



Global AI Trend Tracker

Global Markets Research

27 May 2026

EQUITY: TECHNOLOGY

Google's next-gen AI network for TPUv8t / v8i...

...And implications for the AI networking value chain

TPU v8t & v8i's networking architecture may lead to stronger demand for OCS and benefit Google's TPU supply chain

Google (GOOG US, Not rated) released its eighth-generation tensor processing unit (TPU) at its Cloud Next 2026 in April, and the company for the first time distinguished its TPU products into training (TPU 8t) and inference (TPU vi). TPU 8t improves large-scale pre-training performance with SparseCore cores and Virgo network topology, while TPU 8i is designed for real-time reasoning and complex decision-making, with CAE (Collectives Acceleration Engine) and new Boardfly topology solving the latency bottleneck of long context reasoning to a certain extent. We think the purpose-built networks for TPUv8t and v8i reflect the AI market trend of decoupling network architecture for training and inference, in order to better unlock their performance in different scenarios, while reducing cost and power consumption at the same time. We think **Boardfly network brings incremental demand for optical circuit switches (OCS0 at both the scale-up and the scale-out level)**. Meanwhile, as a technology leader in optical networks, we expect Google's stronger TPU shipments of would accelerate its adoption of all the mainstream optical communication solutions, including **1.6T pluggable transceivers, NPO (near-field packaged optics) and CPO (co-packaged optics)**. We think this is **positive for Google's TPU supply chain, including substrate and PCB suppliers such as VGT (300476 CH/2476 HK, Buy), as well as optical communication solution providers.**

Google introduced Virgo Network and Boardfly topology for TPU 8t and 8i, respectively, aiming at lower latency and higher bandwidth

Regarding networking solutions, TPU 8t still uses 3D Torus topology (similar to TPU v7) to connect chips at the scale-up level to form a pod with up to 9,600 cards, and we think that there are still 48 OCS in each pod, but the number of ports may increase to 300*300 (from 288*288). Google introduced Virgo Network, a flat two-layer non-blocking topology, for TPU V8t scale-out network. According to the company, this new network architecture enables up to 4x increased data center network (DCN) bandwidth and uses high-radix switches that reduce network layers by allowing more ports per switch. Moreover, Google has designed a new network topology, called Boardfly, for the TPU 8i cluster that can connect up to 1,152 of chips. The network connects 8 boards via copper cabling and further scales to 36 groups (up to 1,024 active chips) linked through OCS.

Growth of CPU demand accelerates in agentic era

In addition, both eighth-generation TPU chips will be working with Google's self-developed ARM-based Axion CPU as the main controller. We think CPUs may become the key to running large AI clusters in the future, due to the exponential growth of inference workloads. According to comments from Intel's CEO during the company's 1Q26 earnings call, the CPU-to-GPU attach rate used to be 1:8, which has increased to 1:4 now, and may increase to 1:1 in the future. We think this reflects growing demand for CPUs in AI clusters for the agentic era.

Research Analysts

China Technology

Bing Duan - NIHK

bing.duan1@nomura.com

+852 2252 2141

Joel Ying, CFA - NIHK

joel.ying@nomura.com

+852 2252 2153

Introduction to TPU V8i & V8t

On 22 April 2026, Google released its eighth-generation TPU at the Cloud Next 2026 conference, which is divided into the TPU 8t training chip and the TPU 8i inference chip. From the perspective of Google's chip development path, the current divergence between training and inference architecture requirements is widening, with training clusters focusing on optimizing computing power and inference clusters focusing on optimizing HBM bandwidth and memory capacity. The current mainstream large language models (LLM) are all MoE architectures, which require long context reasoning and the chain of thought scheme, and the importance of communication efficiency, memory access speed, and low-latency synchronization between chips far exceeds that of single-chip computing power. Therefore, Google designed the training chip and the inference chip separately and optimized their internal architectures to achieve high computing performance and low latency for the training chip, and high memory and low power consumption for the inference chip.

Fig. 1: Development path of Google TPU

	TPU	TPU chips per superpod	Topology	ICI bandwidth per TPU chip	ICI optical transceiver	Optical lane rate	OCS
2018	v2	256	2D Torus	800GB/s	None	NA	None
2020	v3	1024	2D Torus	800GB/s	400G AOC	50G	None
2022	v4	4096	3D Torus	600GB/s	400G OSFP	50G	OCS
2023	v5p	8960	3D Torus	1200GB/s	800G OSFP	100G	OCS
2025	v7	9216	3D Torus	1200GB/s	800G OSFP	200G	OCS
2026	V8t	9600	3D Torus	2400GB/s	NA	400G	OCS
2026	V8i	1152	Boardfly	2400GB/s	NA	400G	OCS

Source: Google, Nomura research

TPU 8t improves large-scale pre-training performance with SparseCore and Virgo network topology. TPU 8i is designed for real-time reasoning and complex decision-making, and its CAE acceleration engine and new Boardfly topology solve the latency bottleneck of long context reasoning to a certain extent, allowing AI to evolve from single-next-word prediction to scene simulation and deep logical reasoning, and AI responses will become more timely and coherent. Supported by the computing power of Google's self-developed ARM Axion architecture CPUs, Boardfly topology has achieved a double leap in energy efficiency.

The SparseCore is a specialized accelerator designed to handle the irregular memory access patterns of embedding lookups. The MoE LLM only activates a small number of parameters each time, and although the hybrid expert technology is highly energy efficient, it produces a large number of irregular memory accesses. SparseCore technology is specifically designed to handle this task, working with the Matrix Multiply Unit (MXU) to keep the chip running at high load. Moreover, TPU 8t minimizes exposed vector operation time by implementing more balanced Vector Processing Unit (VPU) scaling. In terms of computing, TPU 8t introduces native 4-bit floating point (FP4) to overcome memory bandwidth bottlenecks, doubling MXU throughput while maintaining accuracy for LLMs even at lower-precision quantization. In terms of memory, TPU 8t adopts de-intermediated TPU Direct technology, and HBM memory can directly talk to the Network Interface Card (NIC), bypassing the CPU and DRAM, resulting in 10x faster memory access speed than the previous generation TPU. Overall, TPU 8t solves several key bottlenecks in training tasks: efficiently processing sparse data, reducing data handling overhead between memory and compute units, and optimizing network communication between large-scale chip clusters. Google claims that compared with the previous generation, the "performance per dollar" of the 8T in training scenarios can be improved by up to 2.7x.

