AI Networking

#AI

Companies: Lumentum, Coherent, EZconn, Corning, Furukawa, Broadcom, Celestica, Arista, Google, Amphenol, TEC, Sumitomo, InnoLight

Skye Chen
+852 3470 0011
skye.chen@aletheia-capital.com

Warren Lau
+852 9181 4766
warren.lau@aletheia-capital.com

Recent related reports:►

►Advanced Materials - PCB & CCL (4 Oct)
►Isu Petasys (4 Oct)
►Elite Material (4 Oct)
►MediaTek (1 Oct)
►AI Accelerators (12 Aug)
►MediaTek (19 Jul)
►Broadcom (28 May)
►MediaTek (2 May)
►Celestica (30, Apr)
►EZconn (26 Apr)
►General Server sector note (11 Apr)
►Lumentum (6 Apr)
►Server CPU sector note (29 Mar)
►AMD (29 Mar)
►AT&S (29 Mar)
►Micron (29 Mar)
►SK Hynix (29 Mar)
►Celestica (29 Jan)
►EZConn (27 Jan)
►Celestica (19 Jan)
►Lumentum (19 Jan)
►ISU Petasys (14 Jan)
►Elite Material (14 Jan)

Upside Surprise from TPU8 Networking Upgrade

Google's networking supply chain — ethernet switches, optical transceivers, optical cables, laser diodes, and OCS systems — is positioned for strong earnings upside in 2027–28, supported by a 2–5x increase in networking-hardware content value per TPUv8 versus TPUv7. These elevated hardware requirements stem from Google's efforts to fully realise the Virgo (East–West) and Jupiter (North–South) architectures, which enable TPUv8 clusters to scale to 134,400 units, reduce latency and hop count, and double ICI bandwidth to 19.2 Tb/s per TPU. Lumentum, EZconn, and Celestica are likely to be key beneficiaries, given their significant sales exposure to Google and dominant market positions.

TPUv8 will see a big network upgrade on hardware and capex

At the Hot Chips Conference in August, Google further disclosed the networking architecture for TPU8t and TPU8i. To support its bandwidth of ICI doubling to 19.2Tb/s, expanding scalability, connecting with more different type of racks and with multi-data center sites, we see Google meaningfully upgrading its network architecture hardware and capex budget, with the aim to realise the Virgo (East-West) and Jupiter (North-South) network system.

- **Ethernet switch:** We estimate that the TPUv8 architecture will require more ethernet-switching capacity across the ToR layer, the Virgo East–West fabric, and the Jupiter North–South network. We estimate bandwidth per TPUv8 rising to roughly 2 Tbps, up from about 0.4 Tbps in TPUv7. In addition to this, we also see adoption of 200Gbps SerDes on TPUv8 network.

- **Transceiver module and laser diode:** TPUv8 node will require roughly three transceiver modules, up from fewer than two in TPUv7. This step-up in optical I/O also implies a substantial increase in laser-diode demand across 100 Gbps, 200 Gbps, and CW (continuous-wave) laser categories.

- **Optical cable:** Given the larger pod and cluster scale in TPUv8 — and the need to interconnect different rack types — the optical-cable count per TPU is expected to more than double, rising from 8 units in TPUv7 to roughly 18 units in TPUv8.

- **OCS:** We are seeing more adoption of OCS in TPU8 for the Jupiter network and TPU8i rack connections. Hence, we estimate total OCS demand for the TPU8 life cycle could be as much as 258,000 units in compared to 20,800 units in TPU7 cycle.

- **Copper cable (DAC/ACC cable):** Although ICI SerDes will upgrade to 200Gbps from 100Gbps in TPU7, the adoption of Boardfly for TPU8i could reduce the number of copper cables per rack. We estimate copper cables per TPU8 will increase slightly to 1.9 from 1.5 for TPU7.

Beneficiaries of TPU v8 network upgrade

Meaningful TPU demand growth—from 3.6m units in 2026 to 8.0–8.5m units in 2027—combined with strong rising networking-hardware content value per TPU, will materially benefit the hardware supply chain. Among them, Lumentum, EZconn, and Celestica warrant particular attention given their high sales exposure to Google and their dominant market position. We rate all three stocks as BUYs and have them in our Google G7 portfolio.

