COMPANY DEEP DIVE

Nvidia

The world's most valuable company is still doubling its revenue. A look at the business, the numbers behind the latest quarter, the financing that now surrounds its customers, and what a $5.6 trillion valuation actually assumes.

Global Markets Research · Technology Desk Prepared for internal & client distribution by ANDRA MARKETS

Coverage summary. Nvidia reported revenue of $96.2 billion for the quarter to late July, more than double the same quarter a year earlier, and guided to $108 billion for the current one. It earns a gross margin of 75 per cent, converts most of its profit into cash and is returning more than $25 billion a quarter to shareholders. On those numbers alone the stock, at around $234 and a market value of about $5.6 trillion, is not obviously expensive. The debate has moved elsewhere: to whether the spending of its customers, now consuming almost all of the hyperscalers' operating cash flow and increasingly funded with debt, can be sustained; to the growing role Nvidia itself plays in financing that demand; and to how long a 75 per cent margin can survive the arrival of credible alternatives. This note works through each in turn.

SHARE PRICE / MKT CAP	Q2 FY27 REVENUE	GROSS MARGIN	P/E, RUN-RATE
$234 / $5.6tn	**$96.2bn**	**75.0%**	**~24x**
world's largest company	*+106% year on year*	*guided lower into FY28*	*~33x trailing, ~15x FY28 plan*

Exhibit A. Key figures (fiscal year ends late January; Q2 FY27 is the quarter to 26 July 2026)

Metric	Figure	Note
Revenue, Q2 FY27	$96.2bn	+18% q/q, +106% y/y
Data center revenue	$89.0bn	93% of total, +117% y/y
Revenue guidance, Q3 FY27	$108.0bn ±2%	excludes China data center compute
Gross margin, Q2 / Q3 guide	75.0% / 74.0%	Q4 trough of 71% to 72% flagged
GAAP net income, Q2	$59.7bn	includes $7.8bn gains on equity stakes
Free cash flow, first half FY27	$69.9bn	operating cash flow $74.4bn
Capital returned, Q2	~$26bn	buybacks and dividend; $99bn authorisation left
Equity investment portfolio	~$95bn	public and private stakes
Market capitalisation	~$5.65tn	share price $233.95 (2 October close)

Figures are from company releases and press reports; multiples and run-rates are Andra Markets calculations and are approximate. Colour convention (charts only): green marks growth or the upside case, red marks risk or the downside case. All body and table text is set in black.

1.0 Executive Summary

Key takeaways

The growth has not slowed; it has accelerated. Revenue in the July quarter was more than double the year before and grew 18 per cent in three months, a remarkable rate for a company of this size. Guidance for the October quarter implies another 12 per cent, and management has indicated it is planning for roughly 70 per cent growth in the next fiscal year, with demand it describes as well ahead of what supply chains for power, memory and advanced packaging can deliver.

The new Vera Rubin platform is entering full production with orders from every major customer, and Nvidia estimates it will capture about $40 billion of revenue per gigawatt of data center capacity, against roughly $25 billion for the current generation.

The quality of that growth is the real question. Three features deserve attention. First, demand is concentrated in a handful of buyers whose capital spending, around $730 billion this year for the four largest, now absorbs nearly all of their operating cash flow and is increasingly debt-funded. Second, Nvidia has become a financier of its own demand: it holds about $95 billion of stakes in customers and partners, has guaranteed up to $105 billion of obligations tied to an OpenAI data center, and gains on those stakes made up between a seventh and a quarter of reported net income in the last two quarters. Third, margins are guided down from 75 per cent towards 72 to 73 per cent, partly on memory costs, while Google's TPUs, Amazon's Trainium and AMD are taking a growing share of new workloads.

The valuation depends on duration, not on the current numbers. At $234, the shares trade on about 33 times trailing adjusted earnings, about 24 times the current quarter's run-rate, and about 15 times earnings for next year if the 70 per cent growth plan is delivered. That spread is the debate in one line: on management's plan the stock is inexpensive for its quality, but if customer spending pauses in 2027 or 2028, as capital cycles historically have, earnings could fall sharply from a high base. The market is pricing something between the two, and the rest of this note sets out what would move it in either direction.

