The Convergence of AI and HPC: A New Hybrid Architecture for High Performance Data Analytics

Andy Grant,
VP, HPC and Big Data, Atos UK&I

Atos

Agenda

1. Factors driving the convergence of HPC, HPDA and AI Architectures
2. Combining the strengths of HPC and AI
3. **BullSequana XH2000**: A new hybrid architecture for HPC and AI
4. **Codex AI Suite:** A new framework for developing cognitive applications
5. **BullExtreme Factory:** HPC within a hybrid cloud services model
6. Expanding Extreme Factory for Deep Learning
7. Conclusions

Atos

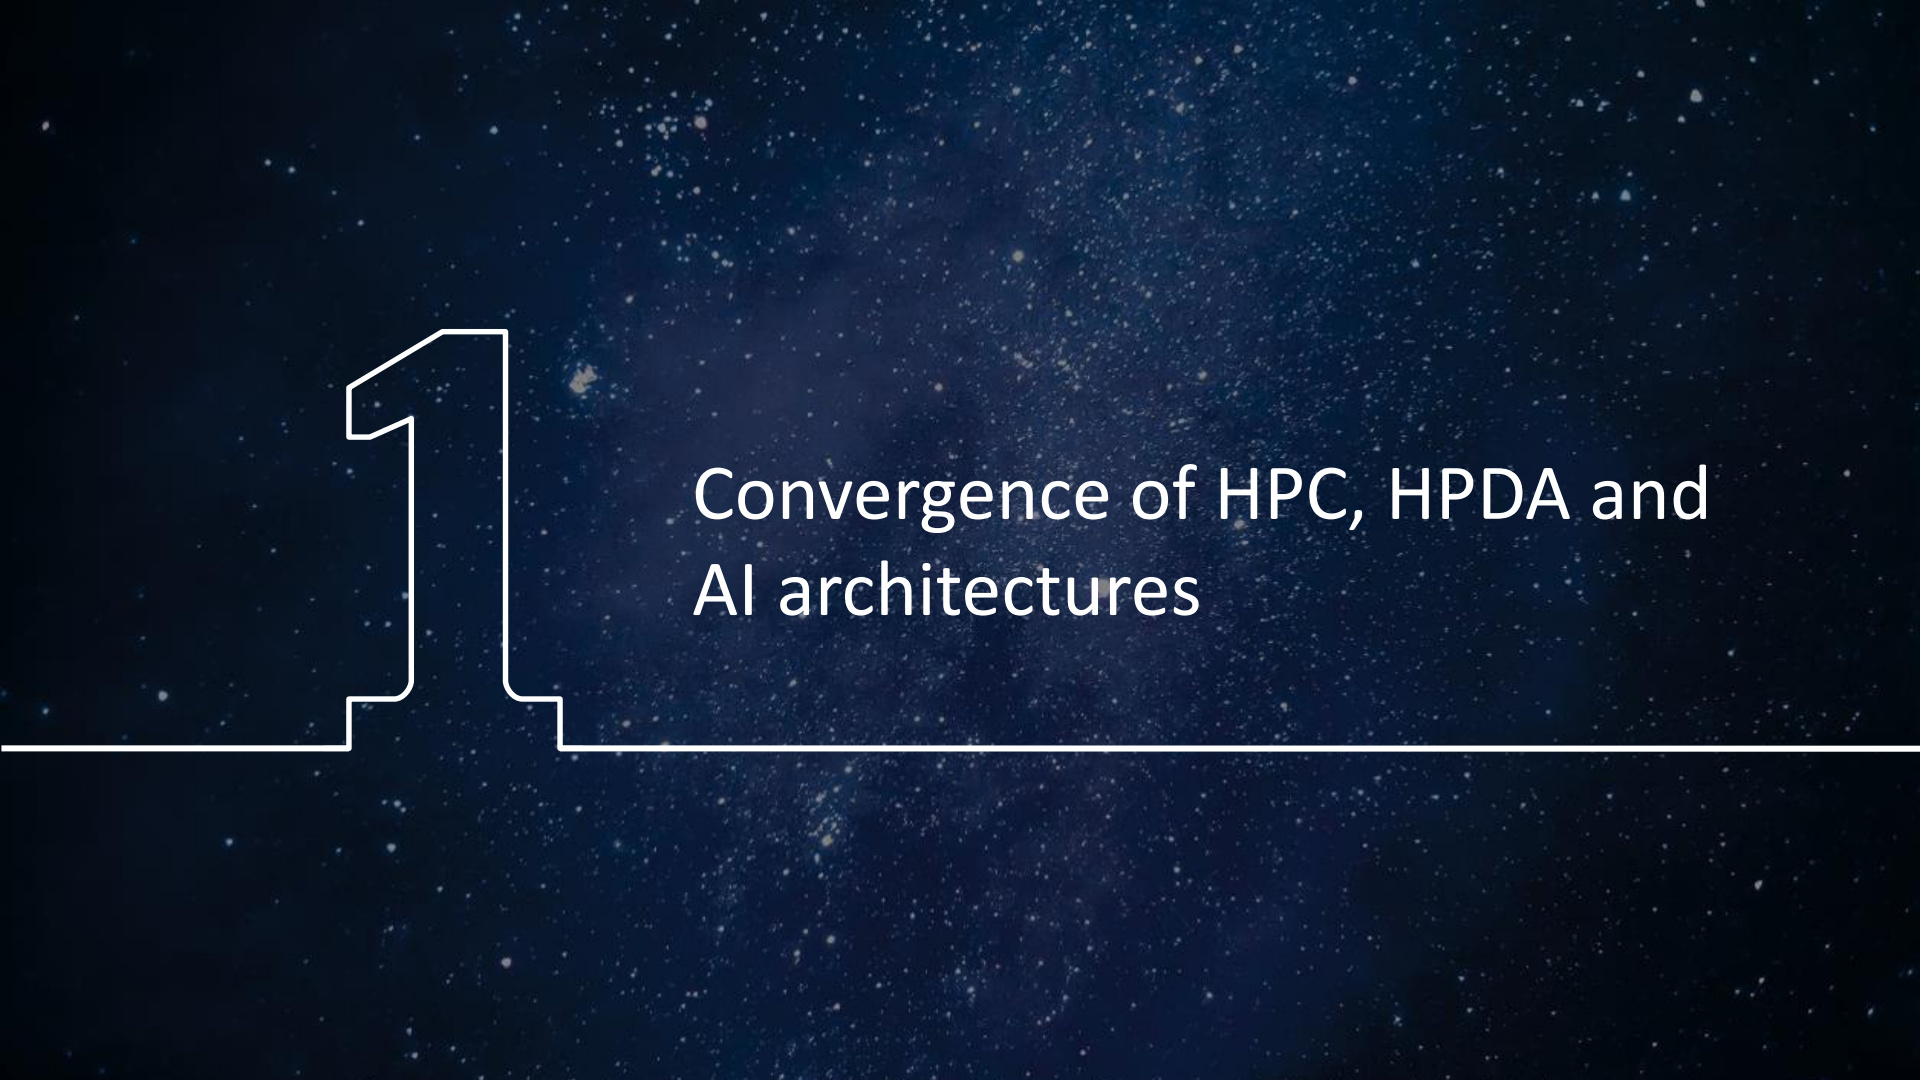
1
Convergence of HPC, HPDA and
AI architectures

