Dynamic DSPM Whitepaper

Secure Data and AI From Code-to-Cloud

Dynamic Data Security Posture Management (DSPM) for AI Era

relyanceai

Table of Content

Executive Summary

Relyance AI's Dynamic Data Security Posture Management (DSPM) platform represents a fundamental leap forward in how organizations can understand and control their data in an AI-driven world. By mapping every data journey across code, infrastructure, and AI systems, it provides a level of transparency and context that was previously unattainable. For CISOs and Security Operations leaders, this translates directly into reduced risk, improved compliance posture, and enhanced ability to enable business innovation. The platform's unique combination of real-time data lineage, predictive risk prevention, and automated enforcement closes the gaps that adversaries and auditors alike would otherwise exploit. It brings privacy, security, and governance together on a single pane of glass, reflecting the reality that these concerns are interconnected in modern enterprises.

Armed with Relyance AI, security leaders can turn unknown data liabilities into well-managed assets. They can walk into the boardroom or regulator's office with confidence, backed by hard evidence of their organization's data governance in action. Equally important, they can partner with engineering and product teams to accelerate the safe adoption of AI and cloud technologies, fueling innovation rather than impeding it. In a landscape where data is the new oil – and oversight of that data is the new mandate – Relyance AI offers an AI-native, future-ready platform to ensure that every data journey is a safe, compliant, and trustworthy one. By embracing this dynamic approach to data security posture management, forward-thinking CISOs can not only protect their enterprises from today's risks but also confidently navigate the opportunities of tomorrow.

The Need for Dynamic Data Visibility

As organizations race to embed generative AI and accelerate cloud adoption, CIOs and security operations leaders face a harsh reality: their sensitive data is constantly in motion—and largely ungoverned. Data no longer remains confined to structured databases or storage buckets—it flows dynamically through a new generation of data assets, developer pipelines, cloud-native services, and increasingly, generative AI systems. This shift has exposed critical shortcomings in traditional Data Security Posture Management (DSPM) tools platforms, which were architected for a simpler, data-at-rest world. Despite advances in adaptive scanning and cloud-native

connectors, these solutions fundamentally lack the depth and contextual visibility needed for today's code-to-cloud security requirements.

1. Static Visibility Is Failing Data Security

The first major limitation lies in the static nature of visibility. Traditional DSPM tools focus on reactive scanning of data stores—often failing to detect sensitive data as it enters pipelines or is transformed within code and SaaS workflows. In contrast, security and privacy by design demands a proactive posture, one that starts at the source (such as code commits or API calls) and continues through to data consumption and sharing. Without deep insight into operational workflows, these tools lack context and can only provide partial and often misleading information that security teams rely on to do their jobs.

2. No Unified Model for Security, Privacy, and Compliance

Secondly, the growing convergence of security and privacy operations demands a unified approach—something legacy DSPM platforms cannot offer. As cloud, AI, and regulatory complexity increase, security teams must collaborate with GRC, legal, and privacy stakeholders to enforce cohesive data governance. Yet most DSPM products lack a shared data model, risk engine, or evidence-as-a-service that spans both security and compliance domains. The result is siloed insights, duplicated controls, and fragmented remediation workflows that undermine accountability and trust. In addition, traditional DSPM tools lack visibility into data provenance, which is critical in building world-class data governance programs.

3. AI Changes Everything: Static Lineage Leaves Organizations Blind to Real-Time Risks

The third and most pressing gap emerges with the adoption of AI. Traditional tools offer basic data lineage capabilities, but these are static—built for legacy reporting needs, not real-time risk decisions. AI introduces novel data flows, such as prompt inputs, model inference outputs, and training pipelines that ingest or expose sensitive information. Without dynamic lineage and

context-aware controls, organizations are blind to how their data is being used—or misused—by internal and third-party AI systems. This lack of visibility also jeopardizes emerging compliance standards for AI accountability and model transparency.

Moreover, legacy DSPM solutions struggle to support explainability and auditability of how data drives outcomes in AI systems. In a regulatory environment increasingly focused on "show your work" mandates—whether it's GDPR Article 30, the EU AI Act, or ISO 42001—static lineage fails to answer the crucial questions: How did this data get here? Who touched it? Was it used according to policy? Dynamic, graph-based lineage and policy-aware guardrails are now prerequisites, not niceties.

