

GLOBAL INFO EDGE

Brand. Build. Grow.

The 50-Skill Google Ads System

Architecture Analysis and Full Build Specification

An AI-assisted operating system for eCommerce paid search — decomposed, specified, and assessed for commercial deployment.



Chandan Kumar

Founder, Global Info Edge · New Delhi

globalinfoedge.in · linkedin.com/in/ckmehta

Revision 1.0 · August 2026

Field	Detail
Document	The 50-Skill Google Ads System — Architecture Analysis and Full Build Specification
Author	Chandan Kumar, Founder — Global Info Edge, New Delhi
Publisher	Global Info Edge — Google Premier Partner 2024, Meta Business Partner
Source material	A publicly published LinkedIn post and infographic describing a 50-skill AI system for eCommerce Google Ads, August 2026
Purpose	Decompose the published system, specify every skill in buildable detail, and assess what it means for anyone running, managing or buying paid search
Audience	eCommerce founders, in-house performance marketers, and agency teams
Licence	Free to read, quote and share with attribution. Not for resale.
Revision	1.0 — August 2026

Contents

1. Executive Summary.....	3
2. What the Source Actually Is.....	5
3. Strategic Read for Global Info Edge.....	7
4. System Architecture.....	9
5. Full Skill Specification.....	12
Pillar 1 — Account Audit · Skills 01–07.....	12
Pillar 2 — Search Terms & Negatives · Skills 08–13.....	15
Pillar 3 — Product Feed & Shopping · Skills 14–19.....	17
Pillar 4 — Structure & Settings · Skills 20–25.....	19
Pillar 5 — Bidding & Budget · Skills 26–31.....	21
Pillar 6 — Margin & Profitability · Skills 32–38.....	24
Pillar 7 — Measurement & Tracking · Skills 39–44.....	27
Pillar 8 — Reporting & Accountability · Skills 45–50.....	30
6. Data Requirements.....	32
7. Limitations and Governance.....	34
8. Build Roadmap.....	36
9. Productising This at Global Info Edge.....	38
10. Recommendations.....	40
Appendix A — Master Skill Index.....	41
Appendix B — Skill File Template.....	43
About the Author.....	44

1. Executive Summary

A growth marketer published a system of fifty AI skills for running eCommerce Google Ads accounts, gave it away free in exchange for a comment, and packaged the whole thing in a single infographic. This document takes that system apart, rebuilds it as a specification detailed enough to actually construct, and assesses what it means for anyone who runs, manages or buys paid media for a living.

The central claim in the source post is correct, and it is worth stating plainly: **most of the work in an eCommerce ad account is not strategy, it is checking things.** Pulling the search terms report and reading it line by line. Reconciling the platform's conversion count against actual store orders. Checking whether feed titles match what people type. That work is unglamorous, it compounds when done and decays when skipped, and it is the first thing to fall off the list when the week fills up.

The system's design answer is to build one narrow skill per checking job, rather than one general assistant. Fifty skills across eight pillars — account audit, search terms and negatives, product feed and shopping, structure and settings, bidding and budget, margin and profitability, measurement and tracking, and reporting and accountability. Each takes an export or a screenshot and returns an answer for that specific account. The human still decides.

The one-line summary

This is a checklist system with an AI reading layer on top. The reading layer is commodity — the checklists are the product. GIE should build its own version, populate it with fifteen years of pattern knowledge, use it as the quality assurance layer under every managed account, and then productise the output.

What we conclude

- **The architecture is sound and worth adopting.** Narrow, single-purpose skills outperform a general assistant on this class of work because each one can carry a fixed checklist, a fixed output format and a fixed set of thresholds. That is exactly what audit work needs.
- **The defensible asset is the checklists, not the technology.** Anyone can wire up an AI assistant. The value sits in knowing that impression share lost to budget and impression share lost to rank need opposite fixes, and in the discipline of checking both every month. That knowledge is what GIE already has from ₹50 Cr+ in managed ad spend.
- **Pillar 8 is aimed at agencies.** Six of the fifty skills exist to help an advertiser interrogate whoever runs their ads — verifying reported numbers against raw exports, auditing change history against claimed activity, generating pointed review-meeting questions. This is not a threat to answer defensively. It is a specification for what a transparent agency should be publishing before the client asks.
- **Pillar 6 is the real differentiator.** Seven skills connect the ad account to the profit and loss statement — contribution margin, CAC ceilings, payback period, profit as conversion value, return rates by market. Very few agencies operate at this layer, because it requires data the client

has never handed over. Building the capability to ask for that data and use it well is a genuine competitive position.

- **This is an operating-system change, not a tooling change.** It only works if it sits on top of documented standard procedure and a fixed data routine. Bolted onto a team with no process, fifty skills produce fifty unread documents.

How to read this document

Sections 2 and 3 assess the published system and what it means for GIE commercially. Section 4 sets out the architecture and the design principles that hold it together. Section 5 is the substance — a full specification of all fifty skills, written to a standard detailed enough to build from directly. Sections 6 to 8 cover what data the system needs, where it can go wrong, and the order in which it should be built. Sections 9 and 10 turn it into service offerings and a decision list.

Readers who want the operating detail should go straight to Section 5 and treat everything before it as context. Readers deciding whether to fund the programme need Sections 1, 3, 8 and 10.

One note on voice. Sections 3 and 9 are written about my own agency rather than in the abstract — what this system threatens, what it is worth, and what we intend to do about it. That is deliberate. A recommendation is easier to weigh when you can see who is making it and what it costs them.

2. What the Source Actually Is

2.1 The artefact

The published material consists of two things: a LinkedIn post describing the system, and a single infographic laying out fifty named skills grouped into eight numbered categories, arranged around a shopping cart motif with the Claude and Google Ads marks at the top. The call to action is a comment gate — comment the word 'skills' and receive the system free.

The infographic does the heavy lifting. It is a good piece of information design, and it is worth understanding why it works, because the mechanics are reusable:

- **Specificity signals depth.** Fifty named skills, each with a verb-noun name that tells you exactly what it does. 'Impression Share Diagnoser' and 'Discount Stacking Auditor' are not vague. A reader who has run an ad account recognises the problem each one names.
- **Grouping makes fifty items legible.** Eight pillars of six or seven items each is comfortably within what a reader can scan. Fifty ungrouped items would be noise.
- **Numbering implies completeness.** Sequential numbering from 01 to 50 reads as a finished system rather than a list of ideas.
- **Symmetry carries authority.** Four groups on each side, connector lines to a central object. The layout looks engineered, which makes the contents feel engineered.

2.2 The commercial mechanic

The post is a lead magnet executed well. The comment gate drives engagement signals, which drives distribution; the deliverable is genuinely useful, which converts distribution into follower growth and inbound; and the author's positioning line — a stated €720M revenue and €150M marketing budget — supplies the credibility that makes a free giveaway read as generosity rather than desperation.

The disqualifier in the post is doing real work too. The line stating that this will not help anyone who has never opened their own ad account filters the audience down to practitioners and budget holders. That is a deliberate narrowing, and it improves the quality of what comes back.

2.3 The stated point of view

Three positions are worth extracting, because they shape every skill in the system:

1. **The bottleneck is reading, not thinking.** The system assumes the operator knows what to do once they see the data. What they lack is time to look at all of it. Every skill is therefore built to compress reading, not to replace judgement.
2. **The human keeps the decision.** Skills prepare; people decide. Nothing in the system is designed to push a change into an account autonomously. That is a defensible design choice and it is also what makes the system safe to hand to a client.
3. **The advertiser is the customer, not the agency.** Skills 45 to 47 exist to check up on whoever is running the account. The system is written from the client's chair.

What is not in the source

The post describes the system; it does not publish the internals. What follows in Sections 4 to 6 of this document is our own specification — inputs, method and outputs for each of the fifty skills, written to a standard detailed enough for GIE to build from directly. Where a skill name is ambiguous, we have specified the interpretation that carries the most commercial value for an eCommerce account.

