anyinsight
Powered by HEARTBOT

HEARTBOT · AnyInsight

ENTERPRISE WHITEPAPER · FOR SECURITY & COMPLIANCE LEADERS

AI Governance & Compliance

A Technical Overview of the AnyInsight Trustworthy AI Workforce Platform

Productive. Governed. Secured.

VERSION	1.7 · August 2026
PUBLISHED BY	HEARTBOT AI Inc.
ONLINE	www.heartbot.com · www.anyinsight.ai

Confidential — © HEARTBOT AI Inc., 2026

Contents

Executive Summary 3

01 Company Brief 4

02 The Enterprise AI Governance Imperative 5

03 AnyInsight Architecture 7

04 The AI Workforce Platform 8

05 The AI Firewall Concept 9

06 Governance Model 10

07 Identity and Workspace Isolation 11

08 Data Flow 11

09 Policy Engine 12

10 Audit Trail 14

11 Certifications 15

12 Compliance Mapping 16

Conclusion 17

© HEARTBOT AI Inc., 2026. All Rights Reserved.

Executive Summary

Generative AI is reshaping every corner of the enterprise at unprecedented speed, promising exponential gains in productivity, new business models, and customer experiences never seen before. The same wave, however, introduces enormous uncertainty: data leakage, loss of intellectual property, an expanding regulatory burden, and an unmanageable "Wild West" of ungoverned AI applications. For a Chief Information Security Officer, the strategic question is no longer whether to adopt AI to raise productivity, but how to adopt it without ceding control of the organization's data, identity, and audit posture — a question made sharper by agentic AI systems that access data across multiple application systems, repeatedly and autonomously, in order to complete the tasks and workflows assigned to them.

At the same time, agentic AI is scaling faster than the guardrails around it. Organizations expect moderate-or-greater use of AI agents to rise from roughly a quarter today to nearly three-quarters within two years. Governance is no longer a compliance afterthought; it is the capability that determines whether AI ever moves from pilot to production.

That pressure is now legal as well as commercial. The EU AI Act became enforceable on 2 August 2026: the transparency obligations of Article 50, the penalty framework, and the enforcement powers of the EU AI Office and national market surveillance authorities all apply from that date. The Digital Omnibus on AI — Regulation (EU) 2026/1744, in force since 27 July 2026 — deferred the high-risk obligations to December 2027 and August 2028, but it moved neither the general date of application nor the duties that reach the widest set of organizations. Section 2 sets out the resulting calendar.

AnyInsight, developed by HEARTBOT AI, is a Trustworthy Secure AI Workforce delivered as a service. It gives enterprises a single governed portal in which employees collaborate with more than 350 built-in AI agents across 20+ large language models (LLMs) and 500+ application connectors, while a cloud-native governance layer and a zero-trust GenAI firewall sit inline between every user, workflow, and model. Productivity and control are engineered as one system rather than traded against each other.

> **THE CORE PROPOSITION**
>
> *Productivity and control are engineered as one system rather than traded against each other — a governed portal above, a zero-trust GenAI firewall inline, and an immutable audit trail beneath every interaction.*

This whitepaper describes how AnyInsight delivers governance and security alongside everyday productivity rather than at its expense. It walks through the five-layer governance architecture, the AI Firewall and policy engine, the three-tier identity and workspace isolation model, the end-to-end data flow, and the immutable audit trail — and closes with a mapping of these capabilities to the EU AI Act, GDPR, and ISO/IEC 27001 obligations that enterprise buyers are being held to. AnyInsight holds ISO/IEC 27001 certification and runs on cloud infrastructure certified to ISO/IEC 27001, 27017, and 27018, on a platform designed to meet the requirements of the EU AI Act and GDPR.

© HEARTBOT AI Inc., 2026. All Rights Reserved.

01 Company Brief

HEARTBOT AI Inc. is a Taiwan-based enterprise AI company founded in 2023 and headquartered at the National Yang Ming Chiao Tung University (NYCU) Innovation Incubation Center in Hsinchu — the region from which many of the world's most advanced technology companies, including TSMC and MediaTek, originated. The company builds trustworthy, secure AI infrastructure for the enterprise under the principle "Think Trusted · AnyInsight," and delivers it commercially through its flagship platform, AnyInsight.ai.

Vision

A world where every business can adopt AI with complete trust.

