How Bull contributes to Meteo Projects?

Xavier VIGOUROUX

Head of « Center For Excellence in Parallel Programming » (CEPP)

27-10-2016

Bull
atos technologies

It was better before

How many humans did you need to optimize a code?

None, I just have to wait for the next generation

How many types of architectures were on the market place?

Much less than today (x86, GPGPU, many cores, ARM, ...)

What occurred?

Circ. 2006 End of dennard's law (constant power density)
End of frequency increase

but Feature size keeps shrinking (14nm today)
Rising of Multicore

➡ ITRS predicts 5nm in 2021

? What will you do with a 8 times larger die?
What will be the impact on your software?

💀 Worse! 2021 is the end of the CMOS road

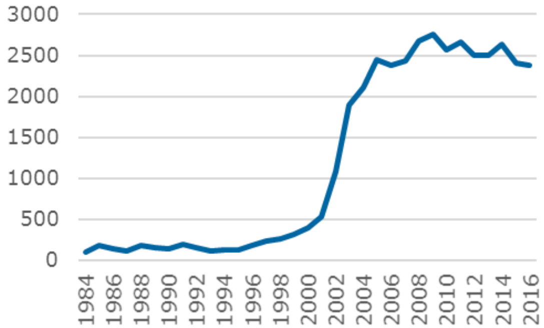

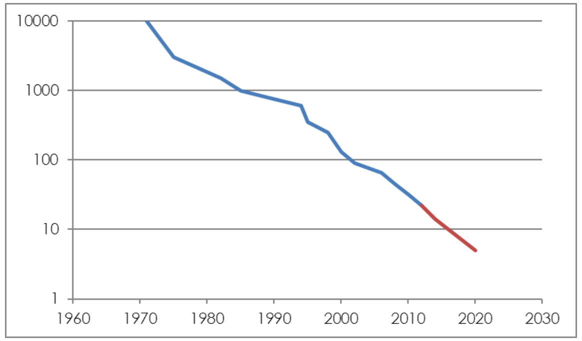

ref: http://www.itrs2.net/

Bull
atos technologies

ITRS: Shrink is not the only way

More Moore

Find other technologies to keep improving (post-CMOS, 3D…)
Quantum, cognitive, beyond von Neumann, in mem. comput….

More than Moore

Mixing with other technologies (sensors, MEMs…)

Process Improvement

More adaptive microelec process

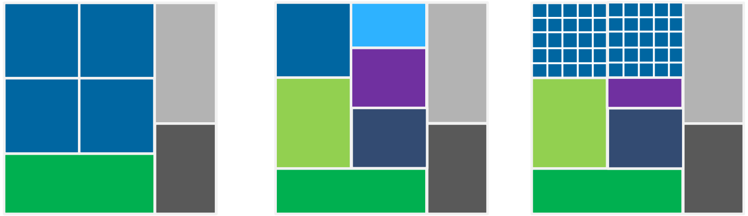

Time Variability and Time underdeterminism

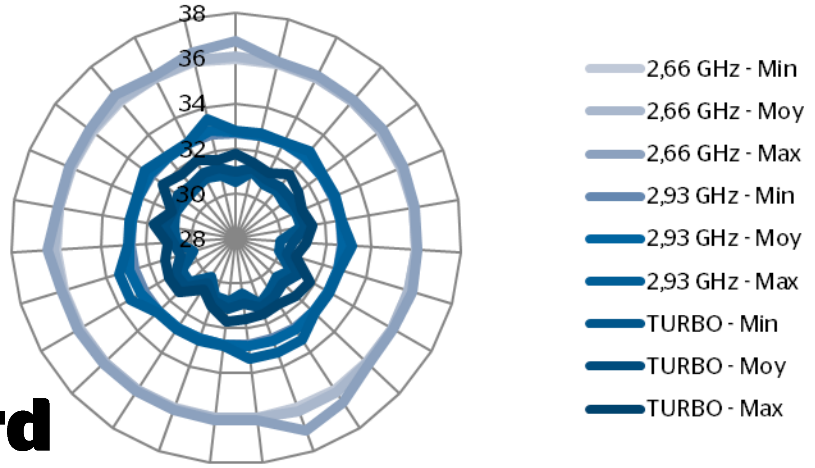

Variability exists for a long time

Turbomode revealed it

Underdeterminism (in time) is now standard

Underlying power optimisation is unknown from application

Applications must be less tightly coupled

The slowest part will pace the whole

Bull
atos technologies

So you are wondering....

How many humans will I need to optimize a code?

3 months for 3% multiplied by 20 codes multiplied by .... Arg!

How will I select the right architectures for my workflow?

x86, GPGPU, many cores, ARM, FPGA
Vector size impact
Flops/Byte decrease impact
Precision impact
Variability impact
Underlying power management impact
....

