
Reaching the Full Potential of Production Machine Learning

The MLOps features and benefits of the Cloudera Machine Learning Platform

CLOUDERA

FEATURING
FORRESTER[®]

Overcoming machine learning's scalability problem

When it comes to deploying and maintaining machine learning models at scale, the majority of enterprises find themselves coming up short.

The process of developing machine learning models is, by its very nature, experimental. However, the goal of any enterprise machine learning initiative is to deliver new business value. For that to happen, models eventually have to move into production—and that makes experimentation only half of the equation.


It's that leap from the lab to the production environment where most enterprises fall short. However, even if an enterprise succeeds in moving a handful of models into production, they can still struggle to manage, monitor, and govern those models. And the thought of scaling to hundreds or even thousands of machine learning models can feel daunting.

That's where machine learning operations (MLOps) comes in. It's a set of defined processes, best practices, and capabilities machine learning teams use from experimentation to production, helping them manage, deploy, monitor, retrain, and govern models at scale.

Cloudera Machine Learning (CML) with MLOps encapsulates those elements of MLOps so enterprises can successfully operationalize the entire machine learning lifecycle. On a single end-to-end platform, teams can seamlessly leverage their data pipelines directly from the source, experiment and train models in any language and with any libraries, deploy models into production—and then manage those models and the underlying data with total transparency and control.

In [MLOps in Enterprise Production Machine Learning](#), Forrester analyst Dr. Kjell Carlsson answered common questions enterprises have about MLOps. This eBook provides a closer look at the critical role MLOps plays in a sustainable, scalable enterprise machine learning production environment and how CML with MLOps is poised to change the enterprise machine learning landscape.

Source
Artificial Intelligence: The Next Digital Frontier, McKinsey Global Institute, June 2017





of enterprise machine learning projects never make it beyond the experimental stage.

Why your enterprise needs MLOps

MLOps helps data scientists and machine learning engineers collaborate seamlessly, iterate quickly and move machine learning models into production efficiently.


By enabling you to move machine learning models into production faster—and by making the management and maintenance of those models easier—MLOps carries with it the promise to unlock the business value of accurate, scalable machine learning models.

CML with MLOps fulfills that promise. It provides a single platform to build, train, deploy, operate, and scale machine learning models. It also enables collaboration between machine learning teams and business users.



“Business stakeholders rarely know exactly what they will need until it has been developed, and by helping iterate faster, MLOps also drives innovation and value in terms of delivering solutions that serve the business better.”

MLOps in Enterprise Production
Machine Learning, Forrester, May 2020





Visibility and control

CML with MLOps provides iterative, full lifecycle monitoring and lineage tracking across teams within the Cloudera Enterprise Data Cloud. From data sourcing to IT to production environments, this infrastructure supports multiple analytic functions working together on the same data while maintaining strict security, governance, and control measures.

Visibility across the full machine learning lifecycle eliminates silos and streamlines workflows across teams, mitigating the machine learning “black box” threat. And with increased visibility and control, teams can set model governance rules to respond quickly to mission-critical changes in production environments, operate models holistically, and continuously deliver results securely.



Power and ease

Data science teams can use CML with MLOps to mitigate the IT and siloed point solution roadblocks data science teams traditionally face when they need to access the necessary resources, IDEs, or libraries.

CML with MLOps supports hybrid deployments and can be utilized across multiple cloud services, such as Microsoft Azure and Amazon Web Services. Once models are deployed, you can scale operations with streamlined tools that ensure models in production are properly maintained and serving use cases across the business. CML with MLOps also implements universal standards—including model metadata, terminology, and prediction tracking—making deployment, monitoring, and governance cost-effective, repeatable, and transparent across your enterprise.



Results and confidence

With the ability to take workflows into production at scale and manage your machine learning projects with total transparency, you can automate business processes, harness predictive analytics, and create new products and services.

Due to the transparent, explainable, and governed nature of CML with MLOps, enterprise leaders can feel confident about using machine learning models and predictions in business decision making. CML with MLOps provides a visual interface for business teams to better understand underlying data, bridging any communications and skills gap between them and data scientists.

Source: MLOps in Enterprise Production Machine Learning, Forrester, May 2020

Key capabilities for production machine learning at scale

To operationalize the machine learning lifecycle, you need an end-to-end ML platform with MLOps capabilities that provides open standards-driven accuracy monitoring, model cataloging, and ongoing governance from data source to business use cases.

