

DISCLAIMERS

FORWARD-LOOKING STATEMENTS

This presentation contains forward-looking statements within the meaning of U.S. federal securities laws, including, but not limited to, statements regarding expectations for: the business outlook, financial condition and financial and operational performance of Sandisk Corporation (the “Company”); industry and market trends and expectations, including the expected impact of artificial intelligence (“AI”) and AI inference workloads on demand for memory and storage; the Company’s new business models (“NBMs”) and the scope, terms and expected benefits of the Company’s multi-year customer partnerships; the financial guarantees supporting the Company’s NBMs; the Company’s technology and product roadmap and leadership, including the development, performance, capabilities, timing, commercialization and market adoption of new and future technologies and products, including its BiCS generations, 3D Matrix Memory and High Bandwidth Flash (“HBF™”) technology; the standardization, commercialization, capabilities, benefits, adoption and industry impact of HBF technology and related products; the Company’s technology transitions, bit growth, scaling roadmap and expected cost and capital efficiency; the Company’s financial guidance, financial targets and long-term financial model; the Company’s operational and capital allocation priorities, including its capital return program; and the Company’s value creation for its customers and stockholders. These forward-looking statements are based on management’s current expectations and are subject to risks and uncertainties that could cause actual results to differ materially from those expressed or implied in the forward-looking statements.

Key risks and uncertainties that could cause actual results to differ materially from those expressed or implied in the forward-looking statements include: adverse changes in global or regional economic conditions, including the impact of evolving trade policies, tariff regimes and trade wars; volatility in demand for the Company’s products; pricing trends and fluctuations in average selling prices; inflation; changes in interest rates and a potential economic recession; future responses to and effects of global health crises; the impact of business and market conditions; the impact of competitive products and pricing; the Company’s development and introduction of products based on new technologies and management of technology transitions; risks associated with strategic initiatives, including restructurings, acquisitions, divestitures, cost saving measures and joint ventures; risks related to product defects; difficulties or delays in product ramps, manufacturing or other supply chain disruptions; our reliance on strategic relationships with key partners, including Kioxia Corporation; risks related to our long-term agreements, or “NBMs”; fluctuation of our operating results, including due to changes in demand, industry cycle and timing of customer deployments, and our ability to accurately forecast demand; the attraction, retention and development of skilled management and technical talent; risks associated with the use of artificial intelligence in our business operations; changes to the Company’s relationships with key customers or consolidation among our customer base; compromise, damage or interruption from cybersecurity incidents or other data system security risks; our reliance on intellectual property; fluctuations in currency exchange rates; actions by competitors; risks related to our share repurchase program; risks associated with compliance with changing legal and regulatory requirements; and other risks and uncertainties listed in the Company’s filings with the Securities and Exchange Commission, including our Annual Report on Form 10-K and Quarterly Reports on Form 10-Q, to which your attention is directed. You should not place undue reliance on these forward-looking statements, which speak only as of the date hereof, and the Company undertakes no obligation to update or revise these forward-looking statements to reflect new information or events, except as required by law.

NON-GAAP MEASURES

This presentation includes references to Non-GAAP financial measures, including, among others, Non-GAAP gross margin, Non-GAAP operating margin, Non-GAAP EPS, Adjusted free cash flow and Adjusted free cash flow margin. Reconciliations of the differences between the Non-GAAP measures provided in this presentation to the comparable GAAP financial measures are included in the appendix and at investor.sandisk.com. We have not reconciled our forward-looking Non-GAAP financial measure guidance, targets and long-term model measures to the most directly comparable GAAP measures because material items that impact these measures are not in our control and/or cannot be reasonably predicted. Accordingly, a reconciliation of these forward-looking Non-GAAP financial measures to the corresponding GAAP measures is not available without unreasonable effort.

BASIS OF PRESENTATION

On February 21, 2025, the Company completed its separation from Western Digital Corporation (“WDC”) and became a standalone publicly traded company. The Company’s financial and operating results after the separation are presented on a consolidated basis. For periods prior to the separation, the Company’s historical combined financial statements were prepared on a carve-out basis and were derived from WDC’s consolidated financial statements and accounting records and prepared as if the Company existed on a standalone basis.

