Data-powered Innovation Review

Wave 11

Capgemini

Foreword

As you go through this new edition of the Data-powered Innovation Review, you'll notice the frost patterns woven through its pages. They're more than decoration. Frost forms when structure and conditions align perfectly. It's a small reminder that innovation, too, depends on precision and connection. That idea inspired the visual identity of this edition: from frost to framework.

You'll see that same sense of alignment reflected in the stories collected here. Each shows how data and AI are becoming part of the framework of how organizations operate: not just powering insights, but driving real outcomes. From agentic systems that collaborate autonomously to responsible AI models that scale with confidence, these examples show what happens when technology meets intent.

One article I'd like to call out in particular is our new Global Data Science Challenge project: **Green Agents of Change.** It aims to explore how agentic AI can do so much more than "just" optimize and transform businesses – it can empower young people to build a greener, more inclusive future. That combination of innovation and purpose is what makes our work meaningful.

So, as you explore these pages, look for the moments where structure meets imagination – where ideas turn **from frost to framework.** That's where meaningful innovation begins.

Niraj Parihar

Chief Executive Officer,
Insights & Data, Capgemini

Editor's note

This edition of the Data-Powered Innovation Review doesn't start with a spark; it starts with a pattern. You'll see it in the frost-like visuals across these pages, in the way AI systems learn and align, and in the growing resonance between humans and intelligent machines. It shows that moment when ideas stop floating in the abstract and finally are made real.

Throughout this eleventh edition, that sense of resonance is everywhere. From articles such as "Blue ocean intelligence" and "Beyond data: Shaping the age of AI products" to "Serendipitous teams" and "The age of anticipation", our authors show how data-powered systems are moving from experimentation to orchestration. Agentic AI doesn't just respond anymore; it plans, reasons, and occasionally outpaces its creators – politely, for now.

Not every story is about algorithms, though. "Bringing 750 years of history to life" from Amsterdam brings things back to the human scale – a reminder that innovation also thrives in culture, collaboration, and community. The city becomes a metaphor for how creativity, data, and design can coexist – proving that great ideas, like canals, flow best when they're connected (and preferably not frozen, for that matter). Of course, as it's the end of the year, we couldn't help coming up with some data and AI predictions for 2026 – against all odds in a highly unpredictable market.

If you prefer your insights with a bit of rhythm, by all means tune in to the Data-powered Innovation Jam podcast – available wherever you get your podcasts. It's the magazine's more talkative sibling: a mix of insight, irreverence, and the occasional odd musical reference – ideal for thinking out loud, or just going with the flow.

Wherever this edition finds you – in winter chill or tropical heat – we hope it offers you a few moments of clarity and ideas that resonate together. Because once the patterns align, innovation doesn't just happen. It crystallizes.

Robert Engels

AI CTO, AI Future Lab Lead,
Insights & Data, Capgemini

Table of contents

Bringing 750 years of history to life – one question at a time

Revolutionizing accessibility using agentic AI

Marc Holtman

Head of Accessibility and Digital Innovation, Stadsarchief Amsterdam

Malgorzata Ilczyk

Microsoft Lead for the European I&D CTO Office, Capgemini

Lucas Puddifoot

GenAI Engineer, Insights & Data, Capgemini

In the heart of Amsterdam lies the Amsterdam City Archive, where you are surrounded by centuries of untold stories. Hidden in sixty kilometers of archive, written in old Dutch, are the stories of everyday Amsterdam: dockworkers, seamstresses, merchants, people who shaped the city we see today. Imagine if there was a way to unlock these histories with a simple question typed on your laptop or spoken into your phone. “Chat with History” makes 750 years of Amsterdam’s history accessible to us all. Innovative Agentic AI technology allows us to remove barriers, making hundreds of years of stories available not just to experts but to schoolchildren, visitors, and anyone curious about Amsterdam’s past. We ensure that the stories that form Amsterdam’s history can be discovered by anyone, transforming the function and reach of the Archive.

The written history of the city, protected by the Amsterdam City Archive, stretches across sixty kilometers of paper and contains records ranging from the earliest history of Amsterdam right up until last week! This data is stored in an enormous range of document types, from medieval charters, church registers, bound correspondence and carbon copies, to construction drawings and all kinds of image material. Between them, they tell stories big and small, with a trove of clues about the people who lived and did business here, and a wealth of knowledge on the city’s administration.

However, these precious documents, which chronicle Amsterdam’s history, have remained largely inaccessible until now, due to the nature of the medium. Those seeking access would need prior knowledge of terminology and the structure of an archive, severely limiting the Archive’s reach.

In recent years, the City Archive has invested heavily in digitizing its paper collection; twenty percent of its documents are now available in digital form, and items that have not yet been digitized can be requested via a scanning-on-demand service. While these advancements boosted the usage of the Archive immensely, the sources remained accessible only to those who could decipher the archaic, handwritten Dutch of the original documents. Handwritten text recognition technology was therefore a key development in the journey of the Archive to our chatbot. It converts historical manuscripts into machine-readable text, opening brand-new possibilities for bringing history directly from the archive to new users.

Chat with History is one of the first projects of its kind, aiming to radically expand the reach of historical archives and connect history with our everyday lives. In this collaboration between the city of Amsterdam, Capgemini, and Microsoft, barriers to education are being removed using the latest AI technology, all in celebration of Amsterdam’s 750-year anniversary. Chat with History makes it possible to ask questions in multiple languages, in natural language and without prior knowledge, to the searchable sources in the archive. Sources that were once only comprehensible to a select few Dutch speakers can now be shared, with responses generated in easy-to-understand language and directly referencing relevant documents.

Users, whether individuals, schools, or travelers, now have the option to interact in written prompts or talk to an avatar. In the future, this will be expanded to include engaging characters that users can communicate with – such as Alewijn, the first Amsterdammer!

> *Innovative agentic AI technology allows us to remove barriers, making hundreds of years of stories available not just to experts but to schoolchildren, visitors, and anyone curious about Amsterdam’s past.*

How have we done it? The solution is built on Microsoft Foundry, Agentic Services, and Microsoft Agent Framework, which work together in a robust multi-agent system.

This system includes agents such as query handlers, document retrievers, and summarizers, all orchestrated to provide seamless interactions with historical data. Microsoft Agent Framework plays a crucial role in coordinating these agents, ensuring efficient processing and accurate responses. This multi-agent framework enables seamless integration of new functionalities in the future. For instance, authentic images such as old city maps could be added to the chatbot's response, with the relevant agent handling the integration of these features.

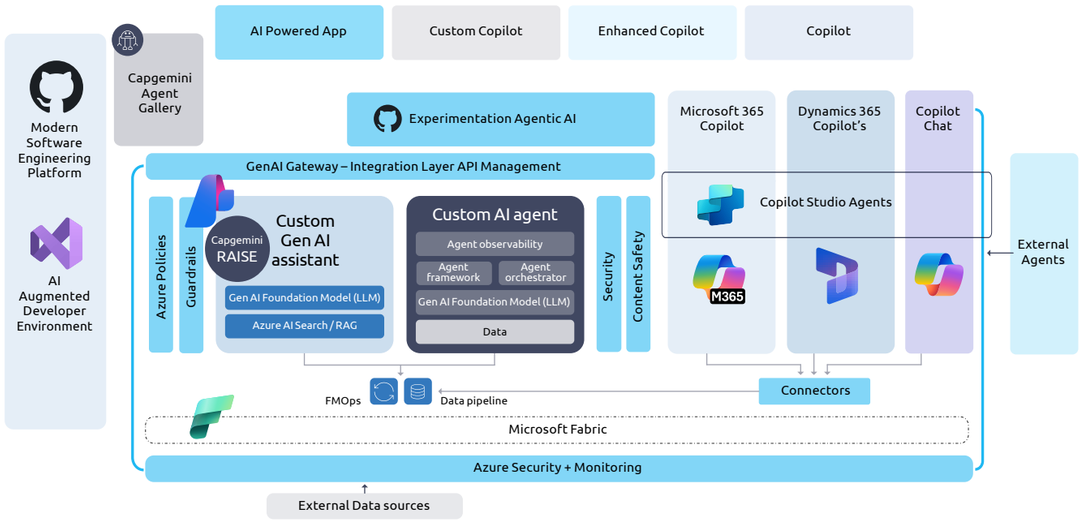

As a shared resource, the Archive embodies the right of everyone to access the wealth of information it holds, a vision that is now being realized more fully than ever before. The most revolutionary change in the Archive's use lies in its ability to reach and include diverse groups, including those who may never have known about it. Chat with History allows them to easily uncover answers about the history of Amsterdam, or perhaps even gain insight into their own family or neighborhood. Moreover, experienced researchers who already use the archive will benefit from a cutting-edge tool supporting them to search more effectively for people, places, and events within the archive, to uncover meaningful connections across data from multiple sources.

What kind of stories has our project uncovered? One letter tells the story of a merchants' hardships and a city standing up for their own in 1574: we came across two brothers, Mertyn and Heynric Gherytsz of Hamburg, who wrote on behalf of their fellow townsman Jan Lambertsz. Lambertsz was a merchant transporting grain by river from the Elbe, destined for sale in Dutch ports. Even before arrival, he suffered heavy losses, losing half of his cargo when his ships loaded with grain were seized. His frustration only increased when he was told he could not sell his wares, despite holding official letters of permission to buy and export grain. Outraged at the treatment, the Gherytsz brothers and the Hamburg council sent a formal letter of protest to the authorities, pointing out that Hamburg's merchants had always paid the required tolls and duties, and had even contributed extra payments for maritime infrastructure such as beacons. They argued that it was unjust to deny Lambertsz permission when consent had already been given, and that he had already borne great financial harm. After nearly 450 years, Jan Lambertsz is alive again through Chat with History – and we can still relate to him and his struggles with Dutch bureaucracy!

The Archive no longer functions solely as a vault; it is opening up, creating a resource where people of all backgrounds can become curious about the history of the world around them. The significance of this work has been recognized far and wide: other archives in the Netherlands and abroad have expressed interest, and the project has earned award nominations. Most notably, Chat with History was presented to Satya Nadella, the global CEO of Microsoft, who highlighted our project in his keynote speech at the Microsoft AI Tour. This recognition underscores the potential of this initiative to redefine how history is accessed, understood, and shared.

This is just the beginning of a journey that reimagines how we engage with history. By transforming the Archive into a living, accessible resource, we are opening doors for curiosity, connection, and discovery for people everywhere. As recognition and interest continue to grow, we look forward to inspiring new ways to explore the past, empowering individuals and communities to uncover stories that shape their identities and understanding of the world. Together, *we are not just preserving history, we are bringing it to life for future generations.*

Start *innovating* now

Digitize to access

Use previously unused sources of information to unlock new potential. As soon as you have saved a document, you have created an archive!

Deploy smart AI agents

Unleash multi-agent AI to answer questions, retrieve documents, and connect the dots instantly. Adapt this framework for customer service, research, or compliance.

Design for everyone, everywhere

Build multilingual, multimodal, and inclusive tools. Let anyone – student, traveler, or pro – explore, learn, and contribute in their own way.

#AIForAccessibility #InclusiveTech #DigitalArchives #HistoryUnlocked #750YearsOfAmsterdam #PreservingHistory #EdTech #DemocratizingKnowledge

Serendipitous teams

The future of agentic marketing teams

Neerav Vyas

Vice President – Customer First Insights & Data, Capgemini

The next great marketing team might not sit around a table but in a shared algorithmic space. Blending human intuition with agentic intelligence, brands can now summon the strategic insight of Ogilvy, the creativity of Bernbach, and the rigor of Drucker – all before the coffee's done brewing.

A handful of scientists solved an impossible problem by making a radical decision: they wouldn't hire anyone. In doing so, they have revolutionized how teams work, especially on Madison Avenue.

Traditional thinking – you can have it good, fast, or cheap; pick two. Yet, what if you could pick all three? What if you could generate a thousand ideas, test them all, and find the perfect one, all in the time it takes to brew a pot of coffee? A handful of researchers at Stanford stumbled on a way to do just that. Their breakthrough is not just a story about science. It's a blueprint for a revolution in any field that relies on collaboration. Its next logical destination is marketing.

***"Good science happens when we have deep, interdisciplinary collaborations where people from different backgrounds work together, and often that's one of the main bottlenecks and challenging parts of research"*— James Zou, PHD Associate Professor of Biomedical Data Science at Stanford.**[1]

Instead of hiring or flying experts in, they decided to build them. They engineered an "agentic squad," where each agent was a cognitive model of world class expertise in a specific domain. There was an AI immunologist, an AI computational biologist, and an AI machine learning specialist.[2] All brought together in a virtual space to think, challenge, and create.

To understand the power of this model, we must shift our thinking from AI as a tool for speed to AI as a vessel for expertise. The Stanford team's success didn't come from simply processing data faster; it came from simulating the collaborative friction and insight of a multidisciplinary team.

Stanford's Virtual Lab of AI Agents to design new SARS-CoV-2 Nanobodies reinvents and accelerates the ability to conduct groundbreaking research

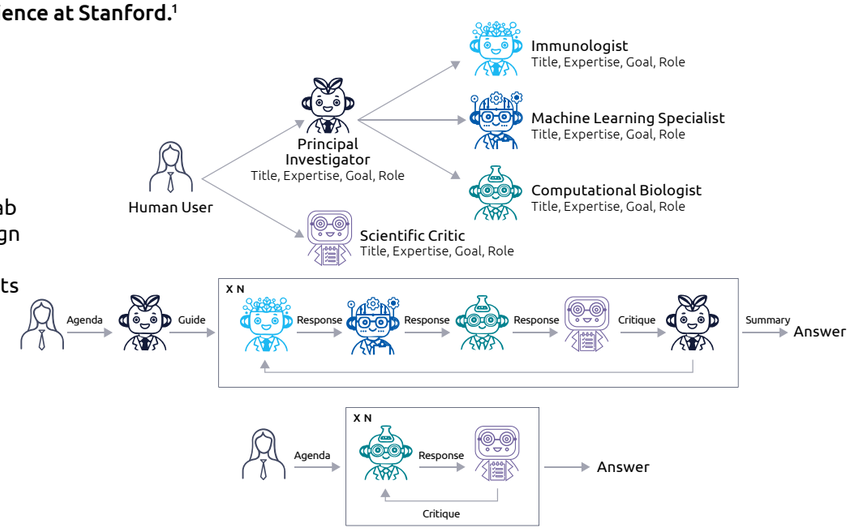

Here's how the agents worked:

- An AI principal investigator (AI PI) managed and ran the team coupled with an AI critic who poked holes in the team's thinking.
- The virtual team held regular meetings where agents generated ideas and engaged in rigorous discourse. They also held one-on-one meetings, allowing human and virtual members to meet individually to discuss ideas or ask questions.

The agents debated, cross-referenced findings, and built upon each other's work instantly, without ego or scheduling conflicts. The true power was their ability to synthesize disparate, world-class points of view. They created a "Davos of data," a summit of the best minds, instantiated on demand.

Now, imagine this same concept applied to the marketing world's oldest fantasy: having the industry's legends guiding your brand. Most companies can't hire the world's greatest brand strategist, a Cannes Lion-winning creative director, and a top MBA with AI expertise. But with an agentic team, they can simulate their collective genius.

[1] https://news.stanford.edu/stories/2025/07/ai-virtual-scientists-lab-llms
[2] https://news.stanford.edu/stories/2025/07/ai-virtual-scientists-lab-llms

The human marketing director of tomorrow won't just manage a team; they will conduct a symphony of virtualized experts:

- **"Ogilvy" – the brand strategist:** This agent is the brand steward. Trained in the philosophies, case studies, and writings of legendary brand builders. It asks the foundational "why" to generate core brand platforms, positioning statements, and strategic narratives that resonate deeply.
- **"Bernbach" – the creative director:** This is the conceptual powerhouse, trained on decades of award-winning advertising and design theory. It takes the strategic platform from "Ogilvy" and translates the big ideas to the narrative arcs and supporting creative concepts, assets, and taglines that define their campaign ideas.
- **"Drucker" – the business strategist:** This agent is the voice of reason and results. Modeled on the analytical rigor of top business minds and data scientists, it pressure-tests the creative concepts against business realities. It projects ROI, analyzes customer lifetime value (LTV), and ensures that brilliant creative is also effective business.

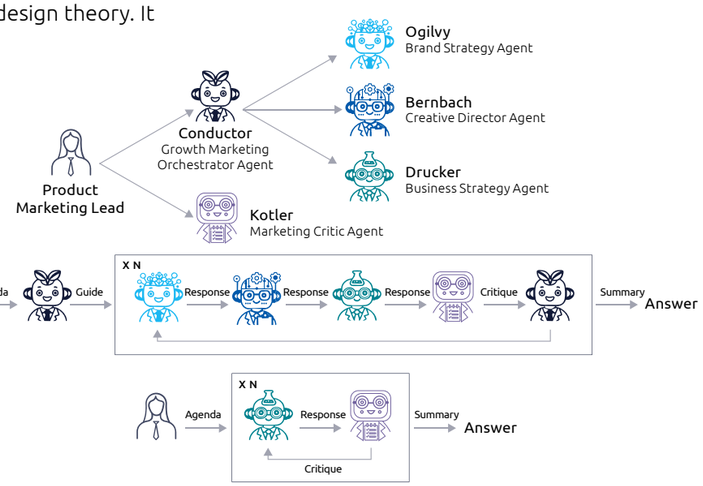

The very same approach will be leveraged by **agentic forward marketing teams** to dramatically accelerate the way they ideate and execute

This new paradigm changes the role of the human marketer. Their job shifts from the slow, linear process of managing people to the dynamic, real-time curation of an expert debate.

