Pools of Virtual Boxes

A Year Later

Craig A. Struble, Ph.D.
craig.struble@marquette.edu

Department of Mathematics, Statistics, and Computer Science
Marquette University, Milwaukee, WI 53201-1881

April 14, 2010

Campus Grids

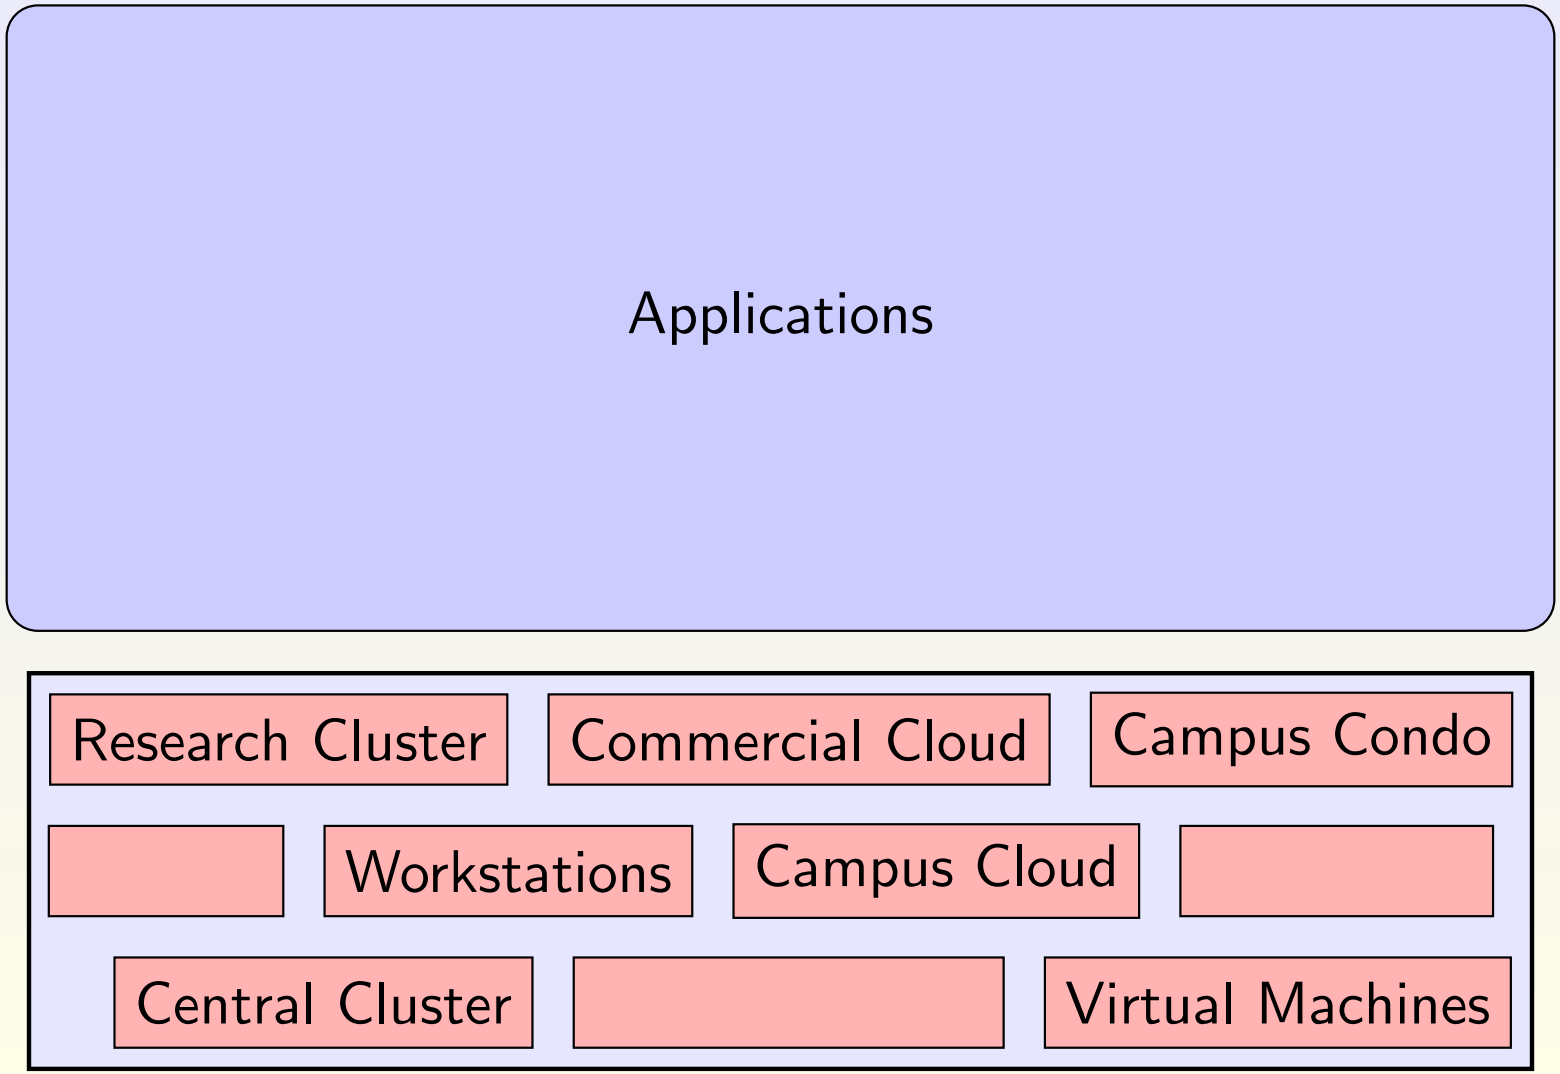

Why Use VMs on Lab Machines?

- Universities of all types have lab machines
- Lab machines are frequently idle
- No additional space or power infrastructure needed
- Decouple lab and research software requirements
- Provide a layer of safety

Overview

Pools of Virtual Boxes

- A *virtual machine* campus grid brick
- Oracle VirtualBox® on Windows host
- CentOS or Fedora Core Linux configured with Condor
- Inspired by coLinux Condor distribution
  - University of Nebraska–Lincoln, University of Oklahoma
- Other possible solutions
  - VMWare (Purdue),
  - Grid Appliance (University of Florida)

Why VirtualBox®?

- Free for academia (both PUEL and GPL versions)
- Excellent command line tools and APIs
- Portable (Windows, OS X, various Linuxes)
- Both hardware and software virtualization

What's New?

- We have a name
- `http://poolsofvirtualb.sourceforge.net`
- Multi-core support (VirtualBox® 3.0.x)
- NAT networking support
- Better integration of packages
- Improved robustness

