BERNSTEIN
SOCIETE GENERALE GROUP

18 May 2026
Initiation

## Japan and EU Semiconductors

# SoftBank, Arm: From GenAI to Agentic AI; Initiating with Outperform Ratings


**David Dai, CFA**
+852 2918 5704
david.dai@bernsteinsg.com

Juho Hwang
+852 2123 2632
juho.hwang@bernsteinsg.com

Jack Lin
+852 2123 2683
jack.lin@bernsteinsg.com

Carmine Milano, CFA
+44 20 7762 1857
carmine.milano@bernsteinsg.com


**The Gen AI paradigm is rapidly shifting from 1.0 (chatbot) to 2.0 (agent).** In Gen AI 1.0, Chatbots are limited to advising, leaving humans to interpret and do the actual job. With Gen AI 2.0, AI Agents are able to plan, act, and execute jobs, and iterate entire workflows autonomously. Agentic AI creates much more value than AI 1.0 and sees much broader adoption; Capgemini estimates the economic impact of Agentic AI to be at least $456bn during 2025-28. Agentic AI consumes 1,000x more tokens compared to AI 1.0, which drives up AI token demand and hence monetization for AI agents, model, infrastructure and chips. Agentic AI will also lead to Gen AI 3.0 (physical AI) and eventually 4.0 (AGI).

**SoftBank benefits greatly from its investment of the full stack of AI and the evolution from Gen AI 1.0 to 4.0**. SoftBank is one of the few companies that invested to the entire stack including infrastructure (SB Energy), Silicon (87% of Arm, Ampere, Graphcore), Models/Agents (11% of OpenAI), and physical AI (ABB Robotics). Agentic AI drove AI adoption and revenue growth of **OpenAI**: its revenue ARR grew more than 10x to $25bn in March 2026 vs. 2 years ago, and its Codex has experienced a "ChatGPT moment" in April 2026 by evolving into a powerful, autonomous agent. **Arm** is the center of the renaissance of CPUs because of Agentic AI, and potentially will capture the GPU accelerator TAM as well. Increased token consumption and AI adoption also means more AI infrastructure build out, including Stargate, which is what **SB Energy** is for. Finally, **SoftBank Robotics** is well positioned for the future rise of Physical AI, especially with the investments in ABB Robotics. Thanks to these investments, SoftBank's NAV has grown $105bn since it started investing in OpenAI in 3Q CY2024.

**Arm is the structural beneficiary of the renaissance of CPUs for agentic AI.** Compared to Gen AI 1.0, Agentic AI involves heavily autonomous task orchestration and execution, which can only be done by CPUs. While traditional AI data centers require ~30 million CPU cores / GW of compute, Agentic AI data centers require ~**120 million CPU cores / GW** — a **4x structural increase.** The ratio of GPU:CPU is shifting from the current 8:1 to 2:1 / 1:1, with server CPU TAM quadrupling to $137bn by 2030. Arm stands out in server CPUs given its unparalleled power efficiency. In addition, Arm is shifting from just IP provider to CPU maker, aiming to capture $15bn revenue by CY2030.

We forecast Arm's revenue to grow more than five-fold to $26 Bn by 2030, with EPS expanding similarly (~5.5x) to $9.83, supported by the growing adoption of Arm CPUs in AI data centers, its own Arm CPU revenue, and the rising royalty driven by the higher value Arm provides to clients through the rapid evolution of CPUs. With 90x P/E on Q5-8 EPS of $3.33, **we rate Arm Outperform with PT= $300.00**. In the bull case, we can see Arm being valued at 40x on $9.8 2030 (FY31) EPS, implying a valuation of $390.

We expect SoftBank's NAV to grow to $390bn in a year, mostly contributed by Arm (based on our PT of Arm). If OpenAI valuation goes up from the latest round of $852bn, that'd be additional upside. With a 25% NAV discount, **we rate SoftBank Outperform with PT = ¥8,200.00.** In the bull case, applying the high end of Arm valuation and 15% discount, we get a valuation for SoftBank at ¥11,200.




First Published: 18 May 2026 08:00 UTC Completion Date: 17 May 2026 15:02 UTC

## BERNSTEIN TICKER TABLE

| Ticker | Rating | Cur | 15 May 2026 Closing Price | Price Target | TTM Rel. Perf. | Cur | Adjusted EPS 2025A | 2026E | 2027E | Adjusted P/E (x) 2025A | 2026E | 2027E |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 9984.JP (SoftBank) | O | JPY | 5,745.00 | 8,200.00 | 152.7% | JPY | 872.47 | 251.35 | 152.44 | 6.6 | 22.9 | 37.7 |
| ARM (ARM Holdings) | O | USD | 228.50 | 300.00 | 34.3% | USD | 1.77 | 2.28 | 3.12 | 129.1 | 100.3 | 73.2 |
| JPL | | | 2,501.77 | | | | | | | | | |
| SPX | | | 7,408.50 | | | | | | | | | |

**COVERAGE INITIATION**

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

ARM base year is 2026;

In the ticker table, 2026 represents FY27/3 for SoftBank and FY26/3 for ARM.

Source: Bloomberg, Bernstein estimates and analysis.

## INVESTMENT IMPLICATIONS

We rate SoftBank (PT=¥8,200.00) and Arm (PT=$300.00) **Outperform**.

### Table Of Contents



## DETAILS

Our financial models can be downloaded here: SoftBank (9984.JP), Arm (Arm.US)

Our ARM Industry models can be downloaded here: ARM Server CPU Industry Model, ARM Automotive TAM Model, ARM Mobile IP TAM Model

## AGENTIC AI AND SOFTBANK / ARM

### AGENTIC AI IS GEN AI 2.0

Since the launch of generative AI, the ability for Gen AI to execute long tasks doubles every 7 months (Exhibit 1). This inevitably leads the paradigm shift from Gen AI 1.0 (chatbot) to 2.0 (Agentic AI). Agentic AI - AI systems that autonomously plan, act, and iterate toward goals - marks the decisive shift from AI as a passive tool to AI as an **active operational participant**. It is **Gen AI 2.0**: the tipping point for enterprise adoption. Agentic AI is not as an incremental upgrade, but as a generational leap - one that closes the execution gap Gen AI 1.0 left wide open.

#### Why Gen AI 2.0 breaks the barrier that Gen AI 1.0 couldn't

Gen AI 1.0 (chatbots, copilots) proved AI could generate value - but stalled adoption because **humans remained the execution layer**. Every ChatGPT output still needed a human to copy-paste, decide, and act. Agentic AI eliminates that bottleneck as it can execute entire workflows instead of just advising on how to execute (Exhibit 2). This is built on the exponential capability of AI to execute long tasks: AI performance in terms of the length of tasks AI agents can complete has been consistently exponentially increasing over the past 6 years, with doubling task capabilities every 7 months. AI agents have reached a tipping point in their capabilities.

As such, Agentic AI generates real economic gains: it reduces up to 86% of time spent on complex workflows previously requiring human execution. Companies report an average 171% ROI, versus far more modest productivity lifts from Gen AI 1.0 copilots. McKinsey estimates agentic AI will contribute $4.4 trillion annually to global productivity. In short, Gen AI 1.0 created demand. Gen AI 2.0 delivers results. The moment AI *acts* rather than *advises*, the ROI crystallizes - and mass adoption follows.

### The tipping point has already arrived

The transition from experimentation to deployment is complete. Enterprise adoption data in 2026 confirms we have reached the inflection point. The adoption curve of Agentic AI is reaching escape velocity. According to BusinessWire, **100% of surveyed enterprises** plan to expand agentic AI usage in 2026, with 65% already actively using AI agents. Fortune 500 agentic AI adoption **doubled year-over-year** - from 32% in Q1 2025 to 68% by Q1 2026. 81% of enterprises have either fully integrated or are actively scaling agentic AI across their teams.

EXHIBIT 1: **The ability for Gen AI to execute long tasks doubles every 7 months**

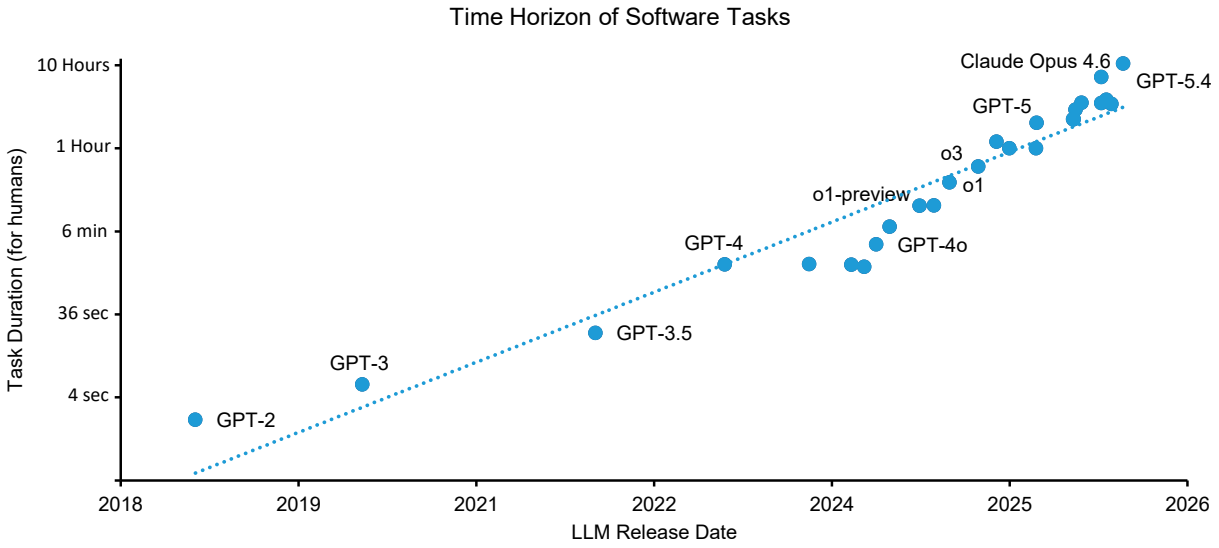


Source: METR, Bernstein analysis

EXHIBIT 2: **Agentic AI marks a shift from passive tools to active participants that plan, act, and iterate toward goals, representing Gen AI 2.0 and a generational leap that closes the execution gap.**

| | Gen AI 1.0 | Gen AI 2.0 |
|---|---|---|
| **Paradigm** | **Chatbots** | **Agentic AI** |
| **AI Capabilities** | Respond to prompts, advises | Plans, acts, iterates autonomously |
| **Interaction Model** | Single turn, prompt-response | Multi-step, goal-directed |
| **Action** | Text output only | Calls APIs, writes & runs code, browses web |
| **Human Dependency** | High — human must execute | Low — agent executes autonomously |
| **Human: AI ratio** | 1:0.1 | 1:1 to several |
| **Token Consumption** | 1x | 1,000x |

Source: Bernstein analysis

## AGENTIC AI DRIVES AI USAGE EXPLOSION, FOLLOWED BY MONETIZATION

With the rise of Gen AI 2.0 (Agentic AI), the economic value of AI now far exceeds the ceiling of Gen AI 1.0 (chatbot).

Gen AI 1.0 had a fundamental commercial flaw: it was easy to demo, hard to monetize. Chatbots and copilots generated impressive outputs, but the value leaked at the last mile - a human still had to review, decide, and act. This made ROI diffuse and hard to attribute, which is why enterprise AI budgets were hesitant and "AI pilots" became a graveyard. Gen AI 1.0 (chatbots) delivered real but bounded gains. The containment rate improvement curve is telling: traditional chatbots improved containment rates by 20-40%, and integrating LLMs pushed that to 40-55% - but that's where the ceiling was hit. Every CX leader that deployed chatbots between 2021-2024 found themselves staring at the same wall: you can only deflect so many conversations before you run out of low-hanging fruit. Chatbot KPIs - containment rate, CSAT, average handle time - are conversation metrics. They don't touch revenue, don't touch operations, and don't touch systems.

Gen AI 2.0 solves this by closing the value loop. When an agent autonomously completes a workflow end-to-end, the output is a *business outcome* - a closed ticket, a qualified lead, an executed trade - not just a text response. Outcomes are measurable. Measurable outcomes are billable. Agents don't optimize conversations - they complete outcomes. This distinction matters enormously for usage because:

- A chatbot handles one interaction at a time, passively
- An agent handles entire workflows, proactively, across multiple systems simultaneously

This means agentic AI can be deployed in every function - sales, ops, finance, HR, R&D, supply chain - not just customer-facing chat windows. The deployment surface is 10-20x larger than chatbots ever were.

Chatbot adoption required a specific trigger: a customer-facing interaction. Agentic AI adoption is triggered by any repetitive workflow anywhere in the organization - and 40% of AI-automated agent workflows today are already focused on customer experience alone, with the rest spanning finance, ops, HR, and beyond. Gartner projects enterprise app AI agent penetration jumping from under 5% in 2025 to 40% by end of 2026 - an 8x increase in a single year, a pace that chatbots never achieved even over a decade.

The usage numbers are unprecedented. OpenClaw, the phenomenal Agentic AI tool, surpassed Linux's 15-year cumulative star count on GitHub in 3 months (Exhibit 3). Ever since the launch of OpenClaw, token consumption on OpenRouter has accelerated (Exhibit 4). In Silicon Valley, companies and individuals compete on who can use more tokens by employing more AI agents to work for them: Meta reportedly runs an internal dashboard called **"**Claudeonomics**"** that aggregates AI token consumption across employees and ranks the top ~250 "AI power users." 87% of companies deploying agentic AI report direct revenue growth as a result. Gartner forecasts Enterprise applications with embedded agentic AI will grow from <1% in 2024 to 33% by 2028. McKinsey case studies report 20% to 60% productivity gains, over 50% lower time and effort in one modernization case, and over $3 million expected annual savings in another (Exhibit 5).

When ROI becomes attributable, budgets unlock - and the market numbers reflect exactly this. Capgemini forecasts that by 2028, AI agents could generate up to $450 billion in economic value through revenue growth and cost savings across surveyed markets (Exhibit 6). Precedence Research predicts that the global Agentic AI market size to increase from USD 7.55Bn in 2025 to approximately USD 199.05Bn by 2034, expanding at a CAGR of ~45% from 2025 to 2034 (Exhibit 7). This is the classic technology S-curve crossing the inflection point. Gen AI 1.0 built awareness and demand. Gen AI 2.0 delivers measurable ROI. Measurable ROI unlocks enterprise budgets at scale. Scaled budgets fund the infrastructure - model improvements, multi-agent frameworks, physical AI systems. The monetization engine is now the R&D engine.

EXHIBIT 3: **Openclaw's GitHub traction accelerates sharply in 2026, surpassing React and Linux in cumulative stars and highlighting exceptional developer interest versus two of open source's most established repositories.**

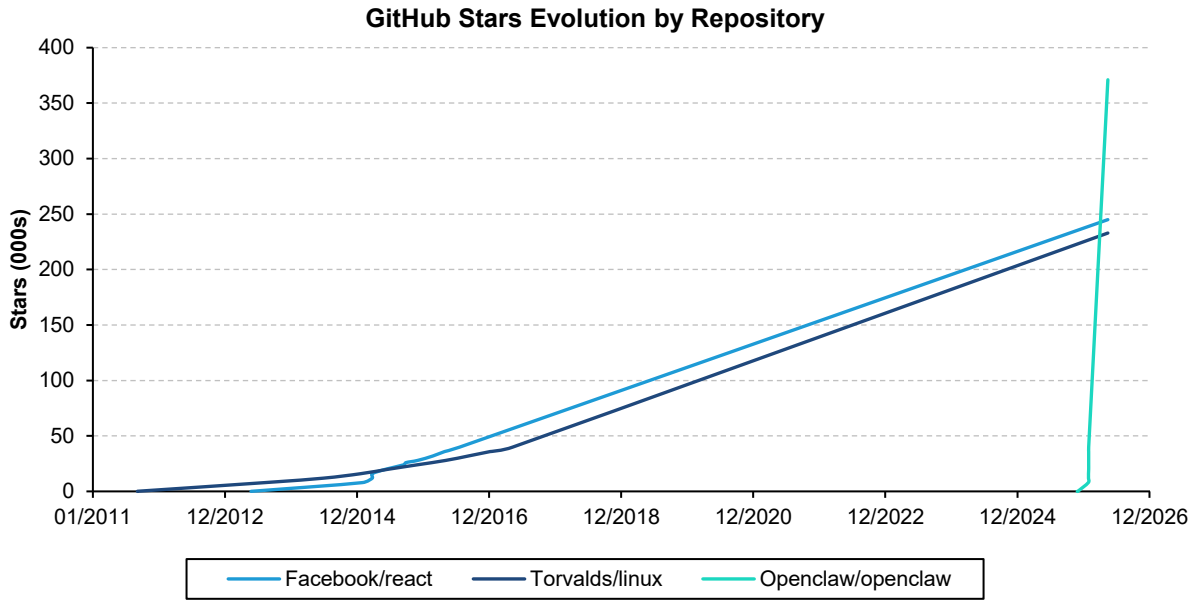


Source: Star-history.com, Bernstein analysis

EXHIBIT 4: **OpenRouter AI token consumption rose ~22x from March 2025 to March 2026, with a clear inflection after OpenClaw's late-November 2025 launch, highlighting how agentic AI workflows are accelerating compute intensity and token demand.**

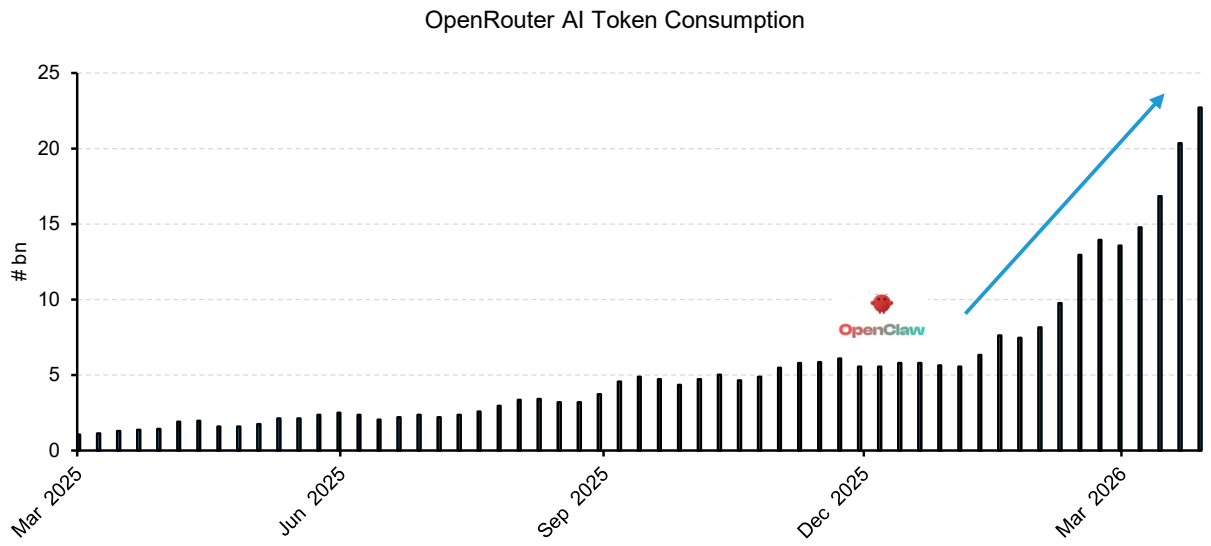


Source: Company reports, Bernstein analysis

EXHIBIT 5: **Agentic ai is delivering measurable enterprise impact across functions, driving productivity gains, revenue uplift, marketing improvement and cost reductions in knowledge work, finance, healthcare, customer service, HR and operations globally**

| Industry | Impact of agentic AI |
|---|---|
| **Knowledge Work** | 33 % per hour productivity gain |
| **Banking & Finance** | $200–340 B nrevenue uplift |
| **Pharma & Medical Products** | $60–110 Bn revenue uplift |
| **Sales & Marketing** | 5–15 % marketing gain |
| **Customer Service** | 30 % cost reduction |
| **Human Resources** | 30 % productivity gain and 10 % cost reduction |

Source: Ahdus Technology, Bernstein analysis

EXHIBIT 6: **Agentic ai is unlocking significant economic value at scale, expanding from early gains to a projected $450 billion opportunity, highlighting strong enterprise potential and accelerating business impact across industries.**

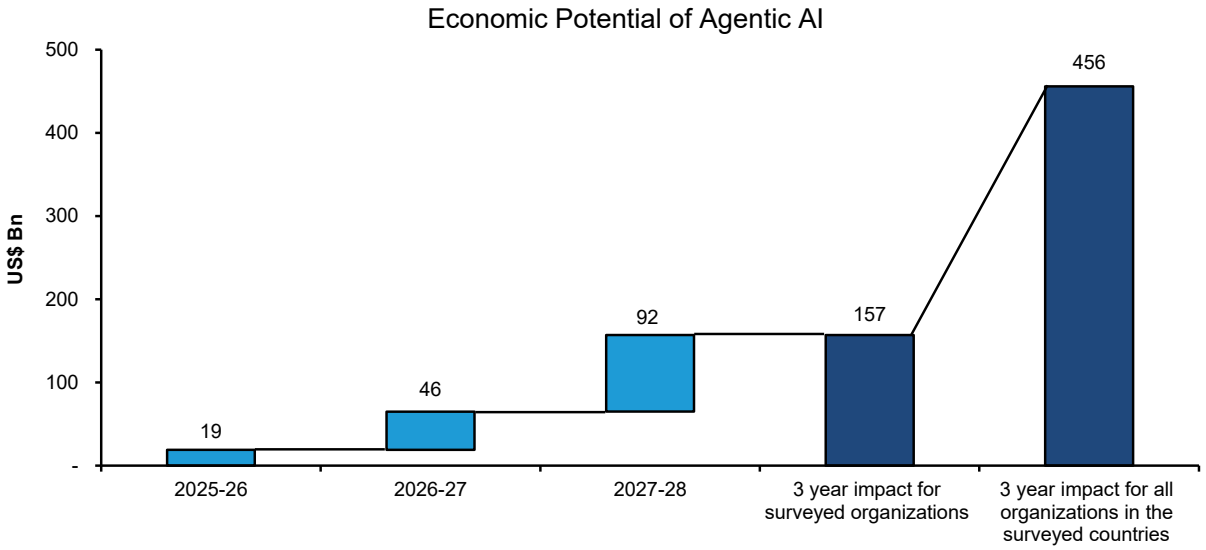


Source: Capgemini Research Institute Agentic AI, Bernstein analysis

EXHIBIT 7: **The global agentic AI market size is predicted to increase from USD 7.55Bn in 2025 to approximately USD 199.05Bn by 2034, expanding at a CAGR of ~45% from 2025 to 2034.**

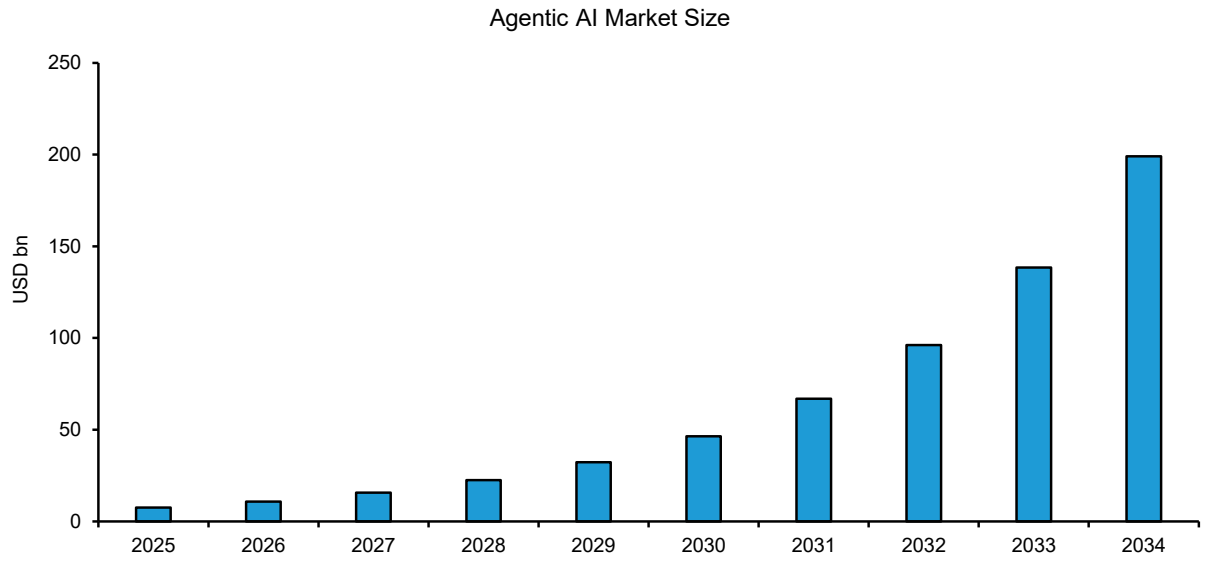


Source: Precedence research, Bernstein analysis

## AGENTIC AI LEADS TO PHYSICAL AI AND EVENTUALLY AGI

Physical AI - robots and autonomous systems powered by LLMs and agentic reasoning - adds the dimension that pure software agents cannot: **operating in an unstructured, unpredictable, physical world**. This is significant for the AGI argument because real-world grounding has long been considered one of the hardest unsolved problems in AI.

Physical AI agents are built around three core blocks that mirror the cognitive architecture of agentic AI: **perception** (sensors, computer vision), **cognition** (LLM-driven reasoning), and **actuation** (physical manipulation and movement). This isn't just robotics - it is agentic AI given a body. While agentic AI learns to reason and act in the software world, Physical AI learns to reason and act **in the real world** - navigating uncertainty, handling failure, and adapting to environments no dataset could fully anticipate. This real-world feedback loop is an irreplaceable training ground for AGI-level generalization.

The entire AI ecosystem is moving towards physical AI as a next step. NVIDIA, Alphabet, and Google are converging on shared platforms (Omniverse, Cosmos, Isaac) that train AI agents in digital twin simulations and deploy them into physical systems - directly closing the sim-to-real gap. Tesla and BMW are already using agentic frameworks to enable robots that self-improve through autonomous learning cycles, minimizing human intervention in manufacturing. China's government has explicitly framed embodied AI (physical AI) as its strategic path to AGI, recognizing that physical grounding produces capabilities that pure language models cannot. Jensen Huang's vision at NVIDIA centers on the convergence of agentic and physical AI as the defining industrial transformation of the decade.

### Agentic and physical AI are bridges to the AGI

The capabilities agentic AI forces AI systems to develop are precisely the capabilities needed for general intelligence. This is not coincidental - it is architectural. Gen AI 1.0 showed us what AI could **say**. Gen AI 2.0 - agentic AI - shows us what AI can **do** in the digital world. Gen AI 3.0 - physical AI - shows us what AI can **be** in the real world. Together, they are converging on the only thing left: a system that can do *all of it*, in any context, without being told how. That is Gen AI 4.0, or AGI, the ultimate form of AI.

Multi-agent systems already command 53.3% of the agentic AI market in 2025, signaling that collaborative AI architectures - a prerequisite for AGI-level coordination - are maturing fast. IDC describes agentic AI as a "keystone technology" for the coming digital labor workforce, expecting it to be embedded in most enterprise operations within a decade.

Physical AI fills the grounding gap that purely digital agents cannot. AGI must understand that a glass will shatter if dropped - not because it read about it, but because it has *acted* in the world. Embodied AI provides exactly this. Research toward "Embodied AGI" already defines a five-level taxonomy (L1-L5) analogous to autonomous driving levels, with L5 being human level proficiency across open-ended real-world tasks.

AGI isn't a single eureka moment - it's the convergence of agent-level capabilities refined at production scale (Exhibit 9).

EXHIBIT 8: **Agentic tasks are uniquely expensive, consuming up to 1,000x more tokens than code reasoning or code chat. This is mainly driven by high input-token usage, making context handling the key cost driver.**

| | Gen AI 1.0 | Gen AI 2.0 | Gen AI 3.0 | Gen AI 4.0 |
|---|---|---|---|---|
| **Paradigm** | **Chatbots** | **Agentic AI** | **Physical AI** | **AGI** |
| **AI Capabilities** | Respond to prompts, advises | Plans, acts, iterates autonomously (digital world) | Plans, acts, iterates autonomously (physical world) | Generalizes across any domain |
| **Human Role** | Interpret & act | Sets goals & reviews | Sets goals & reviews | Set Intent |
| **Human: AI ratio** | 1:0.1 | 1:1 to several | 1:1 to several | ∞ |
| **Token Consumpt** | 1x | 1,000x | 1,000x | ∞ |

Source: "How Do AI Agents Spend Your Money? Analyzing and Predicting Token Consumption in Agentic Coding Tasks", 2026, Longju Bai et al., Bernstein analysis

EXHIBIT 9: **From communication to execution to embodiment: the Gen AI roadmap toward enterprise-scale general intelligence**

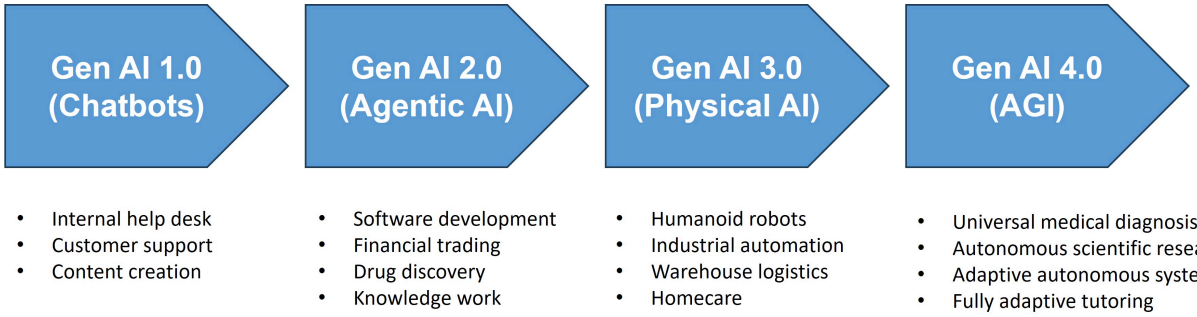


Source: Bernstein analysis

## SOFTBANK INVESTED IN THE FULL STACK OF AI

If the Gen AI 1.0 → 2.0 → 3.0 → 4.0 framework is correct, then **SoftBank is a single investment position to own the entire stack** that each generation runs on. SoftBank has deliberately executed towards this vision. Masayoshi Son has structured SoftBank not as a diversified tech investor, but as a vertically integrated AI holding company - the only entity attempting to capture the entire value chain from silicon to superintelligence.

### The full-stack Gen AI map for SoftBank

SoftBank has invested in every step necessary towards the Gen AI 1.0 → 4.0 roadmap, including (Exhibit 10):

- At the **silicon layer**, SoftBank controls the foundational compute substrate through its ~87% stake in Arm, whose architecture powers 99% of smartphones and increasingly underpins AI inference at the edge and in data centers, alongside Ampere and Graphcore for specialized AI silicon.
- At the **infrastructure layer**, SoftBank has secured the physical backbone of AI through SB Energy (power generation), DigitalBridge (data centers), and most visibly, Stargate, the $500bn data center initiative co-anchored with OpenAI and Oracle, to serve as the sovereign AI compute platform for the U.S.
- At the **model and agent layer**, SoftBank holds an 11% stake in OpenAI with ~$45bn invested, with a further $20bn committed by October 2026, at which point SoftBank's ownership will rise to 13%, and $45bn in unrealized gains, a direct equity claim on the leading frontier model developer and the defining platform for Gen AI 1.0 (chatbots) and 2.0 (autonomous agents).
- At the **physical AI layer**, SoftBank is positioning for Gen AI 3.0, embodied intelligence, through SoftBank Robotics (Pepper, NAO) and the pending $5.4bn acquisition of ABB Robotics, which, upon close, would give SoftBank one of the world's largest industrial robotics businesses as the interface between AI models and the physical world.

EXHIBIT 10: **SoftBank's Investments in the AI ecosystem**

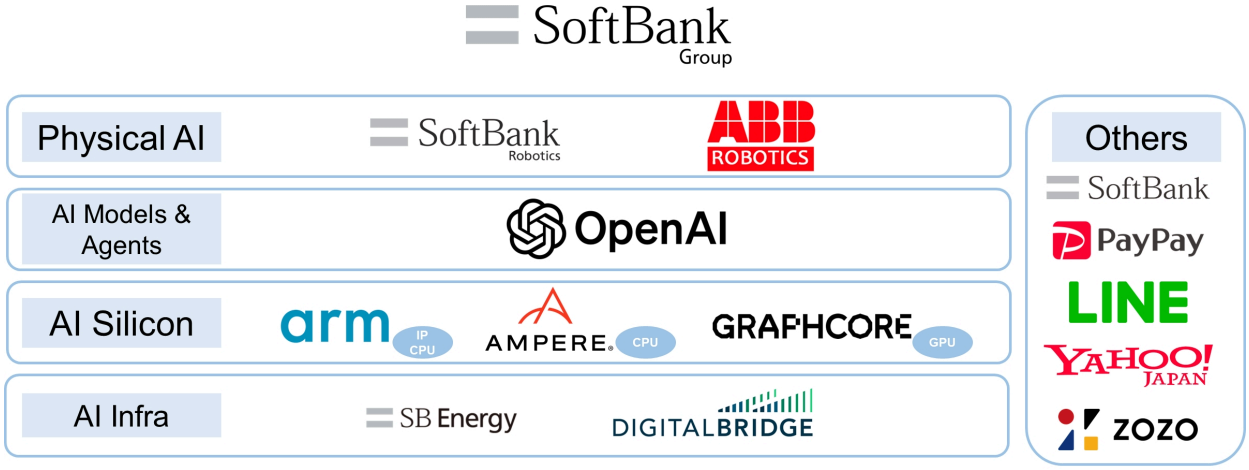


"'SoftBank' within the 'Others' category refers to SoftBank Corp., the Group's telecommunications business (listed separately as 9434.JP; not covered).
Source: Company disclosures, Bernstein analysis.

**The paradigm shift to Agentic AI structurally benefits the full stack of SoftBank investment**

- Agentic AI drove AI adoption and revenue growth of OpenAI, whose revenue ARR grew more than 10x to $25bn in March 2026 vs. March 2024. In May 2026, **OpenAI Codex** has experienced a "ChatGPT moment" by evolving from a simple code autocompletion tool into a powerful, autonomous software engineering agent powered by **GPT-5.5** (Exhibit 11). Thanks to Agentic AI, OpenAI's revenue is accelerating to $25bn in Feb 2026.

- Because of Agentic AI, CPUs are back to the center stage. Compared to Gen AI 1.0 (chatbot), Agentic AI involves heavily autonomous task orchestration and execution, which can only be done by CPUs. The ratio of GPU:CPU is shifting from the current 8:1 to 1:1, making CPU's value growing at least 5-fold by 2030. Arm stands out in server CPUs given its unparalleled power efficiency.

- Increased token consumption and AI adoption means more AI infrastructure build out, which is what **SB Energy** is for. The Stargate project is just the first major project.

- Finally, SoftBank Robotics is well positioned for the future rise of Physical AI, especially with the investments in ABB Robotics.

The thesis is already generating real returns for SoftBank (Exhibit 12):

- SoftBank's NAV has grown 64% since the start of investment in OpenAI in 3Q CY2024.

- SoftBank's OpenAI position alone generated ~$45bn investment gains for the company.

- Arm position has generated ~$180bn investment gains since 2016 (or $94bn since 3Q CY2024) for SoftBank, driven by its leadership in AI-optimized chip architectures. Arm is well positioned to capture incremental market share as agentic AI deployments drive demand for more efficient, low-power inference at the edge, a key use case where Arm's power-efficient architecture continues to dominate.

- Going forward, we expect to see more returns not only from OpenAI and Arm but also from other assets (chips, infrastructure, and physical AI).

EXHIBIT 11: **OpenAI Codex's weekly downloads surged to 86.1mn on 3rd May, highlighting a sharp acceleration in developer adoption.**

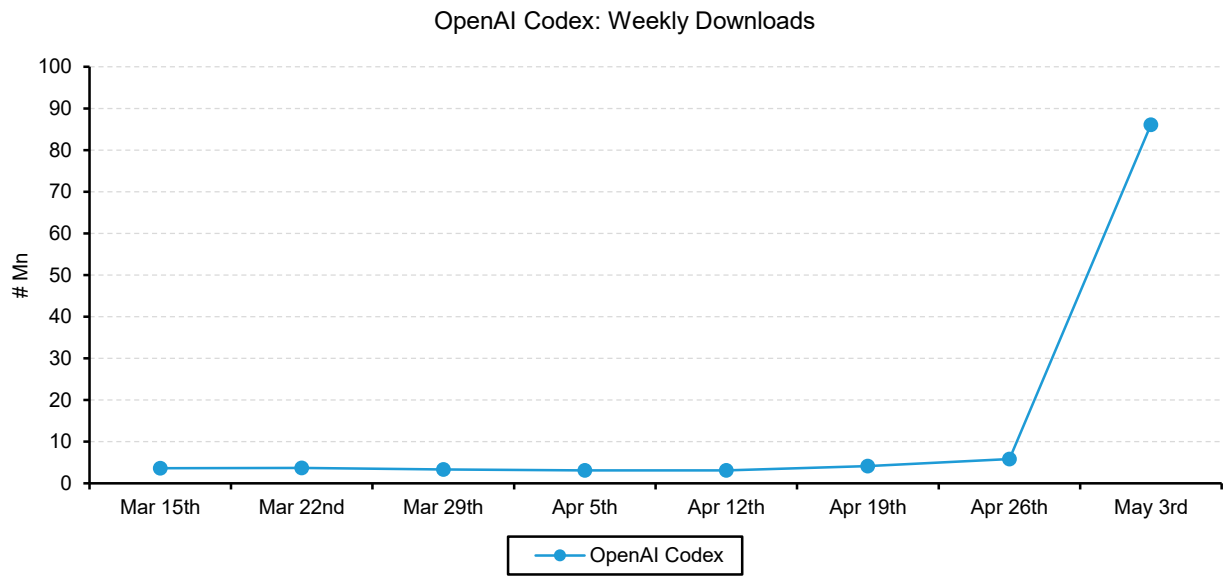


Source: TickerTrends NPM, Bernstein analysis.

EXHIBIT 12: **SoftBank's NAV Growth since First Investment in OpenAI**

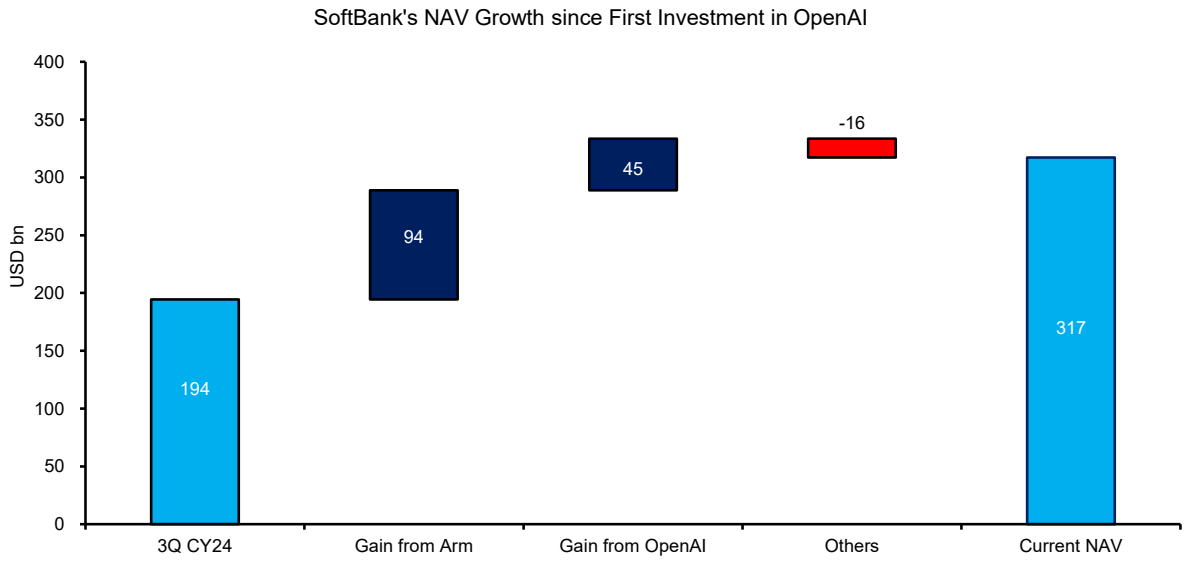


Based on share price as of 13th May 2026.
Source: Company disclosures, Bloomberg, Bernstein analysis.

SoftBank is now trading at 27.5% discount to its NAV. We expect SoftBank's NAV growing to $390bn in a year, mostly contributing by Arm ($274bn; based on our price target of Arm $300.00) and OpenAI ($92bn, based on valuation in the latest funding round). With a 25% NAV discount, we rate SoftBank Outperform with PT = ¥8,200.00. Our PT implies 43% upside to SoftBank's current share price (Exhibit 13).

In our bull case, we apply our bull-case Arm valuation ($390 per share) and OpenAI (still based on valuation in the latest funding round). With a 15% NAV discount, which implies valuation of ¥11,200, while in our bear case, we apply our Arm valuation based

on its share price of $221 and same valuation of OpenAI). With a 45% NAV discount, which implies a valuation of ¥4,900.

Given SoftBank's strategic positioning across the AI value chain, we believe the company will be one of the largest long-term beneficiaries of the AI ecosystem within the Japan semiconductor sector. Moreover, if OpenAI goes public at a higher valuation (*OpenAI Plans Fourth-Quarter IPO in Race to Beat Anthropic to Market - WSJ*), it would represent additional upside for SoftBank. With our PT of **¥8,200.00**, we initiate coverage with an **Outperform** rating.

Our SoftBank financial model can be downloaded here: SoftBank (9984.JP)

EXHIBIT 13: **SoftBank Valuation: Scenario Analysis**

| **Softbank's 1-year NAV to valuation (USDbn)** | | | |
|---|---|---|---|
| | **Bull** | **Base** | **Bear** |
| Arm per share value (USD) | 390.00 | 300.00 | 221.21 |
| Share Count (mn) | 1,052 | 1,052 | 1,052 |
| Market cap (USD bn) | 410,280 | 315,600 | 232,713 |
| % holding | 86.9% | 86.9% | 86.9% |
| | | | |
| OpenAI Valuation (USD bn) | 852 | 852 | 852 |
| % holding | 10.7% | 10.7% | 10.7% |
| | | | |
| **Future NAV (USD bn)** | | | |
| Arm | 356.5 | 274.2 | 202.2 |
| OpenAI | 91.6 | 91.6 | 91.6 |
| Others | 23.7 | 23.7 | 23.7 |
| **Total** | **471.7** | **389.5** | **317.4** |
| *Total (JPY tn)* | *75.1* | *62.0* | *50.5* |
| | | | |
| Discount | 15% | 25% | 45% |
| | | | |
| Share Count (mn) | 5,699 | 5,699 | 5,699 |
| Target price | **11,200** | **8,200** | **4,900** |
| Upside/downside | 94% | 42% | (15%) |

Share price as of 13th May 2026.
Source: Company disclosures, Bloomberg, Bernstein estimates and analysis.

## ARM AT THE CENTER OF CPU RENAISSANCE, DRIVEN BY AGENTIC AI

Agentic AI doesn't just require more GPUs - it structurally increases CPU demand in data centers in a way that Gen AI 1.0 never did. This distinction is critical, and it is the single most important reason Arm's data center opportunity is being repriced upward.

Gen AI 1.0 (chatbots, inference) was GPU-dominant. A user sends a prompt, and GPU executes matrix multiplications and generates output tokens. The CPU was a traffic manager, barely involved in the calculation.

Gen AI 2.0 (agentic AI) breaks this model entirely. An agent doesn't execute one forward pass - it plans, calls tools, reads and writes files, evaluates outputs, retries, coordinates with other agents, and iterates. Every one of those steps requires CPU orchestration (Exhibit 14).

The numbers are staggering:

- Traditional AI data centers require ~**30 million CPU cores per gigawatt** of compute
- Agentic AI data centers require ~**120 million CPU cores per gigawatt** - a **4x structural increase**

As a result, CPU-to-GPU ratios are shifting from the current **1:4-1:8 → to 1:1-1:2** in agentic AI deployments.

**In the age of agentic AI, CPUs matter more than ever** - scaling is no longer just about accelerators, but about efficient orchestration compute. Agentic AI has permanently expanded the CPU TAM inside data centers. This is a structural shift, not a cyclical one - it follows from the architecture of how agents work, not from any single product cycle.

### Arm wins this structural shift

Not all CPUs benefit equally. The agentic AI data center CPU boom rewards **efficiency-first architectures** - and that is Arm's defining advantage.

Arm's Neoverse CPUs deliver **industry-leading performance-per-watt**, reducing energy costs as AI racks scale to hundreds of megawatts. As data centers are power-constrained, not just compute-constrained - scaling AI now requires *more output per watt*, not just more watts. This calls for higher penetration of Arm CPUs.

As an evidence, Arm-based server shipments grew **70% in 2025**, with IDC projecting Arm reaching **21.1% of total server shipments** - up from ~5% just three years prior. Arm CPU penetration is said to have already surpassed 50% in data centers at the top hyperscalers. AWS Graviton4 (Arm-based) is already the **dominant custom silicon choice for cloud AI workloads** at hyperscalers.

The hyperscaler adoption curve tells the story clearly. Every major hyperscaler has independently converged on Arm architecture for their AI infrastructure. This is not a trend — it is a **de facto standard forming in real time**.

### Arm's Next Move: Capturing More of the Value Chain

Arm has historically been a **royalty and licensing business** - it designs architectures and earns a cut of every chip its partners sell. But now it is making a decisive move up the value chain: Arm is building its **own production silicon** for the first time.

