on 13-Jul-2026

BERNSTEIN
SOCIETE GENERALE GROUP

Software: Does "Token Shock" hurt AI adoption?

Mark L. Moerdler, Ph.D. +1 917 344 8506 mark.moerdler@bernsteinsg.com
Richard Nguyen +33 1 42 13 54 22 richard.nguyen@bernsteinsg.com
Peter Weed +1 917 344 8390 peter.weed@bernsteinsg.com
Mark Shmulik +1 917 344 8508 mark.shmulik@bernsteinsg.com
Harshita Rawat, CFA +1 917 344 8485 harshita.rawat@bernsteinsg.com
Stacy A. Rasgon, Ph.D. +1 213 559 5917 stacy.rasgon@bernsteinsg.com
Mark C. Newman +1 212 845 7822 mark.newman@bernsteinsg.com
Nikhil Devnani, CFA +1 917 344 8425 nikhil.devnani@bernsteinsg.com
Madison Rezaei +1 917 344 8622 madison.rezaei@bernsteinsg.com
Gautam Chhugani +91 226 842 1416 gautam.chhugani@bernsteinsg.com
Laurent Yoon +1 917 344 8502 laurent.yoon@bernsteinsg.com
Firoz Valliji, CFA +1 917 344 8316 firoz.valliji@bernsteinsg.com
Shelly Tang, CFA +1 917 344 8342 shelly.tang@bernsteinsg.com

By Bernstein's world-wide software team:

Mark L. Moerdler, Richard Nguyen, and Peter Weed

GenAI became the buzz across industries starting with the broad public launch of ChatGPT v3.5 in late 2022 and v4.0 / v4.0o in H1 2023 (depending on which model you consider really "first"[i] to drive broad awareness / adoption) and since then it has created multiple tsunamis shifting global spending toward building AI datacenters; changing how applications are written; concerns that software is going to be eaten by AI; creating billionaires and trillionaires and so much more. All of this investment and disruption has been driven by the belief that GenAI adoption is going to be huge and while there is no disagreement that adoption is already strong (see Ramp AI Index) the continuing investment cycle argues that Gen AI's impact will be even more massive and will be quite impactful across virtually every industry and in everyone's personal lives.

Experts, pundits and industry leaders almost seem to be competing as to how big a number they can create, whether in GPUs sold over the next few years; datacenter spending; job disrupted; or revenues created or lost. Not surprising this has created some roller coasters in the market as one day a company (tech or not) is going to be the big winner and the next it will be a loser.

Even for those with graduate degrees in the space and dozens of years in the industry the speed of innovation and the even faster speed of commentary is daunting.

In today's Weekend Tech Byte we are going to focus on a new and potentially disruptive trend - "Token Shock" driven by the shift of AI to consumption pricing (tokens) and what it could mean for AI adoption.

We don't have all the answers, and frankly we are not sure if we have all the questions but Bernstein's US and European Software teams have pulled our proverbial heads together and give you our perspective. Hopefully you find this helpful.

WHAT IS THE PROBLEM? HIGH COGS LOW PRICE.

The fundamental problem is that GenAI and Agents are not cost free - someone is paying for it. With the launch of ChatGPT followed very quickly by GenAI capabilities in most existing search tools, all with generous free to use models, consumers started to become accustomed to free access. It wasn't just the AI Labs and search engines… lots of "Born in AI" startups were created with GenAI and Agents built in and, taking a page out of the early days of search, many of the at least early consumer centric AI-tools were free — looking to drive "eyeballs" with a plan (or hope) of converting those eyeballs into revenue at some point in the future. In fact, depending on the tool, there was at least initially no advertisements. Sure, there were premium editions, but most people could do everything they wanted to without the paid-for premium editions. Even if they were paying, it was a fixed unlimited usage subscription (until recently) frequently at recognizable broadly consumer-friendly price levels like $19/29/99 per month — the costs didn't seem excessive for the value delivered. There were discussions of ads, but unless GenAI chatbots were added to existing search tools advertising as a monetization strategy is still quite limited.

It wasn't just consumer looking for free. As GenAI capabilities expanded from consumer into commercial the monetization strategies of well funded startups and many enterprise app vendors (e.g. Adobe) used free or freemium models. Even if there was a paid

See the Disclosure Appendix of this report for required disclosures, analyst certifications and other important information. Alternatively, visit our Global Research Disclosure Website.

First Published: 10 Jul 2026 13:10 UTC Completion Date: 10 Jul 2026 13:10 UTC

on 13-Jul-2026

version or a usage fee for commercial-targeted AI applications it really was not pushed aggressively initially by either the Cloud or "Born-in-AI" vendor. Companies and their users could largely partake of the advantages of AI without having to think about the budget impact. Again, the plan, we believe, was to drive adoption (we don't use the term "eyeballs" in enterprise software) 1st and monetization 2nd.

Well, that is not completely accurate. Most companies and IT organizations expected that there would be fees for AI, but the expectation was that the savings generated by GenAI would more than offset the costs and that they would have enough time to try before having to pay anything significant. This does not mean they were happy to throw money at pilots and internal projects but the expectation, we believe, was that costs would be quite manageable until the technology delivered meaningful value.

Of course, this is not the approach every vendor took - Microsoft started from the beginning with an "all you can eat" subscription approach for their numerous Microsoft Copilots. In fact, we would argue that Microsoft launched Office 365 Copilot before it was really ready (it was almost beta code) and was able to capture not only lots of testers but also revenue - but that is a story for another note.

Coding Co-pilots and Agents, in fact, have been the single most visible and likely effective use of GenAI to date. In both large and small companies, their Product teams (building customer facing applications) and IT organizations (building internal productivity tools), have jumped to use AI coding tools given how much it has accelerated the most basic coding functions. The main coding tools from AI Labs (e.g., Claude Code, Codex), coding startups (e.g., Cursor, Poolside), and traditional enterprise vendors (e.g., Microsoft's GitHub Copilot) were a fixed all-you-can-eat subscription fee and the fee was relatively low compared to the savings that Product teams and IT organizations were perceiving. Even if they could not quantify it easily or decrease developer headcount.

At this point, most development teams have given everyone in the organization responsible for writing code access to one or more coding tools. While a relatively small number have started to adjust their staffing, most did not start to replace headcount due to the productivity gains generated by the new AI coding tools. In some cases Product teams and IT saw this as a chance to finally catch up on the massive backlog of legacy modernization / refactoring / bug fixes — maintenance of existing code can often be 25%+ of a team's time, and the hope was to reduce this cost, thus getting more R&D leverage without sacrificing investment on incremental version+1 new projects (often ~2/3 of a development team's time investment) or planning for even further futures (often ~10% or less of spend).

