News about HPC and Clouds @ Inria

Claude Kirchner
Advisor to the president

24/11/2014

Antoine Petit, new Inria Chairman and CEO

Antoine Petit, CEO - © *Inria / Photo C. Helsly*

By Decree of the President of France on 26 September 2014, Antoine Petit has been appointed as Chairman of the Board of Directors of Inria.

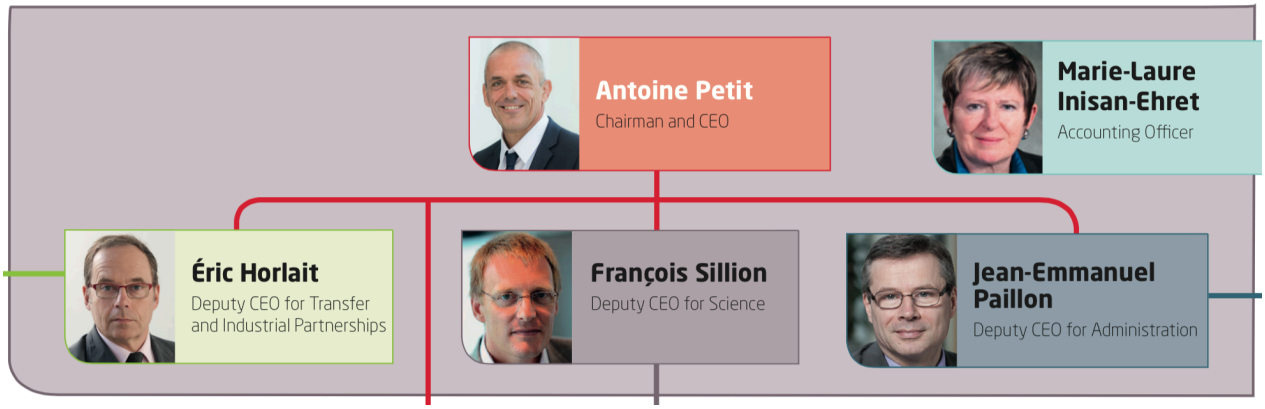

Inria Research Centres

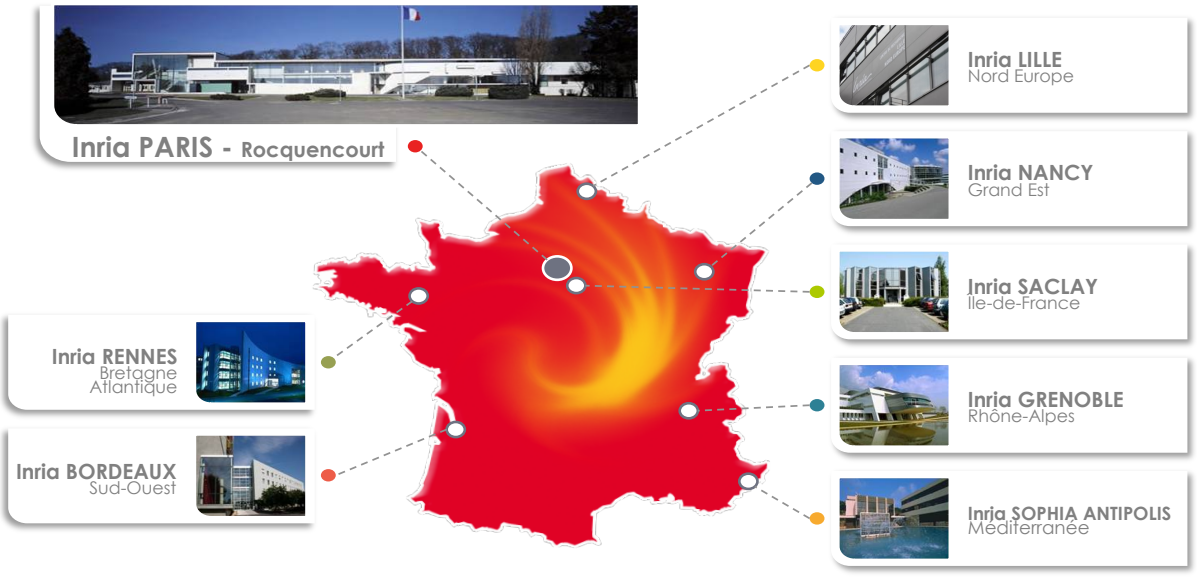

Inria Strategic Plan : 5 years strategy

1- Sciences helpful for humains, society and knowledge

- **Humans as such**: Health and Well Being
- **Humans and their environments**: from the individuals to the society, from the habitat to the planet
- **Humans and knowledge:** emergence of knowledge, scientific mediation and education

2- Key scientific topics at the heart of our sciences

- **Computing the Future**: models, software and digital systems
- **Mastering complexity:** data, networks and flows
- **Interacting with the real and virtual worlds**: interactions, uses and machine learning

Scientific challenges at the heart of our sciences

Computing the Future: models, software and digital systems

- Designing multiscale models integrating uncertainties
- Building very large digital systems, possibly embedded, and systems of systems
- Programming very large software under reliability, safety and security constraints

Mastering complexity: data networks and flows

- Processing data deluges into trustable knowledge libraries
- Generalized, safe cyber-communication, preserving privacy

Interacting with the real and virtual worlds: interactions, uses and machine learning

- Unsupervised machine learning
- Interaction between humans and their digital environments

