

Apple Inc.

Will agentic AI drive the next inflection in Apple's valuation? PO to \$380

Reiterate Rating: BUY | PO: 380.00 USD | Price: 308.82 USD

Agentic AI shifts control from model to trusted endpoint

In an agentic world, value accrues to the platform that controls user intent, personal context, app access, permissions, identity, authentication, payments, and trust. Smartphone is the scaled consumer device where these factors already converge. If AI assistants become the new front door to search, apps, commerce, scheduling, payments, and workflow completion, we think Apple should have meaningful leverage over model providers, app developers, merchants, advertisers, and payment networks. We raise our PO to \$380 on 37x our C27E EPS of \$10.29 from \$330 on 32x C27E of \$10.29. We raise the multiple to account for the significant opportunity in agentic AI. Reiterate Buy rating.

iPhone is well positioned because Apple controls the stack

Apple's agentic AI moat comes from Apple silicon & iOS. Silicon dictates how much inference can happen locally (matters for latency, privacy, reliability, and cost). OS determines whether AI can access user context, call apps, route requests, authenticate identity, request confirmation, & complete tasks in a trusted way. Our analysis around token usage (Fig 16, 17, 19) suggests Apple will need to rely on distributed compute (on device), Private Cloud Compute (PCC) and third party compute to deliver a robust AI experience. Removing the net cash neutral target, in our opinion, is a reflection of a new period of AI driven investments to enable this agentic era.

What AAPL needs to do: make Siri the execution layer

To win, Apple needs to evolve Siri into the orchestration layer of the iPhone. This requires five priorities: 1) make newer iPhones materially better AI devices through stronger NPUs, memory, bandwidth, etc; 2) scale hybrid inference across on-device models & PCC; 3) evolve Siri into an agent that understands intent, retrieves context, calls apps, and completes workflows; 4) make App Intents a broad framework for machine-callable app actions; and 5) preserve trust through granular iOS permissions, and clear user confirmation. The near-term risk is execution which will determine the ability to capture more of the agentic value pool (Figure 5). We estimate agentic Siri could represent a ~\$15-30bn incremental F30E revenue opportunity in the base case and ~\$40-65bn in the bull case or potentially adding up to \$2 in incremental EPS (Fig 23).

Estimates (Sep) (US\$)	2024A	2025A	2026E	2027E	2028E
EPS	6.75	7.46	8.61	9.96	10.84
EPS Change (YoY)	10.1%	10.5%	15.4%	15.7%	8.8%
Consensus EPS (Bloomberg)			8.73	9.62	10.68
Consensus EPS (Visible Alpha)			8.75	9.52	10.52
DPS	0.99	1.02	1.07	1.12	1.18
Valuation (Sep)					
P/E	45.8x	41.4x	35.9x	31.0x	28.5x
Dividend Yield	0.3%	0.3%	0.3%	0.4%	0.4%
EV / EBITDA*	31.8x	29.6x	25.4x	23.0x	21.2x
Free Cash Flow Yield*	2.4%	2.2%	2.3%	2.8%	3.0%

* For full definitions of *IQmethod*SM measures, see page 34.

26 May 2026

Equity

Key Changes

(US\$)	Previous	Current
Price Obj.	330.00	380.00

Wamsi Mohan
Research Analyst
BofAS
+1 646 855 3854
wamsi.mohan@bofa.com

Ruplu Bhattacharya
Research Analyst
BofAS
+1 646 855 0315
ruplu.bhattacharya@bofa.com

Ryan Seungin Choi
Research Analyst
BofAS
+1 646 743 0587
ryan.choi2@bofa.com

Aisling Grueninger
Research Analyst
BofAS
+1 646 855 4273
aisling.grueninger@bofa.com

Stock Data

Price	308.82 USD
Price Objective	380.00 USD
Date Established	26-May-2026
Investment Opinion	B-1-7
52-Week Range	195.07 USD - 311.40 USD
Mrkt Val (mn) / Shares Out	4,535,749 USD / 14,687.4 (mn)
Free Float	98.0%
Average Daily Value (mn)	14502.95 USD
BofA Ticker / Exchange	AAPL / NAS
Bloomberg / Reuters	AAPL US / AAPLOQ
ROE (2026E)	134.6%
Net Dbt to Eqty (Sep-2025A)	85.1%

See Abbreviations at the end of the report

BofA Securities does and seeks to do business with issuers covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision. Refer to important disclosures on page 35 to 37. Analyst Certification on page 33. Price Objective Basis/Risk on page 33.

iQprofileSM Apple Inc.

iQmethodSM – Bus Performance*

(US\$ Millions)	2024A	2025A	2026E	2027E	2028E
Return on Capital Employed	55.0%	59.6%	61.8%	56.7%	51.0%
Return on Equity	187.3%	171.4%	134.6%	102.3%	77.4%
Operating Margin	34.5%	35.0%	36.3%	35.3%	36.1%
Free Cash Flow	108,807	98,767	106,122	128,365	135,276

iQmethodSM – Quality of Earnings*

(US\$ Millions)	2024A	2025A	2026E	2027E	2028E
Cash Realization Ratio	1.1x	1.0x	0.9x	1.0x	1.0x
Asset Replacement Ratio	0.8x	1.1x	0.8x	0.9x	0.8x
Tax Rate	14.4%	14.3%	15.5%	15.8%	15.6%
Net Debt-to-Equity Ratio	134.7%	85.1%	60.3%	34.0%	16.5%
Interest Cover	NM	NM	NM	NM	NM

Income Statement Data (Sep)

(US\$ Millions)	2024A	2025A	2026E	2027E	2028E
Sales	391,035	416,161	468,701	531,398	562,611
% Change	2.0%	6.4%	12.6%	13.4%	5.9%
Gross Profit	182,787	197,489	228,404	254,949	274,320
% Change	6.8%	8.0%	15.7%	11.6%	7.6%
EBITDA	146,349	157,333	183,307	202,367	219,970
% Change	7.1%	7.5%	16.5%	10.4%	8.7%
Net Interest & Other Income	(232)	(330)	(358)	(520)	(414)
Net Income (Adjusted)	103,982	112,010	126,595	143,918	154,114
% Change	7.2%	7.7%	13.0%	13.7%	7.1%

Free Cash Flow Data (Sep)

(US\$ Millions)	2024A	2025A	2026E	2027E	2028E
Net Income from Cont Operations (GAAP)	115,670	124,595	143,521	157,655	171,398
Depreciation & Amortization	11,445	11,698	13,290	14,715	16,592
Change in Working Capital	3,651	(25,000)	(35,112)	(33,354)	(38,648)
Deferred Taxation Charge	0	0	0	0	0
Other Adjustments, Net	(12,512)	189	(4,782)	2,634	0
Capital Expenditure	(9,447)	(12,715)	(10,796)	(13,285)	(14,065)
Free Cash Flow	108,807	98,767	106,122	128,365	135,276
% Change	9.3%	-9.2%	7.4%	21.0%	5.4%
Share / Issue Repurchase	(94,949)	(90,711)	(86,989)	(100,000)	(100,000)
Cost of Dividends Paid	(15,234)	(15,376)	(15,734)	(16,179)	(16,706)
Change in Debt	(9,958)	(6,451)	(7,914)	0	0

Balance Sheet Data (Sep)

(US\$ Millions)	2024A	2025A	2026E	2027E	2028E
Cash & Equivalents	29,943	35,934	15,420	27,606	46,176
Trade Receivables	66,243	72,957	84,038	119,146	148,538
Other Current Assets	56,801	39,066	48,528	51,392	51,737
Property, Plant & Equipment	45,680	49,834	53,772	60,859	67,935
Other Non-Current Assets	166,313	161,450	175,845	176,792	177,749
Total Assets	364,980	359,241	377,602	435,795	492,136
Short-Term Debt	20,879	20,329	10,000	10,000	10,000
Other Current Liabilities	155,513	145,302	123,411	128,919	120,972
Long-Term Debt	85,750	78,328	74,404	74,404	74,404
Other Non-Current Liabilities	45,888	41,549	55,480	55,538	55,530
Total Liabilities	308,030	285,508	263,295	268,861	260,906
Total Equity	56,950	73,733	114,308	166,934	231,230
Total Equity & Liabilities	364,980	359,241	377,602	435,795	492,136

* For full definitions of iQmethodSM measures, see page 34.

Company Sector

IT Hardware

Company Description

Apple Inc. (AAPL) designs, manufactures, and markets consumer electronics and computers, and has developed its own proprietary iOS, Mac OS, tvOS and Watch OS operating systems and related software platform/ecosystem.

Revenues are principally derived from the iPhone line of smartphones, Services, hardware sales of the Macintosh family of notebook and desktop computers, iPad tablets, and wearables.

Investment Rationale

Our Buy rating on Apple is based on 1) expected strong iPhone upgrade cycle in F25, F26 driven by the need for latest hardware to enable Gen AI features, 2) higher growth in Services revenue, 3) higher margins from more internally developed silicon, 4) continuing capital returns, 5) AI features that can drive higher institutional ownership, and 6) risk around legal issues being manageable.

Stock Data

Average Daily Volume 46,962,460

Quarterly Earnings Estimates

	2025	2026
Q1	2.40A	2.84A
Q2	1.65A	2.01A
Q3	1.57A	1.89E
Q4	1.85A	1.86E

Contents


AAPL’s AI narrative today	4
What does an Agentic AI world look like?	5
What changes in an agentic smartphone?	6
Agentic AI stack for smartphones	8
The iPhone moat in agentic AI	11
Monetization structure of AAPL: today	16
Monetization structure of AAPL: future	18
Token consumption and infra economics	20
App store agents: from app to actions	25
2030 incremental revenue analysis	29
Abbreviations	31

AAPL's AI narrative today

Apple may look like an AI laggard today, but in an agentic AI world the key control point shifts from owning the best model to owning the trusted endpoint where user intent, identity, payments, apps, and privacy converge

So far, the AI trade has benefited a wide range of the infrastructure layer that comprised of semis, semi cap, hardware, networking, memory, power, and cloud. When compared to this mix, Apple has remained a relative bystander with a much more muted stock growth over the last two years. We think some of that skepticism is warranted given that Siri has lagged expectations, Apple has not been a visible participant in the compute buildout race, and frontier model building is not where the company has historically differentiated vs. other Mag 7 peers or frontier labs. This has fed the perception that Apple is an AI laggard in a race that is already well underway. We think this view misses where the next layer of value may accrue. If agentic AI shifts the market from answering questions to taking actions, the critical control point becomes centered around owning the endpoint where user context, identity, permissions, payments, apps, and trust converge. In this world of agentic AI, Apple can control the iPhone, Apple silicon, iOS, App Store distribution, Apple ID, Face ID, Apple Pay, Wallet, privacy architecture, and a premium installed base. Apple does not need to own the best frontier model if it owns the trusted interface that routes intent across local models, Apple-controlled cloud models, external models, and app actions. See our recent notes on AAPL's new M5 chip family & implications on inference ([link to M5 chip note here](#)).

Figure 1: Looking at relative performance since 06/10/24 (WWDC 2024), AAPL has outperformed the S&P 500 and broadly matched QQQ
SPX vs. AAPL vs. QQQ stock performance from 06/10/2024 to 05/18/2026



Source: Bloomberg, BofA Global Research

BofA GLOBAL RESEARCH

What does an Agentic AI world look like?

As AI moves from answering questions to completing tasks, smartphones become the key control point because they own the consumer endpoint where context, trust, identity, payments, and app access come together

The defining shift in agentic AI is the move from information retrieval to task completion. The first wave of generative AI was about asking, summarizing, drafting, coding, and creating. The agentic phase adds planning, tool use, context retrieval, workflow execution, and user confirmation. Agents have the ability to compare options, check the user’s calendar, book the ticket, pay, update the itinerary, notify relevant contacts, and store the receipt. This changes where value accrues. In the gen AI world, the model interface captured a large share of value because interactions are query and response based. In an agentic world, models still matter but assets broaden to include distribution, personal context, app access, payments, identity, authentication, and trust.

Smartphones in an agentic AI world

In a world of agentic AI, we see smartphones evolving from a communication and app-consumption device into the execution layer of consumer intent. Across consumer hardware offerings, the smartphone is the only scaled form factor that consistently brings together identity, authentication, payments, location, camera, microphone, contacts, calendar, messages, apps, and real-time context. As AI shifts from answering questions to taking actions, this bundle of context, trust, and control becomes increasingly valuable. We see companies that own the smartphone form factor such as Apple and its smartphone peers as well as the silicon, memory, and foundry layers beneath it (more on the agentic AI stack for smartphones in Figure 5 below), as the next wave of structural beneficiaries as agentic AI proliferates.