Non-coding tools in search, content creation, report creation, and other areas also started proliferate across organizations as the ease of use and power of the tools were tested and used, even if the ROI was not as visible as seen with coding.

Some, including some notable tech companies (e.g. Microsoft announced ~4,800 layoffs or ~2.1% of workforce, Salesforce cut ~4,000 service desk, Meta laid of ~8,000 employees) have announced layoffs that included developers (and other roles) but the total numbers have been relatively small to date. Not to mention, they seem to largely impact companies perceived as "over their skis" on hiring, and perhaps "AI washing" necessary and overdue headcount reductions.

But with Agentic AI, larger frontier models, reasoning models and the maturation of AI as a business with AI companies having to show revenue and not just eyeballs, all-you-can-eat subscriptions is starting to be replaced with some form of consumption. [more on the drivers of the change in monetization in a moment] Suddenly buyers are starting to see real cost associated with the tools and the costs are variable and in some cases blowing up Product team and IT budgets (or corporate AI budgets pools).

EXHIBIT 1: **Until recently the unlimited use / All-you-can-eat subscription offerings from AI apps including coding tools has made the decision to use AI easy.**



Source: Perplexity.AI, Bernstein analysis

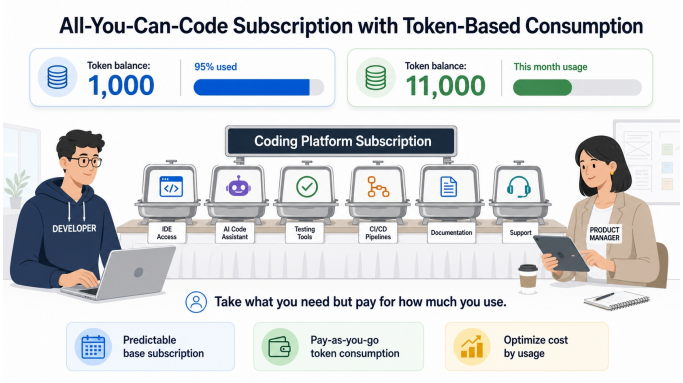
The models' operating costs directly relate to the number of characters input and output (which directly relates to memory and compute use) and maximum memory allowed. To match revenue and cost more closely (and drive customers' incentives to use the resources wisely), AI Labs utilize a consumption meter of "Tokens" (input and output priced separately) and the scale of the "context window" offered. Tokens are roughly four characters of text, and try to provide incentive to limit input/complexity by charging directly for scale of input - they also allow AI companies to monetize what they have created. It also roughly assumes the scale of output is related to the value received (more scale = more value), and if it doesn't, then there is a strong incentive for the buyer to become more thoughtful of their use of this expensive resource. This does not mean that tokens and context window just appeared on the scene — they have been there since the inception of GenAI as model cost has been predominantly measured based on the # of tokens used — whether by the model vendors or the hyperscalers.

Note that this is by no means simply a problem for Coding Agents, Desktop Copilots, or Search tools - **enterprises are, we believe,**

on 13-Jul-2026

starting to grapple broadly with the issue of the cost of AI vs the value. Not trying to jump to far ahead, but this is one of the core issues - cost vs value and how to measure and what is acceptable.

EXHIBIT 2: **But the shift to consumption (tokens) creates new questions for AI tools (e.g. coding tools).**



Source: Perplexity.AI, Bernstein analysis

WHY IS THERE A PROBLEM AND WHY NOW?

The shift in enterprise from free or low cost subscriptions to tokens was inevitable. Consumer internet companies have predominantly monetized via advertising and in theory better ad placements driven by AI and higher usage, driven by AI, would create sufficient revenue to fund many consumer AI solutions. This did not mean that there are no subscription AI offerings, there are (e.g. OpenAI ChatGPT). But in enterprise advertising and usage of data, which is frowned on for consumers but relatively acceptable behavior when used to optimize results and ads, is not going to work in most cases.

At some point it was natural that AI solutions for enterprise would have to charge for their offerings and the charge would have to be more than their cost of inferencing (much more if they were going to generate a profit). With inferencing cost tied to numerous factors which make their cost structure far more variable than for traditional Cloud services, a variable pricing model tied to the level of consumption was a natural outcome.

Add to this the fact that coding, search, support desk software, content creation are all starting to gain traction in enterprise which both drives up the COGs of the AI providers and the opportunity to finally monetize. If customers are using the solution then why not start charging a real price for the solution?

HOW HAS CORPORATE'S APPROACH TO AI CHANGED?

This shift to consumption and with it token-based pricing initially was treated with some surprise by IT, but that time is over and corporate is realizing they need to do something.

TOKEN SHOCK: WHEN AI DEMAND OUTRUNS BUDGET CAPACITY

A growing body of survey data (link and link), combined with our discussions with corporate users, suggests that the next challenge for AI adoption is no longer demand but rather the perceived issue of "cost". The reason we say perceived is that the underlying issue is whether the value created warrants the cost. It is really a Return on Investment issue.

While enterprises continue to increase their use of GenAI, many are finding that operating AI at scale is significantly more expensive than anticipated (in part because expectations were set to low and in part because the ROI is not understood and in part the costs could be too high for the value created). As deployments expand, token consumption is growing faster than token prices are declining, creating what can increasingly be described as a form of "***token shock***" for Product Team and Corporate IT budgets.

Several examples illustrate the issue.

- A major French insurance company recently told us that it had to significantly reduce its deployment of Anthropic Claude only two months after launch because costs had escalated far beyond expectations (workflows and background automation generate far more API calls per task than anticipated).

- Similarly, infrastructure startup Pylon reportedly faced a 3.5x growth in their annual AI bill to ~$1.4MM when its pricing stepped up to enterprise editions (they crossed a key threshold of 150 employees, which meant their seats stopped including tokens and instead all usage is billed at standard API rates), forcing management to become more focused on ROI positive usage. Previously, we understand, they were using AI without any planning / forethought or individual visibility to cost, with an attitude instead about spending tokens aggressively to accelerate growth, and a lot of redundant AI consumption (e.g., rebuilding skills and harness multiple times).

- EY showed us, in one of their projects, that moving from simple search/synthesis direct LLM prompts to Agentic AI where a prompt kicks off a multistep task can increase the cost per interaction about 30-fold (link).

