Unlocking the Power of AI at Ford: A Behind-the-Scenes Look at Mach1ML and Airflow

- Elona Zharri
  - Machine Learning Engineer – Mach1ML

- Nikhil Nandoskar
  - Machine Learning Engineer – Mach1ML

- Prince Bose
  - Tech Lead – Mach1ML

Outline

- Ford's Journey into AI/ML
- Challenges of ML Workflows in Ford
- Intro to Mach1ML and MDK
- Pre-Airflow MDK v1.0
- Airflow Features Enabling Mach1ML
- Various Ford Use Cases, Metrics, Usage
- Conclusion

Ford Motor Company

- Founded in 1903
- 4.4 million vehicles sold in 2023 (6th worldwide)
- ~2 Million Vehicles sold in U.S. in 2023
- F-Series trucks are the best-selling truck lineup in the US for 46 years running [1]
    - » At least one F-Series truck being sold every 49 second on average
    - » Second to only iPhone in terms of product sales
    - » Generate more revenue than the NFL, MLB, NBA, and NHL combined

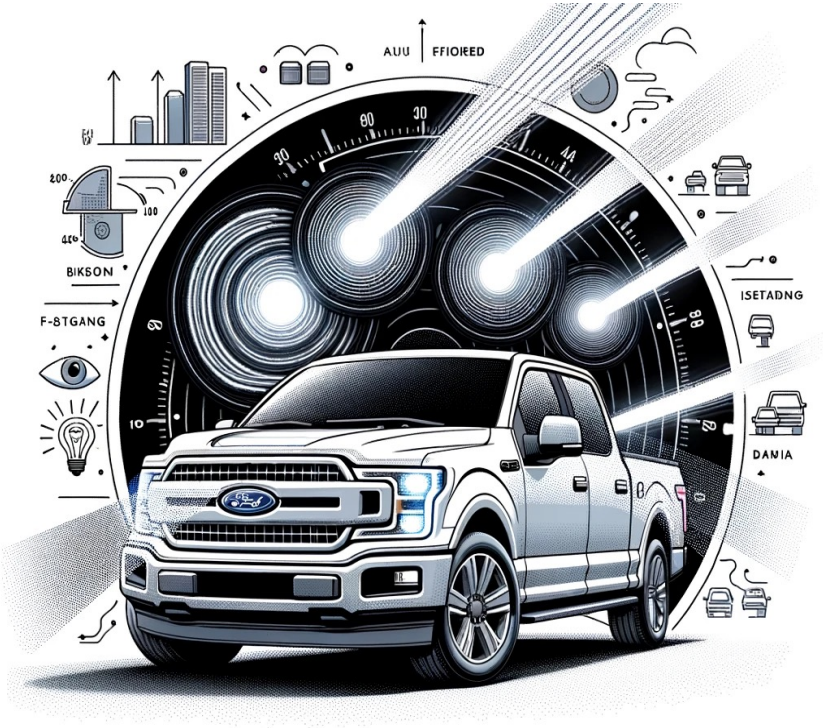

GDI&A
Global Data Insight & Analytics

[1] Based on 1977-2022 CY total sales.

Challenges in Adopting Machine Learning

Organizational

- 100-year-old company needs change management
- Groups duplicate ML work
- Need coordination for common ML patterns and tools
- Lack of established MLOps skillset and practices

Technical Challenges

- Ford infrastructure not designed for ML
- No integrated ML tool stack
- IT rules pose ML hurdles
- Legacy code not ML-optimized

