WHITE PAPER

Planning Your AI Design Journey in RF EDA

How Keysight's Open Capture, Orchestrate and Accelerate Framework Helps Engineering Leaders Re-Imagine Their RF and Microwave Design Workflows

KEYSIGHT

Executive Summary

Artificial intelligence is moving from novelty to necessity in electronic design automation. For RF and microwave teams the pressure is acute: designs are more complex, packaging and chiplet options are multiplying, experienced engineers are scarce, and the institutional knowledge that drives competitive products still lives largely in the heads of a few senior people. The question facing engineering leaders is no longer whether to adopt AI, but how to plan a capability journey that delivers value without forcing a disruptive rebuild of the design flow.

This paper offers a planning framework built on three parallel themes — **Capture, Orchestrate,** and **Accelerate** — which Keysight EDA is enabling beginning with the Advanced Design System (ADS) 2027, RFPro, and RF Circuit Simulation Pro 2027 (“Nexus”) software releases. **Capture** makes expertise visible and reusable. **Orchestrate** takes routine work off the bench and scales it with agents. **Accelerate** expands the design space with faster physics and machine-learning optimization. Underpinning all three is decades of high-frequency design experience and correlation to real measurements. Crucially, these are not sequential phases: organizations can begin each or all of them today, using the EDA tools its engineers already trust.

The sections that follow summarize the three themes, the cost of standing still, the measurement foundation beneath them, the autonomy horizon that shows where orchestration is headed, a proven customer result, and a practical roadmap for capability planning.



Five Questions for your Engineering Team

1. If your three most senior RF engineers retire next year, how much of your design capability would walk out the door with them?
2. When your team explores a design space, do you evaluate one candidate — or a thousand?
3. How long before a new hire can run your standard characterization flow unaided: weeks, or an afternoon?
4. Are your AI investments locked in one vendor model, or can you bring your own and keep your IP?
5. What would it be worth stating what you want to design, and let the tool decide how?

The Cost of Standing Still

For decades, the RF design toolchain has been defined by the underlying technology. Silicon, GaAs, GaN, and PCB laminates each carry their own rigid, siloed pipeline — design, simulate, tape-out — rebuilt from scratch for every process. That model is breaking down. Modern products increasingly center on packaging, 3D-IC, and 3D heterogeneous integration, where a single design may span III-V, SOI, bulk silicon, chiplets, and laminate at once.

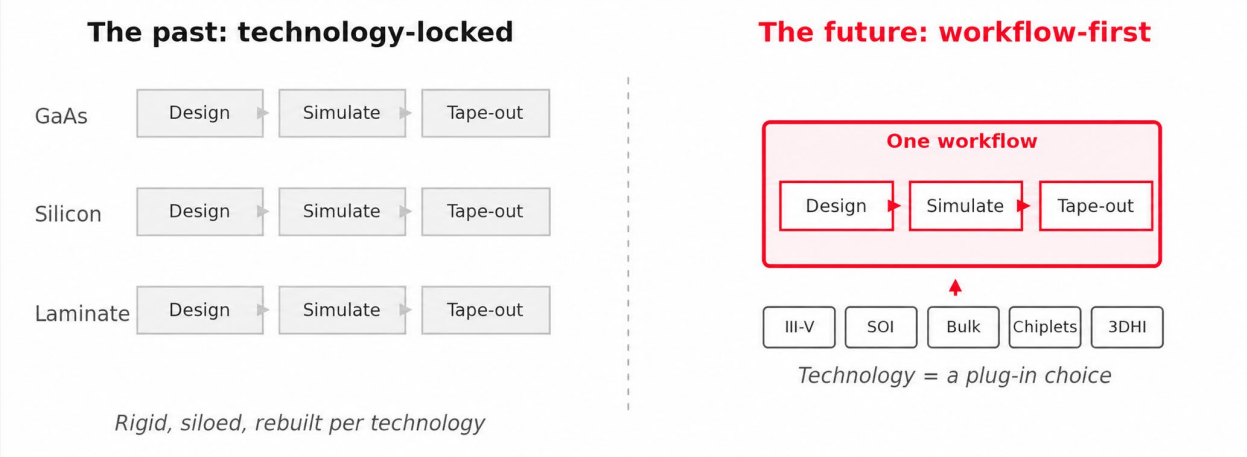


Figure 1. The shift from rigid, technology-locked pipelines to a single, technology-agnostic workflow. (Adapted from the webinar “Flow-First, Not Technology-First.”)

