FACT SHEET

HaystackID® AI Governance Services

Operational AI Governance for Defensible Oversight, Validation, and Compliance at Scale

HAYSTACK®

Service Overview

Strategic guidance and operational execution that embeds responsible AI controls into day-to-day delivery, procurement, and enterprise risk management.

HaystackID® AI Governance Services help organizations **operationalize AI governance as an execution discipline**—moving beyond policy statements to a repeatable operating model integrated into day-to-day delivery. The services are designed for product, engineering, operations, and revenue leaders who need to scale AI into production and commercial offerings while meeting regulator, customer, and stakeholder expectations for transparency, oversight, and accountability.

HaystackID bridges operational and risk stakeholders by translating governance intent into implementable workflows, decision rights, and validation routines. The resulting evidence and reporting reduce deployment friction, shorten customer and partner due diligence cycles, and provide GRC, privacy, and legal teams with **defensible documentation and controls for high-impact AI systems** as requirements evolve.

From Compliance Cost to Business Enablement

AI governance is frequently treated as a necessary expense—reactive documentation assembled for audits, customer questionnaires, or regulator inquiries. Operational AI governance changes the equation by creating an "**evidence factory**": a disciplined set of workflows, controls, and artifacts that can be reused across AI systems, business units, and customer engagements.

When executed as an operating capability, AI governance can **shift from cost center to revenue enabler** by reducing friction in customer and partner assurance, accelerating responsible AI adoption, and enabling market access in regulated jurisdictions and high-stakes use cases.

Common commercial and operational benefits targeted by AI Governance engagements include:

- **Faster approvals** and lower rework by embedding risk classification, documentation, and testing into delivery workflows.
- **Improved sales enablement** through reusable assurance packages (system cards, test evidence, control mappings) for enterprise customers.
- **Reduced procurement friction** by standardizing third-party AI due diligence, contractual controls, and vendor evidence expectations.
- **More predictable cost** per AI system by standardizing controls and evidence capture across the portfolio.
- **Better executive visibility** through governance KPIs that connect AI risk posture to operational performance and business outcomes.

AI Governance Portfolio Overview

AI governance requirements often emerge as AI adoption expands beyond a single team or use case. AI may be developed internally, introduced through enterprise platforms, or embedded in third-party tools adopted by business units. As usage grows, organizations frequently encounter the same challenge: **AI is being used in high-impact contexts** before consistent classification, documentation, approval pathways, testing routines, and oversight reporting are in place. This includes shadow AI—employees using public AI tools or personal accounts to complete work outside approved environments—which can bypass controls and create untracked data, IP, and compliance risk.

HaystackID AI Governance Services are designed to establish visibility, implement operational controls, validate high-risk systems, and produce evidence that supports defensibility. The portfolio is organized around four connected workstreams that **translate governance intent into measurable day-to-day execution:**

Visibility and Risk Classification

Inventory AI systems and use cases, classify risk, and map regulatory exposure so stakeholders can prioritize high-risk systems.

Governance Program Implementation

Define decision rights, standardize workflows, establish documentation and evidence expectations, and enable training and adoption.

System-level Validation

Assess security, robustness, privacy, and fairness characteristics prior to deployment or significant change; translate results into remediation and defensible artifacts.

Sustained Oversight

Operational monitoring and executive/board reporting that track governance maturity, respond to regulatory change, and maintain accountability over time.

HaystackID AI Governance Services are part of HaystackID's expanding offerings that leverage AI to help enable the management and governance of AI. The portfolio supports responsible innovation while delivering the controls and evidence needed to deploy AI with confidence.

High-level outcomes commonly targeted by AI Governance engagements include:

- **Establishing** a clear inventory and classification of AI systems and risk exposure.
- **Implementing** governance structures, policies, and controls aligned with leading frameworks.
- **Validating** security, robustness, privacy, and fairness characteristics of high-risk AI systems.
- **Producing** documentation suitable for audit, regulator engagement, customer assurance, and dispute contexts.
- **Sustaining** governance maturity through ongoing monitoring and executive/board-level oversight.

Why AI Governance Matters

AI governance matters because AI systems increasingly influence decisions that impact customers, employees, patients, citizens, and markets. These systems can introduce significant operational, legal, security, privacy, and reputational risk, and an organization's ability to control that risk varies by implementation approach. When an organization builds and trains its own models, it can govern the full lifecycle—data sourcing, training, testing, fine-tuning, and deployment controls. When an organization implements a third-party model or consumes embedded AI capabilities (for example via APIs or enterprise platforms), it inherits base-layer behaviors and vulnerabilities and often has limited ability to remediate them directly. In those cases, governance must emphasize **risk classification, secure integration, vendor due diligence and contracting, system-level validation in the implemented context, and continuous monitoring**.

