

Global Software

Future of Tech: Does the software world as we know it end?



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



Firoz Valliji, CFA
+1 917 344 8316
firoz.valliji@bernsteinsg.com



Shelly Tang, CFA
+1 917 344 8342
shelly.tang@bernsteinsg.com

Given the worry that Generative AI is so disruptive to everything in the world and especially to software, we lay out our 5+ year view on how AI might impact the software stack and what the industry, layer by layer, will look like beyond 2030 with a focus on IaaS/PaaS and SaaS.

We believe that **Gen AI will increase the TAM of IaaS/PaaS driving solid growth at the hyperscalers and parts of PaaS (e.g. database)**. Within IaaS, both GPU/ASICs and CPU usage will increase, especially as Agentic AI becomes more prevalent. We also believe that **PaaS** usage will grow, but certain functionalities will change, and new vendors take the lead. Within PaaS, database usage will grow and on-prem database (e.g. Oracle) will continue moving to the cloud, with more specialized database likely to be created. On the other hand, development tools will be replaced by AI / AI-enabled platforms, but total usage will increase as traditionally non-IT personnel could use these tools to create agents.

Additionally, IaaS/PaaS is still relatively early in the Cloud transition, which is still a structural tailwind for these layers (as compared to parts of SaaS that are cloud-maturing) which we believe will last beyond 2030.

For **application software (SaaS)**, we believe that by 2030 enterprise application solutions are going to be far more than they are today. AI is going to allow more personalization, customization and is going to support automation of currently manual operations and change certain business processes. This means that AI must be deeply embedded into the applications and apps will also become management layers for complex AI agents and processes.

While there are sub-markets that will be disrupted / replaced by Gen-AI, we disagree with the broad stroke bear thesis that AI eats software. We believe that beneath the AI fear, some of the recent slowing in certain segments was in fact driven by cloud saturation more than anything else.

Although not all SaaS vendors are as insulated from AI disruption, we believe that incumbents have a number of advantages and vendors that are "quick" in innovating and keeping up with competition will survive / prosper. Complex deterministic processes creates a moat around an enterprise application as does domain expertise and deep semantic knowledge. We believe that the most protected areas within enterprise software should be **ERP, followed by HCM and then full-suite CRM**, if the vendors innovate and execute.

Finally, within the next 5+ years, we believe that Generative AI is not an end state or the final iteration of AI technology. Unlike the Cloud tech stack that has solidified, Gen AI is very likely going to be disrupted both by internal technology waves (e.g. reasoning, Agentic AI) and potentially completely new AI technologies, some of which will augment Generative AI and some that might replace parts or all of the present day Generative AI tech stack. This, we believe, is a really important point for anyone investing in what they believe is the tip of the spear in AI - the spear may be replaced by the gun.

BERNSTEIN TICKER TABLE

Ticker	Rating	14 Apr 2026			TTM	Adjusted EPS				Adjusted P/E (x)		
		Cur	Closing	Price	Rel.	Cur	2025A	2026E	2027E	2025A	2026E	2027E
ADBE (Adobe)	O	USD	235.72	447.00	(61.7)%	USD	20.95	24.78	28.52	11.2	9.5	8.3
HUBS (HubSpot)	O	USD	206.44	463.00	(89.8)%	USD	9.70	13.11	16.36	3.0	2.5	2.1
MSFT (Microsoft)	O	USD	393.11	641.00	(27.5)%	USD	13.64	16.84	19.39	28.8	23.3	20.3
MDB (MongoDB)	O	USD	233.60	428.00	18.9%	USD	2,464	2,969	3,474	6.7	5.5	4.7
ORCL (Oracle)	O	USD	163.00	319.00	(7.8)%	USD	6.03	7.50	8.70	27.0	21.7	18.7
CRM (Salesforce.com)	U	USD	171.31	194.00	(61.6)%	USD	12.52	13.33	14.82	13.7	12.9	11.6
SAP (SAP)	O	USD	167.81	322.00	(64.2)%	EUR	6.10	7.29	8.72	23.3	19.5	16.3
SNOW (Snowflake)	M	USD	135.47	195.00	(35.2)%	USD	4,684	6,042	7,492	9.7	7.5	6.1
WDAY (Workday)	O	USD	117.86	214.00	(78.2)%	USD	9.23	10.94	12.87	12.8	10.8	9.2
SAP.GR (SAP)	O	EUR	142.88	273.00	(59.7)%	EUR	6.10	7.29	8.72	23.4	19.6	16.4
SPX			6,967.38									
EDME			1,536.25									

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

MDB, SNOW estimate is Revenues (M); HUBS, MDB, SNOW valuation is EV/Sales (x); MDB, CRM, SNOW, WDAY base year is 2026;

Source: Bloomberg, Bernstein estimates and analysis.

