

KISS.

Runtime fidelity for every agent action.

We execute AI agents' high-impact actions for them — and verify each one is **correct against your own systems of record** before it runs. We close the gap in agent execution at runtime, against your business logic.

THE BLIND SPOT

Access is secured. The *action* isn't.

IDENTITY & ACCESS

Proves the agent is *allowed*

You've done the work: no standing credentials, least privilege, just-in-time access. Authorization confirms the right agent is acting.

But a **fully-authorized** agent can be manipulated — by prompt injection or a poisoned input — into a perfectly in-scope action with the **wrong details**. Credentials never leak. The audit log is clean. The damage is done.

Static policy says: **ALLOWED**

KISS ACTION FIDELITY

Proves the action is the *right one*

Before we execute, we check the action's **actual values** against the trusted record that justifies it — the order, the ticket, the PR. A deterministic check, no LLM guessing.


Match → execute and log. **Mismatch → block or escalate to a human**. Because we hold the credentials and the action runs through us, this is enforced inline, before it happens.

Record check says: **WRONG → BLOCK**

THE FAILURE MODE THAT SCALES

A valid operation with *corrupted values*

Each one passes authentication, authorization, and policy. Only a check against your business record says it's wrong.



WHAT MAKES KISS DIFFERENT

Right agent. Right action. *Verified before it runs.*

01

Agents request — they don't act

Agents hold no credentials. We hold them and execute the action through managed integrations to Stripe, Jira, Salesforce, Workday, GitHub, and more.

- No standing secrets in the agent
- One control point across every system
- The action's only path is through us

02

Verified against your record

The action's parameters must match the independent business record that justifies it. The allowed scope is computed per request from the task — not a fixed config.

- Deterministic check, no LLM guessing
- Anchored to a record the agent can't forge
- No clean anchor → step up to a human

03

Enforced before it happens

Mismatch blocks or escalates; a match executes and logs. It's inline enforcement at runtime — not an alert after the money has already moved.

- Inline, not after-the-fact alerts
- Fail safe — never a silent allow
- Near-zero false positives, trusted to block

Identity tools make sure the **right agent** is acting. We make sure it's doing the *right thing* — verified against your own records, enforced before the **action runs**.

The agent was fooled. The control plane wasn't.

Talk to us → dev@kiss-security.com