

**Open Grid Forum, GridVirt Session
May 7-11 2007, Manchester UK**



Management of Virtual Machines in Grids Infrastructures

Rubén S. Montero
<http://asds.dacya.ucm.es>



**Distributed Systems Architecture Group
Universidad Complutense de Madrid**



Contents

Management of Virtual Machines in Grid Infrastructures

1. Grids & Virtual Machines
2. Case of Study: XMM-Newton SAS
3. Management of Virtual Machines with GridWay
4. Experience & Results
5. The GridHypervisor Project



Grids & Virtual Machines

Management of Virtual Machines in Grid Infrastructures

Grid Infrastructures

- International research projects (EGEE, OSG, TeraGrid)
 - Unprecedented levels of **resource sharing**
 - Increase the **computational and storage resources**
- **High degree of heterogeneity** (software & hardware)
 1. Increases the cost and length of application development cycle
 2. Limits the effective number of resources available to the user
 3. Increases the operational cost of the infrastructure
- **Isolate and partition amount of resources** contributed to the Grid.
- Even worse when using different Grids!



Grids & Virtual Machines

Management of Virtual Machines in Grid Infrastructures

Current Approaches

- **Software-environment configuration systems**
 - For the Users, define what applications they want to use
 - For the Sysadmins, make a applications available to the user
 - Example: *SoftEnv*
 - Does not completely solved none of the problems
- Deployment of **software overlays**
 - Custom software configurations (user-space)
 - Managed by a general-purpose LRM
 - Issues shifted from the applications to the overlaid infrastructure
 - Examples: Condor Glide-in, GridWay-BOINC



Grids & Virtual Machines

Management of Virtual Machines in Grid Infrastructures

Virtual Machines

- Renewed interest on virtualization technologies (performance and VM support)
- 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.



Provide highly distributed environments with a powerful tool to overlay custom infrastructures on top of physical resources



Case of Study: XMM-Newton SAS

Management of Virtual Machines in Grid Infrastructures



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
 - Sysadmin staff
 - Developers
 - Users which may need specific versions



The XMM Newton satellite



Case of Study: XMM-Newton SAS

Management of Virtual Machines in Grid Infrastructures

Goal

Leverage **actual Grid infrastructure** with virtual the use of 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 Virtual Machines with GridWay