In short, traditional DSPM tools were built for risks—anchored to data-at-rest, compliance checkboxing, and passive discovery. But the enterprise is shifting to a real-time, data-in-motion model—where AI, privacy, and cloud-native velocity collide. To meet this moment, organizations need a Dynamic DSPM platform that delivers continuous observability, integrated enforcement, and explainable data governance from code to cloud to AI. That is the foundational gap Relyance.AI is uniquely designed to fill.

This white paper explains how Relyance AI's platform uniquely maps, secures, and governs dynamic data flows across code, infrastructure, and AI models.

Introduction to Dynamic DSPM with Relyance

Relyance.AI introduces a new category: Dynamic DSPM – a platform purpose-built to trace, understand, and control how data flows through modern digital ecosystems. By mapping every "data journey" from code to cloud, and even to contractual obligations, Relyance AI's platform provides continuous, context-aware visibility and control over data flows that were previously opaque.

In practice, data is perpetually in motion – flowing through source code, microservices, pipelines, SaaS apps, and machine learning models. Without real-time visibility into these flows, security leaders are left with blind spots that jeopardize compliance and trust. In the AI-driven era, organizations require a new approach that illuminates how data is collected, used, and transformed across complex ecosystems – in real time and in context.

> "*CISOs need to see the entire data journey, not just point-in-time snapshots... you need to know how data moves, who touches it, and what happens at every step of the journey. That's exactly what Relyance AI delivers.*"
>
> Cybersecurity veteran and CISO, Caleb Sima

Relyance AI's Data Journeys take data lineage to the next level by capturing the full end-to-end story of each data element in your system. Instead of isolated, column-level lineage within a single database, Data Journeys provide an ecosystem-wide view: it traces data from its original point of collection, through every transformation, service, API call, and third-party touchpoint, all the way to storage or deletion.

By starting at the source (code) and ingesting signals from live systems, Relyance AI can illuminate data flows with precision and context. This continuous lineage map allows security teams to see every hop and transformation in context – from an initial data collection in an application, through microservice pipelines and databases, into model training or third-party sharing – along with the purpose and identity associated at each step.

Relyance platform gives a visual view of the entire journey of data across applications, services, infrastructures, third parties – even revealing the underlying purpose and context of processing (the "why" of data usage). Additionally it highlights the risks and violations through the data flow and at every node with actionable remediation to mitigate the risks.

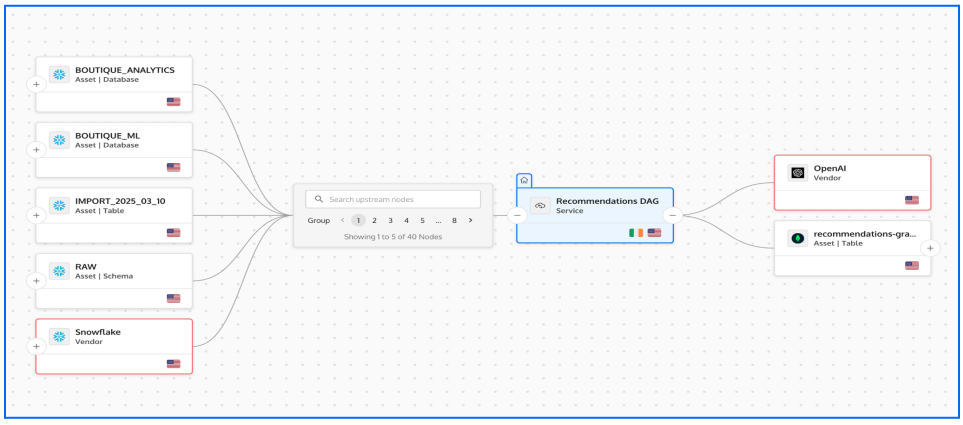

Fig 1. Relyance Data Journeys™

Critically, Data Journeys capture context. The platform doesn't just show that data moved from point A to B; it also records *why* – linking each data flow to the business process, application, or AI algorithm that generated it. This context-aware approach directly addresses compliance questions and regulatory requirements around purpose limitation, lawful processing, and AI transparency. Stakeholders and auditors can finally get visual, audit-ready proof of how personal data is used, combined, and transformed, with explanations that make sense in business terms.

