

EVIDEN



ISC

High Performance

IMAGINE
TOMORROW

Accelerated Hybrid Computing Any way you want to run

Eviden Breakfast Briefing

Eviden HPC Leadership Team
24/05/2023

© Eviden SAS

an atos business

EVIDEN

Content overview

01

**Eviden – the
supercomputing leader**

02

Our recent success

03

Product update

04

As-a-service enablement

05

Conclusion



EVIDEN

01 Eviden – The Supercomputing Leader



Emmanuel Le Roux
SVP, Head of Advanced Computing & AI

Eviden ID & leadership in HPC



Circa €5B revenue

An Atos business that will become an independent company in late 2023


6 Business segments

Digital Transformation, Smart Platforms, Cloud, Advanced Computing, Digital Security and Net Zero

57,000 engineers and problem-solvers in **45** countries.

Supercomputer leadership

#1 Europe, India and S. America
#3 worldwide




43 systems



Eviden – The Most Advanced High-end HPC, AI & Quantum Portfolio

On-prem and in-the-cloud

as-a-Service Enablement




Application & Optimization



AI-embedded HPC Software Environment




High-end Hybrid HPC, AI & QC Hardware



Eviden HPC Hardware – Pioneer, Innovation, Performance @scale

High-end Hybrid HPC, AI & QC Hardware



- New Quantum Computing offering called **Qaptiva** (ex QLM) is
 - **Full-fledged environment** to develop and run QC applications
 - Broad partners ecosystem on **Software** with domain specific libraries and **QPU HW providing** aaS access
 - New **programming model** for hybridization between HPC and Quantum
 - New **Software** to distribute QC simulation on top of any HPC Systems (Q4 2023)
 - Quantum Computing Appliance, the QLM next generation named Qaptiva 800 (Q3 2023)
- Develop, Manufacture and Deliver **Exascale and Post Exascale** production systems
 - Build on strong momentum – 43 TOP500 machines/ better score ever, large wins in Europe, India and Americas
 - Best in class XH3000 with first major deployment starting in Q2 2023 with Barcelona SuperComputing and CEA
- **Sovereign by design** with more and more EU-Components in Eviden SuperComputers:
 - Launch of **BXI V3/V4** to develop next gen European Interconnect for HPC and AI
 - **Investment in SiPearl** to deliver European processor fully integrated into first European Exascale machines

EVIDEN

Eviden HPC Software – Energy Efficiency and Performance @scale

AI-embedded HPC Software Environment



- Continuous improvement: HPC Management Software designed for every HPC needs from small system to Exascale Supercomputer
 - Install, Configure and Operate HPC clusters at all scales
 - Leveraging Artificial Intelligence using Machine Learning algorithms to anticipate Hardware Failures
- Performance Improvement by monitoring critical HPC components:
 - Smart Data Management Suite provides monitoring interface to understand I/O patterns and accelerators to speed up / optimize Data transfer and improve Application performance.
 - Smart IO optimizer & Smart IO Libraries using machine learning algorithms to dynamically change malformed IO, replacing them by accelerators to boost overall application performance by up to 30%, requiring zero application change.
- Major investment to Monitor and Optimize Energy/Performance Efficiency
- Smart Energy Management Suite to understand, control and optimize energy efficiency
 - I/O optimizations is also a key component to reduce energy footprint of HPC applications.

Eviden CEPP – Center for Excellence in Performance Programming

Collaborations and services added values

Application &
Optimization



- Team of Experts in Computer Science and Applied Mathematics dedicated to Applications Performance

- CEPP does have 5 sub-centers of Excellence




For HPC Services:	Trainings, Porting, Profiling, Optimizing, Developing for Partners and Customers
For Collaborations:	Mainly science research funded projects where CEPP does HPC Services & Co-design
For IT Partners:	With Key IT partner in HPC AI & Quantum
For Customers:	Taylor-made collaboration on customer KEY applications and Topics

- And a final one who is most of the time on support of the 4 previous
- For Eviden R&D collaboration: The link between customers and Eviden R&D (HW & SW) for Application driven Co-design

