

Human 2.0 Intelligence Flywheel Roadmap

Knowledge that makes people 10x more capable.

The Intelligence Flywheel — and the roadmap that rides the new AI frontier.



The H2O Flywheel Roadmap

Four complementary pieces — one story, told at the depth each room needs. Start practical, zoom out to vision, then go under the hood.

1

DECK · 12 slides

START HERE

Operators Edition

FOR Operations & L&D, frontline leaders

“What changes on my shift?” The flywheel through one person's day — jargon-free, fear addressed.

2

DECK

ZOOM OUT

The Vision

FOR CIO office, leadership, investors

Why now, the strategy, the moat, the frontier shift. The big-picture case for the whole platform.

3

DECK · 14 slides

GO UNDER

Technical Deep-Dive

FOR Architects, integration, IT / security

How it actually works: architecture, data flows, trust & verification, and systems integration.

★

DOCUMENT · 31 pages

REFERENCE

Product Foundation

FOR The build team & anyone going deep

Per-product spec depth plus worked oil & gas examples — the written reference behind all three decks.

HOW THEY FIT

They don't replace each other — they complement. Present one, two, or all three depending on the audience; the Foundation document answers anything a deck raises. The same package fields field-operations questions, CIO-office vision questions, and deep technical questions.



The Intelligence Flywheel in practical steps

How EON makes your people 10x more capable — without replacing them.

A briefing for operations and learning & development teams.



This makes your people stronger. It does not replace them.


Every tool in this deck does one thing: it puts a senior expert's judgment into your team's hands, right when they need it.

The worker stays in charge. The AI helps — and, crucially, it asks a human for help when it isn't sure, instead of guessing.

10x

more capable

The worker stays in charge. The AI helps, and asks for help when it is not sure.



The knowledge is walking out the door

S

Sam

Lead operator · 30 years · retiring in 6 months

Knows every quirk of the plant — most of it in his head, not the manual. When he leaves, that judgment leaves with him.

M

Maria

Technician · joined 8 months ago

Capable and keen — but can't absorb 30 years of know-how quickly. On the night shift, the expert isn't there to ask.

This gap — between the expert who's leaving and the team who's staying — is what the flywheel closes.



Capture it. Put it to work. Learn from it. Repeat.

1

Capture what experts know

Record how your best people actually do the job.

2

Put it on the worker

Guide anyone through it, hands-free, in the field.

3

See what really happens

Check the work and capture how reality differs.

4

Make it better

Feed that back so the next shift is sharper.

From the manual to the gloves — and back again. Most systems only push knowledge out. Ours also brings reality back, so it keeps getting smarter — every shift makes the next one better.



1

Capture what your experts know

In our product line: [Genesis 3 + Scenario Factory](#)

What happens

An expert demonstrates a job once. We turn it into a clear, 3D step-by-step guide — and we can do it for thousands of procedures at once, not one at a time.

SAM & MARIA

Sam walks through the C-3 separator shutdown one time, wearing a headset. We build it into a guided 3D procedure in days. It even captures his unwritten trick — in cold weather, crack the bypass a quarter-turn first. That know-how is now the plant's, not just Sam's. The same engine can author all 9,000 of the plant's procedures over a weekend.

What it means for you: Your experts' know-how becomes a permanent asset — captured before they retire, not lost when they do.



2

Put it on the worker

In our product line: Field IQ + Brainy-on-the-Shoulder

What happens

The guide rides on safe-for-the-field smart glasses. It shows the next step, hands-free, and answers questions right where the work happens.

SAM & MARIA

3 a.m. Maria has to isolate the separator and has never done it alone. She puts on the glasses; the guide walks her through it step by step. At a tricky gauge read, the assistant isn't sure — so it pauses and calls the on-call senior, who sees her exact view and confirms in 30 seconds. It guided her safely, and it knew when to ask for help instead of guessing.

What it means for you: New people perform like seasoned ones sooner — and the system escalates to a human rather than risk a wrong call at the valve.



3

See what really happens

In our product line: Assess IQ + xAPI


What happens

Cameras and the glasses check the work against the procedure — confirming competency and capturing where reality differs from the manual.

SAM & MARIA

Assess IQ confirms Maria did the isolation in the right order — verified, on the record. It also spots two things: several operators over-torque a flange (unsafe — flag it), and many crack the bypass first in the cold, just like Sam (smart — better than the written SOP). Every action is written back into your own LMS in the standard your team already uses.

What it means for you: You can finally prove what's actually happening in the field — and the gap between the manual and reality becomes valuable data you own.



4

Make it better

In our product line: **Compound IQ**


What happens

Overnight, what we learned in the field improves the guides automatically — the good practices spread, the unsafe ones get fixed.

SAM & MARIA

That night, the cold-weather trick becomes part of the official procedure — so everyone gets Sam's wisdom, not just the people who worked beside him. The unsafe over-torque becomes a quick targeted lesson for the 15 operators who did it, plus an alert to the supervisor. By the next cold-weather shift, the guidance is already better than what Maria got.

What it means for you: Your training improves itself every shift — the knowledge compounds instead of going stale, even after Sam is gone.



Can you trust it? Yes — and here's how.

✓ It checks itself first

Before anything reaches the field or changes a procedure, the system tries hard to prove itself wrong. Only what survives that check goes live.


✓ It asks when unsure

On a safety-critical step, an uncertain assistant stops and calls a human — it never bluffs its way through.

✓ Everything is on the record

Every action and decision is logged, reversible, and mapped to your safety rules (ATEX, API, OSHA) — audit-ready by default.

The AI never gets the final say, on its own, on a step that could hurt someone.



Outcomes your team can measure

Knowledge that outlives retirement

Your experts' judgment becomes a permanent, owned asset.

Faster time-to-competency

New hires reach the bar sooner, guided by the best.

Fewer incidents

Adherence is verified in the field, not assumed.

Audit-ready by default


Every action recorded and mapped to your safety rules.

Works with your LMS

We take in your old SCORM courses and hand back xAPI.

It improves itself

Every shift makes the guidance a little better.



A 90-day start — and you own what we build



1 Look

We baseline how you work and find the highest-value procedures.

2 Choose

Together we pick the first jobs to capture.

3 Build

We author the first wave of guides — fast.

