

AI GOVERNANCE AND AGENTIC SECURITY

Define what your AI agents can access and do.

Map the agents and systems in scope, implement the agreed policies and approvals, and review the results against a signed test set.

-
- 01 DECIDE
Choose the question
 - 02 REDESIGN
Change the workflow
 - 03 GOVERN
This engagement
 - 04 SCALE
Expand, revise, or stop

01 THE DECISION

When agents can access sensitive data or change business systems, written policies need a clear implementation. Your team needs to know which actions are permitted, which need review, and what evidence remains.

We map the agents and systems in scope, agree access and approval rules, and assess whether your existing controls can enforce them. We implement the supported controls, define logging and retention requirements, and document gaps that need a different approach.

You receive a trust-boundary map, implemented controls in your environment, and the AI Governance Pilot Measurement Plan with observed results. Your team can review what passed, what failed, and what to expand, constrain, improve, or stop.

Bring H2H in when

- Your security or technology leaders will not approve broader AI use until data access, actions, and logging are clearly limited.
- AI agents may reach sensitive systems, change records, call tools, or take other consequential actions.
- Your policies do not yet specify how agent actions are checked, reviewed, or recorded.
- An AI workflow needs customer-managed control or formal human approval before it can be adopted.

What we need from you

- A security lead who can approve policies and the limits on what each agent may do.
- The owners of the agents, identities, and systems in scope, available for working sessions.
- An independent reviewer on your side to approve the test set and review results, plus read access to agent configurations under accounts you issue.

AT A GLANCE

DECISION

What can our agents see, decide, and do in production?

PHASE

03 · Govern

DELIVERABLES

Trust-boundary map · policy and approval design · implemented controls · AI Governance Pilot Measurement Plan

ENDS WITH

Expand, constrain, improve, or stop the control path.

OWNER ON YOUR SIDE

Your security lead, who approves data, tool, and action limits

SECURITY SIGN-OFF

Policies approved before any control runs against production agents; results signed by an independent reviewer

DELIVERED BY

H2H Technology (vendor), with your team (customer)

02 HOW IT RUNS

Three steps. A checkpoint after each one where you can stop with something in hand.

01 BOUND Map what each agent can touch

What happens	For each agent in scope, document allowed data and tools, actions requiring approval, logging requirements, and the owner of each decision.
Your involvement	Your security lead and the agent or platform owners in working sessions; read access to agent configurations and identities.
Checkpoint	Trust-boundary map signed by your security lead. You can stop here and keep the map.
Output	Trust-boundary map: agents, data, tools, actions, approvers, and logging requirements.

02 ENFORCE Implement the agreed policy checks

What happens	Implement supported rules for the connected systems: allow, warn, require review, or block. Configure the agreed decision records, access, and retention; document unsupported requirements.
Your involvement	Your security lead approves the policies; identity and system owners issue the accounts and keys the control path needs.
Checkpoint	Policy review with your security lead before any control runs against production agents.
Output	Policy and approval design; agreed controls running in your environment, under your accounts, with known gaps documented.

03 EVIDENCE Test the controls against a signed test set

What happens	Run the signed test set, compare expected and observed outcomes, count unapproved high-risk actions, and check the required records for gaps.
Your involvement	An independent reviewer on your side signs the test set and the results; your security lead attends the results review.
Checkpoint	Results review: expand, constrain, improve, or stop.
Output	AI Governance Pilot Measurement Plan with the observed results.

03 WHAT YOU GET

Four deliverables, each with an acceptance test you can apply.

DELIVERABLE	WHAT IT CONTAINS	FORMAT	ACCEPTED WHEN
1 Trust-boundary map	Every agent in scope: what it may see, which tools it may call, which actions need approval and by whom, and what is logged.	PDF plus editable diagrams	Your security lead signs the map.
2 Policy and approval design	Supported allow, warn, review, and block rules; approvers for consequential actions; exception and escalation handling.	PDF plus policy definitions in the format your controls use	Your security lead approves the policies.
3 Implemented control path	Policy checks, approvals, and agreed logging for the connected agents and systems in scope, using existing controls where they meet the requirements.	Deployed in your environment, under your accounts, with a runbook	Every rule is exercised at least once by the signed test set.
4 AI Governance Pilot Measurement Plan	The signed test set, expected and observed outcomes, unapproved high-risk action count, record completeness, and known gaps, with a recommendation based on those findings.	PDF plus data export	The independent reviewer signs the results.

04 HOW IT ENDS

Expand, constrain, improve, or stop the control path.

Expand

The controls met the agreed test criteria. Assess new agents, tools, or data before extending coverage.

Constrain or improve

An agent reached further than it should, or a control slowed the work. Narrow the agent or tune the rule, and test again.

Stop

The risk outweighs the value in this workflow. Take the agent out of the path and keep the map and test set.

Not included

- A certification or compliance determination. We document the control mapping and evidence for your reviewers.
- A general risk-reduction percentage. The results show what the test set exercised, nothing more.
- A guarantee against loss or misuse, and agents outside the scope named at kickoff.

Exact scope, participants, timing, and fees are confirmed in the statement of work. Changes after kickoff go through a written change note before work continues.

05 HOW WE WORK INSIDE YOUR ENVIRONMENT

What your security and procurement teams will ask.

Access and data

- A mutual NDA precedes discovery.
- You keep your own data.

Controls and frameworks

For each agent in scope, we document data access, permitted tools and actions, human approvals, and logging requirements. Each implemented policy links to its requirement and test evidence. Record access, retention, and any coverage gaps remain explicit. Where relevant, we map the agreed controls and evidence to NIST AI RMF, ISO/IEC 42001, and OWASP guidance for LLM and agentic applications. H2H does not certify compliance or guarantee an audit outcome.

Who does what

ROLE	BOUND	ENFORCE	EVIDENCE
Security lead (customer)	Accountable: trust-boundary map	Accountable: approves policies	Accountable: expand, constrain, improve, or stop
Agent or platform owners (customer)	Consulted: what agents touch	Integrate the control path	Run the agents under test
Identity and system owners (customer)	Provide read access	Issue accounts and keys	Confirm the trace is complete
Independent reviewer (customer)	Informed	Signs the test set	Signs the results
H2H (vendor)	Responsible: boundary map	Responsible: control path	Responsible: measurement plan and results

Assumptions

- A defined set of agents and systems in scope. Adding agents is a written change note or a second engagement.
- Your identity and access systems can issue the accounts and keys the control path needs.
- Someone on your side who did not build the controls can sign the test set and the results.

Where Tutela fits

Optional, not required. Where existing controls leave a gap, we evaluate Tutela Agentic Security by H2H for the required policies and connected systems. Tutela is customer-owned, deployed in your environment, and available on AWS Marketplace. We document supported coverage and remaining gaps; the engagement does not require Tutela.

Thirty minutes. Bring the workflow you're worried about.

No deck. We look at the workflow, the data it touches, and who has to approve it, and tell you whether this engagement is the right next step.



Book a conversation

calendly.com/cortez-h2htech
or scan the code to open the calendar

DELIVERED BY

H2H Technology

H2H brings workflow design, engineering, and AI governance together. We help you decide what deserves investment, implement the agreed work, and review the evidence for what comes next.

OTHER BRIEFS IN THIS SERIES

AI Opportunity & Value Sprint · AI Workflow Transformation and Automation · AI-native Product Development