Enterprise AI vs. Startup AI: Know the Difference

RELIAQUEST

EXPONENT

Not All AI Is Built for the Enterprise SOC

“Agentic-SOC” has flooded the security market.

There’s a new AI Startup promising transformation every day.

80% of so-called agents are glorified workflows.

But many enterprises remain skeptical.

- **What happens when AI is wrong?**
- **Are the agents protecting your data?**
- **How do I know this will work in my environment?**

THE TRUST PROBLEM: Why Enterprises Don’t Trust AI

Broken Promises

Demos impress. Production exposes. Most GenAI tools deliver "prompt fatigue," not results.

Low Accuracy in the Wild

Marketing benchmarks ≠ production reality. 40%+ of Agentic AI projects may be canceled by 2027 (Gartner).

Unproven Across Environments

Validated in one stack or vertical doesn't mean validated in yours. Enterprises average 45+ security tools.

Your Data Funds Their Product

Vendors use customer telemetry to train their models — often buried in terms of service. Always ask: does my data train your product?

Limited Data Normalization

Most competitors 'normalize' data by throwing an LLM at log mappings and hoping. That's not a solution — it's a gamble.

Not Built For Operators

Founder pedigree ≠ production maturity. Thousands of environments — or a handful of early adopters?

5 Questions To Evaluate AI for the Enterprise SOC

1. Who Built it and Why?
2. Can You Trust It at Scale?
3. Does it Operate or Just Assist?
4. How is the AI Secured and Validated?
5. Does it Bring Real Security Context?

Who Built It – and Why?

Operational DNA

Built by teams who've triaged millions of real alerts — not just studied them.

Mission-Driven, Not Hype-Driven

Did the company exist before the AI boom — or because of it?

Designed for Workflows, Not Demos

Shaped by 15+ years of SOC operations across 1,300+ enterprise environments.

If it wasn’t built by people who’ve run a SOC, it’s going to struggle in one.

Can You Trust It at Scale?

In the enterprise, trust is the foundation of any AI-driven decision.

Teams need to know that what they’re relying on has been tested, validated and proven under real business conditions.

Many startups have refocused at the SME market because they struggle at scale.

7 Tests for AI Validation

Prompt guardrails, LLM as Judge, Golden Datasets, Sampling, Expert Validation, Crowdsourcing QA (+Transparency)

Deployed over 1300 Environments

74 Million alerts annually, across 250+ technologies, various industry and attacks

Accuracy Proven in Production

Better than human precision, 11,000 Agentic Memories, thousands of use-cases

Does It Operate – or Just Assist?

Agents

- Little or unknown intent
- No shared memory or operational state
- Single action execution and visibility
- Achieve simple tasks
- Adds cognitive load

Siloed agents, operating independently without shared context or unified goals.

Agentic Systems

- Clearly outlined objective and intent
- Leverage memory for context across environment
- Multi-coordinated actions executed autonomously
- Can take on job roles
- Reduces workload and decision fatigue

Architected for objectives, coordinating multiple agents and AI components toward shared goals.

How is the AI Secured and Validated?

Continuous red teaming against AI attacks including prompt injection.

Prompts tailored to specific models with testing and validation built in.

Adapts as models improve, change or fail.

New models evaluated and integrated without re-architecting

Does It Bring Real Security Context?

Standalone AI is just a capability.

Effective AI grabs context from telemetry sources, uses security operations platform capabilities, and leverages real-time threat research.

Most startups only address a subset of the available context.

Grabs Context from Telemetry

AI needs full environment context to accurately reason to build or tune detections and respond to threats

Connects to Native SecOps Capabilities

AI should pull context from solutions like CAASM, Dark-web Monitoring, Exposure Management.

Boots On The Ground Threat Intel

Enterprises require real-time, actionable, and rich threat intel to drive AI decisions.

The ReliaQuest Difference: Flashy Demos vs. Proven Operations

What We Are Seeing in the Market

- Many "agents" are workflow executors
- Optimized for demo conditions and limited environments, not enterprise complexity
- Data normalization via "throw an LLM at it and hope”
- Security of AI outputs is on the customer
- Vendors claim model-agnostic; customers left to select and validate
- Vendor #1 priority: get acquired
- Requires your data to train and build its product

GreyMatter

- Six AI personas with 200+ skills and 400+ tools
- Deployed across 1,300+ environments with diverse industries and tech stacks
- Patented universal data normalization across 250+ technologies
- Seven-phase AI testing and validation lifecycle with prompt guardrails
- 15+ years of mission-driven enterprise SOC operations
- Doesn't train on your data. Apply context in real time.

Key Takeaways

Majority of 'AI agents' are marketing.
Demand proof of true agentic operation.

AI without data normalization is AI that guesses.
Demand a real data layer.

Vendors who push securing AI to the customer are adding risk, not solving for it.

Company priorities matter
Building for you, or building to exit?

RELIAQUEST

EXPONENT