AI, Machine Learning and Deep Learning for Quantitative Finance with MATLAB

Yi Wang, Application Engineering Manager
MathWorks, Inc. USA

Agenda

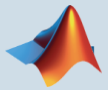

Introduction and Challenges

- Machine Learning
  - Credit Scoring
- Model Management
- Text Analytics and Deep Learning
  - Twitter Sentiment
- MATLAB-Python Integration
- App Building and Deployment
- Summary

AI, Machine Learning and Deep Learning

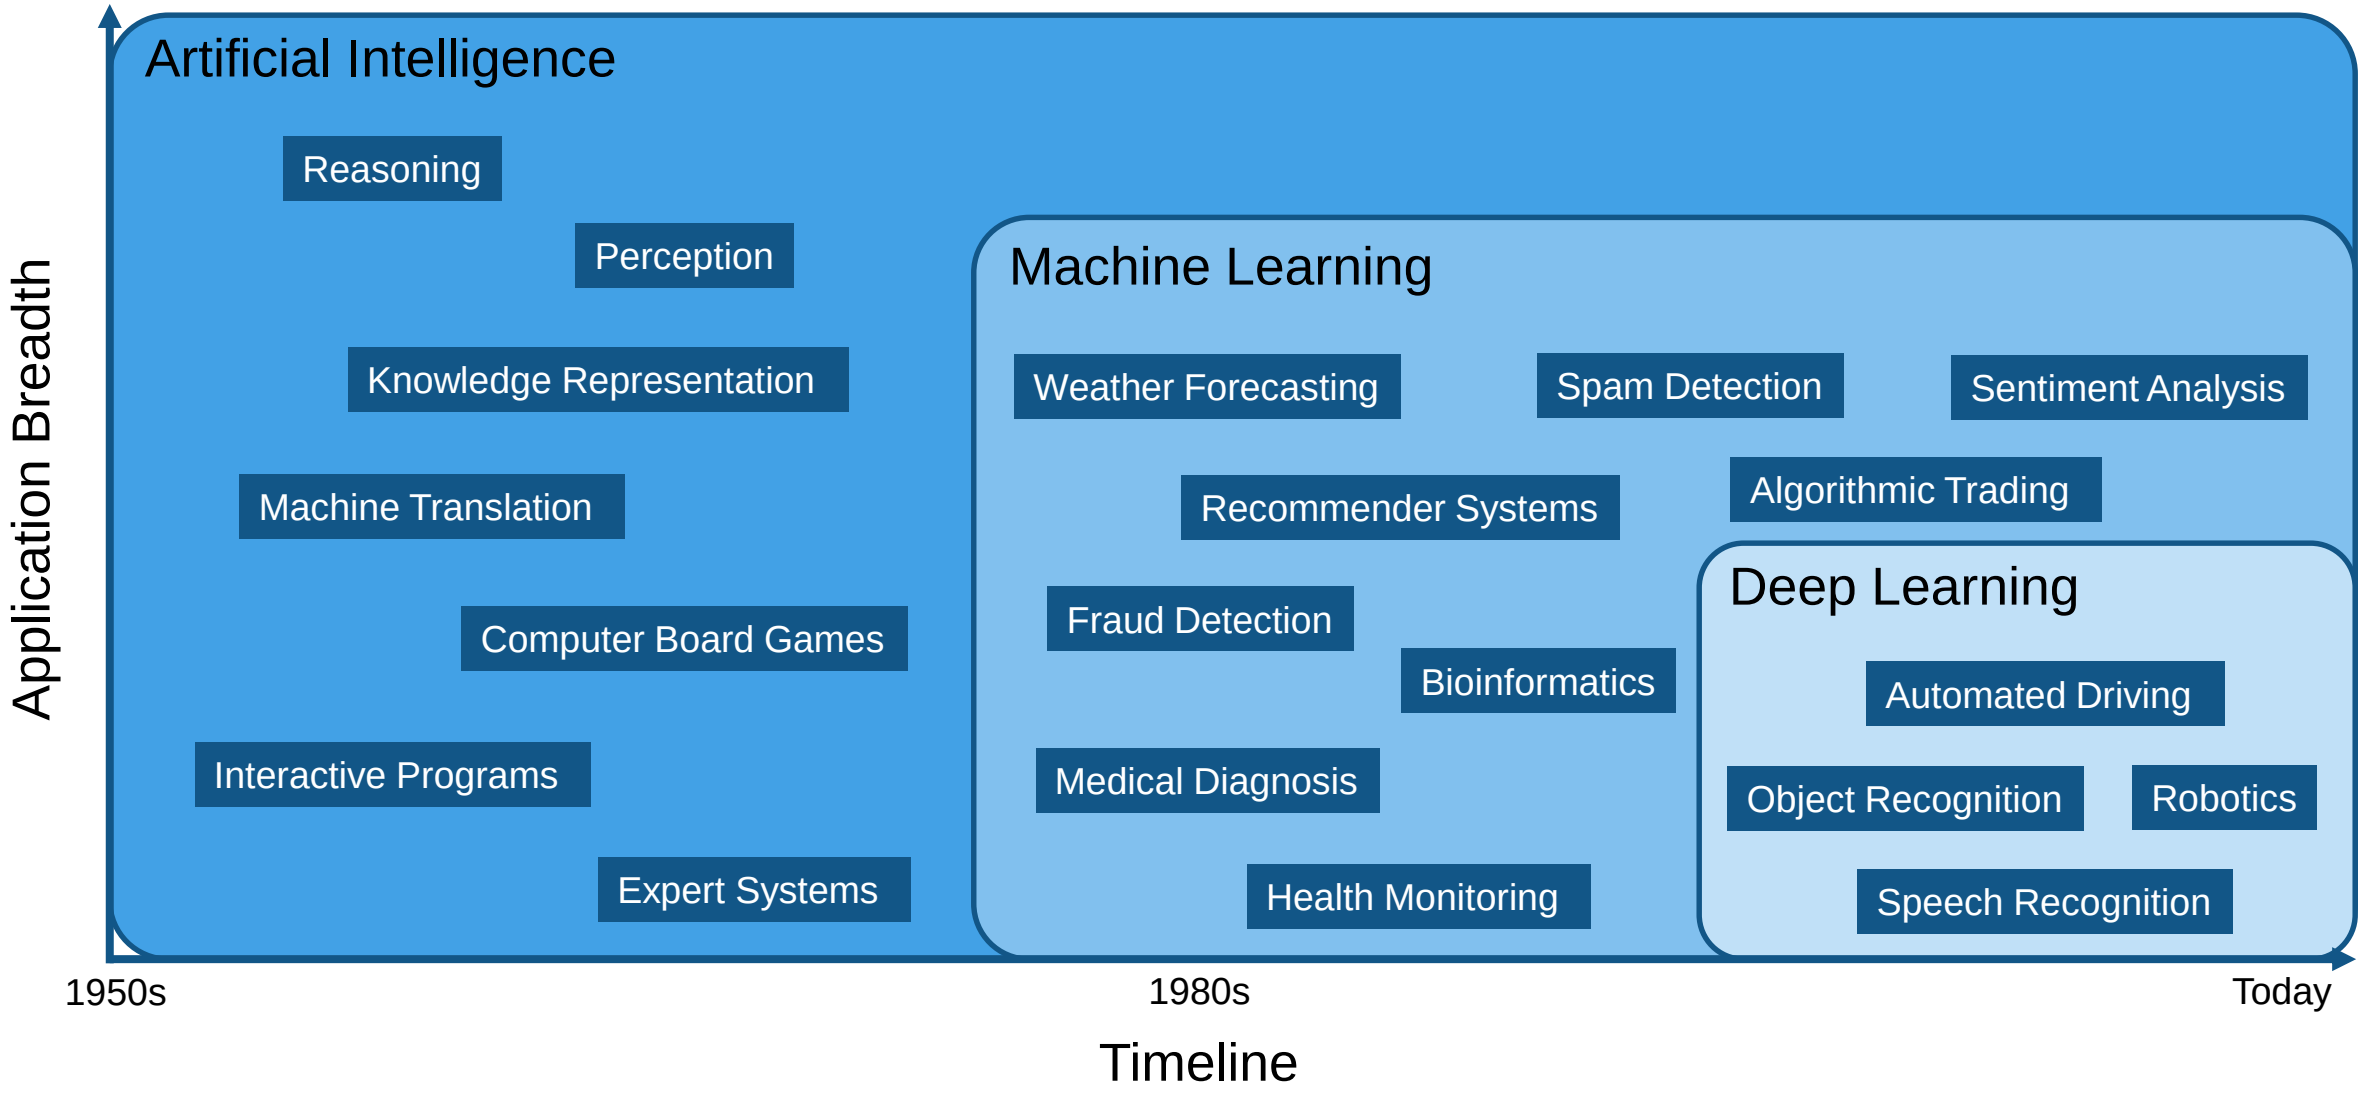

Machine Learning in Practice

Standard Approach

Computer Program

Buy

Hold

Sell

Hand Written Program

If RSI > 70
then “SELL”
If MACD > SIG and RSI <= 70
then “HOLD”
…

Formula or Equation

$$Y_{\text{Trade}} = \beta_1 X_{\text{RSI}} + \beta_2 X_{MACD} + \beta_3 X_{\text{TSMom}} + \dots$$

Machine Learning Approach

Machine Learning

Buy

Hold

Sell

$model$: Inputs → Outputs

$$model = <\begin{matrix}\textbf{Machine}\\ \textbf{Learning}\\ \textbf{Algorithm}\end{matrix}>(factors, trade\ decision)$$

Challenges

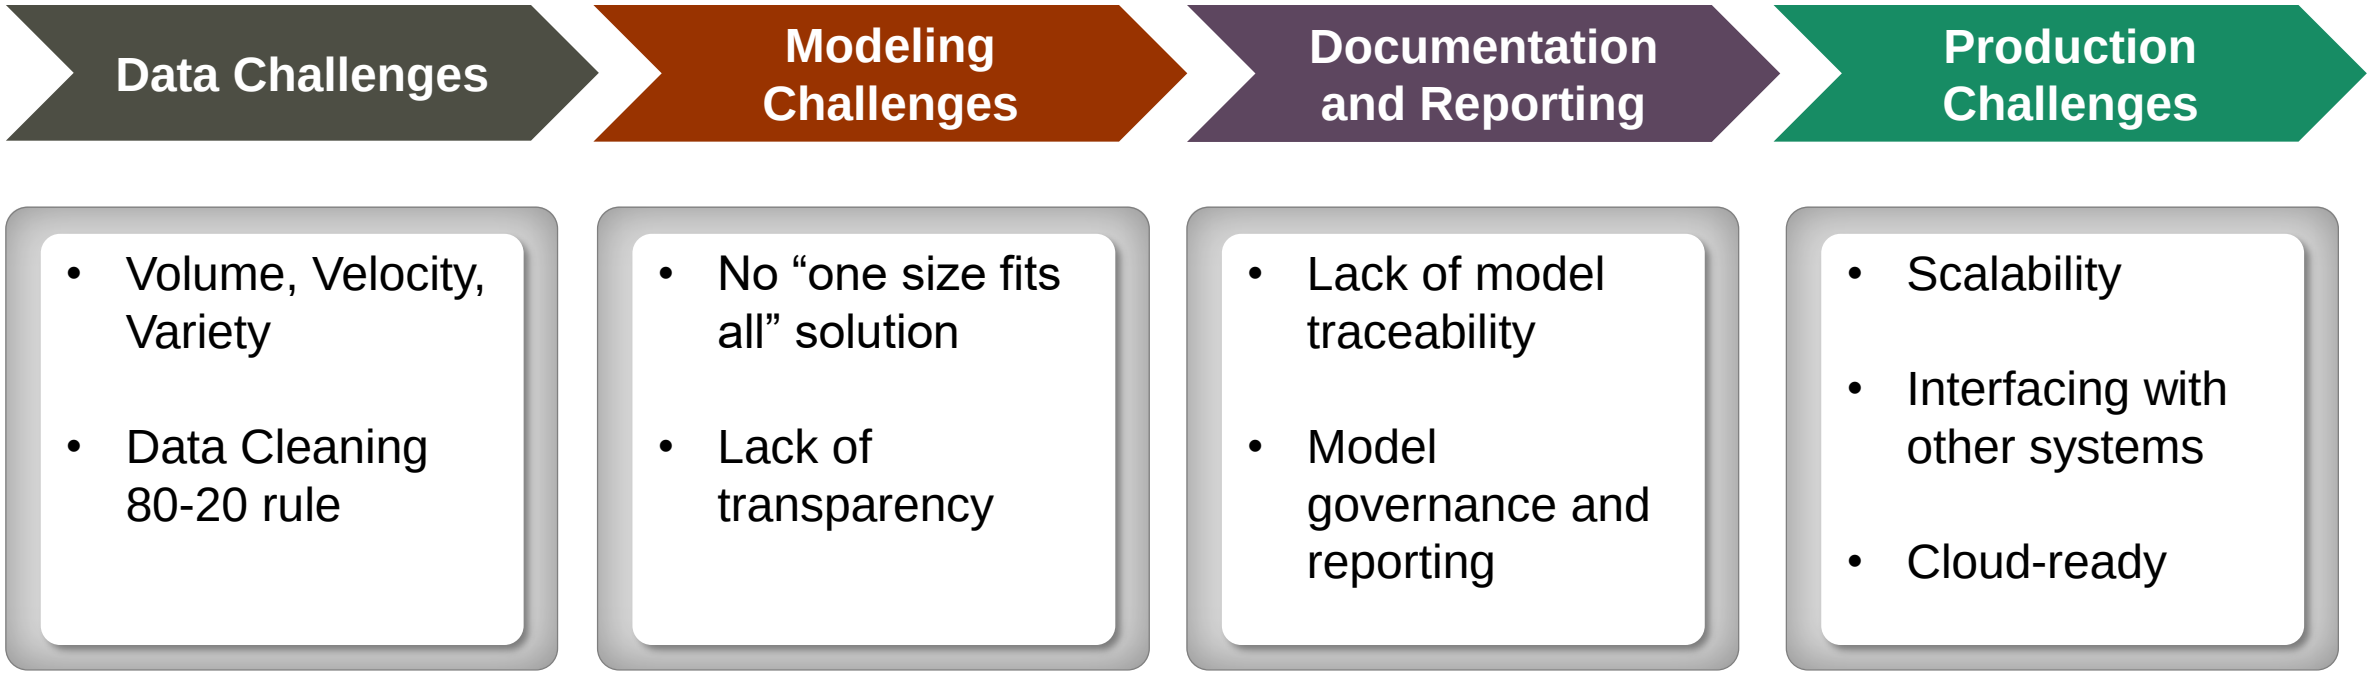

Model Management

Agenda

- Introduction and Challenges
- 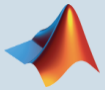 Machine Learning
  - Credit Scoring
- Model Management
- Text Analytics and Deep Learning
  - Twitter Sentiment
- MATLAB-Python Integration
- App Building and Deployment
- Summary

Consumer Credit Risk Modeling

- Create Credit Scorecard to predict Default
- Build a Challenger Machine Learning model

A1 | LoanAmnt

	A	B	C	D	E	F	G
1	LoanAmnt	FundedAmnt	FundedAmntInv	Term	Installment	EmpLength	AnnualInc
2	$15,000	$15,000	$15,000	36	$556	2	$420,000
3	$15,000	$15,000	$15,000	36	$556	2	$420,000
4	$24,000	$24,000	$23,250	36	$800	10	$200,004
5	$24,000	$24,000	$23,250	36	$800	10	$200,004
6	$15,000	$15,000	$14,593	36	$497	8	$189,996
7	$22,500	$22,500	$10,863	36	$774	10	$227,000
8	$21,600	$21,600	$21,002	36	$745	4	$240,000
9	$21,600	$21,600	$21,002	36	$745	4	$240,000
10	$20,000	$20,000	$7,497	36	$665	10	$290,004
11	$18,000	$18,000	$17,875	36	$624	7	$300,000
12	$25,000	$25,000	$24,975	36	$837	10	$102,000
13	$25,000	$25,000	$24,975	36	$837	10	$102,000
14	$28,000	$28,000	$27,975	60	$731	10	$200,000
15	$20,000	$20,000	$19,711	60	$512	10	$305,000
16	$20,000	$20,000	$20,000	60	$503	4	$145,000

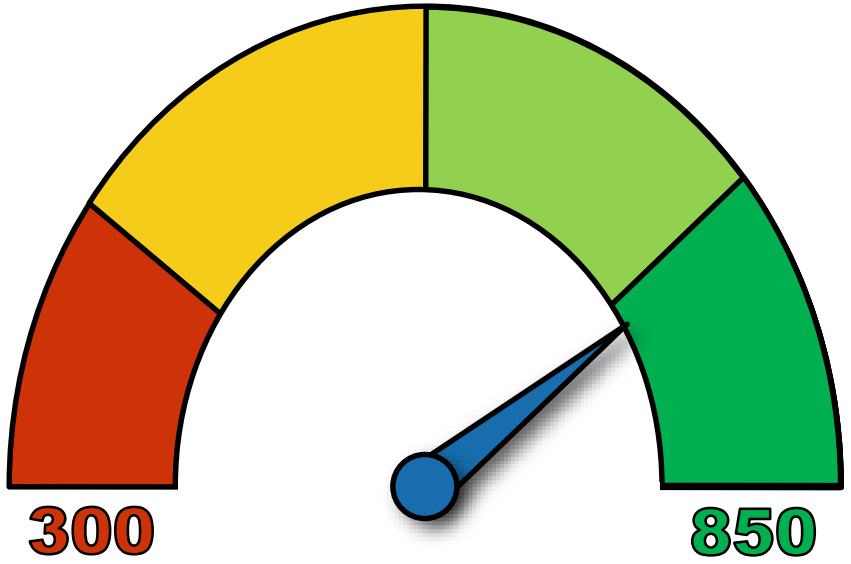

Default / Non-Default

Model Review and Validation techniques for Machine Learning

- Review of model training/testing procedure
- Determination of whether model hyperparameters are appropriate
- Calculation of metrics for determining model accuracy and robustness
- Building traditional models to test challenger models against

Summary of Solutions

- Rich visualizations to extract insight quickly
- Easily tune/fit logistic and other complex Machine Learning models
- Build prototypes and scale without recoding
- Automate reporting

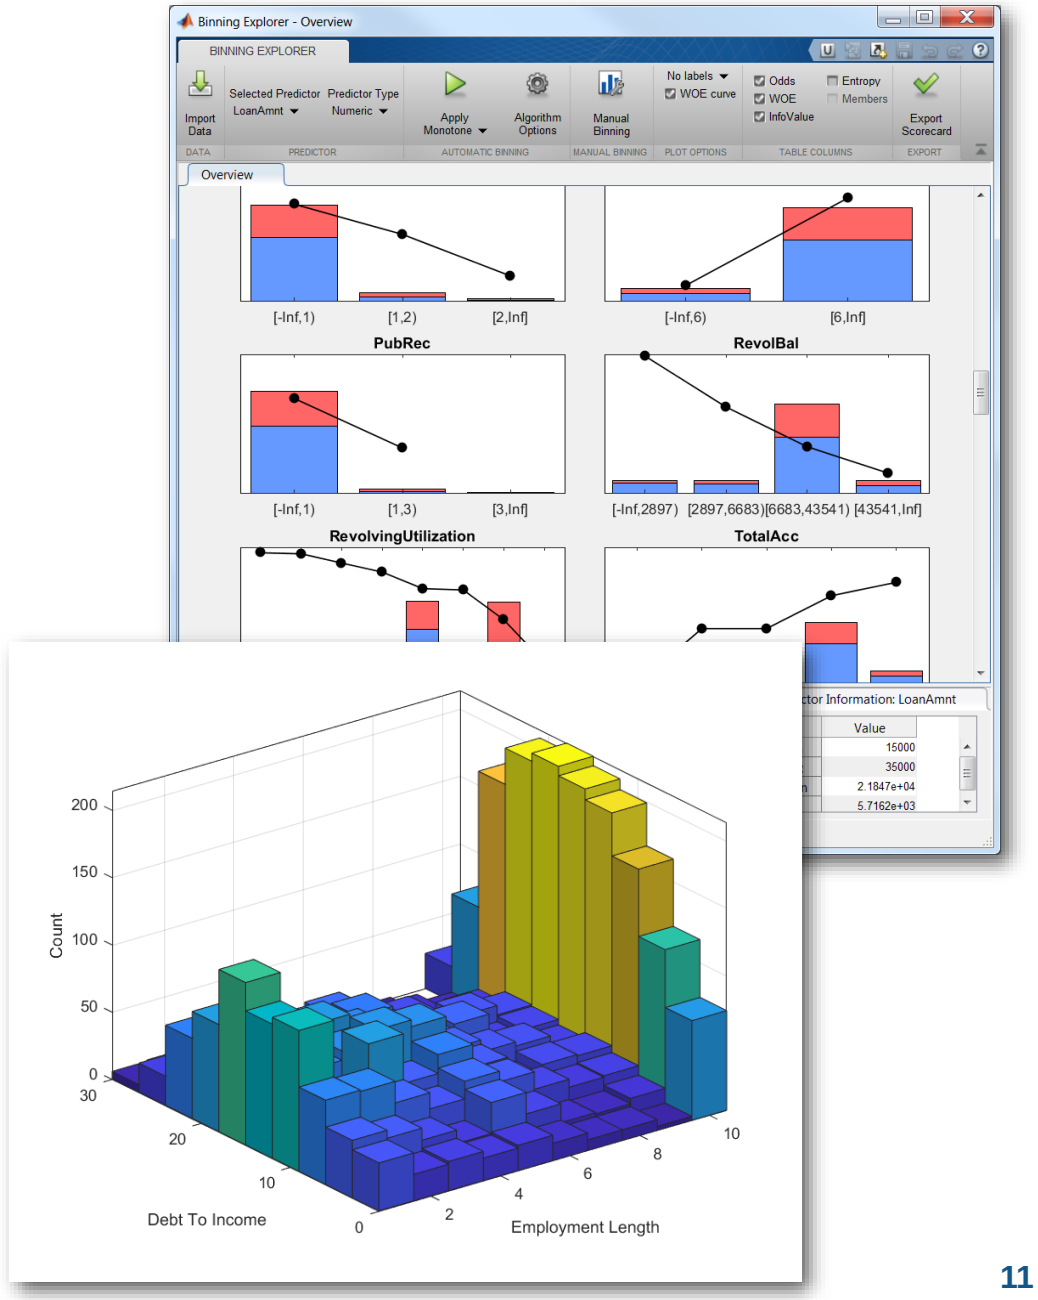

Agenda

- Introduction and Challenges
- Machine Learning
  - Credit Scoring
- 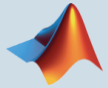 Model Management
- Text Analytics and Deep Learning
  - Twitter Sentiment
- MATLAB-Python Integration
- App Building and Deployment
- Summary

Model Governance Lifecycle

Risk management
Board and stakeholders
Regulator
Model Owners

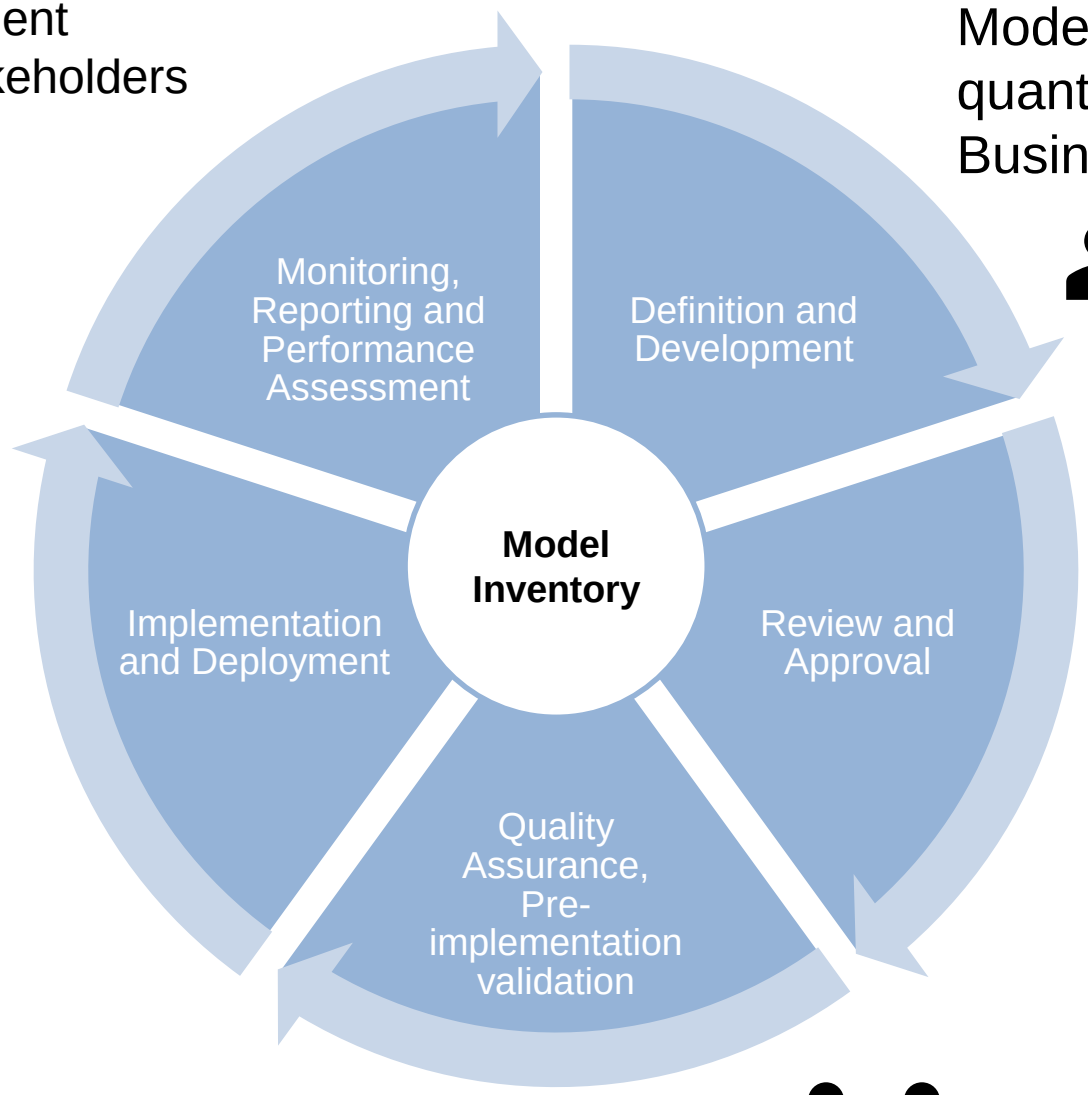

Model developers,
quants, analysts.
Business lines

IT, front-office
End-users

Independent model
review and audit
Regulator

Model Validation, IT

Agenda

- Introduction and Challenges
- Machine Learning
  - Credit Scoring
- Model Management
- 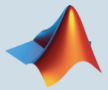 Text Analytics and Deep Learning
  - Twitter Sentiment
- MATLAB-Python Integration
- App Building and Deployment
- Summary

Demo: Sentiment Analysis of Tweets

Goal:

- Analyze the sentiment of tweets for several securities and compare the sentiment scores to stock prices

```
ans = 508×1 string array
    "Walmart: "you wanna destroy Amazon?" Google: "bet" $WMT $GOOG
    "$WMT wants next level customer service w/highly personalized
    "Ironic prelude to $DIS buying $TWTR soon IMO $AAPL $GOOG $SPY
    "$AMZN the $WMT threat grows each and every day  https://t.co/
    "MU Investments Co. Ltd. Sells 30 Shares of Alphabet Inc. $GOO
    "Ad $ are going to $GOOG and $FB away from wppgy  #Advertising
    "Big bullish unusual option activity detected: $SPX, $GOOG, $C
    "REPORT: Apple to build data center in Iowa: https://t.co/jwH6
    "RT @theflynews: REPORT: Apple to build data center in Iowa: h
```

Approach

- Access data from Twitter
- Preprocess to clean-up text and deal with domain-specific terms
- Predict sentiment from word embedding
- Compare sentiment scores to prices

Workflow

Tweets

```
ans = 508×1 string array
    "Walmart: "you wanna destroy Amazon?" Google: "bet" $WMT $GOOG
    "$WMT wants next level customer service w/highly personalized
    "Ironic prelude to $DIS buying $TWTR soon IMO $AAPL $GOOG $SPY
    "$AMZN the $WMT threat grows each and every day  https://t.co/
    "MU Investments Co. Ltd. Sells 30 Shares of Alphabet Inc. $GOO
    "Ad $ are going to $GOOG and $FB away from wppgy  #Advertising
    "Big bullish unusual option activity detected: $SPX, $GOOG, $C
    "REPORT: Apple to build data center in Iowa: https://t.co/jwH6
    "RT @theflynews: REPORT: Apple to build data center in Iowa: h
```

Scoring Algorithm

```
scoreTweet(tweet,emb,svmModel)
```

Score

Positive + Negative Word List

pos — 2006x1 string

	1
1	a+
2	abound
3	abounds
4	abundance
5	abundant
6	accessable
7	accessible

neg — 4783x1 string

	1
1	2-faced
2	2-faces
3	abnormal
4	abolish
5	abominable
6	abominably
7	abominate

Machine Learning Model

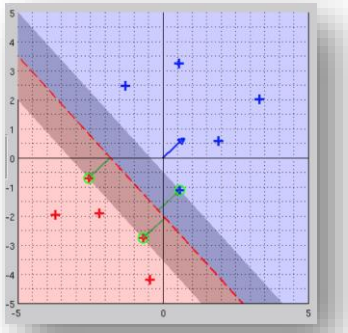

Word Embedding

```
wordEmbedding with properties:

   Dimension: 100
  Vocabulary: [1×1193514 string]
```

What is Deep Learning?

- Type of machine learning in which a model learns to perform classification and regression tasks directly from images, text, or sound.
- Usually implemented using a neural network architecture.
- The term “deep” refers to the number of layers in the network—the more layers, the deeper the network.

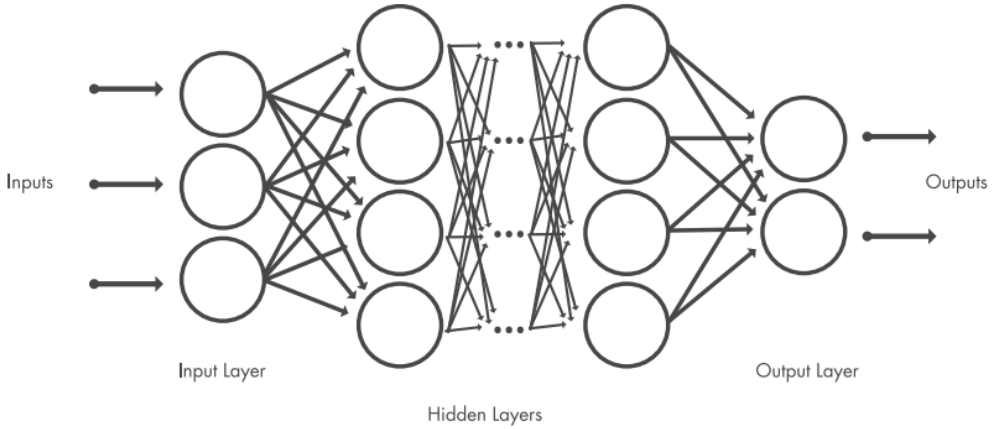

Why is Deep Learning So Popular Now?

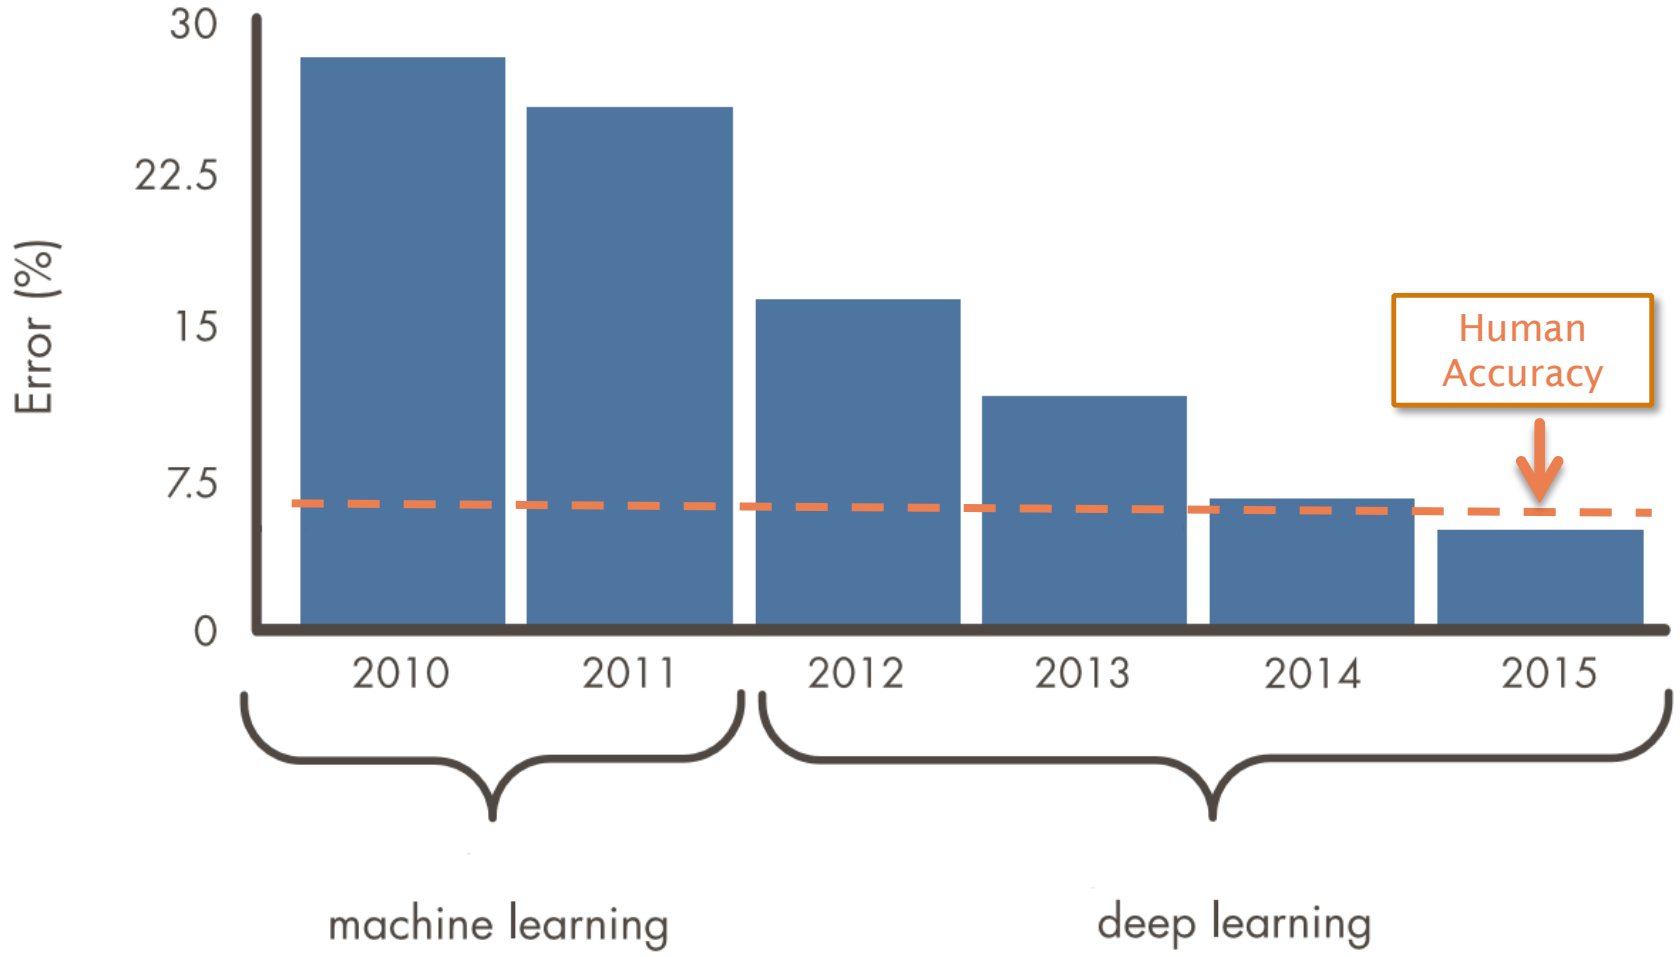

Source: ILSVRC Top-5 Error on ImageNet

Machine Learning vs Deep Learning

Deep learning performs **end-to-end learning** by learning **features, representations and tasks** directly from **images, text, and signals**

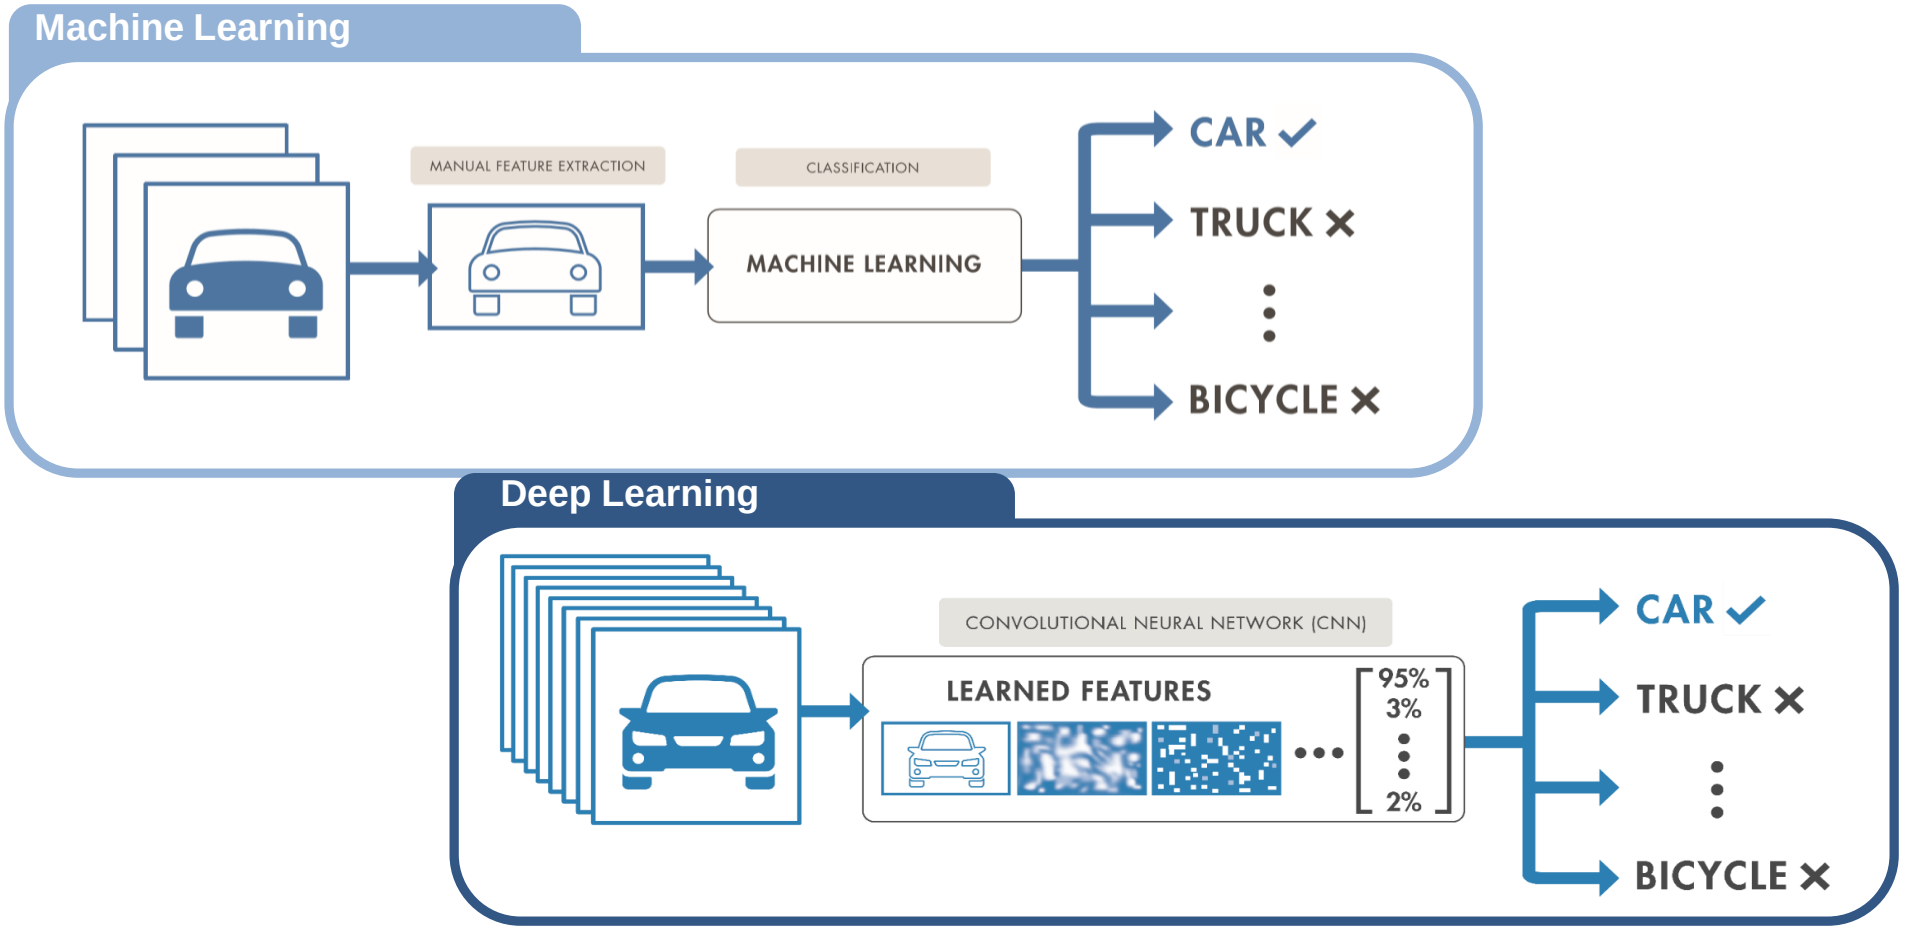

MATLAB Differentiators for AI / Deep Learning

- **MATLAB**
  - makes it easy to learn and use deep learning techniques
  - provides complete workflow from research to prototype to production (enterprise or embedded systems)
- **It enables analysts** to
  - **Access pretrained models** from Caffe and TensorFlow-Keras
  - **Automate ground-truth labeling** with Apps
  - **Visualize** intermediate results and **debug** deep learning models
  - **Accelerate model training** using NVidia GPUs, Cloud and Clusters
  - **Automatically convert** deep learning models to CUDA or C code for cloud or embedded deployment

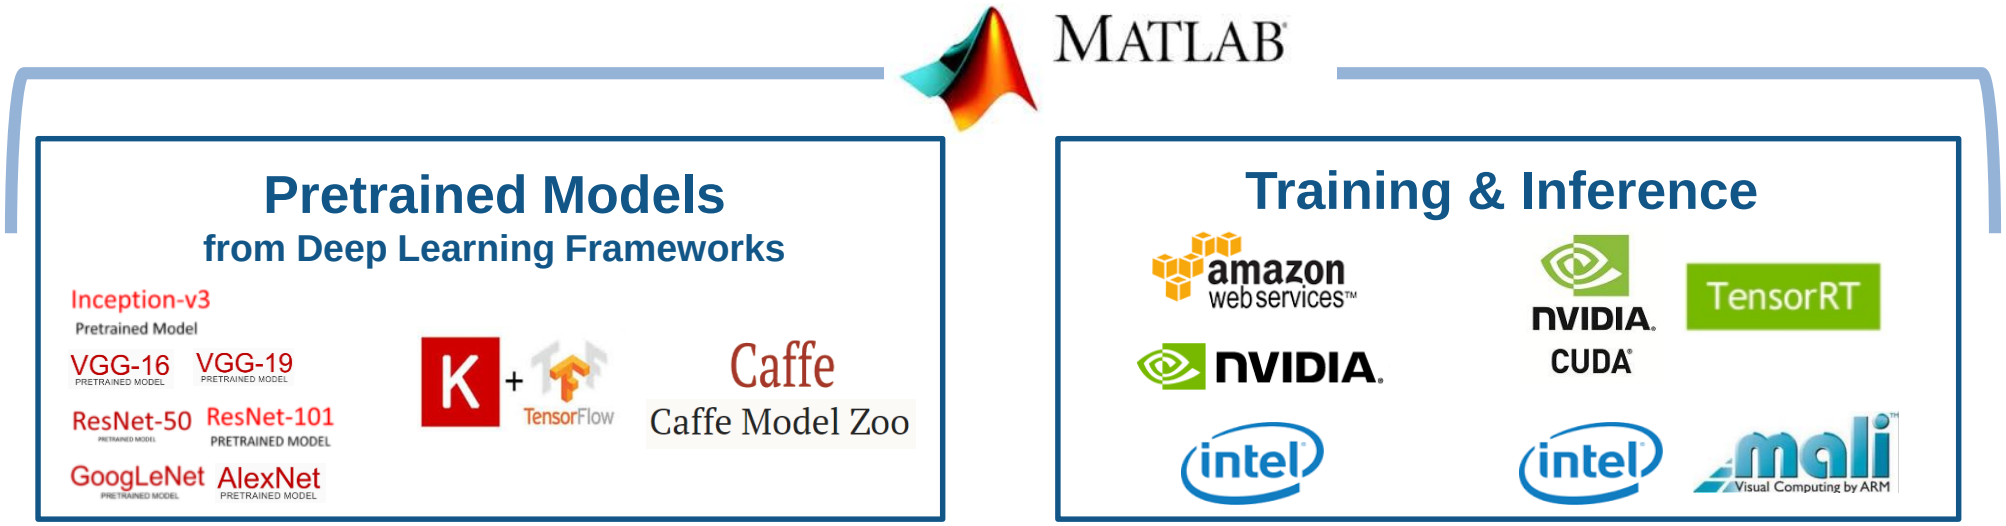

With GPU Coder, MATLAB is fast

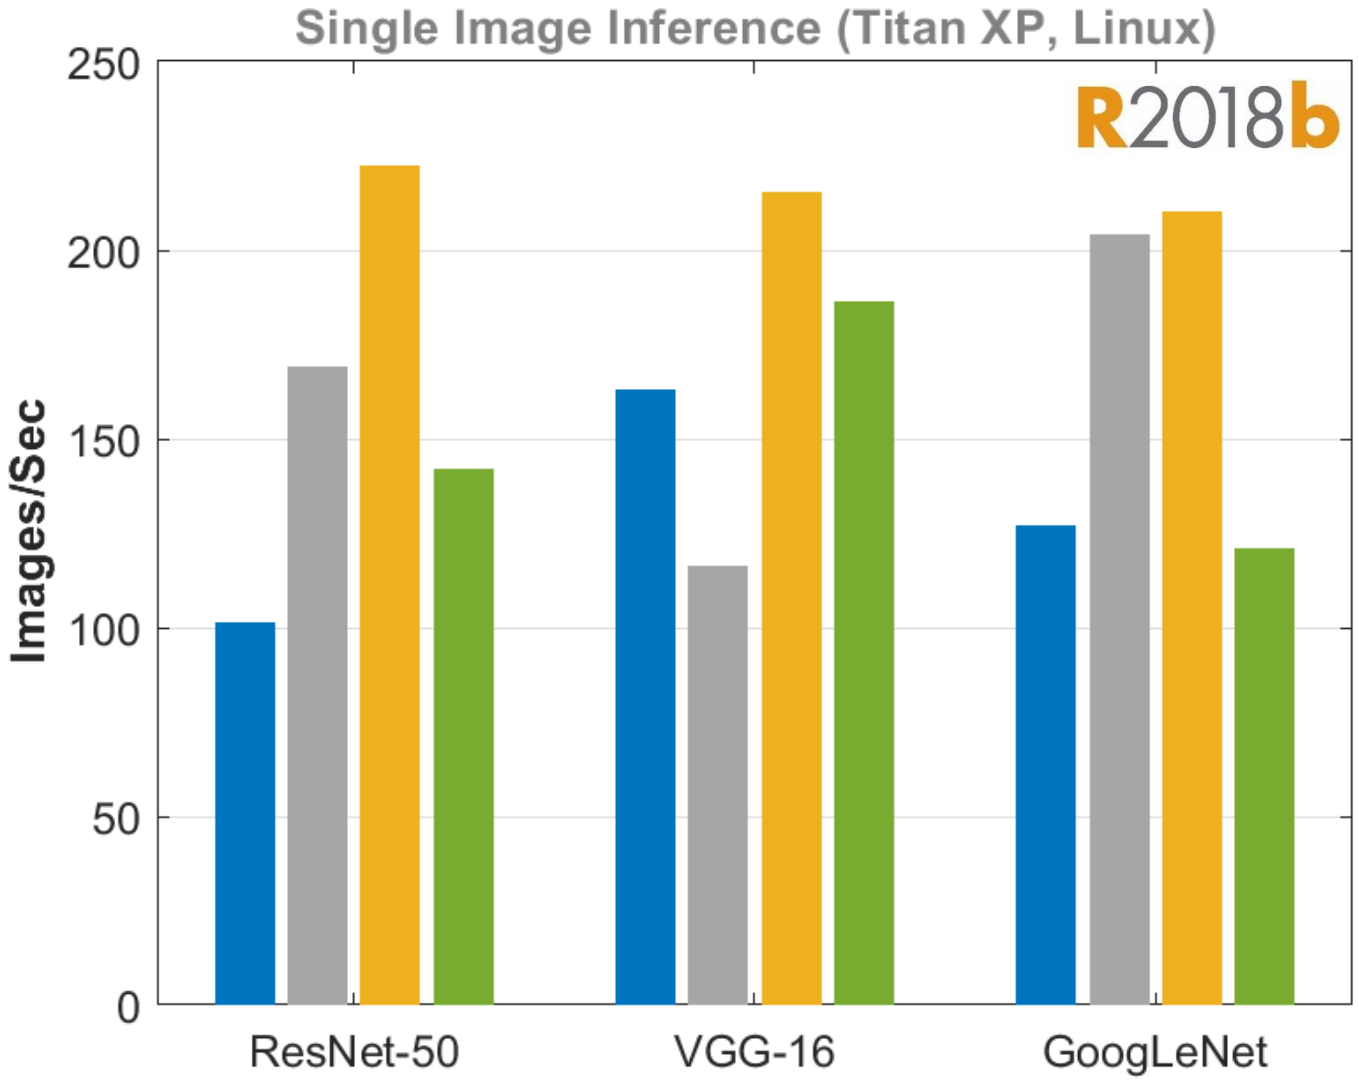

Faster than TensorFlow, MXNet, and PyTorch

Intel® Xeon® CPU 3.6 GHz - NVIDIA libraries: CUDA9 - cuDNN 7 - Frameworks: TensorFlow 1.8.0, MXNet 1.2.1, PyTorch 0.3.1