2.0 Business Overview

From chips to AI factories

Nvidia began as a designer of graphics chips for video games and is now, in practice, the main supplier of the computing infrastructure behind artificial intelligence. It does not manufacture chips itself; production is outsourced, principally to TSMC, and Nvidia's value lies in design, software and system integration. What it sells today is less a chip than a complete data center building block: graphics processors such as Blackwell and the incoming Rubin, its own Grace and Vera CPUs, NVLink and Spectrum-X networking that let thousands of chips work as a single machine, and the CUDA software platform on which most AI models are developed. That integration is the core of its competitive position, because customers buy a tested system that works at scale rather than components they must assemble themselves.

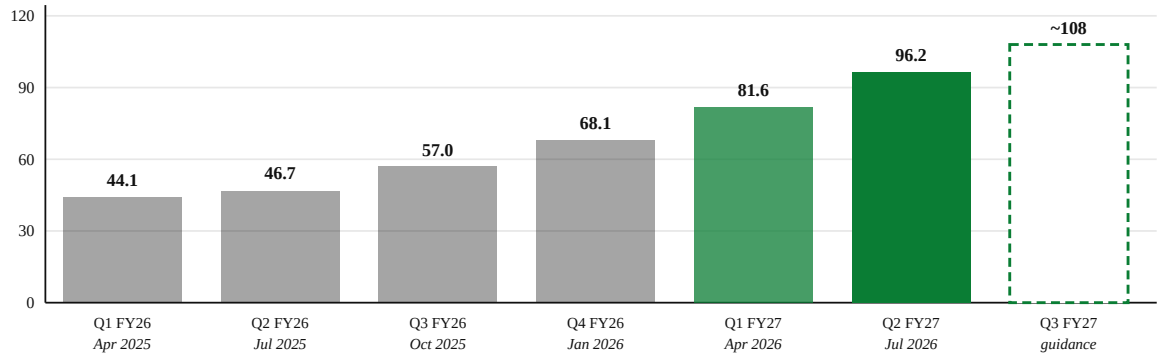
The data center segment accounts for 93 per cent of revenue. Within it, Nvidia now splits its customers into two groups. Hyperscale cloud providers, Microsoft, Amazon, Alphabet, Meta, Oracle and a few others, contributed about $49 billion in the July quarter, or 55 per cent of data center sales. The rest, about $40 billion, came from a group spanning AI labs, specialised cloud providers, enterprises and national governments building sovereign AI capacity, which grew 25 per cent in the quarter and has more than doubled in a year. That broadening of the customer base is strategically important, since it reduces dependence on four or five buyers, but as later sections show, part of it is demand that Nvidia has helped to finance. Networking has become a business of significant size in its own right, with Spectrum-X Ethernet revenue up 2.6 times on the year, while gaming, professional visualisation and automotive, once the core of the company, are now a small edge-computing segment of about $7 billion a quarter.

3.0 Financial Performance

Six quarters of acceleration

The scale of the expansion is easiest to see quarter by quarter. Revenue has more than doubled in five quarters, from $44 billion in the April 2025 quarter to $96 billion in the July 2026 quarter, and the guidance for October implies a further step to $108 billion. Over the last four reported quarters Nvidia generated about $303 billion of revenue; the current quarter's guidance annualises to more than $430 billion.

Exhibit 1. Quarterly revenue, fiscal 2026 to fiscal 2027 ($bn)



Source: Andra Markets; Nvidia quarterly results (fiscal 2026 and 2027). Q3 FY26 derived from reported full-year revenue. The dashed bar is the midpoint of company guidance. Green marks fiscal 2027.

Profitability has kept pace. Gross margin was 75 per cent in each of the last two quarters, an extraordinary level for a hardware business that reflects how much of the value sits in design, networking and software. Operating income reached $63.7 billion in the July quarter, a margin of about 66 per cent, and operating expenses of around $9 billion a quarter mean that almost every additional dollar of gross profit falls through to the bottom line. Cash conversion is equally strong: in the first half of the fiscal year operating cash flow was $74.4 billion and free cash flow $69.9 billion.

