

A FIRST-PARTY INTELLIGENCE INFRASTRUCTURE WHITE PAPER

The State of First-Party Intelligence

Measurement got sharper.
Ownership got rarer.

Why growth-stage companies are making their biggest
decisions on data they don't own.

THE THESIS

Every layer of the modern growth stack has quietly moved from reporting what happened to estimating what probably happened. The estimates keep getting more sophisticated. What almost no growth-stage company has is a reconciled, finance-grade record of its own that the estimates can be checked against. That gap is one we'd seen in client data for years before we had a name for it: the trust deficit. Closing it is not a better dashboard or a smarter attribution model. It's a category of infrastructure that most companies don't yet know they're missing: First-Party Intelligence Infrastructure.

SECTION 1

The number problem in the corner office

Ask a founder what their customer acquisition cost was last quarter and you'll often get a pause, then a question back: "By whose number?"

Marketing has one figure. Finance has another. The ad platforms each report a third, a fourth, a fifth. None of them agree, and the gaps aren't rounding errors. They're large enough to change a hiring decision, a budget reallocation, or a board narrative. The company is spending real money against a version of reality it can't fully stand behind.

We call this the trust deficit. It's the moment a leadership team realizes that the data running their most expensive decisions is fragmented, un-auditable, and rented from vendors whose incentives don't match their own. It rarely shows up as a crisis. It shows up as a slow erosion of confidence, a sense that the dashboards are theater and the real answer is somewhere none of them can see.

This paper is about why that gap exists, why it's getting wider rather than narrower, and what the durable fix actually is.

SECTION 2

What actually changed

For most of the last decade, the growth stack was built on a reasonable assumption: the platforms could see what happened. A click, a view, a conversion, all observed, all reported back with enough fidelity to trust.

That assumption broke, and it broke quietly. Browser privacy changes, third-party cookie restrictions across most browsers, mobile tracking limits, and consent frameworks removed a large

share of the raw signal the platforms used to observe directly. The platforms didn't respond by reporting less. They responded by modeling more.

This is the shift almost nobody named clearly. The industry moved from measurement to estimation, and kept using the language of measurement. Under privacy consent frameworks, a large share of web traffic now goes unconsented and is reconstructed by modeling rather than observed directly. In privacy-regulated markets the gap is stark: one study of more than 1.2 million users found 68.9% withheld consent, leaving the platforms to model the rest.¹ By Google's own documentation, its conversion modeling uses observed conversions to predict the ones it cannot see, then folds the modeled results into your reported conversions at the same granularity as the observed ones, so the estimated and the counted land in the same column, indistinguishable.²

None of that is fraud. It's the rational response to signal loss, and in fairness the modeling got good. Geo holdouts, media mix models, and validated incrementality are genuinely sharper than the last-click reporting they replaced. The problem isn't accuracy. A model can be technically excellent and still hand you a different answer each quarter, because it's recalibrating against a baseline that keeps moving. A number that re-checks itself forever is a defensible estimate. It is not a settled record. Most growth-stage companies have never owned the settled record the estimates are supposed to approximate.

SECTION 3

Why the current category language fails

When a leadership team goes looking for a fix, the market offers five buckets. Each solves a real slice of the problem. None of them closes the trust deficit. Here's where each one stops short.

Attribution and analytics overlays

The well-known dashboard tools sit on top of platform data and apply their own modeling logic to it. They're useful for a media buyer running day-to-day decisions. But the number they produce is their interpretation of your data, living in their environment, and it changes when their model changes. You're renting their lens on your data, and you never own the record underneath it.

Customer data platforms

These route real-time events between systems well. But building a permanent, structured, auditable account of what your business actually did is a different job. A CDP moves the water. It doesn't build the reservoir.

Signal routers and CAPI tools

The privacy-safe plumbing exists to feed the platforms' targeting engines. Their job is to improve what the algorithm learns, not to give your CFO a ledger. Confusing the two is one of the most

common and expensive mistakes in the stack.

