

LAYER 3 OF 3 · FOR BUSINESSES, AGENCIES & PARTNERS

Channel & Execution. The Loop, Closed.

Autonomous execution, within limits you define.

Layer 3 is the full loop — closed. For action types where preferences are fully calibrated through Layer 2, execution becomes autonomous within limits you define. Five safety checks before every action. Complete audit trail. Every boundary under your control.

Marketing Intelligence & Decision Automation System

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

5

Safety Checks
Per Action

8

Execution
Channels

100%

Auditable
& Reversible

∞

Instant Override
Anytime

SECTION 01

What Layer 3 Does

Layer 3 is not 'remove the human from marketing decisions.' It is 'the human has already established, through months of Layer 2 approvals, what types of actions can execute automatically — and within what parameters.' Autonomy in Layer 3 is earned through calibration, not assumed from day one.

When a specific action type — for example, budget reallocation between audience segments within a defined maximum change percentage — has been approved consistently through Layer 2, the client can activate autonomous execution for that action type. Layer 3 then executes it continuously, within the defined limits, without requiring explicit approval each time.

The human stays in control of every boundary. What can execute autonomously. What cannot. The maximum change magnitude. The channels included. The days and times permitted. The conditions that trigger a pause and return to explicit approval. Every parameter is client-defined and can be modified at any time.

The Three-Stage Path to Layer 3 Autonomy

The path from Layer 2 to Layer 3 is earned through demonstrated preference calibration — not assumed from day one.

Stage	What Happens
Stage 1 · Layer 2 Calibration (Months 1–6)	Every recommendation is reviewed and decided by a human. The AI learns. The preference profile builds. Approval rate increases from ~40% toward 75%+. For specific action types that reach consistent approval, evidence accumulates.
Stage 2 · Limit Definition (Month 6+)	For action types with mature preference calibration (typically 80%+ approval rate, consistent parameters), the client defines explicit autonomy boundaries: maximum budget change %, spend caps, channels included, permitted days and times, override triggers.
Stage 3 · Autonomous Execution (Ongoing)	Layer 3 executes approved action types within defined limits — continuously, without manual approval. Daily digest reports every action taken. Real-time alerts for any action approaching a defined limit. Override available instantly.

SECTION 02

Five Safety Checks — Before Every Autonomous Action

Every autonomous action — regardless of how routine, how small, or how many times the same type of action has been executed before — passes the same five-stage validation sequence before it executes. Safety checks are structural: they cannot be bypassed, disabled, or skipped.

Safety Check	What It Prevents
Check 1 · Spend Cap Enforcement	Checks the current cumulative spend against the client-defined spend cap for this period before every paid media action. If executing the action would bring total spend to within 5% of the cap, the action aborts and triggers a human alert for explicit approval.
Check 2 · Change Magnitude Limits	Verifies that the proposed change is within the client-defined maximum change magnitude for this action type. No single autonomous action can change a budget or bid beyond the defined percentage limit — even if the AI's recommendation called for a larger change.
Check 3 · Pre-Execution State Validation	Confirms that the current state of the campaign or asset matches the state that existed when the recommendation was generated. If a human has made a manual change between recommendation generation and execution, the validation fails, the action aborts, and the recommendation is invalidated.
Check 4 · Idempotency Check	Confirms that this specific action has not already been executed — preventing duplicate execution due to technical retries, system restarts, or timing edge cases. The same action is never applied twice.
Check 5 · Immutable Audit Log	Records the full action context — what was executed, when, in which platform, at which parameters, following which recommendation, with which pre-execution state — in an append-only log that cannot be modified or deleted. Complete forensic history of every autonomous action.

If Any Safety Check Fails

The action aborts immediately. An alert is generated. The action is flagged for explicit human review before it can be re-attempted. No autonomous action ever executes through a failed safety check. There is no override for safety validation — only a human re-approval flow.

SECTION 03

What Layer 3 Executes — Eight Action Channels

Layer 3 executes across the channels and action types where autonomous execution creates the most value — those that are time-sensitive, high-frequency, or benefit from continuous optimisation rather than periodic review.