Host-Guest Architecture

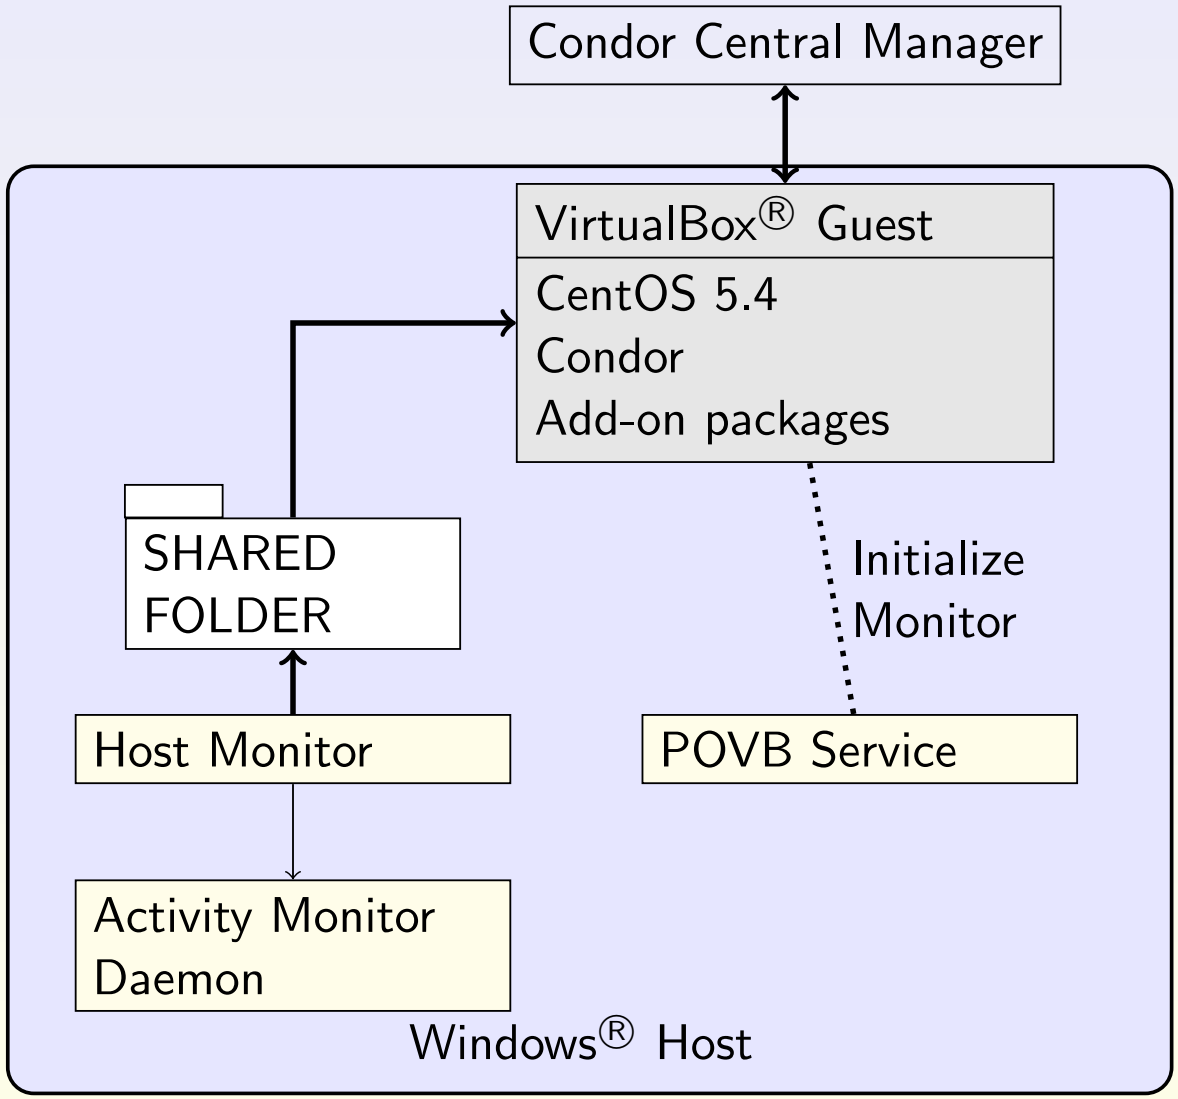

VirtualBox® Configuration

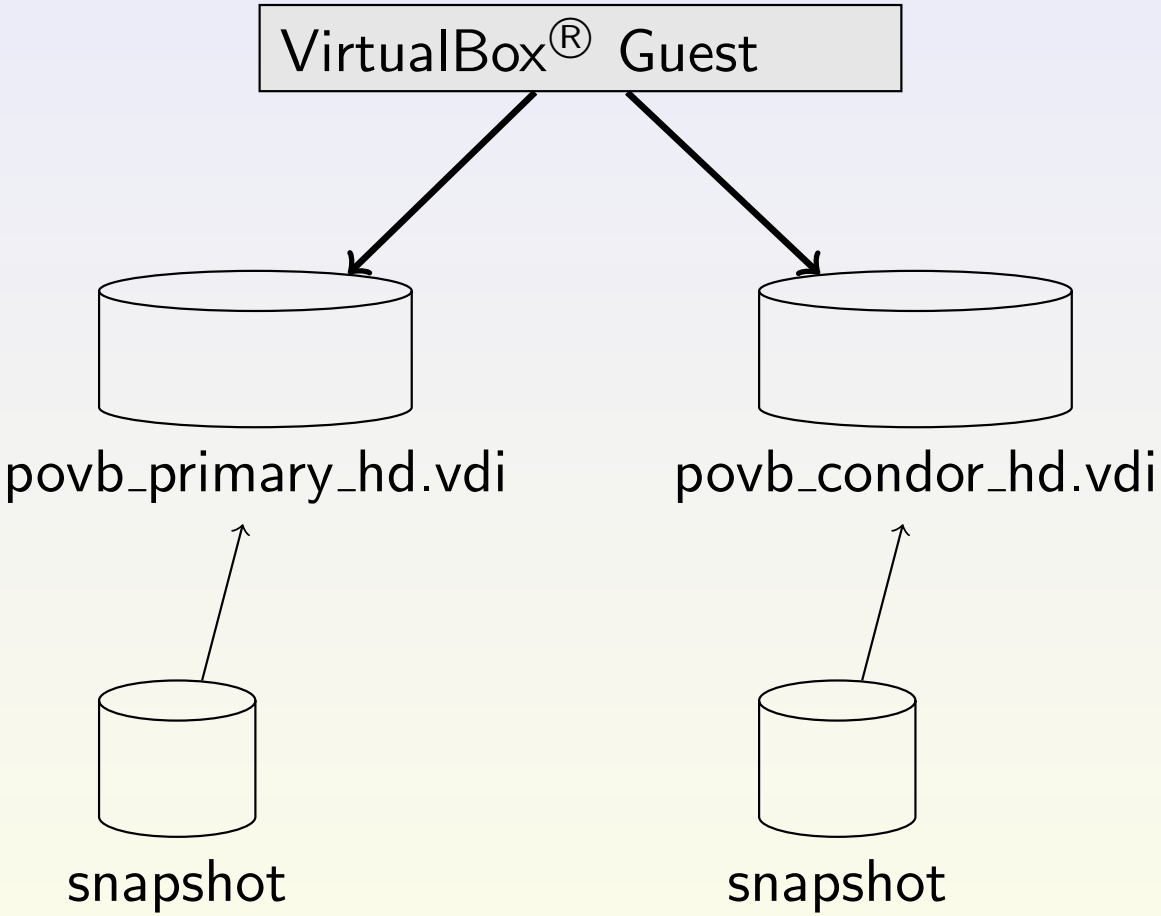

POVB Behavior

- VirtualBox® hard drives are immutable
- Upon reboot
  - Snapshots are removed
  - VM is reconfigured based on host characteristics
  - VM is started
  - VM is monitored for failures
  - Host monitor passes information to VM

Why the POVB Behavior?

- Research support and IT support involve different people
- Common Windows IT practices
  - Install application(s) before semester starts
  - Freeze machine config, reboot to restore
  - Run Windows update nightly, might reboot host
  - Generally want/need hands off solutions

Networking

- Bridged
  - Assigned addresses via DHCP
  - Problem is that it doubles address space
- NAT
  - Made possible by Condor CCB
  - Assigned private address via VirtualBox®
  - Submit nodes must be accessible by POVB machines

Condor Configuration

- Configured to preempt jobs based on host load, user activity
- Each slot is assigned a condor account
  - Removes job files upon completion
  - Kills all processes owned by slot user
- Can be configured for any Condor security model
  - host based, SSL, GSI/X.509, Kerberos, ...

