

---

## 3 WAYS TO ACHIEVE MACHINE LEARNING EXCELLENCE



HYBRID  
CLOUD  
WORLD

## Table of Contents

Introduction	3
Three Challenges of Enterprise Machine Learning	3
Organizational and Political Issues	3
The Problem of Scale	3
Lack of Data Product/Market Fit	4
How Successful Businesses Enable ML Excellence at Scale	4
Conclusion	6

ML work is highly iterative; even well-defined criteria may need to shift throughout different phases of the project as data scientists and business owners alike explore what is possible, ethical, and profitable.

Pipeline & package management, monitoring, alerting, and other DevOps-like activities may not be the most exciting areas to think about, but failure in any of these spheres will ground even the most promising ML/AI efforts.

At Cloudera Fast Forward, we've spent a lot of time thinking and writing about technical breakthroughs and new research that enables Production Machine Learning (ML) and Artificial Intelligence (AI), as well as advising customers on when and where to invest, in order to accelerate their respective enterprise ML journeys. Through that work, we've learned that many organizations don't focus on addressing the underlying causes of the challenges they face when industrializing ML capabilities.

### Three Challenges of Enterprise Machine Learning

Truly excellent industrial-grade ML requires transformation in almost every part of an organization, and as a result, production ML hurdles are often *actually* organization-wide hurdles. Here are three things to keep in mind as your team builds its enterprise ML abilities:



#### 1. Organizational and Political Issues

Organizational and political issues are probably larger and more impactful to your ML/AI outcomes than you may think (or want to admit). Even though ML projects are different from traditional projects from a technical perspective (requiring different skills, platforms, software, and workflows), they should still be regarded as *projects* within the enterprise, which will require adequate budget, sponsorship, cross-functional alignment, and management like any other initiative.

Our research and experience have shown us that even the most committed and mature organizations can sometimes struggle with the people and process complexities germane to large-scale ML projects. Because ML/AI projects blend scientific methodology with product-centric outcomes, traditional corporate decision support tools are often ineffective. Determining the investment level, duration, and evaluation criteria necessary to support a research project is a very different undertaking than developing the same metrics for a SaaS application. Adding to the noise, ML work is highly iterative; even well-defined criteria may need to shift throughout different phases of the project as data scientists and business owners alike explore what is possible, ethical, and profitable. This can lead to confusion on when an ML project is actually "done" or still "in production." Predictable chaos often ensues.

When it comes to the production phase of ML development in particular, it's also easy for organizations to take a "this will be somebody else's problem" approach (and many do). Pipeline & package management, monitoring, alerting, and other DevOps-like activities may not be the most exciting areas to think about, but failure in any of these spheres will ground even the most promising ML/AI efforts.

In summary, although many organizations recognize that developing a production-grade enterprise ML/AI capability is an existential requirement for success, few plan for the (monetary and non-monetary) costs. Even fewer consider what long-term commitment actually means in real world terms.

## 2. The Problem of Scale

Applied machine learning at enterprise scale requires a particular union of cutting edge technology and enterprise expectations. That is, brand new (and at times) untested or poorly supported technology, tools, and systems must be rapidly deployed into enterprise application workflows, and perform at least as well as existing software that has been in place (at times) for years. This means ML applications must be highly responsive with deliberate SLAs in place, for example, as well as support robust integrations with other software.

What's more, enterprise problems require enterprise grade solutions, which are subject to ethics, governance, compliance, and security requirements that Proofs of Concept, prototypes, and experiments can sidestep. This leads to dependencies on other parts of the business, different technologies, and regulatory approvals.

The growing pains of scale add to the organizational and political costs of developing an excellent production ML capability. Many groups do not consider the total cost of ownership for ML applications.

## 3. Lack of Data Product/Market Fit

Not all problems have the same solutions; similar problems across multiple verticals will potentially have multiple solutions. The right solution for your business problem will have several factors, including data, skills, conditions, and time:



- **Right data, right time:** it seems intuitive that having access to quality data would be a fundamental requirement of sound data science, but unfortunately a lot of well-meaning groups don't consider this element until too late. Having the right data means that your organization has both instrumented and collected data properly over time **and** that the right team members can adequately access, understand, and analyze this data in the context of building an ML capability. This can be a massive architectural undertaking for some use cases!
- **Right skills, right time:** At Cloudera Fast Forward, we've made it no secret that we believe excellence in Machine Learning requires a team-oriented approach. What this means in practice is that a diverse range of skills must be applied throughout the lifecycle of the project. For example, in many enterprise ML projects large cross-functional teams including, data stewards, business analysts, data and software engineers as well as ML Operations practitioners may be required for success, no matter the quality of the Data Scientists involved. Because no one individual has all of the skills necessary to build, scale, and maintain an enterprise ML application, multiple teams will need to support the project **cross-functionally**. This, of course, comes with risk. The success or failure of the project hinges on the proper prioritization and utilization of skills.
- **Right conditions, right time:** A supportive environment for innovation/ideation, research, development, deployment, and adoption of ML applications (at scale) is also very important. Though many organizations realize the importance of these factors, and cultivate wonderful environments at the beginning of projects, few are willing to invest what it takes to sustain these conditions over time. ML projects can often last several months or even years, and because they are largely experiment-driven, inherently involve at least some failure.

Experiment with not just the technology but the operational and process considerations that it takes to manage an application at scale.