Please refer to important isclosures and disclaimers at the end of this report. This report is intended solely for the subscriber.
Unauthorized sharing, distribution, or reproduction may violate applicable regulatory requirements and could result in legal action.

TPU8 will see a big network upgrade

At the Hot Chips Conference, a premier annual semiconductor industry event at Stanford University at the end of August, Google further disclosed the networking architecture for its upcoming TPU8t (Zebrafish – MediaTek) and TPU8i (Sunfish – Broadcom). To support its bandwidth of ICI (Inter-Chip Interconnect) doubling to 19.2Tb/s, expanding scalability, connecting with more different racks (such as general server and storage racks) and with multi-datacenter sites, we expect Google will meaningfully upgrade its network architecture and capex budget.

Google is employing two major network architectures — Jupiter (North–South traffic) and the newly announced Virgo (East–West traffic) — to expand TPUv8 scalability and integrate TPU clusters with storage systems, general-purpose servers, and Google's global WAN.

Jupiter handles North–South traffic, connecting TPU compute racks to storage, general servers, and Google's global WAN through a combination of ethernet switching and optical circuit switching (OCS). When paired with the new Boardfly ICI topology, TPU8i benefits from lower latency and fewer hops, improving inference responsiveness and model-weight access.

Virgo, designed for East–West traffic, is a co-design of hardware network architecture, optical switching, and software orchestration. This architecture dramatically increases horizontal scalability, enabling Google to expand the maximum TPU8t cluster size to 1 million units — a tenfold increase from the 100,000-unit limit of TPUv7 (Ironwood).

FIGURE 1: KEY NETWORK UPGRADE IN TPU V8 VS TPU V7

Comparison Metric	TPU 7 (Ironwood)	TPU 8t (Zebrafish)	TPU 8i (Sunfish)	Key Evolution Benefits
Inter-Chip Interconnect (ICI) Bandwidth	9.6 Tb/s	**19.2 Tb/s**	19.2 Tb/s	
Network Topology	3D Torus @ 100G	Enhanced 3D Torus @ 200G + Virgo Network	Boardfly @ 200G + SPOCS	
Single Superpod Scale	Up to 9,216 chips	Up to 9,600 chips	Up to 1,152 chips	Expands the single hardware pool size to deliver up to 121 ExaFlops of compute power.
Data Center Network (DCN) Bandwidth	100 Gbps / chip	**400 Gbps / chip**	400 Gbps / chip	
Max Logical Cluster Scale-out	Over 100,000 chips	Up to **1,000,000** chips (for a single logical workload)	N.A	Combined with Google Pathways and JAX, it enables near-linear scale-out across geodistributed data centers.
Data & Storage Access Optimization	Traditional PCIe Host-managed mode	TPUDirect RDMA & Storage	TPUDirect RDMA & Storage	Bypasses the Host CPU to stream data directly into TPU HBM, delivering 10x faster storage access and maximizing Flops utilization.

Source: Aletheia Capital/TAG

This major upgrade is highly reliant on networking hardware advancements, particularly in optical communications. Our deep-dive research into Google's new TPU network architecture leads us to conclude that we will see meaningful demand increase for ethernet switches, optical transceivers, laser diodes (both EML and CW lasers), optical cables, and optical circuit switches (OCS). For copper cabling within the TPU rack, we see a mixed trend: while TPU8 upgrading the SerDes to 200Gbps (up from 100Gbps in TPU7) would benefit content value, the adoption of the Boardfly topology could reduce the overall volume of copper cables required per rack in TPU8i (accounting for 50% of TPU8 volume).

- **Ethernet switch:** We estimate that the TPUv8 architecture will require more ethernet-switching capacity across the ToR layer, the Virgo East–West fabric, and the

Jupiter North–South network. As a result, the aggregate network bandwidth per TPU is projected to increase sharply — rising to roughly 2 Tbps per TPU, up from about 0.4 Tbps in TPU7. In addition to this, we also see adoption of 200Gbps SerDes in TPU8 network.