INVESTMENT IMPLICATIONS

By 2030, we believe that overall IT spending, which in some cases will not come from traditional IT, will grow strongly with the increasing adoption of Generative AI technology. Over the last 70+ years the speed of innovation has accelerated with each technology cycle and AI not just add to the rate of innovation but the rate of innovation adoption.

Within **SaaS**, two key factors that determine disruption risk are speed of catch-up innovation and the complexity of functionality, especially around the deterministic processes. In addition, incumbents must leverage their domain expertise, semantic knowledge of the interrelationships between data, customer relationships. Those that are not innovating aggressively will be disrupted as are most pure-play on-premise vendors.

Within **IaaS/PaaS**, we believe that AI workloads will increase the TAM for IaaS/PaaS vendors as:

- AI will drive not only GPU/ASICs usage but also CPU usage growth, especially with Agentic AI
- To leverage AI, workloads are going to need to be in the Cloud and by 2030 a much larger portion of critical workloads will have moved driving solid growth for areas such as on-premise databases such as the Oracle database.

We see IaaS/PaaS vendors as beneficiaries of AI revolution. We see **hyperscalers** (**MSFT** and **ORCL** within our coverage), **MDB** and **SNOW** (to some extent) as beneficiaries from incremental database usage driven by AI.

Within SaaS, we see **SAP** and **Oracle** as best positioned and **HubSpot** as relatively unaffected by the AI headwinds investors are concerned about. **Workday** should be fine. AI should be both a small headwind and a mostly offsetting small tailwind (Agentforce) from AI for **Salesforce**.

VALUATION COMPS TABLE

4/14/2026		Net Cash Per Share		Free Cash Flow (LTM)		Revenue Growth	Valuation Multiples				
Bernstein Coverage	Company	Share Price	\$	% of Share Price	Per Share (\$)	% Yield	NTM (%)	EV / NTM Revenue	P / FE	Price / FCF (LTM)	PEG Ratio
	Adobe	\$235.72	0.56	0.2%	24.55	10.4%	8.7%	3.6x	9.8x	9.2x	1.0x
	Microsoft	\$393.11	(4.53)	(1.2%)	10.37	2.6%	15.3%	8.4x	22.4x	37.7x	1.6x
	Oracle	\$163.00	(42.37)	(26.0%)	(8.52)	N/A	28.0%	7.2x	21.4x	N/A	4.0x
	MongoDB	\$233.60	28.61	12.2%	6.16	2.6%	17.7%	5.7x	39.8x	37.5x	2.2x
	Salesforce.com	\$171.31	(8.52)	(5.0%)	15.06	8.8%	11.1%	3.6x	13.0x	11.0x	2.4x
	HubSpot	\$206.44	27.10	13.1%	10.85	5.3%	18.1%	2.6x	16.6x	18.9x	0.6x
	SAP	\$168.51	2.21	1.3%	7.03	4.2%	9.1%	5.1x	20.0x	25.1x	0.9x
	Snowflake	\$135.47	3.82	2.8%	3.32	2.5%	26.3%	7.7x	74.5x	41.8x	1.6x
	Workday	\$117.86	6.09	5.2%	10.36	8.8%	11.5%	2.7x	11.2x	10.9x	0.8x
Other Comp	Autodesk	\$228.59	(0.64)	(0.3%)	11.20	4.9%	13.1%	5.9x	18.4x	20.0x	0.9x
	ServiceNow	\$87.79	3.71	4.2%	4.37	5.0%	20.4%	5.4x	20.9x	19.9x	1.1x
	Amadeus	\$58.76	(6.07)	(10.3%)	4.72	8.0%	6.0%	4.2x	14.7x	12.6x	1.4x
	Datadog	\$110.57	8.79	8.0%	2.52	2.3%	19.4%	8.8x	50.5x	42.8x	7.5x

Source: Bloomberg, Bernstein analysis

Table Of Contents

Portfolio Managers Summary.....	4
What is the right time horizon?.....	5
What happens to the Hyperscalers (IaaS/PaaS) over the next 5+ years?.....	6
Does AI negatively impact IaaS/PaaS and the hyperscalers?.....	9
Does AI increase the TAM for the hyperscalers?.....	10
Does the Enterprise applications industry get disrupted by AI?.....	11
What about the bear concerns?.....	12
Then why are so many Cloud companies' growth slowing?.....	13
Does the technology get disrupted before it completely disrupts the software world?.....	14

DETAILS

Everything today seems to be impacted directly or indirectly by AI - whether wars (autonomous drones and AI driven attack site selection); hiring (can I replace my employees by autonomous agents); or even the software industry. While we would love to pontificate on the global ramifications of AI (and to a small extent I will in this note) I will focus on where we have the most expertise and that is look at how AI will impact the software / Cloud industry over the next 5+ years.

