AI Security & Governance Whitepaper

Secure Data and AI From Code-to-Cloud

AI Security Posture Management & AI Governance

relyanceai

Table of Contents

Executive Summary

Relyance AI's Unified AI Security Posture Management & Governance platform represents a fundamental leap forward in how organizations can secure and govern AI applications in production. As AI adoption accelerates across first-party models, third-party APIs, and AI features embedded in SaaS, the attack surface expands—and so does regulatory scrutiny. These non-deterministic, high-agency systems no longer follow fixed control paths; data itself becomes executable, shaping model behavior and agent actions at runtime.

Relyance AI addresses this challenge with a data-journey-first approach. By mapping every AI data flow—from source code and pipelines to inference and agent actions—the platform delivers the unified visibility and control needed to address following CISO-critical AI risks around:

- **Sensitive Data Leakage** – Exposure of regulated data through AI training or inference flows
- **Autonomous Agent Overreach** – AI agents chaining tools in unsafe, unintended ways
- **Overprivileged AI Access** – Excessive model permissions or weak IAM exposing AI asset
- **Shadow & Unsanctioned AI** – Unapproved AI use creating blind spots and high risk
- **Compliance Drift** – Runtime gaps against security and regulatory controls (NIST AI RMF, ISO 42001, OWASP, EU AI Act etc)

This Live AI Inventory and AI Data Journeys™ capability goes beyond static asset list to capture who or what touched the data, where it went, and why—creating a single source of truth for both SecOps and Privacy teams.

For CISOs, this makes AI Security Posture Management an operational reality: a real-time, explainable view of AI applications, their training and inference data flows, exposure points, and control status. For privacy and compliance leaders, it means AI Governance can be continuously monitored and compliance gaps closed at runtime, with automated mapping and gap analysis to frameworks like the EU AI Act and ISO 42001. In a landscape where AI risk and AI opportunity are inseparable, Relyance AI provides the unified trust layer to secure and govern AI.

Why AI Creates a Never-Before-Seen Risk Landscape

AI isn't just "software with data" – it's software whose behavior is continuously shaped by data. Unlike traditional applications, where logic is fixed in code, AI systems generate decisions, actions, and content in real time based on incoming inputs. This creates a risk surface unlike anything CISOs or Privacy teams have had to secure before.

- **Data as Executable Logic** – In AI systems, the data isn't just used, it's run. A customer record, a support ticket, or a workflow log can directly shape model outputs, prompt behavior, or even trigger downstream system actions.
- **Behavioral Fluidity** – AI doesn't follow pre-coded control paths. The same model, on the same endpoint, may behave differently from one moment to the next depending on inputs and context.
- **Agentic Autonomy** – With the rise of agentic AI, models aren't limited to responding – they can take action, chain tools, call APIs, and iterate on plans. This gives them a form of operational "agency" that evolves as they encounter new data.

Dual Risk Lenses

- SecOps must manage a moving target: dynamic decision logic, unpredictable data flows, and output behaviors that can introduce new security gaps instantly.
- Privacy & GRC teams must govern an environment where data protection rules are harder to enforce – because compliance obligations now apply not just to stored or transferred data, but also to how that data influences model behavior.

In short, AI's risk surface is both behavioral and informational. Securing and governing it means tracking not only where data goes, but how it changes what the system does.

The Unified AI Security & Governance Imperative

SecOps and Privacy teams are staring at the same AI system – but from different angles. The very behaviors that create security exposure – overprivileged actions, sensitive data flows, unverified tool calls – are often the same ones that trigger compliance violations. And in AI, these risks emerge at runtime, not in neat quarterly reports. When these teams operate on separate tools and partial inventories, blind spots multiply, response slows, and risks go unaddressed.

The Power of One Unified Platform:

- **Unified Visibility**: One shared view of AI inventory, data flows, and runtime behavior closes coordination gaps.
- **Shared Controls:** Policies and safeguards are applied consistently at runtime when both SecOps and Privacy operate from the same live context
- **No More Reconciling Reports:** One source of truth means no duplicated analysis or conflict when issues arise – both teams act decisively, aligned.
- **Scales from Development to Runtime:** Whether it's models under development, APIs in SaaS, or autonomous agents in action, there are no blind spots.

Put simply: Unification isn't just a nice-to-have – it's mandatory. When security and governance act from one platform, you not only contain risk – you accelerate AI innovation safely. Anything less puts you at risk of managing only pieces of the truth, and no one needs half security or half compliance in the AI era.

The Unified AI Security & Governance Platform at Relyance

Relyance AI unifies AI security posture and AI governance in a single, runtime platform – closing the gap between SecOps and Privacy teams so risks are seen in context, acted on instantly, and governed under the same controls.

