

EDUCATIONAL INVESTMENT RESEARCH

High-Growth AI & Tech Equity Portfolio

An educational, hypothetical US\$1,000,000 model across the artificial-intelligence supply chain

PORTFOLIO SIZE	TYPE	DATE
US\$1,000,000	Hypothetical / illustrative	June 2026

EDUCATIONAL & ILLUSTRATIVE ONLY — NOT FINANCIAL ADVICE. This is a hypothetical model portfolio created to explain how the AI supply chain fits together. It is not a recommendation to buy or sell any security and is not tailored to anyone's personal circumstances. See the full legal disclaimer at the end of this document.

About this model

All figures are in US dollars (USD). The portfolio is a hypothetical US\$1,000,000 allocation used purely to illustrate how capital could, in principle, be spread across the layers of the artificial-intelligence supply chain. Percentages and dollar amounts are illustrative and do not account for transaction costs, taxes, or any individual's objectives, time horizon or risk tolerance.

1. The hypothetical allocation

The model is organised into three risk tiers. Roughly 64% sits in lower-volatility core anchors, 22% in core-growth/cyclical names, and 14% in higher-risk satellites. The weightings below are **illustrative** — they show one way a high-growth AI book could be structured, not a recommendation to anyone.

Tier 1 — Core buy-and-hold anchors

Ticker	Company	Segment	Weight	USD
NVDA	NVIDIA	AI compute / GPUs	13%	\$130,000
MSFT	Microsoft	Cloud / AI platform	8%	\$80,000
GOOGL	Alphabet	Cloud / AI platform	9%	\$90,000
AMZN	Amazon	Cloud / commerce	7%	\$70,000
AVGO	Broadcom	Custom AI ASICs / networking	6%	\$60,000
TSM	Taiwan Semiconductor	Foundry	6%	\$60,000
ASML	ASML	EUV lithography tools	5%	\$50,000
GEV	GE Vernova	Power generation / grid	5%	\$50,000
CEG	Constellation Energy	Nuclear generation	5%	\$50,000
Subtotal			64%	\$640,000

Tier 2 — Core growth / cyclical

Ticker	Company	Segment	Weight	USD
AMD	Advanced Micro Devices	AI compute / GPUs (stretched)	5%	\$50,000
VST	Vistra	Nuclear + gas power (value)	5%	\$50,000
MU	Micron	HBM / memory (cyclical)	4%	\$40,000
CRWD	CrowdStrike	AI-native security (stretched)	4%	\$40,000
CCJ	Cameco	Uranium / nuclear fuel	4%	\$40,000
Subtotal			22%	\$220,000

Tier 3 — Higher-risk satellites

Ticker	Company	Segment	Weight	USD
VRT	Vertiv	Data-centre power / cooling	4%	\$40,000
ORCL	Oracle	AI cloud infrastructure	4%	\$40,000
TRGP	Targa Resources	Natural-gas midstream	3%	\$30,000
IONQ	IonQ	Quantum computing (speculative)	3%	\$30,000
Subtotal			14%	\$140,000

Grand total: US\$1,000,000 — allocated across 18 holdings (hypothetical).

2. The AI supply chain — who owns which link

AI is not one industry; it is a stack of dependent layers. Money flows from the cloud giants downward into chips and power, and value concentrates wherever a single company is hard to replace. The model deliberately includes a position in every layer. The third column gives the technical picture; the fourth translates it into plain English — what it actually means and why it matters.

Layer (top → bottom)	Holdings	The detail — market control / moat	In plain English — what it means & why it matters
1. Design tools & lithography	ASML	ASML is a true monopoly in extreme-ultraviolet (EUV) lithography — no advanced AI chip below ~5nm can be made without its machines. Effectively single-source.	ASML makes the only 'printing presses' capable of etching the world's most advanced chips. If you want to build a cutting-edge AI chip, you must buy from ASML — there is no rival. That gives it huge pricing power and makes it almost impossible to disrupt: it sits at the very bottom of the chain, so everyone above it depends on it.
2. Chip design — GPUs	NVDA, AMD	NVIDIA holds an estimated ~80–90% of the AI training-GPU market and owns the CUDA software moat. AMD is the credible #2 (MI300/MI450) and the main alternative hyperscalers fund to avoid sole dependence on NVIDIA.	GPUs are the 'brains' that do the actual AI maths. NVIDIA makes the best ones AND owns the software (CUDA) that developers have spent years learning — so switching away is painful and costly, which is why it keeps ~80–90% of the market. AMD is the only serious alternative, which big tech deliberately backs so it isn't 100% dependent on one supplier. Owning both means you have exposure whichever way that contest goes.
3. Custom silicon & networking	AVGO	Broadcom co-designs custom AI ASICs for hyperscalers (Google TPUs, Meta, OpenAI) and dominates high-end networking silicon — the connective tissue between GPU clusters. A near-duopoly with Marvell.	The tech giants increasingly want their own bespoke AI chips rather than buying everything off-the-shelf from NVIDIA. Broadcom is the partner that helps them design those custom chips, and it also makes the high-speed 'plumbing' that lets thousands of chips talk to each other inside a data centre. Few companies can do either job — so Broadcom gets paid as customers try to reduce their NVIDIA reliance.
4. Memory (HBM)	MU	High-bandwidth memory (HBM) is an oligopoly of three — Micron, SK Hynix, Samsung. Micron's HBM capacity has been reported sold out well into the future; memory has repriced as a strategic asset.	AI chips need ultra-fast memory sitting right beside them to feed them data — that's HBM. Only three companies on earth can make it, so it's a tight club with strong pricing. Demand is so high that Micron has effectively pre-sold its output for years, which is great for revenue — but memory is famously cyclical, so the shares swing hard in both directions.
5. Manufacturing (foundry)	TSM	TSMC manufactures the leading-edge chips for NVIDIA, AMD, Apple and others — an estimated ~65%+ of the global foundry market and effectively the only volume source for the most advanced nodes.	NVIDIA, AMD and Apple design chips but don't actually build them — they hand the blueprints to a 'foundry' to manufacture. TSMC is by far the best and biggest, and for the most advanced chips it's essentially the only option at scale. So almost every AI chip on earth passes through TSMC's factories. That's its strength — and a risk, because those factories are in Taiwan.
6. Cloud platforms (where AI is sold)	MSFT, GOOGL, AMZN, ORCL	An oligopoly: AWS ~30%, Azure ~20%, Google Cloud ~13% of cloud infrastructure. Oracle is the fast-growing challenger via the OpenAI/Stargate build-out. This is the layer that monetises AI to end customers.	This is where AI is actually rented out and sold to businesses — you don't buy a supercomputer, you rent it by the hour from Amazon, Microsoft or Google's cloud. These three control most of the market, and it's hard to dislodge them because customers' data and software already live there. Oracle is the fast-rising challenger riding the giant OpenAI build-out. This is the layer that turns all the AI spending into paying customers.