SANDISK™
INVESTOR DAY
IN FOCUS 2026

NEXT WAVE OF TECHNOLOGY INNOVATION

ALPER ILKBAHAR
CHIEF TECHNOLOGY OFFICER

SANDISK BREAKING DOWN THE MEMORY WALL



3D Matrix Memory

Next-Generation, High-Density Memory




HBF™

High Bandwidth Flash

SANDISK 3D MATRIX MEMORY DEVELOPMENT UPDATE


SOLID PROGRESS TOWARDS TECHNOLOGY QUALIFICATION



SANDISK™ INVESTOR DAY
IN FOCUS 2026

SANDISK 3D MATRIX MEMORY DEVELOPMENT UPDATE

SOLID PROGRESS TOWARDS TECHNOLOGY QUALIFICATION



SANDISK™ INVESTOR DAY
IN FOCUS 2026

SANDISK BREAKING DOWN THE MEMORY WALL



3D Matrix Memory

Next-Generation, High-Density Memory





HBF™

High Bandwidth Flash

HBF™: SANDISK'S SOLUTION TO THE MEMORY WALL PROBLEM

- Same read bandwidth as HBM with up to 8-16x capacity
- Invented for AI inference workloads with large MoE based LLMs, long context lengths, and large KV Caches
- Architecture developed with inputs from major cloud/AI customers

HBF



NOTE: Image a representation only – not real image.

©2026 SANDISK CORPORATION OR ITS AFFILIATES ALL RIGHTS RESERVED


RECENT AI MODEL TRENDS

THE NEW AI PARADIGM: MEMORY-CENTRIC INFERENCE

Models	Released	Total Parameters	Mixture of Experts	Active Parameters	Context Length
Grok 4	Jul 9, 2025	~1.7T [†]	Dense (Est.) [†]	~1.7T [†]	256K
GLM 5.2	Jun 13, 2026	744B	256 (8 Active) [†]	40B	1M
GPT 5.6	Jul 9, 2026	~1.7-1.8T [†]	No Public Est.	No Public Est.	1.05M
Kimi K3	Jul 16, 2026	2.8T	896 (16 Active)	~50B [†]	1M


[†] Best-available published estimate, not an official figure - actual specs unconfirmed. Anthropic, OpenAI and xAI do not publish full architecture details.

Estimates - published third-party analyses:
Grok 4 ~1.7T dense (whitex.ai, apxml); Claude Mythos 5 ~10T leak (Medium, AI Magicx); GLM 5.2 256 experts / 8 active (glm5.app, DeepWiki)
GPT-5 ~1.7-1.8T dense-equiv (CometAPI, A2E); Kimi K3 ~50B active (Hugging Face). Official context: GPT-5.6 1.05M (OpenAI), Grok 4 256K (xAI)



RECENT AI MODEL TRENDS

THE NEW AI PARADIGM: MEMORY-CENTRIC INFERENCE



Higher Memory Capacity
(Model Size, Context Lengths)



Lower Compute Requirements
(MoE Architecture)



Memory-Centric AI

HBF™ AS INNOVATION PLATFORM

HBF ARCHITECTURE PROPOSALS FROM ECOSYSTEM

HBF Only



HBF Can Replace HBM Stacks to Add More Read-Optimized Memory in a Similar Footprint

HBM/HBF Cached



HBM is Used for Tiered Cache and HBF Can Store Weights and KV Cache


HBF (+HBM)

(N) stack positions with (k) HBFs and (N-k) HBMs



HBF Can Augment HBM Stacks to Add More Read-Optimized Memory in a Similar Footprint

Disaggregated



HBF is Used for Decode Weights and KV Cache Memory

HBF™ vs HBM SIMULATION CONFIGURATIONS

		VS	
Architecture	HBF Only		HBM Only
GPU Perf	4.5 PFLOPS		
HBM or HBF Total Bandwidth/GPU	12.8TB/s		
GPU to GPU Bandwidth	900GB/s		
HBM or HBF Total Capacity/GPU	4TB		192GB
Per GPU Connection to KV Spillover Memory	128GB/s to LPDDR DRAM		128GB/s to LPDDR DRAM
Number of GPUs	1,4		8
Model	Qwen3-480B-A35B (960GB in bfloat16)		
Workload	Multi-Turn Agentic AI Workload Emulation With Log-Normal Inter-Turn Arrival Time		


HBF™ vs HBM PERFORMANCE COMPARISON FOR AI INFERENCE



HBF Only

1 GPU

4 GPU



HBM Only

8 GPU

8x

CAPEX EFFICIENCY
for minimum configuration
required to run model :
1 HBF GPU vs 8 HBM GPUs¹




