ISSUE #5 - MARCH 2026

EVALUESERVE

ELEVATE

THE MOMENT AI BECOMES ACCOUNTABLE

THE RISE OF THE AI-AUGMENTED ANALYST: EXPERT INTERVIEW

PROCESS CLARITY BEFORE AGENT DESIGN

HOW AI IS EMBEDDING INTELLIGENCE INTO EVERY DECISION

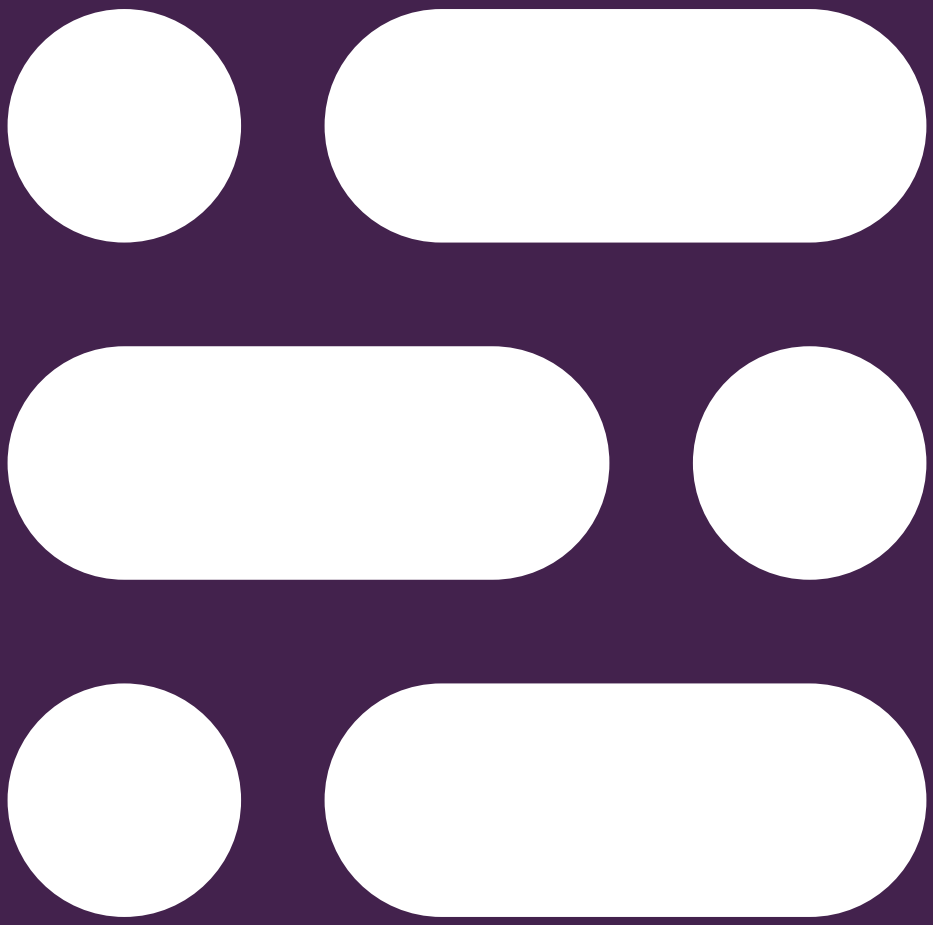

Contents

• Editorial Note

The Moment AI Becomes Accountable

There's a moment every AI team eventually faces.

The demo works. The room claps. The answers are sharp. The workflow looks seamless. The agent books the meeting, drafts the memo, and updates the CRM. Then someone says, "Okay. Let's turn it on."

That's when the real test begins.

In production, the data is messy, users are impatient, attackers are creative, compliance is watching, and mistakes have names attached to them. This is the point where AI must move from impressive demos to dependable infrastructure.

For years, we've celebrated capability, and to be fair, the progress has been extraordinary – 78% of organizations already use AI, and more than half are seeing revenue gains. The appetite is real.

However, that appetite is not architecture.

Securing AI agents is the responsibility of the team that implements them. It's a mindset as much as an engineering choice. When people are left handling prompt injection, permissions, or execution boundaries on their own, it starts to feel less like a finished product and more like a science project with a login.

That's the dividing line between novelty and reliability.

Ownership

Infrastructure starts with ownership. Ownership of identity, of permissions, of what happens when something goes wrong. Demos offer solutions by assuming cooperation, while in real-life production assumes resistance. And resistance is where identity becomes decisive. "Block all bots" is demo thinking. The reality is that some automation is malicious and some is mission-critical, and the difference can't be inferred from behavior alone. If a system can't definitively identify an agent, it can't safely grant autonomy.

That's the missing primitive of the agentic web: **verifiable identity**.

Trust is moving from heuristics to certainty. Agents need provable permissions because once they move inside the enterprise, the stakes change dramatically.

Some investors predict 2026 will mark the breakout of agent-driven enterprise operations.

Agent-driven enterprise operations doesn't mean chatbots anwering questions; rather, it means agents updating financial systems, triggering supply chain orders, enforcing compliance policies, and resolving IT incidents autonomously.

At that point, we're no longer talking about demos. We're talking about machinery.

> **If your AI can't operate inside your ERP, your CRM, your compliance stack — it's not embedded. It's hovering. And hovering doesn't run a company.**

Enterprise leaders increasingly describe this next phase as the moment when agents stop being an experiment and start becoming infrastructure. And infrastructure plays by different rules. It doesn't apologize for hallucinations. It doesn't ask users to double-check its work. It is governed, monitored, permissioned, and continuously operated.

This is where many impressive systems will break down. Autonomy without graduated trust is just automation roulette.

Prompt injection is a system flaw. The agents that fail less aren't necessarily more intelligent; they're engineered with constraints that make them safer and more predictable.

The distinction is critical:

- Demos succeed because they showcase novelty.
- Production systems succeed because they deliver reliability.

The next phase of AI

The companies that win the next phase of AI will be the ones that recognize that **security is the foundation of capability and not a limitation.**

They'll be the ones that treat identity as core infrastructure, embed governance from the start rather than bolt it on later, and absorb complexity themselves rather than offload it onto customers.

When a demo breaks, that's not a failure. That's graduation.

And in 2026, when agents evolve from assistants to true actors inside the enterprise, we'll see the real divide: those who built something impressive, and those who built something you can trust.

This issue of *Elevate* explores that transition from novelty to reliability across several dimensions of AI adoption. We examine how the analyst role is shifting from data gathering toward judgment, interpretation, and decision support. We look at why effective agent design begins with mapping decision workflows and making hidden assumptions explicit. We explore how organizations must rethink knowledge management to embed intelligence directly into everyday work. And we introduce new ways of measuring value in agent-led systems, from cost per insight and time-to-decision to exception management and human–agent leverage.

Together, these perspectives reflect a broader shift: AI is now infrastructure for decision-making, and we must stop treating it as an answer producer.

Because, in the end, the future of AI doesn't rely on what it can do in a demo—but by what it can be trusted to run.

Author:

Marcela Cortez
Editor & Senior Associate
Evalueserve

• Expert Interview

The Rise of the AI-Augmented Analyst

Whitt Ellis
Director of Insights & Advisory

As AI takes over routine analysis, the role of the analyst is being fundamentally redefined. We spoke to Whitt Ellis, Director of our Insights & Advisory division, about what it now takes to be effective in an AI-augmented world, and why judgment has become the most critical skill of all.

What feels fundamentally different about this AI moment compared to previous waves of analytics or automation?

Whitt: At the individual level, every professional needs to be able to use AI independently and effectively, whereas only a subgroup of professionals needed to learn analytics or automation.

At the enterprise level, AI not only democratizes access to information but also makes that information available in real time. So, while analytics and automation were a starting point for transforming how businesses function, they tend to represent large and long investments. AI has radically changed how efficient business transformation can occur, both in terms of cost and time.

How has the analyst role changed in practice, not on paper, over the last 12–24 months?

Whitt: I see five expectations that have evolved or are new.

Analysts must deliver everything faster. GenAI is a great tool for conducting exploratory research, checking for contrarian views, summarizing, polishing professional language, learning, and getting advice.

Analysts must ensure sources are reliable and research outputs are accurate – this expectation existed before AI. What changes is checking for and mitigating bias, hallucinations, and drift.

Analysts must tell impactful and balanced stories around the data generated by AI. This involves explaining why AI-generated outputs are relevant, what underlying assumptions were taken, and what potential risks exist in the output.