> "*Relyance AI's Data Journeys™ provides a dynamic, graph-based view of how sensitive data travels and transforms through code, cloud infrastructure, third-party APIs, and AI models. Such a complete data flow visualization eliminates blind spots that traditional tools miss and forms the foundation for proactive security and compliance.* "

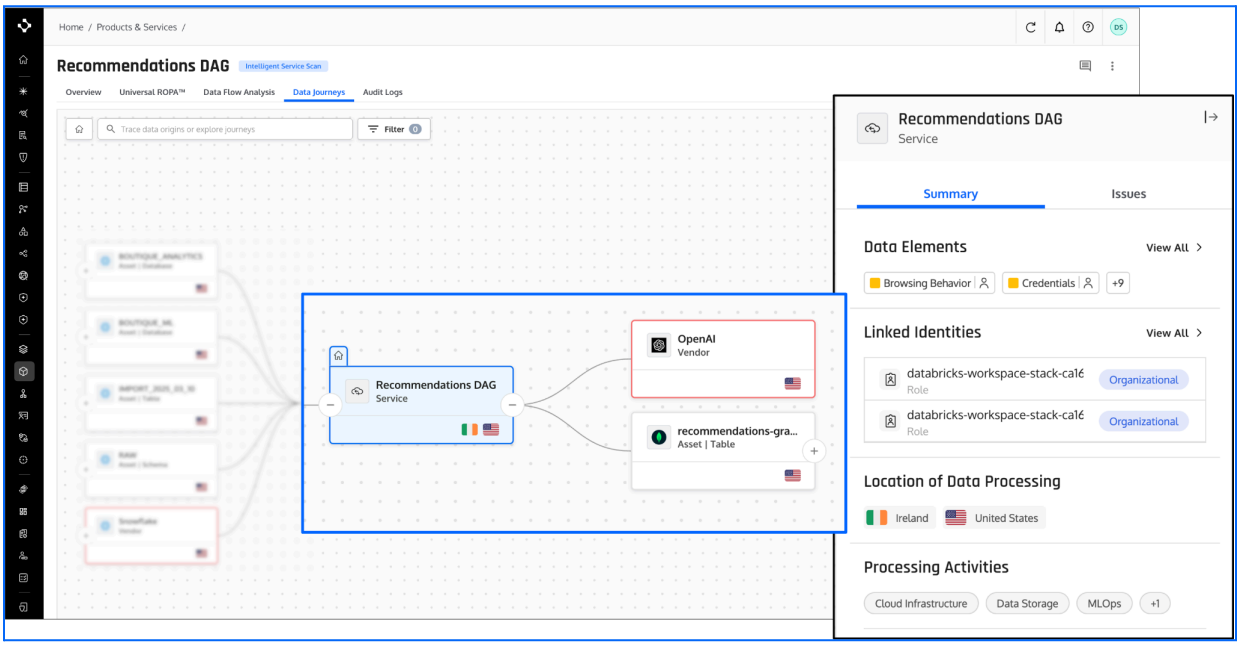

Fig 2. Data Journeys™ with details for each hop of data

In short, Relyance AI turns what was once a murky "black box" of data activity into an open book of explainable data usage. By providing live data flow lineage instead of after-the-fact reports, Data Journeys™ empower organizations to identify risky or unauthorized data use *before* it causes an incident or violation.

Unified Platform For Data Security, Privacy, and AI

Relyance AI distinguishes itself by unifying three traditionally siloed domains – data security posture, privacy compliance, and AI governance – into a single intelligent platform. At its core is an AI-powered Trust Intelligence (TrustiQ™) engine that continuously ingests signals across your environment (via APIs and metadata, no agents required) and updates a comprehensive data graph. This unified data graph underpins all the platform's capabilities, enabling cross-functional visibility and coordinated controls that benefit security and compliance teams simultaneously.