7. Security	CRWD	CrowdStrike is a leader in AI-native endpoint and cloud security — a fragmented but consolidating market where it is one of the clear platform winners.	More AI and more data in the cloud means a bigger target for hackers. CrowdStrike is one of the leaders that protects all of it, using AI to spot attacks in real time. Once a company installs it across thousands of devices, ripping it out is disruptive — so customers tend to stay and spend more each year. That 'stickiness' is what makes it valuable.
8. Power generation & grid	GEV, CEG, VST	No monopoly, but scarce assets: GE Vernova is one of a handful of global suppliers of gas turbines and grid equipment; Constellation owns the largest US nuclear fleet; Vistra is a top independent nuclear/gas generator. AI data centres are power-constrained, making these chokepoints.	AI data centres are astonishingly hungry for electricity — a single campus can use as much power as a small city. The problem is there isn't enough to go around, so whoever can supply reliable power holds the leverage. GE Vernova makes the equipment (turbines, grid gear) few others can; Constellation runs America's biggest nuclear fleet; Vistra is a major independent generator. Power is now the real bottleneck for AI growth, which is why this sleeve is large.
9. Fuel & power infrastructure	CCJ, TRGP, VRT	Cameco is a top-two Western uranium producer (+ Westinghouse stake); Targa is a major US gas-midstream operator; Vertiv is a leading data-centre power & thermal-management supplier and NVIDIA partner.	Behind the power plants sit the fuel and the kit. Cameco digs up the uranium that fuels nuclear plants (and a shortage is looming). Targa moves the natural gas that fires many power stations. Vertiv makes the specialised cooling and power gear that stops racks of red-hot AI chips from melting. These are the unglamorous 'pick-and-shovel' suppliers that get paid no matter which AI company ultimately wins.
10. Frontier compute (optionality)	IONQ	Quantum is pre-commercial. IonQ is a pure-play leader in trapped-ion systems — a lottery ticket on a market that barely exists yet but could be transformative.	Quantum computers are a completely different, experimental kind of machine that could one day crack problems today's computers can't touch. The technology mostly isn't ready for real use yet and the market is tiny. IonQ is one of the front-runners. This is a small, high-risk allocation: if quantum inflects this decade the payoff could be enormous; if it doesn't, the position is kept small enough not to hurt.

Reading the chain: a dollar of AI spending starts at the top (a hyperscaler's budget), flows down through cloud → GPUs/ASICs → memory → foundry → tools, and sideways into the power and security needed to run it all. The strongest monopolies — ASML, NVIDIA's CUDA, TSMC's leading-edge node — are where pricing power is most durable, which is why they carry core weightings.

3. Master comparison table

Analyst consensus, average 12-month price target, and the high/low target range for every holding. *Upside/downside is relative to early-June 2026 prices. Where a figure is shown in red, the consensus average sits below the current price (implied downside). These are third-party analyst estimates, not forecasts by the author.*

