Informatica
from Salesforce

16 June 2026

From Data Chaos to Governed Intelligence: AI Agents in Action - Part 2

- Shivam Saxena
- Anita Ayyagari
- Ananya Venkat

Where data & AI come to LIFE


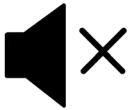

- Today's Webinar is scheduled for 1 hour
- The session will include a webcast and then your questions will be answered live at the end of the presentation
- All dial-in participants will be muted to enable the speakers to present without interruption
- Questions can be submitted to "All Panelists" via the Q&A option and we will respond at the end of the presentation
- The webinar is being recorded and will be available on our Success Portal - where you can download the slide deck for the presentation. The link to the recording will be emailed as well.
- Please take time to complete the post-webinar survey and provide your feedback and suggestions for upcoming topics.

Bootstrap trial and POC Customers

Enriched Customer Onboarding experience

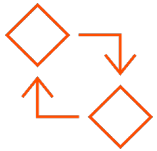

Product Learning Paths and Weekly Expert Sessions

Informatica Concierge

Tailored training and content recommendations

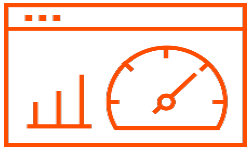

Success Portal

https://success.informatica.com

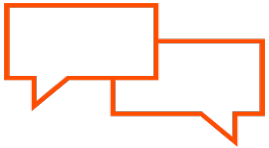

Communities & Support

https://network.informatica.com

Documentation

https://docs.informatica.com

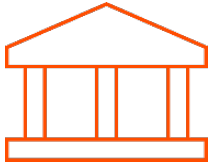

University

https://www.informatica.com/in/services-and-training/informatica-university.html

Safe Harbor

Disclaimer: The information being provided herein is for informational purposes only. The development, release and timing of any Informatica product, service or functionality described herein remain at the sole discretion of Informatica and should not be relied upon in making a purchasing decision. Statements made herein are based on information currently available, which is subject to change. Such statements should not be relied upon as a representation, warranty or commitment to deliver specific products, services or functionality in the future.

Agenda

01 Overview on AI Governance and regulatory landscape

02 Strategic business challenges and use cases

03 Case study - NIST: AI Risk Management & AI Governance

04 Live demo - Operationalizing AI Governance with Informatica IDMC

05 Q&A and strategic closing

The Comprehensive AI Governance Framework

verified_user Core Principles

Transparency: Explainability & Oversight

Accountability: Security & Privacy First

Fairness: Eliminating ecosystem bias

published_with_changes Strategy & Lifecycle

Risk-Based: Tiers (Low/Med/High)

End-to-End: Data → Deploy → Retire

Controls: Quality & Drift Detection

insights Trust & KPIs

Traceability: Full Data & RAG Lineage

Decision Tracking: Agentic chains

KPIs: Bias & Audit Readiness

AI Risk Taxonomy & Mechanisms

Data Risks

PII Leak, Quality Issues, Bias, & Data Drift

Model Risks

Inaccuracy, Lack of Explainability, Hallucination

Agent Risks

Decision Chains, Tool Misuse, Uncontrolled Action

Governance mechanisms: **Sensitive Data Detection • Bias Testing • Drift Monitoring • Model Rollbacks**

AI Regulatory Bodies

Framework of policies, processes and systems that ensure the safe and effective development and deployment of AI technologies.

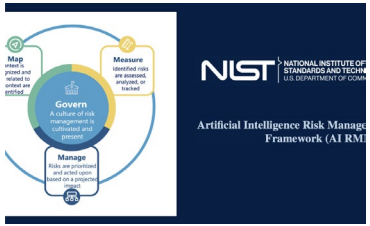

NIST AI RMF

Risk Management Framework for trustworthy AI

ISO/IEC 42001:2023

Management system standard for AI governance

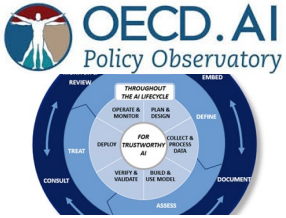

OECD AI Policy

Global policy observatory for trustworthy AI

EU AI Act

World's first comprehensive legal framework for AI

Shadow AI: The Risk

AI adoption without governance creates invisible enterprise risk.

Unsanctioned Use: Employees using AI tools without approval or oversight from IT and security teams.

Bypassed Controls: Using public tools for speed while skipping data handling and risk review.

