The DeployCo Landscape

DeployCos.com cheat sheet. Updated 11-08-2026. Sources at deploycos.com/sources/.

Provider	Type	Money	Model choice	Delivery model
OpenAI Deployment Company	lab-backed DeployCo	$4,000,000,000+ initial investment (announced)	single model	A diagnostic, then a few priority workflows built and run in production by embedded engineers
Ode with Anthropic	lab-backed DeployCo	$1,500,000,000 (press-reported; not in Anthropic's announcement)	multi-model	Small engineering teams build custom Claude-powered systems with Anthropic Applied AI engineers; Claude-first, other models where required
Palantir (forward deployed engineers, AIP bootcamps)	software vendor FDE	Not disclosed	own platform	Forward deployed engineers deploy and extend Foundry and AIP at the customer; AIP bootcamps open the sale
AWS Forward Deployed Engineering	hyperscaler unit	$1,000,000,000	multi-model	Pods of five or six engineers with one customer for about 45 days, ending with an exit plan
Microsoft Frontier Company	hyperscaler unit	$2,500,000,000	multi-model	Embedded engineers and industry experts build, launch and keep improving AI systems
Google Cloud partner-embedded FDEs	hyperscaler unit	$750,000,000 partner fund	multi-model	Google engineers embed alongside systems integrators, who deliver most of the work
Distyl AI	AI-native startup	$175,000,000 Series B at a $1,800,000,000 valuation (09-2025 funding)	any	Forward deployed engineers on site plus the Distyl platform (formerly Distillery); part of the project fee tied to client objectives, plus platform licensing
Global systems integrators (Accenture, Deloitte and peers)	systems integrator	About $11,500,000,000 cumulative advanced AI bookings at Accenture through Q1 FY26 (last disclosed)	any	Consulting-led teams for strategy, integration and change management; also co-deliver for the lab and hyperscaler programs
In-house AI team	in-house	Not disclosed	any	Your own engineers build and run AI systems with no outside dependency