Ticker	Company	Segment	Tier	Consensus	Avg target (12m)	High target	Low target
NVDA	NVIDIA	Compute / GPUs	Core	Strong Buy	+36% to +40%	+92% (\$425)	-17% (\$180)
MSFT	Microsoft	Cloud / platform	Core	Strong Buy	~+27% (\$561)	+97% (\$870)	-9% (\$400)
GOOGL	Alphabet	Cloud / platform	Core	Strong Buy	~+19% (\$430)	+43% (\$515)	-5% (\$340)
AMZN	Amazon	Cloud / commerce	Core	Strong Buy	~+18% (\$313)	+40% (\$370)	-22% (\$207)
AVGO	Broadcom	Custom ASIC / net	Core	Strong Buy	~+1% to +18%	+53% (\$630)	-24% (\$314)
TSM	Taiwan Semi	Foundry	Core	Strong Buy	~+6% (\$468)	+36% (\$600)	-20% (\$354)
ASML	ASML	Litho tools	Core	Strong Buy	-2% to +4%	+28% (\$2,200)	-45% (\$900)
GEV	GE Vernova	Power / grid	Core	Buy / Strong Buy	+12% to +25%	+47% (\$1,424)	-14% (\$836)
CEG	Constellation	Nuclear power	Core	Buy / Strong Buy	+26% to +32%	+50%+ (\$462)	-5% (\$272)
AMD	AMD	Compute / GPUs	Cyclical	Strong Buy	~-10% (\$478)	+25% (\$665)	-58% (\$225)
VST	Vistra	Nuclear + gas	Cyclical	Strong Buy	+40% to +60%	+105% (\$320)	-30% (\$97)
MU	Micron	HBM / memory	Cyclical	Strong Buy	-27% to -31%	+64% (\$1,750)	-77% (\$249)
CRWD	CrowdStrike	Security	Cyclical	Buy / Strong Buy	-24% to -28%	+7% (\$800)	-49% (\$368)
CCJ	Cameco	Uranium	Cyclical	Buy	+14% to +28%	+50% (\$175)	-29% (\$83)
VRT	Vertiv	DC power / cooling	Satellite	Strong Buy / Buy	~+8% to +22%	+60%+ (\$488)	wide
ORCL	Oracle	AI cloud	Satellite	Buy	+1% to +7% (\$247)	+64% (\$400)	-36% (\$155)
TRGP	Targa	Gas midstream	Satellite	Buy	flat to +12%	+26% (\$331)	-20% (\$188)
IONQ	IonQ	Quantum	Satellite	Strong Buy (spec.)	-1% to +10%	+47% (\$100)	-34% to -49%

Note: ranges reflect genuine disagreement between data providers (TipRanks, MarketBeat, Yahoo Finance, Investing.com, GuruFocus, Barchart) and the speed at which several of these stocks have moved. AVGO's spread is wide because of conflicting post-earnings price feeds in early June.

4. The smart money — who actually holds these (Q1 2026 13F filings)

Based on the most recent quarterly 13F disclosures (positions as of 31 March 2026). 13F filings are backward-looking and only show US-listed long positions — they are a clue to conviction, not a recommendation. Bridgewater runs a highly diversified, quant-driven book that holds small slices of most mega-caps, so its 'Yes (broad)' entries are not high-conviction signals.

Stock	Berkshire (Abel/Buffett)	Tepper (Appaloosa)	Ackman (Pershing)	Wood (ARK)	Dalio (Bridgewater)	What the signal says
NVDA	No	Trimmed	No	Buying hard	Small (index-like)	ARK aggressively adding; ~\$5T cap
MSFT	No	Cut ~82%	New buy	Light	Yes (broad)	Ackman initiated; Tepper took profits
GOOGL	Top-5 buy (tripled)	Added +29%	Sold out	Light	Yes (broad)	Berkshire's signature 2026 bet
AMZN	Exited	Largest pos. (~15%)	Added +19%	Light	Yes (broad)	Tepper & Ackman's top conviction trade
AVGO	No	No	No	Some	Yes (broad)	Institutional Strong Buy; Druckenmiller added
TSM	No	Added (~7.6%)	No	Trimmed	Yes (broad)	Tepper core; ARK trimming
ASML	No	No	No	No	Yes (broad)	Held mainly via index / quant books
GEV	No	No	No	No	Some	Popular with power-thematic & long-only funds
CEG	No	No	No	No	Some	Core nuclear-AI holding for utility/thematic funds
AMD	No	No	No	Selling	Some	ARK exiting into strength; OpenAI is a ~10% holder
VST	No	Doubled	No	No	Some	Tepper's clearest AI-power conviction add
MU	No	#2 pos. (tripled)	No	No	Yes (broad)	Tepper's highest-octane semiconductor bet
CRWD	No	No	No	Some	Some	Growth/tech funds; richly valued
CCJ	No	No	No	No	Some	Held by uranium/energy-transition specialists
VRT	No	No	No	Some	Some	Favoured by AI-infrastructure thematic funds
ORCL	No	No	No	No	Yes (broad)	Backlog story; less classic-value ownership
TRGP	No	No	No	No	Some	Income/energy funds; not a growth-fund favourite
IONQ	No	No	No	Held (ARK)	No	ARK's quantum optionality; otherwise retail-driven

Key takeaways: (1) The most-owned name across these funds is **Amazon** — Tepper's largest position and a fresh Ackman add. (2) **Alphabet** splits the room: Berkshire tripled it into a top-5 holding while Ackman, Druckenmiller and Altimeter sold out. (3) **Nvidia** is being aggressively bought by ARK but, notably, is not a classic value-fund holding. (4) Tepper's book is the closest match to this model's thesis — Amazon, Micron, Alphabet, TSMC and a doubled Vistra stake.

5. The modern kingpins' lens — would they approve?

Eight of the most influential investors active in the market today, run through this model. The table is an **interpretive read** of how each one's documented approach — and, where known, their most recent disclosed positioning — would likely treat a portfolio like this. It is a guide to how the smart money is thinking, *not a personalised recommendation*.

