

Virtual Workspace Appliances

Tim Freeman, Kate Keahey

Supercomputing 2006, Tampa, FL
tfreeman@mcs.anl.gov

Required Environments

- Diverse client environment requirements
 - Library versions
 - Application versions
 - Custom applications (with possibly complex installs)
 - OS type, version, modules

vs.

- Provider constraints
 - Security policies
 - Administrator time

What is a Workspace?

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. [[Quality of Life]]

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. [[Quality of Service]]

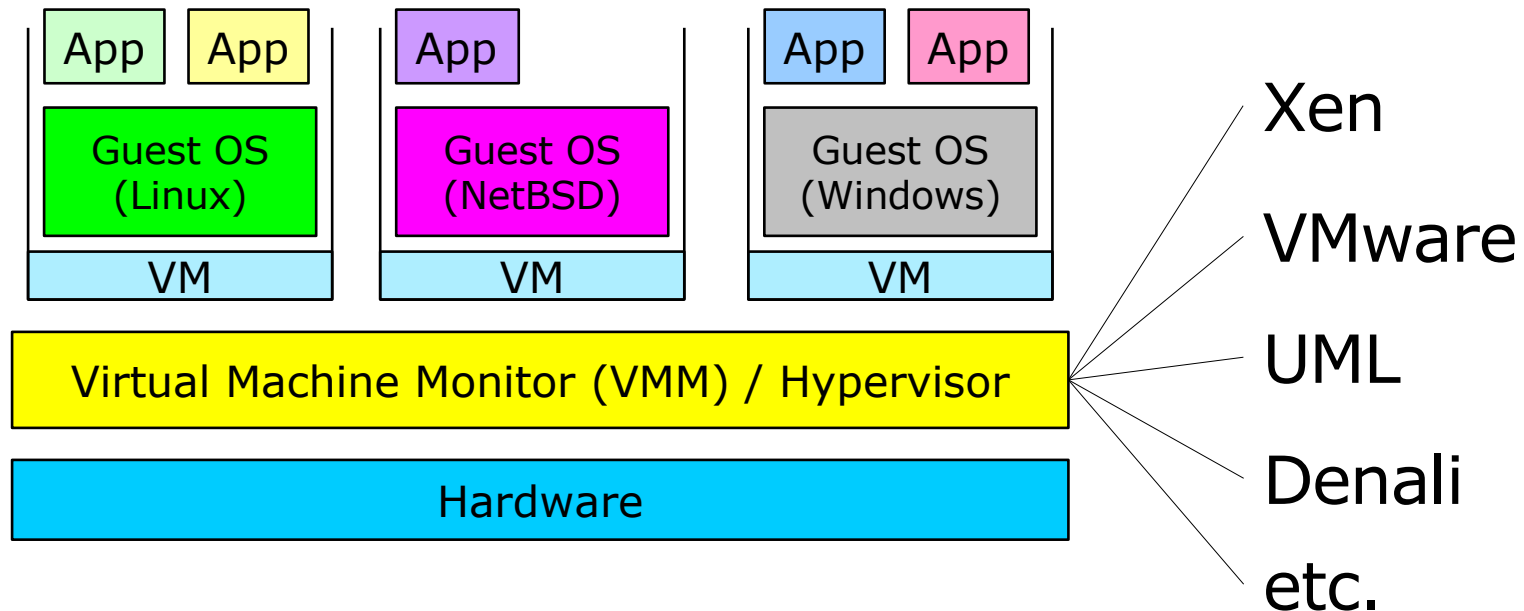
Existing implementations either don't provide both, or...

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.



Virtual Machines

- **Isolation**
 - Security enforced at hypervisor layer
 - Fine grain (alterable) resource allocations
- Flexible **control** and accounting for site
- Customization: **any software** (including legacy)
- Client can have administrator privileges
- Site software requirements reduced to VMM
- **Performance** overhead is becoming acceptable
 - Currently support Xen (studies: *within 5%*)
 - Experimented with VMware in the past


Use Cases

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


GT4 Workspace Service

- The GT4 Virtual Workspace Service (VWS) is a VM-based workspace implementation.
 - GT4 WSRF frontend
 - Xen-based, but other VMMs can be used (interfaces are generic).
 - <http://workspace.globus.org/>

