

Technologie für Big Data

Aktuelle Entwicklungen
und Einsatzbeispiele


Klaus Gottschalk
HPC & AI Architekt
IBM Cognitive Systems




HPC and HPDA towards Cognitive, AI and Deep Learning

HPC Simulation


Oil and Gas



CAE





Material Science



Life Sciences




Cognitive and Deep Learning



The data is trained by 16,000 images and 20 categories with 40,000 iterations accelerated by 88 Spectrum Conductors with Spark

High Performance Data Analysis


Big Data



Business Intelligence



Financial Analytics



Social Analytics



Big Data and AI Examples in Every Industry




Autonomous driving
Accident avoidance



Location-based advertising




Sentiment analysis of what's hot, problems



Market prediction
Fraud/Risk



Experiment sensor analysis



Mfg. quality
Warranty analysis



Clinical trials, drug discovery,
Genomics



Captioning,
search, real time translation



People & career matching



Patient sensors,
medical image interpretation



Drilling exploration
sensor analysis



Consumer sentiment Analysis



Sensor analysis for optimal traffic flows

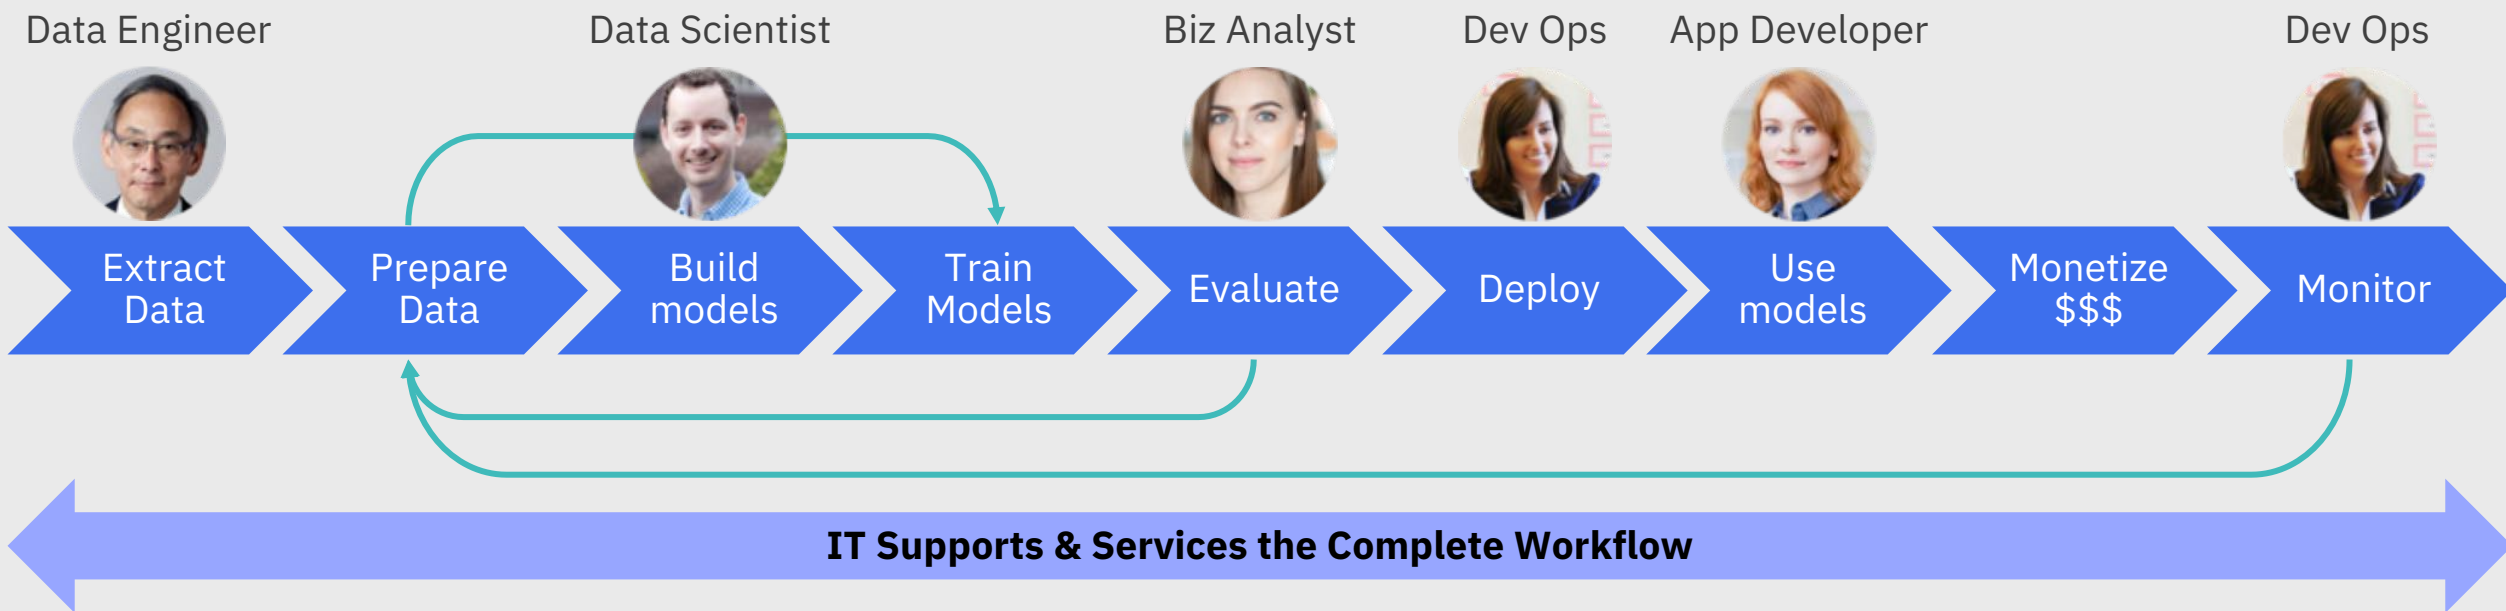


Smart Meter analysis for network capacity,



Threat analysis - social media monitoring, video Surveillance

Data Science is a Team Sport and Iterative

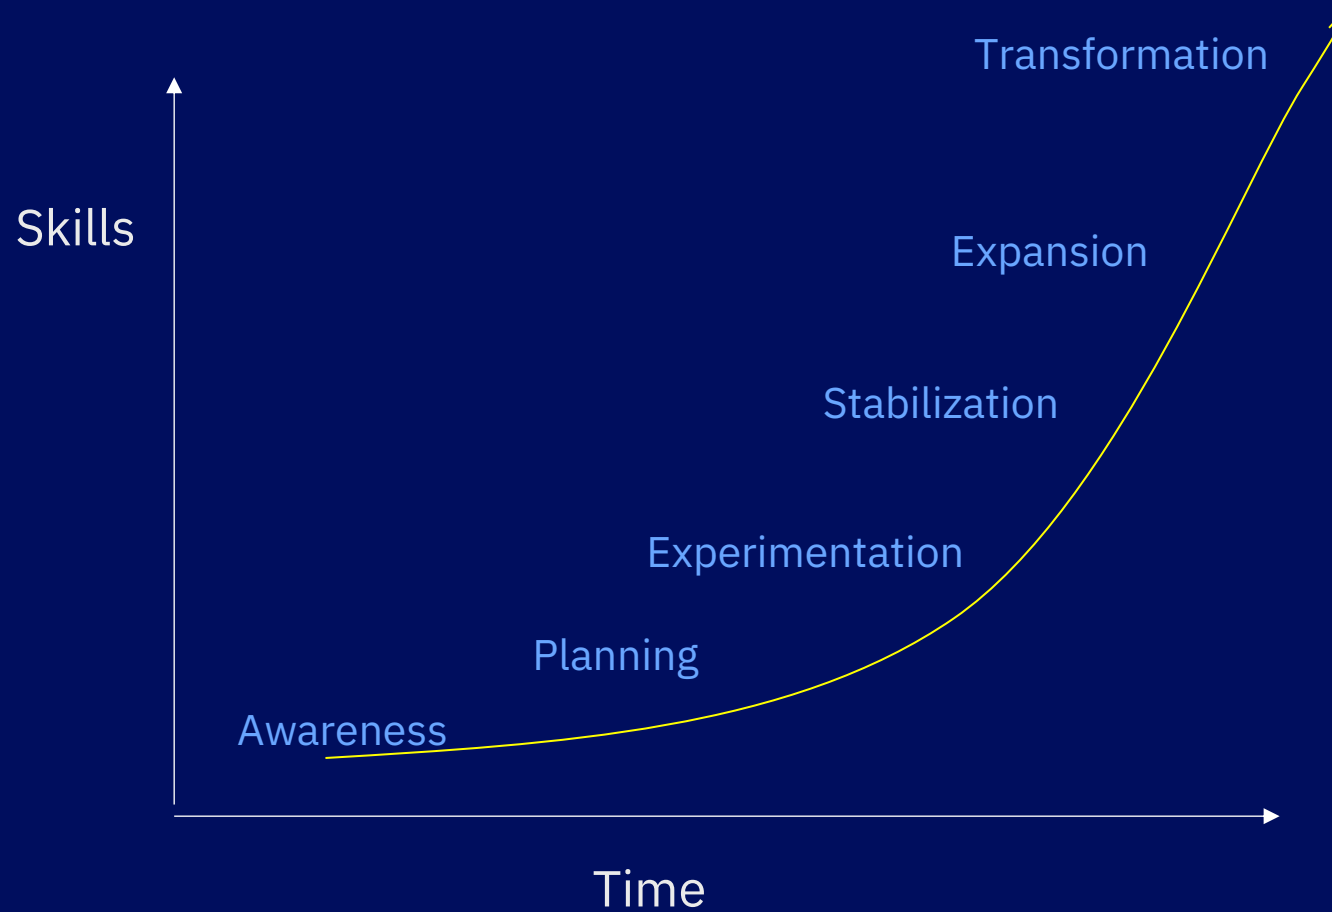


Building cognitive apps using deep learning **requires** multiple skillsets
Connected infrastructure for data, development and iteration.

A common data platform and workflow is crucial for enterprise success.

Data Science starts Small

The first Data Science Project is on a laptop




The goal is to deliver shared services built around shared view of the enterprise.