Our strategy

Understand the needs, constraints and wishes

Be part of projects, spend time with you

Design products

Answer the needs

Accompany users

Be beside users to overcome issues and get the performance you target

2

Understand the needs, constraints and whishes

Bull Role

Optimisations of “dwarfs”, KNL

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 671627

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 675191

Bull Role

Roadmap vision, optimisation, guidance

enes
EUROPEAN NETWORK FOR EARTH SYSTEM MODELLING

European Earth System Modelling Infrastructure Strategy

ENES is holding a community meeting in Reading (UK) on the 25th, 26th and 27th of October, supported by IS-ENES2 to discuss the future infrastructure strategy for earth system modelling. The meeting is expected to run from midday Tuesday, to mid-day Thursday.

In 2012, the European Network for Earth System Modelling (ENES) published an "Infrastructure Strategy for the European Earth System Modelling Community" (Mitchell et.al., 2012[1]) based on meetings held in 2010 and 2011. This strategy addressed the underlying needs for the delivery of the next decade of European research on seasonal to centennial climate prediction. It envisaged a drive towards convective scale global modelling, with improved initialisation and larger ensemble sizes. At the same time, attribution was expected to be addressed with enhanced paleo-climate modelling, and more attention would be focused on climate predictability on regional scales.

The key recommendations were to:

1. Provide a blend of high-performance computing facilities ranging from national machines to a world-class computing facility suitable for climate applications, which, given the workload anticipated, may well have to be dedicated to climate simulations.
2. Accelerate the preparation for exascale computing, e.g. by establishing closer links to PRACE and by developing new algorithms for massively parallel many-core computing.
3. Ensure data from climate simulations are easily available and well documented, especially for the climate impacts community.
4. Build a physical network connecting national archives with transfer capacities exceeding Tbits/sec.
5. Strengthen the European expertise in climate science and computing to enable the long term vision to be realized.

Five years on, ENES is convening a meeting to address a "mid-term" update of this strategy. Since 2010-2012 there has been much progress, often with support from IS-ENES2[2]. In particular, ideas that have been taken forward range from the establishment of a European Centre of Excellence in the Simulation of Weather and Climate[3] and the European engagement and leadership in the Earth System Grid Federation providing access to climate data, to the proposal of a European Programme on Extreme Computing and Climate[4]. However, it is timely to take stock of how much progress, and whether or not these are still the right objectives – both scientifically, and in terms of the infrastructure. An updated strategy will also be important to address the issue of how to sustain the European research infrastructure for climate modelling.

The outcome of this meeting should be both the input for an update of the infrastructure strategy, and community agreement on which new or existing initiatives should be prioritised to address the science requirements of decadal to centennial prediction (including model evaluation, process understanding, and perhaps whether the scope should be widened to include seasonal prediction).

A registration website will be made available shortly. In the mean time, interested participants should hold the dates.

Sylvie Joussaume (ENES Chair and IS-ENES2 coordinator)
Bryan Lawrence (Meeting Organiser)

---

[1] https://is.enes.org/archive/dissemination-documents-about-is-enes/dissemination-activities/ENES foresight.pdf
[2] http://is-enes.org
[3] https://www.esiwace.eu/
[4] https://ec.europa.eu/futurium/en/content/flagship-european-programme-extreme-computing-and-climate

