HQ FIELD GUIDE

Give more people and agents access to AI. Keep control of the company.

Manage who can use which knowledge, skills, tools, credentials, projects, agents, and actions without reducing AI adoption to a policy document.

Manage who can use which knowledge, skills, tools, credentials, projects, agents, and actions without reducing AI adoption to a policy document.

01 For it, security, and operations leaders

02 Broader AI adoption with clearer company boundaries, role-appropriate access, accountable ownership, and review where required.

THE CURRENT COST

The policy says one thing. The workflow is configured somewhere else.

As AI adoption grows, every individual setup becomes another place where context, credentials, customer data, and authority can drift away from company intent.

01 Access becomes all or nothing: Teams either withhold useful context or share too broadly because the operating system does not reflect roles and groups.

02 Credentials enter the conversation: Keys and tokens are copied into prompts, files, or messages when governed tool access is not built into the workflow.

03 Approval happens after the action: Sensitive communication or system changes reach people only after an agent or user has already acted.

THE HQ WORKFLOW

Put governance inside the company workflow.

HQ connects membership, groups, knowledge, projects, skills, tools, secrets, agents, and review boundaries so useful access can expand deliberately.

01 Define the company boundary: Separate companies, clients, or brands and assign people and agents to the right membership and groups.

02 Grant the minimum useful capability: Provide only the knowledge, project, skill, tool, credential, and action scope required for the role.

03 Test access and approval: Verify the authorized workflow succeeds, the unauthorized path fails, and sensitive output reaches the right reviewer.

PROOF AND EXAMPLE

Governance is proven by access behavior, not policy language.

A useful test shows that the right person or agent can complete the job, the wrong one cannot reach the protected context, and the sensitive action still reaches its accountable reviewer.

01 Membership and groups: People and agents belong to the correct company and role-based access group.

02 Scoped capability: Knowledge, projects, skills, tools, and secrets are granted for a specific workflow.

03 Approval boundary: Draft-only, review-required, and escalation behavior are explicit where the business needs them.

Examples are anonymized. Named customer claims and proprietary artifacts require recorded approval before public reuse.

START SMALL

Test one real boundary from both sides.

Start with a single knowledge area and capability, then prove that company access behaves the way the policy intends.

01 Create one group for a real team or role

02 Grant one knowledge area plus one tool or skill

03 Verify authorized use, denied use, and the required review step

TRY HQ

Make useful access the default and excess access the exception.

Start with one group, one shared capability, and one access test that proves the company can move faster without giving up its boundaries.

01 Give teams useful access by role

02 Reduce raw credential sharing

03 Keep accountable people in the approval path

Start free with HQ

https://www.hqforwork.com/welcome?plan=free&utm_source=sales_pdf&utm_medium=pdf&utm_campaign=hq_govern_ai_access_knowledge_and_approvals_v1