451 Research® | PATHFINDER REPORT

A Guide to Enterprise AI

USE CASES AND OVERCOMING KEY BARRIERS TO ADOPTION

COMMISSIONED BY

Hewlett Packard Enterprise

JULY 2019

About this paper

A Pathfinder paper navigates decision-makers through the issues surrounding a specific technology or business case, explores the business value of adoption, and recommends the range of considerations and concrete next steps in the decision-making process.

ABOUT THE AUTHOR

NICK PATIENCE

FOUNDER & RESEARCH VICE PRESIDENT, SOFTWARE

Nick Patience is 451 Research's lead analyst for AI and machine learning, an area he has been researching since 2001. He is part of the company's Data, AI & Analytics research channel but also works across the entire research team to uncover and understand use cases for machine learning. Nick is also a member of 451 Research's Center of Excellence for Quantum Technologies.

JEREMY KORN

SENIOR RESEARCH ASSOCIATE

Jeremy Korn is a Senior Research Associate at 451 Research. He graduated from Brown University with a BA in Biology and East Asian Studies and received a MA in East Asian Studies from Harvard University, where he employed quantitative and qualitative methodologies to study the Chinese film industry.

Enterprise AI is Here

After centuries of theorizing, decades of research and years of advertising, artificial intelligence (AI) has finally begun to make inroads into the enterprise. More than 50% of respondents in 451 Research's recent Voice of the Enterprise: AI and Machine Learning survey said they have deployed an AI initiative, have one in a proof-of-concept stage or plan to within the next year. Other survey data from 451 Research reinforces this trend: 32% of respondents cited the technology as a top IT priority in 2019, and 48% said that budgets for AI and machine learning projects will increase in subsequent budget cycles. AI is thus set to become a pervasive feature of the enterprise. In the same Voice of the Enterprise (VotE) survey, 92% of respondents said they believe AI will have at least a slight impact on their industry over the next two years.

What is driving this rapid pace of adoption? The truth is that recent advances in algorithms, the proliferation of digital datasets and improvements in computing – both increases in processing power and decreases in price – have come together to initiate a new breed of AI technology that is enterprise-ready. Nearly all organizations have an ever-growing mountain of data assets, and AI provides the means to analyze this resource at scale.

But these technical components only provide half of the answer. The other reason AI is set to become a staple in the enterprise is because it is a cornerstone in the digital transformation process: AI is an omni-use technology that can bring improvements in efficiency and insight to almost any business process – from customer service operations and physical and cybersecurity systems to R&D functions and business analytics processes.

This diversity of potential impact points means the technology can produce a multiplicity of organizational benefits. Early adopters of AI are already seeing the impact of the technology: 43% said the technology has improved customer experiences, and 37% cited gains in competitive advantage as a benefit (see Figure 1). It is the allure of these benefits that spurs organizations to adopt AI. Organizations want to capitalize on the advantages of this technology to improve key processes and give themselves a competitive advantage.

