

Sendex AI Support Escalation System

RAG-grounded support triage, controlled routing, engineering handoffs, human review, and audit logging for a fictional B2B SaaS support operation.

HubSpot CRM


Make

Knowledge / RAG

Slack

Airtable

Google Sheets



5 controlled validation cases

3 routing outcomes

RAG explicit knowledge retrieval

Jerry J. Hudson

From escalation request to controlled handoff

The goal was not to let an AI agent run support autonomously. It was to combine retrieval-grounded reasoning with explicit routing rules so the right work reaches the right owner.


THE CHALLENGE

Support teams need consistent escalation quality without sending routine tickets to Engineering or forcing frontline agents to gather Engineering-only diagnostics.

THE SOLUTION

HubSpot flags eligible tickets. Make retrieves relevant Sendex knowledge, aggregates the context, runs the AI Agent, then routes RETURN, ENGINEERING, or HUMAN_REVIEW.

WORKFLOW ARCHITECTURE




DESIGN PRINCIPLES

- Knowledge retrieval happens before AI reasoning, so the agent receives explicit grounded context.
- The AI Agent returns a structured decision, confidence, reasoning, impact scope, and known-issue match.
- Engineering priority is assigned by deterministic rules rather than by MRR or the model alone.
- Every outcome is written back to HubSpot, notified in Slack, and logged in Google Sheets.

RAG first. Agent reasoning second.

The final design makes retrieval visible and testable instead of relying on a hidden optional knowledge lookup inside the agent.

RETRIEVAL FLOW



1 RETRIEVE

Query all five Sendex knowledge files using product area, subject, and symptoms. Return the top 10 semantic chunks.

2 AGGREGATE

Combine the retrieved chunks into one Retrieved Knowledge input so the agent runs once per ticket.

3 DECIDE

The agent compares ticket evidence against retrieved knowledge before setting `similar_ticket_match` and routing.

Why this matters

The controlled campaign test correctly matched KI-1005 - Campaign Processing Queue Degradation only after relevant knowledge was retrieved and passed into the agent.

Operational pattern

Explicit retrieval -> aggregation -> structured reasoning -> deterministic routing. Each stage can be inspected independently during testing and review.

STRUCTURED OUTPUT

- decision: RETURN / ENGINEERING / HUMAN_REVIEW
- confidence, reasoning, missing information, recommended action
- impact_scope, customers_affected, similar_ticket_match
- Known-issue IDs are returned only when retrieved evidence supports the match.

Three outcomes, each with a controlled next step

The agent supplies context. Make routes the ticket and triggers only the actions appropriate for that outcome.

RETURN

Frontline-resolvable or incomplete support work

HubSpot is marked Return, the original support agent is notified in Slack, and the decision is logged. No Engineering issue is created.

AI Review – Ticket Returned to Support

Ticket: CSV import rejects valid date values
Product Area: Contacts & Segmentation

AI Assessment:

The issue appears to be user-format mismatch (DD/MM/YYYY vs required YYYY-MM-DD) and not a platform defect. Frontline can guide the customer to reformat dates and retry; no observed platform error or internal failure suggesting Engineering involvement at this time.

Information / Troubleshooting Needed:

None

Recommended Next Action:

- Have the customer format all date values in the CSV to YYYY-MM-DD as per the import template and retry the import. - Provide the customer with a quick rollback/retry checklist and a sample snippet showing the correct date format (YYYY-MM-DD) for clarity. - Confirm whether any non-date fields in the import also require template-aligned formats and verify a small test batch before a full import.

HubSpot Ticket ID: 214336681935

ENGINEERING

Evidence supports a product or service failure

Airtable creates the Engineering issue, HubSpot receives the issue ID, #engineering is notified, and the action is logged.

New Engineering Escalation – SEN-1020

Issue: Campaign remains stuck in processing after successful validation

Priority: High

Product Area: Campaigns & Sending

Customer Plan: Professional

Customers Affected: 1

Impact Scope: SINGLE_CUSTOMER

AI Assessment: The symptoms show campaigns remain in Processing beyond normal waits and affect multiple campaigns on the same account, which aligns with the documented KI-1005 Campaign Processing Queue Degradation. Evidence supports a product defect requiring Engineering investigation.