3. Strategic Read for Global Info Edge

3.1 Why this matters to an agency

A system like this changes the floor, not the ceiling. It does not make a mediocre operator good — the skills assume you can pull a report and act on what it says. What it does is make it much harder for a supplier to be quietly sloppy, because a client with fifty structured checks can now ask specific questions about specific numbers.

For an agency with weak process, that is a threat. For an agency with strong process, it is the best sales asset available, because it gives the client a vocabulary for the difference between careful management and invoice-generating activity. GIE sits in the second category and should behave accordingly.

3.2 The three responses available

Response	What it looks like	Assessment
Ignore	Continue current process; treat the published system as noise.	Not viable. Clients who use these skills will ask questions we should be answering unprompted. Being surprised in a review meeting is an avoidable cost.
Defend	Prepare rebuttals to the accountability skills; discourage client use.	Actively harmful. It positions transparency as adversarial and confirms the assumption behind Pillar 8.
Adopt and lead	Build our own version, run it as internal QA on every account, and publish the outputs to clients before they ask.	Recommended. Converts an accountability tool into a differentiator, and the marginal cost is process time we should be spending anyway.

3.3 Where the commercial advantage sits

Three of the eight pillars are genuinely hard to copy, and they are where GIE should concentrate:

Pillar 6 — Margin and profitability

Optimising towards revenue when the business needs profit is the most common and most expensive error in eCommerce paid media. The seven skills in this pillar require contribution margin, return rates by market, discount stacking data and repeat purchase curves — data most clients have never been asked for. An agency that asks for it, gets it, and then bids against profit rather than revenue is doing a materially different job. This is also the natural entry point for the automation and productised services track.

Pillar 7 — Measurement and tracking

Conversion reconciliation, attribution windows, new-customer filtering, offline conversion import and consent mode. Every decision in the other seven pillars rests on these being sound, and in most accounts they are not. Running Skill 39 first on a new account frequently changes what everyone

believes about performance before a single optimisation is made. That is a powerful opening move in a pitch.

Pillar 8 — Reporting and accountability

Answering the audit questions before the client asks them is a positioning move, not a defensive one. In our experience, recommending the route that reduces your own share of the work builds more commercial credibility than an optimised pitch ever does — and it survives contact with procurement far better.

3.4 Where the advantage does not sit

Being honest about the limits matters as much as the opportunity:

- The skills themselves are reproducible in an afternoon by anyone competent with an AI assistant. There is no technical moat.
- The published system is free. Charging for access to something equivalent is not a viable business model; charging for the judgement applied to its outputs is.
- Skills read exports; they do not have live account access. Every skill therefore depends on someone pulling the right report. Without a fixed data-pull routine, the system does not run at all — which is why Section 8 treats the data routine as the first build phase, not the last.
- Nothing here removes the need for a competent operator. The value is a floor on quality, not a substitute for skill.

4. System Architecture

4.1 The eight pillars

The pillars are sequenced deliberately. Pillars 1 to 4 establish and correct the state of the account. Pillar 5 governs money. Pillars 6 and 7 connect the account to the business and validate that the numbers can be trusted. Pillar 8 governs how all of it is communicated and challenged.

#	Pillar	Skills	Question it answers	Cadence
1	Account Audit	01–07	What is the current state, and where is money leaking?	New account, then quarterly
2	Search Terms & Negatives	08–13	Are we buying the right queries?	Weekly
3	Product Feed & Shopping	14–19	Can the catalogue be found and bought?	Fortnightly
4	Structure & Settings	20–25	Is the account built to be manageable?	Quarterly
5	Bidding & Budget	26–31	Where does the money go, and how fast?	Weekly / monthly
6	Margin & Profitability	32–38	What can we actually afford to pay?	Quarterly, or on price change
7	Measurement & Tracking	39–44	Can we believe the numbers?	Monthly reconciliation
8	Reporting & Accountability	45–50	Is the work being done and communicated honestly?	Monthly

4.2 Design principles

Six principles hold the system together. They matter more than the individual skills, because they are what make the outputs consistent enough to be trusted.

- 1. One skill, one job.** Each skill answers a single question with a fixed output shape. A skill that does three things produces vague answers to all three. The narrowness is the feature.
- 2. Input is an export, not an integration.** Every skill takes a report, a screenshot or a data file. This keeps the system portable across accounts, avoids API permissions, and means it works for someone auditing an account they do not own.
- 3. Output is a decision list, not a description.** No skill should return a summary of what the data says. Every skill returns actions, ranked by rupee impact, with the specific change named.
- 4. The human decides.** Skills prepare and rank; a person approves and implements. This is a hard rule, not a phase-one limitation.
- 5. Rank by money, never by count.** Twelve small issues and one large one is not thirteen issues. Every output is ordered by spend at risk or revenue exposure.

- 6. Thresholds are written down.** 'Low conversion volume' is unusable; 'below thirty conversions per thirty days' is a rule. Every skill carries explicit numeric thresholds that can be argued with and revised.

4.3 Anatomy of a single skill

Each skill is a self-contained instruction file. The structure below is the standard we recommend GIE adopt, and it maps cleanly onto how skills are packaged for use with an AI assistant.

Component	Purpose	Example — Skill 03, Impression Share Diagnoser
Name and trigger	When this skill should be invoked	Use when a campaign is losing impression share, or as part of the quarterly audit
Required input	The exact report and columns needed, so the operator pulls it once and correctly	Campaign report with Search IS, IS lost (budget), IS lost (rank), Top IS, Abs Top IS, avg CPC, QS components
Thresholds	The numeric rules the skill applies	Flag when IS lost to budget exceeds 10%; flag as structural when both budget and rank losses exceed 20%
Method	Ordered steps, so two operators get the same answer	Split lost IS by cause → model incremental spend against CAC ceiling → decompose Quality Score → flag dual-cause campaigns
Output format	The fixed shape of the answer	Per-campaign verdict table: raise budget / fix relevance / raise bids / leave alone, with expected volume gain
Escalation rule	What the skill must refuse to conclude alone	Do not recommend a budget increase without a current CAC ceiling from Skill 33

The escalation rule is the important one

A skill that will confidently recommend a budget increase without knowing the CAC ceiling is worse than no skill at all, because it produces a fluent and wrong answer. Every skill in Pillar 5 must refuse to conclude without the relevant Pillar 6 input. Building those dependencies in is what separates a usable system from an impressive-looking one.

4.4 Dependency map

The pillars are not independent. Four dependencies must be enforced or the system produces confident nonsense:

- **Pillar 5 depends on Pillar 6.** No target ROAS, budget reallocation or spending ceiling can be set without contribution margin (32) and a CAC ceiling (33).
- **Everything depends on Pillar 7.** If conversion reconciliation (39) shows a material unexplained gap, every performance number in the account is provisional and optimisation should pause until it is resolved.

- **Pillar 3 depends on Skill 19.** Custom labels are the segmentation layer that Skills 17, 21 and much of Pillar 4 rely on. Build the label schema before the campaign structure.
- **Skill 49 depends on Skill 31.** No test enters the backlog without its kill number and scale number already attached.

5. Full Skill Specification

The fifty skills are specified below in pillar order. Each entry states the purpose, the exact input required, the ordered method, and the output shape. Notes flag the skills where our own account experience suggests a specific caution or an outsized return.

These specifications are written to be built from directly. Thresholds stated as numbers are starting positions drawn from our own portfolio and should be revised per account as evidence accumulates.

Pillar 1 — Account Audit

Seven skills that establish the current state of the account before anyone proposes changing it. This pillar answers one question: where is the money going, and is any of it going somewhere stupid? Run this pillar first on every new account and quarterly on every retained account.

01 Account Health Scanner

A single first-pass diagnostic across the whole account that tells you where to look next.

Input required — Campaign-level export for 90 days (cost, clicks, conversions, conversion value, search impression share, CTR, average CPC), plus account settings and current campaign statuses.

