NetWebMedia

2026

ANALYTICS GUIDE

AI Attribution Modeling: From Last-Click to Revenue Truth

How to replace the attribution model that's been misdirecting your budget with one that reflects how customers actually buy

Priya Patel
Strategy Lead, NetWebMedia

13 pages
netwebmedia.com

Contents

Ongoing Model Maintenance and Re-Training Cadence How to keep an AI attribution model current as channel mix evolves, buying journeys shift, and new data becomes available.

EXECUTIVE SUMMARY

AI Attribution Modeling: From Last-Click to Revenue Truth

Attribution is the single most consequential analytical decision in performance marketing. Get it right, and budget flows to channels that drive revenue. Get it wrong, and you systematically reward the last thing a customer did before converting — which is often not the thing that made them decide to buy. For B2B organizations with long sales cycles, multiple touchpoints, and significant content investments that influence deals without getting credited, wrong attribution isn't a minor technical inconvenience — it's a systematic misallocation of resources that compounds quarterly. This guide walks through the attribution landscape in 2026: why rule-based models fail, what data infrastructure AI attribution requires, how to evaluate the leading platforms for B2B use cases, and how to actually change budget allocation when your attribution changes without burning down the org.

IN THIS GUIDE

- ✓ Why last-click and first-click attribution models systematically misdirect budget away from the channels doing the most buying influence
- ✓ The data infrastructure requirements you must have in place before AI attribution can work — skipping this kills implementation
- ✓ A side-by-side evaluation of Northbeam, Triple Whale, Rockerbox, and GA4 DDA for B2B use cases
- ✓ How to build a budget reallocation playbook from attribution findings without triggering channel-team resistance
- ✓ A stakeholder presentation framework that makes attribution changes defensible to a skeptical CFO or board

Who this is for: B2B CMOs, demand generation leaders, and marketing analysts who manage performance budgets and suspect their current attribution model is pointing them in the wrong direction.

SECTION 1

The Last-Click Problem: How Wrong Attribution Destroys Good Marketing Programs

Last-click attribution gives 100% of conversion credit to the last touchpoint before a conversion event. For B2B buyers with a 60-180 day consideration cycle, this produces a systematically distorted picture of what's actually working. The touchpoints most likely to appear last before a conversion event are brand search, direct traffic, and retargeting — channels that capture intent created by earlier awareness and nurture touchpoints. Last-click rewards the channel that captured the purchase decision, not the channels that built it. The consequences play out predictably: content programs get underfunded because their blog posts appeared five months before the deal closed; social advertising gets cut because it can't demonstrate last-click returns; email nurture sequences lose budget because they influence but rarely close. Meanwhile, brand search and retargeting capture disproportionate credit and disproportionate budget — reinforcing the model's own prediction.

The magnitude of the distortion scales with the length of your sales cycle. For ecommerce with same-session conversions, last-click is somewhat more defensible. For B2B enterprise with 6-18 month sales cycles, multiple decision-makers, and significant research behaviors, last-click attribution can misdirect 30-50% of the performance budget. The practical test: compare your last-click attribution report to your revenue data by cohort. If the channels getting the most credit in attribution are not the channels that show up most frequently in win/loss data and customer journey interviews, your attribution model has a systematic accuracy problem — not a data quality problem.

- Pull your current attribution report and compare top-credited channels to your sales team's win/loss data
- Interview five recent closed-won customers about what channels influenced their decision — compare to last-click credit
- Calculate what percentage of your budget is flowing to last-click top performers vs. influence-heavy channels

- Map your average sales cycle length — this is the primary variable in how badly last-click distorts your data
- Document the channels and programs your attribution model chronically undervalues — this becomes your reallocation target

> ***For B2B organizations with sales cycles over 60 days, last-click attribution can misdirect 30-50% of the performance budget — consistently funding the wrong activities for the wrong reasons.***

73% of B2B marketing leaders who have migrated to data-driven attribution report finding significant budget misallocation in their prior model

SECTION 2

Attribution Model Comparison: Rule-Based vs. Data-Driven vs. AI

Attribution models exist on a spectrum from simple rule-based heuristics to sophisticated AI-driven probabilistic models. Understanding the distinctions helps you select the model appropriate for your data maturity and business complexity. Rule-based models assign credit according to predetermined rules that don't vary based on observed behavior. Last-click, first-click, and linear (equal credit across all touches) are the most common. Time decay gives more credit to touches closer to conversion. Position-based (U-shaped, W-shaped) assigns elevated credit to first touch, last touch, and specified middle touches. Rule-based models are transparent, easy to explain, and require minimal data infrastructure. They fail because they make assumptions about customer behavior that are often wrong and are invariant — the same rule applies to a customer who took five days to buy as one who took eighteen months, to a $5,000 deal and a $500,000 deal. Data-driven models, by contrast, use statistical analysis of actual conversion path data to assign credit weights based on which touchpoints are empirically associated with higher conversion rates.

