

A photograph of three business professionals in a meeting. A man with glasses and a beard is leaning over a laptop, looking at the screen. A woman is on the left, looking at the laptop. An older man with grey hair is on the right, also looking at the laptop. The background is blurred, showing office shelves. A large white semi-circle is overlaid on the bottom left of the image.

AH≡AD

Using MLOps & Platform
Engineering to Enhance the
Machine Learning Lifecycle

Executive Summary

This whitepaper explores the integration of MLOps with platform engineering, highlighting its significance in managing the machine learning (ML) lifecycle as more enterprises adopt the technology to drive innovation and see an increased need for better MLOps practices.

We'll discuss the challenges of traditional ML workflows and how MLOps can address them, with a focus on platform engineering's key role in MLOps via automation. We will show how this enables self-service tools for data scientists and ensures consistent infrastructure across environments. With these practices, companies can accelerate time-to-delivery and improve the efficacy of their machine-learning initiatives.

Introduction

Exciting new AI technologies have led to a boom in the creation of ML models, but getting these models to work reliably in real-world applications can be tricky. MLOps, short for Machine Learning Operations, helps bridge this gap. MLOps combines ideas from DevOps (used for software development) with the world of ML workflows. This makes it easier to move ML models from creation to real-world use.



The MLOps Community reported substantial growth in the Platform Engineering MLOps [market](#), with approximately **\$1 billion invested across 43 vendors**.

A recent [Gartner](#) study shows that even enterprises with AI experience struggle to translate prototypes into production, with only [53% achieving success](#). This statistic underscores the critical need for robust MLOps practices.

Moreover, there are certain [platform engineering principles](#) which can be used to build a strong foundation for MLOps. This will help companies to operationalize their machine learning models and empower data science teams to focus on innovation while ensuring consistent and scalable model deployments.

Examining the Machine Learning Lifecycle

Machine learning is playing a significant role in transforming businesses across a growing range of industries. However, developing a strong machine learning model and making it all the way to production can be easier said than done.



According to [McKinsey](#), only 15% of ML projects succeeded.

Here enters MLOps, serving as the practice that automates and manages the entire machine-learning lifecycle. It encompasses each stage of bringing an ML model from conception to production and beyond. MLOps also ensures the ongoing performance and improvement of ML models.

Here's a breakdown of the key phases within the ML lifecycle:

Data Collection & Preparation: This stage involves gathering relevant data for training the model. MLOps practices can ensure efficient data ingestion and transformation to prepare it for model training.

Model Training & Experimentation: Data scientists experiment with different algorithms and hyperparameters to build and train the model. MLOps can automate this process and facilitate version control for models and datasets, enabling easy rollbacks and comparisons.


Model Evaluation & Validation: The trained model is evaluated on unseen data to assess its performance, which also helps to identify any potential biases. MLOps can help automate this testing process to ensure model quality and reliability.

Model Deployment: After validation, the model needs to be deployed into production for real-world use. MLOps practices can automate model deployment pipelines and infrastructure provisioning for seamless integration with existing systems.

Model Monitoring & Feedback: Continuously monitoring the model's performance in production is another important component within the lifecycle. MLOps helps with real-time monitoring of model metrics (accuracy, drift, etc.) and provides alerts for potential issues.

Model Retraining & Improvement: With the increase of real-world data, the model performance can degrade over time. MLOps can automate retraining pipelines to keep models up-to-date and ensure optimal performance.

MLOps: Eternal Knot



[Source](#)

MLOps Landscape & Challenges

MLOps is a rapidly evolving ecosystem filled with numerous tools and frameworks. Still, companies implementing [MLOps often face significant challenges](#). Let's explore both the landscape and the hurdles that need to be overcome:

- **Fragmented Toolset:** MLOps is fragmented with various open-source and commercial tools suitable for different stages of the ML pipeline. This can lead to siloed workflows and complex integrations if tools are not carefully selected at the outset of a project.
- **Lack of Standardization:** MLOps lacks standardized practices and methodologies. This inconsistency can make it difficult to ensure the reproducibility and maintainability of ML models across environments.
- **Data Management:** MLOps relies heavily on high-quality, governed data. Thus, it can be challenging to ensure optimal data versioning and lineage tracking while addressing data drift (data distribution changes over time).

- **Infrastructure Management:** Managing the infrastructure for training and deploying complex ML models requires more budget and resources. This can be particularly challenging for organizations with limited resources.
- **Explainability & Fairness:** It is important to understand how AI models make predictions and ensure that they are unbiased. MLOps tools lack the features to address these concerns.