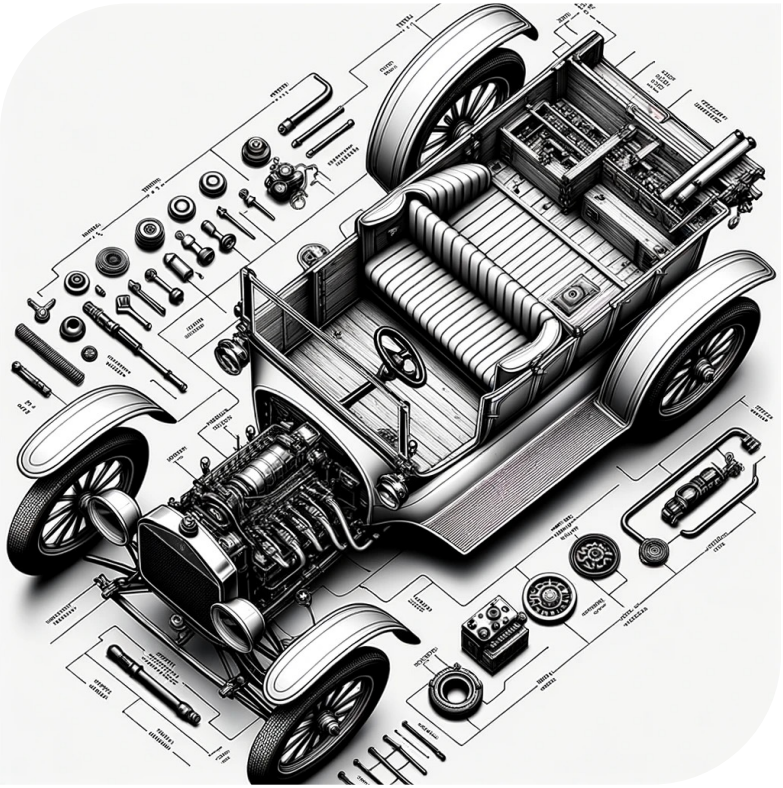

What Did We Do to Address?

- **Coordinating the Talent**
  - Ford's Artificial Intelligence Advancement Center
  - Team of 100+ ML and DevOps experts
  - Consulting on and implementing key use-cases
- **Data Factory**
  - Data Centralization Platform
- Enter: **Mach1ML**
  - **Tool Offerings**
    - » Compute Cluster, Google Cloud, Seldon Deploy, Weights & Biases, etc.
  - **MLOps Platform** (multiple iterations)
    - » Team of 25+ engineers
    - » On-prem low/no code system for new users
    - » Cloud-based codable system for experts
    - » Hybrid system for varying users
- **Generative AI**
  - **FordLLM**
    - » Portal for usage of enterprise LLMs (ChatGPT, Gemini, Llama)
    - » Programmatic access via APIs to multiple models
  - Many applications in development

AI/ML Landscape @ Ford Motor Company

- We are at the cutting edge of applying AI across every use-case imaginable
    - » Connected vehicle, manufacturing, advanced driver assistance systems, marketing, finance, supply chain, etc.

- Navigating complex technical and organizational challenges, beyond a one-size-fits-all solution
    - » *Legacy systems integration* - Working with and integrating many outdated legacy systems, while transitioning beyond them
    - » *Cross-team coordination* - Coordinating across multiple large teams (each the size of a startup), with the intention to effectively serve everyone

Intro to Mach1ML

- **What is Mach1ML?**
  - Ford's AI and Machine Learning tool suite.
  - **Mission -** Simplify AI/ML development, deployment, and scaling for Ford's community.
  - **Goal -** Empower Ford's AI/ML community with tools and infrastructure.

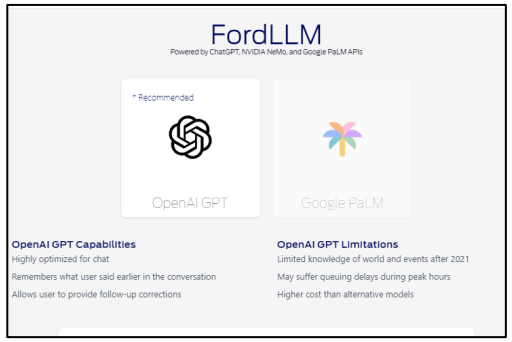

FordLLM

Over 50,000+ unique Ford users!

- **What is MDK?**
  - MDK (**MLOps Development Kit**): Extensible MLOps framework for building "**production-ready"** machine learning pipelines.
  - **Compatibility**: Works with Google Cloud Platform (**GCP**), on-prem High Performing Computer (**HPC**) and **Hybrid** Approaches.
  - **Focus**: Empowering ML engineers to focus on model development, not productionization.
  - **Offerings -** Provides tools, frameworks, and infrastructure for *end-to-end ML model development and deployment*

