Putting AI to work in an Enterprise: Deep Learning as a Service

Nick Werstiuk
Director Offering Management
IBM Spectrum Computing
werstiuk@ca.ibm.com

© 2018 IBM Corporation

Session Outline

Deep learning has moved out of the lab and into production. But how do you provide training environments to all your internal customers working across business units and with different requirements, and how do you avoid provisioning separate clusters for each customer?

IBM has applied decades of HPC experience to build a production ready learning stack, including servers accelerated with NVIDIA GPUs, workload and resource management software, ready to use open source frameworks and it's all covered by IBM support.

The solution provides a secure multi-tenant environment so multiple data scientists can share a common set of resources, eliminating silos, while running multiple instances of the same or different applications. The deep learning effort is enhanced with end-to-end pipeline support from data ingestion and preparation, through model training and tuning, to inference. In this session, we will explore what an enterprise deep learning environment looks like and provide insights into the unique IBM value for accelerating the use of deep learning across a wide variety of industries.

Introduction/Setting the Context

- Describe the **requirements/challenges** we have seen in our client base as they build out an architecture to support enterprise wide AI and Deep Learning
- Focus **on use cases** where the data and infrastructure is typically deployed **'behind the firewall'**, and **shared amongst many users, applications, LOB's**
- **Tactical** starting point vs. **Strategic** end goal and vision
- Approach we have taken to help clients address the challenges to deliver on the vision.

© 2018 IBM Corporation

Client Goal: High Performance Data Science as Service

- Spark and AI centric environment
- Load data from Hortonworks data lake and other enterprise data sources
- Hundreds of users, interactive notebooks, ad-hoc queries & batch reports
- ML and DL Model development training

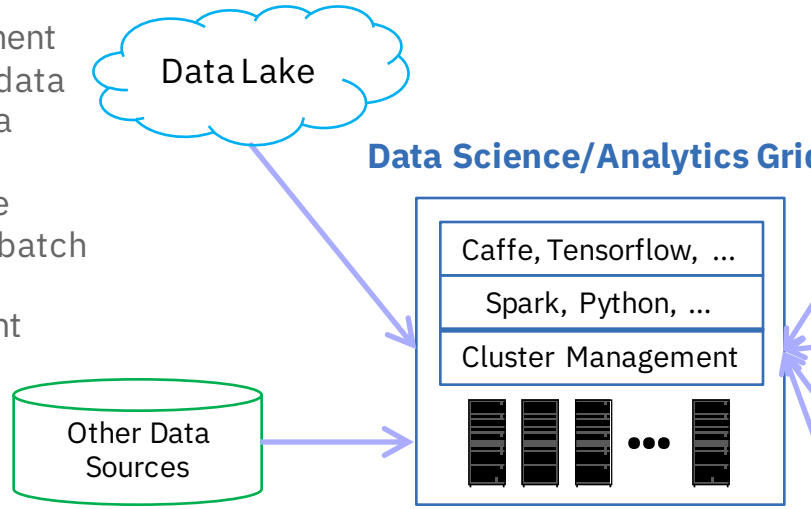

Data Scientists
- Via note book, analyze data and develop new applications, Models

Data Engineering
- Consume ETL as service - develop new ETL apps

Fraud Detection Teams
- Run third party ML application
- Run risk reporting

Risk & Quant Teams
- Develop risk models, train DNN/CNN models

- High performance **POWER and X86** environment for compute & memory intensive workloads
- **Spark in memory** analytics for real time/ad-hoc query/analysis & batch reporting, data transformation, ETL
- **Dis-aggregated compute/storage infrastructure**; Scale compute & storage independently
- **True multi-tenant**: Many LOBs, users, resource, and application SLAs
- **Multiple Framework versions**, multiple notebook versions
- **GPU acceleration**, Shared RDD optimization & **new innovations around ML/DL**
- Support Spark, batch, micro-service application frameworks in the same environment

HPC & HPDA Towards Cognitive, AI & Deep Learning

HPC Simulation

Oil & Gas

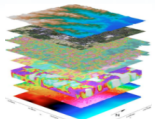

CAE

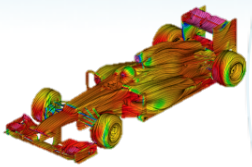

Eledctronics

Life Sciences

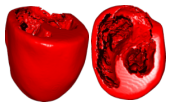

Cognitive and Deep Learning

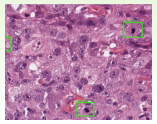

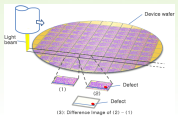

High Performance Data Analysis

Big Data

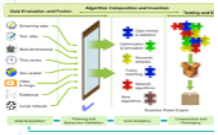

Business Intelligence

Financial Analytics

Social Analytics

IBM Spectrum LSF

IBM Spectrum Conductor

IBM Spectrum Symphony

© 2018 IBM Corporation

1. Diverse Users and Applications

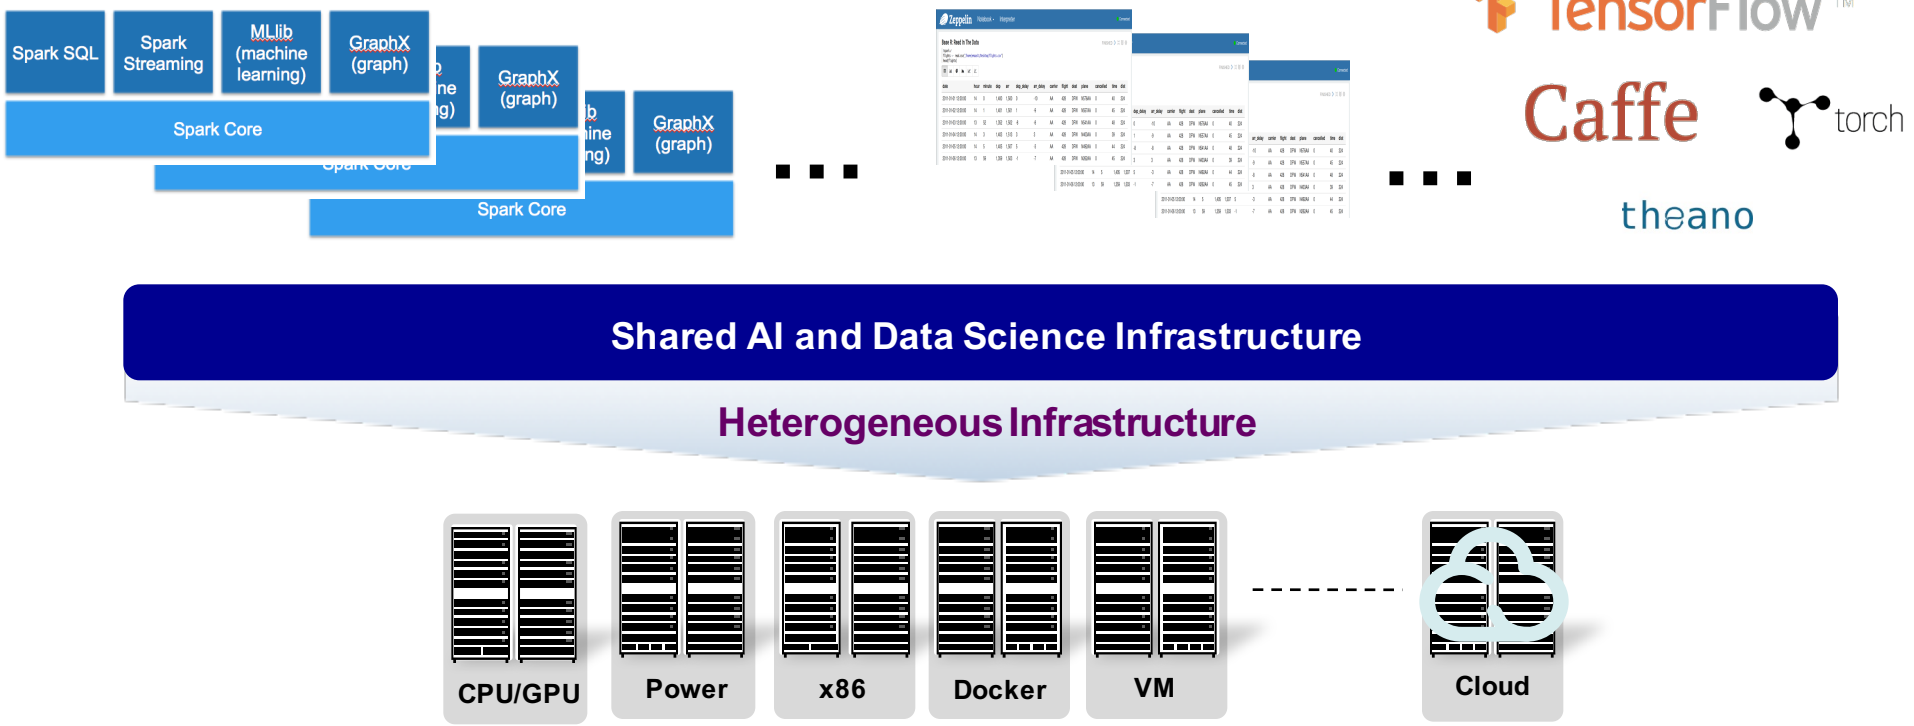

- Multiple groups and lines of business with diverse range of data scientists & analyst end users
- Different application types (notebooks, streaming, batch, interactive query, training, inference)

© 2018 IBM Corporation

2. Dynamic and Agile

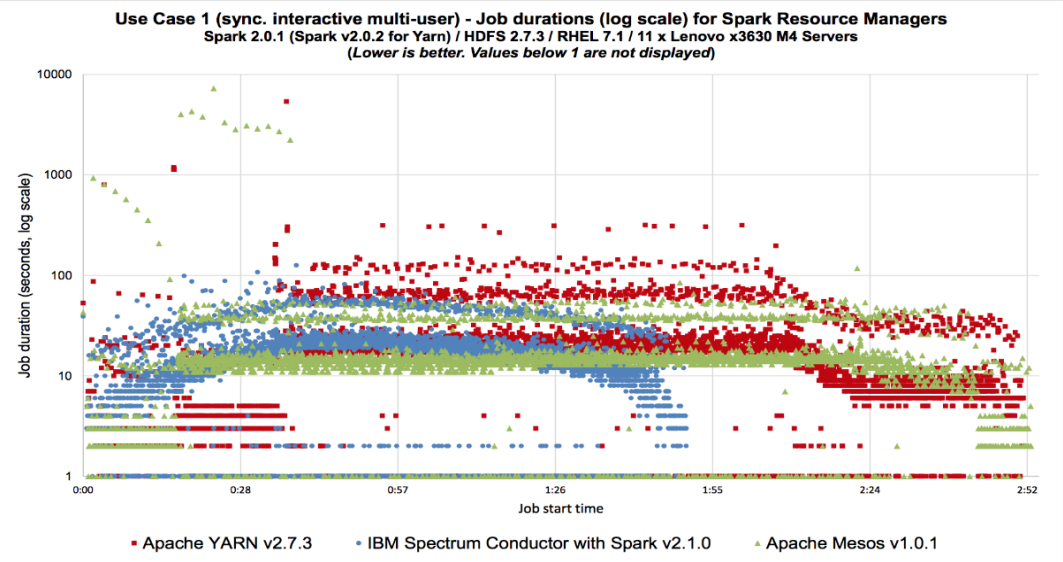

- UC1
  - 56% higher than YARN
  - 57% higher than Mesos
- UC2
  - 30% higher than YARN
  - 62% higher than Mesos
- UC3
  - 55% higher than YARN
  - 88% higher than Mesos
- UC4
  - 224% higher than YARN
  - 25% higher than Mesos