At its core, the AI-powered Trust Intelligence (TrustiQ™) engine ingests signals across the AI footprint – including source code, cloud environments, model registries, datasets, MLOps workflows, and third-party AI vendors – all through agentless integrations, continuously maintaining a unified, always-current data graph. This shared foundation enables cross-functional visibility and coordinated controls across the entire AI lifecycle.

Data Journeys – The Crown Jewel of Unified Platform

At the heart of this unified platform is Data Journeys™ – a dynamic, runtime-aware data map showing exactly how data moves through AI applications for training and inference flows.

- **For SecOps**: Pinpoints sensitive data entry, flow, and exposure; flags unsafe agent actions; and reveals overprivileged access.
- **For Privacy & GRC**: Provides the traceability needed to demonstrate lawful, transparent, and accountable AI operations.

This level of visibility is not optional – it's what leading frameworks and regulations now expect. The **EU AI Act** calls for full documentation of input-to-output data flows, the **NIST AI RMF** emphasizes data provenance and usage tracking, and **ISO/IEC 42001** mandates end-to-end traceability to prove quality and lawful use.

By delivering a single, living data map that meets these demands, Data Journeys™ power both operational defense and regulatory compliance - making it the *crown jewel* of the Relyance AI platform.

Below, we unpack the platform's core capabilities, empowering both security and privacy teams to act with confidence across their entire AI footprint.

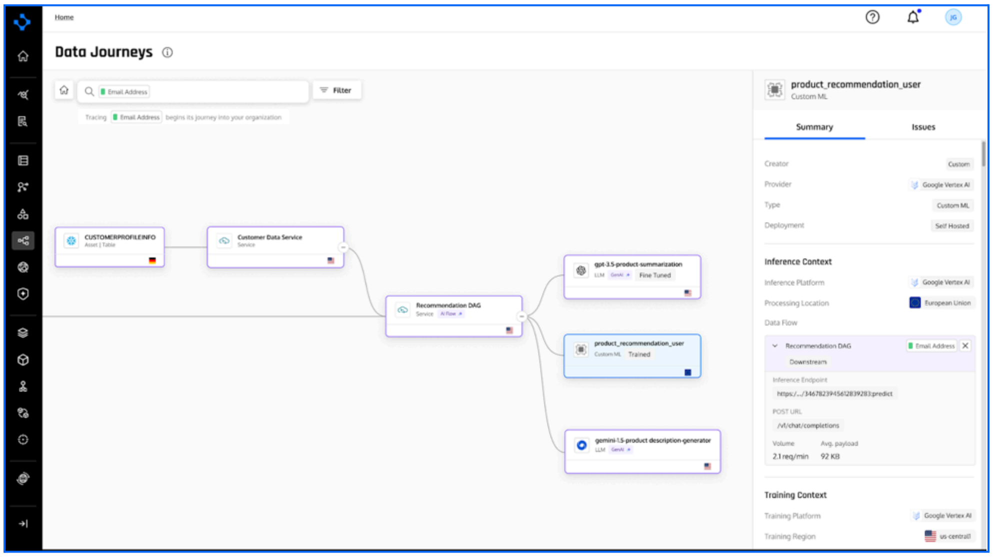

product_recommendation_user
Custom ML

Creator: Custom
Provider: Google Vertex AI
Type: Custom ML
Deployment: Self Hosted

Inference Context
Inference Platform: Google Vertex AI
Processing Location: European Union
Data Flow
Recommendation DAG — Email Address
Downstream
Inference Endpoint: https:/.../3467823945612839283:predict
POST URL: /v1/chat/completions
Volume: 2.1 req/min
Avg. payload: 92 KB

Training Context
Training Platform: Google Vertex AI
Training Region: us-central1

Fig 1. Data Journeys™ map

AI Visibility & Security Posture Management (AI-SPM)

Relyance AI delivers one of the **most comprehensive and differentiated AI-SPM offerings** in the market — unifying full AI asset visibility with deep security posture analysis across **both training and inference workflows**. Every AI asset in the path — models, datasets, features, code, pipelines, APIs, agents, and third-party integrations — is continuously discovered and mapped, giving SecOps and Privacy teams the **single, live source of truth** they've never had.

Key capabilities

- **Comprehensive AI Asset Discovery** – Automated, agentless discovery of AI applications, models, datasets, MLOps workflows, and third-party AI integrations — including unapproved shadow AI — across training and inference.
- **Sensitive Data Classification & Mapping** – AI-powered detection of regulated personal data (PII, PHI, PCI) and enterprise-specific proprietary fields across the entire AI lifecycle.
- **Context-Aware AI Data Lineage** – Runtime and training-time tracing of how sensitive data enters, flows through, and exits AI applications, including detection of unauthorized or non-compliant data flows.
- **Identity & Access Context** – Captures human and non-human identities (services, agents) with access to sensitive data to inform least-privilege reviews and posture scoring.
- **Shadow AI Detection** – Surfaces AI systems, models, and data flows operating outside approved governance processes, ensuring nothing is missed.

