SPONSORED CONTENT | WHITE PAPER



Moving from AI experiments to enterprise impact

Google Cloud and Slalom help organizations transform work by embedding AI across data, systems, workflows, and teams.

CIO

SPONSORED BY slalom | Google Cloud

MANY ORGANIZATIONS are struggling to recognize long-term business value from their AI investments, but that's not because their AI is stalling on launch day.

More commonly, problems emerge weeks later, when people open new tools and try to put them into action for the first time. Too often, employees find that these new AI solutions are disconnected from the data, systems, and processes they rely on to do their jobs. To achieve real value with enterprise AI, organizations need to move beyond standalone tools, embed AI capabilities into real workflows, and connect these capabilities to enterprise data, systems, and people.

This shift requires more than engineering. Many software companies are deploying engineers with new technology purchases, and while this model helps organizations stand up impressive prototypes, its utility is typically more limited for pushing AI into full production across the enterprise.

After a prototype is rolled out, organizations need cross-functional expertise to turn it into a working business capability. In a Forward Deployed Expertise (FDX) model, embedded teams combine an understanding of the organization's business processes with technical expertise to connect AI solutions to existing workflows and provide solutions to real-world problems. These teams also help test solutions in production, refine them based on real-world use, and support employee adoption.

These three lessons from the field illustrate how forward-deployed teams are helping organizations achieve enterprise value across AI models, agents, and experiences delivered through Gemini and Gemini Enterprise.

Lesson #1:
Start with the process, not the model



Artificial Intelligence

Frontier models provide the intelligence that underpins AI applications like chatbots and agents. But access to a powerful model alone does not help leaders decide which business problems to solve, where AI should enter existing processes, and what enterprise data and systems must be incorporated with AI tools.

The real design work lies in deciding which steps need judgment, and which should follow hard-coded business rules. Decisions that require judgment over unstructured information are often a good fit for generative or agentic AI, while those that are determined by explicit business rules are often better handled by code or core systems. For workflows that require both, an agent can interpret the request and then call on appropriate systems to execute rules-based tasks.

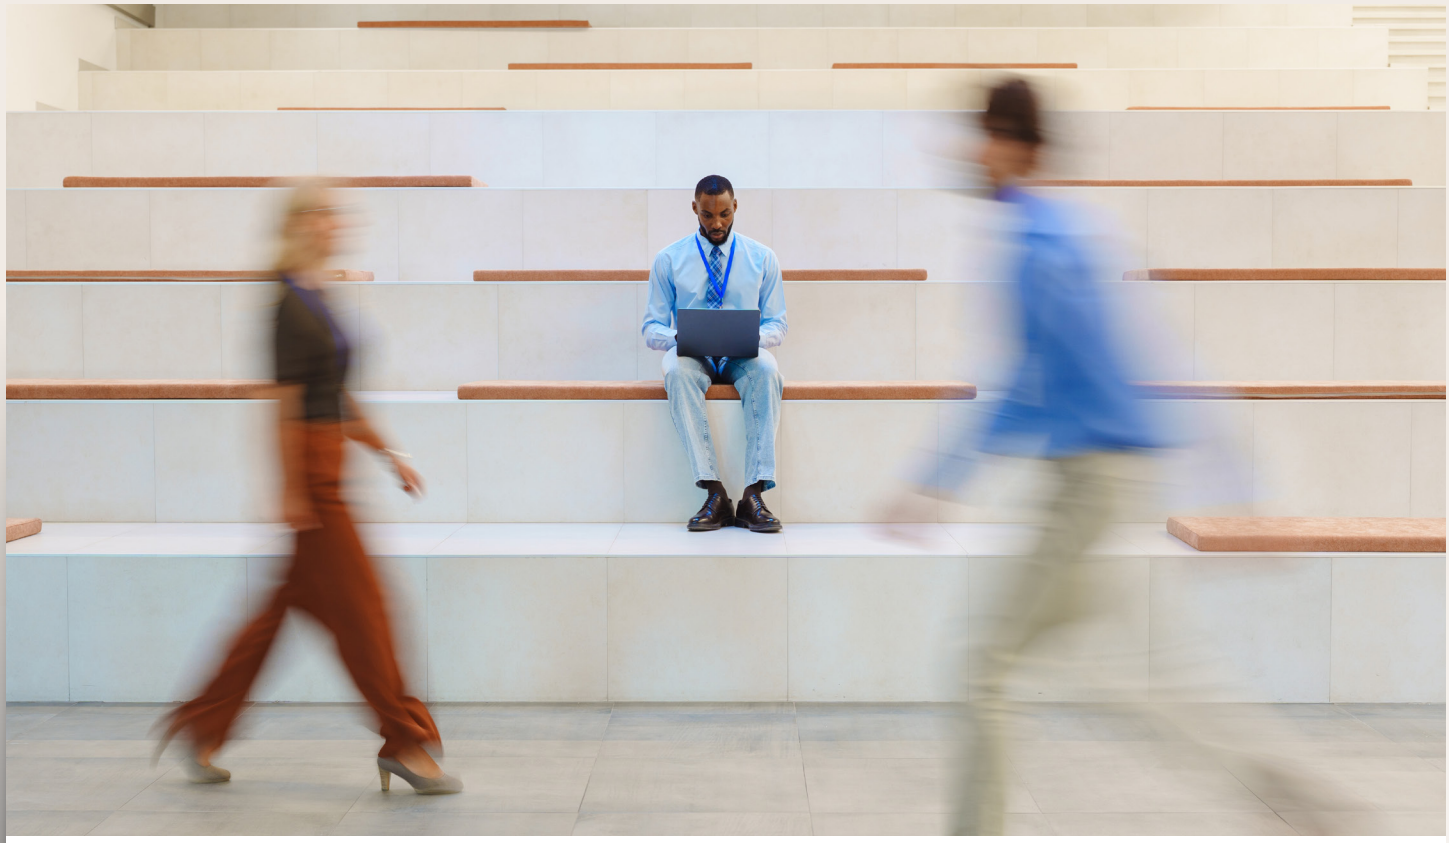
Before handing processes off to AI tools, organizations should map these processes step-by-step.