Figure 1: The benefits AI brings to enterprises

Source: 451 Research's Voice of the Enterprise: AI and Machine Learning 2H 2018

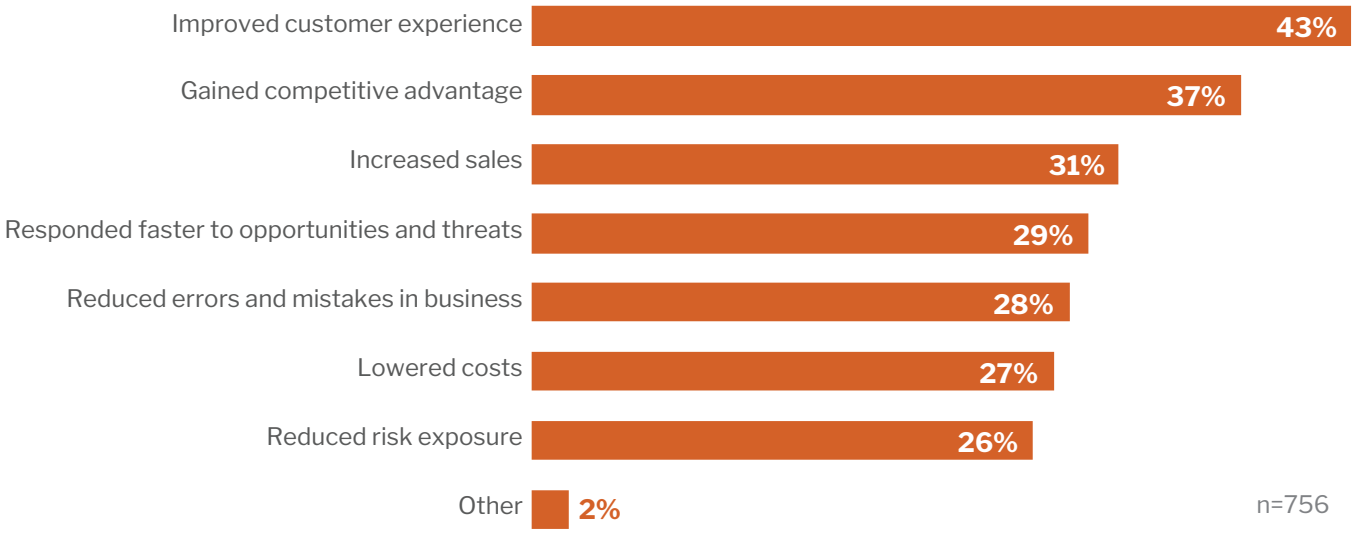

What is AI, and What Can it Do?

It is important to bring some nuance to our discussion of AI to better understand how the technology can transform so many enterprise processes. Artificial intelligence is an umbrella terminology encompassing a broad spectrum of technologies that mirror some aspect of human intelligence – everything from a setpoint-maintaining thermostat to an anthropomorphic robot reasonably can be defined as AI. However, as Figure 2 shows, AI encapsulates two more narrowly defined technologies – machine learning and deep learning. It is machine learning – and increasingly deep learning – that is driving the advances in AI described in the previous section.

Figure 2: A schematic representation of AI and its subsidiaries

Source: 451 Research

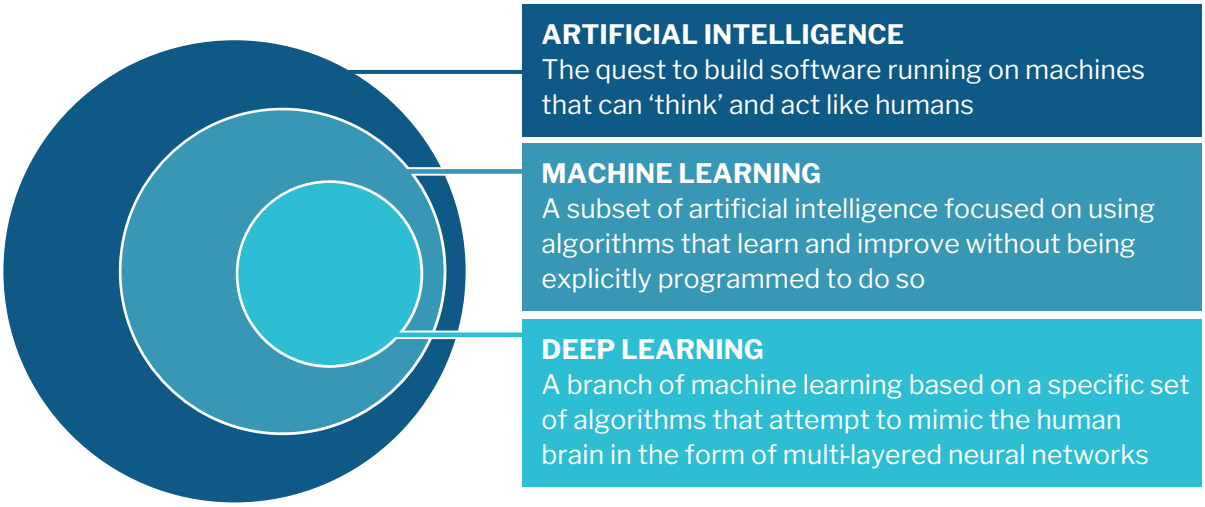

Machine learning is a technique employed toward the goal of creating a system with artificial intelligence. In a sense, machine learning is analogous to the process by which humans gain new capabilities: In the same way we learn a new task through repetition, an algorithm is optimized through repeated exposure to training data. With time, the machine, just like a human, gains enough proficiency in a task that it can apply this capability by producing new outputs based on novel inputs.

Another productive way to think about AI and machine learning is how the technology maps to fundamental human functions:

- Computer vision: Machines that can 'see'
- Natural language processing (NLP): Machines that can 'hear'
- Natural language generation (NLG): Machines that can 'write'
- Text to speech: Machines that can 'speak'
- Knowledge representation: Machines that can 'interpret'
- Robotics: Machines that can 'move' and perform other actions

These functions can then be applied to specific domains or combined in different ways to create a broad set of applications, ranging from facial recognition to virtual assistants or recommendation engines.

