Harvard Business Review
ANALYTIC SERVICES

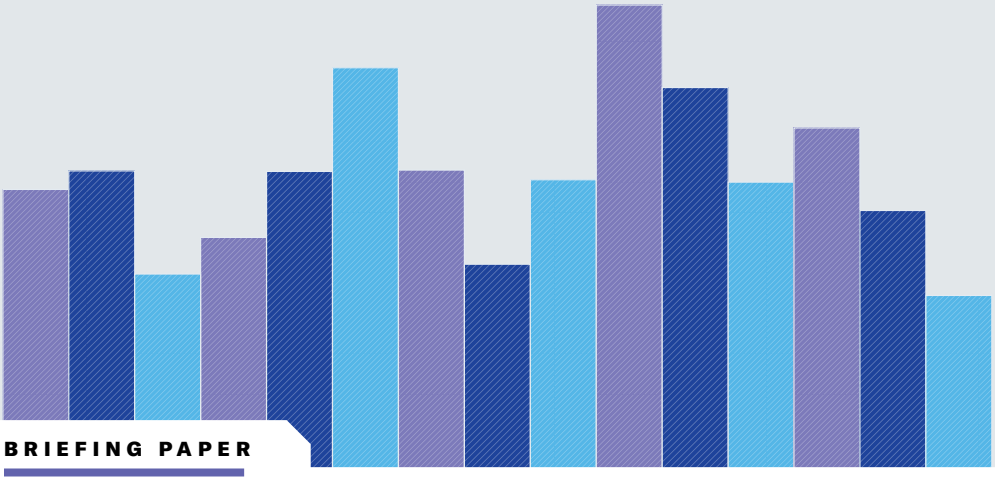

BRIEFING PAPER

Asia Pacific's AI Edge:

Unlocking Growth with Open Hybrid Cloud

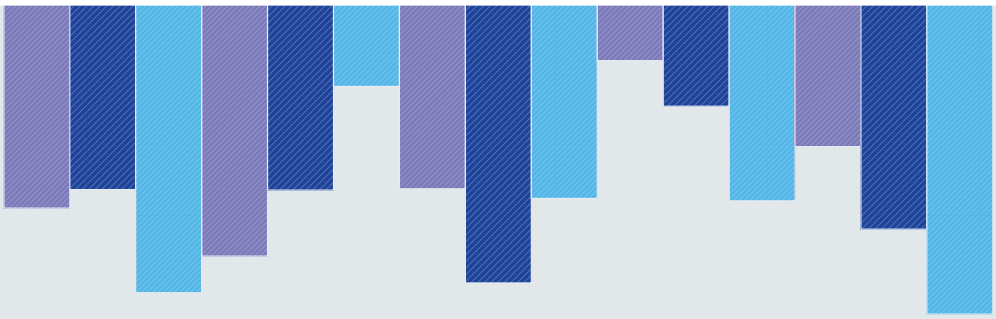

Sponsored by

Red Hat | intel

SPONSOR PERSPECTIVE

Vincent Caldeira
Red Hat
Chief Technology Officer, APAC

Unifying AI and Application Development: Navigating Tomorrow's Challenges Today

At Red Hat, our vision for the future of technology rests firmly on the foundation of integration and innovation. Recently, I've observed the seismic shifts caused by generative AI technologies in the IT landscape. Their integration into software development isn't just an enhancement—it's becoming a necessity.

The fourth pillar of our hybrid artificial intelligence (AI) strategy, "Unify AI and App Development," is particularly crucial. It's not just about incorporating artificial intelligence into software or vice versa; it's about creating a symbiotic environment where both can evolve simultaneously, driving efficiency and innovation.

The Dual Path of Integration
Integrating AI into software development processes has led to tools like AI coding assistants, which have fundamentally changed how developers work, making coding more efficient and less error-prone. However, the integration of software development into AI is equally critical. This approach involves developing platforms that not only support but enhance the capabilities of AI systems through better data management, increased processing power, and more sophisticated algorithms.

Our Approach: The AI Software Factory
At Red Hat, we advocate for an "AI software factory" approach. This isn't just a methodology; it's a transformative framework that integrates AI into the very fabric of software development. This model includes automated decision making and continuous optimization, which leverage AI to enhance operational processes, making them more dynamic and efficient. We focus on continuous improvement of development processes, where AI helps streamline operations and reduce inefficiencies, thereby accelerating innovation.

