A2 — Supporting Analyses Pack

Relay Pulse: the alarm layer. VSM diagnosis, THRIVE-led recommendation, governance and ESG check, KPI tree, and strategy cascade — the evidence behind every claim in the Explainer Video.

Live artifacts: KPI Dashboard — https://alexly0329-cpu.github.io/relay-pulse-dashboard/ · App Prototype — Pulse Console (Google AI Studio).

1. Five-Jobs Diagnostic (W1) — where Relay is failing

The job	Who / what does it	Working? Evidence	Where it fails
S1 Doing the work	3 DCs, product, CX teams	Yes — orders +11%, GM 57% stable, fulfillment shipped what it had in December	Only downstream of missing signals — starved of warning, not capability
S2 Coordinating	Nobody owns cross-function sync	No — pilot launched without Ops or Community informed	December: Marketing, Fulfillment, Community each correct alone; disaster together
S3 Running today	Monthly business review (FP&A + Mktg analytics)	Partly — corrects monthly, not daily	An intra-month demand turn is invisible until the next review
S4 Sensing tomorrow	No owner; market awareness anecdotal; R&D is a product backlog	No — no competitive-intelligence function	Missed Cadence economics, the Dec 10 turn, and Pacer churn drift
S5 Holding identity	Dana personally; values not operationalized	No machinery — Legal flag died 14 months in transit	Jules: unapproved tool, no review process existed to fail

One-sentence autopsy. Relay is failing because the jobs of sensing (S4) and coordinating (S2) are missing or disconnected from action, while identity (S5) has values but no machinery — evidence: eleven days from signal to expedite; fourteen months for a legal flag to reach no one.

Mindset Matrix. Relay today is reactive–narrow (12 scattered initiatives launched to satisfy "AI everywhere" — Resilient Recovery, the vanity posture). Relay Pulse moves the company to proactive–broad (Regenerative Transformation): AI as permanent sensing capability woven into how the company runs, while identity stays protected.

2. Systems Diagnosis (W2) — the coordination traps

Failing system. Not a starved S1 — the link between S1 units: S2 is missing and S4 is unowned. No function failed at its own job; the space between functions had no owner.

Coordination trap 1 — pilot × fulfillment. Marketing ran a conversion pilot (+22%); Fulfillment planned to baseline demand. Each needed: the campaign calendar as a data feed and a demand alert on day one — before launch, not at the weekly report.

Coordination trap 2 — campaign × community. Complaints landed on Pacers who did not know a campaign existed. Each needed: a pre-launch Pacer brief and a same-day trust alert when complaint volume turned.

December, traced. The signal existed Dec 10 (daily order feed). An actor existed the same day (merch planner with expedite authority). The shortest path between them is an anomaly alert routed to a named owner with a clock — hours to build against data Relay already owns, versus the $2.8M proposal that would not have built it at all.

THRIVE T. What must stay human: community relationships, CX conversations, merchandising judgment — the work carrying Relay's distinctive value. AI support (context, drafted options) makes that work more meaningful, not less.

THRIVE H. Hand the machines the watching: continuous monitoring of orders, tickets, forums. Freed capacity is spent on response quality and community investment.

AI SHOULD	AI SHOULD NOT
Watch every owned data stream continuously; detect anomalies; route alerts with owner + severity + clock; draft response options; log outcomes and learn	Talk to customers; generate community-facing content; auto-execute irreversible actions; run outside the tool registry

Completion. The company's work breaks at nothing. Its coordination breaks at signal handoffs (Marketing ↔ Ops ↔ Community). The cheapest structural fix is a defined anomaly-to-action path with a name attached — Relay Pulse.

