

CoreWeave (CRWV US)

Riding the AI Infrastructure Supercycle

Maintained **Buy**
Target price \$162

Initiate at Buy, with TP of \$162: Backed by first-mover GPU deployment, exclusive AI-only focus, best-in-class efficiency, and LTAs with major AI hyperscalers, CoreWeave is positioned as a long-term winner in the AI infra boom. We are positive on **1)** accelerating demand for AI buildouts; **2)** RPO that largely underpin our 2026–2027 revenue forecasts; **3)** declining debt financing costs; **4)** profitability to inflect sharply as large contracts mature; **5)** new LTAs benefiting from higher GPU rental pricing. Historically, CoreWeave's share price has been driven by large contract wins; we expect the next key catalyst to be the announcement of new signed agreements or further disclosure regarding the size of its contract with Anthropic. We forecast revenue of \$12.3/25.6/39.6bn in 2026/2027/2028, with adjusted EBITDA reaching to \$7.6/14.2/18.8bn. Our target price of \$162 is based on 10x 2028E EV/EBITDA.

Large NeoCloud opportunity supported by AI datacenter tailwinds: 1) Driven by exponential growth in token consumption, global AI datacenter power capacity is expected to reach 283GW by 2030. Amid tightening cash flow and capital constraints, we estimate the Top 5 CSPs will directly build ~53% of incremental capacity. After incorporating demand from enterprises and sovereign AI projects, we estimate the remaining addressable market represents a **~\$1.7tn TAM for NeoCloud providers from 2026–2030**. 2) OpenAI and Anthropic have disclosed data center build plans exceeding 20/10GW, and CoreWeave remains the only NeoCloud provider with signed LTAs with both. Existing agreements with OpenAI, Anthropic, Meta, and Microsoft are expected to contribute 60%/56%/47% of 2026/2027/2028 revenue, providing exceptional earnings visibility.

Profitability trajectory better than expected: We project CoreWeave to achieve profitability in 2028, supported by: 1) most large-scale contracts remain in the early construction or initial billing phase, with revenue recognition expected to accelerate as deployments transition into mature monthly billing and post-contract monetization phases. 2) New contracts signed in 2026 account for ~40% of its current RPO. Benefiting from rising GPU rental prices, these new contracts are expected to drive a 20% increase in blended ASP once they begin generating monthly payments. 3) Financing costs continue to decline, with the latest DDTL 4.0 priced at SOFR +2.25%, vs. DDTL 1.0/2.0/3.0 (at SOFR + 9.62% / +4.25% / +4.00%).

Risks: 1) AI demand deceleration; 2) Geo-political uncertainties; 3) Competition.

Profit forecast

(mn USD)	FY2024	FY2025	FY2026E	FY2027E	FY2028E
Revenue	1,915	5,131	12,266	25,582	39,611
Revenue YoY (%)	736.6%	167.9%	139.1%	108.6%	54.8%
Net profit	(65)	(606)	(2,070)	(1,037)	198
Net profit YoY (%)	-	-	-	-	-
EPS (\$)	(0.29)	(1.51)	(3.93)	(1.97)	0.37
P/E	-348.1	-66.1	-25.5	-50.8	266.7
ROE (%)	-1.3%	-41.6%	-10.2%	-1.6%	0.1%

Source: Company data, GF Securities (Hong Kong) Brokerage.

Michelle Jing

SFC CE No. BUK594
michellejing@gfgroup.com.hk

Jeff Pu, CFA

SFC CE No. BNO719
jeffpu@gfgroup.com.hk

Yang Zhou

SFC CE No. BSF949
zhouyang@gfgroup.com.hk

Table of Contents

NeoCloud: a New Paradigm for AI Computing Infrastructure	3
Strong Capex from upstream: CSPs + AI Giants	3
Token Usage Implies a 283GW Global AI Datacenter Power Capacity by 2030	8
Relationship with Hyperscalers: Coexistence of Cooperation and Competition	10
Compute Cost Surge: Memory + Power Cost Inflation	11
Long-Term Take-or-Pay Contracts: Anchor Tenant Strategy	13
The Sustainability of Asset-Backed Debt Model	14
CoreWeave: NeoCloud Leader	15
CoreWeave's Ongoing "Efficiency Gap" Compared to Alternatives	16
Software Stack Along with Bare-Metal to Provide Better Performance	17
Recent New Contracts Diminishing the Risk of Customer Dependence	19
Strong RPO Backed up by LTAs	20
Debt Might Be the Only Concern	21
Rising GPU Rental Price Will Benefit CoreWeave's New Contracts	22
Pace to Profit-making	23
Nvidia Both an Investor and a Supplier	27
Company Overview	28
Earnings Review	29
Valuation and Recommendation	30

NeoCloud: a New Paradigm for AI Computing Infrastructure

NeoCloud refers to a new category of cloud computing providers specializing in GPU computing rental services, optimized specifically for AI large-model training and inference. Unlike traditional cloud vendors (AWS/Azure/Google Cloud) that offer "full-stack services", NeoCloud adheres to a "less is more" principle, focusing on high-performance GPU clusters and delivering raw, software-agnostic computing resources.

Figure 1: NeoCloud vs Traditional Cloud Vendor

	NeoCloud	Traditional Cloud Vendor
Core Resources	Top-of-the-line GPUs	CPU + GPU + Storage + Network
Specialization	AI Training/Inference	General computing
Business Model	Long-term rental	Pay-as-you-go
Customer Group	CSPs, AI giants, scientific institutions	Enterprises across all industries
Deliverables	Server + cooling + operations & maintenance + SLA	GPUs rent by day/month/year

Sources: GF Securities (Hong Kong) Brokerage

Strong Capex from upstream: CSPs + AI Giants


Top CSPs are investing aggressively to meet their AI infrastructure needs to ensure they maximize their strategic flexibility over the coming years. Among the investments, they prioritize the expansion of data center footprint. Apart from their strategic AI focus, CSPs are also receiving record revenue and backlog growth in their cloud business, the strong results reinforce their conviction to invest the capital required to continue capturing the AI opportunity.

- 1) Google: Google revised up its 2026 Capex from the prior \$175-185bn USD to \$180-190bn, where significant amount will be spent on data centers and servers. Management also noted that 2027 capex will show a significant expansion to support ongoing AI platform leadership and fulfillment of committed cloud agreement.
- 2) Meta: During its 1Q26 earnings call, Meta revised up 2026 Capex to \$125-145bn USD (previously \$115-135bn, reflecting increasing component pricing and increasing data center costs. Management guided that the cloud deals signed will come online over the course of 2026/27. These multi-year deals drove their contractual commitments to rise by \$107bn.
- 3) Amazon: During its 4Q25 earnings call, management guided a \$200bn USD capex in 2026, which predominantly in AWS (its cloud computing business). In 2025, AWS added the highest number of data centers than any other company, and the company will continue this trend to sustain its leading position.
- 4) Microsoft: Management guided a higher Capex growth in 2026 than 2025, due to higher spending on GPUs and CPUs for Azure platform demand. Full year 2026 Capex is expected to be around \$190bn USD, where \$25bn was impacted by higher component prices. Even with these investments, management still expects to remain constrained at least through 2026.

- 5) Oracle: Management guided \$50bn USD Capex for FY26. For FY27, we expect Oracle to continue increase Capex according to its additional funding of \$50bn from bond issuance. On the other hand, some of Oracle's strategic operational investment including a 10GW data center and power build, are largely funded by partners, materially decoupling Oracle's Capex requirements from out-of-pocket expenditures.