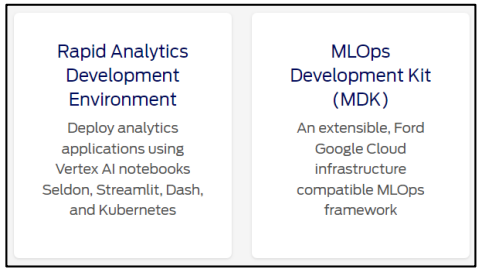

AI / ML Development Frameworks

SELDON

Deployment

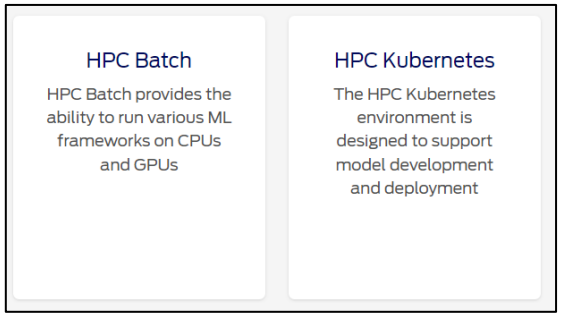

High Performance Computing Center (HPC) GPUs

Weights & Biases

Experiment Tracking

Ford GDI&A Global Data Insight & Analytics

Mach1ML MDK (MLOps Development Kit) v1.0 - Pre-Airflow Era

- **Initial Approach**
  - Utilized a Kubeflow-based approach for orchestration
  - Enabled diverse machine learning workflows on Vertex AI

- **Challenges Faced**
  - Resulted in platform lock-in to Google Cloud
    - Overhead in 3rd Party Integrations
  - High learning curve associated with Kubeflow

- **Demands and Requirements**
  - Need for on-premises orchestration
  - Need for hybrid orchestration (HPC/GCP)

- **Solution Sought**
  - Robust Orchestration tool
  - Rapid adoption

Finding the Right Orchestration Tool - Airflow

- **Low Barrier to Entry:**
  - Runs with well-established tools (Astronomer, Docker).
  - Easy to start a simple dev environment.
  - Pre-existing operators for common tasks.
- **High Customizability:**
  - Written in Python, easy to understand tasks/DAGs.
  - Integrates with services like BigQuery, VertexAI, Cloud Run, Seldon, and HPC.
  - Independent tasks allow easy DAG modifications.
  - Build custom operators for specific use cases.
  - Versatile for ML workflows, ETL, data pipelines, task scheduling.
- **Smooth CI/CD Integration:**
  - Easy local development with Astro CLI.
  - Provision remote workspaces and deployments via MDK project.
  - Easy pipeline promotion across environments.

Mach1ML MDK (MLOps Development Kit) v2.0 - Airflow Era

```yaml
general:
  use_case: !Mach1ml use_case
  cdsid: !Mach1ml generic_cdsid
  project_id: !Mach1ml gcp_project_id
  astronomer_workspace_name: !Mach1ml astronomer_workspace_name
  astronomer_deployment_name: !Mach1ml astronomer_deployment_name
defaults:
  hpc: &HPC
    cdsid: !Mach1ml generic_cdsid
    runtool: runpytorch
    ssh_config:
      ssh_conn_id: "ssh_default"
    gcp_config:
      vault_service_account: !Mach1ml gcpVaultSA
      project_id: !Mach1ml gcp_project_id
      keyfile_location: !Mach1mlFormat /s/{generic_cdsid}/gcp_key.json
tasks:
  model_training: !Task.HPC.SSHJob
    <<: *HPC
    script: pytorch_using_trainer.py
    script_type: python
    runtool_arguments:
      NGPUS: 8
      i: hpcregistry.hpc.ford.com/mach1mlmodel/mach1ml-pytorch_ftenf:0.0.5
    script_args:
      batch_size: 32
      epochs: 2
      bucket_name: !Mach1ml HPCArtifactsBucket
      num_labels: 10
      lr: 0.00005
      folder_name: mnist/trainingSet
      gcs_csv_path: !Mach1mlFormat gs://{HPCArtifactsBucket}/mnist/mnist_test_transformer.csv
      input_csv_file: !Mach1mlFormat /s/{generic_cdsid}/mnist_test.csv
      model_checkpoint: 'google/vit-base-patch16-224'
      remove_unused_columns: false
      evaluation_strategy: epoch
      save_strategy: epoch
      gradient_accumulation_steps: 4
      warmup_ratio: 0.1
      logging_strategy: epoch
      load_best_model_at_end: true
      metric_for_best_model: accuracy
      output_dir_for_each_checkpoint: 'finetuned-mnist'
    wait_for_completion: true
```