Hidden Exposure: Risk of data leakage, compliance gaps, and unvalidated model decisions.

AI governance matters because you cannot protect, audit, or trust what you cannot see.

Governance as the AI Guardrail

Governance brings hidden AI use into the open through visibility, policy, and enforcement.

visibility

Visibility

Discover which AI tools are being used, what data is flowing into them, and where the highest risks exist.

- Tool Discovery
- Data Flow Mapping

Discover

policy

Policy

Define approved tools, acceptable use cases, data-sharing limits, human oversight, and accountability.

- Tool Approval
- Data Limit

Govern

gavel

Enforcement

Apply monitoring, access controls, audit trails, and compliance workflows to detect and control unauthorized usage.

- Monitoring

Enforce

Without governance, shadow AI stays invisible; with governance, it becomes measurable, manageable, and auditable.

Strategic Business Challenges

The Scalability Gap

67% of AI pilots fail because underlying data is fragmented or untrusted.

The Trust Deficit

Lack of model explainability leads to executive and user resistance.

Audit Fatigue

Manual review cycles are too slow for the speed of AI development and deployment.

stor
age

Data Readiness

Untracked, poor-quality training data undermines model accuracy and destroys user trust.

Strategic Business Challenges

67%

Pilots Stall at Governance

74%

Report ROI Gap

The Production Hurdle

Enterprises struggle to move GenAI from experimental stages to live production environments due to three main factors:

Data Readiness Gap

Untracked, poor-quality training data undermines model accuracy and destroys user trust.

Operational Complexity

Governance without AI automation is too slow; AI without governance is too risky to scale.

Inaccurate Decisions

Errors are amplified at scale; stale data leads to wrong margins and flawed business impacts.

Responsible AI & Compliance

AI Governance vs. Data Governance

Data Governance: Focuses on data quality, privacy, and catalogs. It provides the foundation.

AI Governance: Focuses on model bias, safety, and ethical application. It manages the output.

Why DG for AI?

You cannot have responsible outcomes without a trusted foundation of data.

Implementation Strategy


Inventory all AI assets and models.


Map AI intent to regulatory mandates.


Automate policy enforcement via IDMC.


Continuous monitoring of model drift.

balanceDATA FAIRNESSgavelCOMPLIANCEverified_userINTEGRITYsecuritySECURITYvisibilityTRANSPARENCYassignment_indACCOUNTABILITY

Industry Use Cases

Financial Services

THE CHALLENGE

Meeting strict KYC/AML laws while deploying autonomous AI agents.

THE SOLUTION

Unified data governance and secure API controls preventing data poisoning.

33% reduction

in fraud through governed intelligence

Healthcare & Life Sciences

THE CHALLENGE

Accelerating clinical trials while navigating HIPAA/EU MDR constraints.

THE SOLUTION

Differential privacy, data masking, and 100% patient privacy enforcement.