Distribution is a critical control point if not more important than underlying model capabilities. Google’s TAC payments to Apple are evidence that device owners can extract meaningful economics by controlling access to consumers. We think this dynamic can potentially amplify in an agentic era. If the AI assistant becomes the new search box, browser, app launcher, and transaction interface, then the owner of primary interface should have meaningful leverage over players like model providers, app developers, commerce platforms, and advertisers. We expect smartphones to remain the principal consumer monetization surface for agentic AI. New AI-native devices may emerge, but in most cases those devices are more likely to extend the phone than replace it as smartphones already have scale, compute roadmaps, app ecosystems, payment rails, biometric authentication, and daily user habit loops.

Figure 2: In 2023 and 2024 (calendar year), Apple led share in global smartphone shipments
Global smartphone shipment (mn) share data from IDC

2023			2024		
Company	Shipments (mn)	Share	Company	Shipments (mn)	Share
Apple	234.3	20.1%	Apple	232.1	18.7%
Samsung	226.7	19.5%	Samsung	223.4	18.0%
Xiaomi	146	12.5%	Xiaomi	168.5	13.6%
Transsion	94.9	8.2%	Transsion	106.9	8.6%
OPPO	103.4	8.9%	OPPO	104.8	8.5%
Others	358.9	30.8%	Others	402.9	32.5%

Source: IDC, BofA Global Research

Figure 3: In C4Q25, Apple had the highest share of smartphone share although Samsung was the leader throughout C1Q25 to C3Q25

2025 quarterly global smartphone shipment (mn) data from TechInsights

1Q 2025			2Q 2025			3Q 2025			4Q 2025		
Company	Shipments	Share	Company	Shipments	Share	Company	Shipments	Share	Company	Shipments	Share
Others	67.7	22.8%	Others	78.4	26.2%	Others	86.7	27.4%	Others	85.1	25.3%
Samsung	58.3	19.6%	Samsung	55.2	18.4%	Samsung	58.1	18.3%	Apple	83.4	24.8%
Apple	54.7	18.4%	Apple	51.4	17.1%	Apple	53.5	16.9%	Samsung	51.1	15.2%
Xiaomi	39.2	13.2%	Xiaomi	39.3	13.1%	Xiaomi	38.8	12.3%	Xiaomi	33.9	10.1%
Oppo	16.4	5.5%	Oppo	17.0	5.7%	Oppo	19.2	6.0%	Oppo	19.2	5.7%
Honor	14.4	4.9%	Vivo	14.5	4.8%	Vivo	17.3	5.5%	Vivo	17.5	5.2%
Huawei	13.6	4.6%	Honor	14.1	4.7%	Honor	12.7	4.0%	Honor	15.6	4.6%
Vivo	12.4	4.2%	Huawei	12.6	4.2%	Lenovo-Motorola	12.1	3.8%	Lenovo-Motorola	12.2	3.6%
Lenovo-Motorola	11.5	3.9%	Lenovo-Motorola	11.4	3.8%	Huawei	10.4	3.3%	Huawei	11.0	3.3%
Oppo (Realme)	4.5	1.5%	Oppo (Realme)	4.1	1.4%	Oppo (Realme)	4.2	1.3%	Oppo (Realme)	3.6	1.1%
Google	2.0	0.7%	OnePlus	1.8	0.6%	Google	3.3	1.0%	Google	2.8	0.8%
OnePlus	1.5	0.5%				OnePlus	0.6	0.2%	OnePlus	0.8	0.2%
HMD (Nokia)	0.6	0.2%									

Source: TechInsights, BofA Global Research

BofA GLOBAL RESEARCH

What changes in an agentic smartphone?

An agentic smartphone shifts the phone from an app-centric access device to a trusted execution layer that understands user intent, uses personal context, coordinates apps and services, and completes tasks on the user's behalf

An agentic smartphone is a smartphone that can understand user intent, access relevant personal context, reason through the steps required to complete a task, invoke apps and services as tools, and execute actions on the user's behalf with the right level of permission and trust. Today's smartphone is still largely an app-centric device as users decide what they want to do, open relevant apps, navigate interfaces, input information, compare options, and complete tasks. In an agentic model, that workflow inverts as users express intent like booking, buying, summarizing, scheduling, sending, finding, comparing, paying, or reminding.

The smartphone then becomes the orchestration layer that determines how to fulfill it. The key change is that the smartphone moves from being a passive access device to an active execution layer. A traditional smartphone gives users access to apps, content, communication, and commerce while an agentic smartphone uses AI to coordinate across those same tasks. For example, instead of opening Calendar, Mail, Messages, Maps, Safari, and Wallet separately, a user can ask an agentic smartphone to reschedule a meeting, notify attendees, update travel time, and pay a bill. The value of the device shifts from how many apps it can host to how effectively it can translate intent into completed actions. This is where the smartphone becomes uniquely important: it already sits at the intersection of identity, authentication, location, camera, microphone, contacts, calendar, messages, apps, payments, and real-time context.

Figure 4: An agentic smartphone can capture user intent, retrieve personal context, plan and use tools, and execute tasks, all with the trust from the user

Technical layers of an agentic smartphone



Source: BofA Global Research

BofA GLOBAL RESEARCH

An agentic smartphone does five things differently:

1. Understand intent more naturally, allowing users to interact through voice, text, images, or screen context rather than rigid app menus.
2. Retrieve personal context from the device (like calendar, messages, photos, files, location, contacts, and app history) to make the response relevant to the user.
3. Plan workflow by deciding what steps are needed and which apps or services should be used.
4. Execute through app actions, APIs, payments, and device-level permissions.
5. Govern trust by knowing when to act autonomously, when to ask for confirmation, and when to require authentication.

In other words, agentic smartphones are a step up from traditional smartphones into something more useful because they can complete workflows rather than simply provide information. As for apps, in the current smartphone model, the app is the front-facing interface while in an agentic smartphone model, the AI assistant increasingly becomes the front door to digital tasks. The user may not care whether a task is completed through one app or another as long as the outcome is trusted, fast, and accurate (app abstraction). The agent sits between the user and the app interface, calling into apps when needed but reducing the need for the user to manually navigate each app.

However, this does not mean apps will diminish in value. Apps with proprietary data, inventory, network effects, workflows, or transaction rails should remain valuable because agents need fulfillment endpoints. The other major change is that smartphones become more context-aware and more personalized. A generic AI model can answer general questions but a smartphone-based agent can understand the user's specific circumstances such as where the user is, what is on the screen, who the user is messaging, what the user has scheduled, what the user has purchased, what apps the user uses, and what permissions the user has granted. We see this personal context being what turns AI from a general assistant into a useful agent.

For Apple, this distinction matters. If Apple can make the iPhone the trusted environment where agentic AI is executed, it has an inherent advantage as it can control the device, OS, silicon, app distribution, identity layer, authentication, payments, privacy architecture, and a premium installed base. Those assets become more valuable as the smartphone shifts from app consumption to intent execution. The more capable the agent becomes, the more sensitive the data it needs to access. High-value actions require permission, authentication, payment credentials, and user confirmation. That makes trust central to both user experience and monetization. Apple already combines device control, iOS, app distribution, identity, authentication, payments, privacy architecture, and a premium installed base. If Apple succeeds, Siri and Apple Intelligence can evolve from features into the orchestration layer of the iPhone.

Agentic AI stack for smartphones

Agentic AI value should accrue across the smartphone stack, but the biggest winners can be platforms that control multiple layers at once; Apple is advantaged because the iPhone brings together silicon, iOS, app distribution, identity, payments, privacy, and trust at the point where user intent becomes action

Agentic AI should create value across the smartphone stack, but the largest winners are likely to be companies that control multiple layers where compute, context, permissions, identity, payments, apps, and trust converge. As shown in Figure 5 below, the agentic AI smartphone stack comprises foundry and memory, mobile SoC / NPU, operating system, device OEM, foundational models, agent frameworks, and apps / services.


Monetization pools cut across these layers rather than mapping one-for-one to a single layer:

- semiconductor infrastructure monetizes the rising compute and memory intensity
- edge compute providers monetize local inference
- device / OS gatekeepers monetize distribution, context, permissions, and routing
- AI platform and agent owners monetize intent capture and reasoning
- vertical workflow owners monetize fulfillment
- payments / identity / trust rails monetize authorization and trusted execution.

For Apple, the key point is that the company touches multiple monetization pools at once: silicon, device, OS, app distribution, agent framework, identity, payments, privacy, and premium users. If Apple makes Siri and Apple Intelligence the trusted interface, the iPhone can become the consumer execution layer where compute, context, permissions, identity, payments, apps, and trust converge. In that scenario, Apple can own the routing layer, permission layer, trust layer, and default interface where user intent becomes completed action.

Figure 5: Value capture spans across the full agentic smartphone stack by various entities

Agentic AI monetization layers mapped across the agentic AI smartphone stack



Source: BofA Global Research

BofA GLOBAL RESEARCH

Foundry & memory are the foundation for local AI

At the bottom of the stack, foundry and memory benefit as local inference raises the compute and memory requirement per device. Agentic smartphones need more capable NPUs, higher memory bandwidth, more DRAM, better thermal management, and tighter hardware / software optimization to support multimodal workloads across text, voice, images, sensors, and personal context. Foundry and memory suppliers do not own the user relationship, but they should benefit from higher semiconductor content in premium smartphones as more inference moves onto the device. The risk is that memory remains cyclical and parts of the stack can commoditize, but directionally, agentic AI should improve content opportunity even if smartphone unit growth remains mature.

Mobile SoC / NPU: enable local AI at the edge

Above foundry & memory layer is the mobile SoC / NPU layer. The SoC brings CPU, GPU, NPU / Neural Engine, modem, image signal processor, memory controller, and power management all together. In an agentic smartphone, this layer becomes more important because the phone needs to run frequent, low-latency, privacy-sensitive inference without sending every request to the cloud. Qualcomm, MediaTek, and Apple silicon should all benefit from this shift, but Apple is particularly well positioned because it controls its own silicon roadmap. For Apple, the A-series and M-series chips increasingly point in the same direction: more local AI capability, better performance per watt, and tighter hardware / software integration.

OS & OEM: control context, permissions, & distribution

The next control point is the operating system and device OEM layer. The OS determines what the agent can access, what actions it can take, which apps it can call, when the user must approve an action, and how privacy is governed. The device OEM controls the full product experience, including defaults, distribution, hardware / software integration, pricing, preloaded services, and the commercial relationship with the user. This is where we see Apple to be structurally advantaged. Apple controls the iPhone,


The iPhone moat in agentic AI

Apple's iPhone moat in agentic AI rests on two control points: silicon that enables local inference and iOS that governs context, permissions, app actions, payments, and trust; recent third-party reports of a rebuilt Siri point to Apple potentially turning that stack into a true consumer execution layer

We see the iPhone's moat in an agentic AI world through two layers of the smartphone stack: mobile SoC / local compute and operating system. The first layer determines how much AI can happen locally on the device. The second determines whether that intelligence can access context, call apps, authenticate users, and complete tasks in a trusted way. Apple's advantage is that it controls both the silicon roadmap and the operating system environment, which gives the company a differentiated ability to optimize the full path from inference to execution.

Figure 6: We model iPhone installed base (IB) to grow to 1.5bn in 2026

iPhone installed base (mn) history & projections



Source: Company reports, BofA Global Research

Note: Installed base reflects total IB that includes primary and secondary

BofA GLOBAL RESEARCH

Apple's SoC: local inference moat

The first part of the iPhone moat is Apple's control of the mobile SoC. In an agentic AI smartphone, the chip becomes the foundation for local inference. Smartphones need to understand speech, interpret images, summarize text, retrieve personal context, rank notifications, classify intent, and decide what can be handled locally versus what should be routed to a larger cloud model. That requires a combination of CPU, GPU, NPU / Neural Engine, memory bandwidth, power efficiency, and thermal management. This is strategically important because agentic AI increases the value of on-device compute. If every interaction with an AI assistant had to go to the cloud, the user experience would face latency, privacy, reliability, and cost challenges. Local inference helps solve all four. It reduces response time, keeps sensitive context closer to the user, allows lightweight

AI features to work more seamlessly, and lowers the marginal cost of serving high-frequency interactions. As a result, the smartphone SoC becomes a more important value-capture layer as AI shifts from chatbot usage to persistent agentic workflows.

Apple is well positioned, in our view, because it owns the full silicon roadmap. Unlike most Android OEMs, Apple is not dependent on a merchant silicon provider for its flagship smartphone compute architecture. Apple designs the A-series chips and the M-series chips, allowing the company to align hardware capabilities with software priorities over multiple product cycles. Apple’s M5 announcement is relevant because Apple framed M5 around AI performance, including a next-generation GPU architecture with Neural Accelerators in each GPU core, a faster Neural Engine, and higher unified memory bandwidth. Apple silicon is increasingly being optimized around local AI workloads, not just traditional CPU/GPU performance. The recent iPhone 17 family and Air use Apple’s A19 chips that mark a step closer towards a more effective mobile AI performance architecture. Macbook Neo is powered with A18 which we see as encouraging evidence of Apple Silicon’s working scalability across form factors.

iPhone 17 has 8GB of RAM, iPhone Air has 12GB, iPhone 17 Pro has 12GB, and iPhone 17 Pro Max has 12GB. This is a step up from the iPhone 16 family which had 8GB across each of them. This matters for local AI because memory capacity and bandwidth are increasingly important constraints for on-device inference, particularly as models become more multimodal and context-aware. However, based on current silicon, Apple would need to step up the memory significantly (minimum of 16GB but more likely 24GB) to be able to support sufficient on device performance.

