Using GridWay for the Management of Virtual Machines in the EGEE Infrastructure

(2nd EGEE User Forum. May 10th, 2007)

A. J. Rubio Montero.
CIEMAT.
I. M. Llorente, R. S. Montero, E. Huedo.
Universidad Complutense de Madrid.
D. Tapiador
ESA/ESAC.

www.eu-egee.org

Information Society and Media

Contents

1. Grids & Virtual Machines
2. The XMM-Newton Science Analysis Software
3. Management of Virtual Machines with GridWay
4. Experience & Results
5. Conclusions

Grids Infrastructures

- International research projects (EGEE, OSG, TeraGrid)
  - **Resource sharing**
  - Increase the **computational and storage resources**
- **High degree of heterogeneity** (software & hardware)
  - Increase the cost and length of application development cycle
  - Limit the use of the infrastructure
- **Isolate and partition amount of resources** contributed to the Grid.

Virtual Machines

- Renewed interest on virtualization technologies
  - Processor's performance and support for VM (e.g. Intel VT)
- Add a new abstraction layer to the Grid:
  - Natural way to deal with the **heterogeneity** of the infrastructure
  - Allow **partitioning and isolating** of physical resources (access to your HW not to your system)
  - Execution of **legacy applications** or scientific codes.

The XMM-Newton SAS

- SAS Science Analysis Software
  - Analysis of the data provided by XMM-Newton
  - Frequently released
  - Support for several platforms (OS, hardware)
  - Must be deployed in all the Grid resources

- Impose a significant effort
  - System admin staff
  - Developers
  - Users which may need specific versions

The XMM Newton satellite

The XMM-Newton SAS: Goal

- Leverage **actual Grid infrastructures** with the use of virtual machines:
  - **Straight-forward deployment** using existing middleware
  - Based on well-tested and standard services (**production-ready**)
  - Not tied to a given virtualization technology
  - Co-exist within other Grid infrastructures

Management of VMs with GridWay

- Encapsulate a virtual machine in a grid job.
  - Incorporate the **functionality of a general purpose meta-scheduler**
  - Do **not need new middleware**
  - The underlying LRMS is not aware of the **nature of the job**
  - Only suitable to medium/coarse grained **HTC applications**.