Analysts must implement robust AI governance. Again, this builds off traditional research governance activities such as documentation, risk assessment, and version control, with additional concepts like explainability and policy-as-code.

Where do you see analysts over-trusting AI, and where do they still underestimate it?

Whitt: Analysts sometimes over-trust AI when they equate confidence in their prompt to confidence in the output. Analysts need strong problem framing skills to translate business questions into prompts, as well as intentional and critical thinking around bias, hallucination, drift, and comprehensiveness to ensure quality outputs.

Analysts often underestimate AI's ability to orchestrate end-to-end research processes.

It's relatively easy to replace a singular, discrete research task with a hyper-specialized agent, for example, an agent that generates company profiles or an agent that gathers product details. These agents are research tools.

It's harder to orchestrate multiple agents at scale. But this is where transformative impact is born – when multiple hyper-specialized agents collaborate and synthesize insights, AI starts to resemble a high-performing research team rather than a collection of tools.

This orchestrated intelligence is where our most exciting innovation is happening. For example, orchestrating go/no-go decision workflows or sales lead intelligence workflows.

5 Skills Every AI-Augmented Analyst Needs

- Domain expertise
- Critical thinking
- Governance
- Big-picture vision
- Prompt engineering

What mistakes do organizations make when introducing AI into analyst teams?

Whitt: I have heard many clients talk about democratizing access to AI. In theory, this is excellent – AI in the hands of each employee. In practice, however, there needs to be a sharp focus on cost-benefit.

Assess whether team members have sufficient skills to use the AI tools provided. If not, establish the necessary training programs and consider a phased rollout of access based on demonstrated skills. This ensures optimal use of paid tools.

Quantify the baseline before AI implementation. This allows your organization to track efficiency gains and ROI.

Brainstorm how AI might be monetized. This provides team members with a starting point for innovation that generates revenue.

Understand the cost implications of each AI tool and set up governance. This is fundamental for controlling spend, especially with pay-per-run tools.

Is AI narrowing the gap between analysts or widening it?

Whitt: AI will widen existing gaps between analysts that have deep domain knowledge and generalists since AI replaces a significant portion of data gathering that was done by generalists. It will also widen the gap between those who have strong critical thinking skills and those who can't move beyond process-oriented thinking.

On the other hand, AI provides analysts with a tool to narrow the gap. In my opinion, career development now feels far less tied to one's university education, previous work experiences, or even connections. With AI, curious minds can accelerate their learning, innovate, or generate significant financial improvements for their company, and through that, shift into new career paths and move up the ladder.

What advice would you give analysts who are worried about staying relevant in the new AI-driven world?

Whitt: At the individual level, it's important to assess your role and identify what tasks AI will replace, transform, or introduce. Start by learning and then applying AI to those tasks.

At the enterprise level, do the same. Think about how your company will be impacted by AI – what gets replaced, what gets transformed, and what is new. Thinking big picture has been and will always be fundamental for staying relevant.

Understanding your role within the broader company context allows you to take actions that keep you relevant. As a north star, know what your value addition will be in an AI-enabled world. Develop your skills and your company's service and product offerings to deliver on that value proposition.

Why Strategic AI Transformation Requires Tailored Rules

The $2 Billion Question

The CFO asks a simple question: How long does it take to turn an investment proposal into a board-ready recommendation?

Everybody goes silent.

Eventually, someone offers a cautious estimate: "Four weeks… maybe four months." It depends on the business unit, the type of proposal, the time of year, who's available. Around the table, others nod. Different answers, none definitive.

For a team allocating $2 billion a year, there is no clear, shared understanding of how decisions actually move from idea to approval. The process itself is largely invisible—distributed across inboxes, spreadsheets, and institutional memory.

That ambiguity is more common than most executives realize. And it explains why many ambitious AI transformations stall before they begin: you can't automate a process that no one has truly mapped.

Today, most companies have deployed AI agents in some form. Budgets are rising. Vendors promise an agentic future embedded in every application. Yet the leap from pilot to production remains stubbornly wide.

The Problem: "Let's Just Apply AI to Our Process"

Vendors make transformation sound simple: deploy an agentic AI platform, and results will follow. Systems integrators promise implementation in weeks.

What often gets overlooked is that AI can only work as well as the process beneath it. If an organization cannot clearly measure or describe its current process, there is no reliable way to know whether AI is improving performance—or simply accelerating hidden flaws.

The problem becomes even clearer in strategic functions like capital allocation.

Operational processes—like invoice processing—are repetitive, rule-based, and stable. They are ideal candidates for automation.

Strategic processes are fundamentally different. Capital allocation is variable, judgment-heavy, and shaped by shifting priorities and risk appetite. Every decision carries context that does not fit neatly into a workflow diagram.

Treating these decisions like back-office automation is a category error.

What Process Clarity Actually Means (And Why It's Hard)

Process clarity means making the invisible visible—surfacing the judgment calls, unwritten rules, and quiet rework loops that shape real decisions.

In capital allocation, complexity hides everywhere:

Critical knowledge is distributed across people, not systems.

Decision logic lives in habits ("this is just how we do it here") rather than documented standards.

The same investment can follow different approval paths depending on the business unit.

And while the math behind Net Present Value (NPV) or Internal Rate of Return (IRR) is straightforward, the real decision hinges on judgment: risk appetite, strategic fit, timing, and competitive context.

None of that sits neatly in a workflow.

Drop AI agents into this environment without first making those dynamics explicit, and the results may look impressive—but feel untrustworthy: recommendations no one backs, black-box outputs that stall in governance, and local optimizations that miss enterprise trade-offs.

Eventually, the initiative becomes another AI pilot that never scales.