is-enes
INFRASTRUCTURE FOR THE EUROPEAN NETWORK FOR EARTH SYSTEM MODELLING

Bull contribution

Exchanging about technology trends and impacts

3 Design Products

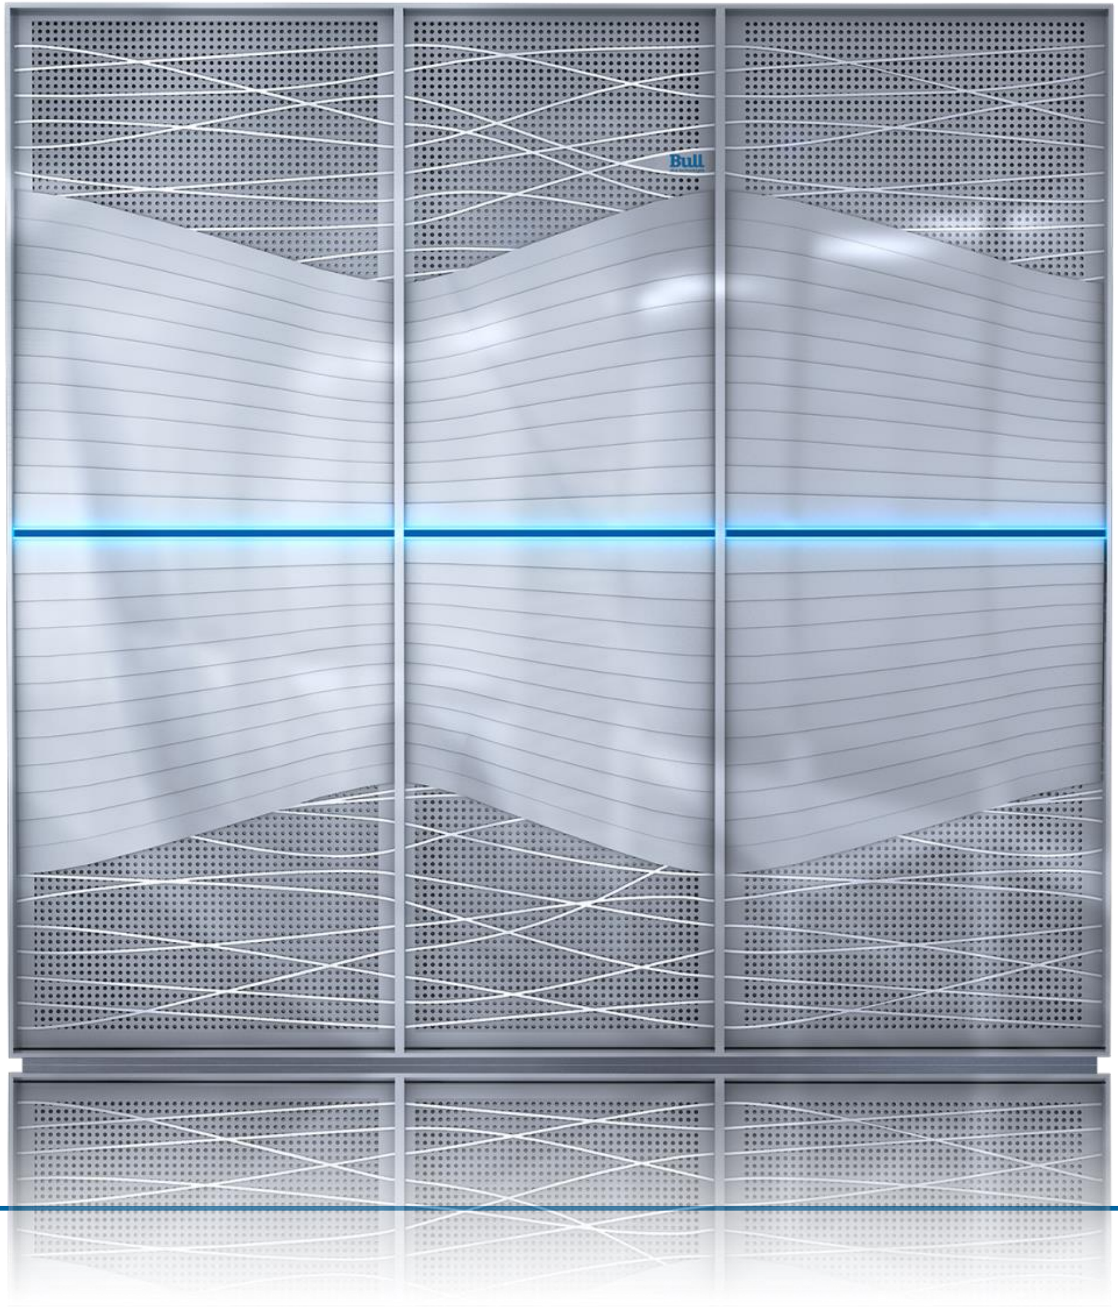

Bull
atos technologies

Front Side

2 x 24 blades

Back Side

2 x 24 blades

Bull
atos technologies

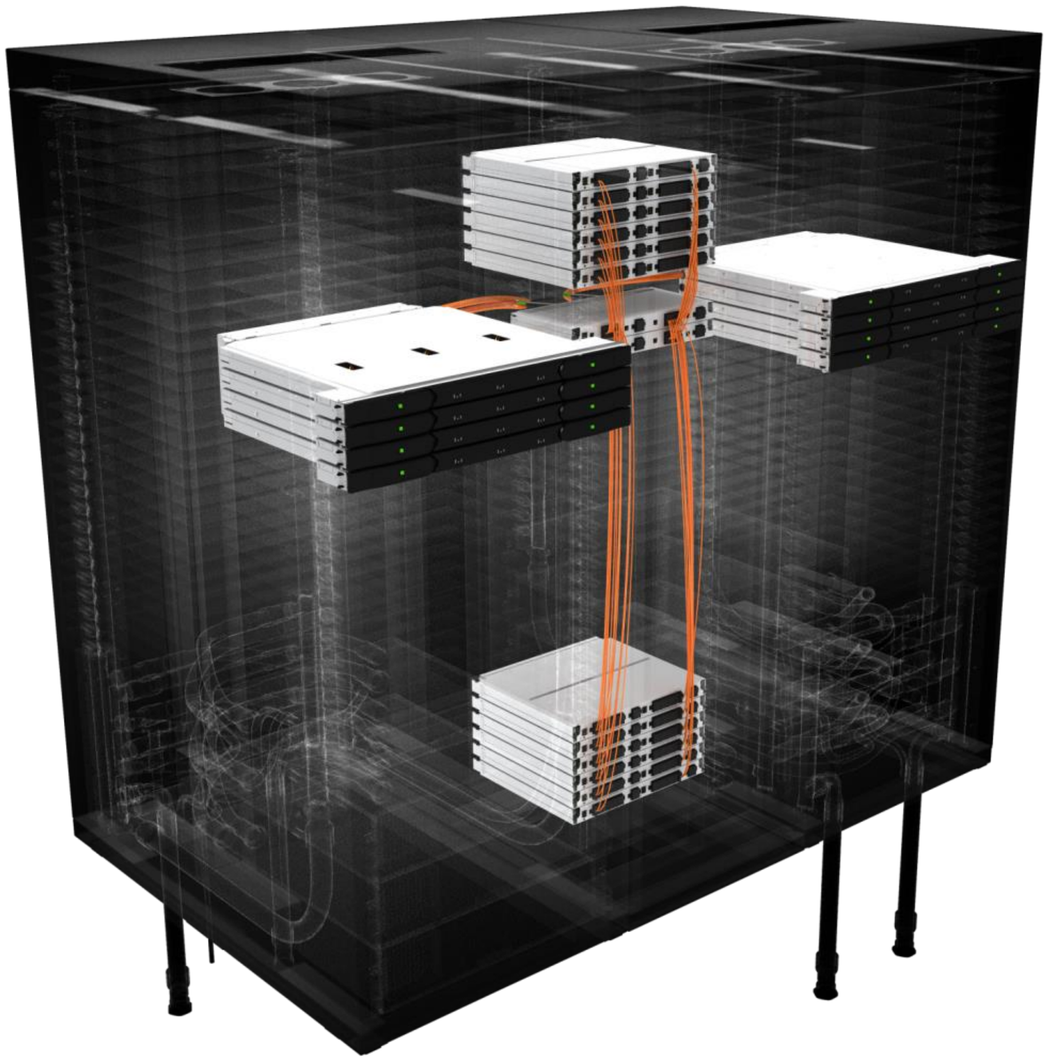

Broadwell

3 x 2 Intel® Xeon® E5-2600 v4 processors
3 x 1 Intel® C610 chipset

3 x 1 optional SATA SSD drive

3 x 8 DDR4 memory slots (max 256 GB with 32 GB DIMMs) with heat spreaders for cooling