Insatiable need for increasing compute power:

E.g: Ensemble Prediction Systems (EPS) in weather forecasting

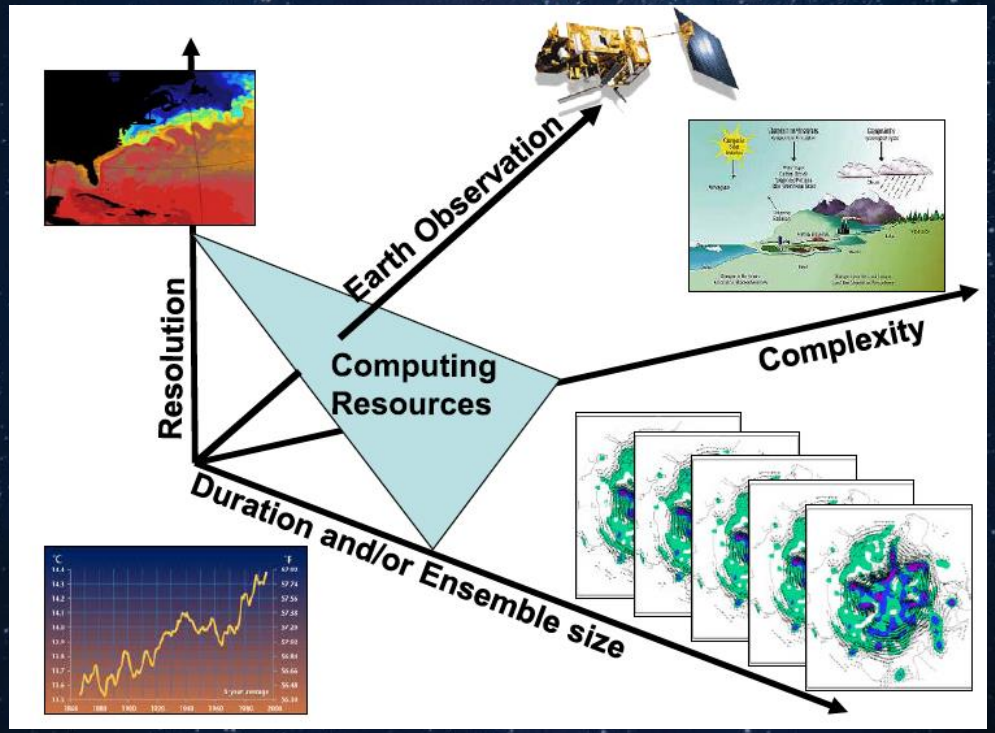

Met Office

Increasing parallelism inevitable

Putting additional demands on the programmer

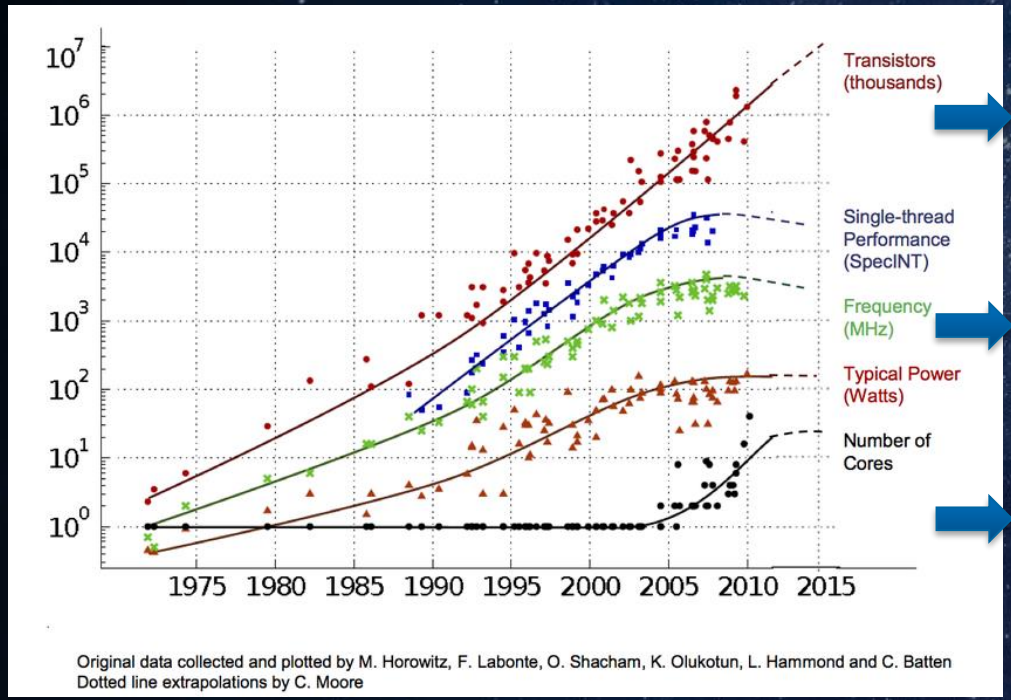

- Number of transistors continues to grow roughly in proportion to Moore's law
- Dennard Scaling (power density roughly constant as transistors get smaller) has broken down
- Single thread performance and processor frequencies have now plateaued
- Computational power now coming from increasing number of cores per processor
- Increasing parallelism inevitable

Atos

Factors Driving Innovation in HPC and DL

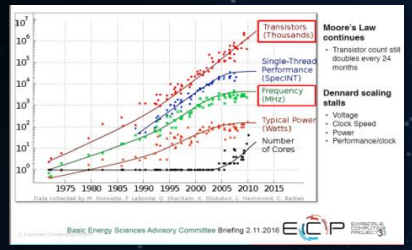

- End of Dennard Scaling places a cap on single threaded performance
- Increasing application performance will require fine grain parallel codes with significant computational intensity

- AI and Data Science emerging as important new components of scientific discovery
- Dramatic improvements in accuracy, completeness and response time yield increased insight from huge volumes of data

- Cloud based usage models, in-situ execution and visualization emerging as new workflows critical to the science process
- Tight coupling of interactive simulation, visualization, data analysis/AI

Atos


2
Combining the strengths of HPC
and Deep Learning

Incorporating DL as part of the HPC Workflow

HPC	Deep Learning
Long history of modelling and simulation of physical phenomena. Track record of enabling grand challenge scientific discovery and proven return on investment in multiple science domains	New methods to improve predictive accuracy, insight into new phenomena and response time with previously unmanageable data sets
• Develop training data sets using first principal models • Apply Bayesian regression methods to expedite/ensure training accuracy • Incorporate AI models in semi-empirical style applications to improve throughput • Validate new findings from AI	• Train inference models to improve accuracy and comprehend more of the physical parameter space • Implement inference models with real time interactivity • Analyze data sets that are simply intractable with classic statistical models • Control and manage complex scientific experiments or apparatus

