Staffbase

# STAFFBASE EMPLOYEE AI

## The IT leader's guide to compliance, security, and enterprise readiness

*A technical overview of how Staffbase Employee AI is purpose-built to exceed enterprise standards for security, compliance, and control*

## WHAT THIS GUIDE COVERS

Staffbase Employee AI is the intelligent foundation of the Staffbase Employee Experience Platform. It powers secure, compliant, and context-aware automation across communication, content, and insights, helping enterprises maximize employee experience without sacrificing control.

This guide gives IT, security, and compliance leaders a clear technical overview of how Staffbase Employee AI is designed to exceed enterprise requirements for data protection, governance, and control. It answers the most critical questions about data residency, model training, permissions, certifications, and legal safeguards — all grounded in Staffbase's T.R.U.S.T. principles: Transparency, Responsibility, User-in-the-Loop, Security, and Traceability.

As AI becomes foundational to the digital workplace, the question isn't whether to adopt it, but how to do so responsibly, strategically, and at scale. This guide explains how Staffbase delivers exactly that: a secure, closed-loop architecture; enterprise-grade certifications like ISO 27001 and SOC 2 Type II; and full compliance with evolving regulations such as GDPR and the EU AI Act.

You'll also learn how Staffbase gives IT leaders strategic control with administrative flexibility, audit-ready governance, and AI features that support, not replace, human communication.

The goal: to help you evaluate Staffbase Employee AI with full confidence in its security, compliance, and enterprise readiness.

## TOP QUESTIONS ABOUT STAFFBASE EMPLOYEE AI ANSWERED FOR IT PROFESSIONALS

1. Where is data stored and processed, and is it used to train AI models?
2. How do you ensure compliance with key regulations like data protection and AI regulations?
3. What language model powers Staffbase Employee AI?
4. Who owns data inputs and generative outputs?
5. What level of control do we have over the AI features?
6. How does Staffbase Employee AI handle user permissions?
7. What measures are in place to address the risks of AI hallucinations and bias?
8. Do you have independent third-party security certifications that cover your AI services?

# STAFFBASE'S AI PRINCIPLES

The use of AI in Staffbase products is based on the principles of **T.R.U.S.T.** — **Transparency**, **Responsibility**, **User-in-the-loop**, **Security**, and **Traceability**. AI in the Staffbase platform is designed to support human communication, not replace it. Our AI T.R.U.S.T. principles make sure that every feature is understandable, controlled, and built with care. This allows organizations to adopt AI without compromising on ethics or employee experience.

T R U S T

### TRANSPARENCY

Your data stays yours. Under your control, never beyond. No surprises.

### RESPONSIBILITY

We design and deploy AI with accountability, fairness, and compliance at its core, ensuring trustworthy, legal, and ethical outcomes.

### USER-IN-THE-LOOP

We keep humans in control. AI assists, but never replaces human decision-making.

### SECURITY

We protect your data with enterprise-grade safeguards. Privacy and security by design.

### TRACEABILITY

We provide clear insight into how AI results are generated and why. Results are reviewable.

# TOP QUESTIONS ANSWERED FOR IT PROFESSIONALS

## 1. Where is data stored and processed, and is it used to train AI models?

Staffbase is committed to data privacy, security, and compliance at an enterprise level. All AI-related features are built on a foundation of data segregation and controlled processing, keeping customer information protected and isolated at all times.

Staffbase uses a self-run model and Azure OpenAI Services. Unlike solutions that rely on shared or public AI endpoints, Staffbase Employee AI operates within a fully isolated, enterprise-managed Azure environment, so no data ever leaves your organization's secure tenant.

### THIRD-PARTY MODELS (AZURE OPENAI SERVICES)

Staffbase does not, and will never, use customer data or content to train third-party models. For AI features that use Azure OpenAI Services, all prompts, inputs, and outputs are processed and stored exclusively within the Microsoft Azure infrastructure. This means your enterprise data is:

- **Encrypted and securely stored** within Azure's enterprise-grade environment.
- **Isolated per customer** — never mixed or shared across Azure tenants.
- **Not transferred to or processed by OpenAI-operated services**, such as ChatGPT.
- **Not used** to train or improve any Microsoft, OpenAI, or third-party models.

This architecture delivers end-to-end data segregation and compliance with global standards, including GDPR.

