The AI Economy

How a single person’s behaviour becomes someone else’s revenue

Keeley Walsh | Associate Analyst, Research

19 May 2026

The consumer

Consumer adoption is the proof point that drives the enterprise demand multiplier

Consumers are the test dummy.

If consumers won't change their habits, businesses won't spend money rebuilding their systems. Every major technology shift in modern history followed the same pattern: consumers adopt first, enterprise demand follows, and capex leads both.

The same pattern, four times:

1990s	**Consumer:** Personal email and the web **Enterprise:** Corporate email, intranets, e-commerce
2000s	**Consumer:** iPhone, mobile apps **Enterprise:** Mobile workforce, BYOD, cloud productivity
2010s	**Consumer:** Dropbox, Gmail, Netflix **Enterprise:** AWS, Azure, GCP and the enterprise cloud era
2020s	**Consumer:** ChatGPT, Copilot, Gemini **Enterprise:** Enterprise AI agents and AI-native workflows

One day. Eight AI interactive moments

From the first scroll to the last stream, every interaction routes money through an AI-exposed company

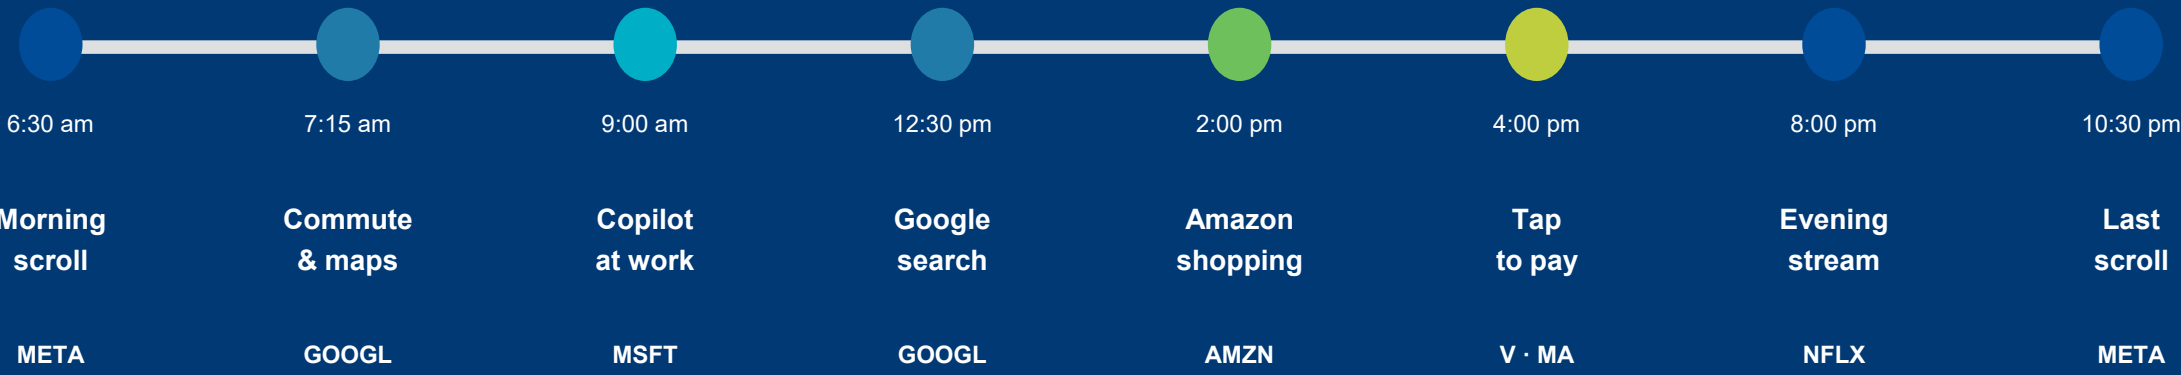

Revenue type: Advertising · Cloud / SaaS · Subscriptions · Payments

Revenue flows through a three-layer stack

Applications — **Cloud & data centres** — **Chips & Hardware**

Follow the dollar

How each tap makes its way from a search bar to revenue

Applications

A single AI interaction looks unremarkable on the surface. A few words in, a few words back. In practice, that one exchange triggers a chain of payments running through four distinct businesses, each charging the one above it.

Cloud & data centres

The prompt is processed by an application, which runs on a cloud provider's servers, which are built around specialised chips, manufactured in a foundry. The cost of a single query works its way down through all four layers.

