Artificial Intelligence: Infrastructure Matters

Data Server Day 2019
May 9th, 2019

Kelly Schlamb
Executive IT Specialist,
IBM Cognitive Systems

kschlamb@ca.ibm.com
@KSchlamb

"*AI promises to unlock the value of your data – in totally new ways*"

MITSloan

270%

increase in Enterprise AI adoption over the past 4 years

91%

expect new business value from AI implementations in next 5 years

$13.5B

2019 market for AI software platforms and AI applications

47%

CAGR over next 5 years

Software is largest & fastest growing AI tech category

Common AI Use Cases *(but the possibilities are limitless)*

Hyper or Radical Personalization

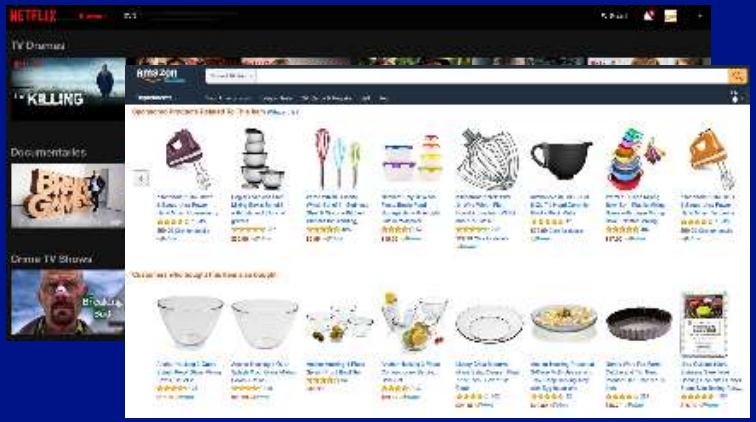

Price and Product Optimization

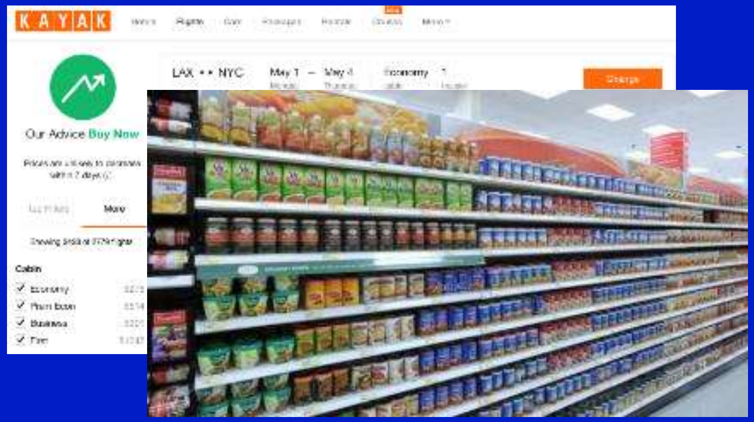

Predictions and Classifications

Resource Allocation and Strategic Planning

Discover Patterns, Anomalies and Trends

Natural Language Processing

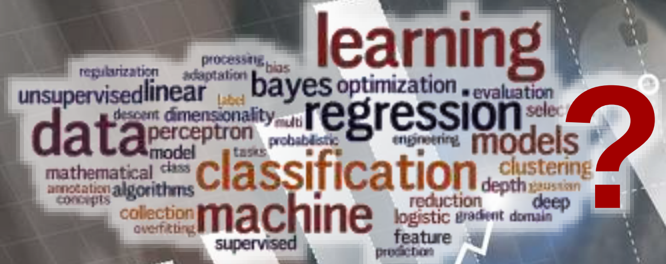

Data Science

The practice of applying various scientific and statistical methods, algorithms, approaches and processes…

using programming languages and software frameworks…

to extract knowledge, insights and recommendations from data…

and deliver them to business users and consumers in consumable applications.

Artificial Intelligence – System that mimics human intelligence (very broad ranging applications, methods and supporting technologies)

Machine Learning – Enable machines to learn from data and make accurate predictions, without being explicitly programmed to do so

Deep Learning – Perform complex tasks (like speech and image recognition) by exposing multilayered neural networks to vast amounts of data

Infrastructure [in-fruh-struhk-cher]

noun

The system of hardware, software, facilities and service components that support the delivery of business systems and IT-enabled processes.

NETFLIX

didn’t put BLOCKBUSTER out of business, ridiculous late fees did

MUSIC

didn’t force change on the music industry ,

being forced to buy full length albums did

isn’t hurting the TAXI business, limited access & fare control are


airbnb
isn’t hurting the HOTEL business, limited availability & pricing are

TECHNOLOGY itself
is not the disruptor

Not being CUSTOMER-CENTRIC is the biggest threat to any business

Technology is an ENABLER

AI is a Science

Machine learning, deep learning, neural networks, predictive analytics, prescriptive analytics, etc. are all technologies that enable AI.

AI is an Enabler

It helps domain experts make faster decisions to drive business outcomes more quickly. AI doesn’t replace experts, it enables them.

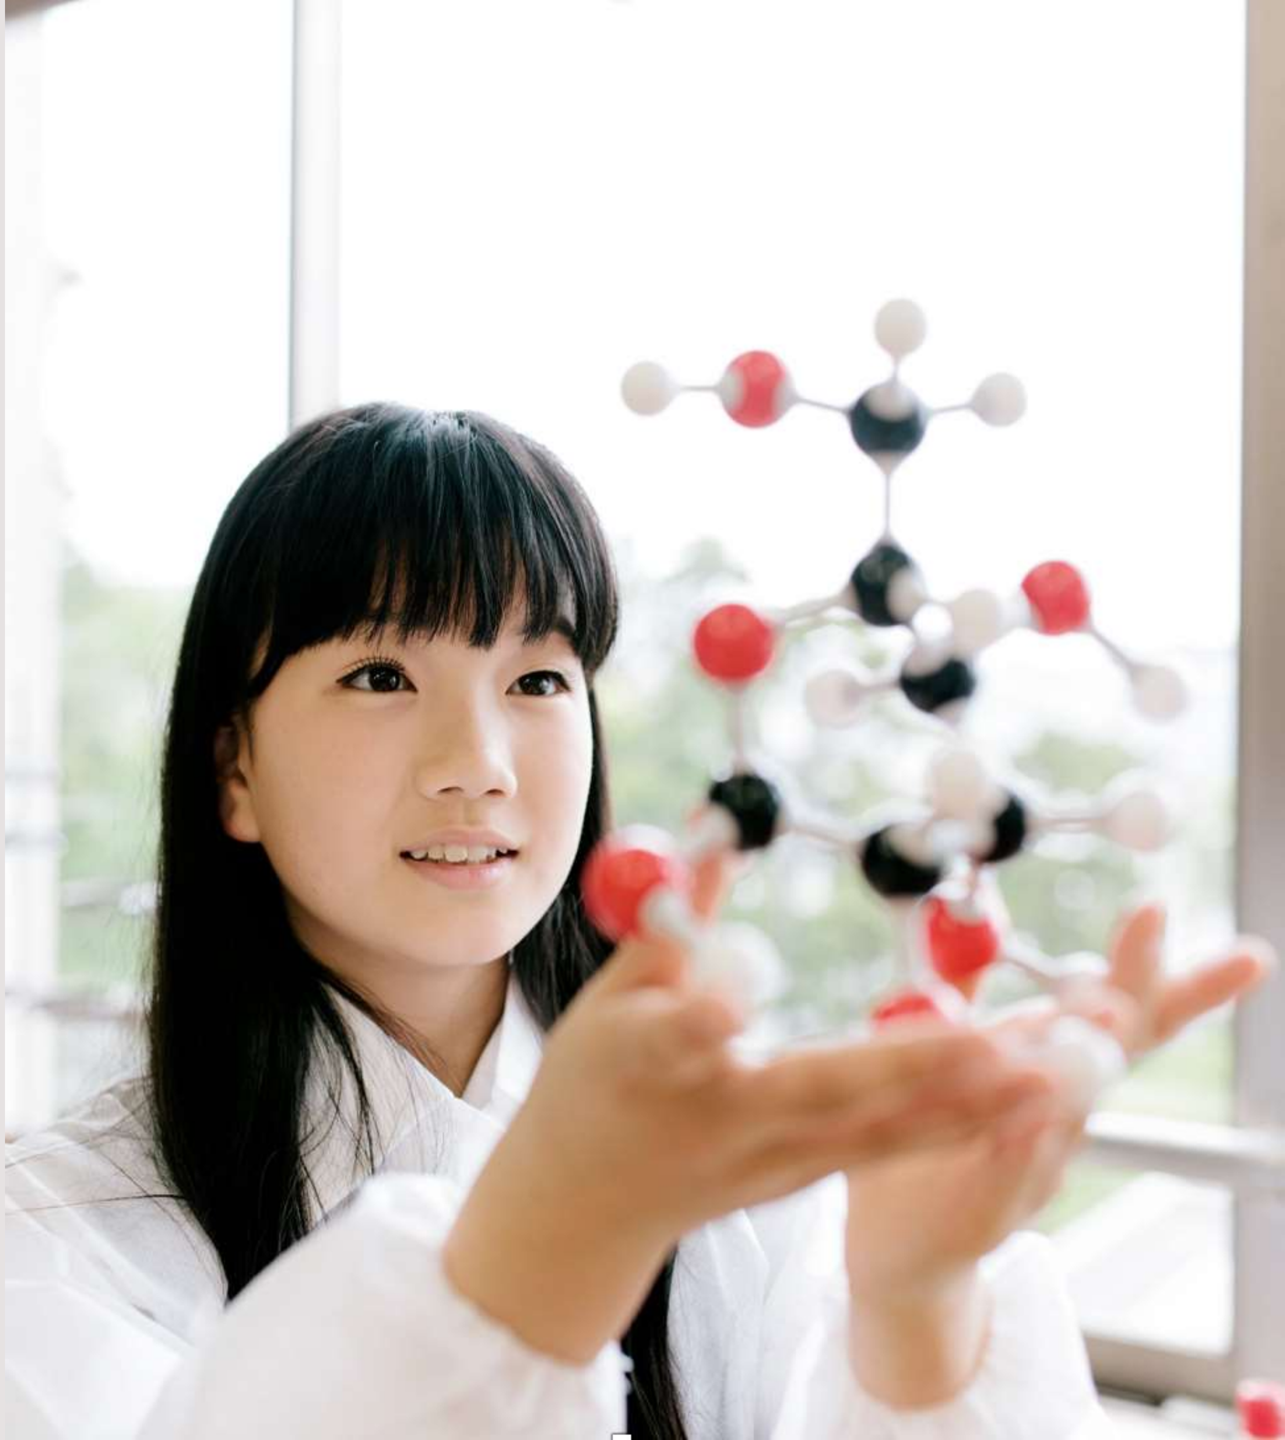

99% say their firms are trying to become insights-driven,

but only one-third report succeeding

NewVantage Partners,
"Big Data Executive Survey 2018 Executive Summary of Findings"

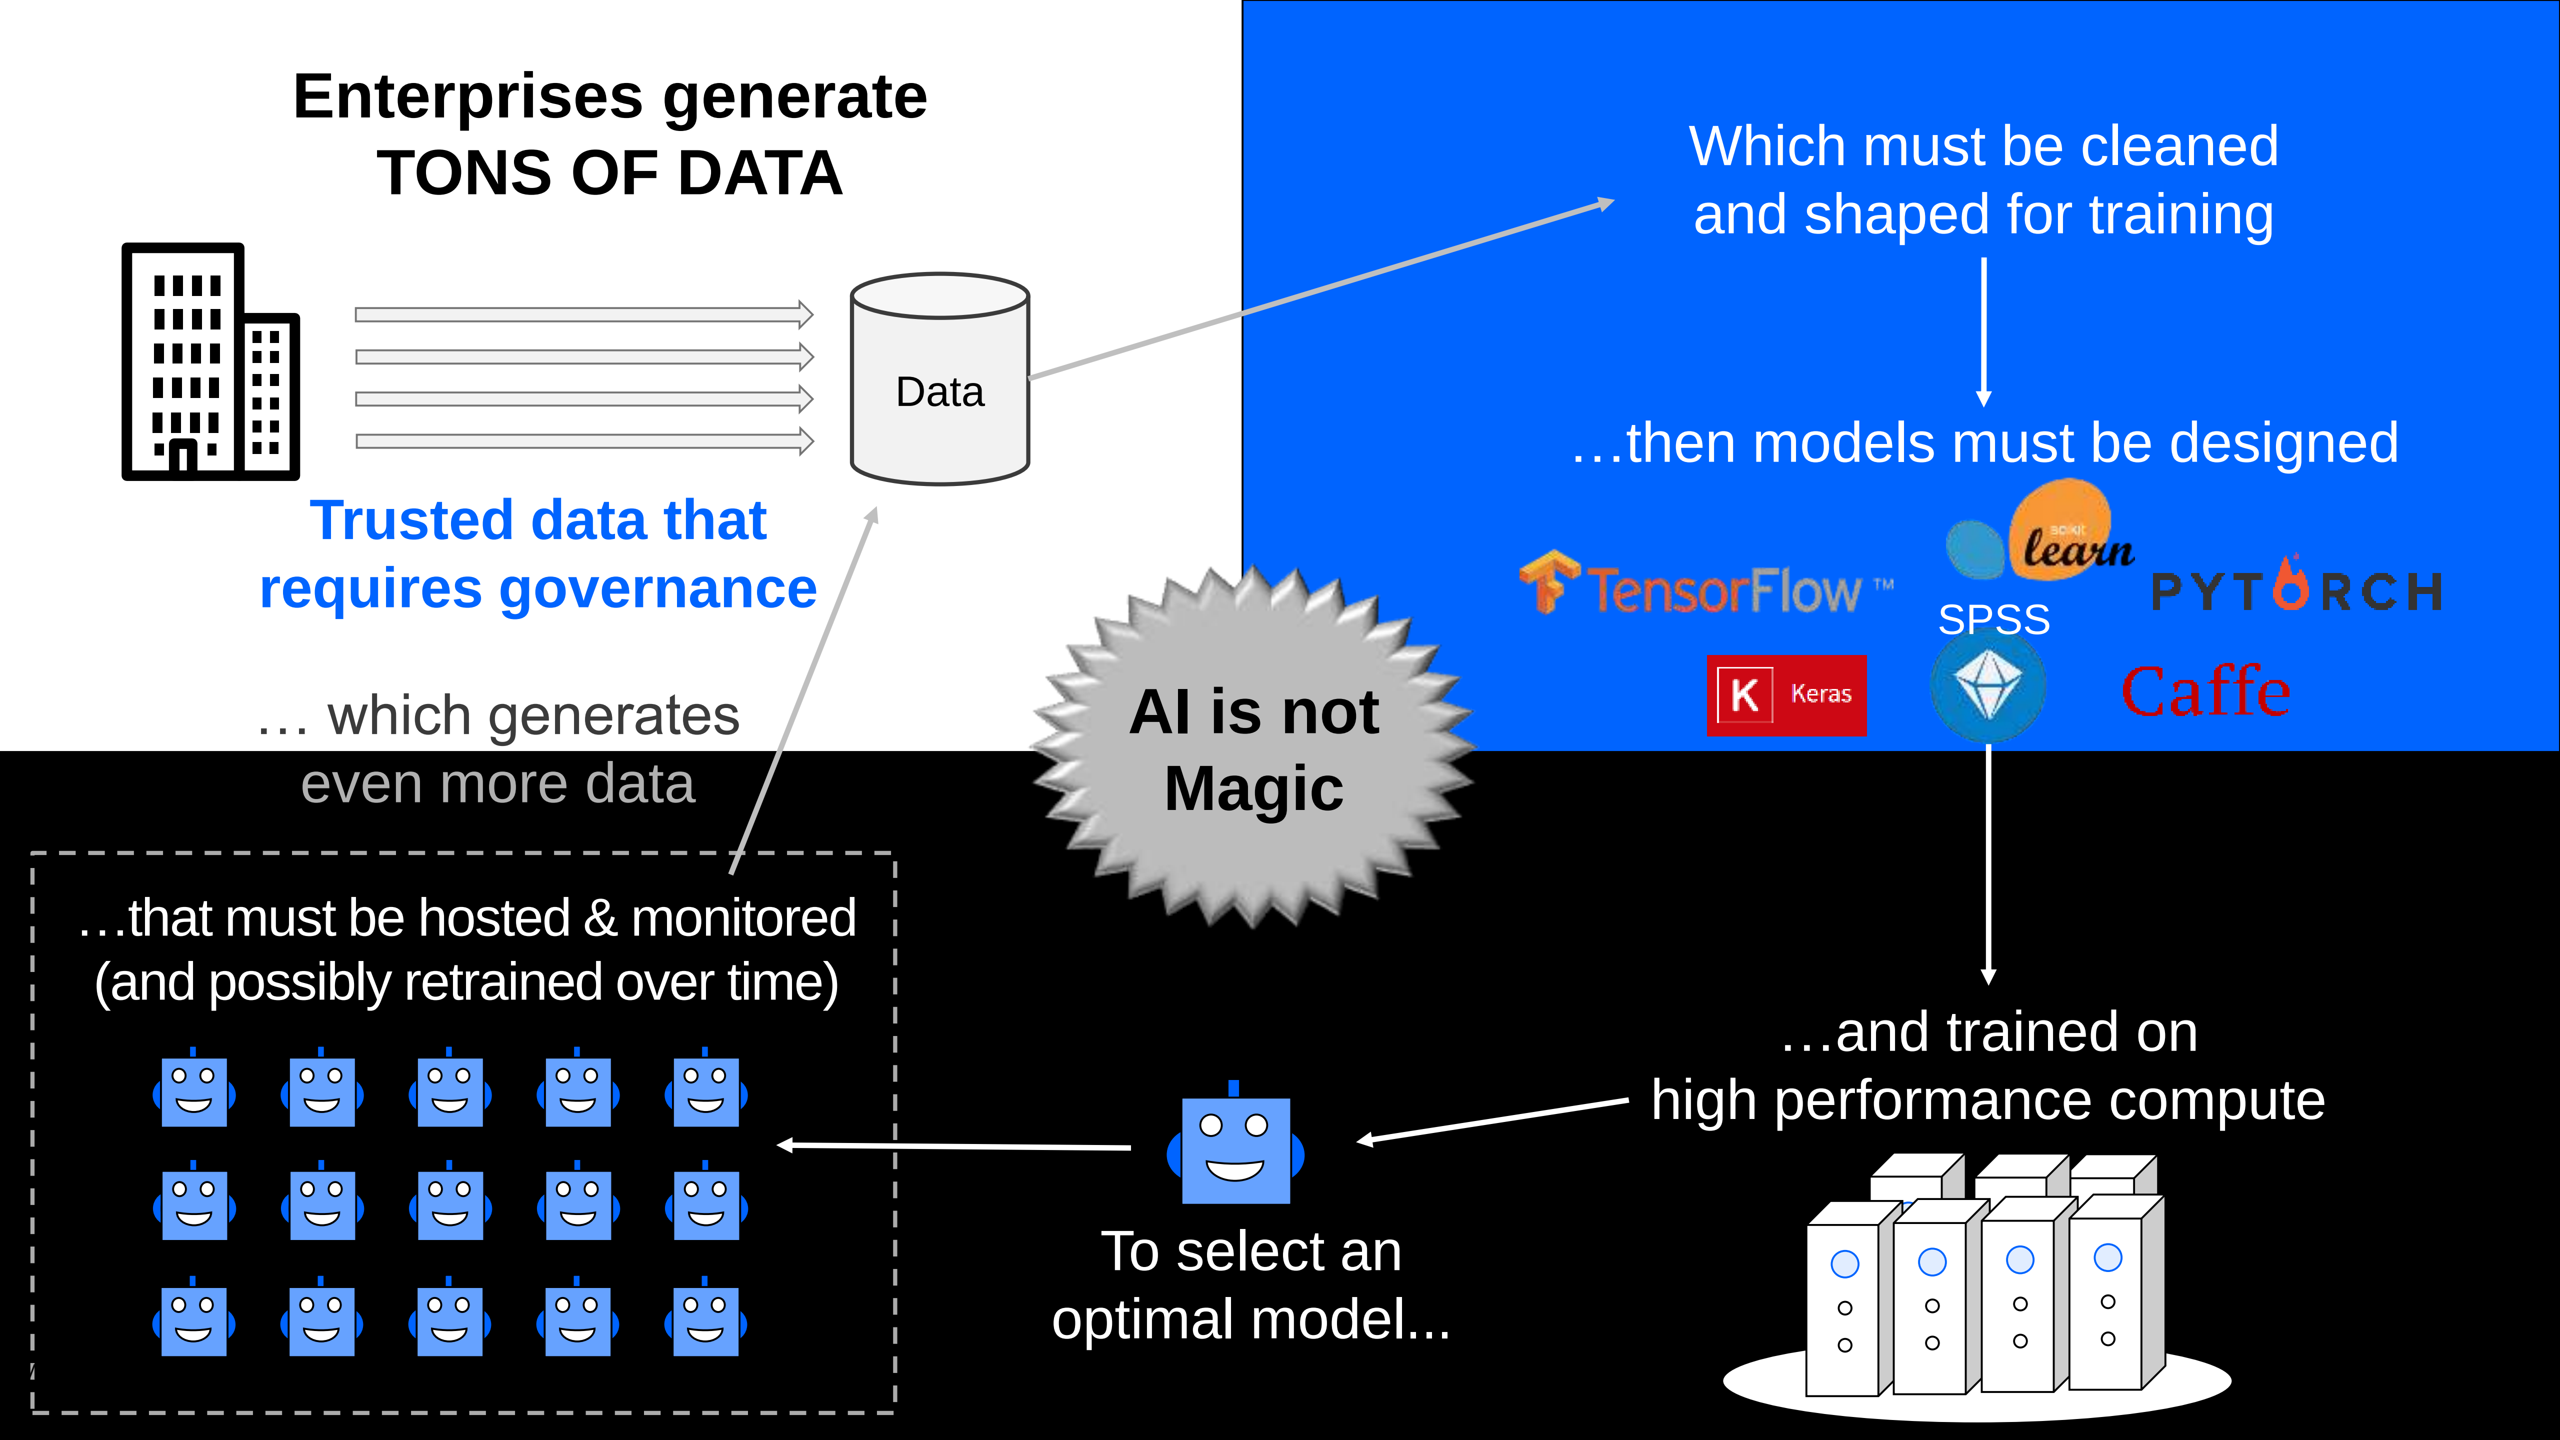
Enterprises generate
TONS OF DATA
Data
Trusted data that
requires governance
Which must be cleaned
and shaped for training
…then models must be designed
TensorFlow
learn
PYTORCH
SPSS
Keras
Caffe
AI is not
Magic
…and trained on
high performance compute
To select an
optimal model...
…that must be hosted & monitored
(and possibly retrained over time)
… which generates
even more data