Below, we delve into the platform's key capabilities, describing how they deliver strategic value and operational efficiency:

Continuous Data Flow Mapping and Risk Visibility

At the heart of the platform is the ability to discover and map 100% of data flows across code, cloud, and third-party systems in real time. Relyance AI automatically scans source code, monitors runtime behavior, and connects to cloud/SaaS logs to build a live model of how sensitive data travels through your architecture. This means security teams get 360° visibility without blind spots – every API call, every microservice output, every database write involving sensitive information is charted on the Data Journey™ map.

Compared to legacy DSPM tools that might list a database containing PII, Relyance AI shows *where that PII came from, which services or AI models touched it, and where it went next*. By illuminating entire data journeys from code to runtime to cloud, the platform exposes hidden risk propagation paths that piecemeal tools overlook.

This deep visibility directly translates into better risk management. Security teams can spot unauthorized or unexpected data flows (for example, an API sending user data to an unapproved external service) and receive alerts about anomalous behavior with the context and lineage necessary to prioritize and address the most pressing risks. The platform's intelligence engine uses pattern recognition and policy rules to flag these issues in context. As a result, CISOs and their teams move from reactive firefighting to proactive oversight – they can address misrouted data or access violations as they emerge, not weeks after the fact. The blind spot

elimination provided by Relyance AI gives leaders confidence that there are no "unknown unknowns" in how their sensitive data is. In practical terms, this means only contextualized and actionable risks will be surfaced, leading to fewer security incidents, faster incident response, and the ability to demonstrate control over data flows to boards and regulators.

AI Governance and Responsible AI Enablement

In an era when AI systems (like large language models and custom ML algorithms) are becoming ubiquitous, governing AI's use of data has become a top concern for security and compliance leaders. Relyance AI extends data security posture management into the realm of AI, providing an "AI-aware DSPM." The platform automatically discovers where AI and ML models are integrated into your business – both the first-party models your engineers build and the third-party AI services (such as OpenAI APIs or SaaS with embedded AI) that your teams may be leveraging. Once discovered, Relyance AI monitors those AI systems to map real-time sensitive data flows into and out of them.

This is crucial because AI pipelines often ingest large volumes of sensitive data (for training or inference) and produce outputs that could inadvertently expose that data. With Relyance, a CISO can see, for example, that a customer dataset was used to train Model X, or that an engineer's prompt to an LLM included confidential source code – all in the unified data journey view.

Armed with this visibility, organizations can implement responsible AI practices with confidence. Relyance AI lets teams set guardrails specific to AI contexts: for instance, preventing an LLM from logging or returning personal information in its responses, or ensuring that data used for AI training is properly anonymized. The platform's shift-left AI security approach helps catch risky AI usage at the source – it will warn if sensitive data is about to be sent to an external AI API or if an AI model is drawing on data in ways that violate policy. It also continuously audits AI model behavior for compliance with emerging regulations and ethical guidelines (e.g., checking that outputs don't inadvertently include prohibited personal data).

As regulations like the EU AI Act loom, having continuous AI lifecycle monitoring and documented data lineage for models becomes a strategic necessity, not just a nice-to-have. Relyance AI builds that auditability and control into AI initiatives from day one. This capability enables CISOs to support fast AI adoption safely – they can green-light innovative AI projects knowing that any misuse of data by the AI will be quickly detected and corrected. In effect,

Relyance AI provides the toolset for security teams to become AI enablers rather than blockers, striking the balance between rapid innovation and rigorous governance.

Holistic Privacy Compliance and Data Governance

One of the strategic advantages of Relyance AI's platform is how it bridges the gap between security operations and privacy compliance. Global privacy regulations (like GDPR, CCPA, and beyond) demand detailed knowledge of how personal data is used, shared, and protected. Relyance AI was built with these requirements in mind, embedding compliance checks and governance metadata into each data journey. Every data flow mapped by the system isn't just a technical path; it is enriched with information like the type of data (e.g. health data, financial data), the purpose of processing, consent status, and applicable legal obligations.

This means privacy officers can leverage the same data journey map to ensure that usage aligns with what users have consented to and what regulations allow. The platform can automatically detect, for example, if a piece of EU customer data is being used in a way that violates GDPR purpose limitation, or if an AI model is trained on data beyond the scope of user consent, and then raise an alert or trigger a mitigation.

