

EON Universal: Facility-Agnostic Field Intelligence

How a composable equipment ontology turns any worker into an expert




Table Contents

- 1. Executive Summary..... 3
- 2. The Problem: The Expertise Cliff..... 3
- 3. The Vision: Any Worker, Any Facility..... 4
- 4. The Core Insight: Composition Over Enumeration..... 5
- 5. The Component Ontology..... 6
 - 5.1 Anchoring to an industry standard..... 6
 - 5.2 The EON Universal six-layer competence record..... 6
 - 5.3 Coverage: how many components..... 7
- 6. The Architecture: The Intelligence Flywheel..... 7
 - 6.1 The four stages..... 8
 - 6.2 The two governing layers..... 8
- 7. How It Works in the Training Center and the Field..... 8
- 8. Realism: What Is Deliverable Today..... 9
 - 8.1 Deliverable now..... 9
 - 8.2 Where the thirty-year claim must be qualified..... 9
- 9. Deployment Realities and Guardrails..... 10
- 10. The Strategic Frame..... 10
 - 10.1 Makers versus takers..... 10
 - 10.2 Work Intelligence as a system of record..... 11
- 11. Roadmap..... 11
- 12. Conclusion..... 11

1. Executive Summary

Heavy industry is losing its experts faster than it can replace them. The engineers and operators who hold thirty years of judgment about how a compressor sounds when a bearing is going, or what a separator does when the interface level drifts, are **retiring** — and the **people arriving to replace them are green**. In a high-consequence environment, that gap is not merely a productivity problem; it is a **safety and reliability problem** measured in unplanned **downtime, incidents, and lost production**.

This paper describes a **practical, deployable answer**:

1. a system that **recognises** *any piece of equipment* a worker is looking at,
2. **understands** *what it is and how it behaves*, and
3. **guides** a low-experience *worker through the correct procedure* with the confidence of a veteran.
4. The worker holds a **FieldIQ Tablet** or wears **FieldIQ** — *assisted-reality glasses*
5. backed by **Genesis**, *the engine that captures and structures expert knowledge* into reusable, machine-readable competence.

The central claim of this paper is an engineering one, and it is what makes the vision tractable rather than aspirational: **a facility is not a monolith to be memorised** — *it is a graph assembled from a finite library of known components*. You do not program five thousand locations.

You **program the vocabulary the locations** are built from — roughly *fifty equipment classes* — and the system composes guidance for whatever configuration it is looking at, including configurations it has never seen.

The bottom line

- The **component library** to start is *~50 equipment classes — not 5,000 locations*. The top 15 classes carry most of the daily field value.
- The **recognise-and-guide** capability is deliverable on today's AI. It realistically *compresses a green hire* toward a competent multi-year technician on routine and many non-routine tasks.
- Full **veteran judgement** in novel upset conditions is an asymptote the *system approaches over time* through a compounding *data flywheel* — not a switch flipped on day one. Designing for that honestly is *what makes the system safe*.

2. The Problem: The Expertise Cliff

Three forces are converging across **energy, manufacturing, aerospace, and process industries**:

1. **Demographic loss.** A generation of *operators and maintenance engineers is retiring*, and with them goes tacit knowledge that was never written down — the diagnostic intuition built over decades of contact with real equipment.
2. **Procedural fragmentation.** What *knowledge* is documented lives in standard operating procedures, P&IDs, vendor manuals, and reliability databases that are *scattered, inconsistent across sites, and rarely available at the point of work*, in the field, at the moment a decision must be made.
3. **Configuration explosion.** A large operator may run thousands of facilities, each individually engineered. *No two are identical, so site-by-site training does not scale*, and a worker proficient at one site is not automatically proficient at the next.

The cost of this gap is concentrated exactly where it hurts most: in **hazardous, high-consequence environments** where a *green worker doing the wrong thing* is not an inconvenience but a *risk to people, assets, and production*.

The **traditional mitigations** —

1. long **apprenticeships**,
2. sending **scarce experts to travel** between sites,
3. **conservative work** scheduling

are **slow and expensive**, and they *do not survive the retirement wave*.

