

Technical Confidence in Six Questions

How the BoundaryScore™ assurance framework applies proven principles of zero trust, independent evidence, traceability and controlled change to Physical AI verification.

1 What makes the BoundaryScore™ assurance framework technically trustworthy?

The BoundaryScore™ assurance framework is designed so that no single participant, machine-generated record, sensor or evidence source is inherently authoritative. It defines exactly what is being evaluated, authenticates and attributes the supporting evidence, preserves uncertainty and provenance, and maintains a reproducible path from a conclusion back to the underlying events and evidence. Every conclusion is bounded to the identified system, configuration, mission, operating conditions and exposure actually evaluated.

2 Does the BoundaryScore™ assurance framework simply trust what the vehicle or another system reports?

No. The BoundaryScore™ assurance framework applies a zero-trust evidence principle. Machine-generated data and independently obtained observations are both treated as evidence to be authenticated and evaluated, not as unquestioned truth. Identity, integrity, provenance, timing, calibration, coverage and independence determine how much confidence each source can support.

3 How does independent motion validation fit into the broader BoundaryScore™ assurance framework?

Independent motion validation is an important evidence layer, but it is not the entire architecture. The Motion Validation Envelope™ (MVE™) reconciles machine-state evidence with independently observed physical behavior, accounting for timing, spatial alignment, calibration, measurement precision, latency and uncertainty. MVE™ establishes the defensible event record; the applicable performance metrics are evaluated only after that evidentiary step.

4 What prevents weak evidence or conflicting data from producing a misleading result?

The BoundaryScore™ assurance framework deliberately separates evidence confidence from measured performance. Missing evidence, contradictions and uncertainty remain visible rather than being averaged away. Authenticated exposure establishes whether enough relevant observation exists, while critical events can be preserved through explicit gates. When the evidence cannot support a defensible conclusion, the framework can qualify or abstain rather than manufacture certainty.

5 Can companies, regulators or other stakeholders use the metrics that already matter to them?

Yes. The BoundaryScore™ assurance framework is designed to verify existing performance metrics. An AV company, customer, insurer, regulator or other authorized party can define the measures relevant to its decision. BoundaryScore™ Verified can then provide independent assurance around those metrics by applying the appropriate evidence requirements, denominators, thresholds, confidence requirements and critical gates while preserving the same underlying verification architecture.

6 How does the BoundaryScore™ assurance framework preserve trust as vehicles, software and operating conditions change?

A conclusion is never treated as a permanent attribute of a vehicle or product. The BoundaryScore™ assurance framework binds each conclusion to an identified system, configuration, mission, operating conditions, exposure period and methodology version. Material changes can trigger reevaluation, while provenance, versioning and the audit trail preserve how the conclusion was produced. The result is designed to remain inspectable, challengeable and reproducible rather than simply trusted.