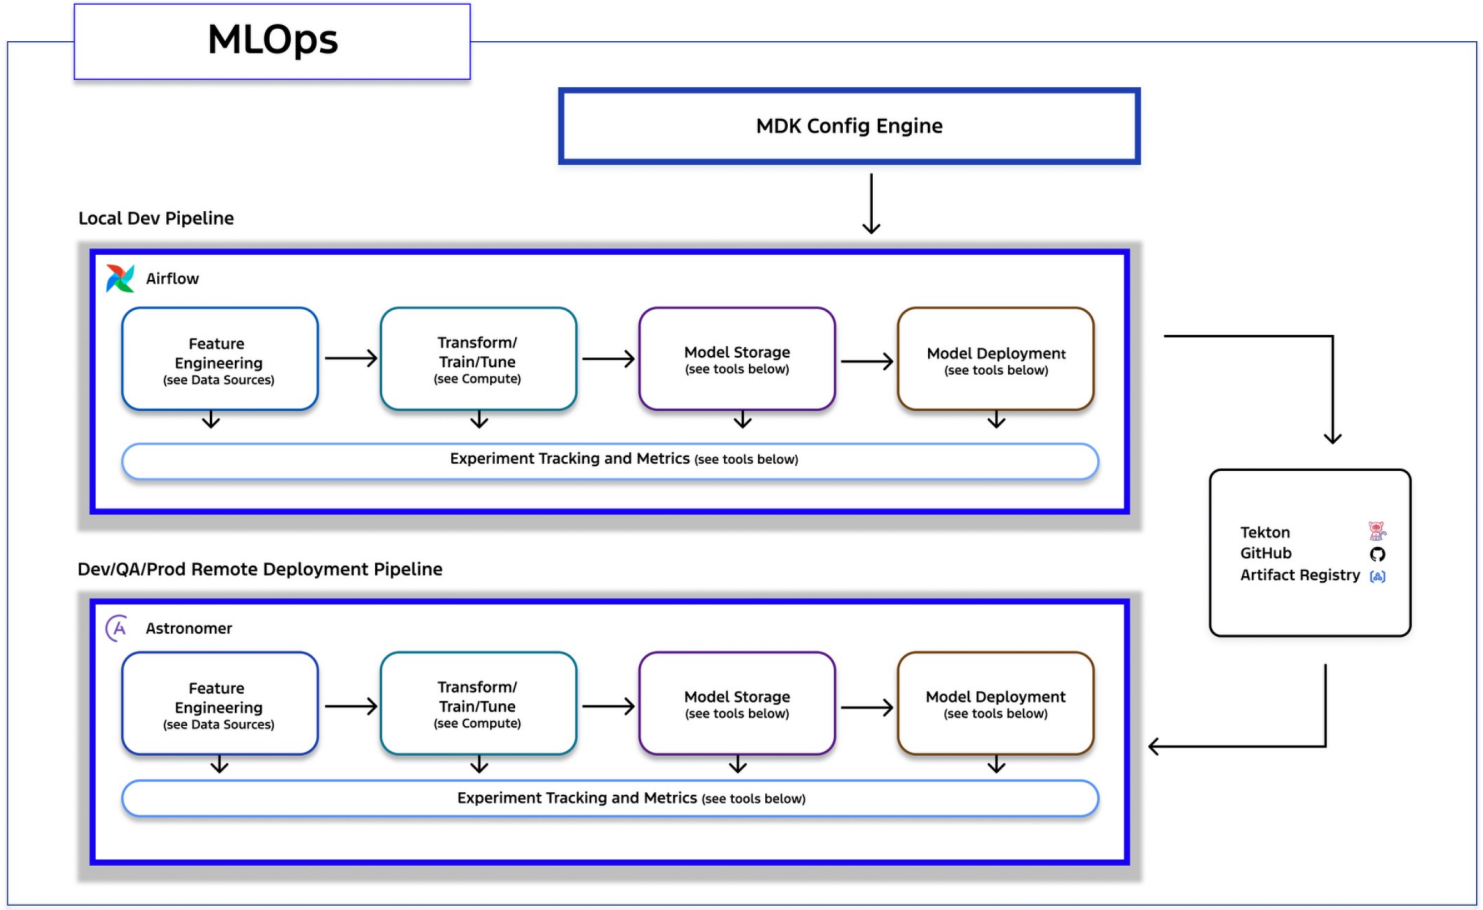

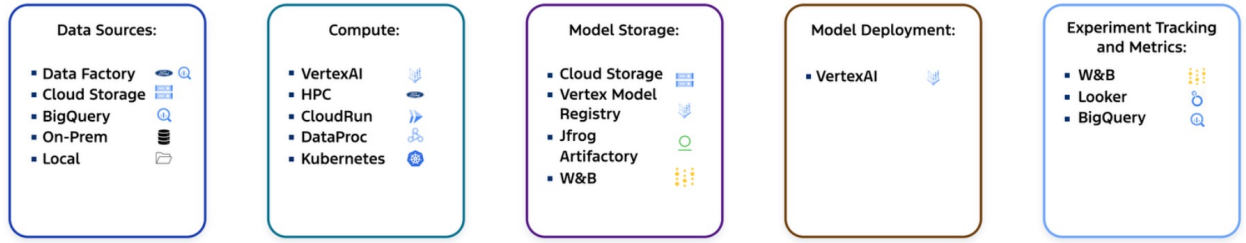

Ford GDI&A
Global Data Insight & Analytics

Custom Built Mach1ML Operators

- Developed a series of provider packages
  - Designed to complement the existing Airflow operators
  - Facilitate access to various internal and external services
- *HPC Operator*
  - Submits jobs onto **HPC** and monitors them for completion.

