AVENLOFT SYSTEMS GROUP
Governance, Risk & Compliance

LABLIST CASE FILE | PROJECT 2

EVALUATE A NEW AI USE CASE

CASE ID
GRC-AI-002

ASSIGNED BY
Aaron Bell, Director - Governance, Risk & Compliance

ISSUED
September 22, 2026

YOUR ASSIGNMENT

Assess a proposed AI-enabled SaaS use case for Customer Success. Determine whether the pilot should proceed and under what conditions.

DELIVERABLES

- AI use-case inventory
- AI risk assessment
- Decision memo

Fiction notice: Avenloft, LumenFerry, all people, records, terms, metrics, and events in this packet are fictional. Reserved example domains are used where applicable.

Assignment Brief

To: Junior GRC Analyst

From: Aaron Bell, Director - Governance, Risk & Compliance

Business sponsor: Lena Ortiz, VP, Customer Success

Subject: Governance Review - Customer Success AI Summarization Pilot

Due: October 2, 2026

The Customer Success AI request documented in the September Security Steering Committee notes remains paused. Lena Ortiz has now submitted a defined pilot proposal for an AI-enabled SaaS service that would summarize customer support calls and tickets and prepare draft internal case notes. Leadership wants GRC to determine whether the proposed use is appropriate for a controlled pilot and, if so, what conditions must be satisfied before customer data is processed.

Your job is not to decide whether AI is “good” or “bad.” Review the intended use, data, vendor dependencies, test evidence, human oversight, existing safeguards, and unresolved questions. Then make a defensible recommendation.

Analyst rule	Do not convert every AI concern into a reason to reject the project. Also do not treat vendor security claims as proof that the use case is safe. Your decision should connect the specific context of use to the risks and controls that matter.

Required Deliverables

- AI use-case inventory record: document what the system is, why Avenloft wants it, who uses it, what data it processes, and who owns it.
- AI risk assessment: identify and rate the material risks using Avenloft's existing qualitative risk method.
- Decision memo: recommend Approve, Conditionally Approve, Defer, or Reject, with required controls, pilot boundaries, and monitoring expectations.

Assessment Boundaries

- Use the supplied facts and the authoritative NIST references. Record material unknowns instead of inventing vendor behavior or legal obligations.
- Assess the proposed Customer Success use case, not every capability the vendor markets.
- Do not provide legal advice. Hannah Kim's notes identify contract and privacy questions that may require validation before deployment.
- The packet is designed so that more than one final decision can be defensible if the reasoning and conditions are supported by evidence.

Avenloft Context for This Review

Avenloft Systems Group, Inc. remains the same fictional B2B logistics SaaS company used in Project 1. Avenloft Relay serves 185 business customers and approximately 22,000 customer users. Customer Success owns several workflows that handle customer support tickets, call transcripts, and customer-provided troubleshooting material.

Company fact	Current case context
Business objective	Scale responsibly: use automation and AI where productivity improves without creating unmanaged customer-data or decision risk.
AI risk appetite	Moderate appetite for controlled pilots; low tolerance for customer or confidential data entering unapproved tools.
Customer data	Business contact information, warehouse locations, purchase-order identifiers, inventory quantities, shipment schedules, support attachments.
Assurance posture	Current SOC 2 Type II report; next examination period begins in Q4 2026.
Relevant prior observation	OBS-12 noted that Customer Success requested an AI summarization tool and that Avenloft lacked a formal AI use-case inventory or standard AI-risk intake process.

Governance Roles

Name	Role in this review
Aaron Bell - Director, GRC	Coordinates the AI risk review, records the use case, and prepares the governance recommendation.
Lena Ortiz - VP, Customer Success	Business sponsor and accountable owner for the proposed Customer Success workflow.
Maya Chen - CISO	Security executive and escalation point for material customer-data or AI risk.
Hannah Kim - General Counsel & Privacy Lead	Reviews contractual, privacy, and data-handling questions that GRC should not decide alone.
Noah Grant - Director, Procurement	Coordinates vendor onboarding, contracting, and the third-party review workflow.
Marcus Holloway - Director, IT & Workplace Technology	Supports SSO, access lifecycle, and workforce technical controls if the pilot proceeds.