Based on this fundamental understanding of AI, it is clear that despite its tremendous upside, the technology is not easy to implement. In fact, there are three primary barriers to enterprise adoption.

1. **Lack of skilled resources:** AI and machine learning require a good deal of technical prowess to do correctly. The paucity of machine learning experts combined with the demand for their expertise means that most organizations experience a skill shortage. To address this issue, many are looking to upskill existing resources. Others are turning to new tools offered by vendors that seek to reduce the complexity of machine learning to facilitate enterprise adoption.
2. **Too many use cases:** Paradoxically, the scope of AI use cases can make it challenging for business decision-makers to evaluate how AI should be imbued into their organizations. In general, AI projects are most successful when they are narrowly focused and leverage high-quality datasets. Many enterprises find that working with a technology partner helps them identify the AI use case that best aligns with their business.
3. **IT infrastructure:** The current IT infrastructure of most enterprises will not be adequate to support future AI workloads. Burgeoning datasets will require enterprises to modify their data management strategy to ensure that the feedstock for AI initiatives is accessible and governable. A growing desire for real-time inferencing, particularly on IoT devices, will stretch AI from the cloud to edge environments. Enterprises can address some of these infrastructure issues by adopting new machine- and deep-learning-specific hardware, but enterprises need to adopt a new paradigm for their IT infrastructure, one that places AI workloads at the core.

Top Use Cases in Key Industries

Financial Services

ONE-YEAR ADOPTION RATE	TOP USE CASE TODAY	OTHER NOTABLE USE CASES	MAJOR HURDLE
61%	Fraud detection	Customer service, compliance	Data access

The financial services industry has traditionally been quick to implement new technologies that offer a competitive advantage, so it's not surprising that the short-term adoption rate for the industry is quite high, at 61%. According to 451 Research's VotE: AI and Machine Learning data, 37% of financial services respondents said the technology helps them reduce risk exposure. Because they are in a heavily regulated industry, many financial firms are using machine learning to shore up their assets and networks.

In particular, financial institutions are excited about applying the technology to the perennial problem of fraud detection. As the digital economy proliferates, financial organizations need to be able to quickly and accurately assess whether a transaction is valid. Old rules-based systems might have done an adequate job, but they still missed a decent proportion of fraudulent activity – or worse, they mistook legitimate transactions as fraud. With machine learning, financial firms can train models on extensive sets of historical data to create next-generation fraud-detection systems.

Financial firms are also interested in using machine learning to improve compliance procedures. In the wake of data-privacy regulations such as GDPR, financial institutions need a way to identify and manage customer data assets. Machine-learning-enhanced tools provide a mechanism for quickly surfacing pertinent information so that organizations can demonstrate or maintain compliance.

But there's more to consider beyond security: 33% of financial services organizations are building AI-based chatbots to improve customer service, and 23% are using the technology to improve the targeting of marketing campaigns. As will be a common theme across industries, machine learning can be used to make interactions with customers and potential customers smoother and more personalized.

A primary barrier to successful adoption in the financial services industry is data access; 37% said this problem negatively impacts their AI initiatives. Given the number of datasets – and their proprietary nature – that financial services companies incorporate into their machine learning projects, it's no surprise that the combination of these assets represents a significant barrier. As AI systems become more complex, the interoperability of these resources will become a bigger concern for all organizations, especially those in financial services.

Healthcare

ONE-YEAR ADOPTION RATE	TOP USE CASE TODAY	OTHER NOTABLE USE CASES	MAJOR HURDLE
54%	Patient monitoring, clinician workflow optimization	Disease diagnosis	Skilled resources

Considering the transformative potential that machine learning could have on healthcare – both to optimize delivery systems and to reshape care practices – it's no wonder there is so much buzz about the integration of AI tools in the industry. According to 451 Research's VotE: AI and Machine Learning data, 54% of healthcare organizations said they plan to adopt the technology within the next year.

Right now, healthcare respondents are focused on the point of care. Patient monitoring is the most popular machine learning use case today at 45%, followed by clinician workflow optimization at 41%. By applying the technology to the growing volume of data generated by IoT medical devices and other lifestyle data, healthcare providers hope intelligent patient monitoring and clinician workflow optimization tools will allow doctors to focus on the higher-level decision-making that improves patient outcomes and produces more personalized care. While these efforts today are most likely at the remedial end of the spectrum, as the technology progresses, AI and machine learning could fundamentally alter the way healthcare is delivered.

