MLOps – Staying on-Track by Providing Value

05/15/2019

Declaration of Interdependence

RAILINC

We **increase return on investment**

by making *continuous flow of value* our focus.

The Pillars of Modern SDLC

Agile Methodology

Deals with Uncertainty

- Strives to create a desirable product and business outcome
- Accepts that we might change our minds once we see it
- Embraces that we will fail at many things

DevOps

Deals with Complexity

- Helps iterate faster
- Delivers higher quality
- Removes siloes because we can’t possibly understand all implications, nor should we

DevOps

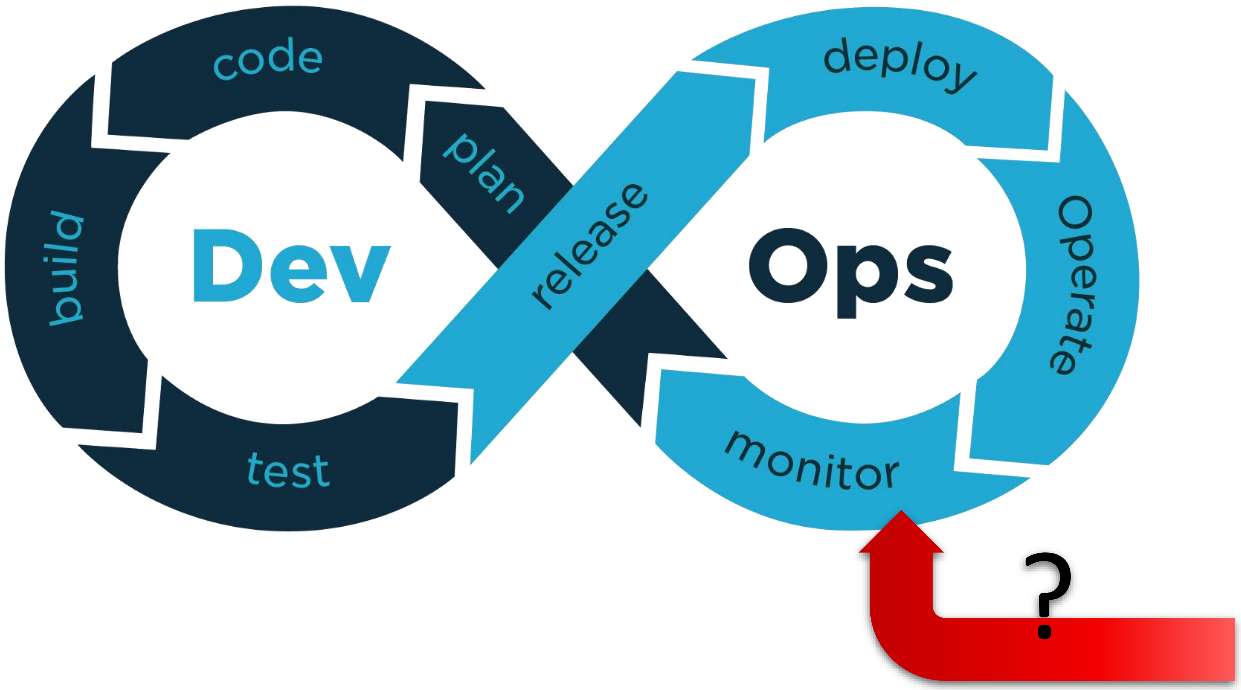

DevOps is a set of practices that combines software development and IT operations to shorten the SDLC while delivering features, fixes, and updates frequently

in close alignment with business objectives.

Operationalizing Machine Learning(ML) Models @Railinc

Railinc built on the methodology and principles of Agile and DevOps

MLOps by Railinc's definition

- Closes the gap on measuring business objectives
- Extends the DevOps principles by incorporating value-driven measurements throughout the development and operations cycles
- Borrows from the research definition of operationalization

Definition: Operationalization is the process of strictly defining variables into measurable factors. The process defines fuzzy concepts (like value) and allows them to be measured, empirically and quantitatively.