4 Go live

The loop runs, improves, and you own the knowledge.

You keep the guides, the data, and the improving system. The knowledge is yours.



10x your people. Own the knowledge.

Capture what your experts know. Put it where the work happens. Learn from what really happens — and get better every shift.



The Vision: Human 2.0 Intelligence Flywheel Roadmap

FOR CIO office and leadership

Why now, the strategy, the moat, the frontier shift. The big-picture case for the whole platform.



10x your people. Don't replace them.

Fuse AI depth with human judgment, dexterity, and practical know-how.

- Turn a **retiring expert into a permanent asset**, not a vacancy.
- A **workforce that compounds** — every shift makes the next one better.
- The promise we sell: **Human 2.0**.



Faster. More honest. More agentic.

- **Faster** — work that took weeks now finishes in hours.
- **More honest** — models flag uncertainty and escalate instead of bluffing.
- **More agentic** — hundreds of sub-agents plan, act, and check each other.
- **A new apex** — Mythos: near-superhuman reasoning for the hardest calls.

EON doesn't build the models — we're built to ride them. Always on the edge.



The expert is retiring. Raw AI can't stand at the valve.

- **The most valuable knowledge isn't written down** — it's in the hands of senior operators.
- **They're retiring faster** than the next generation can be trained.
- **A confidently wrong AI** answer at a high-pressure valve can cost a life.
- **Until now: slow to build**, and impossible to trust in the field.



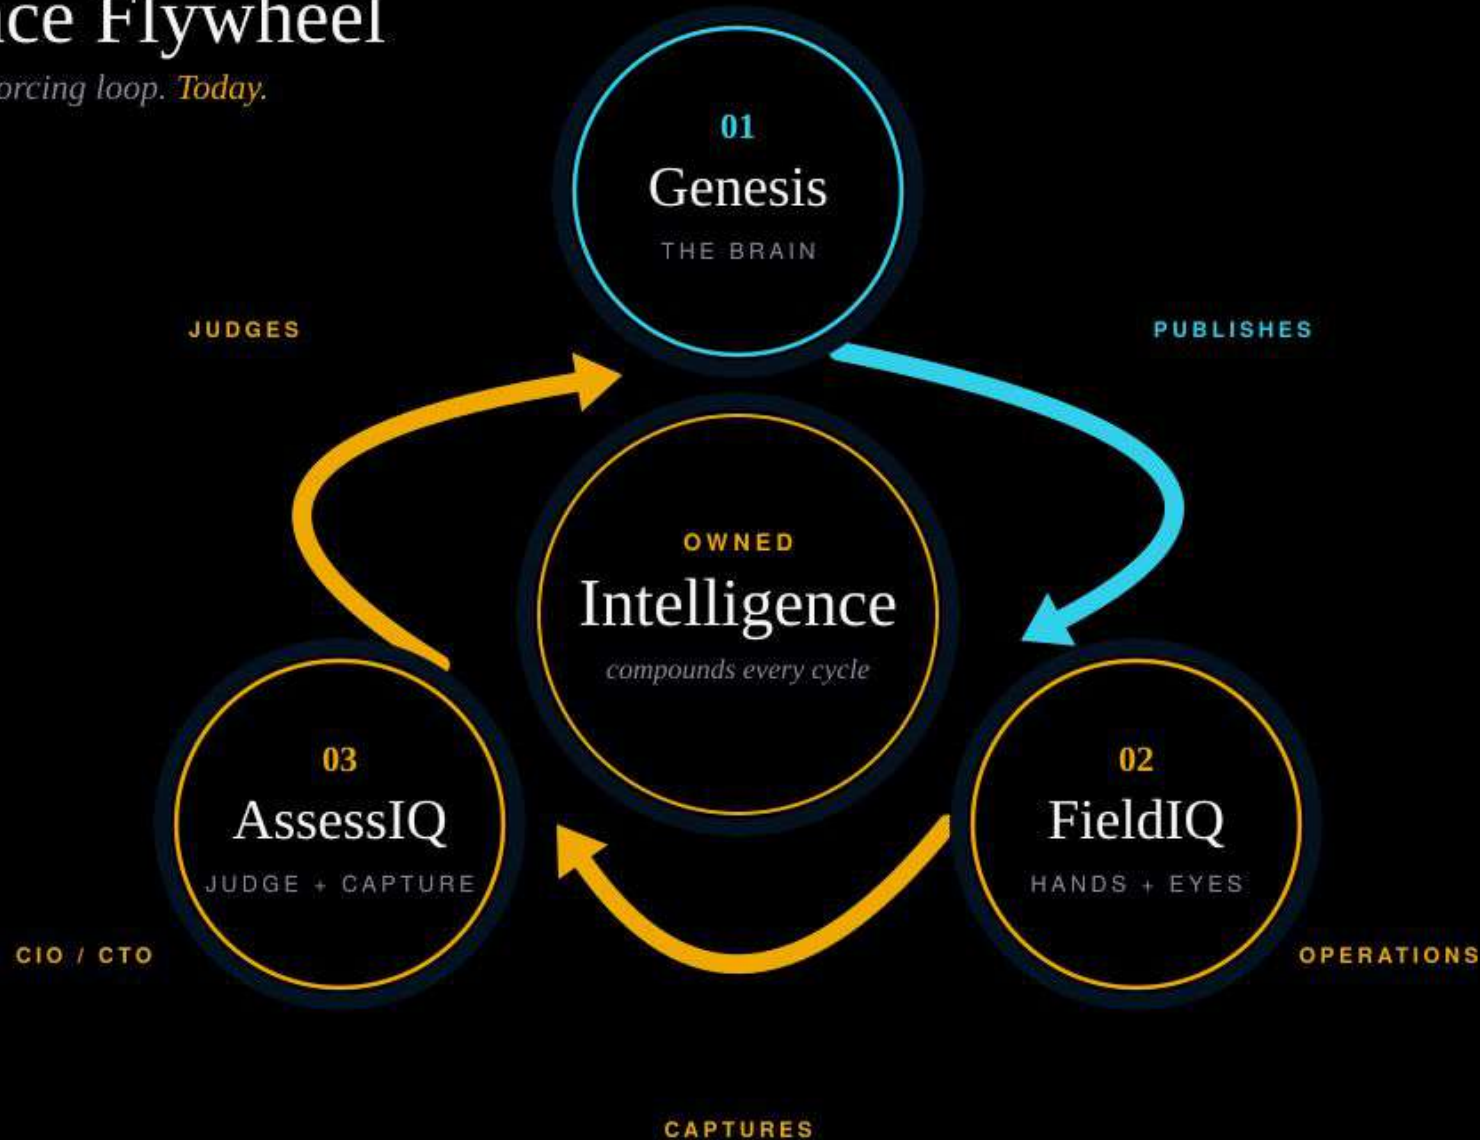
Bytes → Atoms → Bytes

- **Most knowledge is trapped in bytes** — LMS, PDFs, people's heads.
- **EON turns it into atoms** — action in the physical world, on the worker.
- **The field generates new bytes** — captured reality that flows back.
- **EON is the only loop** that runs both directions.