We revise up our CAPEX forecast accordingly from +62% to +82% YoY for Top 5 CSPs in 2026E. We expect another 30% YoY increase in CSP Capex hike in 2027 (vs. consensus +26%), driven by significant upside in the AI accelerator market, despite anticipated memory cost inflation. This is further bolstered by Google's 2027 guidance, which points to a substantial CAPEX step-up fueled by an "explosion" in Cloud backlog and intensifying compute needs across its ecosystem.

Figure 2: Capex of Top 5 CSPs (bn USD)



Sources: GF Securities (Hong Kong) Brokerage


Figure 3: Global AI accelerator market forecast

	2024	2025E	2026E	2027E
GPGPU (Units m)				
NVIDIA	3.9	6.0	8.5	7.6
AMD	0.5	0.5	0.8	1.1
Others	0.7	1.3	1.6	2.5
Total	5.1	7.8	10.9	11.1
YoY	160%	53%	39%	3%
ASIC (Units m)				
Google	2.8	2.9	4.5	8.8
AWS	0.9	1.6	2.6	3.1
Meta	0.2	0.2	0.3	0.5
MSFT	0.1	0.1	0.2	0.5
Other ASICs	0.3	0.8	1.1	1.5
Total	4.2	5.6	8.7	14.3
YoY	-39%	32%	56%	64%
Market (USD bn)				
NVIDIA	88	160	300	451
Total GPGPU	98	176	328	502
Total ASIC	20	30	54	117
Total Market	236	382	710	1121
YoY	645%	62%	86%	58%

Sources: Company data, GF Securities (Hong Kong) Brokerage

As major CSPs keep ramping up Capex, they are facing mounting cash flow pressure, which makes them increasingly reliant on NeoCloud players to share such burdens. Taking Amazon as an example, its cash flow from operations (CFO) stood at \$140 billion in 2025, exceeding its capex of \$132 billion. However, we forecast its 2026 capex to reach \$200 billion, surpassing the projected CFO of \$183 billion. For Microsoft, Google and Meta, although their 2026 operating cash flow will still outpace capex, the divergent growth rates between the two metrics will soon lead to a reversal of this trend. Despite recent record-high revenue posted by Microsoft Azure, AWS, Google Cloud and Oracle Cloud Infrastructure, cloud business revenue growth remains slower than capex expansion across these hyperscalers.

Figure 4: CSP Operating Cash Flow vs. Capex (bn USD)



Sources: GF Securities (Hong Kong) Brokerage

Figure 5: CSP Operating Cash Flow / Capex / Cloud Revenue Growth Comparison

	FY2025 YoY	FY2026 YoY	FY2027 YoY
MSFT			
CFO	14%	27%	22%
Capex	45%	122%	75%
Azure Revenue	13%	39%	40%
AMZN			
CFO	21%	31%	29%
Capex	59%	52%	24%
AWS Revenue	20%	30%	28%
GOOG			
CFO	32%	28%	23%
Capex	72%	107%	40%
Google Cloud Revenue	36%	62%	50%
META			
CFO	27%	21%	26%
Capex	77%	100%	33%
Total Revenue	22%	26%	19%
ORCL			
CFO	11%	24%	45%
Capex	200%	138%	40%
Cloud Infra Revenue	48%	168%	103%


Sources: GF Securities (Hong Kong) Brokerage

Additionally, we expect leading AI companies like OpenAI and Anthropic to expand their investment in AI infrastructure build-out for model trainings.

- 1) OpenAI: In its investor memo in Apr, OpenAI states that they have now confirmed over 8GW of AI compute capacity and plan to scale up to 30GW by 2030. Based on our checks, the AI infrastructure build-out agreements that OpenAI has signed with CSPs and NeoCloud players have already covered its mid-term target of 30GW. This includes contracts with CoreWeave worth up to ~\$22.4 billion, encompassing the initial deal signed in Mar 2025 as well as two subsequent expanded agreements signed in May and Sep 2025.
- 2) Anthropic: With Anthropic's ARR surpassing OpenAI's in this April, we project that Anthropic's mid-term power capacity requirements will be no less than OpenAI's 30 GW target. The company has just announced an agreement to lease xAI's Colossus 1 supercomputer, primarily to address recent stringent usage constraints. Anthropic has intensified its compute capacity expansion through a series of major infrastructure deals, including: 5GW with Amazon, multi-GW with Google and Broadcom, 1+GW with Microsoft and NVIDIA, and undisclosed gigawatt-scale collaborations with Fluidstack and CoreWeave. These expanded compute resources have enabled Anthropic to double its five-hour rate limits, eliminate peak-hour restriction reductions, and raise API rate limits across its Claude product suite. Our channel checks indicate that Anthropic is set to announce an additional ~8 GW compute partnership with AMD in the near term, further validating its ambition to exceed OpenAI's power capacity target. While specific

contract values for Anthropic's agreements with CoreWeave remain undisclosed, we anticipate these will be at least on par with OpenAI's \$22.4 billion contract with CoreWeave.

Figure 6: Anthropic & OpenAI's ARR (bn USD)



Sources: GF Securities (Hong Kong) Brokerage

Figure 7: Anthropic & OpenAI's Major Agreements Overview

Collaborator	Date Signed	Duration	Time Online	Scale	Content
Anthropic					
Amazon	2023 + Apr 2026	10-yr	1H26 end of 2026 Future	5GW / Over \$100bn	over one million Trainium2 capacity nearly 1GW of Trainium2+3 capacity Graviton + Trainium4 capacity
Google	Oct 2025	-	2026	"tens of billions of dollars"	one million TPUs
Google + Broadcom	Apr 2026	5-yr	2027	-	multiple gigawatts of next-generation TPU capacity
Microsoft + Nvidia	Nov 2025	-	2026	1+GW	\$30bn of Azure capacity + 1GW Nvidia's capacity
Fluidstack	Nov 2025	-	2026	\$50bn	Datacenter in Texas and New York
CoreWeave	Apr 2026	Multi-yr	-	-	-
xAI	May 2026	-	2026	0.3GW	Use the capacity of Colossus 1 datacenter (220k Nvidia GPUs)
AMD	Yet not disclosed	-	-	~8GW	-
OpenAI					
CoreWeave	Mar 2025	5-yr	2026	\$11.9bn	Cloud compute deal
	May 2025	5-yr	2026	\$4.0bn	Expanded deal
	Sep 2025	6-yr	2026	\$6.5bn	Expanded deal
Oracle	Jul 2025 Oct 2025	5-yr ~10-yr	2025	15GW / \$500+bn	AI data center for "Stargate"
Microsoft	Oct 2025	-	2025	\$250bn	Azure services
AMD	Oct 2025	-	2H26	Warrant for 160mn AMD shares	6GW of AMD GPUs
Broadcom	Oct 2025	3-yr	2H26	10GW	Accelerators and Ethernet solutions
Amazon	Nov 2025 Feb 2026	7-yr 8-yr	2026	\$100bn	EC2 UltraServers + Nvidia GPUs 2 GW of Trainium 3+4 capacity


Sources: Anthropic, OpenAI, GF Securities (Hong Kong) Brokerage

Token Usage Implies a 283GW Global AI Datacenter Power Capacity by 2030

According to OpenRouter, weekly token usage posted an average MoM growth rate of 30% throughout 2025, translating to an annual growth rate of ~1,600% YoY. Based on our conservative forecasts, we expect token usage to surge by 2,000% in 2026 and 1,500% in 2027, before gradually moderating to a 200% growth pace by 2030. On this basis, we project global AI datacenter power capacity will reach 283 GW by 2030, up from 80 GW at end-2025, representing an incremental demand of 203 GW over the next five years, which corresponds to ~\$8.5 trillion USD in total capital expenditure.