Top risks addressed

- **Sensitive data leakage** – Identifies and maps where regulated and proprietary data is used, enabling targeted controls to avoid fines, lawsuits, and reputational harm.
- **Shadow AI & unauthorized data flows** – Eliminates blind spots by flagging unapproved AI systems and data movements that violate policy or regulations.
- **Overexposed access paths** – Highlights identities and services with risky or unnecessary access to sensitive data, reducing insider and supply chain risk.
- **Continuous compliance support** – Maintains live traceability of sensitive data to meet EU AI Act, NIST AI RMF, and ISO 42001 requirements without manual, point-in-time audits.

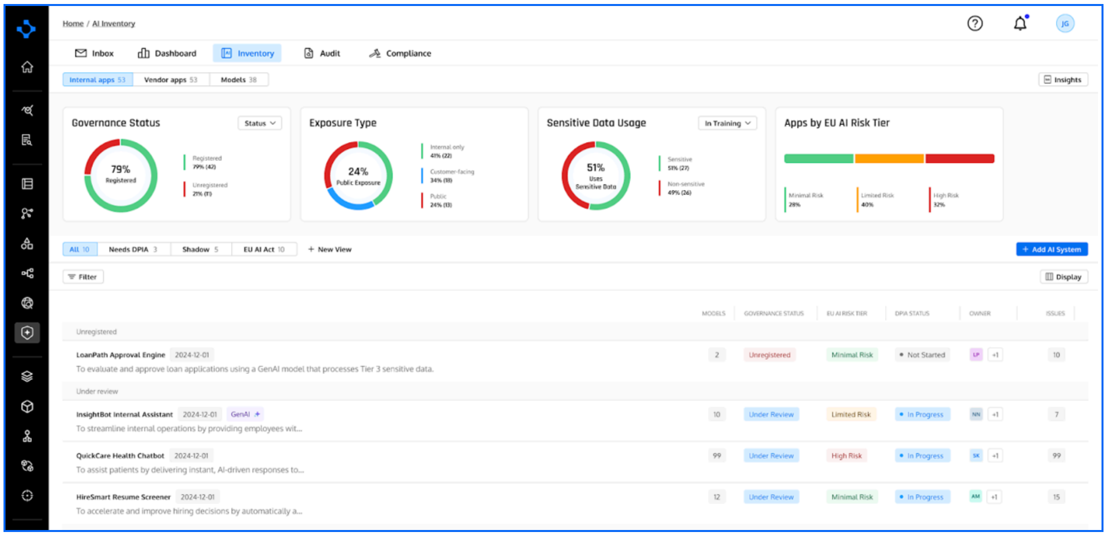

	MODELS	GOVERNANCE STATUS	EU AI RISK TIER	DPIA STATUS	OWNER	ISSUES
Unregistered						
LoanPath Approval Engine 2024-12-01 To evaluate and approve loan applications using a GenAI model that processes Tier 3 sensitive data.	2	Unregistered	Minimal Risk	● Not Started	LP +1	10
Under review						
InsightBot Internal Assistant 2024-12-01 GenAI To streamline internal operations by providing employees wit...	10	Under Review	Limited Risk	● In Progress	NN +1	7
QuickCare Health Chatbot 2024-12-01 To assist patients by delivering instant, AI-driven responses to...	99	Under Review	High Risk	● In Progress	SK +1	99
HireSmart Resume Screener 2024-12-01 To accelerate and improve hiring decisions by automatically a...	12	Under Review	Minimal Risk	● In Progress	AM +1	15

Fig 2. AI Inventory

AI Governance

Relyance AI brings **runtime AI governance** to the forefront — enabling Privacy, Risk, and Compliance teams to see exactly how AI applications handle sensitive data, how models evolve in production, and whether they remain aligned with policy over time. Governance shifts from static, point-in-time reviews to a **continuous, evidence-backed process** that complements upstream lifecycle controls.

Key capabilities

- **Runtime Policy Mapping** – Observes AI application behavior in real time and validates it against stored governance policies, regulatory requirements, and organizational risk thresholds.
- **Data Journey Traceability** – Maps how sensitive data moves through AI applications, capturing full lineage from input to output to support lawful, transparent, and accountable usage.
- **Model Lineage & Drift Tracking** – Monitors model changes, retraining events, and drift that can introduce new governance risks, tying each to associated data flows and policy impact.

Audit & Accountability Framework – Maintains a historical record of governance decisions, policy versions, and runtime deviations, with clear owner attribution for remediation and escalation.

