




BUILDING A SMARTER BANK WITH AI AS A PLATFORM



THE PROMISES AND CHALLENGES OF AI

Artificial intelligence is a once-in-a-generation technology. One of its incredible breakthroughs is the ability to identify patterns hidden in vast amounts of data at levels of precision, speed, and scale that were previously impossible. The ability brings multiple benefits to the financial services industry, from enabling more intelligent trading to expanding credit and services access to underserved people. It can complement traditional calculations for risk to ensure a safer and better managed financial system. It can be used to detect fraud. And it can introduce intelligent assistance and automation to routine tasks, like claims processing or customer support, freeing up humans to handle more complex, high-value, and interactive tasks.

With all the promise that AI offers, according to [Accenture](#), a large percentage of enterprise executives are worried about not digitally transforming fast enough. And they struggle with how to start.

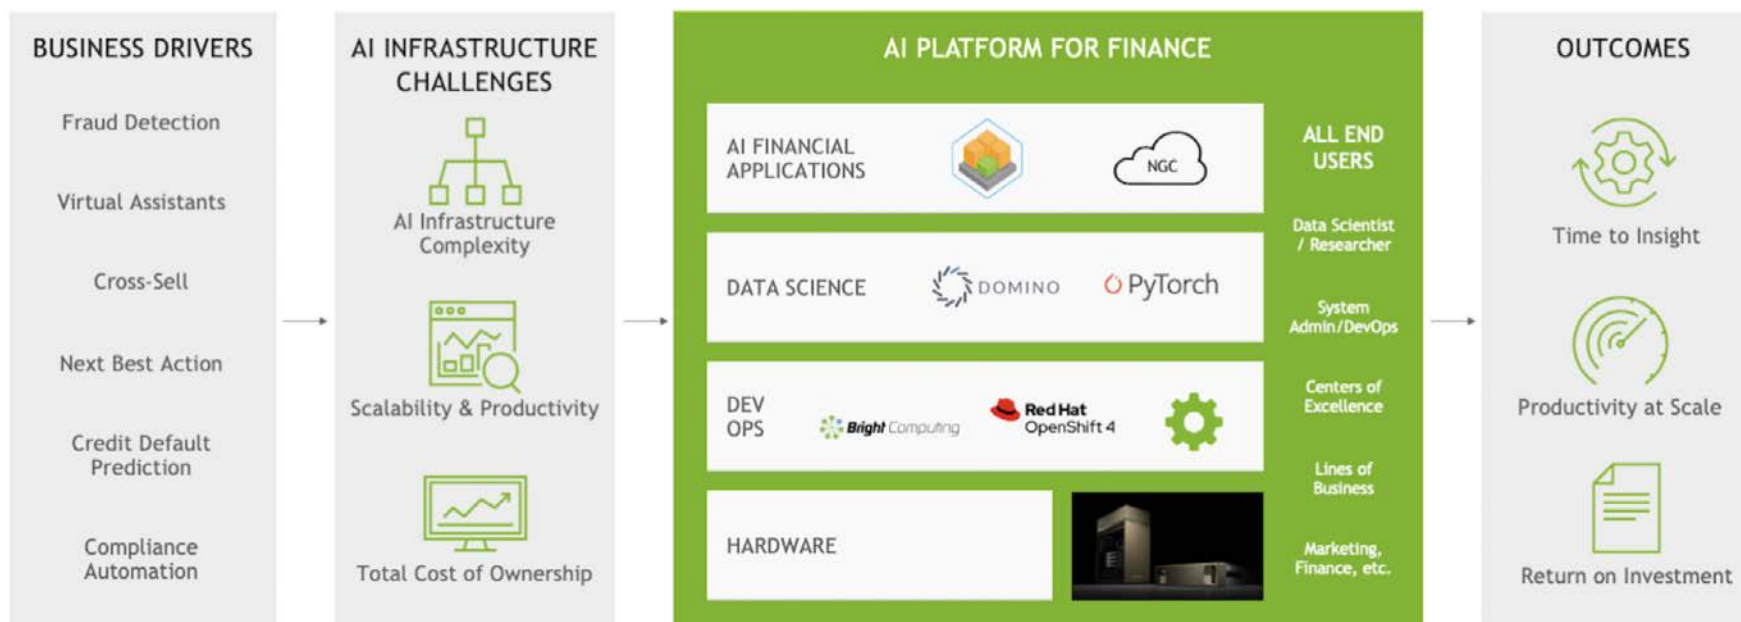


¹ Accenture: "AI: Built to Scale, From Experimental to Exponential." 2019

Enterprise leaders, regardless of whether they need on-premises or hybrid AI infrastructure, find implementing AI challenging for three reasons. First, unless they have long-standing experience in accelerated computing, developing the right data center design for compute, storage, and networking is complex. Second, it can take months to procure, integrate, and troubleshoot all the discrete layers of software and hardware needed to support AI. And third, once an environment has been deployed, supporting all the componentry can be challenging.

On top of the complexities of the AI infrastructure itself, many banks face the challenge of "shadow IT." Lacking a definitive enterprise AI platform, divisions across banks have created their own IT infrastructure and teams of data scientists. The disparate infrastructure estates raise the costs of AI, create a siloed approach to data governance, and ultimately limit the bank's ability to scale AI effectively. As a consequence, many banks fail to fully unlock the value of AI. In contrast, a centralized data and AI platform approach gives data science operations teams the ability to manage and protect data. It also gives management a scalable, cost-effective platform upon which to define and execute a bank's AI strategy.