Mission

Empower businesses to deploy AI securely by embedding governance, compliance, and security into the core of AI transformation — not bolting them on afterward.

HEARTBOT positions AnyInsight as a data aggregator and protector: a governed layer that lets an organization harness public, paid, authorized, industry-specific, and confidential knowledge together through agentic AI, without surrendering control of any of it.

1.1 Company at a Glance

Attribute	Detail
Legal entity	HEARTBOT AI Inc.
Founded	March 2023
Core proposition	Productive. Governed. Secured.
Certifications	ISO/IEC 27001 (certified) AWS-inherited ISO/IEC 27001 / 27017 / 27018 Designed for EU AI Act and GDPR compliance
Recognition	2025 Best Solution AI Award, Taiwan Artificial Intelligence Association
Ecosystem	NVIDIA Inception · AWS Activate · Google for Startups NYCU IAPS program
Achievements	Spotlight partner for the NCHC/TAIDE sovereign-AI programme Joint-defense with network security launched at Early Stage Expo, RSA 2026, San Francisco

© HEARTBOT AI Inc., 2026. All Rights Reserved.

02 The Enterprise AI Governance Imperative

Before describing the platform, it is worth being precise about the problem it solves. Four structural forces currently hold enterprises back from scaling GenAI, and each is a security and governance concern first.

Security gaps

The black-box operation of public models makes it impossible to guarantee that enterprise-sensitive data will not be leaked or absorbed into model training. Without unified identity and privilege control, every AI interaction becomes a potential breach — and as agentic AI moves data between applications autonomously, it creates a new class of "shadow agent" risk.

Management chaos

Different teams independently procure or build AI tools, producing hard-to-track "shadow AI." With no centralized monitoring, logging, or auditing, compliance obligations simply cannot be met.

Development bottlenecks

Teams burn time on environment setup and model maintenance rather than innovation, while fragmented models, tools, and data sources drive up integration cost and delay deployment.

Spiraling cost

The inability to evaluate and manage AI usage leads to budget overruns, and duplicated investments across teams leave resources under-utilized.

Regulation is turning these operational risks into legal ones. The EU AI Act, adopted in 2024 and applicable in staged phases, introduces risk-based classification, high-risk governance mandates, and material fines for non-compliance — and since 2 August 2026 it is enforceable, with the penalty framework operative and supervisory powers held by the EU AI Office and national market surveillance authorities. It creates a control gap that conventional network security was never designed to close: AI prompt manipulation, unsafe model outputs, data leaks through AI channels, and the absence of an AI audit trail.

Closing that gap requires governance to be treated as a strategic capability. Organizations that build strong governance frameworks now will be positioned to scale AI quickly and safely; those that treat it as a checkbox exercise risk being unable to move AI from pilot to production at all.

2.1 Most Concerning Issues While Deploying AI

Concern	Share of organizations
Data privacy and/or security	73%
Legal, IP, or regulatory compliance	50%
Governance capabilities and oversight	46%
Model quality, consistency, and explainability	46%
Workforce impact	30%

© HEARTBOT AI Inc., 2026. All Rights Reserved.

Source: Deloitte, State of AI in the Enterprise: The Untapped Edge, January 2026 (N = 3,199).

2.2 Where the EU AI Act Now Stands

That calendar changed shortly before it took effect. On 27 July 2026 the Digital Omnibus on AI — Regulation (EU) 2026/1744 — entered into force and amended the AI Act days before its main application date. It did not move the general date of application, which remained 2 August 2026, but it deferred the high-risk obligations that had dominated most compliance roadmaps. Two things are therefore true at once: enforcement has begun, and the high-risk regime has not. The table below sets out which obligations are live and which are still ahead.

