Grid Dynamics white paper

Agentic AI: The next evolution in enterprise automation

www.griddynamics.com

Grid Dynamics

Contents

Executive summary

In an era where speed, operational excellence, and client-centricity drive competitive advantage, enterprise leaders must leverage every resource at their disposal and embrace automation as a strategic enabler. What began in the 1990s as scripted automation—where simple scripts handled repetitive tasks with little adaptability and required manual updates—evolved into robotics process automation and AI assistants over the following decades. Today, Agentic AI—a promising new frontier in enterprise automation—is reshaping the way organizations streamline decision-making and manage complex processes through autonomous reasoning and dynamic orchestration.

> ***Globally, the agentic AI market—valued at $5.1 billion in 2024—is projected to reach $47.1 billion by 2030, accounting for a notable subset of overall generative AI spending. Meanwhile, Gartner predicts that by 2028, at least 15% of daily business decisions will be made autonomously through agentic AI but, also by that time, 25% of enterprise breaches will be tied to AI agent abuse.***

Agentic AI autonomously plans and executes complex tasks through iterative cycles of reasoning and action. Unlike traditional AI systems which only respond to queries or follow fixed rules, much like a passive assistant, agentic AI introduces “agents” capable of taking initiative and adapting in real time by understanding objectives, interacting with tools, learning from outcomes, and dynamically adjusting their approach.

It’s like having a qualified team member to whom you can delegate high-level goals—someone who interprets instructions, takes initiative, and coordinates every step needed to meet the objective. This enables organizations across retail, financial services, technology, manufacturing, and more to automate sophisticated business processes while still maintaining oversight and compliance. For example, in a customer service context, an AI agent can move beyond basic question-answering by checking a user’s outstanding balance, recommending which account to draw funds from, and completing the transaction when prompted. Throughout this process, the agent can conduct a natural conversation, adapt to unexpected questions, and seamlessly handle requests that arise in real time.

In this white paper, we explore how organizations can harness agentic technology to drive operational efficiency, enhance customer experience, and boost revenue growth—all while ensuring robust implementation and building trust. We take a deep dive into:

- the core value propositions of Agentic AI;
- the maturity spectrum of Agentic AI from POC to production;
- the key innovation drivers behind these systems;
- multi-agent software development as a win-win for every industry;
- preparing enterprises for Agentic AI adoption; and
- addressing critical security, bias, and safety concerns.

Core value propositions of Agentic AI for enterprises

Whether you're a data and AI leader, an engineering decision-maker, or a senior executive, you likely have several questions on your mind: Why should you adopt this technology? What does it offer that traditional solutions don't? Below is a list of compelling core value propositions to help answer those questions.

Complex multi-step processes

If a process involves multiple steps, traditional rule-based engines or machine learning systems, often referred to as workflow automation, tend to demand significant maintenance or retraining whenever the environment changes. However, Agentic AI handles multi-step processes using dynamic planning and orchestration through problem decomposition and adaptive task execution, continuously observing and learning from each outcome.

Although this approach offers greater flexibility and resilience, it also introduces new challenges. AI agents that make decisions must be auditable and understandable as errors—such as hallucinations or incorrect decisions—can cascade, jeopardizing the entire process. Therefore, solution architects must fully understand both the complexity of the tasks and the potential consequences of errors when designing and deploying these adaptive, AI-driven systems.

Knowledge-intensive tasks

For knowledge-intensive tasks, traditional approaches often require manually retrieving data from various sources—such as search engines, databases, or APIs—and then synthesizing that information before making decisions. Today, agentic knowledge assistants can perform in-depth research, use multi-step reasoning, and integrate data from multiple sources across organizational silos. They analyze large volumes of information and determine subsequent steps—such as identifying additional data requirements or clarifying user intent.

However, these systems face challenges similar to other Generative AI solutions. Hallucinations can result in inaccurate assertions, and tracking the origin of the information becomes essential, especially in high-stakes environments. For example, conflicting versions of a policy might exist in the same repository without a clear indicator of which version is valid. Ensuring trustworthiness and transparency in these knowledge-intensive scenarios demands rigorous validation of both the data and its origins.

