



the globus® toolkit
www.globustoolkit.org

Virtual Workspaces

Dynamic Virtual Environments in the Grid

October 5, 2006 – ESAC Grid Workshop '06

Borja Sotomayor – University of Chicago

Index

- Virtual Workspaces
 - ◆ What is a workspace?
 - ◆ Why are *VM-based* workspaces interesting?
- GT4 Workspace Service
 - ◆ Globus Workspace implementation
 - ◆ Current release: TP1.2
- Research



Before we begin...

- **Workspace Team**
 - ◆ Kate Keahey (ANL)
 - ◆ Tim Freeman (ANL)
 - ◆ Borja Sotomayor (UChicago)
 - ◆ With lots of help and insights from Ian Foster (ANL+UC) and Frank Siebenlist (ANL)
- **Globus is 10 years old!**

Before we begin...

- Totally unrelated to this talk, but some of you might find it interesting:
 - ◆ AstroPortal, being developed at the University of Chicago by Ioan Raicu and Ian Foster.
 - ◆ Using Grid Computing as the main mechanism to enable the dynamic analysis of large astronomy datasets on the TeraGrid.
 - ◆ <http://people.cs.uchicago.edu/~iraicu/research/AstroPortal/>

Index

- Virtual Workspaces
 - ◆ What is a virtual workspace?
 - ◆ Why are *VM-based* workspaces interesting?
- GT4 Workspace Service
 - ◆ Globus Workspace implementation
 - ◆ Current release: TP1.2
- Research


What is a workspace?

- What do I do if I want to “run something in the Grid”?
 - ◆ We have to encapsulate our work in a job abstraction.
 - ◆ This abstraction imposes a lot of restrictions on the resource consumer.
 - For example: Resource consumers can specify what libraries their jobs require, but they (usually) can't have those libraries dynamically installed when the job is submitted. If the library is not already available in the resources, they generally have to make a prior arrangement with resource providers to make sure that software is installed.
 - ◆ However, these restrictions are usually perfectly acceptable.
 - ◆ But... what if they're not?