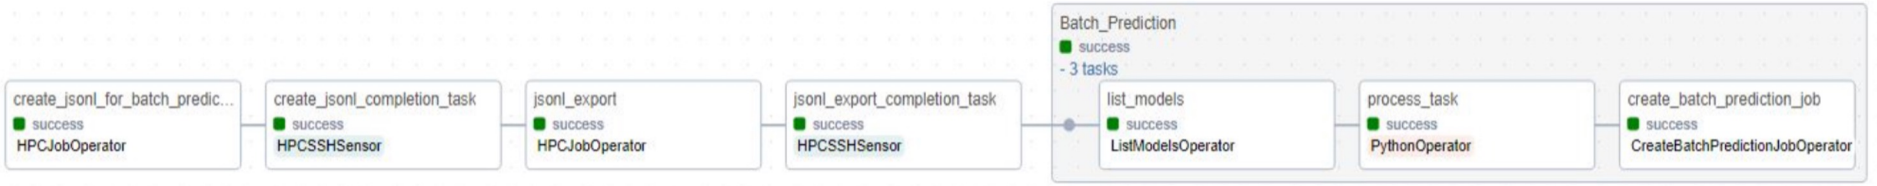

- *VertexAI Operator*
  - Submits jobs onto **VertexAI** and monitors them for completion
- *Weights and Bias Operator*
  - Provides integration with **Weights and Biases** – Artifact and Experiment tracking with Airflow DAGs

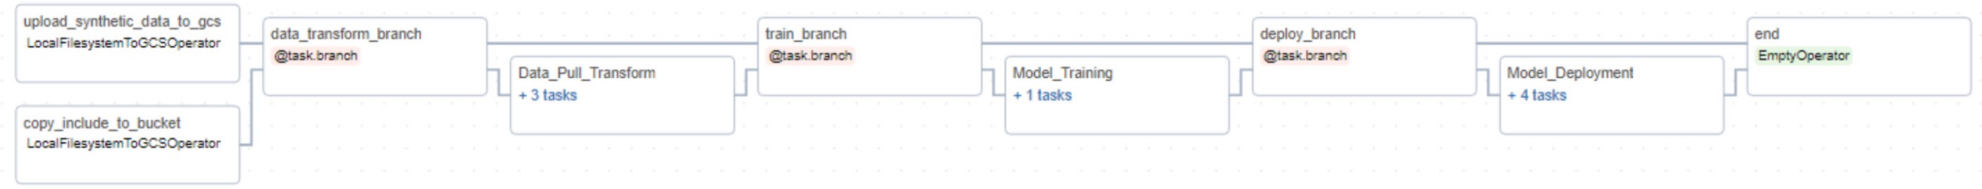

Custom Mach1ML DAG Decorator

- Facilitates better tracking and Debugging of Airflow DAGs
- Enhance logging and error handling for Airflow DAGs
- Generates comprehensive task and DAG level metrics
- Enables insights into DAG execution, including duration, state, and arguments

```python
@mach1ml_dag(dag_id='cv_hybrid_mach1ml_hf_pytorch_inference_pipeline',
             description="inference",
             start_date=datetime.datetime.now(),
             schedule_interval=None,
             tags=["hybrid-hf-pytorch-inference", "dev"],
             mach1ml_general_config=run_config.general,
             )
def dag_inference():
    create_jsonl_for_batch_predictions_task = HPCJobOperator.from_config(run_config, 'create_jsonl_for_batch_predictions')

    create_jsonl_completion_sensor = HPCSSHSensor(
        task_id='create_jsonl_completion_task',
        source_task_id='create_jsonl_for_batch_predictions',
        poke_interval=150,
    )

    jsonl_export_task = HPCJobOperator.from_config(run_config, 'jsonl_export')

    jsonl_export_completion_sensor = HPCSSHSensor(
        task_id='jsonl_export_completion_task',
        source_task_id='jsonl_export',
        poke_interval=150,
    )
```

Reusable Templates - MDK Examples

Hybrid Examples

- Generative AI
  - Multi Label Classification
- Computer Vision
  - Tensorflow
  - Ray
  - Pytorch
  - Huggingface
  - Qnn