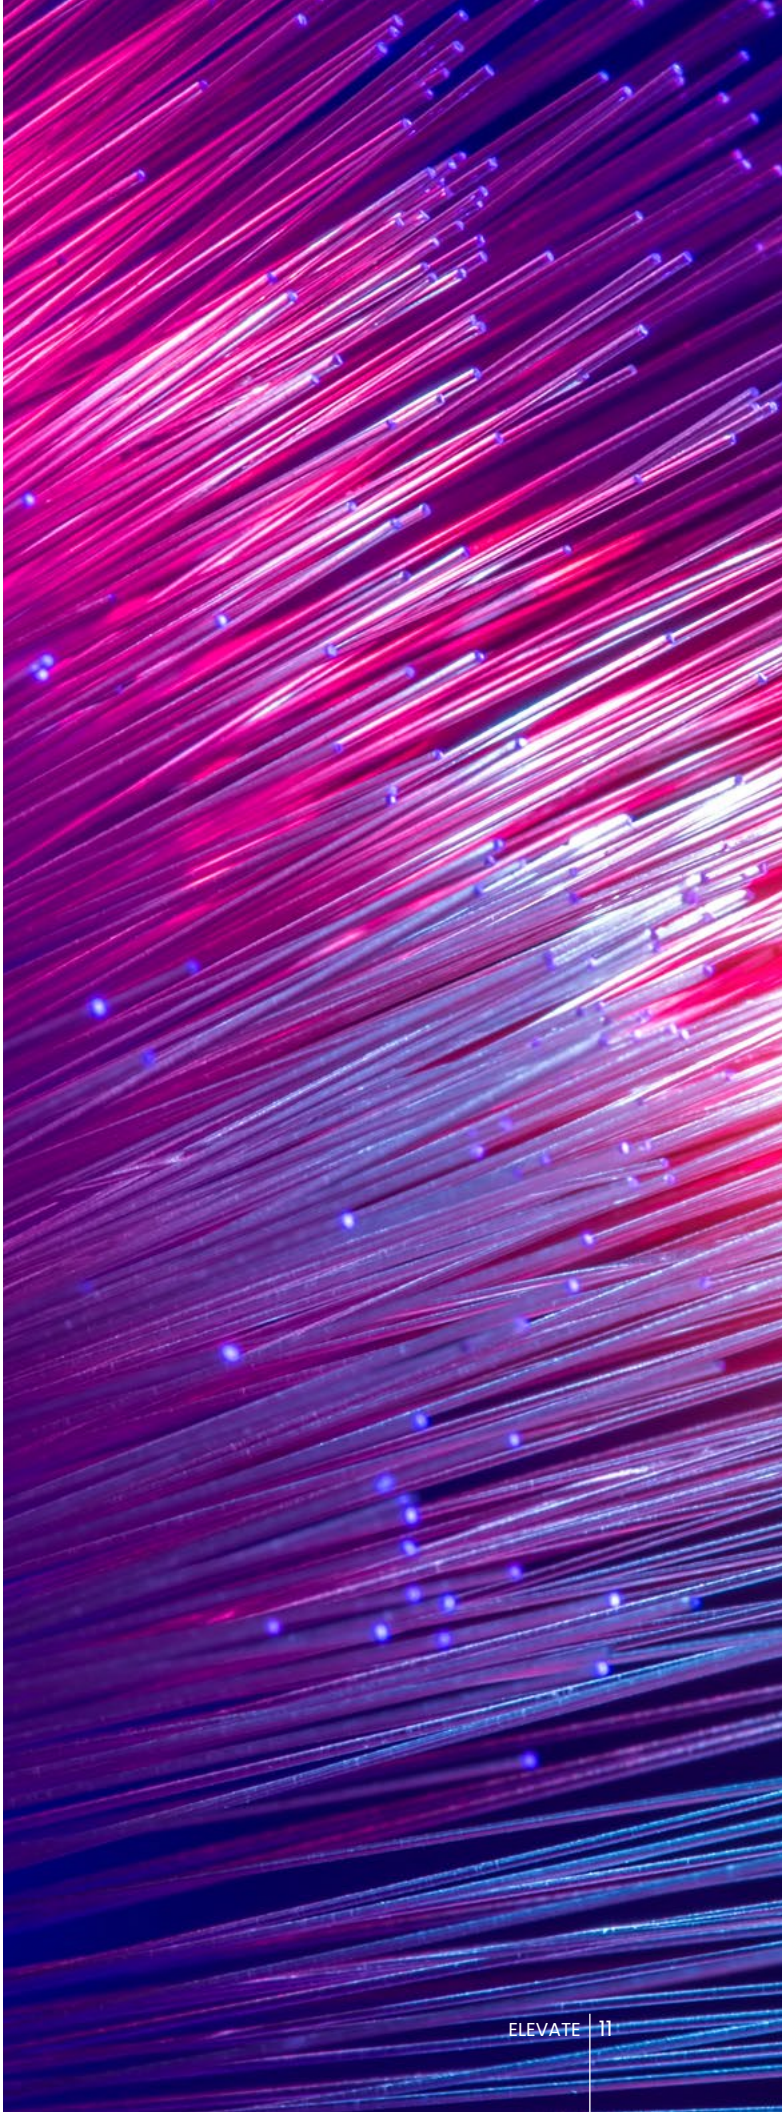

How to Get Process Clarity Right

A Four-Phase Methodology

After many agentic AI deployments in capital allocation, a consistent pattern has emerged: the companies that succeed start with forensic process clarity.

Phase 1: Make the Invisible Visible

Map how decisions actually get made—across business units, investment types, and approval paths.

Track handoffs, rework loops, and "it depends" moments.

In one energy company, this revealed a stark inversion: analysts spent nearly 75% of their time gathering data, rebuilding models, and formatting decks, and less than a quarter on true strategic thinking.

Even more telling were the unwritten rules buried in past approvals: lower IRR thresholds for decarbonization, automatic flags for policy risk, softer scrutiny for certain sponsors. None documented. All real.

Phase 2: Quantify the Current State

Measure cycle times, effort hours, rework rates, and learning gaps.

This is where the "hidden tax" becomes visible. One industrial firm tested only three scenarios on a $250M investment because deeper analysis took days. A missed downside case later contributed to an $80M impairment.

The constraint wasn't intelligence—it was bandwidth.

Phase 3: Design the Agent Architecture

Translate process steps into clear agent responsibilities with defined inputs, outputs, and human checkpoints.

The goal is not autonomous decision-making, but orchestrated intelligence.

Each agent explains its logic. Each recommendation includes its assumptions and sensitivities. Trust is designed into the system from the start.

Phase 4: Prioritize the MVP

Start where pain is high, risk is low, and confidence is strong—standardizing intake, synthesizing data, and automating repeatable modeling.

More complex agents, such as portfolio optimization, follow only after credibility is established.

The Real-World Test: Energy Transformation

Process mapping at one energy company revealed the real issues: inconsistent templates across business units, unquantified decarbonization impact, uneven risk analysis, no feedback loop on investment performance, and 30% of Investment Committee meetings sent back for rework.

Before deploying agents, the organization standardized the fundamentals—common templates, a shared assumption library, three required risk scenarios, and clear decision criteria based on historical approvals. Only then did automation follow. Agents validated proposals, synthesized data,

generated risk scenarios, quantified emissions impact, and supported recommendations with transparent logic.

Twelve months later:

- Cycle time dropped from six weeks to two
- Analyst effort was cut by more than half
- Rework fell to single digits
- Portfolio optimization became a quarterly discipline

But the biggest impact was capacity.

With the mechanics handled, the team shifted from assembling analysis to shaping strategy.

Clarity first. Agents second. Results that scale.

The Anti-Patterns to Avoid

- **Pilot First**
  Starting with a narrow pilot often creates disconnected agents and technical debt. Map the full process before building anything.
- **Governance Later**
  If leaders cannot explain recommendations, they will not trust them. Build explainability, audit trails, and human checkpoints from day one.
- **Everything Is Unique**
  Strategy is unique. Financial math is not. Customize strategic logic; configure standard analytics.
- **Train an LLM on Everything**
  Unstructured historical memos produce confident nonsense. Define process structure first, then use LLMs to enhance insights.

Where the Advantage Begins

The companies gaining ground share one trait: clarity in how capital decisions are made, measured, and improved.

That clarity turns AI from experimentation into advantage.

Start by making the process explicit—cycle times, effort allocation, decision rules, and feedback loops. Once the foundation is visible and measurable, intelligent agents can orchestrate, optimize, and scale.

The result is not just faster analysis, but better capital decisions: faster reallocation, stronger trade-offs, and the ability to ask—and answer—strategic questions that were previously out of reach.

That's where long-term advantage begins.

Authors:

Sujay Dutta
Director,
AI Innovation & Analytics

Gauravgeet Singh
Vice President,
Energy & Natural Resources

How AI Is Embedding Intelligence Into Every Decision

The greatest satisfaction for any analyst comes not from producing insight, but from seeing it used. Few experiences can match watching a client acquire a company you once flagged in your research and realizing that your work helped guide a high-stakes decision. Yet these moments are the exception, not the rule. In most organizations, research, and decision-making follow separate paths – insights are delivered, but decisions happen elsewhere. And in that gap, much of intelligence's potential impact is lost.

For years, commercial intelligence has been treated as a support function—producing reports, dashboards, and analyses that sit adjacent to decision-making rather than inside it. As AI moves into day-to-day workflows, intelligence is starting to shape decisions as they happen, not after the fact. This changes what "commercial excellence" actually means. The advantage is now about who can **translate intelligence into better, more consistent decisions** that drive growth.

This shift is forcing organizations to rethink their decision-making toolkit and the role intelligence plays in how work gets done. Knowledge management, long considered a back-office discipline, is being pulled into this reckoning. Once designed to store and retrieve information, it is now being asked to do something more ambitious: help people decide.

That expectation is being accelerated by generative AI. But the most consequential change is not technological. It is conceptual. Instead of starting with data, tools, or platforms, organizations are beginning to work backward from value.

- Which decisions matter most?
- Where do delays, inconsistencies, or poor judgment create measurable commercial risk?
- What intelligence, delivered at the right moment, would improve those decisions?

This decision-first framing represents **a break from traditional analytics thinking.** In the past, success was measured by output: how many reports were produced, how quickly analyses were delivered, and how much data could be processed. Today, leaders are asking harder questions.

- Did intelligence reduce decision cycle time?
- Did it improve pricing discipline, customer targeting, or forecast accuracy?
- Did teams actually act on the recommendations—and do so consistently?

The limits of legacy knowledge systems have made this shift unavoidable. Most organizations generate thousands of documents every day, spread across teams, regions, and formats. Over time, **knowledge bases became cluttered with duplicates, outdated versions, and poorly tagged content.** Simply layering generative AI on top of that complexity produces faster answers, but not necessarily better ones. Without structured data, clear taxonomies, and reliable context, AI amplifies noise as efficiently as it surfaces insight.

