

Sample — Enterprise AI & Security Review

7 questions · 7 approved · status approved · exported 2026-05-24

QUESTION 1

What AI models or LLM providers do you use in production, and for what specific use cases?

ANSWER

Acme AI uses Anthropic Claude (claude-opus-4-7 and claude-sonnet-4-6) as its primary LLM provider for all production customer-facing AI features. OpenAI GPT-4 is retained as a fallback path behind a feature flag and accounted for less than 2% of inference traffic in Q1 2026. Production use cases are limited to three: ticket triage and auto-categorization; suggested replies for support agents (always human-reviewed before sending); and knowledge base search via semantic retrieval.

HIGH CONFIDENCE

APPROVED

SOURCES

primary-llm-provider — Anthropic Claude (claude-opus-4-7, claude-sonnet-4-6) for all production customer-facing AI features. OpenAI GPT-4 retained as a fallback path behind a feature flag, used <2% of traffic in Q1 2026.

ai-use-cases — Three production use cases: (1) ticket triage and auto-categorization, (2) suggested replies for support agents (human-in-the-loop, never auto-sent), (3) knowledge base search via semantic retrieval.

QUESTION 2

Do you use customer data to train, fine-tune, or improve AI models? If not, what contractual guarantees are in place with your LLM providers?

ANSWER

No. Acme AI does not use customer data to train, fine-tune, or improve AI models. Anthropic's enterprise zero-data-retention (ZDR) amendment is in effect for all API traffic, signed 2025-11-12, which contractually prohibits Anthropic from retaining or training on prompts. The terms of service explicitly carry this prohibition through to all customers.

HIGH CONFIDENCE

APPROVED

SOURCES

training-data-policy — We do NOT train models on customer data. Anthropic's enterprise terms (zero data retention for API traffic) are in effect. Prompts are not used by Anthropic to train their models.

Sample — Enterprise AI & Security Review

7 questions · 7 approved · status approved · exported 2026-05-24

QUESTION 3

How do you ensure human oversight over AI-generated outputs that reach end users?

ANSWER

Every AI-generated reply is shown to a human support agent for review before being sent to the customer. Agents can edit, regenerate, or discard the suggestion. Accept, edit, and reject rates are logged per agent and per category. No AI output reaches a customer without explicit human approval.

HIGH CONFIDENCE

APPROVED

SOURCES

human-oversight — Every AI-generated reply is shown to a human agent for review before being sent. Agents can edit, regenerate, or discard.

QUESTION 4

Where is customer data stored, and what data residency options exist for EU customers?

ANSWER

All customer data is stored in AWS us-east-1 by default, with cross-region encrypted backup to us-west-2. EU customers can opt into eu-west-1 residency under our Data Processing Addendum. Backup snapshots inherit the same KMS-managed encryption as the primary store.

HIGH CONFIDENCE

APPROVED

SOURCES

data-residency — All customer data stored in AWS us-east-1 (primary) with cross-region encrypted backup to us-west-2. EU customers can opt into eu-west-1 residency under our Data Processing Addendum (DPA).

Sample — Enterprise AI & Security Review

7 questions · 7 approved · status approved · exported 2026-05-24

QUESTION 5

How is your AI system classified under the EU AI Act risk tiers, and how do you meet the corresponding obligations?

ANSWER

Acme AI is classified as a Limited Risk system under the EU AI Act (Article 50). The product provides AI-generated content to support agents who review and approve every output before it reaches an end user. We meet the Article 50 transparency obligation through clear user-facing disclosure inside the agent console. The system does not fall under any Annex III high-risk use case.

HIGH CONFIDENCE

APPROVED

SOURCES

eu-ai-act — Acme AI is classified as a Limited Risk system under the EU AI Act (Article 50): provides AI-generated content to support agents who review before sending. Transparency obligations met via user-facing...

QUESTION 6

Do you hold SOC 2 Type II? What is the audit period and audit firm?

ANSWER

Yes. Acme AI holds a SOC 2 Type II report issued 2026-01-15 covering the Security, Availability, and Confidentiality trust service criteria. The audit was performed by Prescient Assurance. The report is available to enterprise customers under NDA.

HIGH CONFIDENCE

APPROVED

SOURCES

soc2 — SOC 2 Type II report issued 2026-01-15 covering Security, Availability, Confidentiality. Audit by Prescient Assurance.

Sample — Enterprise AI & Security Review

7 questions · 7 approved · status approved · exported 2026-05-24

QUESTION 7

Are you HIPAA compliant, and do you sign Business Associate Agreements?

ANSWER

Acme AI is not a Covered Entity or Business Associate under HIPAA and does not sign BAAs. The Terms of Service explicitly prohibit customers from uploading Protected Health Information (PHI).

HIGH CONFIDENCE

APPROVED

SOURCES

hipaa — We are NOT a Covered Entity or Business Associate under HIPAA. We do not sign BAAs and customers must not upload PHI.