
Héméra
Inria Project Lab
July 2010 – June 2014

Final Evaluation
Paris, December 17th 2014

Christian Perez
AVALON
INRIA, France



Agenda

10:00-10:10. Bienvenue et tour de table

10:10-10:35. Présentation et Bilan des 4 ans d'Héméra, Christian Perez (Avalon)

10:35-10:55. Improving Experimentations on Grid'5000, Laurent Poulloux (Avalon)

10:55-11:15. Hemera from Grid'5000 Technical Team Point of View, Simon Delamare (Avalon/G5K)

11:15-11:35. Outils pour la recherche reproductible sur Grid'5000, Lucas Nussbaum (AlGorille)

11:35-11:55. Modeling Large Scale Systems and Validating their Simulators, Martin Quinson (AlGorille)

11:55-13:45. *Déjeuner à 12h au restaurant Assis au neuf 1 rue Godefroy, Paris*

13:45-14:05. Energie, Laurent Lefevre (Avalon)

14:05-14:25. Large scale management of VM - de Flaucher à VM5K, Adrien Lebre (Ascola)

14:25-14:45. Data / MapReduce, Shadi Ibrahim (KerData)

14:45-15:05. Data / Analytics, Reza Akbarinia (Zenith)

15:05-15:25. Challenges in solving large scale combinatorial optimization problems, Bilel Derbel (Dolphin)

15:25-16:00. *Café (et recalage horaire) !*

16:00-17:00. Discussion



Outline of the talk

- Motivations
 - Status in 2010
- Overview of Héméra
 - Scientific Challenges of Héméra
 - Working groups of Héméra
- Héméra Resources
- Conclusion
 - Status in 2014


Motivations – Science

■ Scientific issues

- ❑ Large scale, volatile, complex systems
 - Performance, fault tolerance, scalability, data storage, programming models, algorithms, resource management, etc.
 - Methodological challenges

■ Positioning

- ❑ Mathematics
- ❑ Simulation
- ❑ Emulation
- ❑ *Experimental testbed (Grid'5000)*
- ❑ Production environment



Motivations – Status in 2010

- Many teams using Grid'5000
- No sharing of knowledge of how to use it
- Many local tricks developed
- Huge waste of PhD time for making experiments
- Few exchanges between Grid'5000 and scientists
- Not so many large scale experiments

Overview of Hemera

- (Original) Goals (2010-2012)
 - Demonstrate **ambitious up-scaling** techniques for large scale distributed computing by carrying out **several dimensioning experiments** on the Grid'5000 infrastructure
 - **Animate** the scientific community around Grid'5000
 - **Enlarge** the Grid'5000 community by helping newcomers to make use of Grid'5000
- Open to everyone (not only INRIA)



Overview of Hemera

- (Revised) Goals (2013-2014)
 - Demonstrate **ambitious up-scaling** techniques for large scale distributed computing by carrying out **several dimensioning experiments** on the Grid'5000 infrastructure
 - Improve experimentation tools
 - Support a limited set of research fields
 - Experiments, virtualization, energy, big data
 - **Animate** the scientific community around Grid'5000
 - **Enlarge** the Grid'5000 community by helping newcomers to make use of Grid'5000
- Focus on a subpart of the participants (not only INRIA)

Hemera: Full Participant List

1. *ACADIE - Assistance à la Certification d'Applications Distribuées et Embarquées*
2. *ALGORILLE - Algorithms for the Grid*
3. *APO - Algorithmes Parallèles et Optimisation*
4. *ASAP - As Scalable As Possible: foundations of large scale dynamic distributed systems*
5. *ASCOLA - Aspect and composition languages*
6. *AVALON - Algorithms and Software Architectures for Service Oriented Platforms*
7. *CC-IN2P3 - Equipe de recherche du Centre de Calcul de l'IN2P3*
8. *CEPAGE - Chercher et Essaimer dans les Plates-formes A Grande Echelle*
9. *DOLPHIN - Parallel Cooperative Multi-criteria Optimization*
10. *GRAND-LARGE - Global parallel and distributed computing*
11. *ICPS - Scientific Parallel Computing and Imaging*
12. *KERDATA - Cloud and Grid Storage for Very Large Distributed Data*
13. *OASIS - Active objects, semantics, Internet and security*
14. *MAESTRO - Models for the performance analysis and the control of networks*
15. *MESCAL - Middleware efficiently scalable*
16. *MINC - Micro et Nanosystèmes pour les Communications sans fils*
17. *MYRIADS - Design and Implementation of Autonomous Distributed Systems*
18. *REGAL - Large-Scale Distributed Systems and Applications*
19. *RESO - Protocols and software for very high-performance network*
20. *ROMA - Resource Optimization: Models, Algorithms, and scheduling*
21. *RUNTIME - Efficient runtime systems for parallel architectures*
22. *SAGE - Simulations and Algorithms on Grids for Environment*
23. *SARA - Services and Architectures for Advanced Networks*
24. *SEPIA - Système d'exploitation, systèmes répartis, de l'intergiciel à l'architecture*
25. *ZENITH - Scientific Data Management*