Warehouse-native BI tools

These got the closest, and they deserve credit for teaching this exact buyer that owning your data beats renting a dashboard. That's tailwind, not competition. The limit is in the fine print: the "warehouse" is usually still their managed product, on their platform, structured on their terms. If you cancel, you may get an export. That's Tenant Portability, the vendor returning your things when you move out. Ownership means you never had a landlord, and you control the schema, the reconciliation logic, and what you're allowed to query.

Causal and incrementality measurement

This is the sharpest of the five, and the one growth teams should take most seriously. Done well, it answers a question the others can't: did this channel actually cause revenue. But it's still a modeled estimate of lift, and by the method's own logic that estimate has to be recalibrated on an ongoing basis, because a model trained on last year's behavior no longer describes this year's. In practice, teams that run these models treat quarterly recalibration as a working minimum. A perishable estimate built on inputs you don't own is a different object from a reconciled record that sits still. You want both. Most companies have only the first.

The common blind spot across all five: they improve the quality of the guess. Not one of them gives the company a finance-grade record of its own that the guesses can be reconciled against. That missing layer is not a product feature inside any of these categories. It's a category of its own.

SECTION 4

The category: First-Party Intelligence Infrastructure

First-Party Intelligence Infrastructure is a permanent, client-owned data layer that unifies performance, customer, and product signals into a single, queryable, finance-grade source of truth, living inside the company's own cloud environment.

Three properties define it, and all three have to hold at once:

Owned

The warehouse and the identity layer sit in the client's own cloud account. If the vendor relationship ends tomorrow, the asset stays, structure intact, not as an export but as the thing itself.

Finance-grade

The data is engineered to the standard of rigor, transparency, and auditability a company already applies to its financial statements. CFO-ready, reconciled, and traceable back to source. The number a marketing tool reports and the number that hits the bank account can be put side by side and explained when they differ.

Foundational

It sits underneath the rest of the stack rather than beside it. Attribution tools, BI dashboards, incrementality models, and AI systems all become things you run on top of the foundation. Swap any of them out and the record underneath doesn't move.

What the category is not, so it stays sharp: it is not another dashboard, not a custom-report service, and not an analytics overlay competing with the tools above on modeling sophistication. It doesn't try to win the argument over whose attribution number is most accurate. It removes the need to have that argument by giving every tool the same ground truth to sit on.

Ownership of intelligence beats the convenience of point solutions.

SECTION 5

Why now

Four forces are converging, and any one of them would be reason enough. Together they make the timing hard to argue with.

Signal loss is structural, not cyclical

The privacy and consent changes that pushed the industry from observation to modeling are not reversing. The share of modeled data in a typical stack goes up from here, not down. The case for owning your own first-party record strengthens every quarter the platforms model more.

AI readiness is now a board-level concern, and it runs on data structure

A company can't deploy predictive or generative models against its own business if that business lives in fragmented, black-box SaaS tools. Structured, owned, first-party data is the precondition. In a 2025 survey of more than 1,000 finance leaders, 46 percent named data accuracy concerns as a

critical obstacle to AI adoption in finance,³ which means the buyer for this is no longer only the growth team.

Acquisition costs keep rising, which raises the cost of being wrong

When CAC is low, a fragmented data picture is survivable. When every incremental customer is more expensive than the last, a double-digit disagreement between marketing's number and finance's number is the difference between scaling a channel and quietly bleeding on it.

The market has already been taught the vocabulary

Warehouse-native and causal-measurement vendors have spent real money educating this exact ICP that owning your data matters and that modeled numbers deserve scrutiny. That education is done. What's left unaddressed is the gap between "own an export of your data" and "own the reconciled foundation itself." This paper is written into that gap.

