October 2, 2026 05:32 AM GMT

NVIDIA Corp. | North America

Reinstating Top Pick as NDR supports our positive views

WHAT'S CHANGED		
NVIDIA Corp. (NVDA.O)	From	To
Top Pick Added	-	**NVDA.O**

On our road show this week, NVDA's CEO highlighted that recent developments in land/power/shell, financing, component constraints, and agentic all play to their strengths; we agree and are reinstating the stock as our Top Pick in semis.

Key Takeaways

- Road show highlights the diversity of the company's business beyond hyperscalers and frontier labs
- The focus on more expensive processor per GW seems likely to pay off given land/power/shell constraints
- We see the rise of agentic as positive for CPU - including NVIDIA CPUs - but as a much bigger driver for GPU
- We certainly see NVIDIA as well positioned, with a very undemanding valuation - we are reinstating NVIDIA as our Top Pick

We recently hosted Nvidia CEO Jensen Huang, CFO Colette Kress, and Head of IR Toshiya Hari for investor meetings in NYC and Boston.

What we heard from management was largely reinforcement of recent priorities around shifts in the AI investment climate. Our biggest takeaways include:

1) Diversity of the customer base: The rationale for the resegmentation of the business six months ago was to allow investors to understand the breadth of the business. While 90% of investor conversations are around the biggest hyperscalers and frontier model companies, half of the business - probably the faster growing half - is a broader range of AI model companies, Neoclouds, sovereign entities, ODMs and OEMs, and enterprise customers. That ACIE segment needs the full stack AI factory ecosystem that NVIDIA can provide.

Joseph Moore
Equity Analyst
Joseph.Moore@morganstanley.com +1 212 761-7516

Mason Wayne
Research Associate
Mason.Wayne@morganstanley.com +1 212 761-6012

Ella Tulchinsky
Research Associate
Ella.Tulchinsky@morganstanley.com +1 212 761-2222

Nicole Kozhukhov
Research Associate
Nicole.Kozhukhov@morganstanley.com +1 212 761-1636

Shane Brett
Equity Analyst
Shane.Brett@morganstanley.com +1 212 761-1022

Cate Folan
Research Associate
Cate.M.Folan@morganstanley.com +1 212 296-3520

NVIDIA Corp. (NVDA.O, NVDA US)
Top Pick

Semiconductors | United States of America

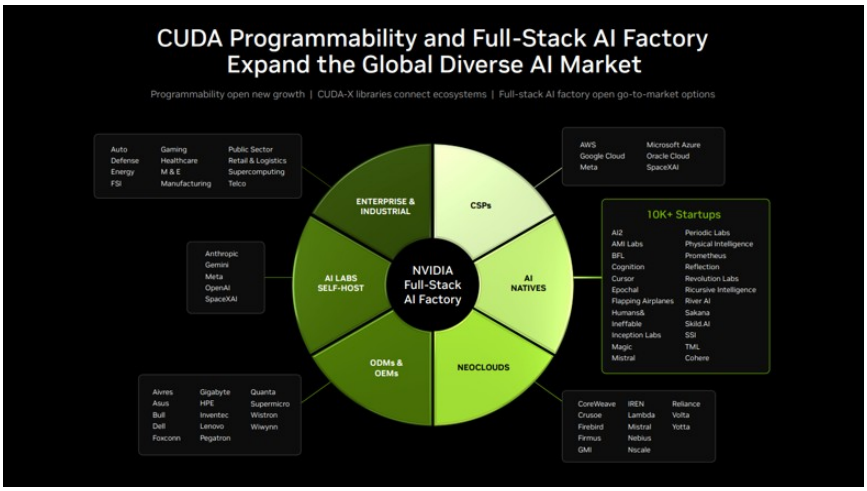
Stock Rating	**Overweight**
Industry View	**Attractive**
Price target	**$300.00**
Shr price, close (Oct 1, 2026)	$230.86
Mkt cap, curr (mm)	$5,628,453
52-Week Range	$236.54-164.27

investors should be aware that the firm may have a conflict of

Research as only a single factor in making their investment decision.

For analyst certification and other important disclosures, refer to the Disclosure Section, located at the end of this report.