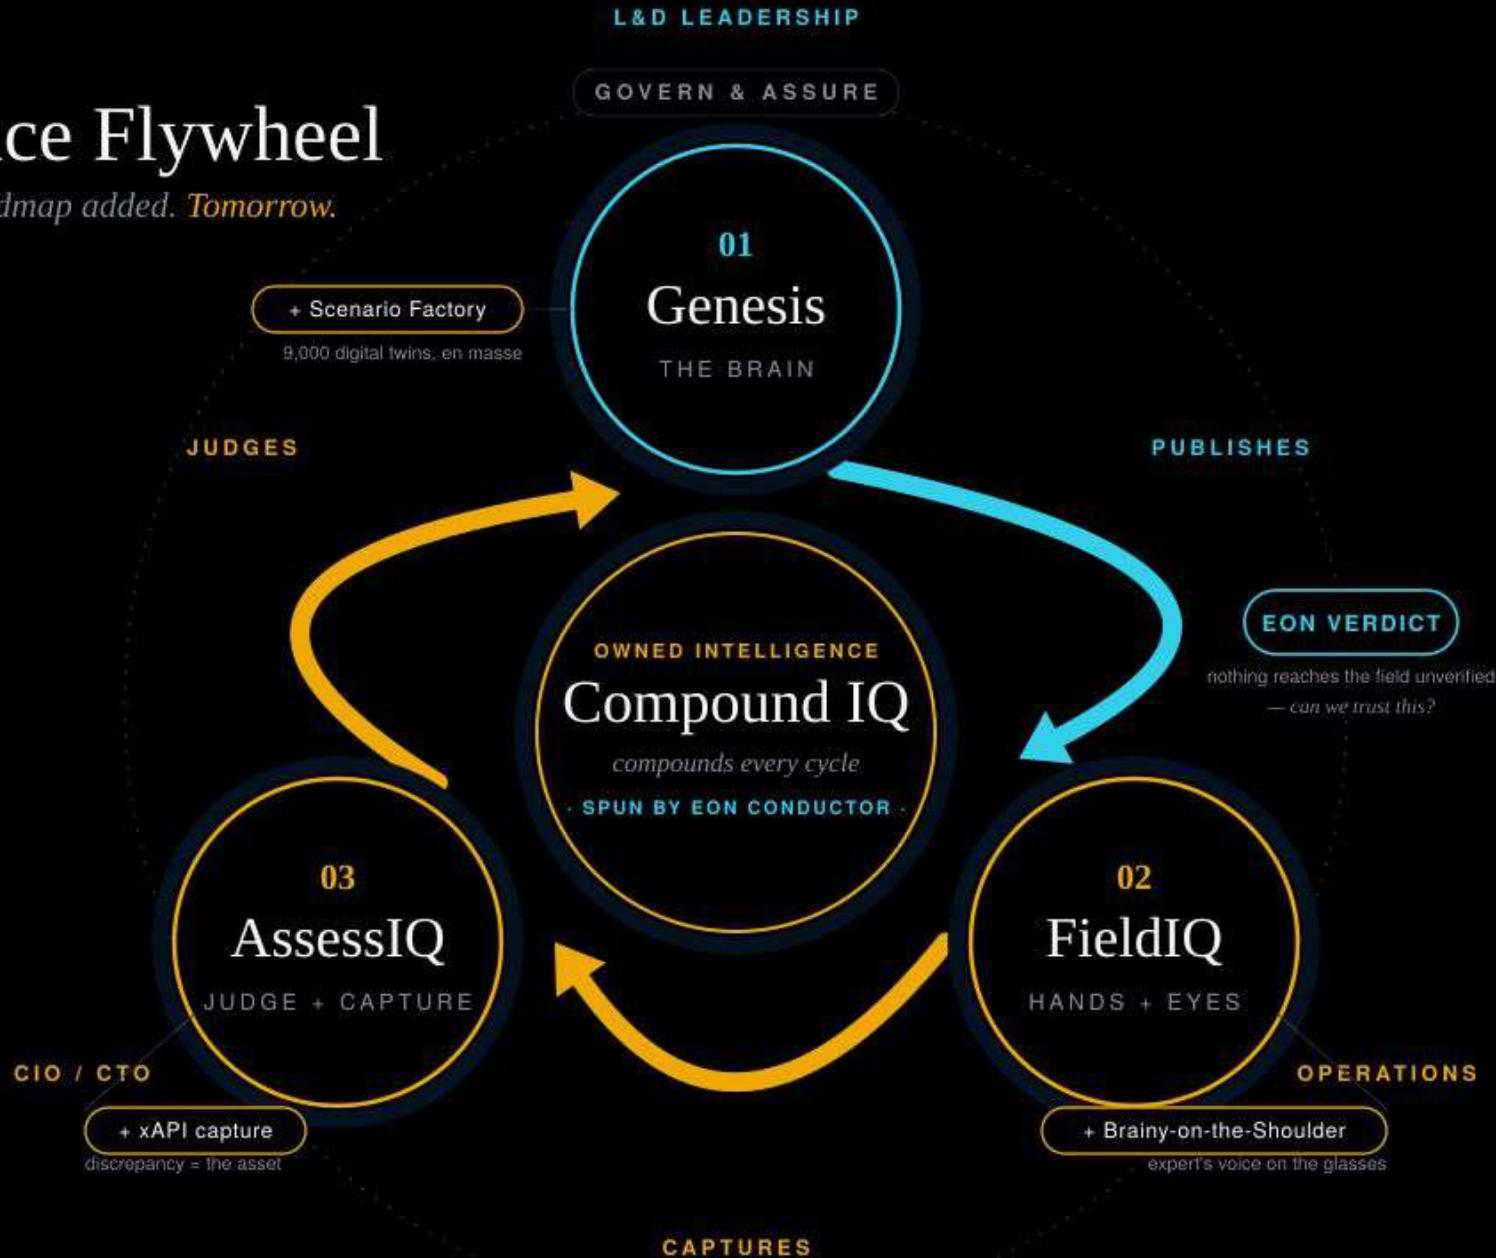
The Intelligence Flywheel

Three products. One self-reinforcing loop. *Today.*



The Intelligence Flywheel

The same loop — with the roadmap added. *Tomorrow.*



Do it en masse.

