STAFFINITY

EU AI Act — Conformity Assessment

Risk Classification & Compliance

Regulation (EU) 2024/1689 — AI Act

Risk Classification & Compliance Assessment | May 2026

security@staffinity.io | trust.staffinity.io | staffinity.io

EU AI ACT — OVERVIEW & APPLICABILITY

The EU Artificial Intelligence Act (Regulation (EU) 2024/1689) entered into force on **August 1, 2024**. It establishes the world's first comprehensive legal framework for artificial intelligence, applying a risk-based approach to regulate AI systems placed on the EU market or put into service within the EU.

Enforcement Timeline

Date	Milestone
August 1, 2024	EU AI Act entered into force
February 2, 2025	Prohibited AI practices (Title II) banned
August 2, 2025	GPAI model obligations (Title VIII) apply; governance rules apply
August 2, 2026	High-risk AI systems (Annex III) fully subject to all obligations
August 2, 2027	All remaining provisions fully apply

Staffinity's Role Under the EU AI Act

The EU AI Act defines distinct roles with different obligations:

- **Provider (Art. 3(3)):** An entity that develops or places an AI system on the market. Staffinity acts as a Provider of the Staffinity AI Agent Platform.
- **Deployer (Art. 3(4)):** An entity that uses an AI system under its authority. Staffinity also acts as a Deployer when operating agents on behalf of clients in managed service configurations.
- **Downstream Deployer of GPAI (Art. 50ff):** When Staffinity deploys Anthropic Claude as the underlying model, Staffinity is a downstream deployer — not the GPAI model provider.

> **Scope of this Assessment: This conformity assessment covers the Staffinity AI Agent Platform (workplace assistant deployed via Microsoft Teams). It does not cover bespoke AI systems developed specifically for individual clients, which would require separate assessment.**

RISK CLASSIFICATION ANALYSIS

The EU AI Act establishes four risk tiers. This section analyzes which tier applies to the Staffinity AI Agent Platform:

UNACCEPTABLE RISK (Title II)	DOES NOT APPLY

Staffinity does not deploy social scoring systems, real-time biometric surveillance in public spaces, subliminal manipulation techniques, or systems exploiting vulnerable groups. None of the Art. 5 prohibited practices apply.

HIGH RISK (Annex III)	DOES NOT APPLY — Standard Deployments

Annex III Category	Assessment
Biometric identification	✗ Not used — Staffinity uses Microsoft Entra ID for identity; no biometric data processed.
Critical infrastructure	✗ Not applicable — Staffinity serves workplace productivity, not infrastructure control.
Education / vocational training	✗ Not applicable to standard deployments.
Employment / worker management	■ Partial consideration — HR policy Q&A; and scheduling assistance are common use cases. Analysis: Staffinity agents provide information and answer questions; they do not make employment decisions, evaluate employee performance, determine compensation, or control access to employment. Does not meet Annex III threshold.
Essential private services (credit, insurance)	✗ Not a standard Staffinity use case. Financial services deployments are analyzed separately.
Law enforcement	✗ Not used.
Migration / border control	✗ Not used.
Administration of justice	✗ Not used.

CLASSIFICATION RESULT:	LIMITED RISK (Art. 50) — with some deployments at Minimal Risk

Rationale: Staffinity agents interact with humans via a chat interface (Microsoft Teams). Under Art. 50(1), providers of AI systems intended to interact directly with natural persons must ensure those persons are informed they are interacting with an AI system. Staffinity complies with this transparency obligation (see Page 4).

ARTICLE 50 — TRANSPARENCY OBLIGATIONS (LIMITED RISK)

As a Limited Risk AI system, Staffinity's primary obligations fall under Article 50 of the EU AI Act, which establishes transparency requirements for AI systems that interact with humans and AI-generated content:

Art. 50(1) — Disclosure that users are interacting with AI	**■ COMPLIANT**

Staffinity agents are named and persona'd as AI assistants (e.g., Sunny, TARS, Kevin, Bishop). Users are informed at deployment onboarding that they are interacting with an AI system. The agent's AI identity is explicit in the system prompt and visible in the Teams interface.