BXI 1 or 2 ports mezzanine board
OR InfiniBand EDR 1 port mezzanine board

Cooling by direct contact with DLC coldplate

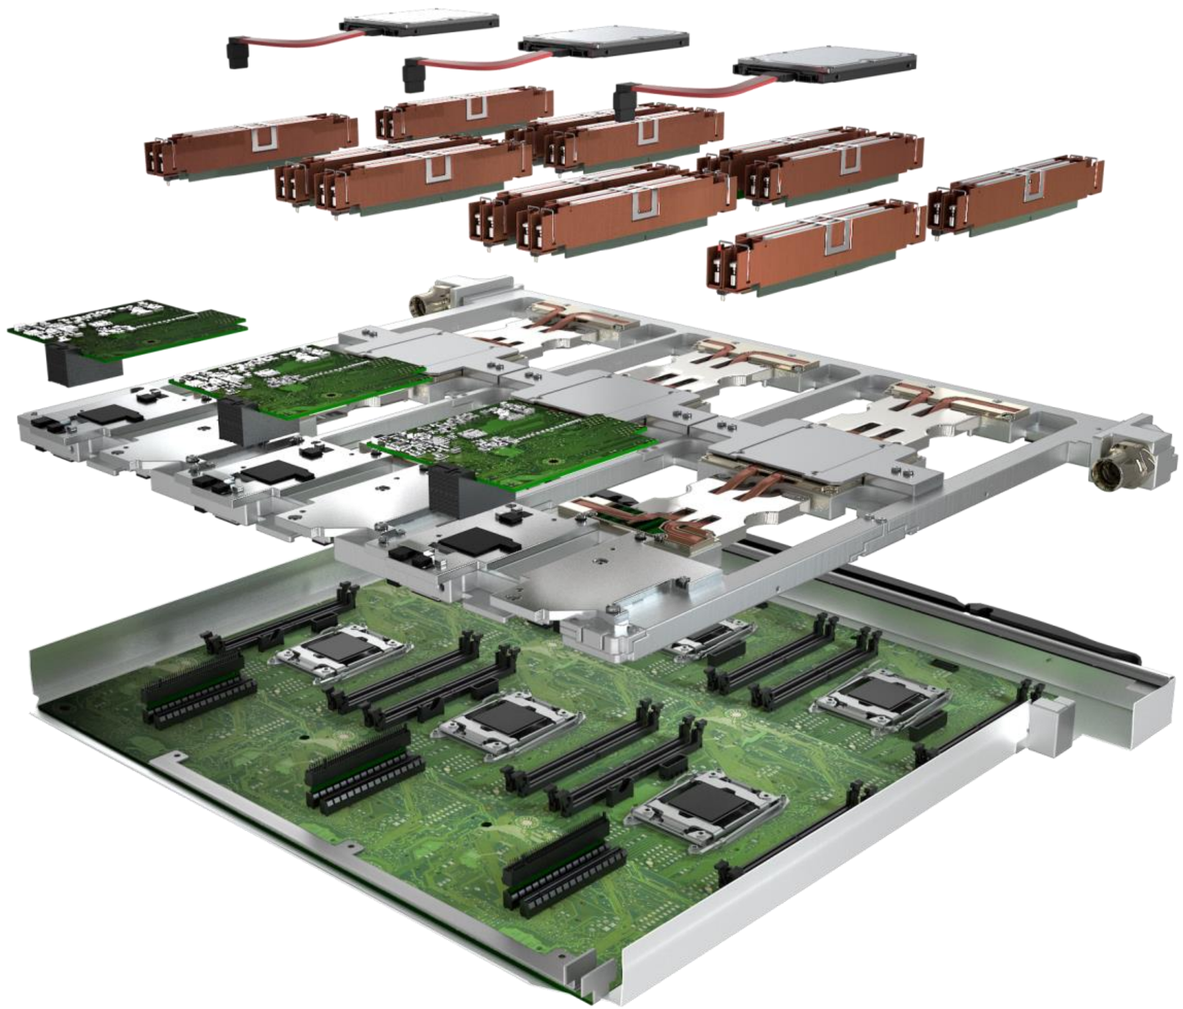

KNL

3 x 1 Intel® Xeon Phi™
(Knightslanding) processor
3 x 1 Intel® C610 chipset

3 x 1 optional SATA drive
3 x 1 optional PCIe SSD drive via PCIe switch
3 x 6 DDR4 memory slots (max 192 GB with 32 GB DIMMs) with heat spreaders for cooling
BXI 1 or 2 ports mezzanine board
OR InfiniBand EDR 1 port mezzanine board

