SAMPLE DELIVERABLE

Apple Search Ads account audit.

A worked example of the audit I run before recommending a single bid change — API and MMP data joined, structure taken apart, then leakage, headroom, brand cover and geo spill in one prioritised pass.

Illustrative client	Window	Storefronts	Scope
Lumen Fitness (fictional)	**90 days**	**US · CA · UK · AU**	**6 campaigns · 1,240 keywords**

This is a sample, not a client report. Every number, keyword, campaign name and finding in this document is synthetic and generated for illustration. Lumen Fitness does not exist. No client or employer data appears anywhere in this file.

Sara Togashi · Seattle, WA

01 · EXECUTIVE SUMMARY

Six problems, ranked by what they cost.

$184.2k	12,480	$14.76	$38.6k
SPEND, 90D	INSTALLS	BLENDED CPI	RECOVERABLE

The account isn't underfunded — it's leaking, and the structure is what's letting it. Roughly **21% of spend** is buying traffic the account already owns or doesn't want, and the reporting layer is hiding it. Nothing below needs a budget increase.

The activation event is broken — read this first

HIGH

Your activation event is mapped to `app_open`, so it fires on every returning session and reports ~3× more activations than installs. Every activation rate your dashboards show is inflated, and any keyword tuned against it has been optimising for retention, not acquisition.

Evidence: 41,200 activations against 12,480 installs (3.30×) · figures in this report re-baselined · see p.8

Structure is doing the damage, not the bids

HIGH

One campaign per storefront with mixed match types and no ad-group separation by intent. Brand, category and competitor traffic all clear against the same budget, so no line item can be judged on its own.

Evidence: 6 campaigns covering 4 storefronts · 3 mixed-match campaigns · see p.4

Discovery is re-buying exact-match traffic

HIGH

No negative list mirroring the exact campaigns. 31% of Discovery spend landed on terms already covered by an active exact keyword — bid twice, credited once.

Evidence: 412 of 1,331 Discovery search terms matched an existing exact keyword · $19.8k

Spend spilling into unsupported storefronts

HIGH

Multi-storefront campaigns are serving in markets with no localised onboarding. Installs land, activation doesn't. The blended CPI looks fine because the cheap installs mask it.

Evidence: 14.8% of spend in storefronts below half the account activation rate · see p.7

Bid headroom misread as demand ceiling

MEDIUM

Nine keywords flagged internally as "capped" and queued for increases. Only three are actually demand-capped; six sit well below their own average CPT and need a bid change, not more budget.

Evidence: CPT-to-bid compared against impression share per keyword · see p.6

Brand terms uncovered at the top of the funnel

LOW

Brand defence runs on broad match and stops at the head term. Competitors are appearing on brand-adjacent queries the account never bids on.

Evidence: 28% of Brand impressions from non-brand queries · 11 uncovered brand-adjacent terms

02 · HOW THE AUDIT RUNS

Two data sources, joined. Neither is enough alone.

The audit is run off the **Apple Search Ads API** paired with whatever **MMP** you already use — AppsFlyer, Adjust, Branch, Singular or Kochava. The API tells me what you *paid*. The MMP tells me what you *got*. Judging keywords on one without the other is how accounts end up optimised toward cheap installs that never activate.

APPLE SEARCH ADS API

- Keyword-level bids, CPT and impression share at full account scale
- Search-term reports well past what the UI will export
- Historical daily series — needed to tell a throttle from a demand cap
- Campaign, ad group and negative structure as configured
- Programmatic bid and negative changes once we agree the actions

YOUR MMP

- Post-install events — activation, first payment, whatever your funnel is
- Cohort retention and payback by acquisition source
- Geo-level attribution, to catch storefront spill
- Cross-channel view, so ASA isn't credited for what other channels drove
- SKAdNetwork and attribution-window sanity checks

The join is the point. Keyword-level spend from the API, matched to downstream value from the MMP, is what lets me say "raise this bid, cut that one" with a reason attached. Without the MMP side, every recommendation collapses back to CPI — and CPI is the metric that got the account here.

What I'll tell you I can't see

Two honest limits, stated up front rather than buried. **SKAdNetwork coarsens iOS post-install data** — some keyword-to-event joins are directional, not exact, and I'll label those rather than dress them up. And if your MMP isn't firing the events that matter to your business, no audit fixes that; getting those events right is usually the first recommendation, not the last.