Atos

Use of DL for uncertainty modelling

Example: Reducing the Parameter sweep in Ensemble Prediction Systems

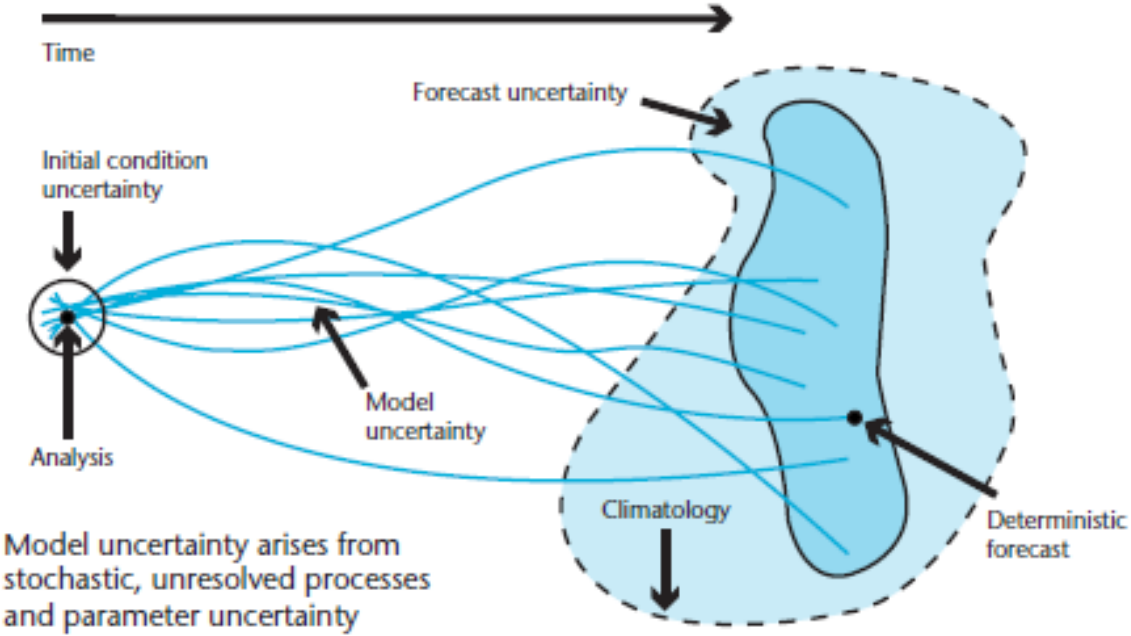

Met Office

Atos

Organising HPC + DL Convergence

Future of HPC

Transformation

HPC + DL couple simulation with live data in real time detection/control system

Experimental/simulated data is used to train a NN that is used to for detection/control of an experiment or clinical delivery system in real time.
The NN is improved continuously as new simulated / live data is acquired

Augmentation

HPC + DL combined to improve simulation time to science > orders of magnitude

Experimental/simulated data is used to train a NN that is used to replace all or significant runtime portions of a conventional simulation.
The NN is improved continuously as new simulated / live data is acquired

Modulation

HPC + DL combined to reduce the number of runs needed for a parameter sweep

Experimental/simulated data used to train a NN which steers simulation/experiment between runs

The steering NN can be trained continuously as new simulated / live data is acquired

Potential for Breakthroughs in Scientific Insight

Courtesy of Nvidia

Atos

Exascale Deep Learning Enabled Precision Medicine for Cancer

CANDLE accelerates solutions towards three top cancer challenges

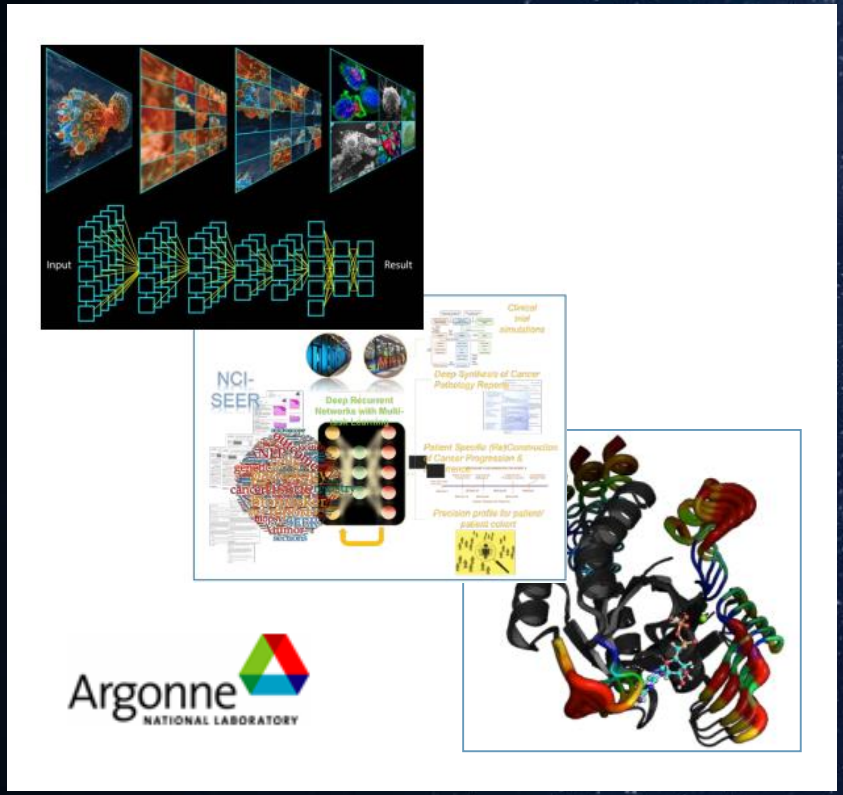

- Focus on building a scalable deep neural network code called the CANcer Distributed Learning Environment (CANDLE)

- CANDLE addresses three top challenges of the National Cancer Institute:
  1. Extraction of information from millions of cancer patient records to determine optimal cancer treatment strategies
  2. Understanding the molecular basis of key protein interactions
  3. Developing predictive models of drug response, and automating the analysis

