MIDAS
BY NETTSCIENCE

FREE TECHNOLOGY GUIDE

The AI Marketing Stack Explained

How AI-Native Marketing Intelligence Actually Works — No Hype

Most marketing AI is a dashboard feature on a SaaS platform.
Real AI-native marketing intelligence is an architecture —
a system of components that work together continuously.
This guide explains what that architecture looks like and why it matters.

6	Local AI	Cloud AI	24/7
Components in the MIDAS stack	Runs inside your infrastructure	Only for complex reasoning tasks	AI running continuously

CUTTING THROUGH THE HYPE

What AI Marketing Intelligence Actually Means

The phrase 'AI-powered marketing' has been applied to almost every category of marketing software in the last two years. Recommendation engines, content generators, chatbots, predictive lead scores in CRMs — all marketed as 'AI'. Most of these are point solutions: AI applied to one narrow task within one platform's ecosystem. They do not talk to each other. They do not share data. They do not reason across your full marketing operation.

Genuine AI marketing intelligence is different in kind, not just degree. It requires an architecture — a system of components that collect, unify, model, and continuously monitor data across every channel simultaneously. The AI is not a feature you turn on. It is embedded throughout the entire data pipeline, performing specific, well-defined tasks at each stage of the intelligence cycle.

AI FEATURE VS AI-NATIVE ARCHITECTURE

DIMENSION	AI AS A PLATFORM FEATURE	AI-NATIVE ARCHITECTURE (MIDAS)
Data scope	One platform's data — bounded by platform walls	*All channels unified — cross-platform by design*
Attribution	Within-platform attribution only — self-reported	*Cross-channel attribution above all platforms*
AI tasks	One task — optimisation, prediction, or generation within that tool	*Multiple specialised AI tasks across the full data pipeline*
Data sovereignty	Data processed on vendor servers — no client control	*Local AI within client infrastructure for most tasks*
Continuity	Runs when triggered or when platform refreshes	*Continuous — monitoring, scoring, classifying 24/7*
Learning	Platform trains its own models — not on your preferences	*Preference learning specific to your business decisions*

THE KEY DISTINCTION

An AI feature makes one tool smarter. An AI-native architecture makes your entire marketing operation smarter — continuously, across every channel, without requiring a human analyst to run the process each time.

THE TECHNOLOGY STACK

Six Components. One Coherent Intelligence System.

The MIDAS stack is built on six proven, production-grade open-source components. Each has a precise, non-overlapping role. None is experimental — each is deployed at enterprise scale globally. Together they form a coherent system that is greater than the sum of its parts. Every component runs inside the client's own infrastructure as a Docker container.

01 **n8n**

ORCHESTRATION ENGINE

Self-hosted workflow automation with 400+ native platform integrations — Google Ads, Meta, LinkedIn, GA4, CRM, email, SEO. Manages data collection, scheduling, AI routing, and approval orchestration. Runs nightly data pulls and triggers all downstream processes.

02 **KNIME**

ANALYTICS ENGINE

Visual data science platform handling ETL transformation, cross-channel attribution, anomaly detection, lead scoring, and lifetime value prediction. Triggered by n8n after each data collection cycle. Used by 1M+ data scientists globally.

03 **PostgreSQL**

UNIFIED DATA WAREHOUSE

Central data warehouse with four logical zones: raw data, modelled data, AI context, and operational logs. pgvector enables semantic AI memory. Row-level security enforces strict client data isolation at the database layer.

04 **Ollama**

LOCAL AI RUNTIME

Open-source LLM runtime serving AI inference locally within the client's VPS. Handles content classification, anomaly scoring, and embedding generation. Zero external data routing. Zero marginal cost per task.

05 **Apache Superset**

DASHBOARDS & VISUALISATION

Open-source BI platform serving client-facing dashboards via secure browser. Connects directly to PostgreSQL — no intermediate data layer. Role-based access control. Originally built at Airbnb.

06 **Caddy**

SECURE GATEWAY

The sole internet-facing component. Auto-provisions TLS certificates, enforces HTTPS, applies rate limiting and IP access controls. All other components sit on a private Docker network — invisible to the public internet.

THE AI LAYER

Local AI and Cloud AI — A Deliberate Architecture

The most important architectural decision in the MIDAS AI layer is the distinction between local AI and cloud AI — and which tasks go where. This is not a cost decision. It is a data sovereignty decision first, a cost decision second, and a quality decision third. The routing is explicit and intentional: every AI task in the system has a documented reason for running locally or in the cloud.

