A Ascensão de MLOps

Mar 2021
Felipe C Penha
#DataScienceBits

Data Science Bits

HACKERNOON

https://linktr.ee/felipepenha

Análitico Vs Operacional

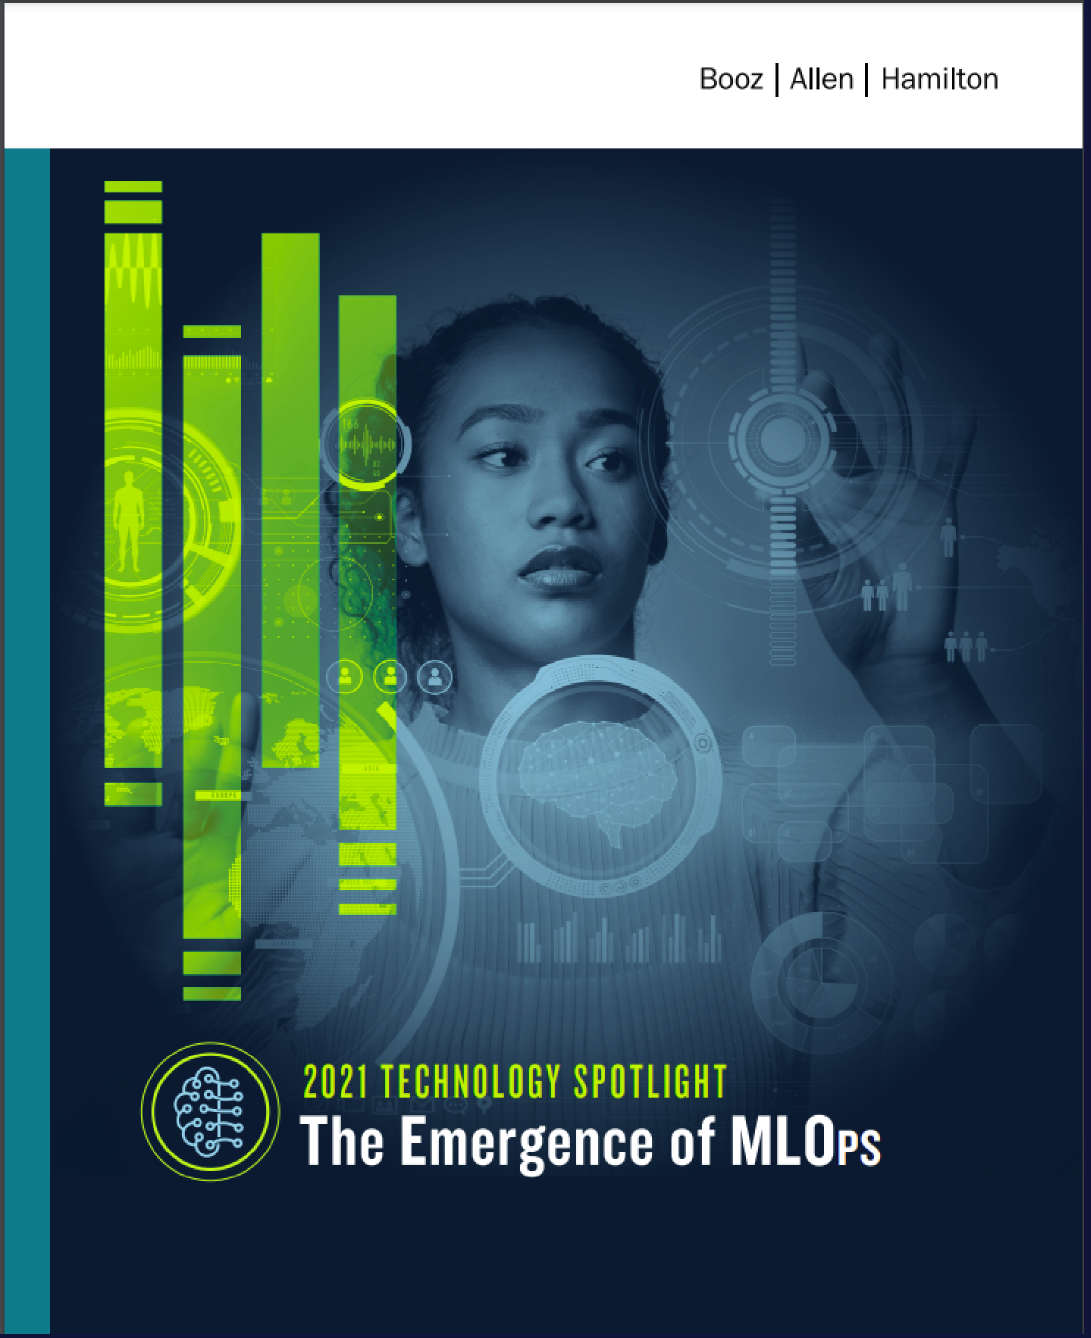

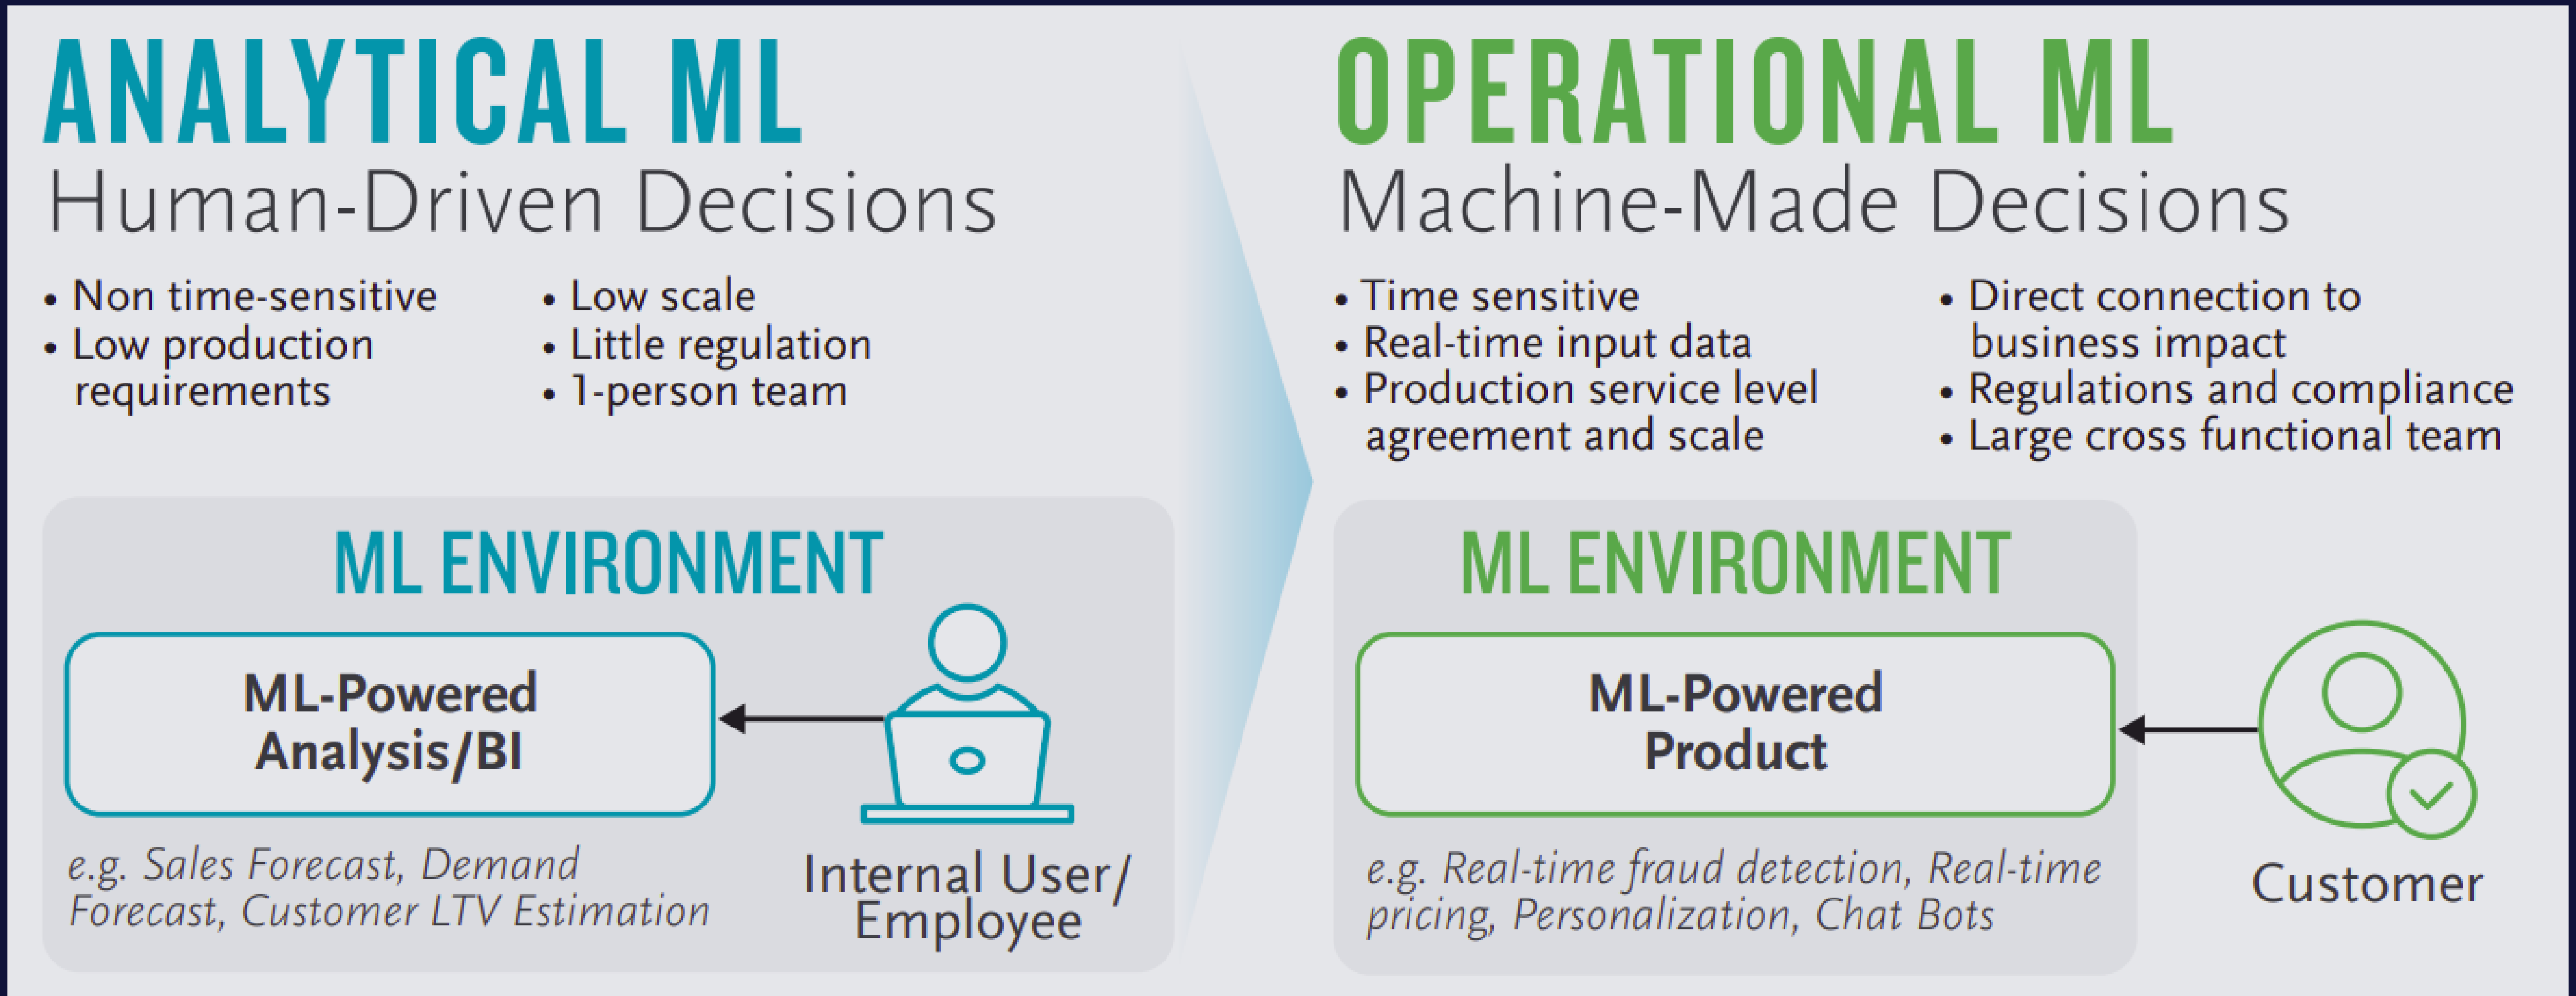

MLOps: "conjunto padrão de práticas para Operações de Machine Learning em grande escala"