This is why many GenAI initiatives stall after early pilots. The problem is the foundation and fit of the model's capabilities. Data annotation, document parsing, and metadata design re-emerged as strategic priorities. The old rule still applies that **systems can only be as intelligent as the information they are built on.**

Where generative AI begins to show real commercial impact is when intelligence moves into the workflow. Instead of requiring users to search for information, AI can now summarize markets, flag risks, generate first drafts, and surface relevant context directly within sales, marketing, and strategy processes. Intelligence becomes ambient rather than optional—present at the moment decisions are made.

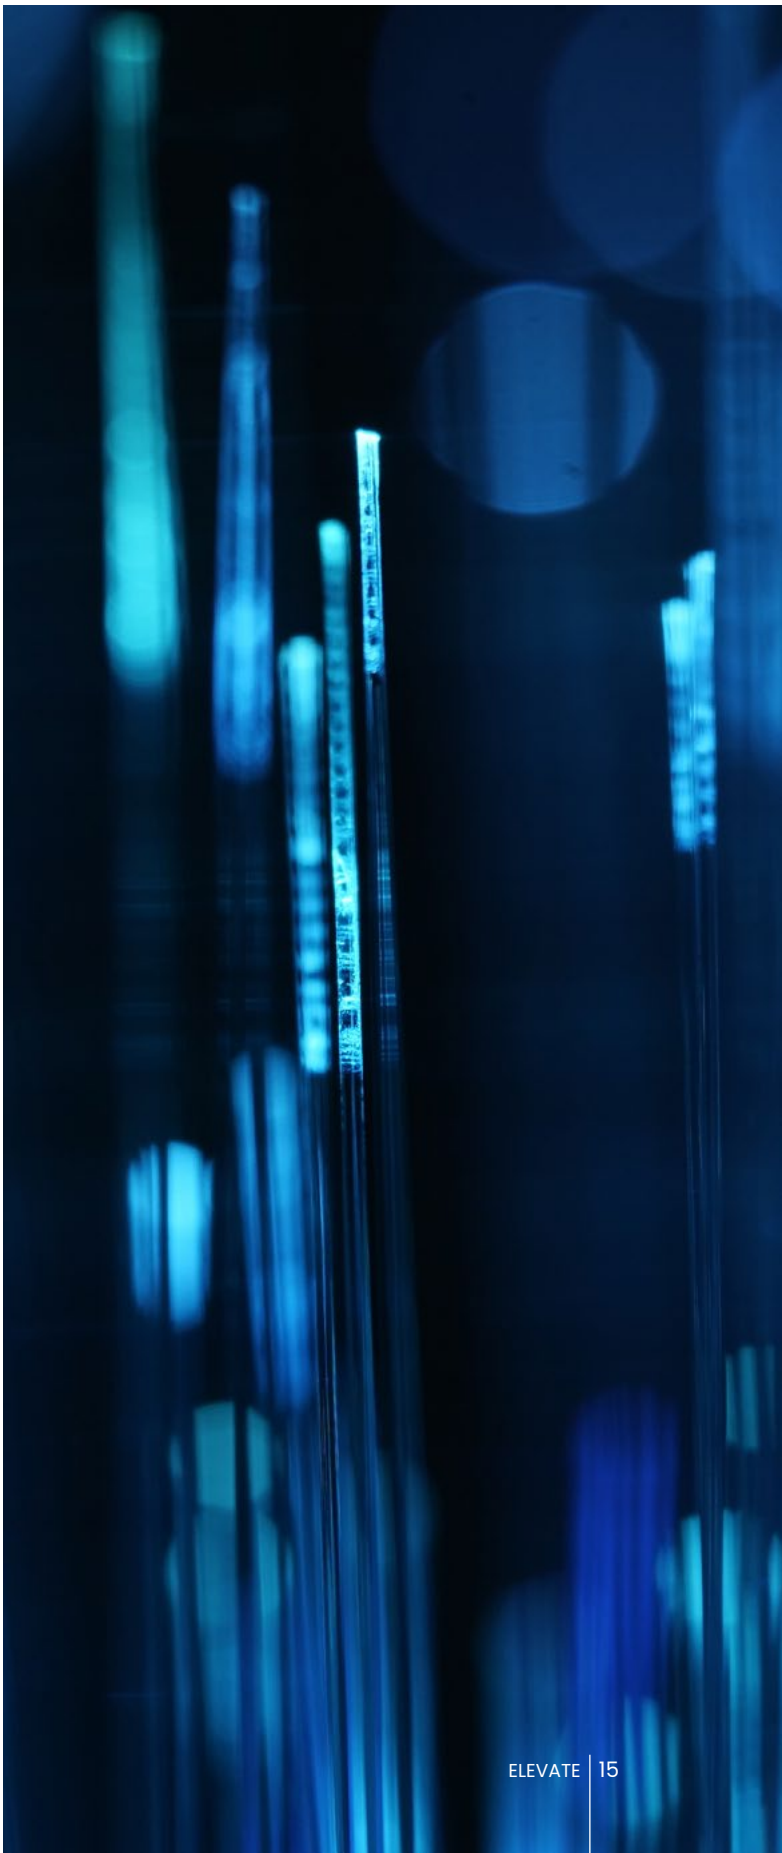

This shift also changes how people interact with knowledge. Early evidence suggests that GenAI-driven systems are used in shorter, more frequent sessions, woven into daily work rather than reserved for formal research. That has implications for design, trust, and adoption. Users need confidence in the quality and security of outputs, clarity on how recommendations are generated, and reassurance that AI augments judgment rather than replaces it.

For many years, our clients relied on daily and weekly newsletters covering topics such as competitive developments, sector performance, and market trends. These products were valued for their consistency and breadth, providing a steady flow of information to support awareness and monitoring.

Over time, however, client expectations have evolved. Today, clients are looking for intelligence that is **personalized, contextualized, and explicitly focused on implications.** They no longer want information simply to read; they want insight that interprets what is happening, explains why it matters, and clarifies what actions should be considered. Intelligence is increasingly expected to function less like a report and more like a **personalized consultant**—one that helps decisionmakers navigate complexity and uncertainty.

We are also seeing a clear shift in how insights are consumed. Clients increasingly want intelligence delivered through face-to-face and virtual interactions, not just static outputs. While slide decks remain important, they are no longer sufficient on their own. Clients want to engage directly with the researchers and analysts who are closest to the subject matter—to hear their perspectives, challenge assumptions, and explore nuances that do not always fit neatly onto a slide.

As a result, the role of the analyst is evolving. Analysts are no longer valued solely for their ability to gather information or produce content. They are increasingly expected to act as **thought partners**—bringing judgment, domain expertise, and point of view to client discussions. This evolution reflects a broader shift from information delivery to decision enablement, where the true value of intelligence lies not in what is produced, but in how effectively it informs action. AI is not automating judgment. It is industrializing it. And commercial excellence is no longer about knowing more, but about deciding better.

Authors:

Michal Radziuk
Director,
Evalueserve

Shobhit Saxena
Director,
Evalueserve

Rethinking Cost, Pricing, and ROI in Agent-Led Processes

Many organizations attempting to quantify the return on investment of agent-led processes are still relying on metrics designed for an earlier generation of automation. Measures such as hours saved or headcount reduced were once useful for evaluating deterministic workflows. Still, they often fail to capture how AI agents actually create value in modern analytical and decision-driven environments. In some cases, these legacy metrics underestimate the impact of agents by misrepresenting them.

As agent deployments move from controlled pilots into real production environments, particularly in professional services and analytics-driven workflows, **a different evaluation lens becomes necessary**. Agents do not simply replace discrete tasks; they alter the structure of work, the pace of decision-making, and the distribution of effort between humans and machines. Traditional productivity accounting rarely reflects these shifts.

This article proposes a practical ROI framework aligned with how agent systems behave once deployed at an enterprise scale. Drawing on observed implementation patterns, it introduces operational metrics that finance and operations leaders can evaluate with confidence: cost per insight, time-to-decision, exception rates, and human–agent leverage.

The goal is not to claim a universal formula for "AI value," but to outline a set of measurable indicators that **move agent adoption from experimentation** toward economic discipline—treating agent-led processes as a business capability that can be managed, optimized, and justified over time.

How to Pivot

The real shift agents introduce is not faster task completion, but **a new unit of value altogether**. Tasks generate outputs; decisions are inflection points. In agent-led systems, value appears as a reconciled dataset that doesn't need manual cleanup. A client-ready hypothesis formed before the meeting; a risk flagged early enough to matter; a prioritization call made upstream rather than too late.