Packages

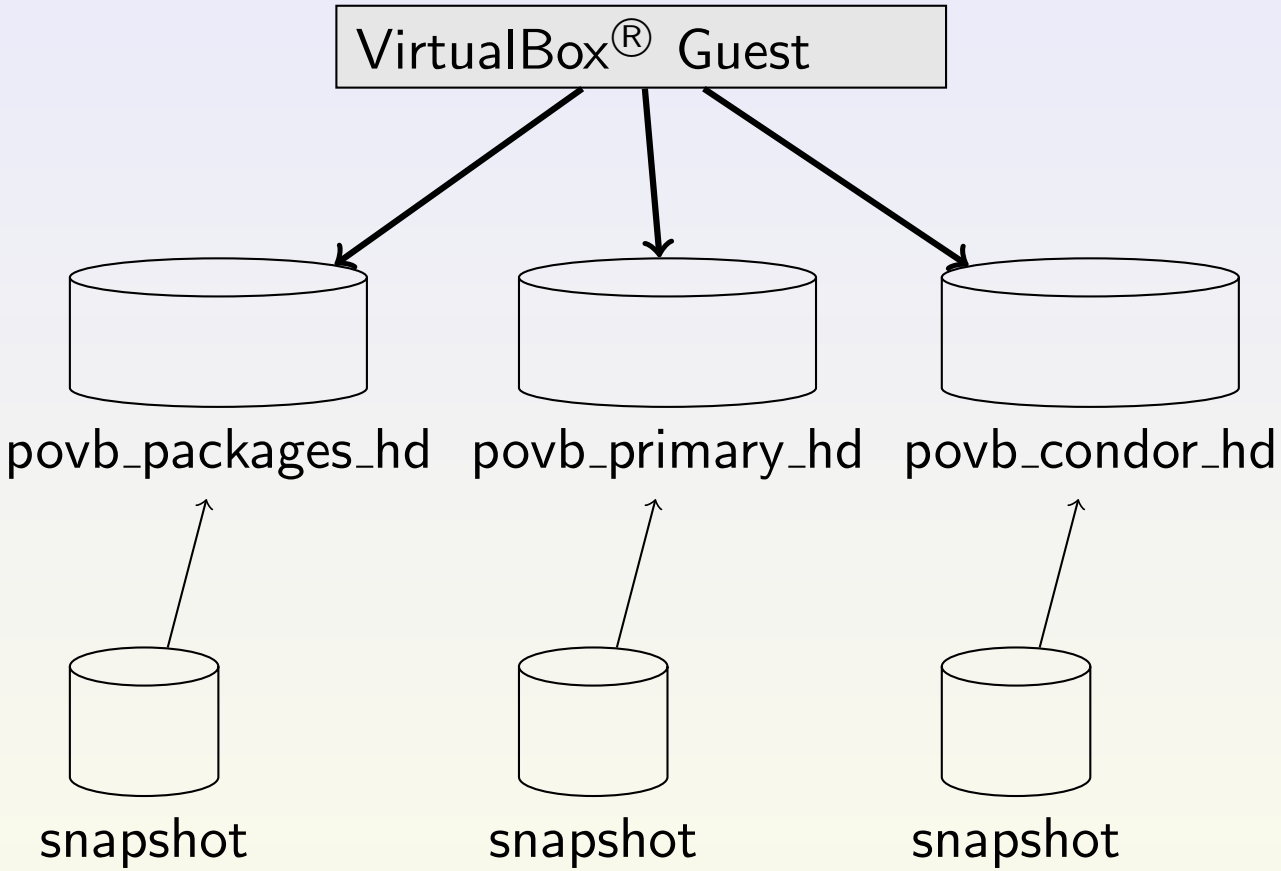
VirtualBox® Guest
povb_packages_hd
povb_primary_hd
povb_condor_hd
snapshot
snapshot
snapshot

Packages

- /packages directory searched in guest environment
- Available packages advertised by Condor
- Condor *job hooks* configure environment for packages

Example

```
+HookKeyword = "POVB_PACKAGES"
+POVB_PACKAGES = "Matlab,Gate"
Requirements = $(Requirements) && HasMatlab =?= TRUE
Requirements = $(Requirements) && HasGate =?= TRUE
```