2021

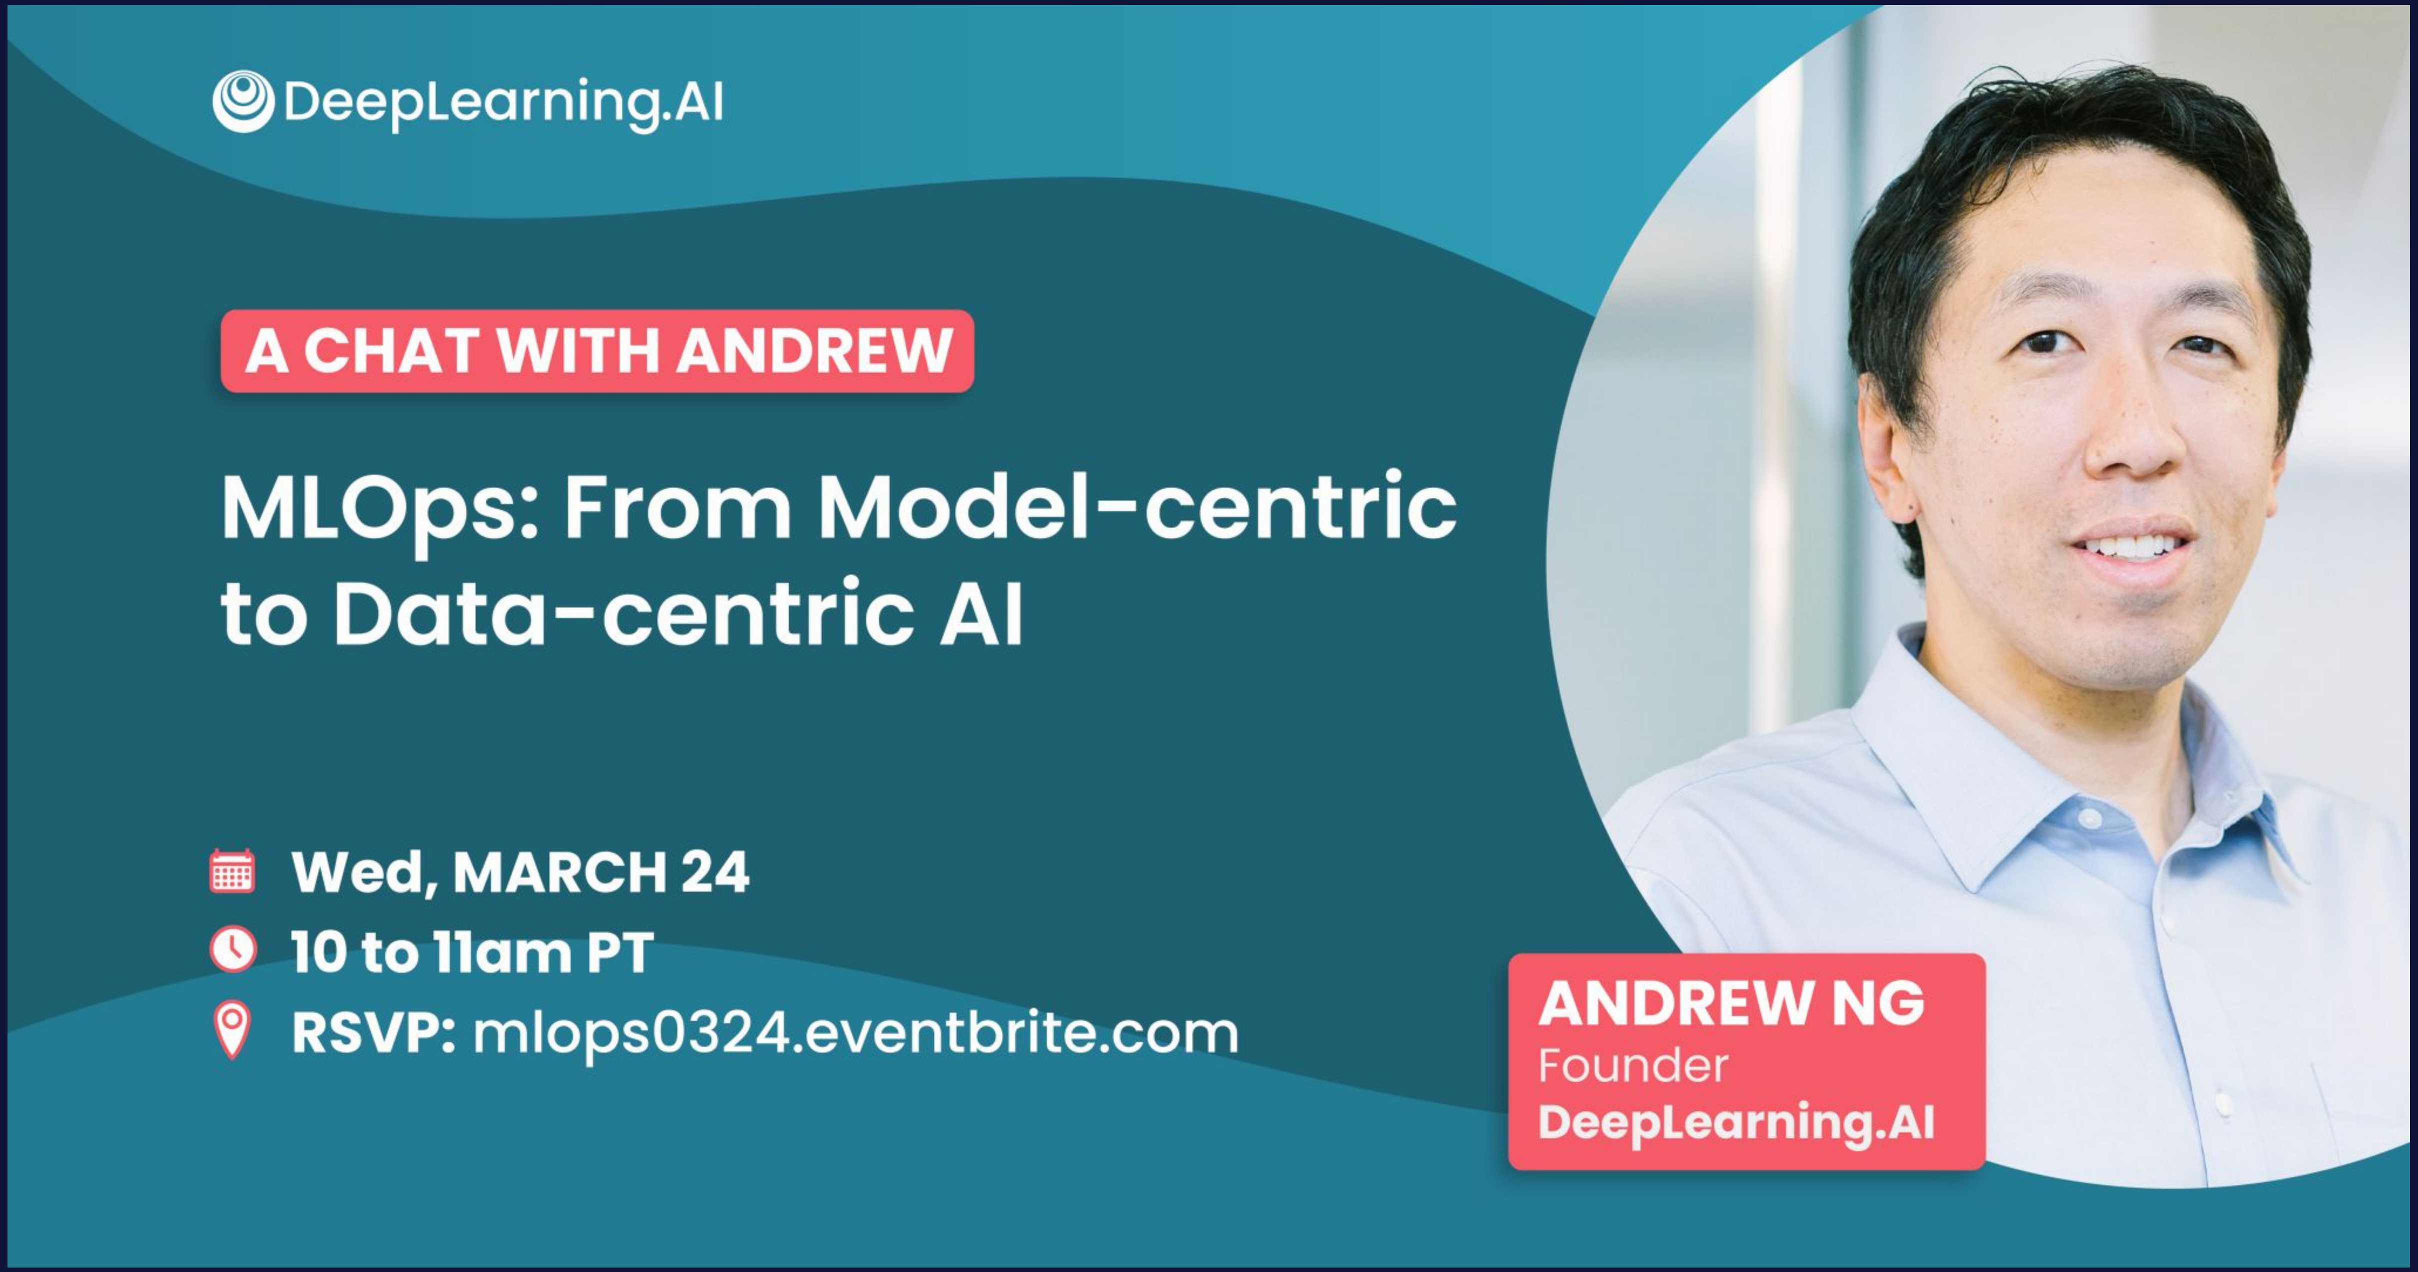

MLOps ~ DataOps ?

ML Eng. ~ Data Eng. ?