ACCESS NEEDED

Read-only, and nothing from engineering.

API credentials scoped to read, plus a read seat on the MMP. The audit itself needs no engineering time. If we move to executing changes, that's a separate scoped permission and a separate conversation — I don't take write access by default.

03 · ACCOUNT STRUCTURE

Most accounts don't have a structure problem. They have no structure.

This is where the majority of recoverable spend hides, and it's the part most accounts have never had audited. Structure decides what you're *able* to measure. If brand, category and competitor intent share one budget and one bid, no report you run afterwards can separate them — and every optimisation is guesswork dressed as data.

CAMPAIGN (SYNTHETIC)	STOREFRONTS	MATCH	SPEND	KEYWORDS	ISSUE
Lumen — Brand	US, CA	Broad + Exact	$31.4k	14	Mixed match
Lumen — Category	US, CA, UK, AU	Broad	$58.9k	186	Geo spill
Lumen — Competitor	US	Exact	$22.7k	41	Clean
Lumen — Discovery	US, CA, UK, AU	Search match	$44.1k	—	No negatives
Lumen — Generic UK	UK	Broad + Exact	$18.2k	312	Duplicates
Lumen — Retargeting	US	Exact	$8.9k	687	Bloated

What's actually wrong

- **No intent separation.** Brand and category share bidding logic in two campaigns
- **Four storefronts, one campaign.** Category can't be bid or budgeted per market
- **Discovery has no negative list** mirroring exact coverage
- **37 duplicate keywords** across three campaigns, self-competing
- **687 keywords in Retargeting** — an unpruned dump, not a strategy

What it should look like

- One campaign per **storefront × intent** — brand, category, competitor, discovery
- Exact-only in intent campaigns; search match isolated to Discovery
- Discovery negatives auto-mirrored from exact coverage
- One owning campaign per keyword, enforced
- Ad groups split by theme so creative sets can be matched to intent

Why this comes first. Restructure before re-bidding. A bid change on a leaking account just moves the leak somewhere harder to see — and you'll have burned the budget proving it.

04 · STRUCTURE, DRAWN

What you have, and what it should be.

The left side isn't a caricature — it's the most common shape I find. It works right up to the point where someone asks "which of these is actually working," and then nothing can answer.

TODAY

6 campaigns, no intent separation

Lumen — Category

$58.9k · 186 keywords · broad

US CA UK AU

Brand Category Competitor

Broad Exact Search match

Three intents, four markets, one bid.

One budget · one bid · one report

"Which of these is working?"

Unanswerable. Nothing is isolated.

- Brand CPT inflated by category traffic
- Cheap AU installs mask expensive US ones
- 37 keywords bidding against themselves
- No negative path from exact to Discovery
- Any bid change moves the leak, not fixes it

RECOMMENDED

Storefront × intent — each cell its own line item

	BRAND	CATEGORY	COMPET.	DISCOVER
US	exact	exact	exact	match
UK	exact	exact	exact	match
CA	exact	exact	exact	match
AU	exact	exact	exact	match

Negatives mirror automatically

Every exact keyword becomes a Discovery negative

16 line items you can actually read

Each independently bid, budgeted and measured

- Brand CPT is brand CPT — nothing else in it
- Per-market bids; leaking storefronts isolatable
- One owning campaign per keyword, enforced
- Discovery only ever finds genuinely new terms
- Bid changes land where you aimed them

Migration is additive, not a teardown.

The new structure is built alongside the old one. Nothing is paused until the replacement is proven.

WEEK 1

Build alongside

16 new line items created at low bids. The old campaigns keep running, untouched.

WEEK 2

Shift the budget

25% → 50% → 75% as the new set's CPA stabilises. Roll back in a day if it doesn't.

WEEK 3

Archive the old

Old campaigns paused, not deleted — history stays intact for year-over-year reporting.

There is no dark period and no lost learning — at every point in the three weeks, the account is fully live.

05 · BID HEADROOM

Check the ceiling before you raise the bid.

A keyword spending its full allowance isn't necessarily demand-capped. If average CPT sits well under the bid, it's winning auctions cheaply and simply isn't being shown more — a relevance or budget constraint, not a bid one. Raising there buys nothing.

