on 22-Sep-2026

BERNSTEIN
SOCIETE GENERALE GROUP

22 September 2026

Global Software, US and China Internet, & US Communications Infrastructure

Cloud in the Quarter: How did the hyperscalers do in 2Q26?


Mark L. Moerdler, Ph.D.
+1 917 344 8506
mark.moerdler@bernsteinsg.com


Mark Shmulik
+1 917 344 8508
mark.shmulik@bernsteinsg.com


Robin Zhu
+852 2123 2659
robin.zhu@bernsteinsg.com


Madison Rezaei
+1 917 344 8622
madison.rezaei@bernsteinsg.com


Firoz Valliji, CFA
+1 917 344 8316
firoz.valliji@bernsteinsg.com


Shelly Tang, CFA
+1 917 344 8342
shelly.tang@bernsteinsg.com


Deeksha Pandey
+1 917 344 8447
deeksha.pandey@bernsteinsg.com


Wenhuan Chang
+1 917 344 8546
wenhuan.chang@bernsteinsg.com


Charles Gou
+852 2123 2618
charles.gou@bernsteinsg.com


Nancy Wu
+1 917 344 8545
nancy.wu@bernsteinsg.com

The hyperscale market has become a picks & shovels basket to support both traditional Cloud and the Generative AI wave even though it continues to be predominantly a CPU driven business. Investors remain concerned about the CAPEX growth and the associated margins / returns of AI infrastructure investment as AI becomes a more meaningful contributor to revenues. Suddenly, the laggard in hyperscale appears the leader in AI cloud with Google's vertically integrated solution now adding close to as muchas Microsoft in incremental cloud dollars added Q/Q and Oracle growing much faster off a smaller base than their peers.

Even with CAPEX stepping up, CSPs all point to capacity constraints, before it was GPU availability, but that has shifted to physical powered up datacenter capacity. All this ushers in a slew of new investor questions: Are Microsoft and Amazon going to see a return on CAPEX? Is speed and cost to bring capacity online the new basis of competition? Do the model wars determine winners and losers? What are the economics of Oracle's AI build-out, and can they rely on their OpenAI commitments? Are NVIDIA and OpenAI losers to Google or kingmakers? When will AI drive increased cloud IT budgets? Will the slowing of frontier model development impact hyperscaler growth and margins? Bernstein's Cloud in the Quarter note, which focuses on the largest hyperscale Cloud providers, should help investors get a better picture of the secular issues and why results may differ between providers.

The hyperscale market is the largest tangible market opportunity in Software / Cloud / Internet ($1.3-1.5T — see Diving into the Cloud TAM) and possibly the largest in tech. Bernstein's Global Hyperscale Cloud Team (Global Software, U.S. Internet, and China Internet and now US Comm. Infrastructure) supplies this quarterly note to help pull all the data together in one place for Amazon, Microsoft, Google, Alibaba, Oracle and now a neocloud - CoreWeave.

Updating our data through this quarter, today's note looks at the metrics each of the companies supply as well as our estimates through Q2 CY2026. The note also supplies our updated thoughts on the different companies' hyperscale businesses, where we are on AI, and how the markets and companies are progressing on what we believe is going to be a very long and valuable journey. This is the 22nd note in our series comparing and contrasting the hyperscale vendors. This many years of historical perspective should help investors better understand the trends and the opportunities.

To further help understand the long-term opportunities, we've published notes diving into the historical trends in IaaS and PaaS in the U.S. and globally in which we give our thoughts on the progression of the industry (see Link). We would also call attention to our notes that look at the capital intensity of IaaS/PaaS based on our proprietary work on capital intensity at Azure and AWS (What does it cost to offer IaaS / PaaS?) and the impact of AI driven CAPEX in Why Azure CAPEX will slow and revenue accelerate, Microsoft Azure AI margins, Oracle: The story is better than we thought. We also look at off balance sheet and long term leases and commitments in Microsoft: Are they overbuilding capacity?

On the Generative AI opportunity we have published a Blackbook: Scaling Intelligence: The Generative AI Handbook and numerous notes including: Shifting from Training to Inferencing impact on datacenters , 5 Takeaways on datacenter build-out.

See the Disclosure Appendix of this report for required disclosures, analyst certifications and other important information. Alternatively, visit our Global Research Disclosure Website.
www.bernsteinresearch.com
First Published: 22 Sep 2026 10:30 UTC Completion Date: 22 Sep 2026 10:30 UTC

on 22-Sep-2026

BERNSTEIN TICKER TABLE

Ticker	Rating	Cur	21 Sep 2026 Closing Price	Price Target	TTM Rel. Perf.	Adjusted EPS Cur	2025A	2026E	2027E	Adjusted P/E (x) 2025A	2026E	2027E
GOOGL (Alphabet)	M	USD	354.97	385.00	24.1%	USD	10.81	20.84	15.26	32.8	17.0	23.3
AMZN (Amazon)	O	USD	258.45	320.00	(3.0)%	USD	7.17	12.65	12.13	36.0	20.4	21.3
MSFT (Microsoft)	O	USD	501.61	660.00	(19.0)%	USD	17.28	19.96	24.01	29.0	25.1	20.9
ORCL (Oracle)	O	USD	148.56	325.00	(71.2)%	USD	7.63	8.37	12.19	19.5	17.7	12.2
BABA (Alibaba)	O	USD	115.75	165.00	(46.0)%	CNY	26.82	44.78	64.65	28.9	17.3	12.0
9988.HK (Alibaba)	O	HKD	114.90	161.00	(53.4)%	HKD	8.82	3.69	6.46	11.1	26.6	15.2
700.HK (Tencent)	O	HKD	450.60	760.00	(55.2)%	CNY	28.09	29.72	31.27	13.7	12.9	12.3
CRWV (CoreWeave)	U	USD	85.43	74.00	(52.4)%	USD	(1.20)	(4.56)	(4.77)	N/M	299	37.9
SPX			7,764.70									
ASIAX			1,983.89									

O - Outperform, M - Market-Perform, U - Underperform, NR - Not Rated, CS - Coverage Suspended

GOOGL, AMZN estimate is Reported EPS; GOOGL, AMZN valuation is Reported P/E (x); CRWV valuation is EV/EBIT (x); MSFT, ORCL, BABA base year is 2026;

Source: Bloomberg, Bernstein estimates and analysis.

INVESTMENT IMPLICATIONS

Microsoft (MSFT, TP $660, Outperform): Microsoft made a big bet on Azure, investing aggressively in a global footprint and technology innovation. Starting in ~CY2019 they doubled down on their existing AI investment as they prepared for and added Gen AI capabilities, built AI datacenters and invested in their own chips. With Azure and Copilot both accelerating growth, AI has become a meaningful contributor to top line growth and not a headwind. We also believe Microsoft is strongly positioned in the case of an AI bubble because their first-party usage of capacity is ramping thanks to their diversified app base, along with third-party contracts which comes from not only OpenAI but a diversified large enterprise customer base, and contract duration is well aligned with useful life of the equipment. While most of software is slowing, Azure is stable and will, we believe, accelerate supporting Microsoft's strong double-digit revenue growth and GAAP margins in the mid 40s. In addition, we are now starting to see substantial seat growth at Office 365 Copilot and usage at GitHub which both show that AI will over time drive not just GPU be increased CPU usage. While concerns continue surrounding the huge CAPEX investment across the hyperscalers Microsoft has slowed their CAPEX spend and seems to be growing CAPEX slower than Azure revenue, while accelerating Azure growth. Even concerns about the growth of long-term non-started datacenter leases are overblown (Microsoft: Are they overbuilding capacity?) as are future hardware commitments. Management has guided to postive FCF and no additional need for debt in FY27 which differentiates it from its US hyperscaler peers.

Oracle (ORCL, TP $325, Outperform): The Oracle story has shifted quickly from a potential Cloud and hyperscaler story to a major hyperscaler and AI training / inferencing provider. Oracle has $664B in Remaining Performance Obligations (RPO) and exceptional OCI growth ($166B in OCI revenue guidance for FY30), driven by multiple multi-billion dollar deals led by ~$300B in RPO from OpenAI. Based on the work we have done on understanding the value of this transition (Oracle: The story is better than we thought, with better upside potential), we believe it would be highly value accretive. While near term this might impact margins and cash, it will likely drive very substantial earnings and FCF growth over the next 5–10 years. The stock has taken a beating over the last year due to various concerns, but we see the overhangs getting gradually clearer as Oracle is securing most, if not all, the financing it will need for current demand by the end of FY27, sentiment around OpenAI continues to improve, and OCI growth demonstrates strong execution as seen in the past few quarters.

Amazon (AMZN, TP $320, Outperform): The quarter AWS bulls have been waiting for. Azure's print raised the bar heading into earnings, but AWS blew right through it, delivering 37% Y/Y growth that was a clearly visible ~9pts Q/Q growth inflection after spending the last year stuck in a 3-4pt acceleration range. It was also the fastest growth in 18 quarters, pushing AWS to the top of the hyperscaler leaderboard for net new dollars added Q/Q. AI is making encouraging progress; Amazon lifted its AI revenue run-rate to $25B from $15B just one quarter ago, while its custom silicon business reached a $25B run-rate, up from $20B last quarter. Add in a backlog that now stands at $496B (helped by the Anthropic deal earlier this year) and margins that reached 39% (aided by an energy credit but still expanding sequentially even excluding the benefit) -- everything is working. Another noteworthy takeaway was Amazon's unusually clear articulation of returns on AI investment. A significant portion of current investment is tied to data center shell buildouts, with spending later shifting toward servers that carry attractive

on 22-Sep-2026

payback profiles of 2-3 years tied to contracts that run 5-6 years. Looking ahead, we are confident on AWS fundamentals to take advantage of a multi-model, multi-product world with everything enterprises need under one roof including the oft-forgotten about non-AI hyperscaler world that's very much needed to deploy and scale AI solutions across organizations.

Google (GOOGL, TP $385, Market-Perform): Maybe an afterthought as the #3 player in the pre-AI cloud world, but Google is proving to be a much stronger competitor in AI cloud. Cloud revenue came in well above Street at +82% Y/Y, margins expanded to ~36%, and net dollars added Q/Q ran neck-and-neck with AWS and well ahead of Azure. The strength of Google's AI cloud is well understood: the only hyperscaler with true vertical integration across its own frontier model and custom TPU silicon (for which 3P sales started monetizing this quarter). But investor worries persist…there are delays on frontier model execution with Gemini 3.5 Pro. Cloud margins can be pressured as Google is increasingly forced to lean on third-party capacity to bridge ongoing compute constraints. And with capacity constraints come allocation challenges: Google needs compute to keep enterprise customers happy (lest they drift to AWS or Azure), developers happy (lest loyalty shifts elsewhere and pressures the token usage momentum Google has grown to highlight in its quarterly commentary), and of course it needs enough capacity internally to sustain AI momentum in Search (where let's be real...any compromise would create far bigger problems for the stock). Overall, long-term fundamentals remain hard to argue against as Cloud momentum is exceptional, backlog remains strong, third-party TPU sales are opening a new monetization vector, and Google's vertically integrated stack remains one of the most differentiated assets in AI. But there may be some clouds to watch in the upcoming months (pun intended).

Alibaba (BABA, TP US$165/HK$161, Outperform): Alicloud's revenue growth accelerated to 44.9% year-on-year in the June quarter, while adjusted EBITA margin improved to 11.6% from 9.1% in the prior quarter. AI-related products delivered triple-digit growth for a twelfth consecutive quarter and accounted for 35% of external cloud revenue. Management indicated that external cloud revenue growth is expected to continue accelerating beyond 50% in the coming quarters, with adjusted EBITA margins expected to improve steadily. ARR of MaaS revenue remains on track to exceed the RMB30bn Q4 FY3/27E target guided last quarter. Alicloud revenue growth should continue to accelerate in the coming quaters, as the Chinese market for compute remains tight.

Coreweave (CRWV, TP $74, Underperform): We are Underperform on CRWV with a price target of $74. CRWV benefits from supply constraint in the data center market, which undoubtedly continues to exist today; however, as capacity eases, which we anticipate, this company will likely be among the first and hardest hit. We value the company on a 25.5x EV/EBIT basis.

on 22-Sep-2026

DETAILS

The hyperscaler story has changed significantly over the last 2 years as AI has become the big driver of sentiment. The biggest debate to date, arguably, is how long can this build-out continue, and when will we have the proof of ROI?.