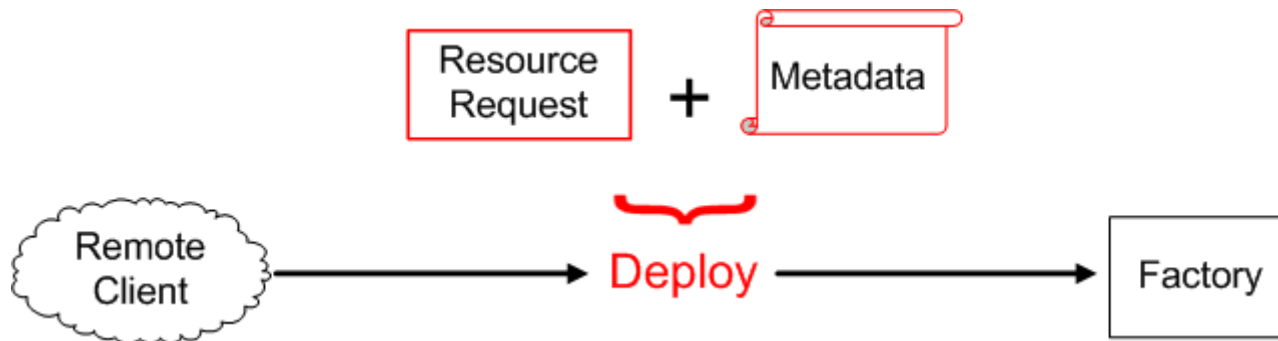
GT4 Workspace Service




Remote Interfaces



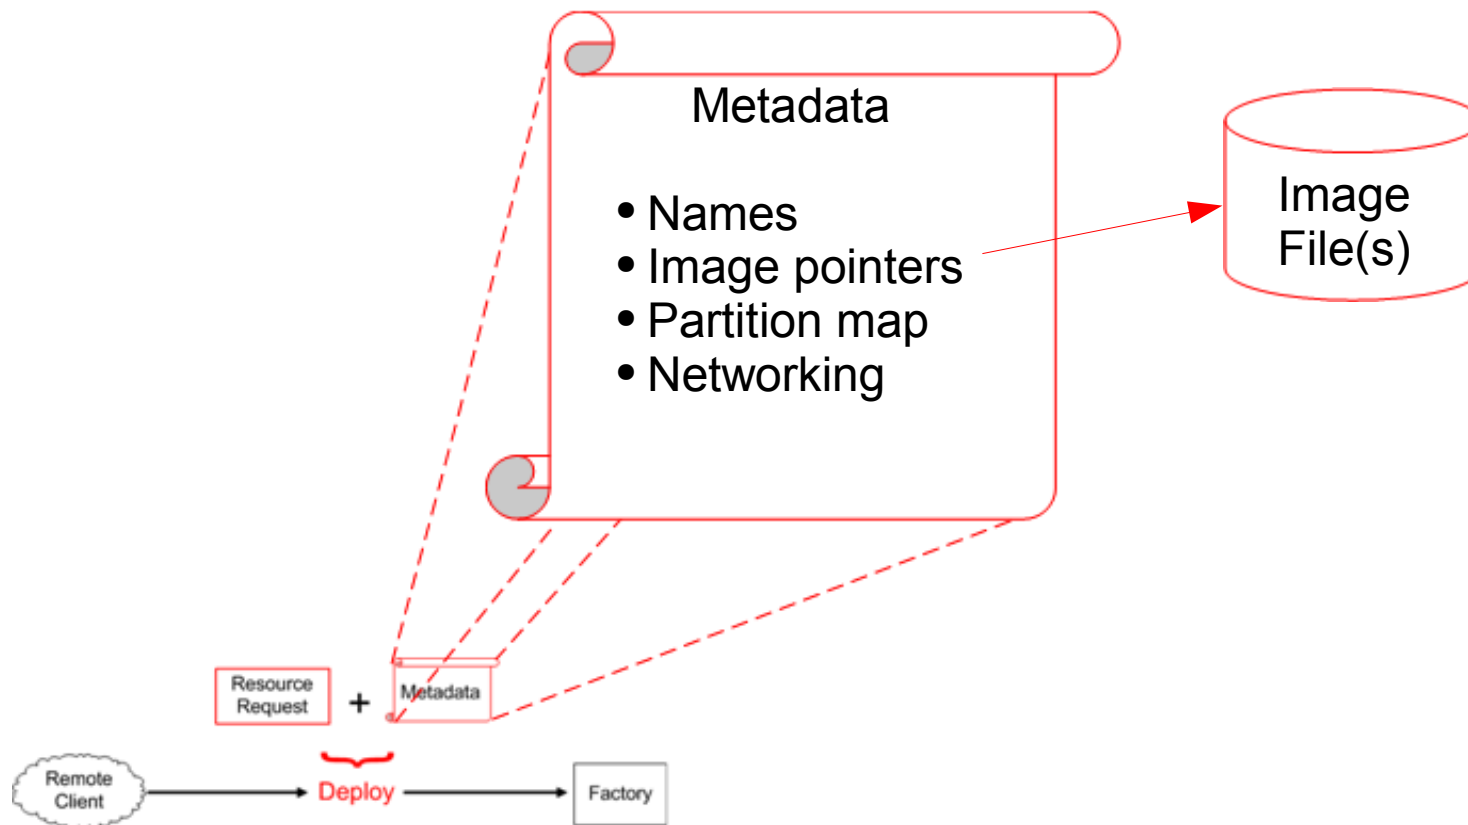
Deployment



Deployment



Deployment



Status

- TP1.2 was released 09/14
- A lot of improvements compared over TP1.1.1
- Highlights
 - ◆ Implements the pool model
 - ◆ More functionality and deployment options
 - ◆ More reliable and manageable internal structure
 - ◆ Staging plugins
 - ◆ State machine (tracks asynchronous state changes and transitions)
 - ◆ Better installation tools
- At least one more release planned by the end of the year, to include (mainly) C client and better IP handling
- To be included in the next VDT release
- VWS is an incubator project in dev.globus

Who's Using VWS?

- **Open Science Grid**
 - Edge Services
 - <http://osg.ivdgl.org/twiki/bin/view/EdgeServices/WebHome>
 - STAR application
 - <http://www.star.bnl.gov/>
 - Demo today
- **Intel**
 - GPE (Grid Programming Environment)
 - Includes VWS as part of a grid stack (to set up an execution environment for jobs)
- **New collaborations**
 - Rpath (rBuilder), image/appliance creation
 - Part of demo today
 - <http://www.rpath.com/rbuilder/>
 - Gridway metascheduler
 - <http://www.gridway.org/>

Workspaces at SC

- **Booth Talks/Demo**
 - Tuesday 3:30pm
 - Wednesday 5:00pm
 - Thursday 10:30am
- **Poster**
 - *To Bid or Not To Bid: A Hybrid Market-Based Resource Allocation Framework.* Elizeu Santos-Neto and Kate Keahey
- **Paper, VTDC06 (Friday)**
 - *Overhead Matters: A Model for Virtual Resource Management.* Borja Sotomayor and Kate Keahey

