

EON Field IQ Recognize: Guidance in the Physical World

How Field IQ Recognize turns a phone camera into a step-by-step guide for equipment it has never seen — and why it is built to say “I can’t tell.”



Table of Contents

1. Executive summary	2
2. The problem	2
3. What Field IQ Recognize does	2
3.1 Three questions	2
3.2 The guidance loop	3
3.3 Lineage: the nine layers of Genesis	3
4. Designed to refuse	4
4.1 Provenance on every step	4
4.2 No invented manual	4
4.3 “Unclear” is a first-class answer	4
4.4 A control’s appearance is not evidence	5
4.5 The enhancement may not fail the job	5
5. Architecture	5
6. What is verified, and what is not	5
6.1 Verified	5
6.2 Not verified	6
7. One camera, every product	6
7.1 Permit IQ — their paper in	6
7.2 Genesis — from once every two years to every day	6
7.3 EON Tacit — the expert’s words at the right step	7
7.4 Assess IQ Live Interrupt — the check made continuous	7
7.5 What it costs — less than the gas meter	7
8. What comes next	8
8.1 Ask what the worker wants to do	8
8.2 Continuous tracking	8
8.3 Head-worn displays	8
9. Conclusion	8

1. Executive summary

Industrial knowledge leaves with the people who hold it. The manual is out of date or missing; the expert has retired; the new worker is standing in front of the machine now. Field IQ Recognize addresses that moment directly: a worker points a phone at real equipment and is walked through the job, one step at a time, with the relevant control highlighted, the instruction spoken aloud, and the motion animated.

The system finds the parts a person can operate, reads the text printed beside them, identifies the make and model, searches for the manufacturer's own documentation, and produces an ordered procedure. It can then check whether a step has been carried out by looking at a fresh photograph.

Its most important design property is that it refuses. Generated procedures are never presented as approved ones. A manual is never invented. A step is never called done on assumption. This paper describes what has been built, how it works, what is verified, and what is not.

2. The problem

Every operations organisation has the same conversation about tacit knowledge: the procedure exists in someone's head, the written version is wrong or unfindable, and the person who needs it is not the person who has it.

Training addresses this in advance. Simulation addresses it in a safe copy of the equipment. Neither addresses the worker who is standing in front of the real machine, now, with a job to do.

A new worker stands in front of a machine nobody trained him on. There is no manual. The guy who knew it retired. Today he calls someone, or he guesses.

The product question is therefore narrow and demanding: can a phone camera, pointed at equipment the system has never modelled, produce guidance a worker can safely follow?

3. What Field IQ Recognize does

The worker opens the app, taps Recognize, chooses what question to ask, and takes a photograph. Everything after that is automatic.

3.1 Three questions

Mode	What it answers
Equipment	What are these plant components? Names them, reads asset tags and rating plates, draws the engineering symbol, and binds them to the authored equipment drawing — which is what unlocks an approved procedure for that exact component.

Anything	What is this? Names any object in plain language. The recognition layer underneath the other two.
Guide me	What do I do? Finds the controls a hand can operate, identifies the make and model, and returns an ordered procedure. The only mode that produces steps.

The first mode finds an approved procedure. The third generates one. The product never blends the two, and the screen states on every step which one the worker is holding.

3.2 The guidance loop

1. Find the controls. Not the object — the switches, dials, levers, handles, doors, drains and gauges on it, and the text printed beside each.
2. Identify the machine. Brand and model, read from the rating plate.
3. Find the real procedure. A second, search-grounded call looks for the manufacturer's documentation for that exact model. If found, its steps replace the generated ones and the document is named. If not found, nothing is invented.
4. Guide. One step at a time. The relevant part pulses green while the rest of the image recedes; a voice reads the step, with any caution spoken before the instruction; an animation shows the motion.
5. Verify. On request, a fresh photograph is taken and one question is asked: has this step been done? Three answers — done, not done, unclear — always with a line describing what was visible.
6. Advance. The worker presses Next. The system never advances a step by itself: a procedure that moves on without the worker is one the worker has stopped reading.

3.3 Lineage: the nine layers of Genesis

EON's Genesis platform guides a worker through a digital twin using nine layers. Field IQ Recognize applies the same layers to physical equipment.