- **Transceiver module and laser diode:** The increased reliance on ethernet switches in Virgo naturally drives higher demand for optical transceiver modules. The reduction in the number of TPUs per rack to 32 units from 64 will also further add to transceiver requirements. Based on our model, each TPU8 node will require roughly three transceiver modules, up from fewer than two in TPUv7. This step-up in optical I/O also implies a substantial increase in laser-diode demand across 100 Gbps, 200 Gbps, and CW (continuous-wave) laser categories.
- **Optical cable:** Given the larger pod and cluster scale in TPUv8 — and the need to interconnect different rack types — the optical-cable count per TPU is expected to more than double, from 8 units in TPUv7 to roughly 18 units in TPUv8.
- **OCS:** We are seeing more adoption of OCS in TPU8 for the Jupiter network and TPU8i rack connections. Hence, we estimate total OCS demand for the TPU8 life cycle could be as much as 258,000 units compared to 20,800 units in TPU7 cycle.
- **Copper cable (DAC/ACC cable):** Although ICI SerDes will upgrade to 200Gbps from 100Gbps in TPU7, the adoption of Boardfly for TPU8i could reduce the number of copper cables per rack. We estimate copper cables per TPU8 will increase slightly to 1.9 from 1.5 for TPU7.

FIGURE 2: ESTIMATE OF OPTICAL NETWORK HARDWARE DEMAND

Component ratio to TPU/Pod	TPU 8	TPU 7	Remarks
Ethernet switch bandwidth – Tb per TP	2.0	0.4	Including both of 100G and 200G SerDes
Transceiver module per TPU	3.0	1.8	Including both 800 Gbps and 1.6 Tbps
Total OCS demand in product lifecycle (mix of different specs)	258,542	20,833	We assume lifecycle volume for TPU 7 is 4m and TPU 8 is 14-15m units
Optical cable per TPU	18.2	8.0	
DAC/ACC cable per TPU	1.9	1.5	TPU 7 uses 100 Gbps per Lane and TPU 8 uses 200 Gbps per Lane

Source: Aletheia Capital/TAG

Meaningful TPU demand growth—from 3.6m units in 2026 to 8.0–8.5m units in 2027—combined with strong rising networking-hardware content value per TPU, will materially benefit the hardware supply chain. Below we outline the major solution providers and manufacturers supporting Google's TPU networking ecosystem. Among them, Lumentum, EZconn, and Celestica warrant particular attention given their high sales exposure to Google and their dominant market position.

FIGURE 3: KEY SUPPLIER OF NETWORK HARDWARE FOR TPU8

Key network hardware	Key Solution provider	ODM/OEM
Ethernet switch	Broadcom and others	**Celestica**, Inventec, Arista
Transceiver module	**Lumentum**, InnoLight, Coherent	
Laser diode	**Lumentum**, Sumitomo	
OCS	Lumentum	Fabrinet, **Celestica**
Optical cable	Upstream material: Corning, Furukawa Optical cable: Amphenol, TEC and **EZconn**	
DAC/ACC cable	Amphenol, TEC	