LOCAL AI — Ollama

Content Classification & Tagging

Every ad creative, social post, email, and competitor activity classified by type, intent, tone, and funnel stage. Runs continuously at zero marginal cost. Data never leaves the client's VPS.

Anomaly Detection & Severity Scoring

Campaign performance anomalies classified by type and severity continuously — not in batches. High-volume and high-frequency, unsuitable for cloud API cost structures.

Embedding Generation

Content, recommendations, and performance data encoded into vector representations stored in PostgreSQL (pgvector), enabling semantic search across institutional memory.

CLOUD AI — Claude / GPT-4

Anomaly Explanation & Diagnosis

For highest-severity anomalies, the system assembles context and generates a plain-language diagnosis. Account managers receive a reasoned narrative — not a raw alert. Enterprise agreements guarantee no training on client data.

Strategic Recommendations

Cross-channel attribution analysis, budget reallocation, audience expansion, and content strategy recommendations requiring nuanced multi-factor reasoning.

Preference Learning & Calibration

Approval, rejection, and modification signals update the AI Context Zone, making recommendations progressively more calibrated to each client's decision patterns and priorities over time.

THE DAILY INTELLIGENCE CYCLE

What Happens Every Morning Before 8am

06:00 AM — Data Collection
n8n pulls all platform data via API

›

07:00 AM — Transform & Model
KNIME runs attribution, anomaly detection, scoring

›

07:15 AM — Local AI
Ollama classifies, scores, generates embeddings

›

07:30 AM — Cloud AI
Frontier models explain anomalies and draft recommendations

›

08:00 AM — Dashboards Live
Superset refreshed — intelligence ready for review

The MIDAS daily intelligence cycle — automated, continuous, zero manual effort

DATA SOVEREIGNTY

Why Your Data Never Has to Leave Your Control

Most cloud-based marketing analytics platforms route your data through their own infrastructure for processing and AI inference. This is not a theoretical concern — it is the default architecture of most SaaS analytics tools. MIDAS is designed around the opposite principle: your data stays in your environment, processed by infrastructure you own, operated by NettScience as a managed service.

01

Client-Owned Infrastructure

The entire MIDAS stack runs on a VPS owned and controlled by the client — not on NettScience's shared infrastructure, not on a third-party analytics cloud. NettScience builds, deploys, and operates the system as a managed service. The infrastructure belongs to the client. If the engagement ends, the infrastructure and all data within it remains with the client.

02

Local AI for High-Frequency Tasks

Content classification, anomaly severity scoring, embedding generation — the AI tasks that run continuously and at high volume — execute via Ollama on the client's VPS. No data leaves the environment for these tasks. Zero external routing. This is not a limitation of the system — it is the system working as designed.

03

Enterprise Agreements for Cloud AI Tasks

When complex reasoning tasks route to frontier AI APIs (Claude, GPT-4), this occurs under Anthropic's and OpenAI's enterprise data processing agreements — which explicitly prohibit using client data for model training. Minimum necessary data is sent per call. The AI Context Zone in PostgreSQL ensures the system has client-specific context without sending sensitive data to external APIs.

04

Database-Level Security

PostgreSQL enforces row-level security at the database layer — not in application code. Each client's data is accessible only through queries that include the correct identity filter. No application-layer failure can expose one client's data to another. Every workflow execution, data quality event, and system action is logged to the Operational Zone — a complete, immutable audit trail.

QUICK REFERENCE

The MIDAS Stack at a Glance

COMPONENT	ROLE	KEY PROPERTY
n8n	Orchestration & data collection	*400+ integrations · self-hosted · Docker*
KNIME	Analytics, attribution & ML models	*ETL · attribution · anomaly detection · lead scoring*
PostgreSQL	Unified data warehouse	*4-zone schema · pgvector · row-level security*
Ollama	Local AI inference	*In-perimeter · zero routing · classification & embeddings*
Apache Superset	Dashboards & visualisation	*Direct PostgreSQL connection · role-based access*
Caddy	Secure gateway	*Auto TLS · rate limiting · sole internet-facing component*

Book a Demo

See the stack in action on real anonymised data.
midas.nettscience.com

Architecture Deep Dive

Full technical documentation: stack, security, and infrastructure.

Talk to the Team

Send a mail to analytics@nettscience.com