Data Science and AI is a Team Sport

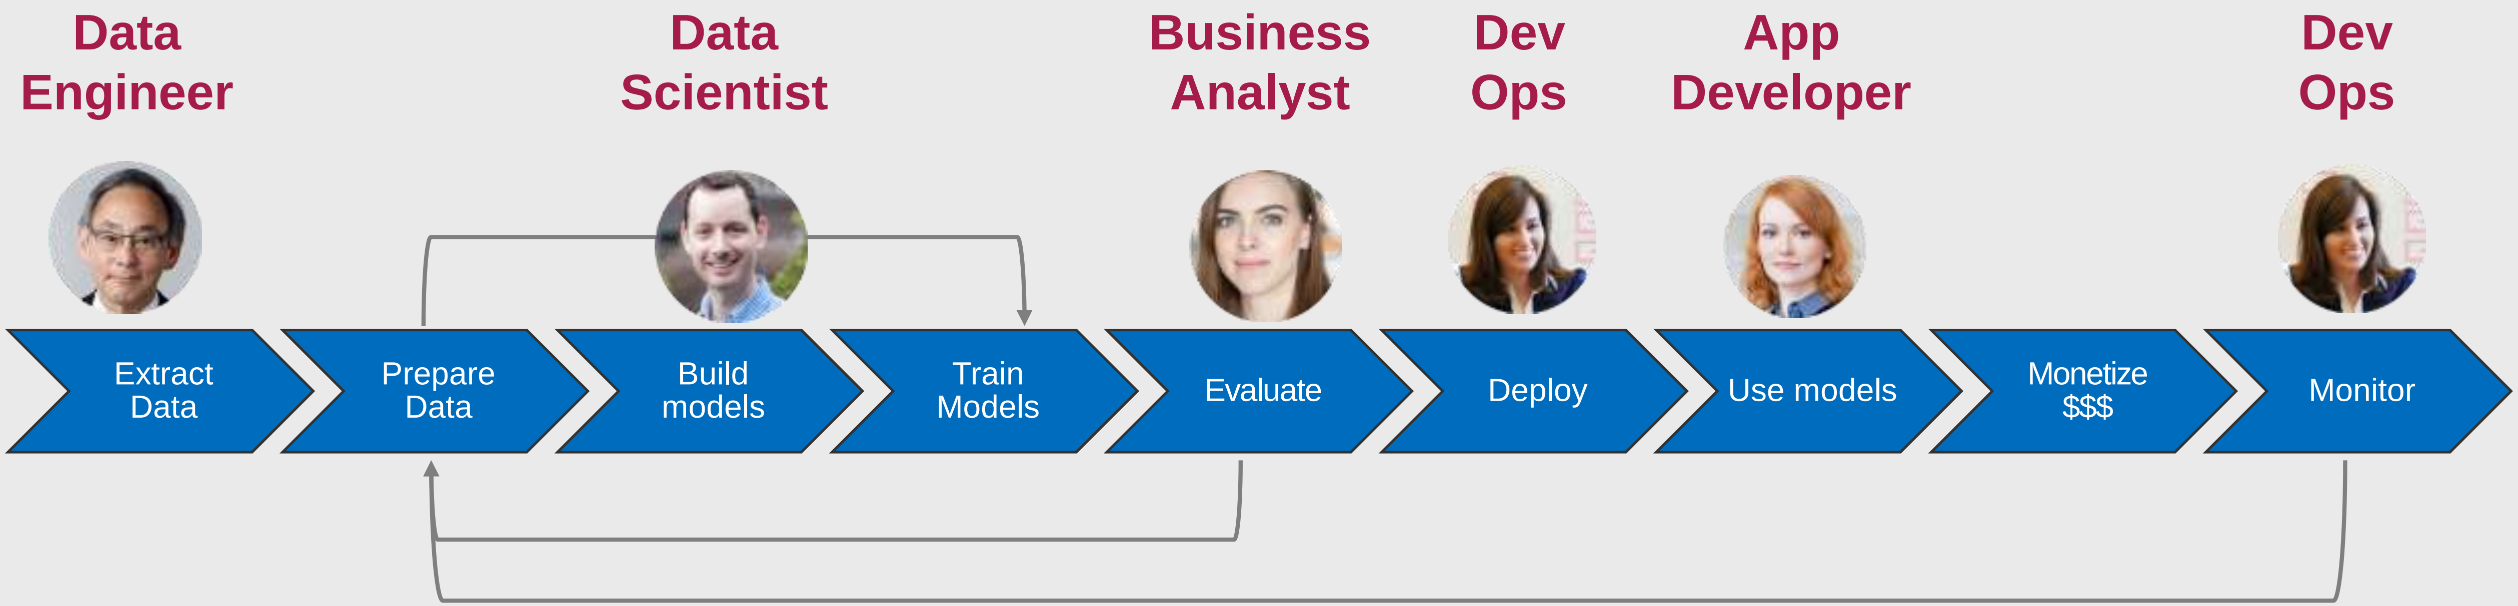

Building cognitive applications using ML and DL
requires multiple skillsets and collaboration

What are the impediments of progress?

1. Data Explosion
2. Self-Service Bottleneck
3. Lack of Compliance & Visibility
4. Lack of Skills, Availability & Collaboration
5. Inability to Monetize Insights

Assembling the building blocks of a data and analytics platform is complex, risky and time-consuming

CAPTURE | **SECURE** | ← **ANALYZE** → | **OPERATIONALIZE**

Collect, Connect, and Access Data

Connect and **discover** content from multiple data sources across your organization.

Provision databases and **virtualized** data access.

Govern, Search, and Find Data

Grant user access levels and **enforce business policies**.

Index for search, visualize consumers and producers of assets with **lineage, metrics, and quality profiles**.

Find data and analytics assets in the **Enterprise Catalog**.

Understand and Prepare Data for Analysis

Understand, cleanse and prepare your data to create data preparation pipelines visually.

Use popular open source libraries to prepare structured and unstructured data.

Build Descriptive, Predictive, and Prescriptive Models

Create Machine Learning, Deep Learning, Optimization, and other advanced mathematical models.

Tools to design your models **programmatically** or **visually.**

Train at scale with support for **distributed** compute and **GPUs**.

Model Management and Deployment

Manage your models across **dev, test, staging, and prod**.

Deploy your models and scale automatically for online, batch or streaming use cases with **SLAs**.

Monitor model **performance** and **automatically trigger retraining** and redeployment as rolling upgrades.

Create Analytics Applications

Incorporate **trusted and governed models** into applications, dashboards, and operational systems.

There is no AI without an IA

(information architecture)

81%

do not understand the data required for AI

80%

of data is either inaccessible, untrusted or unanalyzed

"No amount of AI algorithmic sophistication will overcome a lack of data [architecture]… bad data is simply paralyzing"

MITSloan

The AI Ladder

A prescriptive, proven approach to accelerating the journey to AI

Ladder to AI

Infuse – Deploy trusted AI-driven business processes

Analyze - Scale insights with AI everywhere

Organize – Create a trusted analytics foundation

Collect – Make data simple and accessible

Requires a strong foundation that is built on a modern cloud-native architecture
with underlying infrastructure that is highly performant, secure, reliable and cloud-ready

Data Science Platform Market and Trends

88% of Insight Leaders use a data science platform to overcome tool sprawl and put insights into action

46% of organizations lack an integrated approach to their data science technology stack

Adoption of a single data science platform rising to 69% within next two years

Data Science Platform market estimated to grow to $128.21B by 2022 (CAGR 36.5%)

IBM's AI Portfolio

Everything you need for Enterprise AI, on any cloud

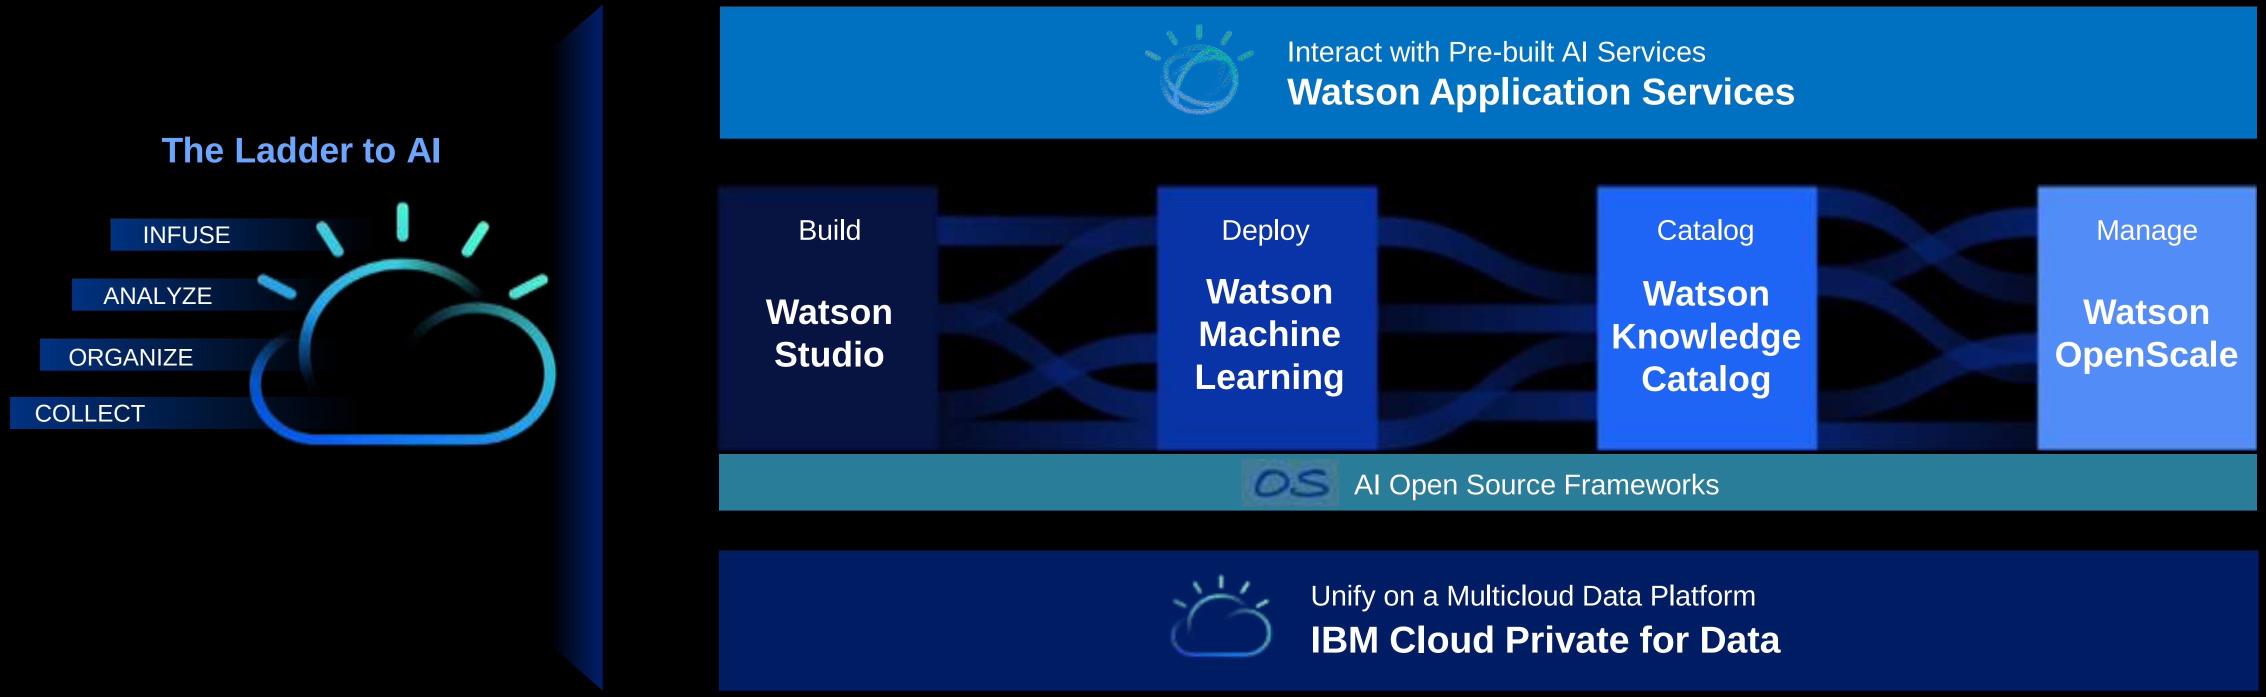

Watson Studio and Watson Machine Learning inject AI firepower into your business

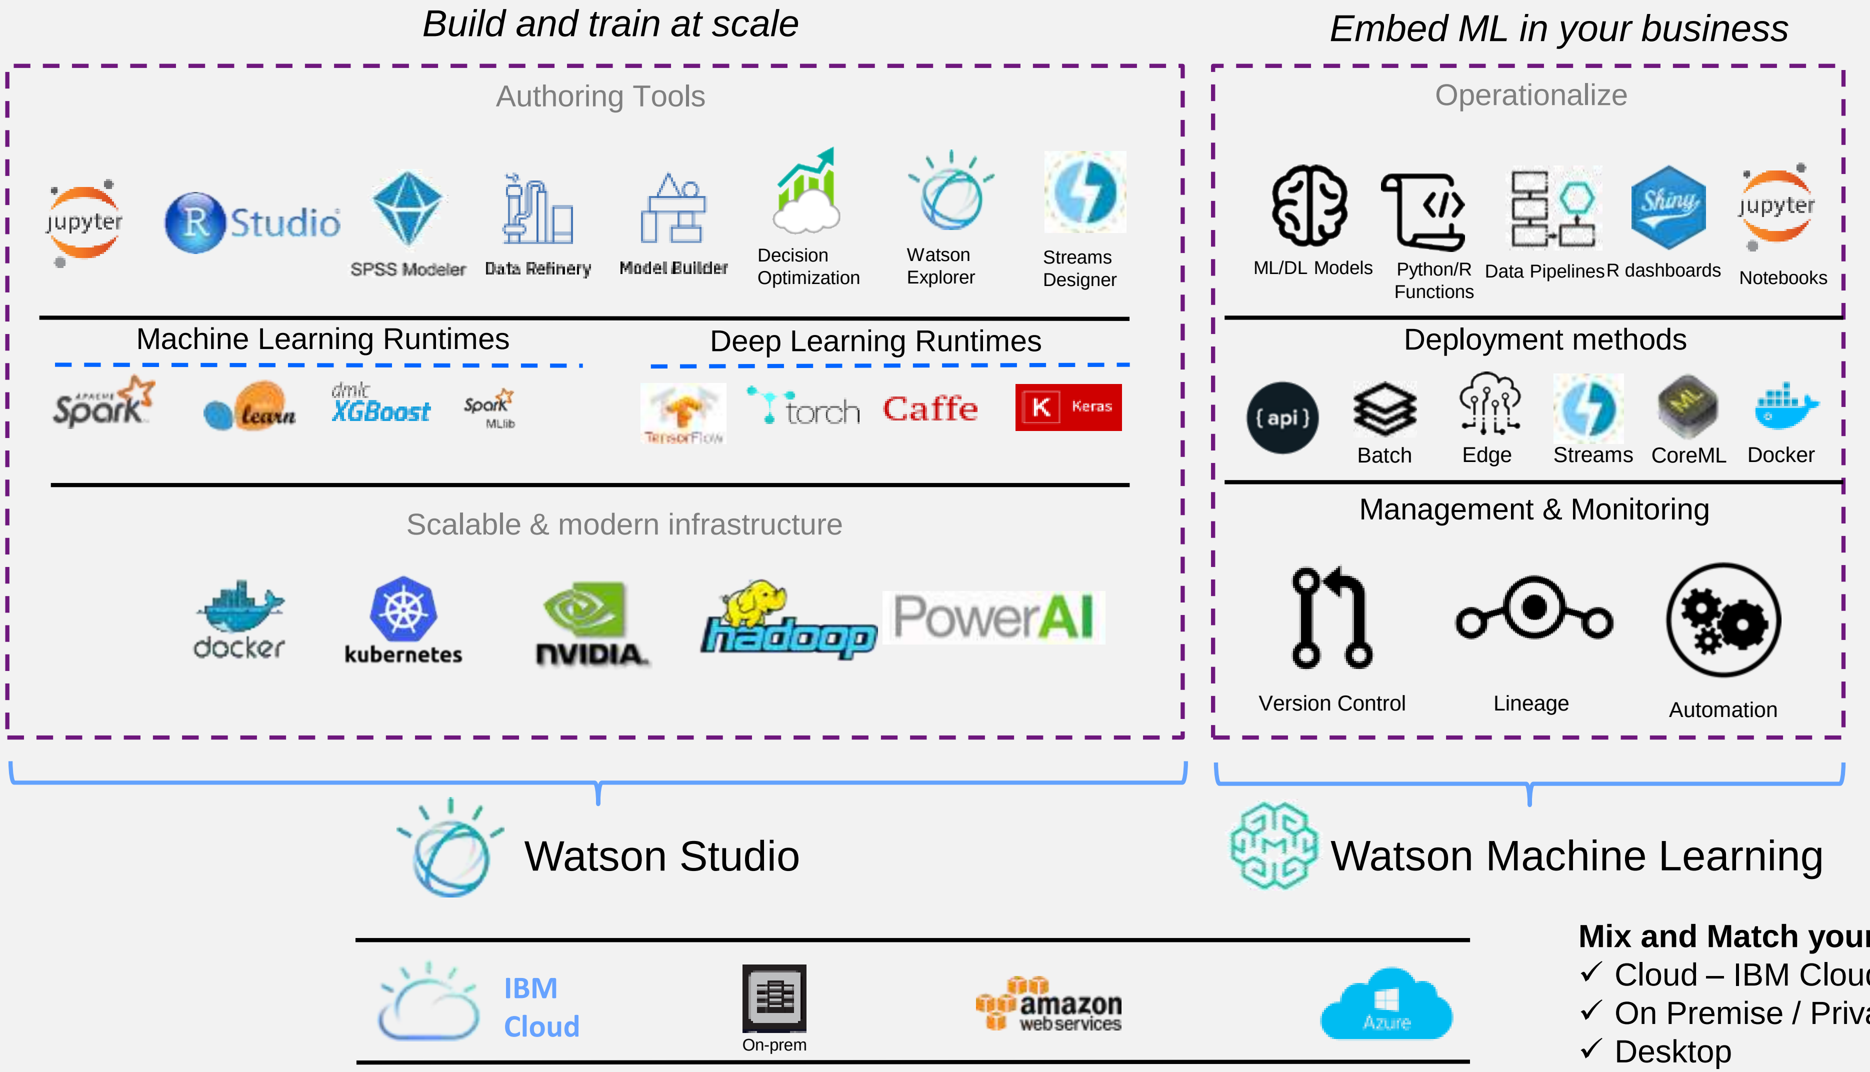

Mix and Match your deployment

✓ Cloud – IBM Cloud, Azure, AWS
✓ On Premise / Private Data center
✓ Desktop

Watson Studio

Solve your business problems by collaboratively working with data

Watson Studio

Cloud, Local, Desktop

Securely Collaborate with Your Team

Watson Studio Projects provide the best tools for collaborating with your team without sacrificing security

Authoring Tools

jupyter · RStudio · SPSS Modeler · Data Refinery · Model Builder · Decision Optimization · Streams Designer

Rapidly Build Models, Find Answers Quickly

Combine the best of open source and IBM Technologies to build and evaluate models

Machine Learning Runtimes

Apache Spark · learn · dmlc XGBoost · Spark MLlib

Deep Learning Runtimes

TensorFlow · torch · Caffe · Keras

Prototype and Develop with Cutting Edge Infrastructure and Tech

Spark, Kubernetes, Docker, NVIDIA GPUs and more power Watson Studio, so you can move from discovery to production with the latest and greatest tech, unencumbered

Scalable & Modern Infrastructure

docker · kubernetes · NVIDIA · hadoop · Power · zLinux

Watson Machine Learning

Embed Machine Learning and Deep Learning into your business

Deploy and Manage Models

Move models to production, in an easy, secure, and compliant way

Intelligent Model Operations

Embed intelligent training services, with feedback loops that constantly learn from new data, regardless where it resides

Accelerate Compute Intensive Workloads

Distribute your deep learning training and Hadoop/Spark workloads with multi-tenant job scheduling