Interim AI Direction

- AI use cases that process company or customer information require a named business owner and documented review before production use.
- Customer or confidential information may not be entered into unapproved AI services.
- AI-generated content used in customer communications or authoritative business records requires human review unless a separate approval explicitly permits automation.
- Third-party AI services must complete the applicable security, privacy, procurement, and contractual review before production use.

AI Use-Case Intake - Submitted by Customer Success

Intake field	Submitted information
Business owner	Lena Ortiz, VP, Customer Success
Proposed service	LumenFerry Resolve, an AI-enabled SaaS support summarization service
Business problem	Support agents spend significant time turning call transcripts and long ticket histories into internal case notes before handoff or closure.
Intended users	25 Customer Success agents during an eight-week pilot
Primary AI activity	Summarize ticket text and support-call transcripts; extract issue themes and draft action items.
Output destination	A draft field in Salesforce. The agent can edit, accept, or reject the generated summary before it becomes the final case note.
Requested data	Ticket body, selected ticket metadata, and call transcript text. Attachments are not required for the proposed pilot.
Out of scope	Automated customer replies, entitlement decisions, refunds, account actions, case closure, or autonomous prioritization.
Requested start	October 12, 2026, to evaluate productivity before Q4 renewal volume peaks.
Estimated spend	$46,000 annually if the pilot converts to a production subscription.

Business Sponsor Estimate

Customer Success estimates that agents create approximately 58,000 case notes per year. Lena Ortiz believes a reliable summarization workflow could reduce repetitive documentation time and improve handoffs between agents, but she does not want AI output sent directly to customers.

What GRC must decide	Is the proposed use sufficiently understood and controlled to begin a limited pilot? If not, what must change before it can proceed?

Proposed Pilot Workflow

The requested integration is intentionally narrow. LumenFerry Resolve would read selected case content, generate an internal summary, and return the result to a draft field for agent review.

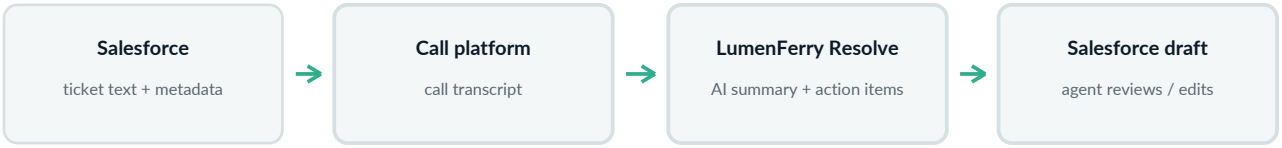

Proposed pilot flow. Attachments and automated customer replies are outside the requested scope.

Step	Proposed operation	Important detail
1	Salesforce sends the ticket body, case ID, product area, and selected metadata to LumenFerry through an API integration.	Customer contact name and business email may appear in the ticket body.
2	The current call-transcription platform sends transcript text for calls linked to the case.	Avenloft already creates transcripts for internal support use; the AI service would be a new recipient of that content.
3	LumenFerry sends the submitted text to a third-party foundation-model service and generates a summary and action items.	The underlying model provider is a subprocessor to LumenFerry.
4	The generated summary is returned to an AI Draft field in Salesforce.	The draft does not overwrite the original ticket or transcript.
5	The assigned agent compares the draft with the source material and chooses Accept, Edit, or Reject.	Only accepted or edited text is moved into the final internal case note.

Explicit Pilot Exclusions

- No support attachments are sent to LumenFerry.
- No AI-generated text is sent automatically to customers.
- The AI does not close cases, change entitlements, issue refunds, or take account actions.
- The pilot does not use AI to score employees or customers.

