[Service Brief]

Agentic AI + Identity Risk Assessment

Secure the Next Generation of Non-Human and Autonomous Identities

OVERVIEW

Agentic AI introduces a fundamentally new identity risk challenge.

Unlike traditional applications or service accounts, AI agents reason, make decisions, interact with multiple systems, and act with little or no human involvement. As organizations adopt AI into workflows, customer interactions, development pipelines, and operations, they create a new class of non-human identity (NHI) with expanded access and elevated risk.

SDG's Agentic AI + Identity Risk Assessment is a structured assessment service that helps organizations identify, measure, and reduce the identity risks introduced by AI agents, NHIs, and autonomous workflows.

The assessment gives you a clearer picture of:

- **What AI agents and NHIs exist**
- **How those identities authenticate**
- **What systems, data, and actions they can access**
- **Whether that access is appropriate and sufficiently controlled**
- **How quickly those identities can be monitored, contained, or disabled if something goes wrong**

The result is a structured Identity Risk Assessment that identifies where organizations are vulnerable, what should be prioritized, and what controls are required to securely scale AI.

WHY THIS MATTERS NOW

Agentic AI is already embedded in enterprise environments through:

- Code generation and software delivery pipelines
- IT ticketing and automation workflows
- Data pipelines and retrieval-augmented generation (RAG)
- Customer service and digital interaction platforms
- Security and operations automation

Most IAM programs were designed for people and not autonomous, persistent, machine-speed identities.

Without stronger controls, the exposure quietly compounds; organizations face the challenge of unmanaged agents, exposed credentials, excessive permissions, and no clear way to contain something that goes wrong.

SDG SOLUTION

SDG's Agentic AI + Identity Risk Assessment is an advisory-led assessment that identifies the identity-related risks associated with AI agents and NHIs.

You will receive:

- A current-state assessment of AI and NHI risk
- An executive-friendly heat map showing the most significant risks
- A maturity score across critical control areas
- A prioritized remediation roadmap
- A future-state Identity Fabric architecture for securing AI-enabled environments

WHAT YOU'LL EXPERIENCE DURING THE ASSESSMENT

During the **Agentic AI + Identity Risk Assessment**, SDG works with your security, IAM, cloud, application, and infrastructure teams to:

- Identify every AI agent, NHI, service account, API credential, workload identity, and machine identity in scope
- Document what each identity can access and whether that access is appropriate
- Identify high-risk patterns such as shared accounts, excessive permissions, static credentials, missing ownership, and absent kill-switch controls
- Evaluate whether existing IAM, PAM, cloud, monitoring, and security tools are sufficient to govern AI agents
- Determine where additional controls, governance, and automation are required

At the end of the assessment, you will come away with a clear picture of:

- Where the highest identity risks exist
- Which issues should be addressed first
- What changes are needed to securely support AI adoption

SCOPE OF THE IDENTITY RISK ASSESSMENT

The assessment spans the full identity and infrastructure landscape where AI agents and NHIs operate.

Identity & Access Domains

- Identity Governance & Administration (IGA)
- Privileged Access Management (PAM)
- Access Management (SSO, MFA, Federation)
- NHI Management
- Secrets and Credential Management

Infrastructure & Runtime Domains

- Cloud IAM (AWS, Azure, GCP)
- Kubernetes, containers, and CI/CD pipelines
- APIs and service integrations
- Application and workload identities
- Certificates, keys, and trust material
- Monitoring, logging, and containment controls

ASSESSMENT DOMAINS

The Agentic AI + Identity Risk Assessment evaluates 10 core domains:

1. **Scope & Coverage**

   Confirm that all relevant IAM, cloud, infrastructure, and runtime environments are included

2. **Agent Inventory & Architecture**

   Catalog all AI agents, service accounts, APIs, machine identities, and execution environments

3. **Agent Identity Model**

   Define ownership, unique identity assignment, and accountability for each AI or NHI

4. **Authentication & Secrets**

   Assess how credentials are issued, stored, rotated, and protected

5. **Authorization & Guardrails**

   Evaluate least privilege, just-in-time access, segregation of duties, and human approval checkpoints

6. **Monitoring & Attribution**

   Ensure all agent actions can be traced back to a specific identity and activity

7. **Containment & Kill Switch**

   Validate the ability to rapidly disable, isolate, or limit an AI agent if it behaves unexpectedly

8. **Hybrid & Multi-Cloud Coverage**

   Ensure controls remain consistent across on-premises, cloud, and multi-cloud environments

9. **Vendor Capability Validation**

   Compare existing vendor capabilities to actual control requirements and identify gaps

10. **Risk Ratings & Roadmap**

    Prioritize findings based on business impact, likelihood, and remediation effort

EXAMPLE HIGH-RISK FINDINGS

Common high-risk patterns identified during the assessment include:

- Shared service accounts used by multiple AI agents
- Long-lived API keys and OAuth tokens
- Excessive permissions across multiple systems
- Static credentials stored in CI/CD pipelines or configuration files
- AI agents with no defined owner or accountability
- Missing segregation of duties between data access and system actions
- No kill-switch or containment procedure
- Limited monitoring of agent actions and decision-making

WHAT YOU RECEIVE

Current-State Identity Risk Assessment

A documented assessment of the client's current controls, risks, and gaps across each assessment domain.

