


Build an effective AI strategy: Overcome four common adoption challenges

Learn how to implement AI across your organization, speed up time to value, and innovate faster. See how Google Cloud customers have driven impactful business results with AI.



Introduction

Improving decision-making, customer experience, and innovation requires faster data insights. Artificial intelligence (AI) and machine learning (ML) increase the value of data by rapidly identifying patterns that are either impossible or too time-consuming to see with manual analysis. By improving data usage with AI and ML, you can instantly give customers, employees, and applications the insights they need.

Whether your organization wants to predict trends or provide customers with personalized services, it's easier to meet these and other requirements with AI and ML. This is why organizations across industries increasingly employ these technologies to improve competitiveness:

By 2026, trends show that 85 percent of enterprises will combine human expertise with AI, ML, natural language processing, and pattern recognition to augment foresight, **making workers 25 percent more productive and effective.**¹



Worldwide revenues for the AI market—including software, hardware, and services—are expected to grow by **19.6 percent in 2022 to \$432.8 billion, and by 2023, the AI market is expected to break the \$500 billion mark.**²



“AI plus human ingenuity is the differentiator for enterprises to scale and thrive in the era of compressed digital transformation.”³

Ritu Jyoti, Group Vice President, Worldwide Artificial Intelligence and Automation Research, International Data Corporation

1. IDC, [Worldwide Artificial Intelligence and Automation 2022 Predictions](#), Doc #US48298421, October 2021.
2. IDC, [IDC Forecasts Companies to Increase Spend on AI Solutions by 19.6% in 2022](#), February 2022.
3. Ibid.

Build an **informed** AI strategy

Despite the surge in AI development, most initiatives currently stay in development. IDC “Survey results show that while AI/ML initiatives are steadily gaining traction with 31% of respondents saying they now have AI in production, most enterprises are still in an experimentation, evaluation/test, or prototyping phase. Of the 31% with AI in production, only one third claim to have reached a mature state of adoption wherein the entire organization benefits from an enterprise-wide AI strategy.”⁴

So why do so many companies that have otherwise efficient software practices, plenty of data, and skilled data science teams struggle to make AI work for them? Four core challenges across people, processes, technologies, and data are most often the culprit:

→ **Challenge 1:**

There’s an ML skills gap.

→ **Challenge 2:**

AI/ML is a multidisciplinary process. Variances in skills, tools, data, and processes often create obstacles.

→ **Challenge 3:**

Slow progress in deploying and maintaining models often undermines momentum, confidence, and support.

→ **Challenge 4:**

Steep compute and data requirements can create infeasible resource demands.

With years of AI development experience, Google has learned how to overcome these challenges. This e-book reviews how some of our customers have successfully implemented AI and ML to drive business transformation by following four **Paradigms**:

- Make AI accessible for more employees by meeting them where they are in terms of skills.
- Adopt a unified data and AI strategy rather than thinking about them as separate entities.
- Standardize the AI development lifecycle while maintaining flexible technology choices.
- Simplify model deployment and infrastructure management.

Challenge 1:

There's an ML skills gap.

Today, the current demand for AI and ML experts exceeds availability. Gartner® reports that, “only 10 percent of organizations have 50 percent or more of their software engineers trained in ML skills.”⁵ When user proficiency in ML skills becomes the gating criterion for ML-project contribution, companies place artificial limits on their own innovation.

Solution: Make AI accessible to more employees.

Rethink the all-or-nothing approaches and remove obstacles for AI and ML adoption by:

- ✓ Equipping your business users and developers with out-of-the-box ML tools so they can easily apply the best of AI to common business challenges.
- ✓ Giving your developers, analysts, and data scientists user-friendly tools for advanced custom modeling and tuning.

“Technical expertise should not be a barrier to implementing AI. Otherwise, use cases where AI can help will languish without modernization, and enterprises without well-developed AI practices will risk falling behind their competitors.”⁶

Andrew Moore, Vice President and General Manager,
Cloud AI and Industry Solutions, Google

5. Gartner Research, [Survey Analysis: AI Adoption Spans Software Engineering and Organizational Boundaries](#), Van Baker and Benoit Lheureux, November 2021.

6. Google, [Making AI more accessible for every business](#), July 2022.

How Google Cloud helps with this challenge.

You can choose from a comprehensive portfolio of AI and ML capabilities that meet users where they are. Employees with limited or no AI and ML skills can easily access powerful pretrained AI/ML capabilities through [AI agents](#) or developer-friendly APIs. Other options simplify and accelerate model training by unlocking ML for data analysts and meeting the specialized needs of advanced data scientists.

[Vertex AI](#) gives all users one interface to access [Google Cloud AI and ML technologies](#), including:



Pretrained models enable developers to directly apply Google's state-of-the-art AI to quickly solve real-world problems without having to build AI models themselves.




[BigQuery ML](#) gives data analysts a familiar SQL interface to develop ML models.



[AutoML](#) provides automated machine learning capabilities so that everyone, from novices to experts, can quickly create custom models with little or no coding.




[Tabular Workflows](#) enable greater control and customization in [AutoML](#) by giving users the ability to interpret and control each step within a workflow.



[Vertex AI custom training](#) gives data scientists and ML experts the option to use open source frameworks along with hardware services and advanced tooling in [Google Cloud](#).

Customers who have made **AI accessible** to more employees using Google Cloud include:

- [InteractiveTel](#)
- [Grupo Folha](#)
- [Twitter](#)
- [YOOZOO Games](#)



Customer: InteractiveTel

Industry: Technology

Location: United States

Google Cloud:

[Speech-To-Text](#)

[Cloud Storage](#)

Telephony service provider **improves transcription accuracy by 30 percent with AI.**

Challenge

Customers who are ready to buy a car often call automobile dealers for information about their inventories. When InteractiveTel learned that dealers don't respond to customers' calls 75 percent of the time, the telephony service provider developed a SaaS application for automobile dealers that tracks, monitors, and reports on phone-based customer interactions. However, the application's transcription system was unable to deliver the required levels of accuracy and scalability, creating customer frustration and business risk.

Solution

InteractiveTel replaced its existing transcription system with Speech-to-Text API. It integrates seamlessly with its application and automatically converts speech to text in near real-time.

The application now stores call recordings in Google Cloud Storage.

Value

InteractiveTel increased transcription accuracy by up to 30 percent, improving customer experience.

The transcription system now automatically scales to support the company's expanded sales of its SaaS application.

For more information about InteractiveTel's use of Google Cloud, please [read this blog](#).

“Google Cloud Speech API [Speech-to-Text API] performs highly accurate speech-to-text transcription in near real-time. The higher accuracy rates mean we can help dealers get the most out of phone interactions with their customers and increase sales.”

