Bricklayer AI

AI IS NOT BUILT FOR CYBERSECURITY:

A GUIDE TO ENSURE YOUR AI INVESTMENTS ARE BUILT FOR LONG-TERM SUCCESS

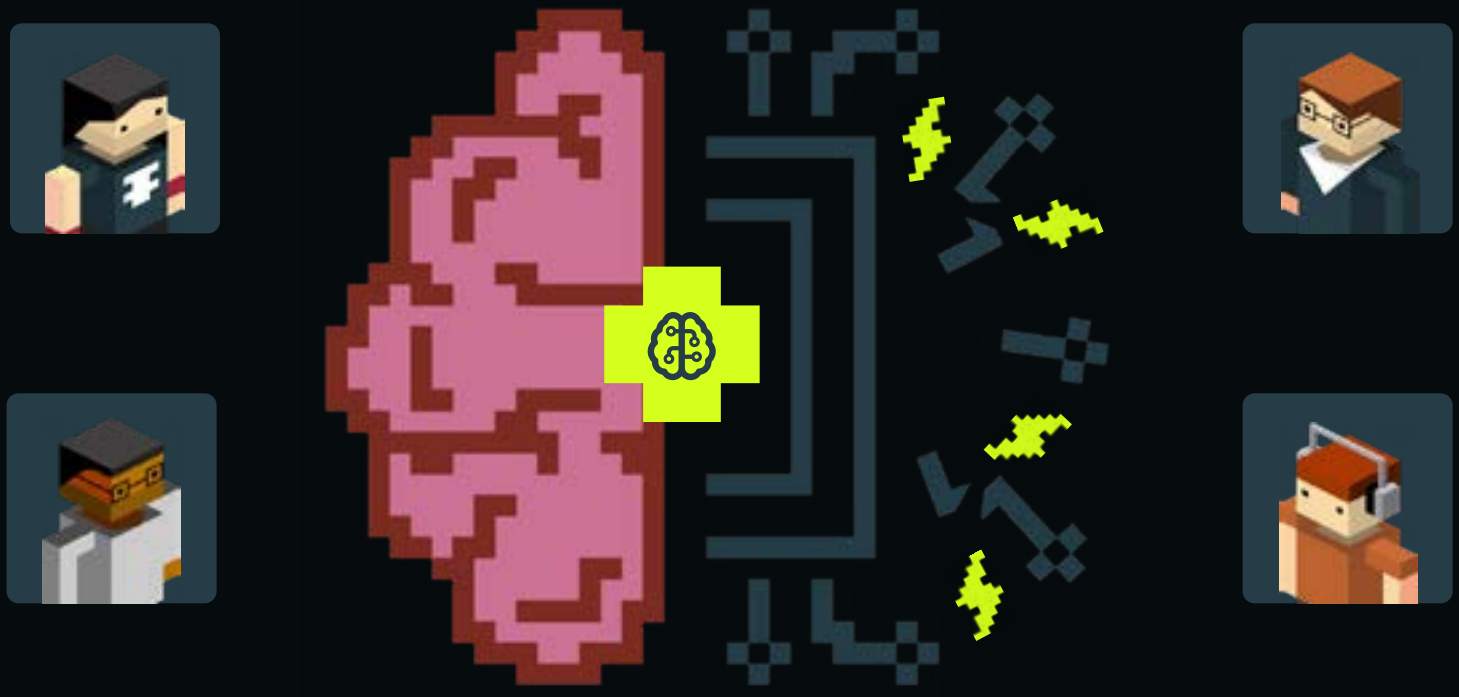

AI. BOTH THE GREATEST OPPORTUNITY AND THE BIGGEST HYPE IN THE BUSINESS WORLD TODAY.

Yet, the vast majority of enterprise AI projects fail (MIT Nanda study), but not because there isn't a need. It's because AI has major issues in most enterprise settings.

Rigidity
Automation that snaps the second the landscape shifts.

Indeterminism
Nobody can explain why the AI said yes & no.

Inaccuracy
AI models only know what they're taught - limited or inaccurate data has major implications.