Management of Virtual Machines in Grid Infrastructures

Idea

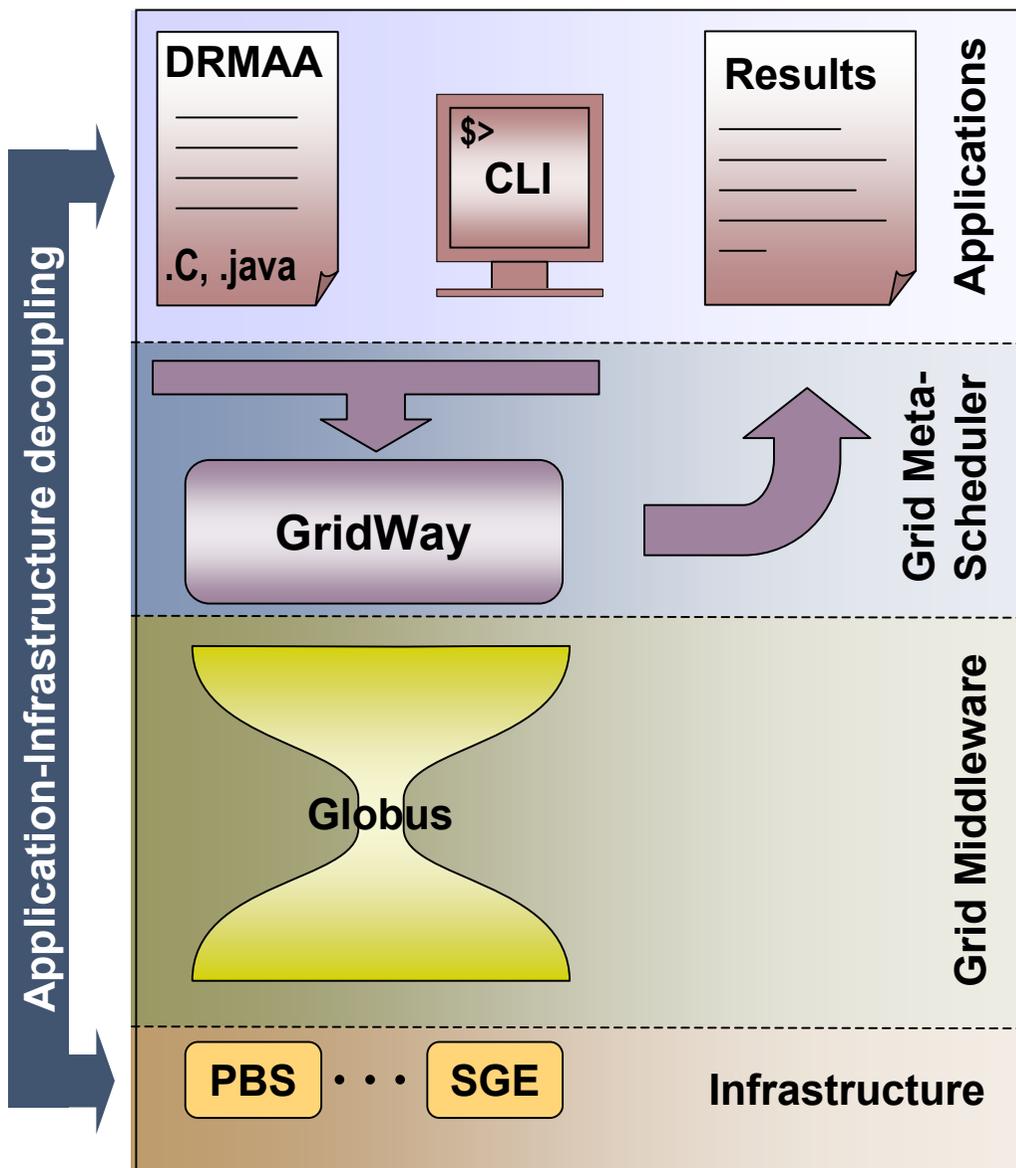
Encapsulate a virtual machine in a grid job.

- 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**.

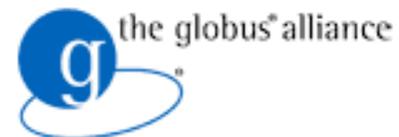


Management of Virtual Machines with GridWay

Management of Virtual Machines in Grid Infrastructures



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



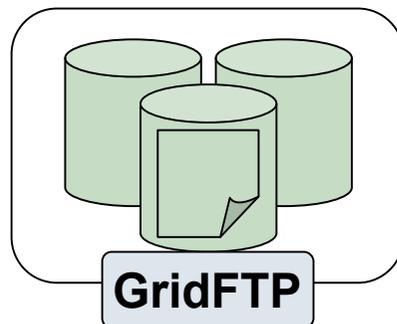
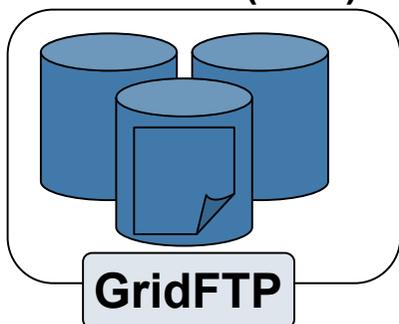