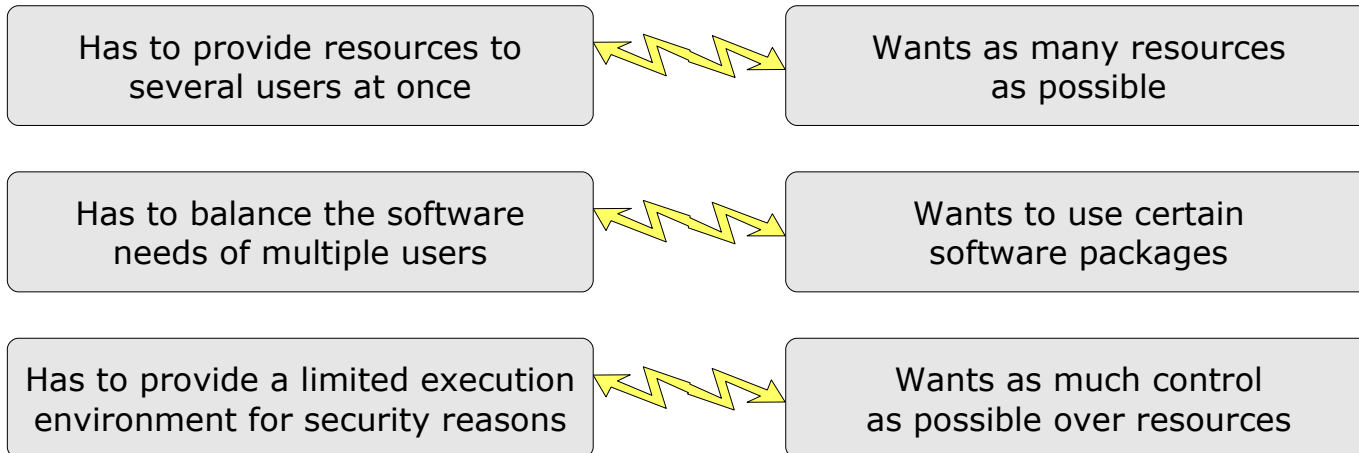
Resource provider

Provides computational, storage, and network resources



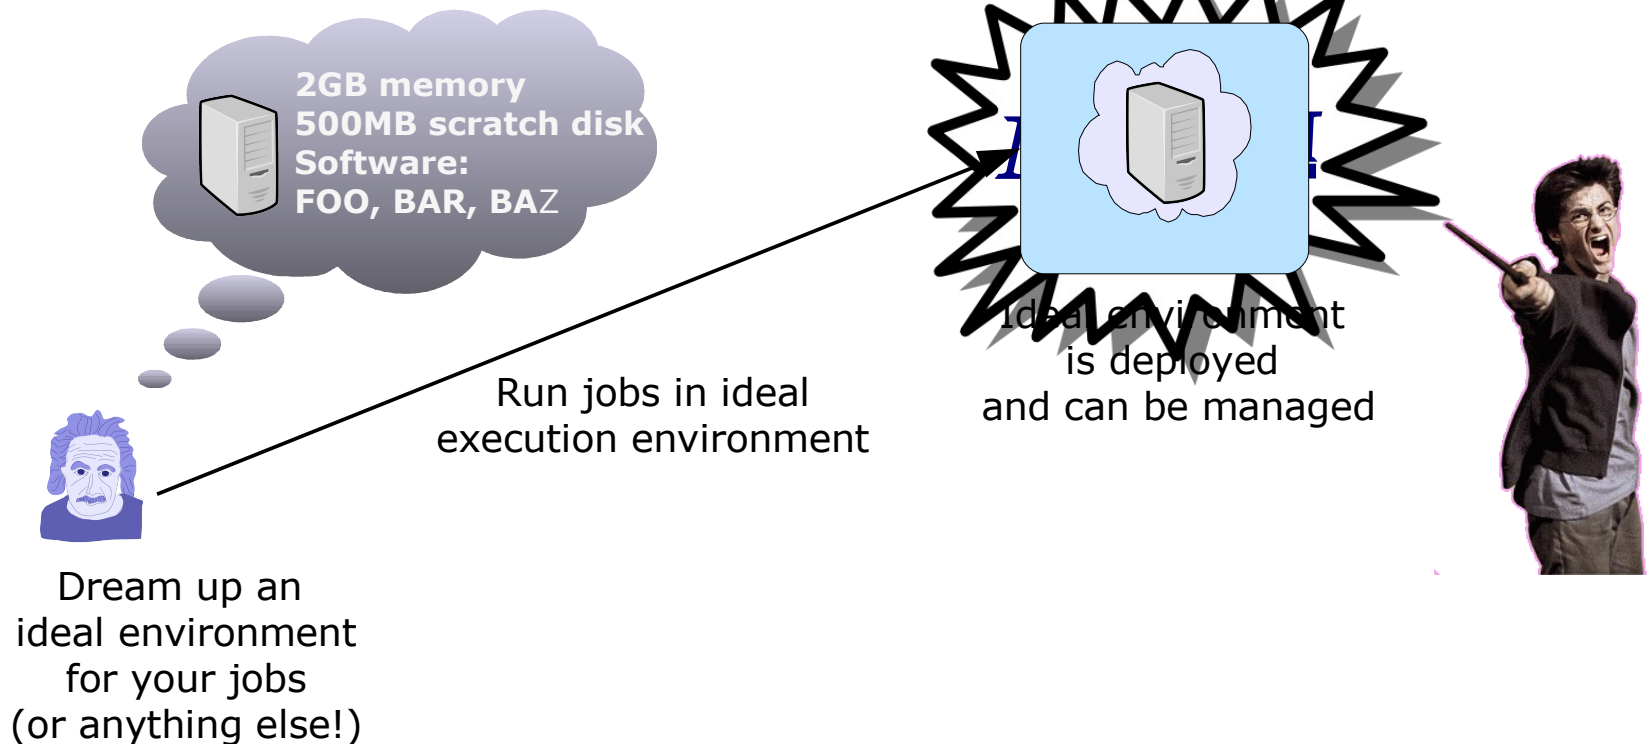
Resource consumers

Want to run experiments on the resources, but they each have different software and hardware requirements



What is a workspace?

- In some cases, resource consumers need *execution environments* to be *dynamically* deployed on *remote* resources.
- The virtual workspace is an abstraction for such an execution environment.



What is a workspace?

- Of course, this is not an entirely new idea...
It is possible to create custom execution environments by:
 - ◆ Dynamically setting up cluster nodes:
 - CoD: <http://www.cs.duke.edu/nicl/cod/>
 - bcfg: <http://trac.mcs.anl.gov/projects/bcfg2/>
 - ◆ Providing access to existing installation
 - Dynamic Accounts: <http://workspace.globus.org/da/>
 - ◆ Refining site configuration
 - Pacman: <http://www.archlinux.org/pacman/>

What is a workspace?

- Two aspects of workspaces:

- ◆ *Environment definition*: We get exactly the (software) environment we need on demand.

