

INTELLIGENCE SERIES • AI INFRASTRUCTURE

Indirect Is Not Diversified

There are two kinds of AI infrastructure, and only one of them diversifies anything. A follow-on to *The AI Trade Hidden in Your Portfolio*, answering the question that came back from allocators: if I want AI exposure without a binary bet on which model wins, what do I actually buy?

Date: July 2026 • **Audience:** Professional investors & allocators • **Basis:** 3HC forward analytics against J.P. Morgan 2026 long-term capital market assumptions • **Prepared by:** 3HC Intelligence

70%+

of family offices say AI is a priority and hold no infrastructure allocation at all

-318bps

the enabler-heavy tilt underperforms in the one shock aimed at its own thesis

15 of 15

percentage points of the rotation went into assets the standard stress test never shocks at all

EXECUTIVE SUMMARY

- **The question is the right one.** 65% of family offices are prioritising AI investments while more than 70% hold no infrastructure allocation. Moving down the stack, away from the model layer and into what it runs on, is a sound answer to a real problem.
- **The constraint is not chips. It is watts.** A current-generation rack draws roughly 120 kilowatts against five to ten for a legacy rack. Northern Virginia's large-load queue holds about 70 GW against a 24.7 GW system peak, with a seven-year wait, and vacancy is at a record low of 1.4%.
- **Moving down the stack genuinely removes selection risk.** You no longer need to be right about which laboratory wins. The landlord is paid by whoever signs the lease. That is a real improvement over buying the model layer.
- **But there are two kinds of infrastructure in that list.** AI-enabler assets are paid if the build continues. AI-agnostic assets are paid regardless. Both sit under the same heading; only the second is a diversifier.
- **We built the enabler-heavy version and measured the difference.** Under a hyperscaler capital expenditure pullback it falls 13.35% against 10.17% for the baseline. The concentration risk flagged at the outset was real and larger than assumed.
- **Indirect is not diversified.** The tilt swapped which company wins for whether the build goes on. That is a different expression of the same outcome, not an offset to it, and nobody is starting from a clean sheet.

A note on method. Forward analytics are estimates computed against third-party long-term capital market assumptions and are not a forecast by Three Horizons Capital. The hyperscaler pullback scenario was built by us and its shocks are analyst assumptions, not calibrated history. Private-asset volatilities are as published and are not adjusted for appraisal smoothing, so risk in the private sleeves is likely understated.

PART I

The answer is to move down the stack.

Our previous piece measured portfolios that had become a single trade by accident. The question that came back from allocators was the practical one: if I want AI exposure without a binary bet on which model wins, what do I actually buy?

Family offices have decided they want AI. They have not bought what it runs on.

J.P. Morgan Private Bank, 2026 Global Family Office Report: 333 family offices across 30 countries, published February 2026.

65%

say they are prioritising AI investments

70%+

hold no infrastructure allocation at all

0.7%

infrastructure and real assets, as a share of the average portfolio

✓ The intent is close to unanimous. The exposure is close to absent.

THREE HORIZONS CAPITAL

In J.P. Morgan Private Bank's 2026 Global Family Office Report, covering 333 offices with an average net worth of \$1.6bn, 65% say they are prioritising AI investments and more than 70% hold no infrastructure allocation at all. BNY Wealth found 83% of single family offices ranking AI the single biggest opportunity of the next five years. The intent is nearly unanimous and the exposure is nearly absent.

If you do not want to pick between laboratories, stop buying the model layer and buy what the model layer runs on. AI compute is an energy problem wearing a semiconductor costume. **The constraint is not chips. It is watts.** A current-generation NVIDIA GB200 NVL72 rack draws roughly 120 kilowatts; a legacy enterprise rack draws five to ten.

The constraint is not chips. It is watts.

Power draw per rack, kilowatts. A twelve to twentyfold step change that has to be supplied, cooled and connected somewhere.

GB200 NVL72 rack
current generation, liquid cooled

120 kW

Legacy enterprise rack
air cooled

7.5 kW

✓ Northern Virginia's large-load queue holds about 70 GW of requests against a 24.7 GW all-time system peak. A new large load now waits about seven years.

NVIDIA GB200 NVL72 documentation; Virginia State Corporation Commission Case PUR-2026-00011 (Feb 2026). As at Jul 2026.

THREE HORIZONS CAPITAL

Those places are running out. Northern Virginia's large-load interconnection queue holds around **70 gigawatts** of requests against an all-time system peak of 24.7 gigawatts, and a new large load waits about **seven years**. CBRE has North American primary-market vacancy at a record low of 1.4%, Northern Virginia at 0.3%, with roughly 80% of space under construction pre-leased. The four largest hyperscalers are on track for roughly \$700bn of capital expenditure in 2026, and three of the four raised guidance in the first quarter.

There are six investable layers beneath the models: power and energy, semiconductors and advanced packaging, digital real estate, cooling and physical plant, fibre and network, and the enabling software layer. The most under-owned is **power**. Most family offices holding anything here hold the data centre REIT and the semiconductor name and stop. Almost nobody owns the grid, which is where the supply constraint is most severe and where power purchase agreements create genuinely long-dated contracted cash flow.

PART II

What moving down the stack really buys you.

It removes selection risk, and that is not a small thing.

You no longer need to be right about which laboratory wins, which architecture prevails, or which model is state of the art in three years. The landlord is paid by whoever signs the lease. The utility is paid by whoever draws the power. You have swapped a binary bet on companies for a claim on the physical plant all of them need, and picked up contracted, often inflation-linked cash flows

along the way.

That is a real and defensible improvement over buying the model layer directly, and anyone telling you otherwise is not paying attention. But removing selection risk is not the same as removing risk, and the difference is where this paper earns its keep.

PART III

There are two kinds of AI infrastructure.

Look again at those six layers and you will notice the list contains two different things wearing the same coat.

THE DISTINCTION THAT DECIDES WHETHER AN INFRASTRUCTURE SLEEVE DIVERSIFIES


	AI-ENABLER INFRASTRUCTURE	AI-AGNOSTIC INFRASTRUCTURE
Paid	If the build continues	Regardless
Examples	Data centre REITs, liquid cooling and plant, development-stage private equity, advanced packaging and wafer fabrication equipment	Regulated transmission and distribution serving all load, existing generation under long-dated contract, water, ports, toll roads
Cash flows are a function of	Hyperscaler capital expenditure	An economy existing
Behaviour if the build slows	Lease-up slows, orders slow, development marks reprice	Largely unaffected
Diversifies your existing AI exposure	No	Yes

The distinction is not academic, and we can show it from our own scenario. Rank every asset class in the book by how hard a hyperscaler capital expenditure pullback hits it, and the four buckets the AI tilt bought are the four worst-hit. Not roughly. Exactly, with a clear gap to the fifth.

The four buckets the tilt bought are the four the shock hits hardest.

Shock applied to each asset class under a hyperscaler capital expenditure pullback. Teal outline marks the four buckets the AI tilt bought.

■ AI-enabler: paid if the build continues ■ Everything else in the book



△ That is what AI-enabler means. Four assets whose cash flows all depend on the same capital expenditure, bought together.

3HC hyperscaler pullback scenario, Jul 2026. Shock assumptions are analyst estimates, not calibrated data. Estimate and forecast.

THREE HORIZONS CAPITAL

That is what AI-enabler means in practice. Four assets bought together whose cash flows all depend on the same capital expenditure. Diversified by asset class on the statement, one dependency underneath.

THIS WAS FLAGGED BEFORE ANYTHING WAS BUILT

The original framework that prompted this work said it in terms: "Concentration risk is asymmetric. Several of these layers are highly correlated. A slowdown in hyperscaler capital expenditure would simultaneously hit data centre REITs, power procurement, cooling and semiconductor equipment. **True diversification requires mixing AI enabler with AI agnostic infrastructure.**" What we did next was find out how much it mattered.

PART IV**So we built the enabler-heavy version and measured it.**

Fifteen percentage points out of generic large-cap equity, core bonds and traditional real estate, into four AI-linked buckets.

BASELINE FAMILY OFFICE BOOK VERSUS THE AI-INFRASTRUCTURE-TILTED VERSION

ASSET CLASS	BASELINE	AI-TILTED	CHANGE
US Large Cap	25%	20%	-5%
EAFE Equity	12%	10%	-2%
Emerging Markets Equity	5%	5%	0%
US REITs	3%	8%	+5%
Global Core Infrastructure	3%	7%	+4%
Private Equity	15%	18%	+3%
Venture Capital	2%	5%	+3%
US Core Real Estate	7%	4%	-3%
US Aggregate Bonds	12%	8%	-4%
US High Yield Bonds	3%	3%	0%
Direct Lending	4%	4%	0%
Diversified Hedge Funds	5%	5%	0%
US Cash	4%	3%	-1%

Against the J.P. Morgan 2026 assumptions it does what the story promises. Expected return goes from **7.10% to 7.35%**, volatility from **11.61% to 12.26%**, and concentration in a single growth factor comes down from **95.9% of variance to 93.8%**. It also holds up better in five of six standard stress tests. That is the sentence we would normally have led with, and it is the one that turned out not to mean anything.


PART V**The stress test was not looking at what we bought.**

A portfolio that improves its downside in five of six scenarios while raising expected return is suspiciously good. So we checked what the scenarios actually shock.

The six standard macro scenarios shock US large cap, emerging market equity, core real estate, aggregate bonds, high yield, hedge funds and cash. They apply **no shock at all** to listed REITs, private equity, direct lending, EAFE equity, global core infrastructure or venture capital. Not in most of the six scenarios. In none of them. Those six asset classes have no rows in the shock table at all.

Every point of the rotation went where the stress test was not looking.

Change in weight, percentage points, coloured by whether the standard six-scenario library applies a shock to that asset class.



3HC scenario engine coverage audit, 21 Jul 2026. Verified line by line against the scenario definitions.

THREE HORIZONS CAPITAL

Every single percentage point of the rotation went into a bucket the scenario library never shocks. Eighty-seven percent of what funded it came out of buckets it shocks in every scenario. The portfolio did not become more resilient. It moved its money to where the stress test was not looking.

THIS IS NOT A QUIRK OF ONE MODEL

Scenario libraries are built from public-market history, because that is where clean mark-to-market series exist. Private and real assets get muted shocks or none, because the data to calibrate them honestly does not exist in the same form. The structural consequence is that **a standard stress test will tend to reward any move from public into private**. If you have been shown one that says your private-heavy reallocation improves your worst case, this is the first thing to check.


PART VI

The test that mattered.

The risk the original framework named was not a generic recession. It was a pullback in hyperscaler capital expenditure. That scenario did not exist in the library, so we built it.

In the one shock aimed at its own thesis, the AI tilt loses.

Modelled portfolio outcome under a hyperscaler capital expenditure pullback. More negative is worse.



3HC scenario built Jul 2026. Shock assumptions are analyst estimates, not calibrated data. Estimate and forecast.

THREE HORIZONS CAPITAL

The concentration risk flagged at the outset was real, and larger than anyone had assumed. Adding four points of gold and long Treasuries recovers 58 of those 318 basis points and leaves the tilt still 260 behind the baseline. It also barely touches the factor concentration, moving the growth share from 93.78% to 93.43%, thirty-five basis points. Ballast is worth holding for an ordinary drawdown. It is not an antidote to this.

PART VII

Indirect is not diversified.

The enabler-heavy tilt did remove selection risk. You are no longer betting on which laboratory wins. That

was the objective and it was achieved.

What it did not do is reduce your dependency on the build continuing. You swapped *which company wins* for *whether the build goes on*. That is a different risk, often a better one, and for some allocators it is exactly the right trade. But it is a **different expression of the same outcome**, not an offset to it.

Indirect and diversified are not the same word. The first describes how you hold the exposure. The second describes whether you have less of it. Moving down the stack does the first and, on our numbers, slightly worsens the second.

AND NOBODY IS STARTING FROM A CLEAN SHEET

If you held no AI exposure at all, an enabler-heavy infrastructure allocation would be a perfectly sensible, cash-flow-backed way into the theme, and we would say so without qualification. But we measured what Western books already hold, and they are already long this outcome through the mega-cap complex, in size, by default. For that reader, adding enabler infrastructure is not diversifying into the theme. It is concentrating further into a dependency they already have and mostly did not choose.

PART VIII

What actually diversifies.

The resolution was in the original framing before any of this was modelled. Our numbers only sharpen it.

Mix the enabler with the agnostic

If the infrastructure sleeve is entirely data centres, cooling and development equity, it is an AI position with an infrastructure label. Regulated networks, contracted existing generation, water and transport carry the characteristics people actually want, contracted and inflation-linked cash flow, without the capital expenditure dependency riding underneath.

Prefer layers where the constraint is physical and the counterparty is not a single theme

Power and grid remain the most under-owned layer for a reason worth acting on. Transmission that serves all load is a different asset from a data centre that serves one tenant.

Watch the entry point, not just the exposure

Listed proxies are priced for the benign case. Private development co-investment and infrastructure debt still carry an illiquidity premium, particularly in power and fibre, and are the difference between paying for the theme and being paid to finance it.