Eviden Nimbix – Any workload, Any application, Any way you want to run

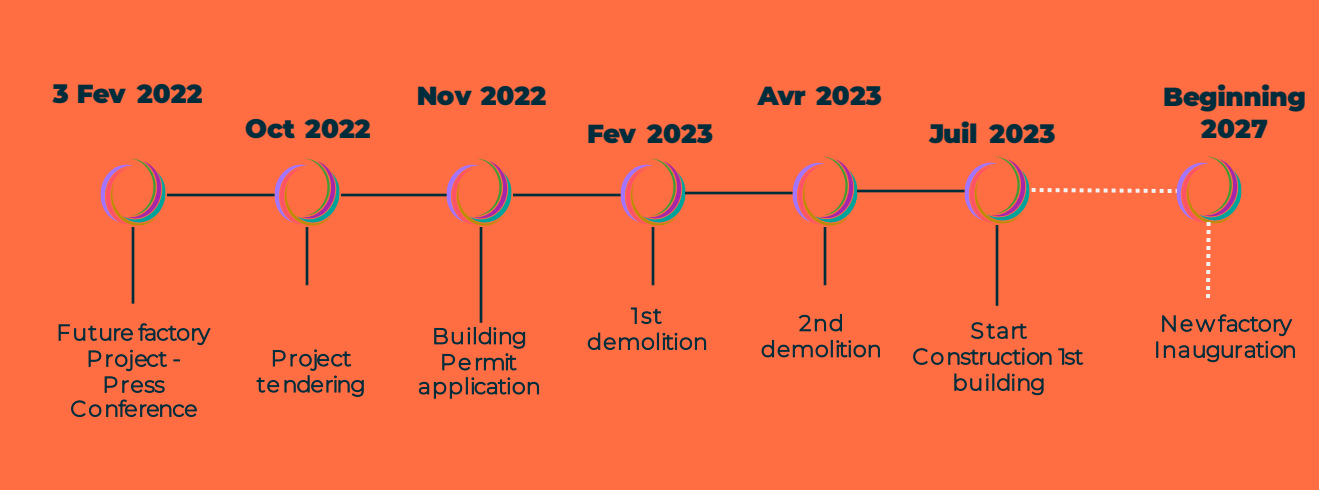
**as-a-Service
Enablement**




- Fully Elastic HPC service across millions of machine enabling **HPC capacity on demand** using large Hyperscaler or BareMetal dedicated instance
- **Federated Bare-Metal Supercomputer** providing option for Public Federated Supercomputing and enabling Academic, research and Industry collaborations.
- **AI and Quantum as a Service** with Industry leading performance systems, advanced development Environment and Quantum emulation platforms
- **BareMetal HPC as a Service** using Subscription-based hardware platforms enabling to shift from CAPEX to OPEX for dedicated bare-metal HPC / Supercomputing capacity.


Re-enforce the European HPC, AI & Quantum Self-reliance

Expanding HPC manufacturing capability in Angers, France







Double capacity with only 30% more staff



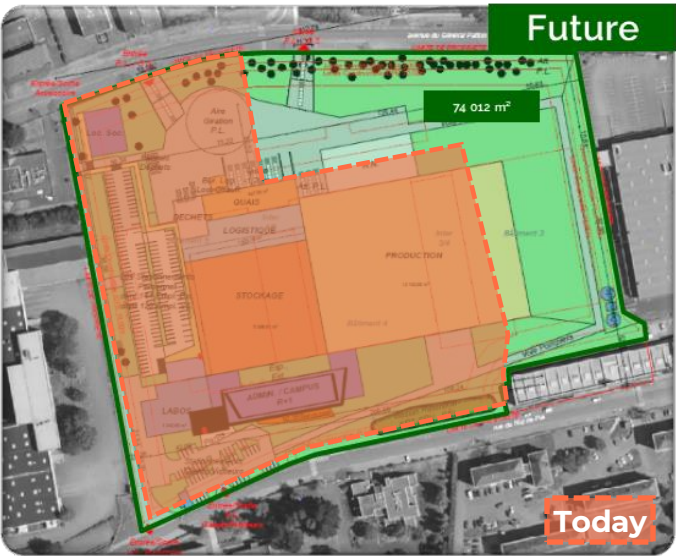
Industry 4.0 and Cybersecure



+30% productivity, -30% maintenance costs, +30% stock reliability



Intelligent and sustainable: recovery of waste heat, particularly from test & operations



EVIDEN

02 Updates and Recent success



Andy Grant
Group VP, HPC, AI and Quantum global sales

New customers orders since SC22

Strong momentum in all Geographics



Hartree Centre
Science & Technology Facilities Council



Mercedes-Benz



MAX PLANCK
COMPUTING & DATA FACILITY




पृथ्वी विज्ञान मंत्रालय
Ministry of Earth Sciences



CINECA – Leonardo (#4 in Top500 since SC22)

Hybrid Pre-Exascale Supercomputer



"The Leonardo supercomputer is the result of our long-term commitment to pushing the boundaries of what a modern exascale supercomputer can be."
Sanzio Bassini, Director of the HPC department at CINECA

10 ExaFlops FP16 AI workload

320 PFlops (Rpeak)

250 PFlops (Linpack)



EuroHPC
Joint Undertaking

>4 PFlops (HPCG)



EVIDEN
an atos business

ParTec
Modular Supercomputing

NVIDIA

intel
XEON

ddn

EVIDEN

CINECA Installation

Fantastic effort to deliver and Build Leonardo



EuroHPC
Joint Undertaking

May 2022

November 2022

Today



174.6 Pflops
#4 Top500



238.7 Pflops

EVIDEN



EuroHPC/Barcelona Supercomputing Center

Mare Nostrum 5 -Empower world-changing scientific breakthrough

MareNostrum 5 (MN5)

The acquisition of MareNostrum 5 will enable world-changing scientific breakthroughs such as the creation of digital twins to help solve global challenges like climate change and the advancement of precision medicine. In addition, BSC-CNS is committed to developing European hardware to be used in future generations of supercomputers and helping to achieve technological sovereignty for the EU's member states.



