BERNSTEIN
SOCIETE GENERALE GROUP

3 July 2026

China Semiconductors

China CPU: Hygon - the key China beneficiary for the CPU renaissance

Qingyuan Lin, Ph.D.
+852 2123 2654
qingyuan.lin@bernsteinsg.com

Francis Ma
+852 2123 2626
francis.ma@bernsteinsg.com

Kai Zhang
+852 2123 2665
kai.zhang@bernsteinsg.com

Agentic AI is leading to a new ear that we call "The CPU Renaissance". Within China, we like Hygon's unique position as the only alternative x86 server CPU vendor to Intel/AMD, allowing the company to gain share in China for the next five years. We discuss the market opportunity and tech comparison for Hygon in this report, our China x86 server CPU model can be downloaded here. **Outperform on Hygon, PT CNY 450.**

Agentic AI is driving the expansion of global server CPU TAM, creating a much larger opportunity for Hygon than investors currently appreciate. We estimate global server CPU TAM will expand from roughly US$39bn in 2025 to US$223bn by 2030, as agentic AI shifts data-center architectures from being GPU-centric to increasingly CPU-intensive. We expect the CPU:GPU ratio to gradually move from 1:8 now to 1:1 in 2030. Within China, x86 server CPU TAM is expected to increase from US$7bn in 2025 to US$27bn by 2030. While growth may lag global markets from 2025-28 due to supply constraints, we expect a meaningful reacceleration from 2028 onward, as domestic supply capacity expands and AIDC investment accelerates.

Accelerating localization and tightening supply of foreign CPUs should drive substantial share gains for Hygon. China's first wave of CPU localization was driven by Xinchuang policies, where Hygon already enjoys a unique position as the only domestic x86 server CPU provider. We believe a second, more commercially driven localization wave will emerge from 2027 onward as Chinese CSPs increasingly adopt domestic CPU. Rapid AI infrastructure expansion is tightening AMD and Intel CPU supply globally, with available capacity increasingly allocated to large U.S. hyperscalers. As global vendors' supply availability in China becomes less predictable, CSPs will have stronger incentives to qualify domestic alternatives, just like what happened to the AI chips. Hygon also enjoys a unique CPU+AI chip ecosystem that they can bundle both of them to improve the sales of both chips. We forecast Hygon's market share by value in China x86 server CPU market to rise from 19% in 2025 to 36% by 2030.

Hygon's CPU performance is only 2-3 years behind Intel/AMD, and was continuously improving even without AMD support. Despite losing access to further AMD technology support after its 2019 Entity List designation, Hygon has continued to advance its x86 CPU roadmap independently. We estimate the technology gap versus global x86 leaders has been maintained at 2-3 years over the past few years. Even with that gap, its products remain sufficiently competitive for a growing range of workloads, able to address roughly half of China market. Looking ahead, we expect C86-G5 to deliver a meaningful performance upgrade, further expanding Hygon's addressable market beyond its traditional Xinchuang customer base and supporting broader adoption among CSP and government supported AIDCs.

We remain Outperform on Hygon with a CNY450 PT. As the only credible domestic x86 server CPU supplier in China, Hygon is uniquely positioned to benefit from both a rapidly expanding CPU TAM and accelerating localization.

First Published: 02 Jul 2026 21:44 UTC Completion Date: 02 Jul 2026 17:22 UTC

BERNSTEIN TICKER TABLE

Ticker	Rating	Cur	2 Jul 2026 Closing Price	Price Target	TTM Rel. Perf.	Reported EPS Cur	2025A	2026E	2027E	Reported P/E (x) 2025A	2026E	2027E
688041.CH (Hygon)	O	CNY	322.29	450.00	99.6%	CNY	1.10	2.04	3.59	293.0	157.7	89.7
ASIAX			1,987.28									

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

Source: Bloomberg, Bernstein estimates and analysis.

INVESTMENT IMPLICATIONS

We rate **Hygon Outperform, PT CNY 450**. As the only credible domestic x86 server CPU vendor, Hygon is well positioned to benefit from China's accelerating AI infrastructure buildout and CPU localization trend. We expect agentic AI to drive an upward revision to China's CPU TAM, with Hygon as a key beneficiary. Hygon's tech gap vs Global leaders has been maintained at 2-3 years, allowing the company to address about half of China's TAM. As the upcoming C86-G5 narrows the performance gap with global peers, we expect Hygon's addressable market to further expand beyond Xinchuang into the much larger CSP segment, supporting sustained market share gains in China over the coming years.

EXHIBIT 1: **Hygon has broadly maintained a stable performance gap versus global peers at 2-3 years over the past few years. We expect C86-G5 to achieve a meaningful performance upgrade, further expanding Hygon's SAM penetration within the TAM from the current ~50% level**

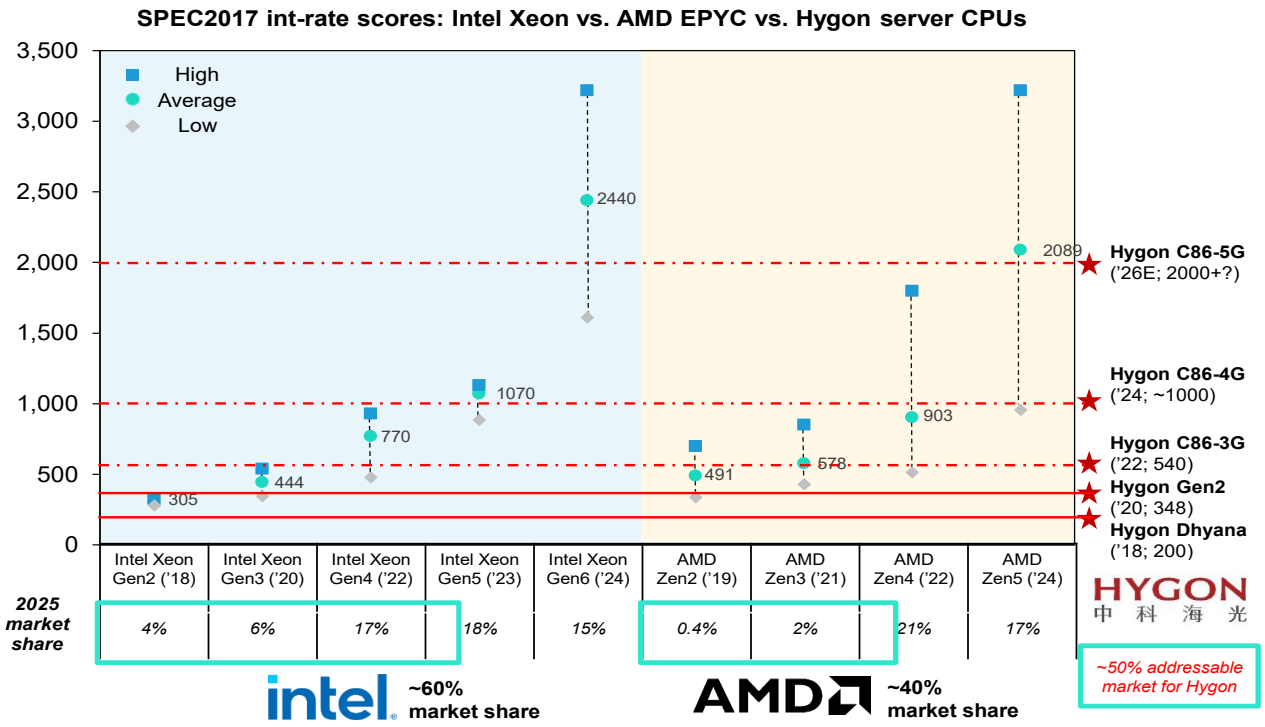

1. Hygon score is estimated base on it's flagship 7000 series. Gen 1 and 2 data are reported by the company and rest are estimation through channel checks.

2. 2025 market share is based on Mercury's analysis on global Intel/AMD sales, excluding Hygon.
Source: Companies disclosure, Mercury, Bernstein analysis and estimates

Table Of Contents

DETAILS

GLOBAL AND CHINA X86 CPU TAM

THE GLOBAL CPU RENAISSANCE BOOSTED BY AGENTIC AI

The global server CPU market is undergoing a structural expansion qualitatively different from prior upgrade cycles. We estimate the global server CPU TAM will expand from $39Bn in 2025 to $223Bn by 2030, driven primarily by the emergence of Agentic AI as a fundamentally CPU-intensive compute paradigm. We have identified this trend as the "CPU Renaissance", mainly driven by Agentic AI.