Future-Proofing Through Open Source
A significant aspect of our strategy involves leveraging open-source technologies. This effort not only prevents vendor lock-in but also ensures that our platforms remain at the cutting edge of technology, adaptable and robust enough to meet the challenges of tomorrow. Leveraging open source, particularly with regard to the access to models and the data generation process for the instruction-tuning phase of model training, helps address innovation challenges that often arise during large language model training by promoting a community-driven approach to development that accelerates the pace of innovation and supports a broader adoption of new technologies.

Conclusion: A Call to Action
As we continue to navigate the complexities of modern technology landscapes, our focus remains on supporting enterprises in their AI adoption journey. The fusion of AI and application development holds the key to unlocking new realms of possibility, making it viable for enterprises to safely tune and integrate AI based on their curated data assets. It's an exciting time to be at the forefront of technological change, and at Red Hat, we are committed to leading the charge, ensuring our clients are well equipped to turn these possibilities into realities.

This unified approach isn't just about enhancing current capabilities—it's about setting a new standard for what technology can achieve. Let's explore together how this integration can revolutionize your operations and help you stay ahead in an ever-evolving digital world.

Asia Pacific's AI Edge:
Unlocking Growth with Open Hybrid Cloud

Generative AI (gen AI) demands mammoth amounts of data and computational power to train and run its models—and that is only possible with a robust cloud infrastructure. The cloud alone provides the scale, flexibility, agility, and speed needed for gen AI data, workloads, and applications. Choosing between public and private clouds, however, is tricky: public clouds are not fit for every purpose, and building a private cloud can be a complicated and expensive endeavor. Going hybrid—using private and public cloud components—allows a single computing platform to operate across multiple clouds.

A January 2024 study by Bangalore, India-based Infosys, an information technology consulting firm, shows that companies in the Asia-Pacific (APAC) region plan to increase spending on gen AI by more than 140% to $3.4 billion in 2024—well ahead of Europe's forecast growth rate of 115% and more than double the growth expected in North America (70%).[1] Those figures suggest that organizations in high-growth gen AI regions such as APAC are now faced with an infrastructure issue as their use of the technology expands.

"I look forward to more open-source software out in the open market, particularly for gen AI," says Richard Turrin, author of the books *Cashless: China's Digital Currency Revolution* and *Innovation Lab Excellence* and a fintech, artifical intelligence (AI), and innovation consultant. "If you look at the growth in digital in APAC since 2015, it's exploded. The region has prospered, and a large portion of this prosperity has come from digital. Without open source powering behind the scenes, the development of countries in the APAC region, particularly in gen AI, will be withheld."

HIGHLIGHTS

76% of survey respondents say their organization is **using**, **piloting**, or **exploring generative AI** (gen AI) for business purposes.

69% expect employees at their organization to be **more productive as a result of using gen AI**, but only 16% have had that experience.

68% say they expect employees to **complete tasks faster as a result of using gen AI**, but only 18% have had that experience.

Gen AI's potential to improve customer experience, drive efficiencies within the business, and assist in the development of innovative products and services is moving C-suite executives to invest in ways to advance their use of the technology. In December 2023, Harvard Business Review Analytic Services fielded a global survey of 500 members of the *Harvard Business Review* audience who are familiar with their organization's current state regarding the use of gen AI and decisions made about it. The survey found that 76% of respondents say their organization is using, piloting, or exploring generative AI for business purposes.

Now the going will get tougher because organizations with legacy technology and IT infrastructure aren't adequately prepared for gen AI. Myriad hurdles, such as finding talent, also complicate the gen AI situation. More benefits from using gen AI have emerged, but so have more barriers to implementing it correctly.

"Eighteen months since ChatGPT entered gen AI into part of the everyday lexicon, people are now getting clearer about what some of the more meaningful 'low-hanging fruit' use cases are, as well as some of the use cases that might need more time to work out," says Kamayini Kaul, group chief data officer at Standard Chartered Bank.

This report examines how gen AI is being used by organizations in the APAC region and how the barriers and challenges they face in their efforts to deploy the technology are emblematic of what companies globally must contend with. The report also delves into gen AI's specific impact on enterprises in the APAC region and how organizations there are taking a consistent platform approach, using open hybrid cloud, to create the platform they need to maximize the opportunities the technology creates.

