



Prifysgol Cymru
Y Drindod Dewi Sant
University of Wales
Trinity Saint David

Automating Business Email Communication: An AI-Driven Response System Built in n8n

Kieran Brown | Supervisor: Nitheesh Kaliyamurthy | Deployed at SHIPMAX Ltd, UK



School of
Applied Computing

ABSTRACT AND AIM

An AI-powered email response system was designed, built and deployed live at SHIPMAX Ltd, a UK logistics consultancy receiving around 100 emails per day. Using n8n, Claude Haiku 4.5 and a RAG architecture, the system classifies emails and generates professional draft replies grounded in company knowledge. All drafts are reviewed before sending. The project proves a real SME can deploy effective AI without enterprise budgets.

Aim: To create an AI-powered email response system using n8n to ensure accurate, professionally toned draft replies in a live business environment.

- Review literature on AI, NLP and workflow automation to inform design decisions
- Justify methodology including a project plan and risk register
- Design and implement the system integrated with the live SHIPMAX Gmail inbox
- Test across four defined levels using real emails during live deployment
- Reflect critically on the project process and skills developed

1. INTRODUCTION

SHIPMAX Ltd is a UK logistics consultancy whose founder was the sole operator managing approximately 100 emails per day. Response times were inconsistent and significant time was being lost to routine correspondence that could realistically be automated.

"AI-assisted tools increased worker output by 14% on average, rising to 34% for less experienced workers in communication-heavy roles." — Brynjolfsson et al., 2023

ONS projected AI adoption among UK firms would double from 9% to 22% between 2023 and 2024. McKinsey found one third of organisations were already using generative AI in at least one business function by 2023.

3. METHODOLOGY

A **post-positivist, deductive** philosophy was adopted using a **development-based strategy** with a cross-sectional time horizon, structured using the Research Onion framework. A mono-method approach was used with all data collection organised around the development and evaluation of the system itself.

Data collection combined unit and system testing against defined pass/fail criteria, qualitative usability feedback from the SHIPMAX founder, and quantitative metrics covering email processing counts, draft quality classifications and pipeline completion timings.

2. LITERATURE REVIEW

AI in the Workplace. McKinsey reported that by 2024, 72% of organisations were using AI in at least one business function, up from 55% the previous year. Productivity gains are most pronounced in communication-heavy roles. Brynjolfsson et al. found AI-assisted communication tools increased output by 14% on average, with the largest gains among less experienced workers. For single-operator SMEs with no team to distribute workload, this represents a significant opportunity.

Large Language Models and NLP. The Transformer architecture introduced by Vaswani et al. (2017) represented a turning point in natural language processing, enabling models to process context at scale. Claude Haiku 4.5 was selected for its speed, strong instruction-following capability and cost-effectiveness. Anthropic's Constitutional AI training emphasises outputs that are safe, reliable and honest.

Retrieval-Augmented Generation. Lewis et al. (2021) demonstrated that combining a pre-trained generative model with a retrieval system produces significantly more accurate and grounded outputs than a language model on training data alone. Without retrieval, a general-purpose model has no knowledge of the specific company, its tone or its services. The RAG architecture directly addresses this by retrieving SHIPMAX-specific knowledge at runtime before generating each response.

Workflow Automation. n8n was evaluated against Zapier and Make. Its open-source licence, self-hosting capability and node-level control made it the strongest choice for a business handling sensitive client correspondence. Existing AI email tools including Gmail Smart Compose and Microsoft Copilot lacked the customisation and business-specific grounding required.

"Without retrieval, a language model has no knowledge of the company it represents. RAG directly addresses this." — adapted from Lewis et al., 2021


4. DESIGN AND IMPLEMENTATION

A seven-node n8n workflow integrates directly with the live SHIPMAX Gmail inbox. The Claude AI agent uses Conversation Memory, RAG-based Knowledge Base Retrieval, and a Response Structure Parser enforcing JSON output. All drafts await human review.

- Gmail Trigger** polls the inbox every minute for new unread emails
- Get Message** retrieves full email content including sender, subject and body
- Extract Email Data** normalises content into clean structured data via JavaScript
- Draft Response with Claude** classifies the email and generates a grounded draft using RAG
- Mark as Read** prevents the same email being processed twice

WORKFLOW: SUCCESSFULLY EXECUTED

All nodes show green ticks confirming successful execution. 1,536 knowledge base chunks retrieved via OpenAI embeddings during RAG retrieval.



5. TESTING

- Unit Testing:** each node validated in isolation against defined pass/fail criteria
- System Testing:** end-to-end workflow tested on real emails. Pipeline under 20 seconds every time
- Usability Testing:** founder assessed draft quality and tone. Feedback consistently positive
- Acceptance Testing:** all five SMART objectives met in full

6. RESULTS

The system ran live on the SHIPMAX Gmail inbox for several days, processing emails every minute.

214

Emails processed during deployment

86

Draft responses generated

14h

Manual drafting time saved

40%

Drafts sent with zero editing

0

Drafts discarded as unsuitable

<20s


Full pipeline completion

12 to 15 3PL Partners

Onboarded the month after deployment — direct commercial impact


DRAFT QUALITY OUTCOMES (N=86)

Over 85% of drafts were sent untouched or with only minor edits. Zero were discarded.



TIME SAVED: MANUAL VS AI-ASSISTED

Manual drafting estimated at 499 minutes. AI-assisted took 28 minutes, saving ~14 hours.



7. CONCLUSION

This project proves a practically valuable and ethically responsible AI email response system can be deployed within a real SME using open-source tooling, without enterprise budgets. The RAG architecture ensured responses were grounded in company-specific knowledge. Human oversight was preserved throughout and no email was ever sent without founder approval.

The three new 3PL partners onboarded following deployment demonstrate commercial value well beyond operational efficiency, directly linked to time freed up by the system.

REFERENCES

- Brynjolfsson et al., Generative AI at Work, NBER, 2023
- Lewis et al., Retrieval-Augmented Generation, arXiv, 2021
- McKinsey, State of AI in Early 2024, 2024
- Vaswani et al., Attention Is All You Need, NIPS, 2017
- ONS, AI adoption in UK firms, Mar 2025
- Saunders et al., Research Methods for Business Students, 8th ed., 2019
- Hossain, AI in Small Enterprises, IJFMR, 2024
- Jauhiainen, Effective Human Oversight in AI, UTUPub, 2025