AI in general and Generative AI more specifically is a disruptor and while it will disrupt hiring in some industries the rate of disruption, we would argue will be different between industry and even sub-industries. This, we believe, is due to both how disruptive (and value creating / destroying) the technology will be in a specific area, and the height of the barriers to adoption.

Today, most "industry experts" and many investors want to / expect to see very fast disruption because of:

- The rate of initial adoption of the technology in some very high profile areas (one need only look at the ramp to 1 billion ChatGPT users);
- How fast the technology is moving in areas such as coding; and
- There is so much money being thrown at the technology that fast adoption is required or else something really ugly could occur.

We will look at each of the major layers in the current technology stack individually and what we believe those layers of software and the companies in those layers will look 5+ years from now.

PORTFOLIO MANAGERS SUMMARY

There is a concern that AI is going to eat software and this concern has driven down valuations across the software industry. But if we make some educated assumptions and look out to 2030 and beyond we believe one will get a different perspective:

- The slow-down in growth that has caused consternation across investors is, we believe, driven by a combination of factors including:
 - The Cloud from a technology perspective is maturing and from a Cloud migration perspective is saturating in many sub-markets of enterprise application software / Cloud. The tailwind to growth that has driven software for 20+ years is abating.

- IaaS/PaaS is much earlier in its transition but the law of large numbers and the fact that the easy to move workloads have moved is slowing growth. That said the transition will last, we believe, beyond 2030.
- To leverage AI, workloads are going to need to be in the Cloud and by 2030 a much larger portion of critical workloads will have moved driving solid growth for areas such as on-premise databases such as the Oracle database.
- Incumbant Enterprise applications have a number of advantages with AI that they lacked with the Cloud transition and these will allow many enterprise app vendors to continue to be successful in the AI world. We lay out in the note what we believe are the key differentiators of those that are most likely to be successful and list a number of well protected sub-markets of enterprise apps/SaaS.
- AI IaaS/PaaS is not fundamentally different from traditional IaaS/PaaS and by 2030 the margins for AI and non-AI workloads will not be meaningfully different. Nor will the stickiness of the business.
- Inferencing and agents are most likely to be run, in the enterprise world, at the same vendor where the data resides for performance and security purposes. This is true for both apps and IaaS/PaaS.

This does not mean that there will not be those companies that do not survive, but the opportunity is there for many incumbents to be successful, especially as the roll-out of AI in enterprise is going to take many years and this should allow the incumbents to be successful if they innovate, move quickly, and execute. Just like with the Cloud transition many of those that are around today will not be in by 2030 and beyond because they could not innovate and execute fast enough.


WHAT IS THE RIGHT TIME HORIZON?

Before going into details about the impact of AI on the “future” of software / Cloud we need to define a time horizon. Historically the major technology disruption cycles occur ([Exhibit 2](#)) over a 10-20 year period of time with the SaaS / Cloud disruption impacting software from the early-2000s and starting to abate relatively recently. Note this can be seen in our Cloud penetration analysis notes which shows that many sub-markets of software are reaching 70%+ Cloud penetration and some even higher - such as CRM and collaboration. Note that IaaS/PaaS is still much earlier (more on this later in the note).

EXHIBIT 2: **Major technology disruption cycles occur every 10 - 20 years**

Major Software Disruption Cycles

From centralized computing to AI-native software architecture



Source: Perplexity.AI, Bernstein analysis

This begs the question of how to think about AI and its impact on software - should one also figure on a 10+ year cycle or has that changed? We would argue that:

- Currently, we are seeing Generative AI disruption, but Gen AI is not the final AI technology - there will be additional AI technologies, some of which cause other disruptions and some will impact Gen AI. Example technology include First-principles AI (which could change Gen AI), and Quantum AI (which could augment Gen AI).
- The rate of Gen AI disruption on software has been different from prior technologies' disruption cycles but not as different as many expected. If we look back to SaaS which started, depending on your definition when Salesforce.com was founded in 1999 or when Salesforce started to have small but relatively meaningful revenue (2005 - \$176M) SaaS took time. Similarly, if we look at Amazon AWS which was founded in 2006, but the company did not break out AWS revenue until Q1 2015 when it showed that 2013 AWS revenue had reached \$3.1B. In other words, Cloud took > 5 years to generate meaningful revenue even though it started to have an impact even earlier.
- Gen AI generated a tsunami of emotions and expected business impacts in 2023 with the launch of ChatGPT even though AI has been around since the 1930s and the underlying technology of Gen AI since 2019, if not earlier. While consumers, we would argue, are already feeling the effect of Gen AI, enterprise AI is still relatively muted in impact and limited to a sub-set of future use cases.