Investor & style	Likely stance on this book
Warren Buffett / Greg Abel — Berkshire Hathaway	Berkshire under Greg Abel tripled Alphabet into a top-5 holding and exited Amazon in Q1 2026 — the lens is durable, moaty compounders bought when reasonably priced. Would likely bless GOOGL, MSFT and TSM; would avoid IONQ and probably the above-target names (MU, AMD) on valuation. Verdict: APPROVES the anchors, avoids the speculative tail.
Stanley Druckenmiller — Duquesne Family Office	Famously sold Nvidia (and later called the early exit a 'big mistake'); in Q1 2026 he bought Broadcom, Intel and Arm instead — a bet that custom silicon and inference, not just GPUs, drive the next leg. Endorses the AVGO / broad-infrastructure tilt here, but would question paying up for NVDA at today's levels. Verdict: APPROVES the infrastructure spread, picky on price.
Philippe Laffont — Coatue Management	One of the most influential technology investors alive; Coatue lives in the AI trade and has held both Nvidia and Broadcom. Would be comfortable with the chip-and-cloud core (NVDA, AVGO, MSFT, GOOGL) and the growth tilt generally, while actively trading around the volatility. Verdict: APPROVES — this is his neighbourhood.
Seth Klarman — Baupost Group	The modern standard-bearer for deep value and margin of safety; happy to hold cash when bargains are scarce. At forward multiples of 30–170x he would find little safety here and would wait for a sell-off, perhaps tolerating only Alphabet and the cash-generative energy names. Verdict: mostly WAIT / AVOID at these prices.
Howard Marks — Oaktree Capital	Famous for asking 'where are we in the cycle?' and for warning when euphoria runs hot. Wouldn't deny AI's importance, but would flag the crowding and circular financing as late-cycle signals — buy carefully and expect deep drawdowns. Verdict: PROCEED WITH CAUTION, respect the cycle.
Michael Burry — Scion Asset Management	The 'Big Short' investor is the loudest bear on this trade: he has placed roughly a US\$1.1bn bet against the AI complex (large put positions on Nvidia and Palantir) and compares Nvidia to Cisco in the dot-com bust. The direct counter-argument to this entire model. Verdict: SHORT IT — he is betting the bubble bursts.
Li Lu — Himalaya Capital (Munger's protégé)	Runs an extremely concentrated, low-disclosure value book in deeply understood businesses; has historically owned the likes of Alphabet and Micron. Would concentrate in just 1–2 of these names (Alphabet most likely) rather than hold 18. Verdict: SELECTIVE — quality, far fewer names.
Ray Dalio — Bridgewater Associates	Trimmed Nvidia but still holds it inside a hugely diversified, risk-balanced book. Wouldn't pick these individually and would warn hard against ~100% exposure to a single theme. Verdict: DIVERSIFY beyond AI.

The pattern is clear: the growth-oriented kingpins (Laffont, Druckenmiller) are comfortable in this neighbourhood and mainly debate which layer to own, while the value-and-cycle voices (Klarman, Marks) would wait for lower prices and Burry is betting against the trade outright. That split — enthusiasm on the layers, caution on the price — is exactly why this model concentrates in the strongest names, keeps the speculative sleeve small, and devotes a full section to the bear case.

6. Where each sector is going — the growth math

High growth requires the underlying markets to keep expanding. Here is what the major research houses (IDC, McKinsey, Goldman Sachs, Gartner, IEA, Grand View, MarketsandMarkets) currently forecast for each layer, with the implied runway. These are third-party projections, not guarantees.

Sector	Size now → 2030 (approx.)	Growth rate & runway
Semiconductors (total)	~\$680bn (2024) → ~\$1.6–1.75tn by 2030	AI has 'reset the demand baseline.' Data-centre chips alone reach ~\$843bn by 2030 — nearly half the market. Multi-year structural growth, no longer tied to consumer cycles.
AI data-centre capex	~\$405bn (2025) → ~\$1.0–1.4tn/yr by 2030	Goldman models ~\$7.6tn of cumulative AI capex 2026–2031. Hyperscalers alone plan ~\$700bn in 2026. The core engine of the whole thesis.
Cloud computing	~\$750bn–1.3tn (2025) → ~\$1.6–2.4tn by 2030	~12–20% CAGR. AWS ~30% / Azure ~20% / Google ~13% share. AI is now the dominant marginal driver of cloud growth.
AI power / data-centre electricity	US DC capacity ~33GW (2024) → ~95–120GW by 2030	Goldman: data-centre power demand +175% by 2030. US data centres rise from ~6% to ~11–12% of national electricity. Power is the binding constraint on AI.
Nuclear / uranium	Demand +~28% by 2030; WNA 5.3% CAGR to 2040	Hyperscalers signing 20-year nuclear PPAs (e.g. Microsoft–Constellation). Structural uranium deficit + government backing. Long-cycle tailwind.
AI in cybersecurity	~\$25bn (2024) → ~\$94–134bn by 2030	~24–28% CAGR — the fastest-growing slice of a ~\$350–500bn total security market. Expanding attack surface from AI/cloud.
Quantum computing	~\$1.7–3.5bn (2025) → ~\$4–20bn by 2030	Very high CAGR (~20–42%) but tiny absolute base. McKinsey sees ~\$97–173bn of value by 2035–2040. Real, but a decade-plus payoff — hence speculative sizing.

The through-line: every layer this model touches is forecast to grow at double digits or faster through 2030, and most analysts now describe the build-out as multi-year and structural rather than a single cycle. The debate is no longer *whether* AI matters, but which layers capture the returns and whether today's prices already embed the growth.

7. Why each major holding is included

For each anchor and key position: its role, the landmark deals it has signed, its recent financial scale, the analyst view, and the one-line case for inclusion in the model. Figures are USD.

NVDA — NVIDIA · Core · 13% · \$130,000

Role: The centre of gravity in AI compute. Sells the GPUs that train and run almost every large AI model, protected by the CUDA software moat. Reached a ~\$5 trillion market cap in late 2025.