IBM AI Architecture from Experimentation to Expansion

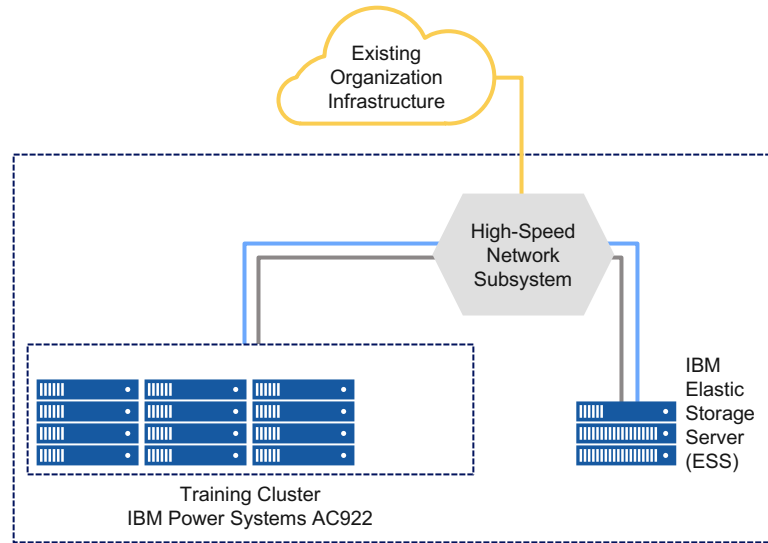
Experimentation

Single Tenant



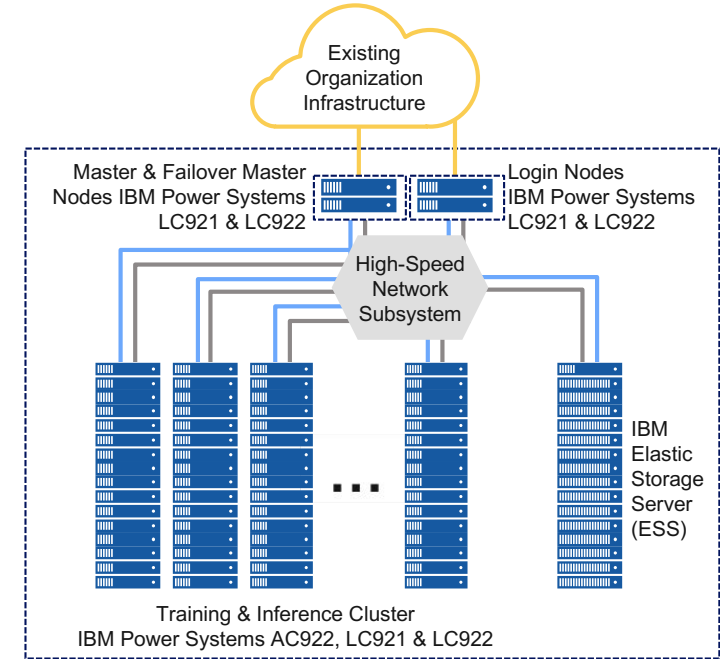
Stabilization & Production

Secure Multitenant



Expansion

Enterprise Scale / Multiple Lines of Business



Services & Support

IBM PowerAI Enterprise

Red Hat Enterprise Linux (RHEL)

IBM Power System & x86 Servers

IBM Spectrum Scale / IBM Elastic Storage Server (ESS)

One software stack from experimentation to expansion

Reference Architecture for AI Infrastructure

From Proof of Concept to Enterprise Scale



Building Block Approach

- Tested, optimized & validated solution
- Integrated tools to facilitate workflow
- Develop according your needs
- Supports the latest Open Source tools