- **Feed an entire plant's processes** — 9,000 digital twins, not one at a time.
- **Simulate** → field → assess at the scale of the whole company.
- **Almost no human touch** — it chews tokens while you sleep.
- **Why now:** faster, cheaper, agentic, and honest enough to trust.




Genesis 3 + Scenario Factory

Problem Your best operators' **judgment is not written down**; it retires with them.

What we solve Capture **tacit expertise as governed digital twins** — and author them en masse, not one at a time.

What it does Turns a **demonstration or SOP into a self-verified 3D procedure**; the Factory runs hundreds of parallel agents overnight.



Field IQ + Brainy-on-the-Shoulder

Problem Knowledge sits in an LMS; the work happens at the valve.

What we solve Put the expert's judgment on the worker — in the field, hands-free, exactly when it is needed.

What it does Brainy rides on AR smart glasses, guiding step by step and escalating to a human when unsure instead of guessing.



Assess IQ + xAPI

Problem You assume procedures are followed; you cannot prove what actually happened.

What we solve **Verify competency and capture reality** — including the gap between the procedure and what the worker really did.

What it does Multi-camera **Assess IQ judges performance and records every action as xAPI** into your own system. The discrepancy becomes the asset.



Compound IQ

Problem Most training is a one-time event; the knowledge never improves.

What we solve **Make every shift improve the system** — owned intelligence that compounds instead of depreciating.

What it does **Feeds captured discrepancies back to sharpen Genesis** and Brainy automatically, so each cycle beats the last.



EON Conductor

Problem Frontier intelligence is powerful but costly, and hard to control at scale.

What we solve Run the **whole loop en masse, with the right model on each task** — and almost no human babysitting.

What it does **Plans a goal, fans it across hundreds of verifying sub-agents**, and routes each to the right tier: fast for volume, Opus for judgment, Mythos for the apex.




EON Verdict + Govern & Assure

Problem You **cannot run autonomous AI in a safety-critical plant** if you cannot trust it.

What we solve **Make autonomy trustworthy** — nothing reaches the field or updates the brain until it has been proven.

What it does **Adversarial agents try to break every result** before a human sees it; passports, audit, rollback and the Trust Ledger record it. Concord extends this safely across firms.




Backcast IQ

Problem **Knowing where to point** all of this is the hardest, highest-stakes decision.

What we solve **Architect the AI-native future state and the path to it**, with the most capable reasoning available.

What it does Runs on Mythos-class intelligence to **backcast from your 2031 operating model** to the moves you make today.



Commercial Intelligence + Top Gun Outpost

Problem The same **engine that runs operations should also grow the business.**

What we solve **Point the loop at revenue** — find, qualify, and win, autonomously and self-verified.

What it does **RADAR, Bullseyes, and Top Gun** run autonomous outbound; Top Gun Outpost packages it as a product your teams — or customers — deploy.



The discrepancy is the IP asset - You own it

- The **gap between procedure and reality is tacit knowledge** — in no training corpus on earth.
- It compounds every shift — **the wheel deepens the more it turns.**
- It is the **embodied data the robotics era is starving for.**
- Nobody can buy it or scrape it. **You own it.**



Fast where it's cheap. Frontier where it counts.

- **Fast** — volume work: RADAR, routine field narration, bulk drafting.
- **Workhorse (Opus)** — judgment: Conductor, Compound IQ, Verdict, capture.
- **Apex (Mythos)** — the rare high-consequence call and Backcast strategy.
- **Conductor routes every task** to the cheapest model that can do it right.



10x your people. Own the intelligence.

Every turn of the wheel turns human work into owned, governed, compounding intelligence — captured from your own operations, deployable where raw AI cannot go.


How to engage: the 5 days workshop and productized 90-day H2O Engagement — pulling through to multi-year platform deployment value



Technical Deep-Dive the Intelligence Flywheel as a system

Architecture · data flows · the pipeline · trust & verification · how it integrates.

Part 3 of the H2O deck package — pairs with the Operators Edition, the Vision deck, and the Product Foundation document.



Ten technical questions, answered

01 **System architecture**
The flywheel as one platform: stations, engine, gate, tiers.

02 **Data model & flows**
The artifacts that move, and where they live.

03 **Encode pipeline**
Genesis 3 + Scenario Factory: 11 steps, parallelized.

04 **Field loop**
Field IQ + Brainy: the sub-3-second loop and escalation.

05 **Capture & standards**
Assess IQ + xAPI/cmi5: schema, discrepancy classes, LRS.