Landmark deals: A landmark partnership to deploy at least 10 gigawatts of NVIDIA systems for OpenAI, with NVIDIA intending to invest up to \$100 billion in OpenAI as systems roll out (first phase H2 2026 on the Vera Rubin platform). At Computex 2026 it unveiled the RTX Spark — its first PC processor — pushing into AMD/Intel's x86 turf.

Recent financials: Quarterly earnings have been growing rapidly (July-quarter EPS estimated ~\$1.88 vs \$0.99 a year prior, ~+90%); full-year FY2027 EPS estimates near \$8.04, up ~76%. Data-centre revenue is the dominant driver.

Analyst view: Strong Buy from ~49–60 analysts; average target ~\$302 vs ~\$222 price (~+36–40%). Bulls: Tigress \$425, Cantor \$300, Barclays \$275. ARK has been aggressively buying.

Why it's here: The only mega-cap where consensus still implies large upside, and the supplier on the other side of nearly every major AI deal — hence the top weight in the model.

MSFT — Microsoft · Core · 8% · \$80,000

Role: The enterprise on-ramp to AI: Azure cloud, Copilot across Office, and the deepest commercial tie to OpenAI.

Landmark deals: OpenAI has committed to purchasing an incremental ~\$250 billion in Azure cloud services over roughly six years. Anchored a 20-year nuclear PPA with Constellation (Three Mile Island / Crane restart) and helped fund a ~\$30bn BlackRock AI-infrastructure vehicle.

Recent financials: Among the largest, most cash-generative companies on earth, with Azure as the key growth engine; the stock sold off modestly post-earnings.

Analyst view: Strong Buy (~53 buys); average target ~\$561 vs ~\$441 (~+27%), high ~\$870.

Why it's here: Lowest-drama way to own AI monetisation at scale; a quality ballast for the model.

GOOGL — Alphabet · Core · 9% · \$90,000

Role: Google Cloud, the Gemini model family, Search and YouTube — a full-stack AI platform that also designs its own TPU silicon (with Broadcom).

Landmark deals: Guiding to ~\$175–185 billion of 2026 capex, almost entirely AI infrastructure; signed a ~\$10bn Google Cloud deal with Meta. Gemini 3 launch cited as a major milestone; annual revenue topped \$400bn for the first time.

Recent financials: Q4 2025 revenue \$113.8bn; Google Cloud +48% YoY to \$17.7bn; ~28–29x trailing earnings — the cheapest mega-cap multiple in the group.

Analyst view: Strong Buy; average target ~\$430 (~+19%). The standout 'smart money' bet: Berkshire tripled its stake into a top-5 holding.

Why it's here: Cheapest quality compounder in the group, with the Berkshire endorsement — sized as the second-largest anchor.

AMZN — Amazon · Core · 7% · \$70,000

Role: AWS is the largest cloud platform (~30% share) and the biggest single AI-infrastructure landlord; retail and advertising add diversified cash flow.

Landmark deals: Projecting ~\$200 billion of 2026 capex (up from ~\$131bn). Acquired Talen's data-centre campus next to a nuclear plant; in talks to invest ≥\$10bn in OpenAI and part of a ~\$15bn Anthropic commitment (with Google).

Recent financials: AWS reaccelerating; the capex leader among hyperscalers.

Analyst view: Strong Buy (45 of 46 analysts buy); average target ~\$313 (~+18%). Tepper's largest position (~15%) and a fresh Ackman add.

Why it's here: The broadest hedge-fund conviction in the universe plus the largest cloud franchise.

AVGO — Broadcom · Core · 6% · \$60,000

Role: Designs custom AI accelerators (ASICs) for hyperscalers and dominates high-end networking silicon connecting GPU clusters.

Landmark deals: OpenAI committed to deploying ~10 gigawatts of custom accelerators co-designed with Broadcom (industry-estimated value ~\$350bn). Q2 FY2026 AI semiconductor revenue +143% YoY to \$10.8bn; AI bookings exceeded \$30bn in the quarter.

Recent financials: Total Q2 FY2026 revenue ~\$22.2bn (+48%); shares dipped after the print on a small software-revenue miss.

Analyst view: Strong Buy; average target ~\$487 implying anywhere from ~+1% to +18% depending on the (conflicting) live price; high ~\$630. Druckenmiller added it in Q1 2026.

Why it's here: The custom-silicon and networking pick-and-shovel play — benefits from hyperscalers diversifying away from sole-sourcing NVIDIA.

TSM — Taiwan Semiconductor · Core · 6% · \$60,000

Role: The foundry that physically manufactures the leading-edge chips for NVIDIA, AMD, Apple and Broadcom — the single most critical manufacturing chokepoint in AI.

Recent financials: Reasonable ~38x P/E for a near-monopoly on advanced nodes; a steady compounder rather than a momentum name.

Analyst view: Strong Buy; average target ~\$468 (~+6%), high ~\$600. A Tepper core holding (~7.6%).

Why it's here: Diversified exposure to every AI chip designer at once, with monopoly-grade manufacturing economics.

ASML — ASML · Core · 5% · \$50,000

Role: Sole global supplier of EUV lithography — no advanced AI chip can be made without its machines.

Analyst view: Strong Buy; JPMorgan Overweight with a \$2,200 target. The average target sits roughly in line with the ~\$1,718 price, so it is arguably fully valued today.

Why it's here: An irreplaceable monopoly at the base of the chain — sized at 5% because the valuation already reflects the moat.