> "We are always looking at patient outcomes, and one of the big things in healthcare is alarm fatigue…Have you ever walked through a hospital or been on one of the intensive care floors, and there's always things beeping and flashing going off? And a lot of times, the nurses just ignore it because they know that that beep doesn't mean anything…They get so used to ignoring the beep that they miss the one that matters…That's alarm fatigue, and I think it lends itself well to AI."
>
> **– HEALTHCARE ORGANIZATION**
> WITH $2.5-5BN IN REVENUE AND 10,000-50,000 EMPLOYEES

Disease diagnosis is another area where machine learning will most likely upend a legacy process. Currently, medical professionals rely on their own expertise and experience when interpreting data points to make diagnoses. In some cases, the process is straightforward. If your blood pressure is over 120, your doctor should diagnose you with hypertension. But other situations are more complicated: either the diagnosis requires weighing several data points, or the data itself is more ambiguous. Using historical datasets to train models should lead to AI systems capable of making more accurate diagnoses by synthesizing more data points than a doctor or interpreting data with more accuracy.

This process, of course, raises the issues of data privacy and security, two preeminient concerns in the healthcare space. Any non-compliant use of data or breach in security protocols could lead to massive financial and reputational repercussions, and the problem is compounded by AI systems that ingest data obtained by various methods often across different organizations.

Somewhat paradoxically, AI itself helps to address these two issues. Models trained to identify and classify document and data types can help organizations ensure patient data is shared compliantly, and models trained on user behavior can help detect and prevent unauthorized access to private data.

In the healthcare space, the predominant barrier to AI adoption is the lack of skilled resources, a problem cited by 39% of the industry. Healthcare delivery already requires significant expertise. Factor in the need for machine learning skills, and companies are looking at a very small set of workers; therefore, the healthcare sector represents a prime target for tools that democratize AI. Expanding the set of end users who can develop AI applications or employ AI systems means the benefits of the technology can reach a larger audience.

Manufacturing

ONE-YEAR ADOPTION RATE	TOP USE CASE TODAY	OTHER NOTABLE USE CASES	MAJOR HURDLE
42%	Prescriptive maintenance	Product design	Deploying AI in operational system

Each link of the manufacturing supply chain is a target for AI integration, which could lead to considerable improvements to an industry used to the disruptive potential of new technologies. Data from the VotE: AI and Machine Learning survey shows that the manufacturing sector lags other industries when it comes to adoption of the technology; only 42% will have integrated AI in some capacity within one year.

Predictive maintenance is the top use case at the moment – 36% of manufacturing respondents currently use the technology for this purpose, a result that makes sense given the problems (and opportunities) manufacturers face as they attempt to move from a pure product focus to a model that focuses on products plus ongoing services. Downtime is a major problem, with cascading downstream effects. If the tire assembly system breaks down, production comes to a halt, leading to reduced output and decreased profits. However, predictive maintenance is a natural use case for AI. Combining historical data about the tire assembly system with related data sources, such as weather, staffing levels, similar machinery results and even order pipelines, can create a model that predicts when a breakdown is likely to occur based on real-time data, leading to significant cost savings through reduced downtime.

A similar application of AI can help manufacturers transform their quality assurance processes. QA is typically a laborious – but important – step in manufacturing that involves a good deal of human intervention. By training a video-recognition system to detect faulty products, a manufacturer can reduce the number of defective products that enter the supply chain, which cuts costs and improves customer satisfaction.

Another interesting use case for AI in the manufacturing space is its application to the production design process. In theory, algorithms could be trained that enable the automated design or optimization of new products. This application could fundamentally change the way research is conducted or new products are designed.

Manufacturing processes are usually well-oiled operations, and any AI system will need to be similarly reliable. It is not surprising, then, that 27% of manufacturers said that operationalizing models is their biggest challenge. There has always been a tension between the operational technology (OT) and IT systems of manufacturers. Making sure the AI models have access to machine data at the right time is no small feat, but it is of the upmost importance in the manufacturing space.

Government

ONE-YEAR ADOPTION RATE	TOP USE CASE TODAY	OTHER NOTABLE USE CASES	MAJOR HURDLE
40%	Physical and digital security	Transportation	Limited budget

As is often the case with the public sector, where risk-taking is not rewarded, the pace of adoption of AI and machine learning technology lags other industries. The survey shows that only 40% of government respondents have plans to implement the technology within the coming year.