3. The Vision: Any Worker, Any Facility

Picture a **new technician** who has been on the job for three weeks, sent to a facility neither they nor the company has worked before. They put on a **pair of glasses or use a tablet**. The system **recognises**

- a) **the equipment** in front of them — that this is a *horizontal three-phase separator*,
- b) that the **valve** on the discharge line is a *pneumatically actuated control valve*,
- c) that the **instrument** on the bridle is a *radar level transmitter*.

It knows:

- I. **what** each *does*,
- II. **how** each *behaves*, and
- III. what the **correct procedure** is for the *task at hand*.
- IV. It **guides** the worker *step by step*,
- V. **watches** what they *do*,

- VI. **flags** when something looks *wrong*, and
- VII. **escalates** to a human when the *situation exceeds what it can safely advise*.

The worker is green. The guidance is not. The vision is to put thirty years of structured expertise behind a person who has none — and to do it in any facility, including one the system has never seen, by recognising the parts even when the whole is new.

4. The Core Insight: Composition Over Enumeration

The **naive approach** to *any facility* is to **model each facility**. That *path is hopeless*: the permutations across thousands of individually engineered sites are *effectively infinite*, and every new site would demand a new modelling effort.

The **correct approach** inverts the problem. A process facility, however complex, is an *arrangement of a finite set of equipment classes* —

- a) **pumps**
- b) **compressors**
- c) **separators**
- d) **exchangers**
- e) **valves**
- f) **instruments**
- g) connected by **piping** and
- h) **instrumentation.**

If the **system knows** each class deeply:

- a) **what it is**
- b) **what it is made of**
- c) **what it does**
- d) **how it behaves**, and
- e) **what to do to it**

then a never-before-seen facility is simply a **new graph over known nodes (P&ID)**. Recognise the nodes, infer the topology, compose the guidance.

Why this is the whole game

Enumeration scales with the number of facilities — an *unbounded, ever-growing problem*.

Composition scales with the number of component classes — a *bounded library of ~50 to start, ~300 for breadth*.

Every **new facility makes the system smarter** without making the modelling task bigger. That is the *difference* between a *project that ends* and a *product that compounds*.

5. The Component Ontology

5.1 Anchoring to an industry standard

The component library is not invented from scratch. It is anchored to **ISO 14224**, the **petroleum, petrochemical** and natural-gas **industry standard** for collecting and exchanging reliability and maintenance data. ISO 14224 defines a *nine-level equipment taxonomy*; equipment classes —

- a) centrifugal pumps,
- b) gas turbines,
- c) heat exchangers,
- d) valves,
- e) compressors
- f) **sit at level seven**, each with a defined boundary specifying its subcomponents.

Anchoring to ISO 14224 (and to the CFIHOS reference-data library for attributes) is a **strategic decision**, not merely a convenient one. The **major operators** already structure their **CMMS** and reliability data against this taxonomy.

When **Genesis** recognises a component, its output **maps directly** into the customer's **existing functional-location** and maintenance-**data structures**. The system **speaks the customer's data language** on day one.

5.2 The EON Universal six-layer competence record

For each equipment class, the **EON Universal captured competence** is complete when **six layers** exist. This is the **ingestion checklist** that defines when a component is fully ‘known’:

Layer	Content	Primary source
1 — Identity	ISO 14224 <i>class</i> , aliases, <i>tag-naming</i> pattern	<i>Nameplates / labels</i>

Layer	Content	Primary source
2 — Geometry	<i>3D model</i> and multi-angle imagery for recognition	<i>3D + photo corpus</i>
3 — Anatomy	<i>Subcomponents</i> / maintainable items and boundary	Engineering <i>teardown</i>
4 — Function	<i>What it does</i> and its normal operating envelope	<i>Foundational skill</i> capture
5 — Behavior	<i>How it responds</i> to inputs and to upset conditions	<i>Simulator</i>
6 — Procedure	<i>Operate</i> , inspect, isolate, troubleshoot (SOPs)	<i>SOP library</i>

Recognition keys on *layers one and two*. *Guidance* composes from *layers three through six*.