- Recently Uber announced that they had blown through their whole 2026 AI coding budget in 4 months (The Information - Uber CTO and coding). It was reported that the company placed a cap of $1,500 per employee and per agentic coding tool (including Anthropic Claude Code and Cursor) in order to manage budgets (TechCrunch: Uber caps employee AI spending).

- US supermarket giant Walmart is capping staff usage of an AI agent called Code Puppy, used to help with spreadsheets and presentations. Workers who previously had unlimited access to tokens now have a set allocation (link).

on 13-Jul-2026

- We have also heard similar concerns from a large European bank, which emphasized that production AI systems can be dramatically more expensive than pilot projects[ii]. The bank noted that inference costs are expected to represent the majority of an AI model's lifetime economics, potentially accounting for approximately 80% of total costs. In our view, this is a critical point for investors. Much of the market's attention remains focused on training costs, yet the long-term economics of enterprise AI are increasingly determined by the cost of serving and running models in production.

Our discussions with major technology and consulting providers, including Accenture, IBM, Deloitte, Capgemini, and PwC, point to the same conclusion. These firms consistently report that many clients have either postponed, scaled back, or canceled at least one generative AI or agentic AI initiative because of cost concerns and uncertain economics. One of these recently told us they see ROI positive results in less than ¼ of their pilot projects, and even the highest return projects are often thin on result.

Survey evidence supports these observations. Surveys from IBM's Institute for Business Value (link) and SAS (link) indicates that computing and large language model costs have risen substantially as deployments move beyond the pilot stage. In a meaningful number of cases, organizations now describe these costs as prohibitive. Deloitte has similarly estimated that at least half of the generative AI initiatives could exceed their planned budgets by 2028, citing poor architectural decisions, implementation challenges, and limited operational expertise.

On the positive side, rising AI costs are also evidence that AI is moving beyond experimentation and into real-world production environments. The fact that enterprises are facing meaningful AI bills confirms that adoption is occurring at scale and that hopefully demand is generating tangible economic activity rather than remaining largely hype-driven. For instance, Fireworks AI Inc., a US inference platform, serves roughly 30 trillion tokens a day for more than 10,000 customers — including Cursor, Uber, Samsung, Notion and Shopify — and its annualized revenue, an estimated $305m at the end of 2025, had reportedly reached $800m by May (link).

While some could argue that this is because the usage is predominantly in coding tools rather than broadly with agents, but we heard a similar issue last year from early adopters of Salesforce's Agentforce (which an Agentic platform and not a coding tool) that periodically costs got out of hand. The underlying issue was that usage did not match to value being created.

Regarding the issue of model degradation, we talked to two large European insurers about their own domain-specific models that they created to speed up and refine risk assessment and underwriting. For instance, one of the insurers uses weather and climate modeling to price agribusiness and climate-exposed risks, building or commissioning specialized models for hazard and loss estimation. They told us that they face high risks on distribution and concept drifts for their own models. They also said that they have seen the same issues with vendor-trained models.

However, the negative implication is that cost is becoming an increasingly important constraint on deployment. Many organizations are responding by narrowing their focus to high-value use cases, imposing usage limits, delaying rollouts, or canceling projects altogether to avoid budget overruns. In our view, this suggests that the key limitation on future AI adoption may not be lack of interest or demand, but affordability combined with leveraging the technology to drive a visible ROI rather than just using it because it is there and easy to use.

EXHIBIT 3: **AI cost overview**

Area	Observation	Remarks
Model acquisition and licensing	Enterprise token and API spend typically lands in mid- to high-five figures per month.	Often 30% to 50% above forecasts due to usage volatility.
Compute and inference demand volatility	Runtime costs scale non-linearly with concurrency, GPU class, and retraining cadence.	Total runtime can arrive at mid-five-figure monthly ranges.
Workflow integration complexity	Multi-system integrations generally fall in the low-six-figure range, with legacy systems.	Often adding 25% to 35% overruns.
AI skills and enablement requirements	Scarce senior ML talents and slow hiring cycles make personnel the most inelastic AI cost block.	Enterprise AI teams typically carry seven-figure annual cost structures.
AI governance requirements	AI governance requires continuous monitoring, retraining controls, and accountability where programs run low- to mid-six figures annually.	More than 90% of models degrade without monitoring.
Vendor dependency	Contract terms, commit levels, and data rights create disproportionate downside risks.	"Take-or-pay" structures regularly generate five-figure annual waste.

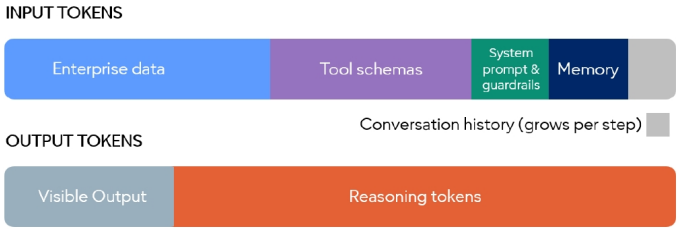
Source: Gartner Group (April 2026), Bernstein analysis

ENTERPRISE AI ADOPTION: ECONOMICS ARE BECOMING THE MAIN CONSTRAINTS

Our conversations with corporate users show that the primary challenge facing enterprise AI adoption is no longer whether the technology works (even though those issues have not gone away), but whether it can be deployed and scaled economically. Unlike traditional software, where costs are largely tied to licenses, subscriptions, or user seats, AI is now shifting to a consumption-based model in which costs depend on the cost per token, token usage, inference activity, API calls, and increasingly autonomous agents. The complexity of modeling the cost can be seen by simply looking at all the components of just the token cost is composed of (Exhibit 4) and that does not include all the other cost factors (e.g. non-GPU costs). As a result, AI spending is far less predictable and much harder to budget.

on 13-Jul-2026

EXHIBIT 4: **Token cost composition per AI agent action**



Source: Gartner Group (March 2026)

A major concern for enterprises is cost uncertainty. Organizations often struggle to forecast how much AI will cost because token consumption depends on user behavior, workflow complexity, and adoption levels. CFOs and CIOs frequently find it difficult to answer basic questions such as: How much will AI cost next year? What happens if usage doubles? How will costs evolve as AI agents become more autonomous? This lack of visibility slows investment decisions and makes organizations more cautious when approving large-scale deployments.

The problem is compounded by the risk of budget overruns and "bill shock." Many companies only discover their true level of token consumption after moving from pilot projects to production environments[iii] and even then ramping usage can add more “bill shock”. While pilot programs often appear affordable, costs can increase sharply once AI is rolled out across thousands of employees or integrated into core business processes. As a result, organizations may find that successful pilots become unexpectedly expensive production systems, creating concerns about cost control and governance.