Figure 7: The iPhone 17 Pro and iPhone Air uses A19 Pro Chip, though the 17 Pro has a 6-core GPU
iPhone 17 family and Air chip specifications





Figure 8: Neo uses A18 Pro (5-core GPU) with 8GB of unified memory & up to 512GB of storage
MacBook Neo technical specifications



The key moat is that Apple can optimize the entire local AI system: chip, memory, OS, model size, app frameworks, privacy policies, and user experience. In a cloud-centric AI model, the main question is who has the largest GPU cluster or the best frontier model. In an agentic smartphone model, the question becomes who can run the right task in the right place at the right cost with the right privacy guarantee. Apple's vertical integration gives it an advantage in that routing problem. Apple can decide what runs on-device, what goes to Private Cloud Compute, and what is sent to external models. This routing advantage matters economically. Local inference can reduce Apple's dependence on third-party model providers and external cloud inference costs. If Apple can handle high-frequency, lower-complexity tasks locally then Apple can preserve margin while making AI more pervasive. More complex tasks can be routed to Apple's own Private Cloud Compute or to third-party models as needed.

Apple's privacy documentation states that Apple Intelligence attempts to run models on-device when possible and uses Private Cloud Compute for more complex requests that require additional compute. The mobile SoC moat also supports hardware monetization. If the best agentic AI experiences require newer Neural Engines, more memory, better thermal performance, and stronger local inference, then AI becomes a reason to upgrade. The smartphone market is mature, and replacement cycles are long. But AI can support premium mix, higher storage and memory configurations, and faster adoption of Pro-tier devices. In other words, Apple can monetize local AI through hardware even before it directly charges for AI services. Apple's strength is controlling the edge device where consumer AI is used repeatedly throughout the day. If agentic AI is consumed through the smartphone, then the chip inside the phone becomes a strategic asset, not just a component.

Apple's OS: execution moat

Apple's OS moat is that iOS can become the trusted governance and routing layer for agentic AI (controlling context, permissions, app actions, authentication, payments, and model selection), but the moat only matters if Siri becomes the default interface for consumer task execution

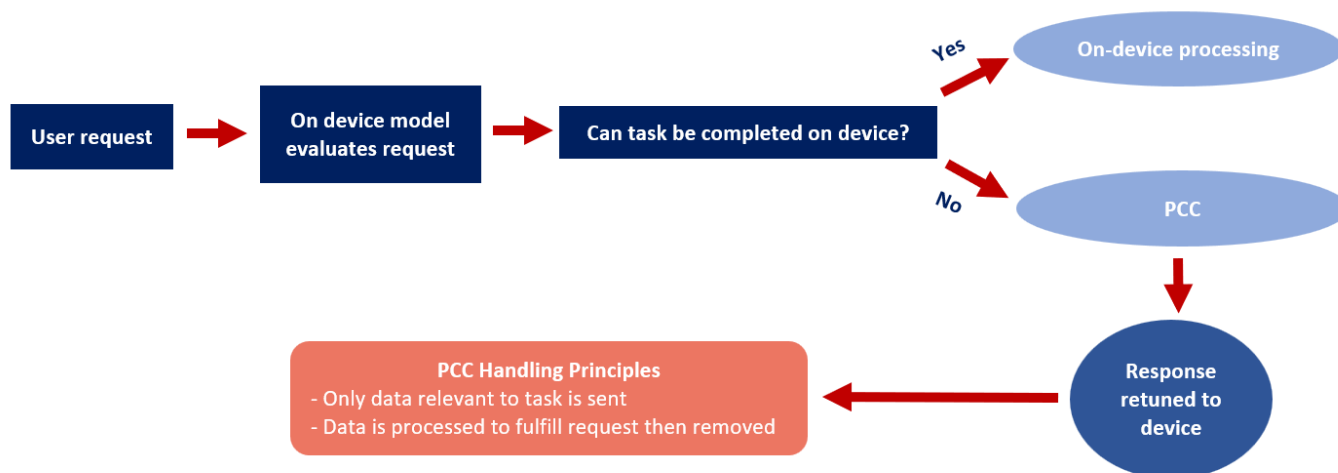
The second part of the iPhone moat is the OS layer. As SoC determines what intelligence can run locally, iOS determines whether that intelligence can become useful. Agentic AI is about taking an action and actions require access, permissions, identity, app invocation, and user confirmation (OS-level functions). iOS sits between the user, apps, device hardware, model layer, and payment layer. It controls permissions, notifications, app entitlements, privacy prompts, background activity, defaults, app discovery, and authentication. Apple's App Intents framework is strategically important because it gives developers a structured way to expose app functionality to Siri, Spotlight, Shortcuts, and system experiences. This supports the shift from Siri as a voice assistant to Siri as an execution layer.

The OS layer is also where Apple can make models interchangeable. This is a critical point because if the user experience is controlled by the OS, then the underlying model can be routed depending on the task. A simple request can run on-device while a more complex private request can go to Apple's Private Cloud Compute (see Figure 9 below). A broad knowledge or frontier reasoning request can go to an external provider. Apple does not need to own every model if it owns the routing layer, the permission layer, and the interface where the user grants trust. Apple's privacy documentation already describes this hierarchy: on-device processing when possible, Private Cloud Compute for

more complex tasks, and user-controlled sharing when external integrations are used. This creates a different kind of moat vs. model ownership. Model quality matters, but models can be swapped, updated, or multi-sourced. The OS relationship is harder to displace because it is embedded in daily device usage, app permissions, identity, payment credentials, and user trust. If Apple can make Siri the default interface for task completion, it can preserve the user relationship even when the underlying reasoning engine is not Apple's own model. This is also where the economics accrue. The OS layer can monetize through multiple paths like services attach, App Store distribution, paid AI features, search or answer revenue share, payments, commerce referrals, app discovery, enterprise management, and developer tools. The more the AI assistant becomes the front door to digital tasks, the more valuable it becomes to control the default interface.

Figure 9: Workloads that can't be processed locally are sent to PCC

Apple's PCC



Source: BofA Global Research

BofA GLOBAL RESEARCH

Apple's OS moat is also reinforced by trust. Agentic AI will require users to delegate more authority to software. The user must trust the assistant to access personal information, understand intent, make recommendations, avoid mistakes, ask for confirmation, authenticate identity, and execute payment. Apple is advantaged because the iPhone already embeds Apple ID, Face ID, Apple Pay, Wallet, App Store billing, app permissions, and privacy controls. The OS moat also protects Apple from app abstraction risk. In an agentic model, apps may become less visible as the assistant becomes the front door. That could weaken the value of some app brands. But for Apple, app abstraction can be positive if Siri becomes the layer that determines which apps are called and how tasks are fulfilled. Apple can preserve the value of the App Store while shifting the user experience from manual app navigation to agent-mediated app execution. In that world, App Intents becomes the equivalent of making apps "machine-callable" by the OS. The risk is execution. Apple's OS moat only matters if Siri and Apple Intelligence become good enough to serve as the trusted agentic interface.

If users instead build daily habits around ChatGPT, Gemini, Claude, or other LLMs, Apple's role could be reduced to hardware distribution and permissioning infrastructure. Apple would still matter, but it would capture less of the agentic value pool. Can Apple can make iOS the default environment for consumer agentic workflows? We hope to hear more around this topic during WWDC in June where we can get further clarification on material improvements and clear AI roadmap execution. On arguments that AI

assistants collapse the need for smartphones due to fewer apps, fewer screens, and fewer explicit interactions, we see that as directionally right for the app interface but wrong for the device category. Agentic AI increases the value of the device (smartphone) that has continuous personal context, trusted identity, payments, secure storage, local sensors and user-granted permissions.

Third-party information on Siri reconstruction

In mid-May, third-party sources suggested that Siri could be materially rebuilt in iOS 27, shifting from a voice assistant powered by a more limited, command-based architecture into an always-on agent that can tap into personal context and take action across apps. If accurate, this would align well with our thesis that Apple's AI opportunity is less about owning the best frontier model and more about making Siri the trusted orchestration layer across local models, external models, app actions, payments, and personal context. It would also support our view that Apple can turn the iPhone into the execution layer for consumer intent, rather than simply another device running AI features. The user interface also appears likely to change meaningfully. Rather than the prior Siri experience, which was largely voice-command driven and episodic, the new interface is expected to be more conversational and back-and-forth, closer to the interaction style users have become accustomed to with frontier AI assistants. Third-party sources also indicate that Siri could become more integrated with the Dynamic Island interface, making the assistant a more persistent on-screen element vs. a standalone overlay.

Figure 10: Siri interface today is command-based

Siri interface today

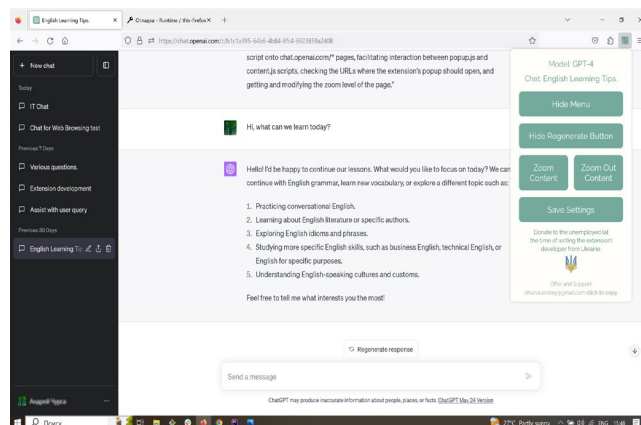


Source: BofA Global Research

BofA GLOBAL RESEARCH

Figure 11: ChatGPT enables a chat-based conversation-like interface

ChatGPT interface




Source: Firefox, BofA Global Research

BofA GLOBAL RESEARCH

When users engage with the new Siri, the “Search or Ask” interface is expected to resemble Spotlight Search (Figure 12 below) on iPhone today, but with more advanced responses and deeper access to data from within apps. Another notable change is the potential for easier toggling between Siri and third-party AI offerings. This would fit with Apple's likely model strategy: Apple can own the interface, permissions, routing, and trust layer while allowing external models to handle tasks where they are better suited. AI-powered search is also expected to be a key addition, allowing Siri to deliver more detailed answers to general knowledge questions. Siri may also become a standalone app for the first time, which would further signal Apple's intent to make Siri a more central consumer AI surface rather than a background voice feature.

Figure 12: Spotlight Search enables a quick search function across an iPhone

Spotlight Search



Source: Apple

BofA GLOBAL RESEARCH

Monetization structure of AAPL: today


Apple monetizes through a hardware-led installed-base flywheel today, but agentic AI could make that base more valuable by shifting Services monetization from app access and transactions toward intent routing, trusted execution, and commerce enablement

Apple's monetization structure today is built around a simple but powerful flywheel: sell premium hardware, expand the installed base, increase user engagement, and monetize that installed base through high-margin services over time. The iPhone remains the center of this model. The key point is that Apple is not purely a hardware company, but the hardware remains the entry point into the ecosystem. Products generated \$307.0bn of FY25 revenue, or ~74% of total sales, while Services generated \$109.2bn, ~26% of total sales. Within Products, the iPhone represented ~68% of Product revenue, evidence we see as why the iPhone remains Apple's important monetization surface even as Services continues to scale at higher margins.

Figure 13: In F25, Product revenue of \$307bn was 74% of total Apple's total revenue

AAPL revenue (\$mn)

) trajectory from F20-F28E




Source: Company reports, BofA Global Research

BofA GLOBAL RESEARCH

Figure 14: In F25, Product GM came in at 37% while Services GM was 75%

AAPL GM trajectory from F20-F28E



Source: Company reports, BofA Global Research

BofA GLOBAL RESEARCH

HW: premium devices, high ASPs, & replacement cycle

The first layer of Apple's monetization is upfront hardware sales. Apple monetizes the user at the point of device purchase through iPhone, Mac, iPad, Apple Watch, AirPods, and accessories. This is still the largest revenue contributor, but it is not a simple unit-volume model. Apple's hardware monetization is driven by premium mix, ASP, storage tiers, Pro adoption, trade-in programs, financing, carrier subsidies, and replacement cycles. The iPhone is the most important part of this structure. In F25, iPhone revenue was \$209.6bn, up from \$201.2bn in F24. This matters because Apple's iPhone model is less dependent on broad unit growth than many consumer hardware categories. Apple can grow through mix shift toward higher-end devices, storage upgrades, and ecosystem-driven retention even when global smartphone units are relatively mature. The hardware gross margin profile is meaningful but lower than Services. In FY25, Apple Products gross margin was 36.8% vs. Services gross margin at 75.4%. Hardware brings users into the ecosystem and creates the installed base, while Services monetizes that base at structurally higher margins.

