

Conversion Architecture

Sequencing Online and Offline Conversions in Attribution

Conversion data arrives from several sources with different confidence levels. Digital deterministic conversions sit at the foundation; independently matched offline conversions sit at a structural ceiling; platform-imported and modeled conversions describe what the platform inferred. A well-built attribution program sequences these tiers by confidence, deduplicates holistically, and discloses which tier produced which credit. These twelve questions surface conversion architecture as a deliberate choice. The reference grid on page 3 describes the five confidence tiers and what each is for.

SECTION 1 — CONVERSION INVENTORY

1. What conversion types feed the attribution model — digital deterministic, online-to-offline matched, platform-imported, modeled, or some combination?

Yes Partial No

Why it matters: Programs commonly conflate conversion types under a single "conversion" label. Each type carries a different confidence profile, match rate, and credible use. The inventory step is the foundation; what follows depends on knowing what's in the inventory.

2. For each conversion type in the program, what data infrastructure produces the conversion record — site tag, CRM import, platform offline import, modeled inference?

Yes Partial No

Why it matters: The infrastructure producing a conversion record defines what the record can defensibly claim. A modeled conversion is not the same artifact as a tag-fired conversion, and treating them as identical in the attribution model produces output that doesn't match what the data actually says.

3. Is the confidence level associated with each conversion type documented, and is that documentation surfaced when reports are produced?

Yes Partial No

Why it matters: Confidence tier is the missing column in most attribution reporting. A report that mixes high-confidence and low-confidence conversions without tier labels obscures which findings to act on with conviction and which to treat as directional.

SECTION 2 — MATCH RATE TRANSPARENCY

4. For online-to-offline matching, what match rate does the program achieve, and against what denominator?

Yes Partial No

Why it matters: Match rates above 70% often reflect online-to-CRM matching across an already-identified customer base. Match rates around 40–50% commonly reflect platform offline-conversion imports. True independent offline attribution carries a structural ceiling around 4–20%, depending on category and data conditions. The match rate without its denominator is uninformative.

5. Has the vendor distinguished between the three match-rate categories — online-to-CRM, platform offline imports, true independent attribution — or aggregated them into a single program-wide figure?

Yes Partial No

Why it matters: A single aggregate match rate hides which method produced which match. Vendors marketing 80% match rates are usually combining types or measuring the easier denominators. The distinction by category is what makes the number defensible to a CFO or operating partner.

6. Are the structural ceilings on each match category disclosed in vendor documentation?

Yes Partial No

Why it matters: True independent offline attribution operates against ceilings set by shared IP addresses, device fragmentation, offline-only households, and the absence of deterministic POS signal. These ceilings are technical realities, not failure modes. A vendor that doesn't disclose them is presenting model output as data.

SECTION 3 — SEQUENCING AND DEDUPLICATION

7. Is the conversion architecture sequenced by confidence — digital deterministic first, independently matched offline second, platform-imported third, modeled fourth?

Yes Partial No

Why it matters: Sequencing by confidence means high-confidence conversions are credited first; lower-confidence tiers fill in around them rather than competing for the same conversions. The alternative — treating all tiers as equally weighted inputs — produces overcount and miscredit at scale.

8. Is holistic deduplication run across all conversion tiers, including modeled and platform-imported records?

Yes Partial No

Why it matters: Deduplication within a tier is straightforward; deduplication across tiers is the harder work. A digital conversion paired to a platform-imported offline conversion for the same buyer should be reconciled, not summed. The cross-tier dedupe step is where most programs leak conversion volume.

9. Where a digital conversion serves as proxy for a weakly-matched offline event, is the proxy relationship documented and surfaced in the report?

Yes Partial No

Why it matters: Proxy reasoning is legitimate when documented. A strong digital signal as proxy for a weakly-matched offline conversion is more defensible than a low-confidence offline match treated as deterministic. The disclosure that the proxy is in use is what separates rigor from sleight of hand.

SECTION 4 — INDUSTRY FIT AND OPERATIONAL DISCIPLINE

10. Has the conversion architecture been adapted to the specific industry — long sales cycles, offline-heavy categories, regulated verticals each warrant different sequencing?

Yes Partial No

Why it matters: There is no universal correct conversion architecture. Pharma, automotive, financial services, retail, and B2B each have different conversion topographies. A vendor running the same architecture across every client is running a template; a vendor adapting to industry conditions is running a method.

11. Is the structural ceiling on true independent attribution disclosed, and is the program operating within it or above it?

Yes Partial No

Why it matters: A program reporting independent offline attribution rates well above 20% should produce questions. Operating within the ceiling is honest. Operating above it requires explanation — usually disclosure that proxy, platform, or modeled data is mixed in.

12. Is the conversion architecture documented, replicable, and auditable — could a successor team replicate it from documentation alone?

Yes Partial No

Why it matters: Architecture that lives in the heads of a few practitioners breaks when those practitioners leave. Documented architecture survives team changes and produces consistent output over time. The replicability test is whether a new analyst could reproduce last quarter's report from the documentation.

Conversion Confidence Tier Reference

Five tiers, sequenced by confidence. A well-built attribution program credits from Tier 1 first and fills in with lower tiers where Tier 1 doesn't cover. Use the grid to map each conversion type in the current program to a tier, and document what each tier is being asked to do.

Tier	Type	Typical match rate	Best used for	Caveat
1	Digital deterministic	Effectively 100% within tagged scope	Foundation of the program; the credit layer everything else fills in around.	Coverage limited to tagged digital touchpoints and consented users.
2	Independent offline attribution	4%–20% (structural ceiling)	True online-to-offline attribution with no platform mediation.	Ceiling is technical; aggregate numbers above 20% typically blend in other tiers.
3	Online-to-CRM matched	40%–70%	Within-customer-base journey attribution; useful for retention programs.	Denominator is the known customer base; not equivalent to acquisition attribution.
4	Platform offline imports	40%–50%	Where platform-mediated offline conversion data is the only signal available.	Match operates on platform identity graph; orientation toward the platform's view.
5	Modeled / inferred	Not directly applicable	Proxy reasoning where direct measurement isn't possible; explicit when used.	Treat as directional; document the proxy relationship when surfaced in reports.

One approach is always wrong: the universal one. The right conversion architecture varies by industry, sales-cycle length, and data conditions. Pharma and automotive — long cycles, heavy offline — sequence differently from retail or B2B SaaS. A vendor running the same architecture across every client is running a template. The starting point is honest inventory of what's in the current program; the work is sequencing what's there by confidence.

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.