

August 2026

Does your analytics platform still fit?

Google and Adobe are changing direction.
Does their roadmap still match your needs?

PREPARED FOR

Data leaders, digital strategists, and senior marketers responsible for the analytics stack of large organisations.

AUTHORED BY

Stitchd - digital analytics consultancy



Executive summary

Google Analytics and Adobe Analytics built the vocabulary of digital measurement for two decades. That vocabulary no longer describes the world it's used in. Cookies are disappearing, data now lives across a dozen systems instead of one, AI is rewriting how people find and interact with brands, and both platforms have slowed the pace at which they solve any of it.

For many organisations, switching is exactly the right call. Getting there well starts with naming the specific limitations your organisation runs into with GA or Adobe today, and then testing whether a given alternative actually resolves them. The market has split into categories that each solve part of the problem, which is exactly why that test matters: the right alternative depends entirely on which limitations you're trying to fix.

How to use this ebook



If you lead the team

You don't need the technical detail in Chapter 1 to use this paper. Read the executive summary and Section 3 (the requirements checklist), then hand this to whoever owns your analytics stack.

Share this framework with your team, and ask them to score your current platform against Section 3.



If you own the stack

Chapters 1 through 5 are for you: the specific technical case against GA4 and Adobe Analytics, and the process for running an evaluation.

Use Section 3 as your scorecard and Section 5 as your project plan.

SECTION 1

1 Why the ground has shifted

Two platforms carry most of this argument, and the case differs in the details even though the shape looks similar. It's worth being specific about each, because the fix your organisation needs - and how urgently - depends on which one you're running, and why it's stalled.

1.1 Google Analytics 4 is legacy

Identity. GA4's measurement model still depends on client-side, cookie-based tracking, and that layer degrades a little more every year. Safari's ITP caps script-set cookies at seven days, and often one, for domains it classifies as trackers¹; Firefox and Brave block third-party cookies by default. Plausible Analytics has documented as much as 58% of visits going unmeasured by GA on developer-heavy audiences like Hacker News² - an extreme case, but directionally real for any site with a technical or privacy-aware visitor base. Regulators have reinforced the point directly: France's CNIL ruled stock GA implementations unlawful under GDPR in 2022, a decision several other national authorities have echoed³.

Data ownership & modeling. GA4's schema is built around GA4's own reporting model - event-scoped parameters, nested repeated fields in the BigQuery export. The export itself is a genuine improvement over Universal Analytics, but turning it into clean, dimensional tables a warehouse can use is still real engineering work, not a plug-and-play feed.

Roadmap investment. GA4 absorbed years of Google's analytics roadmap to ship, and meaningful investment in the core reporting and analysis experience has visibly slowed since. Most of what continues to ship pulls in one direction: tighter integration with Google's advertising stack - audience signals for Google Ads, conversion modelling, Performance Max feeds - rather than the reporting and analysis capabilities an independent analytics function actually needs. GA4's roadmap increasingly serves Google's media business, not yours.

AI-readiness. In May 2026, Google added an "AI Assistant" channel to GA4's Default Channel Group, automatically classifying sessions referred by ChatGPT, Gemini and Claude⁴. It's a genuinely useful, zero-configuration fix for a real blind spot - and a narrow one: it only catches sessions with a matching referrer, isn't retroactive, and says nothing about whether GA4's underlying event data can leave the platform in a form your own AI or agentic systems can use.

1.2 Adoby Analytics is legacy

Identity. Adobe Analytics' classic visitor-ID model is less exposed to third-party cookie deprecation than GA's - it was built first-party from the start. What it isn't shielded from is Safari's ITP capping first-party cookie lifespan to as little as a single day¹, or the compliance overhead of proving, region by region, exactly how consent and residency are enforced.