Agenda

- Introduction and Challenges
- Machine Learning
  - Credit Scoring
- Model Management
- Text Analytics and Deep Learning
  - Twitter Sentiment
- 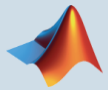 MATLAB-Python Integration
- App Building and Deployment
- Summary

Using MATLAB with Python

- Three ways to integrate:

Algorithm Developers

- Calling Python from MATLAB
- Calling MATLAB from Python

Users

- Packaging MATLAB programs for scalable deployment with Python

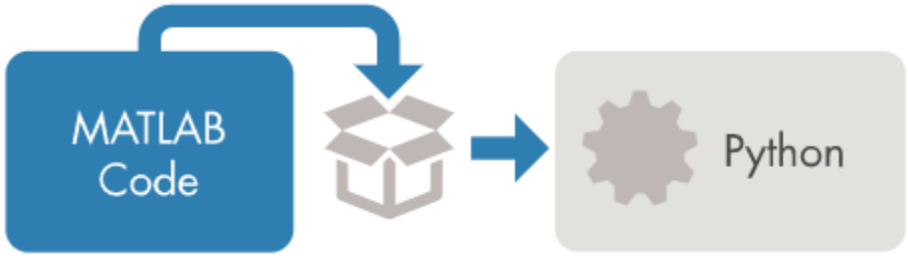

Price Trend Predictor – Regression Modeling

- **Goal:**
  - Cryptocurrency price forecaster

- **Approach:**
  - Use Python libraries for web scraping
  - Rapidly iterate through predictive models

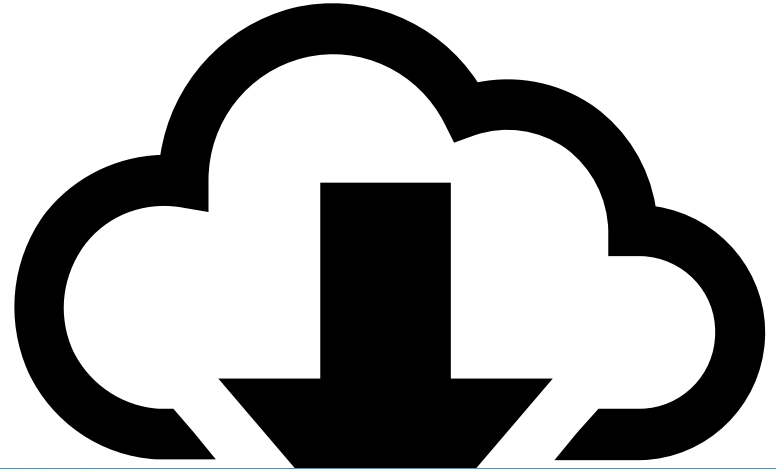

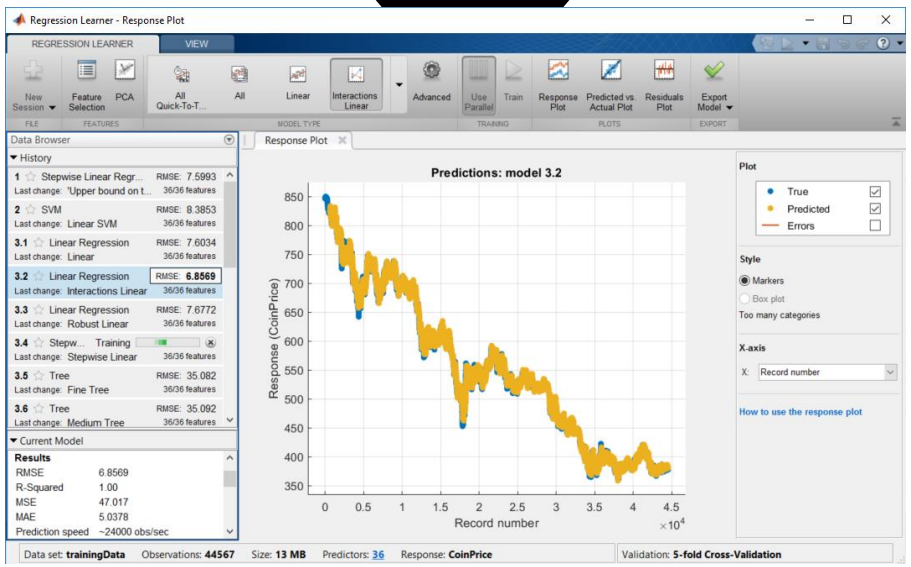

Agenda

- Introduction and Challenges
- Machine Learning
  - Credit Scoring
- Model Management
- Text Analytics and Deep Learning
  - Twitter Sentiment
- MATLAB-Python Integration
- 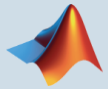 App Building and Deployment
- Summary

Sharing MATLAB Applications

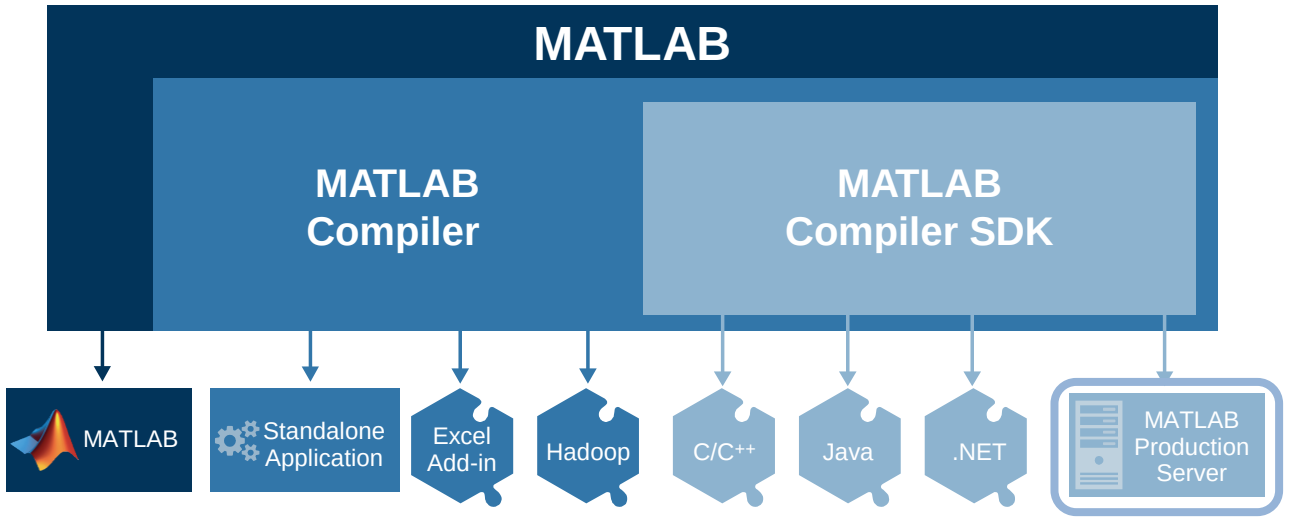

- Share applications with those who do not need MATLAB
- Royalty free
- ***MATLAB Production Server*** provides most efficient path for secure and scalable enterprise applications

MATLAB App Designer Apps – Web Browser Deployment

Authoring

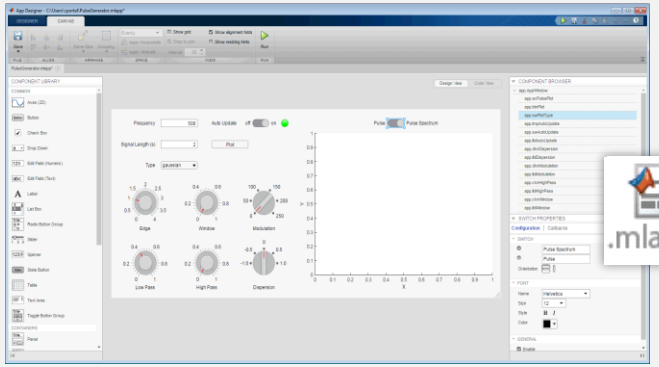

1 Create application in MATLAB App Designer

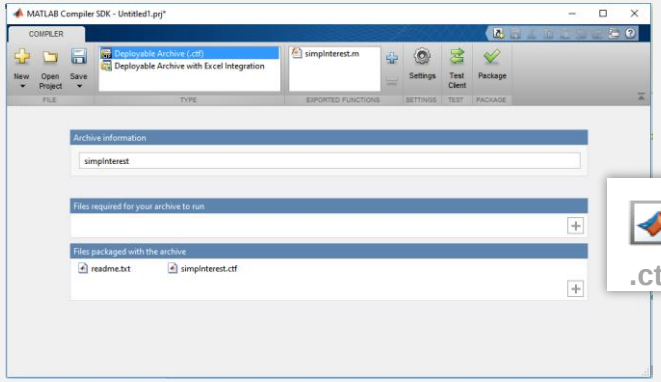

2 Share application as a web app (creates .ctf)

4 Copy .ctf to Web Apps server

One-time server installation

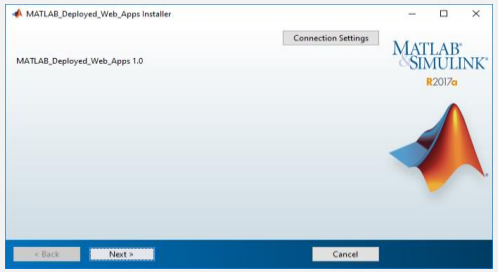

3 Anyone can install the Web Apps server

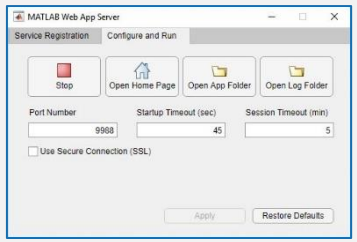

5 Manage the server e.g. start/stop

Access

6 Access the web app through browser (IE 11 or Chrome)

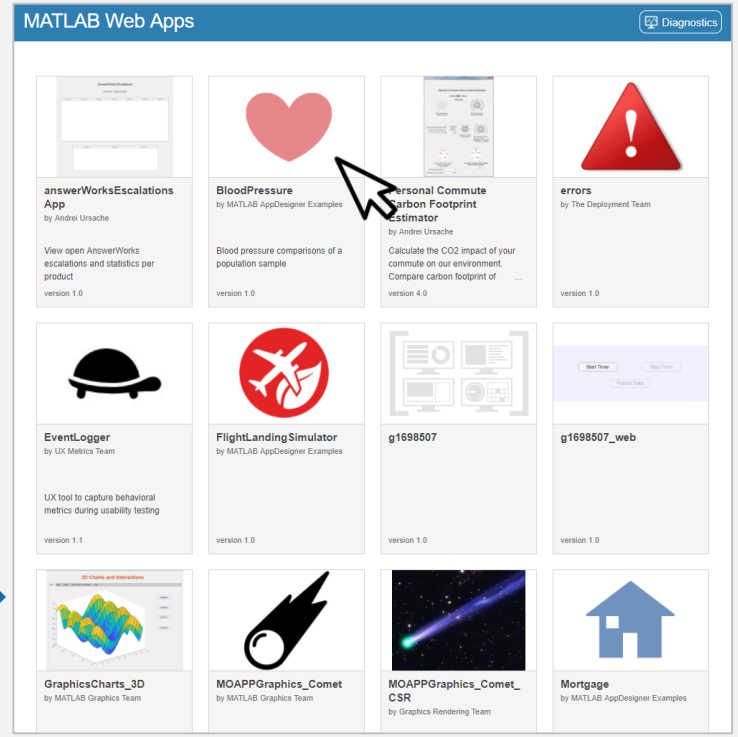

Enterprise Challenges

- Need to provide access to multiple users concurrently
- Need to deploy multiple applications
- Need to integrate with multiple technologies
- Support multiple MATLAB versions

Enterprise Deployment and Scalability of MATLAB

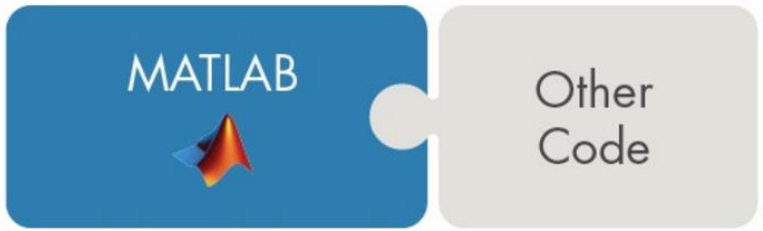

- C/C++
- Java
- C# and .NET
- Python
- Excel
- COM Automation Server
- RESTful, HTTP, and WSDL web services

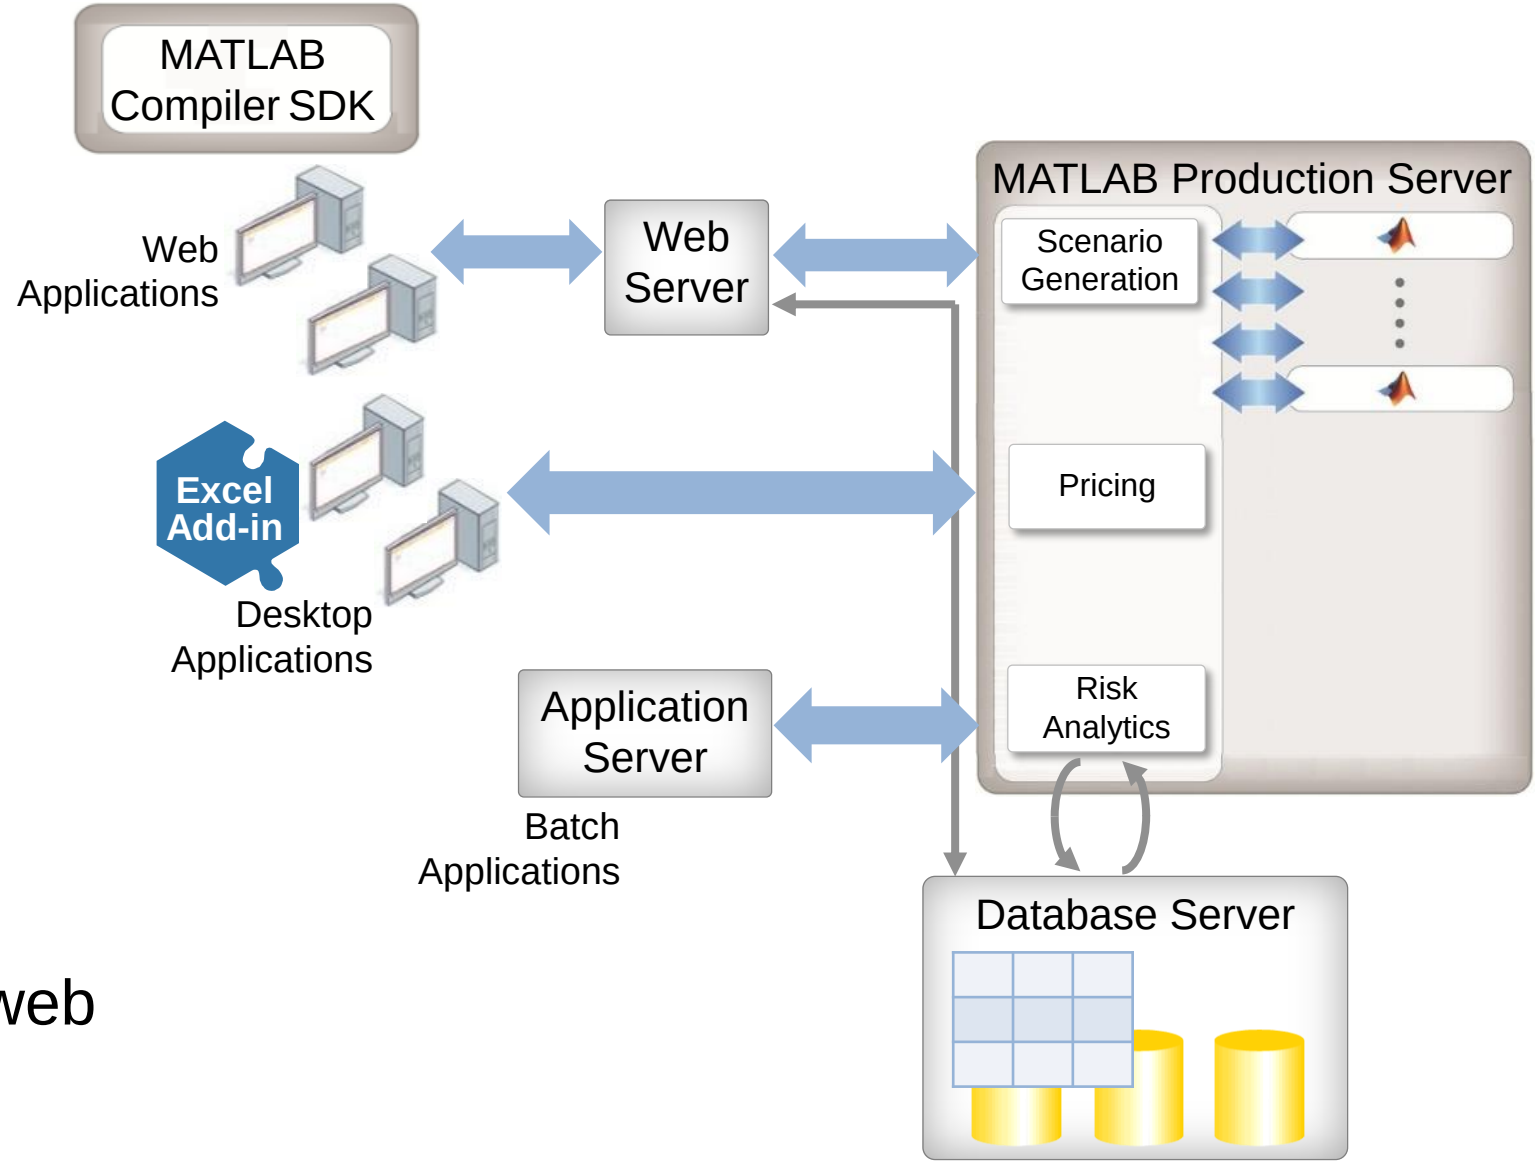

Agenda

- Introduction and Challenges
- Machine Learning
  - Credit Scoring
- Model Management
- Text Analytics and Deep Learning
  - Twitter Sentiment
- MATLAB-Python Integration
- App Building and Deployment

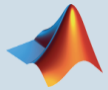

Summary

Challenges

Data Challenges

- Volume, Velocity, Variety
- Data Cleaning 80-20 rule

Modeling Challenges

- No “one size fits all” solution
- Lack of transparency

Documentation and Reporting

- Lack of model traceability
- Model governance and reporting

Production Challenges

- Scalability
- Interfacing with other systems
- Cloud-ready

Model Management

Consulting Services – Analytics

Our products and services cover the entire end-to-end data analytics workflow.

Data Management

- Map data flows
- Connect to data
- Preprocessing

Model Creation

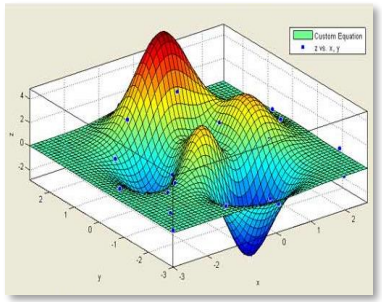

- Analysis
- Predictive modeling
- Scenario analysis
- Data products

App Development

- Architecture
- Custom GUIs
- Component integration

Enterprise Deployment

- Deployment & system integration
- Distributed computing
- Maintenance

Training Services

Exploit the full potential of MathWorks products

mathworks.com/training

- GARP **CPE APPROVED PROVIDER:** Earn one CPE credit per hour of content.
- **Flexible training:** Attend classes available around the world, at your own work site, or on the Web
- **Expert trainers:** Our instructors possess unparalleled knowledge of MathWorks products, gained through the completion of many successful projects throughout multiple industries
- **Effective teaching methods:** Benefit from real-world examples and individualized attention and customized content

Self Paced Training

- Priced dropped by ~ 60%
- Relevant for Basic, Intermediate, Advanced users
- MATLAB Fundamentals: $500
- MATLAB for Financial Applications: $500
- MATLAB for Data Processing and Visualization: $200
- MATLAB Programming Techniques: $350
- Machine Learning with MATLAB: $350
- Deep Learning with MATLAB: $350

Finance Curriculum

Prerequisites
For engineers who are new to MathWorks tools.

Risk Management
For risk analysts, risk managers, portfolio managers, and other financial professionals who analyze, assess, and manage risk.

Portfolio Management
For portfolio managers implementing asset allocation strategies.

Time-Series Modeling
For professionals performing quantitative data analytics.

Application Development
For professionals developing and maintaining complex, standalone applications.

Self-paced Online Course
***** All online courses can be offered in a classroom

Soft Prerequisite

** Onsite trainings can be customized

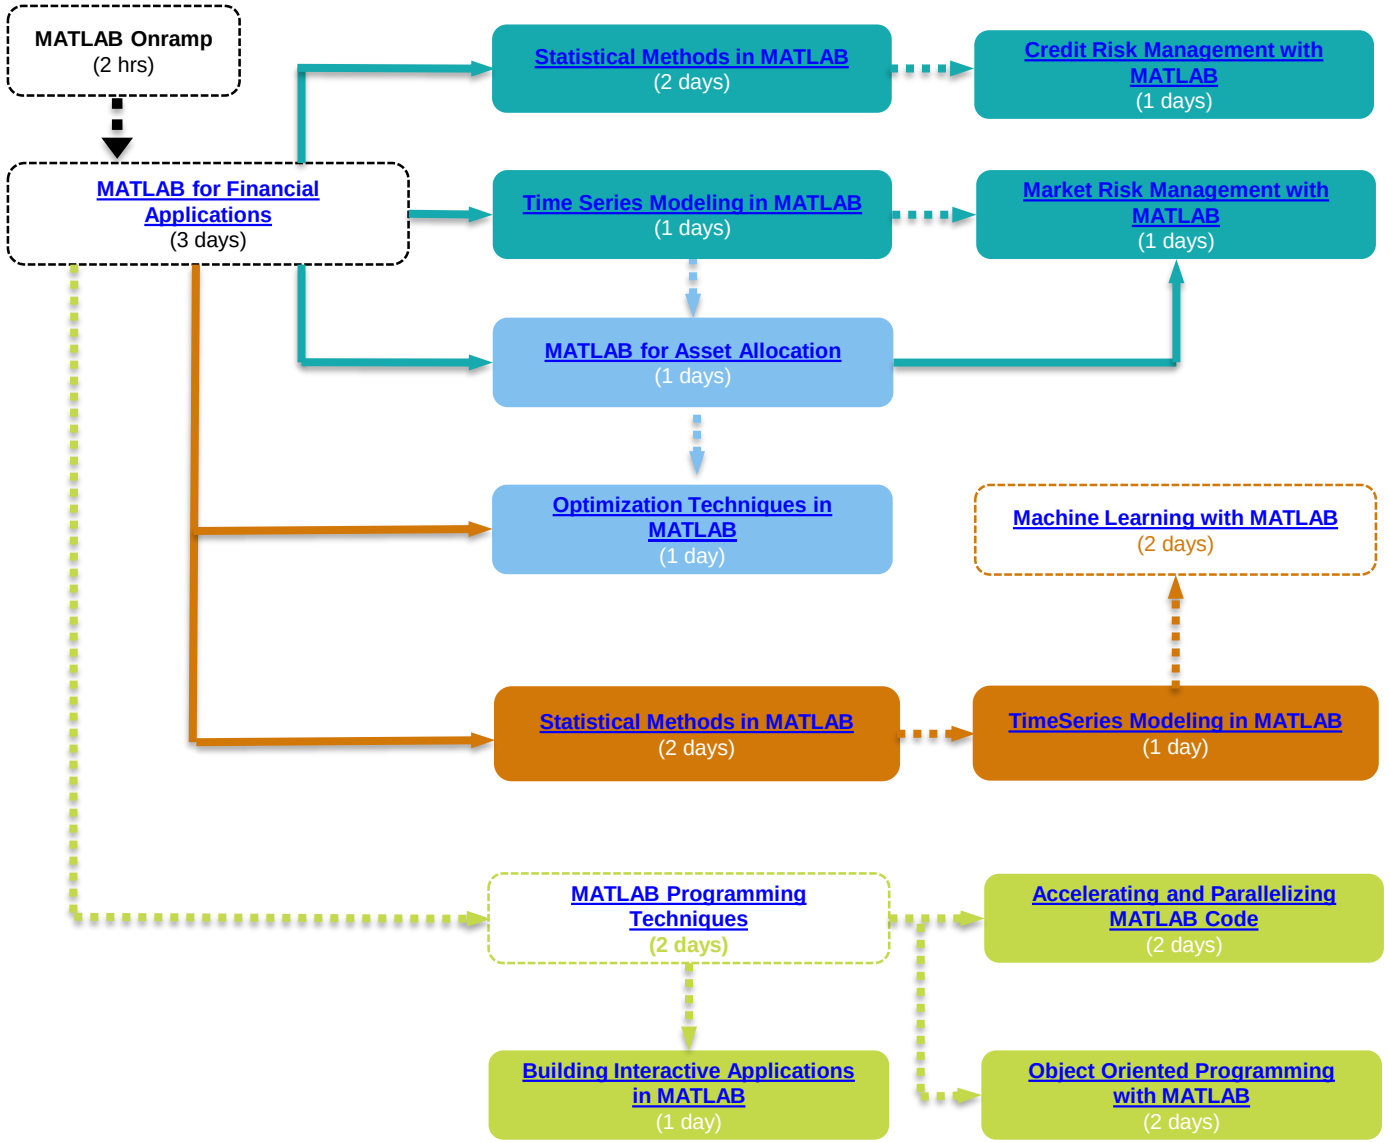

Q&A

Scaling up MATLAB Analytics with Kafka and Cloud Service

Yi Wang
Application Engineering Manager
MathWorks, Inc. USA

MATLAB is only a desktop tool.

That is a myth.

Typical data analytics workflow with MATLAB

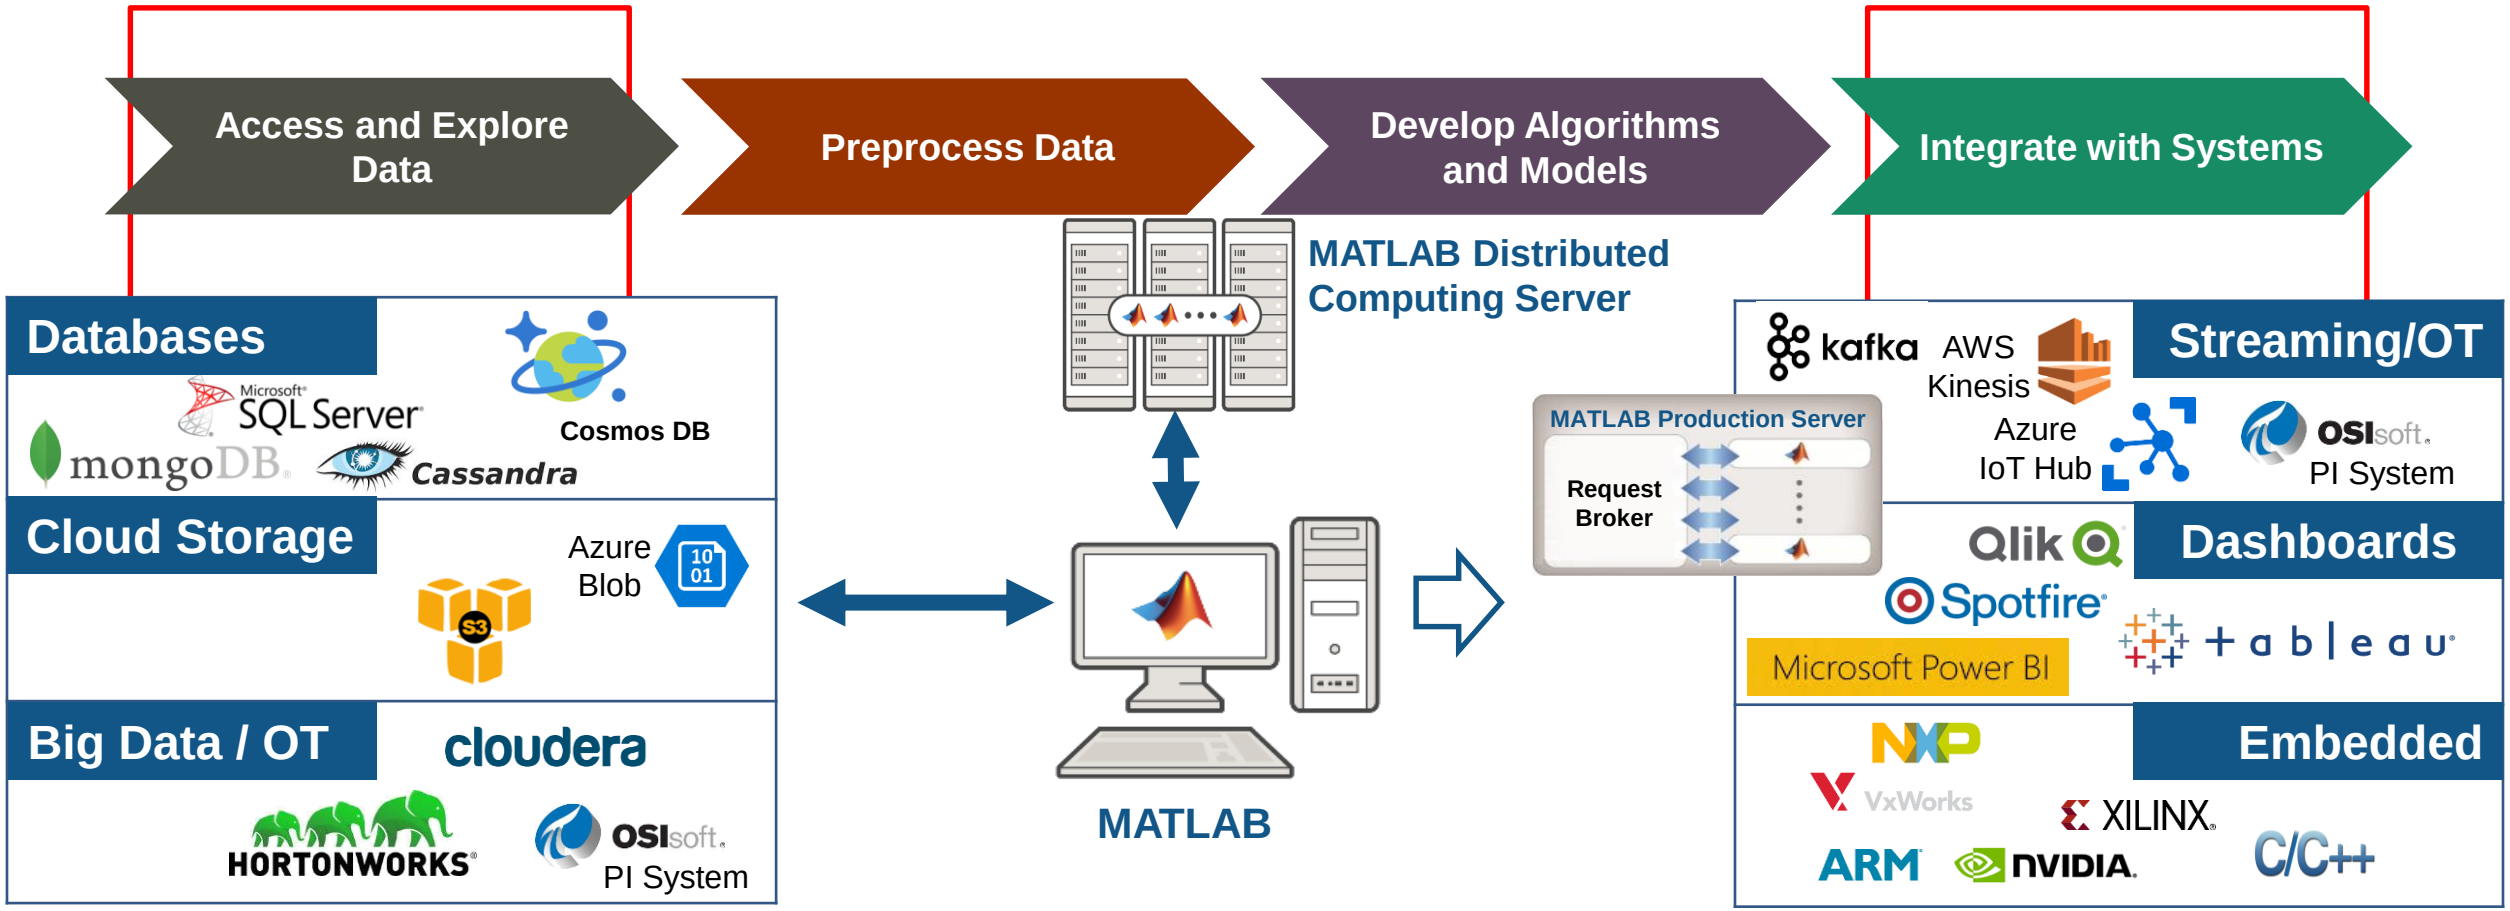

Agenda

- Introduction
- Developing on the cloud
- Scaling to the cloud
- Deploying in the cloud
- Summary

MATLAB Parallel Server

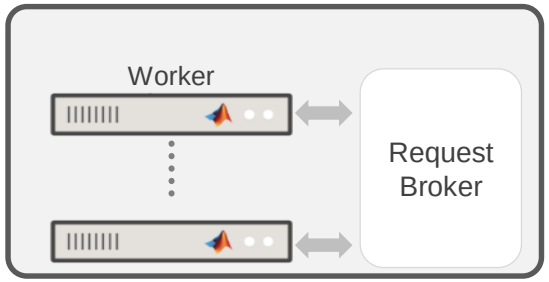

MATLAB Production Server

Cloud Computing

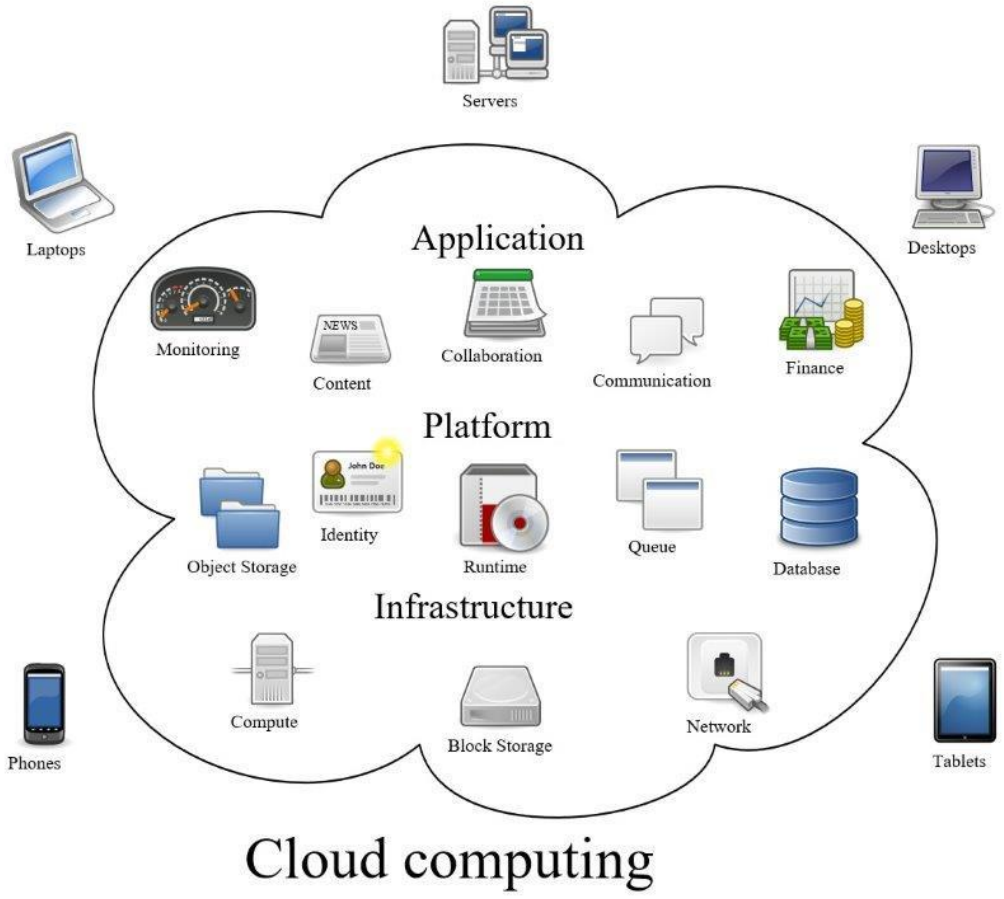

Definition

- Resource pooling
- Broad network access
- Rapid elasticity
- On-demand self-service
- Measured service

***Source**:* National Institute of Standards and Technology

***Source**:* Wikipedia

Cloud adoption is accelerating

- Increasing appeal in the trade-off between capital vs. variable expense with pay-as-you go
  - Agility affords more robust solutions with faster time to market
- Decreasing concerns about all aspects of security
  - Cloud vendors comply with government regulations and laws
  - Internal infrastructure might not be as secure as it was perceived
- Offering continue to evolve
  - Quality of service
  - Storage and databases
  - Application-specific services, e.g., analytics, machine learning, IoT

Increasing demand for MATLAB in the cloud

Four primary use cases:

- Co-locating MATLAB with data stored in the cloud
- Computation on specialized hardware: Multi-core, GPU, FPGA
- Scalability to virtual clusters
- Deployment to elastic infrastructure, integration with enterprise software

Applications:

- Data analytics
- Machine learning
- Condition monitoring
- Predictive Maintenance
- Health monitoring
- Asset Analytics
- Industrial Internet of Things
- ….

BuildingIQ

Adaptive building energy management

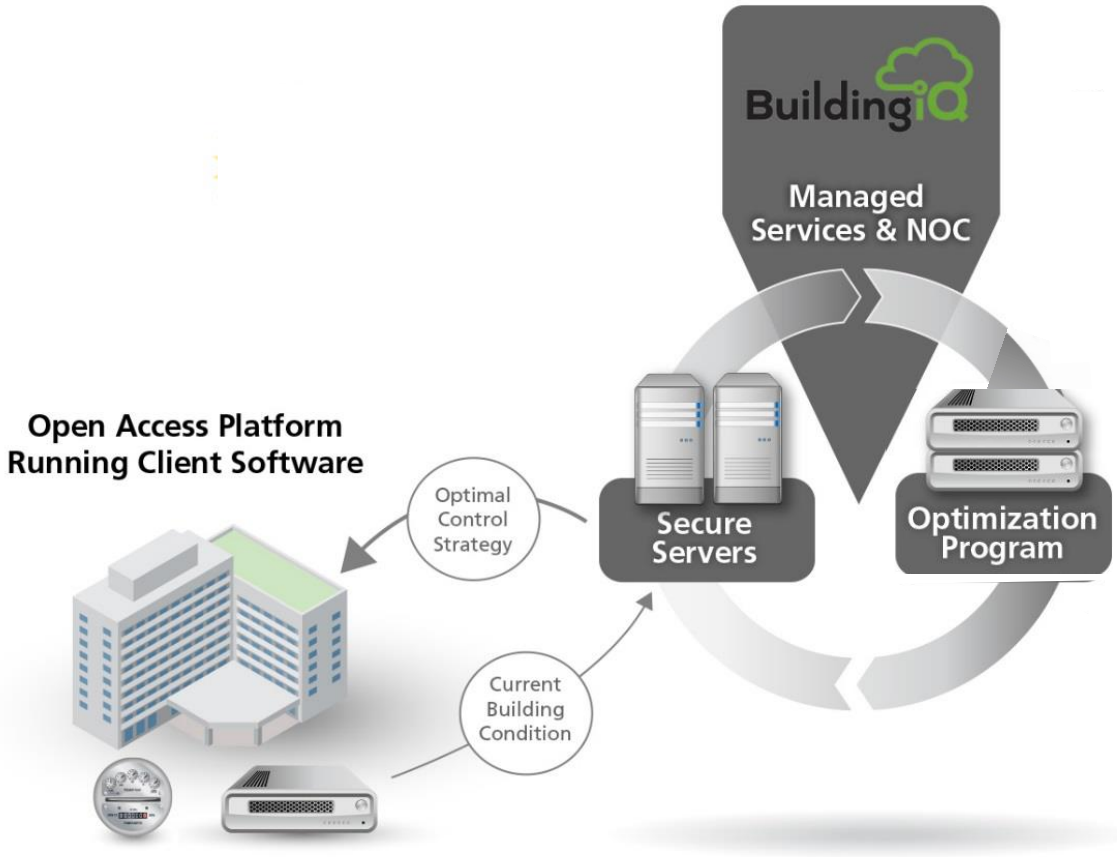

BuildingIQ

Adaptive building energy management

Opportunity

Reduce HVAC energy consumption

Analytics

- Data: Power meters, thermometers, and pressure sensors; weather and energy cost, comprising billions of data points
- Machine learning: SVM regression, Gaussian mixture models, k-means clustering
- Optimization: multi-objective, constrained
- Controls: analyze system-response dynamics
- Hardware: Runs on the cloud, schedule is sent to the building

Result

Typical energy consumption reduced 15-25%

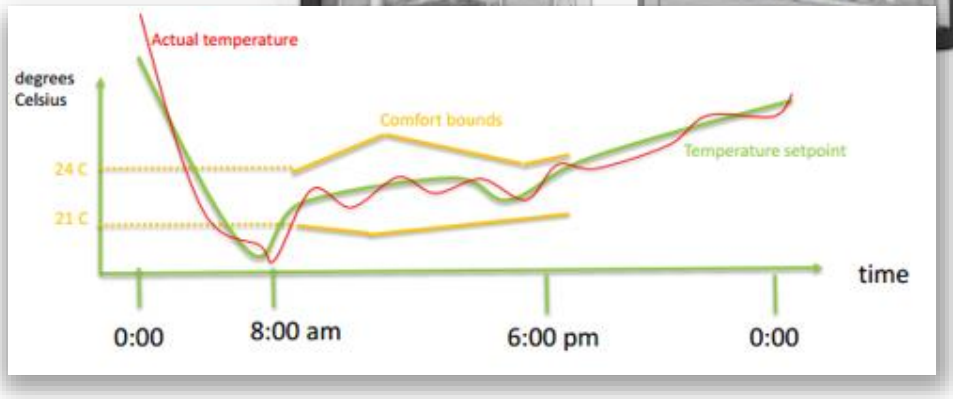

Link to User Story

Agenda

- Introduction
- Developing in the cloud
- Scaling to the cloud
- Deploying in the cloud
- Summary

MATLAB Parallel Server

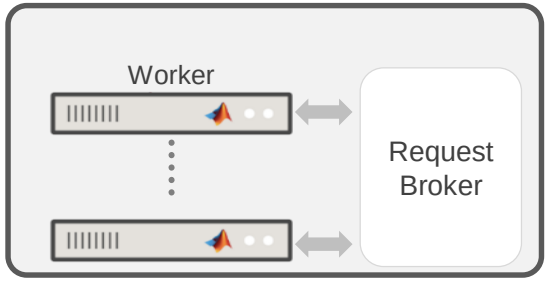

MATLAB Production Server

Developing on the cloud

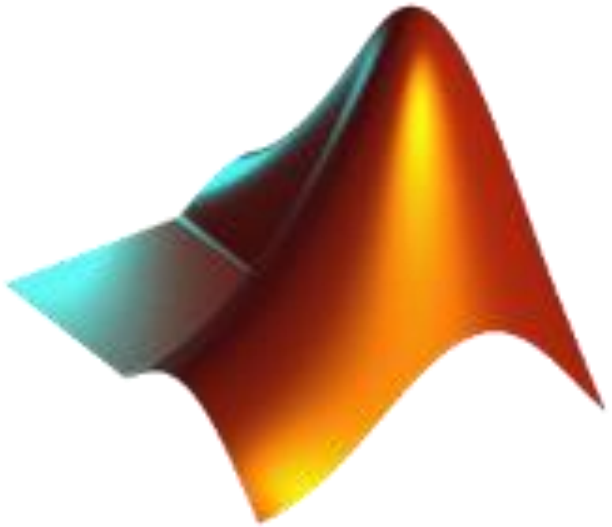

MATLAB users want to get things done ***anywhere***, on ***any device***

- Instant access
- Work and learning to happen *anywhere*, unconstrained by location
- Solutions that play well on *any device*

Access MATLAB from your web browser

MATLAB Online

- Hosted on the MathWorks Cloud
  - Latest version of MATLAB - No downloads, no installs needed
  - Access through your browser – anytime, anywhere
  - Supports desktops, laptops, Chromebooks
  - Share files with others directly
- Available with many license types
  - Including Individual, Total Academic Headcount, Student, Home, etc.
  - Must be current on MathWorks Software Maintenance Service

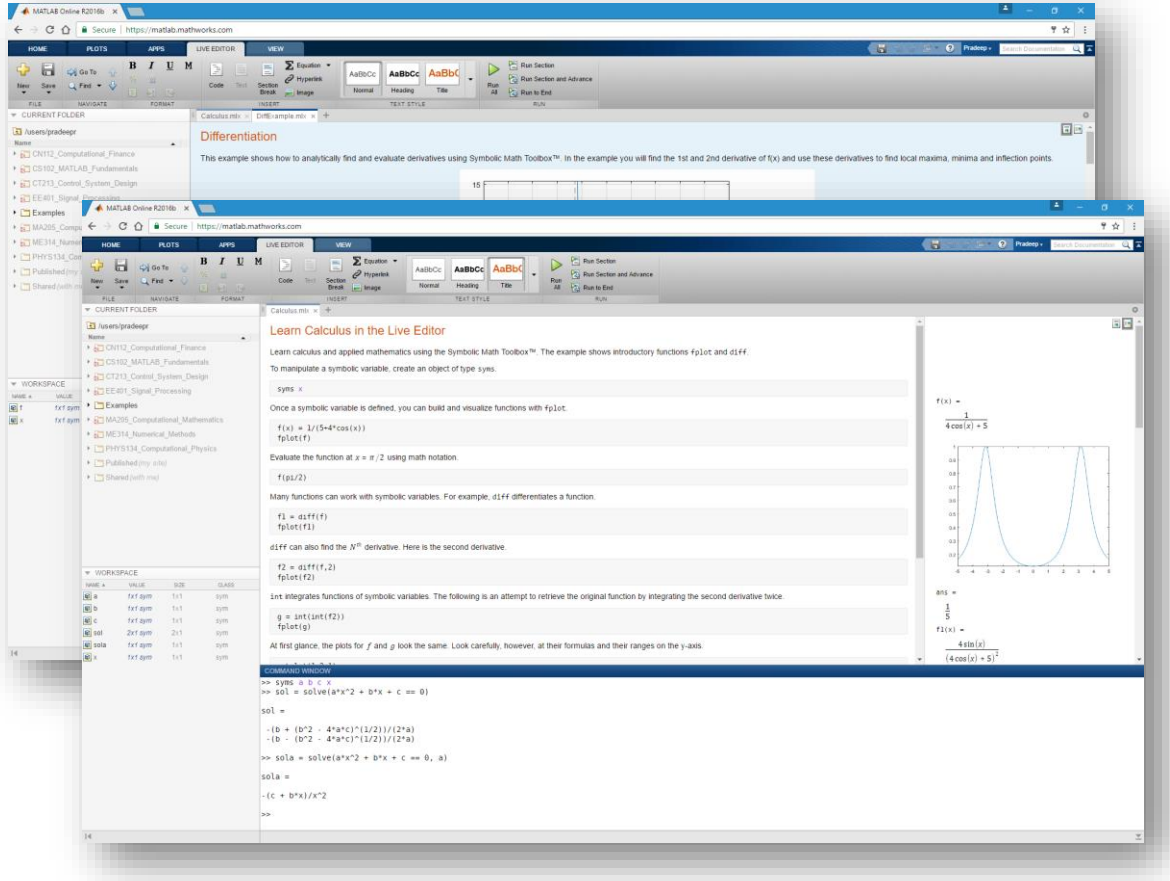

"I love MATLAB Online. It's the most exciting addition I've seen since I started using MATLAB twenty-five years ago."