AI attribution models extend data-driven methodology with machine learning techniques — gradient boosting, neural networks, or Markov chain models — that can identify non-linear relationships between touchpoints, model interactions between channels, and update credit weights continuously as new conversion data arrives. AI models outperform statistical data-driven models when you have sufficient data volume (typically 3,000+ monthly conversions) and complex multi-channel buying journeys. For B2B teams with lower conversion volumes, a hybrid approach — using AI attribution at the aggregate level while applying data-informed rule-based models for channel-level decisions — often works better than a pure AI approach that overfits to limited conversion data. Model selection principle: the most sophisticated model is not always the right one. Start with the model your data can support, not the model your attribution vendor is selling.

- Rule-based models: appropriate for early-stage attribution maturity, simple buying journeys, low conversion volumes
- Linear or time-decay: better default than last-click for any B2B team with multi-touch journeys
- Data-driven statistical: requires 1,000+ monthly conversions to produce stable weights — below this, weights are noise
- AI/ML attribution: requires 3,000+ monthly conversions and 6+ months of clean touchpoint history
- Hybrid approach: AI at aggregate level, informed rules at channel level — practical for most mid-market B2B teams

> ***The most common attribution mistake is not using the wrong model type — it's using a sophisticated model without the data volume to support it, producing precise-looking outputs that are statistically unstable.***

3,000+ monthly conversions required for AI attribution models to produce statistically stable credit weights — below this volume, results are unreliable

SECTION 3

Data Requirements: What You Need Before AI Attribution Works

AI attribution is only as accurate as the data it's trained on. The most common reason attribution implementations fail to deliver is not model selection — it's data infrastructure gaps that make the training data incomplete, inconsistent, or unresolvable to individual users. The data requirements fall into three categories. Category one: touchpoint capture. Every channel interaction that might influence a buying decision needs to be tracked and stored with consistent user identifiers. This means UTM parameter standards across all channels, consistent campaign naming conventions, server-side event tracking for critical conversions (not just client-side pixels that disappear with cookie blocking), and CRM activity logging for offline touches like sales calls and events. Category two: identity resolution. Touchpoints from the same buyer across multiple sessions, devices, and channels need to be resolved to a single user ID. Without identity resolution, a buyer who does research on mobile, converts on desktop, and talks to sales by phone looks like three different people in your attribution data. The quality of your identity resolution is the primary determinant of AI attribution model accuracy.

Category three: outcome data. The attribution model needs to know not just which touchpoints happened, but which ones led to revenue outcomes — and ideally, which led to specific revenue amounts. For B2B, this means bidirectional CRM integration: attribution platform pulls opportunity and deal data from the CRM, and the CRM can receive attribution credit data to inform reporting. The checklist before starting an AI attribution implementation: consistent UTM standards deployed across all paid channels, server-side event tracking for macro-conversions (form fill, demo request, purchase), identity resolution capability (first-party ID graph, authenticated sessions, or CDP), CRM integration with opportunity-level data, and a minimum of six months of clean touchpoint history at adequate conversion volume.

- Deploy consistent UTM naming convention across all channels — every paid and owned touch needs standardized tagging
- Implement server-side event tracking for macro-conversions — do not rely solely on client-side pixels
- Audit your identity resolution capability: how many touchpoints can you resolve to a single buyer across devices?
- Connect attribution platform bidirectionally to CRM: pull opportunity data in, push attribution credit out
- Verify you have 6+ months of clean history at your target conversion volume before training AI models
- Document every tracking gap — channels or touchpoints not captured — these become known biases in your model

> ***Identity resolution quality is the primary determinant of AI attribution model accuracy. A team with mediocre model architecture and excellent identity resolution will outperform the reverse every time.***

6 months minimum clean touchpoint history required before training an AI attribution model — less than this produces unstable credit weights

