

## AI OPPORTUNITY &amp; ROADMAP

# Strategic Review — a sample.

An illustrative example of the board-grade review every K&E Studios AI engagement starts with: where AI actually moves the needle, what to do first, and what it's worth. Client details are anonymized.

**6**

INITIATIVES ASSESSED

**3**

"DO NOW" WINS

**~\$240k**

YR-1 VALUE MODELED

**<90d**

TO FIRST ROI

PREPARED FOR

**"Meridian" · mission-led DTC  
brand**

PREPARED BY

**N.O.V.A. · Studio Strategy**

ENGAGEMENT

**AI Opportunity Review**

# Executive summary

## THE BOTTOM LINE

Meridian doesn't have an "AI strategy" problem — it has a **where-to-start** problem. The team is stretched across support, content, and operations, and treats AI as a threat or a toy. In reality, **three unglamorous, low-risk applications** would return their cost inside a quarter and free roughly a full-time-equivalent of senior time. We recommend ignoring the hype categories entirely for now and shipping those three first — with a human in the loop on everything customer-facing.

## WHAT WE FOUND

- **Support is the bottleneck.** ~70% of tickets are repeats answerable from existing docs; first-response time is 14+ hours.
- **Content is the constraint on growth.** One person owns email, SEO, and social — the channel that drives the cheapest revenue is the most starved.
- **Data is everywhere and used nowhere.** Numbers live in five tools; decisions wait on a manual weekly export.
- **Fear is the real blocker.** No owner, no policy, and a worry that AI means layoffs — so nothing ships.

## WHAT WE RECOMMEND

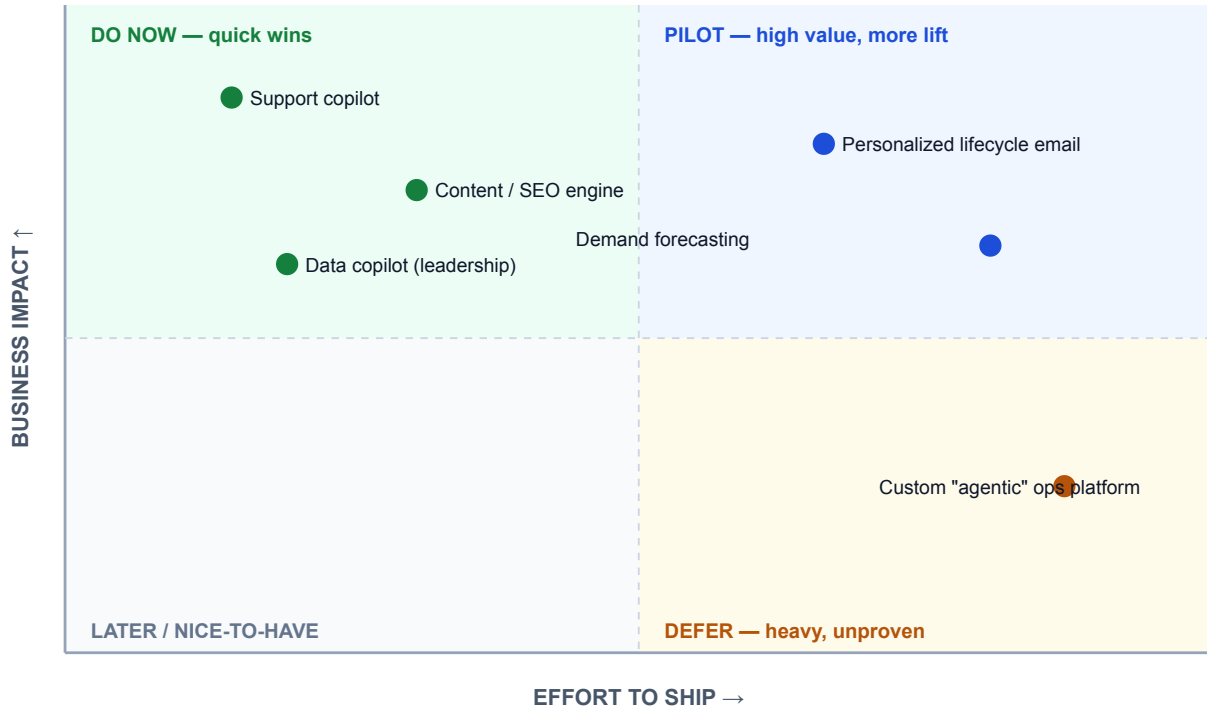
- **Ship a support copilot** (draft-and-approve, grounded in your own docs) — fastest, safest ROI.
- **Stand up a content engine** for email + SEO with brand-voice guardrails and human sign-off.
- **Give leadership a data copilot** — ask plain-English questions of your own numbers.
- **Name one owner + a one-page AI policy** so adoption is safe and accountable.

