EON AI Ventures

Field IQ Edge

A Digital Twin Beyond Connectivity

How EON destills the right Work Intelligence from the Digital Twin and carries it securely to every worker - online, intermittently connected , or fully offline.

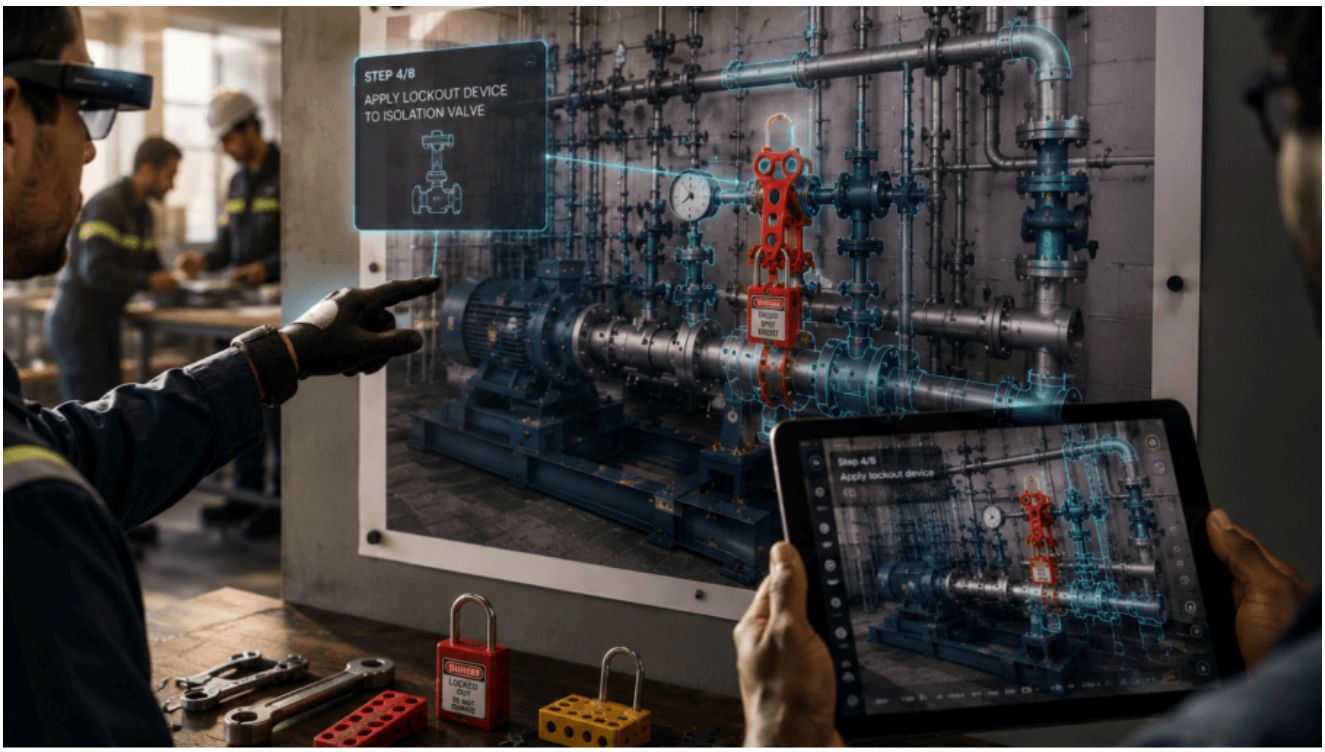

Document purpose and status

This white paper defines the product vision, capability boundaries, reference architecture and recommended development path for Field IQ Edge. It is designed to become the source document for a press release, executive presentation, web experience, product requirements document, technical specification and intellectual-property review after approval.

DOCUMENT STATUS

Strategic and technical concept draft. Performance targets, regulated-industry claims, model choices and patent positions must be validated through engineering tests, customer discovery, standards review and legal counsel before external publication.

Title	Field IQ Edge - A Digital Twin Beyond Connectivity
Primary focus	Offline-capable and intermittently connected field execution
Broader category	Work Intelligence
Enabling process	Work Intelligence Distillation
Existing flywheel	Digital Twin → Field IQ Edge → Assess IQ → Compound IQ → Digital Twin
Priority industries	Oil & gas, aerospace and healthcare; extensible to mining, utilities and manufacturing
Research horizon	Technology and product sources reviewed through July 2026

Terminology used throughout this paper

To preserve the existing product story, this paper does not introduce another platform or product family. The terminology is deliberately constrained:

- The EON Digital Twin consists of Genesis, EON Universal, EON Live and Situation Room.
- Field IQ Edge is the offline-capable deployment profile of Field IQ. It extends the Field IQ experience; it is not a separate ecosystem.
- Assess IQ observes, measures and verifies real-world performance and evidence.
- Compound IQ converts validated operational deltas into proposed improvements to procedures, models and the Digital Twin.
- Work Intelligence is the enterprise system of record for how work is understood, guided, performed, assessed and improved.
- Work Intelligence Distillation is a process, not a new product name.

Table of Content

1. Executive Summary

Enterprise AI should not stop where connectivity ends. Field IQ Edge extends the EON Digital Twin into the physical workplace by carrying a secure, mission-specific portion of the enterprise's Work Intelligence directly on a phone, tablet, smart-glass companion device or local edge computer. The worker can recognize equipment, retrieve trusted procedures, receive context-aware guidance, capture evidence and continue an approved workflow even when there is no internet connection.

The central design decision is to avoid placing the entire Digital Twin or a general-purpose cloud model on every device. Instead, EON should use Work Intelligence Distillation: a governed process that selects, specializes, compresses and packages only what a defined worker, site, asset and mission require. The package combines compact AI models with deterministic workflows, approved knowledge, digital-twin relationships, visual recognition, local retrieval, security controls and synchronization rules.

This approach is technically realistic now. Apple, Google, PyTorch, Qualcomm, Microsoft and NVIDIA provide production frameworks for on-device inference across mobile and edge hardware; Apple exposes an approximately three-billion-parameter on-device foundation model, Google supports fully on-device LLM inference, and ExecuTorch is designed to run compact language models on phones and embedded systems. Research families such as MobileVLM and SmolVLM further demonstrate that useful multimodal capabilities can be built at mobile-scale model sizes.

CORE RECOMMENDATION

Build Field IQ Edge now as a device-first, gateway-optional and cloud-enhanced system. The core safety workflow must remain available locally. A site edge gateway may add larger models and local sensor access without internet. The cloud should provide enterprise-scale simulation, global analytics, model training and continuous improvement - but should not be required to complete the bounded field mission.

Six decisions recommended by this paper

1. Keep Field IQ Edge inside the existing EON product architecture and brand story.
2. Use a modular system of specialized perception, retrieval, reasoning and workflow components rather than one monolithic vision-language model.
3. Keep approved SOP truth, safety logic and authorization outside generative model weights so each decision remains auditable and updateable.
4. Develop reusable industry intelligence for oil & gas, aerospace and healthcare, while isolating each customer's private data and site-specific adaptations.

5. Treat Work Intelligence Distillation, digital-twin-generated training data, asset grounding, safety compilation and the governed feedback loop as the principal EON intellectual-property opportunity.
6. Begin with a 90-day oil & gas pilot covering 20-50 assets, two normal workflows and one bounded abnormal-condition scenario, then reuse the architecture across aerospace and healthcare.