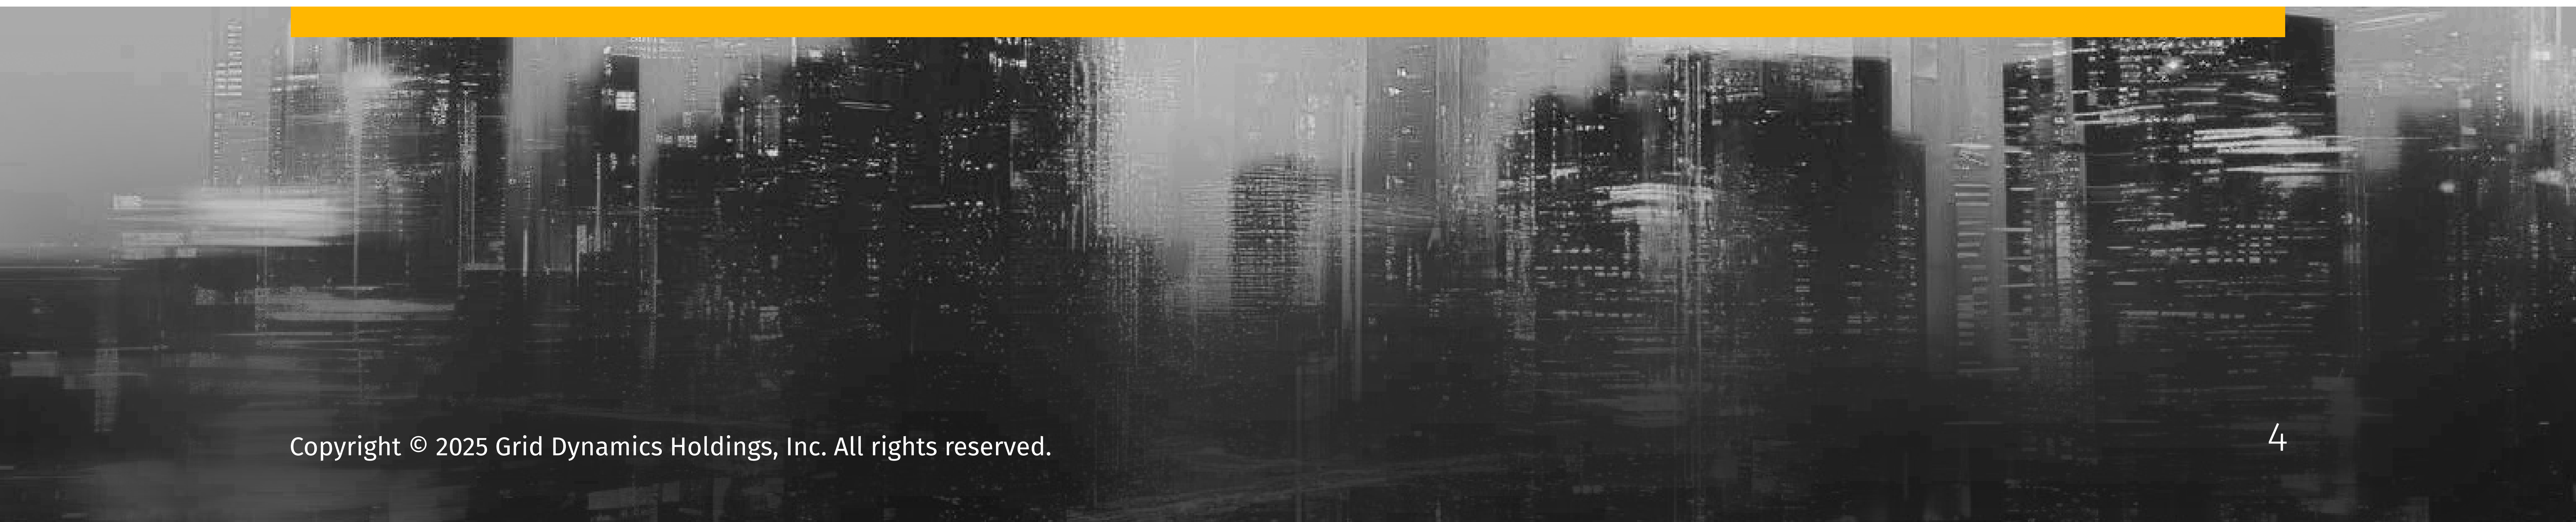

Operational decision-making

Operational decision-making is another key value proposition of Agentic AI. When the challenge is to decide when to trigger human intervention or raise alerts, traditionally, rule-based systems and threshold-based alerts have been used to hand off situations to human operators for further resolution. Agentic AI offers more advanced autonomous capabilities. It can analyze problems, gather relevant information, and apply sophisticated reasoning to determine when human escalation is necessary or when an alternative solution might be more appropriate.

This shift toward deeper automation brings its own challenges. As organizations rely more on AI-driven decisions, questions arise about how to keep humans engaged and maintain their expertise. In industries where human intervention is critical, there is a risk that essential skills may erode if not regularly practiced and reinforced. These considerations underscore the importance of striking the right balance between AI autonomy and human oversight in operational decision-making.

Customer experience personalization

Currently, customer-facing interactions largely rely on rule-based or ML-driven personalization that supports predefined journeys or offers broad recommendations based on crowd behavior. In contrast, Agentic AI has the potential to transform these experiences by providing a nuanced, real-time understanding of customer intent. By analyzing a customer's history, session patterns, and broader context—much like a knowledgeable store associate—this technology can craft highly tailored recommendations and solutions.

While deeper personalization promises to enhance customer engagement and drive business value, it also raises significant privacy concerns. Balancing personalized insights with robust data protection is essential to fully realize the benefits of this advanced technology.

Agentic AI maturity spectrum: From proven products to next-gen pilots and POCs

Having explored the core value propositions, we'll now examine the evolution of Agentic AI along a maturity spectrum—from systems in production to emerging proofs of concept that are helping enterprises like yours transform their value chains.

	Information-retrieval agents	Multi-tool retrieval agents	Agentic deep research	Task execution agents	Autonomous agents
Use case	Conversational information retrieval	Enterprise knowledge agent	Enterprise research agents	Enterprise process productivity	Enterprise workflow automation
Features	Results summarization & follow- up questions	Orchestrates multiple data sources and tools to retrieve comprehensive information	Conducts autonomous multi-step investigation with reasoning and synthesis	Adaptable task automation, human-in-the loop escalation	Autonomous planning & execution of operations
Tech	Vector search, semantic matching, stateless prompting	Multi-shot, reranking, dynamic tool selection, API orchestration, context management, parallel tool execution	Chain-of-thought reasoning, task decomposition, self-criticism	ReAct frameworks, feedback loops, tool grounding	Event-driven architectures, multi-agent coordination, enterprise guardrails
	PRODUCTION		PILOTS	POC/SOTA	

A common theme that emerges from the agentic AI maturity spectrum is the stark contrast between generative AI and agentic AI. Currently, generative AI functions like an assistant: you ask a question, it produces content, and suggests a plan of action—but it doesn't take initiative. The final decision remains entirely with you.

In contrast, the agentic approach takes this model a step further. At the initial level, agents can formulate and execute complex plans without directly changing the environment. At more advanced stages, they not only execute their plans but also begin to assume responsibility for those actions. Although accountability ultimately remains with you, these agents take on a more active role in the decision-making process.

Let's explore each agentic AI maturity level in detail below.

1. Information-retrieval agents

Today, many information-retrieval agents that function as “information search” engines are in production. They use a fixed corpus of data that is analyzed by multimodal large language models—capable of understanding graphs and charts—and techniques such as vector search and semantic matching to deliver a single-shot answer to the customer's question. If needed, the customer can ask follow-up questions, and the process is repeated iteratively. These agents can serve as robust business intelligence tools across industries.

Platforms like Google Agentspace provide pre-built connectors for the most commonly used enterprise applications, so businesses can save time deploying information-retrieval agents when they need quick answers or facilitate efficient actions. However, enterprises must partner with a co-innovation expert who understands their complex ecosystem and can help unlock their data.

2. Multi-tool retrieval agents

Multi-tool retrieval agents are the smarter successors of Generative AI knowledge assistants, functioning as enterprise knowledge agents. They consolidate information from multiple sources and plan which ones to invoke based on the customer's request.

For example, if a user asks about quarterly sales, the agent might query a CRM system, databases, and even Google spreadsheets to gather the necessary context and synthesize a response. Sometimes, the process also involves invoking enterprise APIs. Although these agents typically deliver a “single-shot” answer, they employ a cascading approach by dynamically selecting and calling multiple tools to provide a comprehensive response.

3. Deep research agents

The third stage involves ongoing pilot projects in agentic deep research. When a specific, complex insight is needed, the agent embarks on a multi-step investigation into the inquiry. It makes multiple round trips to various systems and engages in self-reflection—asking whether the current context is sufficient, if the user's intent needs clarification, or if additional tools should be introduced. This iterative loop, which includes task decomposition, external data collection, and critical reasoning, aims to deliver a comprehensive report with recommendations.

Recently released deep research agentic features by Open AI, Google Gemini, and Perplexity[4] help models perform dozens of searches, read hundreds of sources, and reason through material to autonomously deliver detailed reports.

4. Task execution agents

Until now, agents primarily analyzed data while users executed tasks. Now, task execution agents not only gather information but also have the ability to make direct changes in the environment.

For example, they might update tickets or onboard a user across multiple systems—similar to how an airline manages passenger check-ins. If a system fails or becomes unavailable, the agent escalates the issue to a human operator to resolve the blocker. Although these agents are tasked with specific goals, they do not have a view of their own. They operate on a "react" paradigm: they act, observe the results, and adjust their plan if the action is unsuccessful.

5. Autonomous agents

Finally, we come to the latest frontier: autonomous agents. At this stage, the agent no longer waits for explicit instructions when it encounters a blocker; instead, it continuously monitors its environment. Subscribed to relevant business events, the agent takes action as soon as an event occurs. It exercises a degree of independent judgment by planning and executing operations while coordinating with other agents through task delegation. Sometimes, one autonomous agent takes the lead and distributes tasks among others, enabling a coordinated response across the organization.