Date	Obligation	Status
2 February 2025	Prohibited practices (Article 5) and the AI literacy duty (Article 4)	**Applicable**
2 August 2025	Obligations for providers of general-purpose AI models (Chapter V)	**Applicable**
2 August 2026	Transparency obligations (Article 50); penalty framework (Articles 99–101); enforcement powers of the EU AI Office and national market surveillance authorities	**Enforcement begins**
2 December 2026	Machine-readable marking of synthetic content by generative systems already on the EU market before 2 August 2026; new prohibitions on non-consensual intimate imagery and AI-generated CSAM	**Transitional**
2 August 2027	Full compliance for GPAI models placed on the market before 2 August 2025; deadline for Member States to operate a national AI regulatory sandbox	**Pending**
2 December 2027	Obligations for stand-alone high-risk systems (Annex III) — deferred from 2 August 2026	**Deferred**
2 August 2028	High-risk AI embedded in regulated products (Annex I) — deferred from 2 August 2027	**Deferred**

Source: Regulation (EU) 2024/1689 as amended by Regulation (EU) 2026/1744 (Official Journal, 24 July 2026). Dates are indicative and do not constitute legal advice.

Two points deserve emphasis for a security function. First, Article 50 reaches far more organizations than the high-risk regime does: any enterprise operating a branded chatbot or assistant, using emotion recognition or biometric categorisation, or publishing AI-generated content must now disclose that fact, and the ceiling for breaches of the conformity obligations is EUR 15 million or 3% of worldwide annual turnover, rising to EUR 35 million or 7% for prohibited practices. Second, an organization that customises, fine-tunes, or rebrands an AI system can be requalified from deployer to provider and inherit the full provider regime with it — a transition that in practice happens quietly, inside business units, without anyone recording that it has occurred.

© HEARTBOT AI Inc., 2026. All Rights Reserved.

> **DEFERRAL IS TIME TO BUILD, NOT RELIEF**
>
> *The inventory required for December 2027 is the same inventory a market surveillance authority may request this year — and the technical documentation, log retention, and human-oversight evidence that Annex III will demand cannot be retrofitted to systems that were never instrumented to produce them.*

Both questions can only be answered from a complete and current inventory of the AI systems in use, with an owner and a purpose recorded against each. That is the practical case for governing AI through a single platform rather than across a sprawl of independently procured tools.

03 AnyInsight Architecture

AnyInsight is organized as a five-layer enterprise AI governance stack. The design intent is that every layer a user touches sits above a layer that governs it, so productivity and control are never in separate systems. The upper layers deliver the AI workforce; the middle layer enforces governance; the lower layers provide orchestration and a hardened, multi-tenant foundation. The stack is delivered as multi-tenant, multi-AI, multi-region, and multi-cloud.

Layer	Purpose	Primary console / delivery mode
LAYER 1 **AI Workforce**	Productivity, seamless AI workflow adoption, and data integration	AI Workforce Portal (user-facing)
LAYER 2 **Operation**	Identity and access management, asset management, and observability	Management Console (administrator-facing)
LAYER 3 **Governance**	Risk management, compliance, and the AI Firewall	Governance Control Plane (security and compliance-facing)
LAYER 4 **LLM / Orchestration**	Multi-AI architecture for data retention, service orchestration, and alignment across LLMs	Built in, under the zero-trust security framework
LAYER 5 **Cloud Infrastructure**	Multi-tenant AI compute optimization and system protection across regions	Built in, under the zero-trust security framework

Functionally, AnyInsight positions itself as the gatekeeper between users, workflows, and AI resources. User context input, the model's intrinsic knowledge, and reference data — whether from internal applications or from external public, paid, and authorized sources — all converge through AI agents. A cloud-native governance layer and zero-trust GenAI firewall mediate that convergence, so that industry-specific and confidential knowledge can be brought to bear on a task without leaving the governed perimeter.

© HEARTBOT AI Inc., 2026. All Rights Reserved.

The remainder of this paper examines the layers that matter most to a security function: the AI workforce (Layer 1), the firewall and governance model (Layer 3), identity and workspace isolation (Layer 2), the data flow across all layers, the policy engine, and the audit trail.

04 The AI Workforce Platform

The AI Workforce Layer is delivered through the AI Workforce Portal, the interface users reach when they first log in to AnyInsight. It is where employees work with large language models, built-in agents, custom agents, and automated task agents to produce results. Its purpose is to build sustainable human–AI collaborative productivity for everyone in the organization — across productivity, marketing, R&D, legal, HR, finance, supply chain, and ESG functions — while keeping every interaction inside the governed platform. It is built on three components.