Mateo Valero,
Director of
BSC-CNS

Eviden's role

- Prime leader for a consortium of technology suppliers
- Design, deliver, install and configure BullSequana XH3000

EVIDEN
an atos business



314 PFlops, (Rpeak)

163 PFlops Rmax GPUs – Nvidia H100

35 PFlops Rmax CPUs – Intel SPR

200+ Petabytes of storage

EVIDEN
an atos business

Lenovo

IBM



EuroHPC
Joint Undertaking

BullSequana XH3000



EVIDEN

Max Planck Computing and Data Facility (MPCDF)


Leading German Computational Research Lab



The Max Planck Computing and Data Facility (MPCDF) provides high-level support for the development, optimization, analysis and visualization of HPC and HPDA applications to German research institutions with high end computing needs.

This includes:

- Astrophysics,
- Fusion Research
- Materials and Biosciences,
- Polymer Research and
- Theoretical Chemistry



Max Planck Computing and Data Facility (MPCDF)

First Genoa and MI300 Deployment in Europe



Customer Challenge


- **3x** performance boost for German Research Community
- Replace the aging Max Planck Cobra supercomputer Intel Skylake Processors with OmniPath interconnect
- Provide a solution that delivers a significant performance boost for both HPC and AI workloads



Eviden Solution

BullSequana XH3000 with:

- 768 AMD EPYC “Genoa” nodes
- 192 AMD Instinct MI300A APUs
- Nvidia networking
- IBM Spectrum Scale storage solution.
- Maintenance and Support
- Collaboration support with AMD
- >30PF HPL



What the Customer says

- “**The computing power required by scientific research is ever increasing** and we see an unabated need for high-performance computing capacity. We want to **provide the best possible support to our researchers in their work** and have therefore decided to modernize our high-performance computing complex. **With Atos and AMD, we have the right partners for this.** The new solution will certainly meet our demands and once again advance science in leaps and bounds.”

Prof. Erwin Laure,
Director of MPCDF



EVIDEN




© Eviden SAS

Large Southern Hemisphere Weather Customer

Success Stories

Customer Challenge

- Improved climate change simulations,
- Coupled Ocean-Atmospheric Climate Forecast Models
- Prediction for extreme weather event, lows and depressions; local severe storm and cyclones



EVIDEN



BullSequana X

© Eviden SAS

The customer provide services for weather, climate, ocean and coastal state, hydrology, seismology, and natural hazards

x 10 Production Speed-up

~5000 AMD Milan nodes,
18PFlops, ~8PB Storage
The solution will be deployed across two sites

>\$100M

ClimateRean
Climate Change Initiative (CCI)

EVIDEN
an atos business

AMD

NVIDIA


ddn

Eviden ThinkAI - Customer reference

University of Linköping



- AI Center of Excellence
- Research support for distributed trainings
- 12 months of consulting on security and system design
- NVIDIA DGX SuperPOD (94 DGX)
- Eviden Codex AI SUITE




EVIDEN

03 Product Update



Eric Eppe
Group VP – Portfolio and Strategy HPC, AI, & Quantum

EVIDEN BullSequana XH3000 is definitely here !




EVIDEN

Eviden SAS

BullSequana XH3000 is launched with already 10 blades


The most complete portfolio in the industry (and more to come)

CPU Compute blades




BullSequana X3140
(codename TRIO)

intel. Sapphire Rapids




BullSequana X3420
(codename MATO)

AMD Genoa



BullSequana X3610
(codename MILO)


SIPEARL Rhea 1



BullSequana X3620


SIPEARL Rhea 2

GPU Compute blades




BullSequana X3145-H
(codename TIYA)

nvidia HGX 4-H100



BullSequana X3515-H
(codename NEBA)


nvidia HGX CG4



BullSequana X3406-3
(codename MANA)


AMD MI300-A/-C

High-speed Interconnect switch blades




BullSequana X32B2-0100
(codename W3B1M)

EVIDEN BXI V2



BullSequana X35IN-0800
(codename W3NM)

nvidia InfiniBand NDR



BullSequana X35ET-0200
(codename W3EM)

nvidia High-speed Ethernet

BXI: BullSequana eXascale Interconnect

The European born ExaScale Interconnect

The first European Interconnect product with mature technology



Present in the TOP500 and Green500 List on some of the world's fastest supercomputers



10-year of development of superior bandwidth and low latency communication network technology optimised for HPC



A powerful performance of new BXI generation developed by Eviden experts and teams



A combination of great complexity and innovation generated accurate results immediately