Source: Aletheia Capital/TAG

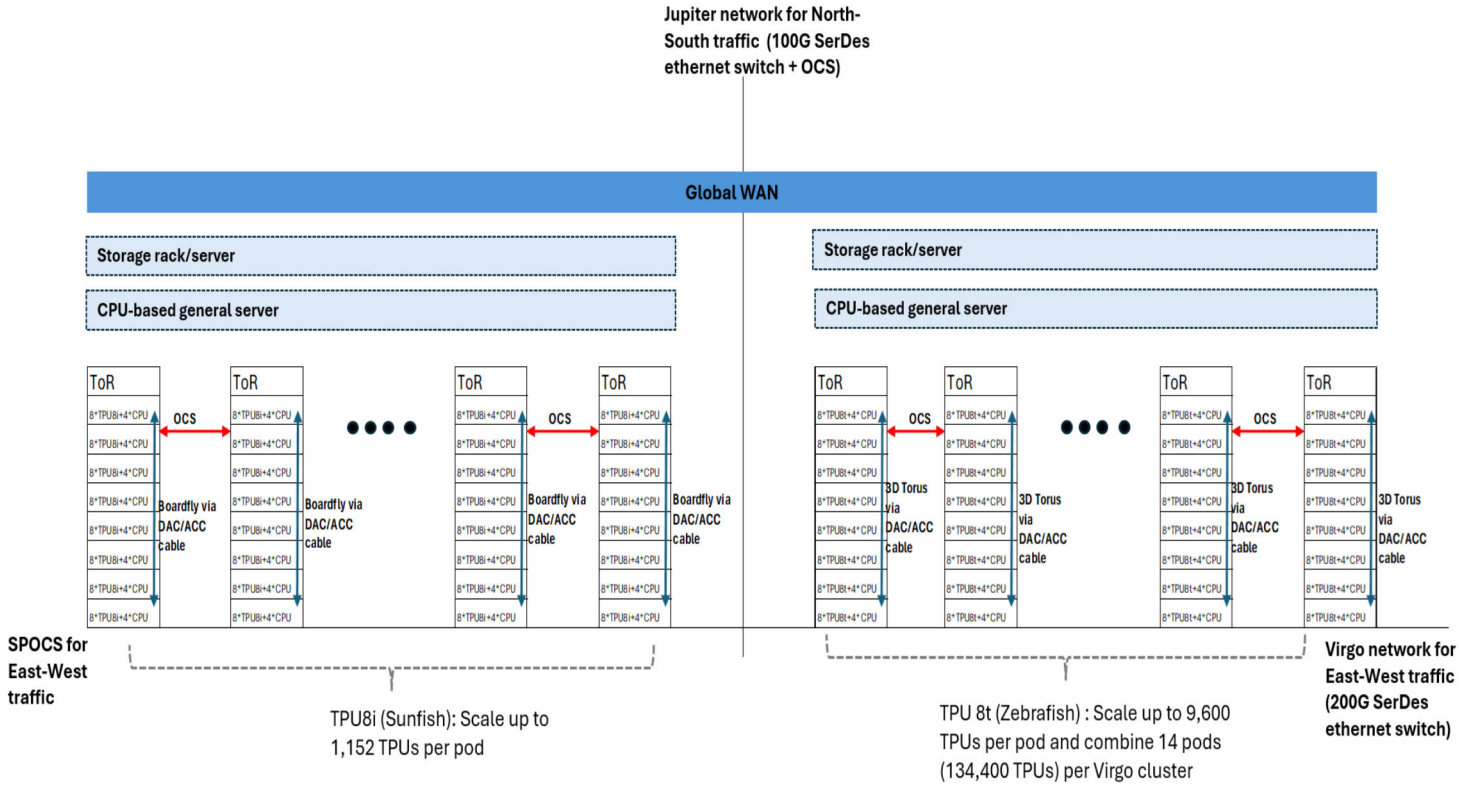
Details of Virgo and Jupiter

Google's upcoming TPUv8 generation is expected to introduce a differentiated networking architecture for 8i and 8t, reflecting their distinct roles within Google's AI stack. While the company has not fully disclosed the boundary between the two, industry signals suggest TPU8i is optimized for inference workloads, whereas TPU8t will lean toward training-optimized performance.

We observe meaningful architectural differences in how TPU chips are interconnected within each rack, as well as in the East–West rack-to-rack topology used to assemble a full TPU Pod. These variations reflect Google's workload-specific optimization for scaling compute domains and fabric bandwidth across the 8i and 8t configurations.

However, both TPU 8i and TPU 8t are expected to leverage Google's Jupiter networking fabric, using North–South traffic flows to interface with the broader data center stack. This architecture enables each TPU domain to connect efficiently to general-purpose servers, storage racks/servers, and ultimately to Google's global backbone network (Global WAN).