More recently adding to the concern was the call by Anthropic's CEO that frontier model innovation needs to slow to assure the models are safe (Link) which was supported by Sam Altman, Elon Musk and most recently Satya Nadella. This only increases investor concern around the sustainability and value of the massive AI investments to date and in the near future.

Over the past few years, we have seen massive CAPEX ramp (Exhibit 1), and that race does not seem to stop: cash CAPEX across the 4 hyperscalers reached roughly $150B in 2Q26, and with the guidances these companies have given, Bernstein is modeling $660B cash Capex across the 4 hyperscalers for 2026 (plus Meta it would be close to $800B). To fund this level of CAPEX all but Microsoft have gone to external funding sources this year, while Microsoft guides to FCF to stay positive for FY27 (They just reported 1H CY27).

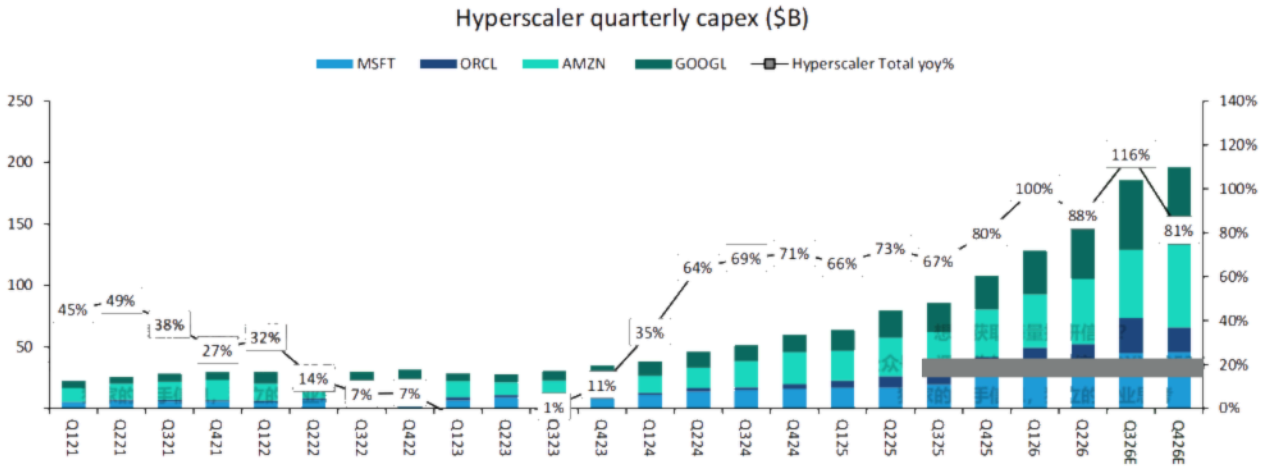
On the other hand, revenue (Exhibit 2) and backlog (Exhibit 3) have started to accelerate in the last few quarters, and despite the recent growth in capacity, demand continues to outstrip supply. Until the recent questions around whether AI model innovation should slow and thus AI adoption could slow, sentiment this quarter seemed to have improved.

In addition, Microsoft has broken away from the pack guiding to being FCF positive in FY27 while growing CAPEX, excluding the recent component price increases, slower than Azure revenue. We also recently published a deep dive on Microsoft's long-term commitments and showed that 1) Microsoft is not overbuilding; 2) Microsoft datacenter SQ FT are not growing different from the historical norms; and 3) Microsoft does not have long term hardware commitments beyond the current fiscal year.

To provide a broader perspective, especially on the conversation of CAPEX / Cloud capacity we have added our recently launched US Communications Infrastructure team lead by Madison Rezaei to this series. We recommend you read their work which focuses on data center capacity and the neoclouds.

Finally, at the same time as investors are worried about an AI bubble, they are also concerned on whether AI is going to eat enterprise software (Link) and now that slowing frontier model development could negatively impact AI adoption (Link). With the quantity, level and frequency of information flow in the space, who is an AI winner and who is a loser is changing quite frequently across tech and this can be seen even across the hyperscalers.

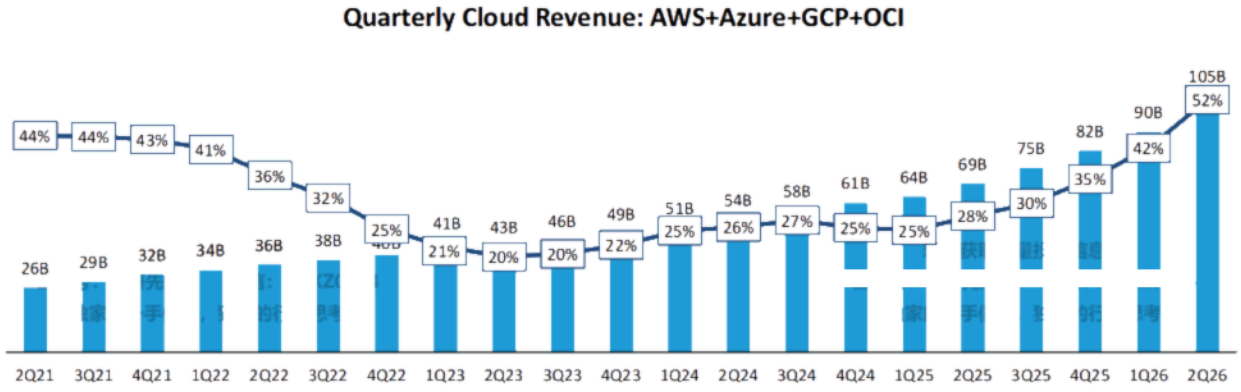
EXHIBIT 1: **Hyperscaler Capex has grown significantly and is expected to continue to grow**



Based on calendar quarters (Calendar Q1 = Microsoft and Oracle FQ3). Forward estimates based on Bernstein estimates
Source: Company disclosures, Bernstein analysis and estimates

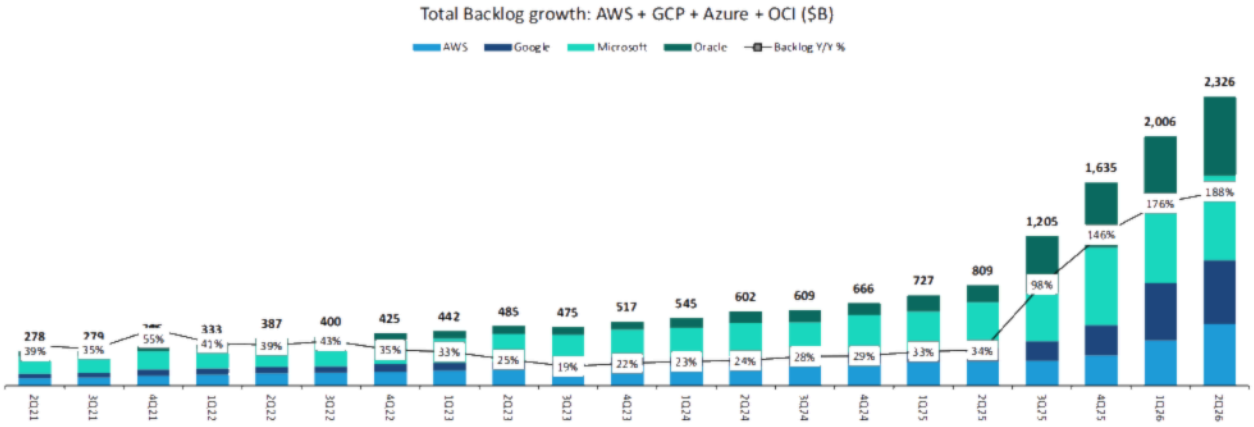
on 22-Sep-2026

EXHIBIT 2: **While cloud revenue growth has accelerated in the recent quarters, it needs to continue accelerating to keep up with the pace of capex growth**



Based on calendar quarters (Calendar Q1 = Microsoft and Oracle FQ3). Azure revenue based on Bernstein estimates
Source: Company disclosures, Bernstein analysis and estimates

EXHIBIT 3: **On the other hand, ignoring contract duration, total backlog (RPO) growth has accelerated significantly across all hyperscalers in recent quarters**

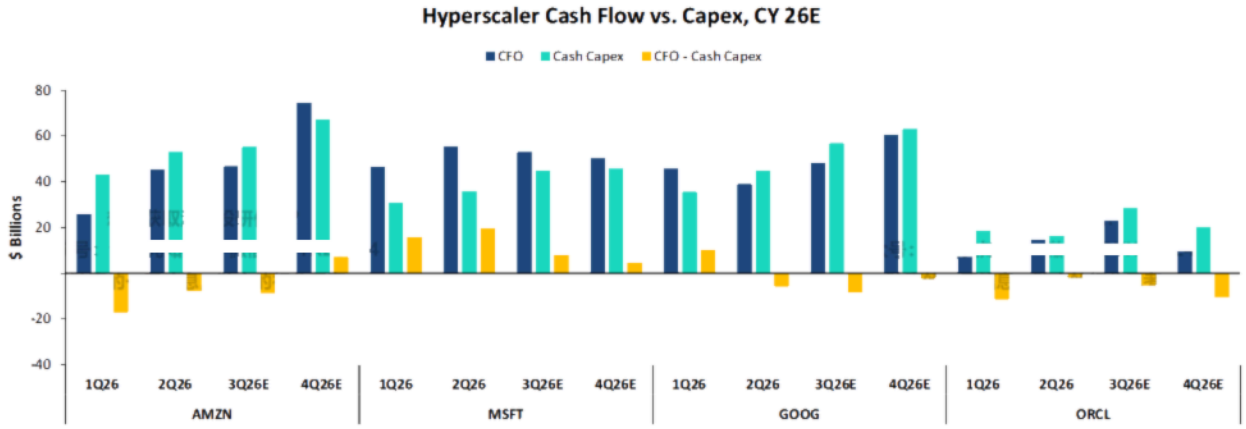


Based on calendar quarters (Calendar Q1 = Microsoft and Oracle FQ3).
Source: Company disclosures, Bernstein analysis

on 22-Sep-2026

EXHIBIT 4: **Meanwhile, as many hyperscalers are running out of internal funding, to increase Capex further many have been raising external funding**

Note that Microsoft has guided for FCF to be positive and no need to raise cash in FY27



Based on calendar quarters (Calendar Q1 = Microsoft and Oracle FQ3). Forward estimates based on Bernstein estimates
Source: Company disclosures, Bernstein analysis and estimates

DETAILS FROM LAST QUARTER'S RESULTS

Azure grew 43% in CC (43% reported) in 4FQ26, beating guidance of 39-40% CC. Management is guiding to 45% CC growth for next quarter and 2H CY26 is expected to accelerate as compared to the first half. CAPEX including finance leases in 4FQ26 was $41.0B with roughly two-thirds spent on short-lived assets (predominantly hardware). FQ4 CAPEX is guided to >$50B, inclusive of the lease reclassification impact from the useful life update. CY26 CAPEX guide is maintained at $190B before the lease reclassification, but with the reclassification impact (some financing lease become operating lease) it is now adjusted to $175B. For a better understanding of the drivers of all of Microsoft's CAPEX we recommend reading Microsoft: Could the bear case in fact be the bull case?. To understand that Microsoft is not overbuilding compared to demand and historical datacenter construction we recommend you read our note on their long term, including not yet commenced / "off balance sheet" commitments: Microsoft: Are they overbuilding capacity?

Please note that Microsoft has re-segmented their revenue (Quick Take: Microsoft's new reporting - is visibility increasing?) and focused the definition for Azure on IaaS/PaaS, removing GitHub Cloud and industry specific solutions. They will also be reporting Azure revenue each quarter. This will make Microsoft's Azure revenue definition the most pristine hyperscaler definition and number (for more on what others include see later in this note).

