



AI IMPLEMENTATION REPORT

Northbrook & Co.

AI-assisted product content at scale

Prepared for: Northbrook & Co.

Prepared by: AgentAya

Date: June 2026

Confidential. Prepared exclusively for Northbrook & Co.

1. Context Summary

This report documents the first AI implementation delivered to Northbrook & Co., covering the period from initial scoping in early May 2026 through handover in early June 2026. It builds directly on the AI Strategy Brief produced for Northbrook in May 2026, which identified five priority opportunities across the business and recommended starting with AI-assisted product content at scale (opportunity 4.1).

The implementation has been built, tested with the content team, and is now in production use. This document captures what was built, how it is operated day-to-day, what early results have looked like, and what we recommend as the next implementation cycle.

Audience for this document: Sarah Eltham and Mira Yates as the operational owners, plus any future external implementer or hire who may pick up subsequent opportunities from the strategy brief.

What was delivered

A structured prompt library, a documented workflow for product description drafting, a lightweight review and approval flow integrated into Northbrook's existing Notion workspace, and three weeks of supported rollout across Mira's team. The system is now used as the default starting point for all new product descriptions and is being extended to email campaign copy.

2. Current Technical Environment

This section describes the technical and operational environment as it stood at the start of the implementation, plus what changed during delivery. It is the reference point for any future implementation work.

2.1 Tools and platforms

E-commerce platform	Shopify (current plan: Advanced). All product data, customer accounts, and online sales reporting lives here
Email marketing	Klaviyo, integrated with Shopify. List size approximately 28,000 subscribers
Wholesale operations	Custom Airtable base, two years old. Standalone, not connected to Shopify
Internal collaboration	Notion (workspace), Slack (communication), Google Workspace (email, docs)

AI tools introduced	Claude (paid team plan, 5 seats), ChatGPT (existing free use by Mira), Notion AI (enabled but underused before implementation)
Customer service inbox	Shared Gmail inbox routed to four team members across retail, online, and wholesale
POS for the shop	Shopify POS, integrated with online inventory

2.2 Data sources used in this implementation

The product content workflow draws on three data sources, two of which already existed and one of which was built during the implementation:

- **Existing Shopify product data:** titles, attributes, supplier metadata, existing description history. Read-only access from the workflow.
- **Existing Northbrook brand guidelines:** a 6-page Notion document covering tone, voice, and product description structure, lightly updated during the implementation to clarify a few edge cases.
- **New prompt library:** a Notion database built during week 1 of the implementation, containing structured prompts for product descriptions, email subject lines, social captions, and seasonal campaign copy. Each prompt is versioned and includes example outputs.

2.3 Infrastructure and access

No new infrastructure was required. The implementation runs entirely inside tools Northbrook already pays for (Shopify, Notion, Klaviyo) plus the new Claude team plan. There is no hosted code, no database, and no external SaaS dependency beyond the AI providers themselves.

Five Claude team plan seats were provisioned: Mira (content lead), the two members of her team, Sarah (read access and oversight), and one shared seat for the wholesale team to start using AI for follow-up emails (a small extension of opportunity 4.1 into 4.3 territory, agreed mid-implementation).

2.4 Internal technical capability

Northbrook has no in-house developer or data analyst. The implementation was therefore designed to be operable entirely through Notion and Claude's web interface, with no command-line tools, no API integrations to maintain, and no custom code. This is a deliberate constraint: any future addition to the workflow must remain inside these boundaries unless Northbrook decides to bring in technical support.

Mira has emerged as the de facto AI operations lead during the implementation. She has good general technical sense, understands what Claude can and cannot do, and is comfortable

maintaining the prompt library. Sarah has explicitly endorsed this remit as a permanent extension of Mira's role.

2.5 Known constraints

- Claude's outputs need human review every time, particularly for tone and brand voice. The team has standardised on a two-pass review (writer + editor) before any AI-drafted content goes live.
- Notion AI was evaluated as an alternative drafting tool but produces less brand-aligned output than Claude with the structured prompts. Notion AI is retained for short-form summarisation tasks only.
- Shopify's content fields have a 5,000-character soft limit for product descriptions. The prompts have been calibrated to produce content within 1,200-1,800 characters, well within this limit.

3. Implementation Specification

This section captures the detail of what was built, in the form a future implementer or auditor would need to understand and extend the work.