2x

GPU EFFICIENCY
Half the number of HBF GPUs
deliver the same token output :
4 HBF GPUs = 8 HBM GPUs¹



HBF™ ROADMAP

2ND-GEN HBF TECHNOLOGY FOR THE EDGE



“AI THAT NEVER FORGETS”

Enabling 100B+ Models on Your *Edge Device*


*BASED ON INTERNAL ESTIMATES

Actively being co-developed with multiple customers

SANDISK™ INVESTOR DAY
IN FOCUS 2026

©2026 SANDISK CORPORATION OR ITS AFFILIATES ALL RIGHTS RESERVED

HBF™ CONSORTIUM MILESTONES



A BOARD BUILT FOR THE AI MEMORY ERA

HBF™ TECHNICAL ADVISORY BOARD



DAVID PATTERSON

Pardee Professor Emeritus, UC Berkeley •
Fellow, Google

Pioneer of RISC and modern
computer architecture



RAJA KODURI

Founder & CEO, Oxmiq Labs •
Co-Founder, Mihira Visual Studios

GPU and visual-compute visionary

NEW MEMBER



JIM KELLER

CEO, Tenstorrent

Legendary chip architect bringing AI-inference
silicon and systems expertise to HBF

SANDISK™ INVESTOR DAY
IN FOCUS 2026

A BOARD BUILT FOR THE AI MEMORY ERA

HBF™ TECHNICAL ADVISORY BOARD



NEW MEMBER

JIM KELLER

CEO, Tenstorrent

DEC ALPHA / AMD K7 & K8 / APPLE A4 & A5 / AMD ZEN / TESLA FSD / INTEL / TENSTORRENT

Four decades of silicon leadership — lead architect behind AMD Zen, Apple's A4 and A5, Tesla's Full Self-Driving chip, and the DEC Alpha

Focused on AI inference at every scale — Tenstorrent's expertise spans AI inference across multiple verticals, from hyperscale data centers to power-constrained edge devices


SANDISK™ INVESTOR DAY
IN FOCUS 2026

©2026 SANDISK CORPORATION OR ITS AFFILIATES ALL RIGHTS RESERVED

TENSTORRENT BUILDS AI COMPUTERS AND RISC-V CPU_s

AI Tensix Cores

- Inference and Training
- Fully Open Source Software Stack with Multiple Entry Points
- Access to Bare Metal (TT-Metalium™)
- Seamless Scaling




RISC-V CPU

- World's Highest Performing Open Standard CPU
- Scalable CPU Design
- Optimized for AI Architectures
- Customizable and Modular for Specific User Workloads

No switches, routers, host computers, or expensive HBM → Ethernet & SRAM

Data flow architecture allows for optimal place-and-route so operations occur in the right place and time

TENSTORRENT GALAXY BLACKHOLE IS IN PRODUCTION AND BUILT FOR SCALE




Highly Flexible Configuration Enables:

- Enable workloads to leverage high-throughput
- Realtime applications leveraging low-latency
- Realtime Video Generation

24 PetaFLOPS



1 TB Capacity, 16 TB/s Bandwidth



11.2 TB/s I/O

"AI Switch"



Tenstorrent uniquely supports both what GPUs and other AI Accelerators offer in a seamless integrated system without the additional overhead

INDUSTRY LEADING PERFORMANCE AT SCALE

DeepSeek R1
400+ t/s/u

supporting batch sizes from 8 to 64 and up to 128k context with DeepSeek-R1-0528 671B

DeepSeek R1
<4 second

time-to-first-token on 100K context with DeepSeek-R1-0528 671B

Wan 2.2
10x faster

than real-time high-quality video generation: 720p, 81-frame video in 2.4 seconds

New

Kimi K2.6
900 t/s/u

Kimi K2.6
100k Context,
3x faster than GPUs

New

LTX 2.3 Fast
4x faster

than GPUs
144 frames,
24 fps,
6 sec video gen with LTX 2.3 Fast


We Run Everything

92% of HuggingFace models, compiling and running on our open-source compiler, TT-Forge today

Run **2.4M+**
models with our compiler

HBF™ ROADMAP

FIRST HBF MEMORY DIE TAPED OUT



FIRST HBF INFERENCE PRODUCT SAMPLES

COMING
SOON
2027