

THE AI SECURITY & GOVERNANCE AUTHORITY

CONNECT, DON'T REPLACE.

GOVERNING AI ACROSS THE ENTERPRISE SECURITY STACK

Why AI governance should connect your stack, not replace it.

FOUNDER-ORIGINATED PUBLIC PERSPECTIVE

ALI AQEL / AISPG™ / OCTOBER 2026

PUBLIC EDITION — CONCEPTUAL RESEARCH

FIELD NOTE / 01

AI cuts across every security domain at once.

Large enterprises do not start AI governance with a blank estate. They already run identity, cloud, data, detection, response and governance platforms. One AI agent can touch all of them in a single task.

01 Every tool sees part of the agent

Identity sees its credentials, cloud sees its runtime, detection sees its events and governance sees its policy. None sees the whole governed action.

02 No tool owns the AI question

Is this AI use authorized, within its boundaries and supported by evidence? That answer spans systems that were never designed to share it.

03 Replacement is the wrong answer

A new platform that duplicates existing controls adds cost and conflict. The gap is shared meaning, not another enforcement point.

THE AISPG VIEW

AISP[™]G is designed as an AI-centric governance umbrella, not a replacement stack. It defines how existing technologies take part in governed AI outcomes and how their contributions combine into one defensible posture.

FIELD NOTE / 02

Excellent systems, fragmented answers.

Each platform is authoritative within its own domain. The fragmentation appears when a single AI decision depends on several domains at once.

01 / IDENTITY

Knows which privileges a service identity holds, but not which model or business decision depends on that identity.

02 / CLOUD

Knows configuration and exposure, but not whether the affected workload makes a consequential automated decision.

03 / GOVERNANCE

Knows the policy and the risk acceptance, but often not the current runtime state of the AI system it applies to.

04 / DETECTION

Sees events and anomalies, but not which governance obligation an event affects or who must approve a response.

THE REFRAME

Do not ask which tool should own AI governance. Ask which system is authoritative for each governance claim, and how those claims are joined into one accountable conclusion.

FIELD NOTE / 03

What stays authoritative.

The umbrella principle is simple: govern and assure the controls the enterprise already owns first. Add new mechanisms only where a required capability is missing.

01 Identity & Zero Trust

Authentication, privilege and access enforcement for people and machines.

02 Cloud security posture

Cloud configuration, workload findings and provider-level controls.

03 Detection & response

Telemetry, correlation, investigation and incident context.

04 Automation & orchestration

Executing remediation, containment and approval workflows.

05 Data security

Classification, exposure and protection of the data AI systems use.

06 Governance, risk & compliance

Policies, obligations, assessments, exceptions and risk acceptance.

07 Model & runtime security

Specialized protection of models, prompts and inference behavior.

FIELD NOTE / 04

Four roles, one governed conclusion.

Every system that takes part in AI governance plays one or more of four roles. Making the role explicit shows why each integration exists and what it is trusted to prove.

01**SIGNAL PROVIDER**

Supplies events, configuration, inventory, findings or state.

02**ENFORCEMENT POINT**

Allows, denies, constrains or requires approval for an action.

03**EVIDENCE AUTHORITY**

Provides attributable proof that a control claim holds.

04**RESPONSE & WORKFLOW**

Carries out remediation, ticketing, containment or approval.

05**THE UMBRELLA**

Connects these contributions into one governed AI state that leadership can defend.

DESIGN PRINCIPLE

Prefer the authoritative source for each claim and keep its provenance, rather than copying every record into a parallel control database.

FIELD NOTE / 05

One alert, many systems: an illustrative decision.

Hypothetical example: detection tooling flags unusual data access by an AI agent. The steps below are illustrative, not the result of a production assessment.

1

01 / IDENTIFY THE AI ENTITY

Link the alert to the specific agent, its version, its owner and the business process it serves.

2

02 / TRACE ITS AUTHORITY

Use the identity platform as the authority for which permissions the agent actually holds.

3

03 / READ THE EVIDENCE

Take data classification, cloud state and policy context from the systems that own each answer.

4

04 / ROUTE THE DECISION

Hand containment or approval to existing response workflows, with the governance context attached.

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 CISOs and architecture boards.

These questions test whether the current security stack can answer AI governance questions together, without adding a parallel stack.

- 01 For our most important AI agent, which system is authoritative for its identity, data and runtime state?
- 02 Can we join those answers into one conclusion without manual reconciliation?
- 03 Which AI governance capabilities are genuinely missing, rather than present but unconnected?
- 04 Where are we copying evidence instead of referencing the authoritative source?
- 05 When detection flags an AI agent, who decides the response, and with what context?
- 06 Which planned AI-security purchase duplicates a control we already own?

START WITH ONE REAL AGENT

Map one production AI agent across your existing platforms. Note which system answers each governance question, and where no system can answer it today.

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 Cybersecurity Framework 2.0

NIST, 2024. Outcome-based framework including the Govern function for cybersecurity risk.

<https://doi.org/10.6028/NIST.CSWP.29>

REFERENCE / 02

NIST SP 800-207: Zero Trust Architecture

NIST, 2020. Principles and logical components of zero trust architectures.

<https://doi.org/10.6028/NIST.SP.800-207>

REFERENCE / 03

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 / 04

ISO/IEC 42001:2023

International standard specifying requirements for an AI management system.

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

PUBLIC EDITION / DISCLOSURE BOUNDARY

© 2026 Ali Aqel / AISPG™. This original paper is provided for education and discussion. AISPG™ is independently owned by Ali Aqel and licensed to Blue Sentinel AI, Inc. 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