Appendix 1: The Evidence Stack

This series is designed to move from theory to proof: from defining AI governance as infrastructure, to mapping where consequential decisions occur, to systematically evaluating the domains where evidentiary control must exist. The introductory "AI Governance Is Not Policy. It Is Infrastructure." article establishes the core thesis that governance is no longer a policy exercise but an operational evidence problem. Project Zero then identifies the enterprise decision surfaces where AI systems materially influence outcomes. Building on that foundation, each subsequent domain in the RSAC Field Notes applies a consistent methodology to assess how vendors support the evidentiary chain from authorization to system behavior, with corresponding battle cards to reflect current market capabilities. Taken together, this series provides a structured framework for evaluating whether an organization can not only state its AI governance posture but prove it.

Series Stage	*Article*	*Core Concept*	*What It Establishes*	*Vendors (Where Applicable)*
Foundation	*AI Governance Is Not Policy. It Is Infrastructure.*	*Governance is infrastructure, not policy*	*Establishes shift from policy to evidence and introduces "pipeline is the policy"*	*N/A*
Phase 0	*Project Zero — Decision Surface Mapping (RSAC Field Notes Part 1)*	*Decision surface mapping*	*Identifies where consequential decisions occur and defines governance scope*	*N/A*
Domain 1	*Governance and Risk Orchestration (Part A — Methodology published)*	*Evidence architecture and legal accountability*	*Defines evidentiary chain from authorization to runtime behavior; introduces evidence objects*	*IBM, Credo AI, Collibra, Monitaur*
Domain 2	*AI Discovery and Security Posture Management (Upcoming)*	*System identification and AI visibility*	*Determines whether all AI systems touching enterprise data can be identified in real time*	*Wiz (AI-SPM), Orca Security, Noma Security, Pillar Security, Cranium*
Domain 3	*Agent Orchestration and Workflow Control (Upcoming)*	*Action authorization and execution control*	*Defines how agent actions are authorized, constrained, and recorded*	*Microsoft (Semantic Kernel, AutoGen, Copilot Studio), Amazon Bedrock, LangGraph, Dagster, Apache Airflow*
Domain 4	*Data Security Posture Management (Upcoming)*	*Data exposure and control boundaries*	*Identifies what sensitive data has reached models and whether exposure was known or controlled*	*Netskope, Cyberhaven, Microsoft Purview, Google Workspace DLP, Nightfall AI*
Domain 5	*Data Lineage and Pipeline Visibility (Upcoming)*	*Output traceability and lineage*	*Enables step-by-step reconstruction of model outputs and system behavior*	*Collibra, Monte Carlo, Apache Atlas, OpenLineage*
Domain 6	*Identity and Access Governance for AI Systems (Upcoming)*	*Identity, access, and accountability*	*Defines what AI systems can access and who is accountable for those permissions*	*Okta, Azure AD, ConductorOne*
Domain 7	*Runtime Protection and Behavioral Monitoring (Upcoming)*	*Live monitoring and intervention*	*Establishes how organizations detect, escalate, and respond to abnormal AI behavior*	*Protectt.ai, Robust Intelligence (Cisco), Lakera Guard, Lasso Security, Arthur AI, MLflow, Giskard*
Domain 8	*AI Supply Chain and Model Integrity (Upcoming)*	*Model provenance and change integrity*	*Addresses how model updates are tracked and whether behavioral changes are understood*	*HiddenLayer, Protectt.ai, ModelScan*