We believe the bulk of new capex will be driven primarily by CSPs and other enterprises. Government direct AI subsidies are already factored into corporate Capex, while most indirect subsidies mainly apply to datacenter operational costs such as utility subsidies. The remaining demand will be fulfilled by NeoCloud players. Accordingly, we estimate the total TAM for NeoCloud will hit \$1.7 trillion USD over the next five years, accounting for 20% of overall AI infrastructure construction demand.

Figure 8: Token Usage During the Final Week of Each Month (tn)



Sources: Grid Strategies, GF Securities (Hong Kong) Brokerage

Figure 9: 2030 Global AI Datacenter Demand Analysis

	2030
Daily token usage (tn)	360,000
Compute offloading	0.8
Compression rate	0.15
Traffic peak factor	1.75
theoretical token throughput per GPU (token/sec)	7500
Datacenter utilization	90%
Power per GPU (W)	1806
Redundancy	1.1
PUE	1.1
Total GW needed	283.2

Sources: OpenRouter, GF Securities (Hong Kong) Brokerage

Figure 10: NeoCloud TAM Analysis

	2026	2027	2028	2029	2030
Incremental GW to build from 2025 to 2030*			203		
Equivalent Capex (bn)			8534		
AI Capex from:					
Top 5 CSPs	568	737	885	1061	1274
China Top 4 CSPs	44	50	60	72	86
Other enterprises	200	300	400	500	600
TAM for NeoCloud (2026-2030)			1698		

*IEA estimates, by end-2025 global datacenter capacity was 114.3GW, 70% for AI

Sources: IEA, Company IR Materials, GF Securities (Hong Kong) Brokerage

Relationship with Hyperscalers: Coexistence of Cooperation and Competition


On the one hand, hyperscalers serve as key customers to NeoCloud companies especially CRWV: due to their own GPU supply shortages and limited AI infrastructure optimization capabilities, they often rent high-performance GPU clusters from CRWV to meet the explosive demand for AI computing power. On the other hand, as traditional cloud leaders, hyperscalers are accelerating their AI computing power layout, launching their own GPU-based cloud services, thus forming direct competition with CRWV in the AI computing market. Facing such a dual relationship, CRWV holds its comparative advantages:

- 1) First-to-market advantage in top-tier GPU clusters: CRWV is the first cloud provider to deliver large-scale clusters of NVIDIA's top-tier GPUs (H100, H200, and GB200). For model trainers, the speed of access to AI compute is essential — delays in accessing advanced computing resources can directly lead to missed opportunities in model iteration and market competition.
- 2) Time value: Time taken for hyperscalers to build a new AI cluster is typically 9–18 months; outsourcing the project to Neocloud, whose supply chain is already well-established, can be completed in 3–6 months.
- 3) Superior cost: Neocloud has access to multiple channels (including servers/optical modules/liquid cooling/network equipment) and can achieve better power and energy consumption. Through bulk GPU purchasing, wholesale data center leasing, and energy-efficient operations, CRWV achieves 15-20% lower unit computing costs than hyperscalers such as AWS and Azure, forming a cost-performance barrier that is difficult for hyperscalers to replicate in the short term.
- 4) Exclusive focus on AI compute, no trade-offs with traditional workloads: Unlike hyperscalers that need to balance traditional workloads (e.g., enterprise storage, general computing, and office services) with AI computing demands, CRWV adheres to a "focused strategy" and devotes all its resources to optimizing AI computing infrastructure. This exclusive focus enables CRWV to avoid the resource allocation conflicts faced by hyperscalers, achieving more efficient GPU utilization (MFU exceeding 50%), lower latency, and higher performance for AI-specific tasks—critical advantages for large-model training and inference. Additionally, CRWV's bare-metal GPU clusters, optimized specifically for AI workloads, deliver ~30% better performance than hyperscalers' AI services based on architectures for general computing.
- 5) Outsourcing heavy assets for financial flexibility: lighter on-balance-sheet Capex, more flexible lease terms, and easier upgrades to next-generation GPUs. Hyperscalers can also transfers supply chain risks, schedule pressure, and power allocation problem to more specialized contractors.

Compute Cost Surge: Memory + Power Cost Inflation

According to our calculation, AI server memory demand (HBM and Server DDR) to account for ~40% of total memory demand in 2026 and will grow by 49% in 2027 and 60% in 2028. Taking the price hike recently and follow on, we expect memory cost to account for ~50% of CSP's Capex in 2027 which will further suppress their Capex utilization efficiency. Against this backdrop, NeoCloud's external financing capacity enables CSPs to outsource heavy-asset investments, which will be more conducive to alleviating their cash flow pressure.


Figure 11: Memory Demand by End Markets



Sources: Company data, GF Securities (Hong Kong) Brokerage


In terms of the variable cost for AI datacenter, electricity cost represents 40-60% of the total operational cost, making data center's operational cost to have even stronger volatility as power shortage worsen. We estimate that newly installed AI datacenters across the U.S. will need an additional capacity of ~203GW between 2025 and 2030, excluding power demand from non-AI data centers. By comparison, the U.S. Energy Information Administration stated that 53 GW of new power capacity was added into the national grid in 2025, marking the largest annual expansion since 2002. We believe the aggressive AI datacenter construction will keep tightening U.S. power supply. In Feb 2026, the average blended natural gas power price in the U.S. stood at \$6.57 per cubic foot, rising 30% YoY. Specifically, industrial power prices surged 42% to \$8.40 in Feb 2026 from \$5.92 in Feb 2025, driving a ~20% increase in overall AI data center operating expenses.

Figure 12: GW-Scale Data Center Plan 2026-2027



Sources: Grid Strategies, GF Securities (Hong Kong) Brokerage

Figure 13: U.S. Natural Gas Electric Power Price (USD/k ft³)



Sources: EIA, GF Securities (Hong Kong) Brokerage

Long-Term Take-or-Pay Contracts: Anchor Tenant Strategy

In the capital-intensive NeoCloud industry, securing large-scale, long-term "anchor tenant" agreements is a foundational prerequisite for validating business models, securing capital, and scaling operations. CoreWeave has cemented its market leadership by securing landmark, large-scale long-term contracts with AI hyperscalers. The strategic value of these "anchor agreements" is multifaceted: 1) they serve as the ultimate market endorsement of a company's technological capabilities and operational excellence, validating its ability to deliver high-performance, reliable AI infrastructure at scale; 2) they effectively lock in long-term demand, providing investors with clear visibility into medium- to long-term revenue trajectories and mitigating growth uncertainty; 3) they underpin a debt-driven growth model—long-term agreements with highly creditworthy tech giants provide robust liquidity support, enabling NeoCloud players to access large-scale debt financing. This debt leverage enables infrastructure expansion at a rate that would be unachievable through equity financing alone, accelerating market penetration and reinforcing scale advantages that are difficult for competitors to replicate.

Figure 14: The importance of "anchor agreements"




Sources: GF Securities (Hong Kong) Brokerage

The Sustainability of Asset-Backed Debt Model

CoreWeave have secured billions of dollars in debt financing by using their GPU clusters as collateral, underscoring that high-end GPUs are no longer merely seen as technological products, but rather as a stable, revenue-generating asset, similar to real estate, securitizable and leveragable. For NeoCloud companies, its ability to repay its debt depends on whether its contracted revenue can consistently cover debt costs; its profitability depends on whether contracted revenue can cover capital costs and operating costs. Under this model, a potential risk could be that technological iteration may devalue NeoCloud's GPU holdings as assets.

