

EON Control

A 360 degree model for knowledge capture and performance evidence



Table of Content

© 2026 EON AI Ventures. All Rights Reserved.....	0
Table of Content.....	1
Executive perspective	2
The scope of the 360 degree model.....	2
The basis for differentiation.....	2
An evidence model that preserves meaning.....	3
Different signals answer different questions.....	3
No universal score by default.....	3
Architecture and authority.....	4
Trust at the integration boundary	4
Corrections and longitudinal review.....	4
Expertise capture and human judgment	5
Field evidence and improvement	5
Responsible information handling.....	5
A staged implementation programme	6
Measuring value without overstating it	6
Current completion boundary	6
References and evidence.....	7
1 Product evidence	7
2 Experience API.....	7
3 EON ecosystem.....	7
4 Expert knowledge	7
Terminology	7

Executive perspective

Organisations need evidence of what people understand and what they can do in context. Assessment results, expert knowledge, simulations and field observations often answer different questions. Connecting them can make the next learning action more precise, provided the system preserves provenance, permissions and the limits of each observation.

EON Control's proposed role is a common administration and evidence layer across this ecosystem. Its local implementation already supports substantial learning administration, immutable accepted records, supported scoring profiles and audited reporting. The full cross-product integration remains a development and acceptance programme.

The scope of the 360 degree model

The intended evidence perimeter includes MCQ, oral interviews and textbook reading; Genesis; EON Live when work goes right and when conditions may go wrong; the Situation Room; tacit-knowledge interviews and learning grounded in experts with 30 years of experience; Field IQ; Permit IQ; and AssessIQ.

These are complementary perspectives, not interchangeable measures. A broad view should help the reviewer identify which question remains unanswered. It should not imply that collecting all available data creates a validated competence judgment automatically.

The basis for differentiation

Learning-record aggregation has established precedents. The ADL Experience API describes statements as evidence of experiences and events, and EON has already published a direction for capturing expert knowledge. A defensible positioning therefore focuses on the proposed EON-specific combination, its operational context and the quality of its traceable evidence. This paper does not assert a proven worldwide first. [2–4]

An evidence model that preserves meaning

Every accepted learning observation needs an attributable actor, source product, activity identity and version. Where applicable, it also needs an attempt identity, event sequence, capture and receipt times, result status, original scale, definition or rubric version, context and evidence references. This is a target integration discipline; the mounted Control contract supports a defined subset rather than every possible canonical source field. [1]

Different signals answer different questions

MCQ and oral interviews can establish different aspects of knowledge under a defined assessment. Reading activity records engagement and requires a separate comprehension assessment. Simulation evidence describes performance in its scenario. EON Live observations distinguish successful actions from deviations and recovery. Situation Room evidence must separate individual contributions from team outcomes.

Tacit knowledge captures the expert's reasoning, conditions and exceptions for review. Field IQ observations describe work execution under a defined method. Permit IQ preparation evidence should distinguish form completion from approved permission. AssessIQ outcomes need their native definition, method and scale. Mapping these signals requires a contract for each source; a product name alone is insufficient.

No universal score by default

Percentages from unrelated assessments are not automatically comparable. A short recall test, an oral rubric and a field observation can have different scales, confidence and consequences. Control preserves native results and supports calculated projections. A combined score, readiness threshold or trend rule requires explicit owner-approved policy and validation. Unknown data should remain unknown rather than being represented as zero.

Architecture and authority

EON Platform remains the shared identity and licensing authority. Each source product owns its execution environment, content and raw media. Control receives approved learning evidence, applies its supported validation and scope rules, and supplies authorised administration and review. Genesis and XR Infinite remain separate integrations. A shared person mapping must not silently merge product identities or grant access. [1]

Trust at the integration boundary

A production feed requires verified working grants, exact product bindings, approved actor mappings, version semantics and clear treatment of duplicate, late and abandoned events. Staged source observations and accepted learning records must remain distinguishable. A transport success is not proof that the result is semantically correct.

Current local controls include scoped database operations, immutable accepted facts and audit records, and supported strict scoring contracts. Original and calculated results remain distinguishable. Required human assessments remain pending until the appropriate review occurs. Local automated tests are valuable evidence of these behaviours, but do not substitute for producer acceptance or managed-database concurrency testing.