Chips & Hardware

Each of those businesses earns revenue differently, competes against different peers, and carries a different risk profile. A thesis built on only the most visible name (typically NVIDIA) overlooks much of where AI capital is actually flowing.

01 Applications

Where consumers interact, the first revenue stream starts with the companies consumers use. Your feed doesn't sleep.

This is where the initial interaction takes place

Before you open your app, Meta's Advantage+ AI runs a real-time auction across 500+ candidate posts and selects a personalised ad in under 200ms. Revenue flows the instant the ad impression is served.

Advertising

Meta

Advertising and feed ranking

- Ranks 500+ posts per feed open in under 200ms using transformer models
- Advantage+ selects the highest-value ad for each person in real time
- AI also powers content moderation, spam detection and Reels ranking
- They hold their own AI infrastructure

Alphabet

Search & navigation AI

- Gemini generates AI Overviews on every eligible search query
- Google Maps uses AI for live traffic prediction and route optimisation
- They use Google Cloud and are their own largest customer

Subscriptions

Netflix

Recommendation and retention

- Recommendation AI scores 15,000+ titles per user to build the homepage
- ~80% of content watched is discovered via the AI recommendation system
- AI also optimises video encoding and thumbnail personalisation

Transaction fees

Amazon

Shopping & ad optimisation

- Recommendation engine scores 1M+ products in 50ms per browse session
- AI-powered sponsored product auctions run on every search result page
- They use their own (AWS) AI infrastructure and proprietary chips
- They also generate revenue from prime subscriptions

02 Cloud & data centres

Where compute happens, the second revenue stream is the companies the applications pay

This is where application interactions flow and live

Amazon, Microsoft and Alphabet are vertically integrated, meaning they are their own largest cloud customers, capturing both the application fee and the compute margin.

AWS

Amazon's own Infrastructure

- Runs AI training and inference for Amazon's own products and for thousands of external customers via SageMaker and billed by usage
- Built proprietary chips including Inferentia to cut inference cost by ~70% vs third-party GPUs
- They are the compute landlord like Azure and rent to other companies, not just Amazon

Azure

Microsoft's own Infrastructure

- Every Copilot prompt fires a GPT-4o API call on Azure GPU clusters, billed per token consumed
- Exclusive cloud partner for OpenAI where all ChatGPT and API traffic runs on Azure
- Azure AI Studio lets enterprises fine-tune and deploy their own models at scale

Google Cloud

Alphabet's cloud

- Runs Gemini, Search AI Overviews and YouTube recommendations on Google's own infrastructure
- Custom TPU chips (now on Gen 6) compete with Nvidia GPUs for training
- Growing fast in enterprise AI. They had the fastest hyperscaler cloud growth in 2024-25

CoreWeave

Pure-play AI infrastructure

- Rents Nvidia GPU clusters to customers who need compute
- Nvidia holds an equity stake
- CoreWeave gets early access to the latest GPU hardware before others
- The first pure-play AI infrastructure company to go public

03 Chips & hardware

Where AI becomes tangible, these are the companies underneath that make all of this possible

This is where AI processing is built

ASML is the only company in the world that manufactures EUV lithography machines which is the equipment TSMC uses to fabricate Nvidia's chips. Hence the three companies form an unbroken monopoly chain. The hyperscalers rent the Nvidia GPU chips to process their compute.

Nvidia

The GPU maker

- Designs but does not manufacture its own chips. All chip fabrication is outsourced to TSMC
- H100/H200/Blackwell GPUs are the primary hardware for training and inference - running AI models globally
- CUDA software ecosystem is backwards compatible and locks developers in because rewriting AI code for a different chip is costly and slow

TSMC

The chip maker

- Manufactures chips for Nvidia, Apple, AMD, Qualcomm and Google's custom TPUs at 3nm precision
- No other can match its leading-edge processing at scale
- Nvidia's sole manufacturer

ASML

The machine maker that builds the tool that makes the chips

- Leading manufacturer of photolithography machines which use ultraviolet light to etch circuits at atomic scale
- ASML is the leading manufacturer of chip making and test equipment. It is a supplier to TSMC, Intel and many other chip makers
- Each machine costs ~$380M, takes years to build, and contains 100,000+ components

One day. Eight AI revenue moments

The individual compute power per average interaction over a day is minute in isolation but massive in aggregate

The math from one person to a capex super-cycle

~1.004 Wh per person per day for AI interactions (99.7% of that is Netflix video). Pure AI inference is **~0.004 Wh/day** - about 14 seconds of an LED light.