Genesis layer	In the twin	On real equipment
SOP parsing	Authored procedure	Generated from the photograph, or read from the manufacturer's manual
Geometric understanding	Model geometry	Detected controls and their boxes; LiDAR depth for 3D placement
Step generation	Authored	Ordered steps with an action verb each
Camera path	Flies the virtual camera	Not applicable — the worker moves; the screen directs the eye

HUD synthesis	Step card	Step card: title, instruction, target part, caution, provenance
Mesh highlights	Highlight the mesh	Green pulse on the part; optional 3D marker on the real surface
Narration	Avatar voice	On-device voice, caution first
Sound design	Sound library	Three system tones, never over the voice
Pacing	Timed steps	Expected duration per step; overdue steps prompt a check

4. Designed to refuse

A worker acts on what the screen says. That makes a confidently wrong answer the failure that matters, and it shapes every safety decision in the product.

4.1 Provenance on every step

Generated steps and manufacturer steps look identical on a phone screen. The step card therefore carries a provenance line on every step: “Suggested from what the camera sees — not an approved procedure,” or “From the manual: <document title>.” The server is the only party that knows which is true, so the server says, and the device renders anything it does not explicitly recognise as generated — the cautious direction to fail in.

4.2 No invented manual

The search-grounded lookup is instructed to return nothing rather than fall back on general knowledge, because a generated procedure already exists and is already labelled as one. An invented procedure presented as a manual is the single outcome that must not happen. This was verified by submitting a fictitious model number: the system searched, found nothing, and reported the steps as generated.

4.3 “Unclear” is a first-class answer

The verification prompt names four cases that must return unclear rather than done: the control is not in frame; the reasoning is about what a person would probably have done rather than what is visible; the effect is internal or electrical and nothing shows it; the image is too dark or distant to read.

A fifth was added after a live test. Asked whether a power button had been pressed, the system answered done at 0.9 confidence because the button was green. A power button is green whether or not anyone has touched it. The rule now states plainly that a control’s colour, styling and labelling describe the control, not its state, and that evidence of a step is a visible change in the thing the step affects. After the change, the same question was answered from a lit indicator lamp instead — with confidence reported as zero. The card renders that as “Looks done — but not sure,” in yellow rather than green. The verdict is never rewritten; only what the worker is told about its certainty.

4.4 A control's appearance is not evidence

Every verdict carries an evidence line describing what was visible in the photograph, not what the instruction said. This is the part a worker can argue with, so it sits beside the verdict rather than behind it. A verdict without evidence is rejected rather than defaulted.

4.5 The enhancement may not fail the job

Web search, the 3D marker and verification are all built as additions that can only improve a result, never break one. Each has a bounded time budget measured from the moment the request was made, not from the moment the worker picked it up; each returns the recognition unchanged on any failure. A worker with a slow connection gets a generated procedure rather than a timeout.

5. Architecture

The phone never holds an AI key. It photographs, and posts the image to a backend. The backend queues the job; a perception worker picks it up, calls the vision model with the prompt for the requested mode, validates the response against a strict schema, optionally runs the grounded manual lookup within the remaining budget, and publishes the result. The backend returns it to the phone.

Component	Role
iOS app	Camera, AR session, LiDAR depth, the step card, voice, sound. Sends a photograph and a mode; receives components, steps, provenance and verdicts.
Backend	Validates the request, queues it, waits for the result within a fixed budget, and returns it. Never calls an AI model directly.
Perception worker	Runs the prompts: industrial, general, guidance, verify. Holds the only AI credential. Validates every answer; drops what it cannot read rather than the whole response.
Manual lookup	A second, search-grounded call, made only when a brand and model were read and only within the time that remains.

Every deploy of the perception worker refuses to ship unless the source carries all four prompts, the bounded lookup, the guard around it, the appearance-is-not-evidence rule, and a passing linter. Each item on that list is a failure that occurred once.

6. What is verified, and what is not

6.1 Verified

- The full chain answers on the production server: from a test image, ten controls and six ordered steps in about thirty seconds, each step referencing a part actually detected, each with a motion the device can animate.
- The manual lookup refuses to invent documentation for a fictitious model.
- Verification answers from visible evidence, reports its confidence, and the device flags a low-confidence done as unsure.
- Ninety-six automated tests across the server and the app, run on every push; the app builds and its tests execute on a cloud Mac on every change.

