A POCKET GUIDE

The path towards ML-ightenment

A framework for optimizing value from machine learning operations

CREDERA

This guide offers a practical machine learning operations (MLOps) framework that captures the concepts and tools to accelerate a machine learning (ML) adoption journey.

We provide an example use case to bring it to life and explore why MLOps has emerged.

Contents

What is MLOps? 4

Why has MLOps emerged? 6

The crucial link to business strategy 8

Which problems does our MLOps framework solve? 10

How could an MLOps framework help with this issue? 12

How to implement the framework 14

Unpacking the framework 17

AI Strategy, use case, and operating model development 18

Rapid prototyping and model building 20

ML engineering and operationalization 22

How do I know if I need to adopt MLOps? 24

Further resources 26

What is MLOps?

MLOps frameworks sit at the intersection of data science, DevOps, and data engineering, and provide standardized capabilities throughout the ML model lifecycle. They encompass a set of methodologies that aim to effectively define, develop, deploy, and operationalize machine learning models.

MLOps frameworks can be seen as an 'enabler of artificial intellingence' in the same way that DevOps and DataOps frameworks provide a methodology for software and data engineering initiatives, but with provisions suited to the unique complexities of machine learning solutions in an enterprise environment.

MLOps exposes the intrinsic link between code and data to ensure that models are effective, relevant, reproducible, explainable, auditable, and traceable.

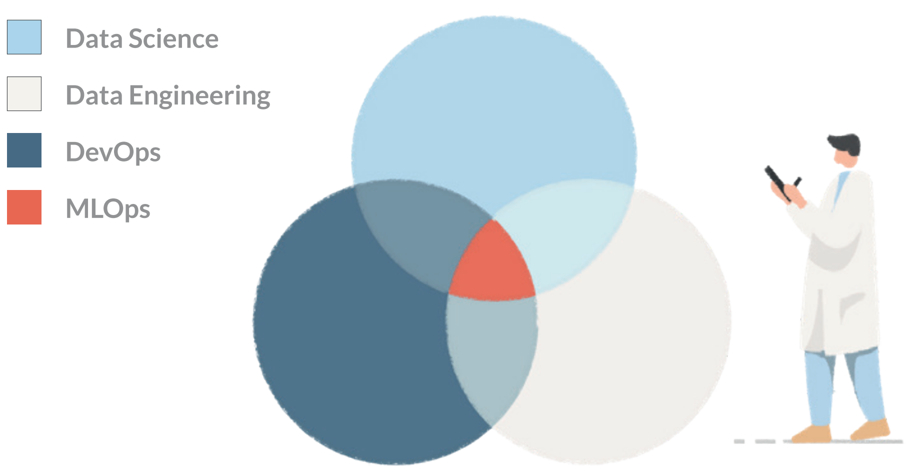

BACK TO CONTENTS

Why has MLOps emerged?

Why has MLOps emerged now? There is a need to transcend the hype of ML, embed methodology, and embrace innovation to create real value for enterprises.

Enterprise applications have been organically using machine learning, causing a need for robust methodology. There has been a shift from ‘siloed’ IT environments, created by data scientists, toward MLOps frameworks that accelerate the progression beyond the experimentation and prototyping stages onto enterprise-scale operationalization of models.

While examples of successful MLOps platform implementations do exist, an often-fragmented approach toward the end-to-end MLOps journey excludes proactive business engagement and a clearly defined operating model. This typically results in the benefits of ML models being missed, despite their high quality.

Credera’s MLOps framework brings ML model consumers into the critical path of model assurance, pre-deployment sign-off, and post-deployment monitoring, which are often afterthoughts in the lifecycle.

“There has been a shift from ‘siloed’ IT environments... toward MLOps frameworks...