Management of Virtual Machines with GridWay

Management of Virtual Machines in Grid Infrastructures

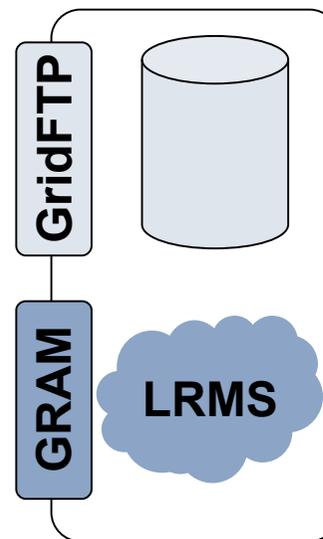
XMM-Newton Science VM Image Repository

Archive (XSA)

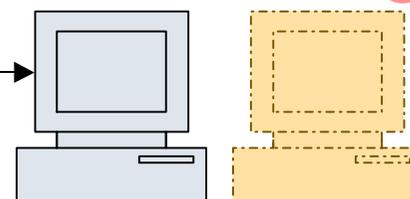


Client Machine

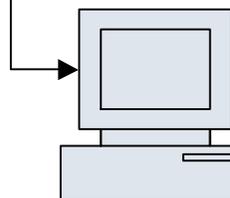
Front-end



SAS



Virtual WNs



Worker Nodes

3. Policy (Stage-in)

- 2.1 Stage-in to virtual WN
- 2.2 Execution in the virtual WN
- 2.3 Stage-out to cluster FS



Experience & Results

Management of Virtual Machines in Grid Infrastructures

Some Implementation Details

- Xen capable resources (Worker-nodes)
- Virtual Images available at remote resources.
- Virtual OS has been hardened