Method

1. Rebuild the account tree — campaign count, ad group count, active vs paused, budget per campaign.
2. Flag any campaign taking more than 5% of account spend with zero conversions in 30 days.
3. Flag campaigns on Smart Bidding sitting below roughly 30 conversions per 30 days — the signal floor.
4. Check for disapprovals, limited-by-budget status, and campaigns that have not been touched in 90 days.
5. Rank every finding by rupees at risk, never by count of issues.

Output — A one-page health card: red / amber / green per campaign, the top five issues ranked by spend exposure, and the specific next skill to run against each issue.

Note. *This is the only skill in the system that should ever be run in isolation. Everything else follows from what it surfaces.*

02 Wasted Spend Finder

Isolate the spend that produced nothing and could be switched off this week.

Input required — Search terms report (90 days), keyword report, placement report for Display and Performance Max, plus device and geography splits.

Method

1. Sum spend where conversions are zero across 90 days and clicks exceed three times the account's average clicks-to-conversion.
2. Segment the waste by cause: irrelevant query, wrong geography, wrong device, junk placement, out-of-catalogue product.
3. Compute recoverable rupees per month for each bucket.

4. Separate 'cut now' from 'cut but watch' — some zero-conversion terms carry assist value.

Output — A ranked wasted-spend table with cause, monthly rupee value, and the exact action required: negative keyword, placement exclusion, bid adjustment, or pause.

03 Impression Share Diagnoser

Distinguish a budget problem from a rank problem, because they need opposite fixes.

Input required — Campaign report including search impression share, impression share lost to budget, impression share lost to rank, top and absolute top impression share, average CPC and Quality Score components.

Method

1. Split lost impression share cleanly into budget-caused and rank-caused.
2. For budget-lost share, model the incremental spend at current CPA and test whether it stays under the CAC ceiling.
3. For rank-lost share, break down Quality Score into expected CTR, ad relevance and landing page experience.
4. Flag any campaign where both exceed 20% — that is a structural problem, not a lever problem.

Output — A per-campaign verdict — raise budget, fix relevance, raise bids, or leave alone — with the expected volume gain attached to each.

04 Learning Phase Auditor

Break the cycle in which every change resets learning and the account never stabilises.

Input required — Change history export (90 days), campaign status history, and conversions per campaign per week.

Method

1. Map every bid strategy, target and budget change against the 7 to 14 day learning window it triggers.
2. Count how many days each campaign has spent in a learning state over the period.
3. Flag every campaign changed more than once per 14 days.
4. Test conversion volume against the threshold Smart Bidding needs to produce a stable model.

Output — A learning-stability score per campaign, a 'do not touch until' date for each, and a consolidation recommendation wherever conversion volume is too thin to support the current structure.

05 Auto-Applied Recommendations Checker

Find the changes the platform made to the account that nobody approved.

Input required — Screenshot of the auto-apply settings page, plus change history filtered to changes made by Google.

Method

1. List every auto-apply toggle currently switched on.
2. Pull all auto-applied changes from 90 days of change history.
3. Classify each as benign, spend-expanding, or targeting-expanding — broad match additions, display expansion, similar audiences.
4. Quantify the spend attributable to auto-applied changes.

Output — A toggle-by-toggle keep-or-kill recommendation, plus a list of already-applied changes that should be reversed.

Note. *In most audited accounts this is the single fastest source of unexplained spend. It costs nothing to check and is almost never checked.*

06 Geographic Targeting Reviewer

Confirm the account is buying attention where the business can actually deliver and get paid.

Input required — Geographic report by location of interest and user location, the location setting on every campaign, and the serviceable or shippable area list from operations.

Method

1. Compare spend by region against revenue by region and return rate by region.
2. Flag every region carrying spend without delivery coverage.
3. Check the presence versus presence-or-interest setting on each campaign — the default is frequently wrong.
4. Identify tier-2 and tier-3 regions with above-average ROAS that are currently under-bid.

Output — An exclusion list, a bid-adjustment table by region, and a settings correction list.

07 Seasonality Checker

Separate a real performance problem from a calendar effect, before anyone panics or fires anyone.

Input required — Twenty-four months of weekly cost, conversions and revenue; the festival and sale calendar; category-level search trend data.

Method

1. Build a week-of-year index for spend, conversion rate and average order value.
2. Overlay the current period against the same weeks in the prior year.
3. Identify structural peaks — festive season, end-of-season sale, payday cycles — and distinguish them from noise.
4. Flag where a seasonality adjustment is justified versus where Smart Bidding should be left to adapt on its own.

Output — A seasonality profile for the account, a list of dated adjustment windows, and a plain verdict on whether the current dip is normal.

Pillar 2 — Search Terms & Negatives

Six skills covering the highest-yield and least-performed task in paid search: reading the search terms report properly. This is the work that gets skipped first when the week fills up, and it is where the majority of recoverable waste sits.

08 Search Term Report Reader

Read the full search terms report line by line — the job that is always skipped.

Input required — Search terms report for 30 to 90 days including cost, clicks, conversions, conversion value, the matched keyword and the match type.

Method

1. Bucket every term: converting, spending-without-converting, informational, competitor, brand, irrelevant, out-of-catalogue.
2. Sort the spending-without-converting bucket by cost descending — this is the working list.
3. Cross-check high-cost terms against the actual product catalogue.
4. Surface converting terms that do not yet exist as keywords in the account.

Output — Four action lists — add as keyword, add as negative, watch, no action — each with the rupee impact attached.

09 Negative Keyword List Builder

Turn scattered ad-hoc negatives into shared, maintained, correctly-scoped lists.

Input required — Existing negative keyword lists, the search terms report, the product catalogue, and brand plus competitor name variations.

Method

1. Audit existing negatives for over-blocking — negatives that are quietly suppressing converting terms.
2. Build themed shared lists: irrelevant intent, jobs and careers, free and DIY, competitor, out-of-catalogue, geography mismatch.
3. Assign the correct match type — phrase for themes, exact for single high-risk terms.
4. Define which list attaches to which campaign type. Never blanket-apply a list to a brand campaign.

Output — Named shared lists formatted for direct upload, an attachment map by campaign, and a list of existing negatives to remove.

10 Brand vs Non-Brand Splitter

Stop brand traffic from flattering non-brand performance.

Input required — Search terms and keyword reports, plus the full set of brand term variations and common misspellings.

Method

1. Classify every term as brand, brand-plus-modifier, competitor, or generic.
2. Recompute ROAS, CPA and revenue with brand stripped out entirely.
3. Quantify what share of reported account performance is brand demand the business already owned.
4. Check whether brand and non-brand currently share campaigns or budgets.

Output — A restated performance table showing blended versus non-brand, a structural split recommendation, and a defensible position on brand-defence budget.

Note. *For agencies, running this before the client does is the difference between a comfortable review meeting and an uncomfortable one.*

11 Query Intent Classifier

Match message and landing page to what the searcher is actually trying to do.

Input required — Search terms report with conversion data and current landing page mapping.

Method

1. Tag each term by intent: transactional, comparison, research, support, navigational.
2. Compute conversion rate and average order value by intent class.
3. Map each intent class to the page it currently lands on.
4. Flag the mismatches — research intent landing on a product page, transactional intent landing on a blog.

Output — An intent-by-performance matrix, a list of routing corrections, and ad copy angles specific to each intent class.

12 Match Type Reviewer

Check that match types are doing the job they were chosen to do.

Input required — Keyword report with match type, search terms mapped back to keywords, and cost and conversions by match type.

Method

1. Measure query-to-keyword drift on broad match — how far the platform is travelling from intent.
2. Compare CPA and conversion rate across exact, phrase and broad for the same keyword theme.
3. Check for internal competition where one term is served by multiple match types.
4. Verify broad match is only running where there is enough conversion signal to steer it.

Output — A match-type change list, duplication fixes, and a written rule for where broad match is permitted in this account.