Give the historical context and if we use 2023 as the "starting" date of its impact on software, we should consider the majority of the impact well understood by 2030 or shortly thereafter, with the full impact likely not felt until more like 2035. But given our expectation that Gen AI is 1st wave of what is likely a multi-wave tsunami of AI technologies we could likely see the next disruption wave before this one is fully registered (e.g. 2035). Therefore, we are going to use 2030 which is ~ 4 years out from now and ~7 years after the launch of ChatGPT as a good point in time to lay out the major impacts of AI.

WHAT HAPPENS TO THE HYPERSCALERS (IAAS/PAAS) OVER THE NEXT 5+ YEARS?

There are 3 main layers to the Cloud ([Exhibit 3](#)): Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS). IaaS and PaaS are often grouped together for numerous reasons, because they are IT focused rather than end-user focused (SaaS) or because there is overlap between the two layers (e.g. some include the operating system in IaaS) or because the hyperscalers bundle both together (with some infrastructure SaaS-style apps). Together IaaS/PaaS is often called the infrastructure layer since applications, whether SaaS or enterprise built Cloud apps, are built often on the combination of IaaS and PaaS. Note calling it the Infrastructure layer in some ways devalues the software component (e.g., database, middleware, security) that is PaaS and in some ways creates, we believe the misunderstanding we are about to discuss relating to the economics and stickiness of the AI Infrastructure layer.



EXHIBIT 3: The main 3 layers of Cloud - IaaS, PaaS and SaaS



Source: Perplexity.AI, Bernstein analysis and estimates

The reasons that IaaS/PaaS have been so successful are numerous but fundamentally, there is an enormous level of efficiency that comes with centralizing the procurement, delivery and management of the plumbing on which applications run. Without this layer SaaS vendors would have had to build out their own datacenters and manage and deploy more of the necessary platform software (databases, development environments, management, security and other platform software). This would have taken each individual SaaS vendor time, effort and money and impacted the speed of adoption especially as they go global.

It is true that in the early days of SaaS the leading SaaS vendors (e.g. Salesforce, Workday, ServiceNow, etc.) rented datacenter capacity from co-location vendors, installed their own server and related hardware, and installed and maintained all the platform software. But we would argue that this work limited the number and diversity of SaaS vendors that came to market and created an environment where smaller vendors were more easily gobbled up by large on-premise and Cloud vendors. It also limited where SaaS was available as there are legal, regulatory and business reasons where the data must reside, and the applications must run in numerous jurisdictions globally.

Enterprises could not have moved the applications they built internally to the Cloud without the hyperscalers because there would not have been Cloud as we know it today. The multi-tenancy and shared resources of IaaS/PaaS made the economics and ease of migration and upgrades far easier and more compelling. Without IaaS/PaaS vendors enterprises would have continued to host in their datacenters or leverage co-location vendors that supplied the datacenters but not necessarily the hardware. How that would have impacted AI adoption in enterprise is a very interesting question.

While the migration from on-premise to IaaS/PaaS and the building of new Cloud native apps on IaaS/PaaS is still relatively early days (possibly <1/2 of the way to completion), the transition itself is well understood and one in which we have seen the formation of a handful of global hyperscalers including Amazon, Microsoft, Google, Oracle, and Alibaba; numerous small/specialized (e.g. government, regulated industries); and regional players driven by local regulations and requirements. Without AI or some other major disruptive trend (e.g. deglobalization) this structure would likely not meaningfully change for the foreseeable future.

To understand the next 5 years we need to consider really 3 trends:

- **Cloud saturation:** As discussed earlier the Cloud transition has been going on since the early teens driven initially by a host of smaller players and then dominated by Amazon AWS. The transition became well entrenched with the launch of 1st

Microsoft Azure and then Google GCP which supplied more locations, more hardware options, and more platform software capabilities. Sovereign Cloud was much more recently, we would argue, enabled by Oracle OCI and Alloy. Based on current IaaS/PaaS saturation we would expect that the IaaS/PaaS Cloud transition should still be continuing, at a slower but still very healthy rate 5 years from now. In fact, Cloud saturation could in theory take 5–10 years depending on the rate of Cloud migration and new workload creation (more on this in our discussion below of the impact of AI on the hyperscalers).