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.

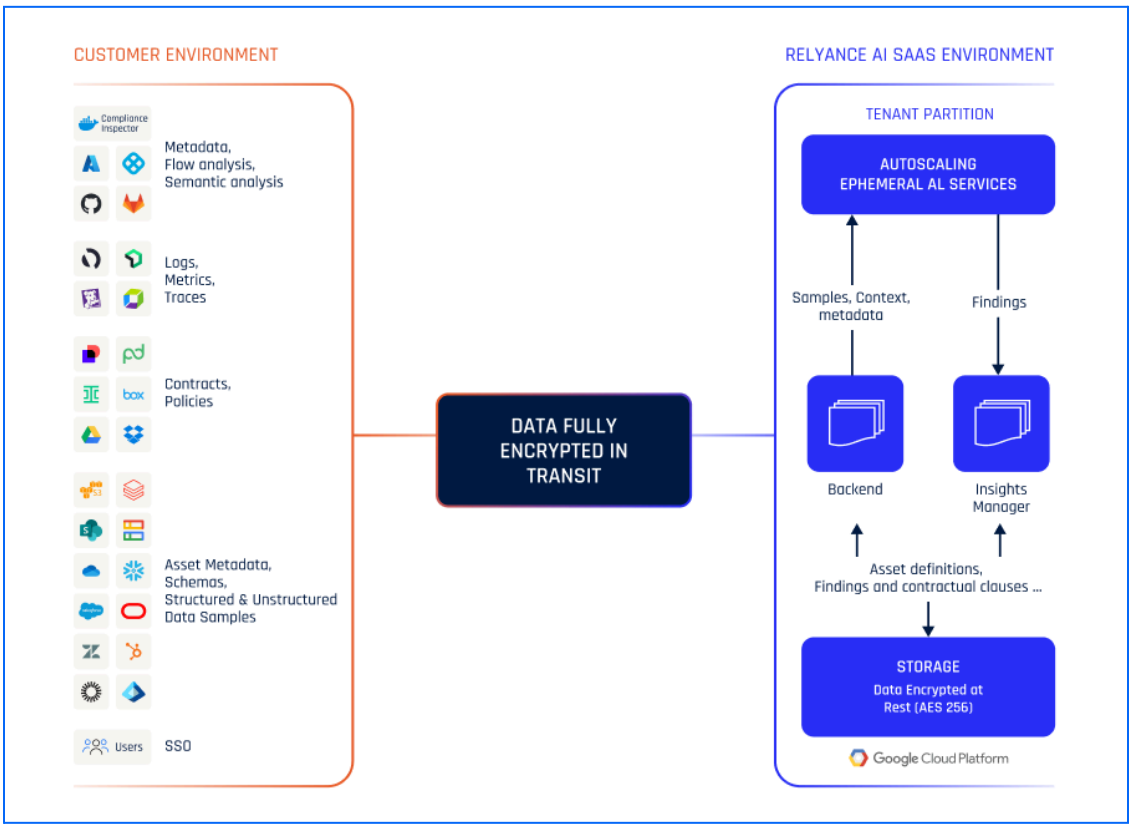

Fig 4. Full SaaS Architecture Diagram

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.

