



HPC and Clouds @ Inria

F. Desprez
Frederic.Desprez@inria.fr

Jun. 12, 2013

INRIA strategy in HPC/Clouds

INRIA is among the HPC leaders in Europe

- Long history of researches around distributed systems, HPC, Grids, and now Clouds
- Several activities virtualized environments
- Culture of multidisciplinary research
- Culture of exploration tools (owner of massively parallel machines since 1987, large scale testbeds such as Grid'5000)

• National initiatives

- Collaboration with Bull on Supercomputer design
- Strategic Partnership with EDF on simulation
- Joint laboratory with CERFACS
- Collaboration with CEA on key system software (Kadeploy) for Supercomputers
- French Strategic Committee on HPC: ORAP, TER@TEC
- HPC-PME (with GENCI and OSEO)
- ...



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

International


INRIA-Illinois joint laboratory

G8 ECS (Towards Exascale Climate simulation)

USA with Inria@SiliconValley: Stanford and Berkeley Universities

Associated teams with key HPC players

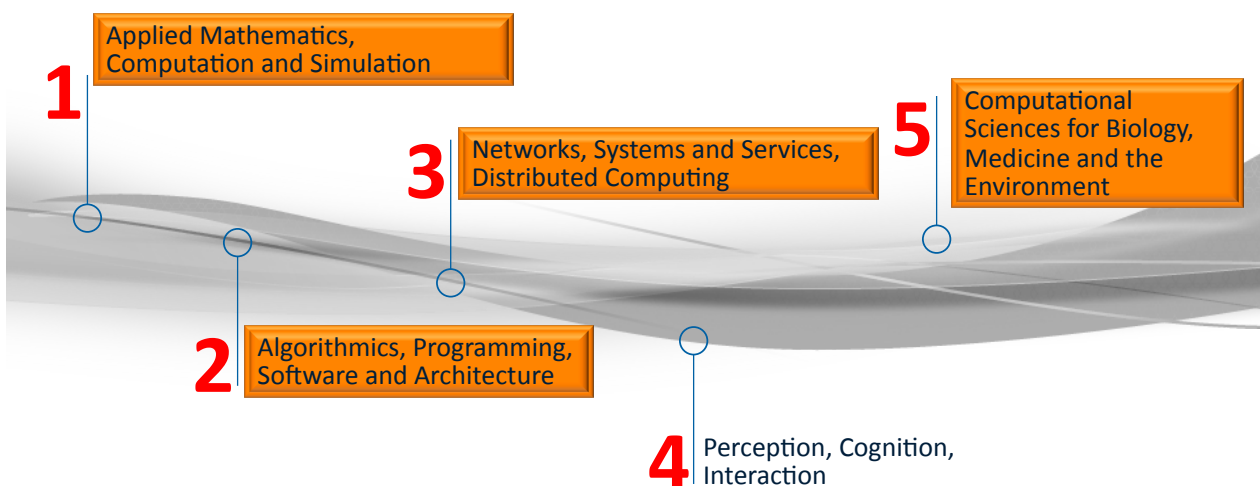
Standardization (DTMF, OGF/OCCI)



HPC/Clouds @ Inria - F. Desprez

June 12, 2013 - 3

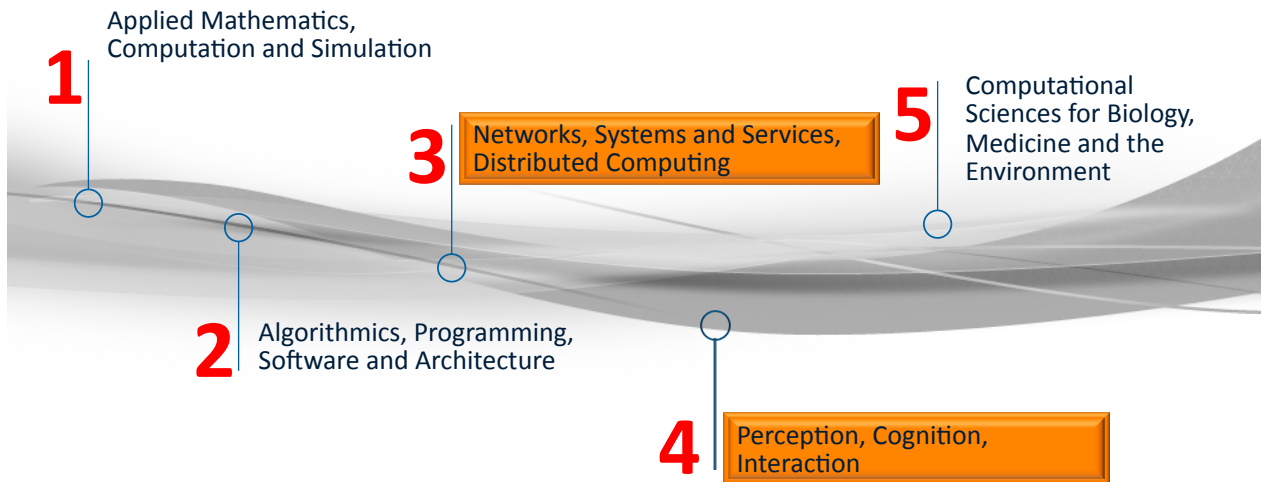
1. HPC : where within Inria ?



