

Meta's Compute-For-Sale Signal: Insurance, Not Retreat

- Meta Compute is a positive monetization strategy rather than an AI overcapacity.
- Compute-for-sale is a hedge that attacks the depreciation line, not a retreat.
- Multiple still ignores the pivot; we maintain BUY on META80, TP THB2.90.

One Headline, Two Stories

Bloomberg's 1 Jul report that Meta is building a unit to rent bare GPU capacity and sell hosted model access sent CoreWeave and Nebius sharply lower on an "overcapacity" read. We think the market pointed the gun the wrong way. The neoclouds losing a future competitor is one story; META gaining a second way for its capex to pay for itself is another, and only the first is bearish. Shares agreed — up almost 9% on the headline, then a further 4.7% on 9 Jul when Muse Spark 1.1, Meta's first paid model API, shipped. The reaction gap between META and the names that sold off is the thesis in miniature.

A Buyer Who Keeps Buying

Our disagreement with the bear case rests on one distinction: is compute-for-sale a confession or a hedge? We read it as a hedge. Zuckerberg has floated the cloud option since the 3Q25 call and never once slowed the buildout to back it up — CoreWeave's commitment grew to \$21b in April, Nebius signed up to \$27b in March. A company genuinely worried its buildout had overshot does not keep signing multi-year capacity while simultaneously building the arm to resell it. Meta is doing both, which reads as a company hedging how it monetizes capacity it was always going to build, not one retreating from the build.

What The Depreciation Buys

That hedge matters because of what the Street is actually pricing: depreciation, not demand. The multiple has stayed compressed all year on the D&A that lands once \$125-145b of 2026 capex goes into service — Evercore flags close to \$19b of 2027 operating income at risk from that alone. A monetization arm that converts idle GPU-hours into revenue attacks precisely that line item. It doesn't erase the depreciation, but it changes what the depreciation buys — from pure internal cost to a partially revenue-generating asset base. That reframing, not any single contract, is what the compressed multiple currently ignores.

The Clock Favors the Seller

The reason to act ahead of the crowd is timing. Meta reports Q2 on 29 Jul, the front edge of the three-to-six-month window Zuckerberg set himself on 2 Jul, and short-duration compute is renting at a premium — 2.6-4.0x a standard 5-year neocloud contract per SemiAnalysis. The sell-side has already built the framework before management has confirmed a word of it: Morgan Stanley sizes 250MW at \$40/W as ~8% to 2028 EPS, scaling toward 33% at 1GW. When the model is built and the price still assumes it won't happen, the asymmetry favors the long into the print.

Recommend a BUY with TP of THB2.90 (\$700, USD/THB at 33.3)

We maintain BUY on META80 with TP THB2.90 (\$700 at USD/THB 33.3). Some re-rating is already in — shares closed 9 Jul at \$631 off a \$520 low — but at roughly 19x consensus 2026E EPS of \$32.8, META still sits below its own ~23x five-year average. The multiple still assumes this pivot mostly doesn't happen. That is the setup we want: a confirmed anchor deal or explicit commentary on 29 Jul has room to re-rate the stock further, on optionality the market is currently pricing at close to zero.

Analyst

Suwat Sinsadok, CFA, FRM, ERP
suwat.s@globlex.co.th,
+662 687 7026


Assistant Analyst

Peerayu Sirivorawong

Meta's Compute-For-Sale Signal: Insurance, Not Retreat

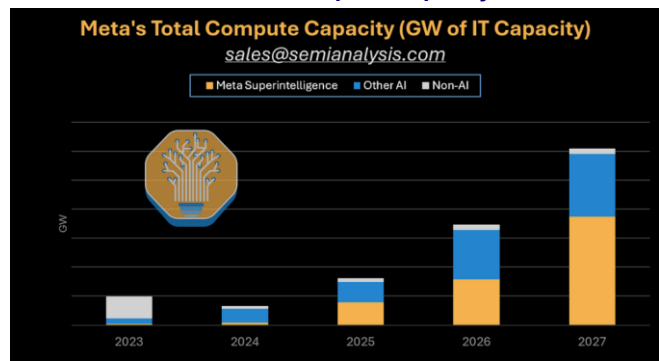
Bloomberg's 1 July 2026 report that Meta is building "Meta Compute" — a unit to rent bare GPU capacity and sell hosted model access — sent CoreWeave and Nebius sharply lower (both down more than 6% initially, extending into double digits within days) and reopened the "overcapacity" debate. We read the reaction as backwards for META itself. A company willing to monetize spare cycles is diversifying how its capex pays for itself, not admitting the buildout failed. Shares rose almost 9% on the same headline, then added a further 4.7% on 9 Jul when the first concrete piece of this shipped: Muse Spark 1.1, Meta's first paid model API.

Exhibit 1: Meta's Compute Deals per Quarter



Sources: Semi Analysis

Exhibit 2: Meta's Total Compute Capacity



Sources: Semi Analysis

Not New, Just Louder

Zuckerberg has said this before. He first flagged a cloud option on the Q3 2025 call, repeated it on the April Q1-26 call, and told the May AGM entering cloud was "definitely on the table" once internal build outpaces internal need. None of this has stopped Meta buying: CoreWeave's commitment grew to \$21b in April (on top of ~\$14b from 2025), and Nebius signed for up to \$27b in March. Meta stays one of neoclouds' largest customers even as it becomes their competitor — that dual role is the point, to us. Overbuilding for MSL while quietly building a second monetization leg reads as a hedge, not a confession that the primary bet failed.