### STAFFBASE MODELS AND DATA USAGE

Staffbase's self-run model is securely hosted on Microsoft Azure. Your enterprise data are processed as follows:

- **Customer-created content**, such as text, images, or files uploaded to an intranet or employee app, is not used to train or fine-tune Staffbase's AI models.
- **Usage data** — such as engagement patterns or interaction data — is a key driver of continuous performance improvements across the Staffbase platform. It powers contextually relevant experiences, like personalized content recommendations, that help employees find what they need faster. This is always done in a closed-loop system: models are trained exclusively on each customer's own metadata and remain fully contained within that customer's secure environment, with no data sharing or mixing.

→ **Learn more about data, privacy and security at Microsoft Azure OpenAI Service.**

→ **Learn more about Staffbase security standards.**

## 2. How do you ensure compliance with key regulations like data protection and AI regulations?

Staffbase's compliance obligations are governed by existing contractual agreements with our customers.

- **Data Privacy:** All personal data processed through AI features is handled strictly in accordance with our **Data Processing Addendum (DPA)** and flows only through approved **sub-processors** listed therein. Data processing remains fully within the defined legal and geographical boundaries (e.g., EU, North America, or regional hosting, as applicable).
- **AI Regulations:** Staffbase reviews all AI-powered features under applicable AI classification frameworks (including the EU AI Act and national or state-level AI acts). Based on their design and purpose, our features are assessed to qualify as **limited-risk or minimal-risk** systems under the EU AI Act. Continuous reviews are conducted to ensure compliance as the regulatory environment evolves.

- **Contractual Safeguards:** The introduction of AI functionality does not alter any existing Staffbase contracts. Specific AI clauses are included in our Product-Specific Terms to provide explicit coverage for AI features, ensuring consistent data-handling, security, and accountability obligations.

### GOVERNANCE AND OVERSIGHT

Staffbase maintains a cross-functional AI governance process that includes legal, security, and product experts to monitor emerging AI and data-protection regulations globally. This ensures that all AI developments remain **compliant by design** and that customers retain **full control and transparency** over their data.

## 3. What language model powers your AI?

Staffbase uses only trusted, enterprise-grade language models that have been reviewed and approved by our Security and Legal teams. We maintain full transparency about the models behind each feature, and customers can review detailed information in our Support Portal.

Most of the generative AI capabilities provided by Staffbase are powered by Microsoft's Azure OpenAI service which uses **OpenAI's large language models (LLMs)**, securely hosted within **Microsoft Azure's enterprise infrastructure**. This setup provides the benefits of OpenAI's advanced model capabilities while keeping all data processing in a controlled, compliant, and regionally governed Azure environment.

For example, features such as **Podcast Summaries**, **Companion Editor**, and **AI Translations** use these Azure-hosted OpenAI models. Staffbase may also leverage other specialized or fine-tuned models for non-generative use cases (e.g., content recommendations or metadata analysis). All models are reviewed for data protection, access control, and risk classification before deployment.

**To review the models behind each AI feature and their specific configurations, please visit our Support Portal.**

## 4. Who owns data inputs and generative outputs?

Customers retain full ownership rights and control over their data, including all prompts, inputs, and any AI-generated content produced within the platform.

All customer data — whether entered manually or produced through AI — is handled in accordance with your existing **Data Processing Addendum (DPA)** and the AI-specific clauses in our **Product-Specific Terms**. This framework guarantees:

Customer content remains **private**, **protected**, and **isolated** within your environment.

All data handling complies with the same security, privacy, and retention standards that govern the rest of the Staffbase platform.

In short: **you own your data**, and Staffbase protects it under the same contractual and technical safeguards that apply to all other platform content.

