HARD TRUTHS: PART 1
EXECUTIVE SUMMIT

A Beaming Future
Over 100 of the world's leading executives gathered not to celebrate what they've built, but to confront what comes next.

LEADING THROUGH THE HYPE

At the 2026 Insigniam Executive Summit, executives separated signals from static.

By Ryan Jones & Photography by Christopher Ginn

In 1899, Charles H. Duell, Commissioner of the U.S. Patent and Trademark Office, reportedly declared that everything that could be invented had been invented. The quote has been disputed by historians, but the underlying impulse—the deeply human tendency to look at the current state of the world and mistake it for the permanent one—has not. It's the same impulse that led AT&T engineers in 1993 to correctly predict GPS navigation, tablet computing, on-demand streaming, and video calls in a famous advertising campaign, and then fail to bring a single one of those products to market. It's the same impulse that leads organizations today to sit on the most powerful general-purpose technology in history and use it to write emails faster.

CHRISTOPHER GINN

Thomas M. Koulopoulos
Futurist, Board Chair, Author & Keynote Speaker

That was the opening provocation of the 2026 Insigniam Executive Summit, held April 21 at The Westin Philadelphia, setting the tone for a day that had little patience for comfort.

Over 100 senior executives gathered for the Summit's highest-attended event yet, structured this year as a venture capital-style forum: live technology demonstrations, industry deep-dives led by founders and investors, and the kind of frank executive dialogue that polite conference rooms rarely produce. The theme, Leading through the AI Hype: Delivering Hard Results, was not aspirational. It was a verdict.

"Part of our philosophy," said Nathan Owen Rosenberg, Insigniam co-founding partner and Elixirr partner, "is that it would be inauthentic not to practice our own methods on ourselves. We are in the same place that you all are, working our way through this incredible new technology." That candor was the last moment of comfort the room would enjoy.

Decoding the Signals Ahead

Thomas M. Koulopoulos, founder and chairman of the Delphi Group, author of *GigaTrends*, and one of the more reliable discomfort-delivery systems in American business thinking, opened not with a slide but with a question: how many people in the room had checked their phones before getting out of bed that morning? Nearly every hand went up. "This is not a technological issue," he told them. "That device has become your oxygen mask." Within five years, he predicted, an AI personal agent would be woven so deeply into every executive's daily navigation of the world that losing signal would feel like losing the ability to function.

The point wasn't the device. The point was the behavior. And behavior, Mr. Koulopoulos argued,

Standing for Leadership
Insigniam co-founding partners Nathan Owen Rosenberg (above) and Shideh Sedgh Bina (top right) challenged executives to stop treating AI as a project—and start building the corporate infrastructure that will matter long after the hype fades.

is where organizations consistently and fatally fall behind. The AT&T story was instructive: engineers at Bell Labs knew exactly what 2026 would look like in 1993. GPS, tablets, streaming, video calls: all of it was in the ads. And yet not one of those products came from AT&T. "How can you predict the technology so well," he asked, "and miss the behavioral trajectory by such an enormous margin?" The answer: the engineers thought about technology as a feature. They could not imagine how quickly behavior would embrace it and make it indispensable. Organizations are making the same mistake today.

"It's not enough to just know the trajectory of technology," he told the room. "That's the easy part. The hard part is predicting behavior."

From there, Mr. Koulopoulos traced the full sweep of computing history, from the 18,000 vacuum tubes of the ENIAC—one of the world's first computers—built in 1945 through the exponential curve of devices that followed, to establish the trajectory. From 1960 to the present day, the number of user computing devices has grown by one order of magnitude per decade. The projection for 2100: one sextillion devices, or 1,000,000,000,000,000,000,000. For reference, that is the computing equivalent of 666,666 IBM 350 disk drives held between two fingertips, or 400 million ENIACs in your pocket.

"By the way," he said, "the reason you should believe that number today is exactly the reason you would have thought me insane if I had made that prediction in 1950."

Beneath the scale of devices lies a more urgent infrastructure reckoning. At current growth rates, the amount of data generated globally is on course to exceed the number of atoms in the Earth, and eventually those in the solar system.

CHRISTOPHER GINN

More immediately: a study Mr. Koulopoulos conducted for the U.S. utility industry found that by somewhere between 2040 and 2050, data centers alone will consume all available energy on the planet. "The technologies that got us here," he said flatly, "simply don't scale to the future we need to build."

This set up his central framework: the shift from infrastructure to intelligence, across three phases.