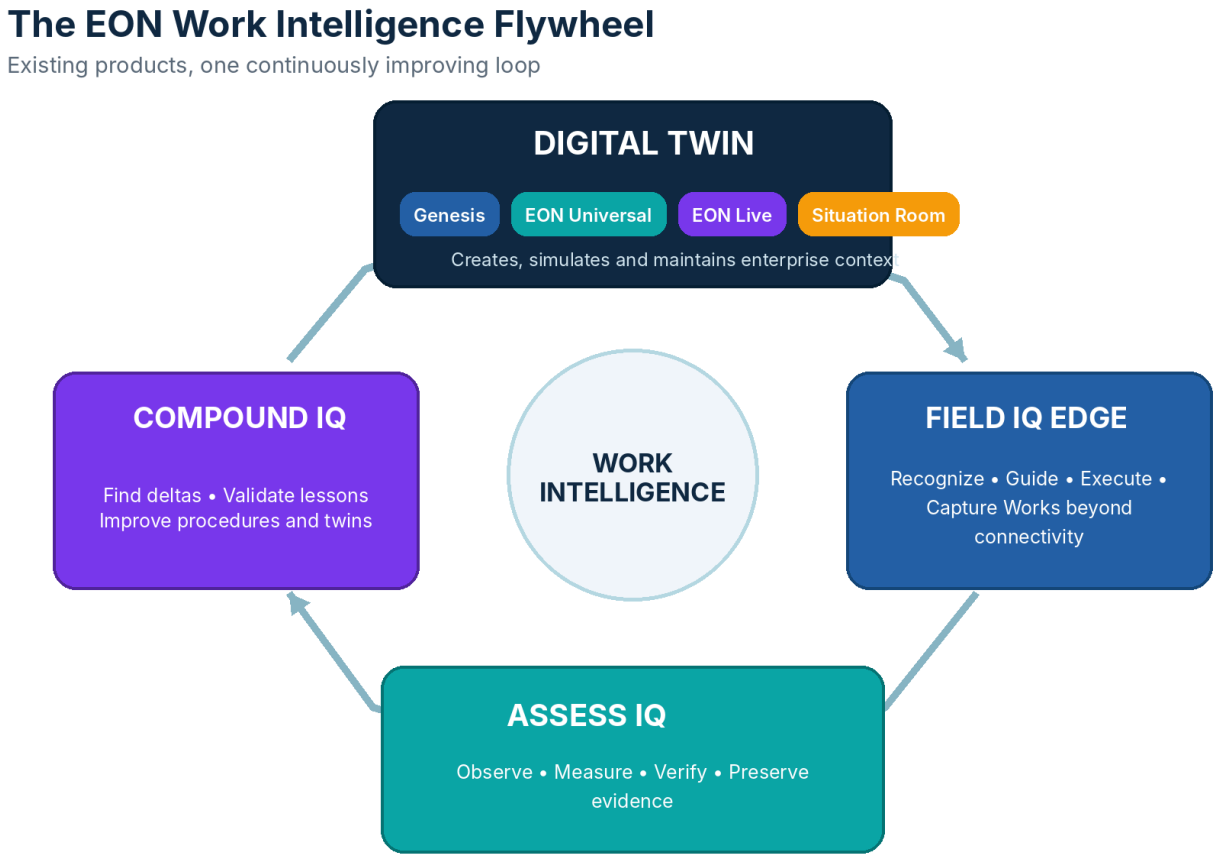

Figure 1. Field IQ Edge extends the existing EON Work Intelligence Flywheel without introducing another platform name.

2. The Connectivity Gap

The physical environments where expertise is most valuable are often the environments where cloud dependence is least reliable. Offshore installations, refineries, processing plants, aircraft hangars, underground mines, utility corridors, hospitals and emergency sites may have weak coverage, intentional network segmentation, congested Wi-Fi, metal structures that block radio signals, cybersecurity restrictions or complete air gaps. Even where a network exists, workers

cannot assume that every camera frame, voice question or safety-critical decision can wait for a remote service.

This creates a structural mismatch. Enterprise knowledge is increasingly centralized in cloud systems, while work remains distributed across assets, rooms, shifts and remote locations. A cloud-only assistant may be impressive during a demonstration yet fail at the exact moment a technician enters a dead zone, a hospital isolates a network segment, an aircraft is serviced on a distant apron, or an industrial site blocks outbound traffic.

Why offline capability matters

- Continuity: an assigned workflow should not stop because the wide-area network disappears.
- Latency: recognition, stop conditions and step confirmation may need immediate local response.
- Privacy and sovereignty: sensitive images, procedures and operational data may be required to stay on the device or site.
- Resilience: local operation reduces dependence on a single provider, region, network path or external API.
- Cost control: continuous video and multimodal cloud inference can become expensive at enterprise scale.
- Trust: a worker needs to know which procedures and model versions are present, when they were approved and what happens when the system is uncertain.

Offline does not have to mean isolated from every local signal. Field IQ Edge should distinguish four operating conditions: zero network; site-local connectivity without internet; intermittent wide-area connectivity; and full cloud connectivity. A device can remain "offline" from the internet while communicating with a Bluetooth instrument, a local industrial gateway or an air-gapped site server.

DESIGN PRINCIPLE

The user experience should remain continuous across connectivity states. The system may change its source of intelligence, but it should not force the worker to change applications, relearn the workflow or guess which features are available.

3. The EON Digital Twin

Field IQ Edge is not an isolated offline assistant. It is an extension of the EON Digital Twin. In the architecture defined for this paper, the Digital Twin consists of Genesis, EON Universal, EON Live

and Situation Room. Together they create the enterprise context from which Field IQ Edge receives the intelligence required for a specific mission.

Digital Twin component	Role in Field IQ Edge
Genesis	Captures and structures expert knowledge, procedures, simulations and behaviors.
EON Universal	Maintains persistent spatial content, assets, components, relationships and reusable knowledge.
EON Live	Provides dynamic scenario, behavior and physics context for what-if reasoning and simulation.
Situation Room	Prioritizes events, operating conditions, risks and organizational situational awareness.

EON already positions Field IQ as an AI-powered field-assistance capability and describes a broader Work Intelligence flywheel in which Genesis captures expertise, Field IQ deploys it, Assess IQ verifies performance and Compound IQ improves the underlying procedures and twins. [1-5] Field IQ Edge adds one missing operating condition to that story: the ability to continue the bounded mission beyond internet connectivity.

What the Digital Twin contributes to the edge

- A structured asset graph rather than a pile of documents.
- Known component names, relationships, locations and allowable interactions.
- Approved SOPs, hazards, prerequisites, stop conditions and escalation points.
- Visual and spatial representations that can train recognition models and drive AR overlays.
- Normal and abnormal scenarios that can be simplified into local decision trees and cached guidance.
- A governed destination for field evidence and validated improvements.

4. Introducing Field IQ Edge

Field IQ Edge is the offline-capable deployment profile of Field IQ. It carries the right Work Intelligence to the point of work and maintains the same field experience across a continuum from disconnected to fully connected operation. It is not intended to reproduce every cloud capability; it is intended to preserve the capabilities that matter most for a defined job.

What the worker can do without internet

- Point the device at a known asset and recognize the equipment class and visible components.
- Confirm the specific asset using a combination of visual evidence, OCR, QR or NFC, location and Digital Twin relationships.
- Retrieve the approved procedure, drawing, checklist, training animation or troubleshooting branch for that asset.
- Ask natural-language questions against a trusted local knowledge package.
- Receive step-by-step visual, spoken or AR guidance with deterministic safety gates.
- Record measurements, photos, short video, voice notes, exceptions and completion evidence.
- Continue the assigned workflow and queue all events for governed synchronization.

What Field IQ Edge is not

- It is not a general chatbot that is expected to know every industrial situation.
- It is not a replacement for a qualified operator, clinician, maintainer or responsible engineer.
- It is not an autonomous control system and should not directly actuate equipment by default.
- It is not an excuse to let a generative model invent safety-critical steps.
- It is not a copy of all enterprise data placed on every worker device.

PRODUCT PROMISE

The Digital Twin does not stop at the enterprise. With Field IQ Edge, it travels with the worker - carrying the right intelligence for the job and returning the evidence needed to make the next cycle better.

5. From Cloud Intelligence to Edge Intelligence

Cloud models remain valuable because they can use the largest reasoning systems, access the full enterprise knowledge base and run computationally intensive simulation. The mistake would be treating cloud intelligence and edge intelligence as competitors. They are different operating tiers of one system.

Tier	Primary responsibility
Device tier	Immediate perception, trusted retrieval, approved workflow execution, safety rules, evidence capture and bounded reasoning.
Site edge tier	Optional larger models, local sensor aggregation, multi-user services and site-wide operation without internet.
Cloud tier	Full Digital Twin, EON Live simulation, enterprise search, cross-site analytics, model training, Work Intelligence Distillation and remote collaboration.

The system should always choose the smallest tier capable of completing the task safely. A local detector should identify a known valve rather than sending a continuous video stream to a general cloud model. A deterministic workflow should enforce lockout-tagout prerequisites rather than asking a language model to remember them. A cloud VLM or expert should be called only when the local package cannot resolve the situation and connectivity permits escalation.

Why a modular system is better than one supermodel

A monolithic vision-language model is attractive because it appears simple: one model sees, reasons and speaks. In a safety-critical field product, however, a modular design is more controllable. Specialized detectors can be trained and measured for specific equipment; OCR can read asset labels; a graph resolver can reconcile identity; a small language model can explain; a workflow engine can control sequence; and a safety policy engine can stop the process. Each component can be tested independently, replaced as technology advances and assigned a clear authority boundary.

- Better latency and battery use for common recognition tasks.
- Higher explainability because identity and next-step decisions have traceable evidence.
- Faster updates because procedures and knowledge can change without retraining the entire model.
- Reduced hallucination risk because generative outputs are constrained by approved retrieval and workflow state.
- Hardware portability because each model can be compiled for the available CPU, GPU or NPU.