THE PLATFORM FOR AI-DRIVEN FINANCE



Building a centralized AI platform starts with the key business drivers. For example, banks are focused on improving fraud detection, enabling virtual assistants, and creating recommenders to suggest next best actions, while insurance companies are looking to automate claims processing and identify fraudulent claims. Regardless of the type of financial institution, however, the challenges of AI infrastructure remain: It's complex, it's difficult to make both scalable and productive, and it's challenging to do on a budget.

NVIDIA's AI platform for finance delivers the entire solution set, from the hardware foundation with GPU compute, to DevOps/MLOps and data science software, to industry-specific applications—all NGC™ certified and tested for interoperability and functionality to deliver the experience end users demand. And it all comes together to produce the outcomes enterprises want to achieve through AI transformation: faster time to insight, productivity at scale, and a positive return on investment (ROI).

AI AS A PLATFORM (AIaaSP) POWERS THE SMARTER BANK

There are hundreds, if not thousands, of applications that AI-powered banks will enable in the coming years. Consequently, banks need a scalable infrastructure and strategic platform to digitally transform with AI. The following are the key application areas.


Measuring risk accurately is critical to the financial stability of banks, insurance carriers, and credit card issuers. With the proliferation of data about consumers, such as social media behaviors, telematics driving data, and open-banking insights, analyzing risk is a priority use case for AI and machine learning. The insights gleaned from AI can be used at all points in the consumer lifecycle, from acquisition to servicing, and can help risk analysts quantify the appropriate marketing, pricing, servicing, and retention actions to take. Additionally, explainability tools will prove valuable when financial institutions need to share the rationale for underwriting decisions with regulators.

Conversational AI is enabling consumers to manage all types of financial transactions, from bill payments and money transfers to opening new accounts. Banks encourage these self-service interactions, as it frees customer service agents to focus on higher-value interactions and transactions. At the heart of conversational AI are deep learning models that require significant computing power to train chatbots to communicate in the domain-specific jargon and language of financial services.

Some of the biggest AI wins are those related to fighting transaction fraud—a multi-billion-dollar problem. Detecting true fraud is critical, but traditional systems have historically generated many more false-positive than true-fraud signals. Now, advanced machine learning and deep learning techniques are improving detection and, at the same time, drastically cutting false-positive rates.

On some of the largest commercial platforms, recommendations account for as much as 30 percent of revenue, which can translate into billions of dollars in sales. That's why banks and insurance companies are using recommender systems to drive every offer they make to consumers, from targeted advertisements to personalized products.

All of these applications need to be supported by a full-stack AI platform that delivers scalability, productivity, and a positive ROI.






DELIVERING AI AS A PLATFORM

Scalable, productive, and cost-effective AI platforms deliver performance across three key dimensions: technology, people, and economics.

For technology, it's not enough to deliver great hardware. After all, that hardware is meaningless without great software running on top of it. Whether built on premises or in the cloud, AlaaP must bring the full stack to the enterprise: all the way from the bottom hardware layers up through DevOps and data science operations software layers to the very top application layers. With a full stack operating in harmony, people can do their best work. The IT team has a platform that can scale to meet the needs of the business. Data scientists can remove themselves from DevOps responsibilities and focus on building AI applications that deliver the outcomes necessary for business leaders to achieve their targets. When business outcomes are met and the platform is operating efficiently from a scale-out and productivity perspective, the economics and ROI for the business are meaningful.



GPU | Servers | Storage
Network | Tools | Applications

TECHNOLOGY



IT | DevOps
Business | Scientists

PEOPLE



Cloud vs. on Premises

ECONOMICS

COMMON QUESTIONS, ANSWERED

What do CIOs and IT leaders need to consider in their AI strategy?

When IT and executive leadership drive the AI platform strategy, they shift AI from a siloed mindset to something that's mission critical and can be resourced and centralized as a shared service. This shared infrastructure approach offers benefits to the business and IT in three key areas, namely people, process, and platform. From a people perspective, a business now has the ability to eliminate siloed approaches to innovation and instead bring line-of-business experts together to share best practices. With this centralization of talent, it's possible to solve for a common problem—lack of data science expertise—since IT can now build a talent pipeline that can groom these capabilities from within. Centralization means having the ability to develop and replicate an IT standard that ensures an optimized balance of data center resources across compute, storage, and networking, with the added benefit of making the platform universally accessible. This democratizes and accelerates the use of machine learning and deep learning across the organization.

How does a purpose built AI platform enable scalability for the enterprise?

The platform gives IT a standardized approach for AI infrastructure and simplifies infrastructure planning. It also helps to provide security at every layer and gives the operations team peace of mind. The integrated reference solution also results in linearly predictable performance with scale.

Why are data scientists not as productive as they want to be?

Data scientists are doing very little data science; the time they need to allocate to underlying hardware, software, and web development is too much. But with AlaaP, they can focus on AI rather than the AI plumbing, and businesses get an integrated hardware and software solution that can run on scalable GPU server hardware. The solution can manage areas that take time away from actual data science, such as orchestration, job scheduling, version control, and deployment.



NVIDIA SOLUTIONS AND PERFORMANCE

AlaaP is the implementation of a shared, centralized infrastructure for AI that consolidates expertise, speeds and scales the lifecycle from development to deployment, and drives down total cost of ownership with more efficient utilization of compute and storage resources. Forward-leaning enterprises have already been doing this, experiencing the positive impact of successful AI implementation. These enterprises also receive an additional benefit that sets them apart—they become the kind of organization that attracts the world's best talent. The people who lead AI innovation come to companies that offer these tools and scale, enabling them to do their life's most important work. In every industry, including financial services, this delivers a profound competitive advantage.