EXHIBIT 2: **We estimate the global server CPU TAM will reach $223Bn by 2030, with ARM's share surpassing x86 from 2027 and rising to 55% by 2030**

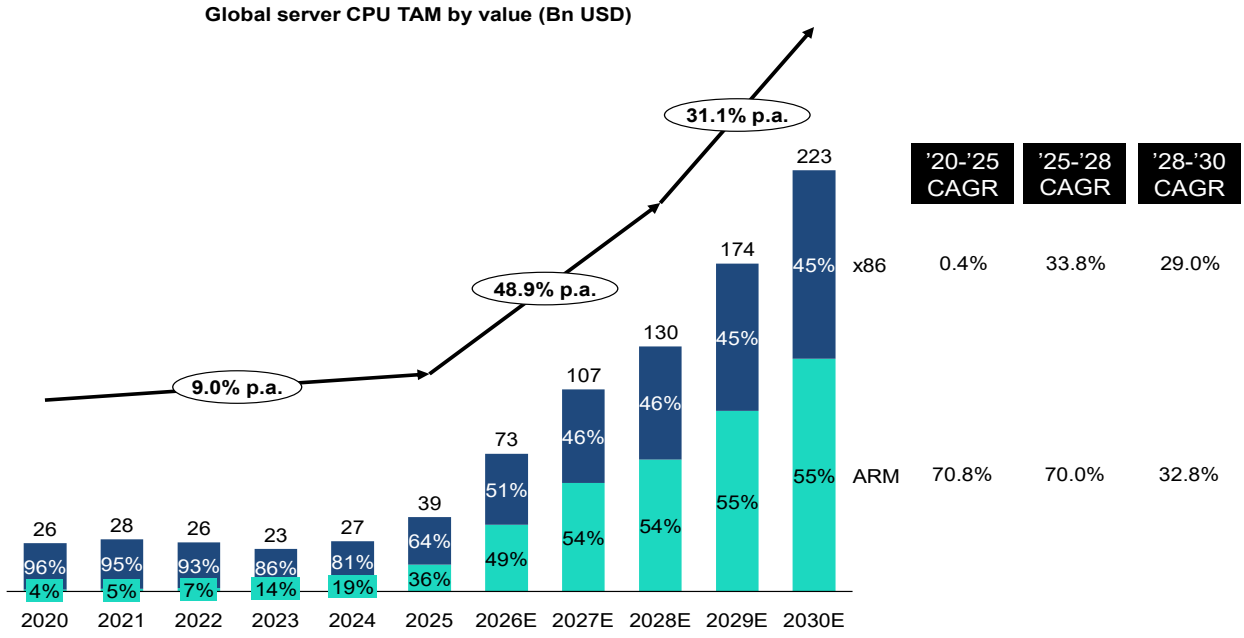

Source: Company disclosures, Mercury, Bernstein estimates and analysis.

EXHIBIT 3: **We estimate global server CPU shipments will reach 95 million units by 2030. ARM's unit share is expected to be ~20%, reflecting its significantly higher ASPs**

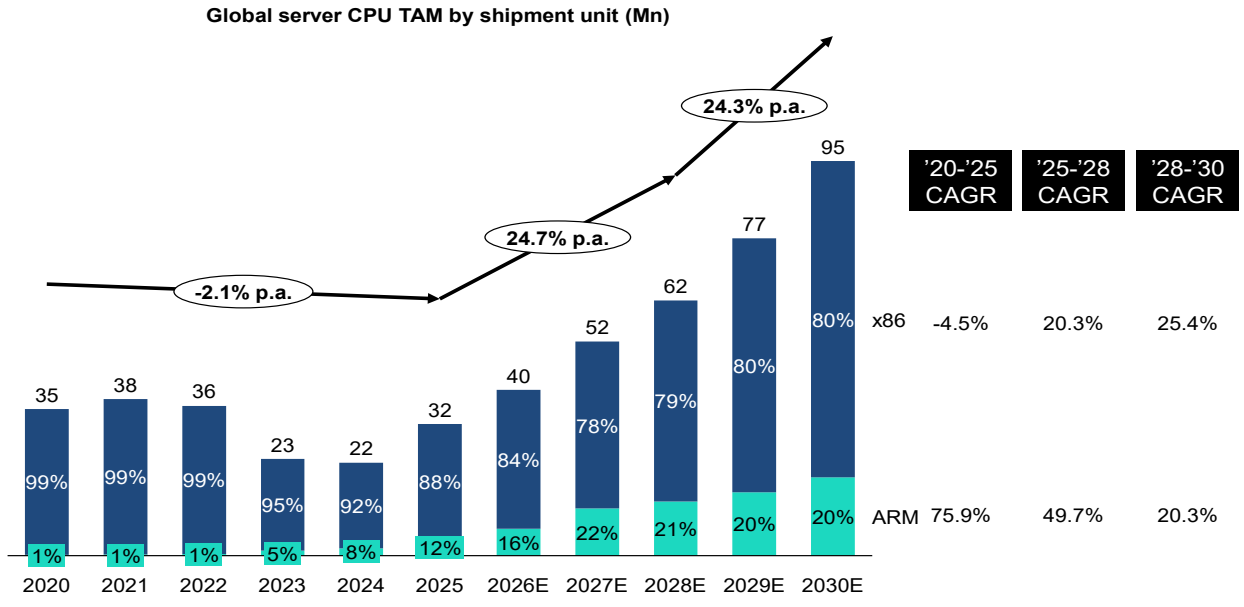

Source: Company disclosures, Mercury, Bernstein estimates and analysis.

Traditional LLM workloads and Agentic AI workloads have different demands from data center hardware. Traditional LLMs are GPU-centric. Training and batch inference both map cleanly onto massively parallel matrix operations that GPUs excel at. During the 2021-25 period of LLM scale-up, the GPU's role dominated, while the CPU was a supporting actor. This was reflected in CSP cluster configurations: the GPU-to-CPU ratio in CSP training clusters rose from roughly 3:1 in 2020 to around

8:1 by 2024-25. During this period, the CPU's share in AI CAPEX was approximately 6.2%.

Agentic AI inverts this dynamic. Agentic AI involves systems of autonomous AI agents that orchestrate multistep workflows, call external tools and APIs, retrieve and synthesize information in real-time, manage persistent memory, and coordinate with other agents to complete complex tasks. The GPU's job in this architecture remains token generation, but the overwhelming majority of everything else is CPU-bound. Agent orchestration, tool execution, context window management, multi-agent coordination, and human-in-loop interfaces all require general-purpose sequential processing with low latency and high memory bandwidth. As Agentic AI scales, the demand for CPU capacity per AI cluster rises sharply relative to the LLM era.

The quantitative evidence for this shift is already in the data. The GPU-to-CPU ratio, which peaked at ~8:1, will reverse to about 2:1 by 2030.

Meanwhile, the secular buildout of hyperscale cloud infrastructure provides a durable second layer of demand, further reinforced by the rise of Agentic AI. In our base case, we forecast cumulative AI capex of $3.5 trillion from major cloud providers, with the potential to reach $4.0 trillion under more bullish scenarios. Cloud buildout and Agentic AI are therefore compounding forces on CPU demand. Accordingly, we raise our 2030 global server CPU TAM estimate to $223 Bn (vs. $137 Bn previously), implying a sixfold expansion from the $37 Bn TAM in 2025.

EXHIBIT 4: **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**

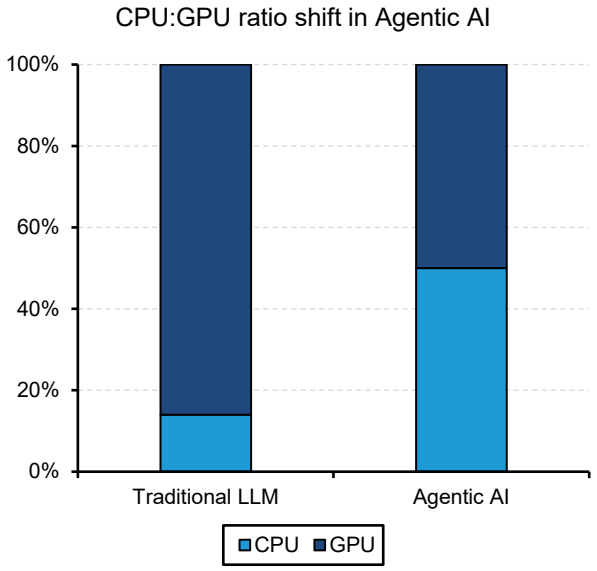

Source: TrendForce, Bernstein analysis

EXHIBIT 5: **CPU is expected to play a more important role within inference, in the agentic area.**

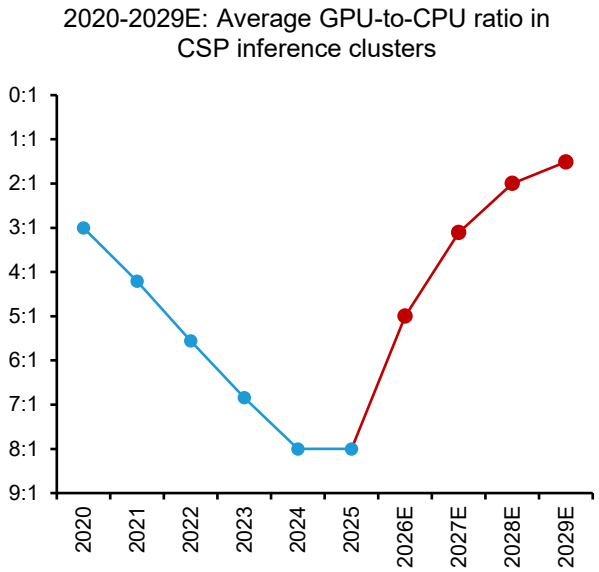

Source: Ciena estimates, Bernstein analysis.

EXHIBIT 6: **We believe 2030 server CPU TAM will be $223bn in the base case of $3.5tn AI capex.**

CY	2025	2030 Base	2030 Bull	2030 Bear
AI GW additions	15	70	80	60
AI capex intensity (bn / GW)	40	50	50	50
AI capex ($ bn)	600	3,500	4,000	3,000
AI GPU/accelerator TAM ($ bn)	240	1,575	1,800	1,350
Inference ratio	35%	70%	70%	70%
AI GPU for inferencing ($ bn)	84	1,103	1,260	945
CPU:GPU ratio (inference)	0.25x	1.0x	1.5x	0.5x
AI GPU for training ($ bn)	156	473	540	405
CPU:GPU ratio (training)	0.25x	0.50x	0.50x	0.50x
CPU:GPU cost ratio	10.0%	13.0%	13.0%	13.0%
CPU for agentic AI ($ bn)	6.0	174.0	280.8	87.8
CPU for general CPU server ($ bn)	31.3	49.4	49.4	49.4
CPU TAM ($ bn)	**37.3**	**223.4**	**330.2**	**137.2**

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

ARM VS. X86: RISING COMPETITION, BUT THE X86 MOAT ENDURES

ARM-based designs are gaining share in the server CPU market, particularly on the ASP-weighted value basis. We expect ARM's share of global server CPU value to rise from 36% in 2025 to 55% in 2030. This is concentrated in hyperscalers custom silicon and NVIDIA CPU where ARM's performance-per-watt advantages are most pronounced.

