CEQUENCE

Top Considerations for Enterprise Agentic AI Projects

AI has rapidly moved from hype to reality. As generative and agentic AI projects proliferate, enterprises are discovering that many initiatives fail to deliver measurable business value. Industry research indicates that a majority of generative AI pilots are stalling or failing due to unclear ROI, governance gaps, security risks, and workflow integration challenges.

These outcomes are not unusual for emerging technologies, but they highlight a critical need: enterprises must approach agentic AI with the same rigor applied to other mission-critical systems. Based on real-world deployments and lessons learned from large enterprise customers, the following considerations outline how organizations can accelerate value from agentic AI while maintaining a strong security and governance posture.

Cequence Security has a unique perspective on how to address this challenge due to its experience with enterprise applications and data and how entities, both human and synthetic, interact with them. The Cequence AI Gateway enables organizations to connect AI agents to their data and applications, securely, without coding. Since its introduction, several of the largest Cequence customers have used the AI Gateway to evaluate and deploy cutting edge agentic AI solutions. Based on that experience, we have developed the following list of things to consider for enterprises undertaking agentic AI projects.

1.

Rapid Prototyping

AI systems inherit the probabilistic nature of large language models. Their behavior cannot be fully predicted in advance, making iterative experimentation essential. Enterprises should adopt enablement processes that support rapid prototyping, frequent testing, and continuous refinement. Faster iteration allows teams to fail safely, learn quickly, and converge on solutions that align with real business needs rather than theoretical assumptions.

2.

Minimize Time to Value

Rapid experimentation alone is insufficient if prototypes cannot transition smoothly into production. Successful programs reduce friction between development and deployment by avoiding re-engineering, security retrofits, or governance workarounds late in the lifecycle. Ideally, moving from a validated prototype to production should be a controlled, low-effort step that immediately extends value to target users.

3.

Enterprise-Grade Authentication and Authorization

Agentic AI projects frequently bypass enterprise identity systems in the name of speed, resulting in isolated authentication silos or entirely unsecured prototypes. This approach undermines governance and violates zero-trust principles. Integrating agents with approved enterprise identity providers enables continuous authentication and authorization, ensuring agents operate within established access controls and audit frameworks from day one.

4.

Continuous Monitoring and Visibility

As enterprises deploy agentic AI, attackers are increasingly using similar technologies to automate abuse, exfiltrate sensitive data, manipulate prompts, and exploit APIs. Development teams cannot address these risks alone. Security teams must have real-time visibility into agent behavior, usage patterns, and data access in order to detect anomalies, prevent misuse, and reduce unintended expansion of the attack surface.

5.
Enforce Guardrails and Policy Controls

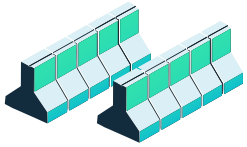

Because agents can act autonomously and at scale, guardrails are essential. Network access controls, rate limiting, and policy enforcement help ensure agents remain within defined operational boundaries. Risk-scoring mechanisms can identify anomalous or dangerous behavior and automatically throttle or halt activity. These controls are especially critical during early experimentation, when agent behavior is least predictable.

6.
Secure and Govern MCP Server Usage

The emergence of the Model Context Protocol (MCP) has simplified agent integration with tools and enterprise data, but it also introduces new risks. Direct access to vendor-provided or public MCP servers reduces enterprise visibility and control, while malicious or poisoned MCP servers create additional attack vectors. In-band monitoring and policy enforcement are required to ensure MCP usage aligns with corporate security standards and data governance requirements.

7.
Insulate Developers from Rapidly Evolving Protocols

MCP and other agent communication protocols are evolving quickly, creating operational risk and technical debt. Rather than embedding protocol logic directly into agents, best practice is to offload protocol handling to a centralized gateway. This abstraction allows protocol updates to be managed in one place, preserving agent functionality and reducing the need for specialized in-house expertise.

8.
Support Flexible Deployment Models

Enterprise environments are rarely uniform. Agentic AI solutions must support multiple deployment models, including public cloud, private cloud, on-premises, and hybrid architectures. Flexibility is necessary not only to meet current business requirements but also to future-proof deployments as infrastructure strategies evolve. Solutions tied to a single deployment model risk obsolescence or abandonment.

9.
Security from the Start

Retrofitting security after deployment has historically proven ineffective. Agentic AI systems must be designed with foundational security controls from inception, based on an understanding of their unique risks. This includes support for human-in-the-loop workflows, recognizing that non-deterministic agents may occasionally behave in unexpected ways. Security architectures should remain modular, allowing defenses to evolve independently of agent logic.

10.
Choose Partners with Deep Domain Expertise

Finally, enterprises should be selective in choosing agentic AI partners. Ideal partners bring deep experience in APIs, application behavior, and enterprise security, along with financial stability and a proven ability to support large-scale deployments. Strong partners help align agentic AI initiatives with business context, regulatory requirements, and long-term operational goals.

The Cequence Advantage

The Cequence AI Gateway enables organizations to unlock the promise of agentic AI productivity by easily and safely connecting agents to enterprise and SaaS applications. It transforms any application or API into a dynamically discoverable MCP-compatible endpoint in minutes, without coding. Go from prototype to production without incurring the technical debt associated with basic solutions that lack core enterprise hosting, authentication, authorization, and monitoring capabilities. Additionally, the AI Gateway integrates with the Cequence UAP platform, offering enhanced protection for enterprise applications and data from malicious agents and users. Visit **www.cequence.ai/products/ai-gateway** to learn more.

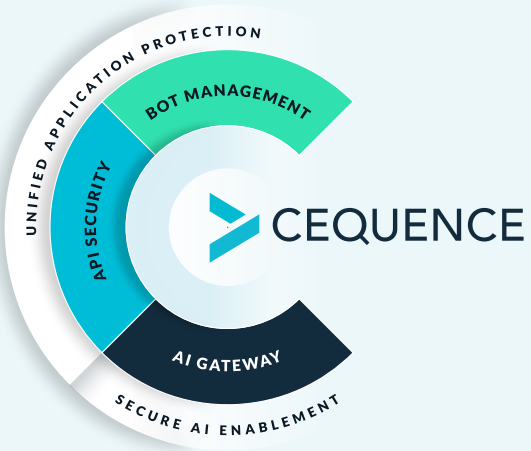

CEQUENCE

5201 Great America Pkwy, Suite 240, Santa Clara, CA 95054 | 1-650-437-6338 | info@cequence.ai | www.cequence.ai

© 2026 Cequence Security, Inc. All rights reserved.