Services: recurring, high-margin ecosystem

The second layer of Apple's monetization is Services. Services includes categories such as the App Store ([see our recent Sensor Tower app store note here](#)), advertising, cloud services, AppleCare, payment services, and digital content subscriptions. Apple does not break out each Services component in detail, but the aggregate category has become the company's most important profit engine. In F25, Services revenue was \$109.2bn, up from \$96.2bn in F24. Services is economically powerful because it combines recurring revenue, high gross margin, and strong linkage to installed base growth. In F25, Services generated roughly \$82.3bn of gross profit, representing 42.2% of Apple's total gross profit despite being only about 26.2% of total revenue. A larger installed base creates more opportunities for App Store transactions, iCloud storage, Apple Music, Apple TV+, AppleCare, Apple Pay, advertising, subscriptions, and other ecosystem services. At the same time, Services increases switching costs because the user becomes more embedded in Apple's ecosystem.

Distribution, payments, identity, and installed base

Apple controls access to a premium installed base through iOS, the App Store, default settings, APIs, privacy rules, payments, and device-level surfaces such as Safari, Siri, Spotlight, Wallet, and notifications. This distribution layer is critical for agentic AI because it shows Apple already monetizes control points around user access, defaults, app distribution, and payments. This gives Apple the ability to monetize distribution directly and indirectly. The App Store is the most visible example, where Apple monetizes digital goods, subscriptions, and developer access to iOS users. Search and advertising economics are another example, where default placement and user access create monetizable distribution value. If the AI assistant becomes a new consumer gateway, Apple's existing distribution economics could extend into model routing, answer placement, commerce referrals, and agent-mediated app discovery. The most important asset behind Apple's monetization structure is the installed base. The installed base converts hardware sales into recurring monetization opportunities across apps, services, payments, subscriptions, search, advertising, and support. Agentic AI matters because it could increase the monetization intensity of that installed base.

On payments, identity, and trust, Apple has built a set of assets around Apple ID, Face ID, Touch ID, Wallet, Apple Pay, App Store billing, subscriptions, and device security. These assets are not always visible as separate revenue lines, but they support Services growth and deepen ecosystem lock-in. This layer matters because Apple monetizes content and apps as well as trusted access and transaction enablement. App Store billing, in-app subscriptions, Apple Pay, Wallet, and AppleCare are all examples of Apple monetizing the user relationship after the hardware sale. Today, Apple's payment and identity assets reinforce the broader monetization flywheel. They make transactions easier, reduce friction for subscriptions and purchases, and create more reasons for users and developers to stay inside the ecosystem.

Monetization structure of AAPL: future

Apple can monetize agentic AI by using better iPhone hardware to drive upgrades and using Siri, App Intents, Apple Pay, Wallet, iCloud, and iOS permissions to increase recurring Services revenue per user

What AAPL needs to do to win

Apple's future monetization opportunity in agentic AI depends on whether the company can evolve the iPhone from a premium device and Services gateway into the trusted execution layer for consumer intent. The opportunity is to deepen both sides of Apple's model: make users more likely to upgrade to AI-capable devices and increase recurring value of the installed base through AI services, app actions, payments, storage, search / answer economics, and commerce routing. We see Apple's future monetization structure depends on five execution priorities:

- Make newer iPhones materially better AI devices through stronger NPUs, memory, bandwidth, thermal design, and battery efficiency
- Scale hybrid inference economically by routing high-frequency tasks on-device, complex private tasks to Private Cloud Compute, and only selected tasks to external models

- Transform Siri into an orchestration layer that can understand intent, retrieve context, call apps, and complete workflows
- Turn App Intents into a broad agent-action framework so third-party apps become callable fulfillment endpoints
- Own the trust layer through iOS permissions, Apple ID, Face ID, Apple Pay, Wallet, privacy architecture, and clear user confirmation

In an agentic AI world, the opportunity is to deepen both sides of the model: make users more likely to upgrade to AI-capable devices and increase the recurring value of the installed base through AI services, app actions, payments, storage, search / answer economics, and commerce routing.

Hardware

For Apple to succeed, the iPhone must become capable enough to run a meaningful amount of AI locally. Agentic AI creates a different compute requirement than traditional mobile apps because the device needs to understand language, interpret images, summarize context, classify intent, retrieve relevant personal data, and decide whether a request should be handled on-device, through Apple's Private Cloud Compute, or through an external model. Apple's M5 announcement is directionally important because Apple highlighted a GPU architecture with Neural Accelerators in each core, a faster Neural Engine, and higher unified memory bandwidth, all framed around AI performance.

For iPhone, the hardware roadmap needs to emphasize four things: NPU / Neural Engine performance, memory capacity and bandwidth, thermal efficiency, and battery life. The more agentic workflows run locally, the more Apple can reduce latency, preserve privacy, and avoid turning every AI interaction into an external cloud inference cost. This is particularly important because high-frequency AI tasks (summarizing messages, ranking notifications, drafting replies, semantic search, etc.) could occur many times per day. If those workloads are mostly cloud-based, the economics become variable-cost heavy. If those workloads are local, Apple can monetize through hardware mix and Services while protecting gross margin. To monetize hardware, Apple needs to make AI capability a reason to buy newer and higher-end iPhones. That does not necessarily mean a classic unit super cycle, but it can support premium mix, Pro adoption, higher memory / storage configurations, and faster replacement among power users. In our view, the goal should be to make the newest iPhones meaningfully better at agentic AI than older devices. In that future world, Apple monetizes AI first through device differentiation.

Siri / OS

Apple must evolve Siri from a voice assistant into an agentic orchestration layer. That means Siri needs to understand intent, retrieve context, use the right model, call the right app action, ask for confirmation when needed, and complete the task. The monetization opportunity only emerges if Siri becomes useful enough to capture the user habit loop. This is where App Intents becomes strategically important. Apple describes App Intents as a way for developers to expose app functionality to Siri, Spotlight, Shortcuts, and system experiences. Apple has also expanded developer frameworks so Siri and Apple Intelligence can take actions in apps and use app functionality more directly. For Apple to win, this needs to scale from a developer feature into a broad execution framework across travel, commerce, productivity, payments, messaging, files, calendar, media, and enterprise workflows. The software monetization implication is that Apple can create the agentic equivalent of App Store distribution. If successful, Apple can monetize through app discovery, App Store economics, subscriptions, paid placement, commerce referrals, Apple Pay, and potentially agent-mediated transaction fees. The app may still fulfill the task, but Siri controls the front door.

Apple also needs to make iOS the trusted permission layer for agentic AI. The agent can only complete high-value tasks if users trust the system to access personal data, confirm intent, authenticate identity, and execute payment. This is where Apple's existing assets matter: Apple ID, Face ID, Apple Pay, Wallet, App Store billing, app permissions, device security, and privacy architecture. Apple will need granular controls around what the agent can see, what the agent can do, what requires confirmation, what can happen automatically, and what must be authenticated. If Apple gets this right, trust becomes a monetization layer. If Apple gets it wrong, users will not allow Siri to act on their behalf. This is why Apple's privacy architecture is a product requirement. PCC is designed to extend device-like privacy into the cloud for more complex AI requests, with Apple stating that data sent to PCC is processed only to fulfill the request and is not retained. That framework matters because agentic AI requires far more sensitive context than traditional search or app usage.

Developer ecosystem

For agentic AI to work, apps need to become machine-callable. Users should not need to open each app manually. The agent should be able to call the relevant app capability in the background, subject to permissions and confirmation. This requires developer adoption of App Intents and related frameworks. Apple's challenge is to make this easy enough for developers and valuable enough that developers participate. If developers expose useful actions, Siri becomes more capable. If developers do not, Siri remains limited to Apple's own apps and a narrow set of workflows. Apple needs the same flywheel it created with the App Store: developer adoption increases user utility, user utility increases engagement, and engagement increases developer incentive. The monetization opportunity is that Apple can become the marketplace for app actions. In today's model, apps compete for downloads and screen time. In the agentic model, apps may compete to be the service an agent calls when a user expresses intent. That could create new economics around ranking, fulfillment, transaction fees, subscriptions, and developer tooling.

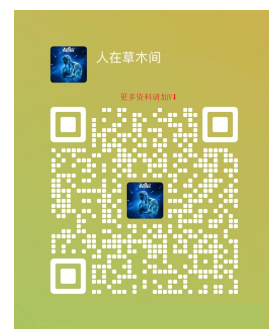
Consumer experience

Apple should make the agent reliable before making it autonomous. The path should likely be staged: assistive AI (summarize, rewrite, search, recommend), confirmable actions (draft, schedule, prepare, compare, then ask for approval), delegated low-risk workflows (complete low-risk tasks within user-defined constraints), and then high-value execution with strong authentication (purchases, payments, bookings, transfers, filings, and enterprise actions with strong authentication). This staged approach matters for monetization because users will only pay for AI that is useful and trusted.

Token consumption and infra economics

iPhone can become a massive agentic AI token surface, but the margin outcome depends on whether Apple owns the inference stack or pays third-party models at scale


If (once) Siri / Apple Intelligence evolves into a true agentic layer, the iPhone could become one of the largest consumer token-generation surfaces in the world. The shift is where many ordinary smartphone actions become model-mediated: summarizing messages, drafting replies, reading calendar context, searching across apps, comparing options, booking, buying, authenticating, and confirming. An agentic request is token-heavier because the system may need to ingest user context, retrieve relevant app data,



reason through steps, call tools, generate a response, and sometimes maintain state across follow-up actions. For example, as shown in in Figure 15 below, a user would ask to a traditional LLM “what are good flights to NYC” to which the response would be an answer with relevant accompanying information. An agentic query would look like “book me the best flight to NYC” to which the agent can account for calendar context, location, airline preferences, payment confirmation, and also look into hotel accommodations and provide travel summaries, etc. This is a very different token consumption profile.

Figure 15: Queries on agentic smartphones would consume more tokens due to the nature of tasks being executed

Regular query vs. agentic query



Source: BofA Global Research

BofA GLOBAL RESEARCH

Figure 16: We expect different types of queries executed on agentic smartphones to use varying levels of tokens

Types of queries on agentic smartphones

Request Type	Input tokens	Output tokens	Total tokens / request
Simple Siri / search-type request	~25-50	~50-150	~75-200
Standard chatbot Q&A (Siri app)	~100-300	~300-800	~400-1,100
Moderate agentic task	~500-1,500	~500-1,500	~1,000-3,000
Heavy agentic workflow	~2,000-8,000	~1,000-4,000	~3,000-12,000

Source: BofA Global Research

BofA GLOBAL RESEARCH

Using ranges from Figure 16 above, Figure 17 below shows what the token usage from a weighted average request on an agentic iPhone could look like initially as agentic Siri ramps. On an annual basis, we think simple Siri tasks and standard chatbot will comprise 60% of total requests (30% each), moderate agentic tasks to be 25%, and heavy agentic workflows to be 15%. However, we note this dynamic in Figure 17 below could gravitate towards heavier workflow mix over time as the complexity of request queries becomes higher.

Figure 17: We think a weighted average AI request on smartphone will use ~1,900 tokens

Mix assumptions on types of AI requests on agentic smartphones

Request Type	Input mid-point	Output mid-point	Sum	Assumed % of total requests
Simple Siri / search-type request	38	100	138	30%
Standard chatbot Q&A (Siri app)	200	550	750	30%
Moderate agentic task	1,000	1,000	2,000	25%
Heavy agentic workflow	5,000	2,500	7,500	15%
Weighted average request on agentic iPhone	1,071	820	1,891	

Source: BofA Global Research

BofA GLOBAL RESEARCH

We note that workflow orchestration itself is largely an OS / CPU function, but agentic workflows become AI-intensive because the system must repeatedly interpret intent, retrieve and rank context, reason through steps, and generate responses. As a result, not every step consumes model tokens, but the model-mediated parts of the workflow can still create meaningfully higher token demand than a traditional command-based smartphone interaction. See Figure 18 below for a step by step of a sample agentic query being executed on an agentic smartphone.