Gary Graves, Chief Technology Officer and Co-Founder, InteractiveTel



Customer: Grupo Folha

Industry: Media & Entertainment

Location: Brazil

Google Cloud:

[Cloud SQL](#)

[Pub/Sub](#)

[Vision AI](#)

One of Brazil's largest newspapers celebrates its 100th anniversary by **indexing 2.5 million historic photos in the cloud.**

Challenge

Folha de S.Paulo has a rich archive of materials, ranging from everyday stories to events that disrupted Brazil and the world. Its photographic archive alone has more than 27 million photographs, representing every event reported in the company's 100-year history.

To protect all its data, the newspaper began digitizing every file but soon found challenges with file indexing and searching. Not only did photographs have handwritten notes on the back but subject matter data also had to be manually entered. And given the massive volume of files, even if the team could process 6,000 images per day, it would take 35 years to digitize the entire archive.

Solution

To create a plan for managing digital assets in the cloud, Folha de S.Paulo partnered with the Google News Initiative (GNI), a global program that helps foster innovation and digital sustainability in news journalism.

Working with GNI, Folha de S.Paulo also engaged third-party Assetway to deploy the Assetway Media Center, a digital media platform based on Google Cloud technologies.

The team migrated digital images to Media Center, which uses Cloud Storage to store images and Pub/Sub to process them so that they are searchable.

Vision AI automatically detects both text and objects in images, and it creates metadata to improve file indexing and search. Vision AI can even recognize handwritten text on the back of the pictures using optical character recognition (OCR).

Value

With its new solution, Folha de S.Paulo processes more than 200,000 images per day.

Staff can now find images that were previously lost in 100,000 folders in just seconds.

Fast and easy access to the company's massive archive enhances stories and improves team productivity.

For more information about Grupo Folha's use of Google Cloud, please [read this case study](#).

“We started talking with Google when our spirits were really low and our projections had been scrapped. **Discovering it could be done and that all the material would finally be indexed and available** to the editorial staff was very encouraging.”

Juliana Laurino, Administrative Manager of the Editorial Staff, General Manager of Folhapress and Grupo Folha's Database



Customer: Twitter

Industry: Technology

Location: United States

Google Cloud:

[AutoML](#)

[Vertex AI](#)

Social network surfaces relevant voice conversations to millions of users in near real-time.

Challenge

Millions of people create and join live audio conversations on Twitter using Spaces. To alert every Twitter user about conversations that align with their interests, the organization needed to add an ML component to Spaces. However, Twitter lacked dedicated ML resources that could build, deploy, and manage this feature.

Solution

Engineers with minimal ML expertise use AutoML and AutoML Tables to train and deploy models running on a multicloud infrastructure.

Value

Global Twitter users receive immediate alerts about live conversations that might interest them.

Twitter maintains its solution without a dedicated team of ML experts.

The solution performs 900 queries per second on the Spaces tab, and it evaluates 50,000 users per second to determine their interest level. Ninety-nine percent of these requests are faster than 100 milliseconds, and 90 percent of requests are faster than 50 milliseconds.

For more information about Twitter's use of Google Cloud, please [read this blog](#).

“After deploying our AutoML Tables solution we saw an increase of 1.96% in Spaces daily active customers, which is one of our key metrics. We also noticed an increase of 1.99% in Spaces join-in rates, and an increase of 8.42% in user clicks to explore a Space.”

Diem Nguyen, Senior Machine Learning Engineer and Data Scientist, Twitter

YOOZOO GAMES

Customer: YOOZOO Games

Industry: Media & Entertainment, Gaming

Location: China

Google Cloud:

[AutoML](#)

[BigQuery](#)

[Cloud SQL](#)

Gaming company increases ad clickthrough rate by 300 percent through ML-powered precision targeting.

Challenge

YOOZOO Games wanted to further improve the profitability of its advertising investments by leveraging the latest ML technology.

Solution

The company partnered with Google gTech gPS, Google Large Customer Sales (LCS), and the Google Cloud team to develop a solution called MOCHA, which delivers end-to-end machine learning pipelines for advertising.

With assistance from Google, YOOZOO Games started uploading its user behavior data to BigQuery and employing AutoML to predict users' inclination to make in-game purchases.

Value

YOOZOO Games can identify potential and whale payers from newly installed users, which guides Google Ads for accurate user targeting and optimized user acquisition.

The company increased ad clickthrough rate by 300 percent.

For more information about YOOZOO Games's use of Google Cloud, please [read this case study](#).

“Google Cloud has helped us greatly in improving gaming experiences...

Seeing the results so far, we're confident that we can offer even more exciting game features to players around the world in the future.”

Songtao Lu, Head of Operations, YOOZOO Games

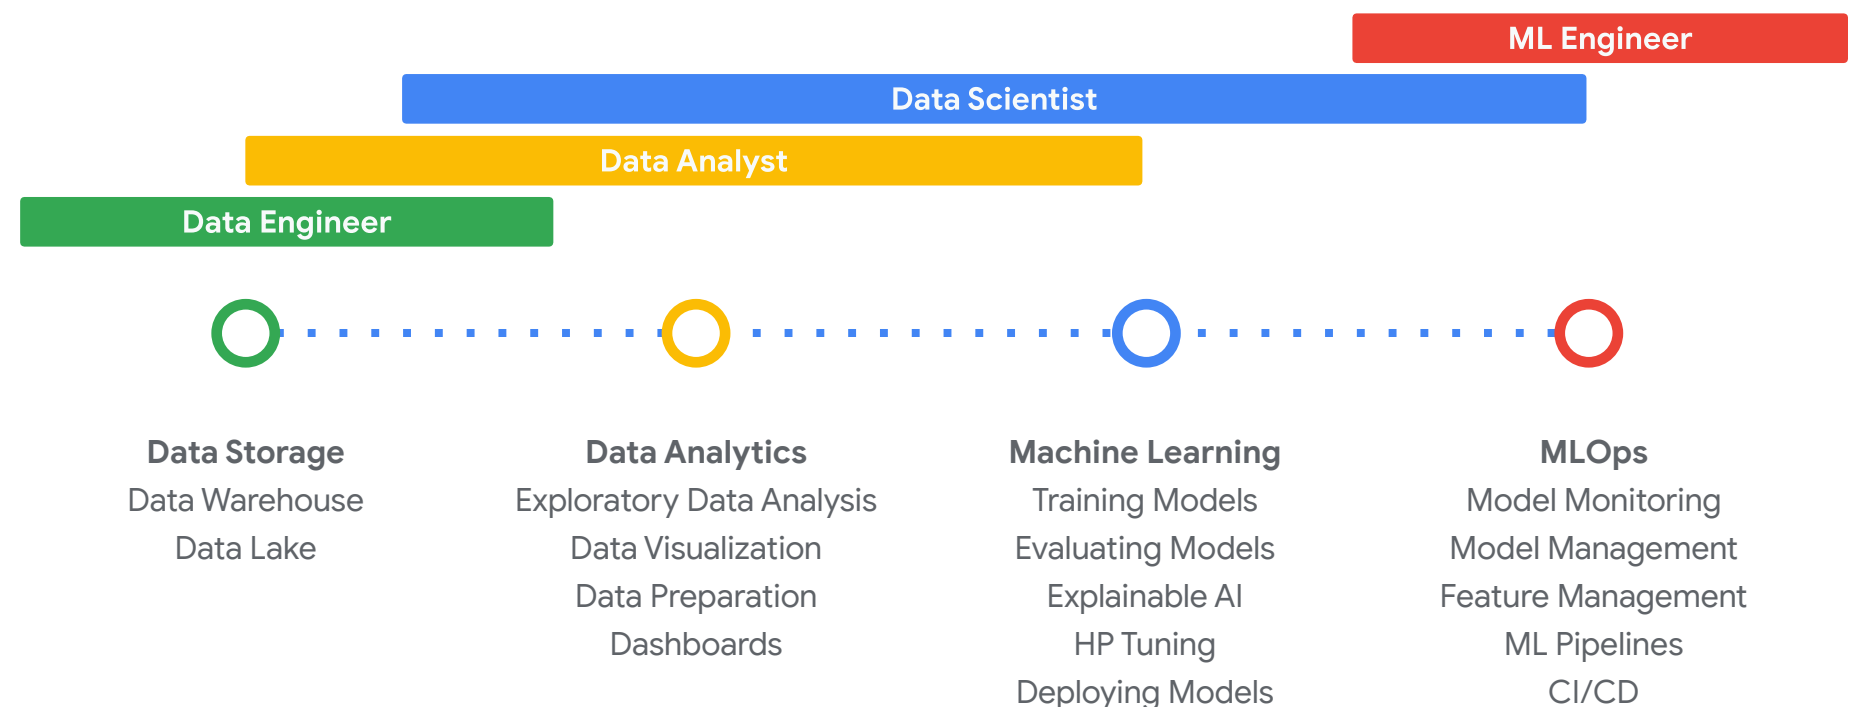
Challenge 2:

AI/ML is a **multidisciplinary process**. Variances in skills, tools, data, and processes can create obstacles.

Traditionally, machine learning starts with data analysis, crosses over into statistics and engineering, and then goes back again to data analysis. Each transition can involve changes in team members, priorities, and approaches, which can slow or entirely stop development momentum. Disparity in teams' technologies as well as sprawled data silos also create complexity, drive up costs, and hinder AI and ML development.

Solution: Adopt a unified data and AI strategy.

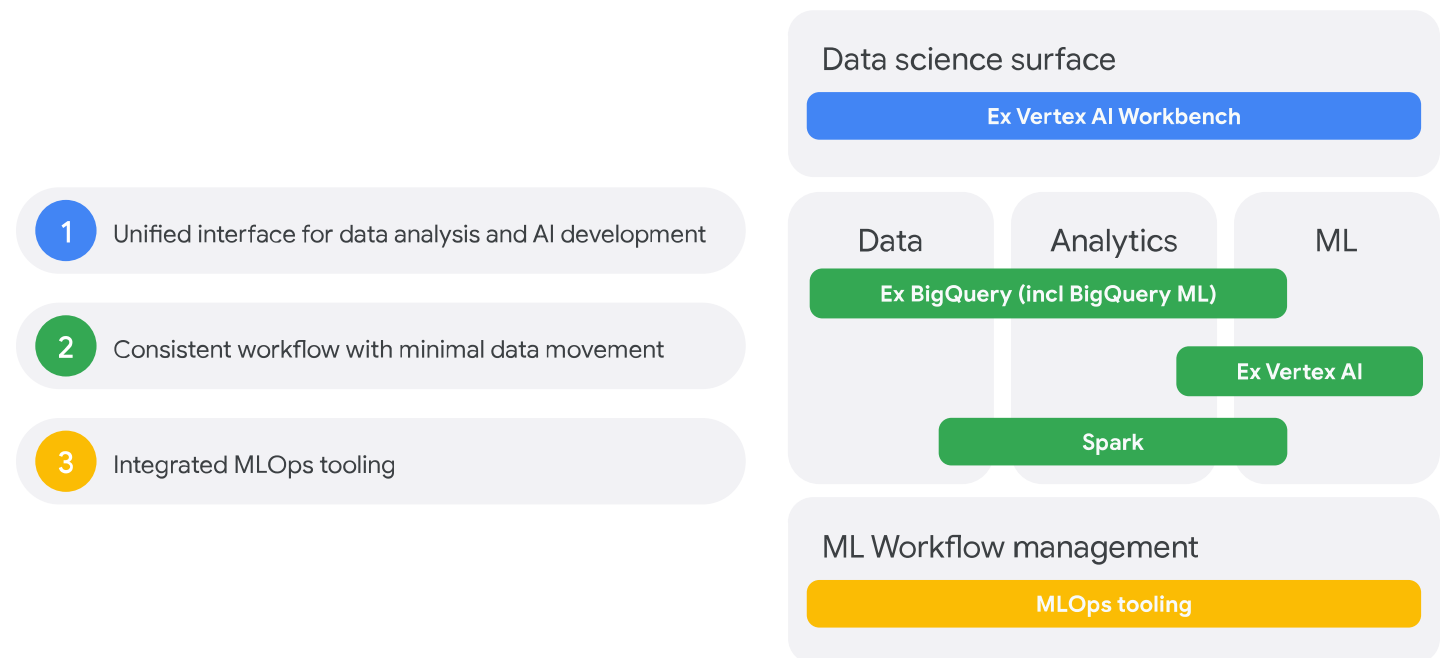
Manage data and AI as one system rather than thinking about them as separate entities. At the same time, protect freedom of choice so individuals can mix and match components and processes as they see fit.



By providing one holistic platform for AI and ML development, organizations can eliminate much of the complexity and inertia that comes from disparate tools, processes, and strategies. That's because data engineers, data analysts, data scientists, business analysts, and ML engineers use a single data science surface to:

- ✓ Minimize context switching.
- ✓ Access the same real-time data, ensuring freshness, reducing silos, mitigating data movement, minimizing administrative requirements, and lowering costs.
- ✓ Simplify the development of AI and ML models, features, applications, and visualizations.
- ✓ Share and reuse artifacts and workflows.
- ✓ Integrate MLOps tools and workflows, which reduces friction between data scientists and IT teams.
- ✓ Give individuals and teams choice and flexibility in ML frameworks, deployment instances, and compute processors.

