Memory Market Update

SOCAMM content noise offers a buying opportunity; thoughts on NVDA-SKH partnership and takeaways from Computex 2026

Technology - Semiconductors

Jay Kwon [AC]
(82-2) 758-5725
jay.h.kwon@jpmorgan.com
J.P. Morgan Securities (Far East) Limited, Seoul Branch

Sangsik Lee
(82-2) 758 5146
sangsik.lee@jpmorgan.com
J.P. Morgan Securities (Far East) Limited, Seoul Branch

Neelay Y Kamath
(91-22) 6157 3764
neelay.kamath@jpmchase.com
J.P. Morgan India Private Limited

Mio Shikanai
(81-3) 6736 1313
mio.shikanai@jpmorgan.com
JPMorgan Securities Japan Co., Ltd.

- **The truth behind the SOCAMM content cut for Vera Rubin.** Multiple media outlets highlighted the SOCAMM content downgrade for the VR NVL72 rack from 192GB to 96GB (or from 1.5TB of content to 768GB per CPU) stoking fears among investors on the cooling CPU-related memory demand (link). CPU-driven memory demand upside was the core driver of memory demand accelerating into 2027E (JPMe: DRAM bit demand +33%/ +34% in 2026/27E - see details in our GMM report). Based on our supply chain checks, we sense growing demand for the 96GB SOCAMM module, **while believe that the shift in demand is predominantly due to limited memory supply rather than performance down-spec**. Despite the higher 96GB-grade SOCAMM SKU, we are seeing no changes to overall SOCAMM bit procurement demand, implying that the content downward revision will likely be compensated by higher SOCAMM module demand (e.g. reshuffle of content and volume based on memory supply and content configurations).

Table 1: Sensitivity analysis on NVDA demand based on Vera CPU content configuration

Unit in K, Content in GB

AI headnode	2026E	50% of 96GB and 50% of 192GB	75% of 96GB and 25% of 192GB	100% of 96GB and 0% of 192GB
# of CPU				
Rubin	1,050	1,400	1,680	2,100
Blackwell	2,568	2,568	2,568	2,568
Content per CPU				
Rubin	1,536	1,152	960	768
Blackwell	512	512	512	512
Bit demand by product				
Rubin	1,613	1,613	1,613	1,613
Blackwell	1,315	1,315	1,315	1,315
Bit Demand (8Gb, Mn. Equiv)	**2,928**	**2,928**	**2,928**	**2,928**
Standalone CPU				
# of Vera CPU	600	600	600	600
Content per CPU	768	768	768	768
Bit Demand (8Gb, Mn. Equiv)	**461**	**461**	**461**	**461**
NVDA bit demand	**3,388**	**3,388**	**3,388**	**3,388**

Source: J.P. Morgan estimates.

- **Strengthening partnership between NVDA and SKH.** Nvidia (OW, covered by JPM analyst, Harlan Sur) and SK Hynix announced a multi-year technology partnership to co-develop the next-generation HBM memory for global AI data centers, with SK Hynix formally designated as Nvidia's largest memory partner (2+ year term with extension options). The pact extends well beyond data center HBM, with SK Hynix co-developing memory across Nvidia's full stack namely Vera Rubin supercomputers (JPMe: AI memory products including HBM/DRAM/NAND), Vera CPUs (JPMe: SOCAMM2), RTX Spark PCs (JPMe: LPDDR5X), and Jetson Thor robotics platforms, effectively diversifying SK Hynix into the AI infrastructure, personal AI, and Physical AI markets Nvidia is creating (see more details link). We believe the multi-year partnership implies extended demand visibility (JPMe: likely until 27E) and highlight that this also acts in favor of other existing memory makers given the likely limited supply. In addition, NVIDIA announced a series of partnerships with leading South Korean technology companies to build large-scale AI

See page 4 for analyst certification and important disclosures, including non-US analyst disclosures.

infrastructure in Asia, seeking to solidify its data-center footprint and expand its AI ecosystem (see more details in the table 2 provided below). Lastly, CEO Jensen mentioned (*link*) that 2x DRAM wafer capacity expansion plans over the next five years may not be sufficient, suggesting potentially stronger AI memory demand ahead (consistent with *JPM view*).

Table 2: NVIDIA partnership announcements with major South Korean firms

Partner	Nature of Deal	Key Details
SK Hynix	Multi-year memory technology partnership	Next-gen memory chips for AI data centers; ensures stable supply of advanced memory. Named Nvidia's largest memory partner; 2+ years with extension options.
SK Telecom	Gigawatt-scale AI cloud	Building GW-scale AI cloud on Nvidia tech; first data center online in 2027.
Naver	GW-scale AI factory roadmap	Starts at Gak Sejong data center, expanding capacity then building additional GW-scale AI factories; begins with 55MW project in 2027.
Doosan	Physical AI / robotics	Develops robotics and makes components for Nvidia GPUs; planning several ventures with Nvidia.
Krafton / NC	Gaming to physical AI	Collaboration with game developers for AI and next-gen robotics; Nvidia views gaming as strategic to physical AI / humanoid robots.
Samsung / Hyundai	AI factories	Both companies to invest in AI factories.

Source: NVIDIA newsroom, Reuters and WSJ.

