

THE DAILY RUN • ISSUE N° 001

Should Jollibee Singapore go plant-based?

Six AI experts argued for thirty-three messages. Five frontier models were asked separately. They agreed on an answer worth exactly eighty cents, and caught an AI engine getting the halal facts wrong along the way.



SUBJECT **JOLLIBEE SINGAPORE**
MARKET **SINGAPORE • QSR**
PREPARED **27 JULY 2026**

METHOD **PLATFORM RUN**
CAPABILITIES EXERCISED **NINE**
PAGES **12**

SECTION 01

Executive summary

One brand, one morning, one decision put through every layer of the platform. The question was a live one: should Jollibee Singapore add plant-based options, and if so, how?

RECOMMENDATION

Introduce a single plant-based swap inside existing family bundles at a flat S\$0.80, opt-in, priced identically across outlets. Pair it with a visible five-second halal handling ritual at the counter. Treat it as a defence of the family bundle, not as a repositioning toward green lifestyle.

S\$0.80

RECOMMENDED UPCHARGE, ARGUED TO
ACROSS THREE PRICED OPTIONS

28 days

PROPOSED PILOT, THREE NAMED
OUTLETS, SINGLE PRICE

+0.22

NET SENTIMENT, THE ONLY POSITIVE
SCORE IN THE CATEGORY

Four findings

- **The commercial problem is not vegetarians.** It is the single flexitarian in a household who redirects the whole group to a competitor. That reframes the initiative from acquisition to retention, and changes which budget it belongs to.
- **The brand holds an unusual asset.** Across 778 public conversations it is third in share of voice but first in sentiment, and the two larger players are both in negative territory. The task is removing friction, not buying affection.
- **Operational trust is the constraint, not appetite.** Halal separation must be seen to be believed; the panel converged on visible handling over faster throughput.
- **One AI engine misstates the brand's halal certification.** Four of five answer correctly. The fifth repeatedly tells users the chain is not certified, which quietly removes an audience the family bundles are built to include.

BASIS OF THIS DOCUMENT

Every figure, quotation and image in this report was produced on the gimmefy platform between 15 and 26 July 2026, from a brand vault built on public material. Jollibee is not a client and has not reviewed this work. Jollibee® is a trademark of Jollibee Foods Corporation. Nothing here is an official brand communication.

SECTION 02

Contents and method

01 · Executive summary	02
02 · Contents and method	03
03 · Discovery: what the market says	04
04 · Discovery: what the machines say	05
05 · The decision and the options	06
06 · The expert debate	07
07 · Independent model review	08
08 · Campaign development	09
09 · Creative output	10
10 · Method notes and limitations	11
11 · Appendix: sources and full records	12

How this work was produced

gimmefy is AI infrastructure for marketing teams, organised in four layers. Each morning one brand is run through all four to verify the system end to end.

EXHIBIT 1 The four layers, and what each contributed to this issue

LAYER	WHAT IT IS	CONTRIBUTION HERE
01 Models	48 frontier models, routed by task	Five were used independently for the model review in section 07; others sat behind individual capabilities
02 Context	Brand vaults, pulses, memory	A vault for Jollibee Singapore, enriched by four environment reports, read by every capability downstream
03 Capabilities	Workflows, media, research, reports	Nine exercised: Beacon, Focus Group, Model Council, Playbooks, Studio, Skills, Visual, Video, Publish staging
04 Governance	Workspaces, roles, audit, white label	The entire run sat in one governed workspace with a durable, inspectable record of every execution

Source: gimmefy platform run records, 15–26 July 2026.

Route

The platform can be driven three ways: directly through its own interface, through Friday (a concierge that plans work from a single instruction and executes on approval), or through MCP from inside Claude, ChatGPT or Gemini. **This issue was produced through the platform interface**, with one campaign kit routed through Friday for comparison. A later issue tests the

SECTION 03

Discovery: what the market says

Before any work was briefed, four environment reports were run against the brand's public world. Their findings were written back into the brand vault, so every later output stands on them.

EXHIBIT 2 Share of voice and sentiment, Singapore QSR chicken category

BRAND	SHARE OF VOICE	NET SENTIMENT
McDonald's	27.2%	-0.16
KFC	26.3%	-0.23
Jollibee	20.3%	+0.22
Popeyes	15.4%	+0.14
Texas Chicken	10.7%	+0.09

Source: Beacon listening report, 778 public conversations across seven source families, frozen 17 July 2026.

IMPLICATION

Third in volume, first in warmth, and the two larger competitors are both under water. The strategic task is to remove reasons not to visit, rather than to manufacture affection the brand already has.