This March, CoreWeave announced it has successfully closed its first investment-grade rated GPU-backed financing of an \$8.5bn DDTL (delayed draw term loan) 4.0 facility which matures in March 2032. The DDTL 4.0 Facility received ratings of A3 by Moody's and A (low) by DBRS, features non-recourse structure, includes a floating rate tranche financed at SOFR + 2.25% and a fixed rate tranche financed at approximately 5.9% (v.s. SOFR + 9.62% for DDTL 1.0 and SOFR + 4.25% for DDTL 2.0 and SOFR +4% for DDTL 3.0).

Figure 15: CoreWeave's DDTL4.0



	Change	Impact to CoreWeave	Commentary
Cost of Debt	S + 225 bps / 5.9% ⁽¹⁾	●	Additional 175 bps reduction in cost of capital vs. DDTL 3.0 and ~750 bps decrease vs. inaugural DDTL 1.0
Parent-Level Debt	Unchanged	●	Parent-level debt is unimpacted given lack of full recourse parent guaranty
Total Debt	+\$8.5 billion	●	Increase to consolidated debt, in-line with previous DDTLs
Additional Capital Available to CoreWeave at Stabilization	+\$1.0 billion	●	Unlocks significant additional capital at investment grade cost, reducing ParentCo's need for financing
Annualized EBITDA	Increases	●	Growth in EBITDA while reducing Parent level leverage as deployments stabilize
Levered Free Cash Flow	Increases	●	Continues to build on strong foundation of contracted levered free cash flow

Sources: CoreWeave, GF Securities (Hong Kong) Brokerage

CoreWeave: NeoCloud Leader

Under tight computing power supply and soaring GPU prices, CRWV announced a series of significant customer partnership agreements in the first quarter of 2026, including a \$21 billion expanded agreement with Meta, a \$6 billion deal with Jane Street (its first major non-CSP/AI giant major customer), and a partnership with Anthropic. We believe the continued upbeat sentiment among CoreWeave's clients indicates that rising compute cost have not diminished their willingness to expand AI infrastructure. More importantly, clients remain confident on CoreWeave's ability to 1) obtain NVIDIA's top-tier GPUs; 2) deploy large-scale of GPUs clusters with high performance and cost efficiency.

CoreWeave maintains leading advantages over its key NeoCloud rivals -- Nebius, Crusoe, and others in the following aspects: 1) secure the exclusive priority supply right of NVIDIA's top-tier GPUs (the first cloud platform to reach NVIDIA's exemplar cloud status for GB200, and the first to bring NVIDIA's new Rubin GPU platform to market in 2H26); 2) the large-scale deployment of GPUs and extreme infrastructure optimization, forming an insurmountable scale effect and performance barrier; 3) the mature Take-or-Pay long-term contract model which brings stable pricing; 4) reduced cost-of-capital backed by strong DDTLs.

Figure 16: CoreWeave vs Competitors

	CoreWeave	Nebius	Crusoe	Together.ai
Core Advantage	Scale effect, leading construction speed, extreme cluster optimization	Full AI ability with AI Cloud and Token Factory	Energy-first approach, Broad enterprise cloud service	Open-source AI platform, diversity of model selection
Key Customers	AI giants, CSPs, large enterprises	Large enterprises and startups	ESG-focused enterprises, mid-sized AI companies	Developers and AI-native startups
Customer concentration	High	Mid	Low	Low
Go-to-market strategy	Focus on billion-dollar deals with tech giants to get high backlogs and high revenue visibility	Billion-dollar deals + diversify customers (different end markets)	Engage in large infra projects (such as "Stargate") + target AI startups	Attract customer by massive open-source models, more product-oriented
Business Model	Long-term Take-or-Pay rental (3-7 years) + value-added services, stable pricing	Short-to-medium term rental (1-3 years) + pay-as-you-go	Long-term rental + carbon-neutral service premium	Short-to-medium + pay-as-you-go


Sources: GF Securities (Hong Kong) Brokerage

CoreWeave's Ongoing "Efficiency Gap" Compared to Alternatives

CoreWeave achieves a 20% higher performance than public foundation model training benchmarks shown by its 50+% Model FLOPS Utilization (MFU), vs. 35%-45% for alternative solutions. The complexity of managing AI infrastructure means that a majority of the compute capacity in GPUs can be lost to system inefficiencies. Bridging the efficiency gap between the typical 35–45% of MFU and the theoretical 100% presents a major opportunity to unlock the full potential of AI infrastructure. Closing this gap can significantly enhance performance, improve model quality, accelerate development timelines. The reduction of overall AI model costs is significant under current computing power scarcity.

CoreWeave's GPU performance has exceeded Nvidia's reference target performance has become the first cloud provider in the world to be named an Nvidia Exemplar Cloud for training workloads running on Nvidia H100 and GB200 NVL72. In figure 18, CoreWeave achieved better H100 MFU performance than two benchmark with the same parallelism hyperparameters. For the comparison between CoreWeave and Nvidia's reference target, CoreWeave also achieved better performance across different models on GB200 NVL72: 1.9%/4.7%/3.8%/2.4% higher MFU than Nvidia reference target on DeepSeek v3 (BF16)/ Grok-1 314B (BF16)/ Llama 3.1 405B (FP8)/ Llama 3 70B (FP8).

Figure 17: CoreWeave's MFU performance vs Nvidia reference target on GB200 NVL72



Sources: CoreWeave, GF Securities (Hong Kong) Brokerage

Figure 18: CoreWeave's MFU vs Benchmark

Benchmark Run	Benchmark MFU	CoreWeave Run	CoreWeave MFU	Difference
Aleph Alpha/ Hasso Platner Institue:				
30B Llama model	40.43%	30B Llama model	51.90%	28% improvement
8K tokens		8K tokens		
128 A100 GPUs		128 H100 GPUs*		
MosaicML/ Databricks:				
30B MPT model	41.85%	30B MPT model	49.20%	18% improvement
2K tokens		2K tokens		
128 H100 GPUs		128 H100 GPUs		

*Comparable even with different GPU types, as MFU is not an absolute score but a percentage of the accelerators' theoretical FLOPs


Sources: CoreWeave, GF Securities (Hong Kong) Brokerage

Software Stack Along with Bare-Metal to Provide Better Performance

CoreWeave leverages a unified stack with performance optimizations across every layer from metal to model, maximizing performance, resiliency, and efficiency with deep integration between hardware and software. CoreWeave Cloud Platform's advanced cluster validation, health monitoring, proactive node replacement, and deep observability mean the workloads can consistently run on healthy infrastructure, significantly reducing the likelihood of disruptions. By minimizing interruptions, it can achieve a goodput rate of 96%, maintaining consistently high MFU throughout training.


- **CoreWeave Kubernetes Service (CKS)** provides the base runtime environment and is fully integrated with CoreWeave Mission Control to minimize overhead and maximize compute performance.
- **CoreWeave Slurm on Kubernetes (SUNK)** enables topology aware scheduling, which optimizes performance and utilization for large-scale training clusters. Based on a battle-tested Slurm distribution scaled to handle tens of thousands of nodes and hundreds of thousands of concurrent jobs.
- **CoreWeave Cabinet Visualizer** monitors the health of our full rack systems. The dashboard allows customers to timely see their infrastructure and managed service health and performance, such as Infiniband bandwidth, GPU temperature, power consumption, and real-time alerts, so that customers can troubleshoot and resolve issues faster and maximize the performance of their workloads.
- **CoreWeave Mission Control** provides infrastructure level monitoring, and Weights & Biases for job level monitoring.
- **CoreWeave GPU Straggler Detection** provides real-time collective metrics, including Bus Bandwidth, and automatically detects any hardware lockups and pinpoints the root cause.