Project-teams involved in HPC

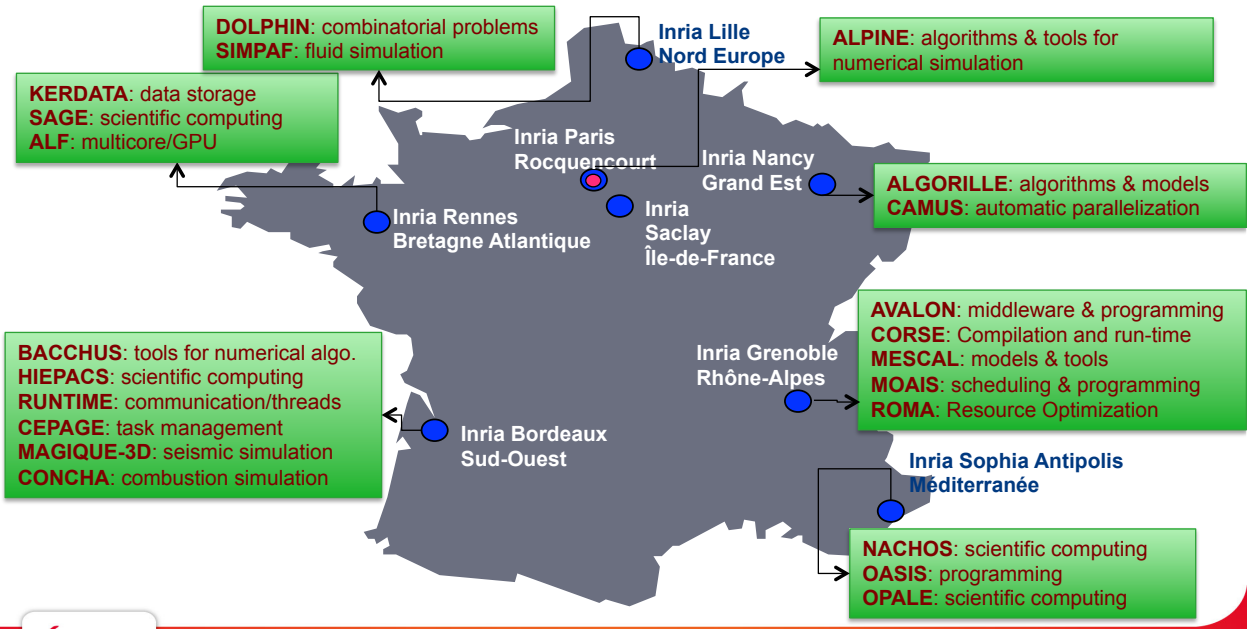

Project-teams involved in Clouds

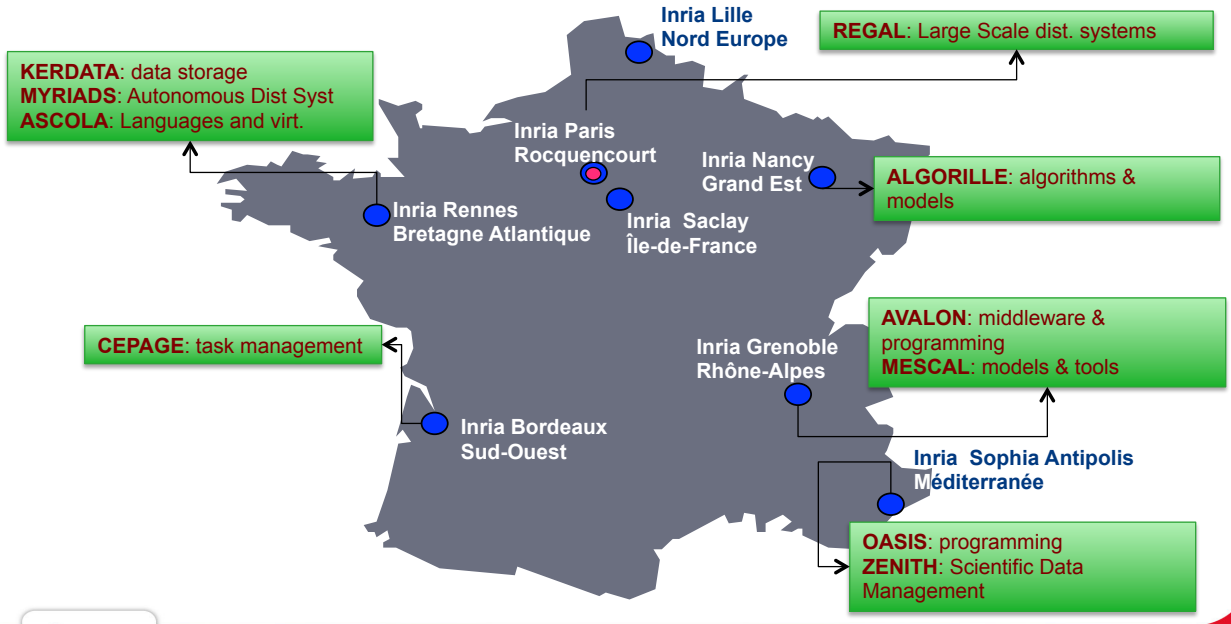

Initiatives to support HPC/Clouds strategy within Inria

Inria Project Labs (IPL)

- Enable to launch ambitious projects of the strategic plan
- Promote an interdisciplinary approach
- Mobilizing expertise of Inria researchers around key challenges
- 4 years duration
- ~1.5 Meuros over 4 years
- ➔ **CS2@Exa**
- ➔ **Hemera**
- ➔ **Fusion, Multicore**

Other (GIS, PIA/ANR)

- ➔ **GRID5000**
- ➔ **ELCI**

Access to Machines

- Curie (GENCI / CEA)
- Fermi, … (PRACE)
- GRID5000
- Blue Waters
- Bull prototypes
- Plafrim (Bordeaux)
- …

IPL: C2S@Exa
Computers and Computational Sciences at Exascale

Contact: Stephane.Lanteri@inria.fr

- **Numerical simulation for high performance massively parallel architectures**
  - **Numerical linear algebra**
  - **Numerical schemes for PDE models**
  - **Optimization of performance of numerical solvers**
  - **Programming models**
  - **Resilience for exascale computing**