- **Old workflow:** Hire a strategy firm and wait weeks for a PowerPoint deck. Brief a creative agency and wait for three ideas. Ask an analyst for a report and get a spreadsheet a week later. Then, struggle to make it all coherent.
- **New workflow:** The marketing director issues a directive: "Team, we're launching a new product line. 'Ogilvy,' give me three revised brand architectures based on consumer values in this sector linked to the fastest growing segments. 'Bernbach,' generate five distinct creative territories and supporting assets. 'Drucker,' run a preliminary ROI and market-fit analysis on each. Conductor, run the meetings with Kotler providing constructive criticism. The goal is to review with our in-house teams and agency partners ASAP."

The human lead remains the ultimate arbiter, the conductor, the editor-in-chief. Their value lies not in execution, but in judgment. They ask insightful questions, spot the flaws in a simulated expert's logic, and provide the final strategic, creative, and ethical sign-off.

The scientists weren't replaced by their AI teammates; they were elevated. The same will be true for marketers. The future of marketing won't be defined by the talent you can hire, but by the expertise you can build, orchestrate, and collaborate with. It won't belong to the teams that can do the most work, but to the teams that can ask the best questions. It will be a world where the most valuable skill is not execution, but judgment. And in that world, the most powerful creative force won't be a single person, but a seamless, symbiotic partnership between human insight and artificial intelligence.

> *The fantasy of having the world's best minds in your conference room is no longer out of reach. The room is now virtual, and the minds are algorithmic.*
>
> **Source:** https://news.stanford.edu/stories/2025/07/ai-virtual-scientists-lab-llms

Start *innovating* now

Instantiate the talent you can't hire

The war for A-list talent is a costly, often losing battle. Instead of competing for the same handful of experts, identify your most critical talent gap and build an agent to fill it. Don't start by trying to create a full virtual dream team; start with a single expert. Is your marketing team held back by a lack of data science resources? Is your supply chain strategy hobbled by a shortage of predictive modeling experts?

Accelerate a single, high-friction process

Every organization has a critical workflow that is notoriously slow, bogged down by manual analysis and long review cycles – be it market research synthesis, campaign performance reporting, or competitive intelligence gathering. These processes are the perfect incubator for your first human-agent team. The goal isn't to automate the entire department, but to dissolve a specific, high-value bottleneck.

De-risk one high-stakes decision

Major strategic initiatives – a new product launch, a market entry, a significant re-branding – are fraught with uncertainty and internal biases. Before committing millions in budget, use an agentic team as a dedicated, virtual "red team" to pressure-test your plan. This moves AI from a simple productivity tool to a powerful strategic sparring partner.

#SerendipitousMarketing #Serendipity #AgenticAI #AgenticSerendipity #AgenticCX #AgenticExperiences #AgenticMarketing #AgenticBrandBuilding #TheFutureOfMarketing

Blue ocean intelligence: Where generative meets agentic AI

Turning AI competition into AI creation

Yashowardhan Sowale

CTIO I&D India, VP, Global Architects
Head Insights & Data, Capgemini

As generative AI becomes ubiquitous, the real frontier lies in **agentic intelligence** – AI that can plan, reason, and act. By fusing the creativity of Gen AI with the autonomy of agentic AI, enterprises can move from competing in crowded "red oceans" to creating uncontested "blue oceans" of value, where machines collaborate, adapt, and innovate alongside humans.

From red oceans to blue intelligence

The past year saw an arms race of chatbots, copilots, and model releases. Everyone talked to their AI – and everyone got roughly the same answers. True differentiation now lies beyond conversation. Agentic AI shifts the focus from **generation to execution** – from producing words to producing outcomes. In a "blue-ocean" world, competition fades because new forms of value emerge. When generative and agentic systems converge, AI stops reacting and starts reasoning. It anticipates problems, orchestrates resources, and creates experiences so personal and adaptive that no single model or human team could replicate them.

> *This isn't about better AI; it's about a new market.*

The four shifts

Escaping the red ocean of commoditized generative AI requires rethinking the entire value chain across four dimensions:

1. **Eliminate static learning cycles**
   Traditional models grow stale within months. Self-optimizing agents with continual learning evolve in real time, updating themselves as the business changes.
2. **Reduce computational drag**
   Move from GPU-hungry giants to agile small language models and hybrid architectures that blend LLM scale with edge efficiency – lighter, faster, cheaper.
3. **Raise autonomy through reinforcement**
   Techniques such as RLHF and RLAIF teach agents to reason, decide, and act – transforming passive generators into proactive decision-makers.
4. **Create new value through collaboration**
   Multi-agent frameworks such as CrewAI, LangGraph, and Model Context Protocol enable specialized AIs to negotiate, plan, and execute together. This is where entirely new service categories appear: autonomous supply chains, self-healing IT, adaptive customer ecosystems.

Differentiation between Red and Blue Oceans

The way to succeed is to produce new value derivation, monetization and innovation focusing on making the competition irrelevant.

• Operate within established market dynamics of commoditized chatbots and minimal differentiation • Exploit existing demand curves (i. e., play in same market space) • Engineer the value-cost equilibrium with larger training costs and static learning • Constrained in semi autonomous AI • Pretrained intelligence within limits	VS	• Redefine market boundaries • Fusion of generative and Agentic AI • Execution over conversation • Adopt strategic learning • Orchestrate strategic coherence across the enterprise via collaborative AI • Agentic intelligence networks

Breakthroughs in action

An American drone company shows what adaptive learning looks like in flight. Its drone-delivery network uses meta-reinforcement learning to optimize every route, weather pattern, and landing – millions of micro-lessons shaping each new mission.

A global automotive company , partnering with Intel's Loihi neuromorphic chip, proves that cognition can be energy-efficient. Spiking neural networks make split-second driving decisions using a fraction of the power of traditional deep learning – distinguishing a plastic bag from a child chasing a ball.

A global cybersecurity company employs collaboration at scale. Its multi-agent cybersecurity fabric uses graph neural networks to model relationships between agents: one tracks network anomalies, another monitors endpoints, a third aggregates threat intelligence. Together they detect multi-stage attacks invisible to any single agent.

Each case shows the same principle: when agents learn, reason, and cooperate, performance compounds.

From experiments to ecosystems

These breakthroughs mark the dawn of an entirely new design philosophy: **intelligence as infrastructure.** Instead of monolithic models locked inside applications, AI becomes a distributed network of capabilities that can be combined, reused, and orchestrated across the enterprise.

It's a quiet revolution – one that turns AI from a differentiator into the foundation of the business itself. The question is no longer how to use AI, but how to organize around it: to build ecosystems where human creativity and machine autonomy reinforce each other in real time.

The enterprise opportunity

For enterprises, this fusion of generative and agentic intelligence changes the economics of innovation. AI ceases to be a tool and becomes a colleague – one that scales expertise, optimizes operations, and continuously learns from every interaction.

Clients are already exploring these blue oceans through our **rAIse AI** platform, building domain-specific agents that manage workflows, reallocate resources, and optimize processes autonomously. The goal isn't just efficiency; it's to redefine what a product or service is.

Imagine autonomous supply-chain conductors negotiating with logistics agents in real time. Finance advisors that not only forecast cash flow but adjust it dynamically. Customer-experience ecosystems where brand, data, and operations all learn and act together.

This isn't about better AI; it's about a new market logic – where adaptability itself becomes the competitive advantage.

Beyond competition: The new logic of enterprise AI

When generative and agentic intelligence combine, the impact goes far beyond efficiency. Enterprises gain a new kind of agility – one where systems learn, reason, and collaborate as seamlessly as the people who guide them.

Instead of racing to build faster copilots, organizations can design **autonomous ecosystems** that continuously adapt, optimize, and create: supply chains that reroute themselves, IT systems that self-heal, and customer experiences that evolve in real time.

The companies that master this fusion will define the standards others must follow. The rest will remain in the red, chasing diminishing returns. In the age of agentic intelligence, the real winners won't fight the tides of the red ocean – they'll learn to sail the blue.

Start *innovating* now

Fuse creativity with autonomy

Combine generative AI's ideation with agentic AI's execution to deliver outcomes that are creative, contextual, and continuously learning.

Build collaborative agent networks

Move from single-model applications to multi-agent ecosystems where specialized AIs coordinate across functions – design, supply, finance, service.

Create your blue ocean

Use this fusion to define entirely new service categories and outcome-based business models that make competition irrelevant.

#AgenticAI #AutonomousAI #BlueOceanStrategy #AIInnovation #EnterpriseAI #NextGenAI #Capgemini #AIEcosystems #ProactiveAI #AIValueCreation #AITransformation #IntelligentAutomation #AIOrchestration #BusinessReimagined #FutureOfAI

Move 37

How agentic AI turns back-office drudgery into bold process breakthroughs

Dinand Tinholt

Vice President, Data and Analytics, Capgemini

Remember when AlphaGo shocked the world with Move 37? That one move nobody saw coming, yet it changed the entire game. Businesses today are stuck buffing their back-office processes to a shine while the competition is learning to play on a different board entirely. The real leap isn't squeezing out another percent of efficiency. It's daring to reimagine the game. Data and agentic AI are the new board. Move 37 thinking is the way to play.

Stop refining old processes. Start reimagining them. Finding your Move 37 means breaking free of efficiency worship and unleashing process innovation that plays a different game altogether.

The world of business loves efficiency. Shave a minute here, cut a cost there, automate just enough to avoid complaints. It's the corporate equivalent of organizing your junk drawer and calling it transformation. But history doesn't remember the people who got to three percent gains. It remembers the bold leaps.

In 2016, AlphaGo, an AI system built by Google's DeepMind, stunned the world by beating Lee Sedol, one of the greatest Go players in history, at a game humans had dominated for thousands of years. Go, unlike chess, has more possible positions than atoms in the universe. It's a game of intuition, not just calculation. Experts had predicted machines wouldn't crack it for another decade. Then came Move 37 in the second game. AlphaGo placed a stone in a position so unexpected that commentators thought it was a mistake. Lee Sedol left the room. Spectators gasped. The move violated centuries of accumulated wisdom about how the game should be played. It wasn't efficient by any traditional measure. It wasn't even logical to human eyes. Move 37 broke every expectation, every model, every best practice. And that's precisely why it won. The machine hadn't learned to play Go better. It had learned to play Go differently.

That's what businesses need now. The real breakthroughs won't come from tuning yesterday's processes. They'll come from redefining what a process even is.

Enter data and agentic AI. For decades, back-office routines have been designed like assembly lines: rigid, rule-bound, and optimized for predictability. But processes aren't predictable anymore. Supply chains bend, customers shift, regulations zigzag, and your CFO wakes up with a new important KPI every quarter. Traditional workflows crack under that pressure.

Agentic AI doesn't just keep up. It changes the game. These systems observe patterns, anticipate needs, and take initiative. Instead of waiting for humans to script every conditional, agents self-adjust and negotiate between competing goals. The expense workflow doesn't just process receipts. It flags anomalies, proposes corrective action, and nudges behaviors before they spiral into losses. Order-to-cash doesn't just move invoices down a pipe. It becomes a living system that optimizes cash flow in real time.

The old playbook would call this optimization. That's the wrong word. Optimization is what you do to an engine built last century. This is reinvention.

Consider finance teams still obsessed with cutting days off closing cycles. Admirable, but irrelevant when your competitors are deploying AI agents that collapse the cycle altogether. Think of HR departments proud of digital onboarding portals while agentic systems are already predicting attrition risk and reshaping talent flows on the fly. Or supply chains that still sweat over visibility while AI-driven agents negotiate freight rates autonomously. These aren't incremental steps. They're category shifts.

Here's the pattern: most companies are still asking how to make their existing processes faster. The right question is whether those processes should exist at all. That three-way match in accounts payable? An artifact of a world where humans had to verify paper trails. Exception handling in procurement? A Band-Aid for systems too rigid to handle variance. Escalation workflows in customer service? Admission that your front line lacks the intelligence to solve problems. Agentic AI doesn't make these processes better. It makes them obsolete.

This is the Move 37 mindset. You don't win by playing the old game better. You win by changing the game so the old players look quaint.

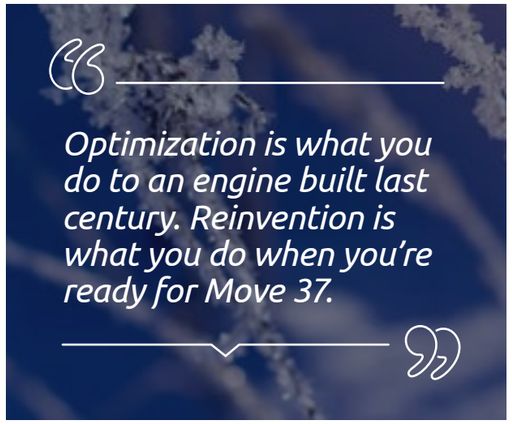

The resistance is predictable. Finance leaders will say their processes are too complex, too regulated, too embedded in the fabric of the business to reimagine. HR will point to compliance requirements. Operations will cite the need for human judgment. These are the same arguments that kept ledgers in bound books long after spreadsheets arrived. The same logic that defended fax machines in the age of email. Comfort disguised as caution.

But the real risk isn't moving too fast. It's moving too slow while pretending you're being prudent. Every quarter spent optimizing legacy processes is a quarter your competition spends building something that doesn't need optimizing because it was designed for a world that actually exists.

The technology is ready. Data platforms are no longer fragile silos but adaptive lakes and meshes. Large language models aren't clumsy autocomplete toys anymore. They're orchestrators of business logic. Agentic frameworks tie it together, giving companies an operating system where processes evolve on their own. Not chaos. Not magic. A system that adapts in ways you didn't script but that you can still steer.

And the economics have flipped. Five years ago, reimagining a core process meant armies of consultants, years of implementation, and budgets that required board approval. Today, agentic systems can be piloted in weeks, scaled in months, and iterated continuously. The barrier isn't technology or cost. It's imagination.

Markets are already moving. The companies pulling ahead aren't the ones with the slickest dashboards. They're the ones brave enough to break the glass, toss the manual, and say: let's rebuild this process as if it were invented today. They're the ones willing to look foolish in the short term to look prescient in the long term.

That's the uncomfortable truth. Efficiency worship feels safe. It looks good in quarterly reviews. It generates tidy before-and-after metrics that executives can present with confidence. But it's the kind of safety that leads to slow irrelevance. Move 37 thinking is messy, unpredictable, and at first, it looks absurd. Until it changes everything.

The choice is stark. You can keep polishing processes designed for a world that no longer exists. You can celebrate marginal gains while the ground shifts beneath you. You can tell yourself that steady progress is strategy. Or you can accept that the game has changed, and the only way forward is to change with it.

Stop refining. Start reimagining.

Start *innovating* now

Play a different game

Don't ask how to shave another percent off your processes. Ask how agentic AI can make the process itself irrelevant by creating a better one altogether.

Unleash your data exhaust

Your messy, unused operational data is fuel for adaptive agents. Stop treating it like leftovers and start treating it like the spark for process reinvention.

Dare to deploy agents where it hurts

The biggest breakthroughs aren't in the shiny front-end but in the messy back office. Finance, HR, supply chain – those are the places begging for their Move 37 moment.

#Move37 #DataPowered #AgenticAI
#ProcessReinvention #BoldInnovation

Anticipating tomorrow from above

The rise of space data as a service

Sudarshan Sahu

Senior Consultant,
Capgemini

Two years ago, space-data-as-a-service (SDaaS) was a bold promise, driven by falling launch costs and swelling satellite constellations. Now, that anticipation has matured into reality: SDaaS is no longer just raw imagery, but actionable intelligence woven into agriculture, disaster response, finance, and urban sustainability. Startups and governments alike are tapping federated constellations and AI-powered analytics to move from observing Earth to predicting it. The "view from above" has become an indispensable layer of our digital infrastructure, enabling climate foresight, resilient supply chains, and even planetary health models. Space data is no longer a distant future; it's part of the operating system of tomorrow's decision-making.

Humanity has always looked up at the stars for guidance. Sailors once navigated oceans with starlight, farmers planted crops by lunar cycles, and philosophers gazed skyward for inspiration. But today, the stakes are different. We're no longer just charting oceans or tracking seasons, we're navigating a world defined by climate volatility, fragile supply chains, and escalating disasters. And now, the cosmos is once again helping us make decisions. But this time through the lens of data. Welcome to the era of **space data as a service (SDaaS),** where insights from orbit are streamed to Earth in near-real-time, transforming how we anticipate and respond to challenges.

Space data as a service (SDaaS) has evolved from raw imagery streams into an anticipatory infrastructure. Compared to just a few years ago, today's systems offer near-real-time insights, AI-driven edge processing in orbit, wider accessibility, and growing ethical frameworks. This shift represents more than technological progress. It signals the dawn of an era where technology doesn't just record the present, it foresees what's coming. Satellites, sensors, and AI now form a global network of observation, gathering data with remarkable speed and accuracy, enabling earlier responses and more informed decision-making. Further, this shift empowers communities, governments, and industries to predict risks, manage crises before they escalate, and make better decisions for sustainability and resilience.