Management of VMs with GridWay

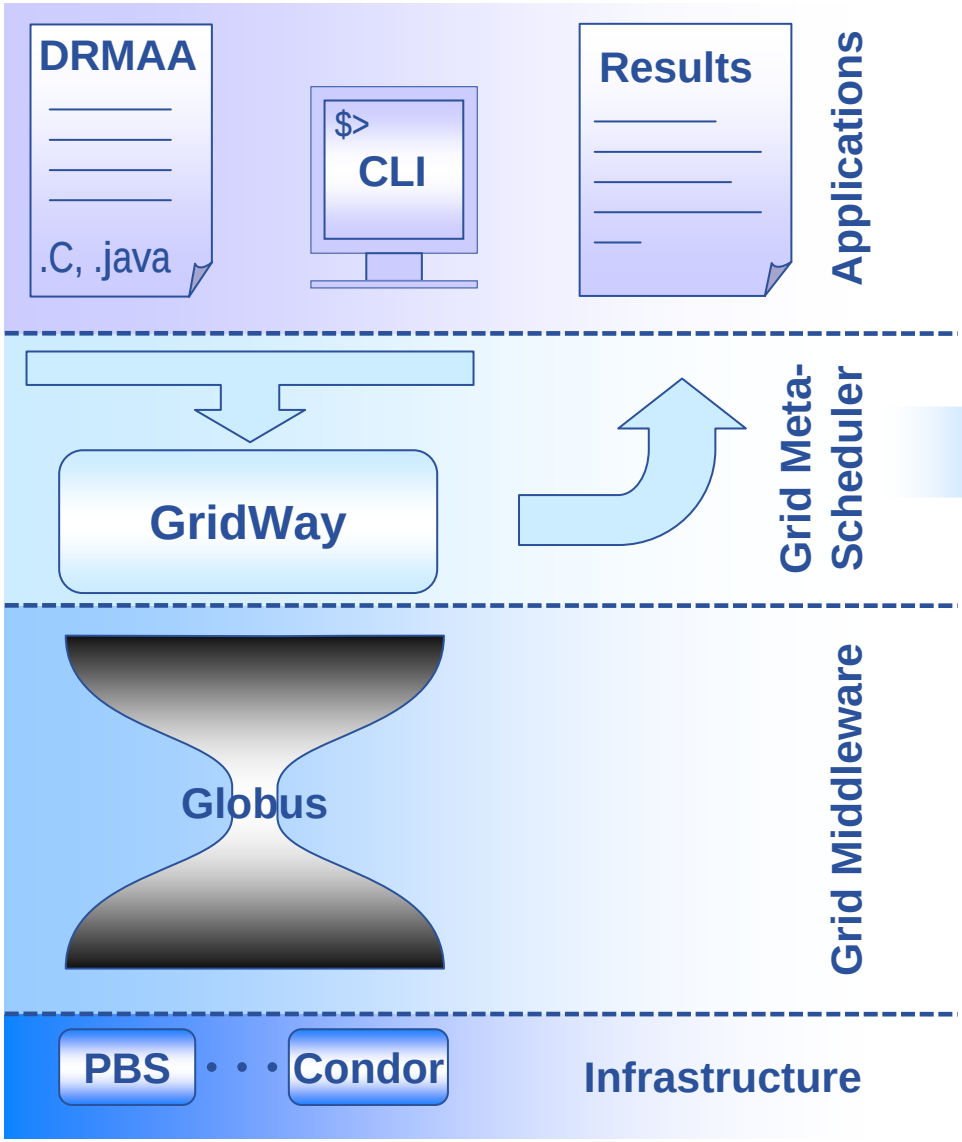

the globus alliance

- **Advanced scheduling**
- Different **application profiles**
- **Fault detection & recovery**
- Job **execution management**
  - Prolog (stage-in)
  - Wrapper (execution)
  - Epilog (stage-out)