6. Work Intelligence Distillation

Work Intelligence Distillation is the core enabling process behind Field IQ Edge. It is broader than model compression. It transforms enterprise-scale Work Intelligence into a secure, optimized and testable deployment package for a defined context.

DEFINITION

Work Intelligence Distillation is the governed process that selects, grounds, specializes, compresses, validates and packages the portion of the EON Digital Twin required for a defined industry, customer, site, role, mission, device and risk profile, so Field IQ Edge can recognize, guide, reason and capture evidence locally.

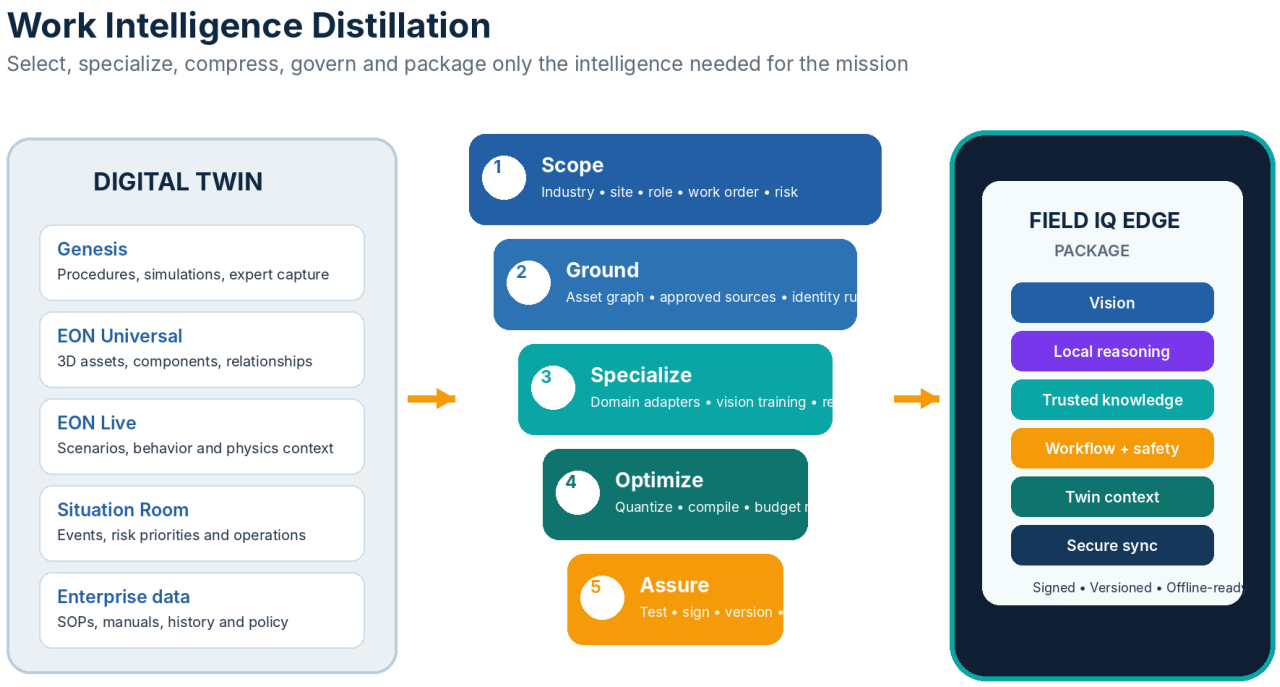

Figure 2. Work Intelligence Distillation converts the full Digital Twin into a mission-scoped, signed Field IQ Edge package.

The distillation pipeline

1. **Scope the mission.** Determine the industry, customer, site, role, work order, asset set, likely scenarios, device profile, connectivity window and risk class.
2. **Ground the truth.** Select approved SOPs, manuals, spatial relationships, component definitions, identity rules, safety policy and escalation paths.

3. **Generate or curate data**. Use expert examples, customer-approved field imagery and synthetic Digital Twin renders to create labeled training and evaluation data.
4. **Specialize the models.** Train or adapt compact vision and language components for stable domain behavior; keep frequently changing customer facts in retrieval and workflow layers.
5. **Optimize for the target.** Quantize, prune or distill models; compile them for the device runtime; budget memory, storage, battery, camera rate and response time. Knowledge distillation, LoRA and quantized adaptation are established methods for reducing deployment cost while retaining task performance. [17-19]
6. **Compile trusted behavior.** Convert procedures, prerequisites, stop conditions and evidence requirements into deterministic workflows and policies.
7. **Assure the package.** Test recognition, retrieval, refusal, safety gates, adversarial inputs, outdated content and loss of connectivity. Sign, version, approve and define rollback.
8. **Deploy and learn.** Field evidence returns through Assess IQ and Compound IQ; approved changes update the Digital Twin and the next package.

Where to put knowledge: weights, retrieval or rules

Layer	What belongs there
Model weights / adapters	Stable domain language, visual features, industry terminology, recurring patterns and general task behavior.
Local retrieval index	Customer SOPs, manuals, asset records, incident summaries, site-specific facts and content that must be updated or cited.
Digital Twin graph	Asset identity, component hierarchy, spatial relations, dependencies, allowable states and task context.
Workflow and safety engine	Required sequence, prerequisites, stop conditions, permissions, evidence requirements and escalation.

This separation is strategically important. Customer data does not need to be baked into model weights to create value. Keeping private facts in an encrypted, customer-specific knowledge layer improves ownership, update speed, deletion, auditability and cross-customer isolation. Model training should be reserved for stable capabilities that genuinely benefit from adaptation.

7. The Field IQ Edge Package

The deployable artifact is a signed Field IQ Edge package. "Package" is a technical description, not a new product brand. A package may range from a small task bundle to a broader site bundle, but it should always be built against explicit device and risk budgets.

Package component	Purpose
Manifest and trust	Package ID, customer, site, role, permitted tasks, version, expiry, signature, dependencies and rollback target.
Perception models	Object detection, segmentation, state cues, OCR and optional compact multimodal model.
Identity resolver	Rules and embeddings that combine visual evidence, tag data, location, proximity and Digital Twin graph context.
Knowledge layer	Encrypted documents, structured facts, local vector index, glossary and source citations.
Twin subset	Relevant assets, components, geometry, relationships, hazards and cached scenario outputs.
Workflow layer	Approved procedures represented as state machines, conditions, required confirmations and evidence contracts.
Safety and authorization	Role permissions, certification status, stop rules, confidence thresholds and escalation policy.
Experience assets	AR anchors, diagrams, animations, multilingual prompts, speech assets and accessible UI resources.
Evidence and synchronization	Append-only event log, photos, readings, hashes, pending uploads, delta updates and conflict rules.

Package design targets

- Mission-complete: everything necessary for the assigned workflow is present before the worker leaves coverage.
- Risk-weighted: safety content and stop conditions take priority over convenience content when space is constrained.
- Incremental: large 3D or media assets can be modular while core guidance remains available.

- Reversible: every update supports version comparison and rollback.
- Expiring: roles, procedures and high-risk packages carry validity windows and mandatory refresh policies.
- Observable: every decision records the package, model and source versions involved.

8. Industry and Customer Specialization

There is substantial value in building compact, specialized intelligence for each major industry. The right architecture, however, is not a completely separate codebase or foundation model for every segment. EON should maintain one Field IQ Edge runtime and a layered specialization system.

How specialization scales without creating a separate product for every customer

One Field IQ Edge runtime; layered intelligence that becomes increasingly specific

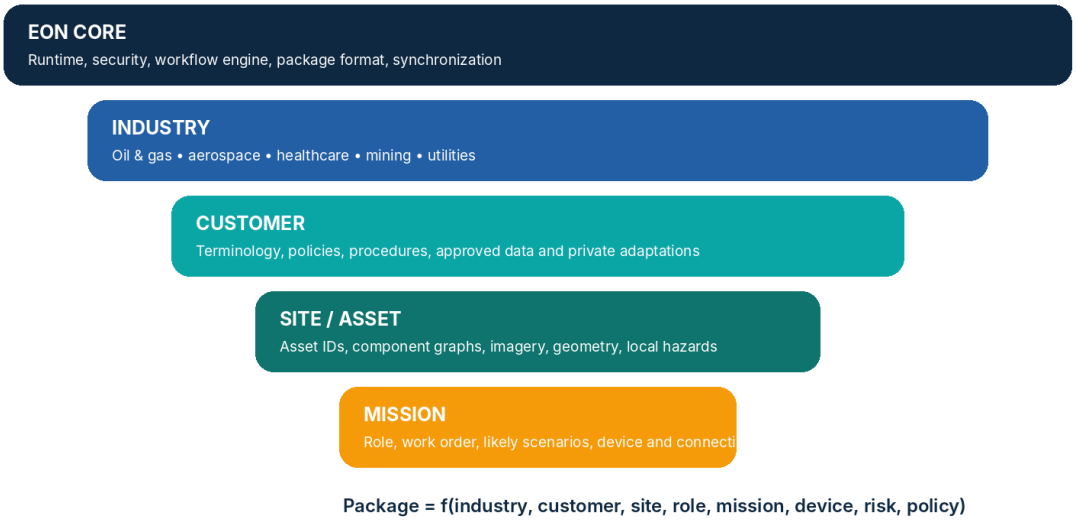

Figure 3. Specialization is layered from shared EON technology to the exact mission, preserving reuse and customer isolation.