3.1 The product description workflow

Opportunity	AI-assisted product content at scale (Strategy Brief opportunity 4.1)
Business requirement	Reduce the time taken to write product descriptions from an average of 35 minutes per product to under 10 minutes, while maintaining the consistency and voice that distinguishes Northbrook from competitors. Scale across 60-80 new products onboarded per quarter.

Current process (post-implementation)

When a new product is added to Shopify, the buying team enters basic attributes (title, category, supplier, key features). Mira's team picks up the new product in a daily Notion view that lists products awaiting description. The content writer opens the relevant prompt from the prompt library, pastes the product attributes into Claude, and reviews the draft output. The draft is edited (typically 2-4 minutes), checked by a second team member for voice consistency, and pasted back into Shopify.

End-to-end time per product is now averaging 8 minutes, down from 35 minutes. The team has retained the same quality standard: every description still goes through human review, but the starting point is no longer a blank page.

The prompt library structure

The prompt library is a Notion database with the following schema:

- **Prompt title** (e.g. "Product description, lighting category, mid-price")
- **Category** (lighting, ceramics, textiles, furniture, seasonal)
- **Use case** (product description, email subject, social caption, etc.)
- **Prompt text** (the actual prompt sent to Claude, with placeholders for product-specific attributes)
- **Example output** (a real example from a recent product to anchor the expected tone and length)
- **Version and last updated** (tracking changes over time)
- **Owner** (the team member responsible for keeping that prompt current)

The current library contains 14 prompts. We expect this to grow to 25-30 over the next 6 months as Mira's team extends the system to new content types and discovers refinements through real use.

Data requirements

The workflow requires no new data sources. All inputs are pulled from the existing Shopify product record and the existing Notion brand guidelines.

Dependencies

- Claude team plan (5 seats, billed monthly)
- Notion workspace (already in place, no plan change required)
- Continued maintenance of the Notion brand guidelines document by Mira
- Quarterly review of the prompt library, scheduled in Mira's recurring calendar

Estimated complexity

Low. The system is operable by anyone on Mira's team after one hour of onboarding. We documented the onboarding flow as a 4-page Notion page that new team members read before their first session.

Open questions

- Whether to extend the prompt library to cover Meta and Google ad copy. We started a pilot for this in week 3 but did not formalise it. Recommended next step in Section 5.
- Whether to integrate any part of this into Shopify directly through an app or API. Not recommended now (see Section 4), but worth revisiting in 12 months.

3.2 The wholesale email assistant (small extension)

Opportunity	Wholesale customer outreach drafting (small subset of Strategy Brief opportunity 4.3)
--------------------	---

Business requirement	Reduce time spent on wholesale follow-up emails for slow-paying or quiet accounts. Improve consistency of tone across the wholesale team.
-----------------------------	---

Current process

The wholesale manager uses the shared Claude seat with a small set of prompts for follow-up emails. The prompts include placeholders for account name, order history summary, and the specific reason for the outreach. Output is reviewed and personalised before sending. Average time per email is now 3-4 minutes, down from 10-15 minutes.

Why this extension was added

Mid-implementation, the wholesale team requested access to the same tooling for their own use. This was a small extension to scope and added no incremental cost beyond enabling the existing shared seat. We agreed to formalise it as part of this handover rather than wait for a full implementation of opportunity 4.3.

This is not the full opportunity 4.3. The full opportunity (wholesale customer prioritisation and segmentation) requires building a scoring model on the Airtable data and is recommended as the next implementation cycle. See Section 5.

4. What We Did Not Build

It is worth being explicit about what was considered and consciously left out of scope, so that future decisions are informed.

4.1 No Shopify integration

We considered building a Shopify app or using the Shopify API to push AI-drafted descriptions directly into product records. We chose not to, for three reasons. First, the human review step is genuinely necessary at this stage; automating the path from AI output to live product would remove the editing pass that keeps quality high. Second, Northbrook has no in-house engineering capability to maintain an app or API integration. Third, the current copy-paste workflow takes seconds and is not the bottleneck.

We recommend revisiting this in 12-18 months if the team has built confidence in the AI output quality and is considering scale that would justify the engineering investment (for example, a wholesale catalogue build-out or seasonal volume spikes).