BACK TO CONTENTS

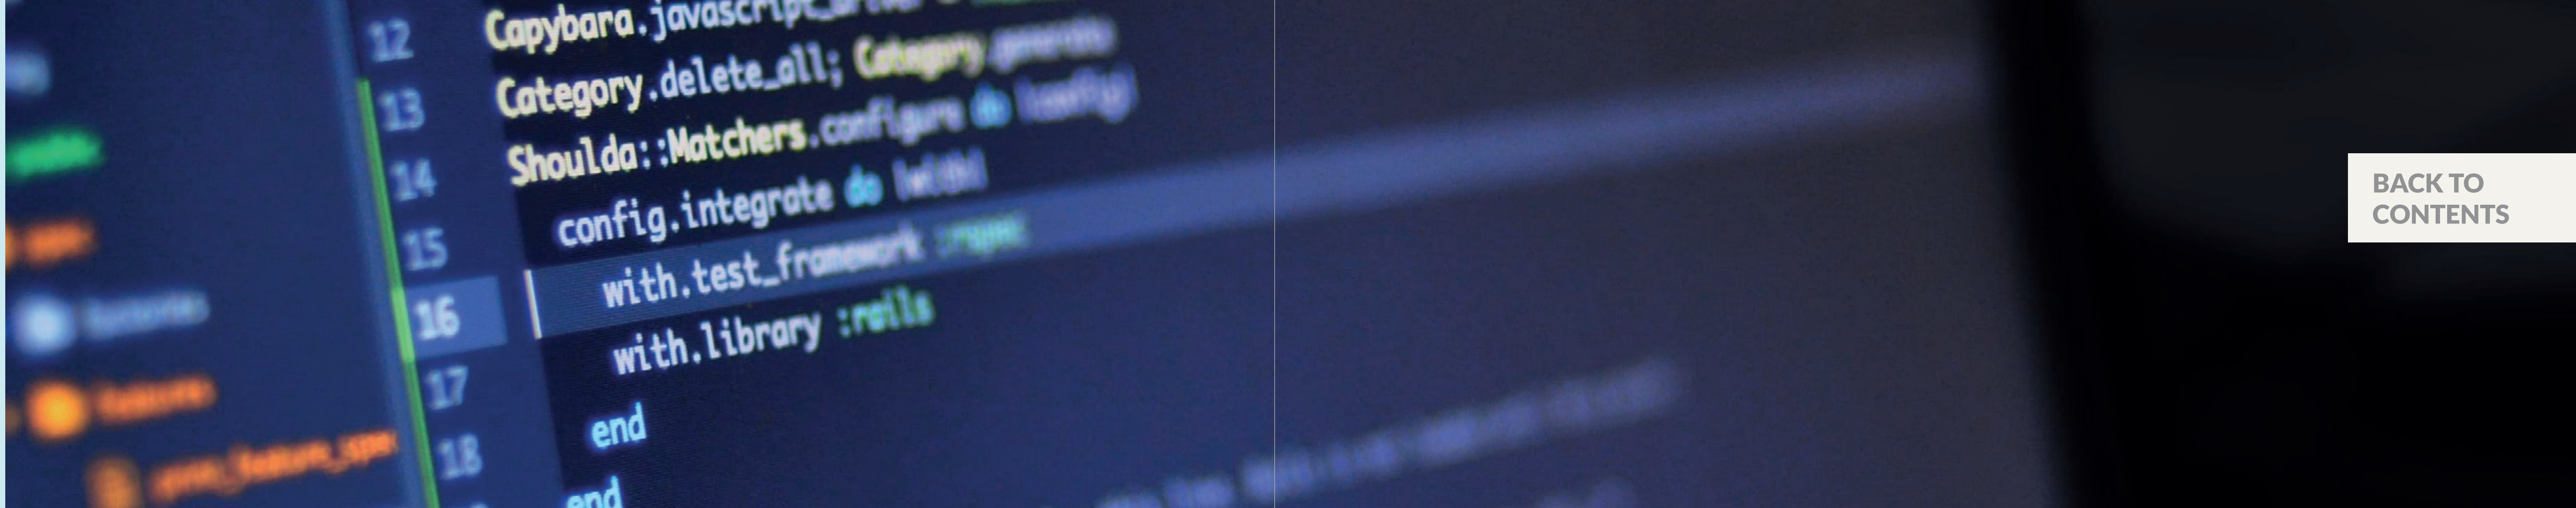

BACK TO CONTENTS

The crucial link to business strategy

The pertinent question to ask is, ‘Are our models truly driving business value?’ At Credera, we converge business and MLOps strategies through the following critical steps:

1

Defining business drivers with appropriate links to key progress indicators and insights. This ensures ML models create insights that contribute real value.

