

THE IDENTITY

The operating system for a private, governed and self-improving AI workforce.

It runs your org's AI, learns from every governed turn, and proves compliance as a byproduct. The trace is the control, the intelligence and the evidence.

THE CORE PROPOSITION

ZeroH creates a private, governed and continuously improving AI workforce that understands how the company works. It builds context from meetings, Teams, email, calendars, CRM and documents; minimizes data exposure through masking, local models and controlled provider routing; learns from successful routines and converts them into company-owned skills; distributes those skills only after approval and security checks; enforces governance during execution; creates audit-ready proof; and manages cost through routing, budgets and task-level economics.

SIX PILLARS OF CORE IP

COMPANY WORK GRAPH



It already knows your work

Answers cite Tuesday's meeting, not the internet.

PRIVATE EXECUTION BOUNDARY



Zero exposure, three ways

Masked, policy-routed, local where it must be.
Undeclared? Masked.

LEARNING FLYWHEEL



IP that compounds

Skills, workflows and knowledge in one provenance-labeled Company Library. Competitors rent the same models, not your learning.

GOVERNED SKILL DISTRIBUTION



Viral AI, wrapped in governance

One governed Company Library, every install job-board visible. Nothing installs, updates or learns behind your back.

COMPLIANCE PROOF



Audit-ready out of the box

Enforced before the model. Sealed after. Proof on demand.

COST OPTIMIZATION



Costs at check

Every task carries its receipt. No task overpays.

Each pillar has an internal IP name and a buyer-facing headline. Both are fixed vocabulary.

☐ sandboxed

☐ policy first

☒ proof

AI employees, not AI freelancers.

THE WORKFORCE

AI employees, not AI freelancers.

Hire agents with identities, jobs, and budgets. Governed policies are enforced before every action. Every job ends in a digital receipt.

Sales

Marketing

Operations

Finance

HR

Compliance

Help Desk

Projects

Chief of Staff

Ask Ali

RUN AGENTS LIKE EMPLOYEES

- 1 **Hire** from the role library. Each agent gets a mandate, jobs, and a budget.
- 2 **Brief** them with your context: meetings, mail, calendars, CRM, documents. Read as you, masked before any model.
- 3 **Supervise**: anything outbound waits for a yes. Approvals, kill-switches and spend guardrails are policy, not promises.
- 4 **Review** performance and cost per agent and per job; every run ends in a signed receipt.
- 5 **Improve**: agents surface what works as ratifiable proposals. Nothing ships on its own.

WHERE IT RUNS

ZeroH Assistant
Teams and the Office pane, always running as you.

ZeroH Governance
Masking policy, approvals, the live ledger, proof packs.

ZeroH Studio
Build agents, skills and governed workflows, no code.

Ask Ali
The domain advisor: cited answers, honest abstentions.

WHAT THE LEADER SEES

JOBS RUN / MO

4,120

up 2.9x this year

SPEND / MO

€982

down 54% this year

INSIDE POLICY

100%

by construction

ONE LEADERBOARD FOR EVERY AGENT: JOBS, SPEND, VALUE, RECEIPTS. SAMPLE DATA.

WHERE YOUR TEAM ALREADY WORKS

Type @ZeroH in Teams and the whole fleet answers. Drafts arrive in your tone, your template, your brand, held for your yes.

STRATEGIC INTELLIGENCE

The Governance Console reads the whole workforce: what needs attention, department output, model concentration risk, approval queues. Compliance is the byproduct; the same trace is the intelligence.

THE COMPANY LIBRARY

Draft a client email in house style
learned from 60 corrections by your sales team · ratified v3 by your DPO · receipt № 2481

LEARNED FROM YOUR WORK

€0.03 / run
142 runs / mo

EVERY LEARNED ROUTINE BECOMES A VERSIONED, OWNER-LABELED, EXPORTABLE ENTRY. YOURS WHEN YOU LEAVE.

ONE STACK, GOVERNED

The whole workforce stack. One governed system.

Teams bolt together eight tools to run agents, and still cannot prove what happened. ZeroH ships the stack as one system, and every component writes to the same signed ledger.

Triggered by key events or signals	like Zapier	Given access to your context	like Zep
Connected to your apps	like Composio	Auto-routed across all LLMs	like OpenRouter
No-code agent and workflow builder	like Dust.tt	Coordinated into departments and teams	like CrewAI
Kept alive through long-running work	like Temporal	Traced at every step of every run	like Langfuse
Performance evaluated, learning ratified	like Braintrust	Masked before any model, on your soil	ONLY ZEROH
Held for a human yes, policy fail-closed	ONLY ZEROH	Sealed on the signed ledger	ONLY ZEROH

THE CONTROLS YOU GET

- Proof packs, on demand**
Per agent, per project, or per time period. The evidence assembles itself from signed turns; you never chase it.
- Regulatory filing and reporting**
The governed trace becomes the report. Frameworks map to your controls, not the other way round.
- ISO 42001 and GDPR, out of the box**
The AI management system and data-protection duties are enforced before the model and sealed after, from day one.
- ISO 27001 and SOC 2, as a plug-in**
Extend the same evidence into your existing GRC provider, Vanta or another, rather than running a second stack.
- Best-in-class regulatory reporting.** We work with regulators not to overwhelm them, but to hand them evidence and the tools to navigate compliance easily. Every governed turn writes its own proof to the append-only ledger.

