Common Romandie

IBM Cloud Private et la plateforme IBM Power

Thibaud Besson

IBM Systems France

thibaud.besson@fr.ibm.com

11 octobre 2018

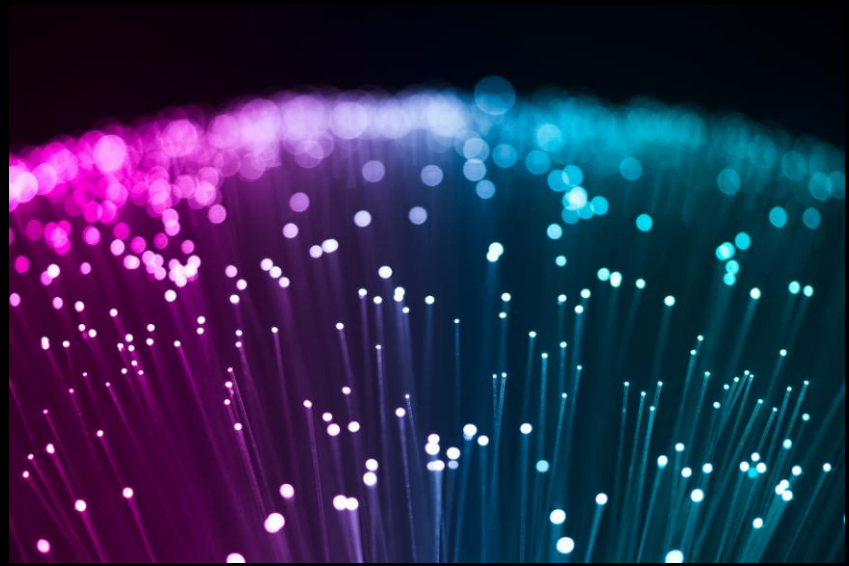

La révolution des données redessine l'informatique

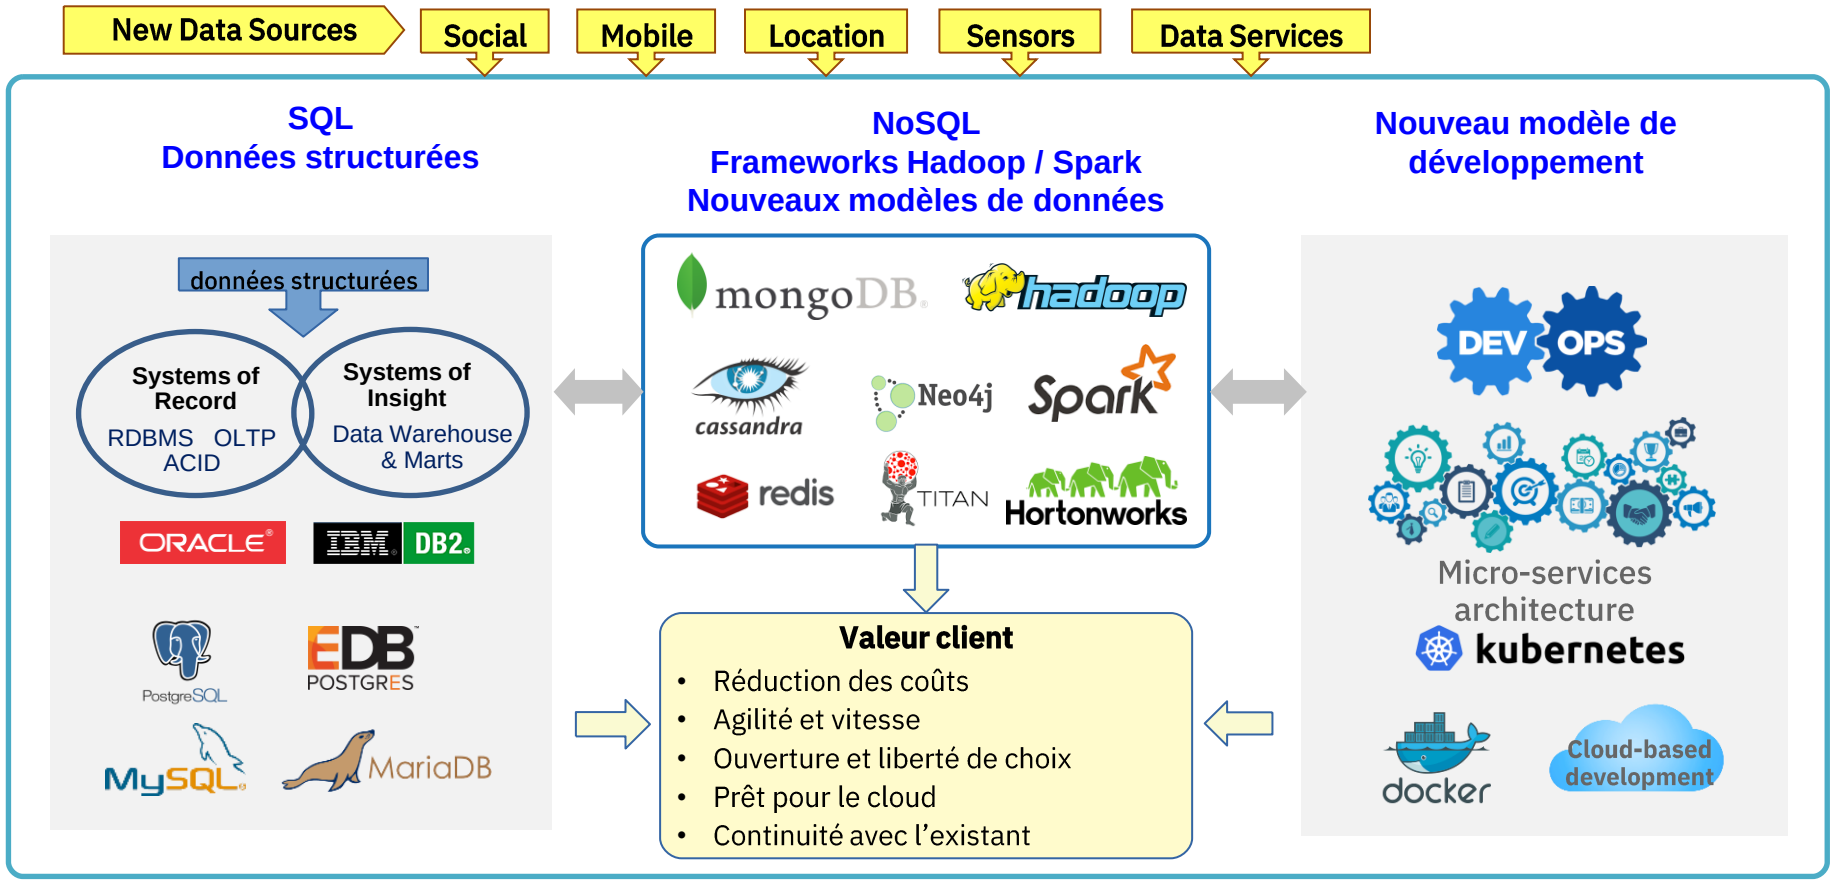

Plus de 20 ans de contribution IBM à l'Open Source

1997 → Aujourd'hui

Initiateur de **80 projets Open Source** et contributeur dans plus de **150 autres**

900 ingénieurs dans le Linux Technology Center et 7000 au sein d'IBM

Versement de logiciels complets : Jikes , HttpServer, Eclipse, CloudScape, PerfHarness, SystemML…

Communauté OpenPower (300 entreprises membres)

Créateur de la Common Public License

Top Contributeur au noyau Linux

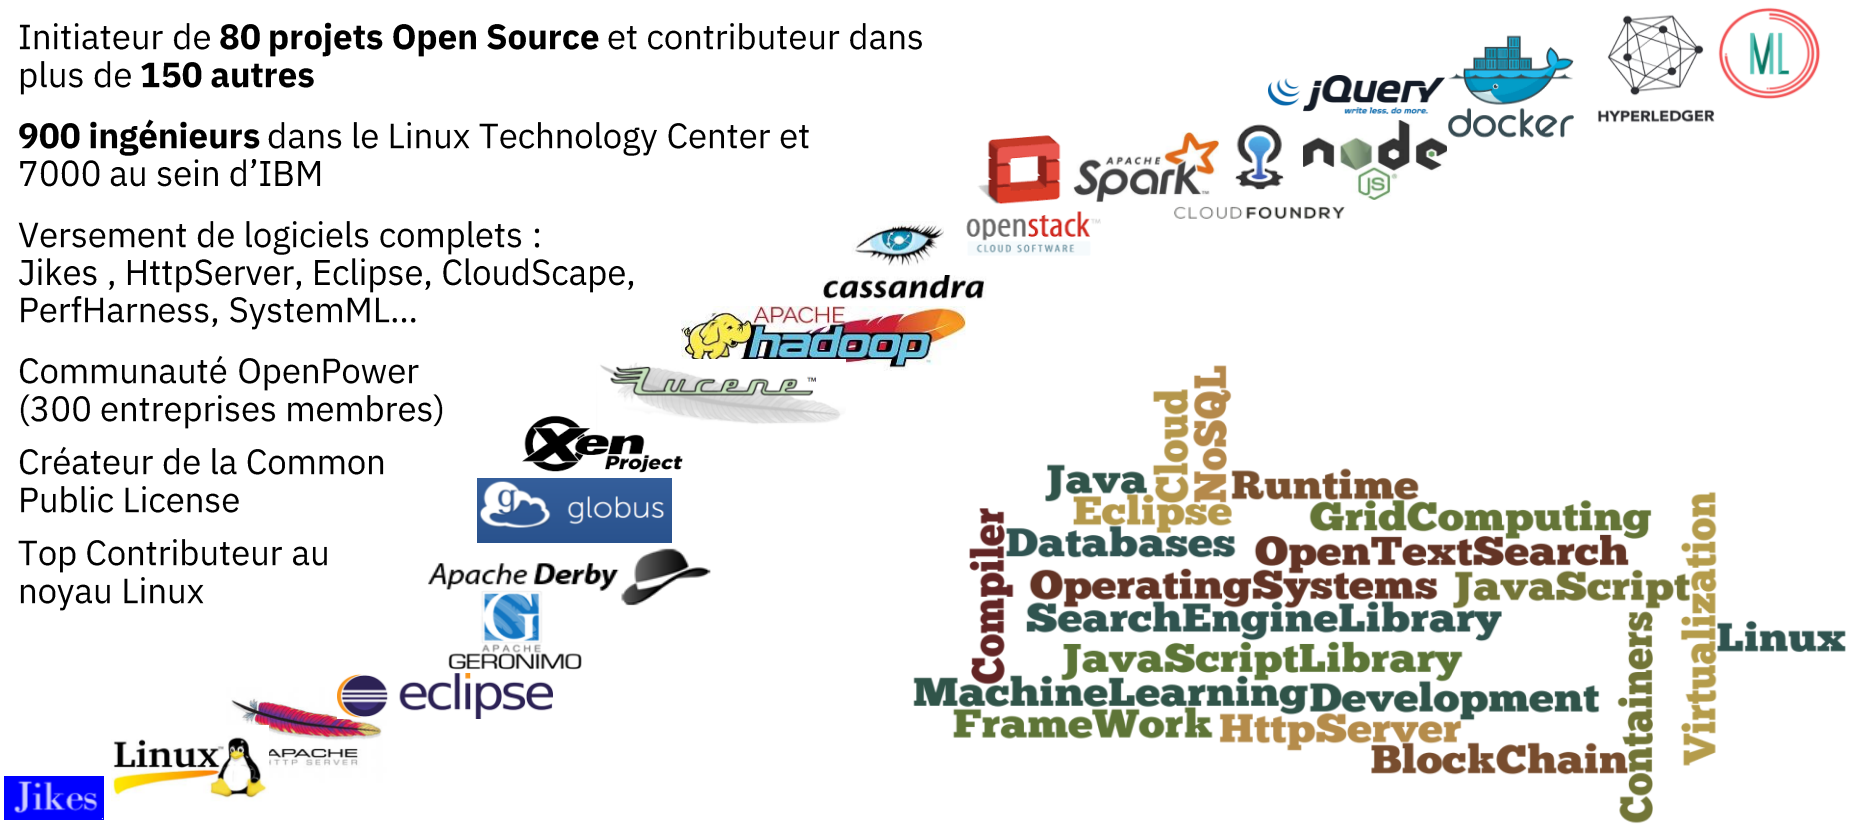

Le modèle économique de l'Open Source et la valeur des distributions

- OPEN :
  - l'accès au code source est garanti
  - Les formats sont ouverts
  - Le code source peut être modifiable ou réutilisable

- Pas toujours FREE :
  - Gratuit : CentOS, postgres, MongoDB Community Edition, etc.
  - Via souscription : RHEL, EnterpriseDB, MongoDB Enterprise Edition, etc.