13 Competitor Query Auditor

Decide with numbers, rather than ego, whether bidding on rivals is worth it.

Input required — Search terms filtered to competitor names, auction insights report, and cost and conversion data on those terms.

Method

1. Isolate competitor-term spend, conversion rate, CPA and average order value.
2. Compare that performance against non-brand generic terms.
3. Read auction insights to establish who is bidding on your brand.
4. Model the cost of defending your own brand terms against the cost of attacking theirs.

Output — A continue, reduce or exit recommendation per competitor, and a stated brand-defence position.

Pillar 3 — Product Feed & Shopping

Six skills on the feed. In eCommerce accounts the feed is the campaign; titles and attributes do the work that keywords do in search. This pillar is consistently the largest untapped source of impression volume in the accounts we audit.

14 Product Title Rewriter

Rewrite product titles so they contain the words buyers actually type.

Input required — Product feed export with title, description, brand and attributes; the search terms report; and top-converting queries by product.

Method

1. Extract high-intent query patterns per category from the search terms data.
2. Apply a title formula appropriate to the vertical — typically brand, product type, key attribute, size or colour, then model.
3. Front-load the first seventy characters, which is what is most reliably read.
4. Strip promotional language and non-compliant text that risks disapproval.

Output — A rewritten title column mapped to item ID, ready for feed upload, with a before-and-after keyword coverage score.

15 Feed Attribute Auditor

Find the missing and malformed fields that quietly suppress impressions without ever showing an error.

Input required — Full feed export and Merchant Center diagnostics.

Method

1. Check completeness across required and recommended attributes: GTIN, MPN, brand, condition, availability, Google product category, product type, colour, size, age group, gender, shipping and tax.
2. Validate formatting — currency, price parity with the live site, image URL validity.
3. Rank every gap by the revenue of the SKUs it affects.
4. Separate hard disapprovals from soft suppressions, which are far more common and far less visible.

Output — An attribute completeness scorecard, a fix list prioritised by revenue exposure, and a clear required-versus-optional split.

16 SKU Coverage Checker

Find the products that receive no impressions at all.

Input required — Shopping and Performance Max product report by item ID, plus the complete catalogue list.

Method

1. Join the catalogue against served items to identify zero-impression SKUs.
2. Classify the cause for each: disapproved, out of stock, excluded by filter, priced out of the auction, or poor title.
3. Compute the catalogue revenue represented by dark SKUs.

4. Check the concentration ratio — what share of spend sits on what share of the catalogue.

Output — A dark-SKU list with cause and fix per item, plus a catalogue concentration report.

17 Shopping Bid Segmenter

Stop treating an entire catalogue as a single bidding unit.

Input required — Product-level performance data — cost, conversions, conversion value, and margin per SKU where available.

Method

1. Segment products into performance tiers: heroes, workers, sleepers and dogs.
2. Cross-cut those tiers by margin band and price band.
3. Recommend campaign or asset group separation for heroes and dogs specifically.
4. Set differentiated ROAS targets per tier rather than a single account-wide target.

Output — A segmentation map, a custom label plan, and a target ROAS per tier.

18 Out of Stock Guard

Prevent spend on items that cannot be bought.

Input required — Feed availability field, live site stock status, spend by item ID, and the feed refresh schedule.

Method

1. Reconcile the availability field in the feed against the truth on the site.
2. Measure spend that landed on out-of-stock or low-stock SKUs in the last 30 days.
3. Compare feed fetch frequency against actual stock turnover speed.
4. Define a supplemental feed rule or an automated exclusion to close the gap permanently.

Output — A discrepancy list, the rupee value leaked, and a stock-sync rule covering feed frequency, supplemental feed logic, or automated rules.

19 Custom Label Builder

Give the account the segmentation dimensions the platform does not provide on its own.

Input required — Catalogue with cost price, margin, stock age, seasonality flag and bestseller flag.

Method

1. Assign the five available custom labels to the five decisions that actually matter — usually margin band, price band, stock age, seasonality and performance tier.
2. Define the exact numeric rule for each band so the labels can be regenerated without judgement calls.
3. Map labels onto campaign and asset group structure.
4. Set a refresh cadence and name an owner.

Output — A custom label schema covering label_0 through label_4, the values and rules for each, and a supplemental feed specification.

Pillar 4 — Structure & Settings

Six skills on architecture. Structure decisions compound: a badly-structured account gets harder to fix every month it runs. This pillar also covers the settings that ship switched on by default and quietly spend money in places nobody chose.

20 Campaign Structure Mapper

Draw the account as it actually is, before anyone proposes changing it.

Input required — Full campaign, ad group, keyword and asset group export with spend and conversions at every level.

Method

1. Build the complete tree with spend and conversion volume attached to every node.
2. Flag thin nodes sitting below the signal threshold for the bid strategy in use.
3. Flag overlapping nodes competing for the same queries.
4. Compare the result against a reference structure for the vertical and business model.

Output — A current-state map, a target-state map, and a sequenced migration path designed to protect existing learning.

21 PMax Asset Group Planner

Impose deliberate structure on the campaign type designed to hide it.

Input required — Performance Max campaign settings, asset group listing, listing group splits, asset performance ratings and configured search themes.

Method

1. Check asset group count against the catalogue segmentation logic — one group for everything is the default failure.
2. Verify full asset coverage in every group: headlines, long headlines, descriptions, square, portrait and landscape images, logo and video.
3. Review listing group splits against the custom label schema.
4. Check search themes, audience signals and brand exclusions.

Output — An asset group blueprint, a per-group asset gap list, and a signal configuration specification.

22 Ad Schedule Reviewer

Buy the hours that convert and stop paying full price for the ones that do not.

Input required — Hour-of-day and day-of-week report with cost, conversions and conversion value, plus support and dispatch operating hours.

Method

1. Build a seven-by-twenty-four heat map of conversion rate, CPA and ROAS.
2. Identify statistically meaningful blocks and discard cells with insufficient volume.
3. Cross-check against operational capacity — call handling hours, dispatch cut-off times.
4. Decide between schedule bid adjustments and simply trusting Smart Bidding.

Output — The heat map, recommended adjustments or restrictions, and an explicit note on where the algorithm should be left alone.

23 Device Bid Analyser

Check that the mobile and desktop split reflects the real customer journey rather than last-click artefacts.

Input required — Device report with cost, conversions, conversion value and average order value; cross-device conversion data; and site speed by device.

Method

1. Compare conversion rate, average order value and CPA across devices.
2. Check assist behaviour — mobile discovery leading to desktop purchase is common and routinely misread.
3. Check mobile page speed and checkout friction on the top landing pages.
4. Decide whether the correct fix is a bid adjustment or a user experience fix.

Output — A device performance table, an adjustment recommendation, and a post-click fix list wherever the real problem is the site.

24 Network Settings Checker

Catch the default settings that quietly spend money in places nobody chose.

Input required — Campaign settings export or screenshots for every campaign in the account.

Method

1. Check whether Search Partners is on, and isolate its performance separately.
2. Check whether Display Network expansion is enabled on Search campaigns.
3. Verify location, language, ad rotation and final URL expansion settings.
4. Compare partner and expansion performance against core search performance.

Output — A settings correction checklist per campaign, with the associated spend attached to each item.

25 Duplicate Coverage Detector

Find where the account is bidding against itself.

Input required — Keyword lists across all campaigns, search terms mapped to keywords, and the coexistence map of Performance Max, Shopping and Search.

Method

1. Find identical or near-identical keywords appearing in more than one campaign.
2. Identify queries being served by more than one campaign type.
3. Check Performance Max versus Search priority conflicts and brand overlap.
4. Check Shopping campaign priority settings where legacy campaigns remain.

Output — An overlap map, a deduplication plan, and explicit traffic-routing rules using negatives, campaign priorities and brand exclusions.

Pillar 5 — Bidding & Budget

Six skills on where the money goes and how fast. The defining principle of this pillar is that targets are derived from the profit and loss statement, never copied from last month or borrowed from a case study.