For many organizations, the most immediate gap is not intent, but execution: inventories, approvals, testing protocols, evidence capture, and incident response procedures for AI - supported by an operating model that remains effective through organizational and technology change.

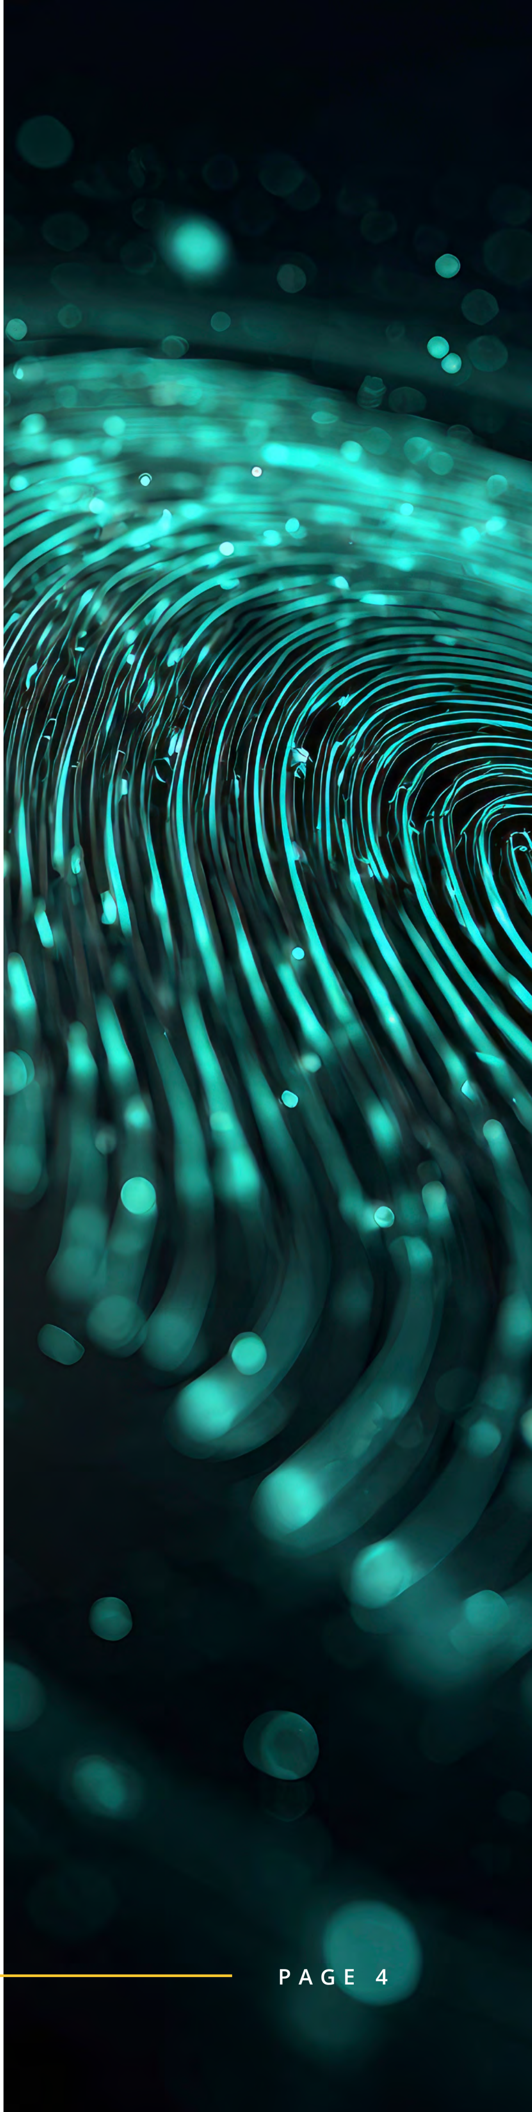

Key drivers shaping the need for AI Governance Services include:

- **EU AI Act Phased Obligations and Penalties:** Prohibited AI practices and AI literacy obligations apply from 2 February 2025; additional obligations phase in through 2 August 2027; administrative fines can reach €35M or 7% of global annual turnover (whichever is higher).
- **U.S. Policy Activity and State Laws:** Overlapping requirements for high-risk or consequential AI use cases (e.g., Colorado's SB 24-205 scheduled to take effect June 30, 2026, as amended).
- **Board Scrutiny:** Rising expectations for AI oversight, accountability, and evidence-based reporting as part of enterprise risk management.
- **Litigation Exposure:** discrimination, consumer harm, privacy, and product liability matters involving AI-driven decisions.
- **Shadow AI and Unsanctioned Usage:** Employees using consumer LLMs or personal accounts, increasing data leakage, IP exposure, and inconsistent controls.
- **Third-party AI Risk:** Vendor governance, contractual controls, and assurance expectations from customers and regulators.
- **Reputational Impact:** public trust and stakeholder confidence are increasingly linked to demonstrated responsible AI practices.

HaystackID AI Governance Primary Offerings

HaystackID delivers six integrated offerings that can be deployed independently or combined into a phased program. Together, the offerings help organizations understand AI risk exposure, build and operationalize governance programs, validate high-risk systems for security and fairness, provide executive and board-level oversight support, and produce independent audit-ready documentation.

AI Governance Primary Offerings: 12-Month Engagement View

Offering	1-3 Months	4-6 Months	7-9 Months	10-12 Months
AI Governance Scoping (3-8 weeks)	■			
AI Governance Advisory (6-12 months)	■	■	■	■
AI Security Testing (4-8 weeks)	■			
AI Fairness Testing (4-8 weeks)	■			
Board Advisory Services (Quarterly)	As Required			
Third-Party AI Compliance Audit	As Required			

AI Governance Scoping (Rapid Assessment | 3-8 weeks)

AI Governance Scoping is a structured assessment designed to establish visibility into the AI landscape, identify regulatory exposure, and produce a **prioritized roadmap for operational implementation**. The engagement maps AI systems, classifies risk, identifies control gaps, and defines practical next steps aligned to business and regulatory timelines.

Deliverables may include:

- Comprehensive AI system inventory and risk classification.
- Regulatory exposure and gap analysis across relevant requirements.
- Governance maturity assessment aligned to leading frameworks (NIST AI RMF, ISO/IEC 42001).
- Executive-ready roadmap with sequencing, priorities, and control gaps.
- Near-term action plan for high-risk systems and urgent obligations.

AI Governance Advisory (Program Implementation and Operations | 6-12 months)

AI Governance Advisory provides hands-on implementation support to build a sustainable AI governance program and embed it into operational delivery. The engagement focuses on establishing repeatable workflows, controls, training, and evidence capture that align with enterprise risk management and technology delivery practices. The objective is **a governance capability that can scale across systems and business units** without disrupting innovation.

Deliverables may include:

- Governance operating model (roles, responsibilities, decision rights, escalation paths).
- Policy and procedure development (approval pathways, documentation standards, change controls).
- AI risk management program design and implementation aligned to SDLC and MLOps and third-party intake.
- Control mappings and compliance documentation aligned to applicable regulations and frameworks.
- Training, communications, and change management materials to drive adoption.
- Metrics, dashboards, and reporting structures aligned to executive and board needs.

AI Security Testing (4-8 weeks per system)

AI Security Testing evaluates AI/ML systems for safety and security issues prior to deployment or significant change. Testing addresses modern AI threats, including prompt-based attacks, model extraction attempts, data leakage risks, and API/integration security issues. Results are translated into **prioritized remediation guidance and defensible documentation** of validation activities.

Deliverables may include:

- Adversarial robustness testing (prompt injection, jailbreak attempts, abuse case validation).
- Model extraction, inversion, and data leakage assessments.
- API security and integration vulnerability testing.
- Findings report with severity ratings and remediation recommendations.
- Re-test support and verification of remediation actions (as scoped).

AI Fairness Testing (4-8 weeks per system)

AI Fairness Testing assesses AI systems for bias and discrimination across protected characteristics and other relevant cohorts. This offering supports organizations deploying AI in **high-impact contexts** such as employment, credit, insurance, healthcare, and public-sector decision-making—where auditability and defensible testing are increasingly expected.

Deliverables may include:

- Disparate impact analysis across protected classes and relevant groups.
- Fairness metric evaluation (including demographic parity and equalized odds, as appropriate for the use case).
- Intersectional bias analysis to identify compounding effects.
- Mitigation recommendations (data, model, workflow, and policy-level options).
- Compliance documentation suitable for audit, customer assurance, and dispute contexts.