The journey from data to AI



“Compared to companies that approach AI in a piecemeal manner, these leaders (just over a third in our survey) are far more likely (36% versus 20%) to report widespread AI adoption. They’re roughly twice as likely to report substantial value from AI initiatives to improve productivity, decision-making, customer experience, product and service innovation, employee experience, and more.”⁷

How Google Cloud helps with this challenge.

[Vertex AI](#) brings data and ML systems together in a single interface so teams have a common toolset across data analytics, data science, and ML, including:

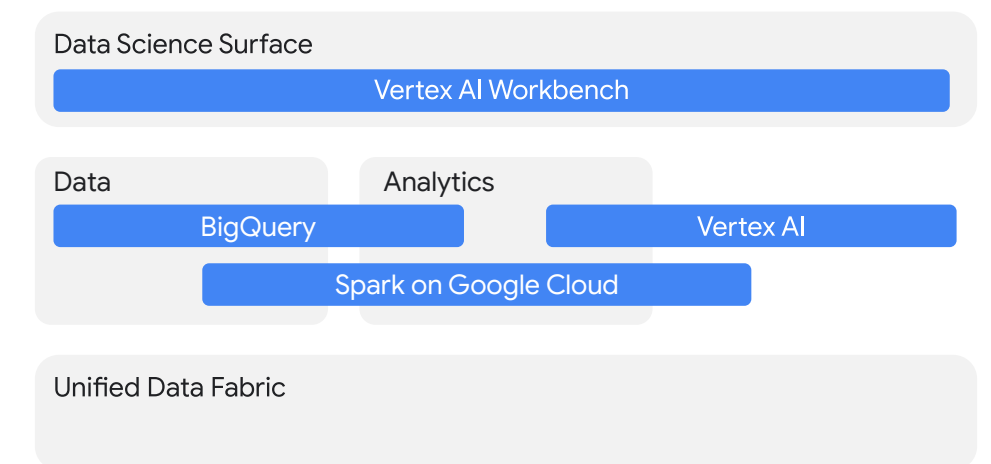
- [Vertex AI Workbench notebooks](#) for exploratory data analysis and model development.
- Serverless AI compute that minimizes data movement.
- End-to-end MLOps tooling for AI deployments.

[Vertex AI](#) is tightly integrated across the [Google's data cloud](#), enabling common governance, MLOps, and AI and ML capabilities directly within other [Google Cloud](#) products:

- [BigQuery ML](#) provides a friendly SQL interface for ML without requiring users to move data or build data pipelines.
- [Spanner](#) and [AlloyDB for PostgreSQL](#) can directly consume [Vertex AI](#) models, greatly simplifying application build processes.
- [Dataplex](#) delivers an intelligent data fabric that enables centralized management across data lakes, data warehouses, and data marts.
- Centralized MLOps capabilities such as prebuilt pipeline components are available across data sources including [Dataflow](#) and [Dataproc](#).
 - [Dataflow](#) enables serverless data processing via a fully managed service.
 - [Dataproc](#) provides fully managed services for running more than 30 open source tools and frameworks.

Vertex AI notebooks

With native integrations across Google [BigQuery](#), [Dataproc](#), and [Dataplex](#), data scientists can build, train, and deploy ML models five times faster with [Vertex AI notebooks](#) than they can with traditional notebooks.



Customers who have adopted a **unified data and AI strategy** with Google Cloud include:

- Swiss International Air Lines (SWISS)
- TELUS
- Coca-Cola Bottlers Japan
- abastece-ai





Customer: Swiss International Air Lines (SWISS)

Industry: Travel & Hospitality

Location: Switzerland

Google Cloud:

[BigQuery](#)

[AI and ML solutions](#)

Airline makes travel more sustainable and efficient with data-driven optimizations.

Challenge

Efficient global air travel hinges on the coordination of multiple stakeholders and processes, including airports, ground handlers, air-traffic controllers, travel agents, and airlines.

To make global air travel even more seamless for its customers and optimize operations, SWISS, a part of the Lufthansa Group, set out to create a joint data repository containing crew, passenger, rotation, and technical information.

Solution

The team created the Operations Decision Support Suite (OPSD), a modular data platform built on Google Cloud that considers all available information on conditions and resources to drive optimal decisions across the air travel ecosystem.

To better manage scenarios and optimize processes across different operational dimensions, the airline replicates all its operational data in BigQuery.

SWISS uses Google Cloud AI and ML solutions to analyze the centralized data in BigQuery to find ways to optimize processes across different operational scenarios and dimensions.

Value

SWISS has optimized about 50 percent of its network flights, elevating the passenger experience.

Operations controllers now have automated decision support, which simplifies decision-making and frees up time for tasks that add value.

The airline decreases carbon footprint and reduces fuel consumption.

For more information about how SWISS uses Google Cloud, please [read this case study](#).

“We can increase efficiency by replacing a larger plane with a smaller one if possible, or enable more bookings on one flight by switching to a larger plane. Those are some of several use cases making us more fuel-efficient, sustainable, and profitable, powered by Google Cloud and OPSD.”

Christian Most, Head of Project Operations Decision Support Suite, Lufthansa Group

Customer: TELUS

Industry: Telecommunications

Location: Canada

Google Cloud:

[BigQuery](#)

[Cloud Bigtable](#)

[Dataflow](#)

[Vertex AI](#)

Telecommunications provider **improves services and business outcomes** by adopting an insights-driven corporate culture.

Challenge

To deliver the insights needed to improve business decisions and customer experiences, TELUS needed to break down data silos, democratize data access across the organization, and speed up the development and deployment of ML models.

Solution

Data scientists cleaned up and migrated siloed data to a common data layer built with BigQuery, Dataflow, Cloud Composer, Cloud Bigtable, and Cloud Storage.

One metadata repository for all assets ensures one source of truth and governance.

Vertex AI unifies data and AI lifecycles including data exploration, aggregation, and cleaning; model building, training, and testing; and ML model deployment.

Cloud Data Loss Prevention protects data without slowing down analysis.

Value

Data scientists develop models faster and now require just a few days—rather than a few months—to test ideas and determine whether a model will be successful.

The organization increases customer understanding and, as a result, improves their experiences.

Teams increase control over data and protect customer privacy.

For more information about TELUS use of Google Cloud, please [read this blog](#).

“In addition to the acceleration of ML model development and experimentation, **with Vertex AI our data scientists will also be able to implement machine learning operations** to efficiently build and manage ML projects throughout the development lifecycle.”

Alexandre Guilbault, Director of Artificial Intelligence and Advanced Analytics, TELUS

Customer: Coca-Cola Bottlers Japan

Industry: Consumer Packaged Goods

Location: Japan

Google Cloud:

[AutoML](#)

[BigQuery](#)

[Vertex AI](#)

Using AI, Coca-Cola Bottlers Japan (CCBJ) processes billions of data records from 700,000 vending machines.

