

Five question AI response example

Fictional worked example | Meridian Desk | Draft Assist v3 | As of 2026-10-28

Use this example to separate supported answers from open factual questions. All companies, people, providers, evidence and approvals are invented for practice. No actual customer result is represented.

The five answers contain two supported responses, two unresolved items and one justified not applicable response. Evidence excerpts follow, so you can check every assertion without another download.

Q01 What data is sent to the model and who sends the final reply?

Supported | Ready for manual export

Draft Assist v3 sends selected ticket text and a style instruction to the fictional Aster inference API. It returns a draft. A support agent must click Send.

Owner: Riley Chen. Evidence: MD-E01 dated 2026-10-20. Approval: Approved by Jordan Park on 2026-10-28.

Limits or applicability reason: Draft Assist v3 only. Other features are outside this answer.

Q02 Do your application logs contain prompt or response bodies?

Supported | Ready for manual export

No bodies appeared in application logs in a five-request staging test. Those logs held request ID, model label and latency.

Owner: Morgan Lee. Evidence: MD-E02 dated 2026-10-22. Approval: Approved by Jordan Park on 2026-10-28.

Limits or applicability reason: Staging sample only. Production behavior and provider logs remain unverified.

Q03 Does the model provider use our submitted data to train models?

Unknown | Unresolved

We have not verified the provider terms for this account. We cannot confirm a training exclusion.

Owner: Morgan Lee. Evidence: None. Approval: Pending. Reviewer: Jordan Park.

Limits or applicability reason: Account-specific provider terms and settings are outstanding.

Open items and applicability

Q04 Can every copy of AI data be deleted within 24 hours?

Unknown | Unresolved

We have not verified deletion across provider systems, backups and ticket storage. We cannot commit to 24 hours.

Owner: Morgan Lee. Evidence: None. Approval: Pending. Reviewer: Jordan Park.

Limits or applicability reason: Scope and deletion evidence are missing. The 7-day draft-store setting does not prove this claim.

Q05 Where do you run training jobs for your own custom models?

Not applicable | Ready for manual export

Not applicable to Draft Assist v3. We do not run custom model training jobs for this feature.

Owner: Riley Chen. Evidence: MD-E03 dated 2026-10-20. Approval: Approved by Jordan Park on 2026-10-28.

Limits or applicability reason: The feature only calls hosted inference. Provider training practices are a separate unresolved question Q03.

Use this on your own request

1. Make a separate folder for one customer request. Copy the five fields below for each question. Use Unknown whenever evidence or approval is missing.

Question and draft	Source and limits	Owner and approval
Your question and proposed answer	Exact file, section, date and limits	Factual owner, reviewer, decision and date
_____	_____	_____
_____	_____	_____

2. Verify the cited source supports the exact wording. Ask the owner for missing evidence. Keep unknowns visible when the deadline arrives.

3. Send only wording that a named reviewer has approved, together with its limits and source reference. If a customer requires yes or no, do not convert Unknown to Yes. Request clarification or retain a qualified unresolved response.

Scope: AI-specific customer questions. This does not complete a whole security assessment or attest to controls. The team supplies facts and owns its response.

Fictional source index

MD-E01 Architecture note v3 section 2

FICTIONAL SOURCE. Dated 2026-10-20. Owner: Riley Chen. Review by: 2026-11-20.

Draft requests send the selected ticket text and style instruction to the fictional Aster API. The API returns a draft. An agent must click Send.

Limits: Applies only to Draft Assist v3. No assertion about other features.

Sharing: Can share excerpt. This excerpt is the complete simulated evidence available in this exercise.

MD-E02 Content logging test LT-04

FICTIONAL SOURCE. Dated 2026-10-22. Owner: Morgan Lee. Review by: 2026-11-22.

In staging, five synthetic draft requests produced metadata-only application logs containing request ID, model label and latency. No prompt or response body appeared in those application logs.

Limits: Five staging requests only. Provider logging and production behavior were not tested.

Sharing: Can share excerpt. This excerpt is the complete simulated evidence available in this exercise.

MD-E03 Feature inventory v3 section 1

FICTIONAL SOURCE. Dated 2026-10-20. Owner: Riley Chen. Review by: 2026-11-20.

Draft Assist v3 has no training job, fine-tuning pipeline or custom model upload. It calls a hosted inference API.

Limits: Does not establish whether the provider uses API data for its own training.

Sharing: Can share excerpt. This excerpt is the complete simulated evidence available in this exercise.

Evidence still needed

Q03: Morgan must obtain the account-specific provider terms and settings. Q04: Morgan must identify every storage location and test deletion timing. Neither item has supporting evidence in this sample.

The paid pack adds a 50-question matrix, owner chase log, two complete ten-question examples and editable response and data-flow documents. This free sample is usable on its own. Product-use help: lattice@shirasusystems.com (Lattice AI customer support). Do not send private questionnaire evidence or credentials.