In March 2026, Arm launched the **"Arm AGI CPU"** - its first proprietary data center chip, explicitly **built for agentic AI workloads**, designed to deliver leading performance per rack. This extends Arm from being an IP licensor into a **systems-level compute provider**, capturing margin that previously went to chip manufacturers. With this move, Arm is set to capture more value: First, it sold the architecture that everyone else built on. Now, as agentic AI expands the CPU TAM by 4x, Arm is positioning to capture not just the royalty stream but the direct silicon margin too.

### Arm is also at the core of edge and physical AI

In addition to the data center AI, which is the core investment thesis, Arm is also at the center stage for edge AI and physical AI. For mobile, application processors are evolving at faster rate than before because of the penetration of AI smartphones, which requires higher performance, low power, and customizations for AI applications. Arm’s new architectures, especially the Compute Subsystems (CSS), are essential to help the customers embrace the edge AI (smartphones, consumer electronics) and physical AI (automotive, humanoid) era. In return, Arm also benefits from the rising royalty rate, from the current ~3% in FY26 to ~5% by FY31.

EXHIBIT 14: **Agentic AI shifts more work back to the CPU: accelerators generate tokens, but CPUs orchestrate the agents, memory and workflows needed to deliver answers, making Arm's efficient CPU architecture increasingly critical in AI data centers**

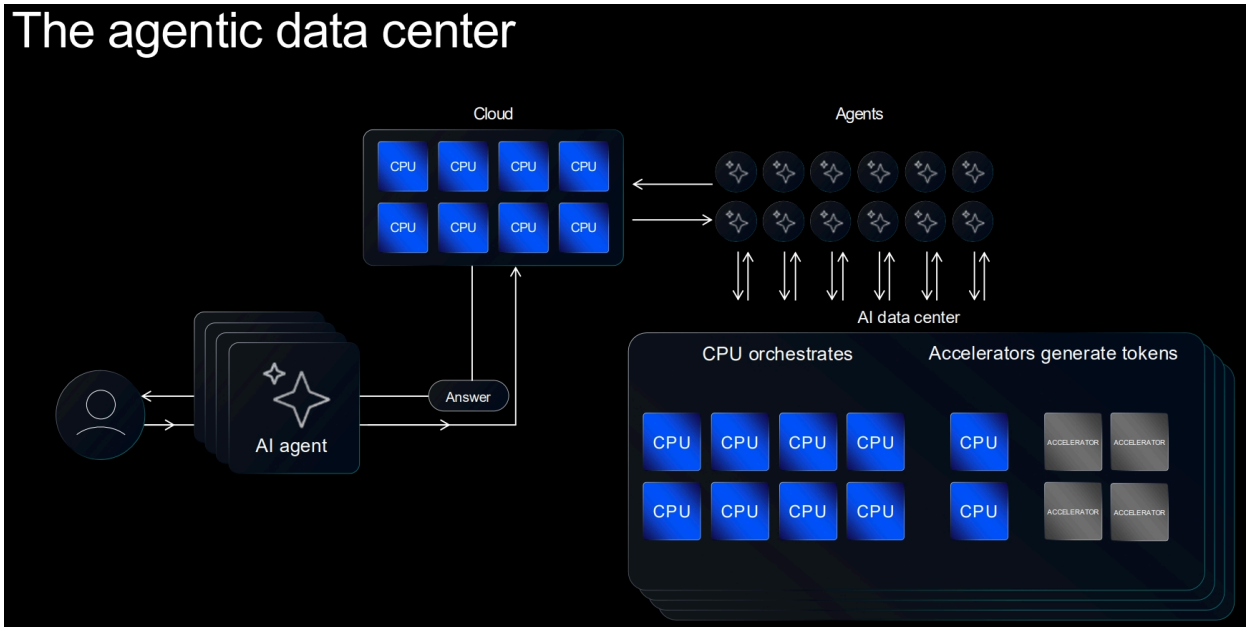


Source: Arm

EXHIBIT 15: **Agentic AI shifts compute balance toward CPUs, with CPU share rising from ~14% in Traditional LLMs to 50%, highlighting CPUs' growing orchestration role alongside GPUs in AI workloads at scale**

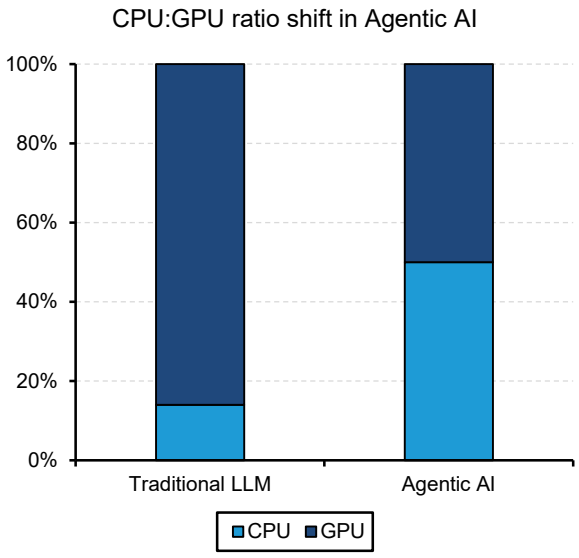


Source: TrendForce, Bernstein analysis

EXHIBIT 16: **AI infrastructure TAM expands sharply by CY30, led by data center accelerators reaching $1T, while data center CPU also quadruples from $33bn to $137bn.**

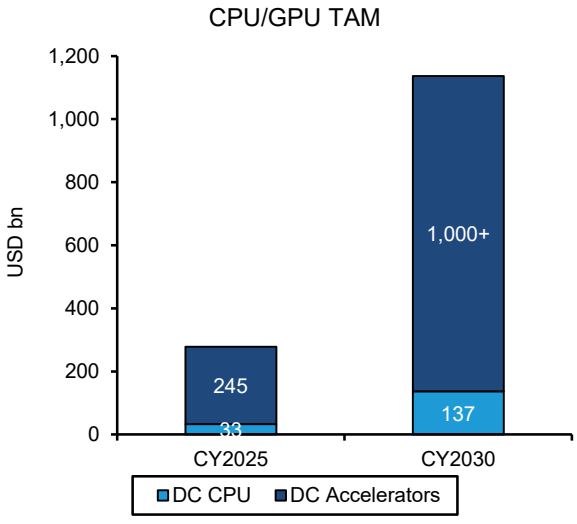


Source: Mercury, Company reports, Bernstein estimates and analysis.

EXHIBIT 17: **Arm Server CPU has seen strong growth since 2023...**

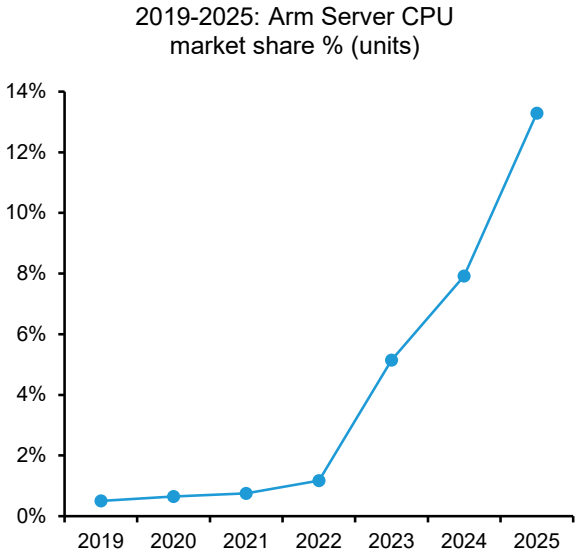


Source: Mercury estimates, Bernstein analysis.

EXHIBIT 18: **We expect the server CPU market to increase from US$33bn in 2025 to US$137bn in 2030, accelerating at a 33% CAGR, driven by Agentic AI adoption.**

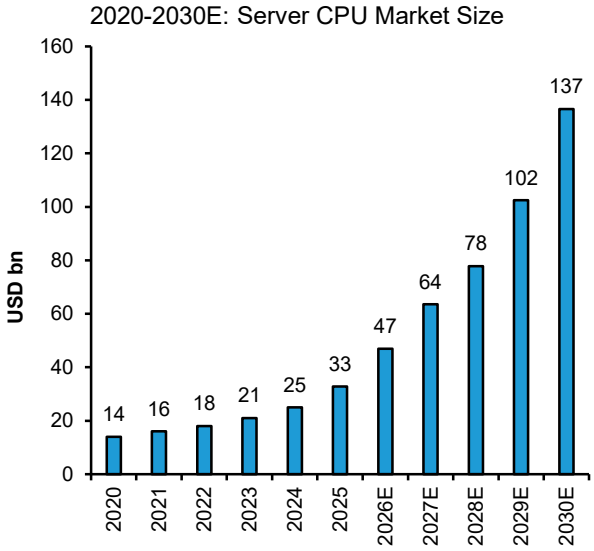


Source: Mercury, Company reports, Bernstein analysis and estimates

## Arm Valuation

Since their IPO in September 2023, Arm has traded at 70x 1-year forward P/E on average. Certainly, part of this is due to Arm's low float % and scarcity value, as SoftBank Group holds 87% of the shares. However, it is worth highlighting that semis IP companies trade, on average, at considerably higher multiples than other companies in the semiconductor sector, with the average FY+1 P/E multiple trading at an impressive ~140x, albeit largely driven by the abnormal multiples of Chinese players. On FY+2 basis, Arm is also below the average. The higher industry multiple can be explained by an asset-light, software-like model with high margins, minimal COGS, recurring royalties, and strong operating leverage.

We believe the high multiple is justified by Arm's position as likely the most reliable growth story in the sector, supported by very strong visibility: ~70% of the royalty forecast is based on existing contracts, while the $15bn AGI CPU estimate is based on actual customer engagement.We expect Arm to deliver a 40% Revenue CAGR and similar EPS growth, with potentially meaningful upside given the massive demand for CPUs driven by Agentic AI. Its monopoly position in the Mobile IP CPU market and rapid market share gains in the server CPU market also support Arm's premium valuation.

We forecast Arm's revenue to grow more than five fold to $26 Bn by 2030 (FY31), with EPS expanding similarly (~5.5x) to $9.83, supported by the growing adoption of Arm CPUs in AI data centers, its own Arm CPU revenue, and the rising royalty driven by higher value provided by Arm for its clients rapid evolution of CPUs. **With 90x P/E on Q5-8 EPS of $3.33, we rate Arm Outperform with PT= $300** (Exhibit 19, Exhibit 20).

**In a bull case scenario** (Exhibit 21), we can see Arm being valued at 40x on FY31E EPS. This 40x multiple is derived from a SOTP analysis, whereby we assign 70x (historical average 1-year forward P/E for Arm) multiple to the traditional IP / CSS business and 30x (historical average 1-year forward P/E for AMD) multiple to the Arm Silicon (CPU) business as AMD would be the most relevant peer for a fabless server CPU business. The blended P/E is 53x, to which we discount back 3 years at WACC of 10%, which yields 40x 1-year forward multiple. **Applying 40x P/E on FY31E EPS of $9.83, we derive the bull-case valuation of $390**.

Our financial model can be downloaded here: Arm (Arm.US). Our ARM Industry models can be downloaded here: ARM Server CPU Industry Model, ARM Automotive TAM Model, ARM Mobile IP TAM Model

EXHIBIT 19: **We expect Arm's revenue to grow at an impressive ~40% CAGR over the next 5 years, reaching US$26Bn in FY31. This growth is primarily driven by the expansion of the new AGI CPU business, which we forecast will reach US$15Bn in FY31, in line with guidance.**

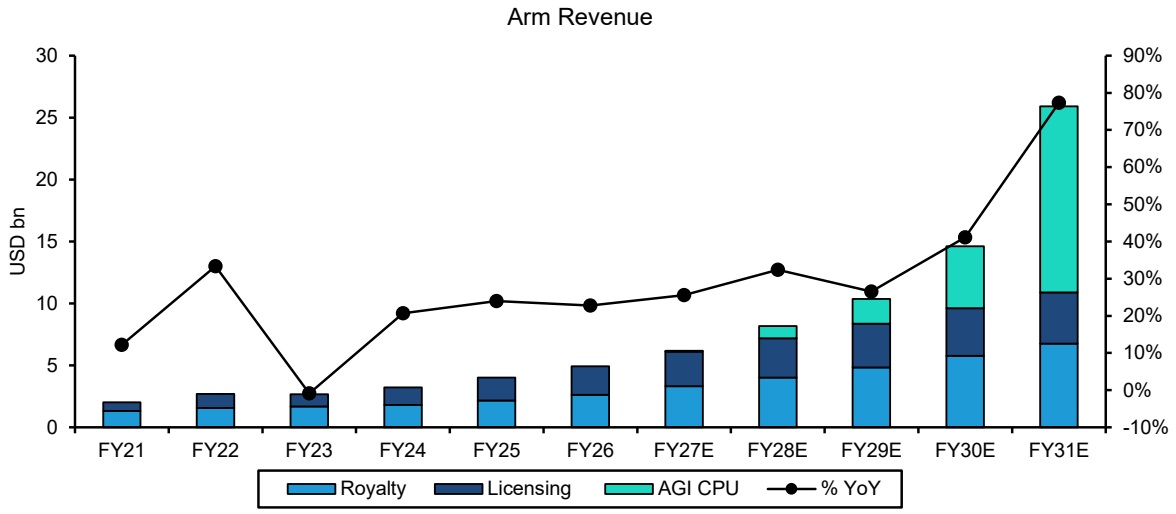


Source: Company reports, Bernstein analysis and estimates

EXHIBIT 20: **Since the IPO in September 2023, Arm has traded at an average 1Y forward P/E of 70x. Although the stock has been highly volatile over the past two and a half years, it is now trading above +1SD of historical range.**

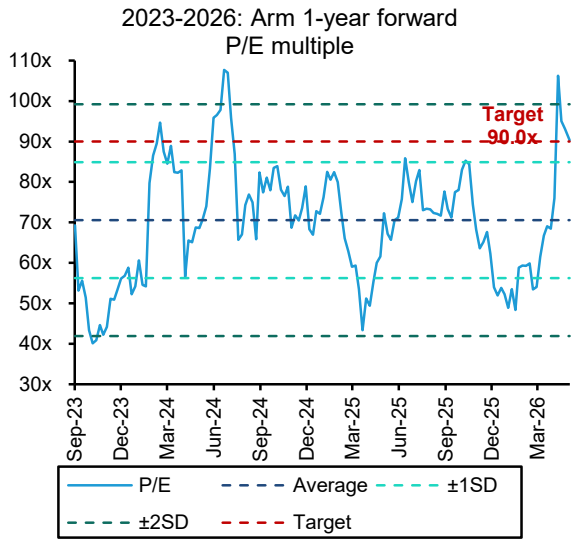


Source: Bloomberg, Bernstein analysis and estimates

EXHIBIT 21: **Our bull case analysis yields an implied target P/E of 40x, and a valuation of $390.**

| **FY31E Financials** | **IP / CSS** | **Arm Silicon** |
|---|---|---|
| Segment OP | 7,122 | 5,253 |
| PF Net Income (USD mn) | 6,054 | 4,465 |
| 1-year Forward P/E | 70x | 30x |
| PF Market Cap (USD mn) | 423,746 | 133,950 |
| | | |
| Arm PF Market Cap | | 557,696 |
| Arm PF Net Income (USD mn) | | 10,519 |
| PF Blended Fwd P/E (FY31E) | | 53.0x |
| Discount Factor (3 years, WACC=10%) | | 1.33 |
| **Implied P/E (Discounted to FY28E)** | | **39.8x** |
| Arm EPS (SCBe) | | $9.83 |
| **Arm Bull Case TP** | | **$390** |

Source: Company disclosures, Bernstein estimates and analysis.

## SOFTBANK: THE FULL AI STACK INVESTOR

Established in 1981 and headquartered in Tokyo, Japan, SoftBank Group (SoftBank) was founded by Masayoshi Son as a domestic software distributor before evolving into one of the world's largest technology investment conglomerates. The company expanded through bold acquisitions, most notably Vodafone Japan (2006), Sprint (2013), and Arm Holdings (2016), while its $20mn early investment on Alibaba (BABA, covered by Robin Zhu) and the landmark $100bn Vision Fund (2017) cemented its reputation as an aggressive, high-conviction technology investor. Today, SoftBank operates as a global AI-focused holding company, with its strategy anchored around Arm Holdings and a growing OpenAI stake.

SoftBank Group stands as one of the most prominent global investors poised to benefit from the accelerating adoption of artificial intelligence across the enterprise landscape (Exhibit 22). With substantial exposure to leading AI infrastructure, semiconductor, and application-layer companies, SoftBank is strategically positioned to capture value as the ecosystem scales. While investor sentiment toward AI remained mixed as recently as a year, or even six months, ago, recent industry developments have clearly signaled a decisive shift: AI is moving beyond pilot programs and isolated proofs of concept into broad enterprise deployment. As organizational access to AI tools expands across the workforce, adoption curves are steepening, monetization opportunities are becoming more tangible, and the companies enabling this transformation, including those backed by SoftBank, are entering a phase of potentially rapid and sustained growth.

EXHIBIT 22: **SoftBank's Investments in the AI ecosystem**

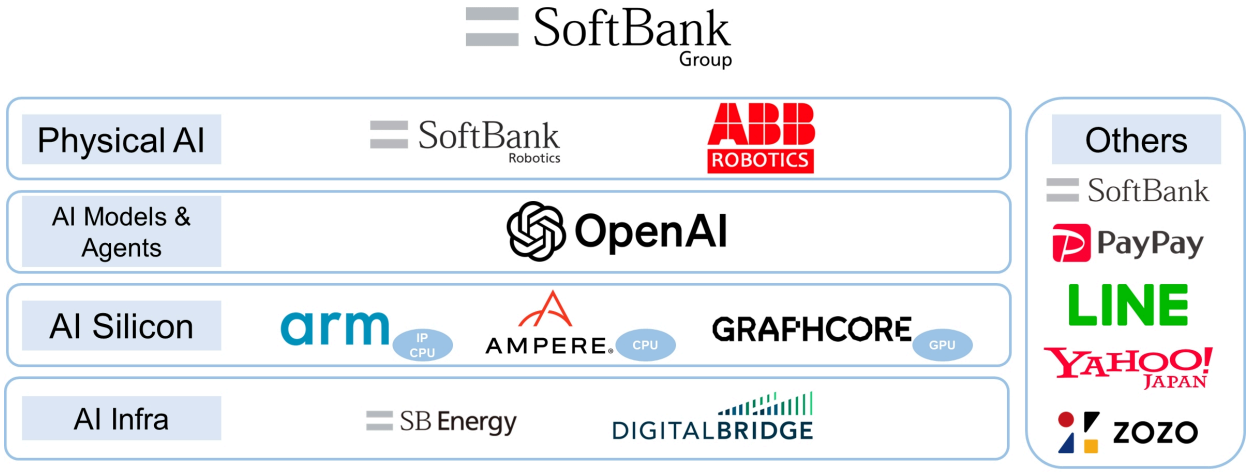


"'SoftBank' within the 'Others' category refers to SoftBank Corp., the Group's telecommunications business (listed separately as 9434.JP; not covered).
Source: Company disclosures, Bernstein analysis.

### STRONG AI ECOSYSTEM CENTERED AROUND OPENAI

OpenAI is SoftBank's second-largest net asset value (NAV) contributor after Arm and is set to become the largest following the closure of the March 2026-announced funding round (31% of NAV as of March 2026, likely rising to ~38% by end of 2026 if the additional $30bn investment closes). OpenAI remains the world's largest AI service provider, and management forecasts revenue to reach $283.5bn in 2030 (Exhibit 23).

Recently, with the strengthening of its agent tool, OpenAI saw an "ChatGPT moment" (Sam Altman) in Codex whose download surged (Exhibit 25). This is because Codex has evolved from a useful coding tool to a broader, faster, more autonomous **work agent** powered by GPT 5.5. When Codex was introduced in May 2025, it was described as a software engineering agent that could work on coding tasks. By April 2026, OpenAI repositioned Codex around GPT-5.5, and built Codex for "real work" and could write and debug code, research online, analyze data, create documents and spreadsheets, operate software, and move across tools until a task is finished. Codex has 3mn (and rapidly growing) paid users, from virtually 0 at the beginning of the year.

The company continues to invest heavily, with management projecting a cumulative cash burn of $665bn before turning cash-flow positive in 2030. However, given that OpenAI has already raised approximately $182bn and is expected to pursue an IPO (*OpenAI Plans Fourth-Quarter IPO in Race to Beat Anthropic to Market - WSJ*), we see limited risk around funding or liquidity (Exhibit 24).

### SoftBank's investments in OpenAI

SoftBank began investing in OpenAI in 2024 with an initial $500mn and has progressively increased its commitment. In the latest funding round announced in March 2026, OpenAI raised approximately $120bn, with SoftBank committing an additional $30bn to be deployed in three $10bn tranches on 1st April, 1st July, and 1st October 2026. The round values OpenAI at $852bn post-money (pre-money $730bn), bringing SoftBank's cumulative investment to $64.6bn and increasing its ownership to 13% (from 11% pre-money) (Exhibit 27). To finance these follow-on investments, SoftBank has also signed a $40bn bridge facility with a consortium of banks (JPMorgan Chase, Goldman Sachs, Mizuho, Sumitomo Mitsui Banking Corporation, and Mitsubishi UFJ Financial Group).

### OpenAI's high growth potential with valuation likely surpassing $1tn amid IPO

Within OpenAI's management forecast of $280bn+ in revenue for 2030, more than half is expected to come from the ChatGPT consumer business (~$150bn), with an additional $70bn from ChatGPT B2B and $45-50bn from platform services and new products/hardware. This compares with today's mix, where roughly two-thirds of revenue is derived from the consumer segment. OpenAI's valuation has continued to rise accordingly: the most recently announced funding round (March 2026) values the company at $852bn post-money (Exhibit 26), and market expectations point to a potential valuation exceeding $1tn at the anticipated IPO *(OpenAI lays groundwork for IPO at up to $1 trillion valuation | Reuters)*.

EXHIBIT 23: **OpenAI Revenue: Management Projection**

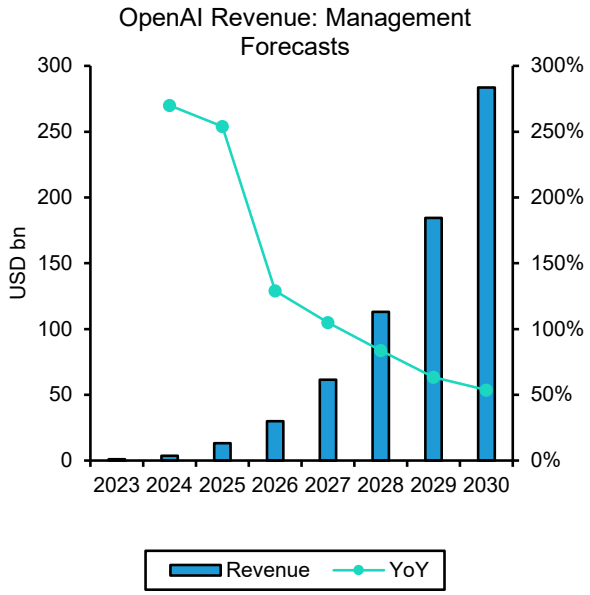


Source: The Information, Bernstein analysis.

EXHIBIT 24: **OpenAI: Cash Raised vs. Cash Flow**

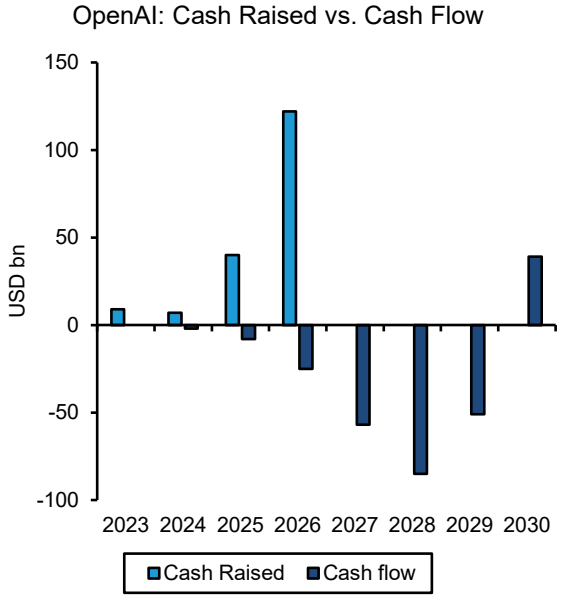


Only includes confirmed financing rounds only; excludes potential IPO proceeds.
Source: The Information, Bernstein analysis.

EXHIBIT 25: **OpenAI Codex weekly download surged to 86.1mn, highlighting a sharp acceleration in developer adoption.**

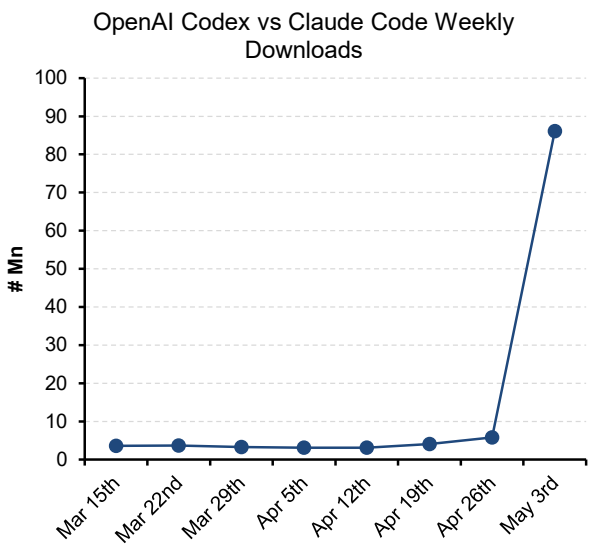


Source: TickerTrends NPM, Bernstein analysis

EXHIBIT 26: **OpenAI's Valuation Trend**

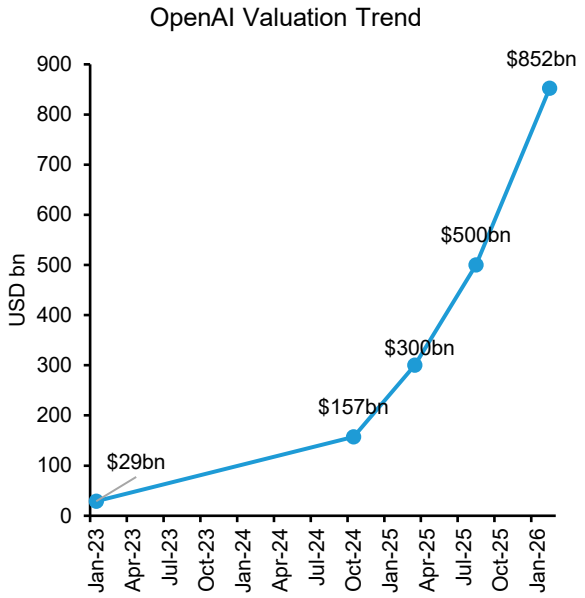


Post-money valuation in each funding round.
Source: Company disclosures, Bernstein analysis.

EXHIBIT 27: **SoftBank's Investments in OpenAI, reaching USD64.6bn by October 2026.**

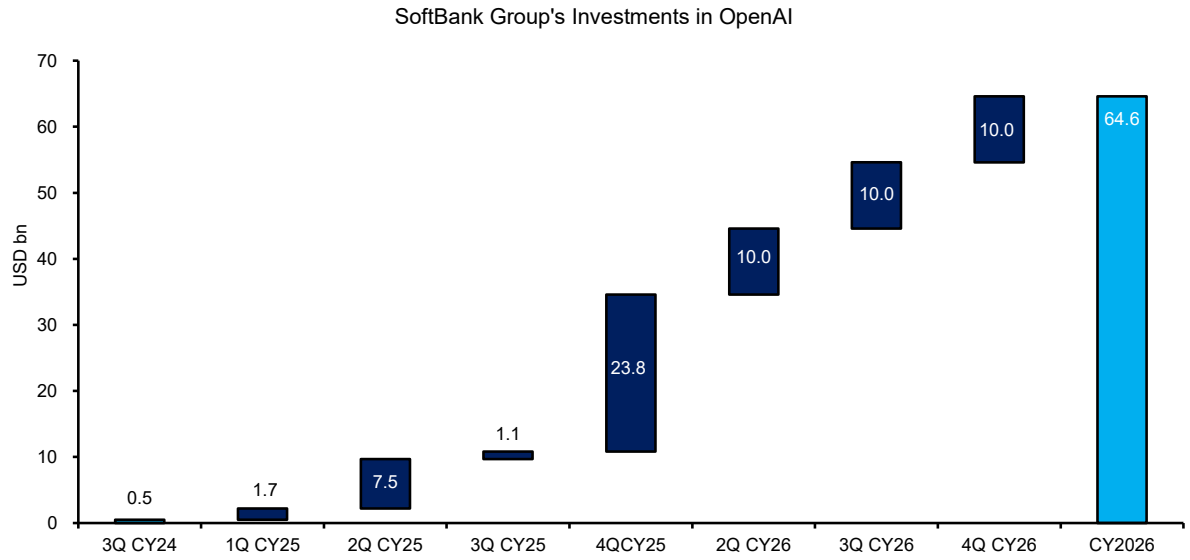


Confirmed investments only; excludes potential IPO investments.
Source: Company disclosures, Bernstein analysis.

**Enterprise adoption of AI is starting to move rapidly, which pushes up enterprise subscription value**

Enterprise adoption of AI is moving from initial trials to broader deployment as employee access expands. According to Deloitte, workforce access to approved AI tools has risen from under 40% to nearly 60% over the past year. Among leading firms, around 11% now provide 80+% of employees with access. However, actual usage still lags, with fewer than 60% of employees who have access using AI in their daily workflows. This indicates that although access is widening, day to day integration remains

limited and much of the productivity and innovation potential is still untapped.

Scaling is beginning to accelerate. Moving AI from trial to full production is essential for unlocking value, yet many firms slow at this stage. Currently, 25% of surveyed companies have moved 40+% of their AI trials into production. Importantly, 54% expect to reach that threshold within the next 3 to 6 months, suggesting that the path to meaningful scale is becoming clearer (Exhibit 28). Early adopters are already demonstrating a faster shift from pilots to enterprise level implementation.

**Agentic AI is expected to fundamentally change how we work**

Agentic AI is set to fundamentally reshape workplace productivity as it shifts from simple conversational tools to systems that execute real day-to-day tasks. Unlike the early ChatGPT applications that focused on text interactions, agentic AI can autonomously perform multistep workflows, driving materially higher productivity gains. This transition also changes the interaction model from human to AI to machine to AI, significantly increasing token consumption as automated agents work continuously to complete tasks. Openrouters has already seen a sharp rise in token usage year to date (weekly token volume surpassed 13tn by mid-February 2026, a more than 10x growth vs. previous year), driven by demand for agentic frameworks such as OpenClaw. Enterprise adoption is poised to accelerate, with 74% of companies planning to implement agentic AI within the next 2 years, vs. 23% today, underscoring the early but rapid shift toward autonomous AI systems (Exhibit 29 and Exhibit 30).

EXHIBIT 28: **Corporate Survey: AI Experiment Moving into Production**

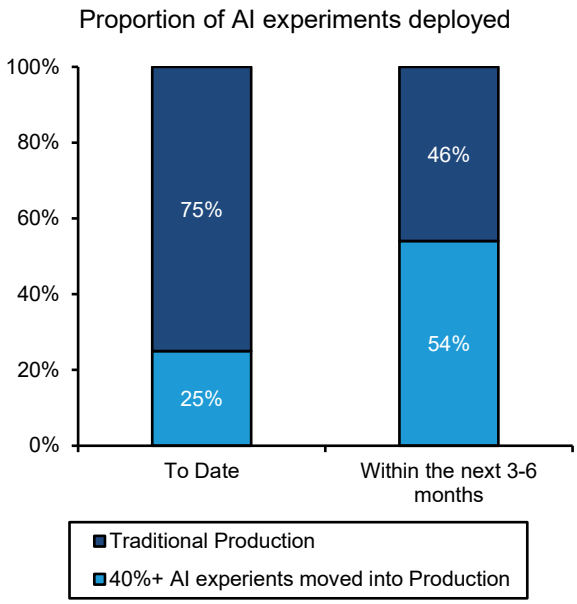


Survey Question: In your estimation, what percentage of your AI experiments (e.g., pilots, test cases, etc.) have been deployed to date into your organization (moved into production)? N=3,235; fieldwork conducted August-September 2025.
Source: Deloitte, Bernstein analysis.

EXHIBIT 29: **Corporate Survey: Extent of Agentic AI Usage**

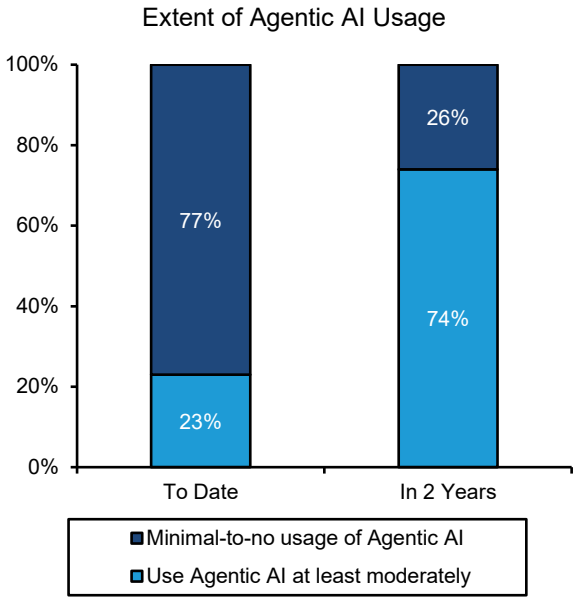


Survey Question: To what extent is your organization utilizing agentic AI in its operations? N=3,235; fieldwork conducted August-September 2025.
Source: Deloitte, Bernstein analysis.

EXHIBIT 30: **Corporate Survey: Current Approach to Transformation with AI**

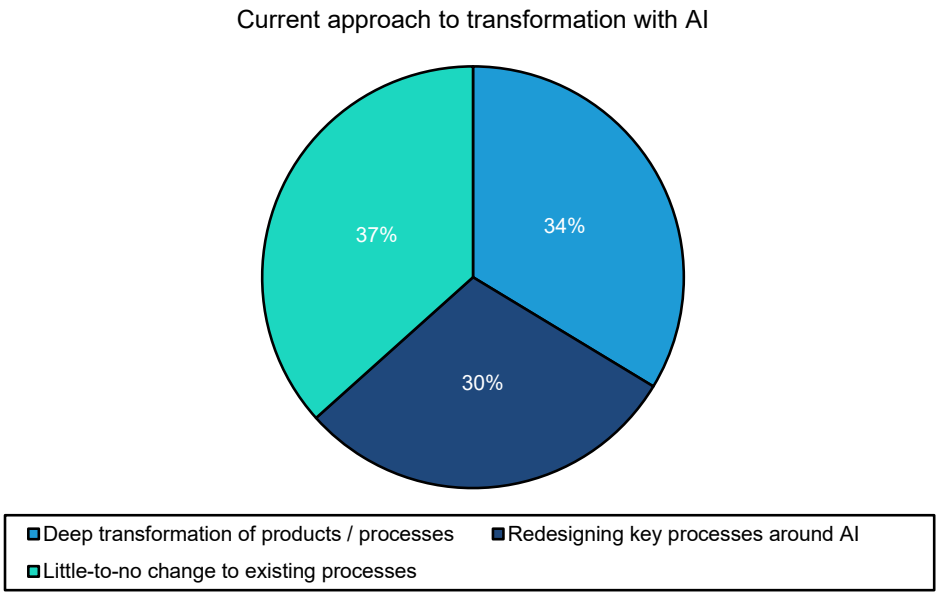


Survey Question: Which best describes your organization's current approach to process transformation during AI implementation? N=3,235; fieldwork conducted August-September 2025.
Source: Deloitte, Bernstein analysis.

**Consumer Monetization: Subscription and Advertising**

OpenAI's consumer monetization strategy is centered on a combination of subscriptions and advertising, supported by a rapidly expanding user base. As of March 2026, OpenAI reported 900+mn weekly active users across consumer AI products, with 50mn+ paid subscribers, implying a ~5% subscription ratio. User engagement continues to deepen, with search usage nearly 3x higher than 2025 levels. At the same time, OpenAI has begun testing advertising, with management indicating that the pilot generated $100mn+ in annual recurring revenue within six weeks (February-March 2026), highlighting a compelling new monetization lever for a platform built largely without ads to date.
ChatGPT currently offers a tiered pricing structure, including $8/month for the "Go" plan, $20/month for the "Plus" plan, and $200/month for the "Pro" plan, while enterprise offerings are contract-based (typically $25-30/month per seat). Based on the current mix of paid users and early advertising contribution, we estimate OpenAI's blended active revenue per user (ARPU) at ~$20-30/month today. Looking ahead, management expects weekly active users to reach 2.75bn and paid subscribers to exceed 220mn by 2030, implying an ~8% subscription ratio and ARPU of ~$57/month, driven by higher-value subscriptions and scaled advertising.

**Competitive Context and Paid Subscription Ratio Upside**

ChatGPT remains the clear leader in consumer AI adoption, with 900+mn users, exceeding peers by a wide margin. Google Gemini follows with 750+mn+ users, benefiting from Google (owned by Alphabet [GOOGL], covered by Mark Shmulik)'s ecosystem, although paid Gemini Advanced is bundled within Google One (~150mn users) and paid subscriber disclosure is limited. Other AI platforms remain smaller, with Perplexity AI at 45+mn users (~4% paid penetration) and Claude at 30+mn users (5-10% paid penetration).

That said, Claude has shown the fastest paid-subscriber growth in 2026, driven in part by the so-called "OpenClaw" usage trend, despite its smaller installed base. However, in early April 2026, Anthropic announced that Claude subscriptions would no longer support OpenClaw usage, citing "outsized strain" on system resources. Since this change, user commentary suggests some OpenClaw-driven demand may migrate back to OpenAI (Exhibit 31).

When benchmarked against mature consumer subscription models, OpenAI's long-term targets (~8%) appear reasonable. Spotify and Amazon Prime have achieved 39% and 80% subscription penetration, respectively, while LinkedIn Premium stands at ~15% (Exhibit 32). Against this backdrop, OpenAI's goal of an ~8% paid subscription ratio by 2030 does not appear

demanding, particularly given ChatGPT's utility-driven use cases and expanding monetization toolkit.

EXHIBIT 31: **AI Services User Base by Provider**

| AI Services | Total Users / MAU (mn) | Paid Subscribers (mn) | Subscription Ratio |
|---|---|---|---|
| ChatGPT | 900+ | 50+ | 5% |
| Gemini | 750+ | n.a. | n.a. |
| Perplexity | 100+ | 4+ | 4% |
| Claude | 30+ | 3-6 | 5-10% |

Gemini is bundled within Google One cloud services; standalone paid-user data is not disclosed. Data of Claude are estimated by Bernstein.
Source: Company disclosures, Bernstein estimates and analysis.

EXHIBIT 32: **Mature Consumer Internet Service User Base by Provider**

| Services | Total Users / MAU (mn) | Paid Subscribers (mn) | Subscription Ratio |
|---|---|---|---|
| Amazon Prime | 310 | 250 | 80.6% |
| Spotify | 751 | 290 | 38.6% |
| LinkedIn Premium | 1,200 | 175 | 14.6% |
| Duolingo | 133 | 12 | 9.2% |
| YouTube Premium | 2,500 | 125 | 4.9% |
| Snapchat+ | 932 | ~25 | 2.7% |
| X (Twitter) Premium | 261 | 3-5 | 1-2% |

Data of X (Twitter) are estimated by Bernstein.
Source: Company disclosures, Bernstein estimates and analysis.

**Significant advertising revenue potential**

OpenAI has begun testing ChatGPT advertising in the U.S., initially targeting larger-budget customers. Early data from FirstPageSage (as of 10th March 2026) indicates that ChatGPT's advertising conversion rates (0.3-6.0% across industries) are already outperforming Google Advertising benchmarks (0.3-2.2% across industries) (Exhibit 33). While these results are based on limited initial testing, we see meaningful potential for ChatGPT's monetization in advertising as the platform scales and advertiser penetration increases.

**Substantial investments with partners for future growth**

OpenAI continues to build out massive AI infrastructure to strengthen its competitive moat and support long-term growth. According to CEO Sam Altman, the company is committed to developing 30+ gigawatts (GW) of compute capacity at an estimated cost of $1.4tn. OpenAI has already secured deals totaling roughly 30-35 GW, including the 10 GW Stargate project with Oracle and SoftBank, the original 10 GW procurement with Nvidia (which might be revised lower following the February 2026 investment), and 6 GW with AMD (Exhibit 34). The company is also developing custom chips (likely in partnership with Broadcom) to reduce reliance on external suppliers and improve performance-per-dollar over time. Collectively, these initiatives position OpenAI to scale faster than peers and reinforce its leadership in infrastructure-driven AI capability.

EXHIBIT 33: **Google Search vs. ChatGPT: Advertising Conversion Rates by Industry**

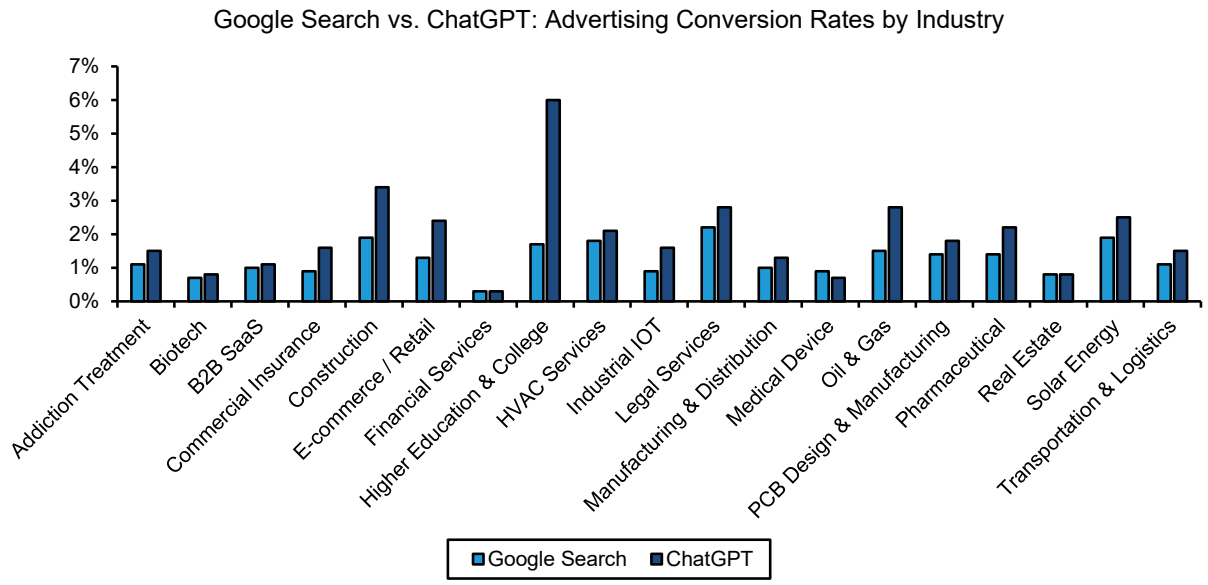


Based on initial ChatGPT dataset covering U.S. larger-budget customers, as of 10th March, 2026.
Source: FirstPageSage, Bernstein analysis.

EXHIBIT 34: **OpenAI: Strategic Partnerships with Major Tech Companies**

| Partners | Announced | Deal Amount | Gigawatts (GW) | Delivery Timeline |
|---|---|---|---|---|
| Microsoft | Original: 2019-2023;<br>Restructured: Oct 2025 | • $250bn compute commitment<br>• ~$135bn equity stake (27%) | N/A | Ongoing since 2019 |
| SoftBank / SB Energy (Stargate) | Jan 2025 (Stargate announced);<br>Jan 2026 (1.5 GW) | $500bn (Stargate as a whole by 2029; original) → $1bn (Jan 2026 project only) | 10GW (original; Stargate as a whole);<br>1.5 GW announced (confirmed) | Multiyear |
| Oracle (Stargate) | Jan 2025 (Stargate announced);<br>July 2025 (4.5 GW); Sept 2025 (expansion; ~1GW) | • $500bn (Stargate as a whole by 2029; original) → tens of $bn for the current 5.5GW announced<br>• $300bn cloud contract over 5 years (starting 2027)<br>• $30bn/year lease | 10GW (original; Stargate as a whole);<br>5.5 GW announced (confirmed) | Abilene, TX: Sept 2025 operational<br>Full 1.2 GW: Mid-2026<br>5 new sites: 2026-2027 |
| CoreWeave | Mar / May / Sept 2025 | • OpenAI acquired $350mn equity stake<br>• Mar 2025: An initial 5-year cloud contract up to $11.9bn<br>• May 2025: An expansion of cloud contract by up to $4bn<br>• Sept 2025: A further expansion of cloud contract by up to $6.5bn | Support part of 7 GW Stargate | Part of Stargate rollout 2025-2027 |
| Nvidia | Sept 2025 (original); Feb 2026 (revised) | $100bn (original) → $30bn (revised) | 10 GW (original) > Suspended | Original deal stalled |
| AMD | Oct 2025 | ~$90bn hardware revenue potential | 6 GW | First 1 GW: 2H26; Full 6 GW: Multi-year |
| Broadcom | Oct 2025 | Undisclosed (custom chip dev) | 10 GW | 2H26-2029 |
| Amazon (AWS) | Nov 2025 / Feb 2026 | • Contract value: $38bn multiyear > $100bn over 8 years<br>• Amazon invested $50bn in OpenAI | 2 GW (Trainium through AWS) | Multiyear |
| Cerebras | Jan / April 2026 | • Jan 2026: $10+bn multi-year agreement till 2028<br>• Apr 2026: Doubled the amount to $20+bn (till 2028) | • Jan 2026: 0.75 GW<br>• Apr 2026: No update | 2026-28 |

All companies mentioned are listed; every company except for Cerebras is covered by Bernstein's global technology research teams.
Source: Company disclosures, Bernstein analysis.