Great Expectations, Occasionally Realized

The possible benefits of gen AI adoption extend beyond individual organizations to entire industries. "The first reason companies should consider using gen AI is because change is inevitable," says Gabrielle Chou, a serial tech entrepreneur and associate professor of generative AI and disruptive business models at New York University in Shanghai. "And then, of course, gen AI holds potential that it is not only about improving productivity or reducing mistakes; it has the potential to disrupt any industry fundamentally."

Among survey respondents whose organization is exploring or using gen AI, the key benefits realized by the organization include improved employee productivity/efficiency (42%), the development of new capabilities/innovations (34%), and better utilization of data (29%). **FIGURE 1**

But some APAC organizations are going beyond these well-established benefits. In China, for example, gen AI is being used in the supply chain. "Gen AI has transformed supply chain and manufacturing through smart decision processes applied on a systematic basis that reduce errors and reduce production time," explains Chou.

Chou goes on to explain that the potential that gen AI holds is not simply about improving productivity or reducing mistakes; it could disrupt industries fundamentally. She uses

FIGURE 1

The Benefits of Using Generative AI

Improved employee productivity outpaces all others

What benefits has your organization realized to date from its use of generative AI (gen AI)? *Select all that apply.* (Among those whose organization is using, piloting or exploring gen AI.)

Base: 380 respondents.

Source: Harvard Business Review Analytic Services survey, December 2023

the metaphor of Spotify to illustrate. “Instead of listening to music on Spotify based on a recommendation that has been made by an algorithm, you will have a gen AI deejay producing a new beat that is only for you. And this will change not only Spotify, but it will also change the way you’re going to consume music. You will get used to having your own beat all the time. This is true for music; it’s true for social media. The giants of yesterday are not necessarily the giants of tomorrow.”

In the retail fashion supply chain, gen AI is used to customize advertising for individuals—and these choices are then filtered into outlets. “The fashion available in the shops is based on the trend that was analyzed and automatically re-created according to consumers’ online shopping preferences. Fashion retailers are using data to create the style that you want, knowing the quantity they need to order, and then produce a variety of visuals that resonate with different audiences,” Chou asserts.

Other new capabilities and innovations include the ability to provide customers with support that previously did not exist, which is something Hangzhou, China-based Alibaba is already doing. “We have Alibaba launching AI chatbots and AI-assisted sales for their clientele,” says Turrin, the author and innovation consultant. “So basically, if you’re a shopkeeper on the e-commerce platform Alibaba, you now have access to gen AI—built into your storefront—so that you can have a gen AI assistant to help you with questions from buyers about what you sell.”

Standard Chartered Bank is one of the largest network banks in the world, with around 90% of its profits coming from APAC, Africa, and the Middle East. Standard Chartered is at the forefront of investigating the potential of harnessing gen AI for both productivity and efficiency gains, as well as focusing on increasing customer engagement, and has developed a few prototypes.

“One prototype relates to streamlining the customer onboarding process and generating onboarding documentation,” says Standard Chartered’s Kaul. “We are also investigating having a question-answering bot on top of internal documentation to help our employees.” The bank is also prototyping the use of gen AI to support its strategy for net-zero carbon emissions in financing by 2050, using gen AI to synthesize large numbers of client reports on sustainable finance or climate risk, she adds.

The abundance of benefits arising from using gen AI makes it easy to understand gen AI’s growing appeal. But expectations do not always translate immediately into realized value. For example, 69% of survey respondents expect employees at their organization to be more productive as a result of using gen AI, but only 16% have had that experience. Similarly, 68% expect employees at their organization to complete tasks faster, but only 18% have had that experience. And while 66% expect employees’ quality of work at their organization to improve, only 11% say that this has happened. **FIGURE 2**

According to Turrin, how organizations choose to use gen AI is the deciding factor in whether they reap the benefits they