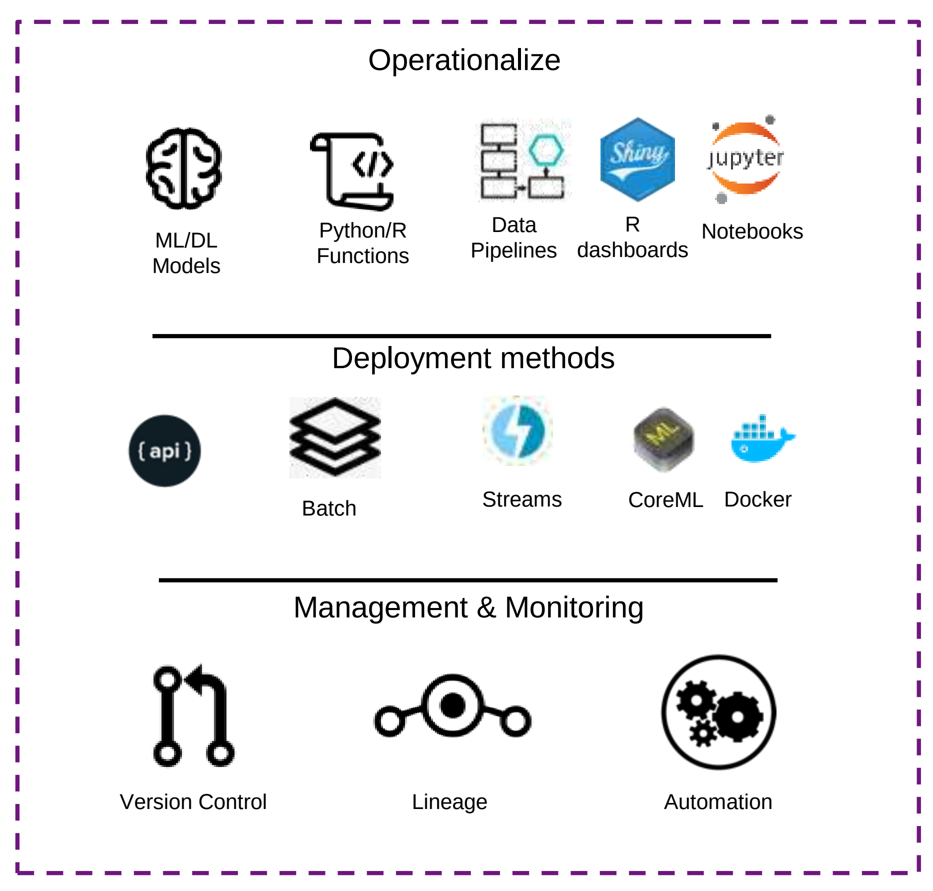

Mix and Match Watson Studio & Watson Machine Learning

across hybrid multi-cloud environments

Watson Studio

Watson Machine Learning

Across Clouds, Public and Private

Private clouds
On-premises
Desktop

Watson Studio and Watson Machine Learning Add-ons

More powerful and flexible tools built for teams

Available as Add-on

Notebooks | RStudio | Decision Optimization | Watson Explorer | SPSS Modeler | Data Refinery | AutoAI (Beta) | Cognos Analytics

Coding tools
Put Python, R, Scala and the most popular Open Source frameworks at your fingertips

Optimization
Run multidecision variable and constraint models with a model wizard

Text Analytics
Parse and extract meanings

Visual Modeling and Automation
Point and click data exploration, data preparation and modeling with productivity and automation

Visualization & Dashboard
Drive and communicate insights

Data Science Skills

High → **Low**

Data science skills are scarce –
maximize productivity with IBM Watson Studio

Watson Machine Learning Accelerator

Scaling deployment & management of ML, DL and AI models in a shared data science environment

High Performance Clustering capability for Machine Learning & Deep Learning

Accelerated analytics execution

Analytics, ML and DL acceleration for faster insights

- Shared, Optimized Apache Spark Execution
- Accelerate training through transparent, Distributed & Elastic Deep Learning
- Enterprise Ready & Proven

Train and Deploy AI at Scale

Expand to distributed models and scale out inference

- Many models, across many users, supporting multi-tenancy at scale
- Optimized scheduling of model training across CPU/GPU
- Distributed & Elastic Inference

Automate Deep Learning Workflows

Drive higher accuracy through faster training iteration

- End to End Deep Learning Workflow
- Extend to Distributed/Parallel Hyperparameter Optimization
- Interactive Monitoring, Analysis of Training

IBM Watson OpenScale

Automate and Operate AI at Scale

- **Trust and Transparency**
  - Intelligently delivers bias mitigation help
  - Provides traceability & auditability of AI predictions made in production applications
  - Tracks AI accuracy in applications
  - Explains an outcome in business terms
- **Automation**
  - Automatically detects and mitigates bias in model output, without affecting currently deployed model or outcomes
- **Open by Design**
  - Monitor and optimize models deployed on third party model serve engines
  - Deploy behind enterprise firewall or on IaaS provider

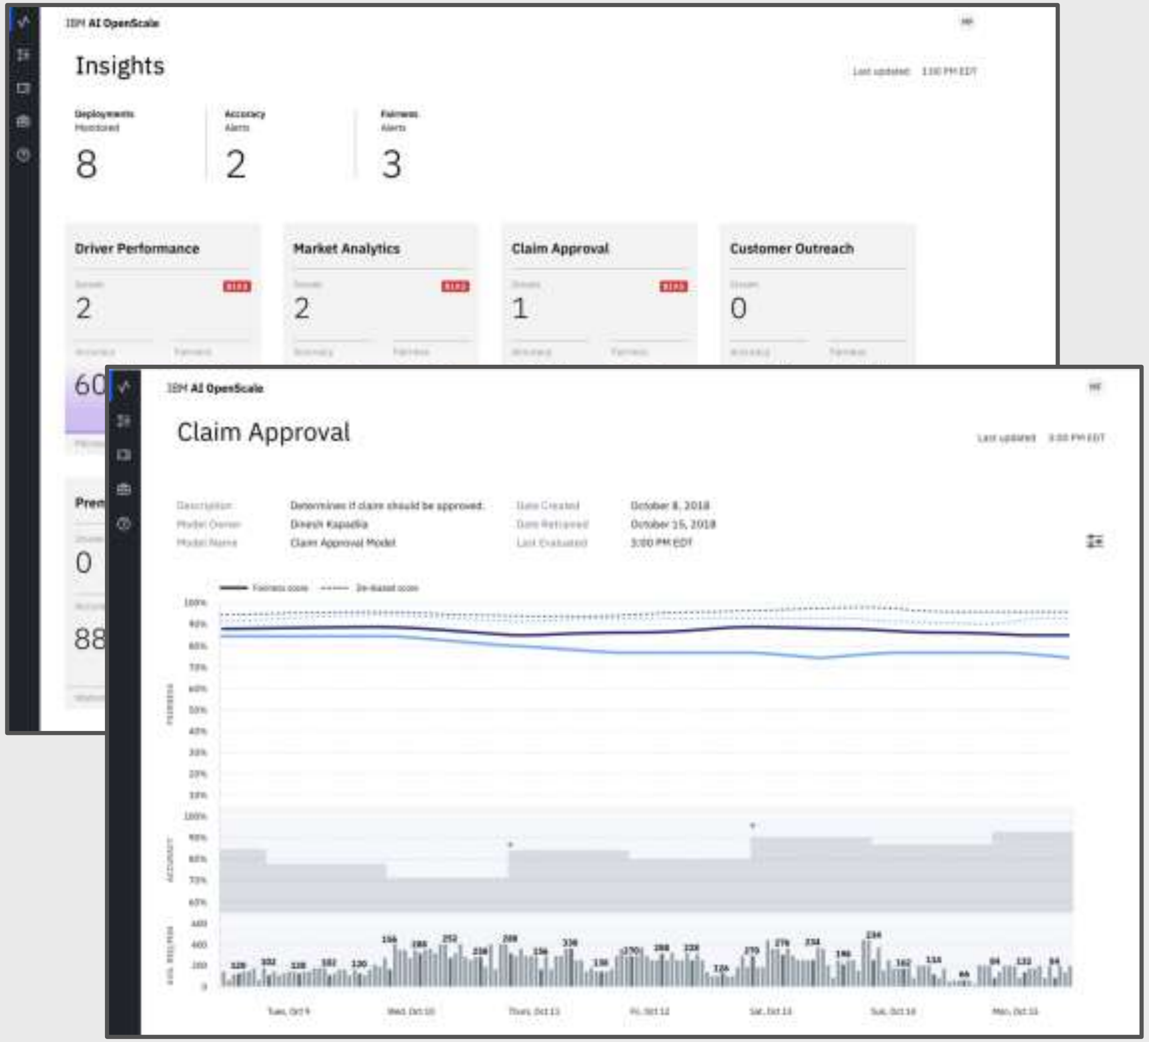

Accelerate your data science lifecycle from discovery to production

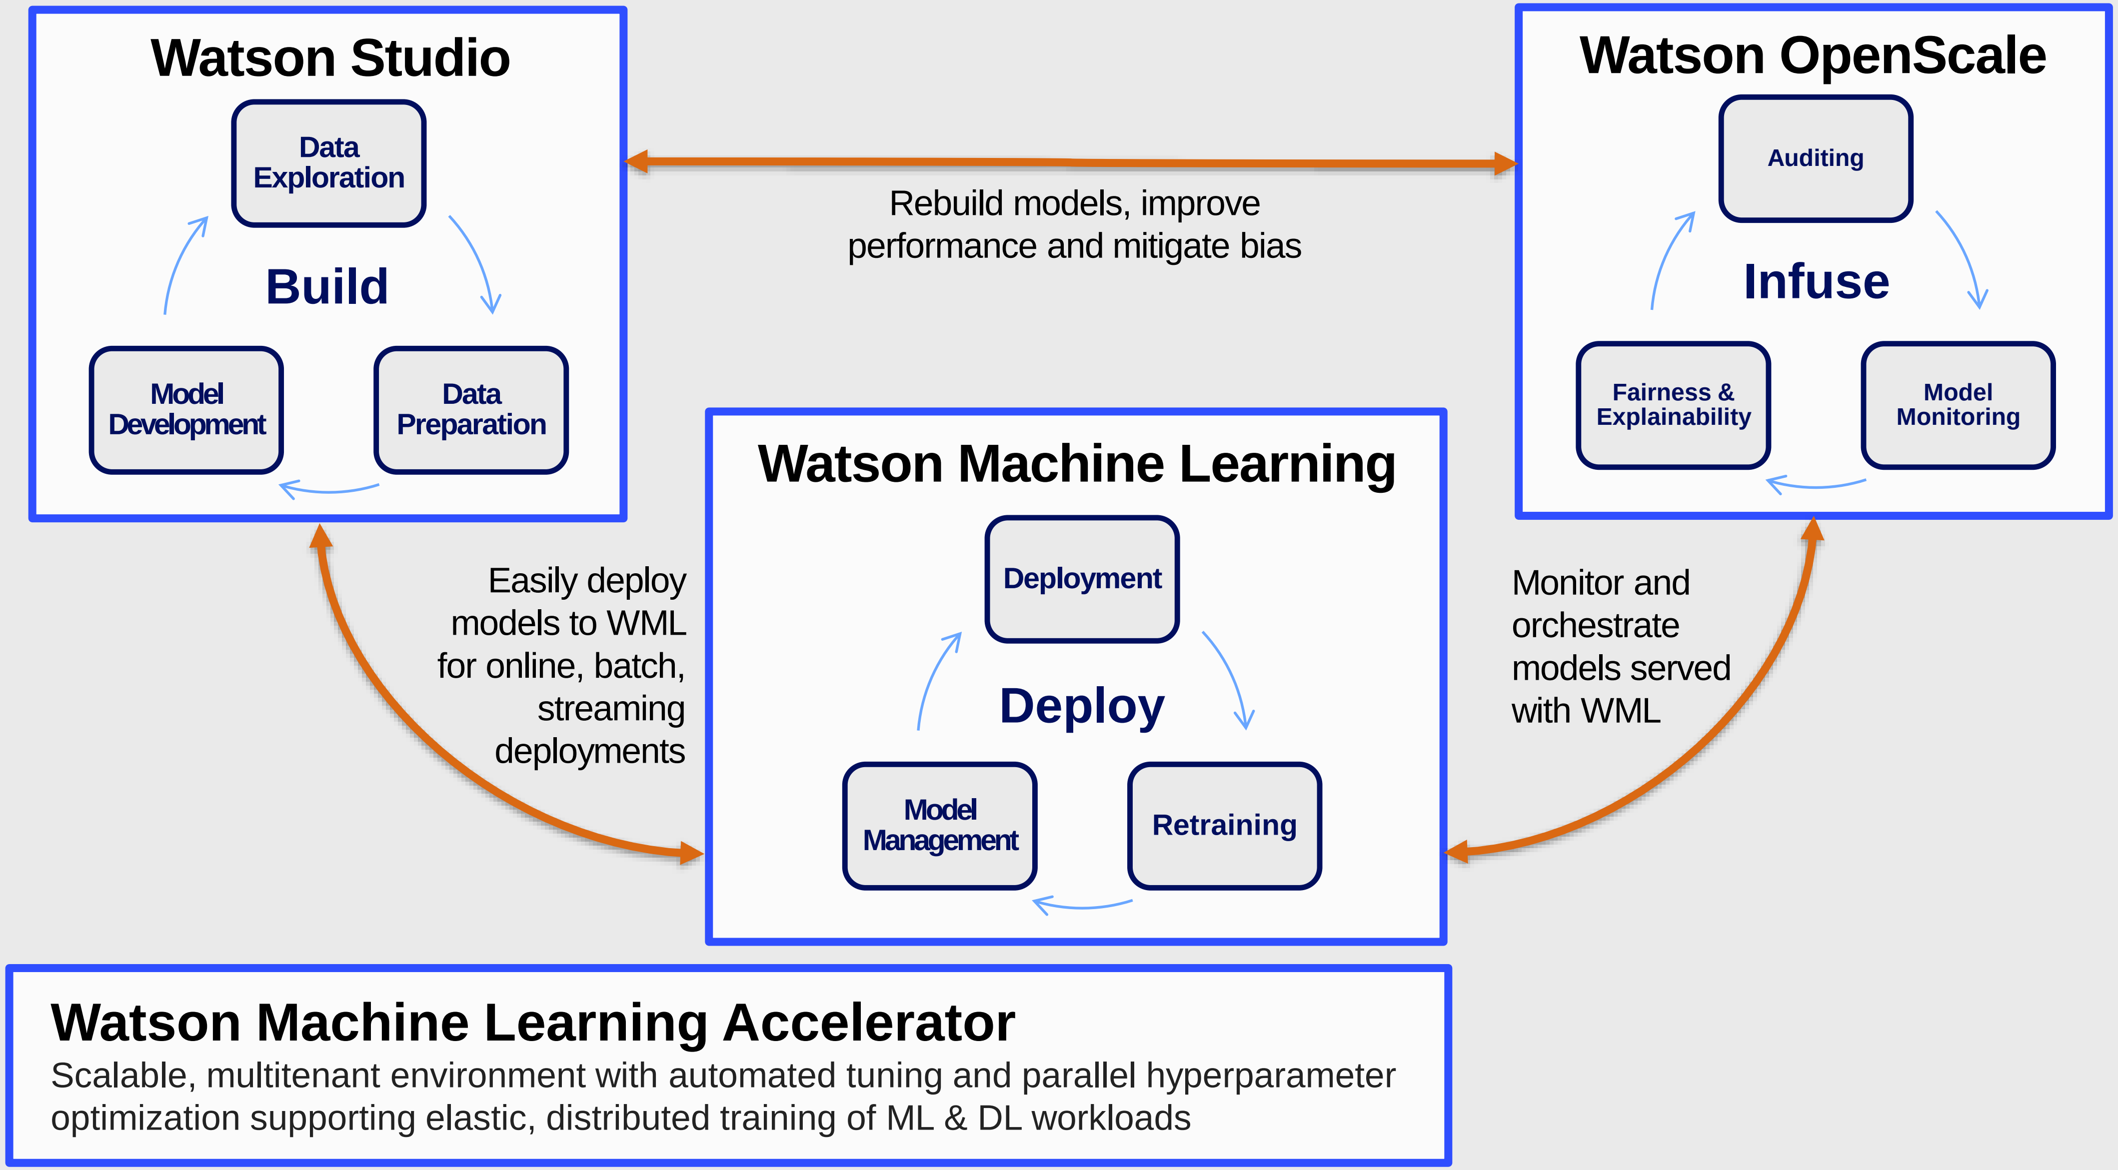

WSL / WML is a self-contained and extendable platform for developing and deploying analytics applications

WSL includes…

Collaborate Using Projects

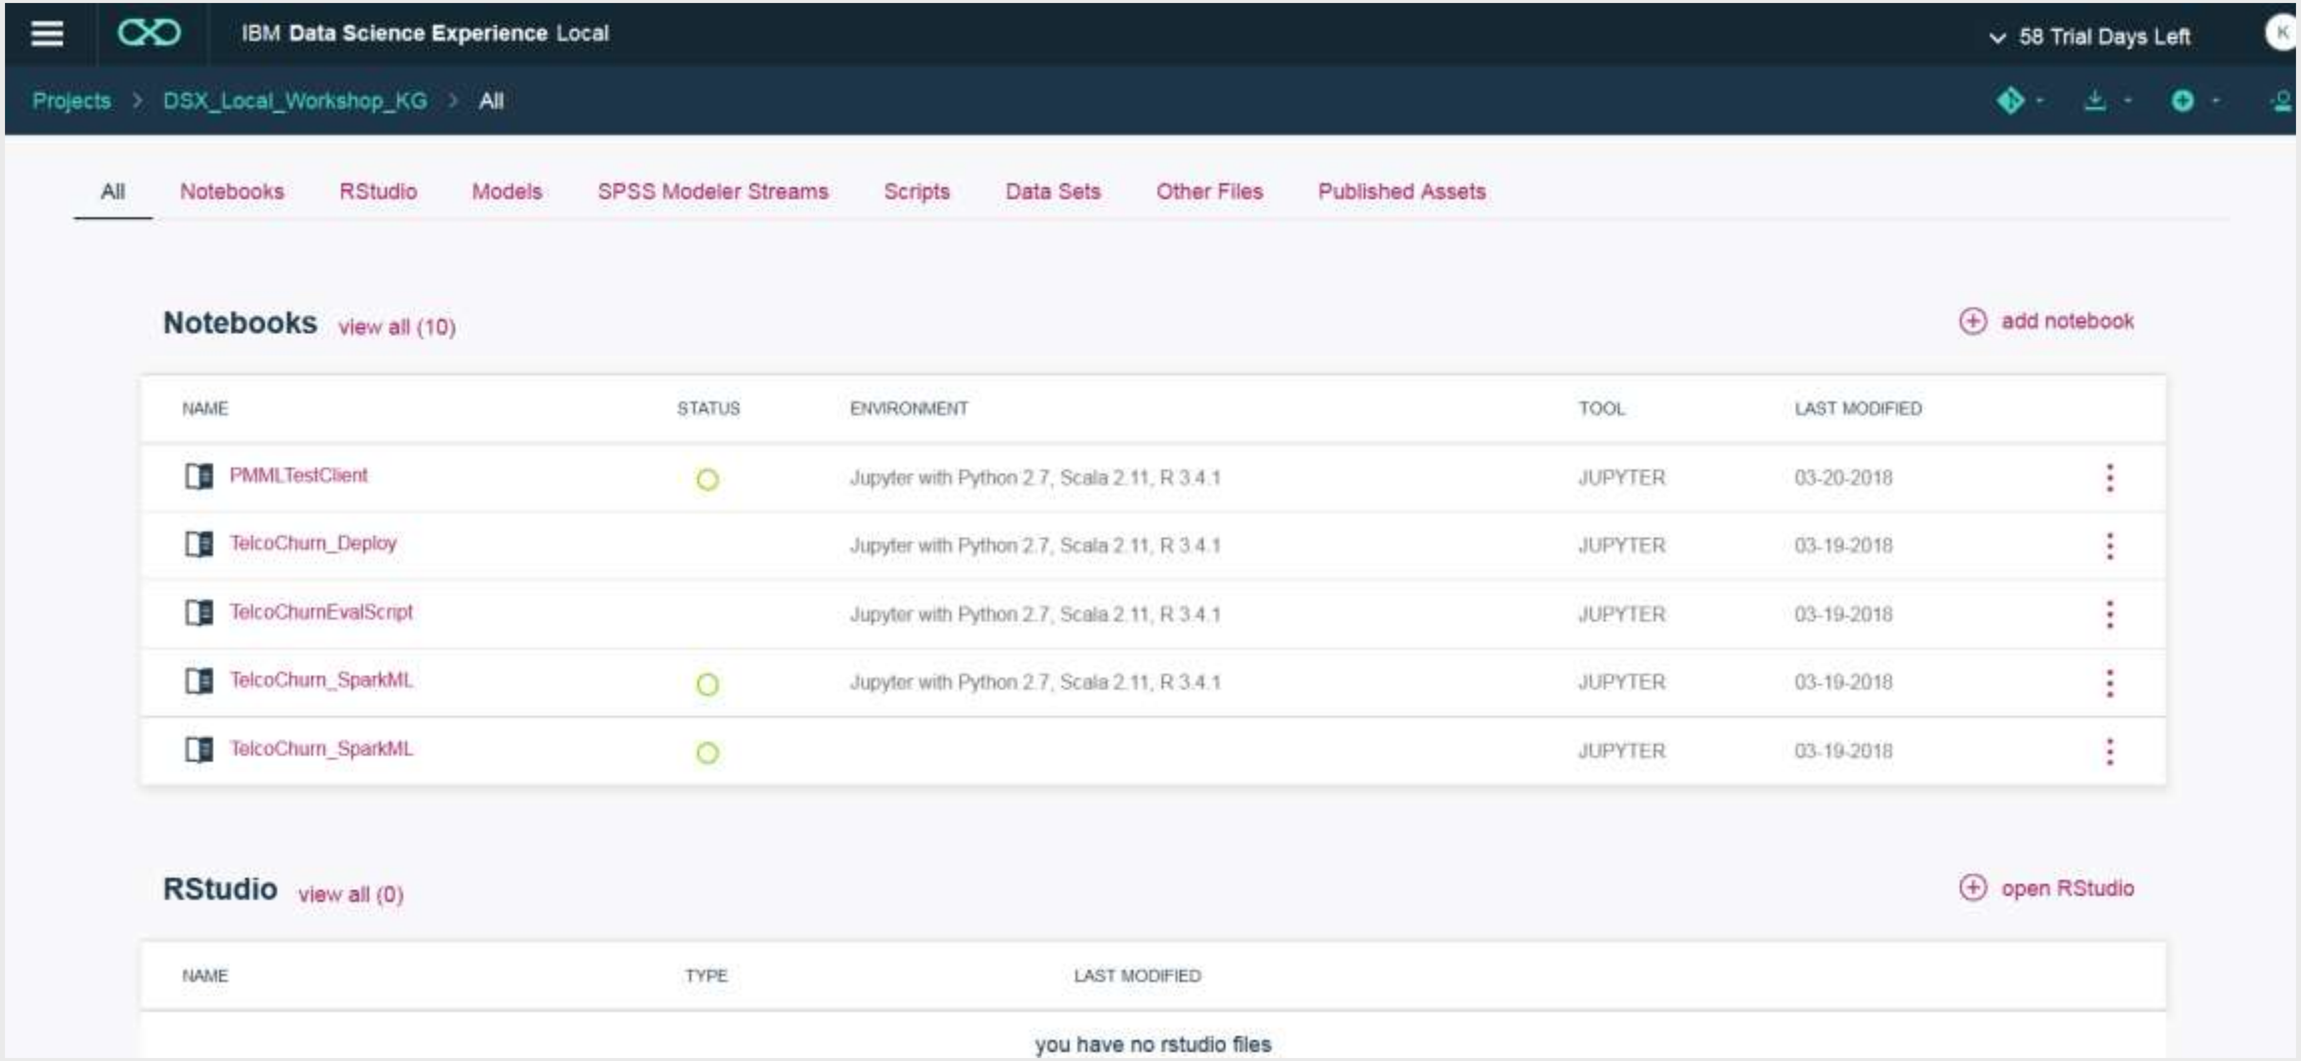

Create Projects

A project is **how you organize your assets** to achieve a particular data analysis goal.