If agents change how value is created, ROI must change with them. The economy lives in how quickly decisions become possible, how often intervention is required, and how much capacity a single human can effectively oversee. Measured this way, agent ROI stops being a proxy for labor reduction and becomes a lens on decision throughout.

An Enterprise ROI Framework for Agent-Led Processes

In practice, evaluating the economics of agent-led systems requires moving beyond generic productivity metrics and toward unit economics grounded in how these systems actually run. In several enterprise deployments across analytics, document processing, and research workflows, a consistent pattern emerges: meaningful ROI only becomes visible once cost, workload, and decision impact are measured at the right operational level.

Cost per Insight (CPI)

Replaces cost per task as the baseline metric. What ultimately matters is not how cheaply an agent can generate output, but how much it costs to produce a usable insight that can support a decision.

In production environments, that denominator expands quickly. It includes retries, orchestration, retrieval of pipelines, preprocessing, post-processing, and the human time required to validate or correct outputs before they become decision-ready. Early deployments often expose a common failure mode: systems that appear inexpensive at the token or API-call level but become costly once human cleanup and orchestration overhead are included.

The practical solution has been to normalize cost around unit economics tied to a specific business output—for example, a profile, an analysis run, or a document extraction cycle. Anchoring the model to a single reference use case forces clarity about the true workload drivers: pages processed, documents ingested, iterations required, and output size. Without that operational granularity, even sophisticated cost models tend to collapse into vague "per request" assumptions that do not survive scale.

CPI therefore reframes the economic question from how cheaply an agent generates text to how efficiently a system produces decision-grade insights.

Time-to-Decision (TTD)

Captures where the largest economic gains often occur. TTD measures the latency between a signal becoming available and a decision being ready to act upon.

In analytics, advisory, and risk workflows, compressing this window frequently creates more value than marginal improvements in output quality. Agents rarely generate value purely by replacing execution; their greater impact comes from pulling moments of judgment forward in time.

In real deployments, this effect shows up in subtle ways. Processes that previously required sequential human review—document reading, extraction, structuring, and synthesis—can run in parallel pipelines. Analysts move from gathering information to evaluating it. As a result, the economic benefit appears less as "automation savings" and more as faster decision cycles, especially in environments where timing materially affects outcomes.

Exception Rates

When properly defined, become an economic signal rather than a reliability score.

The relevant question is what failure costs? To detect and recover from it. High exception rates can be perfectly acceptable if detection is immediate, and remediation is lightweight. Conversely, systems with low apparent error rates can become economically problematic when failures surface late or trigger expensive rework.

Many early pipelines illustrate this dynamic. Cost spikes often emerge from edge cases—failed extractions, incomplete document parsing, or reruns triggered by partial outputs. Once these cases are identified and isolated, targeted engineering changes—better chunking strategies, improved routing logic, or narrower processing scopes—reduce both cost volatility and operational friction.

In this sense, the economic performance of agent systems depends less on perfection than on resilience and controllability.

Human–Agent Leverage Ratio

Reframes the entire productivity discussion.

The relevant metric is not how many tasks agents replace, but how many parallel decision threads a human can effectively supervise. As systems mature, this ratio tends to improve naturally. Agents handle the structured portions of analysis—extraction, normalization, preliminary synthesis—while humans focus on interpretation, exceptions, and judgment.

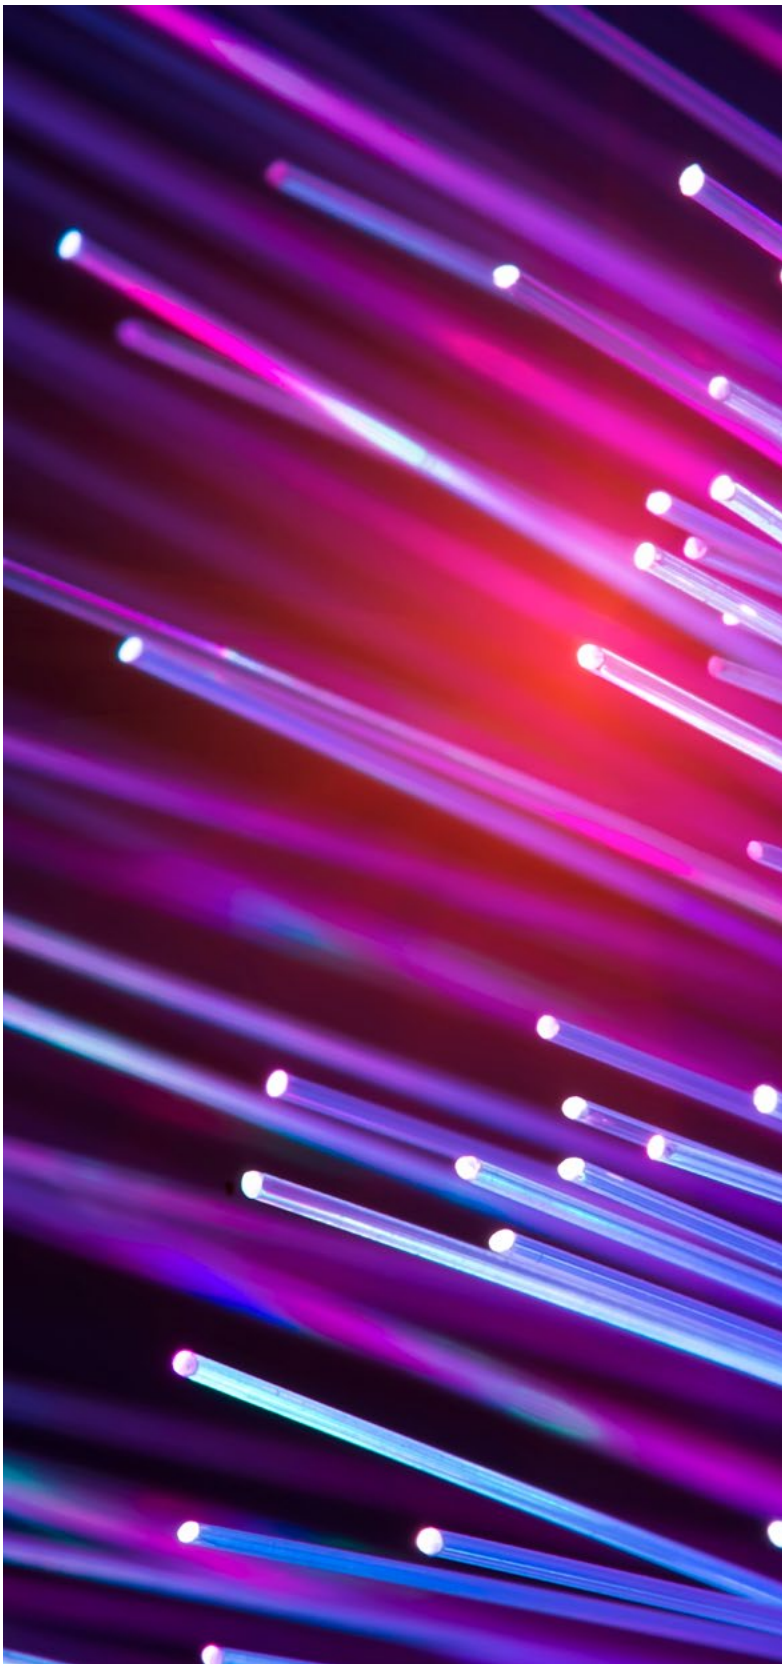

Across knowledge-intensive workflows, this shift changes the economics of scale. Rather than eliminating human roles, agent systems expand the number of analytical processes a single expert can oversee. Value creation therefore moves from substitution to capacity expansion.

This is how agent-led processes ultimately scale: not by removing humans from the loop, but by amplifying their reach across more decisions, datasets, and workflows than was previously feasible.

Changes from ROI Measured Correctly

When ROI is measured correctly, behavior changes. Organizations run fewer pilots, but they run them deeper—designed to survive contact with production rather than to impress in review decks. Investment shifts upstream, toward observability, escalation paths, and recovery mechanics that determine real cost. Decision outcomes gain clear ownership, because value is finally traceable. And finance stops treating AI as an exception, evaluating agent systems with the same rigor applied to any other operating asset. At that point, agents stop being experiments. They become infrastructure.