AWS growth accelerated to 37% Y/Y, the fastest pace in 18 quarters, with backlog surging to $496B and reinforcing the durability of demand. AI revenue reached a $25B run-rate (up sharply from $15B just last quarter) while Amazon's custom silicon business also surpassed a $25B run-rate (up from $20B last quarter). AWS is increasingly positioning itself as the control plane for enterprise AI through Bedrock, with AI acting as a multiplier for broader cloud consumption across storage, databases, networking, and security. Perhaps most importantly, Amazon offered one of the clearest ROI frameworks on AI spending. With server investments generating 2-3 year paybacks against 5-6 year contracts, management made a compelling case that today's elevated capex should translate into attractive long-term growth, margins, and cash flows.

Google Cloud remains one of the biggest AI beneficiaries, delivering 82% Y/Y growth against elevated expectations. Backlog climbed to $514B while third-party TPU monetization officially began this quarter. AI adoption continues to accelerate across the platform, with strong Gemini enterprise uptake, growing developer engagement, and rapidly expanding token volumes. Capacity remains a constraint, and Google is leaning on third-party infrastructure to bridge shortages, creating some margin pressure. Investors are also closely watching Google's ability to maintain frontier-model leadership in light of model delays.

Alibaba cloud reports revenue on a gross basis, including intra-Alibaba transfers. Starting this quarter, it also includes T-head, Alibaba's chip design business, although the contribution to revenue and profit remains small to date. On the newly reported basis, revenue grew 44.9% year-on-year in the June quarter, while EBITA margin improved 250bps sequentially to 11.6%. External revenue growth accelerated to c. 45% year-on-year from 40% in the March quarter, while internal revenue growth

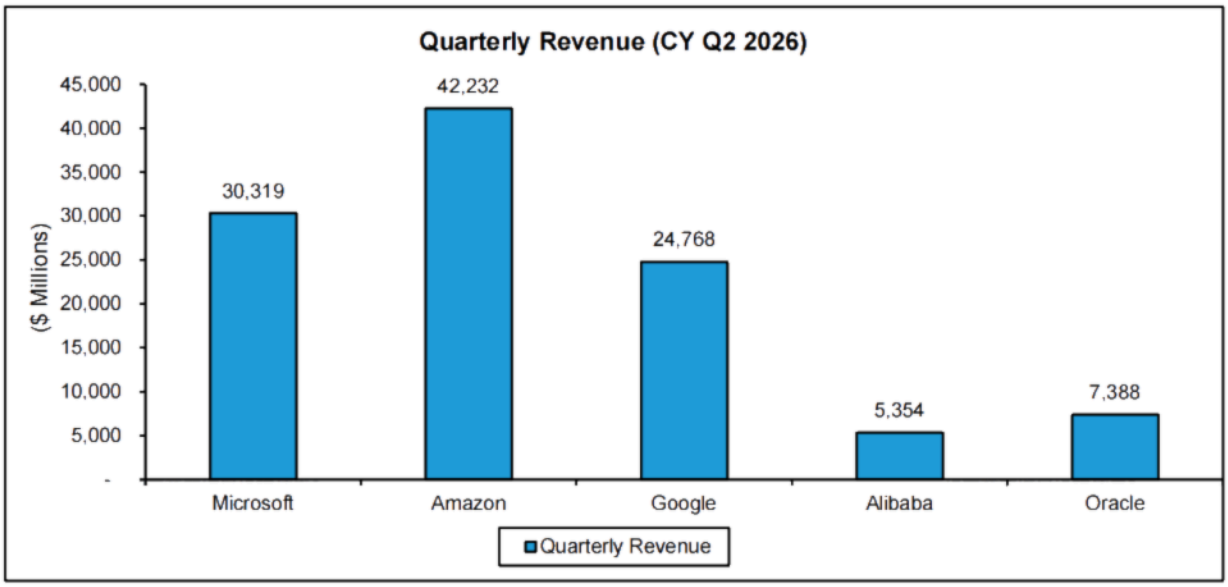
on 22-Sep-2026

also picked up to a similar pace. AI-related product revenue maintained triple-digit growth for a twelfth consecutive quarter and accounted for a higher, 35% share of external revenue. Capex reached RMB68bn this quarter, a large step-up from the RMB30bn quarterly run-rate we had modeled, signaling an inflection in investment as Alibaba ramps up compute capacity to meet growing AI demand, although the company did also warn investors not to extrapolate the most recent quarter's spend.

Oracle OCI continues to grow far faster than its larger peers and is still accelerating. In the recent 1FQ27, OCI revenue reached $7.4 billion, up 121% (120% in CC). Total RPO increased $209B Y/Y to $664B this quarter, up 46% Y/Y. Most of the sequential RPO increase last and this quarter was driven by prepaid/bring-your-own-hardware AI contracts, which requires minimal cash outlay from Oracle and will generate additional cash and FCF. CAPEX in 1FQ27 was $28.5B and $18B net of customer prepayments (cash outlay). For FY27 management reiterated guidance of $90-95B CAPEX, which includes $20-25B customer prepayments, resulting in $70B net cash outlay. This quarter, the company completed their $20B equity raise and has $20B remaining (debt or equity mix unknown) in this FY and has guided there will be no additional cash raised in FY27. Oracle also added +850MW of capacity this quarter and over the next 3 years, Oracle expects to bring 10GW of data center capacity online. We recently updated our 5-10 year Oracle scenario (Link) and our detailed breakdown of what comprises Oracle revenue (Link).

CoreWeave delivered a strong 2Q26, with revenue of $2.58B, up 112% Y/Y and 24% Q/Q, as the company continued to bring contracted GPU capacity online. Revenue backlog reached approximately $104B at quarter-end, up 246% Y/Y, and management disclosed more than $25B of net new customer commitments signed early in 3Q, reinforcing visibility into future revenue even as the exact timing of deployments remains a key variable. CoreWeave raised FY26 revenue guidance to $12.4B-$13.2B and adjusted operating income guidance to $960M-$1.15B, while also raising its CapEx outlook to $35B-$39B, underscoring both the strength of demand and the intensity of the capital required to address it. 2Q CapEx was approximately $9.4B, as CoreWeave continued to invest aggressively in GPUs, networking, and data center capacity; the resulting disconnect between near-term reported revenue and capital deployment remains central to the stock's debate. The company continues to benefit from an advantaged position in large-scale AI capacity, but investor attention is increasingly focused on backlog composition, customer concentration, delivery schedules, financing costs, and the cash-flow conversion of the rapidly expanding installed base.

EXHIBIT 5: **The five hyperscale providers report Cloud revenue, which includes not only IaaS/PaaS but also other revenues**



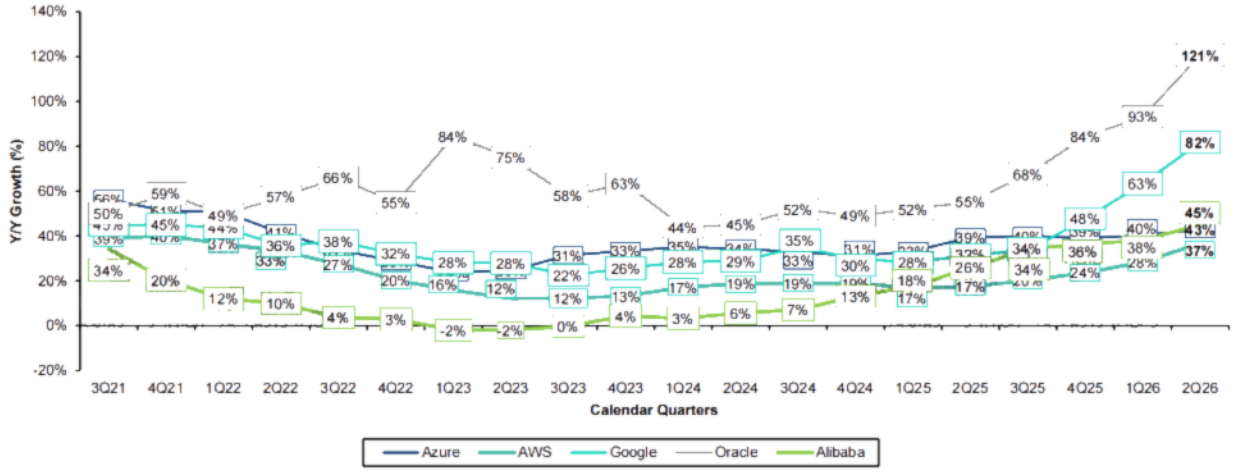
Google includes Workspace. Microsoft Azure historical revenue is based on our estimates. Management has definition for Azure and is now giving Azure numbers. This chart is based on the old reporting.

Please note that Alibaba no longer provides cloud revenue excluding internal transfers. Our revenue number is an estimate based on historical d
Source: Company reports, Bernstein estimates and analysis

人在草木间

on 22-Sep-2026

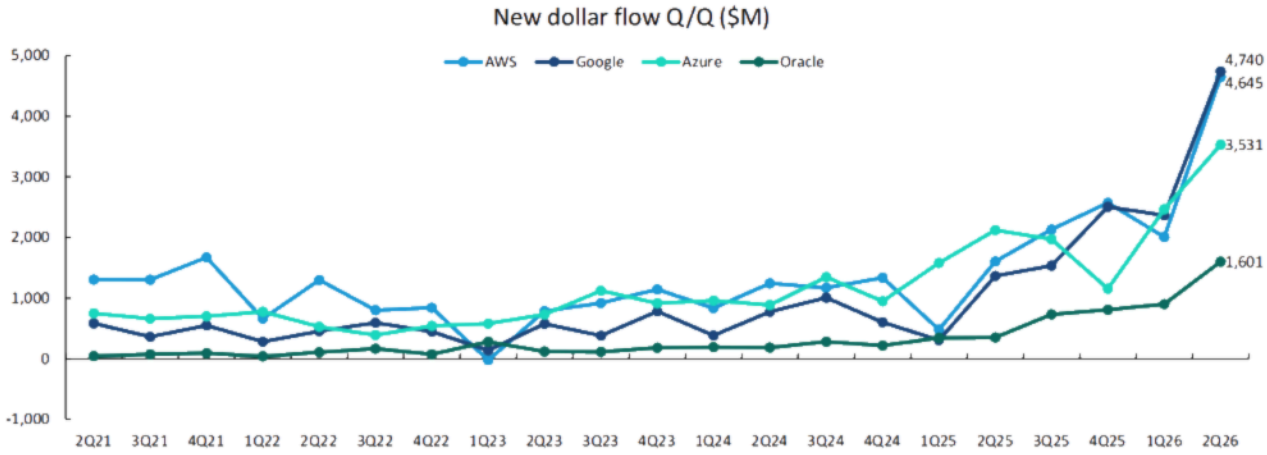
EXHIBIT 6: **In Q2 CY26, Microsoft Azure grew 43% Y/Y, Amazon's AWS grew 37% Y/Y, Google's Cloud grew 82% Y/Y, Alibaba grew 45% Y/Y and Oracle grew 121%**



Microsoft Azure historical revenue is based on our estimates. Management has definition for Azure and is now giving Azure numbers. This chart is based on the old reporting.