Overhead Matters

- Two types of overhead
 - *Preparation overhead*: staging VM images to physical nodes, preparing a software environment.
 - *Runtime overhead*: resulting from the management of the VMs themselves
- Some models already deal with it by:
 - Assuming the preparation overhead away (e.g. assuming that all possible VM images are already predeployed in all physical nodes, which is a reasonable assumption in certain scenarios)
 - Having runtime overhead invade the user's allocation. The user must factor in overhead when requesting resources.

STAR Application

 www.star.bnl.gov

- Time consuming configuration, specific library needs
 - Doug Olson (LBL):
 - *“tends to push the boundaries on what will actually compile”*
 - *“using rarely used features of the language”*
 - *“even just validating a new platform is a big job even when it all compiles.”*
 - *The STAR offline analysis software is about 1.3M lines of code, 2/3 C++, a bit under 1/3 Fortran, and a bit of C.*
 - *(line counts generated using David A. Wheeler's 'SLOCCount')*

Virtual Workspaces: <http://workspace.globus.org>

rBuilder STAR Appliance



- rBuilder to the rescue
 - rPath was founded by ex-RedHat luminaries
 - Software Appliances (and not just VMs)
 - <http://www.rpath.com/rbuilder/>
- Stu Gott, Ken Vandine, and Marty Wesley worked with OSG and Doug Olson to produce a STAR appliance with rBuilder
- Running on the Computation Institute's Teraport Cluster at The University of Chicago. Many thanks to Rob Gardner, Greg Cross, Borja Sotomayor, and the Computation Institute.