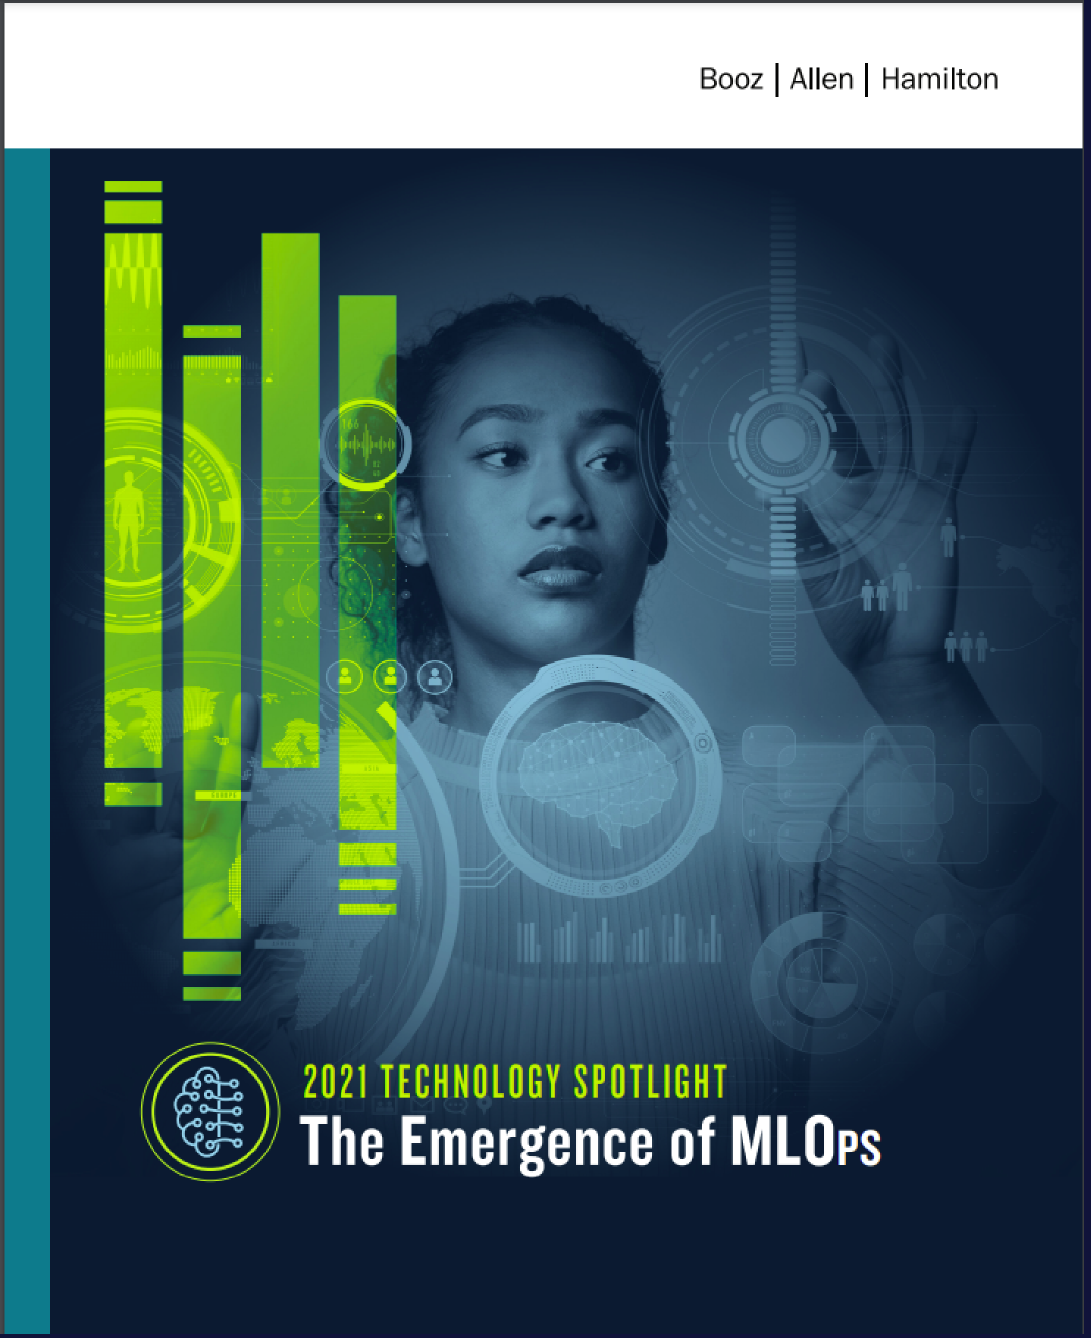

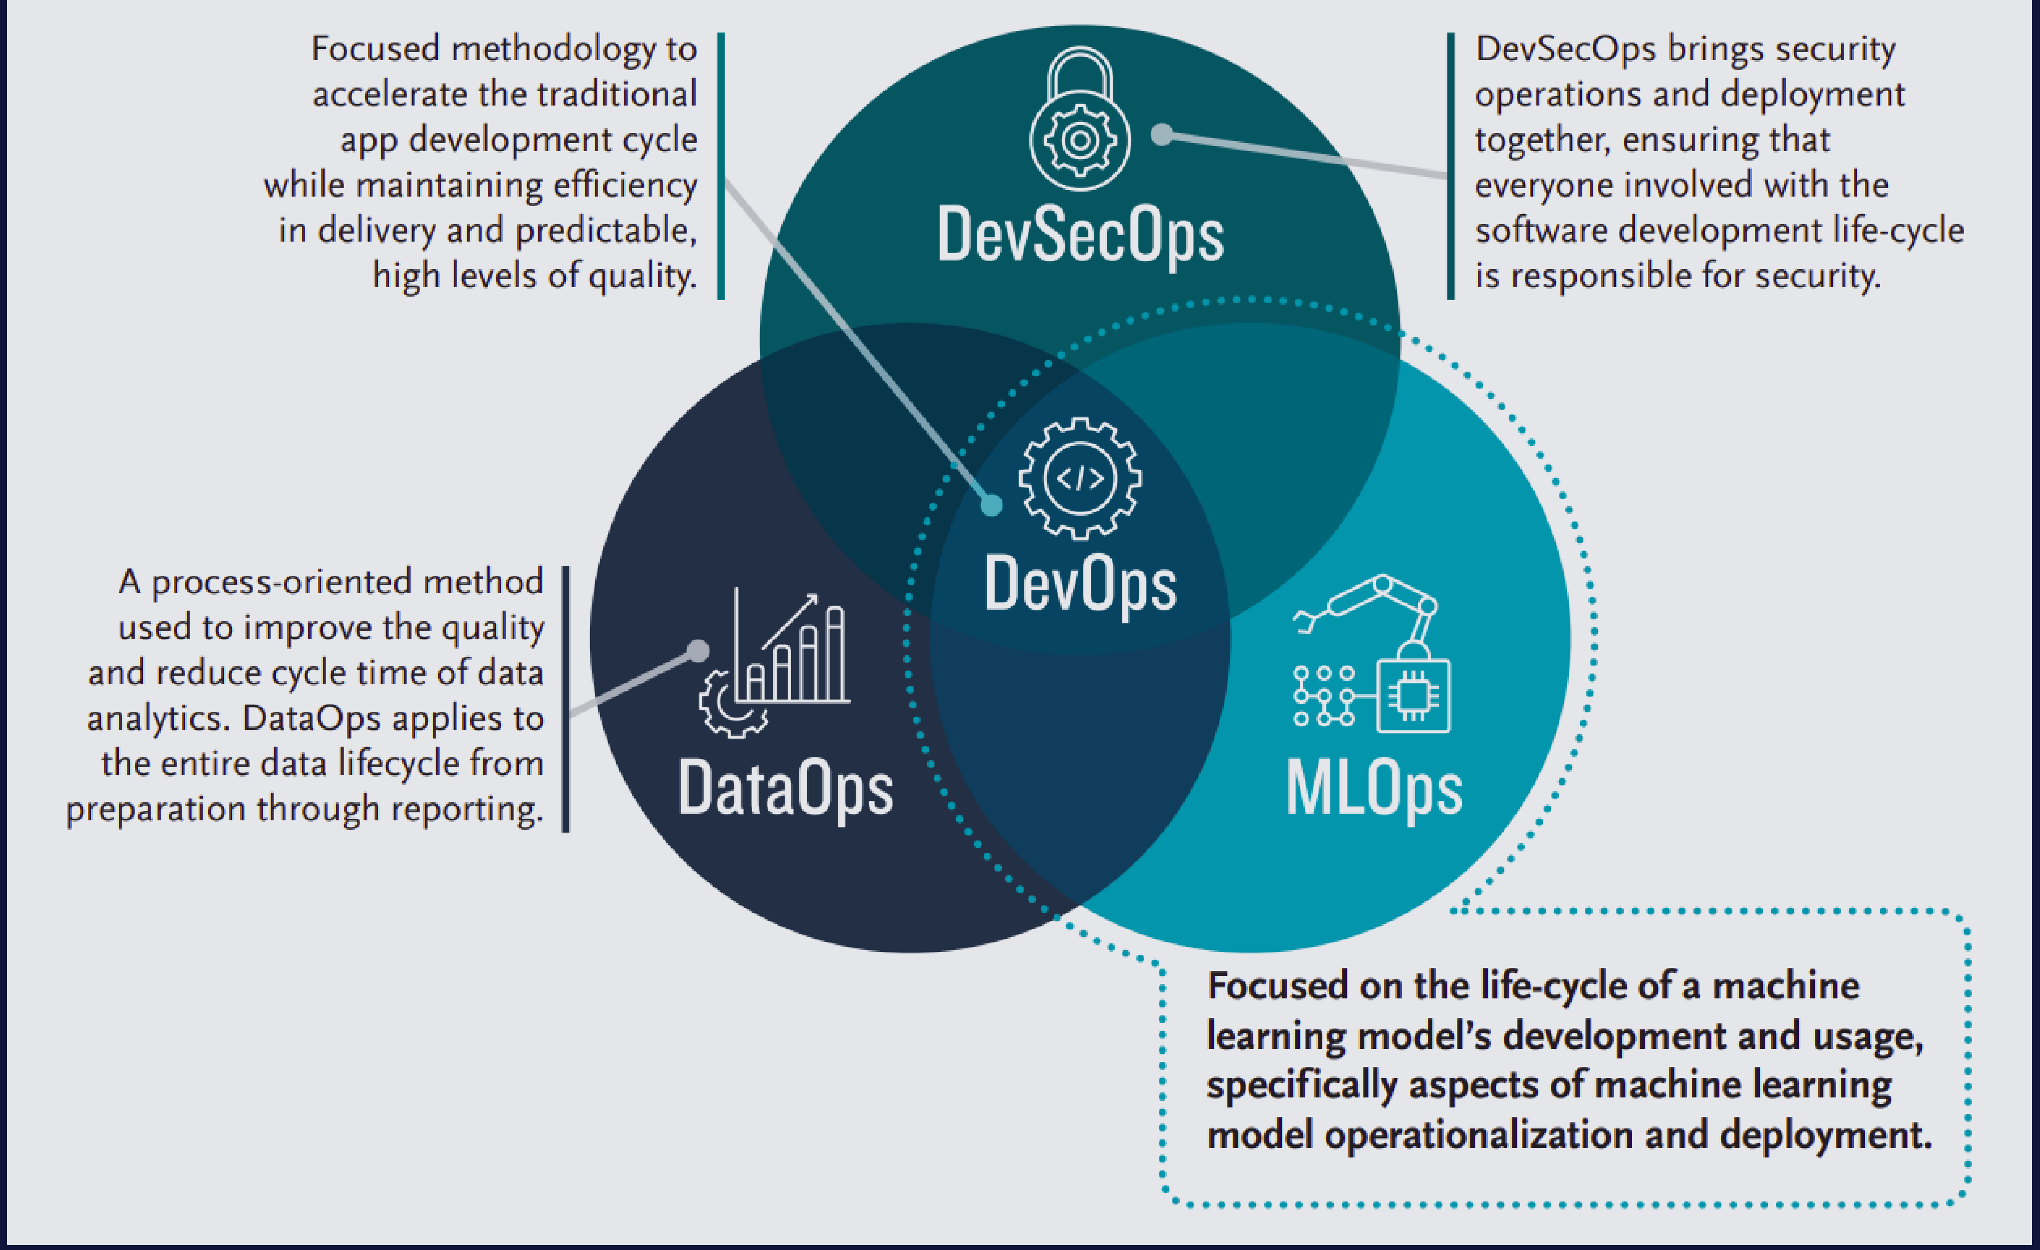

Figure 1–MLOps focuses on intricacies of the ML lifecycle often neglected in the DevOps framework

"Focado no ciclo de vida do desenvolvimento e uso de modelos, especificamente aspectos de operacionalização e deployment"

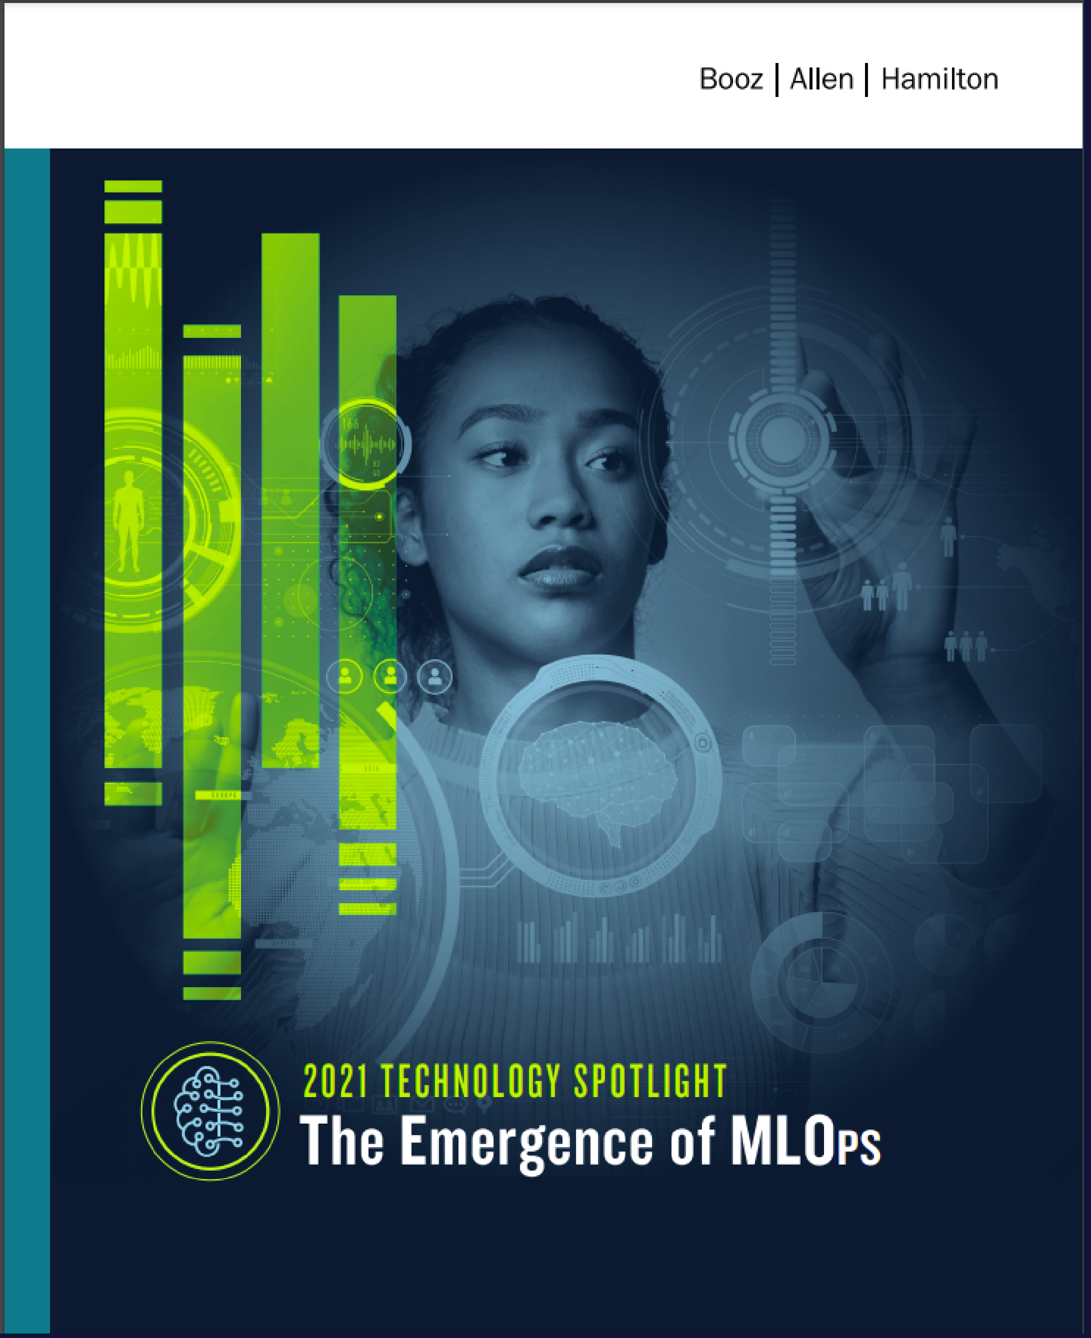

COMMON GOALS OF OPERATIONAL ML

Most MLOps solutions often share common goals born from the problems encountered in Analytical ML solutions. These include the need for auditability, performance across a variety of environments, transparancy into the functioning of a model, technology that enables the AI solution to scale, automation for model updates, and solving the common issue of model drift.