Recommended specialization layers

Layer	Contents and ownership
EON core	Runtime, security, workflow engine, package format, synchronization, evaluation harness and shared multimodal capabilities.

Layer	Contents and ownership
Industry intelligence	Domain ontology, terminology, component families, common hazards, procedures, failure patterns and benchmark datasets.
Customer intelligence	Customer-owned policies, SOPs, naming conventions, asset data, incident history and approved adaptations.
Site intelligence	Specific assets, labels, geometry, locations, visual signatures, local hazards and nearby relationships.
Mission intelligence	Only the role, work order, likely abnormal cases, evidence requirements and media needed for the assignment.

Why the data becomes defensible IP

Small models are especially sensitive to the quality and relevance of their training data. A public foundation model may know what a pump is, but it will not know the exact appearance of a customer's corroded Pump P-204, the relationship between its suction valve and the local line-up, or which photographic evidence the customer requires before closing a task. EON can build differentiated industry intelligence from curated ontologies, expert traces, failure taxonomies, synthetic Digital Twin renders, evaluation suites and governed operational deltas.

The Digital Twin creates an additional advantage: it can generate labeled synthetic images of known equipment across viewpoints, lighting, occlusion, wear and component states. Those images can bootstrap site-specific recognition models before enough field imagery exists. Real field data then calibrates the simulation-to-reality gap. The combination of synthetic twin data and governed field evidence is more valuable than either source alone.

Customer data boundary

- Customer raw data and customer-specific knowledge remain logically and contractually isolated.
- No cross-customer model improvement should occur by default.
- Generic improvements should use explicit rights, review, anonymization and provenance.
- Customer-specific retrieval indexes and adapters should be independently exportable, revocable and deletable.
- The contract should distinguish EON platform IP, EON generic industry intelligence, customer data, customer adaptations and jointly developed improvements.

9. The Offline Capability Pyramid

The capability pyramid provides a practical way to communicate what is easy today, what requires more engineering and what should remain cloud-enhanced. The lower levels should be treated as the dependable offline foundation. Higher levels add adaptive behavior and enterprise reach but should never be prerequisites for the core assigned workflow.

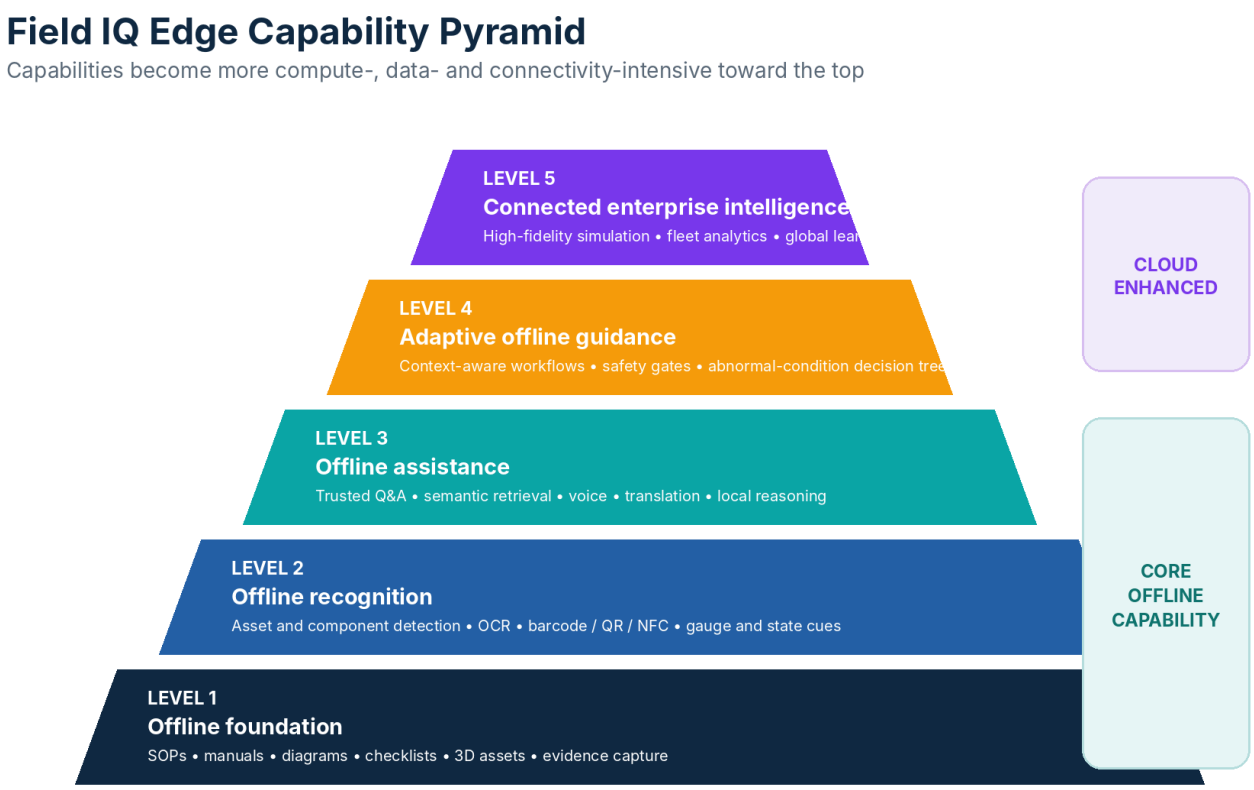

Figure 4. A capability pyramid separates dependable offline functions from advanced connected intelligence.

Level 1 – Offline foundation

Cached manuals, SOPs, diagrams, checklists, approved media, forms, 3D assets and evidence capture. This is the lowest-risk capability and should be available on nearly every supported device.

Level 2 – Offline recognition

Equipment and component detection, OCR, barcode, QR, NFC and basic instrument/state cues. Specialized mobile vision architectures are mature and designed for resource-constrained devices. The main challenge is not recognizing a generic pump; it is resolving the specific asset

identity in a visually repetitive plant.

Level 3 – Offline assistance

Local semantic search, source-grounded question answering, voice interaction, translation and concise explanation. Compact on-device language models are now practical, but the knowledge scope must remain bounded and the answer should cite the local approved source.

Level 4 – Adaptive offline guidance

Context-aware workflows, local assessment, safety gating, cached abnormal scenarios and bounded troubleshooting. This level is feasible now for defined procedures but requires careful engineering, validation and confidence-aware fallback.

Level 5 – Connected enterprise intelligence

Full-fidelity physics simulation, enterprise-wide reasoning, cross-site pattern analysis, fleet optimization, global learning and live expert collaboration. These functions may run in the cloud or on a powerful site edge server, but they should not be promised as phone-only offline capabilities for unrestricted situations.

10. Offline, Limited and Connected Capability Matrix

The following matrix is the central customer-facing boundary statement. "Fully offline" means the device has no internet connection. "Limited / site-local" includes intermittent WAN or a private local network. "Fully connected" includes access to the full Digital Twin and enterprise services.

Capability	Fully offline	Limited / site-local connectivity	Fully connected
View SOPs, manuals, diagrams and cached 3D assets	Full	Full + updates	Full + latest enterprise content
QR, barcode, NFC and label OCR	Full	Full	Full
Known asset and component recognition	Full for trained catalog	Full + gateway models	Full + cloud fallback
Specific asset identity	Full with multi-signal confirmation	Stronger with site systems	Strongest with live enterprise context

Capability	Fully offline	Limited / site-local connectivity	Fully connected
Local semantic search and trusted Q&A	Full within package	Full + selective remote retrieval	Enterprise-wide
Voice input and spoken guidance	Full on supported device profile	Full	Full
Step-by-step SOP execution	Full	Full	Full
Deterministic safety gates and stop rules	Full	Full	Full
AR overlays and component highlighting	Full for cached geometry	Full + local spatial services	Full
Evidence capture and local assessment	Full; queued for sync	Full; partial upload	Full and centrally visible
Live OT / IoT data	Only if cached or directly paired	Yes through site LAN / gateway	Yes through enterprise integrations
Abnormal-condition assistance	Bounded fault trees and approved scenarios	Broader local models and data	Full Digital Twin and expert escalation
High-fidelity EON Live simulation	Cached results only	Possible on powerful site edge	Full
Remote expert collaboration	No	As bandwidth permits	Full
Cross-site analytics and global learning	No	Summary sync only	Full
Model training and Work Intelligence Distillation	No	No; validation data can be staged	Central / governed environment
Package updates, revocation and policy refresh	No until reconnection	Delta updates as bandwidth permits	Immediate

The matrix is intentionally conservative. Some advanced functions may become available on future mobile hardware, but the product should be sold against validated device profiles rather than theoretical model capability. The offline promise is strongest when it is bounded by a known site, asset catalog and procedure set.

