

Work Intelligence

The Understanding Layer for the Physical World



Understand the Work. Compose Any Facility.

Genesis shows the steps. EON Universal understands the work.

The frontier of artificial intelligence has moved from bits to atoms — from the public internet to the physical world of work. At the same moment, the model itself has become a commodity. This paper explains why the durable value now sits in proprietary, understood work data, and how EON AI Ventures built the only platform that turns what a worker sees into intelligence the enterprise owns.

Table of Contents

Executive summary	3
1. The internet was the small dataset	3
2. The model just became a commodity	4
3. The trend everyone sees — and the layer nobody owns	4
4. What EON built: top-down capture meets a bottom-up library	5
The six-layer competence record.....	6
5. The closed loop: capture, understand, act, preserve	6
6. Why this is different	7
7. The dual moat	7
8. Why EON, and why now	7
9. The product	8
Notes and sources.....	9

Executive summary

For two decades, artificial intelligence learned the internet. That corpus — text, images, code — is extraordinary, but it is a thin slice of reality. It is the world reduced to bits. The far larger world, the one where value is actually created and risk actually lives, is physical: the torque on a bolt, the sequence of a lockout, the reason a pump fails, the way an experienced technician actually does the job. That world is made of atoms, and almost none of it is captured as intelligence. It has been scanned, photographed and documented — but documentation is not understanding.

Two shifts in 2026 make this the decisive moment. First, the most credible voices in the field now agree that physical AI is the next wave. Second, and more consequentially for strategy, the model layer has commoditized: open-weight models now rival the frontier at a fraction of the cost, and the industry consensus has moved decisively — when the model is free, the moat is the proprietary data and the learning loop wrapped around it.

EON AI Ventures occupies the layer no one else has built: **the understanding layer for work**. A human wearing glasses is the sensor that captures the physical work as it happens. Underneath sits a bottom-up universal library that already knows every component — what it is, how it is built, how it behaves, how it fails, how it is serviced. Capture meets comprehension; comprehension drives action; action is preserved as the living intelligence of the facility. Intelligence flows out of the physical world and back into it.

This is a category, not a feature. EON calls it Work Intelligence — the fourth enterprise system of record alongside ERP, CRM and PLM. The company brings 25 years of physical-world asset intelligence, the only universal component library of its kind, and validation from two of the largest energy companies on earth — at the exact moment the rest of the market discovered that the physical layer is where the value went.

1. The internet was the small dataset

The first era of AI was built on the assumption that the internet contained the world. It does not. It contains what people wrote down. The operational knowledge that runs refineries, mines, aircraft hangars and hospitals was never written down in any form a model could learn from. It lives in procedures locked in PDFs, in tribal knowledge inside the heads of a retiring workforce, and in the muscle memory of people who have done the job ten thousand times.

The standard response has been to digitize: scan the equipment, photograph the asset, convert the binder to a tablet screen. But a scan is geometry without meaning. A photograph is pixels that do not know what they are looking at. None of it tells a system what a component is, why it matters, how it behaves under load, or what to do when it fails. The gap between data about the physical world and intelligence about the physical world is the entire opportunity.

This is the **“atoms, not bits”** thesis: physical-world operational expertise cannot be captured by public-corpus AI, because it was never in the public corpus to begin with. Whoever captures it,

structures it, and makes it queryable owns something a foundation model can never reproduce from the internet.

2. The model just became a commodity

The strategic ground shifted under the entire industry in the first half of 2026. Open-weight models from Chinese laboratories closed the gap with the Western frontier and gave their weights away. In June, Zhipu AI released GLM-5.2 — a 744-billion-parameter model with a one-million-token context window — completely free under a permissive license, and within hours it topped a leading reasoning benchmark at roughly one-tenth the cost of its closed counterpart. It is not alone: DeepSeek, Alibaba's Qwen and Moonshot's Kimi sit in the same band, matching or beating mid-tier Western models at five to thirty times lower cost, downloadable and self-hostable by anyone.¹

The implication has been absorbed by the people who allocate capital. Gartner now classifies foundation models as “strategic commodities.” Bessemer's 2026 infrastructure roadmap states plainly that as models commoditize, differentiation shifts to the layers that orchestrate them. Andreessen Horowitz's outlook for the year argues that as model capability commoditizes, scarcity moves from the model to the data, and the companies building walls with proprietary data are the ones posting breakout growth. Microsoft's chief executive framed the same point as “token capital”: firms that merely rent models risk having their expertise commoditized, while firms that own a learning loop compound an advantage.²

The conclusion is unavoidable. **The model, once the moat, is now the commodity input.** The value migrates to whoever owns proprietary data the model cannot get anywhere else — and the richest such data is the physical work the internet never saw.

