

WHITE PAPER

The BoundaryScore™ Assurance Architecture

How bounded evaluation, zero-trust evidence, independent observation, provenance and reproducibility strengthen verification for Physical AI

Executive Summary

Physical AI assurance depends on more than checking motion or comparing a reported metric with an external observation. Confidence depends on the architecture around the conclusion: what exactly was evaluated, which evidence was trusted and why, how reported state was reconciled with observed reality, how uncertainty was handled, and whether the result can be challenged and reproduced.

The BoundaryScore™ assurance framework is designed around those requirements. It binds every conclusion to a defined system and operating boundary, treats every signal as evidence rather than unquestioned truth, separates independent observation from machine self-report, reconciles the two through MVE™, and keeps performance distinct from evidence confidence. Provenance, versioning, applicability and the audit trail remain part of the result.

When the applicable requirements are met, BoundaryScore™ Verified can provide an assurance designation around the metrics a company, customer, insurer, regulator or other authorized stakeholder already uses.

IN THIS PAPER

1. Assurance architecture

Why verification requires more than a second measurement.

2. Boundary first

How the evaluation contract limits every conclusion.

3. Zero-trust evidence

How independent and machine-state evidence are handled.

4. Reconcile before evaluation

How MVE™ creates a defensible event record.

5. Preserve confidence and change

How performance, confidence, applicability and traceability stay visible.

6. BoundaryScore™ Verified

How the architecture supports trusted verification of stakeholder-defined metrics.

1. Assurance requires an architecture, not just an extra sensor

An independent physical reference is valuable because it creates an observation path outside the machine being evaluated. But independent observation alone does not establish a trustworthy conclusion. The assurance method also has to establish identity, configuration, timing, provenance, uncertainty, exposure and the rules used to interpret the evidence.

The BoundaryScore™ assurance framework therefore separates the assurance process into distinct functions. Evidence sources provide attributable observations. MVE™ authenticates, binds, aligns and reconciles those observations. The applicable metric profile defines how validated events are evaluated. The result package keeps performance, evidence confidence, applicability and the audit record together. This separation prevents a performance rule from changing what the evidence says and prevents weak evidence from disappearing inside a favorable metric.

2. Start with the boundary

Every BoundaryScore™ evaluation begins with an explicit evaluation contract. It identifies the exact system or subsystem, hardware and software configuration, mission or function, autonomy and oversight model, operating domain and context, exposure, evidence sources, metric profile, evaluation window, dependencies and exclusions.

That boundary is a validity control. A conclusion from one software version, payload, site, weather condition, operating domain or exposure period does not automatically describe another. BoundaryScore™ treats applicability as part of the result rather than an assumption made after the fact.

The assurance path

BOUNDARY	EVIDENCE	MVE™	METRICS	RESULT
Define the exact unit, configuration, mission, conditions and exposure.	Keep machine-state claims and independent physical observations attributable.	Authenticate, bind, align, reconcile and assess uncertainty.	Apply the stakeholder-defined metric profile to validated event records.	Performance + confidence + applicability + audit trail.

3. Zero-trust evidence: every signal must earn confidence

The strongest assurance architectures do not depend on a single unquestioned source. BoundaryScore™ applies that principle directly to evidence: a machine log is not authoritative merely because the machine produced it, and an external sensor is not authoritative merely because it is external. Each source contributes only to the extent that its identity, integrity, provenance, timing, calibration, coverage, uncertainty and independence support the conclusion.

Principle	How the BoundaryScore™ assurance architecture implements it
Bounded claims	No conclusion exists until the evaluated unit, configuration, function, operating conditions, exposure and responsibility model are explicit.
Independent observation	Physics-based evidence is obtained through a materially separate observation path where feasible, providing a check on machine-reported state.
Evidence, not truth	Machine self-report and external observation are both treated as claims with identity, integrity, uncertainty and limitations.
Authentication + provenance	Source identity, configuration binding, timing, calibration and chain of custody determine whether evidence can support a conclusion.
Separation of functions	Evidence acquisition, reconciliation, metric application and reporting remain distinct so one layer cannot silently overwrite another.
Challenge + reproduction	Every conclusion retains a versioned path back to the metric, event, evidence, transformations and calculation method.

