



# HPC and Machine Learning convergence

---

Cédric Bourrasset, PhD

AI Product Manager

Atos Senior Expert & Member of the Atos Scientific Community

# Effective global delivery model in HPC, AI & Quantum

> 500  
Experts  
Worldwide

> 250  
Worldwide  
Customers

400 M€

TOP5  
HPC  
Players



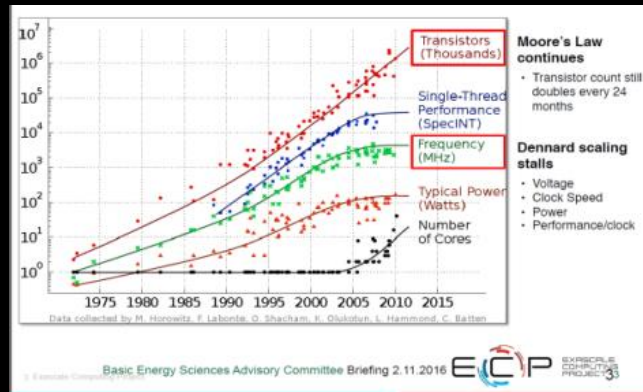
Supplying complex solutions to the most advanced industrial partners



Trusted partner for your Digital Journey



# Factors driving the convergence of HPC and ML



- ▶ Single threaded performance has plateaued
- ▶ Processor core counts continue to rise
- ▶ Increasing application performance now requires increasing levels of parallelism

- ▶ AI and data science have emerged as important new components of scientific discovery
- ▶ Huge volumes of data yielding significant improvements in accuracy, completeness and response time

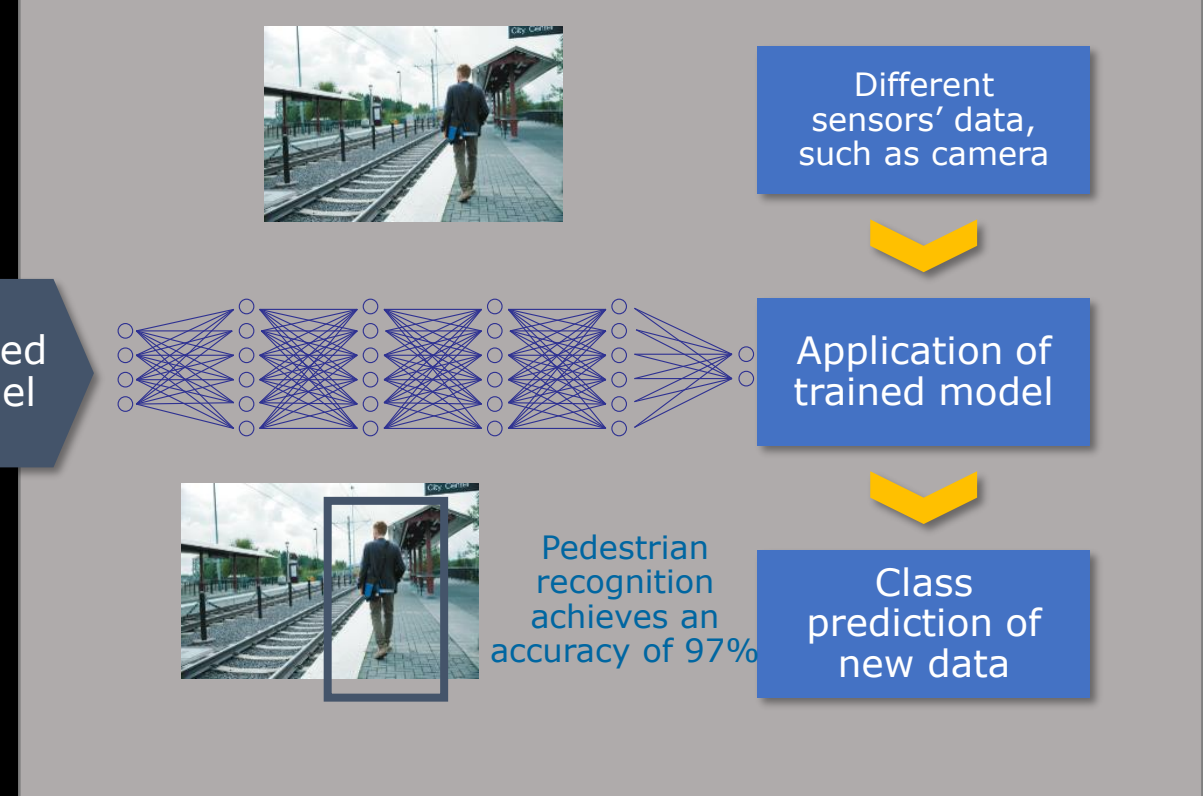
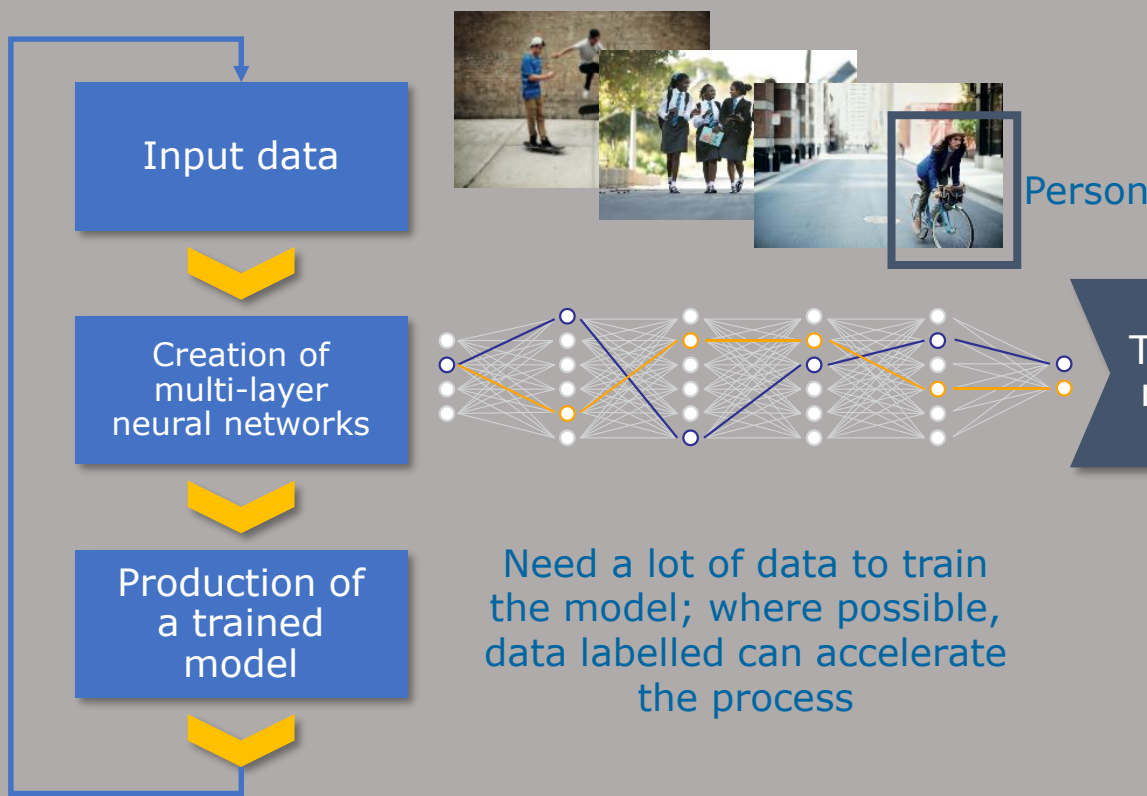