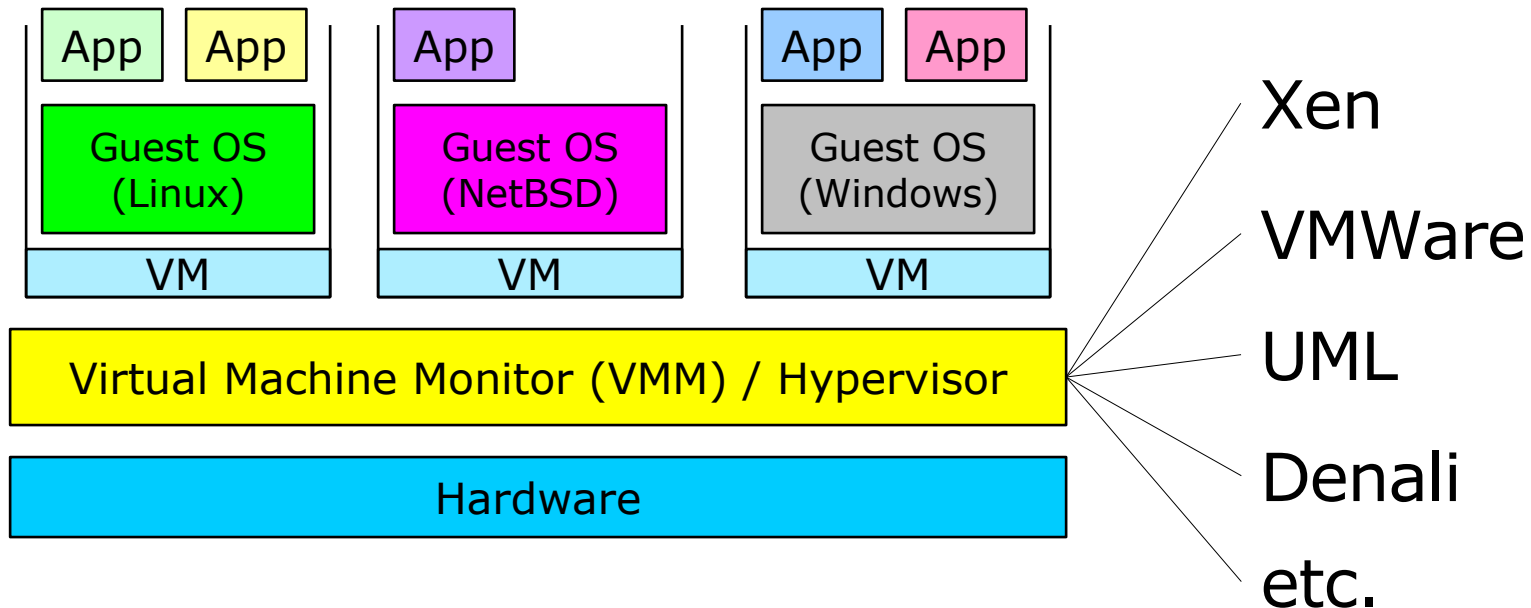
- ◆ *Resource allocation*: Provision and guarantee all the resources the workspace needs to function correctly (CPU, memory, disk, bandwidth, availability), allowing for dynamic renegotiation to reflect changing requirements and conditions.

- Existing implementations either don't provide both, or...

- ◆ *Quality of life*: Setting up a new software environment takes a long time, and still doesn't give the resource consumer *full* control.
- ◆ *Quality of service*: Little or no enforcement.

Virtual Machines

- VM technology is a promising way to achieve higher quality workspaces.



Performance: Paravirtualization (e.g. Xen) is very close to raw physical performance

Virtual Machines

- VM technology is a promising way to achieve higher quality workspaces.
 - ◆ Excellent isolation
 - ◆ Promising enforcement properties
 - ◆ Flexible configuration
 - ◆ The ability to serialize and migrate
- Workspaces can be encapsulated in VM images, and dynamically deployed on VMM-enabled sites.

Use cases

- Science
 - ◆ Batch jobs that require a very specific software environment
 - ◆ Interactive experiments
 - ◆ Resource-hungry event-driven jobs
- Education
 - ◆ Virtual labs

Index

- Virtual Workspaces
 - ◆ What is a workspace?
 - ◆ Why are *virtual* workspaces interesting?
- **GT4 Workspace Service**
 - ◆ Globus Workspace implementation
 - ◆ Current release: TP1.2
- Research

GT4 Workspace Service

- The GT4 Virtual Workspace Service (VWS) is a VM-based workspace implementation.
 - ◆ GT4 WSRF frontend
 - ◆ Xen-based, but other VMMs could potentially be used.
 - ◆ dev.globus Incubator project (<http://dev.globus.org/>)
 - ◆ Current release is TP1.2

More on
this later

GT4 Workspace Service

- A VWS workspace is composed of:
 - ◆ VM image template/s
 - Template: Some configuration (such as the IP address) is not specified in the image itself.
 - ◆ Workspace metadata (XML document). Includes deployment-independent information:
 - VMM and kernel requirements
 - NICs + IP addresses
 - VM image/s to use
- When a workspace is deployed, we must specify a *deployment request*, which includes deployment-dependent information:
 - ◆ Resource allocation: Duration, CPU %, memory, network bandwidth, ...