*Don't buy an "AI platform." Solve three real bottlenecks, keep a human in the loop, and let the wins fund the next move.*

SECTION 02

# The opportunity map

Every candidate initiative plotted by **effort to ship** against **business impact**. The play is the top-left: high impact, low effort. Start there; earn the right to the harder bets.



The three green initiatives share a profile: they automate work the team already does by hand, they're grounded in Meridian's own content/data (so they're accurate and safe), and a human approves anything a customer sees. That's why they go first.

SECTION 03

# Initiative scorecard

Each opportunity, scored on impact, effort, and payback — with a clear call.

INITIATIVE	IMPACT	EFFORT	PAYBACK	CALL	WHAT IT IS
<b>Support copilot</b>	High	Low	~4 weeks	<b>DO NOW</b>	Drafts replies from your own help docs; agent approves & sends. Cuts response time and repeat tickets.
<b>Content / SEO engine</b>	High	Low–Med	~6 weeks	<b>DO NOW</b>	Brand-voice email + blog/SEO drafts with a human sign-off step. Feeds the cheapest growth channel.
<b>Leadership data copilot</b>	Med–High	Low–Med	~6 weeks	<b>DO NOW</b>	Ask plain-English questions of your sales/ops data; kills the manual weekly export.
<b>Personalized lifecycle email</b>	High	Med	~1 quarter	<b>PILOT</b>	Segment + tailor flows by behavior. Real revenue lift, but needs the content engine live first.
<b>Demand forecasting</b>	Med	Med–High	2+ quarters	<b>LATER</b>	Inventory/demand modeling. Valuable, but data needs cleanup before it's trustworthy.
<b>Custom "agentic" ops platform</b>	Uncertain	High	Unclear	<b>DEFER</b>	The shiny option. High build cost, unproven for a team this size — revisit after the wins above.

**Sequencing logic:** the three "do now" items are independent, low-risk, and each frees senior time that funds the next. Pilot personalization once the content engine exists; defer anything that requires a big platform bet until the fundamentals pay for themselves.

SECTION 04

# The 90-day plan & the math

## Now WEEKS 1-4

- Name an AI owner + ship a one-page use policy
- Build the support copilot on your help docs
- Pick the metrics we'll judge success on

## Next WEEKS 5-9

- Stand up the content / SEO engine + brand-voice guardrails
- Launch the leadership data copilot
- Train the team; measure hours saved

## Later WEEKS 10-13+

- Pilot personalized lifecycle email
- Review results vs. baseline; decide quarter two
- Scope forecasting once data is clean

### THE MATH (REPRESENTATIVE)

**~1.0 FTE**  
senior time freed / year

**~\$240k**  
modeled year-1 value

**< 90 days**  
to first measurable ROI

Figures are illustrative for a ~\$6-10M, 12-20 person consumer brand: support hours saved (≈30 hrs/wk), content throughput (3-4x), and faster decisions. Your review uses your real numbers; we agree the baseline before we build so the ROI is yours to verify.

*This is the deliverable: a clear verdict, a map, a scorecard, a dated plan, and the math — then we build the "do now" list, not just hand you the slides.*

**HOW AN ENGAGEMENT RUNS**

- **Week 0 — Review (this document, for you).** Fixed-scope, all-inclusive. Yours to keep whether or not we build.
- **Then — Build.** We ship the "do now" wins on real infrastructure, with your team in the loop.
- **Optional — Embed.** Fractional AI leadership month-to-month, or an expert in your calls by the hour.

Prepared by N.O.V.A. · Studio Strategy for K&E Studios AI Consulting. **This is an illustrative sample** built to demonstrate the format and rigor of a real engagement — "Meridian" is fictional and all figures are representative. Your review is built around your business, your data, and your goals. [kestudios.dev/consulting](http://kestudios.dev/consulting)