

Preventing Incidents Where They Are Born

How the EON Work Intelligence Flywheel makes visible the missing 55% of operational evidence — the half where incidents begin — and compounds every job into a safer next one.




Table of Contents

Executive summary.....	2
1. The facility is not blind — and incidents still happen.....	2
2. Two true timelines that never meet.....	3
3. Where incidents are born: the missing 55%.....	3
4. Defending the numbers: how the 55/45 is derived.....	4
5. The flywheel: how prevention compounds.....	5
6. The economics of prevention.....	6
7. What it prevents.....	6
8. You don't need the whole cake — a phased path.....	6
9. A stacked pilot — prove it one phase at a time.....	7
10. What must be proven.....	7
11. Conclusion.....	8

Executive summary

A modern oil-and-gas facility already collects a great deal about its process — pressure, flow, alarms, historian trends, condition and outcome data. Yet workers are still injured and killed, and facilities still lose tens of millions of dollars a year to unplanned interruptions. The reason is structural: the data facilities collect describes the *consequence* of work, recorded after the fact. It does not preserve the *cause* — the authorized intent, the human judgment, the field execution and the governance where incidents are actually born, and where more than 70% of them originate.

EON models a complete record of one job — a **Verified Work Episode** — as roughly 45% evidence the facility already owns and 55% evidence EON must create. The **EON Work Intelligence Flywheel** manufactures that missing 55%, joins it to the 45%, governs the whole 100%, and — critically — turns each completed episode into a better next mission. The result is a shift from reactive record-keeping to prevention at the point where outcomes are decided.

CENTRAL THESIS

You cannot prevent what you cannot see. Today the evidence that most predicts a serious incident is the evidence no system keeps. EON's opportunity is to make that half visible, verified and reusable — so prevention compounds with every job.




Figure 1. The Work Intelligence Flywheel. Six stations create and compound the Verified Work Episode; EON creates 55% of the evidence, joins the 45% that already exists, and governs the full 100%.

1. The facility is not blind — and incidents still happen

Oil-and-gas operations combine dense instrumentation with constant human observation. Sensors read pressure, flow, temperature, level, gas and vibration. PLCs and RTUs execute local logic; DCS and SCADA let operators supervise plants, pipelines, wells, pumps and valves; historians preserve time-series memory; condition monitoring and outcome systems track asset health, production, quality and downtime. This is a powerful operational nervous system.

It is also, by itself, insufficient. In 2024, upstream oil and gas recorded **32 fatalities across 21 separate incidents**, with explosions, fires and burns causing 41% of the deaths (IOGP, 2024). The average offshore facility loses about **\$38 million a year to unplanned downtime**, and the worst performers exceed \$88 million (Kimberlite). The instrumentation is real, and it is not preventing these outcomes on its own.

2. Two true timelines that never meet

Every governed job produces two records. The process timeline — sensors, alarms, historian — shows what the equipment did. The work timeline — intent, permit, human judgment, physical execution — shows what the people did and why. Today these live in different systems and are rarely reconstructable together. The historian says the process changed; the work order says the job closed. Neither proves the causal story that links them.

What the control room can see (≈45%)	What it usually cannot prove (≈55%)
Pump stopped; pressure and flow changed; alarms appeared; standby equipment carried the process; the asset later stabilized.	Correct physical asset and label; isolation and trapped-pressure check; damage found; part and torque used; procedure revision; worker judgment; safeguards restored.
Captured automatically, after the event.	Mostly uncaptured — and lost when the shift ends.

3. Where incidents are born: the missing 55%

When investigators reconstruct serious events, the root cause is rarely a sensor that failed to read. **More than 70% of oil-and-gas accidents trace to human and organizational factors** — errors in judgment, procedural lapses, gaps in job-safety analysis, weaknesses in permit-to-work — not to the process signals the control room watches (HFACS research). The decisive evidence lives almost entirely inside the half no system preserves.

That missing 55% has a structure, and EON builds each part of it:

- **Human expertise — 15%** (EON Tacit): the cues an expert notices, the option they reject, the moment they stop.
- **Field execution — 18%** (Field IQ / Edge): what physically happened at the equipment — observations, actions, exceptions, photos and readings, even offline.
- **Plan + authority — 12%** (Permit IQ): the authorized job, hazards, isolations, roles, limits and the evidence the work must produce.
- **Simulation + governance — 10%** (Genesis): the approved digital-twin state and the rules under which a lesson may be trusted and reused.

THE CORE PROBLEM

The 45% is the consequence half — it tells you the pressure dropped and the alarm fired, after the fact. It cannot see the cause half: what was authorized, what the worker observed and judged, which procedure revision was truly used, and whether the safeguard was verified. That half evaporates the moment a job closes.

4. Defending the numbers: how the 55/45 is derived

The first question any technical or financial reviewer asks is: *where do these percentages come from?* They are a planning hypothesis — and they are defensible. This section gives the full derivation, the evidence behind each weight, and an honest account of the model's limits.