Re-imagining agentic workflows is how organizations will stay competitive as design complexity explodes. Standing still carries real, if unbudgeted, costs: design spaces explored one candidate at a time while competitors evaluate thousands; onboarding measured in months because workflows live in individuals rather than in shared, runnable form; and re-spins that a wider search would have avoided.

The Planning Insight

Design around workflows, not technology. A workflow captured once becomes an asset you can automate, share, hand to an AI agent, and re-target across processes — turning a rigid pipeline into an explorable design space that expands as physical domains are added.

A Framework for AI Capability Planning: Three Parallel Themes

Keysight EDA organizes its AI strategy around three moves, each tied to a business driver engineering leaders already track. Rather than a single monolithic "AI project," the framework lets an organization assess where it stands on each dimension and invest accordingly. The moves run in parallel and reinforce one another: captured workflows become the raw material for orchestration, and faster physics makes orchestrated exploration worthwhile.

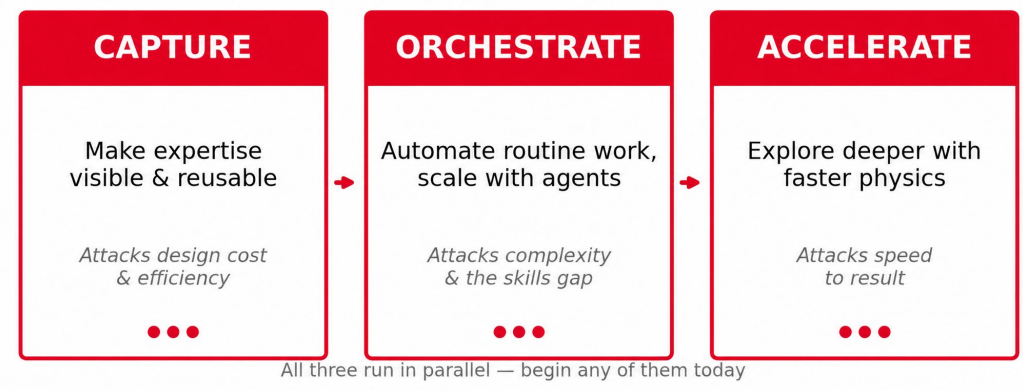


Figure 2. The three moves and the business driver at each address. All three can begin in parallel. (Adapted from "Move Your AI Program Forward in 3 Key Areas.")

1. **Capture —** attacks design cost and efficiency by encoding expertise and more physics into reusable, AI-ready form, so knowledge stops living in one engineer's head.
2. **Orchestrate —** attacks complexity and the skills gap by starting with automation and scaling with agents that take routine work off the plate.
3. **Accelerate —** attacks speed-to-result with better algorithms and surrogate models that cover far more design space at the same time.

A useful planning exercise is to rate your organization on each move independently — many teams are strong on automation yet have captured almost none of their expertise in shareable form. The gaps, not the averages, show where to invest first.

Capture: Making Expertise Visible and Reusable

RF design has resisted automation for a simple reason: large language models excel at code, but RF knowledge does not live in code. It lives in tools, in equations, and in the experience of senior engineers, with no public repository of any team's power-amplifier methodology or foundry IP for a model to learn from.

"Capture" addresses this directly by making soft IP and methodologies visible and composable as pythonic automation tools & skills, so they can be automated and handed to a trusted agent that sits *inside* your IP control.

The entry point is already on the bench. In ADS, an engineer can right-click any schematic or layout and **export Python**, producing a snippet that captures the workspace and can be transformed into a parameterized function — the cell and its parameters, ready to call from a script. Expert users can capture their workflows directly into python using the **python macro recorder**. Sequences of human action in the user interface, windows, and menus enable are converted into a python code stream that can be the basis for automation, or a custom design agent. Knowledge that was once outside the addressable space of AI now becomes code that can be streamlined for automation, shared, and run agentically across ADS, RFPro, and Nexus. These become genuine building blocks: write once, reuse everywhere, composed into workflows that define a plan, run EM, build the circuit, and optimize.