Key innovation drivers behind Agentic AI

Here are three major breakthroughs that help distinguish Agentic AI from traditional AI-driven automation:

1. Inference-time reasoning

Unlike traditional AI models that rely on static pre-training, inference-time reasoning enables AI to dynamically analyze complex problems and generate solutions without retraining. This capability stems from a two-system cognitive architecture that mimics human thought processes—combining fast pattern matching with deliberate step-by-step reasoning. This allows enterprises to automate workflows that require nuanced decision-making rather than simple pattern recognition. Previously, we had to explicitly instruct an agent to think step by step for deep research. Now, high-end models come pre-equipped with such chain-of-thought reasoning. They predict the next token and generate sequences that resemble logical reasoning. Flagship models from OpenAI, Google, and others use powerful optimizations to enable this capability.

> ***[1]A common question among enterprise leaders is whether deeper multi-step reasoning demands more training than earlier Generative AI systems. The short answer is no: these models rely on extensive pre-training by major AI institutions, enabling them to handle new problems without constant retraining—much like humans adapt to new challenges. While fine-tuning can be costly, chain-of-thought reasoning guides the model to predict logical continuations, addressing previously unseen problems without repeated retraining.***

While the results are impressive, there are also notable failures, reminding us that this is not true cognition—it's a language model bridging the gap between a complex problem and its solution. For example, when asked to solve a math problem, such as an integral, traditional LLMs take a "leap of faith" approach, using pattern matching to propose an answer without showing their work. In contrast, advanced reasoning-focused LLMs "build a bridge" by generating intermediate steps and explanations, significantly increasing the likelihood of arriving at the correct solution. This evolution from direct answer generation to structured reasoning represents a fundamental shift in how AI models approach complex problems, even though both approaches ultimately rely on statistical patterns rather than genuine mathematical understanding.

Customer support use case

When a customer submits a complex issue, the Agentic AI system automatically breaks it down into its key components, such as account status, technical requirements, and policy implications. It then evaluates different resolution paths and routes the case to the right specialist, providing comprehensive context along the way. The real advantage is that it can handle new, unforeseen situations without needing constant retraining or rule updates.

2. Multi-agent architectures: Unifying AI, humans, and software agents

Traditional automation frameworks treated AI, human operators, and software agents as separate components, leading to siloed decision-making and rigid workflows. In multi-agent architecture, human operators, software agents, and even databases act as agents within a common framework.

Agentic AI introduces a unified multi-agent architecture where AI agents, human users, and software systems collaborate dynamically within a single event-driven environment. For example, a database can be represented as an agent guarding access to data, while your boss might have a digital twin agent. This unified approach allows agents to communicate naturally, using natural language, coordinate tasks, invoke APIs and business applications, adapt to real-time events, and escalate complex decisions to human experts when needed. When specific protocols are needed—such as for invoking APIs—special guardrails ensure the agents follow a structured schema, creating a resilient enterprise automation ecosystem.

Coordinated claims processing use case

Imagine a claim is submitted—one specialized agent evaluates policy coverage, another analyzes fraud indicators, and yet another coordinates with external services for damage assessment, all at the same time. Human adjusters then receive a comprehensive summary, allowing them to focus on the judgment-intensive decisions that truly require human insight. This connected approach eliminates the need for traditional, sequential, siloed processing, showcasing the true benefits of multi-agent orchestration in real-world applications.

3. Enterprise-grade tool integration

AI-driven automation is only as effective as its ability to interact with enterprise systems securely and efficiently. With Agentic AI, there's now a robust framework that allows AI agents to discover, select, and orchestrate software tools based on the business context.

By combining structured API access with natural language-driven function calls, these systems enable AI to act as an intelligent middleware layer, translating high-level business objectives into precise system interactions. This integration simplifies communication between various systems and tools, ensuring that agents don’t invent non-existent APIs but instead adhere to structured information protocols. On top of that, secure execution environments, role-based access control, and comprehensive audit logging guarantee that the entire automation process remains governed and compliant with enterprise standards.

Compliant financial operations use case

With enterprise-grade tool integration, Agentic AI can seamlessly orchestrate complex workflows while ensuring adherence to regulatory requirements. For example, when processing an international transaction, an agent analyzes the request, determines the necessary compliance checks, accesses sanctions databases, verifies customer identity across multiple systems, and executes the transaction—all within a governed framework. Every decision and action is logged, creating a complete audit trail for transparency and compliance.

Collectively, these innovation breakthroughs form the backbone of Agentic AI systems. They enable autonomous reasoning, collaborative workflows between humans and software agents, and secure, compliant interactions with enterprise tools. The following diagram illustrates how these elements come together in an event-driven environment, highlighting the core components that distinguish Agentic AI from traditional automation approaches:

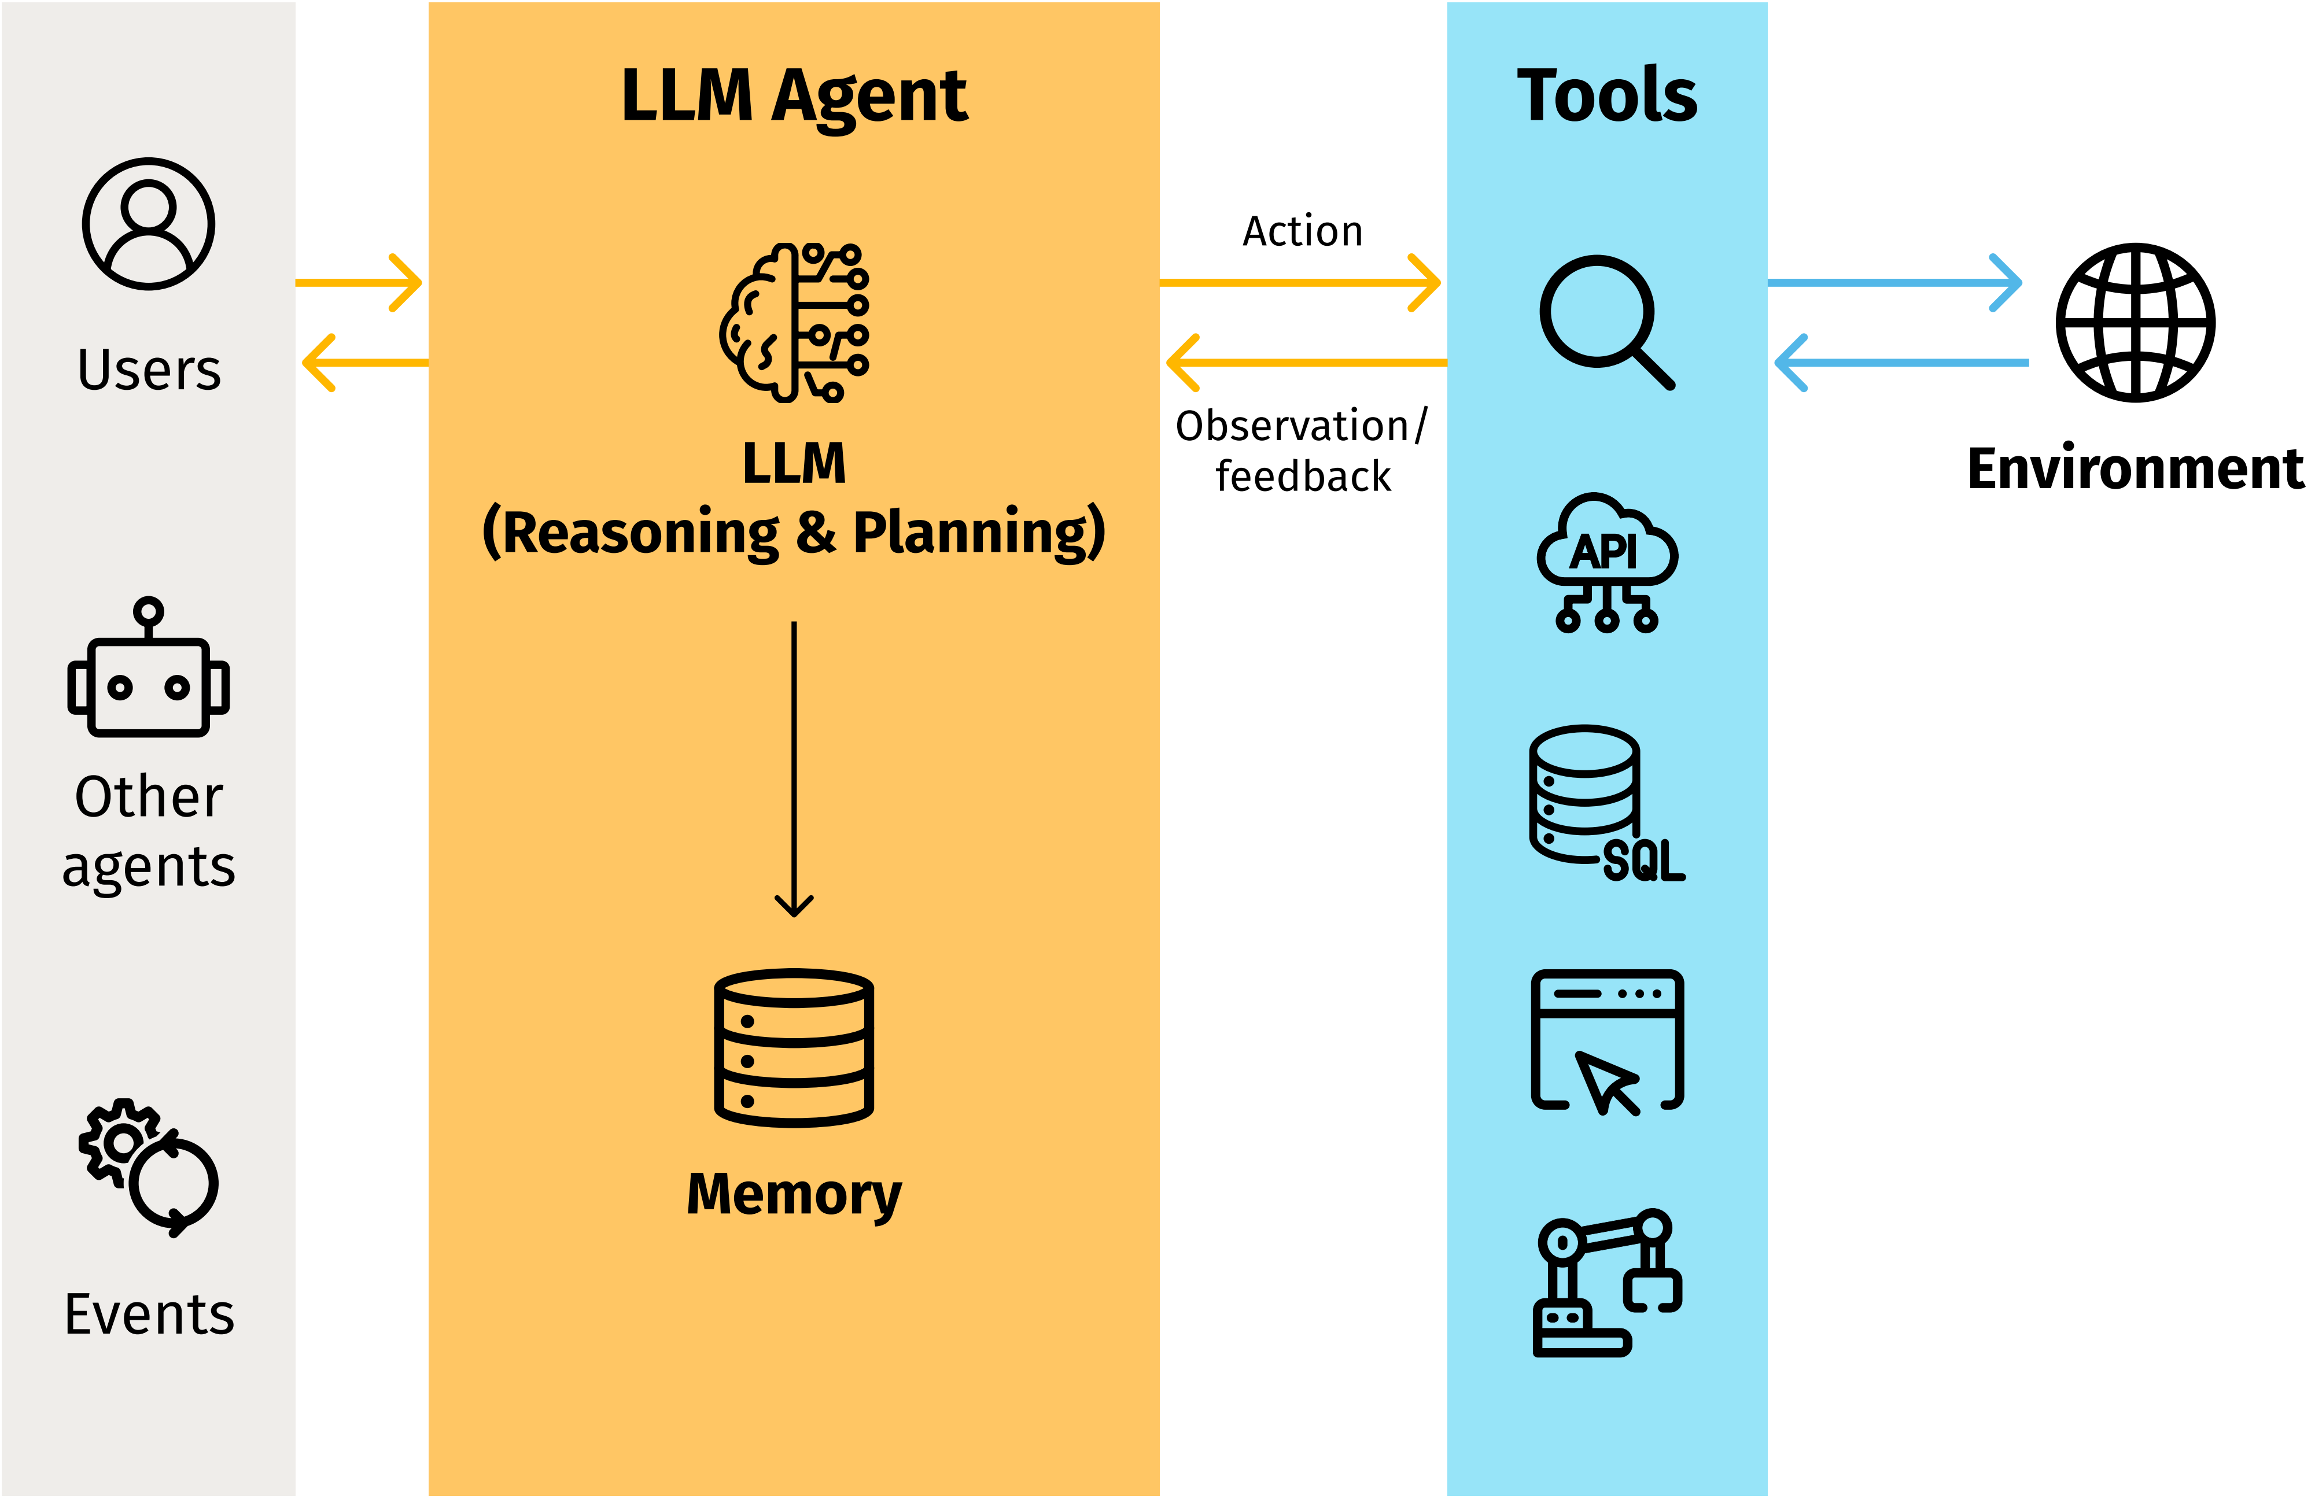

1. **LLM agent core:** Combines reasoning and planning capabilities with persistent memory.
2. **Tool integration:** Enables secure interaction with enterprise systems and external services.
3. **Event-driven design:** Supports dynamic responses to user inputs, system events, and interactions with other agents.

