

Anthropic & OpenAI IPO Supplier Watchlist

Paragraph-format research list: public equities linked directly or indirectly to Claude and OpenAI infrastructure demand. Prepared August 26, 2026.

IMPORTANT

This research list is not investment advice. A relationship with Anthropic or OpenAI does not ensure revenue, earnings, or stock-price gains. IPO timing and terms are unconfirmed; supplier benefits depend on binding contracts, deployment timing, capacity utilization, margins, financing, valuation, and the broader AI-capex cycle.

01 Direct cloud and compute

AMZN

Amazon

EXPOSURE: ANTHROPIC DIRECT; OPENAI INDIRECT

AWS is Anthropic's primary cloud and training partner, with reported commitments exceeding \$100 billion over 10 years and capacity up to 5GW. If Claude continues to scale, Amazon can benefit across AWS compute, Trainium chips, storage, networking, and data-center utilization. This is the most durable Anthropic infrastructure proxy, but Amazon is diversified enough that Anthropic is unlikely to be the sole driver of its stock.

GOOGL

Alphabet

EXPOSURE: ANTHROPIC DIRECT

Google Cloud supplies Claude GPU and TPU infrastructure and Alphabet is a strategic backer. Anthropic's growth can increase Google Cloud utilization and TPU demand. The key caveat is that Alphabet's share price is still more exposed to advertising, its own Gemini model strategy, and broader cloud growth than to Claude alone.

MSFT

Microsoft

EXPOSURE: OPENAI CORE; ANTHROPIC CAPACITY OPTIONALITY

Microsoft is the principal public OpenAI platform proxy through Azure and enterprise distribution, and it also has reported Azure-capacity involvement with Anthropic. It may work if either lab expands workloads because Azure monetizes compute while Microsoft also distributes AI in enterprise software. This broad exposure is less binary than a dedicated GPU cloud provider.

ORCL

Oracle

EXPOSURE: OPENAI DIRECT

Oracle is a main Stargate infrastructure partner, with reported long-term OpenAI infrastructure commitments. If OpenAI's demand scales, Oracle Cloud Infrastructure could capture high utilization and contracted revenue. The major counterweight is execution: Oracle must finance, construct, power, and deliver enormous data-center capacity on time.

CRWV

CoreWeave

EXPOSURE: ANTHROPIC AND OPENAI DIRECT

CoreWeave has a multiyear Anthropic compute agreement and major OpenAI-related capacity commitments. It is the highest-beta listed dual-lab compute proxy because it is focused on selling scarce AI infrastructure. That focus also adds risk from customer concentration, funding requirements, GPU availability, contractual economics, and valuation.

AKAM

Akamai

EXPOSURE: ANTHROPIC DIRECT, BASED ON REPORTED CONTRACT DISCLOSURES

Anthropic reportedly signed a \$1.8 billion, seven-year computing agreement with Akamai. The appeal is a relatively focused Claude-capacity thesis that can alter the company's cloud-growth profile. Investors should monitor the ramp, actual utilization, capital intensity, and economics because the relationship is less transparent than large hyperscaler arrangements.

02 Chips and silicon**NVDA**

Nvidia

EXPOSURE: ANTHROPIC AND OPENAI DIRECT/INDIRECT

Both ecosystems rely on GPU-heavy AI infrastructure, and CoreWeave's Claude work is Nvidia-centric. Nvidia is the clearest liquid hardware read-through for increased training and inference spend. It is not a pure Anthropic/OpenAI trade, however: its valuation and results reflect demand from the entire global AI-capex cycle.

AMDAdvanced Micro
Devices**EXPOSURE: ANTHROPIC DIRECT; OPENAI SUPPLY OPTIONALITY**

Anthropic reportedly agreed to as much as 2GW of AMD Instinct MI450 capacity beginning in 2027, alongside planned AMD investment. AMD can benefit disproportionately if frontier labs seek an alternative accelerator supply chain and negotiate against Nvidia. The critical risks are product delivery, performance, software ecosystem adoption, and whether capacity commitments become recognized revenue on schedule.

AVGO

Broadcom

EXPOSURE: ANTHROPIC AND OPENAI DIRECT

Anthropic, Google, and Broadcom announced next-generation TPU capacity, while OpenAI also has a custom-accelerator roadmap with Broadcom. Broadcom is a compelling dual-lab custom-silicon expression because it can gain when workloads shift from general-purpose GPUs toward purpose-built accelerators and high-speed networking. Lead times are long, so a significant portion of the thesis rests on future deployment rather than near-term sales.