Figure 18: An agentic query of “find a dinner spot near my meeting tomorrow, book a 7pm reservation, and text Wamsi about it,” uses thousands of tokens

Hypothetical agentic query steps – sample of a moderate agentic task

Step	What is being done	Compute/ control layer	Token usage	Illustrative tokens
1. Capture user intent	User speaks or types the request	Microphone / speech model / OS	Sometimes	0-50
2. Parse intent	System identifies task: restaurant search + calendar context + booking + message	On-device model / NPU or cloud model	Yes	~120-280
3. Check permissions	OS checks whether Siri can access calendar, location, contacts, messages, booking apps, etc.	iOS permission layer / CPU	Usually no	NM
4. Retrieve context	Pulls relevant calendar event, location, contact, preferences, prior messages, dietary constraints	OS retrieval + model / embeddings	Yes	~300-1,000
5. Plan workflow	Agent decides sequence: identify meeting location → search restaurants → compare options → book → draft text	Model inference	Yes	~400-1,100
6. Call tools / apps	Calls Maps, Calendar, Messages, restaurant app, Safari, Wallet / Apple Pay, etc	OS / App Intents / APIs + model formatting	Sometimes	~100-600
7. Rank options	Compares restaurants by distance, availability, rating, user preferences, price, schedule fit	Model / embedding / ranking	Yes	~400-1,600
8. Ask for confirmation	“I found XYZ restaurant at 7pm near your meeting. Book and text Wamsi?”	Model inference + UI	Yes	~100-300
9. Authenticate / execute	Face ID / confirmation; booking is submitted; payment method may be authorized	Secure Enclave / OS / app API	Usually no	NM
10. Summarize / update state	Final answer: reservation booked, calendar updated, Wamsi notified	Model inference	Yes	~250-700
Total				~1670-5,630

Source: BofA Global Research

NM: not meaningful

BofA GLOBAL RESEARCH

Token analysis

To get a sense of how much tokens can be used on agentic iPhones in 2030, we use our assumptions for 2030 which is ~2.0bn installed base (including primary and secondary). For both base and high usage case, we keep the installed base same but adjust the mix of AI users within the installed base. We assume at the base case, about two thirds of users use AI, while in high usage case, it is closer to 85% that uses AI. This equates total AI iPhone users to be 1.4bn (base) and 1.7bn (high usage). If we assume average of 10

requests per day (double for high usage case of an average of 20 requests per day) and use ~2000 tokens per average request (~5000 for high usage case), this equates to 27.3tn tokens per day or 9.9qn tokens per year that is consumed through iPhone AI usage (172.9tn tokens per day & 63.1qn tokens per year). If half of queries are executed via cloud, this means 5.0qn tokens are cloud-based at our base case.

Figure 19: In our base case, iPhone AI usage could drive ~9.9qn tokens annually, with ~5.0qn processed through PCC

Token demand scenario on device vs. cloud

Scenario (2030E)	iPhone installed base (mn)	% of AI users	AI users	Avg. Requests / Day	Avg. Tokens / Request	Total Tokens / day	Total tokens / year	Cloud share	Cloud tokens / year
Base Case	2,034	67%	1,363	10	2,000	27.3tn	9.9qn	50%	5.0qn
High Usage Case	2,034	85%	1,729	20	5,000	172.9tn	63.1qn	50%	31.6qn

Source: BofA Global Research

Note: Installed base reflects our model estimates that including primary and secondary

BofA GLOBAL RESEARCH

Token usage per iPhone & implications on device processing power

From the analysis above, we take a closer look at the high usage case and what that implies for Apple's on-device processing requirements. 172.9tn tokens per day on an installed base of 1.7bn (iPhone AI users) implies each iPhone is consuming 100,000 tokens per day or 36.5mn tokens per year. Assuming 50% of tokens are processed locally, the on-device requirement would be ~50,000 tokens per active AI iPhone user per day (172.9tn tokens / 1.7bn AI iPhone users x 50%). Using a simple rule of thumb where 5 tokens/second feels slow, 20 tokens/second is usable, and 40 tokens/second feels fast, if an iPhone processes 20 tokens/second, this implies ~42 minutes of iPhone AI usage per active AI user per day. 50,000 tokens divided by 20 tokens per second = 2,500 seconds of processing per day which equates to ~41.7min of AI usage per day or ~2.1min per task (our high usage case assumes 20 average requests per day). By 2030, we expect Apple's silicon roadmap to cover more efficient NPUs, higher memory bandwidth, more SRAM/cache close to the compute engine, tighter model compression/quantization, and OS-level orchestration that can dynamically split workloads across CPU, GPU, NPU, and PCC. As these hardware developments occur, we see Apple's on-device AI moat increasingly expanding.

Token usage per iPhone & implications on cloud infrastructure needs

At our high usage case, we estimate about 31.6qn tokens will be processed via cloud. To understand what this means for Apple's data center capacity, we used known figures from today and assumed that all Apple servers run M5 Max (the most recent, highest tier of the M-series generation, excluding the Ultra line). For this calculation, we also assume this is all greenfield build-out (assuming Apple has no data center footprint today which it does as shown in Figure 20 below and that 100% of tokens will be processed on Apple data centers vs. none on third-party cloud).

31.6qn tokens annually equate to ~1.0bn tokens/second. If assuming M5 Max can generate ~100-120 token/second on smaller 8B-class quantized models and ~55-80 token/second on larger or heavier models, ~1.0bn token/second requirement implies a range of ~8-18mn M5 Max chips depending on the mix of models used (assume 100% utilization). To get to power per chip, we conservatively assume 60-100W under sustained LLM generation. At the low end of 8mn chips and 60W, we derive 480MW and at the high end of 18mn chips and 100W, we derive 1.8GW of IT load. If we apply a conservative 1.4 power usage effectiveness (PUE), that gives a range of ~0.7GW-2.5GW total facility power. Using a range of \$10-20mn/MW AI data center construction cost range (much lower for Apple using internal Apple Silicon vs having to pay for third party silicon), the implied infrastructure capex could be ~\$6.7-50.4bn. This number could have more variability and could be higher based on usage and overall costs but even at 2x of \$100bn at the high end, it equates to just one year of Apple's FCF and potentially spread across over time.

We note that Apple can accrue benefits as they make incremental advancements with their silicon performance and that Apple already has exiting data centers globally and has optionality to refresh hardware (which cost less than purchasing third party chips) inside which will cost less than greenfield buildout. Also, most token volume does not need to become third-party. Apple's design is explicitly hybrid by running what can be run locally on device, use Apple's Private Cloud Compute for more complex requests, and use external models selectively if or when needed. If Apple uses external models heavily, AI economics look more like a variable COGS model where usage scales, token payments scale, and Services gross margin compress unless Apple charges users, developers, or partners. If Apple uses PCC, the economics shift toward capex, depreciation, power, data center operations, and Apple silicon server costs. That can pressure near-term FCF conversion but preserve longer-term unit economics, privacy, and strategic control. The margin debate is whether Apple can charge for AI faster than AI costs rise. In a free-consumer-AI model with heavy third-party usage, AI can be a margin headwind. In a paid-tier / device-upgrade / transaction-take-rate model with mostly on-device and PCC inference, AI can become margin-accretive over time.

Figure 20: Apple operates data centers globally today, with additional capacity under development and/or via colocation / cloud providers

Apple DC global footprint



Source: Datacenter.com, BofA Global Research

BofA GLOBAL RESEARCH

Apple does not disclose data center nameplate MW capacity in its Environmental Progress Reports for each fiscal year. Apple discloses annual electricity consumption, which we can convert into an implied average MW load by dividing annual kWh by 8,760 hours. This is not the same as built or contracted power capacity as it represents the average power draw over the year. Based on Apple's disclosed data center + colocation electricity consumption, implied average load increased from ~268MW in F23 to ~286MW in F24 and ~295MW in F25 (see Figure 21 below for more details). Actual data center capacity is likely higher than this average-load proxy because facilities require utilization headroom, redundancy, maintenance buffers, and growth capacity.

Figure 21: In F25, Apple's data center locations consumed ~2.6bn kWh

Data center electricity consumption from Apple's annual Environmental Progress Report


Location - Data Center	Electricity (million kWh)- F23	Electricity (million kWh) - F24	Electricity (million kWh) - F25
Maiden, NC	453	466	470
Mesa, AZ	488	530	563
Newark, CA	0	0	0
Prineville, OR	269	255	267
Reno, NV	440	454	422
Viborg, Denmark	40		71
Waukegan, Iowa		59	30
Colo - US	387	422	423
Colo - International	96	105	121
China	171	214	218
Total	2,344	2,505	2,585

Source: Apple, BofA Global Research

BofA GLOBAL RESEARCH

Figure 22: We model AAPL capex stepping up to \$13.3bn in F27, up +23% y/y from \$10.8bn in F26.

AAPL capex trajectory has a potential upward bias



Source: Company reports, BofA Global Research estimates

BofA GLOBAL RESEARCH

App store agents: from app to actions

In an agentic world, Apple's App Store opportunity shifts from monetizing app downloads to monetizing which trusted app capabilities Siri calls to complete user intent

The App Store could evolve from a marketplace of apps users manually open into a marketplace of capabilities agents can call. Today, App Store monetization is largely tied to downloads, subscriptions, in-app purchases, search ads, and developer distribution. In an agentic world, Apple could monetize actions like booking, buying, subscribing, creating, searching, filing, sending, or completing a workflow on behalf of the user. The key issue is that agentic apps challenge the legacy App Store model.

Apple's existing App Store process is built around reviewing apps, updates, in-app purchases, and events for privacy, security, safety, and reliability, which becomes harder when software behavior is partly generated or delegated dynamically by an agent. The most important monetization shift is from app distribution to action distribution. In the current model, Apple monetizes by helping developers distribute apps to users. In an agentic model, the user may not search for an app at all; the user may ask Siri or another agent to complete a task. That means Apple's monetization surface can move upstream from "which app gets downloaded?" to "which app gets called by the agent?" In that world, App Store evolves into a structured registry of trusted app capabilities where apps are ranked by reliability, task completion rate, user satisfaction, privacy posture, and transaction quality.

Monetization Vector 1: Agent-mediated app discovery

The first monetization vector is agent-mediated discovery. Today, App Store search and rankings help users find apps. In an agentic interface, the assistant becomes the discovery layer. If a user asks, "book me a hotel near the meeting," Apple can influence which app or service receives that intent. This creates a new version of App Store search economics. Instead of bidding on keywords, developers may compete for intent categories or task surfaces. Apple could monetize this through promoted action placement, sponsored recommendations, or premium ranking tools, subject to disclosure and regulatory constraints. This would be analogous to App Store search ads, but more powerful because the agent is closer to the transaction.

Monetization Vector 2: App Intent certification

The second monetization vector is certification. If Apple allows agents to act across apps, it needs a trust framework. Not every app action should be equally callable by an agent. Apple could create a certification layer around agent-ready apps, similar in spirit to App Review but more continuous and behavior-based. Apple's value proposition to users would be: "These are agent-safe actions." Apple's value proposition to developers would be: "Certification makes your app more discoverable and callable by Siri / Apple Intelligence." Apple's value proposition to itself is that certification preserves the App Store's trusted environment while opening the door to more dynamic software behavior. This can become a monetization vector through developer program tiers, certification fees, enhanced analytics, enterprise-grade trust tooling, or higher take-rates on agent-mediated transactions. Apple would need to be careful, because over-monetizing certification could create developer pushback. But if certification materially improves conversion, discovery, and safety, developers may view it as a distribution cost rather than a tax.

Monetization Vector 3: Transaction economics on completed tasks

If an agent completes a transaction, Apple can participate in the economics in several ways:

- Digital goods / subscriptions: existing App Store commission structures can continue to apply where relevant.
- Physical commerce: Apple could capture referral fees, affiliate economics, or Apple Pay-related economics.
- Services bookings: Apple could monetize through lead generation or completion fees.

- Enterprise workflows: Apple could monetize through management tools, secure app actions, or device-management integration.
- Payments: Apple Pay / Wallet can become the confirmation and settlement layer.

The key change is that Apple may monetize outcomes beyond just distribution. If Siri routes a user to an app and the app completes a booking, subscription, or purchase, Apple can argue that the App Store / Siri ecosystem contributed the demand and reduced friction. That gives Apple a stronger basis to negotiate economics with developers, merchants, and service providers. This is also why developers may be cautious. Some developers are concerned that integrating app capabilities into Siri could create new ways for Apple to collect commissions. That concern is rational because the agent layer could become a new toll road between user intent and developer fulfillment.

Monetization Vector 4: App Store = governance layer for agents

The fourth vector is governance. Agentic software introduces risks that standard app review was not designed to handle. An app that behaves deterministically at submission is easier to review. An app with an embedded agent that can plan, call tools, generate code, or change workflow behavior is much harder to police. There could be tension between AI agents' complex autonomous behavior and Apple's strict App Store rules around malware, fee avoidance, and app behavior after review. Apple can turn that challenge into a monetization advantage. The company can position the App Store as the safest environment for consumer agents by requiring:

- constrained tool permissions
- sandboxed execution
- user-visible confirmations
- limits on code generation or dynamic mini-app behavior
- auditable action logs
- payment and identity checks
- clear developer accountability

In this model, Apple is productizing a safer version of agentic behavior. That is valuable because users and regulators will likely demand guardrails before agents can act across personal data, payments, and apps.

Monetization Vector 5: Agent analytics and developer tools

The fifth vector is developer tooling. If apps become agent-callable, developers will need a new analytics stack. Traditional App Store analytics answer questions like downloads, conversion, retention, and subscription revenue. Agentic App Store analytics would answer different questions:

- How often is my app called by Siri / Apple Intelligence?
- Which intents drive the most calls?
- Where does the agent abandon the workflow?