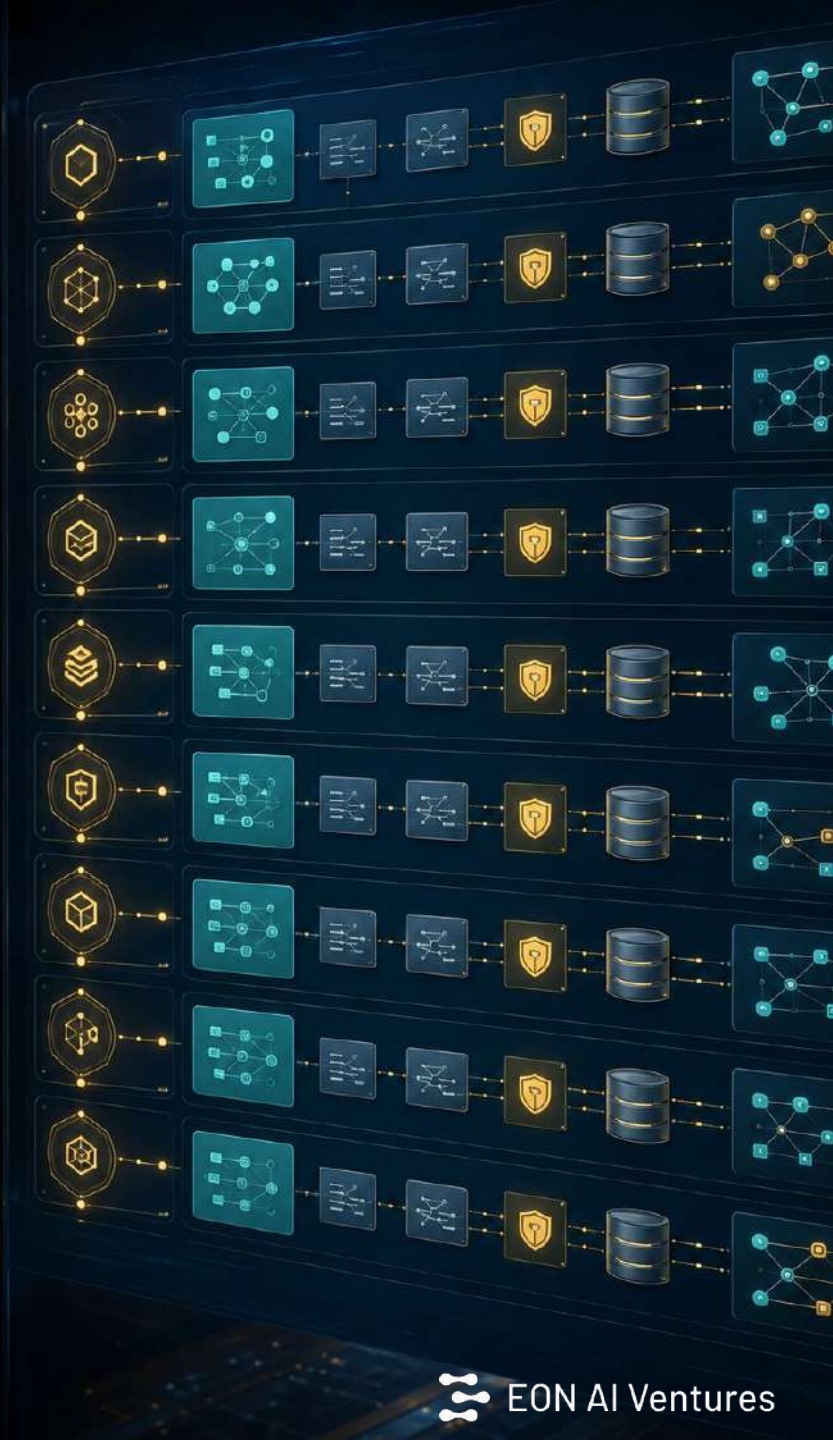
06 **Compound loop**
Compound IQ: the nightly learning batch, gated.

07 **Orchestration & tiers**
EON Conductor: fan-out, routing, execution trace.

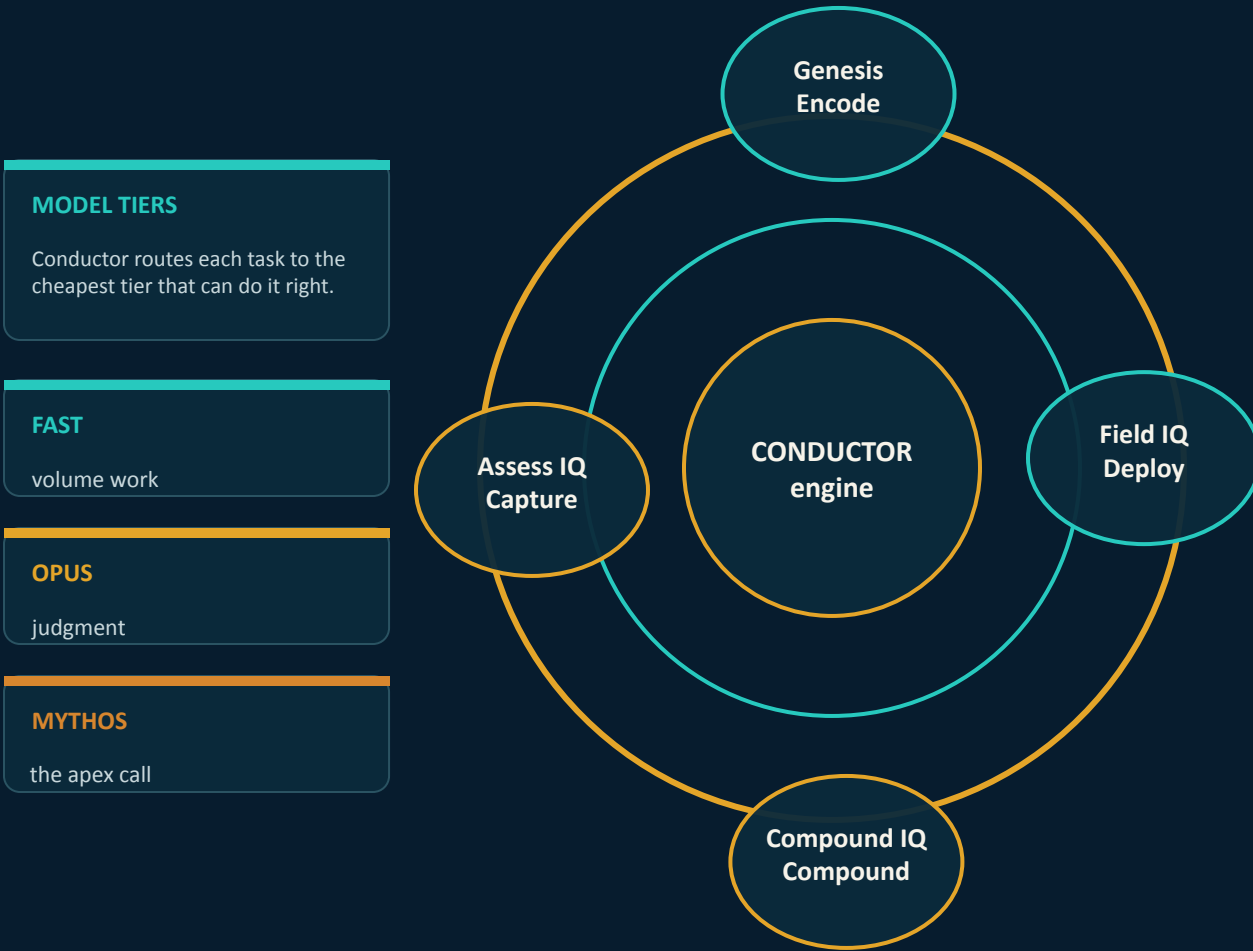
08 **Trust & verification**
Verdict + Govern & Assure + Trust Ledger.