- **Takeaways from Computex Taiwan 2026.** We attended Computex 2026 last week and walked away with positive impressions, with several implications for our sector:

  - **SK Hynix:** We visited SK Hynix's booth and the key takeaways are: 1) Mass production of 12-Hi HBM4 is on track, echoing CEO Jensen's commentary that all three memory makers have qualified HBM4 (link) ; 2) higher interest in 12-Hi 48GB HBM4E which supports up to 16 Gbps vs. 16-Hi HBM4E broadly similar to SEC's recent 12-Hi HBM4E sample shipment to customers (link); 3) Inkling of AI-N B (similar to HBF) which carries throughput like HBM and capacity like SSD and is suitable for inference batches, albeit demand pick-up sounded conservative in 2030 (inline with SKH's HBF demand projection - link).
  - **Solidigm:** We attended the Solidigm Key Note session and the key takeaways were: 1) Storage is becoming a critical infrastructure for various AI data pipelines (i.e. Local SSD (G3) used for data prep, training & inferencing; High performance SSD (G3.5) used in KV cache tier for inferencing; and Shared Storage (G4) used to store inactive KV and for archiving purposes; 2) The value of offloading KV Cache to NAND leads to up to 27x faster time to first token (TTFT) vs. recomputing the context; and 3) Solidigm SSD solutions with VAST's software offers 90% savings in DC space costs vs. HDD (link).
  - **Kioxia:** We attended the Kioxia Key Note session with some information overlapping with the investor day announcement (report). We find the extensive partnership collaboration details between Kioxia and NVDA constructive - Kioxia delivering a cost effective solution for ultra-large scale RAG (Retrieval-Augmented Generation) server mainly leveraging Kioxia's AiSAQ software. The SSD software development was surprising to us given the stickiness of the solution leapfrogging from a hardware vendor. We believe the software partnership and commitment will be a strong differentiating factor for Kioxia against NAND competitors.

- **11% memory share correction offers a buy-into-dip opportunity.** We believe that memory shares corrected 11% in the past week (vs. SOX: -6%) due mainly to the SOCAMM content cut noise and escalating geopolitical issues. We believe the SOCAMM content cut noise is misleading and anticipate Vera CPU volume upside to more than offset the 96GB SOCAMM SKU. Among the multiple partnerships announced between NVDA and SKH, we believe the 2+ year partnership bodes well for extended demand visibility into 2027E which also reads positively to other memory makers. In the NAND space, we reconfirmed eSSD's growing role as the critical infrastructure layer in AGI, which should bode well for its mid-to-long term demand outlook. We do not rule out the possibility of additional memory content cuts or optimizations in the future; however, we view it as a necessary reaction from end customers to handle the severe

Jay Kwon AC
(82-2) 758-5725
jay.h.kwon@jpmorgan.com

shortage of memory supply. We would rather focus on the ongoing customer needs to deploy higher memory content to the system for superior performance. All in all, we see the recent correction as a good buying opportunity from a midterm horizon, and recommend investors add into the pullback and focus on the "higher for longer" upcycle. **Key OW** names in Asia are: **SEC**, **SK Hynix**, and **Kioxia**. Next key catalysts are: a) Long-term supply agreement announcement update by CSP/memory suppliers, b) x86 and ARM CPU related memory sourcing and demand update (incl. Vera CPU), c) 3Q26 contract pricing update, and d) MU May-quarter results and outlook commentary (slated for late June).

Figure 1: Global memory makers' share price performance including SOX (Philadelphia Semiconductor index)

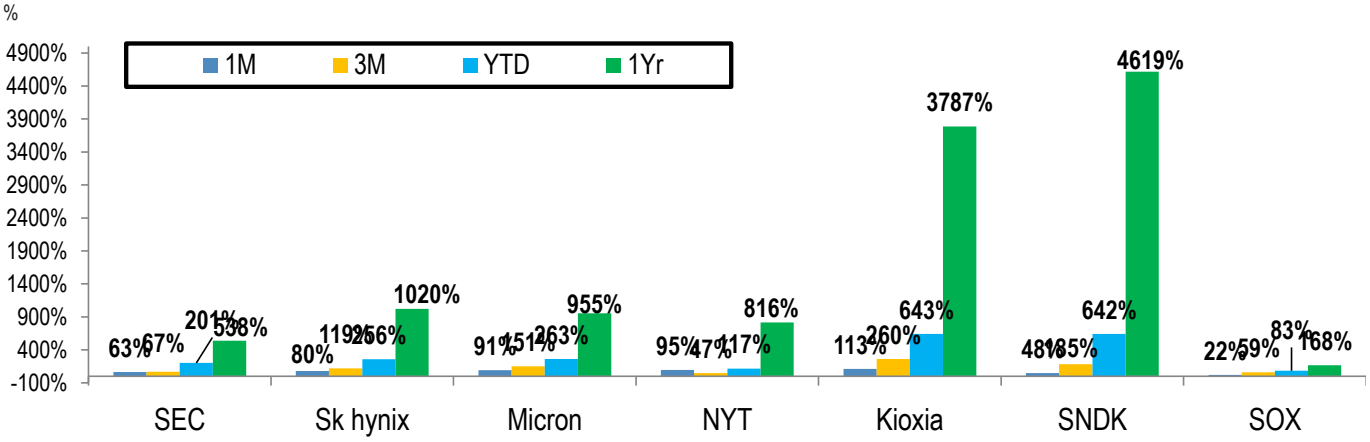

Source: Bloomberg Finance L.P.