### Market Share loss concerns overdone

Despite recent concerns around OpenAI's market share losses to Gemini and Anthropic, we view these shifts as normal given the substantial expansion of the overall AI market. While OpenAI's market share edged down to 73.3% in March 2026 (vs. 74.1% a year earlier) (Exhibit 35), monthly active user (MAU) growth remains solid, with non-Copilot users +4% QoQ in March. As of March 2026, OpenAI still holds ~60% share in chatbots (or ~73% when including Copilot), indicating that ChatGPT is ceding share but not users. At the same time, OpenAI continues to gain traction against traditional internet services such as

Google Search (Exhibit 36). Moreover, the company's annual recurring revenue (ARR) growth remains exponential, reaching $25bn in February 2026 (vs. $6.5bn a year earlier), underscoring resilient monetization despite competitive noise (Exhibit 37).

In addition, OpenAI Codex overtook Claude Code on May 3rd, 2026, with weekly downloads surging to 86.1mn versus 7.2mn for Claude Code, highlighting a sharp acceleration in developer adoption (Exhibit 38). ChatGPT also enjoys the highest 12-month retention rate among major AI tools (Exhibit 39).

EXHIBIT 35: **Generative AI Chatbot Market Share in the U.S.**

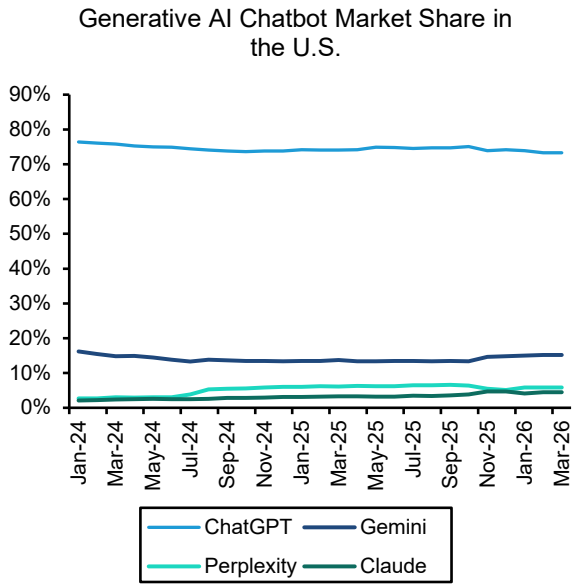


1. ChatGPT market share includes Microsoft Copilot.

2. March 2026 data, updated as of 10th March 2026.
Source: FirstPageSage, Bernstein analysis.

EXHIBIT 36: **ChatGPT has gained share from Google in the global search market**

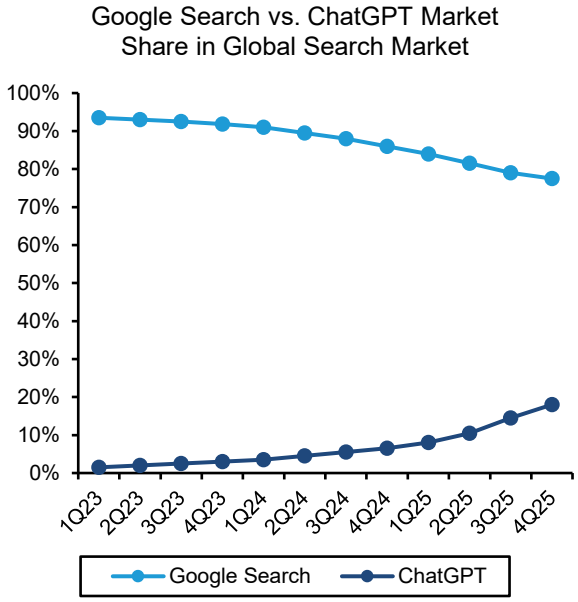


Source: FirstPageSage, Bernstein analysis.

EXHIBIT 37: **OpenAI's ARR has grown exponentially**

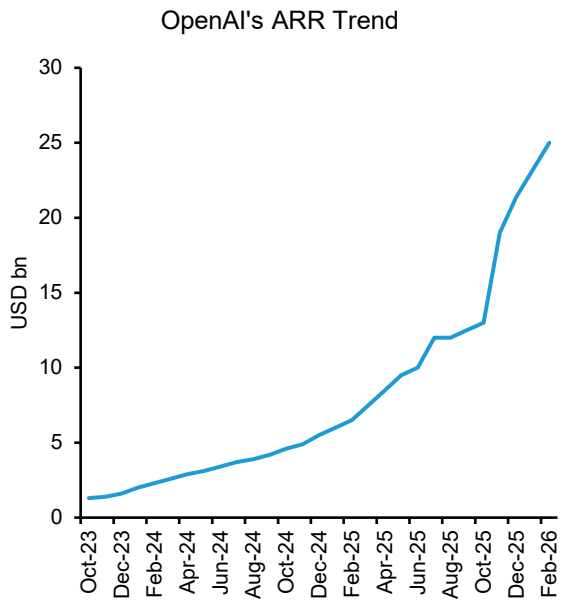


Source: The Information, Bernstein analysis.

EXHIBIT 38: **OpenAI Codex overtook Claude Code on May 3rd, with weekly downloads surging to 86.1mn versus 7.2mn for Claude Code, highlighting a sharp acceleration in developer adoption.**

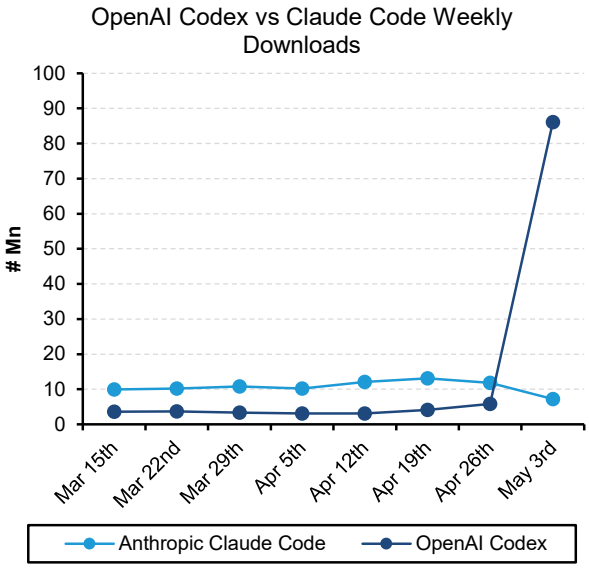


Source: TickerTrends NPM, Bernstein analysis

EXHIBIT 39: **AI Tools: Average User Retention Rate within 12 Months**

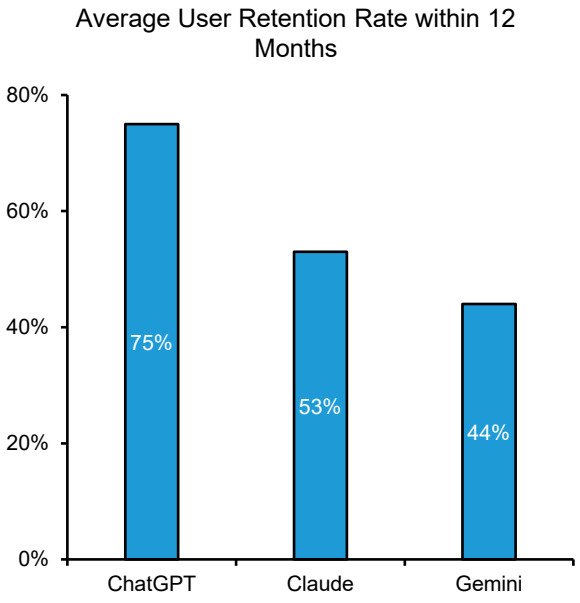


Data represents average retention for premium U.S. Desktop users based on monthly cohorts from December 2023 to March 2026.
Source: SoftBank, SimilarWeb

**Just because it burns cash doesn't mean it's not a viable business model**

It's well reported that OpenAI will burn billions of dollars (The Information estimates more than $200bn burn before cash flow turns positive in 2030). On the other hand, both OpenAI and Anthropic CEO have said that if they stop developing new models, they would be profitable immediately. But the cost of developing a new frontier model is 5-10x that of the previous model. So if they spend $1 to develop a model, next year they make $2 from that model, but spends $5 to develop a new model. The year after, the loss widens as the cost of new model blooms to $25.

However, this is only happening because the new model delivers more than 5x the capabilities, and hence more than 5x the monetization potential compared to the previous model. We believe there are only two ways this can continue:

1. New models continue to scale 5x every year and OpenAI eventually reaches AGI (no matter what the definition is)
2. New models marginal cost slows down, and the company turns profitable immediately and generates enormous cash flow

Scenario #2 is not uncommon in the business history. Netflix, Amazon, Uber (Exhibit 40), Meituan, and numerous other companies have experienced the same period of huge cash burn before becoming a sustainably profitable business, growing in scale and building moat in the process. Those familiar with Chinese food delivery market would understand Meituan's model: It proved its food delivery unit economics in certain areas, then expanded to many more areas with big subsidy, burning cash in the meanwhile, but knowing that the business model is viable. Once the subsidy stopped and market share consolidated, it quickly turned profitable.

Netflix's case is closer to OpenAI's: it also had a timing issue that it had to pay cash upfront to license and produce shows and movies well before those costs were recognized through the income statement. The biggest cash outflow was all on content spending, especially original programming and the build-out of its streaming library. For each of its earlier years, the content it acquired before is generating profit, but it had to spend more on acquiring new content, resulting in more cash outflows. However, once the content acquisition slowed down, it immediately turned positive in cash flow (Exhibit 41). Amazon, Tesla and Airbnb all experienced the same kind of negative cash flow widening before profits blooming (Exhibit 42).

EXHIBIT 40: **Uber: Free Cash Flow**

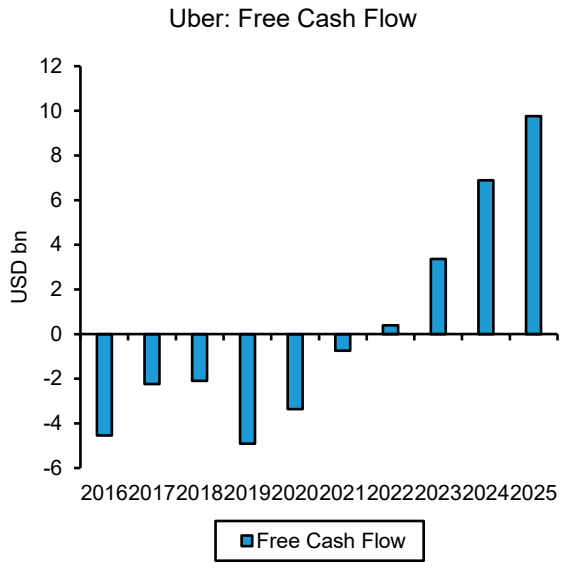


Source: Company disclosures, Bloomberg, Bernstein analysis.

EXHIBIT 41: **Netflix: Free Cash Flow**

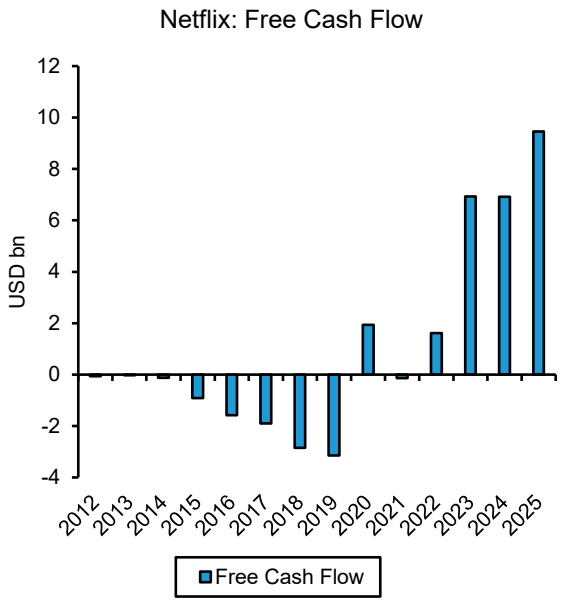


Source: Company disclosures, Bloomberg, Bernstein analysis.

EXHIBIT 42: **"New Business Model" Tech: Free Cash Flow Trajectory**

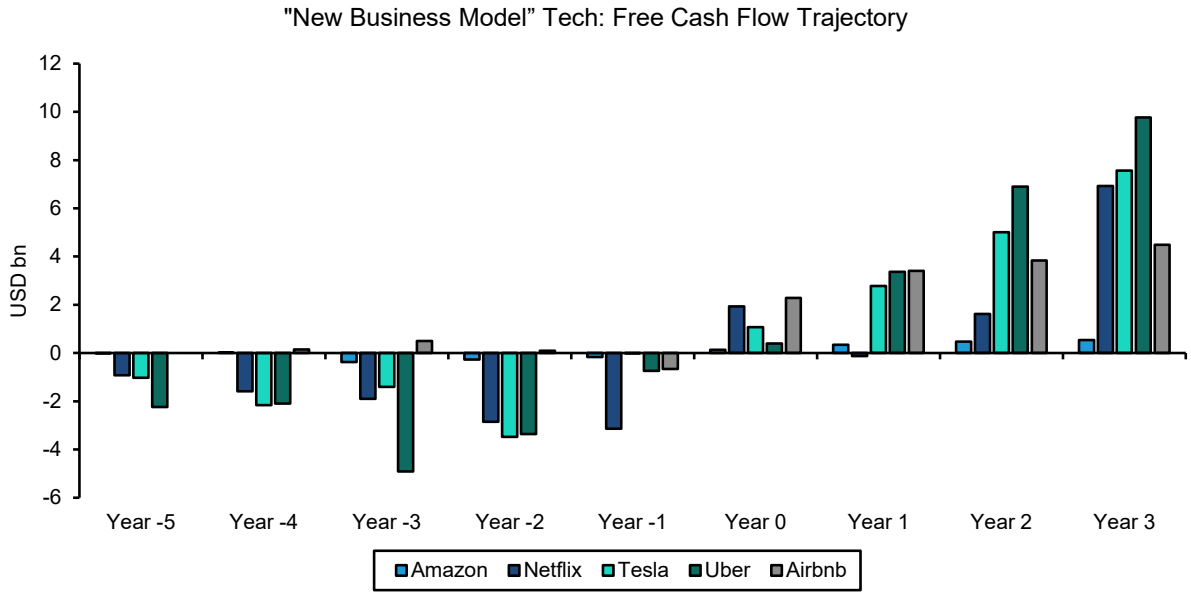


Source: Company disclosures, Bloomberg, Bernstein analysis.

## AI SILICON: ARM, GRAPHCORE, AMPERE (IPS, GPUS, CPUS)

SoftBank consolidated semiconductor-related subsidiaries including Arm, Graphcore, and Ampere, and established a new "AI Computing Segment" in 3Q FY2025 (Dec-25 Q), to provide AI silicon solutions (Exhibit 43). Arm, Graphcore and Ampere plays the most significant roles in SoftBank's AI silicon efforts. Despite still in the preliminary stage, SoftBank plans to coordinate these 3 companies, along with OpenAI partnership, Stargate project to provide comprehensive AI silicon solutions. Although still at an early stage, SoftBank aims to coordinate these three companies, alongside its partnership with OpenAI and the Stargate infrastructure project, to deliver a comprehensive AI silicon solution.

To strengthen execution, Rene Haas (CEO of Arm) was appointed as CEO of SoftBank Group International (SoftBank's international operations) in April 2026, overseeing key overseas subsidiaries, including Arm and Graphcore. This leadership structure is expected to enhance coordination across SoftBank's semiconductor portfolio, supporting the advancement of its integrated AI silicon strategy.

**Arm:** IPs and Silicon (for CPUs, GPUs, ASICs)
Arm is a UK-based global leading provider of semiconductor intellectual property (IPs), designing and licensing energy-efficient CPU cores, GPUs, and other accelerators. Its architecture dominates the smartphone market, powering nearly 100% of modern devices, and underpins 350bn+ chips shipped globally. Arm's technology now extends well beyond mobile into cloud, AI, IoT, automotive, and edge computing. Operating an asset-light licensing and royalty model, Arm enables semiconductor partners to integrate its IP into ASICs across diverse end markets. Its portfolio spans smartphone application processors, high-performance server CPUs, GPUs, and dedicated AI accelerators, forming a foundational compute platform for data center and edge AI workloads. In March 2026, Arm announced its entry into the server CPU market with its new "Arm AGI CPU" product line. The company was acquired by SoftBank in 2016 and subsequently relisted on Nasdaq in 2023.

Beyond its core CPU silicon roadmap, Arm is deepening its collaboration with parent company SoftBank in the accelerator space. Arm confirmed that there are three concurrent NRE projects being run for SoftBank as a related party, which could be for their customers / Crystal Intelligence / Stargate. These seem to be distinct from SoftBank subsidiaries' semiconductor efforts at Ampere and Graphcore, although notably Arm's CEO Rene Haas is expected to take on a broader leadership role overseeing SBG's semiconductor portfolio which will encompass Graphcore and Ampere, further tightening the coordination between the related entities. *For more information on Arm, please refer to the Arm section in this report (Exhibit 44, Exhibit 45, Exhibit 46).*

**Graphcore:** AI Accelerators (GPUs)

Graphcore is a UK-based fabless semiconductor company focused on AI and machine-learning acceleration (the company's own GPU solutions) through its proprietary Intelligence Processing Unit (IPU) architecture. The IPU is a massively parallel processor designed from the ground up for machine intelligence, integrating high on-chip memory and thousands of parallel threads to optimize training and inference performance. Its latest Colossus MK2 GC200 IPU contains 59.4bn transistors and 1,472 cores, offering substantial performance gains versus prior generations. Graphcore provides IPU systems, including the IPU-M2000 and scale-out IPU-PODs, supported by the Poplar SDK for model development and deployment. The company has been recognized for enabling new AI workloads not well-served by traditional GPUs and was acquired by SoftBank in 2024.

**Ampere:** Server CPUs
Ampere is a U.S.-based fabless semiconductor company developing high-performance, energy-efficient Arm-based CPUs optimized for cloud and data center environments. Founded in 2017 and acquired by SoftBank in 2025, Ampere offers cloud-native processors such as Ampere Altra (up to 128 cores) and the newer AmpereOne family, designed for predictable, high-throughput performance, linear scalability, and improved power efficiency versus traditional x86 architectures. Its CPUs are widely deployed across hyperscalers and enterprise cloud platforms for workloads including AI inference, microservices, data analytics, and edge computing. Ampere emphasizes single-threaded cores, high VM/container density, and sustainable compute performance.

EXHIBIT 43: **SoftBank consolidated semiconductor-related subsidiaries including Arm, Graphcore, and Ampere, and established a new "AI Computing Segment" in 3Q FY2025**

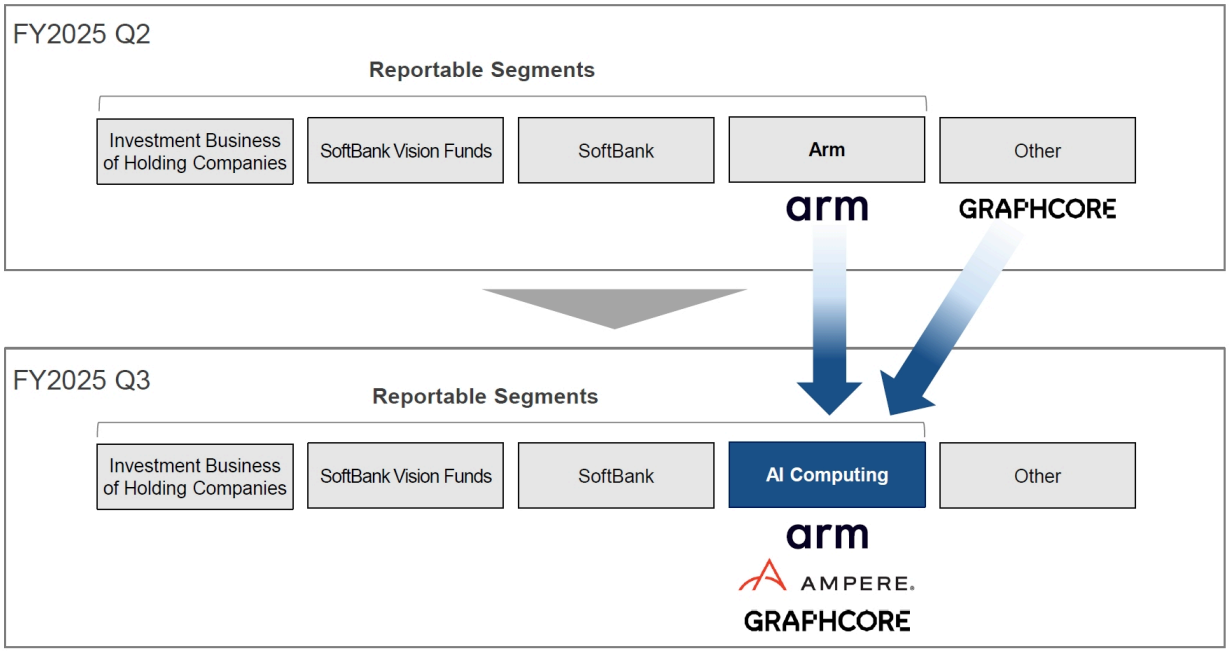


Source: SoftBank

EXHIBIT 44: **We anticipate Arm's revenue to grow at a 40% CAGR over the next 5 years, primarily driven by the expansion of its new AGI CPU business, which we expect to reach US$15bn by FY31. Growth is expected to be further supported by strong momentum in royalty revenue, which we forecast to grow at a 20% CAGR. We expect NG NI to grow at 42% CAGR.**

| End Market | Metrics | FY26E | FY31E | CAGR 26-31E |
|---|---|---|---|---|
| Edge AI | TAM:<br>Market Share:<br>Royalty Revenue: | $174 Bn<br>~60%<br>$ ~2.13 Bn | $198 Bn<br>~66%<br>$ ~4.16 Bn | 3%<br><br>14% |
| Cloud AI | TAM:<br>Market Share:<br>Royalty Revenue:<br>Chip Revenue: | $61 Bn<br>~ 29%<br>$0.38 Bn<br>- | $ 171 Bn<br>~ 54%<br>$ 2.06 Bn<br>$ 15 Bn | 23%<br><br>40% |
| Physical AI | TAM:<br>Market Share:<br>Royalty Revenue: | $25 Bn<br>46%<br>$0.18 Bn | $42 Bn<br>55%<br>$0.52 Bn | 11%<br><br>23% |
| Total Revenue | TAM:<br>Market Share:<br>Royalty Revenue:<br>Chip Revenue:<br>Total Revenue: | $260 Bn<br>50%<br>$2.69 Bn<br>-<br>$4.92 Bn | $411 Bn<br>60%<br>$6.74 Bn<br>$15 Bn<br>$25.90 Bn | 10%<br><br>20%<br><br>40% |
| Net Income | Non-GAAP NI: | $1.89 Bn | $10.96 Bn | 42% |

Source: Company reports, Bernstein analysis and estimates

EXHIBIT 45: **The ~40% CAGR over the next 5 years will be primarily driven by the expansion of the Cloud AI (IP/ CSS + AGI CPU) business, which we forecast will reach US$17bn in FY31, representing 66% of total revenue**

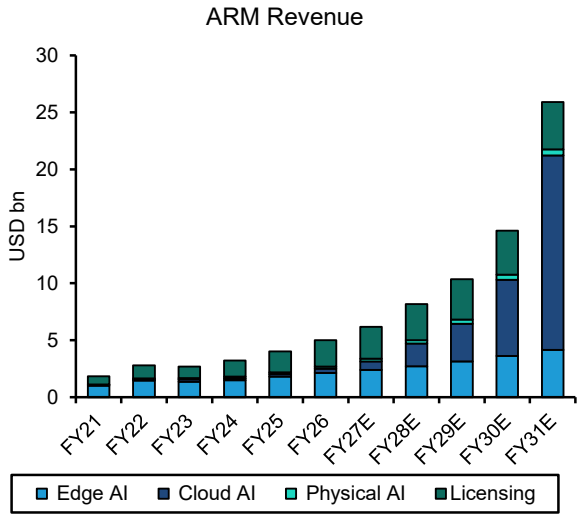


Source: Company reports, Bernstein analysis and estimates

EXHIBIT 46: **We expect Non-GAAP Net Income to surpass US$ 10Bn in FY31, representing a 42% CAGR from US$ 1.9Bn in FY26.**

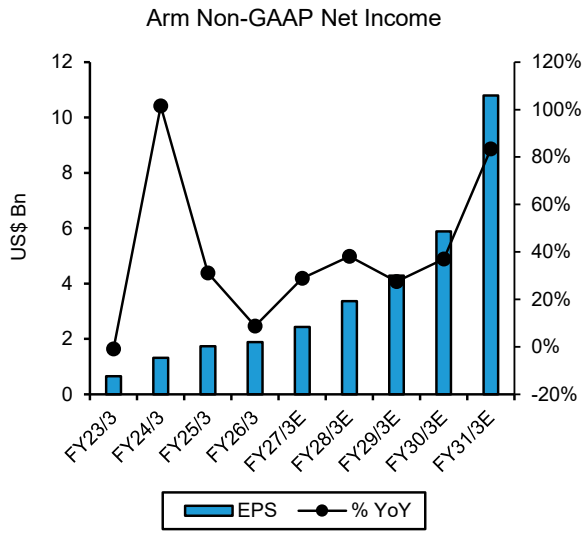


Source: Company reports, Bernstein analysis and estimates

## AI INFRASTRUCTURE: STARGATE & DIGITALBRIDGE

### Stargate: Viable infrastructure model for SoftBank without huge cash burn

The Stargate Project was launched in January 2025 as a $500bn AI infrastructure joint venture between OpenAI, SoftBank, Oracle, and MGX, with SoftBank acting as financial lead and OpenAI as operational lead. The initial plan committed $100bn, targeting the build-out of 10 GW of AI data center capacity across the U.S. **However, the consortium revised and refined the execution plan later in 2025** (Exhibit 47)**.**

Within the new framework, the Oracle-OpenAI partnership is responsible for ~5.5 GW of capacity across 5 sites, representing roughly $300bn of investment over the next five years. SoftBank-OpenAI will jointly develop an additional 1.5 GW across two sites over the next 18 months, supported by SB Energy's powered-infrastructure capabilities. Oracle confirmed 4.5 GW of incremental capacity in July 2025 (5.5 GW together with a project announced earlier in 2025), with active deployments underway in Texas, New Mexico, and the broader Midwest, including partial operations at the flagship Abilene, Texas facility equipped with NVIDIA GB200 racks. SoftBank-OpenAI's 2 sites in Ohio and Texas are similarly designed to scale up to 1.5 GW, with SB Energy enabling rapid-build delivery. Combined, commitments from both partnerships now total ~7 GW and are backed by $400bn+ of secured funding, keeping the project on track to meet the $500bn / 10 GW target.

Critically, and **contrary to persistent investor concerns, SoftBank's capital exposure is limited.** SoftBank's role is primarily financial advisor and coordinator, earning long-term lease revenue, rather than directly funding or operating large-scale data center infrastructure. SoftBank invests only $500mn into SB Energy (matched by OpenAI's $500mn). SB Energy develops the real estate and physical infrastructure, while OpenAI funds and operates the IT stack (including GPUs and racks). With typical data center costs of $10-13bn per GW, the 1.5 GW SoftBank-linked build implies $12-16bn of capex, financed mainly through project financing at the Stargate joint venture level and not consolidated onto SoftBank's balance sheet (i.e., external debt). SoftBank captures 15-year lease payments once operations commence.

In effect, **SoftBank has managed to free itself from the massive capex burden implied by the full $500bn Stargate commitment, a key point that remains widely misunderstood by investors.** Despite market perception, SoftBank is not funding the majority of Stargate's capex, and its cash outlay and balance-sheet risk remain limited. We believe **the long-term lease model offers a more attractive business profile than renting GPU or compute capacity, providing higher visibility and lower operational risk for SoftBank.**

### DigitalBridge: Real estate deal sourcing & financing capabilities enhancement

SoftBank announced in December 2025 that it would acquire DigitalBridge for $4bn, a leading global digital infrastructure asset manager with ~$108bn AUM, to strengthen its capabilities in sourcing, financing, and scaling AI-related real estate and data center assets. The transaction provides SoftBank with deep expertise across data centers, towers, fiber networks, and edge infrastructure, enhancing its ability to originate and structure large-scale AI infrastructure projects. Management highlighted that DigitalBridge's global portfolio and financing platform will directly support SoftBank's next-generation AI initiatives, including large data center developments, by improving access to power-secured land, expanding real estate deal sourcing channels, and enabling more efficient project-finance structures and long-term infrastructure capital.

One strategic asset within the DigitalBridge platform is **Switch**, in which DigitalBridge holds a 55.8% stake. Switch is a leading, renewably powered enterprise data center operator with large campus developments across Nevada, Texas, Georgia, and Michigan. Switch's deep experience in designing, constructing, and leasing data center infrastructure in the U.S. is highly complementary to SB Energy's role in the Stargate project. Although SoftBank previously explored a direct acquisition of Switch, at a reported valuation of ~$50bn before halting discussions in January 2026, its ownership of DigitalBridge enables SoftBank to partner closely with Switch on Stargate-related developments without bearing full acquisition or balance-sheet risk.

EXHIBIT 47: **Stargate Project: Announced Plans**

| Partnership | Capacity (GW) | Site Location | Capex | Funding Structure | SoftBank Exposure | Revenue to SoftBank |
|---|---|---|---|---|---|---|
| Oracle-OpenAI | ~5.5 GW | 5 sites across Texas, New Mexico, Midwest | ~$300bn over 5 years | Oracle + OpenAI financed | None | None |
| SoftBank-OpenAI | ~1.5 GW | 2 sites (Ohio, Texas) | Roughly $12–16bn | Mainly project financing at Stargate JV | Limited consolidated capex; only $500mn equity into SB Energy | 15-year lease revenues from infrastructure/real estate |
| Remaining Capacity | ~3 GW | TBD | TBD | TBD | TBD | TBD |

Source: Company disclosures, Bernstein analysis.

## PHYSICAL AI: SOFTBANK ROBOTICS AND ABB ROBOTICS

Physical AI is the future form of AI, after generative AI and argentic AI (Exhibit 48). The common framework of Physical AI comprising "body-cerebellum-brain-skills-world", through which industry players and their chosen business models are best understood (Exhibit 49 and Exhibit 50). In broad strokes, body and cerebellum are known territories, and brain and world are the unknown. "Skills", like apps in the app store, are far too limited today but can quickly expand when robotic brain plus world models achieve cause-effect reasoning based on physical data. *For more information on Physical AI, please also refer to our Asia Industrial team's research Future of Tech: Physical AI -- Bridging industrial and humanoid robotics.*

Physical AI brings robotics into a new era by combining advanced AI models with machines that can perceive, reason, and act in the real world. Unlike traditional robots that execute predefined rules, physical AI systems use sensors, actuators, and foundation-model intelligence to understand their surroundings, make context-aware decisions, and improve through experience. This enables robots, and broader physical systems, to navigate complex environments, learn cause-and-effect relationships, and generalize skills across tasks, forming the basis for more capable humanoids, autonomous vehicles, and next-generation automated infrastructure. As industry frameworks suggest, progress in "brain" (reasoning models) and "world" (real-world data and simulation) will unlock exponential expansion in "skills," allowing physical AI to move from narrow automation to adaptable, general-purpose physical agents.

The humanoid robot has largely been a research and entertainment tool, with only 4% of shipments in industrial applications in 2025, when total humanoid robot shipments were merely ~20K units. We expect this percentage to quickly increase and peak at 50-60% when the overall market size is 1-10 million units p.a., but drop to ~10% in the far future (Exhibit 51). This curve reflects our view of humanoid killer applications first emerging in structured industrial settings, and proliferating to unstructured consumer settings as the task generalization capability improves. Even in the future, we only expect limited overlap of applications between industrial and humanoid robots (Exhibit 52).

Physical AI has substantial growth potential as enterprise adoption accelerates. According to Deloitte, 80% of companies plan to implement Physical AI within the next two years, up from 58% today, highlighting an early but rapid shift toward AI-enabled physical systems in real-world environments (Exhibit 53). To capture the long-term potential of AI robotics and Physical AI, SoftBank has been investing ahead of the curve, most recently through its $5.4bn acquisition of **ABB Robotics**, a global leader in industrial automation, significantly enhancing its capabilities in advanced robotics, Physical AI, and next-generation AI-driven automation.

EXHIBIT 48: **From communication to execution to embodiment: the Gen AI roadmap toward enterprise-scale general intelligence**

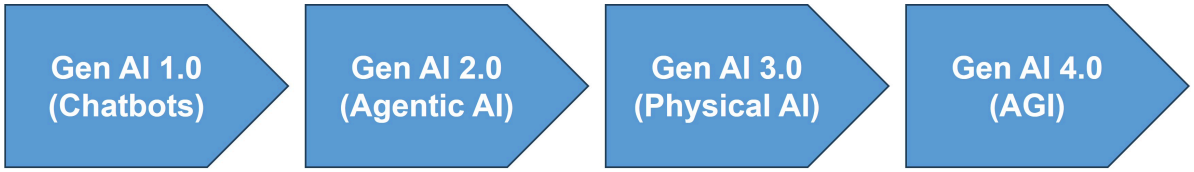


- Internal help desk
- Customer support
- Content creation

- Software development
- Financial trading
- Drug discovery
- Knowledge work

- Humanoid robots
- Industrial automation
- Warehouse logistics
- Homecare

- Universal medical diagnosis
- Autonomous scientific research
- Adaptive autonomous system
- Fully adaptive tutoring

Source: Bernstein analysis

EXHIBIT 49: **A common framework of Physical AI comprising "body-cerebellum-brain-skills-world".**

| | | FANUC | Mech-Mind | Nvidia | Unitree | Helix 02 | Figure AI | Keyence |
|---|---|---|---|---|---|---|---|---|
| World | Simulation | ✓ | | ✓ | | | | |
| Skill | Applications | ✓ | ✓ | | | | ✓ | |
| Cerebellum & brain | Task reasoning | | ✓ | ✓ | ✓ | ✓ S2 | ✓ | |
| | Path/posture planning | ✓ | ✓ | ✓ | ✓ | ✓ S1 | ✓ | ✓ |
| | Motion control | ✓ | | | ✓ | ✓ S0 | ✓ | |
| Body | Robot | ✓ | | | ✓ | | ✓ | |

Mech-Mind, Figure AI, and Unitree are private. Helix 02 is the VLA model developed by Figure AI. NVIDIA is covered by Bernstein U.S. Semiconductors team. FANUC and Keyence are covered by Bernstein Asia Industrial team.
Source: Bernstein Asia Industrial research team analysis.

EXHIBIT 50: **Key industry players and their chosen businesses in physical AI**

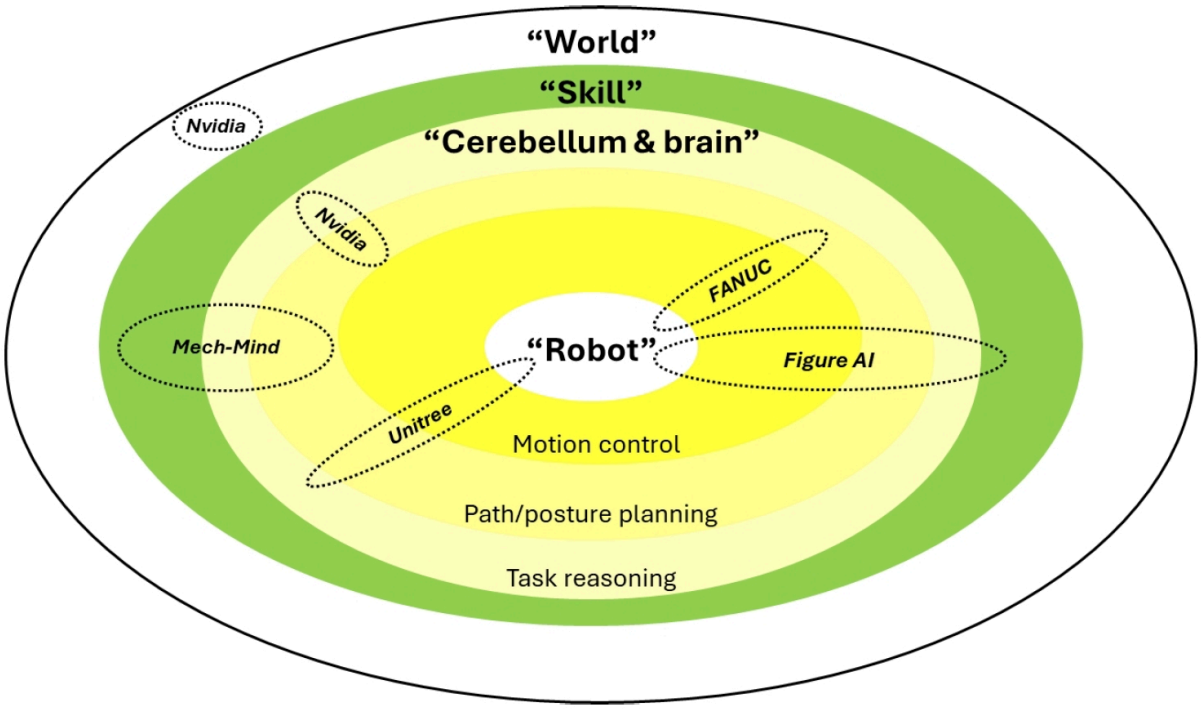


Nvidia is covered by Stacy A. Rasgon. Figure AI, Unitree, Mech-Mind are private.
Source: Bernstein Asia Industrial research team analysis.

EXHIBIT 51: **Only ~4% of humanoid robot shipments were for industrial applications in 2025, when total humanoid robot shipments were merely ~20K units. We expect this share to rise quickly and peak at 50 to 60% when the overall market reaches 1 to 10 million units p.a., but drop to ~10% in the far future.**

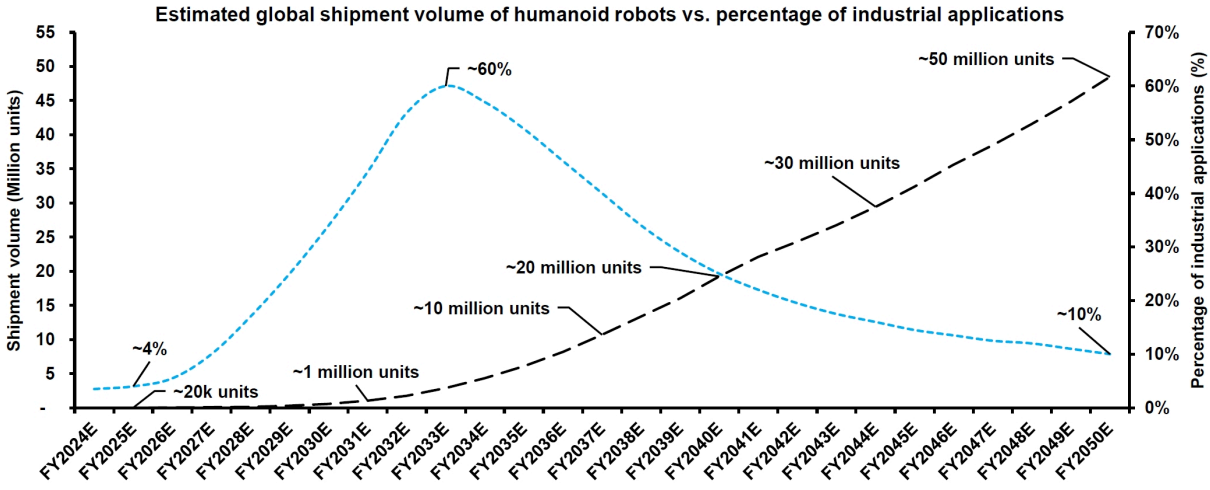


Source: Industry news, channel checks, Bernstein Asia Industrial research team analysis and estimates.

EXHIBIT 52: **We anticipate minimal application overlap between industrial robots and humanoid robots in the overall industrial vertical.**

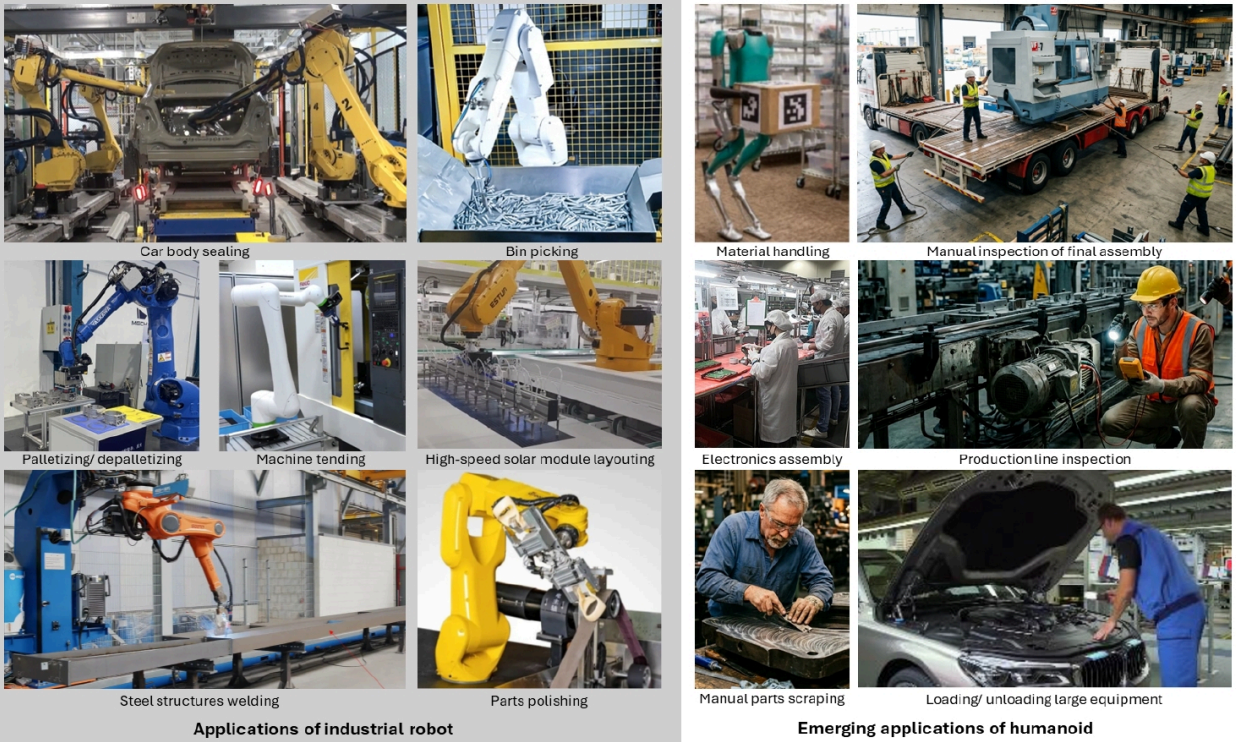


Source: Wiki common, FANUC, Agility, Bernstein Asia Industrial research team analysis.

EXHIBIT 53: **Corporate Survey: Extent of Physical AI Usage**

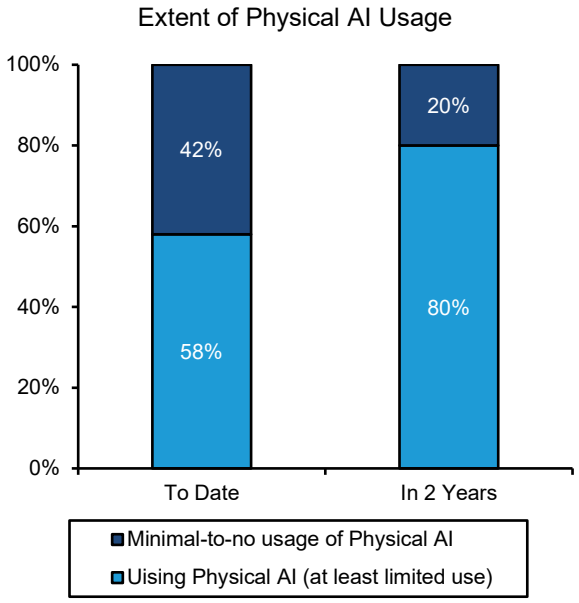


Survey Question: To what extent is your organization currently utilizing physical AI (e.g., robotics, automated machinery) in its operations? N=3,235; fieldwork conducted August-September 2025.
Source: Deloitte, Bernstein analysis.

## VALUATION

Over the past 5 years, SoftBank has traded at a 4-63% discount to reported net asset value (NAV), with a mean discount of ~45%, reflecting asset volatility, valuation opacity of private investments, and balance-sheet leverage (Exhibit 54 and Exhibit 55). Looking ahead, we expect the proportion of non-OpenAI private investments to decline to ~20-21% from CY2026 onward, down from 28% as of end-CY2025 (Exhibit 56). In addition, we view OpenAI's potential IPO as a major catalyst for NAV discount narrowness (according to *OpenAI Plans Fourth-Quarter IPO in Race to Beat Anthropic to Market - WSJ*), leading to a share-price re-rating.

From a relative-valuation perspective, large financial holding companies typically trade at 10-30% discounts to NAV (Exhibit 57). While other investment holding companies such as Naspers exhibit structurally wider discounts (~35-60%) due to high asset concentration, and EXOR N.V. reflects elevated discounts driven by underperformance in down-cycle auto investments (notably Stellantis, share price dropped ~60% over the past 5 years [STLAP.FP, covered by Stephen Reitman]) and a challenging EV transition, we view SoftBank's portfolio as better diversified, with improving valuation visibility. Accordingly, we believe a 30% discount to NAV is appropriate for SoftBank, reflecting its long-term value creation across the AI ecosystem (infrastructure, silicon, models, and Physical AI), alongside the upcoming IPO-driven re-rating catalyst.