Digital and physical security are the primary use cases of AI in the public sector; 39% of government respondents said the technology is being used for digital data security, and 35% cited intelligence or surveillance analysis as a use case. Surveillance and video analysis have been a feature of enterprises since the advent of video camera technology, but artificial intelligence has the potential to supercharge this functionality in two ways. First, it can provide tremendous scale to surveillance systems. In the time it takes a single human to monitor a handful of video feeds or analyze a given amount of footage, AI software can process and review more or less an unlimited amount of video input. Second, artificial intelligence can improve these processes by enabling real-time capabilities; rather than post-processing archival footage, users could get analytical insights in real time.

Although security is obviously an important application of AI and machine learning for government entities, it is by no means the only avenue through which the public sector can benefit from these technologies. An important implementation going forward will be transportation initiatives, especially in the context of the smart city. The interconnection of public and private vehicles with transportation and parking infrastructure can help mitigate urban congestion, reduce fuel consumption and commute time, and improve quality of life for city residents. AI will provide the brain of the operation, taking data collected from various assets, analyzing it in real time and making recommendations.

In addition to lacking skilled resources, public sector respondents are also particularly inhibited by budget limitations, which 42% of respondents said hampers their AI initiatives. These organizations must spend their limited budgets wisely, so they will be drawn to off-the-shelf applications that address key use cases in proven ways. Turning to vendors that provide tools that abstract away much of the complexity of AI offers the best path forward for users in the governmental sector.

Societally Transformative Use Cases

As the number of AI applications grows, the technology will increasingly be an omnipresent feature of modern society. Its impact will be felt at both the home and the office, as well as the commute in between. This section details two AI use cases – autonomous vehicles and natural language processing – and explains how these technologies will transform society.

Autonomous Driving/Connected Cars

One of the most significant ways in which AI could impact society is through autonomous vehicles. Thanks to years of incremental improvements in data-collecting sensors, edge compute resources and decision-making rules engines, autonomy is now within reach of the automotive industry. According to 451 Research data, consumers are bullish on the technology, with 76% saying they expect fully autonomous vehicles to roll out in general availability within five years.

This timeline is probably on the aggressive side, but the idealism is understandable. Fully autonomous vehicles could transform mobility in the broadest sense. As 'driving' becomes less about literally driving, people will likely develop a different relationship with vehicles, which could become spaces of leisure or work in addition to travel. Even if fully autonomous vehicles are a feature of the distant future, semi-autonomous vehicles could, in the near term, lead to improvements in automotive safety and transportation efficiency. And all of these benefits are made possible in part due to artificial intelligence – the engine of autonomous vehicle technology.

More so than other AI applications, the safety and security of autonomous vehicles is a big issue for society. While 67% of respondents said they think fully autonomous vehicles will be safer than human-driven cars, only 16% of respondents said they are comfortable with fully autonomous vehicles, and a further 26% said they are comfortable with limited self-driving. Rationally or not, people are worried about the infusion of artificial intelligence into vehicles, and in the end, this fact might be a bigger barrier to adoption than any technical limitation.

Speech to Text and Natural Language Processing

The ability of AI not only to transcribe but also understand verbal inputs permits new kinds of interactional interfaces between humans and machines. According to 451 Research's Voice of the Connected User Landscape: Digital Transformation survey, 23% of respondents plan to adopt voice-activated interfaces within the next 24 months. The proliferation of smart home devices is a testament to the success of this AI use case: the ability to control connected devices or query information using natural language leads to a more frictionless user experience.

Smart home devices are just the tip of the iceberg for natural language processing; there are several enterprise applications of this technology. Again, one benefit of the integration of AI is improved efficiency. For example, rather than having a human laboriously transcribe interviews or meeting notes, an NLP system can create a transcript of a conversation in real time, potentially with more accuracy. This use case could be extremely beneficial to HR departments or other externally facing workers who need to generate an auditable trail. More generally, this technology could lead to improved employee productivity by allowing knowledge workers to perform higher-quality enterprise search.

Customer service is a natural point of integration for NLP, particularly through chatbot technology. Legacy service systems are not known for providing customers with top-notch experiences, but a key problem is that customers often don't know or can't articulate their issues to service representatives. NLP allows enterprises to better match a customer problem with the right solution.