- Which permissions reduce conversion?
- Which confirmations create friction?
- Which app actions complete successfully?
- How does my app rank for a given intent category?

Apple could monetize this through App Store Connect tools, premium analytics, testing environments, agent simulation, developer APIs, and enterprise support. This would be similar to how cloud platforms monetize developer infrastructure: the more developers build for the agentic App Store, the more valuable Apple's tooling becomes.

Monetization Vector 6: Model and agent marketplace economics

Over time, Apple could also create a broader marketplace around agent providers. If users can choose external models or agent engines, Apple could become the distribution layer for those providers. The App Store already gives Apple the billing relationship, review process, subscription infrastructure, and developer distribution model. Agentic AI could extend that into a marketplace where model providers, app developers, and workflow providers compete to serve user intents:

- which agents can access which app actions
- which models can process which data
- which tasks require Apple-controlled confirmation
- which transactions use Apple Pay
- how developers are paid
- how disputes are handled

This would be strategically consistent with Apple's historical model where Apple needs to own distribution, trust, payments, review, and the user relationship. In agentic AI, Apple needs to own the marketplace where models and apps become useful to users.

Services model implication

This could be a meaningful Services opportunity because agentic App Store monetization builds on Apple's existing strengths. Apple already monetizes developer distribution, billing, subscriptions, payments, and installed-base engagement. Agents could increase the frequency and value of those interactions by reducing friction between user intent and task completion. The current App Store is optimized around apps as destinations. The future App Store can be optimized around apps as capabilities. If Apple executes well, the App Store becomes more relevant in an agentic world.

For Apple to turn agents into an App Store monetization vector, it likely needs to do five things:

- Standardize app actions so developers can expose agent-callable capabilities in a structured way.
- Create a trust and certification layer so users know which actions are safe for agents to perform.

- Build agent discovery economics so developers can compete for intent categories, not just app downloads.
- Protect user control through confirmations, authentication, permission scopes, and transparent logs.
- Avoid over-taxing early adoption so developers are incentivized to expose app capabilities rather than resist integration.

The last point is important. Apple's long-term monetization opportunity depends on developer participation. If developers view agent integration as another commission layer, adoption could be slow. If developers view Apple's agent framework as a new distribution surface that drives incremental demand, adoption could be meaningful. Apple's optimal strategy may be to keep early commissions low or zero, prove that Siri / Apple Intelligence can drive incremental conversion, and then layer in monetization once the agentic App Store has demonstrated clear value.

2030 incremental revenue analysis

Agentic Siri could become a meaningful new Services revenue layer by monetizing paid AI tiers, model routing, app-action commerce, payments, advertising, and iCloud+ AI bundles, with our illustrative framework implying ~\$15-30bn of incremental F30E revenue in the base case and ~\$40-65bn in the bull case after overlap haircuts

We see six concrete monetization paths for an agentic Siri, each independently plausible and most additive. We do not view these revenue vectors as fully additive, given overlap across paid AI tiers, iCloud bundles, model routing, commerce, payments, and advertising. After applying a 20-25% overlap haircut in the base case and a 25-30% haircut in the bull case, we estimate agentic Siri could represent a ~\$15-30bn incremental F30E revenue opportunity in the base case and ~\$40-65bn in the bull case.

Figure 23: In our base case, we expect by 2030, agentic Siri can add \$15-30bn of incremental revenue through six new monetization vectors

Monetization vectors for agentic Siri

Monetization vector	Base case assumption	Base case rev oppty	Bull case assumption	Bull case rev oppty
Apple Intelligence Pro subscription (\$15/mo)	50-75mn subs × \$15/mo × 12	\$9-14bn	100-125mn subs × \$15/mo × 12	\$18-23bn
Default model placement / routing fee	1.36bn AI users × \$0.25-0.50/mo × 12	\$4-8bn	1.73bn AI users × \$0.50-0.75/mo × 12	\$10-16bn
App Store agent commerce take rate	\$200-300bn agent GMV × 1-2% take rate	\$2-6bn	\$400-600bn agent GMV × 2-3% take rate	\$8-18bn
Apple Pay / Wallet routing fees	\$1-2tn payment volume × 15bps	\$2-3bn	\$3-5tn payment volume × 15bps	\$5-8bn
Advertising in agent results	1.36bn AI users × 0.05-0.075 sponsored results/day × \$0.10-0.12 × 365	\$2-5bn	1.73bn AI users × 0.10-0.15 sponsored results/day × \$0.15-0.18 × 365	\$9-17bn
iCloud+ AI tier	50-75mn subs × \$3-5/mo × 12	\$2-5bn	100-150mn subs × \$5-8/mo × 12	\$6-14bn
Gross rev opportunity	Sum before overlap	\$21-40bn	Sum before overlap	\$56-95bn
Overlap haircut	20-25% haircut	(\$4-10bn)	25-30% haircut	(\$14-29bn)
Net incremental 2030E rev	After haircut	~\$15-30bn	After haircut	~\$40-65bn

Source: BofA Global Research

BofA GLOBAL RESEARCH

- **Apple Intelligence Pro subscription:** We assume Apple launches a paid premium AI tier for higher-intensity agentic workflows, longer context, advanced model access, and privacy-preserving compute. We estimate the opportunity using subscriber count × monthly price × 12 months. In the base case, we assume 50-75mn subscribers at \$15/month, implying ~\$9-14bn of FY30E revenue. In the bull case, we assume 100-125mn subscribers at \$15/month, implying ~\$18-23bn.
- **Default model placement / routing fee:** We assume external model providers pay Apple for preferred placement, default routing, or distribution within Siri / Apple Intelligence, similar in concept to search distribution economics. We estimate this using AI iPhone users × monthly model-routing economics × 12 months. In the base case, we assume ~1.36bn AI users × \$0.25-0.50/month, implying ~\$4-8bn. In the bull case, we assume ~1.73bn AI users × \$0.50-0.75/month, implying ~\$10-16bn.
- **App Store agent commerce take rate:** We assume Siri / App Intents route user intent into app-mediated transactions across commerce, travel, subscriptions, digital goods, services, and other workflows. We estimate this using agent-mediated GMV × Apple effective take rate. In the base case, we assume \$200-300bn of agent-mediated GMV × 1-2% take rate, implying ~\$2-6bn. In the bull case, we assume \$400-600bn of GMV × 2-3% take rate, implying ~\$8-18bn.
- **Apple Pay / Wallet routing fees:** We assume agent-initiated purchases increasingly flow through Apple Pay / Wallet as the trusted confirmation and payment layer. We estimate this using agent-initiated payment volume × ~15bps effective fee. In the base case, we assume \$1-2tn of payment volume × 15bps, implying ~\$2-3bn. In the bull case, we assume \$3-5tn × 15bps, implying ~\$5-8bn.
- **Advertising in agent results:** We assume Apple extends search / advertising economics into disclosed sponsored placements or commerce recommendations within Siri / agent results. We estimate this using AI users × sponsored agent results per day × revenue per sponsored result × 365 days. In the base case, we assume ~1.36bn AI users × 0.05-0.075 sponsored results/day × \$0.10-0.12, implying ~\$2-5bn. In the bull case, we assume ~1.73bn AI users × 0.10-0.15 sponsored results/day × \$0.15-0.18, implying ~\$9-17bn.

- iCloud+ AI tier: We assume Apple bundles personal context, storage, memory, retrieval, and privacy-preserving AI features into a higher-value iCloud+ / AI tier. We estimate this using subscribers $\times$ incremental monthly uplift $\times$ 12 months. In the base case, we assume 50-75mn subscribers $\times$ \$3-5/month, implying ~\$2-5bn. In the bull case, we assume 100-150mn subscribers $\times$ \$5-8/month, implying ~\$6-14bn.
- Overlap haircut: These vectors are not fully additive because paid AI tiers, iCloud bundles, model routing, app commerce, payments, and advertising may overlap. We therefore apply a 20-25% haircut in the base case and a 25-30% haircut in the bull case to avoid double counting.
- Net F30E revenue opportunity: Before overlap, the six vectors imply ~\$21-40bn of gross F30E revenue opportunity in the base case and ~\$56-95bn in the bull case. After the overlap haircut, this equates to ~\$15-30bn of net incremental F30E revenue opportunity in the base case and ~\$40-65bn in the bull case.

Abbreviations

AI: Artificial Intelligence

API: Application Programming Interface

AR: Augmented Reality

ASP: Average Selling Price

bps: Basis points

Capex: Capital expenditure

COGS: Cost of goods sold

CPU: Central Processing Unit

DC: Data Center

DRAM: Dynamic Random-Access Memory

FCF: Free Cash Flow

GM: Gross Margin

GMV: Gross Merchandise Value

GPU: Graphics Processing Unit

GW: Gigawatt

HW: Hardware

IB: Installed Base

IDC: International Data Corporation

iOS: Apple mobile operating system

IT: Information Technology

kWh: Kilowatt-hour

LLM: Large Language Model

mn: Million

mo: Month

MW: Megawatt

NAS: Nasdaq exchange

NC: North Carolina

NPU: Neural Processing Unit

NV: Nevada

NYC: New York City

OEM: Original Equipment Manufacturer

OR: Oregon

OS: Operating System

PCC: Private Cloud Compute

PUE: Power Usage Effectiveness

QQQ: Nasdaq-100 ETF

RAM: Random Access Memory

ROE: Return on Equity

SoC: System on Chip

SPX: Standard & Poor's 500 Index

SRAM: Static Random-Access Memory

TAC: Traffic Acquisition Cost

tn: Trillion

VR: Virtual Reality

WWDC: Worldwide Developers Conference

Price objective basis & risk

Apple Inc. (AAPL)

Our PO of \$380 is based on 37x our C27E EPS of \$10.29. Our target multiple compares to the 5-year historical range of 19-35x (median 27x). We believe a multiple above the higher end of the historical range is justified given the significant opportunity in agentic AI as well as a multi-year upgrade cycle, large cash balance and opportunity to diversify into new end markets, improving mix and diversity of services.

Downside risks are a) weaker iPhone cycle on consumer spending risk, b) weaker near-term services trajectory where App Store & Licensing can decelerate, c) execution on agentic Siri, d) iPads/Macs reverting to pre-COVID levels, e) stronger dollar, f) antitrust lawsuit. Other risks are negative impact to licensing revenues from competing AI related search, potential trade conflicts, tariffs, longer iPhone replacement cycles, commoditization in the smartphone market, intensifying competition in the tablet market, ability to manage beat and raise expectations for EPS estimates, and requirement to maintain pace of product innovation.

Upside risks are a) stronger sales of Pro iPhone models which can help average selling price and help the mix of the business, b) est. revisions can be lower vs. peers, c) potential new products (AR/VR), and services (advertising). Other upside risks are stronger than expected iPhone cycle, gross margin upside from better mix of higher end iPhones, tailwinds from lower memory costs, and a faster than expected recovery in emerging markets.

Analyst Certification

I, Wamsi Mohan, hereby certify that the views expressed in this research report accurately reflect my personal views about the subject securities and issuers. I also certify that no part of my compensation was, is, or will be, directly or indirectly, related to the specific recommendations or view expressed in this research report.

