The AI Factory Supply Chain:
The Evolution of WTAI's May 2026 Rebalance

Christopher Gannatti, CFA, Global Head of Research
Blake Heimann, Senior Associate, Quantitative Research

Key Takeaways

- The May 2026 rebalance of **WisdomTree Artificial Intelligence and Innovation Fund (WTAI)** pushes further into the physical supply chain of AI: semiconductors rise from 36% to 43%, while software falls from 16.6% to 13.1%. The portfolio is not merely buying "hardware"; it is buying the constraints that determine whether AI capacity can be built, powered, connected, tested and used.
- The most important additions—Cerebras, KLA, ASE Technology, Advantest and Monolithic Power Systems—point to a new phase of the AI cycle: wafer-scale inference, process control, advanced packaging, optical test and power delivery. In this phase, yield, latency and watts per rack can matter as much as model quality.
- The rebalance also broadens the definition of AI infrastructure beyond chips. CoreWeave adds neocloud exposure, EMCOR adds data center electrical infrastructure, and upweights to GE Vernova, Bloom Energy, and Cummins deepen power infrastructure exposure. Corning, Lumentum and Coherent reflect a simple idea: there are no dark GPUs, and there are no useful GPUs that cannot communicate via networking.

From Bottlenecks to the Bottleneck Stack

Artificial intelligence (AI) investing keeps changing shape. The first clean story was model quality. Then it was graphics processing units. Then it became high-bandwidth memory. Then it became networking. Then it became data center power. The market keeps trying to compress AI into a single bottleneck, but the infrastructure reality keeps pushing in the other direction:

For the Fund's current holdings, click here: WTAI.

AI is a stack of bottlenecks.

That is the most important distinction in this rebalance. The February 2026 WTAI rebalance emphasized the physical constraints of AI deployment—compute, memory, networking and power. The May 2026 rebalance keeps that same philosophy but gets more granular. It asks not only whether the world needs more AI infrastructure, but which parts of the infrastructure chain are becoming harder to source, harder to engineer and harder to replicate.

Said differently, the portfolio is moving from "AI infrastructure" as a category to the AI factory supply chain as an investment map. The AI factory needs accelerators, memory, advanced packaging, optical interconnects, test equipment, grid access, on-site power, cooling, cybersecurity and merchant cloud capacity. A weakness in any one of those layers can delay revenue for the entire system.

Allocating Toward What Can Block Deployment

From a high level, current holdings fall from 64 to 57, with:

- 7 additions
- 14 deletions
- 17 upweights
- 8 downweights

The largest category-level message is clear: semiconductors gained almost seven percentage points. Software, innovation and other hardware were all reduced at the margin.

Figure 1: Evolution of Weight within Four Primary Exposure Groupings

Classification	Before #	Before Weight	After #	After Weight	Weight Change
Semiconductors	19	36.00%	24	42.00%	6.00%
Software	18	16.60%	11	13.10%	-3.60%
Other Hardware	19	26.30%	16	25.90%	-0.40%
Innovation	8	21.10%	6	19.00%	-2.00%
Total	64	100.00%	57	100.00%	0.00%

Sources: WisdomTree, FactSet, Bloomberg and individual company data. Data is as of the quarterly, May 2026 rebalance process for the WisdomTree Artificial Intelligence & Innovation Index. **Holdings subject to change.**

Importantly, semiconductors are not a monolith even if in Figure 1 they are a single line. The incremental weight is going into compute, memory, process control, packaging, test and power semiconductors. Other hardware is being shifted away from lower-conviction servers or autonomy

exposure and shifted toward optics, grid infrastructure and power. Software is being tightened toward cybersecurity, data infrastructure and neocloud exposure, not abandoned.

This is equivalent to moving from the obvious beneficiaries of AI to the suppliers of the suppliers of AI.

The AI Spend Sequence

As a committee, we were looking through the sorts of things that are needed if one is seeking to build a large data center. In 2026, at least so far, a lot of uncertainty has concentrated on software, but any company that has clearly illustrated its importance to any related AI data center supply chain has seen a totally different investment environment.