¹[Tracking Prevention in WebKit](#) - WebKit ; ²[58% of Hacker News, Reddit and tech-savvy audiences block Google Analytics](#) - Plausible Analytics ; ³[CNIL is latest authority to rule Google Analytics violates GDPR](#) - IAPP ; ⁴[Google Analytics adds AI assistant as default channel group](#) - Search Engine Journal

Data ownership & modeling. This is where Adobe's legacy shows most clearly. Adobe Analytics was built around a fixed number of generic slots - evars, props, events - that analysts map to business meaning during implementation. Getting usable data out means working with a Data Feed export: hit-level rows, cryptic generic column headers, and a separate set of lookup and manifest files just to decode what evar12 or prop4 actually means. Anyone who has modelled a Data Feed knows the gap between “the export exists” and “the export is usable” firsthand. Modern analytics tools increasingly expect the opposite: an event-stream flexible enough to shape at the warehouse layer, not a fixed slot system reverse-engineered after the fact.

Roadmap investment. Unlike GA, Adobe has a clearly designated successor. Customer Journey Analytics is where Adobe's product investment is going - a modern data model, cross-channel ingestion, and closer warehouse integration than classic Adobe Analytics ever offered. Classic Adobe Analytics is, by Adobe's own roadmap signals, the platform being migrated away from. Worth being clear-eyed about what that migration actually involves: moving to CJA is a new implementation, not a licence upgrade - a new data model, new tagging, and a new taxonomy to build, on the same order of effort as adopting an outside platform.

AI-readiness. Classic Adobe Analytics carries Sensei-powered anomaly detection and contribution analysis - genuinely useful, and several years old now. The newer generative AI capabilities - natural-language ad-hoc questions, AI-generated narrative summaries - are landing in Customer Journey Analytics specifically, not classic Adobe Analytics⁵. Adobe is telling you, through its own AI roadmap, which of its two products it expects you to be using in three years.

DIFFERENT PLATFORM, SAME PATTERN

GA4 is a product receiving proportionally less investment each year. Adobe Analytics is a product Adobe itself is walking away from. Whichever one your organisation runs, the requirements in Section 3 are how you test whether it still fits or whether it's time to look elsewhere.

⁵[Adobe unveils AI enhancements in customer journey analytics](#) - Adobe blog

SECTION 2

2 From tool selection to capability requirements

If you're unsure whether Google Analytics or Adobe Analytics still holds up for your organisation, the productive place to start is a requirements exercise, not a vendor shortlist.

The question to answer first is: what does any analytics capability need to guarantee for us, regardless of which vendor provides it? That question is specific and answerable. Starting instead from “which tool should replace it” invites a bake-off decided by demos and sales relationships rather than architecture.

Here's the vision we hold at Stitchd: the warehouse will keep taking on a larger role in the modern data architecture. It won't replace digital analytics platforms - the two need each other, and a good integration between them is what actually matters. A digital analytics solution's job is to provide reliable tracking SDKs and a reporting UI that business users - marketers, product owners, UX designers - can use day to day without needing a data team's help. The warehouse's job is to be the shared, governed home where that data lives alongside everything else the organisation collects, ready to be joined, modelled, and reused. What decides whether a given platform fits that picture is whether data moves freely between the two, without a six-month integration project - that's the requirement Section 3 tests for.

Section 3 turns this into six categories of specific, checkable requirements - the questions worth putting to any vendor before the first demo, let alone a signature.

SECTION 3

3 The requirements framework

Six categories, each with requirements specific and checkable enough to raise directly in a vendor conversation.

Treat these six as a layer on top of the baseline capabilities every serious analytics platform already provides - reliable collection, standard reporting, access control. They are the categories most likely to decide whether your architecture still holds up in five years, gathered here as a starting point for your own evaluation, not a substitute for it.

3.1 Data ownership & portability

- ☐ We can export 100% of raw, event-level data - not a sampled or pre-aggregated extract.
- ☐ We own the schema, or can extend it, rather than working inside a fixed vendor taxonomy.
- ☐ We can leave the platform and take our historical data with us in a usable form.
- ☐ Exported data requires no proprietary tooling to read, query or reprocess.

In practice: An export function on a feature list is not the same guarantee as usable portability. Plenty of platforms can technically produce a file; the requirement is that the file needs no proprietary tooling and no costly reprocessing project before your own team or warehouse can use it.

3.2 Privacy & compliance by design

- ☐ The platform can operate at full functionality without third-party cookies.
- ☐ Consent state is enforced at the point of collection, not patched on afterwards.
- ☐ Data residency (EU or otherwise) is a configuration choice, not a workaround.
- ☐ We can demonstrate, on request, exactly what data was collected, filtered, or dropped - and why.