Typical Deployment

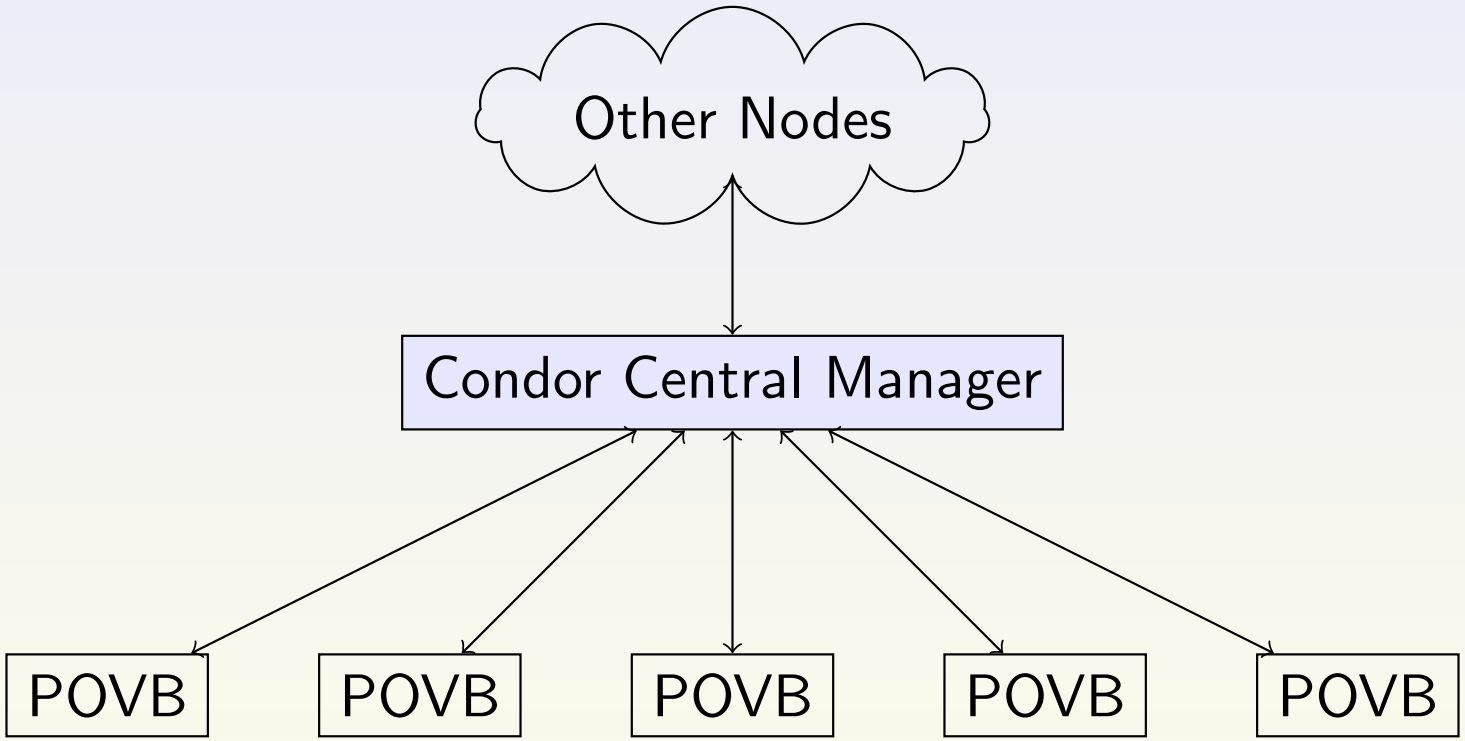

Installation and Distribution

- Installation from a single installer
  - Downloads and installs VirtualBox®
  - Configures VirtualBox®
  - Installs POVB services
- Successfully installed via
  - Network mount
  - Ghost/DeepFreeze
  - Altiris

Bootstrapping

- Relatively "hands free" creation of POVB guests
- Currently CentOS and Fedora Kickstart files
- Sends keystrokes via VirtualBox® command line
- Installs VirtualBox® guest additions and POVB scripts
- Runs on Linux, Mac OS X and (soon) Windows

Marquette University Grid (MUGrid)