Board Advisory Services (Ongoing Oversight | Quarterly)

Board Advisory Services provide structured, recurring support for executive and board-level AI oversight. The engagement translates technical and regulatory complexity into clear reporting, risk monitoring, and maturity progression - supporting accountability and informed decision-making.

Deliverables may include:

- Quarterly board presentations and AI governance dashboards.
- Regulatory change monitoring and impact analysis.
- Emerging risk identification and strategic recommendations.
- Executive education sessions tailored to oversight responsibilities.
- Maturity progression planning and KPI development.

Third-Party AI Compliance Audit (Independent Assessment)

Third-Party AI Compliance Audit provides a rigorous, third-party assessment of AI systems, **delivering an objective verification** of alignment with applicable legislative, regulatory, and contractual requirements. Purpose-built for organizations operating in high-scrutiny environments, this engagement provides essential stakeholder assurance through systematic evidence collection, forensic documentation review, and the production of "audit-ready" substantiation materials tailored for regulators, customers, and external attestations. By serving as a strictly independent evaluator, we provide the specialized expert analysis necessary to validate complex AI frameworks, support formal regulatory submissions, and offer technical clarity in dispute resolution contexts, ensuring that AI governance meets the highest standards of transparency and defensibility.

Deliverables may include:

- Compliance verification mapped to relevant requirements (jurisdiction, industry, use case).
- Technical documentation, model/system card review, and evidence assessment.
- Support for customer and regulator assurance attestations (evidence packaging and response substantiation).
- Audit-ready documentation package and findings summary.
- Expert analysis support for disputes and litigation matters (as scoped).

Industry Focused AI Governance

The same governance objectives apply across sectors, but regulatory emphasis, documentation norms, and AI use cases create meaningful differences in how risk is classified and controlled. HaystackID AI Governance Services are designed for **highly regulated, high-stakes environments** where AI decisions carry material operational, legal, and reputational impacts.

Industry-specific delivery incorporates applicable regulatory obligations, common AI use cases, and risk patterns seen in production deployments. Across industries, AI governance programs typically require alignment between policy expectations and **measurable operational controls**, supported by documentation suitable for audits, regulator engagement, and stakeholder assurance.

Financial Services

Financial institutions deploying AI face overlapping model risk management, fair lending, and consumer protection requirements. Governance typically emphasizes disciplined controls, documentation, and monitoring for AI-enabled decisioning and customer-impacting outcomes.

AI governance elements commonly addressed include:

- Model risk governance alignment.
- Fair lending governance and evidence-based practices.
- Oversight controls for AI-enabled financial products and regulator readiness.
- Bias and disparate impact evaluation protocols and documentation standards.
- Third-party AI governance, vendor assurance, and contractual controls.

Key use cases include credit decisioning, fraud detection, algorithmic trading, AML, and robo-advisory.

Healthcare and Life Sciences

Healthcare AI introduces risk profiles combining patient safety, privacy obligations, and health equity expectations. Governance commonly centers on validation, monitoring, change control, and documentation practices that support clinical confidence and accountability.

AI governance elements commonly addressed include:

- Validation governance for patient-impacting AI.
- Privacy and security controls for AI processing health information.
- Clinical performance monitoring and drift management expectations.
- Health equity and bias evaluation governance across demographic populations.
- Third-party AI oversight for clinical tools and operational platforms.
- Documentation practices for stakeholder assurance and regulator engagement.

Key use cases include diagnostic imaging, clinical decision support, drug discovery, patient risk stratification, biomedical research, and claims fraud detection.

Insurance

Insurers using AI for underwriting, claims, and pricing face increased scrutiny from regulators. Governance typically focuses on explainability, non-discrimination controls, documentation readiness, and oversight of automated workflows.

AI governance elements commonly addressed include:

- Governance alignment with regulatory expectations and industry guidance.
- Non-discrimination oversight for underwriting and pricing decisions.
- High-risk AI governance readiness for applicable requirements.
- Explainability and documentation controls to support filings and inquiries.
- Vendor governance and assurance for third-party AI models and tools.
- Monitoring and reporting practices for model changes and outcomes.

Key use cases include automated underwriting, claims triage, fraud detection, dynamic pricing, unaffordability, and exclusion.

Employment and Human Resources

Employment-related AI use cases are sensitive because they can affect individual rights, access to employment, and claims of discrimination. Adoption and control practices vary widely, and many organizations apply more conservative governance and validation before deploying tools that influence hiring, promotion, performance, and workforce management decisions. Governance typically requires auditability, disclosure readiness, and defensible testing practices for systems that may materially impact candidates or employees.