26 Target ROAS Setter

Derive the target from the P&L instead of inheriting last month's number.

Input required — Contribution margin by product or category, fixed cost base, return rate, target contribution, plus current target and achieved ROAS.

Method

1. Calculate break-even ROAS directly from contribution margin — one divided by contribution margin percentage.
2. Layer the required profit contribution on top of break-even.
3. Adjust for return rate and payment or cash-on-delivery costs.
4. Set targets per campaign by margin band rather than one figure across the account.
5. Define the maximum step size for any change — typically 10 to 15% per adjustment.

Output — Break-even and target ROAS per campaign or segment, with a phased adjustment schedule that avoids resetting learning.

27 Budget Reallocation Planner

Move money to where the next rupee earns the most, not where the average looks best.

Input required — Campaign-level spend, conversion value, marginal CPA trend, impression share lost to budget, and the total budget cap.

Method

1. Estimate the marginal return curve per campaign rather than relying on average performance.
2. Identify budget-capped campaigns still performing above the CAC ceiling — these are the obvious recipients.
3. Identify campaigns past the point of diminishing returns.
4. Produce a reallocation that holds total spend flat unless incremental return genuinely justifies growth.

Output — A from-and-to reallocation table with the expected conversion delta, and a staged implementation plan.

28 Spending Ceiling Calculator

Establish the number above which growth starts destroying profit.

Input required — Contribution margin, fixed cost base, CAC ceiling, the historical CPA-versus-volume curve, and the current cash position.

Method

1. Model CPA inflation as volume scales — CPA is not constant and treating it as constant is the most common planning error.
2. Find the spend point at which marginal CAC meets the CAC ceiling.

3. Sanity-check the result against the cash conversion cycle and inventory availability.
4. Express the answer as a daily, a monthly and a peak-season ceiling.

Output — Three ceiling figures with the assumptions written underneath, and the trigger conditions that require them to be revisited.

29 Bid Change Log Keeper

Make every change reviewable so that learning compounds instead of resetting.

Input required — Change history export, the prior log, and performance data before and after each logged change.

Method

1. Record date, campaign, change made, reason, expected effect and review date.
2. Enforce a minimum observation window before the next change to the same campaign.
3. Measure actual against expected at the review date.
4. Maintain a running hit rate by change type so the team learns what actually works in this account.

Output — A structured change log, an 'under observation, do not touch' list, and a hit-rate summary by change type.

Note. *Unglamorous and disproportionately valuable. Most accounts have no institutional memory of what was tried and what happened.*

30 Q4 Budget Ramp Planner

Enter peak season with the algorithm already warm.

Input required — Prior-year peak data, category trend data, inventory plan, promotion calendar and total peak budget.

Method

1. Work backwards from peak dates to set a pre-peak ramp — typically four to six weeks — so learning completes before demand arrives.
2. Set week-by-week budget steps no larger than 20 to 30%.
3. Apply seasonality adjustments only for known short spikes, not for the whole season.
4. Plan the post-peak taper deliberately to avoid a January performance cliff.

Output — A week-by-week budget and target schedule with dated actions and guardrails.

Note. *For Indian accounts, read 'Q4' as the festive cycle — Onam through Diwali — plus the end-of-season sale window. The mechanics are identical; only the dates move.*

31 Test Kill Number Setter

Decide in advance what will end a test, so that hope does not extend it.

Input required — The test hypothesis, baseline CPA and conversion rate, available traffic volume, and acceptable loss.

Method

1. Compute the sample size needed for a meaningful read at the expected effect size.
2. Set a maximum spend and a maximum number of days.
3. Define the kill threshold and the scale threshold as numbers.

4. Write both numbers down before the test starts, not after it begins to look bad.

Output — A one-line test contract: spend cap, day cap, kill number, scale number, decision date.

Pillar 6 — Margin & Profitability

Seven skills that connect the ad account to the profit and loss statement. This is the pillar that separates a media buyer from a growth partner, and it is the hardest to fake, because it requires data most advertisers have never handed over.

32 Contribution Margin Calculator

Establish the single number that every bidding decision downstream depends on.

Input required — Selling price, cost of goods, shipping, payment gateway and cash-on-delivery fees, packaging, returns cost and discount rate — per SKU or per category.

Method

1. Build the per-unit waterfall from gross price down to contribution.
2. Compute margin at list price and again at average realised price after discounting.
3. Weight by actual sales mix to produce a blended figure.
4. Flag any SKU with negative contribution once returns are accounted for.

Output — Per-SKU and blended contribution margin, plus an explicit loss-making SKU list.

Note. Nothing in Pillars 5 or 6 can be run credibly until this skill has been run once.

33 CAC Ceiling Setter

Set the maximum acceptable acquisition cost with the reasoning attached to it.

Input required — Contribution margin, average order value, repeat purchase rate, gross margin on repeat orders, and payback tolerance.

Method

1. Compute the first-order CAC ceiling — contribution earned on a single order.
2. Compute the LTV-informed ceiling using observed repeat behaviour, appropriately discounted.
3. Choose which ceiling applies based on the current cash position, not on optimism.
4. Translate the chosen ceiling into a target CPA and target ROAS per segment.

Output — Both ceilings, the chosen operating ceiling, and its translation into campaign-level targets.

34 Payback Period Modeller

Answer how long the business waits to get its acquisition money back.

Input required — CAC, first-order contribution, cohort repeat purchase curve, gross margin, and time to cash.

Method

1. Build a cohort contribution curve month by month.
2. Find the month in which cumulative contribution crosses CAC.
3. Test sensitivity to changes in repeat rate and margin.
4. Compare payback duration against actual working capital availability.

Output — A payback curve, the payback month, and the maximum sustainable CAC given the business's real cash tolerance.

35 Profit as Conversion Value Setup

Make the bidding algorithm optimise for profit rather than revenue.

Input required — Margin data by product or category, current conversion value configuration, tag and tag manager setup, and platform capability.

Method

1. Decide the granularity — SKU-level margin, category margin band, or a blended figure.
2. Specify the data layer change or conversion value modification required.
3. Define the implementation route: dynamic value passed in the purchase event, or conversion value rules.
4. Set the reconciliation test and write the rollback plan before deployment.

Output — An implementation specification a developer can act on, a validation checklist, and the restated target ROAS that applies once profit-value is live.

Note. *The highest-leverage change in this entire system, and the one most often deferred because it requires engineering time.*

36 Discount Stacking Auditor

Find the orders where every promotion applied at once and the margin disappeared.

Input required — Order-level data with all discounts, coupon codes, shipping subsidy and gateway fees.

Method

1. Reconstruct realised margin per order after every stacked discount.
2. Identify the specific combinations that push contribution below zero.
3. Quantify what share of ad-driven orders are affected.
4. Model the effect of stacking rules on both volume and margin before recommending them.

Output — A stacking exposure report, the specific rules to implement, and the projected margin recovery.

37 Return Rate by Market Analyser

Price returns into targeting rather than discovering them at month end.

Input required — Order and return data by region, category, size or variant, and payment method.

Method

1. Compute return rate and net contribution by region and by category.
2. Isolate the cash-on-delivery effect on returns and return-to-origin.
3. Restate ROAS net of returns by segment.
4. Identify the segments that look profitable gross and are loss-making net.

Output — A net-of-returns performance table, geography and category bid actions, and a cash-on-delivery policy recommendation.

Note. *In the Indian market this skill frequently reverses the ranking of best and worst performing regions.*

38 AOV and Bundle Planner

Raise the value of each acquired order so that the CAC ceiling rises with it.

Input required — Order composition data, attach rates, margin by SKU, the current AOV distribution, and the free-shipping threshold.

Method

1. Analyse frequently co-purchased items and their combined margin.
2. Model bundle pricing that raises average order value without cutting contribution.
3. Test the free-shipping threshold against the actual AOV distribution.
4. Recompute the CAC ceiling at the new average order value.