- **Auditability:** A given model will have multiple versions, a specific set of training data, and finely tuned hyperparameters. Each must be carefully tracked for versioning and testing purposes
- **Different Environments:** Different environments for data preparation, training, and model deployment reduce speed and scalability of model deployment
- **Model Transparency:** Individual models are difficult to understand and/or explain to others
- **Scalability:** Models created from scratch for individual problems reduce deployment speed because of the inability to reuse and recycle code
- **Model Drift:** Models may lessen in accuracy over time as the statistical properties the model tries to predict change in the real world

Reproducible

Reproducibility and productivity are inextricably linked. It's difficult to be productive when different team members can't reproduce each others' work. This is harder in ML than in software because test & training data and metrics need to be versioned alongside the code and environment.

Accountable

Models that are deployed without full provenance, a record of all the steps taken to create the models, can fail to be compliant, and are hard to debug. Maintaining this provenance record manually slows you down and is error-prone, so automated tooling is needed.

Collaborative

Concurrent collaboration – that is, collaboration without treading on each others' toes – is essential. In ML this is harder than in normal software engineering, because collaboration applies to notebooks, data, models and metrics as well as code.

Continuous

You're not done when you ship. In order to continue delivering value to the business, models must be retrained and statistically monitored to compensate for model drift due to constant changes in your business environment.

Source: mlops.community/manifesto/

Linux Foundation AI Landscape

2020-11-01T21:39:53Z 83fd88a

See the interactive landscape at l.lfai.foundation

Greyed logos are not open source

Machine Learning

Framework
Accord.NET, FedAI, Microsoft LightGBM, Mahout, ML.NET, Ray

Platform
Angel (LF AI & Data), ForestFlow (LF AI & Data), Alibaba Alink, Apache SystemML, Ax, cortex, H2O, Kubeflow, Metaflow, mlflow, PredictionIO, Seldon, Vowpal Wabbit

Library
CatBoost, Spark MLlib, mlpack, scikit-learn, openml, learn, Shogun, Sonnet, Salesforce TransmogrifAI, XGBoost, xLearn

Deep Learning

Framework
SINGA, Chainer, CNTK, dy/net, Alibaba euler, MindSpore, mxnet, ncnn, PaddlePaddle, Pythia, PyTorch, TensorFlow

Platform
DeepDetect, Determined AI, jina, Onepanel, Polyaxon, Alibaba X-DeepLearning

Library
BigDL, Catalyst, DL4J, Keras, OpenMined PySyft, PyTorch Lightning

Tool
BoTorch, Intel Distiller, plaidML, PyTorch BigGraph, PyTorch, tvm

Reinforcement Learning
Chainer RL, Coach, DeepMind Lab, Dopamine, Horizon, OpenAI, Google PlaNet, Google SEED RL, Stable Baselines

Programming
Pyro (LF AI & Data), DASK, Infer.NET, julia, NumPy, NYOKA, pomegranate, PyMC3, SciPy, Stan

Notebook Environment
Apache Zeppelin, BeakerX, colab, Elyra, IPython, jupyter, IBM, Polynote, rmarkdown, Stencila, Streamlit

Data

Versioning
dotscience, DVC, Pachyderm, Quilt

Store & Format
Milvus (LF AI & Data), Alluxio, orc, AresDB, Arrow, Avro, Delta Lake, druid, hudi, JanusGraph, Parquet, pilosa, Vearch, vespa

Operations
Amundsen (LF AI & Data), FEAST (LF AI & Data), Marquez (LF AI & Data, Incubating), Data Hub, great_expectations, snorkel

Stream Processing
NNStreamer (LF AI & Data), brooklin, Faust, Flink, kafka, Pravega, Prefect, samza, Uber uReplicator

SQL Engine
Apache Drill, presto, SQLFlow

Feature Engineering
Featuretools, tsfresh

Visualization
bokeh, plotly Dash, Uber deck.gl, Google Facets, Grafana, RCloud, NYU seaborn, Google TensorBoard

Pipeline Management
Intel Analytics Zoo, beam, DAGSTER, TPOT

Labeling and Annotation
Intel CVAT, Doccano, Label Studio, Labelbox, LabelImg, Hitachi Semantic Segmentation Editor, Microsoft VoTT

Governance
EGERIA (LF AI & Data), Community Data License Agreement

Model

Benchmarking
DAWNBench, MLPerf

Training
Horovod (LF AI & Data), Ludwig (LF AI & Data), Microsoft, Petastorm

Parameter
Hyperopt, talos

Format & Interface
ONNX (LF AI & Data), Uber Neuropod

Marketplace
Acumos (LF AI & Data), IBM

Workflow
Apache Airflow, nifi, argo, Azkaban, Cadence, Couler, Cyclone, DataBolt d6tflow, Flyte, Kedro, Spotify luigi, mleap, Prefect, TRAINS, Volcano

Inference
ADLIK (LF AI & Data), MNN, NVIDIA TensorRT, NVIDIA TensorRT Inference Server, uTensor

Tool
Qualcomm AIMET, Facebook dlrm, Microsoft MMdnn, aws MMS, MDB, amazon Neo-AI, NETRON, ONNX Runtime, PipelineAI, studio.ml, Google, aws TorchServe, turi

Trusted & Responsible AI

Explainability
AI Explainability 360 (LF AI & Data), ALIBI, ELI5, Microsoft InterpretML, University of Washington Lime, Google Lucid, SHAP, SKATER, Bolt TreeInterpreter

Adversarial
Adversarial Robustness Toolbox (LF AI & Data), AdvBox, advertorch, hans, Foolbox

Bias & Fairness
AI Fairness 360 (LF AI & Data), Aequitas, Audit AI, Stanford University DeepLIFT, Fairlearn

Distributed Computing

Computing & Management
EDL Elastic Deep Learning (LF AI & Data), SOAJS (LF AI & Data), Apache Ambari, Apache Bahir, MESOS, Apache Ranger, Apache Storm, GNES, NETFLIX genie, kubernetes, Intel Nauta, Singularity

Interface
sparklyr (LF AI & Data), Apache Toree, LIVY, John Snow Labs Spark NLP

The LF AI & Data landscape explores open source projects in Artificial Intelligence and Data and their respective sub-domains.

LF AI & DATA Landscape

LF AI & DATA

l.lfaidata.foundation

Security & Privacy
Google Differential Privacy, IBM HElib, Microsoft SEAL, SUBSTRA, Google TensorFlow Privacy, TF Encrypted

Natural Language Processing
DELTA (LF AI & Data), Google ALBERT, AllenNLP, Google Bert, CoreNLP, mozilla DeepSpeech, fastText, flair, LUON, Kashgari, Facebook LASER, Intel NLP Architect, NLP, ParlAI, Facebook PyText, RASA, spaCy, Transformers, Facebook XLM, YouTokenToMe, Cisco MindMeld