4.1 MAIA — Multi-AI Architecture

The Multi-AI Architecture (MAIA) lets an organization flexibly leverage the strengths of 20+ large language models — including GPT, Gemini, Claude, Grok, Mistral, Perplexity, and Taiwan's sovereign TAIDE model — matching each task to the most suitable model, or to an agent equipped for the required data access. MAIA supports three collaboration modes that improve output reliability, reduce hallucination, and strengthen cross-domain reasoning, all of which matter for high-impact decisions and in-depth analysis.

Parallel Mode

Two LLMs — or one agent and one LLM — execute the same task independently, increasing decision diversity and guarding against single-model bias and hallucination.

Integrative Mode

An LLM or a custom agent collaborates with a second model to dynamically integrate external enterprise data and knowledge into the response.

Critique Mode

A second model reviews and challenges the conclusions of the first, reducing error and hallucination while surfacing alternative views for the user to weigh and decide.

4.2 AI Enabler

The AI Enabler is the set of AnyInsight features that lowers the barrier to safe, high-quality AI use.

- **Prompt Wizard.** Turns structured user input into validated, high-quality prompts.
- **Prompt Templates.** A library of 100+ templates that makes high-frequency tasks repeatable.
- **Storefront.** More than 350 ready-to-use agents spanning IT, marketing, finance, legal, ESG, and other domains, with support for customization and for preserving team-specific expertise. Users authorized by an Admin, or assigned the Contributor role, can publish additional custom agents organization-wide, to a specific workspace, or for personal use only.
- **Dual support system.** Combines a guided "Get-started Agent" with a built-in "AI Helper," keeping users supported inside the platform and reducing the information-leakage risk that arises when employees seek help from external services.
- **Multi-lingual support.** Presents the entire AnyInsight interface in 32 languages — including English, Spanish, French, Portuguese, German, Italian, Dutch, Traditional Chinese, Korean, Japanese, Thai, and

© HEARTBOT AI Inc., 2026. All Rights Reserved.

other EU regional languages — so that AI terminology is accessible to every employee and adoption is not gated by language.

4.3 Agent-Driven Knowledge Platform

As an agent-driven knowledge platform, AnyInsight turns agents into a durable capability and a source of sustainable productivity, rather than a series of one-off conversations. Context engineering combines personal know-how, private enterprise knowledge, real-time web search, and MAIA multi-model collaboration to move beyond the knowledge limits of any single general-purpose model. Enterprise-grade connectors link to 500+ external systems, including Gmail, Outlook, Jira, Slack, and Salesforce.

Once identity has been verified, agents may use designated connectors to execute assigned workflows automatically — transforming an agent into a cross-department "digital employee" that carries out multi-step tasks repeatedly and autonomously. This same capability is precisely what makes governance non-negotiable: an agent able to move data between applications must be monitored and controlled, and it is that control which allows AnyInsight to serve as a sustainable knowledge platform that compounds productivity over time.

05 The AI Firewall Concept

A conventional network firewall inspects packets at the network boundary. It has no visibility into the semantics of an AI prompt, the sensitivity of the data inside it, or the intent of an autonomous agent. The AnyInsight built-in GenAI Firewall is a zero-trust control that operates precisely where enterprises need it: positioned inline between users, workflows, and AI models, it acts as the gatekeeper and treats every AI interaction as untrusted until policy has evaluated it.

The firewall performs its work by observing AI traffic before it leaves the governed perimeter. When a user or agent submits input to a model, or when a model or connector returns content, the request is inspected in real time against the organization's policies. Sensitive information is automatically identified and masked, non-compliant input is intercepted, and disallowed access patterns are blocked. This design opens the "black box" of AI interaction at its source and ensures that intellectual property — or any other material the firewall's policy rules designate — does not leave the organization through an AI channel.

5.1 Four Firewall Functions

AI content filtering

Inspects user prompts, uploaded data, and content generated by models or agents for sensitive or disallowed keywords, statements, and concepts, using AI reasoning-based rules.

Prompt protection

Guards against prompt manipulation by threat actors and against unsafe inputs at the AnyInsight boundary.

AI activity logging

Records every interaction — timestamp, actor, model, action, and result — to produce an audit trail that reflects both user behavior and end-to-end data movement.

Policy enforcement

Applies the organization's governance policies consistently across users, agents, and models, covering inbound and outbound data throughout AnyInsight.

© HEARTBOT AI Inc., 2026. All Rights Reserved.