The first, already underway, is the agentic age: AI that exercises judgment, identifies resources, and pursues goals without step-by-step instruction. Unlike the programmatic automation of the past 60 years, agentic AI operates with autonomy, and Mr. Koulopoulos made the counterintuitive case that this autonomy is precisely what leaders must learn to embrace.

"If I build an organization that is AI-native, AI-first, and I have perfect fidelity from my agents, I'll never have innovation." The latitude to maneuver, the deviation from strict instruction, is where progress lives.

Phase two, arriving in the 24-to 48-month window, is recursive AI: systems that design and deploy other systems. The agentic cabinet minister in Estonia charged with combating government fraud, who determined the scope was too complex for a single agent and autonomously built a constellation of specialist sub-agents, offered a glimpse of what this looks like in practice. As recursive AI scales, its internal logic will become opaque, a black box whose outputs leaders can evaluate empirically but not fully decode. "We will trust these technologies with aspects of our lives," he said, "and have no idea what they're actually doing." The analogy was precise: it will be like trusting Waze. After one month of being wrong every time he second-guessed it, Mr. Koulopoulos said, he reached a point where he would follow

it to the gates of hell. "That's what we'll be doing with AI."

The warning this generated was his most urgent: the window for putting governance frameworks around agentic AI is 12 to 24 months. After that, the recursive pace of self-improvement will outrun the ability to install guardrails retroactively. The risk slope is rising. The value slope must rise faster. "There has never been a worse time in history to wait for clarity."

The room he invoked to close was not a boardroom but a napkin at an Italian restaurant, some 30 years ago, where Peter Drucker sketched out what he saw as the single most important shift of the coming century: the move from a product-and-ownership economy to a strategy-and-experience economy. Everything held together by strategy. Everything delivered as experience. What struck Mr. Koulopoulos in 2026 was how precisely that napkin sketch described the competitive landscape AI-native organizations are now building, and how completely it renders the product-and-scale organizations of the past century obsolete.

Another Drucker maxim landed with particular weight in the room. From a conversation at Drucker's home in Claremont in 2003, Koulopoulos recalled the management thinker's summary of

"It's not enough to just know the trajectory of technology. That's the easy part. The hard part is predicting behavior."

—Thomas M. Koulopoulos
Delphi Group Chair & Author, *GigaTrends*

the worst mistake he observed in organizations: "The best people are assigned to keeping yesterday alive a little longer." The room was quiet.

"That's the behavioral problem," Mr. Koulopoulos said. "Not the technology."

The 95% Paradox

If Mr. Koulopoulos established the historical reckoning, Adam Hofmann, Elixirr Partner and AI lead, arrived with the current ledger.

The headline: 95% of enterprise AI pilots fail to reach production scale. The footnote that makes it a paradox: AI-native companies built from the ground up are scaling faster than anyone has seen, with fewer people, compressing what used to take decades into months. Both things are simultaneously true. The question Mr. Hofmann had come to answer is *why*.

He opened with a chart showing the pace of AI model releases, 255 in the first quarter of this year alone, alongside the capability curve, which keeps bending steeper. The time it takes a current model to complete what would typically be a 12-hour engineering task, successfully at least half the time, continues to halve. "In 12 months, AI will be 10 times more capable than it is today," he told the room. "Does anybody disagree?" Silence. "So the question for you is: what does it take for you to be twice as effective?"

He then made the gap between capability and adoption tangible. Live. Using Claude, Mr. Hofmann showed what genuine agentic work looks like for a senior leader: synthesizing a quarterly business review presentation, reading a spreadsheet, catching a formula error the human team had missed, building an analytics dashboard from scratch, and, most pointedly, configuring an autonomous agent to check a calendar, locate the relevant data file in email, produce the presentation, schedule review

CHRISTOPHER GINN

No Hype, Only Possibilities

From keynotes to industry deep-dives to candid table conversations, every moment of the day was deliberately designed to turn insight into action. Executives didn't just leave with frameworks and new ideas—they left with the kind of peer relationships and shared perspective that only come from being in the same room, wrestling with the same questions, at the same moment in history.

On Topic, Off Record
Many of the conversations that happened between sessions proved just as valuable as the ones on stage.

time, and draft a summary, all without being re-prompted at each step. A task that would have consumed a team for at least a day. Claude completed it in seven minutes.

The audience had seen AI demos. What landed differently here was the shift in agency. Not AI assisting humans. AI executing, while humans set direction. That shift, from assistant to executor, is precisely what the 95% of failed pilots miss. Mr. Hofmann catalogued the failure modes by name. The bolt-on trap: AI grafted onto existing workflows without questioning whether those workflows should exist at all. The innovation island: a central AI team that becomes a bottleneck, not an accelerator. Governance gridlock: oversight so heavy at the front that nothing ever reaches production. And pilot purgatory: experiments multiplying without ever scaling, creating a tax on the organization rather than a return. What connects these failures, the data says plainly, is not the technology. A stunning 84% of AI project failures are driven by leadership decisions.