ARM

Arm Holdings

EXPOSURE: SECOND-ORDER / POTENTIAL STARGATE PARTICIPANT

Arm can benefit from server CPUs, AI control-plane workloads, and custom-silicon design activity, but it is not a clear direct compute-revenue proxy for either lab. Treat it as an architecture and semiconductor-IP exposure, not as a confirmed Anthropic or OpenAI supplier contract.

03 Networking and data centers**ANET**

Arista Networks

EXPOSURE: THEMATIC

Scaling AI training and inference clusters require high-bandwidth Ethernet switching fabrics. More Claude and OpenAI clusters can therefore lift network demand, but no disclosed direct contract was identified. This is a broader AI-networking trade, and its results will depend on spending by hyperscalers and cloud providers, not simply IPO sentiment.

MRVL

Marvell Technology

EXPOSURE: THEMATIC

Marvell can benefit from custom AI silicon, optical interconnect, and cloud networking as AI clusters grow in size and bandwidth needs. The linkage is logical but indirect; it should be treated as a general AI-infrastructure beneficiary rather than a documented supplier to either Anthropic or OpenAI.

GLW

Corning

EXPOSURE: ANTHROPIC SECOND-ORDER

Amazon has a multibillion-dollar fiber arrangement supporting expanding U.S. AI data centers, making Corning a potential AWS/Anthropic build-out beneficiary. More distributed AI capacity means more fiber connectivity. Still, no direct Anthropic revenue relationship was identified, so this is a picks-and-shovels infrastructure proxy.

COHR

Coherent

EXPOSURE: THEMATIC

Coherent supplies optical components and transceivers needed for high-speed data-center connectivity. It may benefit when AI clusters require faster, denser links within and between facilities. The name is exposed to broader optical, industrial, and networking cycles, so it is not a clean IPO sympathy trade.

CIEN

Ciena

EXPOSURE: THEMATIC

If frontier AI capacity increasingly spans multiple sites, high-capacity optical transport becomes more valuable. Ciena is a longer-duration network-build thesis rather than an immediate Anthropic/OpenAI supplier play, and no specific lab contract was identified.

04 Power and physical infrastructure**VRT**

Vertiv

EXPOSURE: THEMATIC

Dense AI racks require cooling, power distribution, backup systems, and thermal management. If Claude and OpenAI add gigawatts of capacity, those physical requirements should expand. Vertiv is one of the clearest non-chip AI-infrastructure beneficiaries, although no named direct supplier agreement was identified.

ETN

Eaton

EXPOSURE: THEMATIC

AI campuses need switchgear, transformers, electrical distribution, and power-quality equipment. Eaton benefits from construction and energization phases of data-center capex, making it less dependent on an IPO headline and more dependent on actual physical build activity.

GEV

GE Vernova

EXPOSURE: THEMATIC

Anthropic has said it will fund grid upgrades and seek net-new power generation for its data centers. That makes grid and generation equipment companies strategically relevant. GE Vernova is an indirect proxy, though no confirmed Anthropic or OpenAI supply agreement was identified.

CEG**EXPOSURE: THEMATIC**

Constellation Energy

Large, steady AI data-center loads can make reliable carbon-free power more valuable. The investment case relies on specific electricity contracts, local market rules, available capacity, and interconnection timing—not merely Anthropic or OpenAI IPO activity.

PWR**EXPOSURE: THEMATIC**

Quanta Services

AI facilities need transmission, substations, grid upgrades, and construction execution. Quanta is an indirect “the grid has to be built” expression and can benefit from the physical expansion of AI capacity. No direct model-lab agreement was identified.

Practical focus list**MOST DIRECT DUAL-LAB INFRASTRUCTURE LOGIC**

CRWV, NVDA, AVGO, AMD, MSFT, and ORCL. CRWV is the highest-beta compute name; Nvidia and Broadcom are the major hardware/custom-silicon exposures; Microsoft and Oracle are the key platform and cloud expressions. AMZN and GOOGL are the strongest Anthropic-centric platform names.

MORE CONSERVATIVE INFRASTRUCTURE SLEEVE

AMZN, GOOGL, MSFT, NVDA, AVGO, ORCL, VRT, and ETN. Higher-beta additions: CRWV, AMD, AKAM, ANET, and MRVL. The networking names are thematic AI-capex beneficiaries, not confirmed direct suppliers.

RESEARCH BASIS

- Anthropic public materials on AWS/Trainium, Google Cloud, Broadcom TPU capacity, U.S. infrastructure investment, and grid/electricity commitments.
- OpenAI Stargate materials on Oracle, SoftBank, and CoreWeave infrastructure relationships.
- Independent reporting on cloud and chip agreements, including Oracle/OpenAI infrastructure commitments, CoreWeave, AMD/Anthropic, Akamai/Anthropic, and Amazon–Corning infrastructure activity.