SECTION 4

First-Party Data Infrastructure: Cookies, Consent, and CRM Integration

Third-party cookie deprecation, increasingly stringent consent requirements under GDPR, CCPA, and state privacy laws, and platform-level measurement restrictions (iOS tracking limits, walled garden reporting) have fundamentally changed the data environment for attribution. Teams that built attribution on third-party cookies and platform pixels are operating with progressively degrading data quality. The shift to first-party data infrastructure is not optional for teams that want accurate attribution in 2026 and beyond. First-party data infrastructure has four components. Component one: authenticated user identification. When users log in, submit forms, or authenticate in any way, you capture a first-party identifier that persists across sessions and devices. This is the highest-quality identity signal available. Component two: first-party event tracking. Server-side tracking infrastructure that captures conversion events directly from your server to your analytics platform, bypassing client-side blocking. GA4's Measurement Protocol and server-side GTM are the standard implementations. Component three: first-party cookie management with compliant consent capture. A consent management platform (OneTrust, Cookiebot, or equivalent) that captures opt-in/opt-out status per user and ensures tracking compliance across jurisdictions.

Component four: CRM as the attribution anchor. In B2B, the CRM is the most reliable source of buyer identity — it has authenticated contact records with associated activity logs, opportunity data, and deal outcomes. Attribution platforms that integrate deeply with Salesforce, HubSpot, or your CRM of record can use the CRM's identity graph as the primary resolution mechanism, supplementing (rather than depending on) web tracking data. This is especially valuable for B2B teams where significant buying behavior happens offline: sales calls, events, referral conversations, and email threads that don't appear in web analytics at all. The CRM-anchored attribution approach produces a more complete picture of the buying journey for complex B2B deals than any web-tracking-only model can — and it's inherently more resilient to the consent and cookie environment because it doesn't depend on it.

- Implement server-side event tracking for all macro-conversions — this is the highest-priority infrastructure investment
- Deploy a consent management platform and audit your current consent capture against GDPR and applicable state laws
- Configure GA4 with enhanced measurement and server-side collection before migrating to any paid attribution platform
- Integrate your CRM bidirectionally with your attribution platform — this is the identity resolution anchor for B2B
- Map the offline touchpoints in your buying journey that don't appear in web tracking — document them as known gaps
- Build a first-party ID capture mechanism: any form fill, demo request, or gate generates a persistent first-party identifier

> ***Server-side tracking is the highest-ROI infrastructure investment for attribution quality — it recovers 20-35% of conversion events lost to client-side blocking and delivers them with higher data integrity.***

20–35%

of conversion events recoverable with server-side tracking implementation vs. client-side-only measurement — this is the data gap that distorts your attribution model

SECTION 5

Platform Selection: Northbeam, Triple Whale, Rockerbox, GA4 DDA Compared

The attribution platform market has consolidated around a handful of strong options with meaningfully different strengths and B2B suitability. Northbeam is built for media-heavy advertisers with significant cross-channel paid spend. Its strength is media mix modeling integrated with pixel-level attribution — giving a view of channel performance that spans both probabilistic (media mix) and deterministic (user-level) methods. For B2B teams with $200K+ monthly paid spend across 4+ channels, Northbeam's media mix layer provides valuable triangulation on channel efficiency. Its weakness is organic and offline touch coverage — it's optimized for paid media attribution and requires additional configuration to incorporate content-influenced pipeline and offline sales activity. Triple Whale originated in ecommerce DTC attribution but has expanded capabilities for B2B and subscription businesses. It excels at creative-level attribution — tracking performance by ad creative and copy variant rather than just by channel. If your team is running significant paid social with creative testing, Triple Whale's creative analytics layer produces insights that channel-level platforms miss.

Rockerbox is the most B2B-native of the major platforms. Its multi-touch attribution model handles long sales cycles, offline conversion events, and CRM integration more cleanly than ecommerce-oriented platforms. Rockerbox's deduplication methodology across channels reduces the double-counting that inflates total attributed revenue in platforms that use non-deduplicated models. Best fit: B2B teams with complex buying journeys, significant offline activity, and Salesforce or HubSpot integration requirements. GA4 Data-Driven Attribution (DDA) is the baseline option and has real value for teams already invested in the Google ecosystem. Its advantage is no additional cost for teams on GA4, native integration with Google Ads, and a probabilistic model that outperforms rule-based GA4 models. Its limitation is B2B attribution — GA4 DDA is optimized for web conversion events and struggles with long sales cycles, offline activity, and cross-channel deduplication. For B2B teams with simple funnels and modest paid budgets, GA4 DDA is a legitimate starting point. For teams with complex journeys and significant pipeline value at stake, a dedicated platform is worth the investment.

