

Accelerating AI:

Navigating the Future of Enterprise Infrastructure

Featuring exclusive insights from:



Chris Bedi
Chief Customer Officer and
AI Enterprise Advisor
ServiceNow



Shez Partovi
Chief Innovation Officer &
Business Head, Enterprise
Informatics
Philips



Sam Hamilton
SVP, Data & AI
Visa



Raul Martynek
CEO
DataBank

In partnership with



Introduction

For many larger enterprises, Artificial Intelligence (AI) is no longer a technology ‘in the pipeline’ or something worth considering for the future. It has already moved well beyond even the training phase and into the real world of implementation – or inference.

Furthermore, as our survey of IT and tech professionals shows, nearly 25% of enterprises are already seeing consistent annual return on their AI investments with a further 35% expected to do so within the next 12 months. Just 10% of organizations surveyed expect ROI from AI to take over 3 years.

Perhaps, inevitably, most of the initial focus around AI deployment so far seems to be on the ‘quick wins,’ such as improving the customer or employee experience at the same time as saving money.

However, the next stage – and the one which many believe will yield the greatest return – will be what Shez Partovi, Chief Innovation Officer, Philips, calls the ‘third lift’ – when AI allows you to do something that wasn’t previously possible, due to technology or cost restraints.

Key to success is preparing data and people for AI. Whereas our survey shows much work has been done when it comes to the former (poor data quality is cited as a blocker by just 18.5% of respondents), there remains much more of an issue around the latter with ‘integration, scaling and/or deployment issues’ cited as a blocker by over a quarter (26.2%) of respondents.

Our survey shows that currently public and private cloud solutions are favored by nearly two-thirds of respondents (64%) as the starting point for AI adoption. However, it appears that many enterprise are adopting a hybrid approach, combining elements of cloud services and on-premises systems

with specialist hardware such as GPU-based (Graphics Processing Units) infrastructure to support the growing demand for inference-based AI applications as well as incorporating simple SaaS solutions.

Looking at the next five years, what’s clear is that most of the large enterprises expect the distribution of their organizations’ AI workloads to change considerably with only 4% of respondents reporting ‘no significant changes planned.’

With security and compliance as well as performance (latency, compute power) considered the most important factors determining AI workloads, our survey shows that over half of respondents (53%) expect there to be an expansion in both AI-dedicated on-premises as well as colocated data centers, such as those provided by DataBank.

Furthermore, AI adoption will drive the ongoing decentralization and geographic distribution of AI infrastructure. According to our survey, over three-quarters of respondents (76%) expect their infrastructure to expand geographically to be closer to data sources or end users while only one in ten (11%) anticipate consolidation.


Concludes Raul Martynek, CEO of DataBank, which operates 65+ US data centers close to major population hubs: “Data center colocation is an ideal infrastructure option for AI workloads – one that provides geographic diversity and proximity to users and data sets, scalable power and cooling architectures, as well as the latest security and compliance frameworks.”

Achieving ROI: From quick wins to new applications

Our survey of mainly larger enterprises shows that early investment in AI is beginning to pay off with one in four (25%) respondents already achieving consistent ROI and over one third (35%) expecting to make a return on their AI investment within the year.

Fig. 1

Which of the following enterprise data sources would you say are very well prepared for training AI models and generating insights?



For example, Visa – which began using simple predictive AI models for fraud detection and prevention over 30 years ago – now employs deep learning models to make fraud detection more effective and has, more recently, deployed transformer models to prioritize and identify complex fraud patterns. “Last year we prevented \$40 billion of fraud with the enhanced AI models we have,” explains Sam Hamilton, Visa’s SVP Data and AI.

Nor is fraud protection the only AI application where Visa is seeing ROI. Visa has deployed a secure version of ChatGPT for its employees, Microsoft’s Copilot for Office and GitHub Copilot for developers, in order to boost productivity and efficiency. It also has conducted a 3-week hackathon with 6,000 employee participants to generate 2,000 ideas for further AI-driven improvements.



