Re-imagine March 20, 2024

MU: Emerging pure-play in AI – could MU rival AMD's market cap?

LYNX EQUITY STRATEGIES

KC Rajkumar, CFA, PhD
(408) 425-5680
KC@LynxEq.com

We expect MU's earnings event today to be a wake-up call highlighting the central role flash storage plays in the ongoing AI revolution. Whereas investors are generally aware of NAND's gradual ASP increase, we are calling out a trend that is more than mere gradualism. Sometime during the Feb quarter, we believe flash storage demand from AI data centers ***inflected sharply upwards*** *resulting in a ratcheting up of MU's* ***NAND visibility****, and consequently, fab utilization. MU's previous earnings call was typified by management stating strong inflection in DRAM demand due to demand pull from HBM; management set the expectation of industry-wide DRAM pricing power to persist through Cy24. At the earnings call today, we expect a qualitative upgrade in the outlook of its NAND business, driven by all-flash storage arrays.*

We think the Street does not yet appreciate the fundamental role flash storage plays in ***AI data lake creation and AI training****. In addition to the well-understood role of its DRAM/HBM in AI servers, if MU's NAND business too were to become central to AI, we think MU could come to be perceived by investors less as a cyclical commodity play and increasingly as a secular pure play in AI. We expect the company to continue to be prudent with its capex. In the face of sustained demand from AI data centers,* ***at servers and storage****, we think MU is likely to re-prioritize its end markets and trim its exposure to commodity client devices – in other words, become* ***a pure play in AI, second only to NVDA****.*

In our assessment, industry conditions are ripe for MU to revisit the glory days of the FY17/18 super-cycle. That cycle, in our view, was fueled by a mad dash by hyperscale cloud names to grab share of mobile DAUs. If the current AI demand pull has the demand intensity and multi-year longevity Street consensus believes, why wouldn't the demand for memory/storage be just as strong, if not stronger?

Into the previous earnings call, we raised our PT from $80 to the 'high 80s' (link). Going into the call today, the stock is trading close to our expectation set 3 months ago. Over the next 12-18 months we can see a path for the stock to ***levitate to the $150 level****, a ~50% upside from current levels. Whereas the overall memory/storage market continues to be tepid, we think MU is likely to maintain capex discipline and* ***outperform memory peers*** *by gravitating to a richer mix of AI-centric DRAM and storage products.*

As the AI-driven cycle evolves, say 12-18 months out, we expect MU's overall revenue and end-market AI mix to exceed AMD, and gross margin to be comparable to AMD. If we were to think outside the box, we can imagine MU's market cap ***rivaling that of AMD,*** *currently trading at* ***~2x the market cap vs our MU PT****.*

--xx--

All-Flash Array storage – The unappreciated component of AI Data Center: Training and inferencing gets all the attention. However, training of a data set necessarily needs a coherent 'data lake'. Creating a coherent 'data lake' from disparate data sources internal to an enterprise is an arduous and expensive step. This major step in converting enterprise data into AI models has not gained the kind of investor mindshare that

training and inferencing has. We'll jump ahead and state that AI training can **ONLY be done on 'data lakes' residing on all-flash array storage (AFA)**.

According to **NVDA CEO** the amount of data processing prior to data training could represent **30% to 50%** of the total amount of work to get to a working AI model. NVDA should know something AFA storage. Within NVDA-enabled data centers are SuperPOD and BasePOD storage systems, all based on AFA storage.

In a recent blog post, **META** says '*Storage plays an important role in AI training, and yet is one of the least talked-about aspects. As the GenAI training jobs become more multimodal over time, consuming large amounts of image, video, and text data, the need for data storage grows rapidly*'. The blog post goes on to say META's Tectonic storage solution is optimized for Flash media and that META is acquiring the latest high-capacity E1.S SSD drives procured **from the market**.

While META's acquisition of 350K H100 GPUs garnered headlines, little has been said about it acquiring SSD drives in the open market. The last time META went on a buying spree of DRAM modules in the open market, it triggered the FY17/FY18 **super-cycle** at MU.

At its recent earnings call, **CRM CEO** spoke of its Data Cloud offering to enterprise customers, a **precursor to AI training**. The company reported tangible revenue from its Data Cloud product, nearly doubling y/y, but NOT from its AI training product. Company management stated that its annual guidance does NOT factor in AI training/inferencing, but it does factor in Data Cloud aka data processing. While CRM does not state it explicitly, we infer that the output of CRM's **Data Cloud resides on AFA storage,** in anticipation of CRM embarking on AI training once the data processing of the customer's internal data is completed. In other words, AFA storage procurement precedes AI servers. **Flash is just as important as DRAM/HBM.**

NTAP - Secular mix shift to AFA, raises gross margin outlook: NTAP raised annual EPS guidance on higher profitability. The earnings call was centered around a **secular step up** the company was seeing in its AFA business driven by AI data centers, driving richer mix. The company raised its go-forward margin to the **higher end of the historical 50%-60% range**. The company spoke of share gains with their AFA storage systems as enterprise customers prepared internal data ahead of fine-tuning foundational LLMs using domain-specific data. We think NTAP's perennially cyclical business is on the verge of shifting to secular growth. Post-results, the stock hit an all-time high.