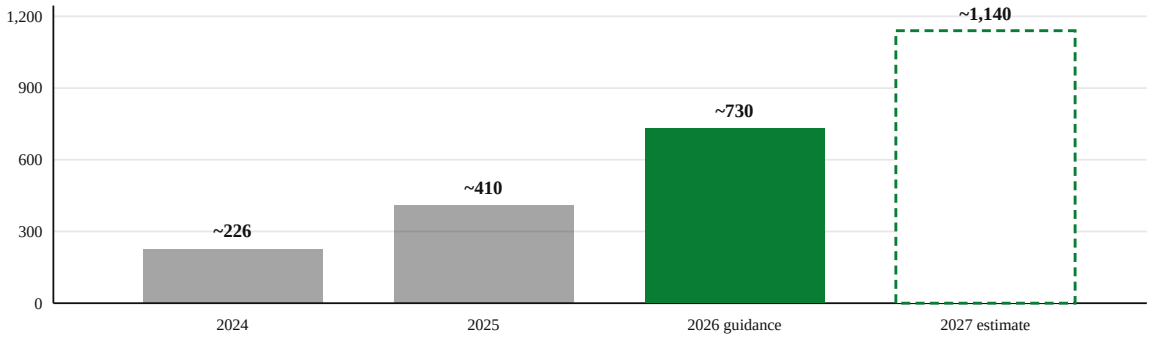
Two details in the accounts qualify that picture. Accounts receivable have roughly doubled to $63.1 billion, around 60 days of revenue, as large customers have been given longer payment terms; that is normal at this scale but worth tracking, because it means a growing share of revenue is recognised before it is collected. And margins are guided to fall: to 74 per cent in the October quarter, a trough of 71 to 72 per cent in the January quarter as memory prices rise and Rubin ramps, and 72 to 73 per cent in fiscal 2028. Those are still exceptional margins, but the direction matters for a stock priced partly on their durability.

4.0 Demand Drivers

Who is paying, and how

Nvidia’s revenue is, in effect, a share of other companies’ capital budgets, and the largest of those budgets have grown at a pace without precedent. Combined capital spending at Amazon, Microsoft, Alphabet and Meta roughly doubled from about $226 billion in 2024 to about $410 billion in 2025, and guidance for 2026 points to around $730 billion. Goldman Sachs projects about $1.1 trillion for the four in 2027. A large share of that spending goes to AI data centers, and within those, Nvidia systems are typically the single largest line item.

Exhibit 2. Combined capital expenditure of Amazon, Microsoft, Alphabet and Meta ($bn)



Source: Andra Markets; company guidance as compiled by public capex trackers; 2027 is a Goldman Sachs estimate for the four companies combined, not company guidance. Green marks the current and projected cycle.

The financing of that spending is changing in a way that matters for Nvidia. In 2025 the hyperscalers' capital expenditure absorbed around 40 per cent of their operating cash flow; in 2026 it is closer to 94 per cent. Alphabet reported its first-ever quarter of negative free cash flow in the second quarter, and the group issued around $136 billion of investment-grade bonds in the first half of the year. Goldman Sachs expects more than a third of 2027's capital spending to be debt-financed. None of these companies is financially stretched, but the shift means AI infrastructure is moving from being funded out of surplus cash to being funded by capital markets, which makes it more sensitive to interest rates and to investors' patience over returns, at a time when long-term Treasury yields are at their highest in over twenty years.

The second engine of demand is outside the large clouds. AI labs such as OpenAI, Anthropic and xAI, specialised GPU cloud providers such as CoreWeave and Nebius, and governments investing in sovereign capacity together now buy almost as much as the hyperscalers. Management says its major customers hold backlogs totalling about $2 trillion and that it could grow significantly faster than its 70 per cent planning assumption if supply allowed. The binding constraints today are physical: electrical power, completed data center shells, high-bandwidth memory and advanced packaging capacity at TSMC. China, once a significant source of data center sales, is now excluded from guidance altogether because of export restrictions, so any resumption of sales there would be upside to current forecasts.

