
Agent-First Company Landscape

A curated guide to the tools, platforms, and services powering AI agent infrastructure — evaluated for OpenClaw / MarkBot.

Prepared by MarkBot • March 11, 2026 • 38 tools across 11 categories
Each entry: what it does, what it enables, and what it costs.

Contents

01	Enterprise Agent Platforms	6 tools
02	Orchestration & Durable Execution	5 tools
03	Tool & Integration Platforms	5 tools
04	Memory & Long-Term State	3 tools
05	Agent-Native Browsers	6 tools
06	Agent Communication	1 tool
07	Observability & Evals	7 tools
08	Search & Document Intelligence	5 tools
09	Voice & Phone Agents	3 tools
10	Specialized Agent Services	1 tool
11	Agent Marketplaces & Registries	3 tools

1

Enterprise Agent Platforms

Full-stack platforms from the major cloud/AI providers. These define agent loops, tool invocation, and multi-agent coordination with enterprise controls.

OpenAI — Agents SDK

Usage-based

Production-ready SDK for building agent workflows with tool use, streaming, and handoffs between specialized agents. Tightly integrated with OpenAI's model and tool ecosystem.

► **Enables:** Multi-agent delegation patterns where MarkBot hands off coding tasks to a specialized agent, with structured handoff metadata. Could complement our Codex pipeline with native OpenAI agent orchestration.

Anthropic — Model Context Protocol (MCP)

Free (protocol)

Open standard for connecting AI agents to external tools, data sources, and resources. Defines a universal interface so any agent can discover and invoke any MCP-compatible tool.

► **Enables:** A single tool interface across all MarkBot agents instead of custom integrations for each service. Instant access to the growing ecosystem of community MCP servers (GitHub, Slack, databases, etc.).

AWS — Bedrock AgentCore

Consumption-based

Managed agent runtime with built-in identity, memory, browser, code interpreter, and observability. Modular — use only the pieces you need, all within AWS infrastructure.

► **Enables:** If MarkBot ever needs cloud-scale deployment (multi-user, enterprise), AgentCore provides managed sandboxes and identity without building from scratch. Useful for ForgeClaw scaling.

Google — Vertex AI Agent Builder + ADK

vCPU/GiB-hour + model fees

Agent lifecycle platform with a model-agnostic open-source framework (ADK). Covers build, deploy, and monitor with explicit multi-agent system guidance and a growing integrations ecosystem.

► **Enables:** A Gemini-based agent tier for MarkBot with first-party Google Workspace integration (Calendar, Gmail, Drive) via ADK connectors, without custom API wiring.

Microsoft — AutoGen

Free (framework)

Event-driven multi-agent framework with a visual studio UI. Supports MCP tools, Docker code execution, and distributed runtime for scalable agent coordination patterns.

► **Enables:** Structured multi-agent workflows where a researcher agent, a coder agent, and a reviewer agent collaborate on complex tasks — with visual debugging in AutoGen Studio.

Salesforce — Agentforce

\$2/conversation or per-user

Autonomous AI agents embedded in the Salesforce ecosystem. Takes actions on CRM data, handles customer interactions, and integrates with Salesforce workflows and voice channels.

▶ **Enables:** If Forge moves CRM to Salesforce, MarkBot could trigger Agentforce for customer follow-ups, lead qualification, and automated service responses directly from CRM data.

2

Orchestration & Durable Execution

Frameworks for building reliable, resumable agent workflows with checkpointing, retries, and human-in-the-loop interruption.

LangChain — LangGraph + LangSmith

\$0–\$39/seat

Complete agent engineering suite: LangGraph provides durable state with checkpoint threads and time-travel debugging; LangSmith adds tracing, evals, and deployment infrastructure.

► **Enables:** [Crash-proof MarkBot workflows that resume exactly where they left off after failures. Full trace visibility into every tool call, token spent, and decision point across multi-step tasks.](#)

LlamaIndex — Workflows + LlamaParse

Free–\$1.25/1K credits

Agent framework optimized for document-centric workflows with first-class memory and state management. LlamaParse extracts structured data from PDFs, invoices, and complex documents.

► **Enables:** [Automated QBO receipt processing — MarkBot ingests PDF invoices, extracts line items, and categorizes them without manual data entry. Could power the bookkeeping pipeline end-to-end.](#)

CrewAI

Free–\$25/mo–Enterprise

Multi-agent platform with visual workflow builder and enterprise controls (SOC2, SSO, secret management, PII masking). Available as SaaS or self-hosted on Kubernetes.