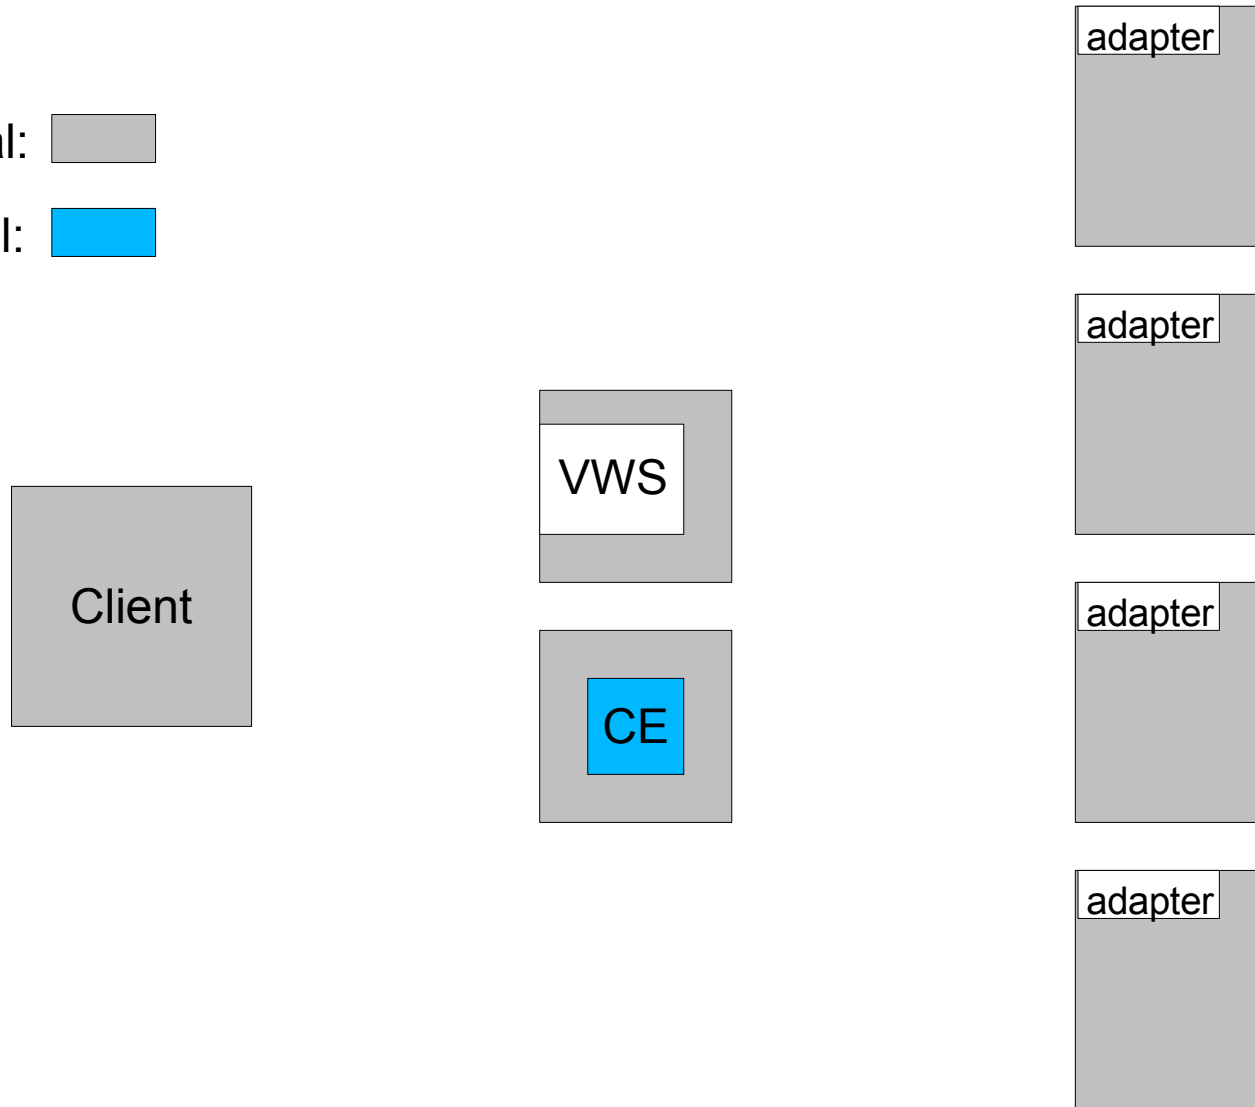


Let's see appliances in action ...



