

THE AI SECURITY & GOVERNANCE AUTHORITY

GOVERNING THE AI CLOUD.

AI CLOUD GOVERNANCE & CONTROL ASSURANCE

Why provider policy is not yet enterprise AI governance.

FOUNDER-ORIGINATED PUBLIC PERSPECTIVE

ALI AQEL / AISPG™ / OCTOBER 2026

PUBLIC EDITION — CONCEPTUAL RESEARCH

FIELD NOTE / 01

Cloud policy enforces. It does not, by itself, govern.

Enterprise AI runs on cloud control planes, managed AI services, identities, networks and data platforms. Providers offer powerful policy engines, but their answers are expressed in provider terms, not in the terms leadership must defend.

01 The cloud is where AI runs

Models, agents, endpoints and retrieval data live in accounts, subscriptions and projects that change continuously through portals, pipelines and automation.

02 Policies speak provider language

Each provider expresses scope, inheritance, exceptions and effect differently. Those differences are real and should not be flattened away.

03 Leadership asks a different question

Not "is the policy assigned?" but "is this AI use governed as intended, everywhere it runs, and how confident can we be?"

THE AISPG VIEW

AISPG[™] treats AI Cloud Governance & Control Assurance as a first-class governance domain: provider mechanisms keep their native meaning, while the enterprise objective, coverage and assurance are expressed in one common governance language.

FIELD NOTE / 02

Four ways a compliant-looking cloud misleads.

None of these conditions is a provider defect. Each is a gap between what a policy object says and what the enterprise assumes it means for AI.

01 / ASSIGNMENT

A policy can be assigned in audit-only mode. An assigned policy is not, by itself, evidence that the control is effective.

02 / INHERITANCE

A rule applied at one level of the resource hierarchy cannot be assumed to produce the same result everywhere below it.

03 / PERMISSIONS

Organization guardrails can limit what identities may do; they do not grant access, and they interact with other authorization layers.

04 / NEW RESOURCES

A new account, region or AI endpoint can appear outside the boundary that existing policies were designed to cover.

THE REFRAME

Do not count assigned policies. Ask which AI governance objective each one serves, which AI workloads it actually reaches and what shows that it works.

FIELD NOTE / 03

Coverage has five dimensions.

A strong control with partial coverage cannot support a strong enterprise-wide conclusion. Coverage should be described, not assumed.

01 Entity

Which AI workloads, agents, models, platforms and non-human identities are actually governed?

02 Environment

Which accounts, subscriptions, projects, regions and runtime boundaries are in scope?

03 Objective

Which enterprise governance objectives have a provider mechanism mapped to them?

04 Time

How current is the observed state, and how quickly would drift become visible?

05 Assurance

Which mapped controls have been validated, rather than only observed?

DESIGN PRINCIPLE

An account no tool can see is not a compliant account. The absence of findings is not evidence of coverage.

FIELD NOTE / 04

One objective, many providers.

The AISPG Assurance Chain applies to cloud the same way it applies elsewhere: observation is separated from assurance, and provider detail is preserved rather than averaged.

01**SIGNAL**

Which provider state, policy evaluation or runtime event was observed, and when?

02**CONTROL**

Which provider-neutral AI governance objective does this mechanism serve?

03**EVIDENCE**

Does the evidence reflect effective behavior, after inheritance and exceptions are applied?

04**ASSURANCE**

Has the control been validated for the AI workloads that depend on it?

05**POSTURE**

What does this mean for each AI use case, and which change would reopen the conclusion?

DESIGN PRINCIPLE

Normalize the objective, never the mechanism. If one provider can enforce a stronger boundary or produce better evidence, that difference should stay visible.

FIELD NOTE / 05

One model, two clouds: an illustrative decision.

Hypothetical example: an approved model is deployed to two cloud providers. The steps below are illustrative, not the result of a production assessment.

1

01 / DEFINE THE OBJECTIVE

Only approved, signed model versions may run, and public invocation is not permitted.

2

02 / MAP EACH PROVIDER

Record each provider's registry, deployment and network mechanisms in their own terms.

3

03 / COMPARE THE EVIDENCE

Note where one environment can prove the running version and the other can only assert it.

4

04 / REOPEN ON CHANGE

A new region, endpoint or exception reopens the conclusion until it is re-examined.

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 cloud, security and AI leaders.

These questions connect cloud engineering detail to the AI governance decisions leadership is asked to approve.

- 01 Which cloud accounts, regions and services host our AI workloads today, and how would we know about a new one?
- 02 For each critical AI objective, which provider mechanism enforces it in each cloud?
- 03 Where is a policy assigned but never validated against effective behavior?
- 04 Which exceptions are open, who owns them, and when do they expire?
- 05 How long does it take for a material cloud change to change our AI governance conclusion?
- 06 Where is one provider's evidence weaker than another's, and do our reports show it?

START WITH ONE REAL WORKLOAD

Pick one production AI workload. Write its governance objective, map the mechanism in each environment it runs in and record what you can prove and what you can only assume.

FIELD NOTE / 07

Context. Sources. Further inquiry.

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

REFERENCE / 01

Microsoft Azure Policy overview

Microsoft Learn. Azure Policy definitions, assignments, scopes and effects.

<https://learn.microsoft.com/en-us/azure/governance/policy/overview>

REFERENCE / 02

AWS Organizations: service control policies

Amazon Web Services documentation. Guardrails that limit, but do not grant, permissions.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

REFERENCE / 03

Google Cloud Organization Policy overview

Google Cloud documentation. Constraints, inheritance and evaluation across the resource hierarchy.

<https://docs.cloud.google.com/organization-policy/overview>

REFERENCE / 04

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/>

PUBLIC EDITION / DISCLOSURE BOUNDARY

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

AISP^g.ai / Research & qualified access