NETACEA

EXECUTIVE BRIEFING | 2026

THE BUSINESS CASE

Agent Trust Management

A growing share of the transactions, queries, and interactions hitting your digital platform are no longer initiated by people. The question for the C-suite is not whether this traffic arrives. It is whether the organisation can distinguish what is valuable from what is extractive, and govern each accordingly.

$190–385bn

US agentic commerce spend projected by 2030

758%

YoY rise in AI-driven traffic to retailers, Nov 2025

53%

of all web traffic is now automated

NETACEA

THE LANDSCAPE

Automation is already commercial reality

AI shopping agents browse product catalogues, compare prices, and complete checkouts. Automated workflows query APIs for inventory and pricing data. Crawlers extract content at scale. The traffic is real, and its commercial consequences are already measurable.

McKinsey puts the global agentic commerce figure at **$3–5 trillion by 2030**. AI-driven traffic to retailers rose approximately **758% year-on-year** during the November 2025 holiday period, according to Adobe Analytics. This is not speculative. Around **23%** of Americans made a purchase using AI in the past month.

The question for the C-suite is not whether this traffic arrives. It is whether the organisation can distinguish what is valuable from what is extractive — and govern each accordingly.

23%

of Americans made a purchase using AI in the past month | Morgan Stanley AlphaWise, Dec 2025

$3–5tn

Global agentic commerce projection by 2030 | McKinsey, Oct 2025

THE GOVERNANCE GAP

From “bot or not?” to “how much do I trust this agent?”

Most digital businesses have some form of bot detection in place. What they do not have is a governance layer for legitimate automation. They can spot and stop a credential stuffing attack, but they cannot answer basic questions about the non-malicious machine traffic hitting their platform: who sent it, what it is authorised to do, and whether it aligns with the organisation’s commercial interests.

FORRESTER OCT 2025 | Renamed the bot management software market to **bot and agent trust management software** — the question has moved from “bot or not?” to “how much do I trust this agent?”

758% Adobe Analytics, Dec 2025

Year-on-year rise in AI-driven traffic to retailers during the November 2025 holiday period

53% Imperva 2026 Bad Bot Report

Automated traffic as a share of all web traffic in 2025, including both benign and malicious bots

NETACEA

THE CASE

7 reasons the C-suite should act now

Agent Trust Management delivers value across every major executive concern — from top-line revenue to regulatory compliance.

01 · GROWTH & REVENUE

A new transaction layer that does not exist today

Agent Trust Management is what lets the business participate in agentic commerce rather than block it. Trusted agents can become a new sales channel and a route to revenue. A business that cannot identify a legitimate shopping agent and govern its access to checkout flows will not capture that demand.

20% of US e-commerce projected to be agentic by 2030, up from 1% in 2026 (Morgan Stanley)

02 · RISK REDUCTION

The same architecture that enables commerce also enables fraud

The same AI agent architecture that enables legitimate commerce also enables credential stuffing, fake account creation, scraping, and fraud at scale. Trust controls reduce fraud losses and operational disruption directly, and limit the regulatory and reputational exposure that follows a breach.

46% of all account takeover incidents in 2025 were in financial services (Imperva, Feb 2026)

03 · COMPETITIVE PROTECTION

Your pricing and product strategy are already exposed

Pricing, inventory, content, and product strategy are high-value assets that can be exposed to agent scraping or automated probing without the organisation knowing it is happening. Trust management protects competitive intelligence while still allowing legitimate customer and partner automation to function.

04 · CUSTOMER TRUST & BRAND EQUITY

The opportunity scales with demonstrated governance

52% of consumers are comfortable sharing data with AI shopping agents, but 53% would require payment approval before any purchase, and 34% would only allow an agent to act with per-transaction sign-off. The businesses that win will show consumers their agent interactions are governed and accountable.

34%

would only allow AI agents to act with per-transaction approval (Radial, Mar 2026)

---

05 · OPERATIONAL CONTROL & DATA INTEGRITY

Protect the quality of commercial data

If AI agents inflate traffic, distort attribution, or pollute analytics, commercial decisions become unreliable. Agent Trust Management gives the business a way to classify, monitor, and govern autonomous activity in real time rather than relying on static rules — protecting forecasting accuracy, marketing ROI measurement, and demand signals.

---

06 · GOVERNANCE, COMPLIANCE & ACCOUNTABILITY

A defensible trust framework is not optional in regulated industries

Only ~30% of organisations have reached maturity level three or higher in agentic AI governance (McKinsey, March 2026). The EU AI Act enforces strict rules from August 2026, with potential fines up to 7% of global revenue. Trust management is the layer that makes agents auditable and explainable enough for enterprise use and legal review.

7%

of global revenue max EU AI Act fine from Aug 2026

---

07 · FASTER, SAFER ADOPTION

Trust reduces hesitation inside the business

Teams can adopt agents with bounded autonomy, monitored behaviour, and clear thresholds for escalation. Trust is what lets organisations move from pilots to production without losing control. Without it, the choice is between moving fast and losing visibility, or moving slowly and ceding ground to competitors who have solved the governance problem.

NETACEA

THE BUSINESS CASE FOR AGENT TRUST MANAGEMENT

HOW NETACEA DELIVERS AGENT TRUST

The trust layer for the automation economy

Agent Trust is a single integrated governance capability that answers every critical question about non-human actors touching your platform — before any business logic executes.

Identify

Surface every machine actor interacting with the platform. Agent Trust resolves humans, declared agents, authorised automation, and undeclared agents into distinct categories in the live traffic stream.

Verify

Establish who each machine actor is and what they are authorised to do. Agent Trust accepts cryptographically signed authority from emerging standards such as Visa TAP and Google AP2, and infers trust from behavioural signals.

Govern

Enforce the operator's policy in real time across every machine actor. Actions extend beyond allow-or-block to include challenge, step-up verification, monitor-only, and custom routing — calibrated by actor type and context.

Deployment | No code changes required

Agent Trust deploys at the edge or server layer with no JavaScript tags, client-side dependencies, or application code changes. Integration is fast, non-disruptive, and does not require involvement from product engineering teams.

Standard Support | Built for emerging agentic standards

As the industry converges on standards for machine-to-machine authority, including Visa TAP and Google AP2 — Agent Trust is already accepting cryptographically signed authority signals, positioning clients ahead of the compliance curve.

Executive Briefing

A tailored working session for your leadership team on the commercial impact of agentic traffic and the governance controls available today. Not a product demo — a working conversation about your platform, your traffic, and the decisions ahead.

REQUEST A BRIEFING