Top risks addressed

- **Compliance & regulatory gaps** – Continuously validates runtime AI behavior against EU AI Act, NIST AI RMF, ISO/IEC 42001, and other obligations, closing the gap between policy intent and production reality.
- **Policy drift** – Detects when deployed AI systems operate outside approved parameters, reducing the risk of hidden compliance failures and costly post-incident remediation.
- **Untraceable data use** – Provides live mapping of data use, eliminating governance blind spots that slow incident response or regulator inquiries.
- **Model drift without governance oversight** – Flags unapproved or unmonitored model updates that can alter risk posture or regulatory compliance status.

Business value: Delivers continuous, provable governance across production AI applications, cutting policy validation time by up to 60%, reducing regulator inquiry turnaround from weeks to hours, and ensuring governance remains synchronized with the actual behavior of deployed systems.

3rd-Party Vendor AI Risk

SaaS vendors with AI capabilities turned on often operate without visibility or runtime validation, creating risk exposure that traditional vendor management processes can't catch. Relyance AI automates the discovery, classification, and continuous risk assessment of these vendors — replacing static, questionnaire-based checks with live operational context.

Key capabilities

- **Automated Vendor AI Discovery** – Finds all SaaS and third-party AI providers with features enabled, including shadow and unsanctioned use.
- **Continuous AI Risk Assessments** – Replaces manual spreadsheets and point-in-time questionnaires with automated, ongoing posture checks.
- **DPIA & Regulatory Support** – Enriches Data Protection Impact Assessments with live data flows and AI-specific risk scoring.
- **Contract & AI Addendum Validation** – Monitors vendor runtime behavior against contractual AI usage clauses to flag non-compliance.

- **Sensitive Data Flow Tracking** – Detects personal and regulated data shared with vendor AI systems during prompts, training, or inference.

Top risks addressed

- **Shadow AI use (incl. vendors)** – Eliminates blind spots by surfacing unsanctioned vendor AI capabilities.
- **Manual AI risk assessments** – Automates continuous evaluation, removing dependency on manual reviews.
- **No audit readiness for AI regulations** – Maintains live context to demonstrate vendor compliance with AI-specific obligations.
- **Inability to prove lawful data use in AI** – Maps and validates data flows to ensure vendors process data in line with legal and contractual requirements.

Business value: Cuts AI vendor risk assessment time by up to 70%, gives security and privacy teams a unified, always-current view of vendor AI risk posture.

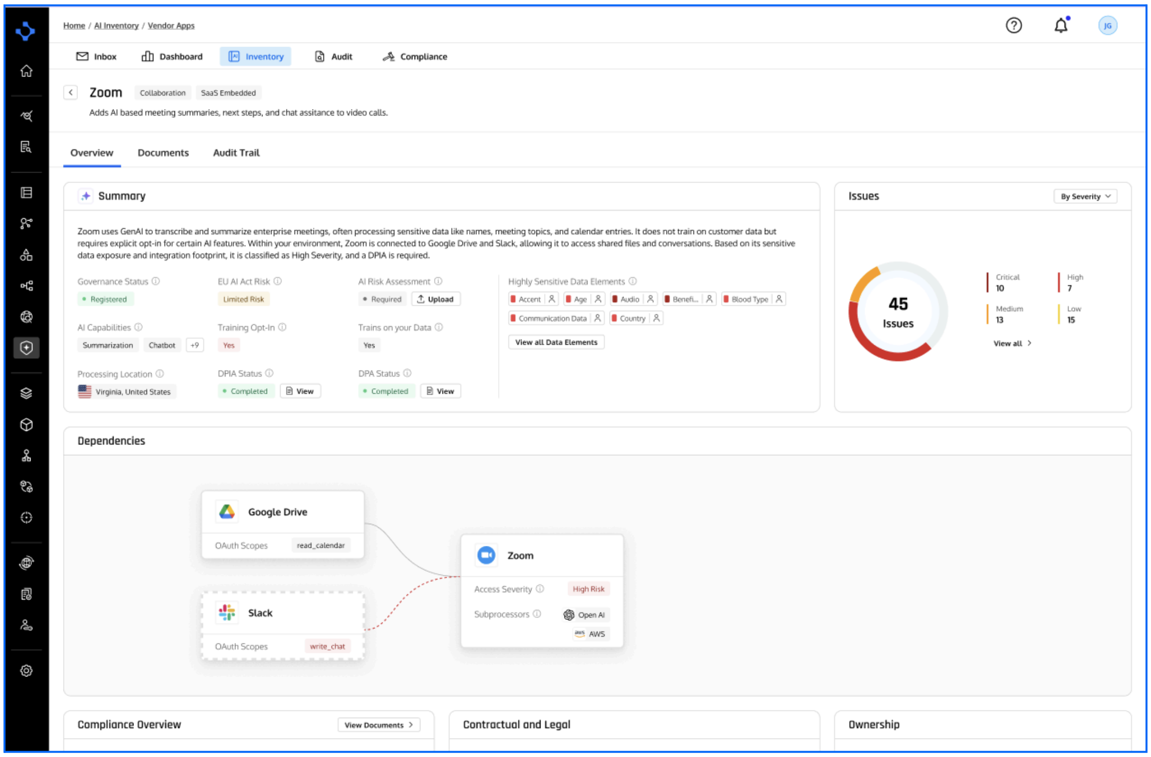

Zoom
Collaboration
SaaS Embedded

Adds AI based meeting summaries, next steps, and chat assitance to video calls.