However, on a unit shipment basis, we anticipate x86 will retain 80% share even by 2030. The software ecosystem moat is the structural moat. Decades of enterprise software, including databases, middleware, ERP systems, security tools, and the full enterprise application stack, have been built, optimized, and certified on x86. Migrating to ARM requires recompilation and recertification cycles that can span years in regulated industries. Especially in China, we expect x86 server CPUs to remain the prevailing server architecture, as export restrictions limit access to NVIDIA's ARM-based CPUs, while U.S. CSP custom silicon remains largely unavailable to the market.

At their analyst day in November, AMD laid out expectations for a ~$60B CPU TAM by 2030 on the rise of agentic AI. Clearly the pace and trajectory of the ramp surprised them as well, as AMD doubled the number in May on their Q126 earnings call to $120B. Our US semi analyst, Stacy Rasgon, found Intel also benefited from the demand environment, that was so strong in fact that they managed to sell previously-written-off parts they had (presumably) trashed, but which still found homes as customers proved willing to buy up anything they could get their hands on.

ACCELERATING LOCALIZATION IN CHINA

CHINA X86 SERVER CPU TAM EXPANSION - A TALE OF TWO GROWTH PHASES

China's x86 server CPU total addressable market (TAM) is projected to grow significantly from USD 7 billion in 2025 to USD 27 billion by 2030. Despite the rapid expansion of the global market, China is expected to maintain a relatively stable share of approximately 25% to 27%.

In terms of shipment volume, China's server CPU demand is forecast to triple over the next five years, reaching approximately 24.6 million units by 2030. This growth reflects both the expansion of cloud infrastructure and the increasing adoption of AI-related workloads.

Although China remains a key growth driver in absolute terms, its CAGR during 2025–2028 is expected to lag behind the global market. This divergence is primarily due to structural constraints on AI chip that put a constraint on the demand growth during this period.

However, beginning in 2028, we expect China's market growth to reaccelerate and outpace global growth, as these constraints are gradually alleviated. On the demand side, we also expect that the government supported AIDC projects will accelerate in spending which also bring upside to the CPU demand growth. At the same time we expect China AI chip and CPU supply will ramp up very quickly with the accelerating capacity expansion for advanced logic.

EXHIBIT 7: **We project China's x86 server CPU TAM to grow to $27Bn by 2030. Despite a temporary decline in global share due to supply constraints, China's contribution will rebound to ~27% by decade-end**

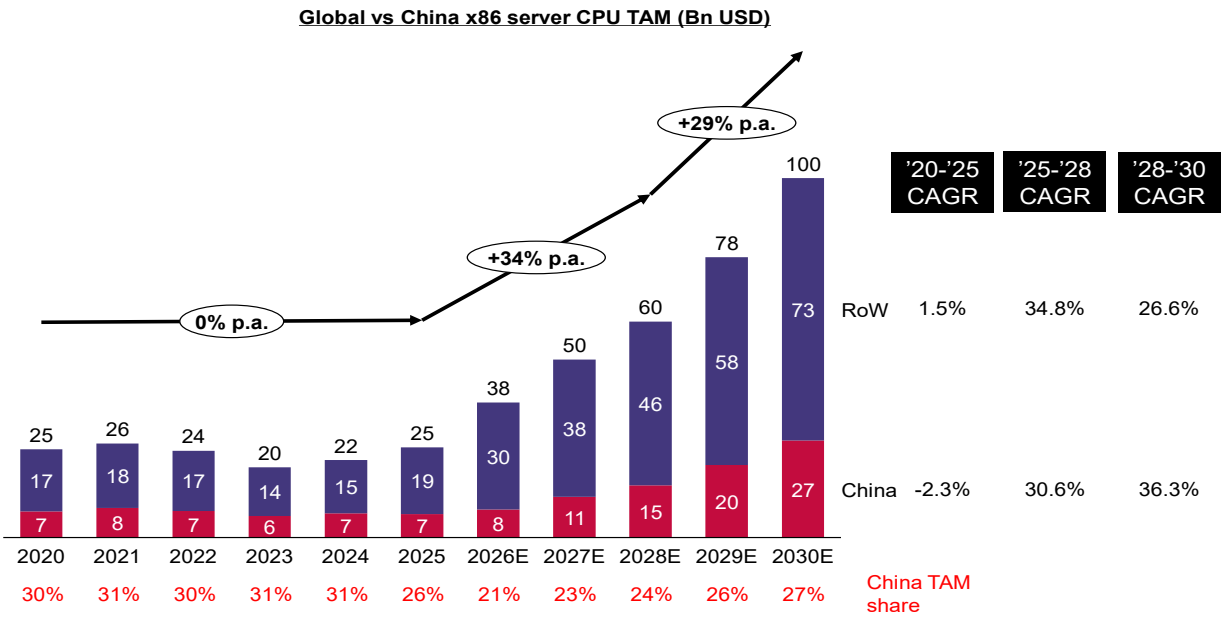

Source: Company disclosures, Mercury, Bernstein analysis and estimates

EXHIBIT 8: **We anticipate China's x86 server CPU shipments to reach 25 Mn units by 2030. China's share of global unit shipments will exceed its share of TAM value due to lower ASPs relative to international markets**

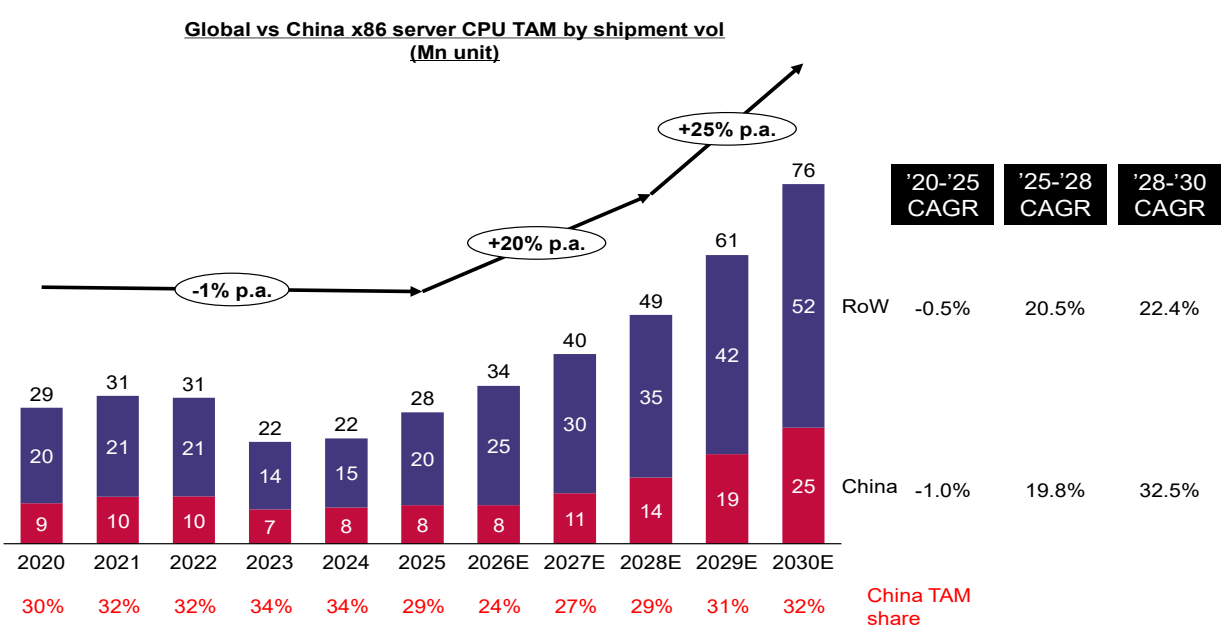

Source: Company disclosures, Mercury, Bernstein estimates and analysis.

Structural supply constraints in the near term (2023-2028)

China's server CPU market has been, and is expected to remain, constrained through the 2023–2028 period, with growth rate lagging behind the rest of the world. Following a period of modest contraction over 2020-25, primarily driven by slower pace in Xinchuang procurement post COVID, China's x86 server CPU market is now recovering. We estimate China's x86 CPU TAM

will grow at roughly 30% CAGR during 2025 to 2028, a healthy absolute growth that reflects the resumption of Xinchuang fulfillment and the build-out of AI infrastructure. However, this recovery still lags behind the RoW (rest of the world, particularly the US) at 35% CAGR, where CSPs are building Agentic AI infrastructure in more aggressive pace. Near-term constraints in China are driven by both global supply dynamics and domestic semiconductor fabrication limitations, which together restrict China's ability to fully meet underlying demand.

- **Global supply allocation to China decreases as AI-driven demand surge in the US**

The rapid buildout of Agentic AI infrastructure among US hyperscalers has created an extraordinary surge in demand for AMD and Intel server CPUs, and supply has not kept pace. AWS, Google, Microsoft, and Meta are expanding their data center footprints at unprecedented speed, absorbing the majority of available AMD EPYC and Intel Xeon production. As a result, AMD and Intel are increasingly allocating their CPU supply toward US and global hyperscaler customers, where demand is stronger and margins are more attractive. Consequently, they de-prioritize China in their delivery queues, leaving China's x86 CPU market growing more slowly than the global average through 2028.