Please note that Alibaba no longer provides cloud revenue excluding internal transfers. Our revenue number is an estimate based on historical data.
Source: Company reports, Bernstein analysis and estimates

EXHIBIT 7: **On a Q/Q basis, AWS added $4.6B vs Azure at $3.5B vs Google Cloud at $4.7B vs Oracle at $1.6B**

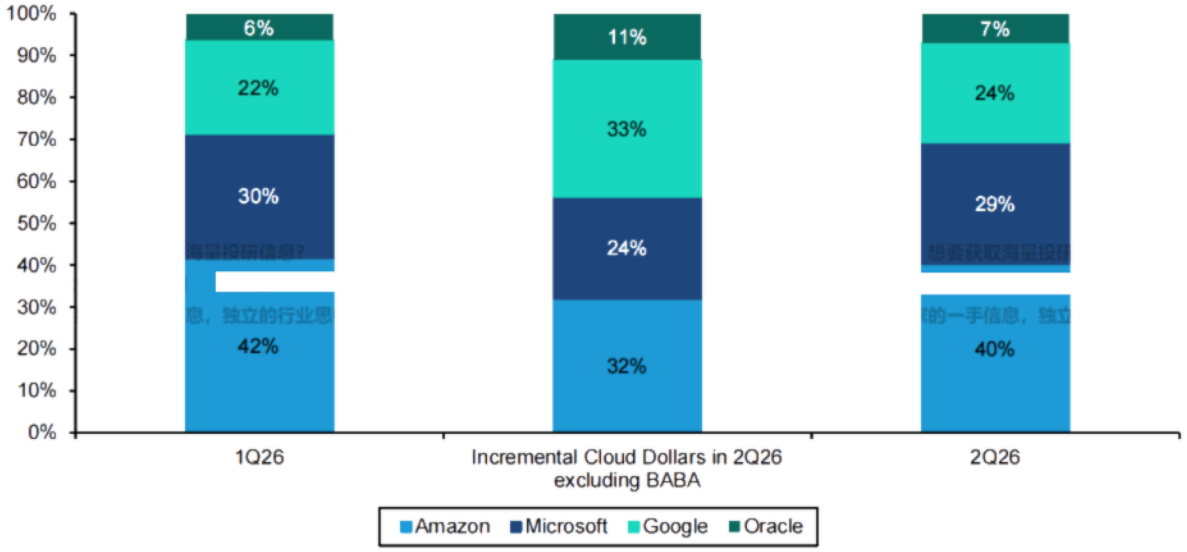


Microsoft Azure historical revenue is based on our estimates. Management has definition for Azure and is now giving Azure numbers. This chart is based on the old reporting.

Please note that Alibaba no longer provides cloud revenue excluding internal transfers. Our revenue number is an estimate based on historical data.
Source: Company reports, Bernstein estimates and analysis

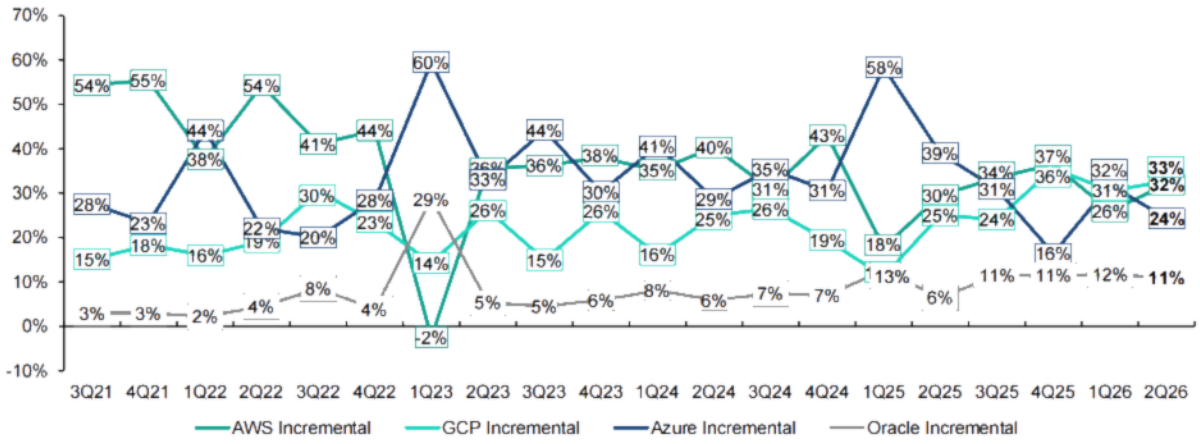
on 22-Sep-2026

EXHIBIT 8: **Of the incremental dollars added Q/Q by the hyperscale Cloud providers (excluding BABA), Google and Oracle have gained incremental shares this quarter.**



Microsoft Azure numbers are based on estimates
Source: Company reports, Bernstein analysis and estimates

EXHIBIT 9: **Q over Q, incremental share gained by each hyperscaler (excluding Alibaba)**



All quarters represent calendar quarters, Google's cloud revenues include both GCP and Workspace.

Microsoft Azure numbers are based on estimates.
Source: Company reports, Bernstein estimates and analysis

With that context in mind, we provide some highlights from each of the five companies' results last quarter.

HIGHLIGHTS FROM THE QUARTER

Amazon AWS

- **Another quarter of AWS acceleration, even as scale continues to grow.** The expectations bar keeps moving higher, and AWS keeps clearing it. Revenue increased 37% Y/Y to an annualized run-rate of approximately $169B, representing a roughly 900 bps acceleration from 28% growth in 1Q26. Looking longer-term, AWS now believes it has a **credible path toward becoming a $1T annual revenue business** while maintaining attractive margins and returns on invested capital.

on 22-Sep-2026

- **Leader in Q/Q net dollar adds**. AWS added approximately $4.6B of revenue sequentially, ahead of the ~$3.5B Q/Q increase at Azure and comparable to Google's roughly $4.7B increase (although with GCP benefiting from likely $1-1.5B of TPU system sales).

- **Backlog growth continues to reinforce the durability of demand.** AWS backlog expanded to $496B from $364B in 1Q26, representing triple-digit Y/Y growth. While the increase benefited from large AI commitments, including the ~$100B Anthropic agreement signed in April, demand remains broad-based across both AI and traditional cloud workloads. Importantly, AWS indicated that much of its 2027 capacity is already committed, with orders extending into 2028.

- **AI is scaling rapidly and increasingly driving the entire AWS platform.** AI revenue surpassed a $25B annualized run-rate, growing triple digits Y/Y and increasing significantly Q/Q. The pace of expansion is particularly noteworthy given Andy Jassy disclosed just one quarter ago that AWS's AI business had reached a $15B+ run-rate. Growing from a $15B+ to $25B + ARR in a matter of months offers optimism on the exceptional pace of AI adoption.

  - **Bedrock progressing as the control plane for enterprise AI.** AWS continued expanding Bedrock's model ecosystem, adding support for leading models from OpenAI, Anthropic, Google, and xAI while further building out agentic capabilities through Bedrock AgentCore. Adoption remains robust, with more customers added during the past six months than during the platform's first two years, and quarterly spending exceeding all prior quarters combined. AWS's messaging increasingly suggests that the next phase of AI adoption will be won not by model access alone, but by helping enterprises deploy, govern, orchestrate, and scale AI systems in production environments.

  - **And more importantly, AI is a multiplier for the whole ecosystem.** There is an increasingly symbiotic relationship between AI and core cloud adoption. Customers scaling AI workloads are simultaneously increasing spend on storage, databases, networking, analytics, and security services. AI can also help accelerate cloud migrations.

  - **Data gravity remains one of AWS's most important advantages in the AI era.** As inference workloads move into production, enterprises increasingly prefer to run AI applications alongside their existing data and mission-critical applications. Given that more enterprise workloads and data already reside on AWS than any other platform, management believes AWS benefits from a powerful structural advantage.

- **Custom silicon is a strategic differentiator.** Amazon's chips business surpassed a $25B annualized revenue run-rate, growing triple digits Y/Y and accelerating from the roughly $20B run-rate announced just last quarter.

  - **Trainium adoption is broadening beyond frontier AI labs.** Sure, Anthropic and OpenAI remain anchor customers, each making multi-year, multi-gigawatt commitments to Trainium infrastructure. However, AWS is increasingly seeing adoption across AI-native startups, unicorns, and large enterprises including Uber and Pinterest. Customer demand for Trainium has become strong enough that AWS is open to selling chips outside its cloud platform in the future.

  - **AI may ultimately become as much a CPU story as a GPU story.** Many of the fastest-growing AI workloads, including reinforcement learning, post-training, memory management, agentic workflows, and tool use, are heavily CPU intensive. This creates a significant tailwind for Graviton, which offers 30%-40% better price-performance than competing alternatives. Adoption remains exceptionally broad, with 98% of AWS's top 1,000 EC2 customers now using Graviton, revenue commitments nearly tripling Q/Q, and Graviton5 growing nearly twice as fast as Graviton4.

- **Amazon laid out one of its clearest frameworks yet** for why today's elevated AI spending should generate highly attractive long-term returns.

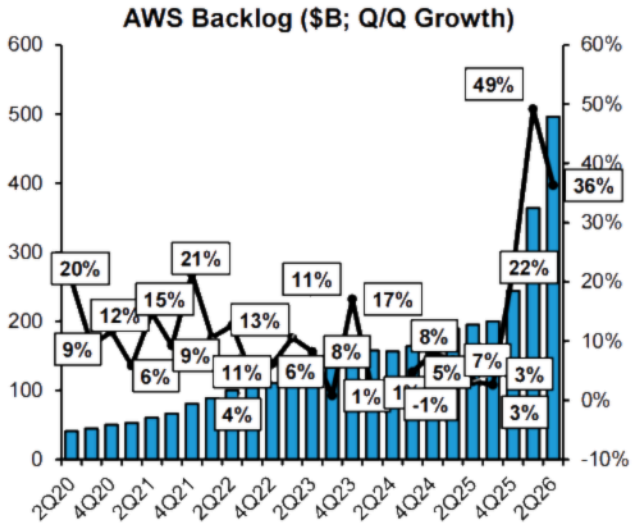
  - Data centers require investment roughly two years before monetization but can generate revenue for more than 30 years once operational.

on 22-Sep-2026

- Servers and networking equipment, by contrast, are typically purchased only months before deployment, allowing AWS to maintain strong visibility into customer demand before capital is committed. These investments generally break even in less than three years while carrying useful lives of five to six years. With most AI infrastructure contracted on five-year or longer commitments, AWS expects to generate meaningful cash flow for years after reaching payback.

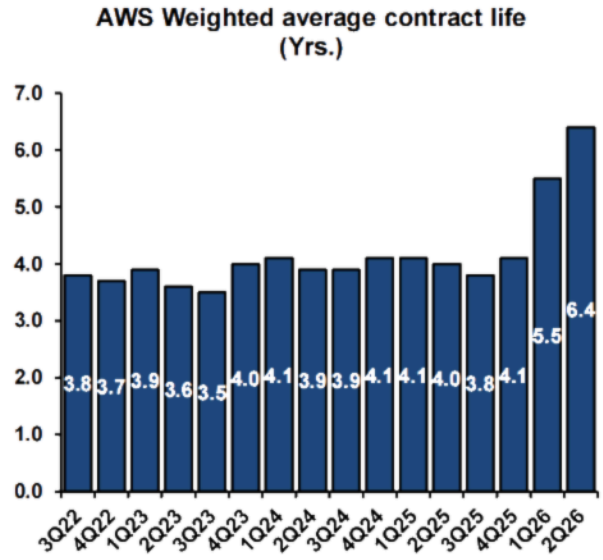
- Just as importantly, each successive generation of infrastructure should carry better economics as new compute capacity can be deployed into existing facilities without repeating the largest portion of the upfront data center investment.

EXHIBIT 10: **AWS Backlog grew to $496B up ~50% from the $364B announced in 1Q26**



Source: Company reports, Bernstein analysis

EXHIBIT 11: **AWS average contract life rose Q/Q to 6.4 years in 2Q26**



Source: Company reports, Bernstein analysis

Microsoft Azure

- **Microsoft Azure** (Exhibit 12)**,** now predominantly PaaS/IaaS and will be purely PaaS/IaaS following the FY27 re-segmentation (link), grew 43% as reported and in CC this quarter, beating consensus by 200bps and guidance by 300bps. Management is guiding to 45% CC growth for next quarter, and continues to expect acceleration in 2H CY26 as compared to first half. This quarter, growth was ahead of management and street expectations, driven by efficiency gains across CPU/GPU fleet as well as process improvements to enable earlier delivery of new capacity. The additional incremental capacity was quickly monetized. Stronger than expected GitHub Copilot consumption following the June business model change also contributed to the results. Again, demand continues to exceed available capacity.

  - Management continues to take the long view in a supply constrained world, prioritizing 1st party applications (Copilots) and R&D while balancing strong customer demand for Azure.