EXHIBIT 54: **SoftBank Group - Market Cap as % of NAV**

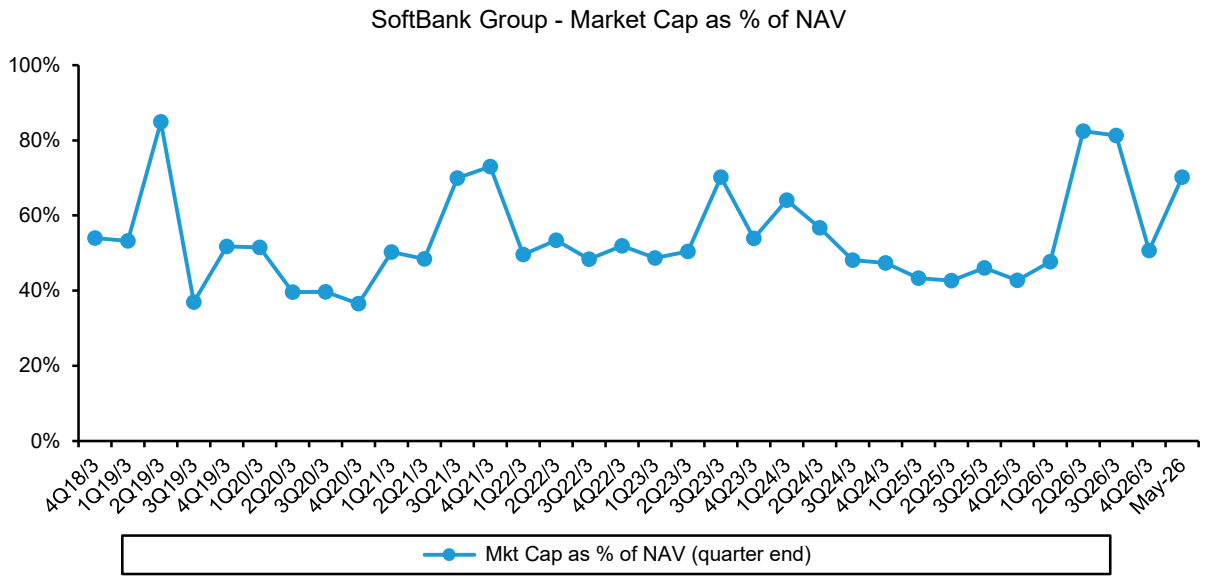


Source: Company disclosures, Bloomberg, Bernstein estimates and analysis.

EXHIBIT 55: **SoftBank's Valuation as % to NAV vs. Public Investment % Correlation**

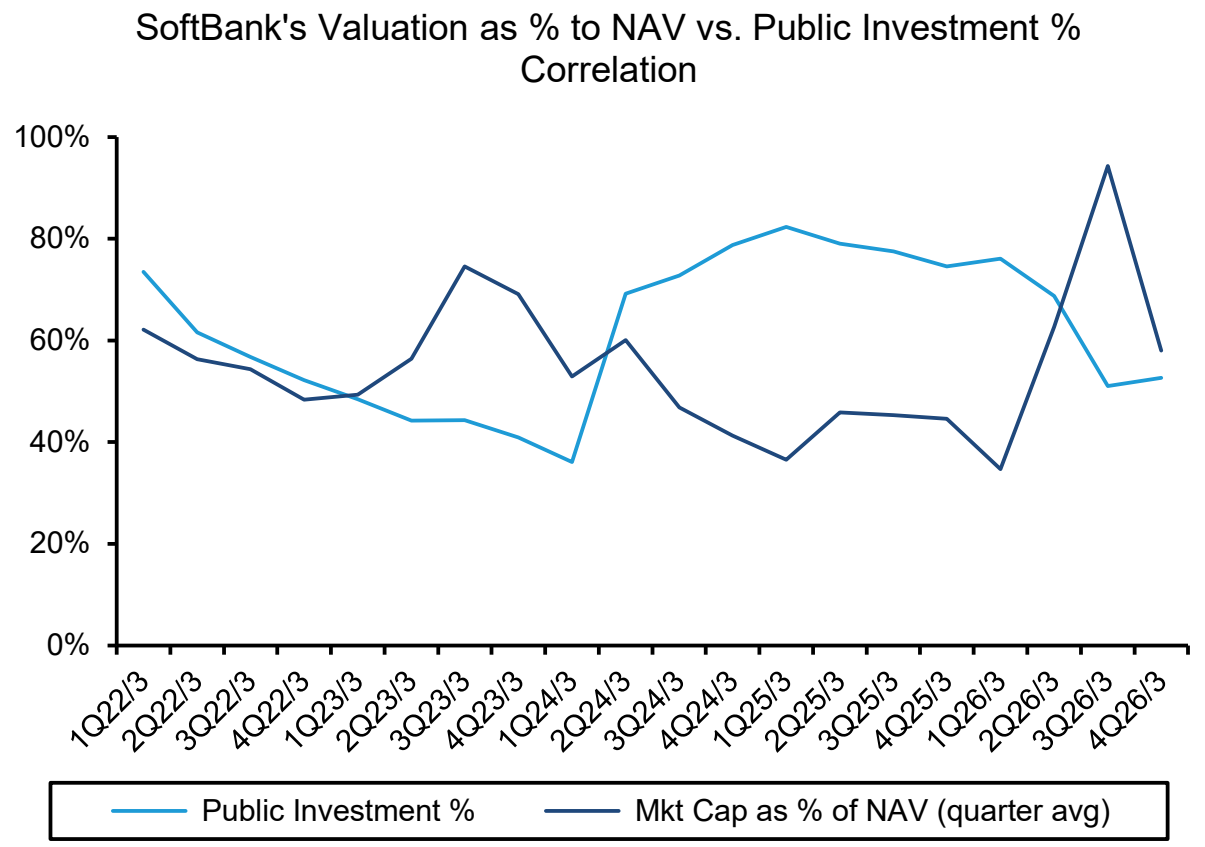


Source: Company disclosures, Bloomberg, Bernstein analysis.

EXHIBIT 56: **SoftBank Group's Investment Portfolio**

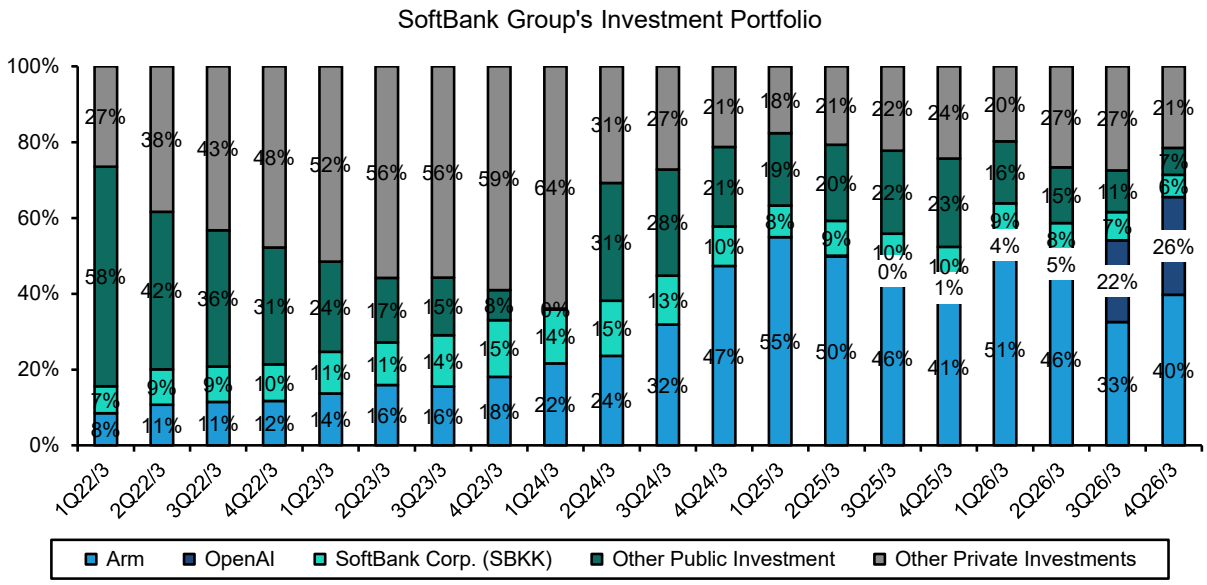


Source: Company disclosures, Bloomberg, Bernstein analysis.

EXHIBIT 57: **Major Investment Holding Companies' Valuation Discount to NAV**

| Investment Holding Company | Ticker | Country | Key Holdings | Discount to NAV |
|---|---|---|---|---|
| SoftBank Group | 9984.JP | Japan | Arm, OpenAI, SoftBank Corp. (SBKK), Alibaba, Intel | 4-63% |
| Naspers | NPN.SJ | South Africa | Tencent (via Prosus) | 35-60% |
| EXOR N.V. | EXO.NA | Netherlands | Ferrari, Stellantis, CNH, Philips, Juventus | 30-62% |
| Wendel Group | WF.FP | France | Bureau Veritas, Stahl, Crisis Prevention | 20-50% |
| AVI Global Trust | ABT.LN | UK | Holdcos & closed-end funds at discount-to-NAV | 6-12% |
| Berkshire Hathaway | BRK/B | U.S. | Apple, BNSF, GEICO, Coca-Cola, AmEx | 15% discount to 5% premium |
| Investor AB | INVEB.SS | Sweden | Atlas Copco, ABB, AstraZeneca, EQT | 5% discount to 10% premium |

Source: Company disclosures, Bloomberg, Bernstein analysis.

## SoftBank valuation

SoftBank is now trading at 27.5% discount to its NAV. We expect SoftBank's NAV growing to $390bn in a year, mostly contributing by Arm ($274bn; based on our price target of Arm $300.00) and OpenAI ($92bn, based on valuation in the latest funding round). With a 25% NAV discount, we rate SoftBank Outperform with PT = ¥8,200.00. Our PT implies 43% upside to SoftBank's current share price (Exhibit 58).

In our bull case, we apply our bull-case Arm valuation ($390 per share) and OpenAI (still based on valuation in the latest funding round). With a 15% NAV discount, which implies a valuation of ¥11,200, while in our bear case, we apply our Arm valuation based on its share price of $221 and same valuation of OpenAI). With a 45% NAV discount, which implies a valuation of ¥4,900.

Given SoftBank's strategic positioning across the AI value chain, we believe the company will be one of the largest long-term beneficiaries of the AI ecosystem within the Japan semiconductor sector. Moreover, if OpenAI goes public at a higher valuation (*OpenAI Plans Fourth-Quarter IPO in Race to Beat Anthropic to Market - WSJ*), it would represent additional upside for SoftBank. With our PT of **¥8,200.00**, we initiate coverage with an **Outperform** rating.

Our SoftBank financial model can be downloaded here: SoftBank (9984.JP)

EXHIBIT 58: **SoftBank Valuation: Scenario Analysis**

| **Softbank's 1-year NAV to valuation (USDbn)** | | | |
|---|---|---|---|
| | **Bull** | **Base** | **Bear** |
| Arm per share value (USD) | 390.00 | 300.00 | 221.21 |
| Share Count (mn) | 1,052 | 1,052 | 1,052 |
| Market cap (USD bn) | 410,280 | 315,600 | 232,713 |
| % holding | 86.9% | 86.9% | 86.9% |
| | | | |
| OpenAI Valuation (USD bn) | 852 | 852 | 852 |
| % holding | 10.7% | 10.7% | 10.7% |
| | | | |
| **Future NAV (USD bn)** | | | |
| Arm | 356.5 | 274.2 | 202.2 |
| OpenAI | 91.6 | 91.6 | 91.6 |
| Others | 23.7 | 23.7 | 23.7 |
| **Total** | **471.7** | **389.5** | **317.4** |
| *Total (JPY tn)* | *75.1* | *62.0* | *50.5* |
| | | | |
| Discount | 15% | 25% | 45% |
| | | | |
| Share Count (mn) | 5,699 | 5,699 | 5,699 |
| Target price | **11,200** | **8,200** | **4,900** |
| Upside/downside | 94% | 42% | (15%) |

Share price as of 13th May 2026.
Source: Company disclosures, Bloomberg, Bernstein estimates and analysis.

## SOFTBANK COMPANY PROFILE HISTORY / MANAGEMENT

### History

Established in 1981 and headquartered in Tokyo, Japan, SoftBank Group was founded by Masayoshi Son as a domestic software distributor before steadily transforming into one of the world's largest technology investment conglomerates. In its early years, the company built scale through software distribution, publishing and trade shows, then broadened its reach as the internet era emerged, taking stakes in high-growth technology businesses and assembling a wider portfolio of telecom and internet assets that ultimately evolved into today's SoftBank Corp., a listed subsidiary of SoftBank Group. That strategic evolution was defined by Son's willingness to make concentrated, long-duration bets on platforms he believed could shape the next phase of global computing.

SoftBank's profile changed dramatically in the 2000s and 2010s as it paired investment activity with large-scale corporate acquisitions. The acquisition of Vodafone Japan in 2006 established SoftBank as a major force in Japanese telecommunications, while the purchase of Sprint in 2013 marked a bold push into the U.S. wireless market. Alongside those operating investments, SoftBank became widely known for its remarkably successful early investment in Alibaba (2000), which turned a relatively small initial outlay into one of the most consequential venture-style investments in modern technology history and provided capital and credibility for future expansion.

The company's modern identity was cemented in the second half of the 2010s. SoftBank acquired Arm Holdings in 2016, giving it ownership of one of the semiconductor industry's most strategically important intellectual property companies, and in 2017 launched the first Vision Fund, a $100bn vehicle that made SoftBank synonymous with large-scale, high-conviction technology investing. After a period marked by portfolio volatility, asset monetization and tighter financial discipline, SoftBank has increasingly repositioned itself as an AI-focused holding company, with its strategy now centered on Arm, selective deployment through its investment platforms and a growing emphasis on AI infrastructure and ecosystem exposure, including OpenAI (since 2024).

### Management

**Masayoshi Son** founded SoftBank Group in September 1981 and has remained the defining figure in the company's leadership ever since. He currently serves as **Chairperson & CEO and has led the company since its establishmen**t. SoftBank credits him with driving the group's expansion across internet, telecommunications and AI, including landmark moves such as the

Alibaba investment, the acquisition of Arm and the establishment of the SoftBank Vision Fund. As the largest shareholder, Mr. Son holds 30% of SoftBank's shares.

On the finance side, **Yoshimitsu Goto** is **the current CFO of SoftBank since 2018**. Mr. Goto joined SoftBank in June 2000, became Head of Finance Department in July 2012, took on the CFO & CISO role in April 2018, and as of April 2026 serves as Board Director, Corporate Officer, Senior Vice President, CFO & CISO, with SoftBank highlighting his long-standing role in financing and business management across the group.

At the **Vision Fund**, the current head is **Alex Clavel**, whom SoftBank appointed as CEO since January 2025 (co-CEO from July 2023 to December 2024). In that role, Clavel is presented as overseeing the Vision Fund platform during a period when SoftBank has emphasized portfolio resilience, disciplined investment management and a strong focus on AI-led growth opportunities. Alex joined SoftBank in 2015, after 19 years at Morgan Stanley in investment banking across New York, Tokyo, Hong Kong, and Shanghai.

## ARM: THE CPU RENAISSANCE BROUGHT BY AGENTIC AI

Arm is in the middle of a critical strategic transition in its history: from a pure-play licensor of CPU and system IP, to a provider of higher-value Compute Subsystems (CSS), and now, for the first time, to a supplier of Arm-designed silicon through the newly launched Arm AGI CPU. In practical terms, the business now relies on three revenue streams. First, licensing revenue is the upfront monetization of Arm's technology designs, software and access rights, which customers use to build their own chips and platforms. Second, royalty revenue is the recurring annuity stream earned on chips shipped with Arm technology inside them, and it remains the economic core of the model because it scales with customer volume, silicon complexity and mix. Third, the newly announced silicon business adds direct chip revenue from Arm-designed CPUs, with the AGI CPU for AI data centers. Arm guided that the chip business will start contributing meaningfully to revenue only from FY28 onward.

Strategically, Arm has now organized the business around three AI-centric domains: Cloud AI, Physical AI and Edge AI.

- **Cloud AI** is the data-center and infrastructure business, where CPUs increasingly act as the orchestration layer for inference, agentic workloads, memory, storage and accelerator coordination. This is also where Arm has extended into silicon with the AGI CPU.
- **Edge AI** is the broad device-side franchise, covering smartphones, PCs, wearables, XR, smart home and embedded systems, where AI inference runs locally under tight power, thermal and privacy constraints.
- **Physical AI** is the autonomous-systems business, spanning vehicles, robotics and machines that must sense, decide and act in the real world under strict latency, safety and reliability constraints.

The common theme across all three segments is the same economic flywheel: a growing total addressable market, higher Arm market share penetration and a structurally higher royalty rate per chip driven by the adoption of Armv9 and rising CSS penetration. **Cloud AI** is where the Arm story changes most rapidly. Arm highlighted AI infrastructure is becoming more CPU-led as inference and agentic workloads require more orchestration, memory management and accelerator coordination, while CSS is increasingly helping customers move faster in the data center. We therefore estimate Cloud AI TAM rises from US$61 Bn in FY26E to US$171Bn in FY31E, with Arm market share increasing from 29% to 54% and **royalty revenue growing from US $0.38 Bn to US$2.06 Bn**, implying a **40% CAGR.**

The next leg of that evolution is the **AGI CPU**, which Arm introduced in March 2026 as its first silicon product and as the next step beyond IP and CSS into chips. Arm guided that the AGI CPU is purpose-built for agentic AI infrastructure and backed by broad ecosystem support, positioning it as a new class of processor for large-scale AI data centers. This will be the most important growth driver of the broader Arm story, and **we anticipate the AGI CPU business can reach US$15bn by FY31E.**

**Edge AI** remains the largest segment today and still the foundation of Arm's royalty model. Edge AI is benefitting from AI workloads moving into smartphones, PCs, wearables and embedded devices, while Armv9 and CSS are increasing content and royalty intensity per chip. **We estimate Edge AI royalty revenue rises from US$2.13bn in FY26E to US$4.16bn in FY31E, implying a 14% CAGR,** and that the segment remains the largest royalty pool in FY31E. Within Edge AI, Mobile AP remains the core profit engine, while Other Mobile is a smaller, mature supporting handset silicon segment; Consumer is the fastest-growing sub-segment, driven primarily by CSS penetration in PCs; and IoT & Embedded provides broad-based, lower-growth expansion across connected edge devices.

**Physical AI** is also a major growth pillar focusing on autonomous vehicles and robotics which must sense, decide and act in the real world under tight latency, safety and reliability constraints, with Zena CSS designed to reduce complexity and accelerate deployment. We estimate Physical AI TAM increases from US$25bn in FY26E to US$42bn in FY31E, market share rises from 46% to 55%, an**d royalty revenue grows from US$0.18 Bn to US$0.52 Bn, implying a 23% CAGR.**

Overall, we estimate **Arm TAM rises from US$260 Bn to US$411 Bn**, market share from 50% to 60%, and royalty revenue from **US$2.69 Bn to US$6.74 Bn**, a **20% CAGR**; including AGI CPU silicon, we anticipate total revenue rises from **US$4.92 Bn to US$25.90 Bn**, implying a **40% CAGR (Exhibit 59, Exhibit 60)**.

EXHIBIT 59: **We anticipate Arm's revenue to grow at a 40% CAGR over the next 5 years, primarily driven by the expansion of its new AGI CPU business, which we expect to reach US$15bn by FY31. Growth is expected to be further supported by strong momentum in royalty revenue, which we forecast to grow at a 20% CAGR. We expect NG NI to grow at 42% CAGR.**

| End Market | Metrics | FY26E | FY31E | CAGR 26-31E |
|---|---|---|---|---|
| Edge AI | TAM:<br>Market Share:<br>Royalty Revenue: | $174 Bn<br>~60%<br>$ ~2.13 Bn | $198 Bn<br>~66%<br>$ ~4.16 Bn | 3%<br><br>14% |
| Cloud AI | TAM:<br>Market Share:<br>Royalty Revenue:<br>Chip Revenue: | $61 Bn<br>~ 29%<br>$0.38 Bn<br>- | $ 171 Bn<br>~ 54%<br>$ 2.06 Bn<br>$ 15 Bn | 23%<br><br>40% |
| Physical AI | TAM:<br>Market Share:<br>Royalty Revenue: | $25 Bn<br>46%<br>$0.18 Bn | $42 Bn<br>55%<br>$0.52 Bn | 11%<br><br>23% |
| Total Revenue | TAM:<br>Market Share:<br>Royalty Revenue:<br>Chip Revenue:<br>Total Revenue: | $260 Bn<br>50%<br>$2.69 Bn<br>-<br>$4.92 Bn | $411 Bn<br>60%<br>$6.74 Bn<br>$15 Bn<br>$25.90 Bn | 10%<br><br>20%<br><br>40% |
| Net Income | Non-GAAP NI: | $1.89 Bn | $10.96 Bn | 42% |

Source: Company reports, Bernstein analysis and estimates

EXHIBIT 60: **We expect Arm's revenue to grow at an impressive ~40% CAGR over the next 5 years, reaching US $26Bn in FY31. This growth is primarily driven by the expansion of the new AGI CPU business, which we forecast will reach US$15Bn in FY31, in line with guidance.**

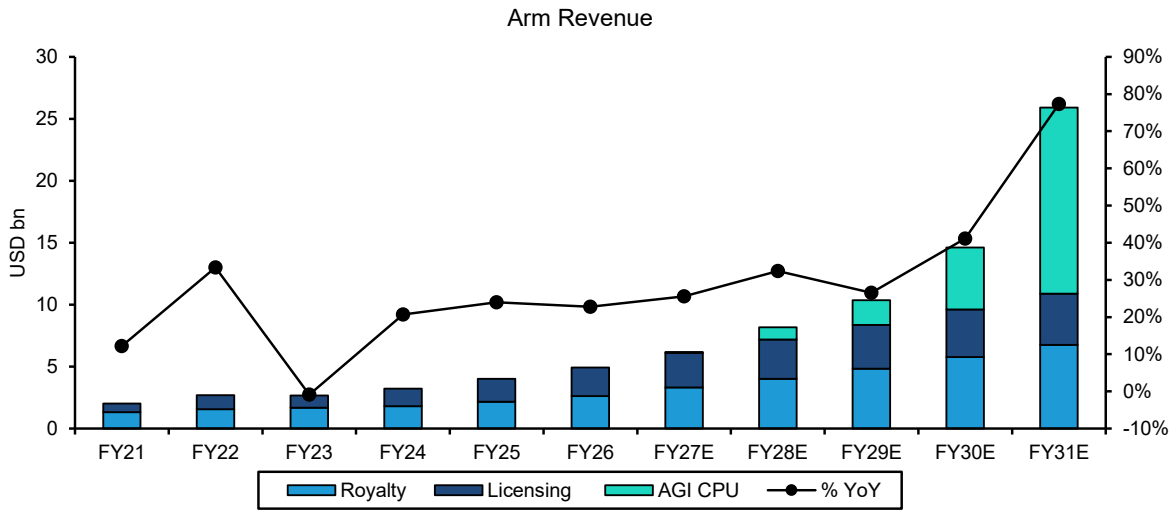


Source: Company reports, Bernstein analysis and estimates

## CLOUD AI: RENAISSANCE OF CPU IN DATA CENTERS

Cloud AI encompasses end markets of cloud compute, networking equipment and other infrastructure (Exhibit 61), of which cloud compute is the most important segment of the three, with an estimated $33bn TAM as of FY26 and also the fastest growing — we estimate FY31E TAM of $137bn and CAGR of 33%.

EXHIBIT 61: **We expect Cloud AI revenue to grow at a 114% CAGR over the next five years, primarily driven by the ramp-up in Silicon revenue, which we forecast to reach US$15Bn by FY31. Royalty revenue should also support growth, expanding at a 40% CAGR.**

| End Market | Metrics | FY26E | FY31E | CAGR 26-31E |
|---|---|---|---|---|
| Cloud Compute | TAM:<br>Market Share:<br>Royalty Revenue:<br>Chip Revenue: | $33 Bn<br>31%<br>$0.27 Bn<br>- | $137 Bn<br>60%<br>$1.82 Bn<br>$15.00 Bn | 33%<br><br>46% |
| Networking Equipment | TAM:<br>Market Share:<br>Royalty Revenue: | $17 Bn<br>32%<br>$0.09 Bn | $21 Bn<br>39%<br>$0.20 Bn | 5%<br><br>18% |
| Other Infra | TAM:<br>Market Share:<br>Royalty Revenue: | $11 Bn<br>16%<br>$0.02 Bn | $14 Bn<br>21%<br>$0.03 Bn | 4%<br><br>9% |
| Total Cloud AI: | TAM:<br>Market Share:<br>Royalty Revenue:<br>Chip Revenue:<br>Total Revenue: | $61 Bn<br>~ 29%<br>$0.38 Bn<br>-<br>$0.38 Bn | $171 Bn<br>~ 54%<br>$2.06 Bn<br>$15.00 Bn<br>$17.06 Bn | 23%<br><br>40%<br><br>114% |

Source: Company reports, Bernstein analysis and estimates

### Arm in AI servers

Arm's cloud and data center business is undergoing its most transformative phase since the company's founding, evolving from an IP licensing pure-play to a CPU solutions company. The Arm Everywhere event held on March 24-25, 2026 in San Francisco was the formal crystallization of this strategy, with CEO Rene Haas announcing Arm's first-ever production silicon product and laying out a road map that extends Arm's total addressable market from the historic $3bn IP TAM to over $100bn in cloud CPU alone by FY31. In terms of revenue, we model it to grow from $135mn in FY25/3 to $15bn in FY31/3.

Arm server CPU royalty revenue is modeled to grow from ~$135mn in FY25/3 to $1.8bn by FY31/3 on the back of two distinct drivers: hyperscalers (CSPs) accelerating their own Arm-based in-house CPU programs, and Nvidia bundling its Arm-based Grace / Vera CPUs with nearly every AI accelerator rack they ship. As more AI servers are deployed, Arm's CPU penetration is forecast to expand from 28% in CY25 to 58% in CY30, per our model assumptions.

### Agentic AI: CPU is back to center stage

Since the rise of LLM, GPU/ASIC accelerators has been the core of AI computing. While training clusters once required a dense 4:1 ratio to handle heavy data-loading, the focus shifted toward eliminating the 'CPU tax' that plagued high-scale inferencing. In custom inference-optimized deployments like Google's TPU v6e and Meta's Grand Teton, the GPU-to-CPU socket ratio moved to 8:1.

Agentic AI is pushing the CPU back to center stage (Exhibit 62) because AI systems are no longer just running a model once and returning an answer. They are increasingly acting like software workers that break problems into steps, call tools, manage memory, interpret results, and decide what to do next, which turns the CPU into the coordinator of the whole process rather

than a background component. Today, forward-looking cluster designs are reintegrating high-core-count CPUs to a leading role in the compute hierarchy, with some expecting GPU-to-CPU ratios narrowing to 1:1 in certain clusters. That's the case of clusters specifically designed to handle reinforcement learning environments, tool invocation, logic gating, and document merging. Specific 2026-era hardware pairings illustrate this narrowing gap:

- AMD Venice: 1 CPU to 4 MI455X GPUs per compute tray.
- NVIDIA Vera: 1 CPU to 2 Rubin GPUs (4 GPU dies) per superchip.
- Google TPU7x: 1 CPU to 4 TPU chips per scale-up unit.

This change matters because the architecture of AI infrastructure is becoming more balanced, not less. GPUs still handle the dense parallel math that makes modern AI possible, but CPUs are becoming more important in preparing data, scheduling work, managing memory, routing requests, handling control flow, and coordinating repeated interactions between models, applications, and external systems, so AI performance depends on the pairing of GPU acceleration with strong CPU orchestration.

The GPU/CPU pairing is especially important in agentic workloads because inference is turning into a loop instead of a single pass. A request may trigger retrieval, planning, tool use, intermediate reasoning, another model call, and then action, which means the GPU does the heavy compute while the CPU keeps the workflow moving efficiently; if the CPU is weak, expensive GPUs can sit underutilized, and the overall system becomes slower and less efficient.

Agentic AI also increases pressure on networking and distributed infrastructure, which strengthens the CPU's role even further (Exhibit 63, Exhibit 64, Exhibit 65). As workloads stretch across servers, clusters, and locations, the system has to move state, manage traffic, and coordinate resources in real time, so the CPU becomes critical not only inside the server but across the wider data-center fabric that supports autonomous AI execution.

Arm CPUs stand out in this environment because the new bottleneck is not only peak performance but efficient orchestration under power and space limits. As operators need more CPU capacity to support growing numbers of AI agents, Arm's pitch around performance per watt, high core density, and scalable data-center compute becomes more compelling, which is why agentic AI is helping bring CPUs back into focus and giving Arm a stronger strategic role in the next phase of AI infrastructure.

EXHIBIT 62: **Arm argues that agentic AI shifts more work back to the CPU: accelerators generate tokens, but CPUs orchestrate the agents, memory and workflows needed to deliver answers, making Arm's efficient CPU architecture increasingly critical in AI data centers**

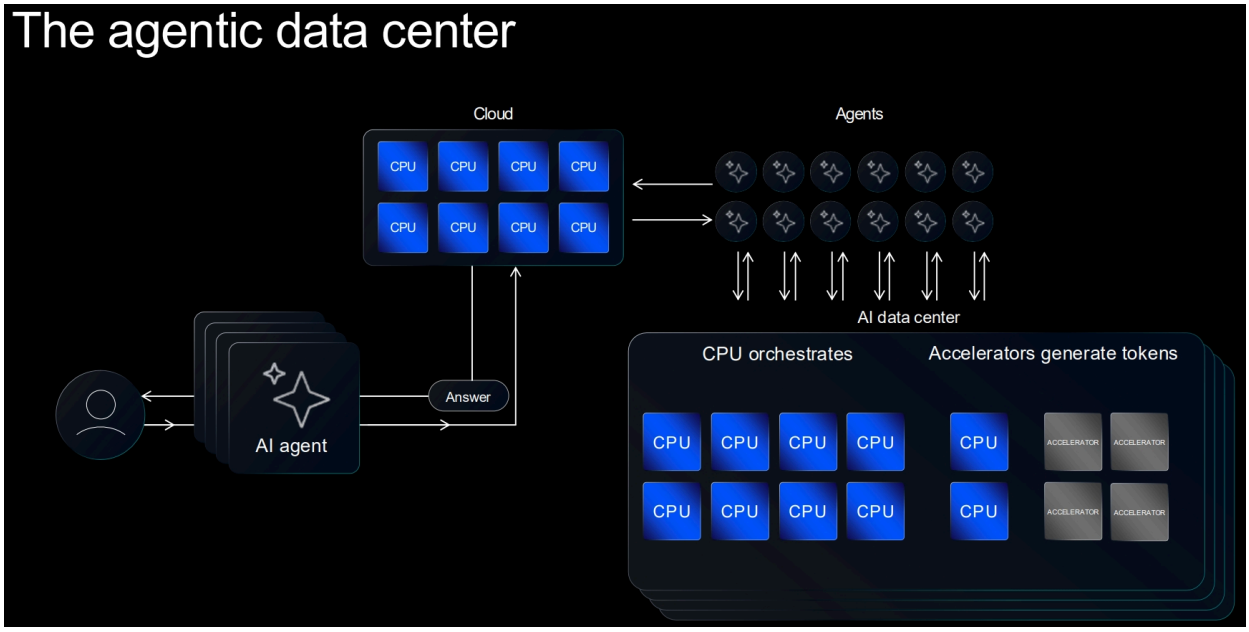


Source: Arm

EXHIBIT 63: **CPU is expected to play a more important role within inference, in the agentic area.**

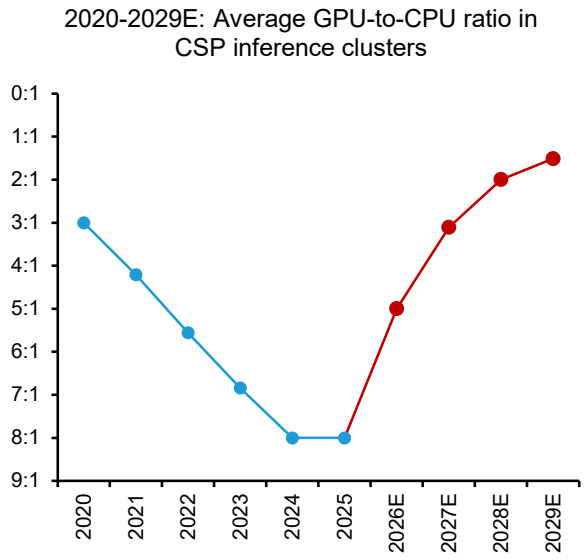


Source: Ciena estimates, Bernstein analysis.

EXHIBIT 64: **Agentic AI shifts compute balance toward CPUs, with CPU share rising from ~14% in Traditional LLMs to 50%, highlighting CPUs' growing orchestration role alongside GPUs in AI workloads at scale**

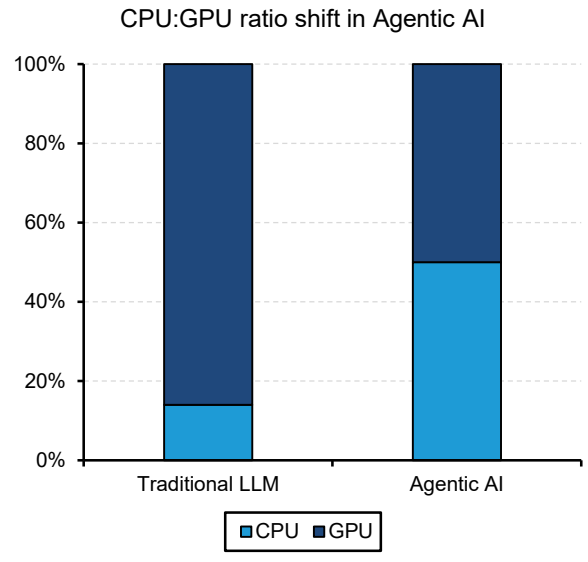


Source: TrendForce, Bernstein analysis

EXHIBIT 65: **We expect the server CPU market to increase from US$33bn in 2025 to US$137bn in 2030, accelerating at a 34% CAGR, driven by Agentic AI adoption.**

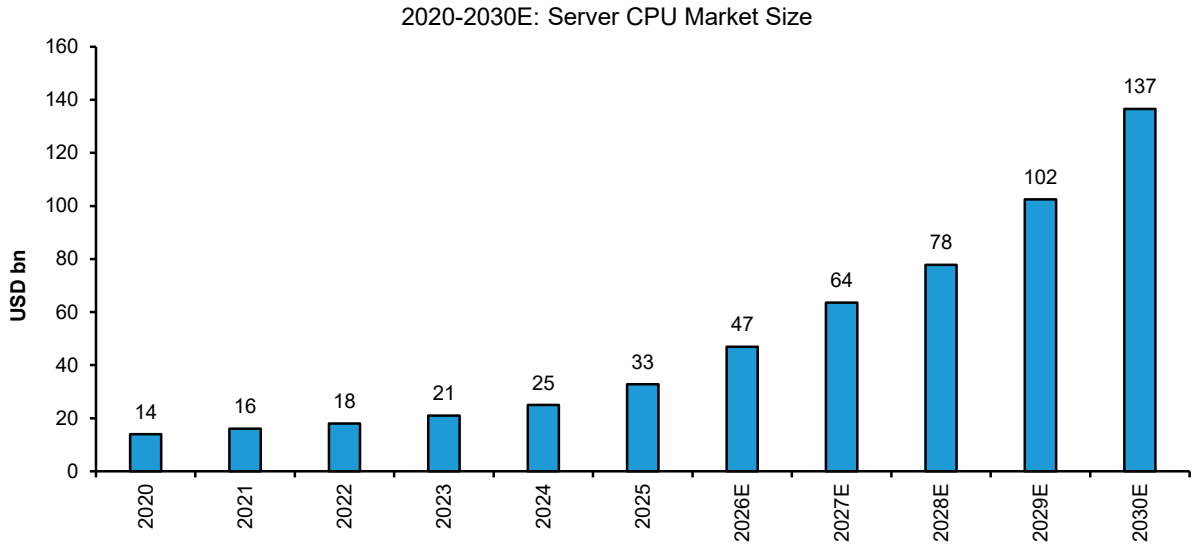


Source: Mercury, Company reports, Bernstein analysis and estimates

### Why Arm over x86 in servers?

The core economic case is performance per watt and total cost of ownership. AWS states Graviton delivers up to 40% better price performance versus comparable x86 instances, up to 20% lower cost, and up to 60% lower energy consumption for equivalent workloads on certain use cases. Microsoft Azure notes that each vCPU on Cobalt 100 maps to a full physical core, making performance more predictable for server workloads versus x86 designs that use hyperthreading to inflate vCPU counts.

Beyond economics, hyperscalers gain architectural control. Designing in-house CPUs lets them shape the memory subsystem, I/O, accelerator interconnects, and security offload to fit their own infrastructure rather than relying on the roadmap of a merchant silicon vendor. This is why each major CSP has now committed to a multi-generation internal CPU program built on Arm Neoverse cores.

Above all, power consumption is one of the biggest reasons that Arm CPUs are much more likely to be adopted in AI data centers. The world is facing an energy issue that not enough electricity is available to power the new data centers. According to our forecast, AI accelerator power will almost 3x to reach 37 GW to be deployed in 2028. Including the power for CPU, memory, networking, etc, the total power consumption of AI data center would surpass 50 GW. As we can't do much to cut down power consumption of GPU/AI accelerators, using Arm server to reduce power consumption becomes the only viable choice.

EXHIBIT 66: **We estimate AI accelerator GPU power reaches ~37GW in 2028 versus 12.8GW in 2025, a 43% CAGR over 2025–2028E. Most of the growth in AI power is coming from higher TDP.**

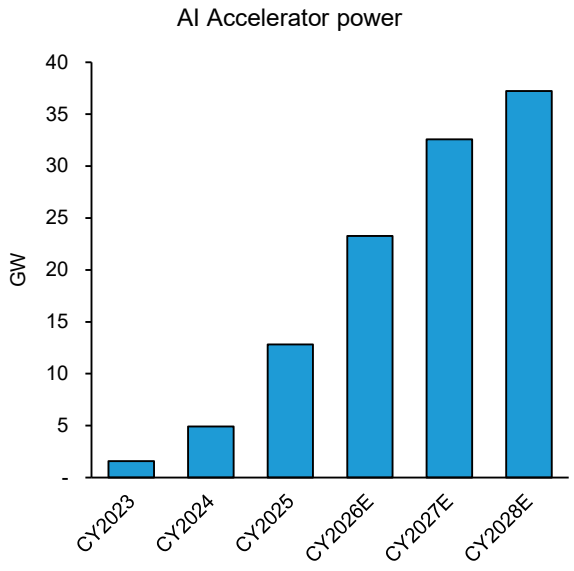


Source: Company reports, Bernstein analysis and estimates

EXHIBIT 67: **CPU represents between 40% and 60% of power consumption in a general server, but only 5–10% in an AI server, where most power is consumed by AI accelerators.**

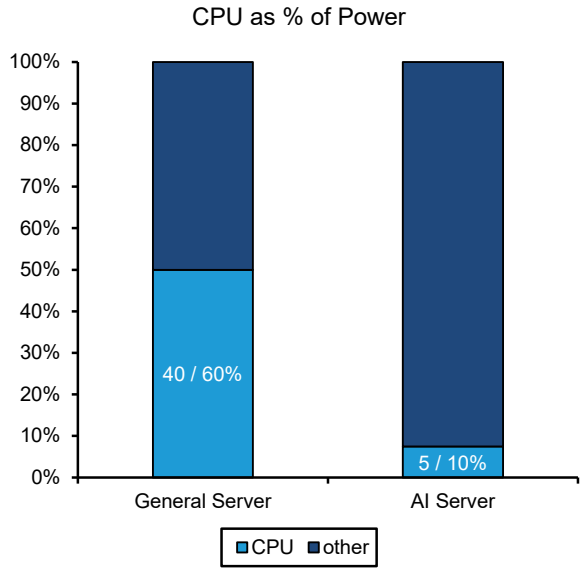


Source: Bernstein analysis and estimates

### CSP CPU Roadmaps and Arm Penetration

AWS, Microsoft, and Google have each built Arm-based CPU programs (Exhibit 68) that are deepening Arm's data center presence across generations. According to Arm, their penetration at data centers has surpassed 50% among top hyperscalers with over 1bn Neoverse cores deployed.

- **AWS:** As the largest IaaS provider with 47% market share as of 2024, AWS is also the dominant deployer of Arm-based instances, with 57.2% of Arm instances across the public cloud as of 2024 year-end being deployed at AWS. AWS was the first mover to deploy Arm processors on public cloud, debuting their first arm-based instance back in 2018. As of December 24, according to Liftr Insights, Arm-based **Graviton** represented 25% of the total instances deployed at AWS.
- **Microsoft Azure:** As the second largest IaaS provider with 17% market share as of 2024, Azure also heavily deploys Arm-based instances, with an estimated 33% of chips being Arm-based as of Q4 2024. This is a rapid ramp considering their first Arm processor **Cobalt** 100 debuted only in 2024. (What performance and efficiency does Microsoft Azure Cobalt 100 VMs deliver using Arm Neoverse performance? - Arm Newsroom)
- **GCP:** Google was the third biggest IaaS provider as of 2024 with 7% market share, and they too have ramped Arm instances rapidly, with an estimated 21% share of all instances being **Axion**, which was introduced in 2024.

### Why Nvidia uses Arm CPUs

Nvidia's motivation for Arm is distinct from the hyperscalers: in AI and HPC systems, the critical bottleneck is more in data movement between the CPU, memory pool, and GPU - rather than raw CPU instruction throughput or enterprise application compatibility. Nvidia says its **Grace** CPU (72 Neoverse V2 cores per die, 144 cores in the Grace CPU Superchip using dual-die NVLink-C2C) delivers up to 1.0 TB/s of LPDDR5X memory bandwidth and benefits from the NVLink-C2C interconnect that carries up to 900 GB/s bidirectional bandwidth between the CPU and Blackwell GPU - versus PCIe Gen 5's ~128 GB/s. The second-generation NVLink-C2C used in GB200 racks achieves 1.8 TB/s of bidirectional bandwidth, 7x faster than PCIe Gen 6. (NVIDIA Grace CPU Delivers World-Class Data Center Performance and Breakthrough Energy Efficiency | NVIDIA Technical Blog, NVIDIA Grace CPU and Arm Architecture | NVIDIA)

Grace was first deployed in the **GH200 Grace-Hopper Superchip** for HPC/AI clusters starting in 2023. The major inflection came in 2025 when Nvidia standardized on the **DGX GB200** rack form factor - bundling one Grace CPU with two Blackwell GPUs per Superchip, with each DGX GB200 rack containing 36 Grace CPUs and 72 Blackwell GPUs, totaling 2,592 Arm Neoverse V2 cores per rack. This shift from the standalone HGX/MGX GPU tray format toward integrated DGX rack-scale deployment materially increased Arm CPU unit volume in AI server shipments.

The Nvidia **Vera** CPU (successor to Grace, launching with Rubin GPUs in 2026) further raises the bar: Vera delivers up to 1.2 TB/s of memory bandwidth - twice the bandwidth at half the power compared to traditional x86 CPUs - and is built on Nvidia's proprietary Olympus CPU cores, a custom Arm architecture designed for Reinforcement Learning and agentic AI workloads. Vera enables software environments to run up to 50% faster with twice the efficiency of conventional CPU infrastructure. (Next Gen Data Center CPU | NVIDIA Vera CPU)

EXHIBIT 68: **Comparison of Server CPU specifications for Arm ASICs, Nvidia custom CPUs and x86.**

| Specification | Arm Cloud ASICs | | | NVIDIA | x86 Flagships | |
|---|---|---|---|---|---|---|
| | Microsoft Cobalt 200 | Google Axion C4A | AWS Graviton 5 | NVIDIA Vera | Intel Xeon 6 6980P | AMD EPYC 9965 |
| ISA | Arm v9 | Arm v9 | Arm v9 | Arm v9 | x86-64 | x86-64 |
| Microarchitecture | Neoverse V3 (CSS) | Neoverse V2 | Neoverse V3 | Nvidia Olympus (Custom) | Granite Rapids | Turin Zen 5c |
| Process Node | TSMC 3nm | TSMC 5nm | TSMC 3nm | TSMC 3nm | Intel 3 (~4nm) | TSMC N4P |
| Deployment | Cloud (Azure) | Cloud (GCP) | Cloud (AWS) | Universal | Universal | Universal |
| Release Date | Nov 2025 | Oct 2024 | Dec 2025 | 2H 2026 | Sep 2024 | Oct 2024 |
| Physical Cores | 132 | 72 | 192 | 88 | 128 | 192 |

Source: Company disclosures, Bernstein analysis.

EXHIBIT 69: **Arm Server CPU has seen strong growth since 2023...**

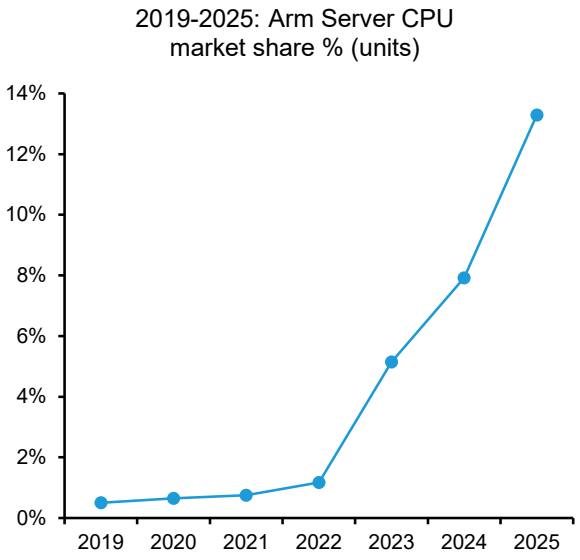


Source: Mercury estimates, Bernstein analysis.