The diagram illustrates the Zebra architecture, showing the connection between a Host and a Zebra package. The Host is connected to the Zebra package via PCIe Gen5 (x16) and gBMC. The Zebra package contains a TensorCore (TCS) with VPU + Vmem, connected to XLU and MXU blocks. It also includes Sparse cores, a Chip manager, and a Memory and DMA interconnect. The interconnect is connected to HBM3E Ctrl and HBM3E stack blocks. The Zebra package is connected to an ICI (Interconnect Interface) block, which is connected to an ICR Router and 6x224G SerDes octals. A legend on the right defines the components: Compute (blue), Memory and DNA (grey), Interconnect (light blue), Host interface (dark blue), Management (dark grey), Logic chiplet (light grey), SerDes chiplet (dashed box), and Zebra package (white box).

TPU 8i is a new chip designed for inference tasks, with the most notable feature being the tilt of resource allocation towards memory. Although its peak hash rate is lower than that of 8t, it is equipped with a 384MB on-chip SRAM cache (3x that of 8t), higher HBM memory capacity (288GB), and greater bandwidth. Google claimed that TPU 8i can host a larger KV Cache entirely on silicon, significantly reducing the idle time of the cores during long-context decoding. In addition, when an LLM performs inference, chips need to frequently synchronize data and summarize results, a process called collective communication. Chains-of-thought require multiple computing cores to frequently synchronize intermediate results, and traditional synchronization operations have extremely high latency. Therefore, TPU 8i uses the CAE, which aggregates results across cores with near-zero latency, specifically accelerating the reduction and synchronization steps required during auto-regressive decoding and chain-of-thought processing. For each TPU 8i chip, there are two Tensor Cores (TC) on-core dies and one CAE on the chiplet die, replacing four SparseCores (SCs) on core dies in the previous-generation Ironwood TPU. CAE further reduces the on-chip latency of collectives by 5x. Lower latency per collective operation means less time spent waiting, directly contributing to higher throughput required to run millions of agents concurrently.

Fig. 3: TPU 8i ASIC block diagram



Source: Google, Nomura research


Fig. 4: Comparison of TPU v7, TPU v8t and TPU v8i

	TPU v7	TPU 8t	TPU 8i
Primary workload	High-performance ASIC for training and inference	Focus on large-scale training	Focus on large-scale inference
Design of die	Dual die	Single die	Dual die
HBM capacity	192GB HBM3E	216GB HBM	288GB HBM
HBM bandwidth	7.3TB/s	6.5TB/s	8.6TB/s
on-chip SRAM	-	128MB	384MB
Core feature	SparseCore, shared memory	SparseCore and LLM decoder engine, native FP4, Vrigo network	CAE, large SRAM, Boardfly network
Pod size	9216	9600	1152
Network topology	3D Torus	3D Torus	Boardfly
Peak FP4 PFLOPs	4.6	12.6	10.1

Source: Google, Nomura research

With the explosive growth of inference workloads, CPUs are the key to determining the efficiency and cost of AI systems. Data orchestration and memory management of inference tasks are highly dependent on the CPU, and the "logical scheduling" and "serial processing" used by AI Agents to perform complex tasks such as reading databases, running code, and parsing documents are the expertise of CPUs. According to comments from the Intel CEO during the company's 1Q26 earnings call, the CPU-to-GPU attach rate used to be 1:8, which has increased to 1:4 now, and may increase to 1:1 in the future. We think this reflects growing demand for CPU in AI clusters for agentic era.

Axion is Google's first custom ARM architecture-based CPU designed specifically for data centers, announced in 2024. For AI inference, Axion's dedicated optimizations can significantly boost performance, enabling AI workloads to run faster and more efficiently,




according to management. The Axion product portfolio now includes three options: N4A, C4A, and C4A metal. Compared to virtual machines using mainstream x86 architecture CPUs, N4A offers twice the cost-performance ratio and leads in energy efficiency by up to 80% per watt, according to Google. In the TPU v8 series, Google is the first to use the Axion CPU as the main computing header, effectively reducing data preprocessing latency and ensuring the TPU computing engine runs at full capacity.

Fig. 5: AI inference performance on Axion CPU

AI Inference Performance on Google Axion

Prompt Encoding & Tokens Generated /Sec



+29%

Better Performance over
AMD Turin
c4d-highmem-48 vCPU

+23%

Better Performance over
Intel Granite Rapids
c4-highmem-48 vCPU

Model Format: w8a8 | Benchmark: Throughput Serving | Serving Framework: vLLM | Concurrency: 1 and 16


(Compared to x86 based Instances)

Source: Google, Nomura research

TPU V8i & V8t networking solutions

TPU 8t still uses 3D Torus topology to connect chips at the scale-up level to form a pod. The pod volume has increased to 9,600 cards (vs. 9216 cards for Ironwood), and the single-card scale-up bidirectional bandwidth is 19.2Tb (2x that of Ironwood), and we think that there are still 48 OCS in each pod, but the number of ports may increase to 300*300. To support the massive data requirements of TPU 8t, Google introduced Virgo Network for scale-out networking. This new networking architecture enables up to 4x increased data center network (DCN) bandwidth on TPU 8t training over DCN. Built on high-radix switches that reduce network layers by allowing more ports per switch, Virgo Network employs a flat, two-layer non-blocking topology. Compared with traditional datacenter networks, this significantly reduces latency by minimizing network tiers. It features a multi-planar design with independent control domains to connect TPU 8t chips. The TPU 8t racks also connect with the Jupiter north-south fabric to access compute and memory services. Virgo Network can connect more than 1 million TPU chips into a single training cluster, with a single network architecture connecting 134,000 chips.