In, addition this does not even consider the increased traditional compute that is associated with production enterprise AI workloads.

EXHIBIT 5: **Token consumption comparison**

Use case	Relative token consumption	Remarks
Single-turn prompt	Low	This approach does not typically require reasoning or iteration. Input and output remain straightforward to measure.
Multi-turn conversation	Medium to high	Each turn typically resends the full prior conversation as input, increasing token usage as the conversation grows. Long sessions can consume more tokens than the task requires.
Chain-of-thought reasoning	High	The model generates internal reasoning tokens and bills them as output tokens. These tokens do not appear in the visible output.
Retrieval-augmented generation (RAG)	Medium to high	Retrieved context can add a significant number of tokens to the user’s prompt or query. Token usage depends on document corpus size and format. For example, PDF files require more tokens than plain text.
Agentic AI systems	High	Agentic systems call models multiple times for planning, tool use and reflection, which compounds token usage as the number of agents and workflow steps increases.

Source: Gartner Group (April 2026), Bernstein analysis

Measuring ROI is difficult

Another major obstacle is the difficulty of measuring the return on investment. Tokens measure computational activity and resource consumption, but they do not directly measure business value. High token usage does not necessarily translate into higher productivity, lower total operational costs, stronger revenues, better decisions, or improved customer outcomes. Since management teams allocate capital based on expected business benefits rather than technical metrics, the inability to clearly link AI spending to tangible business results often makes large-scale deployments difficult to justify.

In fact, while many software / Cloud / AI companies will point to the value created by leveraging AI to automate processes or increase operational efficiency IT is not good at measuring and validating the return on investment in general and especially of AI. It all looks great on paper but if you can not prove that spending $X on an AI solution directly drives $Y in revenue or savings (where Y is much larger than X) then as soon as the costs spike anyone with budgetary responsibility has got to get concerned.

The challenge becomes even more pronounced with agentic AI. Unlike traditional chatbots that perform a single query-response task, AI agents can plan, reason, validate outputs, use external tools, execute multi-step workflows, and launch additional tasks. As a result, they consume significantly more tokens than conventional AI applications. They also consume CPU-centric compute, storage and bandwidth which have real costs in a Cloud world.

While agentic systems may in theory generate greater business value, they can also increase operating costs by multiples of traditional AI workloads. This creates an important paradox: the most

on 13-Jul-2026

powerful and potentially valuable AI use cases are often the most expensive to run.

Regulated industries face additional hurdles. In sectors such as financial services, AI deployments require extensive governance, auditability, explainability, monitoring, and compliance controls. These requirements add significant operational costs beyond the underlying token and compute expenses. Consequently, the total cost of ownership for AI can be substantially higher than the model usage costs alone, raising the threshold required to achieve an attractive return on investment.

There are also concerns about infrastructure capacity. Demand for advanced AI models may grow faster than the industry's ability to supply computing resources. If this occurs, enterprises could face even higher costs, limited access to leading models, or usage restrictions despite being willing to pay. This introduces another layer of uncertainty into long-term AI adoption plans.

Taken together, these factors suggest that the main friction to enterprise AI adoption is economic rather than technological, even though there are still technological hurdles. Organizations today generally accept that AI can deliver meaningful capabilities. However, uncertainty around costs, the risk of budget overruns, the challenge of proving ROI, and the rising expense of agentic AI are causing many companies to proceed more cautiously. Until enterprises can better predict costs and clearly connect AI spending to measurable business outcomes, adoption is likely to remain more gradual than the technology's capabilities alone would imply.

IS THERE AN ANSWER AND IF SO WHAT IS IT?

We have just laid out what could be a bad scenario in which:

- **Subscription shifting to Token / consumption pricing**: Software / Cloud / AI companies are trying to monetize what they have created and fund the services they are delivering (and hopefully a profit)
- **Enterprise costs are rising**: What was an easy to manage free, or subscription cost suddenly is becoming a variable cost that is swamping expectations and budgets.
- **ROI is tough**: Without a visible ROI (or at least route to an ROI) that not only covers the rising token costs, cut backs are natural.

We believe that there are 3 potential parts to the resolution of this problem:

1. **Enterprise must mitigate costs**
2. **Vendors must lower both their costs and what they charge**
3. **Everyone's expectations need to be managed better**

HOW ENTERPRISES ARE MITIGATING AI COSTS AND TOKEN ECONOMICS CHALLENGES

As AI adoption scales, enterprises are increasingly applying the same discipline to AI spending that they previously developed for cloud infrastructure, based on our industry feedback. The focus is shifting from simply deploying AI to maximizing the business value generated for every dollar spent. In this environment, cost management, governance, and operational efficiency are becoming as important as model capability itself.

The first and most important response is the adoption of **AI FinOps and token governance**. Organizations are implementing tools and processes to monitor token consumption, allocate costs to business units, track spending by use case, establish budgets and usage limits, and continuously measure return on investment. The objective is straightforward: enterprises cannot control or optimize costs that they cannot see. Greater visibility into usage patterns helps prevent unexpected spending and gives management teams more confidence when scaling AI deployments.

Many organizations are also introducing **usage limits, quotas, and spending controls**. As AI expenses become material budget items, enterprises are implementing team budgets, departmental quotas, approval processes, spend caps, and real-time monitoring dashboards. These controls help prevent unexpected bill increases and ensure that resources are directed toward the highest-value applications. In some cases, organizations are effectively rationing AI usage while they develop stronger governance frameworks and clearer ROI metrics.

Some are also moving away from measuring success in terms of tokens consumed and toward measuring **cost per business outcome**. The relevant question is no longer how many tokens were used, but rather how much it cost to process a claim, detect a fraud case, resolve a customer issue, complete a workflow, or support an underwriting decision. A workflow that costs significantly more may still be highly attractive if it generates substantially greater business value. For example, Levi's showed us a supply chain business process was reduced from two days previously to 20 minutes now. Conversely, a low-cost AI application may not be economically justified if it delivers little measurable benefit. Over time, the key performance metric is likely to become business value per dollar spent rather than tokens per dollar spent.

This change in measuring from tokens to cost per business outcome brings organizations closer to understanding the ROI of any AI projects, whether internally developed or from a 3rd party. The two problems with the approach is that 1) consumption may differ even between uses which requires IT to create average cost per business outcome and 2) someone still needs to understand the ROI of each business process AI is automating and modeling the average cost based on current model costs.

on 13-Jul-2026

There are also strategies companies are taking to manage their 3rd party AI and their internally developed AI both to manage the costs and to map the cost to the value be generated to combine and deliver an ROI.