- Quelle est la valeur des éditeurs ?
  - Solution clef en main (intégration)
  - Fonctions avancées
  - Service
  - Support
  - influence sur le développement du produit -> besoins clients

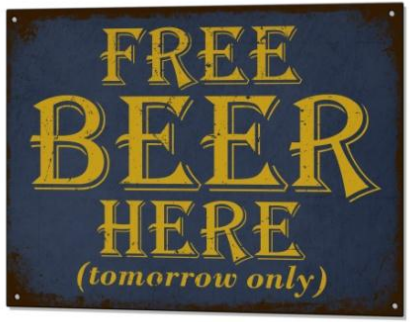

Construire une solution à base d'Open Source peut ressembler à ça :

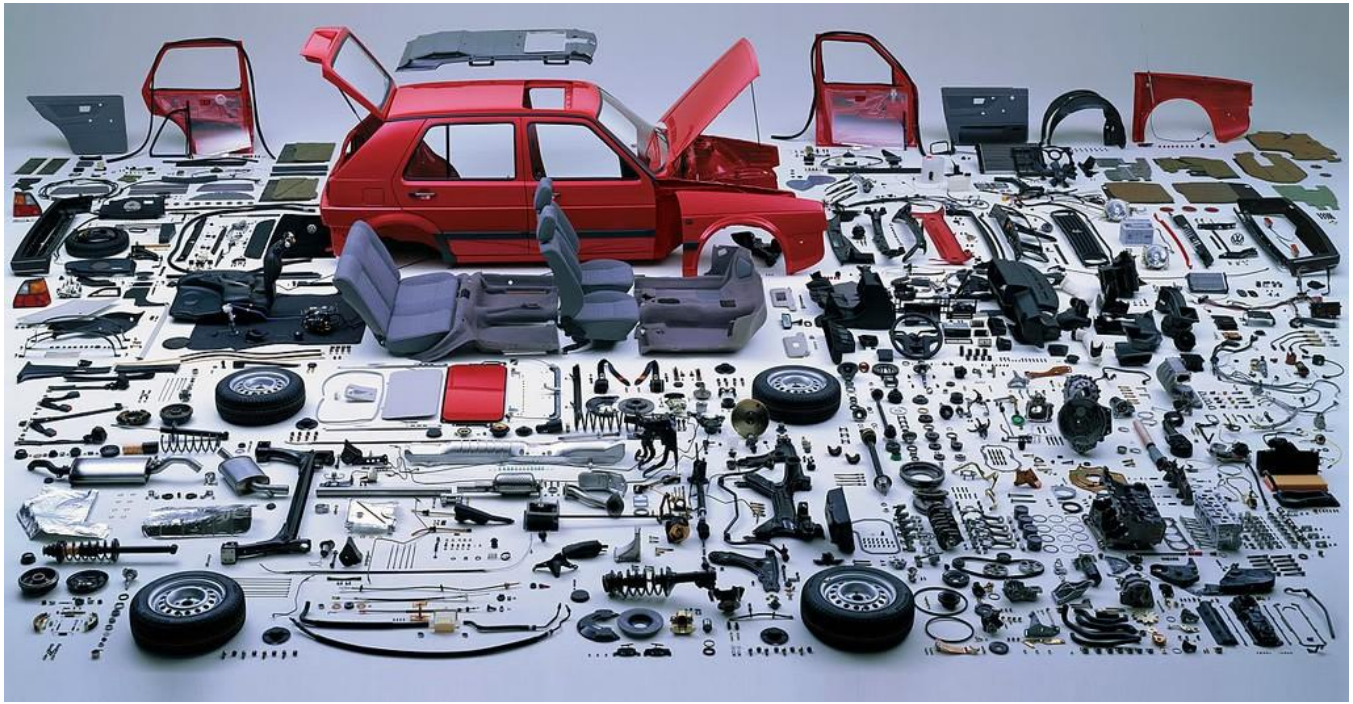

...Pourquoi pas, si vous aimez la mécanique !

Exemples de support IBM intégré pour l'open source

Par exemple

Power with Cloud Private

It's a heterogeneous world – When to leverage Power Systems for Private Cloud?

Examples of Power Benefits

Application & Infrastructure Modernization

Modernize and optimize existing applications with high performance access to enterprise data

Run on same CPU as AIX and IBM i apps for fastest data access

New Cloud Native Applications

Create new cloud-native applications with Open Source apps and databases

Use Open Source DB's with 2X Performance

Deep Learning

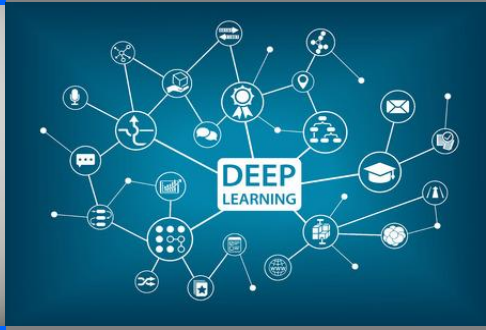

Create new ML/DL apps with PowerAI, DSX, Open Source tools

Train 4X+ Faster on Power9

POWER9 Systems family – run linux everywhere !

OpenPOWER™

IBM

Power S922LC
Power AC922
For Cloud / Cluster
Artificial Intelligence

Power L92x
1-2 sockets, 2U

Power S9xx
1-2 sockets, 2U-4U

Power E9xx
4-8 sockets
Enterprise

Small cores

Big cores

Linux

Linux

Linux AIX i for Business

Linux AIX i for Business

KVM

PowerVM

IFLs

Hyperviseurs sur POWER

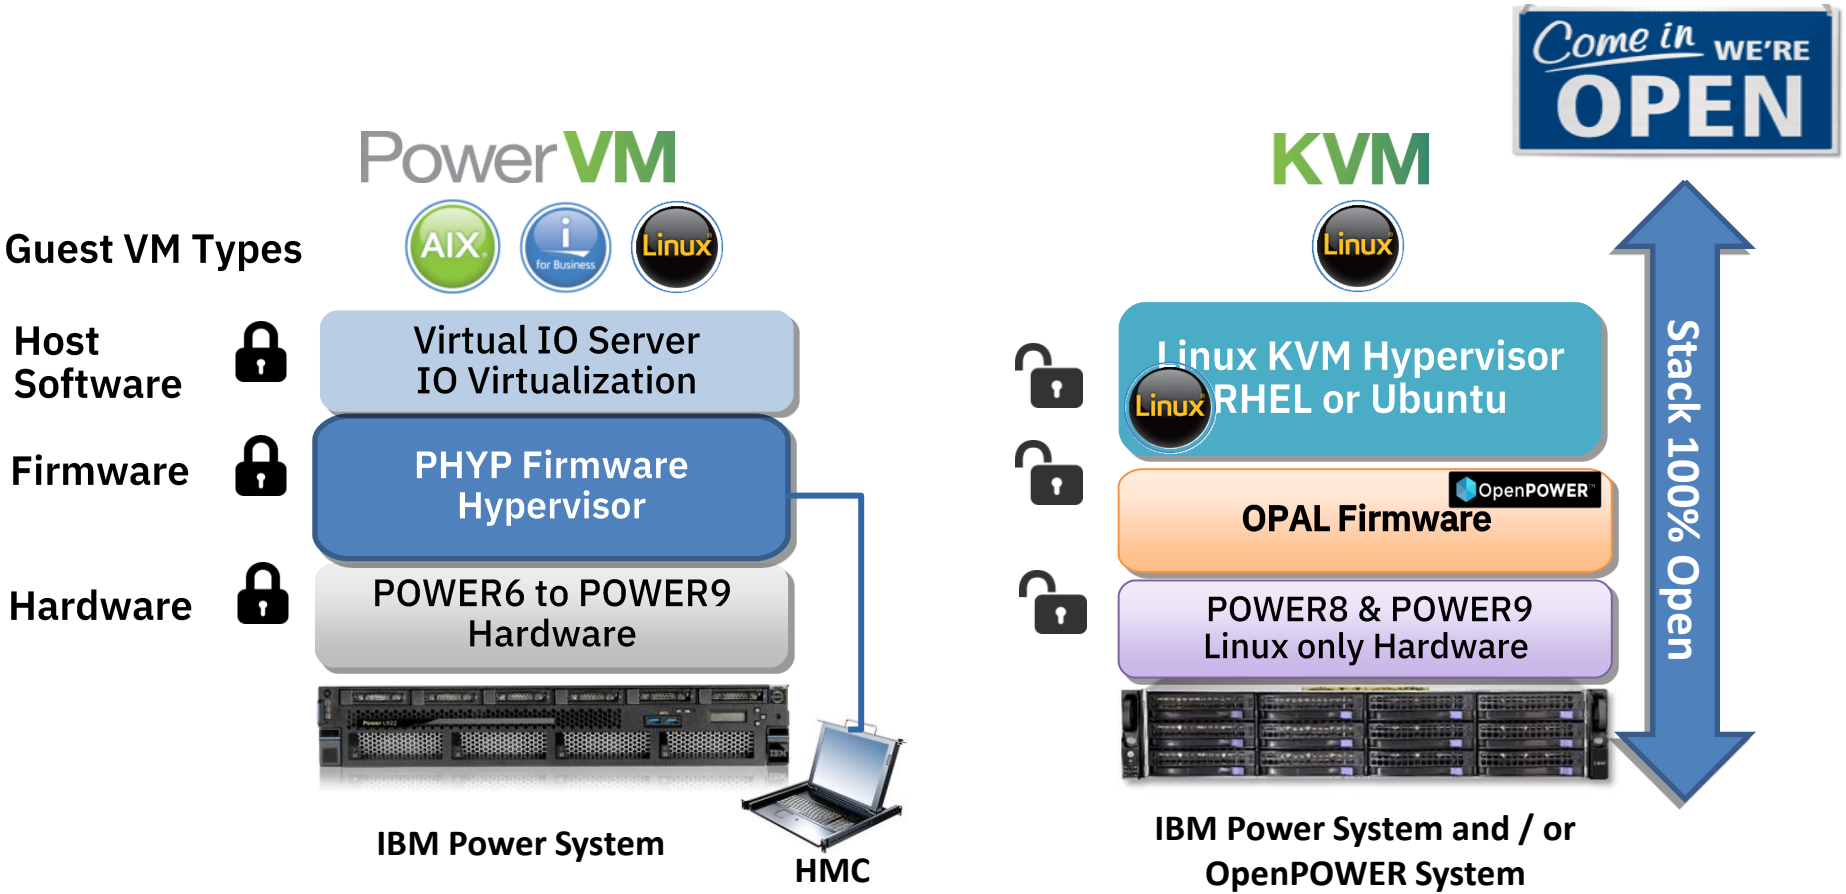

Getting Money in IT

Gartner: Relational OSDBMS has matured and today can be considered as a standard infrastructure choice for a large majority for new enterprise applications.