The Margin Math the Street Is Pricing

2026 capex guidance now sits at \$125-145b, against \$72b spent in 2025; Morgan Stanley models \$175b for 2027E and \$205b for 2028E. FY2025 group operating margin was 41%, with the Family of Apps segment at 52% (down 2ppt y-y). Depreciation from this buildout is the swing factor once assets go into service — Evercore has flagged close to \$19b of 2027 operating income at risk from D&A alone, on the prior, lower capex guide. In our view, a monetization arm that turns idle GPU-hours into revenue is a direct, if partial, offset to precisely that worry.

Optionality, Not A White Flag

On 2 July 2026 Zuckerberg told staff agentic development "hasn't really accelerated" over the prior four months, with benefits still three to six months out. A week later Meta shipped Muse Spark 1.1 as its first paid API — strong on tool-use and orchestration benchmarks, but third on coding and multimodal versus Opus 4.8 and GPT-5.5 on Meta's own comparison table, and priced to undercut rather than to lead. We do not read this as MSL conceding the frontier race. Renting out spare capacity while a model sits mid-pack, the way xAI kept training Grok while renting out Colossus 1, buys optionality without foreclosing the upside case.

Scarce Compute, Premium Rent

Timing favors the seller. Short-duration capacity is fetching 2.6-4.0x the annualized per-GW revenue of a typical 5-year neocloud contract, per SemiAnalysis's tracked deals — against xAI's real-world comps of \$1.25b/month from Anthropic and \$920m/month from Google, both cancellable on 90 days' notice. SemiAnalysis itself now expects a Meta-Anthropic private-Claude arrangement near \$10b on similar terms, though neither side has confirmed it yet. One caveat on scale, in our view: Meta can only rent out its own first-party infrastructure, since CoreWeave and Nebius capacity carries consent-based subleasing restrictions — this is a bolt-on to the model, not a new core business line.

Why 29 July Is the Test

Meta reports 2Q26 results on 29 July 2026, consensus EPS \$7.17 on \$60.19b revenue — sitting at the front edge of the three-to-six-month window Zuckerberg set for himself on 2 July 2026. Morgan Stanley has already sized the mechanism ahead of any management confirmation: 250MW leased at \$40/W adds roughly \$2.97, or ~8%, to 2028 EPS, scaling toward 33% at 1GW. Sell-side is converging on the framework before Meta has said a word on pricing, scale, or timing — usually where the asymmetry sits, since confirmation is a low bar to clear when so little is priced in.

Maintain a BUY with TP of THB2.90 (\$700, USD/THB at 33.3)

Shares closed 9 Jul at \$631, up from a 52-week low of \$520, which already reflects some of this re-rating. At roughly 19x consensus 2026E EPS of \$32.8, META still trades below its own ~23x five-year average — a narrower discount than three weeks ago, but a discount all the same. We maintain BUY on META80 with TP THB2.90 (\$700 at USD/THB 33.3): the multiple still assumes this pivot mostly does not happen, so a confirmed anchor deal or explicit management commentary on 29 Jul has room to re-rate the stock further, even against a still-modest near-term revenue contribution.

GENERAL DISCLAIMER

Analyst Certification

Suwat Sinsadok, Register No. 020799, Globlex Securities Public Company Limited

The opinions and information presented in this report are those of the Globlex Securities Co. Ltd. Research Department. No representation or warranty in any form regarding the accuracy, completeness, correctness or fairness of opinions and information of this report is offered by Globlex Securities Co. Ltd. Globlex Securities Co. Ltd. Accepts no liability whatsoever for any loss arising from the use of this report or its contents. This report (in whole or in part) may not be reproduced or published without the express permission of Globlex Securities Co. Ltd.

RECOMMENDATION STRUCTURE

Stock Recommendations

Stock ratings are based on absolute upside or downside, which we define as $(\text{target price}^* - \text{current price}) / \text{current price}$.

- BUY:** Expected return of 10% or more over the next 12 months.
HOLD: Expected return between -10% and 10% over the next 12 months.
REDUCE: Expected return of -10% or worse over the next 12 months.

Unless otherwise specified, these recommendations are set with a 12-month horizon. Thus, it is possible that future price volatility may cause temporary mismatch between upside/downside for a stock based on market price and the formal recommendation.

* In most cases, the target price will equal the analyst's assessment of the current fair value of the stock. However, if the analyst doesn't think the market will reassess the stock over the specified time horizon due to a lack of events or catalysts, then the target price may differ from fair value. In most cases, therefore, our recommendation is an assessment of the mismatch between current market price and our assessment of current fair value.

Sector Recommendations

- Overweight:** The industry is expected to outperform the relevant primary market index over the next 12 months.
Neutral: The industry is expected to perform in line with the relevant primary market index over the next 12 months.
Underweight: The industry is expected to underperform the relevant primary market index over the next 12 months.

Country (Strategy) Recommendations

Overweight: Over the next 12 months, the analyst expects the market to score positively on two or more of the criteria used to determine market recommendations: index returns relative to the regional benchmark, index sharpe ratio relative to the regional benchmark and index returns relative to the market cost of equity.

Neutral: Over the next 12 months, the analyst expects the market to score positively on one of the criteria used to determine market recommendations: index returns relative to the regional benchmark, index sharpe ratio relative to the regional benchmark and index returns relative to the market cost of equity.

Underweight: Over the next 12 months, the analyst does not expect the market to score positively on any of the criteria used to determine market recommendations: index returns relative to the regional benchmark, index sharpe ratio relative to the regional benchmark and index returns relative to the market cost of equity.