AI and machine learning in financial services

RESEARCHED BY

Omdia was established following the merger of Ovum, Heavy Reading and Tractica with the acquired IHS Markit technology research portfolio.

COMMISSIONED BY

Red Hat

Introduction: Financial services organizations increasingly realize the value of deploying AI workflows in hybrid cloud environments

Use cases

There are many ways AI can be utilized by firms in the financial services industry (FSI) to improve the customer experience, IT and business processes, workflows, and IT operations, some of which are uniquely applicable to the space.

Such use cases include fraud detection and mitigation, credit scoring and loan analysis, claims processing, tax filing and processing, and more. As the technology advances, so too will the number and caliber of use cases within financial services.

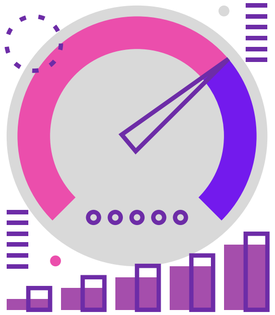

Best practices

Successful AI outcomes require a concerted effort of bringing together people, process, and platform. Many FSI organizations seek to accelerate their AI/ML workflows and scale outcomes. They look to cloud-based platforms that can meet a high level of deployment flexibility and offer the ability to utilize open-source tools and a sandbox environment in which they can develop and test their models. A hybrid cloud platform for AI and machine learning (ML) built upon Kubernetes is well-suited to meet this demand.

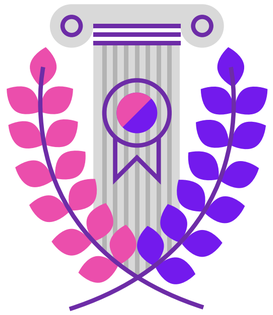

Success stories

This eBook, written on behalf of Red Hat, highlights three key customer success stories from companies across the industry and across the globe – Royal Bank of Canada, AXA France, and İşbank.

These success stories highlight the utilization of AI/ML (Red Hat® OpenShift® hybrid cloud platform for ML in particular) and the broad impact the technology can have on banks, insurance companies, and payments firms alike.

AI in financial services

Financial services companies embrace AI while being some of the most regulated, underscoring the efficacy of the technology

The financial services industry has been on the leading edge of technology adoption for years and that has been no different with Artificial Intelligence (AI) software. Some of the earliest use cases and success stories have come from large banks, investment firms, and insurance providers, which are also some of the most heavily regulated organizations worldwide. While in some industries regulation has historically been an inhibitor to technology adoption, financial services firms have been and continue to embrace AI to improve their internal and customer-facing operations and business outcomes while maintaining compliance with corporate, regional, and global security and governance policies.

Organizations that have shown to properly implement and onboard AI projects are able to scale the usage of AI across users, departments, and functions, furthering the cost and time saving benefits they can provide while also helping to scale business outcomes. While successful adoption of AI software has proven to be valuable for a business, there are many considerations and steps to be taken before adoption to ensure the most successful outcomes. It is imperative that AI is deployed in a trustworthy, governed, ethical, and secure manner so as to protect business and customer data from the ever-growing number of risk factors.

Banking 1

Insurance 2

Payments 3

AI aims to address many of the very challenges facing financial services companies

Financial services organizations are faced with a number of business challenges, especially given the new and persisting obstacles brought on by the COVID-19 pandemic that piled on top of the existing pressures from increased competition, customer demand for increasingly digital ways of doing business, and the rapid pace of technology innovation and change.

The number and caliber of AI use cases in financial services continue to grow at notable rates as the technology, when implemented and utilized properly, has the ability to help many facets of business. Financial services organizations' most common challenges relate to the desire to reduce operating costs, improve and optimize processes and efficiencies, improve customer service, and ultimately business outcomes. Improving efficiency and lowering costs are of top priority to address and many are utilizing AI to help do so, understanding that the long-term benefits outweigh the upfront cost and time investment in the technology.

Business challenges

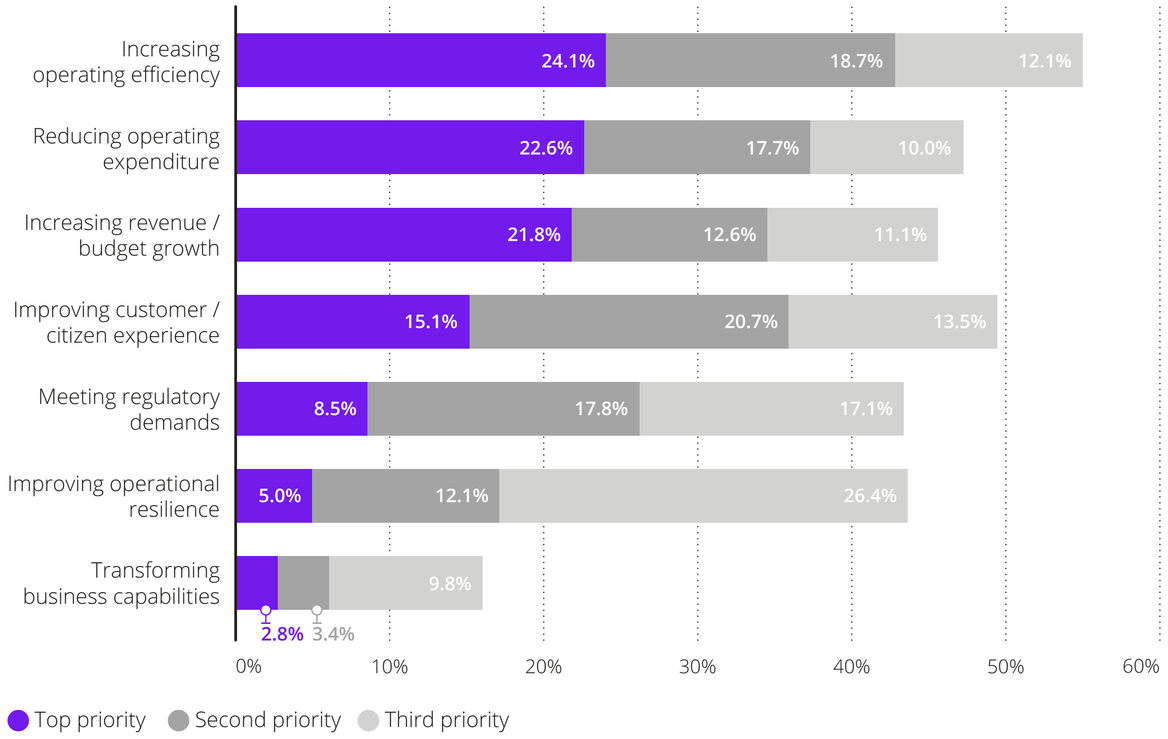

NOTE: N=2,285 QUESTION=WHICH OF THE ABOVE DO YOU PERCIEVE AS THE TOP-THREE BUSINESS CHALLENGES TO YOUR ORGANIZATION?