Output — Bundle and threshold recommendations, with the resulting change in permissible CAC stated explicitly.

Pillar 7 — Measurement & Tracking

Six skills on whether the numbers can be believed. Every decision in the preceding six pillars rests on measurement being sound. In practice it usually is not, and the gap is rarely investigated until someone notices platform-reported revenue does not match the bank.

39 Conversion Reconciliation Skill

Reconcile what the platform reports against what the business actually banked.

Input required — Google Ads conversions and value by day, GA4 data, the backend or store order export, and refunds and cancellations.

Method

1. Align date ranges, time zones and attribution basis — conversion time versus click time is a frequent source of phantom discrepancy.
2. Compare order counts and revenue across all three sources.
3. Attribute the gap to known causes: modelled conversions, cross-device, consent loss, duplicate tag firing, tag misfire, returns.
4. Set an acceptable variance band and an alert threshold for future months.

Output — A three-way reconciliation table, variance explained by cause, and a fix list for the unexplained remainder.

Note. If this skill is run first on a new account, it will frequently change what everyone believes about performance before any optimisation begins.

40 Attribution Window Reviewer

Check that the measurement window matches how people actually buy this product.

Input required — Path length and time lag reports, the configured window on every conversion action, the attribution model in use, and the known consideration cycle.

Method

1. Read the time-lag distribution — what share converts after day 7, day 14, day 30.
2. Compare that against the click-through and view-through windows currently set.
3. Check window consistency across every conversion action, since mismatches are common.
4. Model how reported performance shifts once the window is corrected.

Output — A window recommendation per conversion action, and the restated performance baseline the change creates.

41 Repeat Buyer Conversion Filter

Stop paying new-customer prices for orders the business was going to receive anyway.

Input required — Conversion action configuration, customer type data, new versus returning order split, and the CRM list.

Method

1. Quantify what share of ad-attributed conversions are existing customers.
2. Recompute CAC on new customers only.

3. Specify a new-customer conversion action or a new-customer acquisition goal.
4. Define the value differential between a new and a repeat order.

Output — True new-customer CAC, the configuration specification, and the revised targets that follow from it.

42 Audience Exclusion Builder

Define who the account should stop paying to reach.

Input required — CRM segments, recent purchaser lists, converter lists, high-return customers, and the existing audience configuration.

Method

1. Build the exclusion inventory: recent purchasers inside the repurchase cycle, serial returners, employees, existing subscribers.
2. Check exclusions are applied at the right level and on the right campaign types.
3. Verify Customer Match list health, size and refresh cadence.
4. Separate exclude-entirely from bid-down — they are different decisions.

Output — An exclusion list with scope, a refresh schedule, and a campaign-type application map.

43 Offline Conversion Import Guide

Feed downstream truth back into bidding for anything that does not close online.

Input required — The CRM and order pipeline, GCLID capture status, sales stage definitions, and value per stage.

Method

1. Verify GCLID and GBRAID capture and storage through the full funnel — this is where most implementations fail.
2. Map CRM stages to conversion actions and assign values.
3. Define the upload mechanism and cadence.
4. Set stage values, including net-of-return values where applicable.

Output — An end-to-end import specification, field mapping, cadence, and a validation test.

Note. For lead-generation accounts this is the single change that most reliably improves lead quality, because it teaches the algorithm what a good lead looks like.

44 Consent Mode Checker

Confirm measurement holds up under consent and privacy constraints.

Input required — Consent management platform configuration, Consent Mode v2 implementation, tag firing behaviour, enhanced conversions status, and regional traffic mix.

Method

1. Verify consent signals are transmitted and that modelling is active.
2. Check enhanced conversions for web and, where relevant, for leads.
3. Test tag behaviour under both grant and deny states.
4. Quantify what share of reported conversions are modelled rather than observed.

Output — A compliance and measurement-integrity checklist, the gaps found, and the fix sequence.

Pillar 8 — Reporting & Accountability

Six skills written explicitly from the advertiser's side of the table. This pillar exists to help a business interrogate whoever runs its ads. For an agency, that makes it the most commercially significant pillar in the system — see Section 9.

45 Monthly Report Interrogator

Read a supplier or internal report for what it leaves out.

Input required — The monthly report itself, plus raw platform access or exports covering the same period.

Method

1. Check whether brand performance is separated from non-brand.
2. Check whether the metrics presented are the ones that drive decisions, or the ones that happen to look best.
3. Verify every headline number independently against the raw export.
4. Check whether last month's stated commitments were reported against at all.

Output — A verified and unverified list, a missing-metric list, and the questions that follow from both.

46 Agency Question Generator

Produce the specific questions that make the next review meeting useful.

Input required — The report, change history, current performance against targets, and prior meeting notes.

Method

1. Generate questions tied to specific numbers rather than general challenge.
2. Order the questions by materiality.
3. For each question, define what a good answer looks like and what a deflection looks like.
4. Include at least one forward-looking commitment question.

Output — A ranked question list with acceptance criteria attached to each question.

Note. This skill is aimed squarely at agencies. The correct response is not to resist it but to answer these questions before the client asks them.

47 Change History Reviewer

Establish what was actually done to the account, and by whom.

Input required — Change history export for 90 days, the user list, and the activity that was reported.

Method

1. Count changes by user, by type and by week.
2. Compare that against what was reported as having been done.
3. Flag high-impact changes made without an obvious documented rationale.
4. Identify dormant periods on campaigns that were underperforming at the time.

Output — An activity summary, the reported-versus-actual gap, and a list of changes requiring explanation.

48 Landing Page Post-Click Checker

Check whether the problem sits after the click rather than before it.

Input required — Landing page report with conversion rate, bounce and speed; Core Web Vitals; page URLs; ad copy; device split.

Method

1. Compare conversion rate by landing page against the account average.
2. Check message match between what the ad promised and what the page headline says.
3. Check load speed and Core Web Vitals, mobile first.
4. Audit the checkout path for friction and for costs revealed late.

Output — Page-by-page findings with a fix list ordered by revenue impact.

Note. *A large share of what gets diagnosed as a media problem is a post-click problem. This skill is the check that prevents wasted optimisation effort.*

49 Test Backlog Prioritiser

Turn a pile of ideas into a sequenced queue.

Input required — The list of test ideas, traffic volume per area, estimated effect size, build effort, and current performance gaps.

Method

1. Score each idea on impact, confidence and effort.
2. Check each has enough traffic to produce a readable result at all.
3. Sequence tests so that overlapping tests do not contaminate one another.
4. Attach the kill and scale numbers produced by Skill 31.

Output — A ranked, sequenced test queue with owners, dates and decision criteria.

50 Executive One Pager Writer

Compress the account into what a decision-maker actually needs to know.

Input required — Period performance, targets, decisions taken, risks, and the next-period plan.

Method

1. Lead with the commercial outcome, not with platform metrics.
2. State performance against target with the variance explained in one sentence.
3. Name the three decisions taken and the three that need to be taken.
4. State risks with an owner and a date against each. Hold to one page.

Output — A one-page brief written in commercial language, suitable for a founder or a board.

6. Data Requirements

The system is only as good as the data pulled into it. Fifteen inputs cover all fifty skills. Ten come from the ad platform and can be pulled by anyone with account access; five come from the business and require the client to cooperate. The second group is where most implementations stall, and it is also where the commercial advantage sits.

6.1 Platform inputs

Input	Feeds skills	Cadence
Campaign performance report (cost, clicks, conv, conv value, IS metrics)	01, 03, 20, 27	Weekly
Search terms report	02, 08, 09, 10, 11, 12, 13	Weekly
Keyword report with match type	10, 12, 25	Fortnightly
Change history export	04, 05, 29, 47	Monthly
Campaign settings export or screenshots	05, 06, 24, 25	Quarterly
Geographic and device reports	02, 06, 23	Monthly
Hour and day-of-week report	22	Quarterly
Product feed export and Merchant Center diagnostics	14, 15, 16, 18, 19	Fortnightly
Product-level performance report by item ID	16, 17, 18	Fortnightly
Landing page report and Core Web Vitals	48	Monthly

6.2 Business inputs

These require the client. Getting them is a commercial conversation, not a technical one — and an agency that has them is doing a fundamentally different job from one that does not.