Scripts alone, however, do not transfer expertise — many workflows begin on a whiteboard and need something visual and shareable. **Nexus Connect** answers this with a visual workflow language built directly into the next-generation simulation environment. Engineers draw a simulation flow as a flowchart of inputs, tasks, conditions, and outputs, capturing critical design decision trees where the whole team can read them and a new engineer can run them. One click converts that flowchart (which is a visual test plan) into a documented, executable python script — making every captured workflow ready for AI agents.

These are important IP bridges that enable your best designers to achieve even higher levels of performance.

Orchestrate: From Automation to Agents

Once expertise is captured, orchestration puts it to work. With the ADS 2027 MCP Server, engineers no longer need to write the script at all — they can simply ask. The Model Context Protocol (MCP) lets a Large Language Model (LLM) discover and call the right tool on its own, acting as a single connection point that branches to every tool instead of hand-wiring each integration. Keysight MCPs will work with the LLM and coding agent(s) that your enterprise has already provided for security, performance, and cost. You are not constrained to one model, or even to English: the same engineering requests often work across languages, making Keysight's agentic surface accessible to global engineering teams.

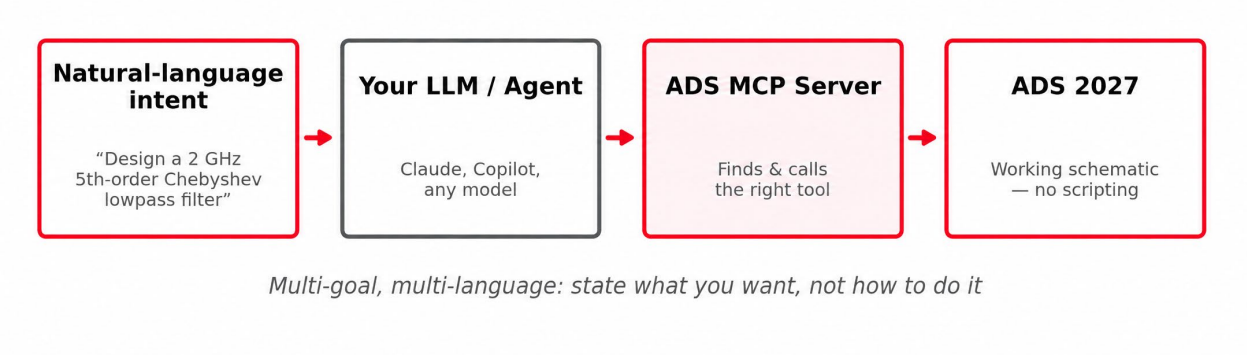


Figure 3. Natural-language intent flows through an LLM and the ADS MCP Server to help produce a working design — no scripting in between. (Adapted from the ADS 2027 MCP Server demo.)

The effect will be a shift in the fundamental culture of design, raising the abstraction from "Tools" to "Intent." An engineer connects an agent through the MCP server and states a desired outcome or requirement: "design a MMIC power amplifier with gain above 20 dB and output power above 28 dBm". The multi-goal definitions describe what is wanted, not how to achieve it. So, the key change is from procedure to outcomes, using plain language. Then protocol finds the right tools, and the tool stack does the work.

The Autonomy Horizon: Where Orchestrate is Headed

Orchestration is best understood against a longer arc. Keysight frames RF design as a five-level scale — a shared vocabulary for locating where a team sits today and seeing the path ahead. It is an autonomy scale, not a maturity score: most teams today sit between Levels 1 and 2, using pythonic automation, and experimenting with human supervision of copilots. The value is not where an organization stands, but that a clear, incremental path forward exists without rebuilding the flow to walk it. Today's natural-language agents move teams toward Level 3; the same orchestration foundation extends toward Levels 4 and 5.