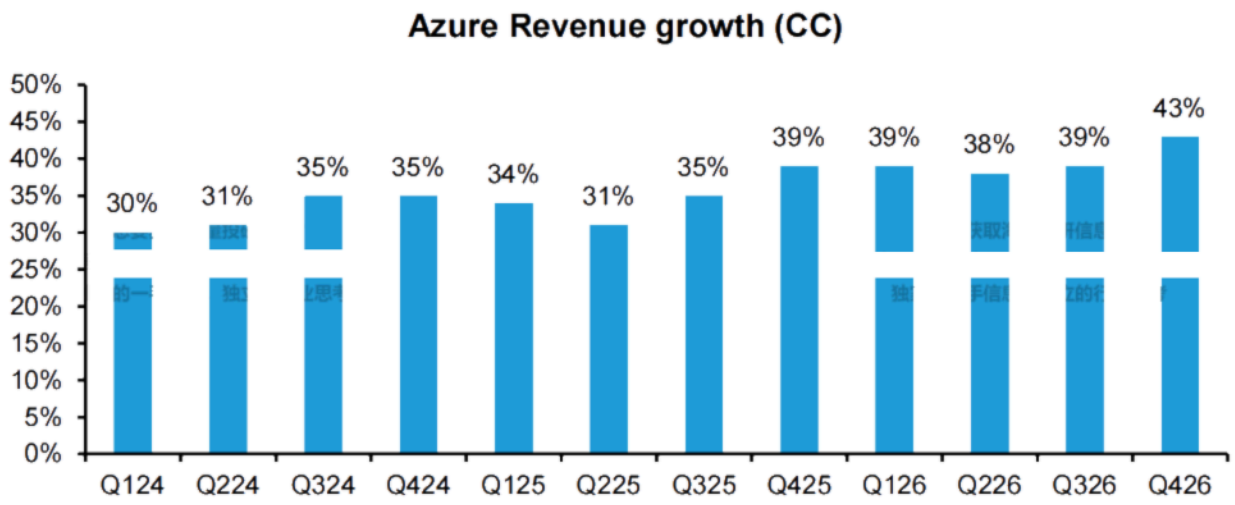
  - Microsoft did not update AI ARR this quarter, but Azure has accelerated and paid Copilot seats increased 50% QoQ to 30M seats. As a reminder, AI ARR reached $37B in 3Q26.

- **Commercial RPO** reached $678B this quarter, growing 84% YoY. All QoQ growth came from enterprise customers and excluded any new commitments from frontier model companies (including OpenAI and Anthropic) and RPO grew 25% YoY excluding OpenAI. Including OpenAI, RPO has a weighted average contract duration is 2.3 years. Roughly 30% will be recognized over the next 12 months, growing 37%. The remainder is recognized > 12 months and grew 112%. We understand that AI contracts are typically longer duration (6 years and in some cases longer), but we note management is not including all the committed revenue especially in outer years for risk management purposes ("net of reserves").

on 22-Sep-2026

- **CAPEX** (Exhibit 13) including finance leases this quarter was $41.0B, up 70% Y/Y. Cash paid for PP&E was $35.8B. Roughly two-thirds of CAPEX was spent on short-lived assets including GPUs and CPUs. Management expects FQ1 total CAPEX to exceed $50B, inclusive of the impact of lease reclassification (see below). Management reiterated the CY26 guide of $190B ($175B after lease reclassification adjustment) and expect FY27 Capex to grow Y/Y.
  - **Management expects to remain free cash flow positive in FY27**. While many of Microsoft's hyperscaler peers have turned FCF negative and some have raised debt and/or sold equity, Microsoft has set the tone that there is a limit to how large cash CAPEX will be next year and how their cash generating capabilities will be strong.
  - **Increasing the useful life of datacenters**: Microsoft has increased the useful life of their datacenters from 15 years to 25 years. This allows the company to recognize some shorter duration financing leases as operating leases, decreasing the "Capex including finance lease" figure, and lower slightly operating and free cash flow. It also should allow Microsoft to amortize the datacenters they build over a longer period of time. The decrease in the company's CAPEX from 190B to 175B was driven by this change in the useful life and not by a change in the quantity or unit cost of hardware, equipment or datacenter facilities. Without this change management was reiterating their FY guidance.
  - Note that higher CAPEX is going to translate into increasing D&A which will incrementally impact margins. However, we continue to believe that the Cloud revenue strength and Microsoft's tight cost control will help offset at least part of this and with CAPEX growth slowing the headwind will likely abate. We deep dived into Azure margins here: Microsoft: Azure AI margins, and its implications for next 5 years
  - CAPEX spending has shifted this year from 60% Land, Buildings and Improvements to 2/3 spending on servers. We see this as a good sign that the rebuilding of datacenters for mixed use (CPU and GPU) is likely decreasing and datacenter spending is now more focused on increasing demand. We note that the company is also leveraging datacenter leasing and shorter term Neocloud contracts to bring capacity on-line quicker. We also believe that as the shift continues from long life assets to shorter term assets (6 years), CAPEX will likely grow much more inline with Cloud revenue growth and then slower than revenue.
  - Many investors have been concerned about the large increase in long-term, not yet started, leases but in fact, if one layers the start of new leases over the next 7 years (as per the disclosures) and use a 15+ year average contract duration (not started leases are from 1 to 20 years) then in fact the growth is in-line with historical growth rates in the high teens (Link).
  - In addition, Microsoft has commitment to procuring hardware including power and cooling systems over FY27 but has very little commitments beyond that which answers the oft asked question if the company is locked into hardware additions they may not be able to utilize (Link).

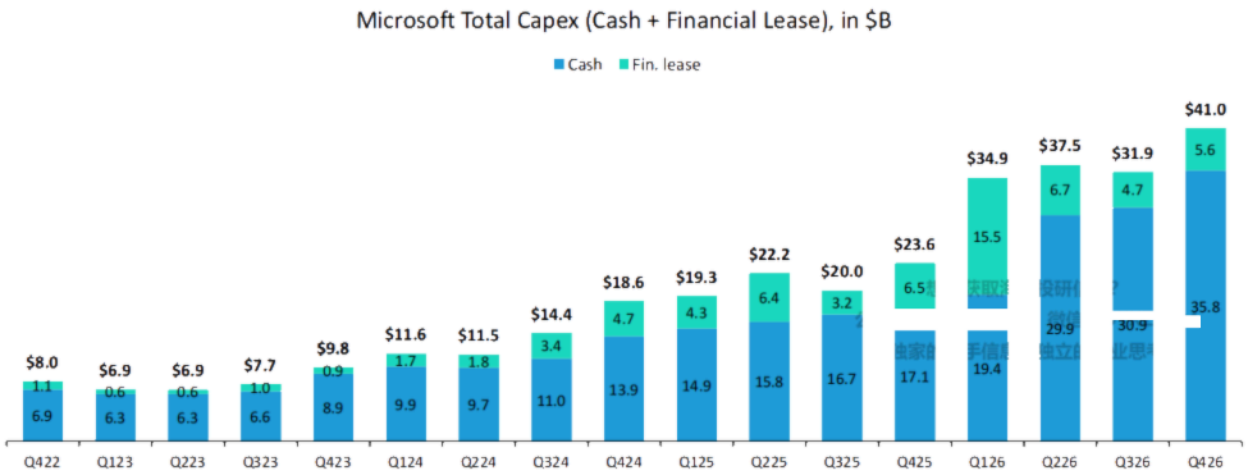
EXHIBIT 12: **Microsoft Azure constant currency growth**



Source: Company disclosures, Bernstein analysis

on 22-Sep-2026

EXHIBIT 13: **CAPEX has gone up as Microsoft continues to invest in AI to meet demand**



Source: Company reports, Bernstein analysis

Google Cloud

- **Google Cloud continues to outpace rising expectations**. Cloud revenue grew 82% Y/Y to $24.8B, driven by strength across GCP, enterprise AI solutions, enterprise AI infrastructure, and core cloud services. Importantly, management noted that Cloud growth would have accelerated meaningfully even excluding the contribution from TPU system sales that began shipping this quarter.

  - Backlog remains one of the strongest proof points. After backlog nearly doubled sequentially last quarter to roughly $462B, Google added another $50B+ of backlog in Q2, bringing total Cloud backlog to $514B. Management attributed the increase primarily to strong demand for enterprise AI offerings. Notably, the vast majority of backlog remains traditional GCP contracts across a broad customer base, and Google expects to recognize just over 50% of backlog as revenue over the next 24 months.

- **TPU monetization begins.** Q2 marked the first quarter in which Google recognized revenue from TPU system sales, although management emphasized that the contribution remains relatively small today. Existing TPU agreements are already reflected within Cloud backlog, with the majority of associated revenue expected to be recognized during 2027 as deployments ramp.

- **Enterprise AI adoption remains broad-based**. Google highlighted continued momentum for Gemini Enterprise, with nearly 90% of the Fortune 100 now using the platform. Adoption is increasingly translating into meaningful usage, with nearly 500 Cloud customers processing more than 1T tokens over the past year and more than 2,000 enterprises consuming over 100B tokens. Growth continues to be driven by a combination of new customer wins, deeper deployments among existing customers, and expanding partner activity through Google Cloud Marketplace.

  - Security highlighted as a growing pillar of the Cloud story. Google Cloud Security is now used by 90% of Fortune 100 companies, nearly 90% of Wiz customers utilize AI-powered security features, and the number of AI workloads scanned and protected by Google's security platform increased more than 45% Q/Q. The company continues to position security as a differentiated full-stack offering that combines threat intelligence, AI-powered monitoring, automated scanning, vulnerability remediation, and Wiz integration.

  - Underlying demand for Google's models continues to ramp. More than 9M developers are building with Google's models each month across APIs and developer platforms, while model APIs are now processing approximately 22B tokens per minute, up from 16B last quarter and 10B two quarters ago.

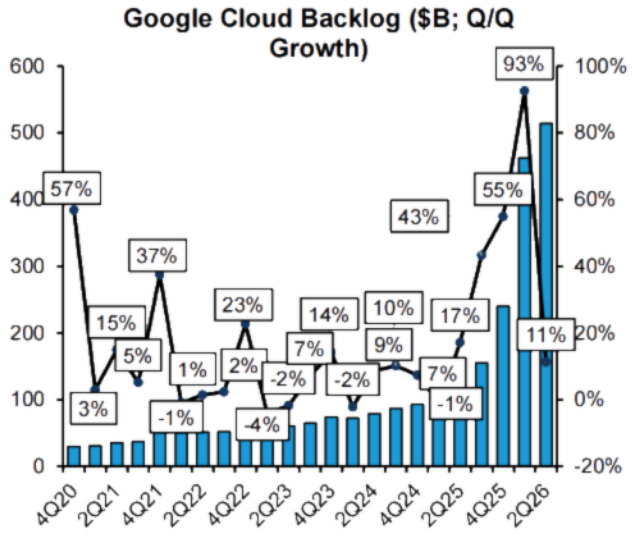
on 22-Sep-2026

- **Investor scrutiny around the frontier increases**. Investor questions increasingly focused on delays surrounding Gemini 3.5 Pro, particularly around coding performance and Google's competitive position relative to other frontier labs. While Google confirmed Gemini 3.5 Pro remains in testing, management highlighted Gemini 4 and called out that it has started its most ambitious pre-training run yet for Gemini 4.

  - In the meantime, Google continues to iterate aggressively across the Flash family of models. The company launched Gemini 3.6 Flash, Gemini 3.5 Flash-Lite, and Gemini 3.5 Flash Cyber during the quarter, emphasizing strong customer demand driven by their performance, latency, and cost advantages. Management also highlighted Flash Cyber's ability to identify and remediate vulnerabilities while delivering performance comparable to significantly larger cybersecurity models.

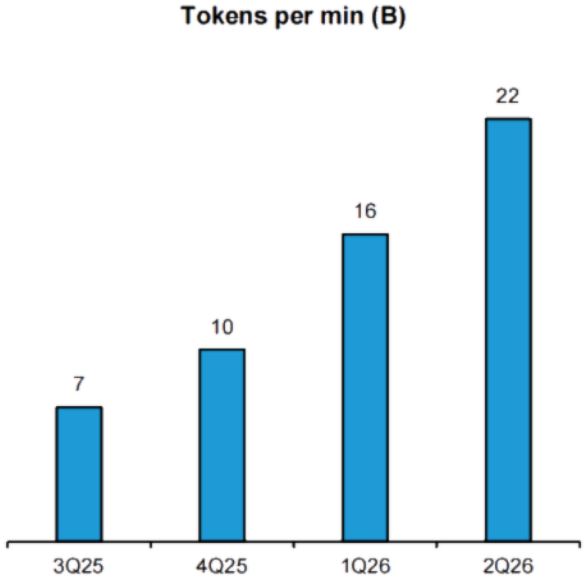
- **Capacity remains the key bottleneck.** Google remains capacity constrained and continues to see demand outpace available capacity. To address near-term demand, the company plans to utilize additional third-party infrastructure capacity in Q3 while expanding internal capacity. While this will create modest pressure on Cloud margins, management framed it as a temporary bridge strategy intended to support customer growth and capture long-term opportunities.