09 **Integration & deploy**
Surfaces, LMS/LRS, identity, residency, the 90-day stand-up.

10 **Security & economics**
Audit posture, and why running it is affordable.



One loop: four stations, an engine, a gate, three tiers



Verdict (the gate) sits on every publish path — nothing reaches the field or the brain unverified.



The artifacts that move — and where they live

Content library

Digital twins · rubrics · SCORM/cmi5 packages

Learning Record Store (LRS)

xAPI statements · discrepancy records (safe / unsafe)

Audit store

Agent passports · immutable logs · execution traces · Trust Ledger

HOW THEY MOVE

Genesis writes twins + rubrics + standards packages → Content library.

Field IQ/Brainy emit field actions → LRS as xAPI; decisions → Audit store.


Assess IQ scores vs rubric; writes discrepancy-tagged xAPI → LRS.

Compound IQ reads LRS telemetry → proposes refinements → (Verdict) → updates twins in the library.

Conductor writes an execution trace for every run → Audit store.

Verdict / Govern & Assure issue passports + logs → Audit store; xAPI handed back to the customer's own LMS.

Standards in / out: SCORM ingested → xAPI + cmi5 emitted into the customer's own system. No proprietary lock-in.



Genesis 3 — 11 steps in 3 phases, then parallelized

Ingest & structure STEPS 1–5

Parse SOP / demo / CAD / photogrammetry (photo-to-3D, Gaussian splatting); build the 3D model; attach hazards, tolerances, PPE.

Assemble & confirm STEPS 6–7

Draft the scenario from 60+ interaction types (valve turns, LOTO, gauge reads, comms, emergency branches); a reviewer confirms.


Assess, branch, publish STEPS 8–11

Add rubrics, error-state branches, localization; publish to XR / glasses / browser / mobile + SCORM/cmi5 + xAPI profile.

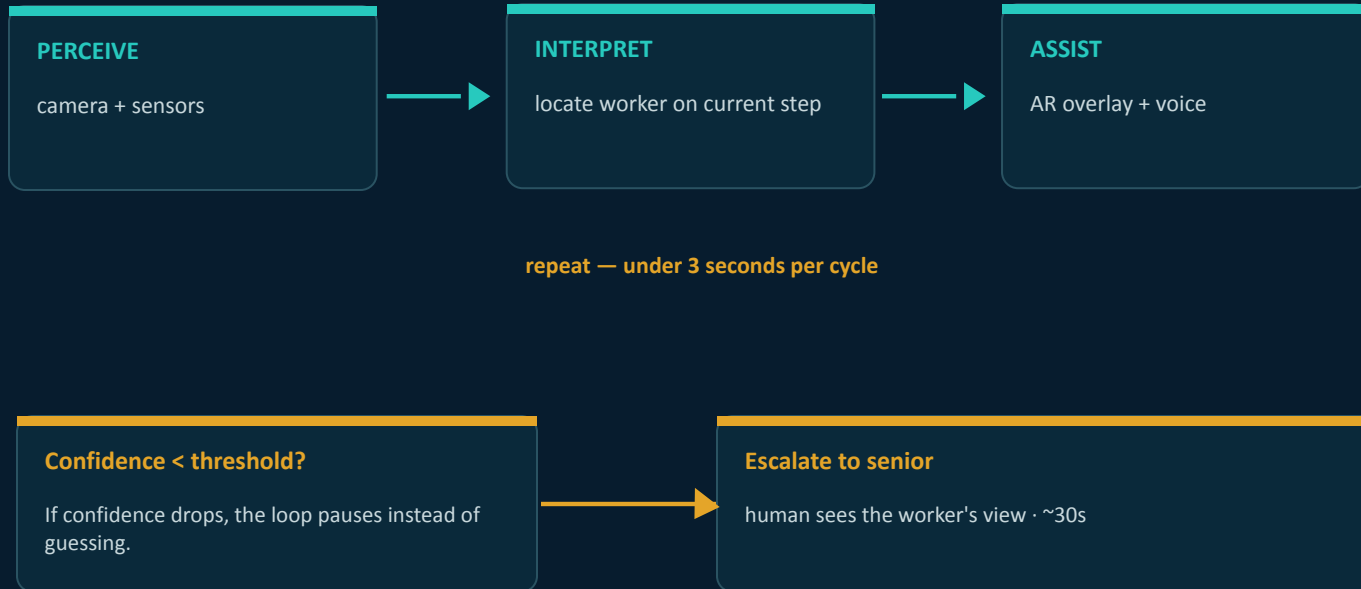
SCENARIO FACTORY — THE SAME PIPELINE, FANNED OUT

Conductor runs hundreds of pipeline instances in parallel (one per procedure) with resumable state. Each agent self-verifies against the source SOP; Verdict's adversarial agents try to break each twin (wrong tolerance, missing hazard, unsafe sequence) until it converges. Only genuine ambiguity routes to a single human queue — so a whole plant's ~9,000 procedures author concurrently, not one at a time.

Inputs: SOPs · manuals · expert capture (video/voice) · CAD/photos · legacy SCORM → Outputs: governed twins · rubrics · SCORM/cmi5 · xAPI profiles



Field IQ + Brainy — a sub-3-second loop with honest escalation



Hands-free, on intrinsically safe (ATEX-rated) glasses. The 30-second review threshold is encoded in the agent passport, so escalation behavior is governed, not ad-hoc. Every step, query, and escalation is emitted as an xAPI activity stream — the raw material the Capture station consumes. Honesty (escalate rather than guess) is the precondition for running this unattended at a high-pressure valve.