Strategically, the AI firewall extends to joint defense with network security, supports EU AI Act compliance, and defends AI-driven business data. For managed service providers and firewall vendors, it is also the mechanism by which an existing network-security practice can extend into joint defense — protecting the AI attack surface, meeting AI compliance obligations, and establishing AI visibility. Sections 9 and 10 describe the policy engine that drives the firewall and the audit trail it produces.

06 Governance Model

The Governance Layer establishes the highest standards of trustworthy AI and auditability. It is delivered through the Governance Control Plane — the surface a CISO or compliance owner works from — and rests on three pillars: the Data Vault, violation monitoring, and security policy enforcement. Together they realize the EU AI Act's core principles of transparency, accountability, security, and risk management in operational form.

6.1 Data Vault

The Data Vault retains all inbound and outbound interaction data — files uploaded by users, content generated by any model or agent, and content produced through a connector — automatically and securely, with traceability and security flags attached, following privacy-by-design principles. A defining design choice is that retention does not depend on compliance status: data is preserved in full whether or not it violated a policy, on the principle that an audit is only as trustworthy as the completeness of its evidence. The result is an immutable, tamper-evident evidence chain available for retrospective review, internal audit, compliance checks, and incident investigation.

6.2 Violation Monitoring

The platform automatically tracks and tabulates policy violations over rolling 24-hour and 7-day windows, surfacing the top violators and the most common violation types so that risk concentrations are visible rather than buried in logs. A proactive alarm mechanism notifies administrators the moment a violation occurs and links directly to the Moderation Record for human review and resolution. This is real-time tracking coupled with proactive alerting, so that governance operates as a live control loop rather than a periodic report.

6.3 Security Policy Enforcement

Enforcement combines a configurable governance framework with custom inspection, allowing an organization to codify its own corporate policies and apply them uniformly across every model and agent. Because enforcement, monitoring, and the Data Vault share the same policy definitions, a rule authored once is applied consistently at the point of interaction, recorded in the audit trail, and surfaced in monitoring — with no gaps between what is intended, what is enforced, and what is auditable.

07 Identity and Workspace Isolation

Identity is the foundation of the whole governance model: a policy is only meaningful if the platform knows precisely who — or which agent — is acting, and what they are permitted to touch. AnyInsight's Operation Layer implements identity and access management (IAM) as a three-tier hierarchy with least-privilege controls throughout.

© HEARTBOT AI Inc., 2026. All Rights Reserved.

Org / Workspace / User hierarchy

Authority is separated across three tiers: the organization, the workspace (typically a department), and the user (an individual or user group). An Org Admin can create multiple workspaces along the lines of the organization's real structure, reducing the risk of lateral movement between departments.

Resource segmentation

From a central resource hub, administrators assign specific models, agents, templates, and connectors to specific workspaces. A workspace dedicated to the sales team receives only the AI resources appropriate to that team, rather than a uniform distribution that would expose sales data to those with no need to access it.

Least-privilege roles

Within a workspace, users may be granted Workspace Owner, Admin, Contributor, Auditor, or User roles. Each role is scoped to least privilege by default, preventing unauthorized access to critical functions.

THE WALLED WORKSPACE

Where a large organization's subsidiaries or departments must not share data — between R&D and legal, for instance — walled workspaces implement zero-trust, least-privilege access that prevents employees from laterally misusing AI resources and eliminates internal shadow-AI exposure at its root.

Because AI resources are managed as inventory and access is scoped per workspace, the organization always knows which models and connectors are in use, by whom, and under what authority.

08 Data Flow

Understanding how data moves through AnyInsight is the clearest way to see where control is applied. AnyInsight positions itself as a data aggregator, for productivity, and a data protector, for governance and security. Every stage of the flow passes through a governed checkpoint.

01 **Authenticated entry**

A user or agent acts within a specific workspace under an assigned role. Identity and privilege are resolved first, so the request carries a known actor and a known permission scope from the outset.

02 **Context assembly**

MAIA assembles the working context — user input, the selected model's intrinsic knowledge, user know-how, private enterprise knowledge, and reference data drawn from internal and external applications through connectors. Only resources assigned to that workspace are eligible.

03 **Inline inspection (egress)**

Before input leaves the perimeter for a model or connector, the GenAI Firewall observes it and performs real-time risk detection. Sensitive data is identified and masked, and violating input is intercepted at source.