FIGURE 4: NETWORKING STRUCTURE FOR TPU8i AND TPU8t



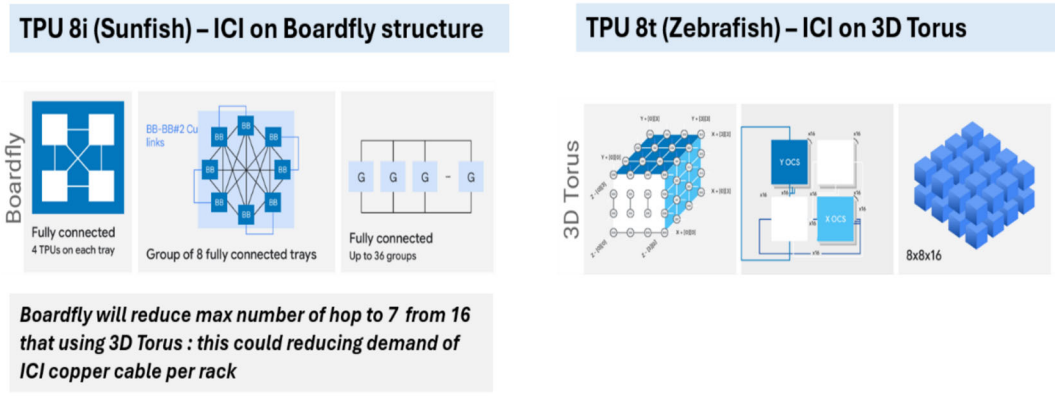
Source: Google at Hot Chips Conference and Aletheia Capital/TAG

TPU connection inside of rack

TPU v7 (Ironwood) uses a 3D-Torus interconnect implemented through ICI copper cabling. Google will continue to adopt 3D Torus for its TPU 8t (Zebrafish), aligning with its training-oriented design.

In contrast, TPU 8i will shift to Google's new Boardfly topology — a high-radix, multi-way hierarchical fabric that behaves more like a high-radix switch than a traditional torus. Boardfly is engineered for instant, any-to-any node communication, reducing hop count and latency, and is optimized for MoE (Mixture-of-Experts) and emerging agent-style inference architectures, where rapid expert routing and low-latency fan-out/fan-in patterns dominate.

FIGURE 5: GOOGLE'S NEW BOARDFLY STRUCTURE FOR ICI



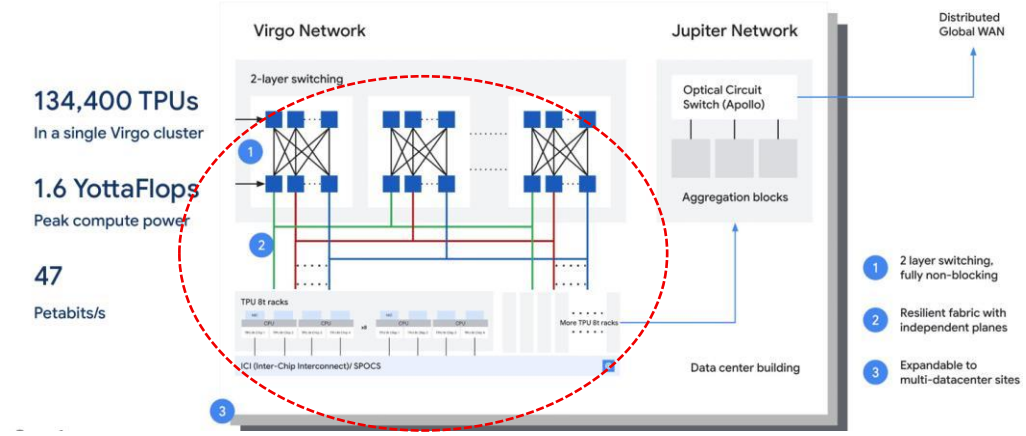
Source: Google at Hot Chips Conference

East-West (Virgo): building a Pod for TPUv8

TPU 8i will use the Boardfly topology inside each rack and extend it across racks through SPOCS (Optical Circuit Switch). This optical fabric allows Google to interconnect 36 racks, enabling 1,152 TPU8i chips to operate cohesively as a single Pod with low-latency, any-to-any communication that is optimized for popular MoE and AI agent.