In practice: Being able to prove data handling, not just describe it, is what regulators and internal legal teams actually test for. A vendor that can state its consent logic and residency setup in one clear sentence has usually built it in; one that needs a workshop to explain it usually hasn't.

3.3 Warehouse & ecosystem integration

- ☐ Data flows into our warehouse continuously, not via a nightly batch job or manual export.
- ☐ The platform can incorporate data imported from our warehouse - CRM segments, product attributes, transactional history - into its own reporting, without a separate integration project.
- ☐ Adding a new data source doesn't require re-architecting the whole pipeline.
- ☐ The vendor's roadmap treats continuous, two-way data flow with the warehouse as a current capability, not a "coming soon."

In practice: This is the requirement that most separates modern comprehensive platforms from legacy ones: whether data flows to and from the warehouse continuously and as a matter of course, rather than through a periodic export somebody has to remember to run.

3.4 AI-readiness

- ☐ Traffic from AI assistants and agentic sessions is identified and attributed as its own channel, distinct from direct and unattributed traffic.
- ☐ AI-assisted analysis - natural-language queries, automated insight, anomaly detection - is available to us today, in production, at our contracted tier.
- ☐ The data the platform produces is clean and structured enough to feed our own AI, personalisation or agentic workflows without a separate modelling project.
- ☐ We have a concrete plan for measuring agentic and AI-referred traffic within the next 12 months, using this platform's own data.

In practice: Ask any vendor to show, concretely, how their platform separates a ChatGPT-referred session from a bot session and from a human one - live, in the product, not on a roadmap slide - and whether you can take that classification logic and the underlying data with you. The quality of that answer says more than any feature list.

3.5 Cost transparency & scalability

- ☐ Pricing is tied to a dimension we actually control (events, tracked users, warehouse compute) - not a black-box “tier.”
- ☐ We know, today, what our cost looks like if usage doubles.
- ☐ There is no material difference between the “evaluation” price and the production price at our real volume.
- ☐ We understand exactly what happens to our access to historical data if we downgrade or leave.

In practice: A pricing model that looks cheap at pilot volume and expensive at production volume isn't a pricing problem to negotiate later. It's a signal about how the vendor expects to make money from you.

3.6 Usability & team fit

- ☐ The primary users of the platform - marketer, analyst, engineer, or all three - can each get a real answer without waiting on someone else.
- ☐ The learning curve fits the team we actually have today, not the team we hope to hire next year.
- ☐ The questions that come up weekly can be answered through self-service reporting, without writing SQL.
- ☐ A new team member can run our standard reports independently within their first week, without vendor-led training.

SECTION 4

4 What these requirements look like by organisation type

The weighting of these six categories is never even. Here is how it tends to break down in practice - deliberately rough archetypes, useful as a starting point for your own weighting.

LARGE E-COMMERCE OR MULTI-BRAND RETAILER

Dominant requirements: 3.1 Data ownership, 3.3 Warehouse integration, 3.5 Cost transparency

Good looks like: a platform that reconciles online and offline conversion inside the warehouse, with a cost model that doesn't punish a Black Friday traffic spike.

REGULATED ENTERPRISE (FINANCE, HEALTHCARE, PUBLIC SECTOR)

Dominant requirement: 3.2 Privacy & compliance

Good looks like: the compliance team gets a straight answer, in writing, about data residency and consent enforcement before procurement even opens a vendor conversation.

MEDIA, PUBLISHING, SUBSCRIPTION CONTENT

Dominant requirements: 3.4 AI-readiness, 3.6 Usability

Good looks like: editorial and subscription teams self-serve engagement and paywall-funnel questions daily, while the platform correctly separates AI-referred readers from organic ones.

DIGITAL-FIRST PRODUCT COMPANY (SAAS, MARKETPLACE, APP)

Dominant requirements: 3.1 Data ownership, 3.3 Warehouse integration

Good looks like: product and marketing analytics share one event taxonomy instead of maintaining two conflicting versions of "what is a user".

DATA-ENGINEERING-LED ORGANISATION BUILDING AI OR PERSONALISATION

Dominant requirements: 3.3 Warehouse integration, 3.4 AI readiness

Good looks like: the warehouse is the source of truth by default, and every analytics interface on top of it is replaceable without touching the underlying data model.

