

AISPG™ / PUBLIC WHITE PAPER 01

THE AI SECURITY & GOVERNANCE AUTHORITY

THE AI ASSURANCE IMPERATIVE.

A NORTH STAR PERSPECTIVE ON ACCOUNTABLE AI

Why enterprise AI governance must become
architecture.

FOUNDER-ORIGINATED PUBLIC PERSPECTIVE

ALI AQEL / AISPG™ / OCTOBER 2026

PUBLIC EDITION — CONCEPTUAL RESEARCH

FIELD NOTE / 01

The assurance gap is an architecture problem.

Enterprises can assemble excellent cloud security, identity, data protection, model oversight and assurance programs - yet still struggle to explain whether a material AI action was authorized, governed and supported by sufficient evidence at the moment of decision.

01 Fragmentation is structural

Signals and decisions span systems and teams. A collection of dashboards cannot, by itself, reconcile ownership and obligations.

02 Proof is contextual

A configuration result, inventory record or log is useful evidence only when linked to a scoped obligation, entity, time and test.

03 Assurance is a decision

Leadership needs a bounded conclusion with known exceptions, confidence limits, responsible approvers and a revalidation trigger.

THE AISPG VIEW

AISPG™ is a founder-originated framework that articulates a conceptual governing architecture across these responsibilities. It is not a government regulator, a credential or an automatic claim of technical implementation.

FIELD NOTE / 02

The enterprise reality: exceptional tools, incomplete answers.

The challenge is rarely a complete absence of controls. It is the missing relationship among what each system can know, what each team owns and what leadership is being asked to approve.

01 / IDENTITY

An AI agent can act through multiple service identities. Its effective permissions may differ from its stated purpose.

02 / MODEL

A model may change version or operating context after evaluation. The governance conclusion needs a precise scope.

03 / DATA

Approved data access for one workflow does not establish permission to reuse, summarize or export the same data.

04 / CLOUD

Cloud and AI-native telemetry may come from different providers, timestamps and control contexts.

THE REFRAME

Do not ask whether each tool has a dashboard. Ask whether the enterprise can reconstruct effective authority, scoped obligations and the evidence supporting a consequential decision.

FIELD NOTE / 03

Seven domains; one public governance vocabulary.

AISPG's public framework distinguishes governing responsibilities from technical products. These domains may all be implicated in a single enterprise AI use case.

01 Identity Governance

Human and non-human identities, access, delegation and accountable authority.

02 Model & Pipeline Governance

Provenance, approved purpose, version movement and evaluation context.

03 Agentic Governance

Tool permissions, runtime boundaries, delegated actions and oversight.

04 Data Governance

Classification, lineage, permitted use and output-handling obligations.

05 Autonomous Decision Governance

Decision rights, explainability, human oversight and escalation.

06 Risk, Assurance & Evidence

Evidence integrity, limitations, test outcomes and defensible posture.

07 AI Cloud Governance & Control Assurance

Cloud-hosted AI control applicability and cross-provider governance context.

FIELD NOTE / 04

From signals to a defensible decision.

The public AISPG Assurance Chain deliberately separates observation from assurance. It is a conceptual method, not a claim that every stage is technically validated in a live deployment.

01**SIGNAL**

What was observed, by which source and at what point in time?

02**CONTROL**

Which obligation, policy expectation or accountable boundary applies?

03**EVIDENCE**

What reliable proof supports or challenges control effectiveness?

04**ASSURANCE**

What conclusion follows, considering contradictory or missing evidence?

05**POSTURE**

What scoped decision can the owner defend, and when must it be revisited?

DESIGN PRINCIPLE

Missing evidence does not become positive proof. Failed tests, exceptions and uncertainty belong in the conclusion - not hidden behind a composite color or score.

FIELD NOTE / 05

A governed AI release: an illustrative decision.

Hypothetical example: an enterprise proposes a retrieval-augmented assistant for an internal regulated workflow. The questions below are illustrative, not the result of a production assessment.

1

01 / DEFINE SCOPE

Identify purpose, population, owner, model version, deployment environment and the actions the assistant may take.

2

02 / TEST BOUNDARIES

Assess identity, sensitive data, vendor services, logging and limits on downstream decisions.

3

03 / EVALUATE EVIDENCE

Relate source records and control tests to the approved use; document incomplete or conflicting signals.

4

04 / RECORD DECISION

Approve, hold or narrow deployment with exceptions, accountable approvers and an explicit review event.

ILLUSTRATIVE ONLY

This is a teaching scenario, not a customer claim, control test, attestation, benchmark or record of AISPG production activity.

FIELD NOTE / 06

Six questions for boards, CISOs and AI leaders.

Better governance conversations start with answerable questions that cross functional boundaries without replacing those functions.

01 What is the exact AI entity, version, workflow and material decision in scope?

02 Who possesses effective authority - human, service, agent or delegated tool?

03 Which obligations apply, and who has accepted accountability for them?

04 Which systems hold source evidence, and what can each genuinely prove?

05 Where are the unknowns, exceptions, failed tests and revalidation triggers?

06 What bounded decision can leadership make today, with what confidence?

START WITH ONE REAL JOURNEY

Start with one consequential AI journey. Map authority and source evidence, name the accountable approver and document the limits of what you can prove.

FIELD NOTE / 07

Context. Sources. Further inquiry.

AISPG's interpretation is independent. The following authoritative publications are external context, not endorsement, adoption, certification or validation of AISPG.

REFERENCE / 01

NIST AI Risk Management Framework 1.0

NIST, 2023. Voluntary risk-governance framework with Govern, Map, Measure and Manage functions.

<https://airc.nist.gov/airmf-resources/airmf/>

REFERENCE / 02

NIST Generative AI Profile

NIST AI 600-1, 2024. A cross-sectoral profile relating generative AI risk to the AI RMF.

<https://doi.org/10.6028/NIST.AI.600-1>

REFERENCE / 03

ISO/IEC 42001:2023

International standard specifying requirements for an AI management system.

<https://www.iso.org/standard/42001>

REFERENCE / 04

OWASP GenAI Security Project

Official open guidance, including the September 2026 LLM risks and agent-control resources.

<https://genai.owasp.org/>

PUBLIC EDITION / DISCLOSURE BOUNDARY

© 2026 Ali Aqel / AISPG™. This original paper is provided for education and discussion. AISPG is a privately owned, founder-originated framework, not a regulator or accredited certifier. Conceptual diagrams and questions do not represent live control validation, licensed operating procedures, legal advice or standards conformity. Proprietary control logic, undisclosed reasoning mechanisms, private architecture and sensitive platform verification remain outside this edition.

EXPLORE THE PUBLIC AUTHORITY EXPERIENCE

AISPG.ai / Research & qualified access