Autonomic Condor Clouds

David Wolinsky
ACIS P2P Group
University of Florida

UF UNIVERSITY of FLORIDA ACIS

So What's the Big Deal

- Support connectivity across the Internet, in constrained locations, and with clouds
- Simplify packaging
- Minimize Condor configuration
- Reduce downtime
- Let's try to make this easy ....

Discussion Goals

- The High-Level Overview
- Self-Configurable Condor Components
- Virtual Networking
- The User Experience (Time allowing)

High level Overview

- Self-contained VM Appliance
- Configuration through virtual floppy
- Fully distributed, decentralized Virtual Private Network via P2P Overlay
- Job scheduling via Condor
- Customization through Debian and Stacked File system

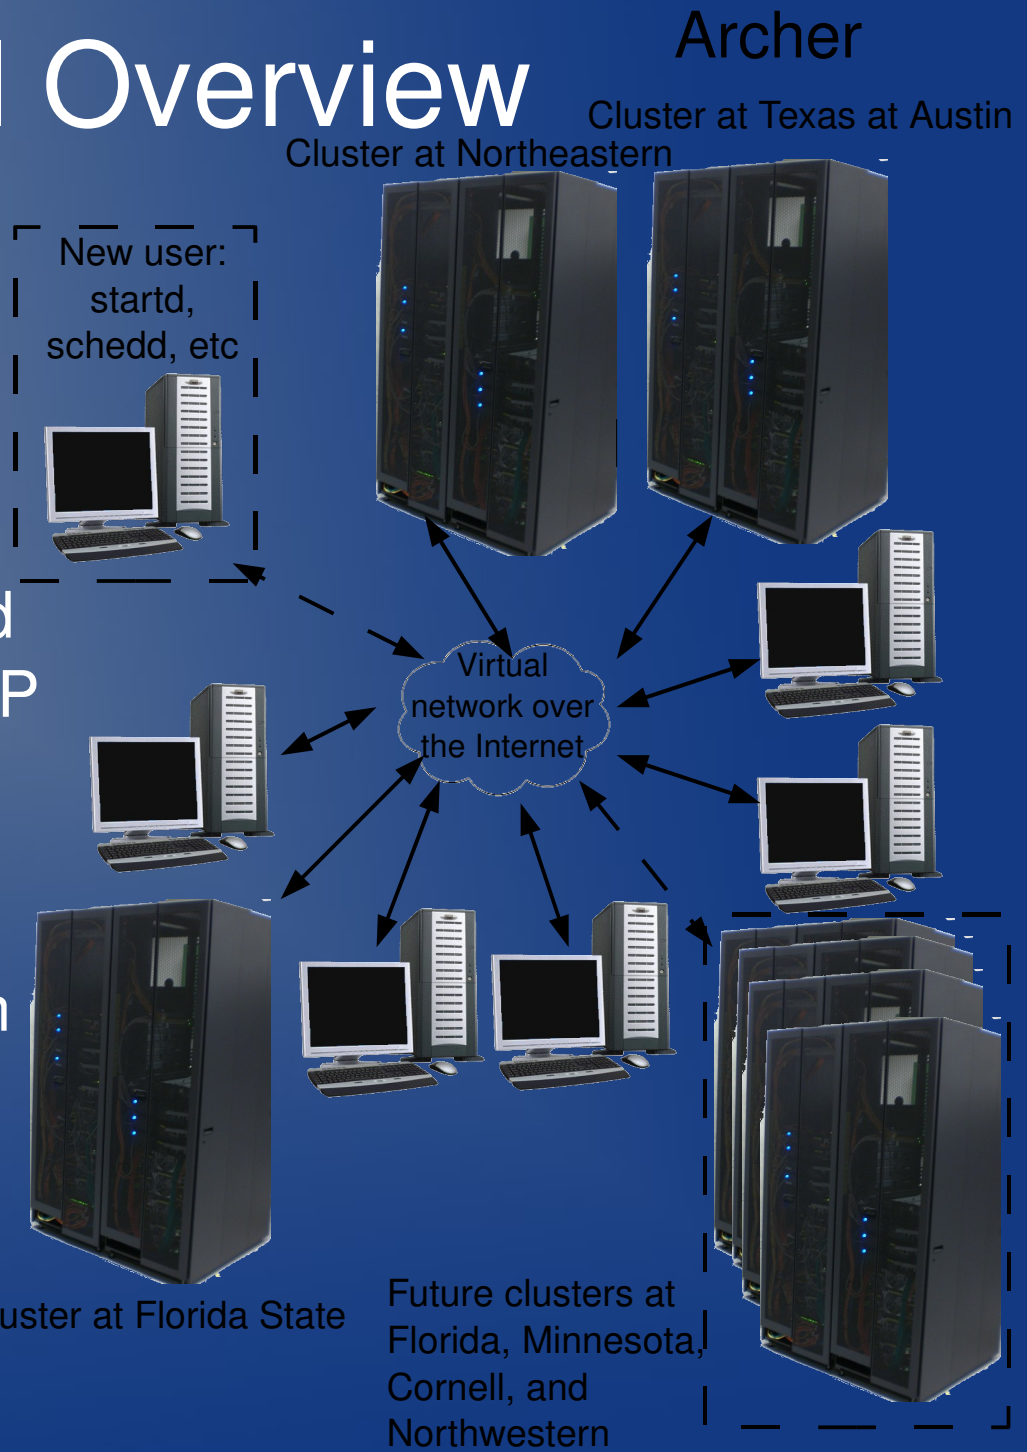

UF UNIVERSITY of FLORIDA ACIS

Everything But Virtual Networking

- Virtual Machines (VMs)
  - Support for VMware, VirtualBox, KVM, and Xen
  - Homogenous environment on heterogenous resources
- Configuration
  - Provided by Virtual floppy, 25 KB download
  - P2P overlay info
  - Users identification certificates
  - Condor machine type
- Customization
  - Based upon Debian 4.0, access to apt repositories
  - Stacked file systems allow users to modify image and share only the changes
- Condor .... :-)

UF UNIVERSITY of FLORIDA ACIS

Configurable Components

- Daemons to start
  - Worker – execution only - startd
  - Client - submission and execution – startd and schedd
  - Master / Server – negotiator and collector
- Which Master / Server to connect or flock to
- User / Group Resource ownership and preemption
- Client can share files via autofs enabled NFS
- Monitoring and binding to a dynamic IP address

More Opportunities

- Support for VM Universe
- Distributed data storage
- Web portal front-end
- Self-policing security system
- Self-sustaining condor cluster
- Portable (decentralized) Condor System Configurer
- Configurable through Social Networks and User Portals

Virtual Networking

- Unified layer 3 (IP) network for all machines
- Cross-site communication without a middleware broker