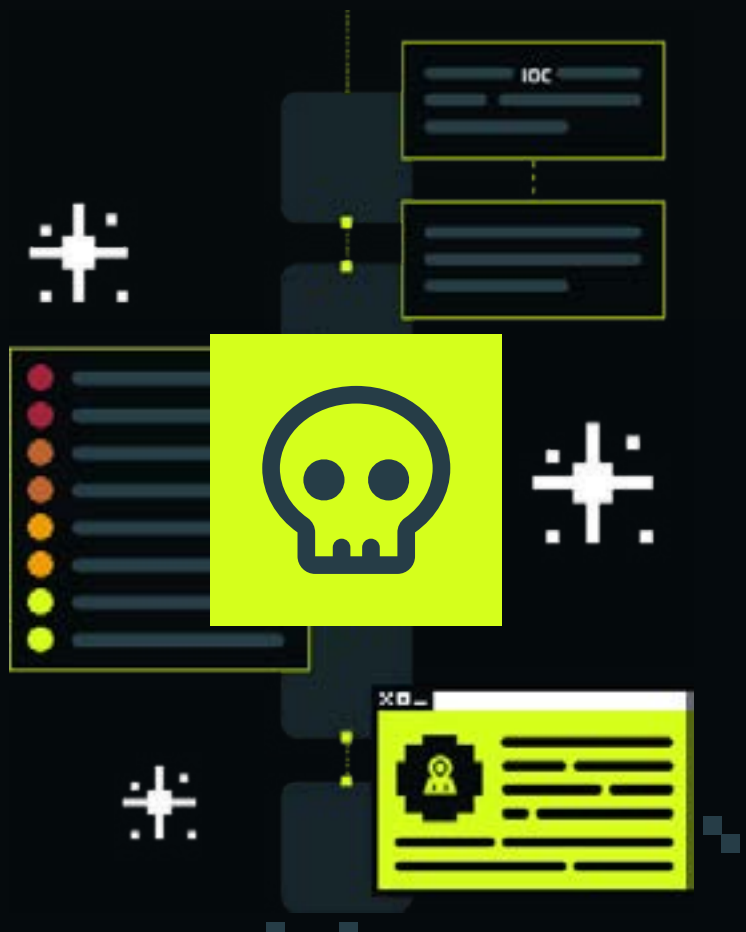

In most industries, these failures are inconvenient. In cybersecurity, they're catastrophic.

SOCS ON FIRE:

A State of the Union

People are burning out.

Security Operations Centers are becoming fatigue factories.
Hire more analysts? Sounds good, but doesn't solve the root problems:

Human Inconsistency: Ten analysts, ten wildly different responses. The more people you add, the more variance you introduce.

Team Burnout: Repetitive, monotonous work drains people's focus and energy.

Hidden Cost: Talent shortage and high turnover lead to increased hiring costs and loss of institutional knowledge.

AI Agents Aren't a Catchall

Drop in a single AI agent and you will run into a different set of problems:

Hallucination: AI systems can produce incorrect or fabricated answers—often stated with total confidence.

Black Box Syndrome: AI makes decisions without explaining how it did it, meaning you can't back it up to your board or recreate it on your own.

Lack of Context: Single AI agents often don't have the context or communication capacity to gather all of the information needed to make a decision.

Legacy Tools are Ancient Artifacts

Security Orchestration, Automation, and Response (SOAR) tools were sold as a savior - automation across your entire cybersecurity team. Save time, less burnout.

But, the reality is, SOAR is flexible, but extremely brittle.

As soon as an attacker sidesteps the script (a frequent outcome), a SOAR playbook doesn't work anymore.

Today, SOCs run on burnout and broken automation. But there's a better way.

Building for
THE LONG GAME

Forget Lone Agents. Look for an Agentic Cybersecurity Platform.

Right now, "agent" is being used loosely to describe everything from LLM-based chat tools to scripted automations. Often these are wrappers around foundation models, or static bots tuned to execute a narrow function.

One agent acting alone is like a team member who follows directions to a tee - but can't deviate from those directions at all.

They may assist an analyst or automate a task, but they don't collaborate, adapt, or scale beyond their single job. They operate in silos – lacking memory, shared context, and extensibility.

Bricklayer was built differently. It's an Agentic Cybersecurity Platform built from the ground up to build, deploy, manage, and scale AI agents across the entire SOC. Bricklayer offers:

Agent Building:
Persistent systems for defining, launching, tracking, and evolving agents, customized for each organization's needs.

Collaborative Workflows:
Tasks are routed, escalated, and continuously improved. Agents learn and evolve inside every customer environment.

Workbench UI:
A dedicated interface for humans to collaborate with, guide, and review agent decisions.

