Saviynt

WHY AI AGENTS NEED THREE GATEWAYS

Two checks passed. One question left unanswered.

Securing an AI agent means answering three different questions for every action it takes. Each gateway answers one, and none can answer for the others.

Gate 1 : Content Security

Is this content safe?

Catches prompt injection, redacts sensitive data, filters harmful output.

Pass

Gate 2 : MCP and LLM Routing

Can this agent reach this tool?

Approves MCP servers, enforces rate limits, routes work to the right model.

Pass

Two green checks. Zero answers about what happens next.

Customer Financial Records

Gate 3 : Runtime Authorization

Is this agent authorized to do this, right now, for this purpose?

Weighs identity, intent and context before the action runs. Continuously, not just at login.

Identity | Intent | Context | Held: Outside the task

Why two out of three fails

A customer support agent has a valid session, a clean prompt and a sanctioned connection. Its task is to update one customer record. Then it asks to export every customer's payment and billing history.

Why two out of three fails	
Content check	Pass
Connection check	Pass
Fits the task?	Never asked
Outcome	Financial records exported

All three gates	
Content check	Pass
Connection check	Pass
Fits the task?	No
Outcome	Action held

Content protects the conversation. Connection protects the route.

Authorization protects the action.

Saviynt Agent Access Gateway, part of Zuma Access

Learn More →