Education
OpenDS4All (LF AI & Data), DataPractices.org

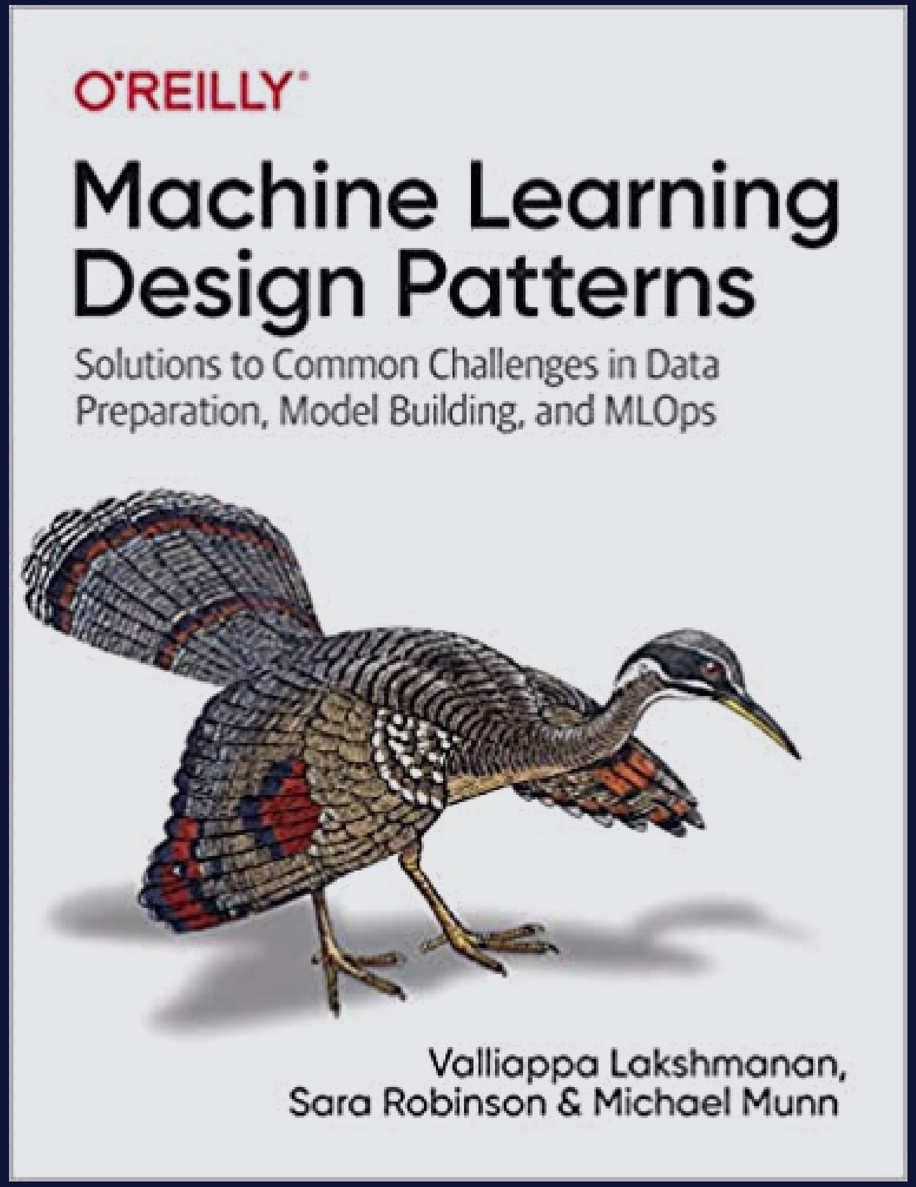

Design Patterns 1-30

- Problema
- Solução
- Por que funciona
- Prós, contras e alternativas

Design Pattern 26

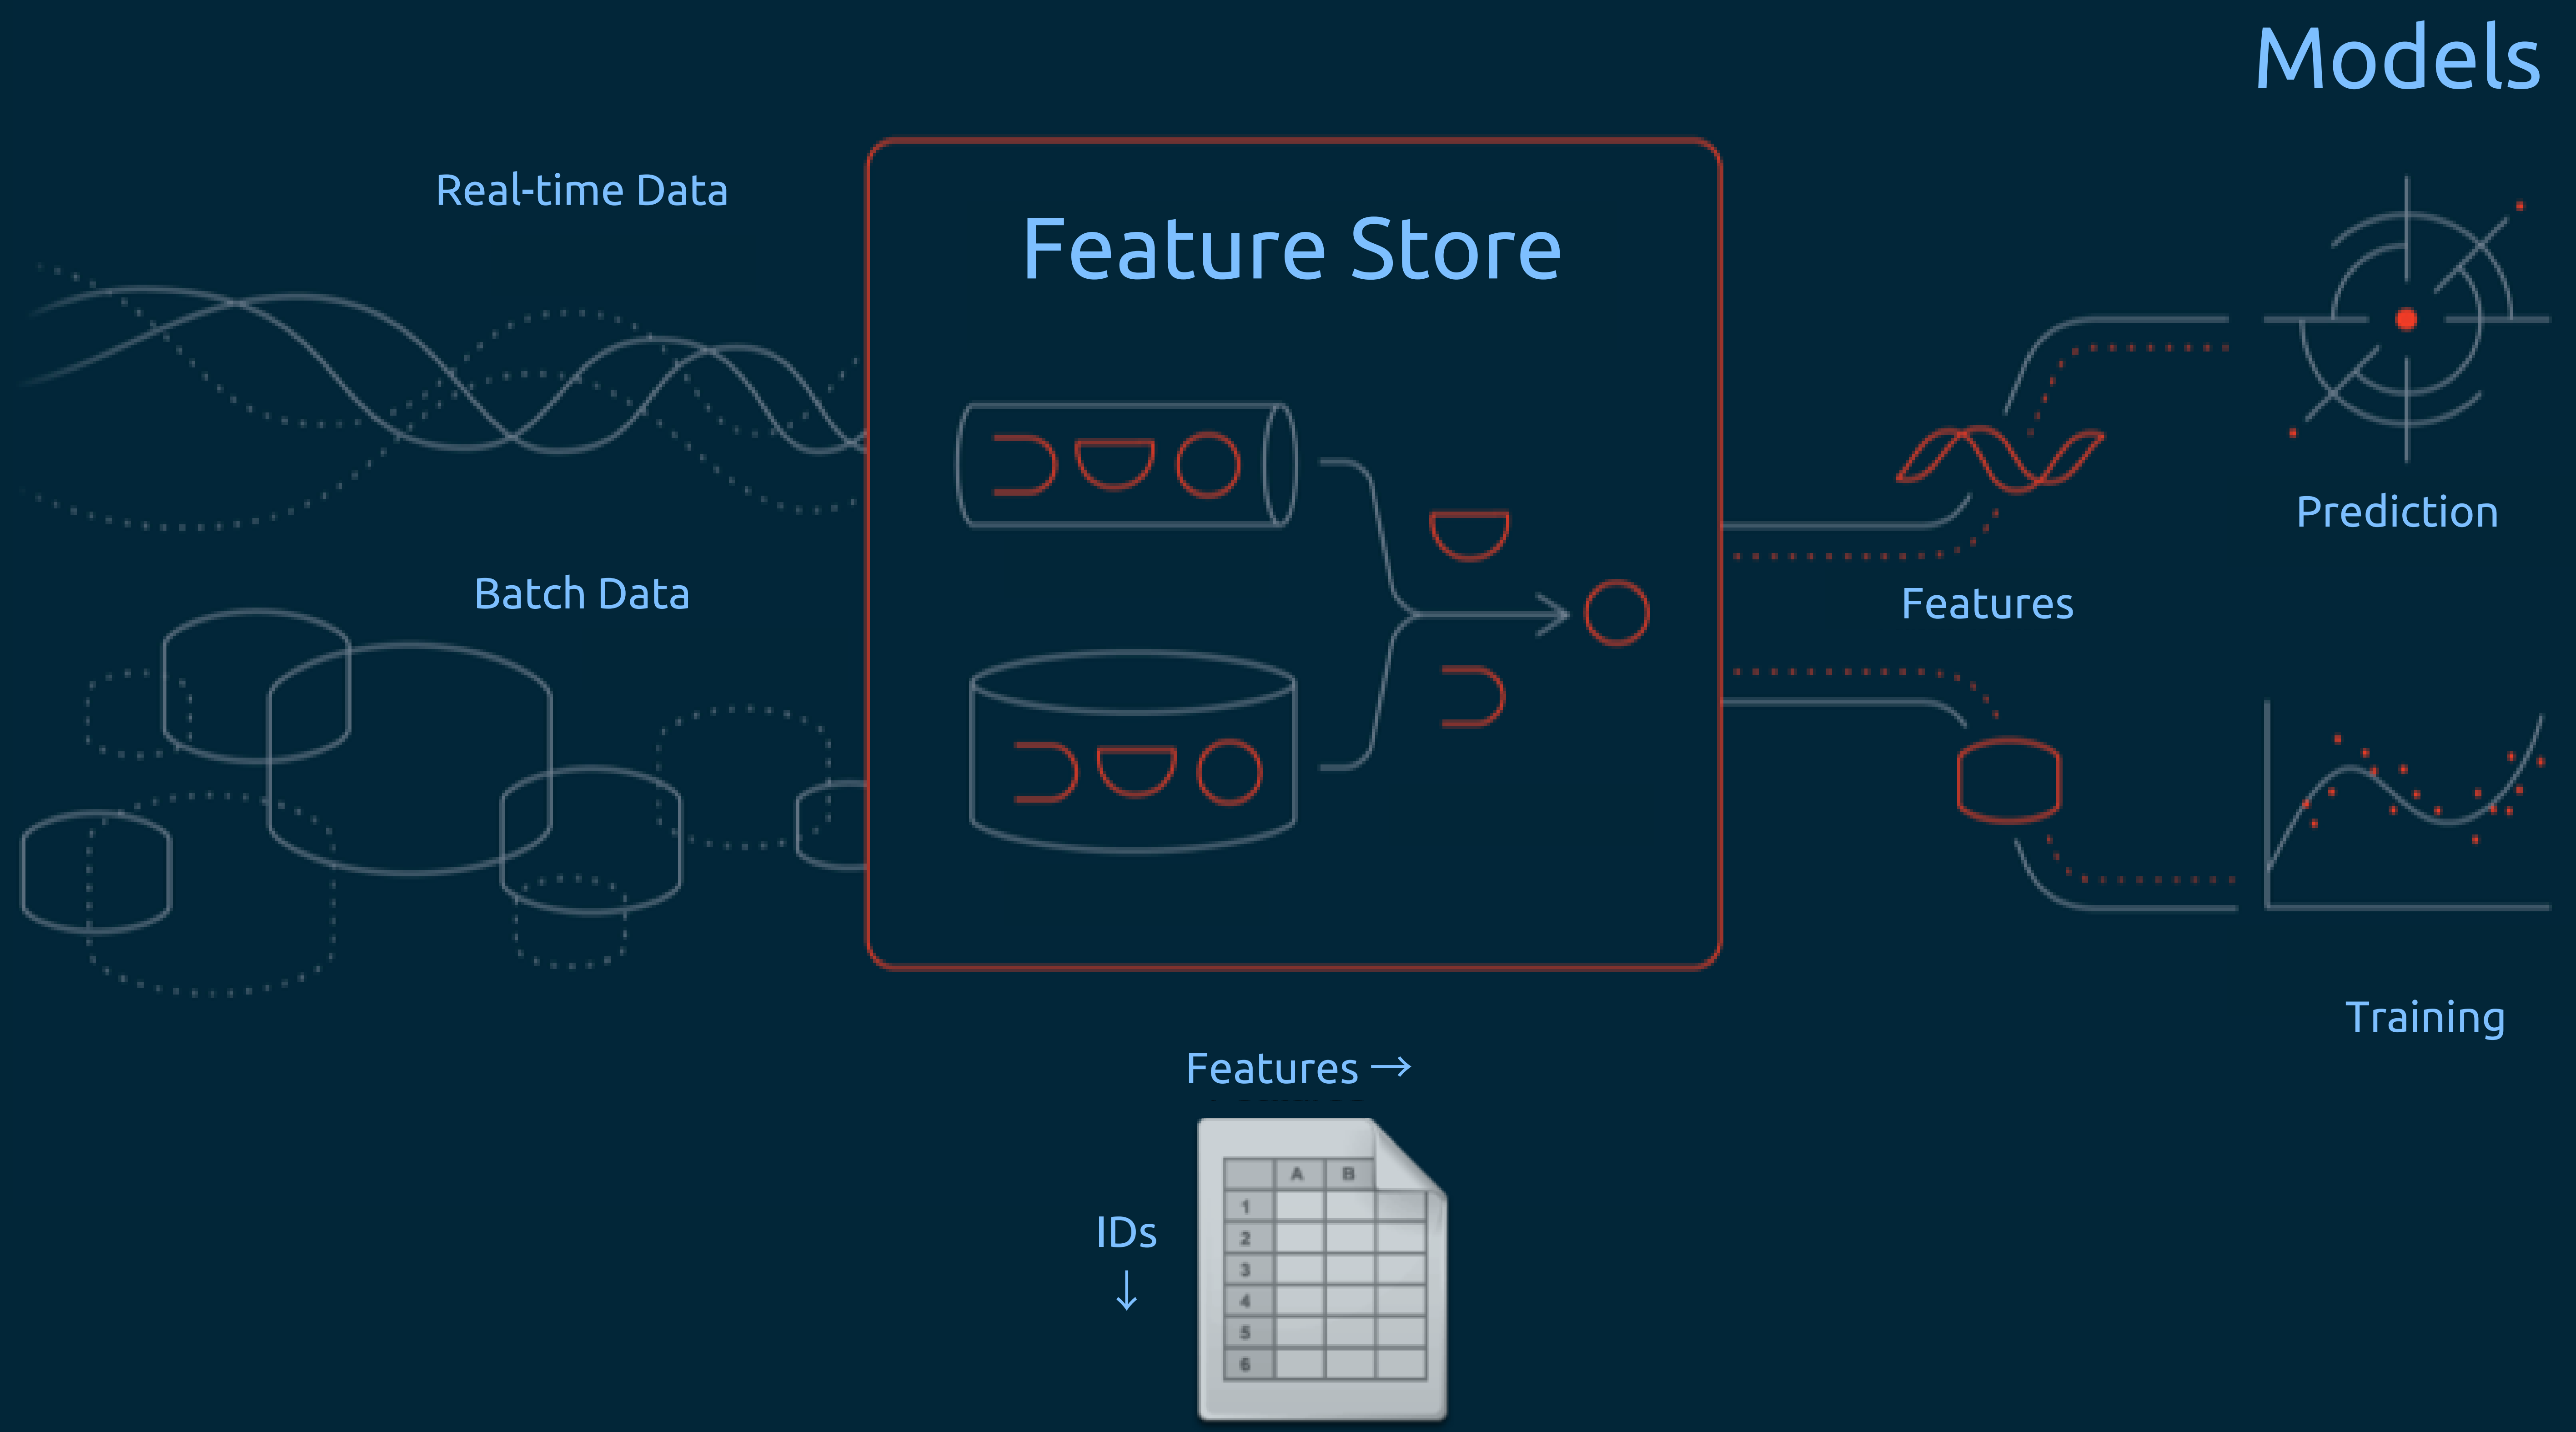

Design Pattern 25

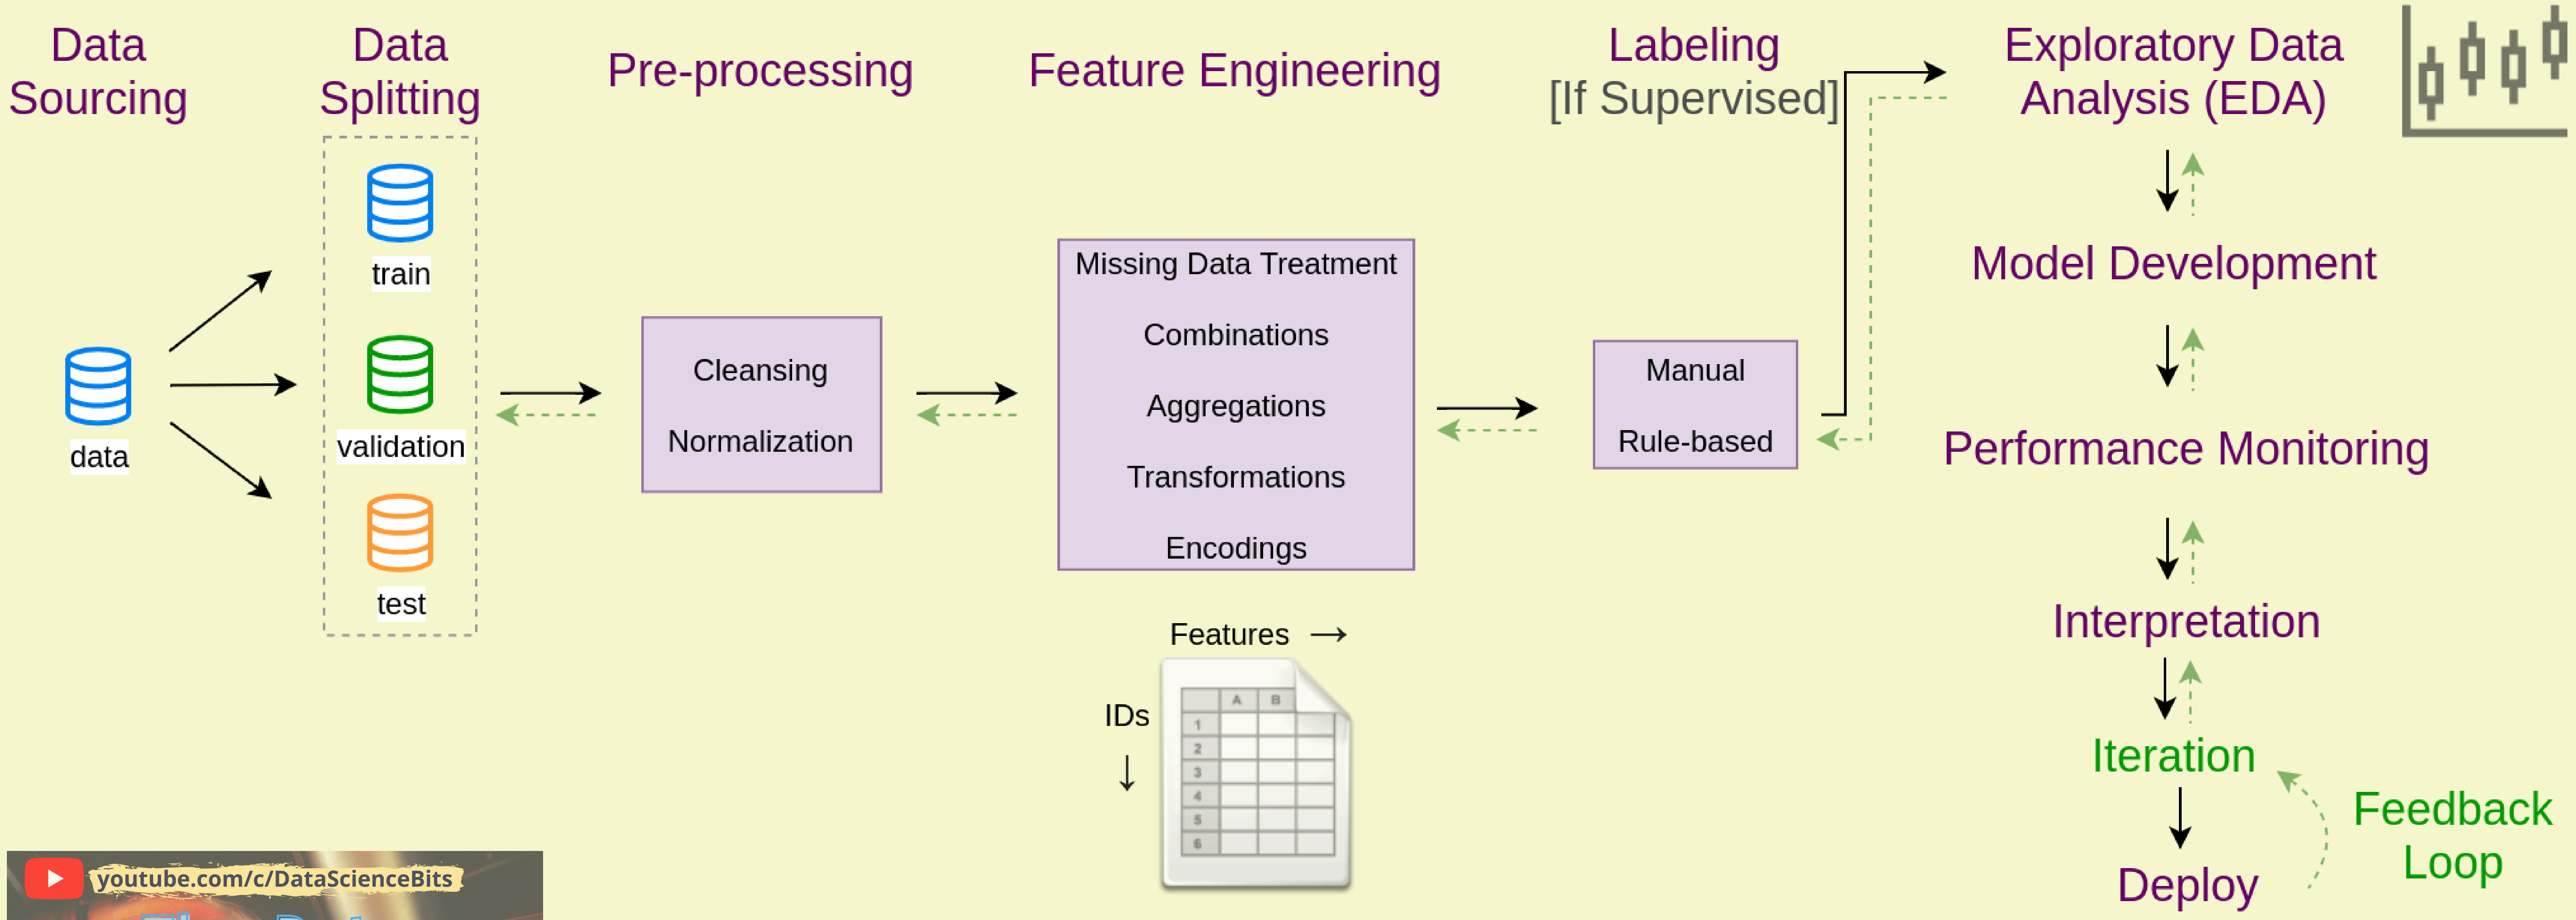

