Is Your Agentic AI Solution “Enterprise-Grade?”

“Agentic AI” is quickly becoming table stakes. What’s rarer (because it’s harder) is enterprise-grade agentic AI.
Enterprise-grade is not a UX flourish. It’s an architectural and operational commitment. For agentic AI, "enterprise grade" is a specific set of capabilities, methods, and mindsets that separate production systems from prototypes.

See some of the most critical points of distinction below and **read our full Evaluation Guide** to learn more.

Deep, Bi-Directional Data Integration

"Integration" means more than read access, what we call ‘Data in.’ Agents must write back, update records, trigger workflows, and create tasks. Read-only access produces insights. ‘Data out’ produces outcomes.

Knowledge Architecture: The Compounding Layer

Raw data integration feeds what matters most: your knowledge graph. A knowledge graph maps entities, relationships, and context and compounds value as the data feeding it grows over time. This context compounds and the system gets smarter, not because the model changes, but because enterprise context grows steadily over time.

Security, Access, and Governance

Enterprise buyers don't ask "can your agent do X?" They ask "who can make your agent do X, and how do we prove they should?" You should have strong policy-driven access control and your agentic system must enforce the proper security policies and governance before the agent acts, not after you discover the violation.

The Audit Imperative

"The agent made a mistake" is not an acceptable incident report. Enterprise-grade systems maintain complete audit trails of every decision, every interaction, every response, and every failure.

Model Operations: The Invisible Infrastructure

Users don't care which model answered their question, they care about accuracy, speed, cost, and reliable answers. Enterprise systems cannot tolerate the same question producing different answers across successive calls. Orchestration isn't about using every model. It's about using the right model for each task, automatically, while staying within operational constraints.

Kana

Access the full **Evaluation Guide here** and learn more at **www.kana.ai**