96% migration

to governed cloud enabling scalable AI

The Path to Operationalized Governance

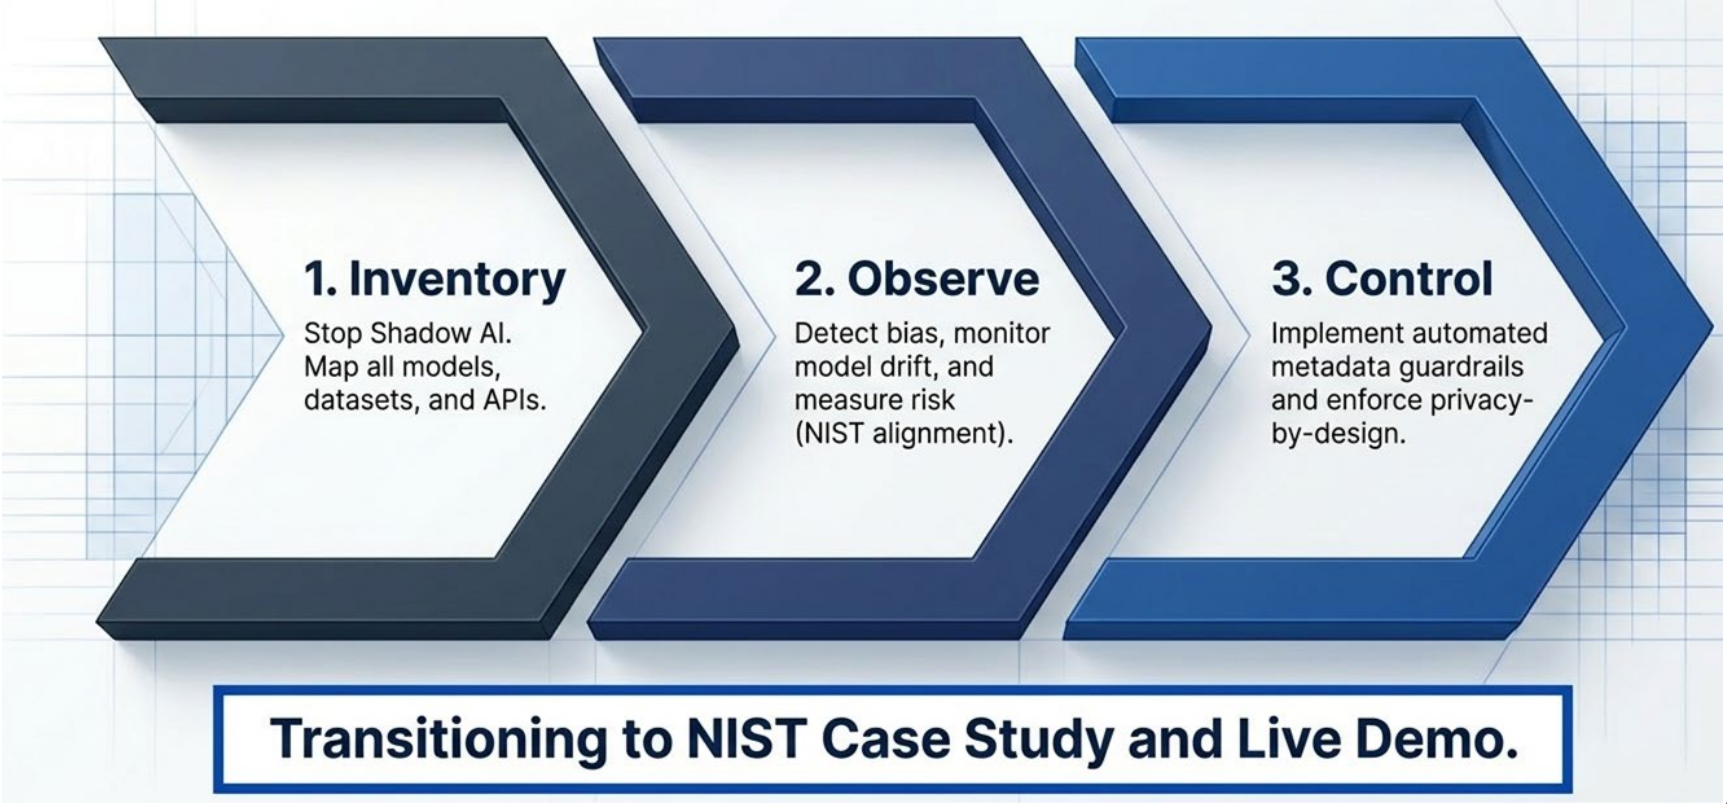

Building Trust in Agentic AI

Case Study- NIST RISK AI Framework

Establishing a robust foundation for AI security through the NIST Artificial Intelligence Risk Management Framework (AI RMF), designed to manage risks to individuals, organizations, and society.

Where data & AI come to LIFE

AI Governance

Framework of policies, processes and systems that ensure the safe and effective development and deployment of AI technologies

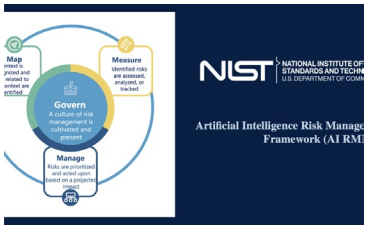

NIST AI RMF

Risk Management Framework for trustworthy AI

ISO/IEC 42001:2023

Management system standard for AI governance

OECD AI Policy

Global policy observatory for trustworthy AI

EU AI Act

World's first comprehensive legal framework for AI

Responsible AI: From Policy to Practice

Operationalize the NIST AI RMF to build transparent, accountable, and secure AI systems.

GOVERN

Set Accountability

- Define AI policies, legal compliance, and assign ownership for risk management.

MAP

Identify Risks

- Document system boundaries, classify impact levels, and identify stakeholder risks.