EXHIBIT 3 Brand claims scored against observed market perception

Product distinctiveness		85
Mainstream appeal		75
Digital convenience		55
Operational consistency		40

Source: Beacon landscape report, 17 July 2026. Scores indicate alignment between brand claim and observed market perception, 0-100.

The landscape report's own summary is blunter than the scores: product superiority is *"actively bottlenecked by operational friction."* The food wins; the queue loses. That finding shaped the operational half of the recommendation.

EXHIBIT 4 Findability audit, jollibee.com.sg

100	87	29
CONTENT	TECHNICAL	PERFORMANCE

Source: Beacon search audit, 17 July 2026. Full evidence set: approximately 150KB. On-page 80, authority 52 omitted for space.

SECTION 04

Discovery: what the machines say

Consumers increasingly ask an AI assistant where to eat before they ask a friend. Five engines were interrogated about the brand and each answer scored against the vault.

EXHIBIT 5 AI engine perception audit, five engines

DIMENSION	BRAND POSITION	ENGINE CONSENSUS	VERDICT
Positioning	An everyday family habit	A nostalgia-led occasional visit, constrained by queues	Drifted
Proof points	Generous portions, bundles saving up to \$10	Portions smaller and pricier per piece than rivals	Contradicted
Audience	Open to all; 100% halal certified	Four engines confirm. One states repeatedly that it is not certified	Contradicted

Source: Beacon GenAI audit, 17 July 2026. Engines: GPT, Claude, Perplexity, Gemini, Grok.

THE MATERIAL FINDING

An AI assistant is informing Muslim diners that a halal-certified restaurant is not certified. This is not a matter of tone or emphasis. It is a factual error that removes an entire audience from consideration, invisibly, at the moment of choosing where to eat.

For completeness: the engine in question is Claude, which is also the model used to compose this report. The error was verified rather than assumed, and is recorded here on that basis.

Remediation identified by the audit

- **Certification correction.** A structured knowledge-graph and content correction targeted at the engine carrying the error, with authoritative certification references.
- **Portion evidence.** A visual programme demonstrating bucket sizes and meal weights, to retrain the value association that engines currently dispute.
- **Convenience narrative.** Content weighted toward neighbourhood outlets and delivery to displace the queue-led framing.

Each remediation is routed by the audit to the capability that would execute it, rather than left as an observation.

SECTION 05

The decision and the options

Three structurally different options were priced against each other, not compared on preference. Each was argued for by a different discipline before being tested against brand economics.

EXHIBIT 6 Options assessment

OPTION	PRINCIPAL STRENGTH	DETERMINING WEAKNESS	ASSESSMENT
Free swap No upcharge	Highest trial rate; removes all financial friction	Margin dilution and cannibalisation of the highest-margin item	Rejected
\$S\$0.80 swap Opt-in, uniform	Protects bundle economics while neutralising the objection	Requires a 28-day window to demonstrate retention effect	Recommended
\$2.00 side Add-on item	No kitchen reconfiguration required	Does not substitute a main; the group objection remains live	Rejected

Source: Focus Group session, 15 July 2026, synthesised verdict. Full transcript in appendix.

Pilot design

- **Duration and scope.** Twenty-eight days across NUS, Pasir Ris Mall and Paya Lebar, at a single uniform price. Regional price variation was rejected on the grounds that perceived unfairness travels faster than the test data.
- **Operational condition.** Dedicated preparation with a visible five-second handling and sealing step at the counter, accepted as a deliberate trade of peak-hour throughput for demonstrable trust.
- **Digital condition.** Native integration into delivery bundle customisers rather than interruptive prompts, and no pre-selected default.

A brand built on warm, generous comfort cannot afford to optimize its way into customer frustration.

BRAND STRATEGIST, IN REJECTING A PROPOSED DEMAND TEST

SECTION 06

The expert debate

Six expert personas were selected and given a single priced decision to pressure-test, alongside five named Singapore customer profiles to argue through. A facilitator opened the session; a synthesiser produced the verdict.