FIGURE 2

Reality Does Not Always Meet Expectations

There are gaps between expected and realized benefits from using gen AI

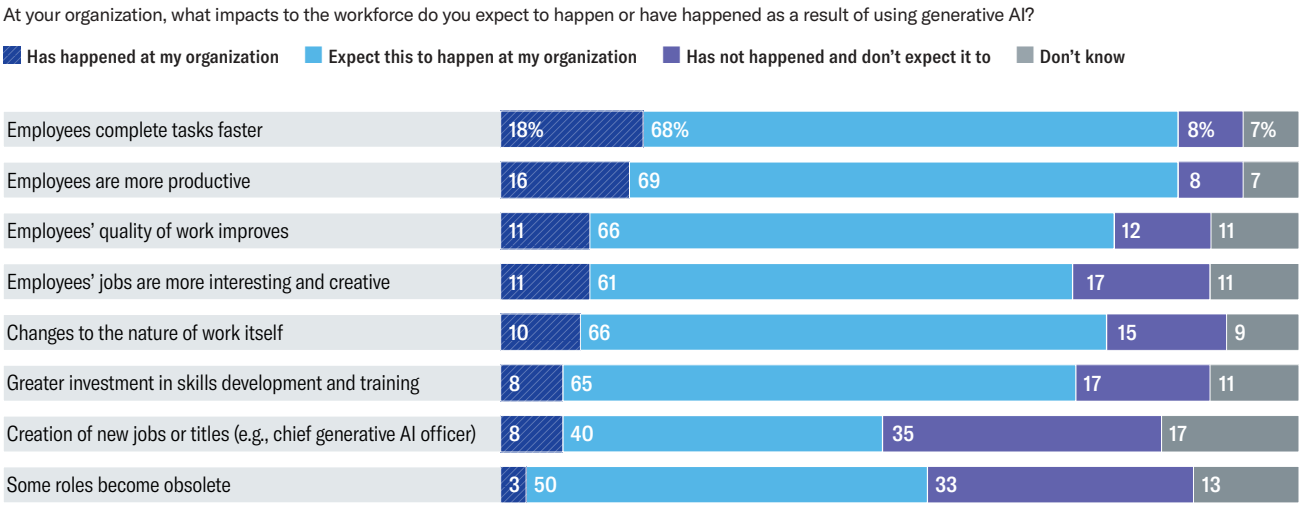

Base: 500 respondents.

Source: Harvard Business Review Analytic Services survey, December 2023

expect. "There is a common misperception that those with the best artifical intelligence will win in this great AI race. But this is missing something. And what it is missing is that it's not just about how good your AI is. It is how you use it, and what you're going to connect it to," he says.

Building a Foundation for Generative AI

All the expectations relating to gen AI mean little if there isn't the proper technology foundation for it. The strategy development process for businesses implementing gen AI needs to include decisions about which products and technologies will best support their enterprise needs. The decision between open-source or closed-source models depends on factors such as security, control over the source code, and the intended use case, and the choice is not always clear-cut. For many companies, building their own model does not make sense unless there are strategic reasons to do so, accompanied by large budgets and vast reams of data they can use. And right now, organizations are only at the start of the strategic journey. Only 5% of respondents say their organization has fully developed an enterprise-wide gen AI strategy with defined objectives, while another 38% are currently developing one.

There is a range of possibilities for utilizing open-source models, which starts at one end of the continuum with off-the-shelf usage, where organizations can simply adopt and implement pretrained models without modification. At the other end of the continuum, organizations with extensive resources and skills can build and train models entirely from scratch. A gradient of increasingly complex customization lies between these options, with the level of customization an organization can achieve depending on their available data, resources, and skills.

Fully open-source models include not just the model weights and code, but also the data that the models were trained on. The inclusion of training data is crucial, because it allows for full reproducibility and transparency. Incomplete transparency means that the training process cannot be reproduced, which limits the ability to verify and build on the model and raises potential bias and compliance issues.

"I would say everyone around the world is really passionate about gen AI, and like everywhere else, in APAC, you have companies trying to be the first to develop a foundational model," says New York University's Chou. "But there is not always a use case for that. And then you have companies that instead are trying to leverage an existing model."

China's technology conglomerate, Tencent, she explains, is an example of a company that has built its own foundational model, Hunyuan, which is available via Tencent Cloud for enterprises in China to use for testing and building apps via application programming interfaces. Hunyuan has both

"There is a common misperception that those with the best artificial intelligence (AI) will win in this great AI race. But this is missing something. And what it is missing is that it's not just about how good your AI is. It is how you use it, and what you're going to connect it to," says Richard Turrin, author and fintech, AI, and innovation consultant.

English and Chinese language processing abilities.[2] This model can help drive advancements in natural language processing technologies, benefiting a wide range of potential applications, from automated translation services to multilingual virtual assistants.