-Dr. Mike Fitzpatrick, Vanderbilt University

https://www.mathworks.com/products/matlab-online.html

Explore and analyze problems in a single, interactive environment

MATLAB Live Editor

- View, create and run Live Scripts
  - Interactive documents that combine code with formatted text, equations, and images
- Eliminate context switches
  - Store and display output alongside the code that creates it
  - Add text, equations, LaTeX
- Available on desktop MATLAB and MATLAB Online

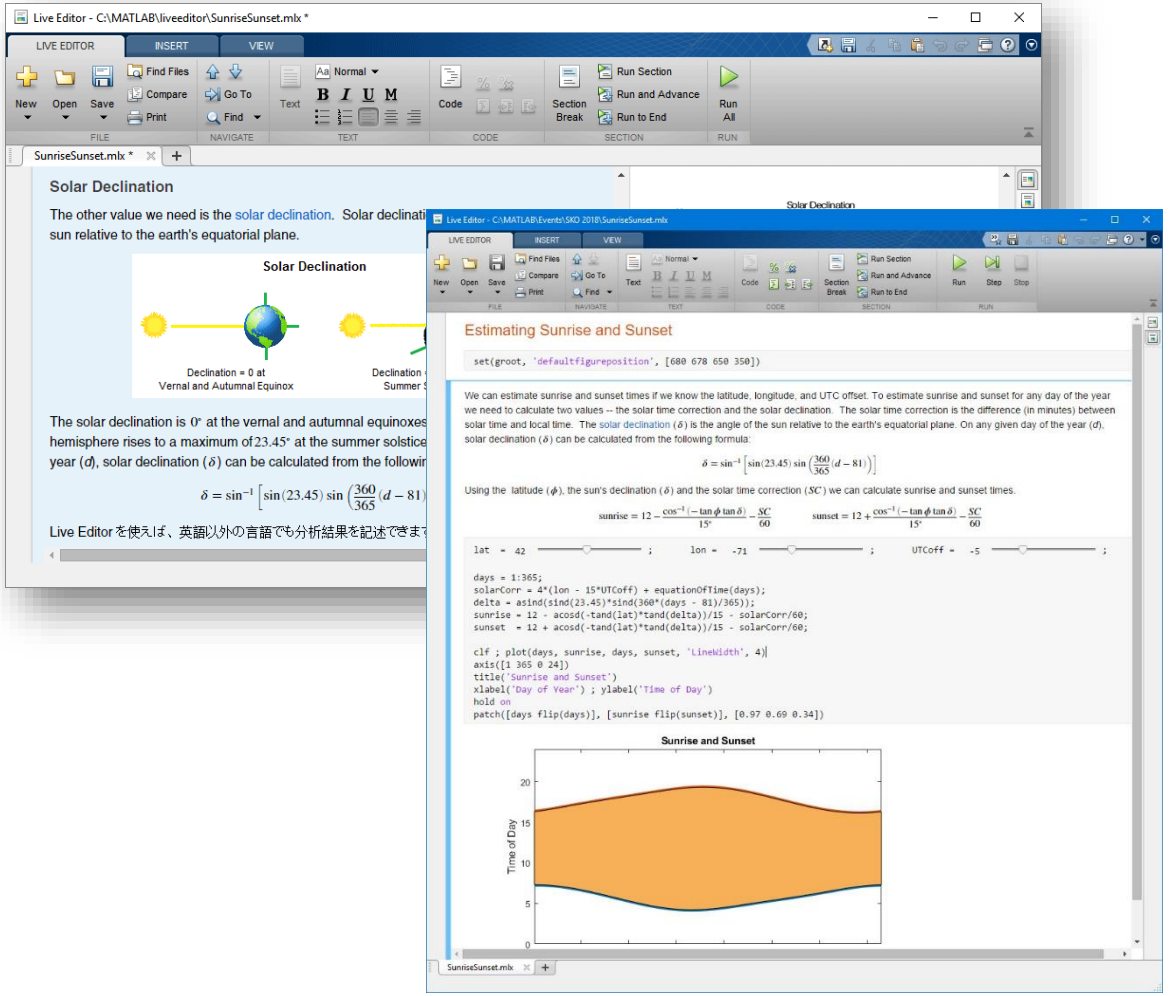

Demo: Using Live Scripts in MATLAB Online

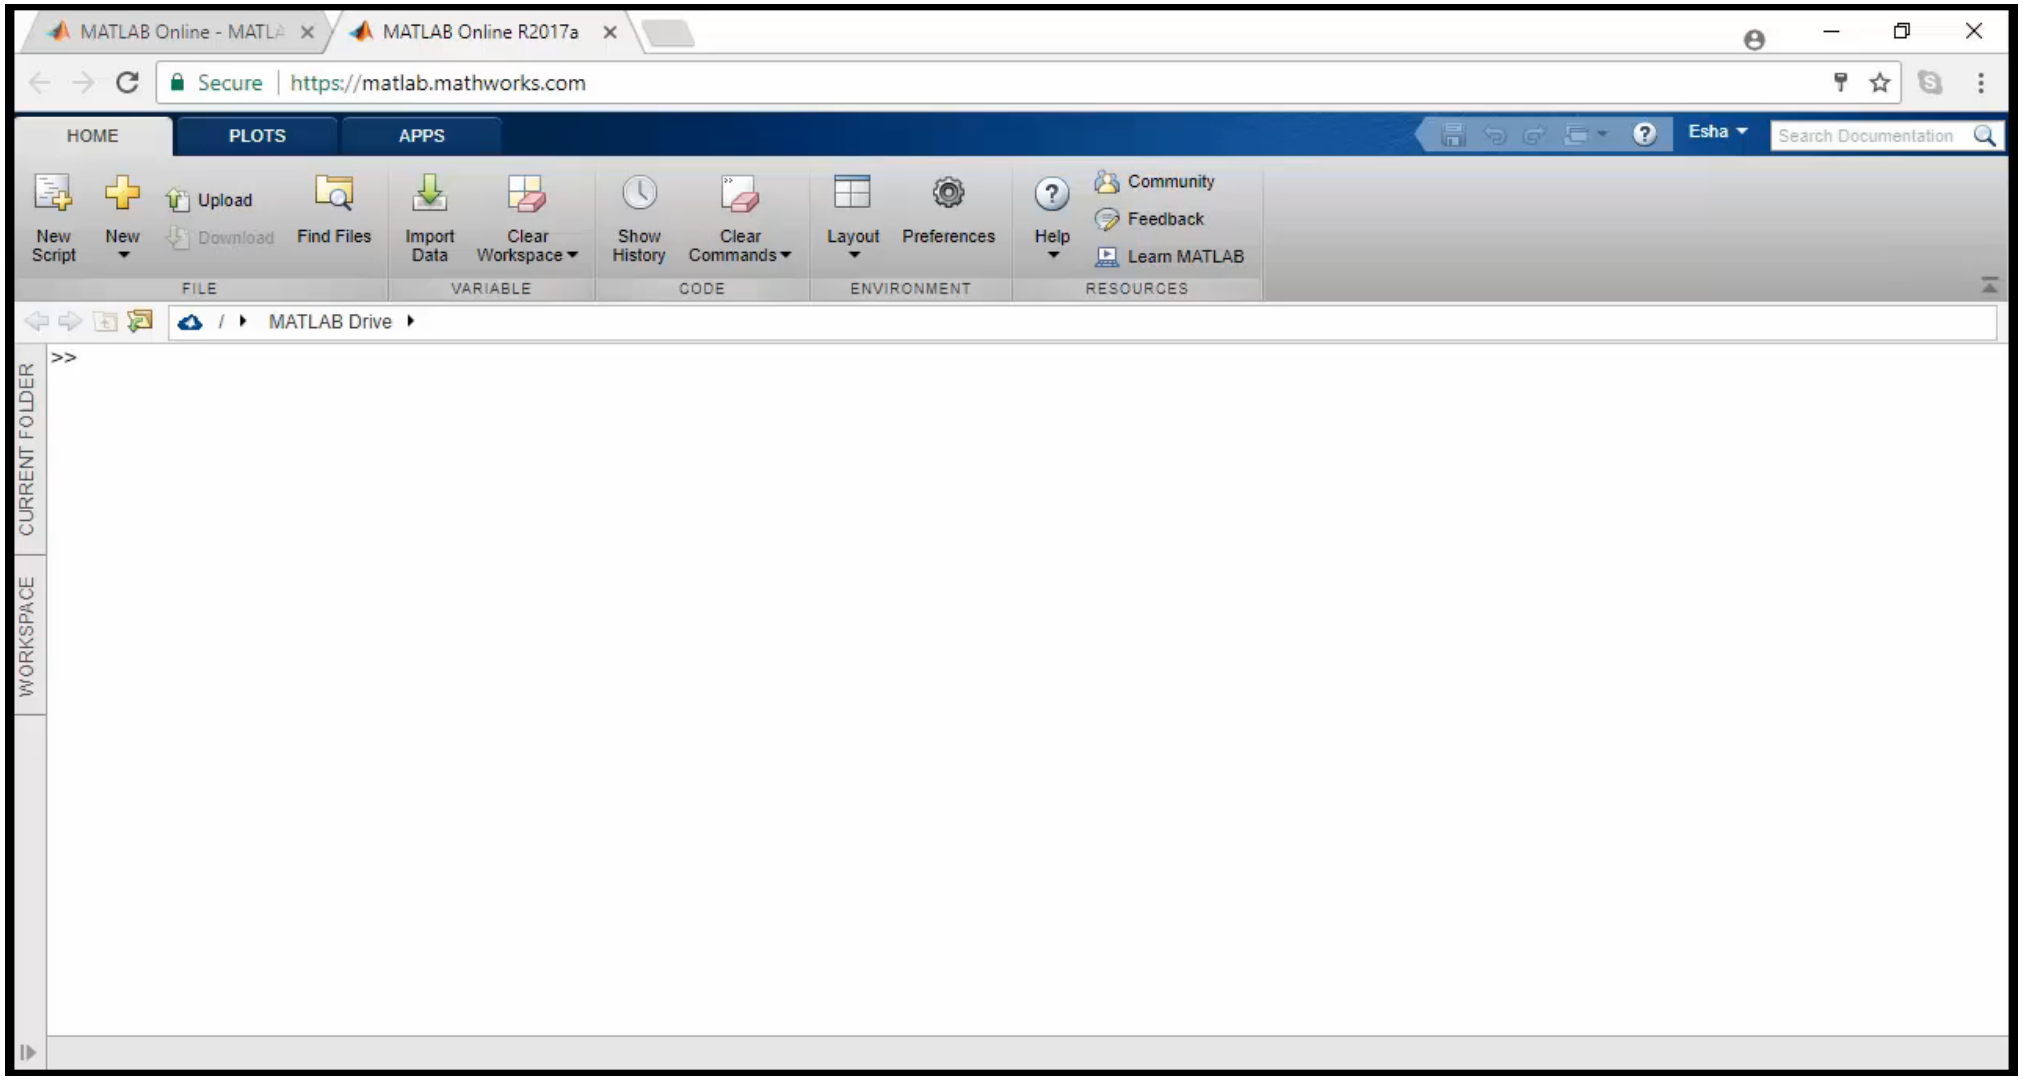

Connect to MATLAB from your iPhone, iPad or Android device

MATLAB Mobile

- Create, edit, view and run MATLAB files, visualize data
- Acquire data from device sensors like accelerometer, GPS etc.
  - Lightweight exploration on your device, or analysis on MATLAB
- Connect to the cloud and to MATLAB running on your computer
- Lightweight app suited for teaching, learning and quick prototyping

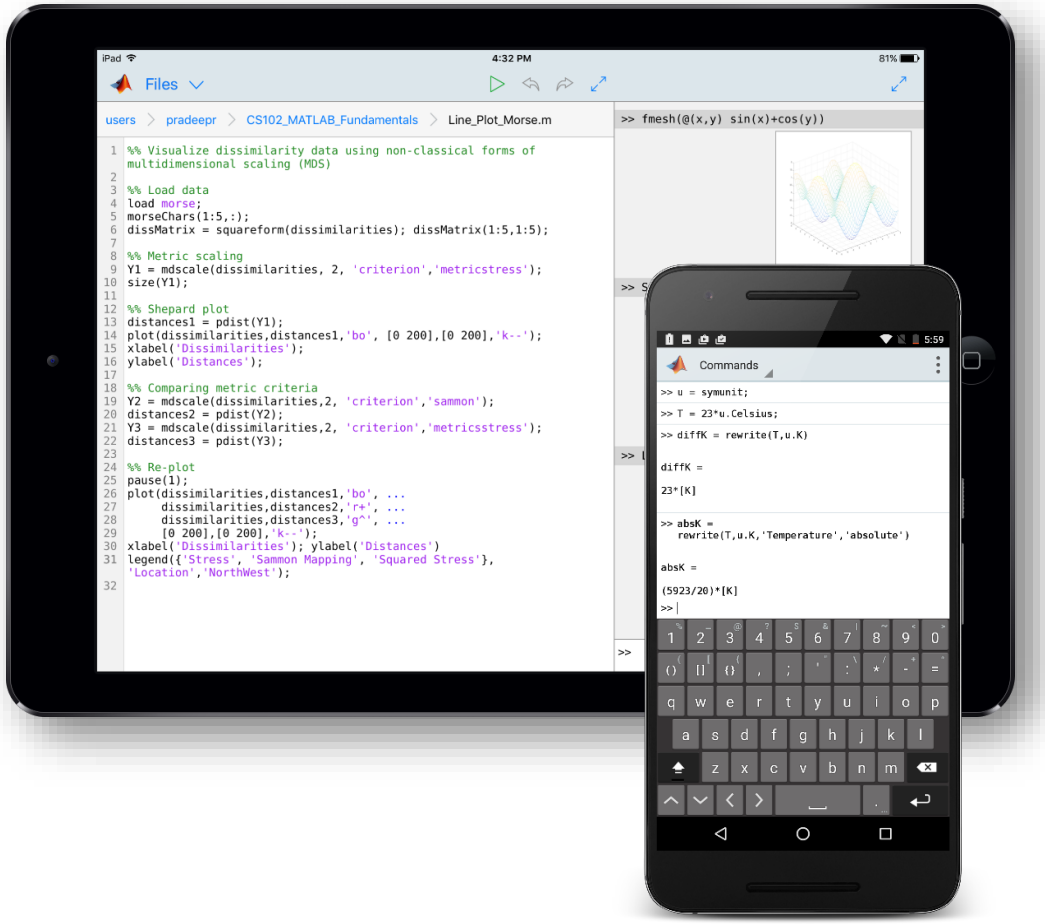

"To someone who thinks in MATLAB, it is great for allowing you to explore a problem on the move."

https://www.mathworks.com/products/matlab-mobile.html

Store, access and manage files from anywhere

MATLAB Drive

- Cloud-based storage for your files
  - Access from MATLAB Online and MATLAB Mobile
  - Access from your desktop with the MATLAB Drive Connector
  - Browser-based access at drive.matlab.com
- Two storage quota options
  - 250 MB free storage with MathWorks Account
  - 5 GB with MATLAB license current on MathWorks Software Maintenance Service

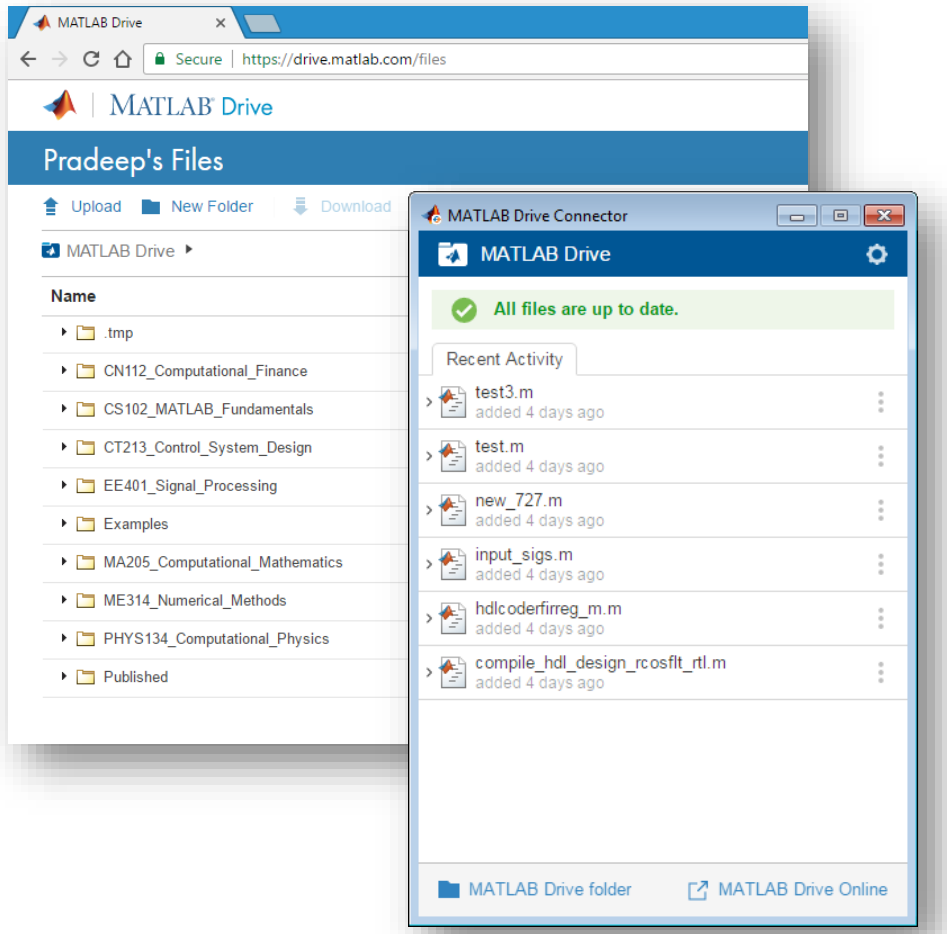

https://www.mathworks.com/products/matlab-drive.html

Built-in MATLAB Language Support for Cloud Data Services

Support Provided By:	Data Services			
Datastore	Blob Storage *(read, out of memory data)*			
Database Toolbox *(standard ODBC/JDBC)*	Azure SQL Database	Azure Database for MySQL/PostgreSQL	SQL Data Warehouse	SQL Server Stretch Database
Support Packages *(provided via pilot engagement)*	Blob Storage: Blob, Table, File *(read/write/delete, encryption, access control)*	Data Lake *(read/write/delete, encryption, access control)*	CosmosDB: MongoDB, Cassandra, Table interfaces	

Datastore	Amazon S3 *(read, out of memory data)*			
Database Toolbox *(standard ODBC/JDBC)*	Amazon Aurora	Amazon RDS for PostreSQL/MySQL/MariaDB/Oracle/SQL Server		
Support Packages *(provided via pilot engagement)*	Amazon S3 *(read/write/delete, encryption, access control)*	Amazon EFS *(NFS/Linux)*	Amazon Athena (*Query of S3 data*)	

Example: Access Data in S3 from MATLAB

athenaQuery.mlx

Access the data in S3

Bring up the AthenaClient

```
athenaClient = aws.athena.Client();
athenaClient.Database = 'trainingdata';
athenaClient.initialize();
```

Create a query and submit

```
athenaClient.submitQuery('SELECT * FROM "trainingdata"."sampledata" limit 100','s3://fleettrainingdata')
```

Fetch data as a table for easy analysis

```
ds = datastore('s3://fleettrainingdata/*.csv');
ds.NumHeaderLines = 2;
data = table(ds);
```

Your usual MATLAB workflow goes here

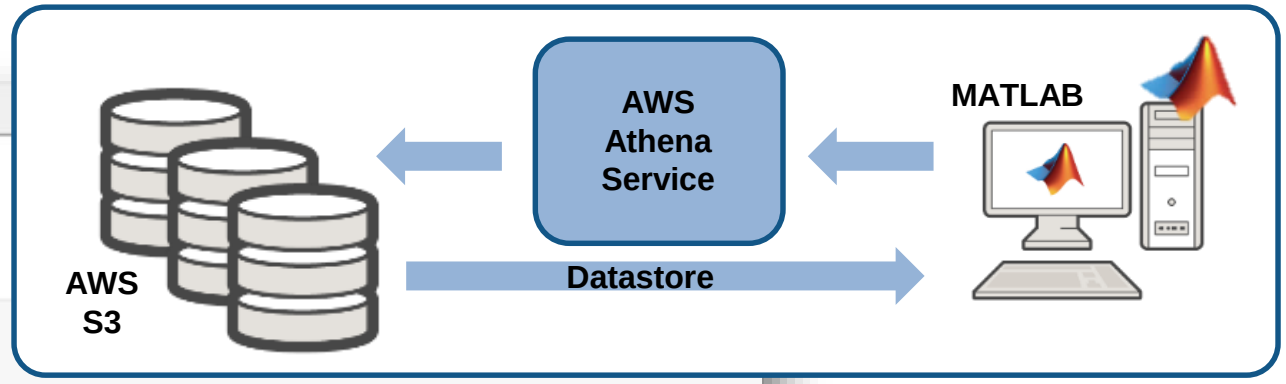

MathWorks provides an ecosystem to develop on the cloud

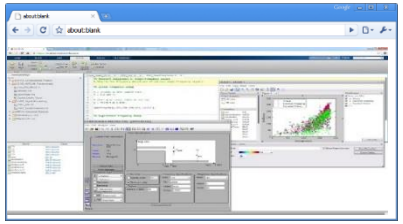

MATLAB Online

Browser-based access to MATLAB

Get started instantly (no downloads, no install, no platform restrictions)

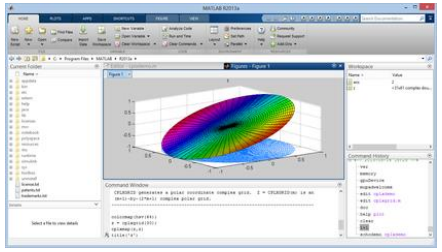

Desktop MATLAB

MATLAB Mobile

Access to MATLAB from iOS or Android devices

Quick, lightweight prototyping, follow along from your mobile device

MATLAB Drive

Cloud storage for your files

Access files anywhere (your computer, MATLAB Online, MATLAB Mobile)

Live Editor

Cloud Reference Architectures

- Published on GitHub, by MathWorks
  - Cloud templates
  - Architecture diagram
  - Step-by-step instructions
- Automates the process of launching MathWorks platform products in the cloud
  - MATLAB
  - MATLAB Distributed Computing Server
  - MATLAB Production Server
- Supports Azure AWS and Microsoft Azure out of the box

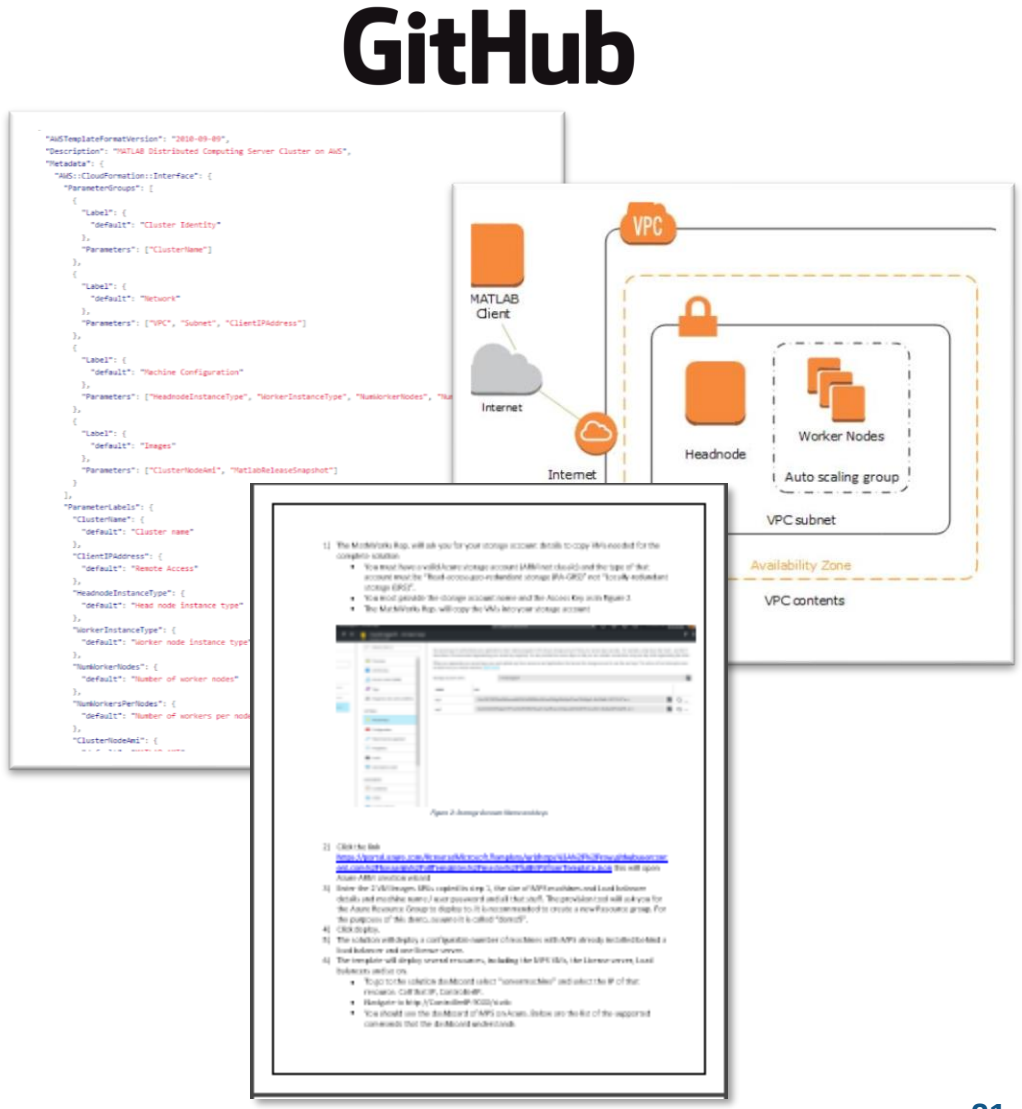

MATLAB & Simulink

Cloud Reference Architecture

Use cases:

- Data analytics on cloud-stored data
- Access to high-end hardware
  - Multi-core VMs
  - GPUs
  - FPGAs
- Prototyping parallel algorithms on one VM before scaling to a cluster
- On-server algorithm testing before production deployment

Benefits:

- Quickly get up-and-running in AWS or Azure
- Leverages MATLAB installed on instances in AWS/Azure, eliminating installation process
  - MATLAB language can access pre-built capabilities that cloud services provide (e.g., data storage)
- Incorporates IT best practices
- Adapt or extend for your specific needs

Scaling to the cloud

MATLAB Parallel Server

Parallel Computing Paradigm

- Prototype on the desktop
- Integrate with infrastructure
- Access directly through MATLAB

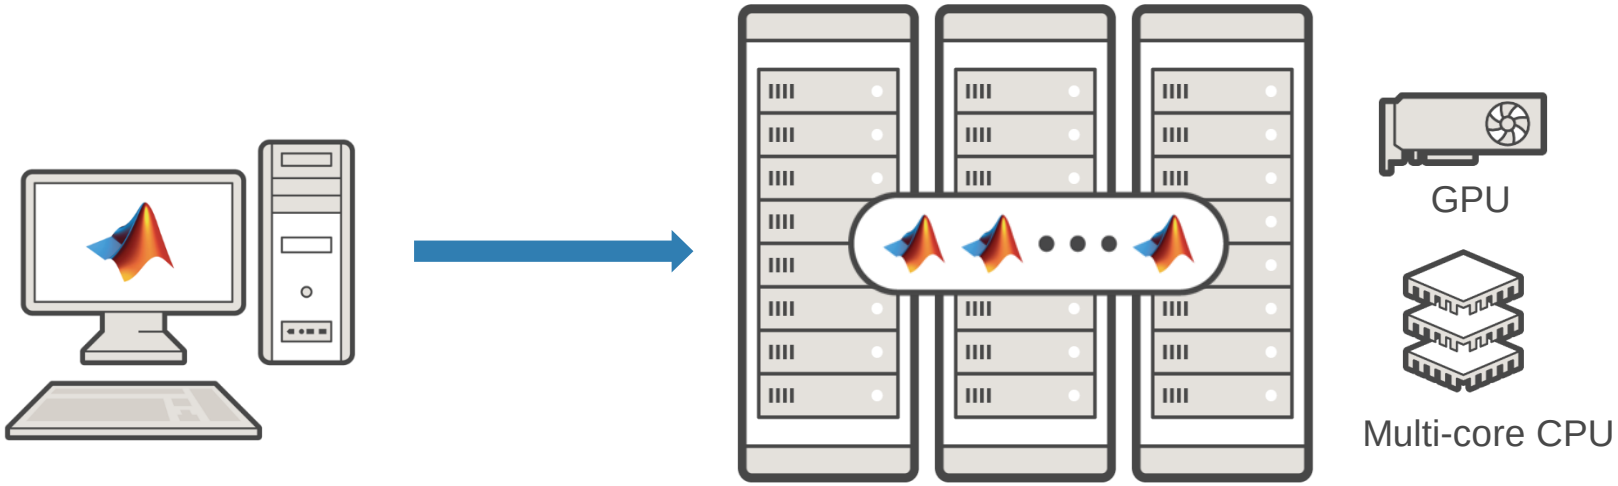

Speeding up in the cloud

Same code, three different environments

MATLAB:
4000 seconds

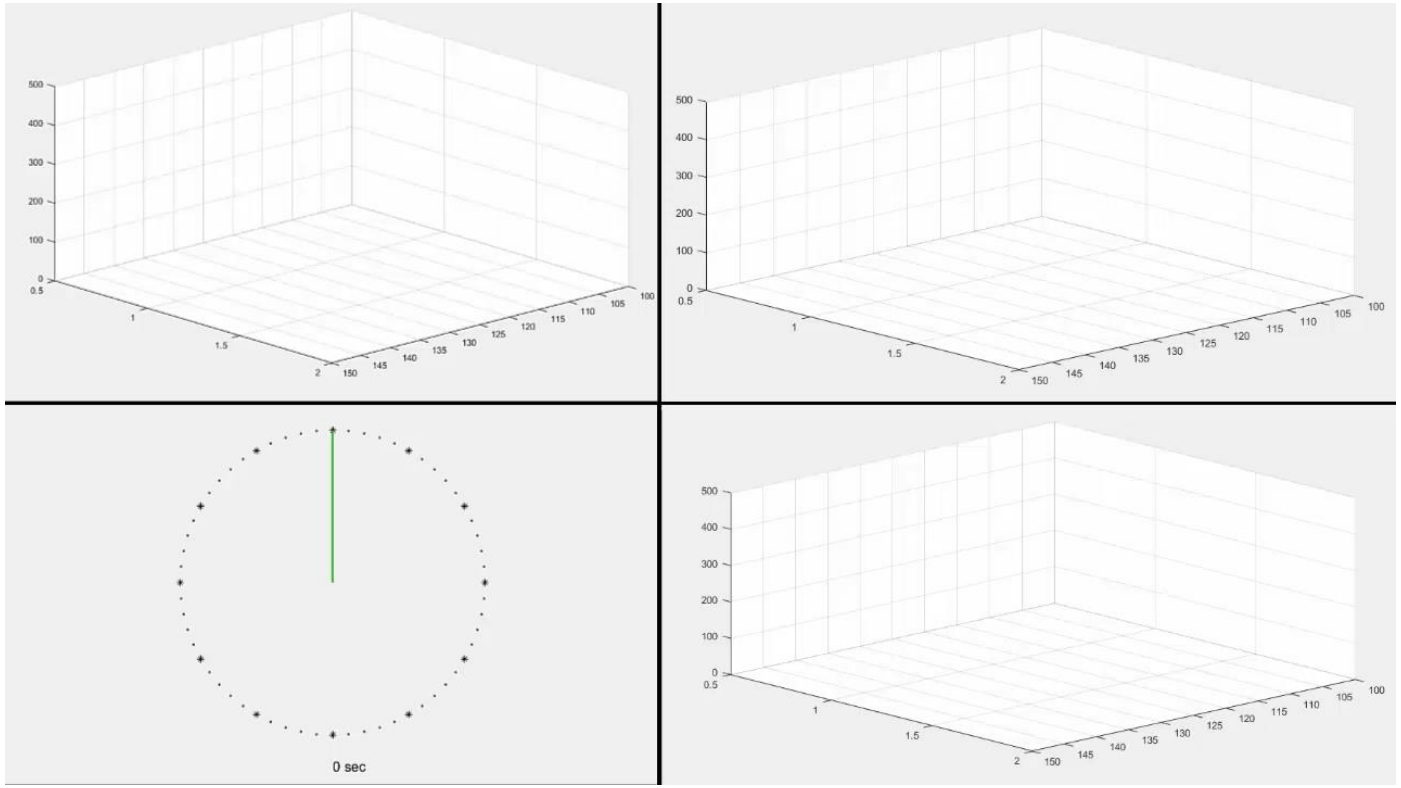

8 local workers:
925 seconds

256 cloud workers:
41 seconds

MATLAB Parallel Server

MathWorks Cloud Center

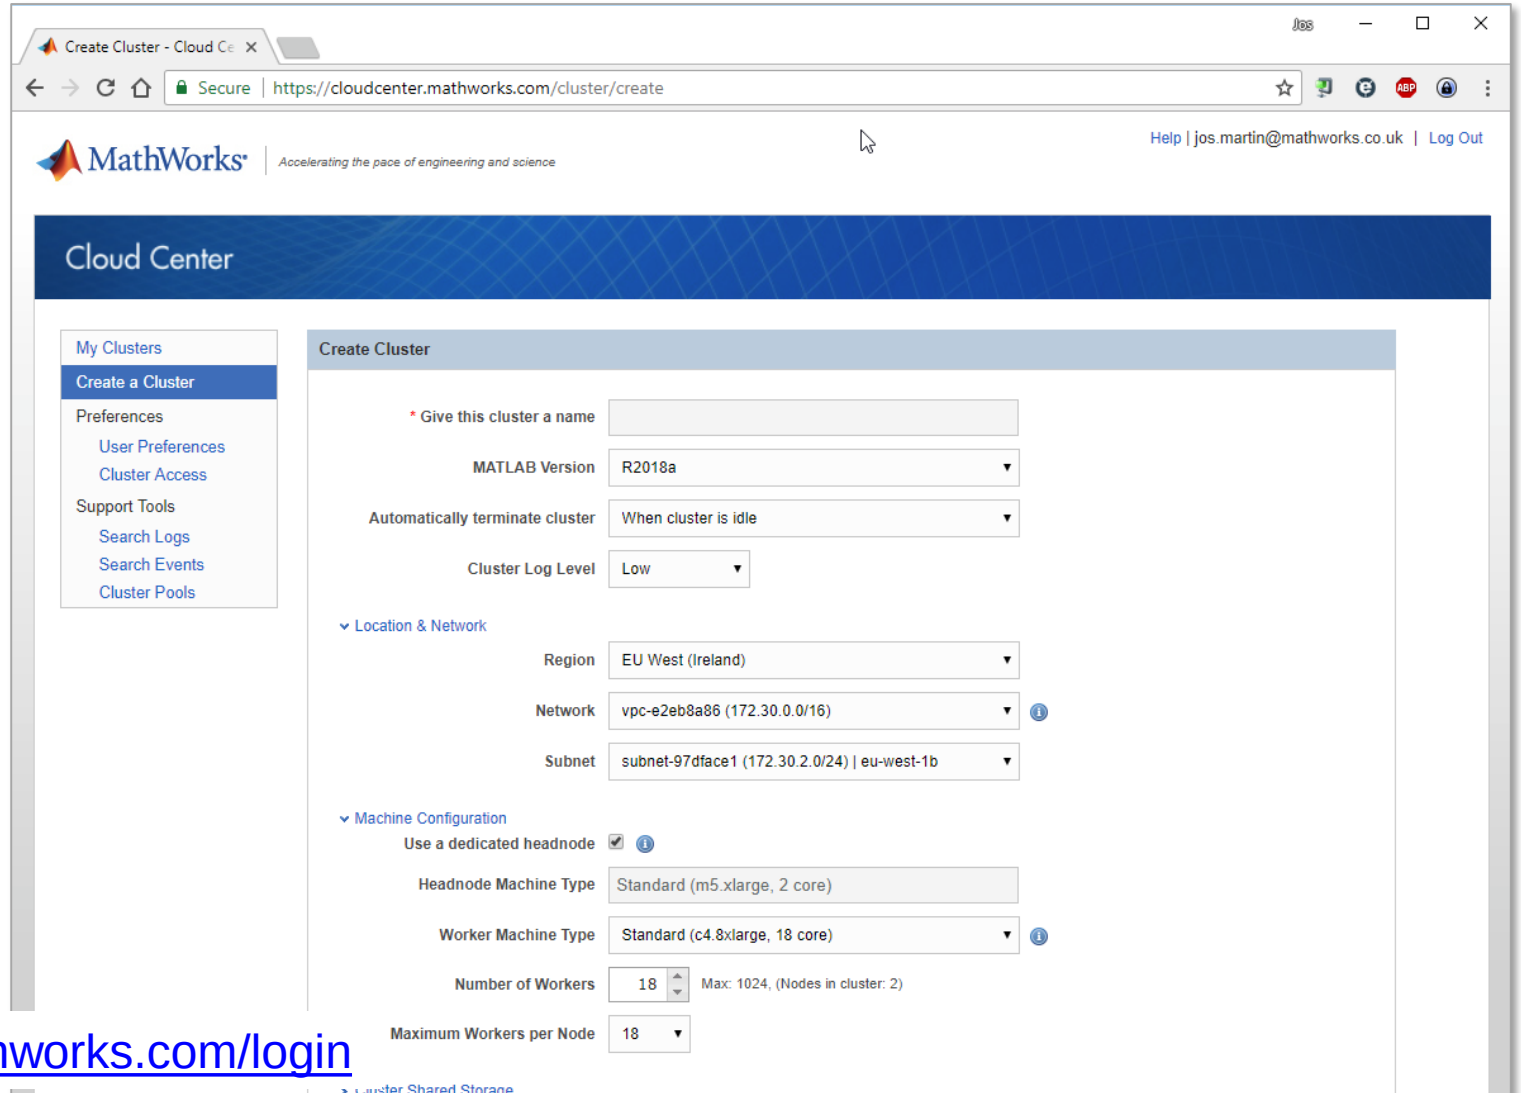

https://cloudcenter.mathworks.com/login

Demo: Using MathWorks Cloud Center

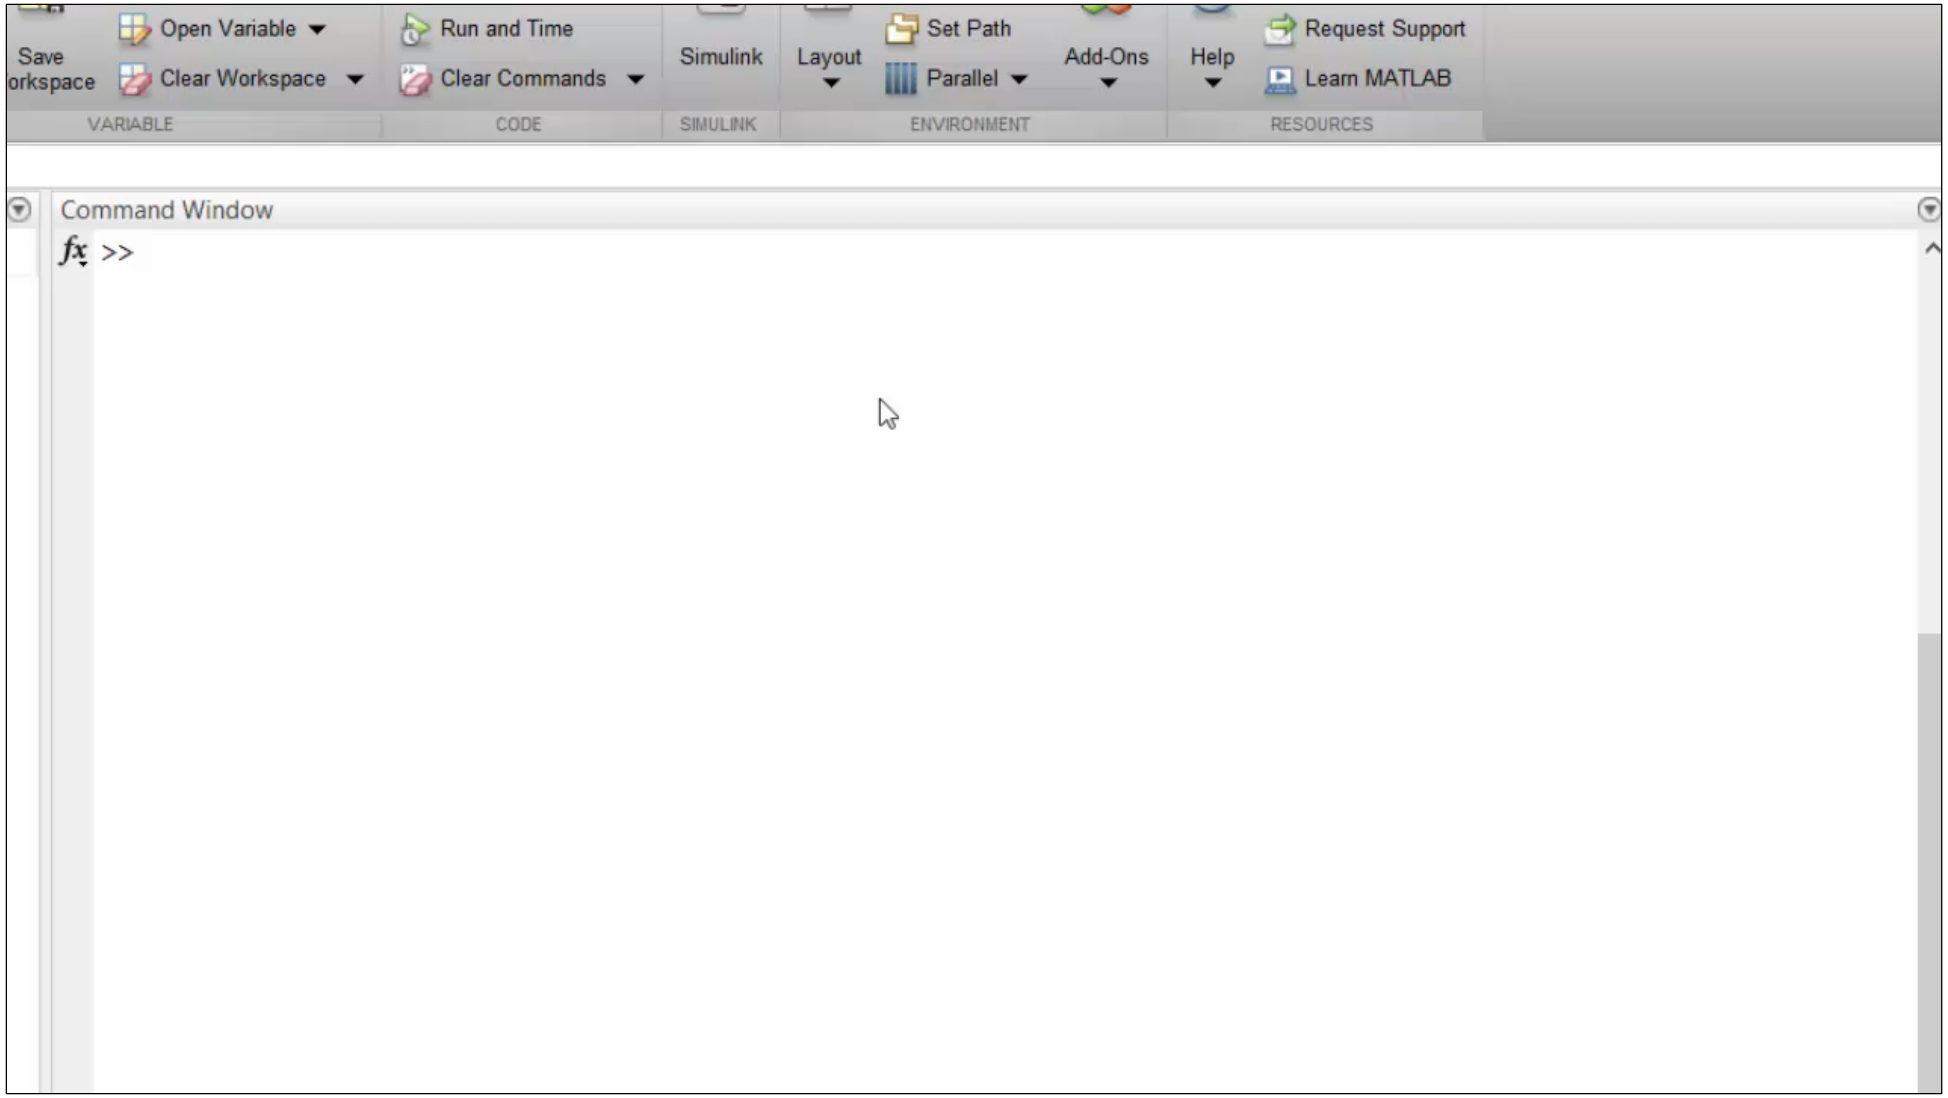

MATLAB Parallel Server

MathWorks Cloud Center

Requirements:

- Workstation that can connect outside of your network
- MATLAB and Parallel Computing Toolbox
- Amazon Web Services account
- MATLAB Distributed Computing Server License
  - Configured for MathWorks Hosted License Manager

Latest functionality:

- Dedicated head node
- 1024 workers per cluster
- Pretrained convolutional neural networks (CNN) for deep learning

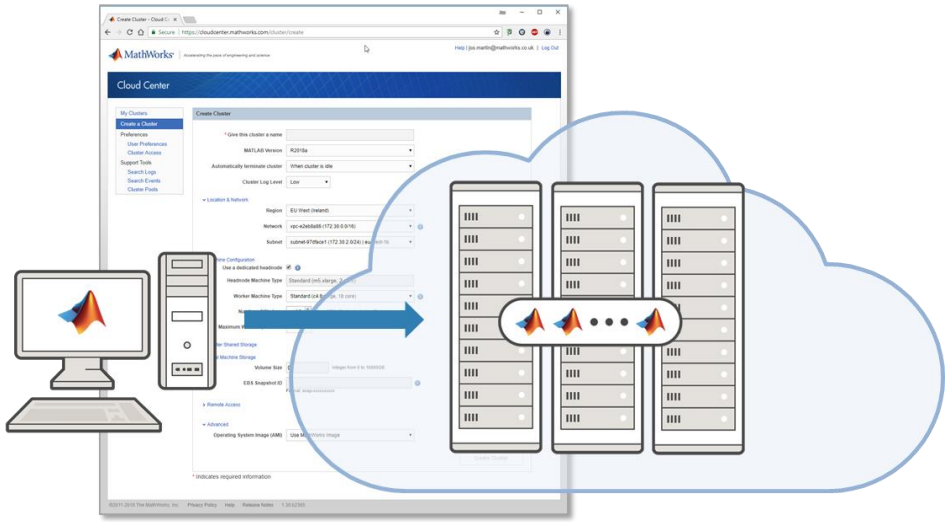

MathWorks Hosted License Manager

Just clicks away from previously activated licenses

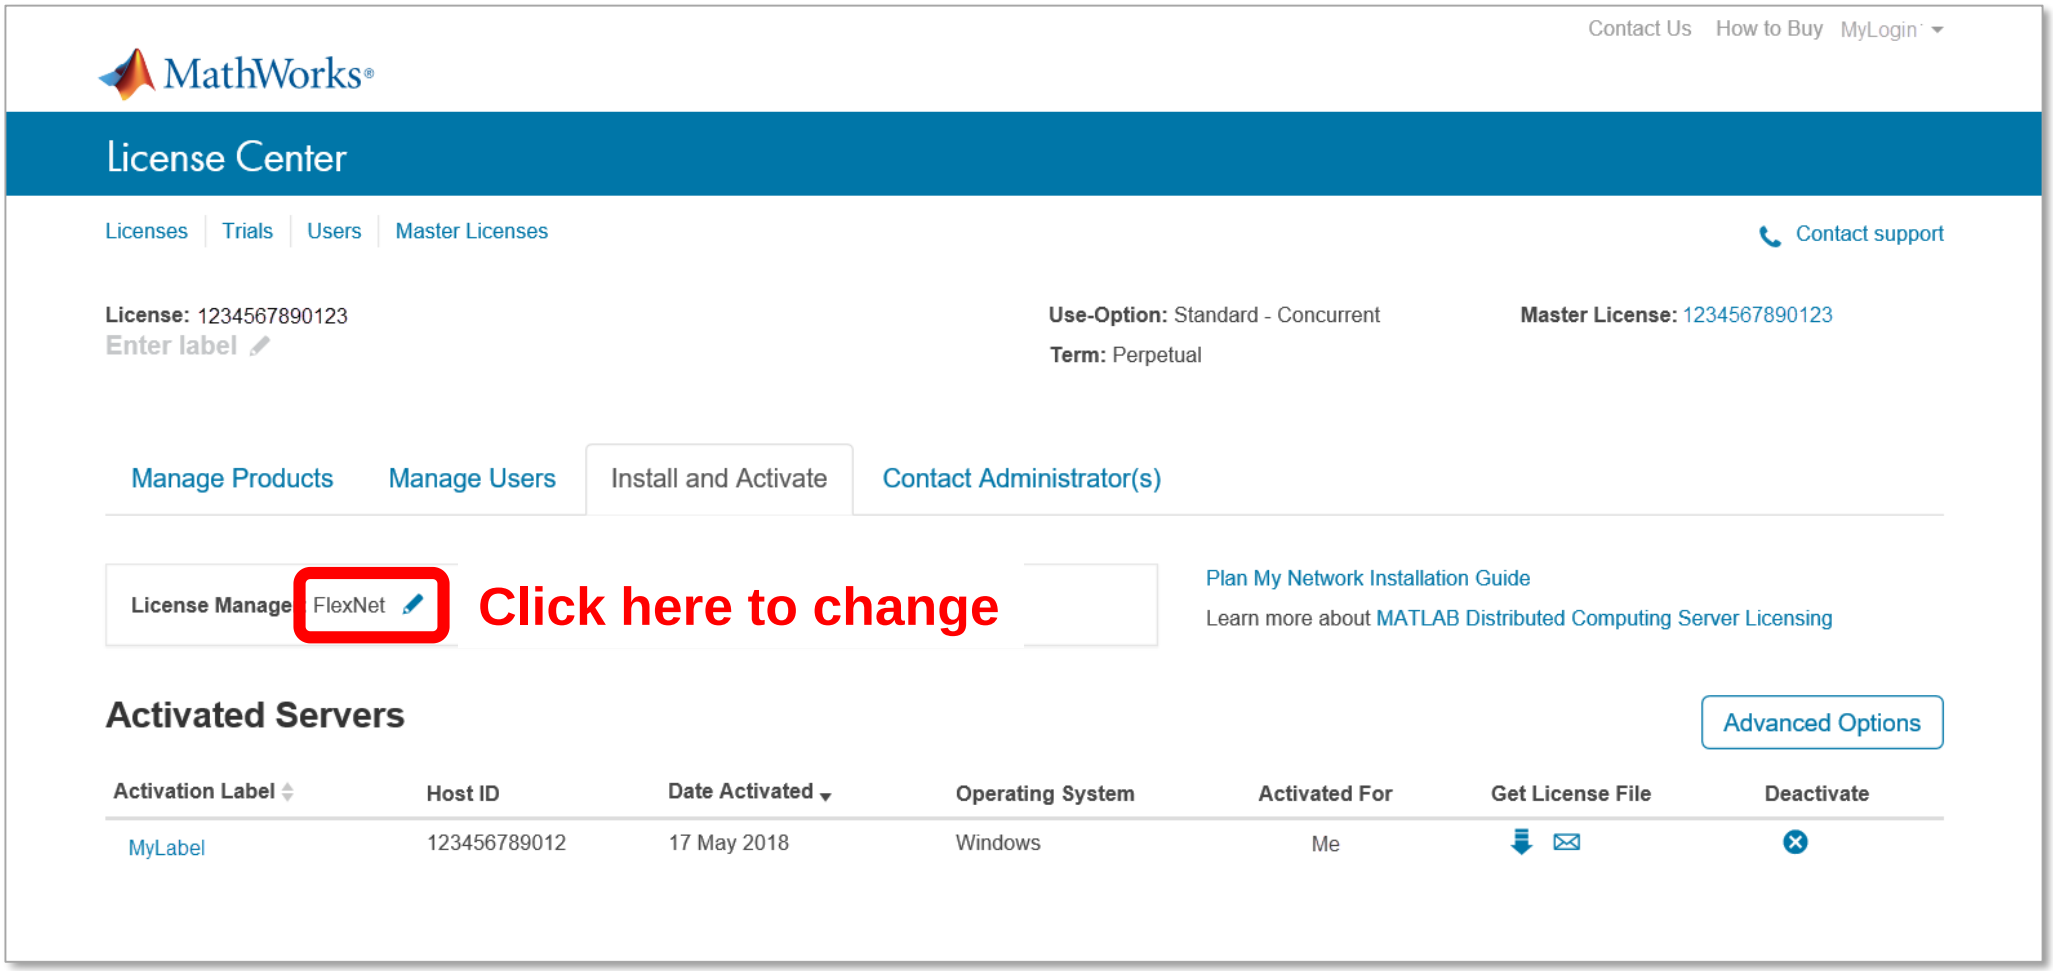

www.mathworks.com/licensecenter

For help with converting trials, contact MathWorks Support

MATLAB Parallel Server

Reference Architecture

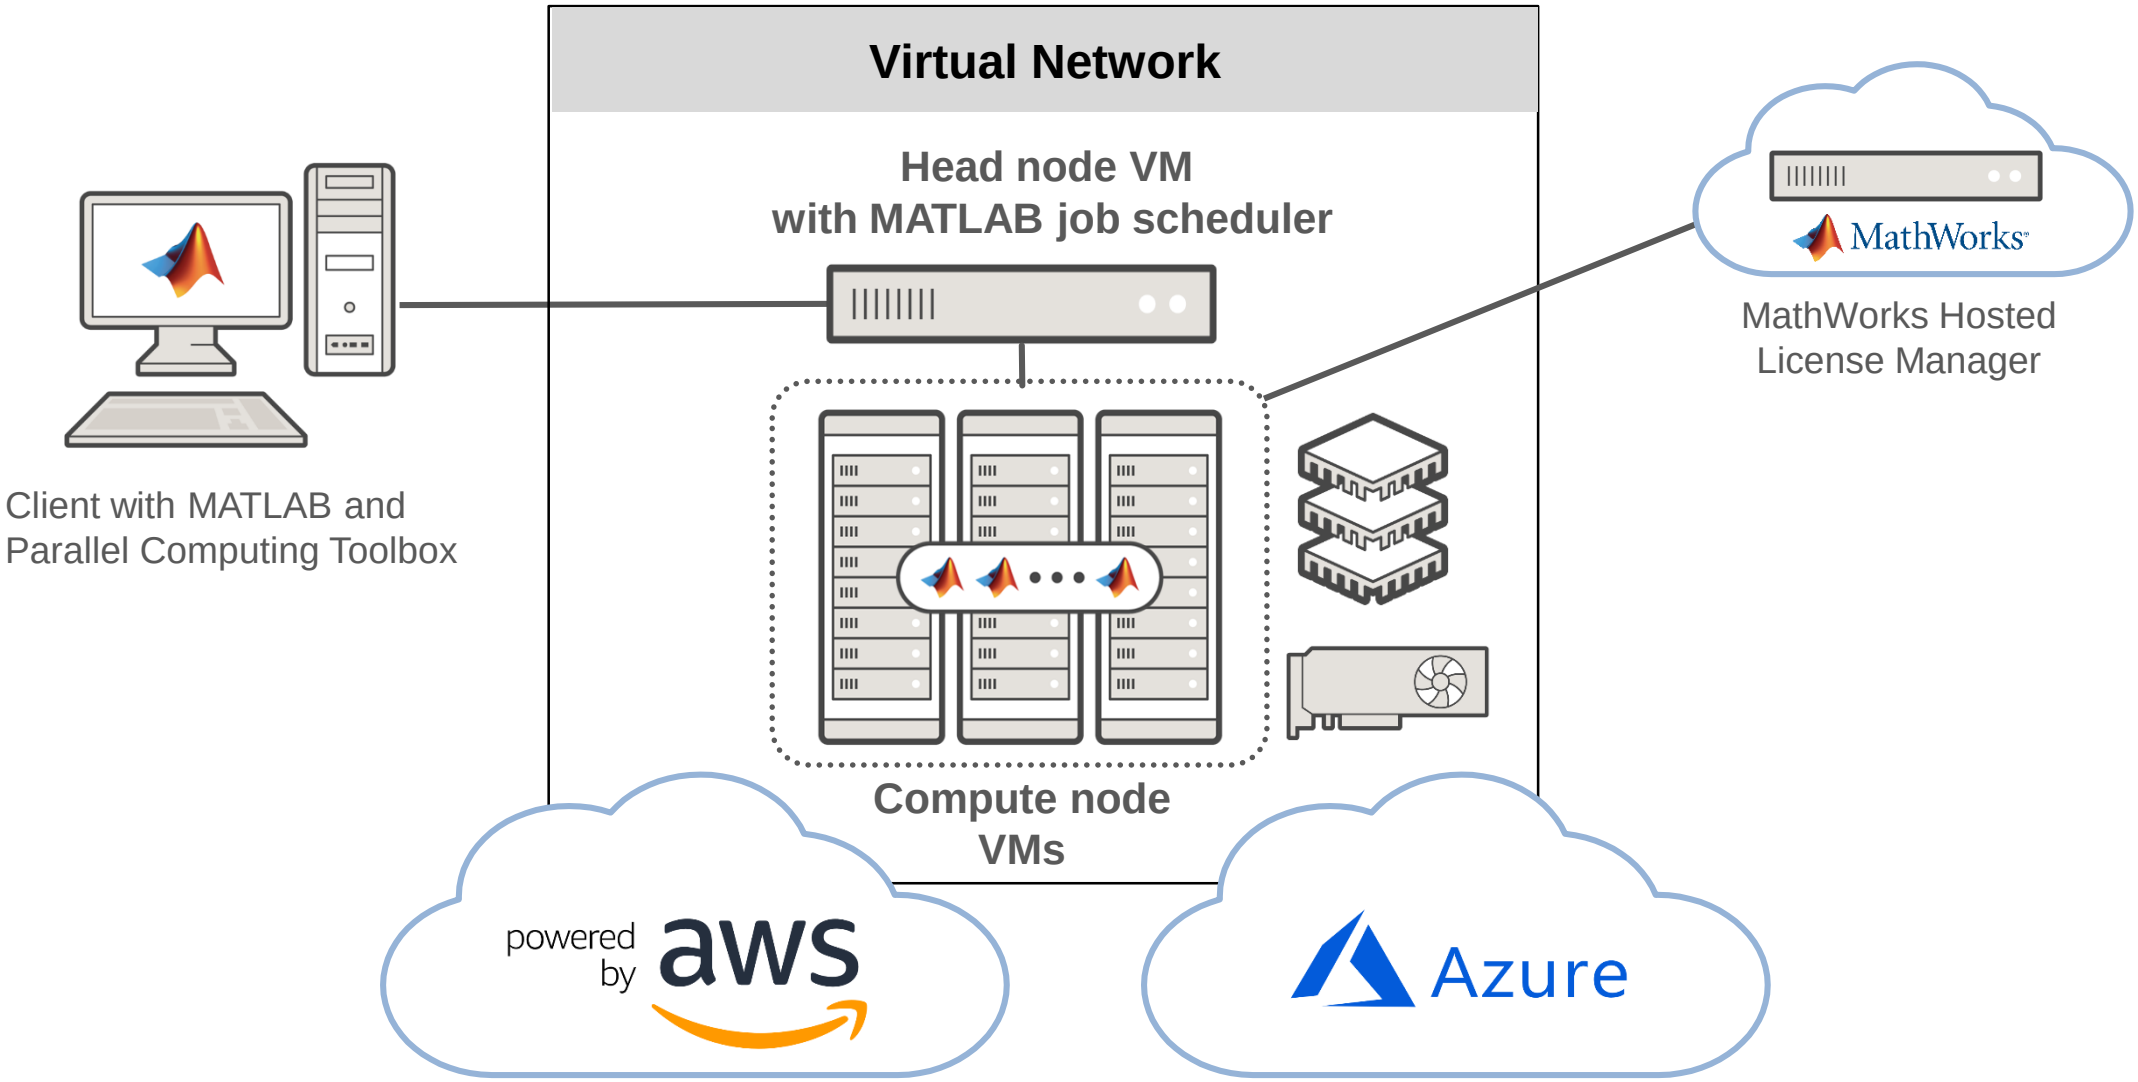

MATLAB Parallel Server

Reference Architecture

Use cases:

- Parameter sweeps
- Monte Carlo runs
- Optimization
- Distributed array calculations

Benefits:

- Quick infrastructure setup in the cloud
- MathWorks software pre-installed
- Incorporates best practices
- Adapt or extend for your specific needs

Azure ARM Template

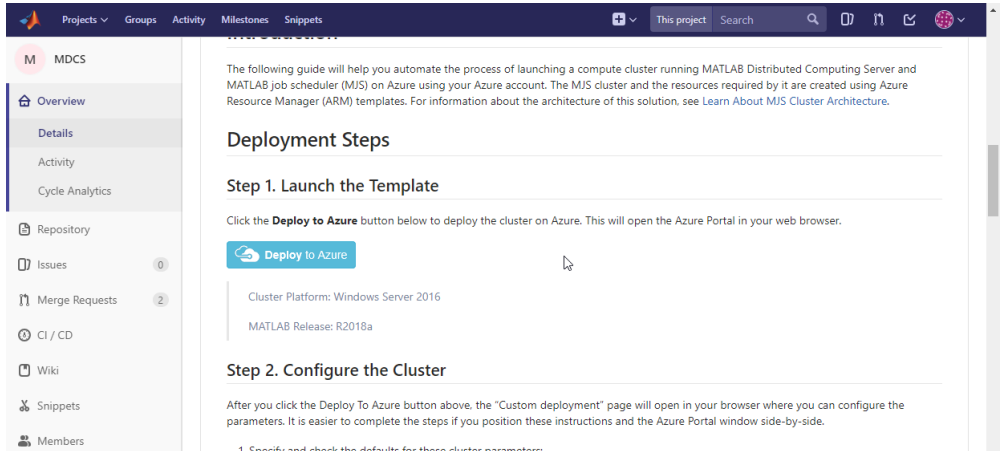

AWS Cloudformation Template

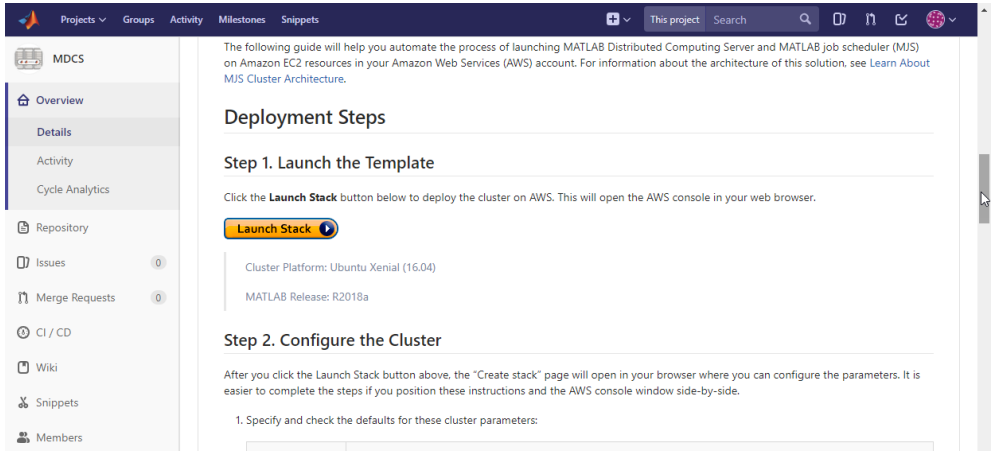

MATLAB Parallel Server

Reference Architecture

Requirements:

- MATLAB and Parallel Computing Toolbox (R2018a) on your desktop
- MATLAB Distributed Computing Server
- Amazon Web Services account
- SSH Key Pair for your AWS account in the US East (N. Virginia) region

Deployment steps:

1. Launch the template 

2. Configure the cloud resources, e.g.
   - Cluster name
   - Number of AWS instances, workers per instance
   - Instance type for the head node

   < creation may take up to 10 minutes>
3. Connect to the cluster from MATLAB

Aberdeen Asset Management Implements Machine Learning Based Portfolio Allocation Models in the Cloud

Challenge

Improve asset allocation strategies by creating model portfolios with machine learning techniques

Solution

Use MATLAB to develop classification tree, neural network, and support vector machine models, and use MATLAB Distributed Computing Server to run the models in the cloud

Results

- Portfolio performance goals supported
- Processing times cut from 24 hours to 3
- Multiple types of data easily accessed

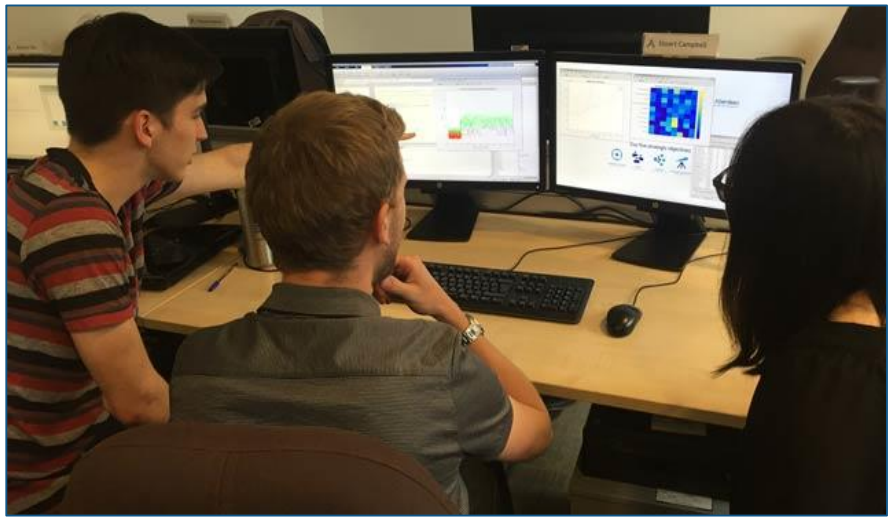

Interns using MATLAB at Aberdeen Asset Management.

"The widespread use of MATLAB in the finance community is a real advantage. Many university students learn MATLAB and can contribute right away when they join our team during internship programs. In addition, the strong MATLAB libraries developed by academic researchers help us explore all the possibilities of this programming language."

- Emilio Llorente-Cano, Aberdeen Asset Management

Link to user story

Aberdeen: Modeling on Azure

Aberdeen Asset management

Investment strategy and portfolio construction

- IT team not able to implement large scale cluster in-house
- Needed elastic resources for sporadic usage

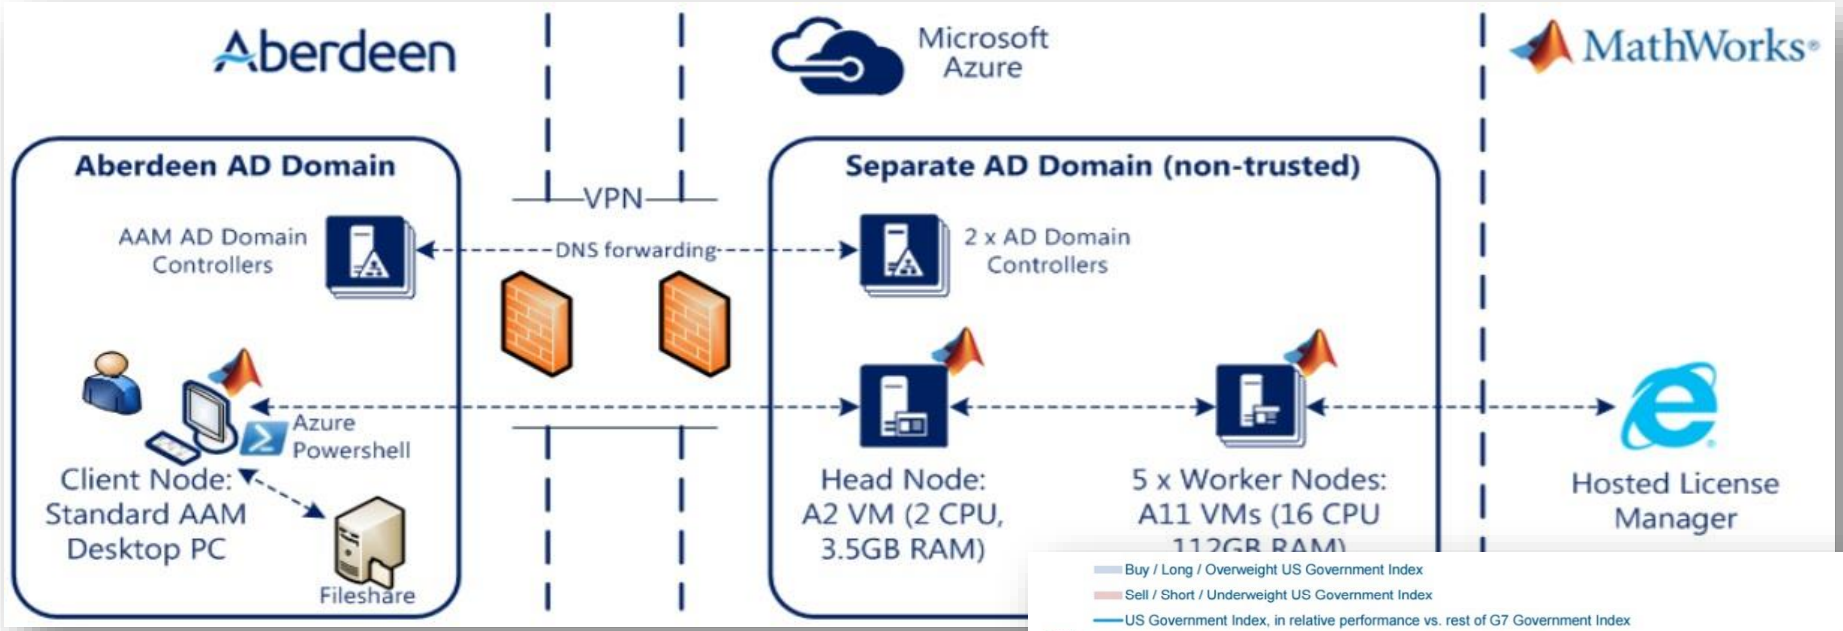

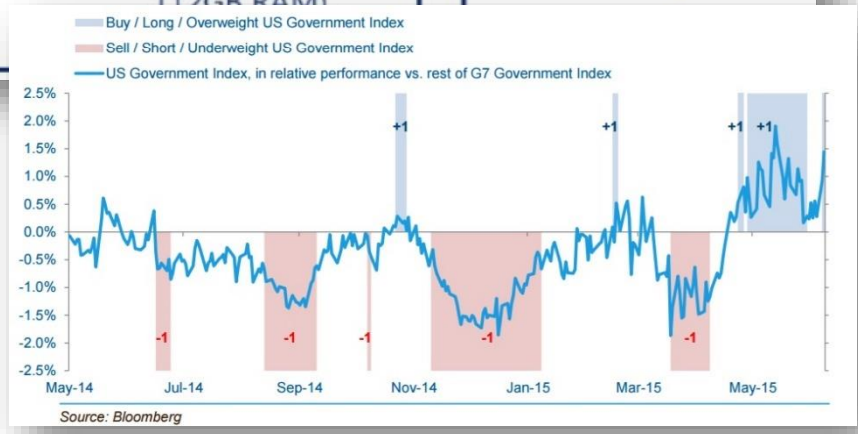

Deploying in the cloud

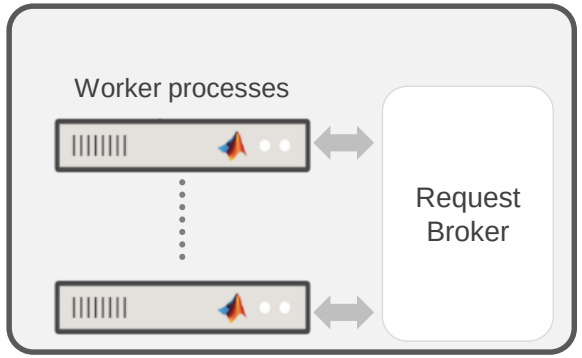

MATLAB Production Server

Share Programs Outside of MATLAB

Application Deployment

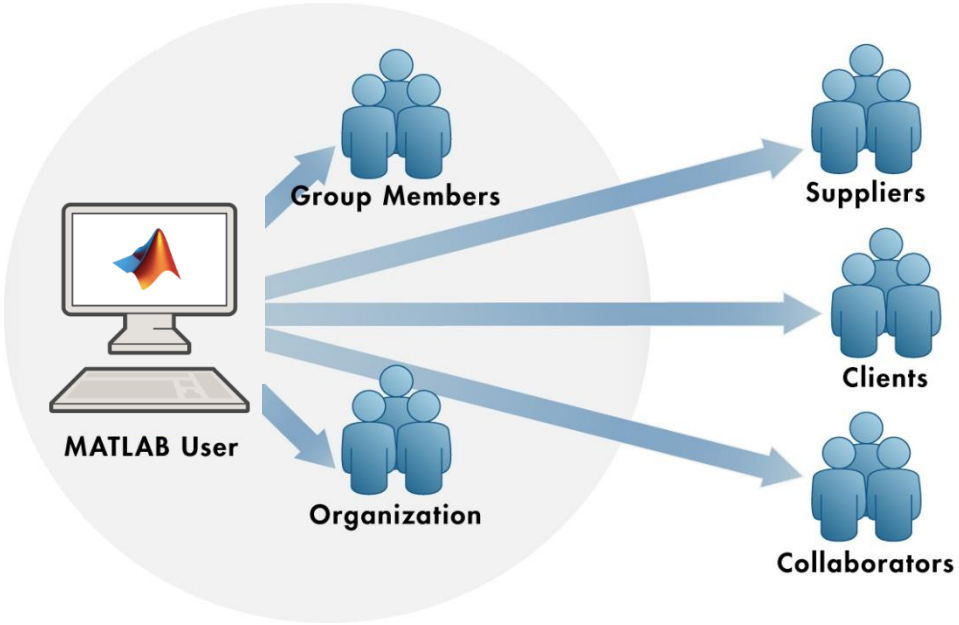

Desktop

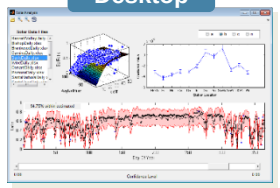

Web Apps

Excel® add-ins

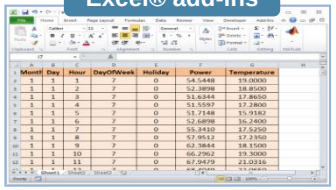

- Royalty-free
- Encryption to protect intellectual property

Enterprise

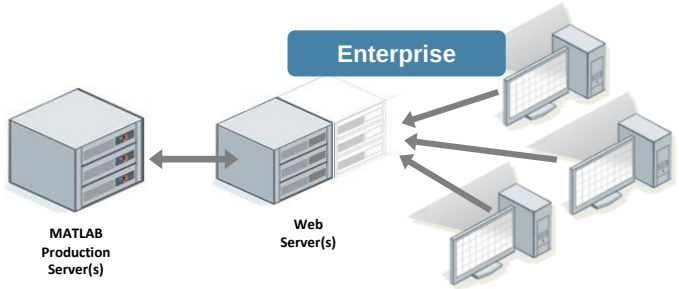

MATLAB Production Server is an application server

Publishes MATLAB code as APIs that can be called by other applications

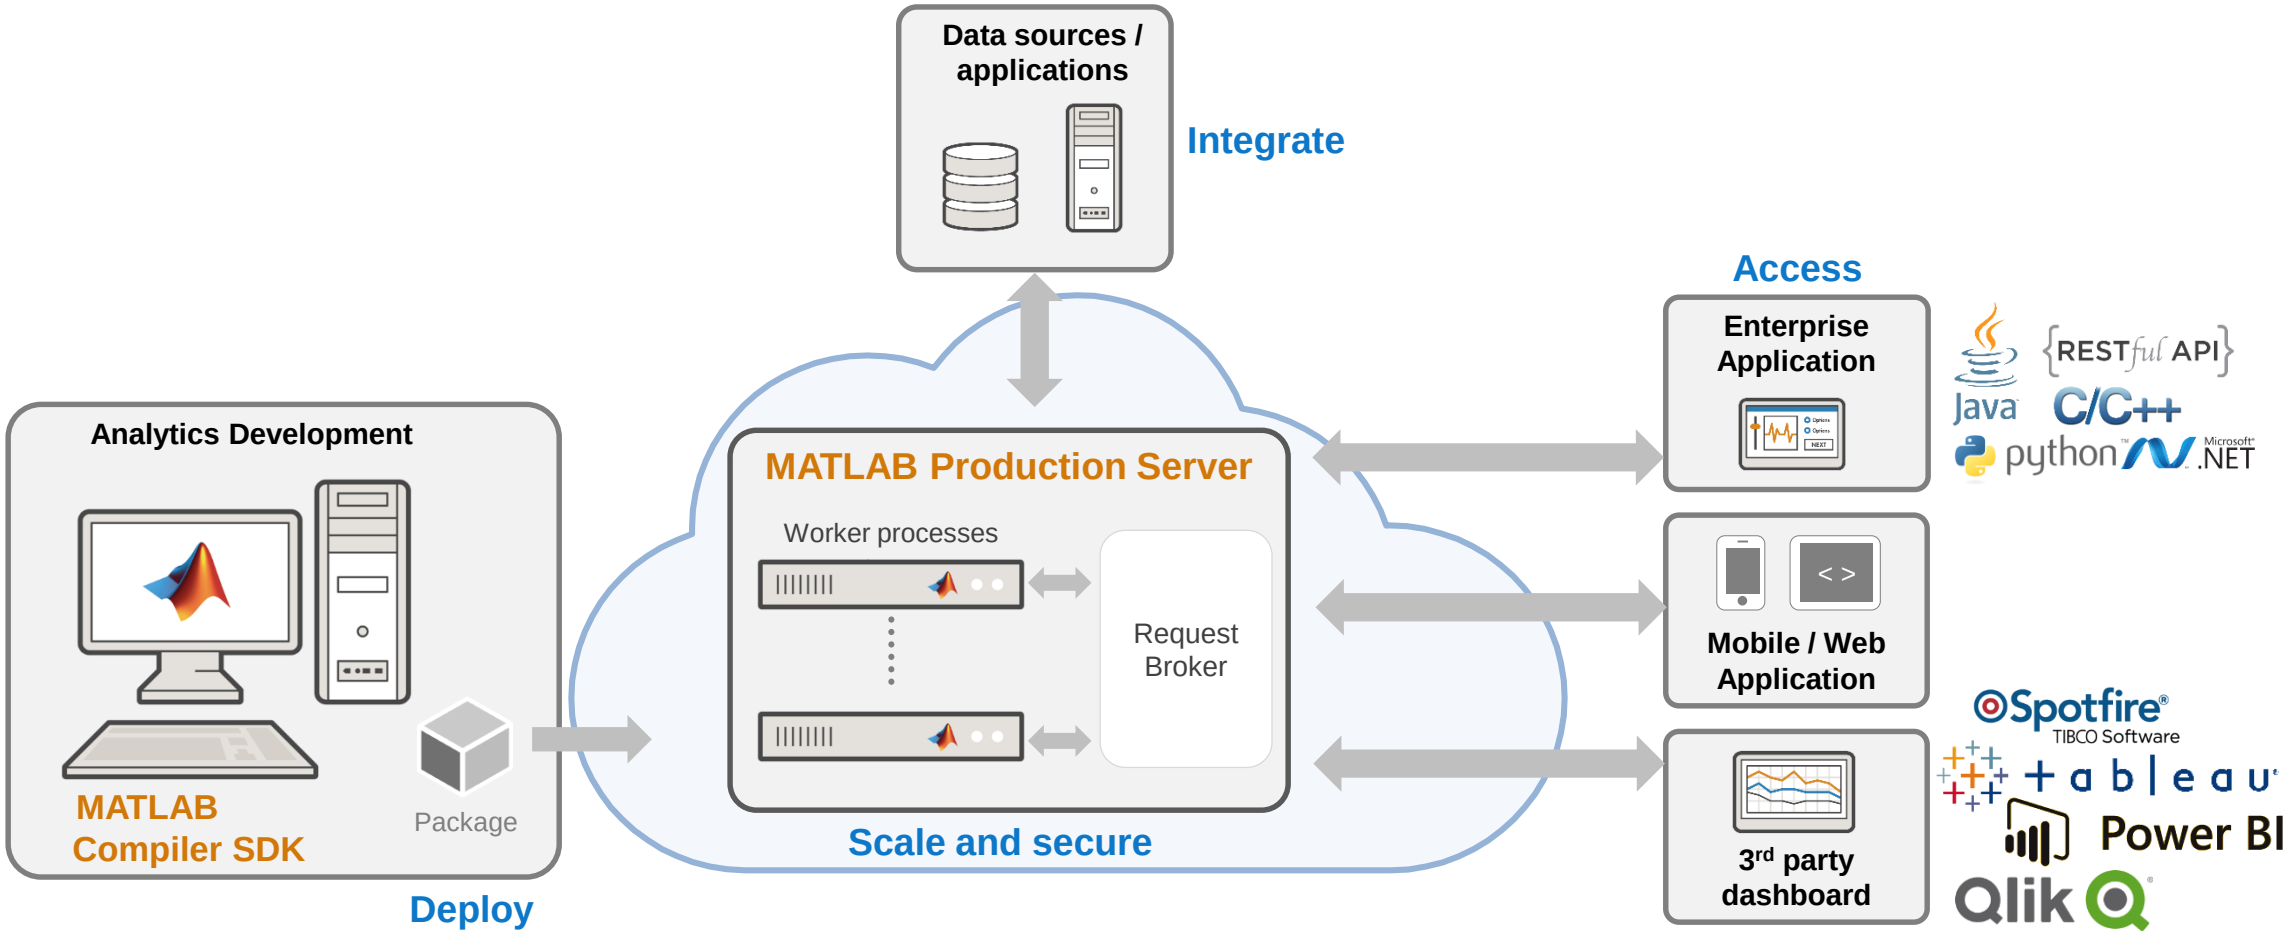

MATLAB Production Server

Reference Architecture

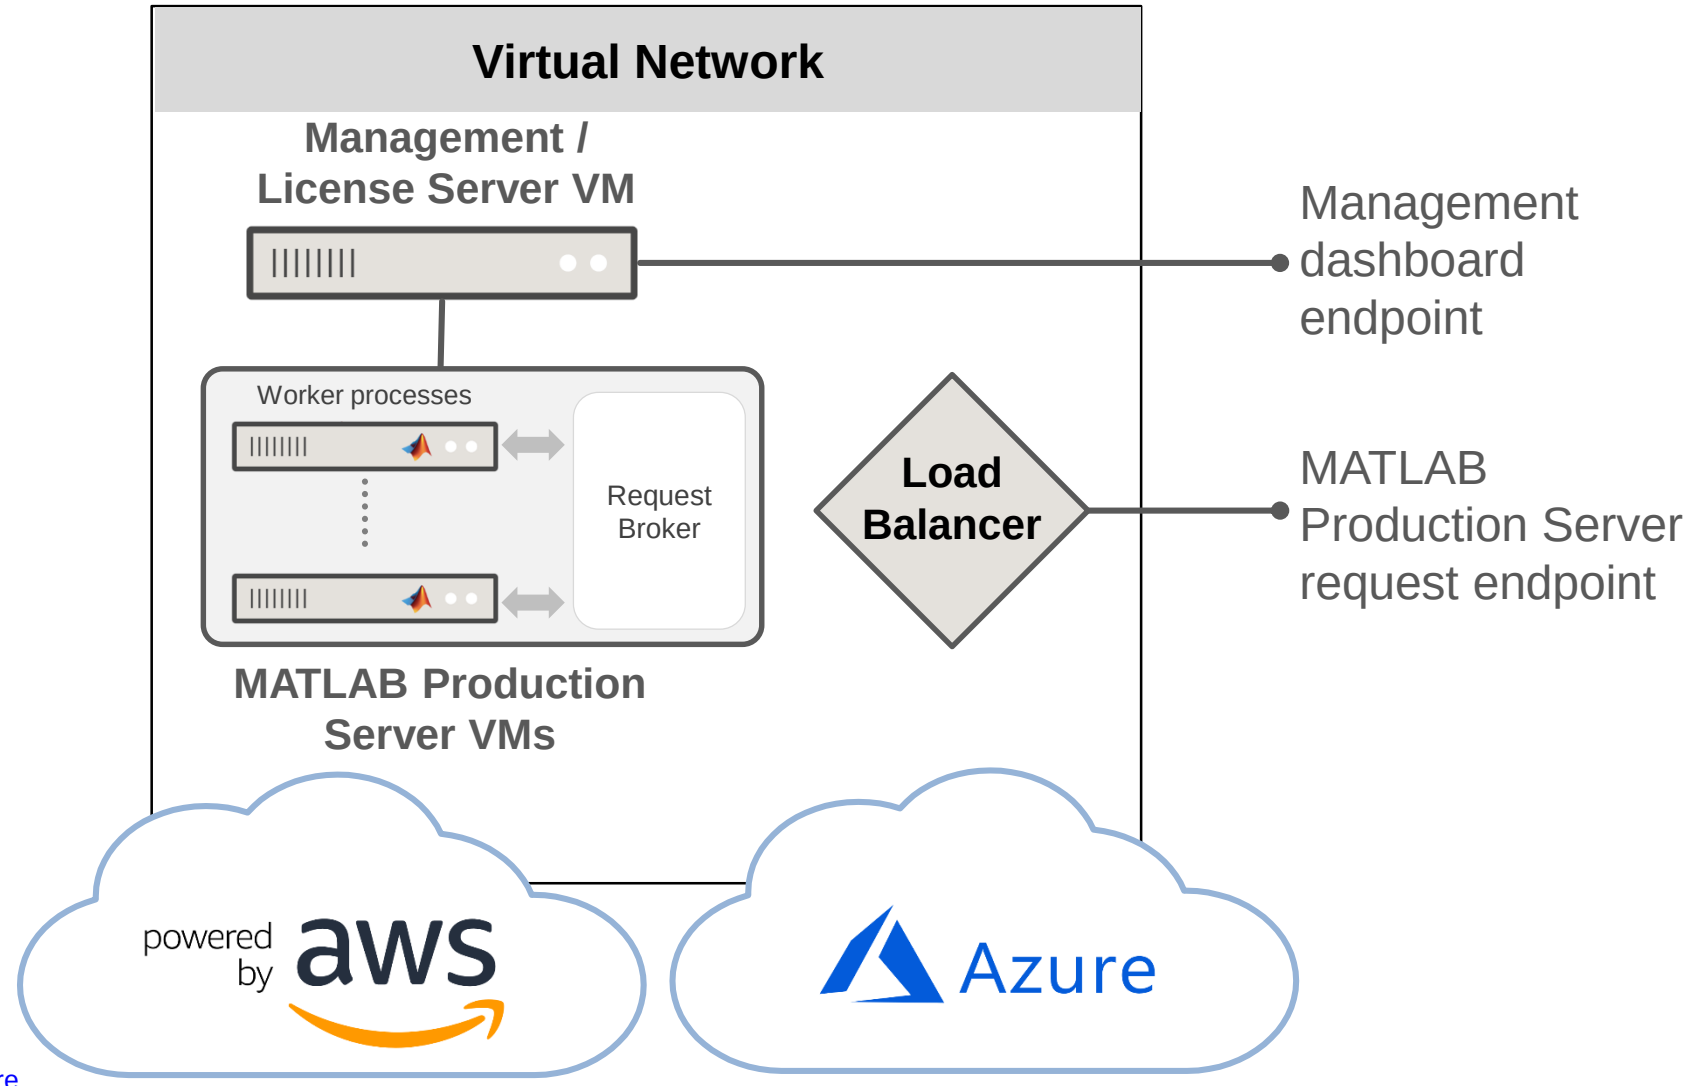

https://github.com/mathworks-ref-arch/mps-on-aws
https://github.com/mathworks-ref-arch/mps-on-azure

MATLAB Production Server

Reference Architecture

Use cases:

- Scalable deployed algorithms and analytics
- Application server for MATLAB algorithms
- IoT / streaming data analytics

Benefits:

- Quick infrastructure setup in the cloud
- MathWorks software pre-installed
- Incorporates best practices
- Adapt or extend for your specific needs

MATLAB Production Server

Reference Architecture

Requirements:

- MATLAB Production Server license
- Amazon Web Services (AWS) account
- SSH Key Pair for your AWS account in the US East (N. Virginia) region

Deployment steps:

1. Launch the template 

2. Configure the cloud resources, e.g.
   - Number of AWS instances
   - Workers per instance

   < creation may take up to 10 minutes>
3. Log in to the Cloud Console
4. Upload license file

Configure and manage server instances on the cloud

MATLAB Production Server Cloud Console

- Log in to the Cloud Console
- Download MAC address
  - Go to MathWorks License Center and get the MATLAB Production Server licens
  - Upload the license file
- Upload ctf files, monitoring logs, etc.