Figure 19: CoreWeave's Layered Architecture Stack




Sources: CoreWeave, GF Securities (Hong Kong) Brokerage

Figure 20: CoreWeave's Cabinet Visualizer dashboard



Sources: CoreWeave, GF Securities (Hong Kong) Brokerage

Figure 21: CoreWeave's Mission Control dashboard




Sources: CoreWeave, GF Securities (Hong Kong) Brokerage

Recent New Contracts Diminishing the Risk of Customer Dependence

In 2024, Microsoft was CoreWeave's largest customer, accounting for 67% of its revenue; in the latest Q1 2026 financial report, Microsoft and OpenAI accounted for 45% and 20% of total revenue, respectively, raising concerns about CoreWeave's customer concentration. However, CoreWeave recently announced new contracts with Meta, Anthropic, and Jane Street, significantly reducing this risk. According to our forecasts, by 4Q27, Microsoft and OpenAI will only account for 16%/17% of total revenue, while Meta, Anthropic, and Jane Street will account for 9%/6%/4%, respectively. By 4Q28 (under the current contract conditions), Meta will become CoreWeave's largest customer, accounting for 16% of revenue, while Microsoft, OpenAI, Anthropic, and Jane Street will account for 11%/11%/4%/3%, respectively, significantly reducing reliance on a single major customer.

Figure 22: CoreWeave's Major Contracts




Sources: GF Securities (Hong Kong) Brokerage

Strong RPO Backed up by LTAs

In April, Meta announced an expansion agreement of \$21bn on top of its \$14.2bn initial agreement from last year. This contract, valid until 2032, will grant Meta the initial deployment rights to NVIDIA's next-generation Vera Rubin chips. As CoreWeave CEO Michael Intrator previously stated, "Despite Meta's massive internal investments in data centers, ...continues to rely on CoreWeave for AI computing power due to the superior quality and specialization of CoreWeave's infrastructure." Following this, CoreWeave won another LTA with Anthropic. Although specific contract amount was not disclosed, we expect it to be no less than the initial agreement CoreWeave signed with OpenAI in March 2025. CoreWeave will provide Anthropic with massive computing capacity to support the development and deployment of its Claude series models. In the same month, Jane Street also announced a \$7bn agreement with CoreWeave, marking CoreWeave's first non-AI-giants LTA customer and signifying its customer expansion into different industries. These three large orders directly lift CoreWeave's RPO to a record \$98.8 billion in 1Q26, +63% QoQ and +572% YoY.

Figure 23: CoreWeave's RPO




Sources: GF Securities (Hong Kong) Brokerage

Debt Might Be the Only Concern

To achieve a \$98.8 billion RPO, CoreWeave will have to pay an extremely heavy financial price. Looking at the 1Q26 financial report, although quarterly revenue more than doubled YoY to \$2.1bn, and adjusted EBITDA reached \$117mn, but adjusted net loss was \$589 million (vs. \$150 million in 1Q25). The culprit behind the compressed profitability is its continuously increasing Capex. To accelerate the construction of data centers and the procurement of chips, CoreWeave burned \$1.12bn in 2025 and announced that Capex will further double to \$30-35bn in 2026. Compared to Microsoft and Amazon's tens of billions of dollars in cash reserves, CoreWeave only had about \$2.3 billion in cash on hand in Q1 2026, and it consistently had negative free cash flow. The company's solution to this problem is to fill the gap through debt.


Currently, the company's total liabilities have reached \$24.9 billion. Recently, the company successfully obtained DDTL 4.0 with an interest rate of ~5.9%, further reducing its finance costs. The company has consistently maintained positive adjusted EBIT, but after deducting interest costs, it enters a loss-making state. If the company can: 1) continuously reduce debt costs; 2) maintain contract revenue growth; and 3) reduce operating and other costs, it can improve its profitability. The company's breakeven point will occur when its EBIT is sufficient to cover interest payments.

Figure 24: CoreWeave's Debt & Interest Rate (mn USD)



Sources: GF Securities (Hong Kong) Brokerage

Figure 25: CoreWeave's Cash Flow from Operations & Capex (mn USD)



Sources: GF Securities (Hong Kong) Brokerage


Rising GPU Rental Price Will Benefit CoreWeave's New Contracts

During the first week of May, total global token usage on OpenRouter platform reached 21.9 trillion, +1176% YoY. Under this tremendous growth in token usage, users have experienced significant 1) service limitations: some Anthropic's Claude users hit their 5-hour usage limits in just 20 minutes; OpenAI has begun shuttering Sora, as the number of developers using its coding assistant Codex has surged to four million per week. 2) API price increase for high-end models: newly launched GPT-5.5 has seen a ~2x price increase over GPT-5.4, leading to estimated cost increases of 49-92% for users.

Due to the shortage in computing power, H100 rental price has risen ~36% in 1Q26 (see figure 26), and B200's rental price has risen 23.5% just in March. The price ratio of B200 to H100 rebounded from the initial 2.63 when B200 first launched to 2.07 in March, suggesting that the shortage of computing price has led to an increase in the rental price of even older generation GPUs.


Benefiting from rising GPU rental prices, we expect pricing for newly-signed contracts in 2026 to be markedly higher than those inked in 2025 and prior years. Contracts secured in 2026 account for 40% of the company's total current RPO. As these new deals start generating monthly payments from 1Q27 onwards, we forecast CoreWeave's blended ASP will rise by 20%. Upon the expiry of existing contracts with major clients, the completed data center capacity can be re-contracted under new multi-year agreements or switched to on-demand consumption models, which will drive another round of price hikes.

Figure 26: 1-year Contract Term Average Rental Price (USD per hour per GPU)



Sources: SemiAnalysis, GF Securities (Hong Kong) Brokerage

Figure 27: H100 80GB Rental On-Demand Price Comparison, May 2026 (USD per GPU per hour)



Sources: Company data, GF Securities (Hong Kong) Brokerage

Pace to Profit-making


When entering into large-scale 3–7 year contracts, CoreWeave pre-factors and incorporates three key cost components into its pricing model: 1) the full procurement cost of GPUs; 2) total interest expenses on borrowings raised to fund these GPU purchases over the contract term; 3) infrastructure costs including data center construction, power, cooling and maintenance.

In essence, customers commit to long-term contract payments that are sufficient for CoreWeave to fully recoup its costs and generate profits over the core economic lifespan of the GPUs. Through this mechanism, even if the GPUs suffer substantial market value depreciation or are outdated by new technologies upon the expiry of the contract, CoreWeave is able to lock in predetermined profits from the transaction.

We believe most of CoreWeave's major contracts are currently at one of the following phases: Phase 1 of "Contract Signing", Phase 2 of "Infrastructure Purchase & Installation", or the early Phase 3 of "AI Cloud Go-Live". As a result, a large portion of the company's current outlays are cash payments for GPUs and other infrastructure upon delivery.

Under its business model, an individual contract only starts to generate profit in the late stage of Phase 3 (Microsoft in 2H26) or Phase 4 (Microsoft in 2H27). Theoretically, profitability will improve materially once the majority of contracts, or a meaningful number of contracts, progress to this profit-generating stage.


Figure 28: CoreWeave's Business Model



Sources: CoreWeave, GF Securities (Hong Kong) Brokerage


In 1Q26, rent and utility costs accounted for 57% and 22% of CoreWeave's cost of revenue, respectively, with the remaining costs—depreciation and amortization for power installation and distribution systems, labor, and stock-based compensation—each below 10%. Based on our forecasts of power shortages in the US and data center land rent prices, we predict that rent and utility costs per WM will continue to rise. We project rent per WM to reach \$0.42/\$0.45 by end-26/end-27, and utility & power costs per WM to reach \$0.15/\$0.18, bringing the cost of revenue for 4Q26/4Q27 to \$0.71/\$0.71, representing YoY increases of 20%/9%, respectively.