### How Successful Businesses Enable ML Excellence at Scale

The good news is that as with all good science, experiments often lead to results (positive and negative). As such, we've captured the most effective solutions and mitigations to the above challenges that we've seen:

- To address organizational and political headwinds, start with an ML/AI strategy, then follow with small—but important—projects. Executing such projects serves a dual purpose. Successful deployments of even small ML applications will prove to decision-makers within an organization that ML is a valuable, strategic, and differentiating asset. Additionally, going through the motions of experimentation, design, development, deployment and maintenance builds both experience and confidence in technology, processes and teams. Experiment with not just the technology but the operational and process considerations that it takes to manage an application at scale. This approach addresses both the motivational and technical challenges of production ML, but allows for rapid iterations while the stakes are still manageable.
- Perform frequent architecture reviews and treat everyone as “customers” in the conversation.
  - *IT* is one of the ML/AI team’s customers. Data Scientists are building an application which has to work in the managed ecosystem provided by IT. If an ML application isn’t designed to work in this environment, or is challenging to work with, it will be difficult to convince IT to properly deploy, scale and maintain the application, once the ML/AI team is no longer involved.
  - *ML/AI Team(s)* are conversely the IT team’s customer. ML practitioners need a stable, flexible and powerful development environment, delivered in an experience that doesn’t negatively impact the research and ML application development process.
  - *Lines of Business Owners/Data Consumers* are the customer of both IT and the ML/AI team(s). It’s not enough to provide an interface and “good news;” LoB owners and consumers need to deeply understand the product side of any production ML application as well as the data products that support them.
  - *End Users* are the traditional customers of most applications, and ultimately are everyone’s responsibility. That said, core principles that inform world-class customer engagement, customer success, and application performance management serve as a useful north-start for how to support end users.
- Take a [portfolio](#) approach to selecting which data products to pursue. Not all enterprise use cases are created equal, and not every great idea is actually possible to execute in practice. To address the challenges of data product development mentioned above, we find that successful organizations consider these approaches.
  - **Data:** Enterprise Data Management excellence is a critical organizational competency that all enterprises should develop and sustain. We like to think of it as the base of the ML/AI “hierarchy of needs” shown below. Many organizations already have a strong technology and skills foundation in data management. It’s important that leaders realize that this capability is central to success in enabling ML outcomes, and continue to invest in it.

When considering the total cost of ownership for an ML capability, leaders and development teams alike should estimate data ingest, storage, and access requirements, and revise those estimates throughout the project lifecycle.

#### 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](https://cloudera.com)

#### Connect with Cloudera

About Cloudera:

[cloudera.com/more/about.html](https://cloudera.com/more/about.html)

Read our VISION blog:

[vision.cloudera.com](https://vision.cloudera.com)

and Engineering blog:

[blog.cloudera.com](https://blog.cloudera.com)

Follow us on Twitter:

[twitter.com/cloudera](https://twitter.com/cloudera)

Visit us on Facebook:

[facebook.com/cloudera](https://facebook.com/cloudera)

See us on YouTube:

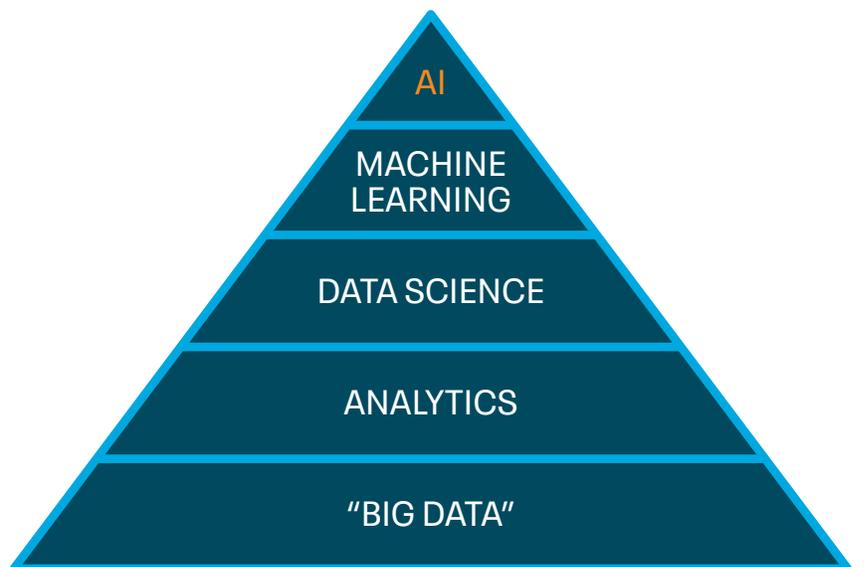
[youtube.com/user/clouderahadoop](https://youtube.com/user/clouderahadoop)

Join the Cloudera Community:

[community.cloudera.com](https://community.cloudera.com)

Read about our customers' successes:

[cloudera.com/more/customers.html](https://cloudera.com/more/customers.html)



- **Skills:** If your data organization is just getting started, your first hires might need to be a bit more flexible and versatile (data scientists will need big data and some full stack application engineering skills). This means that you might have to invest more in early hires, but that your organization will likely get by with fewer of them. As you continue along the ML/AI industrialization journey, more specialized functions and skill sets will likely emerge (i.e., distinct research, applied ML, Data/ML engineering, and MLOps teams).
- **Conditions:** Sponsorship, investment, and strategic alignment should come from the top—and domain experience, innovation, and enterprise management from the bottom. It takes deliberate organizational patience to enable long-term ML/AI success in the enterprise.

When considering the total cost of ownership for an ML capability, leaders and development teams alike should estimate data ingest, storage, and access requirements, and revise those estimates throughout the project lifecycle.

### Conclusion

It's true: excellence in production ML is difficult. Adding in the requirements of enterprise scale leads to even more complexity, but success is not impossible. Those who recognize the ML/AI industrialization journey for what it is—an organization-wide transformation—will not only enjoy positive initial outcomes but will also be set up well for a long-term competitive advantage.

#### Sources

<sup>1</sup> Image originally created by Hilary Mason