Your project assets can include:

- Notebooks
- RStudio files
- Models
- SPSS Modeler Streams (add-on)
- Data sets (local files and remote data sets)
- Scripts
- Published assets

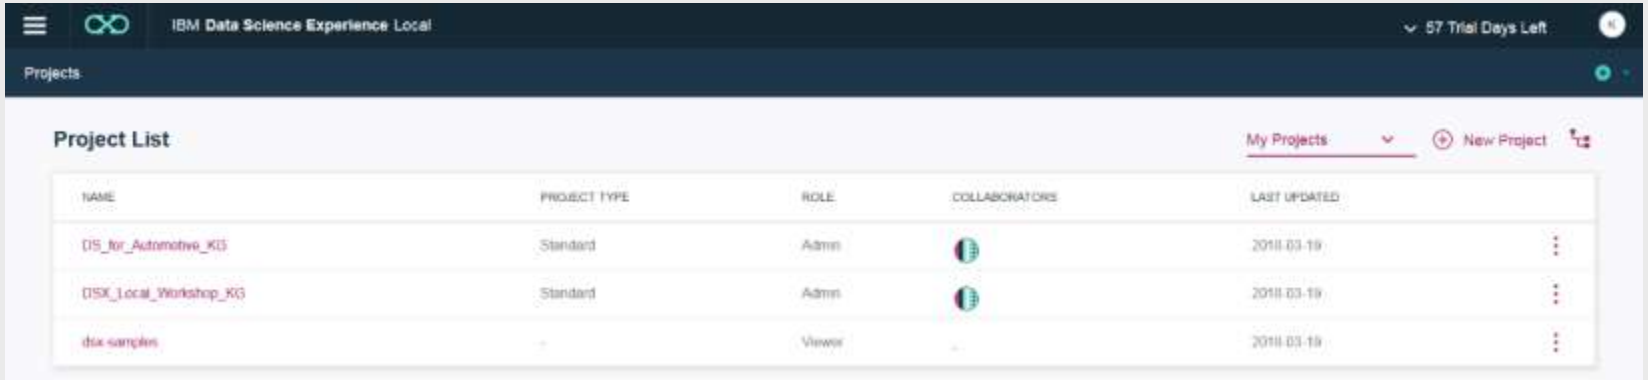

Share, fork and reuse Project assets to increase your data science team's productivity

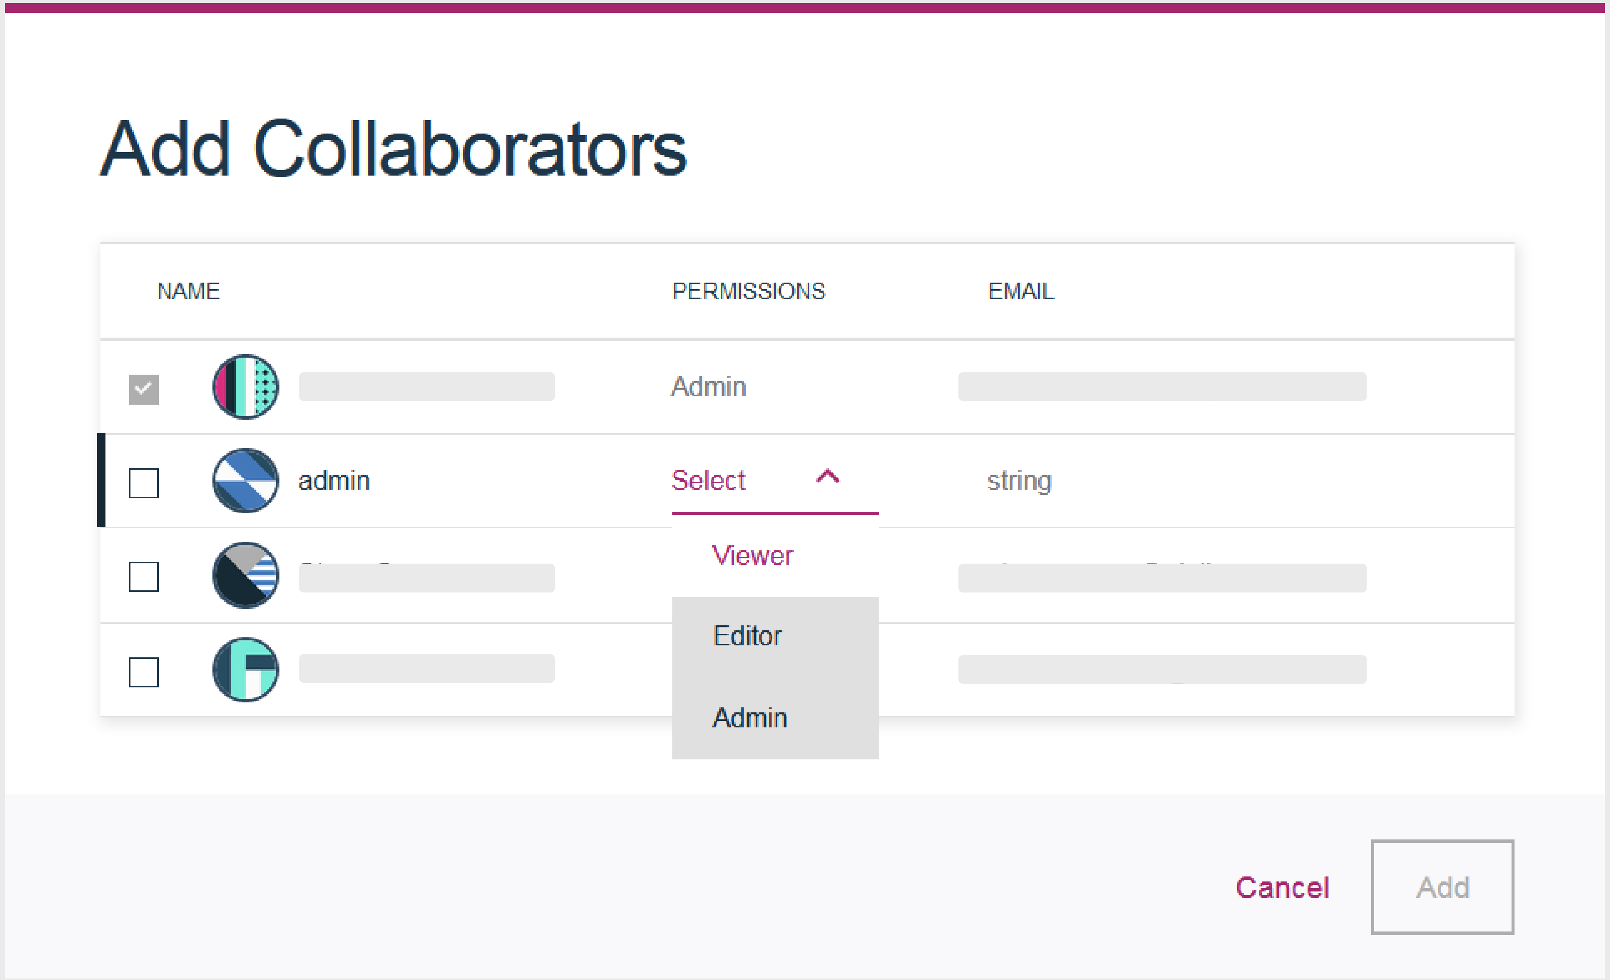

GitHub Integration

When starting a project you can:

- Create new project
- Import a project
- **Connect to a project on GitHub**

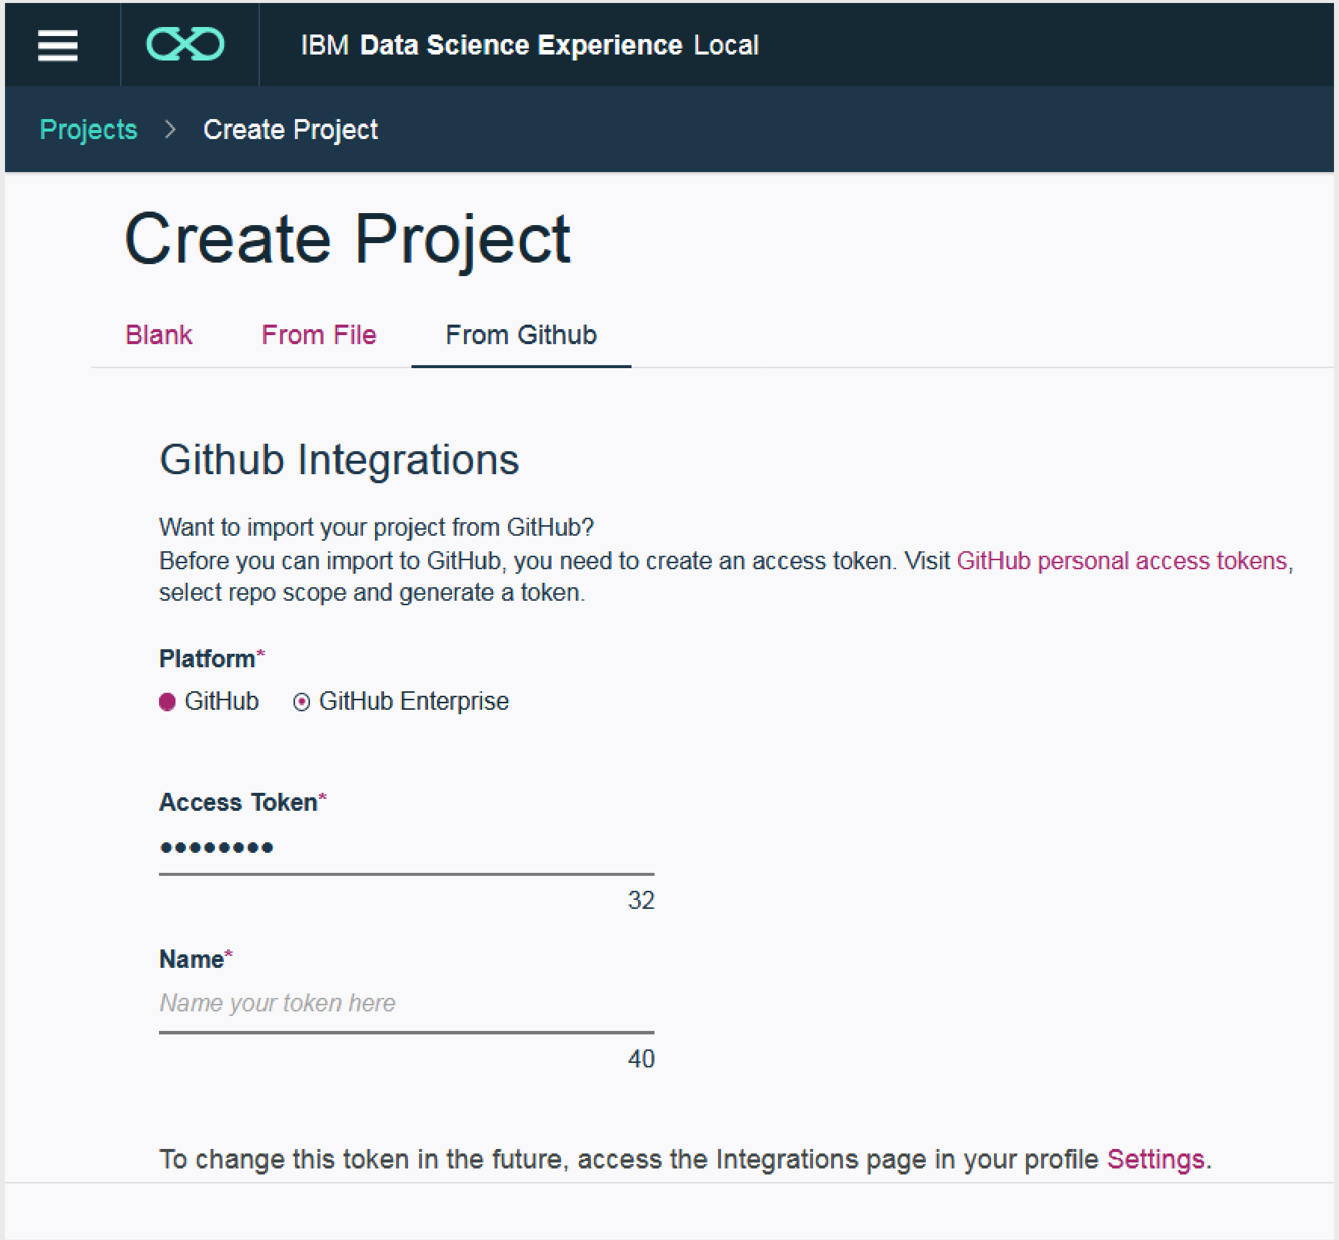

Integrated notebooks for interactive and collaborative environment – seamless Spark execution

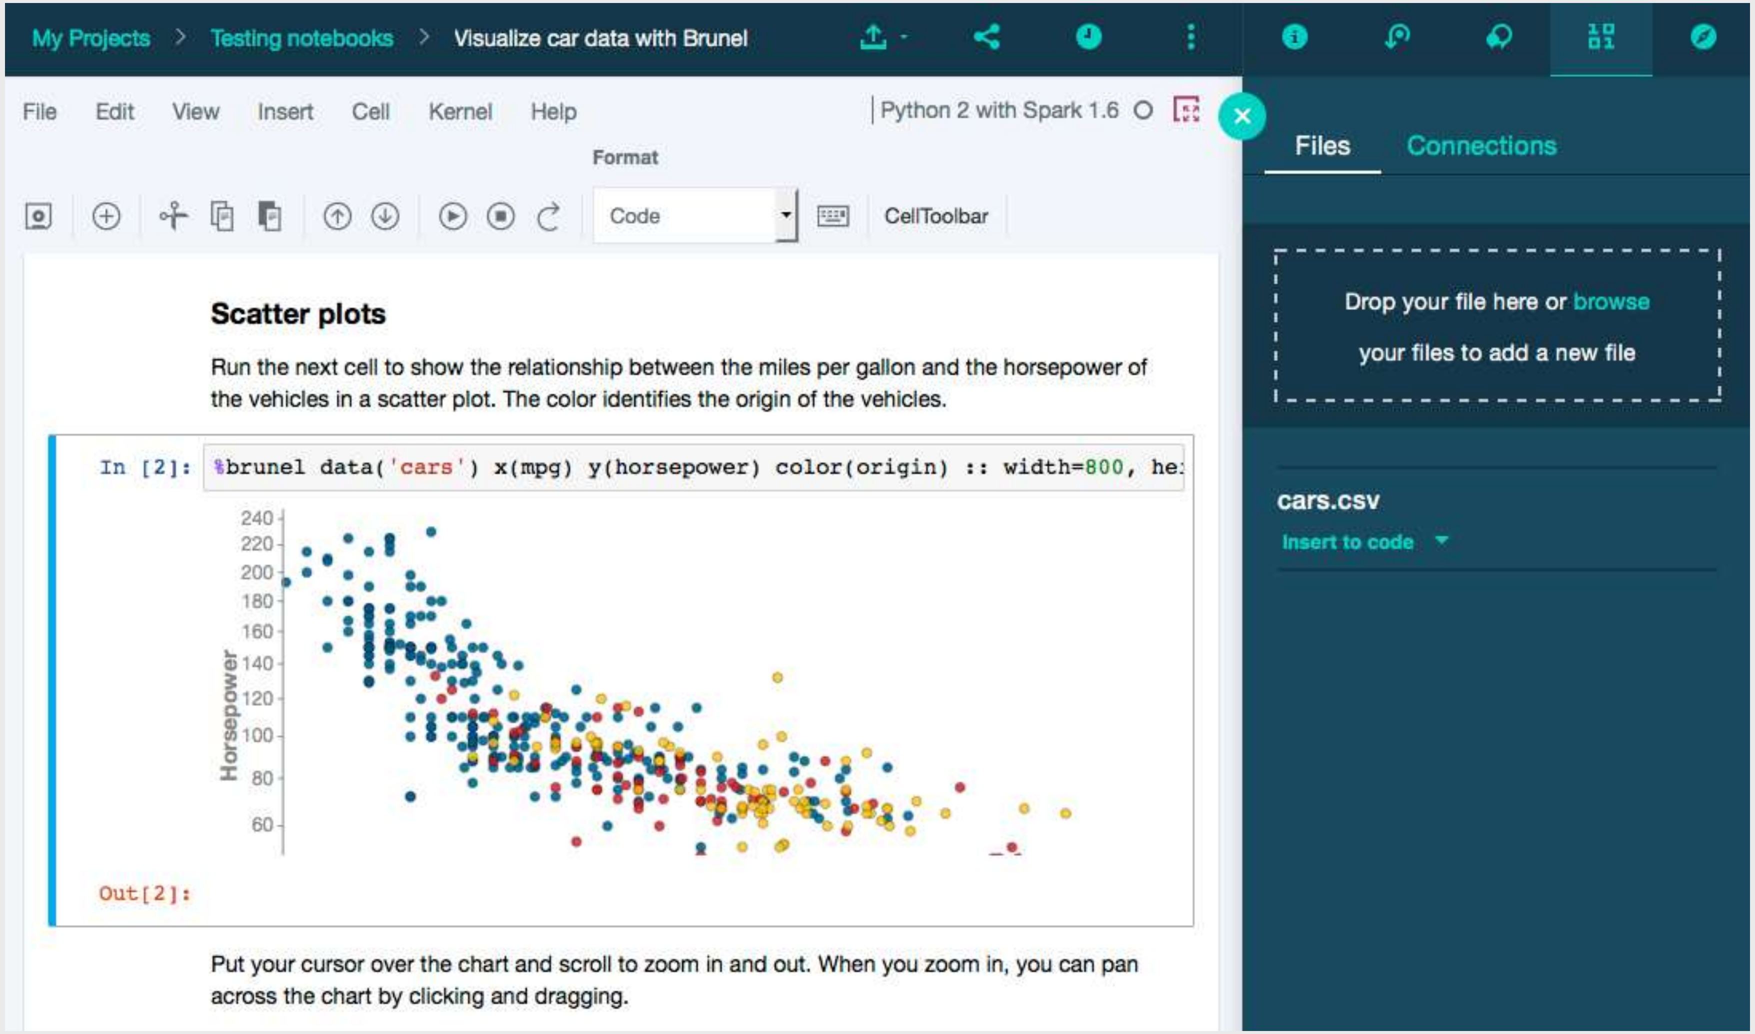

What is a notebook?