► **Enables:** [Visual design of multi-agent "crews" for complex operations — e.g., a Forge Intelligence crew where a research agent, analysis agent, and reporting agent collaborate on competitive analysis.](#)

Inngest — AgentKit

Free tier + usage overages

Orchestration-first framework built on a durable execution platform with production-grade queueing, rate limiting, retries, and idempotency baked into the runtime.

► **Enables:** [Replacing our custom cron/retry logic with a battle-tested execution layer. MarkBot's nightly research pipeline could use Inngest for automatic retries, concurrency control, and step-level observability.](#)

Portia AI

Free tier + per-run pricing

Open-source multi-agent framework emphasizing explicit plans with human approval gates, stateful task execution, and audit trails. Built for compliance-forward use cases.

► **Enables:** [Formal approval workflows for sensitive MarkBot actions \(financial transactions, external communications\) with full audit trails that could satisfy Forge compliance requirements.](#)

3

Tool & Integration Platforms

Managed authentication, tool catalogs, and execution governance. These let agents connect to hundreds of SaaS apps without building custom OAuth flows.

Composio

Free tier + paid by tool calls

Agent-first connector platform with managed OAuth for 250+ SaaS integrations, exposed via MCP or direct API. Handles token refresh, scoping, and least-privilege automatically.

► **Enables:** Instant MarkBot access to HubSpot, Notion, Linear, Slack, and dozens more without building custom auth for each. Could replace our manual GOG/API key management with proper OAuth flows.

Pipedream Connect

Credit-based + per-user

Embeddable integration toolkit with thousands of pre-built connectors, managed auth, and code-level control. Positions itself as infrastructure for adding integrations to apps or AI agents.

► **Enables:** Rapid prototyping of new MarkBot integrations (ConnectTeam, BuildOps, Eight Sleep) without writing auth code. Ship a new integration in hours instead of days.

Nango

Published tiers (varies)

Unified API platform handling auth, execution, scaling, and observability for 700+ API integrations. Supports both real-time tool calls and ongoing background data sync pipelines.

► **Enables:** Continuous background sync of BuildOps job data, ConnectTeam schedules, and HubSpot contacts into LifeOS — always fresh, without polling or manual triggers.

Merge

Sales engagement required

Unified API provider normalizing data across HRIS, ATS, CRM, accounting, and ticketing categories. Provides consistent object schemas regardless of the underlying provider.

► **Enables:** A single "employee" API for Forge that works whether the backend is ConnectTeam, Rippling, or any future HR platform — no rewrite needed when systems change.

Arcade

Free tier + per-execution

MCP-compatible runtime sitting between AI and action, handling user-specific permissions, OAuth/IDP integration, and governance. Think of it as a permission layer for agent tool calls.

► **Enables:** If MarkBot ever serves multiple family members or Forge employees, Arcade provides per-user permissions without building auth plumbing. Each person gets scoped access to their own data.

4

Memory & Long-Term State

Dedicated memory services that give agents persistent, structured recall across sessions — beyond what's in the context window.

Letta (formerly MemGPT)

Unspecified

Agent memory runtime that manages what an agent remembers across conversations with an agent-native mental model. Evolved from the influential MemGPT research on self-editing memory.

► **Enables:** A structured memory layer where MarkBot autonomously decides what to remember vs. forget, with retention policies and the ability to delete/rectify memories on demand.

Zep

Unspecified

Long-term memory layer for AI agents and assistants with deterministic retrieval semantics. Focuses on making memory reliable and auditable for production deployments.

► **Enables:** Replacing our file-based memory system (MEMORY.md, daily logs) with a queryable memory service that guarantees consistent retrieval across all MarkBot sessions.

Mem0

Free tier + paid plans

Persistent memory platform with multiple memory types (episodic, semantic, procedural) and enterprise features including on-prem deployment, SSO, and audit logs. Available as hosted service or open-source.

► **Enables:** Structured memory that distinguishes between facts MarkBot learned (semantic), things that happened (episodic), and how to do things (procedural) — with self-hosted control over all data.

5

Agent-Native Browsers

Cloud browser infrastructure purpose-built for AI agents — with session persistence, CAPTCHA handling, proxies, and agent-facing APIs.

Browserbase

\$99–\$499/mo + overages

Cloud headless browser platform with persistent cookie/state via a Contexts API and deep integrations with agent frameworks. Used by Stripe and others for agent-driven web workflows.

► **Enables:** Offloading browser automation to the cloud when local Camoufox is insufficient (e.g., sites requiring US residential IPs or specific browser fingerprints we can't replicate locally).