- The modern data center is nothing without AI accelerators, whether for training or inference workloads. NVIDIA remains the largest target weight in the rebalance at 5%, and AMD is upweighted as inference and agentic workloads increase the importance of CPU-GPU system balance. The importance of CPUs being widely discussed is new, and it's not lost on us how Intel has been great at CPUs all along and the world is now focusing back towards this area.
- Every accelerator needs memory. Samsung, Micron, SK Hynix, SanDisk and Kioxia remain central holdings, with Kioxia upweighted as NAND and high-bandwidth flash become more relevant to inference economics. We expect to continue to see interesting innovations within memory. This does not mean the demand for HBM suddenly drops, but we do understand the supply constraints paired with the memory hierarchy are inspiring a lot of entrepreneurs to try out different approaches to the so-called memory wall. Cerebras, in our view, represents one such innovative stance on this question.
- Advanced packaging and test are becoming capacity gates. KLA, ASE Technology and Advantest are added because the industry is no longer only asking, "Can we design better chips?" It is asking, "Can we manufacture, package and test them at high enough yield?" One of the notable findings is that as the systems get more and more complex and transition from using copper networking to optical network, for example, a greater and greater premium is being placed on ensuring that the installed semiconductors are up to standard. The cost of installing a so-called 'bum chip' has risen, simply put.
- Lumentum, Coherent and Corning are higher-conviction exposures because optical connectivity becomes more important as AI factories stretch beyond the practical limits of copper. Now, this isn't just us making this statement, as we have seen over the prior months, Nvidia has made investments in each of these firms. Nvidia is sending us a clear signal on the importance of securing supply for optical networking.
- GE Vernova, Bloom Energy, Cummins, Delta Electronics, Texas Instruments, Monolithic Power Systems and EMCOR represent the widening power stack—from the voltage regulator to the

rack, from the generator to the grid and electrical interconnects. One of the things that we force ourselves to remember is that the power envelope is tightly controlled as it traverses the data center and these companies are all performing different functions in that regard.

The result is not a single bet on a single AI supplier. It is an attempt to own the layers that make AI deployment possible.

Figure 2: New Additions – Where the Portfolio Added New Optionality

Ticker	Name	Subclass	New Weight	High Level Rationale
CBRS US	Cerebras	Compute Semi	1.00%	Wafer scale computing solution optimized for inference with SRAM
KLAC US	KLA Corporation	Semicap	1.00%	Increased complexity at leading nodes requires more process control/yield optimization; TSMC capex revision upward
CRWV US	CoreWeave, Inc. Class A	Neo Cloud	1.00%	GPU rental pricing strong in compute constrained environment; recent META deal and Semi Analysis platinum operator
3711 TT	ASE Technology Holding Co., Ltd.	Semicap	1.00%	Advanced packaging demand beneficiary as TSMC CoWoS capacity constrained; 23x forward earnings with 40% growth
6857 JP	Advantest Corp.	Semicap	0.80%	Optical Test for silicon photonics/co-packaged optics for AI datacenter
MPWR US	Monolithic Power Systems, Inc.	Analog / Power Semi	0.80%	Vertical Power Delivery for GPUs in AI Server Racks
EME US	EMCOR Group, Inc.	Thermal / DC Infra	0.80%	Datacenter electrical infrastructure and connectivity

Sources: WisdomTree, FactSet, Bloomberg and individual company data. Data is as of the quarterly, May 2026 rebalance process for the WisdomTree Artificial Intelligence and Innovation Index. **SRAM:** A fast, expensive type of computer memory that doesn't need to be constantly refreshed to hold data. **Semicap:** Short for semiconductor capital equipment; the machines that manufacture chips. **CoWoS**: TSMC's advanced chip-packaging technology that stacks multiple chips closely together. Vertical Power Delivery: Routing electricity to a chip underneath a processor on the opposite side of the motherboard to send power straight up, cutting energy loss. **Holdings subject to change.**

What Is Wafer-Scale Inference?

Cerebras is the clearest example of the portfolio trying to capture a different path through the compute bottleneck. Most AI infrastructure has been built around systems of many separate chips connected through high-speed networking. Cerebras takes a different approach, specifically a wafer-scale processor that attempts to keep more computation and memory closer together on one enormous piece of silicon.

Why does that matter?

In AI inference, the question is often not only “How much compute exists?” but “How quickly can data move?” If a model is constantly waiting for memory to transfer data, then raw theoretical compute can sit idle. SRAM—memory that sits on the processing element—can reduce some of the distance data needs to travel. The economic question becomes whether the architecture can deliver lower latency or better utilization for certain inference workloads.