Input	Feeds skills	Typical owner
Cost of goods, shipping, gateway and packaging costs by SKU or category	32, 33, 35, 38	Finance
Order and returns export by region, category and payment method	36, 37, 39	Operations
Cohort repeat purchase data and customer lists	33, 34, 41, 42	CRM / eCommerce
Inventory and stock status feed	16, 18	Operations
Promotion, discount and festive calendar	07, 30, 36	Marketing

The data conversation is the pitch

Asking a prospective client for contribution margin by category and returns data by region does two things at once. It makes better bidding possible, and it demonstrates a level of commercial engagement that most competing agencies will not have shown. Several of GIE's strongest engagements began exactly this way.

7. Limitations and Governance

A system that produces fluent, confident, well-formatted output carries a specific risk: it is easy to trust. The controls below are not optional additions; they are part of the build.

7.1 Known limitations

- **Arithmetic must be verified.** Language models are unreliable at sustained numerical work. Every skill that produces a calculation — the whole of Pillars 5 and 6 — must have its output checked against a spreadsheet before it informs a decision. Where possible, the skill should output the formula alongside the result.
- **Skills read a snapshot, not the account.** Every output is only as current as the export it was given. Date ranges must be stated on every output, and stale exports must be rejected.
- **Correlation is not diagnosis.** A skill can identify that a segment underperforms. It cannot know that the region in question had a courier strike that month. Local context stays with the human.
- **The platform changes.** Campaign types, settings, reporting columns and bidding behaviour shift without notice. Thresholds and method steps need a review owner and a review date.
- **Small samples produce confident noise.** Every skill must state the volume floor below which it declines to conclude, rather than producing a recommendation from four conversions.

7.2 Client data handling

Running client data through an AI system requires the same discipline applied to any other processor relationship. The following are mandatory before this system touches a client account:

1. Client consent for AI-assisted processing, recorded in the engagement terms.
2. No personally identifiable customer data in any input. Customer lists are referenced by segment definition and size, never by row-level records. Order exports are stripped of names, contact details and addresses before use — region and pincode-level granularity is sufficient for every skill in the system.
3. Margin and cost data treated as commercially confidential, with access limited to the named account team.
4. Where client data crosses borders — an India-based team working on UK or EU accounts, for instance — a data processing agreement must be in place before this system touches that account. UK GDPR and PECR obligations sit with the controller, and an AI-assisted workflow does not change that.

7.3 Review discipline

Control	Requirement
Human sign-off	No skill output is implemented without a named reviewer. The reviewer's name goes on the change log entry.
Calculation check	Any output containing a target, ceiling or budget figure is independently recalculated before use.

Control	Requirement
Threshold review	All numeric thresholds reviewed quarterly against portfolio evidence, with changes versioned.
Output versioning	Every skill file carries a version number and a last-reviewed date. Outputs cite the version used.
Escalation	Where a skill's escalation rule triggers, the output states clearly that the question cannot be answered from the input given.

8. Build Roadmap

Building fifty skills at once produces fifty mediocre skills. The sequence below front-loads the ones that pay back fastest and defers the ones that depend on client data we do not yet hold. Total elapsed time is roughly one quarter of part-time effort alongside live account work.

Phase 0 — The data routine (Week 1)

Before a single skill is written, fix the export routine. Define the exact reports, columns, date ranges and file naming for every input in Section 6.1, and assign the pull to a named person on a fixed day. **Without this, nothing downstream runs.** This is deliberately the first phase and the least interesting one.

Phase 1 — Waste and search terms (Weeks 2–4)

Build Skills 01, 02, 08, 09, 10 and 12. Six skills covering the account health scan, wasted spend, and the search terms cluster. These are the fastest to build, the easiest to validate, and they recover money in the first month, which funds internal support for the rest of the programme.

Validation gate: run all six against three live accounts. The wasted spend identified must be independently verifiable before proceeding.

Phase 2 — Measurement integrity (Weeks 5–7)

Build Skills 39, 40, 41, 44 and 47. Establish whether the numbers can be believed before building anything that acts on them. Expect this phase to surface uncomfortable findings on at least one existing account; that is the point of running it.

Phase 3 — Margin and bidding (Weeks 8–11)

Build Skills 32, 33, 26, 27, 28 and 35. This is the commercially significant phase and the slowest, because it requires client cost data. Begin the data conversation with three willing clients in Phase 1 so the inputs exist by the time the skills do.

Validation gate: the calculation check in Section 7.3 must be running before any target derived from these skills goes live in an account.

Phase 4 — Feed and structure (Weeks 12–15)

Build Skills 14 through 21 and 25. Sequence Skill 19, the custom label builder, first, since the rest of the feed and structure work depends on the label schema.

Phase 5 — Accountability and reporting (Weeks 16–18)

Build Skills 45 through 50, plus the remaining audit and settings skills. Then invert Pillar 8: instead of using these to check others, use them to check ourselves, and publish the results to clients as a standing part of the monthly reporting pack.

Phase	Skills built	Total	Primary outcome
0 — Data routine	—	0	Reliable, scheduled, correctly-specified exports
1 — Waste and search terms	01, 02, 08, 09, 10, 12	6	Recovered spend, funds internal buy-in
2 — Measurement integrity	39, 40, 41, 44, 47	11	Numbers that can be trusted
3 — Margin and bidding	32, 33, 26, 27, 28, 35	17	Profit-based targets; the core differentiator
4 — Feed and structure	14–21, 25	26	Catalogue visibility and account architecture
5 — Accountability	45–50 and remainder	50	Published transparency as a positioning asset

9. Productising This at Global Info Edge

The system has two commercial uses. The first is internal: a quality assurance layer that runs under every managed account and raises the floor on consistency across the team. The second is external: a set of defined, scoped, saleable deliverables. The second is only credible once the first is real.

9.1 Internal use — the QA layer

Embed the skills into standard operating procedure rather than running them as a separate initiative. Each becomes a step in a documented routine with a named owner, a cadence and an output that lands in the account file. This is the real reason to build them: a skill is what turns an SOP from aspirational into checkable.

The measurable outcome is variance reduction. Across a portfolio of any size, the gap between the best-run account and the median-run account is usually the single largest available performance gain — and closing it does not require winning a single new client.

9.2 External use — productised deliverables

Three offerings follow naturally from the pillar structure. All three work best scoped, fixed-price and time-boxed, because the value is in the discipline rather than the hours.

A. Paid Media Diagnostic

A fixed-scope audit built on Pillars 1, 2, 4 and 7. Delivered as a written findings document with a ranked recovery list and a reconciliation statement. Works as a standalone engagement and as the strongest available entry point into a new account — it produces a verifiable number for money currently being wasted, which is a far more persuasive opening than a capability deck.

B. Profit-Led Media Management

A retainer tier distinguished by Pillar 6. Targets derived from contribution margin, CAC ceilings set against payback tolerance, profit passed as conversion value, and performance reported net of returns. Requires client cost data, which is precisely why it is defensible — most competitors will not ask for it, and clients who supply it do not readily switch.

C. Media Accountability Review

An independent review of an incumbent agency or in-house team, built on Pillar 8. Commercially delicate and should be offered selectively, but it is a real need in the market and there are few credible providers. It also positions GIE as the party willing to be measured by the same standard.

9.3 The transparency position

The most valuable move available is also the simplest: publish the answers to Pillar 8 before any client asks the questions. Every monthly report carries brand and non-brand split as standard, a conversion reconciliation statement, a change log with the reason recorded against each change, and last month's commitments reported against.