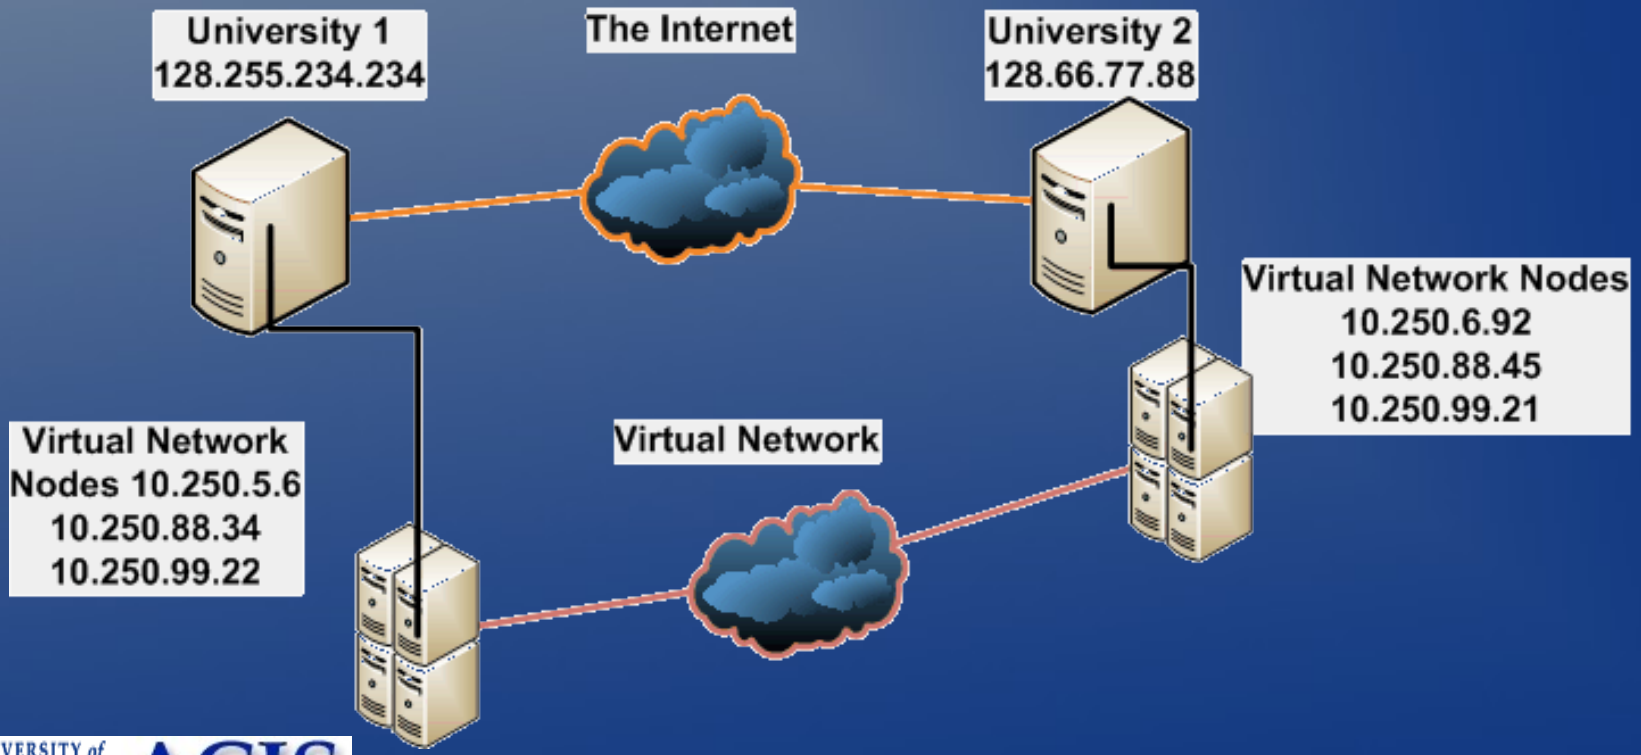

IPOP

- Open Source
- NAT Traversal (STUN)
- Transparent Subnet Gateway
- Structured P2P Network Overlay
- Provides tunneling and direct shortcuts
- Distributed data store (DHT)
- Multiple Virtual Networks Per Overlay