MID-SIZED ORGANISATION, MODEST ANALYTICS AMBITION

Dominant requirements: 3.5 Cost transparency, 3.6 Usability

Good looks like: the team gets traffic, channel and content answers without needing to hire a specialist just to operate the tool.

SECTION 5

5 Running the evaluation

The requirements framework gives you the criteria. Here is the process that turns them into a decision - the same one we run with clients, start to finish.

01. Build the scorecard, then score everyone against it, including your own vendor's "successor."

Turn Section 3's 24 requirements into a simple pass/fail scorecard, weighted by whichever categories matter most for your organisation (Section 4). Score every candidate against that same scorecard - including any migration path your current vendor is steering you toward. If you're on Adobe Analytics, that means scoring Customer Journey Analytics too: moving to CJA is a new implementation - a new data model, new tagging, a new taxonomy - not a licence upgrade, and it should clear the same bar as any outside alternative, not get a pass for being a familiar name.

02. Shortlist to three, and test them on your own data, not a demo.

A vendor demo tests the vendor's best-case scenario. Feed each shortlisted platform a representative slice of your own data and the specific questions your team actually asks weekly, over a fixed two-to-four-week window. Have the people who will use it daily - not just procurement or IT - sign off on whether it produced a real answer.

03. Lock the event taxonomy before you sign, not after.

The taxonomy - what counts as an event, how sessions and users are defined, what counts as a conversion - outlives whichever platform you choose. Assign one owner for it, independent of the vendor, and treat re-doing it later as the single largest hidden cost in most migrations we see.

04. Decide what happens to historical data before go-live, not during it.

Most modern platforms cannot cleanly ingest years of legacy GA or Adobe history. The practical answer is almost always the same: archive the old data in your warehouse, and let the new platform's own history start from day one rather than forcing a clean migration of the old.

05. Run old and new in parallel for one full reporting cycle before cutting over.

A full quarter is the usual minimum - long enough to catch seasonal reporting quirks. The cost of dual collection is far lower than the cost of a missing year-on-year number in a board report.

SECTION 5

06. Budget time for the team, not just the platform.

The platform is rarely the bottleneck three months after go-live; knowing what new questions are now possible to ask is. Plan for a deliberate ramp-up, not just a cutover date.

BEFORE YOU SIGN ANYTHING

Three questions worth asking one more time, out loud, in the room with whoever holds the budget: will this vendor still be here in five years, on terms we'd recognise? Will the specific things we like about it still be there in five years - the Universal Analytics sunset, and Adobe's own move away from classic Analytics, are both proof this isn't hypothetical? And can we still afford it if our usage doubles?

SECTION 6

6 Conclusion

Google Analytics and Adobe Analytics remain mature, familiar products, and for some organisations they are still adequate. The argument in this paper is narrower: the conditions that made them the obvious default no longer describe the world most organisations operate in, and that gap is compounding every quarter it goes unaddressed.

The organisations that come out ahead will be the ones that get specific about requirements before they get attached to a shortlist, and that build an architecture that stays theirs no matter which platform sits on top of it.

STITCHD'S VIEW, PLAINLY STATED

The requirements in Section 3 reflect patterns we see repeatedly across every platform we implement, audit, and migrate. If you'd like a second opinion on how your current stack holds up against them, we'd be glad to talk.

About Stitchd

Stitchd is a digital analytics consultancy. We design, implement, audit, and migrate analytics stacks for data leaders, digital teams, and senior marketers across Europe.

We are platform-agnostic by design: no exclusive partnerships, no resale margin, no incentive to push any one vendor.

info@stitchd.be

AI DISCLOSURE NOTICE

This whitepaper was developed by Stitchd with the assistance of Claude 3.5 Sonnet (Anthropic). AI was utilized to help structure the requirements framework, refine technical phrasing, and optimize readability. All frameworks, methodology, strategic insights, and final conclusions were conceptualized, thoroughly audited, and verified by the Stitchd consulting team, who assume full responsibility for the accuracy and integrity of this content.

© 2026 Stitchd. This whitepaper may be shared internally and externally with attribution.