Figure 29: CoreWeave's Cost of Revenue Analysis (mn USD)




Sources: GF Securities (Hong Kong) Brokerage

Figure 30: CoreWeave's Cost of Revenue per MW (USD)



Sources: GF Securities (Hong Kong) Brokerage


Figure 31: CoreWeave's Non-GAAP GP (mn USD)



Sources: GF Securities (Hong Kong) Brokerage


In 1Q26, technology and infrastructure cost (adjusted) accounted for 90%/59% of CoreWeave's adjusted OPEX and revenue. Adjusted technology and infrastructure expense consists of 1) depreciation and amortization related to servers, switches, networking equipment and internally developed software, 2) research & development cost, 3) other related expenses. We expect as the company achieve economies of scale and operational efficiencies. We expect the Company to boost operational efficiency and capture greater economies of scale going forward. Nevertheless, due to its sustained investment in hardware, the Company's T&I expense as a percentage of revenue will decline, though the drop will not be very significant. We project that the Company's adjusted T&I expense will account for 56% and 52% of total revenue in 4Q26 and 4Q27 respectively. Based on our projections for T&I expense, we expect Coreweave's Non-GAAP operating profit to deliver sustained improvement, rising from 1% in 1Q26 to 6% in 4Q26 and 11% in 4Q27. As discussed earlier, the company's profitability hinges on whether its operating profit can cover interest expenses. As the Company's OP grows to \$1.6bn with a 16% OPM in 2Q28, it will reach its break-even point, turning from losses to a net profit of \$174mn.

Figure 32: CoreWeave's Non-GAAP T&I Expense (mn USD)




Sources: GF Securities (Hong Kong) Brokerage

Figure 33: CoreWeave's Non-GAAP OP (mn USD)



Sources: GF Securities (Hong Kong) Brokerage

Figure 34: CoreWeave's Non-GAAP Net Income (mn USD)



Sources: GF Securities (Hong Kong) Brokerage

Nvidia Both an Investor and a Supplier

Due to considerations of sales channels and production capacity, Nvidia has chosen to back NeoCloud players, namely CoreWeave and Nebius. NeoCloud providers excel at delivering GPU clusters within 3–6 months. This enables NVIDIA to: 1) monetize GPU production capacity more rapidly and narrow the time lag between capacity expansion and revenue realization; 2) obtain timely feedback from frontline AI training and inference scenarios to refine and iterate on its next-generation products; 3) diversify reliance on large-scale CSPs and build out a strong, secondary sales channel. Through these partnerships, Nvidia extends its product advantages into sustained channel and financial advantages. We expect this ongoing support from Nvidia to continue, which empowers CoreWeave to consistently secure access to the latest GPU capacity, deploy asset-backed debt financing, and keep running its proven business model steadily.

Figure 35: Nvidia & CoreWeave's Collaboration

Date announced	Content
Sep 2020	CoreWeave join the Cloud Service Provider Program within the NVIDIA Partner Network (NPN) as a Preferred Partner
Jul 2021	CoreWeave becomes Nvidia's first elite cloud services provider for compute
Dec 2024	CoreWeave is one of the first major cloud providers to bring up an NVIDIA GB200 NVL72 cluster
Feb 2025	CoreWeave is the first cloud provider to make NVIDIA GB200 NVL72-based instances generally available, setting a new benchmark for AI performance and scalability.
Mar 2025	Nvidia anchors CoreWeave's initial public offering at \$40 a share with a \$250 million order
Jul 2025	CoreWeave becomes the first AI cloud provider to offer NVIDIA GB300 NVL72 and RTX PRO 6000 Blackwell GPU at Scale
Sep 2025	Announced \$6.3B Capacity Deal, Nvidia will purchase any unused cloud capacity from CoreWeave through April 2032
Jan 2025	Nvidia invests \$2B in CoreWeave to accelerate the 5GW buildout by 2030. Test and validate CoreWeave's AI-native software and reference architecture (SUNK, CoreWeave Mission Control), and work towards including those offerings within NVIDIA's reference architectures for NVIDIA's cloud partners and enterprise customers. Early adoption of NVIDIA computing architectures, including the Rubin platform, Vera CPUs and Bluefield storage systems.


Sources: GF Securities (Hong Kong) Brokerage

Company Overview

Founded in 2017, CoreWeave is a leading provider of NeoCloud, specializing in high-performance GPU compute, storage, networking, and orchestration solutions tailored to enterprises and research organizations with large-scale AI workloads. CoreWeave differentiates itself by its bare-metal GPU clusters, end-to-end AI model lifecycle management, including training, fine-tuning, deployment, and inference on its cloud platform.

CoreWeave transitioned to cloud infrastructure services from a cryptocurrency mining operation. Since this strategic pivot, the company has accelerated the expansion of its infrastructure footprint, establishing a network of large-scale, GPU-rich data centers across the United States and Europe.

Figure 36: Company Development



Sources: Company IR material, GF Securities (Hong Kong) Brokerage.

Figure 37: CoreWeave's Infrastructure Pace

	1Q24	2Q24	3Q24	4Q24	1Q25	2Q25	3Q25	4Q25	1Q26
Total Active Power (MW)	116	227	322	360	420	470	590	850	1000
YoY %					261%	107%	83%	136%	138%
New CapEx (ex. CIP)	1250	1617	1843	1950	1858	2938	1850	8241	6787
YoY %					49%	82%	0%	323%	265%
Number of Data Center Owned	Undisclosed	Undisclosed	Undisclosed	32	33	33	41	43	49

Sources: GF Securities (Hong Kong) Brokerage

Earnings Review

1Q26 results review

CoreWeave reported 1Q26 revenue of \$2.08bn, +32% QoQ and +112% YoY. Non-GAAP gross profit was \$1.37bn, Non-GAAP GPM at 65.9%, -3.1ppt QoQ and -8.5ppt YoY. The decline in GPM was mainly due to the increase in rent, utility, and D&A cost. Non-GAAP EBIT came in at \$21mn, EBIT margin at 1.0%, -4.6ppt QoQ and -15.6ppt YoY, mainly due to increasing Technology and Infrastructure cost. The company posted a Non-GAAP net loss of \$589mn, +107% QoQ and +294% YoY, due to higher interest cost.

We forecast revenue of \$12/26/40bn USD in 2026/2027/2028

We expect CoreWeave's revenue for 2026, 2027, and 2028 to be 12.3bn, 25.6bn, and 39.6bn USD, representing YoY growth of 139.1%, 108.6% and 54.8%, respectively; Non-GAAP gross profit is forecasted at 8.4bn, 18.2bn, 27.7bn, GPM at 68.8%, 71.1%, and 69.9%; Non-GAAP net income is expected to be -2.1bn, -1.0bn, and 198mn.