Challenge

Coca-Cola Bottlers Japan (CCBJ) wanted to use AI/ML to help process large volumes of data and make strategic decisions about when and where to locate products in 7,000 vending machines across Japan.

Solution

Using BigQuery and Vertex AI, CCBJ created a new data analytics platform that uses ML to improve efficiency by recommending where to place vending machines, suggesting what products to place in the machines and at what price, and predicting how much the company will sell.

The solution also surfaces insights about individual vending machines on a virtual map to help with regional understanding.

Value

Salespeople easily access predictions and insights via their tablets. And by using the placement point suggestions from the company's ML solution, the routing efficiency of salespeople has dramatically improved. In less than a month, CCBJ built and launched its new platform, including prediction model training, onsite proof of concept, and rollout.

With near-real-time data analysis through Google Cloud, CCBJ teams can focus on developing strategies rather than submitting and waiting for data requests from the IT department.

For more information about CCBJ's use of Google Cloud, please [read this blog](#).

“Analyzing data gives us new discoveries and by using catchy storytelling techniques from exploratory data analysis, we are instilling a data culture within our company. **It's worth creating by looking at facts rather than making assumptions!**”

Minori Matsuda, Google Developer Expert and Data Science Manager, CCBJ



Customer: abastece-ai

Industry: Technology

Location: Brazil

Google Cloud:

[BigQuery](#)

[Dataproc](#)

[Vertex AI](#)

Fintech company unlocks the power of ML and **boosts data-based decision-making capabilities.**

Challenge

With more than 36 million registered customers, abastece-ai faced growth challenges. Because the company lacked a structured data environment, employees could not easily cross-reference data from various sources, hindering data analysis and operations.

Solution

In less than five months, abastece-ai migrated to Google Cloud and consolidated data in BigQuery. The data lake currently processes more than 250TB of data every month.

Because Vertex AI and its support for Python makes it easier to use ML and different modeling techniques, abastece-ai uses Vertex AI to test new products and accelerate operational processes.

Value

By improving data handling and analytics, abastece-ai significantly increased the performance of analytical, strategic, and operational deliveries as well as simplified expansion.

Automated dashboards provide segmented customer views and other analytics that meet the different needs of each business area.

For more information about abastece-ai's use of Google Cloud, please [read this case study](#).

“The environment itself drove a culture change. We are now oriented toward data-based decision-making. **We can respond by making improvements and anticipate sensitive changes.**”

Erick Tambori, Head of Data Analytics, abastece-ai

Challenge 3:

Slow progress in deploying and maintaining models often **undermines momentum, confidence, and support.**

According to Gartner, “in the 2019 Gartner AI in Organizations Survey, respondents indicated that, on average, 53% of projects make it to production.”⁸ For those companies that succeed, model deployment times might take a few weeks or even a year or more after initial proofs of concept. That’s because AI and ML model experimentation, formalization, deployment, and serving all have unique requirements not easily met by traditional software engineering and analytics practices, infrastructure, and mindsets. Additionally, once companies launch a model, they often change its underlying strategies and supporting technologies.

Solution: Standardize the AI development lifecycle while maintaining flexible technology choices.

After adopting a culture that embraces AI and ML experimentation, help your teams focus on what they do best by:

- ✓ Automating processes where possible.
- ✓ Maintaining options for data sources, frameworks, and hardware.
- ✓ Providing central repositories for storing, managing, and sharing models, data pipelines, templates, and other artifacts.

“More than half of data scientists’ efforts go into maintaining models after they have launched.”⁹

Andrew Moore, Vice President and General Manager,
Cloud AI and Industry Solutions, Google

8. Gartner Research, [4 Machine Learning Best Practices to Achieve Project Success](#), Afraz Jaffri, Carlie Idoine, and Erick Brethenoux, December 2021

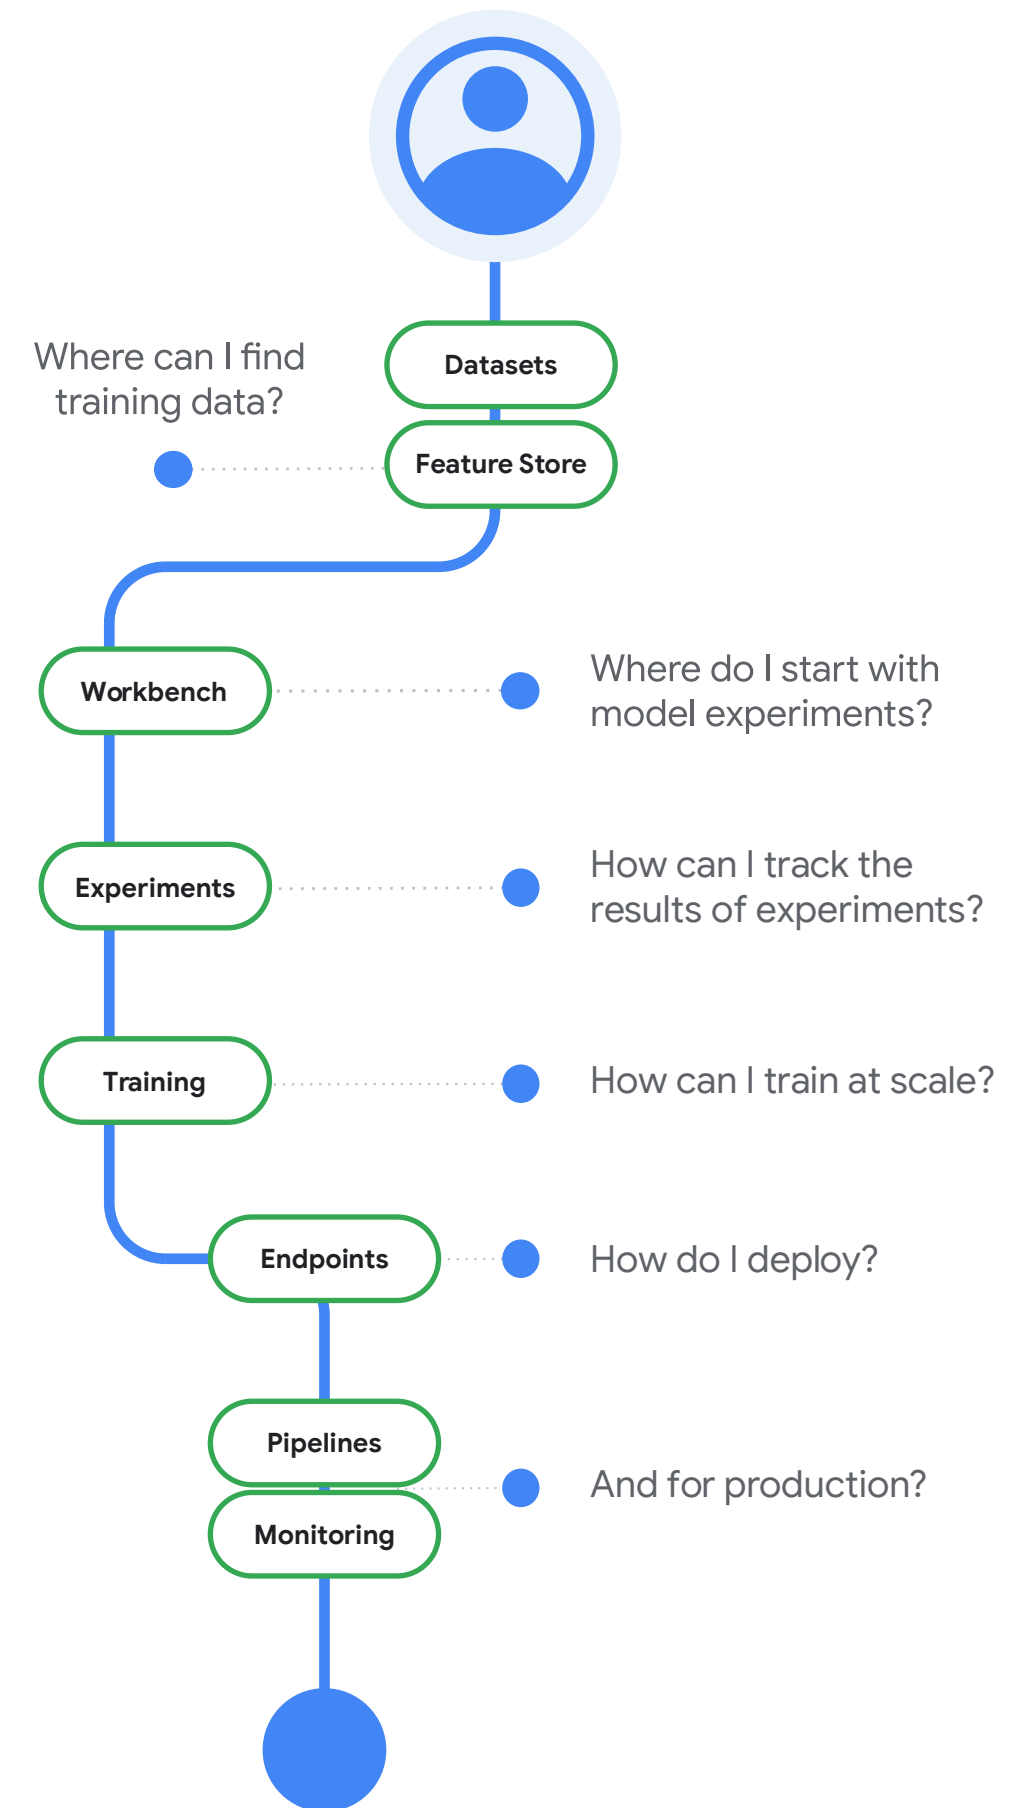
9. Google, [Accelerating the deployment of predictable ML in production](#), July 2022.