GEV — GE Vernova · Core · 5% · \$50,000

Role: Supplies gas turbines, grid equipment and generation that data centres depend on — a scarce global vendor at the power chokepoint.

Landmark deals: Backlog grew to ~\$163 billion (Q1 FY2026), with orders +71% YoY; its Electrification unit booked \$2.4bn of data-centre equipment in one quarter — more than all of the prior year.

Analyst view: Buy / Strong Buy (~29 buys); average target implies ~+12–25%, high ~\$1,424.

Why it's here: The clearest 'sell the picks for the AI power rush' name, with a backlog that gives multi-year revenue visibility.

CEG — Constellation Energy · Core · 5% · \$50,000

Role: Operates the largest US nuclear fleet (~26 reactors, ~22GW) — exactly the carbon-free baseload hyperscalers want.

Landmark deals: 20-year power purchase agreement with Microsoft to restart Three Mile Island Unit 1 (rebranded Crane Clean Energy Center, ~835MW, targeted 2028) — its largest PPA ever. Also the ~\$26.6bn Calpine acquisition.

Analyst view: Buy / Strong Buy; average target implies ~+26–32%, high ~\$462; ~25x P/E — the best value-plus-upside combination in the power bucket.

Why it's here: Direct, contracted exposure to AI-driven nuclear demand at a reasonable multiple.

AMD — Advanced Micro Devices · Cyclical · 5% · \$50,000

Role: The credible #2 in AI GPUs (MI300/MI450) and the main alternative hyperscalers fund to avoid sole-sourcing NVIDIA.

Landmark deals: OpenAI committed to ~6 gigawatts of AMD Instinct GPUs (~\$90bn revenue potential), with a warrant for OpenAI to buy up to 160 million AMD shares (~10% of the company) as deployment scales. First 1GW from H2 2026 on MI450.

Analyst view: Strong Buy, but the stock has rallied ABOVE its ~\$478 average target (~-10% implied); P/E ~169x. Bull case Barclays \$665. ARK has been selling into the strength.

Why it's small: Genuine secular winner, but priced for perfection — held at a contained 5% as a cyclical, not an anchor.

VST — Vistra · Cyclical · 5% · \$50,000

Role: A top independent power producer with nuclear and gas capacity positioned for AI/data-centre demand.

Analyst view: Strong Buy; average target implies ~+40–60% (high ~\$320, ~+105%); PEG ~0.47 — the cheapest high-upside power name. Tepper doubled his position in Q1 2026.

Why it's here: The best value-plus-upside power play, with hedge-fund conviction behind it. Risks: PJM regulation and ~\$20bn debt — hence Tier 2, not Tier 1.

MU — Micron · Cyclical · 4% · \$40,000

Role: One of three global HBM memory makers — the memory that feeds AI accelerators, now repriced as a strategic asset.

Recent financials: Beat recent estimates (EPS ~\$4.78 vs \$4.07); HBM capacity reported sold out well into the future.

Analyst view: Strong Buy (39 of 44 buys, zero sells) — yet the stock trades ABOVE the average target (~-27% to -31% implied) after a parabolic run. Bulls remain loud: Morgan Stanley \$1,050+, UBS \$1,625, a Street-high near \$1,750. Tepper's #2 position (tripled).

Why it's small: Central to the AI memory supercycle with huge bull-case upside, but cyclical and extended — capped at 4%.

CRWD — CrowdStrike · Cyclical · 4% · \$40,000

Role: Leading AI-native endpoint and cloud security platform — the layer that protects AI and cloud workloads.

Analyst view: Buy / Strong Buy; JPMorgan raised its target to \$800, but the stock trades above the ~\$565 average target (~-24% to -28% implied); P/S ~40x.

Why it's small: Best-in-class security compounder in a fast-growing market, but richly valued — a 4% position, not an anchor.

CCJ — Cameco · Cyclical · 4% · \$40,000

Role: A top-two Western uranium producer plus a stake in reactor-builder Westinghouse — leverage to the nuclear build-out.

Analyst view: Buy; average target implies ~+14–28%, high ~\$175. Backed by a structural uranium deficit and rising government support for nuclear.

Why it's here: The fuel layer of the AI-power thesis, with a 20-year demand tailwind from data-centre nuclear PPAs.

VRT — Vertiv · Satellite · 4% · \$40,000

Role: A leading supplier of data-centre power and thermal-management (cooling) systems and a direct NVIDIA partner — critical as GPU racks push power density and heat ever higher.

Analyst view: Strong Buy / Buy; individual targets run up to \$435–488. A favourite of AI-infrastructure thematic funds.

Why it's here: A pure 'arms dealer to the data centre' — owns the cooling/power bottleneck that every GPU build-out must solve.

ORCL — Oracle · Satellite · 4% · \$40,000

Role: Enterprise AI cloud and a central partner in the OpenAI/Stargate build-out.

Landmark deals: Remaining performance obligations (RPO / backlog) reported at ~\$553 billion, up ~325% YoY — much of it the Stargate / OpenAI data-centre build.

Analyst view: Buy (32 of 42 buys); average target ~\$247 (~+1–7%), high ~\$400. Carries heavy debt and reports on 10 June 2026 — a near-term catalyst/risk.

Why it's small: Extraordinary backlog story, but Buy-only consensus, high leverage and binary earnings risk keep it a modest satellite.

TRGP — Targa Resources · Satellite · 3% · \$30,000

Role: Major US natural-gas midstream operator — gas is currently the largest single source of incremental data-centre power (>40% of US data-centre electricity).