Closed-source models allow for greater control, as the source code is kept private by the developer. Hybrid models combine elements of open source and closed source to benefit from the strengths of both approaches, balancing the benefits of open-source accessibility and community-driven innovation with the control and specialized capabilities of proprietary systems. For use with gen AI, closed-source models offer control and security, as enterprises can protect their algorithms, code, and data, which safeguards intellectual property and reduces the risk of data breaches and other security risks.

Kaul outlines the process followed in Standard Chartered when choosing between open or hybrid solutions and models that the bank builds in-house. "We're working to see which use cases will require a purpose-built combination of open-source models and which use cases will require closed-source models that we would build in-house. We are increasingly moving toward a blend of open-source models and then augmenting those open-source models with the specifics of our data in a controlled environment." Often the decision rests on data. "We own the data—that's a core part of our business," Kaul explains. "And having models that really understand our data is crucial for us."

For example, the bank is using some of the publicly available large language models and then enhancing those models. "We are trying to see how far we can get with open source, and then tuning or boosting these models for our specific data," Kaul notes. "We believe that the future is going to need multiple models, with specific models to specialize in particular use cases."

Many enterprises fine-tune smaller AI models for discrete business processes. Fine-tuning smaller models costs less than training large models from scratch and can achieve faster, more accurate results. For example, fine-tuned models are used for fraud detection in financial services and for medical image diagnoses in health care—in both cases, trained on large, domain-specific data sets.

A cloud computing environment that combines multiple types of cloud infrastructures—typically public clouds, private clouds, and on-premises data centers—into a single, unified system allows for integration and management of data and applications across different environments. "For us, the cloud journey will likely stay dual cloud. We're not going to go to free public cloud as it doesn't make sense for us from an economic, regulatory, and complexity standpoint," says Kaul.

Turrin is a strong supporter of a hybrid-cloud approach for enterprises exploring gen AI. "Hybrid cloud is great stuff, keeping some data in your own data centers and pushing some off to the cloud. It saves money; you manage and run some of the data servers, and you ship some of [the data] off on the cloud. And you can balance the amount that's on cloud into a financially viable equation that makes sense for your company," he says. "The reality is that right now, no one is running gen AI on their own server other than the very largest companies. Except for a couple of banks and big tech companies, nobody can afford the chips [graphics processing units, or GPUs]. Somebody who has a hybrid system is in a better position than somebody who is still running on-premises servers for everything. And they would be in a cheaper position than somebody who is fully on cloud."

Beyond the choices between open and closed, cloud and on premises, choosing which vendors to partner with is a key strategic decision. Vendor risk can tend to increase as market concentration increases. For example, the gen AI industry's heavy reliance on a single manufacturer of GPUs creates a vulnerability to any disruptions in the manufacturer's production supply chain, and growing partnerships between large software and gen AI firms raise concerns about vendor lock-in and choice limitations.

There are many things to consider when choosing vendors, according to Kaul. "Classic considerations include the maturity we see from the vendors and the cost versus 'time to internally develop' trade-off," she explains. For example, general-purpose document parsing—the process of analyzing and extracting structured information from unstructured or semistructured documents—is not an activity that makes sense for Standard Chartered to develop in-house. But this choice looks different when it comes to the kinds of documents that are used in the bank's core business, such

as those involved in capital markets, custodial services, and trade financing. When the activities are core to the bank's strategy and work, building in-house solutions makes more sense because these applications are not readily available on the open market.

Other Challenges to Using Generative AI

Enterprises looking to develop a strategy for gen AI implementation will have to tackle a range of challenges over and above the key tech-related strategic choices. They need employees who are equipped with the skills and knowledge to both develop and use gen AI effectively, which, given the newness of the technology, is a tough ask. There are also worries about using data incorrectly or in a way that contravenes legislation—legislation that is constantly evolving. Further, managers and executives who face a number of business priorities do not always see gen AI implementation as a priority.

Navigating this constantly shifting landscape is daunting. Among survey respondents whose organization is not currently using gen AI, many cite a number of barriers to implementation. Topping the list is a lack of talent with the necessary skills/knowledge (48%), followed closely by data privacy/security concerns (47%) and leadership that doesn't see the use of gen AI as critical at this time (44%). Forty-three percent cite current IT/data infrastructure not being prepared for gen AI integration as a close fourth. **FIGURE 3**

APAC organizations are also adjusting their technology strategies according to the high demand for tech talent. Kaul outlines how the geographic presence of Standard Chartered contributes to its gen AI talent strategy. "To quote Nvidia CEO Jensen Huang, increasingly in the next decade, you're going to see China and India become the front door for AI talent, pretty much for the globe. There are about 2.3 billion people in just those two countries with an increasing bankable customer base as well as an avid workforce. For us, that's an opportunity, because Asia is one of our key markets and one where our heritage spans 160 years," she explains. "So, we expect to leverage that on the talent side as well as on the strategic partnerships side."