Chris Bedi, Chief Customer Officer at ServiceNow, being interviewed by Raul Martynek, CEO at DataBank, at the AI CEO Summit 2025.


Meanwhile ServiceNow claims to be achieving over \$350 million in enterprise value from AI, largely due to enablement transformation across the enterprise. “We used to have a team of people in our finance department whose job it was to answer questions and resolve issues. It used to take about four days to get an answer,” explains ServiceNow’s Chief Customer Officer, Chris Bedi. “But thanks to AI it now takes just 15 seconds and we’re repurposing those people to do something more interesting for the company than answering repeated questions.”

According to Mr Bedi, Chief Customer Officer, ServiceNow is heading toward the concept of AI-only service desk departments for functions such as IT support. “We’re leveraging our own platform because it has AI capabilities, data capabilities and workflow execution capabilities, and I do think without all three you can’t be successful,” he adds.

That the early focus for AI adoption seems to be primarily on improving the customer and/or employee experience is reflected in the results of our survey, too. Asked which proprietary enterprise data sources are the most well-prepared for training AI models and generating insights, our respondents rated customer data (including interactions, transactions and feedback) the highest at 71% followed by employee data (HR records, productivity metrics, skill development) at 64%. Indeed, these two types of data were considered much more AI-ready than other forms including product and operational (62%), financial (60%), supply chain/vendor (44%), and legal and compliance data (38%) See Fig. 1 on previous page.

However, according to Philips’ Chief Innovation Officer Shez Partovi, most enterprises are still

One of over 65 DataBank data centers in the US...



at a very early stage of AI adoption, about the same point as ‘64Kb dial-up’ in the journey from dial-up internet connectivity toward full-fiber broadband. He believes that AI adoption can be viewed in three steps or ‘lifts.’ While the first lift is about automating administrative work and the second about augmenting the skills of employees to tackle higher-level tasks and problems they previously couldn’t address, it is the ‘third lift’ where the greatest ROI will be achieved. This is where companies take on tasks that were previously out of reach due to cost or technological limitations.

Currently Philips is using the augmentation approach in various areas, including software engineering and supply chain management. It also is evaluating the potential for full automation in certain supply chain activities and doing early experiments with AI automation in customer service.

Overcoming barriers to entry

Of course, achieving a return on AI investment is far from straightforward. As our survey shows, enterprises face several significant ‘blockers’ along the way – including technical, financial and practical constraints.

Fig. 2

Which of the following would you consider to be a major blocker to achieving ROI from your AI infrastructure?




However, the good news appears to be that poor data no longer appears to be the impediment to AI implementation that it once was. Asked ‘which are the major blockers to your organization achieving ROI from its AI infrastructure investments?’ (See Fig. 2 above), only 20% cited ‘poor data quality, availability and/or labelling.’

Furthermore, when asked ‘how would you assess your organization’s overall data readiness for AI adoption’ (See Fig. 3 on next page), one in three respondents (32%) stated that their ‘data and processes are well organized, governed, and ready for AI training and analysis,’ and a further 60% considered data foundations either ‘OK’ (26%) or ‘good but requiring some improvements’ (34%). Less than one in ten (8%) considered their ‘processes and data immature, fragmented and siloed, making AI adoption challenging.’

Indeed, it seems many enterprises have already invested heavily in ensuring their data is well-prepared for AI. For example, Visa has spent \$3 billion over the past 10 years to rebuild its data platform to be more AI-friendly, creating a “data highway” with multiple lanes to transport data at different speeds. Meanwhile, Philips has established a common data lake in Databricks which it has been operating for the past five years.

