Booz | Allen | Hamilton®

ADOPTING MLOPS

Our team adopted an MLOps solution within the Federal Government to enable Machine Learning for mission critical use cases.

MARCH 2021

OPPORTUNITY

Challenge	“Only 22% of companies using machine learning have successfully deployed a model” - deeplearning.ai
Opportunity	“Three-quarters of enterprises will operationalize AI by 2024” - Gartner

MLOPS AND DEVOPS

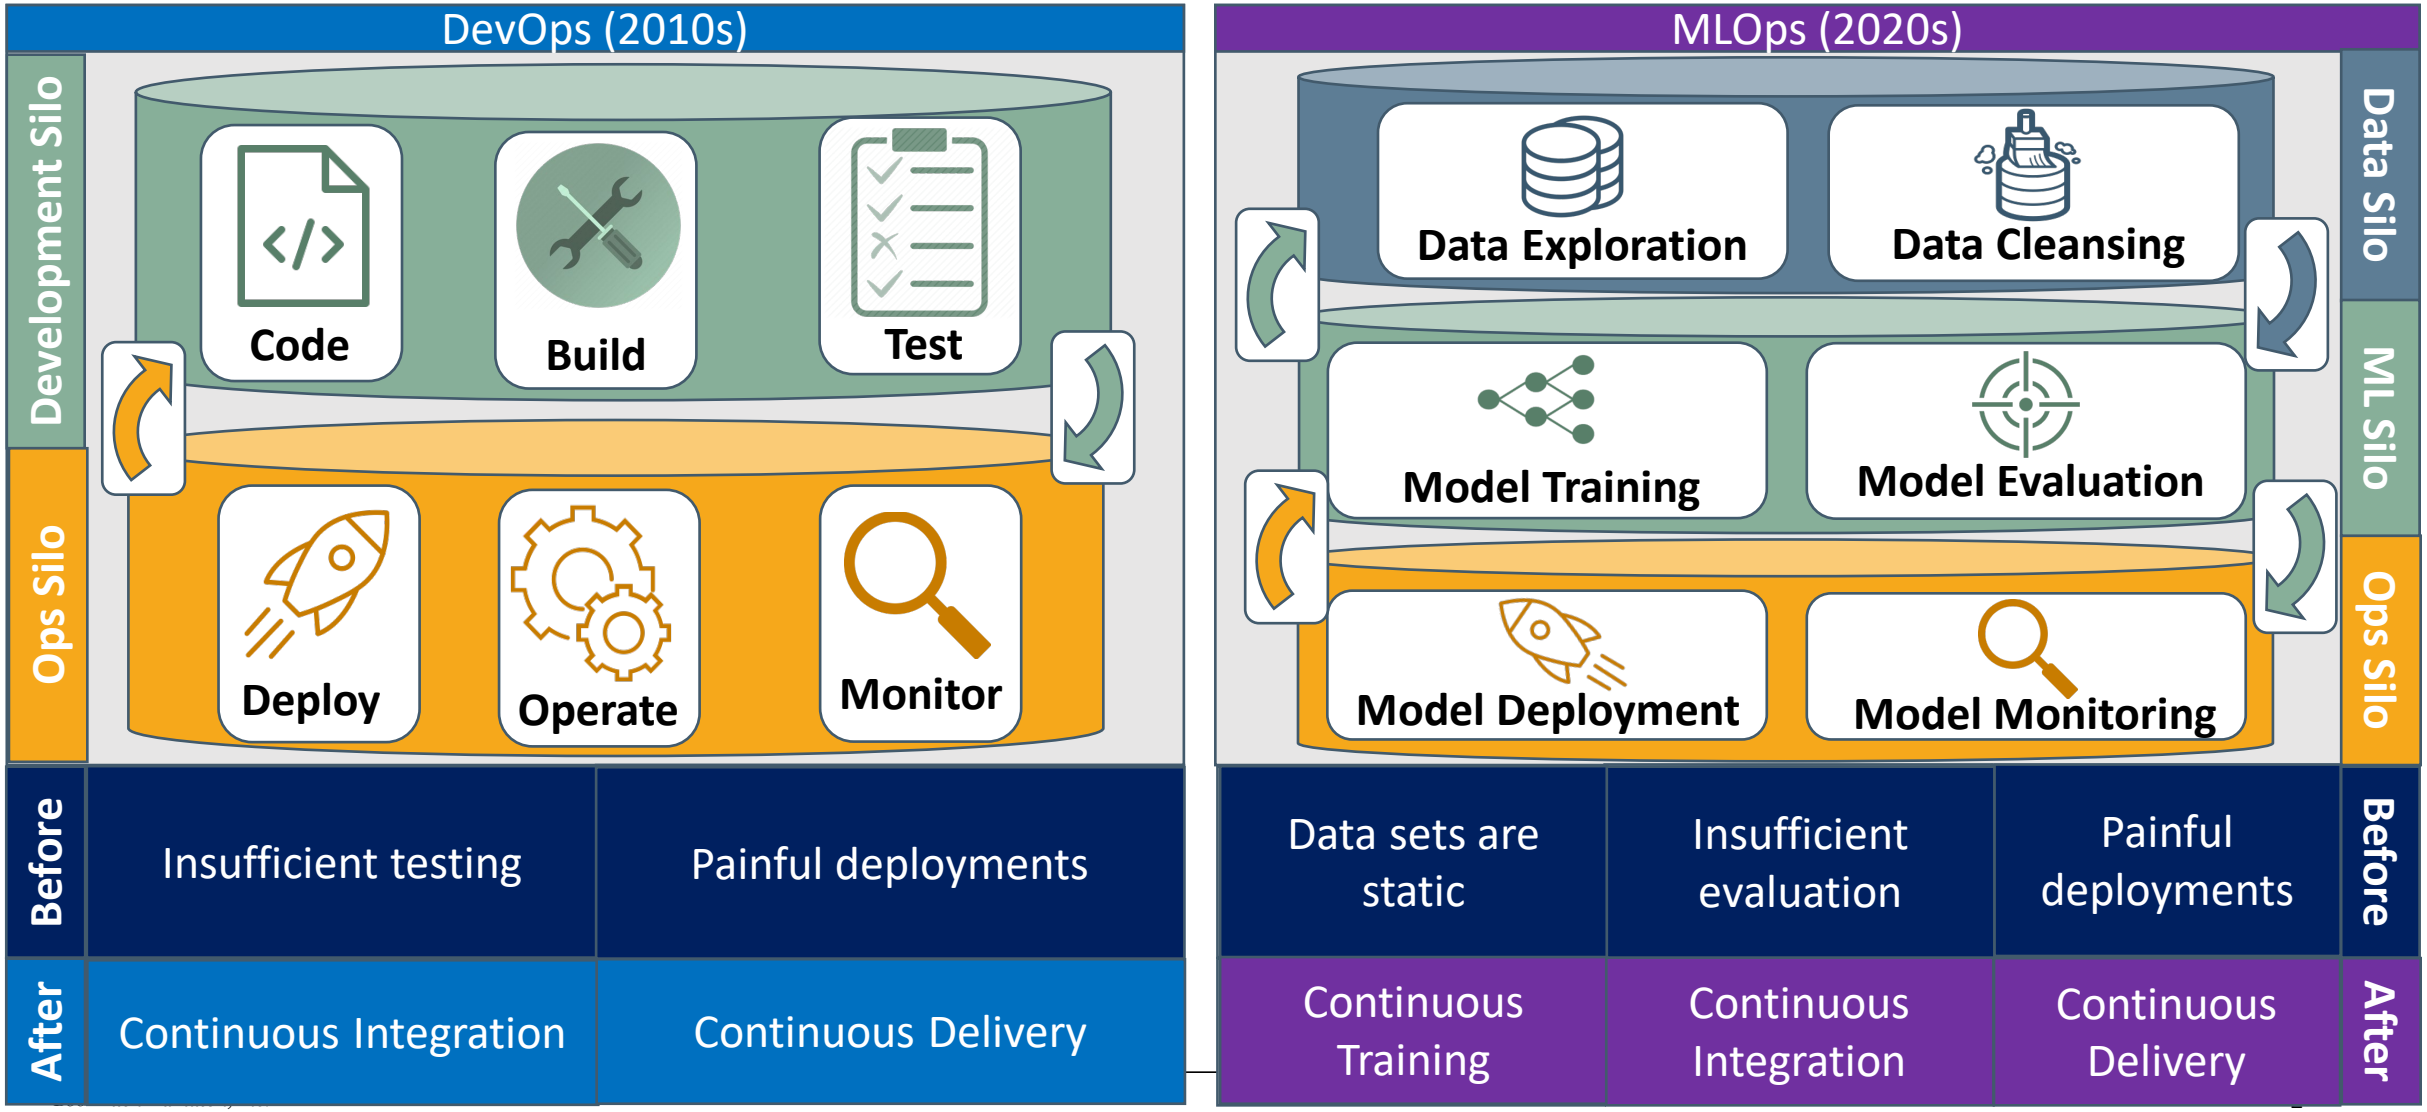

OUR MLOPS SOLUTION

Data Platform

- Automates data ingest and ETL
- Automates data schema validation
- Enables Continuous Training

Modeling Environment

- State-of-the-art ML development environment
- Ensures experiments are reproducible
- Enables Continuous Integration