- **Multidisciplinary approach in between applied mathematics and computer science: From generic building-blocks to large-scale applications**
  - **Nuclear energy production** (fusion): MHD computation with the JOREK simulation software (with CEA)
  - **Radioactive waste management scenario**: Environmental applications (with ANDRA)

GRID'5000

- **Testbed for research on distributed systems**
  - Born from the observation that we need a better and larger testbed
  - High Performance Computing, Grids, Peer-to-peer systems, Cloud computing
  - A complete access to the nodes' hardware in an exclusive mode (from one node to the whole infrastructure): Hardware as a service
  - RIaaS : Real Infrastructure as a Service ! ?
- **History, a community effort**
  - 2003: Project started (ACI GRID)
  - 2005: Opened to users
- **Funding**
  - Inria, CNRS, and many local entities (regions, universities)
- **One rule:** only for research on distributed systems
  - → no production usage
  - Free nodes during daytime to prepare experiments
  - Large-scale experiments during nights and week-ends (no long jobs)

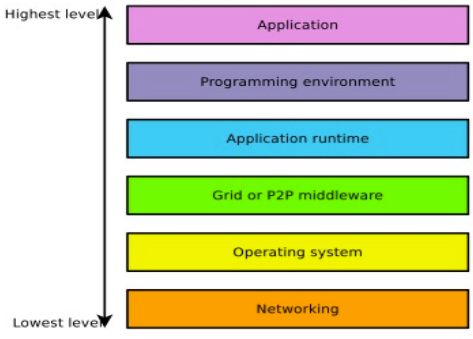

Current Status (Sept. 2014 data)

- 10 sites (1 outside France)
- Dedicated 10 Gbps backbone provided by Renater (French NREN)
- 24 clusters
- 1006 nodes
- 8014 cores
- **Diverse technologies**
  - Intel (65%), AMD (35%)
  - CPUs from one to 12 cores
  - Ethernet 1G, 10G,
  - Infiniband {S, D, Q}DR
  - Two GPU clusters
  - 2 Xeon Phi
  - 2 data clusters (3-5 disks/node)
- More than **500 users** per year
- Hardware renewed regularly

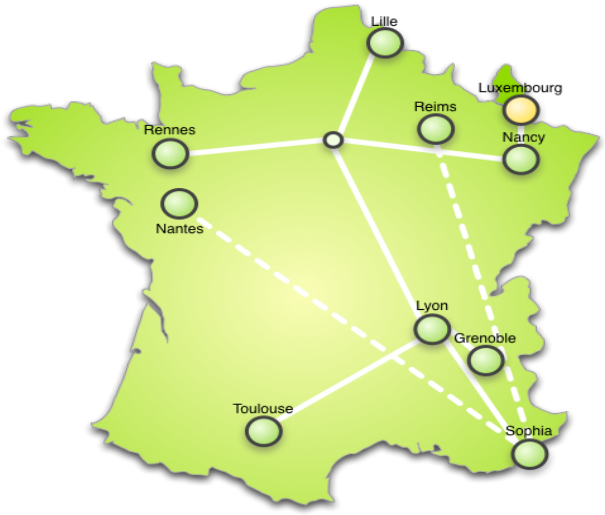

Some recent experiments over Grid'5000

- **Energy monitoring and management**
  - Evaluation of Green Strategies for Energy-Aware Framework in Large Scale Distributed Systems
  - Evaluation of different watt-meters
  - Estimation of energy consumption with or without application expertise
- **Cloud Computing and virtualization**
  - Sky computing between France and US using Hadoop
  - Virtual machines deployment and migration (up to 10240 VMs)
  - Experiments using major CloudKits (Nimbus, OpenStack) and VM stacks (Xen, kvm)
- **High Performance Computing**
  - Replay of the Curie computer traces for resource management systems over an emulated environment
  - Comparison of component based approaches and MPI/threads applications
- **Big data management**
  - Optimization of MapReduce frameworks with high performance data management systems
  - High performance data movements over multicore machines validated with a climate simulation application

