

LAYER 1 OF 3 · FOR BUSINESSES, AGENCIES & PARTNERS

Data & Intelligence. The Foundation Layer.

Six analytical models. One unified truth. Every morning.

Layer 1 is the intelligence foundation of MIDAS. It unifies every marketing channel, runs analytical models automatically, and surfaces plain-English intelligence every morning — without requiring clients to hire data scientists, manage infrastructure, or trust platforms to report their own performance accurately.

Marketing Intelligence & Decision Automation System

A NetScience Product · midas.nettscience.com · nettscience.com

15+

Channels
Unified

7

Intelligence
Capabilities

12 Wks

To Full
Intelligence

100%

Your Data
Stays Yours

SECTION 01

The Problem Layer 1 Solves

Marketing today generates more data than almost any other business function. A mid-market business running campaigns across Google, Meta, LinkedIn, email, and organic search is producing tens of thousands of data points every day — across platforms that use different attribution windows, different conversion definitions, and different reporting standards.

The result is not an abundance of insight. It is a fragmentation problem. Each platform shows you its own performance. None shows you what is actually happening across all of them simultaneously. And every platform — by design — overstates its own contribution.

The Four Failures of Current Analytics

Failure	What It Costs You
Fragmented data	Every platform speaks a different language. One definition of a conversion in Google Ads. A different one in Meta. A third in your CRM. There is no unified view — only a collection of isolated, incompatible reports.
Attribution overlap	Add up what your platforms claim they drove and you will typically reach 150–200% of your actual conversions. Each platform claims every conversion it touched. The result is systematic overstatement and misallocation of budget.
Reactive discovery	Problems are found the next morning in reports — after budget has already been wasted. Creative fatigue, CPA spikes, tracking failures — all discovered in arrears.
No institutional memory	Every analytical task starts from scratch. Insights from last month's campaigns do not inform this month's decisions. There is no system that accumulates marketing knowledge over time.

The Layer 1 Answer

Layer 1 connects every platform, transforms all data into a single unified model, runs six analytical models on that unified data, and delivers intelligence through dashboards and AI-generated summaries — automatically, every morning, before your team starts work. This is not a better dashboard. It is the intelligence layer that makes all your dashboards honest.

SECTION 02

Seven Intelligence Capabilities

Layer 1 operates through seven integrated capabilities, each designed to address a specific dimension of the fragmented data problem. Together they form a continuous intelligence cycle — collecting, transforming, modelling, reasoning, and surfacing — that runs automatically every day.

01. Unified Data Collection

The intelligence foundation begins with collection. Layer 1 connects to the full range of platforms that constitute a modern marketing operation — paid media (Google Ads, Meta, LinkedIn Ads, TikTok), web analytics (GA4), CRM systems (HubSpot, Salesforce), email platforms, organic search (Search Console), SEO tools, social platforms, and e-commerce platforms (Shopify, WooCommerce). Our orchestration engine (n8n) has 400+ native integrations. Data flows in continuously — some sources hourly, others daily — and is stored exactly as received in the raw data zone of the warehouse, creating a complete and unmodified historical record.

02. Cross-Channel Attribution

Platform-reported attribution is structurally biased. Every platform applies its own attribution model, using its own window, claiming credit for conversions that other platforms also claim. Multi-touch attribution built across channels — rather than within them — is the only way to understand which marketing activities are actually driving your best customers. Layer 1 builds attribution models that hold data from all connected platforms simultaneously. MIDAS-calculated attribution is presented side-by-side with platform-reported attribution, making the discrepancy visible. Most clients discover their actual channel contribution is significantly different from what their platforms report — often revealing that one or two channels are being substantially undervalued.

03. AI Anomaly Detection

Campaign performance is dynamic. CPA spikes, budget pacing issues, creative fatigue, tracking failures, and competitive disruptions all emerge continuously — not just overnight. Detecting them requires continuous monitoring against statistical baselines, not a morning report. Layer 1's anomaly detection runs 24 hours a day, monitoring every metric across every connected channel against rolling statistical baselines. When a significant deviation is detected, an alert fires immediately — with a plain-English explanation of the likely cause, not just a number. For high-severity anomalies, the AI reasoning layer generates a diagnostic summary explaining what changed, when, and why.

SECTION 02 • CONTINUED

Seven Intelligence Capabilities

04. Lead Scoring & Segmentation

Most lead scoring is built on demographic or firmographic data — who a person is, not what they have done. Predictive lead scoring built on actual CRM deal outcomes — which leads converted, at what value, over what timeframe — is substantially more accurate, and substantially more useful for the sales team that acts on it. Layer 1 trains AI models on your actual CRM data. Lead scores are written back to HubSpot or Salesforce in real time, updating as new behavioural signals arrive. Segment membership — which audience tier a lead belongs to — is recalculated continuously. Marketing intelligence becomes sales intelligence automatically, without any manual data transfer.