By understanding these challenges, organizations can develop a more targeted approach toward MLOps adoption and select tools that effectively address their specific needs.

Platform Engineering for MLOps

The challenges above highlight the limitations of managing ML workflows – this is where platform engineering principles come into play, offering a structured way to bridge the gap between data science and production.



[50.4% of companies](#) that have implemented platform engineering experienced better standardization and 43.2% noted quicker product delivery.



[Source](#)

HERE'S HOW PLATFORM ENGINEERING EMPOWERS MLOps:

Automation & Self-Service

Platform engineering helps automate the routine tasks involved in building and using machine learning models, like data cleaning, training, and deployment. This saves data scientists time for more important work. Furthermore, data science teams may independently manage their workflows and provision resources thanks to self-service platforms.

Infrastructure as Code (IaC)

IaC practices enable infrastructure configurations for training and serving ML models in code. This makes it easy to set up these systems in different environments, like development or production, and ensures everything works the same way each time.

Standardized Workflows

Platform engineering encourages standardized templates and pipelines for different stages of the ML lifecycle. This promotes code reusability and improves the overall maintainability of ML projects.

Monitoring & Observability

Effective monitoring of model performance and health is vital for reliable MLOps deployments. Platform engineering principles help with the integration of monitoring tools and dashboards, which assists data scientists and operations teams in proactively identifying and addressing issues.

Security & Governance

Platform engineering helps set up security frameworks and access controls for data and models. This ensures compliance with regulations, minimizing the risk of data breaches or unauthorized access.

Streamlined Workflows & Reliable Deployments

Platform engineering provides a solid foundation for successful MLOps. This can help organizations streamline their ML workflows and ensure the reliable and responsible deployment of ML models.

MLOps Platform Design Considerations

Now that we've covered the primary benefits of using platform engineering to augment MLOps, let's explore the key considerations for designing the MLOps platform itself:

Scalability & Elasticity

The platform should be able to handle varying workloads efficiently. This may involve using containerization technologies like Docker for model deployment. Additionally, cloud-based infrastructure can be used with auto-scaling capabilities to handle spikes in demand.

Security Considerations

MLOps systems manage models and sensitive data by putting vulnerability management, data encryption, and access controls into action. For centralized monitoring and threat detection, consider integrating MLOps with current security information and event management (SIEM) systems.

Monitoring & Logging


Granular monitoring of model performance and data pipelines is crucial. Thus, the platform should integrate with monitoring tools to provide real-time insights and facilitate proactive issue identification. Logging best practices will also help with debugging and troubleshooting potential problems.

Integration with Existing Tools & Data Pipelines

A successful MLOps platform should easily integrate with the existing data science and engineering toolchain within the organization. This may involve supporting popular machine learning frameworks (e.g., [TensorFlow](#), [PyTorch](#)) and data warehousing solutions (e.g., [Snowflake](#), [Redshift](#)).

Open Source versus Managed Services

One way to build the MLOps platform is from scratch using open-source tools. The other way is to use managed services offered by cloud providers that depend on factors like organizational expertise, resources, and desired level of control. Open-source solutions offer more flexibility but require more technical expertise to manage, while managed services provide a faster time to value but can limit customization options.



Organizations can implement these factors in the design phase to build an MLOps platform that is scalable and secure. This strategy will also ensure the MLOps platform integrates with the company's existing infrastructure and workflows.

[Source](#)

MLOps Use Cases & Applications

MLOps help businesses use their ML models more efficiently in every industry. Here are a few cases, along with their practical applications, that highlight the potential of MLOps:

01

Fraud Detection and Risk Management

Detecting fraud is of the utmost importance for companies in every industry, which is reflected in the increased adoption of security solutions. According to [Statista](#), the fraud detection and prevention market value jumped from \$19.5 billion in 2017 to \$63 billion by 2023.

Financial institutions and organizations in a wide array of other industries use MLOps to develop and deploy real-time fraud detection models. These models can analyze transaction data and identify suspicious patterns. Automated workflows and continuous monitoring work to ensure that these models stay up-to-date and prevent fraudulent activities.

REAL-WORLD APPLICATION


Every day, [American Express](#) analyzes billions of transactions using MLOps. As a result, they were able to recognize and stop fraudulent activity in real time while minimizing false chargebacks and bringing in millions of dollars for the business.

02

Predictive Maintenance in Manufacturing

MLOps models can predict equipment failures based on sensor data. Proactive maintenance based on these predictions can minimize downtime and optimize resource allocation.

[Deloitte](#) shows that predictive maintenance reduces unexpected equipment failures by 70%, increases operational productivity by 25%, and cuts maintenance costs by 25%.