4.1 What the percentages actually measure

They are not shares of data volume, cost, or headcount. Each is the share of **decision-relevant evidence** in a complete Verified Work Episode — how much a given source contributes to answering the only three questions that make an episode reusable: *did the job succeed, why, and can the lesson be trusted?* Framing each number as evidentiary contribution — not a measured physical quantity — is what makes the model a weighting rather than a claim of fact.

4.2 Why the split is 55 / 45 — cause versus consequence

The existing sensor and control stack captures the *consequence* of work. Incident-causation research consistently attributes the majority of process-safety events to human and organizational factors rather than equipment or sensor failure: HFACS-based analysis of oil-and-gas accidents finds **more than 70% originate in human/organizational factors**. If the majority of what decides an outcome sits in the work-context half, the evidence that half carries must exceed 50%. We set it at 55, not 70, deliberately — the existing 45% is genuinely valuable and necessary, so the split is a conservative majority, not a dismissal of instrumentation.

4.3 Why the sub-weights are ordered as they are

Within each half, sources are ranked by two tests — how often the outcome is decided there, and how uniquely EON captures it. Within the 55% EON creates:

- **Field execution — 18 (highest)**. The physical act is where the plan meets reality and where the largest share of incidents manifest — IOGP 2024 attributes 34% of fatalities to drilling, workover and well operations, with further deaths in maintenance and lifting. It is also the most-uncaptured evidence today.
- **Human expertise — 15**. The judgment layer: the cue noticed, the option rejected, the stop made. High decision value, and essentially absent from enterprise records.
- **Plan + authority — 12**. Permit-to-work and procedure — a recurring named factor in incident investigations (isolation errors, permit failures), and highly structured and improvable.
- **Simulation + governance — 10**. The twin state and the trust rules — weighted lowest because it is an enabler that conditions the other three rather than direct field evidence.

Within the 45% that already exists, sources are ranked by contribution to reconstructing the outcome: process control 15 (the richest continuous signal), outcome systems 10 (the business result), asset condition 8, safety + environment 7, and visual + remote 5.

4.4 Triangulation — three independent checks

(a) The nine categories are the same evidence taxonomy used in EON's Verified Work Episode model, not invented for this paper. (b) The split is consistent with — and conservative against — incident-causation data: if more than 70% of incidents originate in the work-context half, a 55% evidentiary weight could be argued higher, not lower. (c) The weights are treated as **priors to be updated with each customer's data**, never as universal constants.

4.5 How to falsify or calibrate it

The model is testable, and that is the point. Run N episodes; tag every piece of evidence to a category; measure which categories actually resolved “did it succeed and why”; then re-derive the weights from that observed frequency. If, in a given asset class, human expertise resolves more decisions than field execution, the weights change for that class. The 55/45 is a **starting hypothesis, calibrated per customer, facility, asset class and job type** — not a fixed truth.

4.6 Why the exact number does not carry the argument

The strategic conclusion is robust to wide variation. At 50/50, 55/45 or 60/40 the same thing holds: the majority of decisive evidence sits in the half no system keeps, and that is where prevention lives. The argument survives a ± 10 -point swing in any single weight. **The percentages sharpen the story; they do not bear its weight alone.**



Figure 2. Deriving the 55/45 — a plain-English reading and the technical basis. EON-created cause half (55): Field execution 18, Human expertise 15, Plan + authority 12, Simulation + governance 10. Existing consequence half (45): Process control 15, Outcome systems 10, Asset condition 8, Safety + environment 7, Visual + remote 5.

5. The flywheel: how prevention compounds

Capture alone does not prevent anything. The flywheel adds a governed loop that turns each episode into a better next mission. **Assess IQ** measures the delta between the approved plan and what actually happened, and can deliver feedback while the work is still live. **Compound IQ** finds the recurring cue, deviation or failure mode across many episodes. The **Work Intelligence Distillery** filters, validates and packages the lesson under governance, then feeds it back into the digital twin, the permit, the procedure, the training scenario and the next mission package.

THE COMPOUNDING EFFECT

Each turn of the wheel, the next worker arrives already knowing the cue that preceded last year's near-miss. The permit carries a new checkpoint. The procedure closes the gap that caused the last deviation. Safety stops being a poster and becomes a property of a system that improves every cycle.

6. The economics of prevention

The industry's own data shows a data-driven approach already reduces unplanned downtime by about **36%** (Kimberlite). But that gain comes from optimizing the 45% that is already visible. The larger, still-untapped prize is the 55% where more than seven in ten incidents begin. Reducing loss there attacks the root cause rather than the symptom — and it is reachable now, because 45% of the evidence already exists and EON supplies and governs the rest. The question a facility should ask is not whether to buy more software, but how much of the reducible 55% loss — in lives and in dollars — it wants to remove, and in what order.