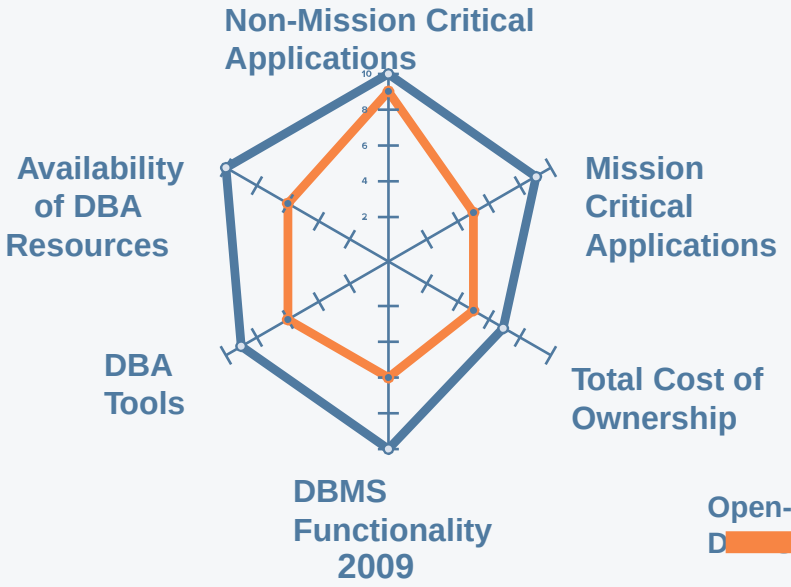

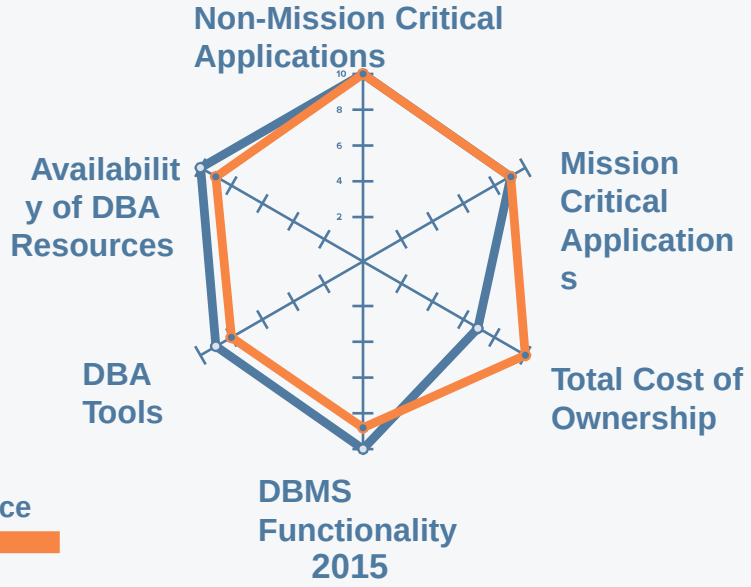

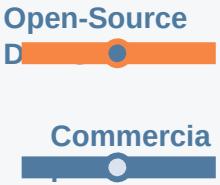

Figure 1
Relational Open-Source DBMS Maturity Evaluation, 2015 Source: Gartner (April 2015)

Figure 2
Relational Open-Source DBMS Maturity Evaluation, 2015 Source: Gartner (April 2015)

Battle 2: High-Density consolidation test

"Get the maximum number of VMs running pgbench while average response-time is < 1ms"

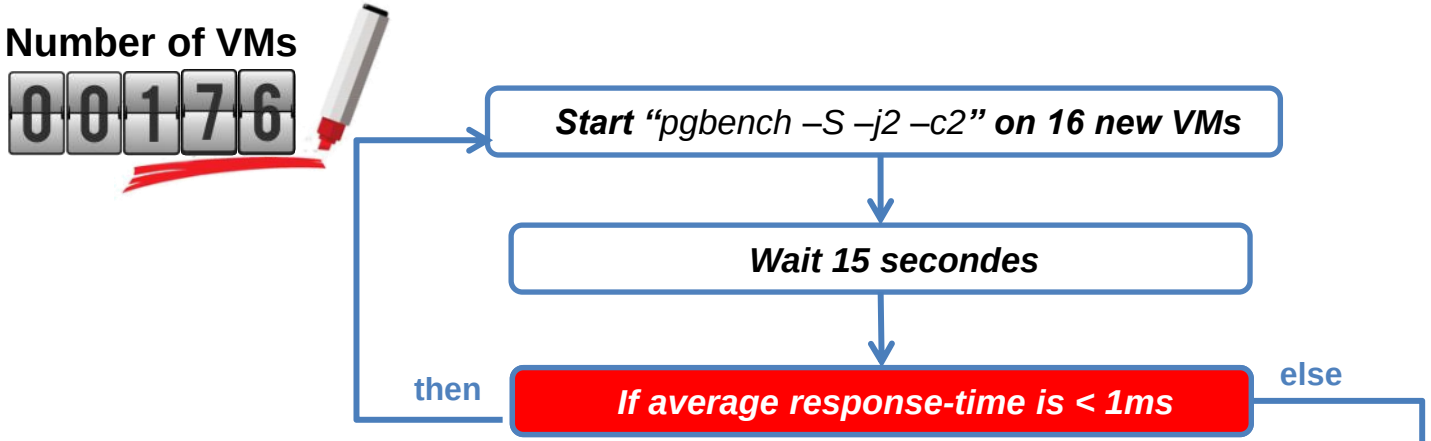

After 5min of stabilized run, the server who runs **the more VMs** wins !

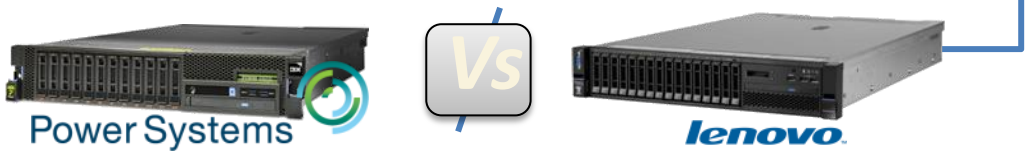

Battle 2: High-Density consolidation test (result)

IBM Power System wins :

x2 more VMs per system

x2.8 more VMs per core

POWER > 12 VMs / core
INTEL = 4.5 VMs / core

Compare to a single all-cores VM test :

POWER keeps **85%** of total system tps compare to a single VM performance test. (600k vs 700k)

Intel drops to **63%** of total system tps compare to a single VM performance test. (300k vs 470k)

YouTube http://ibm.biz/BdrhE9

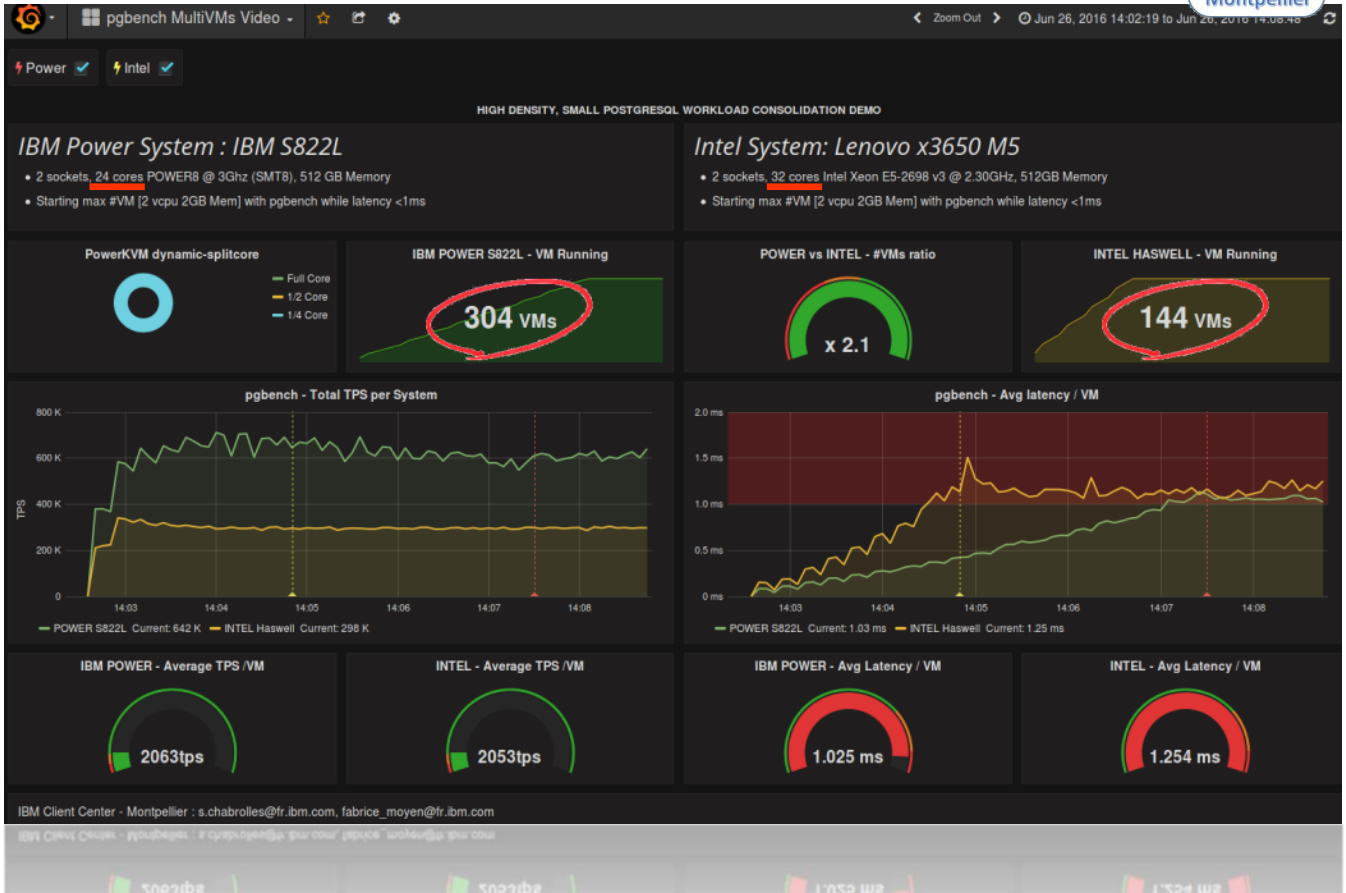

CAS CLIENTS

Les avantages décisifs de Power

Reduce operating costs with Power L922 Server running IBM Cloud Private

1.66X price-performance per rack unit over tested Intel Xeon SP Gold 6130 servers (Skylake)

