Machine Learning (ML) Operations: Overcoming Challenges in Migration to the Cloud

The crucial role of steering teams and templates

Machine Learning (ML) Operations: Overcoming Challenges in Migration to the Cloud

The crucial role of steering teams and templates

By **Amit Sharma**, Senior Director - Data Science, Iris Software

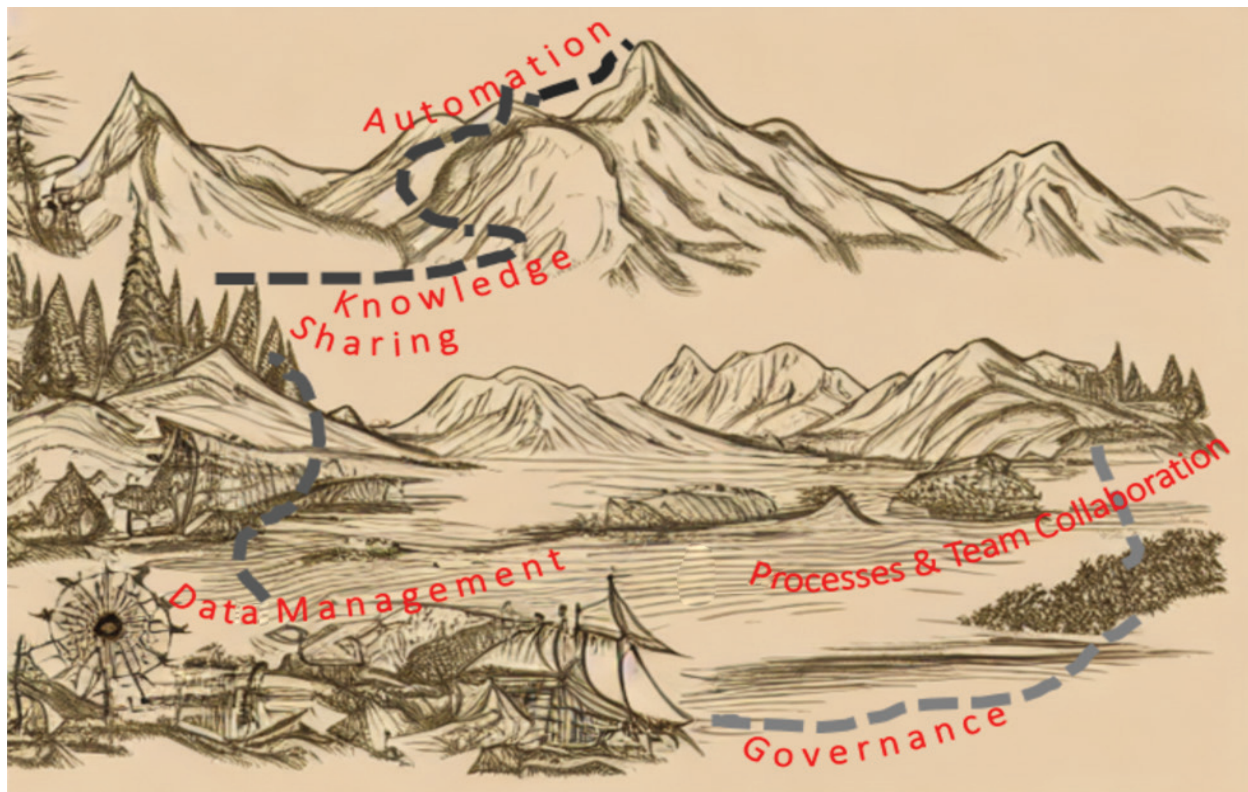"Each part of the system contributes to the strength of the whole."

Figure – ML Journey to Cloud

The promise of artificial intelligence (AI) and machine learning (ML) to deliver significant competitive advantages and contribute to the growth of an organization is now well established. Both big and small enterprises are adopting these to drive their strategic business goals. They are increasingly using these technologies to monetize data, enhance revenue/P&L, improve customer satisfaction, automate processes, improve efficiency, and make better decisions across a wide range of business functions, such as new product/service development, pricing, risk management, marketing, sales, customer service, to name a few.

As ML is adopted more widely across the business world, an increasing number of enterprises are considering cloud as a way to scale their machine learning efforts. However, many organizations are realizing that this has its own challenges as operational and support requirements increase rapidly in the process.

Key Challenges

The following are key challenges that organizations face in effectively supporting the adoption of emerging AI/ML platforms, especially in the public cloud:

- There are multiple teams (Cloud Engineering, Cloud Ops, Infosec, Cybersecurity, LOB Technology, Enterprise Data & Analytics Teams, Design & Architecture Governance teams, etc.) involved in governing, enabling and supporting various aspects of utilizing cloud-based services with their bespoke set of processes, standards, requisitioning & service management platforms. Engaging with these teams, understanding their standards, and navigating their processes can be challenging.