Fig. 6: Virgo Network architecture




Source: Google, Nomura research


Google designed TPU 8t with a new serving optimized network topology called Boardfly. Boardfly is a flat two-layer architecture, significantly reducing the latency. Utilizing a high-radix design, Google connects up to 1,152 of these chips together, reducing the network diameter and the number of hops a data packet must take to cross the system. By slashing the hops required for all-to-all communication, Boardfly achieves up to a 50% improvement in latency for communication-intensive workloads, according to Google. For that same 1,024-chip pod, Boardfly reduces the network diameter from 16 hops down to 7, which directly leads to lower tail latency.

Boardfly consists of the following elements, and its topology is hierarchical by nature:

1. Building Block (BB): Each tray forms a **four-chip** ring using internal ICI links (**connected by PCB**), providing **16** external connections for broader networking.
2. Group (G): **8 boards** are fully connected via **copper cabling** to create a localized group, utilizing **11** of the available external links for intra-group communication.
3. Pod structure: The final architecture scales to **36 groups** (up to 1,024 active chips) linked through **Optical Circuit Switches (OCS)**, ensuring a maximum latency of seven hops for any chip-to-chip communication.

Fig. 7: TPU 8i hierarchical Boardfly topology

Source: Google, Nomura research

Fig. 8: ICI network diameter via optical circuit switch on TPU 8i pod

Source: Google, Nomura research

Impact of TPU v8 on AI networking

The scale-up networking of TPU 8t continues the 3D Torus topology, and the pod volume has increased to 9,600 cards. The number of copper cables has also increased accordingly, while the optical transceiver ratio is about 1:1.5, according to Semi analysis, but the absolute number increases with the number of TPUs.

Fig. 9: TPU v7 3D Torus connection solutions

TPU v7 3D Torus Connection Solutions		
	Per TPU	Rack Total
8 Interior TPUs		
Copper cables	4	32
PCB Traces	2	16
Optical Transceivers	0	0
8 Corner TPUs		
Copper cables	1	8
PCB Traces	2	16
Optical Transceivers	3	24
24 Edge TPUs		
Copper cables	2	48
PCB Traces	2	48
Optical Transceivers	2	48
24 Face TPUs		
Copper cables	3	72
PCB Traces	2	48
Optical Transceivers	1	24
Total for 64 TPU Rack		
Copper cables	1.25	80
PCB Traces	1.00	64
Optical Transceivers	1.50	96

Source: SemiAnalysis, Nomura research

The design of Google TPU 8i shows that in addition to model training requiring high-density computing clusters, large-scale inference involving multi-agent collaboration and complex chain-of-thought is also communication-intensive, which increases the requirements for low latency between chips. Judging from Google's new Boardfly topology, copper cables are still the main interconnect method at the ICI level, and the demand for OCS extends from scale-out to scale-up between chipsets.

Fig. 10: Calculation of usage of DAC/optical transceiver/OCS in TPU 8i pod

	BB	Group	Between group
DAC	8x11x36=3168	-	-
Optical transceiver	-	5x8x36=1440	-
OCS	-	-	20 72*72 OCS / 5 288*288 OCS

Source: Nomura estimates

Theoretically, the switching response speed of scale-up networking for OCS is required to

reach nanoseconds. The MEMS-based OCS relies on the mechanical motion of the micromirror and the DLC-based OCS relies on the rearrangement of liquid crystal molecules, both of which have a switching time in the millisecond range, which cannot meet this requirement in terms of physical principle. However, the SiPh-based OCS can achieve nanosecond-level switching by virtue of the principle of electronically controlled refractive index, which matches the requirements of the scale-up networking, but it is still in the small-batch delivery stage. As a representative manufacturer of SiPh-based OCS, iPrionics (unlisted) has released a 32*32-port SiPh-based OCS in 2025 and expects to release a 64*64-port OCS in 2026. Meanwhile, it has cooperated with Fabrinet (FN US, Not rated) to build the world's first mass-produced SiPh-based OCS production line, which should enter operation from 2Q26, according to company. The new manufacturing line enables iPrionics to scale production rapidly to meet customer demand for SiPh-based OCS in training and inference AI clusters. On the other hand, although MEMS-based OCS switching speed is relatively slow, Google has made up for its switching delay through software scheduling, which is also the current strategy adopted by Google. Therefore, we believe that both the solutions have the potential to be used in TPU v8 clusters, and we will continue to catch up the latest progress in SiPh-based OCS.

In addition, both the eighth-generation TPU chips are equipped with Google's self-developed ARM architecture Axion CPU as the main controller, which solves the host computing power bottleneck caused by data preprocessing delays. The number of CPUs has been upgraded from "1 CPU with 4 TPUs" to "1 CPU with 2 TPUs", and the number of CPU hosts per server has doubled. The CEO of Arm, Rene Haas, mentioned in latest earnings call that as AI shifts from "human-computer interaction" to "agentized continuous workloads", data center CPU capacity will more than quadruple and the market size will exceed USD100bn by 2030. We think that the rapid growth of demand for CPUs will drive revenue growth of CPU PCB suppliers, server foundries and other related supply chain manufacturers.

Google's TPU v8 series chips adopt a dual-vendor strategy, with the TPU 8t being co-designed by Google and Broadcom (AVGO US, Not rated), and the TPU 8i being designed by Google and MediaTek (2454 TT, Buy; covered by Aaron Jeng) for the first time. Broadcom has rich experience in high-performance ASIC design, and the pursuit of extreme computing performance and interconnect bandwidth in training chips is Broadcom's strength, while MediaTek has advantages in power efficiency and cost optimization. Compared to single-chip performance, inference chips need to provide higher inference throughput per-unit power consumption. MediaTek's low-power design experience in the mobile SoC market could be well used for the inference chip.

Based on Counterpoint Research's latest Global AI Server Compute ASIC Shipment Forecast, the combined TPU v8t and v8e shipments are projected to approach 5mn units in 2028, more than 10x growth from the ~400k units shipped in 2026. We believe that the separation of TPU v8 training and inference chips and strong demand will likely help further expand Google's TPU supply chain.