Working with reliable partners & clients



Green technology inside ultramodern HPC at a competitive price

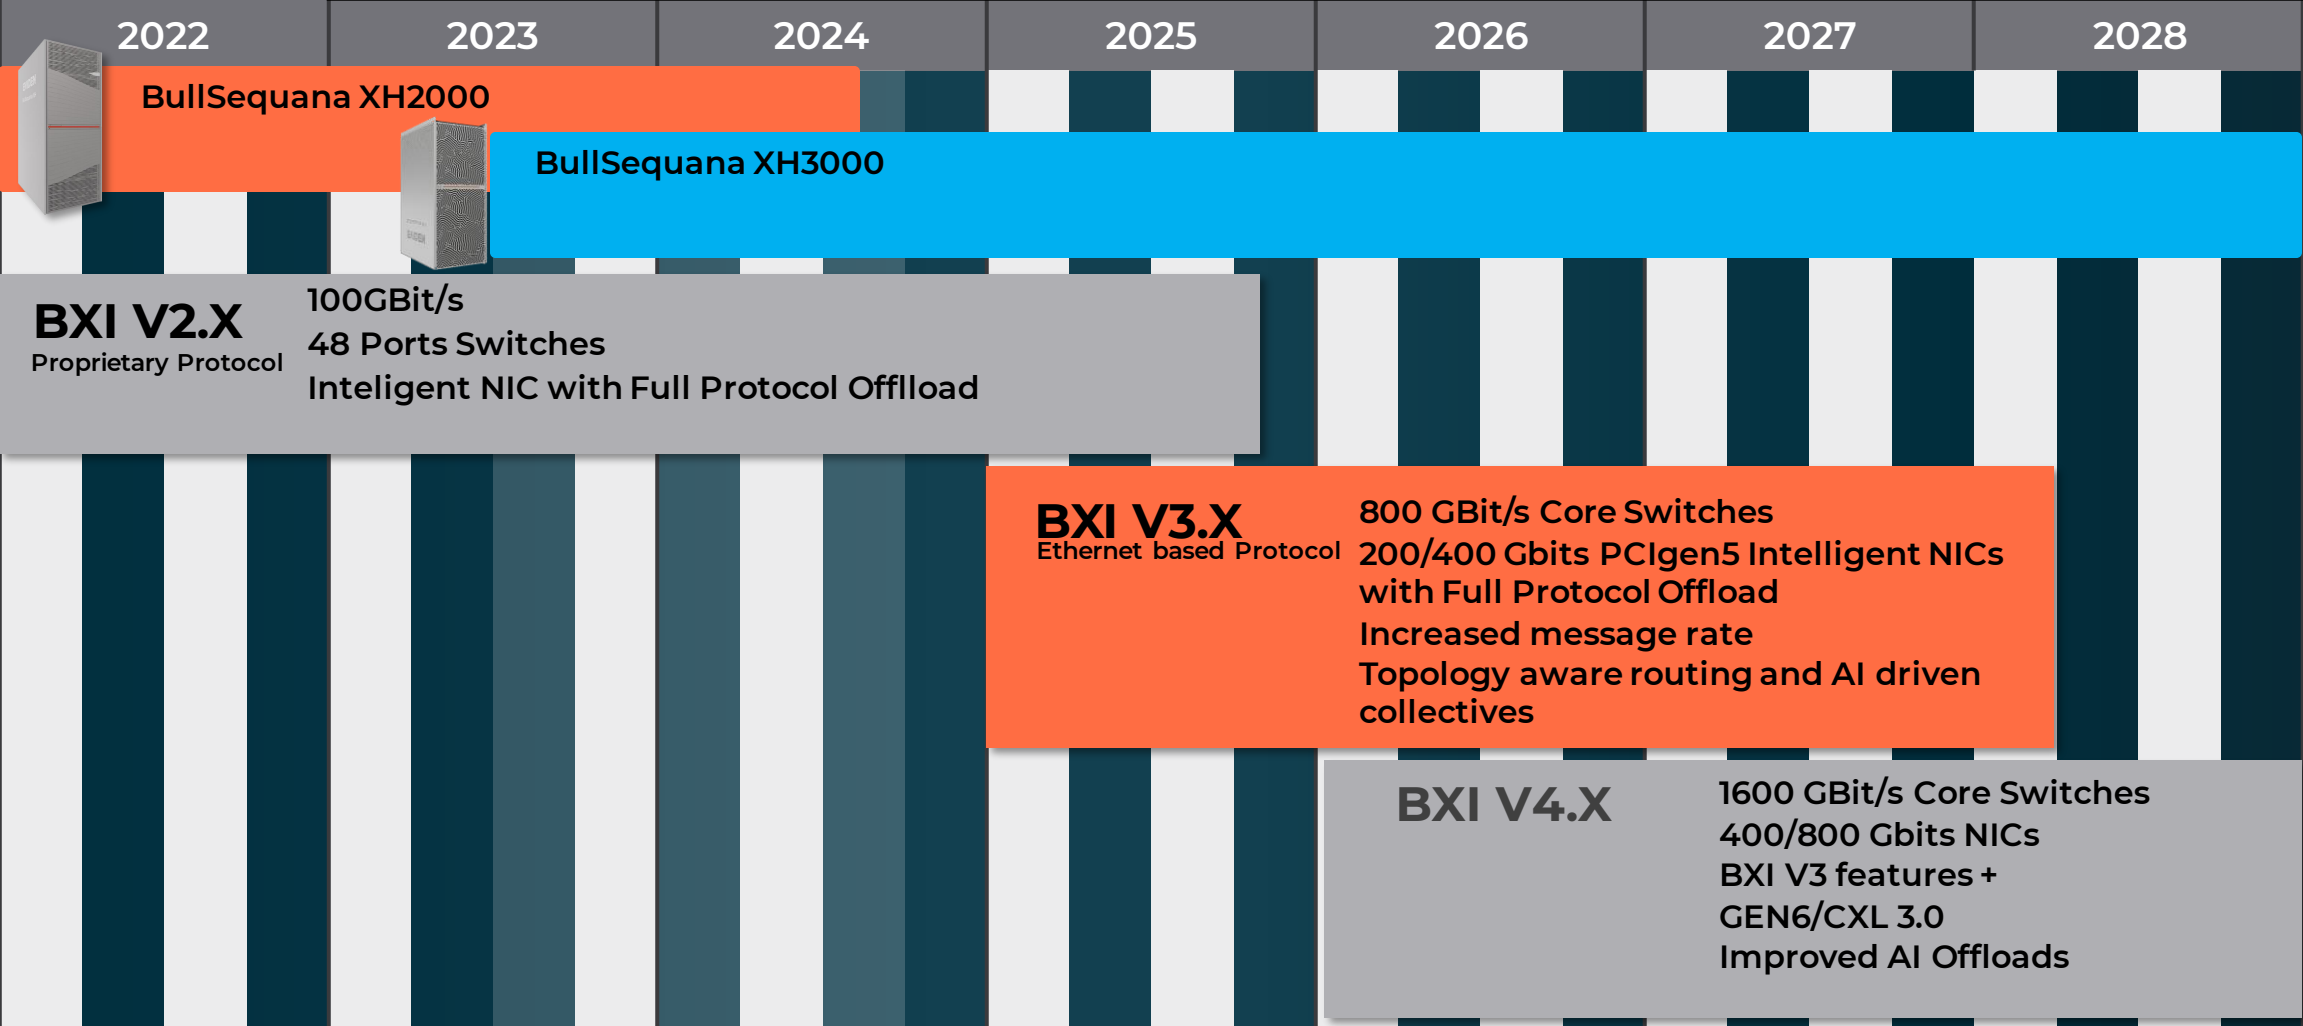


Incorporated with more than 1800 active patents and support a variety of demands



BXI Next Generation

A future for very high performance & low latency Interconnect for HPC & AI workloads



EVIDEN

Qaptiva™

From latin “Captivare”

*“An irresistible attraction
because of their extraordinary
qualities”*



Qaptiva™ : The best route to real world QC applications

Now, not later!

**Hybrid Computing
using Quantum
requires Qaptiva
Commercial Apps are
getting results now.**

EVIDEN is uniquely positioned on Hybrid Computing and on corporate applications with its QC partners. With new appliances, a new hybrid computing framework, partner QC SW & QPUs, and a powerful consulting practice