KEYWORD (SYNTHETIC)	BID	AVG CPT	CPT/BID	IMPR. SHARE	ACTIVATION	READ
home workout app	$3.40	$3.31	97%	61%	38%	Raise
strength training log	$2.80	$2.74	98%	58%	41%	Raise
hiit timer	$2.15	$2.09	97%	64%	35%	Raise
fitness tracker app	$4.10	$2.38	58%	44%	29%	Hold
gym planner	$3.75	$1.92	51%	39%	31%	Hold
calorie counter	$3.90	$1.61	41%	28%	12%	Cut
meal prep tracker	$2.45	$0.88	36%	19%	9%	Cut

How to read it. CPT/bid above ~90% with impression share under ~70% is real headroom — competing and losing on price, so raise. Under ~60% means the bid already clears comfortably and the constraint is elsewhere. The activation column is the MMP join doing its job: the two "cut" rows are cheap *and* convert at a third of the account rate, which is a targeting problem no bid change fixes.

06 · BRAND COVER

Defend the head term, then the shoulder.

Brand defence is the cheapest conversion volume in the account and the easiest to under-build. Two failures here: broad match is letting non-brand queries clear against brand budget, and cover stops at the head term while competitors take the brand-adjacent queries.

Fix now

- Move 4 brand keywords broad → exact
- Add 11 brand-adjacent terms as exact
- Separate brand into its own storefront campaigns

Depends on your tooling

Share-of-voice and competitor-cover depth vary with what you already license. With a third-party SOV tool I can size the gap precisely; without one, I work from search-term reports and impression share, which is directional but enough to act on. I'll tell you which of the two you're getting.

07 · STOREFRONT LEAKAGE

Where the money goes that nobody budgeted.

Multi-storefront campaigns spend where demand is cheapest, not where the product is ready. Cheap installs in unsupported markets drag the blended CPI *down*, which is why this hides so well — the headline metric improves while the funnel gets worse. Only the MMP-side geo view catches it.

STOREFRONT	SPEND	SHARE	CPI	ACTIVATION	VS ACCT	READ
United States	$96.4k	52%	$18.20	37%	1.06×	Core
United Kingdom	$38.1k	21%	$15.40	33%	0.94×	Core
Canada	$22.4k	12%	$13.90	31%	0.89×	Hold
Australia	$14.9k	8%	$12.10	29%	0.83×	Hold
Ireland	$6.8k	4%	$7.40	16%	0.46×	Leak
New Zealand	$3.4k	2%	$6.90	14%	0.40×	Leak
Other (11)	$2.2k	1%	$5.10	11%	0.31×	Leak

$12.4k over 90 days is going to storefronts converting at under half the account rate, none of which were deliberately targeted — they came in with the multi-storefront campaign defaults. Two options: split them out with their own bids and a localisation plan, or exclude them. Either is fine. Leaving them inside a campaign you judge on blended CPI is not.

08 · SEARCH-TERM COVERAGE

What the account isn't bidding on at all.

Mining 90 days of Discovery search terms surfaced 46 queries with real volume and no matching keyword anywhere in the account. Eleven clear the activation threshold and are worth graduating to exact with their own bids.

- **6 queries** — clear intent, activation at or above account rate. Graduate to exact now
- **5 queries** — promising but thin volume. Run as a probing set for 14 days first
- **35 queries** — off-intent or irrelevant. Add as negatives to stop the bleed

09 · EVENT INTEGRITY

One broken event invalidates every number above it.

This check runs **before** any performance read, because if the conversion event is wrong then every CPA, every activation rate and every "good keyword" in the account is wrong in the same direction — and confidently so.

EVENT	FIRES ON	90D VOLUME	VS INSTALLS	STATUS
install	—	12,480	1.00×	OK
activation	**app_open**	41,200	3.30×	Broken
first_workout	workout_completed	4,120	0.33×	OK
subscription_start	purchase_success	892	0.07×	OK
trial_start	not implemented	0	—	Missing

The tell

An activation event cannot fire 3.3× more often than installs. It's mapped to `app_open`, so it counts every returning session as a fresh activation. Consequence: activation rates across the account are inflated roughly 3×, and any keyword optimisation driven off that metric has been rewarding retention, not acquisition. The fix is one line in the MMP event mapping — but every historical read needs re-baselining after it lands, and the numbers will look *worse* before they look right.