For TPU8t, Google is introducing the new Virgo network architecture, which appears to rely primarily on ethernet-switch-based connectivity built on 200 Gbps SerDes. Virgo provides the scale and bandwidth needed to aggregate roughly 300 TPU8t racks (9,600 TPUs) into a Pod and can further interlink 14 Pods into a single Virgo cluster comprising 134,400 TPUs.

FIGURE 6: VIFGO NETWORK ARCHITECTURE

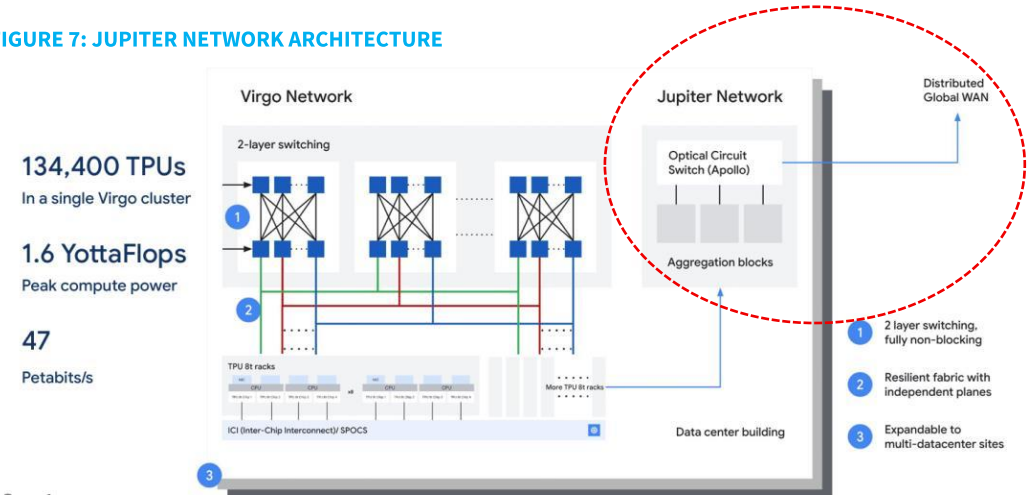


Source: Google at Hot Chips Conference

North-South network (Jupiter): storage, general server and global WAN

Jupiter network handles North–South traffic, including receiving external user inference requests (prompts), returning generated tokens to users, and loading the latest AI model weights from external high-bandwidth storage systems. Hence both TPU8i and TPU8t will connect to Jupiter. This network system is built by both OCS and ethernet switch.

FIGURE 7: JUPITER NETWORK ARCHITECTURE



Source: Google at Hot Chips Conference

Alētheia Research Team

Macro/Strategy	LEAD ANALYST	Sector	LEAD ANALYST
Global Strategy	 Jonathan Wilmot	Technology Hardware	 Warren Lau
Asianomics	 Dr. Jim Walker	Consumer & Internet	 Nirgunan Tiruchelvam
Tech Thematic Strategy	 Keith Woolcock	China Technology	 Eric Wen
Asia Equity Strategy	 David Scott	Telecoms	 Chris Hoare
China Strategy	 Vincent Chan	CrossASEAN	 Angus Mackintosh
Tactical Alpha Strategy	 Justin Collazo	India	 Maulik Patel
Global Commodities	 Steven Schlegel	MENA	 Jaap Meijer

