

Appendix E


Keep the memory governed and the brain replaceable

The candidate portfolio is a working recommendation subject to **industrial benchmarking, licensing review, infrastructure testing, and production-cost validation.**



The EON Work Intelligence Flywheel


From Digital Twin to Field Learning and Back



The flywheel compounds Work Intelligence over time.

Zoom-in: How Field Learning Improves the Digital Twin


From field evidence to approved updates and better mission packages



Better learning goes back into the Digital Twin, and better mission packages go back to the field.

2. What the Distillery Actually Does


Three jobs, one system.



The Distillery does not replace the Digital Twin or the model. It connects them.

3. Open Model vs Databases

The model is the brain. The databases are the memory.



Important: if the model changes, the memory stays.
That means EON can switch models without losing its knowledge.



Closed frontier models can still help build features,
but EON should keep its core knowledge in its own databases and open-model stack.

OPEN-WEIGHT MODEL STRATEGY

Where each model fits in the EON Work Intelligence Platform

EON
APPLICATIONS



Genesis



EON Universal



Field IQ



Assess IQ



Compound IQ



Distillery



EON MODEL GATEWAY

Route • Orchestrate • Guardrail • Observe

1



Primary EON Brain — Qwen 3.6 27B

day-to-day reasoning, multimodal understanding, self-hosted

2



High-End Teacher / Challenger — Kimi K3

difficult reasoning, synthetic data, distillation

3



Agentic / Workflow Challenger — NVIDIA Nemotron

tool use, long context, agent workflows

4



Closed Frontier Helpers — OpenAI / Claude

build plane only, selective use, not the system of record



Customer Digital Twin

Assets • People • Locations • Relationships
History • Context



EON Work Intelligence Core

Documents • Procedures • Knowledge
Events • Lessons • Standards



Mission Intelligence Package

The right knowledge • Right format • Right time



Field IQ Edge

Deliver to the field • Guide • Capture • Verify

Core message



Keep the memory in EON repositories



Keep the brain replaceable




Use the best model for each job



Never store permanent value inside a rented model

Smarter Every Day. Built on Open Models. Secured by EON.



RECOMMENDED MODEL PORTFOLIO

Which models to use, and for what



1

**Primary
Production Model**

Qwen 3.6 27B

- ✓ Best balance
- ✓ Self-hosted
- ✓ Vision + language
- ✓ Strong day-to-day default



2

**Teacher /
Advanced Reasoning**

Kimi K3

- ✓ Very powerful
- ✓ Use selectively
- ✓ Great for distillation
- ✓ Not the everyday default



3

**Agentic
Challenger**

NVIDIA Nemotron

- ✓ Tool use
- ✓ Long context
- ✓ Great for workflows



4

Edge Models
4B–9B small models

- ✓ Runs on device
- ✓ Offline capable
- ✓ Use with recognition models




Closed models still help

OpenAI / Claude remain useful for coding, product development, and difficult non-sensitive tasks



Recommendation: Build around a replaceable open-weight portfolio.
Start with Qwen, test Kimi and Nemotron, deploy small edge models in the field.



PRICING & DEPLOYMENT COMPARISON

Relative cost and control across model options

	 Ownership / Control	 Infrastructure Cost	 Per-Query Cost	 Offline Capability	 Best Role
 OpenAI / Claude API	● Low	● Low	● Medium to High	● No	 Fast build help
 Qwen 3.6 27B self-hosted	● High	● Medium	● Low to Medium at scale	● No	 Primary EON brain
 Kimi K3 self-hosted	● High	● Very High	● Medium at scale	● No	 Teacher / challenger
 NVIDIA Nemotron self-hosted	● High	● High	● Low to Medium at scale	● No	 Agents & workflows
 Small edge model	● High	● Medium upfront	● Very Low after deployment	● Yes	 Field IQ Edge



The cheapest path upfront is rented API usage.
The strongest long-term moat comes from open-weight models plus EON-owned repositories and edge deployment.




OWNERSHIP ADVANTAGE: OPEN-WEIGHT VS CLOSED MODELS

Why open-weight architecture creates a stronger Work Intelligence asset



CLOSED MODELS

You rent the brain from a cloud vendor




- ✗ You may own your data, but you rent the brain
- ✗ Provider controls the model
- ✗ Harder to run offline
- ✗ Behavior, pricing, and availability can change
- ✗ Great for fast help and product development



OPEN-WEIGHT MODELS

EON owns the stack and the intelligence it builds



- APPLICATIONS
- ADAPTERS & AGENTS
- MODEL LAYER
- INFRASTRUCTURE

- ✓ EON controls where the model runs
- ✓ EON keeps the knowledge inside its system
- ✓ Can build domain adapters and edge models
- ✓ Can switch models without losing the memory
- ✓ Creates a compounding enterprise asset

VS



Open-weight does not mean EON owns the original foundation model. It means EON controls the deployment, tuning, packaging, and the compounding Work Intelligence built on top of it.