SECTION 6

What it changes for the buyer

Here is the kind of thing an owned record catches, drawn from a pattern we have watched play out in more than one account. A company scales a channel that genuinely works, and a second platform it never touched starts reporting worse numbers in the very same window. Nothing changed in the second account. What happened is that the first channel drove new visitors to the site, and the second platform's automated bidding quietly reclassified those visitors as its own engagement the moment they searched anything near the brand, then spent budget re-serving them and reported the result as its own discovery. On that platform's own dashboard, the reclassifying channel still looks busy. Its reported activity climbs while the real new-customer contribution it was supposed to be creating falls.

A reconciled, owned record catches this because it sees both platforms against one shared definition of a new customer, instead of taking each platform's word for its own performance. A rented dashboard cannot, because it sees only its own side and has no reason to flag the day its number stops describing reality. And here is the part that makes this infrastructure rather than a one-time audit: once the reclassification was blocked, the platform's reporting did not correct itself. It found a different lever to keep the same story going. The distortion moved, and the record had to stay on to catch where it went.

The pattern above is not a one-off, and it is not rare. Across every account we have audited, spanning Meta, Google, TikTok, and third-party attribution tools, the number the platform reported and the number verified against first-party orders materially diverged. The direction was not consistent, and that matters more than the size.

EVERY ACCOUNT WE HAVE AUDITED — CLAIMED VS. VERIFIED			
ACCOUNT (ANONYMIZED)	PLATFORM	WHAT THE PLATFORM SAID	THE VERIFIED RECORD
DTC quiz-funnel brand	Meta	1,996 conversions	1,406 real orders — 42% overclaim
Subscription pet brand	Google PMax	~\$25 per new customer	\$125 verified — 5x misread
Multi-region retail brand	Google	~2,267 orders implied	2,684 verified — ~18% undercount
DTC hammock brand	Attribution tool	Best channel down	Under-credited by 156%
High-ticket home improvement	Google Demand Gen	Revenue up ~120%, untouched	Re-billed demand another channel created

Overclaim, undercount, misclassification. The divergence runs in every direction, which is exactly why a sharper model cannot fix it. The only constant across the audited book is this: the platform's number and the verified record disagreed, and only one of the two could be traced to source.

The shift from renting visibility to owning infrastructure changes four concrete things.

Decision quality

When marketing, finance, and the platforms all reconcile against one owned record, the standing argument about whose number is right mostly dissolves. Budget calls get made against a figure the whole leadership team can stand behind, not the most flattering dashboard in the room.

Risk

A finance-grade, auditable record is defensible to a board, an auditor, or an acquirer in a way a vendor dashboard never is. The company stops carrying the risk of a critical decision resting on a number nobody can trace.

Defensibility

In a market where acquisition is a commodity and the platforms are increasingly closed, a company's own unified first-party data is one of the few durable advantages left. When you own the data, competitors can't rent the same edge.

Optionality

Because the foundation is separate from the tools on top, the company can adopt the next attribution model, BI layer, or AI system without re-plumbing everything underneath. The record is stable. The tooling stays swappable.

SECTION 7

Where AxiaOS sits

Everything above is true whether or not a company ever works with us. The category is bigger than any one vendor, and a leadership team could, in principle, build First-Party Intelligence Infrastructure themselves. Some have. The independent operators who've spent tens of millions on paid media and eventually hand-rolled their own warehouse to get a straight answer are the clearest proof the category is real: they arrived at the same architecture from first principles because nothing off the shelf gave it to them.

AxiaOS exists because that build is expensive, slow, and easy to get wrong, and because a productized version of it protects the parts most teams underestimate: the reconciliation logic, the identity resolution, the cost controls, and the ongoing operation. We deploy a bespoke, finance-grade foundation in the client's own cloud, unify their core data streams, and deliver a diagnostic that starts from their P&L rather than their ROAS. The client owns the warehouse and the data. We own and operate everything that keeps it alive. Cancel the relationship and they keep the asset, structure and all, in their own account.