- Significant workload variability, throughput and performance across the users and application types
- 'Bring your own' frameworks, libraries and rapid pace of ecosystems evolution drive complex application life cycle management needs

© 2018 IBM Corporation

3. Cost efficiency and SLA's

New modern scale-out applications and frameworks …

…leads to costly, complex, siloed, under-utilized infrastructure and replicated data

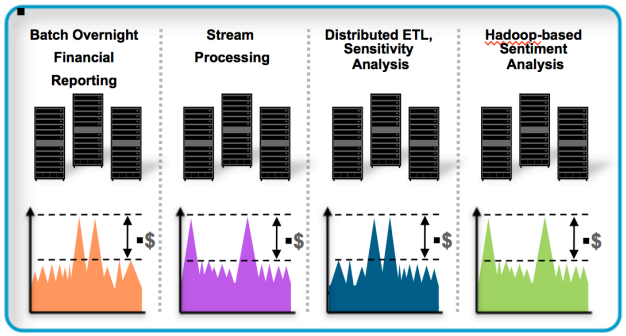

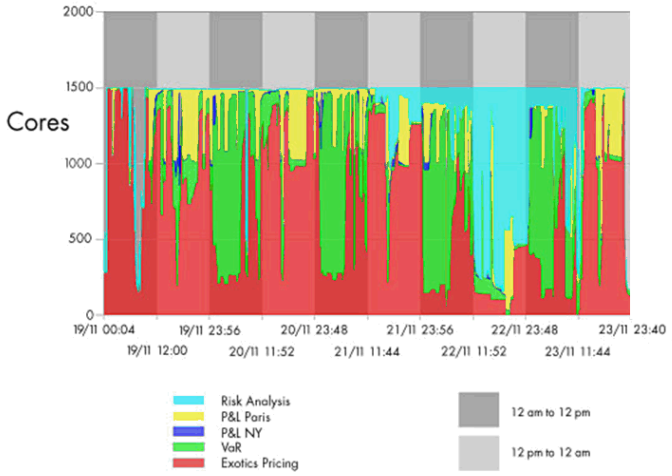

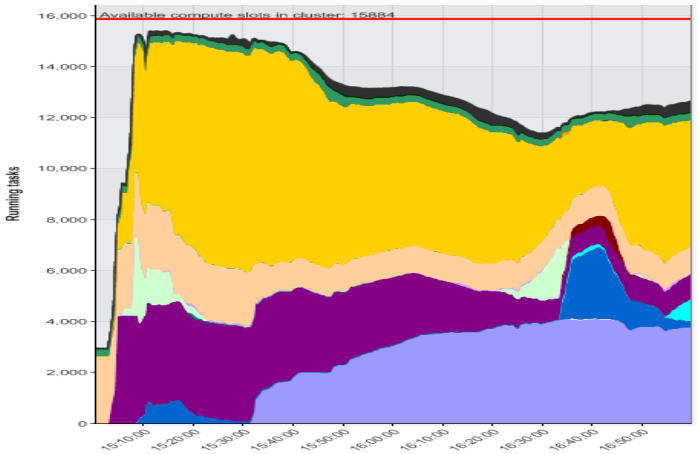

- Avoid silos by user, group, applications
- Apply the right resources to the right workload
- Offer a cost effective shared service to a wide range of users – predictable throughput and performance to ensure SLA

© 2018 IBM Corporation

4. Expansion to Multiple Data Sources and Frameworks

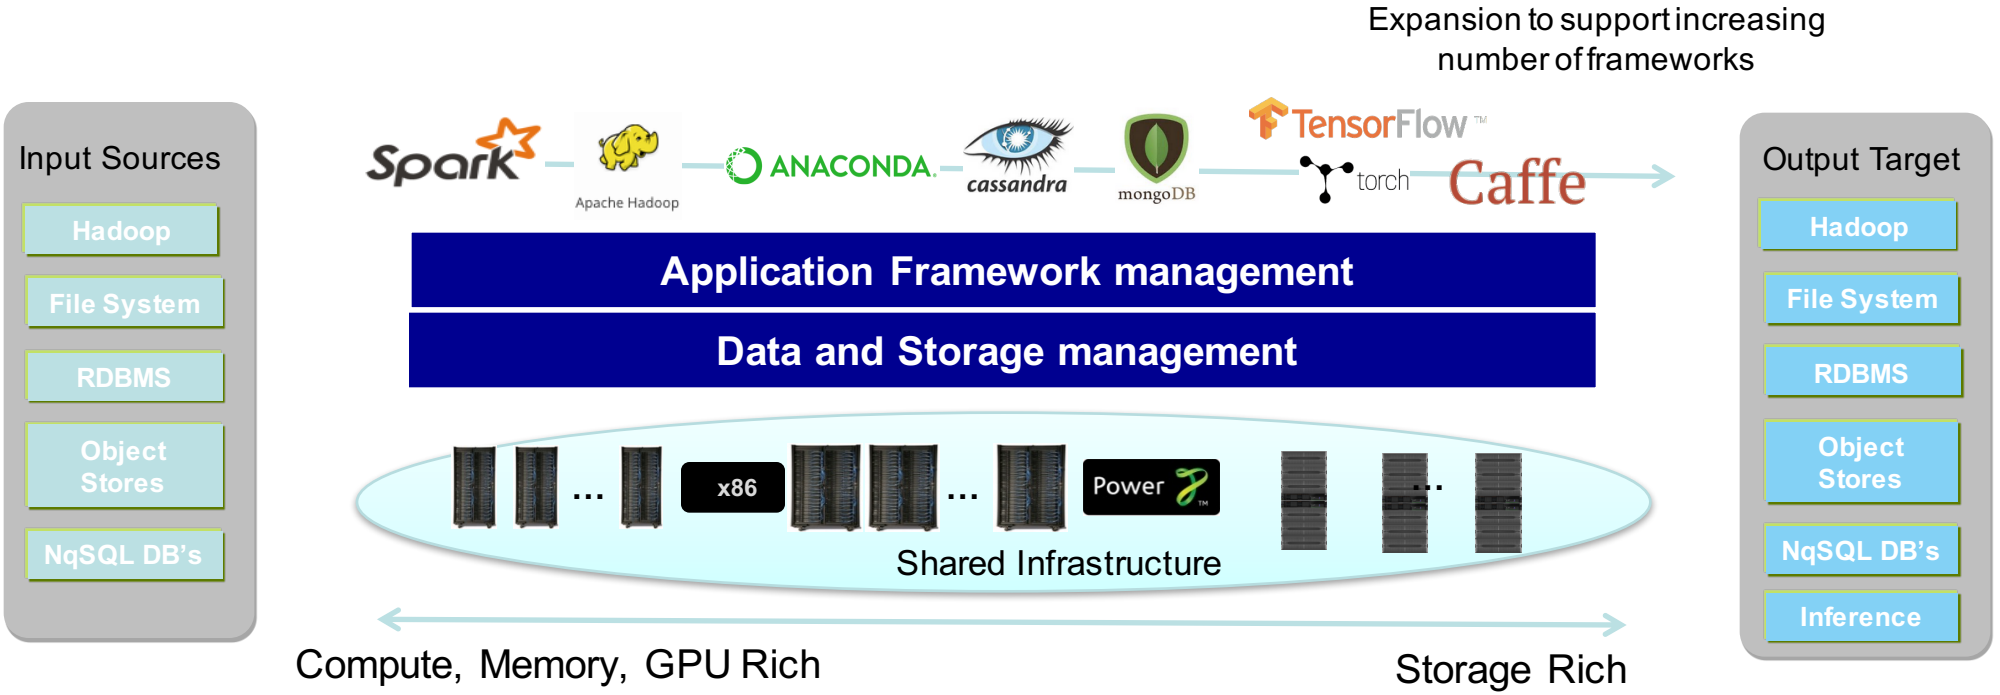

- Data Inputs and Outputs - existing Hadoop/HDFS pools but applications will also depend on a wider range of data outside of existing Hadoop
- Frameworks – ability to expand the shared infrastructure to the ecosystem of frameworks and applications.

© 2018 IBM Corporation

5. Maximize use/Flexibility of GPU Resources

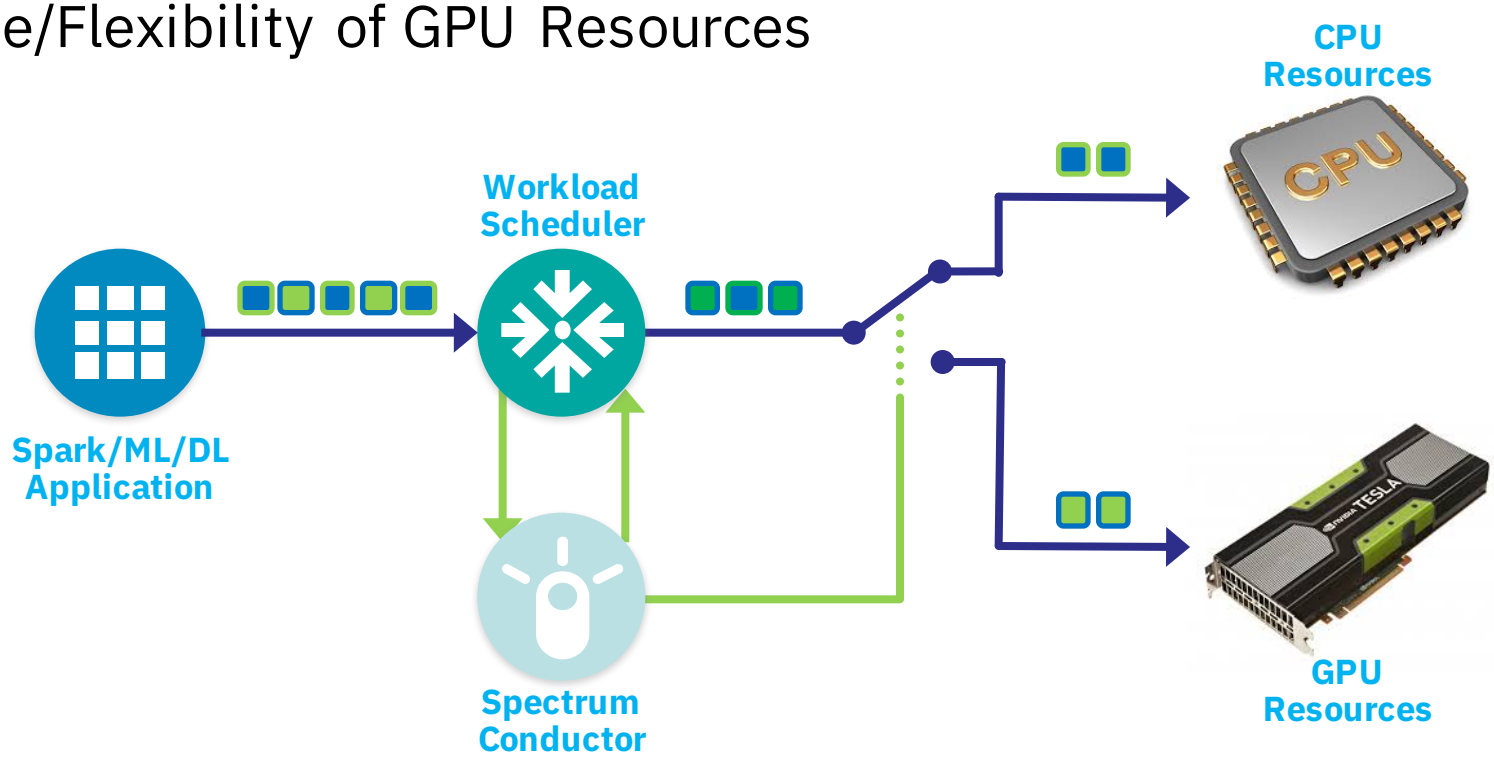

- Resource manager is GPU aware with monitoring and control
- Identify and decouple GPU operations from CPU operations
- Run and optimize application on GPU/CPU cluster transparently
- Share and Utilize both CPU and GPU efficiently to between applications