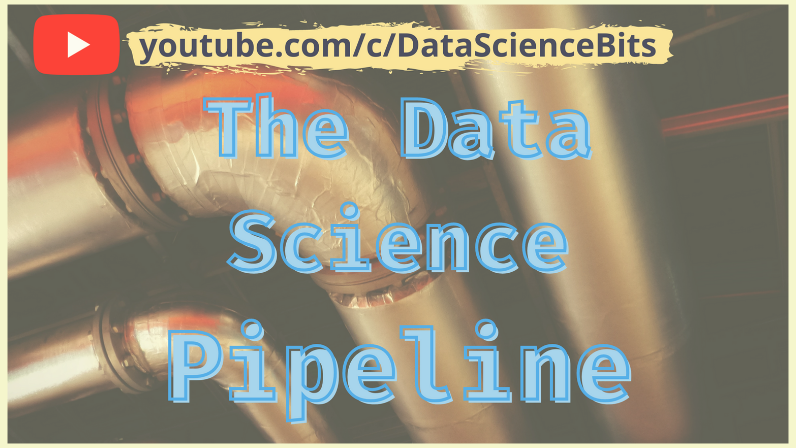

GREENHOUSE

A CONTAINERIZED FRAMEWORK FOR BETTER DATA X

github.com/felipepenha/py-greenhouse

twitch.tv/DataScienceBits

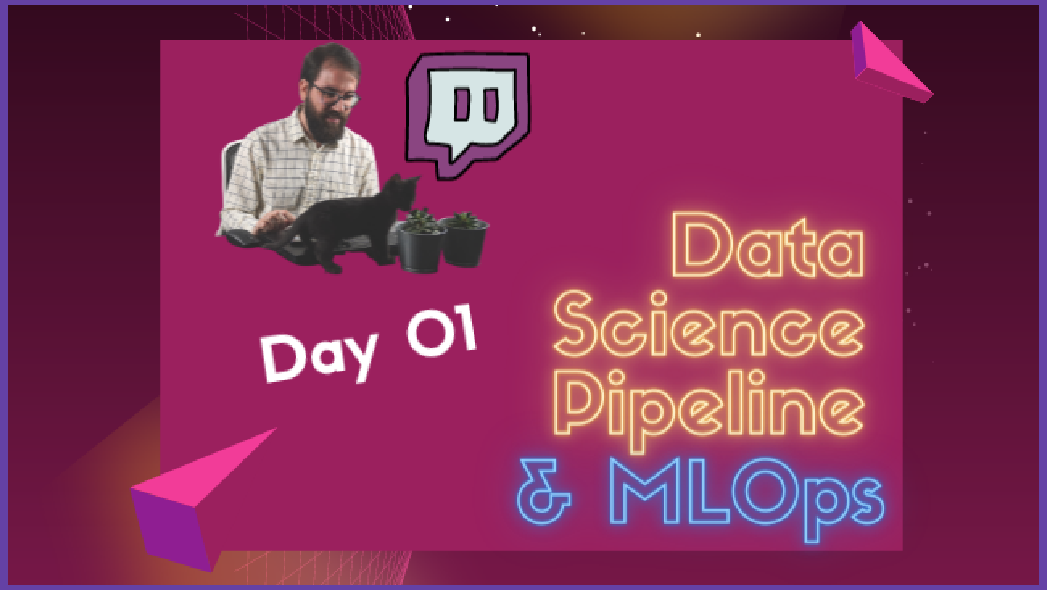

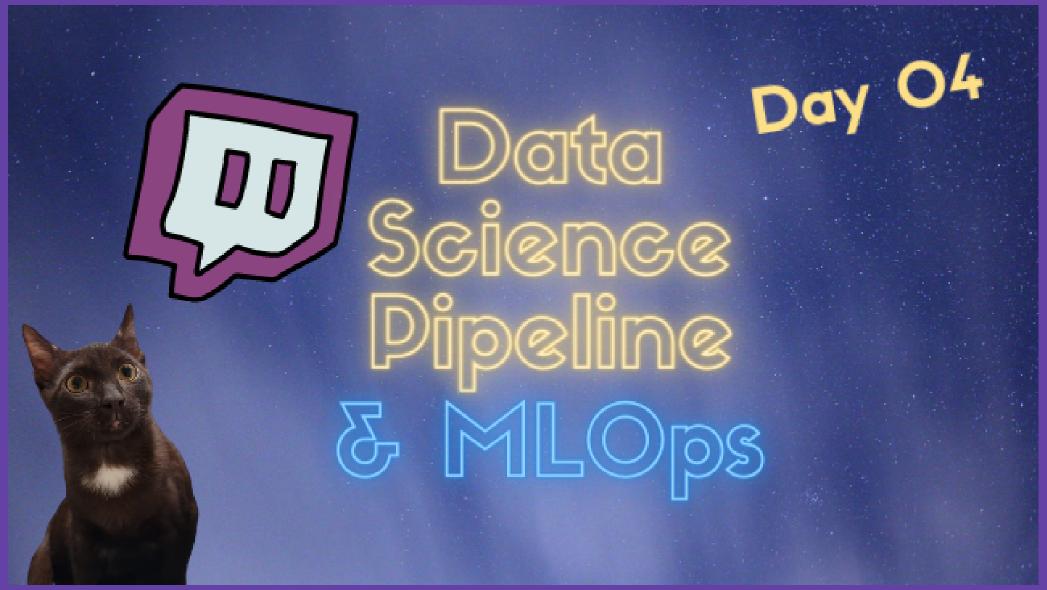

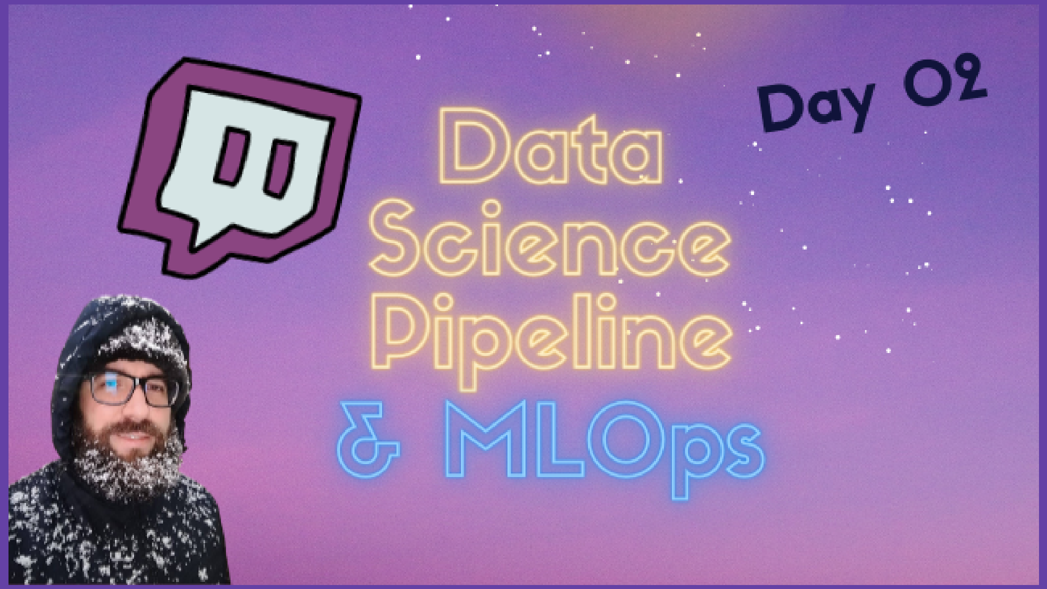

