# Sema4.ai Unveils 3.0 Platform Enabling Team-Based Work and Dramatically Reduces Token Spend

*The latest platform enhancements reduce AI processing costs and improve speed by reusing verified queries and logic, cutting LLM token usage by up to 80% and delivering up to 4× faster results*

**ATLANTA, GA – October 29, 2026** – Sema4.ai, the enterprise AI agent company, today announced the latest release of its Enterprise AI Agent Platform, introducing new capabilities designed to address three of the biggest challenges enterprises face as they move AI agents into production: trusting their output, controlling costs as usage scales, and enabling AI across team-based work.

This release empowers teams to share and manage agents, skills, memories, data, and files within governed workspaces, while introducing new ways to make repeated agent work faster, more consistent, and affordable. It also includes interoperability with personal productivity tools such as Claude Cowork, enabling teams to easily build and engage with agents while automating and sharing repeated work.

“Building an AI agent is no longer the hard part. The real challenge is making sure it can perform critical business work consistently, without relying on an AI model to figure out the same task every time,” said Rob Bearden, CEO and Co-Founder of Sema4.ai. “By turning repeatable work into verified, reusable logic, we’re helping enterprises move toward AI agents that become more efficient and predictable as they scale.”

**Bringing Trust, Efficiency and Governance to Enterprise AI**

AI models are probabilistic, so repeatedly asking one to generate the same query or calculation on every run can produce different answers and burn tokens each time. The latest version of the Sema4.ai platform moves more of each agent's repetitive work out of the model and into code that customers review and approve. The same input returns the same trusted output, and every step is recorded so teams can show auditors what work an agent did and why.

That same shift also cuts costs. Verified Queries reuse approved data-access patterns, using 80% fewer LLM tokens and delivering 4× faster results than repeated and probabilistic AI query generation. Verified Logic reuses approved calculations, checks, and transformations, resulting in 50% fewer tokens and 30% faster runs for repeat work. Enhanced DataFrames allow for calculations across large, structured datasets in memory without sending every row through a model’s context window.

Key updates include:

- **New support for team-based work**: shared workspaces, agents, skills, and files with full control; an improved business catalog built on Data Collections that gives every agent shared context; and Memories that capture customer-specific rules and policy adjustments so agents apply them on every run.
- **Enhanced ecosystem interoperability**: integrations with AI personal productivity tools such as Claude Cowork and Microsoft CoPilot, and Microsoft Fabric support. A new Agent Inbox streamlines Microsoft 365 and Google Workspace email monitoring, handles webhooks in real time, and routes work using plain-language rules.
- **Broader end-to-end lifecycle support**: Sai, Sema4.ai's integrated AI assistant is pervasive throughout the product experience and can suggest ways to improve workflows, instructions,

skills, and data context based on real usage, while forensic audit and reporting add an activity log with before-and-after diffs, an agent timeline, and search and filters.
- **Improved tokenomics**: cost control is built into the runtime, so usage can grow without runaway token spend and savings increase with expanded usage. Enhanced DataFrames process large, structured datasets without sending every row through the model's context window.

Together, these capabilities extend Sema4.ai across the full agent lifecycle, from building and deploying to governing, optimizing, and auditing, as adoption grows.

To learn more or schedule a demo, visit sema4.ai and read the 3.0 platform launch blog.

**About Sema4.ai**
Sema4.ai provides a comprehensive enterprise AI agent platform that enables organizations to build, run, and manage AI agents at scale. The platform empowers business users to create intelligent agents using natural language, connect agents to enterprise applications and data through pre-built actions and universal connectivity, and manage the full agent lifecycle with enterprise-grade security and governance. Sema4.ai is trusted by leading enterprises to transform business operations through autonomous, intelligent agents that work continuously with transparency, auditability, and control.