

Auditing Programmatic Delivery Quality

A 12-Question Framework + Four-Signal Pattern Reference

Platform-reported invalid traffic (IVT) metrics describe what the platform's quality systems saw and chose to surface. Independent delivery-quality auditing reads the raw signal — impression logs, beacon fires, conversion records — for patterns that platform reporting does not flag. These twelve questions surface the audit posture of a program before findings appear, and the four-signal framework on page 3 describes the patterns C3 looks for in client engagements. Bring them to any DSP, agency, or in-house programmatic conversation.

SECTION 1 — DATA ACCESS AND AUDITABILITY

1. Does the buyer have access to line-item-level delivery data — impression logs, beacon fires, conversion records — independent of platform-rolled reporting?

Yes Partial No

Why it matters: Platform-reported summaries hide the patterns that independent auditing surfaces. Raw delivery data is the input layer; summary reporting is a transformation of it. The audit only runs on the layer the buyer can see.

2. Is the chain of custody from impression through conversion documented end-to-end, and is the documentation available to the buyer?

Yes Partial No

Why it matters: A documented chain of custody is what separates a delivery report from an audit. Vendors who decline to walk the trace either lack the infrastructure or have a reason to keep the flow opaque. Both warrant follow-up.

3. Are tag implementations on the buyer's site auditable for completeness and integrity — coverage, fire timing, payload composition?

Yes Partial No

Why it matters: Audit findings rest on the integrity of the data the audit consumes. Untagged placements, partial-page coverage, and inconsistent payload formats compound through the analysis. A pre-audit tag inspection costs little and protects the credibility of every finding downstream.

SECTION 2 — SIGNAL-LEVEL RECONCILIATION

4. Are viewthrough beacon counts reconciled against impression counts at the line-item level, with anomalies flagged when ratios exceed expected bounds?

Yes Partial No

Why it matters: A beacon fires only when an ad is served. Under legitimate delivery, the beacon-to-impression ratio sits at or below 1.0. Ratios above 1.0 surface pixel stuffing, hidden ad stacking, or delivery reporting errors. C3's published audit flagged line items with ratios as high as 1.81 — patterns platform reporting did not surface.

5. Is there a peer-volume comparison framework that flags line items delivering substantially above peer norms under identical format, audience, and platform parameters?

Yes Partial No

Why it matters: Volume spikes happen — budget pacing, auction dynamics, seasonality all produce them. A line item delivering 20x peer volume under identical parameters is a different signal. Peer comparison surfaces outliers more reliably than absolute thresholds.

6. Are impression-spike timing patterns analyzed — late-night and overnight concentration, cross-platform coordination, single-hour peaks well above campaign baseline?

Yes Partial No

Why it matters: Coordinated overnight surges across independent platforms with identical buying parameters indicate supply-side dynamics rather than demand-side audience behavior. The timing signature of the spike is at least as informative as its magnitude.

SECTION 3 — COST AND CHANNEL CLASSIFICATION

7. Is fraud cost calculated as $CPM \times \text{fraud rate} \times \text{impressions}$, rather than fraud rate alone — and is the calculation surfaced in the report?

Yes Partial No

Why it matters: Impression-rate fraud and dollar-cost fraud are different metrics. The most visually dramatic impression-rate days are often not the highest dollar-cost days, because spike days tend to coincide with below-average CPMs. Remediation prioritization should anchor to dollar cost, which is the figure that supports make-good claims.

8. Is the JS/IMG channel split tracked within fraud events, as a diagnostic for source classification?

Yes Partial No

Why it matters: Near-complete concentration of fraudulent impressions in a single channel type — IMG at 99%+, for example — points to a targeted event on a specific inventory source rather than systemic invalid traffic. Channel concentration narrows the investigation from campaign level to specific placements.

9. Are below-the-radar fraud days surfaced explicitly, including days where rate is modest but CPM is high and cost is meaningful?

Yes Partial No

Why it matters: Rate-anchored thresholds miss the highest-cost fraud days when CPMs run above the campaign average. A 3-4% fraud rate at \$40 CPMs produces more dollar cost than an 8% fraud rate at \$12 CPMs. The audit should surface both.

SECTION 4 — INCENTIVE STRUCTURE AND INDEPENDENCE

10. Who runs the audit, and is the auditing party financially independent of the inventory under audit?

Yes Partial No

Why it matters: Audits run by parties with financial relationships to the inventory carry a different reliability profile than audits run independently. A platform auditing its own inventory has financial disincentives to surface findings that reduce confidence in that inventory. An agency auditing its own buys has similar disincentives, in a different direction. Independent reconciliation sits outside that financial chain.

11. If a physical impossibility appears in the delivery data — beacon count exceeding impression count, conversions before impressions — what is the documented process for surfacing and investigating it?

Yes Partial No

Why it matters: The process for handling impossible findings is more diagnostic than the findings themselves. A program with no escalation path for physical impossibilities is a program where impossibilities will accumulate quietly. The escalation path is the audit's operational backbone.

12. Is there a defined remediation workflow for findings — DSP conversations, make-good claims, supply-path adjustments — with documented outcomes?

Yes Partial No

Why it matters: Audit findings without remediation paths produce expensive PDFs. The value of an audit is the dollar recovery, supply-path improvement, and budget reallocation that follow from the findings. A program that finds but doesn't remediate captures only part of the audit's value.

Four-Signal Pattern Reference

The four signals below are the patterns C3 looks for in programmatic traffic-quality audits. Each describes a specific anomaly type, what it indicates, and the kind of delivery data needed to surface it. A program covering all four with documented reconciliation processes has the audit posture the twelve questions on pages 1-2 describe. A program covering one or two has known blind spots.

#	Signal	What it indicates	Data required
1	Viewthrough beacon ratio > 1.0	Pixel stuffing, hidden ad stacking, or delivery reporting errors at the impression level. Physical impossibility — a beacon fires only when an ad is served.	Line-item beacon fires + line-item impression counts, reconciled side by side.
2	Impression-spike timing pattern	Overnight or late-evening concentration coordinated across independent platforms suggests supply-side dynamics rather than audience behavior.	Hourly impression delivery by platform, baselined against preceding 10-14 days.
3	Peer-volume outlier	Line items delivering many multiples of peer norms under identical format, audience, and platform parameters surface as deliverable-suspect.	Comparable line-item set with shared format, platform, and audience targeting.
4	CPM-weighted fraud cost	True dollar cost of fraud is CPM × rate × impressions. Highest-rate days are often not highest-cost days because spike windows coincide with below-average CPMs.	Daily fraud-rate measurement paired with delivered CPM by day.

A fifth pattern worth noting. Single-hour events with near-complete channel concentration (99%+ IMG, near-zero JS) point to targeted incidents on specific inventory sources rather than systemic invalid traffic. Channel concentration narrows investigation from campaign level to placement level — a different remediation conversation than systemic fraud requires.

How to read the result. Patterns 1 and 4 are the most diagnostically useful in C3's experience: viewthrough beacon ratios surface delivery integrity failures the platform reporting layer can't see, and CPM-weighted fraud cost is the figure that supports remediation conversations with DSP partners. Patterns 2 and 3 surface what's worth investigating; patterns 1 and 4 size what's worth recovering.

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.