# Using deep neural networks

New requirements, but different from training to running

Phase 1 : Model training  
(Data center – heavy computation cost)

Phase 2 : Model prediction/recall/inference  
(Edge or data centre – almost instantaneous)



# Incorporating ML as part of the HPC workflow

## HPC


- Modelling and simulation of physical phenomena
- Track record of enabling scientific discovery
- Proven return on investment in multiple domains

- 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

## ML

- New methods to improve predictive accuracy,
- Delivering new insight into complex phenomena
- Allows work with previously unmanageable data sets

- Train models to improve accuracy and understand 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 value for the HPC/AI convergence

## HPC/AI Converged Infrastructure



## Customer benefits

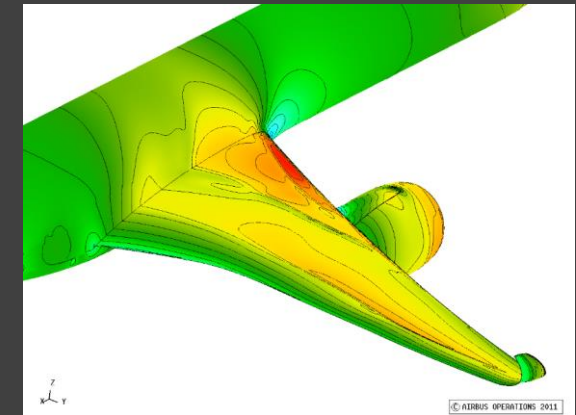
- Reduce TCO
- Drive pre-exascale systems with AI
- Self Learning production tools

## AI driven dynamic optimizations tool



- Scheduling & runtime
- Job Placement
- Communication
- Data Movement

## Start rewriting HPC app to integrate ML



- Reduce time to resolution
- Increase simulation with data-based approach

# BullSequana XH2000

A flexible architecture hybrid HPC/ML needs



Mix & match compute 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



NVIDIA



MARVELL




XILINX



10/25/40/50/100GbE Open Ethernet


Seamlessly integrated with Atos Codex AI Suite and Cloud Platforms



Atos

# BullSequana XH2000

## Increased flexibility and agility



42U Cabinet Back

Front

42U Cabinet Front view

42U Cabinet Rear view

**PDU + Power controller**

**up to 6 x DLC PSU shelves**

**2 x Leaf Eth switches**

**up to 10 IB/BXI/High Speed Ethernet switches**

**4 to 20 compute blades (front)**

**up to 12 compute blades (Back)**

**up to 3 Hydraulic chassis**


90 KW  
n+1 redundancy

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

# BullSequana XH2000 Blades

The best of all worlds, available in one supercomputer

EPYC Rome



Xeon SP  
Skylake / Cascadelake



Volta P100/V100  
NVLink2




ThunderX2  
ThunderX3 when available

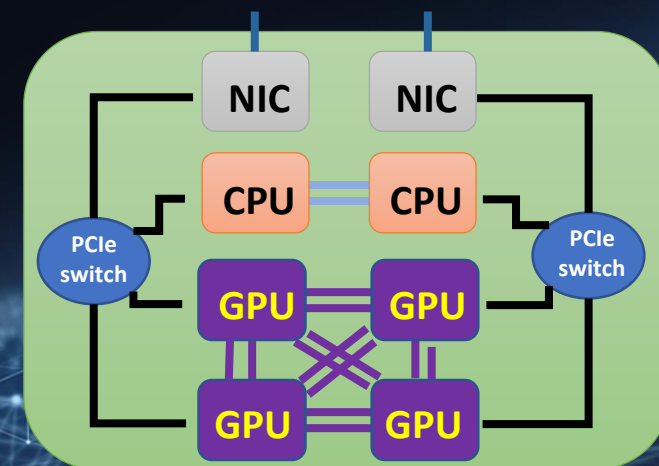


# BullSequana XH2000 Blades

## BullSequana X1125: Nvidia Volta P100/V100 NVLink2



- 1U form factor
- Direct liquid cooling
- 1 compute node per blade with:
  - 1 CPU Board with 2 sockets
  - 1 GPU board supports up to 4 GPUs
  - 2 I/C mezzanine boards (BXI, High-speed Ethernet, InfiniBand EDR, HDR or HDR100)
  - Optional 1 x 2.5" 7 mm SATA SSD