Repositories
= memory

+



Model
= brain

+



Distillery
= factory

+



Edge Package
= backpack

From Digital Twin to Field Learning and Back

Build. Learn. Compound. Repeat.



THE BIG PICTURE – EON WORK INTELLIGENCE FLYWHEEL



Large and small models serve different missions

The **edge model** is a **specialist**, not simply a weaker copy of the enterprise brain.

LARGE HOSTED MODEL

Cross-record analysis, Compound IQ support, procedure structuring, synthetic examples, broad reasoning, and edge-model evaluation.

SMALL EDGE MODEL

Local reasoning for one asset, workflow, permission set, evidence contract, and set of stop conditions.



Three assets remain distinct

Memory, reasoning, and manufacturing play different roles.

CUSTOMER DIGITAL TWIN


Private operational memory: assets, relationships, procedures, history, evidence, and site knowledge.

EON WORK INTELLIGENCE CORE

Authorized reusable knowledge: ontologies, equipment families, patterns, evaluations, and generalized intelligence.

MODEL PORTFOLIO

Replaceable reasoning, perception, speech, and edge capability selected for each workload.



The Knowledge Firewall preserves customer ownership

No customer learning crosses boundaries automatically.

- **Classify** whether evidence is customer-identifiable, facility-specific, sensitive, or permitted for broader use.
- **Generalize** only when technical meaning can be preserved without transferring customer-owned knowledge.
- **Record** permission, provenance, repository destination, and reuse rights before storage or training.



Governance is part of the control architecture

Intelligence authority and **execution** authority remain **separate**.

- **AI may retrieve, compare, recognize, explain, summarize, and propose.**
- **Deterministic workflow and customer policy govern** whether a step is permitted and whether prerequisites are satisfied.
- **Accountable roles approve consequential changes;** the platform preserves reconstruction and rollback.



Governance turns model output into accountable operational decisions

Mission **authorization**, **evidence**, **safety boundaries**, and **escalation** remain **visible** to responsible people.



The Safety Envelope is deterministic

Probabilistic recognition can contribute, but it cannot override mission controls.

PREREQUISITES

Identity, permits, isolation, equipment state, and required evidence.

OPERATING LIMITS

Approved ranges, prohibited actions, and role-specific permissions.

STOP CONDITIONS

Missing evidence, conflicting state, exceeded limit, or uncertain identity.

SAFE-STATE BEHAVIOR

Escalation, hold, rollback, or return-to-safe-location instructions.



Mission authorization and auditability create accountable execution

The Mission Contract **defines who may do what, where, when**, and with which evidence.

- **Verify actor identity and role, device and package status, asset and site context, and required approval before work begins.**
- Link package version, model and adapter versions, source revisions, actions, exceptions, evidence, assessment, proposals, approvals, and released change.
- The purpose is reconstructability and accountability—not surveillance for its own sake.



Packages require expiry, revocation, supersession, and rollback

Offline capability moves controls into the package and device runtime.

- **Know which devices hold each release and prevent further use when a procedure changes or a safety issue is discovered.**
- Encrypt packages and evidence, protect signing keys, minimize edge data, and preserve controlled synchronization.
- Rollback must restore a known approved state without erasing the audit trail.



Required model bake-off

Production fitness must be proven on EON industrial tasks.

- Procedure interpretation, grounded retrieval, equipment-image recognition, and multimodal reasoning.
- Digital Twin graph reasoning, Safety Envelope adherence, refusal, stop, and escalation behavior.
- Tool calling, workflow completion, evidence production, hallucination, provenance, confidence, and adversarial robustness.
- Latency, throughput, hardware footprint, quantization, total cost per successful mission, licensing, data rights, and deployment constraints.



Closing synthesis

Work Intelligence makes experience a governed enterprise asset.



- **Capture operational reality in the Digital Twin.**
- Deliver bounded Mission Intelligence to the executor.
- Measure the Operational Delta and propose evidence-backed learning.
- Govern ownership, safety, authority, and release.
- Make the next mission better because the previous mission happened.