IBM Cloud Private	IBM Power L922 (16-core, 256GB, 2 VMs)	Intel Xeon SP based 2-socket server (32-core, 256GB, 2 VMs)
Server price [2,3,4] 3-year warranty	**$25,932**	**$29,100**
Solution Cost [5] Server + RHEL OS + Virtualization + ICP Cloud Native VPC Annual Subscription @ $250 per core per month x 36 months	**$180,049** **($25,932 + $10,117 + $144,000)**	**$321,019** **($29,100 + $3,919 + $288,000)**
Acme Air workload [1] Total Transactions per Second - With 2 VM's	**36,566 tps**	**39,312 tps**
TPS/K$	**203.1 tps/K$**	**122.5 tps/K$**

1.86X
per core performance

43%
Lower solution costs

1.66X
Better Price-performance

1. Based on IBM internal testing of a VM image running the Acme Air workload (https://github.com/acmeair) with containers bound to a socket including a MongoDB microservice. Results valid as of 3/17/18. and conducted under laboratory condition with speculative execution controls to mitigate user-to-kernel and user-to-user side-channel attacks on both systems, individual result can vary based on workload size, use of storage subsystems & other conditions.
2. IBM Power L922 (2x8-core/3.4 GHz/256 GB memory) 2 x 600GB SATA 7.2K rpm LFF HDD, 10 Gb two-port, 1 x 16gbps FCA, EDB Postgres Advanced Server 10, RHEL 7.4 with PowerVM (2partitions@8-cores each),
3. Competitive stack: 2-socket Intel Xeon Skylake Gold 6130 (2x20-core/2.1 GHz/256 GB memory), 2 x 600GB SATA 7.2K rpm LFF HDD, 1 Gb two-port, 1 x 16gbps FCA , RHEL 7.4, KVM (2 VMs@16-cores each)
4. **Pricing is based on Power L922** http://www-03.ibm.com/systems/power/hardware/linux-lc.html, Typical industry standard x86 pricing **https://www.synnexcorp.com/us/govsolv/pricing/**
5. IBM software pricing for ICP Cloud Native VPC Monthly Subscription .

National institute for subatomic physics

- Interaction and structure of all elementary particles and fields

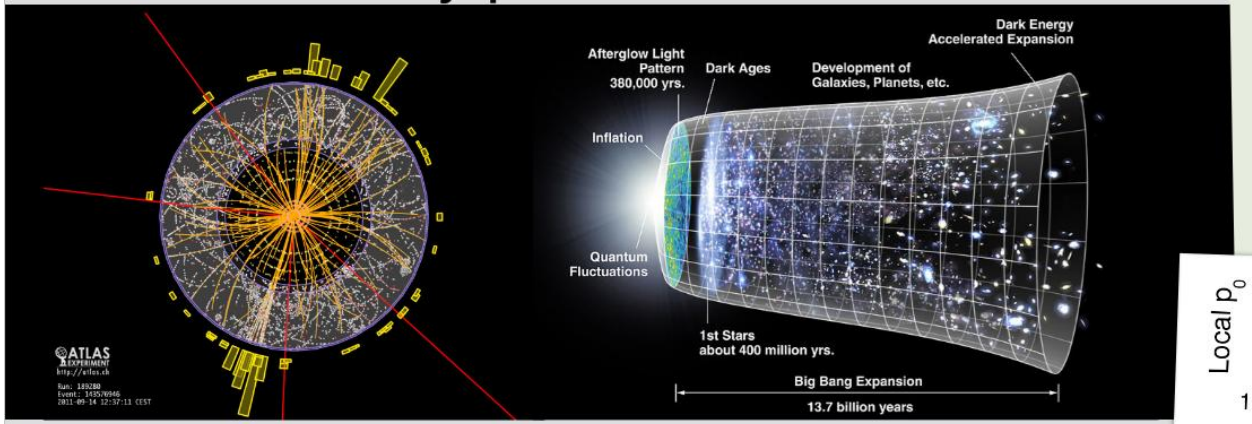

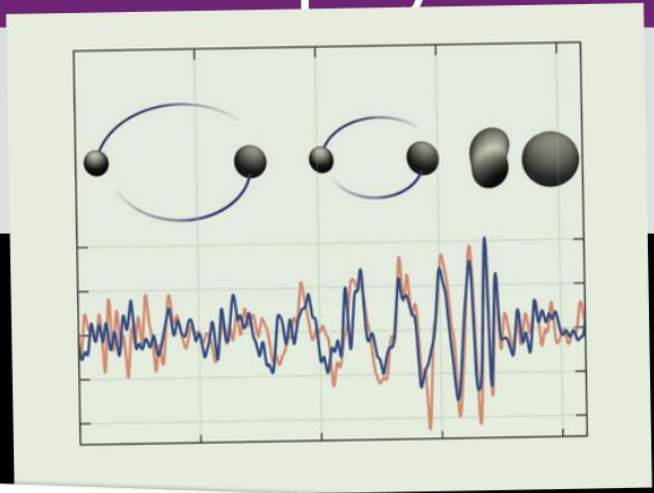

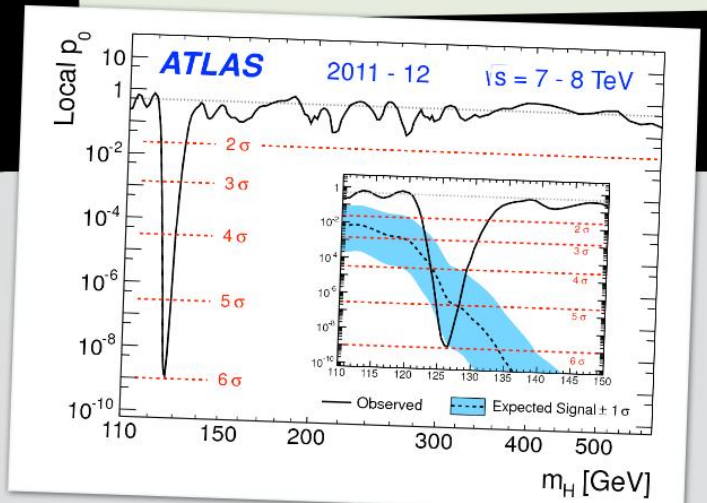

  - Accelerator physics - CERN
  - Astroparticle physics
- Knowledge and technology transfer

LHC programme - Dutch involvement

Long term LHC program and roadmap investment funding (phase1 & phase2)
Excellent scientific programme

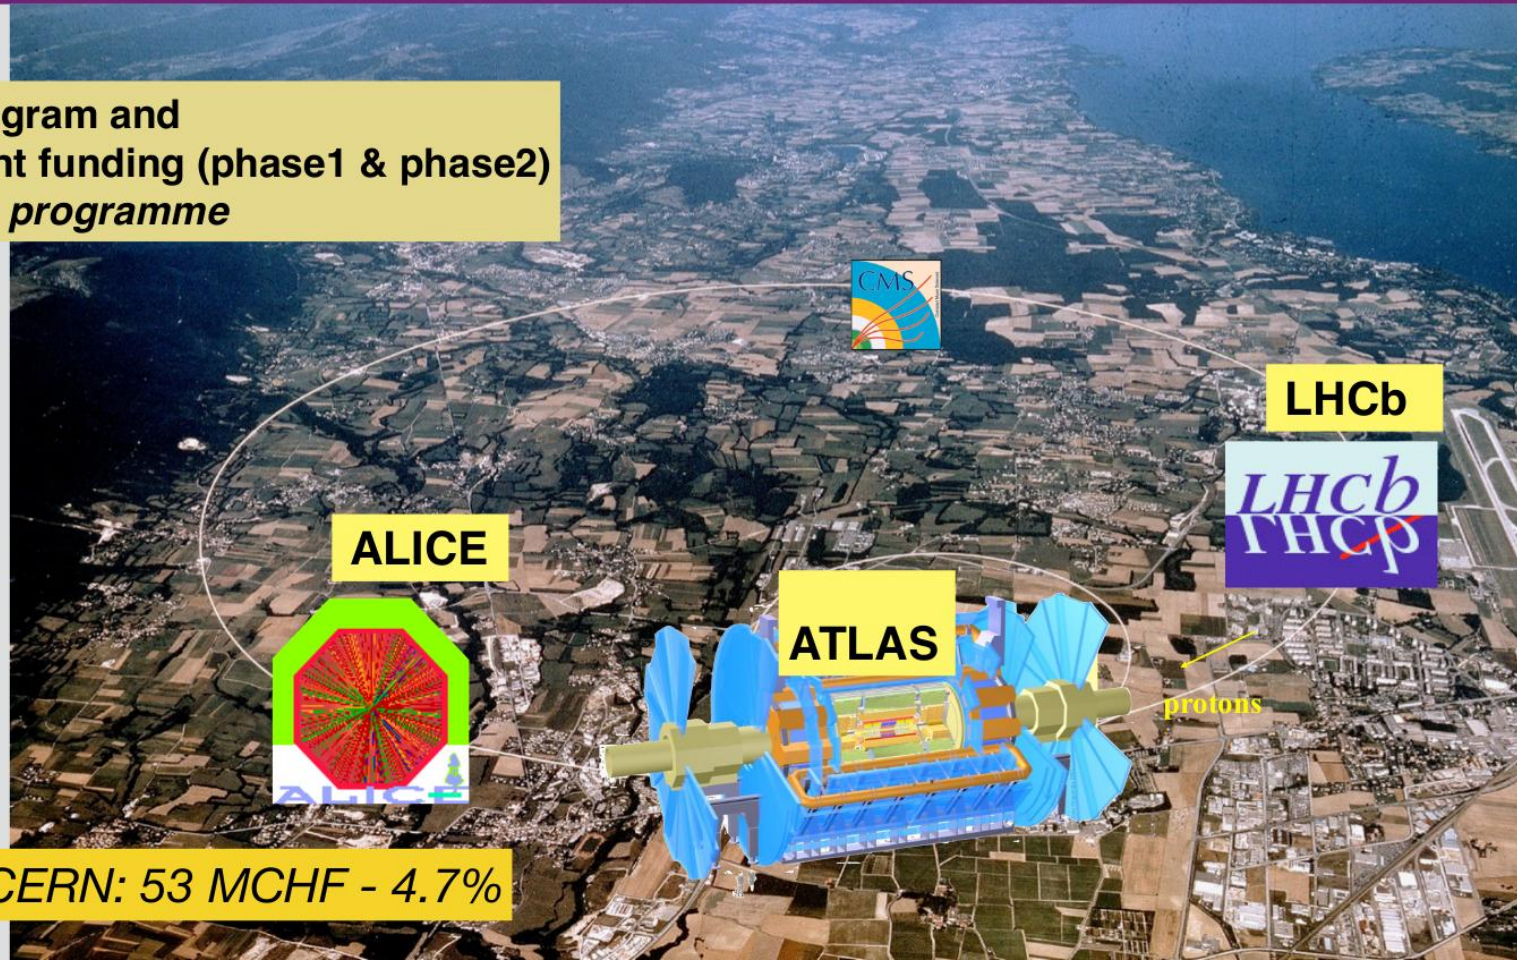

Contribution to CERN: 53 MCHF - 4.7%

Nikhef

1000 LXC containers

- HP DL380 gen 9 vs IBM S822L
- LXC Busybox containers with a dhcp client
- Started all containers in parallel
- HP crashed after several minutes trying…
- IBM did it in 3 minutes

Let's try something fun!