- **Deglobalization:** Our current expectation is that deglobalization is going to increase and there will be a growing list of government enabled / support Cloud platform vendors which will in most cases be far more IaaS vendors than major PaaS vendors since many are going to lack the capabilities to develop, support and deliver their own PaaS but will rather 1) use relatively generic open source PaaS which will have limited applicability for enterprises and existing workloads but will offer the basic functionality that digitally natives will be able to leverage in conjunction with their own internally supported PaaS capabilities; 2) rent the PaaS layer software (databases, development tools, cybersecurity offerings) from point solutions in the PaaS software market¹ or 3) rent PaaS from a 3rd party software centric Cloud provider such as Oracle and its Oracle Alloy and Oracle sovereign Cloud offerings.
- We continue to believe that Oracle is currently the best positioned and has been successful, via their Alloy offering to white box OCI and all its PaaS capabilities as well as Oracle's OCI hardware.
- **AI:** But the big unknown is AI because of how relatively early stage AI adoption is; how fast the technology is changing; and the questions that are arising about the economics of AI IaaS/PaaS. We believe there are two major potential outcomes which will be discussed in greater depth in the remainder of this section of the note.
 - AI is bad for the hyperscalers:
 - AI decreases the margins for the hyperscalers
 - AI lowers the stickiness of the hyperscalers and increases churn - potentially decreasing both revenue growth and the value of these businesses
 - AI completely disrupts the hyperscaler market as we know it today
 - AI is good for the hyperscalers:
 - AI increases the TAM for the hyperscalers
 - AI leverages data so those vendors with lots of critical client data become even more valuable

Today 5 hyperscalers (Amazon, Microsoft, Google, Oracle, and Alibaba) are individually and in combination the largest Cloud infrastructure vendors in revenue, global footprint, customer base and global impact terms, and they have been to date only gaining share at least as a group (the Top 3 AWS, Azure and GCP holding approximately two-thirds of the IaaS market according to IDC - see [Cloud in the Quarter: How did the hyperscalers do in 4Q25?](#)). Currently, the hyperscaler's scale and reach is making it more and more difficult for competitors to breach their moat and were it not for deglobalization and AI, we would argue, they will only become a larger and larger portion of the overall market with a long tail of much smaller local/regional and specialized / industry specific (e.g. regulated industries) competitors.

Deglobalization is having some impact, as discussed above and that will likely continue with regions such as Europe and countries such as China and Russia having the largest local/regional players. A number of Middle Eastern countries are also showing strong signs that they wish to develop their own local IaaS/PaaS industry and with the recent turmoil in the region we would argue that this is likely to become a larger focus moving forward. Again, as discussed above Oracle currently is having success partnering with local providers to deliver a far more feature rich PaaS layer and more flexible IaaS capabilities.

This does not mean that the big 3 are in trouble on a global basis due to deglobalization, but their approach to date has been predominantly focused on partnerships, with the partner delivering the datacenter and selling the services and defining

¹ It is interesting to note that to date the some of the bigger independent PaaS layer software providers such as Snowflake and MongoDB are focused on the 3 largest hyperscalers and are not building meaningful relationships with regional or smaller IaaS/PaaS providers. On the other hand VMware on the other hand has offerings on other IaaS/PaaS providers but their offerings are different and more on-premise centric.

that as sovereign Clouds (e.g., Blue which is a partnership in France between Microsoft, Orange and Capgemini). Thus, it is most effective for large multinational corporations, US (or Chinese centric for Alibaba) centric companies, and digitally native companies (which are mostly US centric today) and some sovereigns where ownership of the tech is less important. It will be interesting to see how this trend continues and whether the big 3 end up having to re-architect their infrastructure to support small footprint datacenters and quasi white box offerings. But we have digressed from the potential impact of AI.

DOES AI NEGATIVELY IMPACT IAAS/PAAS AND THE HYPERSCALERS?

Let us 1st start with the disruptor hypothesis under which AI fundamentally changes, and changes in a negative way, the hyperscaler market. The theory is driven by 3 bear cases:

- AI changes fundamentally how software is written and lowers the moat around application software (more on this later in the note)
- With AI the infrastructure stack is less valuable and the infrastructure vendors are commoditized
- The value shifts to the model companies from the software companies (this has become less prevalent recently)

Let us deal with each one individually and try to show why each in their own way is incorrect:

AI changes how software is written is in fact correct from a coding perspective and this can already be seen in the impact of products such as Anthropic's Claude², OpenAI's Codex, and Cursor. Over time AI coding, agentic development platforms, and related software will only get better. But fundamentally these are all applications and these applications run on Cloud Infrastructure (both CPU and GPU). In other words, while the software products can and will change over time the fundamental requirement that software, whether traditional CPU or AI, requires infrastructure to operate and the more powerful and complex the software the more powerful and complex the infrastructure.

The infrastructure stack becomes less valuable. The 2nd bear case is based on a belief that AI runs on commodity hardware and the difference between one AI IaaS/PaaS vendor and another is relatively nominal. Under this thesis, especially if models are run across all vendors, then workloads can shift around to whomever is delivering the lowest cost or best location (closest to the client) and that once capacity and demand normalize pricing will plummet as low-cost providers enter and drive the market. This thesis, we believe, is wrong for two main reasons:

1. AI is no different from other software and requires platform software to operate, and this software will not be generally available everywhere. For some reason when many people talk about AI they forget all the platform software that runs under the AI applications. In addition, AI software runs in conjunction with and leverages non-AI software which includes traditional applications and infrastructure software (e.g. databases).
2. AI requires data, and it is far better from a complexity, performance, security and governance perspective to run AI where the data resides; and 3) AI is not simply GPUs - it is a combination of GPUs, CPUs, networking, storage and more. There is an enormous amount of knowledge and expertise to build and run datacenters especially mixed use datacenters that will be required for AI as it matures and especially Agentic AI. In fact, looking 5 years out we would argue that the complexity of these datacenters will only increase, and we would not be surprised to see quantum computers starting to be incorporated into the datacenter architecture.


