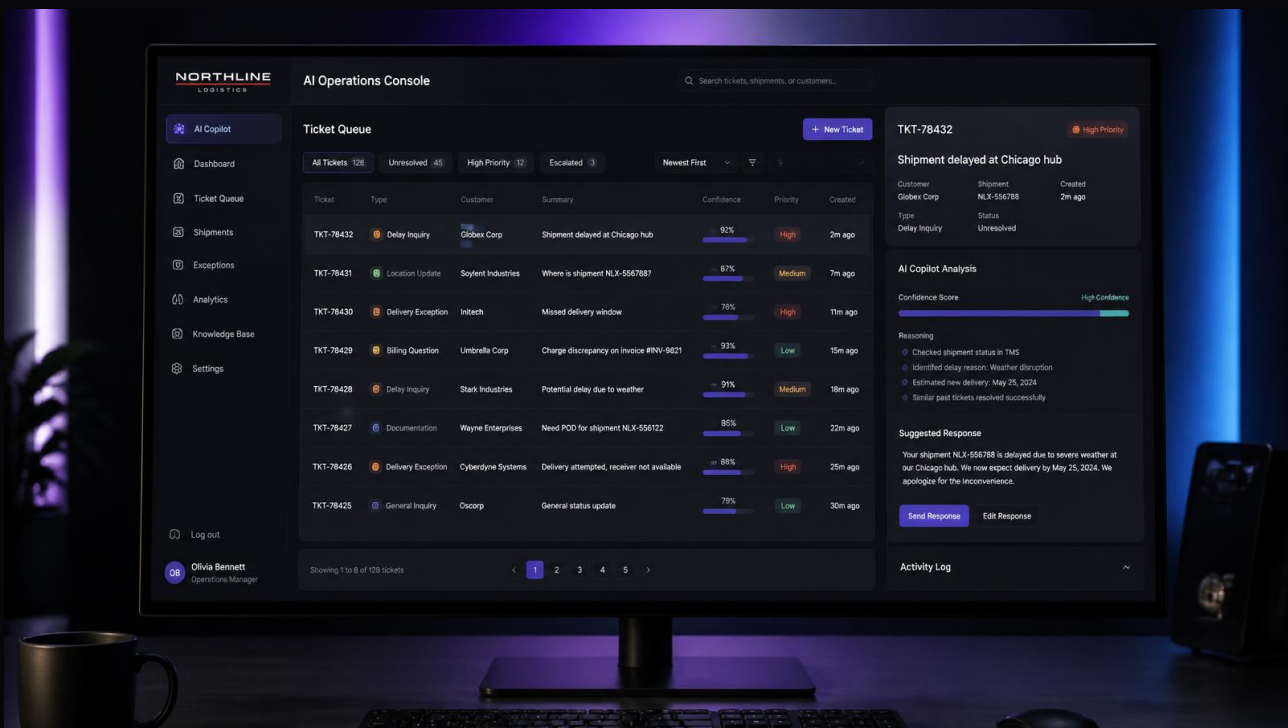


Meridian Logistics — Operations AI Agent

An AI agent that reads shipment exceptions, drafts the customer update and books the recovery — inside the tools the team already uses.



Meridian Logistics · Freight & Logistics
AI Agent Development · 10 weeks · Mar — May 2025

Overview

CLIENT

Meridian Logistics

INDUSTRY

Freight & Logistics

ROLE

AI product partner

DURATION

10 weeks · Mar — May 2025

SERVICES

AI strategy, Agent engineering, Evaluation harness, Change management

TECHNOLOGY

TypeScript, OpenAI + Claude, Postgres + pgvector, Temporal, Zendesk API

The client challenge

Meridian's exception desk handled roughly 1,900 delayed or damaged shipments a month. Median time to the first customer update was 4 hours 20 minutes, most of it spent reconstructing context across four systems. Two experienced coordinators had resigned citing the workload.

Research

Before proposing a model we mapped how the best coordinator on the floor actually resolves an exception.

- Task analysis with 6 coordinators across 40 recorded resolutions
- Classification of 12 months of tickets into 9 exception archetypes
- Risk review with legal and insurance on what an agent may never decide alone
- Baseline measurement of handling time, accuracy and escalation rate

Strategy

Autonomy was granted per archetype and per freight value, not switched on globally.

- Tiered autonomy: draft-only, approve-to-send, or fully autonomous by risk band
- Grounded retrieval over historical resolutions, carrier terms and SLA documents
- Deterministic tool calls for anything that moves money or freight
- Rollout by archetype, each gated on a passing evaluation run

Experience Design

The interface is built for a coordinator reviewing 30 decisions an hour, not for demoing a chatbot.

- Queue view ranked by SLA risk, with the agent's proposal already drafted
- Every claim traceable to its source document in one click
- Single-key approve, edit or escalate
- Confidence and reasoning surfaced without burying the action

Engineering

Durable workflows, hard guardrails and continuous evaluation in production.

- Temporal workflows for retries, timeouts and long-running recoveries
- pgvector retrieval with recency and carrier-specific weighting
- 480-scenario eval suite blocking any deploy that regresses accuracy
- Full audit log of prompts, tools, outputs and human overrides

Final Showcase

The agent runs inside Zendesk and the TMS — no new tool for the team to learn.

- Console: exception queue, reasoning trace, one-key actions
- Zendesk: drafted replies attached to the existing ticket
- Ops dashboard: autonomy rate, accuracy and savings by archetype

Deliverables

- Production agent with tool access to TMS, Zendesk and email
- Retrieval layer over 9 years of resolved exceptions
- Evaluation harness with 480 graded scenarios
- Human-in-the-loop console with full audit trail

Results

68%

Exceptions resolved without a human

-81%

Time to first customer update

+29 pts

Customer satisfaction score

0


Unreviewed actions on high-value freight

Client testimonial


“It behaves like a careful new hire who read nine years of our tickets. The team trusts it because they can see exactly why it said what it said.”

Daniel Vos — Director of Operations, Meridian Logistics

Gallery



Agent console with exception queue



Reasoning trace and human review panel

Conclusion

Meridian Logistics came to DEZMOLAB with a business problem, not a design brief. We diagnosed it, shipped against it and measured the outcome — and the work is still running in production today.

Work with DEZMOLAB

We design, build and grow digital products for startups, businesses and enterprises. If you have a problem shaped like this one, we should talk.

hello@dezmolab.com

dezmolab.com