Figure 38: CoreWeave's P&L


(USD m)	FY25	1Q26	2Q26	3Q26E	4Q26E	FY26E	1Q27E	2Q27E	3Q27E	4Q27E	FY27E
Sales	5,131	2,078	2,669	3,436	4,083	12,266	5,122	6,063	6,824	7,573	25,582
COGS	-1,438	-708	-834	-1,045	-1,235	-3,822	-1,459	-1,692	-1,962	-2,275	-7,388
Gross profit	3,693	1,370	1,836	2,391	2,848	8,445	3,663	4,371	4,862	5,298	18,194
Opex	-3,641	-1,349	1,689	2,188	2,598	5,126	3,087	3,623	4,044	4,430	15,184
Operating profit	668	21	147	203	250	621	576	748	818	868	3,010
Non-op profit	-1,170	-512	-577	-622	-672	-2,383	-901	-939	-978	-1,020	-3,838
Pre-tax profit	-502	-491	-430	-419	-421	-1,762	-325	-190	-161	-152	-828
Income tax	-83	-98	-77	-70	-64	-308	-62	-49	-49	-49	-209
Non-GAAP Net Income	-606	-589	-507	-489	-485	-2,070	-387	-240	-210	-201	-1,037
Diluted EPS (US\$)	-2	-1.12	-0.96	-0.93	-0.92	-4	-0.73	-0.45	-0.40	-0.38	-2
Diluted shares	506	527	527	527	527	527	527	527	527	527	527
Margin analysis											
Gross margin	72%	66%	69%	70%	70%	69%	72%	72%	71%	70%	71%
Operating margin	13%	1%	5%	6%	6%	5%	11%	12%	12%	11%	12%
Pre-tax margin	-10%	-24%	-16%	-12%	-10%	-14%	-6%	-3%	-2%	-2%	-3%
Effective tax rate	-17%	-20%	-18%	-17%	-15%	-17%	-19%	-26%	-30%	-32%	-25%
Growth (YoY%)											
Sales	168%	112%	120%	152%	160%	139%	0%	192%	156%	120%	109%
Operating profit	88%	-87%	-27%	-7%	184%	-7%	-14%	3464%	458%	328%	385%
Net income	837%	294%	287%	1093%	71%	242%	-36%	-59%	-59%	-59%	-50%
EPS	427%	86%	258%	1027%	64%	160%	-51%	-59%	-59%	-59%	-50%

Valuation and Recommendation

Excluding 1Q25, CoreWeave's stock price has consistently trended downward following all other earnings releases. Investors, after marveling at the company's strong RPO, are often feared by its high interest expenses and debt. Following the latest earnings call, the company's stock dropped 9% in after-hours, primarily due to a sequential decline in gross margin and operating margin in 1Q26, as well as the company's 2Q26 revenue guidance (\$2.45bn to \$2.6bn), which came in below the previously market consensus of \$2.69bn.

Despite facing significant near-term profit pressure, we remain optimistic about CoreWeave's business model. We believe that once the majority of the company's contracts, or a meaningful number of them, enter the profit-making phase, its profitability will improve significantly — sufficient to cover the interest expenses. Meanwhile, the company's recent DDTL 4.0 has proven to the market its ability to secure lower financing rates. Moreover, we expect CoreWeave's new contract to drive up its blended ASP by 20% as they ramp into monthly payment phases in 1Q27. We project that CoreWeave will reach its break-even point in 2028.

Figure 39: CoreWeave's Historical Stock Price



Sources: Bloomberg, GF Securities (Hong Kong) Brokerage

Initiate with Buy with TP of \$162

We are positive on 1) accelerating demand for AI buildouts; 2) RPO that largely underpin our 2026–2027 revenue forecasts; 3) declining debt financing costs; 4) profitability to inflect sharply as large contracts mature; 5) new LTAs benefiting from higher GPU rental pricing. Historically, CoreWeave's share price has been driven by large contract wins; we expect the next key catalyst to be the announcement of new signed agreements or further disclosure regarding the size of its contract with Anthropic. We forecast revenue of \$12.3/25.6/39.6bn in 2026/2027/2028, with adjusted EBITDA reaching to \$7.6/14.2/18.8bn. Our target price of \$162 is based on 10x 2028E EV/EBITDA.

Figure 40: Compare Table

Company	Market Cap (bn USD)	EBITDA (mn USD)				25-28 CAGR	Net Profit (Million USD)			
		2025	2026	2027	2028		2025	2026	2027	2028
CoreWeave	55458	3135	7349	15490	24571	96%	(578)	(1688)	(239)	1834
Nebius	48703	(60)	1520	6657	12878	-656%	(79)	(627)	(530)	477
Amazon	2850742	167197	211656	266155	328551	26%	95112	111997	126539	158465
Google	4688479	174183	229638	289532	361161	24%	135738	176764	182870	217783
Microsoft	3127817	157814	198959	239289	295056	42%	100164	127201	144077	170418
Oracle	541157	30137	36412	50161	72224	24%	17158	21853	23821	32613

Sources: Bloomberg, GF Securities (Hong Kong) Brokerage

Risks

1) AI demand deceleration; 2) Geo-political uncertainties; 3) Competition.

Balance Sheet		USD mn				
	FY24	FY25	FY26E	FY27E	FY28E	
Current assets	5,609	5,609	5,691	9,034	12,713	
Cash and cash equivalents	1,398	3,980	1,428	1,509	1,722	
Inventory	0	0	0	0	0	
Accounts Receivable	417	3,169	3,817	7,079	10,546	
Other current assets	3,794	(1,540)	446	446	446	
Non-current assets	15,917	41,814	84,017	163,291	313,450	
Long-term equity investment	0	0	0	0	0	
Fixed assets	11,915	30,557	58,135	95,675	127,061	
Other long-term assets	721	1,506	1,734	1,734	1,734	
Total assets	21,526	47,423	89,708	172,325	326,164	
Current liabilities	4,963	16,440	17,421	22,371	28,483	
Accounts Payable	868	1,623	5,881	10,831	16,460	
Short-term borrowings	2,468	6,708	4,640	4,640	5,123	
Other current liabilities	1,626	8,109	6,900	6,900	6,900	
Non-current liabilities	11,562	29,527	52,039	87,046	114,535	
Long-term Debt	5,458	14,665	36,355	71,362	98,851	
Other non-current liabilities	6,104	14,862	15,684	15,684	15,684	
Total liabilities	16,524	45,967	69,460	109,417	143,019	
Total Equity	5,001	1,456	20,247	62,909	183,145	
Paid-in capital	0	0	0	0	0	
Treasury stock	(34)	(34)	(34)	(34)	(34)	
Retained earnings	(1,476)	(2,643)	(4,172)	(4,286)	(3,002)	
Minority interests	4,789	4,133	24,420	67,195	186,147	
Total Liabilities & Equity	21,526	47,423	89,708	172,325	326,164	

Income Statement		USD mn				
	FY24	FY25	FY26E	FY27E	FY28E	
Revenue	1,915	5,131	12,266	25,582	39,611	
Cost of sales	(492)	(1,438)	(3,822)	(7,388)	(11,923)	
Gross profit	1,423	3,693	8,445	18,194	27,689	
Operating Expense	1,097	3,641	(5,126)	(15,184)	(21,681)	
Operating profit	2,520	7,334	3,319	3,010	6,007	
Interest Income	(361)	(1,229)	(2,479)	(3,934)	(5,850)	
Interest Expense	0	0	0	0	0	
Net other Non-op. Income/(Loss)	(708)	60	96	96	96	
Pre-tax profit	54	(524)	(1,762)	(828)	253	
Income tax	119	83	308	209	56	
Profit for the year	(65)	(606)	(2,070)	(1,037)	198	
Minority interest	0	0	0	0	0	
Net profit to ord. equity	(65)	(606)	(2,070)	(1,037)	198	
EPS (US\$)	(4.30)	(2.81)	(3.93)	(1.97)	0.37	

Cash Flow Statement		USD mn				
	FY24	FY25	FY26E	FY27E	FY28E	
Operating cash flow	2,750	3,058	9,296	13,039	16,006	
Profit for the year	(863)	(1,167)	(1,951)	(1,037)	198	
Depreciation & amortization	863	2,454	6,986	11,160	12,823	
Change in working capital	114	(2,495)	1,158	(189)	(118)	
Others	2,635	4,266	3,104	3,104	3,104	
Investing cash flow	(8,708)	(14,087)	(35,293)	(48,784)	(44,266)	
Capex	(8,702)	(11,199)	(35,245)	(48,700)	(44,209)	
Change in investment	0	0	0	0	0	
Others	(6)	(2,888)	(48)	(84)	(56)	
Free cash flow	(5,952)	(8,141)	(25,949)	(35,661)	(28,203)	
Financing cash flow	7,466	10,197	23,571	35,826	28,472	
Change in Capital	4	87	2,677	923	1,087	
Net Change in Debt	6,433	9,320	20,998	35,007	27,489	
Others	1,028	790	(104)	(104)	(104)	
Exchange influence	0	0	0	0	0	
Net increase in cash	1,507	(832)	(2,425)	81	213	