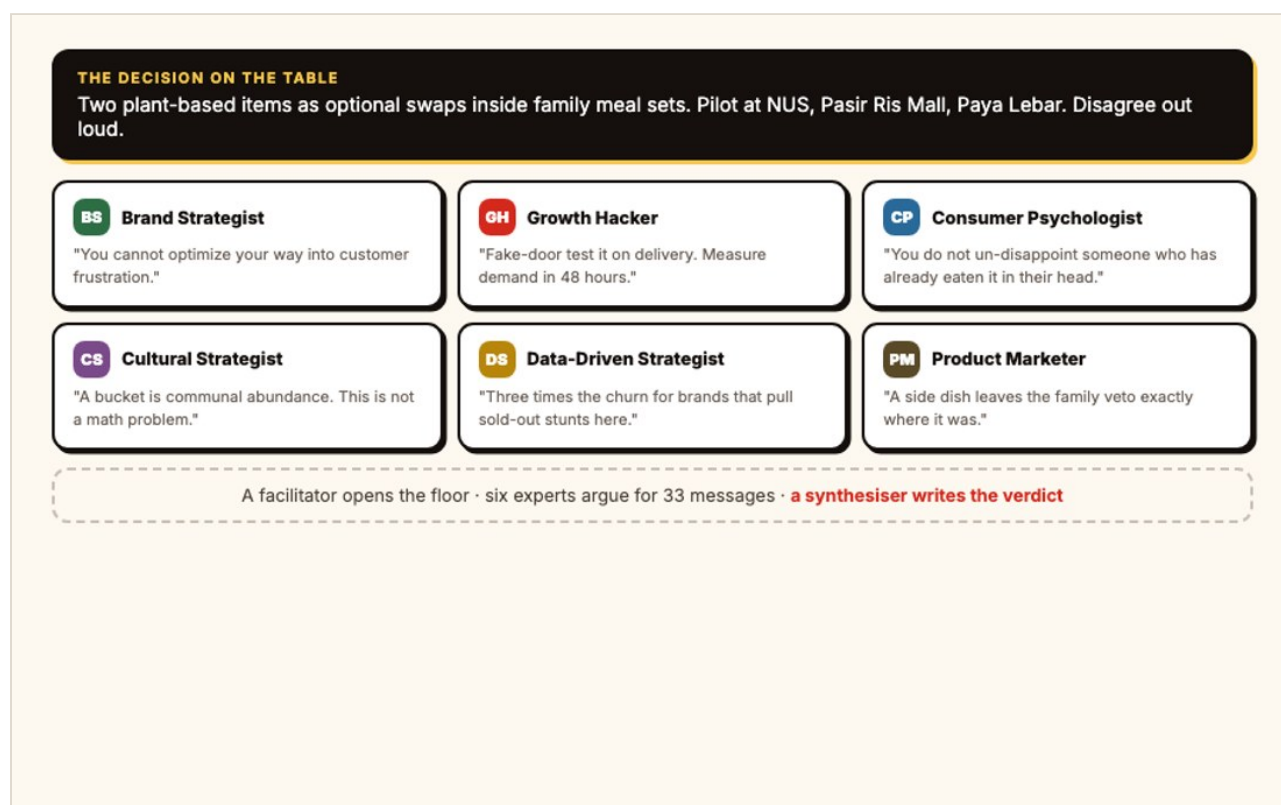


Exhibit 7 • Composition of the session: six selected experts, one decision, thirty-three messages. Source: Focus Group session record, 15 July 2026.

What the debate eliminated

- **A demand test using an unavailable listing.** Proposed to measure appetite in 48 hours without kitchen investment. Rejected on brand-trust grounds and on evidence of elevated churn following such tactics in this market.
- **A zero-cost swap.** Rejected on margin and cannibalisation.
- **A premium side dish.** Rejected on the principle the panel termed dietary dignity: a substitute that is not a main leaves the group objection intact.
- **A pre-selected default at checkout.** Rejected as behaviour inconsistent with a value-led brand.

THE UNRESOLVED TENSION WORTH CARRYING FORWARD

Throughput against demonstrable trust. The operational disciplines argued that added preparation time would erode peak-hour service; the cultural lens argued that unhurried, visible handling is precisely what earns confidence in this market. The recommendation accepts a measured throughput cost.

SECTION 07

Independent model review

The same question, with the same brand context and no shared reasoning, was put to five frontier models from five providers. The purpose is to establish whether a conclusion depends on the choice of model.