What has changed, what enables anticipation

Several shifts have made SDaaS more anticipatory, more pervasive, and more consequential:

Latency, refresh rate and cost efficiency: Satellites now capture more data, more frequently, with better sensors. Miniaturized Earth observation platforms (e.g. CubeSats) have grown in number; reusable launch vehicles have brought down launch costs (fallen from **$65,000 per kilogram to $1,500 per kilogram – more than a 95% decrease**). These combined make it viable to have constellations that cover many locations multiple times per day. Imagery, thermal, radar, spectral data, all updating faster.

Edge computing / on-orbit processing and AI: Previously, a lot of data processing was centralized: download raw imagery, ship to cloud or data centers, then analysis. In 2025, more intelligence is being pushed upstream: satellites or near-space platforms doing pre-filtering, anomaly detection, change detection. Only relevant or priority data is downlinked or alerts triggered in near-real-time. Further, data no longer flows from satellites only; it mixes with IoT, drone imagery, citizen sensors, and ground weather stations. These layers allow better contextualization. In April 2025, Fujitsu launched "**Space Data Frontiers**" focusing on "space data on-demand," which uses AI-equipped satellites for real-time imagery transmission, and "space weather" to improve predictions of solar flare impacts.

> *Space data is no longer about looking down at Earth, it's about looking ahead for humanity. In the age of anticipation, the real power of space data is not in seeing more, but in knowing sooner – and acting together.*

Hyperscale cloud architecture: The vast storage and processing demands are met by planetary-scale cloud data lakes, often leveraging specialized **space cloud** regions near ground stations. These environments offer APIs, machine learning (ML) toolkits, and secure access portals, making data consumption look exactly like calling a standard web service. AWS has boosted its **Ground Station as a Service (GSaaS) Partner Program** by teaming up with Kongsberg Satellite Services (KSAT). Integrating KSAT's global network of 40+ sites and 200 antennas with AWS's cloud platform, this partnership lets satellite operators access scalable ground infrastructure without upfront costs.

Broader access and affordability: Lower costs and more providers means more access. Many users typically were big governments, large corporations, and NGOs with funding. Now, more local governments, smaller

businesses, even community groups can subscribe to services or buy specific SDaaS "modules" (e.g., crop health, disaster risk for a region). In 2025, Israeli startup **CropX partnered with NASA Harvest to deliver soil analytics** by combining satellite data with ground sensors. This collaboration enhances global crop forecasts, cuts waste, and boosts profits – showcasing how legacy space agencies increasingly team up with startups instead of developing all technology internally.

Ethics, regulation, standards: The past few years saw a lot of discussion about data sovereignty, privacy, orbital debris, spectrum allocation; but many gaps remained in enforceable regulation and standard practices. Now, more concrete regulation and standards emerge: countries asserting territorial / sensing rights, agreements around debris mitigation, norms for usage and liability for false alarms. Users care about provenance, accuracy, potential misuse, and regulatory risk more than before.

Shift from proof-of-concept to operational use: Many SDaaS examples used to be proof-of-concepts, academic pilots, or narrow test deployments; now, several are fully operational: recurring contracts, SLAs, integration into disaster management, agriculture planning, insurance underwriting, supply chain optimization, environmental monitoring.

Space data as a service (SDaaS) delivery model

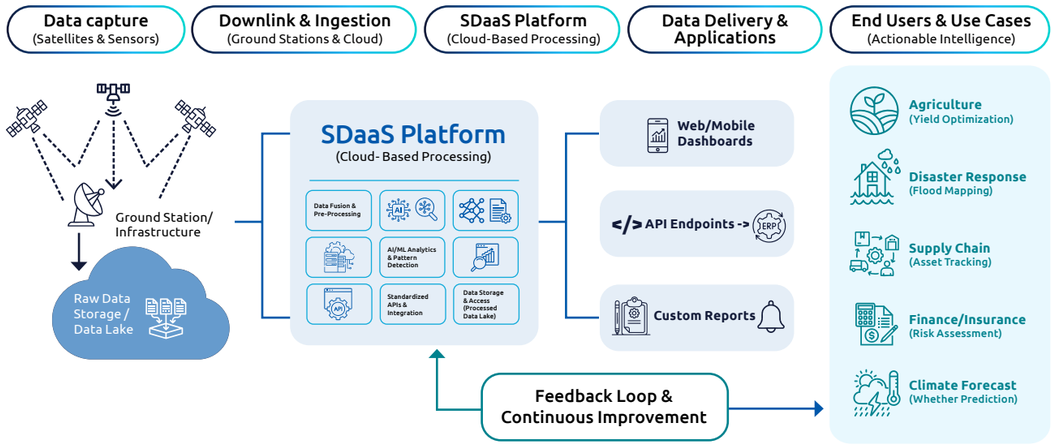

Anticipation isn't coming – it's already here

SDaaS delivers continuous streams of extraterrestrial intelligence: satellite imagery, weather dynamics, crop health data, and maritime traffic – packaged as on-demand services for governments, businesses, and communities. SDaaS has become more anticipatory, pervasive, and impactful, reshaping how decisions are made across critical domains:

Wildfire early warnings (California): Thermal and optical satellites now detect abnormal heat signatures within minutes. AI models cross-reference vegetation and weather data, issuing community alerts before smoke is even visible. **NOAA satellites frequently identify fires prior to 911 reports,** while real-time smoke plume tracking guides airborne firefighting with unprecedented precision.

Carbon monitoring (Europe): The European Space Agency, in collaboration with Capgemini and others, uses **satellite-fed carbon inventories to provide near-real-time emission tracking.** This innovation is transforming climate policy, empowering governments to act faster and enforce industrial accountability.

Agriculture optimization: Farmers now receive daily SMS alerts from hyperspectral satellite scans, predicting drought stress or pest infestations to protect yields. In Odisha, the World Bank-funded Integrated Irrigation Project applies **satellite data and crop modeling to deliver timely irrigation advisories** to nearly 125,000 smallholder households across 15 districts, covering 128,000 hectares, strengthening resilience and productivity against climate shocks.

Supply chain resilience: Maritime constellations track vessels and weather systems simultaneously, enabling shippers to reroute days ahead of storms. This proactive foresight saves millions in damages, delays, and fuel costs while strengthening global trade resilience. **DHL's SmarTrucking program** integrates telematics, sensors, and IoT capabilities with satellite data for optimized routing. The system provides drivers with the most efficient routes based on real-time data, and continuously tracks vehicle and cargo status.

Disaster preparedness and response: In July 2025, NASA and ISRO launched the NISAR mission – one of the most comprehensive Earth observation satellites ever. **Scanning nearly all land and ice surfaces every 12 days,** NISAR will track glaciers, land shifts, forests, wetlands, and agriculture, supporting vital scientific and humanitarian efforts.

Considering all of the facts

Space data as a service doesn't just give us the ability to see more, it gives us the ability to act more wisely. It becomes a foundational layer of society: embedded in how cities plan, how farmers choose, how disaster agencies mobilize, how insurance models price risk, how families decide.

What matters is not only having sensors and satellites, but how we manage uncertainty, build trust, distribute access, and embed ethical guardrails. The difference between a prediction saving lives or being ignored depends on how accurate, timely, accessible, and credible it is.

Start *innovating* now

Smart communities, safer futures

Mobile apps powered by SDaaS deliver early flood and drought warnings to remote areas, while schools empower kids to decode climate data – turning reactive planning into proactive resilience.

Space data as public infrastructure

Policies that treat space data as part of public infrastructure – with open data components, but with clarity on ownership, liability, and privacy.

Next-gen tech for early action

Satellites with AI autonomy, sensor swarms, and smart agents enable real-time monitoring and automatic response, building global climate and disaster anticipation networks.

#DatapoweredSpace
#Data4Good #SpaceDataAsAService
#SpaceDataEcosystem

Beyond data: The rise of AI products

Where intelligence becomes reusable

Niklas Rittmann

AI Solution Architect,
Capgemini

The next era of enterprise value won't be driven by apps or data alone, but by AI products – models, prompts, and agents treated as reusable, governed assets. With the right foundation of data, lifecycle management, and trust, organizations can turn intelligence itself into a scalable product, and innovation into an enterprise habit.

Enterprises that learned to treat data as a product – standardized, governed, and reusable – built the foundation for the next great leap: AI products. Models, prompts, and agents are no longer side experiments; they're reusable, scalable assets that demand the same rigor as data. The next era of enterprise value won't be driven by applications or analytics alone, but by governed marketplaces of AI intelligence – modular, trustworthy, and seamlessly integrated into enterprise architectures. The winners will be those who move beyond pilots and prototypes to build production-grade ecosystems of reusable AI, where innovation compounds and trust scales.

From data to AI products

For years, organizations invested in making data discoverable, reliable, and reusable – a discipline that unlocked new business value through data platforms, catalogs, and marketplaces. Now, with the rise of generative and agentic AI, intelligence itself is becoming reusable. Copilots refine writing, autonomous agents streamline workflows, and AI models analyze data across every corner of the enterprise. But these capabilities can't remain one-off experiments. Like data products before them, AI assets must be standardized, governed, and designed for reuse. The same foundation that made data reliable will now determine the quality and trustworthiness of AI. Enterprises that treat AI artifacts – from fine-tuned models to prompt libraries and autonomous agents – as durable enterprise assets will accelerate innovation without sacrificing control.

Challenges and discipline

AI products introduce new complexity. Unlike static data, they evolve through retraining, feedback, and refinement. Their value depends not just on accuracy but also on interpretability, fairness, and alignment with business outcomes. To manage this dynamic environment, enterprises must move beyond ad hoc experimentation toward formal product lifecycles – complete with ownership, quality metrics, and transparent governance. This discipline ensures every model or agent can be safely reused across functions. Standardization, lifecycle management, and enterprise-wide discoverability turn isolated efforts into building blocks of intelligence. The result is a system where innovation is fast but controlled – one that balances creative exploration with ethical and operational accountability.

Real-world momentum

Across industries, treating AI as a product is already unlocking measurable value. Healthcare organizations are assembling libraries of domain-specific AI models and prompt sets to support GxP risk analysis and regulatory compliance. These assets are versioned, governed, and continuously monitored – not only for accuracy and bias but also for compliance with stringent standards. The outcome: reduced manual effort, fewer errors, and faster, more consistent reporting. In the automotive sector, autonomous-driving innovators are turning perception, prediction, and decision-making algorithms into modular AI components. These can be refined and reused across vehicle platforms, improving safety and speeding up innovation without starting from scratch.

In both examples, the key isn't experimentation – it's reusability, accountability, and lifecycle management.

> *The next wave of enterprise value won't come from using AI – but from producing it.*

Managing the AI product lifecycle

Leading organizations now define clear ownership for every AI product, assigning teams accountable for its evolution. They establish formal lifecycles, treating each AI asset as a living entity that is versioned, monitored, and continuously improved. They implement catalogs and internal marketplaces where approved models, prompts, and agents can be easily discovered and reused. This avoids duplication and ensures consistent quality across the enterprise. Trust becomes a measurable attribute. Beyond accuracy, leaders track robustness, fairness, and explainability, integrating ethical AI into daily operations. The most forward-looking enterprises also design for modularity from the start – creating AI components that can plug into existing workflows and architectures. This composable design ensures interoperability, performance, and scalability across the business.

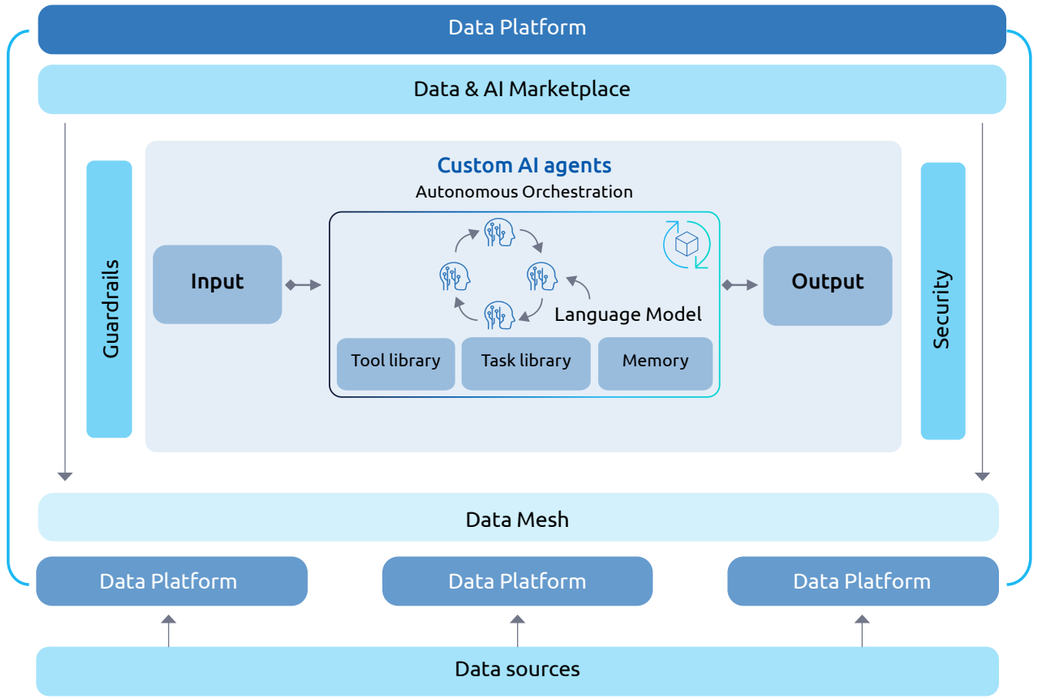

Scaling intelligence: AI factories and marketplaces

The market is rapidly adapting to this new paradigm. A surge of platforms now supports AI product management – bringing together model training, prompt engineering, and agent orchestration in unified environments. Some organizations are already building AI factories: governed environments that produce, test, and distribute AI products at scale. Others are launching enterprise AI marketplaces – internal app-store-like hubs where approved AI assets are shared, monitored, and reused. Meanwhile, composable "AI mesh" architectures are emerging as the connective tissue of this new ecosystem. Intelligence becomes modular, allowing teams to combine pre-approved components instead of reinventing them. This dramatically shortens time-to-value while maintaining oversight and trust.

Together, these developments mark a clear shift: from experimentation to production-grade AI adoption, where governance and lifecycle management are as important as creativity and speed.

Architecting the future of intelligence

The next competitive advantage won't come from who experiments fastest, but from who governs, reuses, and scales intelligence most effectively. Enterprises that once led the data revolution by mastering data products now face a similar opportunity with AI. The same principles apply – clear ownership, lifecycle discipline, and strong architecture – but the impact reaches even further. AI products thrive only on trusted data foundations. The organizations that combine high-quality data with structured AI lifecycles will shape the next era of enterprise value: one defined by ecosystems of modular, explainable, and interoperable intelligence. With nearly every business leader recognizing the competitive edge of scaling AI agents, the clock is ticking. The winners will balance speed with trust, modularity with governance, and experimentation with discipline – creating marketplaces of intelligence that fuel innovation, efficiency, and sustainable growth.

Start *innovating* now

Build your AI product marketplace

Turn scattered AI experiments into a governed, enterprise-wide marketplace of intelligence. Version, approve, and share models, prompts, and agents so teams can innovate faster – without reinventing the wheel.

Establish lifecycles, not one-offs

Treat every AI asset as a living product. Assign ownership, define metrics, and monitor continuously for fairness, explainability, and performance. A disciplined lifecycle keeps AI safe, useful, and aligned with strategy.

Design for modularity and scale

Engineer AI products as composable building blocks that plug seamlessly into existing workflows. Modularity lets you mix, match, and scale intelligence across the enterprise – turning innovation into momentum.

#AIProducts #GenerativeAI #EnterpriseAI #AIatScale #AIGovernance #AIMarketplace #AIInnovation #AITransformation #ComposableAI #AIFactory #ModularIntelligence #DataAsAProduct #ResponsibleAI #AIFuture #AIRevolution

Open standards as the foundation of future data and AI platforms

Sjoukje Zaal

CTO, Insights & Data Europe, Capgemini

Dael Williamson

EMEA CTO, Databricks

Data and AI platforms are at the heart of digital transformation. But too often, they are built on proprietary formats and closed models that create silos, raise costs, and limit flexibility. In a world where organizations need interoperability, sovereignty, and speed, this approach is no longer sustainable. Open standards offer a better path forward.

From lock-in to openness

Data and AI platforms used to be closed worlds. Proprietary formats, vendor-specific connectors, and black-box models accelerated adoption, but at a cost: lock-in, high switching barriers, and reduced sovereignty. Once data was stored in a proprietary format or models were trained in a closed environment, moving them became difficult and costly. Enterprises often found themselves redesigning workflows, rewriting code, or even discarding valuable assets when switching vendors.

That era is ending. Thanks to open standards and an expanding ecosystem of open-source AI models, platforms are shifting from controlled silos to flexible, collaborative ecosystems.

From **Delta Lake** and **Apache Iceberg** for data storage, to **open-source AI models** like **GPT-oss**, **Llama**, and **Mistral**, openness is redefining what's possible. These technologies allow organizations to move data and compound AI model systems (agents) freely, scale AI workloads across clouds, and maintain transparency for regulators and partners.