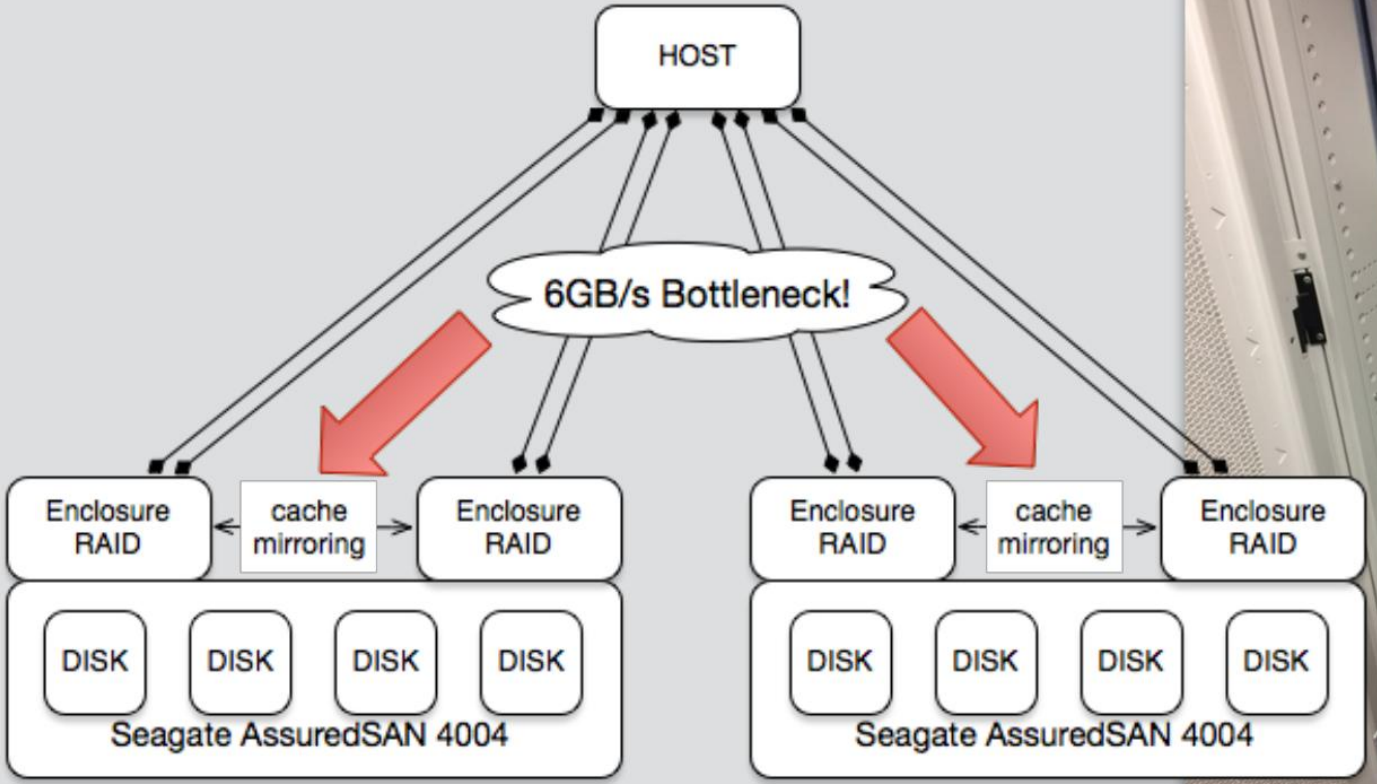

The results!

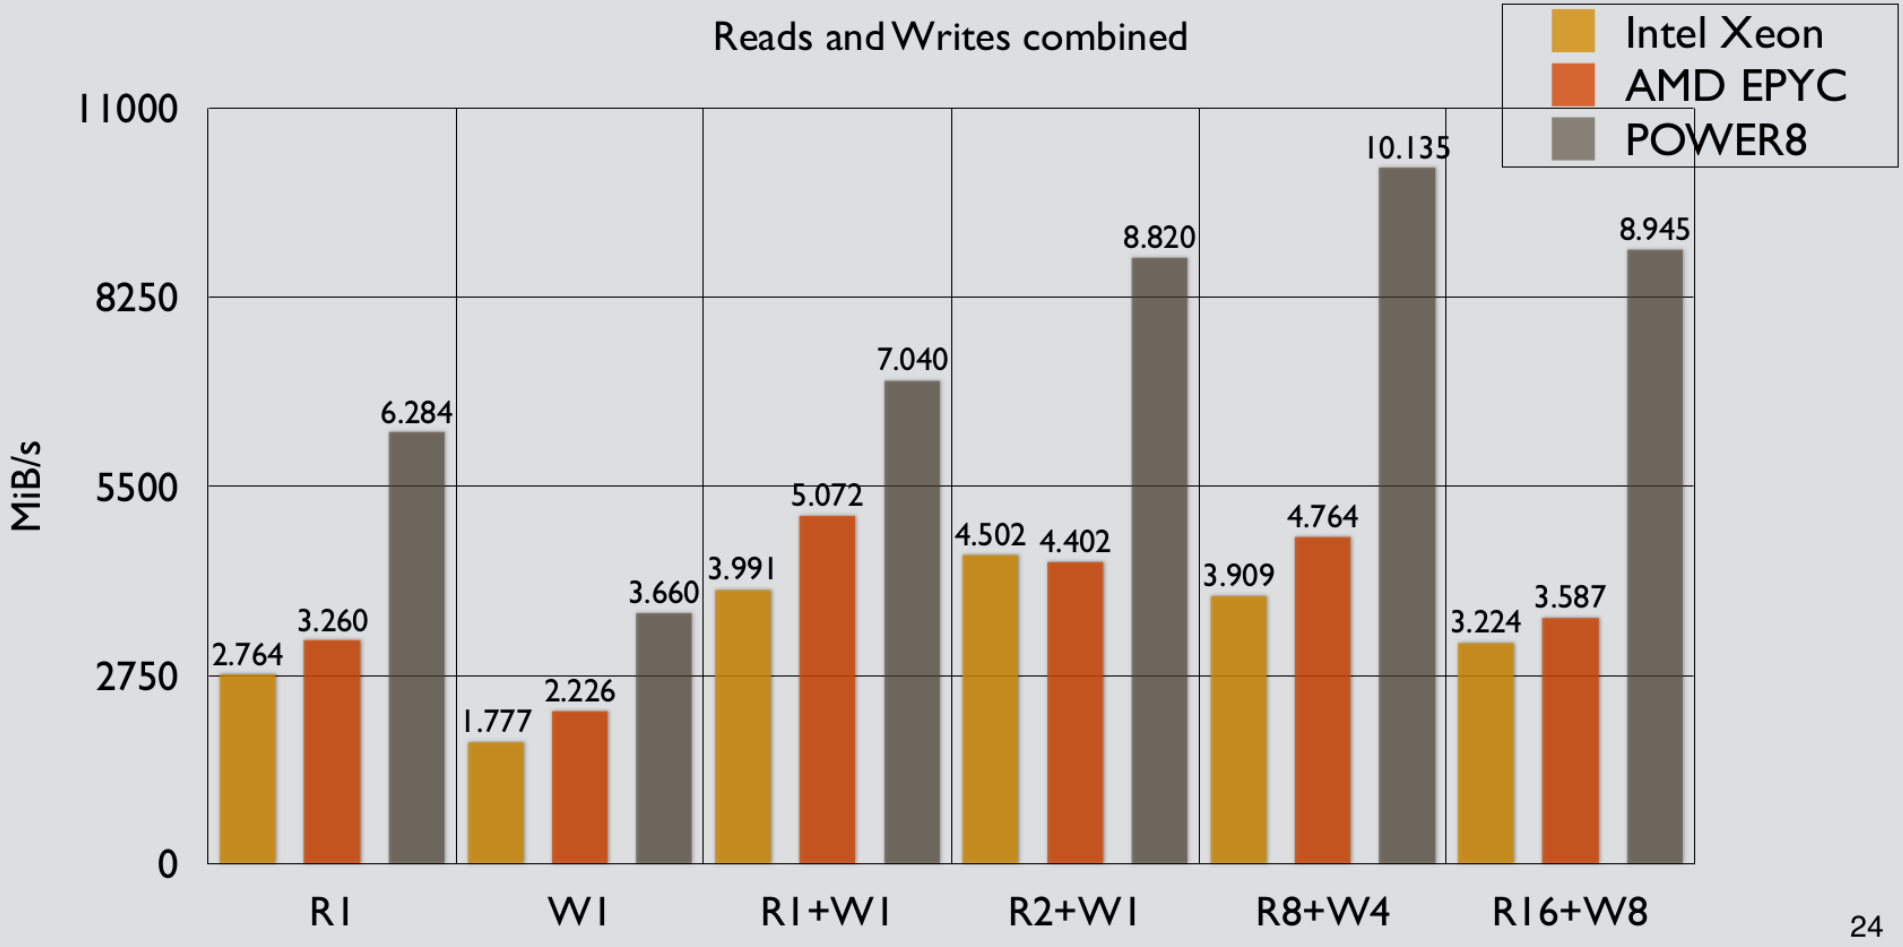
Reads and Writes combined
Intel Xeon
AMD EPYC
POWER8
MiB/s
11000
8250
5500
2750
0
2.764
3.260
6.284
1.777
2.226
3.660
3.991
5.072
7.040
4.502
4.402
8.820
3.909
4.764
10.135
3.224
3.587
8.945
R1
W1
R1+W1
R2+W1
R8+W4
R16+W8

Présentation de Mondial Relay

Le spécialiste de la livraison de colis aux particuliers

49 Millions

de colis livrés

en 2015

36 000

Points Relais

Description du projet applicatif : FAST

Le challenge ? Refondre l'application métier de Mondial Relay

Temps réel

Propager un évènement
En moins de 15 mn

Simple & Efficace

Faciliter l'adoption
15 minutes et c'est parti !

Flexible

Faciliter les évolutions
S'adapter à la croissance

Les enjeux IT ? une plateforme puissante, évolutive et économique pour supporter :

La volumétrie

1 000 000 000 de lignes

Les nouveaux usages

Hyperconectivité
Reporting a tout heure

La croissance des volumes de données

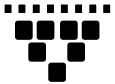

Le colis deviendra t'il
Un objet connecté ?