Airtop

Free tier + credit-based

Browser automation platform for agents with shared authentication, embedded live view for debugging, and explicit concurrency controls. Gives AI agents structured web access.

► **Enables:** Live-view debugging of MarkBot browser sessions — watch what the agent sees in real-time from your phone when troubleshooting a failed web automation.

Steel.dev

Per-hour + CAPTCHA + proxy

Open-source browser API for AI agents with transparent infrastructure-style metering. Provides session APIs, CAPTCHA solving, and proxy bandwidth as separately priced primitives.

► **Enables:** A fallback browser service for CAPTCHA-heavy sites (utility bill payments, government portals) where Camoufox alone isn't enough and you need managed CAPTCHA solving.

Browserless

Free tier + unit-based

Mature managed headless browser platform compatible with Playwright and Puppeteer via WebSocket/CDP endpoints. Widely used in production automation stacks.

► **Enables:** Drop-in cloud Playwright sessions when local browser resources are constrained (e.g., running multiple concurrent browser automations during peak research hours).

Skyvern

Unspecified

Computer-vision-based web automation that uses visual understanding instead of DOM selectors. Significantly more resilient to website redesigns since it "sees" the page like a human.

► **Enables:** Reliable automation of sites that constantly change their layout (BuildOps, vendor portals) without rewriting selectors every time they redesign.

MultiOn

Unspecified

High-level web automation API where you describe tasks in natural language ("book a flight to Boston on March 20") and the agent handles all browser interaction autonomously.

▶ **Enables:** [Delegating entire web tasks to a single API call instead of scripting every click](#) — useful for one-off tasks like "order this item on Amazon" or "schedule this appointment."

Email and messaging infrastructure purpose-built for AI agents to send, receive, and manage conversations.

AgentMail

Free tier + paid plans

Email inbox API that gives AI agents their own email addresses with threading, labeling, and search. Agents can sign up for services, receive OTPs, and maintain long-lived email conversations.

▶ **Enables:** [MarkBot gets its own email identity for service signups and vendor communications instead of using Mark's personal Gmail. OTP verification flows become fully automated.](#)

Tracing, monitoring, evaluation, and cost accounting for AI agent systems. Critical for reliability and spend control.

AgentOps

Unspecified

Agent-specific observability platform designed for tracing multi-step agent workflows rather than individual LLM calls. Provides agent-native dashboards and failure analysis.

► **Enables:** [End-to-end visibility into MarkBot's research pipeline — trace from Brave search through Qwen extraction to Opus editorial to newsletter delivery, spotting failures and bottlenecks.](#)

Helicone

Free tier + paid plans

LLM observability proxy that sits between your app and model providers, logging every request with latency, cost, and token metrics. Works across OpenAI, Anthropic, and others simultaneously.

► **Enables:** [Unified cost tracking across all MarkBot model providers \(Anthropic, OpenAI, Google\) through a single proxy — could replace our custom token tracker with richer analytics.](#)

Langfuse

Free (open-source)

Open-source LLM observability with tracing, evaluations, and prompt management. Self-hostable for full data control, or use the hosted version for convenience.

► **Enables:** [Self-hosted tracing on Big Daddy with zero data leaving the network. Full visibility into every MarkBot conversation, tool call, and token spent — with data sovereignty.](#)

Arize AI — Phoenix

Free (open-source)

Open-source LLM observability framework widely used as a baseline for tracing and debugging AI applications. Provides trace visualization and performance analysis.

► **Enables:** [A lightweight, self-hosted tracing dashboard that complements our audit system — see exactly which tool calls are slow, which models are expensive, and where failures cluster.](#)

Galileo

Unspecified

LLM and agent quality monitoring platform focused on evaluation metrics, regression detection, and governance reporting for production AI systems.

► **Enables:** [Automated quality scoring of MarkBot's outputs — detect when response quality degrades, when the newsletter gets worse, or when the research pipeline produces lower-quality findings.](#)

Traceloop

Published tiers (varies)

LLM tracing built on OpenTelemetry conventions, making it compatible with existing observability infrastructure. Provides correlation IDs across tools and async agent steps.

► **Enables:** Standardized tracing that plugs into any OpenTelemetry-compatible dashboard — useful if we ever want to unify MarkBot observability with Forge's existing monitoring stack.

Braintrust

Free/open tiers available

Evaluation and testing platform for AI systems with regression testing, prompt iteration tracking, and A/B comparison tools. Designed for test-driven AI development.

► **Enables:** Regression testing for MarkBot — verify that prompt changes, model upgrades, or pipeline modifications don't degrade output quality before deploying them.

8

Search & Document Intelligence

Web search, content extraction, and document parsing services built specifically for AI agent consumption.