Assess IQ + xAPI/cmi5 — score, classify, and hand back

WHAT CAPTURE DOES

Multi-camera (fixed + glasses) scores adherence step-by-step vs the Genesis rubric.

Emits each observation as an xAPI statement into the customer's LRS.

Classifies the discrepancy: safe/repeated (candidate better practice) vs unsafe (intervention).

xAPI statement

actor – verb – object

“operator – performed – LOTO step 4 – out of sequence”

SCORM → xAPI / cmi5

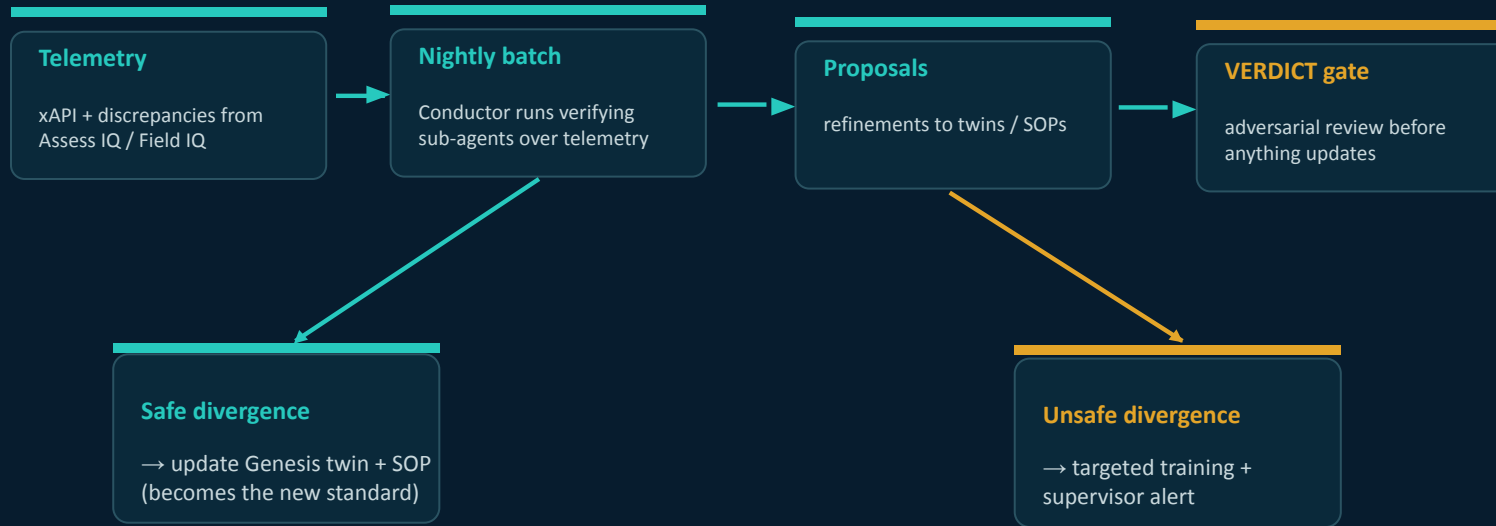
EON ingests legacy SCORM and emits xAPI/cmi5 back into the customer's own LRS.

The discrepancy is the asset — captured here, fed to Compound IQ next.


The result is the richest field-performance data available, in the standard your team already champions.




Compound IQ — the nightly learning batch, gated



Approved refinements flow back to Genesis & Brainy. The honesty gate plus Verdict's adversarial review stop a bad refinement from silently re-entering a safety-critical procedure.




EON Conductor — fan-out, route, verify, trace




Intelligence allocation — the principle

Fast / low effort — RADAR sweeps · routine field narration · bulk drafting
Opus / high effort — Conductor · Compound IQ · Verdict · Assess IQ · Genesis verification
Mythos / max effort — Backcast strategy · novel high-consequence anomalies only

Most of the stack runs at commodity cost; frontier reasoning is spent only where consequence justifies it. Fast mode $\approx 2.5\times$ faster at lower price.



Verdict + Govern & Assure + Trust Ledger — the gate




Trust Ledger

Binds each agent's measured honesty — escalation rate, overconfidence — to its passport, and publishes it into the audit trail. Turns model honesty into a provable control a CRO, insurer, and regulator can sign — not a promise.

EON Concord — across firms

A passport-interchange protocol: one operator's agents coordinate with contractor, regulator, insurer, and OEM agents using role-aware tokens and one shared, audited record. Inter-firm autonomy becomes insurable because every side reliably reports what it did.



How it plugs into what you already run

Delivery surfaces

XR headset · ATEX smart glasses · browser · mobile — same governed twin, rendered to each.

LMS / LRS

SCORM ingested in; xAPI + cmi5 emitted out to your existing LRS/LMS. Your standard, your system of record.

Identity & roles

Role-aware access and agent passports; escalation thresholds encoded per procedure.

Data residency & audit


Immutable logs, rollback, and retention that meet your audit requirements; mapped to ATEX/API/OSHA/ISO.

Edge & connectivity

The sub-3-second field loop runs at the edge; offline capture syncs back when connectivity returns.

Stand-up: 90 days

Score → Backcast → map workflows → cut drag → build edge twins → stand up Govern & Assure → hand off. You own it.



