



nivo

Agentic AI Proof Points

Build campaign from media plan

Map creatives to placements

Test click-throughs and pixels

Create a new QA rule

Hi! What would you like to do today?



Measured results across **seven verticals**

How AI agents are reducing manual work across campaign operations — trafficking, decisioning, QA, and social execution — inside Innovid's advertising infrastructure.

67%

Average reduction in hands-on campaign ops time

98%

Best-case reduction

Telecom: 24–36 hours of Meta trafficking → 30 minutes.

7

Verticals Tested

Standardized enterprise workflow tests, Q2–Q3 2026.

91

Hours saved monthly

Enterprise Technology Meta program, NIVO combined with SAM.



Executive Summary

Agentic AI is moving from promise to measurable operational impact

Advertising operations weren't built for the speed modern marketers need. Campaign setup, creative decisioning, and pre-launch QA still rely heavily on manual processes — often spread across spreadsheets, briefing documents, and point-and-click configuration. As campaigns become more complex, that operational burden can slow teams down precisely when they need to move faster.

NIVO brings agentic intelligence directly into Innovid's advertising infrastructure, enabling purpose-built AI agents to take on repetitive campaign work across trafficking, decisioning, QA, and social execution.

Because Innovid operates independently across the full omnichannel delivery stack — with no media-buying conflict of interest — NIVO agents can work directly within live campaign structures to create placements, apply decision logic, validate tags and creative, and act on the billions of data points the platform already processes.

Across **seven standardized enterprise workflow tests spanning seven verticals**, NIVO reduced hands-on campaign operations time by an average of 67% compared with each client's documented before-state. The results show measurable efficiency gains across a range of campaign types, operating models, and industries.

*Q2–Q3 2026 · VERIFIED SINGLE-TASK TESTS



Why NIVO

Purpose-built agents, trained on how each client actually works.

NIVO's agents aren't generic automation — each is trained on a client's specific campaign structure, creative requirements, and decision logic before it goes to work. That's what makes single-prompt execution possible on tasks that otherwise take dozens of manual steps.





Campaign / Trafficking Agent

Automates campaign shell creation, placements, flighting, and asset mapping directly from media sheets.

"create campaigns and assign creatives"



Decisioning Agent

Builds granular targeting and creative-rotation logic through decision trees, including bulk updates across an entire rotation schedule.

"Add a keyword rule and schedule rotation."



QA Agent

Validates tags, declarative rules, creative assets, and tracking prior to launch, with a human-in-the-loop check.

"Analyze my campaign for quality assurance."



Social Agent

Automates campaign duplication and complex target-set assignment across regional markets, and automates campaign builds and QA on paid social platforms like Meta — increasingly deployed alongside Social Ads Manager (SAM).

"Replicate this campaign across four markets with localized targeting."




Seven industries, one clear pattern

Across seven independent enterprise tests, NIVO reduced hands-on campaign operations time by an **average of 67%** across single-task setup workflows. Every deployment targeted a verified operational bottleneck using agents trained on each client's specific decision logic.

The Scoreboard

% reduction in manual campaign operations time, per vertical.



29–98%

Range Across Verticals

4

Agents Used Across Tests

5 of 7

At Or Above 67%

nivo



Results By Vertical

Every deployment,
a verified bottleneck



Meta campaign trafficking efficiency test

Challenge: Time-consuming campaign builds and trafficking were slowing speed to market.



Social Agent





Decisioning Agent





Trafficking Agent

Travel & Aviation

Price-led campaigns across 7+ markets

Challenge: Frequent short-turnaround rotation between retail sales and always-on, feed-driven pricing.



58h
**Best-case
reduction**


Saved Annually

4 active global markets

Challenge: Multi-market campaign configuration at scale
(960 to 15,120+ creative versions per campaign).



Campaign / Trafficking Agent



67%
Less Time







Social Agent

CPG & Health

4 Southeast Asian markets

Challenge: Manual social campaign duplication and targetset assignment across regional markets.




29%
Less Time

Retail & Grocery

Weekly catalogue across 14 creatives / 28 placements

Challenge: Manual, weekly decision-tree assignment for time-of-day and audience-cluster messaging.



67%
Less Time
2 Hrs / Year,
Full Program



Decisioning Agent







Trafficking Agent




QA Agent



Automotive

8 vehicle campaigns, 300+ active creative & tracking tags


Challenge: Manual trafficking and QA across 100+ standard and dynamic creative files with frequently changing offers.



75%
Less Time

307 placements, 32 concepts, 5 decision trees

Challenge: End-to-end manual execution across campaign creation, placement setup, decision-tree rules, and prelaunch QA.



47%
Less Time



Full 3-Agent Suite





Social Agent



SAM

91h

Ongoing monthly savings

A high-volume Enterprise Technology Meta program measured **27 recurring hours per month** (~0.5 FTE) saved by NIVO on campaign builds (12 hrs) and custom QA (15 hrs) — **91 hours saved monthly combined with Social Ads Manager.**



Extrapolating the agent suite to a **high-volume deployment**

Based on verified test-run efficiency rates, applied to a large-scale enterprise setup. Clearly flagged as a modeled projection rather than a live-measured result.



+60h

Saved Per Setup

3.5x

Speed-Up

1

Agent Suite
Deployment

9,000 placements | **1,600** concepts | **400** decision trees





Less time managing campaigns, more time moving them forward

Across industries and workflows, these tests point to a consistent opportunity: repetitive campaign operations can be completed faster when agentic intelligence is built directly into the systems teams already use.

NIVO is designed to take on the manual, technical work that slows campaign execution — from trafficking and creative decisioning to QA and social campaign builds. Because agents can work independently or together, advertisers can apply automation to individual pain points or connect it across multiple stages of the campaign lifecycle.

The result is more than faster setup. By reducing the hands-on work required to launch, update, and manage campaigns, NIVO gives teams more capacity to focus on strategy, optimization, and the work that requires human judgment.

Ready To Put Agentic Advertising To Work?

Discover how NIVO reduces manual work across the campaign lifecycle.

Meet NIVO →



Methodology

Metrics are derived from independent agentic workflow tests conducted with Innovid clients across seven industry verticals during Q2 and Q3 of 2026, following a standardized measurement framework.



1 Real-world baselines. Measured against each client's own documented, pre-existing operational benchmarks ("before-state") to establish an accurate baseline.

2 Targeted deployment. Each test identified a real, verified operational bottleneck and deployed a specific NIVO agent (or combination) trained directly on that client's unique workflow.

3 Verified single-task tracking. The core 67% average metric is calculated exclusively from live client tests with a documented, single-task "before and after" measurement.

Note On Projections & Ongoing Metrics



- **Recurring operations** — cumulative, ongoing savings (such as the 27 recurring monthly hours saved alongside SAM) are excluded from the 67% average, as they reflect ongoing operational volume rather than a single setup event.
- **Enterprise-scale modeling** — figures for high-volume deployments (e.g., 9,000 placements, 1,600 concepts) are modeled extrapolations based on verified test-run efficiency rates, and are clearly flagged as projections rather than live-measured results.