2

Defining frameworks that guide the development of ML use cases and the associated business case through the AI strategy.

3

Ensuring that ML use case development is compatible with the above frameworks and principles and informs operationalization requirements that feed into the MLOps implementation process.

4

Ensuring ethical considerations, such as bias, are considered and managed. ML models are a representation of the data. As such, data preparation in a way that controls bias, as well as algorithmic moderation and development beyond black-box approaches, is key. Similarly, data should be prepared in alignment with the overarching data governance strategy to ensure compatibility with the values of the enterprise.

Which problems does our MLOps framework solve?

MLOps addresses several persistent and emerging challenges that organizations face when embracing ML and AI. Examples of these challenges include:

- Opportunity costs associated with insufficient business engagement when shaping the future of ML in an organization.
- Cost inefficiencies due to a lack of agility in the ML model deployment process.
- A fragmented approach to assessing the impact of ML on employees and customers, reducing the ability to gain value.

CASE EXAMPLE

To put all of this into context, imagine an organization with fragmented data science and more traditional business intelligence functions, sourcing data from a range of separate sources and struggling to fully productionize meaningful models.

The business needs to urgently adapt in the aftermath of the pandemic and the cost-of-living crisis, which have both caused a spike in demand and an elevated need for personalization. The new CEO of the organization believes that ‘every company is a data company’ and would like to harness the state of the art in ML and analytics.

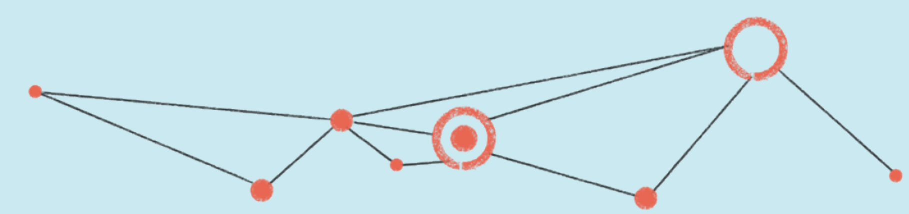

An ML model is proposed as a way of automating customer profiling to tailor existing services and develop new ones, as well as reducing the strain on operational teams. Attempts to successfully address the scale and complexity of deploying ML models and supporting infrastructure within a brownfield enterprise IT landscape while remaining compliant have so far been unsuccessful.

BACK TO CONTENTS

How could an MLOps framework help with this issue?

Should the organization hire ML engineers to respond to the anticipated demand spike? Should new, more efficient models be developed? Could marketing teams create a better view of the customer?

The fundamental principles of MLOps can help teams fully grasp the breadth of the issues and formulate a set of practical steps that aid the journey toward industrialization of the ML model lifecycle. These principles can also serve as a means of surfacing and reconciling the dichotomy of priorities between business, data science, development, and operational teams.