Meanwhile, the data privacy issues surrounding gen AI are particularly vexing in certain industries, such as the banking sector, where customer data is highly confidential. "In New York, certain financial institutions have blocked the IP addresses of gen AI providers because they don't want their employees actually using it, for fear that they will use gen AI on a corporate account or personal account, the personal data [will get] out into the public, and the company is in a world of trouble," says Turrin.

Kaul emphasizes that Standard Chartered takes data privacy extremely seriously. "We're doing a lot of work with risk owners

FIGURE 3

Lack of Talent a Key Barrier

It's the top—but not only—barrier preventing organizations from using gen AI

What barriers are preventing your organization from using generative AI? *Select all that apply.* (Among those whose organization is not currently using gen AI.)

Base: 120 respondents. Not shown: 0% none.

Source: Harvard Business Review Analytic Services survey, December 2023

across the bank to identify new risks—additional risks—that did not exist previously,” she asserts.

Data considerations are taken so seriously at the bank that they drive the strategic choices that the organization is making around gen AI investment. These strategic choices include whether the bank will buy off-the-shelf applications or build applications and where it might partner with companies that build or coinvest, when it comes to gen AI innovations or ventures. “There is AI we are going to buy, AI we are going to build in-house, AI where we are going to partner, and AI [where] we are going to coinvest in solutions,” says Kaul.

She explains that when the bank decides to build in-house or coinvest, it first ensures that close supervision and control requirements are met. This mandate is largely because whatever data is in question is customer data. The bank takes its role as data custodian extremely seriously, Kaul asserts, and building or coinvesting means that the use of data stays within the bank’s governance bounds for compliance and consent.

A crucial tech consideration for implementing gen AI is GPU access, the cost thereof, and the vast amounts of energy consumed. GPUs are important to gen AI because of their ability to handle the massive computational demands of training and running gen AI models. “In APAC, access to a GPU is key,” Chou asserts. “Of course, there are some ways to access the GPUs, but the cost is ramping up. Plus, you need power and energy. This means that there is a battle for energy that is starting to take shape because gen AI now consumes between 40% and 60% of the energy of the data center. What this will mean into the future is that geo-strategically, Asia is not necessarily well-positioned.”

Turrin agrees. “There are major energy issues associated with gen AI; it consumes tons of energy. And right now, with the technology we’ve got, you couldn’t give gen AI to everyone if you wanted to because you couldn’t build electric plants fast enough,” he notes.

Making use of GPU-as-a-service partly addresses this cost challenge by treating GPU resources as a shared, scalable service. This approach allows for more efficient use of expensive GPUs across multiple gen AI projects, thereby improving utilization and providing scaling options. Taking a hybrid-cloud approach means that the solution can work flexibly across both on-premises data centers and public cloud environments.

To be sure, having the scale, flexibility, and other benefits that cloud can offer is a major consideration in the gen AI equation, but from a technology point of view, paucity of data can be a challenge, too. “One of the key problems in APAC is the lack of available data to train the models,” explains Chou. “Asia is relatively younger historically than the West in terms of the tradition of publication or the systematization of organizational information.”

Kaul cites language as another issue in APAC that the bank is tackling, as both a tech challenge and a customer service matter. “Language is relevant for us, because as a network bank, you can either take the approach that you service in English, as the dominant language globally, or you can go

"Early users of software are beta testers, and basically, gen AI has been pushed off into the open arms of society. We are all beta testing this," says Turrin, the author and consultant.

more regional, which is the approach we take," she explains. "For us, this has been particularly challenging because it means there are a lot of gen AI products we cannot just take out of the box."

Kaul makes the case that strategic choices go deeper than just being concerned about regulatory and legal implications. "The legal concerns are largely predicated on infringement, whether that's copyright, trademarks, or other forms of intellectual property," she says. "But on the ethical side, everyone—including the regulators—is trying to figure out how we should capture the ethical considerations beyond just conduct. We are taking a front-runner approach, where we have specifically emphasized AI ethics in the bank's code of conduct. We have been growing awareness and educating our people on how we are going to put the right ethical checks in place as we progressively use more gen AI."

