Cross-Platform Search Allocation

A Methodology Framework + Reference of Typical Findings

Paid search programs at scale typically run across multiple platforms, each reporting conversions on its own attribution basis. The cost-per-conversion figures the platforms report are internally coherent and structurally not comparable across platforms. Cross-platform search allocation analysis surfaces the actual cost differential, sized in dollars, by reconciling conversions independently of any single platform's reporting. These twelve questions describe what a credible search allocation methodology requires. The reference grid on page 3 shows what the analysis typically surfaces.

SECTION 1 — CROSS-PLATFORM VISIBILITY

1. Are cost-per-conversion figures from each search platform comparable on the same basis, or do platform attribution differences make direct comparison structurally impossible?

■ Yes ■ Partial ■ No

Why it matters: Each platform applies its own attribution rules, conversion windows, and credit logic. Comparing the resulting CPAs is comparing internally coherent numbers that describe different events. Cross-platform comparison requires reconciliation onto a common attribution basis before any allocation decision is sound.

2. Is search spend currently allocated across platforms based on platform-reported CPA, or based on independently deduplicated cost-per-acquisition?

■ Yes ■ Partial ■ No

Why it matters: Allocation driven by platform-reported numbers inherits each platform's orientation. Independent deduplication produces a different CPA for the same conversions, often by a substantial margin. The allocation decision is only as good as the inputs it runs on.

3. Is the dollar-volume cost difference between platforms surfaced at the channel level — and also at the day-of-week, daypart, audience, and campaign level?

■ Yes ■ Partial ■ No

Why it matters: Aggregate-level allocation finds aggregate-level inefficiency. Day-of-week, daypart, and audience-level analysis surfaces sub-program patterns that channel-level reporting averages away. Independent cross-platform analysis at this granularity has surfaced CPC differentials above 140% within the same program.

SECTION 2 — ATTRIBUTION DEDUPLICATION

4. When two or more search platforms claim credit for the same conversion, what process resolves the overlap?

■ Yes ■ Partial ■ No

Why it matters: In four-channel programs with overlapping platform attribution, summed reports commonly run two to four times actual conversions. The reconciliation method that resolves the overlap defines the credibility of every downstream allocation decision. A documented dedupe process is the entry point to credible cross-platform comparison.

5. Is the deduplication run by a party with no financial relationship to either platform under comparison?

■ Yes ■ Partial ■ No

Why it matters: Deduplication run by a platform favors that platform; deduplication run by an agency favors the agency's commercial mix. Independent dedupe — by a party with no stake in which platform wins — produces a different number than either of the alternatives. The structural posture of the deduplicator matters more than the dedupe technique.

6. **Is the deduplication output auditable end-to-end — can a CFO or operating partner walk through how a single conversion was attributed to one platform over another?**

■ Yes ■ Partial ■ No

Why it matters: Black-box dedupe survives in stable markets and breaks in disruption. The path from raw conversion event through deduplication logic to attributed credit should be documented and defensible to a non-specialist. An auditable trace is what separates allocation analysis from a vendor's slide.

SECTION 3 — ROLE AND TIMING ANALYSIS

7. **Does the analysis distinguish between Originator (demand creation), Roster (consideration), Assist (mid-journey), and Converter (closing) roles by platform?**

■ Yes ■ Partial ■ No

Why it matters: Standard reporting credits the platform closest to conversion. The same platform may be doing the demand creation work that an upper-funnel platform is being credited for, or the reverse. Role classification surfaces these dynamics; credit-only reporting masks them.

8. **Are day-of-week and daypart efficiency patterns surfaced in the analysis, or aggregated away?**

■ Yes ■ Partial ■ No

Why it matters: Aggregate CPA over a quarter is one number. Day-of-week CPA on the same program can vary by more than 100%. Bidding and scheduling decisions made on aggregate numbers leave the sub-aggregate inefficiency in the program. The analysis should preserve the granularity that the decisions need.

9. **Is platform performance measured against constant audience and offer parameters, isolating allocation efficiency from creative and audience effects?**

■ Yes ■ Partial ■ No

Why it matters: A platform's reported CPA depends on the audience and offer being run on it. Comparing platforms under different audience-and-offer conditions confuses platform efficiency with audience effects. Holding those variables constant is what makes the allocation finding attributable to the platform, not to the campaign setup.

SECTION 4 — ALLOCATION DECISION SUPPORT

10. **Does the analysis surface specific dollar-volume reallocation opportunities, not just CPA differentials?**

■ Yes ■ Partial ■ No

Why it matters: CPA differentials are diagnostic; dollar-volume reallocation is actionable. The translation from "Platform A's CPA is 60% lower than Platform B's" to "shifting $X of monthly spend produces $Y of recovered efficiency" is the work that turns analysis into a budget decision.

11. **Is the reallocation guidance conservative enough to act on without overshooting the inefficiency band?**

■ Yes ■ Partial ■ No

Why it matters: The inefficiency band typically falls between 6% and 23% of program spend. Overshooting that band creates the opposite problem — efficiency lost by moving too much spend onto the lower-CPA platform. Conservative reallocation that captures the documented band protects against the overcorrection that aggressive reallocation can produce.

12. **Does the engagement come with a defined remediation path — implementation guidance, success metrics, time horizon — rather than ending at the report?**

■ Yes ■ Partial ■ No

Why it matters: A finding that doesn't move budget produces no value. The engagement that earns its keep is the one with documented next steps: which platforms to adjust, in what increments, over what time horizon, measured against what success criteria. Analysis without implementation guidance is half an engagement.

Search Allocation Reference Grid

What the cross-platform search allocation analysis typically surfaces, sized in ranges drawn from C3 client engagements. Specific numbers vary by program and category; the directionality and band shape are durable. Use the grid as a calibration reference when interpreting findings on your own program.

Finding type	Typical magnitude	What it means	Action it supports
Independently-attributed CPA differential between platforms	**50%–80%+**	Platform A and Platform B charge meaningfully different amounts per actually-attributed conversion.	Reallocate toward lower-CPA platform, in conservative increments.
Same-spend efficiency recovery	**6%–23%**	Reallocation produces the same conversion total at lower cost, with no budget change.	Capture as efficiency; recoverable dollars become available for other uses.
Day-of-week CPC differential within program	**Up to 140%+**	Some days run substantially more efficient than others; aggregate reporting averages this away.	Daypart-level bid management or scheduling adjustments.
Cross-platform conversion overlap (in 4+ channel programs)	**2x–4x actual conversions**	Summed platform reports count the same conversion multiple times; reported total inflates the picture.	Run independent dedupe before any allocation decision based on summed reports.
Platform role asymmetry (Originator vs. Converter)	**Varies by program**	The platform credited with conversion may not be the one creating demand; role analysis surfaces the difference.	Rebalance top-funnel investment by demand-creation role, not by last-touch credit.

How to read the result. Directionality is not predetermined — which platform runs more efficient varies by program, category, and audience. That variance is the credibility argument for independent analysis. A vendor that returns the same directional answer for every client is running a sales pitch; a vendor that returns different findings by program is running an audit.

For a personalized version of this analysis — sized in your own revenue, channels, and current measurement approach — take the 90-second Fit Assessment at **c3metrics.com/fit-assessment**. Same reference principles applied to your specific program, one page out the other side, ready to share with a CFO or operating partner.