SOURCE: OMDIA, ICT ENTERPRISE INSIGHTS 2021. FILTERED BY: CORPORATE BANKING, FINANCIAL MARKETS, INSURANCE, MERCHANTS – PAYMENTS, PAYMENT ISSUERS/ACQUIRERS, RETAIL BANKING

Process improvements and security, both physical and digital, are key use cases for AI in financial services

AI use cases in financial services: 2021

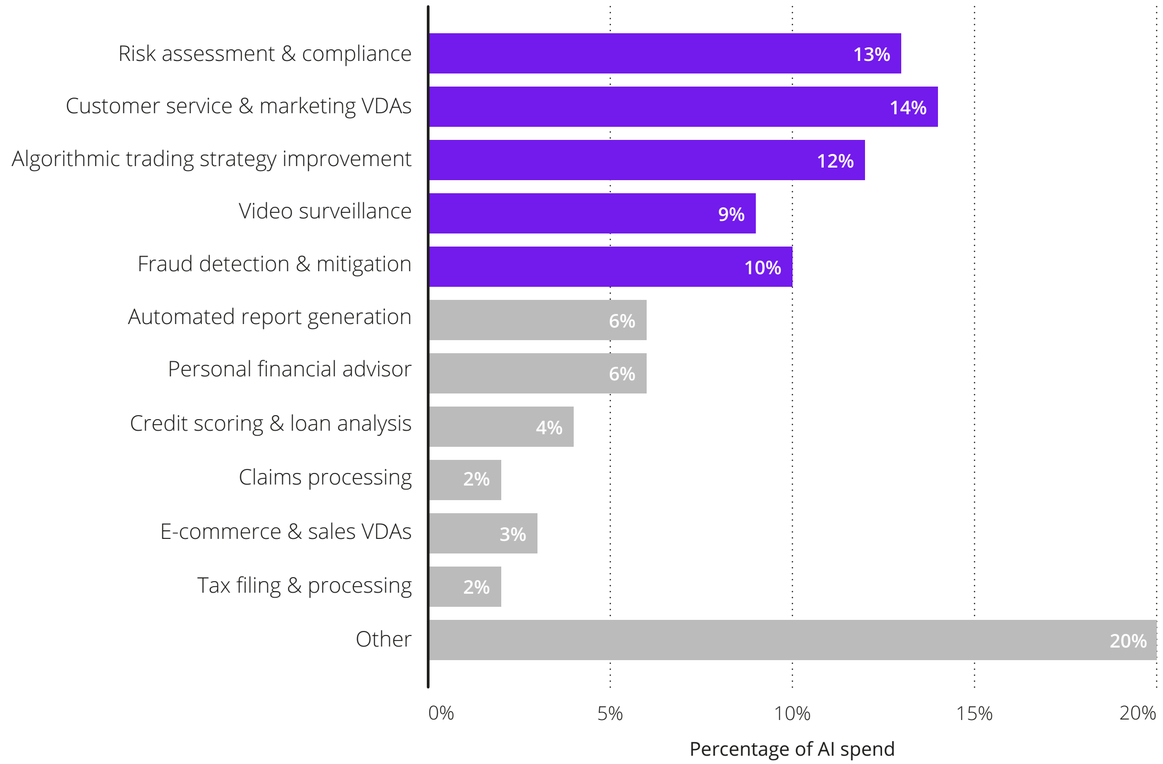

SOURCE: OMDIA ARTIFICIAL INTELLIGENCE SOFTWARE MARKET FORECASTS, 2Q21

Interestingly, the top five use cases for AI in financial services in 2021 are a mix of horizontal and vertical use cases. This underscores that while all businesses face similar overarching challenges such as customer service, marketing, and ensuring compliance, having tailored solutions available, particularly to a heavily regulated industry, is extremely valuable.

AI technologies such as natural language processing (NLP), intelligent document processing (IDP) and capture, and image recognition each play critical roles in the processing of structured, semi-structured, and unstructured data. For example, these technologies help insurance firms ingest, process, classify, analyze, and act upon vast amounts of documentation each day, saving thousands of hours per year – hours that can be reallocated to higher-level tasks, promoting better outcomes.

In order for enterprise IT departments to meet this growing demand within their organizations for AI solutions, infrastructure readiness is imperative.

AI in financial services is poised for strong growth, as are companies that embrace it

70% of financial services companies surveyed in 2020 were hands-on with AI, either piloting projects or scaling across their organizations after finding initial success. That number is only expected to increase in 2021 and onward as organizations embrace AI and the technology continues to advance and get smarter.

- As vendors continue to innovate and expand the capabilities of AI software and infrastructure, it will become increasingly applicable to more functions and processes, more users, and more use cases beyond what it is capable of doing now. As such, we expect the number and type of use cases for AI to expand across businesses and industries, becoming increasingly pervasive across the financial services landscape, with the potential to make lasting changes to how business is done for years to come. Companies that do not accept and onboard the technology soon will face increasingly strong competitive pressures.
- Omdia's ICT Enterprise Insights 2020/21 finds that 69% of financial services organizations view intelligent automation and AI as "significantly more important" (32%) or "more important" (37%) as a result of the COVID-19 pandemic and its impact on their businesses. This underscores the pent-up demand for a platform to build AI outcomes in the enterprise.

What is the state of AI deployment at your company today?

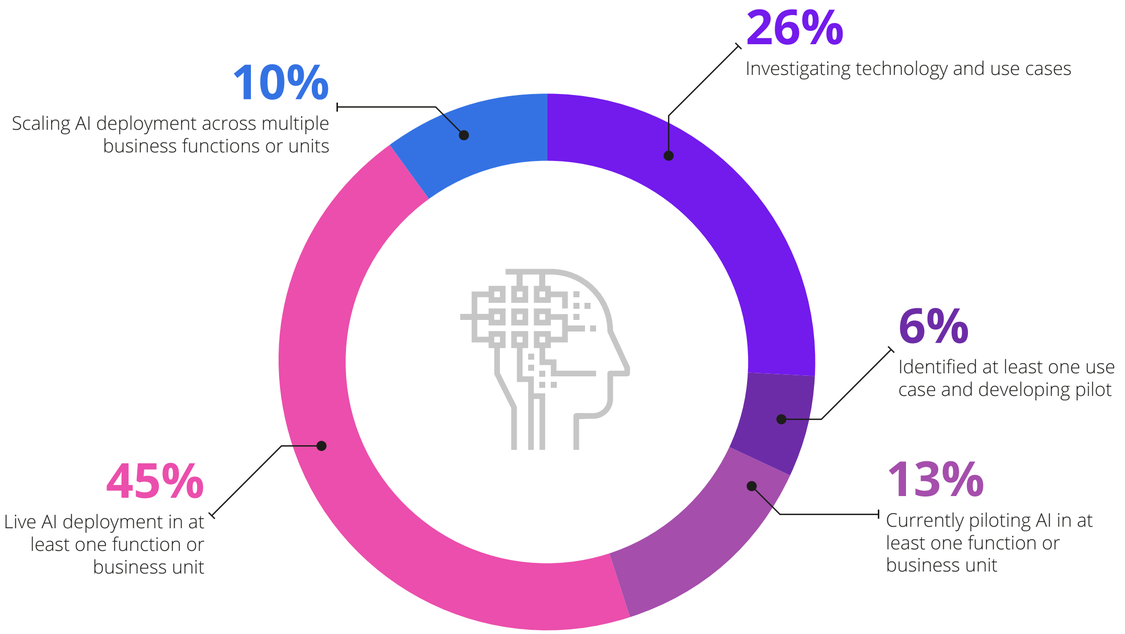

NOTE: N=31, FINANCIAL SERVICES COMPANIES ONLY
SOURCE: OMDIA AI MARKET MATURITY STUDY, JUNE 2020

Technical deep dive

AI goal posts are not yet fully rooted, challenging enterprises and their governing bodies alike to establish best practices