Summary

Zoom uses GenAI to transcribe and summarize enterprise meetings, often processing sensitive data like names, meeting topics, and calendar entries. It does not train on customer data but requires explicit opt-in for certain AI features. Within your environment, Zoom is connected to Google Drive and Slack, allowing it to access shared files and conversations. Based on its sensitive data exposure and integration footprint, it is classified as High Severity, and a DPIA is required.

Governance Status
Registered

EU AI Act Risk
Limited Risk

AI Risk Assessment
Required
Upload

AI Capabilities
Summarization
Chatbot
+9

Training Opt-In
Yes

Trains on your Data
Yes

Processing Location
Virginia, United States

DPIA Status
Completed
View

DPA Status
Completed
View

Highly Sensitive Data Elements
Accent
Age
Audio
Benefi...
Blood Type
Communication Data
Country
View all Data Elements

Fig 3. AI Inventory / Vendor Risk Analysis: Zoom

Agentless, AI-Native Architecture & Deployment Flexibility

A key operational advantage of the Relyance AI platform lies in its architecture: it is both ***agentless*** and ***AI-native***. "Agentless" means that unlike traditional security tools, Relyance AI does not require installing any heavy sensors or software agents on your servers to monitor data flows. Instead, it integrates via APIs, reads from logs, source code repositories, cloud configuration, and other existing data sources to gather its intelligence. This approach dramatically reduces deployment friction and time-to-value – most organizations can connect Relyance AI to their environment and start seeing insights within minutes, not weeks. The platform was designed to work *with* modern development workflows (cloud-native and DevOps practices) rather than slow them down. Because there are no agents consuming resources or needing constant updates, the solution scales easily across cloud and hybrid environments. Security teams appreciate that an agentless design means one less potential point of failure or vulnerability on their systems, and business stakeholders appreciate the rapid, hassle-free rollout.

Being AI-native means that Relyance AI's core is built on advanced machine learning and natural language processing techniques that continuously learn from your data landscape. The platform's intelligence engine (TrustiQ™) uses AI to automatically classify data (including unstructured data and code) and to infer the context of data processing activities. For example, it can distinguish personal health information from log data, or recognize that a piece of data is an email address being used for marketing vs. for authentication, based on how it flows through the code. This intelligent classification and context enrichment is far beyond what any manual rule-based system could achieve at scale. It enables more accurate policy enforcement and fewer false positives – the system "understands" your data environment in a holistic way and adapts as it changes. Moreover, as new patterns (like a new type of cloud service or an emerging AI library) appear, the AI models help Relyance quickly accommodate them, making the platform future-proof by design.

In terms of deployment model, Relyance AI offers flexibility to meet enterprise needs.

1. **Relyance Full SaaS:** Cloud-hosted SaaS deployment
2. **Relyance AI InHost™:** a self-hosted option within your VPC.
3. **Relyance DirectConnect:** private link between your internal network and Relyance platform

Relyance Full SaaS Deployment

Relyance AI's Full SaaS deployment mode is designed to deliver an enterprise-grade solution. In this mode, all data processing occurs within a multi-tenant SaaS environment hosted in Google Cloud Platform (GCP), where each customer is logically isolated through dedicated IAM roles and resource boundaries.

Data from the customer's environment—such as metadata, asset schemas, policies & contracts, logs, and, if configured, sampled data—is securely transmitted to the Relyance environment via encrypted channels (TLS 1.3). Optional structured and unstructured samples, when enabled, are retained only for the duration of the scan and are never persisted. These inputs are processed by ephemeral, autoscaling AI services that extract findings and insights and display them in the Relyance SaaS UI. This architecture enables a secure, low-maintenance deployment model aligned with Zero Trust principles, ensuring data remains protected in transit and at rest, with customer configurations over what is shared. It offers a balance of scalability, automation, and enterprise security readiness—ideal for CISOs seeking a compliant SaaS solution without operational overhead.