Fig. 11: Google PCB and CCL supply chain

Generation	Content	Time	Structure	CCL Material	CCL Supplier(s)	PCB Supplier(s)
TPU v8p	TPU UBB	2H25~	34L PCB (16+18, N+M)	M7, HVLP2	Panasonic, EMC	WUS, ISU, TTM, VGT
	CPU board (New)	4Q25~	22L PCB	M8	EMC, Panasonic	VGT, WUS, ISU, TTM, others?
TPU v7p, v7e	TPU UBB	mid-26~	36~40L+ PCB	M8+M8, HVLP3	Panasonic, EMC	WUS, ISU, TTM, LCS, VGT?
	CPU board (New)	4Q25~	22L PCB	M8	EMC, Panasonic	VGT, WUS, ISU, TTM, others?
	PCIe switch (New)	mid-26~	22-24L PCB	M8?	EMC	WUS, ISU, GCE, ZD T, others?
TPU v8t, v8i	TPU UBB?	2028?	24~26L? HDI?	M8.5? M8Q?	Panasonic? EMC?	WUS, ISU, Unimicron, others?
	More other boards?					

Source: Nomura research

Fig. 12: OCS industry supply chain

MEMS				
Core components	MEMS array	Optical filter	Fiber collimator array	Lens array
Global supplier	Silex		Corning	
Domestic supplier	Yitao Intelligent	Optowide Technologies, DOTI Micro	TFC, T&S, EverProX, Optowide Technologies	Focuslight Technologies
OCS manufacturer/foundry	Lumentum, Accelink, Eoptolink, Triple-stone, Calient, Advanced Fiber Resources (foundry), Taclink (foundry)			
DLC				
Core components	Liquid crystal light module	Fiber collimator array	YVO ₄ crystal	
Global supplier	Coherent	Corning		
Domestic supplier		TFC, T&S, EverProX, Optowide Technologies	Castech, Optowide Technologies	
OCS manufacturer/foundry	Coherent, Celestica (foundry)			
DLBS				
Core components	Piezoelectric Ceramics module	Fiber collimator array		
Global supplier	Polatis	Corning		
Domestic supplier	Luster	TFC, T&S, EverProX, Optowide Technologies		
OCS manufacturer/foundry	Polatis, Luster			
SiPh-based solution				
Core components	SiPh chip	PIC	MPO	PCB
Global supplier	iPronic		Corning	
Domestic supplier			T&S, EverProX	
OCS manufacturer/foundry	iPronic, Taclink			

Source: Company data, Nomura research

Appendix A-1

This report has been produced by Nomura International (Hong Kong) Ltd. (NIHK), Hong Kong.

See [Disclaimers](#) for Nomura Group entity details.

Analyst Certification

I, Bing Duan, hereby certify (1) that the views expressed in this Research report accurately reflect my personal views about any or all of the subject securities or issuers referred to in this Research report, (2) no part of my compensation was, is or will be directly or indirectly related to the specific recommendations or views expressed in this Research report and (3) no part of my compensation is tied to any specific investment banking transactions performed by Nomura Securities International, Inc., Nomura International plc or any other Nomura Group company.

Issuer Specific Regulatory Disclosures

The terms "Nomura" and "Nomura Group" used herein refer to Nomura Holdings, Inc. and its affiliates and subsidiaries, including Nomura Securities International, Inc. ('NSI') and Instinet, LLC ('ILLC'), U. S. registered broker dealers and members of SIPC.


Materially mentioned issuers

Issuer	Ticker	Price	Price date	Stock rating	Sector rating	Disclosures
		TWD				
MediaTek	2454 TT	4,265.00	26-May-2026	Buy	N/A	
Victory Giant	2476 HK	HKD 436.80	27-May-2026	Buy	N/A	
Victory Giant	300476 CH	CNY 379.30	27-May-2026	Buy	N/A	

MediaTek (2454 TT)

TWD 4,265.00 (26-May-2026) Buy (Sector rating: N/A)

Rating and target price chart (three year history)



For explanation of ratings refer to the stock rating keys located after chart(s)

Valuation Methodology Our TP of TWD3,400 is based on 25x our average 2027-28F EPS. Our target multiple of 25x is at its high end of historical range. The benchmark index for this stock is TAIEX.

Risks that may impede the achievement of the target price Key downside risks include: 1) fierce price competition from Qualcomm and Spreadtrum; 2) the company's execution (i.e. a continuous rollout of good products in terms of specification, price and cost); 3) smartphone demand, especially in China and emerging markets, where MediaTek has higher revenue exposure; and 4) ASIC execution and competition.

Victory Giant (2476 HK)

HKD 436.80 (27-May-2026) Buy (Sector rating: N/A)

Rating and target price chart (three year history)



For explanation of ratings refer to the stock rating keys located after chart(s)

Valuation Methodology Our TP of HKD479.00 is based on 27x 2027F EPS of CNY15.44, in line with its A-share historical median P/E. The benchmark index is Hang Seng Index.

Risks that may impede the achievement of the target price Downside risks: 1) lower-than-expected PCB demand in downstream sectors such as servers and auto electronics; 2) more fierce competition in the high-end PCB market leading to pressure on margins; 3) higher-than-expected raw material cost pressure, and 4) worse-than-expected geopolitical tensions in global AI value chain.

Victory Giant (300476 CH)

CNY 379.30 (27-May-2026) Buy (Sector rating: N/A)

Rating and target price chart (three year history)



For explanation of ratings refer to the stock rating keys located after chart(s)

Valuation Methodology Our target price of CNY417.00 is based on 27x 2027F EPS of CNY15.44, in line with company's historical median P/E of 27x. The benchmark index is CSI300.

Risks that may impede the achievement of the target price Downside risks include: 1) key customers' supply chain diversification and more intensified competition from peers; 2) technology changes such as COWOP which may change the competition landscape and threat company's market shares; 3) slower technology upgrade in AI PCB market; and 4) escalations on geopolitical tensions.