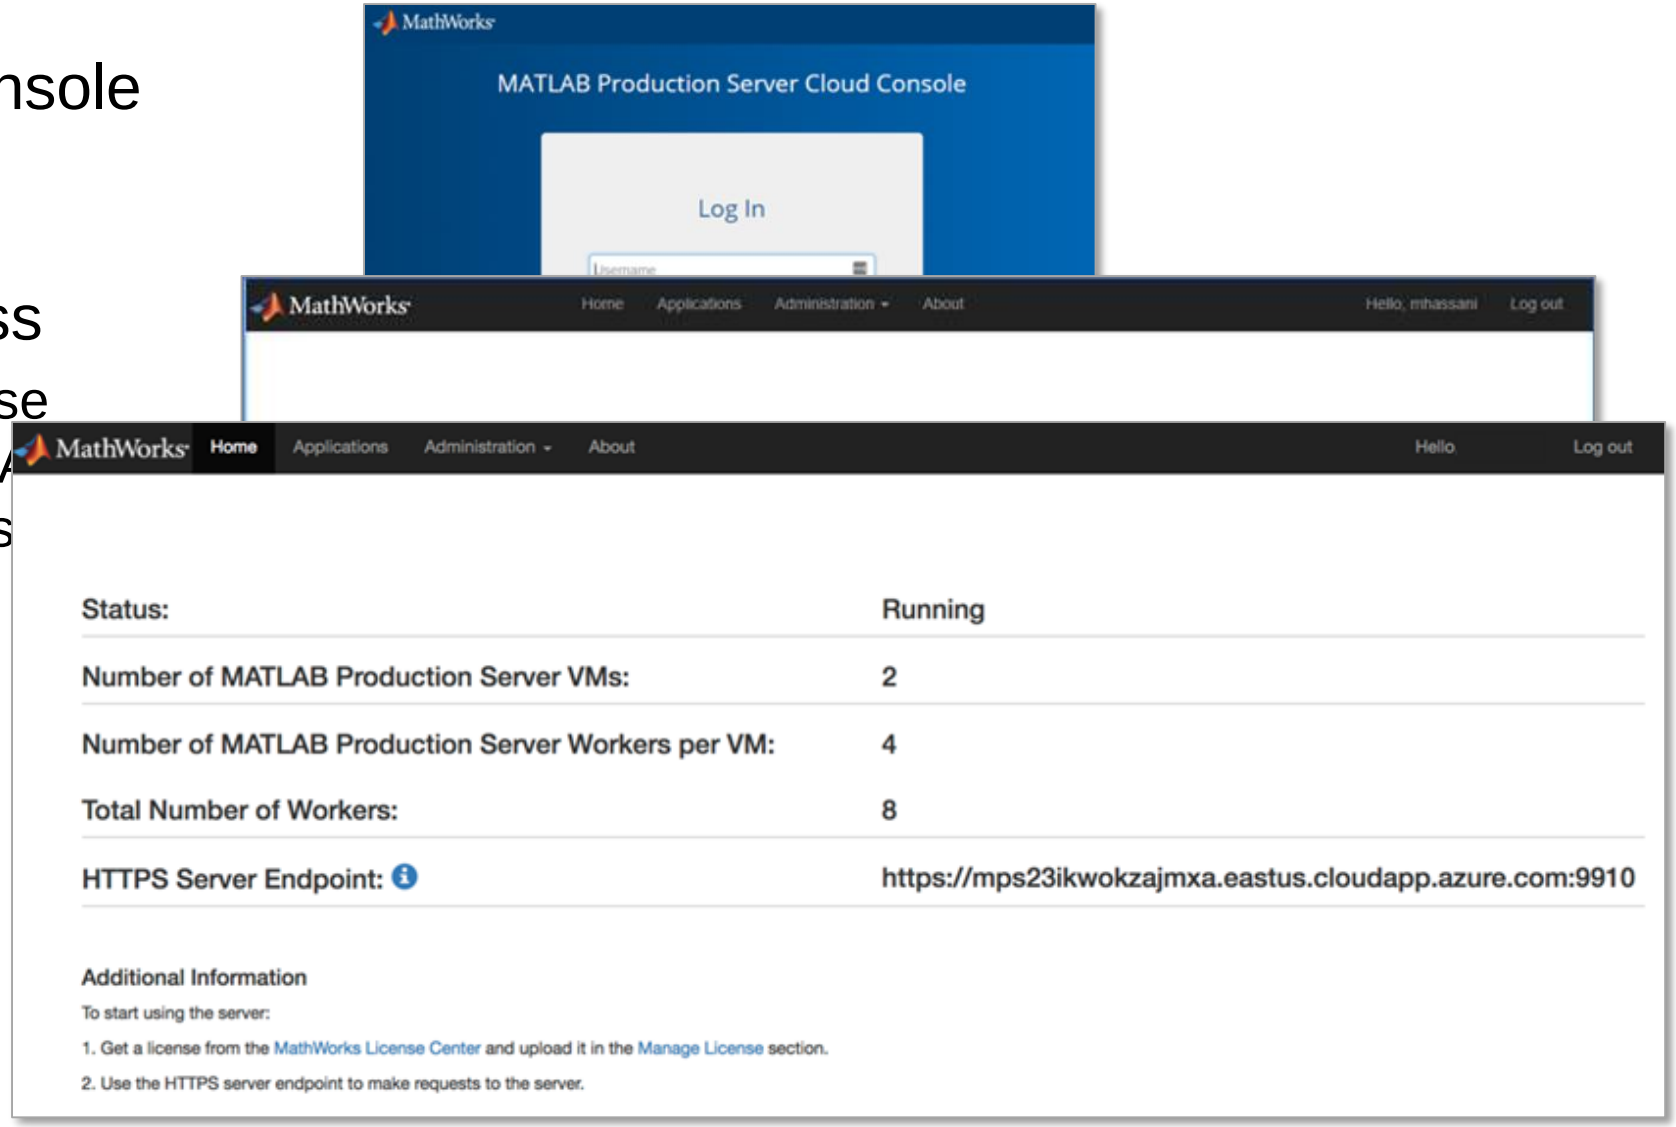

MATLAB Production Server cloud deployment in IoT scenario

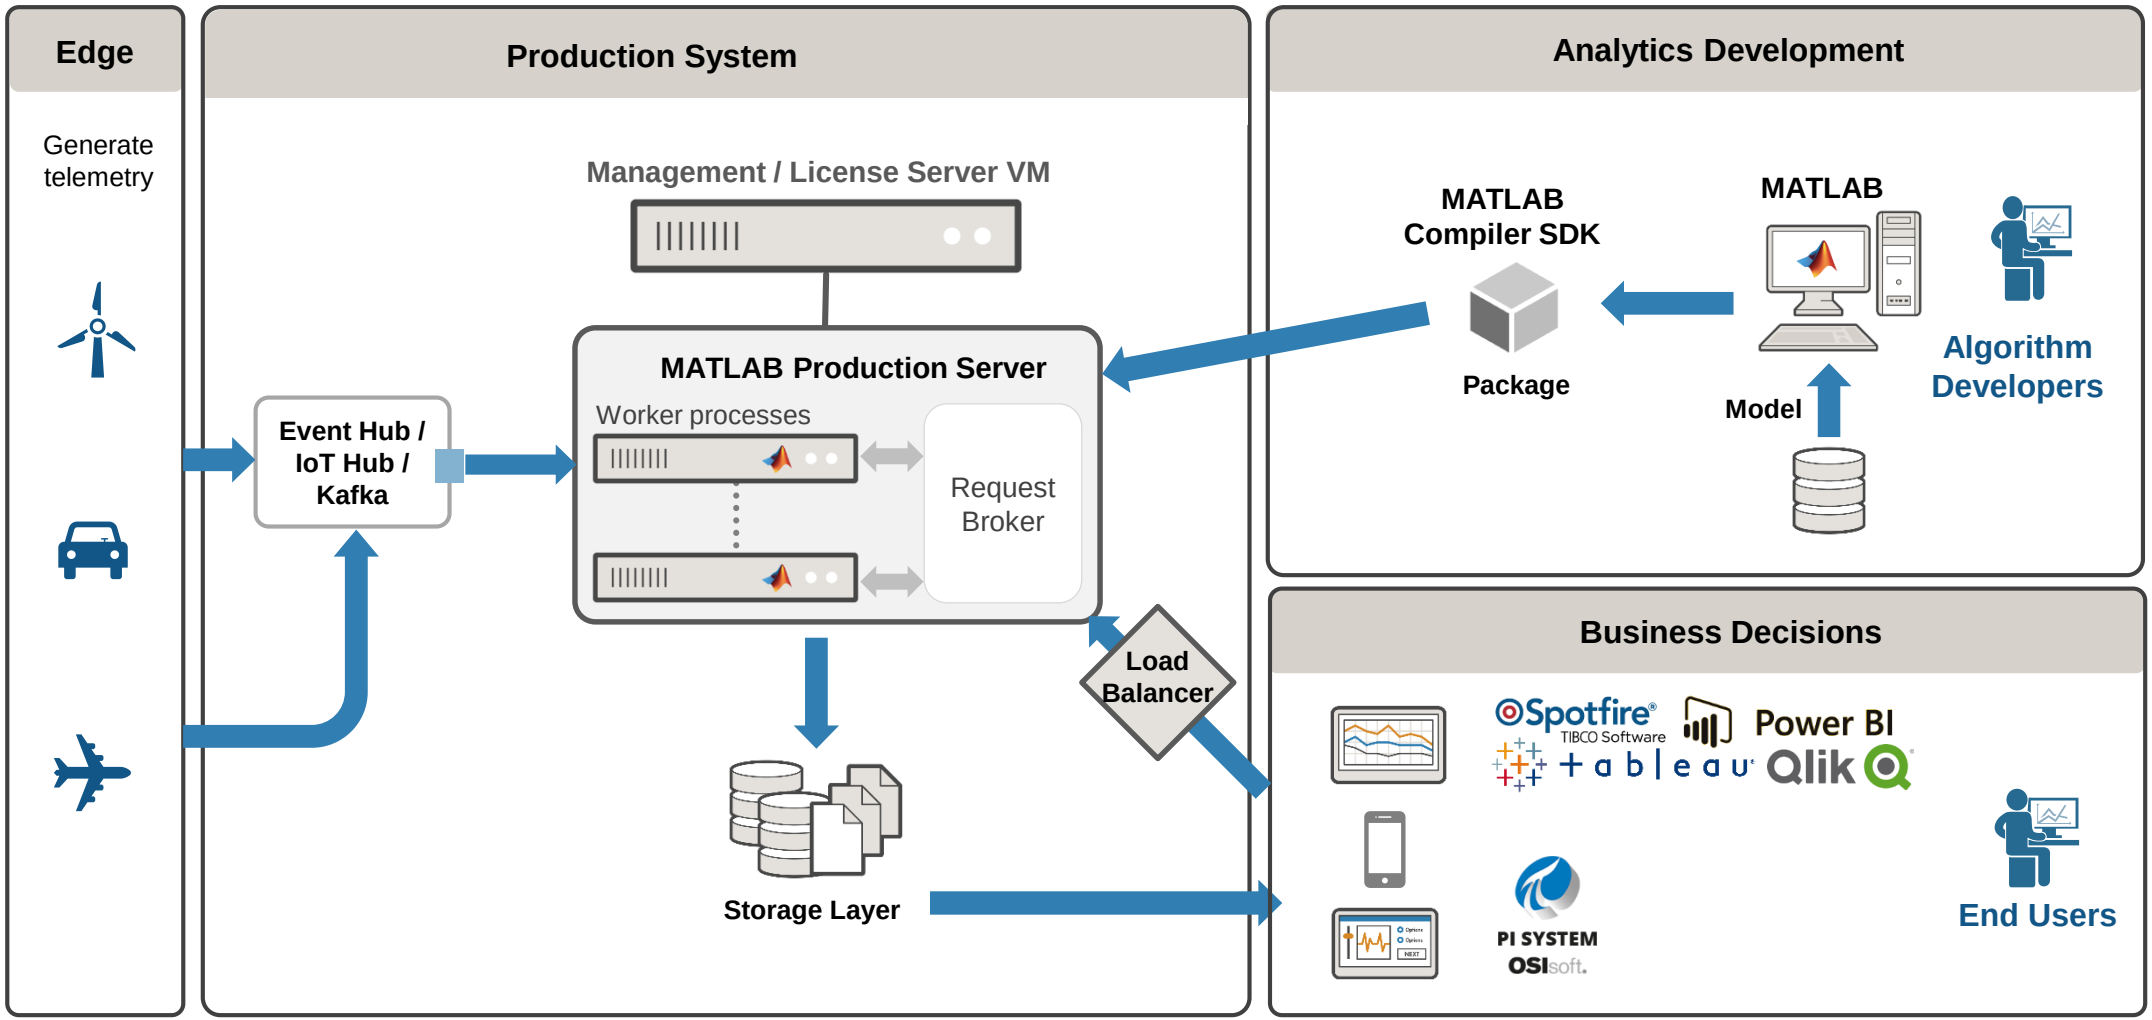

Customer Example: Healthcare Device Vendor

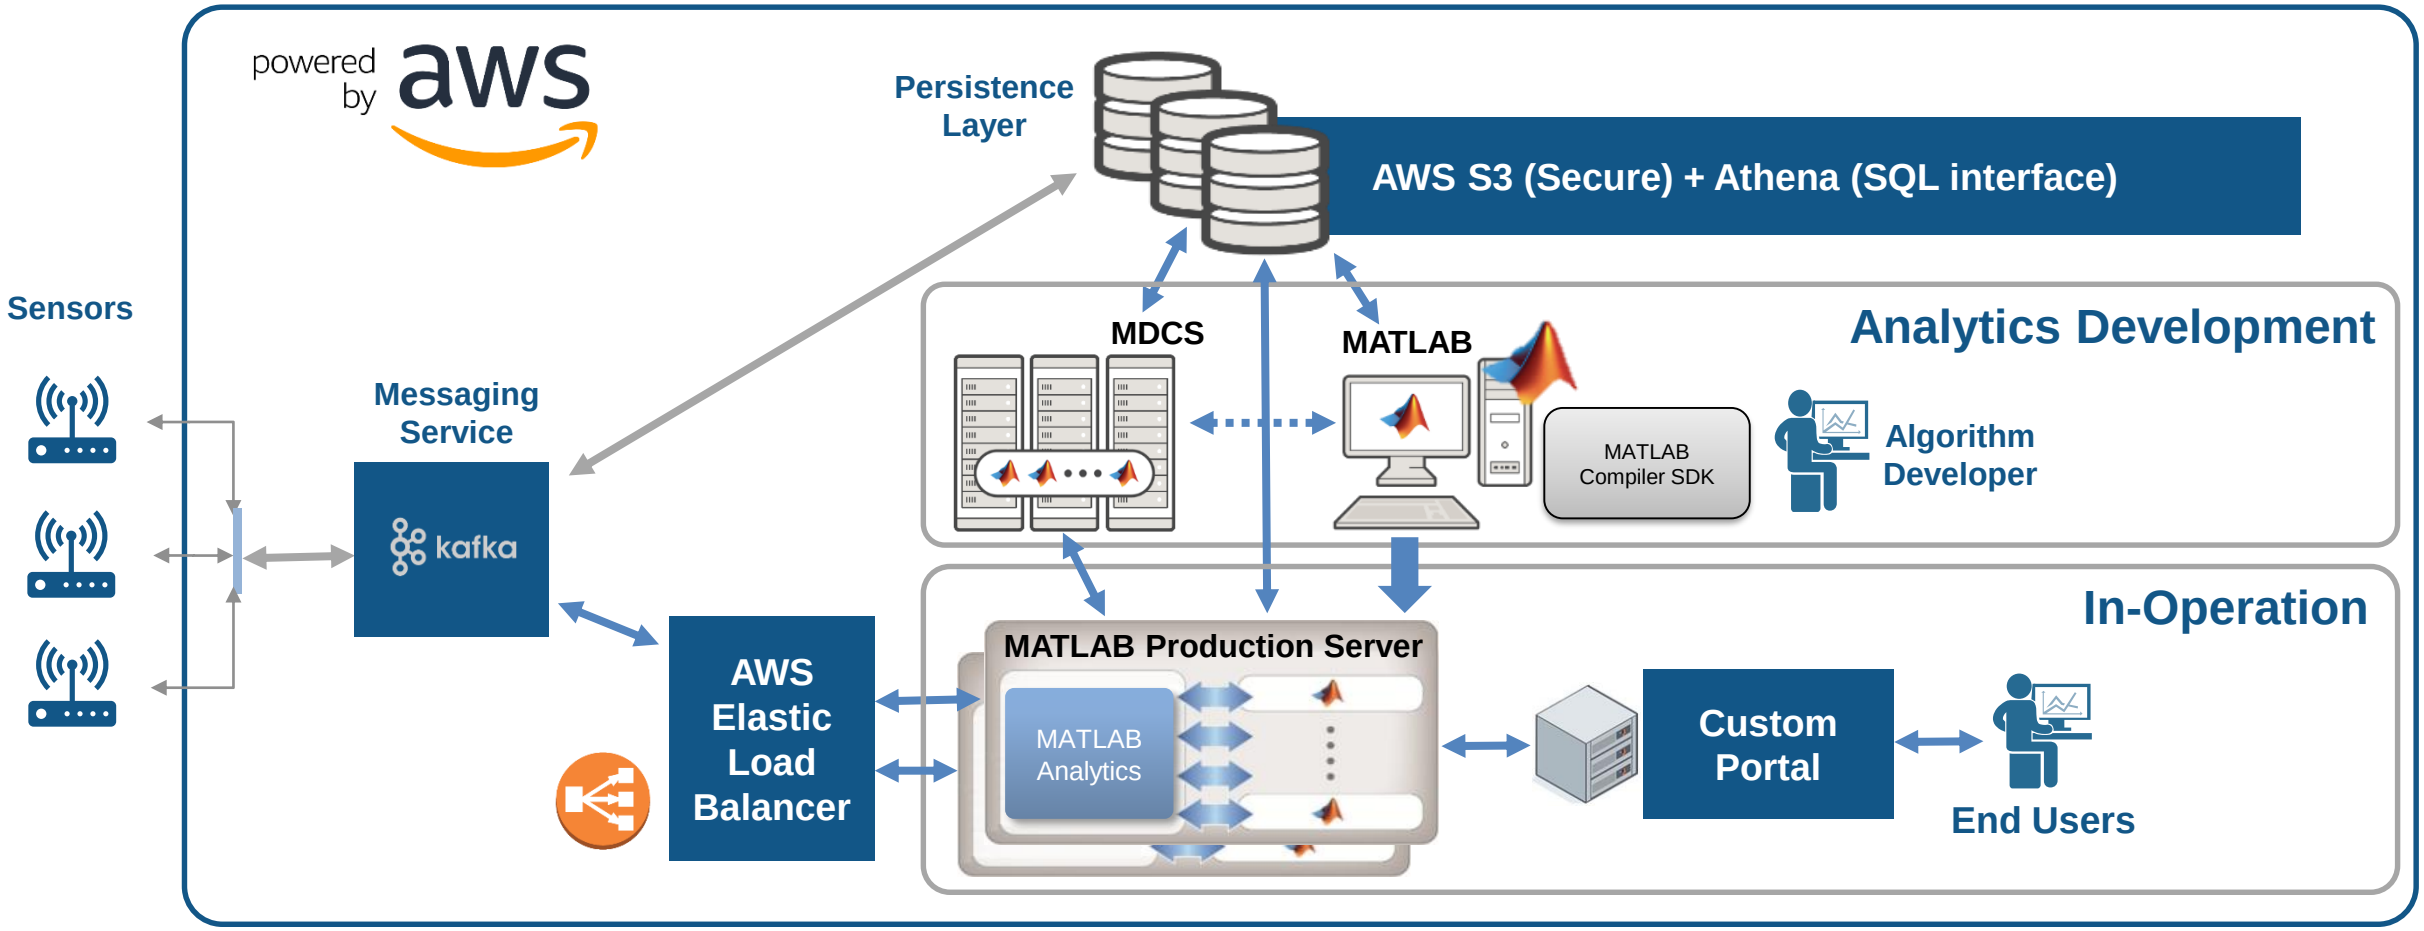

Typical data analytics workflow with MATLAB

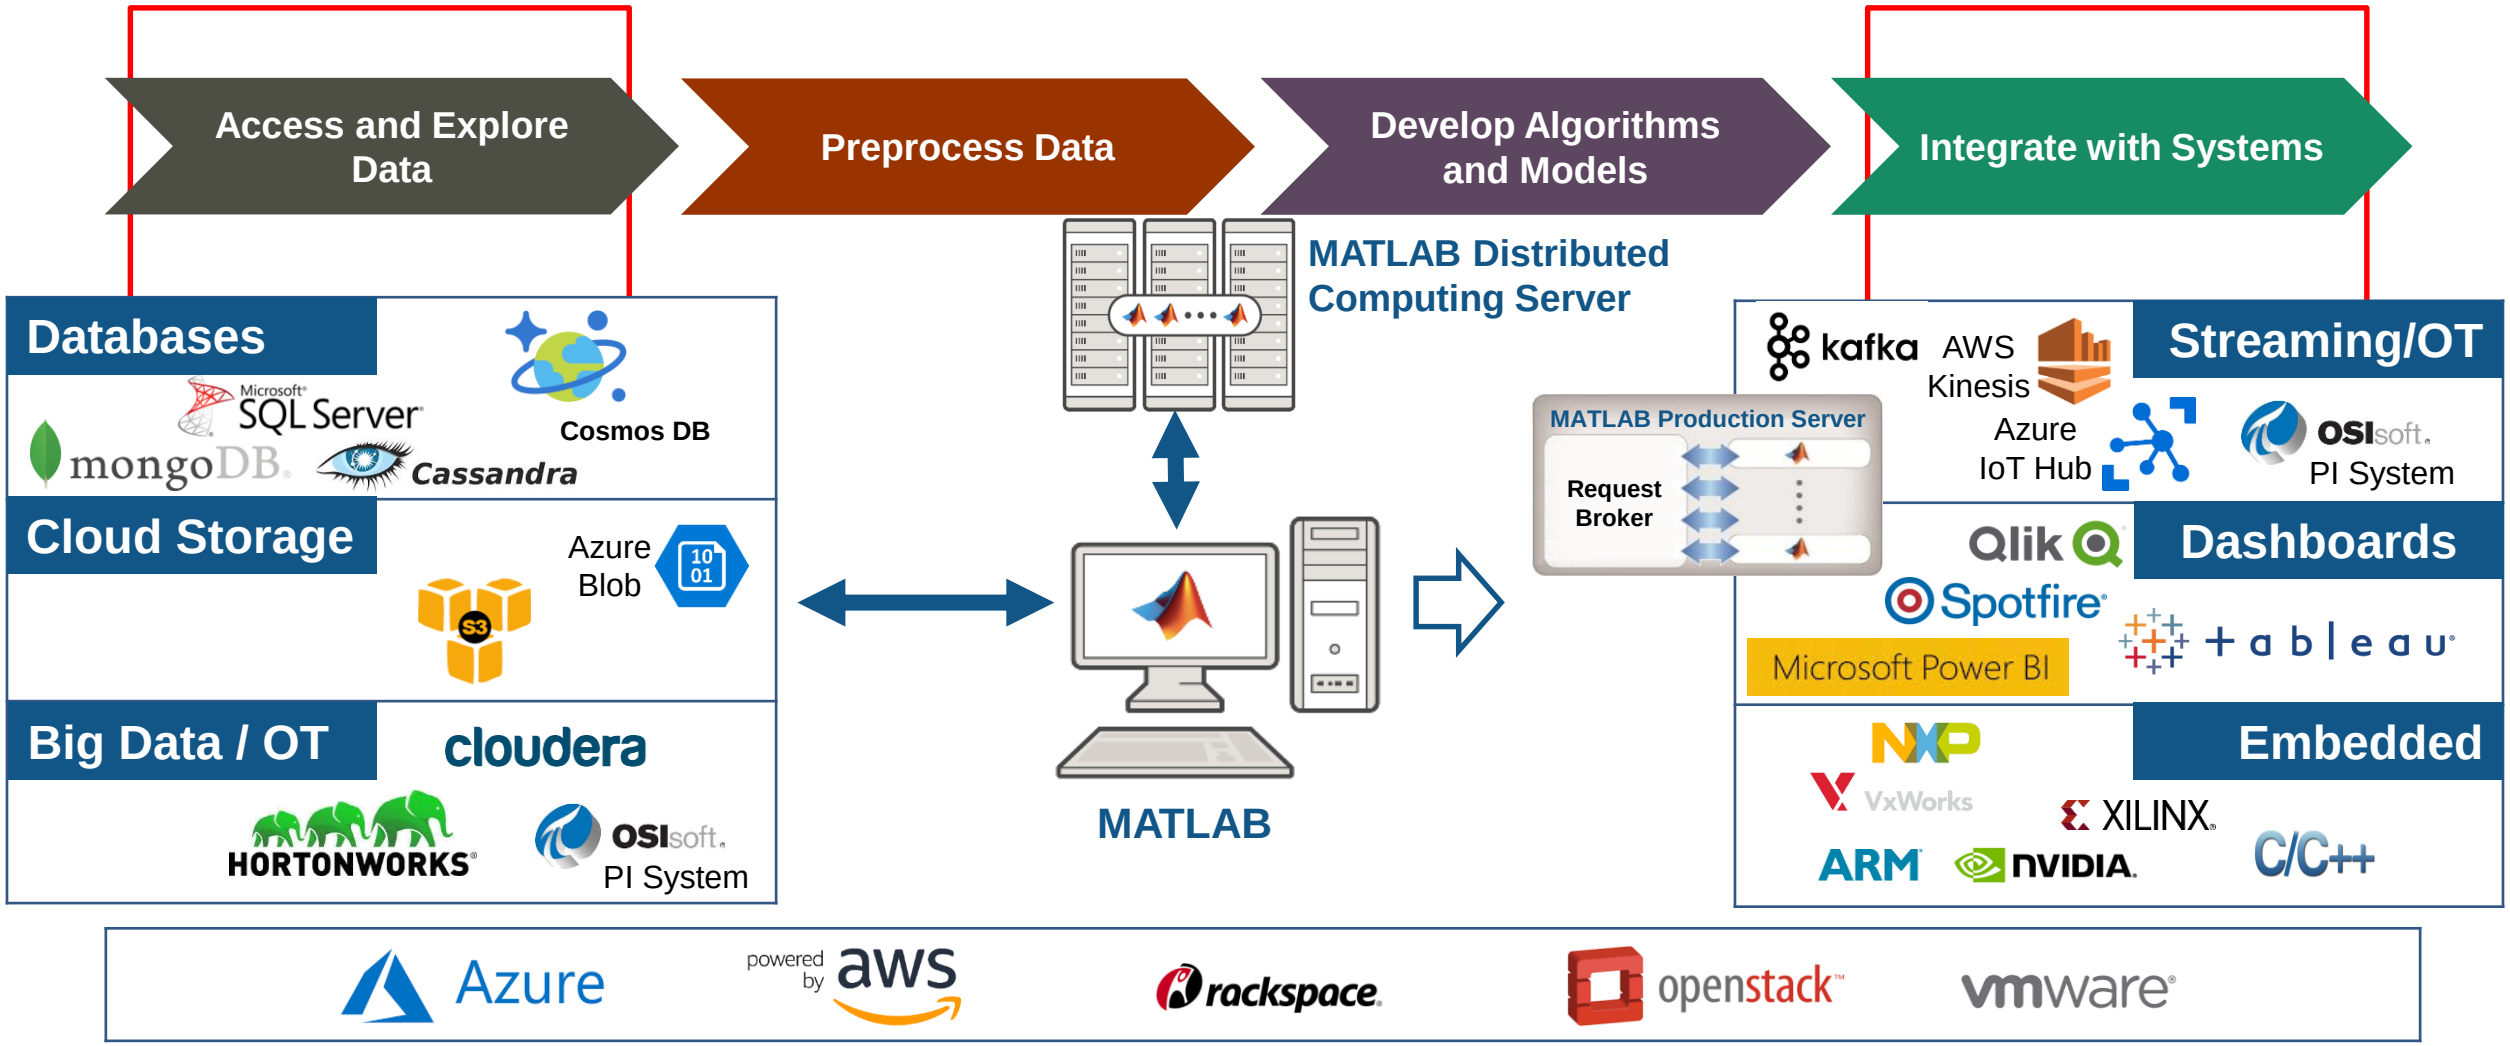

Using the Cloud for Deep Learning

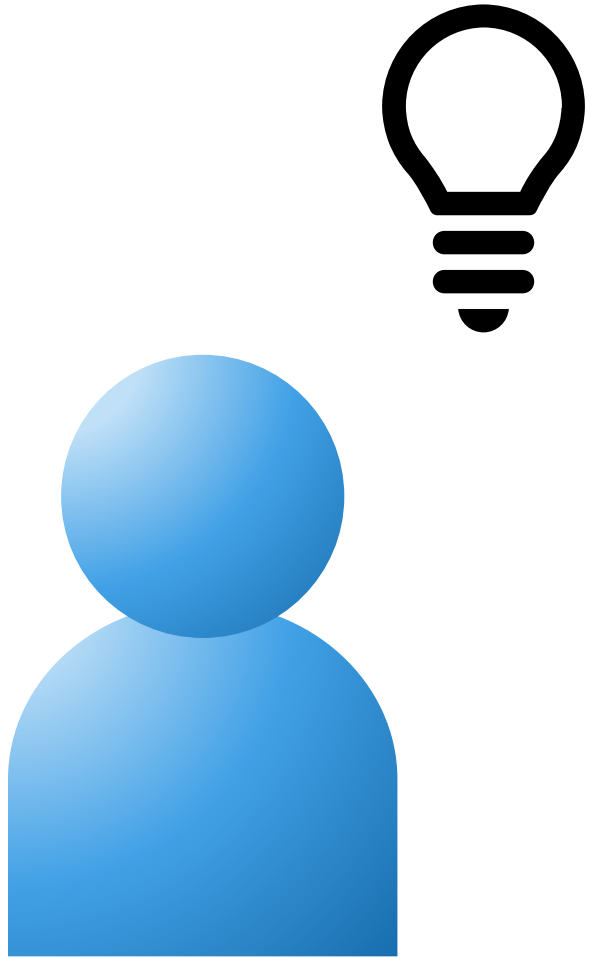

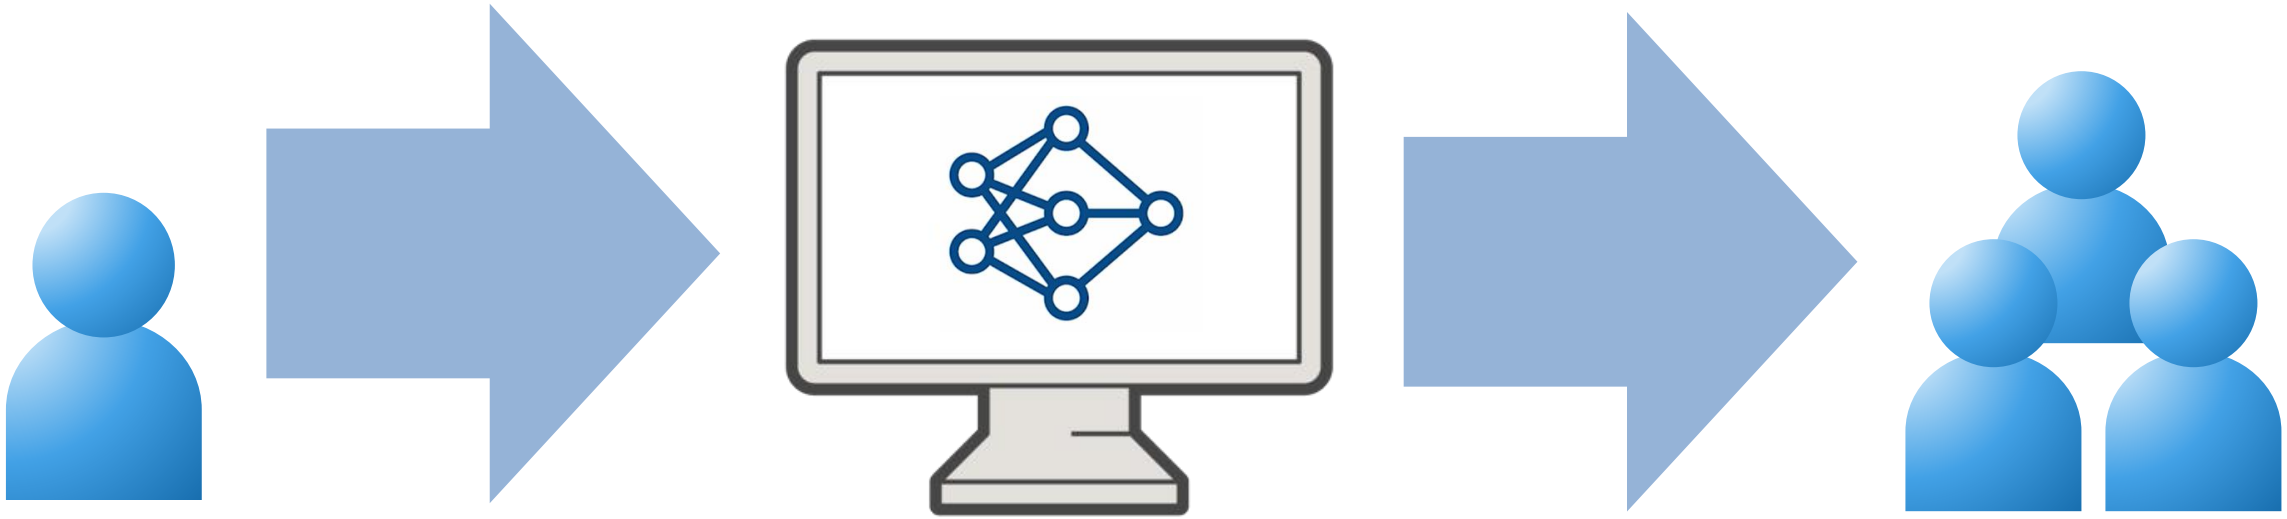

“Bell Pepper”

Quickly **prototype** and **test** your idea

Train the Neural Net model

Deploy algorithm for everyone to use

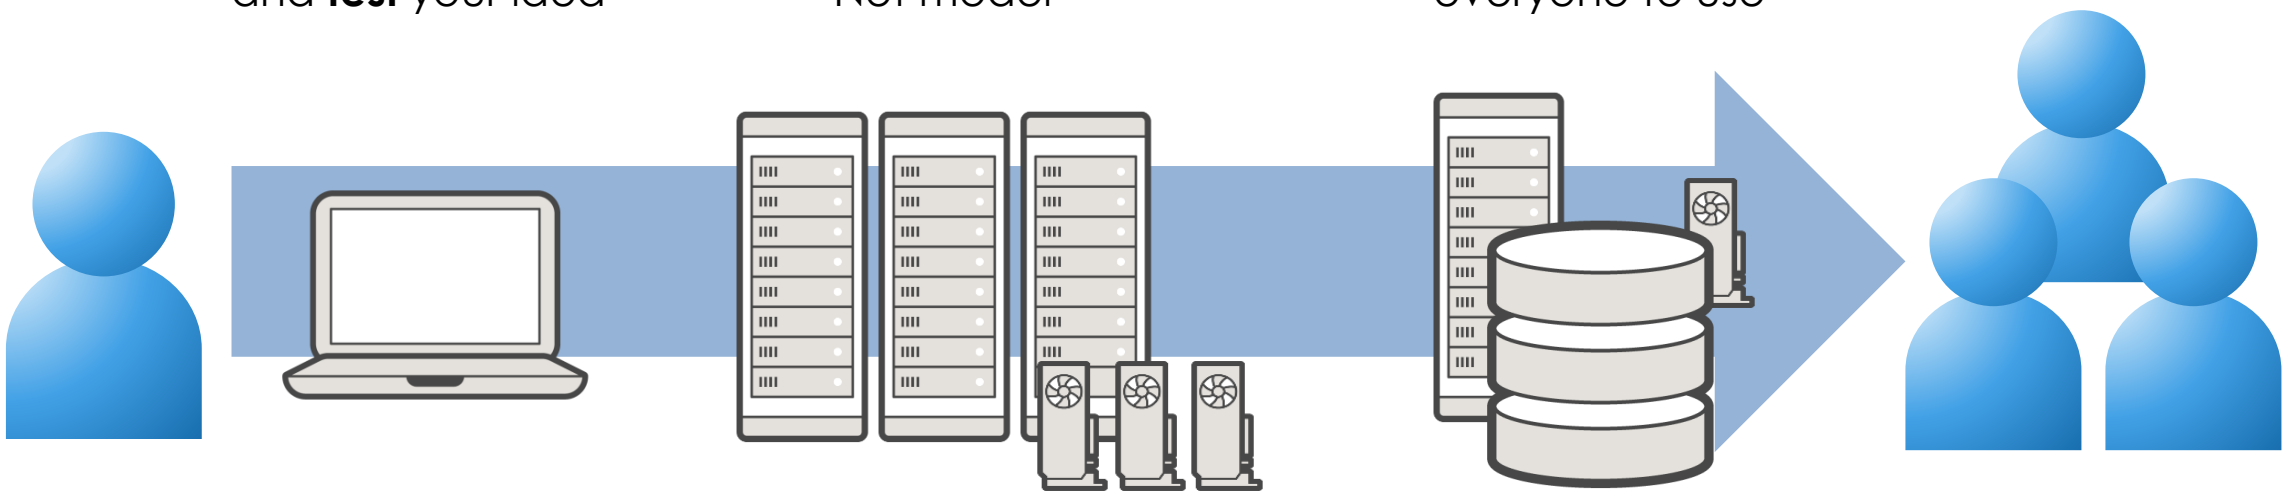

Quickly **prototype** and **test** your idea

Train the Neural Net model

Deploy algorithm for everyone to use

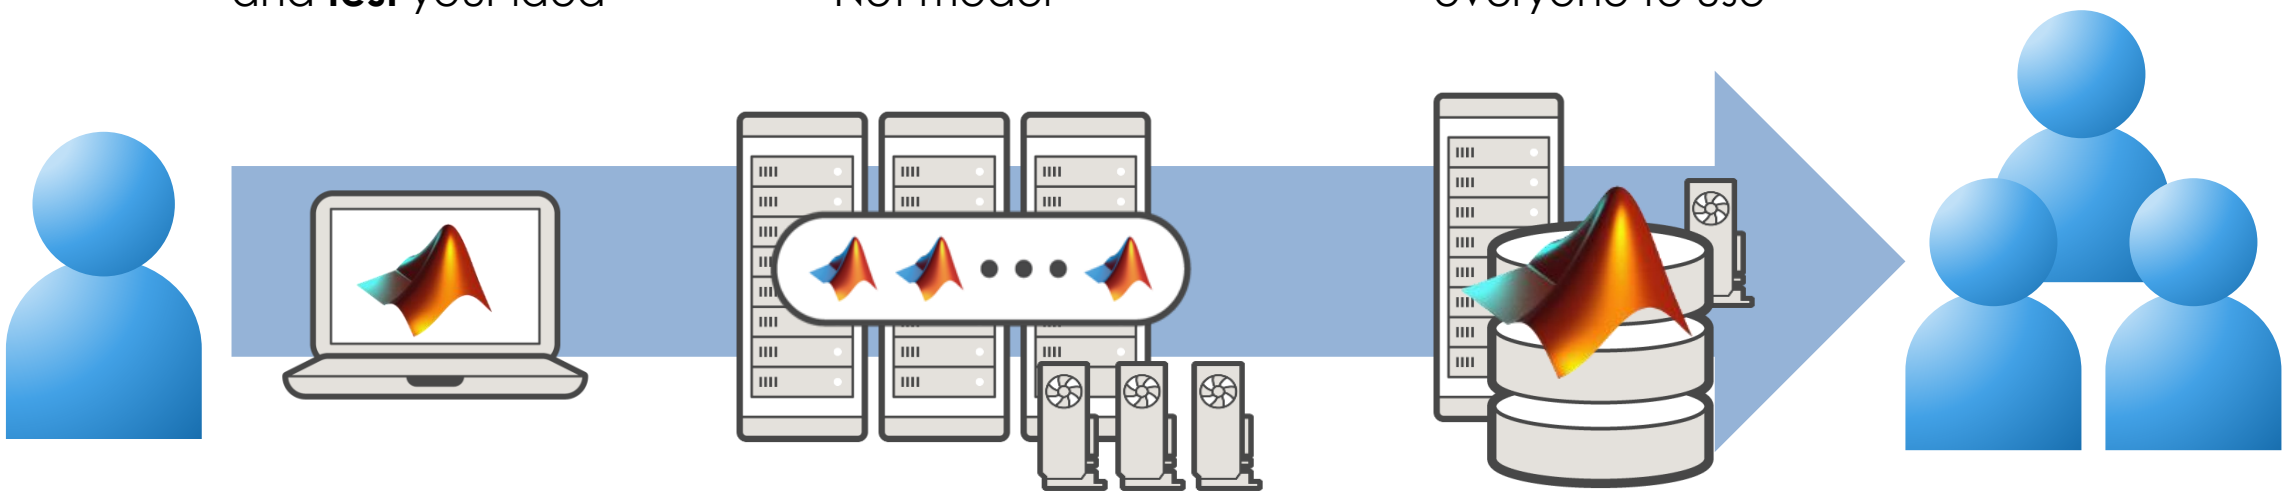

MATLAB Online, MATLAB Drive, MATLAB Mobile

MATLAB Distributed Computing Server

MATLAB Production Server

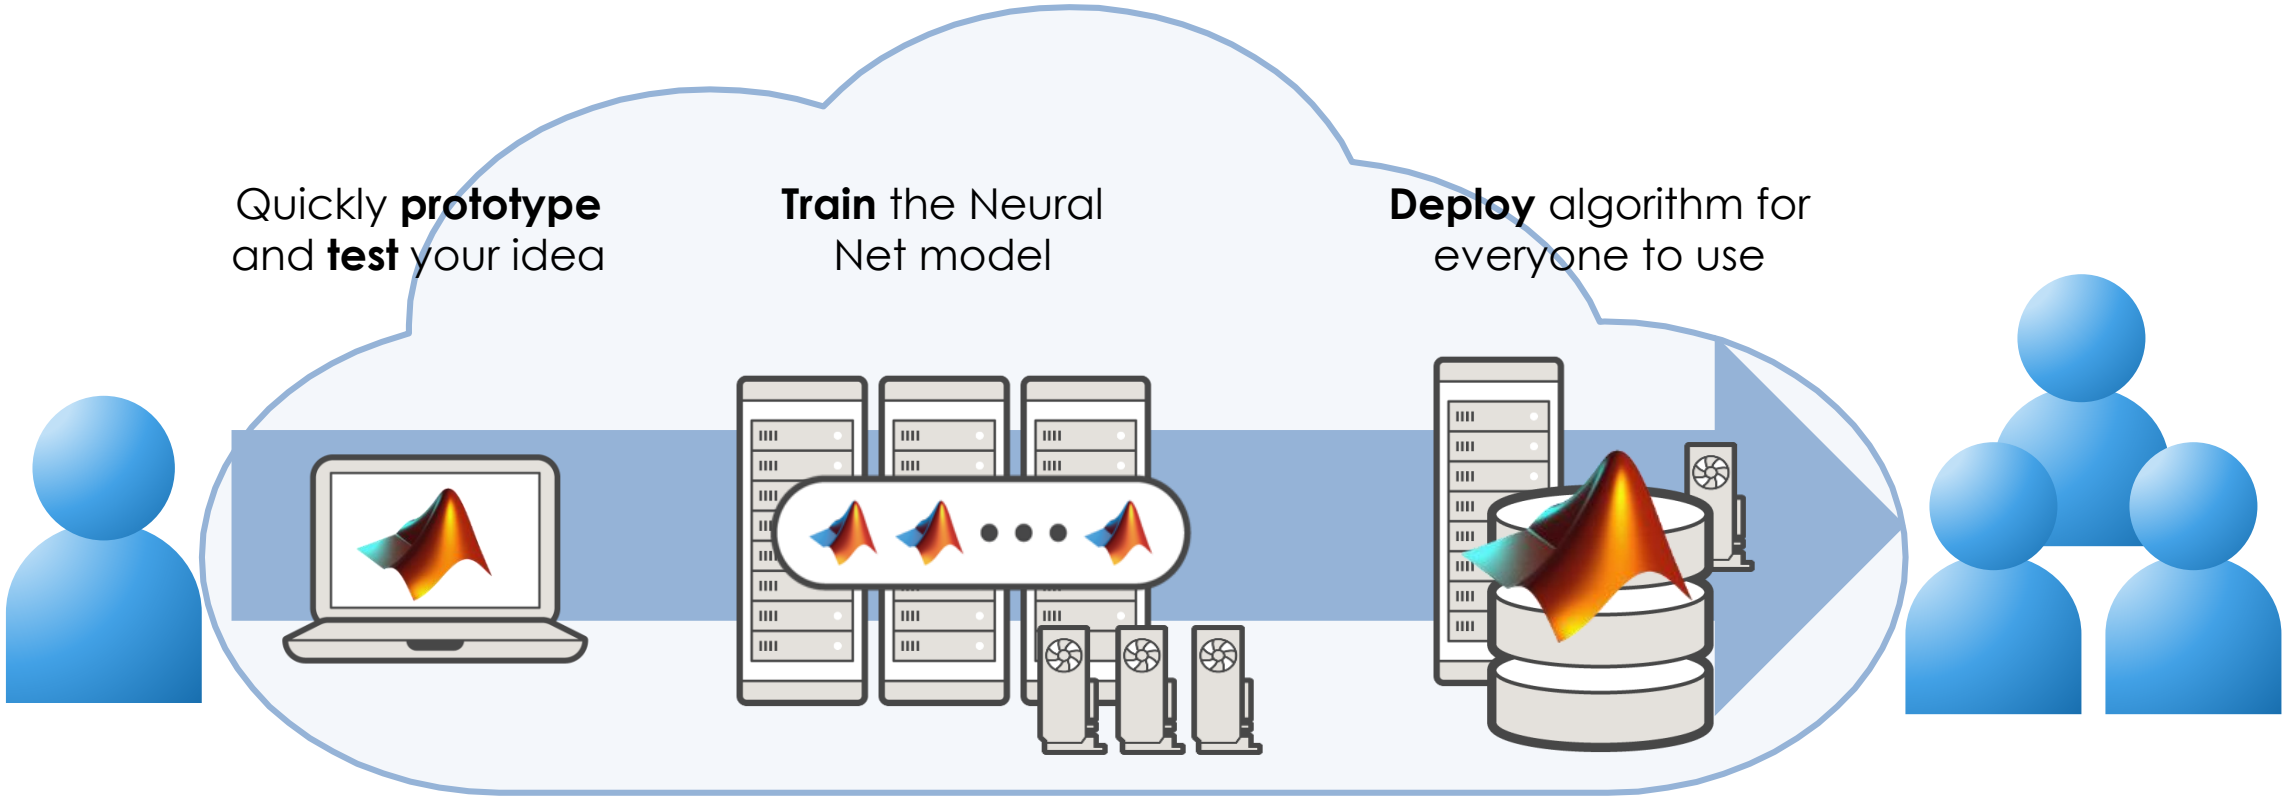

MATLAB Online, MATLAB Drive, MATLAB Mobile

MATLAB Distributed Computing Server

MATLAB Production Server

Prototyping in the Cloud with **MATLAB Online**

Intensive Computing in the Cloud with
MATLAB Parallel Server

Food categories

Category 1

Category 2

Category x

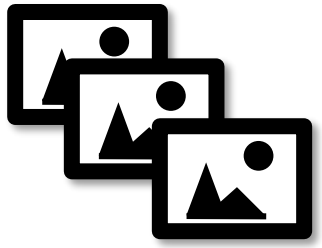

Pictures (1000)