© 2018 IBM Corporation

Summary: Challenges in Managing DL Workflows

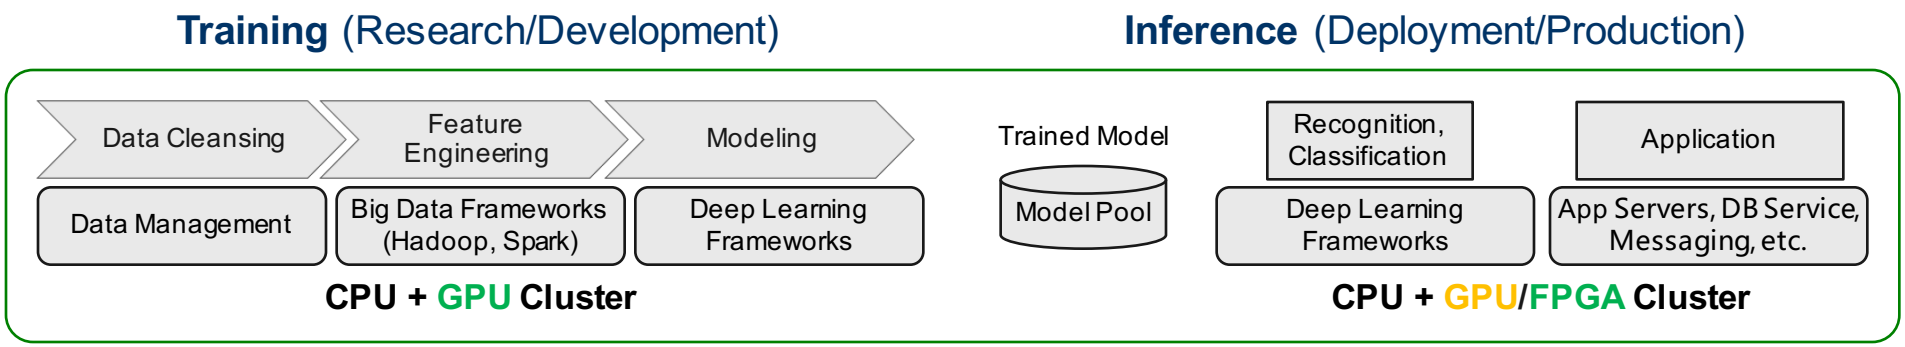

- Scalable data ingest, transformation, labeling, augmentation
- Support of different DL frameworks
- Performance & accuracy
- Parallel distribution & communication bottleneck
- GPU & topology aware scheduling
- Model & data set management
- Hyper-parameter selection & tuning
- Monitoring & execution optimization
- Resilience & elasticity

- Production deployment
- Support data pipeline & workflow
- Mixture of micro-services, batch & real time jobs
- GPU vs. CPU mixed scheduling
- Different data sources - HDFS, Object, Posix, Non-SQL, SQL)

- Multi-tenancy, QoS for data preprocessing, training & prediction
- Shared resource pool of heterogeneous resources – GPU, Power, x86, FPGA
- Multiple & different Spark, DL framework, notebooks, versions & deployment
- Security, governance
- Scalable data management

© 2018 IBM Corporation

Enterprise Deep Learning and AI Environment

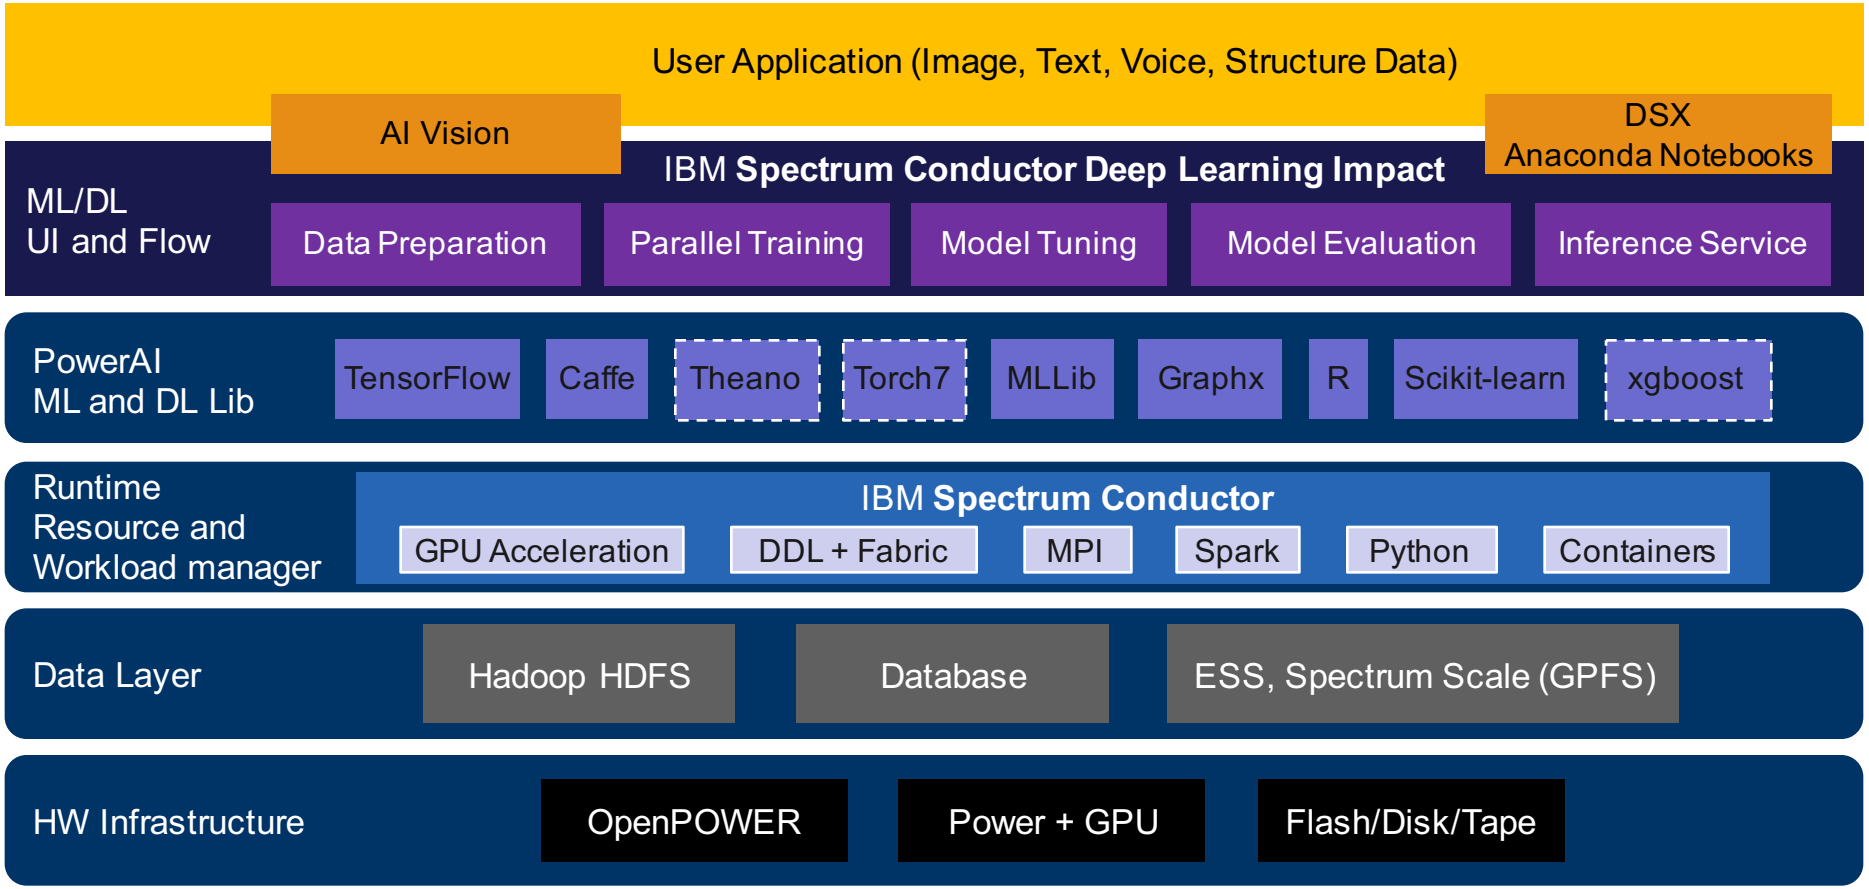

© 2018 IBM Corporation

Enterprise Class Deep Learning Solution
IBM Spectrum Conductor Deep Learning Impact
IBM PowerAI, IBM Storage