Hemera: Core Participant List

1. *ACADIE - Assistance à la Certification d'Applications Distribuées et Embarquées*
2. **ALGORILLE - Algorithms for the Grid**
3. *APO - Algorithmes Parallèles et Optimisation*
4. *ASAP - As Scalable As Possible: foundations of large scale dynamic distributed systems*
5. **ASCOLA - Aspect and composition languages**
6. **AVALON - Algorithms and Software Architectures for Service Oriented Platforms**
7. **CC-IN2P3 - Equipe de recherche du Centre de Calcul de l'IN2P3**
8. *CEPAGE - Chercher et Essaimer dans les Plates-formes A Grande Echelle*
9. **DOLPHIN - Parallel Cooperative Multi-criteria Optimization**
10. *GRAND-LARGE - Global parallel and distributed computing*
11. *ICPS - Scientific Parallel Computing and Imaging*
12. **KERDATA - Cloud and Grid Storage for Very Large Distributed Data**
13. *OASIS - Active objects, semantics, Internet and security*
14. *MAESTRO - Models for the performance analysis and the control of networks*
15. **MESCAL - Middleware efficiently scalable**
16. **MINC - Micro et Nanosystèmes pour les Communications sans fils**
17. **MYRIADS - Design and Implementation of Autonomous Distributed Systems**
18. *REGAL - Large-Scale Distributed Systems and Applications*
19. *RESO - Protocols and softwares for very high-performance network*
20. *ROMA - Resource Optimization: Models, Algorithms, and scheduling*
21. *RUNTIME - Efficient runtime systems for parallel architectures*
22. **SAGE - Simulations and Algorithms on Grids for Environment**
23. **SARA - Services and Architectures for Advanced Networks**
24. **SEPIA - Système d'exploitation, systèmes répartis, de l'intergiciel à l'architecture**
25. **ZENITH - Scientific Data Management**

Hemera: Organization

■ A direction committee

- Aladdin comdir + C. Perez
- Defines research directions around the Grid5000 testbed
- Select & evaluate scientific challenges
- Select & evaluate the working groups

■ Scientific challenges

- A large-scale “experiment” on Grid5000

■ Working groups

- Identified set of teams dealing with scientific challenges

List of Scientific Challenges

- **Experiments**
 - Methods and Tools for Challenging Experiments on Grid'5000
- **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
 - OpenStack on Grid'5000
 - Large Scale Virtual Machine Deployment & Management
 - Virtual Machine Live Migration
 - I/O and in-situ visualization for HPC
- **Programming Paradigm**
 - Large Scale Computing for Combinatorial Optimization Problems
 - Scalable Distributed Processing Using the MapReduce Paradigm
 - Big Data Analytics
 - Evaluating the performance of large-scale triplestores with Big Linked Datasets
 - Low Level Component Model Enabling Performance Portability of HPC Application
- **Application Domain Specific**
 - Multi-parametric Intensive Stochastic Simulations for Hydrogeology
 - Thinking GRID for Electromagnetic Simulation of Oversized Structures

Hemera: Working Groups

■ What

- A group of people

■ Organization

- Manage by two leaders

■ Responsible of

- Leading the working group and its community
- Organizing workshops
- Potentially proposing the organization of schools