100,000's pictures (5+ GB)

MATLAB Distributed Computing Server

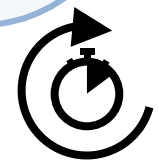

Hours to minutes to train network

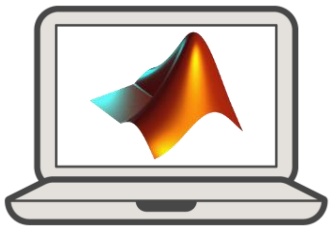

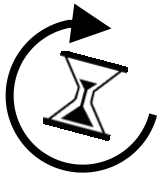

Days to train network

Deploying in the Cloud with **MATLAB Production Server**

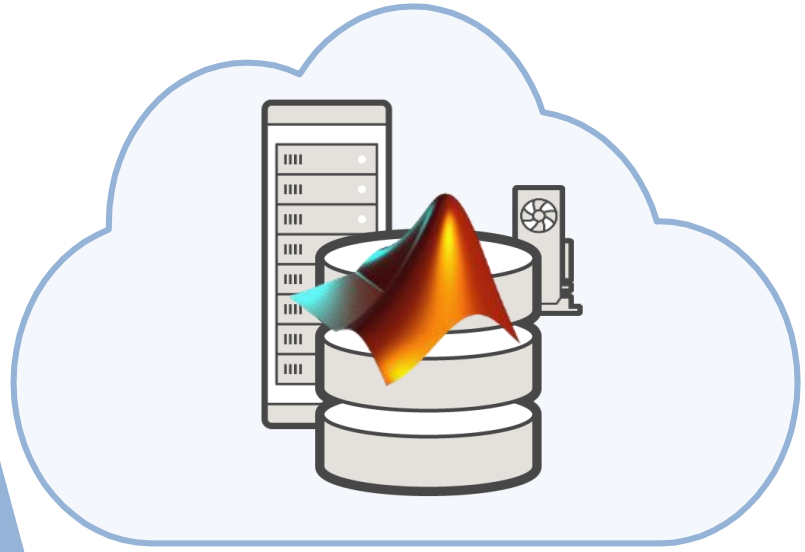

Upload image through web interface

Predicted label

MATLAB Production Server

Image preprocessing and deep learning model to classify food

Take picture of your food

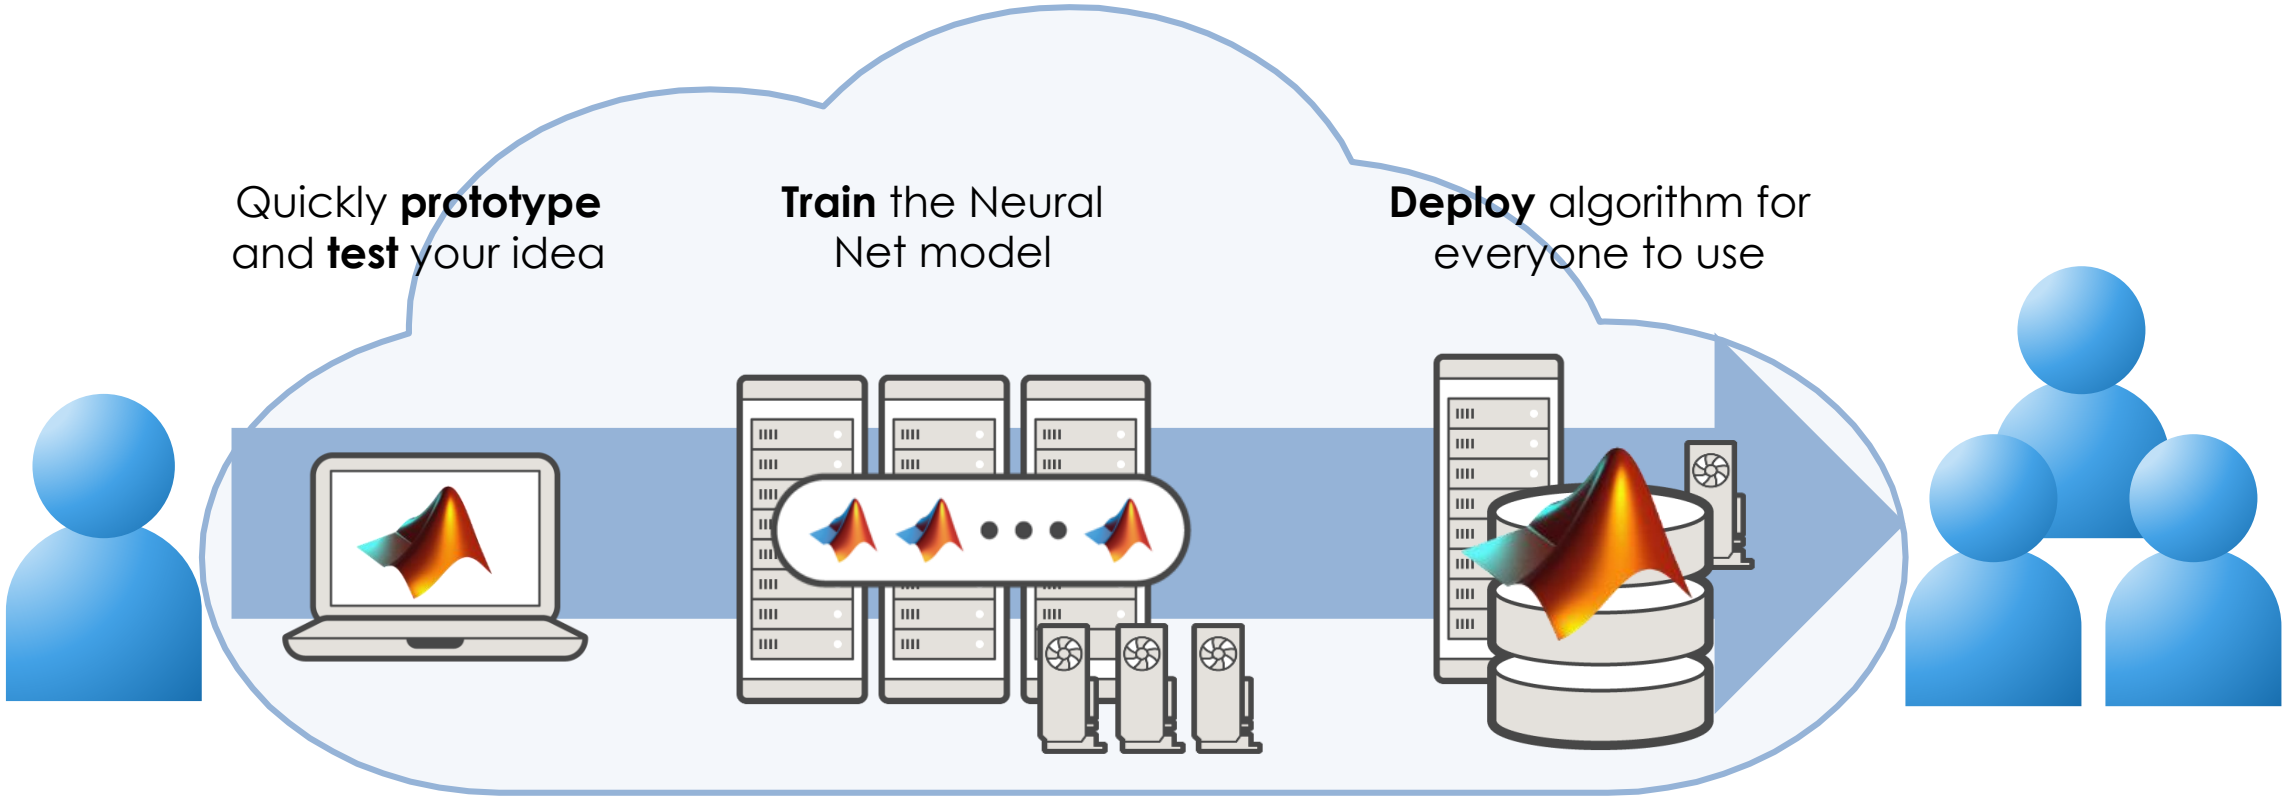
Quickly **prototype** and **test** your idea
Train the Neural Net model
Deploy algorithm for everyone to use
MATLAB Online, MATLAB Drive, MATLAB Mobile
MATLAB Distributed Compute Server
MATLAB Production Server

Deploying in the Cloud with Kafka

The Need for Large-Scale Streaming

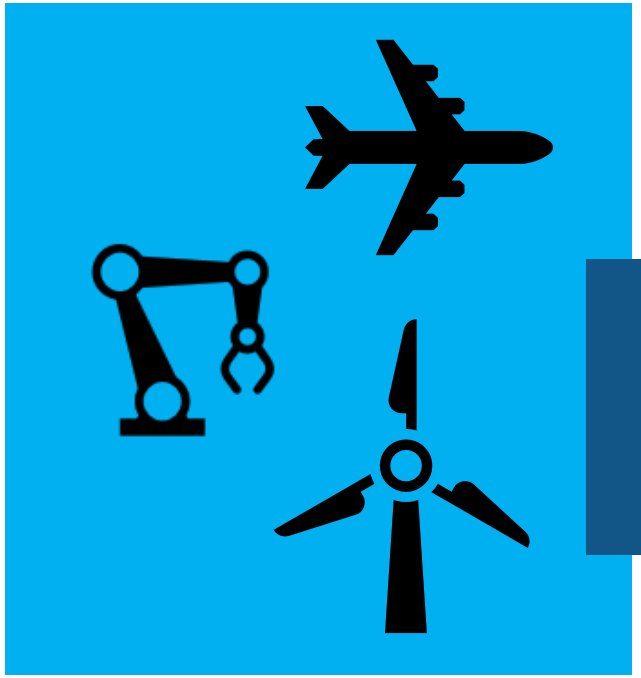

Jet engine: ~800TB per day
Turbine: ~ 2 TB per day

Predictive Maintenance

Increase Operational Efficiency

Reduce Unplanned Downtime

More applications require near real-time analytics

Medical Devices

Patient Safety

Better Treatment Outcomes

Connected Cars

Safety, Maintenance

Advanced Driving Features

Car: ~25 GB per hour

Example Problem: Develop a machine learning model to predict failures in industrial pumps

- We did this for the customer
- We wanted to go further:
  - Create a streaming application based on this real customer request
  - Develop application in a 3-4 week sprint
- We believe this represents a realistic customer situation

Our Project: Develop **and operationalize** a machine learning model to predict failures in industrial pumps

Process Engineer

Develops models in MATLAB and Simulink

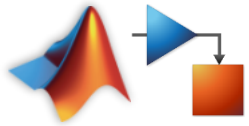

System Architect

Deploys and operationalizes model on Azure cloud

Azure

Operator

Makes operational decisions based on model output

kibana

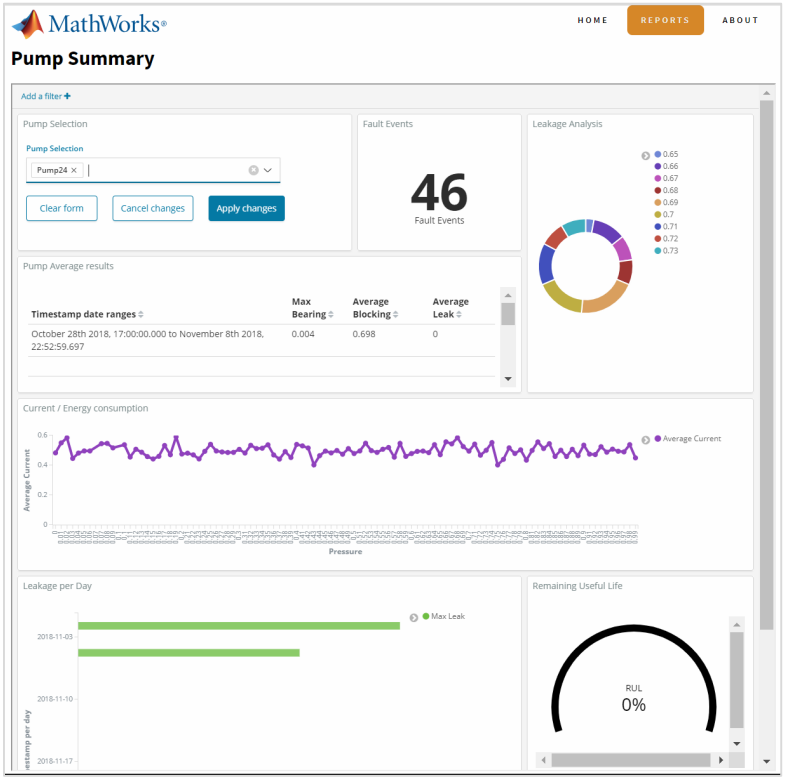

Current system requires Operator to manually monitor operational metrics for anomalies. Their expertise is required to detect and take preventative action


5" 12,000 PSI
5" 12,000 PSI
BAKER HUGHES

Project statement: Develop end-to-end predictive maintenance system and demo in one 3-4 week sprint

Plant Operator

1. Monitor ***flow***, ***pressure***, and ***current*** of each pump so I always know their ***operational state***
2. Need ***alert*** when fault parameters drift outside an acceptable range so I can take ***immediate corrective action***
3. Continuous estimate of each pump's ***remaining useful life (RUL)*** so I can ***schedule maintenance or replace*** the asset

Challenges of AI Deployment

Process Engineer

We don’t have a large set of failure data, and it’s too costly to generate real failures in our plant for this project

Solution: Use an accurate physics-based software model for the pump to develop synthetic training sets

Challenges of AI Deployment

System Architect

We don’t have a large IT/hardware budget, and we need to see results before committing to a particular platform or technology

Solution: Leverage cloud platform to quickly configure and provision the services needed to build the solution, while minimizing lock-in to a particular provider

Challenges of AI Deployment

Process Engineer

Need software for multidisciplinary problem across teams, plus integration w/ IT

Solution: Use MATLAB and integrate with OSS

Predictive Maintenance Architecture on Azure

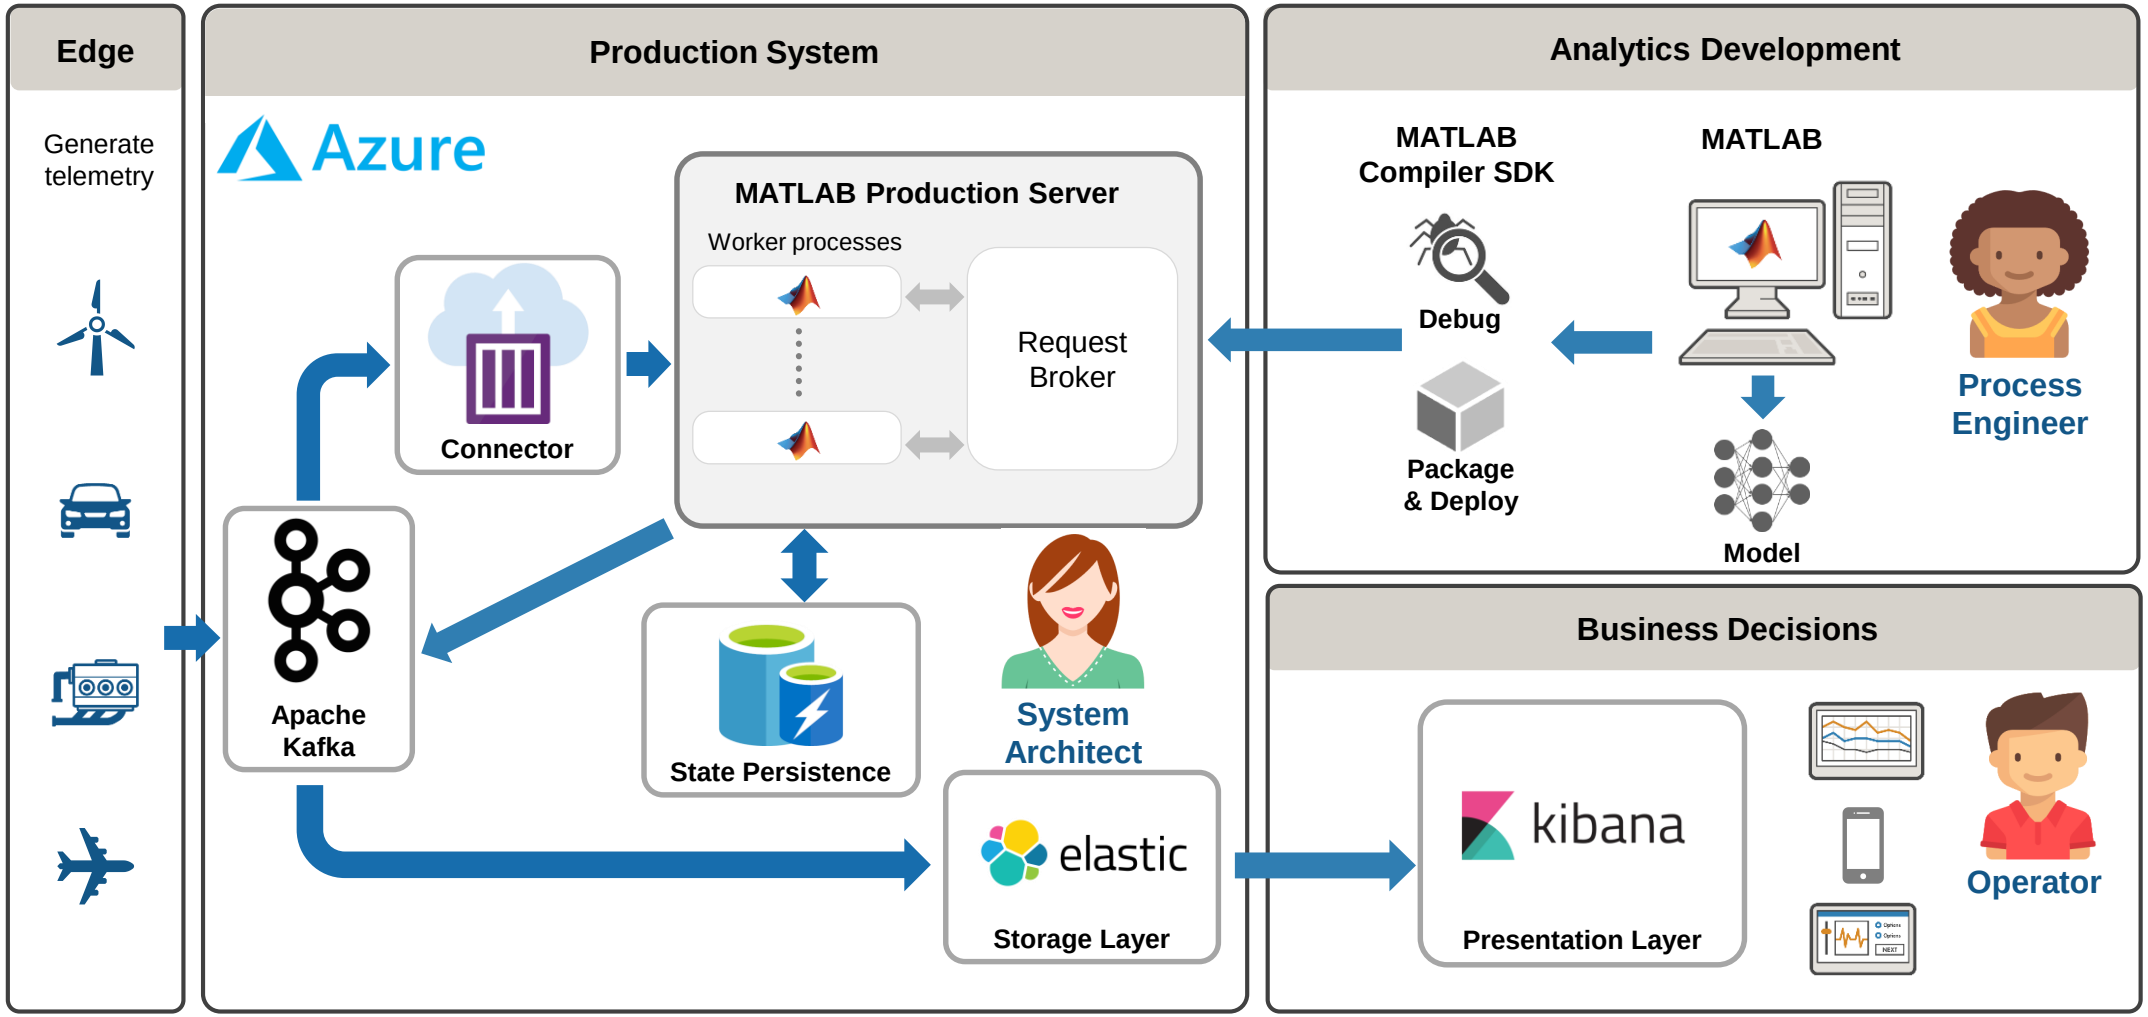

Modeling approach

Process Engineer

1. **Access and Explore Data**
2. **Preprocess Data**
3. **Develop Predictive Models**
4. **Integrate with Production Systems**
5. **Visualize Results**

Files

Databases

Sensors

Working with Messy Data

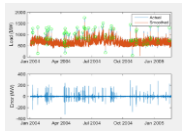

Data Reduction/ Transformation

Feature Extraction

Model Creation e.g. Machine Learning

Parameter Optimization

Model Validation

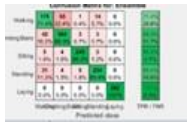

Desktop Apps

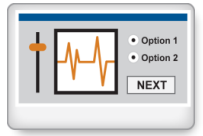

Enterprise Scale Systems

kafka

AWS Kinesis

OSIsoft.

Embedded Devices and Hardware

3rd party dashboards

Spotfire® TIBCO Software

+ a b l e a u

Qlik

Power BI

Web apps

Process Engineer

Review model requirements

Requirements From Operator

- Continuous predictions of type of fault
  - “Blocking”
  - “Leaking”
  - “Bearing”
  - Combination of above
- Continuous predictions of Remaining Useful Life [RUL]

Requirements From System Architect

- Define window for streaming
- Define format of results, intermediate values
- Test code
- Scale code


5" 12,000 PSI
5" 12,000 PSI
BAKER HUGHES

Physics of Triplex Pump

Process Engineer

- Crankshaft drives three plungers
  - Each 120 degrees out of phase
  - One chamber always discharging
  - Three types of **failures**

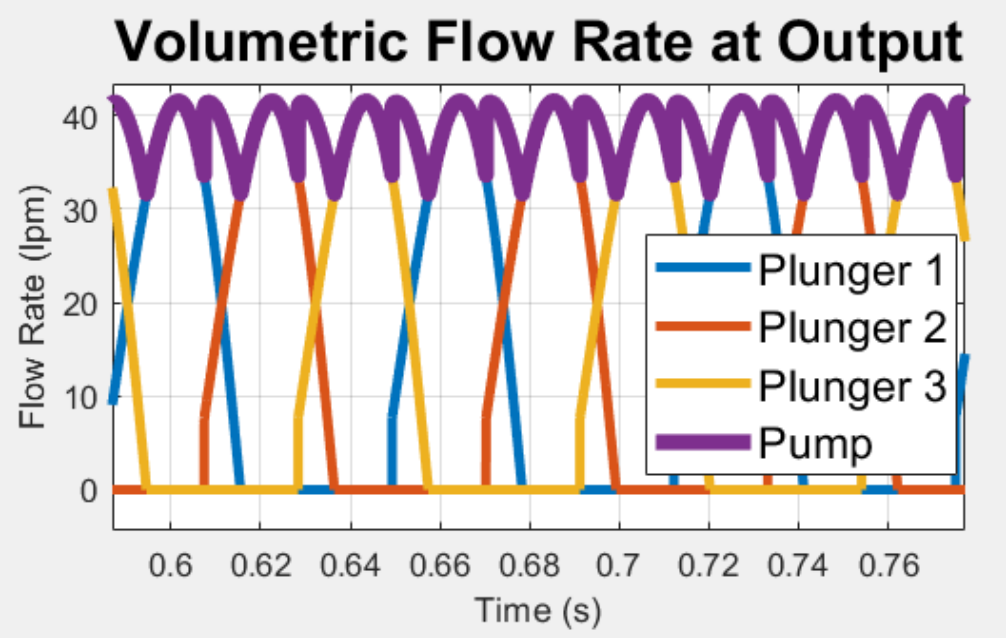

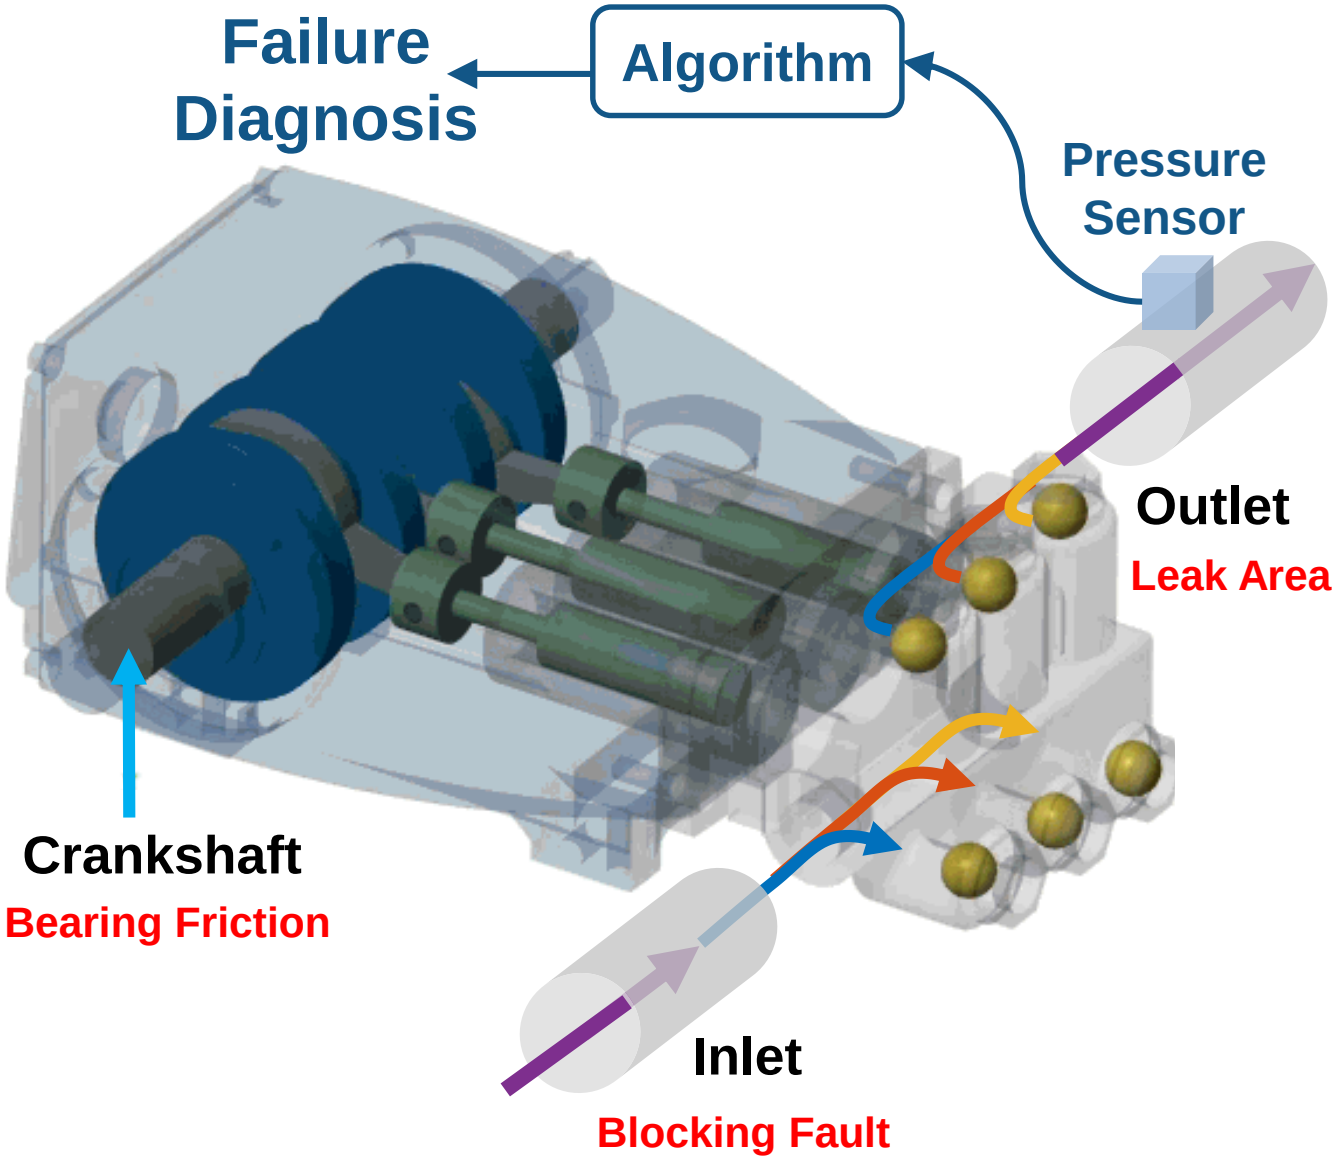

1 **Access and Explore Data**

Use sensor data from pump to identify levels of failure

Process Engineer

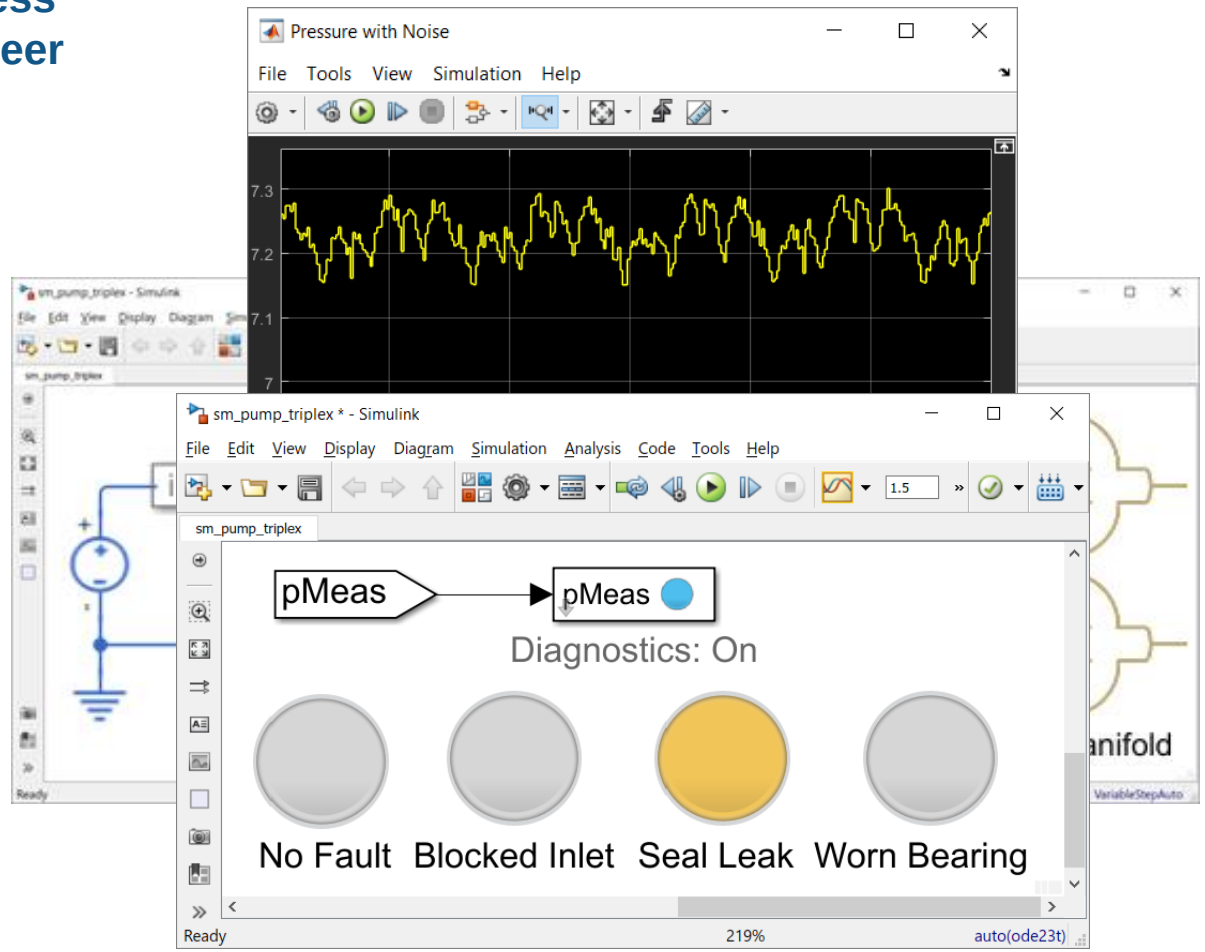

Simulate faults

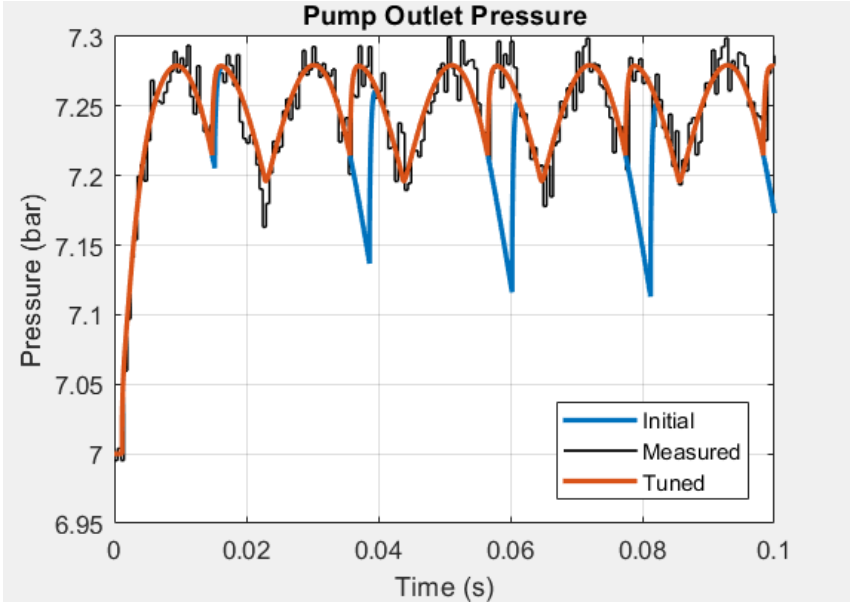

Pump sensor data

1 **Access and Explore Data**

Build digital twin and generate sensor data

Process Engineer

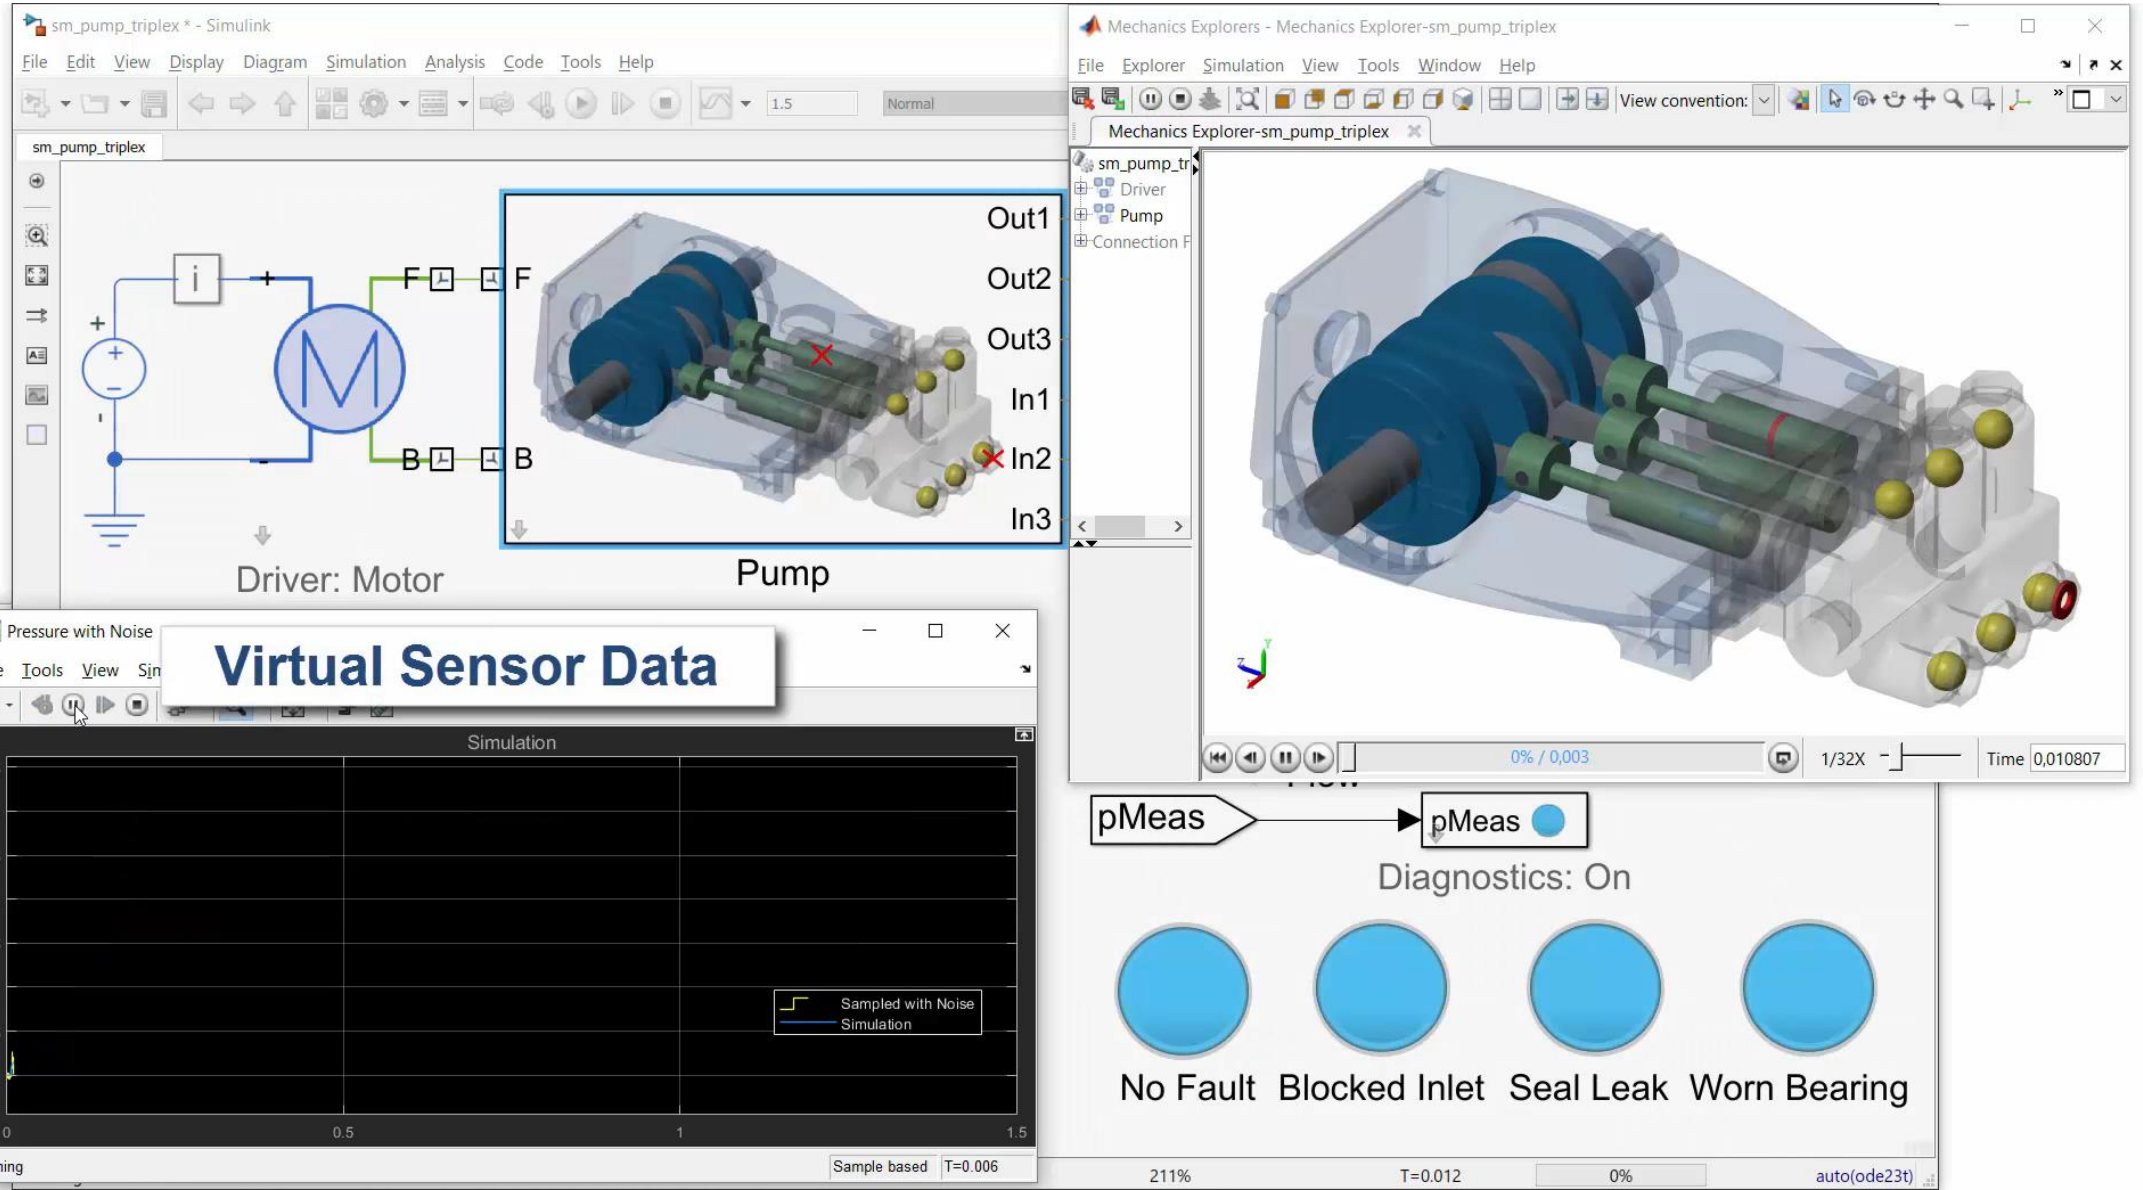

Process Engineer

1 Access and Explore Data

Simulate data with many failure conditions

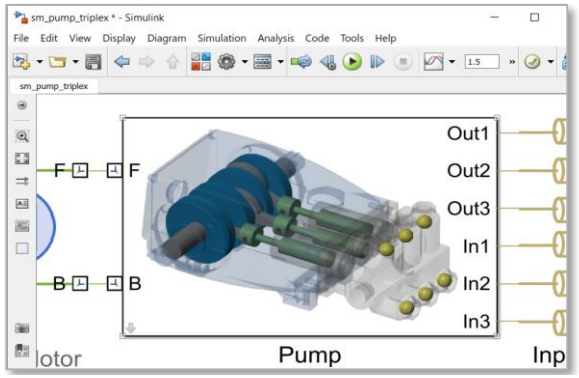

Leak Area = [1e-9 0.036]

Bearing Friction = [0 6e-4]

Blocking Fault = [0.5 0.8]

Simulate data with many failure conditions

Process Engineer

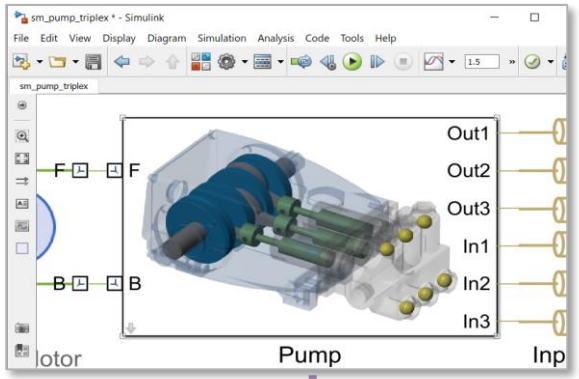

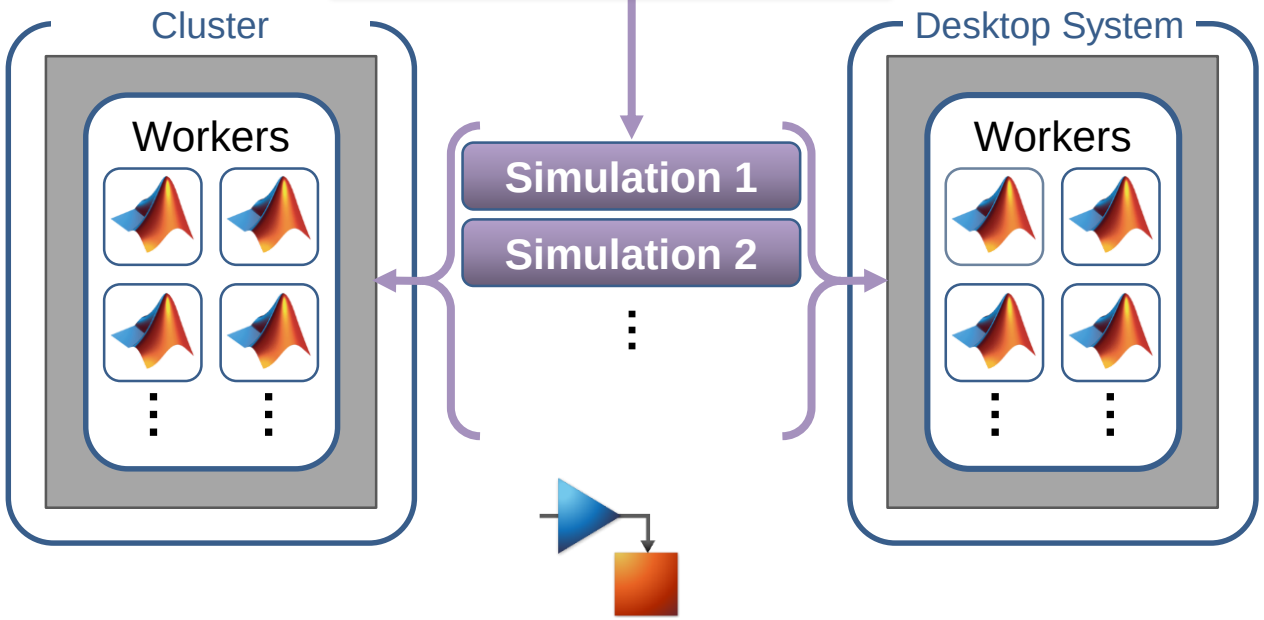

Run parallel simulations

Access Data

```
ens = simulationEnsembleDatastore(location)
```

```
ens =
  simulationEnsembleDatastore with properties:

          DataVariables: [25×1 string]
    IndependentVariables: [0×0 string]
      ConditionVariables: [0×0 string]
       SelectedVariables: [25×1 string]
                ReadSize: 1
              NumMembers: 702
          LastMemberRead: [0×0 string]
                   Files: [702×1 string]
```