Atos

Semi-supervised Deep Learning to detect tissue abnormalities

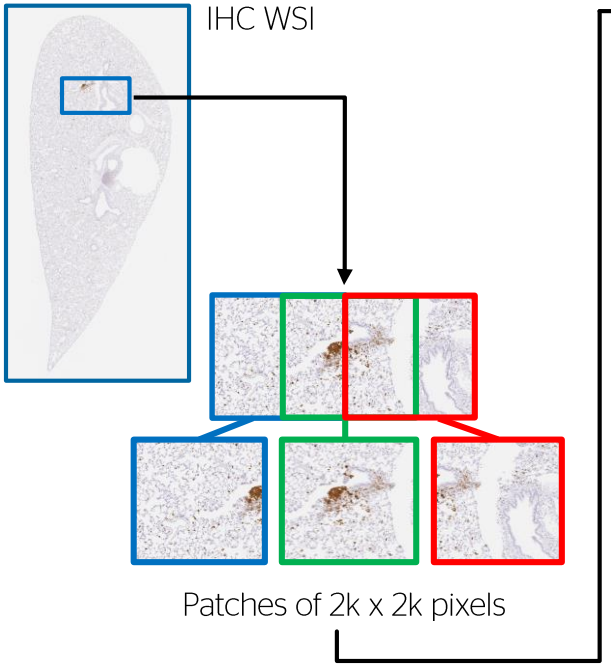

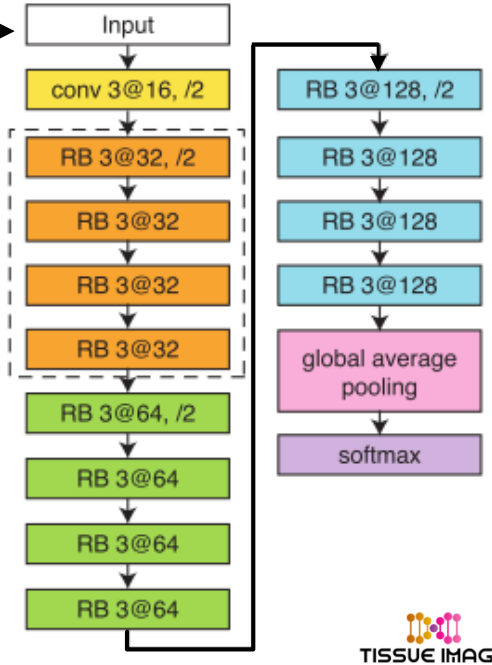

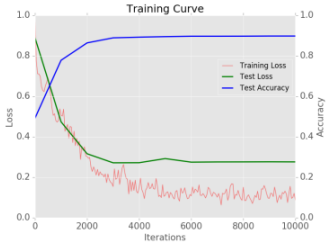

~90% Accuracy

Training on a composite optimization loss function across each batch

TISSUE IMAGE ANALYTICS LAB

University Hospitals Coventry and Warwickshire NHS Trust

intel

Atos

Gravitational Waves – Next steps

Use of Deep Neural Networks to accelerate waveform detection

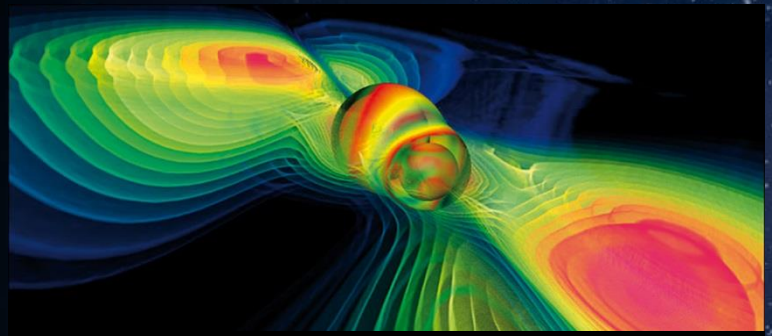

Despite the latest development in computational power, there is still a large gap in linking relativistic theoretical models to observations.
Max Plank Institute

Background

The aLIGO (Advanced Laser Interferometer Gravitational Wave Observatory) experiment successfully discovered signals proving Einstein's theory of General Relativity and the existence of cosmic Gravitational Waves.

Challenge

The initial a LIGO discoveries were successfully completed using classic HPC analysis using hundreds of CPU's where the bulk of the processing was done offline. Here the latency is far outside the range needed to activate resources, such as the Large Synaptic Space survey Telescope (LSST) which observe phenomena in the electromagnetic spectrum in time to "see" what aLIGO can "hear".

Solution

A DNN was developed and trained using a data set derived from the CACTUS simulation using the Einstein Toolkit. The DNN was shown to produce better accuracy with latencies 4500x better than the original CPU based waveform detection.

Atos

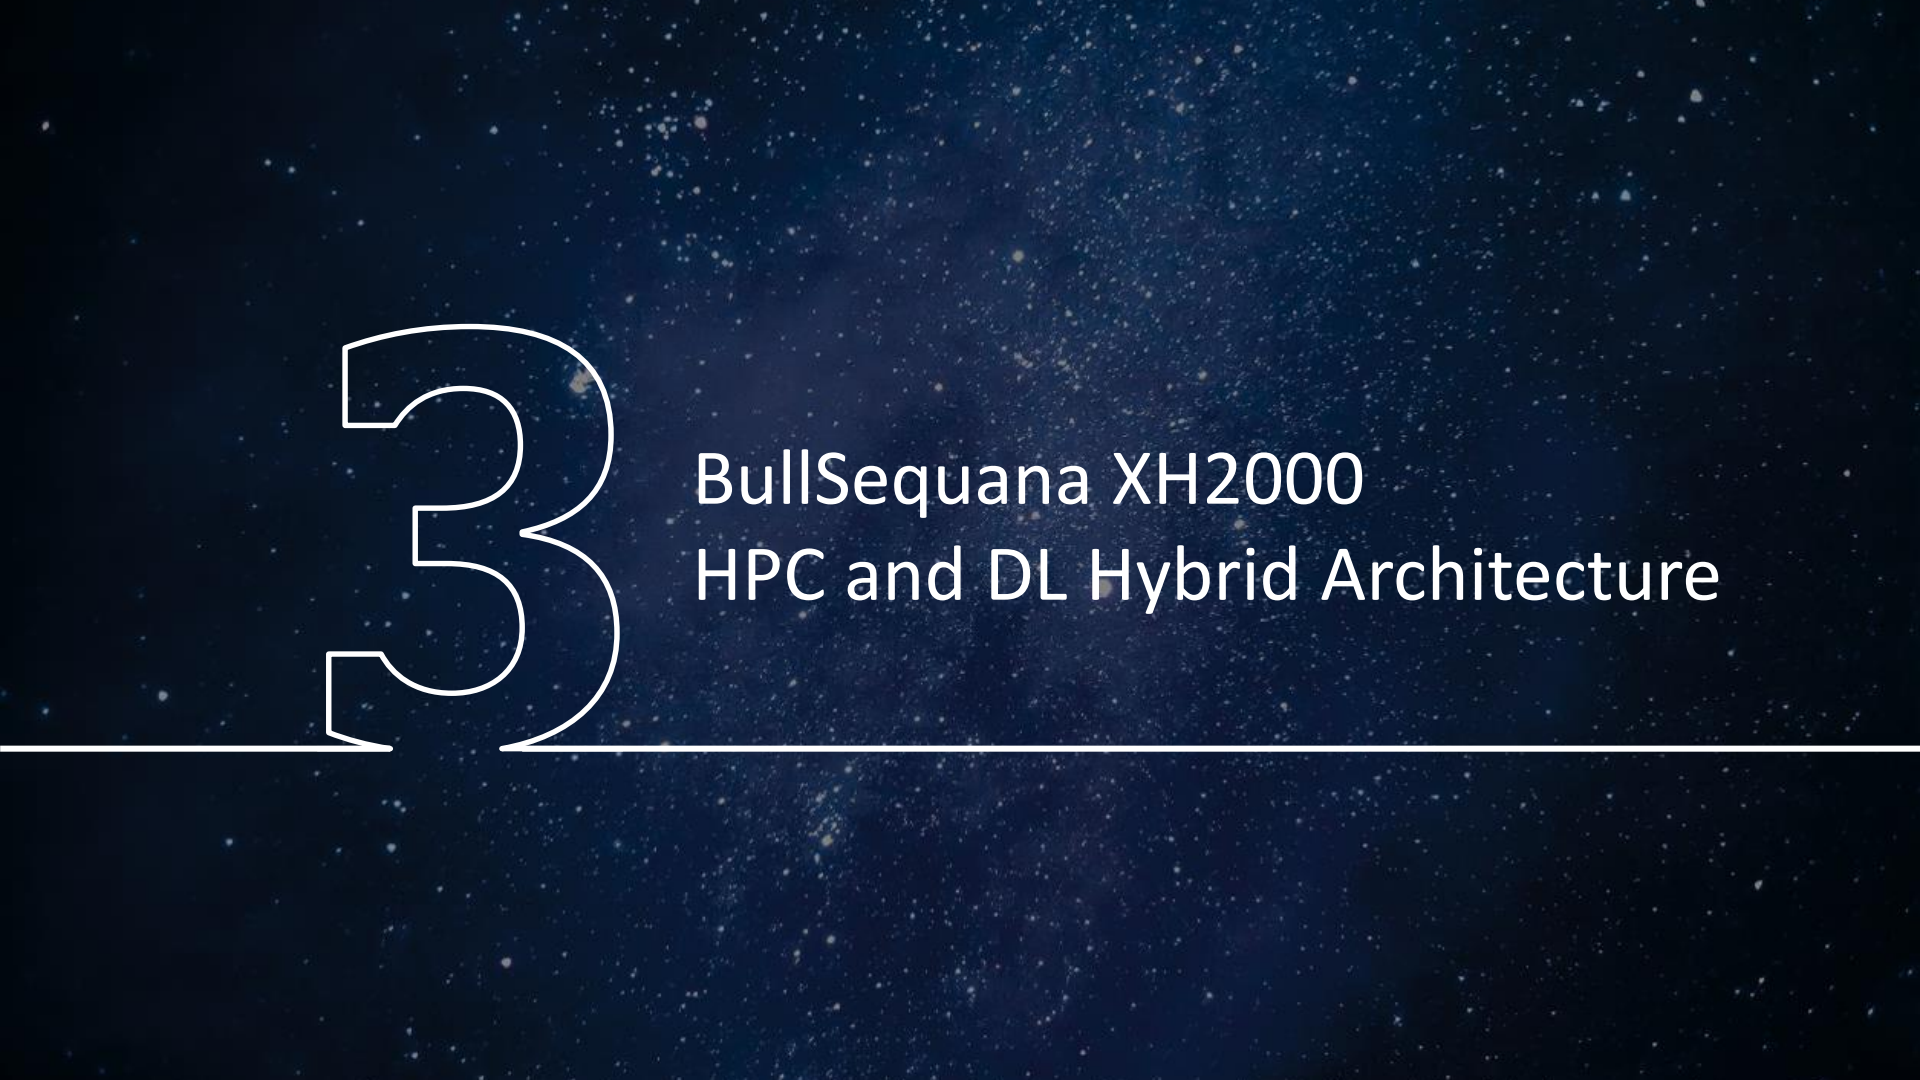
3
BullSequana XH2000
HPC and DL Hybrid Architecture