Product Marketing Teams

Serving 300+ clients, 8,000+ touch points, in 20+ geographies with $15tr+ AUM



- **Al Park**
  1 201 962 0539
  al.park@aletheia-capital.com
  AMERICAS - NEW YORK
- **Linda Gustafsson**
  44 791 91 21 365
  linda@nordlinkcapital.com
  EUROPE - NORDICS
- **Michael Chambers**
  65 9858 9759
  michael.chambers@aletheia-capital.com
  ASIA - SINGAPORE
- **Graeme Bateman**
  852 2534 7437
  graeme.bateman@aletheia-capital.com
  ASIA - HONG KONG
- **Wayne Chang**
  1 646 256 3366
  wayne.chang@aletheia-capital.com
  AMERICAS - NEW YORK
- **Daren Riley**
  44 748 31 59 865
  daren.riley@aletheia-capital.com
  EUROPE - LONDON
- **Richard Wallace**
  852 2534 7429
  richard.wallace@aletheia-capital.com
  ASIA / EUROPE - HONG KONG
- **Terrence Alford**
  1 469 403 9936
  terrence.alford@aletheia-capital.com
  AMERICAS - DALLAS
- **Dhananjay Mahurkar**
  44 777 55 26 870
  dhananjay.mahurkar@aletheia-capital.com
  EUROPE - LONDON
- **Augustine Chen**
  8869 2739 8793
  augustine.chen@aletheia-capital.com
  ASIA - HONG KONG

Firm Disclosures

Aletheia Capital Ltd ("Aletheia") is a limited company registered in Hong Kong, company number 2657920, located at Level 28 AIA, 1 Connaught Road, Central, Hong Kong.

Aletheia Analyst Network Ltd (“AAN”) is a limited company registered in Hong Kong and is a wholly owned by Aletheia and is regulated by the Hong Kong Securities and Futures Commission, is a registered investment advisor with the U.S. Securities and Exchange Commission and is regulated by the Financial Conduct Authority, Firm Reference Number 794762.

Aletheia Capital (Singapore) Pte Ltd (“ACSG”) is a limited company registered in Singapore, UEN 201823248E, and is a wholly owned by AAN and is an Exempt Financial Adviser as defined in the Financial Advisers Act and regulated by the Monetary Authority of Singapore.

This report was published by AAN and is distributed by AAN and ACSG. For investors in Singapore, this material is provided by ACSG pursuant to Regulation 32C of the Financial Advisers Regulations. If there are any matters arising from, or in connection with this material, please contact ACSG, 1 Fullerton Road, #02-01, One Fullerton, Singapore 049213, Tel: +65 6721 7116.

Additional information will be made available upon request.

Analyst Certification and Disclosures

The analyst(s) named in this report certifies that (i) all views expressed in this report accurately reflect the personal views of the analyst(s) with regard to any and all of the subject securities and companies mentioned in this report and (ii) no part of the compensation of the analyst(s) was, is, or will be, directly or indirectly, related to the specific recommendation or views expressed by that analyst herein.

The analyst(s) named in this report (or their associates) does not have a financial interest in the corporation(s) mentioned in this report.

Disclaimer

This report is for the use of intended recipients only and may not be reproduced, in whole or in part, or delivered or transmitted to any other person without our prior written consent. By accepting this report, the recipient agrees to be bound by the terms and limitations set out herein. The report is prepared for professional investors only whose business involves the acquisition, disposal or holding of securities, whether as principal or agent and must not be re-distributed to retail clients. The information contained in this report has been obtained from public sources believed to be reliable and the opinions contained herein are expressions of belief based on such information. Except as expressly set forth herein or in recipient’s subscription agreement, no representation or warranty, express or implied, is made that such information or opinions is accurate, complete or verified and it should not be relied upon as such. Nothing in this report constitutes a representation that any investment strategy or recommendation contained herein is suitable or appropriate to a recipient’s individual circumstances or otherwise constitutes a personal recommendation. It is published solely for information purposes, it does not constitute an advertisement, a prospectus or other offering document or an offer or a solicitation to buy or sell any securities or related financial instruments in any jurisdiction. Information and opinions contained in this report are published for the reference of the recipients and are not to be relied upon as authoritative or without the recipient’s own independent verification or taken in substitution for the exercise of the recipient’s own judgement. All opinions contained herein constitute the views of the analyst(s) named in this report, they are subject to change without notice and are not intended to provide the sole basis of any evaluation of the subject securities and companies mentioned in this report. Any reference to past performance should not be taken as an indication of future performance. Except in the case of fraud, gross negligence or willful misconduct, or except as otherwise set forth in recipient’s subscription agreement, the firm does not accept any liability whatsoever for any direct or consequential loss arising from any use of the materials contained in this report.

© 2026 Aletheia Capital Ltd