Building AI outcomes in the enterprise is no trivial matter. Even with the myriad of operationalized machine learning (MLOps), AI, and intelligent automation platforms available and the maturation of AI more broadly, the technology itself can only take a business so far as tools alone cannot fully address concerns of AI onboarding, deployment, governance, and scaling. Deploying AI at scale demands a concerted effort that blends people, process, and platform equally.

Moreover, enterprises must also factor in those they do business with, including partners, customers, and those down and up-stream. For instance, many U.S.-based businesses submitted Payroll Protection Program (PPP) loan applications amid the COVID-19 pandemic, some of whom utilized RPA to do so. While this sped the process for the businesses, it ultimately crashed the system and businesses are no longer permitted to utilize this method to submit applications to the Small Business Administration (part of the U.S. Department of the Treasury). This is a unique and small example of a greater challenge of onboarding entire industries and their governing bodies to establish governed AI practices with which individual organizations can comply and build towards.

Moving down from the industry level to the enterprise level, building and successfully deploying AI solutions in the enterprise requires involvement from business domain experts, data engineers, developers, data scientists, and IT operations to provide oversight and support from a centralized AI Center of Excellence (CoE) which then is better equipped to operationalize models at scale and keep them up-to-date. Data will not be perfect, but the point of operationalizing ML (MLOps) is to create a level of resiliency. Moreover, the CoE can develop a taxonomy that both IT and business domain experts can understand to optimize AI solutions for the desired business outcomes, establish governance, and adhere to macro-AI standards and ethics.

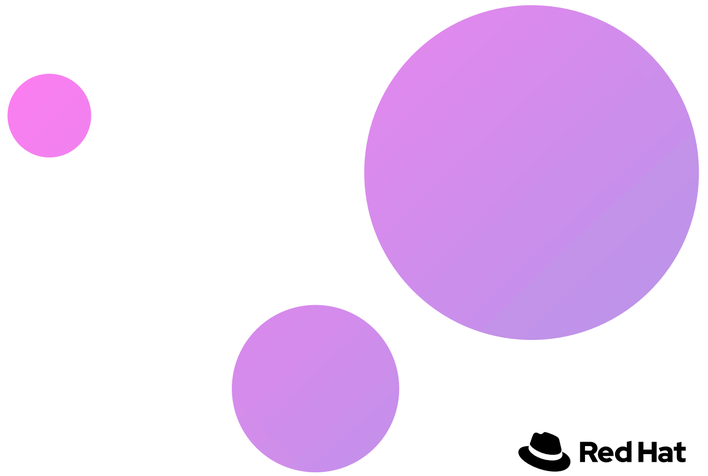

The road to long-term, successful, and scaled AI deployments is filled with obstacles...

Where does AI responsibility and decision-making reside within your company?

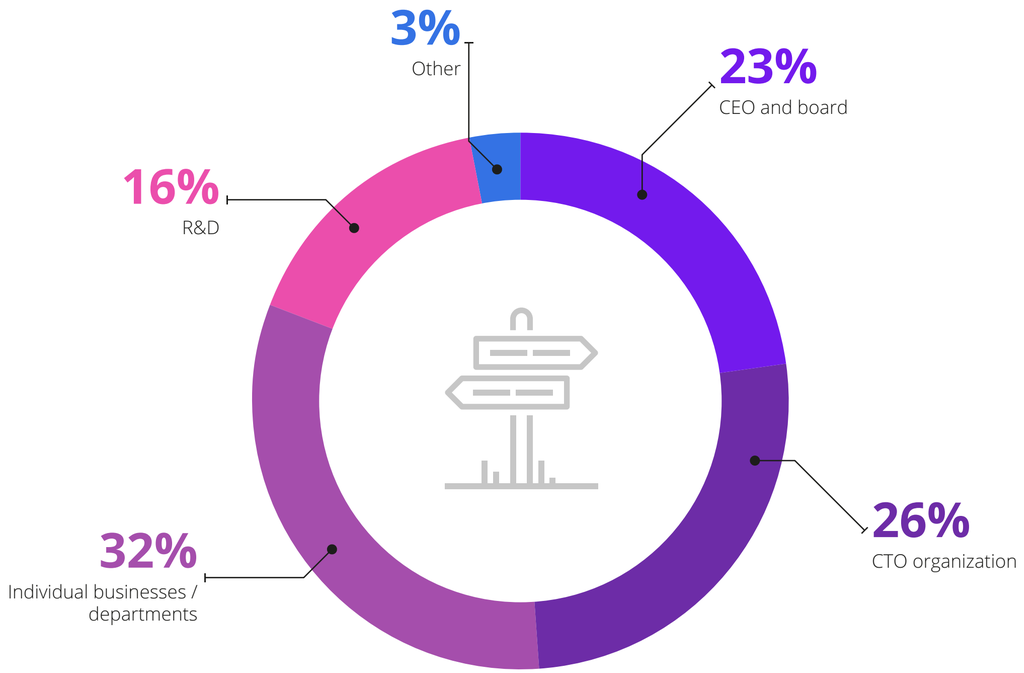

NOTE: N=31, FINANCIAL SERVICES COMPANIES ONLY
SOURCE: OMDIA AI MARKET MATURITY STUDY, JUNE 2020

Among financial services companies, there is little consensus regarding where AI responsibility and decision-making reside and whether or not executive buy-in is a factor as many individual departments are making their own decisions.

Executive buy-in and financial support alone will not guarantee AI project success. Companies face a multitude of challenges when embarking on their AI and ML projects, including:

- Lack of self-service capabilities for developers and data scientists to avoid waiting on IT to provision infrastructure
- Complexity to operationalize models
- Lack of collaboration across siloed teams
- Lack of formal KPIs against which to measure

These, and many other challenges, depending on how severe, delay or negate entirely the benefits and outcomes that can be realized from AI deployments.