Exhibit 1: Nvidia's diverse end markets and customers



Source: Nvidia [https://s201.q4cdn.com/141608511/files/doc_presentations/2026/09/2026-NDR_.pdf]

The importance of open models is also an important element of the narrative, given NVIDA's strong position across both open weight and closed parameter models.

The breadth of that customer base also helps with land/power/shell, as those constraints are having the most impact on US based hyperscalers, with sovereign business creation as a result of countries that want to use their land/power/shell to create businesses. It also affords NVIDIA significant opportunities to drive revenue sharing agreements and valuable software/services revenue streams. MSR: NVIDIA, neoclouds, and rev sharing: the next big debate (14 Aug 2026).

2) Growth within traditional hyperscale & frontier labs. While some of the sentiment is that the resegmentation was to point away from traditional hyperscale, the company showed that they have gained share of spending from the top 5 cloud service providers steadily and consistently, both because AI is a larger portion of the spending but also because the company in a very short term has become the revenue market share leader in markets such as Ethernet, and will rival the leaders in markets such as microprocessors. Custom silicon plays an important role in those hyperscalers, optimizing for specialized workloads, but general purpose processors from NVIDIA play a primary role and offer the longest asset life.

3) The strategy is to drive up revenue per GW substantially, to ensure that NVIDIA solutions continue to drive the most tokens per GigaWatt. NVIDIA showed slides indicating that Feynman (the 2028 architecture) will drive a significant increase in NVIDIA revenue per GW, from $40 bn to something well north of $50 bn. This is going to prove divisive, as competitors highlight their lower revenue per GW - and slower rate of increase - as a significant positive for costs.

But if the biggest constraint becomes land/power/shell, the highest tokens per Gigawatt win, all else being equal. We are seeing this in real time, where the biggest users of ASICs, which have lower cost per GW, are the companies with the biggest data center land/power/shell issues. Vera Rubin is driving up to 25x token improvement vs. Blackwell, which certainly mitigates concerns that dollars per gigawatt are rising.



been highlighting: there isn't enough power to do what everyone wants to do, and there are more semiconductors being built than there are places to put them. He was very clear that this is not true for NVIDIA.

We certainly see the constraints, but AI is so strong that we don't see anyone missing numbers any time soon. But we hear from our industry contacts that NVIDIA tokens per gigawatt are much higher than peers, by multiples, so if gigawatts is the constraint, a strategy of "gigawatt densification" - working within a smaller power envelope - is likely to be the dominant strategy.

4) Persistent agents will indeed drive CPU demand higher - but the biggest portion of spending will still be on GPUs. This seems intuitive to us - a single core of a CPU in an agentic workload such as Muse will spend a few minutes a day at most doing a few process steps for the agentic harness, which would mean that a single CPU with hundreds of cores can handle a massive number of users, but the incremental inference tokens generated is also a huge number; NVIDIA postulated that an agentic user may end up driving 10x the daily tokens that a human user might have.

That said, the CPU shortage is intense - though we do not believe that it has intensified over the last month - and NVIDIA remains confident in its ramp of CPUs from $20 bn in 2026 to $40 bn in 2027, of which we think is a little more than 50% driven by the head node CPU in the AI racks - meaning the balance is standalone CPU racks.

5) The company will continue to use the balance sheet to drive advantages, while living up to commitments to return 50% of cash. NVIDIA has mobilized a half trillion of non NVIDIA capital to invest into AI data centers, and has prepurchased a large amount of semiconductor components and power, as well as data center leases that it intends to transfer to customers. While the leverage inherent in AI remains controversial, NVIDIA has completely avoided stock giveaways or other structures that are actually inherently much more circular than financial support. The company is working towards a longer term goal of establishing a criteria for AI factories as a collateral asset, which we aren't sure about. But we do expect NVIDIA's aggressive support of the ecosystem to be an advantage.