11. The Field Experience

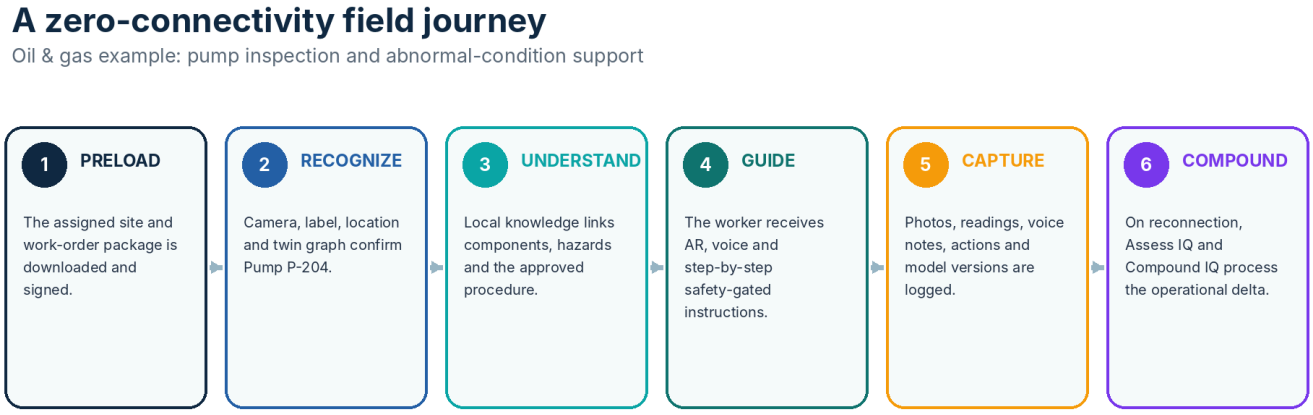

Figure 5. A complete zero-connectivity journey from preload to governed learning.

Oil & gas: pump inspection and abnormal condition

A technician enters a low-connectivity process area and points a tablet at a pump. Field IQ Edge detects a centrifugal pump and visible components, reads the equipment label, checks the device location and compares nearby equipment relationships with the local Digital Twin graph. It presents: "Pump P-204 - identity confidence high. Confirm before continuing." After confirmation, the worker states that vibration is elevated and discharge pressure is low.

The device retrieves the approved abnormal-condition workflow, checks prerequisites, asks for specific observations and highlights the next component. The local reasoning model explains each step, but the workflow engine controls sequence and stop conditions. If the evidence leaves the approved scenario envelope, Field IQ Edge does not invent a repair. It makes the equipment safe according to the available procedure, records the unresolved condition and requests escalation when a connection becomes available.

Aerospace: maintenance task-card execution

A maintainer identifies an aircraft zone, panel, connector and component using visual recognition, OCR and task context. Field IQ Edge retrieves the correct revision of the maintenance task card, presents the sequence, verifies tool and part identifiers, captures pre- and post-task evidence and records deviations. The local package may assist with FOD checks, connector routing, access-panel sequence and required sign-offs. It does not replace the licensed maintainer, approved maintenance data or regulatory release authority.

Healthcare: equipment setup and procedural assurance

A nurse or technician uses Field IQ Edge to identify a supported device, retrieve the hospital-approved setup or cleaning procedure, confirm required accessories, follow a checklist and capture competency evidence. Early healthcare deployments should focus on equipment, training, infection-control workflows, logistics and non-diagnostic procedural support. Diagnostic decisions, medication dosing and autonomous clinical recommendations require separate clinical validation and, where applicable, medical-device regulatory pathways.

Common experience principles

- Show the recognized asset, confidence, evidence and package version before a critical workflow begins.
- Make "not recognized" a successful safe outcome; never force a guess.
- Use hands-free voice and AR where helpful, but preserve a clear touch interface and readable fallback.
- Expose the current connectivity mode and pending synchronization without distracting the worker.
- Always distinguish an approved instruction from an AI explanation or suggestion.

12. How It Works in Simple Terms

Imagine that the Digital Twin is the most experienced team at headquarters. Before a worker goes into the field, that team packs a secure intelligent backpack. The backpack does not contain everything the enterprise knows. It contains the exact maps, pictures, procedures, safety rules and small specialists most likely to be needed for that job.

In the field, the backpack can see the known equipment, look up the right instruction, explain the next step and write down what happened. It does not call headquarters for every question. When the worker comes back into range, the backpack returns the evidence and the differences between the plan and reality. Assess IQ determines what happened; Compound IQ identifies useful, validated lessons; and the Digital Twin improves the next backpack.

ONE-SENTENCE EXPLANATION

Field IQ Edge is the part of the EON Digital Twin that travels with the worker when the network cannot.

13. Technical Reference Architecture

Reference Architecture: Device-first, Gateway-optional, Cloud-enhanced

The core workflow remains local; compute and enterprise reach expand when connectivity is available

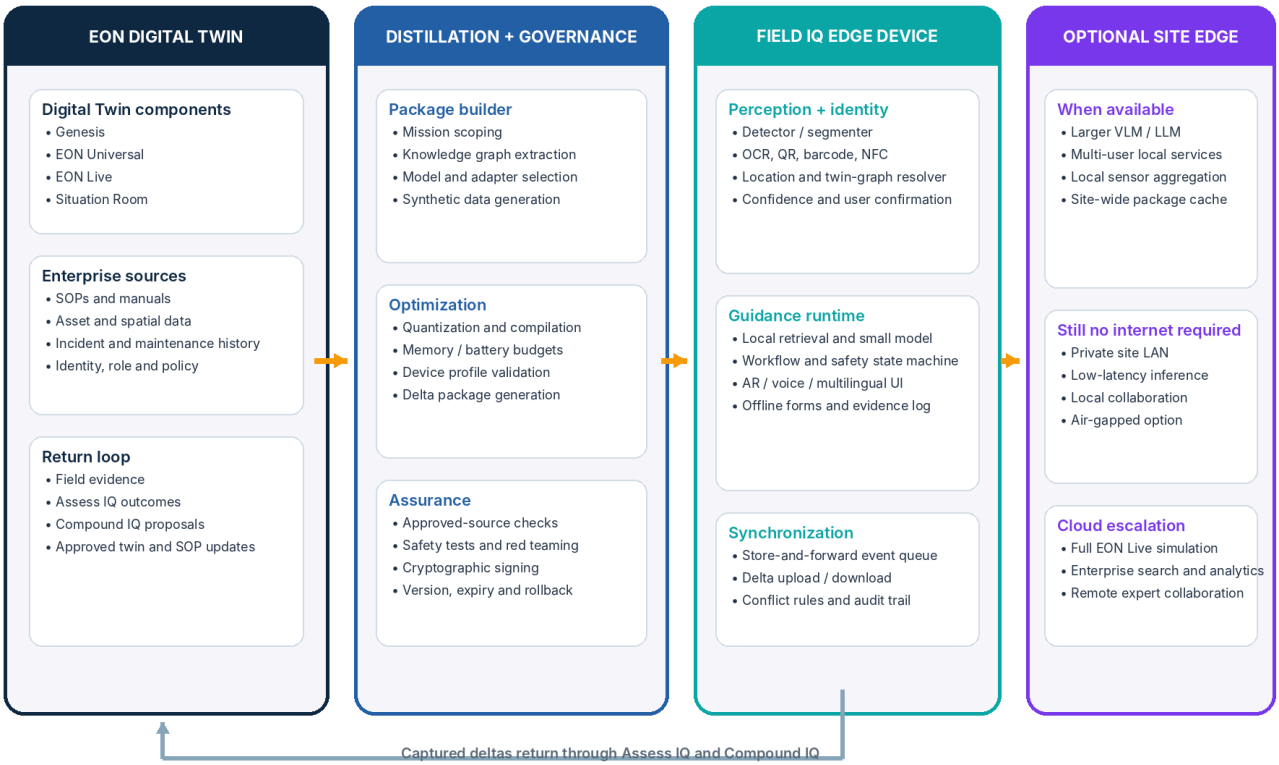

Figure 6. The reference architecture keeps core guidance on-device and adds optional site-edge and cloud tiers.