A notebook is **an open-source web application** that allows you to **create and share** documents that contain **live code, equations, visualizations and explanatory text**

Jupyter Notebooks

Zeppelin Notebooks

Create Jupyter Notebook & Select Environment & Language

Analyze Data and Create Model within Notebook

Evaluate

Batch Deploy

Deploy using RESTful API for Real-Time Scoing

Example WSL Machine Learning Workflow

IBM Watson Studio Local has R built into the experience thanks to our strategic partnership with RStudio

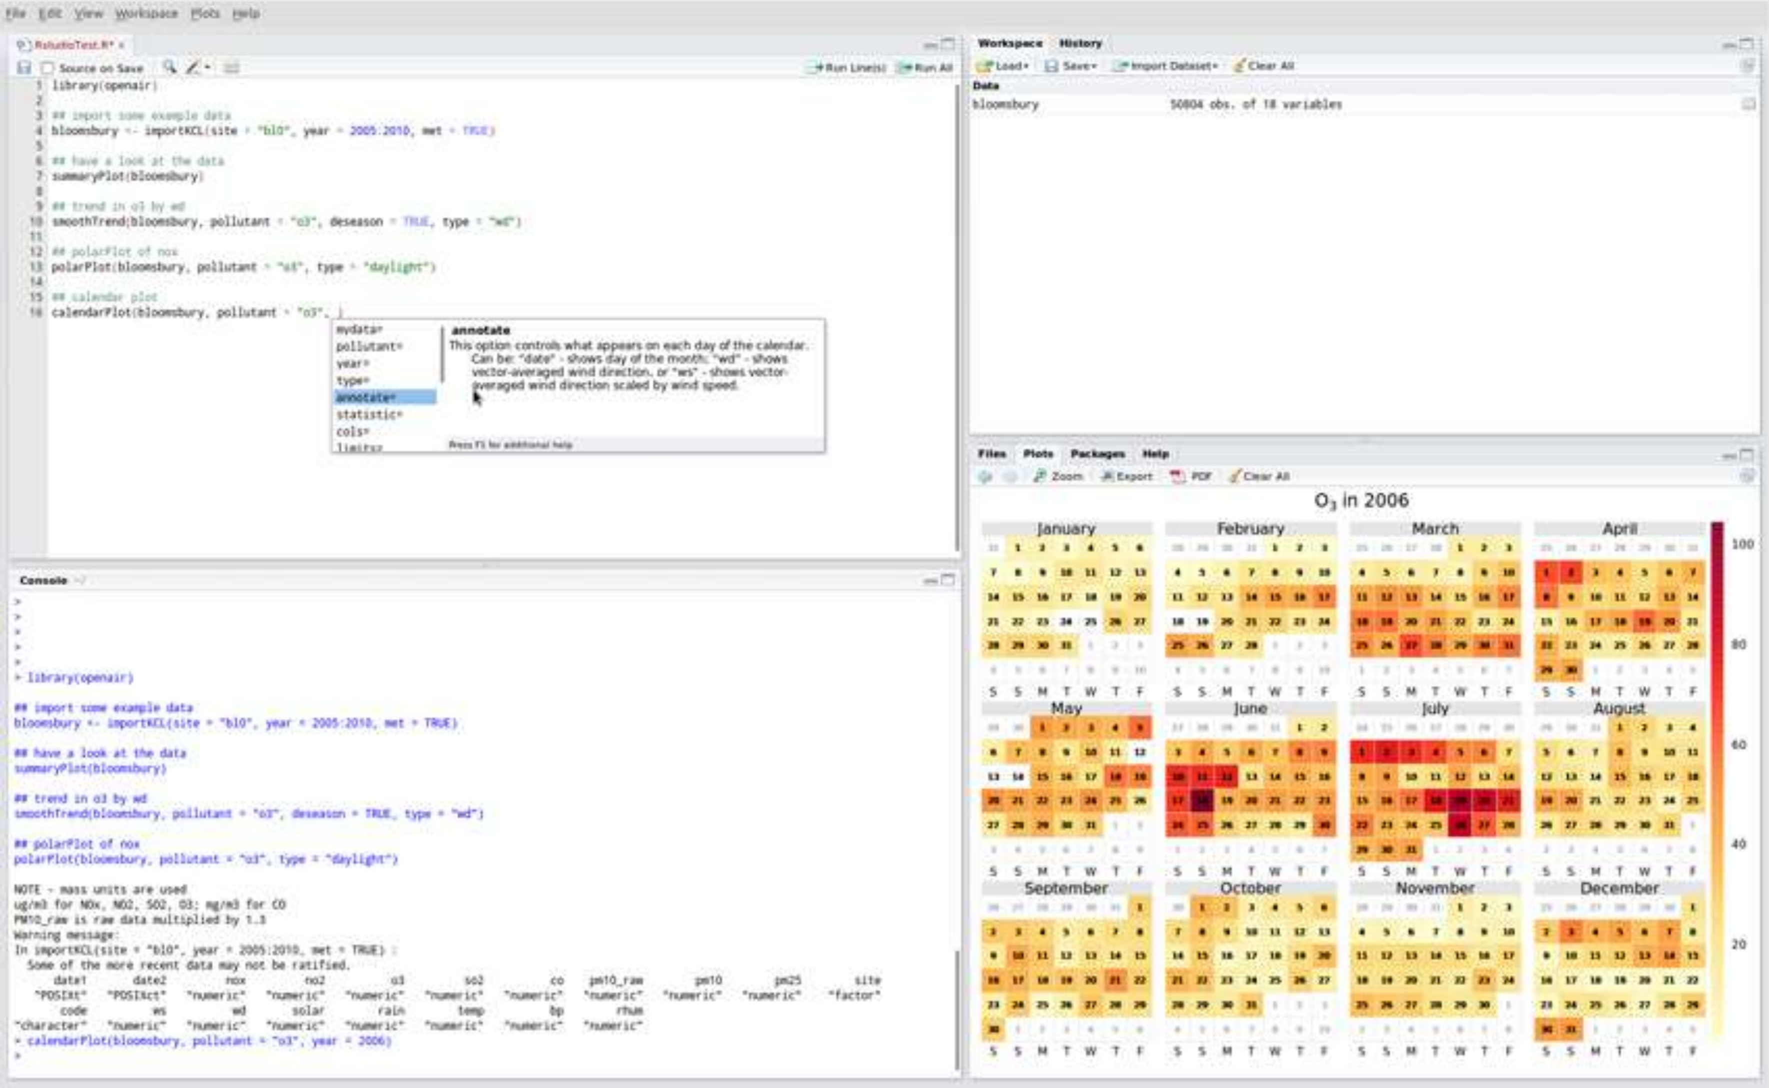

Deploy, monitor and manage

- Monitor models through a dashboard
- Model versioning, evaluation history
- Publish versions of models, supporting dev/stage/production paradigm
- Monitor scalability through cluster dashboard
- Adapt scalability by redistributing compute/memory/disk resources

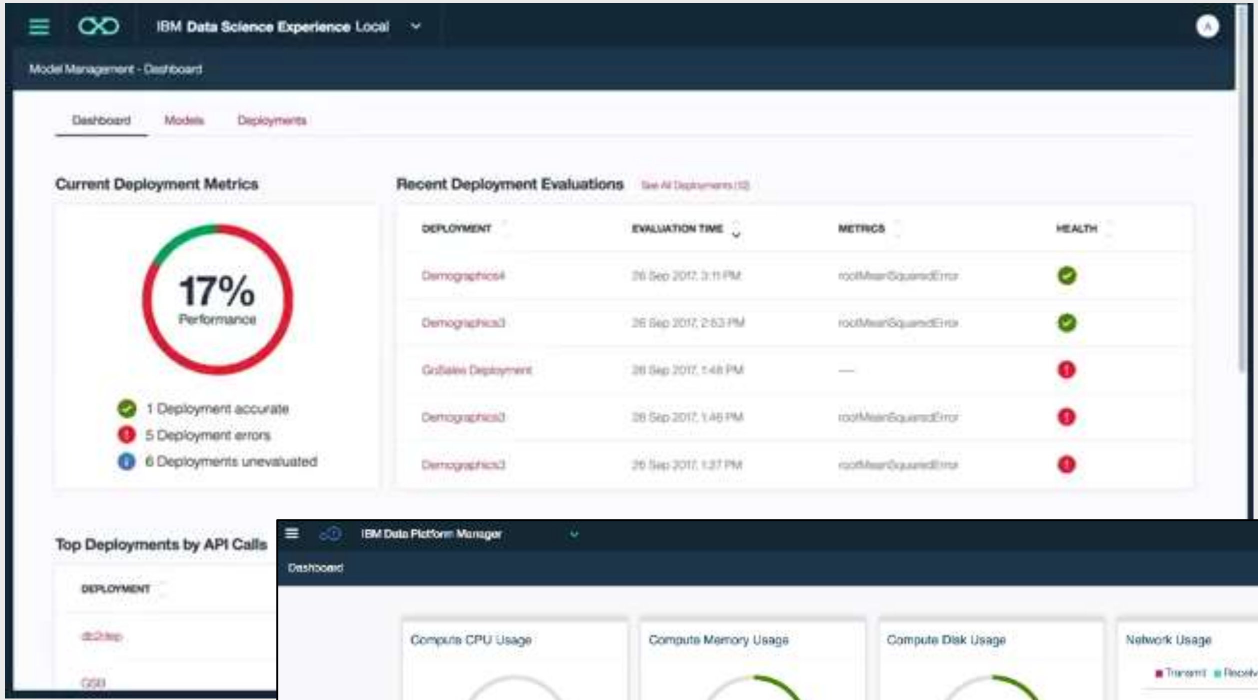

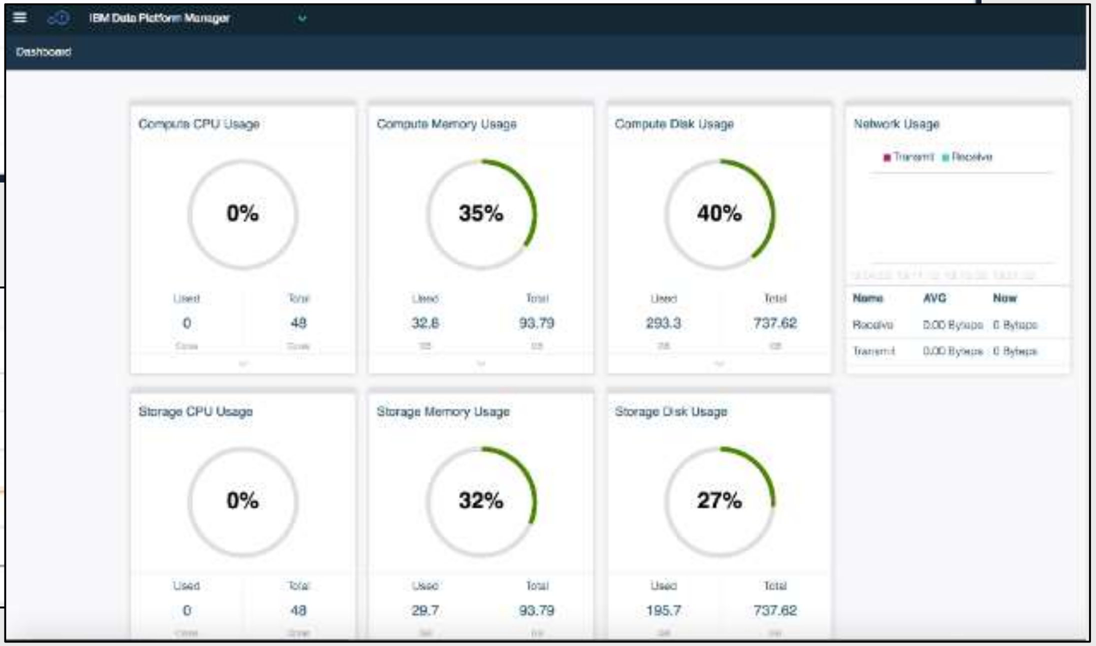

- Watch it in action
  - http://ibm.biz/WSL-Models

Assembling the building blocks of a data and analytics platform is complex, risky and time-consuming

CAPTURE | **SECURE** | ← **ANALYZE** → | **OPERATIONALIZE**

Collect, Connect, and Access Data

Connect and **discover** content from multiple data sources across your organization.

Provision databases and **virtualized** data access.

Govern, Search, and Find Data

Grant user access levels and **enforce business policies**.

Index for search, visualize consumers and producers of assets with **lineage, metrics, and quality profiles**.

Find data and analytics assets in the **Enterprise Catalog**.

Understand and Prepare Data for Analysis

Understand, cleanse and prepare your data to create data preparation pipelines visually.

Use popular open source libraries to prepare structured and unstructured data.

Build Descriptive, Predictive, and Prescriptive Models

Create Machine Learning, Deep Learning, Optimization, and other advanced mathematical models.

Tools to design your models **programmatically** or **visually.**

Train at scale with support for **distributed** compute and **GPUs**.

Model Management and Deployment

Manage your models across **dev, test, staging, and prod**.

Deploy your models and scale automatically for online, batch or streaming use cases with **SLAs**.

Monitor model **performance** and **automatically trigger retraining** and redeployment as rolling upgrades.

Create Analytics Applications

Incorporate **trusted and governed models** into applications, dashboards, and operational systems.

Reality
(“most information architectures”)

- Slow
- Siloed data and workflows
- Multiple disjunct stacks

Ideal State
(“ICP for Data”)

- Fast
- Pre-integrated and governed
- Containerized

What is needed is a single, unified platform

CAPTURE | **SECURE** | ← **ANALYZE** → | **OPERATIONALIZE**

- Collect, Connect, and Access Data
- Govern, Search, and Find Data
- Understand and Prepare Data for Analysis
- Build Descriptive, Predictive, and Prescriptive Models
- Model Management and Deployment
- Create Analytics Applications

Single, Unified Platform

With Added Value:

- Pre-Integrated
- Leverage Existing Investments
- Trusted, Business-Ready Foundation
- Built on Modern, Open Architecture
- Data Virtualization
- Ecosystem Management on a Single Pane

The Building Blocks of AI

Collect	Organize	Analyze
Collect and access data of any type, regardless of where it resides.	Find, integrate, transform and govern data to drive insights while mitigating compliance risks.	Descriptive, predictive, prescriptive: to understand the current, predict the future and change outcomes

Do it yourself – with some assembly required:

- Db2 and Db2 Warehouse
- Db2 Event Store
- MongoDB
- Hadoop

- Information Server
- Data Replication
- Master Data Management
- Optim and StoredIQ

- SPSS & WSL/WML
- Cognos Analytics
- Watson Explorer
- Planning Analytics

Or get these capabilities in an integrated end-to-end platform:

IBM Cloud Private for Data

IBM Cloud Private for Data

Unify on an open, multicloud Data & AI platform

Open, Extensible API Platform

App Developers **Business Partners**

Personalized, Collaborative Team Platform

Data Engineers **Data Stewards** **Data Scientists** **Business Users**

Extensible "add-ons" from IBM's Data & AI Portfolio & 3rd Party Services

IBM | IBM | IBM | IBM | IBM | ISV | OS | Custom | . . .

The Ladder to AI

For data of every type, regardless of where it lives

Collect Data

- Data virtualization
- Data warehousing
- Databases on-demand
- Data source ingestion
- Distributed processing

***Powered by**: Db2 Family*

Organize Data

- Discovery & search
- Data transformation
- Data cataloging
- Business glossary
- Policies, rules & privacy

***Powered by**: InfoSphere*

Analyze Data

- AI: ML and DL
- AI model building & deploy
- AI model management
- Dashboards & reporting
- Visualization & exploration

***Powered by**: Watson & Cognos*

Multicloud Services

- Logging
- Monitoring
- Metering
- Persistent Storage
- Kubernetes
- Security
- Identity Access Mgmt.
- Docker Registry / Helm

Power Systems | IBM Cloud | OpenShift | amazon webservices | Azure | openstack | Google Cloud

ICP for Data embeds Watson Studio
providing open & extensible data science tooling

Extensible Add-on AI Services

Watson Studio & Machine Learning

Data Science Premium
Tools to design, build and deploy AI models

SPSS Modeler · *Data Refinery* · *Model Builder* · *Decision Optimization* · *Watson Explorer* · *Streams Designer*

Watson OpenScale
- *Operational optimization services*

Watson Assistant
- *Conversational AI services*

Watson APIs
- *Interactive AI services (e.g., speech)*

Foundational Data & AI Services

Cloud Private for Data

Open Source frameworks to build, train and deploy AI models powered by **Watson Studio** and **Machine Learning**

TensorFlow · Keras · Caffe · XGBoost · RStudio · jupyter · Apache Zeppelin · Scala · learn · python · Spark ML · torch

- Data visualization
- Machine learning learning
- Model build & deploy
- Model management
- Dashboards & reporting

Data Collection Services | *Data Organization Services*

IBM Cloud · RED HAT OPENSHIFT · amazon web services · Azure · openstack · Google Cloud

ICP for Data Embeds Unified Governance & Integration Capabilities

Add-on Services

InfoSphere Regulatory Acceleration	InfoSphere Data Integration	InfoSphere Data Replication	InfoSphere Adv. Quality	InfoSphere Entity Resolution
Use ML to automate compliance policies	*Simplify enterprise ETL processes*	*Enable fluid multicloud data*	*Tailor quality criteria to unique needs*	*Automate business entity classifications*

Foundational Services

Data Organization Services

Discovery, prepare and activate data/AI models from multiple sources into a business-ready, self-service analytics foundation powered by **Watson Knowledge Catalog** and **InfoSphere**

Catalog Data
Auto discover, classify and catalog data & AI models

Transform Data
Migrate and transform data to be business & AI ready

Govern Data
Define & enforce policies across the Data/AI lifecycle

Multicloud Orchestration & Optimization

Machine Learning Automation

Data Collection Services | *Data Analytics / AI Services*

IBM Cloud · RED HAT OPENSHIFT · amazon web services · Azure · openstack · Google Cloud

Modularized View of ICP for Data

2. Premium Add-Ons (IBM & ISV)

Data Science Premium	**Db2 Event Store**	**Db2 AESE**	**Watson OpenScale**	**Streams**	**plus others, with more coming**
Datameer	**Figure Eight**	**Senzing**	**MongoDB**	**CockroachDB**	

1. ICP for Data

Additional Modules

Data Virtualization

(includes services from : new Db2 Technology)

Data Warehouse

(includes services from : Db2 Warehouse)

Additional Analytics Environments & Frameworks

(Optional Watson Studio content)

Recommended Install

Analytics Dashboards

(Powered by : Cognos CDE)

Core Platform Services & Data Science

Min Install
REQUIRED

(includes services from: Watson Studio & WML)

- User Management
- Roles & Privileges
- Add-on service mgmt.

- Projects
- R Scripts
- Python Scripts
- Jupyter with Python 3.6 & Anaconda 5.2

- Project Deployments
- Models
- Jobs & Scheduling

Unified Governance & Integration

(includes services from : Information Server)

- Discover Data
- Transform Data
- Enterprise Data Catalog
- Business Glossary
- Policies & Rules
- Enterprise Search

ICP Foundation

- Logging
- Monitoring
- Metering
- Persistent volume /Storage
- Kubernetes
- Security
- Identity Access Mgmt.
- Docker registry/Helm chart mgmt.

IBM Cloud Private for Data

Top Use Cases

Accelerate Data & Analytics Monetization

Capitalize on your data as a strategic asset - access all enterprise data **regardless of location and speed**.

Analyze **all data types** at the source to accelerate insights and time to value.

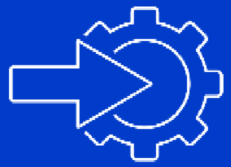

Operationalize Data Science & AI

Build, deploy, manage and govern models with trusted data **at scale**.

Data scientists, developers, engineers, and business experts work **collaboratively** to accelerate innovation and improve business outcomes.

Modernize to Next-gen Workloads

Benefit from a **cloud native** architecture while keeping your data where you want it.

Build once and deploy anywhere, provision and scale in minutes, and achieve the agility you need to evolve your workloads with business needs, without the complexity.

Smarter Governance

Proactively adapt to comply with new or anticipated regulations through **continuous visibility into your data landscape**, universal oversight, and **automated controls** to prevent data misuse.

IBM Cloud Private for Data

Improving confidence in data-driven outcomes

Hear why GuideWell, a rapidly growing healthcare company, is excited about IBM Cloud Private for Data and its potential to help provide better outcomes to their customers: https://www.youtube.com/watch?v=R7P3Ee5n7MQ

> *“ It's a very exciting product. You have a lot of things that a lot of different companies are doing, molded into one technology suite. ”*
>
> —James Wade, Director of Application Hosting, GuideWell

Collect Data

Collect data of every type, no matter where it lives, and achieve freedom from ever-changing data sources

Organize Data

Organize your data into a trusted, business-aligned source of truth to put data to work in new ways.

Analyze Data

Analyze your data in smarter ways, empowering all your teams with Machine Learning Everywhere

Modernize

Make your data ready
for AI and the cloud

“One of the key technologies that the Hub will deploy is IBM Cloud Private for Data. The exciting thing about IBM Cloud Private for Data is how quickly we will be able to drive new innovations in FinTech and SportsTech using the AI and ML microservices within the platform. What makes it especially attractive is that it enables us to develop and deploy new models quickly that brings AI to the data, rather than the other way around.“

— Abdulaziz Al Khalifa
CEO, Qatar Development Bank

QDB
بنك قطر للتنمية
QATAR DEVELOPMENT BANK

Qatar Development Bank (QDB) is a bank in Qatar offering financial services, banking and loans to the development of the industrial, tourism, educational, health care, agricultural, animal resources and fisheries sectors of the Qatari economy. It was created in 1997 by Emiri Decree No. 14 with the objective of diversifying Qatar's economy by promoting developmental projects.

Modernize

Make your data ready
for AI and the cloud

Quote from blog : Trusting AI to save lives in India
https://www.ibmbigdatahub.com/blog/trusting-ai-save-lives-india

“ IBM Cloud Private for Data facilitated the model development and deployment of a predictive model for cardiac care for iKure. IBM’s Data Science Elite team demonstrated the model development process in ICP for Data via Watson Studio with multiple AWS data sources. It also proved model accuracy using patient clinical and demographic variables and physician feedback with the added benefits of rapid model development, publication and iteration.”

— Sujay Santra
CEO and Founder, iKure

iKure

iKure is an award-winning, tech-savvy, rapidly-growing, revenue-positive social enterprise that meets the primary health care and prevention needs through a unique combination of health outreach initiative, skills development, and technology intervention. The venture is poised to rapidly scale beyond its curative model, looking to the future of disease prevention and wellness for 840 million people in rural India. iKure’s healthcare model has acquired extensive support and recognitions from across the world for being innovative, technologically advanced and sustainable.

