

BUYER-READY AI PROCUREMENT EVIDENCE

Turn due diligence into a deal asset.

A synthetic example showing how AssureDeck separates proven facts, partial evidence, and open work for one fictional AI system.

EVIDENCE READY	PROVEN	PARTIAL	OPEN
68%	147	61	112

SYNTHETIC	Northstar Support AI is fictional. All systems, counts, evidence records, and findings are illustrative.
BOUNDARY	Customer-support response drafting for English-language tickets. No automated sending, HR, credit, health, or safety decisions.
METHOD	One claim status per answer. Every PROVEN answer needs current evidence, an owner, a reviewer, and a clear system scope.

This sample does not reproduce the official AI-CAIQ question set. Clients obtain official CSA resources directly from CSA.

What the buyer can trust - and what remains open.

The fictional system has enough evidence to answer many operational questions, but not enough to present every control as operating. The packet blocks unsupported claims and directs the client team to the highest-impact gaps.

STATUS	COUNT	INTERPRETATION
PROVEN	147	Current, in-scope evidence supports the claim and has an accountable owner.
PARTIAL	61	Some evidence exists, but scope, freshness, operation, or approval is incomplete.
OPEN	112	Evidence is missing, planned, unknown, or contradicted and cannot support a positive answer.

BUYER-RELEVANT STRENGTHS

PROVEN	A current model and provider inventory names owners and change triggers.
PROVEN	Customer data-use language is approved and points to provider contract evidence.
PROVEN	Human review is mandatory before any response reaches a customer.
PROVEN	Production traces record system version, retrieval source IDs, output, and reviewer action.

MATERIAL GAPS

P0	No completed AI-specific incident response exercise.
P0	Evaluation thresholds are defined, but regression acceptance is not consistently approved.
P1	Provider change review exists as a checklist but has no completed operating sample.

One AI workflow, clearly contained.

FIELD		SYNTHETIC SYSTEM RECORD
System		Northstar Support AI
Intended use		Draft an English-language response for a support agent to review and edit.
Excluded use		No autonomous sending. No employment, credit, health, identity, legal, or safety decision.
Input		Support ticket, approved knowledge-base excerpts, and limited account context.
Output		Suggested reply plus cited knowledge-base source IDs.
Human checkpoint		A support agent must review, edit if needed, and explicitly send.
Retention		Prompt and output trace retained 30 days; business ticket retained under the support platform schedule.
System owner		VP Customer Operations
Technical owner		Director of Applied AI

ILLUSTRATIVE COMPONENT INVENTORY

ID	COMPONENT	ROLE	OWNER
CMP-01	Support application	Creates request and holds final agent-approved reply.	Support Eng
CMP-02	Retrieval service	Selects approved knowledge-base excerpts and source IDs.	Applied AI
CMP-03	Example model API	Generates a draft from the scoped prompt and context.	Applied AI
CMP-04	Trace store	Records versions, source IDs, output, and reviewer action.	Platform Sec

CLAIM LIMIT	This packet cannot support a claim that responses are always correct, unbiased, safe, or fully automated. The system is explicitly a human-reviewed drafting workflow.
-------------	--

Every positive answer points somewhere real.

ID	ARTIFACT	STATUS	OWNER	SCOPE / LIMIT
EV-001	AI system register	CURRENT	Product Ops	Northstar Support AI only
EV-002	Model and provider register	CURRENT	Applied AI	Version reviewed monthly
EV-003	Customer data-use statement	APPROVED	Privacy	Buyer-NDA disclosure
EV-004	Evaluation plan	PARTIAL	AI Quality	Regression sign-off inconsistent
EV-005	Evaluation results 2026-Q2	CURRENT	AI Quality	English support set only
EV-006	Human review procedure	CURRENT	Support Ops	Mandatory send action
EV-007	AI incident addendum	DRAFT	Security	No completed exercise
EV-008	Production trace sample	CURRENT	Platform Sec	Redacted buyer-NDA sample

TRACEABILITY RULE

A file name is not enough. A reusable evidence record includes the system scope, owner, approver, effective date, latest verification date, confidentiality level, supported claim IDs, and limitations.

NO COPY

Restricted source files should stay in the client's approved repository. The evidence index links to the canonical source and carries only the metadata needed for review.

Missing proof becomes buildable work.

PRI	GAP	OWNER	ACCEPTANCE EVIDENCE	DUE
P0	AI incident exercise not completed	Security	Scenario, attendee list, timed actions, findings, approved remediation	10d
P0	Regression threshold approval inconsistent	AI Quality	Versioned threshold record and two completed release approvals	15d
P1	Provider change checklist has no operating sample	Applied AI	Completed review for one model-version or term change	30d
P1	Retention test omits vector-store deletion	Platform	Deletion test log covering prompt, output, trace, and embedding records	30d
P2	AI literacy completion is not centrally reported	People Ops	Role list, curriculum, completion export, and exception owner	45d

STATUS TRANSITIONS

PLANNED	The client has approved work and a target date. It is not phrased as an operating control.
PARTIAL	Evidence exists but scope, freshness, operation, or approval is incomplete. The limitation travels with the answer.
PROVEN	Current, scoped evidence supports the claim and records an accountable owner and reviewer.

A status changes only after the acceptance evidence exists and an authorized client reviewer approves the external representation.

Approved language keeps its evidence and limits.

The prompts below are original examples for this synthetic packet. They are not copied from the official AI-CAIQ.

PROVEN	Example prompt: Does the product require human review before an AI-generated support response is sent? Approved response: Yes. Northstar Support AI produces a draft only. A support agent must review and explicitly send the final response. Evidence: EV-006 Human review procedure; EV-008 redacted production trace sample. Limit: This answer applies only to the scoped English-language support workflow.
PARTIAL	Example prompt: Are model updates evaluated against approved release thresholds before production use? Qualified response: Evaluation thresholds exist and current results are versioned. Release approval has not been applied consistently across all recent regression runs. Evidence: EV-004 Evaluation plan; EV-005 Evaluation results 2026-Q2. Open work: P0 backlog item requires a versioned approval record and two completed release samples.

CLIENT APPROVAL RECORD

ROLE	APPROVAL	DATE	SCOPE
System owner	Approved	YYYY-MM-DD	System boundary and intended use
Technical owner	Approved	YYYY-MM-DD	Architecture, model, evaluation, logging
Security	Qualified	YYYY-MM-DD	Incident exercise remains open
Legal / privacy	Client review	YYYY-MM-DD	External wording only; no legal opinion by AssureDeck

METHOD REFERENCES

Cloud Security Alliance, AI Controls Matrix v1.1, released 2026-06-22: cloudsecurityalliance.org/artifacts/ai-controls-matrix-v1-1

Cloud Security Alliance, AI-CAIQ v1.1, released 2026-06-23: cloudsecurityalliance.org/artifacts/ai-consensus-assessments-initiative-questionnaire-ai-caiq-v1-1

Cloud Security Alliance, STAR for AI: cloudsecurityalliance.org/star/ai

BOUNDARY	AssureDeck provides evidence-production and deal-enablement services. It does not provide certification, audit, attestation, assurance opinions, legal advice, compliance determinations, or guaranteed buyer outcomes.
----------	--