ELCI: software environment for computation-intensive applications

- Develop numerical simulation for HPC:
  - New generation of SW stack: supercomputer control, numerical solvers, prog. & exec. environment.
  - Validation: better scalability, resilience, security, modularity, abstraction and interactivity of applications.
- Consortium: **Bull**, CEA, Inria, SAFRAN, CERFACS, CORIA, CENAERO, ONERA, Univ. of Versailles, Kitware, AlgoTech
- 36 months (started on 09/2014)
- 150 PY

Inria@ELCI

1. Software environment

1. Intra-node optimization (extending Linux Kernel)
2. Cluster management (batch scheduler)
3. Runtime systems (process placement)
4. Energy efficiency (multi-criteria algorithms)
5. Resilience (replication, impact on energy, etc.)
6. Software integration

2. Solver and numerical methods

1. Mesh tools (mesh refinement, mesh allocation)

3. Programming methods, tools, code optimization

1. tools and prog. models (StarPU/Xkaapi with MPC)
2. High level models for productivity (components)

Inria@ELCI

Myriads (Energy)

Roma (Resilience, energy)
Avalon (components, energy)

Runtime (prog. mod., batch sched, runtime systems)
RealOpt (Resilience)
Bacchus (Meshes)

Moais (prog. models)

Vision

A real opportunity to get some of the very best international actors of HPC to work together: Make the JLESC a success at least as excellent as the JLPC!

- Continue and develop the strong collaboration with NCSA and Argonne
- Develop strong existing and new relationships with BSC, Jülich and Riken
- Develop connections between JLESC and other HPC / Cloud Inria involvements:
  - e.g. Chameleon / GRID5000
  - CS2@Exa
  - …

Have a fruitful workshop!

More technical informations

- **Hélène Kirchner,** Director of International Relations
- **Yves Robert,** JLESC Executive Director for Inria
- **Frédéric Desprez**, Deputy Scientific Director
  - Networks, Systems and Services and Distributed Computing
  - Frederic.Desprez@inria.fr
- **Jean Roman**, Deputy Scientific Director
  - Applied Mathematics, Computing, and Simulation
  - Jean.Roman@inria.fr

Inria strategy in HPC/Clouds

Inria is among the HPC leaders in Europe

- Long history of researches around distributed systems, HPC, Grids, virtualized environments and Big Data
- Multidisciplinary research culture
- Building and using exploration tools (eg massively parallel machines since 1987, large scale testbeds such as Grid'5000)

National initiatives

- Collaboration with Bull on Supercomputer design ➔ ELCI ➔ LECO
- Strategic Partnership with EDF, TOTAL, EADS-ASTRIUM on numerical simulation
- Joint laboratory with CERFACS on robust scalable sparse linear solvers
- Collaboration with CEA on key system software (Kadeploy) for Supercomputers
- Participation to French Strategic Committees on HPC: ORAP, TER@TEC
- Shareholder of GENCI and governance of Maison de la Simulation with CEA and CNRS

INRIA strategy in HPC/Clouds

- **European**
  - PRACE-1IP/2-IP/3-IP (within GENCI)
  - EESI & EESI2 (Exascale initiatives)
  - ETP4HPC
  - FP7 ICT, Challenge 1: Pervasive and Trusted Network and Service Infrastructures
  - XtreemOS , Contrail, BonFire, Fed4Fire
  - H2020

- **International**
  - Joint laboratory on petascale computing and now JLESC
  - Inria@SiliconValley: Stanford and Berkeley Universities
  - HOSCAR with Brazil, CNPq
  - Standardization (DTMF, OGF/OCCI)