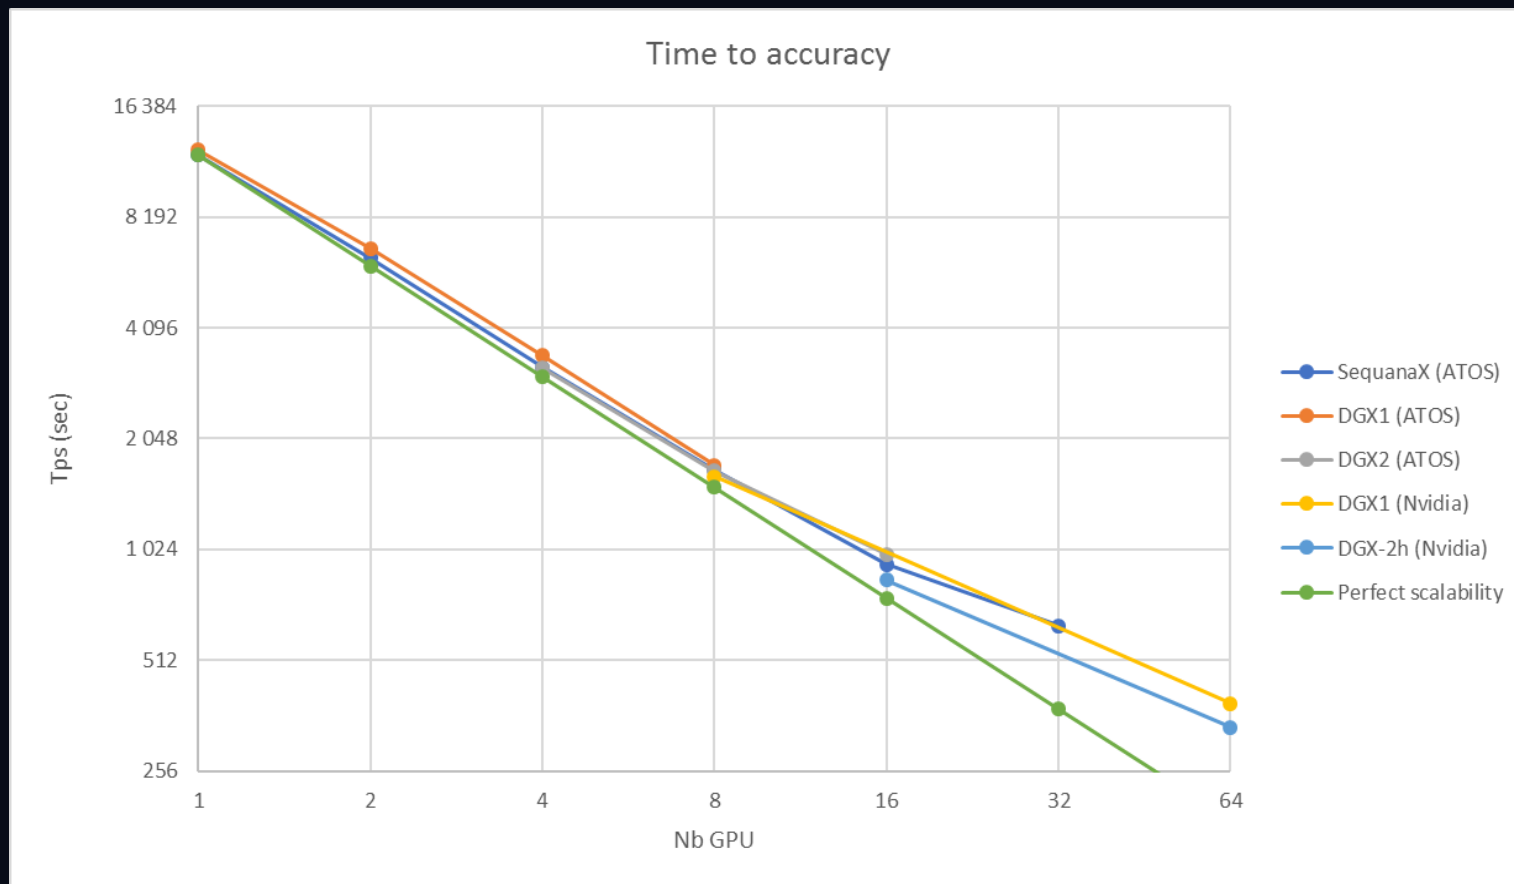




# Enable AI in HPC World

## BullSequanaXH2000 – Some AI performance numbers

MLPERF v0.5 – SSD benchmarks – NVIDIA code – No parameters tuning






- 4% impact 2 nodes SequanaX vs DGX-1 (Mesh NVLINK)

- 9% impact 4 nodes SequanaX vs DGX-2 (NVSWITCH)

- Similar scalability after 4 nodes

# Body forces use case

## Rationale



External aerodynamics simulation with black box engine model



How to perform external aerodynamics simulation with:

- Good engine modelling.
- Computational cost close to classical external aerodynamics simulation.

