

CircuitProof | Automation must earn reality

Executive Proof: AI Agent Exfiltration Proof

Chat, voice, image, document, MCP/tool calls, vector stores, approvals, eval cases, and prompt injection defense.

Board decision: Approve only after owner actions, rollback, connector scope, and residual findings are accepted.

Findings: 1

Simulation throughput: 600/hour

Stress readiness: 100%

Assurance score: 81% assurance evidence readiness

Before: The user requested automation for AI Agent Exfiltration Proof: Chat, voice, image, document, MCP/tool calls, vector stores, approvals, eval cases, and prompt injection defense.

Caught: CircuitProof checked validation, simulation, stress, connectors, and policy boundaries. It caught: No critical failure escaped.

Repaired: CircuitProof links each weak point to a proof gate, repair action, owner decision, and rollback check before execution.

Approved: Release remains approval-gated at 100% readiness with 1 validation findings and 7 simulated items processed.

Exported: The package exports executive summary, proof matrix, failure table, permission map, compliance evidence, rollback check, and human approval.

Human approval is required before real-world execution.

This is assurance evidence, not certification or permission to execute.