3. Horizon Read (W3) — the symmetrical sensing failure

Sensing mechanism	Speed	Owner	Caught / missed last year
Daily sales feed	Real-time	Nobody	Missed the Dec 10 turn for 11 days
Weekly merch report	Weekly	Merch lead	Published Dec 16; read by 4 of 11
CX tickets tagged "stock?"	Daily	Support queue	No route upward — three tickets died
Competitor awareness	Never	Nobody	Anecdotal; missed Cadence CAC at ~half of Relay's
Employee survey verbatims	Annual	HR	Collected, never mined

The Cadence threat, classified. Not cost, and not model quality — models are rentable by anyone. Cadence wins on organization: demand signals reach decision-makers in hours; merchandising, fulfillment and marketing share one picture; a named team evaluates AI capability continuously.

Creativity-crusher patterns at Relay. Fear of being wrong (Legal's early flag died — early warnings are punished by silence); attachment to orthodoxy ("the community will always hold" went unfalsified while

Pacer churn rose); decisions by seniority (the board's "AI everywhere" defined what counted as a signal, producing 12 vanity initiatives).

The two gaps. Reskilling gap (R): alert receivers — merch, ops, community leads, and the leadership layer first — must learn to run on an instrument panel: acknowledge, decide, log. Innovation gap (I): AI-amplified pattern detection over owned unstructured data (transcripts, reviews, forums) feeding product and community imagination, with a named radar team looking.

4. Vanity–Viable Matrix (W4) — three use cases, three verdicts

Row	UC1 · VantagePoint engine	UC2 · Relay Pulse (ours)	UC3 · CX Copilot (phase 2)
What it does	12 pre-built personalization modules; auto-generates customer content	Detects anomalies across owned data; routes alerts to named owners with clocks	Drafts context + replies for human CX agents; human always sends
Metric it moves	Clicks, sessions, opens (leading, unlinked to value)	Signal-to-action latency (leading) → stockout days, returns, LTV:CAC (lagging)	First-response time (leading) → CSAT, retention (lagging)
Vanity test	Fails — clicks can rise while margin falls and returns spike	Alert volume could be gamed → guard-rail: % acted within SLA (81% in pilot)	Deflection could rise while satisfaction falls → guard-rail: CSAT + re-open rate
Costs	$2.8M over 3 years + data rights	Data already owned; small build team; weeks to pilot	Marginal — rides Pulse infrastructure
Risks	Identity: module #7 auto-content; §11.3 trains competitors on Relay data	Low — internal only; never customer-facing	Medium — must stay behind the human (Founder's line)
VSM subsystem	S1 marketing output only	S4 + S2, with an S5 ethics channel	S1 support, augmented
Autonomy level	4–5, ungoverned	2–3: AI recommends and drafts; humans commit	2: AI suggests; human sends
Algedonic check	None anywhere in the proposal	Pulse IS the algedonic channel; owner: CDOO	Escalation path to human lead built in
Verdict	REFUSE as scoped (renegotiate: strike §11.3, disable #7, margin KPIs — then pilot at most)	BUILD — pilot live on the December replay	PILOT after Pulse proves governance

5. CLAIM Remediation Grid (W5) — rebuilding System 5

The dead alarm. Signal sent: Legal's consent-gap memo, fourteen months before the incident. Missed by: everyone — there was no register, no owner, no clock; S5 had no intake. Next time: the ethics

playbook — one-step flag, case ID, named owner, severity clock, auto-escalation to the Founder and board risk committee if unacknowledged.

System 5 verdict. Relay's System 5 failed because it never built machinery to hear, judge, and enforce its own boundaries — identity lived in the Founder's voice instead of in process.

CLAIM	What failed	Structural fix	Owner	Working when…
C — question	Pod never asked "should we?" — no place to ask	One-step ethics flag for any employee or Pacer; raise-amnesty norm	S5 review chair	Raise-rate up in year 1; repeat classes down
L — human judgment	No approval or override gate existed	Pre-launch review gate for any AI touching people; approval required above autonomy level 2	Risk & Governance	Review coverage ~100%; shadow-tool detections → 0
A — align purpose	"The board said everywhere" substituted for purpose	Published charter lines; every deployment maps to mission before build	Founder's office	Zero deployments crossing charter lines
I — explainability	Nobody could say what the tool did or why	Provenance log for all generated content; every alert explains why it fired	Tech & AI Systems	100% outbound content carries provenance
M — visibility	18,000 undisclosed synthetic messages	Disclosure standard (EU AI Act Art. 50); internally visible AI register	Risk & Governance	Zero undisclosed-AI incidents

Charter extract (published). (1) We will never send synthetic words in a real person's voice without written consent. (2) No AI reaches a customer or community member without a named human approving the pathway. (3) Anyone — employee or Pacer — can trip the ethics alarm in one step, and it cannot be closed silently.

