CORAL AI LABS

INVESTOR DATA ROOM / 2026

Market Opportunity and Competition

Enter through software engineering; expand through managed coordination

$500K three-month commercialization sprint

Confidential investor diligence materials • 18 July 2026

The infrastructure gap

Model capability is improving faster than the operating layer around agents. As agents run longer, consume more tokens and cross more systems, buyers need coordination, context selection, spend control, identity, observability and evidence.

Signal	External evidence	Why it matters to us
Coordination penalty	CooperBench reports coding agents average 30% lower success together than alone	Threads, explicit topology, scoped agents and review address the seam between capable workers
Token growth	Goldman Sachs Research forecasts 24x token use to 120 quadrillion tokens per month by 2030	Proxying, model substitution, budgets and receipts become infrastructure
Coding cost pressure	Gartner expects AI coding cost to exceed the average developer salary by 2028	Route smaller tasks to smaller models and govern cost per accepted outcome
Existing spend category	Grand View estimates AI code tools at $9.8bn in 2026 and $26.0bn in 2030	We enter an established budget through Coral Code before broader platform expansion

Wedge and expansion

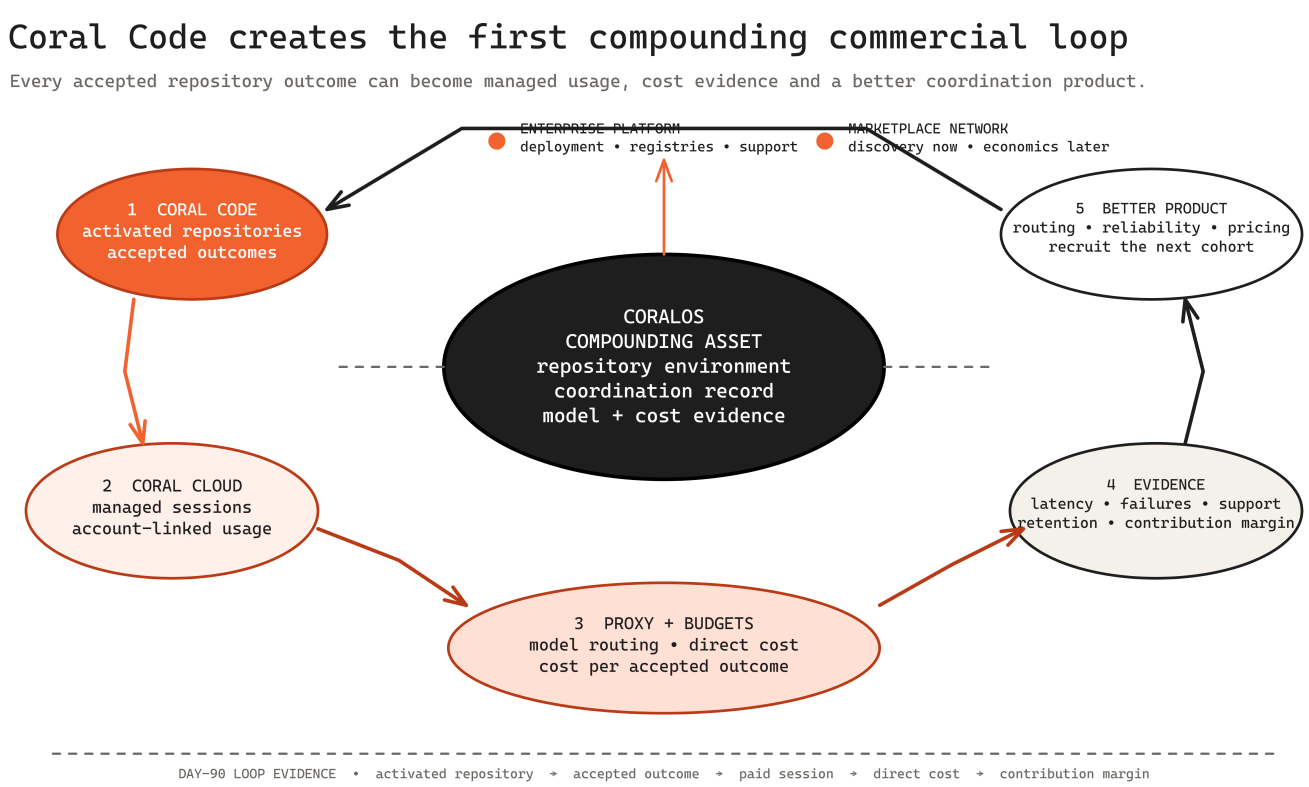

We use Coral Code to create the first paid loop; Cloud and the platform compound the value of each cohort.

Layer	Buyer need	Commercial proof
Coral Code	Useful work across large repositories at controlled cost	Activated and retained repositories; accepted operations
Coral Cloud	Hosted sessions, model access, telemetry and budgets	Paid sessions, direct cost and support by account
Enterprise platform	Deployment, controls, registries and reusable agents	Annual contract, reliability and security evidence
Marketplace network	Discoverable, attributable and eventually payable agent supply	Remote execution and settlement after release

First customer segments

Segment	Pain	Entry offer	Proof
AI-forward engineering teams	Large repositories, context loss, parallel-agent review	Coral Code cohort	Activation, retention, accepted work
AI platform teams	Fragmented agents, credentials, budgets and observability	Cloud API and Server	Paid sessions and integration time
Coding-agent vendors	Need richer repository context and coordination	Embedded Code/Server integration	Outcome lift and cost per task
Regulated software teams	Data boundaries, identity, approvals and audit	Enterprise deployment	Security review and design-partner result