© HEARTBOT AI Inc., 2026. All Rights Reserved.

04 **Model orchestration**

The governed request is routed to one or more models via MAIA, optionally using parallel, integrative, or critique modes to improve reliability.

05 **Inline inspection (ingress)**

Returned content — from a model or a connector — is inspected on the way back, so that unsafe outputs and leaked reference data are caught before reaching the user.

06 **Vaulting and logging**

Every interaction, compliant or not, is written to the Data Vault with traceability and security flags, and captured across the system, dialogue, automation, shared conversation, and connector logs.

The critical property of this flow is that connector traffic — the path by which an autonomous agent could otherwise exfiltrate data between data-driven applications — is inspected and fully audited like every other channel. The “shadow agent” risk described in Section 2 is closed not by trusting the agent, but by placing the same governed checkpoint on the agent’s data movements as on a human’s.

09 Policy Engine

The policy engine is what makes the firewall configurable to a specific organization rather than a fixed filter. Org administrators define governance policies against distinct risk scenarios, and each policy contains multiple rules that can be tuned independently.

9.1 Governance Scenarios and Policy Categories

Policies are organized around three core risk-governance scenarios — Keyword Reasoning, Privacy Protection, and Safeguard — which map in the Governance Control Plane onto a set of concrete, activatable policy categories.

Category	Representative policies	Governance purpose
Content Filter	Reasoning and learning Semantic topic classification	Detect and reason over sensitive keywords in prompts and outputs
Sensitive Information — PII	Contact information, personal identifiers, network identifiers Government IDs, vehicle identifiers	Personal-data protection and PII leakage prevention
Sensitive Information — Financial	Banking details Credit cards	Protect financial data and payment identifiers

© HEARTBOT AI Inc., 2026. All Rights Reserved.

Category	Representative policies	Governance purpose
Sensitive Information — Credentials	Cloud credentials	Prevent secrets and access keys from leaving the perimeter
Sensitive Information — Health	Health records	Protect health-related identifiers
Toxic Content	Safeguard baseline	Baseline content-safety governance
Access Control	Browser, device type, operating system Geolocation, IP address / CIDR, Time windows	Match and constrain sign-in and access conditions to protect AnyInsight service

Access Control policies are at their most powerful when combined with allowlists: restricting sign-in to approved networks, locations, devices, and time windows provides the strongest available protection against unauthorized access to AnyInsight.

9.2 Flexible Rule Configuration

Each policy allows multiple firewall rules, and every rule supports an enabled or disabled state, a priority order, an activated-rule count, and exception conditions expressed as allowlists or blocklists. This lets an organization dynamically tune control strength — tightening enforcement where risk is high, and granting scoped exceptions where a legitimate workflow requires them, without disabling the policy altogether. Access-control policies extend the same model to the sign-in surface, which is frequently the first surface a threat actor probes, constraining access by network, geography, device, operating system, and browser to strengthen the platform's self-protection against external exfiltration. AnyInsight further applies AI reasoning and learning to broaden interception scope over time, so that coverage improves rather than remaining static.

9.3 Enforcement Actions

When a rule matches, the engine applies a graded action — Alert, Block, or Mask — and records the outcome.

Mask

Allows a workflow to continue safely by redacting only the sensitive element, while signalling the issue to the user.

Block

Stops a non-compliant interaction outright, requiring the user to revise the prompt before proceeding.

Alert

Flags an event for review while preserving the interaction in the audit trail.

Every action is written to the Moderation Record with its severity and raised to the administrator, so that the response is proportionate to the assessed risk.

© HEARTBOT AI Inc., 2026. All Rights Reserved.

10 Audit Trail

Auditability is treated as a first-class capability rather than a by-product of logging. The Operation Layer provides end-to-end observability, and the Governance Layer turns that observability into a defensible record for auditors.

10.1 Five Core Log Types

Log type	What it records
System Log	All system events, user operations, and environment changes, supporting problem tracing and resource optimization.
Dialogue Log	Every AI conversation in full — the model or agent used, input and output content, context, and any violation alerts.
Automation Log	Each execution of an automated task: the trigger, the agent and connectors involved, the application systems touched, and the outcome. This historical record also provides the basis for analyzing data trends over time.
Shared Conversation Log	Conversations shared between colleagues. Sharing accelerates communication, but can open a security gap if confidential information reaches unauthorized recipients; AnyInsight restricts sharing to internal members and traces the data exchanged.
Connector Log	A full audit of every operation an agent performs when accessing external application systems. It is the principal control against shadow-agent risk, recording all data flowing between models or agents and external systems where exfiltration could otherwise occur.