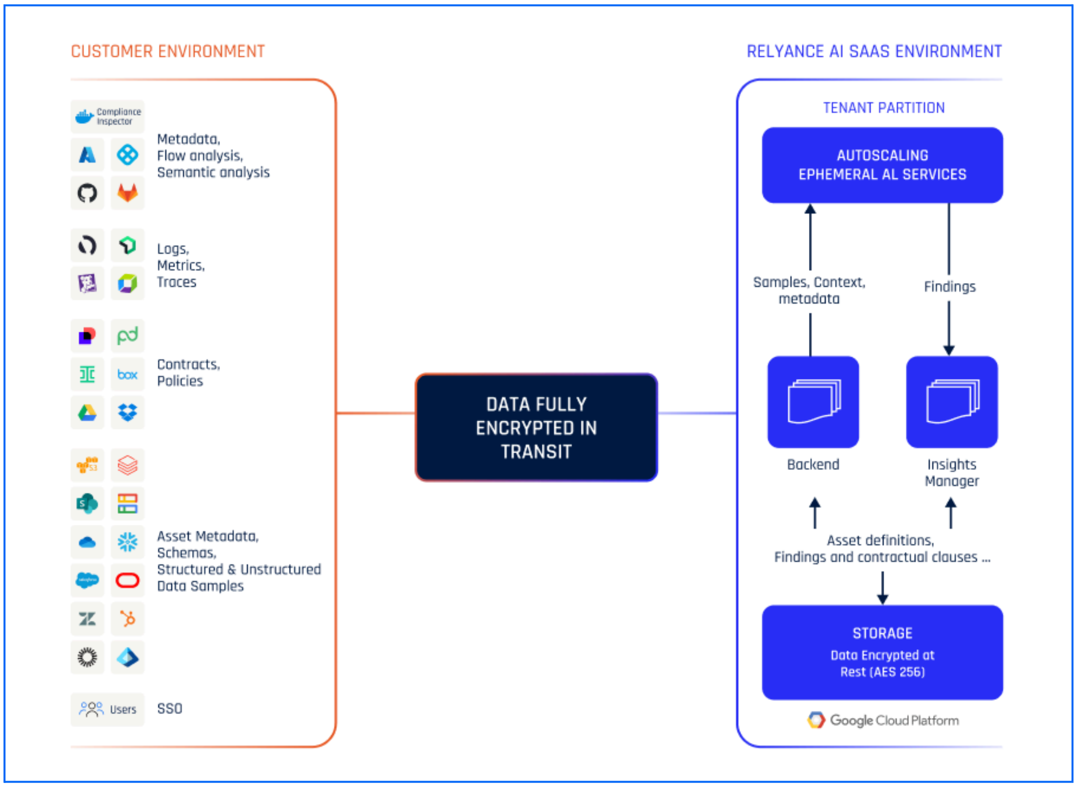
CUSTOMER ENVIRONMENT
Compliance Inspector
Metadata, Flow analysis, Semantic analysis
Logs, Metrics, Traces
Contracts, Policies
box
Asset Metadata, Schemas, Structured & Unstructured Data Samples
Users
SSO
DATA FULLY ENCRYPTED IN TRANSIT
RELYANCE AI SAAS ENVIRONMENT
TENANT PARTITION
AUTOSCALING EPHEMERAL AL SERVICES
Samples, Context, metadata
Findings
Backend
Insights Manager
Asset definitions, Findings and contractual clauses ...
STORAGE
Data Encrypted at Rest (AES 256)
Google Cloud Platform

Fig 4. Full SaaS Architecture Diagram

Relyance InHost™ Deployment

With InHost, the entire platform runs within your own private cloud (VPC), ensuring that sensitive telemetry never leaves your environment. This is especially important for industries with strict data sovereignty and security requirements – you get all the benefits of Relyance AI's intelligence while keeping data under your own roof. Whether SaaS or self-managed, the platform is built with enterprise-grade security and privacy in mind (e.g. strong encryption, role-based access controls, audit logs), as evidenced by Relyance AI's transparent Trust Center and compliance with standards like SOC 2. In all cases, deployment does not require ripping and replacing any existing tools – Relyance AI augments your stack by integrating with developer tools, CI/CD pipelines, cloud accounts, and ticketing systems. This **architectural agility** means faster ROI and the ability to continuously adapt as your data ecosystem and regulatory landscape evolve.