Why openness matters

For enterprises and governments, the stakes are clear:

- **Compliance** with evolving AI and data regulation
- **Sovereignty** in controlling where data is stored and how it is used
- **Interoperability** across hybrid and multi-cloud environments

Closed ecosystems make these challenges harder. They increase dependency on one vendor, complicate audits, and reduce flexibility in adopting new technologies. In contrast, open standards provide a shared foundation that works across providers and platforms. Structured and unstructured data stored in an open format like Parquet, Delta Lake, or Iceberg can be analyzed by multiple engines without conversion. AI models built on open architectures can be audited, retrained, and reused without hidden restrictions.

This interoperability preserves freedom of choice and protects long-term investments, ensuring today's platform decisions remain valid tomorrow.

The Databricks contribution

Databricks has been at the forefront of making openness real in enterprise environments. The company's origins are in open source, and its engineers have played a central role in creating and fostering community growth around some of the most widely adopted standards in data and AI.

- **Apache Spark** transformed large-scale data processing and became the de facto standard for distributed analytics.
- **Delta Lake**, developed by Databricks and now an open-source project managed by the Linux Foundation, introduced ACID transactions and reliability to data lakes.
- **MLflow** emerged as the most widely used open-source platform for managing the machine learning lifecycle, from experiment tracking to deployment.
- **Unity Catalog** serves as the unified governance layer on the Databricks Platform, and its core functionality has also recently been open-sourced.
- **Spark Declarative Pipelines** is Databricks' latest donation to Spark, allowing for easy and fast pipeline development.
- **Delta Sharing** is an open protocol for secure data sharing, making it simple to share data with other organizations regardless of which computing platforms they use.

More recently, Databricks has embraced **Apache Iceberg,** enabling interoperability across engines and vendors, and expanded support for **open-source AI models,** giving enterprises choice beyond proprietary stacks. Databricks is the only vendor that supports multiple formats in Delta Lake and Iceberg.

The **Databricks Data Intelligence Platform** combines these open components with enterprise-grade governance, security, and performance. This ensures organizations can adopt open standards with confidence, scaling them across industries and regions. In short, Databricks shows that openness does not come at the expense of reliability – it strengthens it.

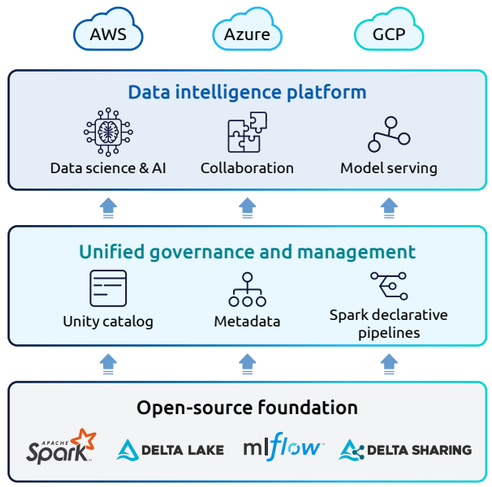

Real-world impact

Openness is already reshaping industries:

- **Financial services:** Banks are under strict scrutiny when it comes to data handling. By building AI models on open data formats, they can meet compliance requirements across jurisdictions while maintaining flexibility. A trading desk in London and a compliance team in Frankfurt can work on the same data without duplicating or reformatting it. This reduces operational risk while speeding up decision-making.
- **Healthcare:** Clinical data, imaging, and genomic research data are often stored in different systems, each with its own rules. Using open standards, healthcare organizations can combine these datasets securely, accelerating drug discovery and enabling more personalized care. Instead of spending months reconciling formats, researchers can collaborate on a trusted, common layer, an approach that became especially critical during the pandemic.
- **Public sector:** Governments are increasingly focused on digital sovereignty. By deploying open AI models on platforms built around open standards, they reduce reliance on single vendors and ensure long-term control of critical digital infrastructure. Whether it's analyzing mobility patterns, monitoring energy grids, or building citizen services, open standards allow governments to meet regulatory requirements while still innovating at pace.

In each case, Databricks provides the foundation for adopting these open standards at scale, adding the governance and reliability needed for mission-critical use.

Looking ahead

The next generation of data and AI platforms will not be defined by single vendors, but by ecosystems built on open standards. The pace of innovation in AI models, storage formats, and regulatory frameworks is simply too fast for closed approaches to keep up.

Enterprises that adopt open data formats and open AI models will have the agility to innovate, comply, and collaborate globally. They will be able to mix and match tools, clouds, and providers with minimal friction. This avoids the costly rewrites and migrations of the past.

With its deep roots in open source and its enterprise platform, Databricks stands at the center of this movement. It demonstrates that openness is not just a technical preference, but a strategic foundation for the future. By combining community-driven innovation with enterprise-grade capabilities, Databricks is enabling organizations to protect their investments, preserve their sovereignty, and accelerate their journey into the data- and AI-driven economy.

Start *innovating* now

Adopt open formats

such as Delta, Apache Iceberg, and Parquet to ensure that data remains portable, interoperable, and future-proof across platforms.

Leverage open AI models

in situations where transparency, sovereignty, and control over data and algorithms are essential.

Combine openness with strong governance

to maintain compliance, build trust, and enable scalable innovation across multiple cloud environments.

#OpenStandards #DataPlatforms #AIModels #DigitalSovereignty #Interoperability #FutureOfAI

Factory Eyes = (human + computer) . vision

How vision systems are transforming machines, automations, and robots into smarter partners

Dr Marc Blanchon

Computer Vision and Robotics
AI Lead, Capgemini

Factories used to rely on hands, then on machines. Now they rely on eyes. Computer vision turns automation into perception, giving machines the ability to see, learn, and adapt. It's not about replacing human judgment, but scaling it. The result: factories that think visually, act intelligently, and never stop improving.

In the beginning, factories placed their trust entirely in humans. Every quality check depended on the trained eye and experience of the worker. This brought intuition, but also variation – quality was never fully uniform or systematic.

Automation marked the next stage. Machines took over repetitive tasks, executing them endlessly with precision. But they were blind. They repeated instructions without comprehension, unable to understand their environment.

That limitation is now disappearing. Today, we are witnessing the rise of intelligent machines – built with the same components, but augmented with cognition and understanding. We are entering the era of perception, where systems don't just repeat; they see, adapt, and calibrate themselves to the world around them.

Quality control evolves too: once manual and fragmented, it is now guided by human expertise through automation. The result is uniformity in acquisition, impartiality in outcomes, and a process anchored in knowledge rather than chance.

Perceptive factories also bring something new: the ability to see and historize. Every check and variation leaves a visual trace – a record that provides proof, resilience, and above all, data. Rich, visual data – the same kind that fuels nearly 80% of human cognition.

> *Automation used to follow orders; now it interprets its environment.*

More perception, less disruption, better scaling

Scaling used to be a challenge. Traditional machines could be duplicated but not adapted. Each variation required a new setup – often a new machine.

That's no longer the case. By adding intelligence – with vision as a cornerstone – machines become both replicable and adaptable. What once required controlled environments can now thrive under constant change. Thanks to AI and computer vision, factories gain invariance: systems that perform even as conditions shift.

The result is versatility. A single machine can take on multiple roles, perceiving new features and learning new requirements without starting from scratch. Sometimes, with just a small dose of data science or targeted learning, adaptation is lightning fast. One machine, multiple values.

A machine that sees is a machine that rarely stops. Vision reduces downtime, accelerates scaling, and avoids the heavy costs of redesign. One camera system can inspect dozens of product variants instead of needing separate setups – a direct path to efficiency and resilience.

A machine that sees is not just an asset – it's a multiplier of value.

Why vision matters to industry

Vision is a strategic enabler for modern industry. It drives quality, uniformity, and flexibility. By observing and recording, vision systems make it possible to act on what is missing, failing, or ready to improve. They not only allow companies to look back – enabling predictive and prescriptive maintenance – but also to move forward, turning insights into better ways to treat defects and optimize processes.

Vision also makes industry more sustainable. Fewer machines are required, but they are smarter – which means less deployment, less waste, and less scrap. Smarter factories reduce the environmental burden while minimizing defects and rework. The outcome is resilience: fewer resources consumed, more value created.

For decades, industries were automation-centric, with humans acting behind machines but without real control. Now we are shifting back toward human-centric processes. Vision matters most when it supplements – not replaces – human expertise.

This requires respect for people and regulation. Vision must align with frameworks like GDPR and the EU AI Act, ensuring transparency, fairness, and trust. Human expertise remains at the center of decision-making. Vision enhances, simplifies, and scales – but it does not replace.

Vision is the enabler that scales human expertise, not the tool that replaces it.

Pathways to autonomy: When vision meets robotics

The journey does not end with machines that see. The next leap is machines that see and move – robots capable of navigating, adapting, and acting autonomously.

Smart machines already reduced the need for countless rigid systems. Robots now bring a new dimension: they interact, adapt, and evolve within human spaces.

Humanoid robots are not a gimmick; they are the convergence of mechanics, AI, perception, and design – all serving human performance. Factories are built for people and will remain so. Efficiency and adaptability demand robots that mirror human form: bipedal to take stairs, two-armed to handle tools, eyes aligned to human vision.

But, like vision, robotics must remain human-centric. Robots are not here to replace, but to enable. In this ecosystem, robots are the hands, vision is the eyes, and AI is the brain. Together, this is physical AI.

Computer vision allows robots not only to monitor processes but to perceive and understand environments – to avoid danger, adapt, and act with purpose. Once cameras are in place, vision unleashes action: screwing, lifting, moving, assembling.

The frontier lies in integration: digital twins and vision provide context, while AI brings reasoning. This is the pathway to true autonomy – machines that see, perceive, and understand.

Robots bring the hands, but vision and AI provide the perception and reasoning that make autonomy real.

Robots bring the hands, but vision and AI provide the perception and reasoning that make autonomy real.

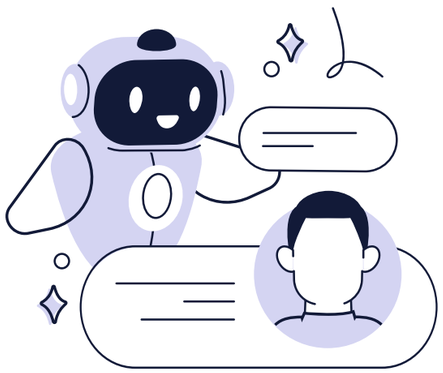

Collaboration in action

At the heart of the perceptive factory lies a new ecosystem built on three pillars: humans, machines with vision, and robots. Each plays a distinct role, but together they perform.

Vision is the catalyst of this collaboration. Just as it makes up 80 percent of human cognition, it acts as the bridge between worlds. Humans see, machines see, robots see – perception becomes the shared layer of decision-making.

From this shared perception come concrete collaborations: adaptive inspection, guided robotics, teleoperation. Safety improves, adaptability increases, and factories become not just efficient but resilient.

The goal is simple: to elevate production. By mastering rich visual data – images and video – factories transform perception into intelligence, and intelligence into performance.

Shared perception is the new language of collaboration between humans, machines, and robots.

Start *innovating* now

See smarter, not harder

Begin by embedding vision into your existing systems. Use cameras and AI models to generate actionable insights from production data – no need to rebuild everything. Quick wins prove value and inspire momentum.

Scale with purpose

Once results are visible, expand strategically. Interconnect vision systems across lines and sites, feeding data into a unified digital backbone. The goal is not more automation, but better awareness and faster decisions.

Keep humans at the core

Innovation succeeds when people trust it. Ensure every AI and vision project enhances operator expertise rather than replacing it. Responsible adoption – aligned with ethics and regulation – builds trust, efficiency, and long-term value.

#AI #ComputerVision #PhysicalAI #SmartFactories #IndustrialAI #Automation #Robotics #HumanCentricAI #PerceptiveIndustry

Strange folders in the cloud

Distributing PDFs to your LLM is no basis for an AI strategy

Joakim Nilsson

Knowledge Graph Lead, Insights & Data Sweden, Capgemini

Johan Müllern-Aspegren

Emerging Tech Lead, Applied Innovation Exchange Nordics and Core Member of AI Futures Lab, Capgemini

In the era of AI-driven decision making, organizations are tempted to shortcut their data strategy by feeding large language models (LLMs) with unstructured documents such as PDFs. While this approach may seem convenient, it often results in poor performance, hallucinated outputs, and a lack of meaningful insight. Throwing PDFs into an LLM is not a valid data strategy.

We propose the need for semantic enrichment and domain-specific context as essential components of a scalable strategy. A key enabler in this transformation is the use of knowledge graphs, which organize information into interconnected entities and relationships. By grounding LLMs in structured, contextualized knowledge, organizations can drastically reduce hallucinations and improve the relevance and accuracy of AI-generated responses. We explore how knowledge graphs bridge the gap between raw data and intelligent reasoning, offering a foundation for trustworthy, enterprise-grade AI systems.

> *Throwing documents at an LLM is like swapping the idol for a bag of sand – looks clever, but you're still triggering all the traps. Structure your knowledge if you want to escape with the prize.*

"Thoughts without content are empty, intuitions without concepts are blind." **– Kant**

"A little error in the beginning leads to a great one in the end." **– St. Thomas Aquinas**

In the era of AI-driven decision-making, organizations are tempted to shortcut their data strategy by feeding large language models (LLMs) with unstructured documents such as PDFs. While this approach may seem convenient, it often results in poor performance, hallucinated outputs, and a lack of meaningful insight. Throwing PDFs into an LLM is not a valid data strategy.

When we provide an LLM with information in the form of text (either the whole text attached in the prompt or matching chunks through a RAG approach), we do reduce the risk of hallucination. The model's responses will be somewhat closer to the source material. But let's be clear: the LLM is still not using the information – it is **interpreting** it through its world model (which is still just a language model), and generating output based on that, i.e., statistical patterns.

That means hallucination risk never truly disappears. Even when the system returns a reference to the chunk from the document that was used as input, the burden of validation remains with you. You must read the referenced passage, interpret it, and verify that the model made the same assessment as you. Ouch. Every answer is still a gamble. Not scalable, not possible to automate, and not agentic compliant.

The alternative is to **remove the knowledge domain from the LLM entirely.** Let the LLM do what it excels at: understanding natural language questions and translating them into structured queries. But do not let it invent the answers. Instead, answers should come from a deterministic, rule-based, yet semantically accessible data layer – a knowledge graph or equivalent structured representation.

We propose the need for semantic enrichment and domain-specific context as essential components of a scalable strategy. A key enabler in this transformation is the use of knowledge graphs, which organize information into interconnected entities and relationships. By grounding LLMs in structured, contextualized knowledge, organizations can drastically reduce hallucinations and improve the relevance and accuracy of AI-generated responses. We explore how knowledge graphs bridge the gap between raw data and intelligent reasoning, offering a foundation for trustworthy, enterprise-grade AI systems.

The reality check: Testing AI with structured knowledge

To make this concrete, we set out to test the hypothesis that using an LLM to generate a knowledge model (and manually validating that), and then using an LLM to query the model for data to answer a question, is superior to letting the LLM ingest and assess the text directly.

To take an example: If you feed the annual report of a major car manufacturer into a cloud LLM and ask about numbers or strategy, you'll often get decent answers. Today's large models can parse tables, charts, even images. But you'll still feel the need to double-check every response – which is exactly the point of this article. To make it clearer (and avoid wading through 14 megabytes of corporate boilerplate), we used a short excerpt with a diminutive local model, where hallucinations are easier to spot.

We tested a local LLM (Gemma 3:1B) with a short excerpt from MegaCorp's annual report . The prompt:

Show all subsidiaries with ESG contributions

The reply came back smoothly:

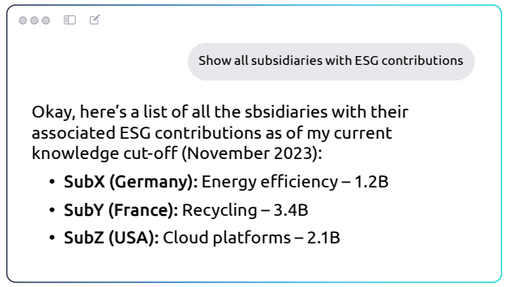

Convincing – but wrong. SubZ had no ESG contributions in the source. The model simply invented a plausible answer.

Next, we asked the LLM to build a graph from the same document: nodes, edges, nothing else. After validating that graph, we let the LLM query it. This time it issued a Cypher query and returned only the subsidiaries with actual ESG contributions. No hallucinated SubZ. And as a bonus, instead of getting merely a reference to a text chunk, you can easily ask the application to also return the query (Cypher/SPARQL) as well as the result.

The point is simple: Querying a document index means validating every answer. Querying a graph means validating the graph once.

This illustrates the key point: **When you query a document index, you validate every answer. When you query a graph, you validate the graph once.**

Conclusion

At its heart, this isn't just about documents versus graphs. It is about decoupling the LLM from the knowledge domain, and where you place the burden of validation.

- **Document-first RAG** forces validation at the very end of the pipeline – every answer must be checked, every hallucination caught, every gap patched. It scales poorly and exhausts human oversight and is not agentic compliant.
- **A separated data/knowledge model approach** moves validation upstream. Once information is structured, linked, and governed in a graph, the downstream systems – LLMs, agents, dashboards – can operate reliably on a deterministic knowledge base.

This shift enables modularity. In a world of autonomous agents and automated decision-making, you want each layer of your architecture to do one thing well and pass on clean, validated outputs. Ontology evolution, knowledge ingestion, reasoning, and generation become separate modules – each improvable, replaceable, and automatable.