Additional Information Needed: None

Engineering Next Step: Escalate to Engineering with reference KI-1005, including: campaigns stuck in Processing beyond expected time for this account and the fact that multiple campaigns on the same account are impacted. Provide all collected details (times, statuses, account timezone, and steps already taken) to enable targeted investigation.

HubSpot Ticket ID: 214436896740

Engineering Issue ID: SEN-1020

HUMAN_REVIEW

Ambiguous or judgment-heavy escalation

The ticket is routed to #senior-support for technical judgment. It is not sent back to frontline and no Engineering issue is created yet.

Human Review Required

Ticket: Intermittent contact count mismatch after segmentation refresh

Product Area: Contacts & Segmentation

Original Support Agent: Sofia

Customer Plan: Enterprise

Customer MRR: \$2200

AI Assessment: The intermittent count mismatch after segmentation refresh is ambiguous and could reflect delayed processing, recalculation behavior, or a product defect. Frontline cannot determine the root cause from available data and a senior technical judgment is required to decide if Engineering involvement is warranted.

Missing / Unclear Information: Exact timestamps of observed counts (with time zones); segment name/ID and filters; whether issue occurs across other segments/accounts; pattern (time of day, load, refresh cycles); account timezone and environment; any recent changes to segmentation logic; any audit/correlation IDs tied to observations.

Recommended Next Action: Review the intermittent count discrepancy to determine if it represents delayed/partial recalculation or a product defect; provide Senior Support with a verdict on whether Engineering should be engaged. If warranted, share the escalation rationale and any high-level guidance for monitoring or reproduction.

AI Confidence: 65

HubSpot Ticket ID: 214336681937

A handoff Engineering can act on

Valid escalations carry the ticket evidence, deterministic priority, and any grounded known-issue context into a tracked internal issue.

HUBSPOT TICKET STATE

Tickets

All tickets

My open tickets

Unassigned tickets

Search (/)

Filter

Sort

Ticket owner

Create date

Last activity date

Priority

Advanced filters

Ticket name	Ticket status	Create date	Priority	Escalation Requested	All Reviewed	Product Area	Ticket description	Source	All Review Outcome	Engineering Issue ID	Customer MRR
Intermittent contact co...	New (Click to update)	Yesterday at 7:12 PM EDT	High	Yes	Yes	Contacts & Segmentation	Customer reports a segment somethi...	Email	Human Review	---	2200
Webhook deliveries fail...	New (Click to update)	Yesterday at 7:11 PM EDT	High	Yes	Yes	Integrations & API	Customer's webhook endpoint is reach...	Email	Engineering	SEN-1002	1800
Campaign remains stuck...	New (Click to update)	Yesterday at 7:09 PM EDT	High	Yes	Yes	Campaigns & Sending	A scheduled campaign remains in Proc...	Chat	Engineering	SEN-1020	899
API returns 403 after A...	New (Click to update)	Yesterday at 7:08 PM EDT	Medium	Yes	Yes	Integrations & API	Customer reports 403 responses after ...	Email	Return	---	699
CSV import rejects vall...	New (Click to update)	Yesterday at 7:07 PM EDT	High	Yes	Yes	Contacts & Segmentation	Customer is importing a contacts CSV...	Email	Return	---	249

ENGINEERING NOTIFICATION

 New Engineering Escalation — SEN-1020

Issue: Campaign remains stuck in processing after successful validation

Priority: High

Product Area: Campaigns & Sending

Customer Plan: Professional

Customers Affected: 1

Impact Scope: SINGLE_CUSTOMER

AI Assessment: The symptoms show campaigns remain in Processing beyond normal waits and affect multiple campaigns on the same account, which aligns with the documented KI-1005 Campaign Processing Queue Degradation. Evidence supports a product defect requiring Engineering Investigation.

Additional Information Needed: None

Engineering Next Step: Escalate to Engineering with reference KI-1005, including: campaigns stuck in Processing beyond expected time for this account and the fact that multiple campaigns on the same account are impacted. Provide all collected details (times, statuses, account timezone, and steps already taken) to enable targeted investigation.

HubSpot Ticket ID: 214436896740

Engineering Issue ID: SEN-1020

SEN-1020

High priority

Campaigns & Sending

1 customer affected

KI-1005 match

Priority is assigned by deterministic impact rules - not by customer MRR alone.