Art. 50(2) — Disclosure when AI generates content that could be mistaken as authentic	**■ COMPLIANT**

Staffinity agents do not generate deepfakes, synthetic audio, or synthetic media. All agent responses are clearly attributed to the named AI agent in Microsoft Teams. The agent interface is clearly distinguishable from human users in the Teams environment.

Art. 50(3) — Emotion recognition / biometric categorization disclosure	**■ N/A**

Staffinity does not use emotion recognition or biometric categorization systems. No biometric data is processed by the platform. This obligation does not apply.

Art. 50(4) — Labeling of AI-generated content	**■ COMPLIANT**

All responses are visibly attributed to the named AI agent in the Teams interface. Users cannot mistake agent responses for responses from human colleagues because the agent appears as a distinct Teams user/bot with its AI name and avatar.

GPAI MODEL OBLIGATIONS — ARTICLES 51–56

Staffinity uses Anthropic Claude (a General Purpose AI model) via API. This section addresses Staffinity's obligations and position under Title VIII (GPAI) of the EU AI Act:

Primary GPAI Obligations

Anthropic, as the GPAI model provider, bears primary obligations under Art. 51–56, including technical documentation, copyright compliance, and (for systemic risk models) adversarial testing and incident reporting. Staffinity relies on Anthropic's compliance with these obligations.

Staffinity as Downstream Deployer

Staffinity's obligations as a downstream deployer of Anthropic's GPAI model are: (1) use the model within Anthropic's permitted use policies; (2) implement appropriate monitoring and oversight; (3) cooperate with Anthropic's obligations upon request.

No-Training Commitment

Staffinity's Enterprise DPA with Anthropic explicitly provides that client data is never used to train Anthropic's models. This is a core contractual protection for Staffinity's enterprise clients and is passed through to client Data Processing Agreements.

Client API Keys

Staffinity clients use their own Anthropic API keys. This means: (1) the client has a direct contractual relationship with Anthropic; (2) Staffinity cannot access client API credentials; (3) clients can independently verify and enforce Anthropic's GPAI compliance obligations.

Systemic Risk Assessment

Anthropic Claude Sonnet is assessed by Anthropic as not presenting systemic risk under Art. 51. Staffinity monitors GPAI classification decisions and will update this assessment if Anthropic's model classification changes.

> **Important: EU AI Act GPAI obligations for providers apply as of August 2025. Staffinity will confirm Anthropic's compliance documentation before that date and update this assessment accordingly. Clients with EU data residency requirements should deploy in the eu-central-1 (Frankfurt) region.**

TECHNICAL DOCUMENTATION & RECORD KEEPING (ART. 11, 12)

For Limited Risk AI systems, full Art. 11 technical documentation is not mandatory. However, Staffinity voluntarily maintains equivalent documentation as a matter of good governance and enterprise client confidence:

System Architecture	Per-client deployment specifications documented in Exhibit A (Development Plan). Architecture includes ECS Fargate, RDS pgvector, ElastiCache Redis, KMS CMK, CloudWatch monitoring.
Training Data	N/A — Staffinity does not train AI models. Training data provenance is Anthropic's responsibility as GPAI model provider.
Intended Purpose	Documented per client in the Development Plan (Exhibit A), including: use cases, out-of-scope topics, user groups, and operating environment.
Human Oversight	Circuit breaker (3-failure escalation), human escalation path in IRP, configurable human-review routing, client ability to disable agent at any time.
Performance Monitoring	CloudWatch Synthetics (1-min canaries), CloudWatch Alarms, monthly metrics reviews, quarterly business reviews with clients.
Incident Logging	7-year immutable audit trail (DynamoDB + S3 WORM Object Lock COMPLIANCE mode). Exceeds any EU AI Act logging requirement for Limited Risk systems.
Change Management	System prompt changes version-controlled in Exhibit A. Model version updates evaluated in staging before production. Clients notified of material changes.