Edge runtime components

Component	Technical responsibility
Camera and sensor pipeline	Controls frame rate, image quality, stabilization, depth, optional thermal input and directly paired instruments.
Perception engine	Runs object detection, component segmentation, OCR and state estimation using platform-optimized models.
Asset identity resolver	Fuses class, label, tag, location, relative geometry, work order and user confirmation.
Local knowledge service	Searches encrypted documents, structured facts, embeddings and the Digital Twin subset; returns sources with answers.

Component	Technical responsibility
Local reasoning model	Interprets questions, summarizes approved information, selects tools and explains workflow state.
Workflow / safety engine	Maintains deterministic state, prerequisites, stop conditions, role permissions and evidence requirements.
Spatial presentation	Renders component highlighting, arrows, animations and 3D guidance using platform AR capabilities.
Evidence ledger	Stores actions, timestamps, media, readings, model outputs, confidence, user confirmations and version provenance.
Sync engine	Performs store-and-forward, resumable delta transfer, conflict handling, revocation and package update.

Platform runtime strategy

EON should remain model- and vendor-portable. On Apple devices, Core ML and the newer Core AI stack provide on-device execution using Apple silicon, while Apple also exposes an approximately 3B-parameter on-device foundation model for supported generative tasks. [8-10] On Android and cross-platform deployments, Google LiteRT and its LLM tooling, PyTorch ExecuTorch and ONNX Runtime Mobile provide viable paths for hardware-accelerated local inference. [11-15] Qualcomm AI Hub can optimize and profile models for Snapdragon-class devices, and NVIDIA Jetson provides an optional site-edge platform when a phone or tablet is not sufficient. [14,16]

The product should implement a model abstraction layer with a common contract for detection, OCR, embedding, language inference and safety evaluation. Each supported hardware profile receives compiled artifacts and benchmarks rather than assuming that one model file performs equally across all devices.

Model architecture

- Primary object detector / segmenter: specialized by asset family or site; optimized for continuous camera use.
- OCR and tag reader: platform-native or open model; used for asset labels, gauges, placards and part numbers.
- Re-identification / embedding model: distinguishes visually similar assets and components.
- Small language model: approximately mobile-scale, quantized and tuned for tool use, explanation and bounded reasoning.

- Optional compact VLM: handles scene-level questions and ambiguity, but does not become the sole authority for precise asset localization.
- Safety classifier / guard: detects unsupported requests, procedure conflicts, missing prerequisites and out-of-domain conditions.

Research demonstrates mobile-oriented VLMs in roughly 1-3B parameter ranges and even smaller multimodal families, but fine-grained industrial localization remains challenging. [20-21] The recommended architecture therefore uses a VLM as a flexible assistant layered over specialized detectors, not as the only perception mechanism.

Asset identity resolution

Specific identity is the most important visual-design problem. Two pumps may look identical. Field IQ Edge should compute identity from multiple independent signals:

- visual class and distinctive features;
- OCR of the asset plate or local signage;
- QR, barcode, RFID or NFC when available;
- GPS, indoor positioning or area selection;
- nearby equipment and relative spatial arrangement from the Digital Twin;
- assigned work order and expected asset set;
- worker confirmation for critical tasks.

The system should present a confidence class and the evidence used. High-risk procedures should require either a strong multi-signal identity or explicit confirmation. When signals conflict, the correct behavior is to stop and resolve identity.

Offline retrieval and workflow

The local retrieval pipeline should index approved content at package-build time. Documents are chunked, linked to assets and procedure steps, embedded, encrypted and accompanied by source metadata. At runtime, the language model receives only the relevant retrieved content, current workflow state and authorized tools. The answer includes a source, revision and validity date. The workflow engine, not the language model, determines whether the next step is allowed.

Synchronization and conflict management

Field IQ Edge should use an append-only event model. Every observation, action and decision is stored locally with an immutable event ID, device time, user identity, package version and model version. Synchronization is resumable and idempotent. If the same work order changes centrally while the device is offline, policy determines whether to merge, require supervisor review or

reject the stale completion. Critical package revocations and safety bulletins receive the highest update priority.

Three deployment patterns

Pattern	Use and trade-off
Device-only	Best for bounded workflows, strongest independence and lowest infrastructure. Limited by device memory, battery and thermal budget.
Device + site edge	Recommended for large sites, air-gapped operations, multi-user services, local sensor integration and larger VLMs. Requires local infrastructure and lifecycle management.
Hybrid cloud	Adds full Digital Twin, enterprise data, high-fidelity simulation and global learning. Core task must still degrade safely when WAN access is lost.

RECOMMENDED DEFAULT

Device-first, gateway-optional, cloud-enhanced. Every critical assigned workflow is validated on the device alone; the gateway and cloud improve breadth and depth rather than basic continuity.

14. The Work Intelligence Flywheel

Field IQ Edge becomes strategically valuable because it is both an actuator of Work Intelligence and a sensor for operational reality. It carries the Digital Twin into the field, but it also captures the delta between the designed procedure and what actually happens.

Flywheel station	Role
Digital Twin	Genesis, EON Universal, EON Live and Situation Room create and maintain the current enterprise representation.
Field IQ Edge	Applies Work Intelligence at the point of work and captures actions, evidence, conditions, questions and exceptions.
Assess IQ	Observes, scores and verifies execution, competency, outcomes and evidence.

Flywheel station	Role
Compound IQ	Aggregates validated discrepancies, identifies patterns and proposes improvements to procedures, models and twins.
Digital Twin - next cycle	Approved improvements become part of the trusted source for the next Work Intelligence Distillation run.

EON publicly describes Compound IQ as the learning loop that converts discrepancy telemetry from Assess IQ into proposed improvements to Genesis twins and SOPs, gated for safety. [4-5] Field IQ Edge strengthens this loop by making the capture station resilient in the exact environments where connectivity is weakest and operational deltas may be most valuable.

The governed delta

Not every deviation is an improvement. A worker may discover a safer sequence, compensate for a site condition, make an error or bypass a required step. The flywheel must classify the delta before it compounds. A recommended lifecycle is: capture -> reconstruct -> compare -> classify -> validate -> approve -> update -> redeploy. Raw field behavior should never be inserted automatically into a trusted procedure or model.

15. Safety, Trust and Governance

Field IQ Edge will be used in environments where a plausible but incorrect answer can cause harm. Trust must therefore be designed into the system architecture rather than added as a disclaimer. NIST recommends lifecycle-based governance, mapping, measurement and management of AI risk, while IEC 61508 provides a risk-based functional-safety framework for relevant electrical and programmable systems. [24,27] The exact standards profile will vary by industry and use case.

Authority hierarchy

1. Applicable law, regulation and certified safety systems.
2. Customer-approved policies, SOPs, task cards and clinical or maintenance instructions.
3. Deterministic Field IQ Edge workflow and safety rules.
4. Human supervisor or qualified professional decision.
5. AI-generated explanation, suggestion or summary.

The AI layer must never silently outrank the approved source or deterministic safety layer. It may explain, retrieve, ask questions, summarize and propose. It may not waive prerequisites, alter lockout-tagout, improvise medication instructions or authorize a maintenance release.

Required trust controls

Control	Implementation principle
Approved-source grounding	Every procedural answer is linked to a package source, revision and validity window.
Confidence-aware UX	Low confidence triggers clarification, alternative identity checks or escalation - not a confident guess.
Safe defaults	Malformed model output, missing assets, expired packages and policy conflicts lead to a defined safe state.
Human confirmation	Critical asset identity, hazardous steps and exception handling require explicit confirmation or supervisor approval.
Separation of duties	Model builders, content approvers, safety approvers and deployment administrators have distinct roles.
Evaluation before release	Each package is tested on representative devices, lighting, PPE, noise, language and failure conditions.
Auditability	Inputs, retrieved sources, workflow transitions, model versions, user actions and evidence are reconstructable.
Rollback and revocation	A defective package or model can be disabled and replaced without waiting for a full application release.

Generative AI attack surface

Prompt injection and insecure model outputs remain relevant even offline because a camera may read malicious text on a label or document and a local model may be asked to ignore policy. OWASP identifies prompt injection, insecure output handling, sensitive-information disclosure and supply-chain risk among the principal LLM application threats. [26] Field IQ Edge should separate trusted instructions from untrusted observed content, validate tool calls against an allowlist and treat external text as data rather than authority.