La créativité métier

Agilité constante

Choix de la solution technique

ROUND 3

2 / Etude de l'architecture

Benchmark à l'IBM Client Center de Montpellier sur l'application de facturation
Même Release, même sizing, même jeu de données

x86 VS POWER8

Performance

Compatibilité

76
Tarification / sec

312
Tarification / sec

1200
Tarification / sec
Après tunning de l'app

Descriptif de l’ architecture technique

- Ouvert sur le monde Open
- Virtualisation avec PowerVM
- 16 Partitions sur 4 serveurs
- OS : Suse SLES 12
- Apps : Java & SpringFramework
- DB PostgreSQL PURE en cluster
- En miroir croisé avec un autre serveur identiquement configuré
- Base de données sur baie stockage SAN virtualisée, avec réplication sur une deuxième baie

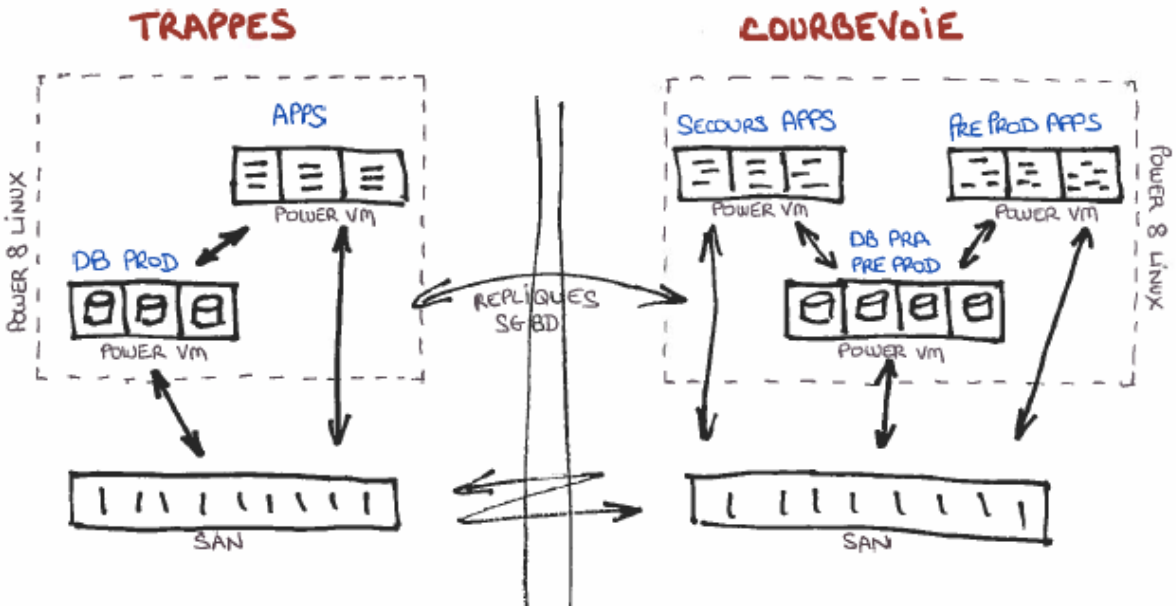

Bénéfices retenus par Mondial Relay

4x

Performance des serveurs

Temps de réponse mesuré 4 fois meilleurs sur la plateforme Power versus x86 :

POWER 4 cores : **3,2 ms** VS x86 4 Cores : **13,1 ms**

Virtualisation & consolidation : flexibilité et meilleure utilisation des ressources/serveurs physiques

- Virtualisation recommandée en PowerLinux alors qu'elle était déconseillée par les experts PostgreSQL sur x86
- A permis de :
  - Passer de 5 serveurs physiques x86 prévus au départ à 2 serveurs physiques Power avec virtualisation
  - Sécuriser le stockage sur SAN virtualisé plutôt qu'en DAS

5x

Evolutivité de la plateforme

- Manque de visibilité du client sur la charge réelle à horizon 2/3 ans
- PowerLinux permet de sécuriser le sizing grâce à une meilleure tenue en charge
- Evolutivité 5x supérieure à celle de la plateforme x86 (grâce au multithreading des processeurs P8)

Data science in ICP / on POWER9

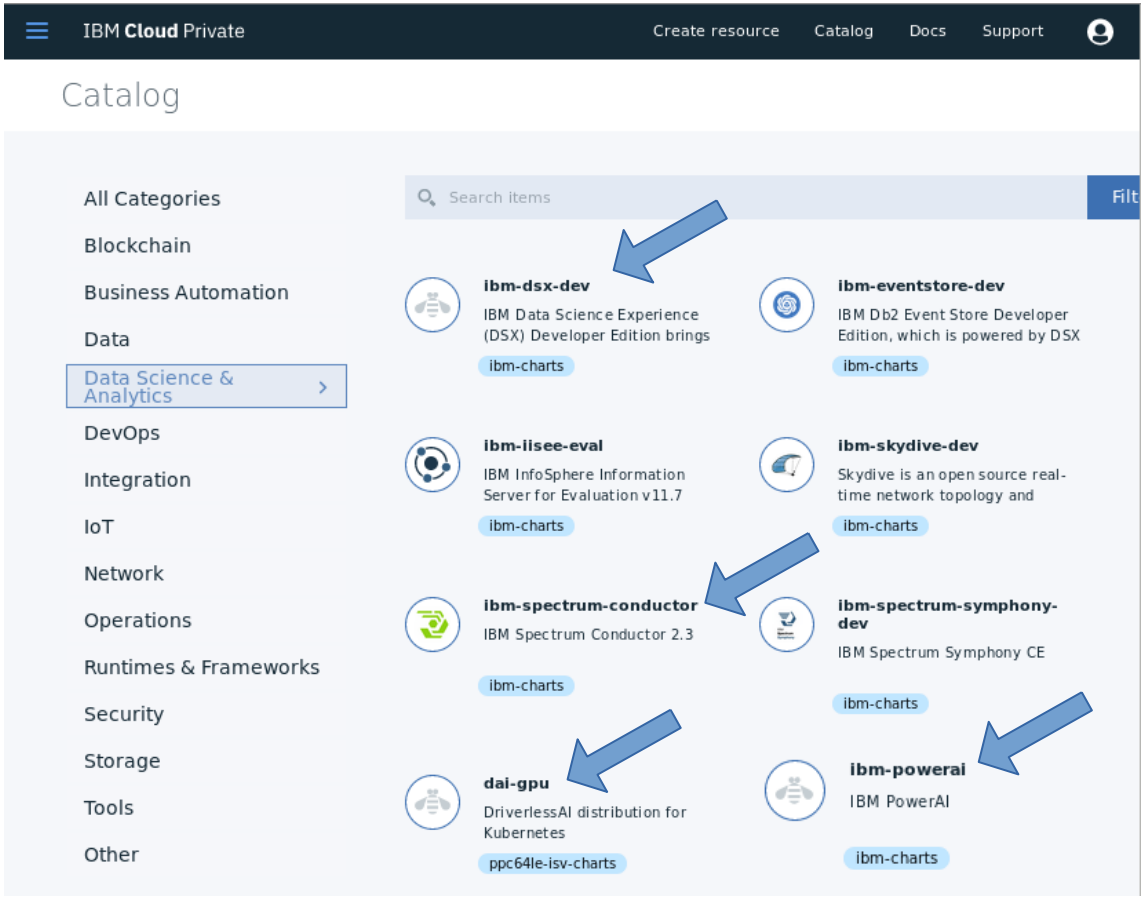

5x Faster Data Communication with Unique CPU-GPU NVIDIA NVLink High-Speed Connection

Store Large Models in System Memory

Fast Transfer via NVLink

Cache layers or tensors between CPU and GPU memory

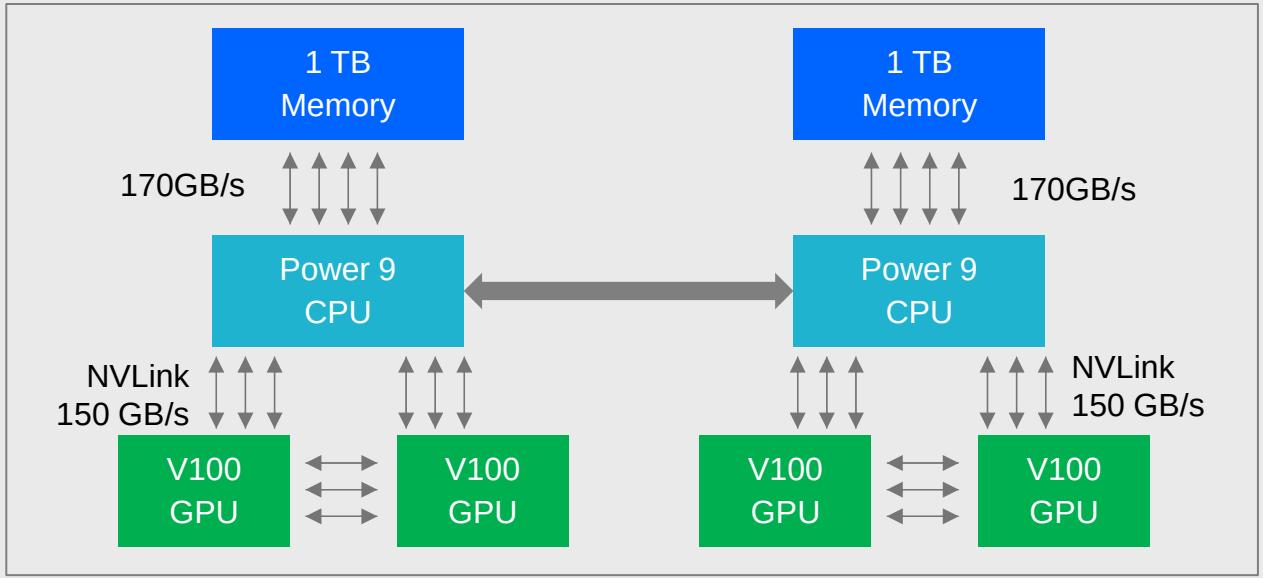

IBM AC922 Power System

Deep Learning Server (4-GPU Config)

IBM PowerAI

- Enable non-Data Scientists to use AI
- Integrated & Supported AI Platform
- Higher Productivity for Data Scientists

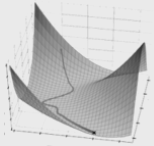

Enable non-Data Scientists to use AI
(Tools for ease of use)

Integrated & Supported AI Platform

TensorFlow Caffe

Keras PYTORCH

Higher Productivity for Data Scientists
(Faster Training with Larger Models)

IBM **Power AI Systems** Performance

IBM PowerAI

Faster training times than x86[1]

- **3.7x** faster running Chainer[2]
- **3.8x** faster running Caffe[3]
- **2.3x** faster running TensorFlow[4] and **35%** more images per second[5]
- Four (4) GPUs on Power is faster than eight (8) GPUs on Intel[6]

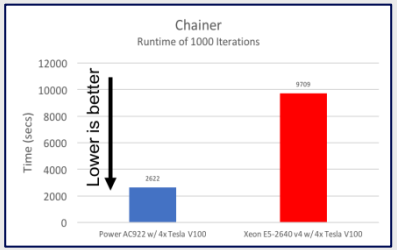

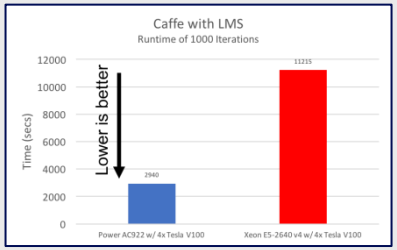

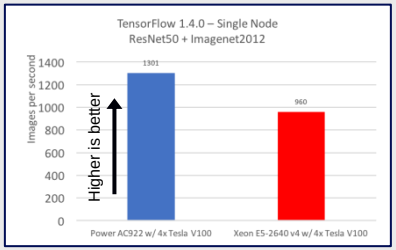

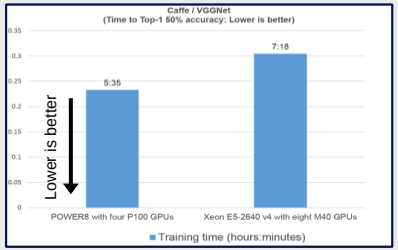

Extended references in notes section [1] https://developer.ibm.com/linuxonpower/perfcol/perfcol-mldl/