As a business capability, agent-led processes require governance, budgeting, and performance management aligned to how decisions are made and amplified.

Author:

Satyajit Saha
Global Head of Technology and Digital Solutions

Client Story

Modernizing Competitive Intelligence: From 6 Days to Minutes for a Global Energy Storage Leader

Overview

A global energy storage firm partnered with Evalueserve to modernize how competitive intelligence was delivered across its commercial organization. They needed to accelerate the process of translating competitive intelligence into timely, decision-ready outputs.

By combining Evalueserve's domain expertise in competitive intelligence with agentic AI embedded in Insightsfirst, the client dramatically reduced turnaround times, scaled intelligence delivery, and provided commercial teams with more actionable, context-rich insights. What once took days of manual synthesis could now be delivered in minutes, without compromising analytical rigor, enabling faster, more confident commercial decision-making.

Challenge

The energy storage firm has operations across multiple regions and a highly competitive market landscape. Its Global Commercial Excellence team plays a critical role in supporting sales, marketing, strategy, and finance with insights that inform commercial strategy, product management, inventory decisions, and capacity planning.

Evalueserve had been a long-standing partner, delivering domain-led market and competitive intelligence (MICI) through Insightsfirst, which served as the central platform for competitive profiles, operational benchmarks, and expert analysis. With growing information depth and faster business cycles, the client saw an opportunity to further enhance intelligence delivery—making it more dynamic, integrated, and directly aligned with business action.

The core challenge was not insight quality, but speed and scalability. When commercial teams needed concise, targeted views for specific situations—such as preparing for a customer meeting or responding to a competitive threat—intelligence managers had to manually navigate multiple profiles, extract relevant updates, and synthesize them into usable formats. Creating a single competitive summary often took five to six days, diverting experienced analysts from higher-value interpretation and strategic work.

Our Approach and Solution

Evalueserve addressed this challenge by pairing deep domain expertise with agentic AI, using technology to amplify, not replace, analyst judgment.

Building on the client's existing Insightsfirst instance, Evalueserve introduced specialized AI agents designed around real competitive intelligence workflows. These agents were grounded in Evalueserve's established research frameworks, taxonomies, and commercial use cases, ensuring that automation respected the nuance and context required for high-quality intelligence.

We worked with the client to prioritize a high-impact use case: on-demand competitive summaries derived from expert-curated intelligence.

The solution enabled:

- AI agents to retrieve the latest, analyst-validated competitive data—such as production capacity and utilization across manufacturing sites—directly from Insightsfirst
- Automated refreshes that preserved Evalueserve's research rigor while eliminating manual extraction
- Intelligent synthesis of insights into concise, decision-ready narratives, tailored to specific commercial scenarios
- Outputs that went beyond static battlecards, incorporating contextual commentary, key implications, and strategic takeaways—reflecting how an experienced intelligence professional would frame the information

This allowed intelligence managers to generate structured summaries—typically delivered in a focused 3–4 slide format—almost instantly, while retaining the depth and relevance expected by internal stakeholders.

Business Impact

By embedding agentic AI within a domain-led intelligence platform, the client transformed how competitive intelligence supported commercial decision-making.

- **Faster insight delivery:** Intelligence that once took 5–6 days to compile could now be generated in minutes
- **Higher analyst productivity:** Intelligence managers spent less time on manual synthesis and more time on strategic interpretation
- **Greater commercial relevance:** Sales and marketing teams received insights framed around real business decisions, not just data points
- **Scalable intelligence operations:** The Commercial Excellence team could support 30–40 competitive profiles annually without increasing effort
- **Stronger ROI on intelligence investments:** Existing analysis within Insightsfirst was more effectively leveraged and monetized

By moving from traditional intelligence reporting to an agent-driven intelligence ecosystem, the client took a significant step toward operationalizing AI, not as an experiment, but as a practical engine for commercial impact.