The principle is simple: **validate early, scale later**. By moving the validation point as far upstream as possible, you gain the ability to automate reasoning chains, orchestrate agentic frameworks, and build enterprise AI systems that are not just powerful, but trustworthy.

Throwing PDFs at an LLM may give you an answer. Building a modular knowledge model gives you a system.

MegaCorp Annual Report (Excerpt)

Page 1 - CEO Letter

"At MegaCorp, our vision is to drive sustainable growth through digital innovation, operational excellence, and a commitment to carbon neutrality by 2030. Across our European subusidiaries, we have seen strong contributions both in technology adoption and ESG impact."

Subsidiary	Region	Focus Area	Revenue (USD)	ESG Contribution
SubX	Germany	Digital Innovation	1.2B	Energy Efficiency
SubY	France	Manufacturing	3.4B	Recycling
SubZ	USA	Cloud Platforms	2.1B	

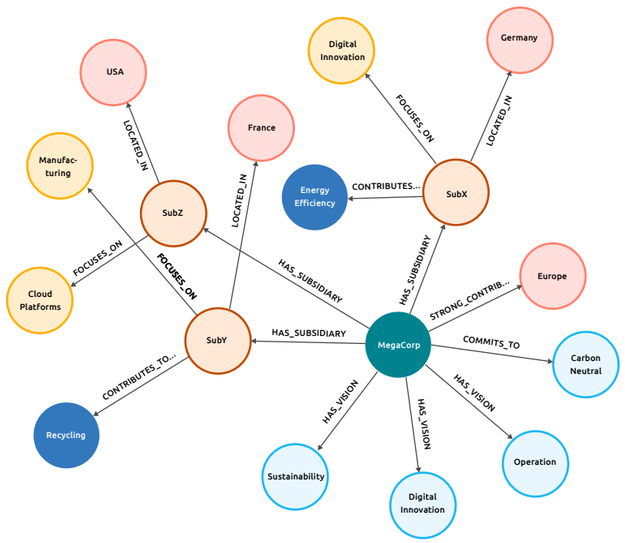


```
neo4j$ MATCH (s:Subsidiary)-[:CONTRIBUTES_TO_ESG] (e:ESG) RETURN s.name, s.revenue, e.name
```

Table

s.name	s.revenue	e.name
"SubX"	1.2	"Energy Efficiency"
"SubY"	3.4	"Recycling"

Text

Code

Start *innovating* now

Stop dumping PDFs into your LLM

Treat unstructured documents as raw input, not a ready-made knowledge layer.

Validate upstream

Structure and enrich your data – supported by LLMs – into a knowledge graph; so hallucinations are caught once, not every time you query.

Let LLMs translate, not invent

Use them for natural language understanding and query generation – but keep answers grounded in deterministic, governed data.

#TrustworthyAI #KnowledgeGraph #LLM #Agentic #EnterpriseArchitecture

Secure AI – Avoid getting famous at DEF CON

Embedding security into the golden path of data and AI

Arne Rossmann

Innovation Lead, Insights & Data, Capgemini

At this year's ***DEF CON***, the world's largest security conference, one theme dominated the conversation: **AI systems are already under attack.** From agentic AI agents exposing sensitive data to misconfigured **model context protocol (MCP)** *servers* granting unauthorized access, to LLM-powered apps tricked by prompt injections – DEF CON made it brutally clear that vulnerabilities in AI aren't theoretical. They're here, and they're being actively exploited.

Does that mean you should halt your agentic AI innovation? **Absolutely not.** But it is a powerful call to action on the awareness of the multiple **threats** in agentic AI ecosystems. To truly unlock the power of data-powered innovation and agentic AI, organizations must weave **DevSecOps principles into the golden path** – turning security from a burdensome afterthought into a seamless, automated foundation for safe, rapid innovation.

Why security-first AI matters

AI brings unprecedented opportunities, but it also **amplifies familiar risks** while introducing genuinely new ones. The convergence of data pipelines, complex model ecosystems, and autonomous agents has created an expanded, more fragile attack surface:

- **Supply chain vulnerabilities:** Just as the **NPM ecosystem** has shown, a single compromised library can ripple through thousands of applications. AI projects depend on frameworks like **LangChain, HuggingFace, PyTorch, TensorFlow** – all containing potential weak links.
- **MCP and integration risks:** MCP servers and agent frameworks connect AI models to APIs and tools. If a connector lacks proper authorization, an agent might gain access to sensitive data or perform unintended actions.
- **AI-specific exploits:** From prompt injection to adversarial triggers, **data poisoning**, and **cross-agent privilege escalation**, attackers now have new vectors to manipulate AI behavior in ways traditional testing never accounted for.

The attack surface is larger than ever. That's why the answer isn't more manual reviews or slowing down delivery. It's **embedding security directly into the developer workflow.**

Two tools are essential for mastering this new security landscape: the **software bill of materials (SBOM)** and robust **role-based access control (RBAC).**

SBOMs: Transparency and trust for AI projects

An SBOM is a complete inventory of every library, framework, and dependency that makes up a system. Just as a BOM is vital in manufacturing, the SBOM is becoming the key resource for the data and AI world. For AI, SBOMs provide critical visibility across:

- **Data products** (ingestion pipelines, storage, transformations)
- **Model training and inference** (PyTorch, TensorFlow, scikit-learn)
- **Agentic AI systems** (like those built with LangChain and MCP connectors).

With SBOMs, organizations gain a clear map of dependencies. When a new vulnerability (like the next NPM security crisis) is disclosed, you can instantly check exposure across all your AI projects. Teams without an automated, up-to-date SBOM are simply left guessing.

Here's a simplified **CycloneDX SBOM** for an LLM-powered agent project:

```
{
  "bomFormat": "CycloneDX",
  "specVersion": "1.4",
  "version": 1,
  "components": [
    { "type": "library", "name": "langchain", "version": "0.1.12" },
    { "type": "library", "name": "openai", "version": "1.10.0" },
    { "type": "library", "name": "fastapi", "version": "0.103.2" },
    { "type": "library", "name": "uvicorn", "version": "0.24.0" },
    { "type": "library", "name": "pydantic", "version": "2.5.2" }
  ]
}
```

This manifest provides transparency. If a vulnerability emerges in **FastAPI 0.103.2,** security teams know immediately which AI agents are at risk.

Role-based access: Controlling what agents can do

While SBOMs protect against vulnerable dependencies, **role-based access control (RBAC)** protects against **misuse of capabilities** when AI agents connect to MCP servers and tools.

A lack of RBAC is a recipe for disaster. If an agent has blanket permissions to query production databases or call sensitive APIs, an attacker only needs a successful prompt injection to trigger a massive data leak or system compromise.

By enforcing RBAC:

- Each AI agent only has the **minimum privileges required.**
- Access to tools and data sources is **segmented by role and scope.**
- Unauthorized tool usage is blocked by design.

For example, a **customer support agent** should only be allowed to ***read*** from a knowledge base and ***create*** support tickets. It should be blocked from deleting data, issuing refunds, or accessing HR systems – no matter how cleverly it's prompted.

The golden path: Security without friction

The golden path in platform engineering is about opinionated, paved roads that make it easy for developers to move fast while adhering to best practices. For security, this means automating guardrails, so compliance is the easiest choice.

The security journey starts the moment a new project is created, using a developer portal like **Backstage.** Instead of manually configuring security tools, developers choose an approved **golden path template** for their "LangChain agentic service."

This template isn't just a starting structure; it's a security guarantor. It automatically scaffolds the project with:

- **Secure framework defaults:** A LangChain boilerplate that enforces secure settings for tool execution and prompt handling.
- **SBOM automation:** A requirements.txt file pre-populated with known-secure dependencies, along with a definition in the project's folder that triggers an **SBOM generation GitHub Action** upon creation.
- **RBAC policy stubs:** Default files defining the minimum required permissions (agent_rbac.yaml) for the **LangChain agent** to connect to an MCP server, ensuring the agent adheres to the **principle of least privilege** from day one. Crucially, the template includes a **policy/rbac.rego** file containing the organizational RBAC rules.

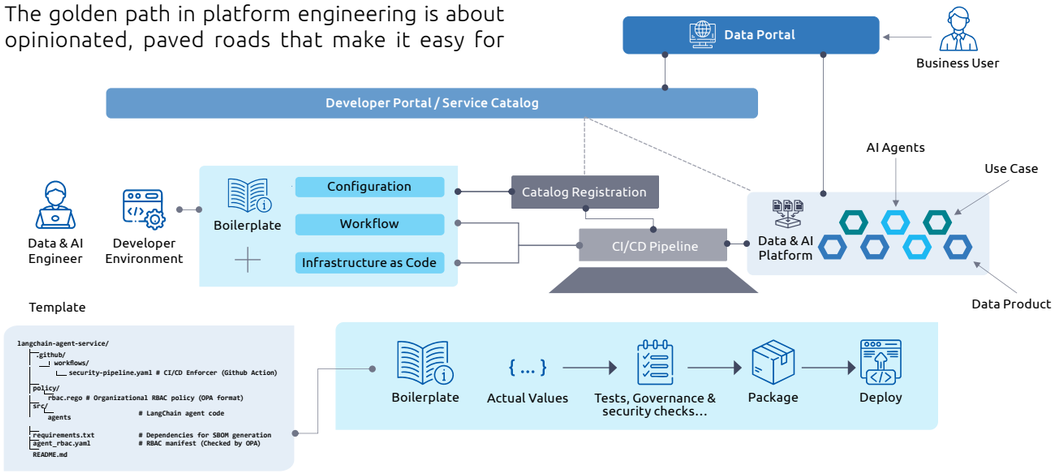

Why this matters for data & AI innovation

For organizations embracing **data mesh, data products,** and **agentic AI,** this dual approach of **SBOM and RBAC** provides:

- **Transparency and trust** in dependencies
- **Control and governance** over agent actions
- **Confidence to innovate** without risking misuse or breaches.

By embedding DevSecOps into the golden path, enterprises not only avoid vulnerabilities but also prevent agents from "going rogue."

Conclusion: Avoid becoming the next DEF CON demo

DEF CON made one thing terrifyingly clear: AI security is already a battleground. Organizations that fail to act now risk becoming the next highly visible case study in how not to adopt AI.

The solution is not to slow down innovation but to secure the golden path. By integrating automated SBOM creation, enforced RBAC, and CI/CD checks into the developer workflow, enterprises not only avoid classic vulnerabilities but also prevent their agents from "going rogue".

Start *innovating* now

Automate SBOMs in every pipeline

Generate and scan SBOMs in your CI/CD workflow. Make vulnerability management a default, not an option.

Adopt security-first templates with RBAC

Use Backstage templates to scaffold AI projects with RBAC policies for MCP tools enforced by default.

Shift security left for AI

Validate dependencies, RBAC policies, and agent behaviors early. Don't wait until deployment to discover vulnerabilities.

#ShiftLeft #SecurityFirstAI #PlatformEngineering #AccelerateInnovation

Gen AI Strategic Intelligence System

The fast track to business results, even before digital transformation

Dnyanesh Joshi

Director, Chief Architect – Large Deals, Middle East, Capgemini

Enterprises are pouring investment into generative AI (Gen AI), but many are missing a critical opportunity. Simply applying Gen AI to traditional business intelligence (BI) systems is like putting a rocket engine on a horse-drawn carriage. The real challenge for most organizations isn't a lack of data; it's the inability to use their existing data to deliver actionable insights and improve key business performance indicators (KPIs).

Presenting an innovative solution that tackles this problem head-on. This **Gen AI Strategic Intelligence System** is a game-changer because it works with a company's data ***as it exists today*** to deliver tangible results, providing significant competitive advantages without waiting for a multi-year digital transformation to complete.

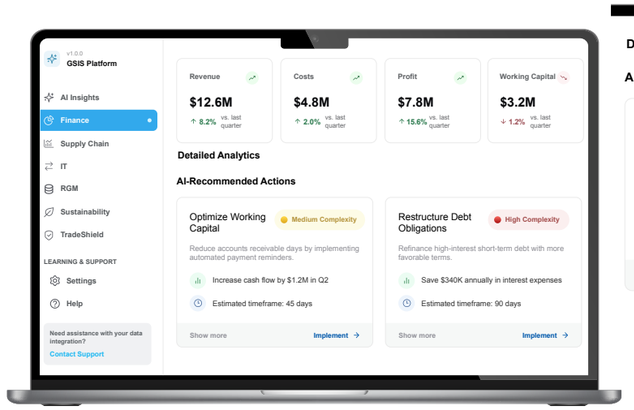

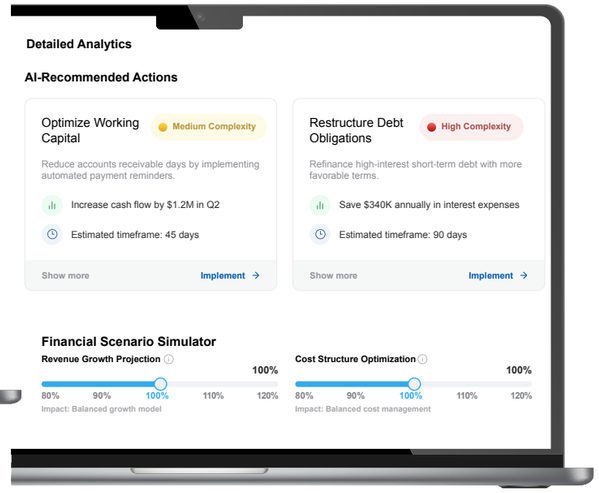

The disconnect between data and insight

Despite widespread investment, most organizations struggle to get true value from their data. Executives understand Gen AI's potential: a Capgemini Research Institute study found that 80% of organizations increased their Gen AI investment in 2024 (Source: Capgemini Research Institute, Generative AI executive survey, May–June 2025, 1,020 organizations that are at least exploring Gen AI capabilities). Yet, legacy BI systems create a major disconnect:

- **Static dashboards:** They often produce static dashboards that fail to provide strategic foresight or transformative insights. They give you a snapshot of what happened, not a roadmap for what to do next.
- **Lack of action:** The "insights" are often nothing more than talking points for the next meeting. They aren't concrete enough for a decision-maker to act on with confidence or speed.
- **One-size-fits-nobody:** These systems rarely personalize output. The same dashboard is shown to a CFO, a marketing manager, and an operations lead, with no context for their specific roles or goals. The result is a generic, unhelpful view.
- **Data silos:** Data is aggregated within silos, meaning a human must still interpret and combine information from different departments to get a complete picture.

In short, many legacy systems miss the big picture, miss actionable meaning, miss the persona – and ultimately, miss the point.

Introducing your agentic AI advisor

These shortcomings are precisely what drove the development of the Gen AI Strategic Intelligence System. It's not just a tool; it's a new kind of intelligent advisor that works with your existing data landscape.

This agentic AI solution leverages cutting-edge technology to analyze and make sense of vast amounts of information. It uses virtualization and governance to capture your organization's data in a "data fabric," eliminating the need for a costly, time-consuming effort to build a unified data platform.

Within this fabric, specialized AI agents get to work. They analyze the data, collaborate with other agents across departments, and generate specific, hyper-personalized recommendations to boost KPIs in a given domain. Imagine having hundreds of domain experts tackling your most pressing challenges, all at once.

One of the solution's key strengths is its ability to provide a complete view of the enterprise – something that is virtually impossible for human experts to achieve on their own. The system also employs **explainable AI (XAI)** and **chain-of-thought reasoning,** so human users can easily understand how the AI arrived at its conclusions. These build trust and confidence, enabling leaders to act on the recommendations immediately.

Real-world impact across business domains

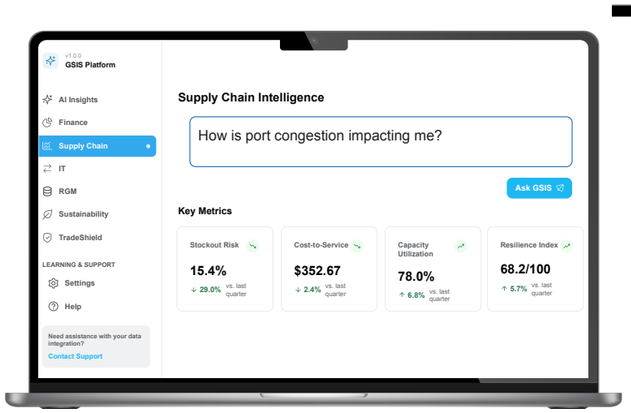

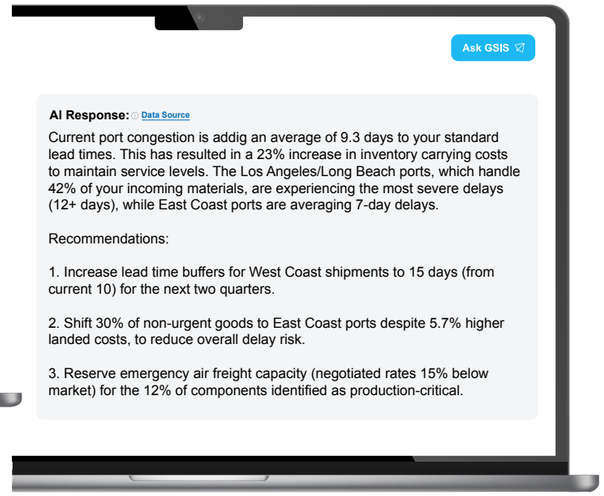

The Gen AI Strategic Intelligence System offers hyper-personalization at scale across various functions.

