

AI-Driven Identity Governance Framework (AIGF)

A Framework for Intelligent Access Management

Presented By:
Neha Asthana
Identity Solutions Architect



Why Identity Governance Needs More Intelligence

The enterprise identity landscape is increasingly distributed and dynamic.

Identity sprawl

Employees, contractors, privileged users, service accounts, applications, workloads, and machine identities.

Changing environments

Cloud, SaaS, APIs, remote work, and distributed applications create more access relationships.

Static governance limits

Periodic reviews, fixed rules, and manual decisions may miss changing context and risk.

Core idea: move from static, periodic access governance toward continuous, contextual, risk-aware decisions.

IGA Remains the Control Foundation

AI augments established governance; it does not replace it.

Core IGA capabilities

- Identity lifecycle management
- Access requests and provisioning
- Role and entitlement management
- Segregation of duties (SoD)
- Access certification and compliance


Where traditional approaches struggle

- Static policies may not reflect changing risk
- Users accumulate access after role changes
- Certification campaigns overwhelm reviewers
- Limited usage and peer-group context
- Manual role engineering is costly

The goal is not just to automate workflows—it is to introduce intelligence into the governance lifecycle.

Four Capabilities: Understand → Analyze → Decide → Act

An operating model for continuous, contextual identity governance.



Feedback from outcomes and human decisions helps refine future recommendations—within defined controls.

Context Makes Access Risk More Meaningful

An example to illustrate the difference.

Raw entitlement view

“User A has access to Application X.”

This alone does not explain whether access is appropriate.



Context-enriched view

Finance employee; recently moved teams; no usage for 180 days; 94% of comparable users lack the privileged entitlement.

Context can prioritize review—not automatically prove that access is improper.

AI Analytics: Match the Technique to the Problem

The proposed framework uses multiple methods rather than one universal model.

Anomaly detection

- Unusual application access
- Unexpected privileges
- Abnormal usage frequency
- Sudden privilege escalation

Clustering & role mining

- Group users by access patterns
- Discover candidate business roles
- Identify peer-group differences


Recommendations

- Suggest access used by similar users
- Prioritize certification decisions
- Explain why access is recommended or flagged

Design principle: recommendations should be explainable and assessed against policy before enforcement.

Governance Decision Engine: Keep a Clear Security Boundary

AI can highlight risk; policy and authorized decision-makers govern the response.



Human-in-the-loop: AI recommendations should not automatically become authorization decisions—especially for privileged access.

Automation & Feedback: Close the Governance Loop

Turn approved decisions into controlled actions and learn from outcomes.

1. Execute

Trigger approved provisioning, deprovisioning, certification, or remediation workflows.

2. Verify

Confirm completion, capture errors, track exceptions, and retain audit evidence.

3. Learn

Use outcomes and human decisions to improve future recommendations and prioritization.

Guardrails

- Approval and policy checks remain mandatory where required
- High-risk actions can require explicit human authorization
- Failed or uncertain automation should fall back to safe workflows

Security, Resilience & Fail-Safe Design

AI becomes part of the security architecture and must be governed accordingly.

Data security

Encryption • least privilege • access controls • data minimization • appropriate retention

Model security

Protect against unauthorized modification, data poisoning, model extraction, and malicious inputs.

Fail-safe behavior

If AI is unavailable, traditional IGA policies continue to operate. AI augments—not replaces—the control plane.

Principle: never make the AI model the sole enforcement mechanism for identity access.

Example for Data Security: If an AI model recommends granting an employee administrative privileges, the system must validate the recommendation against access policies and approval requirements before granting access.

Example for Resilience: If the AI recommendation engine becomes unavailable, the system can continue processing access requests using predefined policies and approval workflows instead of stopping all provisioning activities.

Practical Implementation Roadmap

A phased path helps validate data, governance, and operational readiness.

- 01 Establish foundations** Inventory identity sources; assess data quality, entitlements, ownership, and current IGA controls.
- 02 Start with decision support** Pilot anomaly detection and certification prioritization with explainable recommendations.
- 03 Integrate governance** Connect AI signals to policy checks, SoD, approvals, audit logging, and human review.
- 04 Automate selectively** Enable low-risk, approved remediation; retain stronger approval gates for privileged/high-impact actions.
- 05 Measure & improve** Track precision, false positives, review effort, remediation time, overrides, and control outcomes.

How to Evaluate AIGF

Measure security effectiveness and operational usefulness—not just automation volume.

Risk detection

Detection precision/recall
False-positive rate
High-risk access identified

Governance efficiency

Certification review time
Time to remediate
Manual effort and backlog

Control quality

Policy/SoD violations
Human override rate
Audit completeness and traceability

Compare results with a baseline and validate across application criticality, identity types, and privileged-access scenarios.

Limitations & Future Research

AIGF is a proposed framework; enterprise effectiveness requires empirical validation.

Key limitations

- Inaccurate or incomplete identity data
- Model bias and false positives
- Limited explainability
- Integration complexity
- Privacy and regulatory constraints

Future research

- Validate risk-scoring approaches
- Compare model techniques and thresholds
- Measure governance outcomes over time
- Improve explainability and auditability
- Test safe automation in enterprise settings

Avoid treating a model score as proof: use it as a signal within accountable governance.

Conclusion

AI should make identity governance more contextual, adaptive, and risk-aware—while preserving policy enforcement and human accountability.

UNDERSTAND → ANALYZE → DECIDE → ACT

AIGF combines identity data, contextual intelligence, AI analytics, governance controls, controlled automation, and continuous feedback.

Thank You