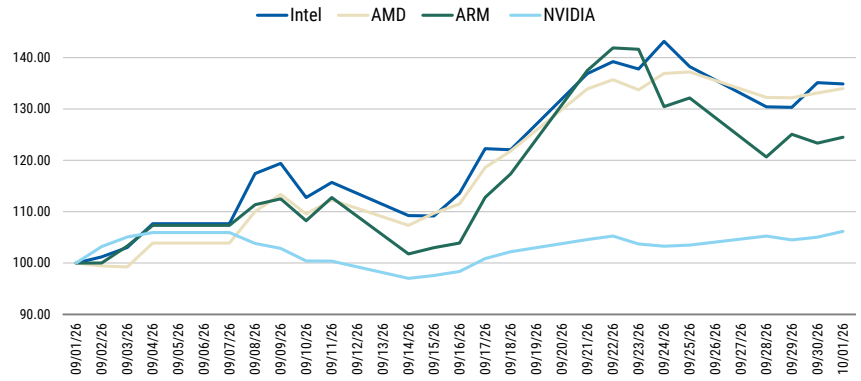
Separately from the NDR, we are reinstating NVIDIA as our Top Pick. We named the stock our Top Pick in March (which was at least three months too early) after an earlier preference for stocks with more leverage in the memory space, but Morgan Stanley process is that the designation is removed after 6 months; we are reinstating it now.

Our Top Pick thesis:

Nvidia remains very early in the ramp of a new product cycle, powered shell constraints play into Nvidia’s strengths, and recent agenetic enthusiasm can accelerate Vera adoption. With the stock still relatively inexpensive at 15x FY28 EPS, there's opportunity for multiple expansion if enthusiasm returns to AI but multiple expansion is also not required for the stock to work as numbers continue to go

higher. Nvidia has also underperformed more CPU stories around the recent Muse enthusiasm (Meta Platforms Inc: 23 Days of Musing...Top Investor Questions and Our Answers (30 Sep 2026)), but we think that outperformance will broaden as consumer agents as a category add to long term AI use cases and token demand.

Exhibit 2: Price performance since 9/1, room for GPUs to catch-up?



The last 12 months have been about semiconductor production limits. Now the bottleneck is quickly shifting to the speed that new datacenter space can be built and how it will be financed, an advantage for the Nvidia ecosystem. With increasing political backlash, the margin for error to build in the US is narrowing, but Nvidia is leading the way when it comes to international deployments, bringing up sovereign spenders and additional neoclouds clouds around the world. Nvidia lists 80 cloud partners on their website, 55 of which are outside the US (although they may have datacenter in multiple regions). This gives them the highest degree of flexibility to go where the power is - which is increasingly outside the US and not available to traditional hyperscalers.

Capital is another important constraint for AI labs, and decisions about compute are increasingly dependent on financing - for which Nvidia has taken large measures to either line-up or provide directly [NVDA $500bn investment mobilization is a positive; brings focus to rev sharing (11 Aug 2026)]. Nvidia's willingness to provide financing and their resegmentation of the business has been taken negatively by investors, but that sentiment can shift in Nvida's favor. That makes NVIDIA a more resilient AI beneficiary within semis today; if there's an expectations reset on the how quickly AI capacity can be deployed, areas of the semis supply chain where earnings are more dependent on a shortage should see the most negative impacts if lead times come in.

Two reasons why we see upside to Nvidia's revenue FY28 guide: On Nvidia's last earnings call management said they expect revenue next fiscal year to grow 70% y/y (implying ~75% for datacenter); that's a supply constrained forecast with demand closer to 100%. One of the primary gating factors on the supply side seems to be the timing of datacenter land power shell availability and while we can't point to specific sites that enable a more robust forecast, there are other reasons to be optimistic.

We would highlight two factors:

(1) Nvidia customers SpaceX and Amazon have both doubled down on their commitment to Nvidia in the last 90 days, and this creates a low bar for the remaining customers. Conservative forecasts for those customers' contribution to Nvidia's 75% datacenter revenue guide implies about 20% growth for the remaining customers in Nvidia's hyperscale segment; significantly below the 39% y/y growth in their capex per consensus estimates.

For SpaceX management's guidance is to exit 2026 at ~2GW of datacenter capacity, growing to 5-10GW by EOY 2027: assuming the low end (in-line with Adam Jonas' SPCX estimates) could imply up to $120bn of Nvidia spend next year at $40bn of GW. In our analysis we assume 79% of $132bn of SPCX AI capex in 2027, 79% being in-line with 1Q26. Each GW SPCX is able to deploy above 3 would add ~5% to Nvidia revenue in FY28.