How Google Cloud helps with this challenge.

[Vertex AI](#) provides seamlessly integrated tools and automated MLOps workflows that help users deploy successful AI models and effectively manage them.

With [Google Cloud](#) and [Vertex AI](#), users can:

- Serve, share, and reuse ML features using the [Vertex AI Feature Store](#).
- Deploy models faster with [AutoML](#) and simplify end-to-end AI/ML processes with [Vertex AI notebooks](#), a fully managed Jupyter Notebook-based environment.
- Track, evaluate, and tune model performance during training runs with [Vertex AI Experiments](#), and tune hyperparameters in complex ML models with [Vertex AI Vizier](#).
- Discover, use, and govern any custom ML model including those built in [BigQuery ML](#) and [AutoML](#) using [Vertex AI Model Registry](#) for centralized model storage.
- Automate, monitor, and govern your ML systems by orchestrating your ML workflows in serverless environments with [Vertex AI Pipelines](#), and by storing your workflow's artifacts with [Vertex ML Metadata](#).
- Maintain optimal model performance and long-term health by using [Vertex Model Monitoring](#) to detect skew (variances between training data and input data) as well as prediction drift (significant changes between feature data in training models and production models).
- Quickly detect underlying model issues that cause inaccuracies with [Vertex Explainable AI](#).



Customers who have sped up AI and ML lifecycles and **protected technology choice** using Google Cloud include:

- CNA Insurance
- Boa Vista Serviços
- Vodafone
- Apna





Customer: CNA Insurance

Industry: Financial Services

Location: United States

Google Cloud:

[AutoML](#)

[BigQuery](#)

[Vertex AI](#)

Insurer builds its ML Model Factory in 12 weeks, accelerating model production from days to hours.

Challenge

Data scientists were spending too much time and money developing ML models to solve business problems.

The company needed faster insight into revenue streams, cash reserves, and other forecast KPIs to improve decision-making.

Solution

CNA Insurance consolidated its data silos in a global data lake with BigQuery.

Engineers built the company's Model Factory. It provides easy-to-use modeling environments with Vertex AI notebooks, reusable ML features using the Vertex AI Feature Store, and model training with AutoML.

Tip

Rather than developing individual ML models for specific projects, CNA Insurance developed its Model Factory. Using this comprehensive, automated ML platform established a framework to allow for continuous business feedback, analysis, and improvement.

Value

Reusable, automated data migration patterns shortened data migration timelines by nearly 50 percent.

The IT team deployed the Model Factory prototype in 12 weeks. Ten weeks later, they had built 14 models.

CNA Insurance accelerates time to market for ML models from days to hours, and data scientists analyze and recalibrate models faster.

The organization is adding a reusable feature layer and more model-monitoring capabilities to its Model Factory to further simplify and expedite model development.

For more information about CNA's use of Google Cloud, please [read this case study](#) and [watch this video](#).

Customer: Boa Vista Serviços

Industry: Financial Services & Insurance

Location: Brazil

Google Cloud:

[BigQuery](#)

[Cloud Storage](#)

[Dataproc](#)

[Vertex AI](#)

Fintech company **increases model processing** by up to 20 times.

Challenge

As one of the main analytical-intelligence and credit-bureau organizations in Brazil, Boa Vista Serviços manages information about 280 million people. This data supports credit activities and services that protect Boa Vista's customers including financial institutions, fintechs, and Brazil's largest corporate and retail groups.

The company's on-premises infrastructure was slowing down innovation, particularly data scientist's ability to train models and implement modern ML techniques.

Solution

Boa Vista created a centralized data lake on Google Cloud to improve data processing, scalability, and governance. Cloud Storage provides the foundation for a data mesh architecture, Dataproc runs data ingestion processes, and BigQuery stores structured data.

Data scientists use Vertex AI to build, train, and implement ML models.