- **Domestic advanced-node capacity constraints, with cascading effects on both AI chips and domestic CPU**

A key bottleneck in China's market lies in limited access to advanced-node fabrication capacity. On one side, the shortage of AI accelerators in China is itself a constraint on server CPU demand: AI racks demand both GPUs and CPUs to function, and with access to NVIDIA's advanced GPUs restricted and domestic AI accelerators still ramping, China's CSPs cannot deploy AI infrastructure at the pace that demand warrants. Server deployments become supply-constrained rather than lack of demand. Hence, the CPU demand is suppressed. On the other side, Hygon's own ability to scale production is bounded by the domestic foundry bottleneck. SMIC's advanced-node capacity is finite and shared across multiple priority programs. This limits Hygon's wafer allocation, caps shipment volumes, and prevents it from fully compensating for the Intel and AMD supply that shifts from China market. Both effects of foundry capacity shortage result in a growth rate for China's x86 CPU that remains below the global pace through 2028.

Reacceleration from 2028 onward

Starting from 2028, we expect a meaningful reacceleration in China's server CPU market, with CAGR increasing to 36.3% during 2028–2030 and outpacing the global growth. This inflection is driven by both supply-side improvements and demand recovery.

- **Domestic advanced-node fabrication capacity ramp up**

The ramp-up of domestic advanced-node fabrication capacity will play a critical role in removing supply bottlenecks. We expect SMIC and other local foundries' advanced-node capacity to scale meaningfully from 2028 onward. This will remove supply ceiling that has suppressed China's x86 market in the near term. Meanwhile, the ramp in domestic leading fabrication capacity also directly supports performance improvements in Hygon's products, as advanced-node availability translates into higher clock frequencies, lower power consumption, and greater core density. Huawei's Tau scaling law is something Hygon can learn from to further improve the CPU performance despite constrain on local manufacturing due to the block on EUV access. All of these improvements can enable Hygon CPU to expand the addressable market.

The scaling of domestic advanced-node fabrication will materially increase the available supply of domestic AI accelerators. Today, the single most binding constraint on China's AI infrastructure buildout is the chip supply, instead of demand or budget. As China's advanced logic capacity scales, this bottleneck eases, enabling CSPs to build out AI server capacity that has been backlogged. The ramp in domestic AI chips supply will translate directly into incremental pull-through of CPU demand.

- **Domestic AIDC investment catch-up**

CSP driven AI infrastructure buildout accelerate domestic CPU demand. As constraints on AI accelerator supply ease, China's CSPs will be able to ramp up AI infrastructure deployment, catching up with global peers. This will directly translate into stronger CPU demand, as server build-outs resume. We believe beyond traditional CSP, local government supported neocloud should also start to play an increasingly important role for the AIDC capex. Assume that by 2028, the vertical AI models are already demonstrating their commercialization potentials, the demand on compute should far exceed what the CSPs can offer, and local government should chime in to build infrastructure to support these demand, just like how China build the electricity infrastructure in the past.

In addition, China may exhibit a structural advantage in inference workloads. Unlike training workloads, which prioritize leading-edge performance and therefore depend on the most advanced chips, inference workloads are more cost-sensitive and scalable. China's ecosystem is well-positioned to deliver lower-cost solutions with acceptable performance, making them competitive for inference deployment. The advantage mainly comes from two angles, lower electricity cost and higher token output per compute power with LLM infra optimization as demonstrated by DeepSeek etc. The broader adoption of inference use cases in China has potential to result in a higher CPU-to-GPU ratio in system configurations vs. in the US. Such a shift would further amplify CPU demand in China, representing a potential upside that is not currently incorporated in our base-case forecasts.

HYGON'S SHARE GAIN THROUGH LOCALIZATION - XINCHUANG NOW AND CSP NEXT

Alongside the absolute TAM expansion, the domestic CPU localization rate in China has been rising continuously since 2020, and we expect this trend to persist and even accelerate through 2030. Moreover, the primary driver of localization is changing: the first wave has been policy-led, anchored in Xinchuang; the second wave will be commercially led, driven by CSP procurement economics and AI infrastructure priorities. The two waves creates a compounding upward trajectory for localization and Hygon's share.

- **Xinchuang as the policy engine since 2020**

The dominant force behind China's CPU localization to date has been Xinchuang（信创）, China's national Information Technology Application Innovation policy mandate. Xinchuang requires government agencies, state-owned enterprises (SoE), and eight key industries — energy and electricity, transportation, telecommunications, finance, education, healthcare, petroleum, and aerospace — to progressively replace foreign IT hardware and software with certified domestic alternatives. The policy is structured as a "2+8+N" rollout: Party and Government entities (the "2") were targeted for near-complete localization by 2022; the eight key industries (the "8") are targeted through 2025-2027; and broader commercial users (the "N") follows thereafter.

Under Xinchuang, Hygon is the only available x86 server CPU vendor at present. This has created a demand floor that was largely insensitive to Hygon's absolute performance relative to Intel and AMD. Especially, purchasing decision of Party & Government, and SoE is driven by compliance, more than benchmarks.

On a value basis, Exhibit 9 shows that the CPU localization rate/ Hygon's market share has risen consistently from just 2% in 2020 to 19% in 2025. The pace of localization has moderated since 2023 relative to prior years, reflecting the temporary softness in Party & Government and SoE procurement in the post-COVID period. The natural question is whether this localization trend will continue to accelerate. We believe it will — and that a second, potentially more powerful wave is on the horizon.

- **CSP adoption accelerates localization from 2027 onward**

We expect a second wave of CPU localization to emerge from 2027 onward as China's CSPs begin adopting Hygon processors at scale. Unlike the first wave, which was largely policy-driven, this phase is expected to be driven primarily by commercial considerations. We believe a structurally tightening supply of AMD and Intel CPUs, combined with meaningful improvements in Hygon's product competitiveness, will be the key catalysts behind this inflection.

The supply outlook for AMD and Intel CPUs in China is becoming increasingly constrained. Rapid growth in AI infrastructure spending, particularly for Agentic AI deployment, by U.S. hyperscalers has significantly increased demand for AMD EPYC and Intel Xeon processors. Competing for the tight leading-edge foundry capacity with other fabless giants, AMD and Intel face growing supply constraints as demand accelerates. Available capacity is increasingly being allocated to their largest and most strategically important customers, particularly U.S. CSPs. Hyperscalers competing aggressively in AI infrastructure are even willing to purchase Intel inventory that had previously been written off, proving the supply tightness. For Chinese customers, the result is a combination of reduced availability and rising procurement costs. In our view, this creates a strong push factor for Hygon adoption, as reliance on foreign CPUs becomes progressively more costly and less predictable.

Hygon's performance and supply capabilities are improving materially. On the performance front, we believe the latest 4th generation Hygon is selling (sales starting from 2024) is already comparable to Intel's 4th- and 5th-Gen Xeon platforms (sales starting from 2022) and AMD's Zen 3 architecture (sales starting from 2021), making it suitable for a meaningful subset of CSP workloads where total core count, memory bandwidth, and total cost of ownership (TCO) matter more than peak single-thread performance. Looking ahead, we expect the upcoming C86-G5 platform to narrow the remaining performance gap with the currently available offerings from AMD and Intel. Meanwhile, China's domestic advanced-node manufacturing capacity continues to expand, providing a more secure and scalable production base for Hygon CPUs.

Unique CPU+GPGPU ecosystem also allow Hygon to bundle more CPU when they are selling GPGPU. The AI chip supply shortage in China means that all local AI chip vendors got massive demand from local AIDC expansion. Hygon is the only player beyond Huawei that has both CPU and AI chip, therefore this allows Hygon to either bundle the AI chip with CPU when the customers are in CPU shortage, or bundle the CPU with AI chip when they are selling AI chip during supply shortage.

Importantly, these trends reinforce one another. Tightening supply from AMD and Intel increases the urgency for Chinese CSPs to qualify and deploy domestic alternatives, while Hygon's improving performance and growing supply readiness reduce the switching costs and execution risks associated with such a transition. As a result, we project Hygon's penetration within the CSP and other AIDCs market to increase from an MSD share today to approximately 20% by 2030. Including demand from other downstream verticals, we forecast Hygon's CPU market share by value to reach 36% in 2030, implying a 40–50% CAGR over the coming years. We estimate its market share by shipment volume can even approach 50% by 2030, given Hygon's lower ASP relative to global peers.

EXHIBIT 9: **We expect a rapid increase in China's x86 server CPU self-sufficiency, with domestic substitution accelerating notably from 2028 onward**

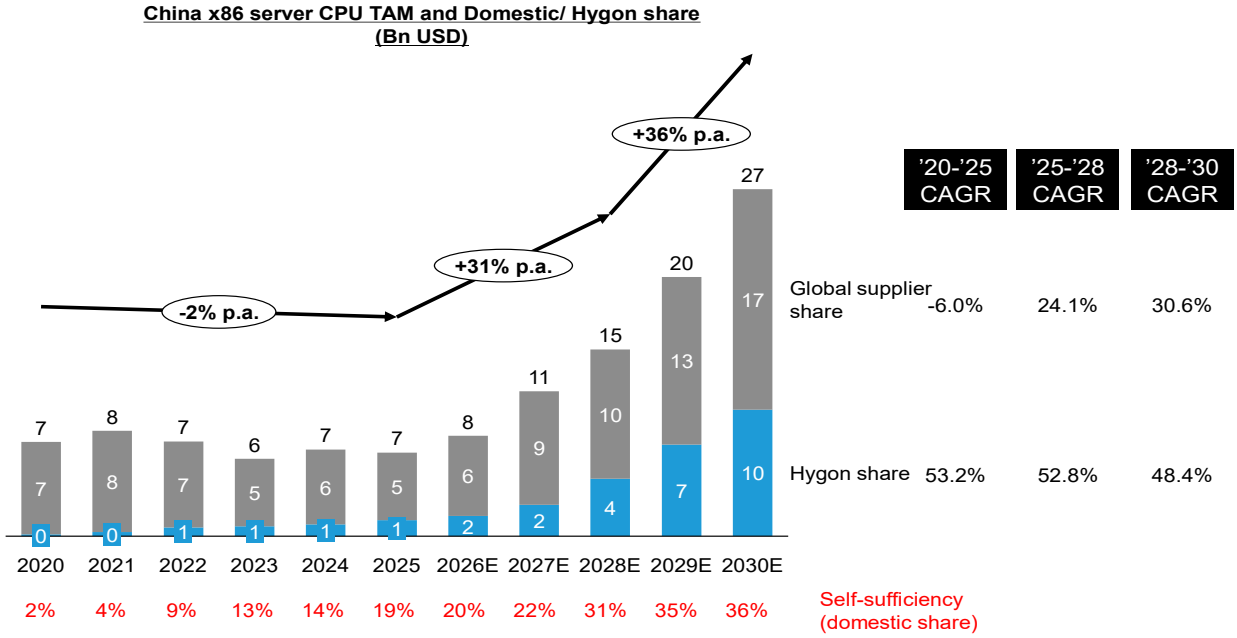

Source: Company disclosures, Mercury, Bernstein estimates and analysis.