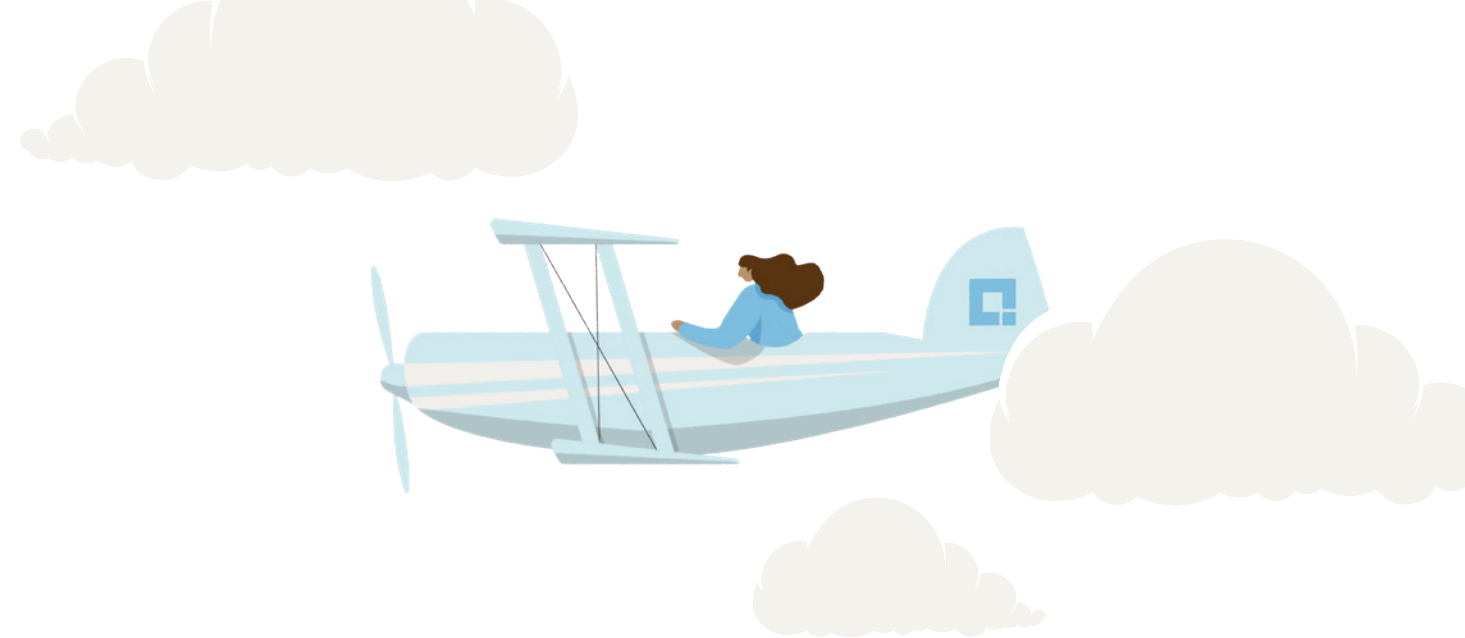

> “ The fundamental principles of MLOps can help teams fully grasp the breadth of the issues...