AI governance elements commonly addressed include:

- Bias audit governance and documentation expectations.
- EEOC-aligned oversight for disparate impact and accommodation considerations.
- Governance readiness for evolving state and federal employment AI requirements.
- Risk classification and governance controls for high-risk employment AI (EU AI Act).
- Disclosure, notice, and documentation practices supporting audit readiness.
- Monitoring controls for model changes, vendor updates, and workflow modifications.

Key use cases include resume screening, candidate ranking, video interview analysis, performance prediction, and transparency reporting.

Government and Public Sector

Public-sector AI deployment requires heightened accountability, transparency, and due-process protections. Governance commonly emphasizes rights-impacting classification, oversight reporting, procurement controls, and public accountability.

AI governance elements commonly addressed include:

- Governance alignment with federal direction and oversight expectations.
- NIST AI RMF implementation and maturity progression within agency frameworks.
- Rights-impacting AI classification and due process governance considerations.
- Community impact analysis and public accountability documentation practices.
- Procurement and third-party AI oversight controls for government use.
- Incident response and escalation procedures for public-impacting AI systems.

Key use cases include benefits administration, fraud detection, public safety, permitting, constituent services, facial recognition, procurement, and contracts.

Technology and AI Providers

Technology providers building AI products face direct compliance obligations and customer-driven governance requirements. Governance typically requires repeatable documentation, assurance processes, and testing evidence that can support enterprise customer expectations.

AI governance elements commonly addressed include:

- Provider governance obligations for general-purpose AI and foundation models (EU AI Act).
- ISO/IEC 42001 readiness and AI management system governance practices.
- Model card and system card documentation governance, and evidence standards.
- Security validation governance, red-team readiness, and customer assurance artifacts.
- Governance controls supporting transparency, accountability, and customer audits.
- Vendor and supply-chain governance for embedded models and third-party components.

Key use cases include foundation models, AI SaaS products, embedded AI features, and AI-powered analytics.

AI Governance Engagement Approach

HaystackID's engagement model is designed to align operations, product, and revenue stakeholders with GRC, privacy, and legal teams, and to bridge the gap between business delivery and compliance expectations through **an implementable AI governance operating model**.

Engagements are commonly structured as staged workstreams that begin with visibility and prioritization, progress through program implementation, and extend into system validation and sustained oversight. This approach enables organizations to **align governance investment to regulatory timelines, operational urgency, and board-level reporting expectations** while maintaining continuity across stakeholders.

AI Governance Phases: 12-Month Engagement View

Engagement	1-3 Months	4-6 Months	7-9 Months	10-12 Months
Scoping Engagement (3–8 weeks)	■			
Advisory Implementation (6–12 months)	■	■	■	■

Scoping Engagement (3-8 weeks): Understand AI landscape and risk exposure.

- AI inventory and risk classification.
- Regulatory exposure and gap analysis.
- Prioritized roadmap, sequencing, and resource plan.

Advisory Implementation (6-12 months): Build comprehensive governance programs.

- Governance operating model, policies, procedures, and controls.
- Risk management workflows, documentation standards, and training.
- Evidence capture aligned to audit and regulator expectations.

This engagement approach is designed to remain modular, enabling organizations to start with the highest-risk systems and expand governance coverage over time without disrupting business adoption.

Consistency and defensibility are strengthened when governance activities map to **recognized standards and guidance** rather than ad hoc policies.

Alignment with Leading Frameworks

HaystackID’s AI Governance methodology aligns with recognized standards and frameworks to support **consistency, defensibility, and audit readiness**. Alignment enables organizations to translate governance objectives into measurable controls and documentation that can be communicated across internal stakeholders and, when required, external regulators and assurance audiences.

HaystackID alignment commonly includes:

- NIST AI Risk Management Framework (AI RMF).
- ISO/IEC 42001 Artificial Intelligence Management System (AIMS) standard.
- EU AI Act risk classification, documentation, transparency, and post-market expectations.
- OECD AI Principles.
- Industry-specific guidance and regulator expectations relevant to the sector and use case.

By grounding governance activities in established frameworks, organizations can demonstrate that governance design, testing protocols, documentation practices, and oversight reporting are anchored to recognized standards rather than ad hoc policy statements. Operational delivery also accounts for geography, stakeholder groups, and enterprise maturity levels so governance remains implementable and repeatable.