Defining Value

Containers are ramped on a train

Where are they headed?

Train is handled by many carriers

How accurate is their reporting?

Stations and terminals are operated by different entities

What is the dwell time today, tomorrow?

Container is de-ramped and then put on a truck

In which order?

When should Joe show up with his truck?

Predicting Where, When and What Event Comes Next

- The Movement data is highly complex and noisy
- We had the idea of mimicking what sports analysts do for playoff predictions
- We had to keep each Origin/Destination pair separate
- We ended up predicting individual event transitions

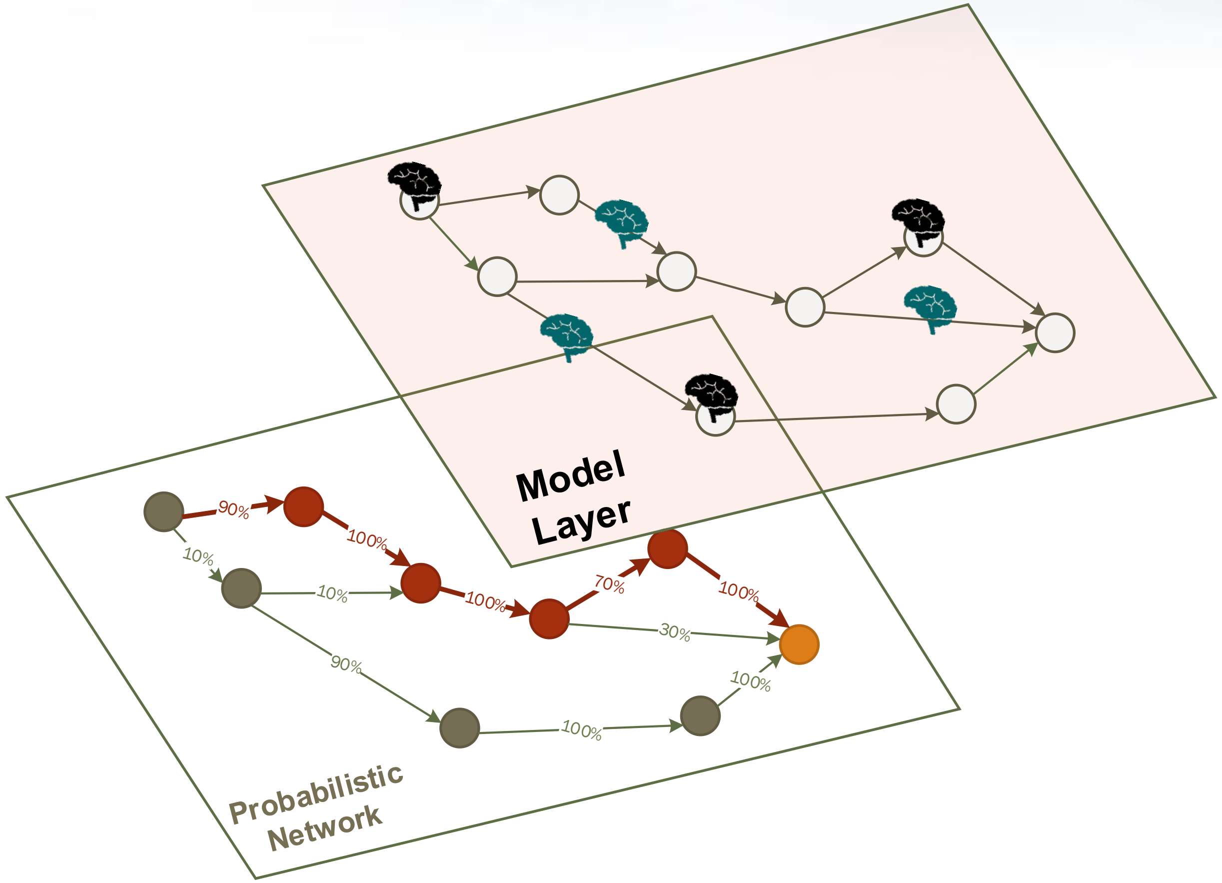

Facing problems of scale and runtime complexity

Data Science: Approximately 10K origin and destination pairs

- How to run the error analysis on thousands of pairs and millions of trips?
- How to understand the impact of model changes on the business value?
- When to retrain?