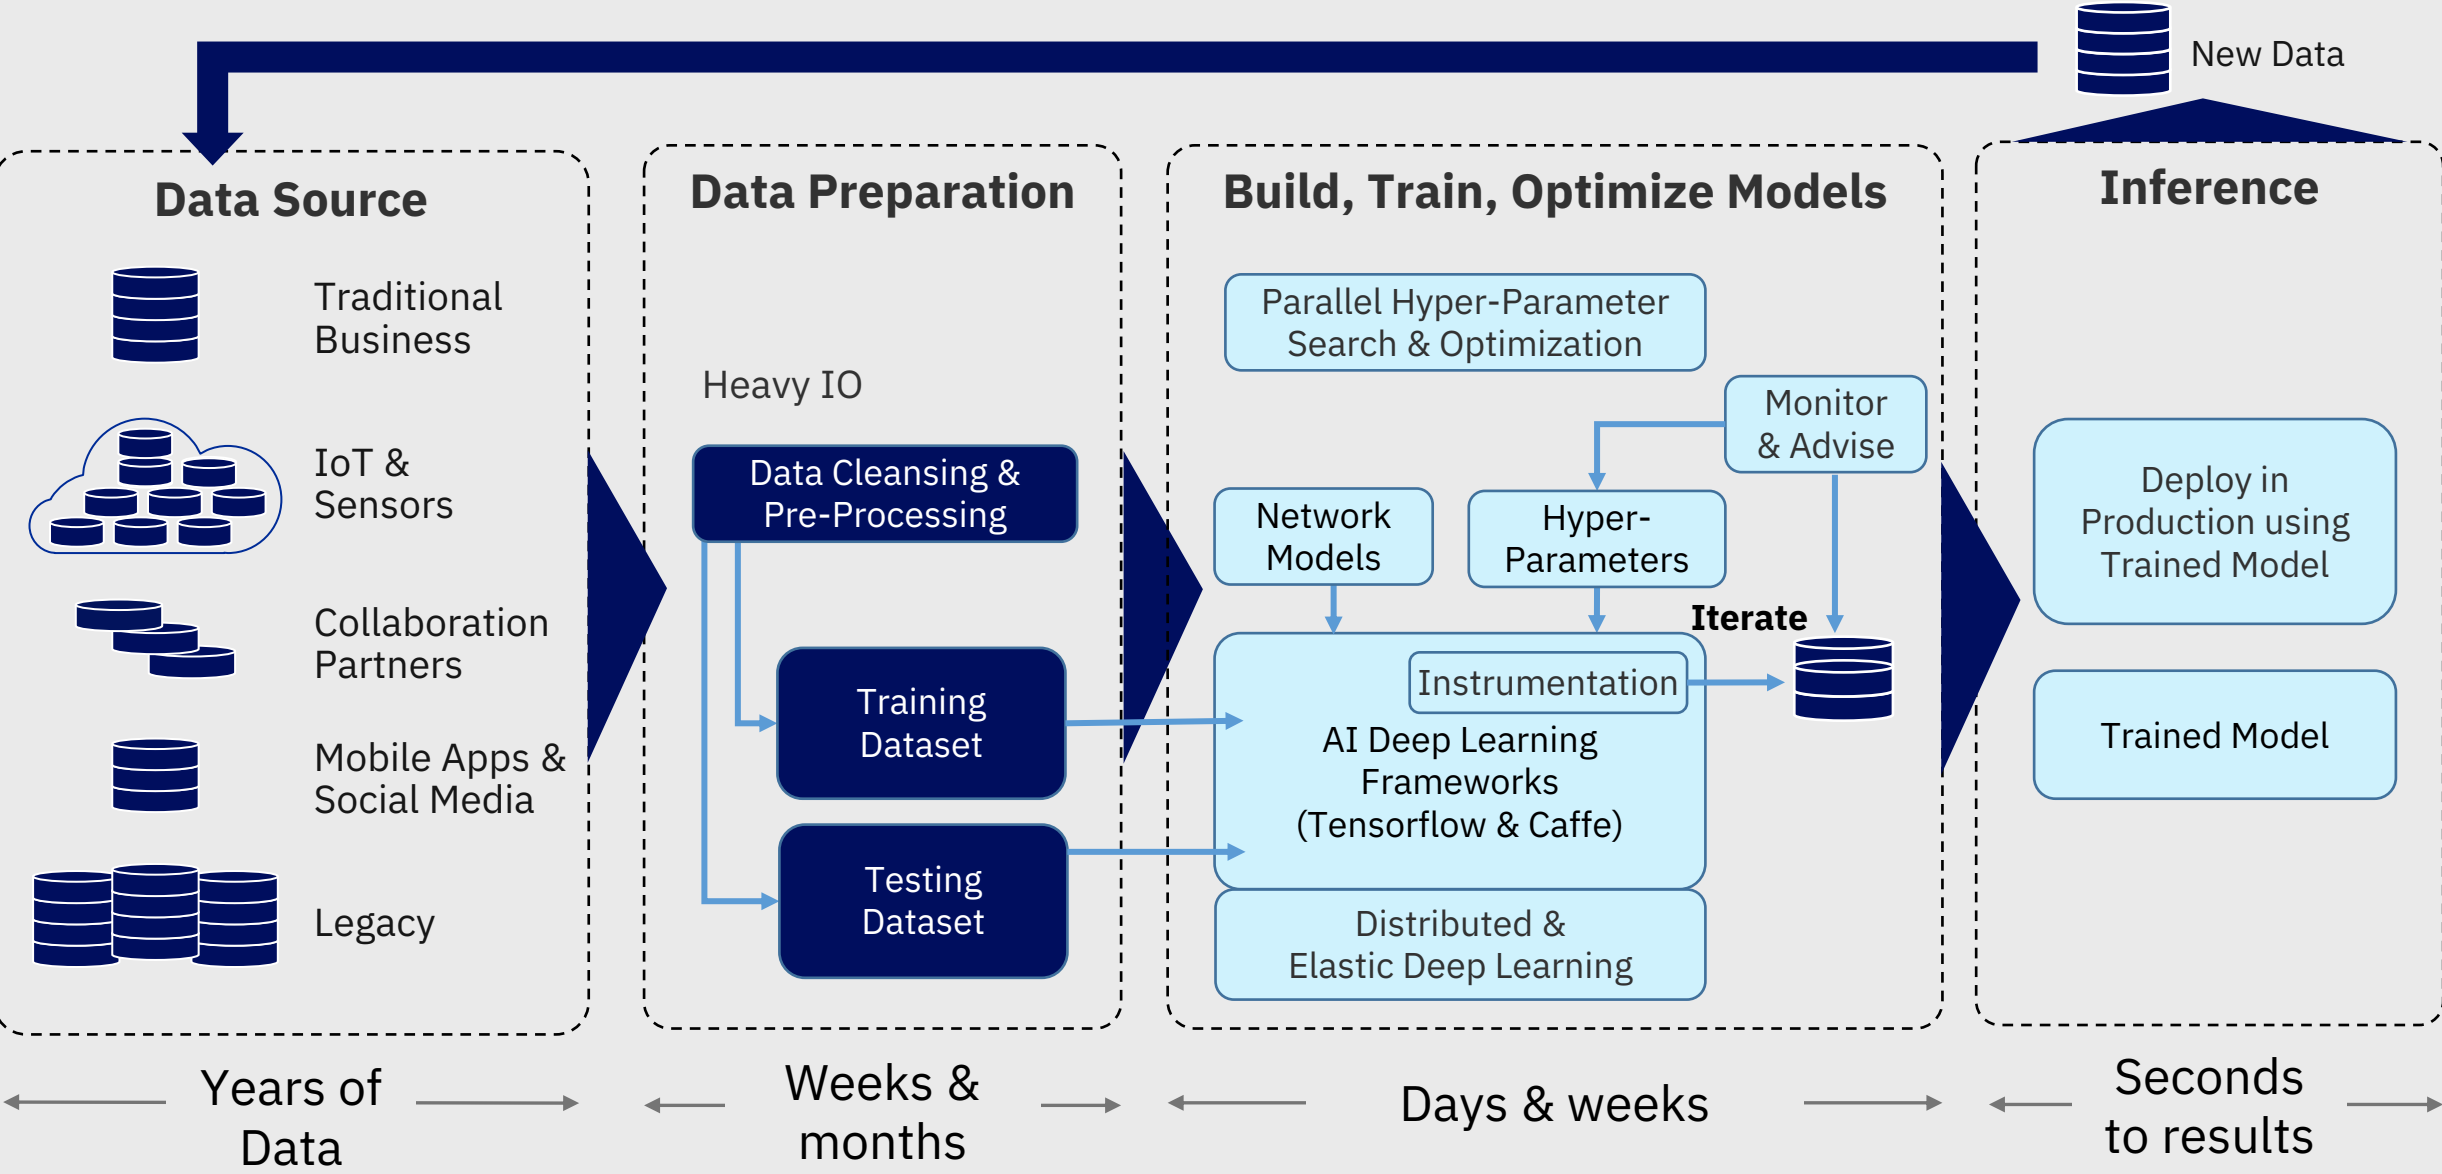
Speeds Time to Accuracy

- Cognitive algorithms for faster model development & optimization
- Reduced training times