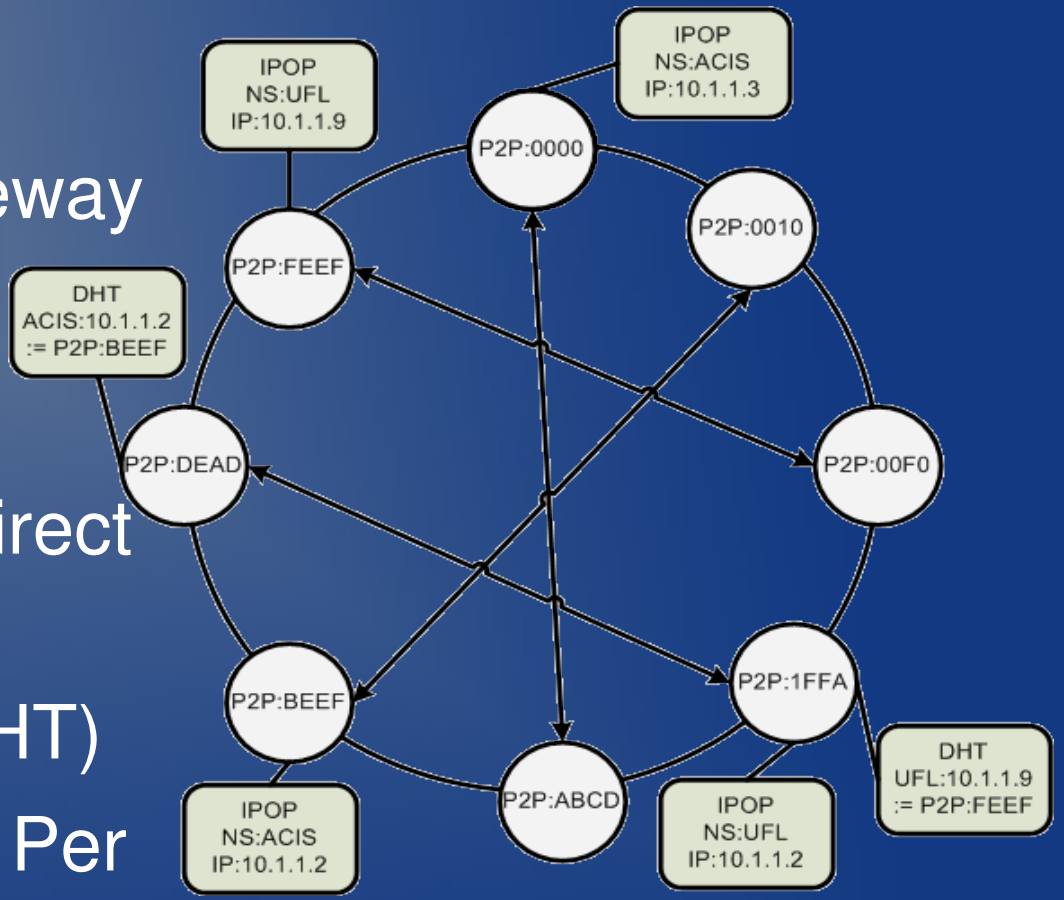

UF UNIVERSITY of FLORIDA ACIS

P2P Overlay

- Several hundred well distributed nodes
- Assist in bootstraping and NAT traversal
- Runs on top of Planetlab

VN Interfaces

- Each machine has VN Interface running locally
- Machine has VN and “Internet” connectivity