Action Channel	What Layer 3 Executes
Paid Media Optimisation	Budget adjustments, bid strategy changes, and audience targeting updates across Google Ads, Meta, and LinkedIn — executed in real time as conditions change, within defined limits. No platform logins required.
CRM Lead Scoring	AI lead scores and segment membership written directly to HubSpot or Salesforce the moment a lead enters or a behavioural signal updates. Sales teams always act on current intelligence.
Content Scheduling & Publishing	Approved content briefs generate publishing actions — scheduled at AI-determined optimal times, across the correct channels, for the right audiences. No manual queue management.
Email Sequence Routing	Engagement-based routing of prospects through email sequences — fast-tracking high-engagement leads, re-routing low-engagement ones, pausing sequences for leads showing disengagement signals.
Content Amplification	When a blog post or content asset performs above its baseline in organic search or social, Layer 3 automatically generates a social amplification series adapting its key points — compounding the value of high-performing content.
Automated Client Reporting	AI-generated performance reports covering paid media and content performance — covering the previous period's outcomes against targets, key anomalies, and forward signals. Generated on schedule, reviewed by the account team, delivered to clients.
Intelligent Alert Routing	Signals requiring human awareness — competitive movements, significant trend shifts, attribution anomalies — delivered to the right person through the right channel with the correct urgency level. No alert fatigue from undifferentiated notifications.
Safety Override Escalation	Any autonomous action approaching a defined limit — spend cap, change magnitude, anomalous condition — pauses execution and escalates to explicit human re-approval before proceeding. The safety floor is structural, not optional.

SECTION 04

The Content Intelligence Learning Loop

Every piece of content executed through Layer 3 generates performance data — engagement rates, traffic driven, leads generated, pipeline influenced — that is indexed and stored against the brief parameters that produced it. This data feeds the next brief generation.

Step	What Happens
Brief Generated	Performance history of similar content — same format, same pillar, similar topic — informs the brief. Not a template. A data-driven starting point calibrated to what has actually worked for this client.
Human Approves	Brief reviewed and approved by the account team or client. Final editorial judgment stays with humans. The brief can be modified before approval — modifications feed the preference model.
Content Published	Approved content executed at AI-determined optimal time — the posting window where this client's audience engagement has been highest over the last 30 days.
Performance Measured	Engagement, traffic driven, leads generated, and pipeline influenced tracked against the brief parameters. Which format, which pillar, which topic, which posting time — all recorded.
Loop Closes	Performance data updates the history for this brief type. Next brief for similar content is better informed. The content strategy improves every month without requiring manual analysis.

The Compounding Value of the Loop

In month one, content briefs are based on limited history. In month six, they are based on 150+ pieces of content with measured outcomes. The AI does not just know that case studies perform better than opinion posts for this client. It knows which case study topics, in which format, posted at which time, have driven the most qualified leads — and calibrates every brief accordingly. This level of content intelligence is not available to a new provider starting from scratch. It belongs to the engagement.

SECTION 05

Data Security at the Execution Layer

Layer 3 introduces execution capability — which raises the security requirements beyond intelligence delivery. Every execution action involves a write operation to an external platform. These write operations are subject to additional security controls.

Security Control	How It Works
API credential isolation	Write-access API credentials for each platform are stored in an encrypted credential store, isolated per client, accessible only to the n8n execution context for that client. No credential is accessible to any other client's execution context.
Execution scope limits	Write access is granted only for the specific action types that have been activated for autonomous execution. A credential that can adjust budgets cannot also create or delete campaigns — scope is minimal by design.
Complete execution audit	Every execution action is logged — the pre-execution state, the action taken, the parameters applied, the platform response (success or error), and the timestamp. Logs are append-only and cannot be modified.
Human override available always	Any autonomous action type can be paused instantly — disabling autonomous execution for that action type and returning it to explicit human approval — without affecting other autonomous action types or the intelligence layers.

SECTION 06

For Agency Partners — Layer 3 as a Service

Agency partners who offer Layer 3 to their clients are offering something genuinely differentiated: a service that operates their clients' marketing — within defined limits — continuously, transparently, and with complete accountability. No other managed service offering in the mid-market provides this.

Dimension	How It Works for Agencies
What your clients get	Autonomous execution of trusted action types. Daily digest of every action taken. Real-time alerts for anything approaching a defined limit. Complete audit trail. Override control at all times.
What your agency gets	The deepest client stickiness available. Once autonomous execution is running, switching providers means losing not just intelligence — but operational continuity that has been managing their campaigns. The switching cost is real and growing.
What we handle	All execution infrastructure. All API integrations. All safety check logic. All audit logging. All alert routing. You configure the limits; we run the execution system.
The commercial tier	Layer 3 represents a premium service tier with its own activation fee and monthly rate — incremental revenue per client, on top of Layer 1 and Layer 2 fees.

Next Steps

Layer 3 activates on top of a mature Layer 2 deployment. To discuss readiness criteria and timelines, book a demo at midas.nettscience.com or contact us at analytics@nettscience.com.

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