While there are many capabilities required for operationalizing machine learning at scale, Forrester categorizes the most important ones into the three main segments of production machine learning: deployment and serving, monitoring, and collaboration and governance.



3X

“Fast-growing firms were 3x more likely to be able to get their models into production and to monitor and retrain them on an ongoing basis.”

MLOps in Enterprise Production Machine Learning, Forrester, May 2020



Deployment and serving

CML with MLOps streamlines model deployment, making it frictionless and predictable. It's also built for intensive enterprise machine learning workloads, so IT administrators and data scientists can ensure high availability serving of machine learning models at scale with no downtime.

Data scientists are able to use the notebook or IDE of their choice to build, package, and deploy machine learning models—including open source tools and frameworks. This means they're never boxed in and forced to use tools and frameworks they may not be familiar with using. CML with MLOps supports Python, R, and Scala ecosystems along with RStudio, PyCharm, and Jupyter Notebook IDEs.

Machine learning teams can also deploy to and scale in on-premises, multi-cloud, or hybrid cloud environments—wherever running a model makes the most sense.



Monitoring

Using CML with MLOps, data scientists and machine learning engineers can track technical metrics—like uptime and throughput—and prediction metrics such as model accuracy and drift on an ongoing basis. Operators can track model performance against operational, accuracy, and business KPIs. Teams will be alerted to mission-critical issues with production environments, so potential problems can be addressed without disrupting models in production.

Additionally, guardrails can be put in place for the automatic scaling of both CPU and GPU compute resources, so usage stays within an acceptable range which helps keep budget under control.



Collaboration and governance

CML with MLOps was built for end-to-end governance—from data pipelines to data science workloads to models in production. Teams can employ repeatable workflows, streamline model operations, and deliver models into production with inherited security and unified authorization.

CML with MLOps also enables model cataloging by supporting the governance of large volumes of machine learning models. Data scientists and machine learning engineers will be able to leverage a first class model catalog and full lifecycle lineage built on the open source governance and metadata framework—Apache Atlas.

MLOps—the future foundation of enterprise machine learning

“If you don’t develop your MLOps capabilities, it is a good bet that your business competitiveness will decline.”¹

MLOps in Enterprise Production Machine Learning, Forrester, May 2020

Forrester predicts that enterprises investing in an end-to-end MLOps platform will yield greater business value from their machine learning efforts. They’ll be able to “detect and adjust to changes in the market faster as they will be reflected in the distribution of the data their models are consuming, the predictions their models are making, as well as the resulting outcomes.”¹

Conversely, Forrester expects enterprises pushing ahead without an MLOps strategy will continue to struggle getting their models into production. By not taking the necessary and diligent steps to operationalize machine learning, these enterprises will likely fail to realize the full business potential of their machine learning efforts.

Achieve production machine learning at scale

If MLOps bridges the gap between experimentation and production, then CML with MLOps paves the way forward for innovative production machine learning at scale. It gives data scientists unfettered access to data for analysis and modeling and it bridges the workflow and skills gap so they can move models easily from experimentation to full-scale production.

Data scientists get access to the resources they need, regulatory compliance can be maintained and proven, and machine learning teams work alongside business users to understand the needs of the enterprise and the business impact of their machine learning models.

Sources:

1 MLOps in Enterprise Production Machine Learning, Forrester, May 2020

2 Artificial Intelligence: The Next Digital Frontier, McKinsey Global Institute, June 2017

15%

higher than industry average profit margins reported for enterprises successfully employing machine learning in the business.²

Learn more

Learn more about this topic from our webinar, [Enabling Production MLOps at Scale—Hands on with Cloudera Machine Learning](#).

About Cloudera

At Cloudera, we believe that data can make what is impossible today, possible tomorrow. We empower people to transform complex data into clear and actionable insights. Cloudera delivers an enterprise data cloud for any data, anywhere, from the Edge to AI. Powered by the relentless innovation of the open source community, Cloudera advances digital transformation for the world's largest enterprises. Learn more at [cloudera.com](#)

© 2020 Cloudera, Inc. Cloudera and associated marks and trademarks are registered trademarks of Cloudera, Inc. All other company and product names may be trademarks of their respective owners.

For more information visit [cloudera.com](#)

US: +1 888 789 1488 | Outside the US: +1 650 362 0488

[Privacy Policy](#) | [Terms of Service](#)

CLUDERA