Batch ML Examples

- *Integrations*
  - PySpark
  - W&B
- *Deployments*
  - VertexAI
  - Seldon

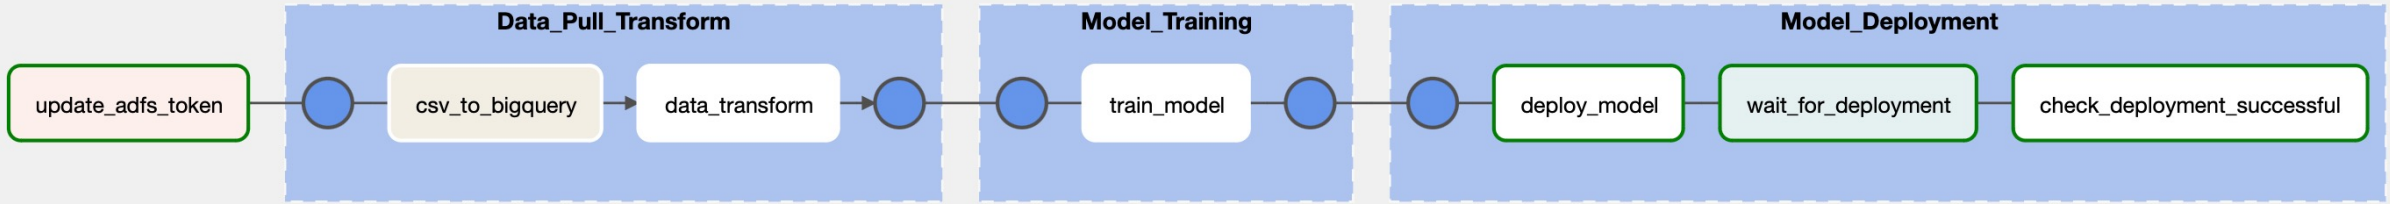

Reporting Analytics

- Tracking Usage metrics across multiple projects
  - Custom Mach1ML Decorator
- Drill Down metrics on operator usage
  - Standard Airflow Operators
  - Custom MDK Operators