Important Disclosures**Online availability of research and conflict-of-interest disclosures**

Nomura Group research is available on www.nomuranow.com/research, Bloomberg, Capital IQ, Factset, LSEG.

Important disclosures may be read at <http://go.nomuranow.com/research/m/Disclosures> or requested from Nomura Securities International, Inc.

If you have any difficulties with the website, please email grpsupport@nomura.com for help.

The analysts responsible for preparing this report have received compensation based upon various factors including the firm's total revenues, a portion of which is generated by Investment Banking activities. Unless otherwise noted, the non-US analysts listed at the front of this report are not registered/qualified as research analysts under FINRA rules, may not be associated persons of NSI, and may not be subject to FINRA Rule 2241 restrictions on communications with covered companies, public appearances, and trading securities held by a research analyst account.

Nomura Global Financial Products Inc. (NGFP) Nomura Derivative Products Inc. (NDP) and Nomura International plc. (Nlplc) are registered with the Commodities Futures Trading Commission and the National Futures Association (NFA) as swap dealers. NGFP, NDPI, and Nlplc are generally engaged in the trading of swaps and other derivative products, any of which may be the subject of this report.

Distribution of ratings (Nomura Group)

The distribution of all ratings published by Nomura Group Global Equity Research is as follows:

57% have been assigned a Buy rating which, for purposes of mandatory disclosures, are classified as a Buy rating; 34% of companies with this rating are investment banking clients of the Nomura Group*. 0% of companies (which are admitted to trading on a regulated market in the EEA) with this rating were supplied material services** by the Nomura Group.

41% have been assigned a Neutral rating which, for purposes of mandatory disclosures, is classified as a Hold rating; 57% of companies with this rating are investment banking clients of the Nomura Group*. 0% of companies (which are admitted to trading on a regulated market in the EEA) with this rating were supplied material services by the Nomura Group

2% have been assigned a Reduce rating which, for purposes of mandatory disclosures, are classified as a Sell rating; 0% of companies with this rating are investment banking clients of the Nomura Group*. 0% of companies (which are admitted to trading on a regulated market in the EEA) with this rating were supplied material services by the Nomura Group.

As at 31 March 2026.

*The Nomura Group as defined in the Disclaimer section at the end of this report.

** As defined by the EU Market Abuse Regulation

Definition of Nomura Group's equity research rating system and sectors

The rating system is a relative system, indicating expected performance against a specific benchmark identified for each individual stock, subject to limited management discretion. An analyst's target price is an assessment of the current intrinsic fair value of the stock based on an appropriate valuation methodology determined by the analyst. Valuation methodologies include, but are not limited to, discounted cash flow analysis, expected return on equity and multiple analysis. Analysts may also indicate expected absolute upside/downside relative to the stated target price, defined as (target price - current price)/current price.

STOCKS

A rating of **'Buy'**, indicates that the analyst expects the stock to outperform the Benchmark over the next 12 months. A rating of **'Neutral'**, indicates that the analyst expects the stock to perform in line with the Benchmark over the next 12 months. A rating of **'Reduce'**, indicates that the analyst expects the stock to underperform the Benchmark over the next 12 months. A rating of **'Suspended'**, indicates that the rating, target price and estimates have been suspended temporarily to comply with applicable regulations and/or firm policies. Securities and/or companies that are labelled as **'Not rated'** or shown as **'No rating'** are not in regular research coverage. Investors should not expect continuing or additional information from Nomura relating to such securities and/or companies. Benchmarks are as follows: **United States/Europe/Asia ex-Japan**: please see valuation methodologies for explanations of relevant benchmarks for stocks, which can be accessed at: <http://go.nomuranow.com/research/m/Disclosures>; **Global Emerging Markets (ex-Asia)**: MSCI Emerging Markets ex-Asia, unless otherwise stated in the valuation methodology; **Japan**: Russell/Nomura Large Cap.

SECTORS

A **'Bullish'** stance, indicates that the analyst expects the sector to outperform the Benchmark during the next 12 months. A **'Neutral'** stance, indicates that the analyst expects the sector to perform in line with the Benchmark during the next 12 months. A **'Bearish'** stance, indicates that the analyst expects the sector to underperform the Benchmark during the next 12 months. Sectors that are labelled as **'Not rated'** or shown as **'N/A'** are not assigned ratings. Benchmarks are as follows: **United States**: S&P 500; **Europe**: Dow Jones STOXX 600; **Global Emerging Markets (ex-Asia)**: MSCI Emerging Markets ex-Asia. **Japan/Asia ex-Japan**: Sector ratings are not assigned.

Target Price

A Target Price, if discussed, indicates the analyst's forecast for the share price with a 12-month time horizon, reflecting in part the analyst's estimates for the company's earnings. The achievement of any target price may be impeded by general market and macroeconomic trends, and by other risks related to the company or the market, and may not occur if the company's earnings differ from estimates.

Disclaimers

This publication contains material that has been prepared by the Nomura Group entity identified on page 1 and, if applicable, with the contributions of one or more Nomura Group entities whose employees and their respective affiliations are specified on page 1 or identified elsewhere in this publication. The term "Nomura Group" used herein refers to Nomura Holdings, Inc. and its affiliates and subsidiaries including: (a) Nomura Securities Co., Ltd. ('NSC') Tokyo, Japan, (b) Nomura Financial Products Europe GmbH ('NFPE'), Germany, (c) Nomura International plc ('Nlplc'), UK, (d) Nomura Securities International, Inc. ('NSI'), New York, US, (e) Nomura International (Hong Kong) Ltd.