Management of VMs with GridWay

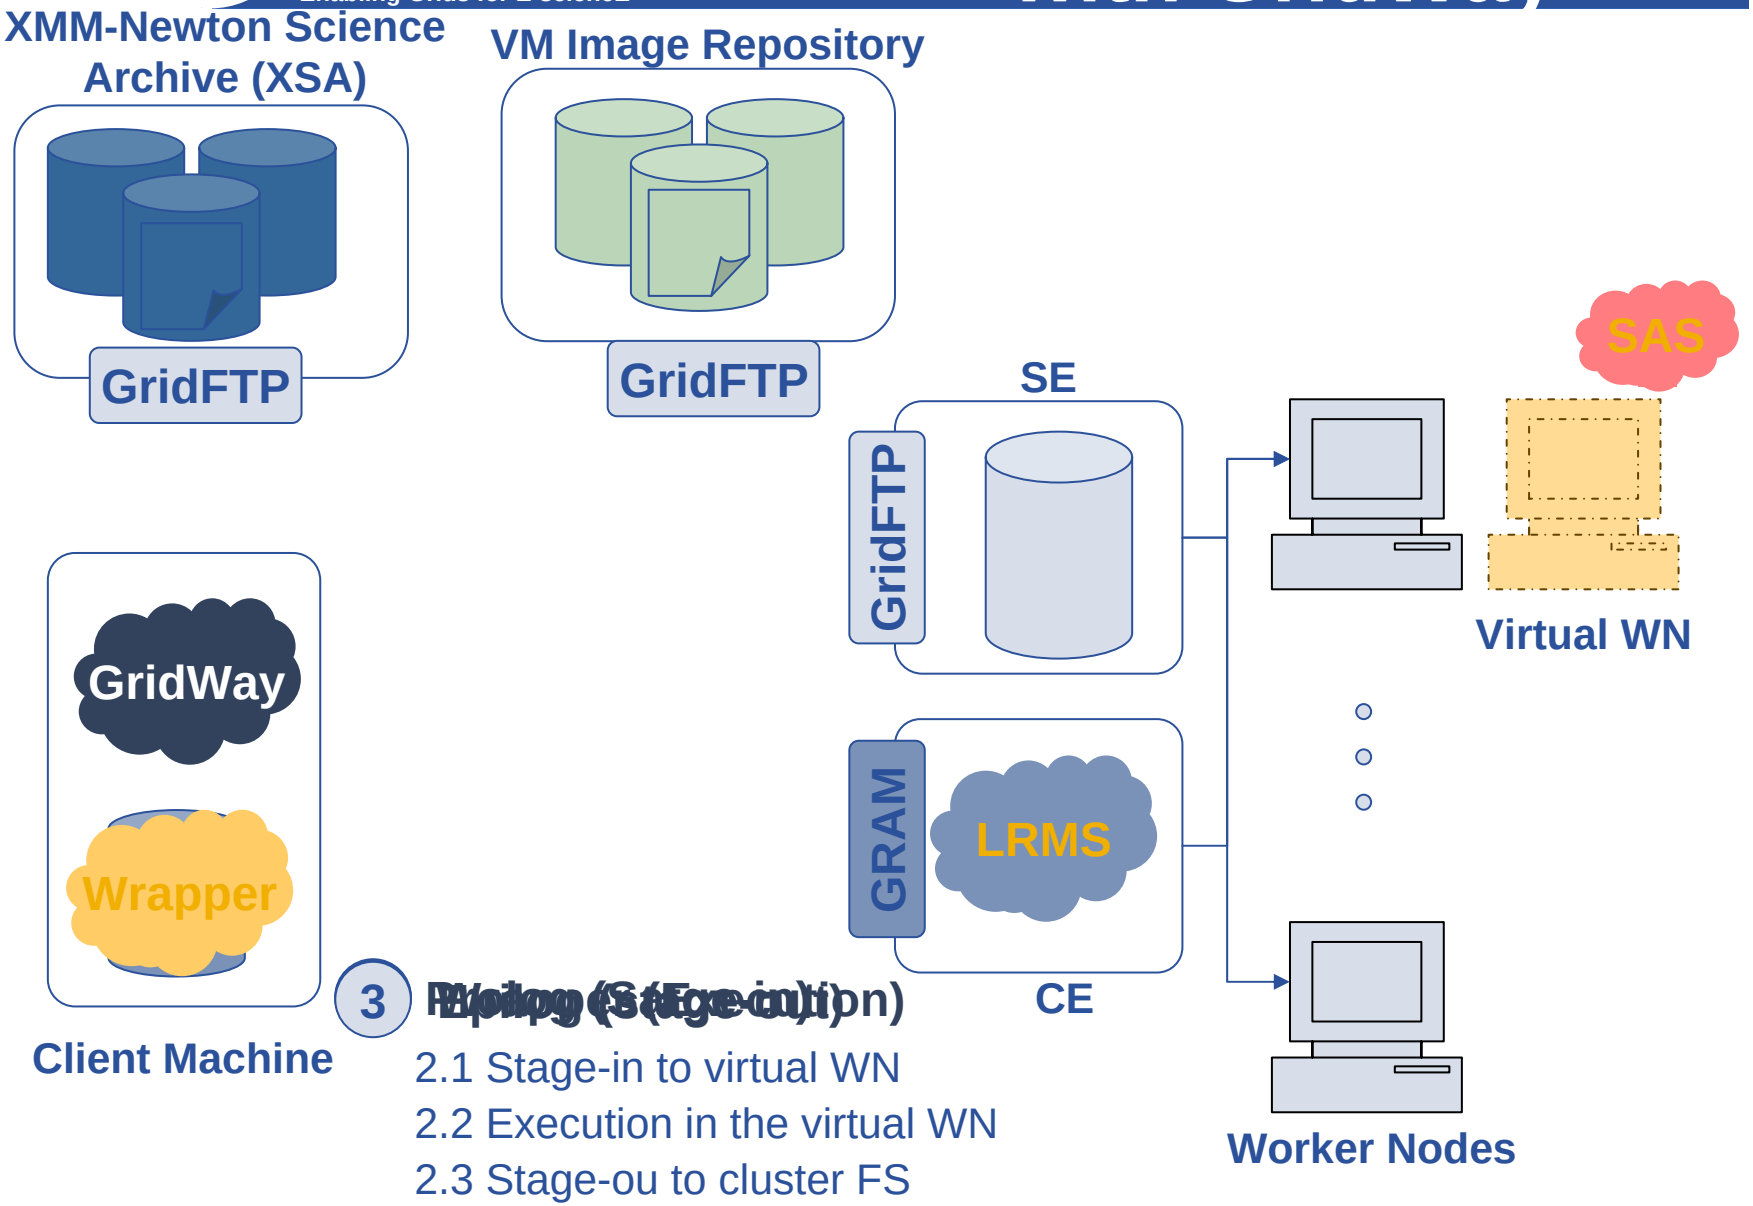

Experience & Results

- Some implementation details
  - Virtual Images available at remote resources.
  - Virtual OS has been hardened

Disk Image layout

Mount Point	Size	Contents
/	500MB	FC4 base system
/usr	650MB	System apps.
/opt	600MB	SAS 6.5.0
/scratch	2GB	SAS tmp