- To effectively support data science and modeling on the public cloud, establishment of pipelines to move data there and enable secure access is critical. However, cybersecurity/infosec requirements around data on the public cloud are still evolving, often being too demanding to deliver acceptable performance.

- While cloud-based modeling platforms support their own version of the modeling lifecycle, the actual firm-specific lifecycle could be more complex. This is because model lifecycle activities such as validation, production-deployment and ongoing monitoring may be spread across infrastructure and applications that are on-prem and cloud, as well as a mix of other cloud-based services and internal or third-party platforms.

- The above result in significant complexity that requires a lot of manual effort, impacting the overall agility and efficiency in adopting emerging platforms.

- Change management for modeling teams becomes significantly harder to implement as various aspects of the above are not productive using cloud-based tools.

Overall, currently, adopting AI in an enterprise holds great potential and promise, but it comes with a number of challenges that must be navigated to realize all the benefits.

Our Learnings and Best Practices

Machine Learning and Operations, or MLOps, is significantly different from traditional software development practices. A different approach is required for development, deployment and maintenance of ML models. This shift can be difficult for the various teams involved in governance, enablement and support of public cloud-based capabilities. They are used to more traditional approaches and would be required to change culture and mindset.

The following are some of our learnings and recommendations for successful cloud migration journeys towards machine learning at scale based on our partnership with clients in this space:

- With many different teams, processes, and standards to navigate, it is important to set up a core MLOps Enablement team. It should have a good understanding of the teams/processes involved and of the differences between a modeling/data-science lifecycle vs. traditional software development lifecycles to be able to bridge the perspective gap where necessary.

- It is also important to establish a DataOps team to ensure that the data required by modelling teams can be pipelined to cloud ahead of modelers being onboarded to cloud-based platforms. This will pre-empt many issues that may arise due to modelers using data from non-enterprise-standard sources and running into problems during validation/vetting and production deployment.

- Since different models may require different resources (by type or by scale, e.g., CPU vs. GPU, single VM vs. clusters of VMs), it would be helpful to define templates to simplify the onboarding of teams, requisitioning, and provisioning resources required.

- In addition to modeling resources, enabling modeling teams on cloud often involves setting up other resources such as Active Directory/LDAP groups, version control repositories, automation pipelines, etc. Simplifying these through templates and automating the requisition, review/approval, and provisioning processes would significantly reduce the typical churn as modeling teams try to adopt cloud-based tools.

- Often, modeling teams need to learn and use different programming languages and modern frame works on the cloud (e.g., SAS vs. Python/PySpark). Establishing model templates to simplify various programming steps to acquire data, engineer features, run the model(s), save interim and final outputs would ease adoption.

Conclusion

The power of AI/ML models and platforms and the potential for deriving significant value from them is accelerating very fast. It is quickly becoming a critical imperative for firms to adopt these rapidly to ensure they remain competitive. However, successfully adopting these at scale would require the use of public cloud technologies. Public cloud adoption in many industries is still in the early stages, with many internal groups to work with and evolving processes and standards. Because model development lifecycles are signficantly different from software develop-ment lifecycles, this journey is even more challenging. However, by establishing clear ownership for MLOps & DataOps and simplifying adoption through the use of templates and automa-tion, firms can overcome these challenges and scale AI/ML on cloud while ensuring overall agility and cost-effectiveness.

For further reading:

- [MLOps: Continuous delivery and automation pipelines in machine learning](#)
- [Amazon SageMaker for MLOps](#)
- [Machine Learning: The High Interest Credit Card of Technical Debt](#)

About the Author

This perspective paper is written by **Amit Sharma**, Senior Director - Data Science at Iris Software. Amit is a seasoned ML professional and has contributed to numerous projects ranging from large-scale personalization to cutting-edge natural language processing. With extensive experience in on-premise and cloud-based solutions, Amit holds multiple patents in machine learning and has over twenty years of software engineering experience. His expertise extends to the critical area of operationalizing models.

About Iris Software

Iris Software has been a trusted software engineering partner to Fortune 500 companies for over three decades. We help clients realize the full potential of technology-enabled transformation by bringing together a unique blend of domain knowledge, best-of-breed technologies, and experience executing essential and critical application development engagements. With over 4,300 skilled professionals, Iris builds mission-critical technology solutions across banking, financial services, insurance, life sciences, and manufacturing. Iris offers comprehensive services and solutions across data and analytics competencies of Data Engineering, Enterprise Analytics, Data Science and Data Governance.

Iris™

When IT matters most.

Thank you for reading this perspective paper.

Click here to read about the milestones of our 30-year journey.

Learn more about our Data & Analytics capabilities at:
https://www.irissoftware.com/services/data-analytics