('NIHK'), Hong Kong, (f) Nomura Financial Investment (Korea) Co., Ltd. ('NFIK'), Korea (Information on Nomura analysts registered with the Korea Financial Investment Association ('KOFIA') can be found on the KOFIA Intranet at <http://dis.kofia.or.kr>, (g) Nomura Singapore Ltd. ('NSL'), Singapore (Registration number 197201440E, regulated by the Monetary Authority of Singapore) (h) Nomura Australia Ltd. ('NAL'), Australia (ABN 48 003 032 513), regulated by the Australian Securities and Investment Commission ('ASIC') and holder of an Australian financial services licence number 246412, (i) Nomura Securities Malaysia Sdn. Bhd. ('NSM'), Malaysia, (j) NIHK, Taipei Branch ('NITB'), Taiwan, (k) Nomura Financial Advisory and Securities (India) Private Limited ('NFASL'), Mumbai, India (Registered Address: Ceejay House, Level 11, Plot F, Shivsagar Estate, Dr. Annie Besant Road, Worli, Mumbai- 400 018, India; Tel: 91 22 4037 4037, Fax: 91 22 4037 4111; CIN No: U74140MH2007PTC169116, SEBI Registration No. for Stock Broking activities : INZ000255633; SEBI Registration No. for Merchant Banking : INM000011419; SEBI Registration No. for Research: INH000001014 - Compliance Officer: Ms. Pratiksha Tondwalkar, 91 22 40374904, grievance email: investorgrievancesra@nomura.com Webpage: [LINK](#)

For reports with respect to Indian public companies or authored by India-based NFASL research analysts: (i) Investment in securities markets is subject to market risks. Read all the related documents carefully before investing. (ii) Registration granted by SEBI, and certification from NISM in no way guarantee performance of the intermediary or provide any assurance of returns to investors. (iii) NFASL terms and conditions for availing research services is disclosed on NFASL webpage.

(l) Nomura Fiduciary Research & Consulting Co., Ltd. ('NFRC') Tokyo, Japan. (m) Nomura Orient International Securities Co., Ltd ("NOI"), is a majority owned joint venture amongst Nomura Group, Orient International (Holding) Co., Ltd, and Shanghai Huangpu Investment Holding (Group) Co., Ltd. In accordance with the laws of the People's Republic of China ("PRC", excluding Hong Kong, Macau and Taiwan, for the purpose of this document), NOI is licensed in the PRC to provide securities research and investment recommendations and it operates independently from the other members of the Nomura Group; in particular, NOI's interests in PRC securities are not disclosed to, or aggregated with the holdings of, any other Nomura Group entities and the interests in PRC securities of other Nomura Group entities are not disclosed to, or aggregated with the holdings of, NOI. An individual name printed next to NOI on the front page of a research report indicates that individual is employed by NOI to provide research assistance to NIHK under a research partnership agreement. 'NSFSPL' next to an employee's name on the front page of a research report indicates that the individual is employed by Nomura Structured Finance Services Private Limited to provide assistance to certain Nomura entities under inter-company agreements. 'Verdhana' next to an individual's name on the front page of a research report indicates that the individual is employed by PT Verdhana Sekuritas Indonesia ('Verdhana') to provide research assistance to NIHK under a research partnership agreement and neither Verdhana nor such individual is licensed outside of Indonesia.

THIS MATERIAL IS: (I) FOR YOUR PRIVATE INFORMATION, AND WE ARE NOT SOLICITING ANY ACTION BASED UPON IT; (II) NOT TO BE CONSTRUED AS AN OFFER TO SELL OR A SOLICITATION OF AN OFFER TO BUY ANY SECURITIES IN ANY JURISDICTION WHERE SUCH OFFER OR SOLICITATION WOULD BE ILLEGAL; AND (III) OTHER THAN DISCLOSURES RELATING TO THE NOMURA GROUP, BASED UPON INFORMATION FROM SOURCES THAT WE CONSIDER RELIABLE, BUT HAS NOT BEEN INDEPENDENTLY VERIFIED BY NOMURA GROUP.

Other than disclosures relating to the Nomura Group, the Nomura Group does not warrant, represent or undertake, express or implied, that the document is fair, accurate, complete, correct, reliable or fit for any particular purpose or merchantable, and to the maximum extent permissible by law and/or regulation, does not accept liability (in negligence or otherwise, and in whole or in part) for any act (or decision not to act) resulting from use of this document and related data. To the maximum extent permissible by law and/or regulation, all warranties and other assurances by the Nomura Group are hereby excluded and the Nomura Group shall have no liability (in negligence or otherwise, and in whole or in part) for any loss howsoever arising from the use, misuse, or distribution of this material or the information contained in this material or otherwise arising in connection therewith.

Opinions or estimates expressed are current opinions as of the original publication date appearing on this material and the information, including the opinions and estimates contained herein, are subject to change without notice. The Nomura Group, however, expressly disclaims any obligation, and therefore is under no duty, to update or revise this document. Any comments or statements made herein are those of the author(s) and may differ from views held by other parties within Nomura Group. Clients should consider whether any advice or recommendation in this report is suitable for their particular circumstances and, if appropriate, seek professional advice, including tax advice. The Nomura Group does not provide tax advice.

The Nomura Group, and/or its officers, directors, employees and affiliates, may, to the extent permitted by applicable law and/or regulation, deal as principal, agent, or otherwise, or have long or short positions in, or buy or sell, the securities, commodities or instruments, or options or other derivative instruments based thereon, of issuers or securities mentioned herein. The Nomura Group companies may also act as market maker or liquidity provider (within the meaning of applicable regulations in the UK) in the financial instruments of the issuer. Where the activity of market maker is carried out in accordance with the definition given to it by specific laws and regulations of the US or other jurisdictions, this will be separately disclosed within the specific issuer disclosures.