Vendor Fact Sheet - LumenFerry Resolve

LumenFerry AI, Inc. and LumenFerry Resolve are fictional. The information below represents the vendor material Avenloft received during procurement. Treat it as vendor-supplied evidence, not independently verified fact.

Area	Vendor-provided information
Service model	Multi-tenant SaaS for support summarization and workflow assistance. Enterprise plan supports API integration and SAML SSO.
Hosting	Primary application and customer-content storage are hosted in a U.S. cloud region.
Encryption	TLS 1.2 or later in transit; AES-256 encryption at rest for stored customer content.
Identity	SAML SSO and SCIM are available on the enterprise plan. Local password accounts can be disabled after SSO configuration.
Audit logs	Administrative events and user generation events are logged. Enterprise customers can export 180 days of administrative audit activity.
Assurance	Vendor states that a current SOC 2 Type II report is available under NDA. Avenloft has not yet reviewed the report.
Incident notification	Standard enterprise terms state that affected customers are notified without undue delay after the vendor confirms a security incident involving customer content.
Business continuity	Vendor states that production backups are maintained and recovery procedures are tested annually. Detailed recovery results were not supplied in the standard packet.

Procurement Routing Note

At the proposed $46,000 annual spend, the current procurement workflow would not automatically route LumenFerry to Security based on spend alone. Noah Grant manually escalated the request because the intake identified customer-data access and an AI capability. This is consistent with the vendor-classification weakness documented in Project 1.

Vendor AI and Data-Handling Responses

Question	LumenFerry response
Is customer content used to train shared or public models?	Under the enterprise terms, submitted customer content and generated outputs are not used to train LumenFerry shared models or the underlying foundation model.
How long does LumenFerry retain submitted content?	Default application retention is 90 days. Enterprise customers may configure retention to 30 days. Zero-retention processing is not available for Resolve.
Does another AI provider receive the content?	Yes. LumenFerry uses a third-party foundation-model subprocessor. The subprocessor may retain API inputs and outputs for up to 30 days for abuse and security monitoring.
Does LumenFerry use product telemetry?	The vendor may use aggregated and de-identified service telemetry to improve reliability and product performance. The standard response does not define whether transformed excerpts of customer content can be included in that telemetry.
Can a single case be deleted?	Yes. Authorized customer administrators can request deletion of stored case content; the vendor states that deletion completes within seven days, subject to backup rotation.
How are model changes handled?	LumenFerry manages the underlying model selection. Material platform changes are documented in release notes; subprocessor changes receive at least 30 days notice under enterprise terms.
Does the product expose confidence scores?	No. Resolve does not provide a calibrated confidence score for individual summaries.
Does the product retain traceability information?	Each generation records the Avenloft user, source case ID, timestamp, product release, and internal model-version identifier.

Do not skip the wording

“Not used for training” does not answer every data-governance question. Retention, subprocessors, telemetry, deletion, access, and change management are separate facts to evaluate.

Avenloft Support Data Profile

Customer Success reviewed a sample of recent support records to identify the types of information that could enter the proposed workflow. The sample below is synthetic but reflects Avenloft's normal support categories.

Data category	Typical presence	Example in support workflow
Business contact information	Common	Customer name, work email, job title, company name.
Operational customer data	Common	Warehouse location, shipment schedule, inventory quantity, purchase-order identifier.
Product and technical data	Common	Avenloft Relay configuration details, API error messages, integration logs, feature settings.
Customer-provided files	Occasional	Screenshots, CSV extracts, log files. Attachments are outside the proposed AI pilot.
Credentials / secrets	Prohibited but observed	Three support tickets in the prior 12 months contained a pasted API key or secret before the value was removed and the customer was instructed to rotate it.
Payment-card or health data	Not intentionally processed	Relay and the support workflow are not designed to collect payment-card PAN or protected health information.