Deployment Framework

- Automates deployment
- Automates monitoring
- Enables Continuous Delivery

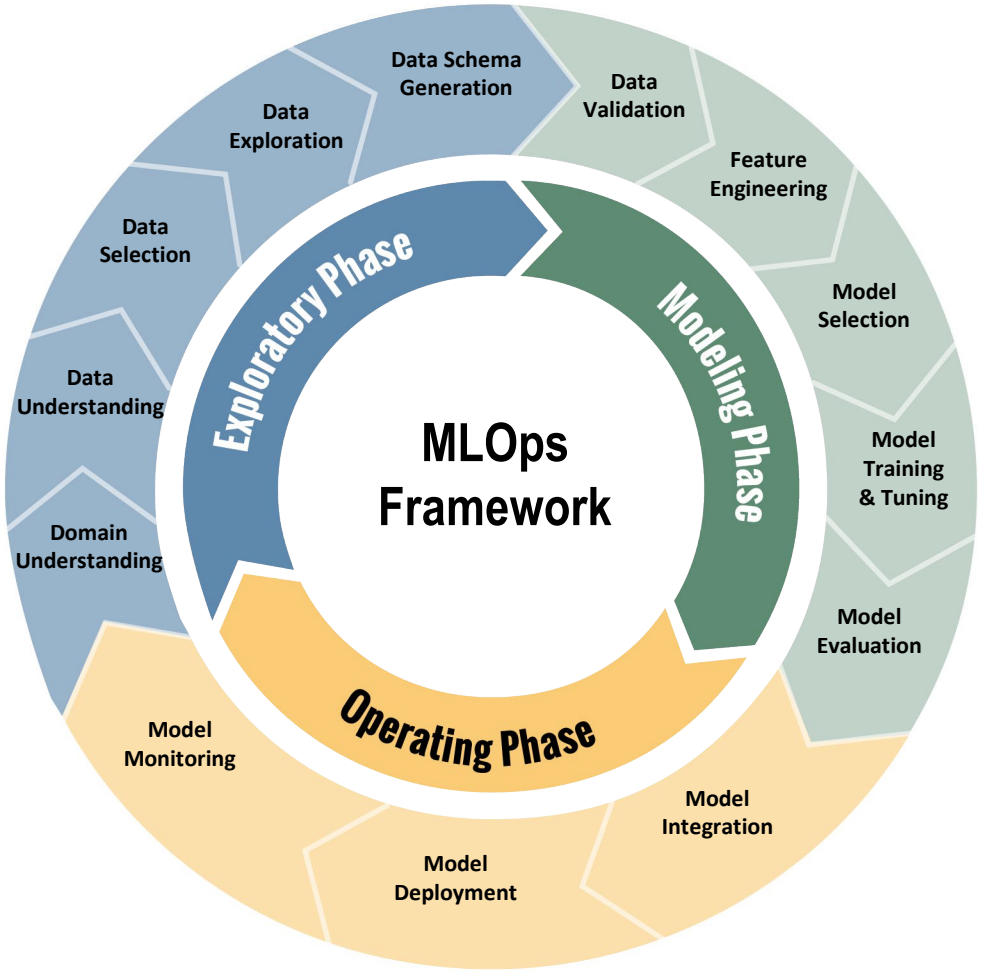

OUR MLOPS SOLUTION

Data Platform

Data Lake: Store data via AWS S3.

Stream: Stream data via Kafka.

Scale: Increase data volume 10x or more.

Data Warehouse: Query TBs of data with our internal solution.

Discover: Discover data through our internal catalog solution.

Modeling Environment

Machine Learning: Use SageMaker.

Develop: Develop software on a virtual desktop.

HPO: SigOpt.

Deep Learning: Run DL on GPUs.

KubeFlow: ML Pipelines.

Deployment Framework

Host: Kubernetes hosting with Rancher.

CD: Continuous Delivery with ArgoCD.

Model Marketplace: Host models with Modzy.

Monitor: Monitor with Grafana.