- Run-time visualization
- Accelerated hardware

Enterprise Resiliency

- Security
- Reliability
- Scalability
- Ease of integration
- A team to help you implement it

Work flow and data flow is complex



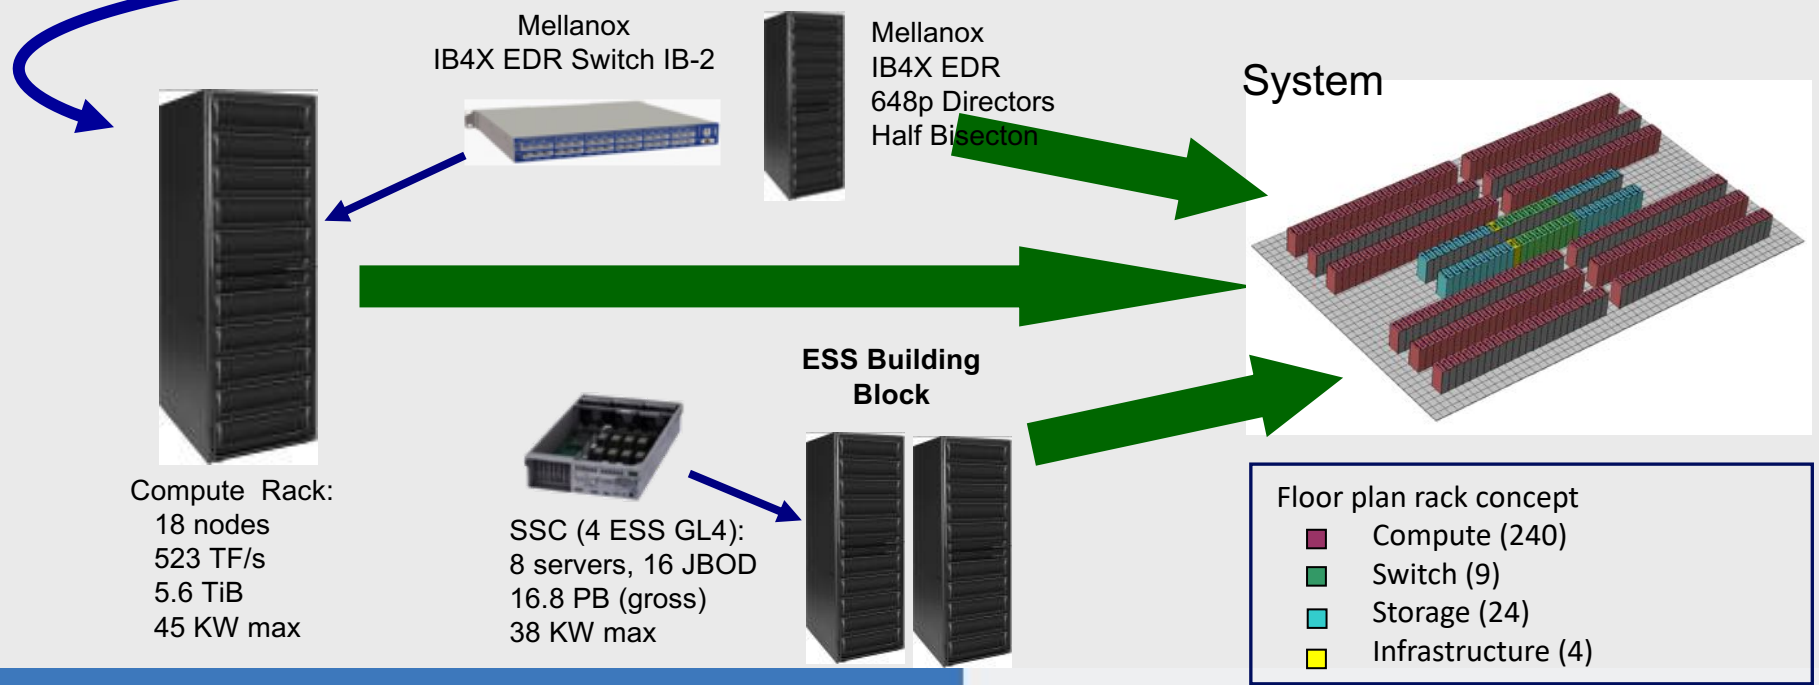
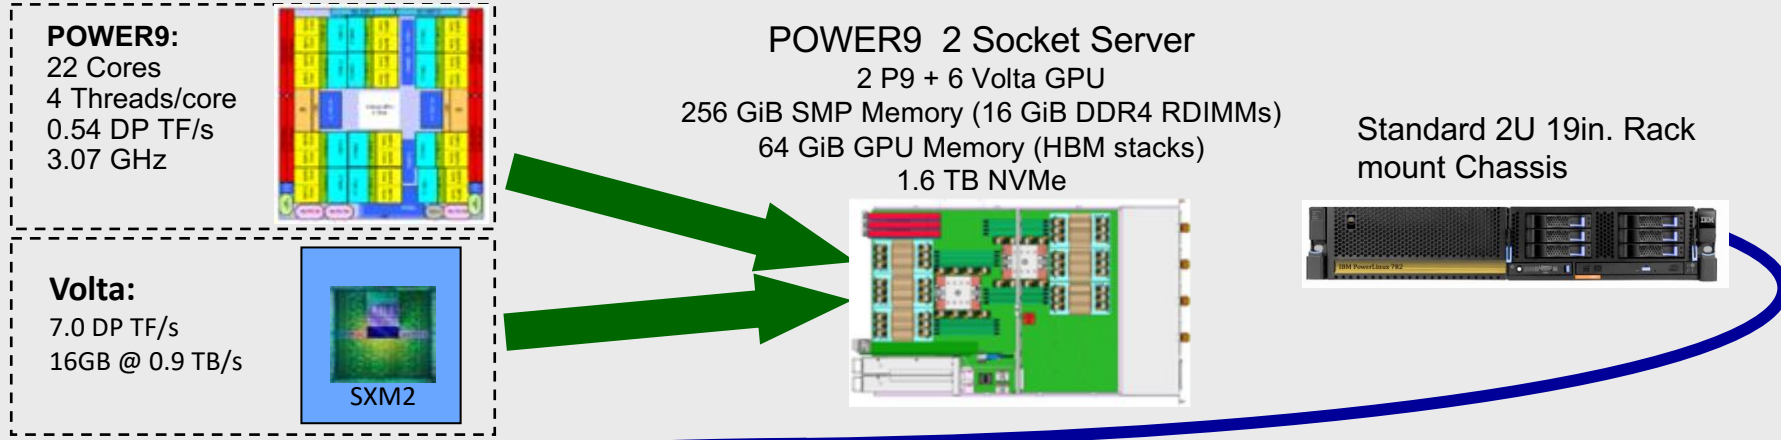
TOP500 List June, 2018



Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	DOE/SC/Oak Ridge National Laboratory United States	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband IBM	2,282,544	122,300.0	187,659.3	8,806
2	National Supercomputing Center in Wuxi China	Sunway TaihuLight - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway NRCPC	10,649,600	93,014.6	125,435.9	15,371
3	DOE/NNSA/LLNL United States	Sierra - IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband IBM	1,572,480	71,610.0	119,193.6	
4	National Super Computer Center in Guangzhou China	Tianhe-2A - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.20GHz, TH Express-2, Matrix-2000 NUDT	4,981,760	61,444.5	100,678.7	18,482



CORAL ORNL 200 PF System



CORAL: ORNL's Summit System



Overhead cooling distribution

Based on real world experience

Different workloads, but built with many of the same building blocks

Wells Fargo: Financial Risk Modeling

Using AI to enhance financial risk models and provide validation to meet regulatory requirements and business goals.

Automotive Sensor IoT: Transforming data from the edge to useful insights

From global data to insight, they manage large data as objects, extracted to run AI.

Top 5 Global Bank: Building a better client profile using Spark and AI

Managing multi-platform data ingest with distributed computing and ML/DL to normalize, clean and tag data to build client behavior profiles .

IBM Global Chief Data Office: One Common Enterprise Data Backbone

The backbone at the core of every business process for a single version of the truth, providing data, computing, analytics & AI.

