

CampaignPilot Support Operations

AI-assisted support triage, risk-aware routing, response drafting, and operational visibility for a fictional email-marketing SaaS platform.

Google Sheets


Google Gemini

Make.com

Airtable

Gmail

END-TO-END WORKFLOW



12 controlled test cases

100% category + queue accuracy

HITL human-in-the-loop

Jerry J. Hudson

AI Support Operations / Workflow Automation

Fictional SaaS environment

From support intake to controlled action

The goal was not to let an LLM run support autonomously. It was to combine AI classification with deterministic business rules so routine work moves faster while higher-risk cases stay under human control.

THE CHALLENGE

Support teams need consistent triage, routing, and faster first responses - but fully autonomous AI can create risk around billing, security, cancellation, data loss, and emotionally charged cases.

THE SOLUTION

Gemini classifies each ticket into a structured schema. Make.com then applies deterministic guardrails, routes the ticket, stores the record in Airtable, and creates a Gmail draft only when the case is eligible.

WORKFLOW ARCHITECTURE



DESIGN PRINCIPLES

- Structured AI output: category, subcategory, priority, queue, sentiment, confidence, risk flags, and summary.
- Deterministic routing controls the operational decision; the model does not decide whether an email is sent.
- Draft-eligible responses are created in Gmail as drafts for review - never auto-sent.
- Airtable provides a single operational view of ticket status, priority, routing, and action queues.

Google Sheets

Google Gemini

Make.com

Airtable

Gmail

AI where it helps. Rules where they matter.

The classifier supplies context. The routing layer enforces safety and operational policy.

HUMAN REVIEW TRIGGERS

Risk flag present

Security/privacy, refund disputes, data loss, cancellation/deletion, and ambiguity are escalated.

Confidence below 0.85

Low-confidence classification does not proceed as routine automation.

Angry sentiment

Emotionally escalated cases are held for human review.

High or Critical priority

Severity alone can move a case out of the automated drafting path.

Account cancellation

Irreversible or sensitive account actions remain human-controlled.

Why this matters

The LLM can be wrong about subjective labels such as sentiment or severity. Deterministic guardrails prevent a single model judgment from silently becoming an irreversible customer action.

Operational pattern

Classification is probabilistic; escalation policy is explicit and reviewable.

MAKE.COM FILTER LOGIC



The screenshot shows the 'Set up a filter' interface in Make.com. The filter is configured with the following conditions:

- Condition 1:** 2. Result: risk_flags[] (Array operators: Array length greater than 0)
- Condition 2:** 2. Result: confidence (Numeric operators: Less than 0.85)
- Condition 3:** 2. Result: sentiment (Text operators: Equal to Angry)
- Condition 4:** 2. Result: priority (Text operators: Equal to critical)
- Condition 5:** 2. Result: category (Text operators: Equal to Account Cancellation)

The conditions are connected by 'or' logic. The interface includes 'Add AND rule' and 'Add OR rule' buttons between conditions, and 'Cancel' and 'Save' buttons at the bottom.

A dashboard built for action, not just reporting


Airtable turns the automation into an operations surface: what came in, what needs attention, and what is ready for review.



Dashboard screenshot captured after cleanup of development-only test records.

Helpful automation without auto-send risk

Only eligible tickets reach the second Gemini step. The output is saved as a Gmail draft so a person can review, edit, or discard it before anything reaches the customer.



DRAFTING RULES

- Acknowledge the customer issue clearly and professionally.
- Do not mention AI, internal classifications, confidence scores, or routing logic.
- Do not promise refunds, credits, cancellation, security actions, or other irreversible outcomes.
- Do not invent product features, troubleshooting facts, or hidden investigation steps.
- Ask a clarifying question only when the missing information is genuinely necessary.
- Return a concise customer-facing draft - not analysis or JSON.

Control point


The automation creates a draft, not a sent message. This preserves the speed benefit of AI-assisted writing while keeping customer communication under human review.

SAFE OUTCOME STATES

Human Review	Escalated
Draft Ready	Review draft
Routed	No draft

A controlled 12-case evaluation

This was a portfolio-scale test, not a production benchmark. Expected outcomes were compared with actual classifier and workflow results, with documented SME adjudication where the test expectation did not match the finalized operating rule.



CampaignPilot — AI Ticket Triage Lab										
File Edit View Insert Format Data Tools Extensions Help										
Q Menu 100% 123 Default... 10 B I A										
O26										
	A	B	C	D	E	F	G	H	I	J
1	Metric	Correct	Total	Accuracy	Notes					
2	Category Accuracy	12	12	100%	All primary ticket categories matched the expected category					
3	Queue Accuracy	12	12	100%	All operational queues matched the expected queue					
4	Priority Accuracy	9	12	75%	T001, T007, and T012 differed from the expected priority					
5	Sentiment Accuracy	8	12	66.70%	Calculated after SME adjudication of T011 from Angry to Frustrated					
6	Risk Flag Accuracy	12	12	100%	All expected escalation risk flags matched and non-risk cases remained unflagged					
7	Human Review Accuracy	12	12	100%	Calculated against the final production routing rules					
8	Draft Eligibility Accuracy	12	12	100%	Calculated against the final production routing rules					
9	Structured Output Validity	12	12	100%	All 12 Gemini classifier responses returned the required structured fields in a usable format					
10										
11										
12	Evaluation Notes									
13	Test ID	Adjudication / Finding								
14	T004	Post-evaluation adjudication: The final production routing rule sends High or Critical priority tickets to Human Review. Operational expected values are Human Review = Yes and Draft Allowed = No.								
15	T011	Post-evaluation SME adjudication: Sentiment revised from Angry to Frustrated. Under the final routing rules, Frustrated does not trigger Human Review, so operational expected values are Human Review = No and Draft Allowed = Yes.								
16	Subcategory	Subcategory was treated as a qualitative field rather than a strict accuracy metric because Gemini frequently returned semantically correct but more specific wording than the expected test label.								
17	Overall Finding	The system performed strongest on category classification, queue assignment, risk detection, escalation decisions, and draft eligibility.								
18		Priority and sentiment showed greater variation because they require more subjective judgment.								
19										
20										

WHAT THE TEST SHOWED

- The workflow was strongest on operational classification, queue assignment, risk detection, escalation decisions, and draft eligibility.
- Priority and sentiment showed more variation, which is expected because both require more subjective judgment.
- Subcategory was treated qualitatively because Gemini often returned semantically correct but more specific wording than the expected label.
- T004 and T011 were documented as post-evaluation adjudications rather than silently rewriting the original test intent.

Evaluation note

These results demonstrate portfolio-scale behavior across 12 controlled cases. They should not be presented as production-scale model performance.

What this project demonstrates

The value of the solution is not that an LLM can label tickets. It is the system around the model: structured outputs, explicit routing policy, operational visibility, safe drafting boundaries, and measurable evaluation.

AI WORKFLOW DESIGN

Structured classification and controlled response drafting

SUPPORT OPERATIONS

Queues, severity, escalation, risk, and human review

AUTOMATION ENGINEERING

Make.com routers, filters, Airtable records, Gmail drafts

EVALUATION

Ground-truth test cases, SME adjudication, measured outcomes

FINAL RESULT

A complete AI-assisted support operations workflow that can classify, route, escalate, draft, track, and surface work - while keeping higher-risk decisions and outbound communication under human control.

Jerry J. Hudson

CampaignPilot Support Operations - Portfolio Case Study

Gemini

Make.com

Airtable

Google Sheets

Gmail