BullSequana XH2000

A natural evolution for BullSequana X1000

A field proven platform

- Blades are backward and forward compatible
- More than 60% of components in common
- Exascale ready
- Platform in production with Governmental, educational and industrial customers

CSC · GENCI · JÜLICH FORSCHUNGSZENTRUM · cea · AWE · IT4I · NSM NATIONAL SUPERCOMPUTING MISSION · INES Centre Informatique National de l'Enseignement Supérieur · SURF SARA · Hartree Centre Science & Technology Facilities Council · MBDA MISSILE SYSTEMS · ccrt · SENAI CIMATEC Sistema FIEB · calmip · Université de Reims Champagne-Ardenne

Enhanced flexibility allows to reduce TCO, from one rack up to Exascale projects

BullSequana XH2000 featuring Intel® Xeon® Scalable processors

Atos

BullSequana XH2000
Designed for Hybrid Computing

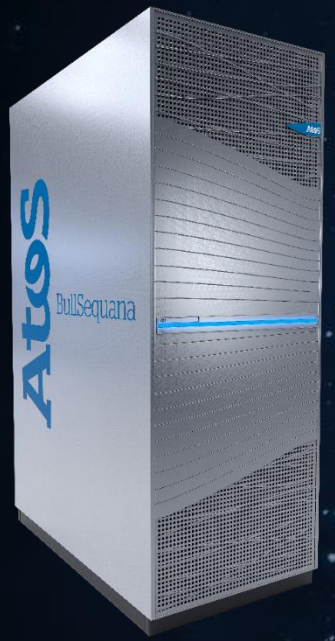

Mix & Match Compute and AI workloads on the same platform…

- Traditional scalar, memory bound or GPU based simulations
- AI augmented simulations & Deep Learning training & inference
- Big Data Analysis

Using best in class CPU/GPU/FPGA & Interconnects

intel AMD NVIDIA MARVELL XILINX

BXI
portals
BXI 1.0
Full Offload

Spectrum·2
10/25/40/50/100GbE Open Ethernet

Seamlessly integrated with Codex AI Suite and Google Cloud Platform

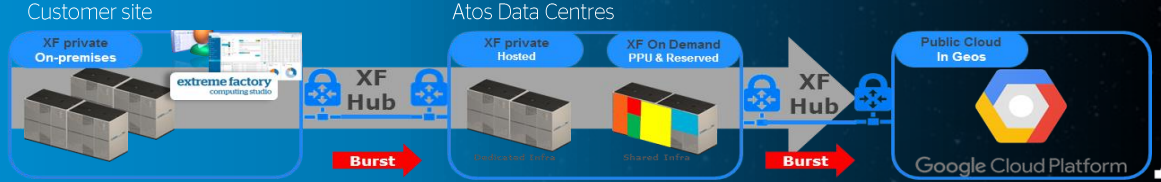

Atos

BullSequana XH2000

Optimized TCO Supercomputer

Best in class Power Usage Effectiveness (PUE)

- 4th generation ATOS DLC using warm water cooling up to 40°C (104°F) Inlet, no single fan.
- Capable of cooling highest TDP/Tcase processors on the market, build for the future

Very high density and usage of available resources

- All-in-one modular design (Compute, Networking, Power, Cooling)
- Up to 96 nodes per rack/in less than $2m^3$ ($68ft^3$)

Atos

BullSequana XH2000 Blades

The best of all worlds, available in one supercomputer

New

EPYC Rome

AMD

intel

intel XEON PLATINUM inside

Xeon SP

Skylake / Cascadelake

NVIDIA

Volta P100/V100

NVLink2

MARVELL

ThunderX2

ThunderX3 when available

Atos

BullSequana XH2000 Networking

Best in class Interconnect flexibility

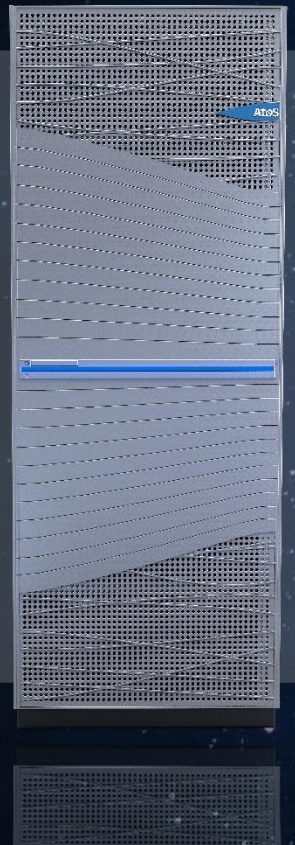

New BXI-2 switch for BullSequana XH2000

- DLC cooled – Atos cold plate technology
- 48 x 100Gb/s ports
- troubleshooting tool : Traffic generator