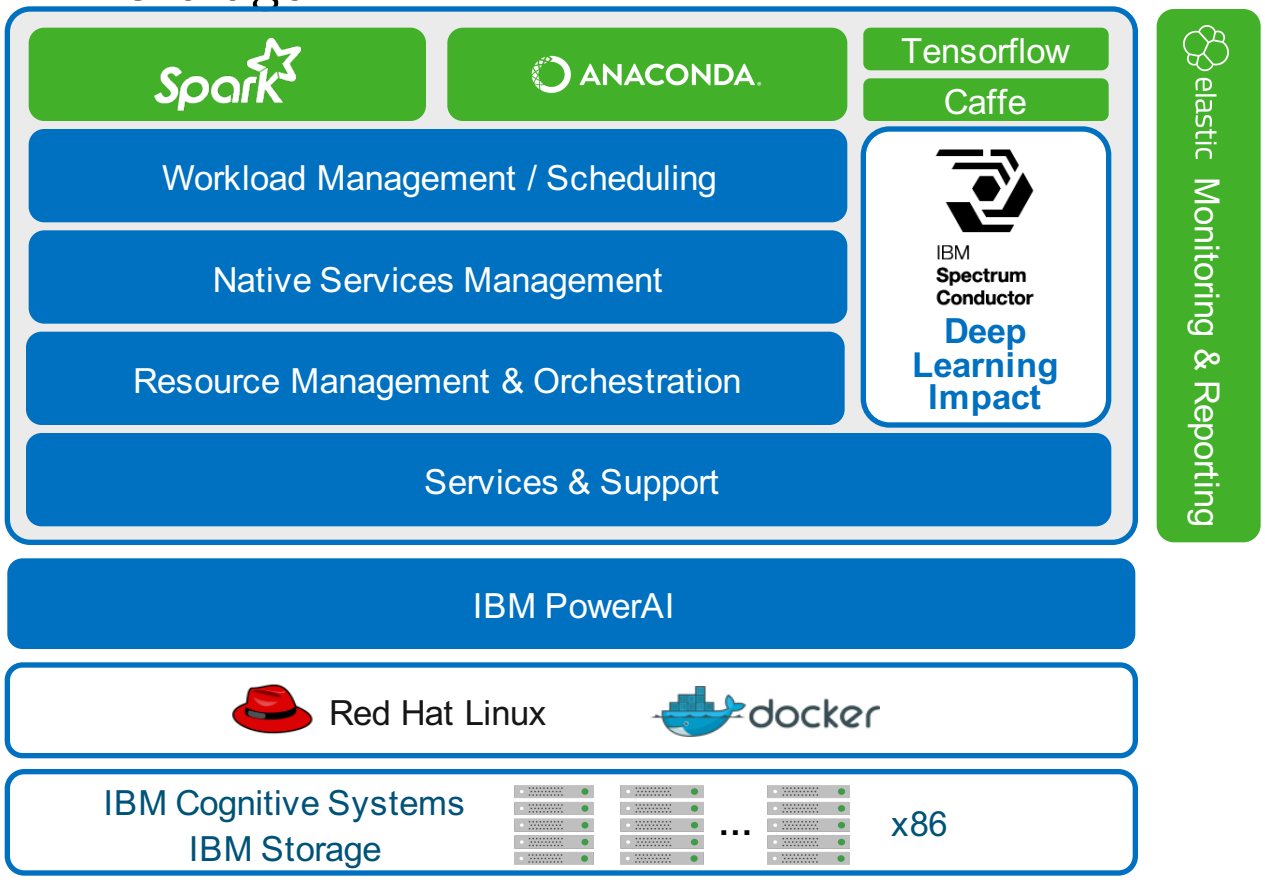

© 2018 IBM Corporation

Deep Learning Pipeline

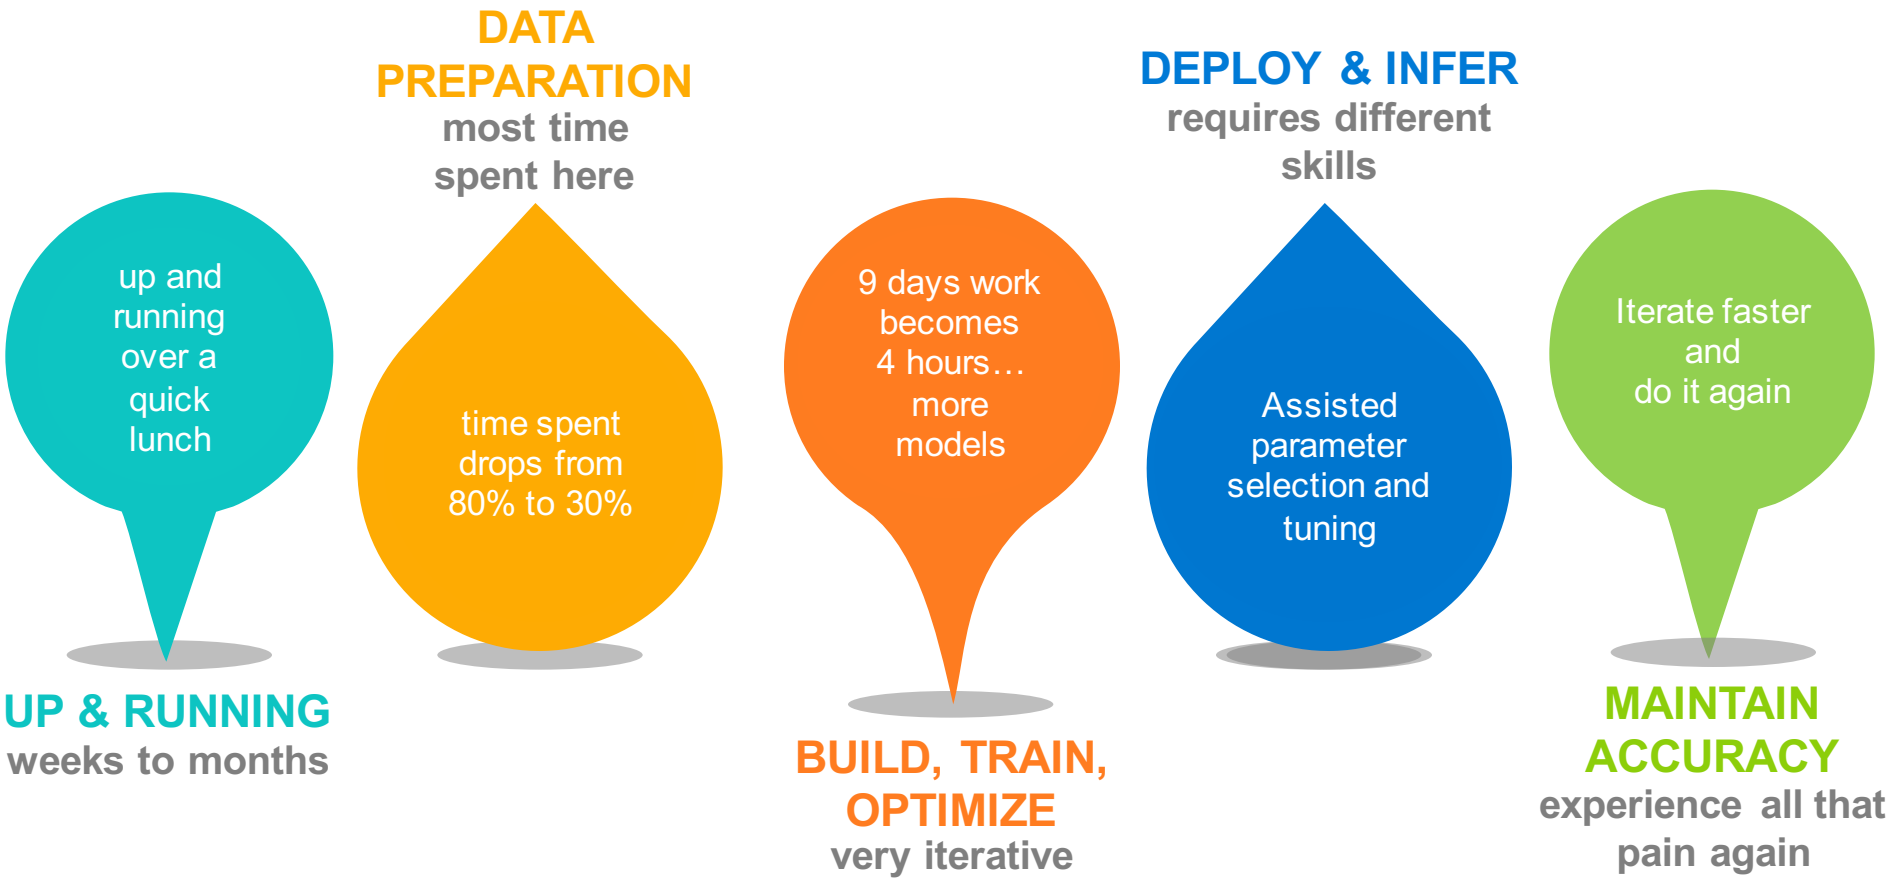

IBM Spectrum Conductor & Deep Learning Impact

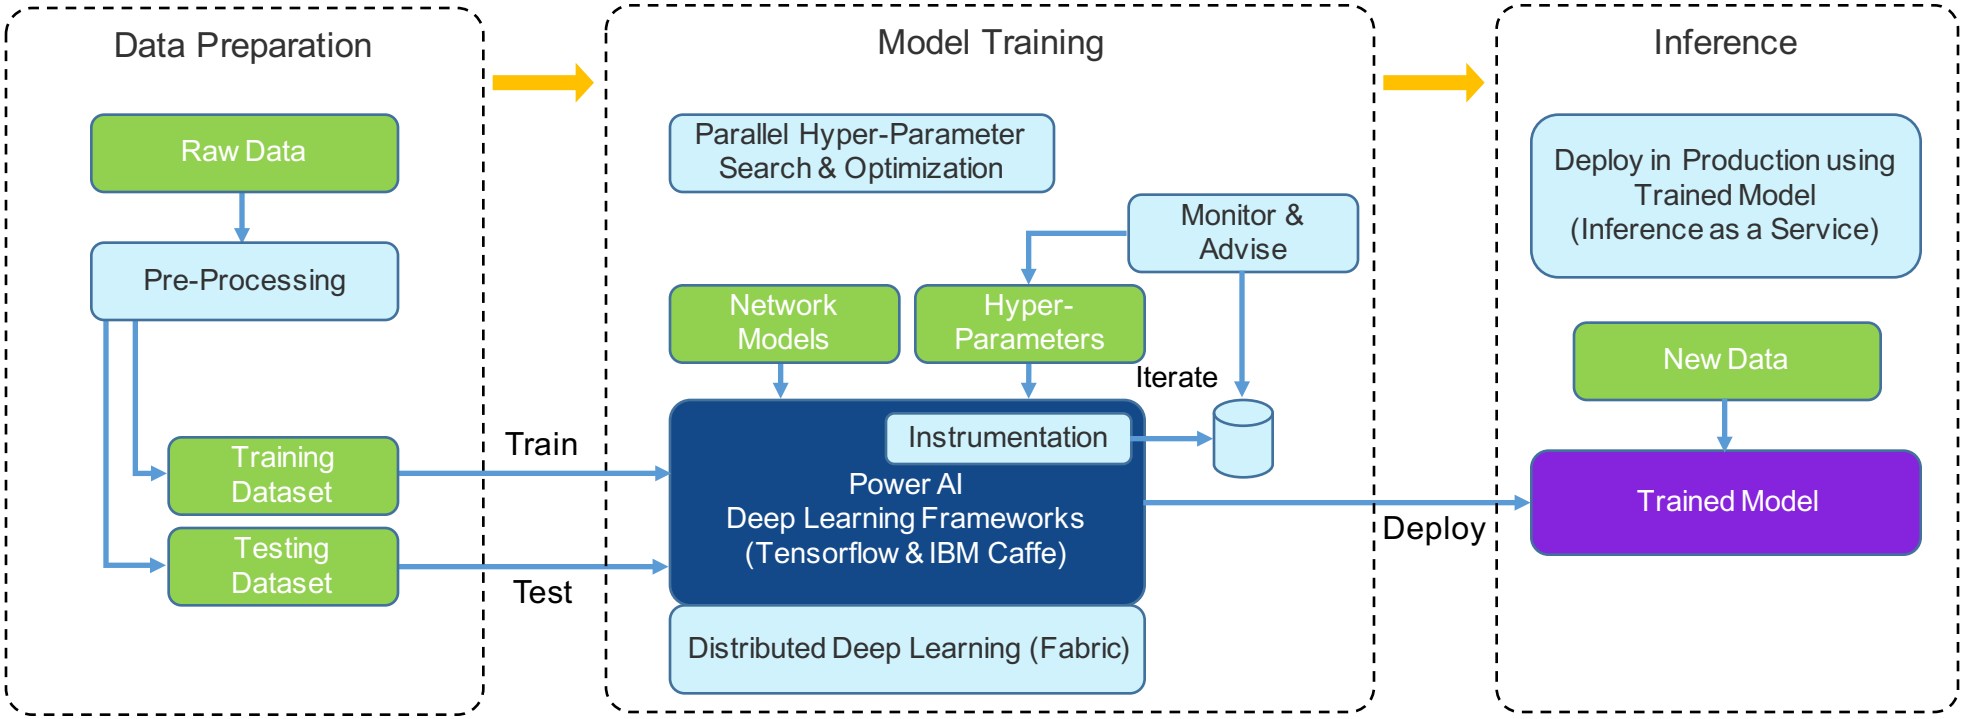

© 2018 IBM Corporation

Data Preparation for Deep Learning

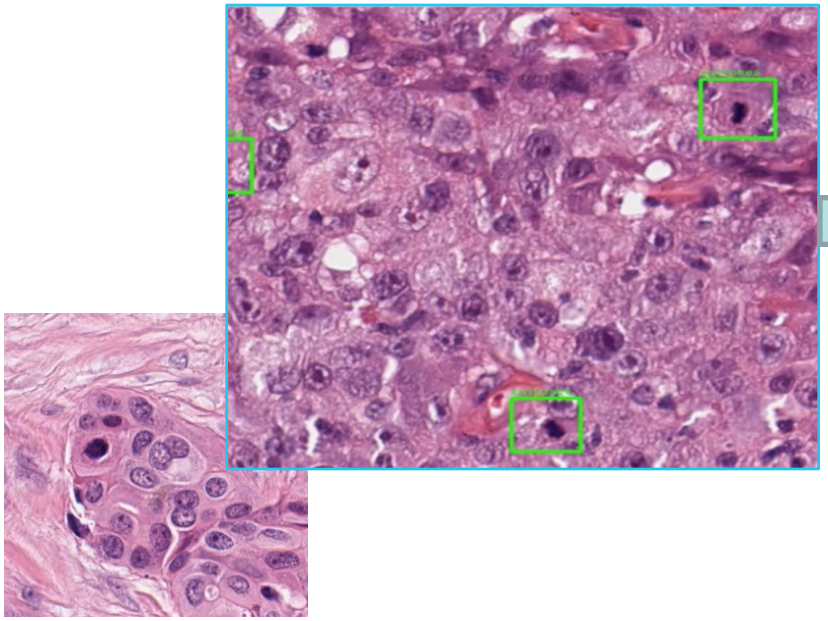

- Tumor Proliferation Assessment – mitosis detection
- Images from electron-microscope
- Size of image - 70K * 60K

Data Transformation

Framework	Format	Input Size (Faster R-CNN)
Caffe	LMDB	1K*1K
TensorFlow	TensorRecord	1K*1K

