



Lessons from One Year of Agentic AI Deployment

MA

Mehir Arora
Head of Engineering

LF

Luke Frey
Sr. Application Engineer

Motivation

Why agentic AI in chip design?

- Increasing design complexity and legacy IP layers
- Need for faster, more accurate data retrieval and decision-making
- Traditional human-driven prompting is inefficient

Mehir Trends (revisit)

- “Amdahl’s Law of AI Agents”
- Coding: Review burden is proportional to output → Human bottleneck
- Search Problems: Verification is constant, search space is exponential
 - Debugging, code gen under oracle feedback, formal specs
- Increasing Human $\longleftrightarrow$ Agent bandwidth
 - The operators won’t prompt the agents, the agent’s will prompt the operators

Transition in Thinking

From “Human Prompting Agent” → “Agent Prompting Human”

- Agents excel at understanding prompt engineering
- Agents can drive:
 - Focused information retrieval
 - Concise, accurate context for tasks
- Human role shifts from “telling the agent what to do” to reviewing and guiding
- Agent pedaling the wheels, human is just steering

Workflow Awareness

Agents need to be workflow- and workspace-aware, not just tool-aware

- Increasing model capability is not enough
- Encoding decision paths is more valuable than raw compute
- Agents must comprehend legacy IP and existing RTL structure

Layered IP Example (Visual)

- Diagram: legacy RTL + new modules + agent analyzing structure
- Show agent navigating layers, highlighting decision paths

Reusable Learning

Externalize knowledge for cross-project reuse

- Use knowledge banks to store learned skills
- Agents can apply prior project learnings to new projects
- Reduces re-learning, accelerates execution

Versatility Beats Specialization

- Single, general-purpose agent handles multiple workflows
- Broad capability and adaptability accelerate adoption
- Reduces maintenance overhead vs. task-specific agents

Human-in-the-Loop as a Feature

Human review and intervention is deliberate and controlled

- Plan mode → review agent reasoning
- Interruptibility → RCA & error handling
- Scoped control → builds trust and adoption

Autonomy Spectrum

- Autonomy is not a binary switch
- Spectrum ranges from:
 - Fully human-driven
 - Assisted agent
 - Semi-autonomous agent
 - Fully autonomous (rare in EDA)
- Show visual: slider or continuum

Key Takeaways

1. Let agents drive prompting; humans guide & review
2. Workflow awareness beats raw AI capability
3. Externalize knowledge to accelerate learning
4. Human-in-the-loop fosters trust; autonomy is gradual