Compliance, privacy & ESG check

Standard	How the recommendation accounts for it
GDPR / privacy	Likeness and voice are personal data: consent registry for all ambassador content; DPIA per deployment; first-party data only — no scraping; the ~18,000 recipients identified and notified in remediation
EU AI Act — Art. 50	All synthetic content disclosed and provenance-tagged; the likeness/consent tripwire enforces it automatically; Pulse itself is internal decision-support (minimal-risk), documented in the register
FTC endorsement rules (US)	Ambassador content must reflect real endorsement — synthetic ambassador messaging prohibited by charter line 1
ESG — compute & energy	Deliberately light: statistical anomaly models + small batched classifier always-on; frontier LLM invoked only to name new clusters — negligible footprint vs. "LLM

	everywhere," reported in the AI register

6. 3-KPI Tree (W6) — the instrument panel, as built

Root sentence. Relay gets healthier when signals reach someone who can act before the customer feels the gap.

Row	KPI 1 · Signal-to-Action Latency	KPI 2 · Emerging-Issue Index (+ trust-weighted attrition pair)	KPI 3 · Flywheel Economics Guard-rail
Name + unit	Median hours, anomaly detected → owner action, by playbook (live: 14h)	0–100 composite from transcripts, reviews, forum text; paired with tenure-weighted Pacer churn velocity (live: 57)	LTV:CAC ratio + effective CAC incl. 24% returns waste (live: 2.81 · $70/$85)
Decision it drives	When latency or SLA slips, escalation rules and staffing change. Owner: CDOO	When a cluster crosses threshold, a response brief goes to the issue owner and the Pacer council. Owner: Head of Community & CX	When ratio trends to 2.5, spend rebalances toward community channels. Owner: CFO
Signal speed	Real-time; unacknowledged alerts escalate automatically	Daily scan; weekly index	Monthly, cohort-based
Data source	alerts log (structured)	support transcripts, review text, Pacer forums — the required unstructured sources	relay-financials cohort economics
Which of the five jobs	S2 coordinating + S4 sensing	S4 sensing + S5 identity	S3 running today
Vanity test / guard-rail	Alert volume could be inflated → paired with % acted within SLA (81%)	Sentiment theater / silent fade → paired with pre-churn behavior (attendance, posting, referral share)	CAC can fall by buying cheap low-LTV cohorts → paired with cohort LTV
E test	Sustains — latency means the same at 10× scale	Sustains — index recomputed as community grows; human judgment names the clusters	Sustains — ratio is scale-independent; threshold reviewed annually

> *Dashboard live at https://alexly0329-cpu.github.io/relay-pulse-dashboard/ — December replay estimates $11.08M revenue protected with Pulse on the Checkpoint line.*

7. THRIVE-led transformation roadmap

Pillar	Applied to Relay	Roadmap phase

T — Transformative Engagement	Community, CX and merch judgment stay human; Pulse hands them context and options, making the human work more meaningful	All phases — the design constraint
H — High Productivity	Machines take the watching (orders, tickets, forums); freed capacity goes to response quality and Pacer investment	Phase 0 (wks 1–4): demand playbook pilot on the Checkpoint replay + tool registry
R — Resilient Adaptability	Reskill the receivers: leadership first, then merch/ops/community leads — acknowledge, decide, log; monitoring roles become response roles	Phase 1 (wks 5–12): trust playbook + Pacer council integration + receiver training
I — Imagination & Creativity	Cluster insights feed product and community innovation; a named radar team looks outward at Cadence and the category	Phase 2 (Q2): ethics playbook full, verbatim mining, outward radar
V — Values & Ethics	CLAIM rail always on; charter lines published; Art. 50 + GDPR machinery automatic	Gate for every phase — nothing ships without review
E — Efficiency & Endurance	KPIs that sustain at 10×; ESG-light models; pilot→scale gates on hard margin evidence	Phase 3 (Q3+): scale decision + CX Copilot pilot gate

Adoption pathway credibility. Each phase has a named owner, a training cohort, and a kill criterion. Phase 0 is deliberately small and real: one product line, one playbook, four weeks — small and real beats broad and imaginary.

