PROJECT 043

Workflow Architecture & Future Roadmap

AI Customer Support Resolution System | Modular architecture, operational controls, and production evolution

Architecture Overview

Project 043 uses four specialized n8n workflows connected through structured handoffs and persistent PostgreSQL state. Immediate ticket handling is separated from scheduled support operations, allowing each workflow to remain focused, testable, and independently maintainable. AI is used for language understanding, knowledge retrieval, and response generation; deterministic rules control verification, routing, status changes, SLA behavior, and escalation.

Workflow Architecture

Workflow	Responsibility	Primary Output / Handoff
WF01 - Intake & Classification	Gmail intake, normalization, ticket ID, AI classification, priority, PostgreSQL storage, safety gate, case routing.	Human review, order investigation, or knowledge-resolution route.
WF02 - Customer & Order Investigation	Validates request, queries customer/order data, verifies customer, order, and ownership, then evaluates investigation result.	Verified resolution context for WF03 or human/exception handoff.
WF03 - AI Resolution & Knowledge Engine	Validates context, retrieves support policy through vector search, generates structured AI resolution, applies safety decision.	Customer response + response_sent, or Slack review + pending_human_review.
WF04 - Follow-Up, SLA & Closure	Scheduled ticket checks, SLA evaluation, escalation, customer follow-up, escalation recording, and case closure.	Operational alerts, follow_up_sent, or closed lifecycle state.

System Layers & Data Flow

Layer	Components	Function
Channels	Gmail, Slack	Customer communication and staff escalation.
Orchestration	n8n	Triggers, routing, sub-workflows, scheduling, validation, integrations.
Intelligence	LLM + Structured Output	Classification and policy-grounded response generation.
Knowledge	Vector Store + Gemini Embeddings	Semantic retrieval of relevant support policy.
Operational Data	PostgreSQL	Persistent tickets, customers, orders, statuses, and lifecycle state.
Control	Rules + Human Review	Verification, safety gates, SLA thresholds, escalation, closure.

End-to-end: Gmail -> WF01 -> PostgreSQL -> WF02 -> WF03 -> Gmail / Slack -> PostgreSQL -> WF04 -> escalation / follow-up / closure.

Future Roadmap

Phase	Recommended Evolution
1 - Production Hardening	Add audit-event tables, explicit last_escalated_at / follow-up counters, retry handling, idempotency, error alerts, credential/secrets controls, backups, and deployment monitoring.
2 - Omnichannel Support	Extend intake to web chat, forms, WhatsApp, CRM/help-desk inboxes, and unify channel history under one customer/ticket identity.
3 - Knowledge Operations	Introduce policy versioning, metadata filters, document owners, review dates, freshness checks, and retrieval-quality monitoring.
4 - Analytics & Management	Dashboards for ticket volume, category mix, SLA compliance, resolution time, escalation rate, automation rate, and human handoff rate.
5 - Controlled AI Agents	Add approved tools for replacements, refunds, order updates, or CRM actions with authorization rules, policy checks, human approval, and auditable tool use.

Recommended Production Design Direction

- Keep identity, authorization, monetary decisions, order changes, and lifecycle state under deterministic controls.
- Use AI where flexibility adds value: language understanding, retrieval, summarization, classification, and response drafting.
- Preserve PostgreSQL as the shared source of truth so workflows can scale independently without becoming one giant automation.

Architecture principle: AI operates inside a controlled support process - verified data and business rules determine what the system is allowed to do.