Price the lock-up honestly

Counting only genuinely locked vehicles, the tilt takes the book from **24% to 34%** in capital you cannot readily retrieve. Capital calls arrive on the general partner's schedule, during drawdowns, which is exactly when your liquid assets are worth least.

And if what you want is protection rather than participation, this is not the list. That answer has not changed since the previous piece: it is the exposures that do not depend on the trade continuing at all.

PART IX

The test you do not run is the risk you do not see.

We could have published the first version of this. Return up twenty-five basis points, growth concentration down two percentage points, better in five of six stress scenarios. Every number in that sentence reproduces exactly against the model. It would have been a good-looking paper and it would have been wrong, because the scenarios did not examine the assets the portfolio had just bought.

The correction did not come from better data. It came from taking a risk that had already been flagged in a footnote and asking what it was actually worth, then building the test that would answer.

TWO PIECES, ONE DISCIPLINE

The first found portfolios that had become a single trade without anyone deciding to. This one finds that the most popular remedy is a different expression of the same trade, and that the standard stress test will congratulate you for buying it. Measure the exposure you have. Then test it against the thing that would actually hurt it, not against whichever scenarios happen to be available.

Data sources. Family office intent and allocation: J.P. Morgan Private Bank, 2026 Global Family Office Report, February 2026 (333 offices, 30 countries); BNY Wealth, 2025 Investment Insights for Single Family Offices (282 single family offices). Rack power: NVIDIA GB200 NVL72 product documentation. Interconnection queue: Virginia State Corporation Commission, Case PUR-2026-00011, February 2026. Data centre vacancy and pre-leasing: CBRE North American data centre trends, 2026. Hyperscaler capital expenditure: company disclosures, first quarter 2026. Portfolio construction, forward analytics, factor decomposition and scenario outcomes: 3HC analytics computed against the J.P. Morgan 2026 Long-Term Capital Market Assumptions (USD, 58-asset covariance matrix), reproduced live 21 July 2026. The long-term capital market assumptions are a third-party forecast over a horizon longer than the five to seven years used as the family office planning horizon here, and are not a forecast by Three Horizons Capital. The hyperscaler capital expenditure pullback scenario was constructed by Three Horizons Capital; its shocks are analyst assumptions and are not calibrated to historical data. Private-asset volatilities are as published and are not adjusted for appraisal smoothing, so volatility and risk in the private sleeves are likely understated relative to true mark-to-market risk. Capital call and drawdown pacing conventions are planning assumptions, not derived from our data. Market figures are as at July 2026. Asset classes and vehicles referenced are illustrative examples, not recommendations.

APPENDIX

The hyperscaler pullback scenario, in full.

Everything else in this paper can be reproduced from named third-party sources or from the published capital market assumptions. This scenario cannot, because we built it. So here is the whole shock vector. It represents our judgement of how directly each asset class depends on hyperscaler capital expenditure continuing, and it is not calibrated to historical data.

SHOCK APPLIED TO EACH BUCKET UNDER A HYPERSCALER CAPITAL EXPENDITURE PULLBACK

ASSET CLASS	SHOCK	REASONING
Venture Capital	-35%	The model layer takes the most direct and least cushioned repricing
US REITs	-30%	Anchor-tenant lease growth assumptions reprice directly
Private Equity	-22%	Development and supply chain marked down, appraisal smoothing softens the near-term print
Global Core Infrastructure	-18%	Capital expenditure demand cut, long-dated power purchase agreements cushion part of the cash flow
US Large Cap	-12%	Still concentrated in the AI complex, but diversified across non-AI mega-caps
EAFE Equity	-8%	Semiconductor supply chain exposure within a diversified international index
US High Yield Bonds	-8%	Fibre and grid private-credit-adjacent spread widening
Emerging Markets Equity	-6%	Foundry and memory supply chain exposure
Direct Lending	-6%	Digital infrastructure private credit, moderate repricing
US Core Real Estate	-3%	Indirect risk-sentiment spillover only
Diversified Hedge Funds	-2%	Modest drag
US Aggregate Bonds	+2%	Modest flight-to-quality bid. This is a sector shock, not a macro recession
US Cash	+2%	As above

Applied to the two portfolios this produces **-10.17%** for the baseline and **-13.35%** for the AI-tilted book, a gap of 318 basis points. Adding four points of gold and long Treasuries, funded entirely from US large cap, gives **-12.77%**, recovering 58 basis points and leaving the tilt 260 behind the baseline.

