

A Full-Stack AI Orchestration Patent Estate

11 USPTO Patent Applications · One Non-Provisional Under Examination · Two New May 2026 Filings

Stratus holds **eleven patent applications** protecting a complete AI orchestration stack — from the workflow language that defines how humans and AI collaborate, through the integration layer that embeds inside third-party platforms, to the federation governance that enables consented data sharing across licensees. **One non-provisional is under active USPTO examination;** ten provisionals establish priority dates spanning January through May 2026.

THE PORTFOLIO

#	PATENT	WHAT IT DOES	STATUS
1	COG™ Notation 63/959,818	A standardized language for writing down human–AI workflows — the way musical notation captures a score. Makes AI work repeatable, auditable, and teachable across teams.	PROVISIONAL
2	CPE™ Framework 63/961,608	Quality management methodology purpose-built for AI-augmented work. Brings process excellence discipline to how organizations govern, measure, and improve AI outputs.	PROVISIONAL
3	CPE-LSS™ 63/969,407	Adapts the classic Lean Six Sigma DMAIC cycle for AI workflows, giving enterprises a proven framework to drive continuous improvement in AI operations.	PROVISIONAL
4	AI Orchestration Layer 19/466,332	The core engine: coordinates multiple AI models in a single workflow — routing tasks, preserving context across handoffs, running models in parallel, keeping humans in the loop, and producing a deterministically replayable audit chain.	UNDER EXAM Non-Provisional
5	Multi-Model Scrutiny 63/975,310	The trust engine: multiple AI models independently evaluate each output, compute trust scores, and flag disagreements for human review. Learns which models are most reliable over time.	PROVISIONAL
6	Data Intelligence (Verantus™) 63/978,108	Generates comparative AI performance intelligence from anonymized workflow data — a Bloomberg terminal for AI model performance across industries and task types.	PROVISIONAL
7	Symphony™ Parallel Eval 63/979,134	A 28-variable engine that assesses AI output quality across five dimensions simultaneously. Delivers faster, more thorough quality assessments at enterprise speed.	PROVISIONAL
8	Predictive Routing 63/998,841	Before any AI model is invoked, predicts which model will produce the best result using a 3D compatibility tensor. Also detects when silent model updates cause performance drift.	PROVISIONAL
9	Closed-Loop Trust 63/998,857	Closes the feedback loop: human accept/override decisions automatically update routing parameters. The system gets smarter with every human judgment — no manual reconfiguration.	PROVISIONAL
10	Embedded Plugin Architecture (COG Switch™ HOST) 64/056,176	The integration layer: lets the orchestration fabric embed inside third-party host platforms with control/data-plane separation, multi-tenant isolation, per-transaction metering, cryptographic host-fabric binding, and audit-continuous fault tolerance.	NEW · PROVISIONAL Filed May 3, 2026
11	Cross-Licensee Data Governance 64/056,270	The federation governance layer: enables consented, jurisdiction-bounded data sharing across licensee deployments through differential privacy, sovereignty boundary controls, coordinated privacy budgets, and cryptographic provenance attribution.	NEW · PROVISIONAL Filed May 4, 2026

WHY THIS MATTERS

These eleven patents are not eleven variations on one idea — they are **eleven layers of a single integrated system**. Each protects a distinct technical function. Together they form a defensive moat that a competitor cannot replicate without infringing multiple independent claims.

The portfolio covers the complete lifecycle: how AI work is *defined* (COG), *governed* (CPE, CPE-LSS), *executed* (Orchestration Layer), *verified* (Symphony, Multi-Model Scrutiny), *optimized* (Predictive Routing, Closed-Loop Trust), *embedded* into third-party hosts (EPA / COG Switch HOST), *federated* across licensees under sovereign data controls (CLDG), and *monetized as intelligence* (Verantus).