Synthetic Support Examples

Case	Excerpt
SR-18421	“Warehouse 14 in Reno is showing tomorrow's carrier pickup at 06:00 instead of 08:00. PO 77142 is attached to the shipment.”
SR-18509	“The API returns 403 after the customer changed the integration service account. Log excerpt shows tenant ID and endpoint path.”
SR-18544	“Please remove the token from the prior comment. Customer confirmed the credential was rotated after the support agent flagged it.”

Data-Minimization Question

The business sponsor requested ticket text and call transcripts because they contain the context needed for summarization. The proposed pilot does not require attachments. The case packet does not establish whether every ticket field currently selected for export is necessary.

Pre-Pilot Evaluation Results

Before submitting the governance request, Customer Success evaluated LumenFerry Resolve with 50 synthetic support cases. No real customer content was used. Two Customer Success supervisors compared the generated summaries with the source records.

Measure	Observed result
Fully accurate and materially complete summaries	42 of 50 cases
Summary omitted a material troubleshooting step or limitation	5 of 50 cases
Summary introduced at least one unsupported factual detail	3 of 50 cases
Case identifier extracted correctly	50 of 50 cases
Generated summary sent externally or took an autonomous action	0 cases; the tested configuration had no customer-send or action permission
Median time to prepare an internal case note	Reduced from 6.8 minutes to 3.1 minutes in the test workflow

Examples from the Evaluation

Test case	Observed behavior
T-17	The source said a sync issue appeared “intermittently after 14:00 UTC.” The generated summary stated that the issue began at 14:00 UTC, removing the uncertainty.
T-31	The summary correctly described the API error but omitted that the customer had already rotated the service-account secret.
T-42	The generated action item suggested “restart the warehouse connector.” The source material did not contain that recommendation.

What the Test Does Not Establish

- Performance with real customer data.
- Performance across every support category or uncommon edge case.
- Whether quality will remain stable after a model or product update.
- Whether agents will consistently detect subtle errors when they are busy.

Stakeholder Notes

Lena Ortiz - Customer Success	“The value is mostly documentation time. I do not want the tool deciding refunds, closing cases, or sending messages. If an agent cannot see the source and correct the draft, the pilot is not useful to me.”
Maya Chen - CISO	“The pilot can be narrow, but customer information cannot quietly become training data or disappear into a vendor chain we cannot track. I also want a way to stop the integration quickly if quality or security becomes unacceptable.”
Hannah Kim - Legal & Privacy	“Several enterprise agreements have different requirements for new subprocessors and use of customer data. I need the final data flow and vendor terms before I can tell the business which customers, if any, require notice or additional approval.”
Noah Grant - Procurement	“The spend-based workflow would have missed this review. If GRC approves the pilot, I can add the security and AI conditions to the order form and make the vendor review part of the renewal process.”
Marcus Holloway - IT	“SAML and SCIM are straightforward if the enterprise plan is purchased. I would disable local logins and tie access to the Customer Success pilot group.”
Aaron Bell - GRC	“I need a decision that another reviewer can understand six months from now: what we approved, what we prohibited, why, and what would cause us to reassess or shut it down.”

NIST AI RMF Working Lens

Avenloft uses the NIST AI Risk Management Framework (AI RMF 1.0) as a voluntary governance reference for this exercise. The framework organizes AI risk management into four functions. The functions are not a checklist and do not have to be performed as a rigid sequence.

Function	How to use it in this case
GOVERN	Look for ownership, accountability, policy, inventory, third-party governance, and a defined review process.
MAP	Document the intended purpose, users, data, affected stakeholders, human role, system boundaries, and third-party dependencies.
MEASURE	Use test evidence and other information to evaluate reliability, privacy, security, traceability, and known limitations. Record what has not been measured.
MANAGE	Prioritize the material risks, decide whether the pilot should proceed, set treatment conditions, define monitoring, and establish change or shutdown triggers.