5.3 Coverage: how many components

Coverage is tiered. The **first tier** is small enough to **build now** and large enough to matter:

Tier	Scope	Coverage of field interactions
Tier 1 — ~50 classes	Core equipment classes a technician physically touches	80%+ of real interactions
Tier 2 — ~300 types	Variants and subtypes for robust breadth	Rarely surprised
Tier 3 — ~1,000–2,000 items	Maintainable-item level (seals, bearings, actuators)	Long tail , accreted over time

The practical number to build first is roughly **fifty equipment classes**; the top fifteen:

1. **pumps,**
2. **compressors,**
3. **separators,**
4. **exchangers,**
5. **control valves,**
6. **relief valves,**
7. **shutdown valves and**
8. **the core instruments**

carry most of the daily value. Everything beyond that is the flywheel doing its job.

6. The Architecture: The Intelligence Flywheel

The component ontology is operated by four stages that form a self-reinforcing loop, governed by two cross-cutting layers.

6.1 The four stages

1. **Genesis** — *captures expert knowledge* and structures it into the six-layer competence record for each component. This is where proprietary expertise becomes machine-readable competence.
2. **Field IQ** — *delivers that competence at the point of work* through assisted-reality glasses: recognition, step-by-step guidance, and in-situ verification.
3. **Assess IQ** — *measures whether the worker performed correctly* and whether the guidance was right, closing the loop between instruction and outcome.
4. **Compound IQ** — *feeds verified field outcomes back into Genesis*, so every job at every facility improves the competence library. This is the compounding mechanism.

6.2 The two governing layers

- **EON Conductor** — *the orchestration layer* that routes each task to the right model and the right knowledge, so the system uses the best available reasoning for the situation in front of it.
- **EON Verdict** — *the trust and verification gate*. Verdict decides when guidance is confident enough to give to a green worker and when the system must hold back and escalate to a human. In a high-consequence environment this gate is not optional; it is the difference between a helpful system and a dangerous one.

7. How It Works in the Training Center and the Field

The green-worker journey composes the layers in sequence:

1. **Recognise.** *The glasses identify the equipment class* and read nameplates, tags, and gauges. Recognition draws on the identity and geometry layers.
2. **Locate in context.** *The system infers how the recognised components connect* — the local topology — from what it sees combined with P&ID logic, so it understands the equipment's role, not just its name.
3. **Compose guidance.** *Drawing on the anatomy*, function, behavior, and procedure layers, the system *assembles the correct step-by-step procedure* for the task — even in a configuration it has not seen before.
4. **Verify and watch.** *As the worker acts, the system checks* that what it observes matches what it expects, and flags deviations.
5. **Escalate when uncertain.** *Where confidence is low or the situation is safety-critical*, Verdict stops and *routes to a human* rather than guessing.
6. **Learn.** The verified *outcome flows back through Compound IQ, improving the competence* library for the next worker at the next site.

8. Realism: What Is Deliverable Today

8.1 Deliverable now

The following stack is achievable on today's AI and constitutes the deployable core of the system:

- **Visual recognition of equipment class and type** from a *labelled 3D and photographic corpus*, with nameplate, tag, and gauge reading.
- **Local topology inference** from what *the tablet or glasses see* combined with engineering logic.
- **Retrieval and composition** of the *correct procedure, grounded* in the captured SOP library.
- Simulator-grounded understanding of how equipment behaves, so guidance reflects cause and effect, not just steps.

This core realistically compresses a green hire toward the competence of a multi-year technician on routine and many non-routine tasks. That capability alone is the business.

8.2 Where the thirty-year claim must be qualified

Honesty about the limits is what makes the system trustworthy. **Three areas remain genuinely hard:**

- **Multi-sensory diagnosis.** *Veterans diagnose using sound, vibration, smell*, and the contextual feel of a reading. Recognition is solved; diagnosis from ambiguous, multi-sensory field signals is only partial.
- **Novel upset conditions.** Recognising what something is composes beautifully. Knowing what to do when *several systems interact in a way no procedure anticipated* is exactly where thirty years pays, and it is the hardest mile.
- **Calibrated uncertainty.** *Field reality is messy — rust, paint, insulation, darkness, occlusion*, undocumented retrofits. The *system must know when it is unsure and say so*, rather than answering confidently and wrongly.