Availability and Deployment

HaystackID AI Governance Services are designed for **rapid activation and scalable delivery** across enterprise maturity levels - from early-stage AI adoption to multi-business-unit programs requiring coordinated governance, validation, and oversight. Engagements can be deployed as discrete assessments, system-level validation sprints, program implementation initiatives, and recurring oversight support, based on organizational priorities and the risk posture of AI use cases.

Services are delivered through a multidisciplinary team model spanning governance, regulatory alignment, risk management, security testing, and fairness evaluation. Delivery can be integrated into established enterprise risk and technology governance structures to support **continuity, evidence capture, and audit readiness**.

Services are available across the six primary offerings described above and can be delivered remotely, onsite, or hybrid based on stakeholder needs. **Primary offerings include:**

AI Governance Scoping

AI Governance Advisory

AI Security Testing

AI Fairness Testing

Board Advisory Services

Third-Party AI Compliance Audit

Each offering can be structured as a fixed-scope assessment, a time-boxed sprint, a phased program with milestones, or an ongoing retainer to sustain governance maturity and audit readiness.

Operational considerations commonly addressed include:

- Governance workflow integration with enterprise change management.
- Evidence capture and documentation standards for audit readiness and customer assurance.
- Third-party AI intake, due diligence, and contractual governance support.
- Reporting cadence and executive/board visibility through dashboards and briefings.
- Coordination with internal counsel and compliance teams to align governance outputs with obligations.

A direct path to initiate an engagement supports **timely risk reduction and alignment with regulatory timelines**.

Appendix: Acronyms and Abbreviations

- **AEDT** - Automated Employment Decision Tool (term used in the context of NYC Local Law 144)
- **AI** - Artificial Intelligence
- **AI/ML** - Artificial Intelligence / Machine Learning
- **AI RMF** - Artificial Intelligence Risk Management Framework (published by NIST)
- **AIMS** - Artificial Intelligence Management System (term used in the context of ISO/IEC 42001)
- **AML** - Anti-Money Laundering
- **API** - Application Programming Interface
- **ECOA** - Equal Credit Opportunity Act (U.S.)
- **EEOC** - Equal Employment Opportunity Commission (U.S.)
- **EO** - Executive Order (U.S.)
- **EU** - European Union
- **FINRA** - Financial Industry Regulatory Authority (U.S.)
- **HIPAA** - Health Insurance Portability and Accountability Act (U.S.)
- **IDC** - International Data Corporation
- **ISO/IEC** - International Organization for Standardization / International Electrotechnical Commission
- **ISO/IEC 42001** - ISO/IEC standard for an Artificial Intelligence management system
- **KPI** - Key Performance Indicator
- **LLM** - Large Language Model
- **ML** - Machine Learning
- **MLOps** - Machine Learning Operations
- **NAIC** - National Association of Insurance Commissioners (U.S.)
- **NIST** - National Institute of Standards and Technology (U.S.)
- **NYC** - New York City
- **OCC** - Office of the Comptroller of the Currency (U.S.)
- **OCC 2011-12** - OCC Bulletin 2011-12 (Supervisory guidance on model risk management)
- **OECD** - Organisation for Economic Co-operation and Development
- **OMB** - Office of Management and Budget (U.S.)
- **OMB M-25-22** - Office of Management and Budget memorandum M-25-22
- **Reg B** - Regulation B (implementing regulation for ECOA, U.S.)
- **SaMD** - Software as a Medical Device
- **SDLC** - Software Development Life Cycle
- **SR 11-7** - Federal Reserve Supervisory Letter 11-7 (Model Risk Management guidance)

Learn More Today.

Contact HaystackID to discuss how AI Governance Services can support **responsible innovation, defensible oversight, and audit-ready compliance** across legal, compliance, regulatory, and cybersecurity requirements.

About HaystackID®

HaystackID solves complex data challenges related to legal, compliance, regulatory, and cyber requirements. Core offerings include Global Advisory, Cybersecurity, Core Intelligence AI™, and ReviewRight® Global Managed Review, supported by its unified CoreFlex™ service interface. Recognised globally by industry leaders, including Chambers, Gartner, IDC, and Legaltech News, HaystackID helps corporations and legal practises manage data gravity, where information demands action, and workflow gravity, where critical requirements demand coordinated expertise, delivering innovative solutions with a continual focus on security, privacy, and integrity. Learn more at HaystackID.com.

Assisted by GAI and LLM technologies.