

Nobulex Progress Update — May 11, 2026

The Shift: From Receipts to Credit

When we applied, Nobulex made AI agents auditable. That's still true. But we've unlocked something bigger.

Bilateral receipts aren't just proof. They accumulate into **Trust Capital**: a machine reputation asset that determines what each agent is allowed to do. Higher Trust Capital = more autonomy, larger transaction limits, lower insurance premiums, enterprise approval.

Think of it as credit scores for AI agents. An agent earns trust the same way a person builds credit. Through verified behavior over time.

New one-liner: The trust economy for autonomous AI agents.

New tagline: Autonomy earned, not granted.

What Shipped Since We Applied

Trust Capital module — `packages/core/src/trust-capital.ts` now implements tier-based capability gating. Five trust tiers (restricted through autonomous), eight capability levels, weighted scoring from receipt chains. Wired into the enforcement middleware so every action automatically builds credit.

Microsoft Agent Governance Toolkit — Bilateral receipt primitive merged (PRs #1302, #1333). Listed as first adopter.

AAIF (Linux Foundation) — Project proposal #20 moved to "Under AAIF Staff Review" on May 11 by staff reviewer christinaharter. AAIF was founded by Anthropic, OpenAI, Google, Microsoft, AWS, and Block. Acceptance would place Nobulex alongside MCP and A2A as foundational agent infrastructure.

TC Sponsorship Confirmed — kenneives (AgentGraph/CTEF v0.3.1 maintainer) confirmed Technical Committee sponsorship. Litepaper "State of Agent Security 2026" (publishing May 12) explicitly cites Nobulex AAIF filing #20 in section 1.8.

Visa trusted-agent-protocol — First responder on Issue #16 (bilateral receipt signing for post-authentication audit trail). Same Ed25519 + JCS stack.

Cross-validation — Four independent implementations byte-match verified (TypeScript + Python). CTEF v0.3.2 spec publishes May 19-22.

Market Context

- Sierra raised \$950M at \$15B valuation (May 4, 2026). 40% of Fortune 50 using AI agents for insurance claims, mortgage originations, and customer service. None of those deployments have independent behavioral verification.
 - 88% of enterprises running AI agents had a security incident (Gravitee 2026).
 - Only 14.4% of AI agents go live with full security approval.
 - The Vercel/Context AI supply chain breach (April 19) cost \$2M on BreachForums because an AI agent had OAuth access with no per-action verification.
-

Revenue Model (Updated)

Transaction rail: 0.15% on verified agent transaction volume.

Trust futures: Enterprises pre-purchase Trust Capital verification capacity.

Compliance autopilot: EU AI Act Article 12 audit trail compliance (enforcement begins August 2026).

Insurance integration: Bilateral receipt chains as actuarial underwriting input for AI agent liability coverage.

Target: \$40-96B revenue by Year 10 as AI agents manage trillions in economic activity.

What We're Building Toward

Nobulex is the neutral trust layer for the agent economy. The same way credit bureaus made lending work at scale by making reputation portable and verifiable, Trust Capital makes agent autonomy work at scale by making behavioral history portable, verifiable, and economically valuable.

Website: nobulex.com **GitHub:** github.com/arian-gogani/nobulex **Blog:** nobulex.com/blog-ai-agents-need-credit-scores