CORAL: National Lab Supercomputers built for AI

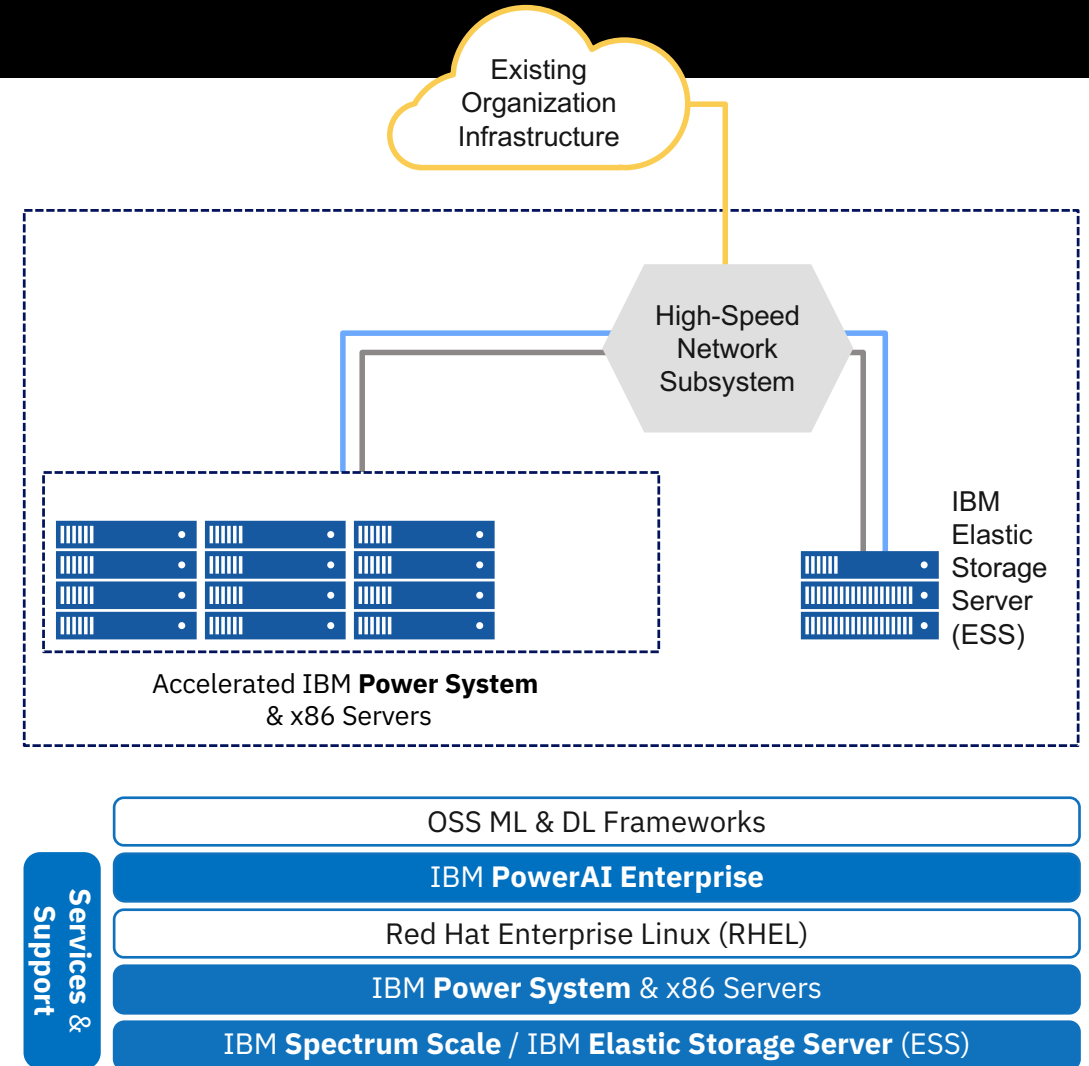
The most powerful and smartest supercomputers in the world, and purpose built for AI workloads.