Industry-specific governance

- **Oil & gas**: align procedures with the operator's process-safety, permit-to-work and functional-safety framework; never bypass certified control or shutdown systems.
- **Aerospace**: treat approved maintenance data, configuration control, tool calibration, sign-off and release authority as controlling records.
- **Healthcare:** maintain clinical governance, privacy, infection-control and device policies; separate training and procedural support from regulated diagnosis or treatment.

16. Security and Intellectual Property

Security architecture

Offline operation improves privacy and resilience, but it also places valuable models and customer knowledge on devices that may be lost, stolen or physically accessed. Security should follow a zero-trust, least-privilege and defense-in-depth model consistent with the organization's NIST Cybersecurity Framework profile. [25]

- Hardware-backed device identity and key storage where available.
- Encryption at rest for models, content, indexes and evidence; encryption in transit for all synchronization.
- Signed packages, verified before activation, with anti-rollback and revocation support.
- Role-based access, cached credentials with expiry, offline authorization policy and remote device management.
- Data minimization: download only the assets and content required for the mission.
- Secure deletion on logout, reassignment, package expiry or device compromise.
- Software bill of materials, model provenance and license tracking for open-source components.
- Air-gapped deployment option with controlled media transfer and offline verification.

The EON IP opportunity

The long-term moat is not the fact that a small open-source model runs offline. Model families and hardware will change rapidly. The defensible value is the system that converts a customer's Digital Twin and operational knowledge into safe, specialized field intelligence, then improves it through governed evidence.

Protection mechanism	Potential role
Trade secrets	Data curation, synthetic-data recipes, task scoping, package optimizer, model routing, evaluation harness, safety compiler and deployment operations.
Copyright / database rights	Industry ontologies, curated examples, structured procedures, synthetic assets, benchmark sets, documentation and package schemas where protectable.
Model and data assets	EON-owned industry adapters, recognition datasets, embeddings, evaluation sets and generic domain models, subject to source licenses.
Trademarks	Field IQ Edge and Work Intelligence Distillation may be evaluated for trademark strategy as terms are finalized.
Patents	Specific novel methods may be candidates, subject to prior-art search, inventorship documentation and legal review.
Contracts	Clear ownership, permitted learning, confidentiality, derivative-data rights, model improvement and customer isolation.

Candidate invention disclosures

- Risk- and mission-aware generation of a device-bounded intelligence package from a multi-component Digital Twin.
- Digital-twin-driven synthetic visual training combined with field evidence to produce site-specific on-device recognition.
- Multi-signal asset identity resolution using vision, tags, location, relative spatial graph and workflow context.
- Compilation of human-readable SOPs, twin relationships and safety policy into a synchronized local AI-plus-state-machine runtime.
- A governed discrepancy loop that converts offline field evidence through Assess IQ and Compound IQ into validated package updates.
- Confidence-aware routing among device, site edge and cloud tiers while preserving workflow state and evidence provenance.

LEGAL CAUTION

These are candidate invention themes, not a conclusion that claims are novel or patentable. EON should conduct a focused prior-art search and file invention disclosures before public launch details eliminate options in relevant jurisdictions.

Open-source model strategy

EON should use open and commercial models through a replaceable model layer. Every model must pass a license, export, security, provenance, performance and data-use review. Customer value should never depend on claiming ownership of the underlying foundation model. EON owns the distillation system, domain assets, package architecture, governance, evaluation and customer experience; the customer owns its defined private content and operational data.

17. Business Value

Field IQ Edge creates value by turning offline resilience into operational continuity while also strengthening the Work Intelligence flywheel. The business case should be measured against specific workflows rather than broad AI claims.

Value area	How Field IQ Edge contributes
Safety and procedural consistency	Approved guidance and stop conditions remain present in dead zones and air-gapped areas.
Reduced search and escalation time	Workers retrieve asset-specific information locally rather than leaving the task or waiting for an expert.
Faster time to competency	The same Digital Twin content supports training, field guidance and objective assessment.
Expertise preservation	Retiring-expert knowledge is delivered in context and refined by evidence from future work.
Lower cloud dependency and inference cost	Common recognition and guidance occur on-device; the cloud is used selectively.
Data sovereignty	Sensitive images and customer knowledge can remain on the device or site until approved synchronization.
Audit and compliance readiness	Evidence, sources, model versions, confirmations and workflow transitions are captured systematically.
Compounding IP	Every governed cycle can improve EON industry intelligence and the customer's private Work Intelligence asset.

Recommended value metrics

- Time required to identify the correct asset and retrieve the approved procedure.
- Percentage of assigned workflows completed without internet or expert intervention.
- Recognition and identity-confirmation accuracy by asset, lighting and PPE condition.
- Critical-step adherence, stop-rule activation and unresolved-condition escalation.

- Time to synchronize and reconcile evidence after reconnection.
- User trust, adoption, cognitive load and preference versus current paper or mobile methods.
- Reduction in repeat questions, incomplete evidence and procedure-search time.

18. Development Roadmap

Phase	Indicative horizon	Primary outcome
Phase 0 - Foundation and data readiness	0-6 weeks	Select the pilot site and device, map assets and procedures, establish ownership, collect representative imagery, define safety envelope and benchmark current process.
Phase 1 - Offline foundation	6-12 weeks	Package delivery, encrypted content, QR/OCR, known-asset recognition, SOP retrieval, workflow state machine, local evidence and synchronization.
Phase 2 - Specialized edge intelligence	3-6 months	Compact local language model, component recognition, local voice, trusted Q&A, role-aware guidance, customer and industry layers.
Phase 3 - Adaptive guidance	6-12 months	Bounded abnormal scenarios, local assessment, spatial overlays, site-edge option, advanced identity resolver and package factory automation.
Phase 4 - Scaled Work Intelligence Distillation	12-24 months	Reusable industry datasets and adapters, automated evaluation, fleet package management, cross-device optimization and governed learning at scale.
Phase 5 - Physical AI expansion	Beyond 24 months	Robots, drones, wearables, multi-device collaboration and deeper local simulation while preserving the same package and governance model.

Build-versus-partner recommendation

Strategy	Recommended scope
Build and own	Work Intelligence Distillation, package schema, asset graph, safety workflow compiler, identity resolver, evaluation, sync lifecycle, flywheel integration and user experience.

Strategy	Recommended scope
Use / adapt	Foundation language and vision models, platform runtimes, OCR, speech engines, vector libraries, device management and commodity cryptography.
Partner where useful	Rugged hardware, smart glasses, local edge appliances, MDM, OT connectors, standards validation and industry-specific certification.

19. Pilot Specification

The first pilot should prove a complete operational loop rather than demonstrate isolated AI features. Oil & gas is the recommended starting point because the environment combines known assets, repetitive equipment families, clear SOPs, high consequence, strong digital-twin value and real connectivity constraints.

Recommended 90-day pilot scope

Pilot dimension	Recommendation
Facility scope	One operating area or training skid with controlled access and representative connectivity loss.
Asset scope	20-50 tagged assets across two or three equipment classes, including visually similar units.
Workflow scope	One inspection workflow, one routine maintenance workflow and one bounded abnormal-condition scenario.
Users	A small group of experienced technicians, newer technicians, supervisors and content approvers.
Device	One reference phone or tablet profile; optional smart-glass display after the core tablet workflow is stable.
Offline duration	Complete core workflows with airplane mode enabled; separately test site-local gateway mode.
Integration	Package build from the Digital Twin; evidence to Assess IQ; selected deltas reviewed through Compound IQ.

Candidate acceptance criteria

Acceptance area	Pilot test
Mission completeness	Every required instruction, source, safety rule and form for the pilot workflow remains available in airplane mode.
Asset identity	Meets a pre-agreed accuracy threshold under defined viewpoints, lighting, PPE and occlusion; critical tasks require multi-signal confirmation.
Source fidelity	All procedural guidance traces to the correct approved source and revision; no unsupported procedural steps are generated.
Latency	Recognition, retrieval and first guidance response remain usable on the reference device under sustained field conditions.
Safety behavior	Expired content, identity conflict, low confidence, malformed model output and out-of-scope conditions produce the defined safe fallback.
Evidence integrity	Actions, images, readings, versions and confirmations are preserved during offline operation and synchronize without duplication or loss.
User outcome	Technicians can complete the workflow with lower search friction and acceptable cognitive load compared with the current method.
Flywheel closure	At least one field delta is reconstructed, classified, reviewed and converted into a proposed Digital Twin or procedure improvement.