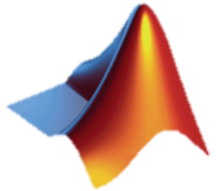

Store data on HDFS

Process Engineer

2 Preprocess Data

Represent signal information

Signal processing

```
[Spectrum,Frequencies] = pspectrum(data.Flow);
[pLow,pHigh] = bounds(Spectrum);
fPeak = Frequencies(Spectrum==pHigh);
qPeak2Peak = peak2peak(data.Flow);
qCrest = peak2rms(data.Flow);
qRMS = rms(data.Flow);
qMAD = mad(data.Flow);
```

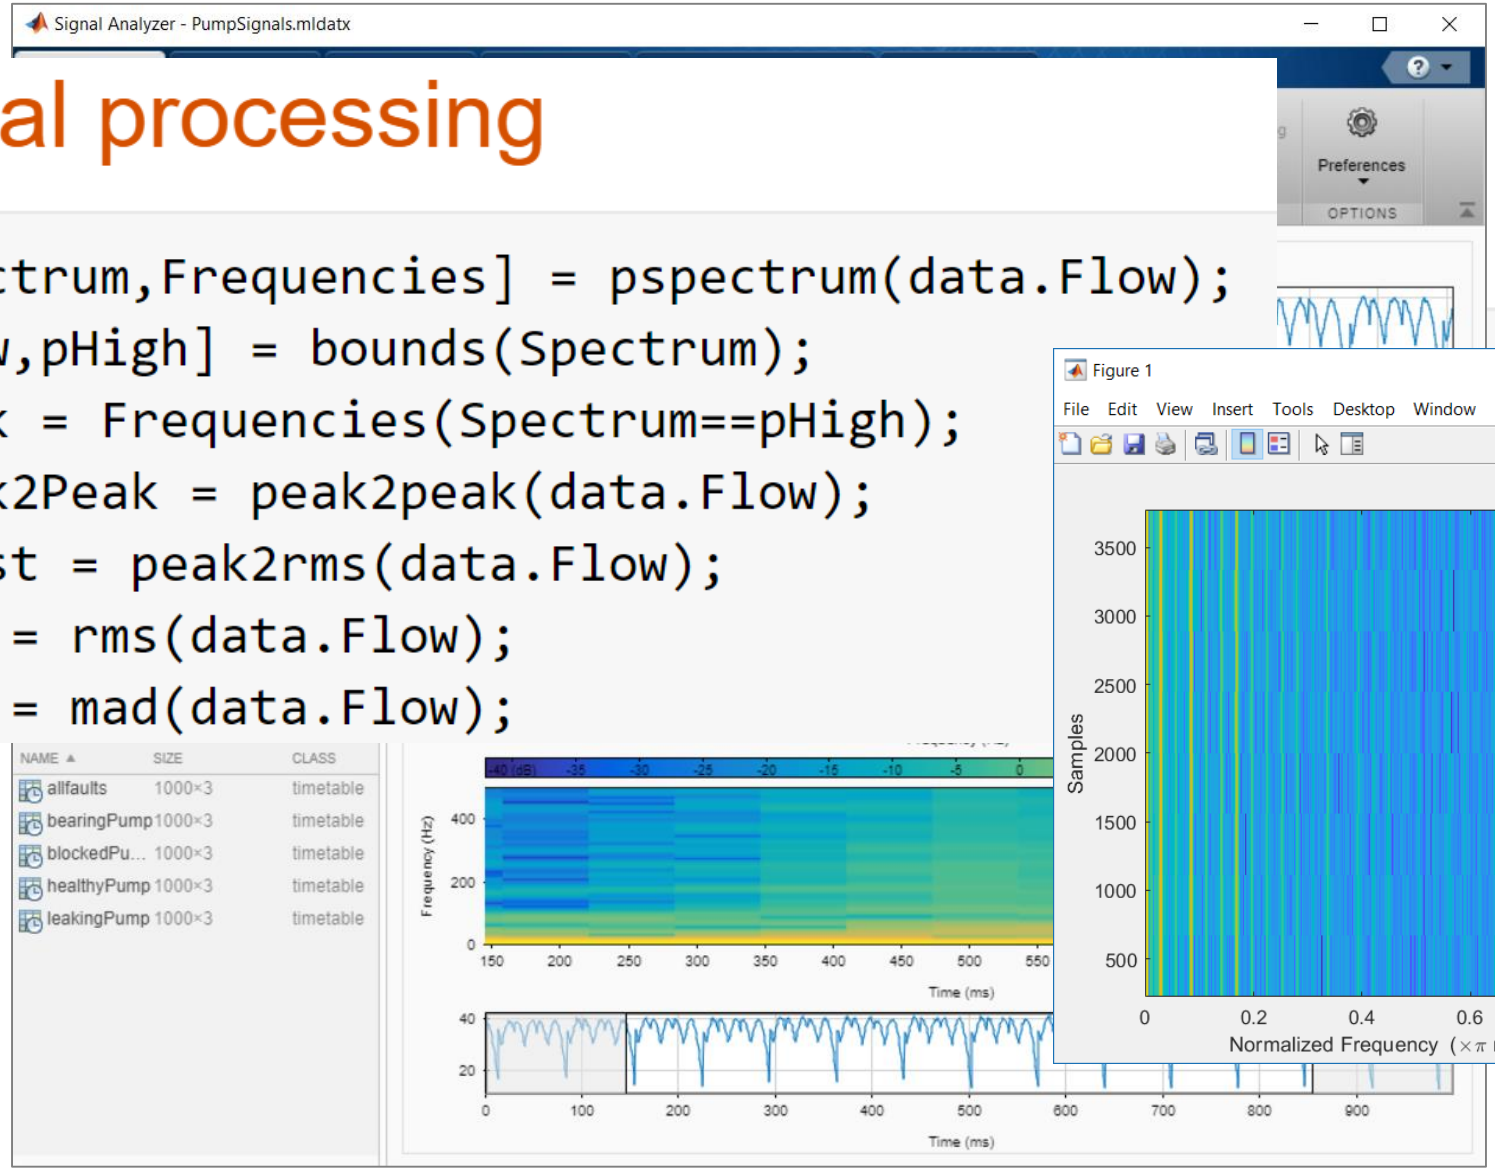

3 Develop Predictive Models

Develop Predictive Models in MATLAB

Process Engineer

	Time	1 LeakFault	2 BlockingFault	3 BearingFault	4 FaultType
1	0 sec	2.8472	-0.1477	1.8000	All
2	0.001 sec	-0.1498	-0.4207	1.3103	Bearing & Blocking
3	0.002 sec	0.6511	1.6521	-0.5557	Leak
4	0.003 sec	0.1469	-0.2775	1.0074	All
5	0.004 sec	-0.6480	0.7065	-0.8878	Blocking
6	0.005 sec	-0.8165	-0.5434	-0.3079	Blocking
7	0.006 sec	-1.0061	1.2083	0.0661	Bearing
8	0.007 sec	1.0125	-1.9098	-0.7027	Leak & Blocking

Label Faults

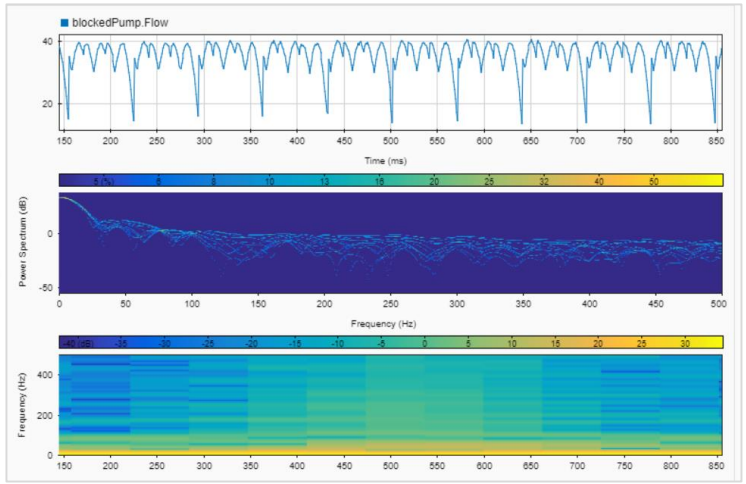

Represent Signals

Train Model

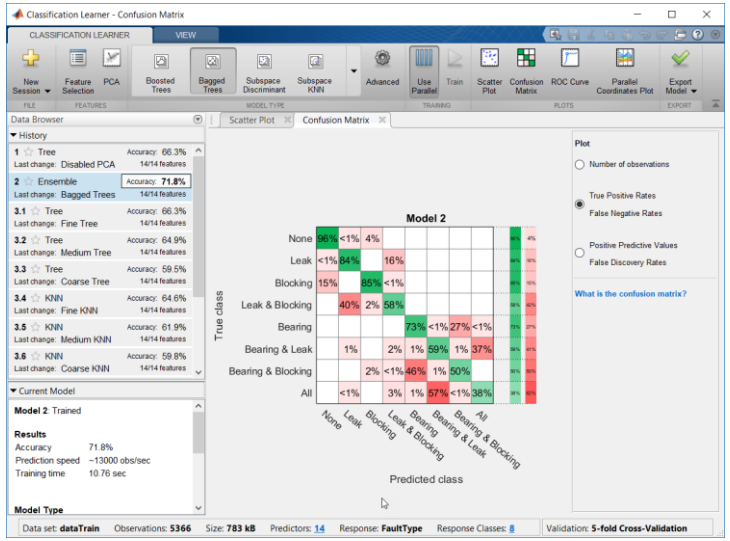

Validate Model

Scale

```
tt = tall(ds);
tt = preprocessData(tt);
model = TreeBagger(50,tt,'Event');
```

```
Evaluating tall expression using the Spark Cluster:
- Pass 1 of 2: Completed in 11 sec
- Pass 2 of 2: Completed in 2.3333 min
Evaluation completed in 2.6167 min
```

3 Develop Predictive Models

Process Engineer

Develop Predictive Models in MATLAB

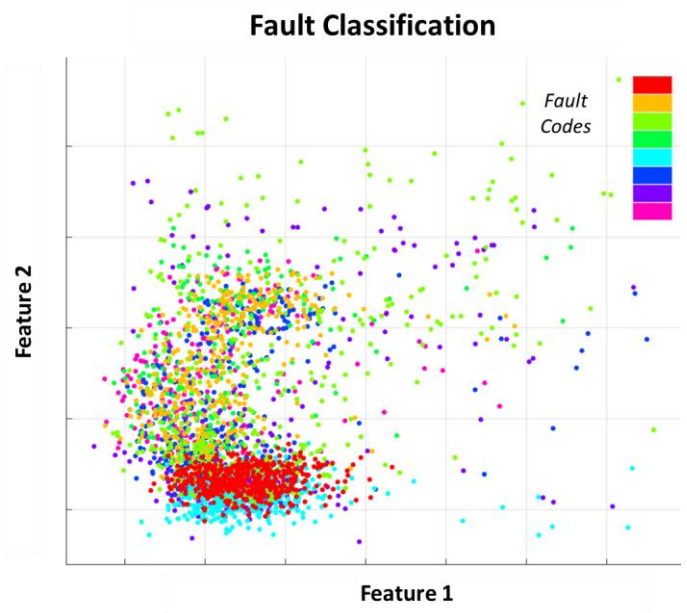

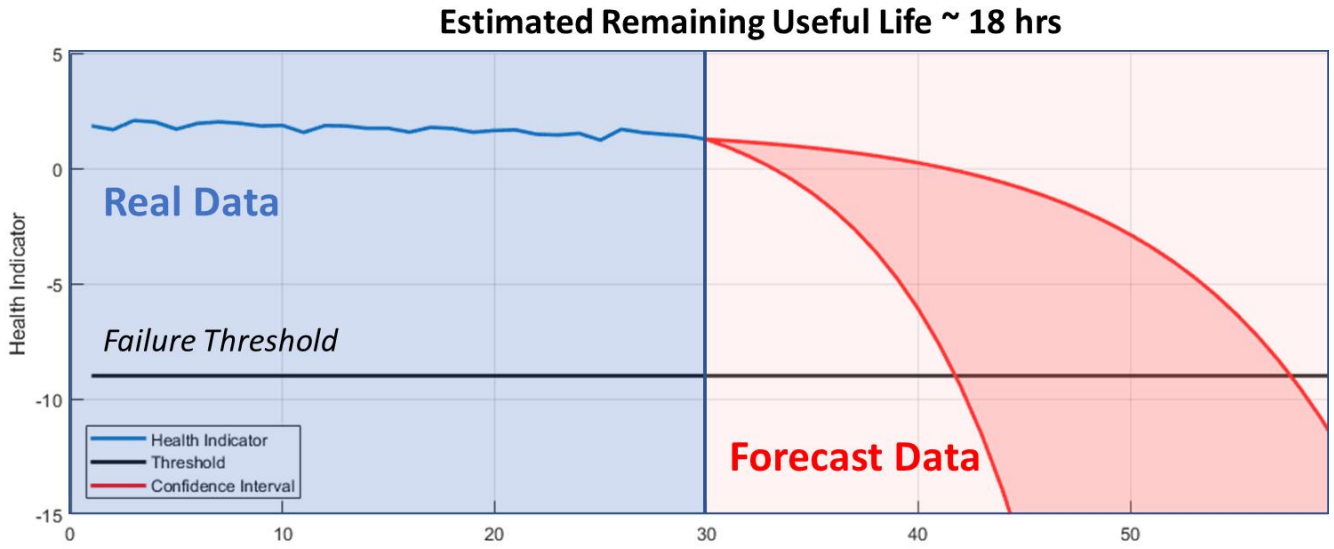

Type of Fault (Classification)

Plant Operator

Remaining Useful Life (Regression)

Process Engineer

3 Develop Predictive Models

Develop Machine Learning Models

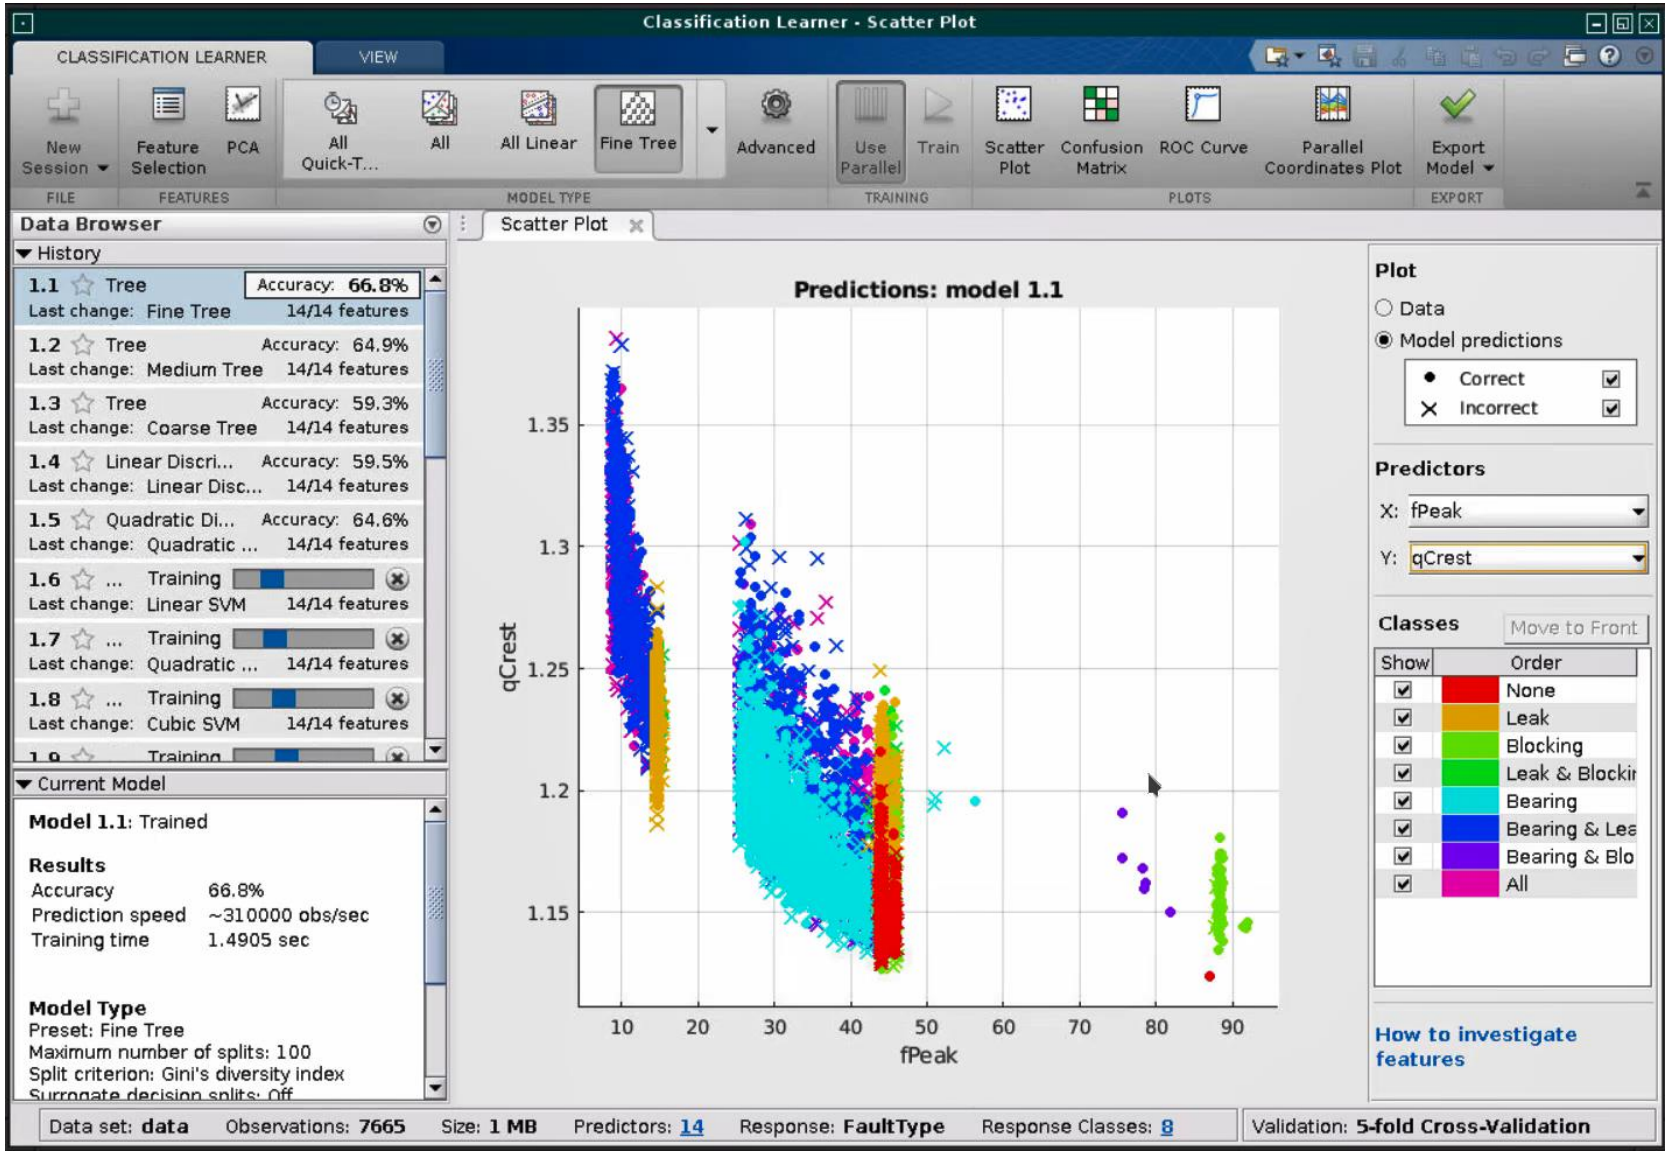

3 Develop Predictive Models

Process Engineer

Estimate Remaining Useful Life

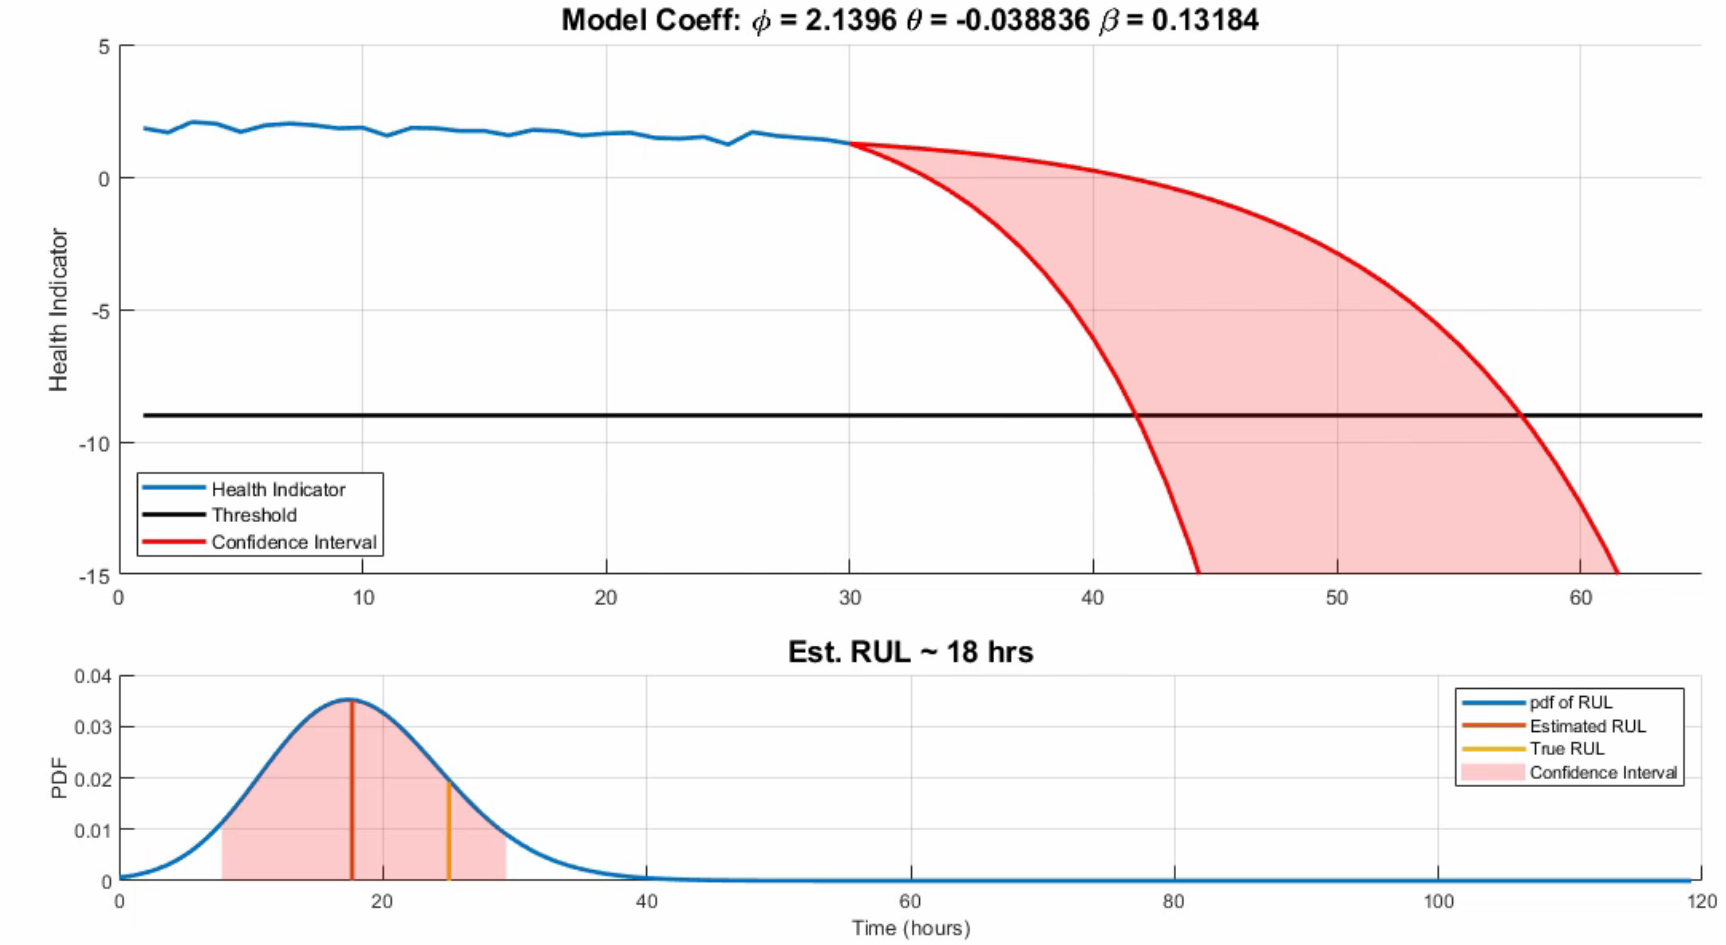

$$S(t) = \phi + \theta(t)\, e^{(\beta(t)t + \epsilon(t) - \frac{\sigma}{2})}$$

Process Engineer

4 Integrate with Production Systems

Develop a Stream Processing Function

- ***Batch Processing:*** Build and test model on simulated data

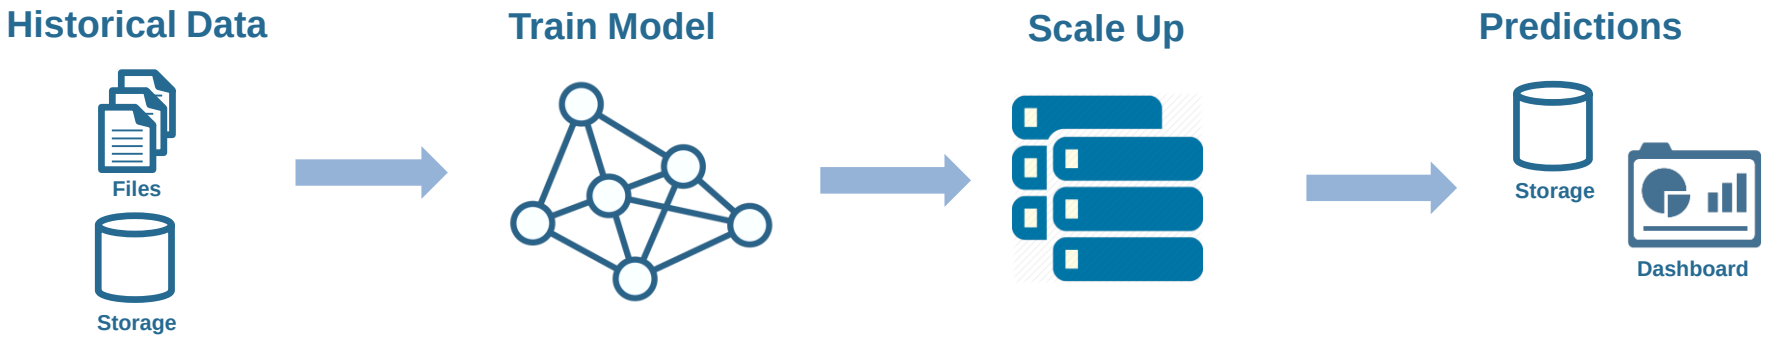

- ***Stream Processing:*** Apply model to sensor data in near real-time

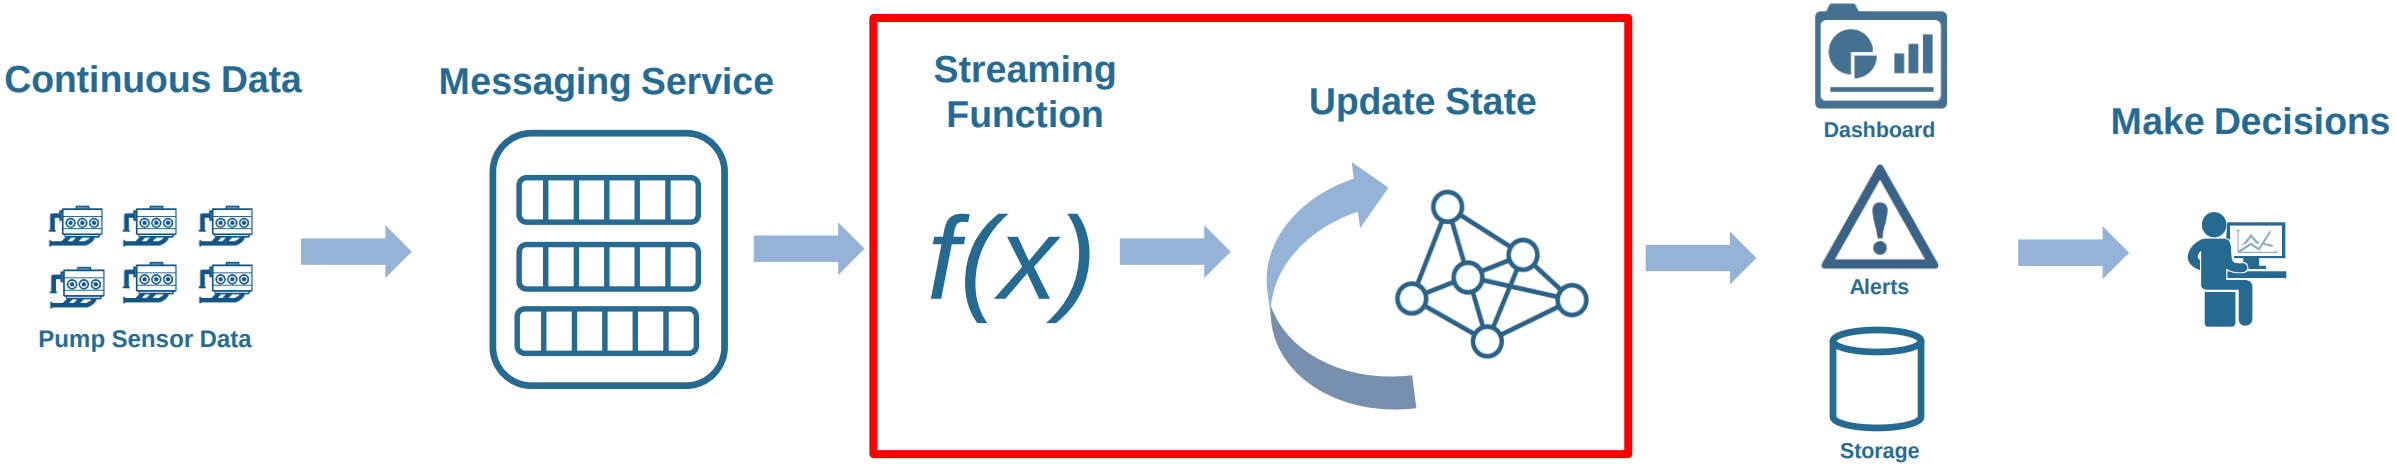

Process Engineer

4 Integrate with Production Systems

Develop a Stream Processing Function

Streaming Function

```
function new_state = streamingFunction(data,old_state)
```

Preprocess signals

```
[data,features] = preprocessData(data);
```

Predict faults

```
[Leak,Blocking,Bearing] = predictFaultValues(features);
FaultType = predictFault(features);
[RUL,Model] = predictUpdateRUL(data.Timestamp,data.Flow,500);
```

Update state

```
new_state = updateState(data,old_state);
```

Write results

```
writeResults(Leak,Blocking,Bearing,FaultType,RUL,Model)
end
```

Process each window of data as it arrives

Previous state

Current window of data to be processed

Process Engineer

4 Integrate with Production Systems

Test Stream Processing Function

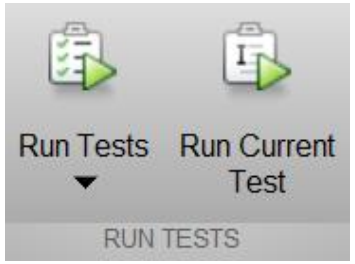


```
results = runtests('predictFaults_tests')
```

```
Running predictFaults_tests
....
Done predictFaults_tests
__________
results =
  1×4 TestResult array with properties:

    Name
    Passed
    Failed
    Incomplete
    Duration
    Details
Totals:
   4 Passed, 0 Failed, 0 Incomplete.
   0.01614 seconds testing time.
```

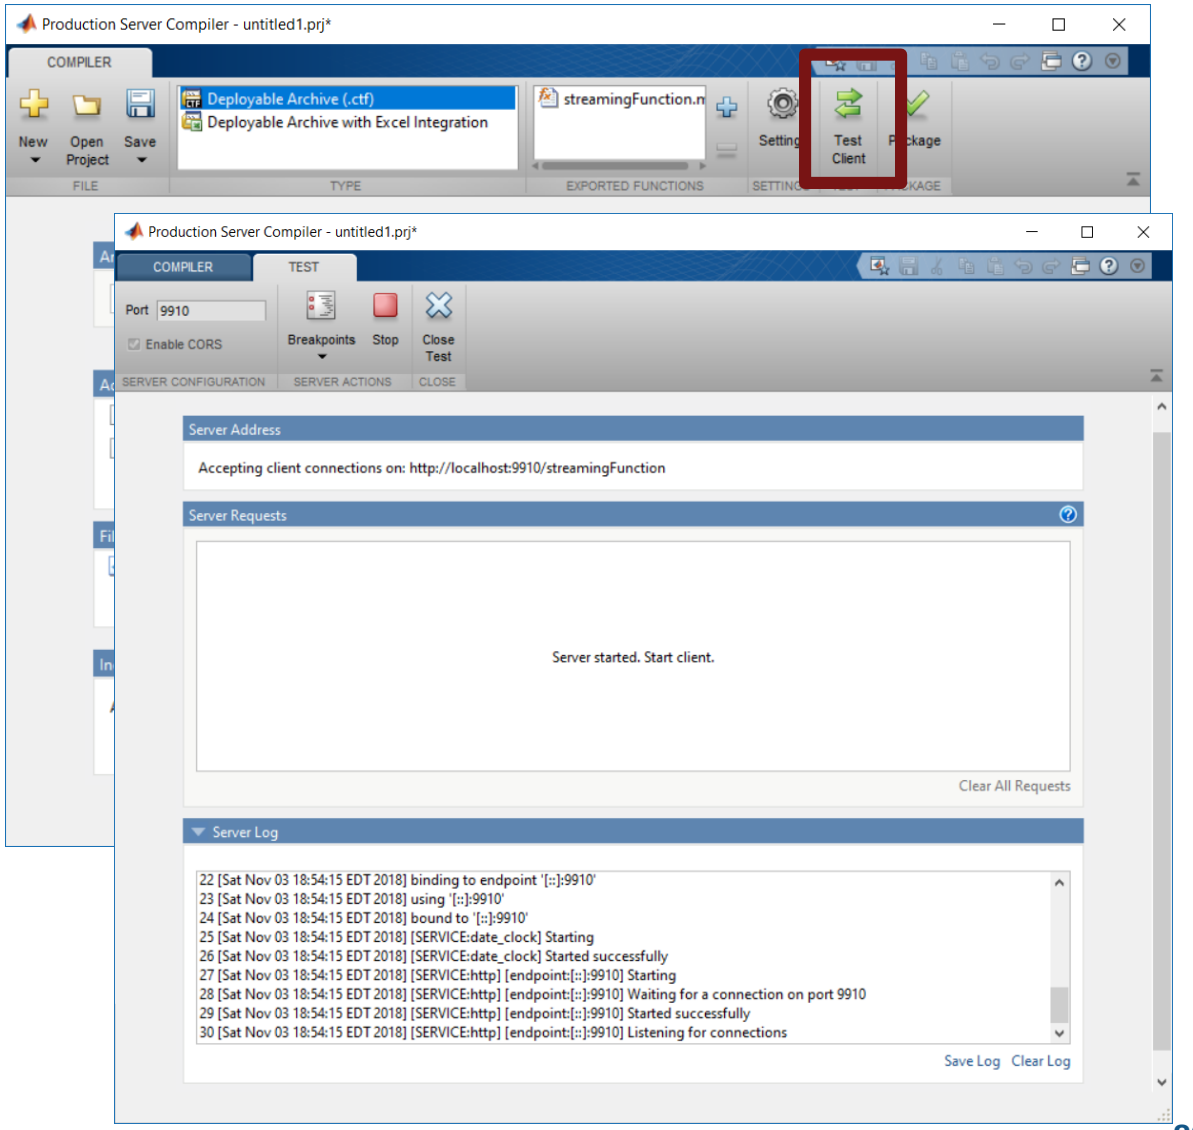

Process Engineer

4 Integrate with Production Systems

Share with the team

Review results with Operator

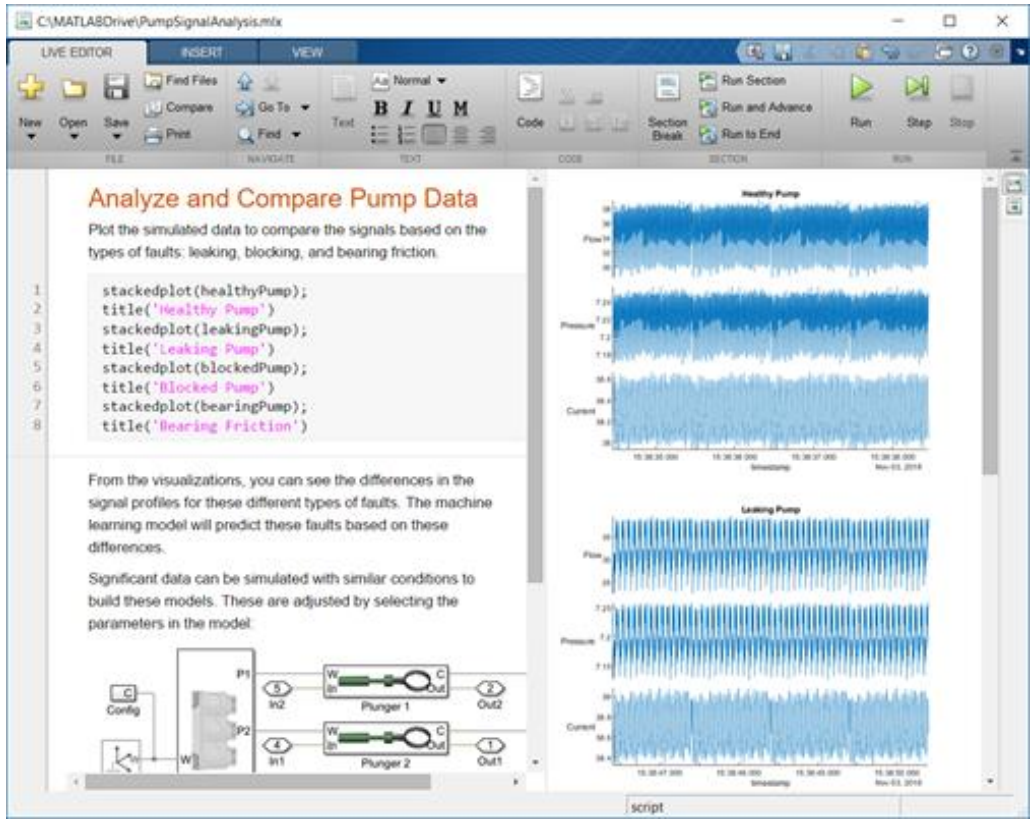

.pdf, html, LaTeX

Share code with System Architect

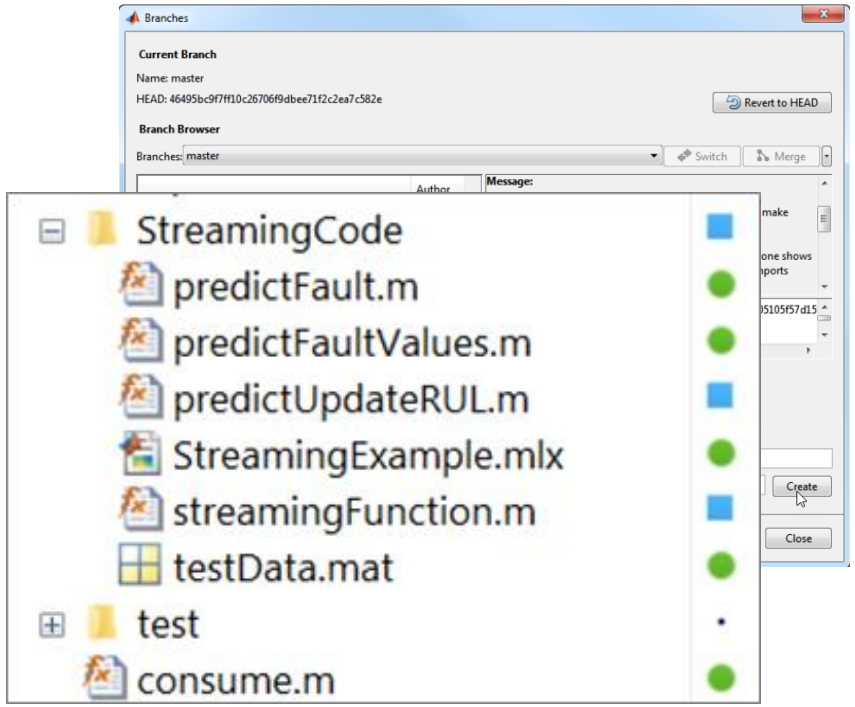

Source Control

Process Engineer

4 Integrate with Production Systems

Package Stream Processing Function

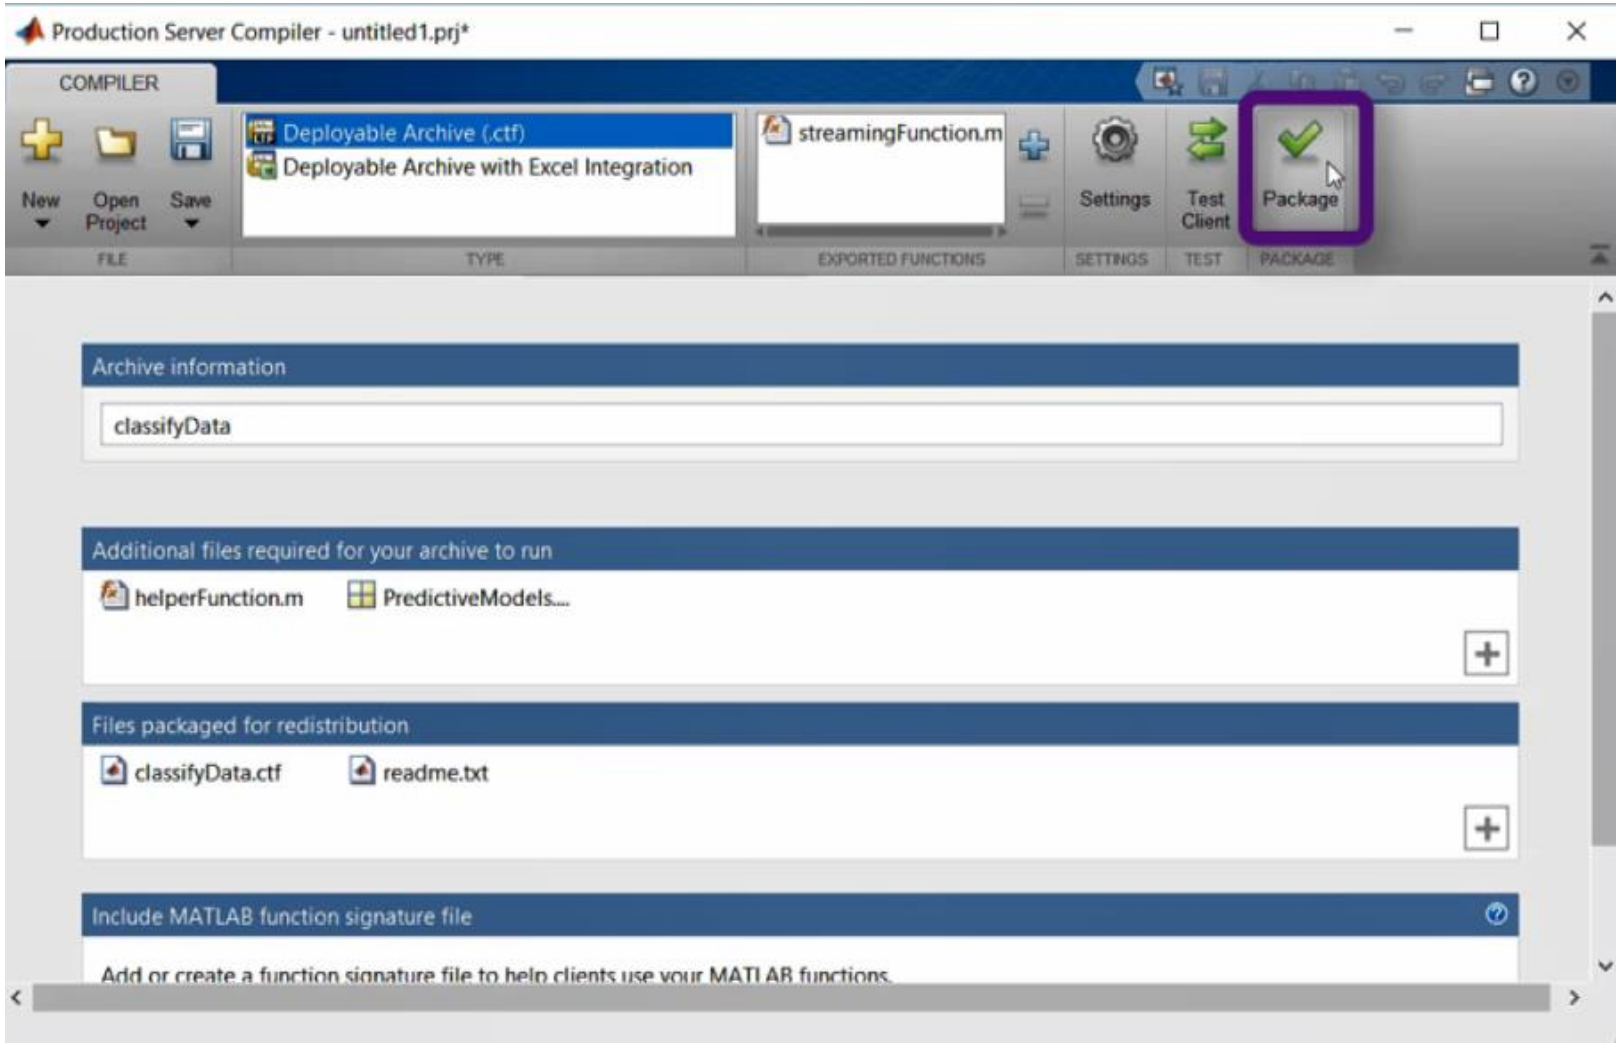

System Architect

4 Integrate with Production Systems

Review System Requirements

- Requirements from the Process Engineer
  - Every millisecond, each pump generates a time-stamped record of flow, pressure, and current
  - Model expects 1 sec. window of data per pump
  - Initially, 1's – 10's of devices, but quickly scale to 100's