Enterprises managing 3rd party AI

Let us start with 3rd party AI apps. A very wide range of 3rd party AI enabled and AI-centric apps are being tested and in some cases used within organizations. Virtually every organization has wanted to prove they are an AI forward company with some mandating from the top aggressive AI adoption. But there have been roadblocks including:

- **Data security issues:** "Are we sure our data is not leaking out?"
- **Data access and integrity issues:** To best use the AI app we need to give the app access to our corporate data, but the data is not clean, it is not setup, so AI can access; and it spread across multiple databases.
- **Access issues**: Our data is on-premise behind all our firewalls and security but the AI is only in the Cloud?
- **Hallucination questions:** Can we be sure that we will not have a problem if we make a decision or answer a customer question based on an app that hallucinates?

Organizations have spent the last 3 years working through all these issues and have taken steps to deal with these and other issues at least in part. Enough progress has been made that AI usage has been growing steadily. For example outside of coding Microsoft, on their last earnings call stated that they had grown their paid commercial Office 365 Copilot seats from 15M to 20M QoQ. But that is small compared to the coding tools, but it is large compared (especially in dollars) to most enterprise non-coding or search AI apps.

But as discussed earlier individual user usage of these tools are growing at the same time that the cost of usage is growing because the apps are become more powerful due to reasoning models, larger frontier models, Agentic AI. Thus, we are seeing numerous vendors switching to a consumption basis or component (e.g. Office 365 E7 which contains Microsoft Copilot and Agent 365 has both a subscription fee and a consumption component related to agent usage).

Thus, as usage is growing and in many cases the pricing is shifting to some form of consumption, organizations are grappling with the cost side even as they realize they do not have a good understanding of the ROI they are getting from the tools.

In our discussions with CIOs and with the team at Gartner that talks to IT organizations about Office Copilot we have found that while users like the product the ROI has been difficult to measure. In part this is because of the large the reach of Copilot and how it saves time which may or may not convert into tangible savings, and in part it is because IT is not good at measuring ROI. Initially enterprise wide solutions such as Office Copilot are adopted by IT and given to the end-users. But with the user counts potentially substantial (and in an increasing # of cases token fees starting) and IT having difficulty in quantifying the value generated by the tool many IT organizations, we understand are handing off the approval process for distribution of Copilot to the lines of business - making it their responsibility to validate and ROI and approve roll-out. What is interesting is the same discussions we held highlighted that the users, while they can not easily prove the ROI generated by the product do not want to give it up and are strong believers that over time the product will improve, and the ROI will become more visible.

Thus, one approach we are finding is **passing the responsibility of quantifying ROI to the lines of business** and therein to teams who have to effectively use their budgets.

Another interesting approach we have seen is in the coding area where the rate of innovation is likely currently more visible than in any other enterprise area. In November at an AI conference the CEO of one AI startup told me that they were on their 3rd set of coding tools in the past 11 months. This points to one of the strategies we are seeing across enterprises - **switching out AI tools based on price and / or value**. This works when the tools are relatively similar but with Apps as compared to development tools (and models), switching products can be much more complex.

The 3rd approach we are seeing for 3rd party apps is some form of a pilot in which validation that the value generated is worth the cost. HubSpot called out on their last earnings (Link) that they were switching to trial periods for their AI tools. **Clients wanted to validate the capability for a couple of months before paying for the AI add-on**. This sounds eminently logical, but initially everyone was rushing to add AI capabilities rather than implementing a pilot, and now it would seem, even for the SMBs that are HubSpot's core customers a more measured adoption approach is starting to happen.

We would also not be surprised to see organizations switching from adding any AI tools to only those with high value returns as they gain expertise in understanding what makes business sense. They are also likely to shift to more limited rollouts giving those with the strongest use case the tool before considering firm-wide rollouts.

While this helps with subscription pricing, consumption is far more difficult to manage and monitor, and we are going to have to see continued evolution of monitoring tools, token caps and other controls to avoid explosive usage that impacts budgets without delivering the level of ROI that validates the spending.

Managing internally developed AI

Next we turn to strategies that organizations are implementing for their internally developed AI apps.

Another strategy we have seen for internally developed AI is **model routing and inference tiering**. Enterprises are increasingly

on 13-Jul-2026

matching model capability to task complexity instead of sending every request to the most powerful and expensive model available. Routine tasks such as classification, retrieval, basic customer support, and simple workflows can often be handled by smaller and lower-cost models. More advanced models are reserved for complex reasoning, strategic analysis, and high-value decisions. This approach improves economics by ensuring that organizations do not pay for unnecessary levels of intelligence where simpler solutions are sufficient. It also shifts usage from frontier models to small specialized models and requires a complex model picker technology (or person) that can help pick the best model. We have seen SAP doing exactly this with their AI use ages since 2024.

Companies are also focusing on **prompt, context, and workflow optimization**. In many cases, AI costs are driven less by model pricing and more by inefficient implementation. Excessively large context windows, poorly designed prompts, redundant inference steps, and inefficient agent workflows can significantly increase token consumption. As a result, organizations are redesigning workflows, compressing context, limiting unnecessary conversation history, and improving retrieval architectures to reduce the number of tokens required to complete each task.

Retrieval-Augmented Generation (RAG) optimization has become another important cost-management lever. Rather than sending large volumes of documents to a model, enterprises are improving retrieval accuracy so that only the most relevant information is included in each query. Better retrieval reduces token usage while often improving output quality.

Another common response is a greater focus on **high-ROI use cases**. Companies are becoming more selective about where they deploy AI, prioritizing applications that generate clear productivity gains, cost reductions, revenue improvements, or customer benefits. As a result, AI pilots are increasingly expected to demonstrate business value rather than simply prove that the technology works. Broad experimentation is gradually giving way to more focused deployment strategies.

To further improve economics, many enterprises are expanding their use of **smaller and open-source models**. Models such as Llama, Mistral, and Qwen are increasingly used for routine and predictable workloads, while premium frontier models are reserved for the most demanding tasks. This model-mix approach allows organizations to reduce costs without materially affecting performance for many everyday applications. In addition, we are seeing increasing usage of small / specialized models which can be far less expensive to run, at least from a GPU usage perspective, than larger / general purpose models.