...many of which are human-centric

Properly deployed, AI solutions can provide significant competitive advantages to financial services firms that cannot be matched by traditional technology or labor approaches. For example, retail banks and institutions can utilize AI to provide a personalized experience, incorporating self-service and guided tools to improve the digital customer experience. Democratization of AI is suited both for developers – both IT and developers with low- and no-code options – as well as for domain experts wherein they can work with developers to better understand what models are doing and measure them according to business KPIs. Enterprises must perform a balancing act between democratization of AI and security, governance, ethics, and control. The more users, the more potential risk, but also the greater the opportunity to generate business benefits as any application generating meaningful data for the organization should be sanctioned by data scientists to ensure the data can be managed and used to make business decisions. Data is a business asset, and for organizations to obtain the best possible outcomes from their AI and ML solutions, enterprises need to marry business domain expertise and IT acumen. An AI CoE can help nurture a digital culture where data scientists and developers work together and better enable MLOps and application development.

By in large, financial services companies are blending in-house and third-party expertise and capabilities to develop and deploy AI solutions, according to Omdia's AI Market Maturity Study. This underscores the importance of involving vendors from the earliest stages in the AI journey to ensure aligned values and the proper installment of an AI CoE and digital culture.

How would you rate your organization's progress in supporting a digital culture and organization structure?

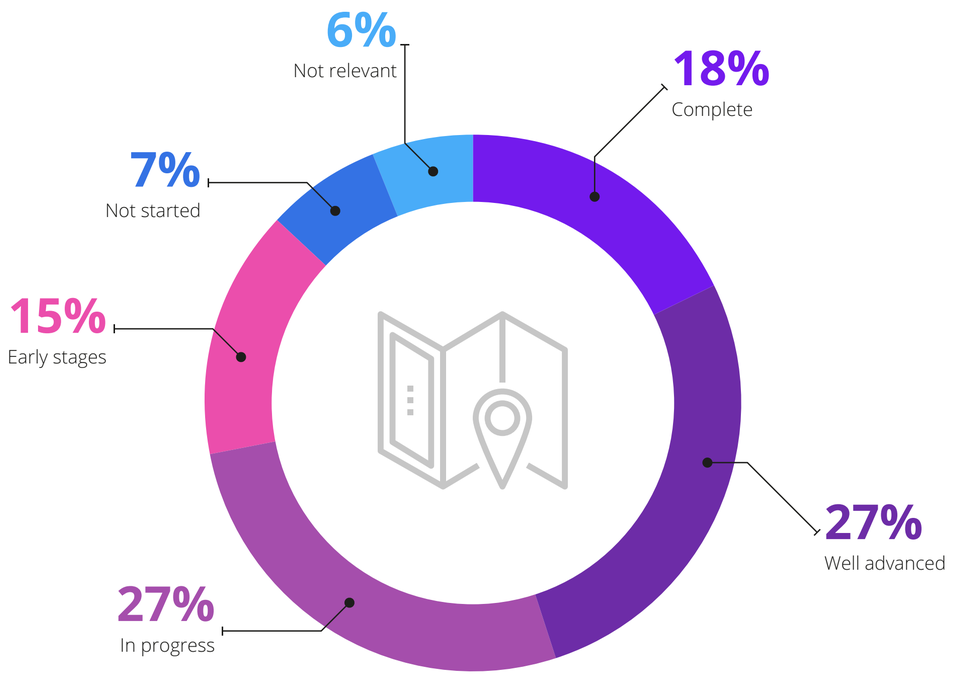

NOTE: N=2,285

SOURCE: OMDIA, ICT ENTERPRISE INSIGHTS 2021: IOT, CLOUD, AI & 5G. FILTERED BY: CORPORATE BANKING, FINANCIAL MARKETS, INSURANCE, MERCHANTS – PAYMENTS, PAYMENT ISSUERS/ACQUIRERS, RETAIL BANKING

Data science runs on open source software but thrives on managed workflows

While a CoE can help nurture a digital culture and better enable the democratization of AI, financial services companies must go further if they are going to deploy AI at scale as a competitive asset. They must make a friend of constant innovation. For enterprise AI practitioners, innovation takes on the form of open source software. Unlike traditional enterprise software practices, which follow a steady development trajectory and stable set of core technologies, data science practitioners rely on an ever expanding and evolving tapestry of open source languages, libraries, frameworks, and tools.

New machine learning ML technologies such as deep reinforcement learning, transfer learning, and general adversarial networks (GANS) emerge within open source projects like TensorFlow and PyTorch. And innovative MLOps capabilities emerge within OSS tools like Kubeflow, Airflow, MLFlow, Argo CD, and Tecton (Feast).

This abundance of innovation encourages AI practitioners to select the best tools for the job at hand – a necessary trait where exploration and experimentation are a necessary part of the development process, particularly as practitioners take on increasingly complex projects.

The challenge for enterprises, therefore, is to enable the free use of open source within a portable, secure, performant, and managed workflow environment where diverse assets can be shared, re-used, and extended with ease. This is the key to preserving innovation while ensuring that data science projects reach production in a more timely manner and at greater scale.

CORE AI OPEN-SOURCE TECHNOLOGIES

python™ R TensorFlow

APACHE Spark™ mlflow™ Kubeflow

SOFTWARE ITSELF RUNS ON OPEN SOURCE

80.6% of the Python language depends upon external open-source projects

Cloud-native development ideals feed data science innovation

For financial services companies seeking to build AI outcomes capable of going beyond the limits of proof-of-concept and serve as a company-wide core competency and competitive advantage, infrastructure means everything. Until very recently, highly regulated companies within the financial services market have had to endure the expense and limitations that come from not just building AI outcomes but also building the datacenter infrastructure upon which those outcomes run. For many highly performant workloads such as fraud detection, this makes sense in terms of fulfilling regulatory, performance, and security requirements.

However, for many ML workloads such as model development and training, and for many use cases such as customer experience management and sales enablement, architectural portability has become crucial. Within the broader IT marketplace, this portability stems from the adoption of cloud-native software orchestrators like Kubernetes which can streamline and even automate the deployment and management of software across disparate containerized platforms.

By investing in software that can run within any containerized environment governed by Kubernetes, enterprise AI practitioners can adopt a hybrid and multi-cloud approach to developing AI outcomes. A containerized, portable hybrid-cloud platform can not only lower AI development infrastructure management costs and bring ML into closer proximity with data, it can also open up new vistas of innovation. Data science teams can more readily work together on the same projects, even across disparate geographies and diverse environments. They can also fine-tune access to the most advantageous underlying hardware, such as AI hardware accelerators from NVIDIA, Google, Intel, et al. The same benefits can even extend to the production-ready ML models themselves. By building ML models within a platform that is environment-agnostic as with Kubernetes-managed containers, AI practitioners can deploy models to premises, edge, and cloud in a consistent manner using DevOps principles of continuous integration and continuous deployment.

Solutions built in this manner can fit effortlessly into existing enterprise environments and can meet customers where they are regardless of their existing levels of investment in AI technologies. And for AI practitioners themselves, working within an open, cloud-native platform, can speed development and boost productivity simply by freeing practitioners to focus not on managing data center technologies but instead on what they do best, building AI outcomes.