850.57
344.22
576.57
289.85
576.66
978.36

Client Story

Scaling Spend Intelligence Through an Agentic Operating Model

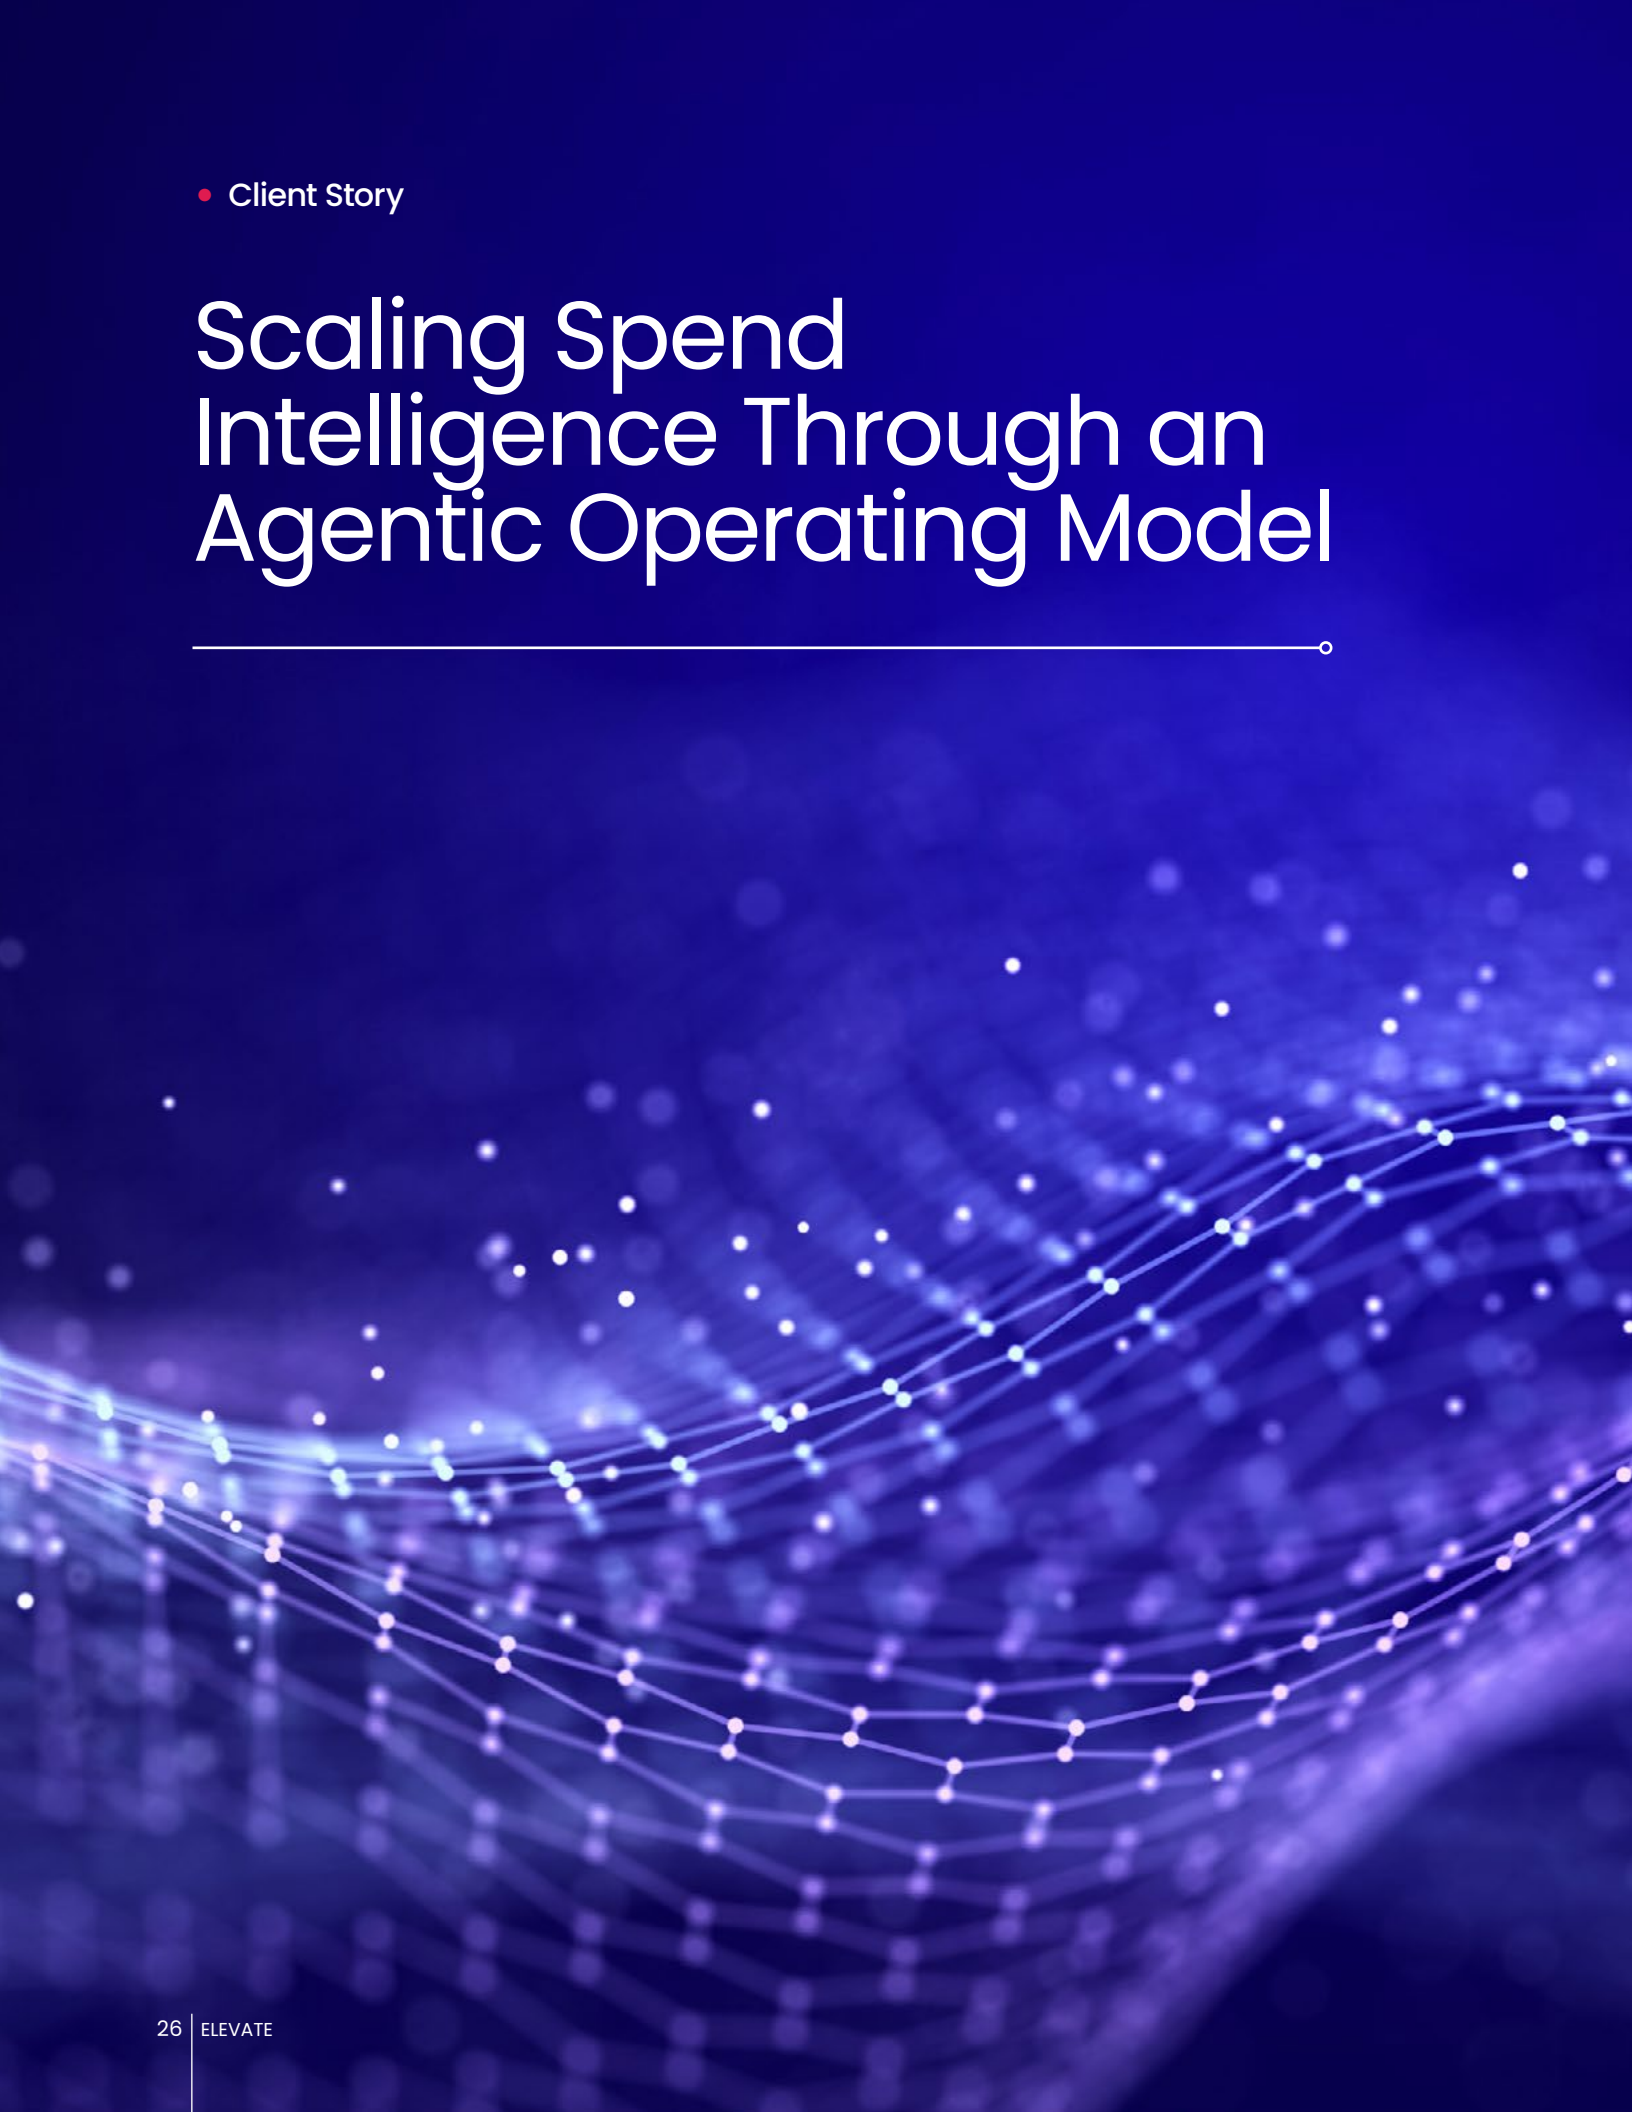

The Background

Leading procurement advisory firms operate in increasingly complex spend environments. Clients expect rapid clarity across fragmented supplier landscapes, deep multi-level classifications, and defensible insights that inform high-stakes sourcing decisions.

For this global professional services firm, spend intelligence had become a core competitive capability—powering client savings, sourcing strategies, and procurement transformation programs.

However, delivering that capability at scale required a significant operational engine:

- Supplier name normalization across fragmented datasets
- Mapping entities to immediate and ultimate parent hierarchies
- Classifying spend into multi-level NAICS and custom taxonomies
- Producing consistent, audit-ready analytical outputs
- Supporting rapid proof-of-concept (PoC) analyses for new client pursuits

For several years, Evalueserve supported this work through a hybrid model of domain analysts and technology accelerators. The approach worked—but demand was growing faster than the model could sustainably absorb.

The Moment of Reassessment

As data volumes increased and client timelines shortened, leadership began asking a more fundamental question:

How should spend intelligence scale in the next phase of growth?

Three realities were becoming clear:

- Data complexity was compounding. Supplier fragmentation and taxonomy depth increased the effort required for normalization and classification.
- Speed was becoming a competitive advantage. Proof-of-concept engagements required faster turnaround without reducing analytical rigor.
- Linear scaling had limits. Expanding capacity solely through additional headcount introduced cost variability and operational lag.

The issue was not performance. The issue was sustainability.

The practice needed a delivery model that could absorb increasing complexity while improving responsiveness and cost predictability.

Redesigning the Workflow

Rather than introducing isolated automation tools, the team chose to redesign the spend intelligence workflow at a structural level.

Evalueserve implemented a coordinated agentic architecture embedded directly into the value chain. Four specialized AI agents were deployed within a governed orchestration model supported by human oversight.