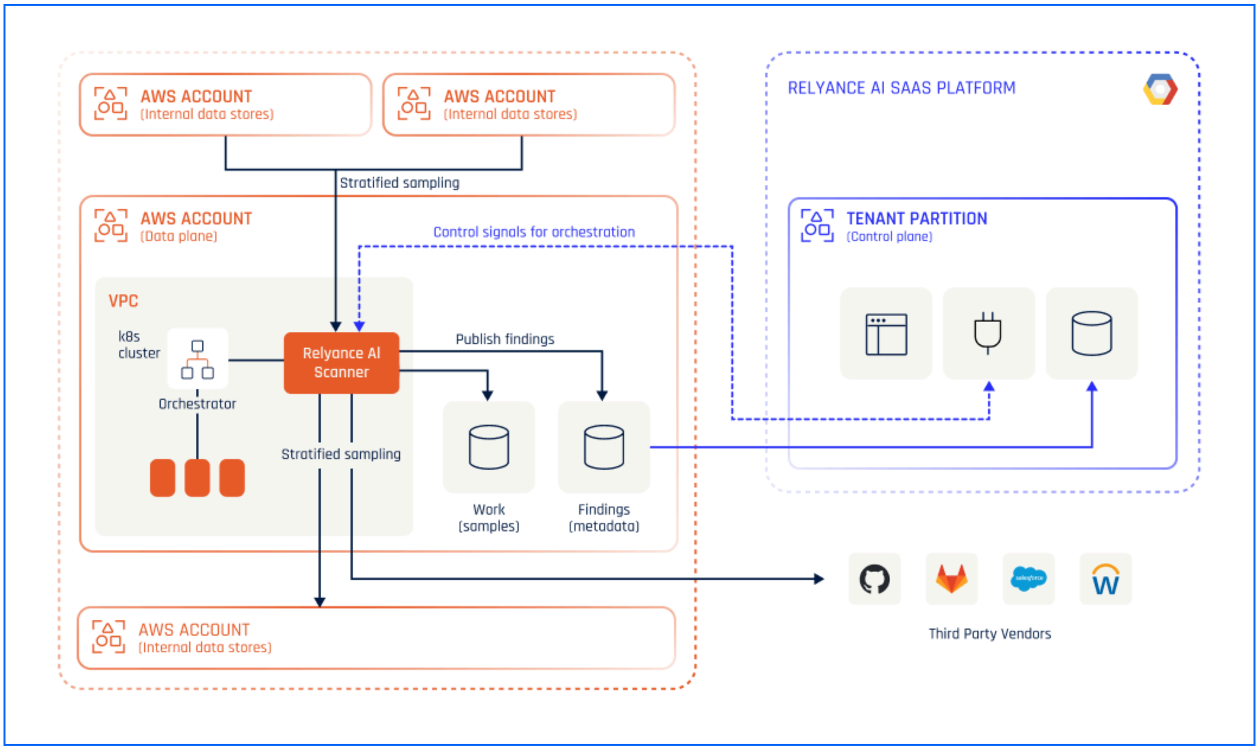

Fig 5. InHost™ Architecture Diagram

Relyance DirectConnect Deployment

DirectConnect is a private, fully managed deployment model that enables secure connectivity between your internal network and Relyance platform - without requiring you to install or manage any infrastructure. By establishing a dedicated Google Cloud project for each customer, DirectConnect ensures strict tenant isolation and complete control over network access. All communication occurs via a secure private link, such as VPN or VPC peering, allowing you to expose only selected internal IPs and services. No public internet exposure or endpoint configuration is necessary, and you maintain full data privacy while we handle the underlying processing infrastructure.

This architecture enables all scans, analytics, and data queries to run entirely within your dedicated cloud environment, offloading all compute and bandwidth costs to us. With no need for agents, VMs, or containers on your end, setup is fast—typically completed within hours. DirectConnect is ideal for enterprises needing to securely scan internal databases or integrate with private cloud data warehouses like Snowflake or Redshift. It's a frictionless, compliant option that offers high security, rapid onboarding, and zero operational burden for your teams.