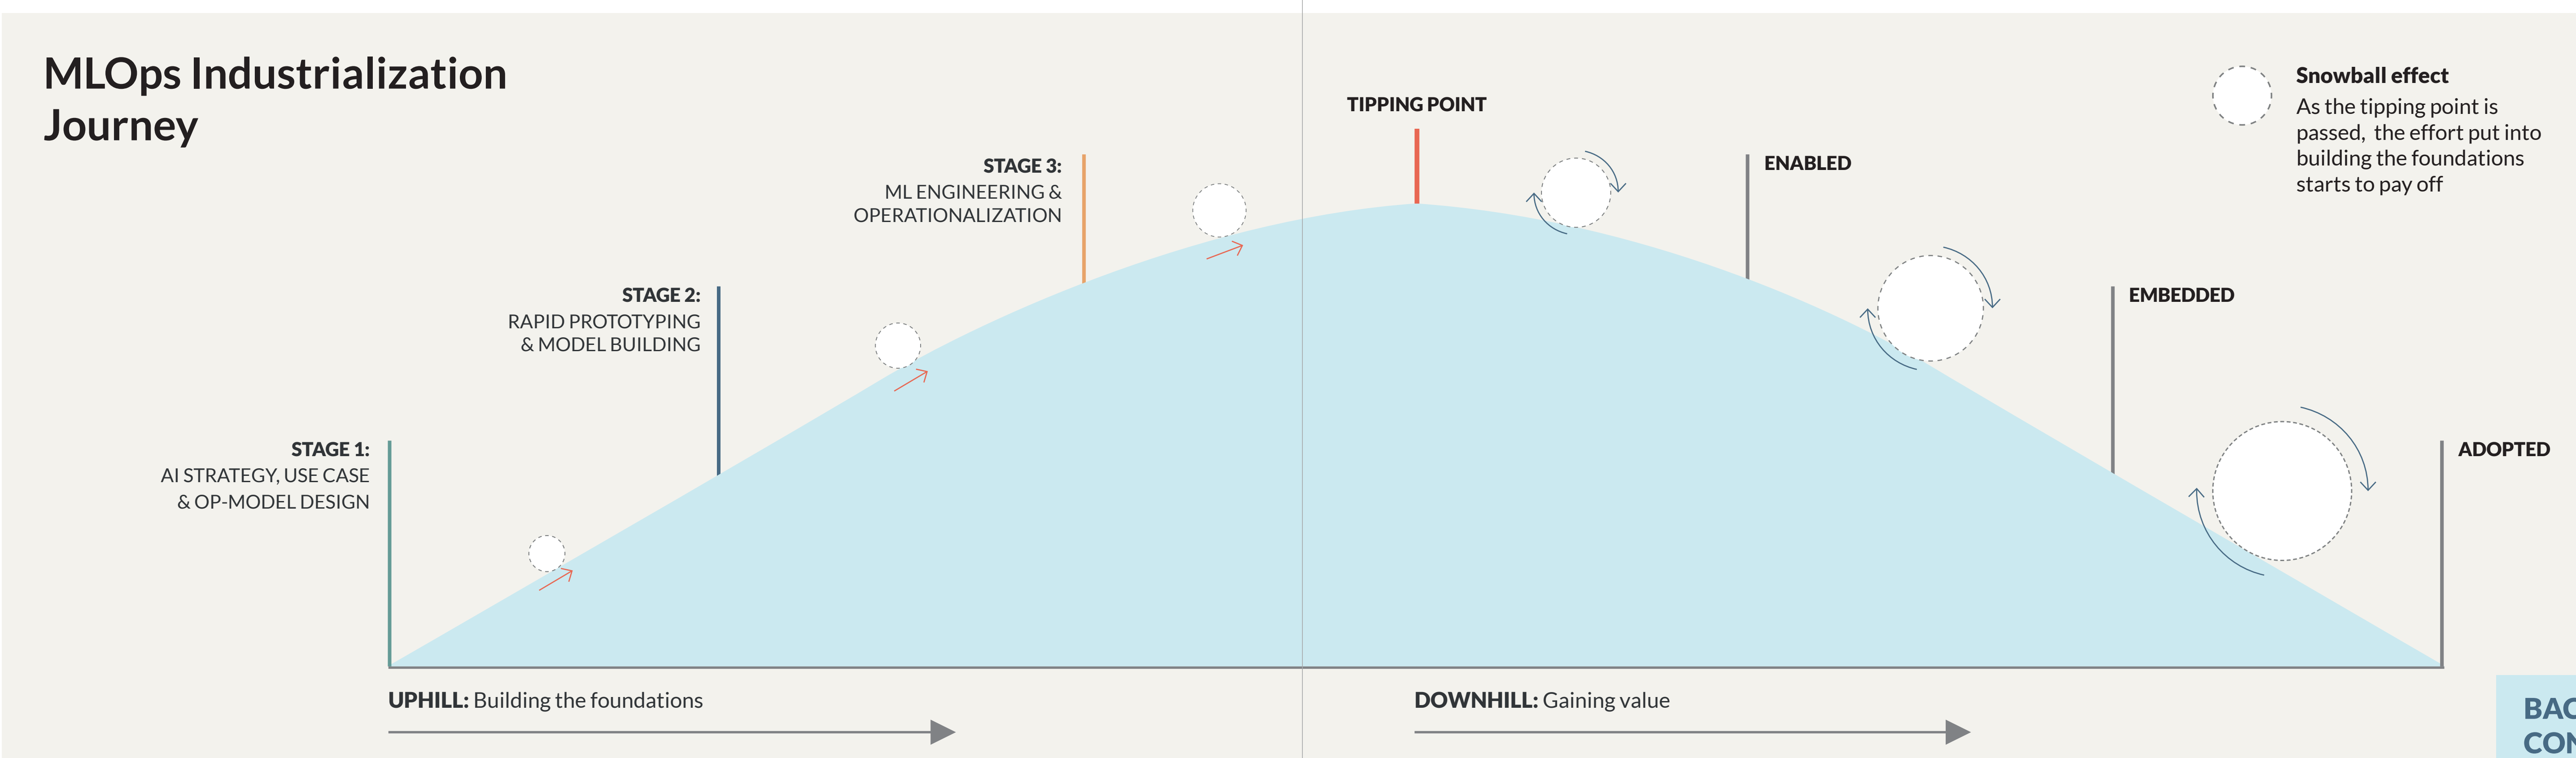

BACK TO CONTENTS

How to implement the framework

In our scenario, the organization has to deal with shifts in market conditions while managing operational overheads. What seems to be missing is a detailed definition of 'what good looks like' to capture nuances of the end-to-end approach.