Infrastructure: Initial models size was 200GB

- Thousands of models, how to version and keep track of release changes?
- Too big to easily run in a container

Monitoring: Data issue or model decay

- How to identify data quality issues at runtime?

MLOps: Start with the Customer Value

- Customer driven measurements: start with their purpose

- Measured throughout the product lifecycle
    - Training time
    - Runtime
    - What the customer sees

- Driving new model releases

- Driving the release of functionality to customers
    - Roll out first the high-performant lanes
    - Added specific measurements for the targeted customers

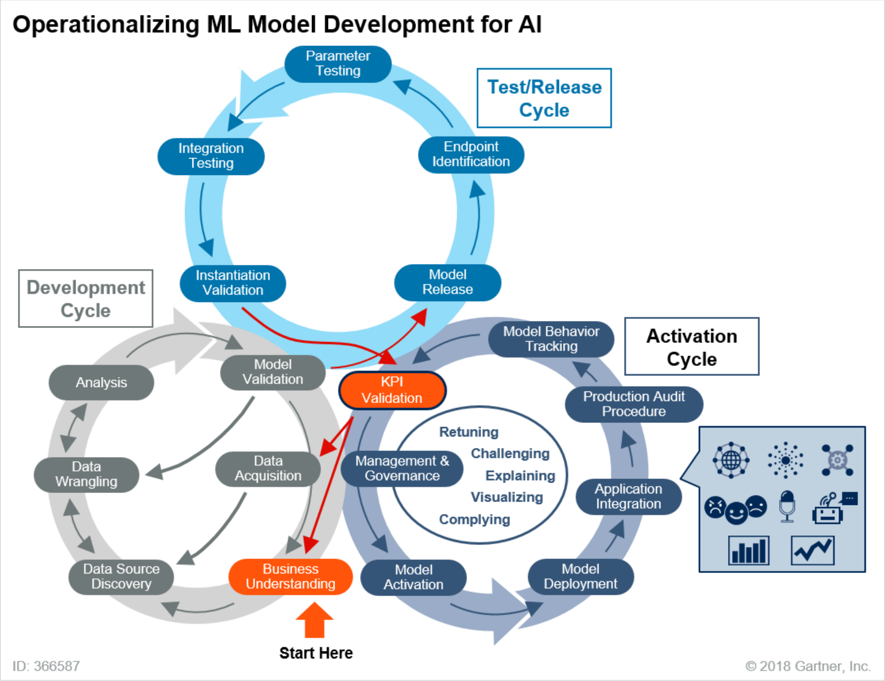

MLOps: Exposing Issues

- Using information at training time that would have not been available at runtime
- Exposing runtime data quality issues
  - Bad data
  - Unusual gaps in the data
- Pinpointing operational variability
- Exposing different understandings or implementations of the metrics definition

Other strategies that made us successful

- Cross-functional collaborative team
  - “If one of us fails, we all fail.”
- Product and outcome-driven
  - Kanban over Scrum
  - Focusing on delivery not planning and control (adaptation vs. anticipation)
  - Festina lente (MVP and Quality)
- Managing uncertainty
  - Parallel experimentation approach
- Creating internal programs to increase critical skillsets

Data Fitness and ML "Magic"

- Many of the data sources have been collected for a different purpose
- Value principles apply also in understanding what products are possible given the available data