Disk Image layout

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



Shared save storage
and ease the deployment



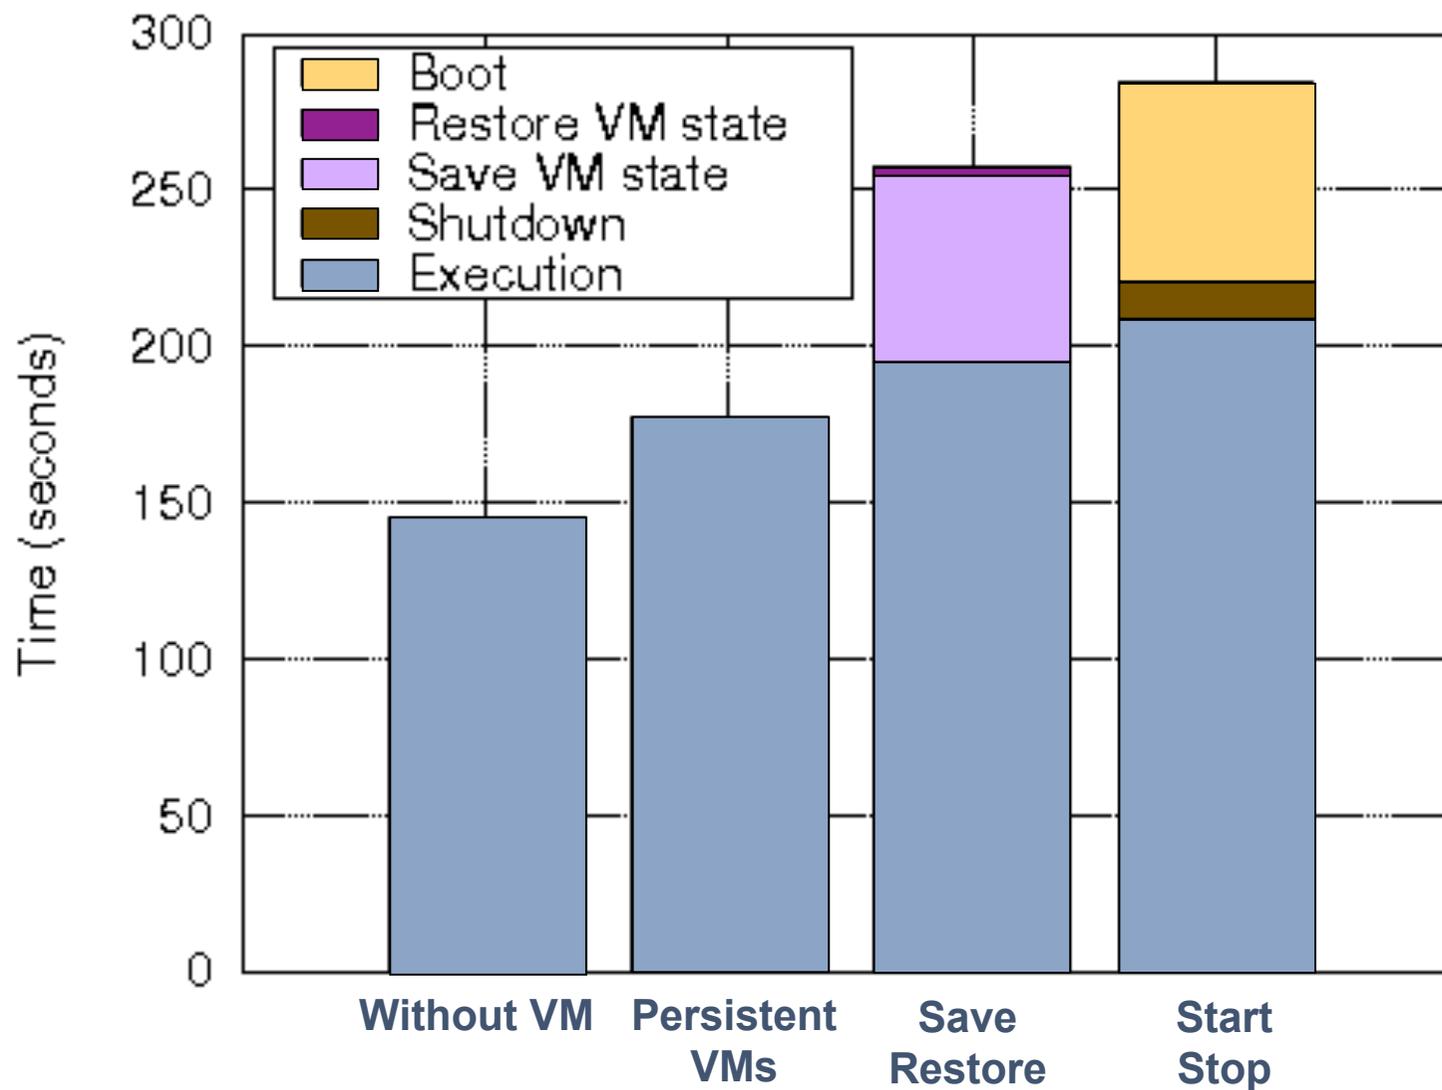
Local for I/O performance



Experience & Results

Management of Virtual Machines in Grid Infrastructures

Overhead Analysis