A Reference Architecture for AI Infrastructure

IBM Systems

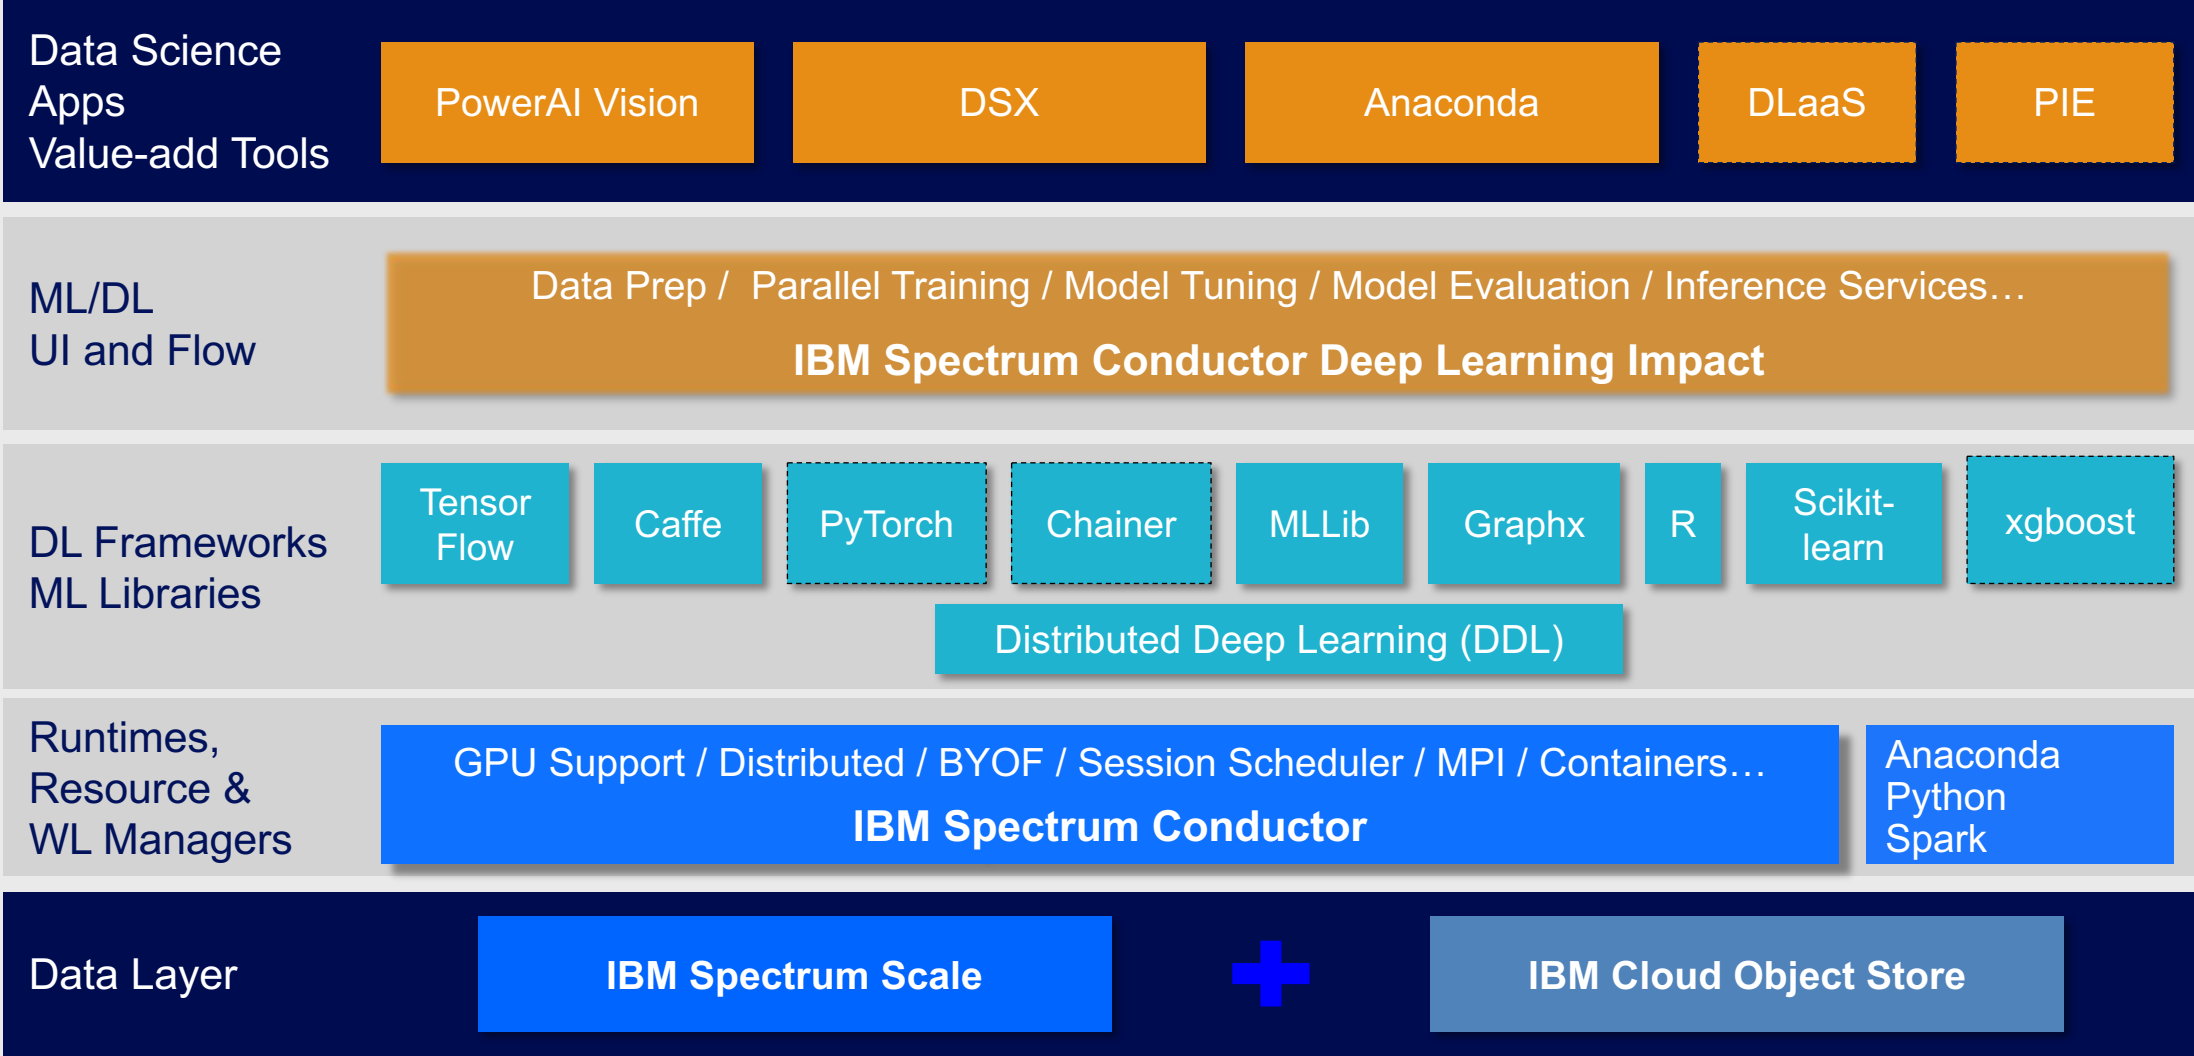
Supports you AI journey from PoC/experimentation into production & then scale-out to support your organization/enterprise. Based on the extensive work done with our clients building their AI environments

- Reduce complexity, time & risk of building & running
- Improves data science efficiency & productivity
- Building block approach
- Speeds time to accuracy
- End-to-end AI workflow support - from data ingestion & preparation through building, training & optimizing models, & into production & inference
- Enterprise resiliency
- Integrated HW & SW solution based on opens source & IBM solutions
- IBM services & support



2018 Reference Architecture for AI Infrastructure: Software

IBM
PowerAI
Enterprise



Thank You !

AI Discovery
Workshops

POC Projects with
IBM Spectrum
Computing

Data Science Elite
Team

Planning your
Storage
Infrastructure