EXHIBIT 14: **Google Cloud's Backlog rose to $514B in 2Q26**



Source: Company reports, Bernstein analysis

EXHIBIT 15: **Token processing grew to 22B token per min**



Source: Company reports; Bernstein analysis

Alibaba

- **Alibaba's cloud revenue** is reported on a gross basis, including intra-company transfers and, from this quarter, Alibaba's chip design business. On this basis, revenue grew 44.9% year on year to RMB48.4bn in the June quarter (Q1 FY3/26: RMB33.4bn). Adjusted EBITA margin improved to 11.6%, up 280bps from 8.8% a year earlier. As a result, Alicloud EBITA rose 132.7% year on year to RMB5.6bn.

  - **Both external and internal revenue growth accelerated.** Alicloud's reported revenue includes an estimated c. 25% contribution from internal workloads. According to Alibaba, AI-related product revenue delivered triple-digit growth for the twelfth consecutive quarter and accounted for c. 35% of external cloud revenue.

  - **ARR of AI-related product revenue reached RMB49.5bn.** Alibaba updated its AI-related product revenue ARR, which increased from RMB36bn last quarter to RMB49.5bn, and is projected to exceed $10bn next quarter. Within this total, ARR from MaaS and AI-native software subscriptions reached RMB16bn as of August and remains on track to exceed previous guidance of RMB30bn by FY3/27E. The company expects AI revenue growth to accelerate further in the coming quarters.

on 22-Sep-2026

- **Incremental cash margins now in the 50% range.** Adding back the group-level delta in depreciation and amortization to Cloud EBITA this year and last implies an incremental cloud EBITDA margin of more than 50%. Management attributed the margin improvement to economies of scale and pricing power amid tight AI compute supply - the latter referring to a series of price increases on compute services during the quarter. The company expects cloud margins to continue improving in the coming quarters as compute supply remains constrained and AI-related revenues carry higher margins.

- **Capex.** Alibaba reported RMB67.6bn of capex at the group level in the June quarter, a meaningful step-up from the RMB30bn quarterly run-rate implied by its previous guidance of RMB380bn in capex over three years. While management downplayed questions whether the RMB380bn target would be raised and cautioned investors against annualising this quarter's figure, we believe the increase reflects a stronger commitment to expanding compute capacity and investing in AI cloud and MaaS, as Alibaba expects cloud revenue to surpass that from e-commerce in five years.

- **T-head.** From this quarter, Alibaba has consolidated its chip design business under cloud reporting. More than 500k previous-generation chips have been manufactured and shipped, serving over 650 Alicloud customers. The latest-generation M890 chip has been commercially deployed on Alicloud and supports inference for models with more than 2tn parameters, and management expects supply to ramp up in H2. The company will also begin development of its Gen-2 chip, which is expected to offer significantly higher compute performance and enhanced interconnect bandwidth.

Oracle

- **Oracle OCI** continues to grow far faster than its larger peers and is still accelerating. In 1FQ27, OCI revenue reached $7.4 billion, up 121% (120% in CC).

  - To get a full view of the impact of the AI datacenter build out on Oracle's financials we have created a 5-10 year view which we updated frequently (Oracle: The story is better than we thought, with better upside potential). We recommend that anyone trying to understand the long term value creation of AI IaaS/PaaS should read this and our associated notes.

- **RPO growth with no additional CAPEX**: Total RPO increased $209B Y/Y to $664B this quarter, up 46% Y/Y. Most of the increase in RPO in recent quarters were driven by customer-prepaid or bring-your-own-hardware AI contracts and requires minimal to no additional funding needs from the company. cRPO represents approx. 13% of RPO. Moreover, management confirmed that the margin profile on prepaid / bring your own hardware contracts are similar and possibly better than the regular AI contracts.

- **CAPEX:** To support this level of OCI revenue growth requires CAPEX. CAPEX in 1FQ27 was $28.5B and $18B net of customer prepayments (cash outlay). For FY27 company reiterated guidance of $90-95B CAPEX, which includes $20-25B customer prepayments, resulting in $70B net cash outlay.

  - Oracle does not build datacenters and the CAPEX will be mostly from hardware (which could include power and cooling systems).

  - The multiple ways Oracle is considering to finance include 1) debt, while maintaining IG rating; 2) equity; 2) vendor financing 3) customers bringing their own hardware.

  - Oracle started to present the net cash outlay metric last quarter (4QF26) to help illustrate cash use from the company's perspective and address some investor concern on the size of the funding needs.

  - Cash requirements and CAPEX will, based on the current fiscal guide, peak in FY28.

  - We hope to gain more information on CAPEX and cash requirements at the upcoming Financial Analyst Day in October.

- **Margin impact of component cost increases:** One of the concerns investors had going into the results was that with the recently announced pricing increases on memory and other components would negatively impact Oracle's OCI margins. In fact, management reiterated on the earnings call that the increase in component pricing will only increase Oracle revenue (and price drops could decrease revenue) but that margins will be stable even if the costs / revenue increases for AI compute. Oracle has fixed priced contracts (for OCI these are the contracts going live shortly and where the hardware is already in the Oracle manufacturing process we believe) but for contracts with longer time durations or where the pricing could vary,

on 22-Sep-2026

Oracle has conditions in their contracts that allow them to pass along those costs. In other words, Oracle's gross profit dollars will not be meaningfully impacted by price increases.

- **Rental of older GPUs are generating 20% more revenue:** Management explained that this quarter a range of older GPUs (3-4 years old) that came up for contract renewal or available for rental to other parties capturing on average 20% more revenue for Oracle.

- **$20B equity raise completed**: Oracle raised $43B in debt and $5B in equity financing in FY26. This quarter (1Q27), Oracle completed $20B equity raise. Of the previously announced fundraising plans, Oracle has $20B remaining which could be a mix of debt and equity (no details yet). No additional debt will be raised in FY27.

- **AI datacenter capacity**: Oracle added +850MW of capacity this quarter and has over 2GW of capacity now. Management stressed that datacenter deliver is accelerating. Oracle has secured more than 10GW of data center capacity (SQ FT plus power) that will likely come online over the next three years, with a vast majority of the datacenter capacity already contracted out. Note that the hardware is not contracted out and is instead bought and assembled prior to the datacenter going live.

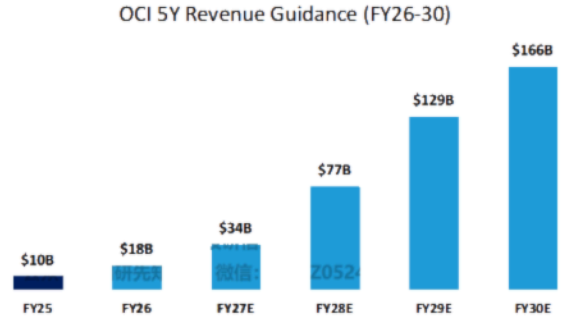
- **Cloud Database:** Oracle's OCI revenue includes the Oracle Cloud database which is a fast-growing business Cloud database. Cloud database grew 26% YoY to $0.9B in 1FQ27 with Multicloud Database revenue up 353% Y/Y and multi-cloud customers grew 180% YoY.

  - Multi-cloud partnerships are very important to Oracle: from a revenue standpoint, the multi-cloud agreements enable customers to run the Oracle database in their preferred cloud and the majority of the database revenue would come to Oracle. From a margin perspective, Oracle is expected to achieve high margins (potentially SaaS level margins) on the database as Azure/GCP/AWS are covering the cost of power, electricity and real estate. This business model also incurs no marketing and very little sales cost to Oracle and Oracle has the sole power to set pricing and quote customers.

  - Oracle has previously guided database cloud to be ~$20B in FY30 vs. $2.4B In FY25, implying 53% CAGR.

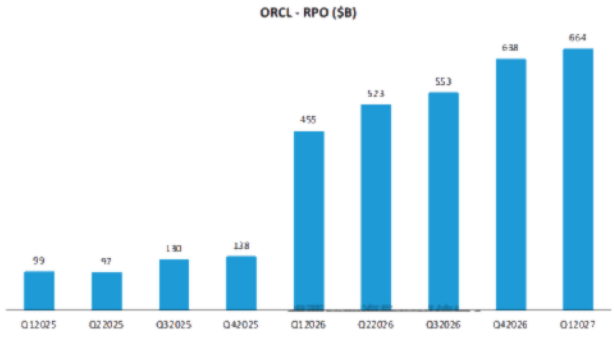
- **Sovereign Cloud:** OCI Gen 2 is technologically different from AWS, Azure and GCP and that difference is allowing Oracle to win in sub-sets of the market that are not well served by their bigger competitors. We discuss the differences at length in our recent notes (see Key takeaways from OCI expert call, Key takeaways from visiting an OCI datacenter and meeting OCI management).

EXHIBIT 16: **Oracle previously guided to significant OCI growth in the next 5 years, driven by RPO. They did not explicitly update this guidance**



Source: Company disclosures, Bernstein analysis

EXHIBIT 17: **Oracle RPO**



Source: Company disclosure, Bernstein analysis

WHAT ABOUT MARGINS?

AI, we believe, is going to be a drag on gross margins at least during the build-out phase and possibly much longer for IaaS/PaaS. Frontier lab consumption (training and inferencing) is a form of IaaS while enterprise inferencing is effectively PaaS. Thus, the mix of customer type and workloads will influence margins, but near term this is a scale game creating headwinds to overall margins. Meanwhile, 2 questions overhang the buildout of AI capacity in general:

on 22-Sep-2026

- **Supply / demand questions for AI capacity.** If there is an overbuilding occurring then pricing for rental of bare metal (or low software) capacity could be hit.

- **What is the real useful life for GPUs?** For the last few years, the change in the useful life of servers and related hardware has created a mixed impact to margins for Amazon, Microsoft and Google, but that may be changing: AI workloads are wearing out servers faster than Amazon had expected. We saw in 2025 that Amazon lowered the useful life of a subset of their servers from 6 to 5 years, likely due to the shorter life of GPU servers. On the other hand, Microsoft increased the useful life of servers and networking equipment to 6 years in 4Q22. The change comes about from improvements in their hardware, software and data center designs, which is feasible, but it is unclear how many additional efficiencies can be gained.

  - We try and answer some of these questions in the note: AI Value Chain: Can you really run a GPU for 6 years?

  - For investors looking over a longer horizon on this topic comparing to CPU-centric workloads, we wrote Global Software Long View: Cloud CAPEX - what does it really cost to offer IaaS / PaaS?.

- Microsoft is now talking about AI margins be similar to non-AI margins (Global Software: Why you want to own Microsoft - Key takeaways from meetings & marketing with management) which is very different from the way many on the street is currently interpreting the long term value of AI and especially AI IaaS/PaaS.

- One interesting takeaway from recent earnings is how the companies are treating the recent increases in component costs. Oracle specifically explained that where they have clarity and line-of-sight to pricing, Oracle has fixed priced contracts (for OCI these are the contracts going live shortly) but for contracts with longer time durations or where the pricing could vary, Oracle has conditions in their contracts that allow them to pass along those costs. In other words, Oracle's gross profit dollars will not be meaningfully impacted by price increases. We do not yet know if how the other hyperscalers will, if at all, pass along to customers the cost of component price increases.

Based on company results and our estimates:

- **AWS** operating income rose to $16.6B, with margins expanding to ~39%, up ~170 bps Q/Q from 38% in 1Q26. While higher infrastructure and depreciation costs tied to AWS's AI build out could weigh on margins near term, AWS remains focused on driving efficiencies through investments in software and process improvements that optimize server capacity utilization and enable a more efficient network architecture, increasingly powered by lower-cost custom silicon and proprietary networking equipment. We note that Amazon benefited from a $600M credit in change in energy contracts value, though even excluding this tailwind margins still grew ~20bps Q/Q.