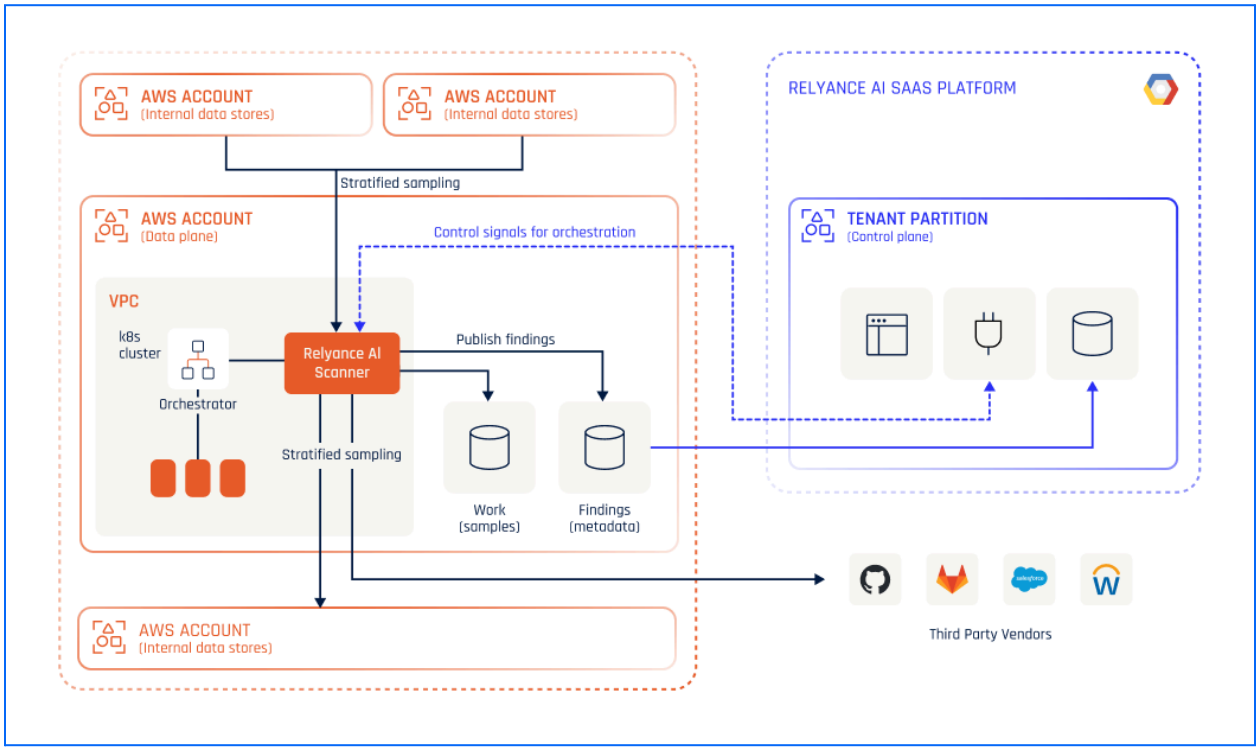

Fig 5.InHost™ Architecture Diagram

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.