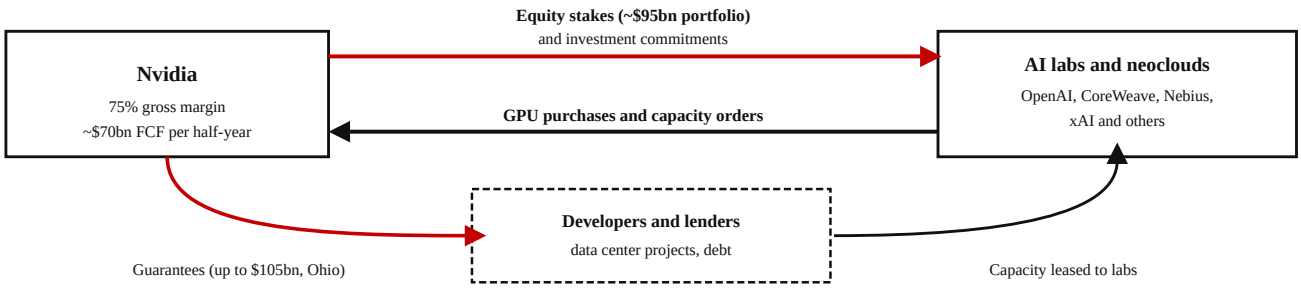
5.0 Ecosystem Financing

Supplier, investor and guarantor

The most significant change in Nvidia's profile over the past year is that it no longer only sells to the AI ecosystem; it increasingly finances it. Its portfolio of equity stakes in customers and partners has grown from almost nothing in 2020 to about $95 billion, including positions in SpaceX, CoreWeave, Nebius, Intel, Nokia, Synopsys and other companies, alongside investments in leading AI labs. In August it signed a guarantee of up to $105 billion supporting the payment obligations of a planned OpenAI data center campus in Ohio, for which Nvidia will be the exclusive initial chip supplier; the guarantee was scaled back from an initially reported $250 billion after investors raised concerns about the exposure. Management has referred to around $500 billion of ecosystem commitments, and Nvidia has launched compute financing platforms with six financial institutions to raise third-party capital for AI infrastructure.

Exhibit 3. How Nvidia's capital and its revenue now flow in a loop

The same dollars can appear as Nvidia investment, customer capital and Nvidia revenue

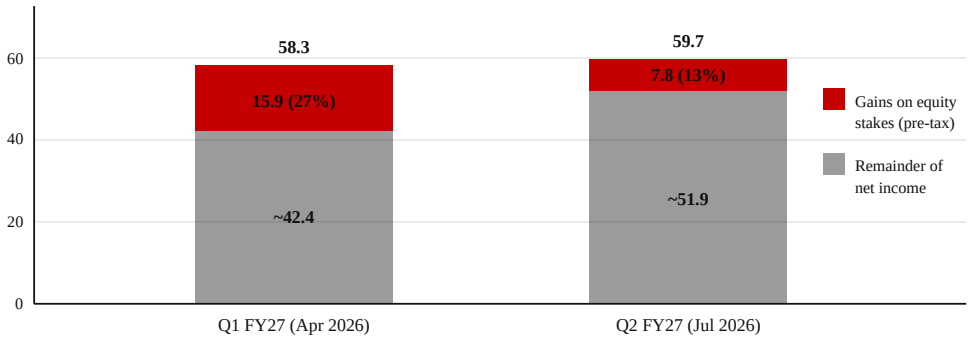


Source: Andra Markets; schematic based on company disclosures and press reports (2026). Not all flows are linked to specific transactions. Red marks flows in which Nvidia takes on financial exposure to its own customers.

There is a coherent business logic to this. Nvidia has more cash than it can usefully reinvest in its own operations, and supporting well-placed customers secures future demand and helps solve the financing bottleneck for data centers that otherwise could not be built. The risk is that it blurs the line between demand and supply. When a supplier invests in or guarantees a customer that then spends the money on the supplier's products, part of the revenue is underwritten by the supplier's own balance sheet, and the true economic demand is harder to measure. Investors reacted to the first reports of the OpenAI guarantee by selling the shares by 4.5 per cent in a day and pushing Nvidia's credit default swaps to their largest intraday move in months, and the chief financial officer addressed the "circular financing" criticism directly on the last earnings call.

The investments also affect reported earnings. Gains on Nvidia’s equity holdings were about $15.9 billion in the April quarter and $7.8 billion in the July quarter, against GAAP net income of $58.3 billion and $59.7 billion. That is why GAAP earnings per share ($2.46) were higher than the adjusted figure ($2.22) last quarter, an unusual pattern, and why the trailing GAAP price-to-earnings ratio understates the true multiple on the operating business. Nvidia has also begun to borrow, with total debt rising to about $33 billion from $8.5 billion at the start of the fiscal year, modest against its cash generation but a change of direction for a company that had been running down its debt.

