CEQUENCE

Top Considerations for Enterprise Agentic AI Projects

AI has rapidly moved from hype to reality. As 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 protecting 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 with a Short, Smooth Path to Production

Agentic AI systems don't behave predictably. You can't spec your way to a correct implementation. You have to run it, see what works, what doesn't, and iterate, all the way to production. Many agentic AI prototypes sit in a queue waiting on security sign-off, governance review, or operational tooling that wasn't considered early enough. When the prototype is ready, shipping it should be a small step, not a monumental undertaking.

2.

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 both users and agents operate within established access controls and audit frameworks from day one.

3.

Overly-Broad Agent Access

While identity is critical, it only tells you who an entity is – not what privileges are appropriate. Agents, by default, should not have all the privileges that their user possesses. Like human workers, agents need a clear job description that frames their assigned task. Done right, having this plain-English articulation can be effectively used to generate a tailored "Agent Persona" that grants only the tools and permissions it needs to operate. The result is minimized risk, better performance, and lower operational costs.

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 support audits, detect anomalies, prevent misuse, and reduce unintended expansion of the attack surface.

5.

Enforce Guardrails and Policy Controls

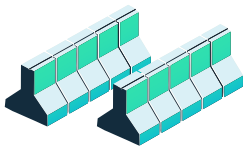

Because agents can act autonomously and at scale, guardrails are essential. Network access controls, rate limiting, and policy enforcement help ensure agents remain within defined operational boundaries. Risk-scoring mechanisms can identify anomalous or dangerous behavior and automatically throttle or 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. Organizations need controls that support access to approved MCP servers and that usage aligns with corporate security standards and data governance requirements.

7.

Protect Sensitive Data

Agentic access to enterprise applications threatens to expose sensitive internal data. Organizations must be able to monitor, redact, and block sensitive data flowing through MCP tool calls, both inbound requests and outbound responses. Ideally, this capability integrates with existing DLP infrastructure and is implemented in a way that is scalable, reliable, and doesn't introduce latency.

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 makes applications agent-ready while securing and controlling agentic AI workflows, enabling organizations to unlock AI-driven productivity and growth. Built-in governance and guardrails constrain agent behavior using capabilities that include least privilege access, rate-limiting, and sensitive data protection. AI Gateway enables organizations to swiftly innovate, going from prototype to production without incurring the technical debt and scalability limitations associated with basic solutions. Visit **www.cequence.ai/products/ai-gateway** to learn more.

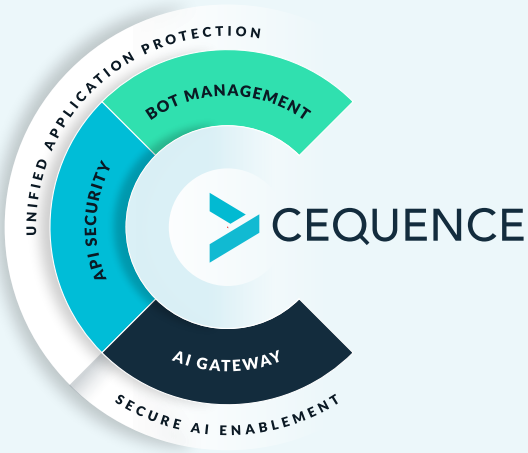

CEQUENCE

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