Shared save storage and ease the deployment

Local for I/O performance

Experience & Results: Overhead

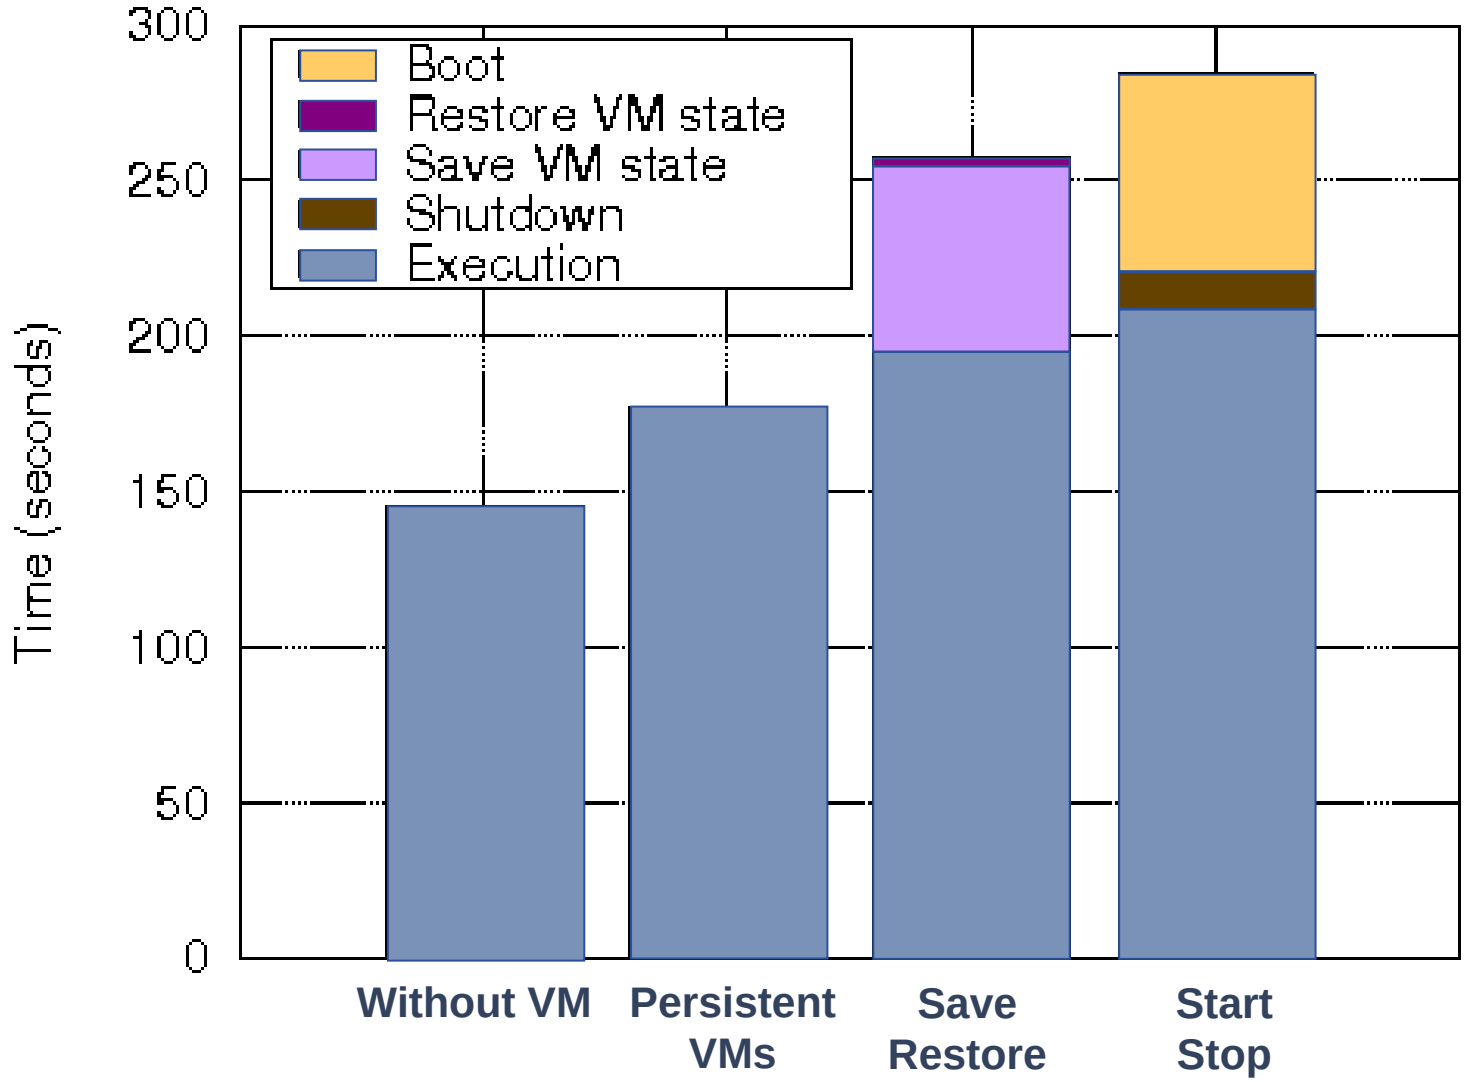

Experience & Results: Application Performance

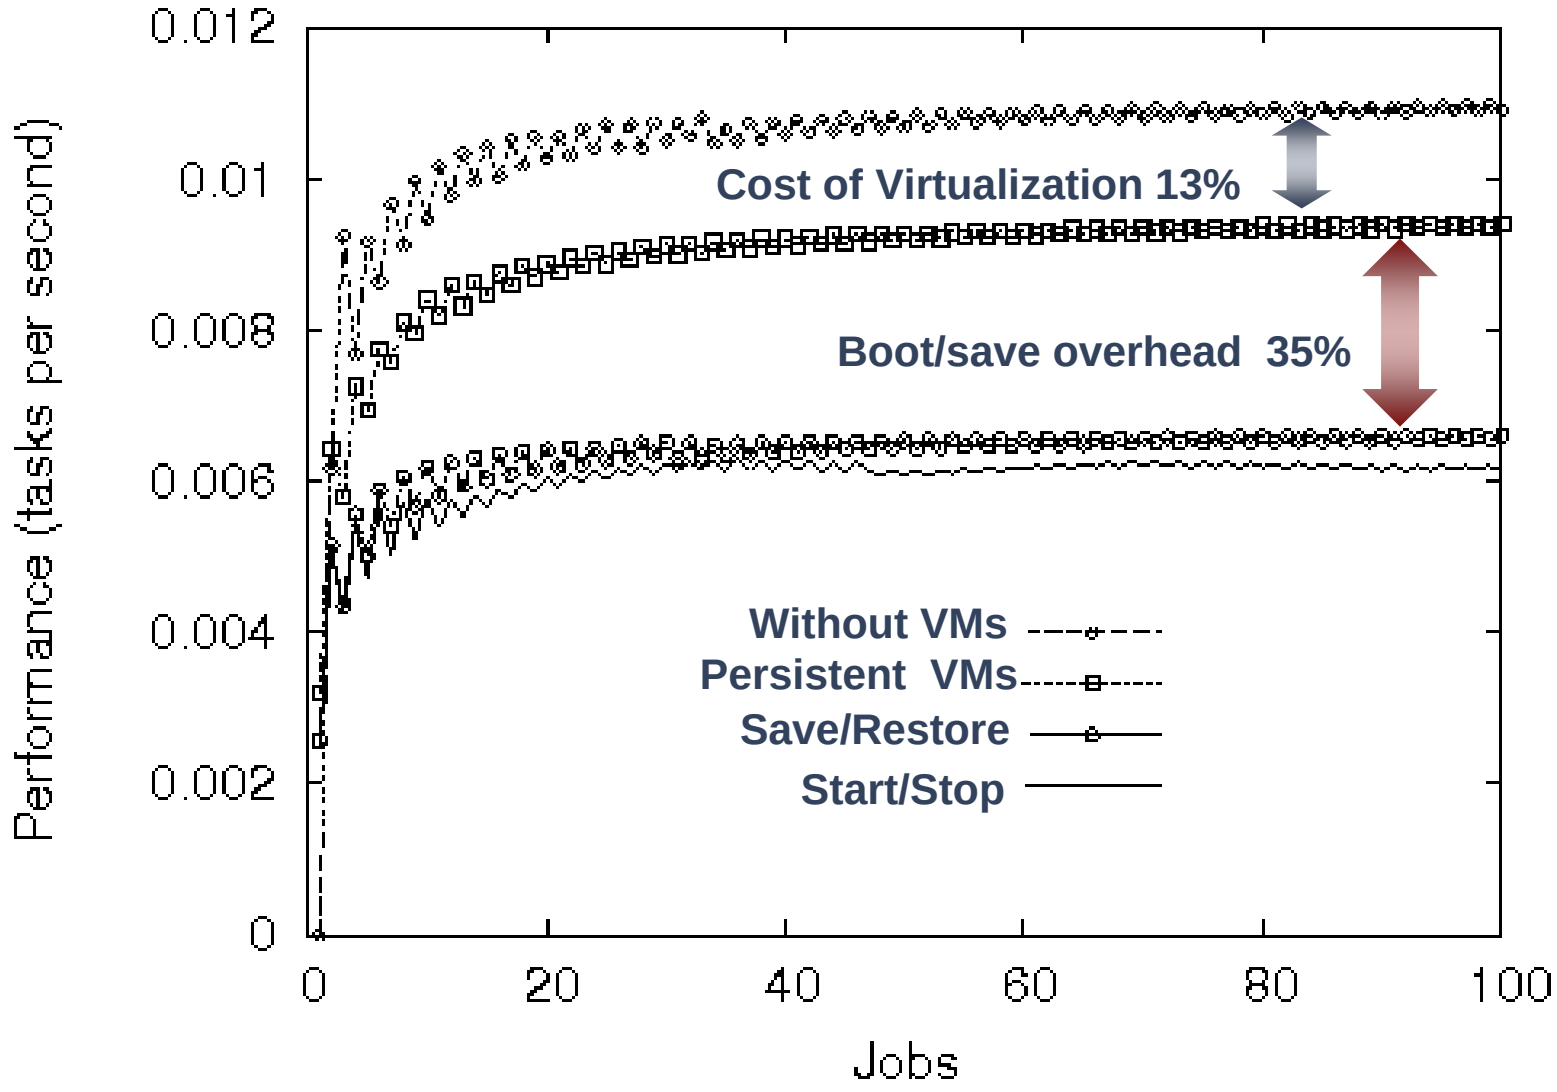

Conclusions & Current Work

- Straightforward deployment of VMs on Grids.
  - **Ready to work on existing infrastructures**
  - **Limited overhead** for some deployments
  - Significant increase of the *quality of life* in the Grid
- However:
  - Does not fully exploit virtualization
  - Limited to medium to coarse grained batch applications

- Provide persistent virtual overlay management system
- **Is building a EGEE/EELA VO Testbed now**.