Tavily

Free tier + usage-based

Web search API explicitly built for AI agents, returning clean, structured results optimized for LLM consumption rather than human browsing. Provides real-time web grounding.

► **Enables:** [Already in use for MarkBot's interactive research. Could expand to real-time fact-checking during conversations and automatic source verification for newsletter claims. \(Currently deployed\)](#)

Firecrawl

Free tier + paid plans

Web crawling and extraction service that converts any webpage into clean markdown or structured data optimized for LLMs. Handles JavaScript rendering, pagination, and anti-bot challenges.

► **Enables:** [Deep website analysis for Forge competitive intelligence — crawl competitor sites, extract pricing/service pages, and feed structured data into MarkBot's research pipeline.](#)

Exa

Per-operation pricing

Semantic search engine for agents with dedicated "research" operations that combine search and page reading into a single workflow. Designed for high-recall research tasks.

► **Enables:** [Higher-quality nightly research by replacing keyword search with semantic understanding — "find articles about HVAC companies adopting AI" instead of matching exact keywords.](#)

Perplexity — Search API

~\$5/1K requests

Search API from Perplexity with predictable per-request pricing and no token costs. Returns search results backed by Perplexity's crawl and index infrastructure.

► **Enables:** [A cost-predictable search alternative for high-volume research — fixed per-request pricing makes budgeting straightforward compared to token-based search APIs.](#)

Jina AI — Reader API

Free (rate-limited)

Simple web-to-text conversion: prefix any URL and get clean, LLM-ready content back. No setup, no API key for basic usage. Pairs well with any search API for content extraction.

► **Enables:** [Already powering our nightly research extraction pipeline. Zero-cost content extraction that converts any webpage into clean text for Qwen/Sonnet processing. \(Currently deployed\)](#)

Programmable voice and telephony platforms that let AI agents make and receive phone calls with natural conversation.

Vapi

~\$0.05–0.15/min + model costs

Developer-first voice agent platform with per-minute pricing covering STT, LLM, and TTS. Programmable telephony with low-latency conversation and webhook-based control.

► **Enables:** [MarkBot answering phone calls on your behalf](#) — screen calls, take messages, schedule callbacks. Could handle Forge customer calls after hours with proper escalation.

Retell AI

Per-component + concurrency

Full-lifecycle AI phone agent platform with SDKs for Node/Python, explicit concurrency controls, and built-in monitoring. Covers build, test, deploy, and monitor for voice agents.

► **Enables:** [Production-grade voice escalation for Forge](#) — when MarkBot detects an urgent customer issue, it calls the right technician or manager with context, tracks the outcome.

Bland AI

~\$0.07–0.12/min

High-volume voice automation platform with dedicated infrastructure and enterprise positioning. Optimized for outbound calling campaigns with per-minute billing.

► **Enables:** [Automated Forge appointment confirmations and follow-up calls at scale](#) — MarkBot calls every customer the day before their service appointment to confirm.

10 Specialized Agent Services

Domain-specific agent products that handle entire categories of work autonomously.

Cognition — Devin

Team plans + usage (ACUs)

Autonomous software engineering agent that can plan, code, test, and deploy across full codebases. Generally available with API access for programmatic delegation of coding tasks.

▶ **Enables:** [Delegating entire coding tasks to Devin as an alternative to Codex](#) — e.g., "build the Mission Control cost charts" as a single API call that returns working, tested code.

11 Agent Marketplaces & Registries

Discovery layers for finding, installing, and paying for agent tools and capabilities.

Smithery

Free (registry)

Community registry and installer for MCP servers. Browse, discover, and install tool servers with one command. Think "npm for agent tools."

► **Enables:** [One-click installation of community MCP servers – instantly add GitHub, Postgres, Slack, or any other MCP tool to MarkBot's toolkit without writing integration code.](#)

Coinbase – x402 Bazaar

Free (protocol)

Machine-readable catalog of payable API endpoints with an open payments protocol (x402). Designed for agents to autonomously discover and purchase tool calls using cryptocurrency.

► **Enables:** [A future where MarkBot can autonomously find and pay for specialized services \(data lookups, premium APIs\) without pre-configured API keys – with spend controls you define.](#)

OpenAI – GPT Store

ChatGPT subscription

Marketplace for custom GPT agents within the ChatGPT ecosystem. Distribution channel for packaged agent experiences, though limited to OpenAI's platform.

► **Enables:** [If you ever want to distribute a MarkBot-lite experience to Forge employees or clients, the GPT Store is a zero-infrastructure distribution channel with built-in billing.](#)