This is the important conceptual point, and the addition is not a call that every model will run on wafer-scale systems. It is an acknowledgment that inference is becoming sufficiently large, diverse and latency-sensitive that alternative architectures deserve index exposure. In a world where every hyperscaler wants more tokens, every credible approach to reducing bottlenecks matters. In our view, it’s an interesting innovation dealing with how data and compute are interconnected, and it won’t be the last new approach that we’ll see. There is so much focus on this area, and many companies are seeking to explore new ways to approach the so-called ‘memory wall.’

Process Control, Packaging and Test: Yield Becomes Supply

KLA, ASE Technology and Advantest are not consumer-facing AI names. That is exactly why they are interesting. The more advanced the semiconductor stack becomes, the more the economics shift toward companies that can help the industry manufacture complexity reliably.

At older nodes, one could still think in relatively simple terms, meaning something akin to ‘smaller transistors, better chips’. At leading-edge nodes and in advanced packaging, that is too simplistic. The AI accelerator is increasingly a system of chips, memory, interposers, packaging technologies, power delivery layers and networking components. A defect in one piece can lower yield, delay supply and raise cost.

That is why process control and test move closer to the center of the thesis. KLA sits in inspection and metrology, the tools that help foundries identify defects, improve yields and manage complexity. ASE adds outsourced semiconductor assembly and advanced packaging exposure. Advantest adds test exposure, including relevance to silicon photonics and co-packaged optics. The more difficult the chip gets, the more valuable the equipment that makes the chip manufacturable becomes.

The phrase “capacity constraint” can sound abstract. In semiconductors, it often shows up as yield. If a wafer or package cannot be manufactured at scale, supply does not exist in economic terms. The May rebalance adds more exposure to that practical reality.

Figure 3: Subclass Changes – Where the Weight Moved

Sub-Classification	Before Weight	After Weight	Change
Compute Semi	6.8%	9.9%	**3.1%**
Semicap	4.4%	7.2%	**2.8%**
Optics	4.6%	6.7%	**2.1%**
Power / Utilities	4.4%	5.8%	**1.3%**
Analog / Power Semi	1.4%	2.7%	**1.3%**
Thermal / DC Infra	2.1%	2.9%	**0.8%**
Compute Semi	4.5%	5.1%	**0.6%**
Data Infra	3.4%	2.5%	**-0.9%**
Robotics	2.7%	1.8%	**-0.9%**
Memory Semi	17.8%	16.6%	**-1.2%**
Enterprise Software	3.6%	2.3%	**-1.3%**
MLOps / Automation	1.4%	0.0%	**-1.4%**
Servers	1.5%	0.0%	**-1.5%**
Other	5.9%	4.3%	**-1.7%**
Autonomy	3.8%	1.0%	**-2.8%**

Sources: WisdomTree, FactSet, Bloomberg and individual company data. Data is as of the quarterly, May 2026 rebalance process for the WisdomTree Artificial Intelligence & Innovation Index. Subcategories with absolute weight change > 50bps shown. Semicap: Companies supplying the equipment, tools and process technology needed to manufacture increasingly advanced AI semiconductors; Compute Semi: Designers or producers of processors, accelerators and related chips that supply the core compute for AI workloads; Optics: Companies enabling high-speed data movement through lasers, optical components, fiber and photonic networking technologies; Power / Utilities: Electricity providers and grid-exposed companies positioned to support rising data center power demand; Analog / Power Semi: Semiconductor companies focused on power management, signal conversion and control chips essential for efficient AI systems; Thermal / DC Infra: Providers of cooling, electrical, mechanical and physical infrastructure used to operate high-density AI data centers; Memory Semi: Semiconductor companies producing DRAM, HBM, NAND or storage technologies critical for training and inference workloads; Autonomy: Companies developing self-driving, intelligent mobility or machine-perception systems that apply AI to physical-world decision-making; Servers: Companies assembling, integrating or supplying AI servers, racks and systems used inside accelerated computing clusters; MLOps / Automation: Software companies helping enterprises build, deploy, monitor, automate and govern machine-learning models or AI workflows; Enterprise Software: Application and platform software companies embedding AI into business processes, productivity tools, security or data systems; Robotics: Companies applying AI, sensors and automation to machines that perform tasks in industrial or physical environments. **Holdings subject to change.**

Copper, Glass and the Grid

The portfolio continues to treat optics as a central AI infrastructure pressure point. Lumentum and Coherent are upweighted; Corning is upweighted; Advantest is added in part because silicon photonics and co-packaged optics require more sophisticated test infrastructure. The reason is not difficult to understand. AI factories are not useful if the GPUs are isolated. The larger the cluster, the more the bottleneck shifts from compute to communication.