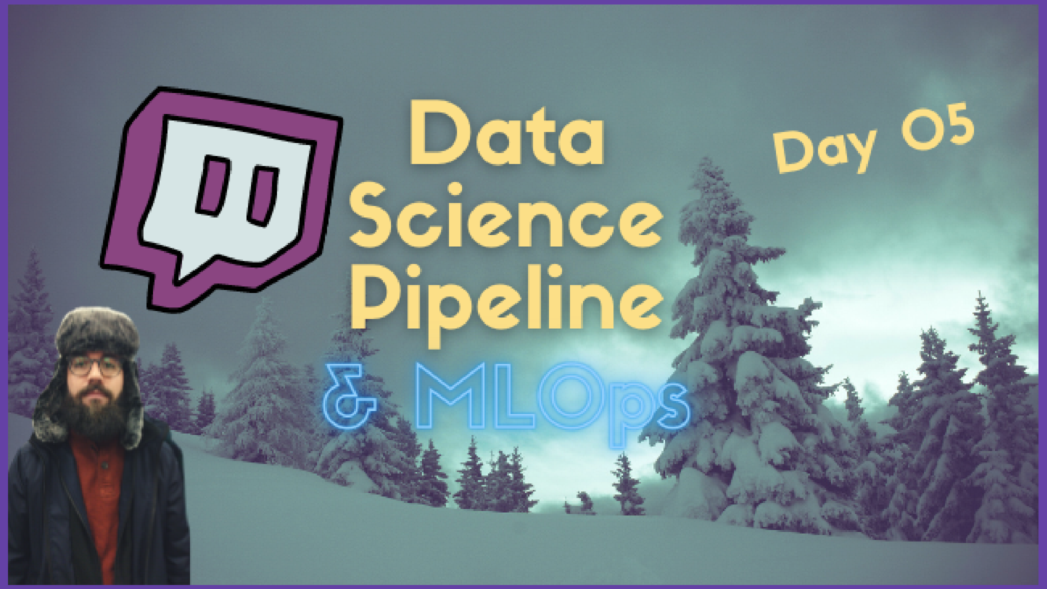

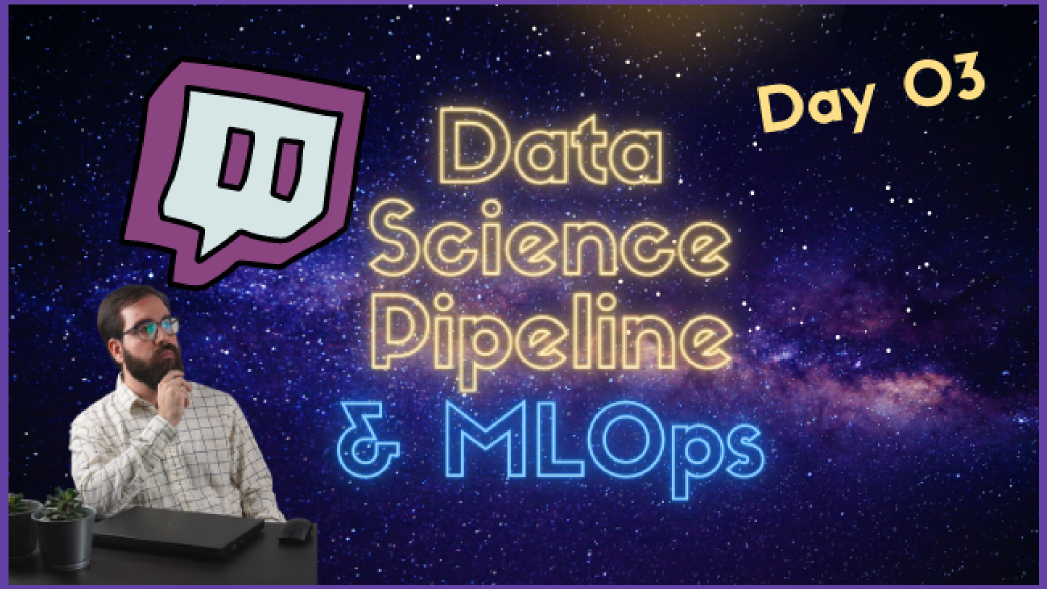

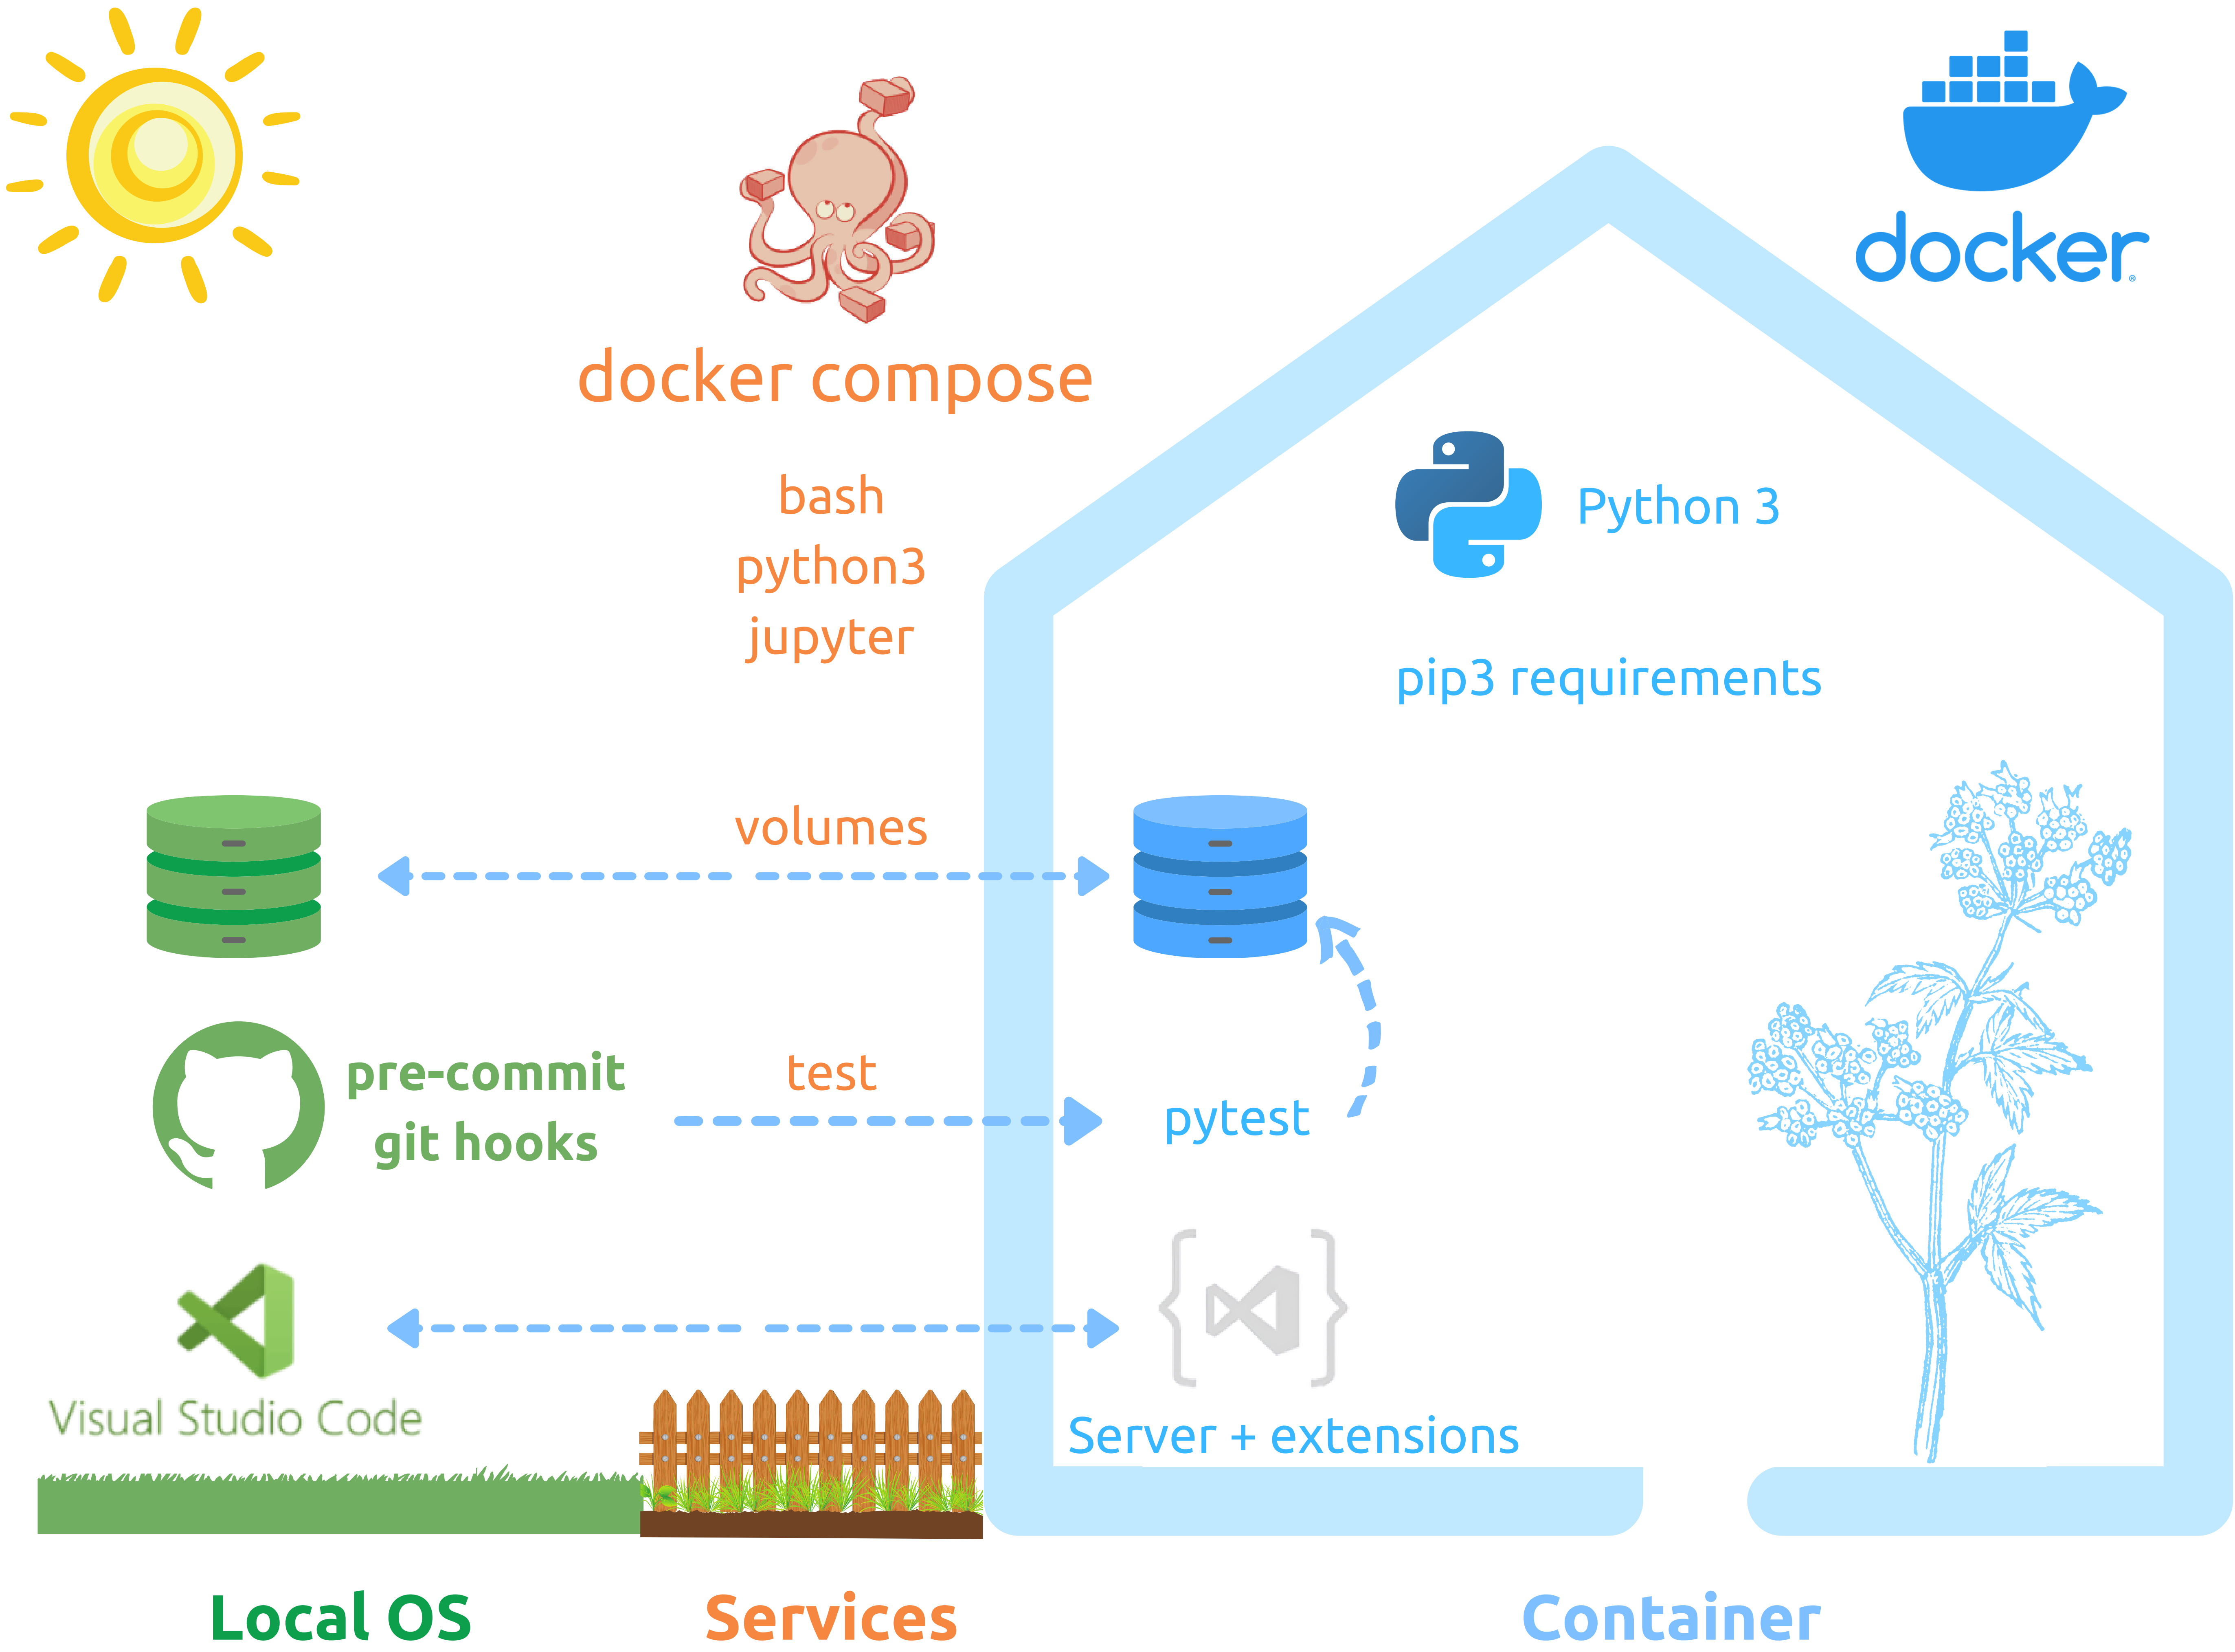
docker
docker compose
bash
python3
jupyter
Python 3
pip3 requirements
volumes
pre-commit
git hooks
test
pytest
Visual Studio Code
Server + extensions
Local OS
Services
Container