IBM Cloud Private for Data

Cloud-Native by Design

Agility ● **Efficiency** ● **Cost Savings**

IBM Cloud Private for Data

Microservices

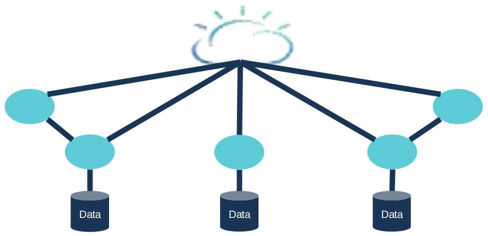

An architecture of loosely coupled data services, easily refactored to create containerized workloads

Containerized Workloads

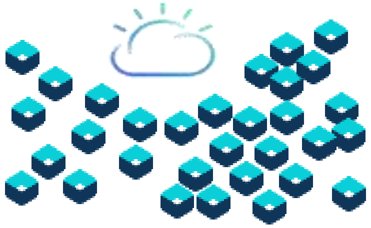

Stand-alone workloads composed of microservices & data that are flexibly deployed, orchestrated and managed

Multicloud Provisioning

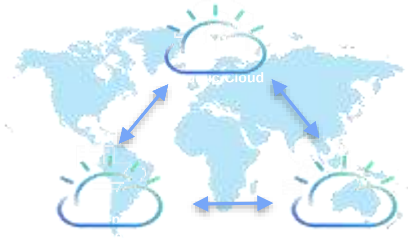

Agile provisioning of containerized workloads in multicloud environments and consumption of cloud services

Underpinned by an Open Ecosystem

Where we partner, co-create & lead

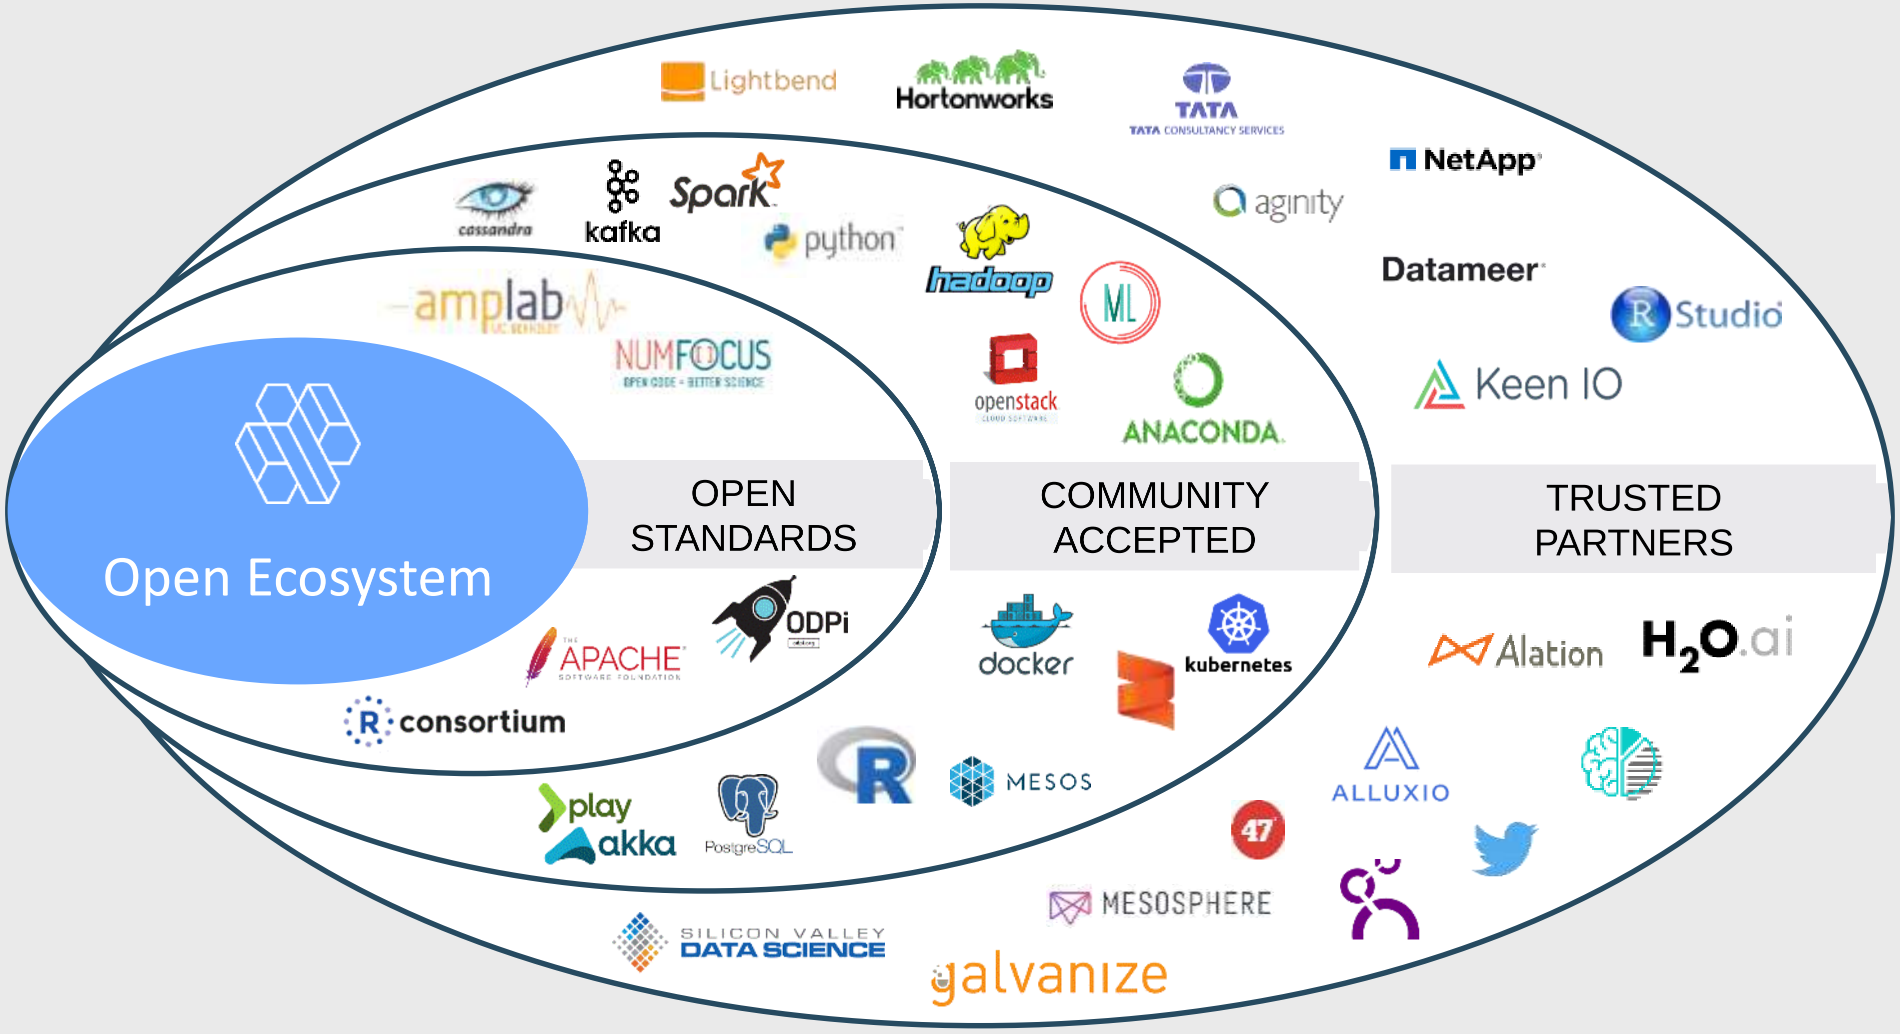

Integrating Industry Leading AI Technology

Forrester Wave™ Reports

Leader:
Watson Studio

Predictive Analytics & Machine Learning

Leader:
Watson Knowledge Catalog

Machine Learning Data Catalogs

Leader:
Cloud Private for Data

Enterprise Insights Platforms

Leader:
Watson Assistant

Conversational Computing Platforms

Leader:
Watson Discovery

AI Text Analysis Platforms

IBM #1 in AI Market Share

IDC

Industry Awards

Reddot Design Awards

The AIconics

iF WORLD DESIGN GUIDE

A'DESIGN AWARD

IBM Cloud Private for Data
A new approach to Data Virtualization

1 **Query anything, anywhere.**

Query **many heterogenous data sources as one** across cloud, on-premise and mobile with advanced analytics using the most popular languages and tools

2 **Simplicity and scalability.**

Automatically discover, and connect **few to many devices and data stores** into a single self balancing constellation. Avoid the complexity of centralized copies. Data only persists at the source.

3 **Execution speedup.**

Many times acceleration using the power of every device to compute and aggregate results.

4 **Security.**

Fully secure and encrypted communication and preservation of data access rights at source.

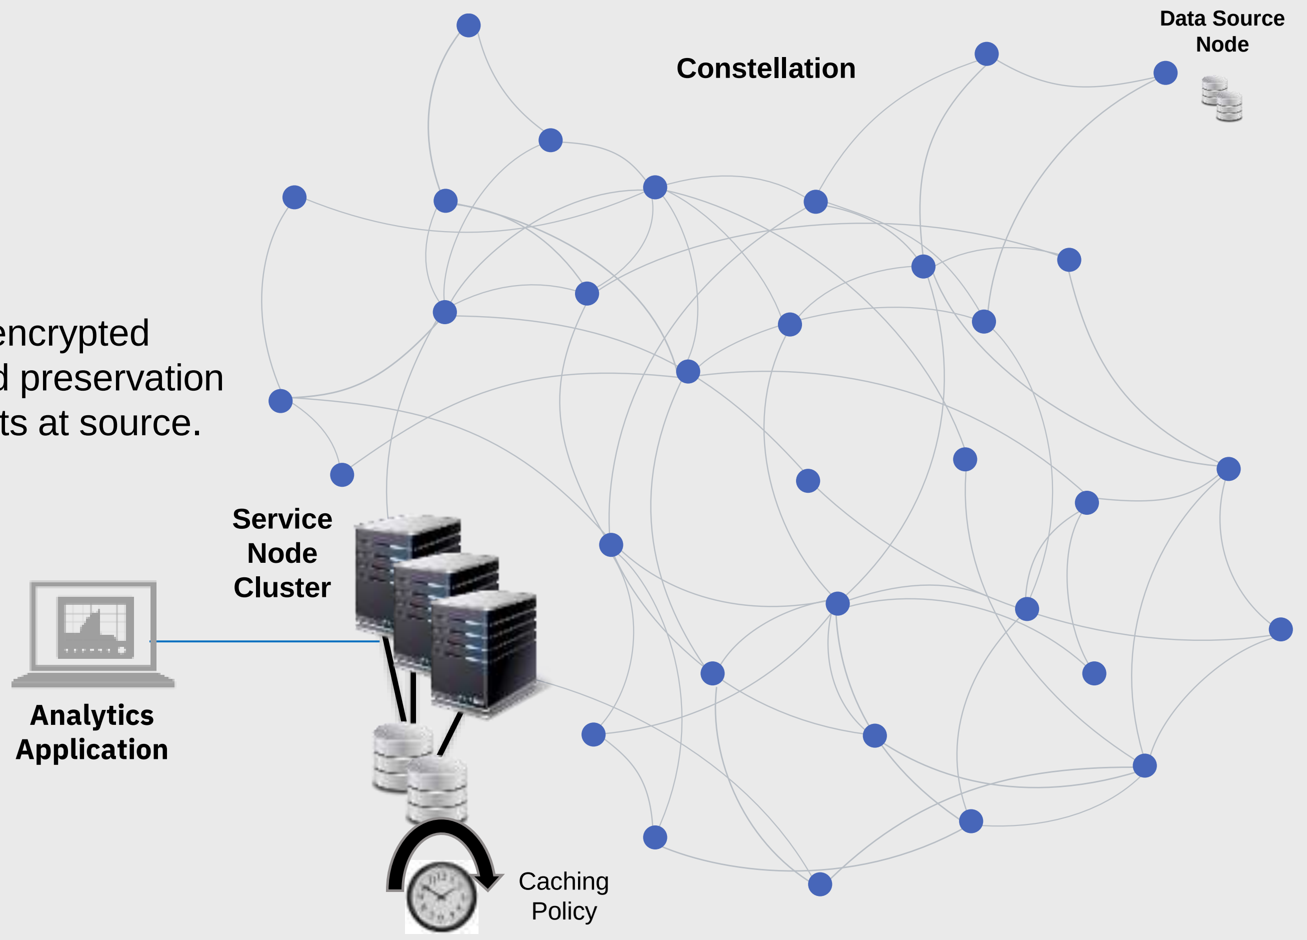

IBM Watson in Multicloud Environments

Deploy "Watson Anywhere" with IBM Cloud Private for Data

I want to use Watson to build and deploy AI on the clouds of my choice:

1. Provision ICP/D which includes base Watson Studio, WML & Knowledge Catalog
2. Optionally provision the Data Science Premium add-on from ICP/D catalog to add premium features a well as optionally provision the Watson APIs add-on from the ICP/D catalog

Watson APIs

Data Science Premium

Cloud Private for Data

I want to use Watson Assistant on the clouds of my choice:

1. Provision Watson Assistant on the cloud which also provisions the embedded ICP/D platform
2. If you already have ICP/D deployed, Watson Assistant will use that instance

Watson Assistant

I want to use Watson OpenScale on the clouds of my choice:

1. Provision ICP/D on the cloud
2. Provision add-on from ICP/D catalog for Watson OpenScale
3. If you already have ICP/D deployed, Watson OpenScale will use that instance

Watson OpenScale

Cloud Private for Data

I want to progressively deploy the full Watson stack on the clouds of my choice

1. Starting with provisioning ICP/D, progressively provision the Watson offerings from the ICP/D add-ons catalog

Watson Assistant

Watson OpenScale

Watson APIs

Data Science Premium

Cloud Private for Data

Provision & Use On Any Cloud:

IBM Cloud · OPENSHIFT · amazon web services · Azure · openstack · Google Cloud

Choose the Right Edition For Your Needs

Cloud Native Edition

IBM
Cloud Private
for Data

Enterprise Edition

- All features are supported across both editions
- Cloud Native Edition has the following limits:
  - # of records in data set on which you can run automated discovery or data transformation jobs (1 million)
  - # of jobs that can run each day to transform data or assign terms (500)
  - # of accepted assets in the enterprise data catalog (3,000)
  - # of concurrently published data science model deployments (50)
  - # of concurrently running data science batch jobs (60)
  - # of concurrently running Shiny apps (40)
  - # of active nodes across all data warehouses (4)

IBM Cloud Private for Data enables you to:

1) Deploy an information architecture for AI
2) Modernize your data estate for a multicloud world
3) Make your data ready for AI
4) Infuse AI everywhere, with confidence
5) Put open source to work

ICP for Data: Additional Sources of Information

Learn more

Resource Center: https://docs-icpdata.mybluemix.net
Product Page: https://www.ibm.com/products/cloud-private-for-data
Community: http://ibm.biz/ICP4D-Community
Product Roadmap: http://ibm.biz/ICP4D-Roadmap

Watch it in action

Product Walkthrough: https://www.youtube.com/watch?v=TiYCZfCZx_o
Collect & Organize: https://www.youtube.com/watch?v=vY3jPXwrhyE
Analyze: https://www.youtube.com/watch?v=-Z4LAe5OkfU
Stock Trader Demo: https://www.youtube.com/watch?v=aD4O3MMhlgo

---

Try it

ICP for Data Experiences: http://ibm.biz/experienceICP4D
DTE Assets: http://ibm.biz/ICP4D-DTE

Get help

IBM Cloud Garage & Systems Co-creation Workshops: http://ibm.com/cloud/garage
- Various engagement offerings available

IBM Data Science Elite Team: http://ibm.biz/DSEliteTeam
- 30 minute consult (free), 6 week kick-start (free), 12 week build (co-invest)

Infrastructure Matters for Private Clouds and Enterprise AI

What do businesses need?

Reliability & Availability	*98% surveyed said that a single hour of downtime costs over $100,000 (and much more for some)* [1]
Performance	*Meet the most stringent application SLAs and scale to handle growth over time*
Security	*A top concern for CIOs – the average cost of a single data breach globally is $3.86M* [2]
Enterprise-cloud ready	*Agile with elastic growth and flexible consumption models*
Cost-effectiveness	*Make most efficient use of resources and do more with less (29% of organizations are reporting a decline in IT capital budgets for 2019* [3]*)*

IBM Power Systems

When data-intensive and AI workloads are the bottom line

Built-in PowerVM virtualization, IBM POWER9-based Power Systems are cloud-ready, enabling you to deploy the right cloud environment to meet your needs.

Enterprise cloud-ready

Power Systems easily integrates into your organization’s private or hybrid cloud strategy to handle flexible consumption models and changing customer needs.

#1 in reliability

Ranked #1 in every major reliability category by ITIC, IBM Power Systems deliver the most reliable on-premises infrastructure to meet around-the-clock customer demands.

Industry-leading value and performance

With Power Systems, clients can take advantage of superior core performance and memory bandwidth to deliver both performance and price-performance advantages.

Modernization To Private Cloud

with IBM Power Systems

Simplifies Cloud Application Deployment

Over **50% faster** to deploy applications versus traditional infrastructure [1]

And is **27% less** expensive than the public cloud to run a typical workload mix [2]

With **88% more** containers per core supported on Power versus x86 [3]

Ranked Number 1 in every major reliability category by ITIC

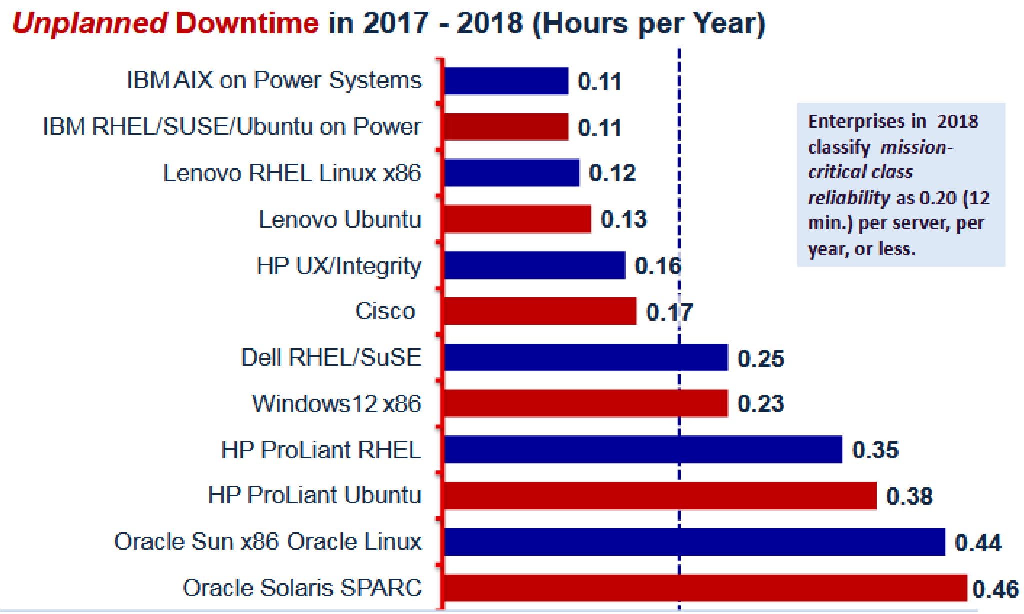

"IBM POWER8-based processor systems and the latest POWER9 servers provide several key feature/function advantages that advance reliability and enable customers to lower Total Cost of Ownership (TCO) and achieve near-immediate ROI."

Modern Infrastructure For Modern Data Workloads

2X Price/Performance Leader versus x86

For data-intensive, scale-out databases and analytics workloads such as MongoDB, EDB Postgres and IBM Db2

Accelerate your AI Journey

Train machine learning & deep learning models faster and accelerate the data science workflow

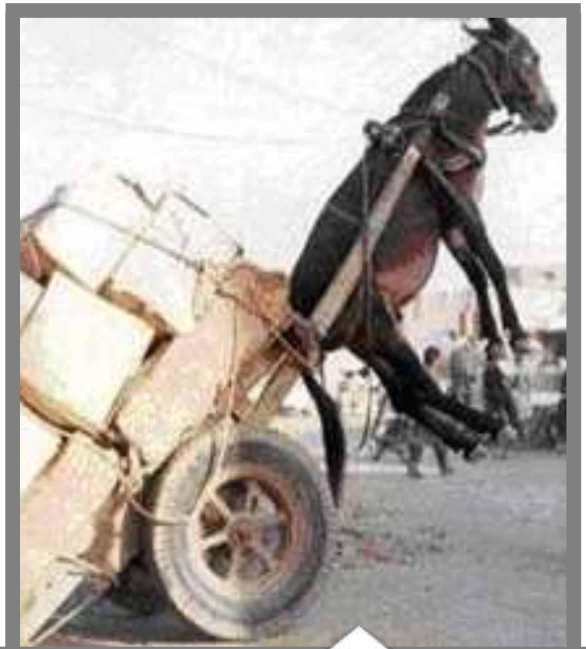

Other

Industry leading container and VM density

In a MongoDB benchmark running on IBM Cloud Private (with open-source Docker on Linux), IBM Power Systems outperformed a comparable Intel system by having…

3.6x

more containers/core

3.3x

better price-performance

POWER9 processor

1st chip with PCIe4

~1TB/s BW into chip

4GHZ PEAK FREQUENCY

7TB/s on chip BW

8 BILLION TRANSISTORS

>15 MILES OF WIRE

>24B VIAS

17 LEVELS OF METAL

POWER9 vs. x86 Xeon SP (Skylake)

2x[1]	performance per core
2.6x[2]	more RAM per socket
1.8x[3]	memory bandwidth per socket

POWER9 with NVLink vs. x86 Xeon

9.5x[4]	CPU to accelerator bandwidth

IBM POWER SYSTEMS for AI

AC922

An Acceleration Superhighway
Unleash state of the art IO and accelerated computing potential in the post “CPU-only” era

Designed for the AI Era
Architected for the modern analytics and AI workloads that fuel insights

Delivering Enterprise-Class AI
Flatten the time to AI value curve by accelerating the journey to build, train, and infer deep neural networks

IBM Power Systems LC922 & LC921

For the Analytics that feed Enterprise AI

The IBM Power Systems LC922 enhances the LC product line's open heritage while delivering superior performance in a cost optimized design needed in today's AI Era.