Value

The analytics department delivers nine times more models and trains each model between three and five times longer than with the previous environment.

Ingesting large volumes of data now takes a few minutes instead of more than 24 hours.

The infrastructure team automated more than 1,000 hours of operational work.

Operational savings enabled Boa Vista to cut nearly every product's price. Bluebox now costs 15 percent of its original value.

For more information about Boa Vista Serviços' use of Google Cloud, please [read this case study](#).

“We have gained a lot of processing power, which has significantly increased the potential of the people working here. The team’s performance capacity as a whole has greatly increased as a result.”

Ricardo Orlando, CTO, Boa Vista Serviços



Customer: Vodafone

Industry: Telecommunications

Location: United States

Google Cloud:

[BigQuery](#)

[BigQuery ML](#)

[Dataproc](#)

[Vertex AI](#)

Global telecommunications provider moves AI and ML models from proof of concept into production 80% faster.

Challenge

The telecommunications provider wanted to use AI and ML at scale to improve customer experiences, increase network performance, and accelerate research in new technologies and services.

Solution

The company's Global Big Data and AI organization moved all of Vodafone's data onto one platform built with BigQuery, Dataproc, and Cloud Data Fusion.

This unleashed a huge number of use cases and increased demand for AI/ML capabilities, so next it built AI Booster, an internal ML platform powered by Vertex AI.

Value

By improving MLOps, the company can move models from proof of concept to production in as little as four weeks, which is 80 percent faster than before.

Data scientists and ML engineers now have access to on-demand ML learning, and within minutes, they can obtain a fully functional AI Booster environment with all the right guardrails, controls, and approvals to experiment.

Vodafone minimizes costs, enhances flexibility, meets security requirements, and improves the efficiency of hundreds of global data scientists and ML engineers.

For more information about Vodafone's use of Google Cloud, please [read this blog](#).

“We've created a common standard for deploying ML models at scale on Google Cloud. **The benefit to one data scientist alone is significant, so scaling this across hundreds of data scientists can really change the business.**”

Jamie Curtis, Practice Lead for MLOps, Vodafone

The Apna logo consists of the word "apna" in a bold, lowercase, sans-serif font. Below the text is a horizontal bar divided into three segments of equal length: green on the left, blue in the middle, and yellow on the right.

Customer: Apna

Industry: Professional Services, Education

Location: India

Google Cloud:

[BigQuery](#)

[Cloud SQL](#)

[Google Kubernetes Engine](#)

[Pub/Sub](#)

[Vertex AI](#)

India's largest jobs and professional networking platform **unlocks global growth potential** with Vertex AI.

Challenge

Apna is one of India's largest job and professional networks, with 22 million users in more than 70 cities. Valued at over \$1 billion within two years of launching, the unicorn startup now supports 18 million job interviews and more than 35 million peer-to-peer conversations every month.

To "democratize the Indian job hunt," Apna sought an advanced AI platform for nimble creation of ML models that could facilitate continuous platform improvements amid rapidly evolving socioeconomic conditions.

Solution

Rather than creating machine learning models from scratch with traditional data analytics engines, Apna adopted Vertex AI along with an agile cloud infrastructure built on Google Cloud.

Using Vertex AI, data scientists deployed an Apna algorithm with little manual work within weeks.

Apna uses CloudSQL as its fully managed database, Pub/Sub for its data integration service, and BigQuery as its scalable data warehouse supporting scalable analytics.

Value

Apna builds AI models 20 percent faster than it could using a traditional data analytics engine.

Data scientists launch up to seven AI experiments per day to fine-tune Apna's job-matching platform and support employers' and candidates' fluctuating needs.

The company's Vertex AI-enabled algorithms analyze up to 500 million user interactions per day, improving user experience and bolstering platform safety and security by uncovering fraud through keyword detection.

hundreds of global data scientists and ML engineers.

For more information about Apna's use of Google Cloud, please [read this case study](#).

“With Vertex AI we're able to build models and deploy them fast. It's so simple to use that we can run multiple ML experiments on a daily basis, testing them live on a small proportion of users, gauging metrics, and making rapid tweaks. **That's how we constantly fine-tune our platform.**”

Suresh Khemka, Head of Platform Engineering and Infrastructure, Apna

Challenge 4:

Steep compute and data requirements can create infeasible resource demands.

The strategies companies adopt in managing their AI and ML infrastructure can expedite deployment and minimize costs—or hurt both.

Challenges that often limit ROI include:

- Increasing friction between data scientists and IT infrastructure teams over cost and timelines.
- Expecting data scientists to be infrastructure or operations engineers, tasked with keeping models accurate, explainable, scalable, disaster resistant, and secure.
- Deploying models on hardware that is not optimized for AI and ML jobs.

Solution: Simplify model deployment and infrastructure management.

Enable your employees to focus on what they do best, ease the burden on infrastructure teams, and ensure right-sized provisioning by using ready-to-go infrastructure services and automated operational tools.

To maximize data usage, control costs, simplify scale out, and reduce data proliferation, ensure that the infrastructure platform you choose:

- ✓ Provides serverless compute and hardware designed for specific AI tasks like training and serving ML models.
- ✓ Minimizes data movement.
- ✓ Scales automatically based on your traffic so that you pay only for what you use.

“Of all the spending in the various AI market segments, AI hardware is by far the smallest. What this should tell organizations is that nickel-and-diming purpose-built hardware for AI is absolutely counterproductive, especially given the fast-growing compute demand from increasing AI model sizes and complexities.”¹⁰

Peter Rutten, Research Vice President, Performance Intensive Computing, International Data Corporation

How **Google Cloud** helps with this challenge.

[Google Cloud](#) provides scalable, high-performance, and cost-effective infrastructure for every ML workload. Choose from GPUs, TPUs, and CPUs to support a variety of use cases including high-performance training, low-cost inference, and large-scale data processing.

The fully managed AI infrastructure in [Vertex AI](#) also makes it easy to quickly provision and scale ML environments, automate orchestration, manage large clusters, and set up low-latency applications with minimal infrastructure knowledge.

- [Vertex AI notebooks](#) provides fully managed, scalable, enterprise-ready compute infrastructure with security controls and user management capabilities.
- The managed training service in [Vertex AI](#) automatically provisions and manages large clusters, with access to on-demand GPUs and TPUs, queue management, and built-in hyperparameter tuning.
- The [Vertex AI](#) managed prediction service enables low-latency serving to autoscaling endpoints with right-sized GPUs and CPUs.
- [Vertex AI Pipelines](#) helps you to automate, monitor, and govern your ML systems by orchestrating your ML workflow in a serverless manner.

By storing data in [Google's data cloud](#), teams can also minimize data movement by reducing infrastructure sprawl and complexity.