Cooling by direct contact with DLC coldplate

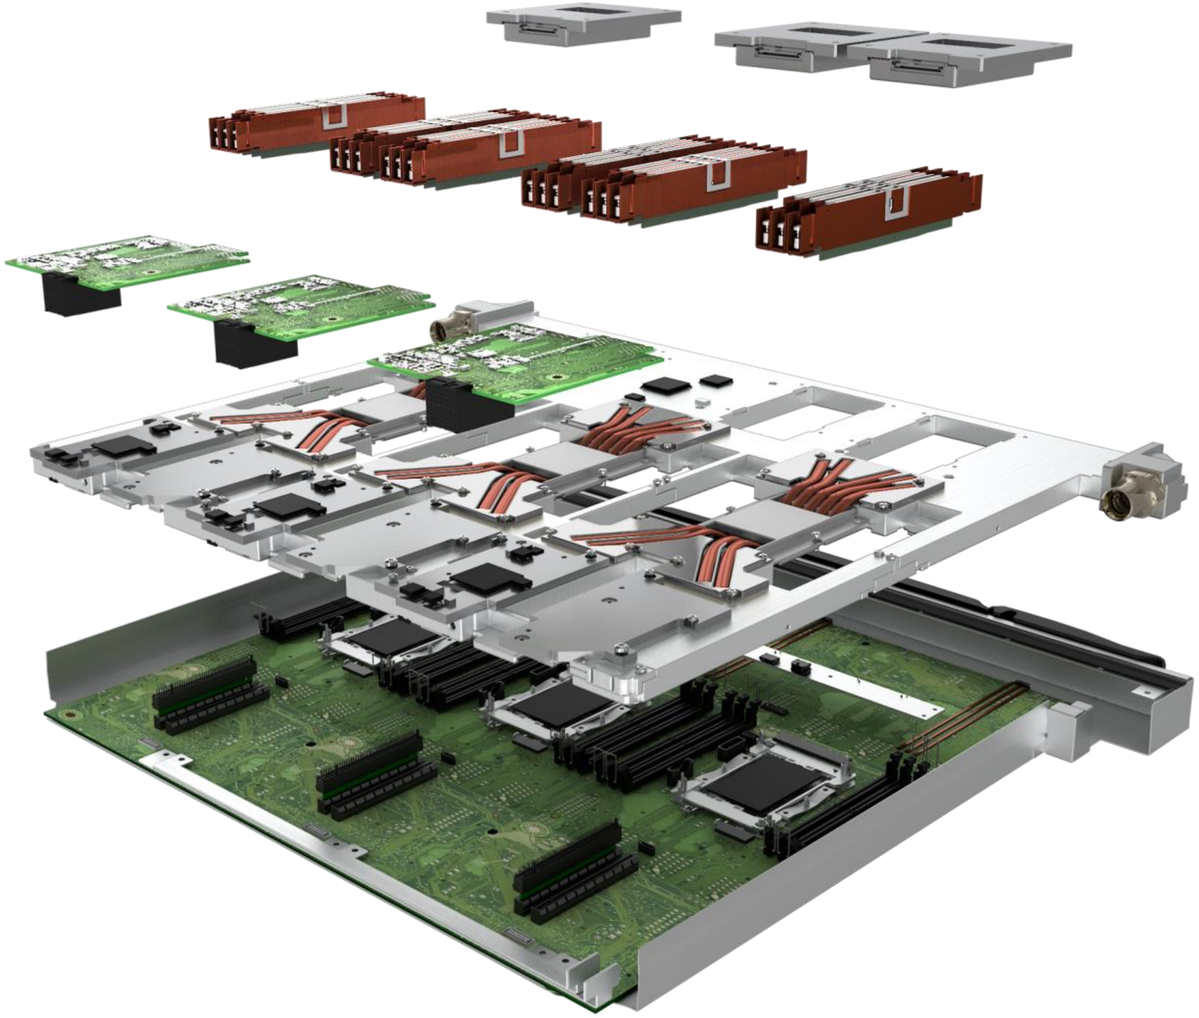

Pascal

4 Nvidia® Pascal GPUs
2 Intel® Xeon® E5-2600 v4 processors

1 optional SATA SSD drive

8 DDR4 memory slots (max 256 GB with 32 GB DIMMs) with heat spreaders for cooling

BXI 1 port mezzanine board
OR InfiniBand EDR 1 port mezzanine board

Cooling by direct contact with DLC coldplate

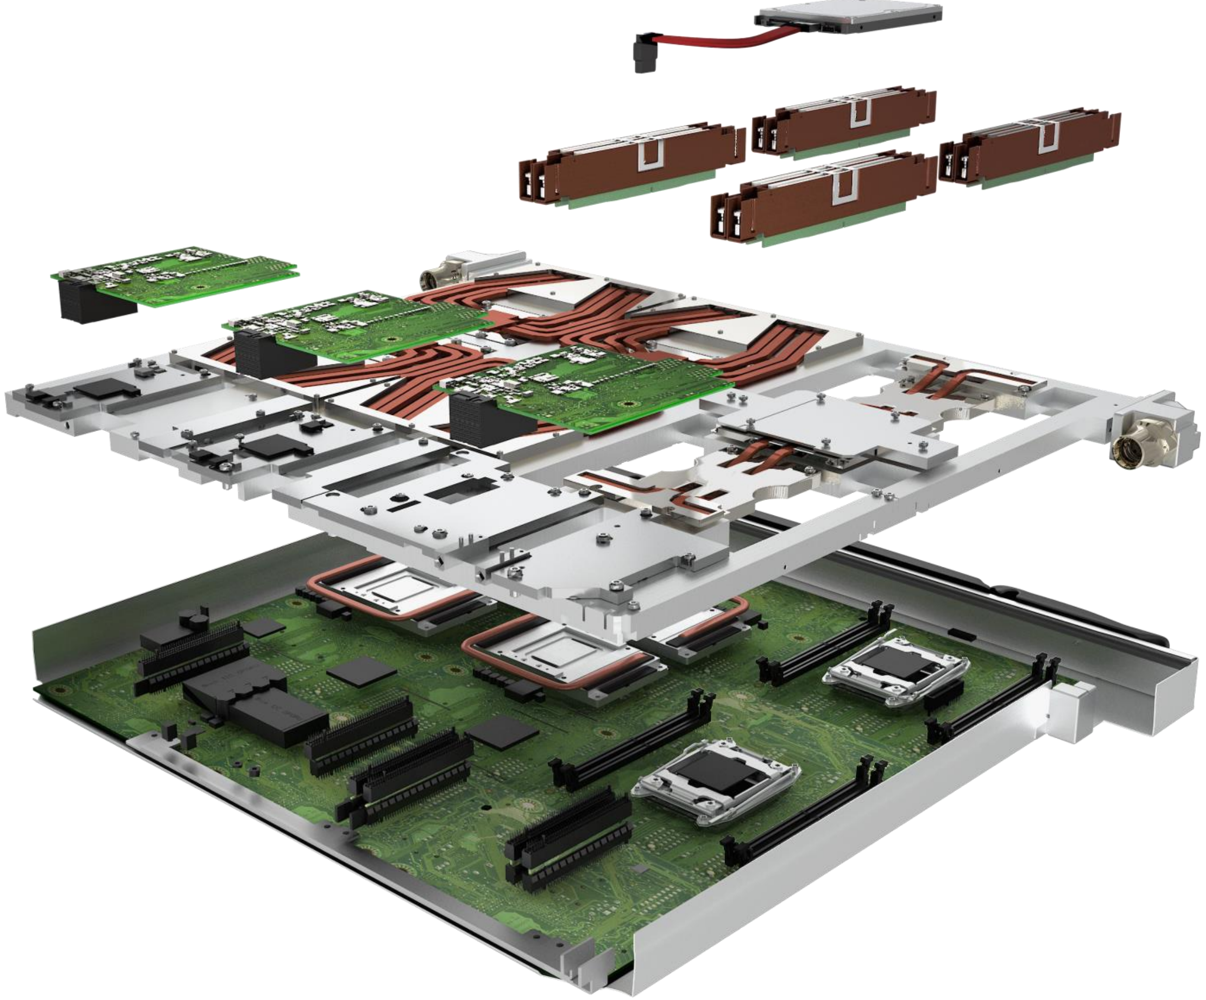

4 Accompany Users

Laboratório Nacional de Computação Científica
Cardiff University / Prifysgol Caerdydd
Royal Netherlands Meteorological Institute
University of Liverpool
Dassault Aviation
EZ Biometrics
cea
NVIDIA
Horizon 2020
calmip
CNI Centro Nacional de Inteligencia
ESCAPE
ILL Neutrons for Science Institut Laue-Langevin
SEISCOPE
ONERA The French Aerospace Lab
esiwace Centre of Excellence in Simulation of Weather and Climate in Europe
DKRZ Deutsches Klimarechenzentrum
GENCI Grand Equipement National de Calcul Intensif
ISTerre Institut des Sciences de la Terre
intel
Université de Reims Champagne-Ardenne
LOCEAN
INERIS
Bull atos technologies