EXHIBIT 70: **...and we expect Arm's market share to continue the strong trend of growth.**

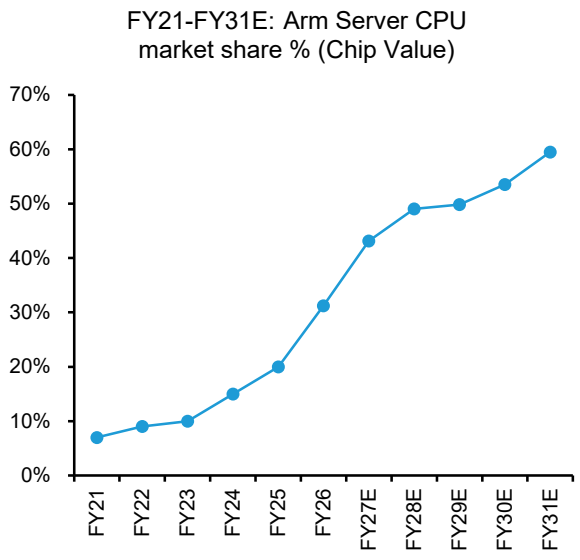


Source: Company disclosures, Bernstein estimates and analysis.

EXHIBIT 71: **Roadmap of CPUs based on Arm IP**

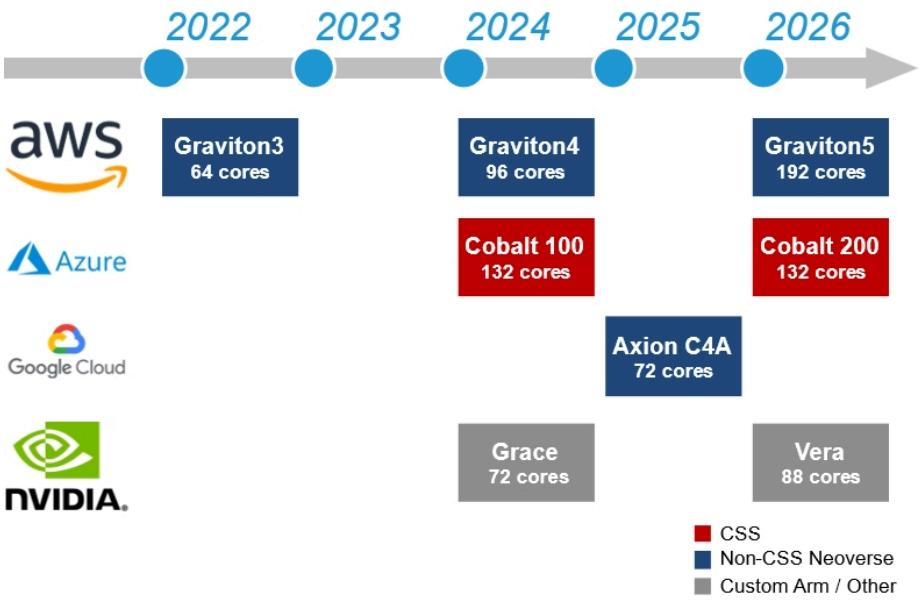


Source: Company disclosures, Bernstein analysis

### Royalty increase in Arm server CPUs driven by core count and increased royalty rate

Even before the AGI CPU contribution, Arm's cloud CPU royalty model is already becoming materially more attractive. Arm said in its Q3 2026 earnings that cloud AI / general-purpose data-center royalty revenue grew more than 100% YoY, Neoverse deployments surpassed 1bn cores, and Arm's share among top hyperscalers was expected to approach 50%.

The first driver is that **core counts** are rising fast (Exhibit 72). AWS Graviton 5 now has 192 cores, up 2x versus Graviton 4; NVIDIA Vera has 88 Arm-based cores versus 72 for Grace; and Microsoft's Cobalt 200 has 132 cores. At Arm Everywhere event, management indicated that the next wave of cloud CPUs could move into the 200-300 core range and, in some cases, toward 500 cores over time.

We think the underlying reason is the same one driving the AGI CPU opportunity: the workload mix is changing. Agentic and always-on inference shifts more of the bottleneck toward orchestration, memory handling and coordination work, which in turn increases the need for higher core count, power-efficient CPUs in AI clusters. Arm themselves said the shift toward agent-based inference is redefining AI data-center designs and raising demand for CPU chips with even more power-efficient cores.

EXHIBIT 72: **Core count for server-side CPUs is increasing.**

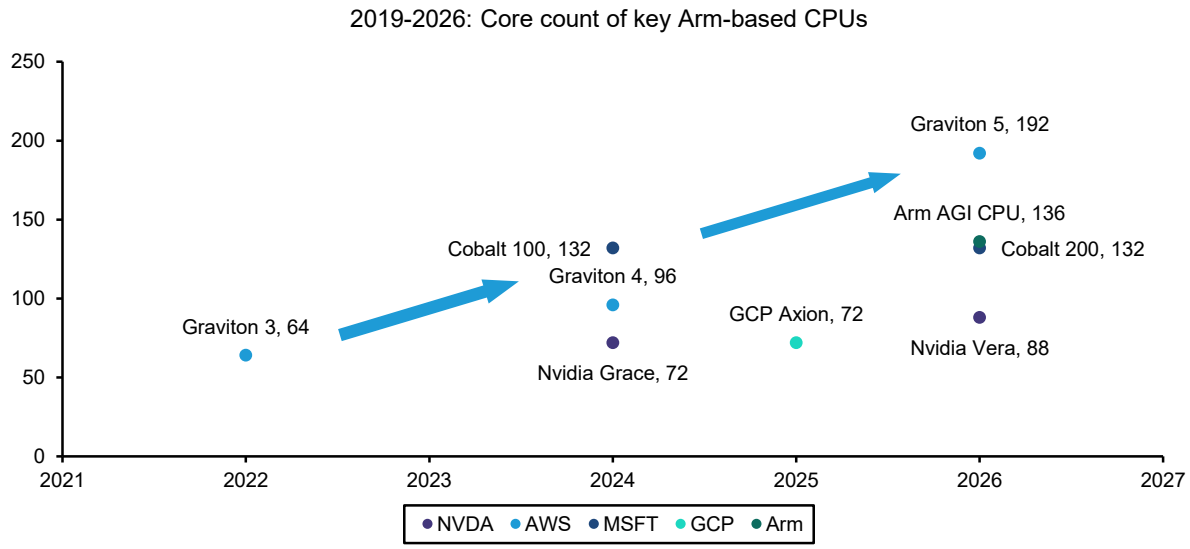


Source: Company disclosures, Bernstein analysis.

The second driver is **higher royalty per core**. Currently, Arm's licensees pay roughly $0.50/core for first-generation v9 server CPUs, around $1/core for first-generation CSS (Compute Subsystem), and potentially $1.50/core for later CSS generations, with more upside in future generations. In other words, Arm is not only getting paid on more cores, but on more valuable cores, and CSS is central to that step-up.

Arm describes CSS as pre-integrated, pre-validated, performance-validated subsystem products that help partners reach production silicon faster. In other words, CSS sits above the core level: it wraps CPU cores with coherent interconnect, MMU/NOC system IP, memory channels, and I/O so a chip company does not have to assemble the whole server platform from scratch.

The company has said CSS adoption is a major tailwind to royalty growth, and customers can save substantial engineering effort and accelerate time-to-market by up to roughly a year. Arm also says CSS customers have reported getting from kickoff to working silicon in 13 months and saving 80 engineering years by offloading non-differentiated design and validation work. That higher value proposition is what allows Arm to charge materially higher royalties than for core IP alone.

CSS includes more than Neoverse cores alone, bundling the cores with CMN mesh, system IP, system management, power management, software, and development tools needed to bring a platform to market quickly. Arm positions this as the fastest path to production silicon for cloud, AI, 5G, and networking chips because partners can focus on customization instead of rebuilding the common infrastructure around the cores.

The cadence of product refreshes also matters. Management's notes suggest server CPUs used to refresh every two to three years, but increasingly move on an annual cadence, which in practice makes CSS more valuable because customers need faster turn times to stay on the latest node and architecture. That dynamic should support both more frequent licensing activity and a faster mix shift toward higher-royalty generations.

In cloud specifically, the early evidence of that shift is already visible. Microsoft's Cobalt family is explicitly built on Neoverse CSS. Company also mentioned that CSS should exceed 10% of cloud AI royalty by FY27, and rise above 30% by FY29 and exceed 50% by FY31. If that progression is broadly right, the royalty mix should keep skewing toward higher-value server designs over the next several years.

In conclusion, we see strong growth in Arm's server royalty revenue (Exhibit 74), reaching $2bn in FY31 from ~$135mn in FY25, driven by the increasing core count as well as higher per-core royalty rate, resulting in as structural server CPU royalty growth even before giving Arm any credit for its own AGI CPU silicon revenue.

EXHIBIT 73: **We believe per-chip royalty dropped in FY26 due to Grace, but expect strong growth going forward due to increase in royalty rate.**

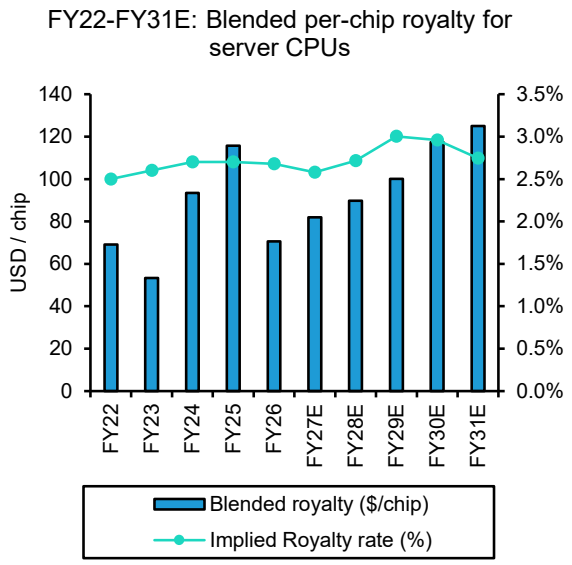


Source: Company disclosures, IDC, Mercury, Bernstein estimates and analysis.

EXHIBIT 74: **We expect a strong growth in cloud royalty revenue.**

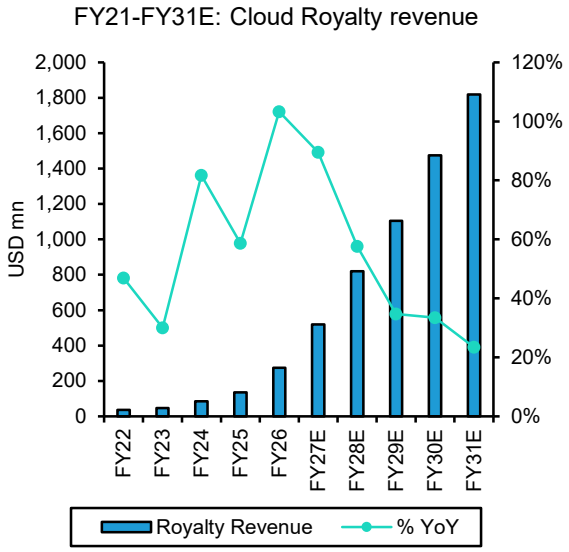


Source: Company disclosures, Bernstein estimates and analysis.

### Arm AGI CPU: Drastic change in business model to capture the growing opportunity

At Arm Everywhere event in March 2026, Arm announced an important strategic move to provide their own silicon, the Arm AGI CPU, after decades of monetizing primarily through IP licensing and royalties, making clear that the silicon offering will complement - rather than replace - its existing IP and CSS businesses.

Their rationale was that not every client can design CPUs of their own, so companies like Meta, - which Arm identified as the first customer, lead partner and co-developer - and OpenAI can just buy Arm's CPUs, while other CSPs such as Amazon, Google and Microsoft who can keep licensing Arm IPs for their CPU endeavours. Alongside Meta and OpenAI, they also confirmed many other partners including Cerebras, Cloudflare, F5, Positron, Rebellions, SAP and SK Telecom. These customers are expected to deploy Arm AGI CPU for agentic usage. Arm partners with OEMs and ODMs, as well as other supply chain partners across cloud, memory, networking, manufacturing and many other aspects, such as AWS, Broadcom, Google, Marvell, Micron, Nvidia, Samsung, SK hynix and TSMC among many others.

Financially, the model change is clearly margin dilutive in percentage terms, but very likely accretive in gross profit dollars and operating profit dollars. Management's event target was for the AGI CPU business to reach $15bn of revenue by FY31, with total company revenue of $25bn and non-GAAP EPS of $9, with AGI CPU OPM is expected to exceed 30% by FY31. We believe the key point is that Arm is choosing to trade some margin rate for a dramatically larger profit pool. Just as importantly, management said the required engineering investment is already largely embedded in today's cost base, meaning the near-term P&L is already carrying much of the R&D burden, while the revenue opportunity from silicon only starts to appear from much later on to bring the operating leverage once volume scales.

Our estimate based on a more conservative assumption and we forecast $15bn chip value as of FY31 (Exhibit 75) which roughly equates to ~2.5mn CPU units based on $4.8k ASP in FY31. We believe Arm silicon should generate 42% of its OP in FY31 in our forecast (Exhibit 76).

**Our Server CPU industry model can be downloaded here: (ARM Server CPU Industry Model)**

EXHIBIT 75: **We forecast Arm AGI CPU sales to reach $15 Bn by FY31.**

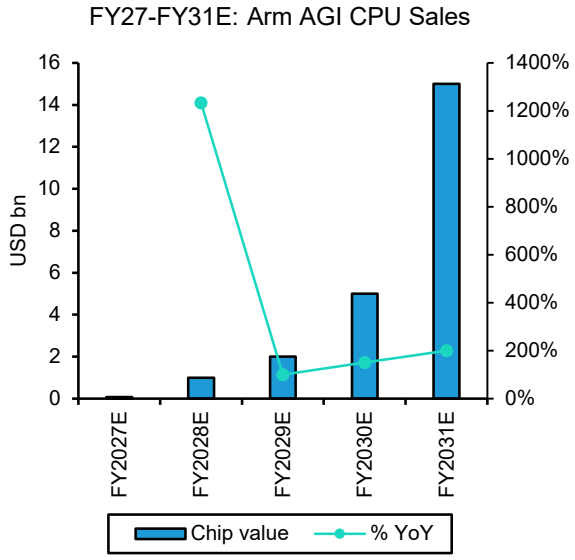


Source: Company disclosures, Bernstein estimates and analysis.

EXHIBIT 76: **We believe OP contribution from Arm silicon could go as high as 42% as of FY31.**

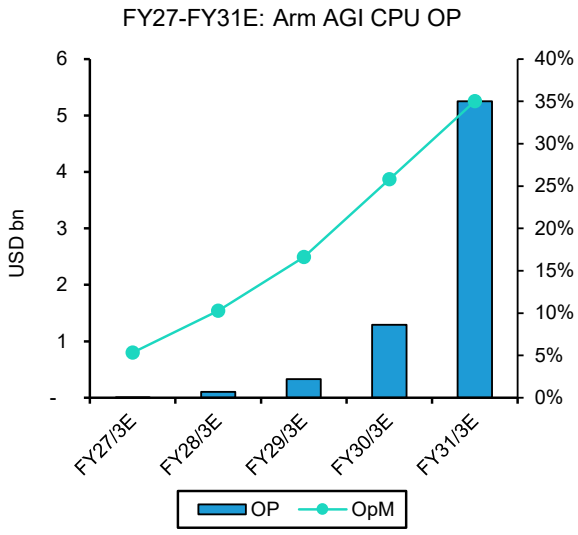


Source: Company disclosures, Bernstein estimates and analysis.

## EDGE AI: ARCHITECTURE UPGRADES TO DRIVE ROYALTY INCREASE

Edge AI is currently the largest segment, with our estimated US$2.13Bn of royalty revenue in FY26E, accounting for ~80% of Arm royalty revenue. While Edge AI will likely grow at considerably slower pace compared with Cloud AI and Physical AI, we still expect it will grow meaningful at 14% CAGR over the next five years and it will remain the largest royalty end-market also in FY31, with US$ 4.16Bn sales and representing ~60% of royalty revenue (Exhibit 77).

Edge AI can be divided into four different end-markets:

1. **Mobile Application Processor (AP):** This is the main application processor SoC in smartphones and tablets, i.e. the central compute engine that runs the rich OS, apps, graphics, camera stack, and increasingly on-device AI workloads. It remains the core of Arm's royalty revenue franchise, supported by a very deep software ecosystem, the Armv9 transition, and growing CSS adoption. It accounts for ~60% of total Edge AI% and 45% of total royalty revenue. We estimate Mobile AP to grow from **US$ 1.21 Bn in FY26E to US$ 2.58 Bn in FY31E**, implying a **16% CAGR.** We anticipate Arm will maintain its monopoly thanks to their solid ecosystem.

2. **Other Mobile:** Mobile silicon content that sits outside the main AP, such as connectivity and modem related chips, as well as touchscreen and image-sensing related subsystems. In other words, it is the supporting semiconductor layer that enables the handset to connect, sense, and interface with the user and network. We estimate this is a smaller and more mature segment, with royalty revenue increasing from **~US$170 Mn in FY26E to ~US$190 Mn in FY31E**, equivalent to a **3% CAGR**. TAM remains broadly stable at **US$17bn**, while Arm's share stays at roughly 70%.

3. **Consumer:** This covers higher-performance consumer electronics such as laptops and PCs, gaming devices, smart TVs, and XR systems. These platforms require silicon that can support rich operating systems, advanced graphics and media processing, and increasingly sustained local AI inference within constrained power and thermal envelopes. We estimate Consumer is the fastest-growing Edge AI segments, with royalty revenue rising from **~US$340 Mn in FY26E to ~US$838 Mn in FY31E**, implying a **20% CAGR**. Growth will be primarily driven by the fast penetration of CSS in PC. TAM increases from **US$68 Bn to US$76 Bn**, while Arm's market share expands from roughly 36% to 50%.

4. **IoT & Embedded:** Purpose-built connected edge devices, including MCU, sensors, smart-home products, industrial controllers, and medical or embedded systems. The silicon here is optimized for low power, low cost, real-time operation, long product lifecycles, and inference close to the sensor or control loop rather than general-purpose consumer computing. We estimate royalty revenue grows from **~US$400 Mn in FY26E to ~US$557 Mn in FY31E**, implying a **7% CAGR**. TAM

expands from **US$46 Bn to US$54 Bn**, with Arm's market share increasing from roughly 50% to 56%.

EXHIBIT 77: **We expect Edge AI revenue to grow at a 14% CAGR over FY26–31E, below the overall royalty CAGR of 20%, primarily driven by v9 and CSS penetration. The consumer segment is expected to deliver the strongest growth, supported by a ramp-up in CSS penetration in PC.**

| End Market | Metrics | FY26E | FY31E | CAGR 26-31E |
|---|---|---|---|---|
| Mobile AP | TAM:<br>Market Share:<br>Royalty Revenue: | $42 Bn<br>>99%<br>$ 1.21 Bn | $51 Bn<br>>99%<br>$ ~2.58 Bn | 4%<br><br>16% |
| Other Mobile | TAM:<br>Market Share:<br>Royalty Revenue: | $17 Bn<br>~70%<br>$0.17 Bn | $17 Bn<br>~70%<br>$0.19 Bn | 0%<br><br>3% |
| Consumer | TAM:<br>Market Share:<br>Royalty Revenue: | $68 Bn<br>~36%<br>$0.34 Bn | $76 Bn<br>~50%<br>$0.84 Bn | 2%<br><br>20% |
| IoT & Embedded | TAM:<br>Market Share:<br>Royalty Revenue: | $46 Bn<br>~50%<br>$ 0.40 Bn | $54 Bn<br>~56%<br>$ 0.56 Bn | 3%<br><br>7% |
| Total Edge AI | TAM:<br>Market Share:<br>Royalty Revenue: | $174 Bn<br>~60%<br>$ 2.13 Bn | $198 Bn<br>~66%<br>$ 4.16 Bn | 3%<br><br>14% |

Source: Company reports, Bernstein analysis and estimates

### Arm in smartphone: reaping higher royalty by enabling edge AI

Even as smartphone unit volumes peaked around 2017 and declined further after Covid, the mobile semiconductor market has continued to grow, driven by increasing semis content per device. This reflects the growing compute intensity of mobile SoCs, along with increasing DRAM and NAND content. We expect this structural rise in silicon intensity to persist and likely accelerate, as on device AI is becoming a key differentiator in next generation smartphones.

With the rise of Generative AI in 2023, smartphone brands have leveraged the opportunity to launch new 'AI phones' since early 2024. While AI has long existed in smartphones through functions such as natural language processing and computational photography, the newer wave of AI phones is increasingly defined by **on-device GenAI capabilities**. In practice, this means more AI workloads running locally on the handset, making metrics such as SoC AI performance, memory content, and model size more important.

What makes on-device AI important is not only raw performance, but also the user experience it enables. Arm emphasizes that enhanced photography, immersive gaming, smart assistants, and real-time translation can increasingly run **directly on device** with **low latency, strong privacy, and efficient power use**. That matters because many of the most compelling smartphone AI features require immediate response times and need to work reliably in everyday usage. Arm's latest Lumex platform is explicitly designed around this goal, combining CPUs, GPUs, interconnect, physical implementations, and software integration for heterogeneous AI compute at scale.

### What on-device AI can do

- **Smart AI-powered photography:** Enhances photos in real time by applying filters and correcting image details directly on the device.

- **Immersive mobile gaming:** Improves graphics quality and power efficiency, enabling smoother and more realistic gaming experiences.
- **Context-aware personal assistants:** Helps manage tasks, optimize schedules, suggest routes, and surface relevant app actions with low on-device power use.
- **Real-time voice translation:** Enables fast, private, and always-available translation directly on the handset, even without internet access.

We believe the industry is still in the early stages of the AI phone era, with several rounds of AI feature and hardware upgrades still ahead. As more workloads shift on device, smartphones will require stronger SoC AI capability, higher memory content, and tighter hardware-software integration. In our view, this should support continued ASP expansion and further growth in semiconductor content per device. **We see this structural trend as a key driver of Arm's growth in mobile, as OEMs move quickly to adopt the latest Armv9 and CSS platforms in order to deliver the highest-performance and most energy-efficient on-device AI experiences (Exhibit 78).**

EXHIBIT 78: **AI smartphone penetration was ~30% in 2025, or ~390Mn units, with the vast majority in the premium segment. Gartner estimates this will rise to ~70% of shipments by 2030, with AI becoming standard across the premium segment and reaching ~50% in low-end devices.**

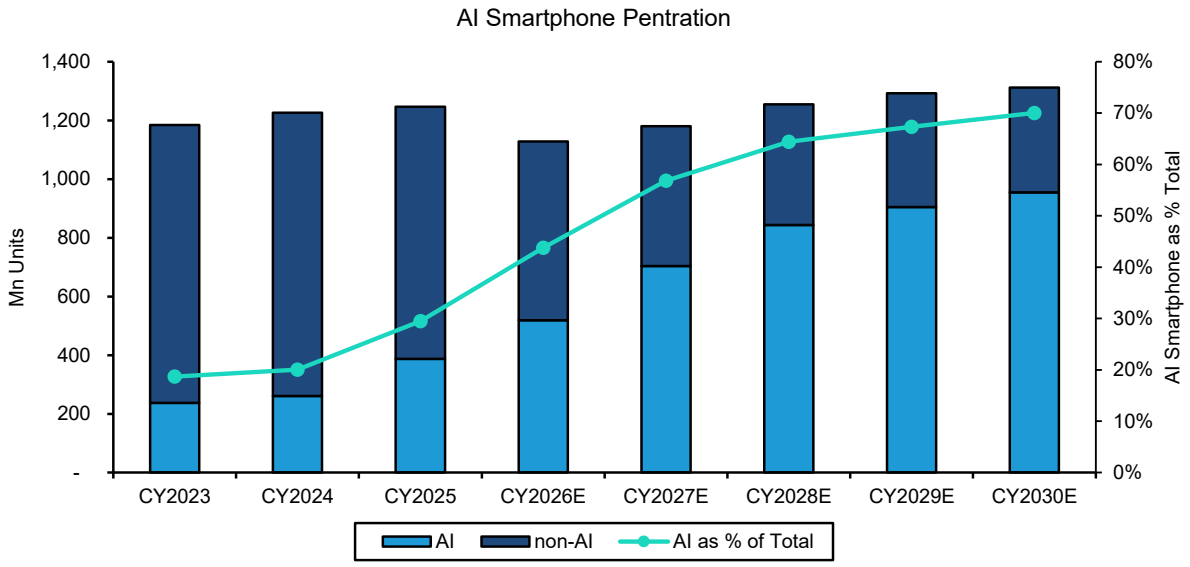


Source: Gartner March 2026, Bernstein analysis

## Arm Mobile Business: History

Arm's mobile business is the historical core of the company and, even after Arm's expansion into cloud, automotive and IoT, it remains its largest end market as percentage of royalty revenue. The mobile story is also the story of Arm itself: the company's licensing model took shape in the 1990s. Its designs became the flagship architecture for mobile phones through products such as the Nokia 6110 launched in 1998. By the end of the decade, Arm had become the leader in embedded RISC, with mobile already a major part of its licensing base.

Arm's origins are also closely tied to Apple (covered by Mark C. Newman). In 1990, Apple was the first external investor in the joint venture that became Arm, providing capital to help commercialize the architecture. Apple also created the first major product based on an Arm processor, the Newton MessagePad, launched in 1993 with the ARM610 at its core. That early strategic role helped establish Arm's commercial credibility and gave Apple a uniquely privileged position in the ecosystem. As a result, Apple is still widely believed to benefit from exceptionally favorable economics with Arm, including materially lower licensing fees and royalty rates than most other customers.

In the 2000s, Arm became central to the smartphone transition. Early Arm-based smart devices included Ericsson's R380, Palm devices, and BlackBerry products. Under CEO Warren East, Arm launched the Cortex product family, which strengthened its position across mobile and embedded markets. At the same time, Intel (Covered by Stacy Rasgon) exited mobile, while Android and iOS emerged as the winning smartphone operating systems. Apple played a major role in Arm's next phase of growth. After launching the iPhone in 2007 with an Arm-based processor, Apple acquired PA Semi, signed an architecture license with Arm, and later developed its own A-series chips. This cemented Arm's position at the heart of the modern smartphone era.

EXHIBIT 79: **The Apple Newton was one of the first consumer products to use the ARM architecture.**


Source: Wikimedia Commons photo by Felix Winkelnkemper

EXHIBIT 80: **The Cortex product family was introduced in 2004 and continues to be used today across multiple end-applications, including mobile devices through the Cortex-A and Cortex-X series.**

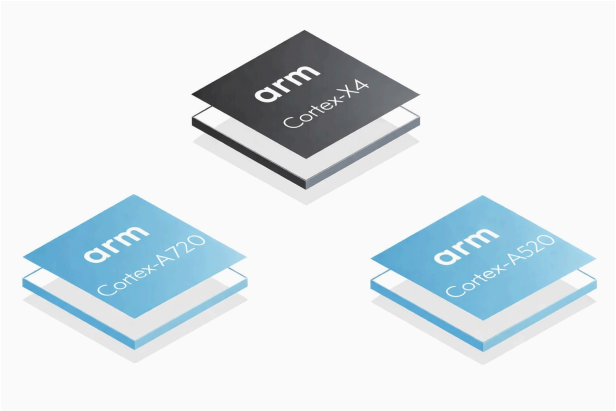


Source: Arm, Bernstein analysis

## Arm Mobile Business: TAM and market share

Arm's early design win became a structural industry advantage because, rather than manufacturing chips itself, it licensed CPU architecture and core IP, later expanding into GPUs, system IP and software, to a broad ecosystem of semiconductor vendors and device makers, while collecting royalties as chips shipped. **Since 2020, Arm estimates it has more than 99% share in mobile applications by chip value**, and that smartphones application processors accounted for about 45% of royalty mix in FYE25, with another 8% coming from other mobile, making mobile by far the largest royalty bucket even after years of diversification.

The strategic importance of Arm's position lies not just in its role in CPUs, but in the fact that the two dominant mobile operating systems (iOS and Android) and their application ecosystems were built around Arm architecture. This is the key reason Arm was able to achieve and sustain a total monopoly in mobile application processors for many years, making mobile a uniquely durable royalty stream. At the same time, the **handset market is now mature, so the story is less about unit growth and more about rising content and rising value per device**. This is why premium smartphones matter disproportionately to Arm. Arm's monetization is increasingly tied to richer SoCs, AI features and higher royalty-bearing silicon content rather than simple volume growth.

ARM Mobile Application Processor (AP) TAM increased from USD 21 Bn in 2015 to more than USD 40 Bn in 2025. This strong growth was driven by a significant rise in processor semiconductor content per smartphone, which increased from an average of USD 15 in 2015 to USD 32 in 2025. The expansion in semiconductor content was fueled by greater penetration of high-end smartphones and the growing need for processing power, which pushed up the ASP of both integrated and discrete application processors. While we expect ARM Mobile TAM to decline next year due to considerably lower volumes, we anticipate a return to strong growth thereafter, driven by continued increases in semiconductor content per device as smartphones require more processing power to support edge AI (Exhibit 81, Exhibit 82).

On the volume side, we expect smartphone unit shipments to decline by around 13% this year, with a significantly sharper contraction in low-end devices as higher memory prices put further pressure on the already limited profitability of this segment.

We do not anticipate a meaningful correction in memory prices in FY2027 and therefore expect weakness in the low-end segment to persist. Over the medium term, we do not expect smartphone volumes to revisit the record levels seen prior to Covid. However, we expect the premium segment to continue rising as a share of total shipment benefitting Arm (Exhibit 83).

EXHIBIT 81: **While we expect ARM Mobile TAM to decline next year due to rising memory prices, we anticipate a return to strong growth thereafter, driven by robust increases in semiconductor content per device.**

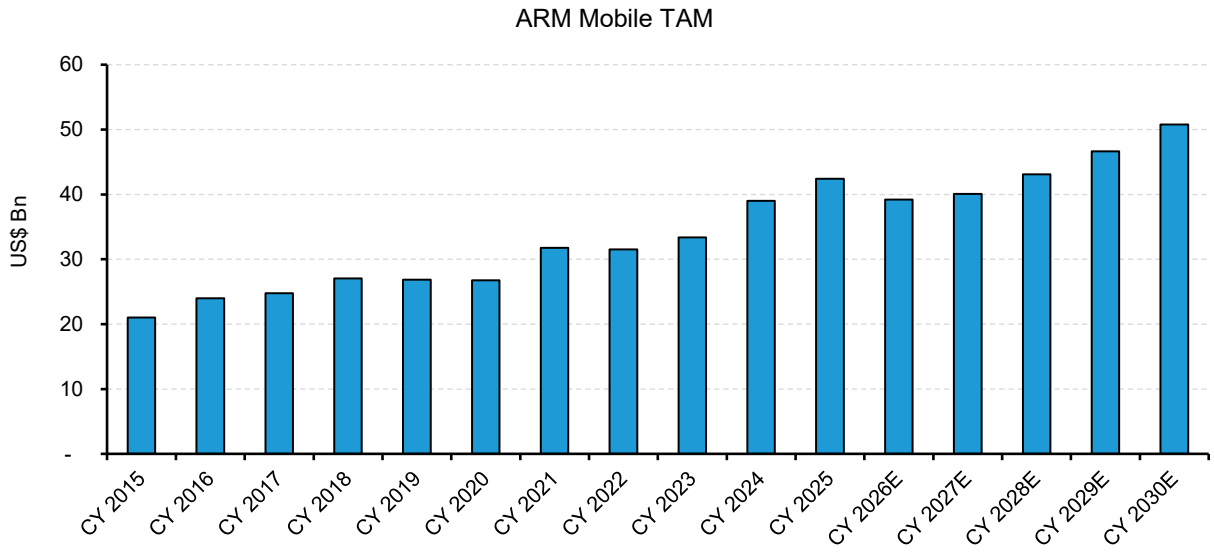


Source: Company reports, Gartner, Bernstein analysis and estimates

EXHIBIT 82: **Processor semi content in smartphones has grown significantly over the past ten years, more than doubling and driving ARM's TAM expansion. We expect semi content to continue increasing as devices require greater processing power to support edge AI.**

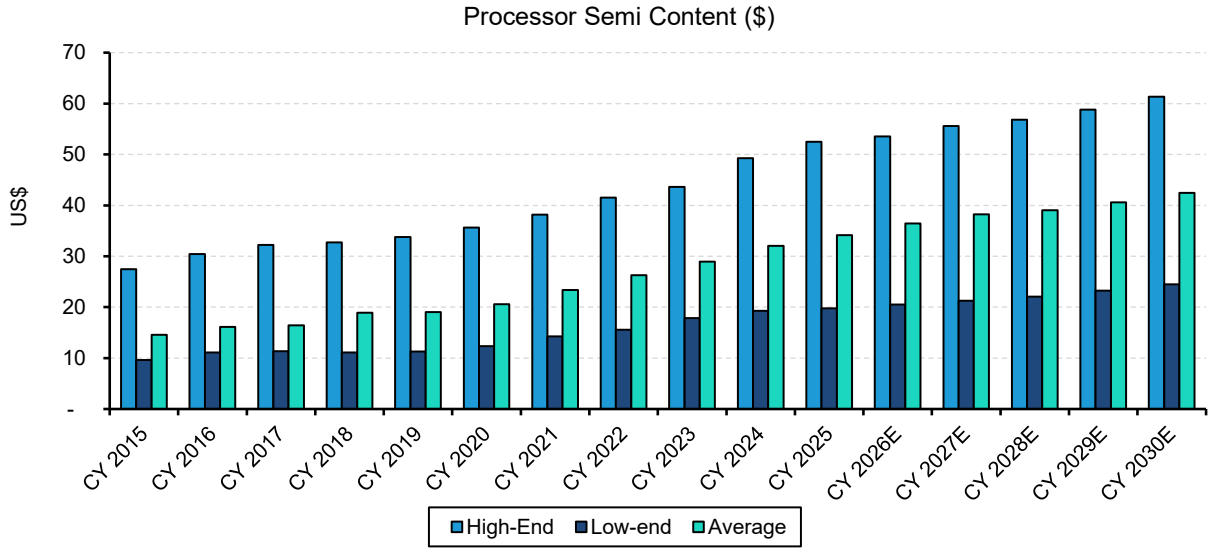


Source: Gartner, Company reports, Bernstein analysis and estimates

EXHIBIT 83: **We expect smartphone units to decline 13% this year, with a much sharper drop in low-end devices due to the significant rise in memory prices. In the years ahead, we do not expect smartphone volumes to return to the record levels reached before Covid.**

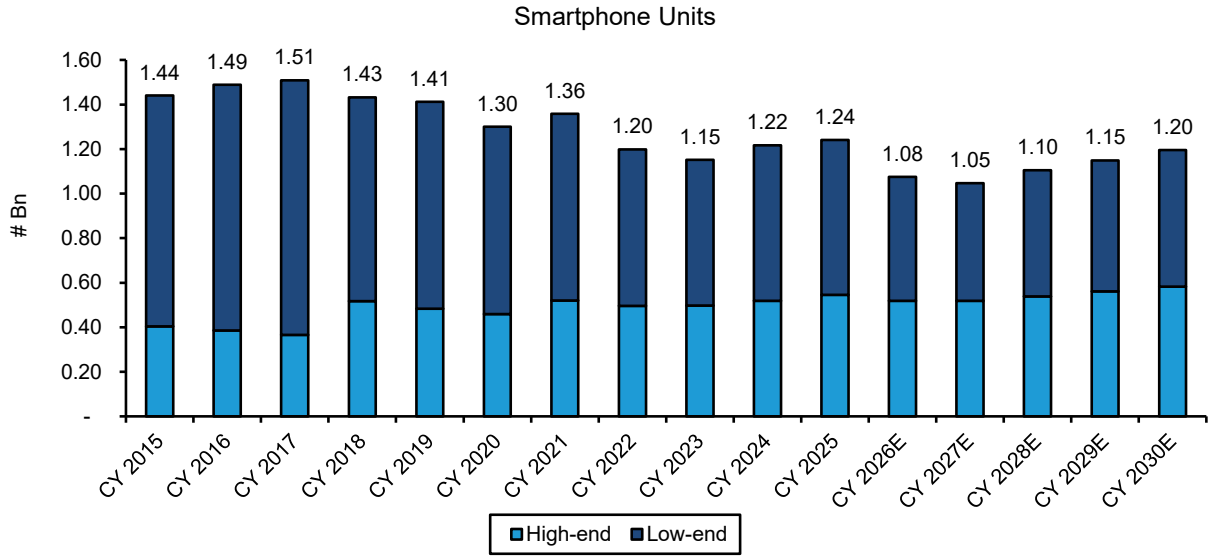


IDC, Strategy Analytics, Gartner, Bernstein analysis and estimates.

## Architecture evolution

**The architecture transitions from Armv7 to Armv8 and now Armv9 explain how Arm has kept deepening that mobile value capture.** Armv7, launched with Cortex-A8 in 2005 was the foundation of the first smartphones in the late 2000s; it expanded multimedia and signal-processing capability through NEON and helped establish modern mobile computing. Armv8, introduced in 2011, was the next major inflection, bringing 64-bit processing to smartphones and tablets and enabling a much more powerful application environment. Armv9, introduced in 2021, is the current monetization engine: it adds new security and AI-oriented capabilities, including SVE2 and SME-related extensions. Management highlighted several times that Armv9 carries a higher royalty per chip than prior architectures, double than V8. That is visible in the numbers, while smartphone royalties rose massively over the last few years largely because application processors were becoming increasingly Armv9-based, units declined.

While Arm does not disclose one single Armv8 versus Armv9 benchmark, its **first Armv9 mobile CPUs still show a clear generational benefit across the smartphone CPU stack (Exhibit 84):**

- **Cortex-X2** is Arm's first Armv9 flagship core and is designed for peak performance in the most demanding workloads such as app launch, browsing responsiveness and heavy single-thread tasks. Here, Arm highlights 30% single threaded performance improvements over the Android flagship smartphones of the time, and 2x machine learning (ML) performance versus the Armv8-based Cortex-X1. In the laptop market, Cortex-X2 delivers 40% single threaded performance improvements over 2020 mainstream devices.

- **Cortex-A710** is Arm's first Armv9 "big" core and is aimed at sustained performance within a smartphone thermal and battery envelope. Versus the Armv8-based Cortex-A78, Arm discloses 10% higher performance at the same power, up to 30% better energy efficiency, and 2x ML performance.

- **Cortex-A510** is Arm's first Armv9 "LITTLE" core and is built for lighter, background and always-on workloads. Versus the Armv8-based Cortex-A55, Arm discloses 35% higher performance, up to 20% better power efficiency, and 3x ML performance.

Arm disclosed that **v9 already accounted for 50% of mobile royalty revenue in FY25**, which we estimate corresponds to a unit penetration of below 20%. This significant gap reflects the fact that v9 carries roughly double the royalty rate of v8-around 5% versus 2.5% -although the blended rate for both architectures is lower due to significant discounts granted to strategic

partners such as Apple. CSS, which is also based on the v9 architecture, adds further upside given its royalty rate of roughly 10%, about twice that of standard v9. As a result, **we expect v9 to represent the vast majority of mobile royalty revenue by FY31, approaching around 90%**. In unit terms, we estimate ARMv9 penetration reached ~25% in FY2026, covering most of the premium segment. We expect ARMv9 to penetrate deeply into the low-end smartphone segment as well, representing close to three-quarters of total units shipped in FY31 (Exhibit 85, Exhibit 86)

EXHIBIT 84: **The first Armv9 mobile CPUs show a clear generational leap across the smartphone CPU stack, delivering substantial improvements in performance, energy efficiency, and machine-learning capability.**

| Frist CPU Armv9-based | Previous CPU Armv8-based | Performance Uplift | Energy Efficiency | ML Uplift |
|---|---|---|---|---|
| Cortex-X2  | Cortex-X1 | 30% smartphone<br>40% laptop | Not disclosed | 2X |
| Cortex-A710  | Cortex-A78 | 10% | 30% | 2X |
| Cortex-A510  | Cortex-A55 | 35% | 20% | 3X |

Source: Company reports, Bernstein analysis

EXHIBIT 85: **We estimate ARMv9 penetration reached ~25% in FY26, representing the vast majority of the premium segment. We also expect ARMv9 to penetrate deeply into the low-end smartphone segment, representing close to 3/4 of total units shipped.**

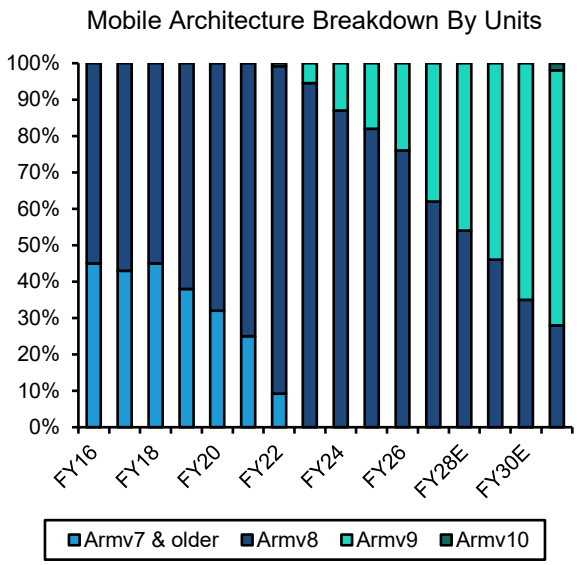


Source: Company reports, Bernstein analysis and estimates

EXHIBIT 86: **Arm disclosed that v9 accounted for 50% of mobile royalty revenue in FY25. v9 carries 2x the royalty rate of v8, with further upside from CSS, which is based on v9 and earns 2x the standalone v9 rate. We expect v9 to represent 95% of royalty revenue in FY31.**

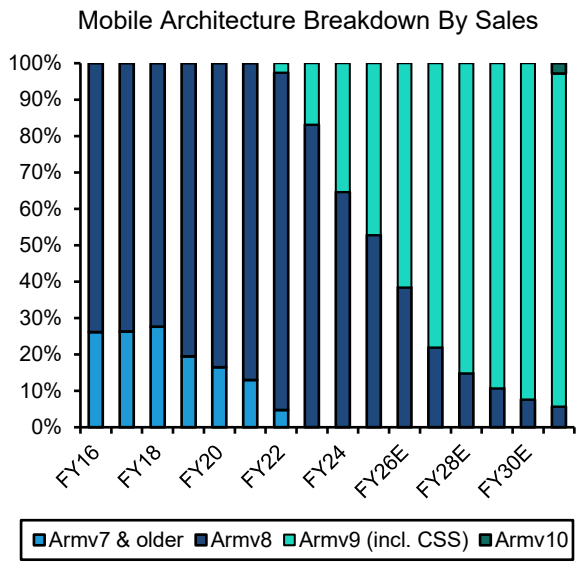


Source: Company reports, Bernstein analysis and estimates

### Arm's Strategic Shift in mobile: From IP Building Blocks to Full Compute Subsystems (CSS)

For more than three decades, Arm built its mobile business around providing individual IP components such as CPU cores, system IP, and other building blocks for partners to assemble into their own smartphone SoCs. But rising design complexity,

shorter timelines, and the growing demands of Edge AI pushed Arm to rethink this model. In early 2023, Arm introduced **Compute Subsystems (CSS)**, targeting primarily the cloud and mobile end-markets. This marked a major strategic shift from supplying isolated IP blocks to delivering **pre-integrated compute platforms** that combine the latest Armv9 CPU cluster, Immortalis/Mali GPU, CoreLink interconnect, system-level cache, memory-management blocks and production-ready physical implementations on leading-edge nodes such as 3nm. **CSS dramatically reduces engineering effort, lowers design risk, shortens time-to-market (in some cases by up to 12/18 months), and increases royalty revenue per chip. Arm has repeatedly argued that the cost savings customers achieve with CSS more than justify the substantial increase in royalties per device.**

Strategically, CSS has been a critical intermediate step in Arm's evolution from a pure IP provider to a company that also supplies silicon products such as the Arm AGI CPU. It shifted Arm from being merely a IP component supplier to more of a platform provider for flagship smartphones, and now a producer of finalized CPU products, though currently only for the cloud market. During the March 2026 Investor Day, management stated that Arm is not entering the mobile CPU business at this time, but it is something they may consider in the future.

### Evolution of Arm's client platform: from TCS to CSS to Lumex

The first platform-type product was Arm's Total Compute Solutions (TCS), effectively the ancestor of CSS, first launched in 2021 to give SoC designers a complete package of IP optimized to work together and reduce integration complexity. By 2023, **TCS23** had already become a materially more complete mobile platform. Arm says TCS23 integrated the latest Armv9.2 compute cluster with Cortex-X4, Cortex-A720, Cortex-A520 and DSU-120, together with Immortalis-G720 GPUs.

### 2024: CSS for Client was the inflection point

Arm CSS for Client, launched in 2024, was the first CSS-branded product and marked an important strategic evolution from Arm's earlier TCS approach into a formally branded compute subsystem offering for consumer devices. Building on the foundation established by TCS, Arm positioned CSS for Client as the compute platform for AI-powered experiences across smartphones and PCs, integrating the latest Armv9.2 CPU cluster, Arm Immortalis and Mali GPUs, CoreLink interconnect system IP, and production-ready physical implementations for CPUs and GPUs on leading-edge 3nm process nodes. The significance of CSS for Client was that it packaged these technologies into a more complete, scalable and silicon-ready subsystem, giving partners a faster path to production while enabling flexible and customizable SoC designs for the AI era**.** That is what made CSS revolutionary versus TCS23. TCS23 was already system-level, but CSS for Client added a much more productized, silicon-ready layer: Arm described it as the **fastest path to production silicon (Exhibit 87).**

The performance delta versus TCS23 was considerable. Arm says CSS for Client delivered **36% higher peak performance**, **33% faster application launch times**, **60% faster web browsing**, and **30% higher peak graphics performance** versus TCS23. On AI specifically, it delivered **42% faster time-to-first-token for Llama 3**, **46% faster TTFT for Phi-3**, **59% faster AI inference on Cortex-X925**, **36% faster AI inference on Immortalis-G925**, and **24% better AI camera bokeh performance.** For gaming, Arm cites **37% average performance uplift at the same power** and **30% GPU power reduction** at 120fps (Exhibit 88).