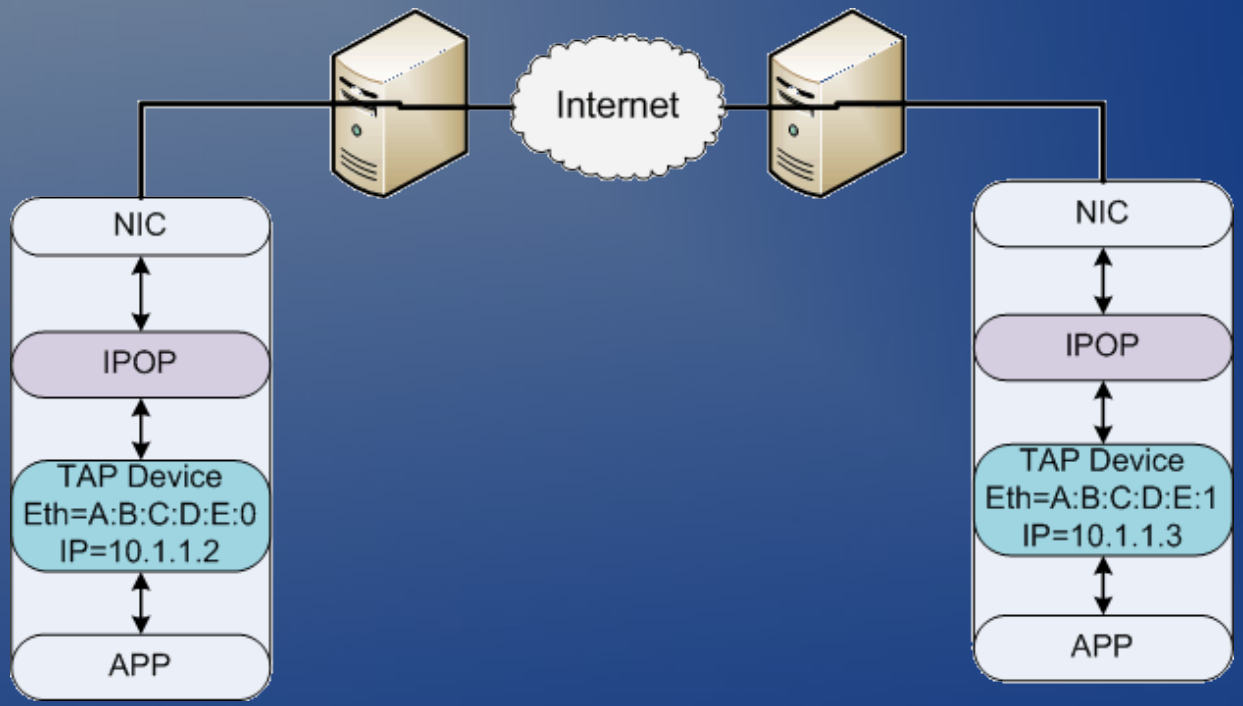

VN Routers

- Single VN instance (Router) for entire cluster
- Limited to no resource configuration
- Isolated VN overhead
- May have “Internet” connectivity if there is an “Internet” router

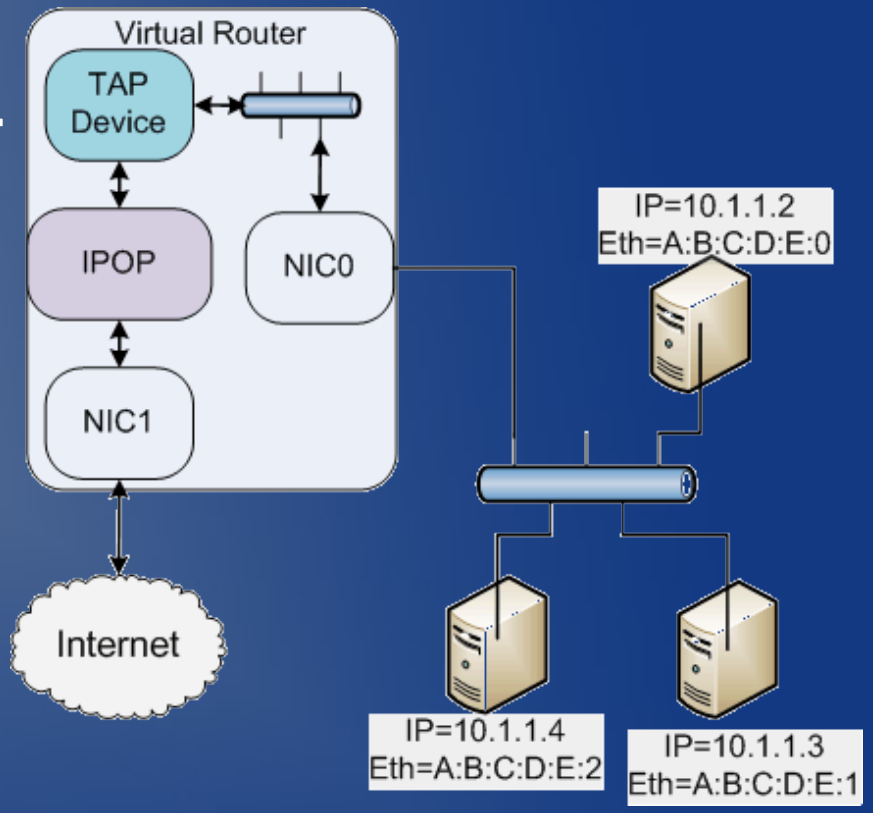

The User Experience

- Time permitting video - http://www.youtube.com/watch?v=1XDvITdayhs
- Otherwise slides –
  - Boot VM and obtain IP Addresses
  - Condor access
  - Direct connectivity (i.e. low ping overheads)