The NIST AI RMF Generative AI Profile (NIST AI 600-1) identifies risks that may be unique to or intensified by generative AI. Not every risk in that publication is relevant to this support-summarization use case. Consider which of the following materially apply to the facts in this packet.

Risk area	Question for this use case
Confabulation	Could the system produce unsupported or incorrect statements that are accepted into the support record?
Data privacy	Could customer or sensitive information be retained, disclosed, reused, or processed beyond Avenloft's intended scope?
Human-AI configuration	Could agents over-rely on generated summaries or fail to provide meaningful review?
Information integrity	Could inaccurate AI-generated case notes affect future support work or customer decisions?
Information security	Are the AI service, integration, identities, logs, and customer content adequately protected?
Value chain / component integration	What risks come from LumenFerry's reliance on a third-party foundation-model provider and future model changes?

Framework note

NIST AI RMF 1.0 is the currently published framework used for this case. NIST has announced that AI RMF 1.0 is being revised. Use the framework as a risk-management lens, not as a compliance certification.

Risk Assessment Method for This Case

Use the same qualitative likelihood-and-impact method established in Project 1. It is based on NIST SP 800-30 Rev. 1, Appendix I, Table I-2. This keeps Avenloft's enterprise risk language consistent across cybersecurity and AI-governance reviews.

Likelihood	Working interpretation
Very Low	Adverse impact is highly unlikely based on the available evidence.
Low	Adverse impact is unlikely but plausible.
Moderate	The scenario is somewhat likely to occur and produce adverse impact.
High	The scenario is highly likely to occur and produce adverse impact.
Very High	The scenario is almost certain to occur and produce adverse impact.
Impact	**Working interpretation**
Very Low	Negligible effect on operations, customers, data, or objectives.
Low	Limited effect handled through normal operations.
Moderate	Serious effect requiring management attention.
High	Severe effect such as significant customer harm, contractual consequence, data exposure, or material operational impact.
Very High	Multiple severe or sustained consequences across Avenloft and its stakeholders.

Likelihood / Impact	Very Low	Low	Moderate	High	Very High
Very High	Very Low	Low	Moderate	High	Very High
High	Very Low	Low	Moderate	High	Very High
Moderate	Very Low	Low	Moderate	Moderate	High
Low	Very Low	Low	Low	Low	Moderate
Very Low	Very Low	Very Low	Very Low	Low	Low

Rate inherent risk before giving credit to current or proposed safeguards. Then determine residual risk after considering controls that are already operating or that you make explicit conditions of approval. If a proposed control is not yet implemented, distinguish it from a current control in your assessment.

Required AI Governance Artifacts

1. AI Use-Case Inventory Record

Field	What to document
Use-case ID / title	Create a simple identifier and business-readable name.
Business owner	Accountable owner for the use and outcome.
Purpose / expected benefit	Problem the AI is intended to solve and expected business value.
Users / affected parties	Who operates the system and who may be affected by its outputs.
AI capability	What the AI does in the approved context.
Data inputs / outputs	Information sent to the system and what it generates.
System / vendor dependencies	Third-party service, underlying model dependency, integrations, and major subprocessors.
Human oversight	Where a person reviews, overrides, corrects, or approves output.
Deployment status	Proposed, Pilot, Approved, Restricted, Suspended, or Retired.
Review owner / date	Who monitors the use and when it must be reassessed.

2. AI Risk Assessment

Field	What to document
Risk statement	Scenario linking an AI-related failure, condition, or dependency to business/customer impact.
NIST AI RMF / GAI lens	Relevant function, trustworthiness concern, or GAI risk area when useful.
Evidence	Case facts supporting the assessment.
Inherent L / I / Risk	Before current safeguards or approval conditions.
Current controls	Safeguards that already exist and are evidenced.
Required conditions	Controls that must be implemented before or during the pilot.
Residual L / I / Risk	Expected remaining risk if required conditions are satisfied.
Owner / monitoring	Accountable owner and evidence that should be monitored.