EXHIBIT 87: **In 2024, Arm launched CSS for Client, featuring the Cortex X925 extreme core, the Cortex A725 performance core, the Cortex A520 efficiency core, and the Immortalis G925 graphics engine with ray tracing support. Arm stated that the new cores provide meaningful improvements in performance.**

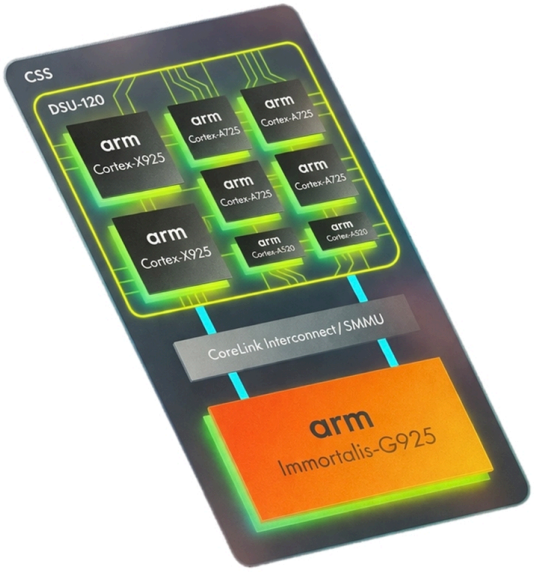


Source: Company reports, Bernstein analysis

EXHIBIT 88: **CSS for Client demonstrated significant improvements across key benchmarks compared with the TCS23 platform. In AI-specific tests, the platform delivers a 42% faster time-to-first-token with the Llama 3 LLM and a 46% with the Phi-3 LLM.**

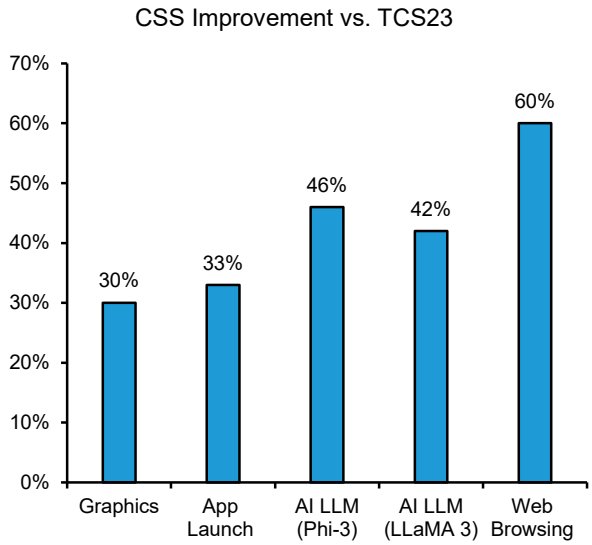


Source: Company reports, Bernstein analysis

## Lumex: the latest and most efficient CSS generation

**Arm Lumex CSS** is the latest platform, and it is the clearest expression yet of Arm's shift toward an **AI-first subsystem architecture**. Arm describes Lumex as an advanced compute subsystem that combines leading-edge CPU clusters, a new GPU, ultra-efficient interconnect, the latest physical implementations, and seamless software integration, all designed for heterogeneous AI compute at scale. Importantly, Arm says Lumex is intended to make AI pervasive not only in flagship devices but also across more cost-sensitive tiers.

Lumex is the first mobile CSS platform to adopt Arm's **new C1-branded CPU family**, built on an **Armv9.3** CPU cluster comprising C1-Ultra, C1-Pro, and C1-Premium, replacing the previous Armv9.2-based Cortex cluster of Cortex-X925, Cortex-A725, and Cortex-A520 used in the prior CSS generation. Lumex also changes the GPU branding, with the new platform featuring the **Mali G1-Ultra GPU** rather than the Immortalis. Notably, the C1-CPU cluster is also available as standalone CPU IP, rather than only being integrated with the CSS Lumex platform (Exhibit 89).

**The key change versus the prior CSS generation is the much stronger focus on on-device AI acceleration at the CPU level, led by the integration of Arm Scalable Matrix Extension 2 (SME2).** Lumex combines SME2-enabled Armv9.3 C1 CPUs with system-level optimizations and software enablement, which Arm says delivers up to 5x AI performance within the same power budget. This is important because SME2 is not simply an incremental CPU upgrade, it is the architectural feature that allows the CPU complex to handle matrix-heavy AI workloads far more efficiently, improving both throughput and latency for use cases such as speech, vision, LLM inference, assistants and real-time personalization. Arm specifically highlights **4.7x lower latency for speech-based workloads** and **2.8x faster audio generation**, underscoring that Lumex is designed for responsive, always-on AI running directly on the device.

Just as importantly, Arm's software stack appears to be aligned tightly around SME2. **KleidiAI support is integrated across major mobile OSes and AI frameworks**, including **PyTorch ExecuTorch, Google LiteRT, Alibaba MNN and Microsoft ONNX Runtime**, allowing developers to benefit from SME2 acceleration with minimal friction. In practice, this makes Lumex not only a faster AI platform, but also a more deployable one, since developers can access the performance uplift without having to rebuild their software stack from scratch.

From an efficiency standpoint, Lumex materially strengthens Arm's performance-per-watt proposition. Arm says the platform provides **up to 3x energy savings versus previous generations**, which is particularly relevant as more AI workloads move on device and begin to run continuously in the background. That combination of higher AI throughput and lower energy consumption is what makes Lumex especially well suited for battery-constrained mobile devices.

**Graphics is the second major upgrade. Mali G1-Ultra GPU** delivers **up to 2x ray tracing performance** through its next-generation graphics architecture, improving immersive gaming and advanced visual workloads. For AI workloads, the G1-Ultra enables **up to 20% faster inference performance**, enhancing responsiveness across real-time applications.

Finally, Lumex looks stronger operationally for partners. Arm highlights ready-to-use physical designs to shorten development cycles and integrated tools for profiling, debugging and optimizing AI applications, which should reduce development time, improve developer productivity, and accelerate launches.

EXHIBIT 89: **C1-CPUs, used in the Lumex CSS as well as offered as standalone CPU IP, provide the flexibility to balance peak performance, sustained efficiency, and silicon area for products ranging from high-end smartphones and PCs to emerging AI-first form factors.**

| C1-CPU Family (Armv9.3-based) | Key Benefit | Performance and efficiency gains | Ideal use cases |
|---|---|---|---|
| C1-Ultra | Flagship peak performance | +25% single-thread performance vs Cortex-X925 | LLM inference, computational photo, content creation, Gen AI |
| C1-Premium | C1-Ultra performance with greater area efficiency | 35% smaller area than C1-Ultra | Sub-flagship mobile segments, voice assistants, multitasking |
| C1-Pro | Sustained efficiency | +16% sustained performance vs Cortex-A725 | Video playback, streaming inference |
| C1-Nano | Extremely power-efficient | +26% efficiency, using less area vs Cortex-A520 | Wearables, smallest form factors |

Source: Company reports, Bernstein analysis

EXHIBIT 90: **The Lumex CSS platform, built for on-device AI, integrates advanced CPU clusters, a next-generation GPU, an efficient interconnect, and modern physical IP under a unified software stack, designed to scale across heterogeneous AI workloads.**


Source: Company reports, Bernstein analysis

## Smartphone adoption of Arm CSS: fast ramp from 2024, with Lumex moving even quicker

Arm's smartphone CSS ramp has been remarkably fast. After only starting to show up meaningfully in flagship Android phones in late 2024, CSS had already expanded to the top four Android smartphone vendors by Q3 FY2026, while Lumex was generating royalty revenue within months of its September 2025 launch. That is a strong indication that CSS is helping customers bring new smartphone platforms to market faster by cutting design complexity and shortening development cycles.

**The clearest first-wave smartphone adoption came through MediaTek's Dimensity 9400**, which powered flagship launches such as vivo X200, vivo X200 Pro, OPPO Find X8 and OPPO Find X8 Pro. Vivo's official specifications list the X200 Pro with Dimensity 9400, while OPPO's official product pages confirm both Find X8 and Find X8 Pro use the same SoC. On performance, MediaTek (Covered by Mark Li) disclosed a meaningful step-up for the Dimensity 9400 generation, delivering up to 35% faster single-core CPU performance, 28% faster multi-core CPU performance, 41% higher GPU peak performance, up to 44% better GPU power efficiency, and up to 40% higher overall power efficiency versus the prior generation. OPPO repeated those gains in its Find X8 series launch communication, underscoring that the CSS-based platform translated directly into flagship smartphone positioning around gaming, imaging and AI performance. The first CSS wave has expanded beyond MediaTek. **Samsung's Exynos 2500** is already shipping in the Galaxy Z Flip7, and confirmed it adopted CSS across its Exynos

lineup. The Exynos 2500 delivers a 15% improvement in multi-core CPU performance, a 39% boost in NPU performance, and a 28% increase in ray-traced gaming frame rates compared with the previous generation.

Lumex adoption has been even faster than first-generation CSS. Arm said in Q2 FY2026 that Lumex had already started generating royalty revenue from an early licensee only a couple of months after launch, and management explained that the speed came from a customer that was already on its second CSS generation. That is a strong indication that **once customers adopt CSS, follow-on generations can move through the smartphone cycle materially faster.**

The most visible Lumex smartphone example so far is MediaTek's Dimensity 9500. MediaTek said the chip delivers up to 32% higher single-core CPU performance, 17% higher multi-core CPU performance, 33% higher GPU peak performance, 42% better GPU power efficiency, and up to 55% lower peak power consumption for the ultra core versus the previous generation. Arm disclosed MediaTek was designing Lumex configurations into next-generation flagship smartphones from OPPO and Vivo, which lines up with the early device ramp seen in market (Exhibit 91). **In other words, Lumex moved from its September 2025 launch to royalties and shipping flagship smartphones in a timeframe that appears considerably shorter than the usual smartphone silicon adoption cycle. This strongly supports the view that CSS is becoming a critical tool for Android OEMs and SoC vendors to shorten product cycles and bring new flagship platforms to market faster.**

EXHIBIT 91: **MediaTek was an early adopter of Arm's CSS-for-Client in the Dimensity 9400, released in 2024 and based on Armv9.2. Its successor, the Dimensity 9500, released in 2025, adopts the new Arm's Lumex-CSS based on Armv9.3 and the new C1 CPU cluster. Both SoCs delivered major improvements in performance, AI capabilities and gaming.**

| Mediatek SoC | Performance | On-device AI | Gaming |
|---|---|---|---|
| **Dimensity 9400** | • 2nd Gen all big core design with **Arm Cortex-X925, Cortex-X4 Cortex-A720**<br>• Based on **Armv9.2 CPU**<br>• +35% single-threaded performance<br>• +28% sustained multicore performance<br>• Up to 40% more power-efficiency* | • 8th Gen NPU with new Agentic AI capabilities<br>• World's 1st Agentic AI engine<br>• World's 1st on-device LoRA training<br>• >50 token MLLM generation performance<br>• +100% diffusion generation performance<br>• +80% LLM prompt performance | • Arm 12-core **Immortalis-G925 GPU**<br>• 41% peak performance<br>• 40% raytracing performance<br>• 44% power saving<br>• Realistic effects from opacity micromaps<br>• MediaTek Frame Rate Converter |
| **Dimensity 9500** | • Octa-core with Arm **C1-Core Cluster**<br>• Up to 32% faster single core performance*<br>• Up to 55% improved Ultra-core peak power efficiency*<br>• Up to 37% lower CPU peak power use* | • NPU 990 with 2nd gen Gen-AI Engine<br>• Leading **Armv9.3 support with new SME2** for faster AI perception, and ML efficiency<br>• Up to 56% reduced peak power use<br>• Over 2x faster token generation speed<br>• 4K resolution text-to-image generation | • Arm **Mali-G1 Ultra** GPU with 2X raytracing units<br>• Up to 33% better GPU performance*<br>• Up to 119% better raytracing performance*<br>• 42% better power efficiency at peak performance* |

*compared to previous generation Dimensity flagship
Source: Company reports, Bernstein analysis

## CSS Penetration

Arm disclosed that they expect CSS royalties to reach 25% of total mobile royalty revenue in FY26, which we estimate corresponds to only about 8% of units, given that CSS royalty rates are roughly double those of standalone V9 and four times those of V8. **Arm guided CSS to exceed 50% of royalty revenue in FY29 and 60% in FY31, translating to around 30% unit penetration** (Exhibit 92).

Importantly, Arm also stated that **they do not expect Apple or Qualcomm to adopt CSS** anytime soon. Since these two companies account for roughly 50% of the SoC market, CSS unit penetration may ultimately peak at around 50%. Additionally, while Arm expects CSS to expand into the low-end segment, adoption there will progress more slowly because low-end designs are often re-spun into cost-reduced variants.

**We expect CSS penetration to be materially lower, reaching around 20% of units in FY31, or roughly 45% of mobile royalty revenue**. Slower adoption is driven primarily by the fact that most of the premium market uses Qualcomm or Apple processors, leaving premium faster adoption mainly to MediaTek, Samsung, and Xiaomi's high-end devices, many of which have already moved to CSS. Additionally, the sharp increase in memory prices will add cost pressure on low-end smartphone makers, making them reluctant to adopt a more expensive architecture and delaying some AI-capabilities to later generations.

Arm noted that the key variable for CSS penetration is the OEM/SoC mix across MediaTek, Samsung, Xiaomi, and others, rather than uniform market adoption. They also plan to release a new CSS generation every year, which should enable higher royalty rates, with an expected 20% revenue uplift per generation on average. Thus, beyond penetration, CSS benefits from significant per-generation monetization.

In terms of revenue, we estimated CSS to reach around 300Mn in FY26. Even with our more conservative assumptions on penetration and royalty-rate expansion, **we still expect CSS revenue to grow at a ~30% CAGR between FY26 and FY31E, reaching USD 1.1Bn by the end of the decade (Exhibit 93).**

EXHIBIT 92: **We anticipate CSS to reach ~45% of total mobile royalty revenue in FY31, which is below Arm's upper-end guidance of approximately 60% provided at the March 2026 Investor Day.**

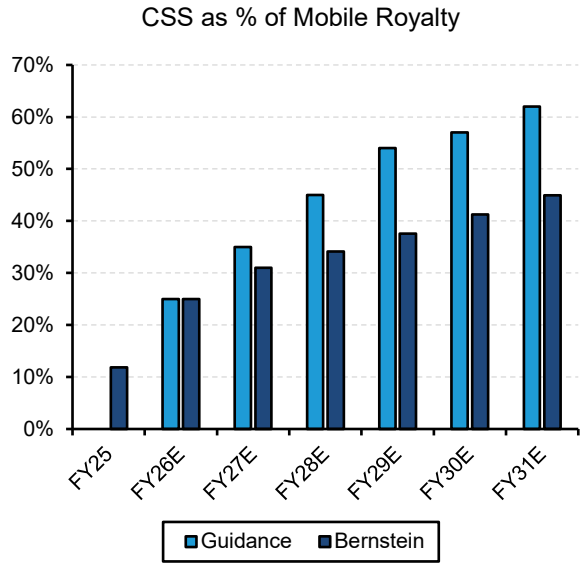


Source: Company reports, Bernstein analysis and estimates.

EXHIBIT 93: **While we expect CSS to penetrate relatively slowly in unit terms, CSS revenue will grow significantly thanks to its royalty rate, which is twice that of v9 and four times that of v8. We anticipate CSS revenue to exceed US$ 1 Bn in FY2031.**

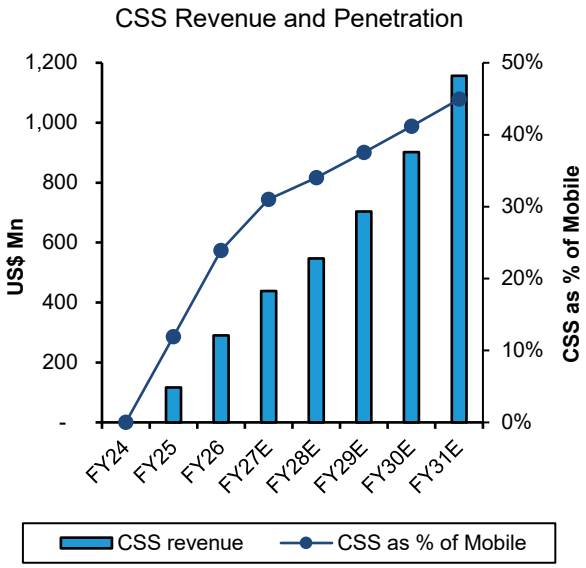


Source: Company reports, Bernstein analysis and estimates

## Mobile Royalty Rates

**Arm disclosed that the V8 royalty rate is about 2.5%, and that the rate for V9 is roughly double at 5%**. We estimate that V7 was around 1-1.5%, although it is now irrelevant because Arm has stated that it disappeared from mobile several years ago.

The actual royalty rate is meaningfully lower in practice because strategic partners receive very large discounts. In mobile, Apple pays royalty levels that are far below those applied to other customers, which pushes the blended royalty rate at least 50-100 bps below the stated figures. Apple does not sell chips, so it is not charged as a percentage of chip value. Instead, it pays a fixed dollar amount with highly preferential terms. This is also visible in the disclosed numbers: Arm provides market share, TAM and mobile royalty revenue, which makes it easy to calculate the implied royalty rate. This implied rate is significantly lower than the weighted average of the stated royalty rates.

**Arm has also disclosed that CSS royalty rates are around 10%, which again is roughly double the V9 level.** The rate increases further with Lumex CSS (V2), reaching the low teens. Arm has guided that royalty rates will continue to rise as new CSS versions are adopted. **Although our assumptions for blended royalty rates and the adoption of V9 and Lumex CSS remain conservative, we project that the average mobile royalty rate will more than double over the next five years to approximately 5%. The rising royalty rates are the primary engine powering Arm's mobile revenue growth (Exhibit 94, Exhibit 95).**

**Our Mobile IP TAM industry model can be downloaded here: (ARM Mobile IP TAM Model)**

EXHIBIT 94: **Arm disclosed that V8 royalty rates are ~2.5%, increasing to 5% with V9 and rising again to around 10% with CSS. Lumex CSS (V2) is even higher, in the low-teens. These higher royalty rates are the key growth driver behind Arm's mobile revenue expansion.**

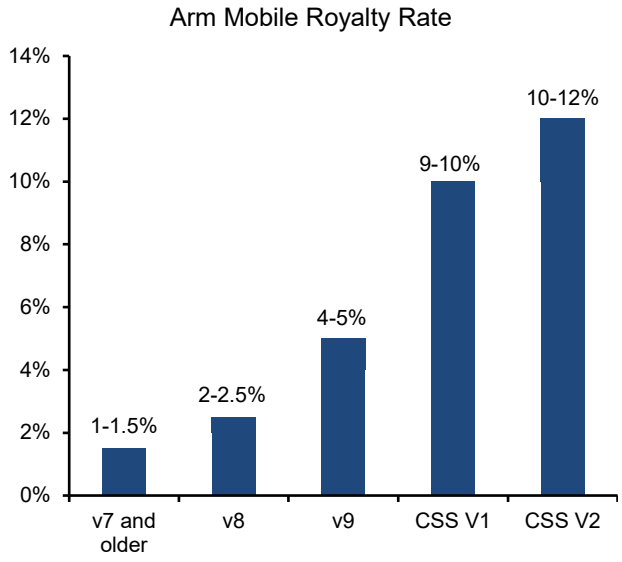


Source: Company reports, Bernstein analysis and estimates

EXHIBIT 95: **Even under our conservative assumptions, we expect the royalty rate to more than double over the next five years, reaching approximately 5%. This strong growth will be driven by increased adoption of v9 and the expanding penetration of CSS.**

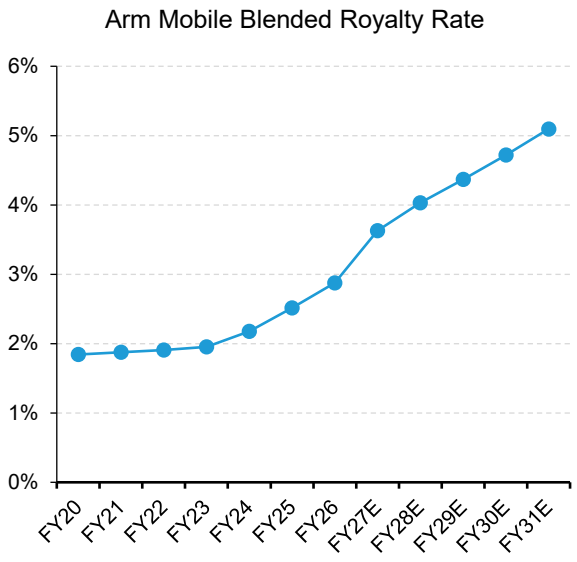


Source: Comapany reports, Bernstein analysis and estimates

### Mobile Royalty Revenue: Growing at 16% CAGR over the next five years

Based on our estimate, mobile royalty revenue reached US$1.2Bn in FY26, representing around 45% of overall royalty revenue. Arm v9, including CSS, accounts for more than 60% of this revenue. Within that, the CSS upside (the 5% additional royalty rate applied to CSS compared with v9 standalone) represents 13% of sales, or about US$120Mn, while v9 standalone contributes 47%. Armv8 makes up the remaining 40%. We estimate total CSS reached ~300Mn or 25% of royalty revenue in FY26.

Driven by improved performance, better power efficiency and enhanced AI-capabilities, we expect Armv9 to exceed 70% of shipments by the end of the decade. This translates into roughly 90% as a share of revenue, with the CSS upside contributing about 30%. We also expect Arm to introduce a new architecture, Armv10, by FY30, although adoption should remain limited in FY31, accounting for only 3% of revenue. We expect total CSS to be 45% of FY31 royalty revenue

The higher penetration of v9 and CSS, combined with their higher royalty rates-double and quadruple those of v8, respectively-will drive mobile royalty growth. **We project mobile royalty revenue to reach US$2.6Bn in FY31E, implying a FY26-31E CAGR of 16%, despite the TAM growing only 4% (Exhibit 96).**

EXHIBIT 96: **We estimate Arm's mobile royalty revenue will rise from US$1.2Bn to 2.6Bn, at 16% CAGR. Growth will be driven primarily by wider adoption of V9 and by the incremental uplift from CSS, which adds roughly a 5% royalty increase compared with V9 standalone.**

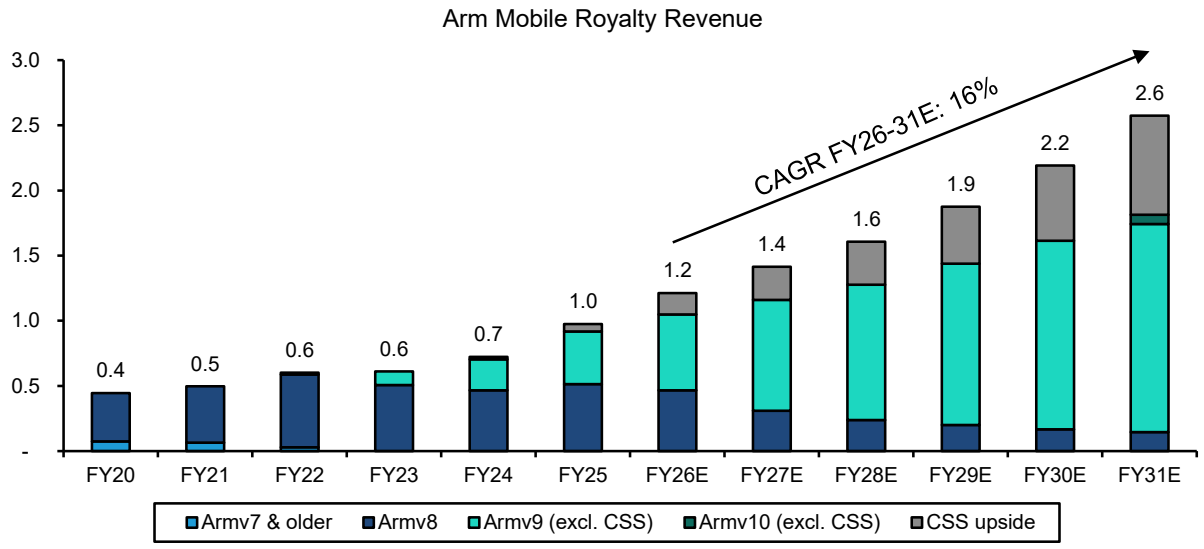


Source: Company reports, Bernstein analysis and estimates

## PHYSICAL AI: CSS TO BRING AI INTO VEHICLES

Physical AI refers to intelligence embedded directly into machines that must sense their environments, interpret continuous streams of data, and act reliably and safely in real-time. It is becoming essential because AI is moving from cloud-based inference into the physical world where decisions immediately affect motion, safety and long-term operational reliability.

**Automotive is the core of Physical AI because vehicles now depend on on-device intelligence for ADAS, automated driving, intelligent cockpit systems and continuous software evolution**. AI-Defined Vehicles extend software-defined architectures by embedding AI across perception, prediction and user interaction, enabling cars to better understand their surroundings, anticipate events and offer more intuitive in-cabin experiences. This shift strengthens real-time decision-making by allowing systems to learn continuously from sensor and environmental data.

Meeting these demands requires high-performance, safety-capable compute able to run real-time perception, planning and sensor fusion with strict determinism. As architectures consolidate hundreds of electronic functions into zonal or centralized controllers, compute loads grow substantially and must operate with guaranteed latency and isolation. Arm's AE technologies respond to these needs through predictable real-time execution, safety readiness and efficient scaling across domains. They support ADAS and autonomy with low-latency inference and robust mixed-criticality separation, while in infotainment and cockpit systems they deliver rich interfaces, multi-OS coexistence and AI-enhanced user experiences within a secure, consolidated platform.

Robotics provides another avenue for Physical AI, though much smaller in scale for Arm's business. Here, Arm technologies enable robots to adapt to human environments, navigate autonomously and perform tasks through on-device perception and efficient real-time compute.

Over the past decade, **Arm's automotive presence has expanded rapidly, significantly outperforming other end-markets as OEMs accelerated adoption of the Arm architecture. Arm's automotive market share increased from 33% in FY21 to 44% in FY25.** Modern vehicles have become complex software-driven systems requiring far more compute than in the past, and Arm's long-standing strength in energy efficient CPU designs positioned it to become a foundational supplier. Arm's Physical AI product portfolio is built on its AE technologies which combine IP and software designed specifically for real-time, safety-critical autonomous machines. **Neoverse V3AE** delivers Armv9.2 server class performance tuned for automotive AI workloads. Armv9 A class processors such as Cortex A720AE and Cortex A520AE scale compute for software defined and AI defined vehicle domains including intelligent cockpit and infotainment. The **Cortex R family**, including Cortex R52, R52 plus

and R82AE, provides deterministic real-time control that is essential for safety islands, powertrain control and timing critical automotive functions. I**n 2025, Arm launched Zena CSS**, a standardized and safety-capable compute subsystem designed to accelerate automotive silicon development.

EXHIBIT 97: **Arm's Neoverse V3AE is designed for maximum performance in automotive central compute and ML. It supports large, high-speed memory and fast die-to-die links, delivering Arm's highest Neoverse performance. This makes it well-suited for ADAS.**

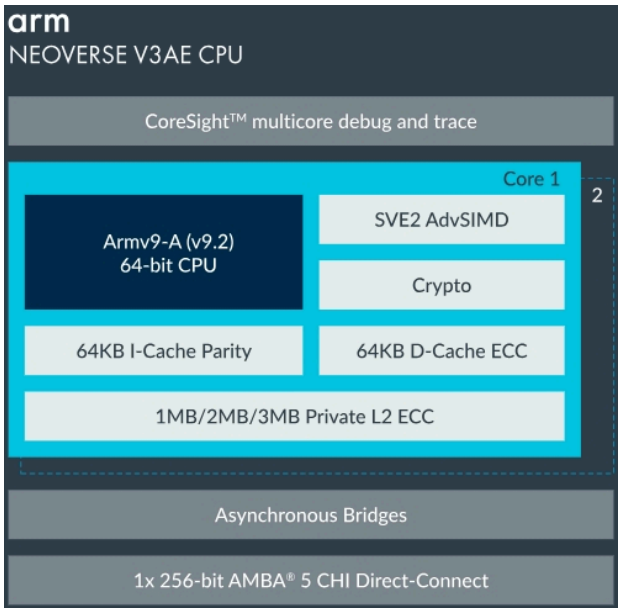


Source: Company reports, Bernstein Analysis

EXHIBIT 98: **Arm Cortex-A720AE is the highest-performance Cortex-A processor designed to meet the complex and demanding safety needs of the next generation of SDVs.**

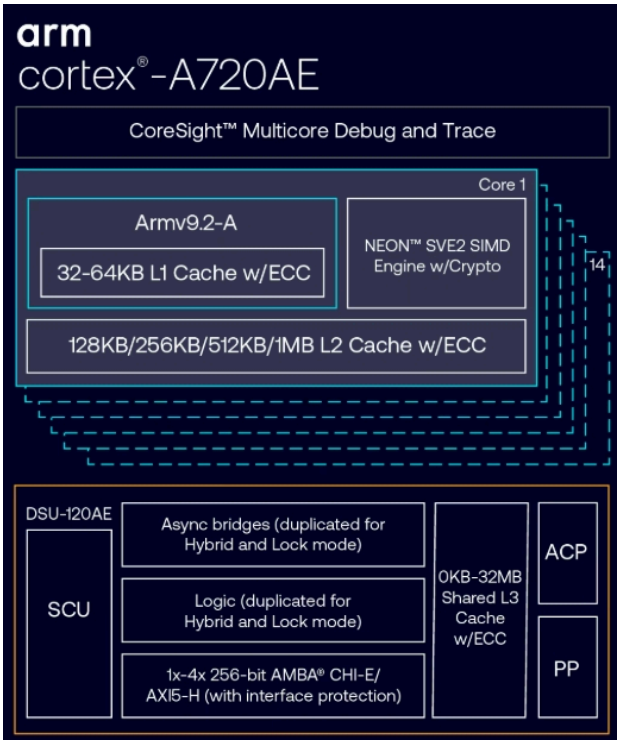


Source: Company reports, Bernstein analysis

## Physical AI TAM

We estimate Arm's Physical AI TAM in FY26 was almost entirely automotive, with robotics contributing only about 2%. The automotive TAM, driven mainly by processors and MCU/MPU, expanded from US$ 11Bn in FY20 to US$ 25 Bn in FY26, growing at a 15% CAGR. This increase occurred despite flattish global vehicle volumes, meaning growth was propelled by rising semiconductor content per vehicle. The uplift in silicon value was driven by the adoption of ADAS (L0-L2), richer infotainment and digital cockpit systems, and the expansion of safety and sensing workloads. More advanced ADAS L3+ capabilities and high-performance compute (HPC) platforms have only recently begun penetrating the market.

**Arm expects its Physical AI TAM to continue growing at roughly 15% CAGR, reaching around US$ 50 Bn by the end of the decade.** This expansion will be driven by AI-Defined Vehicles, where AI becomes embedded across perception, prediction, driver monitoring and user interaction, enabling continuous learning and more personalised in-cabin experiences. Increasing consolidation toward zonal and centralised compute architectures will further accelerate demand for deterministic, safety-capable platforms. Arm disclosed that they expect their long-term Physical AI TAM, beyond FY31, to reach around US$200Bn, driven primarily by growth in the robotics market and the development of fully autonomous vehicles.

**We are slightly more conservative than guidance and estimate the Physical AI TAM to reach US$ 42Bn in FY31, growing at 11% CAGR**. Within this, automotive represents US$ 39Bn and robotics US$ 3Bn (~6% of Physical AI). Although we assume a strong increase in processor and MCU dollar content per vehicle, rising at 9% CAGR due to deeper ADAS adoption and greater penetration of high performance computing, we expect auto unit volumes to increase by only 2% CAGR over the next five years (Exhibit 99).

EXHIBIT 99: **Arm's Physical AI TAM increased from US$11Bn in FY20 to US$25Bn in FY26, reflecting a 15% CAGR. Arm expects the TAM to continue growing at a similar pace, reaching US$50Bn in FY31E. We are slightly more cautious and assume lower growth in auto.**

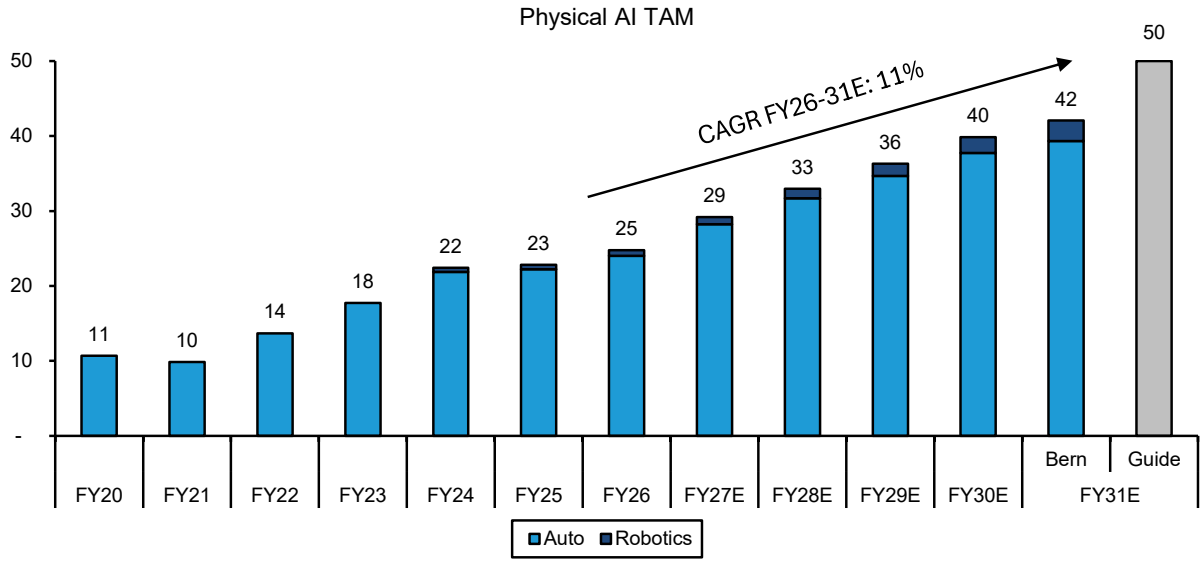


Source: Company reports, Gartner, Bernstein analysis and estimates

### Arm Zena CSS: Unlocking Faster Growth for AI-Defined Vehicles

**Zena CSS is Arm's pre-integrated, pre-validated compute subsystem for the next generation of AI-defined vehicles, designed to scale from digital cockpit systems to advanced driver-assistance systems (ADAS)** as the automotive industry shifts toward centralized compute architectures. Arm positions it as a foundation for "physical AI," where vehicles must sense, reason, and act in real time under tight power, safety, security, and reliability constraints.

What Zena CSS does is reduce the complexity of building automotive SoCs by packaging together a compute foundation with high performance, functional safety, cloud-native readiness, and built-in security. Rather than forcing automakers, Tier 1s, and silicon providers to assemble and validate every core platform element themselves, Zena CSS gives them a ready starting point so they can focus more of their effort on differentiated features and user experiences.

In performance terms, Arm describes the improvement less as a single benchmark number and more as a system-level gain: **Zena CSS is built to support AI-driven workloads with deterministic response, strong energy efficiency, functional safety, and secure operation over long automotive lifecycles.** In practice, that means it is intended to support demanding use cases such as real-time perception for ADAS, intelligent voice interaction, adaptive infotainment, OTA updates, and safety-critical synchronization across distributed software workloads. Its security architecture is also part of that performance story. Arm says **Zena CSS Security provides layered, defense-in-depth protection, rooted in the Runtime Security Engine and Armv9 defensive security architecture**, with secure boot and OTA protection. This is important because in AI-defined vehicles, secure and reliable operation is inseparable from overall platform performance.

The biggest strategic value of Zena CSS is how it shortens the product cycle. Arm's approach combines a pre-integrated hardware platform with SOAFEE-based cloud-native development and virtual prototyping from partners such as AWS, Cadence, Siemens, and Synopsys. Because those virtual platforms model Zena CSS before physical hardware is available, software teams can validate workloads, integrate middleware, and refine system behavior much earlier in the lifecycle. That is the reason **Arm says Zena CSS can reduce time to silicon by up to 12 months and accelerate software development by up to two years: hardware and software no longer need to wait on each other, and integration problems can be found far earlier (Exhibit 100).**

EXHIBIT 100: **Zena CSS helps automakers launch new vehicle models at least one year faster than traditional programs by accelerating both software and silicon development, enabling quicker and more efficient delivery of AI features. The pre-verified, safety-capable platform also saves ~20% of engineering resources, reducing development cost and complexity.**

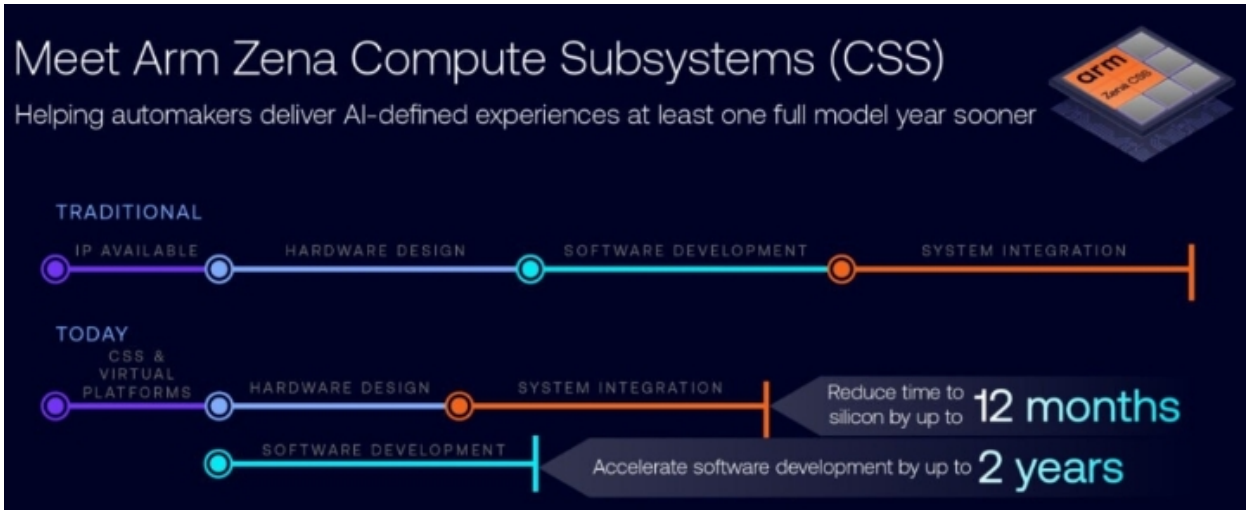


Source: Company reports, Bernstein analysis

Arm signed its first automotive CSS agreement with a leading electric-vehicle OEM as early as May 2025, although royalties will take time to flow given the long licensing-to-production cycle in automotive. During the March 2026 investor day, **Arm stated that CSS should represent only 2% of revenue in FY26 and that adoption in Physical AI will ramp more slowly than in Cloud and Edge. However, they still expect uptake to accelerate meaningfully, reaching 40% by FY31E (Exhibit 101).**

Arm also highlighted that ADAS royalty rates are materially higher than those in traditional automotive. Traditional automotive primarily relies on older Arm technology for lower-value functions such as powertrain, steering control, airbags, and safety systems. ADAS, by contrast, uses much more advanced platforms, mostly v8 transitioning to v9, with some CSS licensees as well. Chip pricing is also higher due to additional testing and safety requirements, along with dedicated automotive CPU development and more extensive safety documentation, all of which support higher royalty per chip (Exhibit 102).

EXHIBIT 101: **While Arm has stated that CSS will penetrate more slowly in Physical AI than in Cloud and Edge, it still expects adoption to ramp up quickly, rising from 2% of revenue in FY26 to 40% in FY31E.**

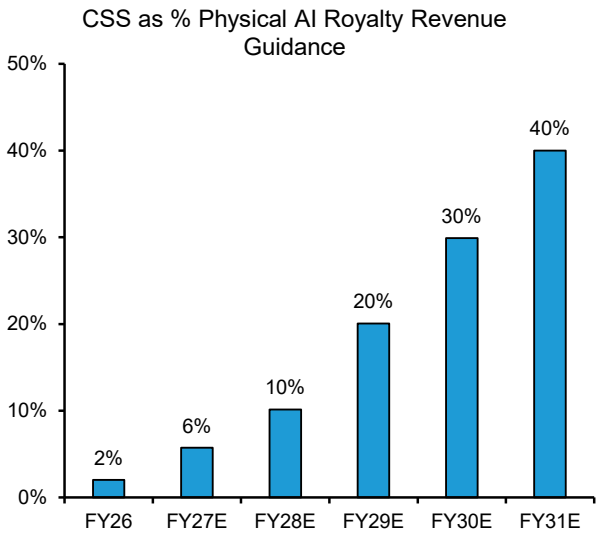


Source: Company reports, Bernstein analysis

EXHIBIT 102: **Arm Zena CSS Security provides a secure foundation for AI-driven compute, with layered protection built on Armv9's defensive security architecture.**

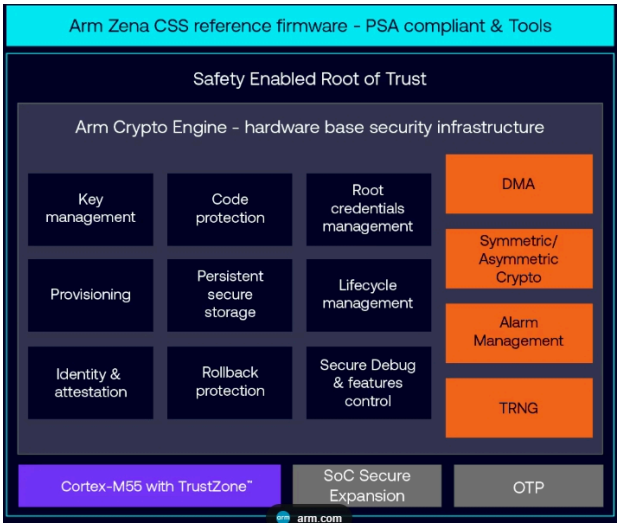


Source: Company reports, Bernstein analysis

## Physical AI Royalty Revenue: Fast Growth with Strong Visibility

At their investor day last month (March 2026), **Arm disclosed that they expect Physical AI revenue to grow at a 25% CAGR over the next five years, increasing by more than 3x**. This growth is expected to be driven by higher CSS penetration and a sharp increase in semiconductor content in autos. Arm believes this will be fueled by autonomous and ADAS vehicles continuing to adopt increasingly sophisticated silicon for both the digital cockpit and driver assistance systems, and they anticipate continued market share gains in this sector. Their confidence in this strong growth outlook is very high, supported by the long lead times and extended product cycles typical of the automotive industry. Notably, an impressive 95% of their royalties are already under contract through FYE31. Robotics is expected to remain a relatively small contributor over the next five years but is anticipated to play a critical role over the longer term (Exhibit 103).

We also expect Physical AI revenue to continue growing at a faster pace than the overall royalty business. Specifically, **we forecast Physical AI revenue to grow at a 23% CAGR over the next five years, reaching above US$ 520Mn in FY31, up from US$ 180Mn in FY26. Our forecast is slightly below Arm's 25% guidance, reflecting more conservative assumptions around TAM expansion. We also expect Arm's market share in autos to increase from 44% in FY25 to 55% in FY31E (Exhibit 104).**

**Our Automotive TAM industry model can be downloaded here: (ARM Automotive TAM Model)**

EXHIBIT 103: **Arm has very high visibility in its automotive royalty business, with 100% of royalty revenue over the next four years based on contracted rates and 95% in FY31E, supported by the sector's long product cycles.**

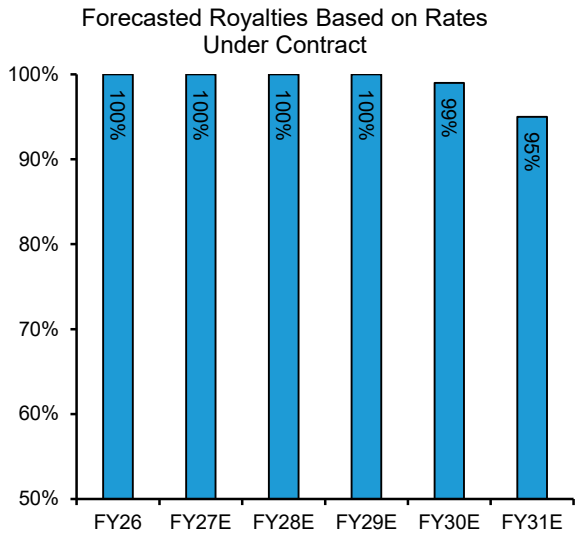


Source: Company reports, Bernstein analysis

EXHIBIT 104: **We forecast Physical AI revenue to grow at a 23% CAGR over the next five years, reaching over US $520Mn in FY31 from US$180Mn in FY26. This is slightly below Arm's 25% guidance due to more conservative TAM assumptions.**

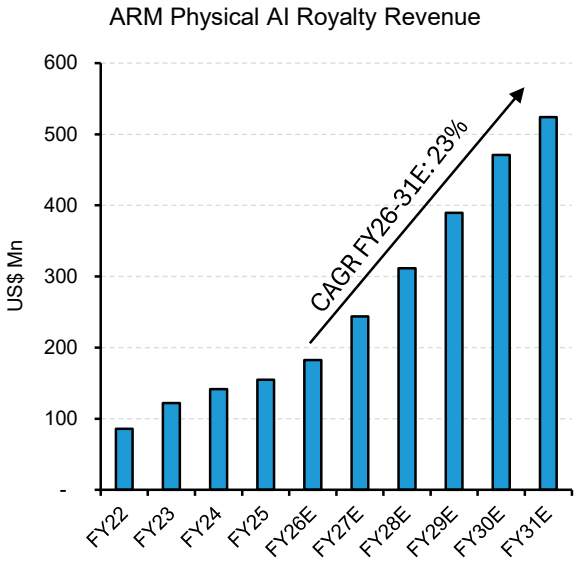


Source: Company reports, Bernstein analysis and estimates

## ARM BUSINESS MODEL

To fully understand the Arm story, it is worth going one layer deeper into the business model and how the company generates revenue. Arm is not a traditional chip company at its core. It monetizes its compute platform at different points in the customer value chain, starting with IP access, then through recurring royalties on shipped chips, and now increasingly through higher-value subsystems and, more recently, direct silicon.