New IB HDR switch for BullSequana XH2000

- DLC cooled – Atos cold plate technology
- 40 x 200Gb/s ports / 80 x 100Gb/s ports

New High-Speed Ethernet network for BullSequana XH2000

BullSequana XH2000

Increased flexibility and agility

42U Cabinet Back Front

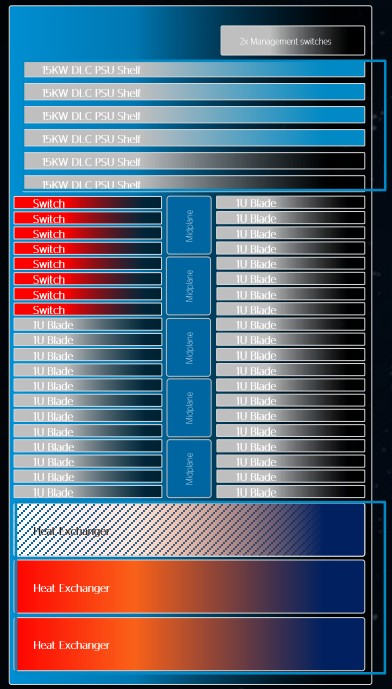

90 KW
n+1 redundancy

Increased flow
HYCs
1+0 redundancy
1+1
2+0
2+1
3+0
SmartThrottle or
Shutdown modes

42U Cabinet Front View

PDU + Power
controller

up to 6 x 15KW DLC shelves

4 to 20 compute blades (front)

up to 3 Hydraulic chassis

42U Cabinet Rear View

2 x Leaf Eth switches

up to 10 IB/BXI/High Speed
Ethernet switches

up to 12 compute blades (Back)

Atos

BullSequana XH2000 Flexibility

Unlimited solutions in real world

Entry level simulation system

60 Cascade Lake Nodes
HDR100
Adjusted Power configuration
with 3+1 R. PSU Shelves
1+1 HYC Redundancy

ATOS SCS5 SI

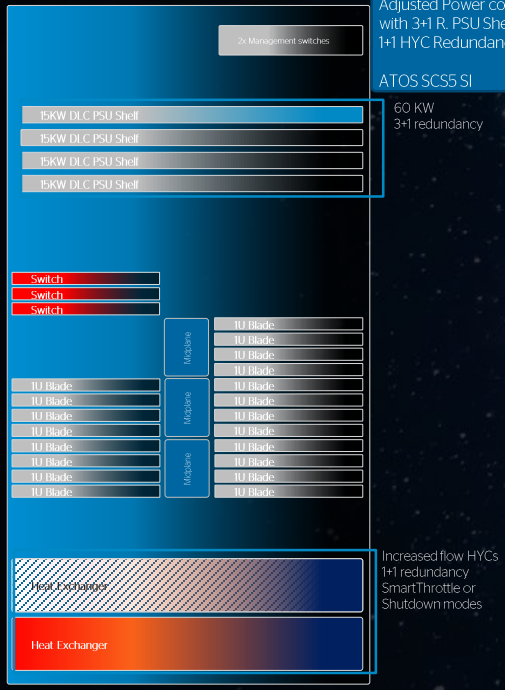

BullSequana XH2000 Flexibility

Unlimited solutions in real world

Entry level simulation system

Deep Learning Training System

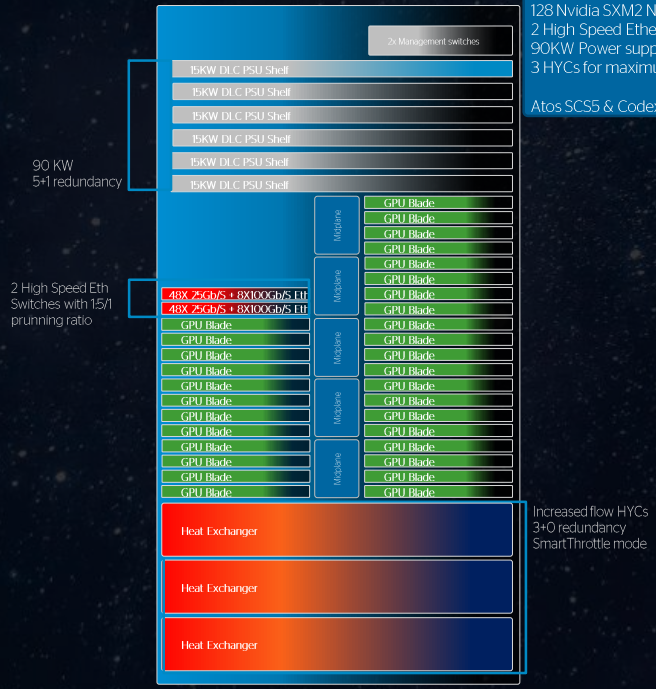

32 Intel® Xeon® CascadeLake
128 Nvidia SXM2 NVLink2 V100
2 High Speed Ethernet switches
90KW Power supplies with 5+1 R.
3 HYCs for maximum efficiency