EXHIBIT 10: **On a unit shipment basis, Hygon's share will grow faster than on a value basis, driving China's x86 server CPU self-sufficiency rate to approximately 50% by 2030**

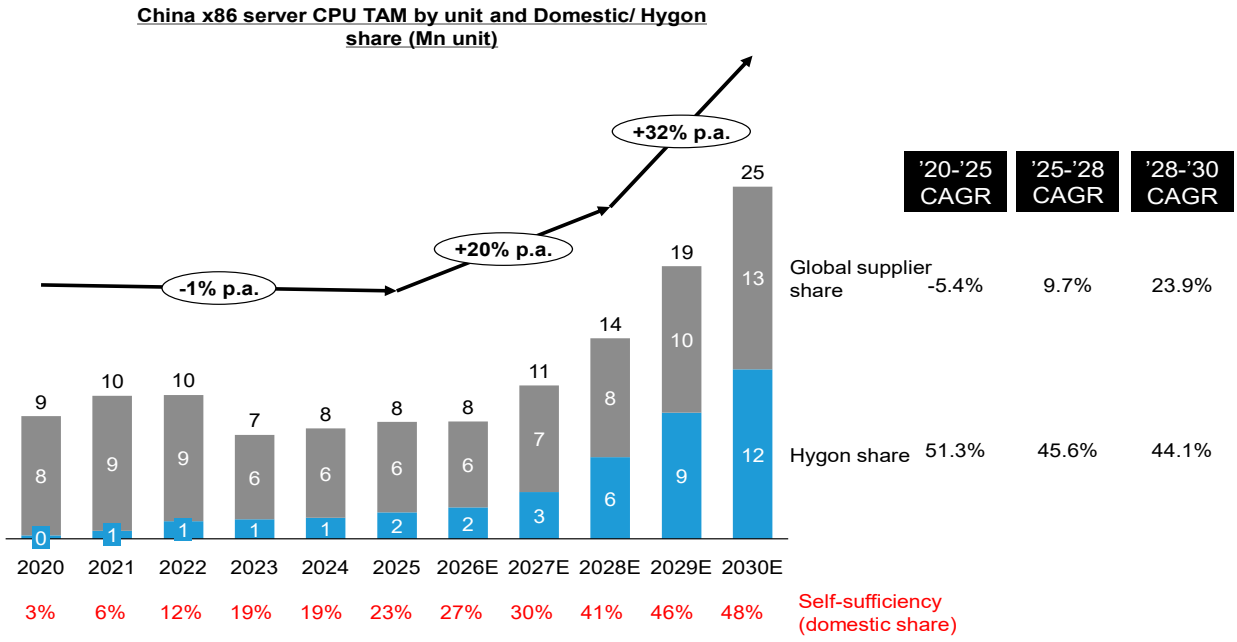

Source: Company disclosures, Mercury, Bernstein estimates and analysis.

HYGON: THE ONLY X86 ALTERNATIVE TO INTEL/AMD GLOBALLY

HYGON'S TECHNOLOGY FOUNDATION-THE AMD DEAL

Hygon was established in 2016 as the commercial product vehicle of a technology licensing arrangement between AMD and Tianjin Haiguang Advanced Tech Investment (THATIC), a Chinese state-affiliated vehicle. Through the joint venture structure, AMD licensed its Zen 1 x86 microarchitecture to the China-based design entity (CHIC), along with SoC design assets for server and workstation applications. The license was geographically restricted to Greater China. AMD held a 51% stake in THATIC, providing economic exposure without controlling the product company. For China, the JV provided immediate access to x86 IP enabling domestic processors compatible with the existing enterprise software ecosystem.

The partnership was curtailed in June 2019, when Hygon was placed on the US Entity List, restricting access to US-origin technology including EDA tools and AMD design services. The Zen architecture license was effectively capped at Zen1. Despite this, Hygon retained the legal rigth to build upon the Zen1 IP already transferred, and the license itself cannot be retroactively revoked by the Entity List designation.

Xinchuang: Hygon's core customer segment today

Hygon's primary sales channel is anchored in the Xinchuang policy ecosystem, under which Hygon is the sole qualified x86 server CPU provider for Party & Government and SoEs in key industries. In-scope buyers need a specific waiver, if they want to procure from global vendors.

End-customer industry mix has diversified from an education-heavy skew in 2019 to a more balanced profile across telecom, finance, internet, and transportation at present.

We estimate that Hygon's current revenue mix is approximately 60-65% CPU and 35-40% DCU. While the DCU business has historically been subscale, it is growing at a faster rate than the CPU segment and increasingly **strengthens Hygon's positioning as a potential full-stack domestic computing platform over the long term**. Drawing on the experience of global leaders such as NVIDIA and AMD, we believe a full-stack compute offering can generate meaningful synergies for the CPU business. A broader product portfolio enhances customer stickiness, enables platform-level optimization across CPU and accelerator workloads, and increases Hygon's strategic relevance in large-scale data center deployments. Over time, these benefits could support higher CPU adoption and reinforce Hygon's competitive position within China's domestic computing ecosystem.

EXHIBIT 11: **Hygon's CPU+DCU hardware is supported by the software ecosystem**

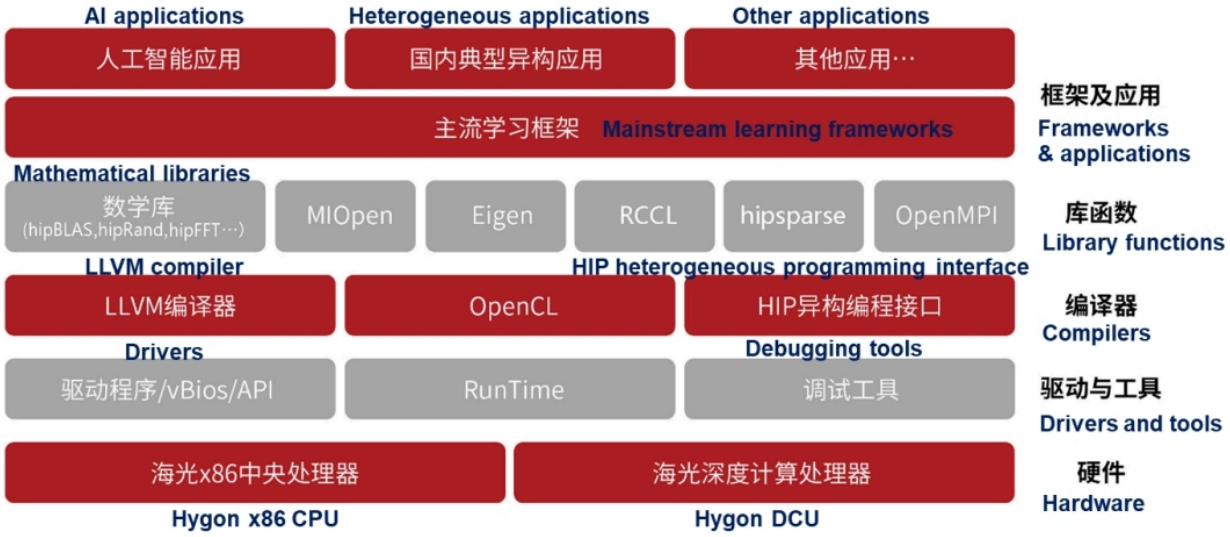

Source: company website, Bernstein analysis

HYGON'S PRODUCT ROADMAP-INNOVATION BEYOND AMD

Hygon has executed four commercial CPU generations since its founding, with a fifth - the C86-G5 - currently in development. Each generation represents a meaningful step in technical capability, design independence, and manufacturing self-sufficiency.

Generation 1 (Dhyana, 2018): the AMD foundation

The first Hygon server CPU, launched in 2018 at GlobalFoundries 14nm, was the direct commercializiton of AMD Zen1 license. Architecturally equivalent to AMD's EPYC Naples in core design, the Dhyana established Hygon's foundational x86 compatibility credentials with Xinchuang OEMs and secured initial government certifications.

Generation 2 and 3 (Gen 2 and C86-G3, 2020-22): the transition generations

Gen2 and Gen3 were originally planned at 7nm using Samsung and TSMC capacity. US sanctions after the 2019 Entity List designation cut off access to leading-edge overseas foundries, forcing a reversion to 14nm manufacturing in China. Despite this setback, Hygon delivered 30-50% per-generation performance improvements through innovations in architecture, memory, and I/O, rather than core count or node scaling.