Corrections and longitudinal review

A correction should append its reason and authority without rewriting the original event. Reports can then show the current projection and its history. Comparisons across time need stable definitions or an explicit explanation of changed definitions. Observed first-pass reports do not establish lifetime attempt counts when earlier source history is unknown. Queue reconciliation likewise requires a producer backlog and watermark contract.

Expertise capture and human judgment

An expert interview should ask what was noticed, why a decision was taken, which alternatives were rejected and under what conditions the advice would fail. Thirty years of experience can provide rich material, but tenure does not make every statement correct or current. The owning knowledge workflow should preserve the original contribution, its interpretation, review state and the approved learning version.

A learner interview based on that material tests the learner's reasoning. It does not reproduce the expert's entire competence. Where an AI-supported interaction contributes evidence, its role and limitations should be explicit and consequential judgments should follow approved assessment and review policy.

Field evidence and improvement

An observation of successful work can reveal useful decision patterns. An observation of a deviation may reveal a training need, a confusing procedure, equipment conditions or a need for supervision. The reviewer should consider those explanations before attributing the outcome solely to the worker.

The proposed loop links reviewed evidence to a focused assignment, then to a later observation under a stated criterion. Field IQ, Permit IQ, AssessIQ and simulation sources can contribute to that loop after their mappings and acceptance are established. The system should preserve assistance and intervention rather than presenting assisted performance as independent execution.

Responsible information handling

Capture only the evidence needed for the approved purpose. Define access, retention, residency and source-media permissions before connecting a feed. Keep raw media in its owning system unless an approved architecture says otherwise. A reference or hash alone does not prove access, authenticity or continued availability. Operational authorisation remains with the designated authority; neither an interview nor a learning certificate automatically grants permission to perform hazardous work.

A staged implementation programme

First, select one defined learning and work scenario and its participating products. Agree the person mapping, activity and rubric versions, permitted fields and expected result states. Next, verify the source contract through approved synthetic trials and accepted-record readback. Then complete the learner, instructor and administrator journey in a managed test environment.

Acceptance should include duplicate and retry behaviour, revoked access, late results, missing evidence, human assessment, corrections and export completeness. Multi-connection PostgreSQL tests must cover conflicts and fencing. A production pilot also needs deployment, security, recovery and operational monitoring evidence. Expand only after the appropriate owners accept the bounded workflow.

Measuring value without overstating it

Useful pilot measures include the proportion of expected records received and reconciled, time to resolve an evidence gap, assessor agreement on a defined rubric, and improvement on comparable follow-up tasks. Define the population, denominator, observation period and exclusions before measuring. Avoid attributing a change to the product without considering changes in task difficulty, assistance and participants. No customer impact percentages are claimed in this paper.

Current completion boundary

The latest packaged local release passed 1,677 tests and a frontend build. The current source checkpoint adds requirement reconciliation and maximum-page test evidence. Full live 360 degree integration is incomplete. Production identity/licensing, producer contracts, owner policies, interactive acceptance and managed deployment remain open. Async evidence requests and receipts have internal foundations, but no mounted worker, storage adapter or pollable retrieval service. [1]

References and evidence

Sources support the architectural context and current completion boundary. Public ecosystem descriptions do not establish live integration into Control.

1 Product evidence

Local release 57cf4aa, 27 September 2026; CURRENT_REQUIREMENT_COVERAGE.csv at 3f766ff; 1,677 passing tests at release, followed by eight focused vocabulary tests.

2 Experience API

ADL xAPI specification, <https://github.com/adlnet/xAPI-Spec/blob/master/xAPI-Data.md> (accessed 27 September 2026). Describes statements as evidence of experiences and events.

3 EON ecosystem

<https://eonreality.com/> (accessed 27 September 2026). Public product positioning includes MCQ, oral exams, Assess IQ and Field IQ. Product positioning does not prove a working Control integration.

4 Expert knowledge

<https://eonaiventures.com/eon-ai-ventures-unveils-eon-tacit-spatial-brain/> (19 August 2026; accessed 27 September 2026). Describes the expanded direction for capturing expert experience with context and evidence.

Terminology

Accepted record: an event or result that passed the implemented Control contract.

Staged observation: source material retained for review without implying accepted learning credit.

Native scale: the source-defined measurement range or rubric, preserved with its definition.

Current projection: the present interpretation after supported review or correction, with original facts retained.

360 degree evidence: complementary perspectives across learning, judgment and work, with their boundaries visible.