Pilot deliverables

- Validated reference architecture and device profile.
- Initial oil & gas industry intelligence layer and site package.
- Recognition dataset, benchmark and failure analysis.
- Safety case, threat model and package approval checklist.
- Measured offline capability matrix and performance report.
- Prioritized backlog for the product requirements document and engineering specification.
- Invention-disclosure package for the strongest novel implementation methods.

20. Strategic Vision

The future of enterprise AI is not one giant model answering every question from a distant cloud. It is a coordinated system that delivers the right intelligence at the right place, with the right authority, on the right device. Field IQ Edge makes that future operational.

The EON Digital Twin - Genesis, EON Universal, EON Live and Situation Room - remains the source of enterprise context. Work Intelligence Distillation converts that context into compact, governed field capability. Field IQ Edge recognizes, guides and captures beyond connectivity. Assess IQ verifies what happened. Compound IQ turns validated differences into improved Work Intelligence. The Digital Twin becomes richer, and the next worker starts from a stronger position.

NORTH STAR

The Digital Twin should not end at the screen, the control room or the network boundary. Field IQ Edge carries it into the work itself.

This is where the strategic value resides. Offline capability is the visible product news. The deeper asset is EON's ability to manufacture, govern, deploy and continuously improve specialized Work Intelligence for the physical world while respecting customer ownership and safety boundaries.

Appendix A. Design Principles

Principle	Meaning
Bound the mission	A smaller, defined problem produces a safer and more useful offline system than unrestricted general reasoning.
Ground identity	Use multiple signals and the Digital Twin graph; never rely on a generic visual guess for a critical asset.
Separate truth from language	Documents, graphs and workflows are authoritative. The model explains and orchestrates within constraints.
Prefer local for the common path	Recognition, retrieval, sequence and evidence should not require WAN connectivity.
Escalate uncertainty	Low confidence is a workflow event, not an invitation to improvise.
Design for updates	Procedures change more often than models. Keep mutable facts in updateable packages.
Make every action reconstructable	Capture source, version, identity evidence, user confirmation and outcome.
Treat data rights as architecture	Customer isolation, permitted learning and deletion must be implemented, not merely promised.
Benchmark on the real device	Model size alone does not predict latency, heat, battery or camera throughput.
Close the flywheel deliberately	Only validated operational deltas should improve the trusted Digital Twin.

Appendix B. Technical Options and Model Strategy

The final model choices should be selected by task benchmark, license and target hardware rather than brand recognition. The following categories illustrate the current option space.

Category	Selection guidance
Mobile language model	A compact instruction-following model, typically quantized, for explanation, retrieval orchestration and structured tool use. Candidate families change rapidly; keep the interface replaceable.
Mobile vision-language model	Useful for scene-level interpretation and flexible questions. Validate fine-grained localization separately; compact research families include MobileVLM and SmolVLM. [20-21]

Category	Selection guidance
Detector / segmenter	MobileNet-, EfficientDet-, MobileViT- or equivalent efficient architecture, specialized to the known equipment catalog. [22-23]
Embedding / re-identification	Compact visual and text encoders tuned for asset and component similarity.
OCR / speech	Prefer platform-native engines where they meet offline, language and licensing requirements; retain replaceable open alternatives.
Inference runtime	Core ML / Core AI, LiteRT, ExecuTorch, ONNX Runtime Mobile or vendor acceleration selected per device. [8-15]
Site edge runtime	NVIDIA Jetson or industrial x86 / ARM edge computer for larger models and local services when needed. [16]

Recommended training approach

1. Start with a capable small base model under an acceptable license.
2. Train EON industry adapters on curated domain language, taxonomies, tool calls and safety-oriented examples.
3. Use retrieval rather than fine-tuning for most customer-specific procedural facts.
4. Train specialized vision models from Digital Twin synthetic renders plus customer-approved field images.
5. Use teacher models to annotate, generate hard negatives and create candidate scenarios, followed by expert review.
6. Quantize and compile for each hardware profile, then evaluate the complete system rather than model benchmark alone.
7. Retain a golden offline test set and regression suite for every package release.

Appendix C. Validation Framework

Test domain	Minimum evaluation coverage
Perception	Detection precision/recall, identity accuracy, OCR accuracy, state recognition and robustness to lighting, occlusion, dirt, PPE and camera motion.
Retrieval	Correct-source rate, citation correctness, coverage, stale-content detection and refusal outside package scope.

Test domain	Minimum evaluation coverage
Reasoning	Tool-selection accuracy, structured-output validity, bounded scenario performance, hallucination rate and uncertainty calibration.
Workflow	Prerequisite enforcement, step order, stop rules, evidence completeness, resume-after-interruption and stale-package behavior.
Human factors	Task time, cognitive load, trust calibration, readability, voice usability, multilingual quality and hands-free ergonomics.
Device	Cold and sustained latency, memory, storage, battery, thermal throttling, camera throughput and crash recovery.
Security	Package tampering, key compromise, prompt injection, malicious visual text, role escalation, data exfiltration and rollback.
Synchronization	Offline duration, queue integrity, resumability, deduplication, conflict resolution, revocation and disaster recovery.
Governance	Approver traceability, model and data provenance, change control, customer rights and incident response.

Appendix D. Glossary

Term	Definition
Digital Twin	The EON enterprise representation consisting in this paper of Genesis, EON Universal, EON Live and Situation Room.
Work Intelligence	The system of record for how work is understood, guided, performed, assessed and improved.
Work Intelligence Distillation	The process that creates a mission-scoped, governed and optimized Field IQ Edge package from the Digital Twin.
Field IQ Edge	The offline-capable deployment profile of Field IQ.
Edge package	The signed technical bundle containing models, knowledge, graph context, workflows, safety policy, experience assets and sync rules.

Term	Definition
Site edge	A local server or gateway available within a facility without requiring internet access.
Grounding	Connecting an AI output to a specific asset, source, location, workflow state and approved context.
Operational delta	The difference between the designed procedure or Digital Twin expectation and what occurred in the field.
Safe fallback	A predefined behavior used when identity, model output, content freshness or system state is not trustworthy enough to continue.

References

[1] EON Reality. "EON Introduces Work Intelligence (WI): the Enterprise's Missing System of Record." 2026.

[2] EON Reality. "EON Reality Launches Field IQ to Redefine Industrial Safety and Precision." 2026.

[3] EON Reality. "EON Field IQ." Product overview.

[4] EON Reality. "Assess IQ 2.0: Real-Time AI Safety Observation with Cryptographically Signed Evidence." 2026.

[5] EON Reality. "EON Launches Compound IQ - the Learning Engine that Makes Industrial Knowledge Appreciate Every Shift." 2026.

[6] EON Reality. "EON Enterprise Hub Revolutionizes Workforce Competency for High-Stakes Environments." 2026.

[7] EON AI Ventures. Company and platform overview.

[8] Apple. Core ML Developer Documentation.

[9] Apple Machine Learning Research. "Apple Intelligence Foundation Language Models" technical report. 2025.

[10] Apple Developer. "What's New - AI & Machine Learning" / Core AI. 2026.

[11] Google AI Edge. LiteRT overview.

[12] Google AI Edge. LLM Inference guide.

[13] PyTorch. ExecuTorch documentation and on-device deployment guidance.

[14] Qualcomm. Qualcomm AI Hub.

[15] Microsoft. ONNX Runtime Mobile.

[16] NVIDIA. Jetson for Robotics and Edge AI.

[17] Hinton, G., Vinyals, O., and Dean, J. "Distilling the Knowledge in a Neural Network." 2015.

[18] Hu, E. et al. "LoRA: Low-Rank Adaptation of Large Language Models." 2021.

[19] Dettmers, T. et al. "QLoRA: Efficient Finetuning of Quantized LLMs." 2023.

[20] Chu, X. et al. "MobileVLM V2: Faster and Stronger Baseline for Vision Language Model." 2024.

[21] Hugging Face. "SmolVLM2: Bringing Video Understanding to Every Device." 2025.

[22] Howard, A. et al. "Searching for MobileNetV3." 2019.

[23] Tan, M., Pang, R., and Le, Q. "EfficientDet: Scalable and Efficient Object Detection." 2019.

[24] National Institute of Standards and Technology. AI Risk Management Framework.

[25] National Institute of Standards and Technology. Cybersecurity Framework 2.0.

[26] OWASP. Top 10 for LLM Applications 2025.

[27] International Electrotechnical Commission. IEC 61508 Functional Safety overview.

All external technical capabilities and product descriptions should be rechecked immediately before public release because model, operating-system and product capabilities change rapidly. Accessed July 2026.

FIELD IQ EDGE

A Digital Twin Beyond Connectivity

The Digital Twin should not end at the network boundary.

With Field IQ Edge, it travels with the worker.

EON AI Ventures | Draft for Review | July 2026