

EON Compound IQ

A Verdict-gated nightly learning loop that turns the gap between the written procedure and what the worker actually did into a sharper twin — or a lesson that leaves the twin alone.



Table of Contents

1. Abstract.....	2
2. The problem, in plant language.....	2
3. Where Compound IQ sits on the Intelligence Flywheel.....	2
4. How one night works.....	3
4.1 Ingest.....	3
4.2 Cluster.....	3
4.3 Classify.....	4
4.4 Propose.....	4
4.5 Gate and write.....	4
5. Safety rules enforced in code.....	4
6. The Distillery classroom.....	5
7. Isolation and compounding.....	5
8. What is real today, and what is waiting.....	5

1. Abstract

Industrial procedures decay over time. A crew learns a job once, and the official book never changes. Meanwhile, a safer extra step lives only in one expert's head, and a dangerous shortcut quietly spreads from worker to worker. Assess IQ already captures the difference between what the manual says and what really happens. Until Compound IQ, that information was simply thrown away.

This paper describes Compound IQ as it is built today. Every night, it runs a batch job that reads records from Assess IQ (delivered as xAPI statements) and groups them by site, procedure, step, and signature. It then checks each group against a published policy and proposes one of two things: a refinement to the twin, or a short micro-lesson. Nothing is written to Genesis unless Verdict returns a pass.

A few boundaries apply throughout. Distillery can bring in teaching material drawn from public accident reports, but that material can never change a twin. One site can never change another site's book. The practice environment is live today, while the connections to real plant systems are waiting on credentials stored in a password manager.

2. The problem, in plant language

Consider hot work. The written step tells the crew to gas-test the work face. The best crews go further and also test two metres upwind. That extra control isn't in the book, so tomorrow's new hire never learns it.

Lockout works the same way. The book calls for one personal lock, but the best crews add a second lock after proving zero energy. Again, the better practice stays with the people who already know it.

The problem also runs in the other direction. A fire watch leaves when the sparks stop instead of staying thirty minutes longer. A try-start after isolation gets skipped. If shortcuts like these become "how we do it here," someone will eventually get hurt.


None of the usual fixes solve this. Training delivered as a one-time event won't do it, and neither will a brochure. Nor will a model that rewrites the official twin simply because it feels confident. What's missing is an engine that runs every night, with a referee that decides what is allowed to change.

3. Where Compound IQ sits on the Intelligence Flywheel

Compound IQ doesn't replace Genesis, Assess IQ, Brainy, Conductor, or Verdict. Think of it as the feedback station in a loop those products already form. Here is how the pieces fit together:

- **Genesis 3** turns an expert's demonstration into a governed 3D procedure.
- **Field IQ and Brainy** deliver that procedure to workers on smart glasses, and escalate when they're unsure.
- **Assess IQ** verifies what actually happened and writes a discrepancy record.
- **Compound IQ** reads those records overnight and proposes a sharper book, or a warning.


- **Conductor** owns the schedule, running the job compound-iq.nightly-batch at 07:00 UTC.
- **Verdict** is the gate. Without a pass, nothing gets written.



EON calls this completed loop the Human 2.0 OS. Compound IQ earns its place by making the loop compound, getting better each night. It isn't meant to be sold as a standalone app.

4. How one night works

Each nightly run moves through five stages, in order.



4.1 Ingest

Records arrive in the compoundiq.discrepancy.v1 format. When they travel through a Learning Record Store (LRS), they are packaged as xAPI 1.0.3 statements.

If no LRS is configured (LRS_URL is unset), the engine packages the Houston Energy Training Center seed data and sends it through the same xAPI path. That way, the connection being exercised is always the real one. During live ingest, the engine follows xAPI "more" links to page through results, and it retries on 5xx server errors.

4.2 Cluster

When different people on different shifts show the same gap, those records form one cluster. A cluster is defined by four things: site, procedure, step, and delta.signature. The principle is simple: one person doing something once doesn't justify a new official step.

4.3 Classify

The classification policy, `compound-iq.classifier.v1`, is published rather than hidden. Each cluster lands in one of three categories:

- **Unsafe.** A cluster is unsafe if a hold point was skipped, PPE failed, a control was removed, an intervention fired, or Assess IQ flagged a possible `safety_violation`. An unsafe cluster never writes to the twin.
- **Safe.** A cluster is safe if it isn't unsafe and meets the minimum evidence thresholds: 5 occurrences, 3 unique workers, and 2 unique shifts. On a safety-critical step, it must also add something, such as a control, a sample point, a step, or a dwell time.
- **Insufficient evidence.** Anything else falls here, and the cluster is put on hold.

4.4 Propose

What happens next depends on the category. A safe cluster produces a twin refinement along with an SOP diff, which contains actual replacement text rather than a slogan. An unsafe cluster produces a micro-lesson and an alert to the supervisor. Distillery capsules only ever produce lessons.

4.5 Gate and write

Verdict reviews every proposal that would change a twin. Its adversarial reviewers try to show that the proposal removes a control, rests on thin evidence, or skips a hold point.

If Verdict returns a fail, or says a human must decide, the twin stays at its current version. If Verdict returns a pass, Genesis bumps the version number and Brainy pushes the new recipe out to the field.

Rollback works by copying an older version forward as the newest one, so history is never deleted. Compound IQ is the only component allowed to write these changes.

5. Safety rules enforced in code

The following rules are enforced directly in the code:

- The Verdict result must be a pass, and it must match this specific proposal.
- Only two proposal types, `twin_refinement` and `sop_diff`, are allowed to write to a twin.
- The classification must be safe. Unsafe and insufficient-evidence clusters cannot write.
- The evidence must come from Assess IQ field records. Distillery material and cross-customer patterns are for teaching only.
- Every evidence record must belong to the same site as the proposal. Houston cannot write Aberdeen's twin.
- Even a forged "safe + pass" result coming from Distillery is still blocked at the Genesis write gate.

These rules are covered by automated tests; they aren't just a slide. For example, if a test forces a fail verdict, `twins.write()` throws an error and the twin's version stays exactly where it was.

6. The Distillery classroom

On day one, many customers have no "lessons learned" file to draw from. Distillery fills that gap by bringing in public incident reports that are lawful and rights-cleared, and turning them into teaching capsules. Each capsule covers what happened and what to notice or ask about. It never tells anyone what to do as an operating instruction.

If a capsule resembles one of tonight's field clusters, that similarity is treated as a clue, not as proof of cause. Material whose rights are quarantined stays out of the classroom entirely. Compound IQ may publish a lesson based on a capsule, but a capsule can never cause Genesis to change a twin.

7. Isolation and compounding

Each twin is keyed to a specific site and procedure, so evidence from site A can never trigger a write on site B.

There is also an opt-in path for compounding learning across customers, and it is turned off by default. Even when it's switched on, it cannot write to another tenant's twin. Cross-customer compounding is planned as later-phase commercial work, not something that happens silently by default.

8. What is real today, and what is waiting

A good deal of Compound IQ is already working. Today, the following pieces are in place:

- The nightly engine, running in-process
- The published classification policy
- A Verdict mock, plus a live HTTP client for Verdict
- Live clients for Genesis and Brainy, each with a mock fallback
- xAPI packing and unpacking
- The Conductor schedule
- The Distillery classroom
- Per-site isolation
- An overnight console
- Twin versioning and rollback
- A Why / What / How guide
- A go-live board
- Audit export of the most recent batch
- Continuous integration (CI), plus 34 automated tests
- A practice site on the public web

The connections to our live systems will be completed once the secure access details are available. Later, the project will be hosted on EON Enterprise Hub.