

AI-NATIVE PRODUCT DEVELOPMENT

Build AI features with clear release criteria.

Define, build, and validate the first release of an AI feature with your product and security teams in the same room from the start.

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

01 THE DECISION

An AI feature needs a clear job in the product. When user needs, data access, integrations, and review requirements remain unresolved, a promising prototype can become difficult to release.

We define the first release with your product, engineering, and security teams, then design, build, and test it against agreed criteria. Your security lead reviews its data access and permitted actions before it uses production data.

You receive the AI Product Definition Pack, a working first release within the agreed scope, and the test results and operating handoff. Your team uses that evidence to approve the release, narrow it, or choose another path.

Bring H2H in when

- The result you need requires a new customer-facing product, internal application, or differentiated capability.
- Existing software cannot support the workflow, decision, or experience you need to create.
- An AI prototype has promise, but ownership, integration, adoption, and production controls are unresolved.

What we need from you

- A product leader who owns the release boundary and can say what is in and what is out.
- An engineering lead who can confirm data and integration constraints and build alongside us.
- Access to target users for interviews and testing, and to the systems the feature will use, under accounts you issue at your least-privilege standard.

AT A GLANCE

DECISION

Which AI feature can we ship that security will sign off on?

PHASE

02 · Redesign

DELIVERABLES

AI Product Definition Pack · experience and workflow design · working first release · validation and handoff

ENDS WITH

Approve the first release, narrow it, or choose another path.

OWNER ON YOUR SIDE

The product leader who owns the release boundary

SECURITY SIGN-OFF

Data and action limits reviewed with your security lead before the release touches production data

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 DEFINE Agree what the first release does

What happens	Define the user task, expected value, AI behavior, and release criteria. Map the data and integrations the feature needs, including what it must not access or do.
Your involvement	The product leader and target users in working sessions; your engineering lead confirms data and integration limits.
Checkpoint	Definition pack signed by the product leader and your security lead. You can stop here and keep the pack.
Output	AI Product Definition Pack: user job, feature boundary, AI behavior, data path, integrations, release criteria.

02 BUILD Build the release with the review points in it

What happens	Build the interface, AI behavior, data access, and integrations together. Implement the agreed logging and human reviews for actions identified as consequential.
Your involvement	Your engineering lead builds alongside us; your security lead reviews the data and action limits.
Checkpoint	Security review before the release touches production data.
Output	Experience and workflow design; the working first release in your environment, under your accounts.

03 VALIDATE Test representative tasks with users

What happens	Test usability, AI behavior, exceptions, and the measures agreed in Define. Record results and remaining gaps, and name the owners who will operate and review the release.
Your involvement	Target users run real tasks; the product leader, engineering lead, and security lead attend the release review.
Checkpoint	Release decision: approve, narrow, or choose another path.
Output	Validation results and operating handoff.

03 WHAT YOU GET

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

DELIVERABLE	WHAT IT CONTAINS	FORMAT	ACCEPTED WHEN
1 AI Product Definition Pack	The user job, the feature boundary, what the AI may recommend, generate, or do, the data path, the integrations, the review points, and the release criteria.	PDF plus the spec in your tracker	The product leader and your security lead sign the definition.
2 Experience and workflow design	Screens and flows, where review and exceptions happen, and how the AI behaves in each state, including when it declines.	Editable designs and specs in your tools	The product leader approves the design for build.
3 Working first release	The agreed feature built in your codebase, with defined logging, access limits, and human approvals for consequential actions.	Deployed in your environment, under your accounts, with a runbook	It passes the behavior and exception tests agreed in Define.
4 Validation and operating handoff	Usability, behavior, exception, and adoption findings against the agreed measures; remaining gaps; named owners for operation and review after launch.	PDF plus test results	Your engineering lead accepts the handoff.

04 HOW IT ENDS

Approve the first release, narrow it, or choose another path.

Approve

Your team accepts the test results and release review. Deploy the agreed scope with named owners and operating instructions.

Narrow

Part of the release meets the agreed criteria. Review a smaller release and document what still needs work.

Choose another path

The capability is better bought, configured, or deferred than built. Stop before the build cost grows.

Not included

- Guaranteed market adoption. The validation record shows how users responded; it does not promise a number in advance.
- Production data access your security lead has not approved.
- A technical path chosen before the fit is known, and features outside the release boundary named in Define.

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 AI steps in the agreed release, we document data access, permitted tools and actions, human approvals, and logging requirements. We agree record access and retention, and link these controls to release tests your team can review. 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	DEFINE	BUILD	VALIDATE
Product leader (customer)	Accountable: definition pack	Accountable: release boundary	Accountable: approve, narrow, or another path
Engineering lead (customer)	Consulted: data and integrations	Builds alongside, reviews code	Accepts the operating handoff
Security lead (customer)	Consulted: data and action limits	Approves limits before production data	Reviews behavior and exception results
Target users (customer)	Interviews	Test the experience	Run real tasks
H2H (vendor)	Responsible: definition pack	Responsible: design and build	Responsible: validation record

Assumptions

- One release, one product leader. A second feature set is a second engagement or a written change note.
- Your engineering team can host the code and run the release in your environment.
- The data and systems the feature needs can be reached under accounts you issue.

Where Tutela fits

Optional, not required. If the release involves agent tool use or sensitive data, we evaluate Tutela Agentic Security by H2H against the required policies and integrations. Tutela is customer-owned, deployed in your environment, and available on AWS Marketplace. The definition pack specifies the controls whether or not Tutela is selected.

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 Governance and Agentic Security