AIRTABLE ENGINEERING ISSUE TRACKER

	Issue Number	R	HubSpot Ticket ID	Product Area	Customers Affected	Issue Summary	Troubleshooting Co...	Related Tickets	Business Impact	Created At	Resolution
1	Drag to reord	12	\$1,800.00 21428807871	Integrations & API	1.0	The customer reports Sendex-initiated web...	None		Impact Scope: SINGLE_CUSTOMER Customers Affec...	3/10/2026 8:31pm	
2	Double-click to wrap	0	\$899.00 214436896740	Campaigns & Sending	1.0	The symptoms show campaigns remain in P...	None		Impact Scope: SINGLE_CUSTOMER Customers Affec...	4/10/2026 12:54am	

Write-back


HubSpot stores the final outcome and Engineering Issue ID so the support record remains the operational source of truth.

Governance

Slack communicates the handoff; Google Sheets records the decision, confidence, priority, and action for auditability.

A controlled 5-case validation

This was a portfolio-scale validation, not a production benchmark. The final set exercised all three routing outcomes plus both positive and negative known-issue matching behavior.



AUDIT LOG

	A	B	C	D	E	F	G	H	I	J	K	L	M	N
1	Timestamp	HubSpot Ticket ID	Customer Plan	Customer MRR	Support Agent	Product Area	AI Decision	AI Confidence	Similar Tickets Found	Customers Affected	Engineering Priority	Engineering Issue ID	Action Taken	
2	2026-10-04T00:30:50.914Z	214336681935	Growth	\$249	Maya	Contacts & Segmentation	RETURN	82		1			Returned to Support	
3	2026-10-04T00:31:09.011Z	214286807871	Enterprise	\$1,800	Maya	Integrations & API	ENGINEERING	72		1	High	SEN-1012	Escalated to Engineering	
4	2026-10-04T00:31:18.972Z	214336681937	Enterprise	\$2,200	Sofia	Contacts & Segmentation	HUMAN_REVIEW	65		1			Sent to Senior Support for Human Review	
5	2026-10-04T02:50:27.346Z	214287367119	Professional	\$699	Daniel	Integrations & API	RETURN	85	None	1			Returned to Support	
6	2026-10-04T04:54:48.059Z	214436896740	Professional	\$899	Sofia	Campaigns & Sending	ENGINEERING	78	KI-1005 — Campaign Processing Queue Degradation	1	High	SEN-1020	Escalated to Engineering	
7														

FINAL VALIDATION SET

Case	Decision	Confidence	Known issue / action
CSV import rejects valid date values	RETURN	82	Returned to Support
Webhook deliveries fail with Sendex 500 responses	ENGINEERING	72	SEN-1012
Intermittent contact count mismatch after segmentation refresh	HUMAN_REVIEW	65	Senior Support review
API returns 401 after API key was regenerated	RETURN	85	No known-issue match
Campaign remains stuck in processing after successful validation	ENGINEERING	78	KI-1005 / SEN-1020

WHAT THE TEST SHOWED

- All three operational routes were exercised in the final clean validation set.
- The negative control returned None for a ticket without a grounded known-issue match.
- The positive control retrieved and matched KI-1005 from the Sendex knowledge base.

Evaluation note: portfolio-scale validation only; these results should not be presented as production-scale model performance.

What this project demonstrates

The value is not simply that an AI agent can label tickets. It is the system around the model: explicit retrieval, structured outputs, deterministic routing, human review, and an auditable operational trail.

RAG WORKFLOW DESIGN

Explicit semantic retrieval across five knowledge files, chunk aggregation, and grounded known-issue matching.

SUPPORT OPERATIONS

Escalation quality, ownership boundaries, support-to-Engineering handoffs, and senior-support review.

AUTOMATION ENGINEERING

Make routers, filters, API write-backs, Airtable issue creation, Slack notifications, and audit logging.

GOVERNANCE & OBSERVABILITY

Deterministic priority rules, structured AI output, confidence capture, and a clean validation trail in Sheets.

FINAL RESULT

A complete AI-assisted support escalation workflow

The system retrieves relevant knowledge, classifies the ticket into one of three controlled outcomes, creates Engineering handoffs when warranted, writes decisions back to HubSpot, routes Slack notifications, and preserves an auditable record in Google Sheets.

PROOF OF IMPLEMENTATION

HubSpot write-back

RAG match

Airtable issue

Slack routing

Audit log

The attached screenshots show each of these system states in the finished portfolio build.

Jerry J. Hudson