"Two years ago we were skeptical that AI could do this," Mr. Hofmann said. "Spoiler alert: it's all on you now."

The prescription: four shifts, executable not by next quarter but next week. First, stop treating AI as a project and treat it as an operating principle. Second, move from humans doing work with AI assistance to AI doing the work with human direction and oversight. Third, stop optimizing broken processes and start eliminating them. Design workflows assuming AI runs the whole thing, then decide where humans belong.

CHRISTOPHER GINN

Fourth, abandon annual planning cycles for anything AI-related. At the pace models are improving, a 12-month planning window for AI investments is already obsolete before it's approved.

He closed with a structural framework developed from studying what separates the 5% from the 95%: what he called the AI Transformation OS, organized across 10 dimensions spanning strategy, culture, data, talent, and operating model. The insight underneath all ten: productive individuals do not make productive firms. The only organizations achieving the outsized results of AI-native companies are the ones that have stopped bolting technology onto old structures and started redesigning the structures around the technology.

Like a factory owner at the turn of the last century who, rather than replacing the steam engine with an electric one, realized the electrical motor could be placed anywhere, and so rebuilt the entire production floor from scratch.

"You don't win by adding AI to what you already do," Mr. Hofmann said. "You win by asking what you would build if you started today."

"The headwind we get is not 'this tool will replace my job.' It's 'I've been doing this for 20 years, I know how to get it right—and if I use this tool, I might get it wrong."

—Dr. Radha Iyengar Plumb
VP, AI-first Transformation, IBM

The Moat That Matters

By the afternoon, the room had absorbed the premise: AI is accelerating, the stakes are real, and the organizations that treat it as a project will be outrun by those that treat it as an operating system. But Shideh Sedgh Bina, Insigniam co-founding partner and Elixirr partner, and Stuart Stern, Elixirr partner and former CIO of one of the world's largest insurance companies, arrived with the question the other sessions had only approached: what actually creates durable advantage when everyone eventually has the same tools?

Their core argument was bracing in its simplicity. Siebel Systems held 45% of the global CRM market. Today, less than 2%. BlackBerry commanded nearly half of U.S. smartphone market share. Today it makes no devices. In both cases, the technology was real, the lead was real, and the moat evaporated, not because the companies stopped innovating, but because the technology itself became the floor, not the ceiling. As Stern put it, drawing on a framework from *Harvard Business Review*: as a technology's ubiquity and power increase, its strategic importance diminishes. It becomes a cost of doing business rather than a source of distinction. "Right now," he said, "firms are saying 'game changer.' But AI is likely to become like cloud: necessary, even essential, but not sufficient."

Ms. Sedgh Bina and Mr. Stern's session was designed not to deliver this verdict and leave the room with it, but to begin the harder work: if technology alone can't be the moat, what can? They call the answer "competitive weapons": the complements that make technology powerful and, crucially, hard to replicate. The organizations that create enduring advantage, their research suggests, are the ones that build a reinforcing system of these complements around a focal technology, such that even a competitor with access to the same AI platform cannot easily copy what they've built.

Their framework gave executives a working vocabulary and a structured canvas to test their own organizations against, and the table discussions that followed were, by the room's own accounting, among the most generative conversations of the day.

"A moat," Ms. Sedgh Bina said, "is not a product. It is not a platform. It is not a technology. It is a way an organization operates."

CHRISTOPHER GINN

While You Were Piloting
AI's power is compounding on an exponential curve, says Elixirr's Adam Hofmann. Organizational adoption is barely linear. The distance between those two lines is where competitive advantage is being won and lost.

From Ambition to Ledger

The afternoon's final session was designed to put hard numbers on the table.

Dr. Radha Iyengar Plumb, IBM's Vice President of AI-first Transformation and former Chief Digital and Artificial Intelligence Officer of the U.S. Department of Defense, has operated at a scale few executives will encounter. IBM is 250,000 to 300,000 employees, publicly traded, running tens of billions of dollars in annual transactions, with all the legacy systems, data silos, and institutional immune responses that entails. What the company has produced from its AI transformation program, which it calls Client Zero, treating IBM as its own first client, is $4.5 billion in realized ledger savings reported to the Street, with a commitment to deliver another billion in 2026.