6.2 Not verified

- Any of it running on a phone pointed at a real object. Every check above is a compiler, a test, or a direct call to the server. The failures that occur when a person points a camera at equipment — poor light, an awkward angle, a hand in frame, AR tracking losing the room — are a different category and none are yet covered.
- The 3D marker's geometry. It anchors a ring to the real surface using LiDAR depth. The coordinate mathematics was found to be wrong and corrected against the platform's own transform, but corrected without a device to check it on. It is opt-in, LiDAR-only, and the two-dimensional highlight remains the default.
- Real-world manual hit rate. The lookup returns nothing far more often than it returns a document. That is the refusal working; it is also a measure of how often a worker gets generated rather than documented steps, and it has not yet been measured on real appliances.

7. One camera, every product

Field IQ Recognize stands on its own. It is also the front door to a family, and the family is what turns a guided job into a corrected one — and a product bought once into a utility used every day.

7.1 Permit IQ — their paper in

Isolation permits are written the morning of the job, on paper, and stay in the pack. Permit IQ makes the permit a living record; Recognize makes it a guide. The permit for this morning's job becomes the procedure: point 7 on the form becomes step 7, lit up on the real hardware it describes. The customer keeps their process. We connect it — their paper in, competence out.

7.2 Genesis — from once every two years to every day

Today a worker meets guided procedure at a training centre, on a schedule measured in years. That is the economics of the classroom: the guidance is excellent and almost never present.

With the morning's permit as the lesson and the phone as the guide, guidance runs every day, for every worker, by default. It is the same nine-layer brain throughout: Genesis guides the worker through the digital twin; Recognize guides them through the real machine, with no 3D model and no wait for a twin to be built. The training-centre skid and the morning permit finally meet.

7.3 EON Tacit — the expert's words at the right step

Captured expertise is only useful at the moment it applies. A step card has exactly one such moment: the retired expert's remark about this valve, this trap, this trick of the trade, spoken at the step where it matters and nowhere else.

7.4 Assess IQ Live Interrupt — the check made continuous

Recognize's verdict — done, not done, unclear — is asked once, when the worker requests it. Assess IQ Live Interrupt is the same judgment made continuous: every observation scored in milliseconds against the current step, the permit's flags, the coded checks and what the camera sees, and stamped Stop, Continue, or Call the supervisor. High confidence acts. Uncertainty interrupts and escalates to a human.

Speed is the property that makes it a different product rather than a faster report. Feedback that arrives while the hand is still on the valve can change what happens next; a score next month is an audit. The refusal culture of section 4 is also what makes an interrupt trustworthy: a system that has learned to say "unclear" has earned the right to say "stop."

7.5 What it costs — less than the gas meter

Every field worker already carries a personal gas detector. Industry cost-of-ownership guides put calibration and maintenance alone at 40–50% of a gas-detection program's lifetime cost, with sensors replaced every two to three years; between purchase, calibration gas, bump testing and sensor replacement, the meter costs an operator several hundred dollars per worker, every year. Nobody questions that spend, because it saves lives.

Field IQ is priced against that anchor: the whole solution, per worker per year, below the meter on the worker's belt. Same justification — it saves lives, and it saves cost. The arithmetic is then simple, and recurring:

Field workers	Per worker per year*	Recurring revenue
1,000	\$400	\$0.4M per year
10,000	\$400	\$4M per year
50,000	\$400	\$20M per year

**Illustrative price point, set below typical gas-detector cost of ownership. The shape is the argument: priced like safety equipment, bought per worker, renewed every year — not a licence that sleeps between training cycles.*

8. What comes next

8.1 Ask what the worker wants to do

Today the system infers the task from the photograph. The obvious next step is to ask: recognise the machine, then offer the jobs a person might want to do with it — defrost, replace the filter, reset after a fault — and generate or retrieve the procedure for the one chosen. This is the difference between a system that guesses the job and one that is told it.

8.2 Continuous tracking

The step marker is placed once, from the capture pose. Tracking the part as the worker moves, and noticing a step's completion without being asked, both depend on the same anchoring work and are the natural extension of it.

8.3 Head-worn displays

Everything the phone shows is designed to move to a display in the worker's line of sight, so that the light, the voice and the arrow arrive with both hands free.

9. Conclusion

The commercial claim is simple: a worker who does not know the machine, cannot find the manual, and may not read the language, follows a green light, a voice and an arrow — and does the job right.

The engineering claim underneath it is narrower and, we think, more important: the system knows what it does not know. It labels what it generated. It declines to invent a manual. It says “unclear” rather than “done.” Those refusals are what make it possible to put in someone's hands.

What remains is to put it in someone's hands.