We are not trying to be a better dashboard, and we don't sell custom reports as a product. Those are the things that pull a company back into the rented model. The point of the foundation is that the client can run whatever dashboard, attribution model, or AI system they want on top of a record that's genuinely theirs. The same holds for an agency partner: they run their own tools and campaigns on top of the same owned record, which gives them cleaner inputs to work from and leaves the brand holding the asset no matter who executes the media.

Questions to ask any vendor, including us

Before you sign with anyone in this space, five questions separate an owned foundation from a rented view:

- 1 If we cancel tomorrow, what exactly do we keep, and is it the working asset or an export we have to rebuild somewhere else? Get the answer in writing.
- 2 Whose cloud account does the data live in, and who controls the schema and what we're allowed to query?
- 3 Can every number you report be traced back to a source record we can inspect, or does it arrive already modeled?
- 4 When your attribution logic or model changes, does our historical record change with it, or does it stay fixed?
- 5 Who reconciles the platform's reported number against our actual orders, and do we get to see that reconciliation?

A vendor confident in what they're selling will answer all five without flinching. The answers you get will tell you which category the product is really in.

SECTION 8

The honest counterarguments

Framing a category this way invites pushback. Here are the strongest objections and where each one actually lands.

"This is just an argument against measurement tools."

No, and the distinction matters. Geo holdouts, MMM, and good attribution are the right instruments for the question "did this cause revenue." The argument is that a recalibrating estimate and a reconciled record are different objects, and the estimate runs better on top of an owned foundation than instead of one. Own the record, then measure all you want on top of it.

"We already own our data, it's in our warehouse."

Worth checking whose warehouse, on whose platform, structured by whom. If canceling the vendor turns your "owned data" into an export you have to re-model somewhere else, you had portability, not ownership. The two get marketed with the same word and they are not the same thing.

"We can't justify the cost versus another tool subscription."

Fair, if you frame it as a tool. It isn't one. A tool subscription buys a rented view that resets when you cancel. This builds a permanent asset that survives every tool decision you make after it. The right comparison isn't against a dashboard's monthly fee, it's against the cost and risk of the internal build that gets you the same owned foundation, which is materially higher.

"Isn't this overkill for our stage?"

Below a certain spend level, honestly, yes. That's why the category has a floor. For a company spending \$50,000 a month and up on paid acquisition through fragmented tools and making real decisions on the output, the fragmentation is already costing more than the foundation would. Above that line, "we'll fix the data later" usually means "we'll keep paying for the gap indefinitely."

"The platforms will just fix their own reporting."

They won't, and the reason is structural rather than a matter of effort. A platform grading its own performance has no incentive to flag the day its reported number stops matching reality. A record you own has no such conflict. That gap closes only when the owner changes, which is not something a vendor can sell you.

SECTION 9

The question to take into your next budget meeting

The next time your leadership team debates a budget, notice how much of the argument is really an argument about whose number to believe. That argument is the trust deficit, in real time, costing you a decision.

Sharper measurement won't end it, because sharper measurement is still someone else's estimate of your business. The end of it looks like a reconciled, finance-grade record you own outright, with every tool you like running on top. That's the category. The companies that build the foundation now will spend the next few years making decisions from ground truth while their competitors are still arguing about the dashboard.

If the platform's number and your bank account disagreed tomorrow, do you own the record that could tell you which one is wrong?

Sources

1. Advance Metrics AG, "Cookie Behaviour Study (5 Years After GDPR)," 2023. More than 1.2 million users across international B2B sites; 68.9% withheld cookie consent.
2. Google Ads Help, "How conversion modeling works." Modeled results are included in total reported conversions at the same granularity as observed conversions.
3. Kyriba 2025 CFO survey of more than 1,000 finance leaders: 46% cite data accuracy concerns as a critical challenge to AI adoption in finance. Corroborating: McKinsey (November 2025), roughly one-third of organizations reporting AI issues cite inaccuracy.

AxiaOS

First-Party Intelligence Infrastructure. A permanent, finance-grade data foundation your company owns outright, in your own cloud.