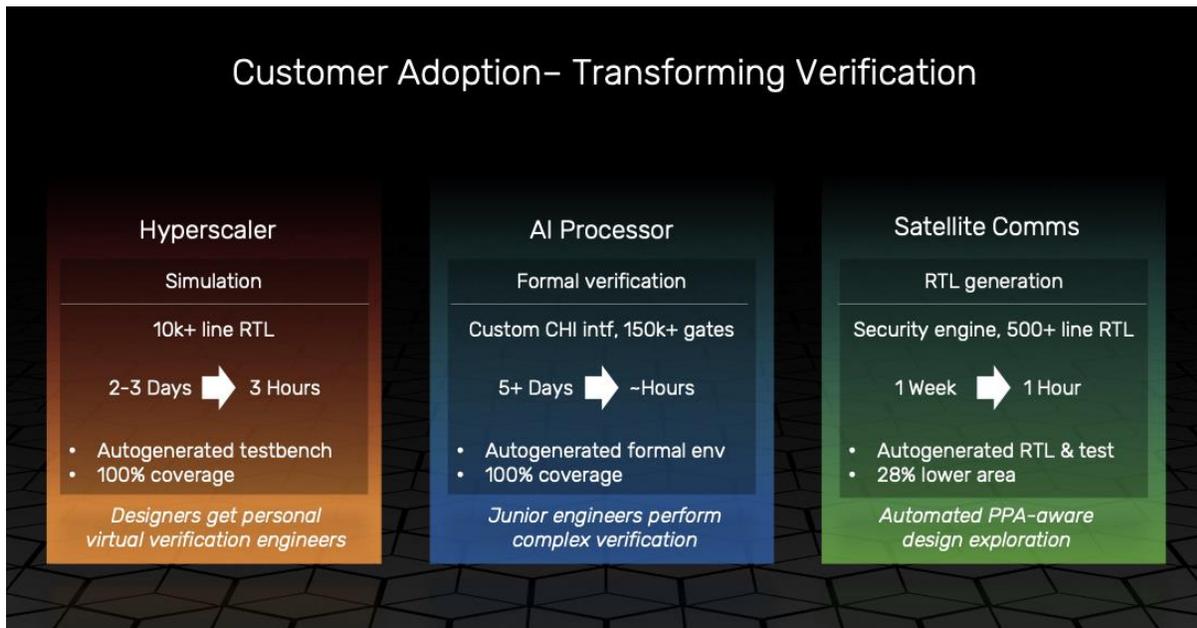
The progression of automation beyond digital chip design includes PCBs, 3D-ICs, multi-physics, system, data center and even biology, where Cadence has been building solutions for several years.

In the custom design Super Agent, Cadence clients have seen significant benefits, with faster closure and higher layout productivity as seen in the example below for an Analog PHY. This technique can also automate the migration of an existing design to a new process node, speeding time to market and allowing senior engineers to focus on new product development.



Benefits and Risks: A Realistic View

As AI continues to expand across the design flow and beyond, engineering teams are already realizing productivity gains and increased throughput. AMD and Nvidia have already doubled their release frequency from two years to a single year, in part due to improved efficiency of their design teams utilizing AI equipped tools.



However, while AI can help tremendously, limitations still demand humans-in-the-loop to oversee and apply governance oversight. Our research indicates that AI design agents like

Cadence ChipStack AI Super Agent meaningfully speed up front-end design and verification, but are constrained by technical, methodological, and organizational limits that keep them as “force multipliers,” not replacements, for senior engineers and offload new-engineer bring-up.