- **License revenue:** The upfront monetization of Arm's technology. Customers pay for access to Arm IP, software, tools, and design rights so they can develop their own chips and platforms around the Arm architecture. This is effectively the entry point into the model, because once a customer signs a license, Arm gets paid before any chip is shipped, and that design can later become a royalty-bearing product if it reaches production. Because the timing of large agreements can make reported license revenue volatile from quarter to quarter, Arm also discloses **ACV (Annualized Contract Value)** and **RPO (Remaining Performance Obligations)** to help investors track the underlying commercial momentum. ACV is Arm's measure of normalized license-and-other revenue. RPO represents contracted revenue that is still unrecognized or yet to be invoiced, so it is effectively a view of future licensing backlog (Exhibit 105, Exhibit 106).

- **Royalty revenue:** Annuity stream of the model. Once a customer ships a chip containing Arm technology, Arm earns a royalty on each unit shipped, so revenue scales with customer volumes, end-market growth, and, importantly, the amount of Arm content embedded in each chip. This is the economic core of the business because it is recurring, highly scalable, and can last for many years once a design win is secured. The main royalty growth drivers today are the transition to **Armv9**, and faster adoption of **CSS**. In simple terms, Arm is not only trying to get into more chips, it is trying to earn more per chip.

- **CSS:** Compute Subsystem is Arm's way of moving up the value stack within the traditional IP model. Instead of licensing only individual IP blocks such as CPUs, GPUs, or system IP, Arm offers a **pre-integrated and pre-validated subsystem** that gives customers a much better starting point for chip design. That reduces engineering effort, lowers execution risk, shortens time to market, and improves end-product performance. CSS is attractive financially because it improves monetization in both revenue streams. It can support larger and stickier license deals upfront, and it also increases downstream royalty capture because a CSS-based chip contains more Arm technology and commands a higher royalty rate

than a chip built from narrower IP content alone.

- **AGI CPU:** Next step in the Arm business evolution. Arm disclosed in March 2026 that it is extending beyond IP and CSS into Arm-designed production silicon, starting with a data-center CPU built for agentic AI workloads. Strategically, this adds a new monetization layer, because Arm can now earn not only from licenses and royalties on partner chips, but also from **direct chip revenue** by selling its own processor into AI infrastructure. For an investor, the key point is that the AGI CPU does not replace the traditional model. It broadens it. Customers can still license Arm IP, adopt CSS, or buy Arm-designed silicon depending on their needs. That gives Arm participation across a larger part of the value chain, from architecture, to subsystems, to full chips, and should allow the company to capture more of the economics of AI infrastructure as CPU orchestration becomes more important in the data center.

EXHIBIT 105: **ACV reached US$1.660bn in Q4 FY26, highlighting sustained licensing momentum and providing a cleaner view of normalized contract activity levels.**

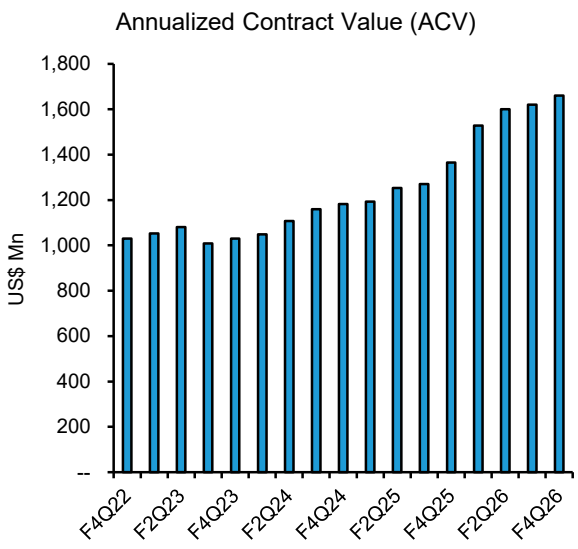


Source: Company reports, Bernstein analysis

EXHIBIT 106: **RPO ended Q4 FY26 at US$2.07 Bn, underscoring strong contracted backlog and giving visibility into future licensing revenue recognition timing patterns.**

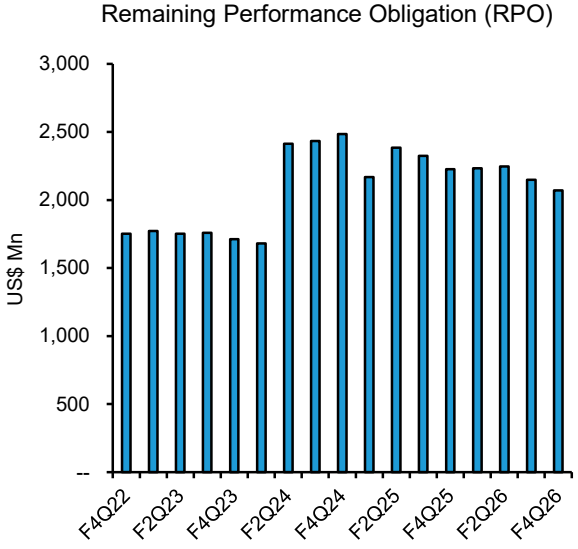


Source: Company reports, Bernstein analysis

## ARM FINANCIALS

**We forecast Arm's topline to accelerate over the next five years, growing at 40% CAGR from just under US$5 Bn in FY26 to slightly below US$26 Bn in FY31**. This growth is driven primarily by the new AGI CPU business, which we expect to reach US$15 Bn in FY31, in line with guidance. Arm views this guidance as conservative, as adoption is expected to extend well beyond the two announced clients, OpenAI and Meta, supported by significantly higher performance per watt versus x86. We anticipate revenue growth to accelerate toward the end of the decade, with CPU revenue more than doubling from FY29 to FY30 and tripling from FY30 to FY31.

**Our royalty revenue forecast is in line with guidance, rising from US$2.7Bn to US$6.7Bn at ~20% CAGR in FY26-31E.** Similarly, we expect licensing revenue to grow at 12% CAGR from US$2.3Bn in FY26 to ~US$4.2Bn in FY31, slightly above the company's outlook. Consistent with Arm's commentary, we anticipate growth to moderate toward the end of the decade, from high-teens over the next two years to high single-digits thereafter (Exhibit 107).

EXHIBIT 107: **We expect Arm's revenue to grow at an impressive ~40% CAGR over the next 5 years, reaching US $26Bn in FY31. This growth is primarily driven by the expansion of the new AGI CPU business, which we forecast will reach US$15Bn in FY31, in line with guidance.**

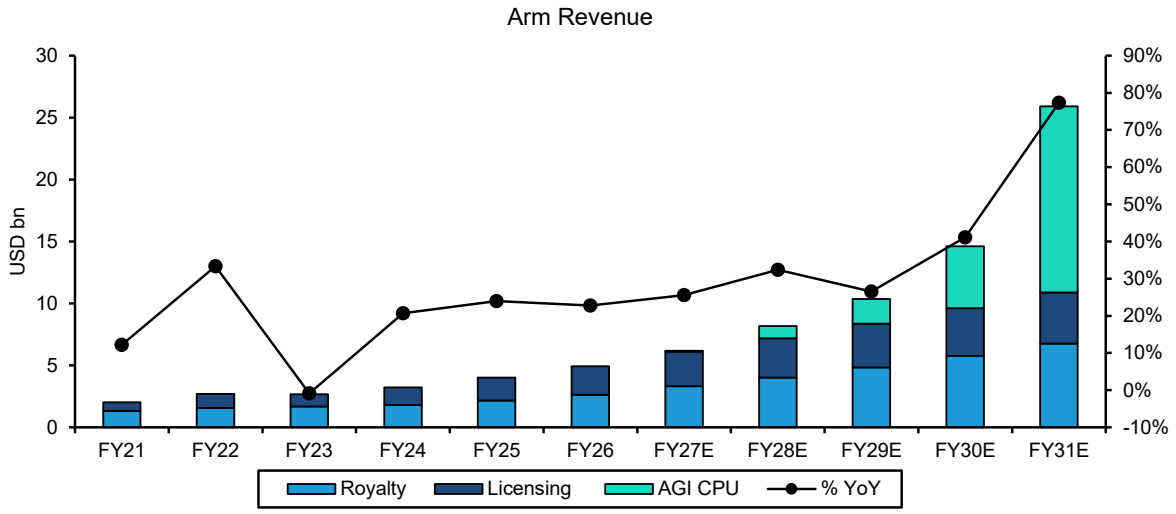


Source: Company reports, Bernstein analysis and estimates

### Royalty Revenue

Looking in more detail at royalty revenue, we expect sales to grow at a 20% CAGR over the next five years, primarily driven by Agentic AI, which we believe will significantly increase CPU demand and support growth in both IP/CSS revenue and AGI CPU revenue. **We expect Cloud AI royalty revenue to grow at a 40% CAGR from FY26 to FY31E, rising from US$0.38Bn in FY26 to US$2.06Bn in FY31E**. Within Cloud AI, we expect Cloud Compute to be the main growth driver, accelerating at a 46% CAGR, with Networking and Other Infra growing at 18% and 9% CAGR, respectively. We expect Cloud AI to reach more than ~30% of overall royalty revenue in FY31E, in line with with Arm guidance, up from ~14% in FY26 (Exhibit 108, Exhibit 109, Exhibit 110, Exhibit 111).

**While we expect Edge AI to grow more slowly at 14% CAGR over the next 5 years, it will remain the largest contributor at ~60% of royalty revenue in FY31**, down from ~79% in FY26. We expect this growth to be driven primarily by Consumer and Mobile, growing at 20% and 16% CAGR, respectively. We expect IoT and Other Mobile to grow only 7% and 3% CAGR, respectively. Finally, we expect Physical AI to slightly outperform the overall royalty business, growing at 23% CAGR. This is just below company guidance of ~25%, reflecting our more conservative TAM assumptions. As a share of royalty revenue, Physical AI should remain relatively stable over the next 5 years.

Notably, **Arm highlighted that its royalty revenue guidance is based on high confidence. More than 70% of forecasted royalties are tied to rates already under contract, giving them strong visibility.**

EXHIBIT 108: **Edge AI, which includes Mobile, IoT, and Consumer, accounts for ~79% of Arm's royalty revenue, while Cloud AI and Physical AI represent ~14% and 7%, respectively.**

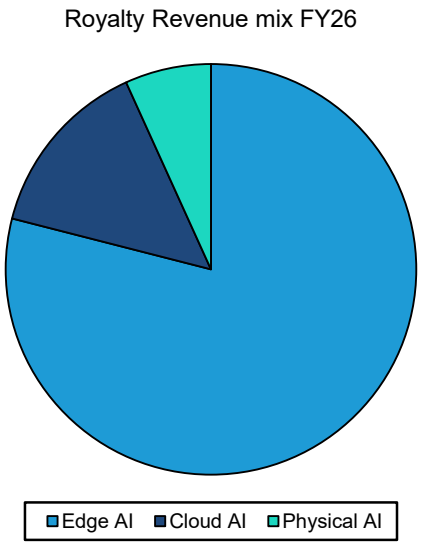


Source: Company reports, Bernstein analysis and estimates

EXHIBIT 109: **We expect Cloud AI to reach above 30% of overall royalty revenue. Edge AI will remain the largest contributor at ~60%, with the remainder coming from Physical AI.**

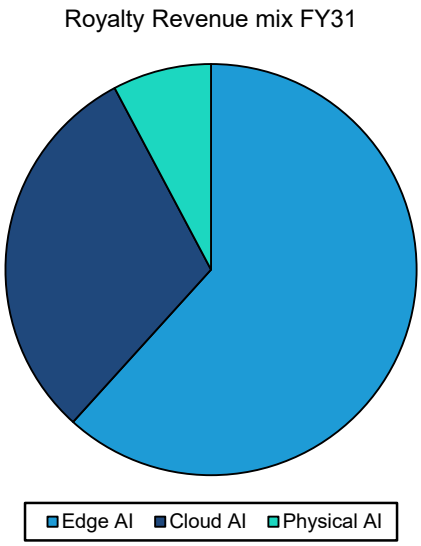


Source: Company reports, Bernstein analysis and estimates

EXHIBIT 110: **We expect Arm's royalty revenue to grow at 20% CAGR in FY26-31E, driven mainly by Cloud AI, which we expect to grow at 40% CAGR. We also expect Edge AI to grow at 14% CAGR and Physical AI to grow at 23% CAGR.**

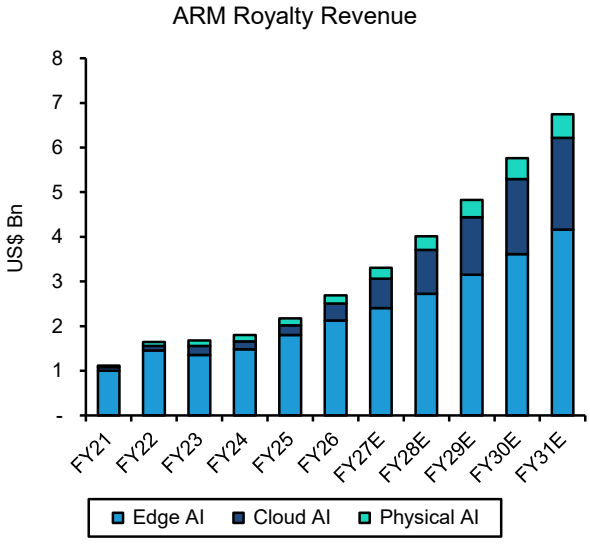


Source: Company reports, Bernstein analysis and estimates

EXHIBIT 111: **However, if we include AGI CPU revenue of US$15Bn, Cloud AI accounts for the vast majority of revenue at ~79%, with Edge AI at ~19% and Physical AI at only ~2%.**

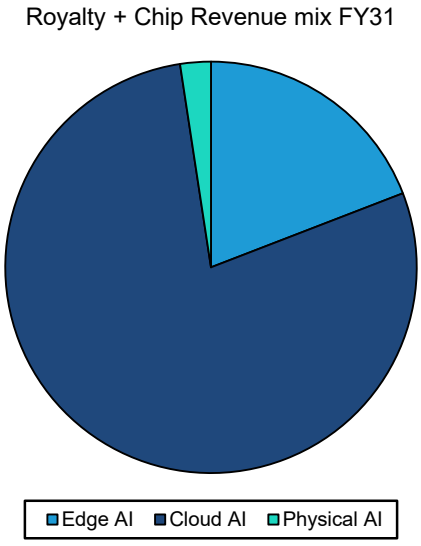


Source: Company reports, Bernstein analysis and estimates

## Profitability and EPS

Arm non-GAAP gross margin reached a record level of above 98% in FY26, as the IP/CSS business carries negligible COGS. However, as AGI CPU revenue ramps, GM will inevitably decline because the chip business has materially higher COGS. Arm

guided that GM for the AGI CPU segment should begin in the high-30s to low-40s range during the first one to two years, move toward 50% around year three, and reach at least 50% by FY31, if not sooner. Based on this, **we forecast GM to trend lower over the next five years, reaching ~72% in FY31, with a more pronounced decline toward the end of the decade as the AGI CPU business gathers momentum (Exhibit 112, Exhibit 113).**

Turning to non-GAAP operating margin, although the AGI CPU ramp will clearly pressure profitability, we expect this to be more than offset by margin improvement in the IP/CSS segment. **We expect AGI CPU operating margin to reach 35% and IP/CSS at ~65% in FY31. This results in an overall non-GAAP operating margin of 48% in FY31**. Arm also guided that non-GAAP operating cost growth will slow from a 26% CAGR over the past three years to a mid-teens CAGR over the next five years. We forecast non-GAAP opex to grow at a 18% CAGR from FY26 to FY31E. GAAP operating margin will be materially lower at 41%, as we expect SBC expense to continue rising, consistent with Arm's own expectations (Exhibit 114).

On the bottom line, we expect Arm's non-GAAP EPS to grow at a 41% CAGR over the next five years, with stronger acceleration toward the end of the decade. This growth should be driven entirely by topline expansion, as we expect OpM to remain broadly flat over the next five years. **We forecast EPS reaching US$9.83 in FY31, ahead of guidance at US$9 (Exhibit 116).**

EXHIBIT 112: **We anticipate Arm's GM to decline to 72% by FY31 as the AGI CPU business ramps up. While we expect the chip business to also weigh on operating margin, we forecast that this impact will be offset by improving profitability in the IP business.**

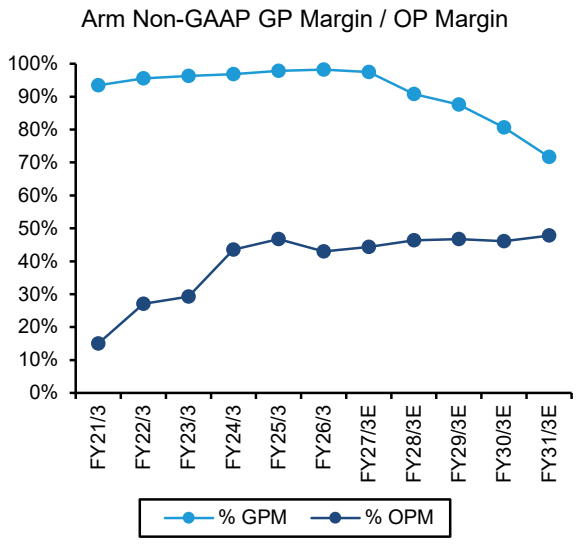


Source: Company reports, Bernstein analysis and estimates

EXHIBIT 113: **While we expect IP/CSS GM to remain flat at ~98%, we forecast Silicon GM to increase from ~30% this year to slightly above 50% by FY31.**

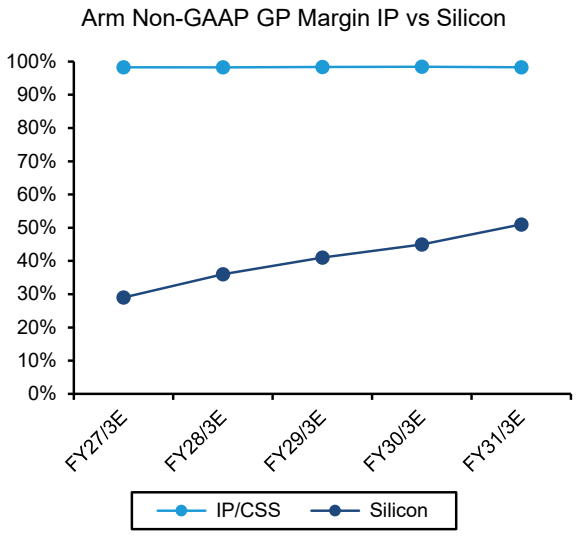


Source: Company reports, Bernstein analysis and estimates

EXHIBIT 114: **Both IP/CSS and Silicon OpM are expected to increase, but as the product mix shifts more toward Silicon by the end of the decade, total OpM is expected to increase only slightly.**

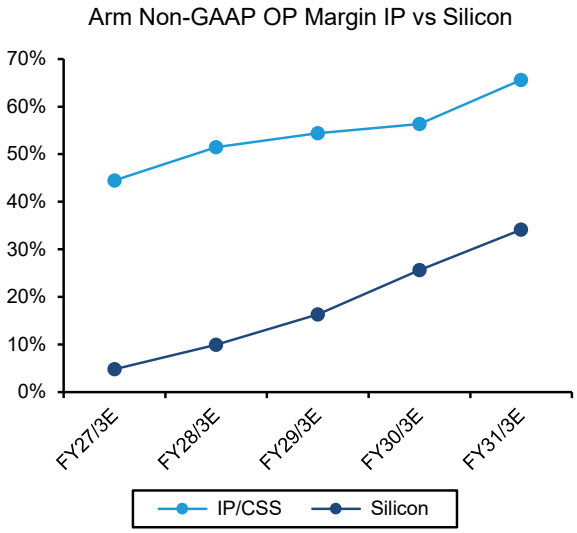


Source: Company reports, Bernstein analysis and estimates

EXHIBIT 115: **We anticipate Arm's Non-GAAP EPS to grow at a 41% CAGR, driven mostly by topline expansion, as we expect OpM to increase only slightly over the next 5 years. We forecast EPS reaching US$9.8 in FY31, above guidance of US$9.**

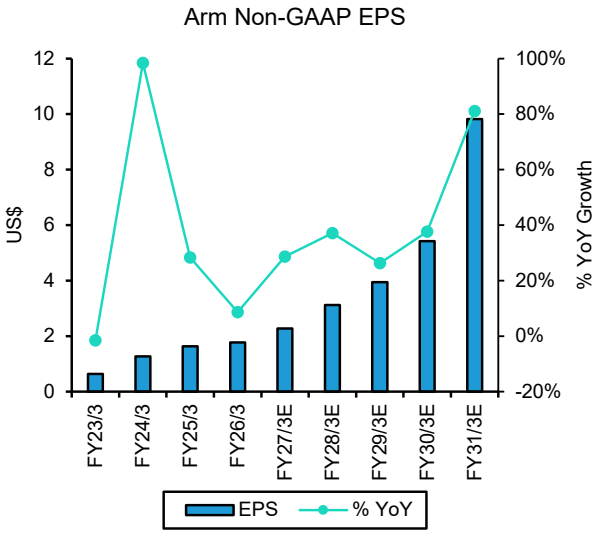


Source: Company reports, Bernstein analysis and estimates

At the topline level, our forecasts are on average mid-single digit above consensus over the next five years, with **our estimates ahead of street across all three segments: AGI CPU, Licensing, and Royalty.** In particular, we are 12% above consensus on royalty and 2% above on licensing. While no explicit consensus estimates exist yet for the new AGI CPU business, the implied Consensus figure (total revenue minus licensing and royalty) suggests we are slightly above there as well.

Similarly, on margins, we are slightly above consensus on operating margin, as we expect the initial AGI CPU revenue to carry relatively modest profitability, reaching ~35% operating margin in FY31. **This results in our Non-GAAP EPS accelerating sharply toward the end of the decade, reaching US$9.83 in FY31E, compared with Consensus at US$9.43 (4% above) and guidance at US$9** (Exhibit 116).

Our financial model can be downloaded here: Arm (Arm.US). Our ARM Industry models can be downloaded here: ARM Server CPU Industry Model, ARM Automotive TAM Model, ARM Mobile IP TAM Model

EXHIBIT 116: **Our topline estimates are above street, particularly toward the end of the decade. We are also above Arm FY31 guidance of US$25Bn. We estimate EPS will reach US$9.8 in FY31, growing at ~40% CAGR.**

| Arm Comparison Table | FY2027E | FY2028E | FY2029E | FY2030E | FY2031E |
|---|---|---|---|---|---|
| **Revenue (US$ Mn)** | | | | | |
| **Bernstein** | **6,179** | **8,181** | **10,352** | **14,606** | **25,898** |
| *YoY* | *24%* | *32%* | *27%* | *41%* | *77%* |
| Consensus | 5,985 | 7,945 | 9,998 | 14,385 | 24,121 |
| *YoY* | *20%* | *33%* | *26%* | *44%* | *68%* |
| BERN vs. consensus | 3% | 3% | 4% | 2% | 7% |
| **Non-GAAP OpM** | | | | | |
| Bernstein | **44.4%** | **46.4%** | **46.7%** | **46.1%** | **47.8%** |
| *YoY* | 3% | 5% | 1% | -1% | 4% |
| Consensus | 42.1% | 42.5% | 43.2% | 43.2% | 45.4% |
| *YoY* | -2% | 1% | 2% | 0% | 5% |
| BERN vs. consensus | 2% | 4% | 4% | 3% | 2% |
| **Non-GAAP EPS (US$)** | | | | | |
| **Bernstein** | **2.28** | **3.12** | **3.94** | **5.43** | **9.83** |
| YoY | 29% | 37% | 26% | 38% | 81% |
| Consensus | 2.19 | 3.00 | 3.73 | 5.13 | 9.43 |
| YoY | 24% | 37% | 24% | 38% | 84% |
| BERN vs. consensus | 4% | 4% | 6% | 6% | 4% |

Source: Company reports, Bloomberg, Bernstein analysis and estimates

## VALUATION

Since their IPO in September 2023, Arm has traded at 70x 1-year forward P/E on average. Certainly, part of this is due to Arm's low float % and scarcity value, as SoftBank Group holds 87% of the shares. However, it is worth highlighting that semis IP companies trade, on average, at considerably higher multiples than other companies in the semiconductor sector, with the average FY+1 P/E multiple trading at an impressive ~140x, albeit largely driven by the abnormal multiples of Chinese players. On FY+2 basis, Arm is also below the average. The higher industry multiple can be explained by an asset-light, software-like model with high margins, minimal COGS, recurring royalties, and strong operating leverage (Exhibit 117)

We believe the high multiple is justified by Arm's position as likely the most reliable growth story in the sector, supported by very strong visibility: ~70% of the royalty forecast is based on existing contracts, while the $15bn AGI CPU estimate is based on actual customer engagement. Arm is expected to deliver a 40% Revenue CAGR and similar EPS growth, with potentially meaningful upside given the massive demand for CPUs driven by Agentic AI. Its monopoly position in the Mobile IP CPU market and rapid market share gains in the server CPU market also support Arm's premium valuation.

We forecast Arm's revenue to grow more than five fold to $26 Bn by 2030, with EPS expanding similarly (~5.5x) to $9.83, supported by the growing adoption of Arm CPUs in AI data centers, its own Arm CPU revenue, and the rising royalty driven by higher value provided by Arm for its clients rapid evolution of CPUs. **With 90x P/E on Q5-8 EPS of $3.33, we rate Arm Outperform with PT= $300.** (Exhibit 118, Exhibit 119, Exhibit 120, Exhibit 121)

**In a bull case scenario** (Exhibit 122), we can see Arm being valued at 40x on FY31E EPS. This 40x multiple is derived from a SOTP analysis, whereby we assign 70x (historical average 1-year forward P/E for Arm) multiple to the traditional IP / CSS business and 30x (historical average 1-year forward P/E for AMD) multiple to the Arm Silicon (CPU) business as AMD would be the most relevant peer for a fabless server CPU business. The blended P/E is 53x, to which we discount back 3 years at WACC of 10%, which yields 40x 1-year forward multiple. **Applying 40x P/E on FY31E EPS of $9.83, we derive the bull-case valuation of $390 (Exhibit 123)**

EXHIBIT 117: **Semiconductor IP companies comparison table**

| | Last price | Crncy | Market cap (USD Mn) | Adj P/E FY+1 | Adj P/E FY+2 | Adj P/E FY+3 | P/B FY+1 | P/B FY+2 | P/B FY+3 |
|---|---|---|---|---|---|---|---|---|---|
| **Semiconductor IP companies** | | | | | | | | | |
| Arm | 221.2 | USD | 234,925 | 97.1x | 70.9x | 56.1x | 20.5x | 14.9x | 11.2x |
| Cadence | 354.6 | USD | 97,791 | 44.6x | 37.9x | 32.7x | 13.9x | 11.8x | 10.0x |
| Synopsys | 509.3 | USD | 97,559 | 35.2x | 29.8x | 25.2x | 2.9x | 2.9x | 2.6x |
| Verisilicon | 277.5 | CNY | 21,492 | 483.0x | 186.3x | 86.0x | 40.7x | 34.2x | 23.4x |
| Rambus | 134.9 | USD | 14,582 | 45.8x | 36.9x | 28.1x | 9.0x | 7.4x | 5.7x |
| eMemory Tech | 3,920.0 | TWD | 9,289 | 108.8x | 76.1x | 54.9x | 56.9x | 43.2x | 30.1x |
| Empyrean | 96.0 | CNY | 7,709 | 236.5x | 144.0x | 107.5x | 9.8x | 9.4x | 8.5x |
| Ceva | 38.2 | USD | 1,063 | 72.0x | 46.3x | 37.9x | 3.0x | 2.9x | 2.7x |
| **IP Peers Average** | | | | **140.4x** | **78.5x** | **53.5x** | **19.6x** | **15.8x** | **11.8x** |

*Based on 13th May 2026 price

Arm estimates based on Bernstein. All the other companies are not covered and estimates are based on Bloomberg
Source: Bloomberg, Bernstein analysis and estimates

EXHIBIT 118: **Since the IPO in September 2023, Arm has traded at an average 1Y forward P/E of 70x. Although the stock has been highly volatile over the past two and a half years, it is now trading above +1SD of historical range.**

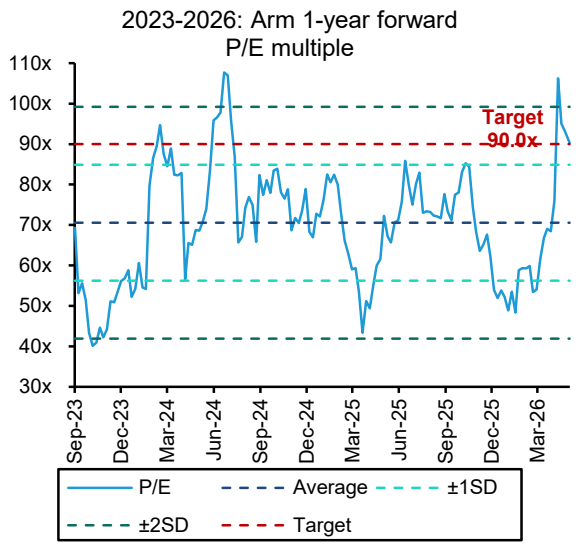


Source: Bloomberg, Bernstein analysis and estimates

EXHIBIT 119: **Arm's premium over the IP semis sector has averaged 1.1x, and it is now trading slightly higher around 1.2x.**

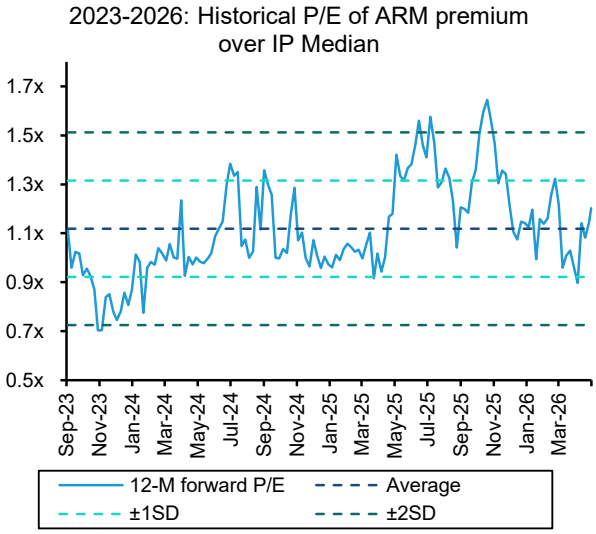


*IP companies included here are the same companies shown in the comparison table above
Source: Bloomberg, Bernstein analysis

EXHIBIT 120: **Arm's historical premium to the SOX has averaged 2.9x, but re-rated recently after the AGI CPU update.**

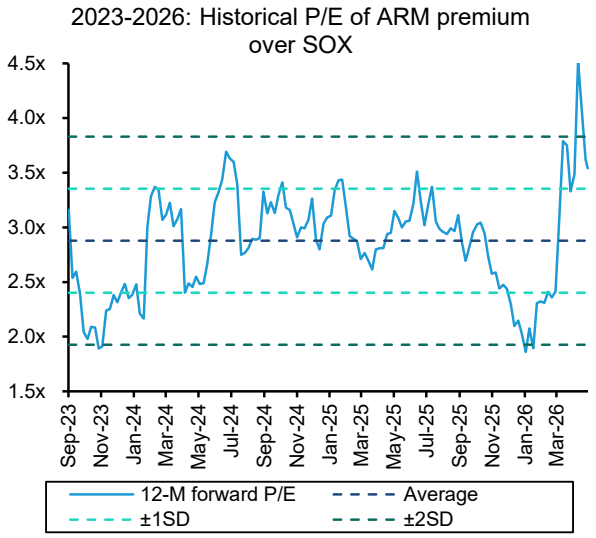


Source: Company reports, Bernstein analysis

EXHIBIT 121: **Arm's premium to the NASDAQ has averaged 2.8x and has been highly volatile. After the 3Q26 results, Arm was trading close to a historical low premium. However, following the sharp rally after the investor day, it is now trading above +2SD of range.**

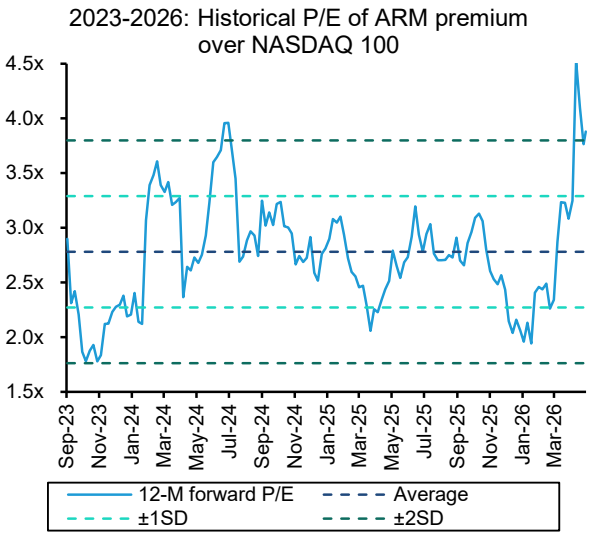


Source: Bloomberg, Bernstein analysis

EXHIBIT 122: **Our bull case analysis yields an implied target P/E of 40x, and a valuation of $390.**

| FY31E Financials | IP / CSS | Arm Silicon |
|---|---|---|
| Segment OP | 7,122 | 5,253 |
| PF Net Income (USD mn) | 6,054 | 4,465 |
| 1-year Forward P/E | 70x | 30x |
| PF Market Cap (USD mn) | 423,746 | 133,950 |
| Arm PF Market Cap | | 557,696 |
| Arm PF Net Income (USD mn) | | 10,519 |
| PF Blended Fwd P/E (FY31E) | | 53.0x |
| Discount Factor (3 years, WACC=10%) | | 1.33 |
| **Implied P/E (Discounted to FY28E)** | | **39.8x** |
| Arm EPS (SCBe) | | $9.83 |
| **Arm Bull Case TP** | | **$390** |

Source: Company disclosures, Bernstein estimates and analysis.

EXHIBIT 123: **AGI CPU comparison table**

| | Last price | Crncy | Market cap (USD Mn) | Adj P/E FY+1 | Adj P/E FY+2 | Adj P/E FY+3 | P/B FY+1 | P/B FY+2 | P/B FY+3 |
|---|---|---|---|---|---|---|---|---|---|
| **AGI CPU companies** | | | | | | | | | |
| Broadcom | 416.8 | USD | 1,973,362 | 37.0x | 23.3x | 18.2x | 18.2x | 11.4x | 7.8x |
| AMD | 445.5 | USD | 726,433 | 61.2x | 34.6x | 25.6x | 10.3x | 8.5x | 6.8x |
| Intel | 120.3 | USD | 604,578 | 109.0x | 77.5x | 53.7x | 5.1x | 4.7x | 4.0x |
| Qualcomm | 213.2 | USD | 224,681 | 19.9x | 20.2x | 18.1x | 9.5x | 8.9x | 7.9x |
| MediaTek | 3,495.0 | TWD | 177,855 | 51.8x | 28.9x | 19.2x | 12.9x | 10.5x | 8.0x |
| Marvell | 178.0 | USD | 159,680 | 46.5x | 32.4x | 23.8x | 9.8x | 8.4x | 6.7x |
| Global Unichip Corp | 5,330.0 | TWD | 22,663 | 108.7x | 54.4x | 42.4x | 40.8x | 19.7x | 19.4x |
| Alchip Tech | 4,780.0 | TWD | 12,365 | 35.4x | 26.9x | 19.9x | 8.1x | 6.9x | 5.7x |
| **AGI CPU Average** | | | | **58.7x** | **37.3x** | **27.6x** | **14.3x** | **9.9x** | **8.3x** |

*Based on 13th May Price

All estimates are based on Bloomberg. Marvelly, GUC and Alchip not covered.
Source: Bloomberg, Bernstein Analysis

## ARM COMPANY PROFILE HISTORY / MANAGEMENT

### History

Arm's story begins in 1978 with the founding of Acorn Computers in Cambridge, and the key technical breakthrough came in 1985, when Sophie Wilson and Steve Furber designed the first ARM1 processor. That chip established the core design philosophy that still defines Arm today: high performance with exceptional power efficiency. The first commercial system built around that architecture was the Acorn Archimedes in 1987, widely regarded as the first RISC-based home computer.

The next decisive milestone came in 1990, when Arm was formally created as Advanced RISC Machines as a joint venture between Acorn, Apple and VLSI Technology. In the early 1990s, two events shaped the future business model. First, the Apple Newton showed the value of Arm's low-power architecture, even if the device itself was not a major commercial success. Second, Arm moved to an IP licensing and royalty model, allowing many semiconductor companies to build products on the same architecture rather than relying on a single end device. That model became the foundation of Arm's economics.

Arm's real scale-up began with mobile. In 1993, Arm signed its deal with Texas Instruments, which helped bring Arm into Nokia's GSM phones. The later success of the Nokia 6110 and the broader mobile market made Arm the dominant architecture in handsets, and that leadership ultimately extended into the smartphone era, where Arm today powers more than 99% of smartphones. Arm then broadened from mobile into embedded systems, consumer devices, automotive, PCs and data center infrastructure, building what is now the most pervasive compute platform in the industry.

In the early 2000s, Arm moved from feature phones into the first smartphone era, supporting devices such as the Ericsson R380, early Palm OS products and BlackBerry devices. The mid-to-late 2000s were even more important. Arm launched the Cortex family in 2004, Apple introduced the iPhone in 2007 on an Arm-based processor, and then secured an architecture licence in 2008, laying the foundations for its own custom silicon. By 2010, with the iPhone, iPad and the rise of Android and iOS, Arm had become the standard architecture of the smartphone era. Another key phase began in 2016, when SoftBank acquired Arm and took the company private. That period allowed Arm to invest more aggressively with a longer-term horizon, particularly around AI, IoT and new end markets, before returning to public markets in 2023 through its Nasdaq IPO. SoftBank retains 87% of Arm post IPO.

The modern phase of the story began in 2023, when Arm returned to the public markets through its Nasdaq IPO. Since then, the company has accelerated the transition from a pure IP vendor to a broader compute platform company. Arm highlighted in its post-IPO update that Armv9 had already reached 25% of royalty revenue, and that CSS had expanded across multiple end markets. By Q4 FY26, Arm had reached 23 CSS licenses across 12 companies, with five customers already shipping CSS-based chips, including the top four Android smartphone vendors.

The latest milestone came in March 2026, when Arm announced the AGI CPU, its first production silicon product and a historic step beyond IP and CSS into Arm-designed chips. That launch marks the beginning of a new phase in which Arm participates not only through architecture and royalties, but also through direct silicon in AI infrastructure.

### Management

Arm's management is led by **Rene Haas**, who has been **CEO since February 2022** after joining Arm in 2013 and running the IP Products Group from 2017 to 2022. Under his leadership of IPG, Arm increased investment in the software ecosystem and expanded into growth markets such as infrastructure and automotive, and as CEO he has since led the company back to the public markets and into its first production silicon era. Haas succeeded Simon Segars, who served as CEO from 2013 to 2022 after spending 30 years at Arm.

On the finance side, **Jason Child** has been **CFO since November 2022**. Arm appointed him specifically as it prepared for a public listing, highlighting his public-company and IPO experience, and he replaced Inder Singh, who had served as CFO from 2019 to 2022.

Below the CEO and CFO, the most visible operating leaders today are Chris Bergey, who has been at Arm for a little over six years and now runs Edge AI, Mohamed Awad, who leads Cloud AI, and Drew Henry, who joined Arm in 2018 and was the founding GM of the infrastructure business before taking on Physical AI.

## APPENDIX - FINANCIAL FORECASTS

EXHIBIT 124: **SoftBank: Income Statement**

| JPY bn | FY23/3 | FY24/3 | FY25/3 | FY26/3 | FY27/3E | FY28/3E | FY29/3E |
|---|---|---|---|---|---|---|---|
| **Revenue** | **6,570.4** | **6,756.5** | **7,243.8** | **7,798.7** | **8,419.5** | **9,054.1** | **9,369.3** |
| **Gross profit** | **3,328.0** | **3,542.4** | **3,754.2** | **4,016.1** | **4,393.1** | **4,781.2** | **4,932.9** |
| SG&A | 2,695.3 | 2,982.4 | 3,024.4 | 4,020.9 | 5,724.4 | 3,938.5 | 4,075.6 |
| Non-operating income | -266.8 | 57.1 | -2,726.2 | -1,146.8 | -2,102.6 | -2,086.1 | -2,086.1 |
| **Pretax Income** | **-469.1** | **57.8** | **1,704.7** | **6,134.9** | **2,012.5** | **1,220.6** | **1,447.9** |
| Income tax expense | 320.7 | -151.4 | 101.6 | 502.9 | 503.1 | 305.1 | 362.0 |
| Net income before minority interest | -789.8 | 209.2 | 1,603.1 | 5,632.0 | 1,509.4 | 915.4 | 1,085.9 |
| **Net Income** | **-970.1** | **-227.6** | **1,153.3** | **5,002.3** | **1,438.9** | **872.6** | **1,035.1** |
| **Diluted EPS (JPY)** | **-154.6** | **-44.2** | **194.0** | **872.5** | **251.3** | **152.4** | **180.8** |
| **Margins** | | | | | | | |
| Gross profit | 50.7% | 52.4% | 51.8% | 51.5% | 52.2% | 52.8% | 52.6% |
| EBITDA | 59.9% | 60.3% | 121.3% | 153.7% | 108.8% | 95.8% | 97.2% |
| Profit before tax | -7.1% | 0.9% | 23.5% | 78.7% | 23.9% | 13.5% | 15.5% |
| Net income | -14.8% | -3.4% | 15.9% | 64.1% | 17.1% | 9.6% | 11.0% |
| **Growth YoY** | | | | | | | |
| Revenue | 5.6% | 2.8% | 7.2% | 7.7% | 8.0% | 7.5% | 3.5% |
| Gross profit | 1.9% | 6.4% | 6.0% | 7.0% | 9.4% | 8.8% | 3.2% |
| EBITDA | 261.7% | 3.6% | 115.7% | 36.4% | -23.5% | -5.3% | 4.9% |
| Profit before tax | N.M. | N.M. | 2849.3% | 259.9% | -67.2% | -39.4% | 18.6% |
| Net income | N.M. | N.M. | N.M. | 333.7% | -71.2% | -39.4% | 18.6% |

Source: Company disclosures, Bernstein estimates and analysis.

EXHIBIT 125: **SoftBank: Balance Sheet**

| JPY bn | FY23/3 | FY24/3 | FY25/3 | FY26/3 | FY27/3E | FY28/3E | FY29/3E |
|---|---|---|---|---|---|---|---|
| Cash and cash equivalents | 7,105.4 | 6,437.8 | 3,976.2 | 5,362.2 | -1,926.0 | -133.9 | 1,920.6 |
| Accounts receivable | 2,358.9 | 2,562.1 | 2,679.7 | 3,302.6 | 3,565.9 | 3,834.7 | 3,968.2 |
| Inventories | 163.8 | 161.9 | 198.3 | 240.2 | 253.7 | 269.2 | 279.5 |
| Other current assets | 958.4 | 2,279.6 | 2,578.8 | 2,910.7 | 2,910.7 | 2,910.7 | 2,910.7 |
| **Total current assets** | **10,586.5** | **11,441.4** | **9,432.9** | **11,815.6** | **4,804.3** | **6,880.7** | **9,079.0** |
| Property, plant and equipment, net | 2,639.7 | 2,642.2 | 3,688.1 | 4,368.2 | 6,493.1 | 5,349.4 | 4,494.7 |
| Long-term investments | 18,196.2 | 20,076.5 | 19,597.5 | 27,760.3 | 37,919.5 | 37,919.5 | 37,919.5 |
| Other non-current assets | 12,513.9 | 12,564.2 | 12,295.2 | 16,805.4 | 16,805.4 | 16,805.4 | 16,805.4 |
| **Total assets** | **43,936.4** | **46,724.2** | **45,013.8** | **60,749.5** | **66,022.3** | **66,955.0** | **68,298.6** |
| Accounts payable and accrued expenses | 3,081.1 | 3,232.2 | 3,038.1 | 3,799.2 | 4,013.0 | 4,258.8 | 4,421.6 |
| ST Debt and Lease Obligations | 5,313.2 | 8,420.9 | 5,795.0 | 7,436.3 | 7,436.3 | 7,436.3 | 7,436.3 |
| Other current liabilities | 2,186.5 | 2,367.1 | 3,768.6 | 3,598.1 | 3,598.1 | 3,598.1 | 3,598.1 |
| **Total current liabilities** | **10,580.7** | **14,020.3** | **12,601.7** | **14,833.5** | **15,047.4** | **15,293.2** | **15,456.0** |
| LT Debt and Lease Obligations | 15,002.0 | 12,941.1 | 13,118.3 | 18,227.3 | 21,602.3 | 21,602.3 | 21,602.3 |
| Other non-current liabilities | 7,704.4 | 6,525.7 | 5,340.7 | 7,220.3 | 7,220.3 | 7,220.3 | 7,220.3 |
| **Total liabilities** | **33,287.2** | **33,487.1** | **31,060.7** | **40,281.1** | **43,870.0** | **44,115.7** | **44,278.6** |
| **Shareholders' equity** | **10,649.2** | **13,237.2** | **13,953.0** | **20,468.4** | **22,152.3** | **22,839.3** | **24,020.0** |
| **Total liabilities & shareholders' equity** | **43,936.4** | **46,724.2** | **45,013.8** | **60,749.5** | **66,022.3** | **66,955.0** | **68,298.6** |

Source: Company disclosures, Bernstein estimates and analysis.