2x

Price performance advantage for data intensive applications such as MongoDB

59%

Improved Spark price-performance for efficiency across the AI data leveraging the P9 thread density for large amounts of concurrent Spark queries

2X

more data scientists on a single server at FASTER RESPONSE TIMES with Watson Studio Local (WSL)

Train larger more complex models with WML-A and AC922

Traditional Model Support

Limited memory on GPU forces tradeoff
in model size / data resolution

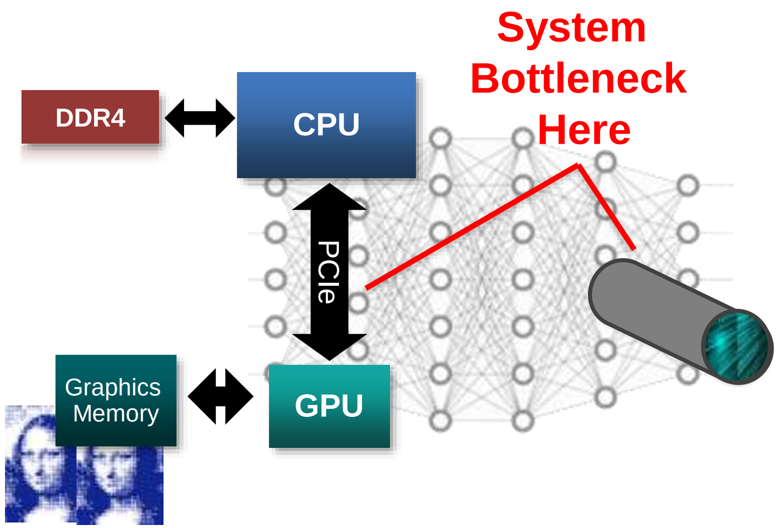

Large Model Support

Use system memory and GPU to support more
complex and higher resolution data

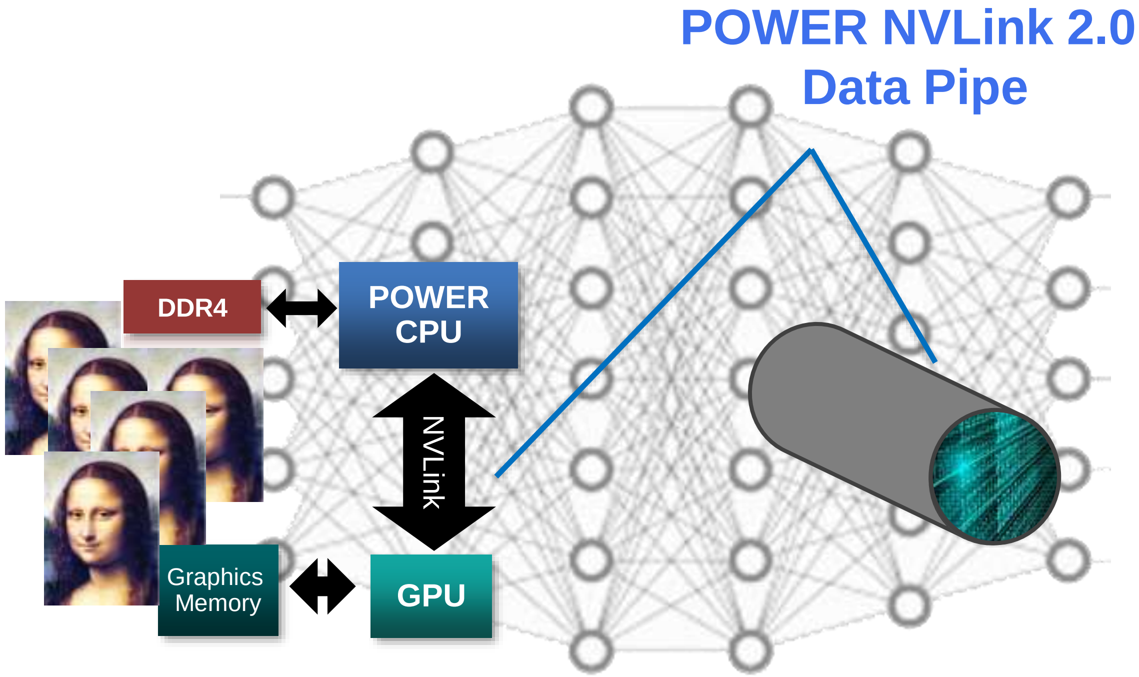

Typical GPU connectivity

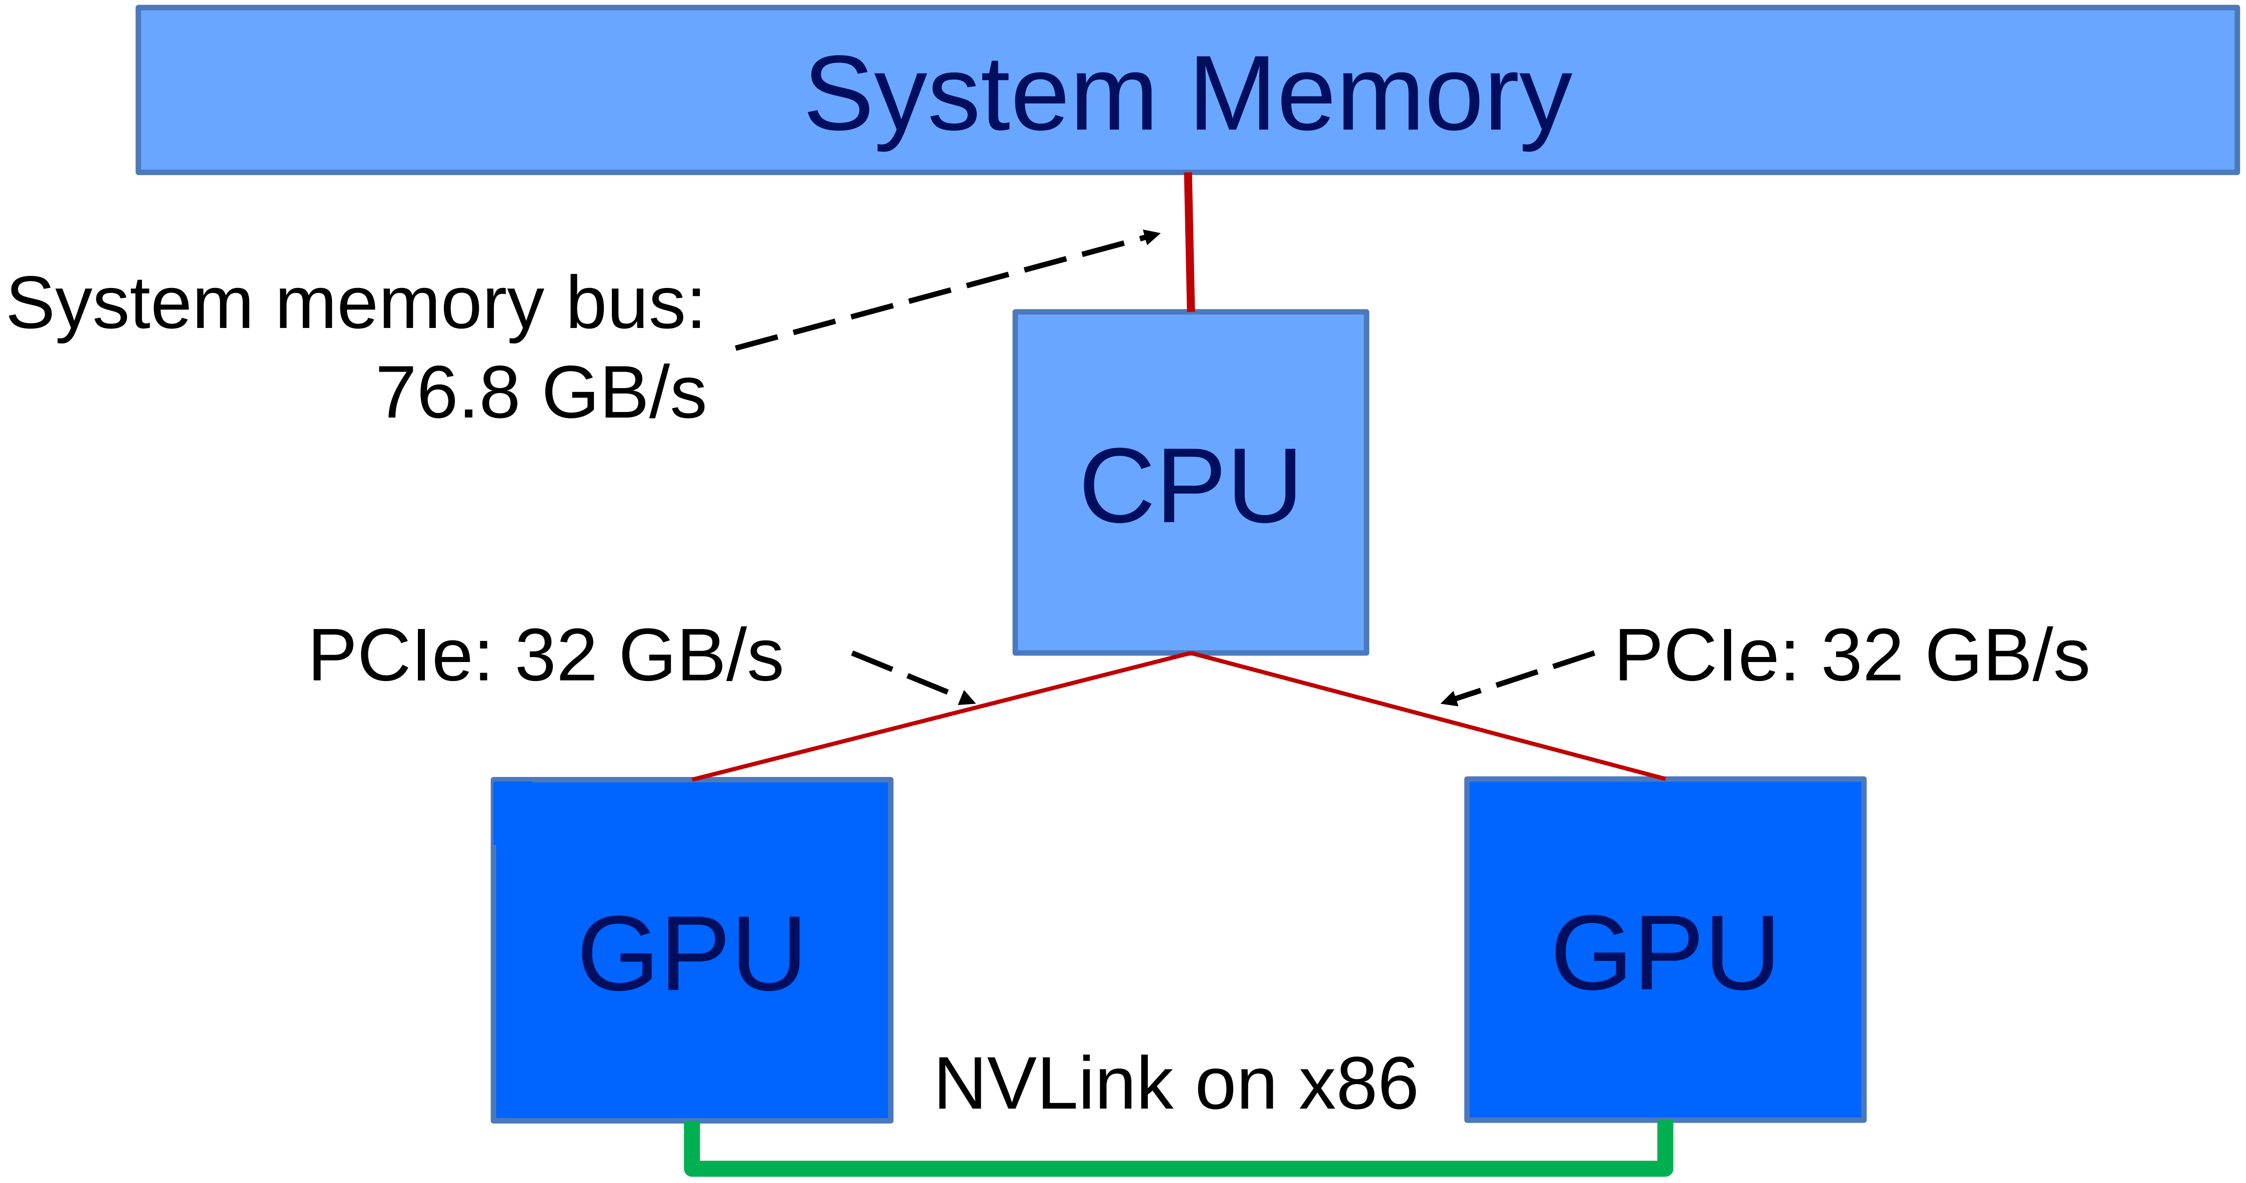

NVLink CPU to GPU connectivity only on POWER9

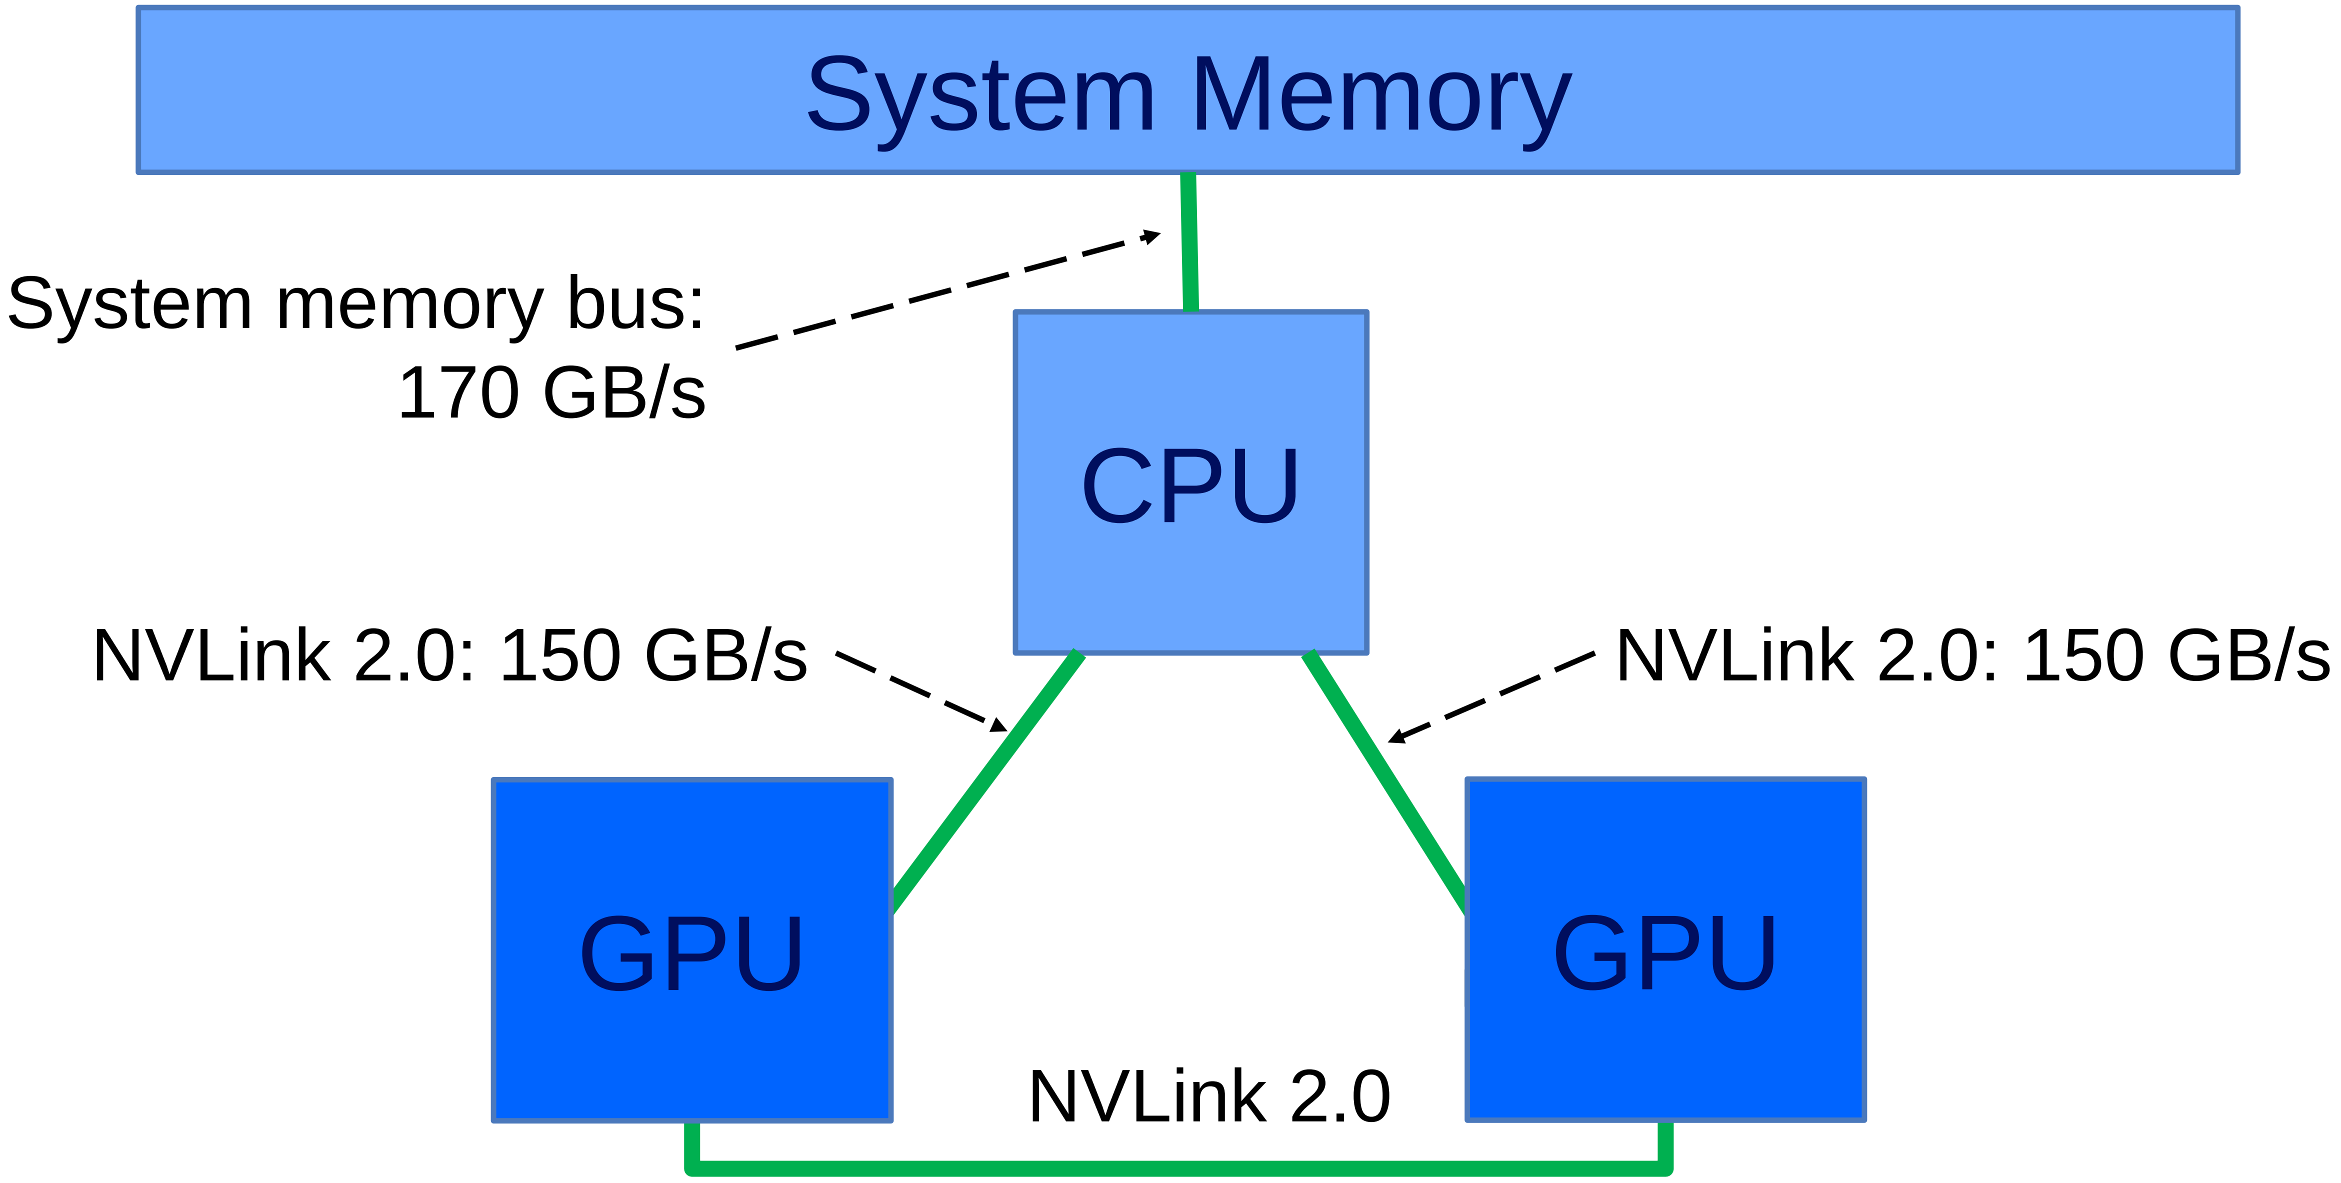

Effects of NVLink 2.0 on Large Model Support (LMS)

DGX1: PCIe connected GPU training one high res 3D MRI with large model support

AC922: NVLink 2.0 connected GPU training one high res 3D MRI with large model support