Amazon and Nvidia recently announced an expansion of their partnership (here), which includes adding 2mn additional GPUs between 2027-2028, Vera CPUs, and NVlink Fusion. With Tranium 4 not expected to ship until late 2027 and the reinforcement of the relationship between AWS/Nvidia, we think it's reasonable to think Nvidia gains some share of AWS capex in 2027. Our internet team currently estimates ~$40bn of amazon spend at Nvidia in 2026 and $78bn in 2027.

Exhibit 3: 3/4th of the growth we expect for Nvidia's hyperscale segment in FY28 can be explained by just 2 customers, creating a low bar for the remainder

Hyperscaler Capex → Nvidia Share	1Q26	2026	2027	Notes
SPCX Capex	$6,528	$53,726	$132,132	*MS Estimate*
Nvidia %	79.4%	79.4%	79.4%	
Nvidia $ Spend	$5,181	$42,640	$104,868	
Impled SPCX GWs	0.21	1.52	2.62	*SPCX management expects to deploy 3-8GW in 2027*
AWS Capex		$180,000	$270,000	*MS Estimate*
Nvidia %		22.2%	28.9%	
Nvidia $ Spend		$40,000	$78,000	*MS Estimate*
Nvidia Datacenter	$75,246	$361,636	$630,917	*MS Estimate (annuals calendarized)*
y/y growth			74.5%	*Nvidia guidance is 70% total revenue growth in FY28*
Nvidia Hyperscaler Segment	$43,050	$198,900	$337,541	
% of datacenter	57.2%	55.0%	53.5%	*Nvidia expects ACIE segment to outgrow hyperscale*
SPCX %	12.0%	21.4%	31.1%	
AWS %		20.1%	23.1%	
SPCX + AWS %		41.5%	54.2%	
SPCX + AWS % of growth			72%	
Implied Other Hyperscaler Spend	$36,522	$105,174	$127,409	
y/y growth			21.1%	*SPCX, AMZN numbers above imply Nvidia significantly*
Cons Capex for (MSFT, ORCL, META, GOOG)		$569,597	$789,852	*undergrows remaining hyperscaler capex ests...*
y/y growth		94%	39%	*...and that they lose share as a % of capex*
Implied Nvidia %		18%	16%	*...vs Nvidia's expecation that they will gain share of the top 5 CSP capex in 2027*

(2) Separating out the contribution from stand alone Vera CPUs, and Groq LPUs implies mid-single digit to high-single digit growth for core GPU sales before including additional content gain in areas like spectrum, NVlink fusion, DPUs/ storage etc. Current CoWoS expectations for Nvidia's Blackwell+Rubin from our greater China semis team is closer to flat, but we feel comfortable that with Nvidia's scale relationship with TSMC, they will see a healthy allocation and given lead times for backend capacity, upside next year is still possible if 3nm wafers free up.

We also continue to be excited about the potential contribution from Groq and the growing fast inference category. It's early to tell what the attach rates will looks like but we assume 30% of deployments will be "Groq enabled" with a ratio of 1 Groq rack for every 4 Rubin racks, per Jensen's commentary at GTC. We also assume $35k ASP LPX tray - we don't know with precision what the price will be but each tray has 8 LPUs, 1 CPU, a NIC/DPU as well as additional DRAM. So we don't view an

average of 3.5k per logic chip as particularly aggressive. We admit that the analysis is somewhat circular, as assuming a certain number of racks requires assuming a certain number of chips. To avoid that issue, to the extent possible, we assume VR racks next year are fairly close to what our ODM analyst Howard Kao is assuming for 2026.

Exhibit 4: Back of the envelope groq + stand alone CPU revenue in FY28 creates a low bar for units

Step 1		Notes
Groq FY28 Build		
ASP per LPU Tray ($)	$35,000	MSe
Trays per LPX rack	32	
VR racks per LPX rack	4	
Attach rate (% of Rubin racks w/ LPX)	30.0%	Nvidia Commentary
Rubin NVL72 racks	85,000	MSe
Rubin racks with LPX attached	25,500	
LPX racks required (1 per 4VR)	6,375	Nvidia Commentary
Groq LPU revenue		
LPU revenue ($M)	**$7,140**	
% of total DC revenue	1.1%	
Step 2		
Vera FY28 Build		
Vera CPUs per NVL72 rack	36	
Vera CPU ASP ($)	$5,500	
Implied Vera CPU units	3,060,000	
Vera CPU revenue ($M)	**$16,830**	MSe
Implied stanalone Vera in guide		
FY28 CPU guide ($M)	$40,000	
Vera CPU revenue from racks ($M)	$16,830	
Implied Standalone Vera ($M)	**$23,170**	
% of total DC revenue	3.5%	