EVIDEN
Qaptiva
Access

EVIDEN
QAPTIVA Access

EVIDEN
Qaptiva
800

EVIDEN
QAPTIVA 800

EVIDEN
Qaptiva HPC

Simulation SW
running on
BullSequana
SuperComputers

EVIDEN
BullSequana XH

EVIDEN
Qaptiva aaS

Application
Platform



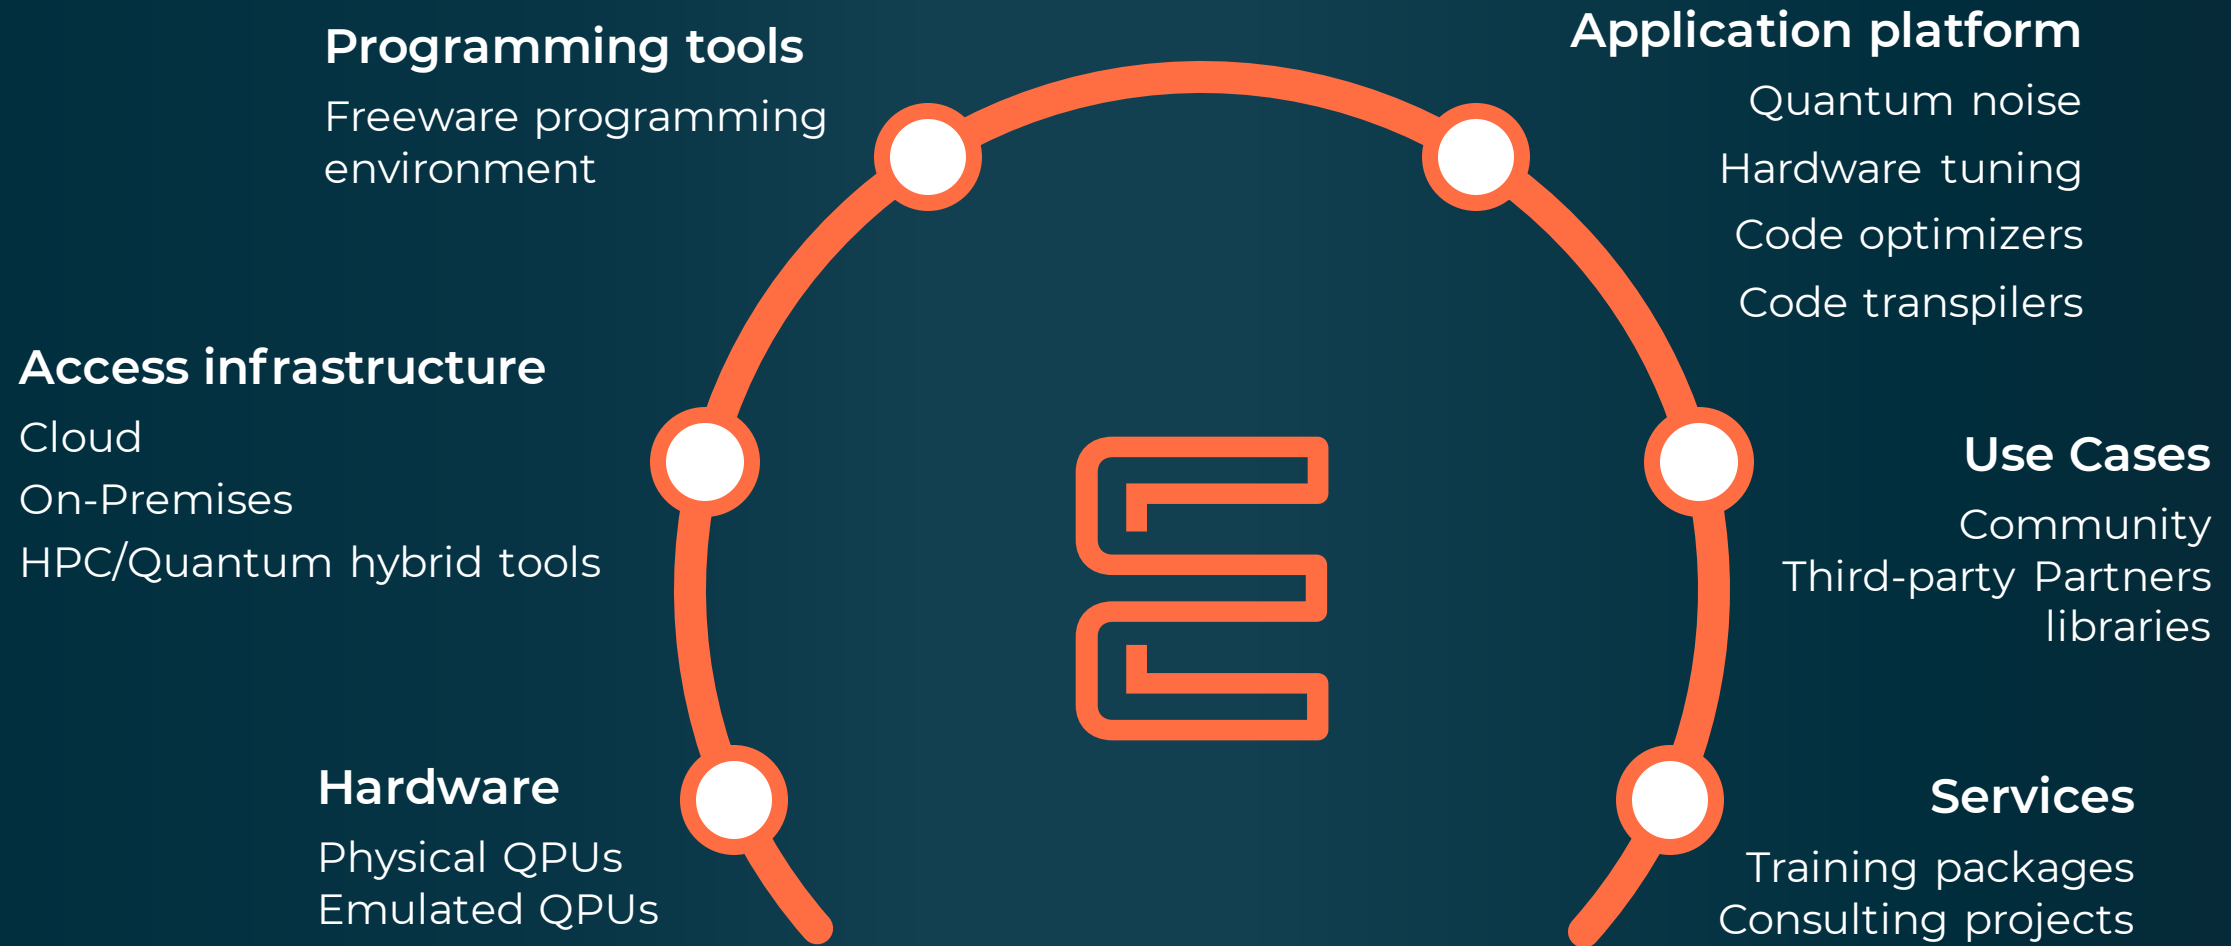
Partner QPU

Real World Quantum
Hardware



Qaptiva™ offering

End-to-end portfolio

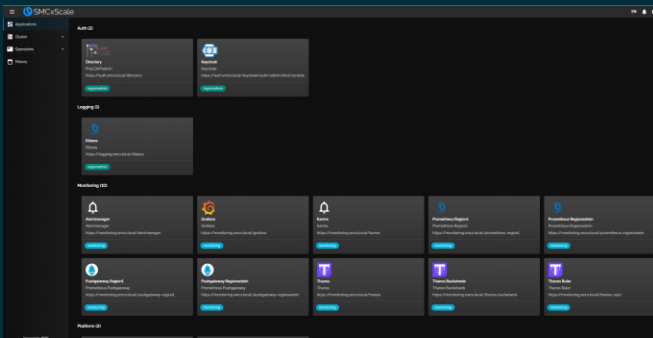


HPC/AI software portfolio

A solution for all customers needs

Cluster essentials

- Hardware tooling to bring HW to life
- 2 management stacks depending on the cluster size: SMC / SMC xScale




- Workload scheduler: Slurm & add-ons
- Parallel programming: OpenMPI & add-ons
- Containers: Sylabs Singularity


Cluster optimization

1. Understand
2. Control
3. Optimize


Fabrics
(IBMS)



Energy
(SEMS)




Data & I/Os
(SDMS)