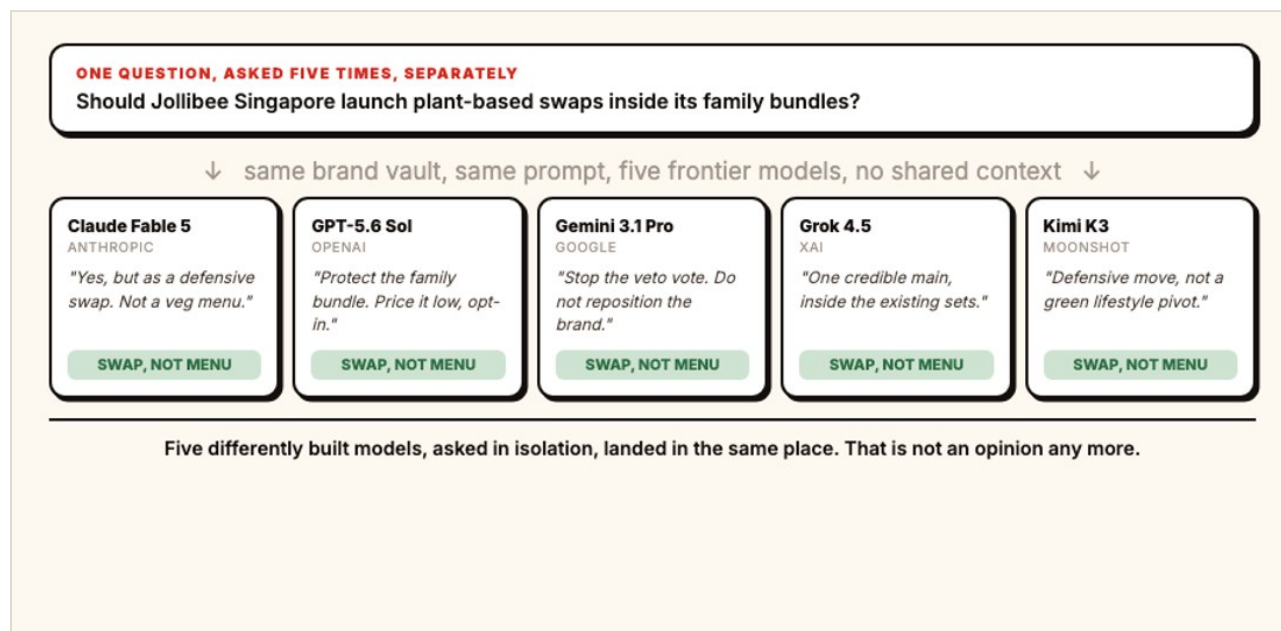


Exhibit 8 • Independent responses from five frontier models. Source: Model Council session, 21 July 2026.

The goal isn't to pivot toward a green lifestyle positioning; it's to stop losing profitable family bundle sales to the flexitarian veto-vote.

REPRESENTATIVE OF THE CONSENSUS POSITION, FIVE OF FIVE MODELS

All five converged: proceed with a plant-based swap, decline a vegetarian menu, and treat the initiative as a defence of existing bundle revenue. Convergence across models that share no training pipeline is a reasonable check on a conclusion. A divergence would have been equally informative, and earlier than a market test.

LIMITATION

Convergence indicates robustness to model choice. It does not validate the underlying commercial assumptions, which remain subject to the pilot described in section 05.

SECTION 08

Campaign development

A separate brief, a month-end payday campaign for the Filipino community in Singapore, was developed through a twenty-one stage workflow. Each stage is a distinct specialist that receives the preceding stage's output.

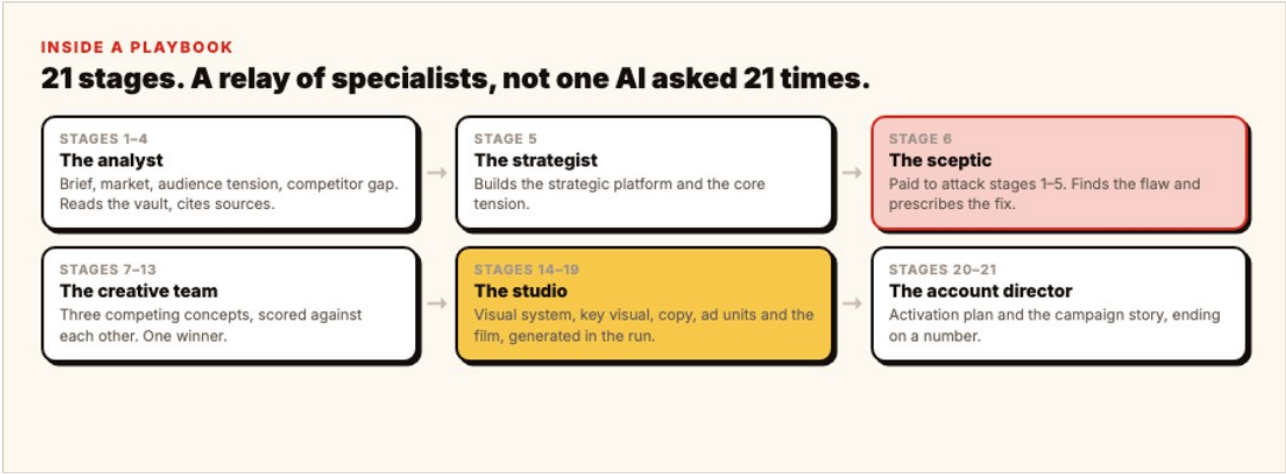


Exhibit 9 • Workflow structure, twenty-one stages grouped by discipline. Source: Campaign Creator run record, 15 July 2026.