For finance

This system fundamentally changes the DNA of performance reporting. It delivers a unified view of the truth by embracing both ERP and non-ERP data sources without any data re-engineering. It enables dynamic, persona-based KPI reporting and provides independent performance tracking and recommendations based on natural language queries from a CFO, chief risk officer, or financial controller. Key benefits include:

- Improved cash flow and working capital
- Increased operational efficiency and asset utilization
- Reduced material and operational costs (as a percentage of revenue)
- Decreased manual effort in financial processes by 20–30%.

For supply chain

The solution provides insights into inventory, logistics, procurement, and more. This can lead to more accurate inventory tracking, improved order management, and optimized transportation.

For revenue growth

It delivers insights on customer interactions, intelligent profiling, sentiment analysis, and marketing campaigns. The expected benefits are clear: sales growth, higher profit margins, and a higher conversion rate. The system can even help formulate highly targeted revenue growth campaigns.

For IT operations

By providing insights into automation and digital employees, the system can significantly increase productivity within the software development life cycle (SDLC) and broader IT operations.

Start *innovating* now

Pursue a best-in-class data strategy

To unlock the full, long-term value of Gen AI, organizations must still pursue a best-in-class data strategy. This includes building a secure and scalable platform, adopting an enterprise-wide governance framework, and implementing appropriate guardrails.

Empower your organization with GSIS

The Gen AI Strategic Intelligence System empowers your organization today. It delivers goal-oriented, actionable insights that produce business value almost immediately, generating new revenue that can then be reinvested in those larger, longer-term Gen AI initiatives.

Don’t wait for your full digital transformation

Embrace an analytics solution that works with your existing landscape, all while you plan for the future. By focusing agentic AI on core KPIs, even a small improvement can provide enormous benefits – a perfect starting point for every domain within an organization.

#Datapowered #AgenticAI #ReasoningAgents #GenAI #AutonomousAgents #AgenticAdvisors

The age of anticipation

From thinking to doing – innovation begins where planning ends

Per Poulsen

Senior Manager, Innovation Lead, Capgemini Invent

Innovation doesn't start with a meeting; it starts with motion. At our AIE innovation center for Bayer, "anticipation" means acting before certainty arrives – cutting timelines, building prototypes, and learning in real time. This is how ideas evolve from slides to solutions, and how anticipation transforms from a mindset into a method.

For years, "anticipation" meant predicting the future – drawing charts, writing roadmaps, and hosting strategy workshops that looked far ahead but rarely moved the needle. At the AIE (Applied Innovation Exchange) for Bayer, we discovered that true anticipation is not about forecasting what might come next, but about building it through action. Innovation doesn't happen in PowerPoints. It happens when people try something. When they test, tinker, and sometimes fail fast. Only through doing do ideas gain momentum – and only through momentum does innovation become real.

From plans to prototypes

From the outside, we seemed to be doing everything right. We had the governance, the funding, the structure, and a beautiful innovation space. Teams were aligned, and clients were engaged. Yet despite the impressive setup, the breakthroughs we expected didn't materialize. Our client kept asking for "more innovation," even as our team continued presenting the newest ideas. Projects were discussed, celebrated, and then quietly faded away. The problem wasn't creativity – it was inertia. Too much time spent thinking about what to do, and too little time actually doing it. So we made a deliberate shift. We stopped organizing around technology for its own sake and started organizing around value creation. We cut project timelines from months to weeks. We moved from decks to demos, from talking about possibilities to building proofs of value. Every initiative had to show tangible progress – not in six months, but in six weeks.

The shift to doing

We also transformed our physical space. Instead of a polished showroom, we built a living lab: visitor tracking at the entrance, a 3D printer humming, a digital twin in motion, a podcast studio recording in the corner, and yes, an air hockey table in the middle of it all. None of it was decoration – it was a message. Ideas here don't stay hypothetical. They take shape. Action created its own energy. Momentum followed. Within a year, the innovation funnel grew tenfold, generating a potential 50:1 return on investment for our client. The mindset shifted too. As one leader summed it up:

"We moved from four ideas in 2023 to fifty in 2024."

What worked wasn't complicated. We stopped planning for perfection and started moving with intent. **"Stop planning, start moving"** didn't mean chaos; it meant acting with clarity and learning in motion.

What makes it work

Three elements proved essential in turning anticipation into execution.

1. **Air cover for action.** Innovation needs protection. Trust from leadership gave us the freedom to test ideas without fear of failure. Not everything worked – but every iteration taught us something useful.
2. **Visibility.** We made the work visible. Prototypes were built, filmed, demoed, and shared. A space buzzing with real experiments inspires belief far more than any slide ever could. People trust what they can see.
3. **Community.** Momentum spread as more people joined in. Every prototype sparked new ideas, collaborations, and questions. The culture began to reward experimentation, not just outcomes. Innovation became contagious.

Through these principles, anticipation stopped being passive and became active – not about waiting for the future to arrive, but building it one small proof at a time.

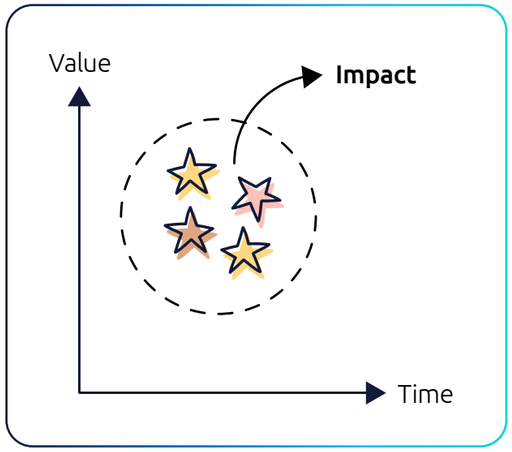

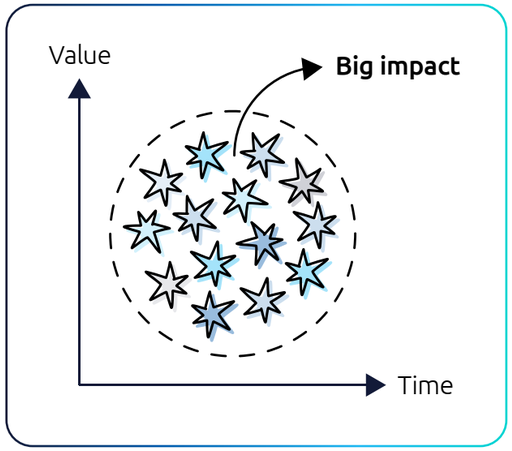

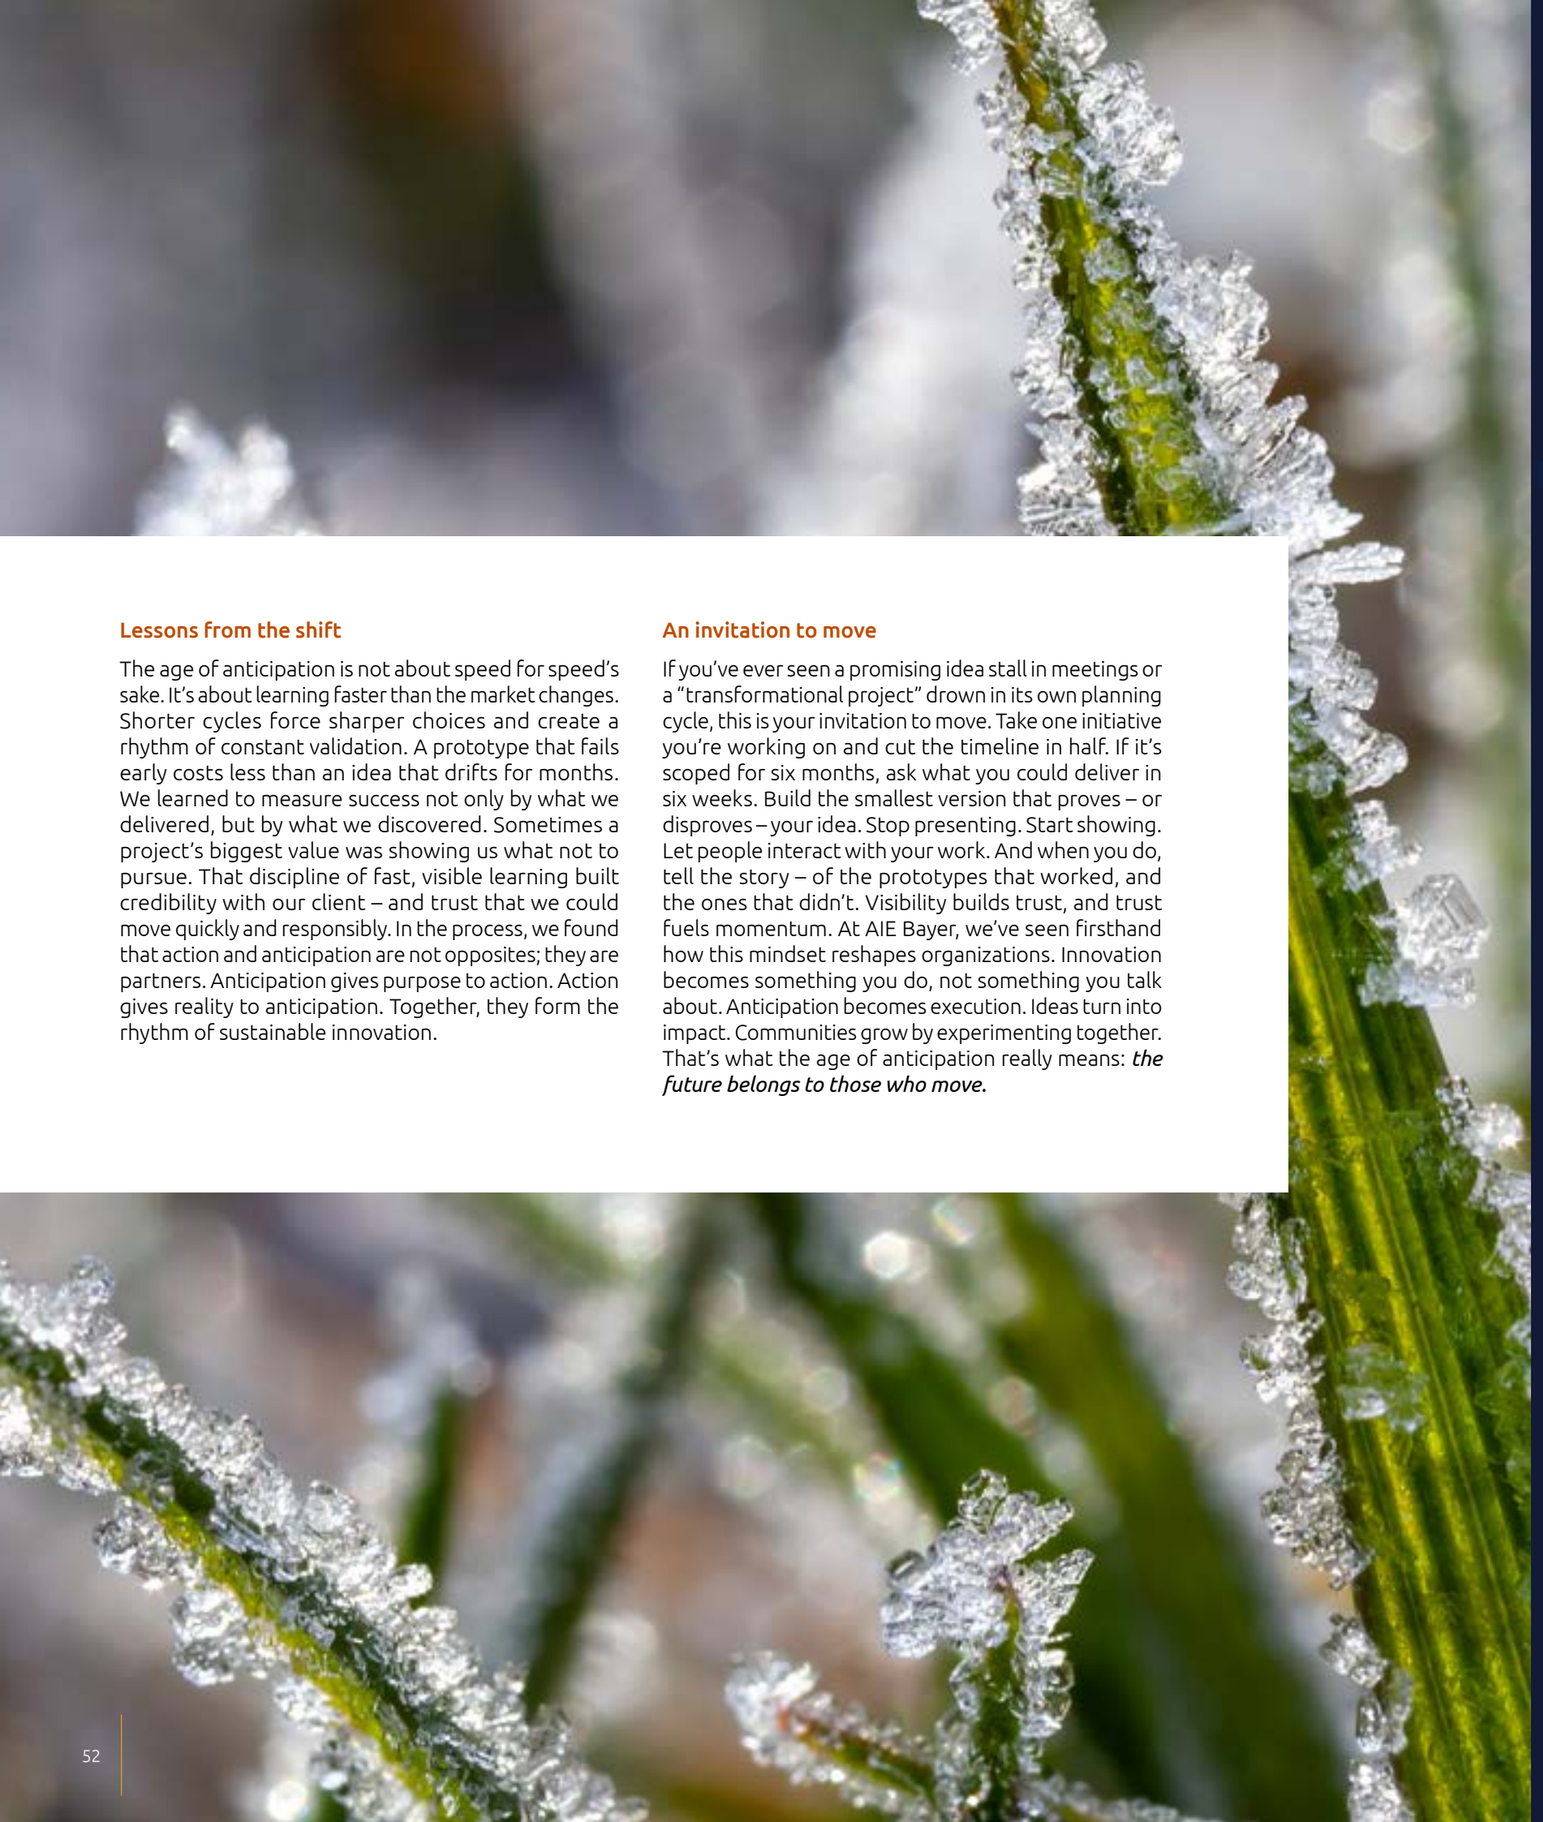

Lessons from the shift

The age of anticipation is not about speed for speed's sake. It's about learning faster than the market changes. Shorter cycles force sharper choices and create a rhythm of constant validation. A prototype that fails early costs less than an idea that drifts for months. We learned to measure success not only by what we delivered, but by what we discovered. Sometimes a project's biggest value was showing us what not to pursue. That discipline of fast, visible learning built credibility with our client – and trust that we could move quickly and responsibly. In the process, we found that action and anticipation are not opposites; they are partners. Anticipation gives purpose to action. Action gives reality to anticipation. Together, they form the rhythm of sustainable innovation.

An invitation to move

If you've ever seen a promising idea stall in meetings or a "transformational project" drown in its own planning cycle, this is your invitation to move. Take one initiative you're working on and cut the timeline in half. If it's scoped for six months, ask what you could deliver in six weeks. Build the smallest version that proves – or disproves – your idea. Stop presenting. Start showing. Let people interact with your work. And when you do, tell the story – of the prototypes that worked, and the ones that didn't. Visibility builds trust, and trust fuels momentum. At AIE Bayer, we've seen firsthand how this mindset reshapes organizations. Innovation becomes something you do, not something you talk about. Anticipation becomes execution. Ideas turn into impact. Communities grow by experimenting together. That's what the age of anticipation really means: ***the future belongs to those who move.***

Start *innovating* now

Cut the frame, build the proof

Shrink project timelines. If an initiative is planned for six months, ask: what can we deliver in six weeks? Build the simplest version that proves (or disproves) your idea. Momentum comes from seeing it in action.

Make action visible

Bring ideas into the real world. Prototype, demo, film, or stage them. Visibility creates belief, inspires collaboration, and sparks momentum no meeting can replicate.