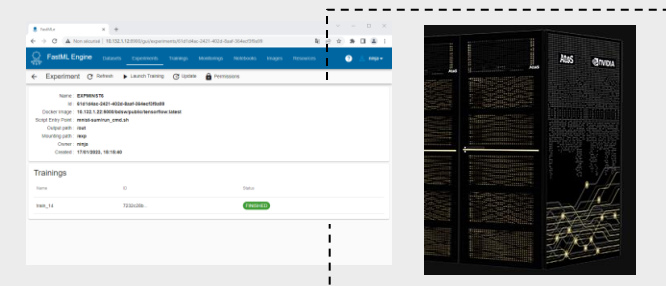


Cluster openness

- Simplify and expand cluster capabilities and location – JARVICE XE/NSS



- Remove complexity for data scientists - FMLe



EVIDEN

04. As-a-service

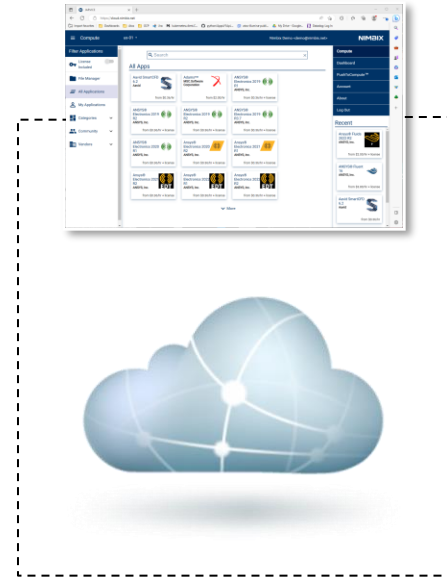


Leo Reiter
CTO, Eviden Nimbix Cloud

Commercial Application of JARVICE™ XE Cloud Pack

A large North American consumer products company seamlessly enables its global staff of simulation engineers on its own hyperscaler tenancy.

- JARVICE™ XE Cloud Pack connects customer's cloud tenancy to Nimbix Cloud
- Professionally managed, point-and-click SaaS simulation platform
- Enables ease of use and convenient “burn down” of hyperscaler usage commitments
- Single sign-on and full compliance with company's security policies
- Cost controls and business-level accounting



JARVICE™ XE Federation Pack

Accelerate Adoption of Supercomputer Federations

JARVICE™ XE Federation Pack is a new software module to enable Nimbix Federated capabilities on managed deployments.

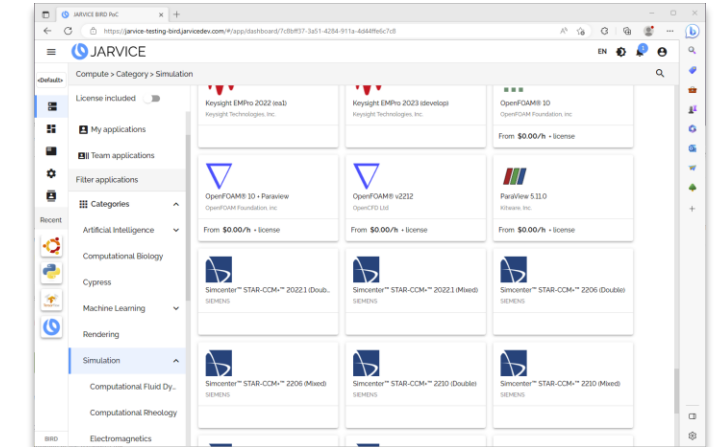
- Connect any Slurm or Kubernetes infrastructure, on-premises or cloud, to a private HPC/AI federation
- “Single pane of glass” for engineers/researchers and operators alike
- Out of the box support for Eviden HyperHub Application Marketplace with 100’s of click-to-run workflows
- Secure, container-native approach to federated HPC and AI