Process Engineer

- Requirements from the Operator
  - Alerts when parameters drift outside the expected ranges
  - Continuous estimating of RUL for each pump

Operator

System Architect

4 Integrate with Production Systems

Integrate Analytics with Production Systems

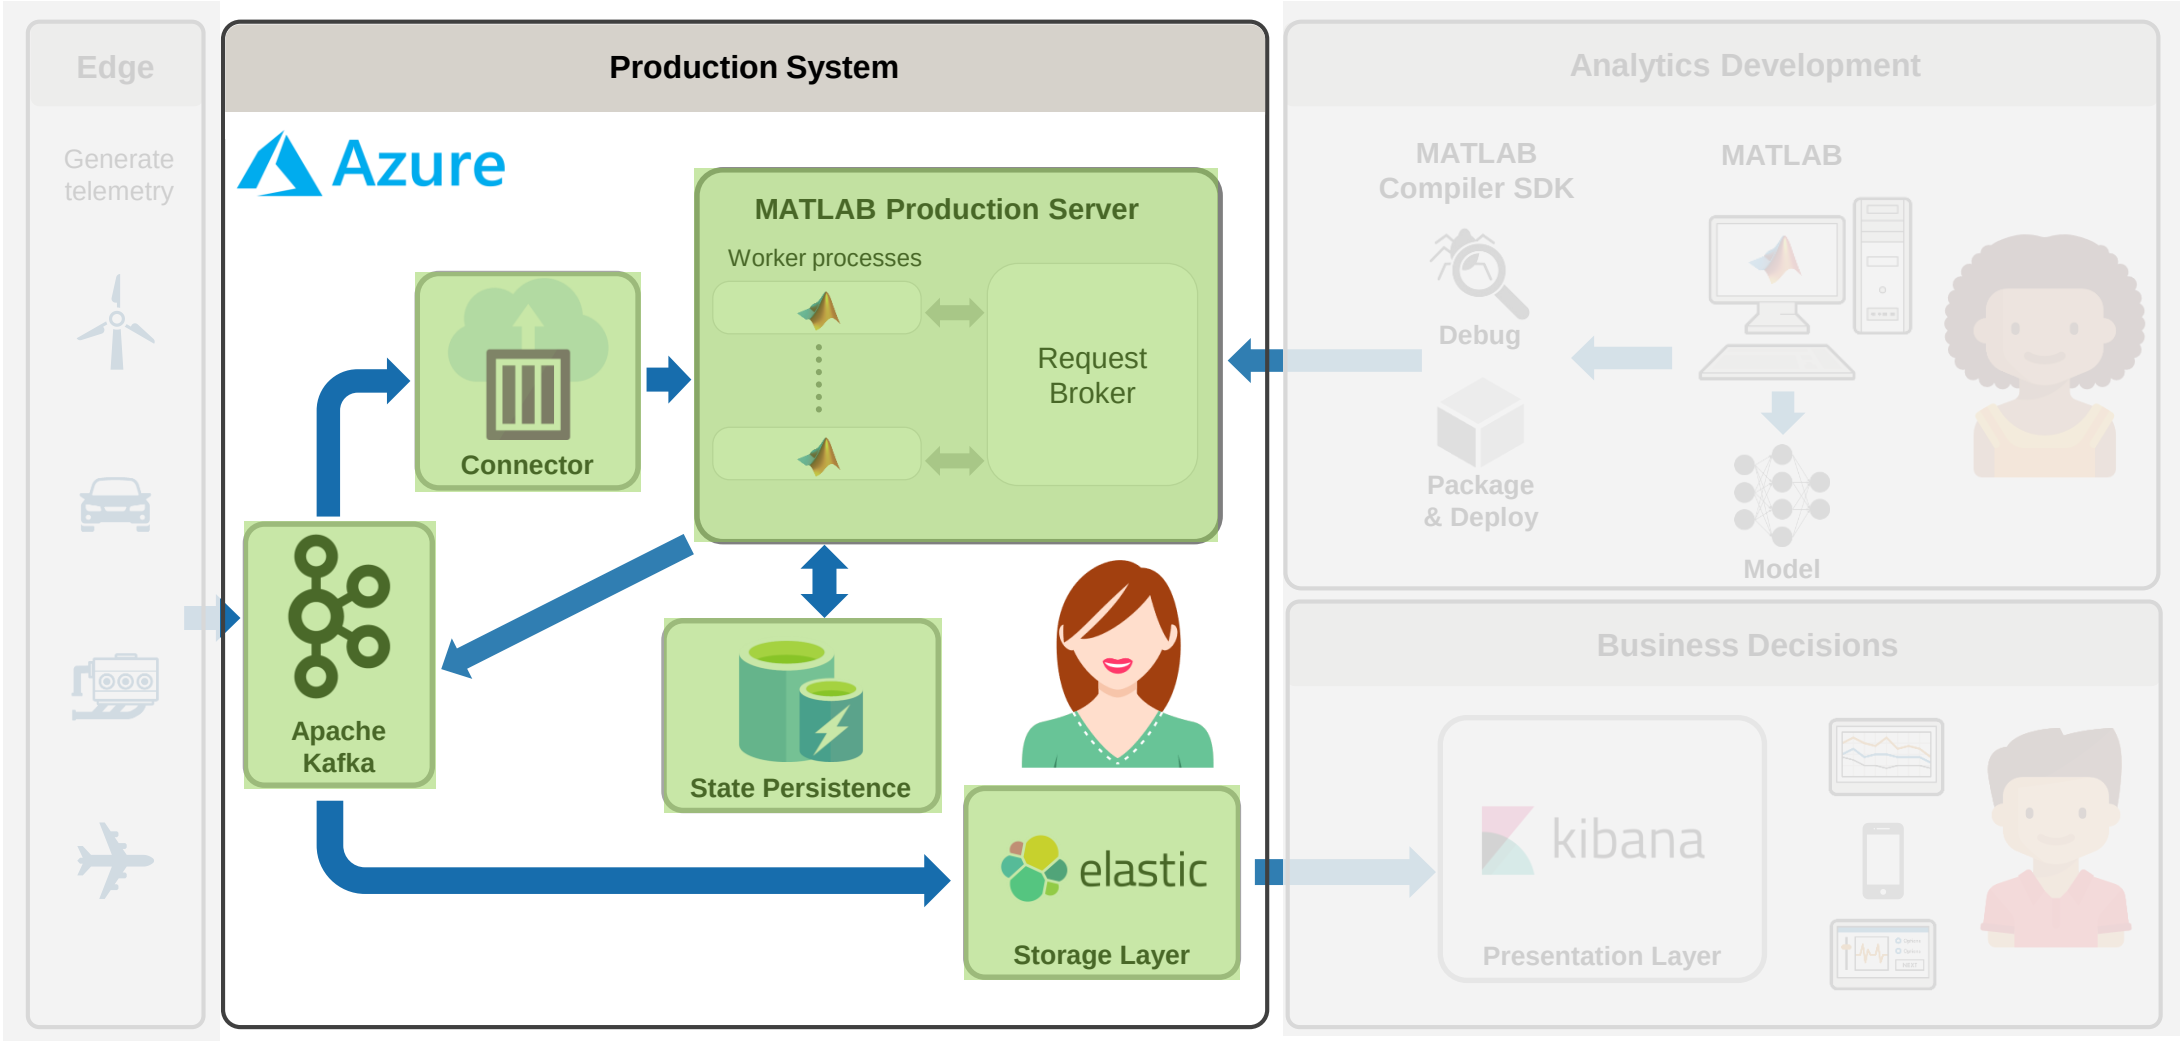

System Architect

4 Integrate with Production Systems

MATLAB Production Server on Azure

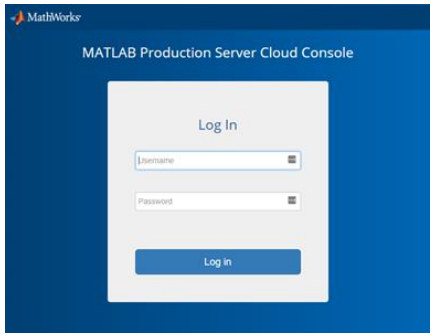

https management endpoint

Production System

Azure

Virtual Network

Management Server

MATLAB Production Server(s) scaling group

Application Gateway Load Balancer

Enterprise Applications

Connectors for Streaming/Event Data

State Persistence

Connectors for Storage & Databases

System Architect

4 Integrate with Production Systems

Connecting MATLAB Production Server to Kafka

- Connector feeds single Kafka topic to a MATLAB function
- Publisher library for MATLAB for writing to a results stream
- Connector Features:
  - Deploy as a micro-service with Docker
  - Drive everything through config
  - Group data into time windows and pass to MATLAB as a timetable
  - Use Kafka's check-pointing (i.e. at-least-once)

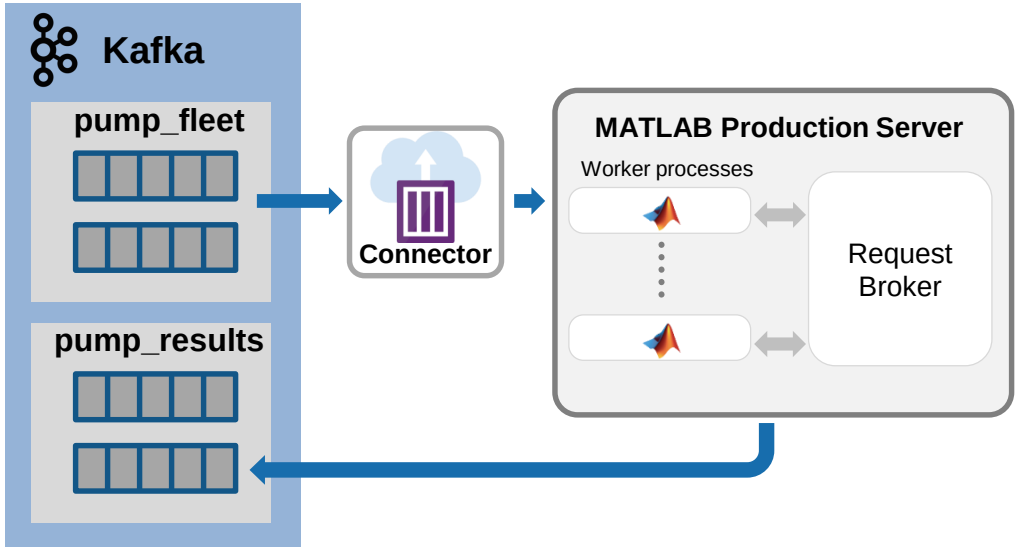

System Architect

4 Integrate with Production Systems

Messaging adapter for Production Server

- Bridges streaming data and Production Server Async Java Client
- Batches incoming messages and sends them via HTTP request/response
  - Time windows, event time processing, and out-of-order data
- Uses Asynchronous pipeline model with back-pressure
  - Kafka consumers are automatically paused when server is busy
- Supports sequential (stateful) and unordered (stateless) processing
  - Provide unique stream ID/topic/partition info for persistence layer
- Pass data as MATLAB timetables
- Partition aware – enables full exploitation of partition-based parallelism

System Architect

4 Integrate with Production Systems

Kafka connector architecture

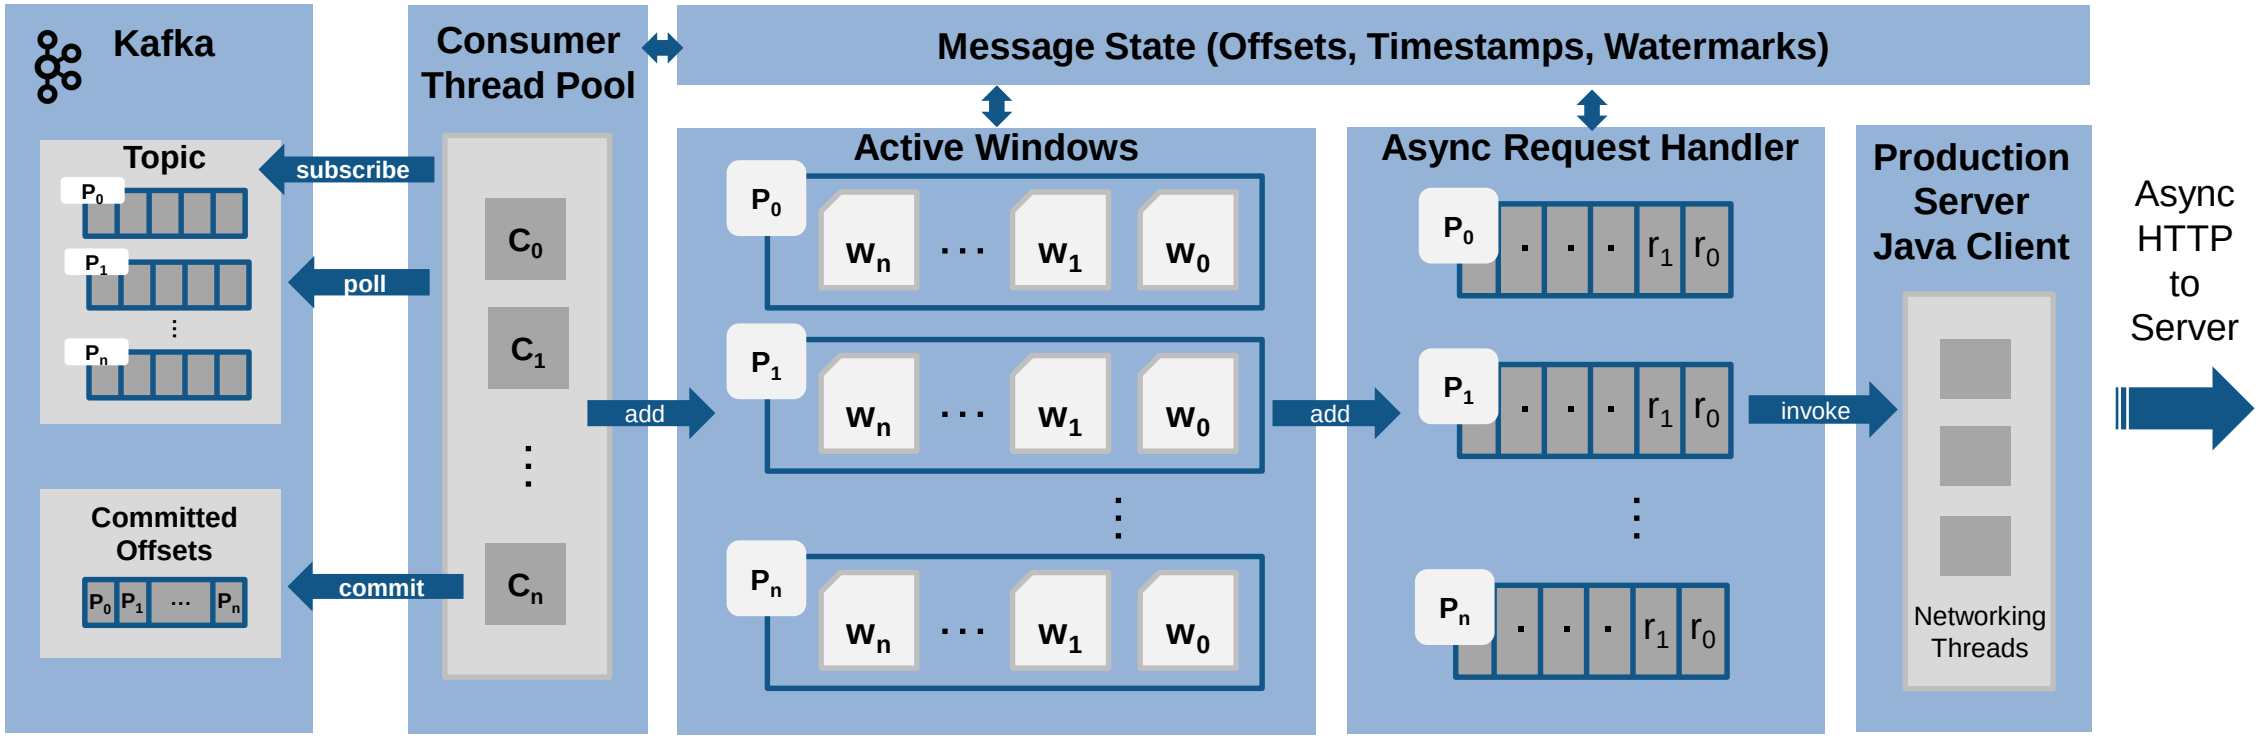

System Architect

4 Integrate with Production Systems

Streaming data is treated as an unbounded Timetable

Input Stream

Event Time	Pump Id	Flow	Pressure	Current
18:01:10	Pump1	1975	100	110
18:10:30	Pump3	2000	109	115
18:05:20	Pump1	1980	105	105
18:10:45	Pump2	2100	110	100
18:30:10	Pump4	2000	100	110
18:35:20	Pump4	1960	103	105
18:20:40	Pump3	1970	112	104
18:39:30	Pump4	2100	105	110
18:30:00	Pump3	1980	110	113
18:30:50	Pump3	2000	100	110
…	…	…	…	…

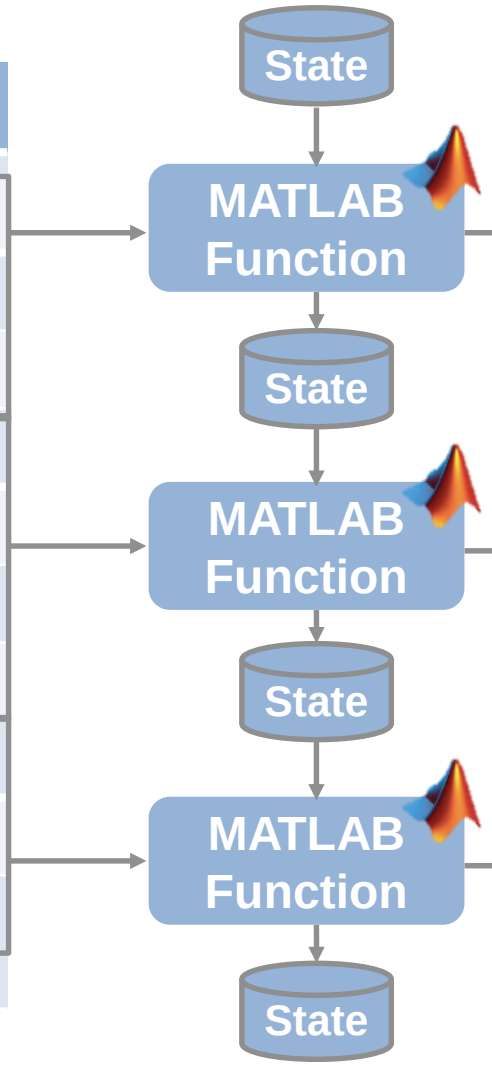

Output Stream

Time window		Pump Id	Bearing Friction
…		…	…
18:00:00	18:10:00	Pump1	5
		Pump3	…
		Pump4	…
18:10:00	18:20:00	Pump2	7
		Pump3	3
		Pump4	…
18:20:00	18:30:00	Pump1	…
		Pump3	4
		Pump4	…
18:30:00	18:40:00	Pump5	…
		Pump3	5
		Pump4	8

System Architect

4 Integrate with Production Systems

Debug your streaming function on live data

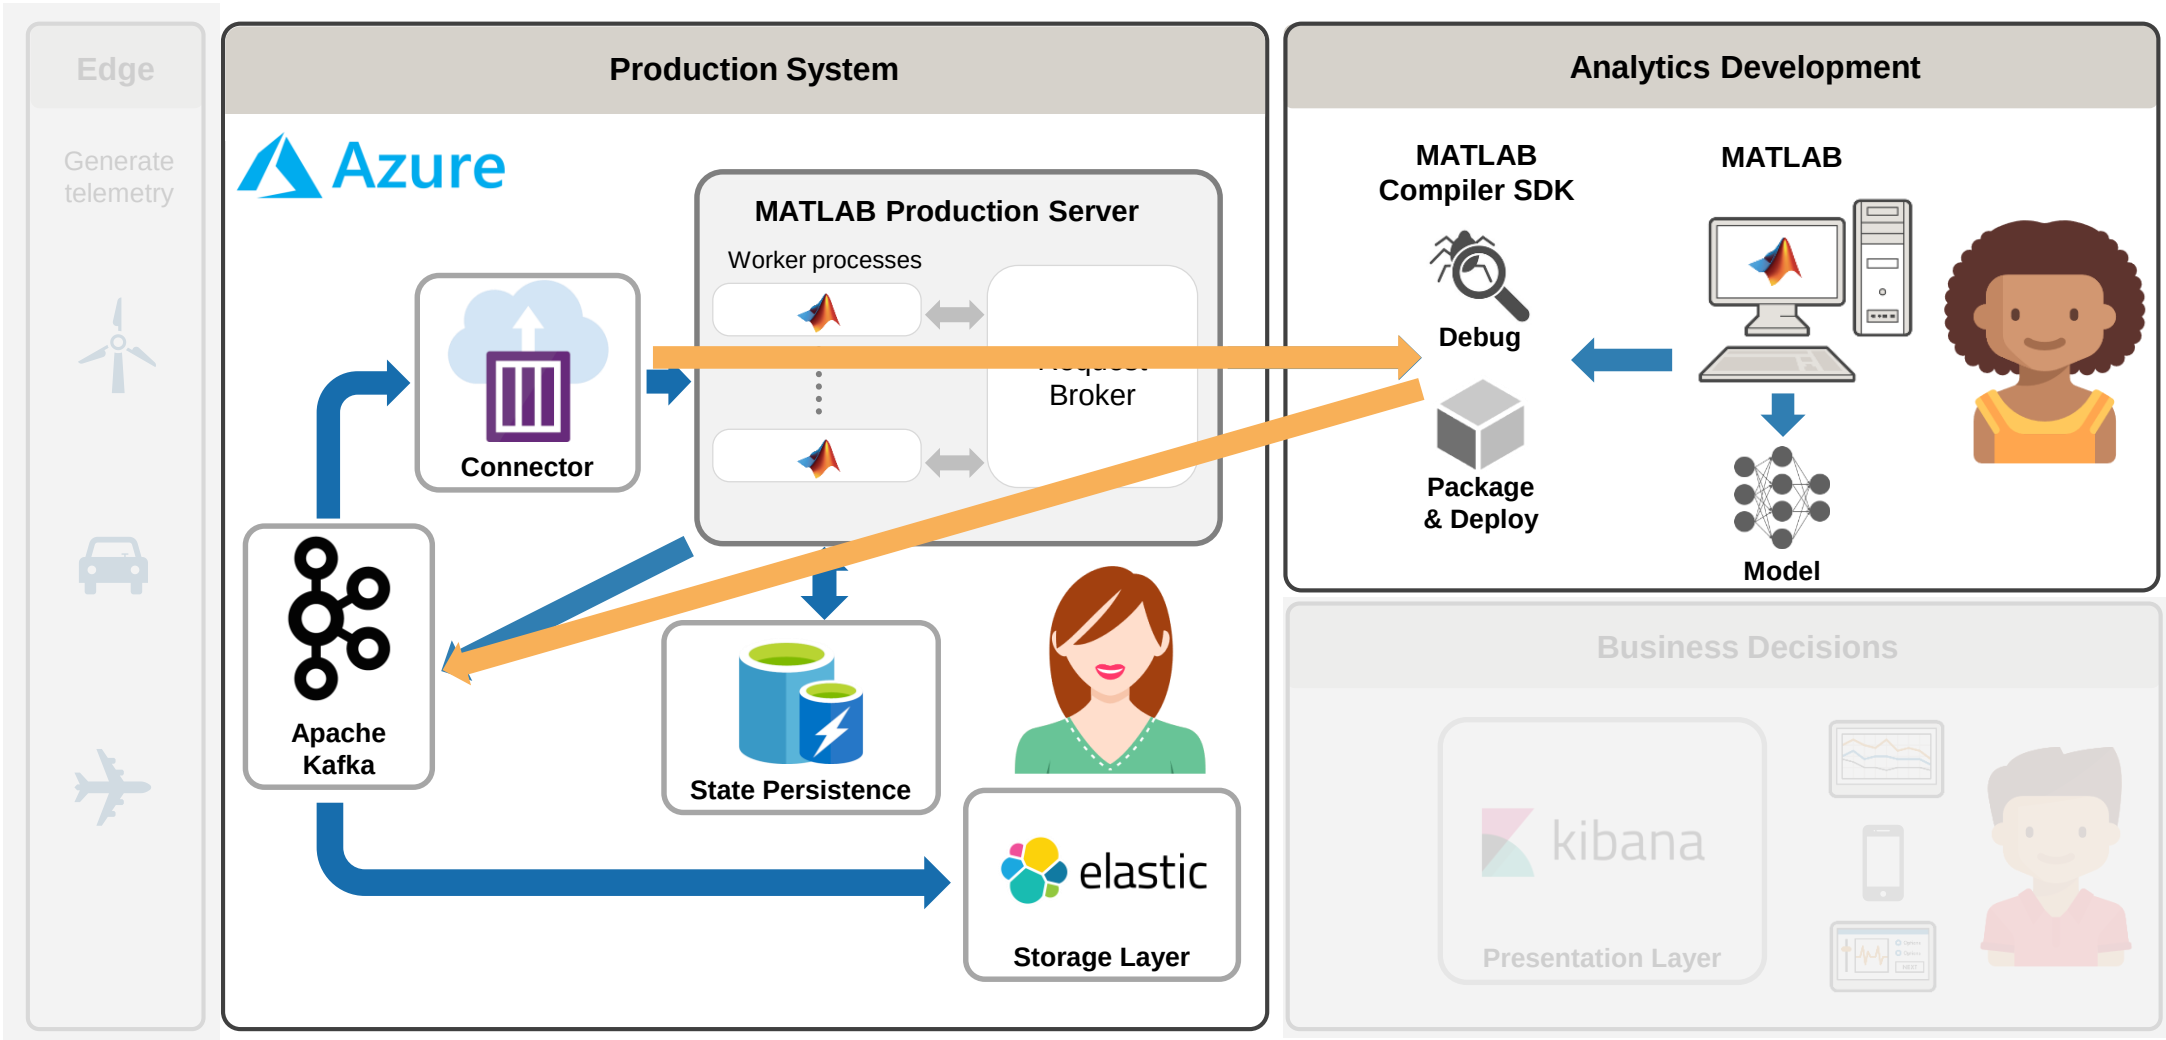

System Architect

4 Integrate with Production Systems

Debug a Stream Processing Function in MATLAB

MATLAB R2018b

HOME | PLOTS | APPS | EDITOR | PUBLISH | VIEW

Search Documentation | Sign In

New | Open | Save | Find Files | Compare | Print | Go To | Find | Insert | Comment | Indent | Breakpoints | Run | Run and Advance | Run Section | Advance | Run and Time

FILE | NAVIGATE | EDIT | BREAKPOINTS | RUN

C: ▸ Users ▸ jstewart ▸ Projects ▸ Kafka ▸ Connector ▸ StreamingExample ▸

Current Folder

Name

+mps
PumpFault_Linux
StreamingExample
test
consume.m
consumeResults.m
fakeResultsToKafka.m
make_cache_name.m
PumpFault_Linux.prj
TheKnowledge.txt

PumpFault_Linux.prj (Production Server Compiler Project)

No details available

Editor - C:\Users\jstewart\Projects\Kafka\Connector\StreamingExample\StreamingExample\streamingFunction.m

streamingFunction.m | analyzeResults.m

```
%#function mlearnapp.internal.model.adapterlayer.TrainedClassificationEnsemble
%#function mlearnapp.internal.model.adapterlayer.TrainedLinearRegressionModel

if nargin < 2
    old_state = [];
end



% Model doesn't work well with too few data points.
limit = 100;
if height(data) < limit
    fprintf("Too few rows (%d < %d) to generate effective model.\n", ...
            height(data), limit);
    new_state = old_state;
    return;
end


% Get the keys present in the data, use categorical for performance
data.key = categorical(data.key);
pumps = string(categories(data.key));

% Load models
persistent leakModel bearingModel blockingModel trainedModel

if isempty(leakModel)
    x = load('MLModels.mat');
    leakModel = x.leakModel;
    bearingModel = x.bearingModel;
    blockingModel = x.blockingModel;
    trainedModel = x.trainedModel;
end
```

Workspace

Name	Value
ctrl	1x1 Contr...

Command Window

New to MATLAB? See resources for Getting Started.

fx >>

System Architect

4 Integrate with Production Systems

Complete your application

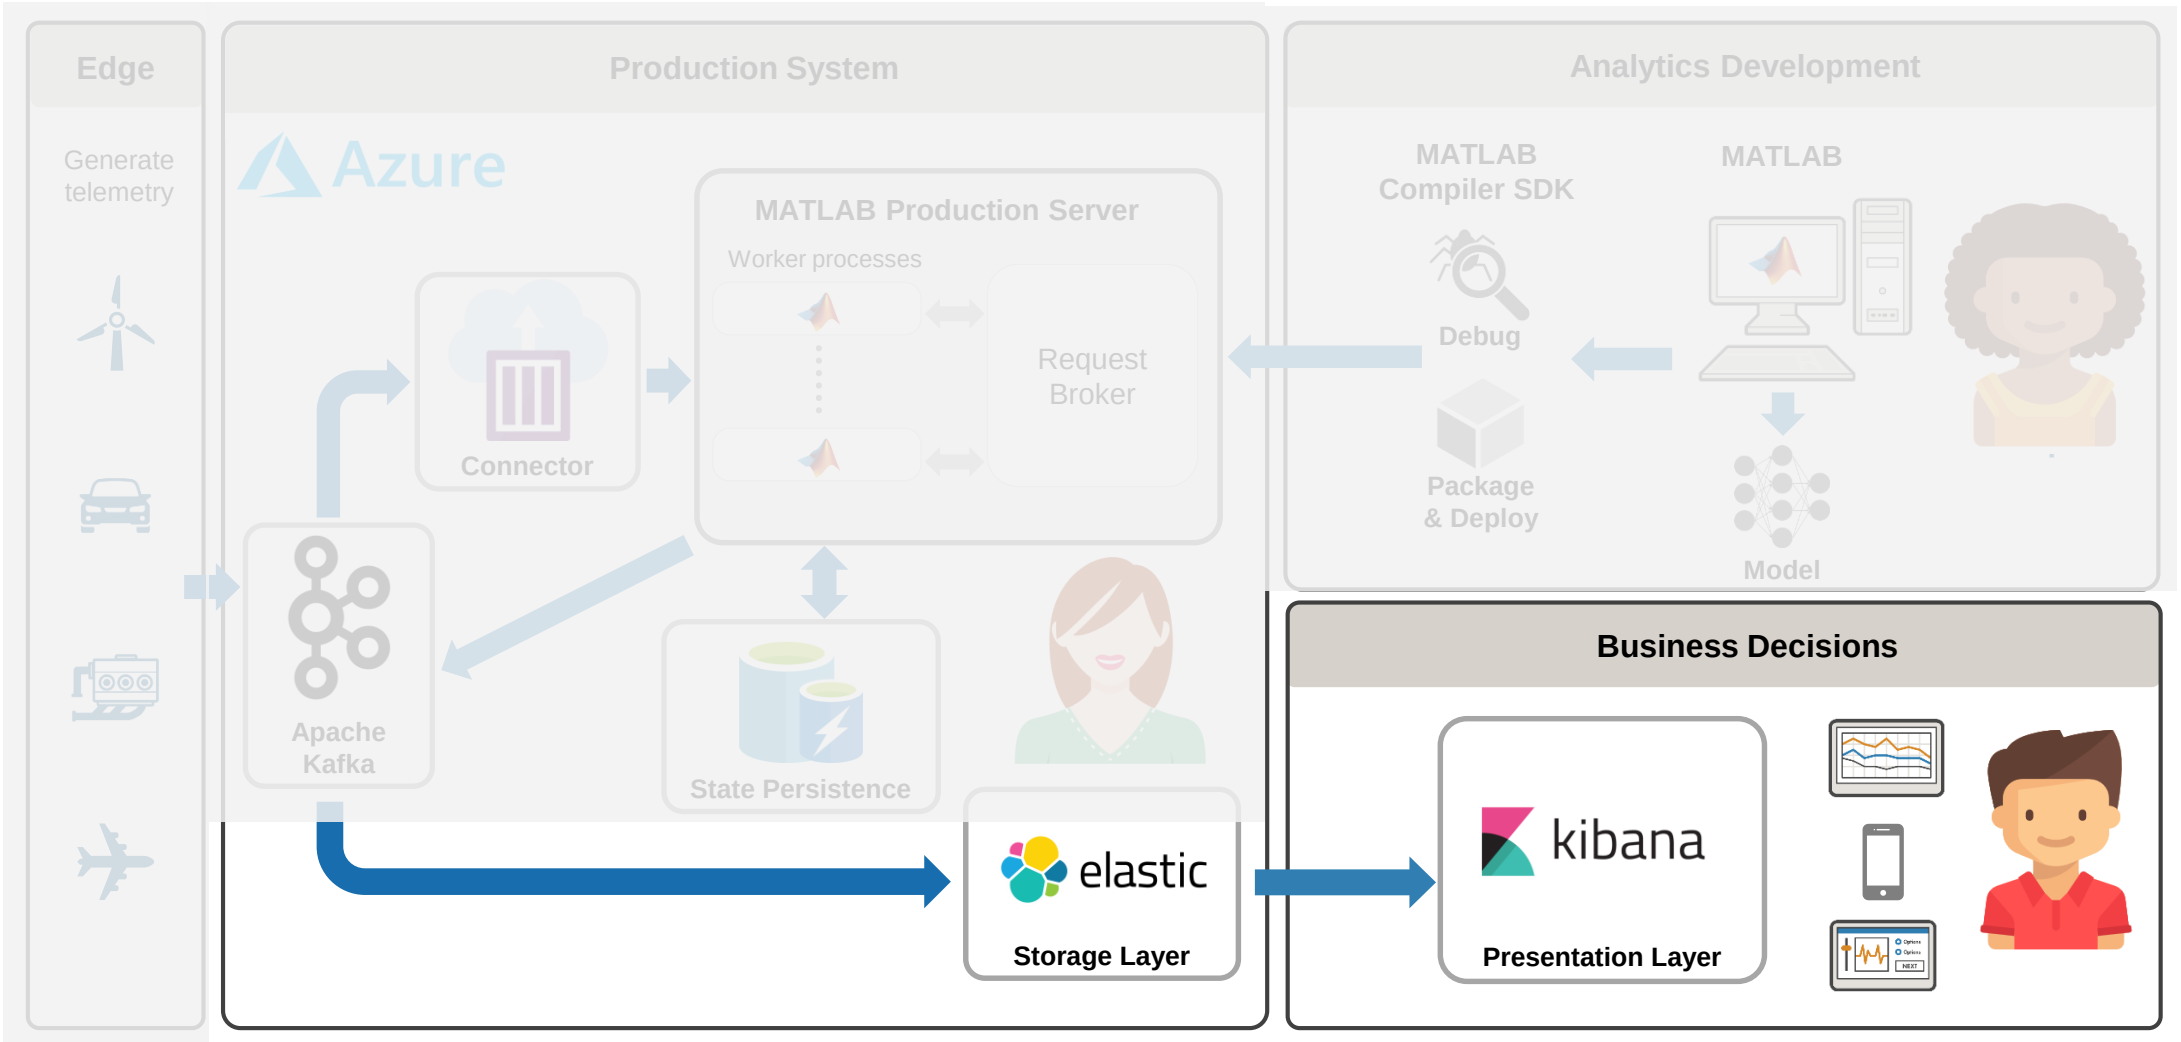

5 Visualize Results

Plant Operator

Complete Your Application

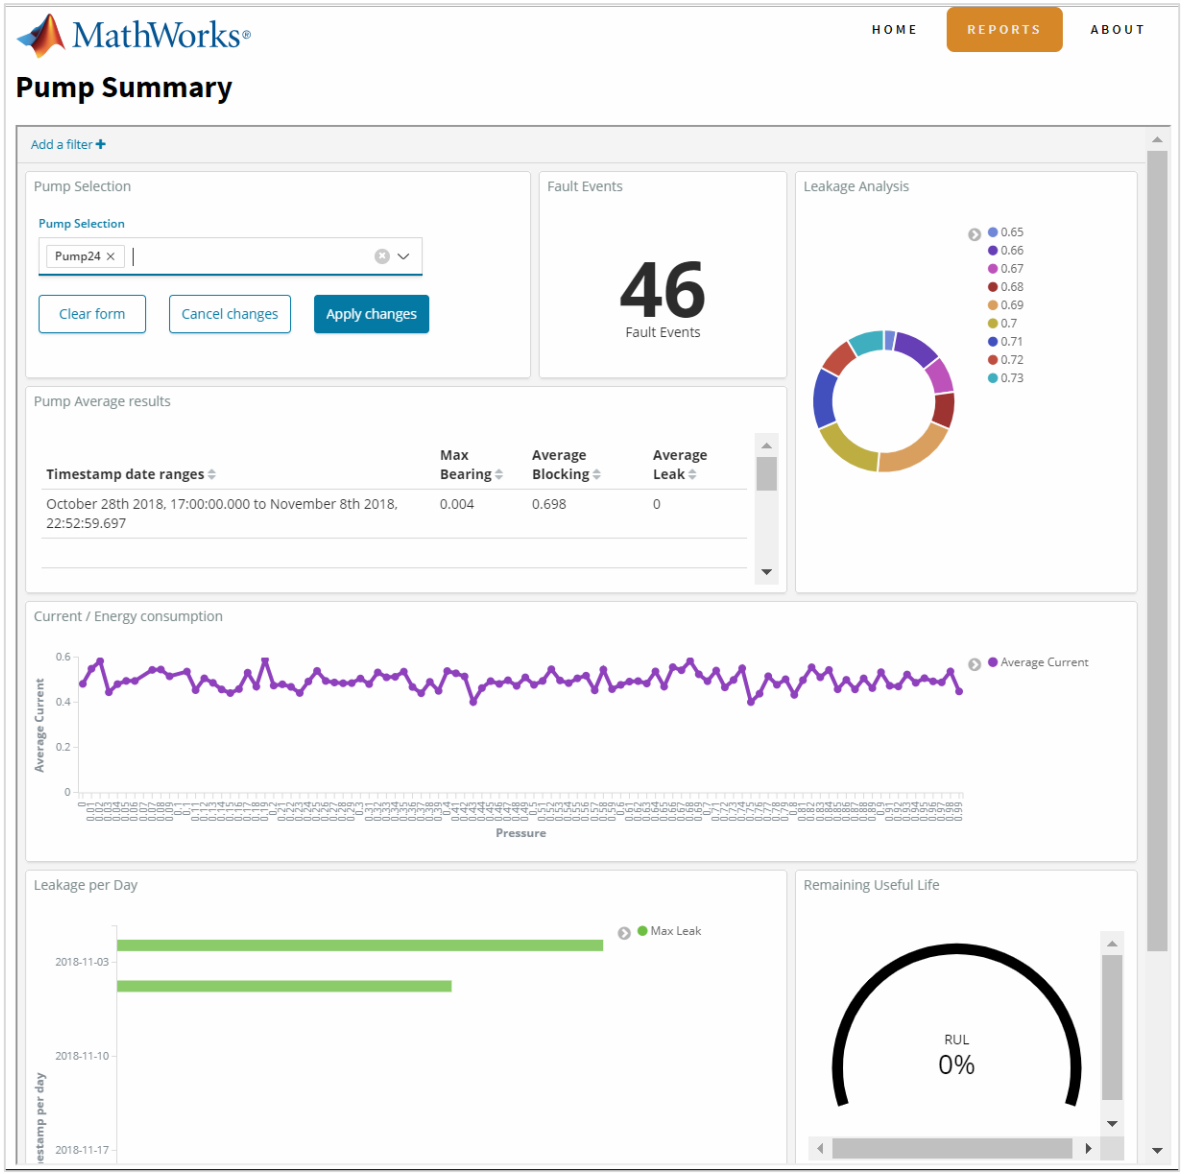

Team Retrospective

- Completed demo of full system in 3 week sprint
- Successfully used digital twin to generate faults and train models
- Fast prototyping of physical and AI models with MATLAB and Simulink. Easy integration with OSS
- Cloud platform enabled faster IT setup
- Next steps:
  - Make model adjustments
  - Test against real pump
  - Customize dashboard for Operator’s needs

Resources to learn and get started

- GitHub: MathWorks Reference Architectures
- Working with Enterprise IT Systems
- Data Analytics with MATLAB
- Simulink

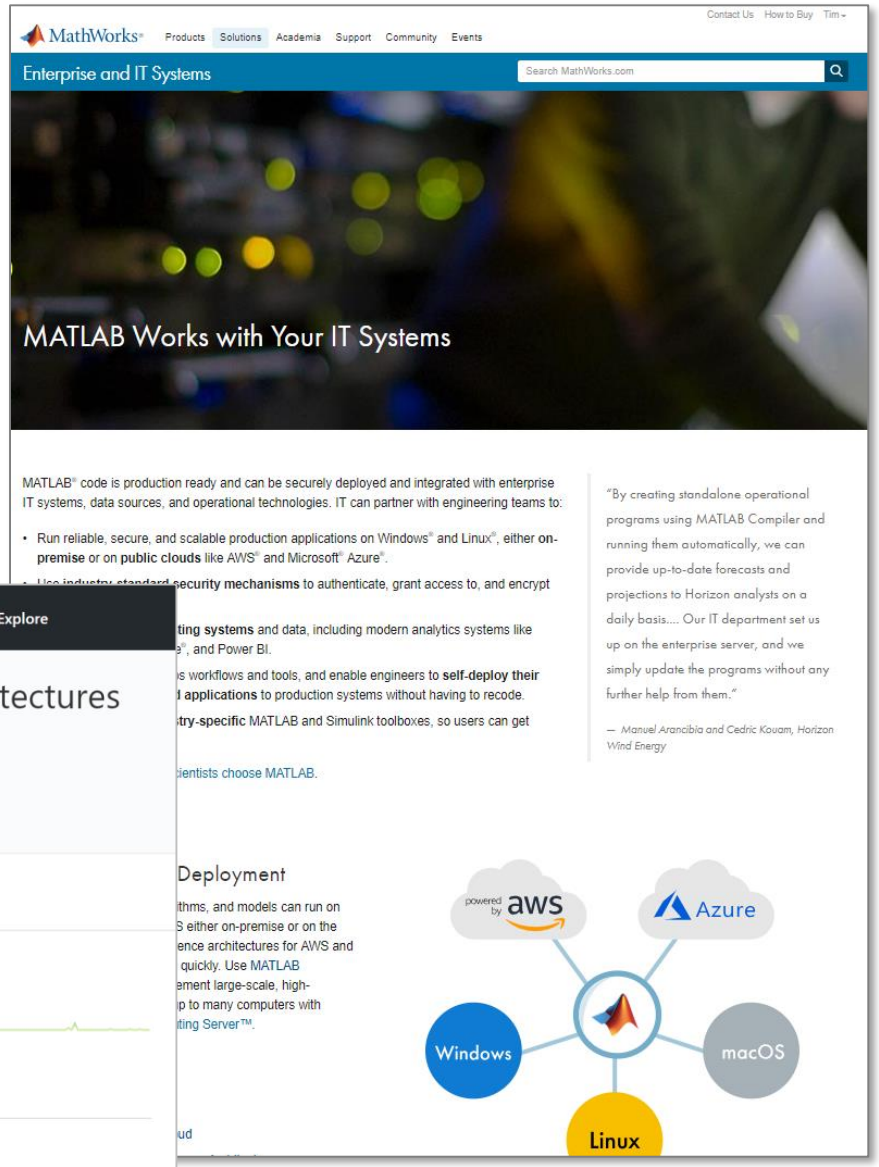

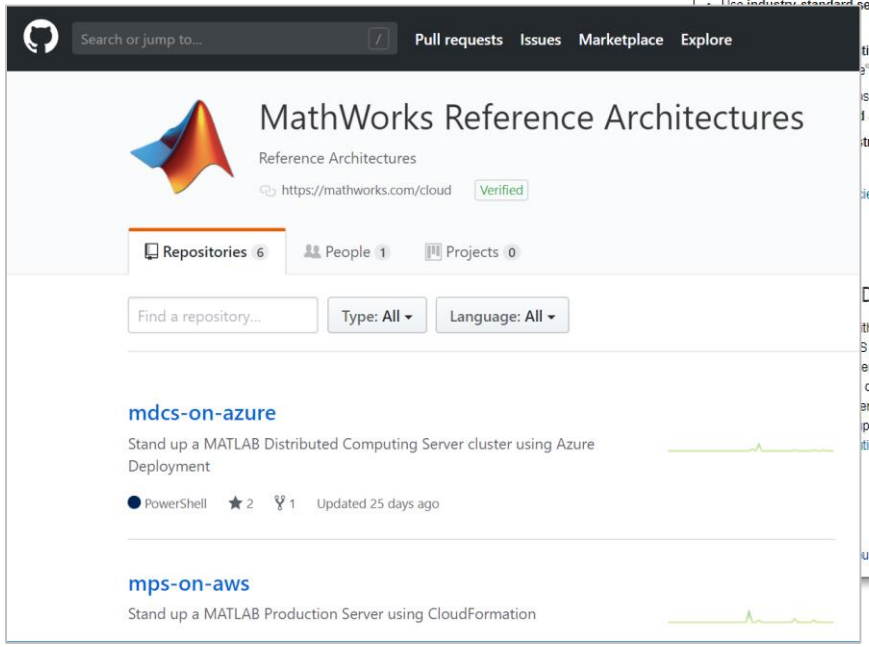

Summary: MathWorks Cloud Offerings

		Pre-Built Solution	Pre-Built Platform	Build your Own
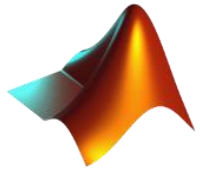	**MATLAB**	MATLAB Online MATLAB Mobile MATLAB Drive		Reference Architectures • Virtual Machine • Formation Scripts • Networking • Documentation
	**MATLAB Distributed Computing Server**		Cloud Center	
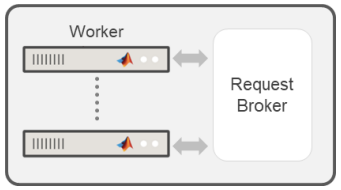	**MATLAB Production Server**			

Learn more

www.mathworks.com/cloud

Q & A