NVIDIA Supply Chain: Every Critical Dependency Mapped from SEC Filings

Supply chain - Supply Chain Analysis - October 7, 2026

How TSMC concentration, HBM memory scarcity, and packaging constraints shape NVIDIA's risk profile - and what the 10-K filings actually disclose

NVIDIA's fiscal year 2024 revenue reached $60.9 billion, up 122% year-over-year, almost entirely on the back of a single product family: H100 and H200 data center GPUs. That kind of vertical demand acceleration stress-tests any supply chain. NVIDIA's in particular has three structural concentration points that the company discloses - with varying degrees of candor - across its annual 10-K filings.

Understanding those dependencies requires reading the filings carefully, because the word "TSMC" does not appear in NVIDIA's 10-K. The company refers to "Taiwan-based manufacturers" and a "limited number of foundry partners." The relationships are there in the data - you just have to know where to look.

The Foundry: 100% TSMC at Advanced Nodes

NVIDIA designs chips but manufactures nothing. Every GPU that powers a data center - A100, H100, H200, the Blackwell B100 and B200 - is fabricated at Taiwan Semiconductor Manufacturing Company. NVIDIA's 10-K Risk Factors section has disclosed this concentration clearly since at least fiscal 2021, using language like:

"We do not have long-term supply agreements with TSMC... we are subject to risks related to our dependence on Taiwan-based manufacturers, including geopolitical and business uncertainties."

The H100 uses TSMC's N4 process (4nm class). The H200 and Blackwell GPUs use TSMC's N3 and CoWoS-L advanced packaging. No other foundry on earth offers comparable capacity at these nodes. Samsung's 3nm gate-all-around process has yield and performance gaps at this scale. Intel Foundry is still in ramp-up mode. TSMC is not a vendor NVIDIA can swap if relations between the U.S. and Taiwan deteriorate - it's a structural dependency with no short-term alternative.

Filing language to watch: NVIDIA's 10-K consistently uses the phrase "Taiwan-based manufacturers" rather than naming TSMC directly. TraderTools extracts these relationship signals from the underlying filing text and maps them to named counterparties using entity resolution across the full EDGAR corpus.

CoWoS: The Packaging Bottleneck That Actually Constrained Supply

In 2023 and through much of 2024, the binding constraint on H100 shipments was not chip wafer capacity - it was CoWoS packaging. Chip-on-Wafer-on-Substrate is TSMC's advanced 2.5D packaging technology that connects the GPU die to High Bandwidth Memory stacks on a silicon interposer. Without CoWoS, you cannot assemble an H100.

TSMC's CoWoS capacity in 2023 was approximately 10,000-12,000 wafers per month across all customers. NVIDIA's H100 demand alone required more than that. The constraint forced NVIDIA to allocate supply and contributed directly to the 8-12 month lead times that hyperscalers were quoting through 2023.

NVIDIA's 10-K Risk Factors disclosed this indirectly: "Our product supply is significantly dependent on our manufacturing partners' ability to procure raw materials and components, and their

manufacturing capacity and production yields." The CoWoS-specific bottleneck was discussed on analyst calls but not explicitly named in filings - a common gap between what's said and what's formally disclosed.

HBM Memory: Three Suppliers, One Dominant

High Bandwidth Memory - HBM3 in H100, HBM3E in H200 and Blackwell - is the second critical dependency. HBM is stacked DRAM manufactured by only three companies globally: SK Hynix, Samsung, and Micron. Each H100 requires six HBM3 stacks totaling 80GB; the H200 uses 141GB of HBM3E.

SK Hynix was first to market with HBM3 in volume and held an estimated 50%+ share of NVIDIA's HBM supply through 2023 and into 2024. Samsung ramped HBM3E but faced qualification delays; reports indicated yield issues pushed more volume toward SK Hynix. Micron qualified its HBM3E with NVIDIA in mid-2024 and began ramping as a third source.

NVIDIA does not name its memory suppliers in 10-K filings. The Risk Factors section uses language like: "We depend on third parties to supply components, including memory, to our manufacturing partners." The concentration is visible only through cross-referencing counterparty filings and supply chain data - exactly what TraderTools is built to do.

Substrate and PCB: The Hidden Tier-2 Layer

Beneath the chip and memory layer, GPU production depends on advanced flip-chip BGA substrates. The primary suppliers are Ibiden (Japan) and Unimicron (Taiwan), with secondary capacity at Nan Ya PCB and AT&S. Substrate lead times for advanced AI GPU designs run 20-26 weeks. NVIDIA's 10-K does not name substrate suppliers - this tier-2 layer typically surfaces only in an 8-K when a disruption causes a material impact.

Customer Concentration: The Hyperscaler Dependency

On the demand side, NVIDIA's data center business is concentrated in a small number of hyperscale cloud customers. The company's 10-K FY2024 disclosed that a single distributor accounted for more than 10% of revenue. The actual end customers - Microsoft Azure, Google Cloud, Amazon Web Services, Meta - are not individually disclosed by name because NVIDIA sells through distributors and ODMs.

The concentration is visible in aggregate: "Our top ten customers represented approximately 57% of total revenue." (NVDA 10-K FY2024, Item 1.) Given that AI infrastructure buildout is dominated by four hyperscalers, the implied customer concentration is substantial even if no single end-customer crosses the 10% disclosure threshold.

What TraderTools maps for NVDA: The NVIDIA chain page shows verified customer and supplier relationships extracted from 10-K, 10-Q, and S-1 filings across the EDGAR corpus - including counterparties disclosed in NVIDIA's customers' own filings. Microsoft's 10-K references NVIDIA as a critical hardware supplier; that relationship is mapped from both sides.

The Taiwan Risk Factor: What the 10-K Actually Says

NVIDIA's Taiwan concentration risk language has grown more specific over successive annual filings. The FY2024 10-K contains a dedicated risk factor covering: military conflict between China and Taiwan, natural disasters affecting Hsinchu and Southern Taiwan Science Parks where TSMC's advanced fabs are concentrated, export controls that could disrupt cross-border component flows, and TSMC's own ability to maintain production yields at N3 and below.

This is not boilerplate. Companies that have genuinely assessed Taiwan dependency write specific, multi-paragraph risk factors. Companies that haven't write one-sentence generic language. NVIDIA's disclosure indicates genuine internal assessment of the dependency - and genuine inability to mitigate it in the near term.

The Full Picture

NVIDIA is simultaneously a customer (to TSMC, SK Hynix, Ibiden) and a supplier (to the hyperscalers building AI infrastructure). Microsoft's 10-K FY2024 discloses NVIDIA as a "key hardware provider." Google's 10-K contains similar language about GPU availability. This cross-filing visibility - where NVIDIA's customers name NVIDIA as a critical dependency in their own filings - is one of the most valuable signals TraderTools extracts from the EDGAR corpus. It gives you the relationship from two perspectives simultaneously.

For investors analyzing NVIDIA or the companies dependent on its products: single-source foundry with no near-term alternative, packaging constraint that eases slowly, three memory suppliers with one dominant, customer base concentrated in four hyperscalers. The NVIDIA supply chain map on TraderTools shows every verified relationship with a primary filing citation.

Generated 2026-11-16 by TraderTools.pro