Also missing: `trial_start` is never sent, so the funnel jumps straight from install to paid subscription. That's the single highest-value event for a subscription app and its absence is why nothing upstream can be optimised toward revenue.

10 · SEASONALITY

Flat budgets fight the calendar and lose.

Run only where there's **12+ months of history** — below that the index is noise. This account has 26 months, so the pattern is real. Monthly demand index, US storefront, three-year weighted:

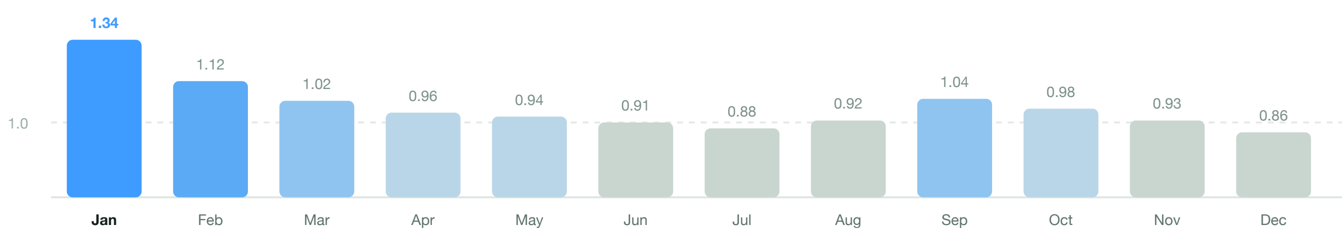

Index = demand relative to the annual mean. Fitness category; your curve will differ by vertical.

What it costs you. Budgets are set flat month to month, so January — at **1.34× demand** — is capped at the same daily allowance as July at 0.88×. You are rationing supply in the one month buyers come to you, and overpaying for scarce demand in midsummer. Same annual budget, redistributed on this curve, is free volume. Note this is a *demand* index, not a spend target — the pacing still has to respect payback.

11 · WHAT I'D ACTUALLY DO

Ranked by impact per hour of work.

Ordered so the first three ship in week one and pay for the rest. Structure before bids. Nothing here needs new budget.

1. **Split Category into storefront × intent campaigns** — Structural · 4 hrs
   Unlocks per-market bidding and stops brand, category and competitor sharing one budget.
2. **Mirror exact keywords as Discovery negatives** — ~$19.8k · 2 hrs
   Stops Discovery re-buying traffic the exact campaigns already own. Largest single recovery.
3. **Exclude or isolate the three leaking storefronts** — ~$12.4k · 1 hr
   IE, NZ and the 11-market tail. Split with their own bids, or cut — but stop blending them in.
4. **Brand: broad → exact, plus 11 brand-adjacent terms** — ~$6.1k · 1 hr
   Cleans competitor bleed out of brand budget and closes the shoulder-term gap.
5. **Raise bids on the three genuine-headroom keywords only** — Volume · 1 hr
   Cancel the six queued increases that would buy nothing. Re-read after 3 days of ramp.
6. **Graduate 6 proven queries, probe 5, negate 35** — Growth · 3 hrs
   Net-new coverage the account has never bid on, plus waste removal on the same pass.
7. **De-duplicate 37 keywords, prune Retargeting** — Hygiene · 2 hrs
   Reporting hygiene — makes every read after this one honest.

12 · KEEPING IT FIXED

An audit is a snapshot. Accounts drift back.

Every finding above will re-accumulate within a couple of quarters unless something maintains it. For smaller teams without an in-house paid-search engineer, that's usually worth automating rather than re-buying an audit every six months.

INCLUDED

The audit

This document, the keyword-level workbook, and a 60-minute walkthrough. Actions come as a ranked list you can hand straight to whoever runs the account.

ADD-ON

API access, wired up

Your ASA API and MMP connected into a repeatable pull, so your team can re-run these same checks themselves instead of waiting on me.

SCOPED SEPARATELY

ADD-ON

Custom automations

Built to your account: self-competition detection, negative-list maintenance, discovery mining, and bulk bid changes pushed through the API — so a course-correction is one approval, not three weeks of manual edits.

SCOPED SEPARATELY

Pricing. Scoped to the problem, not sold off a rate card — account size, storefront count and whether you want the automations drive most of it. Send me a rough shape and I'll come back with a number, usually within a couple of days. · linkedin.com/in/sarasanantonio