Methodology here is key: MLOps can act as a blueprint for amalgamating tools, technology, processes, and roles with the right operating model for scaled delivery. The following framework can help decipher underlying issues and resolve the business problem:

STAGES

1 AI STRATEGY, USE CASE & OP-MODEL DEFINITION

KEY STEPS

- Defining business vision & helping with alignment
- Defining acceptance criteria
- Defining (or aligning to) ML monitoring & governance frameworks

KEY ROLES

Management Consultants | Domain Experts | Sponsors & C-suite leaders

2 RAPID PROTOTYPING & MODEL BUILDING

KEY STEPS

- Feature engineering & model build
- Model distillation & AutoML
- Performance, fairness & vulnerability testing

KEY ROLES

ML Engineers | Data Scientists

3 ML ENGINEERING & OPERATIONALIZATION

KEY STEPS

- Serving infrastructure design, continuous integration & development, infrastructure as code
- Model monitoring & logging
- Drift scenario & alert automation

KEY ROLES

ML Engineers | DevOps Engineers | Ops/IT Support | Data Scientists

SEQUENCING

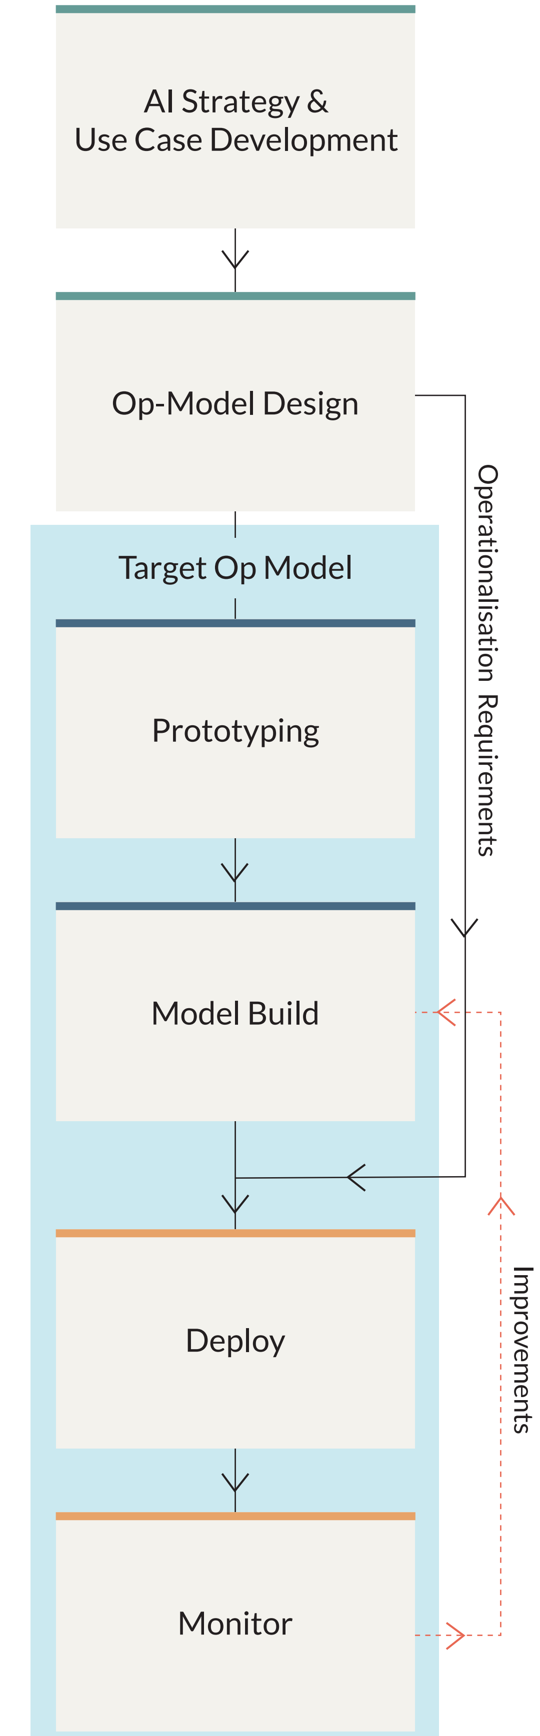

Unpacking the framework

In this section, we will unpack the framework, describe its **three key components**, and share some insights on how they can turn into actionable initiatives relevant to our scenario.