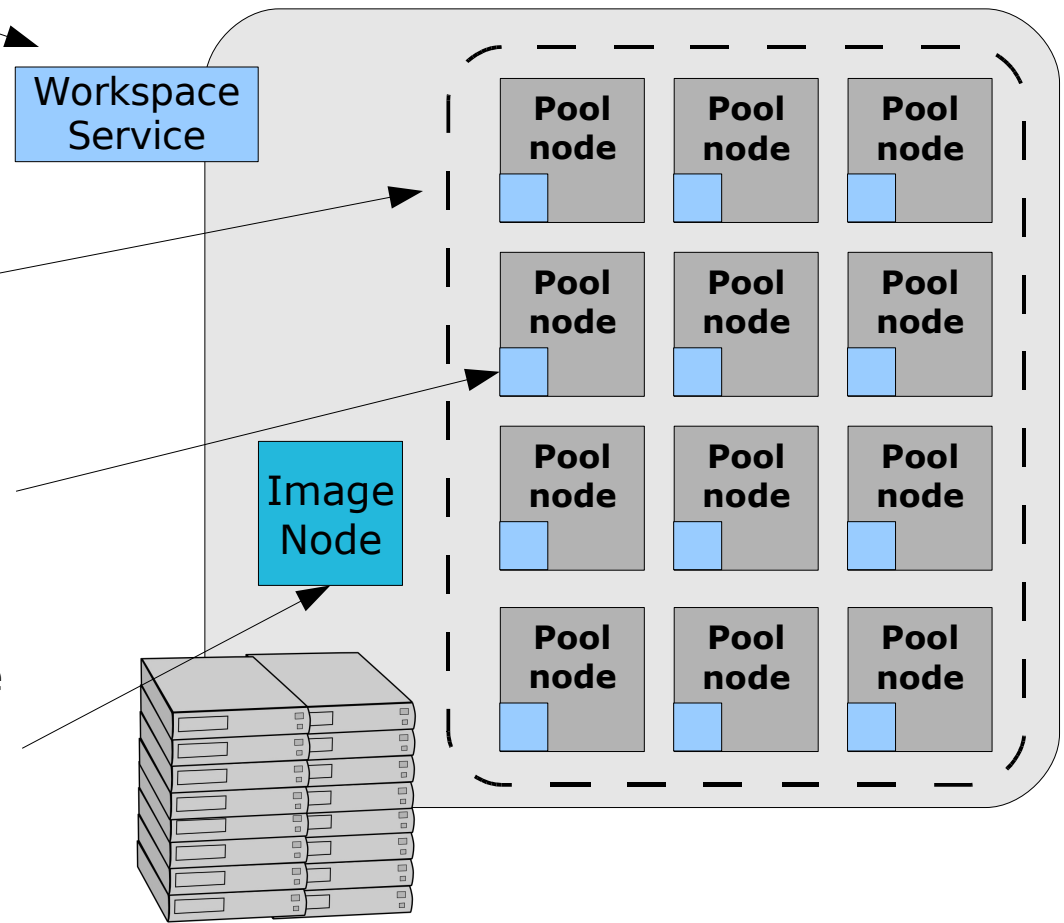
GT4 Workspace Service

The workspace service has a WSRF frontend that allows users to deploy and manage virtual workspaces

The VWS manages a set of nodes inside the TCB (typically a cluster). This is called the *node pool*.