8. Choice Cascade (W7) — the strategy, linked

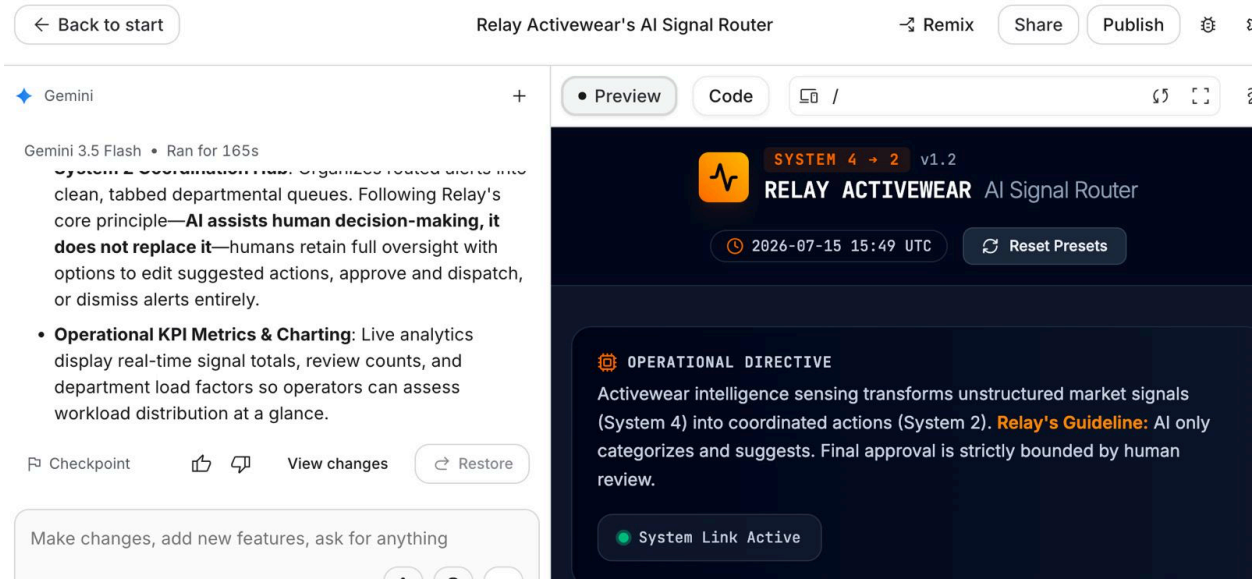
- **Winning aspiration.** Dana's, sharpened: the most human company in activewear, powered by the best operational AI in apparel — measured by health and trust, not applause.
- **Where to play.** Owned channels and the community: the signal-rich ground Cadence does not have. Latent demand evidence: the adaptive-running and Spanish-language communities (Jules's and Eddie's chapters) as under-served growth segments.
- **How to win.** Respond at machine speed while staying human at the boundary — belonging Cadence cannot copy, with the operational reflexes it already has.
- **What capabilities.** Owned first-party data + the Pulse sensing layer + consent-clean community data. Models are rentable by anyone; the data and the listening position are not. No vendor clause may leak them (the §11.3 lesson).
- **What management systems.** The three-KPI instrument panel, the CLAIM rail, the S5 review with clocks, and the monthly Pacer council — capability turned into governed, scalable practice.

Where we correct the Founder. The vault needs windows, not just walls: Pacer churn rose unseen inside the protected community — instrumenting its health is not renovating it. We keep the no-AI-touches-

customers line for now, and propose a Phase-3 evidence gate to revisit CX augmentation at autonomy level 2 with full disclosure — challenged openly, as asked, with the pilot data as evidence.