Physical: 


Virtual: 





Physical: 

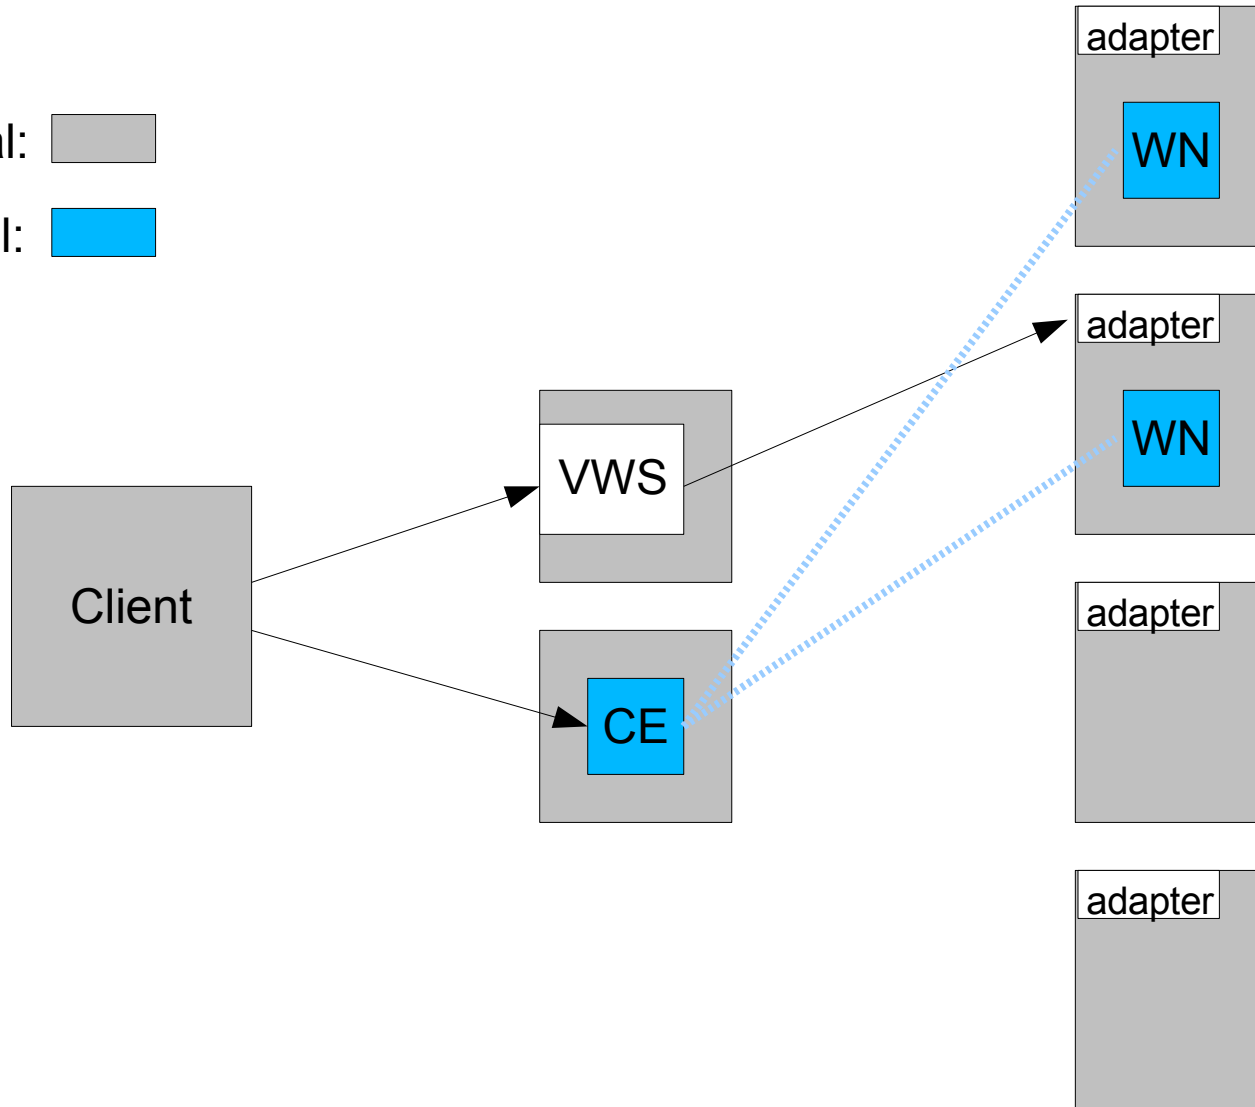
Virtual: 