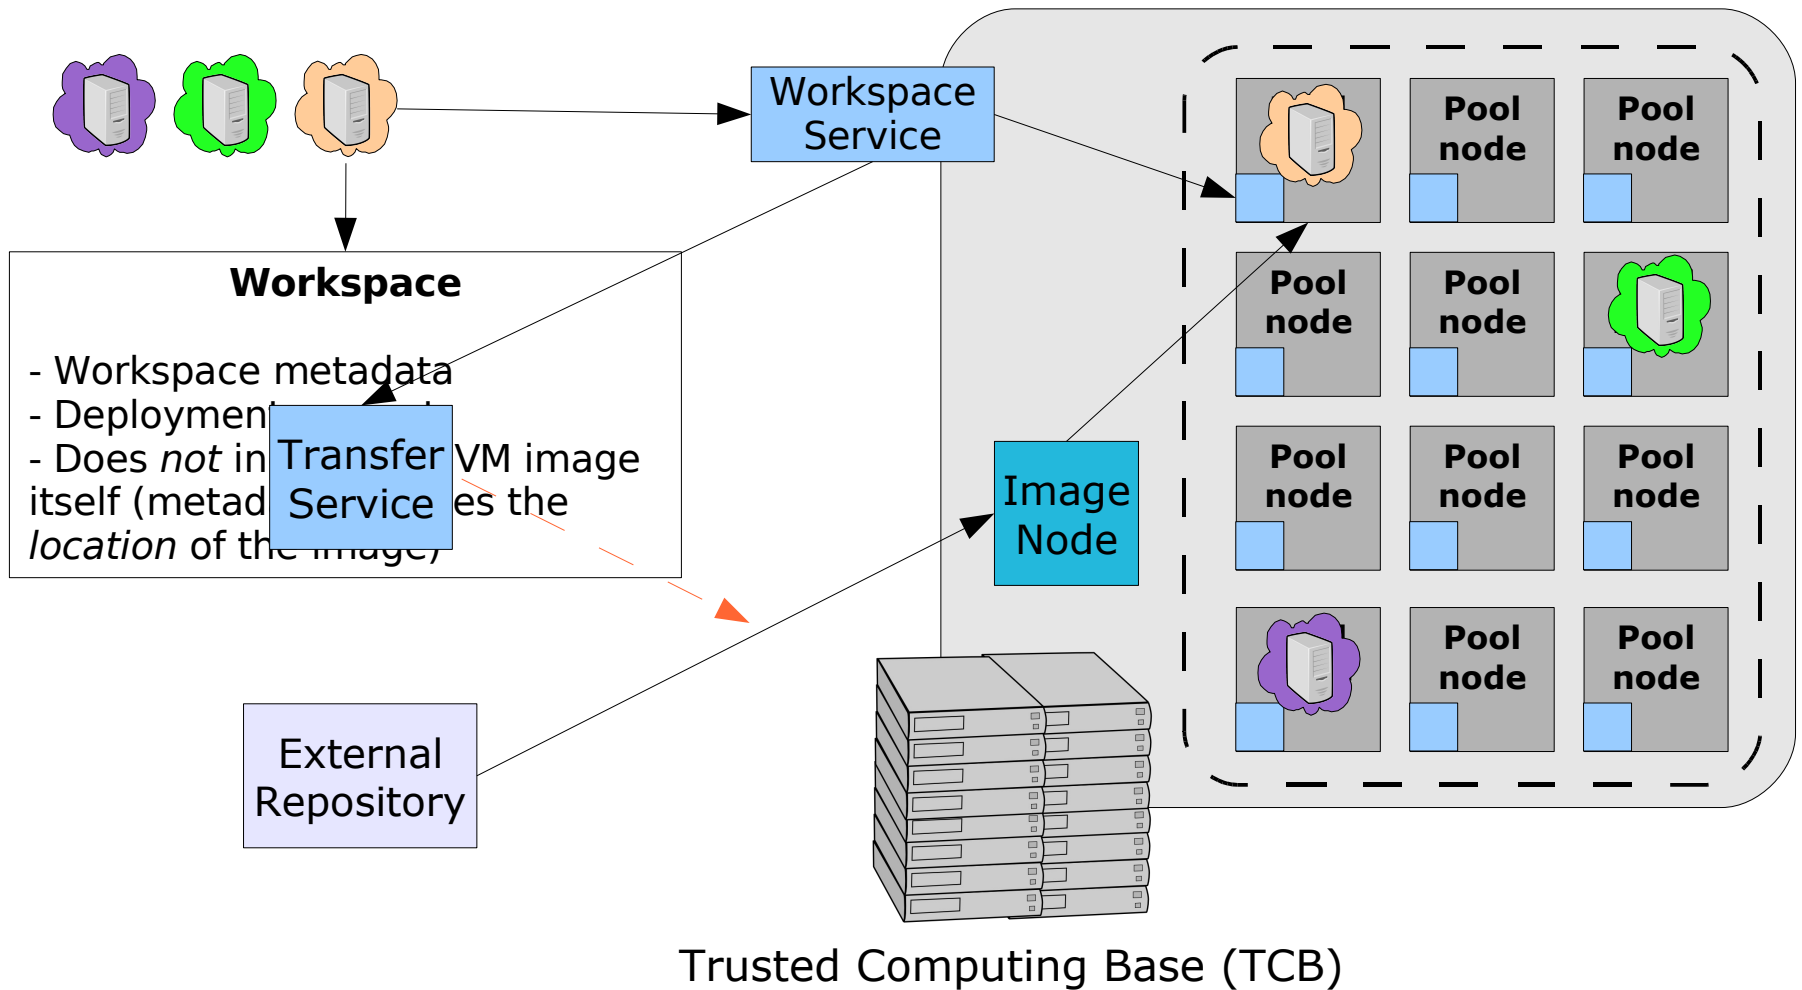
Each node must have a VMM (Xen) installed, along with the *workspace backend* (software that manages individual nodes)

VM images are stored in a separate node inside the TCB (they can also be staged in from an external site)



Trusted Computing Base (TCB)