Key Financial Ratios		FY24	FY25	FY26E	FY27E	FY28E
Growth						
Revenue growth	-	-	-	-	-	-
Operating profit growth	-	-	-	-	-	-
Net profit growth	-	-	-	-	-	-
Profitability						
Gross profit margin	74.3%	72.0%	68.8%	71.1%	69.9%	
Operating Profit Margin	131.6%	142.9%	27.1%	11.8%	15.2%	
Net profit margin	-3.4%	-11.8%	-16.9%	-4.1%	0.5%	
Key Ratio						
ROA	-1.3%	-41.6%	-10.2%	-1.6%	0.1%	
ROE	-0.5%	-1.8%	-3.0%	-0.8%	0.1%	
Stability						
Gross debt/equity	330.4	3157.1	343.1	173.9	78.1	
Interest Coverage	-	-	-	-	-	
Current Ratio	1.1	0.3	0.3	0.4	0.4	
Quick Ratio	1.1	0.3	0.3	0.4	0.4	
Net debt/equity	130.52	1194.57	195.42	118.42	55.83	

Rating definitions Benchmark: Hang Seng Index (Hong Kong)

Company ratings	Buy	Stock expected to outperform benchmark by more than 10%
	Hold	Expected stock relative performance ranges between -10% and 10%
	Underperform	Stock expected to underperform benchmark by more than 10%

Sector ratings	Positive	Sector expected to outperform benchmark by more than 10%
	Neutral	Expected sector relative performance ranges between -10% and 10%
	Cautious	Sector expected to underperform benchmark by more than 10%

Hong Kong Company	GF Securities (Hong Kong) Brokerage Limited
Address	27/F, GF Tower, 81 Lockhart Road, Wan Chai, Hong Kong
Telephone	(852) 37191111
Email	jasminebai@gfgroup.com.hk

Disclaimer

This report has been prepared by GF Securities (Hong Kong) Brokerage Limited ("GF Securities (Hong Kong) Brokerage"). According to the laws, regulations and regulatory requirements in different countries and regions, this report is distributed by GF Securities (Hong Kong) Brokerage with relevant legal and compliant operation qualifications in these countries and regions.

GF Securities (Hong Kong) Brokerage is licensed by the Securities and Futures Commission of Hong Kong ("SFC") to conduct Type 4 Regulated Activity "Advising on Securities". It is regulated by the SFC, and is responsible for the distribution of this report in Hong Kong. Information about the qualifications of the research analyst(s) who is(are) the author(s) of this report as licensed by the SFC are disclosed in the section where research analyst names are shown.

The research analyst(s) primarily responsible for the content of this report, in whole or in part, certifies that with respect to the company or relevant securities that the analyst(s) covered in this report: 1) all of the views expressed accurately reflect his or her personal views on the company or relevant securities mentioned herein; and 2) no part of his or her remuneration was, is, or will be, directly or indirectly, in connection with his or her specific recommendations or views expressed in this report.

This report is published solely for information purpose and does not constitute an offer to buy or sell any securities or a solicitation of an offer to buy, or recommendation for investment in, any securities.

The securities mentioned in this report may not be allowed to be sold in certain jurisdictions. No action has been taken to permit the distribution of the research report to any person in any jurisdiction that the circulation or distribution of such research report is unlawful. This report is distributed solely to clients or designated institutions authorized by GF Securities (Hong Kong) Brokerage, and is not distributed publicly. It is distributed to certain clients based on the conclusion that they are able to assess investment risks independently, execute investment decisions independently and assume corresponding risks independently.

This report has been issued and based on information obtained from sources generally available to the public and believed by the research analyst(s) to be reliable but which has not been independently verified. No representation or warranty, either express or implied, is made by GF Securities (Hong Kong) Brokerage as to their accuracy and completeness of the information contained in this report. GF Securities (Hong Kong) Brokerage accepts no liability for all loss arising from the use of the materials presented in this report, unless is excluded by applicable laws or regulations. Please be aware of the fact that investments involve risks and the price of securities may be fluctuated and therefore return may be varied, past results do not guarantee future performance. Any recommendation contained in this report does not have regard to the specific investment objectives, financial situation and the particular needs of any individuals. This report is not to be taken in substitution for the exercise of judgment by respective recipients of this report, where necessary, recipients should obtain professional advice before making investment decisions.

GF Securities (Hong Kong) Brokerage may have issued, and may in the future issue, other communications that are inconsistent with, and reach different conclusions from, the information presented in the research report. The points of view, opinions and analytical methods adopted in the research report are solely expressed by the analysts but not that of GF Securities (Hong Kong) Brokerage or its affiliates. The information, opinions and forecasts presented in the research report are the current opinions of the analysts as of the date appearing on this material only which may subject to change at any time without notice. The salesperson, dealer or other professionals of GF Securities (Hong Kong) Brokerage may deliver opposite points of view to their clients and the proprietary trading division with respect to market commentary or dealing strategy either in writing or verbally. The proprietary trading division of GF Securities (Hong Kong) Brokerage may have different investment decision which may be contrary to the opinions expressed in the research report. GF Securities (Hong Kong) Brokerage or its affiliates or respective directors, officers, analysts and employees related to research report business may have rights and interests in securities mentioned in the research report. Recipients should be aware of relevant disclosure of interest (if any) when reading this report.

GF Securities (Hong Kong) Brokerage and its affiliates may be seeking or building business relationships with company(ies) mentioned in this report. Therefore, investors should consider the impact on the independence of this report by GF Securities (Hong Kong) Brokerage and its affiliates due to potential conflicts of interests. Investors should not make any investment decisions based solely on the contents of this report. Investors should make their own investment decisions and bear their own risk. No written or verbal commitment of sharing gains or losses from securities investments in any form shall be effective.

This report may contain and/or describe/present factual historical information on prices of Futures contracts (the "information"). Please note that this information is solely for the purpose of forming part of the argument/grounds/evidence in our research methodology/analysis to support our conclusion on our view of the relevant industry/company mentioned. It does not, by any means (express or implied) to be associated with or constituted as SFC Type 5 Regulated Activity (Advising on futures contracts).

Disclosure of Interests

- 1) The research analyst(s) and his/her associate has not served as an officer of the company(ies) mentioned in this report, and does not have any financial interests in the company(ies) mentioned in this report.
- 2) GF Securities (Hong Kong) Brokerage and/or its affiliates does not have any financial interests and has not invested in interests aggregate to an amount equal to or more than (i) 1% of the market capitalization; or (ii) 1% of the issued share capital, or issued units, in the company(ies) mentioned in this report.
- 3) GF Securities (Hong Kong) Brokerage and/or its affiliates does not have any market making activities, and has not employed any individual(s) serving as officer(s) of the company(ies) mentioned in this report.
- 4) GF Securities (Hong Kong) Brokerage and/or its affiliates does not have any investment banking relationship in Hong Kong with the company(ies) mentioned in this report in the past 12 months.

Copyright © GF Securities (Hong Kong) Brokerage

Without the prior written consent obtained from GF Securities (Hong Kong) Brokerage, any part of the materials contained herein should not (i) in any forms be copied or reproduced or (ii) be re-disseminated.