GitOps

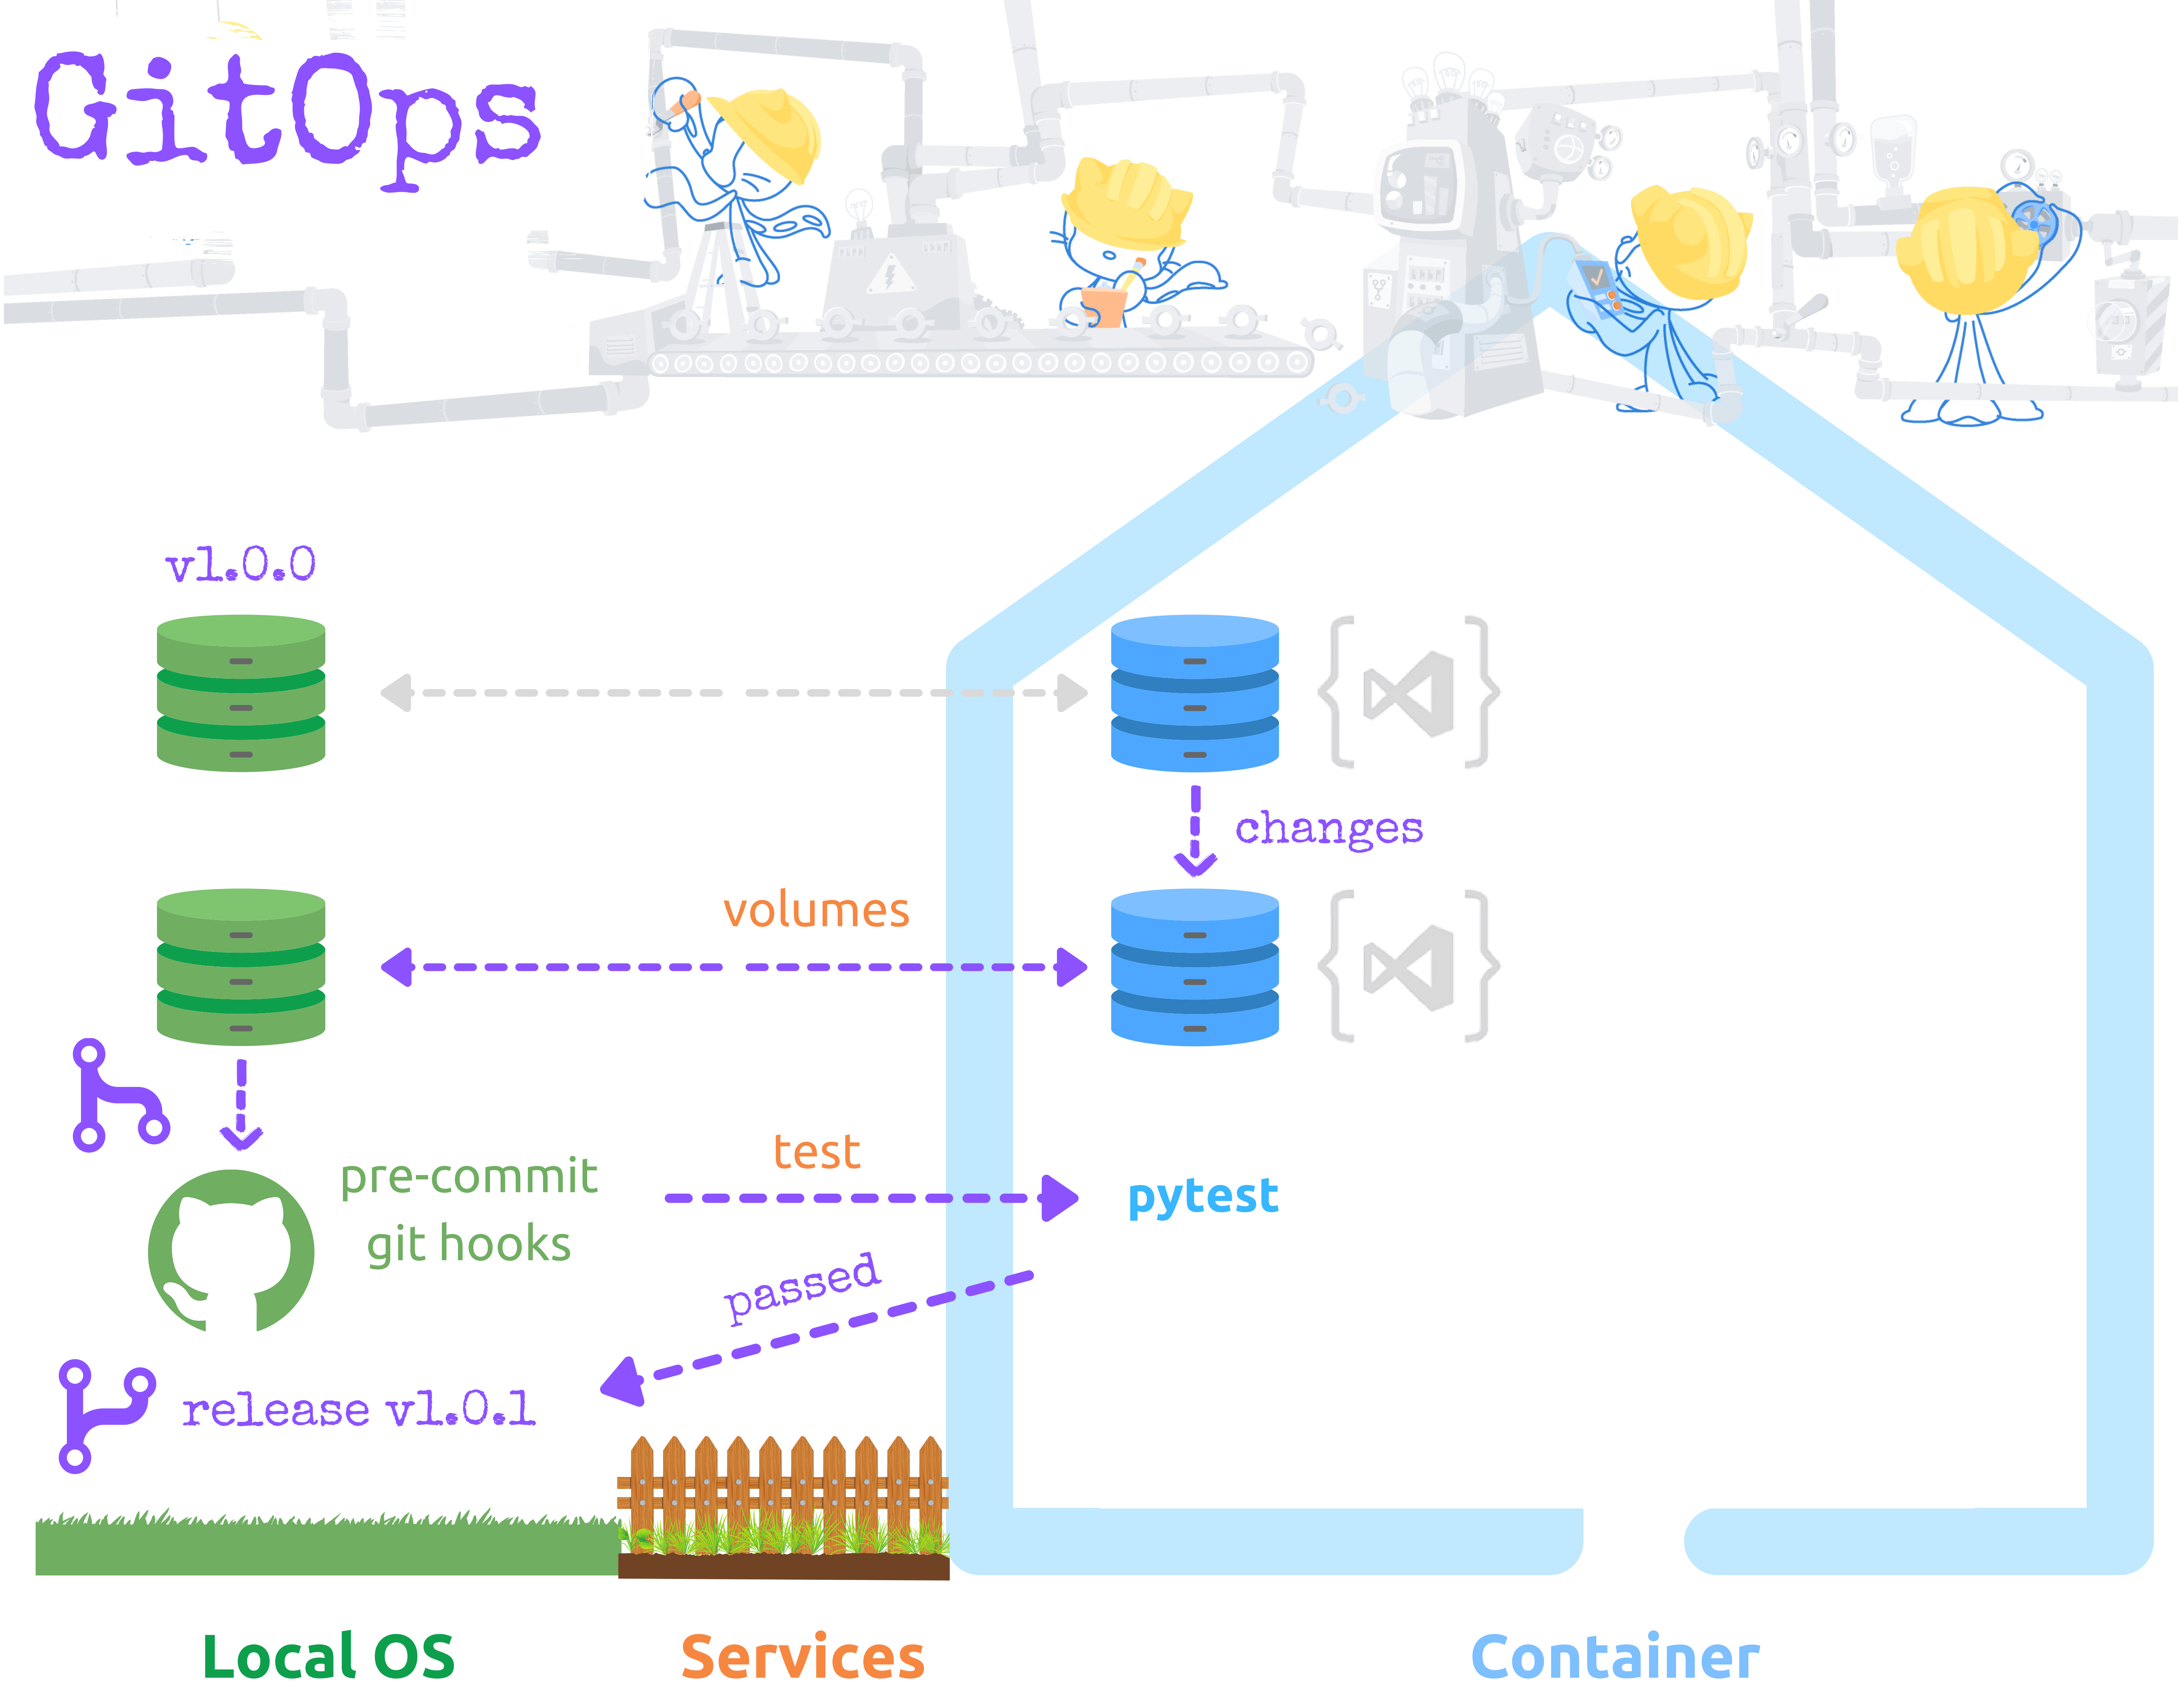
v1.0.0
changes
volumes
pre-commit
git hooks
test
pytest
passed
release v1.0.1
Local OS
Services
Container

https://linktr.ee/felipepenha

Muito Obrigado!

Fique Ligadx!
12 de Maio
Participação no Podcast

MLOps.community

A place to discuss MLOps
An open community where all are welcome

Slack com 2.5k+ usuários