



AI Governance

Adopt AI agents across your organization with the visibility and control to prove they are safe.

AI agents are already in your SDLC. 84% of new code is AI-generated and 79% of companies run agents across roles. Every other system in production has an audit trail, access controls, and boundaries. Your agents do not.

And it is never one agent. Claude Code, Cursor, Copilot, Codex, whatever ships next quarter, each with its own closed-source sandbox that changes overnight. Blocking agents stalls delivery. Ungoverned agents create exposure. Docker AI Governance is the third path: adopt fast with controls that hold from day one.

Docker spent a decade making infrastructure trustworthy for 20M+ developers. Isolation is the job we already do, and we own the primitives. Docker AI Governance puts one sandbox standard under every harness across three surfaces: filesystem, network, and MCP. Vendor-neutral across any model, agent, and tool.

What's in it

01

MicroVM sandbox for every session

Each session runs in its own microVM with its own kernel, a hard boundary to the host that is gone when the task ends. Network and filesystem isolation enforced at the runtime. No bypass path.

02

Org-wide sandbox policy

One policy for all employees, applied across the org from the admin console.

03

Team policies and audit logging

RBAC via SAML and SCIM.
Structured audit events for every evaluation, streamed to your SIEM.

04

MCP Gateway and policies

Managed catalog. Allow or deny by server, tool, and transport.
Default-deny on unvetted servers.

05

End-to-end sandbox and MCP

One console, one policy model.
Sandbox and MCP are governed in one place.

Docker governs the full agent lifecycle, tools, actions, and infrastructure, in one backbone from local to cloud.

Your security boundary should not live inside a tool that rewrites itself every two weeks.
Govern the machine once, not ten harnesses.