This document may contain information obtained from third parties, including, but not limited to, ratings from credit ratings agencies such as Standard & Poor's. The Nomura Group hereby expressly disclaims all representations, warranties or undertakings of originality, fairness, accuracy, completeness, correctness, merchantability or fitness for a particular purpose with respect to any of the information obtained from third parties contained in this material or otherwise arising in connection therewith, and shall not be liable (in negligence or otherwise, and in whole or in part) for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including lost income or profits and opportunity costs) in connection with any use or misuse of any of the information obtained from third parties contained in this material or otherwise arising in connection therewith. Reproduction and distribution of third-party content in any form is prohibited except with the prior written permission of the related third-party. Third-party content providers do not, express or implied, guarantee the fairness, accuracy, completeness, correctness, timeliness or availability of any information, including ratings, and are not in any way responsible for any errors or omissions (negligent or otherwise), regardless of the cause, or for the results obtained from the use or misuse of such content. Third-party content providers give no express or implied warranties, including, but not limited to, any warranties of

merchantability or fitness for a particular purpose or use. Third-party content providers shall not be liable (in negligence or otherwise, and in whole or in part) for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including lost income or profits and opportunity costs) in connection with any use or misuse of their content, including ratings. Credit ratings are statements of opinions and are not statements of fact or recommendations to purchase hold or sell securities. They do not address the suitability of securities or the suitability of securities for investment purposes, and should not be relied on as investment advice. Any MSCI sourced information in this document is the exclusive property of MSCI Inc. ('MSCI'). Without prior written permission of MSCI, this information and any other MSCI intellectual property may not be duplicated, reproduced, re-disseminated, redistributed or used, in whole or in part, for any purpose whatsoever, including creating any financial products and any indices. This information is provided on an "as is" basis. The user assumes the entire risk of any use made of this information. MSCI, its affiliates and any third party involved in, or related to, computing or compiling the information hereby expressly disclaim all representations, warranties or undertakings of originality, fairness, accuracy, completeness, correctness, merchantability or fitness for a particular purpose with respect to any of this material or the information contained in this material or otherwise arising in connection therewith. Without limiting any of the foregoing, in no event shall MSCI, any of its affiliates or any third party involved in, or related to, computing or compiling the information have any liability (in negligence or otherwise, and in whole or in part) for any damages of any kind. MSCI and the MSCI indexes are services marks of MSCI and its affiliates.

The intellectual property rights and any other rights, in Russell/Nomura Japan Equity Index belong to Nomura Fiduciary Research & Consulting Co., Ltd. ("NFR") and FTSE Russell ("Russell"). NFR and Russell do not guarantee fairness, accuracy, completeness, correctness, reliability, usefulness, marketability, merchantability or fitness of the Index, and do not account for business activities or services that any index user and/or its affiliates undertakes with the use of the Index.

Investors should consider this document as only a single factor in making their investment decision and, as such, the report should not be viewed as identifying or suggesting all risks, direct or indirect, that may be associated with any investment decision. Nomura Group produces a number of different types of research product including, among others, fundamental analysis and quantitative analysis; recommendations contained in one type of research product may differ from recommendations contained in other types of research product, whether as a result of differing time horizons, methodologies or otherwise. The Nomura Group publishes research product in a number of different ways including the posting of product on the Nomura Group portals and/or distribution directly to clients. Different groups of clients may receive different products and services from the research department depending on their individual requirements.

Figures presented herein may refer to past performance or simulations based on past performance which are not reliable indicators of future or likely performance. Where the information contains an expectation, projection or indication of future performance and business prospects, such forecasts may not be a reliable indicator of future or likely performance. Moreover, simulations are based on models and simplifying assumptions which may oversimplify and not reflect the future distribution of returns. Any figure, strategy or index created and published for illustrative purposes within this document is not intended for "use" as a "benchmark" as defined by the European Benchmark Regulation. Certain securities are subject to fluctuations in exchange rates that could have an adverse effect on the value or price of, or income derived from, the investment.

With respect to Fixed Income Research: Recommendations fall into two categories: tactical, which typically last up to three months; or strategic, which typically last from 6-12 months. However, trade recommendations may be reviewed at any time as circumstances change. 'Stop loss' levels for trades are also provided; which, if hit, closes the trade recommendation automatically. Prices and yields shown in recommendations are taken at the time of submission for publication and are based on either indicative Bloomberg, LSEG or Nomura prices and yields at that time. The prices and yields shown are not necessarily those at which the trade recommendation can be implemented.

The securities described herein may not have been registered under the US Securities Act of 1933 (the '1933 Act'), and, in such case, may not be offered or sold in the US or to US persons unless they have been registered under the 1933 Act, or except in compliance with an exemption from the registration requirements of the 1933 Act. Unless governing law permits otherwise, any transaction should be executed via a Nomura entity in your home jurisdiction.

This document has been approved for distribution in the UK as investment research by Nlplc. Nlplc is authorised by the Prudential Regulation Authority and regulated by the Financial Conduct Authority and the Prudential Regulation Authority. Nlplc is a member of the London Stock Exchange. This document does not constitute a personal recommendation within the meaning of applicable regulations in the UK, or take into account the particular investment objectives, financial situations, or needs of individual investors. This document is intended only for investors who are 'eligible counterparties' or 'professional clients' for the purposes of applicable regulations in the UK, and may not, therefore, be redistributed to persons who are 'retail clients' for such purposes.

This document has been approved for distribution in the European Economic Area as investment research by Nomura Financial Products Europe GmbH ("NFPE"). NFPE is a company organized as a limited liability company under German law registered in the Commercial Register of the Court of Frankfurt/Main under HRB 110223. NFPE is authorized and regulated by the German Federal Financial Supervisory Authority (BaFin).

This document has been approved by NlHK, which is regulated by the Hong Kong Securities and Futures Commission, for distribution in Hong Kong by NlHK. This document is intended only for investors who are 'professional investors' for the purposes of applicable regulations in Hong Kong and may not, therefore, be redistributed to persons who are not 'professional investors' for such purposes.

This document has been approved for distribution in Australia by NAL, which is authorized and regulated in Australia by the ASIC.

This document has also been approved for distribution in Malaysia by NSM.