1. Trust, correctness, and explainability:
 - Chip design cannot tolerate probabilistic errors; every automated decision must be traceable, which is hard for LLM-style agents that naturally “hallucinate” or generalize beyond their training data.
 - Even with Cadence’s “Mental Model” grounding, users still need to validate that generated RTL, assertions, or tests match actual spec intent, which limits how far you can push full autonomy on critical functions.
2. Coverage and corner cases
 - AI agents are very good at grinding through typical patterns (standard bus protocols, common micro-architectures, regression orchestration) but may miss rare corner cases or novel failure modes that weren’t well represented in historical data.
 - Verification sign-off still depends on coverage metrics, formal proofs, and human judgment about “did we think of the right scenarios,” so AI can accelerate test creation and triage but not fully guarantee coverage closure.
3. Specification and intent capture
 - Tools like ChipStack rely on structured specifications, design databases, and a consistent “truth layer” in the Mental Model; ambiguous, incomplete, or outdated specs directly limit what the agent can safely automate.
 - Translating fuzzy architectural intent, tradeoffs, or cross-IP constraints into machine-readable form is still a hard human task, so the quality of AI output is bounded by how well the design team maintains that intent model.
4. Scope of automation
 - Today’s ChipStack deployment focus is front-end design and verification workflows, not the full chip lifecycle; physical design, sign-off, and package/system co-design remain heavily tool-driven but not fully agentic.
 - Complex system-level tradeoffs (architecture partitioning, power-performance-area decisions, safety/security features) still require experienced humans; AI can suggest options but not take final accountability.
5. Integration, data, and IP constraints
 - To work well, AI agents must integrate deeply with existing EDA flows, CI, and design repositories; legacy or fragmented environments reduce the attainable productivity gains.

- Many companies have strict IP security and on-prem compute requirements, which constrain model size, training data aggregation, and telemetry, limiting how “smart” and adaptive the agent can be in practice.
6. Security and adversarial risks
 - Research shows AI tools can be abused to insert subtle hardware Trojans or security vulnerabilities into designs or code, making automated code and constraint generation a potential attack surface if governance is weak.
 - Any agent that can modify RTL, constraints, or testbenches at scale must be wrapped in strong access control, auditability, and change-review processes, which adds friction and reduces “one-click autonomy.”
 7. Human factors and organizational limits
 - Over-reliance on AI suggestions can erode human skills and situational awareness, degrading expert performance over time if teams stop doing deep independent reasoning.
 - Cultural resistance, union of expert pride and justified skepticism (“I won’t sign off based on an AI’s assertion”), and the need for process re-engineering all cap how aggressively ChipStack-like tools can be deployed on real, high-risk tape-outs.

The Net-Net

Cadence’s acquisition of ChipStack provides a significant step forward in applying AI beyond physical design optimization with agents across the design flow with proper human experience and oversight. Each step of the process can be managed by an agent, which iteratively calls Cadence EDA tools and checks with the design team to ensure convergence to meet design goals. A good place for customers new to applying AI is to start in process node migration, helping teams build confidence in their new AI agent colleagues.

ChipStack can remove a lot of friction in front-end design and verification and plausibly deliver multi-X productivity on well-structured tasks, but it remains somewhat limited by verification assurance, spec/intent quality, integration/security constraints, and the need to keep humans firmly in the loop for architecture, sign-off, and accountability. These are primarily human-driven barriers which will be overcome as teams learn to handle these limitations and realize significant speed up and productivity gains.

While adoption hurdles remain, the ROI of applying AI across the design workflow are compelling and the teams will learn to trust but verify their new AI agents.

Disclosures: This article expresses the opinions of the author and is not to be taken as advice to purchase from or invest in the companies mentioned. My firm, Cambrian-AI Research, is fortunate to have many semiconductor firms as our clients, including Baya Systems BrainChip, Cadence, Cerebras Systems, D-Matrix, Esperanto, Flex, Groq, IBM, Infleqtion, Intel, Micron, NVIDIA, Qualcomm, Graphcore, SImA.ai, Synopsys, Tenstorrent, Ventana Microsystems, and scores of investors. I have no investment positions in any of the companies mentioned in this article. For more information, please visit our website at <https://cambrian-ai.com>.