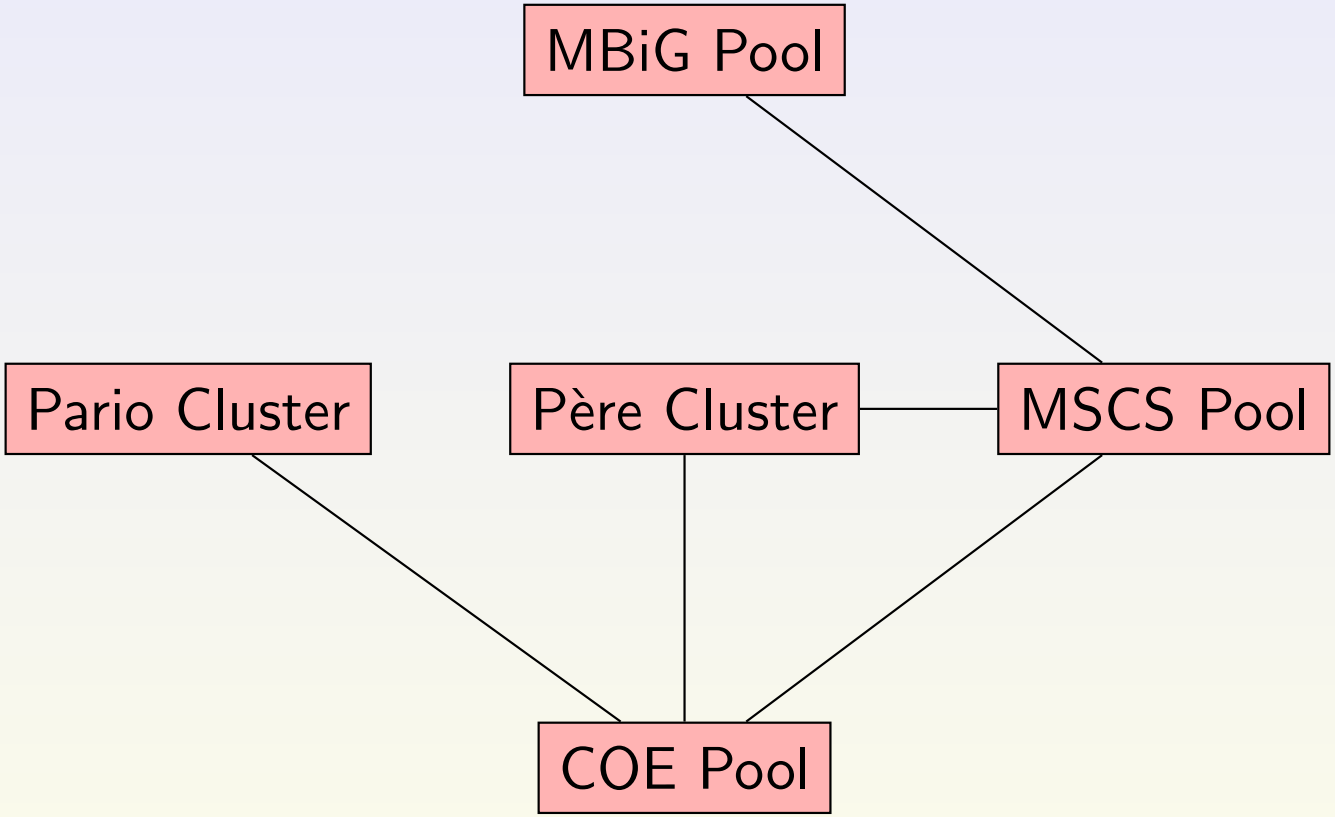

MSCS and COE pools contain POVB instances.

Condor Status (POVB Only)

	Total	Owner	Claimed	Unclaimed	Matched	Preempting	Backfill
INTEL/LINUX	401	107	0	294	0	0	0
Total	401	107	0	294	0	0	0

Varies between 350–420 slots at any given time. Estimated compute capacity is around 0.6 TFLOPS (for free!).

Sample Applications

- Stochastic MRI image reconstruction
- Protein docking for drug discovery
- Genome wide association studies
- Phylogenetic analysis of environmental samples

Summary

- A virtual machine based distribution of Condor execute nodes
- Installs as a single Windows application
- Deployed via standard Windows tools
- Bootstrapping tools for custom distributions

Challenges We've Identified

- Distributing small updates to guest operating system
- Identifying some virtual machine failures
- Detecting incorrect BIOS settings
- Dynamically removing "bad" nodes from the pool
- Remote administration of guest machines

Acknowledgements

- David Herzfeld, POVB mastermind
- Lars Olson, Patrick Blume, Scott Bolte, Karl Stamm, Byron Galbraith, Tim Blattner, Adam Mallen, John Leutke, Kyle Garms, Andrew St. Jean, Andrew Kiel, Gabe Van Eyck, David Hartman, Franco Rupcich, and those I've probably left out
- NSF Award 0923037
- Department of Mathematics, Statistics, and Computer Science, College of Engineering (Marquette University)
- Marquette University HPC Seminar
- SeWHiP
- Condor Team