What is your delivery strategy for BI/analytics applications now and in 18 months' time?

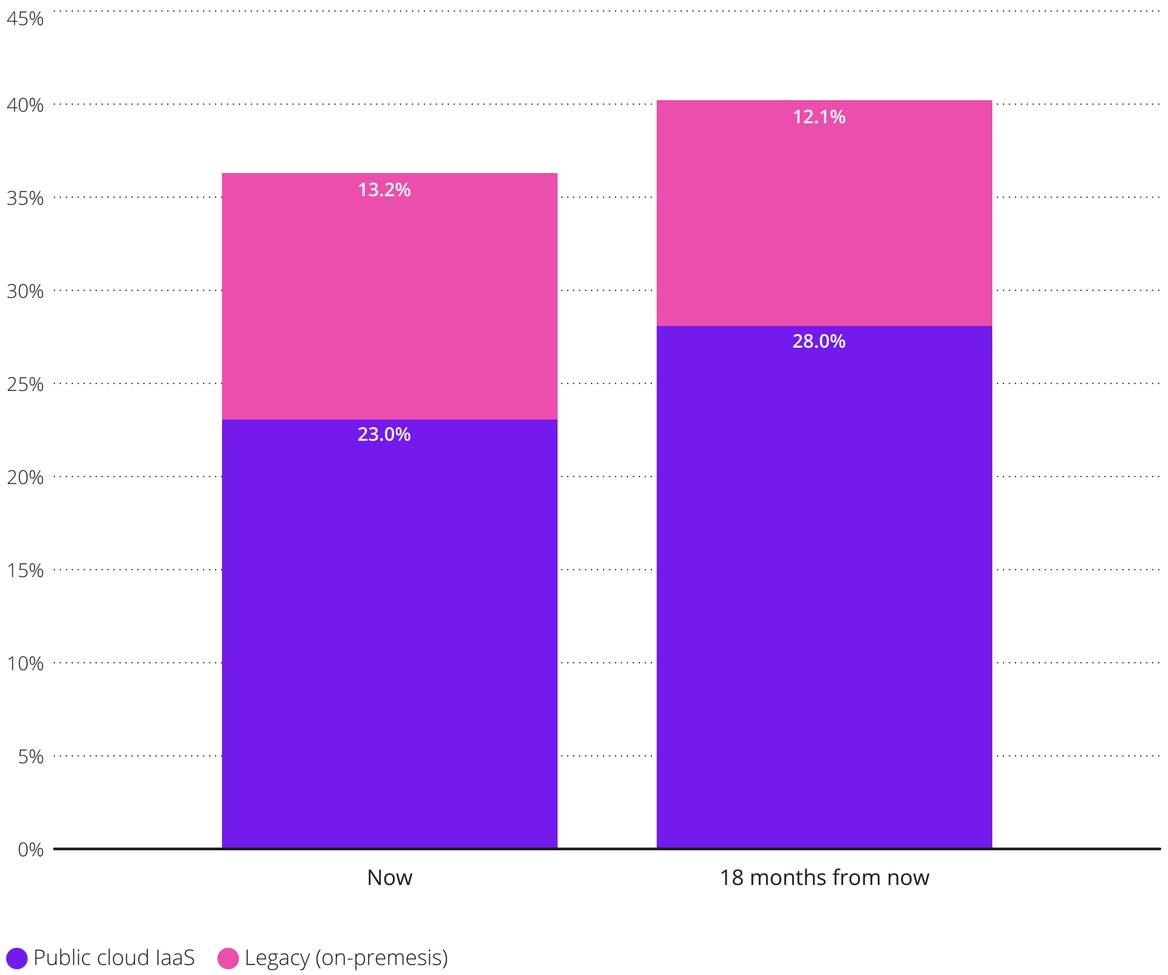

NOTE: N=257

SOURCE: OMDIA, ICT ENTERPRISE INSIGHTS 2021

Cultivating corporate resilience and agility through AI governance

Three keys to leveraging AI outcomes in the enterprise

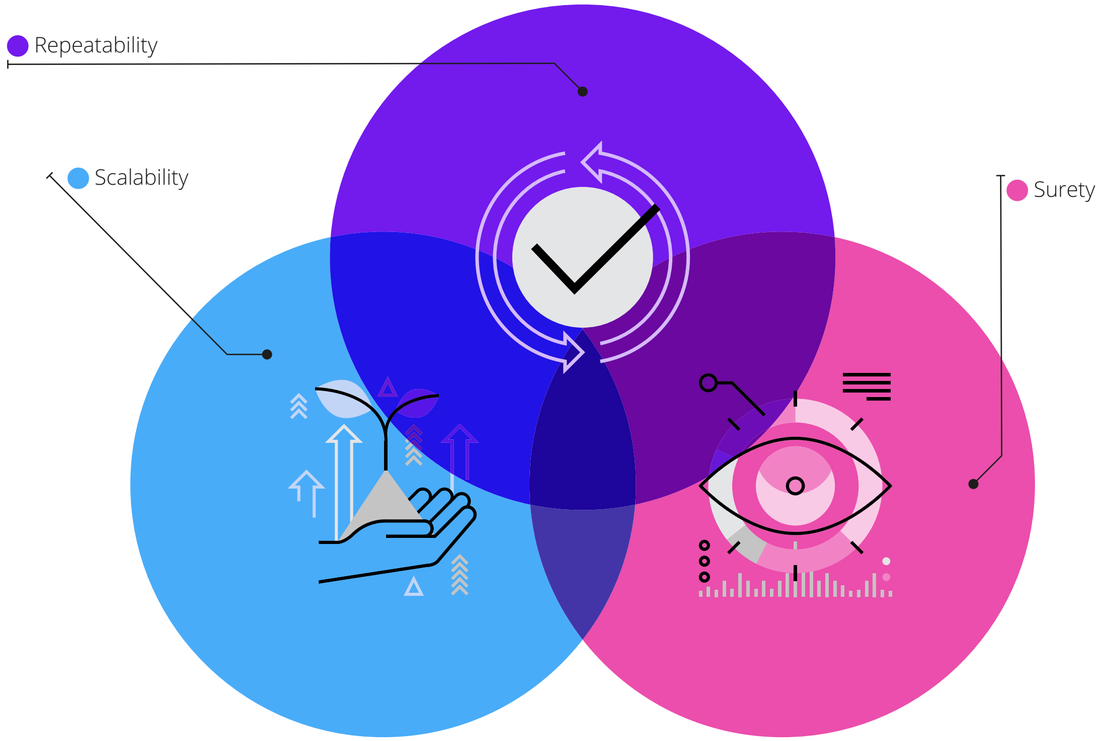

SOURCE: OMDIA

The adoption of a cloud-native and MLOps-savvy AI development platform can pay serious dividends for both inexperienced and experienced enterprises alike. It can help those new to AI rapidly come up to speed by orchestrating what is traditionally a very complex and multifaceted development lifecycle. More importantly, it can help experienced practitioners apply the controls necessary to repeat, scale, and most importantly, trust the use of AI outcomes across the enterprise.