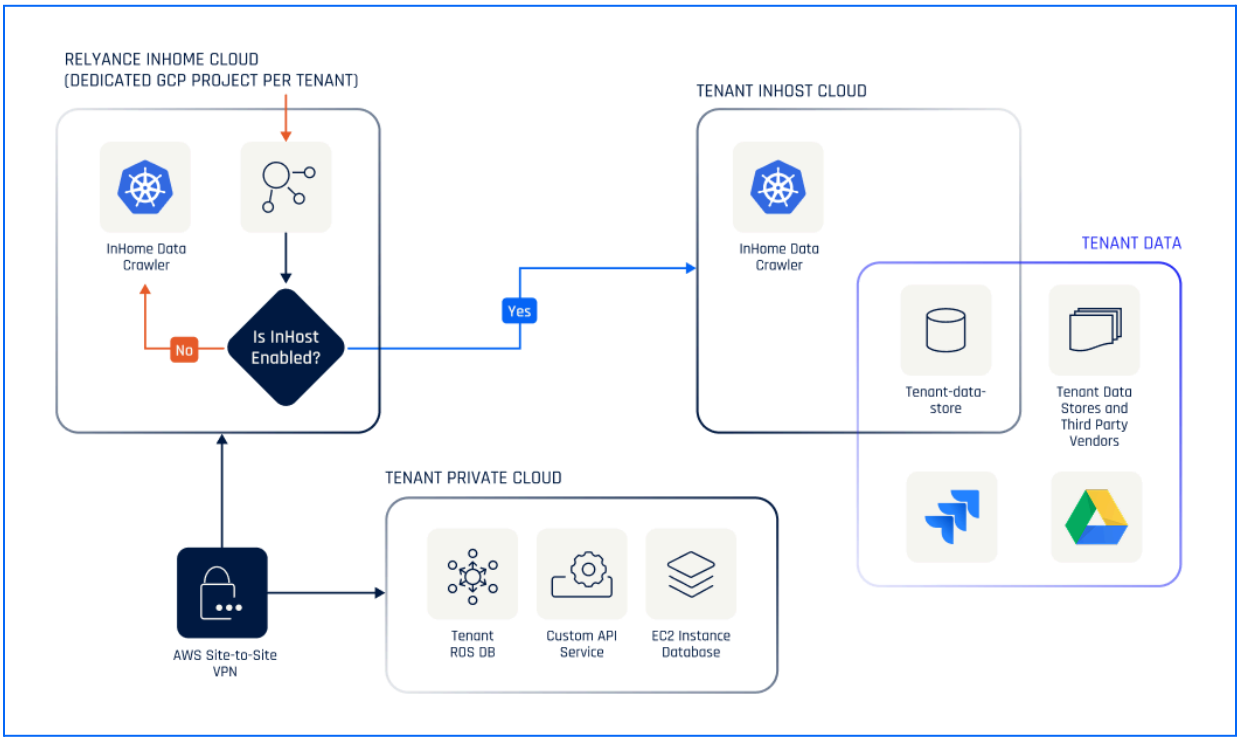

Fig 6.DirectConnect Architecture Diagram

Operational, Business and Strategic Benefits

Relyance AI's Dynamic DSPM platform delivers a range of strategic and operational benefits that align closely with corporate data security priorities.

Firstly, it provides a single source of truth for data usage across the organization. Instead of different teams maintaining separate spreadsheets, dashboards, or assumptions about how data flows, everyone from Security Engineering to Privacy to DevOps can rely on one unified data graph. This not only improves collaboration (breaking down silos between security, privacy, and engineering) but also means that when an incident or question arises, teams are working off the same factual information. Such alignment is crucial when responding to breaches or compliance inquiries under tight deadlines.

Secondly, the platform enables evidence-based assurance. In board meetings or audits, CISOs can move beyond high-level promises and actually *demonstrate* how the organization's sensitive data is controlled. Need to prove compliance with a new AI regulation or with GDPR's Article 30 requirements? With a few clicks, Relyance AI can produce audit-ready documentation showing data journeys, applied controls, and policy compliance status. This capability turns compliance from a costly, anxiety-ridden effort into a more automated, continuous reporting function. It also provides peace of mind to executive leadership that the company can withstand scrutiny from regulators or customers on questions of data use and AI ethics. In essence, Relyance AI operationalizes the principle of "trust by design", giving companies a way to build and verify trust at every step of their data lifecycle.

Operationally, one of the biggest benefits is efficiency and time savings. By automating data discovery, mapping, and monitoring, the platform relieves security operations and privacy managers from tedious manual work. Tasks like chasing down the owner of a dataset, figuring out which APIs touch a certain type of PII, or compiling quarterly compliance reports are handled by the platform's continuous intelligence. This allows skilled staff to focus on higher-value activities such as analyzing risks and improving security posture, rather than playing data detective. The reduction in manual effort also translates to cost savings – organizations can do more with the same headcount, and avoid the fines or legal costs that often come from undetected issues. Moreover, faster risk mitigation (thanks to real-time alerts and guardrails) means lower likelihood of data breaches, which averts the substantial costs associated with incident response and damage control. In a world where data breaches and privacy missteps can severely damage brand reputation, the investment in preventative capability pays for itself many times over.

Strategically, Relyance AI positions the security team as an enabler of innovation. Rather than the proverbial Department of "No," armed with this platform, CISOs can confidently say "Yes" to digital transformation initiatives – whether it's adopting a new AI tool, migrating to a multi-cloud setup, or partnering with a data-intensive third party – because they know they have the oversight to manage the risk. This shift is invaluable: it means security and compliance are no longer roadblocks but competitive advantages. Companies that harness data and AI quickly (while keeping them secure) will outpace those that move slowly out of fear. Relyance AI helps enterprises find that sweet spot where they move fast *and* build things securely. As one investor put it, organizations at this moment demand "unified oversight" across privacy, AI, and security, and Relyance AI is solving this critical challenge in a way that allows innovation to thrive responsibly.

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: