JDP
CAPITAL MANAGEMENT

Survivor & Thriver Fund, LP

2Q 2026 Investor Letter

"The future is already here. It's just not evenly distributed."

— William Gibson

In the second quarter the fund was up 8.5% on a net basis bringing our first half return to -7.8%. This compares to the S&P 500 which was up 10.2% in the first half, assuming dividends had been immediately reinvested back into the index.

The first half of 2026 felt less like six months and more like several years. I have not experienced another period where I could see this much technological change, and the financial debate around this change, crammed into two quarters.

Every week seems to bring another breakthrough. New, AI-native businesses built on top of proprietary models are growing much faster than the early software that moved work flows away from pen and paper.

For long-term investors, periods like this are both exciting and extraordinarily difficult. Our challenge is not deciding whether AI will reshape the economy but where the enduring economics will ultimately reside.

We are also witnessing an extraordinary effort to reduce the cost of AI infrastructure by innovating around today's bottlenecks. The stronger the bottleneck, the greater the incentive to eliminate it. That is why I have been cautious about assuming today's bottlenecks will enjoy durable pricing power.

Markets tend to reward what is easiest to measure first. So far in 2026 that has been the companies supplying the scaffolding for AI, where revenue appears quickly in financial statements and converts to after-tax profit.

Far harder to measure are the businesses quietly becoming stronger because AI is making their products better, lowering costs and deepening customer relationships. That is where my work has been focused.

A boom in leveraged-single stock ETFs

The market's obsession with the bottleneck trade has created a lopsided environment concentrated in what is perceived to go up the most in the shortest period of time.

According to Interactive Brokers there are more than 100 leveraged single stock ETFs today, a product category that *did not exist* four years ago. There are now an estimated 400 leveraged ETFs including baskets of the most popular trades.

The Economic Times reported that 75% to 85% of leveraged ETFs are concentrated in chips and AI bottlenecks. For example the Direxion 2x Micron ETF (MUU) increased assets from just $516 million at the beginning of 2026 to more than $5.4 billion by the end of May, a tenfold increase in less than

five months. The CSOP SK Hynix 2x Leveraged ETF surpassed US$5.1 billion of assets within months of launch, overtaking Tesla's TSLL to become the largest single-stock leveraged ETF in the world.

The harder to measure AI beneficiaries are being cast aside

What is far more difficult to measure are the businesses quietly becoming stronger and more valuable as a result of intelligent AI capex spend. Many of these stocks have not participated in the market rally this year.

There is no line item on Spotify's income statement called "AI-related revenue." No company reports how much customer retention improved because recommendations became smarter, how many engineering hours were saved through AI-assisted development, or how much more quickly new products reached customers because software teams became dramatically more productive.

Warren Buffett once remarked that **"no one ever says, 'I want a Coke only if it comes in an Alcoa aluminum can.'"** The can is indispensable, but it is not what customers are buying. AI infrastructure may prove similar. While commodity inputs are essential, I believe the greatest long-term value will accrue to the businesses that use them to build products customers value.

Selected portfolio update

Over the past year I have mentioned the fundamental acceleration taking place across our portfolio. That trend has only become more pronounced even as stock prices lag behind. Product development cycles have compressed, and features that once required a year or more of engineering are being shipped in weeks. Products are becoming more intelligent, more personalized and more valuable to customers.

As with many modern companies, our holdings are accomplishing more every quarter with the same or fewer resources as a year ago.

The two exceptions to that are businesses we have chosen to own precisely because they are doing the opposite, spending aggressively today to earn a disproportionate return tomorrow.

Tesla and Amazon

Both Tesla and Amazon are in the middle of multi-year capital investment cycles that are difficult to underwrite on a near-term earnings basis but, in our view, represent some of the most compelling long-term value creation opportunities in the portfolio.

Tesla and Amazon are rare in that they are willing and able to sacrifice near-term earnings in order to invest behind very large long-term opportunities. Having the strength to all in on investing for the future when the market punishes your stock has created trillions of dollars in value for shareholders, even as many investors remained skeptical at every turn along the way.

Unlike at Tesla where the return on capital spending is still taking shape, Amazon's $200 billion AI capex has a scoreboard you can read in real time.

AWS grew 17% in Q1 2025, accelerated to 24% in Q4 2025, and hit 28% in Q1 2026, its fastest growth rate in 15 quarters, on a business now running at $150 billion in annualized revenue with 37% operating margins. That acceleration is increasingly the result of Amazon's AI investments.

Amazon's custom silicon division, Trainium and Inferentia combined, hit $20 billion in ARR in the second quarter, growing triple digits year over year and representing roughly 13% of total AWS revenue. Amazon consumes most of its own chips internally and CEO Andy Jassy has said the chip business would be closer to **$50 billion ARR** if sold to third parties, making it one of the largest semiconductor businesses in the world. AWS growth is constrained by its ability to build infrastructure fast enough to serve the backlog.

Despite all this great news, Amazon trades near the lowest valuation in its history, a disconnect that we believe will not last.

On the other hand, Tesla may be the most difficult to value and the most important to hold. Tesla's return on recent capex spend is less visible on a spreadsheet than Amazon's, but one example where it is starting to show up is in the rapid improvements of FSD (Full Self Driving) software.

FSD materially changes the economics of the vehicle business

The ability to invest aggressively has allowed Tesla to improve its FSD software faster than at any point in the company's history. Continuous software releases have materially reduced unnecessary interventions, improved decision-making in complex urban environments, expanded the range of situations the system can handle confidently, and accelerated the rate at which new capabilities reach customers.

Every improvement benefits from billions of miles driven by its fleet, creating a feedback loop that compounds with scale. While regulatory approval and customer adoption will not move in a straight line, the underlying pace of technical progress continues to accelerate.

In the second quarter Tesla reported 1.48 million FSD subscriptions, up 56% year over year, with more than half of North American deliveries now including FSD and the highest-ever quarterly net subscription additions on record.

I expect the pace of FSD subscription growth to accelerate. Tesla continues to expand the product internationally with subscriptions now available in 11 countries, and growing with regulatory approvals.

The present value of a $99 monthly software subscription over the life of the car materially increases lifetime revenue per unit, and as that software becomes a larger percentage of total revenue, Tesla gains significantly more flexibility in how it prices the underlying hardware than manufacturers whose economics depend almost entirely on the initial vehicle sale.

FSD pricing power also grows with adoption. Once a customer becomes accustomed to being driven rather than driving themselves, what began as a discretionary software purchase becomes a habitual service, with no real competition today.

FSD subscription revenue should not be viewed in isolation today. The simple math of current subscription run rate revenue relative to Tesla's market cap materially understates what is happening. The growing installed base of FSD subscribers creates a flywheel in which improving software changes consumer behavior, increases the desirability of owning a Tesla, generates more real-world driving data to further improve the product, and ultimately expands the future Robotaxi opportunity.

The thread that connects Tesla's FSD progress to its upfront spending is the same one that runs through nearly every great investment in history. Very few businesses have both the opportunity set and the leadership willing to exchange years of visible, measurable costs for a leadership position in a large new industry that is still being defined. Tesla and Amazon are two of them.

New position: Cerebras Systems (CBRS)

The fund recently established a position in Cerebras Systems at $178 per share, approximately 50% below CBRS's post-IPO peak. At an enterprise value of roughly $40 billion, the market might be materially undervaluing the company's earnings power.

Based on the revenue already under contract, Cerebras should be worth more than $1,000 per share within the next few years as they deliver on a $20 billion contract with OpenAI and other large-scale partnerships announced just in the last few months (AWS, AMD, CrowdStrike).

Background

Cerebras develops AI infrastructure designed specifically for the inference phase of artificial intelligence, the process of generating answers after a model has been trained. The company first became known for building the Wafer Scale Engine, the world's largest commercially deployed processor, but it is increasingly evolving beyond selling hardware.

Today, Cerebras is building a network of AI data centers that provide inference as a service, allowing customers to purchase computing capacity. This transition shifts the business toward recurring, consumption based revenue that more closely resembles a cloud infrastructure provider than a traditional semiconductor company.

Customers can access this capacity directly through Cerebras Cloud or through partners such as AWS, allowing developers to use the platform without purchasing dedicated hardware. Management expects cloud inference to become the primary driver of revenue and profitability as utilization across its installed infrastructure increases.

What's not in the stock price

While investors continue to debate quarterly margins and near term execution, they are assigning relatively little value to what could become one of the fastest growing inference infrastructure platforms over the remainder of the decade.

- **The OpenAI contract alone should drive more than 8x revenue growth over the next 24 months.** Based on management's deployment schedule, the annualized revenue could increase from ~$750 million exiting Q2 2026 to more than $6 billion by 4Q of next year, before any contribution from AWS or future enterprise customers.

- **OpenAI economics seem underestimated.** OpenAI has committed to deploy up to 750 megawatts of Cerebras inference capacity through 2028 under a contract management values at more than $20 billion. Because deployments require both wafer production and new data center capacity, revenue recognition naturally lags contract announcements by several quarters.

- **The data center buildouts are progressing as planned.** Cerebras currently has approximately 410 MW of inference capacity either operational, under construction or publicly committed,

including production deployments across North America, 40 MW in Alabama, 160 MW at Bell Canada's Saskatchewan AI campus, and the recently announced 200 MW European expansion across France and the Nordics. This represents approximately 55% of the 750 MW OpenAI deployment already operational or committed, leaving approximately 340 MW (45%) still to be secured over the next several years. Nearly all of the currently announced expansion is expected to come online by early 2027, while management has repeatedly stated that additional data center announcements will continue as customer deployments progress. Given the company's consistent execution to date and the long lead times built into the OpenAI rollout through 2028, the market could be materially underestimating Cerebras' ability to deliver the remaining capacity.

- Summary of MW schedule as of July 2026
  410 MW deployed or committed by early 2027
  55% complete
  340 MW remaining
  45% left to secure

- **Beyond OpenAI.** Since the OpenAI announcement in January 2026, Cerebras has announced a multi-year inference partnership with AWS (March 13, 2026), bringing its technology to Amazon Bedrock, a strategic partnership with CrowdStrike (July 22, 2026) to power Falcon AI Detection and Response, and a collaboration with AMD (July 23, 2026) to integrate the Cerebras Wafer-Scale Engine with AMD's Helios AI infrastructure.

- **Current margins materially understate normalized profitability.** To satisfy initial OpenAI demand, Cerebras subcontracted inference capacity from another customer while its own infrastructure is being deployed. As owned capacity replaces subcontracted compute, management expects gross margins to normalize at approximately 60%, versus the high 30% gross margin guide for 3Q 2026

- **The valuation remains compelling relative to earnings power.** Applying a 12-15x revenue multiple to $15 to $20 billion of annual revenue implies an enterprise value of approximately $240 to $300 billion, 4.5x to 6.5x upside over the next three to four years. With slightly more scale Cerebras should be capable of generating returns on invested capital above 25% and producing EBITDA margins in excess of 20%. This valuation already assumes the exercise of the OpenAI warrants, increasing the share count by approximately 11%, and does not include any value from additional large-scale customer wins, which we believe are increasingly likely as the OpenAI deployment validates the platform.

New posistion: Eli Lilly (LLY)

We paid $1,100 per share for our position in Lilly, a price that screens as fully valued for a traditional pharmaceutical company. The stock looks like a permanent compounder when you consider it is running the largest supercomputer in the pharmaceutical industry, training proprietary models on decades of irreplaceable scientific data, and has partnered with Nvidia to re-imagine the drug discovery process.

Drug discovery may be one of the largest beneficiaries of AI of any industry on earth, and the companies with the most proprietary data, the deepest clinical histories, and the strongest market positions are positioned to widen their advantages in ways that are genuinely difficult for competitors to replicate. Lilly is a business where AI is simultaneously reducing the cost of discovery, accelerating time to approval, and compounding its lead over competitors.

Similar to what we described in transportation, AI is restructuring the economics of the entire drug development process in ways that disproportionately favor scale. Preclinical costs are expected to fall 30 to 70%, development timelines to compress by 30 to 40%, and the industry as a whole could save $75 to $125 billion annually by 2030.

In addition to generating nearly $20 billion in quarterly revenue growing at 56% year over year, Lilly is running what may be the most powerful supercomputer in the pharmaceutical industry, operating a platform connecting 1,300 biotech companies, and deploying AI across discovery, manufacturing, and clinical operations.

Company background

Eli Lilly is the definition of a Survivor and Thriver business. Founded in 1876 and public since 1952, the company has spent nearly 150 years repeatedly reinventing itself at the frontier of medicine. That willingness to embrace scientific risk has created one of the greatest long-term compounding businesses in history, generating more than 131,000% in total returns since listing, and approximately 1,300%, or nearly 30% annually with dividends reinvested, over the past decade.

Today Lilly's revenue today is built around its GLP-1 and incretin platform, with Mounjaro for type 2 diabetes and Zepbound for obesity generating $19.8 billion in Q1 2026, up 56% year over year, against full year guidance of $82 to $85 billion.

Alongside those franchises are two additional platform assets with multi-year runways: donanemab, the Alzheimer's treatment that slowed cognitive decline by up to 35% in Phase 3 trials, and orforglipron, the newly approved once-daily oral GLP-1 pill that removes the injection barrier and meaningfully expands the addressable market.

Not your average large cap pharma stock

The market values Lilly primarily as a pharmaceutical company with a leading position in obesity. Beneath that sits a company using 150 years of proprietary experimental data, molecular structures, clinical trial outcomes, and research failures to build an AI-driven competitive advantage that should compound for many years.

Traditional drug discovery is expensive and slow. Screening millions of potential compounds to find a handful worth advancing can cost hundreds of millions of dollars *before* a human trial begins. Lilly is restructuring that model from the ground up.

TuneLab

In September 2025, Lilly launched TuneLab, a platform comprised of 18 AI models trained on hundreds of thousands of proprietary molecules and preclinical data. The platform is open to outside biotech partners and now connects 1,300+ biotech companies. Partners fine-tune Lilly's models against their own compound libraries while neither party sees the other's raw data. Each partner trains locally and Lilly combines those updates into a shared model. Every new partner makes the models smarter, attracting more partners, creating a network effect around Lilly's data that has no precedent in the pharmaceutical industry.

By becoming the standard AI infrastructure for drug discovery, discoveries from partners on the platform become potential opportunities for Lilly to finance or distribute.

LillyPod

LillyPod is the hardware that runs TuneLab. Launched in February 2026, it was built as the first Nvidia DGX SuperPOD powered by 1,016 Blackwell Ultra GPUs, delivering over 9,000 petaflops of AI compute across 290-plus terabytes of GPU memory. Nvidia called it the most powerful AI supercomputer in the pharmaceutical industry.

The B300 architecture is purpose-built for drug discovery workloads, running molecular simulations, protein structure prediction, and reinforcement learning loops that iterate on candidate molecules in real time. Each experiment feeds back into the models, improving the predictions that drive the next one.

Lilly owns one of the world's most valuable scientific datasets, is investing aggressively to compound its advantage through AI infrastructure. Those forces reinforce one another in a way that becomes increasingly difficult to replicate.

Exited Roku after buyout offer

We exited our Roku position at approximately $144 per share following the announcement of a sale to Fox at $160 per share, comprised of $96 in cash with the remainder in Fox Class A stock. We have no interest in owning Fox stock, and the outcome, while profitable on paper, was far from the result this business deserved. Our cost basis was approximately $65 per share, or roughly $8 billion in enterprise value, and while we did okay, the real story is what we are leaving behind.

Roku had penetrated over half of all US broadband households and built one of the most defensible positions in streaming infrastructure. Since going public in 2017, platform revenue grew from $225 million to approximately $5 billion in 2026, compounding at 41% annually. That figure excludes hardware, which Roku deliberately prices as a loss leader to grow its installed base. Hours streamed on the platform grew from 14 billion at IPO to an estimated 172 billion hours in 2026, a CAGR of 32%. Roku carried no debt, held over $2 billion in cash, and was on track to generate $1 billion in free cash flow next year.

At a takeout enterprise value of approximately $23 billion the deal undervalues a business with accelerating growth in Latin America and a valuable leadership position in Mexico and Brazil.
Like our experience with Caesars Entertainment earlier this year, Roku is a painful reminder of the take-private risk public market investors have when a strategically important business trades at a depressed price for long enough. In Roku's case I suspect founder Anthony Wood, who will join the Fox board, was ready to step back and found a structure that allowed him to do so.

The lesson is that being right and being able to collect on being right are two different things in public markets.

New position: Berkshire Hathaway (BRKB) as cash equivalent

Using the Roku proceeds to purchase Berkshire Hathaway at $488 per share. Berkshire is a high-quality placeholder while we wait for better opportunities from our watchlist ideas.

With the broader market trading at elevated valuations, we are in no immediate rush, and Berkshire is one of the few places we are comfortable parking capital while we wait.

The irony is that Berkshire also one of the most undervalued large-cap businesses in the market today. Approximately $717 billion of Berkshire's $1 trillion market cap is in cash and investments. Additionally, the operating businesses, one of the most durable collections of American businesses ever assembled, generate about $360 billion in revenue.

The stock has not recovered since Warren Buffett stepped down as CEO. New CEO Greg Abel inherited a portfolio of 75-plus operating businesses that each ran autonomously for decades. That structure produced extraordinary results but also left significant operational efficiency on the table. My guess is that Greg has a big opportunity to build new incentives for the subsidiaries to modernize, share resources, and benchmark themselves against one another.

Abel is a builder and an operator in a way Buffett never needed to be and his ability to drive even modest improvements in capital efficiency and technology adoption could unlock an enormous amount of value that the current stock price does not reflect.

Berkshire could reasonably double over the next few years simply by allowing the operating businesses to be run with the same discipline and focus that the holding company has always applied to capital allocation.

Added: Alibaba (BABA)

We added to our position at $95 per ADR, above our average cost of $80, paying roughly 12 times operating earnings excluding cash and investments for a company growing its most important division 40% annually.

BABA could have upside of 200% to 300% simply based on the business as it exists today. At a current market cap of approximately $250 billion this is one of the most undervalued collections of technology assets in the world.

In June BABA sold off more than 30% when the Pentagon added the company to its list of Chinese military-linked companies on June 8, 2026, alongside BYD and Baidu, pushing the total list to 188 Chinese companies. On July 5, a US District Judge issued a temporary order blocking the Pentagon restrictions while the broader legal challenge proceeds. The case remains ongoing, but none of it changes what Alibaba is building.

With $38 billion in net cash and more than $70 billion across listed securities and private investments, including stakes in some of the biggest winners of China's AI buildout: Moonshot AI, parent of the Kimi model and targeting a $50 billion valuation; Zhipu and MiniMax, both of which listed on the Hong Kong Stock Exchange in early 2026 at returns of 70x on Alibaba's cost basis; and a position in memory chipmaker CXMT that has delivered over 500% in three years.

Qwen, BABA's owned open source LLM has surpassed 600 million cumulative downloads and birthed over 113,000 derivative models, making it one of the most widely adopted open-weight AI models in the world.

BABA's core business includes $140 billion in revenue anchored by China's dominant e-commerce platform and a cloud and AI infrastructure business. BABA's cloud market share of 22.5% in Asia Pacific is ahead of every global competitor. Alibaba Cloud grew external revenue 40% last quarter is set to accelerate for the foreseeable future.

BABA's Nvidia competitor in China: T-Head Semi

One of the more exciting BABA assets is T-Head Semiconductor, Alibaba's internally built chip division that has spent seven years developing a full-stack silicon. T-Head's Zhenwu AI accelerator is one of China's designated country champions in the national effort to replace Nvidia GPUs. The first generation shipped to 400 customers with 560,000 units delivered and matched the performance of Nvidia's then-available export chip. The second generation, released four months later, claimed three times that performance. The next generation, planned for 2027, targets three times the second generation.

Closing thoughts

Every technological revolution starts with a period where investors overly focus on what enables the change before appreciating what becomes possible because of it. The greatest value ultimately accrues to the businesses that use those foundations to build products, deepen customer relationships and strengthen competitive advantages.

I believe we are approaching that transition today. If we are right, the businesses we own will look fundamentally different and substantially more valuable with the passing of time.

Thank you for your patience and trust. My family and I are fully invested alongside your capital.

Sincerely,

Jeremy Deal