All five logs are available for enterprises to download for backup, anomaly analysis, or internal model training, and their availability and retention are designed to satisfy the audit and traceability requirements of ISO/IEC 27001, GDPR, and the EU AI Act.

10.2 Moderation Record

The Moderation Record makes every issue traceable in a single view. Each entry captures the violator context (workspace, user, and email), the policy and violation rule that fired, the severity, the action taken — Alert, Block, or Mask — and the offending content itself, alongside a resolution status. Administrators can filter across rolling windows of 1 day, 7 days, 30 days, 3 months, or 6 months to investigate patterns and close out incidents. Because the alarm mechanism links directly into this record, detection, investigation, and resolution live in one continuous workflow.

10.3 Visualization and Observability

A configurable dashboard provides real-time monitoring of AI usage, user behavior, environment resource use, and security status, consolidating anomalies, rule-violation alerts, and security indicators. Combined with the downloadable logs and the immutable Data Vault, this gives an organization the complete, transparent

© HEARTBOT AI Inc., 2026. All Rights Reserved.

evidence chain it needs to face strict external audits in finance, healthcare, or multinational manufacturing with confidence.

11 Certifications

For an enterprise buyer, certification is not a badge; it is the mechanism that converts a vendor's assertions into independently audited evidence. An accredited third party has examined the controls, which means the burden of proving they exist no longer rests with the customer's security team. That distinction matters in three practical ways.

Shortens due diligence

A certified information security management system answers, in a single artifact, questions that would otherwise require months of questionnaires and evidence requests.

Narrows assessment scope

Inherited cloud certifications mean controls already covered by the platform's ISO/IEC 27017 and 27018 attestations do not need to be re-verified.

Clears procurement gates

In regulated sectors certification increasingly functions as a gate: without it, a vendor may not be eligible for evaluation at all, regardless of product merit.

Certification therefore accelerates adoption rather than merely documenting it. AnyInsight's compliance posture combines certifications the platform itself holds with certifications inherited from its underlying cloud infrastructure, and with regulatory frameworks the platform is built to satisfy.

Standard / framework	Scope	Status
ISO/IEC 27001	Information security management system	**Certified**
ISO/IEC 27001 (cloud)	Underlying AWS infrastructure ISMS	**Inherited from AWS**
ISO/IEC 27017	Cloud-services information security controls	**Inherited from AWS**
ISO/IEC 27018	Protection of personally identifiable information (PII) in the cloud	**Inherited from AWS**
EU AI Act	Risk-based governance of AI systems; enforceable since 2 August 2026	**Aligned by design**
GDPR	Protection of personal data of EU data subjects	**Compliant by design**

ISO/IEC 27001 certification establishes that AnyInsight operates a formal information security management system. Running on AWS infrastructure certified to ISO/IEC 27001, 27017, and 27018 extends that assurance to the cloud layer specifically — 27017 for cloud-service security controls, and 27018 for the protection of personally identifiable information in the cloud. On top of that certified foundation, the platform is architected to meet the requirements of the EU AI Act and GDPR, which the next section maps to concrete capabilities.

© HEARTBOT AI Inc., 2026. All Rights Reserved.

12 Compliance Mapping

Regulation is written in the language of legal obligation; security teams operate in the language of controls. Compliance mapping is the bridge between the two, and the absence of that bridge is where most AI governance programmes fail an audit. A mapping allows an assessor to trace each obligation to the specific mechanism that satisfies it and to the evidence that mechanism produces, turning the assertion "we are compliant" into something that can be tested. It accelerates the work that gates deployment — vendor risk assessments, data protection impact assessments, and internal control reviews — because the analysis has already been structured rather than reconstructed under deadline. Equally important, it exposes gaps while they remain inexpensive to close, rather than during an external audit or in the aftermath of an incident.

The table below maps the principal obligations of the EU AI Act — as amended by Regulation (EU) 2026/1744 — together with GDPR and ISO/IEC 27001, to the specific AnyInsight capabilities that address them. It is intended as a starting point for a security team's own control assessment, not as a legal opinion; specific obligations depend on an organization's role — provider or deployer — and on its use cases under each framework.