Step 3		Notes
Implied growth in 'other' DC revenue		
FY27 DC revenue ($M)	$378,608	
FY28 DC revenue @ 75% growth ($M)	$662,563	
Target DC revenue growth	75.0%	Nvidia Commentary
Less: Groq + Vera ($M)	$23,970	
Implied 'other' FY28 revenue ($M)	**$638,593**	
Implied 'other' y/y growth	**68.7%**	
Step 4		
GPU price per chip: Blackwell vs Rubin		
Blackwell ASP ($)	$35,000	MSe
Rubin ASP ($)	$55,000	MSe
Growth in price per chip	**57.1%**	
Implied unit growth to reach guide		
Implied 'other' revenue growth	68.7%	
Growth in price per chip	57.1%	
Implied unit growth	**7.3%**	

Two reasons why we see 72-73% gross margin as a solid floor for FY28: Nvidia's gross margins in-light of rising input costs (primarily DRAM) have been a key overhang for the stock, and after guiding them lower for FY28 the next question on investors' minds is whether there's further downside ahead. We expect margins to hold up in-line with guidance due to two levers Nvidia is pulling to offset rising costs: pricing and de-specing. Looking at a simple framework for GM based just on the Rubin transition, NVIDIA can potentially absorb a 30% percent increase in COGS with a 15% price increase and 44% with an HBM de-spec assuming 75% margin as a baseline. If that entire increase in COGS shows up as HBM, the price increase could be 100-200%. We don't know exactly what the price of HBM was in 2026 but that high end of $43 for future Rubin for HBM pricing in the de-spec case is well above the current $30 per GB spot DDR5, which seems conservative to us.

Exhibit 5: 15% Rubin price increase, and de-spec impact on Nvidia's gross margin

GPU Margin analysis	R200 288GB	R200 192GB (de-spec)
GPU ASP ($)	$55,000	$55,000
Nvidia GPU gross margin (288GB input)	75.0%	77.6%
Implied GPU COGS ($) — from margin	$13,750	$12,310
HBM capacity per GPU (GB)	288	192
HBM ASP ($/GB)	$15.00	$15.00
HBM cost per GPU ($)	$4,320	$2,880
Non-HBM COGS ($)	$9,430	$9,430
Total GPU COGS ($)	**$13,750**	**$12,310**
HBM as % of GPU COGS	**31.4%**	**23.4%**
Path to target GM:		
Price increase	15.0%	15.0%
New GPU ASP ($)	$63,250	$63,250
Target Nvidia GPU gross margin	72.0%	72.0%
Allowed GPU COGS at target GM ($)	$17,710	$17,710
Less: non-HBM COGS ($)	$9,430	$9,430
Allowed HBM cost per GPU ($)	**$8,280**	**$8,280**
Required HBM ASP ($/GB)	$28.75	$43.13
% increase in HBM ASP required	**91.7%**	**187.5%**