In Singapore, this document has been distributed by NSL, an exempt financial adviser as defined under the Financial Advisers Act (Chapter 110), among other things, and regulated by the Monetary Authority of Singapore. NSL may distribute this document produced by its foreign affiliates pursuant to an arrangement under Regulation 32C of the Financial Advisers Regulations. This document is intended for accredited, expert or institutional investors as defined by the Securities and Futures Act (Chapter 289). Where the recipient of this document is not an accredited, expert or institutional investor, NSL accepts legal responsibility for the contents of this document in respect of such recipient only to the extent required by law. Recipients of this document in Singapore should contact NSL in respect of matters arising from, or in connection with, this document. THIS DOCUMENT IS INTENDED FOR GENERAL CIRCULATION. IT DOES NOT TAKE INTO ACCOUNT THE SPECIFIC INVESTMENT OBJECTIVES, FINANCIAL SITUATION OR PARTICULAR NEEDS OF ANY PARTICULAR PERSON. RECIPIENTS SHOULD

TAKE INTO ACCOUNT THEIR SPECIFIC INVESTMENT OBJECTIVES, FINANCIAL SITUATION OR PARTICULAR NEEDS BEFORE MAKING A COMMITMENT TO PURCHASE ANY SECURITIES, INCLUDING SEEKING ADVICE FROM AN INDEPENDENT FINANCIAL ADVISER REGARDING THE SUITABILITY OF THE INVESTMENT, UNDER A SEPARATE ENGAGEMENT, AS THE RECIPIENT DEEMS FIT. Unless prohibited by the provisions of Regulation S of the 1933 Act, this material is distributed in the US, by NSI, a US-registered broker-dealer, which accepts responsibility for its contents in accordance with the provisions of Rule 15a-6, under the US Securities Exchange Act of 1934. The entity that prepared this document permits its separately operated affiliates within the Nomura Group to make copies of such documents available to their clients.

This document has not been approved for distribution to persons other than 'Authorised Persons', 'Exempt Persons' or 'Institutions' (as defined by the Capital Markets Authority) in the Kingdom of Saudi Arabia ('Saudi Arabia') or a 'Market Counterparty' or a 'Professional Client' (as defined by the Dubai Financial Services Authority) in the United Arab Emirates ('UAE') or a 'Market Counterparty' or a 'Business Customer' (as defined by the Qatar Financial Centre Regulatory Authority) in the State of Qatar ('Qatar') by Nomura Saudi Arabia, Nlplc or any other member of the Nomura Group, as the case may be. Neither this document nor any copy thereof may be taken or transmitted or distributed, directly or indirectly, by any person other than those authorised to do so into Saudi Arabia or in the UAE or in Qatar or to any person other than 'Authorised Persons', 'Exempt Persons' or 'Institutions' located in Saudi Arabia or a 'Market Counterparty' or a 'Professional Client' in the UAE or a 'Market Counterparty' or a 'Business Customer' in Qatar. Any failure to comply with these restrictions may constitute a violation of the laws of the UAE or Saudi Arabia or Qatar.

For report with reference of TAIWAN public companies or authored by Taiwan based research analyst:

THIS DOCUMENT IS SOLELY FOR REFERENCE ONLY. You should independently evaluate the investment risks and are solely responsible for your investment decisions. NO PORTION OF THE REPORT MAY BE REPRODUCED OR QUOTED BY THE PRESS OR ANY OTHER PERSON WITHOUT WRITTEN AUTHORIZATION FROM NOMURA GROUP. Pursuant to Operational Regulations Governing Securities Firms Recommending Trades in Securities to Customers and/or other applicable laws or regulations in Taiwan, you are prohibited to provide the reports to others (including but not limited to related parties, affiliated companies and any other third parties) or engage in any activities in connection with the reports which may involve conflicts of interests. INFORMATION ON SECURITIES / INSTRUMENTS NOT EXECUTABLE BY NOMURA INTERNATIONAL (HONG KONG) LTD., TAIPEI BRANCH IS FOR INFORMATIONAL PURPOSES ONLY AND IS NOT BE CONSTRUED AS A RECOMMENDATION OR A SOLICITATION TO TRADE IN SUCH SECURITIES / INSTRUMENTS.

This material may not be distributed in Indonesia or passed on within the territory of the Republic of Indonesia or to persons who are Indonesian citizens (wherever they are domiciled or located) or entities of or residents in Indonesia in a manner which constitutes a public offering under the laws of the Republic of Indonesia. The securities mentioned in this document may not be offered or sold in Indonesia or to persons who are citizens of Indonesia (wherever they are domiciled or located) or entities of or residents in Indonesia in a manner which constitutes a public offering under the laws of the Republic of Indonesia.

An individual name printed next to NOI on the front page of a research report indicates that this document is a translation of a research report issued by NOI in the PRC. In all other cases, this document is prepared by Nomura Group or its subsidiary or affiliate (collectively, "Offshore Issuers") that is not licensed in the PRC to provide securities research. This research report is not approved or intended to be circulated in the PRC. The A-share related analysis (if any) is not produced for any persons located or incorporated in the PRC. The recipients should not rely on any information contained in this research report in making investment decisions and Offshore Issuers take no responsibility in this regard. NO PART OF THIS MATERIAL MAY BE (I) COPIED, PHOTOCOPIED, REPRODUCED OR DUPLICATED IN ANY FORM, BY ANY MEANS; OR (II) REDISSEMINATED, REPUBLISHED OR REDISTRIBUTED WITHOUT THE PRIOR WRITTEN CONSENT OF A MEMBER OF THE NOMURA GROUP. If this document has been distributed by electronic transmission, such as e-mail, then such transmission cannot be guaranteed to be secure or error-free as information could be intercepted, corrupted, lost, destroyed, arrive late or incomplete, or contain viruses. The sender therefore does not accept liability (in negligence or otherwise, and in whole or in part) for any errors or omissions in the contents of this document, which may arise as a result of electronic transmission. If verification is required, please request a hard-copy version.

The Nomura Group manages conflicts with respect to the production of research through its compliance policies and procedures (including, but not limited to, Conflicts of Interest, Chinese Wall and Confidentiality policies) as well as through the maintenance of Chinese Walls and employee training.

Additional information regarding the methodologies or models used in the production of any investment recommendations contained within this document is available upon request by contacting the Research Analysts of Nomura listed on the front page. Disclosures information is available upon request and disclosure information is available at the Nomura Disclosure web page:

<http://go.nomuranow.com/research/m/Disclosures>

Copyright © 2026 Nomura International (Hong Kong) Ltd., Hong Kong. All rights reserved.