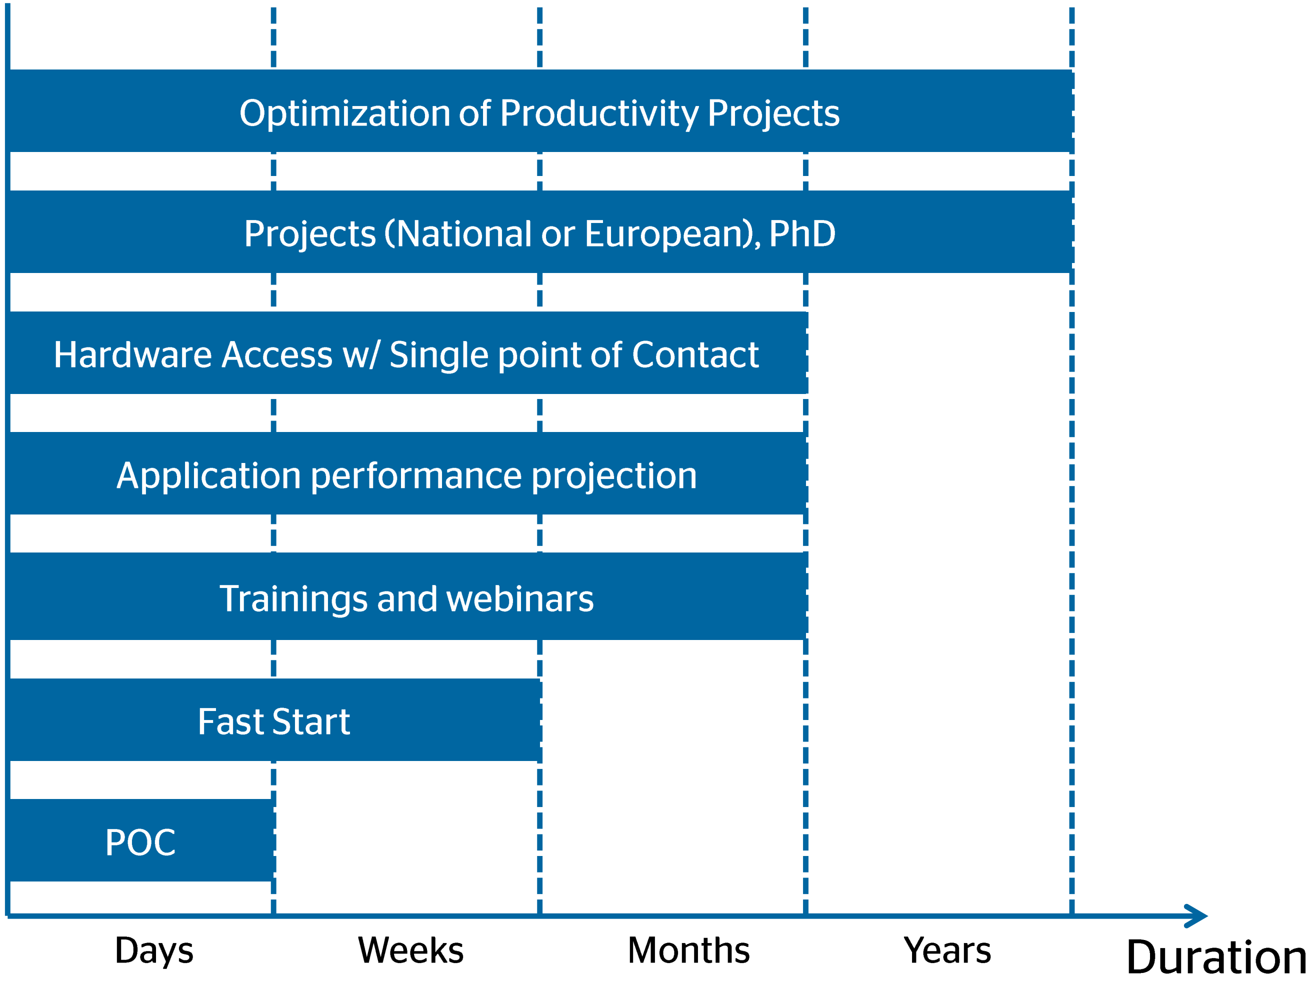

Bull
atos technologies

Oil & Gas

Architecture

Accelerators

DL, KNL, FPGA

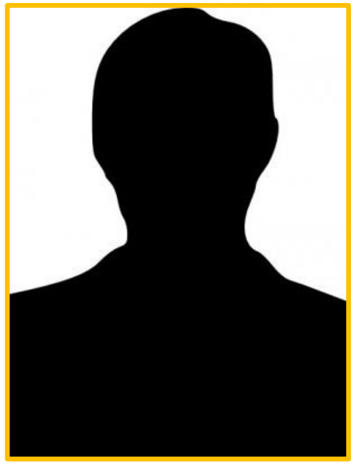

Finances

Oceano

Performance

ARM

Docking

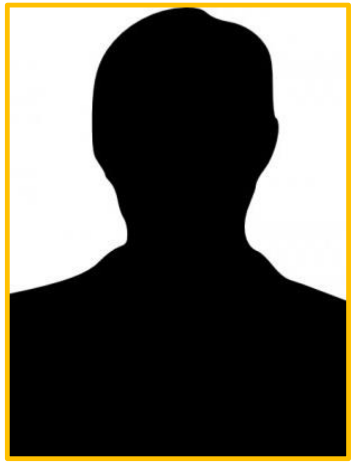

Fluid Dynamics

Bull
atos technologies

Deeply involved in the WF and Climate fields

AEMET, the Spanish meteorological agency

Improving weather forecasting: a direct impact on the safety of people and property