List of Working Groups

- **Transparent, Safe and Efficient Large Scale Computing**
 - Stéphane Genaud (ICPS), Fabrice Huet (OASIS)
- **Energy Efficient Large Scale Experimental Distributed Systems**
 - Laurent Lefèvre (AVALON), Jean-Marc Menaud (ASCOLA)
- **Bring Grids Power to Internet-Users thanks to Virtualization Technologies**
 - Adrien Lèbre (ASCOLA), Yvon Jégou (MYRIADS)
- **Efficient exploitation of highly heterogeneous and hierarchical large-scale systems**
 - Olivier Beaumont (CEPAGE), Eric Vivien (ROMA)
- **Efficient management of very large volumes of information for data-intensive applications**
 - Gabriel Antoniu (KERDATA), Jean-Marc Pierson (SEPIA)
- **Completing challenging experiments on Grid'5000**
 - Lucas Nussbaum (ALGORILLE), Olivier Richard (MESCAL)
- **Modeling Large Scale Systems and Validating their Simulators**
 - Martin Quinson (ALGORILLE), Arnaud Legrand (MESCAL)
- **Network metrology and traffic characterization**
 - Paulo Gonçalves (RESO), Konstantin Avrachenkov (MAESTRO)

Héméra

Resources

Resources

- **PhD Students**
 - For scientific challenges requiring deeper studies.
- **Post-doc/Internships**
 - For short term contributions to challenges.
- **Engineers**
 - For helping making experiments on Grid'5000.
 - Fill the gap between the Aladdin ADT and end users
 - Delegated to partners for a short period of time on a well-defined goal
 - For developing/improving experimentation tooling
- **Missions**
 - For organizing technical meetings, workshops, PhD mobility, etc.

Resources - PhD

- 2010–2013: Energy – Astre/Reso
 - Energy profiling and green leverages for services and applications in large scale distributed systems
- 2011–2014: COPS – Dolphin
 - Robust Peer-to-Peer Algorithms for Large Scale COPs
- 2011–2014: Experiments – Mescal/Laas
 - Methods and tools for challenging experiments on Grid'5000 : a use case on electromagnetic hybrid simulation
- 2012–2015: Data – Zenith
 - Optimizing a Cloud for Data Mining primitives
- 2013–2016: Data – KerData
 - Scalable, Power-efficient Big Data Analysis on Geographically Distributed Clouds
- 2013–2016: Energy – Astre/Avalon
 - Energy proportionality in large scale virtualized environments

Resources: Post-doc

- 2010–2011: COPs – Dolphin
 - Fault-Tolerant Distributed Branch-and-Bound on the Grid
- 2011–2012: Data – Kerdata
 - Consistency, availability, scalability : Building the infrastructure for geo-distributed and geo-replicated cloud storage
- 02/2013–01/2014*: Energy – Ascola/IRIT (*: on 2012 budget)
 - Energy monitoring, from the VM to the room
- 2013–2014: Energy – Myriads/Avalon
 - Evaluating the energy consumption of large-scale applications
- 8/11/2013–13/02/2014: Experiments – Algorille
 - Open Science and Reproducible Research on Distributed Systems

Hemera Engineers

■ People

- IJD – Daniel Balouek (2/2011 – 1/2013)
- IC – Laurent Poulloux (10/2012 – 10/2014)

■ Link to Grid'5000 Technical Committees

■ Missions

- D. Balouek: managed several challenges
 - Hydro, Electro, DynVM, L2C, Diet
- L. Pouilloux
 - Help scientists in setting up their challenges
 - Improve G5K tooling

Conclusion

- **Experimental platforms (and observation instruments) are essential** in the CS methodology - like in other sciences!
- Many research fields need experimental testbed such as Grid'5000
 - HPC, Grids (Classical/Desktop), Clouds, Distributed, Green, Big Data, etc
 - A validation tool for applications/middleware before going to production
- **Hemera**
 - Enable large-scale experiments
 - Improve experimentation tooling + dissemination of good practice
 - Stabilization of Grid'5000 (corner case + advanced feature tests)
 - 16 challenges, rather dynamic, some failures, a lot of successes
 - Evolution of theme like cloud, virtualization, big data
 - 8 working groups, different organizations, low support of Hemera
 - Three of them stop after 2 years
- **Related follow up projects**
 - IPL Discovery and HAC-SPECIS
 - NSF Chameleon

Status in December 2014

- **Bring new teams to use Grid'5000**
 - In particular, around big data
- **Simplify making experiments**
 - **Shared** knowledge of using the platform
 - **More solid tools** – many less tricks needed
 - **Little** PhD time needed for making experiments
- **Fruitful exchanges between Grid'5000 and scientists**
 - Impact both
- **More large scale experiments**
 - Tooling ahead of teams needs?
 - Large scale + virtualization + energy + data