New Capability – Slurm Support

Connect Slurm clusters to JARVICE™ XE Deployment

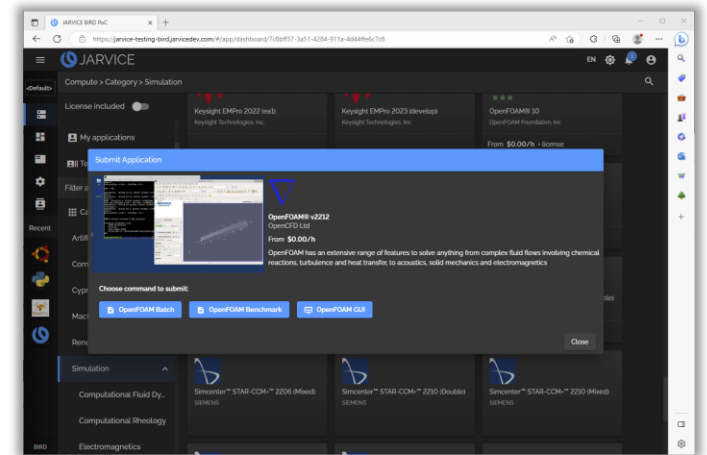
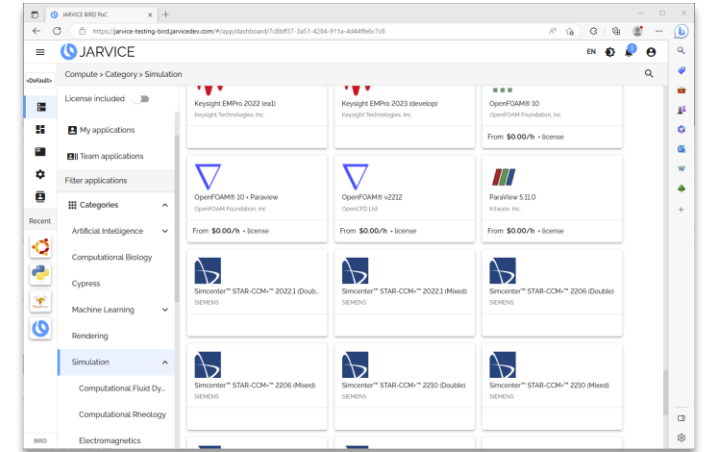
- Combine Cloud-native Kubernetes Capabilities with On-premises Slurm Infrastructure for seamless **hybrid** deployments
- Join existing Slurm-powered clusters to Nimble Federated or private Federation Pack to **drive utilization**
- **Same point-and-click user experience** and compatibility with Eviden HyperHub application marketplace
- **Modernize and extend** complex capabilities to new generations of HPC and AI domain experts
- **Secure**, policy-driven access and storage management
- Single sign-on (**SSO**)
- Seamlessly and securely **container-native** leveraging Singularity containers for workloads
- Batch, **3D** visualization (with GPU acceleration), and **Notebook** support
- **Nothing to install** on existing clusters



JARVICE™ XE HPC Platform Web Console

Latest advancement in HPC user interface technology

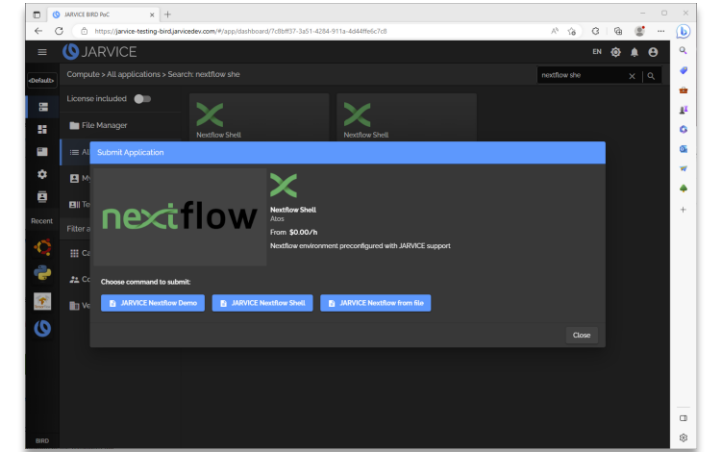
- Same great user experience, but...
 - More **modern**
 - More **ergonomic**
 - **Consistent** with the broader Eviden Software Suites
 - New baseline for **continued evolution**
 - Keycloak integration for **improved security** and SSO



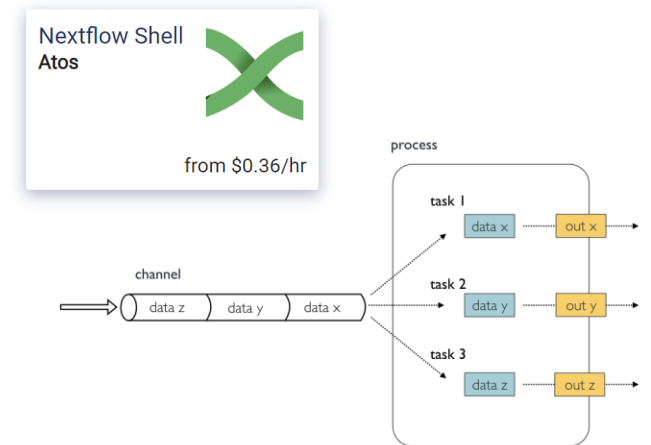
Support for Nextflow Pipelines

For bioinformatics, machine learning, and other scientific domains

- Out-of-the box support for pipeline-oriented scientific **workflows**
- **Eviden-contributed** open-source code for JARVICE™ XE API compatibility to Nextflow
- Easy **cloud enablement** for existing pipelines into Nimbix Elastic or Nimbix Federated
- Easy **cloud bursting** of existing pipelines from on-premises JARVICE™ XE clusters



Search For "Nextflow"



EVIDEN

05. Conclusion

As-a-service

- Any workload: HPC, AI & AC
- Any where: new endpoints (U.S. & Europe)
- Any way: private, public or hybrid cloud

Applications

CEPP: from fast-start program to applications optimization, workload efficiency improvement

AI-empowered Software

- Monitoring -> Capping -> Optimizing workload & application efficiency

HPC Hardware

- Scalable, flexible architecture
- Direct Liquid Cooling
- High-speed, low latency interconnect

Global system efficiency

Performance



European Sovereignty



Global Ecosystem

Greener HPC: The backbone of our strategy execution



BullSequana XH3000

EVIDEN

Questions





EVIDEN

Thank you!

Confidential information owned by Eviden SAS, to be used by the recipient only.
This document, or any part of it, may not be reproduced, copied, circulated
and/or distributed nor quoted without prior written approval from Eviden SAS.

© Eviden SAS