



EON WORK INTELLIGENCE

The compounding intelligence platform for physical
work

“Every completed task can strengthen the next one.”



The model can be replaced. The Work Intelligence cannot.

EON's category is not training software and not a frontier-model lab. It is the **governed mission and intelligence layer** for physical work.

INDUSTRIAL WORK INTELLIGENCE



Revenue proves demand. Compounding Work Intelligence determines how defensible, scalable, and strategically valuable that revenue can become.

Physical work still depends on fragmented and vanishing expertise

Knowledge exists across documents, systems, and experts—but often not in trusted form at the point of work.

CONTEXT

Which asset, condition, and mission?

The field needs specific, approved guidance.

EVIDENCE

What happened, and did it meet the standard?

Execution must become observable and reviewable.

1 MISSING LAYER

Work Intelligence

Generic AI lacks operational grounding; traditional systems rarely close the loop from work to approved learning.

EON connects asset context, approved work, field evidence, assessment, and governed improvement.

Enterprise footholds create access; expansion creates compounding value

PAID ENTERPRISE FOOHOLDS



MULTI-SITE

potential across people, assets, sites, and missions

Opportunity depends on approved scope; this is not a claim of full-platform adoption.


VALUE-CREATION LOGIC

The foothold is the beginning, not the value conclusion.

Sustainable value comes from daily use, validated evidence, governed learning, measurable outcomes, and recurring expansion.

EON grounds AI in specific assets, work, behavior, and outcomes


A governed world model is not one enormous neural network. It is an **auditable system** of **records, graphs, simulations, procedures, evidence, and replaceable reasoning models**.



Exact **operational truth** remains in Customer Digital Twins and versioned repositories. **Models reason** over that memory; **governance** determines what may become trusted.


The integrated system—not the foundation model—is the moat

Switch the model. Keep the memory, evidence, mission architecture, evaluations, distribution, and customer relationships.



Every completed mission can improve the next one


The field never rewrites operational truth directly. Measurement, learning, governance, and signed release remain separated.



The proof of the category is simple: the second mission should be better because the first mission happened.

The **Validated Work Episode** is the atomic unit of intelligence

Raw images, readings, or procedure clicks are not enough. A VWE connects operational context to physical consequence and assessment.



context → action → evidence → physical consequence → assessment → approved Knowledge

A bounded deployment begins with disciplined proof

Start narrow, define acceptance criteria, measure the mechanisms, then expand.

ONE

MISSION FAMILY

ONE

LIVE FACILITY

VALIDATED WORK
EXPANSION




MISSION QUALITY AND LEARNING



The first proof is one asset family, one mission, one execution, one approved delta, and one improved package returned to the field.

Usage creates a larger learning surface over time

More approved missions can create more evidence, better evaluations, and faster reuse—when governance and outcomes improve with scale.



CUSTOMER DEPTH

MORE

sites • assets • missions • evidence

PLATFORM BREADTH

FASTER


sites • assets • missions • evidence

The intelligence asset earns value only when it improves the next mission and the economics of the next deployment.


Governed learning—not raw data volume—creates enterprise value

The Distillery must deduplicate, measure novelty, classify ownership, and release only approved improvements.

CUSTOMER-PRIVATE LEARNING



AUTHORIZED GENERALIZATION




Customer-private operational truth stays private. Reuse requires explicit rights, technical validation, provenance, and governance approval.

One intelligence foundation expands across four executors

The mission still requires identity, context, workflow, safety, evidence, stop conditions, and assessment.


- 1 **HUMAN** users • sites
- 2 **AI AGENT** workflows • tools
- 3 **ROBOT** endpoints • fleets
- 4 **DRONE** routes • coverage

Commercial surface expands from seats to endpoints, fleets, assets, missions, and verified outcomes.



Customer ROI is a portfolio, not a heroic claim


A 90-day pilot should prove leading mechanisms—not claim ownership of every safety or downtime outcome.



Illustrative relative contribution index; every customer must build its own baseline and benefit model.

The commercial model expands with customer value

Fixed recurring platform economics combine with variable usage or outcome components.



COMMERCIAL SURFACE

PLATFORM +

users • assets • sites • missions • intelligence services


Scope and pricing vary by deployment

VALUE-ALIGNED MODEL

Recurring platform + Digital Twin coverage + mission packages + governance + machine usage / verified outcomes

Category proof changes how the platform is understood

EON combines Digital Twins, field execution, governed learning, and a pathway to physical AI.



CATEGORY DIFFERENTIATION

PROOF

comes from customer outcomes, recurring expansion, deployment reuse, governance, and a closed learning loop

PLATFORM EVOLUTION

Immersive learning



Work Intelligence

↓ with proof

Bounded machine execution

Observable proof drives the EON value-creation staircase


Customer outcomes + recurring expansion + compounding intelligence + disciplined governance



Each stage must preserve the evidence and trust earned in prior stages. Expansion follows proof; it does not substitute for it.

One closed loop creates the largest near-term re-rating

The Minimum Complete Flywheel is narrow by design: one facility, one asset family, one critical mission, one execution, one measured delta, one approved change, one better package.



90–120 day design partnership: prove identity, workflow, offline execution, evidence, stop behavior, delta approval, redeployment, security, tenancy, ownership, and auditability.

Execution confidence rises only through proof and governance

BEFORE COMPLETE-LOOP PROOF

HIGH

execution uncertainty

Confidence can rise after one customer closes the loop—and again when another deployment becomes faster, repeatable, and measurable.

Product integration	One frozen hero asset + end-to-end acceptance test
Customer conversion	Paid design partner + sponsor + expansion option
IT/OT + cybersecurity	Private deployment + signing + isolation + evidence
Learning rights	Learning Rights Schedule + enforceable Knowledge Firewall
Safety + model performance	Bounded missions + deterministic gates + rollback
Services + concentration	Reuse metrics + multi-customer pipeline + no exclusivity

Risk does not disappear because the architecture is compelling; it falls as the **evidence becomes operational, commercial, repeatable, and governable.**

Watch operating evidence—not model benchmarks

Every investor metric should answer one of three questions: Is it used? Is it improving? Is it scaling economically?

- | | | |
|---|--------------------------------|--|
| 1 | Distribution | active workers • machines • sites • mission families |
| 2 | Intelligence production | VWEs Validated Work Episodes • candidate deltas • approvals • time-to-release |
| 3 | Product quality | identity accuracy • package success • evidence completeness |
| 4 | Customer outcomes | task time • rework • escalations • safety interventions |
| 5 | Deployment economics | days to first value • reuse % • implementation hours |
| 6 | Commercial | ARR • backlog • gross margin • NRR • expansion • concentration |
| 7 | Autonomy | machine missions • fleet reuse • safe completion rate |
| 8 | Governance | ownership exceptions • rollback • audit findings • incidents |

The intelligence asset earns a premium only when these metrics improve together.

The second mission should be better because the first happened.

When that loop is proven in real operations, enterprise footholds can expand into a compounding Work Intelligence platform—and the same governed mission architecture can progress from people to agents, robots, and drones.

The model can change. The operating memory, evidence, governance, mission architecture, and customer relationships remain.

External platform overview. Product capabilities, availability, scope, and results vary by deployment. This presentation does not identify or imply endorsement by any customer.