Some enterprises are also evaluating **hybrid and on-premise deployment models** for high-volume, repetitive workloads. For stable use cases with predictable demand, the economics of running open-weight models on private infrastructure may be more attractive than continuously paying usage-based fees to external providers. While this approach requires greater upfront investment, it can provide better cost control over time if demand is spread out throughout the day to optimize use of the hardware. For example, a major French bank told us that it runs some AI use cases on its own private infrastructure instead of Microsoft Azure.

The long-term strategic shift

As AI capabilities become more widely available, competitive advantage is likely to shift away from simply having access to the best models. Most enterprises will eventually have access to frontier models, open-source models, specialized models, and AI agents. The more important question will be who can generate the greatest business value from AI spending.

In this environment, AI governance, FinOps discipline, model routing, workflow optimization, and cost-per-outcome measurement are likely to become strategic capabilities rather than administrative functions. Organizations that can efficiently convert token consumption into business outcomes should achieve superior economics, stronger ROI, and better long-term returns from AI investments.

The key insight is that falling token prices alone do not solve the economics challenge. Even if models become cheaper and more efficient, token consumption may grow even faster as agentic AI, automation, and enterprise adoption expand. As a result, total AI spending can continue rising despite lower unit costs.

Ultimately, the most important question, in our view, is not **who has access to the cheapest tokens**, but **who can convert tokens into business value most efficiently**. That may become one of the defining competitive advantages of the enterprise AI era.

EXHIBIT 6: **Insurance claims processing agent - Context budget comparison per action**

Context component	Without cost management practices	With token management practices	What changes
Policy document	8,000 - 15,000	1,000 - 4,000	Pruned to relevant sections
Tool schemas	15,000 - 35,000	2,000 - 3,500	On-demand loading
Claim form and docs	1,500 - 3,500	1,500 - 3,500	No change
System prompt and guardrails	2,000 - 4,000	2,000 - 4,000	No change
Prior claim history	2,000 - 5,000	1,000 - 2,000	Managed memory
Total per action (average)	**~45,000**	**~10,000**	
Reasoning tokens	13,000	6,500	Reasoning tuned
Output tokens	2,000	1,500	Output limits set

Cost per claim (4 actions)

Model	Without token management practices	With token management practices
OpenAI GPT-5.4	$1.35	$0.58 (~57% reduction)
Anthropic Claude Opus 4.6	$2.40	$1.00 (~57% reduction)

Source: Gartner Group (May 2026), Bernstein analysis

WHAT ARE AI COMPANIES DOING?

Then there are the software / Cloud companies, whether called AI companies or not, that are seeing the concerns around the cost of consumption as an issue impacting both their revenue opportunity and margins. Until recently many had been more focused on driving adoption and for those early enough in their individual company maturation process and having sufficient backing to continue to burn cash we would argue they continue with free or freemium models, leverage subscription rather than consumption, or negotiate

on 13-Jul-2026

low cost pilots and early adopter programs. The problem, we would argue is that some point investors are going to want to see these companies cover at least their COGs which means the inferencing costs generated by their customers will need to be covered pushing them toward consumption.

For more mature businesses a combination of 3 methods will likely be used:

- **Optimization:** When competing in a quickly changing valuable market the area of focus for most companies (especially startups) is delivering product and getting customer adoption. If demand is strong that continues to be the focus for an extended period of time as seen in the Cloud where the focus on margins (operational and product optimization) happens when the business slowed. With AI, if the cost of the AI solution is a major pain point then we would argue that many will focus on optimization earlier. In fact, we believe that the Deepseek moment (Link) in January 2025 caused many AI labs to put focus on optimization and for some that has become a meaningful priority - we have seen this focus in the Microsoft AI models announced (Link) at Build which were much less expensive to run and as powerful as peers.
    - With enterprise apps we have been watching since 2024 some enterprise software vendors shifting to small specialized models (e.g. SAP) and targeting point AI functionality / solutions with smaller models and switching out models when appropriate from a quality of result / price perspective. Other vendors (e.g. Microsoft Office 365) are using multiple models including small / specialized models depending on the specific query or action. This shift can drive a meaningful decrease in the cost of inferencing which can be shared with the clients or dropped to the bottom line.
- **Forward Deployed Engineers:** We have written about this trend (Link) focusing on how it is being used to drive customer success, and we continue to believe this is a key driver. But it has a 2nd beneficial result and that is the FDEs can leverage their expertise to 1) help the customer pick projects with visible ROIs; 2) make sure the customer models the ROI and monitors they are delivering value for the AI; and 3) optimize the agents and apps lower the total cost of ownership.
- **Outcome-based pricing:** An interesting variant on the consumption pricing model that we have seen recently announced by a number of software companies (e.g. HubSpot) is to charge per outcome (e.g. $5 per sales lead) rather than by tokens. This allows the software vendor to charge more if the app is used more, and it sets a cap on how much the client will pay.
- **Pricing**: We have seen recent news articles that some AI labs are considering lowering the cost of tokens to drive adoption (Link). We expect that there will be an increasing focus on pricing

WHAT'S THE BOTTOM LINE

One of the key issues, we believe is that enthusiasm has been extremely high as it relates to AI. Everyone is expecting AI to have a massive impact relatively quickly and while we continue to believe that Gen AI will be very impactful - it does not do everything, and it does not fit everywhere. In addition, Gen AI adoption in enterprise is going to be different from consumers as discussed earlier in the note.

Corporate is going to take time to adopt while AI is going to take time to mature. App developers have a lot to learn; the technology has to continue to improve; buyers have to set reasonable goals and rather than boil the proverbial ocean target implementing AI where it is going to be most impactful.

What does this mean:

- **ROI:** Corporate needs to focus not on AI adoption but rather on adoption of AI that has a visible ROI that more than covers the cost of the AI. For corporate this means picking the right problem to solve and for AI companies it could mean outcome based pricing.
- **Optimization:** Everyone needs to make operational optimization a key component in developing AI apps given the cost of tokens and the volume of tokens being consumed.
- **Model picking:** There is as they say a time and a place for everything and with models this is very true. Both AI companies and corporate IT needs to be much more focused on what models they use when. We have been watching the trend of picking models based on how good a fit they are for the use case and the of using the model (e.g. SAP has been doing this for years) and we believe there will be increased focus in enterprise on this issue.
- **Forward Deployed Engineers**: AI and especially agentic AI are difficult / complex solutions that need to designed and implemented correctly from the beginning and whether supplied by the software / AI provider or from a consultant this approach is really valuable until such expertise could be built internally.