Engine simulation without the airplane

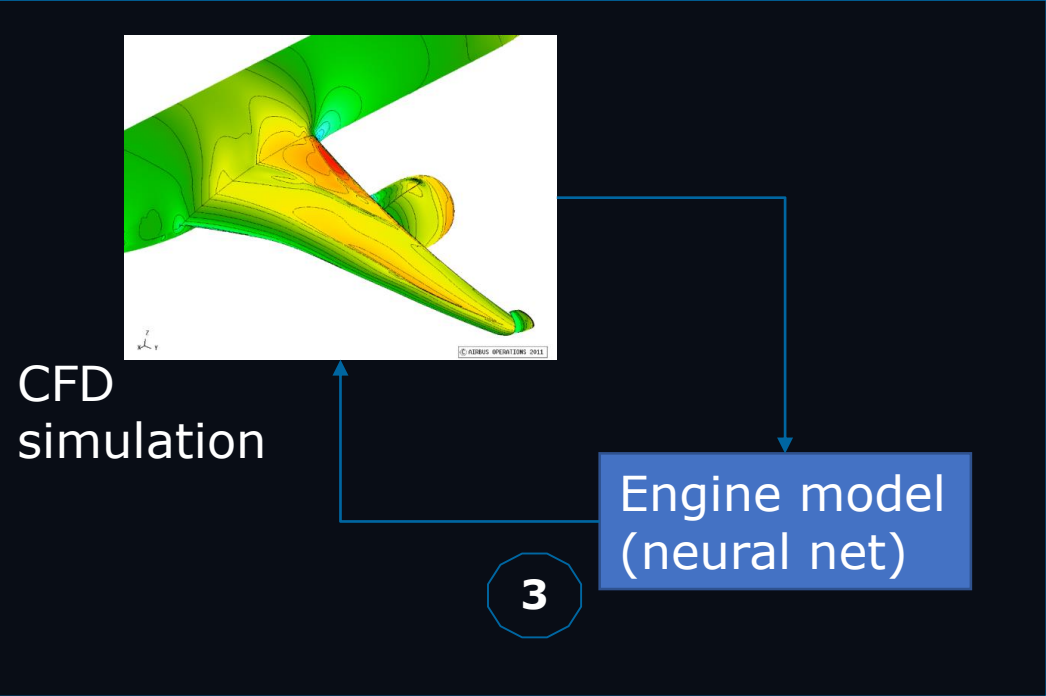


Body force modelling

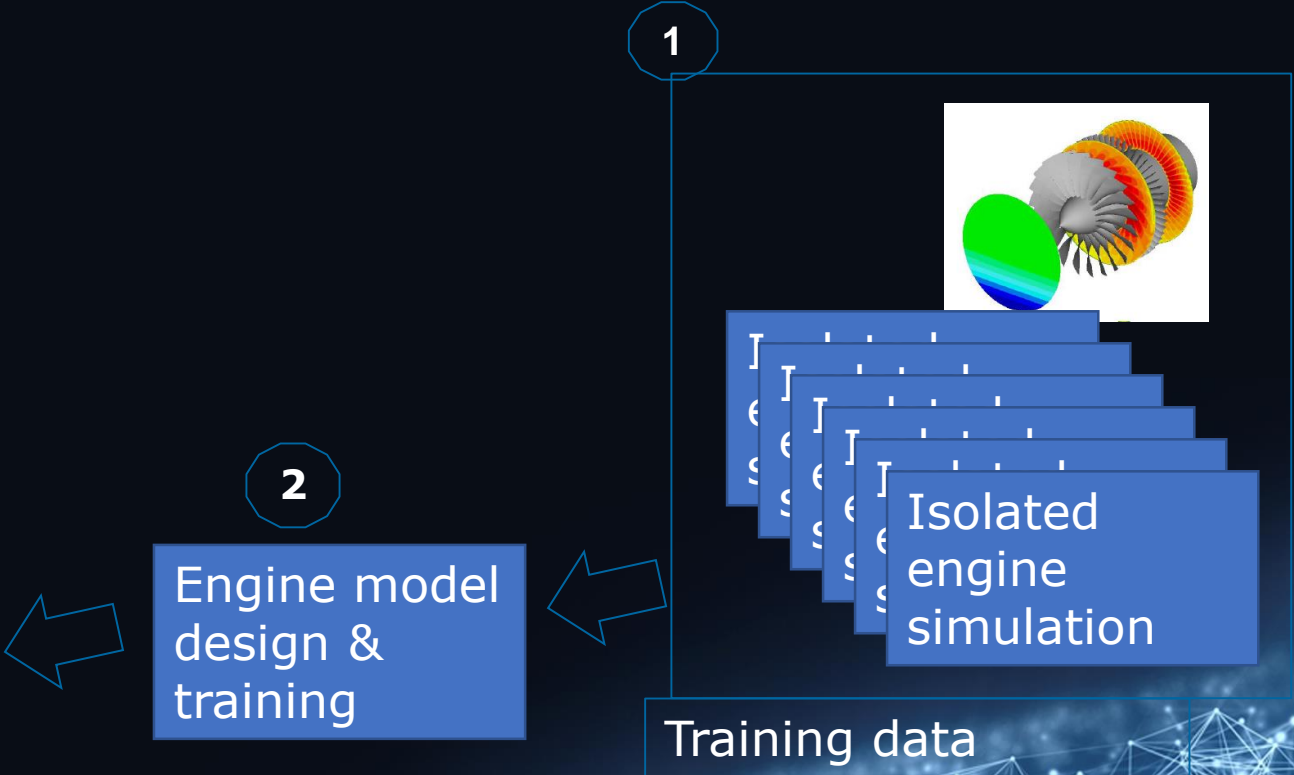


# Body forces use case







## Overview & Challenges



Simulation iterative process



# ML Frameworks – typical usage today

	Framework packaging	Orchestration
Data Science		
HPC	   python™	


There is a need for a ML Toolkit that works for both HPC and enterprise customers