US - IT Hardware and Technology Supply Chain Coverage Cluster

Investment rating	Company	BofA Ticker	Bloomberg symbol	Analyst
BUY				
	Amphenol	APH	APH US	Wamsi Mohan
	Apple Inc.	AAPL	AAPL US	Wamsi Mohan
	Celestica Inc.	CLS	CLS US	Ruplu Bhattacharya
	Celestica Inc.	YCLS	CLS CN	Ruplu Bhattacharya
	Corning Inc.	GLW	GLW US	Wamsi Mohan
	Dell Technologies Inc.	DELL	DELL US	Wamsi Mohan
	DigitalOcean	DOCN	DOCN US	Wamsi Mohan
	Flex Ltd.	FLEX	FLEX US	Ruplu Bhattacharya
	Hewlett-Packard Enterprise	HPE	HPE US	Wamsi Mohan
	Ingram Micro Holding Corporation	INGM	INGM US	Ruplu Bhattacharya
	International Business Machines Corp.	IBM	IBM US	Wamsi Mohan
	Jabil Inc.	JBL	JBL US	Ruplu Bhattacharya
	Nutanix Inc	NTNX	NTNX US	Wamsi Mohan
	Sandisk Corporation	SNDK	SNDK US	Wamsi Mohan
	Seagate Technology	STX	STX US	Wamsi Mohan
	TD Synnex Corp	SNX	SNX US	Ruplu Bhattacharya
	TE Connectivity Plc.	TEL	TEL US	Wamsi Mohan
	Western Digital Corporation	WDC	WDC US	Wamsi Mohan
NEUTRAL				
	Arrow Electronics Inc.	ARW	ARW US	Ruplu Bhattacharya
	Avnet Inc.	AVT	AVT US	Ruplu Bhattacharya
	CDW Corp	CDW	CDW US	Ruplu Bhattacharya
	Concentrix Corporation	CNXC	CNXC US	Ruplu Bhattacharya
	Enovix Corporation	ENVX	ENVX US	Ruplu Bhattacharya
	Everpure Inc	P	P US	Wamsi Mohan

US - IT Hardware and Technology Supply Chain Coverage Cluster

Investment rating	Company	BofA Ticker	Bloomberg symbol	Analyst
	NetApp Inc.	NTAP	NTAP US	Wamsi Mohan
	Sanmina Corporation	SANM	SANM US	Ruplu Bhattacharya
	Sensata Technologies Holdings Plc	ST	ST US	Wamsi Mohan

UNDERPERFORM

	HP Inc.	HPQ	HPQ US	Wamsi Mohan
	Super Micro Computer Inc.	SMCI	SMCI US	Ruplu Bhattacharya
	Teradata Corporation	TDC	TDC US	Wamsi Mohan
	Vishay Intertechnology, Inc.	VSH	VSH US	Ruplu Bhattacharya

iQmethodSM Measures Definitions**Business Performance**

Return On Capital Employed

Return On Equity

Operating Margin

Earnings Growth

Free Cash Flow

Quality of Earnings

Cash Realization Ratio

Asset Replacement Ratio

Tax Rate

Net Debt-To-Equity Ratio

Interest Cover

Valuation Toolkit

Price / Earnings Ratio

Price / Book Value

Dividend Yield

Free Cash Flow Yield

Enterprise Value / Sales

EV / EBITDA

Numerator

NOPAT = (EBIT + Interest Income) × (1 – Tax Rate) + Goodwill Amortization

Net Income

Operating Profit

Expected 5 Year CAGR From Latest Actual

Cash Flow From Operations – Total Capex

Numerator

Cash Flow From Operations

Capex

Tax Charge

Net Debt = Total Debt – Cash & Equivalents

EBIT

Numerator

Current Share Price

Current Share Price

Annualised Declared Cash Dividend

Cash Flow From Operations – Total Capex

EV = Current Share Price × Current Shares + Minority Equity + Net Debt +

Other LT Liabilities

Enterprise Value

Denominator

Total Assets – Current Liabilities + ST Debt + Accumulated Goodwill

Amortization

Shareholders' Equity

Sales

N/A

N/A

Denominator

Net Income

Depreciation

Pre-Tax Income

Total Equity

Interest Expense

Denominator

Diluted Earnings Per Share (Basis As Specified)

Shareholders' Equity / Current Basic Shares

Current Share Price

Market Cap = Current Share Price × Current Basic Shares

Sales

Basic EBIT + Depreciation + Amortization

iQmethodSM is the set of BofA Global Research standard measures that serve to maintain global consistency under three broad headings: Business Performance, Quality of Earnings, and validations. The key features of iQmethod are: A consistently structured, detailed, and transparent methodology. Guidelines to maximize the effectiveness of the comparative valuation process, and to identify some common pitfalls.


iQdatabase[®] is our real-time global research database that is sourced directly from our equity analysts' earnings models and includes forecasted as well as historical data for income statements, balance sheets, and cash flow statements for companies covered by BofA Global Research.

iQprofileSM, **iQmethodSM** are service marks of Bank of America Corporation. **iQdatabase[®]** is a registered service mark of Bank of America Corporation.

Disclosures

Important Disclosures

Apple Inc. (AAPL) Price Chart



B: Buy, N: Neutral, U: Underperform, PO: Price Objective, NA: No longer valid, NR: No Rating

The Investment Opinion System is contained at the end of the report under the heading "Fundamental Equity Opinion Key". Dark grey shading indicates the security is restricted with the opinion suspended. Medium grey shading indicates the security is under review with the opinion withdrawn. Light grey shading indicates the security is not covered. Chart is current as of a date no more than one trading day prior to the date of the report.

Equity Investment Rating Distribution: Technology Group (as of 31 Mar 2026)

Coverage Universe	Count	Percent	Inv. Banking Relationships ^{R1}	Count	Percent
Buy	234	58.50%	Buy	123	52.56%
Hold	90	22.50%	Hold	43	47.78%
Sell	76	19.00%	Sell	23	30.26%

Equity Investment Rating Distribution: Global Group (as of 31 Mar 2026)

Coverage Universe	Count	Percent	Inv. Banking Relationships ^{R1}	Count	Percent
Buy	1993	55.76%	Buy	1186	59.51%
Hold	821	22.97%	Hold	509	62.00%
Sell	760	21.26%	Sell	400	52.63%

^{R1} Issuers that were investment banking clients of BofA Securities or one of its affiliates within the past 12 months. For purposes of this Investment Rating Distribution, the coverage universe includes only stocks. A stock rated Neutral is included as a Hold, and a stock rated Underperform is included as a Sell.

FUNDAMENTAL EQUITY OPINION KEY: Opinions include a Volatility Risk Rating, an Investment Rating and an Income Rating. VOLATILITY RISK RATINGS, indicators of potential price fluctuation, are: A - Low, B - Medium and C - High. INVESTMENT RATINGS reflect the analyst’s assessment of both a stock’s absolute total return potential as well as its attractiveness for investment relative to other stocks within its Coverage Cluster (defined below). Our investment ratings are: 1 - Buy stocks are expected to have a total return of at least 10% and are the most attractive stocks in the coverage cluster; 2 - Neutral stocks are expected to remain flat or increase in value and are less attractive than Buy rated stocks and 3 - Underperform stocks are the least attractive stocks in a coverage cluster. An investment rating of 6 (No Rating) indicates that a stock is no longer trading on the basis of fundamentals. Analysts assign investment ratings considering, among other things, the 0-12 month total return expectation for a stock and the firm’s guidelines for ratings dispersions (shown in the table below). The current price objective for a stock should be referenced to better understand the total return expectation at any given time. The price objective reflects the analyst’s view of the potential price appreciation (depreciation).

Investment rating	Total return expectation (within 12-month period of date of initial rating)	Ratings dispersion guidelines for coverage cluster ^{R2}
Buy	≥ 10%	≤ 70%
Neutral	≥ 0%	≤ 30%
Underperform	N/A	≥ 20%

^{R2} Ratings dispersions may vary from time to time where BofA Global Research believes it better reflects the investment prospects of stocks in a Coverage Cluster.

INCOME RATINGS, indicators of potential cash dividends, are: 7 - same/higher (dividend considered to be secure), 8 - same/lower (dividend not considered to be secure) and 9 - pays no cash dividend. Coverage Cluster is comprised of stocks covered by a single analyst or two or more analysts sharing a common industry, sector, region or other classification(s). A stock’s coverage cluster is included in the most recent BofA Global Research report referencing the stock.

Price Charts for the securities referenced in this research report are available on the [Price Charts website](#), or call 1-800-MERRILL to have them mailed.

BofAS or one of its affiliates acts as a market maker for the equity securities recommended in the report: Apple Inc.

The issuer is or was, within the last 12 months, an investment banking client of BofAS and/or one or more of its affiliates: Apple Inc.

BofAS or an affiliate has received compensation from the issuer for non-investment banking services or products within the past 12 months: Apple Inc.

The issuer is or was, within the last 12 months, a non-securities business client of BofAS and/or one or more of its affiliates: Apple Inc.

An officer, director or employee of BofAS or one of its affiliates is an officer or director of this issuer: Apple Inc.

BofAS or an affiliate has received compensation for investment banking services from this issuer within the past 12 months: Apple Inc.

BofAS or an affiliate expects to receive or intends to seek compensation for investment banking services from this issuer or an affiliate of the issuer within the next three months: Apple Inc.

BofAS or one of its affiliates is willing to sell to, or buy from, clients the common equity of the issuer on a principal basis: Apple Inc.

The issuer is or was, within the last 12 months, a securities business client (non-investment banking) of BofAS and/or one or more of its affiliates: Apple Inc.

BofA Global Research personnel (including the analyst(s) responsible for this report) receive compensation based upon, among other factors, the overall profitability of Bank of America Corporation, including profits derived from investment banking. The analyst(s) responsible for this report may also receive compensation based upon, among other factors, the overall profitability of the Bank's sales and trading businesses relating to the class of securities or financial instruments for which such analyst is responsible.

Other Important Disclosures

From time to time research analysts conduct site visits of covered issuers. BofA Global Research policies prohibit research analysts from accepting payment or reimbursement for travel expenses from the issuer for such visits.

Prices are indicative and for information purposes only. Except as otherwise stated in the report, for any recommendation in relation to an equity security, the price referenced is the publicly traded price of the security as of close of business on the day prior to the date of the report or, if the report is published during intraday trading, the price referenced is indicative of the traded price as of the date and time of the report and in relation to a debt security (including equity preferred and CDS), prices are indicative as of the date and time of the report and are from various sources including BofA Securities trading desks.

The date and time of completion of the production of any recommendation in this report shall be the date and time of dissemination of this report as recorded in the report timestamp.

Recipients who are not institutional investors or market professionals should seek the advice of their independent financial advisor before considering information in this report in connection with any investment decision, or for a necessary explanation of its contents.

Officers of BofA or one or more of its affiliates (other than research analysts) may have a financial interest in securities of the issuer(s) or in related investments.

Refer to [BofA Global Research policies relating to conflicts of interest](#).

"BofA Securities" includes BofA Securities, Inc. ("BofAS") and its affiliates. Investors should contact their BofA Securities representative or Merrill Global Wealth Management financial advisor if they have questions concerning this report or concerning the appropriateness of any investment idea described herein for such investor. "BofA Securities" is a global brand for BofA Global Research.

Information relating to Non-US affiliates of BofA Securities and Distribution of Affiliate Research Reports:

BofAS and/or Merrill Lynch, Pierce, Fenner & Smith Incorporated ("MLPF&S") may in the future distribute, information of the following non-US affiliates in the US (short name: legal name, regulator): Merrill Lynch (South Africa): Merrill Lynch South Africa (Pty) Ltd., regulated by the Financial Sector Conduct Authority; MLI (UK): Merrill Lynch International, regulated by the Financial Conduct Authority (FCA) and the Prudential Regulation Authority (PRA); BofASE (France): BofA Securities Europe SA is authorized by the Autorité de Contrôle Prudentiel et de Résolution (ACPR) and regulated by the ACPR and the Autorité des Marchés Financiers (AMF). BofA Securities Europe SA ("BofASE") with registered address at 51, rue La Boétie, 75008 Paris is registered under no 842 602 690 RCS Paris. In accordance with the provisions of French Code Monétaire et Financier (Monetary and Financial Code), BofASE is an établissement de crédit et d'investissement (credit and investment institution) that is authorised and supervised by the European Central Bank and the Autorité de Contrôle Prudentiel et de Résolution (ACPR) and regulated by the ACPR and the Autorité des Marchés Financiers. BofASE's share capital can be found at www.bofam.com/BofASEdisclaimer; BofA Europe (Milan): Bank of America Europe Designated Activity Company, Milan Branch, regulated by the Bank of Italy, the European Central Bank (ECB) and the Central Bank of Ireland (CBI); BofA Europe (Frankfurt): Bank of America Europe Designated Activity Company, Frankfurt Branch regulated by BaFin, the ECB and the CBI; BofA Europe (Zurich): Bank of America Europe Designated Activity Company, Zurich Branch, regulated by the Swiss Financial Market Supervisory Authority FINMA, the ECB and CBI; BofA Europe (Madrid): Bank of America Europe Designated Activity Company, Sucursal en España, regulated by the Bank of Spain, the ECB and the CBI; Merrill Lynch (Australia): Merrill Lynch Equities (Australia) Limited, regulated by the Australian Securities and Investments Commission; Merrill Lynch (Hong Kong): Merrill Lynch (Asia Pacific) Limited, regulated by the Hong Kong Securities and Futures Commission (HKSF); Merrill Lynch (Singapore): Merrill Lynch (Singapore) Pte Ltd, regulated by the Monetary Authority of Singapore (MAS); Merrill Lynch (Canada): Merrill Lynch Canada Inc, regulated by the Canadian Investment Regulatory Organization; Merrill Lynch (Mexico): Merrill Lynch Mexico, SA de CV, Casa de Bolsa, regulated by the Comisión Nacional Bancaria y de Valores; BofAS Japan: BofA Securities Japan Co., Ltd., regulated by the Financial Services Agency; Merrill Lynch (Seoul): Merrill Lynch International, LLC Seoul Branch, regulated by the Financial Supervisory Service; Merrill Lynch (Taiwan): Merrill Lynch Securities (Taiwan) Ltd., regulated by the Securities and Futures Bureau; BofAS India: BofA Securities India Limited, regulated by the Securities and Exchange Board of India (SEBI); Merrill Lynch (Israel): Merrill Lynch Israel Limited, regulated by Israel Securities Authority; Merrill Lynch (DIFC): Merrill Lynch International (DIFC Branch), regulated by the Dubai Financial Services Authority (DFSA); Merrill Lynch (Brazil): Merrill Lynch S.A. Corretora de Títulos e Valores Mobiliários, regulated by Comissão de Valores Mobiliários; Merrill Lynch KSA Company: Merrill Lynch Kingdom of Saudi Arabia Company, regulated by the Capital Market Authority. This information: has been approved for publication and is distributed in the United Kingdom (UK) to professional clients and eligible counterparties (as each is defined in the rules of the FCA and the PRA) by MLI (UK), which is authorized by the PRA and regulated by the FCA and the PRA - details about the extent of our regulation by the FCA and PRA are available from us on request; has been approved for publication and is distributed in the European Economic Area (EEA) by BofASE (France), which is authorized by the ACPR and regulated by the ACPR and the AMF; has been considered and distributed in Japan by BofAS Japan, a registered securities dealer under the Financial Instruments and Exchange Act in Japan, or its permitted affiliates; is issued and distributed in Hong Kong by Merrill Lynch (Hong Kong) which is regulated by HKSF; is issued and distributed in Taiwan by Merrill Lynch (Taiwan); is issued and distributed in India by BofAS India; and is issued and distributed in Singapore to institutional investors and/or accredited investors (each as defined under the Financial Advisers Regulations) by Merrill Lynch (Singapore) (Company Registration No 198602883D). Merrill Lynch (Singapore) is regulated by MAS. Merrill Lynch Equities (Australia) Limited (ABN 65 006 276 795), AFS License 235132 (MLEA) distributes this information in Australia only to 'Wholesale' clients as defined by s.761G of the Corporations Act 2001. With the exception of Bank of America N.A., Australia Branch, neither MLEA nor any of its affiliates involved in preparing this information is an Authorised Deposit-Taking Institution under the Banking Act 1959 nor regulated by the Australian Prudential Regulation Authority. No approval is required for publication or distribution of this information in Brazil and its local distribution is by Merrill Lynch (Brazil) in accordance with applicable regulations. Merrill Lynch (DIFC) is authorized and regulated by the DFSA. Information prepared and issued by Merrill Lynch (DIFC) is done so in accordance with the requirements of the DFSA conduct of business rules. BofA Europe (Frankfurt) distributes this information in Germany and is regulated by BaFin, the ECB and the CBI. BofA Securities entities, including BofA Europe and BofASE (France), may outsource/delegate the marketing and/or provision of certain research services or aspects of research services to other branches or members of the BofA Securities group. You may be contacted by a different BofA Securities entity acting for and on behalf of your service provider where permitted by applicable law. This does not change your service provider. Please refer to the [Electronic Communications Disclaimers](#) for further information.