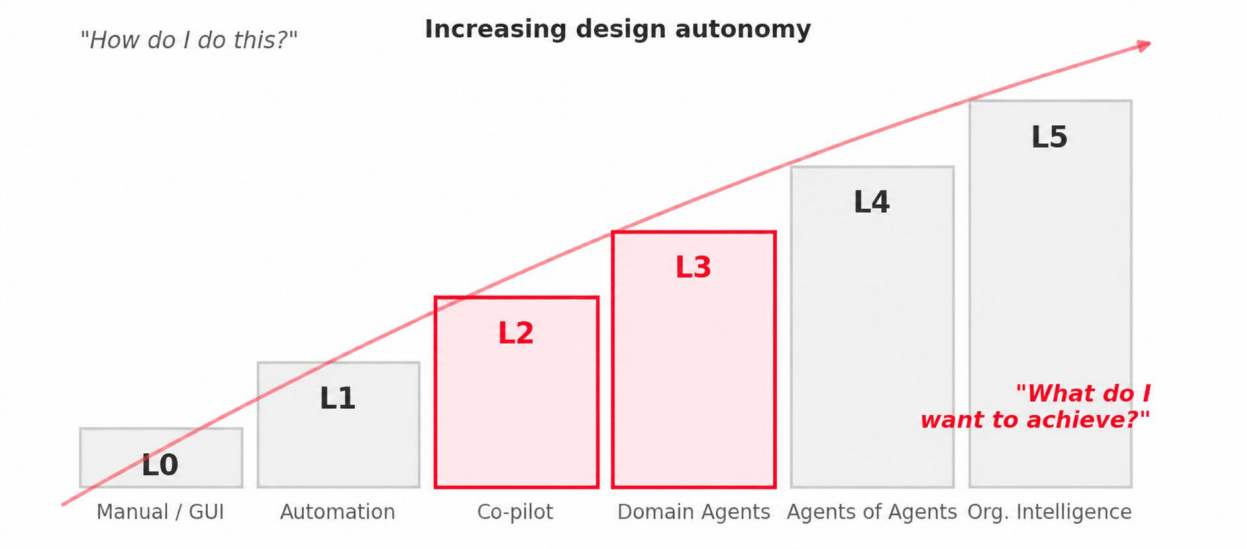


Figure 4. The five-level autonomy scale, from manual GUI work to organizational intelligence. Most teams sit at L1–L2 today; agentic features target L3 and beyond. (Adapted from "Evolution of EDA.")

Level	Stage	What it Means for the Engineer
L0	Manual/GUI	Engineers execute every step by hand using the graphical user interface
L1	Automation	Scripts and headless API's run execution; humans still initiate design tasks
L2	Co-pilot	Natural-language assistants grasp intent while engineers supervise
L3	Domain agents	Specialized agents handle multi-step tasks with partial delegation
L4	Agents of Agents	Orchestrated agents coordinate whole workflows; humans set objectives
L5	Organizational Intelligence	Learning organizations that self-improve; Continuous, always-on operational intelligence

Accelerate: Physics at the Speed of Exploration

The third parallel AI theme brings faster, more realistic physics into the creative phase of design so teams can explore deeper, larger problem spaces. Two engines drive it. The first is **AI-enabled optimization**: a reinforcement-learning framework connects an agent to existing simulators, setting goals and earning rewards to learn effective design strategies. Where conventional optimizers are typically applied to in the order of ten parameters, the ML-enabled optimizer is designed to handle forty or more at once and to learn from the data as it goes. Just as important, the framework is open — teams can bring their own Python optimizer or surrogate and combine it with continuous learning and warm starts.