AI-native application builders	Need multi-agent infrastructure without building it	Cloud sessions and agent registry	Time to production and retained use

Competitive position

Category	Primary strength	How we relate	Our distinction
Frontier coding agents	Reasoning, tool use and distribution	Execution partner	We preserve the outer agent while adding repository-scoped peers and controls
Agent frameworks	Application composition and workflow logic	Complement or integration source	We supply a live multi-agent runtime, communication graph and operating plane
MCP	Agent-to-tool interoperability	A core interface we use	We add agent-to-agent coordination, lifecycle, credentials and state
Model gateways	Provider routing and billing	Adjacent infrastructure	We link model use to agent topology, work claims, sessions and outcomes
Code intelligence	Search, indexing and static analysis	Input to our product	We turn evidence into bounded agents, questions, edit authority and durable results
Agent communication protocols	Interoperable message exchange	Protocol peer	We combine communication with runtimes, API control, budgets, Console and marketplace discovery

We expect competitors to add coordination features. Our defence is the repository evidence model, coordination record, model economics and production control plane operating as one system.

Demand measured from use

We will build the commercial forecast from qualified repositories, retained teams, managed sessions, direct model cost and enterprise contracts. Market estimates define the available budget category; our cohorts will define serviceable demand.

Sources

[S3] CoralOS documentation. https://docs.coralos.ai/

[S5] Coral Server v1.4.0 release. https://github.com/Coral-Protocol/coral-server/releases/tag/v1.4.0

[S8] Coral Marketplace status. https://docs.coralos.ai/concepts/marketplace

[S9] Coral LLM proxy. https://docs.coralos.ai/concepts/proxy

[M1] CooperBench. https://arxiv.org/abs/2601.13295

[M2] Gartner AI coding cost forecast. https://www.gartner.com/en/newsroom/press-releases/2026-06-24-gartner-predicts-ai-coding-costs-will-surpass-average-developer-salary-by-2028-as-token-consumption-surges

[M3] Goldman Sachs Research agent usage forecast. https://www.goldmansachs.com/insights/articles/ai-agents-forecast-to-boost-tech-cash-flow-as-usage-soars

[M4] Grand View Research AI code tools market. https://www.grandviewresearch.com/industry-analysis/ai-code-tools-market-report