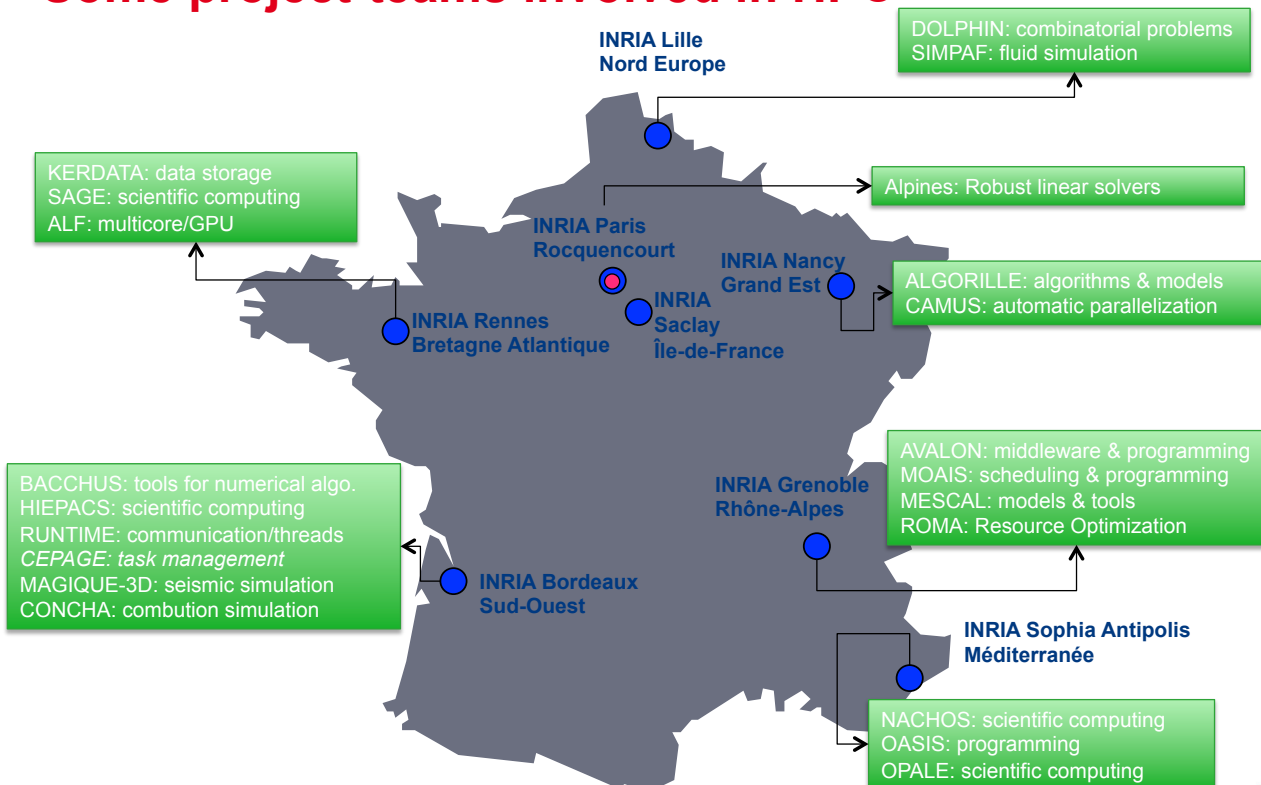
HPC/Clouds @ Inria - F. Desprez

June 12, 2013 - 4

2. Clouds: where within Inria ?



Some project-teams involved in HPC

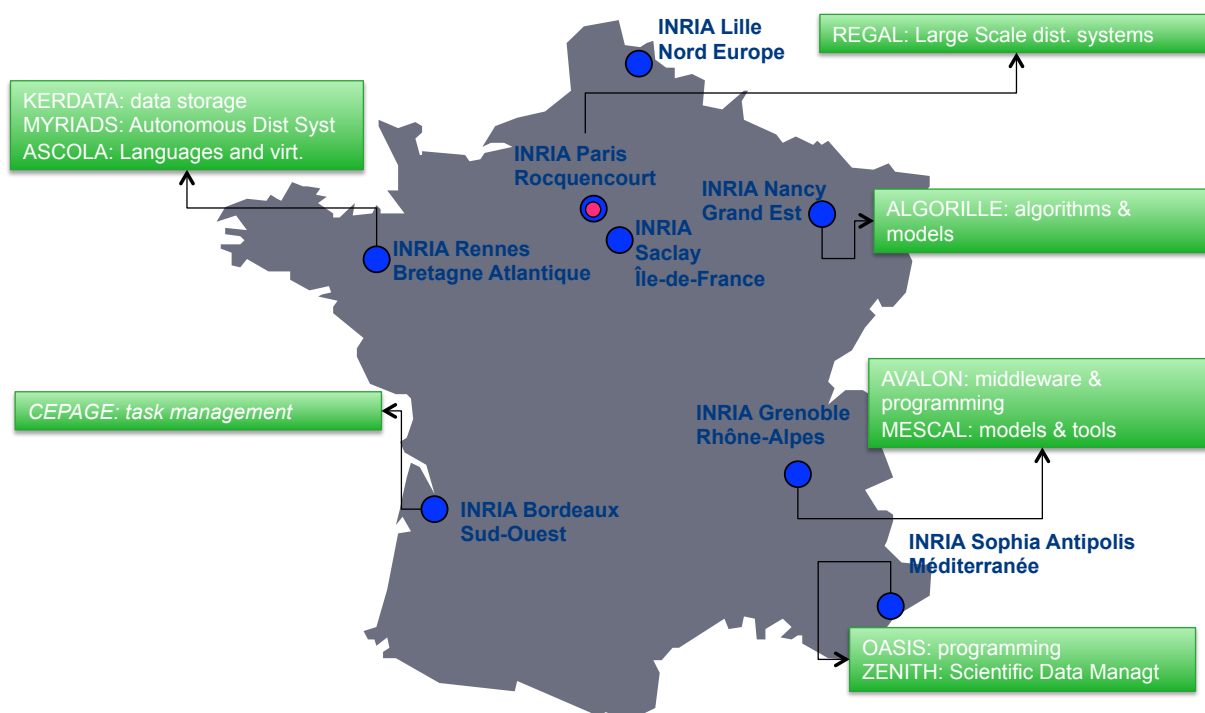


Current activities around HPC@ INRIA

- **Resource allocation and management**
Avalon, Cepage, Roma
- **GPU/Manycore**
Algorille, Runtime, Mescal
- **Parallel Solvers**
Alpines, Hiepacks, Roma
- **Simulation**
Algorille, Avalon, Mescal
- **Fault-tolerance/Resilience**
Cepage, Hiepacks, Roma, Mescal, Algorille
- **Programming models**
Avalon, Moais, Runtime
- **Visualization**
Hiepacks, Moais
- **Data Management**
Avalon, Mescal, Algorille
Kerdata
- **Desktop Grid**
Avalon, Mescal



Some project-teams involved in Clouds/Big Data



Current activities around Clouds@ INRIA

- **Resource allocation and management**
Avalon, Cepage, Regal
- **Energy management**
Avalon, Myriads, Ascola
- **VM management**
Ascola, Myriads
- **Security**
Avalon
- **Adaptative Dist. Applications**
ADAM
- **OS**
Ascola, Myriads
- **Simulation**
Algorille, Ascola, Avalon, Mescal
- **Programming models**
Avalon, OASIS
- **Unconventional paradigms**
Myriads
- **Model Driven Engineering**
Triskell
- **MapReduce paradigm**
Avalon, Kerdata
- **PaaS**
Avalon, Myriads, ADAM
- **IaaS**
Ascola, Myriads
- **Scalable data analysis**
Zenith
- **Big Data management**
KerData, Zenith



Initiatives to support HPC/Clouds within Inria

Why dedicated initiatives to support HPC/Clouds ?