GT4 Workspace Service

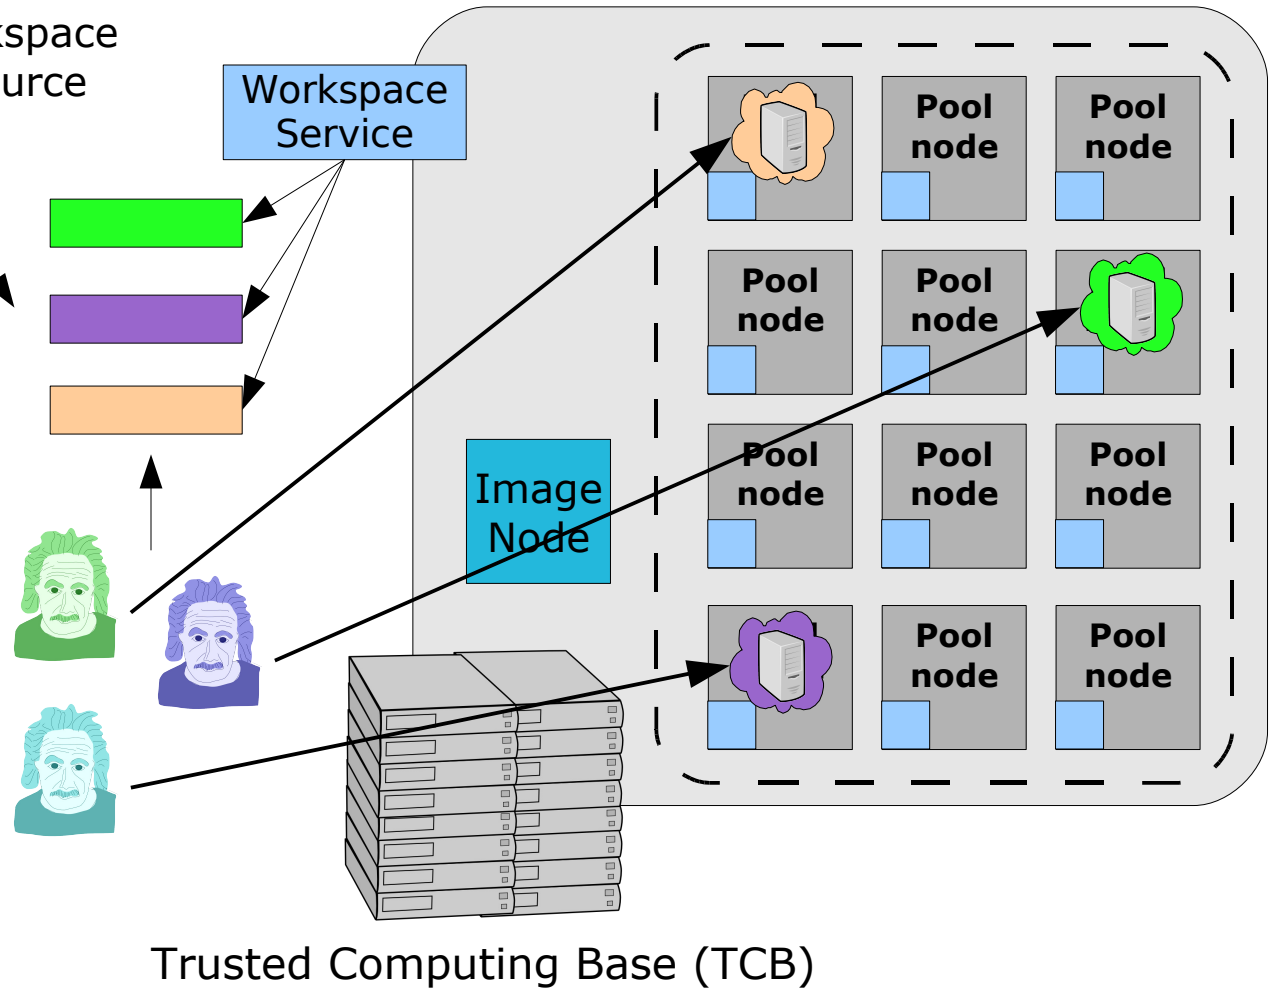


GT4 Workspace Service

The workspace service publishes information on each workspace as standard WSRF Resource Properties.

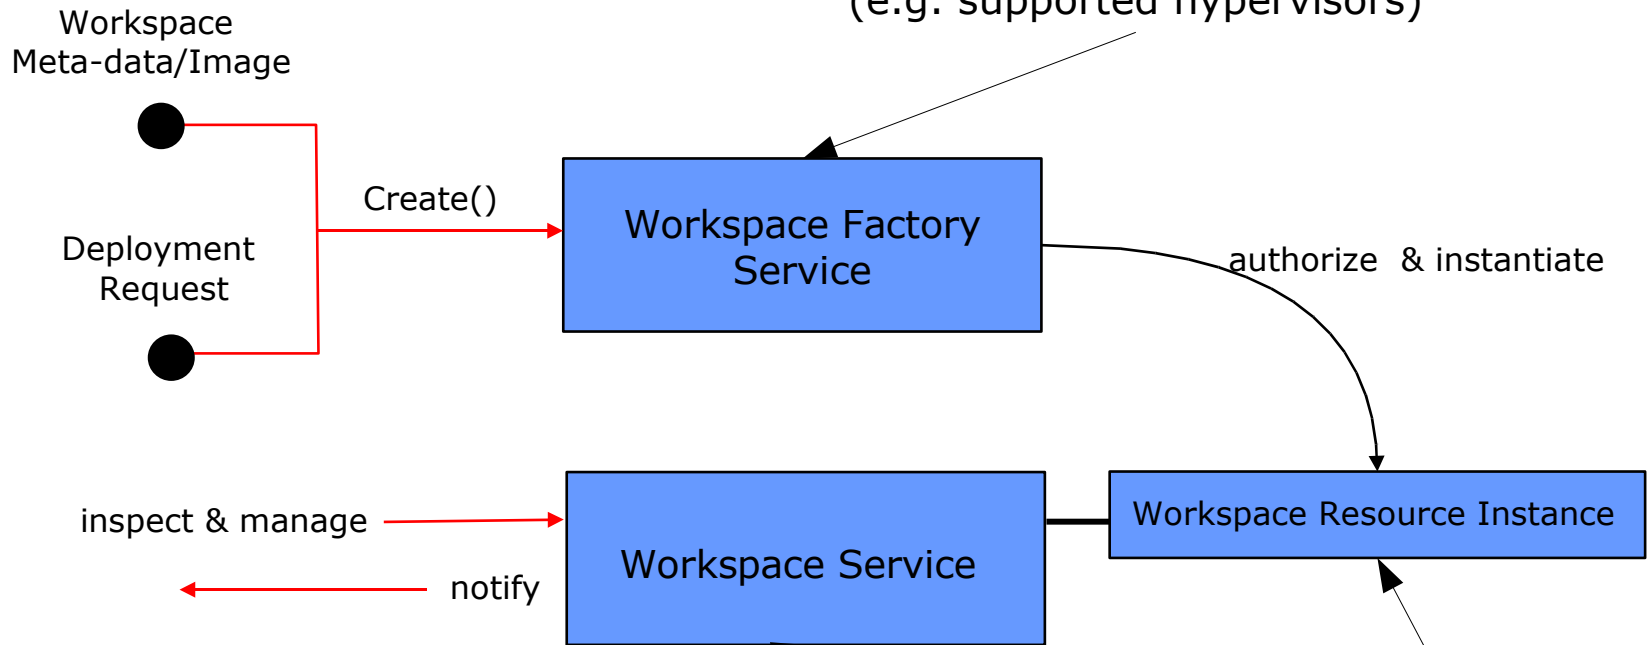
Users can query those properties to find out information about their workspace (e.g. what IP the workspace was bound to)