BACK TO CONTENTS

STAGE 1

AI strategy, use case, and operation model definition

First, the organization must review its current use of ML and determine how the business must adapt. In our case example, the CEO's vision of a data-driven organization requires interpretation and an impact assessment on business and technical capabilities.

A task force is appointed to bring together C-suite, marketing, data scientists, ML, and DevOps engineers, and IT, with the following initial findings:

- **Define business vision and help with alignment**: There is a need to define a clear link between the proposed marketing use cases and their near-term benefit, while reconciling conflicting priorities.
- **Defining acceptance criteria:** Formulate clearer model accuracy and operational success metrics for each ML model.
- **Model explainability:** Proactively plan for and avoid retrofitting explainability features.
- **Business outcomes and insights definition:** Fine-tuning of use cases is required to ensure model performance is more in line with internal community and end-user expectations.
- **Defining operationalization requirements:** Critically, functional and non-functional requirements need to be further refined to address historic effectiveness and efficiency challenges.
- **Defining (or aligning to) ML monitoring and governance frameworks:** Tighter integration to the existing data ecosystem is required, alongside a revision of the operating model to accommodate new use cases and ML platform technology.

BACK TO CONTENTS

STAGE 2

Rapid prototyping and model building

With the foundations set, ML specialists can now start building on the preliminary findings to add a flavor of technical implications and configure the systems and platforms, which will act as the delivery engine.

The organization quickly opts to consolidate previously fragmented tools and platforms to build a central rapid prototyping capability and decides on an 'off the shelf wherever possible' technology approach.

By building models and testing them with real data, the organization can disqualify a subset of options and focus on the ones that matter. This process allows for faster iterations and ultimately more bespoke and better-quality services.

> "By building models and testing them with real data, the organization can disqualify a subset of options and focus on the ones that matter.

The teams set the following direction:

- **Exploratory data analysis and pre-processing:** Limitations of upstream data platforms need to be addressed and more rigorous data cleansing and deduplication processes should be put in place.
- **Feature engineering, transformation, and selection:** Upskill the team on new cloud-based data platform capabilities that support more efficient data augmentation and feature engineering.
- **Model building and high-performance modelling:** Package and register the model making use of modern ML platform custom model support.
- **Model performance, fairness, and vulnerabilit**y: Optimize the model for performance in view of drafted non-functional requirements and deploy data privacy and balancing technology.

BACK TO CONTENTS

STAGE 3

ML engineering and operationalization

The broader engineering and IT operations teams should be called upon to validate the expected operating model changes and provide a view of supportability and infrastructure requirements.

Working collectively with the task force, the teams realize that they would be under-resourced without deploying modern cloud-based infrastructure, pockets of which only existed.

High-priority focus areas for an in-depth review are identified to help make a decision:

- **Serving infrastructure design:** More robust, real-time, and scalable integrations with downstream content generation business systems would be a critical success factor in improving user experience and meeting KPIs.
- **Model monitoring and logging:** New ways of working need to be launched to better align developers, operational teams, and data scientists as part of business as usual. More sophisticated monitoring tools are necessary and operational teams would need to be upskilled on cloud-native monitoring automation capabilities.
- **Drift scenario and alert automation:** Establish ML knowledge transfer mechanisms between data scientists and operational teams to develop common understanding of ML model limitations and intricacies. Develop an evaluation framework for specialist tools and procure technical solutions that integrate with the broader service management ecosystem.

BACK TO CONTENTS

BACK TO CONTENTS

How do I know if I need to adopt MLOps?