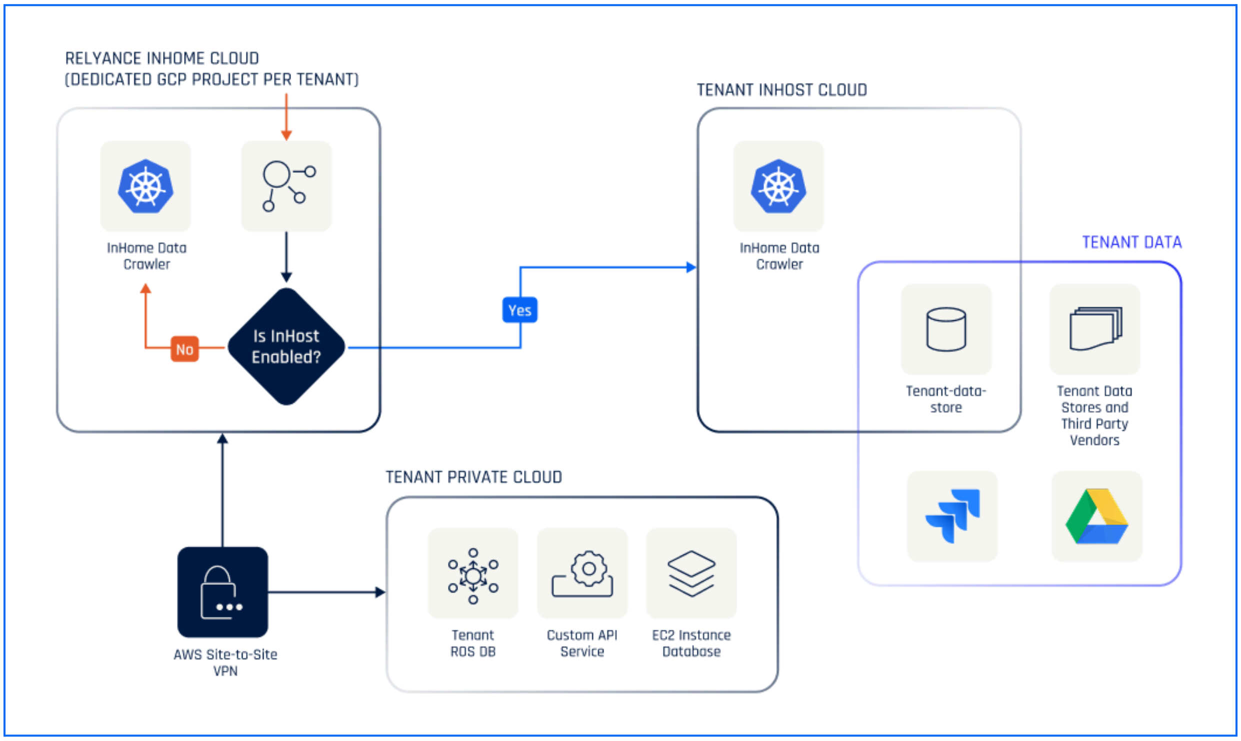

Fig 6. DirectConnect Architecture Diagram

Operational, Business, and Strategic Benefits

Relyance AI's unified **AI Security & Governance Platform** delivers operational, business, and strategic advantages that align directly with the way enterprises are adopting. securing – and governing – AI.

A single source of truth for AI assets and posture.
Instead of security, privacy, and engineering teams maintaining separate spreadsheets, dashboards, or assumptions about where AI is used and how data flows, Relyance AI provides a continuously updated AI asset and posture inventory. Covering the full training and inference lifecycle – from source code to deployed models, datasets, workflows, and vendor AI integrations – this live operational view ensures every stakeholder is working from the same, verified information. That alignment is critical when responding to regulator questions, incident investigations, or high-stakes board reviews.

Evidence-driven assurance for AI compliance and governance.
Whether facing the EU AI Act, NIST AI RMF, ISO/IEC 42001, or internal privacy and acceptable use policies, Relyance AI equips CISOs and governance leaders with concrete, runtime-driven evidence of compliance. Instead of point-in-time assessments, the platform provides continuous validation of AI behavior against obligations, mapping sensitive data use, model lineage, and runtime conditions in real time. This transforms compliance from a manual, episodic effort into an ongoing, automated process that can withstand regulatory and customer scrutiny.

Operational efficiency through automation.
By automating AI asset discovery, sensitive data mapping, and posture monitoring, Relyance AI eliminates the tedious and error-prone manual processes that consume security and privacy teams. Tasks like finding all AI-enabled SaaS vendors, tracing PII through training pipelines, or validating model changes against policy become push-button actions. This reduces assessment cycles from weeks to hours, lowers the cost of audits, and frees up skilled staff to focus on high-value remediation and risk reduction – rather than chasing down undocumented data flows.

Proactive risk reduction.
Continuous monitoring means issues – such as shadow AI tools, policy drift, or unapproved model updates – are detected and addressed before they cause regulatory violations or reputational harm. This early detection reduces the likelihood and cost of breaches, fines, and contract disputes, while strengthening trust with customers and partners.

Strategic enablement of innovation.
With unified oversight across internal AI Applications and vendor AI services, Relyance AI allows security, privacy, and compliance leaders to say "yes" to new AI initiatives with confidence. Teams can adopt cutting-edge AI capabilities, onboard vendors faster, and scale automation without sacrificing control. This positions governance not as a blocker, but as a competitive advantage – enabling the organization to move quickly and responsibly in a market where AI speed is often the differentiator.

As one industry analyst put it, enterprises now require "real-time, unified oversight" across privacy, AI, and security. Relyance AI delivers exactly that – combining the breadth of AI visibility with the depth of continuous posture management to help organizations innovate at speed, without losing control.

About Relyance.AI

Relyance.AI is the leading Dynamic DSPM platform built for the AI era. It was founded by security and privacy experts to meet the demands of modern enterprises facing explosive data growth, AI adoption, and regulatory scrutiny.

With customers across financial services, retail, SaaS, and healthcare, Relyance.AI is redefining how organizations discover, secure, and govern data at the speed of innovation.

Trusted by:

logitech Bolt ZUORA PLAID coinbase Notion samsara
ClickUp Grafana Canva dayforce dialpad NAVAN Fivetran