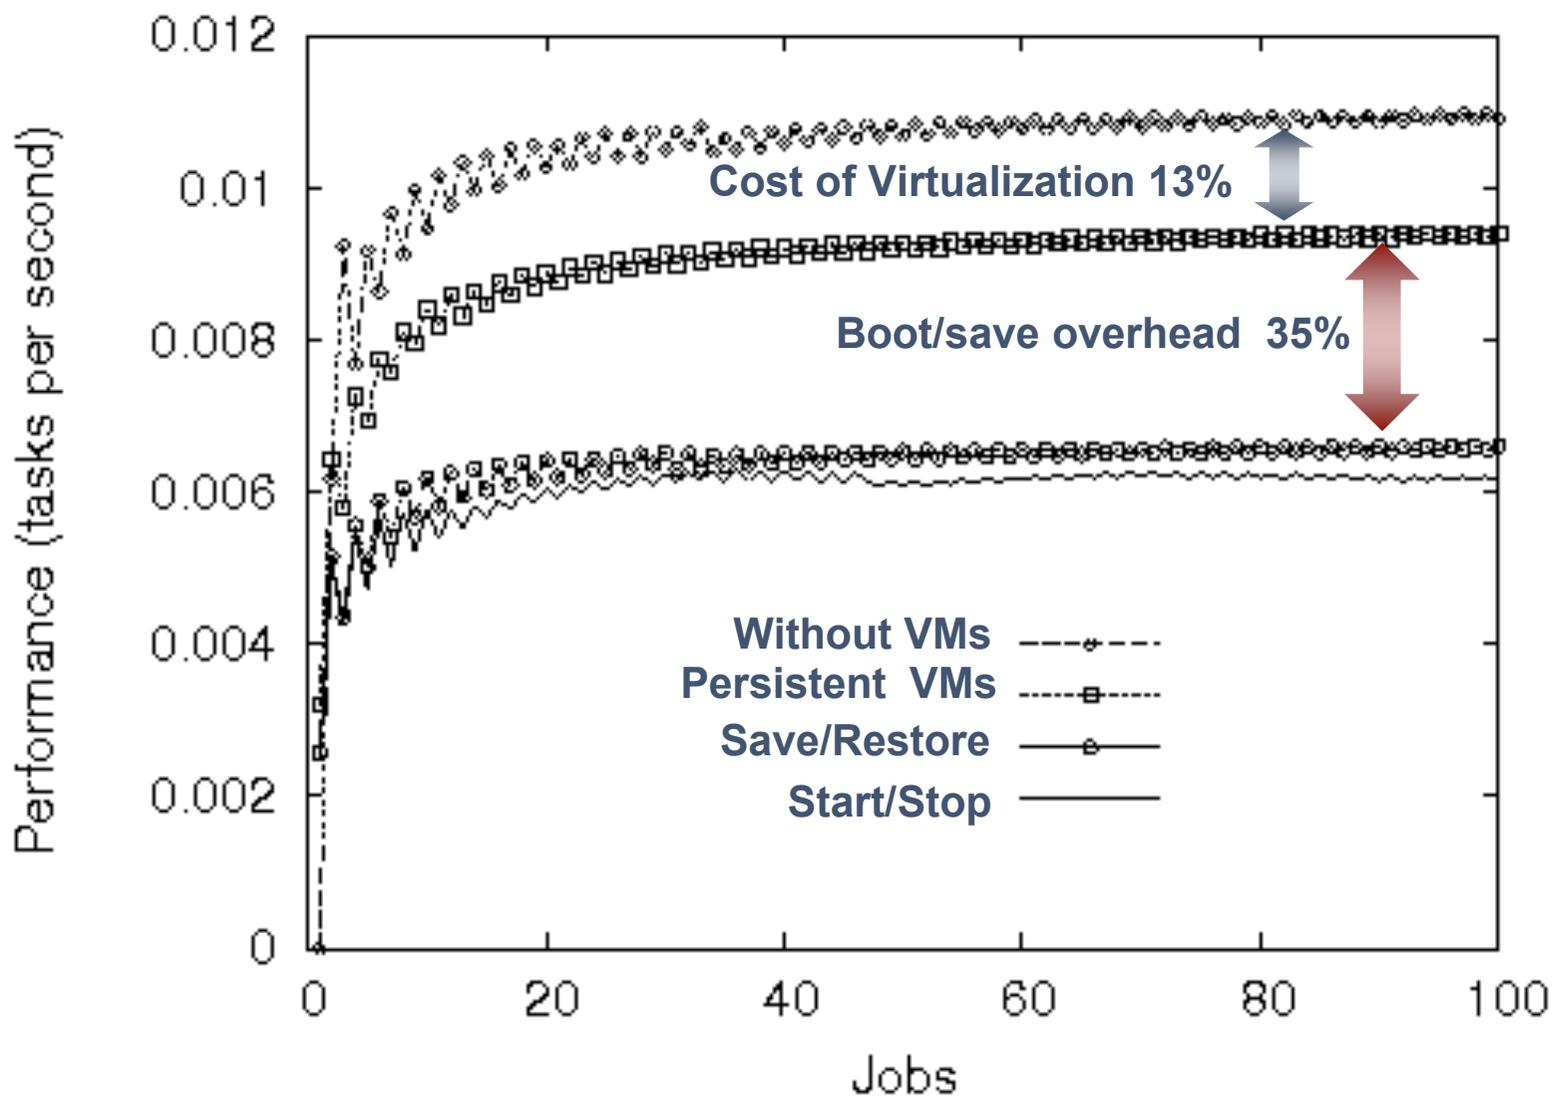




Experience & Results

Management of Virtual Machines in Grid Infrastructures

Application Performance





Experience & Results

Management of Virtual Machines in Grid Infrastructures

Experience

- 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



The GridHypervisor Project

Management of Virtual Machines in Grid Infrastructures

Goal of the Project

Enable large-scale, reliable and dynamic deployment and reallocation of VMs between different administration domains

- The GridHypervisor will manage *Virtual Infrastructures (VI)* in the same way a Hypervisor manage Virtual Machines (VM)