Analyst view: Buy; average target implies roughly flat to +12% — the weakest 'high-growth' fit in the model, with more of an income tilt.

Why it's smallest: A diversifying, cash-generative play on gas-fired AI power — the most expendable name in the model.

IONQ — IonQ · Satellite · 3% · \$30,000

Role: Pure-play quantum computing (trapped-ion) — a frontier optionality bet on a market that is still pre-commercial.

Recent financials: Pre-profit (FY2025 loss ~\$510m) with a large cash cushion (~\$3.3bn); valued on potential, not earnings.

Analyst view: Strong Buy but explicitly speculative; Street-high ~\$100 (Jefferies). ARK holds it as quantum optionality; otherwise largely retail-driven.

Why it's small: A capped 3% lottery ticket — if quantum inflects this decade the payoff is convex; if it doesn't, the loss is contained.

8. Why a model like this is structured this way

- **It owns the whole stack.** Tools, chips, foundry, memory, cloud, security and power are all represented, so no single bottleneck — silicon, electricity or software — can sink the model on its own.
- **Conviction is concentrated in the strongest names.** The nine core anchors all carry Buy/Strong Buy consensus with positive average upside; the names whose average targets imply downside (MU, CRWD, AMD) are deliberately kept small.
- **Risk is sized on purpose.** ~64% sits in lower-volatility anchors and only ~14% in higher-risk satellites, capping the damage any single high-volatility loss could do.
- **There is real optionality.** IONQ, VST and CCJ offer convex payoffs (capped dollar risk, large upside) appropriate for an explicitly high-risk illustration.
- **The smart money corroborates the thesis.** Tepper, Ackman, Berkshire and Wood are all leaning into AI infrastructure — and Tepper's book in particular mirrors this model's chip-memory-power tilt.
- **Every sector has a multi-year runway.** Each layer is forecast to grow double-digits-plus through 2030, so the holdings are riding structural expansion rather than a single-quarter trade.

9. The other side — contrarian views & risks

A fair analysis argues against itself. Here are the strongest reasons to be cautious — or to wait.

- **The 'circular financing' / bubble critique.** NVIDIA invests \$100bn in OpenAI, which buys NVIDIA chips; AMD grants OpenAI warrants for ~10% of itself in exchange for GPU orders; Microsoft, Oracle and Broadcom are all entangled in the same loop. Critics compare this self-reinforcing money-go-round to the late-1990s dot-com build-out and question whether revenue reflects organic demand or recycled capital. If the loop breaks, multiple holdings fall together.
- **Three holdings already trade above consensus.** MU, CRWD and AMD have run past their average analyst targets. Buying them means buying on momentum and bull-case upside, not on a margin of safety — exactly what value investors like Seth Klarman and Howard Marks would warn about, and what Michael Burry is actively betting against.
- **Extreme thematic concentration.** This is ~100% one theme — AI infrastructure. A single macro shock (rates, a hyperscaler capex cut, a Taiwan/China event hitting TSMC, an AI 'capex winter') would hit nearly every position at once.
- **Valuations are stretched across the board.** Forward multiples of 30–170x leave little room for disappointment. Even small earnings misses or guidance cuts can trigger 20–40% drawdowns in high-volatility names.
- **Hyperscaler capex could disappoint.** The entire thesis assumes ~\$700bn+ of 2026 data-centre spend keeps rising. Any pullback in capex guidance would cascade through chips, power and memory.
- **Single-name execution risk.** Oracle's debt and binary earnings; Constellation/Vistra nuclear-restart timelines and PJM regulation; IONQ's pre-profit cash burn; TSMC's geopolitical exposure — each can move independently of the theme.

Bottom line: a model like this is built for a high-risk, high-growth profile and could lose a large portion of its value in a downturn. Anyone considering a similar approach should weigh these risks against their own circumstances and seek independent, licensed advice.

10. Sources & data notes

Analyst consensus & price targets: TipRanks, MarketBeat, Yahoo Finance, Investing.com / S&P Global Market Intelligence, GuruFocus, Barchart, Zacks (early June 2026). **Fund positions:** Q1 2026 13F filings (as of 31 Mar 2026) via SEC EDGAR, Seeking Alpha, ValueSider, Grufity, Yahoo Finance and BBAE coverage. **Sector forecasts:** IDC, McKinsey, Goldman Sachs Global Institute, Gartner, IEA, IoT Analytics, Grand View Research, MarketsandMarkets, BCC Research. **Deals & financials:** company releases and SEC filings via OpenAI, NVIDIA, DataCenterDynamics, TechCrunch, Bloomberg, Utility Dive, CNBC.

Data caveats: (1) Several stocks moved so fast that aggregator 'average targets' lag the price and show paper downside — both a sign of analyst catch-up and a genuine valuation warning. (2) There is a live source conflict on Broadcom's exact early-June price, which widens its implied-upside range. (3) Price targets are 12-month third-party estimates, not guarantees; high-volatility names

here can move $\pm 50\%$ +. (4) The 'verdicts' in the kingpins' section are an interpretive read of public approach and disclosed positioning, not personalised endorsements.

11. Important legal disclaimer

Educational and informational purposes only. This document and the hypothetical portfolio it describes are provided solely for general educational and informational purposes. Nothing in it is intended to be, and nothing in it should be construed as, financial, investment, trading, tax, accounting, legal or other professional advice. It does not constitute a research report within the meaning of any applicable regulation.