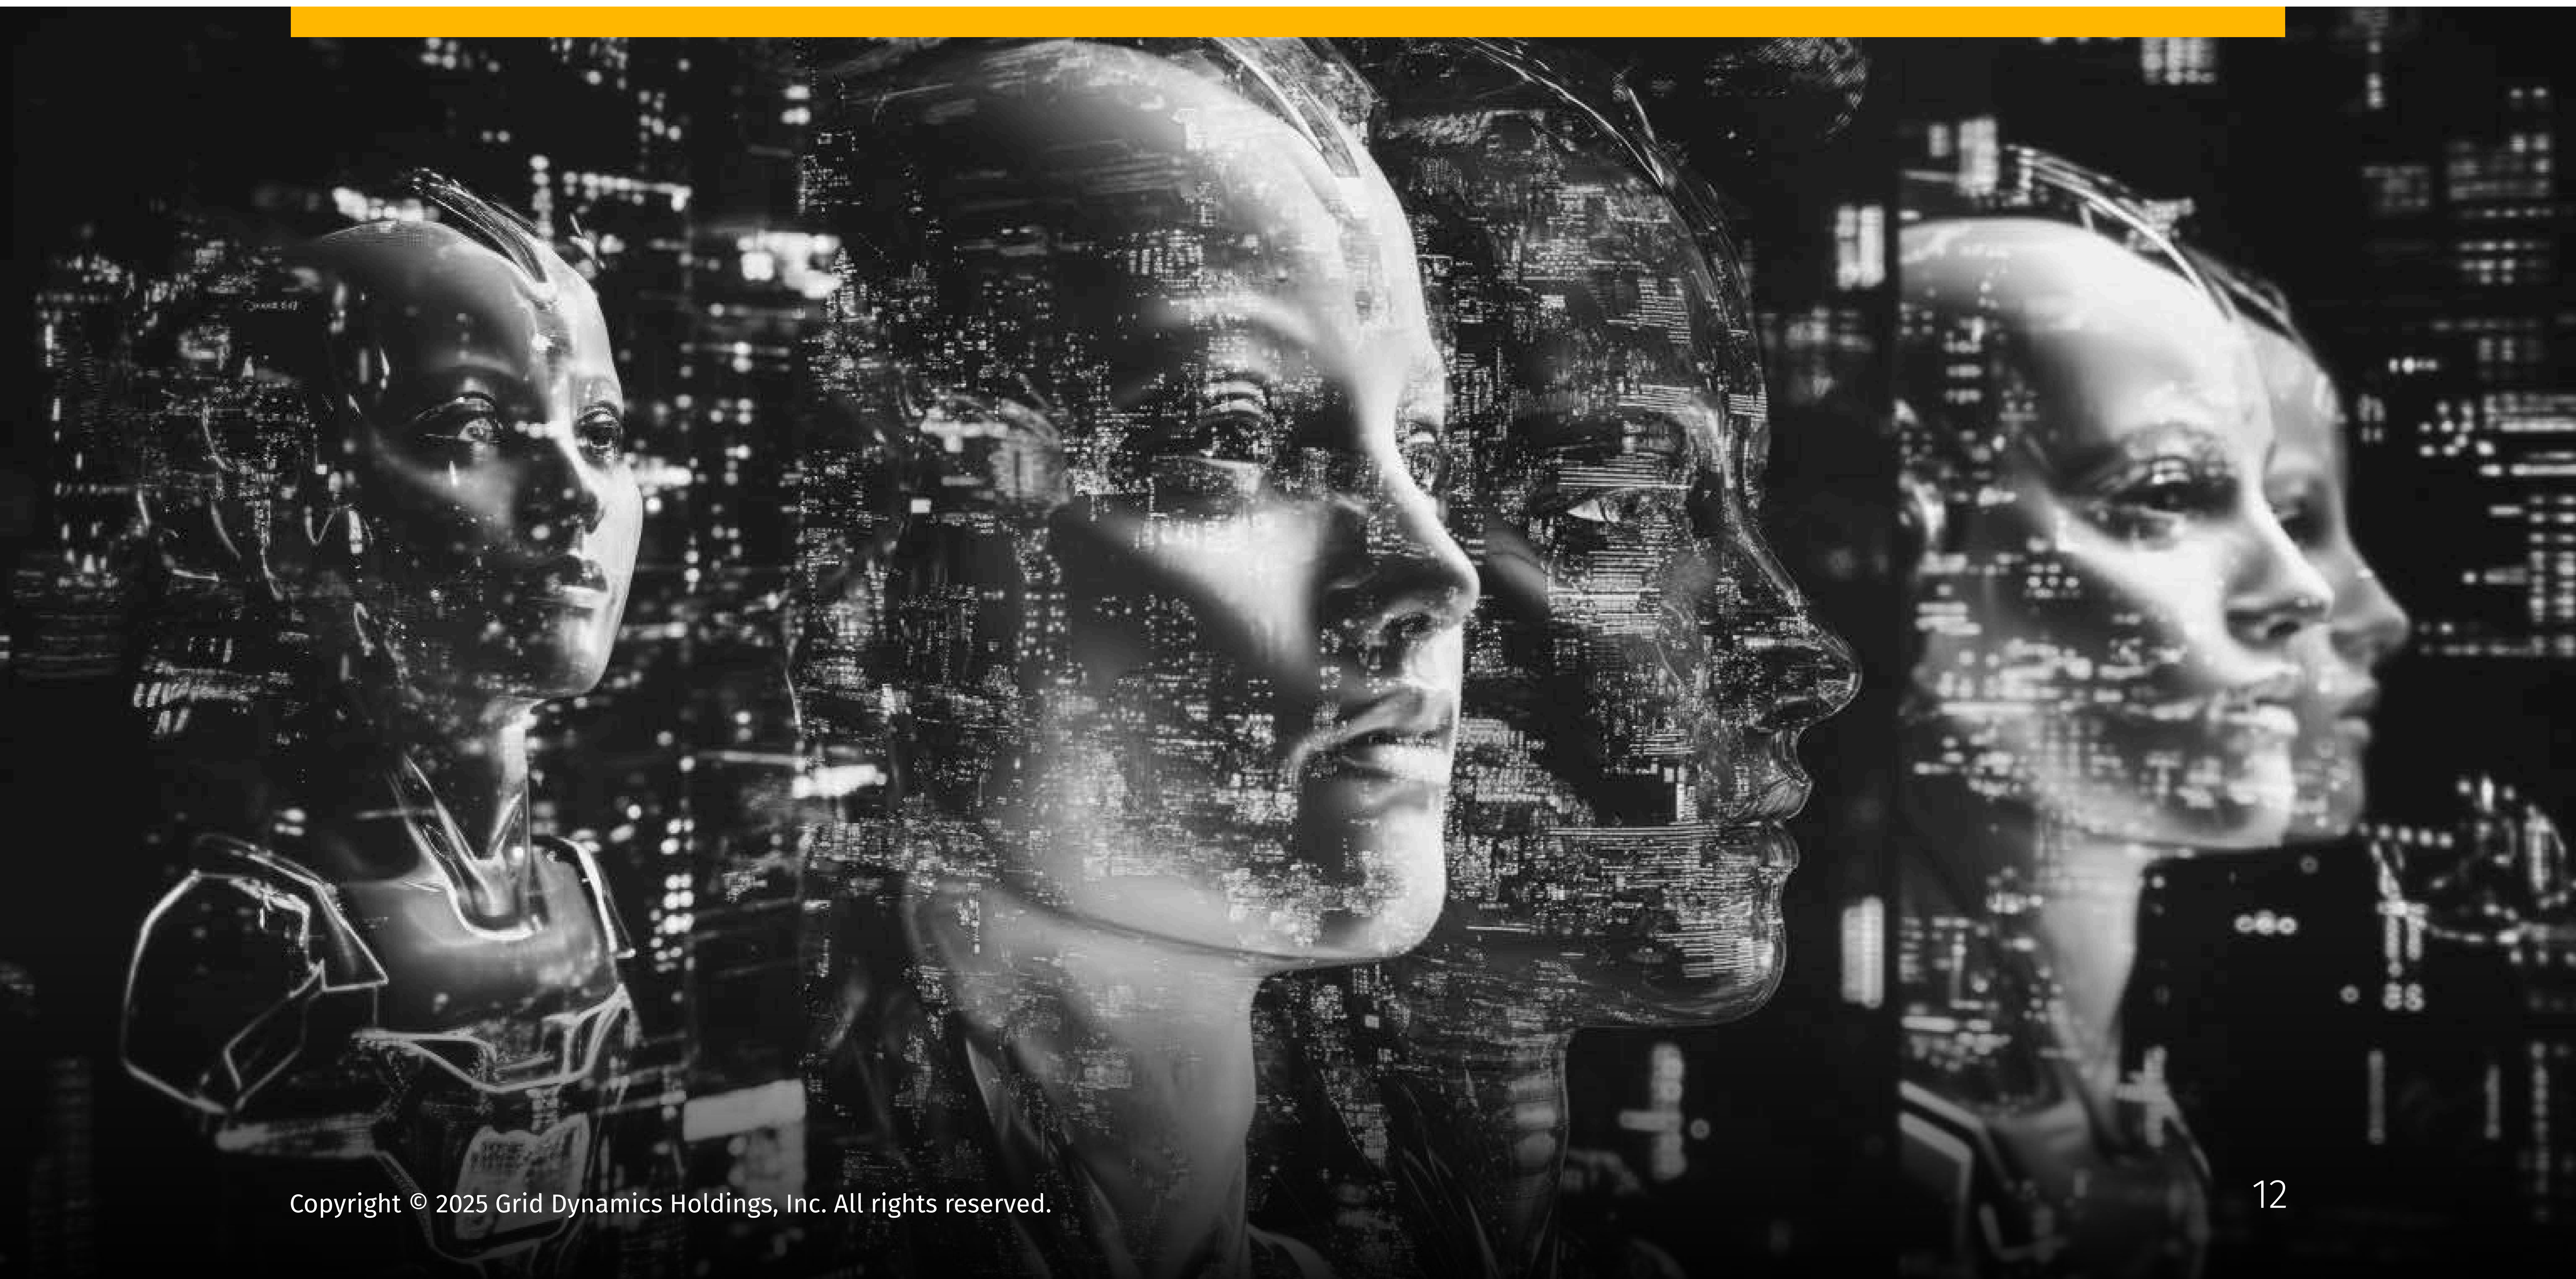

Innovation Canvas: Multi-agent software development

At Grid Dynamics, our AI experts have a deep understanding of the key innovation drivers behind robust agentic AI applications. We're leveraging this expertise to develop cutting-edge solutions across the maturity spectrum that deliver the core value propositions of agentic AI. In our quest to constantly innovate to build AI applications that resonate with business and technology leaders across industries, our AI experts have developed Innovation Canvas—a multi-agent platform foundation that enables multiple AI agents to interact with each other to perform complex tasks.

In our example use case, agents take on development roles to collaboratively build a software application. This isn't about merely assisting developers with coding; instead, these agents actively participate as team members, supporting the entire development process from architectural design to continuous testing and deployment.

Human engineers orchestrate a team of four specialized AI agents: one handles requirements and architecture, another manages full-stack development, a third ensures quality through automated testing, and the fourth oversees cloud deployment and scaling to develop a prototype of a Customer Data Platform (CDP). The end result is a dashboard displaying all metrics and charts, fetching live data from a Customer 360 backend.

Check out this demo to experience this advanced framework that demonstrates how AI-assisted development can transform the traditional software lifecycle while maintaining human strategic oversight throughout the process.

Key considerations for enterprises implementing Agentic AI

