


DRIVING BETTER BUSINESS OUTCOMES THROUGH MULTI-AGENT SYSTEMS



Increased AI mandates are the new normal in the enterprise, and organizations are fueling agentic AI with enterprise data to turn vast amounts of data into actionable knowledge and enhance productivity and operations. The promise and progress of agentic AI is immense and real. It's not plug and play; it's about designing systems that drive business value—through cost, quality, speed, and scale.

According to a [Gartner](#)[®] prediction, "over 40 percent of agentic AI projects will be cancelled by the end of 2027, due to escalating costs, unclear business value or inadequate risk controls." We hold the view that agentic AI can be a game-changer, just not everywhere. That cancellation rate says less about the technology and more about how organizations are approaching it. In situations where I've seen agentic AI succeed, the conversations start with teams who understand when to use AI agents, when to stick to automation, and when to rethink the workflow.

How AI agents differ

[Agentic AI](#) uses sophisticated and iterative planning to autonomously solve complex, multi-step problems. A system of AI agents for IT operations (ITOps), for instance, can triage and isolate the root cause of a service issue by assessing the situation and identifying potential causes based on contextual data from the service topology, telemetry signals, and change events. There are different types of agents for different job functions:

- **Non-deterministic agents** use reasoning to figure out which tasks need to be completed and

track what's already been done. BMC HelixGPT Situation Observer is a good example of this type of agent. It learns to dynamically analyze context across infrastructure, application, and network resources; reduce event noise; and optimize mean time to resolution (MTTR) based on the situation-driven workflow—all without human intervention.

- **Deterministic agents** execute workflows that have pre-determined steps. BMC HelixGPT Employee Navigator is an example of a deterministic agent, and it takes a big step toward delivering better, more natural, and human-like self-service support experiences by “seeing and hearing.” When called upon, the agent accesses a list of tools and uses the best one to provide the right information to speed resolutions and fulfillment.

Oftentimes, you'll need a combination of AI agent types to address complex problems. AI agents can do a lot of interesting things, but they offer the highest value when they work together in a multi-agent system to resolve process-intensive situations.

BMC Helix's multi-agent system approach

Agentic AI isn't about a single “magic” agent—it's about orchestrating the right mix of models, techniques, tools, data, and structured workflows with autonomous decision-making to solve complex problems. Tasks that used to require multiple specialized tools can now be handled by a single, well-designed multi-agent system. This approach multiplies productivity and reduces costs by using the right tools for the job.

BMC HelixGPT Best Action Planner is a good example of an orchestrator or planner that gives you the best remediation steps to fix a problem, along with relevant IT planning and optimization information. This multi-agent system takes an organization's unique data and processes into consideration when formulating an action plan for resolving an incident.

When a problem occurs, the multi-agent system follows the same process as human site reliability engineers (SREs) or DevOps engineers in diagnosing and resolving it. It goes through a plan of analysis by calling the domain-specific agents to process the following information and make decisions based on them:

- Logs
- Metrics
- Traces
- Change (recent, past, similar changes)
- Documents
- Workflows

Trained on your incident data, chat scripts, root cause analysis, and its ability to understand and correlate service level objectives (SLOs), dependency graphs, and change logs, the multi-agent system can replace a lot of the manual investigating that SREs and DevOps engineers must typically do to narrow down the scope of a problem. This orchestration and planning layer will be increasingly important as AI agents are implemented across the business.

As different parts of the organization implement AI agents and automate manual processes, you can quickly end up with patchwork-like architectures that become too complex to rethink. CIOs need to develop a multi-agent framework from the outset to break AI agents into sub-agents, where each one can be added to the system, dictated by new workflows. Asking questions such as, “What decisions can we safely offload?” and “Where can actions be standardized?” are critical to ensure

successful outcomes.

Realizing business value

The real value of agentic AI is in driving enterprise productivity, rather than just automating individual tasks. That means adopting a multi-agent system across service management and operations that aligns to strategic goals like:

Revenue growth: A better employee experience can significantly drive revenue growth. Engaged employees are happier and more productive and innovative, which in turn leads to increased sales and revenue. With BMC HelixGPT Employee Navigator, employees are empowered to resolve tickets in a timely manner through faster access to enterprise-wide knowledge and the routing of automated escalations to appropriate teams.

Cost optimization: Optimizing cost is a critical part of organizational sustainability. CIOs can use AI and integrated workflows to automate repetitive tasks, lower ticket volume, and streamline incident response. By creating a single incident for many correlated events, along with the impacted service and configuration item (CI) details, BMC HelixGPT Situation Observer significantly reduces the number of tickets created, as well as their corresponding costs and incident noise.

Risk management: CIOs can use AI agents to help speed vulnerability remediation. BMC HelixGPT Vulnerability Resolver is designed to accelerate and strengthen your vulnerability management process, making it easier than ever to visualize, prioritize, and act on vulnerabilities at scale. With the flexibility to choose between SaaS or an on-premises deployment, CIOs can align deployment with security requirements, regulatory needs, or infrastructure preferences.

Conclusion

AI, specifically agentic AI, isn't a nice-to-have anymore — and the right [enterprise agentic orchestration platform](#) provides the governance and scalability that multi-agent deployments require to move from pilot to production. As AI continues to transform IT work, CIOs and IT leaders must take a strategic approach to implementing agentic systems into their processes. Don't just plug agents into legacy workflows and expect transformation—it's time to rethink workflows entirely.

*Gartner Press Release, Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027, June 2025, <https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027>.

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.