Shared Memory & Context:
Agents retain institutional knowledge and share it, working together as a coordinated, evolving team.

Performance Management:
Effectiveness and ROI are continuously tracked, highlighting improvements and new opportunities.

Integrated Tooling:
Agents connect directly to tools, data sources, and response systems—so they don't just observe, they act.

Governance & Security:
Every action is policy-bound, monitored, and auditable.

Because it's a platform, Bricklayer doesn't just automate today's workflows – it adapts to tomorrow's challenges.

Modularity Means Transparency

A black-box answer might sound convincing, but it collapses under scrutiny. That's why Bricklayer was built with modular components.

Every objective is broken down into discrete tasks. A task is an atomic unit of work: a prompt, an agent, and one or more inputs. Subtasks can be spawned by agents acting as managers for the work involved within a single task.

In Bricklayer, every objective begins with a goal - not just a fixed script. AI agents dynamically break that goal into discrete, manageable tasks, assigning roles and collaborating based on context, expertise, and constraints.

This mirrors how high-performing human teams operate: smart people come together, assess what's needed, and self-organize to get the job done - within the structure of shared tools, data, and objectives.

Suddenly, "AI said so" becomes "Here's how we got here." CISOs can point to evidence, your board can follow the logic, and analysts can finally trust automation that shows its work.

Embrace Multi-Agent Context Engineering (MACE)

Enterprises are racing to deploy AI into complex investigative domains such as security operations, fraud detection, and compliance. Specialized agents are emerging - SOC Analyst Agents, Threat Intel Agents, Incident Response Agents – each capable of executing role-based tasks.

But there's a gap, and context sharing is the solution.

These agents don't naturally share context. One agent might collect evidence, another might enrich it, and a third might summarize findings - yet without a container to carry context forward, the process is fragile. Investigations lose continuity, teams lose trust, and results are inconsistent.

Because investigations are uncertain, iterative, collaborative, accountable, and dynamic, they cannot be reduced to rigid scripts.

Traditional automation frameworks were built for repeatable transactions; investigations require a framework that engineers and preserves context.

Multi-Agent Context Engineering provides the missing discipline. It treats procedures not as simple workflows, but as context engineering containers.

HOW IT WORKS

MACE defines a lifecycle:

01 INPUT CONTEXT SHAPING

Normalize raw alerts or signals into usable investigative inputs.

02 CONTEXT BUILDING

Layer evidence, reasoning, and annotations as agents act.

03 CONTEXT HANDOFF

Ensure each agent inherits not just data but meaning.

04 CONTEXT RESOLUTION

Guide branching and decisions with a clear rationale.

05 CONTEXTUAL OUTPUT

Produce auditable, reusable artifacts of the investigation.

MACE distinguishes between different kinds of context, each with its own flow:

- Evidentiary Context — facts and artifacts collected.
- Investigative Context — reasoning, hypotheses, and interpretations.
- Procedural Context — what steps were taken and in what order.
- Decision Context — logic behind escalations or remediations.

By engineering these context flows, MACE creates procedures that behave less like scripts and more like investigative scaffolds – capturing and amplifying meaning at every step. Every handoff builds meaning rather than loses it, every decision leaves an auditable trail, and every investigation contributes to institutional memory.

Dynamic Reporting > Static Dashboards

AI-first cybersecurity teams need more than just tasks checked off of a list. They need full transparency into what happened, why decisions were made, and where context came from.

With Bricklayer, analysts no longer need to sift through hundreds of log lines, relationships, and fields. Every Bricklayer agent outcome includes an automatically generated Agent Debrief: a clear, single-page, multi-card summary that includes what a human needs to know – but nothing else.

Static dashboards are the graveyard of security insight: dead on arrival, outdated the moment they land.

Agent Debriefs are dynamically built at runtime by the Agents themselves – tailored to each procedure type and even the unique data within it.

The outcome distills everything down to what the Agent believes is most relevant to make a decision or take action. This may include things like attack timelines, text-based reporting, data visualizations, code snippets, related events, and much more.

The result is a purpose-built, human-consumable report ready for analystsand executives alike.

Score.