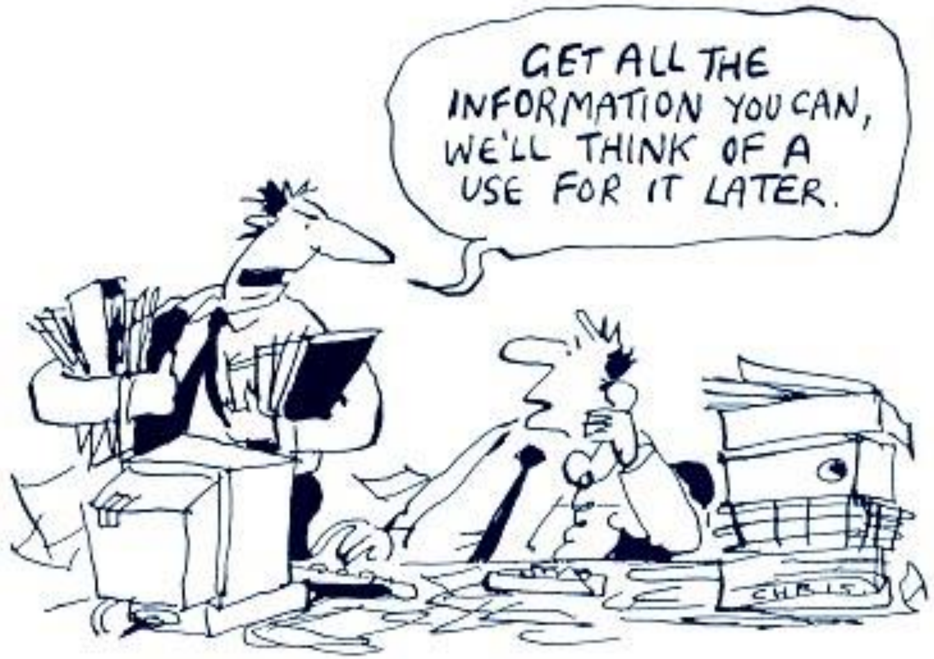

What we learned

- Many times success is not about the best model
- The nature of operationalization of machine learning lends itself to a value-driven, adaptation approach
  - Uncertainty of goals
  - Many unknowns about data
  - Technology complexity and immaturity
- Following Agile and DevOps principles is a must