However, despite the data maturity shown by many of these larger enterprises, there are still significant challenges that need to be overcome before achieving ROI from AI infrastructure investments. Cited as the main blocker to success by the respondents to our survey are ‘integration, scaling, and or deployment issues’ (24%), indicating the challenges are not just technical, but cultural, with leaders needing to get their employees’ support when it comes to introducing new AI technology.

Fig. 3



How would you assess your organization's overall data readiness for AI adoption. Select the option that best describes your data maturity.

- We have good data foundations but need some improvements in standardizing processes and data governance/structure.....34%
- Our data and business processes are well-organized, governed, and ready for AI training and/or analysis
- Our data and processes are ok, but significant investments are needed in clean-up, governance, and standardization before AI can be fully leveraged
- Our processes and data are immature, fragmented and siloed, making AI adoption challenging


"Not unlike any other big change – whether it's technology related or not – it comes down to leadership," says ServiceNow's Chief Customer Officer, Chris Bedi. "It takes leadership courage to redefine a new mission for your department, new roles in your department and new ways of working. It also takes leadership conviction to upskill the existing workforce."

Nor is it just upskilling the existing workforce that is the only challenge. Cited as a blocker to achieving ROI from its AI infrastructure investments by 19% of respondents to the survey is a general shortage of AI talent and/or expertise, perhaps explaining why we are seeing '[AI talent wars](#)' breaking out between some of the big AI companies, including Meta and OpenAI.

"When I think of the current greatest headwind when it comes to AI, it's about having the right talent in place, as well as having the right training and the right upskilling," says Philips' Chief Innovation Officer Shez Partovi. "It's people, process, platform in that order," he adds.

Finally, and perhaps not surprisingly, 'regulatory, governance and/or ethical challenges' are cited as a blocker to return on investment by 21% of respondents to our survey. ServiceNow's Chief Customer Officer Chris Bedi highlights the challenges of different 'sovereign data requirements' in each of the territories in which it operates, as well ensuring compliance with new legislation, such as the EU AI Act.

For Visa, too, there also are challenges of operating across multiple territories, especially when it comes to payment regulations. "We operate in over 200 countries around the globe," says Visa's Senior Vice President Data and AI, Sam Hamilton. "As you can imagine, the payment regulations are very different for each of them." However, he believes that AI can actually help them, and their clients, make sense of local payment regulations. "If an international client wanted to understand the rules and regulations of a particular region, they used to have to search through quite a few documents which took a lot of time," he explains. "But now they've been trained in an LLM, it's possible to get answers to niche questions very quickly, indeed."



Inside a DataBank data center. Our survey shows that over half of respondents expect an expansion in AI-dedicated co-location capacity...

AI evolution: Shift from third-party AI/LLM models to bespoke solutions

Looking at how enterprises are adopting AI it is clear there isn't a 'one size fits all' solution. Instead, it seems companies are adopting a multi-faceted approach – choosing different types of AI to address different areas of their business and different levels of automation.

Fig. 4

Over the next three years, how likely is it that your organization will adopt AI in the following ways?



Asked 'Over the next three years, how likely is it that your organization will adopt AI in the following ways' (See Fig. 4 above), the largest number of respondents (46%) said they would use third-party applications powered by AI – such as AI/LLM-enhanced CRM (Customer Relationship Management) systems and analytics solutions including ServiceNow's own platform.

Meanwhile, the simplest form of AI adoption – using third-party LLM-powered chatbots such as OpenAI's ChatGPT and Anthropic's Claude – is the most favored option of 43% of respondents.

Anecdotally, it seems, larger enterprises are adopting more sophisticated bespoke solutions along their AI journey as the use-cases move from simple automation tasks to more complex tasks that augment the roles of humans within the company.

According to our survey, licensed-versions of third-party LLMs with company-specific fine tuning/customization is the most likely form of AI to be adopted over the next three years

by 37% of respondents, while just over two in five (41%) favor the most complex form of AI adoption, 'developing custom AI models using proprietary data to create AI applications unique to your organization's needs.'