3. The trend everyone sees — and the layer nobody owns

Physical AI is no longer a thesis on the horizon. At CES in January 2026, Nvidia's Jensen Huang declared that “the ChatGPT moment for physical AI is here — when machines begin to understand, reason and act in the real world,” and sized the physical opportunity in the tens of trillions of dollars. The whole industry has turned toward the world of atoms.³ But it has turned toward adjacent layers — and left the center empty.

Map the field and the white space is obvious. Each layer below has real, well-funded players. Only one layer has no one in it.

Layer	Who is building it	What it lacks
Bodies — robots, AVs, humanoids	Nvidia Isaac/GR00T, Boston Dynamics, Tesla, Physical Intelligence	No knowledge of how the work is actually done
Brains — world models	Google DeepMind, World Labs, General Intuition	Spatial reasoning, but not equipment-specific procedure
Eyes from above — capture	Planet (public; daily imaging of Earth's surface)	Captures the surface; analytics bolted on after the fact
Eyes on the worker — capture	Amazon “Amelia” glasses; RealWear, Vuzix hardware	Capture and delivery, with no understanding model beneath
Digitizers — connected worker	Augmentir, Tulip, Parsable, Poka, Dozuki	Digitize the procedures you already have; no model of the equipment
Understanding layer for work	EON Universal	— the lane no one else occupies

The connected-worker platforms deserve a specific note, because at a glance they look closest to EON. They are not. They are authoring tools: they take the procedures a customer already has — in Excel, Word, PDF or video — make them digital, deliver them to a glass or tablet, and capture how the work went. The intelligence is the customer's, carried in as documents. The platform is the pipe, the personalization and the analytics. None of them brings a pre-built understanding of the equipment itself.⁴

4. What EON built: top-down capture meets a bottom-up library

EON's architecture is the marriage of two things that have never been combined before.

Top-down capture. The sensor is not a satellite and not a fixed camera. It is a human wearing glasses, capturing the work from the inside, at eye level, in the moment it happens. This is the same trend now visible everywhere — from Amazon putting heads-up glasses on its delivery associates to the data flywheel that makes a fleet smarter every mile it drives. The worker is the sensor for work.

Bottom-up understanding. Underneath the capture sits the Work Intelligence Library — a universal, pre-built model of the components that make up industrial facilities. When the worker looks at a pump, EON does not merely see pixels of a pump. It knows it is a centrifugal pump: it knows its anatomy, how it behaves under load, the ways it fails, and the procedure to service it. The capture *resonates* with the library. That match — what you see meeting what we already know — is the product.

The library is built by composition, not enumeration. Rather than attempting to model every machine that exists — an infinite task — EON models roughly fifty component classes that, in combination, account for more than eighty percent of real field interactions. Any facility is then composed from these classes rather than enumerated machine by machine. This is what makes the platform facility-agnostic: it does not need to have seen your plant to understand it.

The six-layer competence record

Every component class carries a six-layer record. Together these layers are the difference between a 3D asset and an understood one:

- **Identity** — what the component is, and the standards taxonomy it belongs to.
- **Geometry** — its form and spatial structure.
- **Anatomy** — its parts and how they fit together.
- **Function** — what it does and why it is there.
- **Behavior** — how it acts under load, and how it fails.
- **Procedure** — how it is operated, maintained and repaired.

The taxonomy anchors are real, recognized standards, vertical by vertical: ISO 14224 for oil and gas, ISO 55000 for mining, ATA iSpec 2200 for aerospace MRO, and GMDN / UMDNS (ISO 15225) for healthcare biomedical. The library speaks the language each industry already audits against.

5. The closed loop: capture, understand, act, preserve

Capture that is understood can do something capture alone never could: it can close the loop back into the physical world.

- **Capture.** The worker, wearing glasses, records the work as it happens — the true state of the asset, not a document about it.
- **Understand.** The capture is matched against the library and becomes comprehension: this component, this condition, this step.
- **Act.** Because the system understands the work, it guides and assists the worker in real time to operate, maintain or repair — correctly, the first time.
- **Preserve.** Every action enriches the facility's living record, and feeds back into the library. The asset accumulates intelligence; the platform gets smarter with every use.