05. Content Marketing Intelligence

Most analytics systems stop at the ad or the keyword. Layer 1 goes further: which blog posts, social content, and downloadable assets are actually generating qualified leads? Which content formats and topics drive the best outcomes for this specific client? Where are the SEO gaps a competitor is exploiting? The content intelligence module connects content assets to business outcomes through multi-touch attribution across the full customer journey. Pillar balance monitoring tracks actual vs. target content mix over a 90-day rolling window. SEO gap analysis identifies high-value keywords competitors rank for that you do not target. Optimal posting time signals are updated continuously as audience behaviour evolves.

06. Competitive Intelligence

Competitive shifts do not wait for quarterly reviews. Competitor search rankings change week by week. Ad library activity reveals strategic direction. Content publishing frequency and topic patterns signal positioning moves. Significant competitor movements need to be surfaced proactively, not discovered accidentally. Layer 1 monitors competitor search rankings, ad library activity, content publishing patterns, and share-of-voice trends continuously. When a significant competitor movement is detected — a major keyword ranking shift, an unusual spike in ad activity, a new content push into your core topics — it is surfaced in the daily intelligence output with context about what the movement likely signals.

07. Intelligence Dashboards

All intelligence produced by Layer 1 is delivered through Apache Superset dashboards — configured for each client's specific business context and updated automatically every morning. Role-based access controls ensure each person sees the right level of detail: executive KPI summaries for leadership, campaign-level diagnostics for marketing teams, attribution breakdowns for strategists. Dashboards are not the output of a reporting process that someone has to run. They are live views into a warehouse that is continuously updated — when you open your dashboard at 8am, it reflects yesterday's complete data, already modelled and already reasoned over.

SECTION 03

The Technology Stack

Layer 1 is built on six production-grade open-source components, running on a private VPS in a professional data center. No vendor lock-in. No data routing through third-party analytics clouds. Every component is under our operational control, auditable, and replaceable if a better alternative emerges.

Component	Role and Production Pedigree
n8n — Orchestration	Workflow automation engine with 400+ native platform integrations. Manages all data collection schedules, API calls, transformation triggers, and delivery routing. Deployed at Airbnb and Delivery Hero at scale.
KNIME — Analytics	Visual data science platform. Runs cross-channel attribution, anomaly detection, lead scoring, LTV prediction, and content performance models. Used by 1M+ data scientists globally.
PostgreSQL — Warehouse	Central data warehouse with four logical schemas: raw data, modelled outputs, AI context, and operational logs. Row-level security for client isolation. pgvector extension for semantic AI memory.
Ollama — Local AI	Runs open-source LLMs locally on the VPS. Classification, embeddings, anomaly summaries, content tagging — all processed within the platform. Zero data routing to external APIs for high-frequency tasks.
Apache Superset — Dashboards	Client-facing BI dashboards with role-based access. Connects directly to the PostgreSQL warehouse. Used in production at Airbnb, Lyft, and Twitter.
Caddy — Gateway	Automatic HTTPS gateway. Sole external-facing component. Auto-provisions TLS certificates, handles rate limiting, IP access controls, and internal routing. All other components are invisible to the internet.

Data Sovereignty by Architecture — Not by Promise

High-frequency AI tasks (classification, content tagging, anomaly summaries, embeddings) run locally on the VPS using Ollama. Your data never leaves your environment for these tasks — zero external routing. Complex reasoning tasks that route to frontier AI APIs (Claude, GPT-4) do so under enterprise agreements that explicitly prohibit using your data for model training. Minimum data is sent per call. PostgreSQL row-level security enforces client isolation at the database level — not application code.

SECTION 04

What You Receive — Daily Intelligence Outputs

Layer 1 is a managed service. You do not manage infrastructure or run analytical processes. You receive intelligence outputs — every morning, automatically, before your team starts work.

Intelligence Output	What It Tells You
Unified performance dashboards	All connected channels in one view, updated with the previous day's complete data. Paid media, organic, email, content, and CRM — one consistent model across every source.
Cross-channel attribution view	MIDAS-calculated attribution presented alongside platform-reported attribution. See where the discrepancies are, and understand which channels are actually driving your best customers.
Real-time anomaly alerts	AI-generated alerts with plain-English diagnosis, firing as anomalies are detected — not the next morning. CPA spikes, budget pacing issues, creative fatigue, tracking failures.
AI insight summaries	Weekly intelligence digests summarising performance trends, significant shifts, and forward-looking signals — generated by AI, reviewed for accuracy, delivered in plain English.
Lead scores to your CRM	AI-calculated lead scores written directly to HubSpot or Salesforce in real time. No manual export, no data transfer process — scores update as new signals arrive.
Content performance intelligence	Which blog posts, social content, and assets are generating qualified leads. Pillar balance monitoring. SEO gap analysis. Optimal posting time signals. Competitor content tracking.
Competitive intelligence	Competitor search ranking movements, ad library activity, content publishing patterns, and share-of-voice trends — surfaced proactively when significant movements occur.
Attribution vs. platform comparison	Side-by-side comparison showing how your budget allocation would change if based on MIDAS attribution vs. platform-reported attribution. Typically the most commercially impactful output in the first 90 days.