The value shifts to the model companies. This bear case argument is based on the belief that the model companies of today will become full tech stack companies offering the full range of development technologies and the hardware that it runs on top of. Herein, one is arguing that the model companies will one day own their own training and production datacenters. We believe that the most successful companies focus on the few areas they are best at and thus only cross-industry companies such as the hyperscalers plus Meta will try to own IaaS/PaaS and offer up models, but that is in fact the wrong discussion. For this bear case, similar to the prior 2, one needs to understand that AI is not a standalone complete technology stack. It is built on and into non-AI PaaS functionality (e.g. AI does nothing to change fundamentally the databases market) and will integrate into traditional CPU centric applications. Therefore, the best place to run AI is going to be where applications, data, etc. are run and stored.

² We recently hosted a webinar in which my son, who runs an AI software company, used Claude to create a very basic sales management tool and showed what it can and what it currently can not do and where the limitations are.

DOES AI INCREASE THE TAM FOR THE HYPERSCALERS?

We believe that AI and non-AI are tightly intermingled and that will only increase. In [Exhibit 4](#) we show a simplified graphic representation of who AI will be tightly integrated into each layer of the Cloud. Companies, and especially enterprises, are going to want to run their AI where their applications and data reside and that this will lead to an increase in TAM for the hyperscalers.

EXHIBIT 4: **Our view of the tight interrelationship between AI and non-AI in the Cloud/AI tech stack**



Source: Perplexity.AI, Bernstein analysis

When we consider a new Cloud / AI TAM there are a couple of interesting trade-offs:

- **Infrastructure usage is going to increase** as GPUs or by 2030 mostly ASICs and specialized inferencing chips are used to perform inferencing. CPU utilization will overall increase, especially as Agentic AI becomes prevalent by 2030.
 - Note that some application functionality that is run on CPUs will be better automated with a mix of GPUs and CPUs. While this will impact specific workloads the overall effect, we believe, will be an increase in CPU and GPU utilization because agents leverage both chip sets.
 - Database usage will increase especially as enterprises build / buy AI apps and Agents which will leverage corporate data and drive non-GPU infrastructure usage up. This will also increase storage requirements due to creation of vector and other databases and the possible storage of more data. But directly AI does not create more data.
- **PaaS usage will grow** but what PaaS functionality is used for will change.
 - Database usage will grow and likely more specialized databases will be created.
 - Development tools will be replaced by AI / AI-enabled development tools, but the total usage will increase. Right now IT organizations are changing AI tools very frequently (we have heard 2–3 times per year at some organizations) which is bad for the software development market, but that will likely stabilize by 2030 and dev tools will become even more prevalent as non-IT positions will use what was historically considered as dev tools. Note there is a possibility that some of the dev tool market will commoditize due to AI but that still has to be seen and there are lots of factors to consider.

- System management and cyber solutions will have to change considerably. At this time we are not calling out winners and losers, but the total dollars spent on system management and cyber solutions will likely increase given the larger volume on apps that will be run and the increased requirements of AI enabled IT organization.

IT will become more prevalent by 2030 and while individual products and vendors will likely change overall spending on what today is called IT will increase and with it the demand for more infrastructure to run the IT solutions and more

DOES THE ENTERPRISE APPLICATIONS INDUSTRY GET DISRUPTED BY AI?

There is a movie entitled the "Quick or the Dead" starring amongst others Sharon Stone, Gene Hackman, Russell Crowe and Leonardo DiCaprio which was a 1995 western about a quick draw competition and those that were not quick were dead. While there are definitely sub-markets of software (e.g. coding tools) that are going to be massively disrupted by Gen AI, we believe that the movie title is a much more appropriate definition of what happens to apps - Vendors that are "Quick" will survive / prosper and those that are not will be dead (or at least on the long road to oblivion). The question which will be discussed in this section of the note will be "*What does it take for an enterprise solution / company to survive/prosper?*"

The answer may be easier than one would think. If we look out to 2030 we believe that there will be a number of key characteristics of those software / Cloud products / vendors that will be most successful:

- **Speed of innovation** in some ways is a new / novel concept for software / Cloud companies and especially enterprise apps vendors. Traditional on-premise vendors delivered major upgrades every 2–3 years and IT could take another 1–2 years or longer to implement. The Cloud and agile development techniques accelerated the pace of innovation to weeks (and in some cases faster for incremental updates). AI is further accelerating the rate of innovation as code writing and testing accelerates. Successful software companies will need to maintain a fast pace of innovation with AI making it not just possible but also economically feasible.
- Please note that the pace of innovation has acceleration since the start of software over 70 years ago. AI further accelerates the pace of innovation, but it also should increase the pace of implementation of innovation ([Artificial Intelligence... Sooner or Later: AI is already here - The structural reset of IT Services & Software](#))
- **Keeping up with the competition** becomes even more imperative in the fast pace AI enabled, and the AI embedded software world. Fortunately incumbents don't necessarily have to lead, but they must be fast followers. In addition, incumbents have 3 major advantages in the new AI world: domain expertise, access to large customer base, and resources (i.e., R&D, own IP and long-standing partnerships with other leading tech vendors and major services providers³). If anyone (startup, competitor or client) creates a great new AI enabled feature or function the incumbent has better access to the knowledge, expertise and access to the data required to ground AI capabilities than the smaller even more nimble competitor. It is going to be much more important going forward for application vendors to monitor what their competitors, including AI centric startups, are offering and where appropriate including similar functionality.
- It is important to note that AI is a double-edged sword and that the incumbents have in almost every case access to the same models that their competitors (whether traditional competitors or AI natives) are using. But they also have far more domain expertise, semantic knowledge that details how data for that specific applications market inter-relates, access to a broad customer base to test and potentially monetize any new functionality. They also have the resources, when appropriate to acquire small competitors rather than building similar capabilities. But this requires streamlined operations and focus, which not every company will do.
- **Embedding AI:** By 2030 enterprise application solutions are going to be far more than they are today. AI is going to allow more personalization and customization. It is going to support automation of currently manual operations, and it is going to change business processes. This means that AI must be deeply embedded into the applications and apps will also become management layers for complex AI agents and processes.

³ There is a generally held belief that AI removes the need for service providers - end users and IT can do it themselves. This is in fact incorrect for the foreseeable future - see our note on Forward Deployed Engineering (FDE): Where AI software meets the real world.

- **Complex capabilities and tight integration will be far more important:** Complexity including complex deterministic processes creates a moat around an enterprise application. While the app must support complex processes they must not look to the user as complex, as AI must be used to personalize, streamline and automate.

But not all enterprise sub-markets are equally protected. We will dive into what areas are more protected and those that are most exposed in follow-up notes. That said, we are confident that the most protected area of enterprise software should be ERP (including supply chain software), if the vendors innovate and execute, followed by HCM and then full suite CRM. Within the infrastructure software space databases are not going to be replaced by AI (at least as we know AI today) and are in most cases going to see upside. Areas such as application development tools are amongst those will be most obviously impacted.

WHAT ABOUT THE BEAR CONCERNS?

There are many in the technology and investment world that believe that software is dead - killed by AI, and we believe it is important to address the concerns and not just layout our thesis for the application software in 2030 and beyond. There are 3 parts to the bear case on enterprise apps:

- Startups will abound because AI is going to make it so easy to build competitive apps which will increase churn and /or drive pricing down.
- Rather than buy / subscribe to additional functionality from one's current vendor, IT will build their own add-ons.
- Organizations of all sizes will use the latest AI tools to build their own applications rather than buy / subscribe to apps

All of these revolve around the correct belief that AI makes application development far easier and that the technology will only improve. But there are two glaring issues, we believe, when trying to connect the dots between development becomes easier / faster and application software is dead.

Building an application is not the same as building a business.

Even if we were to assume that by 2030 one can snap one's fingers and AI would build a complete enterprise class application with all the functionality and capabilities a company would require, that is simply not enough. If one looks at any SaaS startup (or even fast-growing public SaaS vendor) the largest portion of spending is S&M. Having built a successful software company and having worked for one of the largest software companies, selling is not instantaneous and enterprise sales is even more complex ([Exhibit 5](#)). Of course, AI will streamline the process, but the reality is that if today you gave us a fully developed, feature rich enterprise application we would have to:

- Find potential customers
- Demonstrate the product,
- Show why it is not just much cheaper but much better (you don't replace important business applications to save a couple of dollars)
- Sign up the client
- Implement the solution (which will require consultants even with AI assistance and likely require meaningful changes to processes and procedures)
- Make the client happy enough to be a reference
- Take that reference and find new potential customers - as they say Rinse and Repeat.


That process takes years to ramp, and it simply will not, we believe, be meaningfully different in 2030. Please remember that it takes time - one need only look at how long it took Salesforce or Workday or ServiceNow to ramp their businesses, and they were offering major technology innovation and far better operational efficiency than the on-premise incumbents.

Meanwhile, as discussed above, the incumbents should not be standing still and waiting (as they did during the early days of Cloud). They will add competitive functionality - whether the competitor is a full-blown enterprise app or an AI Agent that solves

one problem better.

