

● Executive Roundtable Takeaways | Chicago, May 2026

From Pilot Purgatory to Production: What AI Leaders Are Really Thinking

Evalueserve brought together senior executives from Life Sciences, Healthcare, Logistics, and Manufacturing for an evening of candid conversation about what it actually takes to scale AI inside large, complex organizations.

AbbVie · Abbott · Amgen · Astellas · Bayer · BD · Danaher · GE Healthcare · GSK · JCI · Organon · Pfizer · Philips · Regeneron · RXO Inc

Introduction

On a lovely May evening in Chicago's Fulton Market District, Evalueserve hosted a gathering that felt a little different from the average AI conference. No vendor theatrics, no scripted demos, just twenty-plus senior leaders from some of the world's most recognizable companies, sharing what's actually working, what's blocking them, and what keeps them up at night when it comes to enterprise AI.

The guest list spanned a wide range of functions and seniority levels: AI strategy leads, commercial insights managers, directors of market access and data science, chief marketing officers, clinical ecosystem leads, and technology executives, all representing organizations at the forefront of AI adoption across Life Sciences, Healthcare, Logistics, and Manufacturing.

The dinner, moderated by Evalueserve's Gabe Keeler, featured Clara Buenker, Google Cloud AI GTM Lead, as the evening's keynote voice. Together they guided the room through a wide-ranging conversation that moved from cultural roadblocks and security fears to multimodal use cases and the future of the workforce.

Five Key Themes

1. The “Pilot Trap” is real, and leaders know it

The evening opened with a telling ice-breaker: what's one task you'd hand off to an AI agent tomorrow, at home? Answers ranged from planning family vacations to using AI to alter and enhance family photos, setting up competing LLMs head-to-head just to see which performs best, and mapping out full landscaping projects. The room was clearly full of active, curious AI users.

But Buenker was quick to draw the contrast. Trusting an AI to plan a long weekend is a very different proposition from trusting it to manage a clinical trial pipeline or a global logistics network. That gap, between consumer convenience and enterprise accountability, framed much of the evening's discussion.

The dominant enterprise challenge, she argued, isn't building AI. It's about escaping what she called "Pilot Purgatory": organizations sitting on hundreds of proofs of concept that never make it into production workflows where real ROI is captured.

"The conversation has completely shifted from 'How smart is the model?' to 'How do we manage it?'"

Any solution needs to be tested not only at the technical level, but at the people level, to ensure it actually gets adopted

"The conversation has completely shifted from 'How smart is the model?' to 'How do we manage it?'" says Buenker. The prescription she offered: treat AI agents like products, not experiments. Build them deliberately, publish them to a shared internal gallery where employees can discover and use them, and wrap them in governance, including role-based access controls, audit trails, and data loss prevention, from day one.

Evalueserve's Gabe Keeler added a practical framework for driving that adoption from the ground up, pointing to the value of structured cohorts as a bridge between experimentation and scale. "Adoption cohorts, small groups that use AI together and discuss what works and what doesn't, help build the muscle. But any solution also needs to be tested not only at the technical level, but at the people level, to ensure it actually gets adopted."

2. Shadow AI and cultural risk aversion: the unspoken tension

Before the dinner, several attendees submitted candid reflections on their challenges. One from a senior leader in pharma landed squarely at the center of a broader conversation: "My biggest roadblock is cultural. My company is very risk-averse... AI adoption is not incentivized. The only AI tool available for all employees is Copilot at the moment, with all others blocked by VPN."

This resonated across the table. When asked for a show of hands on whether their teams were likely using unvetted consumer AI tools, the response was telling. The phenomenon, employees quietly pasting proprietary data into public chatbots because approved tooling isn't keeping pace, was widely recognized.

Buenker's advice for AI champions in risk-averse organizations: stop framing the ask as "we want a new tool" and start framing it as "we need a governed platform." Build a Center of Excellence that partners with IT rather than circumventing it. Target the tasks that everyone agrees are painful and administrative. And prove value first on synthetic or low-sensitivity data before asking permission to touch regulated information.

A separate concern surfaced around quality control: as organizations push pilots into production, many are doing so without adequate QC processes in place. The rush to ship is outpacing the rigor needed to ensure outputs are accurate, safe, and fit for purpose.

3. Security and trust: Google's acquisition of Wiz enters the conversation

Not every attendee arrived as a Google Cloud convert. A thread of skepticism ran through parts of the room, with some companies remaining cautious about trusting any major cloud provider's AI solutions with sensitive enterprise data.

Buenker addressed this directly, pointing to Google's acquisition of Wiz as a meaningful signal of its commitment to enterprise-grade security and governance. The message: this isn't a cloud AI product layered on top of consumer infrastructure. It's a deliberate build-out of the security, compliance, and data governance architecture that regulated industries actually require.

The core commitment she returned to repeatedly: customer data is never used to train public models. Combined with proactive data loss prevention and granular role-based access controls, the platform is designed to bring IT and legal stakeholders into the conversation, not alarm them.



4. Where Gemini's multimodal edge actually shows up

A senior marketing leader in pharma put it plainly: they were more familiar with Claude and ChatGPT than Gemini and wanted to understand what was actually possible. Buenker's answer focused on native multimodality, the ability to process text, images, audio, and video simultaneously, as the differentiator that matters most in the industries represented at the table.

In pharma, R&D teams are using Gemini to simultaneously analyze visual chemical structures alongside scientific literature, a task that previously required toggling between entirely separate tools and workflows. In logistics, warehouse operators are cross-referencing live camera feeds with shipping documents in real time. In manufacturing, a supervisor can point the AI at a camera feed and ask it to diagnose a stalled robotic arm by cross-referencing the machine's manual.


The pattern across all of these: multimodal agents aimed directly at the organization's most expensive operational bottleneck.

"When everyone has the exact same AI baseline, your differentiators become your proprietary data, your cultural velocity, and your taste."

5. AI won't replace your experts, but their roles will look different

The workforce question came up organically, and Buenker didn't shy away from it. Her framing: AI automates tasks, not relationships. As the technology absorbs routine cognitive work, the premium on human judgment, empathy, and strategic taste should actually increase.

One anecdote from the startup world sparked particular interest: a founder whose top salespeople and customer success managers are evolving into what he called "hybrid forward-deployed product managers," people who don't just relay customer requests to engineering but can now build features themselves and ship exactly what customers need, instantly. Roles as we know them, she suggested, are going to look fundamentally different within a few years.



Looking Ahead

The evening closed not with tidy answers but with the questions that are actually keeping leaders busy. How do you maintain agents as regulations, data sources, and governance requirements evolve? How do you get AI genuinely embedded in a workflow rather than bolted onto individual users? How do you track qualitative AI adoption in a meaningful way? And perhaps most urgently: as pilots move into production, who is responsible for quality control, and what does rigorous QC even look like for AI-generated outputs?

These are exactly the questions worth gathering around a table for. Evalueserve looks forward to continuing the conversation.

Thank You to Our Guests

Evalueserve's Executive Exchange events provide an unparalleled opportunity for senior financial services executives to connect, share insights, and explore innovative ideas in a relaxed setting over dinner and drinks. Our exchanges feature engaging expert-led discussions on trends such as digital transformation and AI.

These intimate, invitation-only events are carefully curated to bring together leaders in banking, lending, and private capital in global financial hubs including New York, Toronto, London, and more.

[Learn More](#)