Decision Memo Requirements

Your final memo should be short enough for the Security Steering Committee to use. It should make a decision rather than simply restating the evidence.

Decision	Meaning in this case
Approve	The proposed pilot may proceed as submitted because the material risks are already within tolerance.
Conditionally Approve	The pilot may proceed only after specified controls or boundaries are implemented and verified.
Defer	Do not begin the pilot yet; material information or controls must be completed before a decision can be made.
Reject	The proposed use is not appropriate within Avenloft's risk tolerance or cannot be reasonably controlled in the proposed form.

Your Memo Should Cover

- Decision: one of the four outcomes above.
- Approved scope: exactly what users, data, integrations, and AI functions are allowed.
- Required conditions: controls that must exist before or during the pilot.
- Prohibited uses: capabilities or data that remain outside approval.
- Monitoring: quality, security, privacy, user-behavior, incident, and vendor-change evidence that should be reviewed.
- Reassessment triggers: events such as a material model change, new data type, expanded automation, unacceptable error rate, vendor incident, or change in contractual requirements.
- Open items: questions requiring Legal, Security, Procurement, the vendor, or the business owner to resolve.

Questions Your Analysis Should Be Able to Answer

1. What business outcome is Avenloft trying to improve, and is generative AI actually necessary for the proposed task?
2. What information will enter the system, and which fields or data types are unnecessary for the pilot?
3. Who is accountable for the use case and who performs human oversight?
4. What do the evaluation results tell you about reliability, confabulation, and the need for review?
5. What customer-data risks remain after considering retention, subprocessors, telemetry, access controls, and deletion?
6. Which vendor statements still need validation or contractual clarification?
7. What risks come from the underlying foundation-model provider and future model changes?
8. What should Avenloft measure during the pilot?
9. What would cause you to pause, narrow, or terminate the use case?
10. Should the pilot proceed, and under what conditions?

Quality Bar for the Analyst

Strong analyst behavior	Weak analyst behavior
Defines the approved use case narrowly and documents what is out of scope.	Reviews “AI” as one generic risk without considering the actual workflow.
Separates vendor claims, test evidence, existing controls, and proposed conditions.	Treats every vendor statement as verified or every unknown as proof of failure.
Uses NIST AI RMF to organize judgment without turning it into a checklist.	Copies framework categories without connecting them to the evidence.
Connects confabulation, privacy, human oversight, and third-party dependencies to business impact.	Lists buzzwords such as hallucination or bias without explaining relevance.
Makes a clear decision and defines monitoring and reassessment triggers.	Says “use AI responsibly” without boundaries, owners, or evidence.

Authoritative Reference Material

- NIST AI 100-1 - Artificial Intelligence Risk Management Framework (AI RMF 1.0)
  https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10
- NIST AI RMF Playbook
  https://airc.nist.gov/airmf-resources/playbook/
- NIST AI 600-1 - Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile
  https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence
- NIST SP 800-30 Rev. 1 - Guide for Conducting Risk Assessments
  https://csrc.nist.gov/pubs/sp/800/30/r1/final

Current-version note	As of September 2026, AI RMF 1.0 remains the published NIST AI RMF used in this training case, and NIST states that a revision is in progress. NIST AI 600-1 was updated on April 8, 2026.

Fiction and Safety Notice

Avenloft Systems Group, Inc., Avenloft Relay, LumenFerry AI, Inc., LumenFerry Resolve, all employees, customer records, vendor terms, test results, metrics, contract excerpts, and events in this packet are fictional and created for training. Commercial products named in the continuing Avenloft environment are real products used only to make the fictional environment recognizable.

The domains avenloft.example and lumenferry.example use the reserved .example namespace. This packet contains no credentials, personal data, active infrastructure, or production customer information.