This costs process time and no margin. It makes you the difficult supplier to replace, because the replacement will not be offering it, and the client will notice within one reporting cycle. The principle generalises: the recommendation that visibly costs you something is the one that gets believed.

Positioning line

Every agency says they are transparent. Very few publish the reconciliation between what the platform reported and what the client actually banked. That single artefact, delivered monthly without being asked for, does more competitive work than any credentials slide.

9.4 What this is worth

Two effects, both realistic:

- **Retention.** The Pillar 8 transparency position and the Pillar 6 data dependency both raise switching costs materially. A client whose margin data is embedded in our bidding logic does not move easily.
- **Margin.** Skills compress analyst hours on the reading and preparation work, which is the majority of servicing time. The saved hours should be redeployed into the strategic layer clients actually pay for, not billed out as reduced scope.

10. Recommendations

#	Recommendation	Owner	Timing
1	Fix the export routine before building anything. Define reports, columns, cadence and naming for all fifteen inputs in Section 6.	Performance media lead	Week 1
2	Build Phase 1 (six skills) and validate against three live accounts. Do not proceed until the wasted spend identified is independently verified.	Performance media lead	Weeks 2–4
3	Run Skill 39, conversion reconciliation, across the top twenty accounts by spend. Treat any material unexplained gap as a stop-work item.	Account management	Weeks 5–7
4	Open the cost data conversation with three willing clients now, so Phase 3 has inputs when it starts.	Founder / account management	Week 2 onward
5	Add brand versus non-brand split and a reconciliation statement to the standard monthly report for every account.	Reporting	Week 6
6	Fold the skills into the existing SOP track rather than running them as a parallel initiative.	Operations	Ongoing from Week 4
7	Scope and price the Paid Media Diagnostic as the first productised offering — fixed scope, fixed fee, fixed turnaround.	Founder / commercial	Week 8
8	Put a data processing agreement in place before this system touches any cross-border client data.	Founder / legal	Blocking

Closing position

The published system is a good piece of work and a better piece of positioning. Its real contribution is not the fifty skills; it is the argument that the checking work is the work, and that it deserves to be systematised rather than squeezed into whatever time is left at the end of the week.

That argument is correct. And the raw material these checklists need is not technology — it is pattern knowledge, accumulated one account at a time. Fifteen years, ₹50 Cr+ of managed ad spend, 2.5 lakh+ leads and 250+ clients is what fills in the thresholds. The skills are only the delivery mechanism. The patterns are the product.

Build it internally, use it to raise the floor across every account you touch, and then let clients see the outputs — including the ones designed to hold you to account. That last part is the whole position.

Appendix A — Master Skill Index

All fifty skills in numerical order, with pillar and build phase.

#	Skill	Pillar	Phase
01	Account Health Scanner	P1	1
02	Wasted Spend Finder	P1	1
03	Impression Share Diagnoser	P1	5
04	Learning Phase Auditor	P1	5
05	Auto-Applied Recommendations Checker	P1	5
06	Geographic Targeting Reviewer	P1	5
07	Seasonality Checker	P1	5
08	Search Term Report Reader	P2	1
09	Negative Keyword List Builder	P2	1
10	Brand vs Non-Brand Splitter	P2	1
11	Query Intent Classifier	P2	5
12	Match Type Reviewer	P2	1
13	Competitor Query Auditor	P2	5
14	Product Title Rewriter	P3	4
15	Feed Attribute Auditor	P3	4
16	SKU Coverage Checker	P3	4
17	Shopping Bid Segmenter	P3	4
18	Out of Stock Guard	P3	4
19	Custom Label Builder	P3	4
20	Campaign Structure Mapper	P4	4
21	PMax Asset Group Planner	P4	4
22	Ad Schedule Reviewer	P4	5
23	Device Bid Analyser	P4	5
24	Network Settings Checker	P4	5
25	Duplicate Coverage Detector	P4	4
26	Target ROAS Setter	P5	3

#	Skill	Pillar	Phase
27	Budget Reallocation Planner	P5	3
28	Spending Ceiling Calculator	P5	3
29	Bid Change Log Keeper	P5	5
30	Q4 Budget Ramp Planner	P5	5
31	Test Kill Number Setter	P5	5
32	Contribution Margin Calculator	P6	3
33	CAC Ceiling Setter	P6	3
34	Payback Period Modeller	P6	5
35	Profit as Conversion Value Setup	P6	3
36	Discount Stacking Auditor	P6	5
37	Return Rate by Market Analyser	P6	5
38	AOV and Bundle Planner	P6	5
39	Conversion Reconciliation Skill	P7	2
40	Attribution Window Reviewer	P7	2
41	Repeat Buyer Conversion Filter	P7	2
42	Audience Exclusion Builder	P7	5
43	Offline Conversion Import Guide	P7	5
44	Consent Mode Checker	P7	2
45	Monthly Report Interrogator	P8	5
46	Agency Question Generator	P8	5
47	Change History Reviewer	P8	2
48	Landing Page Post-Click Checker	P8	5
49	Test Backlog Prioritiser	P8	5
50	Executive One Pager Writer	P8	5

Appendix B — Skill File Template

The standard structure for every skill file built under this system. Copy, complete, version, and store in the shared skills library.

Section	Content
SKILL NAME	Verb-noun. Names the job, not the topic.
VERSION / REVIEWED	v1.0 · reviewed DD-MM-YYYY · owner: [name]
WHEN TO USE	The trigger condition. One or two sentences.
REQUIRED INPUT	Exact report name, required columns, date range, and file format. Refuse to proceed if any required column is missing.
VOLUME FLOOR	The minimum data volume below which this skill declines to conclude.
THRESHOLDS	Every numeric rule this skill applies, stated explicitly.
METHOD	Ordered steps. Two operators running this must reach the same answer.
OUTPUT FORMAT	The fixed shape of the response. Ranked by rupee impact. Action named, not described.
ESCALATION RULE	What this skill must refuse to conclude without, and which skill supplies it.
DEPENDENCIES	Skills that must be run first, and skills that consume this output.
KNOWN FAILURE MODES	Where this skill has produced wrong answers before, and the check that catches it.

About the Author



Chandan Kumar

Founder — Global Info Edge

New Delhi, India · globalinfoedge.in · linkedin.com/in/ckmehta

Chandan Kumar is the founder of **Global Info Edge**, a full-funnel digital marketing and growth agency he established in New Delhi in 2012. He has spent over fifteen years building lead generation and performance marketing systems for B2B, B2C and D2C brands, with a working focus on Google Ads, Meta Ads, PPC and search engine marketing.

He holds an MBA in Marketing and Information Technology, is a Microsoft Bing Ads Accredited Professional and a certified Digital Marketing Specialist. Outside the business he serves with Lions Clubs International through the Lions Club of Garhwa Green, working on education and community development for underprivileged students.

Track record	
Ad spend managed	₹50 Cr+
Qualified leads generated	2.5 lakh+
Client revenue influenced	₹126 Cr+
Clients served	250+
Years in performance marketing	15+

About Global Info Edge

Global Info Edge is a **Google Premier Partner 2024** and **Meta Business Partner**, operating from Saket, New Delhi, with teams in Noida, Garhwa and Ranchi. The agency runs paid media, SEO and local search, conversion-focused web and eCommerce build, online reputation management, and AI and WhatsApp automation for growing brands.

If this was useful

This paper is free to read, quote and share with attribution. If you would like the accompanying build tracker, or you want a second pair of eyes on your own ad account against the fifty checks in Section 5, connect with me on LinkedIn at linkedin.com/in/ckmehta or write to Support@globalinfoedge.com.



Chandan Kumar

Founder, Global Info Edge

If you run, manage or buy paid media, I would rather you
argue with this document than agree with it quietly.

[linkedin.com/in/ckmehta](https://www.linkedin.com/in/ckmehta)

GLOBAL INFO EDGE

Saket, New Delhi · Support@globalinfoedge.com · +91 91559 93311
globalinfoedge.in · globalinfoedge.com

Brand. Build. Grow.