Agentic AI represents a paradigm shift in enterprise automation, empowering organizations with autonomous reasoning, adaptive workflows, and seamless AI-human collaboration. By taking a strategic approach to adoption, enterprises can unlock efficiencies, enhance decision-making, and gain a competitive edge. However, to fully harness its potential, companies must address challenges in governance, infrastructure, and operational readiness.

To successfully integrate Agentic AI, enterprises should focus on these foundational pillars:

Technology & infrastructure readiness

Modernize IT ecosystems to support event-driven, multi-agent architectures. This means adopting cloud-native environments, implementing AI orchestration frameworks, and integrating AI with enterprise systems for real-time decision-making.

Data foundation & access patterns

Ensure enterprise data readiness through structured documentation, real-time access patterns, and comprehensive governance frameworks. This nurtures AI agents to effectively access and act upon enterprise knowledge.

AI governance & risk management

Develop a structured governance model to maintain compliance, security, and ethical decision-making. Implement AI audit trails, bias detection tools, explainability frameworks, and maintain human-in-the-loop oversight to mitigate risks.

Workforce & process evolution

The success of Agentic AI hinges on effective human-AI collaboration. Reskill employees to work alongside AI, redesign workflows to maximize AI augmentation, and set clear escalation paths for critical decisions.

Addressing security, bias, and safety concerns

While Agentic AI offers compelling business benefits, its autonomous nature and access to enterprise systems introduce important considerations around security, bias, and governance. Organizations must balance innovation with appropriate safeguards to ensure reliable, ethical, and compliant deployments. Unlike traditional automation, Agentic AI operates with real-time reasoning, tool access, and workflow execution, which introduces new vulnerabilities that require proactive mitigation.