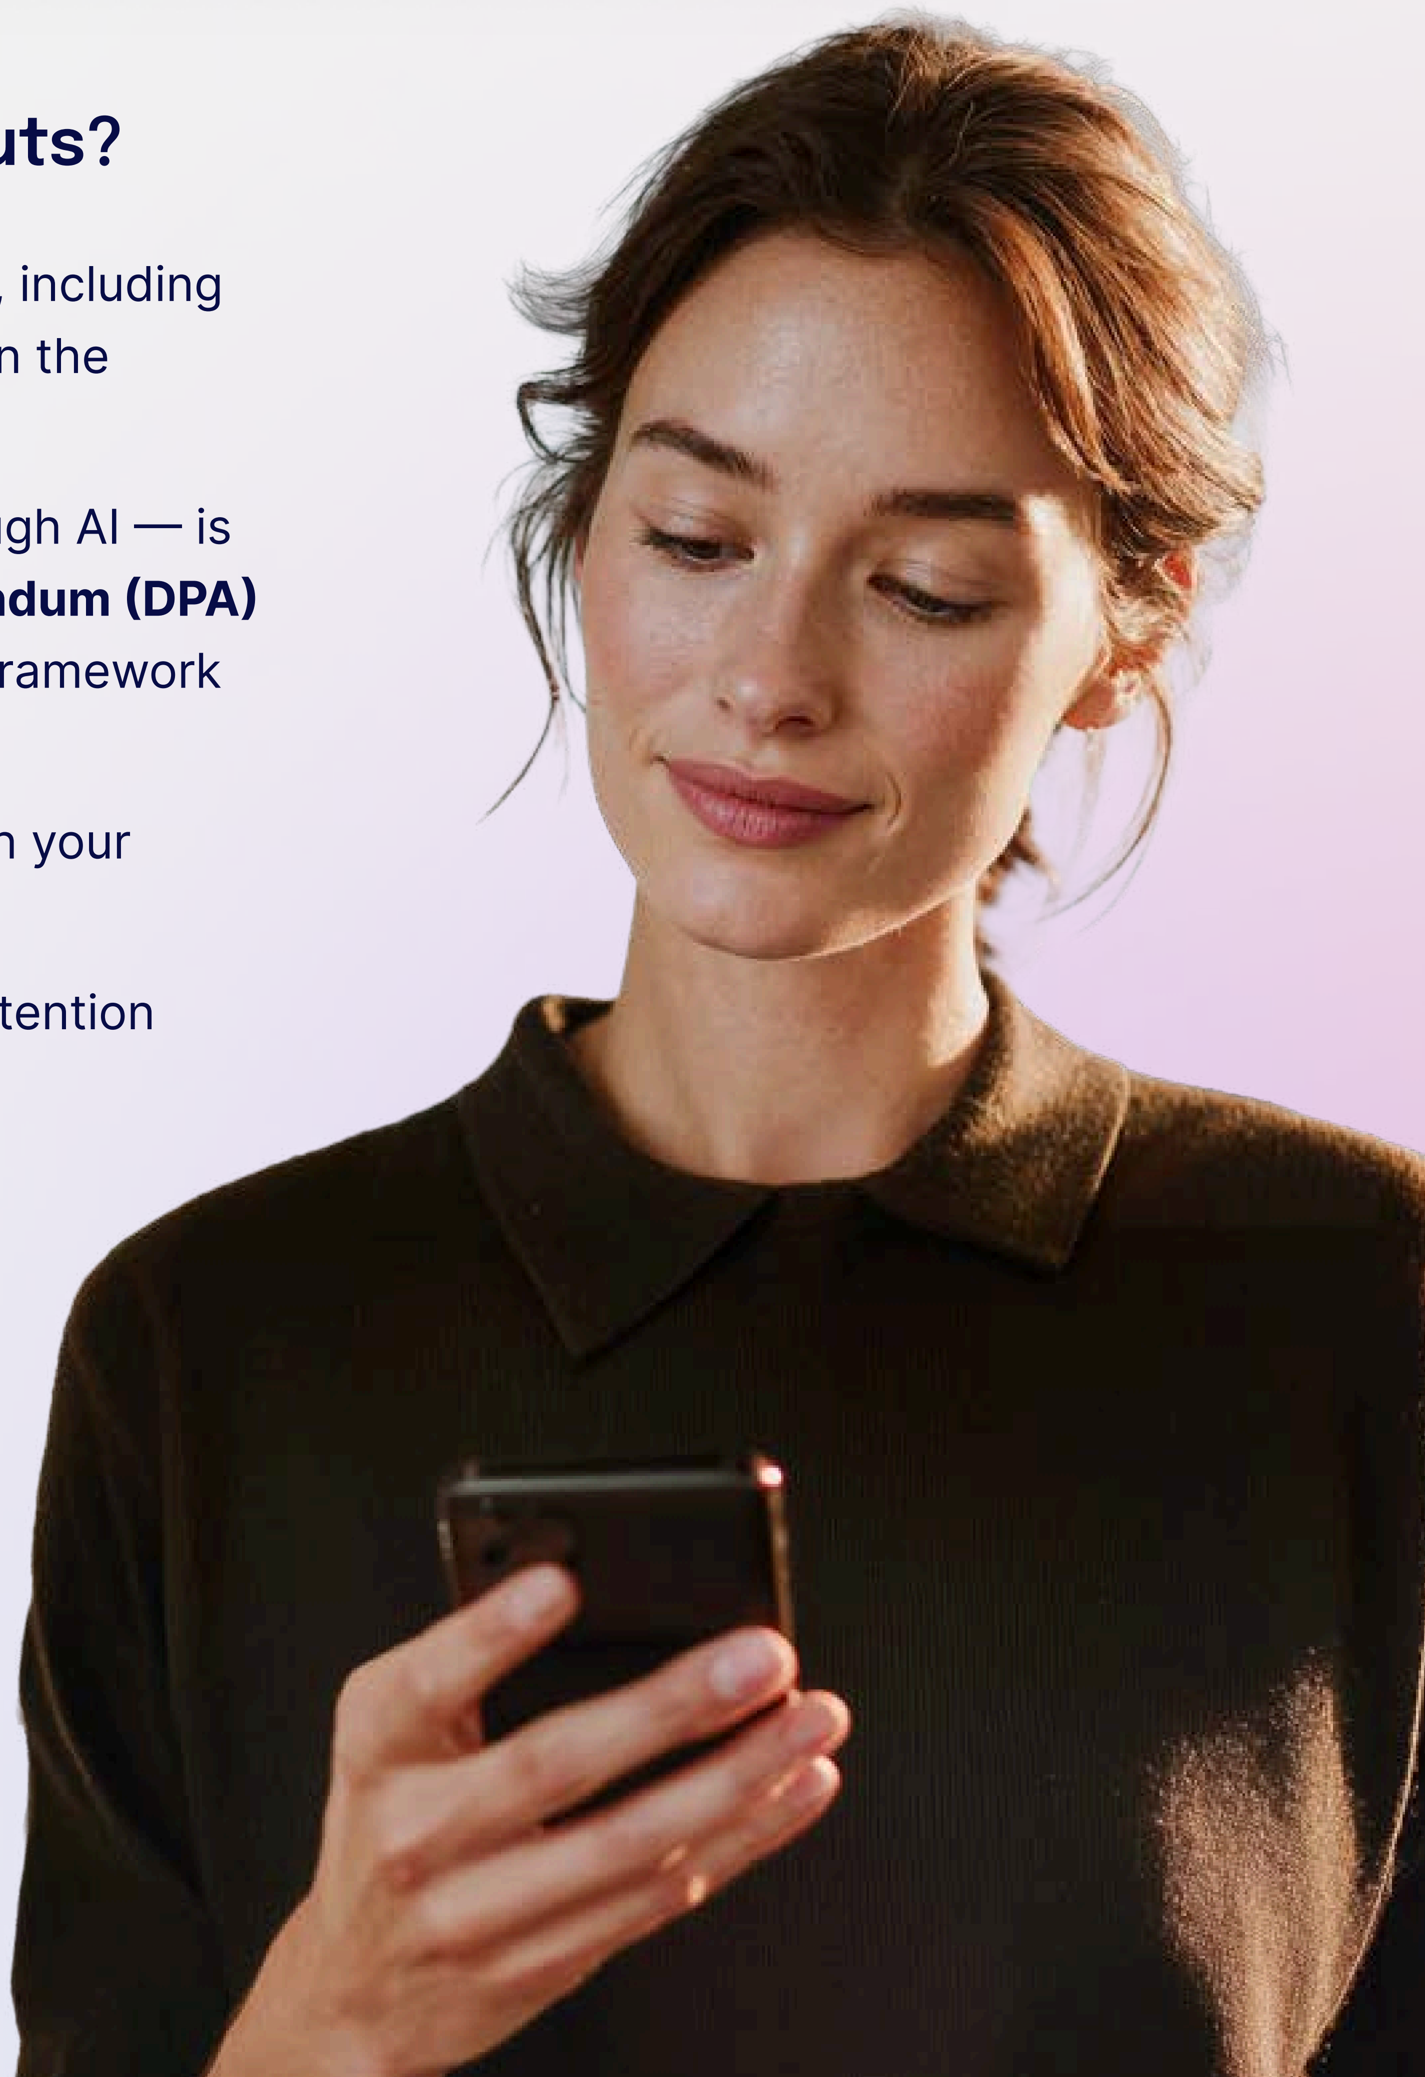

## 5. What level of control do we have over the AI features?

Staffbase provides full administrative control so organizations can decide exactly how and when AI is used within their environment.