Copper is not disappearing, but as distance, bandwidth and power-efficiency requirements rise, photons gain share. The language changes from "chips" to lasers, fiber, photonics, optical circuit switching and co-packaged optics. In this rebalance, optics become less of a side exposure and more of an enabling layer for scale.

Power is a parallel story. GE Vernova rose from 0.94% to 2.00%. Bloom Energy rises from 0.69% to 1.50%. Cummins rose from 0.85% to 1.25%. Delta Electronics, Texas Instruments and Monolithic Power Systems point to power management inside the rack and around the accelerator. EMCOR brings exposure to the electrical infrastructure required to connect new data centers to the grid.

This is the part of the AI cycle that looks least like software and most like industrial history. The limiting question is no longer only whether a model can improve. It is whether the power, thermal and electrical infrastructure can arrive before demand moves elsewhere.

There Are Still No Dark GPUs

The phrase remains useful because it forces discipline. If an accelerator is unpowered, it produces nothing. If it is poorly networked, utilization suffers. If memory cannot feed it, theoretical Floating-Point-Operations-Per-Second (FLOPS) do not become realized throughput. If the data center cannot be permitted, connected or cooled, none of the silicon matters.

That is why the May rebalance makes power a multi-layer exposure rather than a single utility or industrial allocation. GE Vernova and Cummins represent generation and backup power. Bloom Energy represents on-site power via fuel cells. EMCOR represents grid connectivity and electrical infrastructure. Delta, Texas Instruments and Monolithic Power Systems represent the conversion, regulation and efficiency problem closer to the rack. This is not simply "more electricity." It is the power architecture of AI deployment.

The practical implication is that the AI trade has become less ethereal. In the early days, investors could debate whether a chatbot would disrupt a workflow. Today, they can count racks, watts, backlog, interconnects, memory, packaging capacity and orders. That does not eliminate cyclicality or valuation risk, but it does anchor more of the thesis in physical deployment.

Neoclouds: Merchant Capacity Becomes a Signal

CoreWeave is added at 1.00%, while DigitalOcean is trimmed from 3.01% to 2.00%. The combined message is subtle but important. The portfolio is not exiting cloud infrastructure. It is differentiating between general cloud exposure and specialized AI capacity exposure.

Neoclouds exist because hyperscaler capacity is not infinitely elastic. If frontier labs and enterprises need GPUs, power, networking and operational expertise now, specialized providers can become a pressure-release valve. CoreWeave brings exposure to that merchant capacity model. DigitalOcean remains but is trimmed as the rebalance adds a more directly AI-scaled platform.

This can be thought of as a capacity-rental thesis. When compute is abundant, rental economics compress. When compute is scarce, the providers with powered sites, GPUs, networking and customer commitments can monetize the shortage. The addition of CoreWeave acknowledges that AI infrastructure does not have to be owned only by the largest hyperscalers to be economically relevant.

Hyperscalers: Trimmed, Not Dismissed

Amazon and Alphabet are both modestly downweighted, while Oracle is held at roughly 3% and Meta remains close to 3.5%. That might look counterintuitive if one only looks at cloud growth, custom silicon or AI product adoption. But the rebalance is not saying the hyperscalers are less important. It is saying the supply chain under the hyperscalers is becoming more investable.

Amazon remains a critical AI infrastructure company through AWS, Trainium, Graviton and Nitro. Alphabet remains central through TPUs, Google Cloud, Gemini and the broader full-stack AI strategy. Oracle remains meaningful because of its AI infrastructure and OpenAI/Stargate-related exposure. The portfolio simply reallocates some weight from platform owners to the components and infrastructure that may benefit regardless of which platform wins the next contract.

This is a useful framing for investors:

If the hyperscalers are the demand signal, the rebalance is increasingly focused on the supply response.

Software: Tightened, Not Abandoned

Software falls from 16.6% to 13.1%. That is a meaningful reduction, but it is not a philosophical exit. Palo Alto Networks is upweighted. CrowdStrike remains. Snowflake and other data-layer names remain. The difference is that the portfolio is less willing to pay for broad sentiment-driven software exposure when the near-term monetization path is less visible.

The deletes tell the story. JFrog and GitLab are removed from MLOps/automation. SentinelOne and Zscaler are removed as the cyber allocation is consolidated into higher-conviction leaders. Shopify is

removed as agentic payments and commerce dynamics become more ambiguous. Palantir is removed, explicitly flagging valuation risk if revenue decelerates. IBM is removed to reallocate toward purer AI infrastructure exposures.