MLOps seeks to address persistent issues common to the organic adoption of machine learning, but it is critical to consider whether such frameworks are relevant to your organization by asking:

- Do only a small subset of developed models make it to production?
- Is there lack of synchronization between business, data science, and development teams?
- Is there a complex tools landscape with overlapping capabilities and features?
- Are there challenges in sourcing data in adequate form and currency?
- Do models consistently fail to positively impact the business in meaningful ways?
- Is there fragmented understanding of the organization's strengths, weaknesses, and opportunities?
- Do your teams have varying degrees of knowledge on this transformative technology and its potential limitations?

All of the above are strong indicators of possible benefits associated with a review of your MLOps strategy. If you were to start your MLOps journey tomorrow, the following tips can help you stay on course:

Focus on the big picture...

- **Far-reaching impact:** There are a plethora of considerations that need to be agreed and defined as part of a comprehensive MLOps framework. This includes a fit-for-purpose AI strategy and adequate engagement of C-suite leaders to inform use-case development and operating model requirements.
- **Focused innovation:** Business processes in transformation scope should be prioritized by time to value. This provides quick wins and will support your team's wider-scale rollouts.

But it may not be a sprint...

- **AI as a business imperative:** Avoid treating ML as a pure 'technology implementation' initiative. Adding structure and embedding methodology is key in creating true value and overcoming the 'hype.'
- **MLOps as an AI adoption journey:** Enterprises that recognize the full spectrum of MLOps considerations and the need to iteratively and incrementally release value to the business and its customers are more likely to reap the benefits of AI.

At Credera, we help our clients navigate their ML adoption journey by drawing on our collective experience across geographies, as well as the broader Omnicom network of agencies.

> "At Credera, we help our clients navigate their ML adoption journey by drawing on our collective experience across geographies.

Meet the authors

Georgios Sakellariou

PRINCIPAL

Adarsh Panda

DIRECTOR

Harry Tsangari

ANALYST

Here are some further resources:

- MLOps best practices: Part 1
- Six practices for successful AI/ML investments
- Understanding the benefits and dangers of large language models
- Benefits of AI: increase efficiency, enhance safety, and improve reliability
- Cookie-less future: How data-driven marketing will change
- Building a MLOps System for minimal maintenance

If you're interested in having a conversation about MLOps at Credera, please get in touch:

findoutmore@credera.com

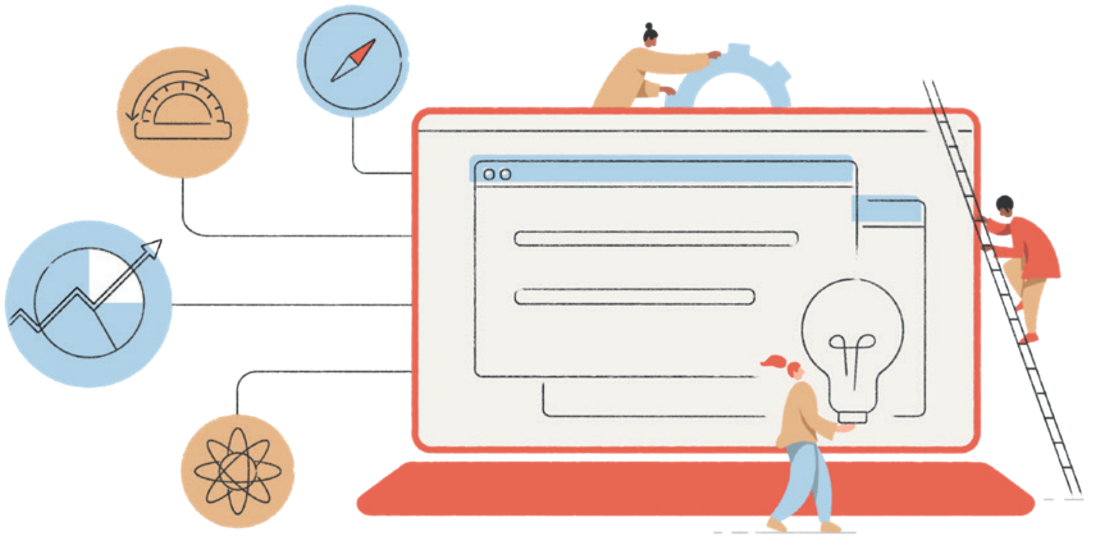

BACK TO CONTENTS

CREDERA

CREDERA.COM