Data Distribution among training, validation and testing

Data Shuffle

© 2018 IBM Corporation

Data Preparation for Deep Learning

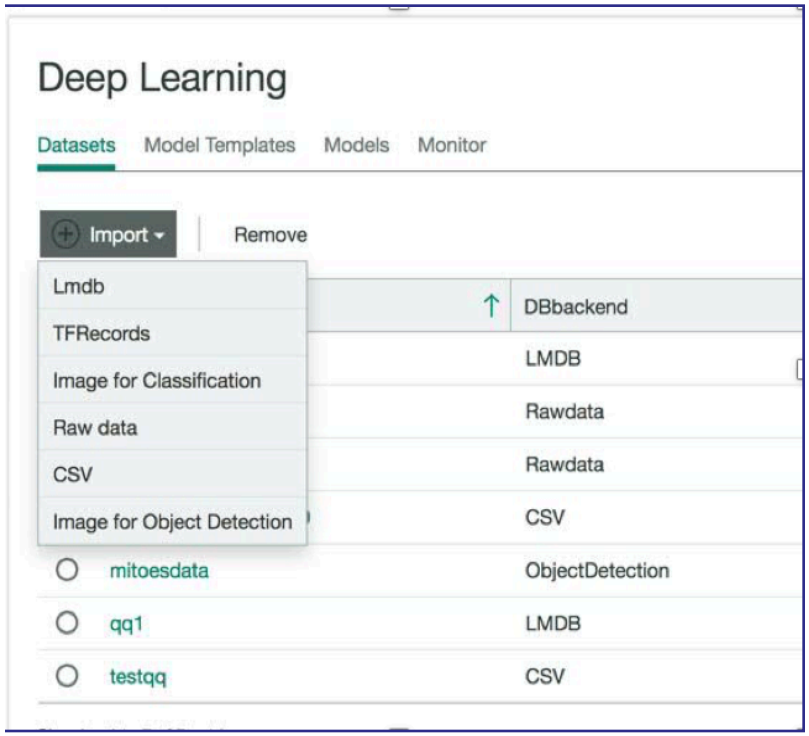

Import data from different formats

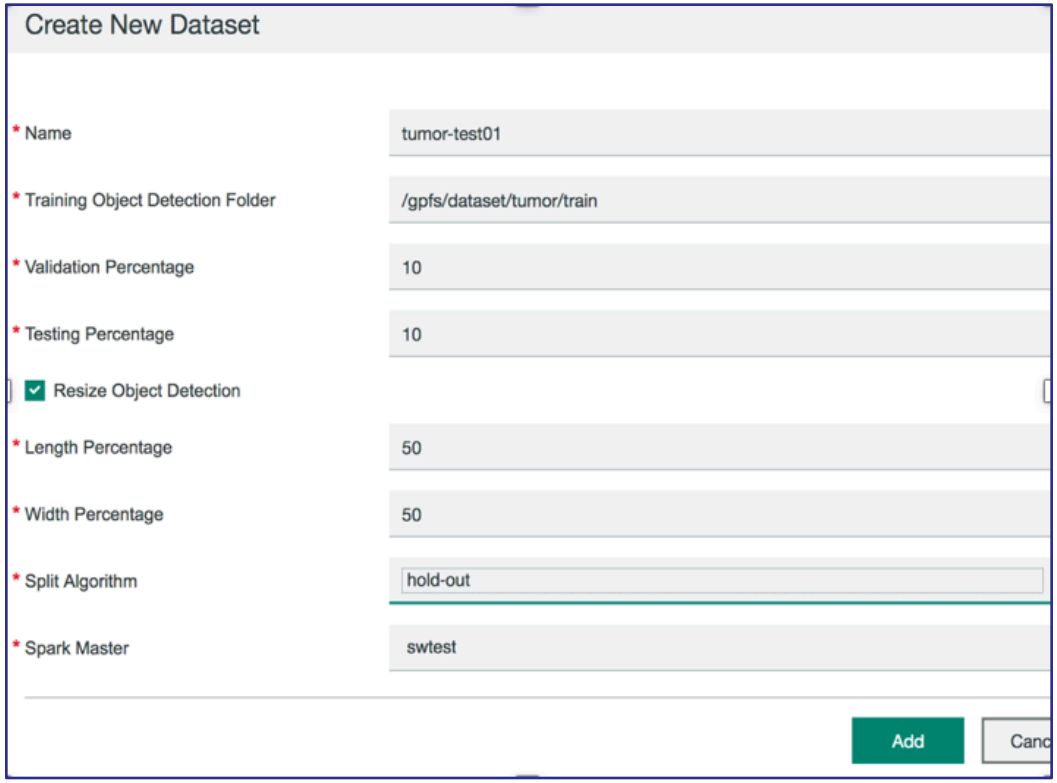

Transform, split and shuffle data

© 2018 IBM Corporation

Hyper-Parameter Tuning/Search

Hyper-Parameters

- Learning rate
- Decay rate
- Batch size
- Optimizer:
  - GradientDecedent,
  - Adadelta,
  - Momentum,
  - RMSProp
  - …
- Momentum (for some optimizers)
- LSTM hidden unit size (for models that use LSTM)