Given these challenges, an investment in an underlying hybrid cloud platform based on containers and Kubernetes capable of addressing these three, intertwined concerns can dramatically upscale the value of AI development efforts. This is of particular importance for companies leveraging open source software across either a self-managed cloud or managed AI cloud service, or both, in combination. Doing so greatly improves access to data and compute resources, while orchestrating the use of interdependent libraries and package versions, all while supporting AI governance, security, transparency, and explainability requirements. With these challenges kept in check, companies can actively build on past efforts, maximize performance (of both practitioner and platform), or prove compliance with regulatory demands. Freed from such concerns, financial services practitioners can more actively and aggressively pursue AI opportunities and adapt to market disruptions.

Repeatability – It is difficult for AI practitioners to repeat what is a highly investigational methodology filled with stops and starts, dead ends, and unforeseen avenues of exploration. Unmanaged code lives ungoverned within various Jupyter notebook implementations, making it nearly impossible for anyone but the original author to track and manage the code over time. Add to this the challenge of maintaining versioning across a wide array of libraries, frameworks and data itself, and it is easy to see how lessons learned in one project do not readily carry over into future endeavors.

Scalability – It can be notoriously difficult for companies to effectively manage resource requirements, such as AI acceleration hardware for both development (training) and deployment (inference) because these tasks are themselves dependent on a myriad of malleable conditions including infrastructure latency, data batch sizes, as well as ML model size and type. The same holds for all supportive storage and processing resources, such as database instances, data pipeline processing, and inference engine execution. The high degree of entanglement makes it difficult for IT managers and CTOs to predict and, therefore, manage costs, a difficulty that grows exponentially as new AI projects enter development and production.

Surety – It can be difficult to trust AI business outcomes altogether because of a lack of transparency within many deep learning (DL) models, unchecked biases lurking in both data and model alike, poor code documentation across the project lifecycle, and inadequate testing of models before deployment. Maintaining a level of confidence over time, therefore, demands a high degree of vigilance, monitoring models to ensure that their efficacy does not diminish because of changes in the supporting data or surrounding systems. For highly regulated industries such as financial services, this kind of monitoring can demand the actual replication of a model's output at a given time. This can be an impossible task for organizations that are unable to fully document the entire ML lifecycle, including data ingestion, data preparation, feature engineering, model selection, parameter tuning, model validation, integration of models with application development process, application deployment, as well as model inferencing model.

Red Hat capabilities & success stories

Hybrid and open-source at its core, Red Hat OpenShift accelerates AI and ML lifecycles

Red Hat's increasingly predominant hybrid cloud platform Red Hat OpenShift, is an open platform built upon containers and Kubernetes. There are many flexible consumption options with Red Hat OpenShift, including a self-managed option across a hybrid cloud environment or as a managed cloud service on Microsoft Azure, AWS, Google Cloud, and IBM Cloud. Moreover, the Red Hat Marketplace extends buying options to certified enterprise software, including from the AI/ML ISV ecosystem, that can run on Red Hat OpenShift-based clouds. Red Hat's broad certified ISV and technology partner ecosystems and participation in open source communities help influence the vendor's product roadmap. Red Hat OpenShift, alongside Red Hat Application Services for both application and data services, provides a hybrid cloud environment in which developers can take more control of their AI and ML workflows and build AI and ML-based applications faster.

The open-source project Open Data Hub is based on Kubeflow and provides developers, data scientists, and data engineers open-source tools for data storage, distributed AI/ML workflows, and a Jupyter Notebook development environment, all atop the Red Hat OpenShift. Red Hat OpenShift Data Science, a managed cloud service (initially available on AWS) that extends upon Open Data Hub, is designed to support the development, training, and testing of ML models, ensuring those models are packaged for export within a container-ready format. Offered as an add-on service to Red Hat OpenShift managed cloud services, the new platform provides a place in which users can sandbox and train their ML models with access to Red Hat and partner solutions from its ISV partner network.

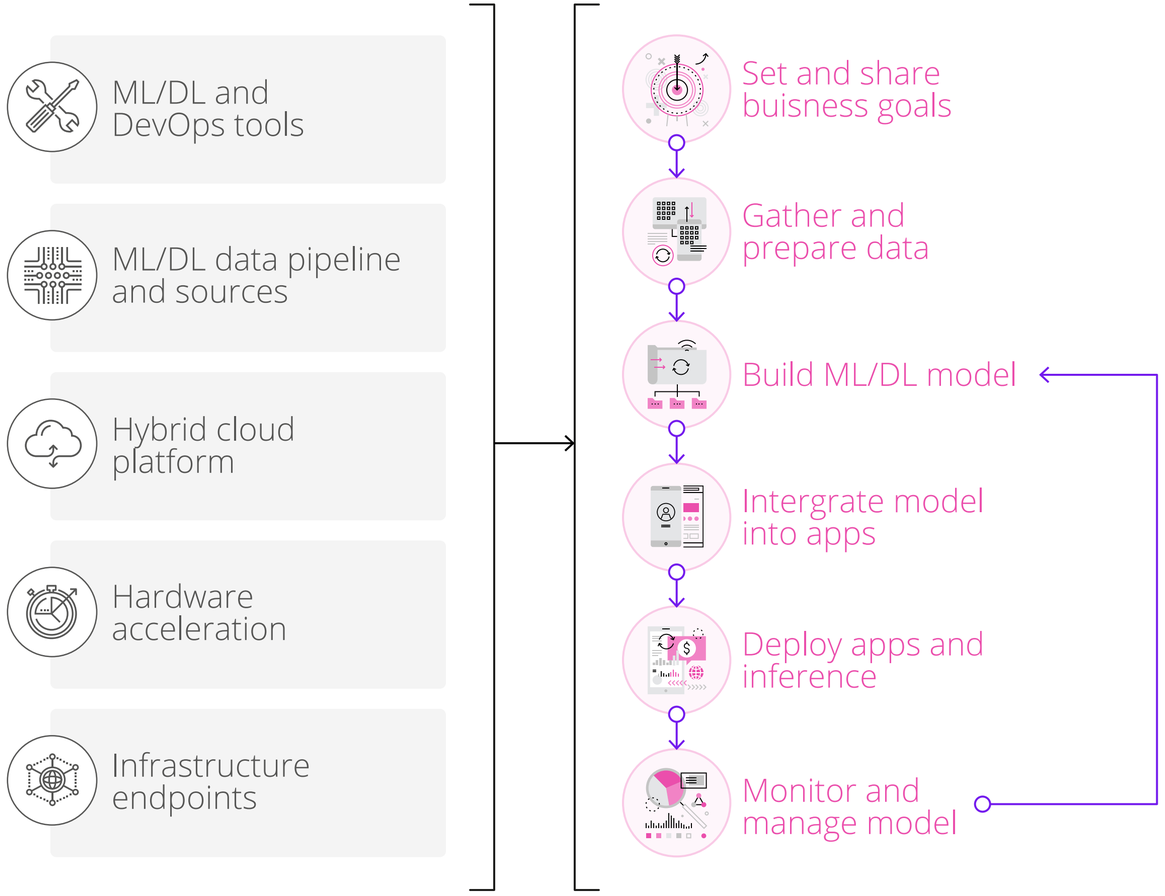

Red Hat is a value-added partner for financial services firms looking to become more agile and resilient to market changes and disruption

Vendors garnering the most attention from financial services firms typically have a strong level of domain expertise. Whether via past successful client engagements or by employing experienced professionals with a background in financial services, specific financial services-related insights have become a key vendor selling point.

