What a Two-Hour Conversation About Wafers and Watts Tells Us About the AI Supply Chain

Published September 22, 2026

Christopher Gannatti, CFA
Global Head of Research

Key Takeaways

- Anthropic and OpenAI could account for roughly 40% to 50% of incremental global compute additions in 2027, according to Dylan Patel's estimate.
- AI infrastructure bottlenecks extend beyond accelerators into memory, lithography, interconnects, power generation, cooling and advanced packaging.
- The WisdomTree Artificial Intelligence and Innovation Fund (WTAI) spans multiple layers of the AI supply chain, from NVIDIA and TSMC to GE Vernova and Vertiv.

A recent long-form conversation between podcast host Dwarkesh Patel and Dylan Patel, founder of the semiconductor research firm SemiAnalysis, is less an interview than a working model of where AI compute economics are headed.1 Below, we walk through the core claims from the conversation and connect them, sector by sector, to names investors will recognize.2 We will also specifically note the exposures within the WisdomTree Artificial Intelligence & Innovation Index, which underlies the WisdomTree Artificial Intelligence and Innovation Fund (WTAI).

The Central Claim: Compute Is Concentrating in a Handful of Labs

The conversation's throughline is that an increasing share of the world's newly built AI compute is being absorbed by a small number of frontier AI labs, principally Anthropic and OpenAI, rather than being distributed broadly across the economy. Patel estimated that these two labs alone could account for roughly 40–50% of *incremental* global compute additions next year (in 2027), up from about 30% currently, with frontier lab compute said to be tripling annually against a global compute base that is merely doubling.

If that thesis holds even directionally, it has a straightforward implication for the chip and equipment makers that supply the buildout.

- NVIDIA (5.10%) remains the most direct beneficiary of any scenario where compute demand keeps compounding, given its position supplying the accelerators that make up the bulk of AI training and inference clusters.

- Advanced Micro Devices (2.10%) and Broadcom (3.30%) occupy adjacent positions, with AMD as a second-source GPU and custom silicon supplier, and Broadcom through its custom AI accelerator (ASIC) partnerships with hyperscalers building their own chips, a trend Patel specifically noted OpenAI is pursuing.
- Taiwan Semiconductor Manufacturing Company (3.00%) sits underneath all of this as the foundry actually fabricating the advanced logic these designs depend on; Patel's framing of compute economics ultimately traces back to wafer output at TSMC's leading-edge nodes.

The Memory Bottleneck

A recurring theme is that memory, not just logic, has become a genuine constraint and a source of significant margin expansion for suppliers who were, in Patel's words, "making no money off of high bandwidth memory (HBM)" as recently as a couple of years ago. That repricing shows up across the memory complex.

- Micron Technology (4.60%) and Samsung Electronics (4.50%) are the two largest global suppliers of high-bandwidth memory, with SK Hynix (2.20%) rounding out the oligopoly that effectively prices HBM globally.
- SanDisk (2.60%) and Kioxia Holdings (2.90%) represent the NAND flash side of the storage stack, which benefits from a related but distinct dynamic as data pipelines for AI training scale up.
- Within the WisdomTree strategy, SK Square (2.50%) is a holding company with a stake in SK Hynix that provides indirect exposure to the same memory-pricing story.

Fab Equipment: The "Mirrors" Bottleneck

One of the more specific claims in the conversation concerns lithography equipment, particularly extreme ultraviolet (EUV) tools from ASML (2.00%). Patel argued that the physical constraint on scaling out chip production isn't capital but the supply chain for componentry, and his example was that EUV mirror production (through supplier Carl Zeiss) simply cannot expand fast enough even if funded generously, and that the industry's current planning assumption of roughly 100 EUV tools by 2030 may already be a conservative anchor if AI demand continues on its current trajectory. That bottleneck framing extends to the broader wafer fab equipment (WFE) supply base.