- **Microsoft** does not report Azure margins. We estimate the Azure gross margins (GAAP) to be 50.7% in 4FQ26, down YoY where we estimated it to be 52.6%. Microsoft Cloud margins were 65%, a 3pp decrease Y/Y driven by sales mix shift to Azure, as well as scaling their AI infrastructure and growing AI product usage which the company is managing to offset through managing OPEX.

  - Microsoft is managing margins. While AI has been an incremental headwind, Microsoft will likely offset this near term via managing OPEX. Operating margin came in flat Y/Y in FY26, and management is guiding to FY26 operating margin to be down less than 1pt Y/Y. We have written over the years how successful Microsoft has been at driving operation leverage. This has been visible in the fact that GAAP margins improved as the company shifted to the Cloud (a feat few have accomplished).

  - We recently did extensive work on Azure AI and non-AI margins (Microsoft: Azure AI margins, and its implications for next 5 years). Based on our estimate of gross margin for total Azure and AI revenue, we estimate that core Azure (non-AI) revenue has a gross margin of ~65%, which is roughly inline with historical trends but lower than the maximum gross margins of ~70%. To understand the decrease in peak margins we recommend you read the note.

- Microsoft has started recently to talk about how AI margins will not be dissimilar to non-AI margins across the business (Global Software: Why you want to own Microsoft - Key takeaways from meetings & marketing with management).

- One of the areas of concern by investors are the lease and purchase commitments that have not yet commenced, what some call "off balance sheet" amounts. We recently looked at Microsoft's lease and purchase commitments and determined that

on 22-Sep-2026

the annual growth in SQ FT is roughly in-line with the prior 8+ years and therefore should not negatively impact margins (Microsoft: Are they overbuilding capacity?).

- **Google Cloud** operating income at $8.8B tripled Y/Y, and operating margin landed at 35.6%, up from ~20.7% in 2Q25 and ~33% in 1Q26. Margin expansion is being driven primarily by strong top line growth, which creates operating leverage complemented by improved cost discipline and operational efficiency. While higher depreciation from ongoing data center CapEx and now third-party bridge compute capacity are headwinds, there is a lot of focus on balancing these headwinds with efficiency gains.

- **Alibaba** reports an adjusted EBITA margin for its Cloud segment, inclusive of internal workloads and, from this quarter, its chip design business. In the June quarter, Alibaba Cloud's adjusted EBITA margin improved 250bps sequentially to 11.6%, or 440bps from a year earlier.

- One of the key debates with investors has been the impact of Alibaba's sizable investment plan and whether incremental contribution margins can offset the higher depreciation burden associated with increased capex. Speaking around the quarter, management reiterated its long-term EBITA margin target of 20% and indicated that the payback period for its capex investments is around three years, which could shorten to 2.5 years with efficiency gains from its in-house chip design business.

- Looking ahead one of the main debates will be the speed at which in-house silicon improves in terms of cost efficiency. Our channel checks point to the latest M890 chipset entering the market.

- **Oracle** is entering a massive investment phase as they build out AI Cloud capacity. This will weigh on gross margins, but the company is driving operating efficiency and operating leverage to offset the impact of all the CAPEX. We detailed how all of Oracle's key financial metrics will be impacted by the step-up in investment in Oracle: The story is better than we thought, with better upside potential. This note is based on Oracle's latest financial results and guidance and should be used, we recommend, to understand the shape and direction of key metrics as the company invests in the build out and beyond.

- Oracle supplied at their 2025 Financial Analyst Meetings (Oracle: Placing in perspective recent news articles and what to expect from the Financial Analyst Day) guidance on the margins of the different parts of their OCI business:

  - OCI AI infrastructure segment is expected to have non-GAAP gross marinas of 30-40%. The margin range is driven by the fact that Oracle is delivering a different product mix to each customer (e.g. training, inferencing, use of Oracle infrastructure software etc.) and by location driven by datacenter costs including power.

- OCI distributed Cloud offering will have non-GAAP gross margins in the range of 40-60%. Again the margin range depends on the services used by the customers.

- Cloud native customers will also have non-GAAP gross margins in the range of 40-60%

COMPARING THE 5 GLOBAL HYPERSCALE CLOUD PROVIDERS

Most of the 5 hyperscalers (Amazon, Microsoft, Google, Oracle and Alibaba) report some form of Cloud revenue. Based on reported data and our estimates, Amazon AWS continues to be the largest (see Exhibit 5) but many of the others are attempting to catch up, growing faster off a smaller base.

Please note that most of the five include revenue other than IaaS/PaaS revenue within their Cloud business. For example:

- Amazon includes within AWS 3rd party software, professional services, and other services (e.g. Alexa services).

- Microsoft historically did not report Azure revenue but supplied dollar figures during the FY27 re-segment and will be supplying each quarter.

  - Following the FY27 re-segment (link), Microsoft moved all the SaaS components out of Azure and the new Azure now is a pure PaaS/IaaS business.

on 22-Sep-2026

  - Please note that Microsoft reports Commercial Cloud revenue but that includes Azure, Office 365, Dynamics 365, and LinkedIn.
- Google includes Google Workspace (formerly known as G-Suite) within the Google Cloud.
- Alibaba includes 3rd party software and enterprise services, cloud desktop, etc. within the Alibaba cloud marketplace.
- Oracle OCI includes the company's old hosting business, which is now de minimus as well as Cloud @ Customer and Multi-Cloud database (Oracle database at Microsoft, Amazon and Google).

While each of the four global hyperscale Cloud providers supply some data or metrics on their Cloud businesses, the information is not complete (e.g., margins are not always supplied) nor easily comparable. Exhibit 18 shows the limited data that the five largest IaaS/PaaS vendors supply. Adding to the complexity of comparing the businesses, revenue (when disclosed) is not pure IaaS / PaaS revenue, and therefore not comparable since it includes other revenues (e.g. SaaS, professional services, 3rd party sales, other products).

EXHIBIT 18: **Each of the IaaS/PaaS vendors supply limited data on the size and economics of their IaaS/PaaS businesses**

	IaaS / PaaS revenue	IaaS / PaaS margins
Amazon	Revenue supplied	GAAP EBIT
Microsoft	Size of Commercial Cloud with some disclosure on Azure revenue	Limited margin commentary but no data
Google	Revenue supplied for Google Cloud (which includes Workspace)	GAAP EBIT supplied (which includes Workspace)
Alibaba	Alibaba Cloud Revenue supplied	Adjusted EBITDA margin supplied
Oracle	OCI Revenue supplied	Not Supplied

Source: Company reports, Bernstein analysis

WHAT DATA INVESTORS ARE LOOKING FOR AND WHERE TO FIND SOME HELP ANSWERING THESE QUESTIONS

We believe that investors want to compare and contrast the different vendors (especially using the US vendors as a proxy for the Chinese hyperscale providers and the new entrants into the market) and understand the trajectory of each company's revenue and margins given that disclosures are often incomplete and different companies report at different times.

We believe that the six (formerly 4) main client questions are:

- IaaS/PaaS revenue
- Current and long-term stable margins
- Competitive capabilities
- How much revenue could each vendor capture?
- How much of revenue is AI?
- What are the margins / ROIC of AI?

To help clients better understand these complex businesses the Global Software team published a series of notes looking at the Total Addressable Market (TAM) for IaaS/PaaS and estimated Microsoft Azure's long-term revenue opportunity. We pulled many of our Cloud and especially IaaS/PaaS notes into a Blackbook: Cloudification of Tech: The Sky's the Limit. We then wrote a note looking at how long it could take for Cloud saturation and whether everything moves to the Cloud (see Weekend Tech Byte: Can everything move to the Cloud overnight?) — for investors trying to get a perspective on how long this march to the Cloud is going to be accretive to revenue, we recommend this note.

on 22-Sep-2026

We also published a note exploring the question: "how should one think about long-term capital intensity and COGs for IaaS/PaaS?" in Global Software Long View: Cloud CAPEX - what does it really cost to offer IaaS / PaaS?. This note helps explain the drivers of COGS. We recommend this note for anyone interested in IaaS/PaaS and wanting to understand long-term margins, and capital intensity.

Additionally, we looked at how the shift to the Cloud is impacting the rest of the tech stack — given the expectations that the shift to Cloud would negatively impact IT hardware. In our note Global Software: If software / Cloud wins in the shift to the Cloud then who loses?, we discussed our surprising results that to date at least, spending on IT hardware and related datacenter costs is not being negatively impacted. Instead, we believe that the margins generated by Cloud are coming from internal and external IT services. A thought-provoking note and one worth reading — if you have thoughts on how/ where we could be wrong, we would appreciate the feedback.

The Bernstein Tech team also published a note Global Tech: Big Tech's big spending... big returns?) looking at CAPEX spending and how it related to revenue generation as well as how the large CAPEX spend will impact the leading tech companies .

Please note that some of the companies report results off calendar.

HOW DOES GPU AS-A-SERVICE COMPARE TO TRADITIONAL CLOUD?

We recently expanded our scope to include CoreWeave, whose rapid emergence as a GPU-first neocloud highlights the increasing importance of specialized infrastructure providers in alleviating capacity constraints and reshaping the competitive dynamics of hyperscale AI deployment. Per our CRWV initiation, our view of CoreWeave's business model is that it serves as an interim capacity provider to hyperscalers and large AI customers—benefiting from near-term power and GPU scarcity, but facing longer-term sustainability challenges given a concentrated customer base, evolving buyer insourcing incentives, and structurally constrained returns relative to scaled cloud incumbents. We see three key themes from the 2Q print:

Although long-dated contracts continue to underpin CoreWeave's business model, the company is increasingly trying to embrace shorter-term agreements that carry both higher revenue potential and greater residual-value risk. Management continues to describe a typical contract as a five-year agreement, with the initial term structured to repay asset-level financing and generate attractive returns before considering any renewal or resale value. The company also highlighted an A100 contract extending into 2029, reinforcing that older GPU capacity can still support multi-year commitments at attractive pricing. That said, CoreWeave is expanding into shorter-duration enterprise agreements. Its second publicly syndicated term loan was the first financing to include shorter-duration customer contracts, which management views as an important milestone in broadening the company's customer addressable market. The company noted that enterprise customers often prefer 2-3 year procurement cycles rather than five-year commitments.

Management indicated that the new financing structure allows it to support customer contracts at durations that are more consistent with enterprise purchasing preferences, while also increasing exposure to shorter-dated agreements that tend to carry higher ASPs and margins. Thus, CoreWeave's strategy intends to preserve long-term contracts as the base of its revenue and financing model, while using shorter-duration commitments to target enterprises, monetize infrastructure rolling off legacy contracts, and expand managed inference. The financing structure expands CoreWeave's ability to serve enterprises seeking 2-3 year commitments, potentially supporting higher ASPs and margins. However, the benefit comes with greater residual value risk, as shorter terms expose CoreWeave more quickly to re-leasing, utilization, and technology-obsolescence risk if newer GPU generations reduce demand for prior-generation capacity.