Large AI Models Train ~4 Times Faster

POWER9 Servers with NVLink to GPUs

vs

x86 Servers with PCIe to GPUs

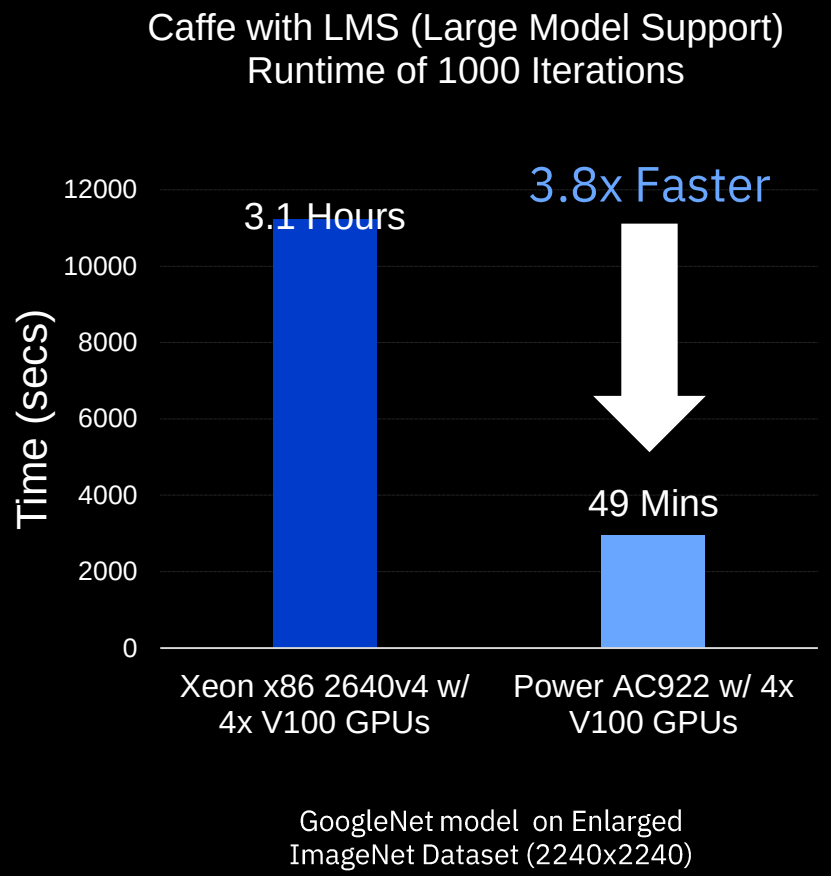

Distributed Deep Learning (DDL)

Deep learning training takes days to weeks

Limited scaling to multiple x86 servers

PowerAI 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

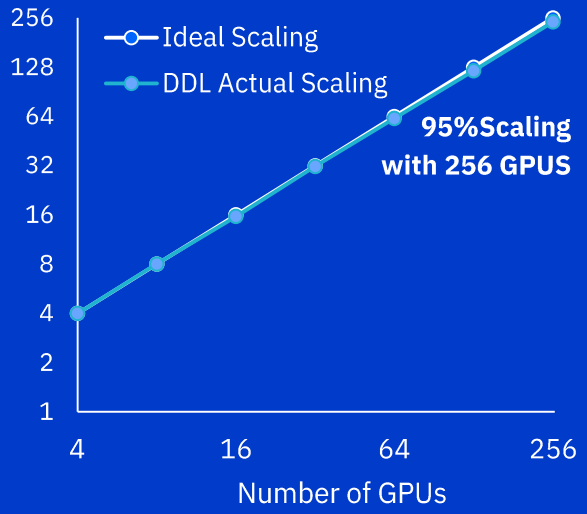

ResNet-50, ImageNet-1K

Caffe with PowerAI DDL, Running on Minsky (S822Lc) Power System

What data science methods are used at work?

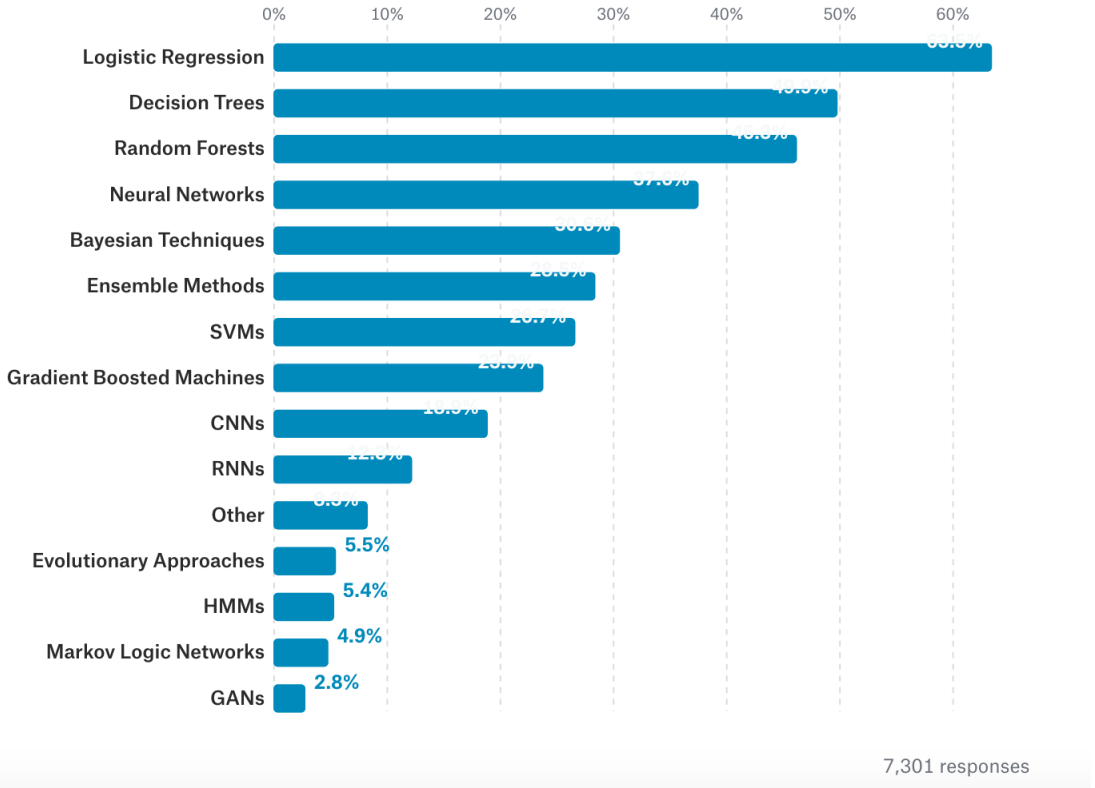

Linear Regression

- When there is a linear relationship between independent and dependent variables.

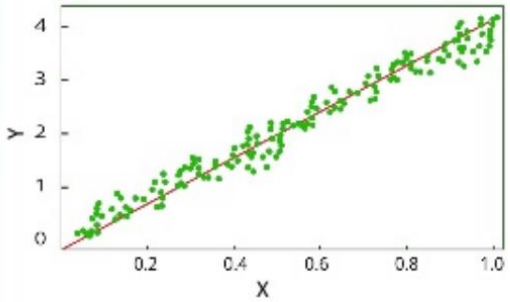

Logistic Regression

- When the dependent variable is categorical (0/ 1, True/ False, Yes/ No, A/B/C) in nature.

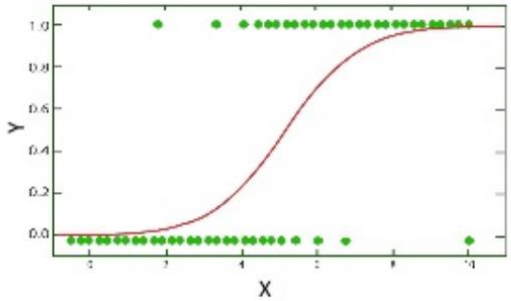

Polynomial Regression

- When the power of independent variable is more than 1.

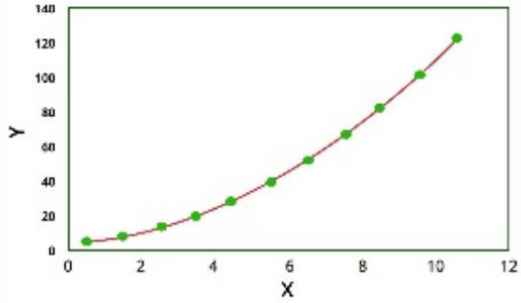

Snap ML Technology Preview

Distributed GPU-Accelerated Machine Learning Library

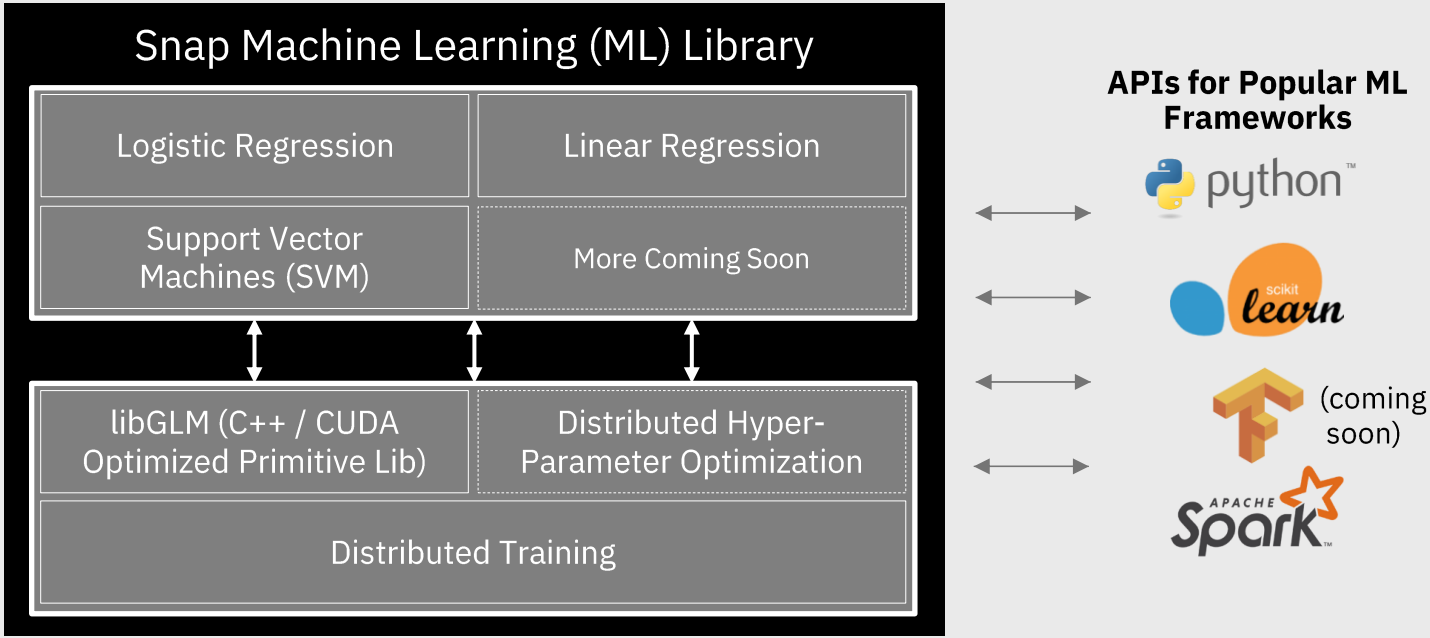

SnapML is 46x faster than best TensorFlow code so far

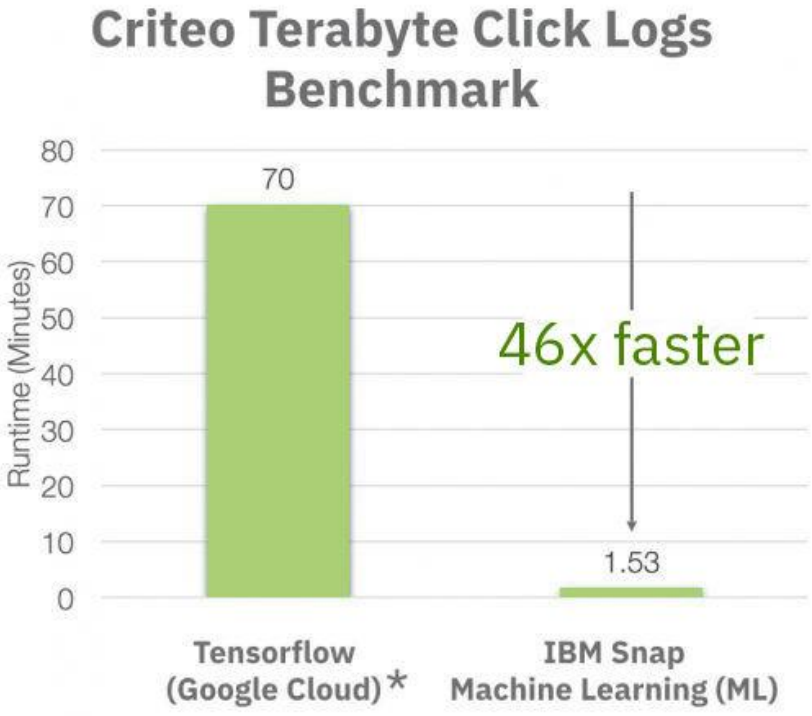

Comparison of Tensorflow on Google Cloud with Snap ML on POWER9 (AC922) cluster

Workload: Click-through-rate prediction for computational advertising, using Logistic Regression