Lam Research (2.40%), Applied Materials (1.70%), KLA Corporation (1.10%), Tokyo Electron (1.30%), Advantest Corporation (1.80%) and Teradyne (1.90%) on the test side. Patel's rough estimate that roughly $6 billion of annual fab capital expenditure can generate a gigawatt of compute capacity capable of producing over $100 billion in AI-related revenue over a multi-year period is a back-of-the-envelope calculation, not a modeled forecast, but it illustrates why this equipment tier has re-rated over the past two years.

Hyperscalers, Neoclouds, and the Economics of Renting Compute

A key economic point in the discussion is the divergence between what compute costs to build (Patel cites roughly $10–15 million per megawatt as a base cost) and what frontier labs can monetize it for, a figure Patel puts as high as $50 million per megawatt currently for Anthropic, with a forecast (again, explicitly speculative) of $70–100 million-plus per megawatt by the end of 2027 if labs keep releasing their best models externally. That spread is what has allowed compute owners without a lab's software stack to extract unusually high rents by simply leasing capacity. Amazon (4.20%), Meta Platforms (3.60%) and Alphabet (3.40%) are repeatedly referenced in the conversation as the hyperscalers historically funding the majority of AI infrastructure buildout, increasingly through debt issuance rather than pure cash flow, a point Patel raised in the context of Meta's recent bond offerings. Oracle (3.00%) has taken a similar posture in aggressively contracting compute capacity to AI labs. Newer "neocloud" entrants without hyperscaler balance sheets, such as CoreWeave (1.30%) and Nebius Group (1.20%), are direct plays on this same rental economics dynamic, monetizing GPU clusters at the premium pricing Patel describes rather than building their own foundation models.

Power and Physical Infrastructure

Patel's estimate that full-stack AI capital expenditure (chips, servers, data centers and power generation together) could approach $10 trillion annually by the end of the decade rests heavily on the point that power infrastructure has a multi-decade asset life and must be contracted years ahead of compute deployment. That has direct relevance for GE Vernova (1.90%), a major supplier of gas turbines and grid equipment, Bloom Energy (1.60%), which supplies fuel-cell power generation increasingly marketed for data center applications, Vertiv Holdings (1.60%), which supplies data center power and thermal management systems, and Cummins (1.10%), a power-generation equipment maker. Cisco Systems (2.00%) and Cloudflare (1.40%) round out the networking and edge infrastructure layer that supports these buildouts.

The Interconnect Layer

Patel's framework implies that as training clusters scale, the bottleneck increasingly shifts to interconnect and data movement rather than raw compute. That is the operating thesis behind optical and interconnect suppliers such as Lumentum Holdings (1.70%), Credo Technology Group (1.40%), Astera Labs (1.20%) and Coherent Corp (1.20%), all of which supply optical transceivers, connectivity chips, or photonics components used to link GPU clusters at scale. Marvell Technology (1.90%) supplies custom silicon and networking chips for similar interconnect and custom accelerator use cases.

China, Export Controls, and the Compute Gap

Patel was notably direct in describing the effect of export controls.

U.S. data center compute deployment has risen to roughly 70% of the global total since 2022, with China's incremental share falling to under 10%, and he projected China's total AI compute could remain under 30 gigawatts by 2028 even accounting for a "hockey stick" acceleration from domestic fabrication capacity at firms like SMIC and CXMT starting around 2027.3 He was careful to caveat that domestically produced Chinese chips are likely to lag Western equivalents in capability, meaning even a large gigawatt figure

would not be directly comparable in compute quality. This is a live policy debate rather than a settled outcome, and export control policy could shift the trajectory in either direction.

The Assembly and Substrate Layer

Behind the chips and the power plants sits a less-discussed but essential tier: the companies that actually assemble servers and package silicon. Patel's framing of AI capex as spanning "critical IT," such as servers, networking and fiber, rather than just accelerators points toward original design manufacturers and packaging specialists such as Wiwynn Corporation (1.20%), which builds AI server racks for hyperscalers, and Delta Electronics (1.00%), a major supplier of power supplies and thermal solutions for data centers.