This ethical mission matters in banking, which is an industry that touches everyone's lives. "Banks have a higher social responsibility for trust and privacy than most others and need to get gen AI right, the first time and every time," explains Turrin. "Early users of software are beta testers, and basically, gen AI has been pushed off into the open arms of society. We are all beta testing this."

The banking industry hasn't been as caught up in the gen AI hype as other sectors. "If you read the management consulting reports, you may think that gen AI is ready to go. But financiers and bankers are taking their sweet time on gen AI because they're being asked to beta test a software," Turrin explains. "They are the guinea pigs. And if the software goes badly or has an error or an issue, the bank is the one that pays, not the consulting company and not the large tech companies that built the AI. Banks are taking gen AI on their own time schedule and their plan, not on management consultants' or big tech's plan, and that's a good thing—for all of us."

Looking to the Future

The excitement and the hype surrounding gen AI have spurred great expectations of value creation that are not immediately realizable because of a number of challenges that organizations need to tackle, especially when it comes to laying the technological foundation gen AI needs to succeed. Organizations developing their strategy for using gen AI successfully should carefully evaluate their key choices, especially regarding software, applications, and IT infrastructure. Then they have talent, customer privacy, and other issues to resolve.

To unlock the greatest value from gen AI, senior executives in businesses can wisely choose the platform that best meets their enterprise needs. Using a combination of cloud environments (public, private, and on-premises infrastructure) provides the flexibility to run gen AI workloads where it makes the most sense, whether on premises for custodial data and compliance reasons or in the public cloud for scalability and cost-efficiency. Open source allows for community-driven problem-solving at speed—and holds wider societal benefit, too, as it can contribute to regional advancement.

Given the forecast that APAC's spending on gen AI will likely surpass that of other regions in the world, the strides the region makes in implementing cloud technology and overcoming other barriers involving the use of gen AI could be telling. "In combination with Asia's global role, and as the highest GDP generator on the planet right now, I think that we're looking at something very significant," says Turrin. "Even if the U.S. is leading in the gen AI race at the moment, I think how quickly gen AI gets put to work, adopted, implemented, and tied into other digital systems in Asia will be something to watch."

Endnotes

1 Infosys Knowledge Institute, "Generative AI Radar: APAC," January 2024. https://www.infosys.com/services/data-ai-topaz/gen-ai-radar-apac.pdf.

2 Raj, Aaron, "Introducing Hunyuan, Tencent Cloud's Proprietary Foundation Model," Techwire Asia, September 2023. https://techwireasia.com/09/2023/how-can-business-benefit-from-tencent-clouds-proprietary-foundation-model-hunyuan/.

METHODOLOGY AND PARTICIPANT PROFILE

Harvard Business Review Analytic Services surveyed 500 members of the *Harvard Business Review* audience via an online survey fielded in November 2023. Respondents qualified to complete the survey if they were familiar with their organization's current state regarding the use of gen AI and decisions about it.

Size of Organization

22% 10,000 or more employees

32% 1,000–9,999 employees

12% 500–999 employees

21% 100–499 employees

12% 50–99 employees

Seniority

26% Executive management/ board members

43% Senior management

20% Middle management

11% Other grades

Industry Sectors

13% Technology

12% Education

10% Business/ professional services

10% Financial services

9% Government/ not-for-profit

9% Health care

All other sectors less than 8% each

Job Functions

17% General/executive management

12% HR/training

9% IT

9% Marketing/ communications

All other functions less than 8% each

Regions

51% North America

23% Europe

13% Asia Pacific

6% Latin America

5% Middle East/Africa

1% Other

Figures may not add up to 100% due to rounding.

Harvard Business Review
ANALYTIC SERVICES

ABOUT US

Harvard Business Review Analytic Services is an independent commercial research unit within Harvard Business Review Group, conducting research and comparative analysis on important management challenges and emerging business opportunities. Seeking to provide business intelligence and peer-group insight, each report is published based on the findings of original quantitative and/or qualitative research and analysis. Quantitative surveys are conducted with the HBR Advisory Council, HBR's global research panel, and qualitative research is conducted with senior business executives and subject-matter experts from within and beyond the *Harvard Business Review* author community. Email us at hbranalyticservices@hbr.org.

hbr.org/hbr-analytic-services

MCCRE52280724