4.2 No analytics integration

We did not build any tracking of which AI-drafted descriptions perform best on Shopify (conversion rate, time on page, add-to-cart rate). This was deliberate scope reduction to keep the first implementation small. Northbrook’s existing Shopify analytics show product-level performance but cannot distinguish AI-assisted from manually-written descriptions.

If Northbrook wants to test the hypothesis that AI-assisted descriptions convert as well as or better than manually-written ones, the cleanest way is a simple A/B test on a single product category over 6-8 weeks. This is a sensible mini-project for the next cycle but was out of scope for the initial implementation.

4.3 No social media or community automation

Northbrook's social presence remains a deliberate manual exercise: Sarah and Mira both feel the personality of the channels would suffer if it became templated. We respected this view and limited the scope to a small set of prompts for occasional caption assistance, used only when the team chooses to.

5. Prioritisation and Sequencing for Next Cycle

The strategy brief identified five opportunities. With opportunity 4.1 now implemented and a partial extension into 4.3, the recommended sequence for the next two implementation cycles is as follows.

Next implementation: Customer service triage (Strategy Brief opportunity 4.2)

Recommended scope: connect the shared customer service Gmail inbox to a lightweight AI workflow that classifies incoming emails by intent (order query, returns, product question, wholesale enquiry) and drafts a suggested response. Drafts appear in the inbox or in a queue for human approval before sending.

Estimated time to implement: 3-4 weeks. Estimated complexity: medium. Estimated time saved across the customer-facing team: 8-10 hours per week. This is the natural next step because it directly benefits the same team that is already using the content workflow, and the technical pattern (prompts plus human review) is the same.

Following cycle: Full wholesale prioritisation (Strategy Brief opportunity 4.3)

Build the scoring model on the existing Airtable wholesale data, identify accounts by segment, and produce tailored outreach for each segment. This is a bigger piece of work (estimated 6-8 weeks) but has the highest revenue impact of the five opportunities and builds on the small extension already running.

Hold for later

Marketing attribution (opportunity 4.4) and buying intelligence (opportunity 4.5) are both worth doing but should not be started until the first two are bedded in. Attempting all three at once would create more disruption than value. Recommended review point: 6 months from now.

6. Notes for the Implementer

These observations are for whoever picks up the next implementation cycle, whether that is AgentAya again or someone else.

Stakeholder dynamics

Sarah is the decision-maker but is not heavily involved in day-to-day operations of the AI workflow. She wants quarterly reviews showing impact, but does not want to be brought into operational decisions. Mira has full delegated authority for content-related AI work and should be treated as the primary client contact for any extension in that area.

The wholesale manager is more conservative about AI and was initially sceptical of the small extension to her team's workflow. She has since become an advocate after seeing the time savings. Future work touching the wholesale function should include her early and explicitly, not as a stakeholder to be managed but as a partner in design.

Cultural context

Northbrook has a strong brand-first culture. Decisions are weighed against "does this make the brand stronger or weaker?" rather than purely against efficiency. This means any automation that risks the brand voice will be resisted, regardless of time savings. The successful implementation of opportunity 4.1 was partly because the team retained creative control through the review step. Future implementations should preserve this principle.

Budget posture

Northbrook is commercially disciplined but not budget-constrained. Sarah is willing to invest where the return is clear. She is wary of recurring software subscriptions that add up over time, so any new tool brought into the workflow needs a clear justification. The Claude team plan was an easy approval because the time savings made the maths obvious within the first week.

Communication style

Sarah prefers written updates over meetings. A monthly one-page email summary of what is running, what is being changed, and what is being considered works well. Meetings should be requested only when a decision is needed.

Practical considerations

- All implementation documentation lives in a dedicated Notion section called "AI Operations". Mira is the owner.
- Prompt library changes follow a simple versioning convention (date plus initials in the title) but there is no formal change control process. This is appropriate at current scale; revisit if the library grows beyond 50 prompts.

-
- The team has Slack channel #ai-experiments where new ideas are shared informally before being formalised into the library. Use this channel rather than email for ongoing collaboration.

About AgentAya

AgentAya helps small and mid-sized businesses navigate AI with clarity and honesty. We publish independent reviews of AI tools, run a free academy for business owners getting started with AI, and are developing an advisory practice focused on practical, hands-on AI implementation for companies that do not have dedicated AI teams.

You can find us at agentaya.com.