Appendix

Architecture

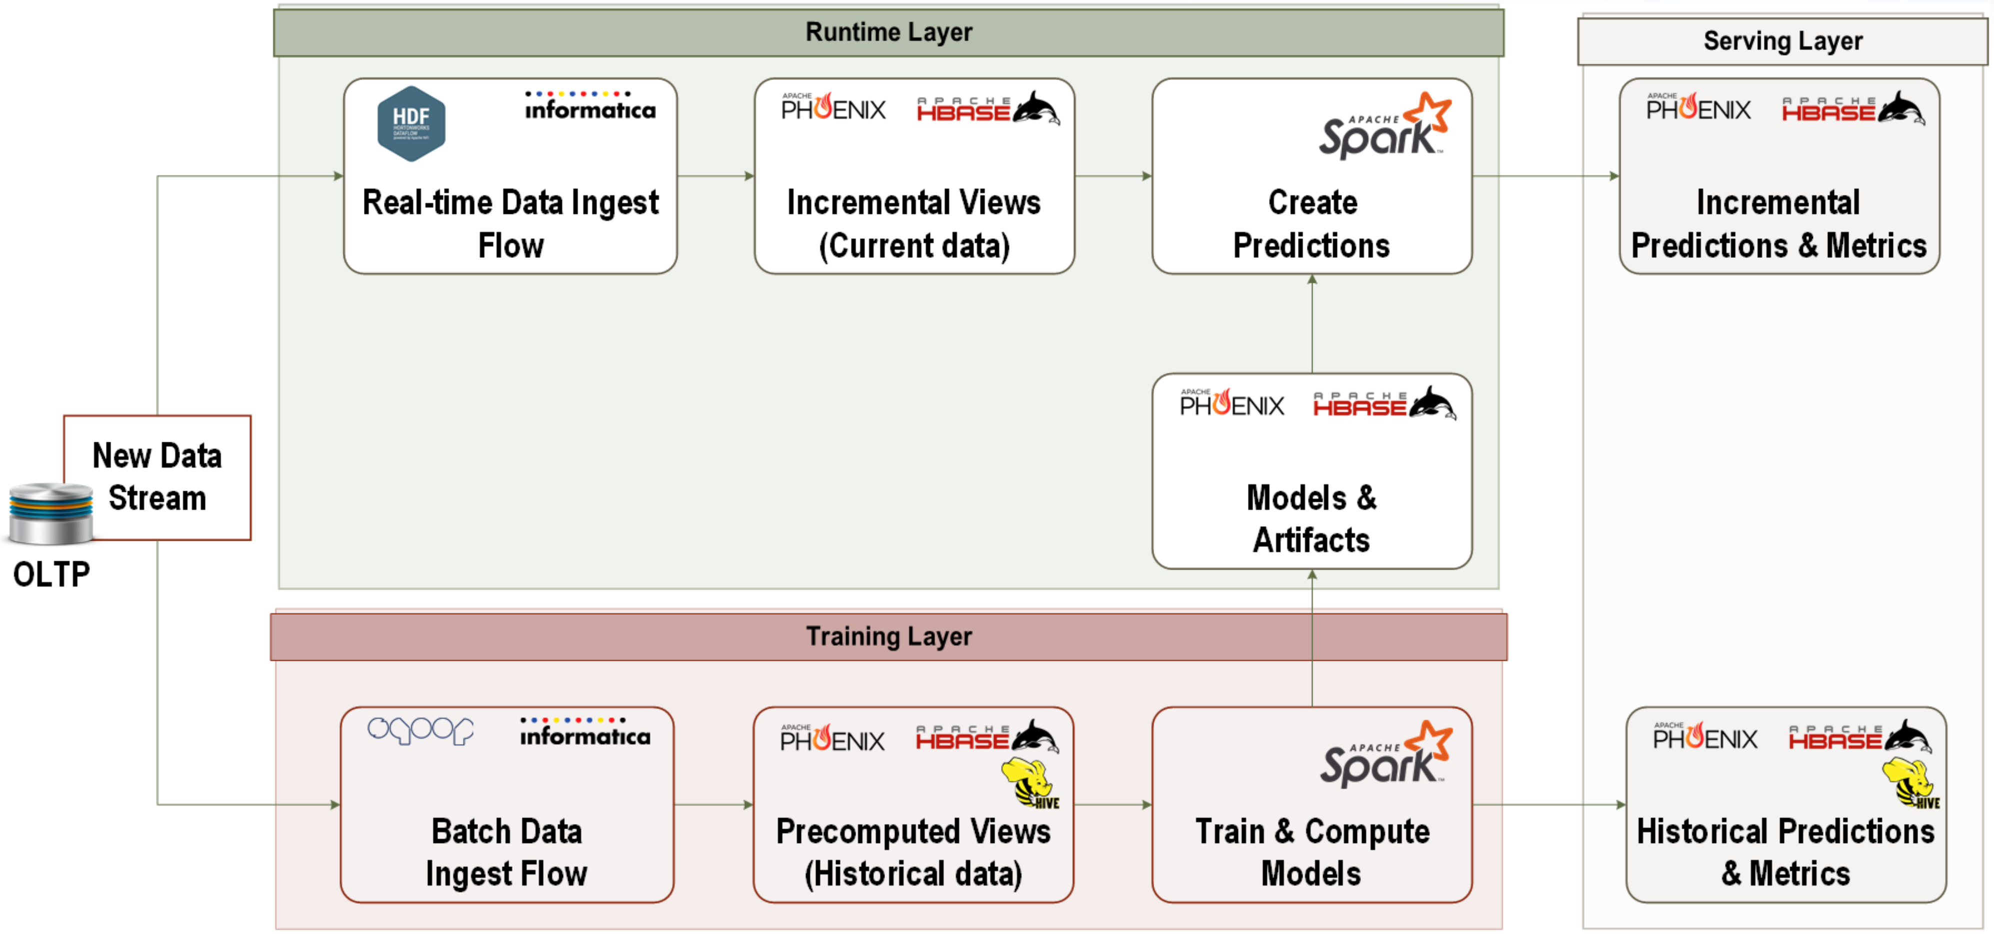

What are we facing now

- Managing the speed of change
  - The case for cloud and changing your mind … often
- Driving benefits
  - Product strategy to drive revenue
  - Use-case selection strategy
- Data Management maturity
  - Data at Forefront
  - Data Fitness
- Data Science maturity
  - Integration into engineering practices
  - Organizational literacy

Analytics Development Lifecycle

Operational Stage
- Streamline and automate the deployment of models
- Monitor model performance