"You don't win by adding AI to what you already do. You win by asking what you would build if you started today."

—Adam Hofmann
Partner, Elixirr

Dr. Plumb was precise about how that happened and careful not to make it sound cleaner than it was. IBM's guiding principle was not to begin with AI. It was to begin with the work: eliminate what doesn't need to exist, simplify end-to-end workflows, automate what remains, and then embed AI everywhere. She distilled it into four pillars, each with specific outcomes: data, workflows, technology, and people. Governing and democratizing data produced billions in insight-driven business value. Breaking down workflow silos and integrating across functions led to a roughly 50% reduction in handoffs. Deploying AI-driven technology automated millions of work hours. And reinventing the workforce through continuous learning drove more than 20 percentage points of improvement in employee engagement.

The case studies made the abstraction concrete. AskHR, IBM's AI-powered HR platform, now handles 11.5 million employee interactions and resolves 94% of inquiries without a human. AskIT, built and deployed in 100 days from scratch, deflects 82% of IT support requests. In contract analysis, IBM has ingested approximately 700,000 contracts across the enterprise and achieved 100% coverage in analyzing customer contracts during M&A due diligence, a task that previously required triage and prioritization simply because the volume was impossible. In finance, a variance detection tool paired with an insight and error resolution agent has the potential to cut reporting cycle times by 50%, freeing analysts for work that actually requires judgment. In procurement, an invoice exception management agent now provides visibility and control over more than $20 billion in annual spend.

None of this happened cleanly. IBM's methodology ran on 90-day sprints: two weeks to identify pain points, two weeks to design a blank-slate solution, a rapid build, then repeated cycles of user testing and iteration.

First versions were, as Dr. Plumb put it, "not exactly right." The biggest headwind was not, as many might expect, fear of job displacement. It was fear of being wrong. Workers who had spent years mastering their own spreadsheet macros and dashboard logic were being asked to trust a system that might reach a different answer. IBM's response combined co-design with the people actually doing the work, and at a certain point, removing the safety net altogether. "It was painful," she said, "until it wasn't."

A Time for Action
Executives who attended the summit were left with the same question to answer: am I building for the future, or keeping yesterday alive a little longer?

The lesson she left the room with was not the savings figure or the methodology. It was the mindset beneath it: if you build a tool that nobody uses, you haven't built a solution. You've built a widget.

Drilling Down: Industry Deep-Dives
Between the morning keynote and the afternoon sessions, attendees moved into four industry-focused breakout rooms alongside venture capital firms and technology founders, a format new to the 2026 Summit and consistent with its venture-style design. The tracks examined financial services, healthcare, manufacturing and production, and retail and consumer goods. In each room, outside investors and founders shifted the register from internal strategy to external reality: what is actually being built, funded, and scaled, and how quickly does it threaten current operating models?

Mr. Rosenberg had framed the intention at the outset. The Summit was designed to function as a compressed version of what Insigniam calls an executive immersion, the deliberate collision with possibility that happens when leadership teams are brought to Silicon Valley, Tel Aviv, or London to engage the ecosystem reshaping their industries from the outside. "If we get our job done," he said, "your world will be bigger at the end of today than it was at the beginning."

A Mandate, Not a Moment
By the closing session, the room had earned its exhaustion. What it had also earned was a sharper frame. The executives who gathered in Philadelphia had arrived with varying degrees of AI maturity. Some had active deployments at scale. Others had pockets of adoption and the familiar backlog of stalled pilots. A few were genuinely early. All of them left with the same challenge in hand: the capability of AI is compounding faster than organizational adoption. That gap is not closing on its own. And the window for establishing the practices, governance, and cultural habits that will determine who wins the next decade is measured in months, not years.

Around the closing tables, one thread kept surfacing. An executive at a stealth-stage company put it clearly: "There's an inherent advantage right now in being small and agile. But it's ephemeral. Getting ahead of the large incumbents isn't going to create long-term value by itself. You also have to build the things that are hard to copy."

Others named the moat session as the piece they hadn't expected, the framework they would bring back to their teams. Several left with specific commitments: next week, not next quarter.

The Summit's final slide showed a dim lightbulb over a plain background, a callback to the Commissioner of the Patent Office and his infamous 1899 certainty. The lesson wasn't that Duell was foolish. The lesson was that the trap he fell into is structural, persistent, and available to any of us at any moment: looking at the current state of the world and mistaking it for the permanent one.

As Mr. Koulopoulos put it early in the morning, in words that proved to be the day's real summary: "The winners are not the ones who get it right. They're the ones who learn fastest in motion." **IQ**