HUMAN OVERSIGHT (ART. 14) & ACCURACY / ROBUSTNESS (ART. 15)

Human Oversight (Art. 14)

Article 14 requires that high-risk AI systems are designed to allow effective human oversight. While Art. 14 is mandatory for high-risk systems, Staffinity voluntarily implements equivalent human oversight mechanisms for all deployments:

- Agents operate within strict system prompt guardrails defined by the client and Staffinity
- Circuit breaker halts agent operation after 3 consecutive tool/API failures; human notification triggered
- Human escalation path documented in Incident Response Plan (IRP v1.0)
- Clients can modify agent behavior, restrict topics, expand capabilities, or disable the agent entirely at any time — no Staffinity approval required
- ECS service can be scaled to zero (complete shutdown) immediately if needed
- Per-user rate limiting prevents individual abuse and protects system availability

Accuracy (Art. 15)

RAG Grounding	Retrieval-Augmented Generation grounds agent answers in the client's own documents, significantly reducing hallucination risk compared to base model responses.
Source Citation	Agents cite source documents by name in responses, allowing users to verify accuracy independently.
Prompt Injection Protection	Boundary markers applied to all external content retrieved during tool calls, preventing adversarial content from manipulating agent behavior.
Performance Baselines	Response latency, error rate, and circuit breaker trigger rate monitored continuously. Deviations from baseline trigger alerts.

Cybersecurity (Art. 15)

- AWS GuardDuty threat detection and CloudWatch anomaly monitoring
- Per-client AWS account isolation — no cross-client data access possible
- AES-256 at rest (per-client KMS CMK); TLS 1.3 in transit
- MFA required for all administrative access
- Regular security updates included in Staffinity managed service
- AWS Well-Architected Framework review completed with 0 high-risk findings

EU AI ACT READINESS SUMMARY

EU AI Act Requirement	Staffinity Position	Evidence	Status
Risk classification	Limited Risk (Art. 50)	Classification analysis (p.3)	Compliant
Prohibited practices (Title II)	None applicable	Risk analysis (p.3)	Compliant
Art. 50(1) — AI disclosure	Named AI agents; onboarding disclosure	System prompt; Teams UI	Compliant
Art. 50(2) — Authentic content	No deepfakes or synthetic media	Platform architecture (p.4)	Compliant
Art. 50(4) — Content labeling	Agent name visible in Teams	Teams interface design	Compliant
GPAI downstream deployer obligations	Client API keys; no-training DPA	Anthropic Enterprise DPA	Compliant
Technical documentation	Exhibit A per client; architecture docs	Exhibit A; Deployment Plan	Implemented
Human oversight	Circuit breaker; client override; IRP	IRP v1.0; architecture docs	Implemented
Accuracy & robustness	RAG grounding; injection protection	Platform architecture	Implemented
Cybersecurity	GuardDuty; per-client isolation; KMS CMK	AWS WAF review; SOC2 in progress	Implemented
Post-market monitoring	CloudWatch Synthetics; monthly metrics	CloudWatch dashboards	Implemented
ISO 42001 AI Mgmt System	Planned 2027	Roadmap	Planned 2027

Classification Confirmed: Limited Risk — Article 50 Compliant Staffinity monitors EU AI Act regulatory guidance, NIST AI RMF harmonized standards, and forthcoming EU AI Act implementing acts. This assessment will be updated as harmonized standards are published and as the platform evolves. ISO 42001 certification (planned 2027) will provide additional documented evidence of Staffinity's AI management system compliance.

Document Information

Document Title	EU AI Act — Conformity Assessment & Risk Classification
Regulation	Regulation (EU) 2024/1689 (EU Artificial Intelligence Act)
Version	1.0
Prepared	May 2026
Owner	Brent Sudeck, CEO & AI Privacy and Security Officer
Contact	privacy@staffinity.io | trust.staffinity.io