Generation 4 (C86-G4, 2024): in-house microarchitecture and core count breakthrough

Key spec includes up to 64 cores, DDR5 memory support, and PCIe5.0 connectivity. We estimate it uses SMIC N+1 (10nm equivalent), reducing exposure to future escalation in the US export controls and establishing a manufacturing baseline scalable with SMIC's ongoing capacity expansion.

Generation 5 (C86-G5, in development): the performance step-change

The C86-G5 is Hygon's most ambitious product and the centerpiece of its medium-term growth thesis. It represents a fundamental architectural inflection. Key specifications include:

- 128 cores maximum (up form 64 in G4), achieving parity with AMD EPYC Bergamo (Zen 4C, 2023) and Intel Grantie Rapids (Xeon 6, 2024) in raw core count per socket
- 512 threads via SMT4, substantially increasing per-socket throughput for multi-threaded enterprise and AI inference workloads
- 16-channel DDR5 memory support, providing materially higher memory bandwidth for data-intensive applications
- CXL2.0 integration, enabling memory expansion, coherent memory pooling, and composable infrastructure architectures aligned with the next-generation DC designs
- A fully-self-developed microarchitecture, representing a clean break from the AMD Zen architecture licensed heritage
- We estimate it will use SMIC N+2 (7nm equivalent)

Together, these specifications position C86-G5 to be broadly competitive with current available offerings from AMD and Intel nowadays.

EXHIBIT 12: **From Gen4 onward, Hygon's CPU roadmap features a doubling of core count with each successive generation, alongside a transition to a self-developed microarchitecture**

Source: company reports, Bernstein analysis

From AMD-licensed to fully self-developed: a fundamental architecture shift

Hygon's CPU architecture should no longer be regarded as a direct derivative of AMD's Zen architecture. Although the company initially licensed AMD's x86-based technology in its 1st generation CPU, years of independent redesign and architectural innovation have resulted in a substantially differentiated processor platform with proprietary security and system features. A notable validation of this technological independence was highlighted in a South China Morning Post report, which noted that Hygon processors were not affected by the recently disclosed "StackWarp" hardware vulnerability that impacted multiple generations of AMD Zen processors. Industry analysts attributed Hygon's immunity to its self-developed CSV3 security architecture rather than AMD's SEV-SNP implementation, demonstrating that Hygon has evolved beyond simple technology adoption toward genuine indigenous innovation. This case provides tangible evidence that Hygon has established a largely self-reliant x86 CPU ecosystem, enhancing both security resilience and technological autonomy while reducing exposure to risks associated with upstream foreign technologies.

COMPARING HYGON'S X86 CPU VS INTEL/AMD

Hygon has made steady progress in narrowing "technology years" behind global x86 peers. The lag has improved from approximately 2–4 years in the C86-G3 generation to around 2 years in C86-G4.

In terms of SPECrate 2017 score, we estimate that the **C86-G4 (2024)** can get mark roughly comparable to **Intel Xeon 4th/ 5th Generation (2023) and AMD Milan (Zen 3, 2021–2022) and Genoa (Zen 4, 2022).** This reflects a gradual closing of the gap compared to earlier generations, where Hygon trailed Intel and AMD more measured by years.

However, it is important to note that global competitors have accelerated their innovation cycles. AMD and Intel have significantly increased both the frequency of product iteration and the magnitude of performance improvements per generation, perhaps driven by advancements in leading-edge fabrication and advanced packaging technologies, as well as intense competitive pressure.

Due to ongoing advanced-node fabrication constraints, Hygon's current-generation products (C86-G4) still exhibit a wide raw performance gap relative to the latest platforms launched in 2024, such as Intel's Granite Rapids and AMD's Turin.

Looking ahead to C86-G5, the combination of 128 cores/ 512 threads, IPC improvements from the self-developed microarchitecture, a 16-channel DDR5 memory subsystem, and SMIC advanced-node fabrication can **push G5's SPEC2017 int_rate score to 2000 range, closing the performance gap to current AMD and Intel best offerings to approximately one product generation by the time of launch**. By the time of G5 ships, AMD and Intel will highly likely have advanced to their own next generations; but Hygon's relative competitiveness for general CSP workloads improves materially, expanding its addressable market beyond the current Xinchuang floor.

EXHIBIT 13: **Hygon is maintaining technology gap in years with global x86 peers to be around 2-3 years.**

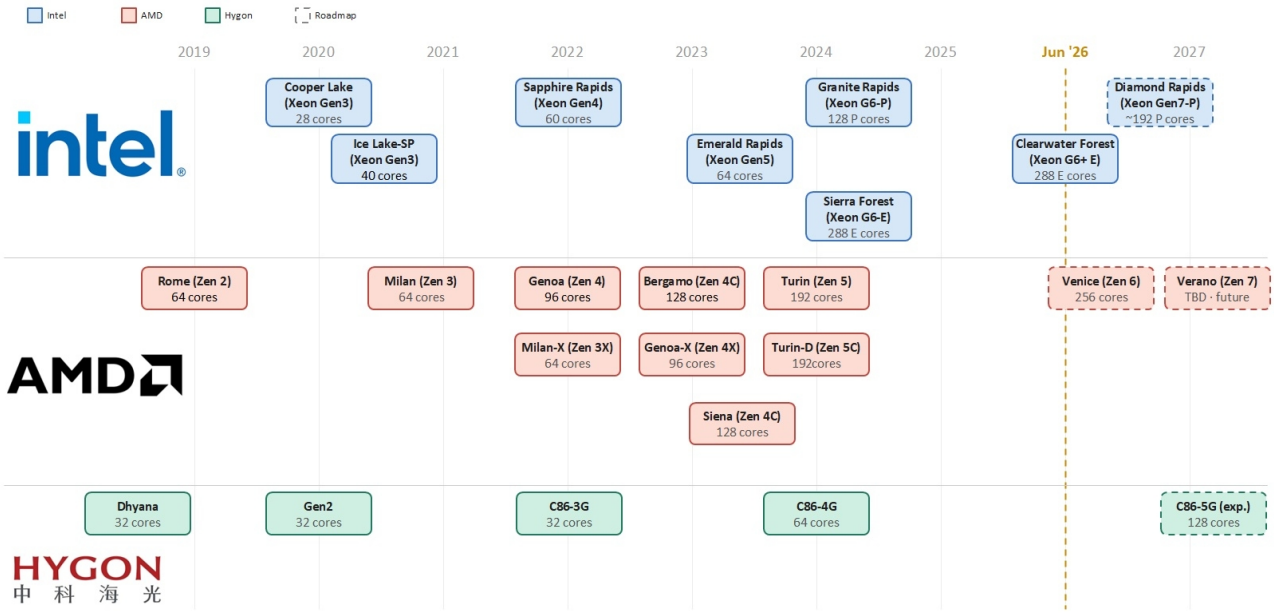

Source: companies disclosures, Bernstein analysis and estimates

EXHIBIT 14: **Hygon has broadly maintained a stable performance gap versus global peers at 2-3 years over the past few years. We expect C86-G5 to achieve a meaningful performance upgrade, further expanding Hygon's SAM penetration within the TAM from the current ~50% level**

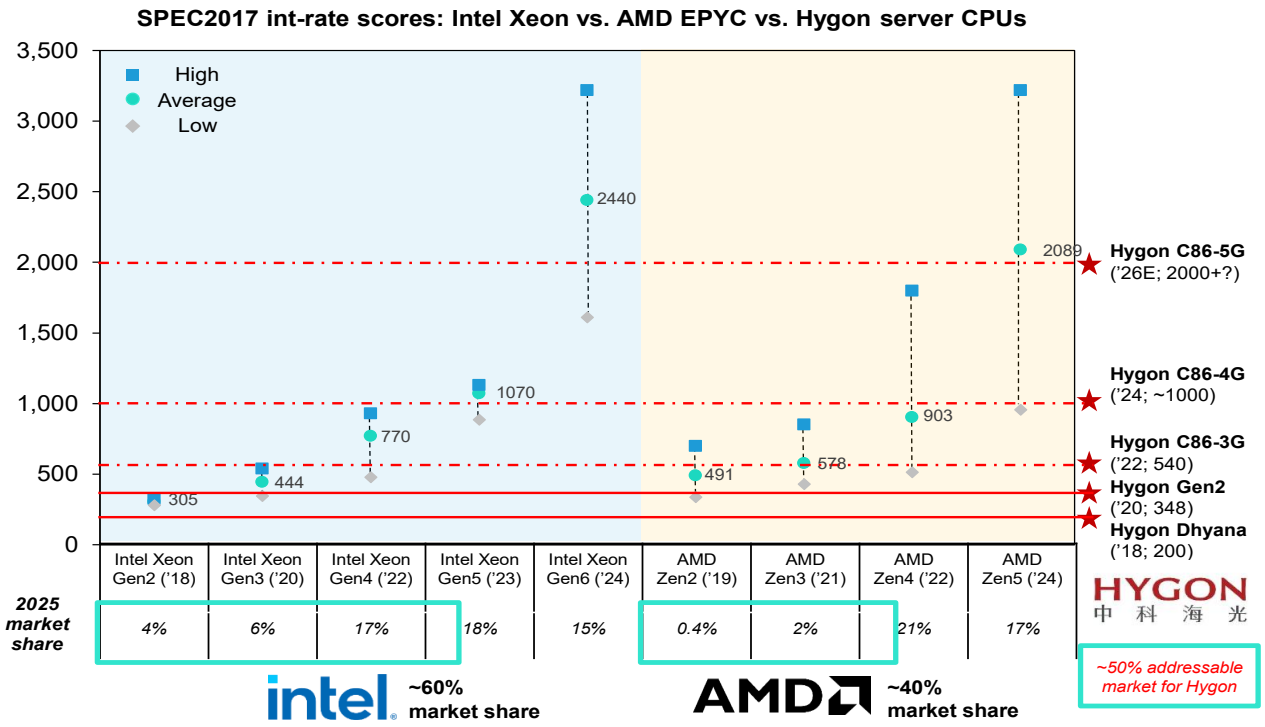

1. Hygon score is estimated base on it's flagship 7000 series. Gen 1 and 2 data are reported by the company and rest are estimation through channel checks.