9. AI Use Appendix (complete per Hult policy)

Template — declare per artifact: tool used, purpose, prompts/scope, human review applied. [Team fills: e.g., Gemini in Google AI Studio for dashboard vibe-coding and CSV mock-data generation; Claude for analysis structuring and drafting support; all frameworks applied and verified by the team; all case figures traced to course artifacts.]



You are Relay Activewear's AI Signal Router. Your job is to convert raw signals sensed by System 4 (market intelligence) into actionable tasks for System 2 (department coordination).

For each incoming alert, output:

1. Category: Demand / Trust / Ethics
2. Responsible party: Merch Ops (R. Ito) / Community (P. Sharma) / Legal-Founder Office
3. Suggested action (one sentence, specific and actionable)
4. Needs human review before acting (Yes/No) — mark "Yes" for anything involving public communication, ethics, or legal risk

Stay aligned with Relay's principle: AI assists human decision-making, it does not replace it. Only classify and recommend — never issue a final directive. TimestampAlertCategoryResponsible PartySuggested ActionHuman Review2026-07-14 09:12West Coast repeat-purchase rate dropped 15% in one weekDemandMerch Ops (R. Ito)Check for pricing/inventory anomalies, intervene within 48hNo2026-07-14 14:35Spike in product-quality complaints on social media (40+ in one day)TrustCommunity (P. Sharma)Issue public response statement, open dedicated support channelYes2026-07-15 08:02Competitor Cadence social mentions surged 300% in one dayDemandMerch Ops (R. Ito)Assess need to adjust this week's marketing spend pacingNo2026-07-15 11:47AI chatbot gave misleading sizing recommendation to customersEthicsLegal/Founder OfficePause the AI feature module, manually review recommendation logicYes2026-07-15 16:20South China warehouse inventory-demand gap reached 22%DemandMerch Ops (R. Ito)Trigger emergency restock process, notify supply chain teamNo

Gemini 3.5 FlashRan for 165s

summarize_autoAction historychevron_right

Here are key actions taken for the app:

editEdited 6 files

metadata.json

check_circle

package.json

check_circle

server.ts

check_circle

src/index.css

check_circle

src/types.ts

check_circle

src/App.tsx

check_circle

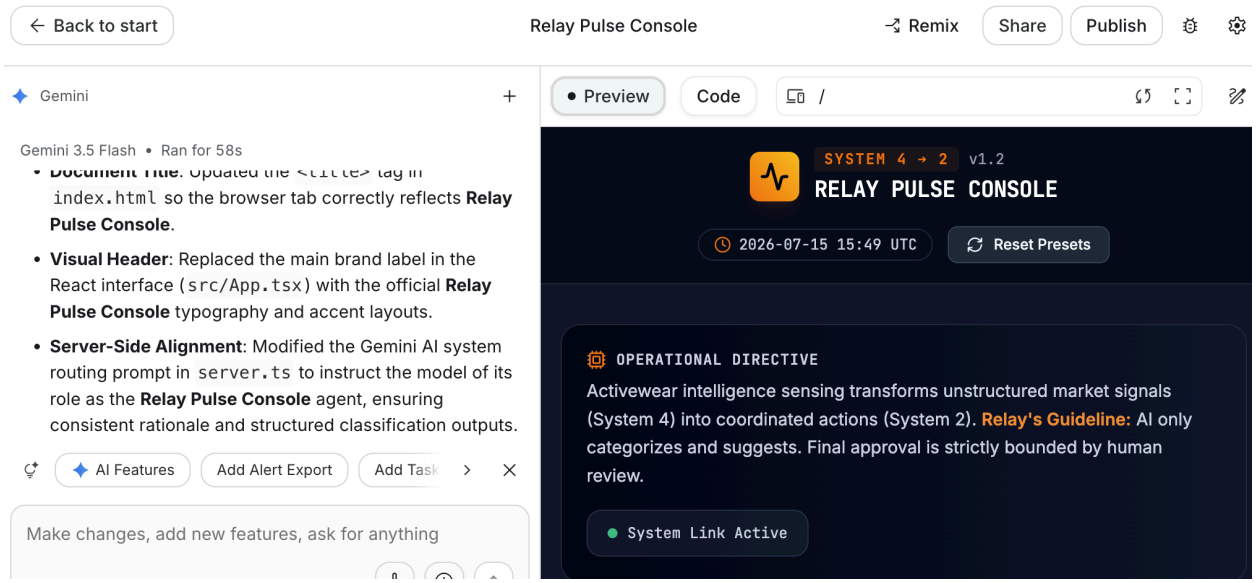
buildBuilt

I have built and configured **Relay Activewear's AI Signal Router** using a secure, full-stack Express + React architecture powered by the Gemini API (`gemini-3.5-flash`). The application converts unstructured raw market intelligence from System 4 into human-actionable directives for System 2.

Here is a summary of the design choices and capabilities:

- **Branding & Theme Selection**: Inspired by premium athletic apparel, the application is styled in a sporty, high-contrast dark theme using Slate and JetBrains Mono with vibrant neon orange and teal highlights.
- **System 4 Raw Signal Ingestor**: Users can write custom market intelligence signals or click interactive preset buttons to immediately inject new incidents into the active stream.

- **Dynamic AI Signal Routing**: The server-side Gemini API dynamically categorizes signals (Demand, Trust, Ethics), assigns the designated department owner (Merch Ops, Community, Legal/Founder), drafts a specific suggested action, and flags any potential legal/public risk for human review.
- **System 2 Coordination Hub**: Organizes routed alerts into clean, tabbed departmental queues. Following Relay's core principle—**AI assists human decision-making, it does not replace it**—humans retain full oversight with options to edit suggested actions, approve and dispatch, or dismiss alerts entirely.
- **Operational KPI Metrics & Charting**: Live analytics display real-time signal totals, review counts, and department load factors so operators can assess workload distribution at a glance.



这份文档已经有功能描述和截图了,可以直接用。按团队文档第9节的模板格式,帮你把这些整理成一段可以直接放进 AI Use Appendix 的正式文字:

App Prototype — Relay Pulse Console

Tool used: Gemini in Google AI Studio (Gemini 3.5 Flash)

Purpose: Built the App Prototype implementing the alert-routing layer for UC2 (VSM System 4 + 2, per team diagnosis), converting raw sensed signals into coordinated departmental actions.

Prompts/scope: System-side prompt instructed the model to act as Relay's AI Signal Router: classify incoming signals into three categories (Demand / Trust / Ethics), assign a designated department owner (Merch Ops, Community, or Legal/Founder Office), draft a specific suggested action, and flag items requiring human review. Model was explicitly constrained to categorize and suggest only — never issue a final directive.

Key features built:

- Dynamic AI Signal Routing — categorizes signals, assigns owners, drafts suggested actions, flags legal/public risk for human review
- System 2 Coordination Hub — organizes routed alerts into departmental queues; humans retain full oversight (edit, approve, dispatch, or dismiss)
- Operational KPI metrics — live signal totals, review counts, department load

Human review applied: Tested against sample alerts (repeat-purchase decline, social complaint spikes, competitor mention surges, AI misclassification, inventory gaps); team reviewed classification accuracy and routing logic before inclusion. Final routing/action decisions remain human-approved — autonomy level 2–3 per the Vanity-Viable Matrix, consistent with the charter principle "AI assists human decision-making, it does not replace it."

Gemini in Google AI Studio was used to build the Relay Pulse Console (App Prototype) — an alert classification and routing tool. Prompted with a system instruction defining Relay's four-signal categories (Demand/Trust/Ethics) and human-review gating logic; tested against 5 sample alerts derived from the case. All classification logic and routing recommendations were reviewed by the team before inclusion; AI output is advisory only — final routing decisions remain human-approved per the app's design (autonomy level 2-3).