The portfolio still recognizes that software ultimately captures value when AI is deployed into workflows. But the timing is harder to underwrite. The physical layers have more visible order books, capital spending signals and bottlenecks. The software payoff may come later, after infrastructure catches up.

Figure 4: Deletes, Pruning the Less Direct AI Transmission Mechanisms

Name	Classification	Subclass	Rationale
MongoDB, Inc. Class A	Software	Data Infra	Increasing competition from private firms and alternative data infrastructure
International Business Machines	Software	Enterprise Software	Reallocate to purer AI infra
Hyundai Motor Company	Other Hardware	Autonomy	Boston dynamics key catalyst has occurred; deployments in 2028
Celestica Inc.	Other Hardware	Servers	Low margin commoditized/competitive AI server business
Schneider Electric SE	Other Hardware	Power / Utilities	Reallocate w/ purer DC power infra play
JFrog Ltd.	Software	MLOps / Automation	Increasing competition from private firms and alternative data infrastructure
Sompo Holdings, Inc.	Innovation	Investor	AI operational efficiencies and PLTR exposure reduction
Shopify, Inc. Class A	Software	Enterprise Software	Sentiment towards agentic payments became headwind rather than tailwind
Deere & Company	Other Hardware	Robotics	Agricultural autonomy becoming secondary to macro and agricultural cycle
SentinelOne, Inc. Class A	Software	Cyber	Consolidate cyber exposure in leaders PANW, CRWD
Kroger Co.	Innovation	NA	Proprietary data advantage realized over long period
Palantir Technologies Inc. Class A	Software	Enterprise Software	High-valuation software trading at 35x revenue ripe for multiple contraction as revenue growth decelerates
Zscaler, Inc.	Software	Cyber	Consolidate cyber in leaders PANW, CRWD
Gitlab, Inc. Class A	Software	MLOps / Automation	Increased competition across public/private peers and model labs

Sources: WisdomTree, FactSet, Bloomberg and individual company data. Data is as of the quarterly, May 2026 rebalance process for the WisdomTree Artificial Intelligence & Innovation Index. **Holdings subject to change.**

The 2026 Horizon: What the Rebalance Is Really Saying

It might be easy or tempting to say that WTAI is becoming more semiconductor heavy. However, we would indicate that the more precise framing is that WTAI is becoming more bottleneck heavy. We think the distinction matters.

A semiconductor overweight alone could simply be momentum exposure. This rebalance is more specific. It is adding process control, packaging, optical test, wafer-scale inference, analog and power semiconductors, grid infrastructure and neocloud capacity. These are not all the same business, but they rhyme. Each one is a potential constraint between AI demand and AI deployment.

The market may continue to debate whether AI application software will monetize fast enough. That debate is real. But while it continues, the physical buildout is creating measurable demand for chips, memory, packaging, power equipment, optical components and specialized cloud capacity. In the May 2026 rebalance, WTAI is leaning into the places where imagination has to become infrastructure.

The next phase of AI may not be won by the company with the best demo. It may be won by the ecosystem that can manufacture the most reliable compute, attach the most memory, move the most data, deliver the most power and bring capacity online fastest.

The Concluding Summary

A simple way to explain the rebalance is this:

- WTAI is moving deeper into the AI factory supply chain.
- The largest increase is in semiconductors, but the real story is broader: packaging, process control, optics, power and neocloud capacity.
- Software remains in the index, but the exposure is more selective and more connected to security, data infrastructure and AI deployment rather than broad SaaS multiple expansion.
- The portfolio aims to follow bottlenecks rather than headlines. If AI demand remains strong, the companies that solve the bottlenecks may have more durable revenue visibility than the companies simply promising productivity gains.

IMPORTANT INFORMATION

Please see the WisdomTree Glossary for definition of terms.

Investors should carefully consider the investment objectives, risks, charges and expenses of the Fund before investing. For a prospectus or, if available, the summary prospectus containing this and other important information about the fund, call 866.909.9473 or visit WisdomTree.com/investments. Read the prospectus or, if available, the summary prospectus carefully before investing.

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.

You cannot invest directly in an index.

Additional information is available at WisdomTree.com/investments.

Christopher Gannatti is a registered representative of Foreside Fund Services, LLC.

WisdomTree Funds are distributed by Foreside Fund Services, LLC.