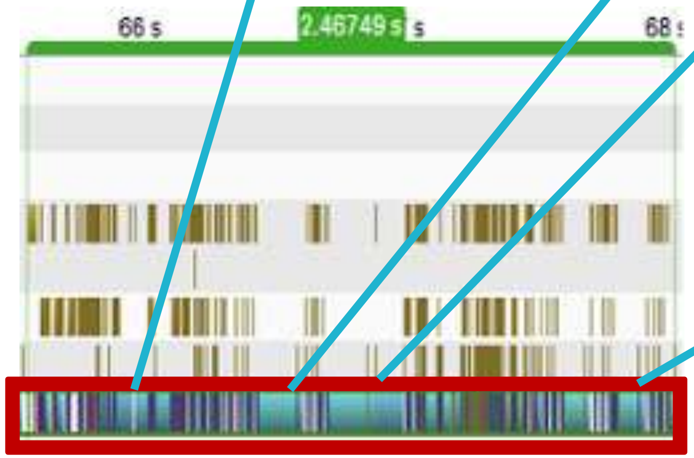

POWER9 with 4 GPUs is **2.1x faster** than x86 with 8 GPUs

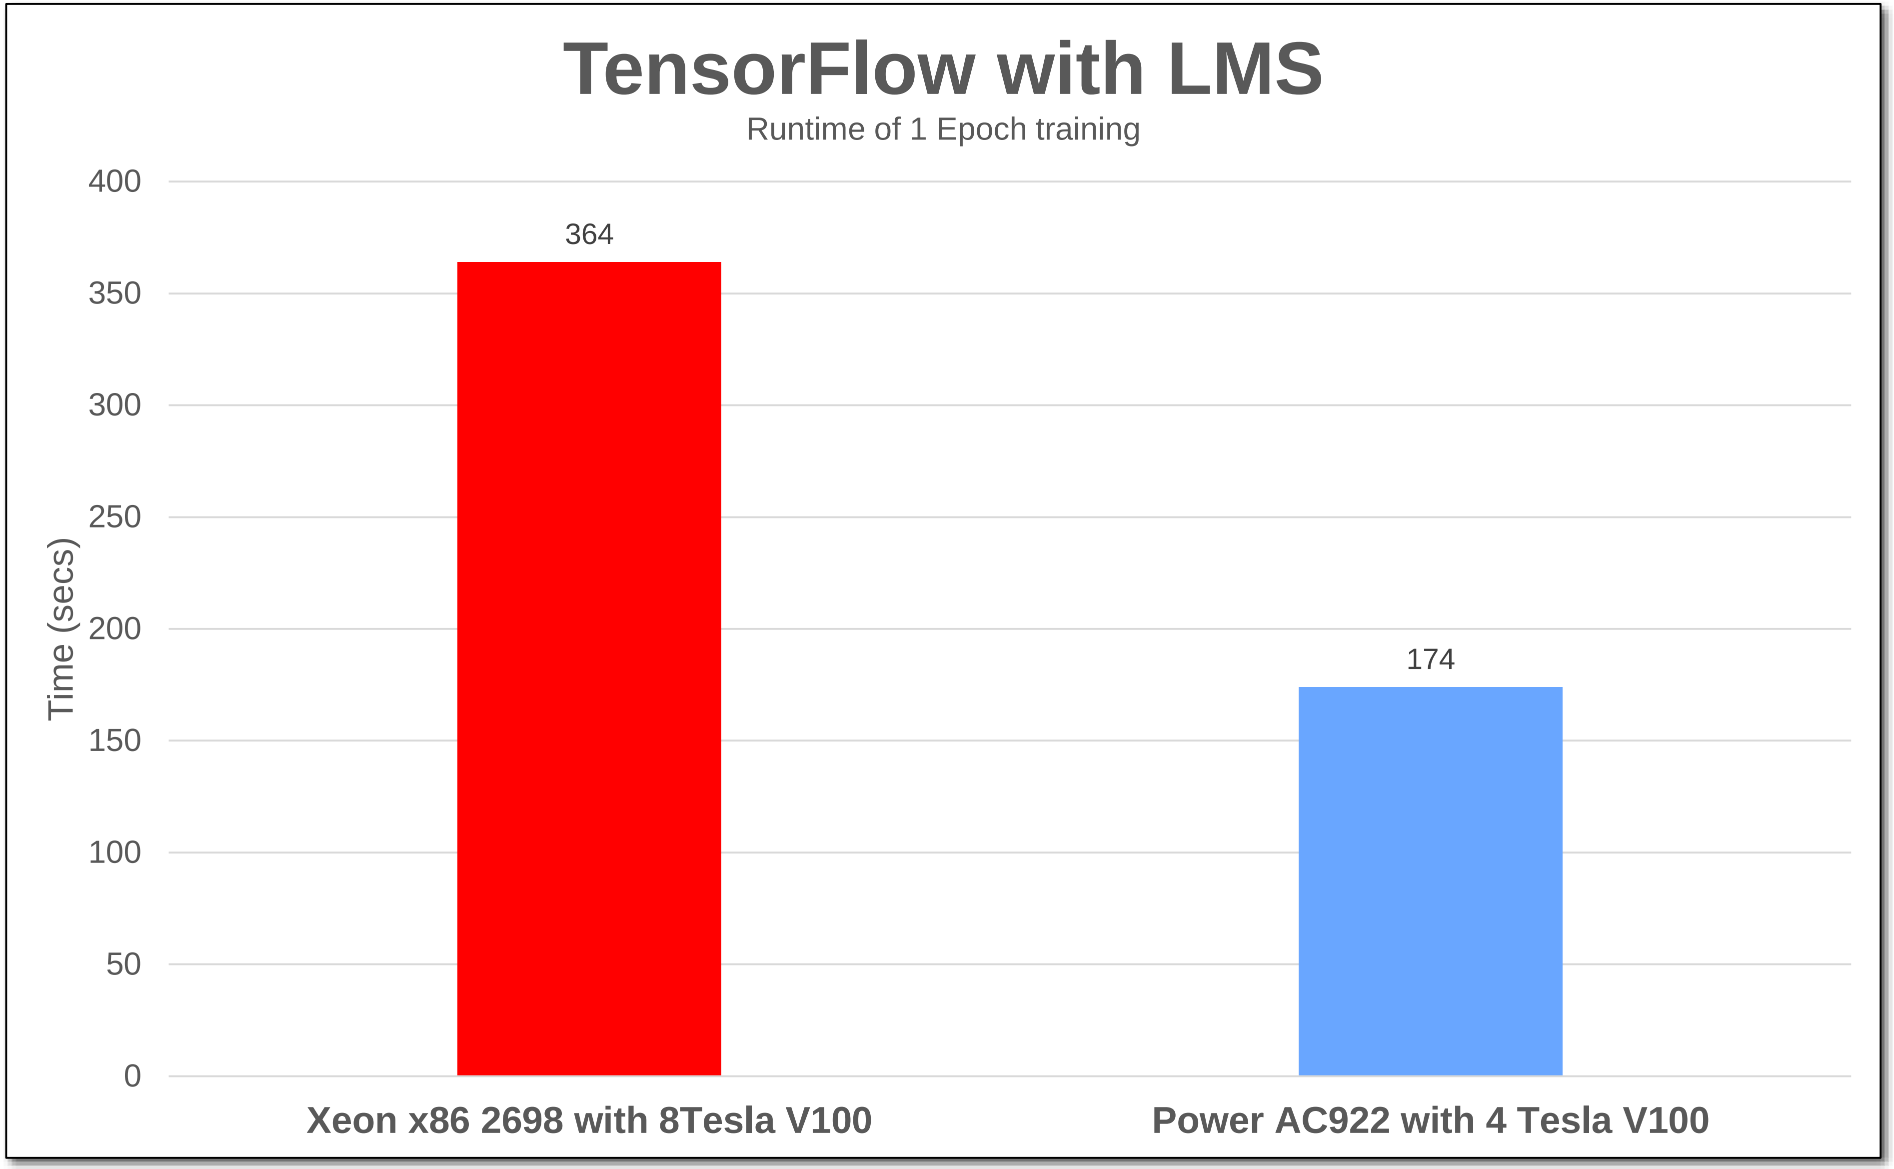

Distributed Deep Learning (DDL)

Deep learning training takes days to weeks

Limited scaling to multiple x86 servers

WML with DDL enables scaling to 100s of GPUs

16 Days Down to 7 Hours
58x Faster

ResNet-101, ImageNet-22K

Near Ideal Scaling to 256 GPUs

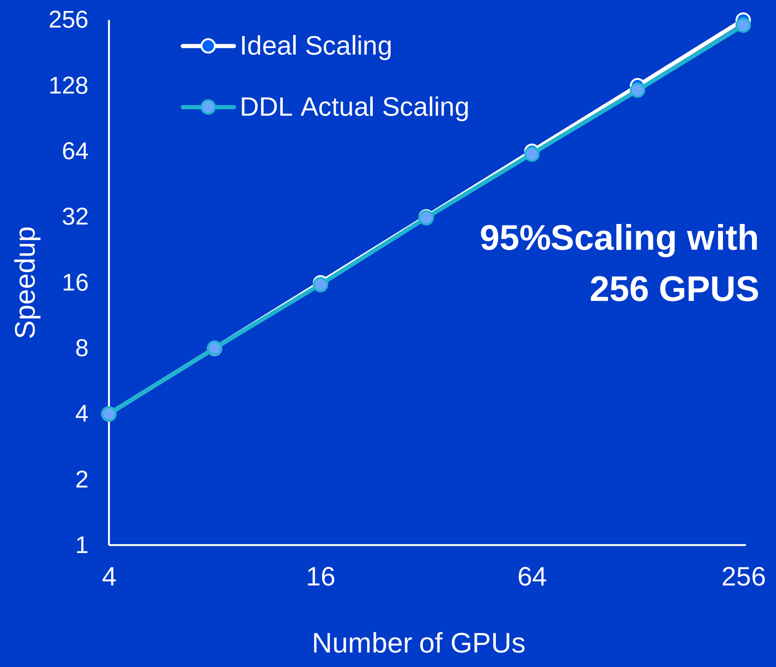

ResNet-50, ImageNet-1K

Caffe with WML DDL, Running on Minsky (S822LC) Power System

Power Systems Value for Data and AI Workloads

Accelerate data scientist productivity and drive insights faster

41%

FASTER Insights

LC922

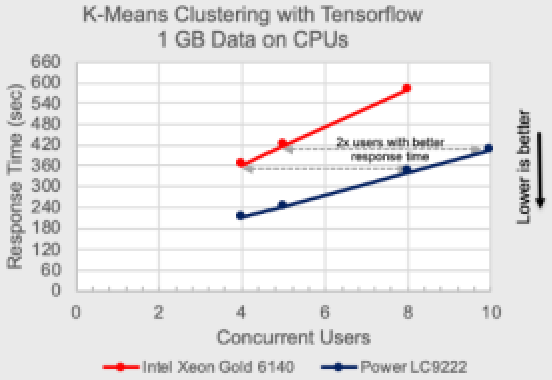

AC922

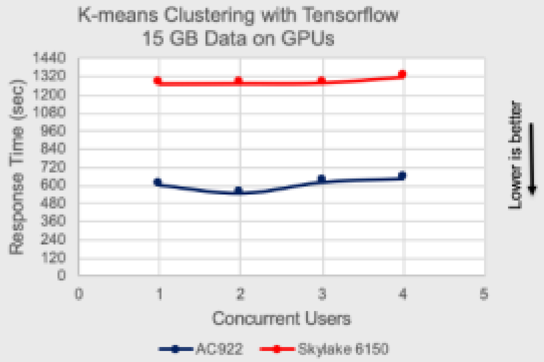

2x

FASTER insights

Greater container density

72 containers | **146 containers**

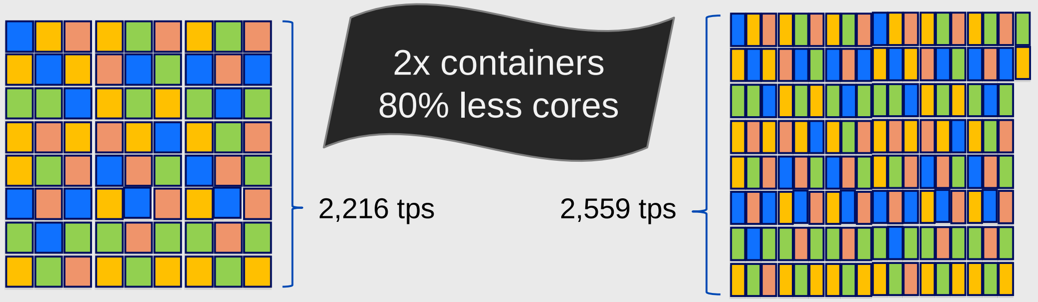

Intel Skylake Xeon
36 cores

POWER9
20 cores

MongoDB -- L922

3.6x

more containers per core

Higher performance for data management workloads

Db2 Warehouse – L922

2.54x

per core performance

IBM Power Systems

The high-performing, reliable and secure foundation for Enterprise AI & ICP for Data

AI

True Cognitive Systems are about co-optimized

software + hardware

which ***just work*** for machine learning, deep learning and artificial intelligence

Infuse – Deploy trusted AI-driven business processes

Analyze - Scale insights with AI everywhere

Organize – Create a trusted analytics foundation

Collect – Make data simple & accessible

POWER9

Modern infrastructure for cutting-edge data & AI workloads

- Enterprise-cloud ready
- #1 in reliability
- Secure to the core
- Unparalleled container density
- Industry leading value & performance

Power Systems:
The bedrock supporting the AI ladder

Simplified Multicloud

IBM Power Systems enable the most data intensive and mission critical workloads in private and hybrid cloud environments.

Scale Performance Affordably

IBM POWER9 processor drives the world’s fastest supercomputers and is ready to accelerate your enterprise.

Delivered with Security

IBM Power Systems have security built-in at all layers, from processor to the OS, designed to deliver end-to-end security.

Proven Reliable

IBM Power Systems ranked the most reliable for the 10th straight year delivering 99.9996% uptime.

Watson Machine Learning Community Edition

Free !!

Curated, tested and pre-compiled binary **software distribution** that enables enterprises to quickly and easily deploy deep learning for their data science and analytics development

- Order it from IBM (support available for a fee)
- Download it from here: http://ibm.biz/download-powerai
- Get Docker container here: https://hub.docker.com/r/ibmcom/powerai/

Includes all of these frameworks:

Available for:

POWER9 With GPUs

NEW x86 64-bit

Snap ML

IBM Spectrum MPI

HDF5

Bazel

ONNX

NVIDIA CUDA

RAPIDS

Caffe Caffe2

BVLC Caffe
IBM Enhanced Caffe
Caffe2

TensorFlow
TensorFlow Probability
TensorBoard
TensorFlow-Keras

PyTorch

H2O Driverless AI Delivers Automatic ML for the Enterprise

- Performs the function of an expert data scientist
- Create models quickly with GPUs and Machine Learning automation
- Delivers insights and interpretability
- Created and supported by world renowned AI experts from H2O.ai
- Award-winning software

21 day free trial for [Driverless AI]

PowerAI Vision: “Point-and-Click” AI for Images & Video

Label Image or Video Data

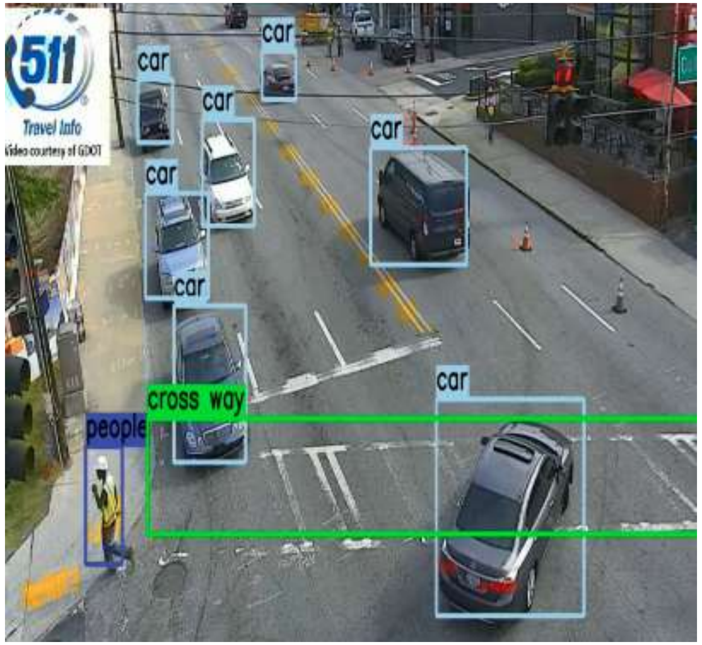

Auto-Train AI Model
(no coding, just point and click)

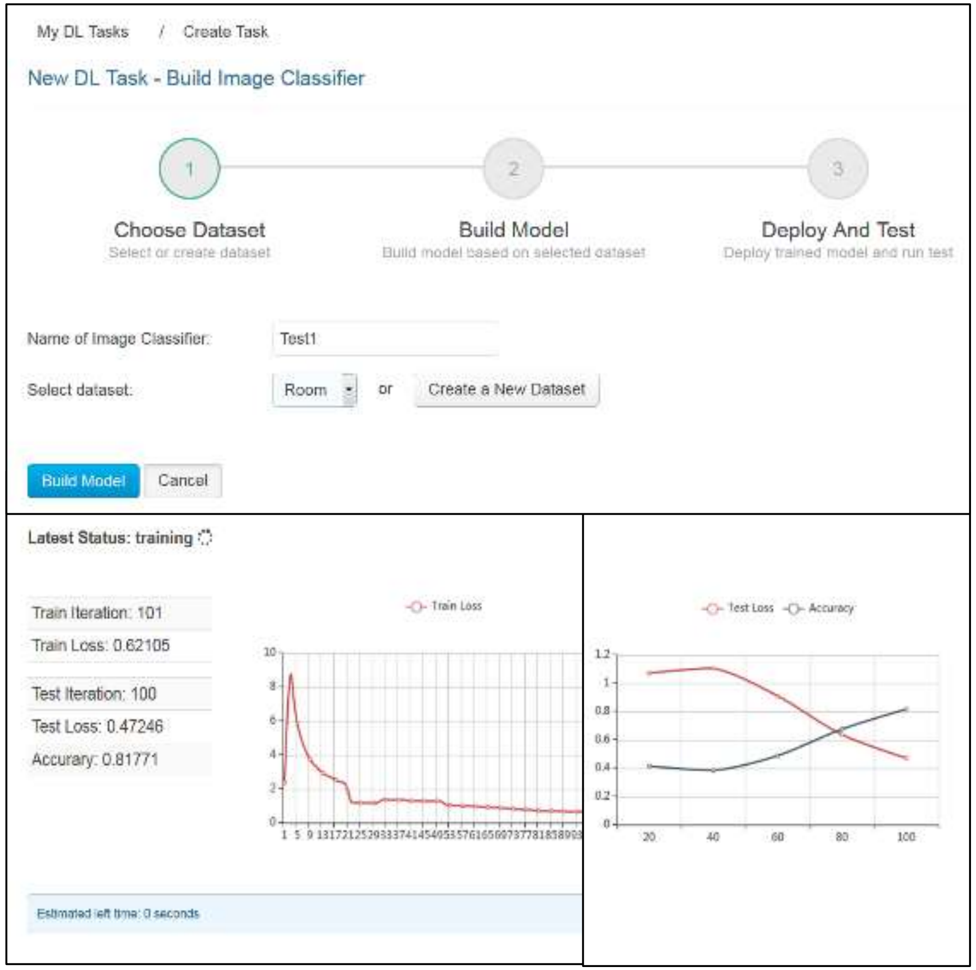

Package & Deploy AI Model

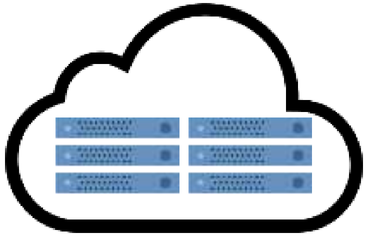


My Data Sets
My Trained Models
My Web APIs
Total: 102, Page Count: 17
1 2 3 4 5 6 ... 16 17 ›
Upload pictures (jpg / png) by dropping them here Or Select some
Tags
PowerAI Vision
NON TECHIES ... CAN
• upload training data from on-site video or cameras
• create detection categories to match the safety use case using business domain expertise
• label data using acquired business domain expertise
• build model & expose it as a RESTful API endpoint for hybrid cloud consumption
• use the model

Watson Machine Learning for z/OS, V2.1

WHAT

An end to end machine learning platform to build, deploy and monitor predictive models

HOW

Integrates and operationalizes AI / machine learning models within transactional systems

WHY

To ensure the lowest latency and the highest performance, security and resiliency to deliver real-time business insight and automation

Watson Machine Learning for z/OS V2.1 Key Features

Flexible model development
Build, train, and evaluate models using your IDE of choice or the WMLz extensive model building features based on enterprise-grade open source software.

Improved productivity
Optimize data science productivity through extensive model building features/modes including notebooks, visual builders, wizards, and enhanced intelligence

Enterprise-ready AI model deployment
Operationalize predictive models within transactional applications, without significant overhead, enabling real-time insight at the point of interaction.

Enhanced model accuracy
Enable data scientists and engineers to schedule periodic re-evaluations of new data to monitor model accuracy over time and be alerted when performance deteriorates. Automatically refresh models to maintain model accuracy with confidence.

Production-ready machine learning
Deliver essential model versioning, auditing and monitoring as well as high availability, high performance, low latency, and machine learning model automation (ML as-a-Service).

Quick-start solution templates
Offer essential foundational templates for common business requirements to bootstrap your machine learning efforts and add value to key business areas including fraud detection, loan approval and IT Operational Analytics (ITOA).

Adopt and Expand AI
Shorten Time to Value with IBM Storage

AI Data Workflow

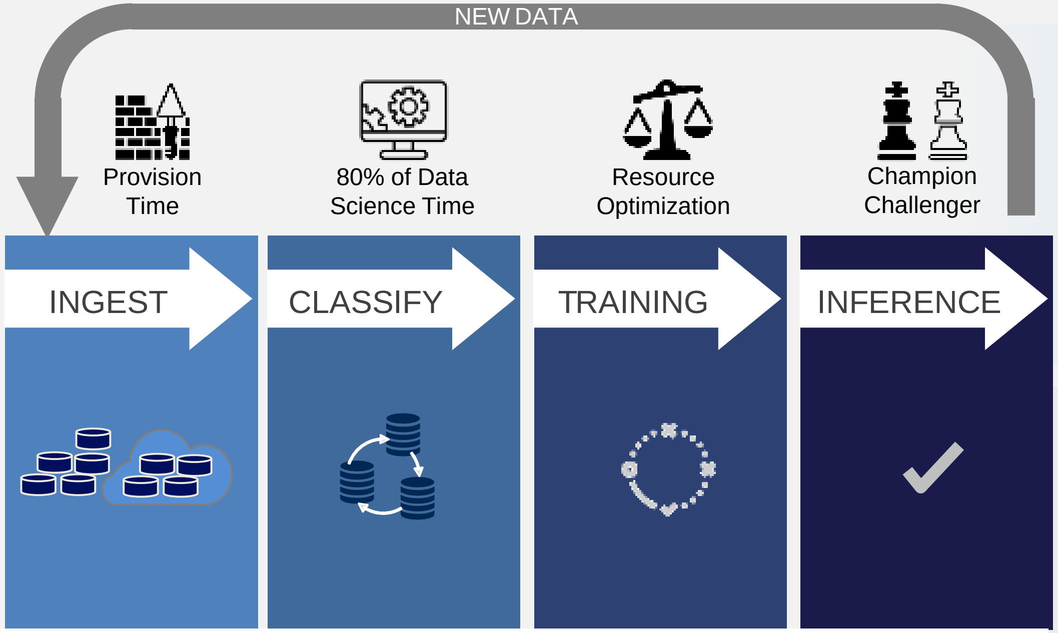

PowerAI

Look for dynamically adaptable, simple, flexible, secure, cost-efficient, and elastic infrastructure that can support high capacity along with high throughput and low latency for high performance training and inferencing

IDC

Business Value

AI Workflow

Industry Standard Approach
- *Deliver consistency and efficiencies*

Uses Technology advances
- *GPU, Open Source Frameworks*

Headwinds Challenge time-to-value

Why IBM?

Data Scientist Productivity
Reduce Time to Accuracy, Improve Provisioning Time, Increase Cycles, Reduce Human Error
- *Improve velocity by getting to your data faster using tools, not trial & error*

Minimize data movement
Increase performance, automate storage processes and reduce cost
- *Using the leading portfolio of Software-defined storage*

Optimized Economics
- *Balance performance and cost with system choices*

Proven Reference Architecture
- *Higher performance, more confidence, lower costs*

Faster Time to Insight

Improve Model Quality

Lower CAPEX

Lower OPEX

Business Agility

Automation Savings

Higher Client Experience

IBM