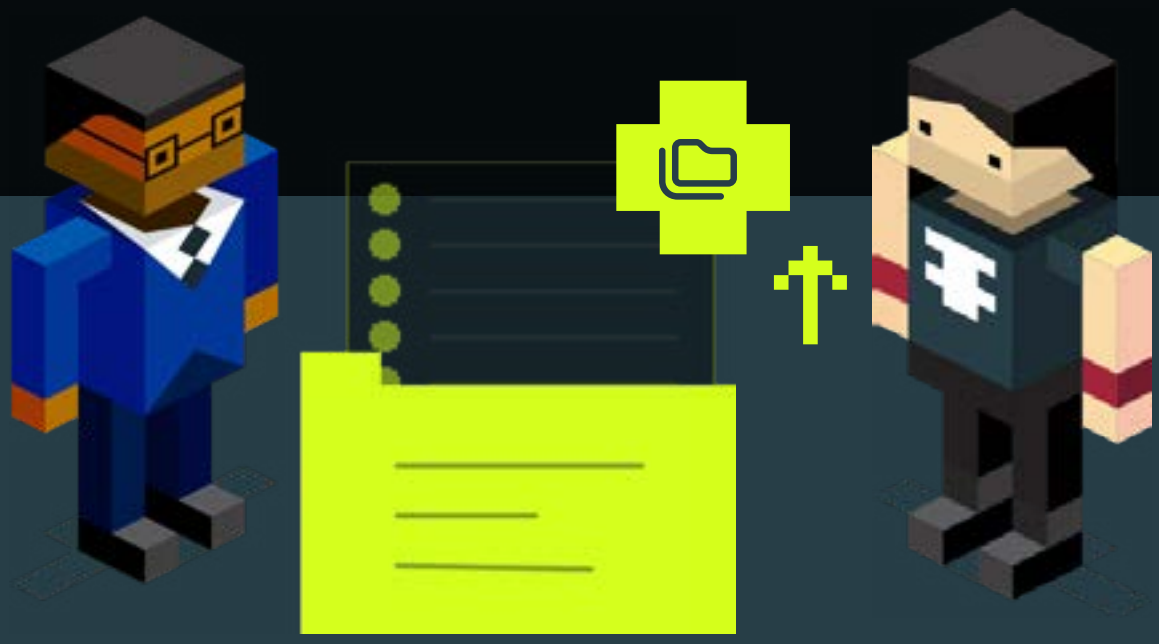

Measure Coverage & Consistency, Not Just Speed

When people talk about applying AI in the Security Operations Center (SOC), the conversation almost always starts with speed. The common comparison goes something like this: "What takes a human analyst 45 minutes, an AI agent can do in 2 minutes and 45 seconds."

That's true, and speed does matter. But it's hardly the most valuable metric. If speed is all we measure, we miss the bigger transformation AI is bringing to the SOC.

The real value lies in **coverage** - ensuring every alert is investigated thoroughly – and **consistency** – ensuring every investigation follows the same reliable process.

Metrics like MTTD and MTTR will always matter, but they aren't sufficient on their own.
They measure speed, not quality.
To capture the true value of AI in the SOC, we need to also measure:

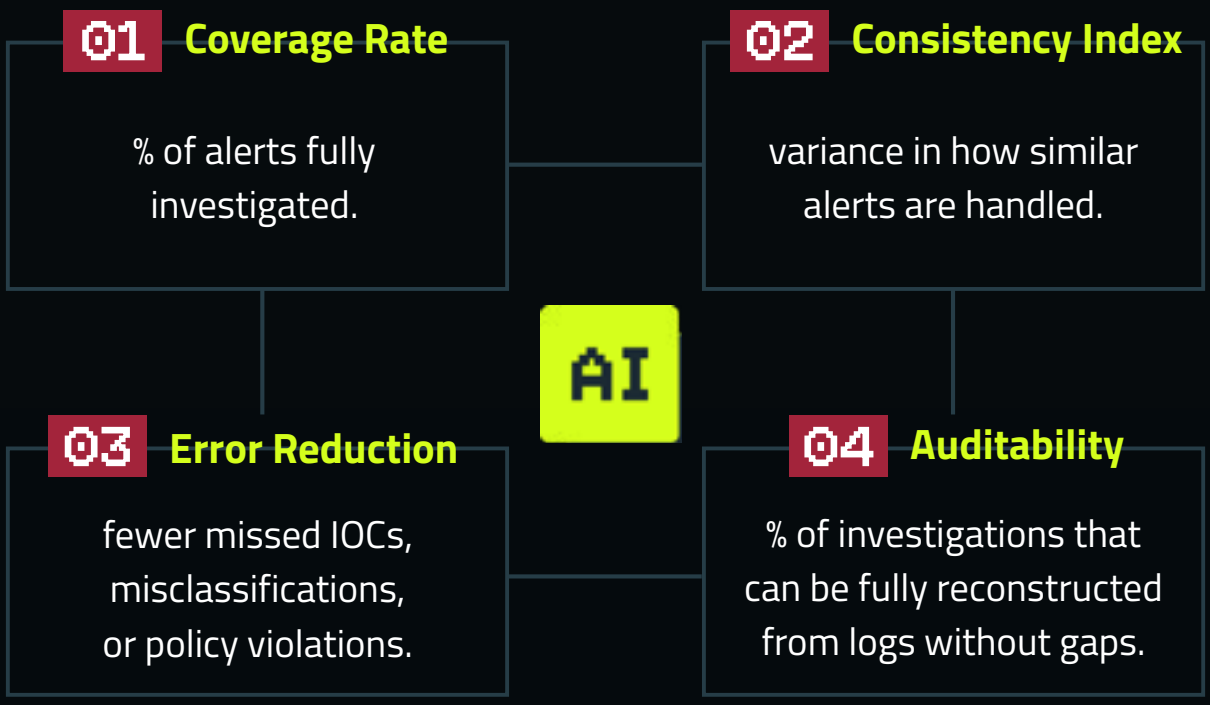

These metrics reflect outcomes that directly reduce organizational risk – outcomes speed alone cannot guarantee.

THE CISO'S AI EVALUATION GUIDE

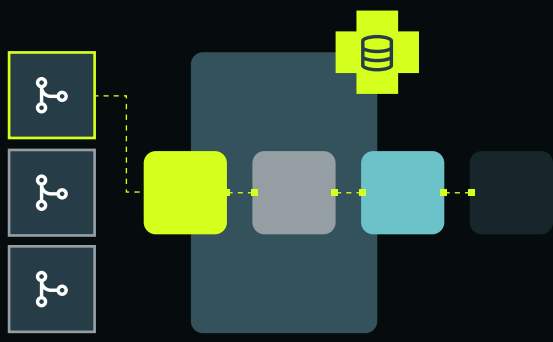

When vendors DM their "AI-powered" pitch, here's how to torch-test it:

01

Can you build/create new agents if you see a gap, allowing you to customize the collaboration to align with your business requirements?

02

Do AI agents collaborate? Do they communicate, delegate, and manage tasks amongst themselves?

03

Is each agent role flexible, meaning that they can act as both a manager and a contributor, just as their human teammates can?

04

Do agents retain and share knowledge with one another?

05

Are complex tasks broken down into tangible, visible steps done by different agents?

06

Are each of these steps traceable and reportable?

07

Can agents modify how they complete tasks based on what they've learned?

08

Do agents gather evidence to make correct decisions like humans do - do they use logic, loops, backtracking, and pattern recognition to ensure they take the correct action?

09

Do agent workflows have flexibility not only in how they are executed, but also in their outcome (for example, investigations may conclude with a confirmed incident, false positive, investigation, ongoing monitoring, etc)?

10

Do agents provide you 'reports' or information that you need in a concise, digestible way?

11

Are agent reports dynamic enough to support ever-changing threat environments and internal processes?

FROM HYPE TO REALITY

The AI hype machine will tell you the future is here. But hype doesn't stop breaches. Bricklayer does.

Bricklayer is an Agentic Cybersecurity Platform built specifically for cybersecurity teams. Every action - whether an agent running enrichment, a plugin tapping third-party tools, or a report for leadership - is part of a transparent, collaborative system. Not a black box.

At the core is our patent-pending model, engineered for today's SOC:

Bricklayer Patent Pending Model

Agent-to-Agent Collaboration:
Enables agents to communicate, delegate, and manage tasks among themselves.

Dual-Role Architecture:
Each agent is role-flexible, with the ability to act as both a contributor and a manager.

Context-Aware Execution:
Agents retain and share knowledge, adapting to real-time needs.

Autonomous Procedure Execution:
Complex tasks are broken into steps and completed across agents with full tracking.

Built for Enterprise Security:
Designed specifically to meet the needs of modern SOC environments.

Want to experience it for yourself?

Bricklayer helps you become an AI-first SOC. See it for yourself.

REQUEST A DEMO
bricklayer.ai/book-a-demo