Agentic AI's ability to autonomously execute actions and interact with multiple systems creates risks of unintended consequences, data exposure, and regulatory violations. Here are some steps enterprises can take to control agent behavior, ensure data privacy and compliance:

- **Execution boundaries & control mechanisms:** AI agents should operate in sandboxed environments with explicit permissions and rate limits to prevent unauthorized actions.
- **Data access restrictions:** Agents must retrieve and process only the data necessary for their tasks, with role-based access control (RBAC), encryption, and anonymization in place to protect sensitive information.
- **Regulatory compliance & auditability:** To meet standards like GDPR, CCPA, and industry-specific regulations, organizations must maintain detailed logs, deploy explainability frameworks, and ensure that human oversight is part of the decision-making process.

As AI models inherit bias from training data, reinforcement learning, and optimization strategies, fairness and decision transparency become critical concerns. Here are some ways to address bias, reliability, and lack of human oversight concerns:

- **Bias mitigation strategies:** Use diverse training datasets, bias auditing tools, and ongoing model evaluations to minimize systemic discrimination.
- **AI decision validation:** Implement self-correction mechanisms, human verification checkpoints, and fallback protocols for ambiguous or high-stakes decisions.
- **Explainability & trust mechanisms:** Ensure that AI provides clear, understandable reasoning behind its decisions, so human operators can audit, adjust, and, if necessary, override outputs.

In short, security, privacy, and fairness in Agentic AI are not automatic—they require structured governance, robust containment strategies, and ongoing validation. AI must be designed to enhance enterprise capabilities without compromising ethical responsibility, regulatory compliance, or operational stability. Organizations that prioritize safety, oversight, and transparency will build AI systems that are trusted, scalable, and sustainable.

Final thoughts

Today, there is a lot of hype around Agentic AI, and enterprises must cut through the noise to identify the most impactful use cases that drive speed, operational excellence, and client-centricity, ultimately boosting competitiveness. To transition from experimentation to enterprise-wide adoption, organizations should take the following strategic actions:

Identify high-impact use cases

Focus on workflows where AI-driven reasoning, decision-making, and orchestration add clear value—such as enhancing customer experience, automating complex knowledge work, or optimizing supply chains.

Run pilot programs in a controlled environment

Begin with sandboxed AI deployments to test performance, integration, and governance before scaling to mission-critical processes.

Implement AI governance & compliance controls

Establish transparency frameworks, audit mechanisms, and ethical guidelines to ensure responsible AI adoption.

Develop an AI-augmented workforce strategy

Invest in training programs to upskill employees, foster a culture of AI adoption, and build effective models for human-AI collaboration.

Monitor, iterate, and scale AI initiatives

Recognize that AI adoption is an iterative process—continuously monitor performance, refine strategies, and expand AI-driven automation across departments.

While many businesses are exploring AI self-service studios that allow customers to build autonomous and task-execution agents, low-code and no-code platforms may only serve as a starting point. The complexity of enterprise ecosystems and unique business needs call for a future-proof, custom-developed Agentic AI platform—an orchestration solution that evolves with the business and manages all agents within the organization. This also prevents vendor lock-in and ensures long-term adaptability.

With 8 years of experience in delivering AI applications across industries, our experts at Grid Dynamics deeply understand the ground reality of your complex value chains. We're happy to serve as your co-innovation partner with the ambitious goal to deploy Agentic AI applications that directly tackle your business problems.

Ready to transform your enterprise with adaptive, AI-driven automation? Contact us to schedule a free consultation.

References

1. Barry, L. (2024, December 19). Agentic AI accelerates network effects—Corporate and customer value. Forbes. https://www.forbes.com/sites/libertbarry/2024/12/19/agentic-ai-accelerates-network-effects--corporate-and-customer-value/
2. Coshow, T. (2024, October 1). Intelligent agents in AI really can work alone: Here's how. Gartner. https://www.gartner.com/en/articles/intelligent-agent-in-ai
3. Gartner. (2024, October 22). Gartner unveils top predictions for IT organizations and users in 2025 and beyond. Gartner. https://www.gartner.com/en/newsroom/press-releases/2024-10-22-gartner-unveils-top-predictions-for-it-organizations-and-users-in-2025-and-beyond
4. Anthem Creation. (2025, February 20). Deep research: Perplexity, OpenAI, Gemini – Who's the best? Anthem Creation. https://anthemcreation.com/en/artificial-intelligence/deep-research-perplexity-openai-gemini-who-is-the-best/

About Grid Dynamics

Grid Dynamics empowers enterprises to drive growth, boost efficiency, and transform their digital capabilities through expert technology consulting, platform engineering, AI, and advanced analytics.

Fusing technical vision with business acumen, Grid Dynamics (NASDAQ: GDYN) fuels a relentless drive toward sustainable digital transformation. As a forefront provider of technology consulting, platform and product engineering services, and bespoke software development, we draw from over 8 years of leadership in Enterprise AI, coupled with profound expertise in cloud, data, and advanced analytics.

Our commitment to engineering excellence, R&D leadership, a co-innovation ethos, globally efficient follow-the-sun delivery model, and an unwavering "whatever it takes" dedication to client success empower us to solve even the most complex enterprise challenges, ensuring profitable business outcomes and future-proof growth.

Founded in 2006, Grid Dynamics is headquartered in Silicon Valley and has a global talent pool of intellectually curious problem solvers in offices across the Americas, Europe, and India.