SECTION 05

Onboarding — From Zero to Full Intelligence

Layer 1 onboarding is structured to deliver intelligence value progressively rather than requiring a long setup period before any output is available. First dashboards are live in weeks 3–4. Full channel coverage by week 12.

Phase	What Happens & What You Get
Weeks 1–4 · Foundation	Google Ads, GA4, and Search Console connected. First dashboards live. Paid search and organic intelligence immediately active. Cross-channel attribution begins accumulating data.
Weeks 5–8 · Attribution	Meta Ads and CRM connected. Cross-channel attribution modelling active. Lead scoring operational. Revenue-linked intelligence available for the first time.
Weeks 9–12 · Full Coverage	Email, e-commerce, competitive intelligence, and content performance active. Full channel coverage for most client profiles. All seven capabilities operational.
Months 4–6 · Intelligence Deepens	AI preference models begin accumulating. Semantic content library builds. Attribution baselines mature. Intelligence becomes progressively more calibrated to this client's specific context.
Month 6+ · Layer 2 Ready	Intelligence baselines established. Preference profile accumulating. Ready to activate Layer 2 decision automation — structured AI recommendations for human approval.

Your Time Commitment During Onboarding

API credentials for each platform (read-only access for collection). Typically 15–30 minutes per platform to create a service account or token. 2–3 hours of briefing conversations with the NettScience team: your business objectives, target KPIs, key audience segments, and priority channels. That is the extent of the time commitment required from your side. We handle all integration configuration, model calibration, and dashboard build.

SECTION 06

The Path to Layer 2

Layer 1 is independently valuable — many clients operate on Layer 1 intelligence indefinitely. But Layer 1 is also the foundation that makes Layer 2 possible. The intelligence baselines, the preference profile, and the calibrated models that Layer 1 builds over the first 3–6 months are exactly what Layer 2's recommendation engine is built on.

Layer 1 Builds	Layer 2 Adds
Attribution baselines calibrated to your specific channel mix. Anomaly detection thresholds tuned to your campaign patterns. Lead scoring models trained on your actual CRM outcomes. Content performance history indexed and embedded. A growing understanding of your marketing's normal behaviour.	Structured AI recommendations — situation, diagnosis, action, expected outcome, confidence, expiry — presented for human approval via Slack or Notion. Every recommendation is grounded in the Layer 1 intelligence baselines. The human stays in control; the AI does the diagnostic and formulation work.

One Platform. Three Layers. Each One Building on the Last.

Layer 1 gives you the intelligence. Layer 2 turns intelligence into approved action. Layer 3 closes the loop — executing approved actions autonomously within limits you define. You do not need to commit to all three layers. Start with Layer 1. Progress to Layer 2 when the intelligence is mature and you are ready to accelerate decision-making. Activate Layer 3 when preference calibration through Layer 2 gives you the confidence to set explicit autonomous limits. Each layer is independently valuable. Each one makes the next more powerful.

SECTION 07

For Agency Partners — White-Label Layer 1

Agency partners access the full Layer 1 capability and deploy it for their clients under their own brand. Every dashboard, every alert, every report carries the agency's name. MIDAS is invisible to end clients.

Dimension	How It Works
What your clients receive	Branded intelligence dashboards. Branded anomaly alerts. Branded weekly insight summaries. Branded attribution analysis. Your agency's name on every output.
What your agency receives	A recurring analytics service tier layered on top of your existing billing. A capability that execution-only agencies cannot match. A retention mechanism that deepens with every month of engagement.
What we handle	All platform integrations. All model configuration. All dashboard builds. All data pipeline maintenance. You brief us on the client; we configure the platform.
The commercial model	Per-client monthly wholesale fee (volume-tiered). You set your own client pricing. The margin between your client fee and your wholesale cost is yours — typically 40–55% at scale.

Why Layer 1 Creates Retention

Every month of Layer 1 engagement builds a richer attribution baseline, a more calibrated anomaly detection model, and a more complete content performance library for this specific client. After 12 months, the accumulated intelligence — the baselines, the content history, the calibration to this client's specific marketing patterns — cannot be replicated by a new provider without starting from scratch. This is not contractual lock-in. It is compounding value that makes staying the obvious rational choice.

Next Steps

MIDAS Layer 1 is live. To see it running on real data, book a demo at analytics@nettscience.com. To discuss the agency partner programme, visit us at midas.nettscience.com.

midas.nettscience.com · nettscience.com · +91 98459 05041 · analytics@nettscience.com