Red Hat, known for its open-source prowess and hybrid capabilities particularly through the Red Hat OpenShift enterprise Kubernetes platform, understands that there are added complexities of doing business in the heavily regulated financial services space. Many of Red Hat's solutions, including Red Hat OpenShift, Red Hat Integration, Red Hat Ansible® Automation Platform, Red Hat Enterprise Linux, and more, can be utilized by financial services organizations to ensure compliance and improve resilience and agility, particularly in responding to customer needs and demands to do business more digitally than ever before, two things that are top of mind for so many organizations.

Moreover, Red Hat has industry-specific solutions catering to retail and corporate banking, insurance, and payments organizations and works with key FinTech ISV partners to bring together open source and industry-specific software for best-of-breed solutions for financial services organizations. Being able to speak both data science and the language of the financial services industry makes Red Hat a strategic partner with which financial services organizations have the opportunity to grow their technical and AI prowess to stand out among the competition and generate better and lasting business outcomes.

KEY FSI PARTNERS

FINASTRA, Infosys Finacle, avaloq, Thought Machine, TEMENOS, intellect, Swift

KEY AI PARTNERS

C3.ai, H2O.ai, NVIDIA, IBM, CLOUDERA, Starburst, SELDON

Red Hat OMDIA

The Royal Bank of Canada relies on Red Hat OpenShift to be increasingly agile and generate better business outcomes

The Royal Bank of Canada (RBC) worked with Red Hat and NVIDIA to modernize and improve the banking experience for its customers. Together, they developed and built a private cloud infrastructure for AI and ML on Red Hat OpenShift and NVIDIA DGX AI for the bank to deploy and run AI and intelligent automation applications.

Those at RBC involved in the project underscored the benefits of working with Red Hat, including the ability to build advanced capabilities from which its customers benefit in an in-house environment. The containerized nature of the new cloud environment offers RBC flexibility and reliability and had, already in its early days in July 2020, provided RBC with faster delivery and caliber of applications to its clients, faster customer service, and an improved trading experience.

RBC has over 1,000 AI models operationalized to help the bank improve its capabilities in areas such as risk mitigation and fraud detection. The new Red Hat and NVIDIA cloud set up helps the bank process more data faster and with improved accuracy when compared to its older on-premises environment.

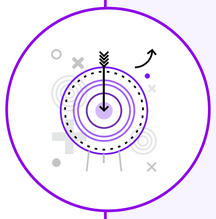

"We're essentially applying better methods to an old problem to achieve a better accuracy," said Foteini Agrafioti, chief science officer at RBC and head of Borealis AI, the bank's research unit. *"An increase of a couple of percentages in one of our businesses is very, very meaningful at the scale that we operate,"* she added.

SOURCE: WALL STREET JOURNAL

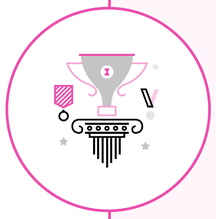

Omdia takeaway: It is pretty easy to build an AI outcome that is 85% what you want, but eking out those last few percentage points, whether those are measured as accuracy, latency, or time to market, can prove to be the hardest won.

AXA France is scaling its utilization of Red Hat OpenShift to further improve its time to market and cost effectiveness

AXA France, part of AXA Group and a well-known insurance firm in France, is making notable strides in its AI/ML efforts to scale its usage of the technology and the associated data science and business outcomes that result.

As an insurance company, AXA France deals with vast amounts of documents each day and much of its document processing was done in highly manual processes, which were also error prone. The company sought to utilize AI and ML to reduce the amount of manual tasks, automatically validate data, reduce back-office process times, and ultimately, improve client satisfaction.

AXA France introduced Red Hat OpenShift and went into production in two months' time and was able to roll it out across a wider swath of users within six months. The company underscored the lack of a DevOps-like model for data science as a challenge that had to be overcome and has since built a new operational model using GitOps and identified three key personas that needed to be involved in the project – ML engineers, platform and automation engineers, and data scientists – in order to operationalize and run AI and ML workloads.

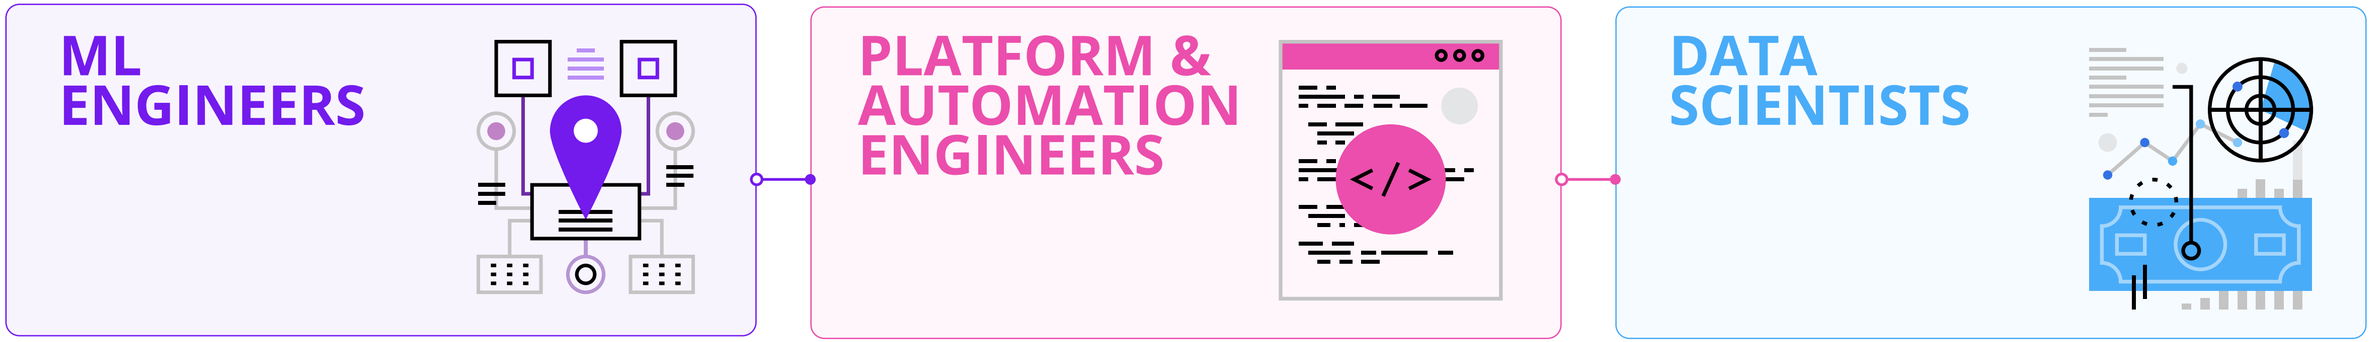

İşbank has accelerated its AI application development and deployment timelines as a result of its utilization of Red Hat solutions