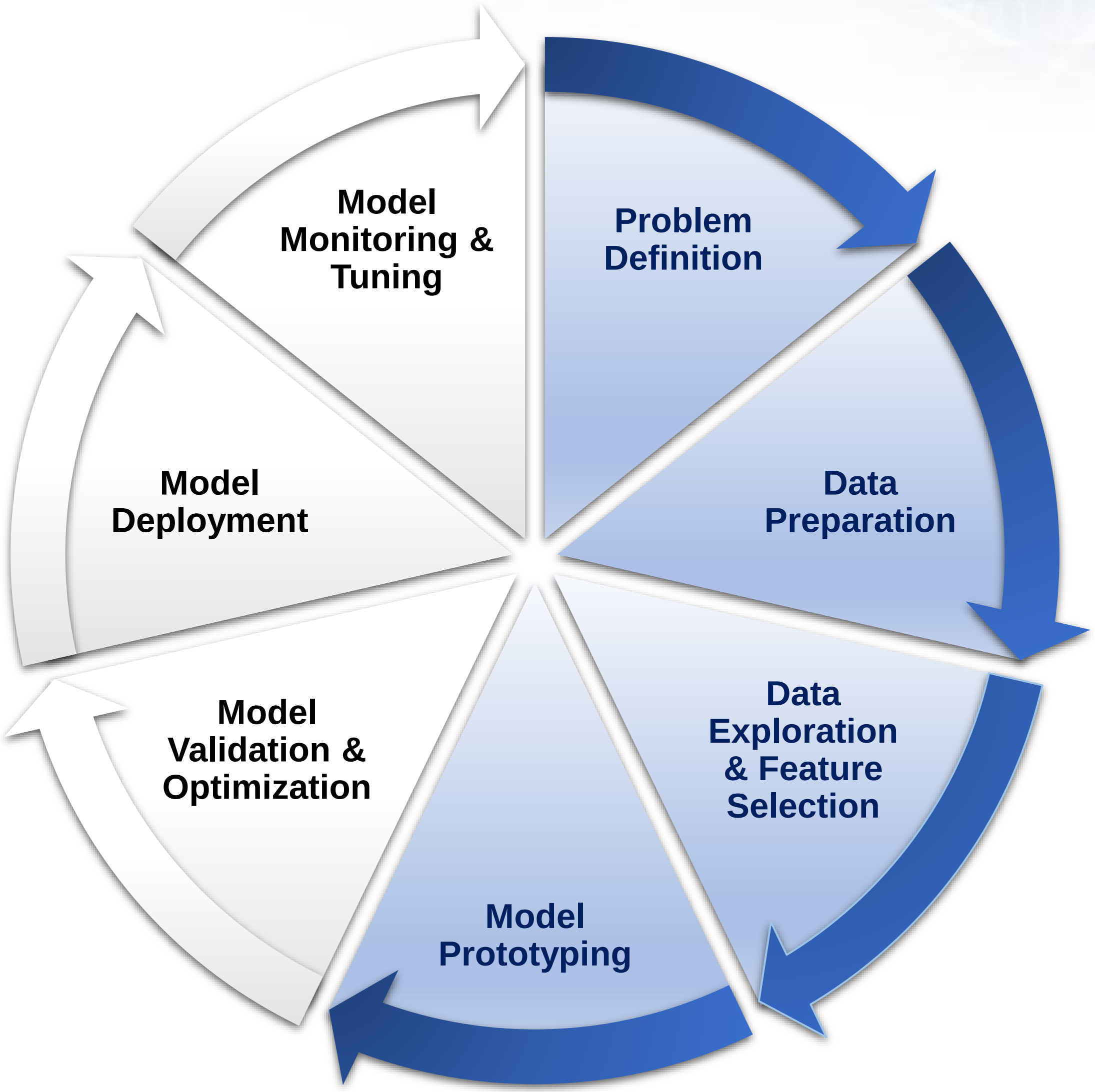

Research Stage
- Enhance data analysis and preparation
- Streamline model creation
- Leverage data platform toolsets