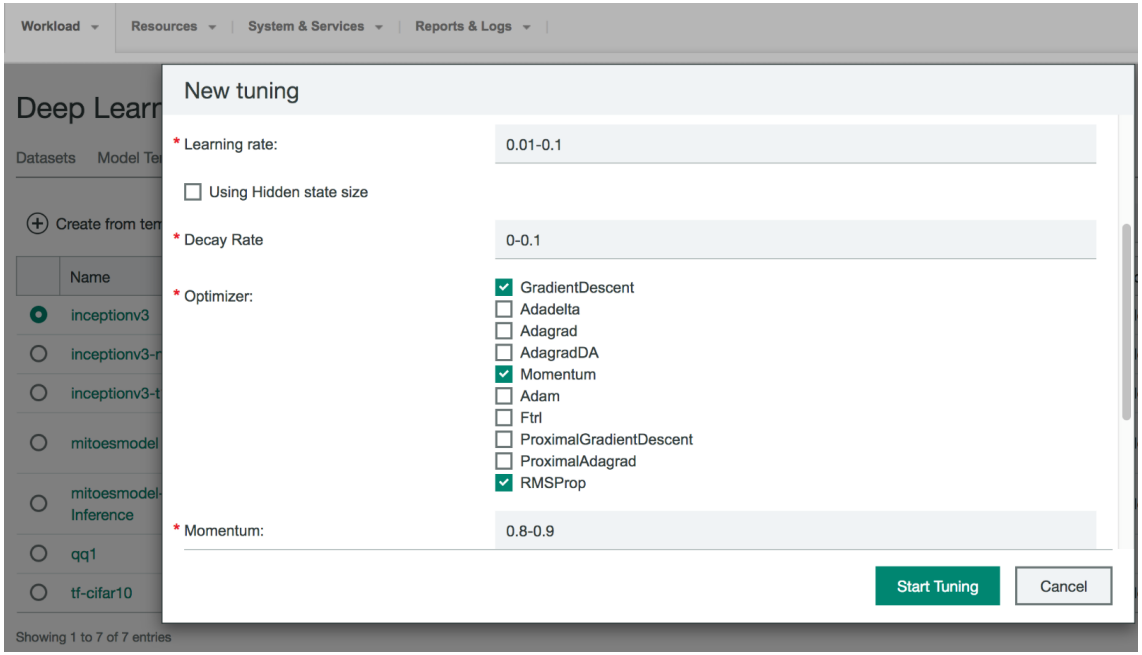

© 2018 IBM Corporation

Hyper-Parameter Search Algorithms

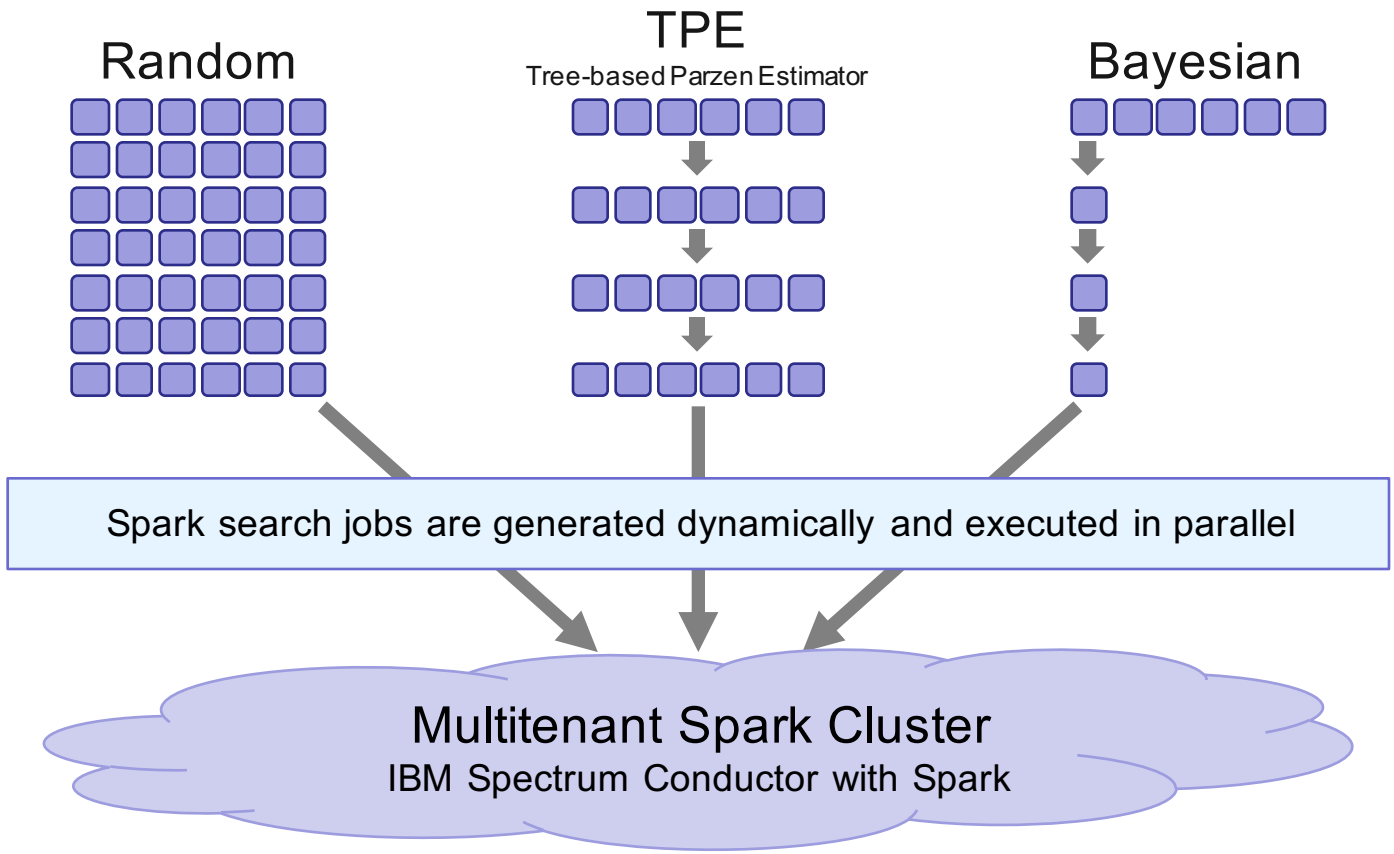

© 2018 IBM Corporation

Interactive Monitoring & Optimization

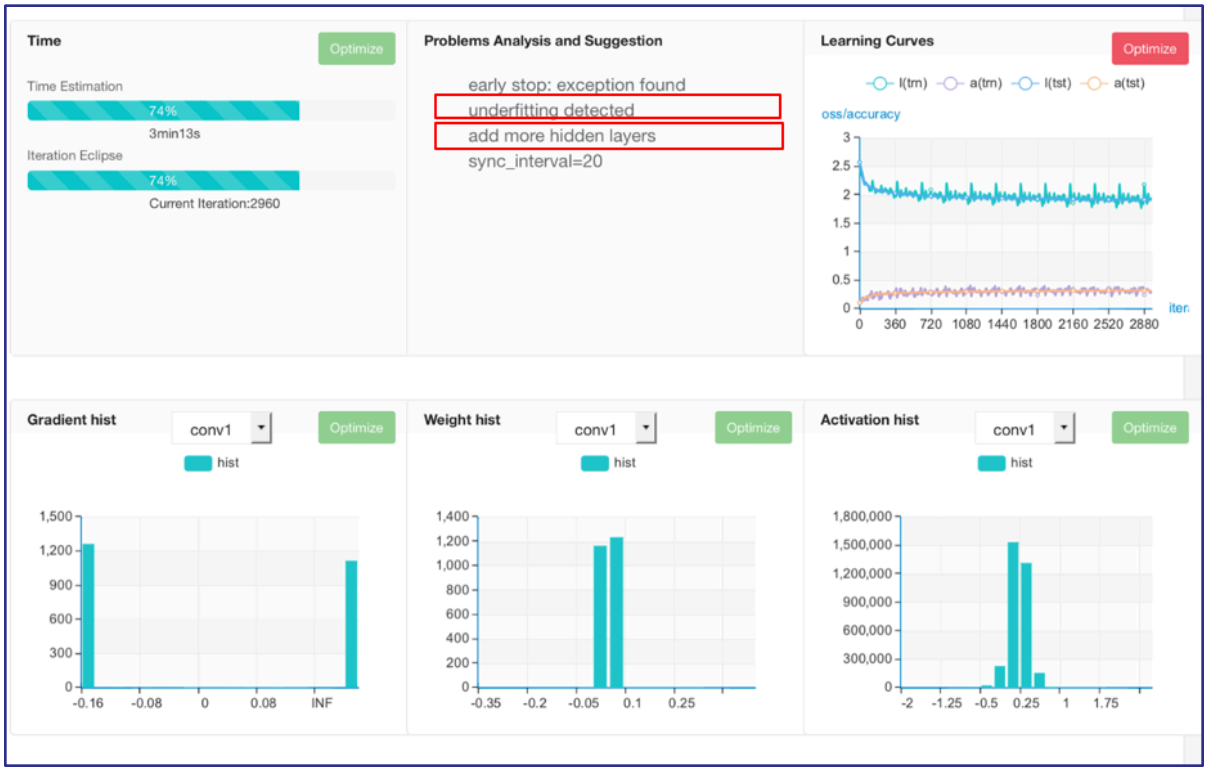

- The nan of weights is disappeared, and the overflow error is away.
- The learning curves find the problem that the curve is converged but the accuracy is still so low about 0.3 on both training and testing.
- The MAO suggests user to stop the training early because underfitting detected, and give the user advisor to add more hidden layers.

© 2018 IBM Corporation

GPU Acceleration

GPU-Enable Spark Apps

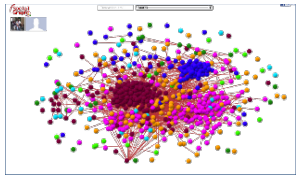

Graph Analytics
Security, Fraud Detection
Social Network Analytics
GraphX

Spark MLlib

Machine Learning
Predictive analytics,
Logistic regression, ALS
Kmeans, etc

…but Spark apps are CPU intensive

GPU Accelerated Deep Learning Apps

Financial Risk Analytics
Market simulation
Credit risk, home-grown, apps
from Murex, Misys

Video/Speech Analytics
Object Recognition
Dialog

…and most GPU-accelerated apps will need to handle more data and bigger models

© 2018 IBM Corporation

Elastic Training – Quality of Service
Transparent, Elastic, Resilient

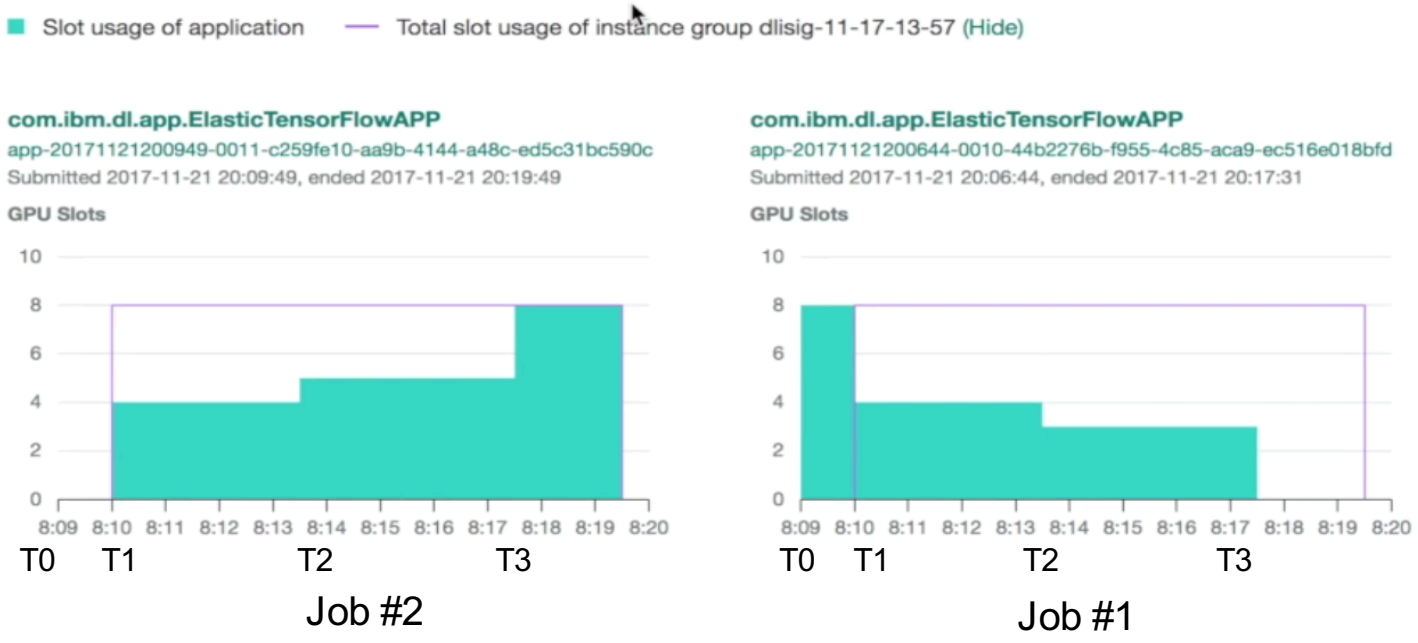

T0: Job 1 start

T1: Job 2 start
Preemption

T2: Priority change

T3: Job 1 finish

© 2018 IBM Corporation

Spark & DL framework Life Cycle Management
Flexible Spark & Notebook Version Management