REAL-WORLD APPLICATION

[GE Aviation](#) uses MLOps to develop and deploy models that predict engine failures in airplanes. This makes it possible for airlines to plan maintenance and prevent any mishaps during flights.

03

Inventory Management and Demand Forecast in Retail

ML models can forecast future demand for specific products based on historical sales data and external factors. This enables optimized inventory management, reducing stockouts and overstocking, and ultimately improving overall supply chain efficiency.

REAL-WORLD EXAMPLE

[Walmart](#) uses ML algorithms to predict demand for thousands of products across their stores. This allows them to optimize their inventory levels and reduce stockouts by 30%. The strategy led to cost savings and improved the customer experience.

04

Personalized Customer Experiences in E-commerce

E-commerce platforms use MLOps to develop recommendation engines that suggest relevant products to customers based on their past behavior and preferences. Automated model training and deployment pipelines make sure that these recommendations are personalized, driving better customer engagement and revenue.

REAL-WORLD EXAMPLE

[Amazon](#) is a leader in using MLOps for personalized recommendations. Their MLOps platform, which is built on open-source tools like TensorFlow, analyzes vast amounts of customer data to offer personalized product suggestions that result in higher conversion rates and customer satisfaction.

05

Churn Prediction in Telecommunications

Telecommunications organizations use MLOps to predict customer churn. Proactive interventions based on these predictions can improve customer satisfaction and loyalty that will eventually result in reduced churn rates.

REAL-WORLD EXAMPLE

[Verizon](#) uses MLOps to develop churn prediction models. They identify customers at risk of leaving and offer them targeted promotions and loyalty programs. This strategy results in churn reduction and increased customer lifetime value.

Considerations & Best Practices for Successful MLOps

Building a successful MLOps practice goes beyond simply utilizing specific tools. It requires a comprehensive approach that considers various factors, such as:

MLOps Maturity Assessment

Before anything, it is important to understand your MLOps maturity level. By identifying areas for improvement, this self-evaluation will assist in developing an adoption roadmap for MLOps. Cloud providers like Microsoft and Google offer MLOps maturity models for this assessment.

Automation is Key

Streamlining the ML lifecycle through automation is a core principle of MLOps. This can involve automating repetitive tasks like data cleaning and feature engineering. Automation can also help with data splitting for training and testing and even writing model training code. Automating these tasks frees up your time for more strategic work and ensures consistency in the ML pipeline.

Experimentation & Tracking

Data science is a process that involves experimentation and exploration. MLOps should facilitate this by helping data scientists track their experiments. Tracking experiment details like parameters used and code versions can lead to reproducibility and simplify the model auditing process.

Continuous Integration/Delivery (CI/CD) & Beyond

Automated [CI/CD](#) is a widely used DevOps technology, with 81% of companies reporting its use. MLOps builds on the idea of [CI/CD](#), which automates testing and deployment for software, by adding more 'continuous' steps:

Continuous Training (CT):

Models are regularly updated with new data to remain accurate as things change.

Continuous Monitoring (CM):

Model performance is constantly checked to catch problems early and retrain or retire the model as necessary.

Cultural Shift & Collaboration

MLOps works best when everyone works together. Data scientists, engineers, and operations teams need to communicate clearly, and each team should know their role and what they're responsible for. An open learning environment helps everyone involved in machine learning to succeed.



Wrapping Up

Organizations' management of the machine learning lifecycle is being rapidly transformed by MLOps, which enables scalability, dependability, and efficiency. Businesses can ensure consistent and high-quality results by streamlining the creation, deployment, and maintenance of ML models via the combination of MLOps and platform engineering. With the correct tools and best practices in place, MLOps' problems — such as dispersed toolkits, lack of standardization, and complicated data management — can be solved, opening the door for effective ML projects.

New tools and technologies are always developing, causing the MLOps landscape to change quickly. Maintaining awareness of these developments and iteratively improving your MLOps procedures will guarantee that your enterprise stays at the forefront of data-driven innovation.

Contributing Author:

Stephen D. King, Chief Architect, Digital Solutions

AHEAD

Combining cloud-native capabilities in software and data engineering with an unparalleled track record of modernizing infrastructure, we're uniquely positioned to help accelerate the promise of digital transformation.

[Contact Us](#)

National Hubs

CHICAGO

444 W. Lake Street
Suite 3000
Chicago, IL 60606

NEW YORK

500 Fifth Avenue
Suite 1500
New York, NY 10110

ATLANTA

1117 Perimeter Center
Suite W406
Atlanta, GA 30338

SAN FRANCISCO

2000 Crow Canyon Place
Suite 250
San Ramon, CA 94583