HOW TO DISAGREE WITH THIS

If you think a given shock is wrong, change it and re-run. The conclusion is not sensitive to any single line: it holds because the enabler-heavy tilt concentrates four buckets that all depend on the same capital expenditure, and no plausible set of shocks makes four correlated bets safer than not making them. What would change the conclusion is a view that hyperscaler capital expenditure cannot fall, and that is a forecast, not a portfolio construction argument.

IMPORTANT INFORMATION

Disclaimer

Please read this notice. It governs your use of this material.

Publisher and purpose. This material is published by Three Horizons Capital Limited, a company registered in Dublin, Ireland ("Three Horizons Capital", "3HC", "we", "our") as part of its Intelligence Series. It is provided for general information and educational purposes only and reflects our views as of the date of publication, which are subject to change without notice. We are under no obligation to update it.

Not advice, not a recommendation, not an offer. Nothing in this material constitutes, or should be construed as, investment, legal, tax, accounting or other advice, a research recommendation, or an offer, invitation or solicitation to buy, sell, subscribe for or transact in any security, fund, financial instrument or strategy. It is not a personal recommendation and does not take account of the objectives, financial situation, knowledge, experience or needs of any person. Any securities, indices, exchange-traded funds, asset classes or vehicles named are referenced solely as illustrative examples to explain an analytical framework, and their mention is not a recommendation to buy, hold or sell.

No reliance; take your own advice. You should not rely on this material as the basis for any decision. Conduct your own due diligence and obtain independent financial, legal, tax and regulatory advice appropriate to your circumstances before acting. Any decision to invest is made solely at your own risk.

Intended audience and distribution. This material is intended only for professional investors, institutional investors, eligible counterparties and their advisers, and is not intended or suitable for retail investors or the general public. It is not directed at, and must not be acted upon by, any person in any jurisdiction where its publication, distribution or availability would be contrary to law or regulation, and it is the reader's responsibility to observe all applicable laws. Nothing here is intended to be a financial promotion, marketing communication or investment research where that would require an authorisation, approval or regulatory classification that has not been obtained. Three Horizons Capital is a non-regulated entity; it is not authorised or regulated by the Central Bank of Ireland or by any other financial regulator, and this material is not a financial promotion of any regulated product or service.

Data, sources and methodology. This material draws on proprietary Three Horizons Capital data and analytics together with third-party data believed to be reliable. We have not independently verified all third-party data and make no representation or warranty as to its accuracy, completeness, timeliness or fitness for purpose. Certain figures are model outputs, estimates or derived statistics and are inherently uncertain; they depend on assumptions, definitions and data windows that are stated where practicable and may differ from figures produced on other bases.

Forward-looking statements and performance. This material may contain forward-looking statements, scenario analyses and statistical relationships that involve risks and uncertainties; actual events or results may differ materially. Statistical relationships such as correlations are period-dependent, may be sensitive to methodology, and may not persist. Past performance and historical relationships are not a reliable indicator of future results. The value of investments and any income from them can fall as well as rise, and investors may not recover the amount originally invested.

Third-party marks; no affiliation. Third-party names, indices, benchmarks and marks referenced (which may include, without limitation, MSCI, FTSE Russell, S&P, Russell, iShares/BlackRock and Morningstar) are the trademarks or service marks of their respective owners, used for identification and commentary only. Their use does not imply any affiliation with, or sponsorship, endorsement or approval by, those owners.

Conflicts of interest. Three Horizons Capital, its affiliates, and their respective officers, employees and connected persons may from time to time hold positions in, provide services relating to, or have other interests in, the securities, funds, instruments, issuers or markets referenced.

Limitation of liability. To the maximum extent permitted by applicable law, Three Horizons Capital and its affiliates accept no liability whatsoever for any direct, indirect or consequential loss or damage arising from any use of, or reliance on, this material or its contents.

© 2026 Three Horizons Capital. All rights reserved. This material may not be reproduced, redistributed or transmitted, in whole or in part, without our prior written consent. Three Horizons Capital Limited is a company registered in Dublin, Ireland; it is not authorised or regulated by the Central Bank of Ireland or any other financial regulator and does not provide regulated financial services, investment advice or a regulated investment research service. Questions: threehorizonscapital.com.

Three Horizons Capital • Intelligence Series • AI Infrastructure

Prepared by Three Horizons Capital for professional investors and allocators. Data is proprietary unless otherwise attributed. Nothing in this document constitutes investment advice or a recommendation to buy or sell any security. Securities and funds named are illustrative examples, not recommendations. Past performance and historical correlations are not indicative of future results.