Exhibit 4. GAAP net income and the contribution from gains on equity stakes ($bn)



Source: Andra Markets; Nvidia results and press reports. Gains are shown pre-tax against after-tax net income, so the split is illustrative and slightly overstates their after-tax share. Red marks earnings from investments rather than operations.

6.0 Competitive Position

The margin as the measure of the moat

Nvidia’s share of the AI accelerator market by revenue peaked at around 87 per cent in 2024 and was about 80 per cent in 2025; most estimates put it closer to 75 per cent in 2026. The competition comes from two directions. The first is custom silicon built by its largest customers: Google’s TPUs, now offered to outside customers and used to train some of the most capable models, Amazon’s Trainium, and in-house chips at Microsoft and Meta. These are cheaper per unit of computation for the specific workloads they are designed for, and every workload that moves to them is one Nvidia does not supply. The second is AMD, whose MI-series accelerators have won large commitments from OpenAI and others and offer a credible alternative for inference at lower prices.

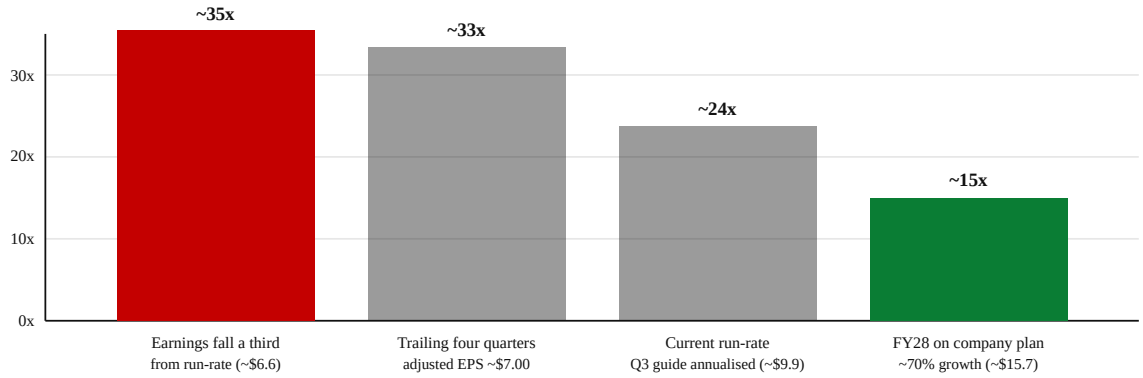
Nvidia’s defence rests on three things. Its systems are full-stack and proven at a scale no rival matches, which matters most for training the largest models. Its CUDA software ecosystem, built over nearly two decades, makes switching costly. And its product cadence, a new architecture roughly every year, keeps it ahead on performance per watt, which has become the critical metric as power, not chips, becomes the scarcest input. Falling market share in a rapidly growing market is not in itself a problem; Nvidia’s revenue can keep rising while its share slips. The more telling indicator is gross margin. A 75 per cent margin reflects pricing power, and the guidance towards 72 to 73 per cent, even if mostly explained by memory costs, is the first sign that this power may have peaked. If competition, rather than input costs, begins to drive margins lower, it would change the valuation case materially.

7.0 Valuation

What the price assumes

Because Nvidia’s earnings are growing so fast, the multiple depends heavily on which earnings are used. On trailing adjusted earnings of about $7.00 per share, the stock trades at around 33 times. On the current quarter’s run-rate, annualising the earnings implied by October guidance at about $2.47 per share, it is about 24 times. If management’s plan for roughly 70 per cent revenue growth in fiscal 2028 were delivered at the guided margin of 72 to 73 per cent, earnings would approach $15 to $16 per share and the multiple would fall to around 15 times, cheaper than the semiconductor sector median and below most of the other largest US companies. On the same run-rate basis, enterprise value is roughly 13 times annualised sales.

Exhibit 5. Price-to-earnings at $234 under different earnings assumptions (x)



Source: Andra Markets calculations on company results and guidance. Run-rate and FY28 earnings are illustrative estimates (FY28 assumes a Q4 FY27 of about $118bn, 72.5% gross margin and a 16% tax rate), not forecasts. Red marks the downside case, green the plan case.