× **1 billion users** = ~4 million kWh of AI inference daily and that is ***today***, before agentic workflows (×100) and video generation (×1,000+) arrive.

Today's interaction is not tomorrow's interaction

However, it signals where demand is heading. Compute intensity per user is growing exponentially, justifying the capex

x1	x10	x100	x1,000+
Yesterday (2022)	Today (2026)	Emerging (2026-27)	Next (2027-28)
Text-only prompts	**Multimodal AI**	**Agentic workflows**	**Video generation & long agents**
A ChatGPT question. A Google search. One model call, a few hundred tokens, response in 1-2 seconds.	AI Overviews, image generation, voice mode. The model processes pixels, audio and text in the same call.	Agents that book your flight, reply to your inbox, or run multi-step research. Minutes of compute per request.	Generative video, simulated environments, autonomous agents running for hours. Compute scales with screen-time, not click-count.

The equation: more users × more interactions per user × more compute per interaction × longer sessions. Four multipliers, all heading up at once.

Why it matters

Combined annual capex of Microsoft, Meta, Alphabet and Amazon. The supply being built to meet today's demand and tomorrows

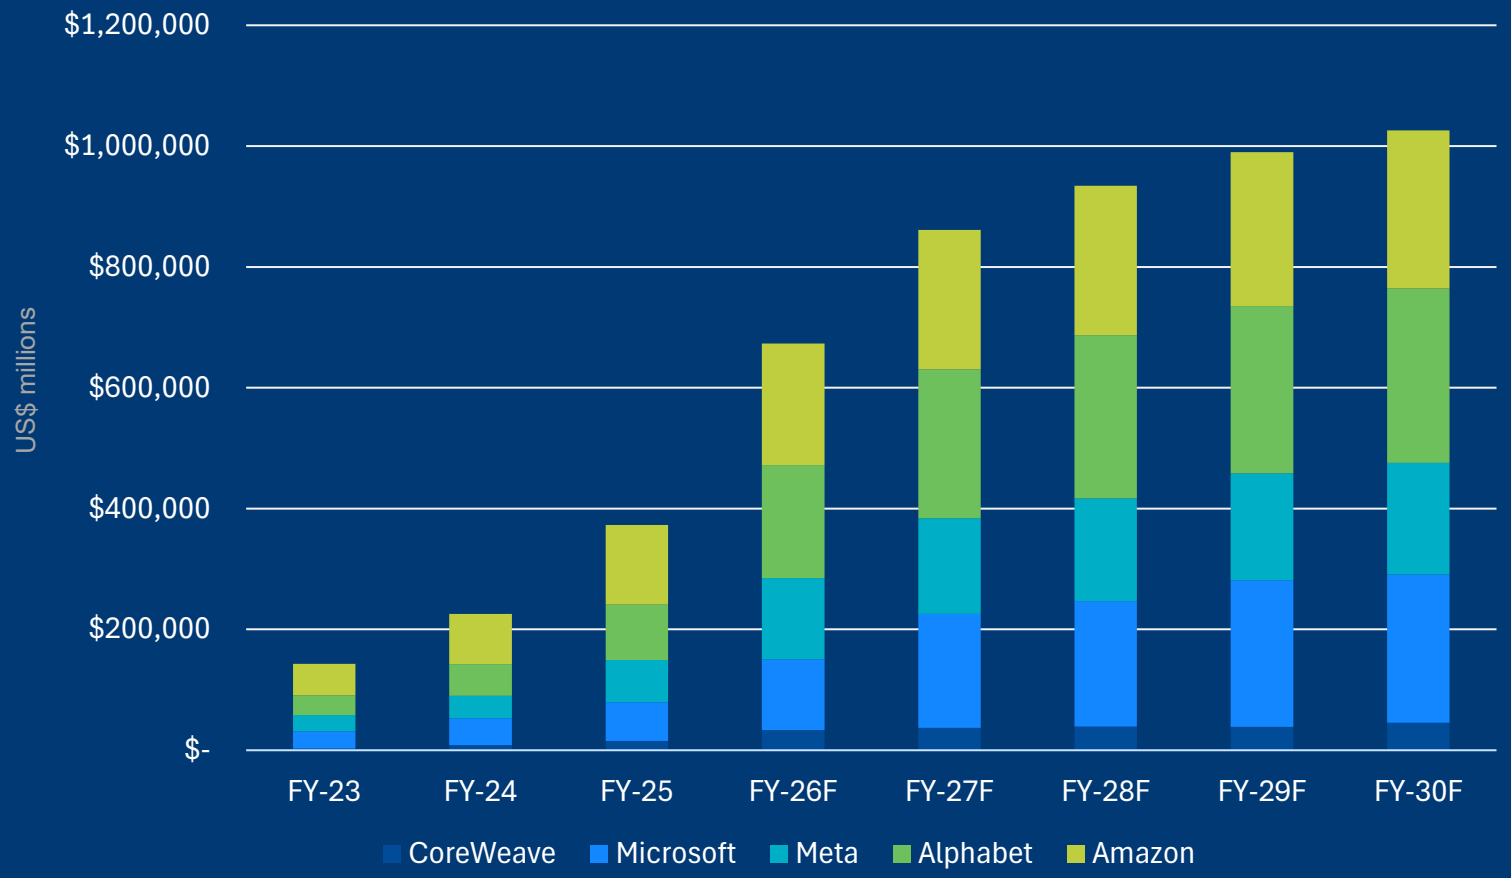

Source: Visible Alpha

~US$673bn

estimated combined Big 4 hyperscaler capex (+ CoreWeave) for FY26 alone

The eight moments earlier in this deck are the demand.

The capex is the supply being built to meet them.

Bringing it together

How consumer-level behaviour feeds the AI capital cycle

Observation

Each AI interaction generates revenue at four points in the chain. The volume of those interactions supported by future enterprise adoption is what the hyperscalers cite when justifying current capital expenditure. Consumer-level demand and infrastructure-level supply are visibly tracking each other.

Implication

AI exposure is not a factor affecting a single business but a series of businesses across the chain, each with different unit economics, margin profiles, and competitive risks. The compounding effect of companies implementing AI within their work-flows is the demand driver.

Outlook

The capex now being committed across the hyperscalers is sized for record order books (customer pre-purchases/RPOs) and use cases that are still emerging. Compute per user is already rising, and the multipliers from agentic workflows and video generation come on top of today's run rate.

Disclaimer