Atos SCS5 & Codex AI Suite

Atos

BullSequana XH2000 Flexibility

Unlimited solutions in real world

Entry level simulation system

Deep Learning Training System

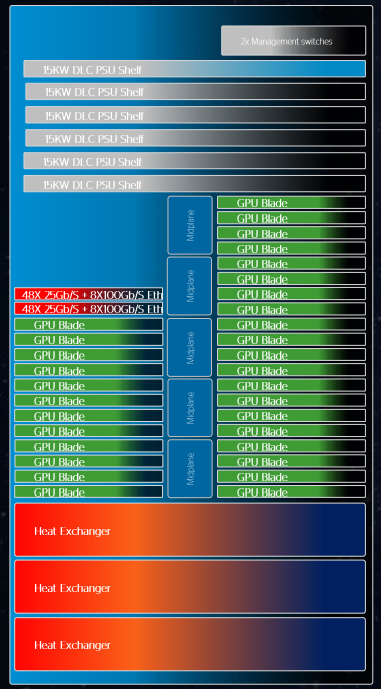

AI augmented HPC System

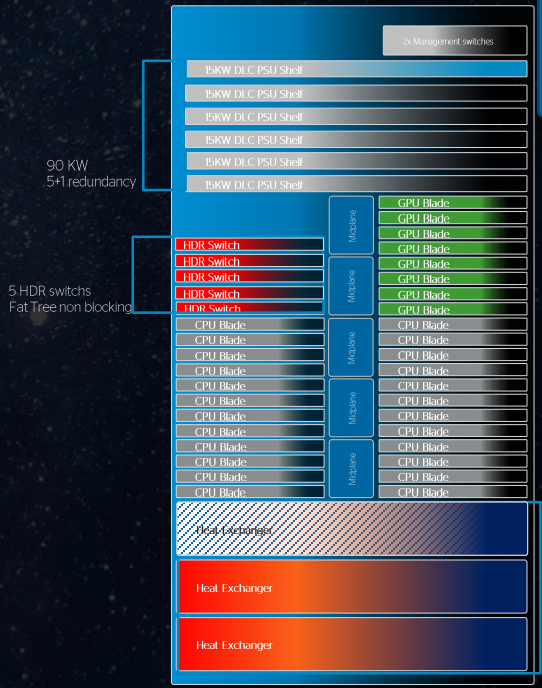

32 Nvidia SXM2 NVLink2 V100
88 AMD EPYC Rome CPU
5 HDR 200 Gb/S switches
90KW Power supplies with 5+1 R.
2+1 HYCs

Atos SCS5 & Codex AI Suite

AMD
NVIDIA
Atos

BullSequana XH2000 Flexibility

Unlimited solutions in real world

192 AMD EPYC Rome CPU
4 HDR 100 Gb/S switches
90KW Power supplies with 5+1 R.
2+1 HYCs

Atos SCS5
Atos Smart Data Management Suite

AMD

Atos

Entry level simulation system

Deep Learning Training System

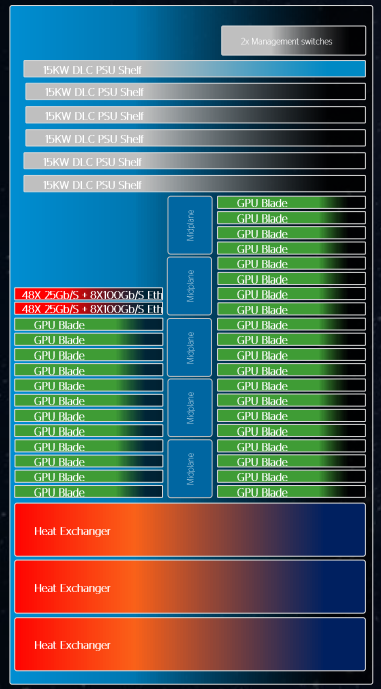

AI augmented HPC System

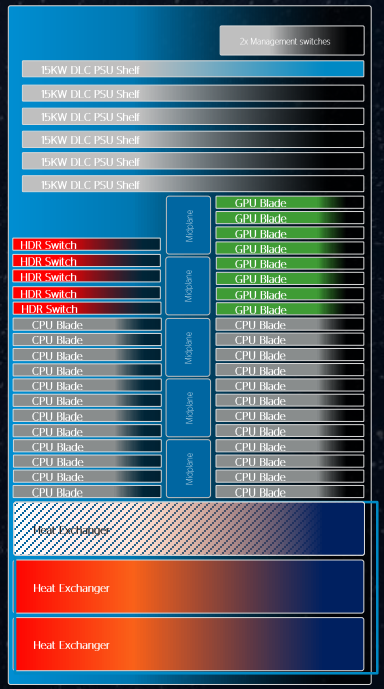

Memory Bound Apps Sim. System

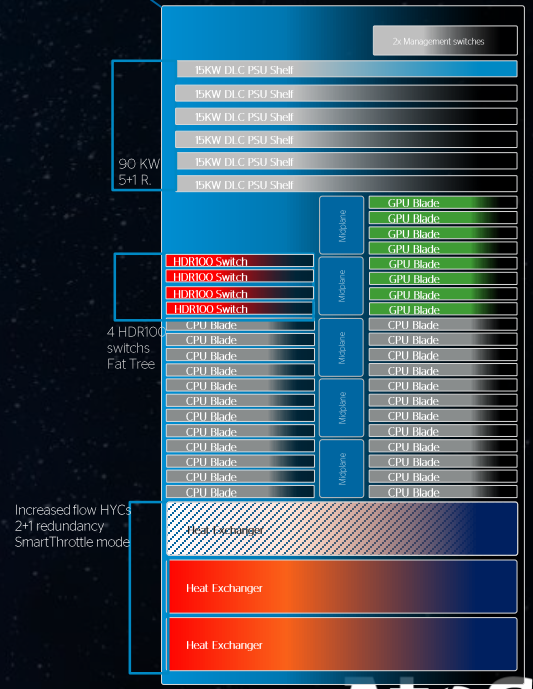

Atos

BullSequana XH2000 Flexibility

Unlimited solutions in real world

Entry level simulation system

Deep Learning training System

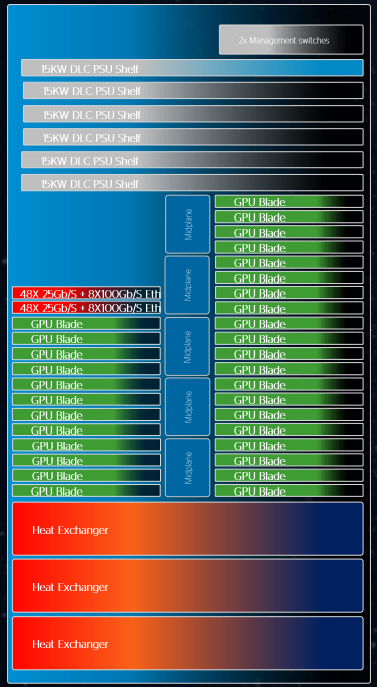

AI augmented HPC system

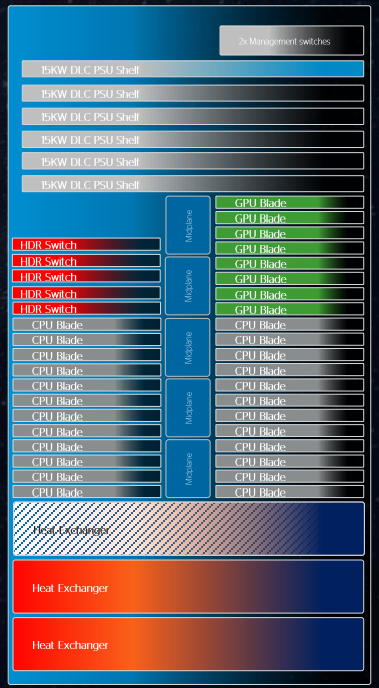

Memory bound simulation

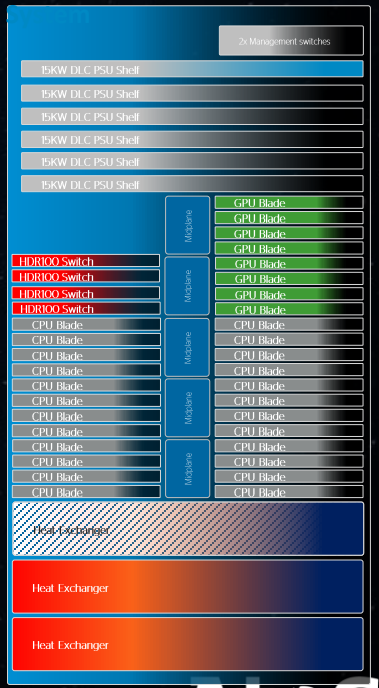

Atos

4

Codex AI Suite:
A Framework for Developing Cognitive Applications

Codex AI Suite tackles the large-scale Deep Learning Problem

New requirements, but different from training to running

Phase 1 : Model Training
(Data Center – heavy computation cost)

Input data

Creation of multi-layer Neural networks

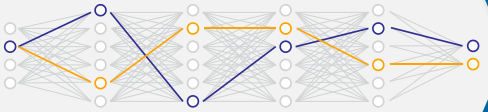

Production of a trained model

Need a lot of data to train the model; if possible, data labelled by human could be better

Trained model

Phase 2 : Model Prediction/Recall
(Local applications or Data Centre – Almost Instantaneous)