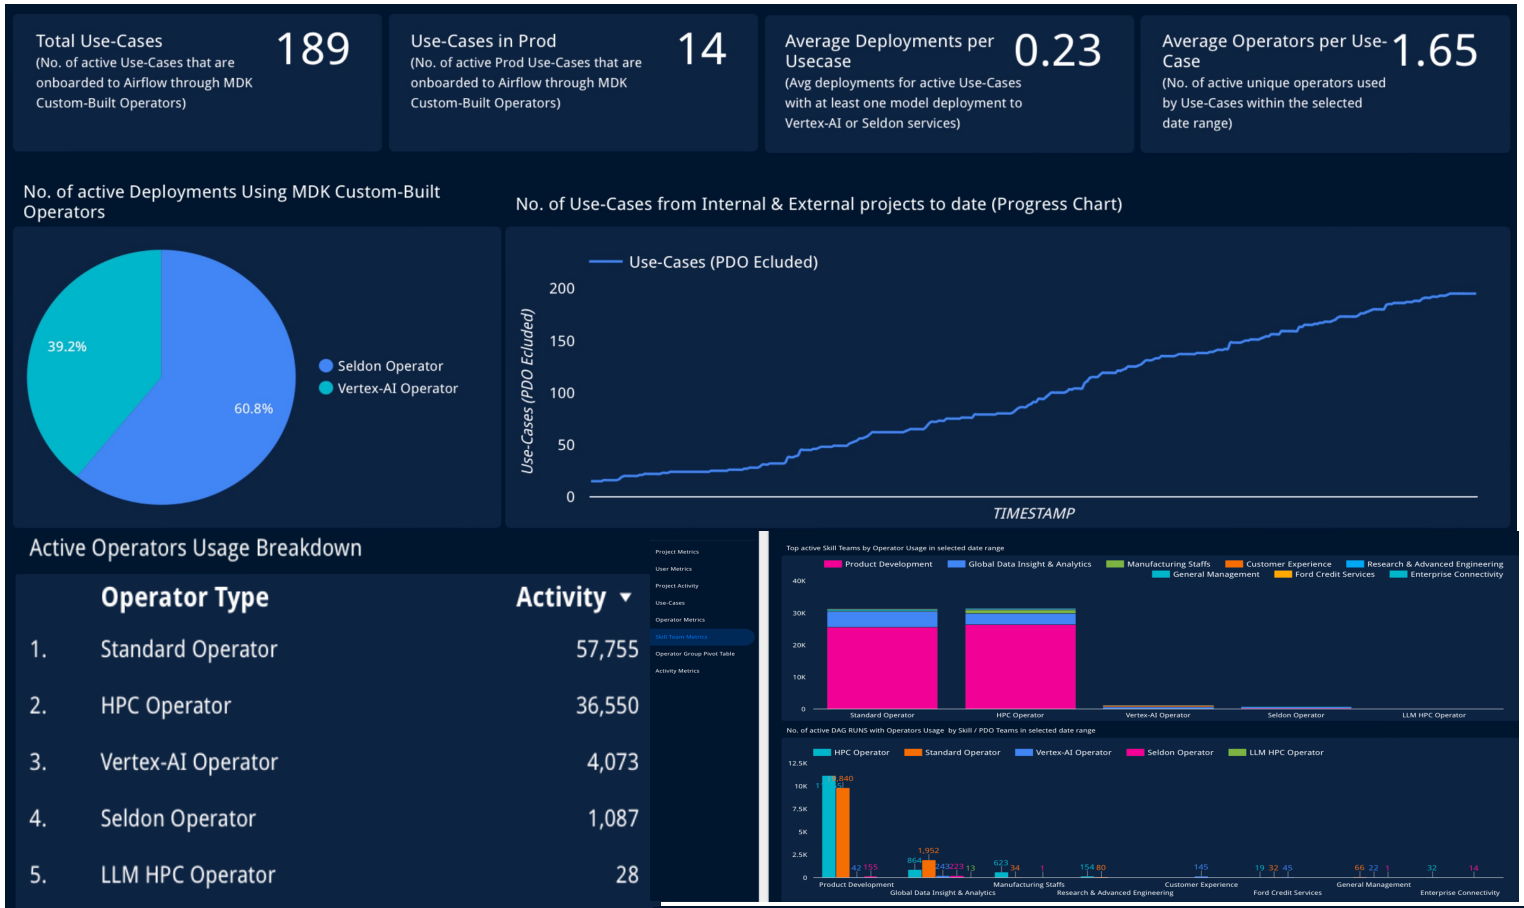

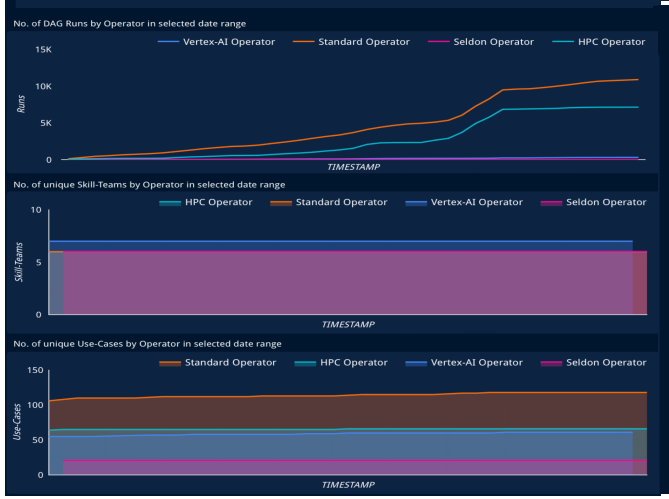

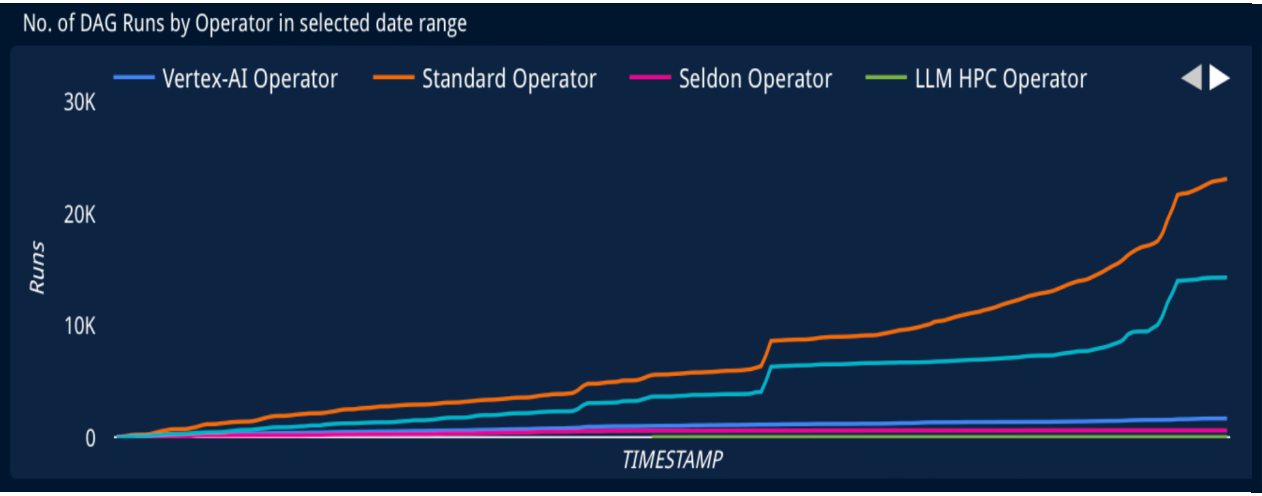

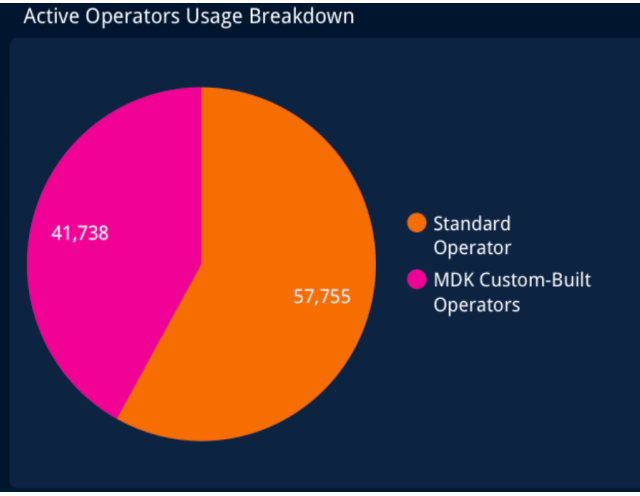

Mach1ML's Airflow Implementation : Key Features

- *Reusable Templates*
  - Pre-built DAGs for common ML tasks, customizable for specific needs.
- *Streamlined Deployments*
  - Flexible operators and plugins for diverse deployment environments.
- *Centralized Configuration Management*
  - Consistent configurations across environments.
- *Effortless Environment Promotions*
  - Seamless workflow promotion from development to production.

Accelerating Time-to-Value : Key Benefits

- *Faster Time-to-Value*