İşbank, known for being Turkey's first national bank and largest private bank with 20 million customers, utilizes Red Hat OpenShift and Red Hat Ceph Storage to enable GPU usage for machine learning, benefitting from Red Hat's increasingly strategic partnership with NVIDIA. Utilizing these solutions enables the bank to accelerate its AI/ML workflows and scale its AI/ML outcomes.

By the end of 2018, the bank chose Red Hat OpenShift as its hybrid cloud platform upon which to run and in early 2021, the company migrated to the latest version of OpenShift. At the time its migration was 15% complete, the bank already reportedly saw benefits such as reducing provisioning time from weeks to seconds.

The bank handles vast amounts of data and transactions each day and as such, needs a strong foundation. Its infrastructure, including Red Hat OpenShift and Red Hat Ceph Storage, provides the basis upon which the bank's applications are developed and run. Speed is of notable importance in this particular case as these applications provide customers with an enhanced banking experience. İşbank now has over 30 AI applications in production spanning areas such as ATM cash optimization, loan underwriting, and more.

"We selected Red Hat OpenShift for our AI applications, including ATM optimization, income estimations, pricing and other AIOps and NLP-based apps. For these applications, we are running artificial intelligence workloads on OpenShift with more than 30 AI/ML apps in production. Red Hat OpenShift helps us manage the MLOps pipeline including AI model development, deployment, and monitoring. We chose Red Hat OpenShift because it is easy to integrate with DevOps and GitOps tools and increases our application development speed so we can deploy faster. We also use Red Hat OpenShift Data Foundation to store our model and the data in a more secure manner."

Çağlar Gülşeni, AI architecture chapter lead, İşbank*

UTILIZING RED HAT OPENSHIFT AND CEPH STORAGE HELPS İŞBANK:

*SOURCE: RED HAT 'RED HAT EXPANDS WORKLOAD POSSIBILITIES ACROSS THE HYBRID CLOUD WITH LATEST VERSION OF OPENSHIFT'

Conclusion: FSI organizations are increasingly realizing the value of AI through the adoption of open, hybrid cloud environments

Use cases

- Financial services use cases for AI and ML span end customer, business user, and IT user benefit
- Risk mitigation, fraud detection, loan processing, and more are all important areas of FSI in which AI and ML can make noteworthy impact

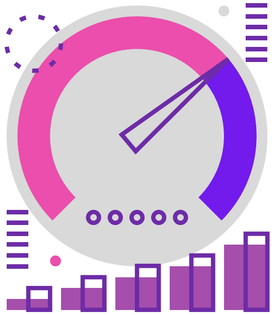

Best practices

- Successful AI and ML outcomes bring together people, process, and technology
- Hybrid cloud-based AI/ML platforms provide speed and flexibility
- Open source communities provide documentation, toolkits, forums, and support

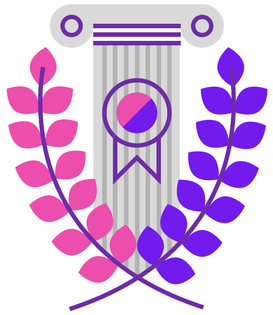

Success stories

- RBC's Red Hat and NVIDIA private cloud helps scale its AI outcomes
- AXA France quickly realized AI/ML efficiencies from its Red Hat OpenShift deployment and is promoting change management to further enable its usage
- İşbank utilizes Red Hat products (including through the NVIDIA partnership) to optimize GPU usage for ML for ML workflow acceleration

Appendix

Red Hat

We're the world's leading provider of enterprise open source solutions, using a community-powered approach to deliver high-performing Linux, cloud, container, and Kubernetes technologies. We help you standardize across environments, develop cloud-native applications, and integrate, automate, secure, and manage complex environments with award-winning support, training, and consulting services.

Red Hat

W redhat.com/en/contact
RedHat
Red-Hat

Omdia

Omdia is a global technology research powerhouse, established following the merger of the research division of Informa Tech (Ovum, Heavy Reading, and Tractica) and the acquired IHS Markit technology research portfolio*.

We combine the expertise of more than 400 analysts across the entire technology spectrum, covering 150 markets. We publish over 3,000 research reports annually, reaching more than 14,000 subscribers, and cover thousands of technology, media, and telecommunications companies.

Our exhaustive intelligence and deep technology expertise enable us to uncover actionable insights that help our customers connect the dots in today's constantly evolving technology environment and empower them to improve their businesses – today and tomorrow.

*The majority of IHS Markit technology research products and solutions were acquired by Informa in August 2019 and are now part of Omdia.

Appendix

Analysts

Bradley Shimmin
Chief Analyst, AI Platforms, Analytics and Data Management

bradley.shimmin@omdia.com

Cassandra Mooshian
Senior Analyst, RPA and Intelligent Automation

cassandra.mooshian@omdia.com

Methodology

This eBook references two Omdia enterprise end-user studies, ICT Enterprise Insights 2020/21 (n = 4,961) and AI Market Maturity (n = 196), published in September 2020 and June 2020, respectively. The charts indicate which study they are from and all are filtered to only include financial services organizations (n = 2,285 and n = 37, respectively). It also references Omdia's AI Software Market Forecast, 2Q21 which published in May 2021.

Further reading from Omdia:

- "Red Hat introduces OpenShift Data Science with a platform-savvy take on machine learning in the enterprise" (May 2021)
- Omdia Universe: Selecting an Enterprise MLOps Platform, 2021 (April 2021)

Further reading from Red Hat:

- Discover how organizations speed AI/ML adoption with Red Hat OpenShift
- OpenShift for AI and ML
- Red Hat OpenShift Data Science
- Red Hat Expands Workload Possibilities Across the Hybrid Cloud with Latest Version of OpenShift

The Omdia team of 400+ analysts and consultants are located across the globe

Americas

Argentina
Brazil
Canada
United States

Asia-Pacific

Australia
China
India
Japan
Malaysia
Singapore
South Korea
Taiwan

Europe, Middle East, Africa

Denmark
France
Germany
Italy
Kenya
Netherlands
South Africa
Spain
Sweden
United Arab Emirates
United Kingdom

Omdia

E insights@omdia.com
E consulting@omdia.com
W omdia.com

OmdiaHQ
Omdia