GPU SMITH
COMPUTE ENGINEERING

GPU SMITH / RESEARCH & PRACTICAL GUIDANCE

NVIDIA China Revenue by Fiscal Year: Billing vs HQ Basis

> Plan and build private AI compute with GPU Smith. We connect workload requirements with GPU hardware, networking, deployment and operating decisions so your infrastructure fits the work it must perform.

Published by GPU Smith · 2026-11-21 · nvidia china revenue · nvidia china revenue by fiscal year · nvidia revenue by geography · customer billing location · customer headquarters basis · nvidia 10-K geographic revenue · h20 export license charge · h200 china license · nvidia data center china

Original article: https://gpusmith.com/articles/en/nvidia-china-revenue-billing-location-vs-headquarters



In this article

1. Executive Summary
2. Introduction and Background
3. The Numbers People Quote and Where Each Comes From
4. Billing Location Versus Customer Headquarters: What Changed
5. The Filing-by-Filing Ledger
6. Events Joined to Quarters
7. Data Analysis and Evidence
8. What the Filings Do Not Disclose
9. How to Read and Cite NVIDIA's China Number: An Analyst Checklist
10. Implications and Future Directions
11. Frequently Asked Questions (FAQs)
12. Conclusion

Executive Summary

NVIDIA does not publish one "China revenue" figure. Since the third quarter of fiscal year 2026 (Q3 FY2026) it publishes China (including Hong Kong) revenue on a **customer headquarters** basis, and its earlier filings used a **customer billing location** basis. The two bases give very different answers for the same year. For FY2025 (ended January 26, 2025) the billing-location table shows **$17,108 million** of **$130,497 million** total revenue, or **13.1%** (GPU Smith calculation) (www.sec.gov). After NVIDIA recast prior periods, the same year shows **$25,048 million**, or **19.2%** (www.sec.gov). The **$7,940 million** gap is a change in how sales are assigned to places, not a change in what was sold.

On the headquarters basis, China including Hong Kong was **$12,330 million** in FY2024 (20.2%) (www.sec.gov), $25,048 million in FY2025 (19.2%) and **$19,677 million** in FY2026 (**9.1%** of $215,938 million) (www.sec.gov) (investor.nvidia.com). In the first half of FY2027 it was **$12,430 million** (www.sec.gov) of $177,837 million (www.sec.gov) (7.0%), made up of **$4,550 million** in Q1 (www.sec.gov) and **$7,880 million** in Q2 (8.2%) (www.sec.gov). The Q2 FY2027 figure was **$3,985 million** a year earlier (www.sec.gov).

Two statements that look contradictory are compatible. NVIDIA's filings say the company was "effectively foreclosed" from China's data center compute market (www.sec.gov), and its outlook assumes no China data center compute revenue (nvidianews.nvidia.com). Yet the geographic table counted $7.9 billion of China-headquartered revenue in one quarter. The table covers all products and segments, and the Q2 FY2027 10-Q says Hopper data center shipments to China were under 1% of Data Center revenue (www.sec.gov). That leaves at least about **$7.0 billion** of the quarter's China-headquartered revenue outside that category (GPU Smith arithmetic); the filing does not itemize it.

NVIDIA took a **$4.5 billion** charge in Q1 FY2026 associated with H20 excess inventory and purchase obligations (nvidianews.nvidia.com). It is a cost item, not China revenue.

The practical rule is to state the basis, the fiscal period and the product scope with every China number. A billing-basis figure from before late 2025 understates the recast FY2025 figure by about six percentage points, and no filing supports treating either figure as data center compute revenue.

Introduction and Background

Investors, credit analysts, supply-chain risk teams and enterprise buyers often quote "how much of NVIDIA's revenue comes from China" to support a view on export controls, GPU supply or vendor concentration. The quoted figure is usually a single percentage with no basis attached. This report treats that as a data-quality problem and answers three questions: which figure is correct for which claim, how far a stale figure can mislead, and what the filings do and do not disclose.

All dollar figures below are company-reported, taken from NVIDIA's Form 10-K and Form 10-Q filings with the US Securities and Exchange Commission (SEC), NVIDIA's earnings releases, and its 8-K current report. Shares, year-over-year changes, sums and differences are GPU Smith calculations and are labelled as such. The report is a disclosure ledger, not investment advice, and not a forecast. Readers who rely on a number for a financial decision should open the cited filing and, where the stakes warrant it, consult a qualified financial or legal professional.

GPU Smith is an independent private AI infrastructure engineering consultancy. It advises on enterprise GPU procurement, workload sizing, networking, inference optimization, power and cooling, and air-gapped deployments. It does not manufacture chips or sell cloud capacity, and this report is independent analysis of public filings. Its interest in the topic is practical: buyers planning private AI capacity ask whether China-related disclosures change allocation, lead times or vendor risk. The rules themselves (which chips, which destinations, which licences) are covered in an earlier GPU Smith explainer on NVIDIA GPU export restrictions. That explainer quotes the older billing-location figure of about $17 billion and 13% for FY2025; this report supplies the later basis for the same year.

A word on fiscal calendars. NVIDIA's fiscal year ends in late January. FY2025 ended January 26, 2025, FY2026 ended January 25, 2026, and FY2027 will end in late January 2027. The most recent quarter reported as of **November 21, 2026** is Q3 FY2027 (ended October 25, 2026), with revenue of $115.8 billion (investor.nvidia.com). The Q3 FY2027 Form 10-Q, which carries the geographic table, had not appeared in the SEC's list of recent NVIDIA filings when this report was prepared, so Q3 FY2027 China revenue on the filing basis is not yet available.

The Numbers People Quote and Where Each Comes From

Four distinct figures circulate under the label "NVIDIA China revenue". They differ in basis, period and scope, and they are not interchangeable.

- **Billing-location China revenue.** The figure NVIDIA reported through the Q2 FY2026 10-Q. For FY2025 it is $17,108 million, **13.1%** of total revenue (GPU Smith calculation) (www.sec.gov). NVIDIA's filing defines the basis as "the billing location of the customer" (www.sec.gov).

- **Headquarters-basis China revenue.** The figure NVIDIA has reported since the Q3 FY2026 10-Q, with prior periods recast. For FY2025 it is $25,048 million (19.2%), and for FY2026 $19,677 million (9.1%) (www.sec.gov). The basis is "the location of the customers' headquarters" (www.sec.gov).
- **China as a share of Data Center revenue.** NVIDIA disclosed this directly only in older filings. The FY2024 10-K says sales to China fell from 19% of Data Center revenue in FY2023 to 14% in FY2024 (www.sec.gov), and described China as a mid-single-digit percentage in Q4 FY2024 (www.sec.gov). Later filings give no such percentage for China in total.
- **China share of the AI accelerator market.** A market-share concept, not a revenue concept. The 95% starting point comes from statements by NVIDIA's chief executive, and a 55% figure for 2025 comes from IDC data relayed by trade press. These are discussed in the data section below; they cannot be converted to a share of NVIDIA revenue.

Table 1 below lays the first two series side by side for each fiscal year in which a figure exists. It is the core of the ledger.

Fiscal year (ended)	Total revenue	China incl. HK, billing basis (share)	China incl. HK, headquarters basis (share)	HQ minus billing
FY2023 (Jan 29, 2023)	$26,974M (www.sec.gov)	$5,785M (21.4%) (www.sec.gov)	Not published	n/a
FY2024 (Jan 28, 2024)	$60,922M (www.sec.gov)	$10,306M (16.9%) (www.sec.gov)	$12,330M (20.2%) (www.sec.gov)	+$2,024M
FY2025 (Jan 26, 2025)	$130,497M (www.sec.gov)	$17,108M (13.1%) (www.sec.gov)	$25,048M (19.2%) (www.sec.gov)	+$7,940M
FY2026 (Jan 25, 2026)	$215,938M (investor.nvidia.com)	Not published (basis retired)	$19,677M (9.1%) (www.sec.gov)	n/a

Table 1. NVIDIA China (including Hong Kong) revenue by fiscal year on both reporting bases. Dollar figures are company-reported; shares, the gap column and all percentages are GPU Smith calculations. Billing-basis FY2023 to FY2025 values appear in the FY2025 10-K table and the earlier 10-Ks; headquarters-basis FY2024 to FY2026 values appear in the FY2026 10-K.

Read the table by row, not by column. In FY2025 the two bases differ by **6.1 percentage points** (13.1% against 19.2%). In FY2024 they differ by 3.3 points. A reader who takes the billing series through FY2025 and the headquarters series for FY2026 will see a fall from 13.1% to 9.1%, about 30% in relative terms, and dollars rising from $17,108 million to $19,677 million (up 15.0%). On a consistent headquarters basis the share fell from 19.2% to 9.1%, and dollars fell from $25,048 million to $19,677 million, a **21.4% decline** (GPU Smith calculations). The spliced reading understates the share decline (4.0 points against 10.1 points) and gets the direction of the dollar change wrong.

Billing Location Versus Customer Headquarters: What Changed

The exact filing wording

Through FY2025, NVIDIA assigned geographic revenue by where the customer was billed. The FY2025 10-K states that "revenue by geographic area is based upon the billing location of the customer" and adds that the end customer and shipping location may differ (www.sec.gov). The FY2024 and FY2023 10-Ks use similar billing wording, and the FY2023 table is described as based on the invoicing address (www.sec.gov).

The change was announced in the Q3 FY2026 10-Q (quarter ended October 26, 2025). That filing says revenue by geographic area is now based on the location of the customers' headquarters, notes that it was "previously" reported by billing location, and says prior period information was recast (www.sec.gov). The FY2026 10-K repeats the change and gives NVIDIA's stated reason: it "believe[s] it provides a better representation of the geographic profile of our revenue" (www.sec.gov). The MD&A states that regions are designated by "the headquarters of direct customers" (www.sec.gov).

Neither basis is an end-customer or shipping-location basis. The FY2026 10-K says the end customer and shipping location may differ from the headquarters location (www.sec.gov). The Q2 FY2027 10-Q repeats that caveat for direct customers (www.sec.gov).

Why the FY2025 China number moved by $7.9 billion

The FY2025 10-K billing table has five lines: United States **$61,257 million** (www.sec.gov), **Singapore** $23,684 million (www.sec.gov), Taiwan $20,573 million (www.sec.gov), China including Hong Kong $17,108 million (www.sec.gov) and Other $7,875 million (www.sec.gov), summing to $130,497 million. The recast headquarters table has four lines for the same year: United States **$77,482 million** (www.sec.gov), Taiwan $23,600 million (www.sec.gov), China including Hong Kong $25,048 million (www.sec.gov) and Other $4,367 million (www.sec.gov). Singapore no longer appears.

Comparing the two tables line by line, as GPU Smith arithmetic and not as a decomposition the filings provide:

- **United States:** up **$16,225 million** (from $61,257 million to $77,482 million).
- **Taiwan:** up **$3,027 million** (from $20,573 million to $23,600 million).
- **China including Hong Kong:** up $7,940 million (from $17,108 million to $25,048 million).
- **Singapore plus Other:** down **$27,192 million** (from $31,559 million to $4,367 million).

The four changes net to zero because total revenue is unchanged at $130,497 million. The filings explain the direction of the shift without apportioning it. The FY2025 10-K says that customers use Singapore to centralize invoicing while products are almost always shipped elsewhere, that Singapore was 18% of FY2025 revenue on a billing basis, and that shipments to Singapore were under 2% of revenue (www.sec.gov). It also says revenue is designated by billing location "even if the revenue may be attributable to indirect customers in a different location" (www.sec.gov). The filings do not say how much of the Singapore and Other revenue moved to China. Any number a third party attaches to that split is its own estimate.

The Taiwan footnote

The headquarters basis does not remove the gap between a customer's address and the place a system is used. The FY2026 10-K footnote on Taiwan estimates that 76% of Data Center revenue from Taiwan-headquartered customers was attributed to end customers elsewhere (www.sec.gov). The Q3 FY2026 10-Q gave 86% for that quarter (www.sec.gov). The consequence for China readers is symmetrical: a China-headquartered customer can buy for deployment outside China, and a customer headquartered elsewhere can buy for use in China. NVIDIA's table captures neither.

The Filing-by-Filing Ledger

Annual series and the older Data Center disclosures

Older filings contain more China context than recent ones. The FY2023 10-K called the impact of the August 2022 licence requirements "difficult to quantify" (www.sec.gov). It reported Data Center revenue of $15.01 billion (www.sec.gov). The FY2024 10-K recorded the October 2023 rule update and the 14% China share of Data Center revenue noted above (www.sec.gov). The FY2025 10-K said Data Center revenue in China grew in FY2025 but remained well below the levels before export controls tightened in October 2023, and it did not give a percentage (www.sec.gov).

The share of revenue from customers headquartered outside the United States fell from 48% in FY2024 to 41% in FY2025 and 31% in FY2026 on the headquarters basis (www.sec.gov). That is a broader statement than China. It shows that the denominator is increasingly dominated by US-headquartered customers, which is one reason China's share falls even when China's dollars rise.

Quarterly series on the headquarters basis

Table 2 below gives the quarterly figures that the 10-Qs make available. Q1 and Q2 FY2026 were originally reported on the billing basis, so each row shows both where both exist.

Quarter (ended)	Total revenue	China incl. HK, headquarters basis (share)	China incl. HK, billing basis as first reported (share)
Q1 FY2026 (Apr 27, 2025)	$44.1B (investor.nvidia.com)	$9,659M (21.9%) (www.sec.gov)	$5,522M (12.5%) (www.sec.gov)
Q2 FY2026 (Jul 27, 2025)	$46.7B (investor.nvidia.com)	$3,985M (8.5%) (www.sec.gov)	$2,769M (5.9%) (www.sec.gov)
Q3 FY2026 (Oct 26, 2025)	$57.0B (investor.nvidia.com)	$2,973M (5.2%) (www.sec.gov)	Not published
Q4 FY2026 (Jan 25, 2026)	$68.1B (investor.nvidia.com)	$3,060M (4.5%), derived	Not published
Q1 FY2027 (Apr 26, 2026)	$81.6B (investor.nvidia.com)	$4,550M (5.6%) (www.sec.gov)	Not published
Q2 FY2027 (Jul 26, 2026)	$96.2B (investor.nvidia.com)	$7,880M (8.2%) (www.sec.gov)	Not published

Quarter (ended)	Total revenue	China incl. HK, headquarters basis (share)	China incl. HK, billing basis as first reported (share)
Q3 FY2027 (Oct 25, 2026)	$115.8B (investor.nvidia.com)	Not yet available (10-Q not located as of Nov 21, 2026)	Not published

Table 2. Quarterly China (including Hong Kong) revenue. The Q4 FY2026 value is a GPU Smith derivation: the FY2026 total of $19,677 million less the nine-month total of $16,617 million (www.sec.gov). Shares are GPU Smith calculations. Headquarters-basis Q1 and Q2 FY2026 values are the prior-year comparatives in the FY2027 10-Qs.

Two observations follow from the table. First, the basis effect is not small even within a year. In Q1 FY2026 the headquarters figure (**$9,659 million**) is **$4,137 million** above the figure first reported on the billing basis ($5,522 million), and in Q2 FY2026 it is **$1,216 million** above ($3,985 million against $2,769 million). Summed, the first half of FY2026 is **$13,644 million** on the headquarters basis against **$8,291 million** as first reported (www.sec.gov). Second, the quarterly series is volatile. It fell from $9,659 million in Q1 FY2026 to **$2,973 million** in Q3 FY2026, then rose to $7,880 million in Q2 FY2027, a **98% increase** on the prior-year quarter (GPU Smith calculation) (www.sec.gov).

The Q1 FY2027 10-Q gives one quantified link between the series and a product category: no Data Center Hopper products shipped to China in Q1 FY2027, compared with $4.6 billion in Q1 FY2026 (www.sec.gov). Taking that $4.6 billion as a rounded figure, the Q1 decline in the headquarters series ($9,659 million to $4,550 million, down $5,109 million) is of the same order, and China-headquartered revenue other than Hopper shipments to China was roughly $5.1 billion in Q1 FY2026 against $4.55 billion in Q1 FY2027, when no such shipments occurred (GPU Smith arithmetic, illustrative only and subject to rounding). The 10-Q does not itself tie the table to that statement.

Events Joined to Quarters

The events below are the ones NVIDIA itself discloses or that government and wire sources date precisely. ***Table 3*** places each in the fiscal quarter in which it occurred.

Date	Fiscal quarter	Event
Aug 2022	Q3 FY2023	First licence requirements affecting exports to China (including Hong Kong and Macau) and Russia of A100 and H100 integrated circuits (www.sec.gov)
Oct 2023	Q3 FY2024 (announced); the 10-K places their effect in Q4 FY2024	Updated licensing requirements announced (www.sec.gov)
Jan 28, 2024 (quarter end)	Q4 FY2024	China a mid-single-digit share of Data Center revenue in the quarter (www.sec.gov)
Apr 15, 2025	Q1 FY2026	8-K anticipates up to about $5.5 billion of Q1 charges (www.sec.gov)

Date	Fiscal quarter	Event
May 28, 2025	Q1 FY2026	Release reports a $4.5 billion H20 charge (nvidianews.nvidia.com) and $2.5 billion of H20 revenue that could not be shipped (nvidianews.nvidia.com); Q2 outlook reflects an H20 revenue loss of about $8.0 billion (investor.nvidia.com)
Aug 27, 2025	Q2 FY2026	No H20 sales to China-based customers in Q2 (nvidianews.nvidia.com); about $650 million of H20 sold to a customer outside China (investor.nvidia.com)
Nov 19, 2025	Q3 FY2026	10-Q switches to headquarters basis and recasts prior periods (www.sec.gov)
Jan 14 to 15, 2026	Q4 FY2026	White House fact sheet describes a 25% tariff on certain advanced chips including H200 (www.whitehouse.gov); Commerce rule moves H200-class licence review to case-by-case (www.govinfo.gov)
Jul 26, 2026 (quarter end)	Q2 FY2027	Hopper shipments to China under 1% of Data Center revenue (www.sec.gov); $0.4 billion H200 charge in the first half (www.sec.gov)

Table 3. Policy and disclosure events joined to NVIDIA fiscal quarters.

Table 3 shows why event timing matters for the ledger. The $4.5 billion charge is a Q1 FY2026 cost-of-revenue item, not a China revenue figure. NVIDIA's Q1 FY2026 10-Q notes the charge was lower than initially anticipated because certain materials could be re-used (www.sec.gov), and the release says that excluding it the non-GAAP gross margin would have been 71.3% (investor.nvidia.com). A reader who treats "$4.5 billion" as lost China revenue has confused a write-down with a sales figure. The revenue effect NVIDIA quantified is the **$2.5 billion** it could not ship in Q1 and the roughly **$8.0 billion** loss reflected in the Q2 outlook.

Data Analysis and Evidence

How far a stale figure can mislead

Using only the filings above, *Table 1* and *Table 2* imply the following magnitudes (all GPU Smith calculations):

- **FY2025 understated by 6.1 points.** Quoting the billing-basis 13.1% for FY2025 instead of the recast 19.2% understates China's share by 6.1 percentage points and its dollar value by $7,940 million.
- **FY2024 understated by 3.3 points.** Billing 16.9% against headquarters 20.2%, a $2,024 million gap.
- **Q1 FY2026 understated by 9.4 points.** Billing 12.5% against headquarters 21.9%, a $4,137 million gap in one quarter.
- **Spliced series misstate FY2026.** Comparing the FY2026 headquarters share (9.1%) with the FY2025 billing share (13.1%) shows a 4.0-point fall and a 15.0% dollar increase. The consistent headquarters series shows a 10.1-point fall and a 21.4% dollar decline.
- **Quarterly range.** Since Q1 FY2026 the headquarters-basis quarterly share has ranged from 21.9% (Q1 FY2026) to 4.5% (Q4 FY2026, derived), a spread of **17.4 points**, so a single quarter is a poor proxy for "the" China share.

Customer concentration is a different problem again. The FY2026 10-K reports that one direct customer represented 22% of total revenue and another 14% (www.sec.gov), and the Q2 FY2027 10-Q reports three direct customers at 16%, 15% and 13% of first-half revenue (www.sec.gov). Neither filing identifies a China-

specific customer concentration. Anyone using China revenue as a proxy for customer concentration is using the wrong table.

The foreclosure statement and the $7.9 billion quarter

NVIDIA's risk factors say that, as of the end of Q2 FY2027, it was effectively foreclosed from competing in China's data center computing market (www.sec.gov). The same wording appears in the FY2026 10-K (www.sec.gov). The earnings releases for Q4 FY2026 through Q3 FY2027 each say the outlook assumes no Data Center compute revenue from China (investor.nvidia.com) (s201.q4cdn.com). These statements concern data center compute.

The geographic table is broader. It allocates all revenue (all segments, all product lines, including graphics, networking, software and non-compute products) by the headquarters of the direct customer. Q2 FY2027 total revenue was $96,221 million, of which $7,880 million was assigned to China including Hong Kong (www.sec.gov). Data Center revenue in the quarter was $89.0 billion (www.sec.gov). If Hopper shipments to China were less than 1% of that figure, they were below about $890 million (GPU Smith calculation), which leaves at least about $7.0 billion of China-headquartered revenue that is not Hopper data center shipments to China. The filings do not itemize it. The Q2 FY2027 10-Q names one category: it says that "while we were able to ship uncontrolled products to China, such as gaming and workstation GPUs," the company was effectively foreclosed from China's data center compute market, but it does not quantify those uncontrolled-product sales. Revenue from China-headquartered customers who take delivery outside China is another plausible category, but no filing quantifies it and it should not be asserted as fact.

The 95% to 55% market-share claim

A frequently repeated claim says NVIDIA's China AI accelerator share fell from about 95% to about 55%. Tracing the numbers shows two different measurements joined together. The 95% comes from NVIDIA's chief executive. Reuters reported in May 2025 that he said NVIDIA's China market share had dropped to 50% from 95% (www.reuters.com), and TrendForce reported the 95% as the share held four years earlier (www.trendforce.com). By October 2025 Fortune quoted him as saying NVIDIA went from 95% market share to 0% (fortune.com). The 55% figure comes from IDC data relayed by Tom's Hardware, which reports an estimate that NVIDIA still led with 55% of the China market, shipping an estimated 2.2 million cards (www.tomshardware.com).

These are market-share estimates in units or in an unspecified measure, from different dates and sources. They are not shares of NVIDIA's revenue and cannot be reconciled with the 9.1% or 19.2% figures above. The 95% starting point is NVIDIA's own claim and the report found no independent verification of it. This report therefore does not adopt the 95% to 55% pairing as a fact.

What the Filings Do Not Disclose

The disclosure gaps matter as much as the figures.

- **China data center compute revenue.** No filing reports it. The geographic table covers all products, and the Q2 FY2027 10-Q gives only a ceiling for Hopper shipments to China (under 1% of Data Center revenue) (www.sec.gov).
- **End-customer geography.** The tables are by direct-customer headquarters, and NVIDIA says end-customer and shipping locations may differ (www.sec.gov).
- **A decomposition of the basis change.** The filings give no bridge from the $17,108 million billing figure to the $25,048 million headquarters figure for FY2025 (www.sec.gov).
- **Singapore after the change.** Singapore is not a separate line on the headquarters basis; the FY2026 10-K table has only four lines (www.sec.gov).
- **Q3 FY2027 China revenue.** The earnings release carries total and Data Center revenue (investor.nvidia.com) but not the geographic table. That table will appear in the Form 10-Q.
- **H200 dollar amounts.** The 10-Qs give only percentages and a **$0.4 billion** H200 charge (www.sec.gov).
- **Future policy.** The Q3 FY2027 call noted that new US export controls were released the same week and that management was still studying them (www.nasdaq.com). Nothing in this report assumes their content.

One timing risk applies to every series here: NVIDIA's outlook for Q4 FY2027 excludes China Data Center compute revenue (nvidianews.nvidia.com), so forward statements from the company do not forecast the geographic table either way.

How to Read and Cite NVIDIA's China Number: An Analyst Checklist

The following checklist is a minimum standard for quoting a China figure. It uses only information in the filings.

1. **State the basis.** Write "customer headquarters basis" or "customer billing location basis". The two series are not comparable.
2. **State the fiscal period and end date.** FY2025 means the year ended January 26, 2025. Quarters are fiscal quarters, not calendar quarters.
3. **Never splice bases.** If a series crosses Q3 FY2026, use the recast headquarters series from the FY2026 10-K and FY2027 10-Qs.
4. **Say what is included.** The line is "China (including Hong Kong)", all segments and products.
5. **Do not call it data center compute.** No filing says it is.
6. **Label derived values.** Shares, differences, quarterly derivations (such as Q4 FY2026) and half-year sums are the analyst's calculations.
7. **Name the filing.** Cite the specific 10-K or 10-Q, because the same year carries two values.
8. **Check the filing date.** Press articles and databases often retain billing-basis numbers.
9. **Separate charges from revenue.** A $4.5 billion inventory and purchase-obligation charge is not a revenue loss (www.sec.gov).
10. **Note the unknowns.** End-customer geography is not disclosed.

Implications and Future Directions

For buyers of private AI infrastructure, the first implication is that "China exposure" in a vendor's financial statements is a poor guide to supply allocation. Headquarters-basis revenue allocated to China tells a buyer little about where capacity is shipped, and the **Taiwan footnote** shows that a large share of revenue booked to one headquarters location is attributed to end customers elsewhere (www.sec.gov). A buyer concerned about allocation should ask a vendor about delivery timelines and licence conditions, not read a geography table as a queue position.

The second implication concerns Chinese-headquartered customers buying overseas. The headquarters basis assigns such purchases to China even if systems are deployed elsewhere, and the billing basis assigned them to wherever the invoice went. Buyers and compliance teams should treat both as address-based labels. NVIDIA's FY2025 10-K describes licence requirements that reach a party headquartered in, or with an ultimate parent headquartered in, a restricted country group (www.sec.gov); the earlier GPU Smith explainer on export restrictions covers that scope.

The third is vendor concentration. The Q2 FY2027 10-Q shows three direct customers at 16%, 15% and 13% of first-half revenue (www.sec.gov), and the share of revenue from customers headquartered outside the United States was 38% in Q2 FY2027 and 30% for the half (www.sec.gov). For a risk team, concentration among a few US-headquartered buyers is a larger number than China's share.

Looking forward, two items should be rechecked as filings appear: the Q3 FY2027 10-Q table; and any further 10-Q wording on H200 sales and the 25% tariff condition (www.sec.gov).

Frequently Asked Questions (FAQs)

What is NVIDIA's China revenue by fiscal year?

On the customer headquarters basis: FY2024 $12,330 million, FY2025 $25,048 million, FY2026 $19,677 million (www.sec.gov). On the billing basis: FY2023 $5,785 million, FY2024 $10,306 million, FY2025 $17,108 million. All figures include Hong Kong.

What percentage of NVIDIA's total revenue is China?

FY2026: 9.1% on the headquarters basis; FY2025: 19.2% (headquarters) or 13.1% (billing); first half of FY2027: 7.0% on the headquarters basis (all GPU Smith calculations from company-reported figures) (www.sec.gov).

Where is NVIDIA revenue by geography in the 10-K?

In the segment information note, as a table of revenue by geographic area with United States, Taiwan, China (including Hong Kong) and Other lines in the FY2026 10-K (www.sec.gov); the FY2025 10-K table has a separate Singapore line.

What is the difference between billing location and customer headquarters?

Billing location is the address where the customer is invoiced; headquarters is the location of the customer's headquarters, and NVIDIA says it adopted it as a better representation of the geographic profile of revenue (www.sec.gov). Neither identifies where systems are used.

What was NVIDIA's China revenue in FY2025 and FY2026?

FY2025: $17,108 million as first reported and $25,048 million as recast; FY2026: $19,677 million (www.sec.gov).

How did H20 export restrictions affect data center revenue?

NVIDIA disclosed a $4.5 billion Q1 FY2026 charge, $2.5 billion of H20 revenue it could not ship in that quarter and an outlook that reflected a roughly $8.0 billion H20 revenue loss in Q2 (nvidianews.nvidia.com) (investor.nvidia.com). It reported no H20 sales to China-based customers in Q2 FY2026 (www.sec.gov).

Is China included in NVIDIA's outlook?

NVIDIA says its outlook does not assume Data Center compute revenue from China (nvidianews.nvidia.com). That is a statement about data center compute, not about all China-headquartered revenue.

Does the filing show end customers?

No. NVIDIA warns that end customer and shipping location may differ from the customer's headquarters location (www.sec.gov).

Conclusion

The correct China figure depends on the claim. For market exposure measured by customer headquarters, use the recast series: $25,048 million (19.2%) for FY2025, $19,677 million (9.1%) for FY2026 and $12,430 million (7.0%) for the first half of FY2027. For continuity with filings before late 2025, the billing-basis series ($17,108 million, 13.1%, for FY2025) remains correct for its stated basis but is stale as a measure of headquarters exposure. For data center compute exposure, no filing provides a figure, and NVIDIA's own wording is that it was effectively foreclosed from that market. For customer concentration, the relevant disclosure is the direct-customer table, not the geography table.

A stale figure misleads in a predictable way. It can understate the recast FY2025 share by 6.1 points, it can reverse the apparent direction of year-over-year dollars when spliced to the new basis, and it can attach a policy event, such as the $4.5 billion H20 charge, to a revenue number that it never measured (nvidianews.nvidia.com). The quarterly series is volatile, with Q2 FY2027 at $7,880 million after $3,985 million a year earlier, so a single quarter is no substitute for a stated period.

Analysts should state the basis, the period and the product scope, cite the filing, and label derived values as their own. Readers making financial decisions should consult the filings directly and a qualified professional. The Q3 FY2027 10-Q will extend the ledger, and this report should be rechecked when it appears.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://www.sec.gov/Archives/edgar/data/1045810/000104581025000023/nvda-20250126.htm
2. https://www.sec.gov/Archives/edgar/data/1045810/000104581026000021/nvda-20260125.htm
3. https://investor.nvidia.com/news/press-release-details/2026/NVIDIA-Announces-Financial-Results-for-Fourth-Quarter-and-Fiscal-2026/default.aspx
4. https://www.sec.gov/Archives/edgar/data/1045810/000104581026000075/nvda-20260726.htm
5. https://www.sec.gov/Archives/edgar/data/1045810/000104581026000052/nvda-20260426.htm
6. https://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-third-quarter-fiscal-2027
7. https://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-first-quarter-fiscal-2026
8. https://investor.nvidia.com/news/press-release-details/2026/NVIDIA-Announces-Financial-Results-for-Third-Quarter-Fiscal-2027
9. https://www.sec.gov/Archives/edgar/data/1045810/000104581025000230/nvda-20251026.htm
10. https://www.sec.gov/Archives/edgar/data/1045810/000104581024000029/nvda-20240128.htm
11. https://www.sec.gov/Archives/edgar/data/1045810/000104581023000017/nvda-20230129.htm
12. https://investor.nvidia.com/news/press-release-details/2025/NVIDIA-Announces-Financial-Results-for-First-Quarter-Fiscal-2026/default.aspx
13. https://www.sec.gov/Archives/edgar/data/1045810/000104581025000116/nvda-20250427.htm
14. https://investor.nvidia.com/news/press-release-details/2025/NVIDIA-Announces-Financial-Results-for-Second-Quarter-Fiscal-2026/default.aspx
15. https://www.sec.gov/Archives/edgar/data/1045810/000104581025000209/nvda-20250727.htm
16. https://investor.nvidia.com/news/press-release-details/2025/NVIDIA-Announces-Financial-Results-for-Third-Quarter-Fiscal-2026/default.aspx
17. https://investor.nvidia.com/news/press-release-details/2026/NVIDIA-Announces-Financial-Results-for-First-Quarter-Fiscal-2027/default.aspx
18. https://investor.nvidia.com/news/press-release-details/2026/NVIDIA-Announces-Financial-Results-for-Second-Quarter-Fiscal-2027/default.aspx
19. https://www.sec.gov/Archives/edgar/data/1045810/000104581025000082/nvda-20250409.htm
20. https://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-second-quarter-fiscal-2026
21. https://www.whitehouse.gov/fact-sheets/2026/01/fact-sheet-president-donald-j-trump-takes-action-on-certain-advanced-computing-chips-to-protect-americas-economic-and-national-security
22. https://www.govinfo.gov/content/pkg/FR-2026-01-15/html/2026-00789.htm
23. https://s201.q4cdn.com/141608511/files/doc_financials/2027/q1/NVDA-Q1-2027-Earnings-Call-20-May-2026-5_00-PM-ET.pdf
24. https://www.reuters.com/world/china/nvidia-says-us-export-controls-ai-china-were-a-failure-2025-05-21
25. https://www.trendforce.com/news/2025/05/21/news-jensen-huang-nvidias-china-market-share-plummets-from-95-to-50-amid-u-s-export-curbs
26. https://fortune.com/2025/10/19/jensen-huang-nvidia-china-market-share-ai-chips-trump-trade-war

27. https://www.tomshardware.com/tech-industry/nvidia-market-share-in-china-falls-to-less-than-60-percent-chinese-chip-makers-deliver-1-65-million-ai-gpus-as-the-government-pushes-data-centers-to-use-domestic-chips
28. https://www.nasdaq.com/articles/nvidia-nvda-q3-2027-earnings-call-transcript

THE PUBLISHER

About GPU Smith

GPU Smith provides independent engineering and research for private AI infrastructure. We help technical buyers and operators reason about GPU workload sizing, hardware selection, procurement, deployment and inference operations, from the compute system to the networking, power and cooling requirements around it.

Start with the workload

Useful infrastructure decisions begin with the models, data, throughput, latency and operating constraints a team actually has. GPU Smith helps connect those requirements with system design and deployment choices. The goal is a privately controlled AI environment whose capacity and operational demands are understood before a purchasing decision.

Hardware and supplier research

Our public reference library covers GPUs, complete systems and networking components. The server-vendor directory and manufacturing research help teams investigate suppliers, compare options and follow sources behind technical claims. We distinguish manufacturer specifications from measured benchmarks and advertised capabilities from tested configurations. Reference pages are research resources, not a live inventory listing or a binding equipment quote.

Deployment and operations

GPU Smith's engineering scope includes infrastructure integration, networking, power, cooling and the practical operation of inference workloads. Our published articles explain the tradeoffs behind deployment, procurement and ongoing operations so teams can ask better questions and document decisions.

Work with GPU Smith

Explore the hardware library, GPU references, networking references and vendor research. Contact GPU Smith with your workloads and deployment constraints to discuss a project and confirm current scope, pricing and availability.

Inclusion of a manufacturer or product in our research does not imply a partnership, certification or endorsement.

Website: https://gpusmith.com

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.