This information has been prepared and issued by BofAS and/or one or more of its non-US affiliates. The author(s) of this information may not be licensed to carry on regulated activities in your jurisdiction and, if not licensed, do not hold themselves out as being able to do so. BofAS and/or MLPF&S is the distributor of this information in the US and accepts full responsibility for information distributed to BofAS and/or MLPF&S clients in the US by its non-US affiliates. Any US person receiving this information and wishing to effect any transaction in any security discussed herein should do so through BofAS and/or MLPF&S and not such foreign affiliates. Hong Kong recipients of this information should contact Merrill Lynch (Asia Pacific) Limited in respect of any matters relating to dealing in securities or provision of specific advice on securities or any other matters arising from, or in connection with, this information. Singapore recipients of this information should contact Merrill Lynch (Singapore) Pte Ltd in respect of any matters arising from, or in connection with, this information. For clients that are not accredited investors, expert investors or institutional investors Merrill Lynch (Singapore) Pte Ltd accepts full responsibility for the contents of this information distributed to such clients in Singapore.

General Investment Related Disclosures:

Taiwan Readers: Neither the information nor any opinion expressed herein constitutes an offer or a solicitation of an offer to transact in any securities or other financial instrument. No part of this report may be used or reproduced or quoted in any manner whatsoever in Taiwan by the press or any other person without the express written consent of BofA Securities.

This document provides general information only, and has been prepared for, and is intended for general distribution to, BofA Securities clients. Neither the information nor any opinion expressed constitutes an offer or an invitation to make an offer, to buy or sell any securities or other financial instrument or any derivative related to such securities or instruments (e.g., options, futures, warrants, and contracts for differences). This document is not intended to provide personal investment advice and it does not take into account the specific investment objectives, financial situation and the particular needs of, and is not directed to, any specific person(s). This document and its content do not constitute, and should not be considered to constitute, investment advice for purposes of ERISA, the US tax code, the Investment Advisers Act or otherwise. Investors should seek financial advice regarding the appropriateness of investing in financial instruments and implementing investment strategies discussed or recommended in this document and should understand that statements regarding future prospects may not be realized. Any decision to purchase or subscribe for securities in any offering must be based solely on existing public information on such security or the information in the prospectus or other offering document issued in connection with such offering, and not on this document.

Securities and other financial instruments referred to herein, or recommended, offered or sold by BofA Securities, are not insured by the Federal Deposit Insurance Corporation and are not deposits or other obligations of any insured depository institution (including, Bank of America, N.A.). Investments in general and, derivatives, in particular, involve numerous risks, including,

among others, market risk, counterparty default risk and liquidity risk. No security, financial instrument or derivative is suitable for all investors. Digital assets are extremely speculative, volatile and are largely unregulated. In some cases, securities and other financial instruments may be difficult to value or sell and reliable information about the value or risks related to the security or financial instrument may be difficult to obtain. Investors should note that income from such securities and other financial instruments, if any, may fluctuate and that price or value of such securities and instruments may rise or fall and, in some cases, investors may lose their entire principal investment. Past performance is not necessarily a guide to future performance. Levels and basis for taxation may change.

This report may contain a short-term trading idea or recommendation, which highlights a specific near-term catalyst or event impacting the issuer or the market that is anticipated to have a short-term price impact on the equity securities of the issuer. Short-term trading ideas and recommendations are different from and do not affect a stock's fundamental equity rating, which reflects both a longer term total return expectation and attractiveness for investment relative to other stocks within its Coverage Cluster. Short-term trading ideas and recommendations may be more or less positive than a stock's fundamental equity rating.

BofA Securities is aware that the implementation of the ideas expressed in this report may depend upon an investor's ability to "short" securities or other financial instruments and that such action may be limited by regulations prohibiting or restricting "shortselling" in many jurisdictions. Investors are urged to seek advice regarding the applicability of such regulations prior to executing any short idea contained in this report.

Foreign currency rates of exchange may adversely affect the value, price or income of any security or financial instrument mentioned herein. Investors in such securities and instruments, including ADRs, effectively assume currency risk.

BofAS or one of its affiliates is a regular issuer of traded financial instruments linked to securities that may have been recommended in this report. BofAS or one of its affiliates may, at any time, hold a trading position (long or short) in the securities and financial instruments discussed in this report.

BofA Securities, through business units other than BofA Global Research, may have issued and may in the future issue trading ideas or recommendations that are inconsistent with, and reach different conclusions from, the information presented herein. Such ideas or recommendations may reflect different time frames, assumptions, views and analytical methods of the persons who prepared them, and BofA Securities is under no obligation to ensure that such other trading ideas or recommendations are brought to the attention of any recipient of this information.

In the event that the recipient received this information pursuant to a contract between the recipient and BofAS for the provision of research services for a separate fee, and in connection therewith BofAS may be deemed to be acting as an investment adviser, such status relates, if at all, solely to the person with whom BofAS has contracted directly and does not extend beyond the delivery of this report (unless otherwise agreed specifically in writing by BofAS). If such recipient uses the services of BofAS in connection with the sale or purchase of a security referred to herein, BofAS may act as principal for its own account or as agent for another person. BofAS is and continues to act solely as a broker-dealer in connection with the execution of any transactions, including transactions in any securities referred to herein.

Copyright and General Information:

Copyright 2026 Bank of America Corporation. All rights reserved. iQDatabase® is a registered service mark of Bank of America Corporation. This information is prepared for the use of BofA Securities clients and may not be redistributed, retransmitted or disclosed, in whole or in part, or in any form or manner, without the express written consent of BofA Securities. This document and its content is provided solely for informational purposes and cannot be used for training or developing artificial intelligence (AI) models or as an input in any AI application (collectively, an AI tool). Any attempt to utilize this document or any of its content in connection with an AI tool without explicit written permission from BofA Global Research is strictly prohibited. BofA Global Research utilizes AI, including machine learning and other technologies, to enhance the services we provide to our clients. These technologies assist our analysts in various aspects of their work, including but not limited to data analysis, content extraction, content creation, data aggregation and summarization and identifying relevant information from diverse sources. All AI-driven processes are subject to review by BofA Global Research employees. BofA Global Research information is distributed simultaneously to internal and client websites and other portals by BofA Securities and is not publicly-available material. Any unauthorized use or disclosure is prohibited. Receipt and review of this information constitutes your agreement not to redistribute, retransmit, or disclose to others the contents, opinions, conclusion, or information contained herein (including any investment recommendations, estimates or price targets) without first obtaining express permission from an authorized officer of BofA Securities.

Materials prepared by BofA Global Research personnel are based on public information. Facts and views presented in this material have not been reviewed by, and may not reflect information known to, professionals in other business areas of BofA Securities, including investment banking personnel. BofA Securities has established information barriers between BofA Global Research and certain business groups. As a result, BofA Securities does not disclose certain client relationships with, or compensation received from, such issuers. To the extent this material discusses any legal proceeding or issues, it has not been prepared as nor is it intended to express any legal conclusion, opinion or advice. Investors should consult their own legal advisers as to issues of law relating to the subject matter of this material. BofA Global Research personnel's knowledge of legal proceedings in which any BofA Securities entity and/or its directors, officers and employees may be plaintiffs, defendants, co-defendants or co-plaintiffs with or involving issuers mentioned in this material is based on public information. Facts and views presented in this material that relate to any such proceedings have not been reviewed by, discussed with, and may not reflect information known to, professionals in other business areas of BofA Securities in connection with the legal proceedings or matters relevant to such proceedings.

This information has been prepared independently of any issuer of securities mentioned herein and not in connection with any proposed offering of securities or as agent of any issuer of any securities. None of BofAS any of its affiliates or their research analysts has any authority whatsoever to make any representation or warranty on behalf of the issuer(s). BofA Global Research policy prohibits research personnel from disclosing a recommendation, investment rating, or investment thesis for review by an issuer prior to the publication of a research report containing such rating, recommendation or investment thesis.

Any information relating to sustainability in this material is limited as discussed herein and is not intended to provide a comprehensive view on any sustainability claim with respect to any issuer or security.

Any information relating to the tax status of financial instruments discussed herein is not intended to provide tax advice or to be used by anyone to provide tax advice. Investors are urged to seek tax advice based on their particular circumstances from an independent tax professional.

The information herein (other than disclosure information relating to BofA Securities and its affiliates) was obtained from various sources and we do not guarantee its accuracy. This information may contain links to third-party websites. BofA Securities is not responsible for the content of any third-party website or any linked content contained in a third-party website. Content contained on such third-party websites is not part of this information and is not incorporated by reference. The inclusion of a link does not imply any endorsement by or any affiliation with BofA Securities. Access to any third-party website is at your own risk, and you should always review the terms and privacy policies at third-party websites before submitting any personal information to them. BofA Securities is not responsible for such terms and privacy policies and expressly disclaims any liability for them.

All opinions, projections and estimates constitute the judgment of the author as of the date of publication and are subject to change without notice. Prices also are subject to change without notice. BofA Securities is under no obligation to update this information and BofA Securities ability to publish information on the subject issuer(s) in the future is subject to applicable quiet periods. You should therefore assume that BofA Securities will not update any fact, circumstance or opinion contained herein.

Subject to the quiet period applicable under laws of the various jurisdictions in which we distribute research reports and other legal and BofA Securities policy-related restrictions on the publication of research reports, fundamental equity reports are produced on a regular basis as necessary to keep the investment recommendation current.

Certain outstanding reports or investment opinions relating to securities, financial instruments and/or issuers may no longer be current. Always refer to the most recent research report relating to an issuer prior to making an investment decision.

In some cases, an issuer may be classified as Restricted or may be Under Review or Extended Review. In each case, investors should consider any investment opinion relating to such issuer (or its security and/or financial instruments) to be suspended or withdrawn and should not rely on the analyses and investment opinion(s) pertaining to such issuer (or its securities and/or financial instruments) nor should the analyses or opinion(s) be considered a solicitation of any kind. Sales persons and financial advisors affiliated with BofAS or any of its affiliates may not solicit purchases of securities or financial instruments that are Restricted or Under Review and may only solicit securities under Extended Review in accordance with firm policies.

Neither BofA Securities nor any officer or employee of BofA Securities accepts any liability whatsoever for any direct, indirect or consequential damages or losses arising from any use of this information.