The redesigned workflow deployed four specialized AI agents operating under a unified orchestration layer:

1. Supplier Ingestion & Normalization Agent

- Standardized supplier names at scale
- Enriched supplier records with business attributes
- Mapped entities to immediate and ultimate parents
- Reduced ambiguity and rework at the data foundation level

2. Hierarchical Classification Agent

- Applied NAICS/USAIC/custom taxonomies up to Level 6
- Ensured consistency across large, multi-client datasets
- Eliminated classification drift over time

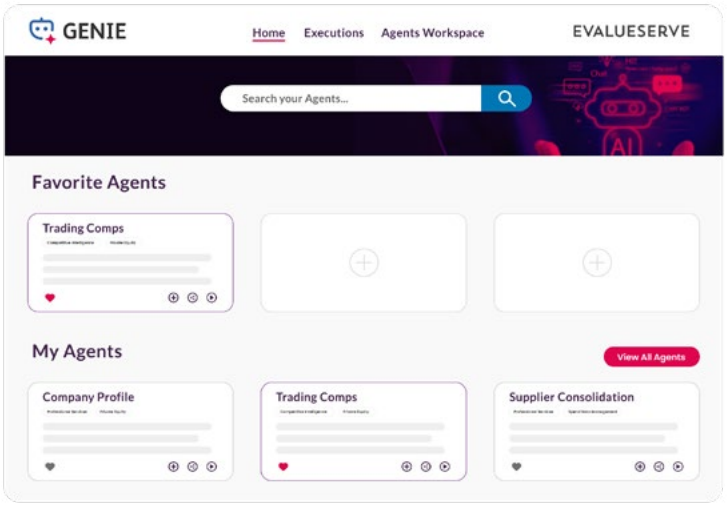

3. Spend Insights Agent

- Generated structured, taxonomy-driven insights
- Enabled benchmarking, opportunity sizing, and optimization analysis
- Supported rapid analytical iteration

4. Market Intelligence Agent

- Identified alternative suppliers and sourcing options
- Enabled faster exploration of consolidation and diversification strategies

Each agent operated within a human-in-the-loop governance framework. Analysts reviewed edge cases, refined outputs, and ensured auditability and client standards were consistently met.

The outcome was a deliberate reallocation of effort:

- Repetitive, rules-based work shifted to orchestrated agents
- Human expertise moved closer to interpretation, judgment, and advisory impact

Client Story

Industrial Solutions Provider Achieves 60% Faster Opportunity Identification with Agentic AI–Led Intelligence

Overview

A global industrial solutions provider partnered with Evalueserve to strengthen its competitive position by identifying commercial opportunities earlier in the sales cycle. While the organization had active lead-generation programs in place, it faced persistent challenges with lead quality and timing.

Evalueserve deployed an agentic AI-led managed services solution, combining autonomous AI agents with deep domain expertise to help the client detect early market signals, prioritize high-potential opportunities, and equip sales teams with decision-ready intelligence. The solution was piloted across two industries and demonstrated the potential to significantly accelerate opportunity identification and improve sales conversion outcomes.

Challenge

The client's commercial teams were dissatisfied with the quality of leads being generated by internal efforts and existing technology-focused partners. Most leads surfaced too late, often after opportunities were publicly announced, leaving little room to shape customer requirements or influence outcomes early in the sales cycle.

Key challenges included:

- **Low-quality and poorly timed leads**, limiting conversion potential
- **Inability to preempt opportunities** before they became widely visible in the market
- **Lack of contextual intelligence**, such as customer-specific drivers, competitive dynamics, and buying triggers, to support early-stage sales engagement
- **Fragmented outputs** from existing vendors, which emphasized AI-driven data signals without sufficient expert interpretation or commercial relevance

The client needed a more proactive, trigger-based approach that could identify and assess early signals of potential demand, before formal announcements, and translate those signals into actionable guidance for sales teams.

Our Solution

Evalueserve implemented an AI-led managed services engagement built on its Insightsfirst platform, anchored by the Opportunity Radar (ORAD) capability. The solution combined agentic AI with Evalueserve's deep domain expertise in industrial markets.

Key elements of the solution included:

- **Agentic AI workflows** designed to continuously scan, identify, and correlate early signals, such as regulatory shifts, investment activity, sustainability initiatives, and customer-specific developments, that could indicate emerging opportunities
- **Domain expert validation and enrichment**, ensuring that AI-identified signals were assessed for commercial relevance, feasibility, and competitive context
- **Prioritized opportunity leads**, delivered at the earliest stages of the buying journey, enabling sales teams to engage customers before requirements are fully defined
- **Decision-ready sales enablement outputs**, including opportunity briefs, competitive positioning insights, and tailored sales narratives that reflected both market intelligence and industry-specific nuances

Unlike other vendors that focused solely on AI-generated outputs, Evalueserve's "mind + machine" approach resonated strongly with the client by pairing automation and scale with analyst-driven interpretation and strategic context.

The initial deployment was launched as a six-month proof of concept across a few industries, with a clear path to scale across the client's broader, multi-sector portfolio. oversight.

Business Impact

Within a few months of deployment, the client began to see meaningful improvements in how opportunities were identified and qualified:

- **Earlier visibility into high-potential opportunities**, allowing sales teams to enter conversations at the start of the sales cycle
- **Significant improvement in lead quality**, driven by trigger-based identification and expert prioritization
- **Up to 60% faster opportunity identification**, reducing time spent chasing low-value or late-stage leads
- **Potential to increase sales conversion rates by up to 3x**, by focusing effort on the most actionable and strategically aligned opportunities

The scalable nature of the agentic AI solution positions the client to extend this capability across additional industries and geographies, supporting sustained competitive advantage in an increasingly complex and fast-moving market.

• In Action

Driving AI Conversations Globally

Evalueserve is actively shaping the dialogue on how AI is moving from experimentation to real business impact. Across industries and regions, these conversations reflect a shared focus on scaling intelligent solutions, strengthening decision-making, and turning innovation into measurable outcomes.

AI Workshop with Google Cloud

At the Google Bogotá AI leadership forum, Evalueserve joined industry and technology leaders to explore how enterprises in LATAM are moving from AI experimentation to scalable, agentic AI solutions across core business functions.

AI Leadership Forum in Boston

Evalueserve convenes Life Sciences & Healthcare leaders at the historic Omni Parker House, Boston, for a focused roundtable on AI-driven competitiveness, operational efficiency, and next-generation forecasting across commercial and R&D functions.

Google Cloud Summit Saudi Arabia

At the Google Cloud Summit Saudi Arabia in Riyadh, Evalueserve joined innovators and enterprise leaders to explore how AI, cloud, and data transformation are reshaping business strategy and growth across the Middle East.

Future Minerals Forum 2026

At FMF26 in Riyadh, our Evalueserve team connected with industry and policy leaders to showcase data-driven strategies accelerating productivity, ESG performance, and decision intelligence across the mining value chain.

Transformation Roundtable

Senior financial services leaders gathered in London, engaging in focused, high-level discussions on geopolitical disruption, risk, and market volatility. The atmosphere reflected candid dialogue, strategic insight, and meaningful peer exchange.

TINtech

Evalueserve at TINtech London Market 2026. Bringing together insurance innovators and underwriting experts to exchange ideas on AI-powered solutions, practical transformation, and collaborative progress across specialty insurance markets.

What's next?

AI is redefining where value sits in organizations. As routine research, synthesis, and drafting are absorbed by intelligent systems, the role of the analyst is shifting decisively toward judgment, context, and decision support.

The real advantage has moved beyond speed, towards the ability to frame the right questions, interpret signals, and translate outputs into clear, actionable insight. At the same time, organizations are realizing that AI does not fix broken processes.

Without clarity on how decisions are made, automation simply scales inefficiency. To lead this space, companies must take a different path: starting with decisions, mapping workflows, surfacing hidden logic and only then deploying AI to orchestrate intelligence across tasks.

Now, how value is measured is also changing, and becoming less about cost per task and more about cost per insight, time to decision and the ability of one expert to oversee greater complexity. Ultimately, AI is embedding intelligence directly into daily workflows, moving organizations from producing reports to enabling better, faster decisions. Embracing this shift will improve efficiency, facilitate more strategic operations, and redefine the role of AI.

Elevate Mag
Issue #5- March 2026

marketing@evalueserve.com
www.evalueserve.com