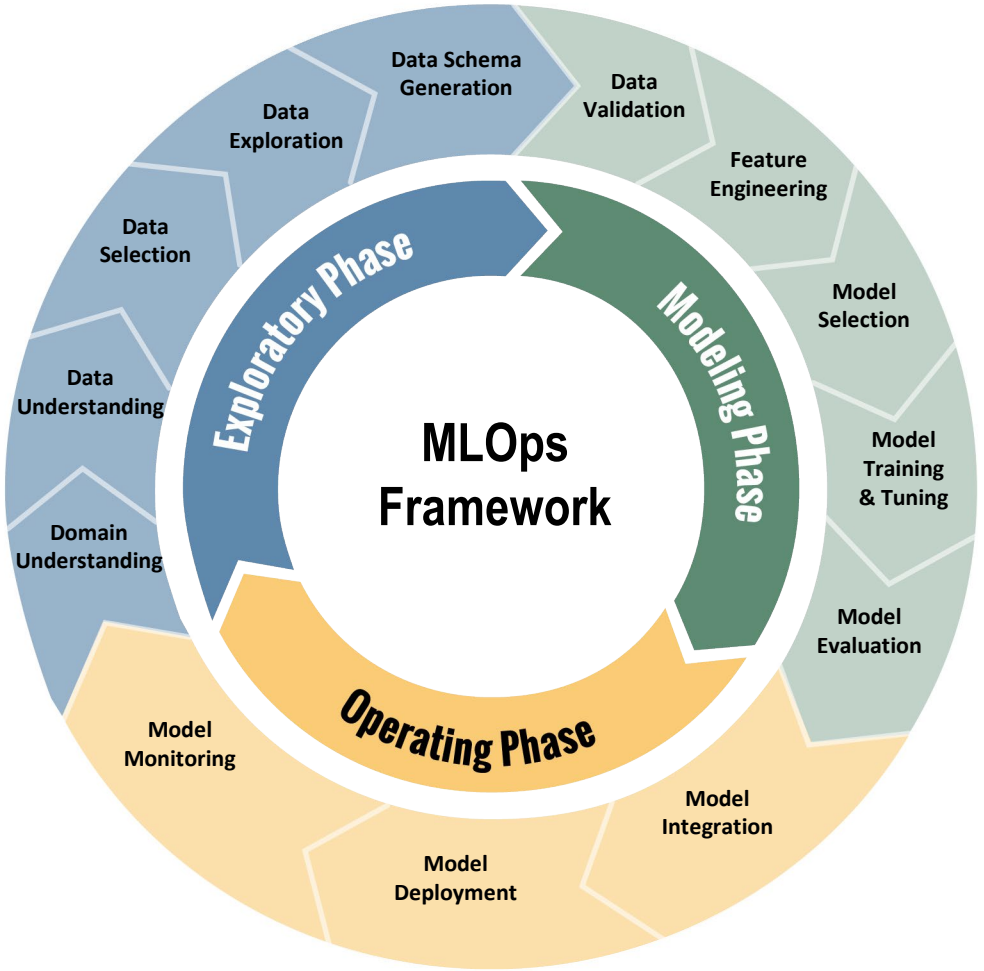

TAKEAWAYS

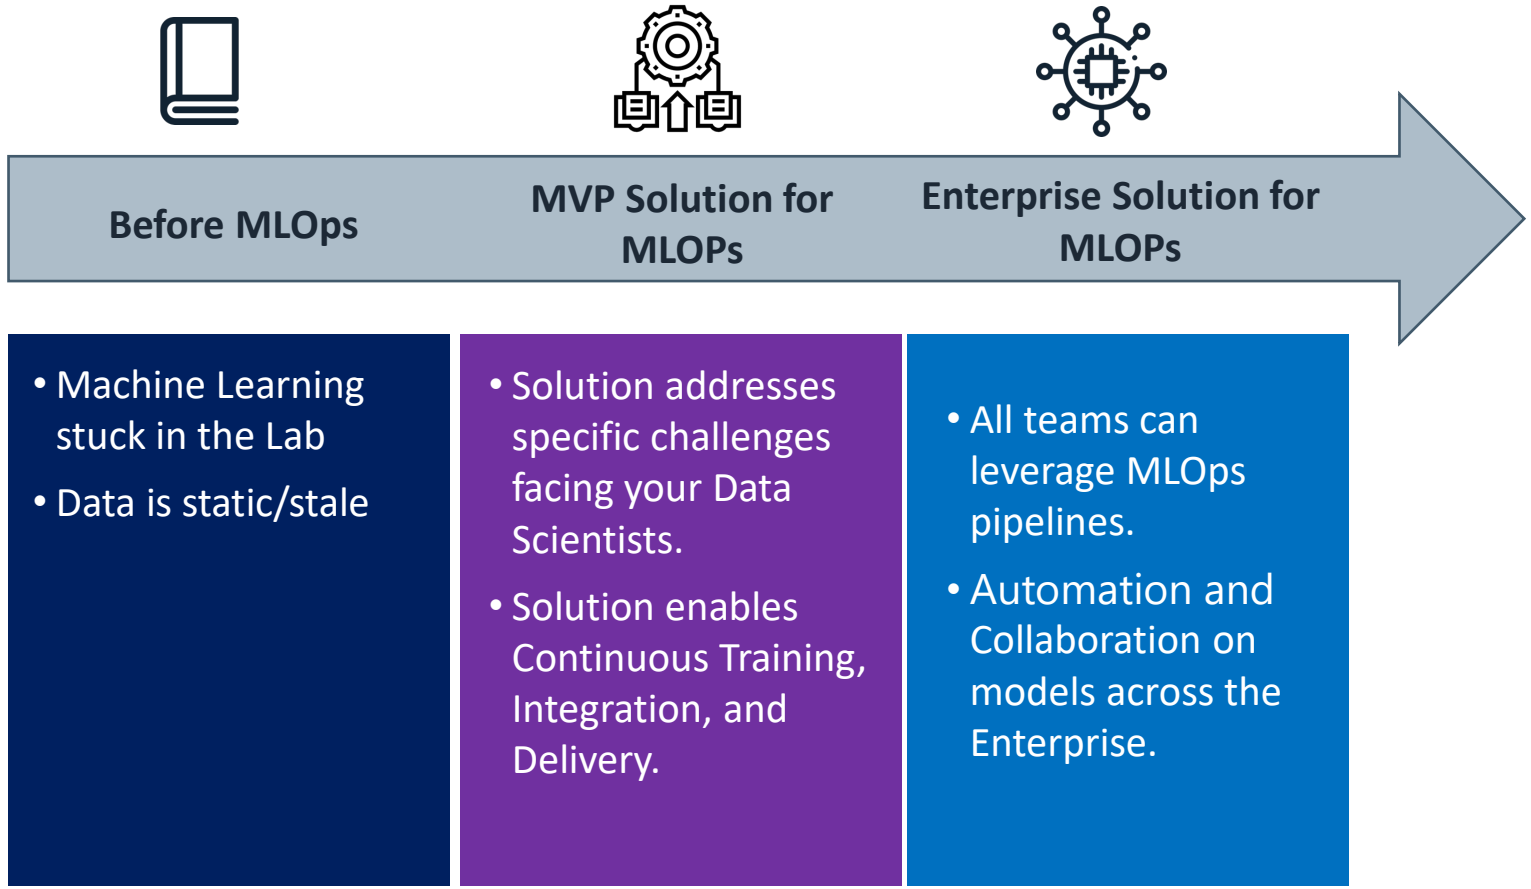
Before MLOps
MVP Solution for MLOPs
Enterprise Solution for MLOPs
• Machine Learning stuck in the Lab
• Data is static/stale
• Solution addresses specific challenges facing your Data Scientists.
• Solution enables Continuous Training, Integration, and Delivery.
• All teams can leverage MLOps pipelines.
• Automation and Collaboration on models across the Enterprise.