NTAP garnered a valuable **shout-out from NVDA CEO** at his recent GTC keynote. And not for nothing. NVDA provides enterprise customers with a suite of libraries to accelerate data processing, the internal data set at NVDA-enabled data centers most likely resides on NTAP storage system. Given the centrality of data storage, we think there is a symbiotic relationship between the two companies.

We note that NTAP raised its margin outlook despite an increase in input commodity costs, aka NAND pricing. The latter, in turn, allowed **WDC to raise current quarter EPS** to the high end of guidance, due to 'a couple of points of additional gross margin'.

MU – what is the Street missing? That the commodity NAND pricing has been rising, from the doldrums of last year, is well understood. We think the Street expects MU's NAND business to mirror that of WDC – cyclical improvement in margin driven by higher pricing of commodity NAND. As the AI cycle evolves, **we expect trend in MU's NAND margin to resemble more like NTAP than WDC**.

Into growing demand for NAND-based AFA storage from the explosive growth in Gen AI data processing, training and inferencing, we expect MU to **trim its NAND business in the client device end market and pivot to a higher growth, higher margin AI end market**.

And into this market, what does MU have its competitors in the NAND business may lack? Two items: 1) customized flash controllers and 2) geographical proximity.

Customized flash controllers: While SATA and SAS based flash drives have standardized flash controllers, high-end AFA storage systems are based on highly optimized, proprietary NVME/PCIe flash controllers. This is especially so for AI storage systems which deal with enormous data sets and require ultra-high bandwidth and ultra-low latency. AI-focused AFA systems, we think, are based on proprietary flash controllers and are inherently NOT commoditized. We think MU has a leg up over its NAND competitors – Kioxia, WDC and SK Hynix. Samsung though may offer comparable solutions to MU.

Geographic proximity: In the still-evolving field of AI, the design of crucial components is in active development. As with the case of **SMCI**, geographical proximity to the brain trust of key AI players is a plus – it enables cross-pollination of ideas across the design teams at partner companies. Even though MU is headquartered in Idaho, its brain trust – company management and key design teams - reside a few miles from their counterparts at NVDA, AMD, NTAP, CRM, MSFT, META, GOOG, AMZN/AWS. And this gives MU an advantage over its NAND competitors in Asia.

Financials and PT: We model DRAM and NAND revenue up 60% each in the current Fy24 vs. consensus up 50%/46% respectively. We model DRAM/NAND revenue growth in Fy25 as up 50% and 70% respectively ahead of consensus up 45% each. We model gross margin for Fy24/25 at 20%/52% vs. consensus of 16%/39%. Our earnings estimates for FY24/25 are $24.3bn/$0.52 and $37.9bn/$11.83 ahead of consensus estimates of $22.7bn/($0.33) and $32.7bn/$7 respectively. At a 12x multiple over Fy25 EPS, we derive PT of $150.

P/S considerations – NVDA/AMD: Versus consensus estimate of Fy24 revenue, **MU currently trades at P/S multiple of 4.4x**. Versus our Fy25 revenue estimate, our **$150 PT implies P/S remains at 4.4x**. Versus the consensus CY24 estimates, NVDA trades at P/S of 20x and AMD at 11x.

Let us re-imagine - where is MU market cap headed? At our $150 PT, MU's cap comes within spitting distance of INTC's current cap. How about AMD?

Based on consensus estimates, **MU's FY25 revenue** of $33bn runs significantly **ahead of AMD's** CY24 of ~$26bn. It would take a 25% bump at AMD in CY25 to get to MU's next FY estimated revenue.

What about gross margin? MU's Feb gross margin, coming out of an extraordinarily difficult post-pandemic period, is sub-15%. But what if the AI revolution is as sustainable as the Street assumes? The consensus estimate is for gross margin to get to ~40% in Fy25; we are modeling 52%, ahead of consensus. A 50% level brings MU **closer to AMD's mid-50% gross margin**. That would still be below MU's 59% gross margin from Fy18.

And what about **AI revenue mix**? AMD guided MI300x to $3.5bn for FY25. Even if we assume the company raises the outlook to say $5bn, that would still be only **~20%** of overall corporate revenue. Is it likely that

over time MU's AI revenue mix could trend to the **50% level?** We think it could. We'll have to watch for MU management commentary.

AMD's revenue and EPS gets an AI multiple even though its AI mix is low. AMD is bogged down with legacy business, which is likely to weigh down corporate growth rates. If the above three items – revenue growth, gross margin and mix - were to go MU's way, **isn't it possible MU could challenge AMD's market cap?**

Net/Net: Over the next 12-18 months, we think MU has the potential to emerge as a pure play in AI, second only to NVDA. With MU's 4-handle P/S vs. NVDA's 20-handle multiple, we think MU offers superior risk-reward to AMD and NVDA. We will be **long into print** and will add to position on weakness. We have a line of sight to PT $150 and a ~$150bn market cap. And if the stars line up just right, MU could rival AMD's market cap, which is currently running at ~$300+bn.

Disclosures and Disclaimers

Limitations of Information

Copyright, Trademarks, Intellectual Property