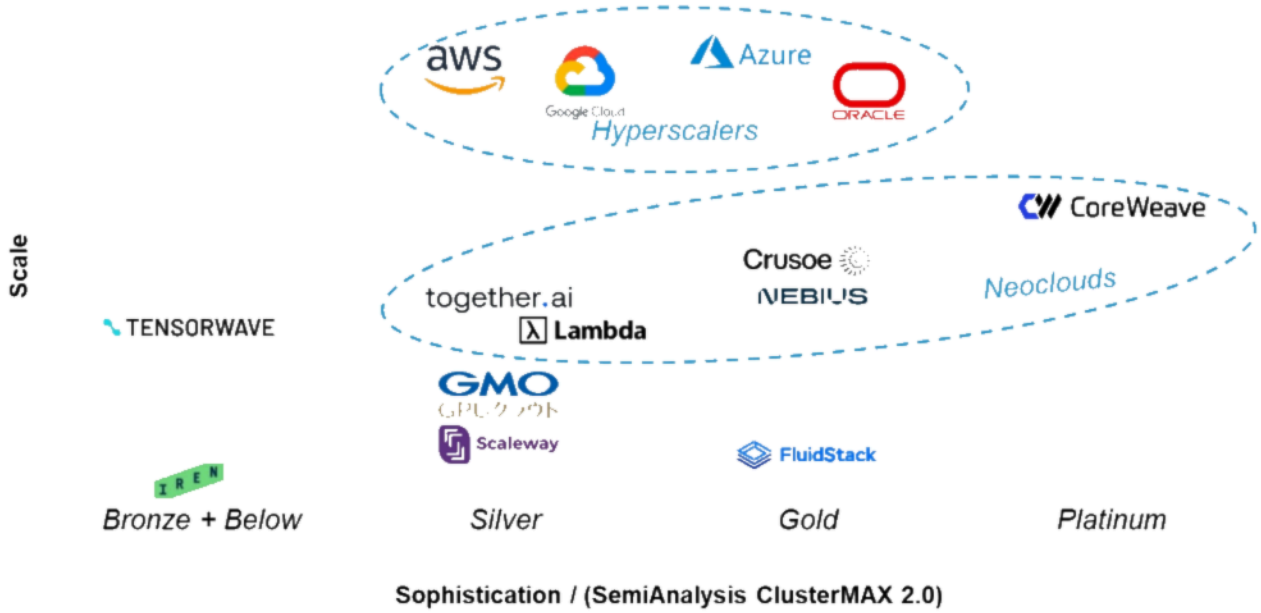
2Q reinforced CoreWeave's ability to secure and energize capacity, but also highlighted the substantial execution required to transform CoreWeave's growing contracted-power portfolio into revenue-generating infrastructure. CoreWeave ended 2Q26 with 1.5 GW of active power, establishing the largest operational neocloud platform among its peers, following the addition of nearly 500 MW during the quarter. More than 300MW was energized in June alone, meaning the capacity ramp was materially back-end-weighted and should contribute more fully to revenue through 3Q and 4Q. CoreWeave raised its FY26 year-end active power target to more than 1.85GW, from a prior target of more than 1.7GW. The company also reported 3.7GW of contracted power at quarter-end and approximately 4.2GW after adding roughly 500MW subsequent to quarter-end. This contracted-power figure excludes more than 1.5GW of additional potential capacity associated with powered land, site expansion options, and executed LOIs. CoreWeave's long-term target of at least 8GW of active power by 2030 remains highly execution-dependent. While the company has secured 4.2GW of contracted power, this is not equivalent to operational or revenue-generating capacity and excludes the remaining development, construction, interconnection, hardware procurement, customer deployment, and financing steps required to convert power rights into active capacity. The scale of the

on 22-Sep-2026

gap between its current 1.5GW active base and its 2030 target leaves the company exposed to delays, cost inflation, funding constraints, supply-chain disruptions, and any deterioration in customer demand or pricing before capacity is delivered.

CoreWeave emerged from 2Q with a more aggressive investment outlook, supported by recent customer wins, rising pricing, and increasing capacity commitments—a theme we are seeing across the neocloud space. CapEx accelerated across the neocloud sector, with CoreWeave providing the clearest near-term upward revision, increasing its 2026 CapEx outlook to $35-39B driven by a higher expected level of capacity delivered to customers during the year and several significant recent customer wins. Q2 CapEx was $9.4B, modestly above the high end of its prior quarterly guidance, while construction-in-progress increased to $11.9B from $9.6B sequentially. The increase reflects accelerating physical deployment, as CoreWeave added nearly 500MW of active power during 2Q, including >300MW in June alone—ending the quarter with 1.5GW. CoreWeave's capital deployment is occurring against a backdrop of rising contract economics, and management noted that contracts signed in Q2 are expected to carry contribution margins 5-10ppt above those added in recent quarters, while July pricing increased by approximately 25% across SKUs. This supports the view that higher CapEx is being deployed into a stronger pricing and margin environment, although the expanded investment level will also increase the company's financing needs and execution sensitivity.

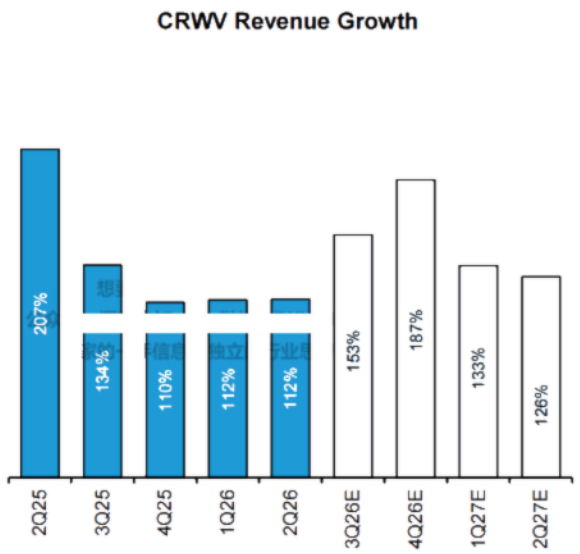
EXHIBIT 19: **GPU Cloud Landscape (Scale v. Sophistication)**



Source: SemiAnalysis ClusterMAX 2.0, Company fillings, U.S. Department of Energy Solstice announcements, Top500 / Green500, PitchBook, Bloomberg, Reuters, Financial Times, The Information, Bernstein Analysis & Estimates

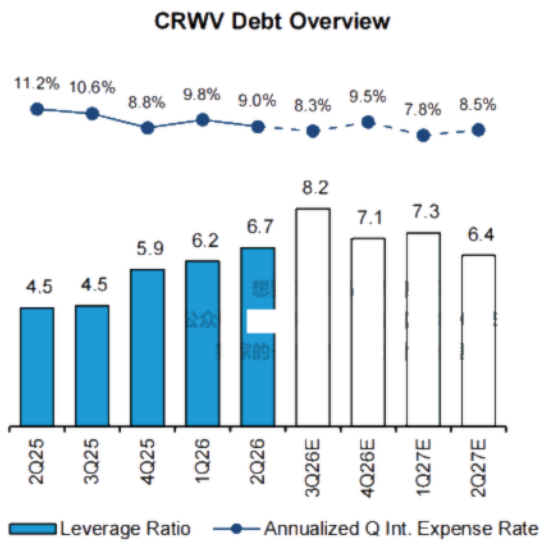
on 22-Sep-2026

EXHIBIT 20: **CRWV Revenue Growth**



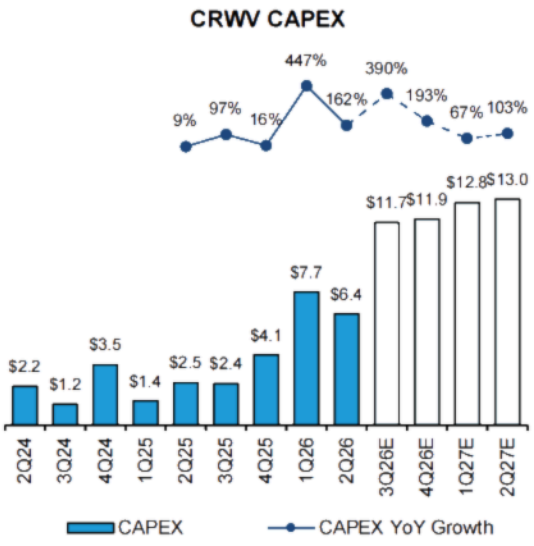
Source: Company filings, Bernstein Analysis and Estimates

EXHIBIT 21: **CRWV Debt Overview (Times, %)**



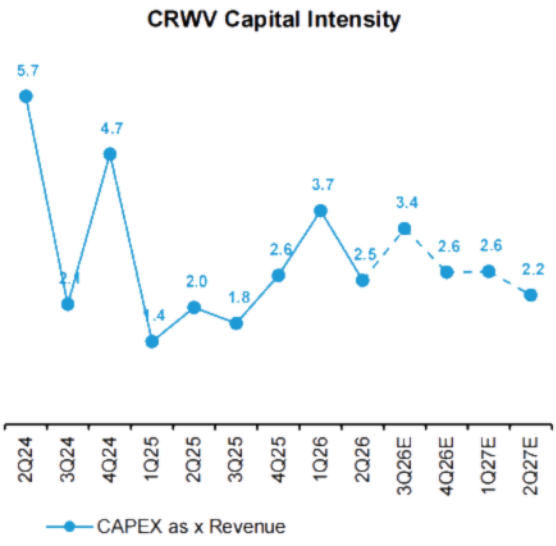
Leverage calculated as LTM Adj. EBITDA over net debt
Source: Company filings, Bernstein Analysis and Estimates

EXHIBIT 22: **CRWV CAPEX Expense ($B, %)**



Source: Company filings, Bernstein Analysis and Estimates

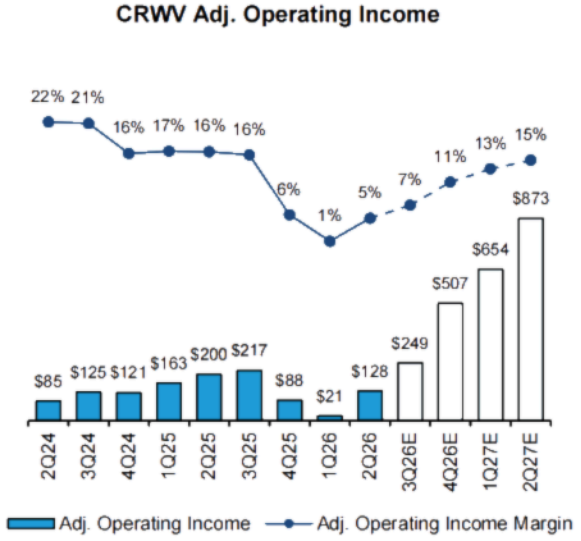
EXHIBIT 23: **CRWV Capital Intensity (Times)**



Source: Company filings, Bernstein Analysis and Estimates

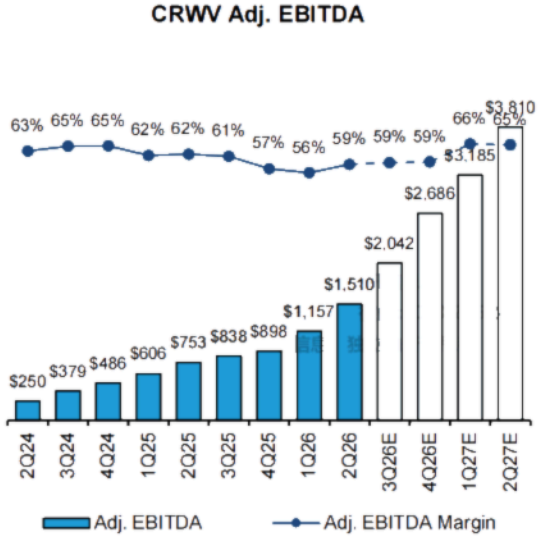
on 22-Sep-2026

EXHIBIT 24: **CRWV Adj. Operating Income ($M)**



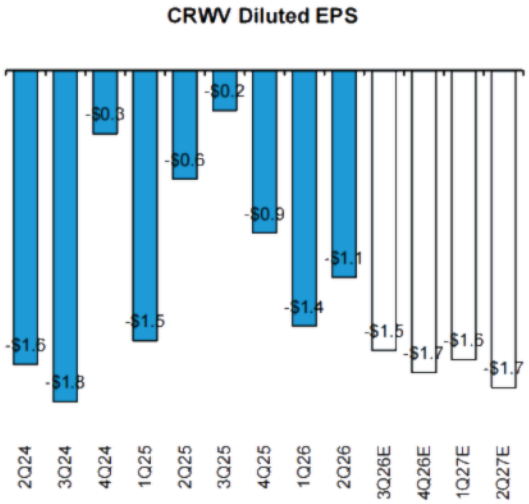
Source: Company filings, Bernstein Analysis and Estimates

EXHIBIT 25: **CRWV Adj. EBITDA ($M)**



Source: Company filings, Bernstein Analysis and Estimates

EXHIBIT 26: **CRWV Diluted EPS ($)**



Source: Company filings, Bernstein Analysis and Estimates