Users can interact directly with their workspaces the same way they would with a physical machine.



VWS Interfaces

Handles creation of workspaces.
Also publishes global information
as Resource Properties
(e.g. supported hypervisors)



Handles management of
each created workspace
(start, stop, pause, migrate,
inspecting VW state, ...)

Resource Properties publish the
assigned resource allocation, how
VW was bound to metadata (e.g.
IP address), duration, and state

GT4 Workspace Service

- VWS is an *incubator project* in dev.globus
- dev.globus
 - ◆ The Globus Toolkit recently shifted to a governance model and infrastructure similar to Apache Jakarta.
 - ◆ Open development model -> Easier for non-Globus developers to contribute and become committers.
 - ◆ All software is released under an Apache License 2.0
 - ◆ <http://dev.globus.org/>
- Workspace Service: <http://workspace.globus.org/>

Latest Release

- TP1.2 was released 09/14
- Lots of improvements compared to TP1.1
- Highlights
 - ◆ Implements the pool model
 - Including basic resource management functions
 - ◆ More functionality and deployment options
 - Broader range of operations
 - IP address assignment
 - Certificates
 - ◆ More reliable and manageable internal structure
 - ◆ Staging
 - ◆ State machine (keeps track of state changes)
 - ◆ Better installation tools

Who's using VWS?

- Open Science Grid
 - ◆ Edge Services
- Intel
 - ◆ GPE (Grid Programming Environment)
 - ◆ Includes VWS as part of a grid stack (to set up an execution environment for jobs)
- Exploring collaborations with...
 - ◆ rPath (rBuilder)
 - ◆ GridWay