Functionality

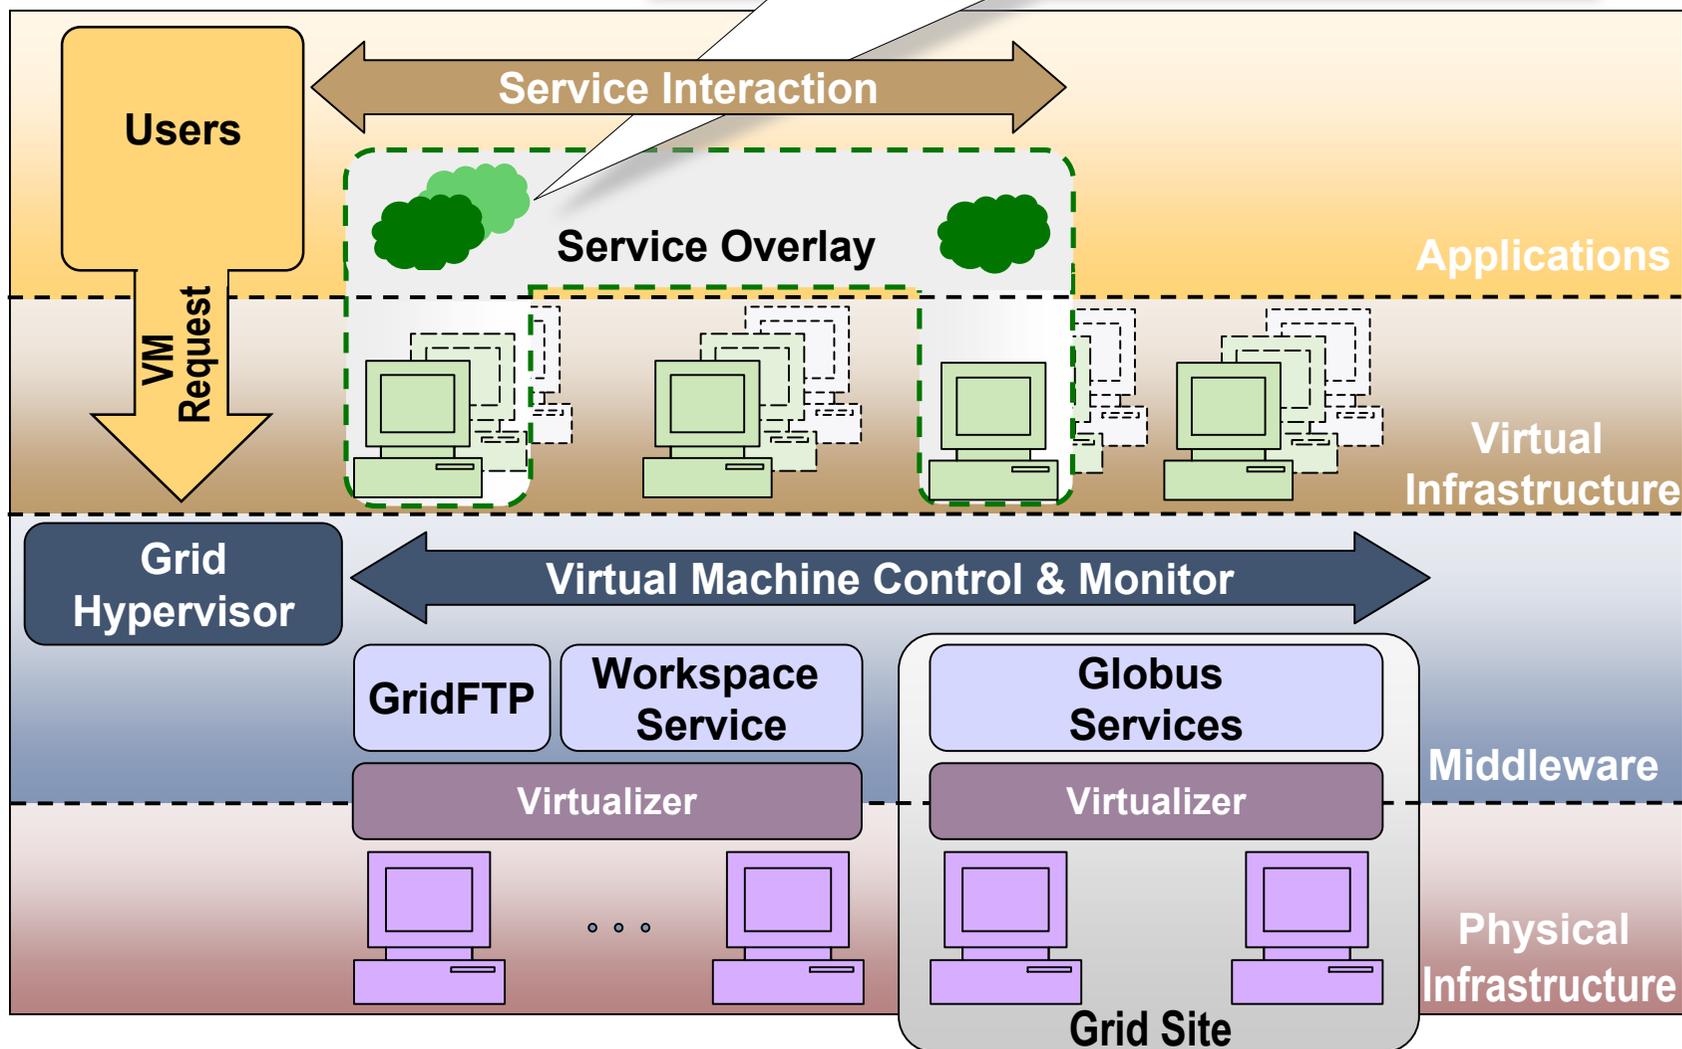
- Based on existing middleware, Globus Workspace Service:
 - WSRF interface
 - Start, monitor and control a **Worksapce**
- VI requests are handled by the GridHypervisor:
 - Physical resource Discovery and Selection
 - Remote System Preparation
 - VM Start/Migration
 - VM Monitoring
 - Termination



The GridHypervisor Project

Management of Virtual Machines in Grid Infrastructures

- Service process
- Service-ready VM Images
- Re-configured on each site
- Examples: a SGE cluster, a virtual classroom...





**Thank you for
your attention!**

More Information at...

- www.gridhypervisor.org
- www.gridway.org