Regulatory requirement	Framework	AnyInsight capability
Risk-based classification and governance of AI systems	**EU AI Act**	Policy engine and Governance Layer codify per-scenario risk policies; graded Alert / Block / Mask enforcement
Transparency toward users of AI systems and AI-generated content (Art. 50, applicable 2 Aug 2026)	**EU AI Act**	Every interaction is attributable to a named model or agent and recorded in the Dialogue Log, giving providers and deployers the record needed to disclose AI use and identify AI-generated output
Transparency of how AI systems process data and reach decisions	**EU AI Act**	Dialogue Log records model or agent, inputs, outputs, and context; MAIA critique mode adds explainable review
Human oversight and accountability	**EU AI Act**	Violation monitoring, proactive alerts, and the Moderation Record route every event to a named human owner for resolution
AI literacy of staff and operators (Art. 4, applicable since 2 Feb 2025)	**EU AI Act**	Prompt Wizard, prompt templates, the Get-started Agent and AI Helper, and a 32-language interface deliver in-platform enablement to every user
Inventory and accountable ownership of the AI systems in use	**EU AI Act; ISO 27001 (A.5.9)**	Central resource hub and three-tier Org / Workspace / User assignment maintain an authoritative register of which models, agents, and connectors are in use, by whom, and under what authority
Record-keeping and traceability	**EU AI Act; ISO 27001**	Immutable Data Vault retains all interactions with security flags; downloadable, tamper-evident evidence chain

© HEARTBOT AI Inc., 2026. All Rights Reserved.

Regulatory requirement	Framework	AnyInsight capability
Security and robustness of AI outputs	EU AI Act	Zero-trust GenAI Firewall inspects egress and ingress; prompt protection and content filtering
Readiness for high-risk obligations (Annex III, from 2 Dec 2027)	EU AI Act	Automatic log retention, human oversight through the Moderation Record, and the immutable Data Vault instrument systems now for the documentation and traceability the high-risk regime will require
Lawful processing and PII protection	GDPR; ISO 27018	PII, financial, credential, and health policies detect and mask personal data before it leaves the perimeter
Data minimization and leakage prevention	GDPR	Inline masking, connector-path inspection, and in-platform support that avoids external tools
Access control and least privilege	ISO 27001 (A.5 / A.8); GDPR	Three-tier Org / Workspace / User IAM with least-privilege roles and walled-workspace isolation
Logging, monitoring, and audit	ISO 27001 (A.8.15 / A.8.16)	System, dialogue, automation, shared conversation, and connector logs; configurable dashboard; downloadable and traceable records
Cloud information security and PII in the cloud	ISO 27017 / 27018	AWS-certified infrastructure; multi-tenant isolation and system protection at the infrastructure layer
Incident detection and response	ISO 27001 (A.5.24–A.5.28)	Real-time violation tracking, alarm notifications, and the Moderation Record investigation workflow

Conclusion

AI is not a tool to be bought; it is a capability to be built. A successful transformation requires a plan that unifies technology, people, and principles — and for a security function, that means governance, identity, and audit cannot be separate projects bolted onto an AI rollout after the fact. They have to be the substrate the rollout stands on.

AnyInsight is designed to be exactly that substrate. It gives employees a genuinely productive AI workforce — 20+ models, 350+ agents, 500+ connectors — while placing a zero-trust GenAI firewall, a configurable policy engine, three-tier identity isolation, and an immutable audit trail on every interaction. Productivity, governance, and security are delivered as one system. For the enterprise buyer weighing how to move AI from pilot to production without inheriting unacceptable risk, that integration is the point: it is what lets an organization adopt AI with complete trust, and it is why HEARTBOT positions AnyInsight as the blueprint for a trustworthy AI transformation.

© HEARTBOT AI Inc., 2026. All Rights Reserved.

HEARTBOT AI Inc.

EMAIL customer_success@heartbot.com

WEB www.heartbot.com · www.anyinsight.ai · www.geniebench.ai · www.GenAIsecurity.ai

This document is provided for evaluation purposes and does not constitute legal advice. Specifications and features are subject to change without notice.

© HEARTBOT AI Inc., 2026. All Rights Reserved.