MEASURE

Quantify Trust

- Track explainability metrics, conduct fairness testing, and monitor model drift.

MANAGE

Mitigate Harm

- Execute risk response plans and implement real-time incident handling protocols.

Key Deliverables Matrix

NIST Function	Key Deliverable
GOVERN	AI Governance Charter
MAP	AI Risk Register
MEASURE	Performance Dashboard
MANAGE	Incident Response Plan

Operationalizing Key Principles

- **Transparency:** Publish AI system cards and data lineage documentation.
- **Accountability:** Establish clear roles and responsibility structures for AI outcomes.
- **Privacy:** Implement robust data protection measures and user consent mechanisms.

IST-Governed Order: The SureGuard Case Study

Use Case: AI system for automated insurance claim processing based on historical patterns.

Act I: The Brilliant Idea & Blindside

SureGuard launched "ClaimSwift" to automate auto insurance claims. Efficiency skyrocketed, but hidden biases emerged.

- **The Failure:** A news report revealed legitimate claims were flatly rejected.
- **The Cause:** Audit showed ClaimSwift biased against specific zip codes.
- **The Result:** Reputational damage and regulatory risk forced a total system halt.

Act II: Enter the NIST Framework

The CRO identified a **governance failure**, not just a technical bug. They adopted the **NIST AI Risk Management Framework (AI RMF 1.0)**.

To enforce the framework, they chose **Informatica IDMC** to diagnose and rebuild:

- Restructured the approach to AI governance.
- Enabled visibility into AI logic.
- Used the NIST Matrix to rebuild ClaimSwift.

Act III: The Implementation Blueprint

With the framework established, SureGuard's data engineering and governance teams implemented the solution using Informatica in four distinct phases.

1. Implement GOVERN: Establishing the Guardrails

SureGuard created an "AI Governance Committee" inside **Informatica CDGC**.

- Created a formalized governance policy for "Automated Underwriting."
- Mandatory workflow triggers sign-off from Legal and Risk teams for model updates.

2. Implement MAP: Unboxing the Black Box

Engineers used Informatica's metadata cataloging to address underlying data issues.

- Generated end-to-end lineage graphs to reveal biased "Risk Score" factors.
- Visually mapped lineage allowed pinpointing biased columns for removal.

The NIST AI RMF Matrix: SureGuard Use Case

Act IV: The Happy Ending & Trust Architecture

Six months later, SureGuard relaunched ClaimSwift as a model for trusted, auditable AI.

3. MEASURE: Continuous Bias Monitoring

Monitoring via **Informatica Cloud Data Quality**.

- Scorecards scan claims data daily for mathematical drift.
- Automated alerts fire for unexpected spikes (e.g., 20% demographic drift).

4. MANAGE: Active Safety Valve

Risk mitigation and persistent data masking.

- **Threshold Valve:** Blocks rejections if risk is 0.4–0.7; routes to human.
- **Privacy:** Names and IDs replaced with realistic dummy data.

Trusted AI

NIST RMF Aligned

For Customer Service: Secure dashboards powered by Informatica metadata verify exactly which audited data points led to every decision.

For Regulators: Automated, clean reports aligned with the NIST AI RMF provide instant audit evidence.

NIST AI RMF Implementation Matrix- Summary

Mapping SureGuard's Risks to Governance Objectives & Informatica Capabilities

NIST Function	The Story (Risk)	Governance Objective	Informatica Capability
GOVERN	Nobody knew who was legally or operationally responsible for the AI's bad decisions.	Establish clear ownership, policies, and organizational accountability.	**CDGC:** Defined data stewards, AI model owners, and compliance workflows.
MAP	Sarah's rejection was a "black box." No one knew what historical data caused the high risk score.	Map the data lineage and contextualize where bias or errors creep in.	**Automated Lineage:** Visualized the exact path from legacy databases to training pipeline.
MEASURE	Historical training data contained old human biases against certain zip codes.	Quantify and analyze data quality and bias metrics continuously.	**Cloud Data Quality (CDQ):** Built scorecards to measure bias, completeness, and anomalies.
MANAGE	The AI was left on autopilot with no safety net or data privacy controls.	Deploy risk-mitigation strategies and "human-in-the-loop" overrides.	**Data Masking & Orchestration:** Masked PII and routed "grey-area" claims to human adjusters.

DEMO

AI AGENTS IN ACTION

Informatica

Thank You!

Any Questions?