

The Great **AI** Build-Out **Part 2**: When the Build-Out Meets the Balance Sheet

Executive Summary: *The AI infrastructure buildout continues to accelerate, but the market is increasingly focused on whether the extraordinary level of capital investment can generate adequate returns. While the technology and demand are real, rising capex, cash-flow pressure, depreciation and increasingly complex financing structures are creating a higher hurdle for investors. Historical parallels to telecom and mortgage cycles suggest that even transformative technologies can produce disappointing investment returns when capital deployment outpaces near-term economic benefits. The recent selloff in AI infrastructure stocks reflects this growing skepticism, with the market demanding greater evidence of sustainable utilization, cash-flow generation and returns on invested capital. **The buildout isn't necessarily slowing—but the hurdle rate is rising.***

Two months ago, we wrote about the extraordinary progression of the artificial intelligence investment cycle: 2024 was the year of the AI leaders, 2025 brought accelerating infrastructure spending, and by 2026 the market had shifted decisively toward the “picks and shovels” of AI. Infrastructure companies became the strongest-performing segment of the market, while the gains of the major AI platforms moderated.

The question now is not whether the AI build-out is real. It clearly is. The question is whether the economic returns from that build-out can keep pace with the extraordinary amount of capital being committed to it. That distinction is becoming increasingly important.

The Build-Out Is Still Accelerating

The most striking feature of the current cycle is that spending has not yet begun to slow. In fact, the opposite is true. S&P Global estimates that the five largest hyperscalers have already spent roughly \$1.1 trillion on capital expenditures over the past five years and could spend another \$5.3 trillion through 2030. Capital intensity is already exceeding the peaks reached during both the dot-com and Global Financial Crisis periods.

Our Three Cycles, One Pattern paper put the scale of the current cycle into historical perspective. The five largest hyperscalers are on pace to spend approximately \$725 billion on infrastructure in 2026, with roughly \$545 billion directed toward AI compute, data centers, power and networking. That compares with an inflation-adjusted peak of approximately \$213 billion for the entire telecom infrastructure build-out.

In other words, **AI infrastructure spending in a single year is now several times larger than the peak annual spending of the late-1990s telecom boom.**

This is not necessarily a problem. Transformative technologies often require enormous upfront investment. The problem comes when capital spending begins to grow faster than the economic returns generated by that capital. And there are signs that the gap is widening.

The Next Phase Is a Test of Capital Efficiency

The market is increasingly focused on a simple question: **What is the return on the next dollar of AI infrastructure spending?**

J.P. Morgan estimates that AI capital expenditures will consume roughly 93% of hyperscalers' operating cash flow in 2026, compared with just 33% in 2023. Epoch AI reaches a similar conclusion from a different perspective: operating cash flow has been growing at roughly 23% annually, while cash capital expenditures have been growing at approximately 70%. On current trajectories, the two converge around the third quarter of 2026.

Again, this does **not** mean the hyperscalers are running out of money. They remain among the strongest companies in the world financially. It does mean that the financial burden of the build-out is becoming large enough that investors will increasingly demand evidence that the spending is producing attractive economic returns.

S&P Global Ratings makes the same point from a credit perspective. It expects combined hyperscaler capital expenditures to exceed \$1.3 trillion by 2027 and forecasts negative free operating cash flow across the six largest hyperscalers during 2026 and 2027, before an anticipated recovery beginning in 2028.

The distinction is important: **negative free cash flow during an investment cycle is not inherently bad.** A company can temporarily sacrifice cash flow to build assets that generate substantial future returns. The issue is what happens if the returns arrive later than expected—or are lower than expected.

This Is Where History Becomes Useful

This is the connection to our recent Three Cycles, One Pattern paper from a few months ago. We argued that the current AI infrastructure boom shares an important structural characteristic with the telecom build-out of the late 1990s. In both cases, the underlying technology is real, the demand is real, and the infrastructure is genuinely valuable.

The problem is that **capital can arrive faster than the underlying economics can absorb it.** That creates a recurring sequence:

Genuine innovation → capital floods in → investment outruns near-term economics → financial and accounting structures bridge the gap → complexity increases → eventually, returns and valuations matter.

That was the central lesson of the telecom cycle, and it is the reason we believe the comparison is worth making today.

There is also an important accounting feature to a capex supercycle. When an equipment manufacturer sells a GPU cluster, server rack or networking system, the seller generally recognizes the revenue immediately. The buyer, however, capitalizes the purchase and recognizes the cost through depreciation over the asset's estimated useful life.

As a result, **the seller's earnings benefit today while the buyer's expense is pushed into future periods.** During a powerful investment cycle, that can make aggregate earnings appear stronger before the full economic cost of the investment has been recognized.

The accounting isn't wrong, just misleading. **Reported earnings and economic returns sometimes move on different timetables.**

The Depreciation Question

This brings us to one of the more important—and underappreciated—questions in the current AI cycle: **How long will today's AI infrastructure actually remain economically useful?**

The traditional data-center model assumed relatively long-lived infrastructure. AI hardware is different. Semiconductor performance is improving rapidly, and newer generations of accelerators can deliver dramatically more computing power per unit of energy.

If a GPU cluster is depreciated over six years but becomes economically obsolete much sooner, reported earnings will initially overstate the economic return on the investment. Conversely, if older equipment remains useful because demand continues to grow faster than supply, the longer depreciation schedule may prove appropriate.

The issue is not whether management teams are following accounting rules. They are. The issue is whether the assumptions embedded in those rules accurately reflect the unusually rapid pace of technological change. This is precisely the kind of question that tends to matter more as an investment cycle matures.

The Physical Limits Are Becoming More Visible

There is another constraint that did not exist to the same degree in previous technology cycles: **the physical infrastructure required to support AI is enormous.**

AI data centers require electricity, transmission capacity, cooling, land, construction labor and increasingly complex connections to the power grid. Across the U.S., permitting restrictions, grid constraints and community opposition are beginning to slow or limit new data-center development. Of course, this may ultimately prove healthy. Investment cycles rarely proceed in a straight line, and physical constraints can force capital deployment to become more selective.

There is also a potential change taking place within AI itself. Recent discussions among leading AI companies about pacing frontier-model development highlight the possibility that the industry could gradually shift some resources from increasingly expensive model training toward inference and commercial deployment. That would not mean the AI opportunity is ending. It could simply mean that the **next phase of the cycle looks different from the first.**

The Market Is Beginning to Ask the Question

Perhaps most interestingly, the market's behavior is beginning to reflect some of these concerns. After an extraordinary first half of 2026, the leadership of AI infrastructure stocks has become much less uniform. The market appears increasingly willing to reward companies that can demonstrate actual demand and cash generation while becoming less willing to pay simply for exposure to the AI spending cycle. That is a meaningful change in character.

Our Three Cycles, One Pattern paper emphasized that speculative cycles do not normally end because investors suddenly discover that the underlying technology is worthless. Telecom fiber eventually carried enormous amounts of data. The technology was not the problem. The problem was that too much infrastructure was built too quickly, financed on assumptions that proved too optimistic, and valued as though the expected economics would arrive immediately.


The same distinction applies to AI.

The build-out does not need to collapse for investment returns to disappoint. It only needs to generate cash more slowly than the market expects.

Credit Markets Are Providing Another Signal

One of the more interesting developments is that this question is beginning to show up outside the equity market.

The UBS HOLT/Merton analysis we have been following shows the share of IT companies with modeled probabilities of default above 5% has risen to levels exceeding those reached during the early-2000s technology bubble.



We would not interpret this as evidence that widespread defaults are imminent. A Merton-style probability of default is a **market-implied measure** that is sensitive to equity prices, volatility and leverage. It can therefore rise well before an actual credit event—and can also rise simply because equity markets become more volatile.

But that is precisely why we find it useful. It is telling us that, beneath the headline strength of the technology sector, **financial stress is becoming more visible in portions of the market.**

That is consistent with what we would expect as a capital-intensive cycle moves from the initial enthusiasm phase toward the much more important question of economic returns. The smaller infrastructure companies that lever up for growth will feel this more acutely as the cycle matures.

What We Are Watching

We do not believe the answer is to declare the AI investment cycle over. The technology is real, the opportunity is substantial, and the strongest companies involved have balance sheets that are vastly healthier than those of the telecom carriers of 1999. That is an important difference.

Instead, we believe the next phase of the cycle will be defined by **capital discipline.**

There are five questions we believe are increasingly important:

1. Are incremental returns on AI capital improving or deteriorating?
2. Is free cash flow beginning to catch up with capital expenditures?
3. Is utilization keeping pace with capacity?
4. Are depreciation assumptions keeping pace with technological change?
5. Is financing becoming more complicated?

As the capital requirements grow, debt, leases, joint ventures and other financing structures become increasingly important. Complexity itself is not a problem, but history suggests that it deserves greater scrutiny as a cycle matures.

Conclusion

The first phase of the AI cycle was about **proving that the technology works**. The second phase may be about **proving that the economics work**. That is a much higher bar.

Transformative technologies can create enormous economic value while simultaneously producing disappointing investment returns when capital is deployed too quickly, valuations become detached from near-term economics, or the expected returns arrive later than investors anticipate.


The AI build-out is not slowing. The hurdle rate is rising.

After two years of extraordinary investment and extraordinary market returns, we believe that may be the most important transition to watch as the cycle moves forward. And we believe our attention to asset utilization and ROIs, coupled with our higher yields and downside protection, position us to continue to outperform in the coming years.

Van Hulzen Covered Call Strategy

The Van Hulzen Covered Call strategy invests in US companies that we consider to have high shareholder yield (dividends and share repurchases) and uses call options with the goal of reducing portfolio volatility and creating incremental income. The goal is a portfolio that has equity exposure while seeking higher than average annual income (target of 6-8% annual), although there is no guarantee that the strategy will achieve its objective, generate profits or avoid losses. Below you will find the graph of the Van Hulzen Covered Call Strategy and the Covered Call Index BXM.

Covered Call Strategy Performance (gross as of 08/31/2026)



Returns (annualized)*	Aug 2026	3M	6M	YTD	1 Year	3 Years	5 Years	7 Years	10 Years	Inception
Van Hulzen (Gross)	1.2%	2.7%	5.1%	5.5%	9.6%	11.9%	7.2%	7.8%	7.2%	8.1%
Van Hulzen (Net)	1.2%	2.5%	4.7%	5.1%	9.0%	11.1%	6.5%	7.1%	6.5%	7.2%
BXM	1.6%	4.8%	7.8%	10.0%	19.3%	13.4%	8.4%	8.4%	7.9%	6.2%
Difference (Gross-BXM)	-0.4%	-2.0%	-2.7%	-4.5%	-9.7%	-1.6%	-1.2%	-0.6%	-0.7%	1.8%
S&P 500	2.7%	1.7%	12.4%	13.1%	20.4%	21.0%	12.8%	16.5%	15.4%	10.1%

*Inception date : 12/31/2001. Figures greater than one year are annualized. Van Hulzen returns represent actual returns from composite of accounts. Source: Bloomberg, Tamarac

The foregoing content reflects the opinions of Van Hulzen Asset Management and is subject to change at any time without notice. Content provided herein is for informational purposes only and should not be used or construed as investment advice or a recommendation regarding the purchase or sale of any security. There is no guarantee that these statements, opinions or forecasts provided herein will prove to be correct. Past performance is not a guarantee of future results. Indices are not available for direct investment. Any investor who attempts to mimic the performance of an index would incur fees and expenses which would reduce returns. All investing involves risk including the potential for loss of principal. There is no guarantee that any strategy will be successful. The CBOE S&P 500 BuyWrite Index (BXW) is a benchmark index designed to track the performance of a hypothetical buy-write strategy on the S&P 500 Index. The BXW is a passive total return index based on (1) buying an S&P 500 stock index portfolio, and (2) "writing" (or selling) the near-term S&P 500 Index (SPXSM) "covered" call option, generally on the third Friday of each month. The S&P 500 Index consists of 500 stocks chosen for market size, liquidity, and industry group representation. It is a market-value weighted index (stock price times number of shares outstanding), with each stock's weight in the Index proportionate to its market value. It is widely used as a benchmark of U.S. equity performance. It is not possible to invest directly in an index. All numbers presented are sourced through Bloomberg and Tamarac.