Dataset: Criteo Terabyte Click Logs, 4.2 billion training examples, 1 million features

Model: Logistic Regression

Test LogLoss: 0.1293 (Tensorflow), 0.1292 (Snap ML)

Platform: 89 machines (Tensorflow) compared to 8 Power9 CPUs + 16 NVIDIA Tesla V100 GPUs (Snap ML)

H2O Driverless AI Delivers “Expert Data Scientist in a Box”

- Award-winning software
- Created and supported by world renowned AI experts from H2O.ai
- Empowers companies to accomplish AI and ML with a single platform
- Performs the function of an expert data scientist and adds more power to both novice and expert teams
- Details and highlights insights and interpretability with easy to understand results and visualizations

21 day free trial for [Driverless AI]

Driverless AI Delivers “Expert Data Scientist in a Box”

Automation:
- Visualization
- Feature Engineering
- Model Tuning
- Time Series

Deployment:
- Automatic Pipelines
- Low latency inferencing

Trust:
- Explainability and reason codes
- Automatic documentation

Enterprise Ready:
- Secure – LDAP, Kerberos
- Scalable – Scale with GPUs
- Support for enterprise data sources

21 day free trial for [Driverless AI]

Data Science is a Team Sport

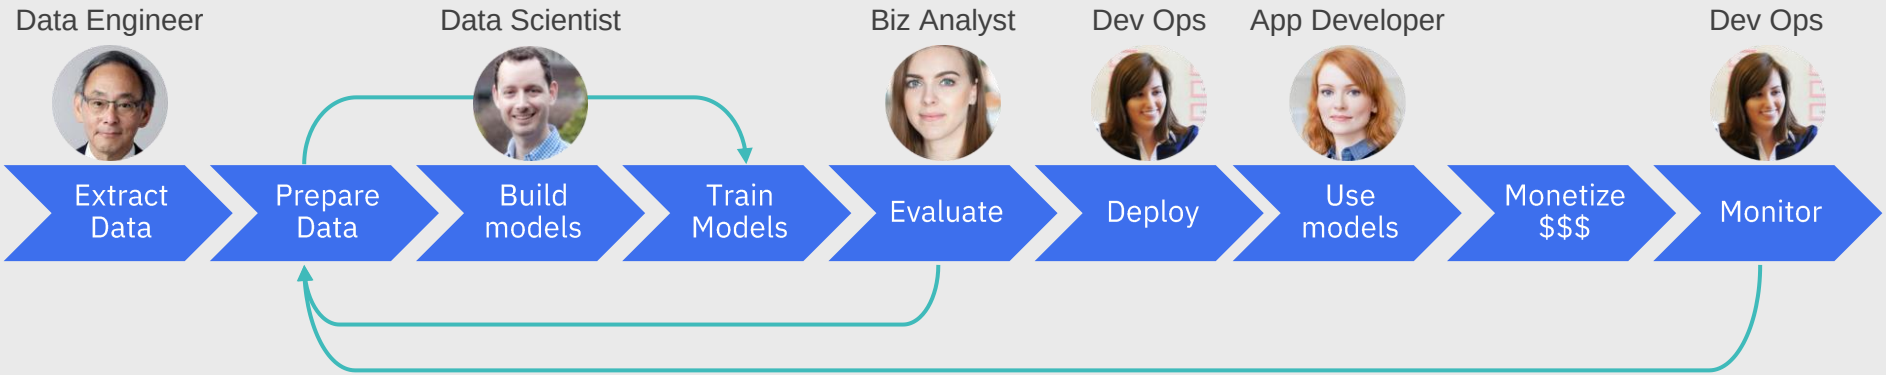

Building cognitive apps using deep learning **requires** multiple skillsets

Data science challenges

datascience.ibm.com

Tools:
- Open Source
- GUI
- Scalability

Data:
- Formats
- Sources

Collaboration :
- Projets
- Teams
- Versioning
- Share
- assets

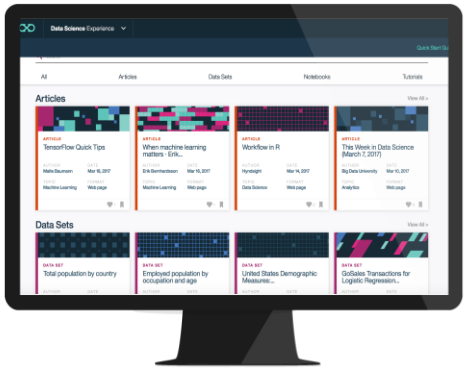

Deployment:
- Monitoring
- Applications

Administration :
- Installation
- Configuration

Data Science Experience

IBM Data Science Experience (DSX) is a collaborative platform for developing and deploying analytics applications

Supports development, deployment, and model management

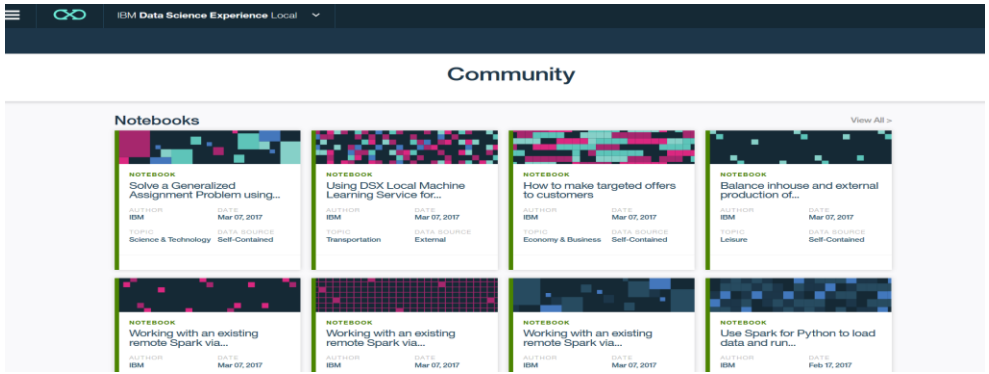

- **DSX is available**
  - As an on-premise solution (Local)
  - As a cloud offering

DSX Local Cluster Architecture

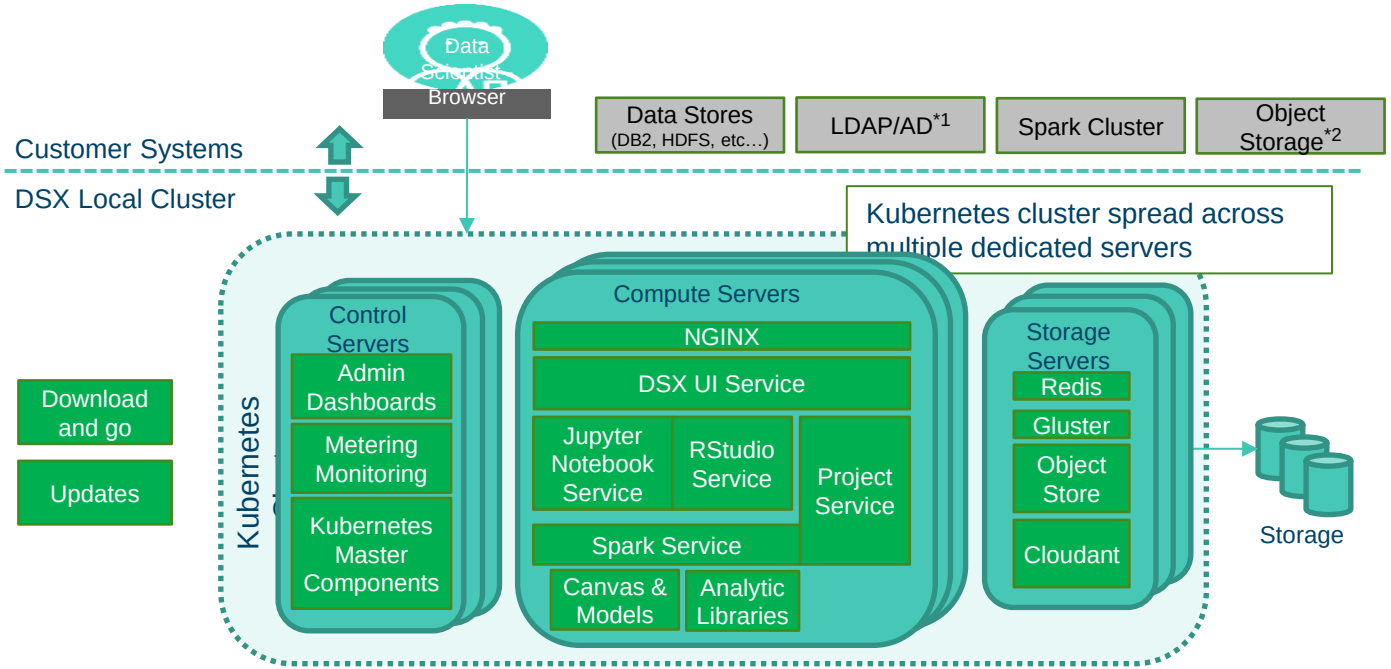

*1 Only for user authentication in initial release
*2 Support planned for a future release

Power with Cloud Private

It's a heterogeneous world – When to leverage Power Systems for Private Cloud?

Examples of Power Benefits

Application & Infrastructure Modernization

Modernize and optimize existing applications with high performance access to enterprise data

Run on same CPU as AIX and IBM i apps for fastest data access

New Cloud Native Applications

Create new cloud-native applications with Open Source apps and databases

Use Open Source DB's with 2X Performance

Deep Learning

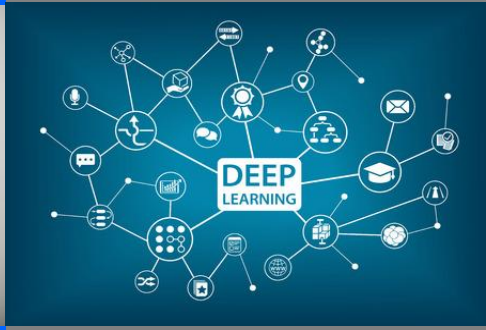

Create new ML/DL apps with PowerAI, DSX, Open Source tools

Train 4X+ Faster on Power9