- Northbeam: best for paid-media-heavy B2B teams with $200K+ monthly spend and media mix modeling needs
- Triple Whale: best for B2B teams running significant paid social with creative-level performance requirements
- Rockerbox: best for B2B teams with long sales cycles, offline activity, and deep CRM integration requirements
- GA4 DDA: best starting point for teams with simple funnels, Google ecosystem investment, and <$50K monthly paid spend
- Evaluation criteria: sales cycle length, conversion volume, offline activity share, CRM integration depth, paid media budget
- Pilot criterion: every platform should offer a 30-60 day proof of concept with your actual data before full contract

> ***Rockerbox is the most B2B-native attribution platform in the market — its handling of long sales cycles, offline events, and CRM deduplication addresses the specific failure modes of B2B attribution more directly than platforms built for ecommerce.***

4 meaningful platform options for B2B attribution in 2026 — but the right choice depends entirely on sales cycle length, conversion volume, and CRM integration requirements

SECTION 6

Custom Attribution Modeling for Teams with Data Warehouses

Teams with established data warehouse infrastructure (Snowflake, BigQuery, Databricks) and in-house data engineering capability have an alternative to third-party attribution platforms: building custom attribution models directly on their own data. Custom attribution has significant advantages. You have full control over the model architecture and can build it specifically for your buying journey rather than adapting a generic model to your context. You own the data and the model — no vendor lock-in, no proprietary black box. The model can incorporate data sources that third-party platforms can't access: product usage data, customer success interactions, community engagement, and any other signal relevant to your specific buying process. The build cost is real but not prohibitive for teams with existing data infrastructure: a data engineer and an analyst can build a functional multi-touch attribution model in Snowflake or BigQuery in 4-6 weeks using open-source frameworks (Shapley value attribution, Markov chain models) that are well-documented and production-proven.

The Shapley value approach — borrowed from cooperative game theory — is the most theoretically defensible attribution model for multi-touch scenarios. It assigns credit to each touchpoint based on its marginal contribution to conversion across all possible orderings of the touch sequence. Shapley value attribution in a data warehouse context requires a clean touchpoint event table with user IDs, channel, timestamp, and conversion outcome — exactly the data you've already built if you've implemented the first-party infrastructure described in the previous section. The practical decision: use a third-party platform if you lack data engineering resources or if the sales cycle length and buying journey complexity make the proof-of-concept window too short to validate a custom build. Build custom if you have the infrastructure, the engineering resources, and data maturity that makes a platform's data collection layer redundant.

- Prerequisite: clean touchpoint event table in your data warehouse with user ID, channel, timestamp, conversion outcome
- Recommended model: Shapley value attribution — theoretically defensible, well-documented, open-source implementations available

- Alternative: Markov chain model — computationally efficient, handles channel removal analysis for budget reallocation scenarios
- Build timeline: 4-6 weeks for a data engineer and analyst to deliver a production-ready Shapley value model
- Validation step: compare custom model outputs against a pilot third-party platform on the same data — divergences reveal data gaps
- Documentation requirement: every custom model must have a documented methodology, data lineage, and refresh schedule

> ***Shapley value attribution is the most theoretically defensible multi-touch model and is implementable in any SQL-capable data warehouse — it's the right custom build target for teams with sufficient data infrastructure.***

4-6 weeks typical build time for a production-ready Shapley value attribution model in Snowflake or BigQuery with a data engineer and analyst

SECTION 7

The Budget Reallocation Playbook: What Typically Changes When Attribution Is Right

When B2B teams migrate from last-click to data-driven attribution, predictable patterns emerge in what gets credited differently. Understanding these patterns in advance helps anticipate stakeholder reactions and prepare the reallocation strategy. Pattern one: paid search brand terms get less credit. Last-click over-rewards brand search because it appears last in many journeys. With accurate attribution, brand search typically sees 20-40% credit reduction, though it rarely loses all credit because it does provide real incrementality. Pattern two: content and SEO get more credit. Long-form content that ranks for research-phase terms often appears early in long buying journeys and gets zero last-click credit despite driving initial awareness. Data-driven models typically allocate 15-25% of total credit to organic content for B2B teams with mature content programs. Pattern three: display and awareness channels get more credit. Particularly for B2B with account-based awareness campaigns, display touchpoints that precede intent-stage activity show measurable lift in data-driven models that last-click entirely ignores.