There are multiple possible views to the current Token Shock issue:

- **Optimistic**: Agentic use cases are promising technology, and we are seeing real (but early) adoption. When this is combined with the cost of Token falling could drive large scale adoption which in turn will drive down costs and focus more resources on leveraging AI which becomes a virtuous cycle. Note that key is that clients / users feel they are capturing an ROI consistent with their spending.
- **Pessimistic**: The pessimistic view is that Agentic use cases are still struggling to generate sufficient ROI and as Agents become more complex they are using a lot more tokens. In addition, the cost of the biggest / the best models are higher than prior models. The net effect is that a lot more (potentially more expensive) tokens are being used and this is more than

on 13-Jul-2026

offsetting declines in Price/Token. This will decrease the rate of adoption which will push how reaching large scale adoption

- **Still too noisy and unclear**: Then there is view that the future is simply unclear to know. Agentic use cases are promising **BUT** Agents are using lot more tokens (token volume increasing) – offsetting the decline in Price/Token. In other words, we need more time for the market will go through optimism/pessimism hype cycle.

If we had to define our position: we are guardedly optimistic with a concern that adoption is going to take longer than expected and that this will cause some bubble like reactions in the industry.

We hope this was helpful. Any feedback / comments are appreciated.

on 13-Jul-2026

BERNSTEIN TICKER TABLE

			9 Jul 2026		TTM	Adjusted EPS				Adjusted P/E (x)	
Ticker	Rating	Cur	Closing Price	Price Target	Rel. Perf.	Cur	2025A	2026E	2027E	2025A	2026E	2027E
ADBE (Adobe)	O	USD	222.65	379.00	(60.5)%	USD	20.95	24.99	29.15	10.6	8.9	7.6
HUBS (HubSpot)	O	USD	205.95	381.00	(83.3)%	USD	9.70	13.16	15.79	2.9	2.4	2.0
MSFT (Microsoft)	O	USD	384.36	646.00	(43.8)%	USD	13.64	16.85	19.16	28.2	22.8	20.1
MDB (MongoDB)	O	USD	362.87	449.00	53.0%	USD	2,464	3,035	3,626	10.9	8.8	7.4
ORCL (Oracle)	O	USD	144.22	325.00	(59.1)%	USD	7.63	8.57	12.86	18.9	16.8	11.2
QNT (Quantinuum)	O	USD	73.51	94.00	NA	USD		(3.00)	(2.63)		(24.5)	(28.0)
CRM (Salesforce.com)	U	USD	162.50	173.00	(58.9)%	USD	12.52	14.53	15.60	13.0	11.2	10.4
SAP (SAP)	O	USD	158.46	323.00	(68.7)%	EUR	6.10	7.54	9.02	22.7	18.4	15.4
SNOW (Snowflake)	M	USD	267.49	250.00	4.3%	USD	4,684	6,228	7,927	19.8	14.9	11.7
WDAY (Workday)	O	USD	138.34	216.00	(60.1)%	USD	9.23	11.15	13.95	15.0	12.4	9.9
SAP.GR (SAP)	O	EUR	138.44	276.00	(63.1)%	EUR	6.10	7.54	9.02	22.7	18.4	15.3
CAP.FP (Capgemini)	O	EUR	89.00	208.00	(54.5)%	EUR	12.86	13.00	13.95	6.9	6.8	6.4
NEM.GR (Nemetschek)	M	EUR	56.45	68.00	(70.9)%	EUR	1.88	2.26	2.75	30.0	25.0	20.5
SGE.LN (Sage)	M	GBp	830.60	1,050.00	(48.2)%	GBP	0.44	0.50	0.57	18.8	16.7	14.5
IDR.SM (Indra)	O	EUR	49.60	66.00	9.7%	EUR	2.48	2.53	3.10	20.0	19.6	16.0
TMV.GR (TeamViewer)	M	EUR	5.24	5.70	(59.9)%	EUR	0.89	1.05	1.06	5.9	5.0	4.9
SPX			7,543.64									
EDME			1,593.15									

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

MDB, SNOW estimate is Revenues (M); QNT estimate is Reported EPS; HUBS, MDB, SNOW valuation is EV/Sales (x); QNT valuation is Reported P/E (x); MDB, ORCL, CRM, SNOW, WDAY base year is 2026;

Source: Bloomberg, Bernstein estimates and analysis.

on 13-Jul-2026

I. REQUIRED DISCLOSURES

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

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

VALUATION METHODOLOGY

This research publication covers six or more companies. For valuation methodology and other company disclosures:
Please visit: https://bernstein-autonomous.bluematrix.com/sellside/Disclosures.action.
Or, you can also write to the Director of Compliance, Bernstein Institutional Services LLC, 245 Park Avenue, New York, NY 10167.

RISKS

This research publication covers six or more companies. For risks and other company disclosures:
Please visit: https://bernstein-autonomous.bluematrix.com/sellside/Disclosures.action.
Or, you can also write to the Director of Compliance, Bernstein Institutional Services LLC, 245 Park Avenue, New York, NY 10167.

RATINGS DEFINITIONS, BENCHMARKS AND DISTRIBUTION

EQUITY RATINGS DEFINITIONS

Bernstein brand

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

The Bernstein brand has three categories of ratings:

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

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

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

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

on 13-Jul-2026

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

Autonomous brand – common stocks

The Autonomous brand rates common stocks as indicated below. As our benchmarks we use the Bloomberg Europe 600 Financials Price Return Index (E600BK) and Bloomberg Europe Dev Mkt Financials Large and Mid Cap Price Ret Index EUR (EDMFI) index for developed European banks and Payments, the Bloomberg Europe 600 Insurance Price Return Index (E600IN) for European insurers, the S&P 500 and S&P Financials for US banks and Payments coverage, S5LIFE for US Insurance, the S&P Insurance Select Industry (SPSIINS) for US Non-Life Insurers coverage, and the Bloomberg Emerging Markets Financials Large, Mid and Small Cap Price Return Index (EMLSF) for emerging market banks and insurers and Payments. Ratings are stated relative to the sector (not the market).

The Autonomous brand has three categories of common stock ratings:

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

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

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

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

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

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

Autonomous brand – preferred stocks

The Autonomous brand has three categories of preferred stock ratings:

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

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

AUTONOMOUS CREDIT RESEARCH

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

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

CREDIT RATINGS DEFINITIONS

The Autonomous brand has three categories of credit ratings:

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

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

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

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

DISTRIBUTION OF EQUITY RATINGS/INVESTMENT BANKING SERVICES

Equity Rating

on 13-Jul-2026

Bernstein Autonomous LLP or BSG France SA, beneficially, has either a net long or short position of 0.5% or more of the total issued share capital of a class of common equity securities of the following MiFID eligible securities: Quantinuum Inc..