ADLC and SDLC Integration

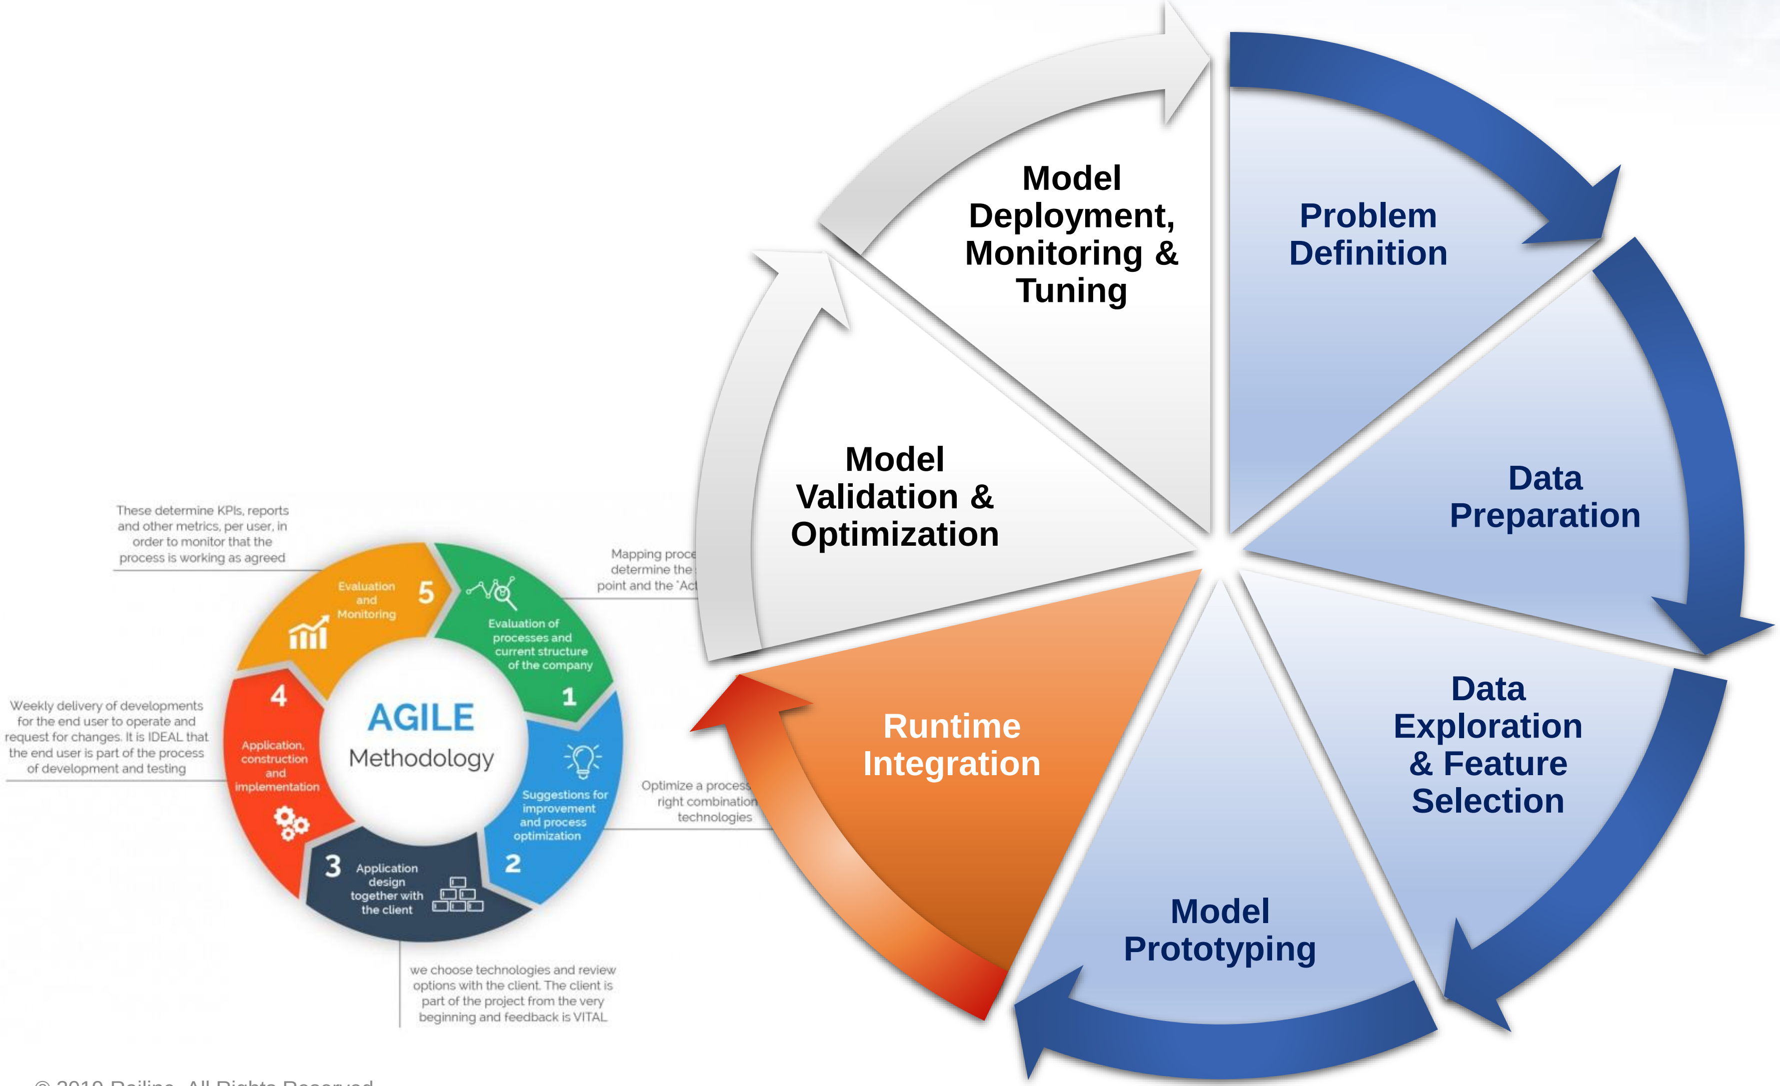

4 tracks to integrate

- Data Science
- Data Engineering
- Application Development
- Infrastructure