BACK-UP

ABSTRACT

- Within the last decade, the concept of DevOps has become an integral part of the software development lifecycle, a necessary shift in practices to provide continuous integration and delivery of software into operations. Industry has continued to uphold DevOps practices for any system that wants to improve their mean time to recovery and lower failure rates of new releases. Machine Learning operations demand the same level of commitment and enthusiasm as DevOps, if not more. Machine Learning models are only as good as the data they are trained on. In most cases, operational data changes over time leading to increasing Concept Drift and decreasing model performance. The IC cannot afford for critical operational ML models to break, and yet operational ML models are often built as monolithic scripts not set up for monitoring, retraining, or replacement. MLOps practices—like DevOps—can help lower failure rates and reduce mean time to recovery for ML operations by adding capabilities for model testing and quality assurance, deployment variants, monitoring, model retraining, hardware and software optimization, and versioning and reproducibility. In order to efficiently and effectively adapt critical Machine Learning models and applications across Exploratory, Modeling and Operational phases, the Intelligence Community must adopt MLOps best practices. This session will cover how our teams have developed a data analytics platform that enables MLOps and associated best MLOps practices to support mission critical operations.

REFERENCES

- NVIDIA: https://blogs.nvidia.com/blog/2020/09/03/what-is-mlops/
- Towards Data Science: https://towardsdatascience.com/ml-ops-machine-learning-as-an-engineering-discipline-b86ca4874a3f
- DataRobot: https://www.datarobot.com/platform/mlops/
- Google: https://cloud.google.com/solutions/machine-learning/mlops-continuous-delivery-and-automation-pipelines-in-machine-learning
- MLOps: https://ml-ops.org/content/mlops-principles
- Point72 Model Driven Business: https://www.youtube.com/watch?v=XtT0za9_sT4&ab_channel=MITClubofNorthernCalifornia
- Hidden Tech Debt in ML Systems: https://papers.nips.cc/paper/2015/file/86df7dcfd896fcaf2674f757a2463eba-Paper.pdf
- Uber MLOPs: https://eng.uber.com/michelangelo-machine-learning-platform/
- A16z: https://a16z.com/2020/10/15/the-emerging-architectures-for-modern-data-infrastructure/
- Tools: https://huyenchip.com/2020/12/30/mlops-v2.html
- Gartner: https://www.ciodive.com/news/gartner-AI-enterprise-projections-2024/580240/
- DeepLearning.ai: https://blog.deeplearning.ai/blog/the-batch-companies-slipping-on-ai-goals-self-training-for-better-vision-muppets-and-models-china-vs-us-only-the-best-examples-proliferating-patents