SG and/or its affiliates beneficially own 0.5% or more of [the total issued share capital/any class of common equity securities] with a net [long/short] position of: Workday Inc, Nemetschek, Sage and Indra.

SG and/or its affiliates beneficially own 1% or more of a class of common equity securities of the following companies: Sage and Indra.

Bernstein and/or affiliates have received compensation for investment banking services in the past twelve months from Adobe Inc, Microsoft Corp, Quantinuum Inc. and Capgemini.

Bernstein and/or affiliates have received compensation for non-investment banking securities-related products or services in the previous twelve months from the following clients: Adobe Inc, Microsoft Corp, Quantinuum Inc., SAP SE, Capgemini, Nemetschek and Indra.

Bernstein has received compensation for non-investment banking securities-related products or services in the previous twelve months from the following clients: Microsoft Corp.

An affiliate of Bernstein has received compensation for non-investment banking securities-related products or services in the previous twelve months from the following clients: Microsoft Corp.

Bernstein and/or affiliates expect to receive or intend to seek compensation for investment banking services in the next three months from Quantinuum Inc. and Indra.

Affiliates of Bernstein managed or co-managed in the past twelve months a public offering of securities of Adobe Inc, Quantinuum Inc. and Capgemini.

Bernstein and/or affiliates had an investment banking client relationship during the past twelve months with Adobe Inc, Microsoft Corp, Quantinuum Inc., Capgemini and Indra.

Certain affiliates of Bernstein act as market maker or liquidity provider in the equities securities of: Adobe Inc, HubSpot, Microsoft Corp, Oracle Corp, Salesforce.com Inc, Snowflake Inc, Workday Inc, SAP SE, Capgemini, Nemetschek, Sage, Indra and TeamViewer.

Certain affiliates of Bernstein act as market maker or liquidity provider in the debt securities of: Oracle Corp, Salesforce.com Inc, SAP SE and Capgemini.

Mark L. Moerdler has a financial interest in the equity or debt securities of Adobe Inc, Salesforce.com Inc and Workday Inc (ADBE, CRM and WDAY).
Shelly Tang has a financial interest in the equity or debt securities of Microsoft Corp and Snowflake Inc (MSFT and SNOW).

OTHER MATTERS

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

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

+44 Bernstein Autonomous LLP; London UK

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

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

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



on 13-Jul-2026

+41 Bernstein Autonomous LLP; Geneva, Switzerland

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

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

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

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

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

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

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

CERTIFICATION

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

II. ADDITIONAL GLOBAL CONFLICT DISCLOSURES

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

III. OTHER IMPORTANT INFORMATION AND DISCLOSURES

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

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

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

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

on 13-Jul-2026

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

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

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

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

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

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

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

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

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

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

To our readers in Singapore: This publication is being distributed in Singapore by Sanford C. Bernstein (Singapore) Private Limited, only to accredited investors or institutional investors, as defined in the Securities and Futures Act 2001 of Singapore ("SFA"). Recipients in Singapore should contact Sanford C. Bernstein (Singapore) Private Limited in respect of matters arising from, or in connection with, this publication. Sanford C. Bernstein (Singapore) Private Limited is regulated by the Monetary Authority of Singapore and licensed under the SFA as a capital markets services licence holder for dealing in capital markets products that are securities and collective investment schemes and an exempt financial adviser for advising on, issuing and promulgating analyses and reports on securities. Sanford C. Bernstein (Singapore) Private Limited is registered in Singapore with Company Registration No. 20213710W and located at 8 Marina Boulevard, #12-01, Marina Bay Financial Centre, Singapore 018981, +65-6326-7000.

on 13-Jul-2026

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

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

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

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

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

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

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

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

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

on 13-Jul-2026

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

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

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

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

LEGAL

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

on 13-Jul-2026

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

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

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

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

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

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

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

This publication was prepared and issued by the entity referred to herein for distribution to eligible counterparties or professional clients. The information in this report is intended for general circulation and does not constitute an offer to buy or sell any security, investment, legal or tax advice nor a personal recommendation, as defined by any of the aforementioned regulators. It does not take into account the particular investment objectives, financial situations, or needs of individual investors. The report has not been

on 13-Jul-2026

reviewed by any of the aforementioned regulators and does not represent any official recommendation from the aforementioned regulators. The investments referred to herein may not be suitable for you. Investors must make their own investment decisions in consultation with advice sought from a financial adviser regarding the suitability of the investment product, taking into account the specific investment objectives, financial situation or particular needs of any recipient of the recommendation, before the recipient makes a commitment to purchase the investment product.

The analysis contained herein is based on numerous assumptions. Different assumptions could result in materially different results. The information in this report does not constitute, or form part of, any offer to sell or issue, or any offer to purchase or subscribe for shares, or to induce engagement in any other investment activity. The value of any securities or financial instruments mentioned in this report may fluctuate subject to market conditions. Information about past performance of an investment is not necessarily a guide to, indicative of, or assurance of future performance. Estimates of future performance mentioned by the research analyst in this report are based on assumptions that may not be realized due to unforeseen factors like market volatility/fluctuation. In relation to securities or financial instruments denominated in a foreign currency other than the clients' home currency, movements in exchange rates will have an effect on the value, either favorable or unfavorable. Before acting on any recommendations in this report, recipients should consider the appropriateness of investing in the subject securities or financial instruments mentioned in this report and, if necessary, seek independent professional advice.

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

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

No part of this material may be reproduced, distributed or transmitted or otherwise made available without prior consent of the entities noted herein. Copyright Bernstein Institutional Services LLC Bernstein Autonomous LLP, BSG France S.A., Sanford C. Bernstein (Hong Kong) Limited 盛博香港有限公司, Sanford C. Bernstein (Canada) Limited, Sanford C. Bernstein (India) Private Limited (SEBI registration no. INH000006378), Sanford C. Bernstein (Singapore) Private Limited and Sanford C. Bernstein Japan KK (サンフォード・C・バーンスタイン株式会社). All rights reserved. The trademarks and service marks contained herein are the property of their respective owners. Any unauthorized use or disclosure is strictly prohibited. The entities noted herein may pursue legal action if the unauthorized use results in any defamation and/or reputational risk to the entities noted herein and research published under the Bernstein and Autonomous brands.