AEMET is Spain's meteorological agency. AEMET's activities include taking meteorological observations in Spain and archiving them, weather monitoring and forecasting, and scientific research in numerical weather prediction models

Business challenge

Extend AEMET's high performance computing resources to:

- improve weather forecasting (finer resolution),
- improve severe weather phenomena forecasts,
- expand AEMET's work in different areas such as climate change or wave prediction.

Solution

- Bull's patented Direct Liquid Cooling system, to reach a PUE < 1.1
- A bullx supercomputer with a peak performance of 168 Teraflops in its fin
  - 338 servers (1
  - R servers) i.e.
  - high-speed net
- 360 Terabytes sto

Benefit

- AEMET can tackle services

DKRZ, Deutsches Klimarechenzentrum

The German Weather Prophet

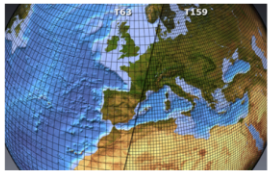

The German Climate Computing Centre (DKRZ) is a national facility that offers customized services to support climate researchers. Climate evolution is a question that arouses a great deal of controversy. To answer it, climate simulations are an essential tool. This involves replicating the climate system on a computer with the help of digital models. These climate simulations demand a huge compute capacity, and produce large quantities of data, which in turn demand the capacity to handle large data sets.

Business challenge

- Climate simulations demand a huge compute capacity
- Climate simulations produce large quantities of data
- Well designed techniques for data management and storage are an important prerequisite for climate research

Solution

- More than 60,000 computing cores in bullx B700 Direct Liquid Cooling blades distributed over 60 racks
- A supercomputer with a peak performance of 3 petaflops
- The corresponding Lu up to 45 petabytes - in the world

Benefit

- Improve climate fore
- Energy efficiency (PU
- Cooperation on appli
- Capacity to handle la

MARIN, the Maritime Research Institute Net

A new HPC facility to initiate a co-operation program with the

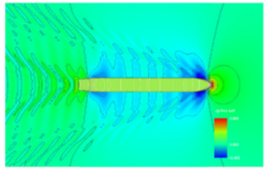

MARIN is one of the leading institutes in the world for hydrodynamic research and maritime technology. The offered services incorporate a unique combination of simulation, model testing basins, full-scale measurements and training programmes. MARIN provides services to the ship-building and offshore industry and governments. Customers include ship builders, fleet owners, naval architects, classification societies, oil and LNG companies and navies all over the world.

Business challenge

- Extend their HPC facility
- Expendable solution (2-phase pr
- Strengthen co-operation with the maritime industry
- Strong constraints on footprint and consumption

Solution

A 4080-core bullx system with a peak performance of almost 170 TFlops:

- Fat-tree interconnect based on InfiniBand FDR
- 204 bullx R424 E4 nodes (Haswell EP)
- Shared NetApp E2600 storage

Benefit

- The detailed benchmarking of MARIN's code by Bull experts convinced the customer
- A new co-operation program with the maritime industry was initiated by sharing with them the new HPC facility with dedicated maritime CFD code

KNMI

A bullx supercomputer for the Dutch weather forecast agency

Koninklijk Nederlands Meteorologisch Instituut
Ministerie van Verkeer en Waterstaat

KNMI is the Dutch national institute for weather, climate research and seismology. It disseminates weather information to the public at large, the government, aviation and the shipping industry in the interest of safety, the economy and a sustainable environment. To gain insight into long-term developments, KNMI conducts research on climate change. Making the knowledge, data and information on hand at KNMI accessible is one core activity.

Business challenge

- More computing power to
  - be able to issue early warnings in case of extreme weather
  - enhance capabilities for climate research

Solution

A system 40x more powerful than the previous one:

- 396 bullx B500 compute nodes, equipped with Intel® Xeon® processors, for a total of 4,752 cores
- 9.5 TB memory
- peak performance 58.2 Tflop/s

Benefit

"The hardware, combined with Bull's expert support, gives us confidence in our cooperation"

Bull
atos technologies

Many others

- Arome on ARM
- BRAMS performance for LNCC
- Full compilation of Harmonie on PGI, test on GPU
- Port of MESONH and DYNAMICO on KNL
- Several kernel of NEMO on KNC
- ...

Contact us : cepp@atos.net

Thanks

Be efficient at scale

If you ware plowing a field, which would you rather use? Two strong Bull or 1024 chickens?

Bull
atos technologies