Exhibit 6: ...and rack level impact

	GB200 NVL72	GB300 NVL72	VR200 NVL72	VR200 (Less Mem)
1. Nvidia price increase (100% captured)				
Total rack price (context)	$3,329,523	$3,994,551	$7,803,148	$7,379,788
Nvidia-sold content	$3,061,440	$3,646,065	$6,958,305	$6,534,945
Price increase on Nvidia content	15.0%	15.0%	15.0%	15.0%
New Nvidia revenue per rack	$3,520,656	$4,192,975	$8,002,051	$7,515,187
Incremental revenue = incremental GP	**$459,216**	**$546,910**	**$1,043,746**	**$980,242**
2. Nvidia margin impact				
Nvidia Gross Margin	73.0%	73.0%	73.0%	73.0%
Nvidia GP/Rack	$2,234,851	$2,661,627	$5,079,563	$4,770,510
Nvidia COGS/Rack	$826,589	$984,438	$1,878,742	$1,764,435
New Nvidia GP per rack	$2,694,067	$3,208,537	$6,123,308	$5,750,752
New GM	**76.5%**	**76.5%**	**76.5%**	**76.5%**
3. COGS headroom at constant margin				
New Nvidia COGS/rack (at constant GM)	$950,577	$1,132,103	$2,160,554	$2,029,100
Increase in Nvidia COGS/rack	$123,988	$147,666	$281,811	$264,665
% increase in Nvidia COGS	**15.0%**	**15.0%**	**15.0%**	**15.0%**
4. Implied blended memory price increase				
Main memory content (GB/rack)	17,280	17,280	55,296	27,648
HBM content (GB/rack)	13,824	20,736	20,736	20,736
Total memory content (GB/rack)	**31,104**	**38,016**	**76,032**	**48,384**
Main memory ASP ($/GB)	$3.50	$5.00	$8.50	$8.50
HBM ASP ($/GB)	$10.80	$10.80	$13.50	$13.50
Main memory cost/rack	$60,480	$86,400	$470,016	$235,008
HBM cost/rack	$149,299	$223,949	$279,936	$279,936
Total memory cost/rack	**$209,779**	**$310,349**	**$749,952**	**$514,944**
Blended memory ASP ($/GB)	**$6.74**	**$8.16**	**$9.86**	**$10.64**
New blended ASP to absorb COGS increase	**$10.73**	**$12.05**	**$13.57**	**$16.11**
% increase in blended memory ASP	**59.1%**	**47.6%**	**37.6%**	**51.4%**
Total memory as % of rack price	6.3%	7.8%	9.6%	7.0%

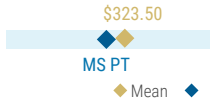
Risk Reward – NVIDIA Corp. (NVDA.O) Top Pick

OW as large language model enthusiasm is transforming cloud capex

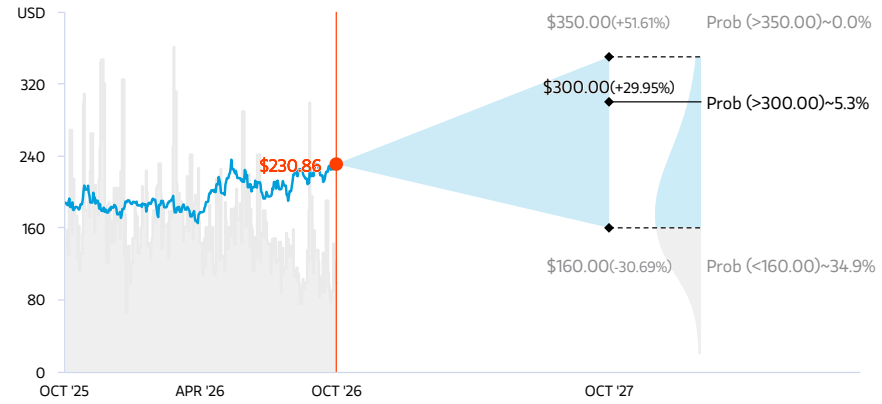
PRICE TARGET $300.00

~20 our MW CY27 EPS estimate of $15.01, in-line with the broader market and a discount to compute semis peers (AMD/AVGO/INTC) as high marketshare and gross margins leave limited levers for multiple expansion in the near term.

Consensus Price Target Distribution



RISK REWARD CHART AND OPTIONS IMPLIED PROBABILITIES (12M)



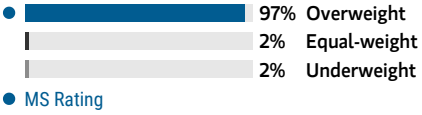
Key: Historical Stock Performance Current Stock Price Price Target

Base, and Bear case scenarios playing out were estimated with implied volatility data from the options market as of 1 Oct 2026. All figures are approximate risk-neutral probabilities of the stock reaching beyond the scenario price in either three-months' or one-years' time. View explanation of Options Probabilities methodology here

OVERWEIGHT THESIS

- Blackwell remains the premiere solution for gen-AI workloads, where compute demand continues to outstrip supply
- We see continued upward pressure to estimates as demand strength continues, with Rubin expected to maintain Nvidia's performance leadership position