This presentation is provided for general information purposes only and is not intended as an offer to enter into any transaction. The information contained within this document is not necessarily complete and its accuracy can not be guaranteed. We have prepared this presentation without consideration of the investment objectives, financial situation or particular needs of any individual investor.

Before a client makes an investment decision, a client should, with or without Morgans' assistance, consider whether any advice contained in the presentation is appropriate in light of their particular investment needs, objectives and financial circumstances. It is unreasonable to rely on any recommendation without first having spoken to your adviser for a personal recommendation.

The information contained in this presentation has been taken from sources believed to be reliable. Morgans Financial Limited does not represent that the information is accurate or complete and it should not be relied on as such. Any opinions expressed reflect Morgans' judgment at this date and are subject to change. Morgans and/or its affiliated companies may make markets in the securities discussed. Further Morgans and/or its affiliated companies and/or their employees from time to time may hold shares, options, rights and/or warrants on any issue included in this presentation and may, as principal or agent, sell such securities.

The Directors of Morgans advise that they and persons associated with them may have an interest in the above securities and that they may earn brokerage, commissions, fees and other benefits and advantages, whether pecuniary or not and whether direct or indirect, in connection with the making of a recommendation or a dealing by a client in these securities, and which may reasonably be expected to be capable of having an influence in the making of any recommendation, and that some or all of our representatives may be remunerated wholly or partly by way of commission.

The presentation is proprietary to Morgans Financial Limited and may not be disclosed to any third party or used for any other purpose without the prior written consent of Morgans.

Morgans Financial Limited ABN 49 010 669 726 AFSL 235410
A Participant of ASX Group

Principal Office: Level 29, Riverside Centre, 123 Eagle Street, Brisbane QLD 4000

This document has been prepared by Morgans Financial Limited in accordance with its Australian Financial Services License AFSL no. 235410. The views expressed herein are solely the views of Morgans Financial Limited.

Stockbroking | Wealth Management | Corporate Advice