Getting the design work right is becoming even more important as organizations embrace "tokenomics" principles designed to control ballooning AI costs. Previously, many organizations defaulted to frontier AI models, with little consideration given to expense. Today, though, leaders must ensure that they are using not only the right tool for the job, but also the most cost-effective method that will get the result they want.

Real-world results

Slalom recently worked with one large manufacturer to leverage Google's frontier model for a new AI solution that streamlines warranty claims. Previously, the organization used large offshore teams to apply business rules to assess whether claims should be approved or denied. With a multidisciplinary team, Slalom built a group of AI agents to support the claims process, combining deterministic and generative agents.

Lesson #2:
Build smaller, reusable AI agents



Artificial Intelligence

Organizations are increasingly looking beyond built-in generative AI tools like chatbots or agents and are instead seeking to build their own agents that can make decisions and take actions across multistep workflows. Early agentic systems often simply recreated old software patterns as large, monolithic agents. While these are interesting as pilots, they are often difficult to test, govern, reuse, and maintain.

Rather than building large, complex agents for each workflow, Google and Slalom are focused on developing "atomic agents." These are smaller agentic components that each handle one small, discrete task and can be composed into larger workflows.

One major benefit of atomic agents is that they can be reused across different departments and business processes. For example, an agent that extracts clauses from contracts can be used to support functions like legal review, procurement, and compliance, rather than being rebuilt anew for each application. Atomic design also makes it easier for teams to isolate failure and apply controls at the level where work is performed.

To create an ecosystem of atomic agents that will deliver maximum value

across the enterprise, organizations should first look for tasks that repeat across departments. Agents should be assigned an owner, a registry entry, and retirement criteria before they multiply. Agents also need access to current, relevant enterprise data with surrounding context, but organizations must grant this access while preserving existing permissions.

The engineering work is often not the most difficult part of building atomic agents. Organizations need experts who deeply understand the roles and workflows of business units, along with the technology. These experts can use this knowledge to determine what work should be handled by agents, break that work down to an "atomic" level, build the agents, and construct repeatable, reliable agentic workflows.

Real-world results

Slalom has worked with companies to build libraries of 250 atomic agents across operations, HR, finance, and IT helpdesk. At one large energy provider, voice-enabled agents built by Slalom are supporting crews that maintain remote grid infrastructure, resulting in millions of dollars in operational savings.

Lesson #3:
Design AI around the people who will use it



Artificial Intelligence

There is no such thing as a one-size-fits-all rollout plan for AI.

Admins need governance, power users need to build agents, everyday users need to know when to use which tool, and frontline users often simply need access to one reliable agent on their phone. To successfully deploy AI across their organizations, leaders need to carefully consider adoption from the outset of their efforts. This may mean embedding AI capabilities into familiar applications or designing new AI solutions around how employees already work. Successful deployment also requires training, communication, and governance.

It is important for organizations to carefully track usage of new AI applications alongside the usage of the tools they are replacing, as well as the performance of the AI solutions. Rising AI adoption means little if legacy tool usage isn't declining, or if the new tools are ineffective. A production rollout should include testing of performance, monitoring of results, gathering of user feedback, and adjustments over time.

Retire legacy tools only when the data says the organization is ready, rather than adhering to a rigid date.

Real-world results

A global retailer used Gemini Enterprise to take AI well beyond a single team. Working with Slalom, 16 departments identified 301 use cases grounded in their own day-to-day work. **Employees now get back roughly 111,000 hours a year, and the program has returned more than $9 million in value.**

Creating customer-facing AI experiences

PUMA and Manchester City wanted fans to do more than react to a new jersey, so they invited fans to design it.

Slalom built the AI backend and infrastructure on Google Kubernetes Engine for a browser-based app where fans described a design in a short prompt, refined it, and submitted it for the club's 2026/27 third kit. Fans in 206 countries created 180,000 designs and cast 1.7 million ratings.

Slalom's engineering team cut generation time from 60 seconds to 17 and nearly halved the number of GPUs required, keeping the experience fast even under more than 10,000 simultaneous requests.

Slalom + Google =
Expertise across the AI ecosystem

Google gives organizations a connected path from advanced models to agentic workflows and enterprise-class AI experiences. Slalom supplies the embedded, cross-functional expertise needed to put that technology into action for the business.

An AI initiative should be staffed in much the same way an organization would staff a product launch. While engineers can stand up a working prototype, full-production AI also requires professionals who own product decisions and someone who owns adoption. Most programs are understaffed on these roles.

The standard Slalom FDX model features a four-person squad, with a product owner/tech lead, a senior AI engineer, a full-stack cloud architect, and a customer experience lead, with this team sized to fit the scope of the engagement.

During these engagements, FDX teams prove the use case, put agents into live production, promote internal adoption, scale platform resources to support growing use, and manage tools and measure value to accelerate future projects. Each successful deployment gives the next team a stronger starting point, with field-tested components, a clearer understanding of users' needs, and concrete evidence of where AI can deliver the most value to the business.

With Slalom's FDX model supporting Google's frontier models, agents, and experiences, organizations can identify important problems, design reusable agents, and help employees adopt new ways of working.

Turn AI potential into enterprise impact with Slalom and Google Cloud.