So, for example, to enable enterprise-wide AI, Philips has built an AI suite on top of the AWS (Amazon Web Services) cloud platform using Amazon Bedrock. This tool suite, along with the data layer, is used for various business activities including finance, HR and the supply chain.

One specific use case is in post-market surveillance, where Philips uses AI to identify trends in patients' health records to ensure high-quality products. Another is in revenue prediction, where models are built to forecast sales pipelines.

Similarly, Visa has rolled out an internal secured version of ChatGPT which Sam Hamilton, Visa's Senior Vice President Data and AI, says is used by '82% of employees.' This is in addition to Microsoft's Copilot for Office and GitHub Copilot for developers.


Cloud vs. Co-location of AI Workloads: Navigating the Infrastructure Landscape

As with the diverse range of AI models being adopted, enterprises are increasingly taking a varied approach to the underlying infrastructure that hosts them. This strategic flexibility allows businesses to optimize performance, cost, security and compliance based on the specific demands of each AI workload.

Understandably, perhaps, the public cloud and third-party SaaS (software-as-a-service) are easy starting points for AI adoption. Indeed, nearly two in three (64%) of respondents to our survey say their AI workloads are currently running in either the public cloud (49%) or in a private cloud (15%), with a further 15% favoring third-party SaaS/web-based platforms, such as Salesforce (See Fig 5). Conversely, on-premises, company-controlled, and external colocated data centers currently represent a much smaller part of the market, approximately 22% of an organization's AI workloads according to our survey (15% on-premises/company controlled and 7% co-location data centers).

Looking to the future, however, what's clear is that while cloud will remain an important part of the hybrid approach, the reliance on colocation as well as owned data centers will increase substantially, too. Asked 'how do you expect the distribution of your organization's AI workloads to change over the next five years?' (See Fig 6. on next page), over a half of respondents are planning to build either more AI-dedicated private data centers (31%) or expanding co-location data center deployments (22%).

Fig. 5




Where are your organization's AI workloads primarily running today?

- Public cloud (AWS, Azure, Google Cloud)49%
- Private cloud15%
- 3rd party SaaS/Web-based platforms (e.g. Salesforce, etc)15%
- On-premise, company-controlled data centers15%
- Colocation data centers7%



JP Laqueur, SVP of Marketing at Databank, chats with Chris Bedi and Raul Martynek at AI CEO Summit 2025.

Fig. 6



How do you expect the distribution of your organization's AI workloads to change over the next five years?

- Increase reliance on cloud for AI workloads.....**43%**
- Build more AI-dedicated private/on-premises data centers**31%**
- Expand colocation data center deployments for AI**22%**
- No significant changes planned**4%**

Clearly, one of the advantages of enterprises adopting a hybrid approach is that while less sensitive workloads – such as those leveraging public datasets – can reside in the public cloud, more sensitive applications can remain on-premises or in a colocation data center. Not only do colocation facilities, such as those offered by DataBank, provide robust physical security features – including 24/7 surveillance, restricted access controls (biometric scanners, security guards), and secure equipment cabinets – they are also certified with various industry and regulatory compliance standards such as ISO 27001, SOC 1/2, and PCI DSS.

Indeed, asked 'how critical are each of the following factors in your choice between cloud v colocation data center for AI workloads', 'security and compliance requirements' were deemed to be the most critical factor by 37% respondents to our survey – above other important factors such as performance (33%) and scalability/flexibility (27%).

As Philips' Chief Innovation Officer Shez Partovi notes: "In a de-globalized world, there is an increasing need to ensure that data is housed and processed in compliance with the specific country or jurisdiction, which is leading to a more decentralized approach." He highlights the example of China, where the data and operations must remain within the country.

ServiceNow's Chief Customer Officer Chris Bedi agrees: "We operate in every corner of the world and there are obviously sovereign data requirements," he says. "To make sure we are serving customers in each region, we need to have a geographically dispersed data center strategy."