  - Accelerated delivery of ML solutions to production.
- *Increased Innovation*
  - Empowers Ford's AI practitioners to focus on developing cutting-edge solutions.
- *Enhanced Customer Experiences*
  - Drives the development of AI-powered features for next-generation vehicles.

Conclusion & Future Steps

- Mach1ML, powered by Apache Airflow, enhances Ford's AI capabilities by streamlining ML workflows, empowering AI practitioners, and accelerating AI-driven solutions for Ford's smart mobility vision

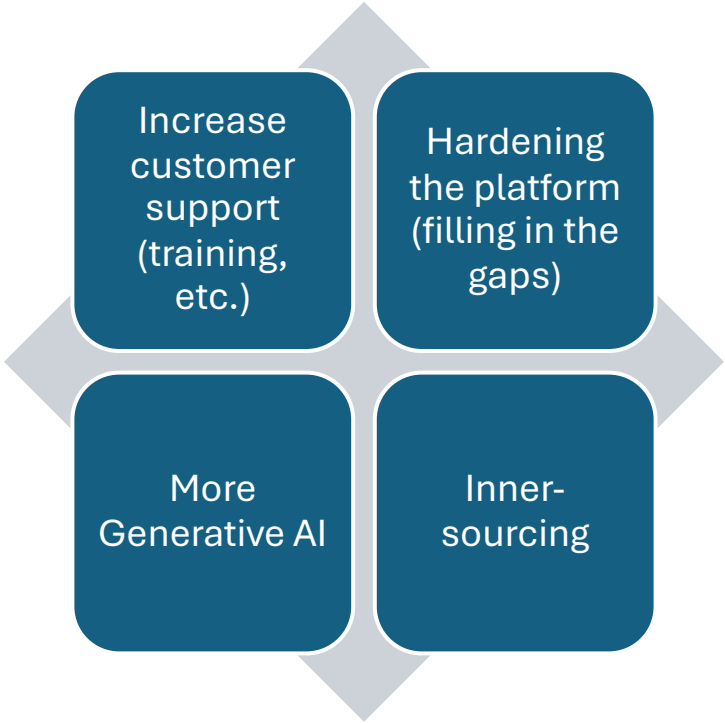

Questions?

Reach out to us:

Elona Zharri
www.linkedin.com/in/elonazharri/

Nikhil Nandoskar
www.linkedin.com/in/nikhil-nandoskar/

Prince Bose
www.linkedin.com/in/prince-bose/