Embrace experimentation with backing

Ask for trust and air cover to experiment. Accept that some things will fail, and treat those as data points, not disasters. With support, iteration accelerates and innovation compounds.

#FromThinkingToDoing
#StopPlanningStartMoving
#BusinessTransformation #ClientValue

vCon (virtualized conversation): Standardizing human and machine conversations

Portable, governed, and AI-ready conversation data

Rajesh Iyer

Vice President, Global Head of AI for FS, Capgemini

Kumar Chinnakali

Portfolio Manager, Hands-on Architect, Capgemini

Conversations are the lifeblood of business, yet most remain locked away in silos or proprietary tools. vCon (virtualized conversation), an emerging *open standard from the Internet Engineering Task Force (IETF)*, changes that. Think of it as a "PDF for conversations," a portable, structured, and auditable format that captures voice, video, and text into one machine-readable container. It turns every dialogue – human or machine – into a reusable, governed enterprise asset.

Ref - Contact Center Use Cases and Requirements for vCON

Today's enterprises span voice, chat, video, and messaging. Each system records interactions differently, leaving architects to stitch fragments across channels. vCon provides the missing blueprint a unified schema built on JSON that organizes every conversation into five layers: **participants, dialog, analysis, attachments,** and **integrity.**

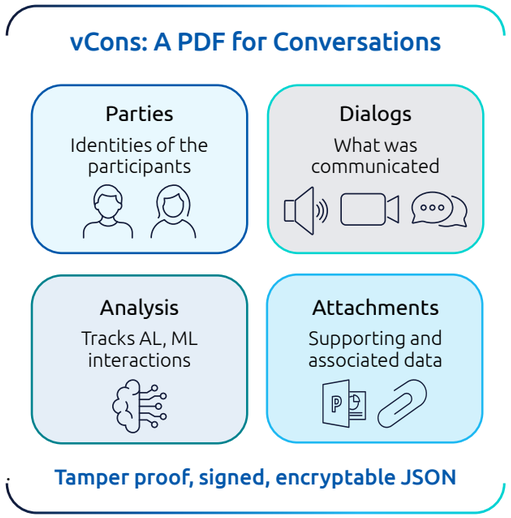

This structure ensures conversations can move securely between vendors, clouds, and systems without losing meaning, metadata, or compliance context.

Why it matters

Because unstructured conversational data is the next frontier of digital transformation. vCon gives AI systems structured memory identities, timestamps, and semantic markers, ready for large-scale learning and model evaluation. For **contact centers,** it enables omnichannel visibility, QA consistency, and real-time coaching. For **risk and compliance,** it provides tamper-evident, audit-ready trails that meet GDPR and CCPA mandates. For **AI and data teams,** it becomes the foundation for training, summarization, and retrieval pipelines.

Architecturally, vCon invites a new design pattern: emit, store, and evolve. Systems should **emit a vCon** at every conversation close, store it in a governed object store or **Iceberg table**, and **activate** it across analytics, model evaluation, and RAG pipelines. **Security** and governance must follow suit: encrypt on write, redact on read, and version with precision.

This standard redefines how enterprises build conversational infrastructure shifting from recording for reference to recording for reuse.

Examples: Industry momentum

Standards in motion: The **IETF vCon Working Group** is actively shaping the specification with contributions from global experts in telecom, cloud, and AI. The draft is now evolving toward interoperability testing, version control, and schema maturity (mirroring the early lifecycle of standards like SIP ***(Session Initiation Protocol)*** and STIR/SHAKEN ***(Secure Telephone Identity Revisited / Signature-based Handling of Asserted information using toKENs).***

Regulated industries: Banks and insurers are assessing vCon formats for "conversations of record" to meet FINRA, MiFID II, and HIPAA requirements. These industries view vCon as a natural successor to legacy call-recording systems, enabling full-lifecycle data governance.

Open-source and developer ecosystem: The **vCon Python SDK** and community implementations on GitHub are accelerating experimentation. Developers can now create, merge, redact, and validate vCon files locally or via APIs, making enterprise integration practical, not theoretical.

Emerging paradigms and future directions

The rise of vCon signals a shift from *channel-based communication* to *data-based conversation architecture.* Future contact centers and collaboration systems will emit vCons as digital artifacts that are structured, governed, and context-rich.

This creates **a shared conversation fabric** connecting humans, bots, and AI systems seamlessly.

Over time, we'll see **vCon-native ecosystems** where CRM, compliance, and analytics tools consume and produce vCons natively. With trusted schemas, embedded encryption, and semantic indexing, every interaction becomes queryable, interoperable, and value-generating. Conversations won't vanish after they end – they'll evolve into enterprise memory.

Why it matters

vCon represents more than a format; it's a foundation for **governed conversational intelligence.** It bridges silos, preserves trust, and transforms fragmented dialogue into structured data that powers compliance, empathy, and innovation. For architects, it's the next layer of digital infrastructure where every conversation becomes a governed, portable asset.

In the bigger picture, vCon is not just a new data standard: it's a new social contract between humans and machines. It ensures that every exchange, whether whispered between colleagues or spoken to an algorithm, carries structure, context, and accountability. In a world where AI learns by listening, governed conversations become the raw material of trust. With vCon, enterprises don't just record words; they preserve intent, emotion, and understanding in a form that both humans and machines can build upon. The result is not more noise, but shared intelligence – a dialogue that doesn't end when the call does.

Start *innovating* now

Map where conversations live

Audit every point where dialogue happens: contact centers, CRMs, chat tools, and collaboration apps. Identify ownership, data flow, and how conversation artifacts are currently stored or lost.

Design for emission

Define the trigger: when a conversation ends, a vCon should be generated. Architect APIs or webhooks to emit vCons that are automatically signed, encrypted, and stored with metadata for easy retrieval.

Build the governance loop

Implement the policy layer encryption on write, redaction on read, and schema versioning. Embed access controls, retention logic, and auditability so vCons remain trusted, compliant enterprise assets.

#vCon #ConversationalAI #ContactCenterAI #AIwithPurpose #EmpathyByDesign

From code to climate

How Gen Garage is shaping the future with AI

Aishwarya Kulkrni

Program Manager, Gen Garage – Strategic Talent transformation program, I&D India, Capgemini

Inside Capgemini's Gen Garage, ideas don't just spark – they scale. It's where emerging technologies meet purpose, and bold prototypes become real-world solutions. In this edition, we take a closer look at three shining examples: Vulcanus AI reimagining software creation, Resume Pilot reinventing careers, and Energi.AI powering a greener, smarter future.

Capgemini's Gen Garage is where bold ideas meet brilliant minds. It's a dynamic innovation space powered by emerging technologies designed to spark transformative solutions that are not only smart, but deeply purposeful. At its heart, Gen Garage is about collaboration. It brings together diverse teams with rising talent and seasoned experts to co-create intelligent, data-powered tools that improve how businesses operate and make decisions. But it's not just about performance; it's about progress. Every solution is built with inclusivity, sustainability, and ethical innovation in mind. Here are three recent examples:

Vulcanus AI: Powering the future with agentic intelligence

Vulcanus AI is redefining how we build software and make sense of data. At the forefront of Agentic AI, this powerful platform enables autonomous systems to be faster and smarter. Whether it's streamlining workflows, optimizing code across languages and cloud environments, or pulling insights from complex, multi-source data, Vulcanus AI is built to elevate how developers, analysts, and architects work.

From answering technical questions and parsing documents to generating wireframes from plain language, Vulcanus AI turns everyday tasks into intelligent, intuitive experiences. Its modular design fits effortlessly into existing systems, adapting and learning continuously to boost productivity and sharpen decision-making.

But Vulcanus AI isn't just another tool – it's a trusted co-pilot. One that helps teams build scalable, resilient solutions with confidence, ready for whatever the future brings.

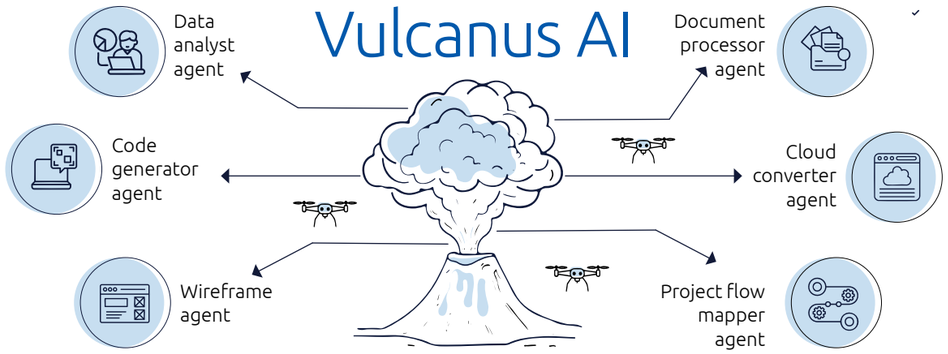

From static to strategic: Gen AI is rewriting the resume rulebook

In today's competitive job market, standing out starts with how you tell your story – and that's where generative AI steps in. The Resume Pilot isn't just a smart editing tool; it's a career companion that helps professionals craft resumes that truly resonate.

By analyzing job descriptions and tailoring content with precision, it fine-tunes tone, corrects errors, and brings key skills to the forefront. The result? Resumes that are not only polished but strategically aligned with what recruiters and applicant tracking systems are looking for.

Whether you're applying for your first role or pivoting to a new industry, Resume Enhancer personalizes your profile at scale, making sure your strengths shine through. It's more than a tool; it's your co-pilot in building a resume that doesn't just inform – it leaves an impression.

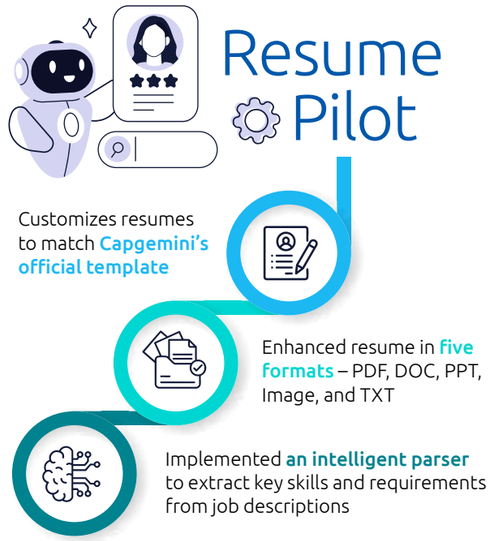

Energi.AI: Smarter energy for a greener tomorrow

The fusion of generative and agentic AI is redefining energy management and Energi.AI is at the forefront. This intelligent platform autonomously monitors, analyzes, and optimizes energy usage in real time, helping organizations cut carbon emissions while boosting performance. With the ability to simulate energy-saving scenarios, adapt to diverse environments, and deliver predictive insights, Energi.AI transforms buildings, factories, and data centers into smart, sustainable ecosystems. Its modular, self-learning architecture ensures continuous optimization and scalability, making Energi.AI a powerful ally in the global push for operational efficiency and climate-conscious innovation.

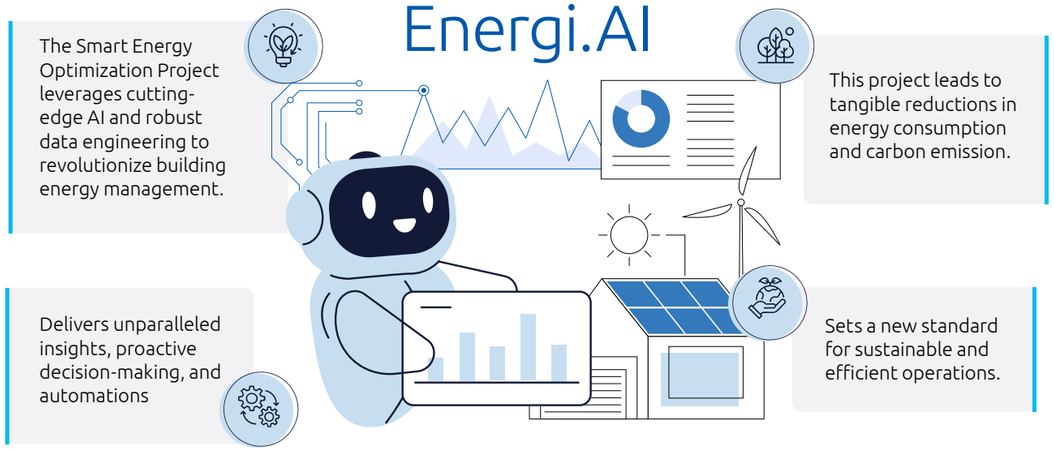

Market trends / key opportunities and developments

Gen Garage aligns its innovation strategy with evolving market needs, driving impact across software development, HR tech, and sustainability.

In IT, the surge in intelligent automation and multi-cloud adoption presents a major opportunity for Vulcanus AI. As agentic and generative AI reshape developer workflows, Vulcanus AI offers a unified platform that autonomously generates, optimizes, and translates codes, and supports low-code environments, intelligent documentation, and real-time data visualization.

In HR, Resume Pilot taps into the growing demand for personalized, AI-driven recruitment. With Gen AI at its core, it tailors resumes to job descriptions and segments skills, and adapts formats that help candidates stay competitive, while enabling HR teams to streamline evaluation in an increasingly remote and gig-based workforce.

In sustainability, Energi.AI responds to rising ESG commitments and regulatory pressures by delivering autonomous energy optimization. With real-time insights and predictive analytics, it supports smart buildings and green infrastructure, which is a key enabler of carbon reduction and operational efficiency.

To summarize, Gen Garage is empowering the next wave of AI innovators and redefining innovation by equipping emerging professionals to lead high-impact projects across IT, HR, and sustainability. By fostering agentic thinking and aligning with global trends, Gen Garage is shaping a future where technology drives smarter, greener, and more inclusive outcomes.

Start *innovating* now

Build AI that matters

Start by identifying real-world challenges in areas such as sustainability and workflow automation, and creating AI-powered solutions that drive meaningful change.

Embrace agentic intelligence

Move beyond static models and experiment with autonomous agents that plan, reason, and take action. Vulcanus AI is already showing what adaptive systems can achieve.

Empower diverse innovators

Bring together cross-functional teams to co-create inclusive, scalable solutions that shape the future of work and sustainability.

#AIForGood #AIPowered #Innovations

Green agents of change

Empowering youth through agentic AI

Guillaume Gérard

Head of AI France, Capgemini

Nadi Albino

Deputy Director of Partnerships, UNICEF (Generation Unlimited)

Mike Miller

Senior Principal Product Leader, AWS Agentic AI

Michael Hoff

Head of Alliance and Partnership, Mistral AI

What if AI could help young people shape a greener future? Capgemini's 8th Global Data Science Challenge (GDSC) – Green Agents of Change – is turning that question into action. At a time when artificial intelligence is transforming every aspect of life, GDSC 2025 focuses on one of its most meaningful applications: empowering the next generation to lead the green transition.

As part of Capgemini's global partnership with UNICEF*, this year's challenge invites participants to design AI-powered assistants that guide young people in Brazil toward sustainable career and learning opportunities. These assistants will engage in realistic, natural conversations with simulated personas, each representing a young individual with unique backgrounds, skills, and aspirations.

Through empathy-driven interaction and data-driven intelligence, participants will explore how agentic AI can support youth empowerment, climate action, and equitable access to green opportunities.

Technology with purpose: AI for youth empowerment

The challenge is a building block of Capgemini's global partnership with UNICEF and its support to the Green Rising initiative, a global effort to mobilize more than 20 million youth by 2026 with the skills and opportunities to lead climate action and shape a sustainable future. Brazil, with its vibrant youth population and vast environmental potential, serves as a compelling backdrop for this innovation.

"Young people have the ideas, the drive, and the urgency. What they need are the tools," says **Nadi Albino, Deputy Director of Partnerships at UNICEF.** ***"Through this challenge, we're co-creating with youth, giving them agentic AI as a lever that empowers them to act, innovate, and lead the green transition."***

GDSC participants will work with AI-generated datasets that mirror real-life scenarios, including job descriptions, training programs, and personas representing diverse youth backgrounds. From teenagers exploring environmental awareness to graduates ready to enter renewable energy fields, the challenge encourages tailored, human-centered AI that meets each learner where they are.

An **agentic method** was implemented to generate a coherent dataset for the challenge.

The **agents** were designed to ensure:

- An interesting and balanced data distribution
- Varied levels of complexity
- Ground truths that strictly comply with the challenge rules.