By offering region-specific storage and processing, data centers ensure compliance with local data sovereignty laws, while integrating security protocols, such as encryption, monitoring, and identity management to protect against breaches.

Shez Partovi, Chief Innovation Officer at Philips, talks AI adoption with Raul Martynek and JP Laqueur at AI CEO Summit 2025.



Expanding geographical infrastructure

Inevitably, as AI tasks become more sophisticated, so too do the AI strategies of the enterprises using them. So, while the cloud is an obvious starting point for AI (with solutions such as ChatGPT, Anthropic Claude and Google Gemini), the reliance on colocation and privately owned data centers is set to increase over the coming years. As noted earlier, over half of respondents expect to build more AI-dedicated private/on premises data centers or expand colocation data center deployments for AI.


Furthermore, when asked 'how AI will impact your organization's IT infrastructure deployment' (See Fig. 7), over three quarters of respondents expect their infrastructure to expand geographically into new regions, to be either closer to data sources (44%) or to end-users (32%). Only 13% claim there will be 'no significant impact on geographic distribution' while 11% say there will be 'consolidation into fewer, but larger, data hubs' (11%).

Conversely, at the same time as inference is becoming more decentralized, AI training is becoming more centralized. Asked whether the 'evolution from AI training and inference will require different strategies for your IT infrastructure' (See Fig. 8 on next page), nearly three quarters of respondents (73%) said that 'training will be centralized while inference will be more distributed'. Less than one in four (22%) said that 'both training and inference will follow the same geographic strategy,' and 5% 'weren't sure yet.'

As touched upon in the previous section, part of the reason for this geographically distributed AI inference infrastructure is inevitably down to legal requirements around data sovereignty and compliance in different countries. Asked 'which of the following would you consider to be major challenges in scaling your AI infrastructure geographically' (See Fig. 9 on next page), just under 1 in 4 (23%) cited 'data sovereignty and compliance.'

By distributing AI infrastructure, businesses can process and store data locally, helping them comply with data residency laws and strengthen their privacy postures, thus

Fig. 7



How do you expect AI adoption will impact your organization's IT infrastructure deployment?

- Require expansion into new regions to be closer to data sources **44%**
- Require expansion into new regions to be closer to end users..... **32%**
- No significant impact on geographic distribution **13%**
- Require consolidation into fewer but larger data center hubs **11%**

reducing the need to transmit sensitive information across geographic borders.

Nor is data sovereignty/compliance the only factor determining the geography of AI infrastructure. 'Network latency concerns for inference' are cited as the most critical challenge by 20% of respondents to our survey. Indeed, with many advanced AI applications, such


as those involving real-time decision-making (autonomous vehicles, smart city sensors, Industry IoT etc.) requiring super low-latency, processing data much closer to where it's consumed and generated is becoming mission critical. Importantly, by reducing the time it takes the data to travel to a data center and back, it's possible to enable much quicker responses for these types of advanced applications.

Meanwhile, other challenges mentioned in our survey include 'security risks in distributed AI' (27%), 'cost of infrastructure' (19%) and 'power and energy constraints' (18%). Finally, when it comes to training AI, being able to access GPU infrastructure is also seen as exceptionally

important, cited as the most critical factor by one quarter of respondents (25%).

According to DataBank's Martynek: "It's not just access to GPUs that matters though. You also need somewhere to put them. Not all data centers can accommodate the power and cooling requirements of these infrastructures and new facilities take time to deploy. That's why we developed a Universal Data Hall Design for all our new data center builds. It allows us to configure each data hall more quickly in a new facility for whatever infrastructure need that market might have – from traditional air-cooled configurations capable of 15-30kW per cabinet, to liquid-cooled systems capable of 100-200kW per cabinet."

Fig. 8




Will the evolution from AI training and inference require different geographic strategies for your IT infrastructure?