Hypothetical and illustrative. The US\$1,000,000 portfolio is entirely hypothetical and illustrative. It does not represent any actual account, holding, trade or recommendation, and no real money has been invested in the manner described. Hypothetical allocations have inherent limitations and do not reflect actual trading, liquidity constraints, fees, commissions, taxes, slippage or the impact of market conditions.

Not personalised advice; no recommendation. The content is general in nature and does not take into account your objectives, financial situation, knowledge, experience or needs. It is not a recommendation, offer, solicitation or invitation to buy, sell, hold or deal in any security, financial product or investment strategy, in any jurisdiction. You should not act or refrain from acting on the basis of this document.

No advisory or fiduciary relationship. Neither the author, the Miss AI brand, nor any associated person is acting as your financial adviser, broker, dealer or fiduciary. No advisory, fiduciary or other relationship is created by your reading or use of this document. The author is not a licensed or registered financial adviser, investment adviser, broker-dealer or financial-product adviser in any jurisdiction unless expressly stated in writing.

Risk warning — you can lose everything. All investing involves risk, including the possible loss of all capital invested. The securities discussed are high-volatility, high-risk equities; their prices can fall as well as rise, sometimes very sharply and without warning. Past performance and forward-looking estimates are not reliable indicators of future results. Currency movements, leverage, illiquidity, concentration and macroeconomic events can amplify losses. Do not invest money you cannot afford to lose.

Forward-looking statements. This document contains forward-looking statements, projections and third-party estimates (including price targets and market-size forecasts). These are inherently uncertain, are based on assumptions that may prove incorrect, and may differ materially from actual outcomes. No representation is made that any forecast will be achieved.

Accuracy, completeness and timeliness. Information has been compiled from third-party sources believed to be reliable as of early June 2026, but no representation or warranty, express or implied, is made as to its accuracy, completeness, reliability or timeliness. Markets, prices, ratings and fund positions change constantly and may already be out of date. The author is under no obligation to update this document.

Possible conflicts of interest. The author and/or associated persons may now or in the future hold positions in, or have other interests relating to, the securities mentioned. No assurance is given that any such position is consistent with the views expressed here. Any potential conflict should be assumed rather than inferred from its absence.

Do your own research; seek licensed advice. Before making any financial decision, you should conduct your own due diligence and obtain independent advice from a professional who is licensed or authorised in your own country, state or territory and who can take account of your personal circumstances. References to named investors or funds are illustrative and do not imply their endorsement of, or association with, this document.

Jurisdiction-specific notices

United States — The author is not registered as an investment adviser under the U.S. Investment Advisers Act of 1940 or with the SEC or any state, and is not a broker-dealer registered with the SEC or FINRA. This material is not a research report under SEC/FINRA rules, is not personalised investment advice, and is not an offer or solicitation in any state where such an offer would be unlawful. Securities investing involves risk of loss.

United Kingdom — This communication has not been approved by an authorised person for the purposes of section 21 of the Financial Services and Markets Act 2000 (FSMA). It is provided on the basis that it is directed only at persons for whom it is lawful to communicate it, is for general information, and does not constitute a financial promotion, personal recommendation or advice under FCA rules. Capital is at risk.

European Union / EEA — This material does not constitute investment advice, investment research, a personal recommendation or a marketing communication within the meaning of MiFID II (Directive 2014/65/EU) or related regulations, and has not been prepared in accordance with legal requirements designed to promote the independence of investment research. It is general information only.

Australia — The author does not hold an Australian Financial Services Licence (AFSL). This material is general information only, is not financial product advice under the Corporations Act 2001 (Cth), and does not consider your objectives, financial situation or needs. Consider obtaining advice from an AFSL holder and reviewing any relevant disclosure document before deciding.

Canada — The author is not registered as an adviser or dealer with any Canadian provincial or territorial securities regulator. This material is general information only, is not advice, and is not an offer or solicitation in any province or territory where that would be unlawful.

New Zealand — This material is not a regulated financial advice service under the Financial Markets Conduct Act 2013 (FMC Act) and is not financial advice for the purposes of New Zealand law. It is general information only and does not consider your particular financial situation or goals.

All other jurisdictions — This document is not directed at, and is not intended for distribution to or use by, any person in any jurisdiction where such distribution, publication or use would be contrary to local law or regulation, or would subject the author to any registration or licensing requirement. It is your responsibility to inform yourself about, and to observe, all applicable laws and regulations in your own jurisdiction. Where any provision of this disclaimer is unenforceable in a given jurisdiction, the remaining provisions continue to apply to the fullest extent permitted by law.

Limitation of liability

No warranty; use at your own risk. This document is provided “as is” and “as available” without warranties of any kind, whether express or implied, including but not limited to warranties of merchantability, fitness for a particular purpose, accuracy or non-infringement.

Exclusion of liability. To the maximum extent permitted by applicable law, the author, the Miss AI brand and any associated persons accept no liability whatsoever for any direct, indirect, incidental, consequential, special or exemplary loss or damage (including, without limitation, loss of profits, trading losses or loss of data) arising from or in connection with any use of, or reliance on, this document or its contents, even if advised of the possibility of such loss. Nothing in this disclaimer excludes or limits any liability that cannot lawfully be excluded or limited.

Trademarks. Company names, ticker symbols and logos referred to are the property of their respective owners and are used for identification and commentary purposes only. Their use does not imply any affiliation with or endorsement by those owners.

This disclaimer is a general template and does not itself constitute legal advice. For full protection you should have it reviewed and tailored by a qualified lawyer in each jurisdiction in which you publish or distribute this material.