Different sensors' data, such as camera

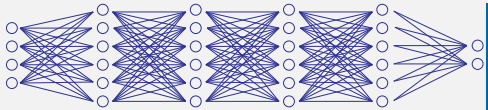

Application of trained model

Pedestrian recognition achieves an accuracy of 97%

Class prediction of new data

Atos

Main Requirements for a Deep Learning Self Service model

- Training Management
- Models Management
- Dataset Management
- Inference Management
- Deep Learning Framework Management
- Optimised Resource Management

Atos

Codex AI Suite at a glance

STUDIO

Cognitive application development self-service

- Application Builder
- FastML Engine

FORGE

common workplace where to store, share, retrieve and update

- Components
  Blueprints
  DL frameworks
- Data sets
  Trained models

FastML ENGINE

- High Level API

ORCHESTRATOR

Fast application deployment on multiple environments

Atos

The Codex AI Suite leverages our partners to accelerate AI solutions

Codex AI Studio is the common workbench to develop use cases

Deep Learning engine
Google ML services

Google Machine Learning

Forge

Partners
(components, algorithms, applications, models)

Orchestrator

amazon web services

Google Cloud Platform

Enterprise, Edge, HPC, On-premise, Cloud

Atos

3 key Codex AI Suite benefits

01

Increases data scientist productivity

- Day-to-day mundane tasks are performed by Codex AI Suite
- Data scientists leverage the use case experience
- Use case policy management

02

Provides a complete cognitive eco-system ready to work

- Customers benefit from an integrated & cost-effective solution
- Applications take immediate benefit of latest technologies with no additional investment

03

Delivers optimal performance at low cost

- High-performance resources are allocated only when need be
- Hybrid cloud resources are allocated whenever possible to share resources.

Atos

Expanding Extreme Factory for Deep Learning

HPCaaS architecture

With the XCS web GUI

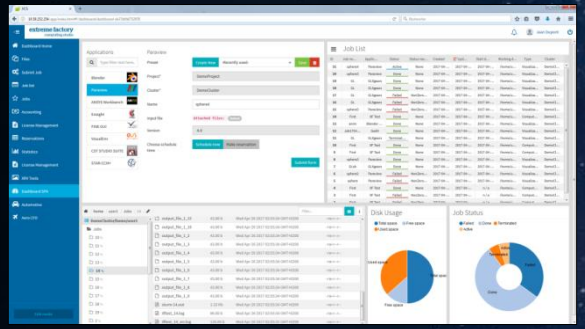

XCS web User Interface

HTTPS

XCS GUI
web server
- Dashboards
- Web Design

HPC
DCs*

HTTPS

HPC
REST API
web server

SSH

HPC
integration layer
- Batch schedulers
- HPC applications

HPC
Data base

Authorization
OAuth

Directory
service

*DC = Dashboard Component

Atos

HPCaaS and DLaaS architecture

With the XCS web GUI

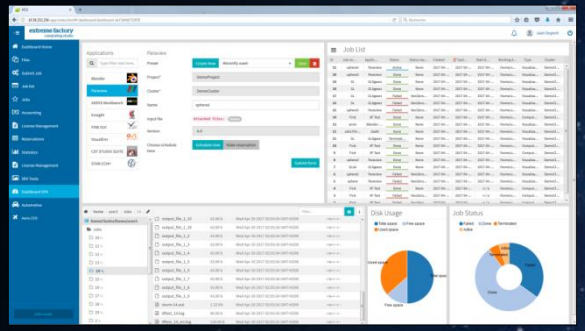

XCS web User Interface

HTTPS

XCS GUI web server
- Dashboards
- Web Design

HPC DCs

DLSS DCs

HTTPS

HPC REST API web server

HTTPS

DLSS REST API web server

SSH

HPC integration layer
- Batch schedulers
- HPC applications

SSH

DLSS integration layer
- DL frameworks
- DL datasets

HTTPS

Authorization OAuth

HPC Data base

Directory service

DLSS Data base

*DC = Dashboard Component

Atos

Deep Learning Self Service (DLSS)

Bringing the two threads together

Bull Extreme Fcatory

Codex AI Suite

Portal / XCS GUI

AI services API

Others...

Dataset Mngt API
Upload/annotate/delete

Model Mngt API
Create/update/delete model
Monitor & compare (DL insight)
Transpiler for model format
Optimise/self-tune a model

Training API
Schedule & run a training job

Inference API
Execute a model (tensorflow, n2d2)
Export (e.g. FPGA)

Catalogue
- Datasets
- Models (trained, pre-trained...)
- DL Frameworks

Orchestrator

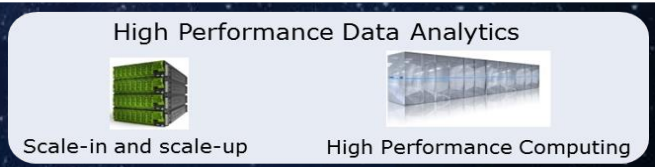

Atos


6
Summary

Summary

1. To continue drive scientific endeavour through simulation and modelling, increasing parallelism and the use of AI are inevitable
2. **The BullSequana XH2000:** a new hybrid architecture for HPC and AI which combines leading technologies in a highly dense water cooled system, delivering market leading TCO
3. **Codex AI Suite:** a new open framework for developing cognitive applications which couples HPC and AI for on-premise or cloud deployments
4. **BullExtreme Factory:** A 3rd generation portal for HPC and AI application deployment within a hybrid cloud services model
5. Visit our booth (24) to find out more

Atos

Thanks to our partner Intel

Proud to be Gold Sponsor of Computing Insight 2018

For more information please contact:

+44 7966 826667
andy.grant@atos.net

Codex AI Suite differentiators

End-to-end, fast development of complex enterprise use cases

Enable to easily combine ML, DL, analytics and data management to deliver complex and accurate use cases

Studio

Forge

Applications are infrastructure-agnostic

Enable applications to run on HPC, enterprise and Edge servers, on clouds and on-premises

Orchestrator

HPC business can deliver NG applications

From precision medicine, to imaging diagnosis, driver assist, autonomous maintenance, ...

ML Engine

Atos

Extreme Factory UK – working with STFC Hartree Centre

Public HPCaaS offering to UK academics and Industry

Key scientific areas include:

- Molecular modelling and material science
- Life sciences
- Virtual Engineering and Digital simulation
- Deep Learning

Atos Solution:

- BullSequana X1000 supercomputer
  - ~4 Pflop/s - Intel® Xeon® Scalable processors and Intel® Xeon Phi™, GPUs
- Extreme Factory (XCS and XRV)
  - Professional services to tailor the interface
- Business Development resources to co-sell

Image courtesy of STFC

Atos

JADE – working with Oxford University and The Hartree Centre

Public DLaaS offering to UK academics and Industry

Key scientific areas include:

- National Deep Learning Service
- GPU enabled Computing
- DL enabled HPC application development
- Prototyping and Algorithm Development

Atos Solution:

- BullSequana X1000 supercomputer
  - ~4 Pflop/s - Intel® Xeon® Scalable processors & Intel® Xeon Phi™, GPUs
- Extreme Factory (XCS and XRV)
  - Professional services to tailor the interface
- Business Development resources to co-sell

Image courtesy of STFC

Atos