TWO WAYS TO RUN IT

- OFFER 1 · SAAS**
Start free today
\$0.00 a month; you pay provider cost at 0% markup, on your Microsoft bill. Live on Microsoft Marketplace.
- OFFER 2 · MANAGED APP**
Move on-soil
The same stack deployed into your own Azure subscription when data residency or the regulator asks.

EVERY JOB ENDS IN THIS RECEIPT

ZEROH · JOB RECEIPT

SAMPLE DATA

Send into mail · Jonas W

ZeroH Governed ✓ VERIFIED

- ✓ **Private data was masked before any AI saw it**
14 items protected · raw personal data reached the model: no
- ✓ **A person said yes before anything went out**
approved by J. Weiss · the send waited 41 minutes
- ✓ **Real names were restored only after the yes**
on your servers · the AI provider never saw them
- ✓ **Sealed in the tamper-evident record**
receipt № 2481 · signed by your organization · ES256

€0.09

what this run cost · metered, never estimated · budget checked before it ran

Meets your frameworks: ISO 42001 · ISO 27001 · GDPR · SOC 2 · AICM

Rent the intelligence. Keep the learning.

Models are the rented part; swap them freely. What the workforce learns from your work stays yours: versioned, owner-labeled, exportable from the Company Library.

WHAT IT REACHES, AND WHAT IT MAY DO

Connected to your stack. Checked before every call.

Breadth is table stakes. What decides whether you can put an agent near real systems is what happens between the agent deciding to act and the action landing.

FROM A FOLDER TO CITED ANSWERS

- 1

Put the files in a folder
A SharePoint folder you already have. Nobody converts anything, nobody uploads a copy elsewhere.
- 2

Create the knowledge base in Studio
Point at the folder, say who may read it, pick how its documents are cut.

New knowledge base

Everyone can read

Clause-numbered

Keeping it current
- 3

Attach it to any agent, or all of them
The help desk gets the product base, compliance gets the standards. Each agent answers from what it is attached to, and from nothing else.

PROSE, OR CLAUSE-NUMBERED

General documents
Handbooks, policies, decks. Cut along the document's own headings.
Remote-Work-Policy.pdf · §3.2

Clause-numbered
Regulation and standards. The clause number and page survive into the citation.
AAOIFI FAS 28 · clause 4.2(b) · p. 17

For anything a regulator reads, that second row is the whole game: an answer your compliance team can check, rather than "it is in this PDF somewhere".

THE CHANNEL IS ALREADY THE PROJECT

A project rides the Teams team you already have: same membership, same files, same threads, nothing to create. Ask the project manager what is waiting on who, or ask marketing to write the newsletter, in the same channel. Each answer is read under the identity of the person asking, so it can only ever contain what they could already open. The thread comes with you: start in Teams, open the pane in PowerPoint, and the same masked conversation resumes.

REACH

22,000+ apps already ship an MCP server. Yours are one governed connection away: ZeroH connects as the person asking, so the server's own permissions still apply, and the result is masked and recorded before any model sees it.



Seven systems connected natively today plus the Microsoft estate, 36 governed tools in the box. Not on the list? Ask: we ship new tools every day, and a connector you need is usually days away.

BEFORE EVERY CALL

A tool is granted to an agent, for a job. Not "the AI can use your CRM": the sales agent may look up an account because that is the work it was hired to do.

Account name	tokenized	the model reasons over [ORG-77be03]
Billing address	never sent	redacted at the seam, it leaves your tenant for nobody
Deal value	banded	"100k to 250k" instead of the figure
Stage	in the clear	the model needs it, and it identifies no one

A field you never declared is masked, not guessed. The person reading sees the real name, rehydrated on your servers after the model is done.

YOUR DPO
decides the agent may reach the system at all.

YOU
decide it may act as you, the first time it asks.

ANY TIME AFTER
turn it off.
Revoking is always allowed.

Reading changes nothing. Sending, writing and spending leave a mark outside ZeroH, so a write goes one of three ways:

a person approves

the agent may decide

never unattended

SKILLS ARRIVE BECAUSE YOU NEEDED THEM

Your team keeps doing the same thing by hand. The workforce notices the routine and proposes the skill for it, either one already scanned and sitting in your library or a new one drafted from your own work. A candidate whose scan came back anything but clean is never even proposed. And nothing installs or updates itself: an upstream change re-enters the same queue and waits for a person.

Nothing reaches a model unmasked. Nothing acts without a mandate.