2. 2025 market share is based on Mercury's analysis on global Intel/AMD sales, excluding Hygon.
Source: Companies disclosure, Mercury, Bernstein analysis and estimates

EXHIBIT 15: **AMD CPU product roadmap: AMD leads in the CPU upgrade, especially aggressively pursuing larger core count. The spec of Milan and Genoa should be comparable with Hygon's C86-G4**

Gen	Brand	Launch time	Max core	Threads	DRAM spec	I/O supporting	IPC vs. prior Gen	Process node	SPECrate 2017 INT (2P)
EPYC 7001 (Zen 1)	Naples	2017 Jun	32	64	8 channels; DDR4-2666MT/s	PCIe3.0 (128L)	baseline	GlobalFoundry 14nm	~380
EPYC 7002 (Zen 2)	Rome	2019 Aug	64	128	8 channels; DDR4-3200MT/s	PCIe4.0 (128L)	+15%	TSMC 7nm	491 avg., 700 flagship
EPYC 7003 (Zen 3)	Milan	2021 Mar	64	128	8 channels; DDR4-3200MT/s	PCIe4.0 (128L)	+19%	TSMC 7nm	578 avg., 850 flagship
EPYC 7003X (Zen 3)	Milan-X (3D V-cache version)	2022 Mar	64	128	8 channels; DDR4-3200MT/s	PCIe4.0 (128L)	0% IPC, upgrade to 3D cache	TSMC 7nm + SoIC	NA
EPYC 9004X (Zen 4)	Genoa	2022 Nov	96	192	12 channels; DDR5-4800MT/s	PCIe5.0 (128L) + CXL1.1	+13%	TSMC N5 (CCD); N6 (IOD)	903 avg., 1330 flagship
EPYC 9004X (Zen 4)	Genoa-X (3D V-cache)	2023 Feb	96	192	12 channels; DDR5-4800MT/s	PCIe5.0 (128L) + CXL1.1	0% IPC, upgrade to 3D cache	TSMC N5 + SoIC	NA
EPYC 97X4 (Zen 4c)	Bergamo	2023 Jun	128 Zen4c	256	12 channels; DDR5-4800MT/s	PCIe5.0 (128L) + CXL1.1	Density focus	TSMC N5 (CCD); N6 (IOD)	1800
EPYC 8004 (Zen 4c, 1P)	Siena	2023 Sep	64 Zen4c	128	6 channels; DDR5-4800MT/s	PCIe5.0 (96L) + CXL1.1	Density focus	TSMC N5 (CCD); N6 (IOD)	NA
EPYC 9005 (Zen 5)	Turin	2024 Nov	192	384	12 channels; DDR5-6000MT/s	PCIe5.0 (128L) + CXL2.0	+17%	TSMC N3E (CCD); N6 (IOD)	958 for 32-core (9355P); 3220 for 128-core (9575F)
EPYC 9005 (Zen 5C)	Turin-D	2024 Nov	192 Zen5c	384	12 channels; DDR5-6000MT/s	PCIe5.0 (128L) + CXL2.0	Density focus	TSMC N3E (CCD); N6 (IOD)	3230 for 192-core
EPYC 9006 (Zen 6/6c)	Venice	2026 H2	256	512	12-16 channels; DDR5 MRDIMM 12800MT/s	PCIe6.0 (exp) + CXL3.0 (exp)	+17-20%; +1.7x throughput vs. Turin	TSMC N2P (CCD); N6 (IOD)	NA
EPYC (Zen 7)	Verano	2027 exp.	256+	512+	TBD	PCIe6.0 (exp) + CXL3.0 (exp)	TBD	TSMC N2P or future	NA

Source: company reports, SPEC, Bernstein analysis

EXHIBIT 16: **Intel CPU product roadmap: Intel also has been accelerating its product iteration in recent years, partially benefiting from the tech advancement of its foundry. We estimate Intel's Grantie Rappids spec is quite comparable with that of Hygon's C86-G5**

Gen	Brand	Launch time	Max core	Threads	DRAM spec	I/O supporting	IPC vs. prior Gen	Process node	SPECrate 2017 INT (2P)
Xeon Scalable 1st Gen	Skylake-SP	2017	28	56	6 channels; DDR4-2666MT/s	PCIe3.0 (48L)	~11%	14nm+	315 (platinum 8180; est.)
Xeon Scalable 2nd Gen	Cascade lake-SP	2019	28	56	6 channels; DDR4-2933MT/s	PCIe3.0 (48L)	~5%	14nm++	395 (platinum 8280; est)
Xeon Scalable 3rd Gen (4/8S)	Cooper lake	2020	28	56	6 channels; DDR4-2933MT/s	PCIe3.0 (48L)	~0% (BF16 add)	14nm++	NA
Xeon Scalable 3rd Gen (2S)	Ice lake-SP	2020-21	40	80	8 channels; DDR4-3200MT/s	PCIe4.0 (64L)	~17-20%	Intel 10nm SF (≈TSM N7)	590 (platinum 8380; est)
Xeon Scalable 4th Gen	Sapphire Rapids	2022	60	120	8 channels; DDR5-4800MT/s or HBM2e	PCIe5.0 (80L); CXL1.1	~10-13%	Intel 7 (10nm ESF, ≈TSM N7)	673 (Gen avg.), 920 (flagship)
Xeon Scalable 5th Gen	Emerald Rapids	2023	64	128	8 channels; DDR5-5600MT/s	PCIe5.0 (80L); CXL1.1	~3-5% per core + core count	Intel 7	1070 (Platinum 8490H)
Xeon 6 P-series	Grantie Rapids	2024 Sep	128 P core	256	12 channels; DDR5-6400MT/s or MRDIMM-8800MT/s	PCIe5.0 (136L); CXL2.0	~10-15%	Intel 3	1610 (Platinum 6960P) by 72-core; 2440 by 128-core (6980P)
Xeon 6 E-series	Sierra Forest	2024 Q2	288 E core	288	8 channels; DDR5-5600MT/s	PCIe5.0 (96L); CXL2.0	Density optimized	Intel 3	3500 (6780E; est.)
Xeon 6+ E-series	Clearwater Forest	2026	288 E core	288	12 channels; DDR5-8000MT/s	PCIe5.0 (96L); CXL2.0	+2.26x throughput vs. Sierra Forest	Intel 18A for CCD; Intel 3 for L3 cache & MC; Intel 7 for IO	NA
Xeon 7 P-series	Diamond Rapids	26H2-2027 exp.	~192 P core	~384	16 channels; DDR5-8000MT/s or MRDIMM-12800MT/s	PCIe6.0; CXL3.0 (exp.)	~50% more cores than Grantie Rapids	Intel 18A for CCD	NA
Xeon next-gen P-series	Coral Rapids	2028 exp.	TBD	TBD	DDR5 MRDIMM	PCIe6.0+; CXL3.0 (exp.)		Intel 14A or future	NA

Source: company reports, SPEC, Bernstein analysis

FEASIBILITY OF G5 PERFORMANCE BREAKTHROUGH

The expected performance improvement in C86-G5 is supported by progress across the three key drivers of CPU performance: **core count, instructions per cycle (IPC), and memory subsystem capability**.

1. **Core count scaling through chiplet architecture:**

   The step from 64 to 128 cores is made feasible by **chiplet-based multi-die packaging**. Rather than attempting a 128-core monolithic die, which would face prohibitive yield losses at any available process node, we expect Hygon to tile multiple smaller compute chiplets on a common substrate using die-to-die interconnect, following the approach pioneered by AMD. This decouples core count scaling from the fabrication node bottleneck in China. Compute chaplets can be manufactured at SMIC's 7nm equivalent node in manageable die sizes with acceptable yields, rather than putting increasing number of cores on a monolithic die that can be only manufactured by TSMC in the most leading node. Through the chiplet approach, Total socket core count is determined by how many compute dies are packaged together.

2. **IPC improvement through self-developed microarchitecture**:

   Hygon's internal CPU design team has accumulated experience across multiple full design cycles from product iterations from Gen2 through Gen4, a critical prerequisite for the IPC improvement +17% over previous Gen. We think the domestic EDA ecosystem should be a critical enabler of this progress. Domestic EDA tools are becoming mature to support advanced digital logic synthesis and physical design at 7nm.

3. **Memory and I/O ecosystem advancement**: The step to 16-channel DDR5 and CXL2.0 integration relies on progress in China's domestic memory and interconnect ecosystem. Contributions should include CXMT's (private company) scaling DDR5 production, providing a domestic DRAM source that reduces dependence on global vendors, who are the most critical bottleneck in today's semi industry. Montage Technology (our coverage) is developing memory expansion controller and CXL IP, enabling G5 to materially increase memory bandwidth and capacity without reliance on foreign IP licenses.

Taken together, these three pillars make the C86-G5 spec technically feasible rather than aspirational.

HYGON'S SAM EXPANSION: FROM XINCHUANG TO CSP

In earlier product generations, Hygon's addressable market was effectively synonymous with Xinchuang procurement. Its primary customer base consisted of Party and Government entities and SoEs across eight designated strategic industries, with demand largely dictated by government budget cycles and localization mandates. Looking ahead, however, we believe Hygon's next phase of growth will increasingly be driven by deployments in data centers operated by commercial hyperscalers and government-backed sovereign AI infrastructure projects. This represents a fundamentally different customer profile, where purchasing decisions are driven by performance, availability, and total cost of ownership (TCO), rather than by policy considerations alone.