EXHIBIT 5: The traditional sales process will not be meaningfully different by 2030 due to AI and thus AI enterprise adoption takes a lot of time and effort

Enterprise-grade process diagram for high-end SaaS company. Full lifecycle of winning and expanding new software Clients



Source: Perplexity.AI, Bernstein analysis and estimates

You don't just build an enterprise app you have to care and feed it forever

Applications are not static. They require maintenance, support and upgrades. Every quality software application is evolving to meet changing business, legal, regulatory, accounting and IT requirements. Software vendors learn what they need to do from their large list of clients that are always asking for additional functionality and by working within their industry (e.g. Accounting systems from accountants, auditors, and regulators). Any IT organization that decides to build their own app or even their own complex agents is taking all this responsibility onto themselves. The past 70+ years of software has shown that this is sub-optimal unless the application requirements can not be purchased elsewhere. There is also a greater security risk if an organization lets their various departments/entities build their own apps with AI, security risks that will require more resources from the central IT department to monitor and repair any breaches. Also create shadow budget - this is what we have seen with the cloud transition.

This is why most organizations would much rather buy (or today subscribe) to an off-the-shelf app or Cloud service than build it themselves. While AI speeds up the coding process and decreases the time and cost of building apps it does not replace all the other requirements necessary for the care and feeding of critical business process applications (and they can't benchmark their solutions against the best-in-class in their industry). We expect to see a couple of companies try this route, and we will wait for what is likely to be epic failures (even though few of those will be publicly announced with the fanfare that the move to internally built apps will be announced).

As a proof point we would recommend the reader consider the large AI companies which are predominantly using off-the-shelf enterprise apps rather than building them themselves, even though they have far more expertise than virtually any IT organization.

THEN WHY ARE SO MANY CLOUD COMPANIES' GROWTH SLOWING?

One of the drivers of the shift in investor sentiment from AI is going to help software to AI is going to kill software is that many growth software companies slowed meaningfully. While there are likely company specific reasons, we believe there are two

major drivers:

- **The Cloud transition is slowing:** Application software has been moving to SaaS for > 20 years and applications have been moving aggressively to IaaS/PaaS for > 10 years. As with many adoption cycles they follow an S curve and as most of the workloads have moved, growth slows. This occurs both because of saturation and because the last workloads to move are often the most difficult to move.
- **The pandemic pulled forward the Cloud transition:** During the worst of the pandemic, we believe, many Cloud transition projects were delayed or not completed but as soon as the world started to open up IT moved more aggressively to move workloads to the Cloud. Helping drive at least procurement of Cloud capabilities was the fear that another lock-down could occur and being in the Cloud was far better than in an on-premise datacenter that was not easily accessible during a lock-down.

In other words, the tailwind that has sustained strong growth in software / Cloud is ending for many application segments. IaaS/PaaS is much earlier in the transition but between larger / more critical workloads now moving (which takes more time) and the law of large numbers growth has slowed.

This does not mean that AI has not slowed some parts of the Cloud transition as IT organizations have tried to grapple with new and disruptive technology that is moving quickly. This has led to some projects slowing or being delayed adding to the natural deceleration that is occurring. But workloads are going to have to move to the Cloud in order for the data to be available for AI and for agentic and other AI technologies to be implemented to streamline organizational operations and automate currently manual processes.

DOES THE TECHNOLOGY GET DISRUPTED BEFORE IT COMPLETELY DISRUPTS THE SOFTWARE WORLD?

One other point that we believe is really important, which we discussed when we were discussing time horizons, is whether Gen AI is itself disrupted? We believe that Generative AI is not an end state or the final iteration of AI technology. We are not a point in AI similar to where we are with Cloud (SaaS/PaaS/IaaS). While Cloud technology will continue to change and improve, the overall technology scheme has solidified and changes will be incremental rather than disruptive at least over the next 5+ years.

AI is different. Generative AI is not the end state technology no matter what many would like to believe, and it is very likely going to be disrupted both by internal technology waves (e.g. It's now all about Agentic AI) and changes in the technology (reasoning models changing how models are used and deliver "better" results). But that is not all we would argue that we will see. We believe that within the next 5+ years there will be completely new AI technologies, some of which will augment Generative AI and solve new problems and some will replace a part or all of the areas that Generative AI is operating today. This, we believe, is a really important point for anyone investing in what they believe is the tip of the spear in AI - the spear may be replaced by the gun. A better technology may (and we believe likely will) be developed that disrupts Gen AI.

This is I am sure a really scary thought because it could mean that so much of the investment cycle going on today may not recoup their investment plus profits depending on what the new technology is and how /where it will operate. Does it run on GPUs? CPUs? Quantum computers? or something else (e.g. Photonic compute).