US Technology

Key takeaways from our Silicon Valley tech tour

Equities
Semiconductors & Equipment

United States

- Key takeaways from our meetings with 27 tech companies in Silicon Valley along with a group of investors
- No slowdown in demand as supply chain bottlenecks are expected to continue through 2027 or possibly longer
- Monetization of AI remains a key debate; AI exposure led re-rating unlikely with focus now on earnings upside potential

Frank Lee*
Global Head of Tech Hardware & Semi Research
The Hongkong and Shanghai Banking Corporation Limited
frank.lee@hsbc.com.hk
+852 2996 6916

Pulkit Aggarwal*
Analyst, Semis and Tech Hardware
HSBC Securities and Capital Markets (India) Private Limited
pulkit.aggarwal@hsbc.co.in
+91 86 5529 3272

** Employed by a non-US affiliate of HSBC Securities (USA) Inc. and is not registered/ qualified pursuant to FINRA regulations*

We held meetings with 27 tech companies in Silicon Valley, along with a group of investors, across five days. We highlight the key takeaways from these meetings in this note.

Supply chain bottlenecks to persist through 2027 at least: There is a broad, AI-led supply crunch rather than a single isolated constraint. Demand continues to run ahead of supply across much of the semiconductor and data center ecosystem. The major bottlenecks are across 1) data center capacity and power; 2) substrates, foundry wafers, and advanced packaging; 3) optical components; and 4) memory. Companies are responding by pre-booking capacity, diversifying suppliers, vertically integrating to the extent possible, and holding more inventory. However, the capacity bottlenecks are expected to persist through 2027, with a possibility of extending beyond that too.

No signs of demand slowdown; focus on open-source models: There are no signs of AI-related demand slowing down. However, demand is broadening from training towards inference, agents, enterprise, sovereign AI, storage, and physical AI. The focus is also shifting towards open-source models due to cost, flexibility, and sovereignty. China's push in open-source models has triggered an urgency to ramp up open-source models in the US. However, both open-source and proprietary / closed-source models are expected to co-exist.

AI monetization remains a key debate: AI monetization remains a key debate among investors with software companies acknowledging that though AI monetization is moving beyond experimentation, they are still unable to quantify the monetization story. The revenue contribution remains modest relative to the underlying usage and investment. Companies that sit inside the workflow or communications layer rather than the ones trying to build frontier models themselves seem to be better set-up for monetization.

Bottom-up earnings upside remains key for re-rating: The demand is expected to remain strong but constrained by supply capacity. While hyperscaler capex momentum remains critical to support growth of semi companies, enterprise and sovereign AI adoption are also becoming more important as they help support a wider set of beneficiaries. With this backdrop, we expect the multiples of semi and tech companies to remain capped with upside driven by bottom-up earnings performance.

Disclosures & Disclaimer
This report must be read with the disclosures and the analyst certifications in the Disclosure appendix, and with the Disclaimer, which forms part of it.

Issuer of report: The Hongkong and Shanghai Banking Corporation Limited

View HSBC Global Investment Research at:
https://www.research.hsbc.com

Ratings and target prices

Company	Ticker	Currency	Current price	TP	Rating	Implied upside	Mkt cap USDm	EPS growth 2026e	EPS growth 2027e
NVIDIA	NVDA US	USD	227.38	365.00	Buy	61%	5,479,858	94%	58%
Intel	INTC US	USD	121.78	200.00	Buy	64%	643,616	245%	70%
Marvell	MRVL US	USD	257.38	325.00	Buy	26%	225,696	52%	79%
Qualcomm	QCOM US	USD	194.23	165.00	Hold	-15%	203,941	-13%	-4%
GlobalFoundries Inc.	GFS US	USD	48.57	55.00	Hold	13%	26,653	14%	31%

Source: HSBC estimates, Bloomberg. Priced at close of 21 Sep 2026

Supply chain bottlenecks

As highlighted above, the supply crunch is broad-based and not an isolated constraint. We summarize some of the key areas and major bottlenecks discussed by the companies.

Data center capacity and power

- **DeepInfra:** DeepInfra can build an inference cluster in as little as 90 days, but securing suitable data center space is becoming increasingly difficult.

  The main constraints are:

  1. Limited availability of data center space.
  2. Removal of tax incentives in some US states.
  3. Securing liquid cooled clusters to accommodate next-gen chips

- **Celesta Capital**: Celesta described major obstacles to building terrestrial data centers in the US. Around 60% of projects are reportedly facing organized protests or political challenges. This creates delays around permitting, power availability, and local acceptance. The firm sees space-based data centers as a possible long-term alternative, although transport and deployment costs are currently around 20 times those of terrestrial facilities.

- **NVIDIA**: NVIDIA said many older data centers cannot support the power and liquid-cooling requirements of newer systems such as Vera Rubin. As a result, much of current customer spending is effectively greenfield, requiring new infrastructure rather than simple equipment upgrades.

Substrates, foundry wafers, and advanced packaging

- **NVIDIA:** NVIDIA's demand remains well above available supply. Management referred to next year's growth of 70% as supply-constrained, while saying unconstrained demand is materially higher. The company's ability to secure wafer allocation from TSMC is strategically important.

- **Intel:** Intel said server demand is strong but that it can currently fulfil only around 50% of potential server demand. The principal constraints are:

  1. ABF substrate availability is more of a constraint than wafer supply
  2. Yield challenges associated with large reticle sizes for EMIB-T packaging
  3. Supplier learning curves for advanced packaging substrates

- **GlobalFoundries:** GlobalFoundries described industry capacity as tigh processes linked to data centers.

- **SoftBank:** SoftBank highlighted continuing foundry tightness and is wo while exploring other capacity options. Its broader strategy is tied to loc semiconductor and packaging supply chains, particularly in the US and



- **Qualcomm:** The USD5bn FY27 data center revenue target is described as supply-constrained rather than demand-constrained. The company has secured good visibility for FY2027 but still faces tight supply in advanced packaging, substrates, testers, and memory stacking.
- **Applied materials:** Applied Materials described a more complex manufacturing environment, driven by 1) increasing logic content in DRAM; 2) advanced 300mm packaging; 3) through-silicon via (TSV) and wafer-level packaging; and 4) higher requirements for yield, power efficiency, and variability control.

Optical components

- **Arista Networks:** Arista is supply-constrained and expects the situation to continue through at least 2027. Demand is exceeding supply, with constraints particularly affecting 3nm chips, transceivers, and high-speed optical components. The company believes that if supply unlocks earlier than expected, it could exceed its current 40% growth guidance.
- **Lightmatter:** Lightmatter identified structural networking bottlenecks rather than simple component shortages:

    1. Insufficient bandwidth per switch.
    2. The need for longer reach than what is provided by one-meter copper cables.
    3. The physical "shoreline" limitation of conventional chips, where I/O can only be placed around the chip perimeter.

    According to Lightmatter, the broader ecosystem is also dealing with delayed CPO adoption, with expected timelines moving towards 2028-30.

- **Lumentum:** Lumentum described one of the clearest component-level bottlenecks: a severe shortage of 200G-lane continuous-wave lasers. Lumentum also highlighted limited supply of six-inch laser substrates, with some sulphur-doped substrates reportedly achieving yields as low as 5%.
- **Marvell:** Marvell identified bottlenecks across the optical supply chain, particularly in lasers, silicon photonics and assembly.

Memory

- **Cohu:** Cohu sees growing demand for HBM inspection because the cost of a failed die in a completed GPU board can reach approximately USD25k. According to Cohu, the bottlenecks are primarily operational like 1) a large number of simultaneous customer qualifications; 2) limited engineering resources; 3) the need to maintain lead times and product availability for existing customers; and 4) increasing complexity in testing HBM3, HBM4 and eventually HBM5
- **Seagate:** The main supply-chain issue is the long production cycle for heat-assisted magnetic recording (HAMR) components. Heads can take three to four quarters to produce, creating a transition lag between investment and revenue.
- **Western Digital:** Western Digital identified constraints in NAND and DRAM components used in bill of materials. Inflation and input-cost pressure are making the traditional 10% annual cost-reduction target more difficult to achieve. The company expects to pass some of these costs through while using long-term agreements to support more predictable pricing.

Demand sustainability

The comments from the companies support sustained AI demand with demand broadening beyond training towards inference, agents, enterprise, sovereign AI, and physical AI.

Structural arguments supporting durable demand

- **Inference and AI agents could materially expand usage: DeepInfra** expects spending to shift from training to inference. AI agents may consume tens of millions of tokens in task loops, particularly for coding, creating much higher usage than human chat. **NVIDIA** also expects inference eventually to become the larger market.
- **Demand is moving beyond a small group of hyperscalers: NVIDIA** said growth is broadening into AI clouds, enterprise, and sovereign customers. Its ACIE segment already receives more than half of its revenue from AI clouds and "neo-clouds", while enterprise and sovereign adoption is still at an earlier stage and could outgrow hyperscale demand.
- **AI infrastructure demand remains supply-constrained:** Multiple companies acknowledged that demand remains strong but remains supply-constrained capping further growth.
  - **Arista Networks:** demand exceeds supply, with constraints expected through at least 2027.
  - **Intel:** able to fulfil only around 50% of potential server demand, with substrates a key constraint.
  - **Qualcomm:** USD5bn data center target is supply-constrained rather than demand-constrained
- **AI is creating new infrastructure requirements, not simply replacing existing workloads: Seagate** argued that AI agents, longer context windows, and physical AI will drive more storage for video, logs, sensor data, and compliance records. Hard disk drives (HDDs) remain economically attractive because flash can cost 20-25 times more per terabyte.
- **Enterprise and sovereign AI could provide a more durable demand base: HP Enterprise** sees enterprise and sovereign deployments as less exposed to hyperscaler capex cycles and less vulnerable to white-box commoditization. Its Private Cloud AI offering is positioned as delivering 30-60% cost savings versus public-cloud, token-based models.

Evidence of confidence

- **SoftBank** dismissed AI bubble concerns, saying demand for computing power remains well above supply.
- **Monolithic Power Systems** described AI/data center demand as healthy and increasingly diversified, with ASIC revenue likely already exceeding GPU revenue.
- **GlobalFoundries** reported that recent silicon-photonics demand signals were "flashing green lights", with more than 40 active customer engagements and a target of over USD1bn in silicon-photonics revenue by 2028.
- **Intel** believes its addressable market is now around 15x larger than historical, supporting confidence in sustained double-digit growth.
- **Cohu** noted customer equipment utilization has reached the 80% threshold at which customers historically begin buying additional capacity.

Areas of caution

- **Demand sustainability depends on economics:** Companies repeatedly referenced power, cooling, data center capacity, financing costs, and supply availability as constraints.
- **Some demand is still greenfield and dependent on infrastructure build-out: NVIDIA** said much current customer spending is greenfield because older data centers cannot support the power and liquid-cooling requirements of newer systems. This supports short-term demand but also makes timing dependent on construction, energy, and supply-chain execution.

Focus shift towards open-source models

The discussion favors open-source or open weight models for cost, flexibility, and sovereignty, while proprietary/closed-source models retain advantages where frontier performance, safety, workflow integration, or enterprise-grade capabilities are critical. Previously the US companies were more focused on frontier models. However, since our Silicon Valley Tech Tour in June, things seem to have changed, especially after the Chinese open-source models reset industry expectations about compute price. As a result, open-source models have become a key debate in the US too.

Cost advantages of open-source models

- **DeepInfra** focuses on Llama, Mistral, DeepSeek, and GLM. It said these models can run around ten times cheaper than leading closed-source models while maintaining competitive quality.
- **Pinterest** said open-source models cost less than 8% per transaction compared with proprietary models of a similar parameter size.
- **HPE** uses open-weight model inferencing and intelligent routing in Private Cloud AI, claiming 30-60% lower costs than public-cloud token-based models.

Flexibility and control

- **Pinterest** can securely post-train open-source models in-house and swap or test multiple models side-by-side, including US-based models such as Gemma and models from Asia.
- **DeepInfra** offers shared serverless hosting and dedicated private deployments, allowing customers to run custom models or select specific performance profiles.
- Open models support data sovereignty and deployment flexibility, particularly for enterprise, government, and regulated customers.

Adoption and ecosystem benefits

- **NVIDIA** is developing open-source models such as NeMo as ecosystem enablers rather than direct threats to customers. Management said open source is important for:
    1. Broadening AI adoption.
    2. Ensuring access to high-performance models.
    3. Strengthening the US AI ecosystem.

AI monetization still in nascent stages

The meetings broadly suggested that AI monetization is moving beyond experimentation, although the revenue contribution is still modest relative to the underlying usage and investment.

- **Twilio was the clearest voice of caution:** AI-driven voice and agent adoption is promising, but monetization remains early. Customers often begin with limited production pilots, especially in regulated industries. AI has not yet materially changed adoption or revenue mix, and the pace of conversion from pilots to large-scale usage remains uncertain.
- **Pinterest sees strong engagement but monetization still has to catch up:** The platform has 640 million monthly active users and sustained user growth, but international monetization remains below peers. AI-related GPU investment of USD50m is a margin headwind, even though the business remains highly cash generative.
- **Adobe** has an AI-first ARR of approximately USD650m, still less than 2% of its roughly USD28bn revenue base. It is moving from seat-based subscriptions towards a hybrid model combining subscriptions with generative credits, allowing revenue to rise with usage.
- **HP Enterprise** is monetizing AI through integrated infrastructure and private-cloud deployments rather than token-based public-cloud consumption.
- **Cadence** is still in the "price discovery" phase for AI design tools. It has more than 20 customer engagements for agentic AI offerings, but monetization is not yet mature.

Strong earnings performance key for re-rating

AI-led demand is expected to remain robust, although growth will continue to be constrained by available supply capacity. While capital expenditure momentum among hyperscalers remains a critical driver of growth for semiconductor companies, the increasing adoption of artificial intelligence by enterprises and sovereign entities is becoming an increasingly important source of support.

This broader adoption is expected to benefit a wider range of companies across the semiconductor and technology ecosystems, extending beyond the largest hyperscalers. Against this backdrop, we expect valuation multiples for semiconductor and technology companies to remain capped in the short-term. Further re-rating is therefore likely to depend primarily on company-specific factors, particularly the delivery of strong earnings performance and the ability to exceed market expectations.

Companies mentioned in this report

Company	Ticker	Currency	CMP	TP	Rating
TSMC	2330 TT	TWD	2,480.00	3,400.00	Buy
Applied Materials	AMAT US	USD	464.24	641.00	Buy
Pinterest	PINS US	USD	19.36	23.40	Hold
Twilio	TWLO US	USD	266.06	211.00	Hold
Adobe	ADBE US	USD	249.52	310.00	Buy
SoftBank	9984 JP	JPY	6,247.00	NR	NR
Arista Networks	ANET US	USD	205.42	NR	NR
Lumentum	LITE US	USD	954.49	NR	NR
Cohu	COHU US	USD	60.62	NR	NR
Seagate	STX US	USD	877.33	NR	NR
Western Digital	WDC US	USD	448.17	NR	NR
HP Enterprise	HPE US	USD	61.75	NR	NR
Monolithic Power Systems	MPWR US	USD	1277.68	NR	NR
Cadence	CDNS US	USD	294.9	NR	NR
Lightmatter	Not listed				
DeepInfra	Not listed				
Celesta Capital	Not listed				

Source: Bloomberg, HSBC estimates

Priced as of close of 21 Sep 2026

Valuation and risks

		Valuation	Risks
NVIDIA NVDA US **Buy**	Current price: USD227.38 Target price: USD365.00 Up/downside: +60.5%	**Methodology:** FY28e target P/E multiple **Assumptions:** Target FY28e P/E of 25x applied to FY28e EPS of USD14.56. Our 25x target P/E is at a 15% premium to 1SD below the five-year historical average (21x). Given the recent rise of agentic AI, stocks exposed to server CPUs are trading at a higher premium and NVIDIA, which is mainly a GPU play, has traded at an average of 19x through the year. Therefore, we take 1SD above five-year historical average (21x) as the base. We apply a 15% premium, as we believe NVIDIA has upside to current earnings estimates; better supply capacity procurement vs peers; and a new narrative emerging with its increasing focus on open-source models. Hence, we believe our 25x target P/E is justified. Our TP and all our assumptions are unchanged. **Frank Lee*** | frank.lee@hsbc.com.hk | +852 2996 6916	**Downside risks:** 1) Slower-than-expected mix shift towards higher ASP AI GPUs; 2) Lower investment from CSPs towards generative AI; 3) Weaker-than-expected overall data center momentum; and 4) Geopolitical uncertainty
Intel Corp INTC US **Buy**	Current price: USD121.78 Target price: USD200.00 Up/downside: +64.2%	**Methodology:** P/E-based SOTP **Assumptions:** Target FY27e P/E multiple of 30x (based on 2SD above the historical average P/E) applied to our core business 2027e EPS of USD4.25 implying core business value of USD125. Target FY27e P/B of 3.1x (in line with second tier foundries) applied to foundry BVPS of USD24.42 implying foundry business value of USD75.00. Our TP and all our assumptions are unchanged. **Frank Lee*** | frank.lee@hsbc.com.hk | +852 2996 6916	**Downside risks:** 1) Worse-than-expected server CPU unit shipments; 2) weaker-than-expected server CPU pricing hike; and 3) higher margin dilution due to foundry ramp.
Marvell MRVL US **Buy**	Current price: USD257.38 Target price: USD325.00 Up/downside: +26.3%	**Methodology:** FY28e target P/E multiple. **Assumptions:** Target FY28e P/E of 42x applied to FY28e EPS of USD7.70. Our 42x FY28e target P/E multiple is the multiple between 2SD above the three-year historical average and the peak P/E multiple (52x). Marvell trades at c33x FY28e P/E based on our EPS estimates, between the three-year historical average and 1SD above three-year historical average. We believe the optical interconnect revenue opportunity is still severely underestimated by the market and there could be significant earnings upside surprise driven by stronger-than-expected optical interconnect revenue growth. In addition, we think CXL demand is also underestimated, hence we expect Marvell to re-rate higher than 2SD above the 3-year historical average P/E and towards the peak P/E multiple. Our USD325 TP implies c26% upside; accordingly, we have a Buy rating on the stock. Our TP and all our assumptions are unchanged. **Frank Lee*** | frank.lee@hsbc.com.hk | +852 2996 6916	**Downside risks:** 1) Slower-than-expected growth of optical transceivers market; 2) Higher-than-expected market share loss in DSPs; and 3) Delays in custom silicon projects
Qualcomm QCOM US **Hold**	Current price: USD194.23 Target price: USD165.00 Up/downside: -15.0%	**Methodology:** FY27e target P/E multiple. **Assumptions:** Target P/E multiple of 16.3x applied to FY27e EPS of USD10.10. Our target P/E multiple of 16.3x is 1.25SD above the 5-year historical average. Given the aggressive data center ambitions, we believe the market is willing to give it a higher premium now. Our TP and all our assumptions are unchanged. **Frank Lee*** | frank.lee@hsbc.com.hk | +852 2996 6916	**Upside risks:** 1) Better-than-expected data center segment performance; 2) smartphone replacement cycle driven by AI smartphones; and 3) higher-than-expected premium handset mix **Downside risks:** 1) Weaker-than-expected smartphone demand; 2) slowdown in autos demand; and 3) geopolitical risks

		Valuation	Risks
GlobalFoundries GFS US	Current price: USD48.57 Target price: USD55.00 Up/downside: +13.2%	**Methodology:** We value the stock on a target P/B multiple. **Assumptions:** Target P/B multiple of 2.6x applied to 2027e BVPS of USD21.78. Our target P/B is higher than the 5-year historical average of 2.3x due to improving supply-demand mechanics of mature foundry. We maintain our Hold rating as we expect a delay in ASP realization. Our TP and all our assumptions are unchanged.	**Upside risks:** 1) faster improvement in economies of scale at fabs than expected; 2) faster margin improvement as a result of stronger pricing upside and better cost management than anticipated; and 3) higher government subsidies than expected. **Downside risks:** 1) lack of economies of scale and increasing variable costs; 2) pricing power capped by high exposure to LTAs; and 3) increasing competition from mainland China, arising from any policy change.
Hold			

Frank Lee* | frank.lee@hsbc.com.hk | +852 2996 6916

Priced at close of 21 Sep 2026

*Employed by a non-US affiliate of HSBC Securities (USA) Inc, and is not registered/ qualified pursuant to FINRA regulations
Source: HSBC estimates