This is a zero-trust evidence principle in a specific assurance sense. It does not claim cybersecurity certification. It means the architecture requires evidence to earn trust through attribution, corroboration and explicit limitations rather than inheriting trust from its source.

4. Reconcile reported state with observed reality before applying metrics

The Motion Validation Envelope™ (MVE™) is the bounded reconciliation layer between raw evidence and performance evaluation. It compares machine-state evidence with independently obtained physical evidence while accounting for latency, temporal and spatial uncertainty, calibration, measurement precision and expected system dynamics.

MVE™ does not simply ask whether two signals are numerically identical. It determines whether the available evidence supports a congruent event record, preserves meaningful discrepancies, records uncertainty and provenance, and identifies when the evidence is insufficient. A discrepancy is not automatically a failure; it is an evidence-supported fact that must be interpreted under the applicable metric profile.

The result package keeps four outputs together

Result element	What it preserves
Performance	How the system performed against the defined metric or requirement, using validated event records and authenticated exposure.
Evidence confidence	How strongly the available evidence supports that conclusion, including limitations, coverage and uncertainty.
Applicability	The exact unit, configuration, mission, domain, context, exposure period and profile to which the conclusion applies.
Audit record	The reproducible path from the conclusion to the metric, event, evidence, transformation and calculation method.

5. Preserve confidence as systems and conditions change

Physical AI systems are not static. Software changes, sensors drift, payloads vary, environments change, humans intervene and adaptive systems may evolve. BoundaryScore™ therefore does not treat a verified conclusion as a permanent attribute of a product name.

Material changes in hardware, software, configuration, mission, operating domain, evidence quality or methodology can trigger reevaluation. Versioning preserves what changed and which evidence supported each conclusion. Critical events remain visible rather than being averaged away, and insufficient or materially contradictory evidence can produce qualification or abstention rather than false precision.

6. BoundaryScore™ Verified

BoundaryScore™ Verified is the assurance designation applied when a defined metric or performance claim meets the applicable verification requirements. The underlying metric can remain the one already used by the company or stakeholder. BoundaryScore™ Verified adds the evidence architecture around it: a bounded evaluation, attributable and reconciled evidence, explicit confidence, applicability and a reproducible audit trail.

Greater confidence inside the company

Engineering, safety, quality, operations and executive teams can use the same verification record to see where a metric is independently supported, where it depends on reported inputs, where uncertainty remains and exactly which configuration and operating conditions the conclusion covers.

Greater credibility outside the company

Customers, insurers, regulators, partners and other authorized stakeholders can review a conclusion supported by evidence they did not generate themselves. They can distinguish independently observed behavior from participant-supplied state, see how discrepancies were handled and follow the evidence trail supporting the result.

Receive-only by design

The assurance architecture is designed to remain outside the machine control loop. Evidence collection, MVE™ and BoundaryScore™ do not command, steer, brake, accelerate or reconfigure the system being evaluated. That separation protects observation integrity and keeps verification distinct from operational control.

What the verification record makes clear

Verification record	What BoundaryScore™ Verified makes explicit
Defined metric	The exact measure or performance claim being verified, including the applicable definition and evaluation boundary.
Independent support	Which parts of the conclusion are supported by independently obtained physical evidence and which inputs remain participant- or context-supplied.
Evidence confidence	The strength, coverage, uncertainty and limitations of the evidence supporting the conclusion.
Applicability	The exact system, configuration, mission, conditions, exposure period and methodology version to which the conclusion applies.
Auditability	The reproducible path from the conclusion back to validated events, source evidence, transformations and calculation rules.

The core proposition

Bound the claim. Authenticate the evidence. Reconcile observed reality. Separate performance from confidence. Preserve the record.

The value of the BoundaryScore™ assurance architecture is not that an outside party substitutes its judgment for the company or stakeholder defining the metric. Its value is that the resulting conclusion can be inspected, challenged and reproduced from evidence whose boundaries, provenance, uncertainty and applicability remain explicit.