Pattern four: some retargeting overperformance gets corrected. Retargeting captures credit for buyers who would have converted regardless — in data-driven models, retargeting channels often see 15-30% credit reduction when incremental lift analysis removes the conversion that would have happened without the retargeting touch. The reallocation playbook: lead with the data, not the conclusion. Present the attribution comparison (old model vs. new model, channel by channel) before presenting budget implications. Frame it as the model getting more accurate, not as a channel winning or losing. Give channels with reduced credit 60-90 days to demonstrate whether their budget can produce its credited results at a reduced spend level before making major cuts. Reallocate incrementally — 10-15% per quarter — to allow measurement of the impact before committing to larger shifts.

- Pattern 1: Brand search loses 20-40% of last-click credit — prepare this channel team before revealing attribution results
- Pattern 2: Organic content and SEO gain 15-25% credit — quantify this to justify content investment increases
- Pattern 3: Awareness display gains credit for early-journey influence — use this to defend account-based campaigns
- Pattern 4: Some retargeting credit reduces after incrementality analysis — test incrementality before cutting budget
- Reallocation sequence: present data → 60-90 day observation window → incremental 10-15% quarterly reallocation
- Channel team communication: frame as the model getting more accurate, not as anyone's channel being devalued

> ***Reallocating budget based on new attribution findings in a single planning cycle is a common mistake — it destroys trust before the new model has proven itself and creates political resistance that outlasts the attribution project.***

60–90 days

recommended observation window after attribution migration before executing budget reallocations — gives new model time to prove itself and builds stakeholder confidence

SECTION 8

Stakeholder Presentation: How to Explain Attribution Changes to Leadership

Attribution changes generate organizational anxiety — channel teams worry about losing budget, finance worries about restatement implications, leadership worries about whether they can trust the new numbers. The stakeholder presentation needs to address each of these concerns directly rather than leading with methodology and expecting the audience to arrive at the business implications themselves. The presentation structure for senior leadership has four components. Component one: the accuracy problem — show why the current model produces inaccurate results using concrete examples. The most effective example is always a customer story: take a recent closed-won deal, pull the full touchpoint history from your CRM and marketing tools, show the full journey, and demonstrate how last-click credited the last touch while ignoring eight months of touchpoints that are obviously influential. This makes the problem tangible in a way that statistical arguments don't. Component two: the methodology — explain what the new model does differently in two or three sentences. Don't go deep on the technical architecture. 'The new model uses machine learning to measure the conversion lift that each touchpoint actually contributed, rather than assigning all credit to the last touch before the sale' is sufficient for most leadership audiences.

Component three: the comparison — show the old model vs. the new model channel by channel. Include both the credit changes and the budget implication of acting on the new model. Do not present this as a recommendation — present it as the data. Component four: the confidence section — address the data quality, the validation methodology, and the reallocation plan. Leadership needs to know that this isn't a single analyst's opinion but a validated model, that you've cross-checked it against other data sources (win/loss interviews, CRM activity), and that the reallocation plan is incremental and reversible. End with a clear recommendation and a proposed decision timeline. Attribution presentations that leave leadership with 'we need to think about this' instead of a decision have failed — the goal is a budget reallocation decision, not a methodology discussion.

- Open with a customer story, not a methodology explanation — show a real deal's touchpoint history to make the problem tangible
- Limit methodology explanation to 2-3 sentences: what the new model measures and why that's more accurate
- Present old vs. new attribution comparison channel by channel before introducing budget implications
- Include the validation section: cross-check against win/loss data, CRM activity, and any incrementality tests
- Present the reallocation plan as incremental and time-bounded: 10-15% quarterly shifts, 60-90 day observation windows
- End with a specific recommendation and proposed decision timeline — attribution presentations need to produce decisions

> *The most persuasive element of an attribution change presentation is a specific customer story showing the journey that the old model ignored — concrete narrative beats statistical argument every time with senior leadership.*

82% of attribution migration approvals happen in the first presentation when the story opens with a specific customer journey example vs. 41% when it opens with methodology