- IBM Certified Apache Spark/Python/PowerAI versions
  - Spark 1.4.1, 1.5.2, 1.6.0, 1.6.1, 2.0.1
- Notebook versions are published in dev-ops site
  - Out of box: Zeppelin 0.5.6
  - Jupyter (iPython)
- Setup Instance Group with desired combination

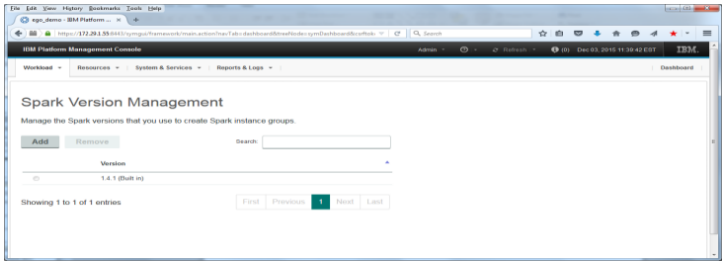

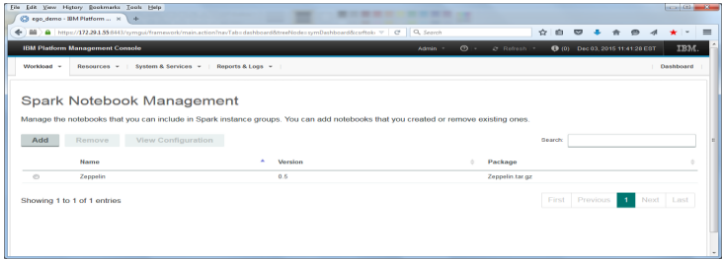

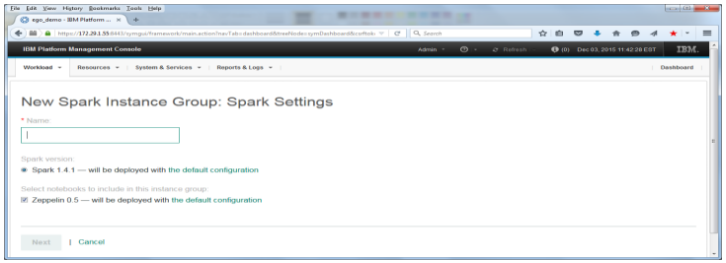

Monitor & Reporting with ELK

Integrated Elastic Search, Logstash, Kibana for customizable monitoring

Built-in monitoring Metrics

- Cross Spark Instance Groups
- Cross Spark Applications within Spark Instance Group
- Within Spark Application

Built-in monitoring inside Zeppelin Notebook

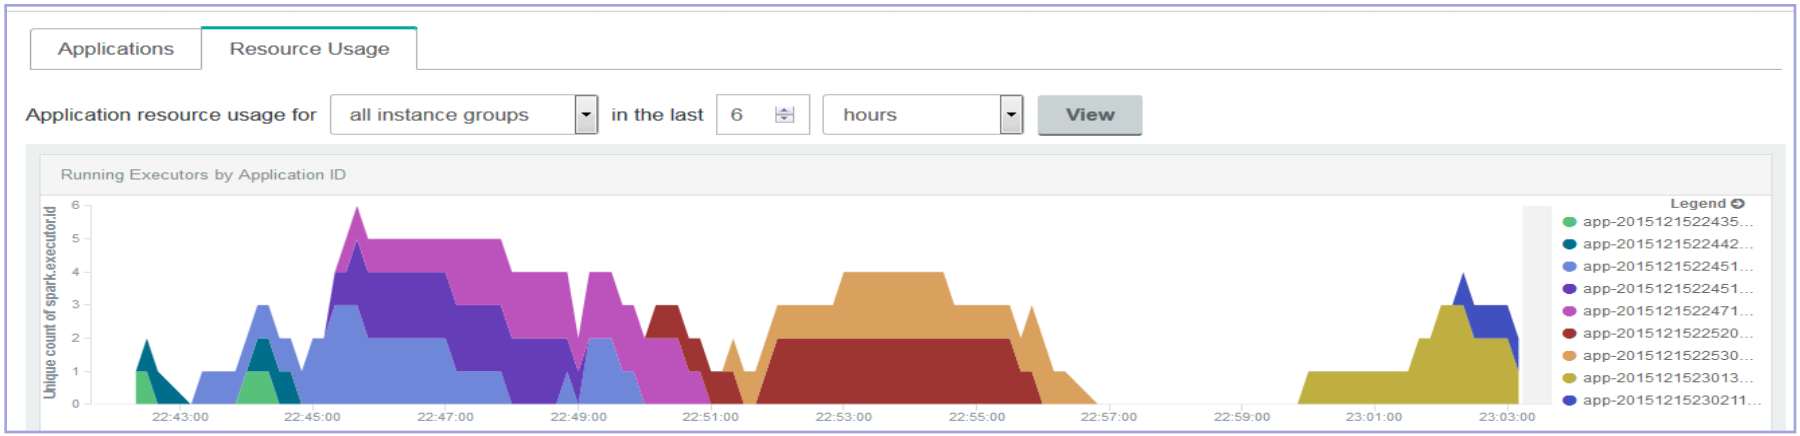

Production Proven with Enterprise Security

- Production proven support for Kerberos, Siteminder, AD/LDAP and OS auth

- Allow different LOB to define production execution users

- Fine grained access control
- ACL/role based control (RBAC)
- Spark binary life cycle, notebook updates, deployments, resource plan, reporting, monitoring, log retrieval and execution

- SSL & daemon authentication
- Data/ Storage encryption

© 2018 IBM Corporation

Multitenant Shared Services

- **Physical view** Single cluster with Spectrum Conductor on each Server
- **Logical view** Each group has **their own Spark/ML/DL cluster,** isolated, protected & secured. Dynamically managed by SLA

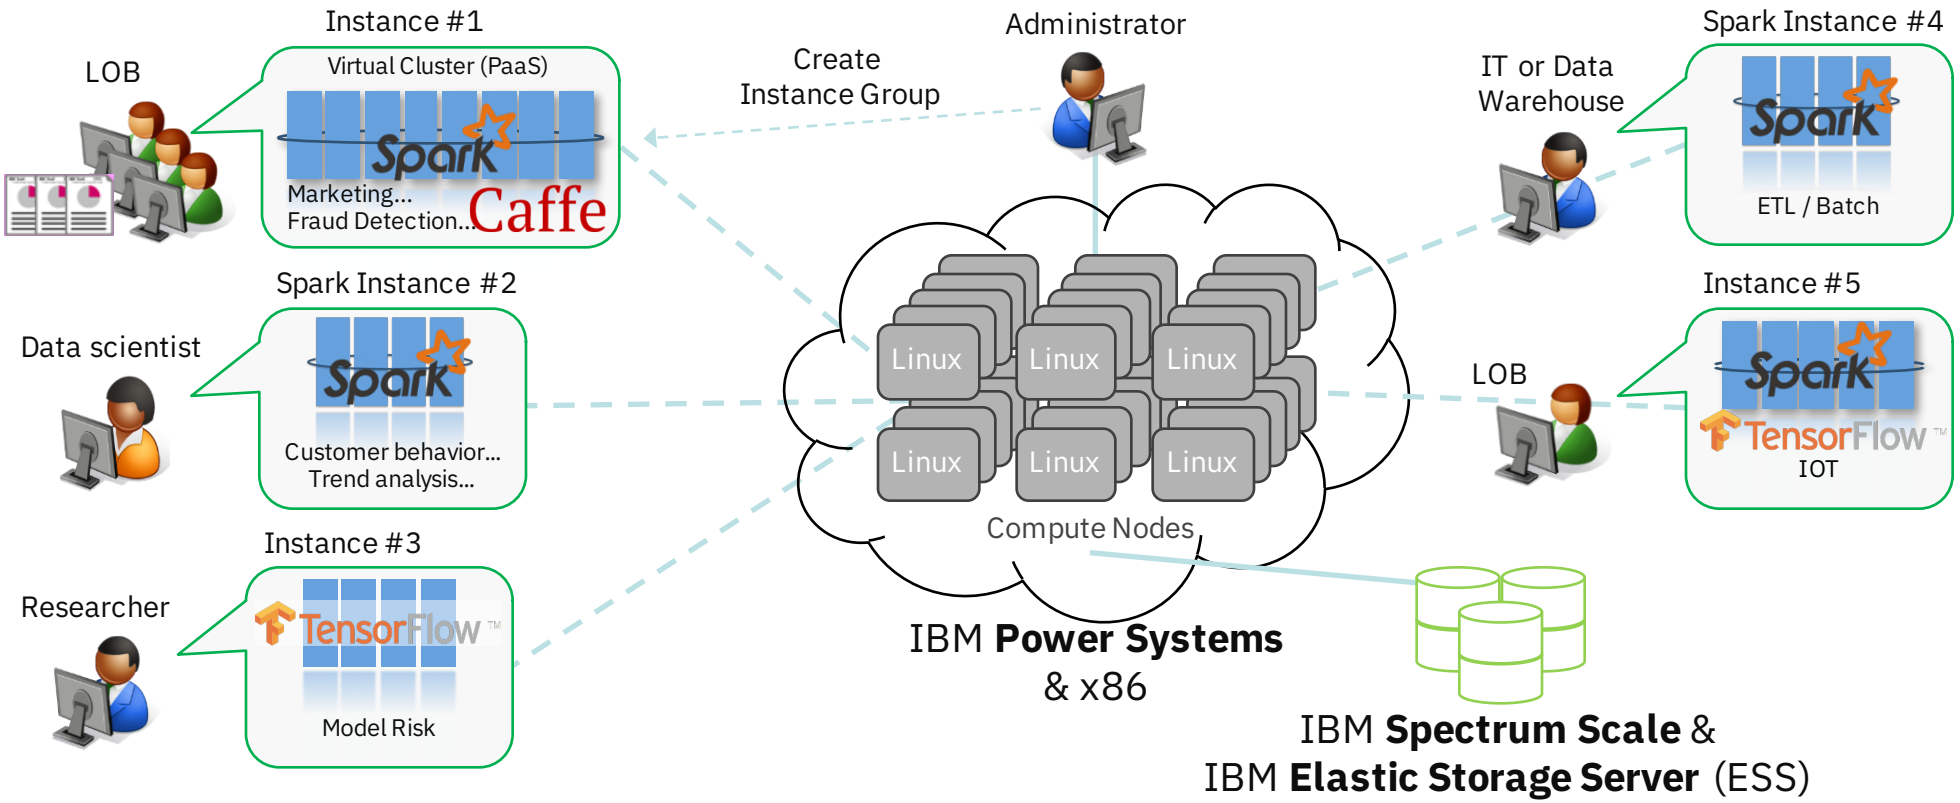

© 2018 IBM Corporation

Major Bank – AI Deep Learning Grid-as-a-Service

Major Bank

- The bank had experience using POWER processors plus NVIDIA GPUs through a POC. They really liked the power and performance delivered. They wanted to put in place an environment that could service their 35+ AI research scientists. Initially and for security reasons, the cluster was going to be deployed in an isolated research facility, when the client assessed the security capabilities of the Spectrum Computing solution they opted to deploy in their production data center running real client data.