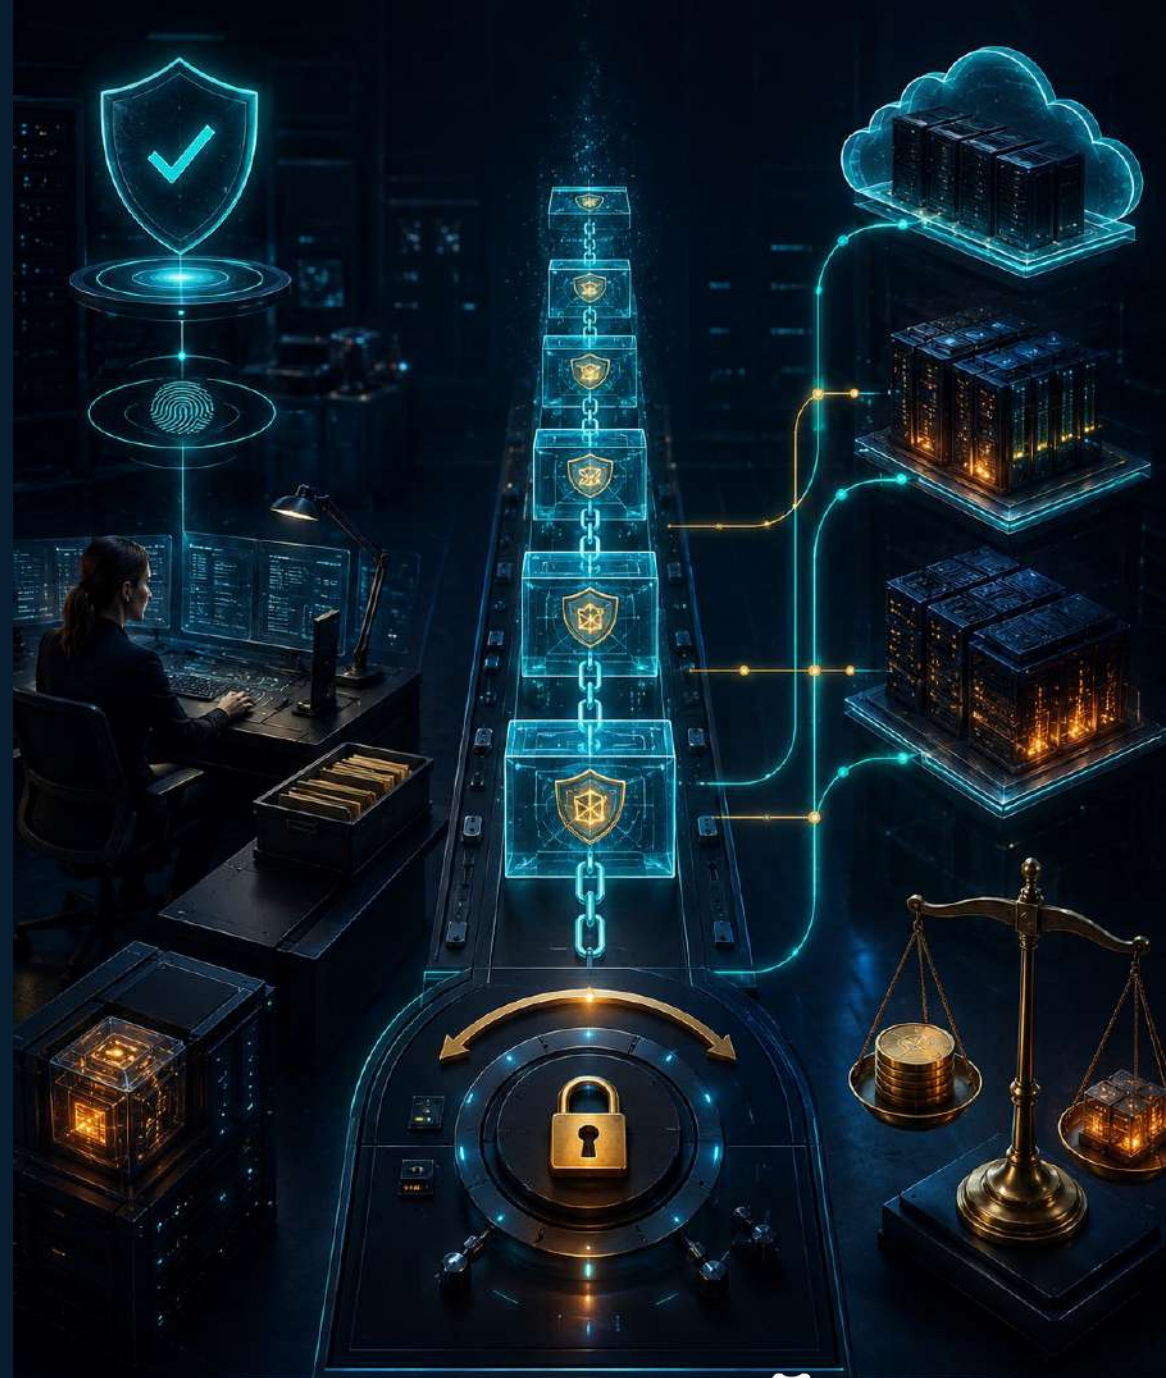
What the CISO sees — and why it's affordable to run

Governance & audit posture

Every agent carries a passport mapped to its regulatory frame.
Every decision is logged immutably and is reversible (rollback).
Exceptions route to a human queue; retention meets audit needs.
Adversarial verification gates every publish path.
Trust Ledger publishes measured honesty into the trail.
xAPI handed back to your own LRS — data stays in your system of record.

Why running it is affordable


Tier routing — most work runs on the cheapest model that's good enough.
Effort modes — fast mode $\approx 2.5\times$ faster at lower price; trade depth for cost.
Parallelism — hundreds of agents author a plant concurrently — overnight.
Falling unit cost — it “chews tokens” while you sleep; cost drops as compute improves.
Apex only where it pays — Mythos reserved for the rare nine-figure decision.



Priority is build order — the loop needs all of it to turn

P0	Genesis 3 + Scenario Factory	Encode · new
P0/P1	Verdict + Integrity Suite / Trust Ledger	Gate · new
P0/P1	Commercial Intelligence (RADAR · Bullseyes · Top Gun)	Growth · strengthen
P1	EON Conductor	Engine · new
P1	Field IQ + Brainy-on-the-Shoulder	Deploy · strengthen
P2	Compound IQ	Core · new
P2	Concord · Top Gun Outpost · Assess IQ · Backcast IQ	Gate / Growth / Capture / Apex

Conductor unblocks the P0 factory and the compound batch; Verdict gates everything before it ships. That is why the engine and the gate come early even though the visible value is the stations.



One loop. Owned, governed, compounding.

Encode → Deploy → Capture → Compound, spun by Conductor, gated by Verdict, recorded by Govern & Assure — integrating with the systems and standards you already run.

Full per-product detail and worked oil & gas examples are in the H2O Product Foundation document.