Physical: 


Virtual: 








Physical: 

Virtual: 

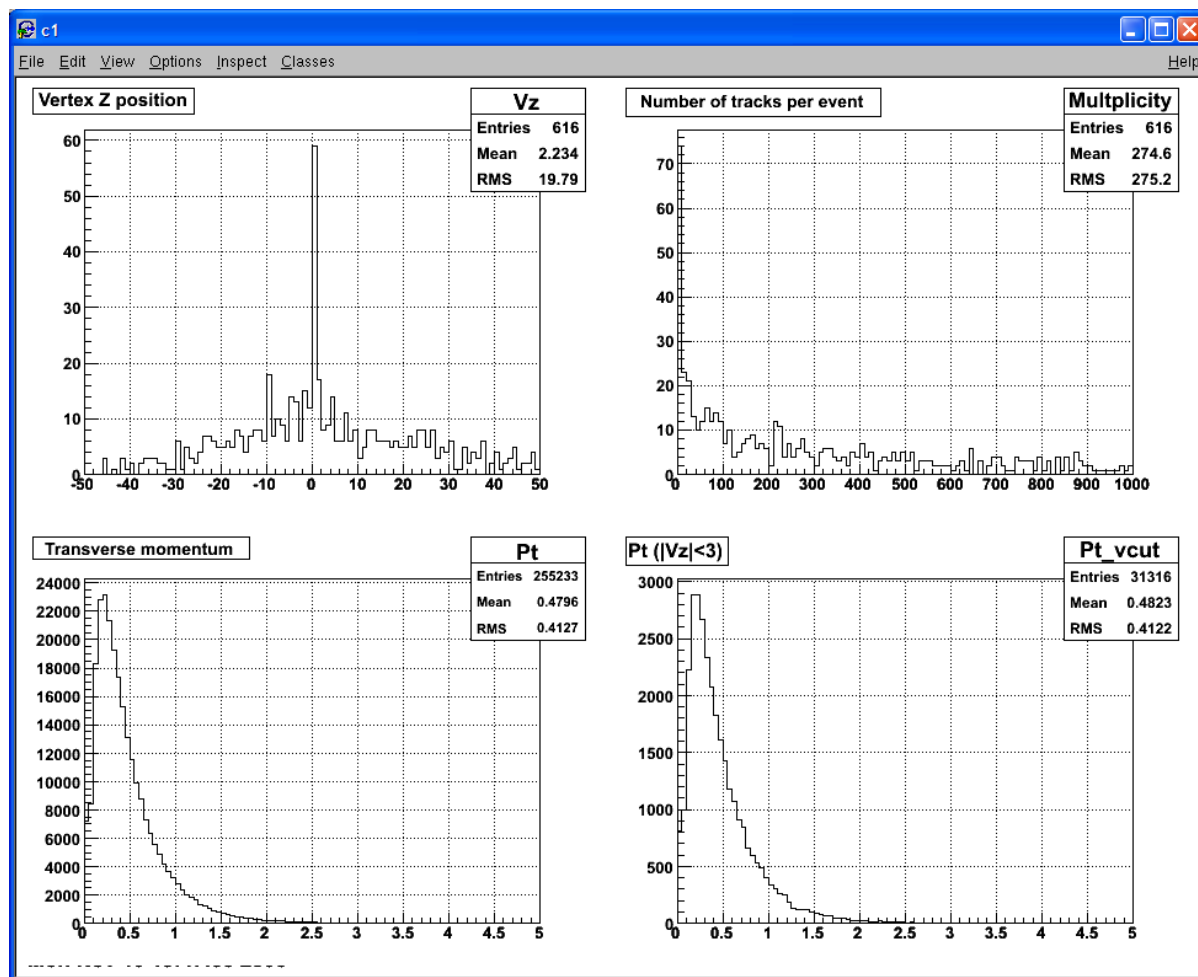


Physical: 
Virtual: 



Plots from analyzing STAR data on Teraport Workspaces

 www.star.bnl.gov



Virtual Workspaces: <http://workspace.globus.org>

Thankyou

<http://workspace.globus.org>

- » Code
- » Documentation
- » Support (mailing lists)
- » Publications