Index

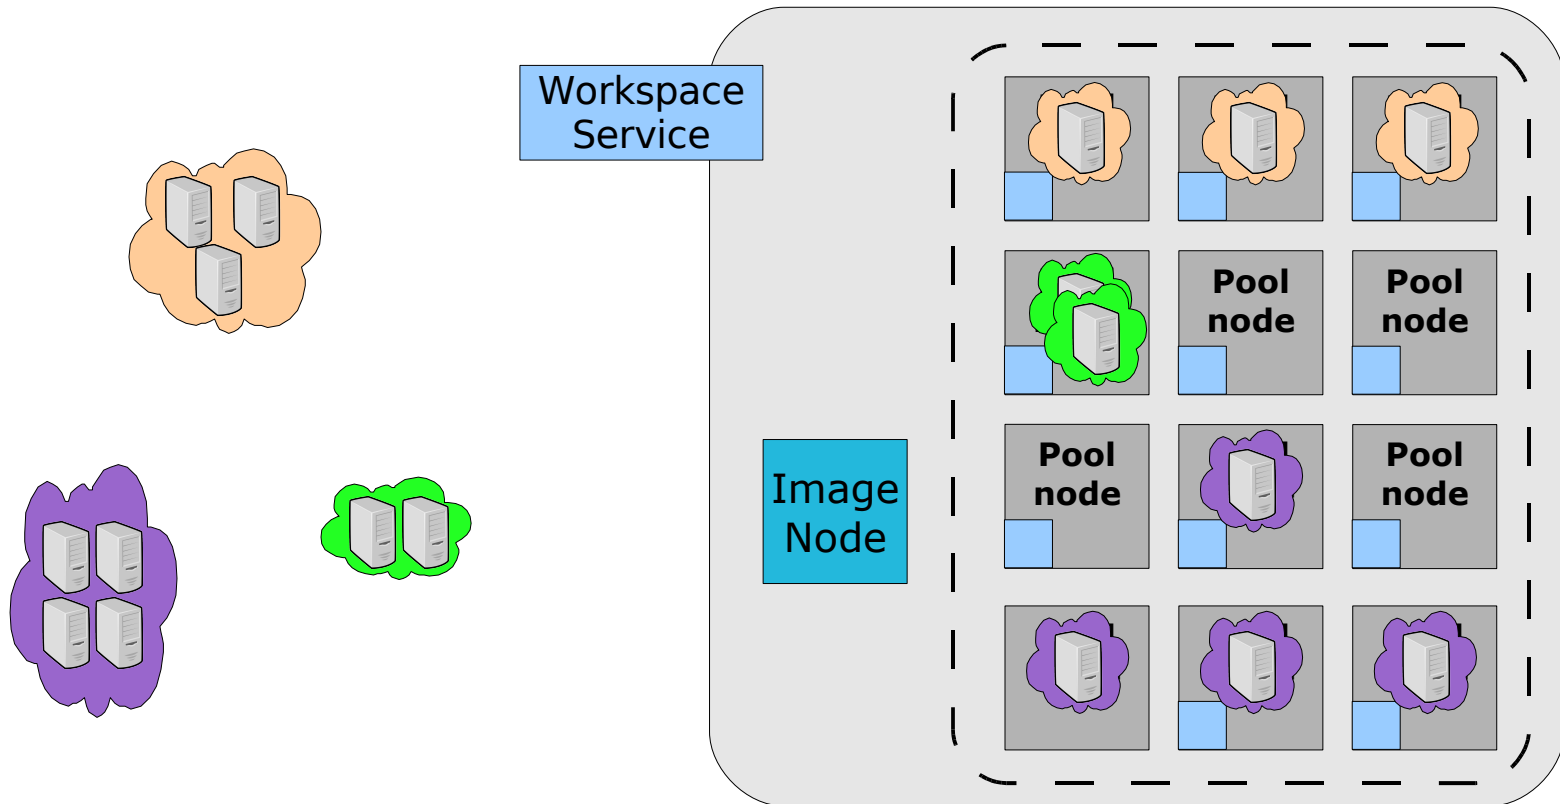
- Virtual Workspaces
 - ◆ What is a workspace?
 - ◆ Why are *virtual* workspaces interesting?
- GT4 Workspace Service
 - ◆ Globus Workspace implementation
 - ◆ Current release: TP1.2
- Research

Ongoing Workspace Research

- Virtual clusters
- Fine-grain enforcement
- Security
- Economic brokering
- From a virtual sandbox to a virtual playground

Virtual clusters

- Extending the workspace abstraction to a virtual cluster



Virtual Clusters

- Virtual nodes and their resource allocations are considered in aggregate.
- Dynamically deploying a virtual cluster increases the complexity of the problem.
 - ◆ Scheduling
 - ◆ Managing overhead: clients want the resources they requested (they shouldn't worry about VW overhead), and providers want to minimize overhead.
- So far, promising results.
- Papers
 - ◆ "Virtual Clusters for Grid Communities", CCGrid06
 - ◆ "Overhead Matters: A Model for Virtual Resource Management", VTDC 2006 (in SC06) [submitted]

Negotiating and enforcing fine-grained resource allocations

- “Division of labor” between resource providers and consumers
 - ◆ Independence of configuration
 - ◆ Isolation
 - ◆ Application-independent resource consumption
- Exploring the feasibility of using workspaces to provide it
- Refining resource allocation and its enforcement
 - ◆ Can get tricky: interdependencies between resources
- Paper:
 - ◆ *“Division of Labor: Tools for Growth and Scalability of Grids”, ICSSOC 2006*

Handling Security

- Handling workspaces securely
 - ◆ Identifying workspaces
 - ◆ Admitting workspaces to a site
- Workspace attestation and signing
- Workspace partitions
- Encrypting secret partitions
 - ◆ An alternative to signing
- Workspace probing

SC05 Poster: “Making your workspace secure: establishing trust with VMs in the Grid”

Economic Brokering

- Two prevalent approaches:
 - ◆ Client auctions
 - Quick turnaround on availability but hard to ensure reliable availability
 - ◆ Provider auctions
 - Difficult from availability perspective
- Investigating a hybrid approach

SC06 Poster: “To bid or not to bid: a hybrid approach to economic brokering”

Virtual Playgrounds

- Virtual Networks
- Deploying virtual clusters across sites.
- Deploying *virtual grids*.

Questions?

Kate Keahey¹, Tim Freeman¹, Borja Sotomayor²

¹Mathematics and Computer Science Division – Argonne National Laboratory

²Department of Computer Science – University of Chicago

{keahey,tfreeman}@mcs.anl.gov

borja@cs.uchicago.edu