# Enable the ML in HPC World

## Codex AI Suite - Fast ML Engine

HPC compatible

Deploy AI Workloads on top of HPC Scheduler



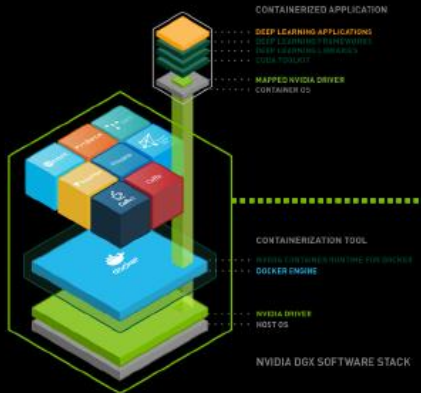
Codex AI Suite  
Fast ML Engine

NVIDIA NGC certified

Leverage the best of AI frameworks



NGC DL SOFTWARE STACK





Ready for hybrid computing

Extend HPC/AI capabilities to the cloud & the edge

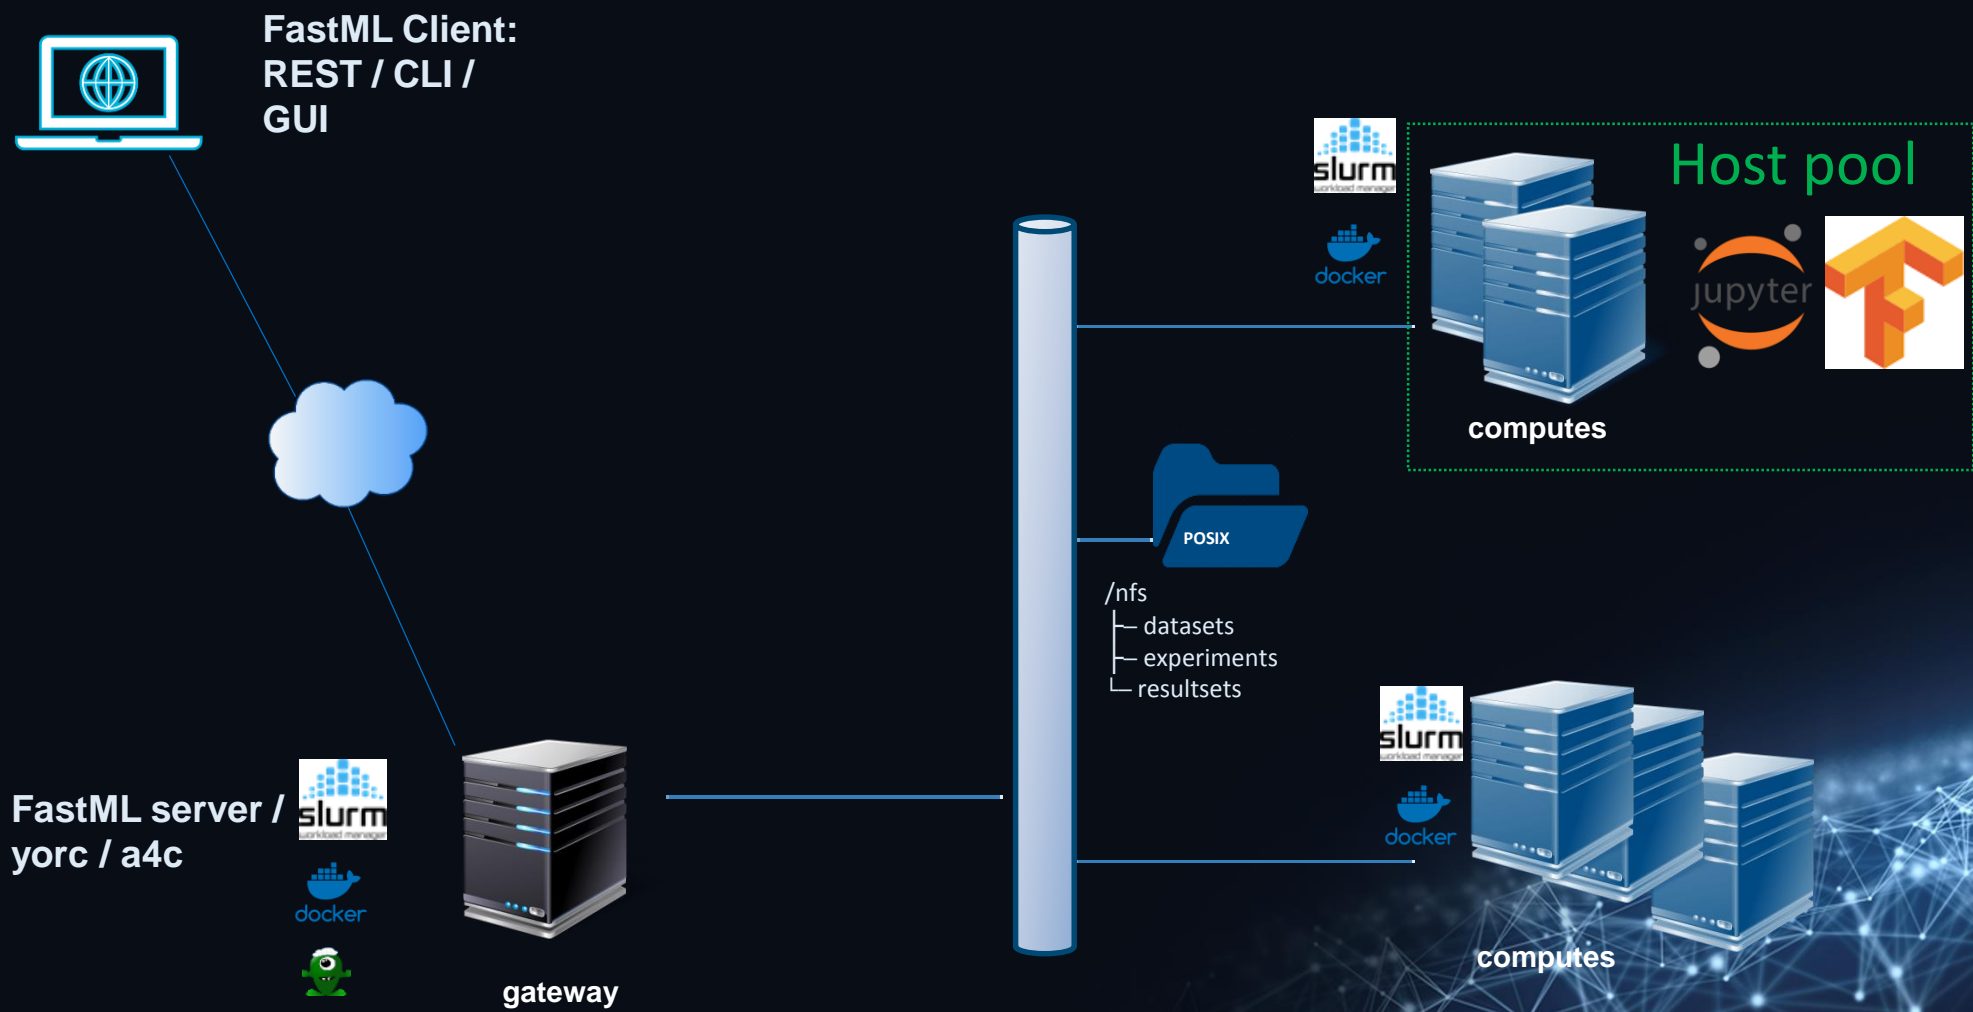


Codex AI Suite  
Orchestrator

# Fast ML Engine in ML activity



# Typical AI Suite Configuration on HPC cluster



# Hybrid Orchestration for AI and more

- ▶ Manage application lifecycle on multiple environments
- ▶ K8S: CaaS
- ▶ YORC : IaaS/CaaS/HPC
- ▶ Similar in Acumos (<https://www.acumos.org/>) LF project on AI
  - Open source: <https://github.com/ystia/yorc>