SECTION 9

Ongoing Model Maintenance and Re-Training Cadence

An AI attribution model is not a one-time implementation. The model is trained on historical conversion data — and as buying journeys, channel mixes, and customer behavior evolve, the historical training data becomes less representative of current reality. Without a maintenance cadence, attribution model accuracy degrades over 12-18 months. The re-training cadence depends on your conversion volume and the rate of change in your marketing mix. For teams with 3,000+ monthly conversions and an active channel mix, quarterly re-training is appropriate. For teams with lower volumes or more stable channel mixes, semi-annual re-training is sufficient. Re-training is not just rerunning the model on updated data — it requires a validation step that checks whether the updated model produces coherent results and that the updated weights are directionally consistent with other performance signals (win/loss data, customer interviews, incrementality tests).

Beyond re-training, attribution models require two types of ongoing maintenance. First, tracking integrity monitoring: automated checks that confirm UTM parameters are being captured correctly, that conversion events are firing consistently, that CRM integration is syncing on schedule, and that new channels or campaigns are being tagged in compliance with naming conventions. Tracking integrity degrades silently — a UTM parameter drops off a campaign launch, a CRM sync breaks, and the attribution model starts producing results based on incomplete data without any visible error. Second, methodology review: an annual structured review of whether the model architecture itself still fits the buying journey. As new channels are added, as the sales cycle evolves, or as new first-party data sources become available, the model architecture may need to be updated, not just retrained.

- Set re-training cadence: quarterly for 3,000+ monthly conversions, semi-annual for lower volumes
- Every re-training requires a validation step: compare updated weights against win/loss data and incrementality tests
- Build automated tracking integrity monitoring: daily checks on UTM capture rate, conversion event firing, CRM sync status
- Alert threshold: if UTM capture rate drops below 90% or conversion event count drops more than 15% week-over-week, trigger review
- Annual methodology review: does the model architecture still fit the buying journey, or do structural changes require a rebuild?
- Document every re-training event with input data range, weight changes, and validation results — this history is critical for debugging

> ***Tracking integrity monitoring is the most important ongoing maintenance investment — silent tracking degradation produces model drift faster than buying behavior changes do, and it's entirely preventable with basic automated monitoring.***

12–18 months

typical timeframe for an unmonitored AI attribution model to degrade significantly — re-training cadence and tracking monitoring prevent this

AI Attribution Modeling Implementation Checklist

Phase 1 — Foundation

- [ ] Audit current attribution model: what model type, what data sources, and what channels are tracked
- [ ] Pull last-click attribution vs. win/loss data comparison to quantify the accuracy gap
- [ ] Deploy consistent UTM naming convention across all paid and owned channels

- [ ] Implement server-side event tracking for macro-conversions
- [ ] Audit identity resolution capability and document known tracking gaps
- [ ] Configure bidirectional CRM integration with your analytics platform
- [ ] Deploy consent management platform and validate compliance with applicable privacy laws

Phase 2 — Launch

- [ ] Select attribution platform or custom build approach based on conversion volume and data maturity
- [ ] Pilot selected platform with 60-90 days of data before committing to full migration
- [ ] Run old model vs. new model comparison and document channel credit changes
- [ ] Cross-validate new model against customer journey interviews and CRM activity data
- [ ] Build the stakeholder presentation with customer story example and channel comparison
- [ ] Present to leadership with reallocation plan: 10-15% quarterly incremental shifts

Phase 3 — Optimize

- [ ] Set re-training cadence calendar: quarterly or semi-annual based on conversion volume
- [ ] Build automated tracking integrity monitoring with alert thresholds
- [ ] Execute first budget reallocation after 60-90 day observation window
- [ ] Schedule annual methodology review for model architecture assessment
- [ ] Document every re-training event with input range, weight changes, and validation results

NetWebMedia

Stop Funding the Wrong Channels Because Your Model Says So

NetWebMedia helps B2B marketing teams diagnose attribution model accuracy, build the first-party data infrastructure that makes AI attribution work, and select or build the attribution solution that fits their conversion volume and buying journey complexity. We've run attribution migrations across Northbeam, Rockerbox, and custom Snowflake implementations — and we've helped marketing teams present the findings to finance and leadership in a way that produces budget decisions rather than more methodology debates.

- AI Marketing Automation
- AEO & AI-First SEO
- Autonomous AI Agents
- Paid Media + AI Creative
- CRM + AI Workflows

netwebmedia.com/contact