Customers who have **simplified model deployment** and infrastructure management using Google Cloud include:

- Limepay
- Wayfair
- SUBARU Corporation





Customer: Limepay

Industry: Financial Services

Location: Australia

Google Cloud:

[AutoML](#)

[BigQuery](#)

[Cloud Identity](#)

[Google Kubernetes Engine](#)

[Looker](#)

Fintech company manages millions of dollars in daily financial transactions and **speeds development of ML-based services** with a single, reliable platform.

Challenge

To ensure that its merchant customers' payment services are continuously fast and available, Limepay had to improve the resilience and scalability of its global cloud platform.

The company also wanted to speed up the development of new offerings.

Solution

In 18 months, the organization migrated its business off a legacy cloud platform and rebuilt it in Google Cloud, improving service reliability, platform scale, and administrative simplicity.

To automate provisioning and infrastructure maintenance, and ensure high availability for all customers, Limepay uses the fully managed Kubernetes Engine and BigQuery services.

Technical and business employees create insights and visualizations with Looker.

Teams make use of pretrained models in AutoML to speed up ML development.

Value

Limepay now builds, tests, and continuously deploys new features for its merchants such as dashboards that provide real-time insights about their customers.

Developers can quickly create new clusters for bespoke projects in independent environments to accelerate development without interfering with shared infrastructure.

For more information about Limepay's use of Google Cloud, please [read this case study](#).

“AutoML offers pretrained models, so we can harness the power of AI to enrich data. **Now, we're exploring AutoML on BigQuery to uncover hidden patterns**, learn more about our consumers, and improve lifetime value for our merchants.”

Andy Britz, Chief Technology Officer, Limepay



Customer: Wayfair

Industry: Retail & Consumer Goods

Location: United States

Google Cloud Platform:

[BigQuery](#)

[Looker](#)

[Spark on Google Cloud](#)

[Vertex AI](#)

Online retailer trains models up to **10 times faster** with tens of millions of datasets and little custom code.

Challenge

To drive growth, Wayfair must deliver personalized recommendations for more than 31 million active customers across more than 22 million products regardless of traffic spikes.

The company also wanted to expedite the development of new services, expand its catalog, and improve how it partners with suppliers.

The only way the organization could meet these requirements was by scaling its use of data and AI, and building ML into the fabric of its decision-making.

Solution

Engineers built one core data layer on BigQuery to support model development and training because the service can balance scalability and performance.

To further expedite model development and training, Wayfair also uses Vertex AI features including Vertex AI notebooks, Vertex AI Feature Store, and Vertex AI Pipelines.

Tip

To simplify model maintenance, Wayfair built a documented, easy-to-scale CI/CD pipeline using Vertex AI Pipelines and best practices. This motivates people to adopt a standardized way of working, which improves efficiency.

For data processing, teams use Spark on Google Cloud, and for data insights they use Looker.

Value

Wayfair accelerates innovation and the delivery of better customer and supplier experiences.

Data scientists train models with tens of millions of datasets distributed across clusters five to ten times faster, with little or no custom code.

With less friction in building and deploying models at scale, data scientists require less time to move models into production and they experiment more.

Because there are fewer lines of custom code, data scientists have greater trust in their model ecosystem.

For more information about Wayfair's use of Google Cloud, please [read this blog](#).



Customer: SUBARU Corporation

Industry: Manufacturing

Location: Japan

Google Cloud Platform:

[Cloud Storage](#)

[Dataflow](#)

[Vertex AI](#)

Automotive company uses AI, ML, and deep learning to improve its advanced driver-assistance system.

Challenge

SUBARU Corporation aims to achieve the target of zero fatal traffic accidents caused by its own cars by 2030. To accelerate AI research and development, and improve its EyeSight driver support system, the company opened an AI development lab.

The lab's workstations proved inadequate for full-scale deep learning research, so the company opted to develop and train AI models in the cloud.

Solution

The company chose Google Cloud as its AI and ML platform because it provides managed services such as Vertex AI that include Vertex AI notebooks for managed notebooks and Vertex AI custom training for development.

To speed up model performance, Subaru provisions the right number of NVIDIA A100 GPUs in each virtual machine using Google Compute Engine, and the service automatically scales resources as needed.

SUBARU controls access to its Google Cloud services using Google Identity-Aware Proxy.

Value

SUBARU reduces pre-processing time for model data from several days to 30 minutes.

Regardless of hardware availability, engineers can prepare development environments for new staff in one day.

The company prevents the leakage of sensitive information.

Engineers spend significantly less time scaling infrastructure and coordinating budgets.

For more information about SUBARU's use of Google Cloud, please [read this case study](#).

“I chose Google Cloud from many platforms because, at the time of selection, it had multiple managed services such as Vertex AI. ... It was also fascinating to have high-performance hardware that could handle large-scale machine learning operations”

Toshimi Okubo, Senior Engineer of AI R&D Section, ADAS Development Department, Engineering Division, SUBARU Corporation

Your AI Journey.

Wherever you are in your data and AI journeys, it's always best to:

- Start with a well-defined business problem and take advantage of AI to drive business impact.
- Establish best practices around governance and management processes as early as possible to save time, minimize costs, and avoid potential issues.
- Democratize AI access to empower employees across your organization.

As you gain proficiency, the sky's the limit for AI and ML innovation. Google aspires to create technologies that solve important problems and help people in their daily lives. We are optimistic about the incredible potential for AI and other advanced technologies to empower people, drive progress for current and future generations, and improve the common good.







For more information about overcoming your business challenges with AI and ML, please go to the [Google Cloud AI](#) page and read our [AI & machine learning blog](#).

“AI is for everyone, and it should be easy to harness in workflows of all kinds and for people of all levels of technical expertise. We see our customers’ accomplishments as validation of this philosophy and a sign that we are taking away the right things from our conversations with business leaders.”¹¹

Andrew Moore, Vice President and General Manager,
Cloud AI and Industry Solutions, Google

Appendix A: AI industry use cases.

AI and ML helps organizations across industries meet their requirements. This table highlights just some of the popular use cases that we're seeing.

 Retail	 Healthcare	 Financial Services	 Media, Entertainment & Gaming	 Industrial & Manufacturing	 Public Sector
Demand Forecasting	Telehealth / Virtual Care	Anti-Money Laundering (AML)	Media Asset Management	Industrial Adaptive Controls	Public Health Platform
Retail Search	Interoperability Accelerator	Know Your Customer (KYC)	Transcoding & optimization	Manufacturing Visual Inspection	Transit Reopening
Recommendations	Hospital Impact Forecasting	Lending Doc Processing	LTV Optimization	Logistics Optimization	Digital Social Safety Nets
Inventory Optimization	Biomedical Data Analytics	Risk Analytics	Virtual opponents	Connected Operations	Virtual Government

For more information about AI and ML use cases, please visit the [Vertex AI page](#).