"Chatbots give more personalized answers, more direct, rather than just throwing up an FAQ…The best ones that we've seen could seamlessly transfer control to a human and then back to the machine…The human can intervene if some of the questions are beyond the capabilities of the system to answer."

– COMPANY IN EDUCATION SECTOR
WITH $1-2.5BN IN REVENUE AND 10,000-50,000 EMPLOYEES

Conclusions/Recommendations

AI and machine learning are the next wave of transformative technologies. Similar to the way the internet changed the world, AI will continue to upend the way people work, communicate and play. The omni-purpose nature of AI means its impact will reach far and wide. There are no industries – and likely even no companies – that could fail to reap its benefits. Therefore, as organizations consider what their AI initiatives will include – and not whether they should have an AI initiative – here are some important recommendations:

1. **Carefully consider which AI and machine learning use cases to pursue:** We've already detailed AI's potential across a variety of industries and use cases. Rather than boiling the ocean, organizations need to select the best use cases for their situation. In particular, enterprises should consider what data assets they have and which areas could produce the largest ROI. It's always a great idea to utilize vendor expertise to help evaluate the best AI strategy.
2. **Leverage the skills and services of technology vendors to bridge the skills gap:** The technical proficiency often required for high-quality machine learning has created a talent deficit. Most vendors not only understand this issue but are also addressing it by providing simplified AI products that abstract away at least some of the complexity of machine learning. Enterprises should at a minimum embrace these tools, and if need be, develop deeper partnerships with technology and services vendors to overcome this problem.
3. **Upgrade/future-proof IT Infrastructure:** Enterprise AI and machine learning will change both the nature and magnitude of compute. Training algorithms at scale will require compute resources to be fast, flexible and accessible. Data environments will need to be more interoperable to accommodate machine learning projects. There is currently no panacea that addresses all of these issues, but organizations should investigate the advantages of high-performance compute technologies as they come to market. In addition, as edge data grows in volume and importance, AI will need to stretch to the network periphery to produce insight in real time. Finally, enterprises will need to secure this whole system to ensure the safety and privacy of users.

Competition – and the need for AI – is everywhere. Whether it's employing new strategies to stay ahead of business adversaries, scoring victories against life-threatening diseases, or unlocking the secrets of the cosmos, data is growing exponentially beyond the ability of humans to make sense of it. The pressure to harness AI for competitive advantage is increasing in nearly every industry. As mentioned earlier, over 50% of respondents in 451 Research's recent Voice of the Enterprise: AI and Machine Learning Survey said they either have deployed an AI initiative, have one in a proof-of-concept stage or plan to within the next year. Other survey data from 451 Research reinforces this trend: 32% of respondents cited the technology as a top IT priority in 2019, and 48% said that budgets for AI and machine learning projects will increase in subsequent budget cycles.

With HPE AI, whether we are working with top automotive manufacturers on building their autonomous vehicle programs with AI or bringing AI into our own operations to improve supply chain forecasting, HPE knows that getting AI right is hard. It starts with your data strategy and navigating a rapidly growing ecosystem of vendors to build complete solutions. The HPE AI approach is simple: We can help get you to the solution faster and extract value from all your data now because:

- We speed the design and deployment of your AI strategy with **expertise** to accelerate your AI journey.
- We give your data science teams instant access to the tools and data they need with the industry's only turnkey, container-based **software platform** that's purpose-built for AI.
- We build your models at scale with less complexity from edge to cloud with the industry's most complete, **AI-optimized infrastructure.**
- We help you execute your AI strategy fast, cost-effectively and with less risk leveraging our **pay-per-use consumption.**

Engage HPE experts to discover practical ways to apply AI for your enterprise and quickly identify how to maximize your data's value with the HPE AI Transformation Workshop. To learn more about HPE AI, visit *www.hpe.com/AI*.

Content provided by

Hewlett Packard Enterprise

About 451 Research

451 Research is a leading information technology research and advisory company focusing on technology innovation and market disruption. More than 100 analysts and consultants provide essential insight to more than 1,000 client organizations globally through a combination of syndicated research and data, advisory and go-to-market services, and live events. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

NEW YORK
1411 Broadway
New York, NY 10018
+1 212 505 3030

SAN FRANCISCO
505 Montgomery,
Suite 1052
San Francisco, CA 94111
+1 212 505 3030

LONDON
Paxton House
30, Artillery Lane
London, E1 7LS, UK
+44 (0) 203 929 5700

BOSTON
75-101 Federal Street
Boston, MA 02110
+1 617 598 7200

451 Research®