EXHIBIT 126: **SoftBank: Cash Flow Statement**

| JPY bn | FY23/3 | FY24/3 | FY25/3 | FY26/3 | FY27/3E | FY28/3E | FY29/3E |
|---|---|---|---|---|---|---|---|
| **Net Income** | **-970.1** | **-227.6** | **1,153.3** | **5,002.3** | **1,438.9** | **872.6** | **1,035.1** |
| D&A | 893.5 | 858.6 | 866.8 | 918.8 | 1,020.7 | 1,097.6 | 1,135.8 |
| Changes in working capital | 4,241.5 | -1,870.1 | -124.1 | -1,448.9 | -62.9 | -38.6 | 19.1 |
| Other adjustments | -3,423.5 | 1,489.7 | -1,692.5 | -4,901.0 | -5,446.4 | 0.0 | 0.0 |
| **Cash flow from operations** | **741.3** | **250.5** | **203.6** | **-428.8** | **-3,049.9** | **1,931.6** | **2,190.0** |
| Capex | -633.8 | -622.6 | -854.2 | -1,733.8 | -2,827.9 | -271.6 | -281.1 |
| Other | 1,181.3 | -218.8 | -777.4 | -2,773.3 | -4,712.7 | 0.0 | 0.0 |
| **Cash flow from investments** | **547.6** | **-841.5** | **-1,631.5** | **-4,507.2** | **-7,540.6** | **-271.6** | **-281.1** |
| **Cash flow from financing** | **191.5** | **-606.2** | **-1,116.4** | **6,377.3** | **3,302.3** | **132.0** | **145.6** |
| Net effect of exchange rate changes | 211.7 | 529.5 | 82.7 | -55.3 | 0.0 | 0.0 | 0.0 |
| **Net change in cash** | **1,692.1** | **-667.6** | **-2,544.3** | **1,441.3** | **-7,288.1** | **1,792.0** | **2,054.6** |

Source: Company disclosures, Bernstein estimates and analysis.

## EXHIBIT 127: **Arm Income Statement**

| Non-GAAP (USD Mn) | FY22/3 | FY23/3 | FY24/3 | FY25/3 | FY26/3 | FY27/3E | FY28/3E | FY29/3E | FY30/3E | FY31/3E |
|---|---|---|---|---|---|---|---|---|---|---|
| **Revenue** | **2,703.0** | **2,679.0** | **3,233.0** | **4,007.0** | **4,920.0** | **6,179.5** | **8,180.8** | **10,352.2** | **14,606.4** | **25,898.1** |
| **Gross profit** | **2,582.0** | **2,580.0** | **3,131.0** | **3,920.0** | **4,832.0** | **6,021.7** | **7,428.0** | **9,064.7** | **11,780.8** | **18,566.6** |
| SG&A | 875.0 | 700.0 | 595.0 | 709.0 | 806.0 | 959.4 | 1,057.7 | 1,200.0 | 1,416.1 | 1,720.5 |
| R&D | 976.0 | 1,095.0 | 1,128.0 | 1,340.0 | 1,911.0 | 2,320.7 | 2,576.0 | 3,027.5 | 3,634.2 | 4,471.4 |
| Other opex | 0.0 | 2.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 |
| **Operating profit** | **731.0** | **783.0** | **1,408.0** | **1,871.0** | **2,115.0** | **2,741.6** | **3,794.3** | **4,837.1** | **6,730.6** | **12,374.7** |
| D&A | 185.0 | 170.0 | 162.0 | 183.0 | 249.0 | 294.5 | 389.8 | 493.3 | 696.0 | 1,165.4 |
| **EBITDA** | **916.0** | **953.0** | **1,570.0** | **2,054.0** | **2,364.0** | **3,036.1** | **4,184.1** | **5,330.4** | **7,426.6** | **13,540.1** |
| Non-operating income | 70.0 | 41.0 | 121.0 | 126.0 | 107.0 | 123.6 | 163.6 | 207.0 | 292.1 | 518.0 |
| **Pretax income** | **801.0** | **783.0** | **931.0** | **1,997.0** | **2,222.0** | **2,865.2** | **3,957.9** | **5,044.2** | **7,022.7** | **12,892.7** |
| Income tax expense | 110.0 | 147.0 | -94.0 | -94.0 | 253.0 | 429.8 | 593.7 | 756.6 | 1,053.4 | 1,933.9 |
| Net income before minority interest | 661.0 | 656.0 | 1,324.0 | 1,737.0 | 1,889.0 | 2,435.4 | 3,364.2 | 4,287.6 | 5,969.3 | 10,958.8 |
| **Net Income** | **663.0** | **657.0** | **1,324.0** | **1,737.0** | **1,889.0** | **2,435.4** | **3,364.2** | **4,287.6** | **5,969.3** | **10,958.8** |
| **Diluted EPS (USD)** | **0.65** | **0.64** | **1.27** | **1.63** | **1.77** | **2.28** | **3.12** | **3.94** | **5.43** | **9.83** |
| **Margins** | | | | | | | | | | |
| Gross profit | 95.5% | 96.3% | 96.8% | 97.8% | 98.2% | 97.4% | 90.8% | 87.6% | 80.7% | 71.7% |
| Operating profit | 27.0% | 29.2% | 43.6% | 46.7% | 43.0% | 44.4% | 46.4% | 46.7% | 46.1% | 47.8% |
| EBITDA | 33.9% | 35.6% | 48.6% | 51.3% | 48.0% | 49.1% | 51.1% | 51.5% | 50.8% | 52.3% |
| Profit before tax | 29.6% | 29.2% | 28.8% | 49.8% | 45.2% | 46.4% | 48.4% | 48.7% | 48.1% | 49.8% |
| Net income | 24.5% | 24.5% | 41.0% | 43.3% | 38.4% | 39.4% | 41.1% | 41.4% | 40.9% | 42.3% |
| **Growth YoY** | | | | | | | | | | |
| Revenue | 33.3% | -0.9% | 20.7% | 23.9% | 22.8% | 25.6% | 32.4% | 26.5% | 41.1% | 77.3% |
| Gross profit | 36.3% | -0.1% | 21.4% | 25.2% | 23.3% | 24.6% | 23.4% | 22.0% | 30.0% | 57.6% |
| Operating profit | 140.5% | 7.1% | 79.8% | 32.9% | 13.0% | 29.6% | 38.4% | 27.5% | 39.1% | 83.9% |
| EBITDA | 81.4% | 4.0% | 64.7% | 30.8% | 15.1% | 28.4% | 37.8% | 27.4% | 39.3% | 82.3% |
| Profit before tax | 124.4% | -2.2% | 18.9% | 114.5% | 11.3% | 28.9% | 38.1% | 27.4% | 39.2% | 83.6% |
| Net income | 220.3% | -0.9% | 101.5% | 31.2% | 8.8% | 28.9% | 38.1% | 27.4% | 39.2% | 83.6% |

Source: Company reports, Bernstein analysis and estimates

## EXHIBIT 128: **Arm Balance Sheet**

| USD Mn | FY22/3 | FY23/3 | FY24/3 | FY25/3 | FY26/3 | FY27/3E | FY28/3E | FY29/3E | FY30/3E | FY31/3E |
|---|---|---|---|---|---|---|---|---|---|---|
| Cash and cash equivalents | 1,004.0 | 1,554.0 | 2,923.0 | 2,825.0 | 3,601.0 | 6,233.7 | 9,863.0 | 14,256.5 | 19,601.7 | 28,331.6 |
| Accounts receivable | 631.0 | 661.0 | 781.0 | 1,107.0 | 1,300.0 | 1,275.2 | 1,688.2 | 1,985.4 | 2,801.2 | 4,257.2 |
| Contract assets - Current | 166.0 | 154.0 | 336.0 | 642.0 | 977.0 | 1,268.2 | 1,364.4 | 1,810.9 | 2,515.6 | 4,289.6 |
| Other current assets | 1,291.0 | 1,168.0 | 157.0 | 256.0 | 358.0 | 362.0 | 362.0 | 362.0 | 362.0 | 362.0 |
| **Total current assets** | **3,092.0** | **3,537.0** | **4,197.0** | **4,830.0** | **6,236.0** | **9,139.1** | **13,277.6** | **18,414.7** | **25,280.4** | **37,240.4** |
| Property, plant and equipment, net | 188.0 | 185.0 | 420.0 | 714.0 | 1,220.0 | 1,419.9 | 1,520.9 | 1,545.2 | 1,579.6 | 1,320.6 |
| Long-term investments | 736.0 | 723.0 | 741.0 | 565.0 | 387.0 | 387.0 | 387.0 | 387.0 | 387.0 | 387.0 |
| Other non-current assets | 2,494.0 | 2,421.0 | 2,569.0 | 2,823.0 | 2,860.0 | 2,955.4 | 2,986.9 | 3,133.1 | 3,363.9 | 3,945.0 |
| **Total assets** | **6,510.0** | **6,866.0** | **7,927.0** | **8,932.0** | **10,703.0** | **13,901.4** | **18,172.4** | **23,480.0** | **30,610.9** | **42,893.0** |
| ST Debt and Lease Obligations | 31.0 | 26.0 | 27.0 | 30.0 | 39.0 | 39.0 | 39.0 | 39.0 | 39.0 | 39.0 |
| Other current liabilities | 1,364.0 | 1,337.0 | 1,478.0 | 899.0 | 1,001.0 | 1,001.0 | 1,001.0 | 1,001.0 | 1,001.0 | 1,001.0 |
| **Total current liabilities** | **1,395.0** | **1,363.0** | **1,505.0** | **929.0** | **1,040.0** | **1,040.0** | **1,040.0** | **1,040.0** | **1,040.0** | **1,040.0** |
| LT Debt and Lease Obligations | 230.0 | 193.0 | 194.0 | 316.0 | 393.0 | 393.0 | 393.0 | 393.0 | 393.0 | 393.0 |
| Other non-current liabilities | 1,337.0 | 1,259.0 | 933.0 | 848.0 | 984.0 | 984.0 | 984.0 | 984.0 | 984.0 | 984.0 |
| **Total liabilities** | **2,962.0** | **2,815.0** | **2,632.0** | **2,093.0** | **2,417.0** | **2,417.0** | **2,417.0** | **2,417.0** | **2,417.0** | **2,417.0** |
| **Shareholders' equity** | **3,548.0** | **4,051.0** | **5,295.0** | **6,839.0** | **8,286.0** | **11,484.4** | **15,755.4** | **21,063.0** | **28,193.9** | **40,476.0** |
| **Total liabilities & shareholders' equity** | **6,510.0** | **6,866.0** | **7,927.0** | **8,932.0** | **10,703.0** | **13,901.4** | **18,172.4** | **23,480.0** | **30,610.9** | **42,893.0** |

Source: Company reports, Bernstein analysis and estimates

## EXHIBIT 129: **Arm Cash Flow Statement**

| Non-GAAP (USD Mn) | FY22/3 | FY23/3 | FY24/3 | FY25/3 | FY26/3 | FY27/3E | FY28/3E | FY29/3E | FY30/3E | FY31/3E |
|---|---|---|---|---|---|---|---|---|---|---|
| **Net Income** | **663.0** | **657.0** | **1,324.0** | **1,737.0** | **1,889.0** | **2,435.4** | **3,364.2** | **4,287.6** | **5,969.3** | **10,958.8** |
| D&A | 185.0 | 170.0 | 162.0 | 183.0 | 249.0 | 294.5 | 389.8 | 493.3 | 696.0 | 1,165.4 |
| Changes in working capital | 669.7 | -491.8 | -757.8 | -1,483.0 | 190.0 | -361.8 | -540.7 | -889.8 | -1,751.4 | -3,811.2 |
| Other adjustments | -1,059.7 | 403.8 | 361.8 | -860.0 | -804.0 | 1,036.6 | 1,232.1 | 1,385.8 | 1,578.1 | 1,797.9 |
| **Cash flow from operations** | **458.0** | **739.0** | **1,090.0** | **-423.0** | **1,524.0** | **3,404.7** | **4,445.4** | **5,276.9** | **6,492.1** | **10,110.9** |
| Capex | -18.2 | -64.0 | -92.0 | -219.0 | -545.0 | -494.4 | -490.8 | -517.6 | -730.3 | -906.4 |
| Other | -28.7 | 37.3 | -424.0 | 184.0 | 220.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 |
| **Cash flow from investments** | **-46.9** | **-26.7** | **-516.0** | **-35.0** | **-325.0** | **-494.4** | **-490.8** | **-517.6** | **-730.3** | **-906.4** |
| **Cash flow from financing** | **-68.7** | **-42.0** | **-208.0** | **-202.0** | **-548.0** | **-277.6** | **-325.2** | **-365.8** | **-416.6** | **-474.6** |
| Net effect of exchange rate changes | 35.6 | -647.3 | 3.0 | 2.0 | 15.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 |
| **Net change in cash** | **378.0** | **23.0** | **-568.0** | **-660.0** | **651.0** | **2,632.7** | **3,629.3** | **4,393.5** | **5,345.2** | **8,729.9** |

Source: Company reports, Bernstein analysis and estimates

## I. REQUIRED DISCLOSURES

References to "Bernstein" or the “Firm” in these disclosures relate to the following entities: Bernstein Institutional Services LLC (April 1, 2024 onwards), Sanford C. Bernstein & Co., LLC (pre April 1, 2024), Bernstein Autonomous LLP, BSG France S.A. (April 1, 2024 onwards), Sanford C. Bernstein (Hong Kong) Limited 盛博香港有限公司, Sanford C. Bernstein (Canada) Limited, Sanford C. Bernstein (India) Private Limited (SEBI registration no. INH000006378), Sanford C. Bernstein (Singapore) Private Limited, Sanford C. Bernstein Japan KK（サンフォード・C・バーンスタイン株式会社）and analysts employed by Société Générale Africa Technologies & Services to produce Bernstein research under a Global Services Agreement in place between Bernstein and Société Générale.

Bernstein is part of a joint venture between Société Générale (SG) and AllianceBernstein, L.P. (AB). Unless specifically noted otherwise, for purposes of these disclosures, references to Bernstein’s “affiliates” relate to both SG and AB and their respective affiliates.

### VALUATION METHODOLOGY

#### SoftBank Group Corp

*For SoftBank, we apply a 25% discount to our one-year forward net asset value (NAV) of $390bn, with a USD/JPY of 159.21, to derive our PT of ¥8,200.*

#### ARM Holdings PLC

*For Arm, we apply a 90x multiple to our Q5-Q8 EPS estimate of $3.33 to get a PT of $300.*

### RISKS

#### SoftBank Group Corp

Downside risks to our price target for SoftBank include: (1) lower-than-expected valuation for Arm, driven by slower growth in its emerging AGI CPU business and limited market share gains in the server CPU segment; (2) potential market share loss at OpenAI amid intensifying competition; (3) execution risk in OpenAI’s monetization strategy, particularly its ability to scale revenues to offset rising operating expenses and computing costs; and (4) financial stability risks, including potential credit-rating downgrades, stemming from continued large-scale investments in OpenAI and other strategic initiatives.

#### ARM Holdings PLC

Downside risks to our price target for Arm include slower growth in the new AGI CPU business and a lack of market share gains in the server CPU market; a slower increase in royalty rates across different end-markets; slower v9 and CSS penetration; higher-than-expected costs to ramp up the AGI CPU business; and higher OPEX required to sustain the massive projected growth.

### RATINGS DEFINITIONS, BENCHMARKS AND DISTRIBUTION

### EQUITY RATINGS DEFINITIONS

#### Bernstein brand

The Bernstein brand rates stocks based on forecasts of relative performance for the next 12 months versus the S&P 500 for stocks listed on the U.S. and Canadian exchanges, versus the Bloomberg Europe Developed Markets Large and Mid Cap Price Return Index EUR (EDME) for stocks listed on the European exchanges and emerging markets exchanges outside of the Asia Pacific region, versus the Bloomberg Japan Large and Mid Cap Price Return Index USD (JPL) for stocks listed on the Japanese exchanges, and versus the Bloomberg Asia ex-Japan Large and Mid Cap Price Return Index (ASIAX) for stocks listed on the Asian (ex-Japan) exchanges -unless otherwise specified.

The Bernstein brand has three categories of ratings:

- Outperform: Stock will outpace the market index by more than 15 pp
- Market-Perform: Stock will perform in line with the market index to within +/-15 pp

- Underperform: Stock will trail the performance of the market index by more than 15 pp

Coverage Suspended: Coverage of a company under the Bernstein research brand has been suspended. Ratings and price targets are suspended temporarily, are no longer current, and should therefore not be relied upon.

Not Rated: A rating assigned when the stock cannot be accurately valued, or the performance of the company accurately predicted, at the present time. The covering analyst may continue to publish research reports on the company to update investors on events and developments.

Not Covered (NC) denotes companies that are not under coverage.

Bernstein brand stock ratings are based on a 12-month time horizon.

**Autonomous brand – common stocks**

The Autonomous brand rates common stocks as indicated below. As our benchmarks we use the Bloomberg Europe 500 Banks And Financial Services Index (BEBANKS) and Bloomberg Europe Dev Mkt Financials Large and Mid Cap Price Ret Index EUR (EDMFI) index for developed European banks and Payments, the Bloomberg Europe 500 Insurance Index (BEINSUR) for European insurers, the S&P 500 and S&P Financials for US banks and Payments coverage, S5LIFE for US Insurance, the S&P Insurance Select Industry (SPSIINS) for US Non-Life Insurers coverage, and the Bloomberg Emerging Markets Financials Large, Mid and Small Cap Price Return Index (EMLSF) for emerging market banks and insurers and Payments. Ratings are stated relative to the sector (not the market).

The Autonomous brand has three categories of common stock ratings:

- Outperform (OP): Stock will outpace the relevant index by more than 10 pp
- Neutral (N): Stock will perform in line with the market index to within +/-10 pp
- Underperform (UP): Stock will trail the performance of the relevant index by more than 10 pp

Coverage Suspended: Coverage of a company under the Autonomous research brand has been suspended. Ratings and price targets are suspended temporarily, are no longer current, and should therefore not be relied upon.

Not Rated: A rating assigned when the stock cannot be accurately valued, or the performance of the company accurately predicted, at the present time. The covering analyst may continue to publish research reports on the company to update investors on events and developments.

Those denoted as 'Feature' (e.g., Feature Outperform FOP, Feature Under Outperform FUP) are our core ideas.

Not Covered (NC) denotes companies that are not under coverage.

Autonomous brand common stock ratings are based on a 12-month time horizon.

**Autonomous brand – preferred stocks**

The Autonomous brand has three categories of preferred stock ratings:

- Outperform (OP): The total return of the preferred instrument is expected to outperform preferred securities of other issuers operating in similar sectors or rating categories over the next six months.
- Neutral (N): The total return of the preferred instrument is expected to perform in line with preferred securities of other issuers operating in similar sectors or rating categories over the next six months.
- Underperform (UP): The total return of the preferred instrument is expected to underperform preferred securities of other issuers operating in similar sectors or rating categories over the next six months.

Autonomous preferred stock ratings are based on a 6-month time horizon.

**AUTONOMOUS CREDIT RESEARCH**

Where this report contains investment recommendations for credit instruments, as defined in article 3(1)(35) of the Market Abuse Regulation, the information below is presented to comply with its disclosure requirements.

The report may also include reference(s) to published opinions by other Autonomous or Bernstein analysts covering the equity securities of the issuer(s) referenced herein. Please note an investment recommendation for credit instruments published by the author(s) of this report may differ from the published view of the analyst covering equity securities for the issuer(s) contained in this report and vice versa.

**CREDIT RATINGS DEFINITIONS**

The Autonomous brand has three categories of credit ratings:

- Credit Outperform (C-OP): The total return of the Reference Credit Instrument is expected to outperform the credit spread of bonds of other issuers operating in similar sectors or rating categories over the next six months.
- Credit Neutral (C-N): The total return of the Reference Credit Instrument is expected to perform in line with the credit spread of bonds of other issuers operating in similar sectors or rating categories over the next six months.
- Credit Underperform (C-UP): The total return of the Reference Credit Instrument is expected to underperform the credit spread of bonds of other issuers operating in similar sectors or rating categories over the next six months.

Autonomous credit ratings are based on a 6-month time horizon.

A list of all investment recommendations produced by the author(s) of this report alongside credit ratings history are available upon request.

It is at the sole discretion of the Firm as to when to initiate, update and cease research coverage. The Firm has established, maintains and relies on information barriers to control the flow of information contained in one or more areas (i.e. the private side) within the Firm, and into other areas, units, groups or affiliates (i.e. public side) of the Firm

**DISTRIBUTION OF EQUITY RATINGS/INVESTMENT BANKING SERVICES**

| Equity Rating | Market Abuse Regulation (MAR) and FINRA Rating Category | Global Rating Distribution | Investment Banking Relationships* |
|---|---|---|---|
| Outperform | BUY | 51.1% | 16.5% |
| Market-Perform (Bernstein Brand)<br>Neutral (Autonomous Brand) | HOLD | 36.3% | 17.8% |
| Underperform | SELL | 12.6% | 14.9% |

* These figures represent the percentage of companies within each equity rating category for which affiliates of Bernstein have provided investment banking services within the previous 12 months.
As of March 31, 2026. All figures are updated quarterly.

## PRICE CHARTS / RATINGS AND PRICE TARGET HISTORY

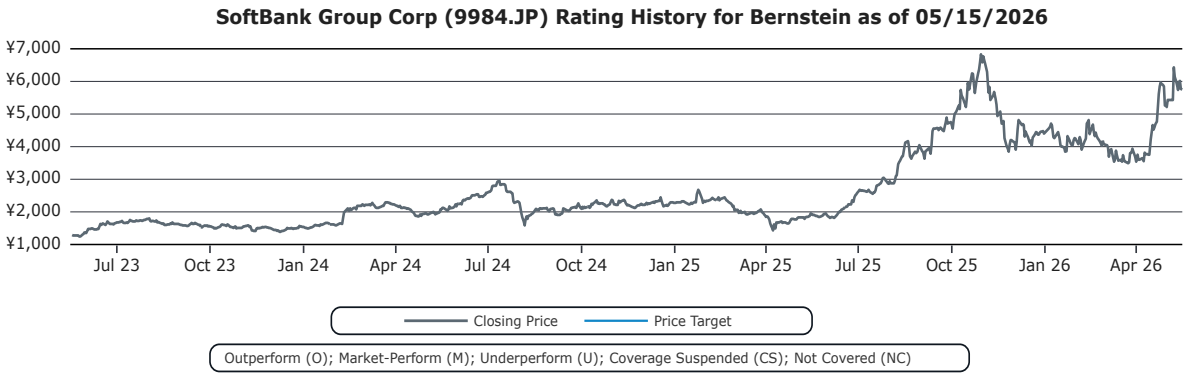


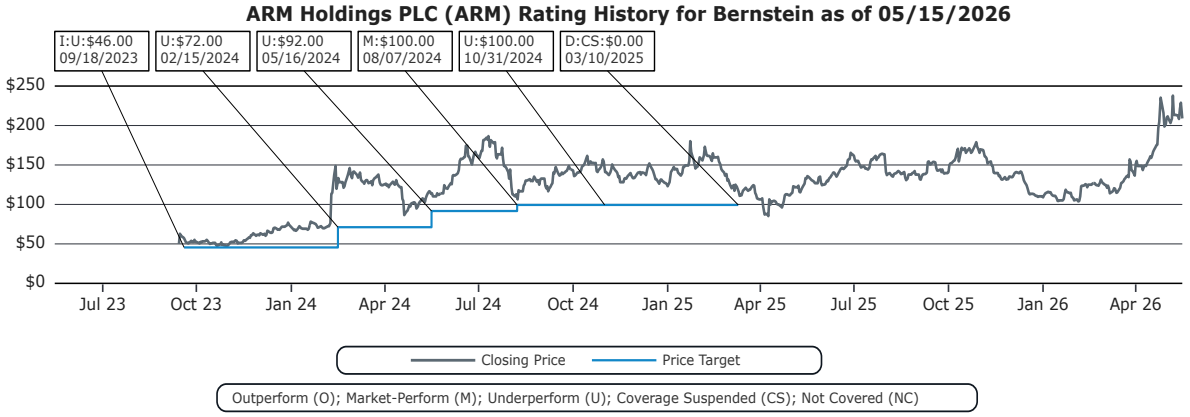


All price target and closing price data in the chart(s) above are denominated in the currency noted in the Ticker Table of this report.

## CONFLICTS OF INTEREST

Bernstein and/or affiliates have received compensation for investment banking services in the past twelve months from SoftBank Group Corp.

Bernstein and/or affiliates have received compensation for non-investment banking securities-related products or services in the previous twelve months from the following clients: SoftBank Group Corp and ARM Holdings PLC.

Bernstein and/or affiliates expect to receive or intend to seek compensation for investment banking services in the next three months from SoftBank Group Corp.

Bernstein and/or affiliates had an investment banking client relationship during the past twelve months with SoftBank Group Corp.

Certain affiliates of Bernstein act as market maker or liquidity provider in the debt securities of: SoftBank Group Corp and ARM Holdings PLC.

Affiliates of Bernstein managed or co-managed in the past twelve months a public offering of securities of SoftBank Group Corp.

Certain affiliates of Bernstein act as market maker or liquidity provider in the equities securities of: ARM Holdings PLC.

## OTHER MATTERS

The legal entity(ies) employing the analyst(s) listed in this report, and their location, can be determined by the country code of their phone number, as follows:

+1 Bernstein Institutional Services LLC; New York, New York, USA

+44 Bernstein Autonomous LLP; London UK

+212 Société Générale Africa Technologies & Services; Casablanca, Morocco

+33 BSG France S.A.; Paris, France

+34 BSG France S.A.; Madrid, Spain

+41 Bernstein Autonomous LLP; Geneva, Switzerland

+49 BSG France S.A.; Frankfurt, Germany

+91 Sanford C. Bernstein (India) Private Limited; Mumbai, India

+852 Sanford C. Bernstein (Hong Kong) Limited 盛博香港有限公司; Hong Kong, China

+65 Sanford C. Bernstein (Singapore) Private Limited; Singapore

+81 Sanford C. Bernstein Japan KK; Tokyo, Japan

Where this report has been prepared by research analyst(s) employed by a non-US affiliate, such analyst(s), is/are (unless otherwise expressly noted below) not registered as associated persons of Bernstein Institutional Services LLC or any other SEC-registered broker-dealer and are not licensed or qualified as research analysts with FINRA. Accordingly, such analyst(s) may not be subject to FINRA's restrictions regarding (among other things) communications by research analysts with a subject company, interactions between research analysts and investment banking personnel, participation by research analysts in solicitation and marketing activities relating to investment banking transactions, public appearances by research analysts, and trading securities held by a research analyst account.

Where this report has been prepared by research analyst(s) employed by Société Générale Africa Technologies & Services (part of the Société Générale group of companies), it has been prepared on behalf of a Bernstein company under a Global Services Agreement in place between Bernstein and Société Générale.

**CERTIFICATION**

Each research analyst listed in this report, who is primarily responsible for the preparation of the content of this report, certifies that all of the views expressed in this publication accurately reflect that analyst's personal views about any and all of the subject securities or issuers and that no part of that analyst's compensation was, is, or will be, directly or indirectly, related to the specific recommendations or views in this publication.

## II. ADDITIONAL GLOBAL CONFLICT DISCLOSURES

It is at the sole discretion of the Firm as to when to initiate, update and cease research coverage. The Firm has established, maintains and relies on information barriers to control the flow of information contained in one or more areas (i.e., the private side) within the Firm, and into other areas, units, groups or affiliates (i.e., public side) of the Firm.

## III. OTHER IMPORTANT INFORMATION AND DISCLOSURES

Separate branding is maintained for "Bernstein" and "Autonomous" research products.

- Bernstein produces a number of different types of research products including, among others, fundamental analysis and quantitative analysis under both the "Autonomous" and "Bernstein" brands. Recommendations contained within one type of research product may differ from recommendations contained within other types of research products, whether as a result of differing time horizons, methodologies or otherwise. Furthermore, views or recommendations within a research product issued under one brand may differ from views or recommendations under the same type of research product issued under the other brand. The Research Ratings System for the two brands and other information related to those Rating Systems are included in the previous section.

- Autonomous operates as a separate business unit within the following entities: Bernstein Institutional Services LLC, Bernstein Autonomous LLP, Sanford C. Bernstein (Hong Kong) Limited 盛博香港有限公司 and Sanford C. Bernstein (India) Private Limited. For information relating to "Autonomous" branded products (including certain Sales materials) please visit:

www.autonomous.com. For information relating to Bernstein branded products please visit: www.bernsteinresearch.com.

Analysts are compensated based on aggregate contributions to the research franchise as measured by account penetration, productivity and proactivity of investment ideas. No analysts are compensated based on performance in, or contributions to, generating investment banking revenues.

This report has been produced by an independent analyst as defined in Article 3 (1)(34)(i) of EU 596/2014 Market Abuse Regulation ("MAR") and the same article of MAR as it forms part of United Kingdom domestic law by virtue of the European Union (Withdrawal) Act 2018.

**To our readers in the United States**: Bernstein Institutional Services LLC, a broker-dealer registered with the U.S. Securities and Exchange Commission ("SEC") and a member of the U.S. Financial Industry Regulatory Authority, Inc. ("FINRA") is distributing this publication in the United States and accepts responsibility for its contents. Where this material contains an analysis of debt product(s), such material is intended only for institutional investors and is not subject to the US independence and disclosure standards applicable to debt research prepared for retail investors.

Bernstein Institutional Services LLC may act as principal for its own account or as agent for another person (including an affiliate) in sales or purchases of any security which is a subject of this report. This report does not purport to meet the objectives or needs of any specific individuals, entities or accounts.

**To our readers in Canada:** If this publication pertains to a Canadian domiciled company, it is being distributed in Canada by Sanford C. Bernstein (Canada) Limited, which is licensed and regulated by the Canadian Investment Regulatory Organization. If the publication pertains to a non-Canadian domiciled company, it is being distributed by Bernstein Institutional Services LLC, which is licensed and regulated by both the SEC and FINRA, into Canada under the International Dealers Exemption.

This document may not be passed onto any person in Canada unless that person qualifies as "permitted client" as defined in Section 1.1 of NI 31-103.

**To our readers in Brazil:** This report has been prepared by Bernstein Institutional Services LLC, and Banco BTG Pactual S.A. ("BTG") is responsible for the distribution of this report in Brazil.

**To readers in the United Kingdom**: This publication has been issued or approved for issue in the United Kingdom by Bernstein Autonomous LLP, authorised and regulated by the Financial Conduct Authority and located at 60 London Wall, London EC2M 5SH, +44 (0)20-7170-5000. Registered in England & Wales No OC343985.

This document is for distribution only to persons who (i) have professional experience in matters relating to investments falling within Article 19(5) of the Financial Services and Markets Act 2000 (Financial Promotion) Order 2005 (as amended, the "Financial Promotion Order"), (ii) are persons falling within Article 49(2)(a) to (d) ("high net worth companies, unincorporated associations, etc.") of the Financial Promotion Order, (iii) are outside the United Kingdom, or (iv) are persons to whom an invitation or inducement to engage in investment activity (within the meaning of section 21 of the FSMA) in connection with the issue or sale of any securities may otherwise lawfully be communicated or caused to be communicated (all such persons together being referred to as "relevant persons"). This document is directed only at relevant persons and must not be acted on or relied on by persons who are not relevant persons. Any investment or investment activity to which this document relates is available only to relevant persons and will be engaged in only with relevant persons.

**To our readers in the member states of the EEA**: This publication is being distributed by BSG France SA, which is authorised and regulated by the Autorité de Contrôle Prudentiel et de Résolution (ACPR) and Autorité des Marchés Financiers (AMF).

**To our readers in Hong Kong:** This publication is being distributed in Hong Kong by Sanford C. Bernstein (Hong Kong) Limited 盛博香港有限公司, which is licensed and regulated by the Hong Kong Securities and Futures Commission (Central Entity No. AXC846) to carry out Type 4 (Advising on Securities) regulated activities and subject to the licensing conditions mentioned in the SFC Public Register (https://www.sfc.hk/publicregWeb/corp/AXC846/details)). This publication is solely for professional investors, as defined in the Securities and Futures Ordinance (Cap. 571). The purpose of this report is solely to provide an analysis of the issuers referred to in this report and is not intended for any purpose contrary to the laws of Hong Kong.

**To our readers in Singapore:** This publication is being distributed in Singapore by Sanford C. Bernstein (Singapore) Private Limited, only to accredited investors or institutional investors, as defined in the Securities and Futures Act 2001 of Singapore ("SFA"). Recipients in Singapore should contact Sanford C. Bernstein (Singapore) Private Limited in respect of matters arising from, or in connection with, this publication. Sanford C. Bernstein (Singapore) Private Limited is regulated by the Monetary Authority of Singapore and licensed under the SFA as a capital markets services licence holder for dealing in capital markets products that are securities and collective investment schemes and an exempt financial adviser for advising on, issuing and promulgating analyses and reports on securities. Sanford C. Bernstein (Singapore) Private Limited is registered in Singapore with Company Registration

No. 202137 10W and located at One Raffles Quay, #27-11 South Tower, Singapore 048583, +65-6230-4612.

**To our readers in the People's Republic of China:** The securities referred to in this document are not being offered or sold and may not be offered or sold, directly or indirectly, in the People's Republic of China (for such purposes, not including the Hong Kong and Macau Special Administrative Regions or Taiwan, the "PRC") in contravention of any applicable laws of the PRC.

This document does not constitute an offer to sell or the solicitation of an offer to buy any securities in the PRC to any person to whom it is unlawful to make the offer or solicitation in the PRC.

We do not represent that this document may be lawfully distributed, or that any securities may be lawfully offered, in compliance with any applicable registration or other requirements in the PRC, or pursuant to an exemption available thereunder, or assume any responsibility for facilitating any such distribution or offering. In particular, no action has been taken by us which would permit a public offering of any securities or distribution of this document in the PRC. Accordingly, the securities are not being offered or sold within the PRC by means of this document or any other document. Neither this document nor any advertisement or other offering material may be distributed or published in the PRC, except under circumstances that will result in compliance with any applicable laws and regulations.

**To our readers in Japan**: This publication is being distributed in Japan by Sanford C. Bernstein Japan KK (サンフォード・C・バーンスタイン株式会社), which is registered in Japan as a Financial Instruments Business Operator with the Kanto Local Finance Bureau (registration number: The Director-General of Kanto Local Finance Bureau (FIBO) No.3387) and regulated by the Financial Services Agency. It is also a member of Japan Investment Advisers Association. This publication is solely for qualified institutional investors in Japan only, as defined in Article 2, paragraph (3), items (i) of the Financial Instruments and Exchange Act.

For the institutional client readers in Japan who have been granted access to the Bernstein website by Daiwa Securities Group Inc. ("Daiwa"), your access to this document should not be construed as meaning that Bernstein is providing you with investment advice for any purposes. Whilst Bernstein has prepared this document, your relationship is, and will remain with, Daiwa, and Bernstein has neither any contractual relationship with you nor any obligations towards you.

**To our readers in Australia:** Sanford C. Bernstein (Hong Kong) Limited 盛博香港有限公司 is responsible for distributing research in Australia. It is regulated by the Securities and Exchange Commission under U.S. laws, by the Financial Conduct Authority under U.K. laws, which differs from Australian laws. Sanford C. Bernstein (Hong Kong) Limited 盛博香港有限公司 is exempt from the requirement to hold an Australian financial services license under the Corporations Act 2001 in respect of the provision of the following financial services to wholesale clients:

- providing financial product advice;
- dealing in a financial product;
- making a market for a financial product; and
- providing a custodial or depository service.

**To our readers in India:** This publication is being distributed in India by Sanford C. Bernstein (India) Private Limited (SCB India) which is licensed and regulated by Securities and Exchange Board of India ("SEBI") as a research analyst entity under the SEBI (Research Analyst) Regulations, 2014, having registration no. INH000006378 and as a stock broker having registration no. INZ000213537. SCB India is currently engaged in the business of providing research and stock broking services. Please refer to www.bernsteinresearch.in for more information.

- SCB India is a Private limited company incorporated under the Companies Act, 2013, on April 12, 2017 bearing corporate identification number U65999MH2017FTC293762, and registered office at Level 3A, 4th Floor, First International Financial Centre, Plot Nos C-54 and C-55, G Block, Near CBI Office, Bandra Kurla Complex, Bandra (East), Mumbai 400098, Maharashtra, India (Phone No: +91-22-68421401).
- For details of Associates (i.e., affiliates/group companies) of SCB India, kindly email MUM-BERNSTEIN-InCompliance@bernsteinsg.com.
- SCB India does not have any disciplinary history as on the date of this report.
- Except as noted above, SCB India and/or its Associates (i.e., affiliates/group companies), the Research Analysts authoring this report, and their relatives
  - do not have any financial interest in the subject company

- do not have actual/beneficial ownership of one percent or more in securities of the subject company;
- is not engaged in any investment banking activities for Indian companies, as such;
- have not managed or co-managed a public offering in the past twelve months for any Indian companies;
- have not received any compensation for investment banking services or merchant banking services from the subject company in the past 12 months;
- have not received compensation for brokerage services from the subject company in the past twelve months;
- have not received any compensation or other benefits from the subject company or third party related to the specific recommendations or views in this report; and
- do not currently, but may in the future, act as a market maker in the financial instruments of the companies covered in the report.
- do not have any conflict of interest in the subject company as of the date of this report.

- Except as noted above, the subject company has not been a client of SCB India during twelve months preceding the date of distribution of this research report. Neither SCB India nor its Associates (i.e., affiliates/group companies) have received compensation for products or services other than investment banking, merchant banking or brokerage services from the subject company in the past twelve months.
- The principal research analyst(s) who prepared this report, members of the analysts' team, and members of their households are not an officer, director, employee or advisory board member of the companies covered in the report.
- Our Compliance officer / Grievance officer is Ms. Rupal Talati, who can be reached at +91-22-68421451, or MUM-BERNSTEIN-InCompliance@bernsteinsg.com / Scbin-investorgrievance@bernsteinsg.com
- The Research investor charter and Terms & Conditions of SCB India are available on its website and may be accessed at Sanford C. Bernstein (India) Private Limited (https://bernsteinresearch.in/) for your reference.
- Disclaimer: Registration granted by SEBI, and certification from NISM, is in no way a guarantee of performance of the intermediary or provide any assurance of returns to investors. Investments in securities market are subject to market risks. Read all the related documents carefully before investing.

**To our readers in Switzerland:** This document is provided in Switzerland by or through Bernstein Autonomous LLP, and is provided only to qualified investors as defined in article 10 of the Swiss Collective Investment Scheme Act ("CISA") and related provisions of the Collective Investment Scheme Ordinance and in strict compliance with applicable Swiss law and regulations. The products mentioned in this document may not be suitable for all types of investors. This document is based on the Directives on the Independence of Financial Research issued by the Swiss Bankers Association (SBA) in January 2008.

**To our readers in the Middle East:** Bernstein Autonomous LLP, DIFC branch has its principal office at Gate Village 06, DIFC, Dubai, UAE. Bernstein Autonomous LLP, DIFC branch is regulated by the Dubai Financial Services Authority (DFSA) with the registration number CL10040 and is provisioned for Arranging Deals in Investments and Advising on Financial Products. All communications and services are directed at Professional Clients and Market Counterparties only (as defined in the DFSA rulebook). Persons other than Professional Clients and Market Counterparties, such as Retail Clients, are not the intended recipients of our communications or services.

**LEGAL**

All research publications are disseminated to our clients through posting on the firm's password protected websites, bernsteinresearch.com and autonomous.com. Certain, but not all, research publications are also made available to clients through third-party vendors or redistributed to clients through alternate electronic means as a convenience.

This publication has been published and distributed in accordance with the Firm's policy for management of conflicts of interest in investment research, a copy of which is available from Bernstein Institutional Services LLC, Director of Compliance, 245 Park Avenue, New York, NY 10167. Additional disclosures and information regarding Bernstein's business are available on our website

www.bernsteinresearch.com.

The securities described herein may not be eligible for sale in all jurisdictions or to certain categories of investors where that permission profile is not consistent with the licenses held by the entities noted herein. This document is for distribution only as may be permitted by law. This publication is not directed to, or intended for distribution to or use by, any person or entity who is a citizen or resident of, or located in any locality, state, country or other jurisdiction where such distribution, publication, availability or use would be contrary to law or regulation or which would subject any of the entities referenced herein or any of their subsidiaries or affiliates to any registration or licensing requirement within such jurisdiction. This publication is based upon public sources we believe to be reliable, but no representation is made by us that the publication is accurate or complete. We do not undertake to advise you of any change in the reported information or in the opinions herein. This publication was prepared and issued by entity referred to herein for distribution to eligible counterparties or professional clients. This publication is not an offer to buy or sell any security, and it does not constitute investment, legal or tax advice. The investments referred to herein may not be suitable for you. Investors must make their own investment decisions in consultation with their professional advisors in light of their specific circumstances. The value of investments may fluctuate, and investments that are denominated in foreign currencies may fluctuate in value as a result of exposure to exchange rate movements. Information about past performance of an investment is not necessarily a guide to, indicator of, or assurance of, future performance.

This report is directed to and intended only for our clients who are "eligible counterparties", "professional clients", "institutional investors" and/or "professional investors" as defined by the aforementioned regulators, and must not be redistributed to retail clients as defined by the aforementioned regulators. Retail clients who receive this report should note that the services of the entities noted herein are not available to them and should not rely on the material herein to make an investment decision. The result of such act will not hold the entities noted herein liable for any loss thus incurred as the entities noted herein are not registered/ authorised/ licensed to deal with retail clients and will not enter into any contractual agreement/arrangement with retail clients. This report is provided subject to the terms and conditions of any agreement that the clients may have entered into with the entities noted herein. All research reports are disseminated on a simultaneous basis to eligible clients through electronic publication to our client portal.

The information in this report was prepared by Bernstein solely for the internal business use of our clients. Clients may store, display, analyze, reformat and print the information in this report for this limited use only. Clients may not copy, alter, create derivative works, resell, reverse engineer, commercially exploit, share or distribute any part of the information contained herein for any purpose without Bernstein's express written consent. These restrictions include extracting data or using the content to develop indices or other products. Further, you may not use this report, or any portion of this report, to train or finetune any third-party machine learning or artificial intelligence system, or as a prompt or input into any such system. You also may not, without Bernstein's express written consent, do any of the foregoing in connection with your own internal machine learning or artificial intelligence system.

Bernstein may use artificial intelligence tools in the preparation of its materials. Any such materials are reviewed by Bernstein's research analysts prior to publication.

This report has been prepared for information purposes only and is based on current public information that we consider reliable, but the entities noted herein do not warrant or represent (express or implied) as to the sources of information or data contained herein are accurate, complete, not misleading or as to its fitness for the purpose intended even though the entities noted herein rely on reputable or trustworthy data providers, it should not be relied upon as such. Opinions expressed are the author(s)' current opinions as of the date appearing on the material only and we do not undertake to advise you of any change in the reported information or in the opinions herein.

This publication was prepared and issued by the entity referred to herein for distribution to eligible counterparties or professional clients. The information in this report is intended for general circulation and does not constitute an offer to buy or sell any security, investment, legal or tax advice nor a personal recommendation, as defined by any of the aforementioned regulators. It does not take into account the particular investment objectives, financial situations, or needs of individual investors. The report has not been reviewed by any of the aforementioned regulators and does not represent any official recommendation from the aforementioned regulators. The investments referred to herein may not be suitable for you. Investors must make their own investment decisions in consultation with advice sought from a financial adviser regarding the suitability of the investment product, taking into account the specific investment objectives, financial situation or particular needs of any recipient of the recommendation, before the recipient makes a commitment to purchase the investment product.

The analysis contained herein is based on numerous assumptions. Different assumptions could result in materially different results. The information in this report does not constitute, or form part of, any offer to sell or issue, or any offer to purchase or subscribe for shares, or to induce engage in any other investment activity. The value of any securities or financial instruments mentioned in this report may fluctuate subject to market conditions. Information about past performance of an investment is not necessarily a guide to, indicative of, or assurance of future performance. Estimates of future performance mentioned by the research analyst in this report are based on assumptions that may not be realized due to unforeseen factors like market volatility/fluctuation. In relation to securities or financial instruments denominated in a foreign currency other than the clients'

home currency, movements in exchange rates will have an effect on the value, either favorable or unfavorable. Before acting on any recommendations in this report, recipients should consider the appropriateness of investing in the subject securities or financial instruments mentioned in this report and, if necessary, seek for independent professional advice.

The securities described herein may not be eligible for sale in all jurisdictions or to certain categories of investors where that permission profile is not consistent with the licenses held by the entities noted herein. This document is for distribution only as may be permitted by law. It is not directed to, or intended for distribution to or use by, any person or entity who is a citizen or resident of or located in any locality, state, country or other jurisdiction where such distribution, publication, availability or use would be contrary to law or regulation or would subject the entities noted herein to any regulation or licensing requirement within such jurisdiction.

Source: Bloomberg Index Services Limited. BLOOMBERG® is a trademark and service mark of Bloomberg Finance L.P. and its affiliates (collectively "Bloomberg"). Bloomberg or Bloomberg's licensors own all proprietary rights in the Bloomberg Indices. Neither Bloomberg nor Bloomberg's licensors approves or endorses this material, or guarantees the accuracy or completeness of any information herein, or makes any warranty, express or implied, as to the results to be obtained therefrom and, to the maximum extent allowed by law, neither shall have any liability or responsibility for injury or damages arising in connection therewith.