Executive Risk Heat Map

A simple, executive-friendly heat map showing where the most significant identity risks exist, how serious they are, and which issues should be addressed first.

The heat map highlights issues such as:

- Shared or unmanaged accounts
- Excessive access
- Exposed credentials
- Missing ownership
- Lack of containment or monitoring

Maturity Scorecard

A maturity rating across key areas such as identity governance, secrets management, cloud IAM, monitoring, and containment.

Prioritized Remediation Roadmap

A practical roadmap that includes:

- 30-day quick wins
- Medium-term improvements
- Long-term transformation initiatives

Future-State Identity Fabric Architecture

A target-state architecture showing how to secure:

- Human identities
- NHIs
- AI agents and autonomous workflows
- Monitoring and response capabilities

SDG ACCELERATORS & DIFFERENTIATORS

SDG brings proven frameworks, templates, and assessment tools that accelerate delivery and make the results more actionable.

Identity Risk Assessment Framework

A pre-built assessment model covering the critical domains of AI and NHI risk, including identity governance, privileged access, secrets, cloud IAM, monitoring, and containment. Why it matters:

- Reduces time spent deciding what to assess
- Ensures consistent, repeatable results
- Prevents important risks from being overlooked

Agent & NHI Taxonomy

A structured classification model for AI agents, service accounts, APIs and tokens, workload identities, secrets and certificates, and privileged machine identities. Why it matters:

- Quickly identifies what types of identities exist
- Prioritizes which identities represent the greatest risk
- Creates a foundation for governance and ownership

Risk Heat Map & Maturity Scoring Model

A scoring approach that rates each issue as Critical, High, Medium, or Low and measures the maturity of the client's current controls. Why it matters:

- Provides executives a simple way to understand significant issues
- Helps security teams prioritize remediation
- Makes it easier to justify follow-on projects

Identity Fabric Reference Architecture

A future-state architecture showing how to secure AI agents and NHIs across IGA, PAM, Cloud IAM, secrets management, PKI, and monitoring and containment. Why it matters:

- Provides the client a practical target state
- Creates a roadmap for future transformation
- Differentiates SDG from firms that stop at findings and recommendations

Remediation Roadmap Templates

Pre-built remediation plans showing quick wins, medium-term improvements, and long-term strategic initiatives. Why it matters:

- Accelerates delivery of actionable next steps
- Helps clients move from assessment to execution at speed
- Creates a bridge to SDG advisory, implementation, and managed services

DELIVERY OPTIONS

Tier 1: Quick Assessment (1–2 Weeks)

- High-level agent and NHI inventory
- Executive risk heat map
- Preliminary findings and recommendations

Tier 2: Standard Assessment (3–5 Weeks)

- Full Identity Risk Assessment across all domains
- Detailed control evaluation
- Maturity scoring and remediation roadmap

Tier 3: Deep Dive Assessment (6–10+ Weeks)

- Threat modeling and control validation
- Detailed containment and kill-switch testing
- Implementation backlog and future-state architecture

WHAT SDG DELIVERS | AI VISIBILITY, GOVERNANCE, AND SECURITY

SDG helps clients see AI, govern AI, and secure AI — with identity at the core.

01 VISIBILITY

Know Every AI Agent and its Access

Inventory sanctioned and shadow AI across agents, copilots, models, identities, data, and connected systems.

Client Value: A prioritized baseline for ownership, risk reduction, and faster decisions.

02 GOVERNANCE

Assign Ownership and Enforce Control

Define policy, decision rights, lifecycle governance, and least-privilege access aligned to NIST AI RMF, EU AI Act, and ISO 42001.

Client Value: An audit-ready AI operating model leadership can scale with confidence.

03 SECURITY

Protect AI from Misuse and Attack

Assess, adversarial test, and continuously monitor AI across prompts, agents, models, tools, and human interaction points.

Client Value: Reduced exposure to prompt abuse, deepfakes, agent misuse, and control drift.

WHY SDG

SDG has spent more than 30 years working inside enterprise identity environments across IGA, PAM, access management, NHI, and cloud security. That experience informs how we assess, what we prioritize, and what we consider a real risk versus noise.

Our Identity Fabric architecture methods come from work done in complex, real-world environments. And because SDG delivers advisory, transformation, and managed services, clients stay with the same team from assessment through execution.

SDG

75 North Water Street
Norwalk, CT 06854
203.866.8886
sdgc.com

Contact Us: solutions@sdgc.com