- Yes, training will be centralized while inference will be more distributed **73%**
- No, both training and inference will follow the same geographic strategy **22%**
- Not sure yet **5%**

Fig. 9

Which of the following would you consider to be major challenges to scaling your AI infrastructure geographically?





Conclusion: How advances in AI applications are shaping infrastructure

What's clear is that, when it comes to AI, we are still in the early stages of implementation – or inference. As noted earlier, if we use the analogy of the transition from dial-up to broadband we're still at the '64K stage,' in other words a long way from the high-speed, full-fibre broadband many of us enjoy today.

Clearly, despite many enterprises having data that is more than ready for AI training and inference, much of the focus so far around has been on the quick wins, the low-hanging AI fruit that early implementations of the technology can provide. Says Raul Martynek, the Chief Executive Officer of DataBank, "The ability to eliminate administrative tasks and allow employees to focus on higher-level work is a clear sign of AI adoption."

Going forward, though, it's expected that the role of AI within enterprises will develop from simple automation tasks toward much more complex 'augmentation' tasks. Not only will this enable employees to do much higher-level work, but also allow them to become more productive. Many commentators believe that it's this stage – referred to as the 'third lift' – where organizations will see the great return on their investment as they take on tasks previously thought impossible, or, perhaps, not thought of at all.

When it comes to AI, it seems many enterprises are already taking a hybrid

approach with regards to data infrastructure, often using public/private cloud for less sensitive workloads and colocated or private data centers for more sensitive applications.

This looks set to continue over the coming years as the geographic distribution of AI workloads continues apace. Indeed, the growing need for data sovereignty, coupled with the technical need for infrastructure to be closer to data sources and end users for advanced AI applications requiring low-latency, will inevitably increase the demand for on-premises and colocated data centers.

Concludes DataBank's Raul Martynek, "This is why it's so critical to have data center partners with a wide footprint and not just facilities in two or three availability zones like Santa Clara, CA and Ashburn, VA, and it's why DataBank intentionally targeted 25+ underserved markets like Dallas, Salt Lake City, Denver, Atlanta, Minneapolis, and Kansas City, allowing our customers to put their infrastructure within 100 miles of more than 60% of the U.S. population."

Methodology

Results for the survey were collated from a national (US) online poll of 304 senior professionals (director level and above) working at companies with 501 or more employees.

A survey of 277 senior professional was conducted by Propeller Insights between June 4 and June 13, 2025, on behalf of Kisaco Research. A further 27 respondents were surveyed by email directly from the Kisaco research database, making a total sample size of 304.

Respondents were sourced through a vetted online research panel and selected based on demographic criteria as well as company size and job title. All participants self-verified their qualifications within the survey to further confirm seniority and organizational fit. The margin of sampling error is ± 6 percentage points at a 95% confidence level.

In addition to the survey, the report also contains interviews with industry executives including Shez Partovi, Chief Innovation Officer, Philips, as well as Chris Bedi, Chief Customer Officer ServiceNow, and Sam



Shez Partovi and Raul Martynek at AI CEO Summit 2025.

Hamilton, Senior Vice President Data and AI, Visa. These were recorded in conversation with DataBank CEO Raul Martynek at the Kisaco AI CEO Summit in Palo Alto in April 2025.

About the survey respondents

This survey gathered responses from 300+ individuals, all of whom are full-time employed.

The age distribution shows a concentration between 41 and 50 (38%). The vast majority of respondents (85%) work in IT/Technology departments, followed by Supply Chain/Manufacturing (7%), and Procurement/Sourcing (7%).

In terms of company size, a substantial 38% work for organizations with 1,001 to 5,000

employees, and another 24% are in companies with 5,001 to 10,000 employees. Regarding job titles, "Director" is the most prevalent, accounting for 72% of responses, with "C-Level" positions making up 22%. The primary industries represented are Science/Technology (21%) and Manufacturing (13%), with High Tech also significant (11%).