The exercise makes clear what the market is actually pricing. It is not paying 33 times for a mature company; it is paying a moderate multiple for earnings that are expected to keep rising sharply, with a discount for the risk that they do not. The relevant historical comparison is not a software company but a capital-goods supplier at the top of an investment cycle, whose earnings can fall far faster than revenue when customers pause, as Nvidia's own did in 2022 when a smaller crypto and gaming downturn cut its profits by more than half. A digestion phase in 2027 or 2028, in which hyperscalers slow spending to absorb capacity already built, would leave the shares on a demanding multiple of lower earnings. Continued growth on the scale of the plan would leave them looking cheap. The current price sits closer to the second outcome than the first, but not so close that it requires everything to go right.

8.0 Risk Considerations

Where the case could break

The central risk is cyclical. Nvidia's revenue is a function of its customers' capital budgets, and those budgets are now being stretched to the limit of operating cash flow and beyond. If the return on AI investment disappoints, or if higher interest rates make debt-funded data centers harder to justify, spending could plateau or fall, and with operating leverage this high the effect on earnings would be magnified. Customer concentration amplifies this: a decision by two or three of the largest buyers to slow orders would be visible in Nvidia's results within a quarter or two.

The second risk is the financing loop. The more Nvidia invests in and guarantees its customers, the more its revenue depends on the financial health of companies that are themselves loss-making or heavily indebted, and the more its own earnings depend on the market value of those stakes. A setback at a major AI lab or GPU cloud provider could hit Nvidia three times, through lost orders, investment write-downs and potential calls on guarantees. The third is competition and margins, discussed above, with custom silicon the most direct threat. Beyond these sit supply-chain concentration, given Nvidia's reliance on TSMC in Taiwan for its most advanced chips; regulatory risk, from export controls that have already removed China from its forecasts and from antitrust scrutiny of its dominance; and the physical constraints on power and memory that, while currently limiting supply rather than demand, also cap how fast the business can grow.

9.0 Outlook and Monitoring Points

What to watch

The next test is the October quarter, due in late November, where the market will want to see revenue at or above $108 billion and confirmation of the Rubin ramp. Beyond the headline numbers, five indicators will say most about the durability of the cycle: hyperscaler capital spending guidance for 2027, which will start to emerge with their October results; the gross margin trajectory, and in particular whether the decline stops at the guided 71 to 72 per cent; the size and composition of gains on Nvidia's equity portfolio relative to operating earnings; the growth of accounts receivable relative to revenue; and the progress of custom silicon, especially any shift of large training workloads to TPUs or Trainium. On the macro side, long-term interest rates matter more than usual, because they now set the cost of financing

much of the demand Nvidia depends on. The company has never been stronger operationally; the questions that will decide the next leg for the shares are about its customers' balance sheets as much as its own.

Methodology and Source Notes

Method. This note is a fundamental review of Nvidia based on company results for fiscal 2026 and the first half of fiscal 2027, the second-quarter earnings call, market data as of the 2 October 2026 close and press reporting. Run-rate and fiscal 2028 earnings, multiples and ratios are Andra Markets calculations using stated assumptions and are illustrative, not forecasts. Industry capital spending figures combine company guidance with third-party estimates, and market share figures are third-party estimates that vary by source. This note is analysis, not investment advice.

Principal sources. Nvidia second-quarter fiscal 2027 and fourth-quarter fiscal 2026 results releases and earnings call coverage (Yahoo Finance, StockTitan); Yahoo Finance, on gains in the equity portfolio; Axios, Fortune and the Wall Street Journal via Yahoo Finance, on the OpenAI financing guarantee (July to August 2026); public hyperscaler capex trackers and Goldman Sachs estimates (2026); Silicon Analysts, on accelerator market share; CompaniesMarketCap and GuruFocus, for market data.

Disclaimer

This document is produced by **Andra Markets** Global Markets Research for informational and educational purposes only. It is not investment, legal, accounting or tax advice, and it is not an offer or solicitation to buy or sell any security or financial instrument. Figures are drawn from public sources believed to be reliable but are not independently verified. Earnings scenarios and multiples are illustrative calculations, not price targets or forecasts. Views are those of the desk as of the date of publication and are subject to change without notice. Past performance is not indicative of future results. Recipients should conduct their own analysis and consult their own advisers before acting. Andra Markets accepts no liability for any loss arising from the use of this material. Redistribution outside the intended recipient list is not permitted without consent.