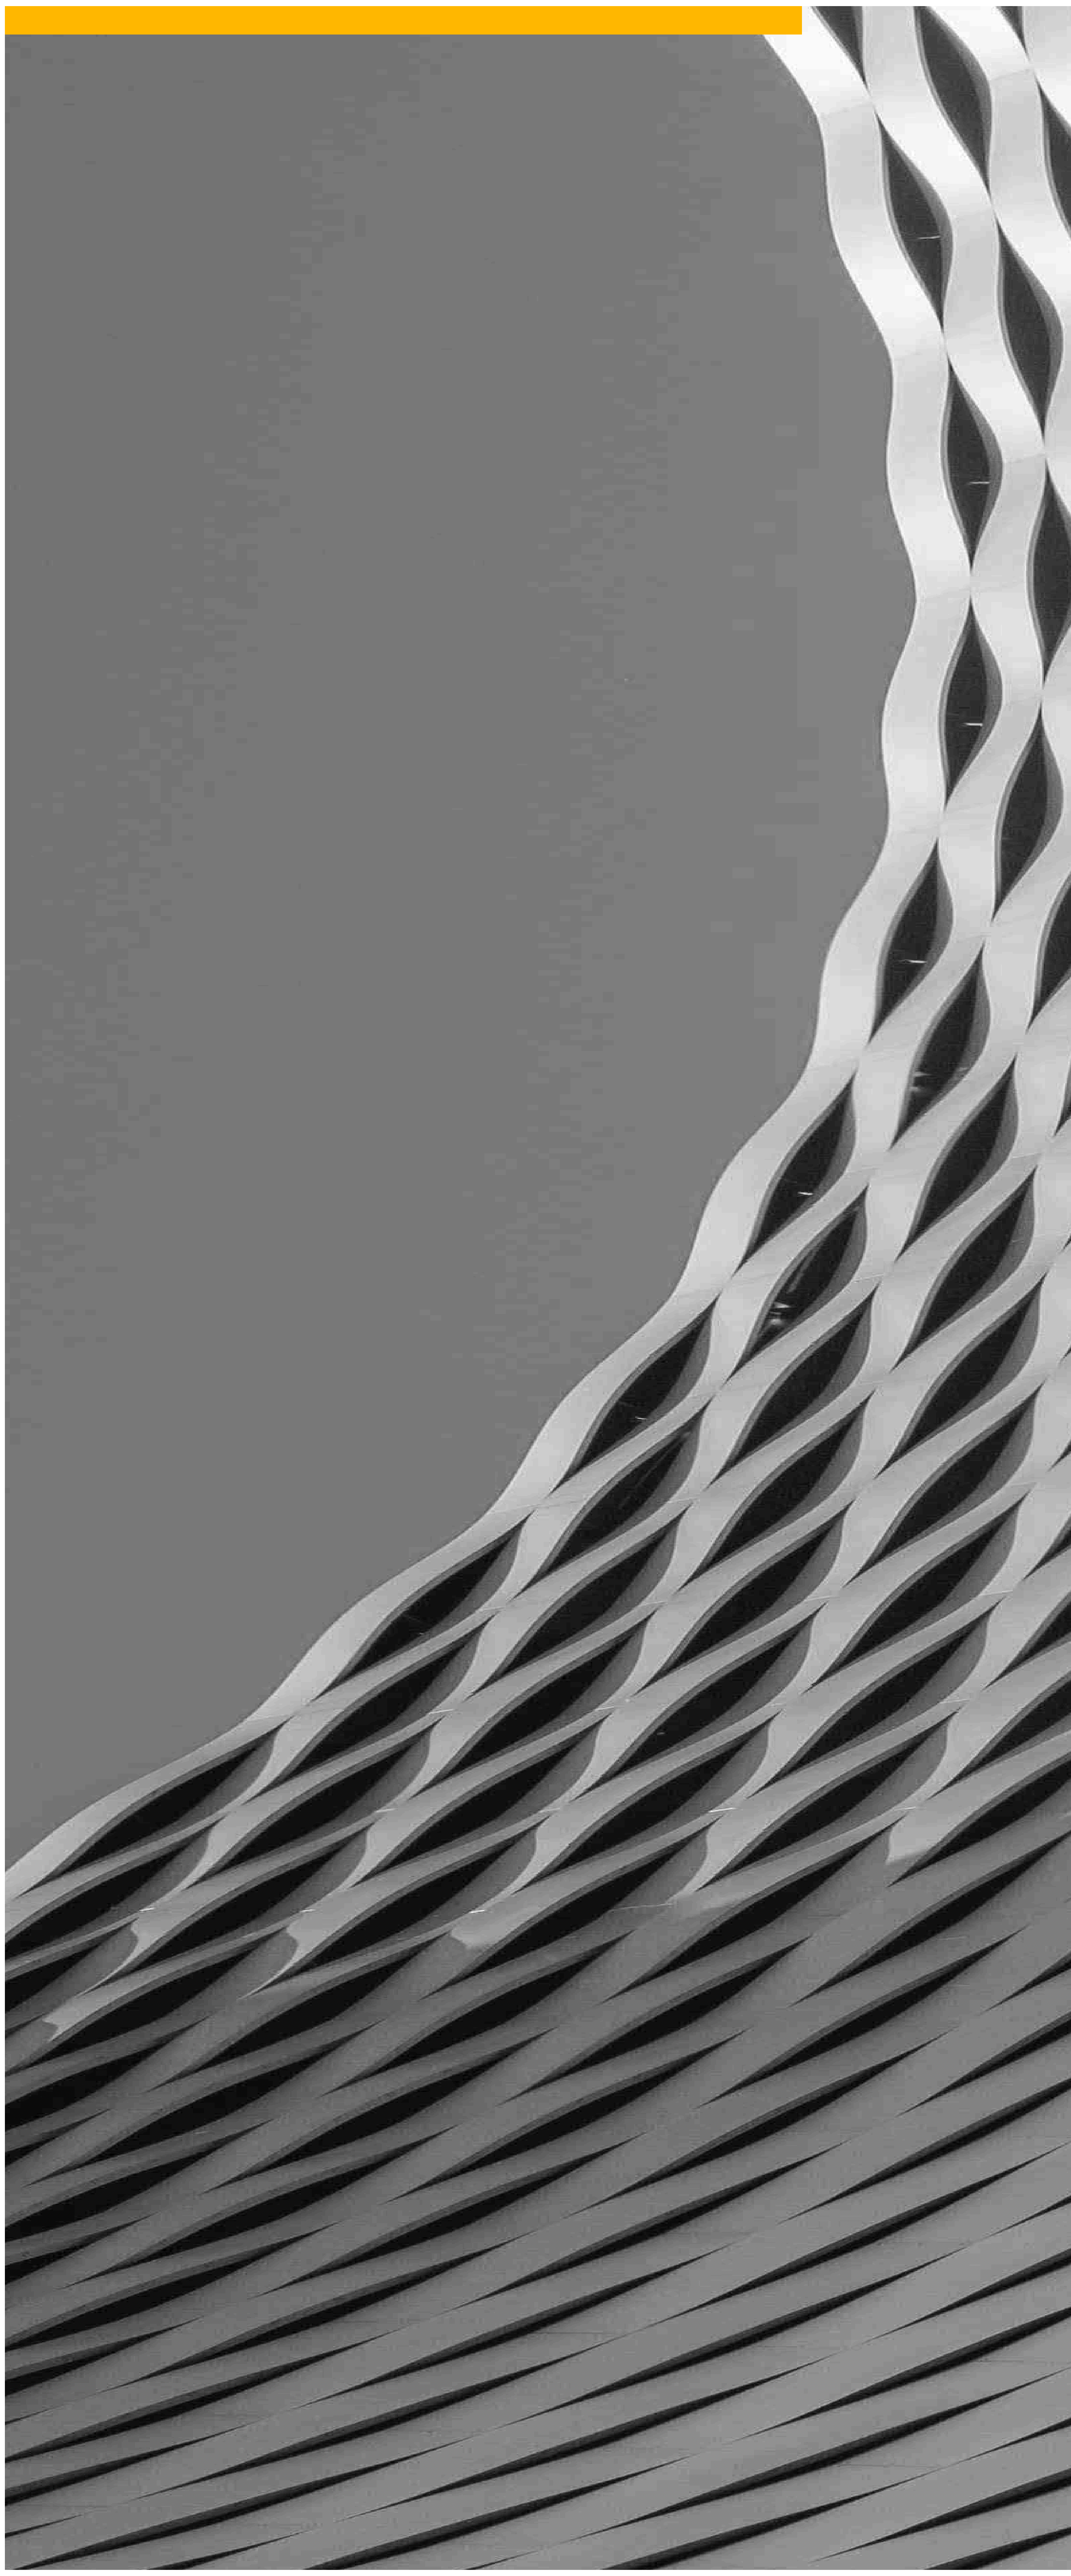

LEARN MORE AT GRIDDYNAMICS.COM →

Grid Dynamics

trusted engineering partner for digital transformation

Grid Dynamics Holdings, Inc.

5000 Executive Parkway,
Suite 520 / San Ramon, CA
650-523-5000
www.griddynamics.com