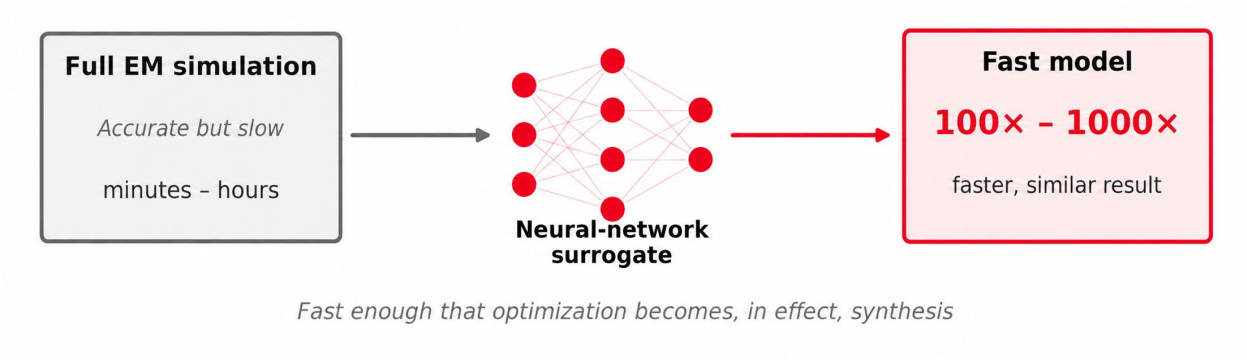


Figure 5. A surrogate model replaces a slow EM simulation with a fast neural network approximation. (Adapted from "Acceleration Through Surrogate Modeling.")

The second engine is **surrogate modeling.** A surrogate is a faster approximation to a computationally expensive simulation (such as a full EM run that might take minutes or hours), which might use a compact neural-network model. On representative structures such as spiral inductors, Keysight reports speedups of two to three orders of magnitude. But raw speed is only valuable if the answer can be trusted, which is where Keysight's measurement heritage matters. Physical domain simulation speed becomes unlocked when a design evaluates in milliseconds; an organization can explore thousands of candidates and let agents search the space. Through pySMI, Keysight's published Smart Modeling Interface, a team's own surrogates, Keysight's, or a third party are all used the same way — in context, inside the existing flow. Such an open system where you can "**Bring Your Own Surrogate**" means you are using models your team already trusts, and you can create your own enduring competitive advantage using your own proprietary IP.

The Foundation: Physics You Can Trust

These three themes rest on a foundation that is easy to overlook and hard to copy. Any vendor or talented individual can wrap a language model around a simulator; the differentiator is whether the underlying physics is right, and can a workgroup share them in an organizationally consistent way. Keysight's simulation models are grounded in decades of correlation to real hardware measurement — the same measurement science that underpins the instruments on the bench. At the same time, having a unifying platform such as Keysight ADS is necessary to rein in the freedom and innovation of personal modeling to create *organizational* benefits that scale, remain manufacturable, and transcend the risk of relying on only one expert or technology.

For a planning conversation, this reframes the question of Trust. The risk in agentic design for RF is not that the tools are slow — it is that a fast, confident answer might be wrong, and that these silent faults will not be traceable in a Root Cause Analysis. "Grounding the loop" using measurement-correlated physics is how an organization can safely delegate more of the design to agents over time. When comparing vendors, the sharper question may be "how well does your physics correlate to measured hardware?" due to packaging, interconnects, coupling, power and ground integrity, moding, launches, 3D currents, and other effects. Furthermore, the data lineage of where models originated, how the data was tagged or cleaned, and limitations associated with AI/ML artifacts should be versioned and inspectable. Visibility and benchmarked evaluations also build Trust.

Proof Point: AI-Driven RFIC Design at Scale

These capabilities above are not aspirational, and the pattern behind them is repeatable. Sphere Semi, an RFIC design company, faced the classic wall: RF and analog design are manual and talent-constrained, exploring one design at a time. Their response — a fully code-defined, Python-based flow in which an AI engine generates hundreds to thousands of candidates, evaluates each with circuit and EM co-simulation in ADS, and closes the loop — is exactly the generate-simulate-rank-refine pattern the three Keysight AI themes are designed to enable. While this organization succeeded at the time using the Python automation of a previous release of Keysight tools, new generations of Keysight users can take advantage of newer AI features from Keysight to create their own customer AI-driven workflows.

