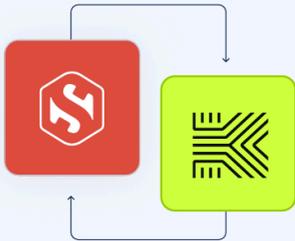


MFA for AI

Prevent unauthorized AI actions – even with valid tokens



A valid token should not be enough for an AI agent to act.

Today, if an attacker gets an agent's credentials, they can execute actions through your own infrastructure. The system cannot tell the difference between a real agent and an attacker using a valid token.

MFA for AI prevents token reuse from untrusted environments. Every action requires both a verified agent and a trusted machine.

MFA for AI enforces that every action comes from a verified agent running on a verified environment.

The system enforces a dual check before execution:

Identity Plane (Keycard): Verifies who the agent is and what it is allowed to do. Each agent gets a unique identity and scoped permissions.

Trust Plane (Smallstep): Verifies where the agent is running using hardware-bound credentials. Only trusted devices or workloads can act.

Enforcement Plane: Blocks execution unless both identity and environment pass. Every decision is logged for audit.

KEY FEATURES

Dual Identity Proof

Verify both the agent and the machine it runs on. Both must pass.

Eliminate Static Secrets

Replace tokens and API keys with short-lived, hardware-bound credentials that cannot be reused.

Task-Level Authorization

Limit what each agent can do for each task.

Frictionless Adoption

No additional steps for developers. Policies are enforced automatically.

Standards & Scale

Built on OAuth, OIDC, and ACME device attestation. Works across environments.

BUSINESS BENEFITS

Reduce Risk Without Slowing Innovation: A stolen token cannot be used to execute actions from an untrusted machine.

Unified Governance & Compliance: Every action is tied to an agent identity, a specific action, and a verified environment.

Future-Proof Security: Works with existing identity standards without adding new secrets or custom logic.

Without MFA for AI: A stolen token can execute actions.

With MFA for AI: Execution can be restricted to requests where both agent identity and environment are verified.

USE CASES

Secure Data Access: Agents query sensitive data without exposing credentials. Access is limited by policy, and execution is allowed only from trusted environments.

Agentic DevOps & CI/CD: Only agents running on verified infrastructure can build, deploy, or push code.

Internal Automation & Bots: Every action traces back to a known agent and a trusted machine.

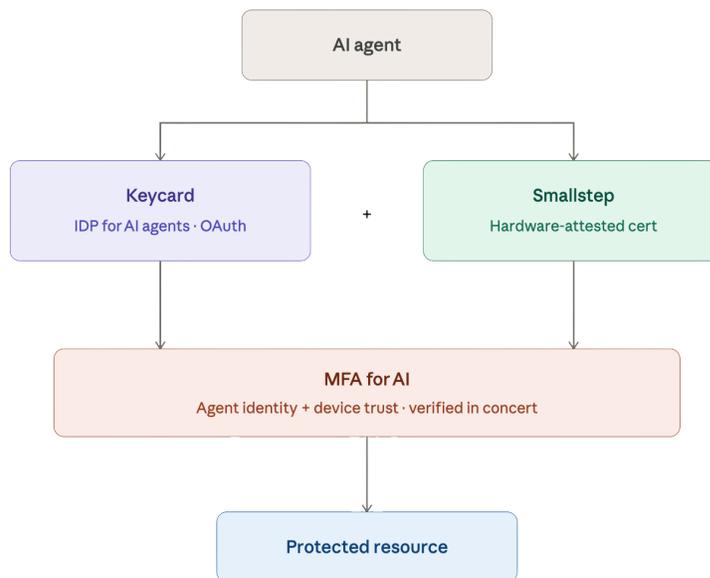
ADOPTION & IMPLEMENTATION GUIDANCE

Inventory & Scope: Inventory agents and the systems they can act on. Prioritize high-risk workflows.

Onboard Agents to Keycard: Register each agent with Keycard. Define least-privilege scopes and expiration per task.

Enroll Infrastructure with Smallstep: Enroll the machines running agents. Issue hardware-bound credentials to devices, VMs, and containers.

Enforce & Monitor: Enforce that both agent and machine must pass before execution. Log and monitor all decisions.



Smallstep + Keycard

With **MFA for AI**, **Keycard** authenticates the agent and its permissions, while **Smallstep** attests the device or workload. Together they provide high-assurance identity and enforceable control for every agentic action.