When your monthly income clusters around SGD 560, a \$10 saving isn't a promotion. It's permission.

STAGE 12, CITING AN AUDIENCE STUDY SURFACED THROUGH THE BRAND VAULT

The strategic insight is that the barrier is not price but guilt: for an audience whose income is committed before it arrives, discretionary spending requires permission rather than persuasion.

EXHIBIT 10 Selected stage outputs

STAGE	DISCIPLINE	OUTPUT
05	Strategy	Core tension: payday functions as a checklist day, not a spending day
06	Challenge	Identified that the strategy suppressed the brand's value proof; severity critical; prescribed reinstating the stated saving
16	Copy system	Campaign line, channel variants, and an audience-specific headline in Tagalog-English
21	Planning	Activation plan with a stated target of 15-20% uplift in bucket and family-bundle sales across month-end weekends

Source: Campaign Creator stage outputs. Full record in appendix.

SECTION 09

Creative output

The workflow produced finished creative within the same run, without a separate production brief.



Exhibit 11 • Master key visual, generated at stage 15. Source: Campaign Creator run record.

WORTH NOTING

The headline carries the stated saving. That element was absent from the strategy until stage 06 identified its omission as a critical weakness. The correction propagated into finished creative nine stages later without further instruction.



SECTION 10

Method notes and limitations

This exercise serves two purposes. It produces the work above, and it verifies that the platform behaves correctly end to end. Both outcomes are reported.

EXHIBIT 14 Run metrics

171

MESSAGES EXCHANGED

103

WORKFLOW STAGES, EIGHT RUNS

191KB

EVIDENCE GATHERED, FOUR REPORTS

Source: platform execution records, 15–26 July 2026.

Issues identified during the run

- **Video sharing gap.** Images generated on the platform receive a public share link; video does not. Logged as a defect and assigned. Video references in this report link to a direct file as a result.
- **Brand vault rebuilt.** The initial vault was judged insufficiently detailed and replaced before substantive work began. All outputs derive from the second vault.
- **Input tolerance.** The model review question contained two typographical errors. All five models returned correct, complete responses.

Limitations of this document

- **No proprietary data.** The brand vault was constructed entirely from public sources. Internal margin, traffic and cost figures were unavailable, so financial conclusions are directional.
- **Not exercised in this issue.** Large-format production, including single-visual resizing at scale and full advertising sets, was not run. Those capabilities are scheduled for subsequent issues rather than represented here.
- **Unreviewed by the subject.** Jollibee has not seen this work and has not been contacted. Recommendations are an independent exercise.

SECTION 11

Appendix: sources and full records

Every output cited in this report is publicly accessible. No account is required.

EXHIBIT 15 Full records, by capability

CAPABILITY	RECORD	LOCATION
Beacon Listening	778 conversations, share of voice, sentiment	activasia.gimmefylabs.com /chat-result/dr001-listening
Beacon GenAI audit	Five-engine perception audit and remediation	/chat-result/dr001-genai
Beacon Landscape	Claim-versus-perception scoring	/chat-result/dr001-landscape
Beacon Search	Findability audit, full evidence set	/chat-result/dr001-beacon
Focus Group	Thirty-three message transcript and verdict	/chat-result/dr001-debate
Model Council	Five independent model responses	/chat-result/dr001-council
Playbooks	Twenty-one stage campaign development	/chat-result/dr001-campaign
Friday	Single-instruction campaign kit	/chat-result/dr001-friday
Skills	Channel copy and activation checklist	/chat-result/dr001-captions
Studio	Framing session preceding the briefs	/chat-result/dr001-studio
Visual Hub	Brand-vault visuals	/visual-hub/share/dailyrun001a /visual-hub/share/dailyrun001b

All records accessible without authentication. Web edition: gimmefylabs.com/journal/daily-run/001-jollibee-singapore

ABOUT THE DAILY RUN

Each morning we select one brand we admire, run it through every layer of the gimmefy platform, and publish the result in full, including what did not work. The subject changes daily. Featured brands are not clients, have not commissioned or reviewed the work, and no affiliation is implied. Jollibee® is a trademark of Jollibee Foods Corporation.

gimmefy is AI infrastructure for marketing teams: forty-eight frontier models, brand context that persists, marketing-native capabilities, and the governance to run it across a team. Available directly, through the Friday concierge, or from inside Claude, ChatGPT and Gemini via MCP.