Consensus Rating Distribution



Risk Reward Themes

New Data Era:	*Positive*
Pricing Power:	*Positive*
Secular Growth:	*Positive*

View descriptions of Risk Rewards Themes here

BULL CASE $350.00

~22x bull case MW CY27 EPS of $16

Bull case has DC revenues surprising to the upside as supply becomes more available and other products gain steam (CPUs, Groq)

- Higher margin data center and AI-focused software and services growth accelerates

- Revenue sharing agreements add billions in 100% margin revenue

- GPU based AI PC gains traction, widely increasing the client TAM

BASE CASE $300.00

~20x our MW CY27 EPS of $15.01

~20x PE in-line with the broader market and a discount to compute semis peers (AMD/AVGO/INTC) as high marketshare and gross margins leave limited levers for multiple expansion in the near term..

- Revenue grows by 88% in 2026 and 70% in 2027

BEAR CASE $160.00

~13x bear case MW CY27 EPS of $13

Two key debates both go the wrong direction, causing investors to question future prospects for growth

- Growth in DC slows substantially as supply catches up to demand faster than anticipated

- AI development costs come down materially, a strong competitor enters the market to take market share, or customers begin insourcing custom hardware solutions

- Greater than expected impact from tariff headwinds and export controls

Risk Reward – NVIDIA Corp. (NVDA.O)

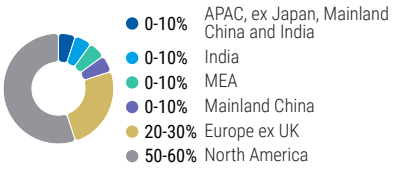
KEY EARNINGS INPUTS

Drivers	Dec 2026	Dec 2027e	Dec 2028e	Dec 2029e
GAAP Revenue ($, mm)	215,938	405,968	690,168	970,222
MW Gross Margin (%)	71.2	73.7	72.5	72.7
MW EPS ($)	4.60	9.18	15.01	22.00
Inventory ($, mm)	21,403	44,841	65,460	79,121
DOI	123.3	151.1	123.9	68.3

INVESTMENT DRIVERS

- Growth in AI capex from customers
- Next gen GPUs continue to outpace the competition
- Systems approach allows for higher monetization over time
- New drivers emerge for Nvidia such as AI PCs, autonomous vehicles, robotics, and software

GLOBAL REVENUE EXPOSURE



View explanation of regional hierarchies here

MS ALPHA MODELS

4/5 BEST	24 Month Horizon	4/5 MOST	3 Month Horizon

the highest favored Quintile and 5 is the least favored Quintile

RISKS TO PT/RATING

RISKS TO UPSIDE

- Growth in inference and training propel data center revenue
- Gaming sales accelerate as GPU based AI PCs gain traction
- Nvidia can recapture lost revenue in China

RISKS TO DOWNSIDE

- AI end markets don't materialize as expected, customers sharply reduce GPU purchases
- AMD reemerges as a viable GPU competitor
- Cloud customers outside of Google are able to develop competitive custom hardware

OWNERSHIP POSITIONING

Inst. Owners, % Active	51.2%
HF Sector Long/Short Ratio	2.2x
HF Sector Net Exposure	30.6%

Refinitiv; MSPB Content. Includes certain hedge fund exposures held with MSPB. Information may be inconsistent with or may not reflect broader market trends. Long/Short Ratio = Long Exposure / Short exposure. Sector % of Total Net Exposure = (For a particular sector: Long Exposure - Short Exposure) / (Across all sectors: Long Exposure – Short Exposure).

MS ESTIMATES VS. CONSENSUS

FY Dec 2027e

Sales / Revenue ($, mm)	405,968	695,569	802,617
EBITDA ($, mm)	269,866	460,890	542,215
Net income ($, mm)	223,210	381,211	448,289
EPS ($)	9.21	15.77	18.71

◆ Mean ◆

Risk Reward Reference links

1. View explanation of Options Probabilities methodology - Options_Probabilities_Exhibit_Link.pdf

2. View descriptions of Risk Rewards Themes - RR_Themes_Exhibit_Link.pdf

3. View explanation of regional hierarchies - GEG_Exhibit_Link.pdf

4. View explanation of Theme/Exposure methodology - ESG_Sustainable_Solutions_External_Link.pdf

5. View explanation of HERS methodology - ESG_HERS_External_Link.pdf