On the packaging and substrate side, ASE Technology Holding (1.10%), Elite Material Co. (1.00%) and Unimicron Technology (1.00%) supply advanced packaging and substrate materials that have become increasingly load-bearing as chip packaging complexity rises alongside compute density. Tower Semiconductor (1.00%) and Intel (0.90%) round out the specialty and legacy logic manufacturing base, while Texas Instruments (0.90%) supplies the analog and embedded chips that, while less central to Patel's frontier-compute narrative, remain necessary inputs across the broader data center and power infrastructure buildout. DigitalOcean Holdings (1.00%) offers a smaller-scale cloud analogue to the neocloud dynamic described above, aimed at a different customer segment than the hyperscale contracts Patel focuses on.

Software, Security, and the Application Layer

Patel's observation that "the app layer has generated very little value" relative to the model layer and the hardware supply chain is a useful lens for the software names in this space. Snowflake (2.80%) and Datadog (0.90%) sit at the data infrastructure and observability layer that AI workloads increasingly run on top of, benefiting from higher compute utilization even if they are not themselves capturing frontier-model economics. Adobe (1.10%) represents an application-layer incumbent navigating the same value-capture question Patel raised, specifically, whether incumbents monetize AI features directly or see value migrate toward the model providers embedded in their products. Security spend tends to scale with infrastructure buildout regardless of where AI value capture ultimately settles, which is the more durable argument for Palo Alto Networks (3.60%) and CrowdStrike Holdings (1.20%), both providers of the cybersecurity tooling that hyperscalers and enterprises alike require as their AI-driven infrastructure footprint expands. SoftBank Group Corp (1.40%), through its investment arm, has direct exposure to several of the frontier AI labs and chip designers referenced in the conversation, making it something closer to a diversified capital allocator into the exact "arms race" dynamic the two Patels are describing than an operating company within the supply chain. Cerebras (1.20%) represents an alternative compute architecture bet, which is a reminder that the accelerator market Patel frames as an NVIDIA-dominated duopoly-plus-hyperscaler-silicon story is not fully closed to novel entrants.

Conclusion: Following the Continued AI Buildout with WTAI

We have written numerous times about the 'bottlenecks' thesis that drives much of the overall exposure within WTAI as of September 2026. Companies are currently in the process of building and scaling up data centers, and the physical things needed to convert these inputs into tokens are in short supply. As of the

release of this conversation on August 25, 2026, we think the conversation on this podcast offers direct support and clarity for much of that thesis, and this piece is meant to tie a lot of what people might hear in that discussion to positions held in the WisdomTree strategy.

We will continue to update our views as this story that is driving a lot of the world's equity market performance continues to evolve.

1 Source for the AI economics details in this piece, unless otherwise noted: Patel, D. (Host). (2026, March 13). *Dylan Patel — Deep dive on the 3 big bottlenecks to scaling AI compute* [Audio podcast episode]. In *Dwarkesh Podcast*.

2 For each company, there will be a figure in parentheses following, which represents the announced targeted weight from August 13, 2026, with implementation after the close on August 21, 2026, from the most recent rebalancing of the WisdomTree Artificial Intelligence & Innovation Index. **Holdings subject to change.**

3 As of August 24, 2026, the WisdomTree Artificial Intelligence & Innovation Index had 0% exposure to any Chinese company. **Holdings subject to change.**

Important Risks Related to this Article

There are risks associated with investing, including possible loss of principal. The Fund invests in companies primarily involved in the investment theme of Artificial Intelligence (AI) and Innovation. Companies engaged in AI typically face intense competition and potentially rapid product obsolescence. These companies are also heavily dependent on intellectual property rights and may be adversely affected by loss or impairment of those rights. Additionally, AI companies typically invest significant amounts of spending on research and development, and there is no guarantee that the products or services produced by these companies will be successful. Companies that are capitalizing on Innovation and developing technologies to displace older technologies or create new markets may not be successful. The Fund invests in the securities included in, or representative of, its Index regardless of their investment merit and the Fund does not attempt to outperform its Index. The composition of the Index is governed by an Index Committee and the Index may not perform as intended. Please read the Fund's prospectus for specific details regarding the Fund's risk profile.