This is the decisive contrast with capture-only businesses. Earth-imaging proved that capturing the physical world at scale is a real, public-market category — but it images the surface and adds analytics afterward. EON's loop does not end in a dataset. **It ends in changed physical reality and a compounding asset.** Intelligence is mapped out of the physical world and injected back into it.

6. Why this is different

The single most important question a serious evaluator will ask is: how is this not a tool we have already seen? The honest, defensible answer is one sentence. Everyone else digitizes the procedures you already have. EON already understands the equipment. It is the difference between an authoring tool and a knowledge model.

Compared with	What they give you	What EON adds
CAD / 3D asset libraries	Geometry — accurate shapes with no meaning	The other five layers: function, behavior, failure, procedure
LMS / courseware	Generic training, decoupled from the live asset	Guidance bound to the specific component in front of the worker
Process simulators	Modeled behavior of one designed system	A universal library that composes any facility
Connected-worker platforms	Your procedures, digitized and delivered	Pre-built understanding of the equipment, before any document

7. The dual moat

Because the model is now a commodity, the defensibility question matters more than ever. EON's answer is two moats that reinforce each other, and that a free model cannot replicate.

- **The universal library (horizontal).** The bottom-up Work Intelligence Library — fifty component classes, six layers each, built over 25 years of physical-world asset work — is the cross-customer asset. It is what makes capture meaningful on day one, and it deepens with every deployment across the network.
- **The facility record (vertical).** Every customer accumulates its own living intelligence about its own plant — proprietary, irreplaceable, and compounding with use, exactly the data flywheel that makes a fleet or a warehouse smarter over time. The operator owns this. No vendor and no foundation model can recreate it from outside.

This is precisely the structure the moat literature now says wins: proprietary data plus the software that makes it indispensable in the user's actual work. A freely downloadable model has neither the library nor the capture, and cannot acquire either by being clever. It is the commodity engine; EON is the data, the understanding and the loop around it.

8. Why EON, and why now

EON did not pivot into this position; it spent a quarter of a century arriving at it. The lineage includes more than 4,400 institutional customers across 80-plus countries, over 136 million

platform downloads, and a 9,000-course virtual campus — a body of physical-world, asset-level intelligence that has no public-corpus equivalent.

That heritage now meets the market at the precise moment the value moved. Two of the largest energy companies in the world are already engaged: a global oil-and-gas supermajor, with a signed proof of concept and a train-the-trainer deployment underway, and a national energy company in advanced commercial engagement with in-country hosting confirmed. When the most safety-critical, most conservative operators on earth move first, the category is real.

The timing is not a coincidence. The model went free in the same months the industry agreed that physical AI is the next wave and that proprietary data is the only durable moat. **The understanding layer for work is the asset that satisfies all three at once — and EON is the platform the enterprise builds it on.**

9. The product

The thesis is delivered through a single, coherent product spine:

- **Genesis** authors the simulation and shows the steps — the procedural backbone.
- **EON Universal** is the understanding layer — the universal library that turns capture into comprehension, facility-agnostic by composition.
- **Field IQ** delivers guidance and capture at the point of work, on the glass, in the worker's hands.
- **Assess IQ** verifies competence — confirming the work was understood and done correctly.

These run on the Intelligence Flywheel: each turn of capture, guidance and verification feeds the next, so the system — and the facility's own intelligence — compounds over time. Genesis shows the steps. EON Universal understands the work.

Understand the Work. Compose Any Facility.

Notes and sources

1. GLM-5.2 release and benchmark results: Zhipu AI / Z.ai, June 2026, reported across industry coverage (e.g. explainX, June 2026). Comparative cost and capability of Chinese open-weight models (DeepSeek, Qwen, Kimi): industry benchmark roundups, April–June 2026.
2. Foundation models as “strategic commodities”: Gartner, 2026. Differentiation shifting to orchestration layers: Bessemer Venture Partners 2026 AI infrastructure roadmap. Scarcity moving from model to data: Andreessen Horowitz, Big Ideas 2026. “Token capital”: Satya Nadella, June 2026 essay, as reported in market commentary.
3. “The ChatGPT moment for physical AI is here...”: Jensen Huang, NVIDIA, CES, January 2026. Physical-opportunity sizing: Huang public remarks, 2026.
4. Connected-worker platforms generate digital work instructions from Excel, Word, PDF, images or video: vendor product documentation and independent comparisons, 2026 (e.g. Augmentir, Tulip, Parsable, Poka, Dozuki).
Customer references are described in general terms in this document; named attribution requires customer communications approval prior to external use.