

## **MARKET CONTEXT: FROM RISK IDENTIFICATION TO OPERATIONAL DECISION MAKING**

AI adoption is accelerating across the enterprise, but governance maturity is not keeping pace. Most organizations have made meaningful progress in understanding where AI is used and aligning systems to internal policies and emerging regulatory frameworks. Capabilities such as system inventory, classification, and policy mapping are now well established.

A more fundamental challenge is emerging.

The constraint is no longer identifying risk. It is determining, in a consistent and defensible way, what actions are required when issues arise and how those decisions should be documented.

As AI systems move into production, organizations are encountering bias, drift, unexpected outputs, and failures in automated decision-making. These events require coordination across legal, privacy, compliance, and technical teams. The challenge is not detection. It is interpretation and action.

Organizations are struggling to answer key questions:

- What does this issue mean from a regulatory perspective?
- Are obligations triggered, including notification requirements?
- What actions must be taken, and how should they be documented?

This gap creates inconsistency across teams and increases exposure under regulatory scrutiny. It is increasingly recognized in practice, even if it is not yet fully defined as a category.

## **RADARFIRST APPROACH: A DECISIONING LAYER WITHIN AI GOVERNANCE**

RadarFirst is designed to address this gap by introducing a structured decisioning layer within AI governance.

Rather than focusing on generating signals or alerts, the platform determines what actions are required when AI systems produce outcomes that need evaluation. This includes whether regulatory notifications, escalations, or remediation steps are triggered, and ensures that those decisions are documented in a consistent, repeatable way.

The result is a shift from identifying issues to producing decisions that can be acted on and defended.

## BRIDGING THE GAP BETWEEN GOVERNANCE AND ACTION

The current AI governance landscape comprises several adjacent categories. Governance platforms provide centralized oversight and policy alignment. Monitoring and observability tools detect technical signals and model behavior. GRC platforms support controls, documentation, and reporting.

Each addresses part of the problem, but none directly addresses the operational step between identifying an issue and determining the required actions.

RadarFirst fills this gap by interpreting events from a regulatory and organizational perspective. It determines what an issue means, whether obligations are triggered, and how decisions should be documented.

This represents a shift within AI governance from observation to decision-making.

## PLATFORM ARCHITECTURE: A CLOSED-LOOP GOVERNANCE MODEL

RadarFirst is structured around two core areas of regulatory risk management: privacy incident management and AI incident management. Together, these serve as systems of record for evaluating, responding to, and documenting incidents.

At the foundation of the platform is a deterministic model.

- Regulatory requirements and internal standards are codified into structured logic so that identical inputs produce consistent, explainable, and repeatable outcomes across teams, use cases, and jurisdictions.

This approach addresses a fundamental tension.

- AI systems are probabilistic, while regulatory compliance requires consistency and defensibility. The deterministic engine removes ambiguity from decision-making and creates a system of record for decisions, not just signals.

Supporting this foundation are capabilities such as system inventory, classification, and ongoing risk evaluation, along with customizable workflows that extend decision-making into adjacent regulatory use cases. An agentic layer operates across workflows to improve intake, triage, investigation, and coordination, while preserving deterministic decisioning at the core.

## THE ROLE OF AI: SUPPORTING THE WORKFLOW

Artificial intelligence plays an important but clearly defined role within the platform.

AI is used to improve efficiency across intake, triage, and investigation. It helps identify missing information, prioritize high-risk cases, and coordinate across stakeholders.

It helps identify missing information, prioritize high-risk cases, and coordinate across stakeholders. In this model, AI prepares and accelerates work, while the deterministic engine produces structured guidance, and humans make the final decisions.

This approach allows organizations to scale operations without introducing variability into outcomes, preserving accountability while improving speed and efficiency.

## FINAL PERSPECTIVE & BUYER CONTEXT

AI governance is evolving from visibility to accountability. Organizations are moving beyond understanding what systems exist toward being responsible for how those systems behave and how issues are evaluated, resolved, and documented.

RadarFirst provides a structured approach to this challenge.

With system context, incident evaluation, deterministic decisioning, and workflow orchestration, the platform helps organizations move from manual, inconsistent interpretation to repeatable, defensible decision-making at scale.

The goal is not just to identify risk. It is to ensure organizations can consistently determine what that risk means and what must be done about it.

RadarFirst is designed for organizations operating in regulated or high-risk environments where consistent decision-making is critical. Primary stakeholders include leaders across privacy, legal, compliance, risk, and AI governance, all of whom are increasingly responsible for evaluating AI system behavior and ensuring defensibility under regulatory scrutiny.



### ABOUT RADARFIRST

*RadarFirst is an AI-forward regulatory risk management platform for privacy and AI incidents. Trusted by leading organizations, RadarFirst enables teams to manage incidents with speed, consistency, and defensibility by standardizing how incidents are captured, assessed, and actioned.*

*With more than a decade of experience supporting real-world privacy incidents, including over 1 million incidents managed, 4 million regulatory decisions delivered, and more than \$2 trillion in revenue protected, RadarFirst brings proven operational discipline to modern regulatory risk. As AI drives greater incident volume, complexity, and scrutiny, RadarFirst helps organizations prepare before risk becomes operational, using AI-assisted capabilities to support intake, assessment, and prioritization while preserving human judgment.*