- Project-teams are geographically dispersed
- Project-teams belong to different domains
 - Researchers from scientific computing need access to the latest research results related to tools, libraries, runtime systems, ...
 - Researchers from “computer science” need access to applications to test their ideas as well as to find new ideas !

Concept of “Inria Large Scale Initiatives” (now Inria Project Labs)

- Enable the launch of ambitious projects linked with the strategic plan
- Promote an interdisciplinary approach
- Mobilizing expertise of Inria researchers around key challenges



C2S@Exa Large-Scale Initiative

Computers and Computational Sciences at Exascale

Contact: Stephane.Lanteri@inria.fr

- Development of numerical simulation tools taking full benefits of processing capabilities of emerging high performance massively parallel architectures
- Establishment of a continuum of skills in the applied mathematics and computer science fields for a multidisciplinary approach
- Activities and contributions are organized along a three-level structure from generic building-blocks to large-scale applications:
 - Nuclear energy production (fusion) from CEA
 - Environmental applications from ANDRA



C2S@Exa thematic areas

- **Numerical linear algebra**
 - Core numerical kernels, sparse direct solvers, preconditioned iterative solvers, continuous solvers
- **Numerical schemes for PDE models**
 - Efficient numerical schemes to exploit massively parallel systems
- **Optimization of performance of numerical solvers**
 - Resource management and scheduling strategies, runtime systems, static and dynamic processing of numerical data sets
- **Programming models**
 - Component models for code coupling
 - High level parallel programming models to abstract the architecture
- **Resilience for exascale computing**
 - Energy effective fault tolerant protocols, algorithm-based fault tolerance, performance execution models for fault-tolerant applications, resilience for sparse linear algebra.



Hemera Large-Scale initiative

Scientific Challenges using Grid' 5000

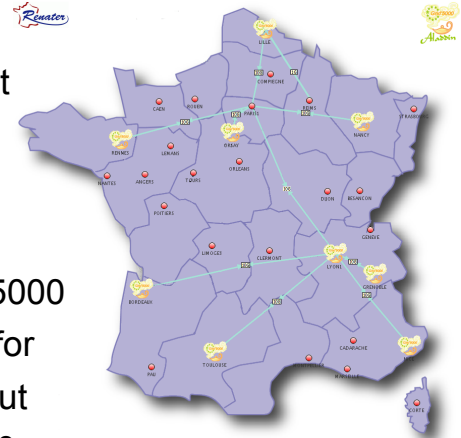
Contact: Christian.Perez@inria.fr

Grid'5000 is a scientific instrument designed to support experiment-driven research in all areas of computer science related to parallel, large-scale or distributed computing and networking.

- 9 main sites connected through 10G Ethernet
- 20 clusters, 1500 nodes totalling 7244 cores

Hemera goals

- Animate the scientific community around Grid'5000
- Demonstrate ambitious up-scaling techniques for large scale distributed computing by carrying out several dimensioning experiments on Grid'5000



HPC/Clouds @ Inria - F. Desprez

June 12, 2013 - 13

Hemera challenges

Network

- Traffic Awareness

System

- Energy Profiling of Large Scale Applications
- Robustness of Large Systems in Presence of High Churn
- Orchestrating Experiments on the gLite Production Grid Middleware
- Large Scale Virtual Machine Deployment & Management

Programming Paradigm

- Large Scale Computing for Combinatorial Optimization Problems
- Scalable Distributed Processing Using the MapReduce Paradigm

Application Domain Specific

- Multi-parametric Intensive Stochastic Simulations for Hydrogeology
- Thinking GRID for Electromagnetic Simulation of Oversized Structures




HPC/Clouds @ Inria - F. Desprez

June 12, 2013 - 14

ADT SimGrid: Simulator of Distributed Applications

Scientific instrument for the study of large scale distributed computing



Main Features

- Versatile: Grid, P2P, HPC, Volunteer Computing, Clouds, . . .
- Valid: Accuracy limits studied and pushed further for years
- Scalable: 3M chord nodes; 1000× faster than other (despite precise models), thousands of VMs (and Amazon Cloud)
- Usable: Tooling (generators, runner, vizu); Open-source, Portable, . . .