Unique value propositions of Hygon in Agentic AI and inference-centric era:

We believe Chinese CSPs will have growing incentives to adopt domestic server CPUs as the performance gap with global competitors continues to narrow while supply constraints on foreign processors become more pronounced. As discussed in previous sections, tightening availability of AMD and Intel processors, coupled with improving competitiveness from domestic vendors, is reshaping procurement economics. In our view, the upcoming C86-G5 will be sufficiently competitive for the majority of AI inference workloads deployed in Chinese data centers.

More importantly, **C86-G5's distinctive value proposition lies in its SMT4 (Simultaneous Multithreading 4) architecture, which enables one of the highest thread counts among server CPUs currently available**. **With up to 512 threads per socket, Hygon enjoys a structural advantage in several emerging AI workloads where thread-level parallelism is more valuable than peak compute performance**. These include multi-tenant inference serving, agentic AI orchestration, and memory-bound large language model (LLM) inference, areas where competing AMD and Intel offerings have limited comparable capabilities.

- **Multi-tenant inference serving**. Inference service providers increasingly operate shared infrastructure serving large numbers of concurrent users and applications. SMT4 effectively doubles logical CPU density relative to conventional SMT2 designs, enabling more simultaneous request handling and lowering cost per inference request.

- **Agentic AI orchestration**. Agent-based AI systems involve multiple sequential and parallel processes, including tool execution, retrieval-augmented generation (RAG), context assembly, workflow coordination, and inter-agent communication. These tasks are primarily memory- and I/O-bound rather than compute-intensive. SMT4's ability to hide latency and maintain high utilization of execution resources is particularly well suited to these workloads.

- **Memory-bound LLM inference**. During LLM inference, CPU threads frequently stall while waiting for model weights and KV-cache data to be retrieved from memory. SMT4 enables additional threads to utilize execution units during these stalls, improving overall throughput per socket without requiring additional silicon area or significantly higher power consumption.

CPU + DCU: Building a Complete Domestic Compute Stack

Hygon's opportunity within CSPs extends beyond CPUs through its DCU product portfolio, which targets AI training and inference workloads. We believe the combination of CPU and accelerator offerings positions Hygon to evolve into a comprehensive domestic compute platform over time.

A full-stack architecture can generate meaningful synergies across the product portfolio. A broader platform offering enhances customer stickiness, supports tighter optimization between CPUs and accelerators, and increases Hygon's strategic relevance in large-scale AI and cloud infrastructure deployments. Over time, Hygon's DCU franchise could therefore serve not only as a growth engine in its own right but also as a catalyst for broader CPU adoption.

A Broader CSP Demand Base: Commercial Hyperscalers and Sovereign AI

The CSP opportunity extends well beyond commercial hyperscalers. China's AI infrastructure build-out increasingly includes a substantial number of government-invested computing centers, funded by provincial and central authorities and often referred to as computing power centers or national AI hubs. While these facilities operate at hyperscale and provide computing services to commercial users, their procurement priorities differ from those of purely commercial CSPs.

Specifically, procurement decisions at sovereign AI projects tend to place greater emphasis on technology sovereignty, supply-chain security, and domestic ecosystem development. Although domestic hardware is not always formally mandated, it is generally preferred. As a result, we believe Hygon's competitive positioning within this segment is even stronger than in the commercial CSP market.

Taken together, commercial hyperscalers and sovereign AI infrastructure projects represent the full scope of Hygon's CSP opportunity. **We project Hygon's penetration of CSP server procurement to increase from an MSD today to around 20%**

by 2030. Given our expectation that CSP-driven demand—including both commercial hyperscalers and national AI hubs—will account for roughly 75% of China's x86 server deployments by 2030, this segment should represent Hygon's most important growth driver in the second half of the decade. Any incremental share gains within this core market would provide meaningful upside to our forecasts.

EXHIBIT 17: **We expect CSP-led demand to dominate server market expansion, contributing ~76% of incremental installations by 2030. As such, successful penetration of the broader CSP segment is a key growth lever for Hygon**

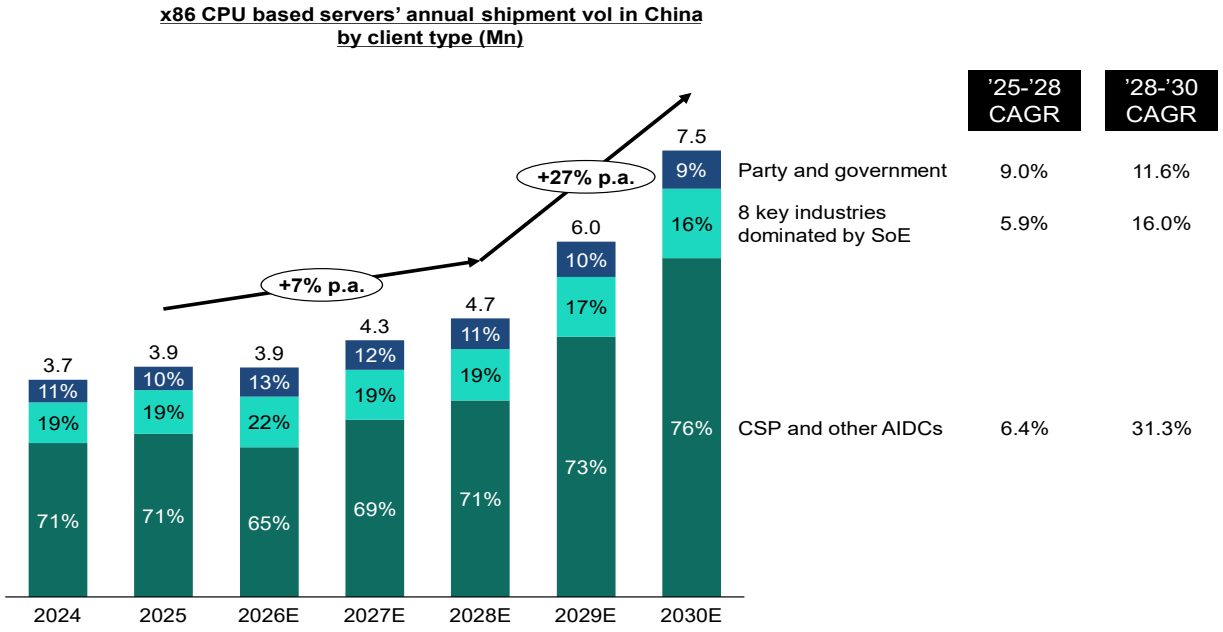

Source: Bernstein analysis and estimates

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

Hygon Information Technology Co Ltd

We value Hygon (PT of CNY 450), based on 2028E EPS of CNY 6.30 and 71x P/E .

RISKS

Hygon Information Technology Co Ltd

Downside risks:

- Failure to develop future generation of server CPU products or keep up with the technology iteration with Intel/AMD.
- Supply chain risks associated with AMD technology authorization and advanced node foundry capacity due to Hygon being added to the US BIS' Entity List.
- The risks associated with slowing down of China’s Xinchuang server deployment in coming years, due to a weaker than expected China macroeconomics.
- Given Hygon is on the Entity List, the US and its Western allies may continue to impose increasingly stringent sanctions on the company.

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	51.2%	15.3%
Market-Perform (Bernstein Brand) Neutral (Autonomous Brand)	HOLD	35.8%	16.2%
Underperform	SELL	13.1%	13.6%

* 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 June 30, 2026. All figures are updated quarterly.

PRICE CHARTS / RATINGS AND PRICE TARGET HISTORY

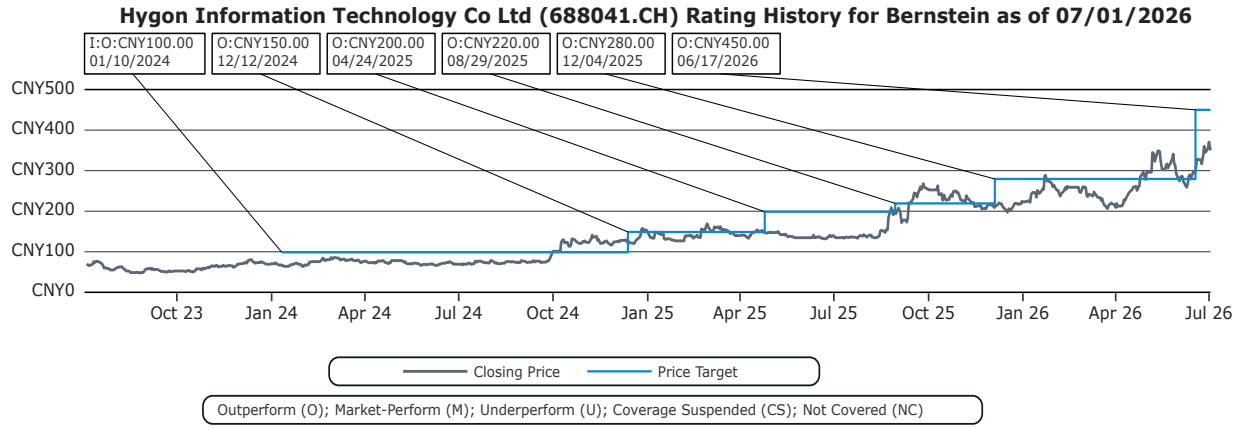

All price target and closing price data in the chart(s) above are denominated in the currency noted in the Ticker Table of this report.

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. 20213710W and located at 8 Marina Boulevard, #12-01, Marina Bay Financial Centre, Singapore 018981, +65-6326-7000.

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 Investment Management Association of Japan. 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.

Any references to Societe Generale herein are purely factual, based upon publicly available information, and included for comparative purposes only. They do not constitute an opinion or recommendation with respect to the securities of Societe Generale.

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 engagement 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 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.