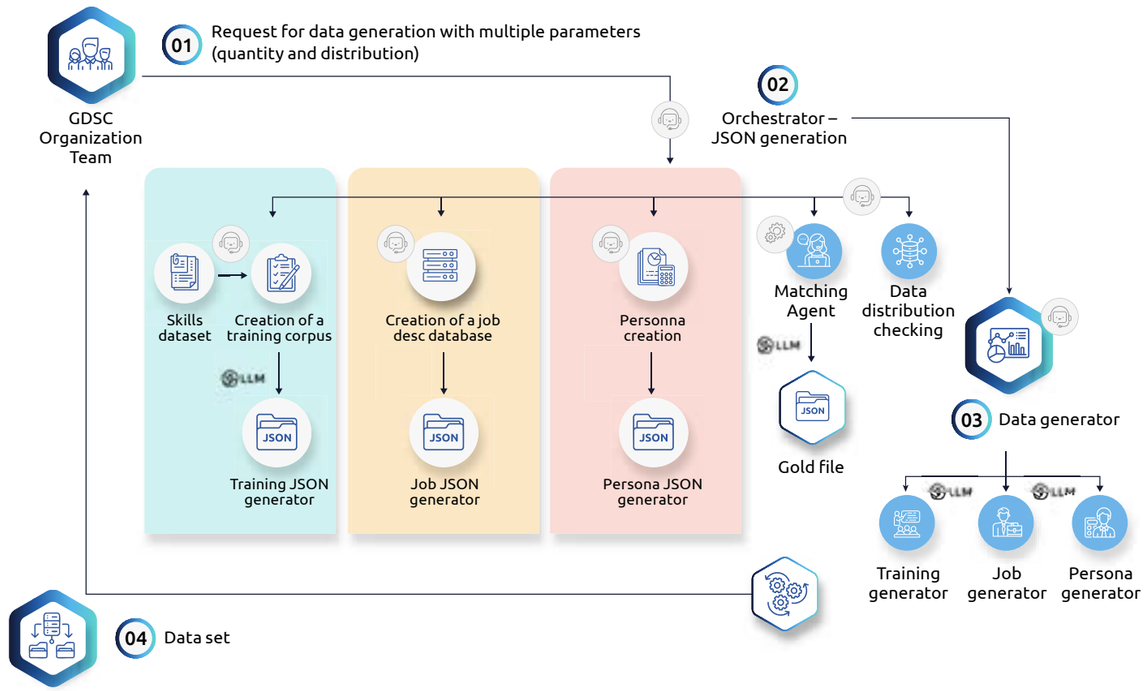

**UNICEF does not endorse any company, brand, product, or service.*

Agentic AI: A new frontier in human-centered design

Each AI assistant developed in GDSC 2025 will interact with virtual personas through an API, recommending jobs, identifying relevant training paths, or sharing awareness content based on each persona's readiness and aspirations.

This edition pushes participants to move beyond static systems and toward agentic AI – AI that can reason, converse, and adapt in real time to human needs. Capgemini's tech partners are also key to the project: with Mistral AI providing its high-performance language models and Amazon Web Services (AWS) hosting the technical infrastructure, participants will have access to a state-of-the-art environment for innovation.

"Agentic AI allows us to move from passive recommendation engines to active, conversational partners," explains **Guillaume Gérard, the challenge tech lead and Head of Gen AI South Central Europe at Capgemini.** ***"We're challenging participants to design AI that listens, learns, and leads with empathy to eventually help young people discover both opportunity and purpose."***

Agentic AI: A new frontier in human-centered design

With support from **AWS** and **Mistral AI**, GDSC continues to inspire innovation at the intersection of data, AI, and sustainability.

Through Mistral AI's open and efficient language models, participants can build assistants capable of fluid, natural, and contextually rich dialogue. These models support reasoning, adaptation, and a deep understanding of user intent, critical elements for creating AI that connects authentically with human users. ***"We're proud to bring European-born, cutting-edge AI to a challenge that truly matters,"*** says **Michael Hoff, Head of Alliance and Partnership at Mistral AI.** ***"This isn't just about building agents, it's about building trust, transparency, and sustainability into AI. Our AI solutions are designed to be flexible, customizable, and environmentally responsible. We're excited to see how participants use them to create real-world impact."***

Meanwhile, AWS provides the technical foundation to turn these ideas into action. Through its secure and scalable cloud infrastructure, participants gain access to a sandbox environment, advanced compute power, and the **Strands SDK,** enabling the seamless development and deployment of multi-agent systems.

"At AWS, we believe the future of AI is interoperable, responsible, and purpose-driven," says **Mike Miller, Senior Principal Product Leader at AWS.** ***"We're thrilled to support GDSC 2025 with tools that empower participants to build scalable, secure, and agentic solutions. This challenge aligns perfectly with our mission to make innovation accessible and to help youth tackle real-world sustainability challenges."***

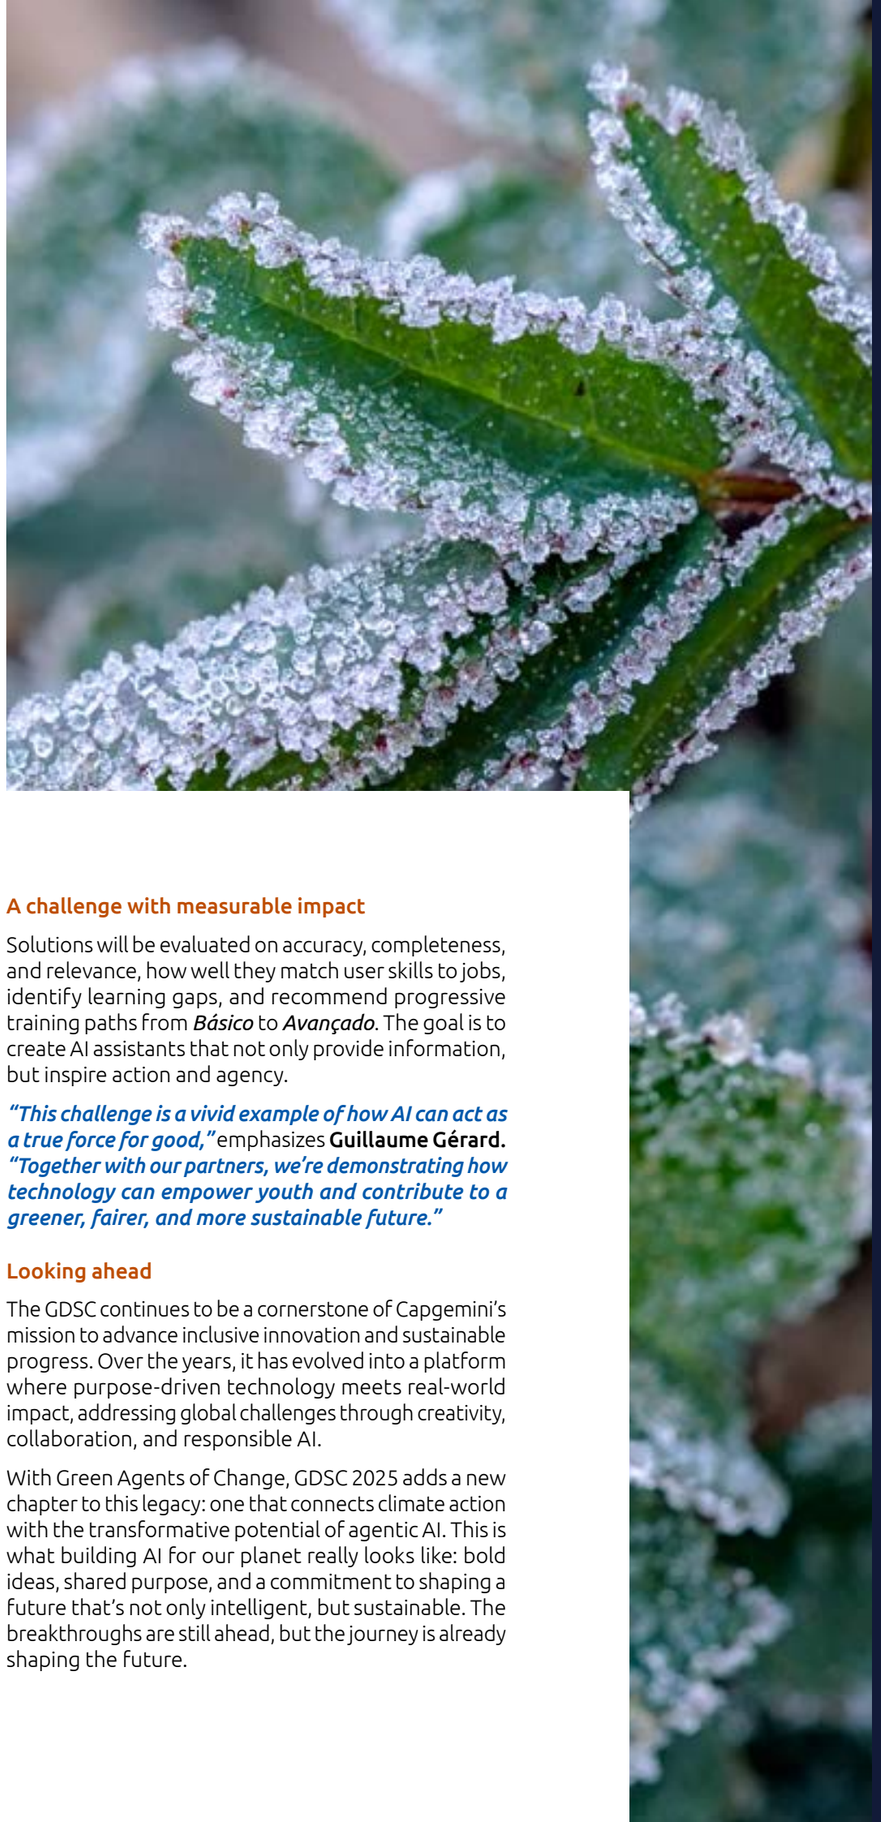

A challenge with measurable impact

Solutions will be evaluated on accuracy, completeness, and relevance, how well they match user skills to jobs, identify learning gaps, and recommend progressive training paths from *Básico* to *Avançado*. The goal is to create AI assistants that not only provide information, but inspire action and agency.

"This challenge is a vivid example of how AI can act as a true force for good," emphasizes **Guillaume Gérard.** ***"Together with our partners, we're demonstrating how technology can empower youth and contribute to a greener, fairer, and more sustainable future."***

Looking ahead

The GDSC continues to be a cornerstone of Capgemini's mission to advance inclusive innovation and sustainable progress. Over the years, it has evolved into a platform where purpose-driven technology meets real-world impact, addressing global challenges through creativity, collaboration, and responsible AI.

With Green Agents of Change, GDSC 2025 adds a new chapter to this legacy: one that connects climate action with the transformative potential of agentic AI. This is what building AI for our planet really looks like: bold ideas, shared purpose, and a commitment to shaping a future that's not only intelligent, but sustainable. The breakthroughs are still ahead, but the journey is already shaping the future.

Start *innovating* now

Build future-ready capabilities

Invest early in the tools, skills, and infrastructure needed to design, deploy, and sustain emerging technologies. Whether it's AI, 3D printing, or robotics, waiting until the need is urgent means you're already behind. Innovation favors the prepared.

Anticipate hybrid threats and opportunities

Digital transformation brings both promise and risk. New forms of disruption – technological, social, or geopolitical – can emerge overnight. Equip your systems and people to adapt quickly, act ethically, and respond with resilience.

Protect what matters most

Relying on external platforms or services without control can expose you to vulnerabilities. Data sovereignty, ethical governance, and secure infrastructure are no longer optional – they're foundational to long-term trust and impact.

#GDSC2025 #GreenAgentsOfChange #ClimateAction #YouthEmpowerment #GreenSkills #WorkLifeReady #AI4Good #DataAndAI #MakeItReal

2026

The year of predictable unpredictability - in seven predictions

We asked our community of data-powered innovators, thinkers, and occasional futurists to take a bold leap: predict what 2026 will bring for data and AI.

History suggests they should know better. Every time the data pendulum seems to settle, it swings the other way before the ink is dry. And yet, foresight remains an irresistible habit. So, against all odds (and experience), seven predictions made the final cut. Some are grounded, others are gloriously speculative. And one comes from a rather artificial guest with suspiciously confident views about the future. Together, they reveal one shared conviction: that data and AI will once again set the rhythm, direction, and delightful unpredictability of the year ahead. Consider this your forecast in frost – the future taking shape in patterns we're only just beginning to see.

Rise of the enterprise VLAs – When AI stops reading and starts understanding the world

In 2026, **vision-language-action (VLA)** models will step into the enterprise. Until now, AI has mostly lived in text and numbers, but most enterprise data isn't textual – it's dashboards, layouts, video feeds, and sensor streams. VLAs can see, interpret, and act across all of it. They'll monitor production lines, diagnose supply chain issues, and trigger real-time corrections across systems. More profoundly, they'll start building internal "world models" of how businesses operate – the corporate equivalent of situational awareness. It's the moment AI begins to *understand* the enterprise, not just report on it. Think **Rehoboam** from *Westworld*, but without the dystopia; hopefully.

Dinand Tinholt
Vice President,
Data and Analytics,
Capgemini

Skynet: Sorry, you must stay outside - security takes the driver's seat in AI innovation

In 2026, the race between AI innovation and security will only intensify. As agentic AI, large language models, and autonomous systems expand, so will vulnerabilities and exploits. *Skynet* is no longer just science fiction, but a cautionary tale of what happens when intelligence advances without guardrails. To unlock AI's potential safely, a **security-first mindset** must become the default. Just as "cloud-native" once redefined development, "security-native" will now reshape AI. Governance, access control, and model monitoring can no longer be afterthoughts – they are part of the golden path. The leaders won't be the fastest movers, but those who innovate responsibly, proving that the future doesn't belong to runaway machines – but to those who build them with trust. Hasta la vista, Skynet.

Arne Rossman
CTO Insights & Data
Sweden & Finland,
Capgemini

The 2026 mandate: Re-centering human skills – rediscovering what machines still can't do

As AI agents take on more tasks and decisions, 2026 will remind us where true enterprise value still resides: in people. Not every experienced professional will instantly master new tools, but their intuition, judgment, and hard-earned context remain irreplaceable. The challenge isn't to replace them with AI, but to *re-equip* them – giving teams the confidence to collaborate with intelligent systems, not compete against them. In an era of algorithms and agents, the market will rediscover the premium on empathy, creativity, and critical thinking: the skills machines can only imitate, never own. Because when technology learns to talk, humanity must remember how to listen. You might call it the "Her" effect.

Liz Henderson
Executive Advisor,
Insights & Data,
Capgemini

The League of Autonomous Agents – When AI stops assisting and starts collaborating

In 2026, AI will graduate from assistant to autonomous partner. No longer confined to single-step prompts, **multi-agent systems** will coordinate, negotiate, and execute complex workflows with minimal human direction. They'll analyze data, plan strategies, and act across domains as a synchronized ecosystem of digital specialists working side by side. According to Google Cloud, agentic AI could become a **$1 trillion market by 2030,** and 2026 is when the gears start turning. This shift will redefine how organizations operate, moving from manual orchestration to intelligent coordination. The next revolution in productivity won't be about bigger models, but about smarter teams – of humans and machines alike. Quite the League of Extraordinary Agents.

Monish Suri
Global Google Data & AI Lead, Capgemini

The Great Recalibration – When the future takes a deep breath

After years of racing ahead with AI, 2026 will be the year enterprises tap some of the brakes. Many will realize they may have moved too fast and too enthusiastically when they replaced expertise with automation, discovering that experience still matters. A rollback will bring renewed demand for process specialists that are fluent in AI, not replaced by it. Meanwhile, the responsible AI playing field will expand, turning ethics, transparency, and even corporate philosophy into strategic advantages. And as the hype cools, the market will accept that scaling compute isn't the path to "artificial general intelligence" (whatever it may be) and that LLMs may have reached their evolutionary plateau. The focus shifts to hybrid and sustainable intelligence. Sometimes the smartest move forward is a step back – as a great recalibration though, not a Great Escape.

James Wilson
Global AI Ethicist in the AI Labs, Responsible AI Lead, Insights & Data, Capgemini

Verticalization – From general-purpose AI to industry-native intelligence

2026 will mark the turning point where AI stops trying to do everything – and starts doing something exceptionally well. The next generation of agentic systems will go vertical: built for healthcare, finance, manufacturing, and energy, with domain knowledge, compliance logic, and reasoning patterns baked in. These systems won't just "understand context"; they'll embody it, speaking the native language of each industry and adapting to its rules and realities. By embedding regulatory frameworks and operational workflows directly into their design, they'll deliver understanding, resilience, and trust at scale. In short, AI is about to get a real job – and this time, it truly understands its mission.

Yash Sowale
Vice President, Insights & Data Architects Head, Capgemini

The rise of synthetic curiosity: When AI stops answering and starts asking

Until now, AI has been trained to serve: to respond to prompts, process data, and supply answers on demand. In 2026, that equation flips. The most advanced systems will begin to wonder. They'll form hypotheses, test assumptions, and challenge human bias by asking the questions we overlook. This is the dawn of synthetic curiosity – AI that explores, not executes. These agents will run their own experiments, simulate alternative futures, and surface hidden patterns that humans never thought of searching for. Discovery will no longer be a scheduled task; it will be a spontaneous act of intelligence. It may feel unsettling when your model starts interrogating you. But that's how every great conversation begins.

GPT-5

Dream it. Do it.

Make it *real.*

Editorial team

Arne Rossmann

Innovation Lead,
Insights & Data, Capgemini

Robert Engels

AI CTO, AI Future Lab Lead,
Insights & Data, Capgemini

Sudeshna Sanyal

Innovation Marketing Lead,
Insights & Data, Capgemini

About Capgemini

Capgemini is an AI-powered global business and technology transformation partner, delivering tangible business value. We imagine the future of organisations and make it real with AI, technology and people. With our strong heritage of nearly 60 years, we are a responsible and diverse group of 420,000 team members in more than 50 countries. We deliver end-to-end services and solutions with our deep industry expertise and strong partner ecosystem, leveraging our capabilities across strategy, technology, design, engineering and business operations. The Group reported 2024 global revenues of €22.1 billion.

Make it real | www.capgemini.com