The honest framing

We do not ship a thirty-year veteran on day one. *We ship a reliable multi-year technician on day one* — and the *flywheel compounds it toward the veteran*, facility by facility, month by month. That is both the more *truthful story and the stronger one*: a **system that demonstrably gets smarter** is a moat that generic, public-knowledge AI cannot build.

9. Deployment Realities and Guardrails

We propose to start in a training facility first to prove EON Universal. Moving from training center capability to field deployment *surfaces constraints that are easy to underestimate*. Naming them is part of doing this responsibly.

9.1 Hazardous-area hardware

Consumer AR hardware cannot simply be carried into a hydrocarbon-processing area. Classified hazardous zones require intrinsically safe, certified equipment (ATEX / IECEx). Therefore after proving EON Universal in a training facility when we go *in the field initially, we propose a ruggedised intrinsically safe tablet*. Any field-worn device deployed in a live process area must meet the applicable zone rating, or be confined to safe areas. This is a hard regulatory gate, not a soft preference, and it shapes hardware strategy from the outset.

9.2 Connectivity and edge compute

Many facilities are remote or RF-restricted. The system must *degrade gracefully and run meaningful capability at the edge*, without assuming a reliable cloud connection at the point of work.

9.3 Accountability and liability

When a worker acts on AI guidance, responsibility must be unambiguous. The system integrates with the customer's process-safety-management regime, *keeps a human accountable for safety-critical decisions, and maintains an auditable record* of guidance given and actions taken.

9.4 Brownfield drift

Facilities rarely match their drawings. Topology inference *must tolerate undocumented modifications and as-built reality*, and surface discrepancies rather than assume the drawing is correct.

9.5 Change management

The experts whose knowledge the system captures must be partners, not subjects. *Capture works best when it is designed around the people who hold the expertise*, with incentives that reward contribution rather than treating it as extraction.

10. The Strategic Frame

10.1 Makers versus takers

General-purpose AI averages publicly available knowledge. It is, in this sense, *a taker — useful, but bounded* by what everyone already knows. The *competence captured here is the opposite: proprietary, equipment-specific, procedure-specific expertise that exists nowhere on the public internet*. **A system built from that expertise is a maker.** The moat is not the model; it is the captured competence and the flywheel that compounds it.

10.2 Work Intelligence as a system of record

Enterprises run on systems of record: ERP for resources, CRM for customers, PLM for product lifecycle. *None of them captures how work is actually performed* — the operational expertise that makes a facility run safely and well. **Work Intelligence is that fourth system of record:** the structured, living record of operational competence, captured by Genesis and delivered by FieldIQ.

11. Roadmap

1. **Prove EON Universal Pilot** – Start *in the training Center with the existing traing system hardware* to test, prove and improve the company specific EON Universal platform
2. **Recognition core.** *Build the identity and geometry layers across the ~50 classes* so the glasses reliably name what they see. This is the capability that proves the concept.
3. **Procedure layer on the top fifteen.** *Add full guidance for the highest-frequency classes*, where daily value concentrates.
4. **Behavior via simulator.** *Ground the rotating and fired equipment in simulated cause-and-effect*, where understanding behavior matters most.
5. **Verdict gates on safety-critical items.** *Wire hard escalation* to every relief valve, shutdown valve, fire-and-gas detector, pig trap, wellhead, and electrical lineup.
6. **Topology inference and compounding.** *Compose guidance for unseen configurations* and turn the Compound IQ loop so the library improves with every job.

12. Conclusion

The vision — any worker, any facility, veteran-level guidance — is realistic, and the reason is structural. *Facilities are graphs assembled from a finite component vocabulary*, so the *problem to solve is bounded*:

1. **capture** roughly *fifty equipment classes deeply*,
2. **recognise** them *in the field*, and
3. **compose** *guidance* from known parts.

The **recognise-and-guide core** is *deliverable today* and immediately valuable; **full veteran** judgement is an asymptote the flywheel *approaches over time*. Built honestly — with hard safety gates, calibrated uncertainty, and an honest account of its own limits — **this is not a demo**. It is the *fourth*

system of record for the industrial enterprise, and a moat that compounds with every component learned and every job done.

© 2026 EON AI Ventures. This document describes a technical approach and product direction and does not reference specific customers. Forward-looking capability statements describe intended product behavior and are subject to development.