## Orchestrator

Public Cloud



Google Cloud Platform



Virtual  
Machines



openstack™  
CLOUD SOFTWARE



Container as a  
Service



kubernetes



HPC Job  
Scheduler



slurm  
workload manager




# Fast ML Engine Release v3.2.0 (July, 2019)

## Existing Features

**Training Management**



**Model Management**



**Framework Management**



**Datasets Management**




## New Features

**Distributed Training**



**Monitoring Management**



**Notebooks Management**



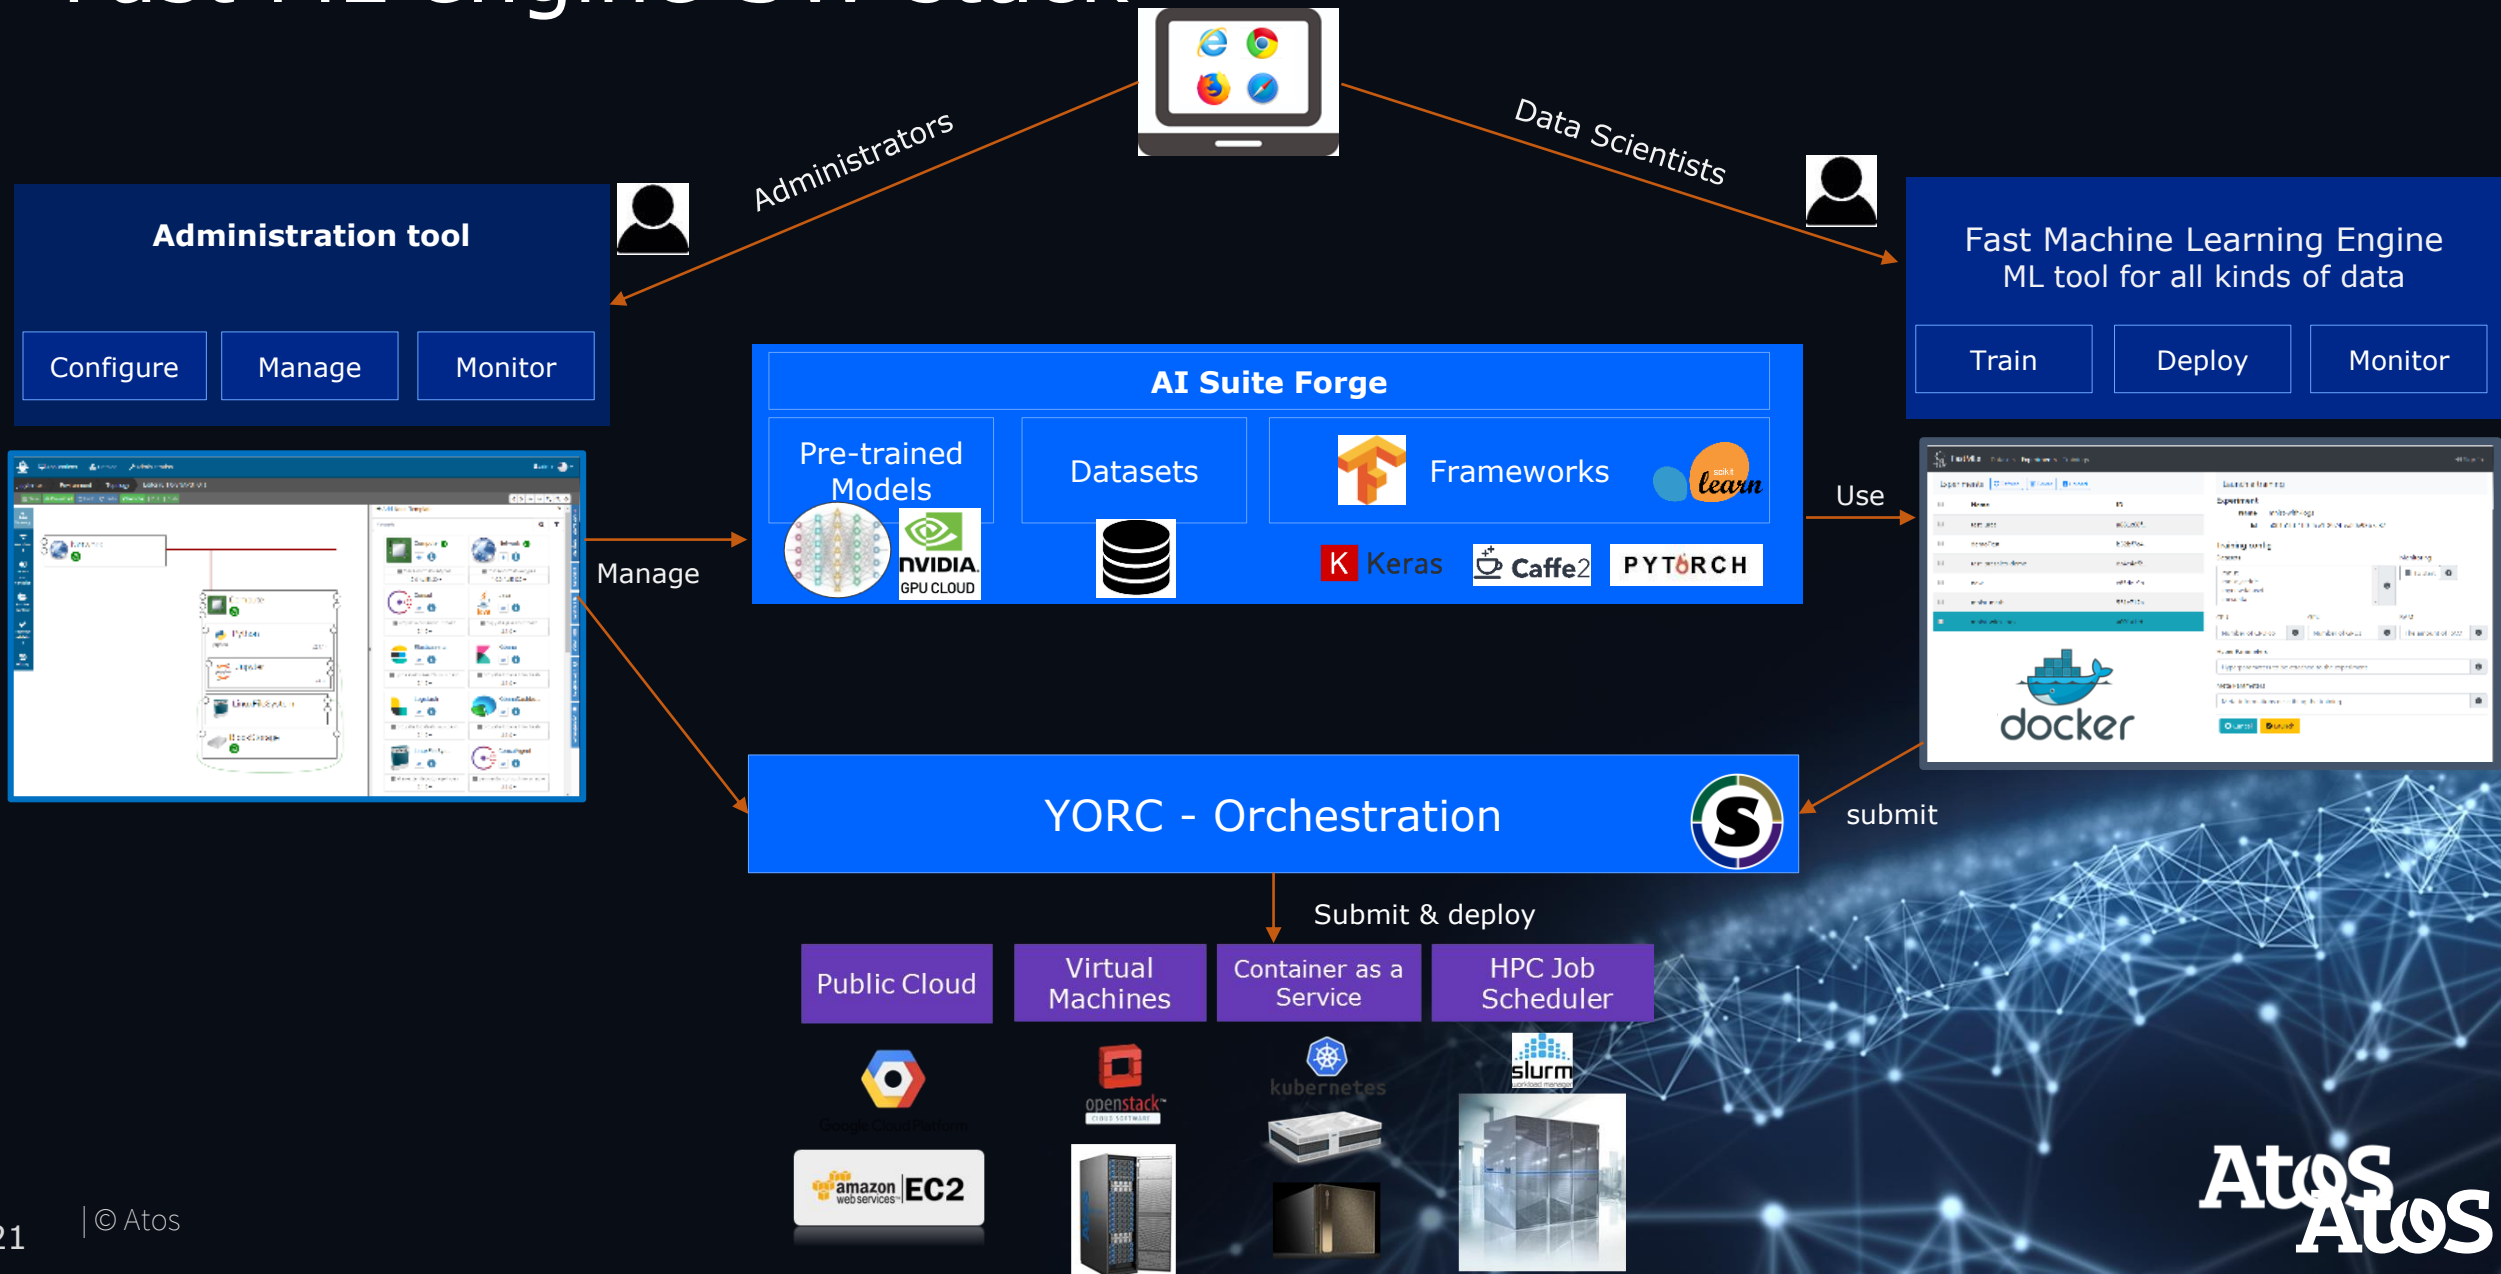
**Security**



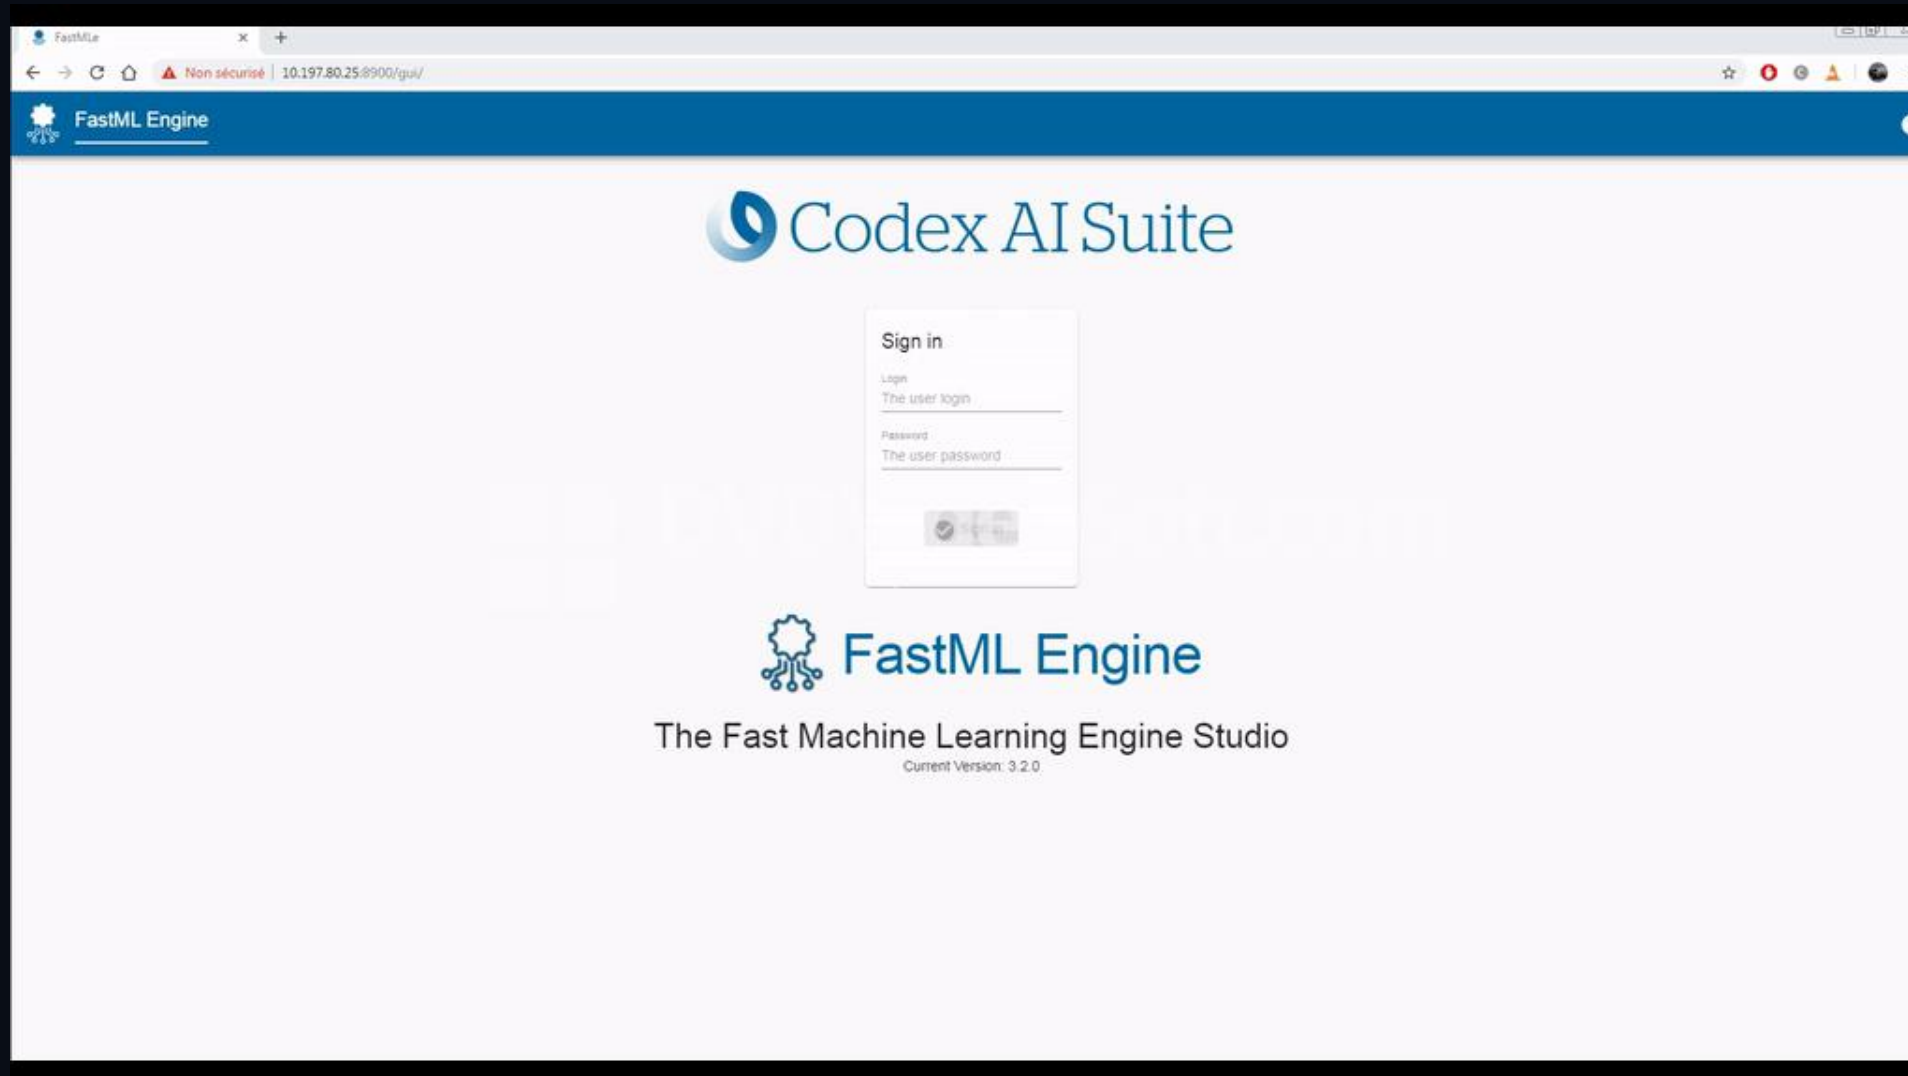
**Containers support**



# Fast ML engine SW stack



# Demo



# Summary

1. To continue drive scientific endeavour through simulation and modelling, increasing parallelism and the use of AI are inevitable
2. **BullSequana XH2000**: a new hybrid architecture for HPC and AI which combines leading technologies in a highly dense water cooled system, delivering market leading performance and 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. Visit our booth to find out more

# Thank you

## and Questions

**Cédric Bourrasset, PhD**

**AI Product Manager**

Atos Senior Expert & Member of the Atos Scientific Community



**Atos**