7. What it prevents

Moving safety upstream — to the point where incidents are born — changes what an organization can expect from every job:

- **Fewer fatalities and serious incidents** — by catching the human, procedural and authority conditions that cause most of them, before they combine.
- **Fewer unplanned interruptions and lower downtime cost** — by proving physical work was done right the first time, cutting repeat visits and ambiguous closeouts.
- **Expertise that no longer walks out the door** — the retiring expert's judgment captured, verified and reusable.
- **Faster, more confident execution** — the crew arrives with asset context, evidence requirements and relevant experience already in hand.
- **A digital twin that gets smarter every cycle** — not only what the asset is, but how work on it actually succeeds.

8. You don't need the whole cake — a phased path

Because 45% of the evidence already exists, EON can be added in phases that each remove a measurable share of the 55% loss — no single large program required. The sequence is ordered by value-per-effort, so wins are felt before cost is.

Phase 1 — Foundation: Genesis first, Universal to scale

Genesis stands up the twin — the “what good looks like” of simulation and governance (10%). It is the best place to start, and it assumes the customer already has models or facilities to work from. **Universal** is the scaling enabler: it turns twin-building from a manual, semi-manual effort into a repeatable one. Without it, every later phase is hand-built and will not scale across sites and asset families.

Phase 2 — Low effort, high return: Tacit + Permit

This phase front-loads **27 of the 55 points at low effort**. EON Tacit (15%) is mostly structured expert interviews — accelerated by EON's synthetic data and feeding Situation Room — capturing the human-expertise half before it retires. Permit IQ (12%) digitizes plan and authority; it is a painkiller rather than a vitamin, because crews already suffer permit-to-work and form-filling friction. Together they capture the majority of the cause half at the lowest implementation cost in the plan.

Phase 3 — The big block: Field execution

Field IQ / Edge (18%) is the single largest slice of the 55% and the highest-effort to deploy, because it involves edge capture and a live connection to EON Live. The largest recoverable value is here, so it should not be left for last — but a full deployment before Phases 1–2 are proven would re-create the very overwhelm the phased approach avoids. The recommended move is to **pull a thin, single-crew Field IQ Edge pilot forward**, capturing the highest-value signal early at contained effort, then scale it once proven. Assess IQ, Compound IQ and the Distillery then close and compound the loop.

9. A stacked pilot — prove it one phase at a time

The phases map directly onto a low-commitment commercial path. Each tier stacks on the one before and “eats” a further share of the 55%, so the customer expands only as value is proven. Prove it first where ground truth is known and no one is at risk — a training center with physical skids, a simulated process and coordinated control-room and field teams can manufacture verified episodes, including rare and abnormal conditions, before any move to bounded, customer-controlled field pilots.

Tier (each stacks on the prior)	What it unlocks · illustrative pricing
Tier 1 · Foundation — Genesis	Stands up the twin (10% of the 55%). Illustrative from ~\$50K (published Genesis Pilot list price).
Tier 2 · Scale — + Universal	Makes twin-building repeatable so later phases scale. Illustrative + ~\$100K.
Tier 3 · Cause-half context — + Tacit & Permit	Adds human expertise (15%) and plan + authority (12%) — cumulatively 37% of the 55%. Illustrative + ~\$250K.
Tier 4 · Field execution — + Field IQ / Edge	Adds the largest slice (18%) and closes the loop — the full 55% captured. Scoped per engagement.

Against ~\$37,500/day of recoverable loss on an average facility, **each tier pays for itself in days, not quarters**. Dollar figures are illustrative, built from EON's published list pricing; final pilot scope and price are confirmed per engagement.

10. What must be proven

The 55/45 split is a hypothesis, and the prevention claim must be earned, not asserted. A credible pilot should quantify both the immediate job value and the quality of the learning asset created: correct-asset and correct-step rates, first-pass completion, evidence completeness, closeout time and repeat visits; recognition, escalation and stop-work quality; and episode verification rate, reusable lessons produced, and recurrence of the failure modes those lessons address. Each release should be benchmarked to show it measurably improves the next mission.

11. Conclusion

Physical AI will not be won only by the most capable machine or the strongest model. It will be won by whoever can supply the trusted, enterprise-specific knowledge required to perform complex work safely. That knowledge lives in the 55% no system keeps today. EON's Work Intelligence Flywheel exists to make it visible, verify it, govern it, and compound it — so that prevention happens where incidents are born, and every completed job makes the next one safer.

CLOSING PRINCIPLE

The 45% tells you what happened. EON adds the 55% that tells you why — and compounds both into the one thing that prevents the next incident: a verified, governed lesson delivered to the next crew before they start.

NOTES AND SOURCES

- IOGP 2024 safety data — 32 fatalities in 21 incidents; fatal accident rate 0.77 per million hours; explosions/fires/burns = 41% of fatalities.
- Kimberlite oil & gas research — average offshore facility $\approx$ \$38M/year unplanned downtime, worst performers > \$88M/year; data-driven approaches reduce downtime $\approx$ 36%.
- HFACS-based analysis of oil-and-gas accidents — human/organizational factors implicated in > 70% of accidents; procedure, job-safety-analysis and permit-to-work among top contributors.
- The 55/45 split and per-category weights are EON's planning-hypothesis model of a complete Verified Work Episode, to be validated by customer, facility, asset class and job type — not audited statistics.