AI features are opt-in by default. Your team decides what to activate, when, and for whom. Reach out to Staffbase Support or your Customer Success Manager for more information. Customers stay in control of the deployment of AI features. This is guided by several guardrails:

- **Administrative control:** Organizations can disable individual AI features — or all AI functionality — at any time by contacting **Staffbase Support** or their **Customer Success Manager**.
- **Self-service settings:** A self-service activation and deactivation option is currently in development for **Staffbase Studio**, which will allow administrators to immediately enable or disable AI features directly from the platform interface.
- **Consistent governance:** Whether AI is active or not, all data handling remains subject to the same **contractual, privacy, and security protections** defined in the DPA and Product-Specific Terms.

This approach gives IT leaders **complete control and oversight** of AI usage, aligning deployment with their organization's policies and risk management requirements.

## 6. How does Staffbase Employee AI handle user permissions?

**Staffbase AI features fully respect existing user roles and content permissions within the platform.** Users can only generate or access content that aligns with their current permissions and visibility settings.

AI models operate strictly within the same access boundaries defined for each user. This is achieved by:

- **Isolating AI** prompts from protected content when the model doesn't require contextual data.
- **Using the active user's permissions** when content retrieval is necessary, so only information the user is already authorized to view is included in the AI prompt.

This ensures that Staffbase Employee AI operates under the same **principle of least privilege** as the core platform, maintaining consistency, security, and compliance across all AI interactions.

## 7. What measures are in place to address the risks of AI hallucinations and bias?

Staffbase recognizes that generative AI systems can produce inaccurate or biased outputs, as responses are generated based on statistical likelihood rather than guaranteed facts. To minimize these risks, Staffbase applies a combination of **model selection**, **product design**, **transparency**, and **human oversight**.

### HOW STAFFBASE MITIGATES HALLUCINATION AND BIAS

- **Model vetting:** Staffbase uses only enterprise-grade language models that are reviewed by our Security, Legal, and Product teams for quality, risk level, and compliance with industry best practices.
- **AI-native infrastructure:** Generative features are embedded within a secure, closed-loop architecture that minimizes uncontrolled model behavior and keeps prompts contextually grounded in the user's existing environment.
- **Transparency by design:** Wherever feasible, **AI-generated content is clearly labeled** (for example, with a magic wand icon or labels such as "AI Summary"). Labeling and contextual cues help users identify dynamically generated content.
- **Human oversight:** Staffbase strongly encourages users to review and validate all AI-generated outputs before publishing or distributing them, particularly for critical communications or sensitive topics.

These measures reflect Staffbase's T.R.U.S.T. principles, designed to keep AI features helpful, accountable, and under user control.

## 8. Do you have independent third-party security certifications that cover your AI services?

Staffbase maintains independent third-party certifications to leading international security and privacy standards, supporting both our platform and AI features with proven, enterprise-grade controls.

### STAFFBASE CERTIFICATIONS

Staffbase is certified to:

**ISO/IEC 27001** — Information Security Management
**SOC 2 Type II** — Security, Availability, and Confidentiality
**TISAX Level II** — Automotive industry data security
**TX-RAMP Level II** — Texas Risk and Authorization Management Program
**California Consumer Protection Act** — Data Rights and Transparency for California Residents

As a company headquartered in Germany, Staffbase also operates under the **EU General Data Protection Regulation (GDPR)** and provides a **Data Processing Addendum (DPA)** that meets GDPR requirements.

### AI-SPECIFIC FRAMEWORKS

The Microsoft Azure infrastructure used to host our generative AI models (Azure OpenAI service) is certified under **ISO/IEC 42001:2023**, the new international standard for AI management systems.

By building on Azure's certified AI governance, risk, and security controls, Staffbase ensures that its AI features inherit the same enterprise-grade protections, alongside the safeguards already established by our own ISO and SOC frameworks.

Staffbase maintains robust, continuously updated resources on our Support Portal, where IT and security teams can find the latest information on data handling, governance, and certifications.

## Ready to unlock AI productivity for the whole workforce?

Visit staffbase.com/employee-ai — or scan the QR code — to see how Staffbase Employee AI transforms employee experience.