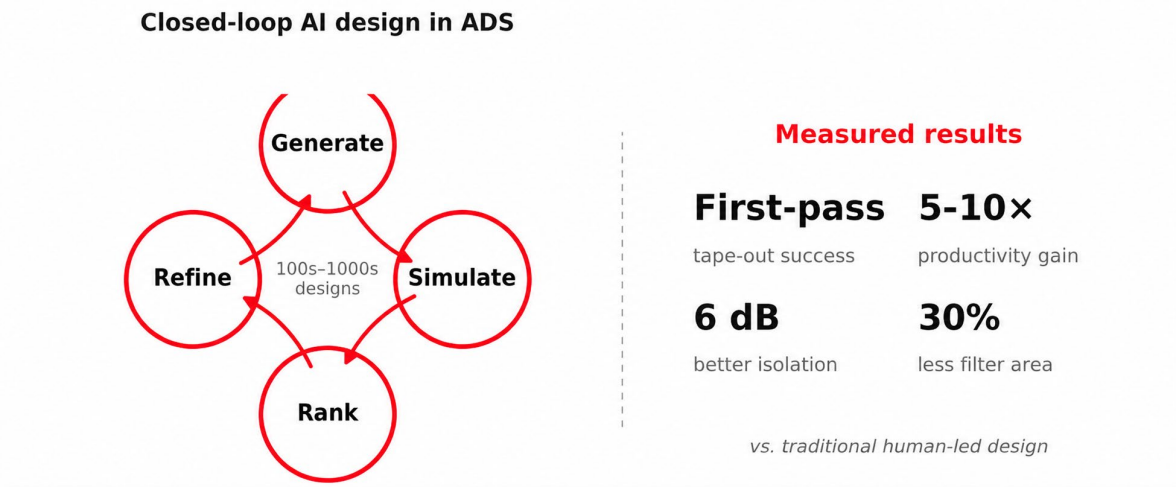


Figure 6. The closed-loop AI design pattern and Sphere Semi's measured results versus a traditional human-led design. (Adapted from the Sphere Semi case study.)

In the Words of Sphere Semi

"We were thrilled to learn that the ADS 2026 release reinforced Python extensibility and the API for its programmatically callable co-simulation engine, which made ADS a clear winner for us — there was nothing else with the combination of features we needed."
— Mitchell Nahmias, CTO, Sphere Semi

A Planning Roadmap for Engineering Leaders

Because the three AI themes run in parallel, capability planning is less about sequencing and more about deliberate, simultaneous investment aligned with the status of each team. The horizons below offer a way to structure that investment without stalling on any single dimension.

Horizon	Planning Focus	Where to Invest
Now	Make expertise visible	Export designs to Python; document flows in Nexus Connect. (Capture)
Near term	Take routine work off the bench	Stand up natural-language agents on the ADS MCP Server (Orchestrate)
Next	Expand the design space	Adopt RL optimizers and surrogate models; bring your own. (Accelerate)
Ongoing	Compound the advantage	Govern data and models at enterprise scale toward higher autonomy

Three principles should guide the plan.

1. First, be workflow-first: capture the engineering process as technology-agnostic actions before optimizing any single technology.
2. Second, stay open: a bring-your-own approach to models, optimizers, and surrogates protects the organization's freedom to combine its own IP with Keysight's and avoids lock-in.
3. Third, treat learning as the asset: as surrogates and workflows accumulate, the discipline of versioning and governing data and models — not just the models themselves — becomes the enterprise's durable competitive advantage.

Conclusion: Start from where you are

Everything described here builds on tools RF teams already run — ADS, RFPro, Nexus, and Python — and every building block, from knowledge capture and the MCP interface to RL frameworks and surrogates, lives in one open, integrated platform grounded in measurement-correlated physics. The first step is modest, and none of the three themes requires finishing another first. The real question for engineering leaders is not whether their organization will adopt agentic design, but what their workflow will look like in three years — and how much of that future they choose to start building today.

Take the Next Step

Request a free software trial of ADS and Nexus, explore the Python automation and AI capabilities firsthand, and read the full Sphere Semi case study on how a custom, end-to-end AI system multiplied productivity. The journey toward higher design autonomy is well underway — and it begins wherever your teams are today.

Webinar: Future Directions in AI/ML for RF

Sphere Case Study

Nexus Connect Solution Brief

Nexus free Software Trial

ADS free Software Trial

As a mission-critical design enablement partner, Keysight connects design, emulation, and test so customers can solve the world's most complex engineering challenges. Accelerate innovation at www.keysight.com.

KEYSIGHT

This information is subject to change without notice. © Keysight Technologies, 2026, Published in USA, August 10, 2026, 3126-1495.EN