Solution

- Deep Learning Cluster for neural network training - The Deep Learning Cluster provides the bank with a premier platform for Distributed Deep Learning to support their ground-breaking AI research initiatives.

What was sold

- 11 x 8335 GTG compute nodes
- 1 x 8001-12C login nodes
- 1 x ESS GS2S - 730 TB SSD
- 1 Gb, 10 Gb, IB switches
- IBM Spectrum Conductor
- IBM Spectrum Conductor Deep Learning Impact

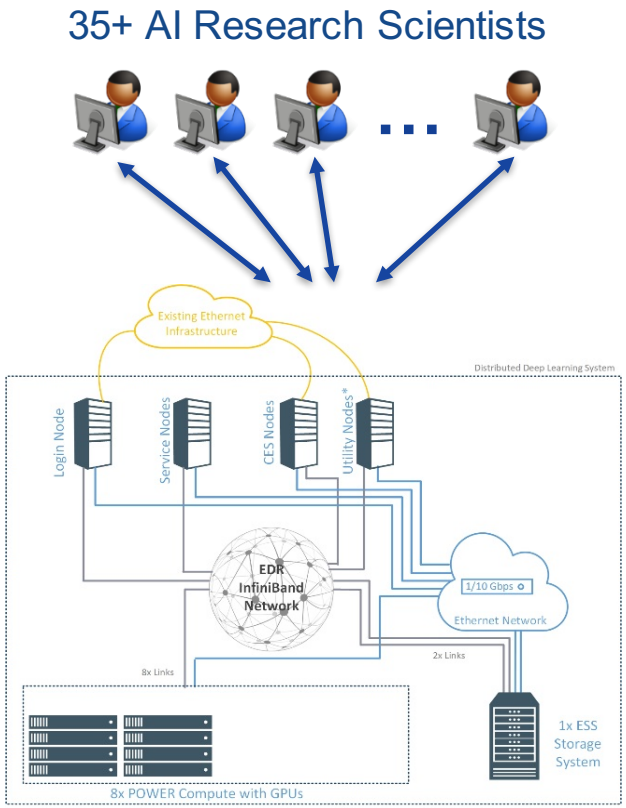

Retail Bank – Masked Face Detection

Challenge

- Unattended ATMs become target of crimes
- Masked face detection could help recognize potential criminal actions and then trigger alarm or limit function
- Traditional pattern identification algorithms are not so effective to resolve such many diversified and changing possibilities

Benefits

- Masked Face detection for ATM proved production ready and significantly improve the security of ATMs and banks

Solution

- A cluster of IBM servers with NVIDIA GPU and NVLink
- PowerAI stack for the deep learning framework, running over Spectrum Conductor with Spark and Spectrum Scale
- Modeling services of face occlusion detection thru Caffe
- Sampled 150+ face occlusion videos to generate 1500+ images as the training and testing dataset
- Real time video recording and auto-detect face occlusion, support multiple ways of recognition simultaneously

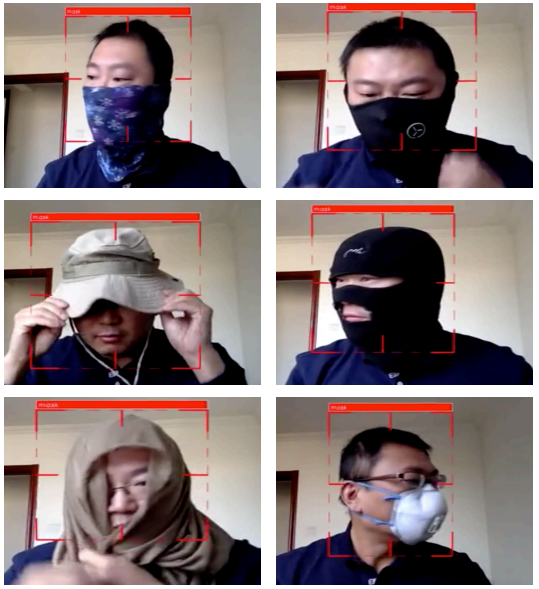

Futures Index Prediction

Business Requirement

- Predict the next day futures index status
- Previous 1000 minutes trade record as input
- P&L (Profit Loss), AR (Annualized Revenue), SR (Sharp Ratio), MDD (Max Drawdown) measure the accuracy

Challenges

- Traditional data format transfer to the deep learning required vector
- Small input may cause gradient vanish
- Model accuracy

Our Value

- Handle the ETL data transform to training data
- Support hyper-parameter search for RNN model
- Visualize the training process and provide optimizer solution

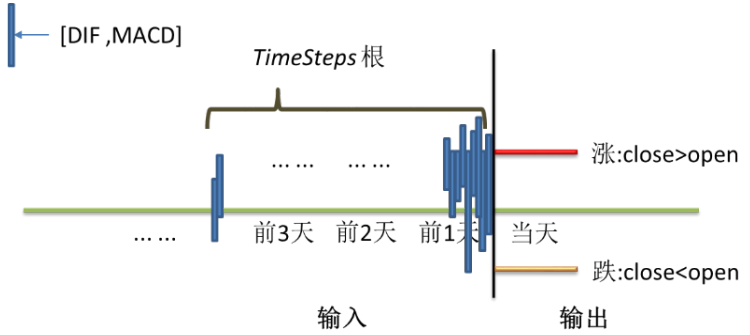

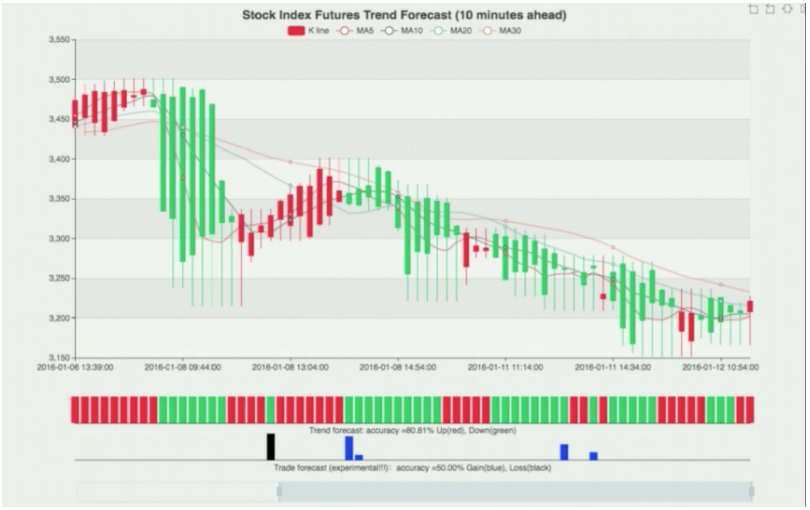

Tumor Proliferation Assessment

Business Requirement

- Predict tumor proliferation
- Mark the images from electron-microscope
- Support high resolution ratio
- Above 75% accuracy

Challenges

- Classification vs Object Detection
- Choose the correct model: CNN vs Faster R-CNN vs R-CNN
- Parallel training to improve the performance and keep convergence

Our Value

- Object detection
- Data preparation
- Change picture size and label of coordinate at the same time
- Show the ground truth without changing the original image

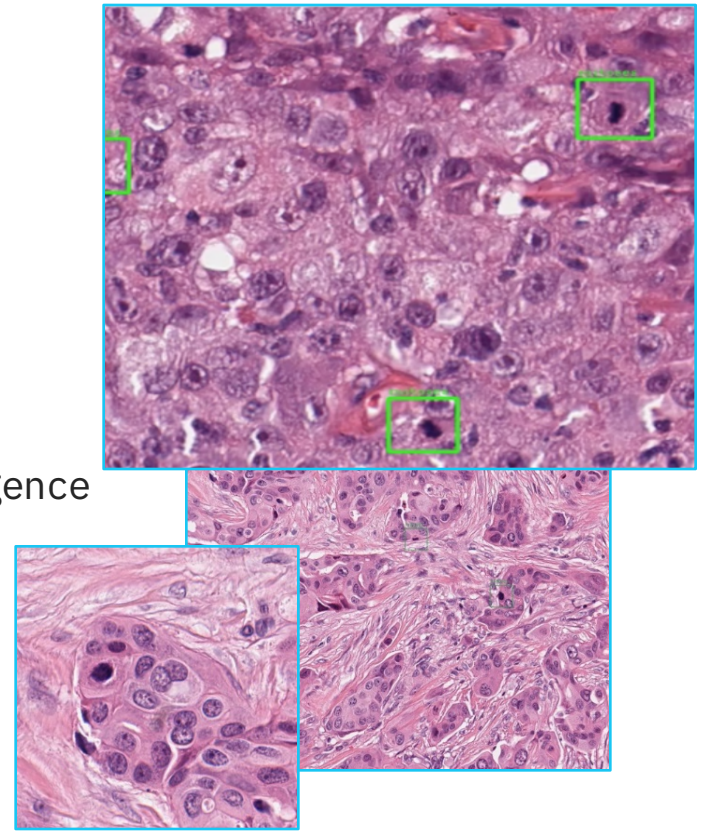

Solar Panel Quality Classification
Industrial Manufacturing

Business Requirement

- Automatic solar panel quality classification and interact with the product line
- Picture sizes are not same and the resolution is around 3k*3k which is not good for direct training

Challenge

- Shorten the training duration
- Data distribution is not uniform

Value

- Bottle-neck training achieve 86% accuracy (100G data is around 4 hours)
- VGG19 + Caffe fully training achieve 89% accuracy (Full training around 4 hours)
- Combine with machine learning method to improve the yield rate

Q0

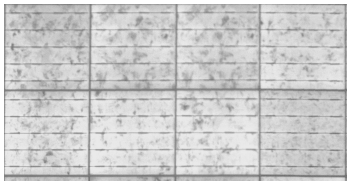

Q1

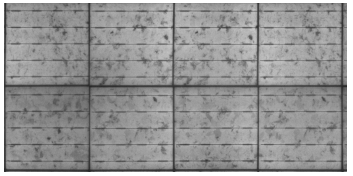

Q2

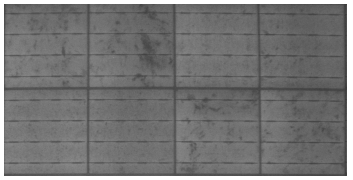

NG

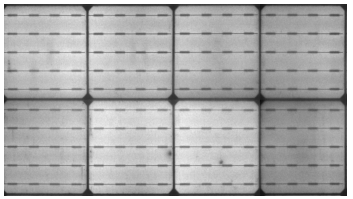

Conclusions

Enterprises have an expectation of cloud like, multi-tenant AI experience

- End users are getting that type of experience from public clouds – need to have similar capabilities internally
- Deliver flexibility of a personal/dedicated environment for users while managing as a cost effective shared service

Performance and SLA are key due to wide range and variety of workloads

- Sophisticated resource and workload management to utilize the infrastructure in the best way across a wide variety of Spark use cases and applications
- Ensure SLA's while managing cost and TCO of infrastructure

'Future proof' the environment so that it can evolve and scale in multiple dimensions

- Versions of Frameworks, notebooks, libraries, scale of underlying infrastructure, new data sources, resource types and application frameworks