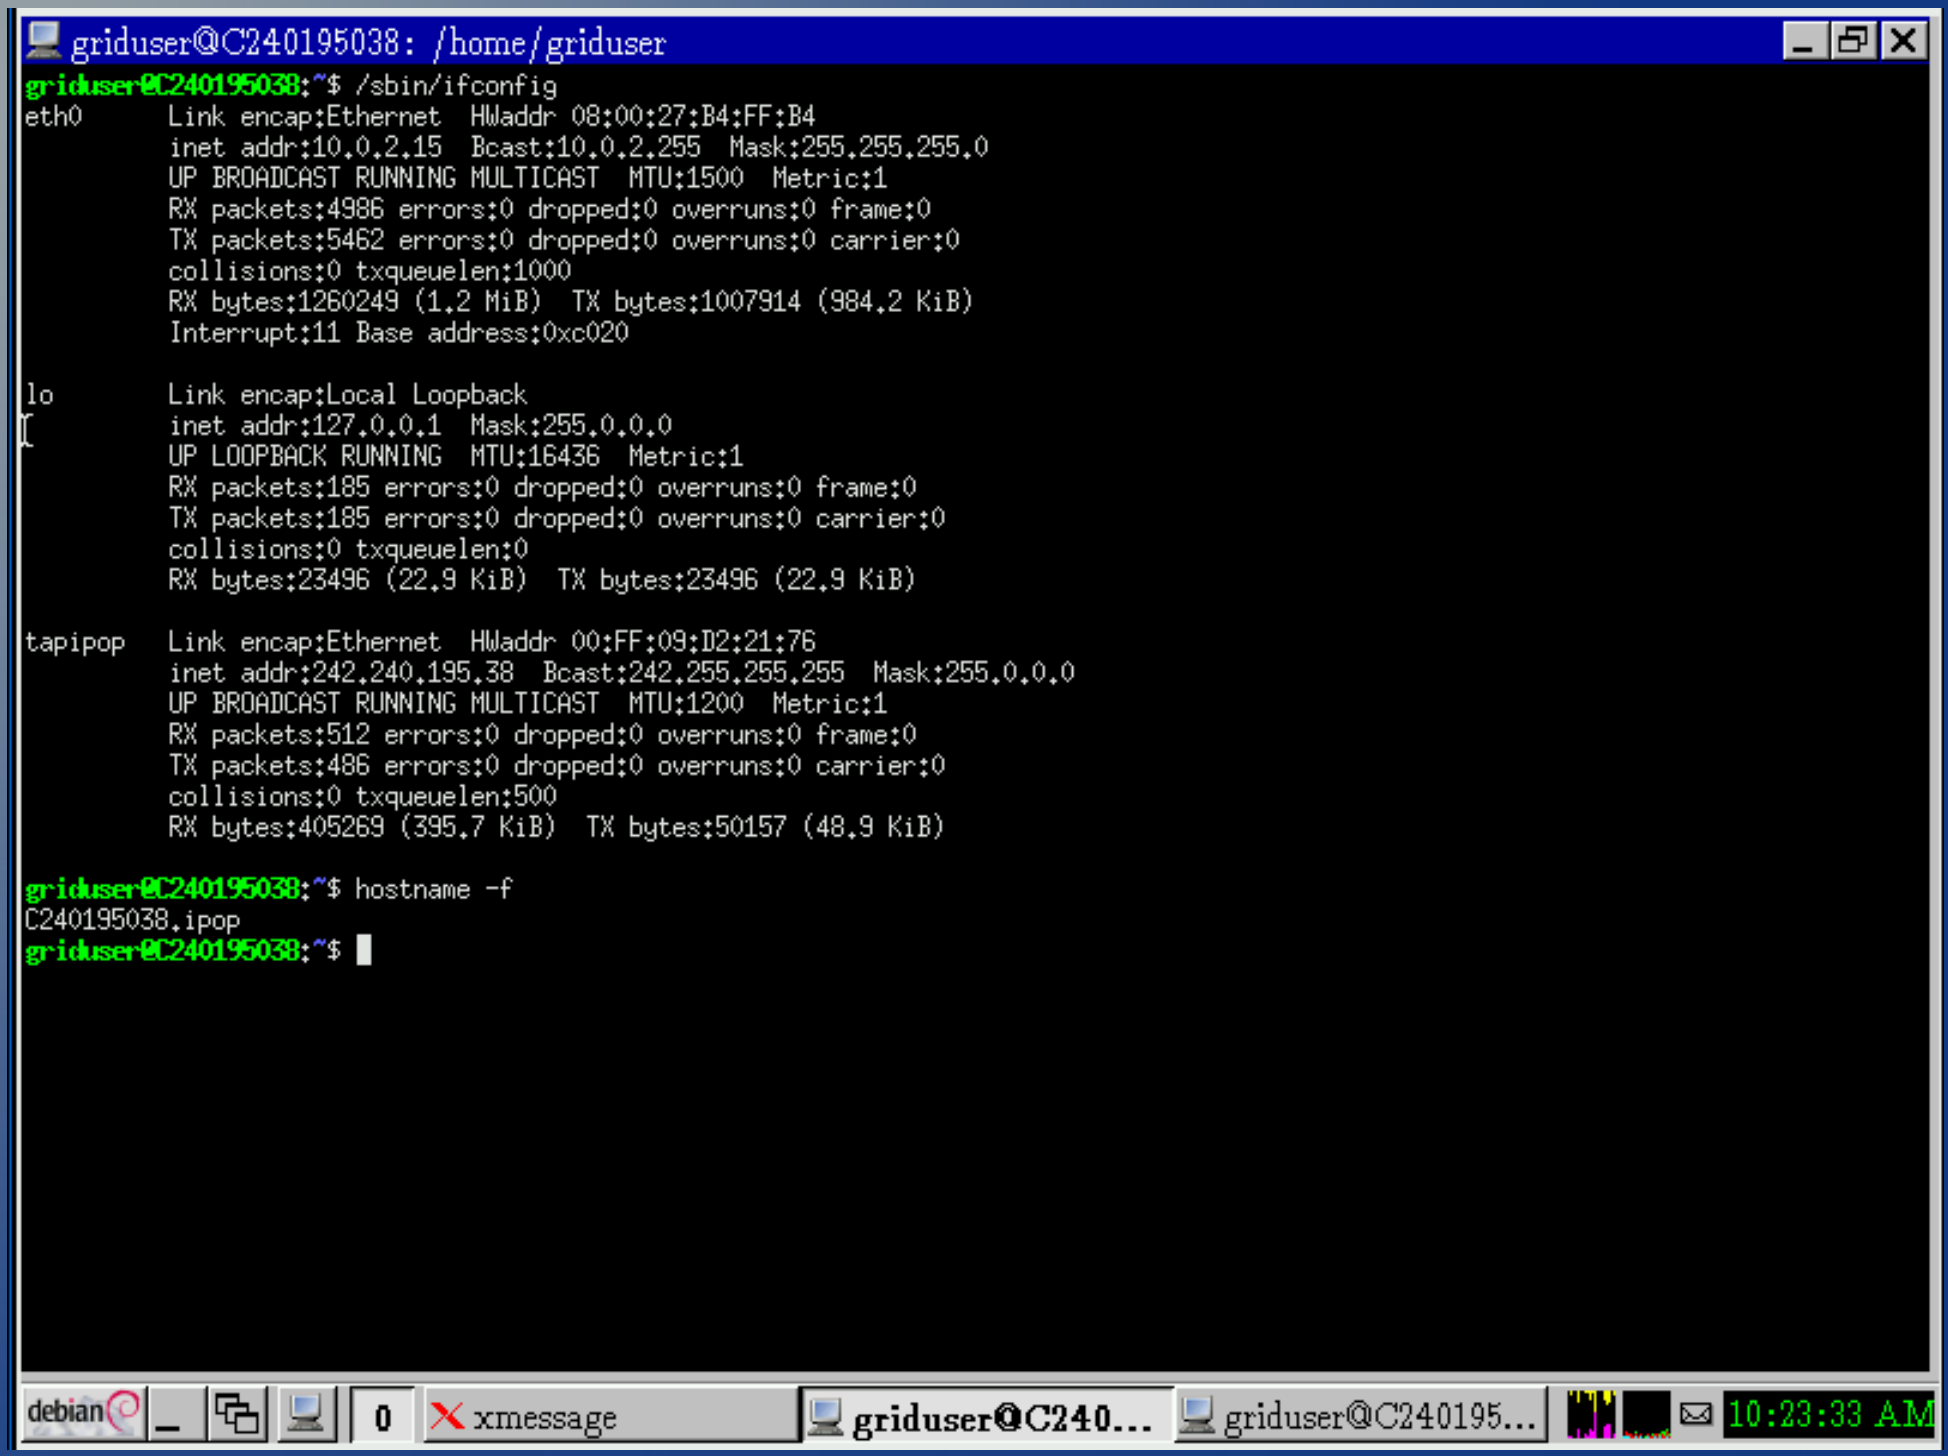

UF UNIVERSITY of FLORIDA ACIS

UF UNIVERSITY of FLORIDA ACIS

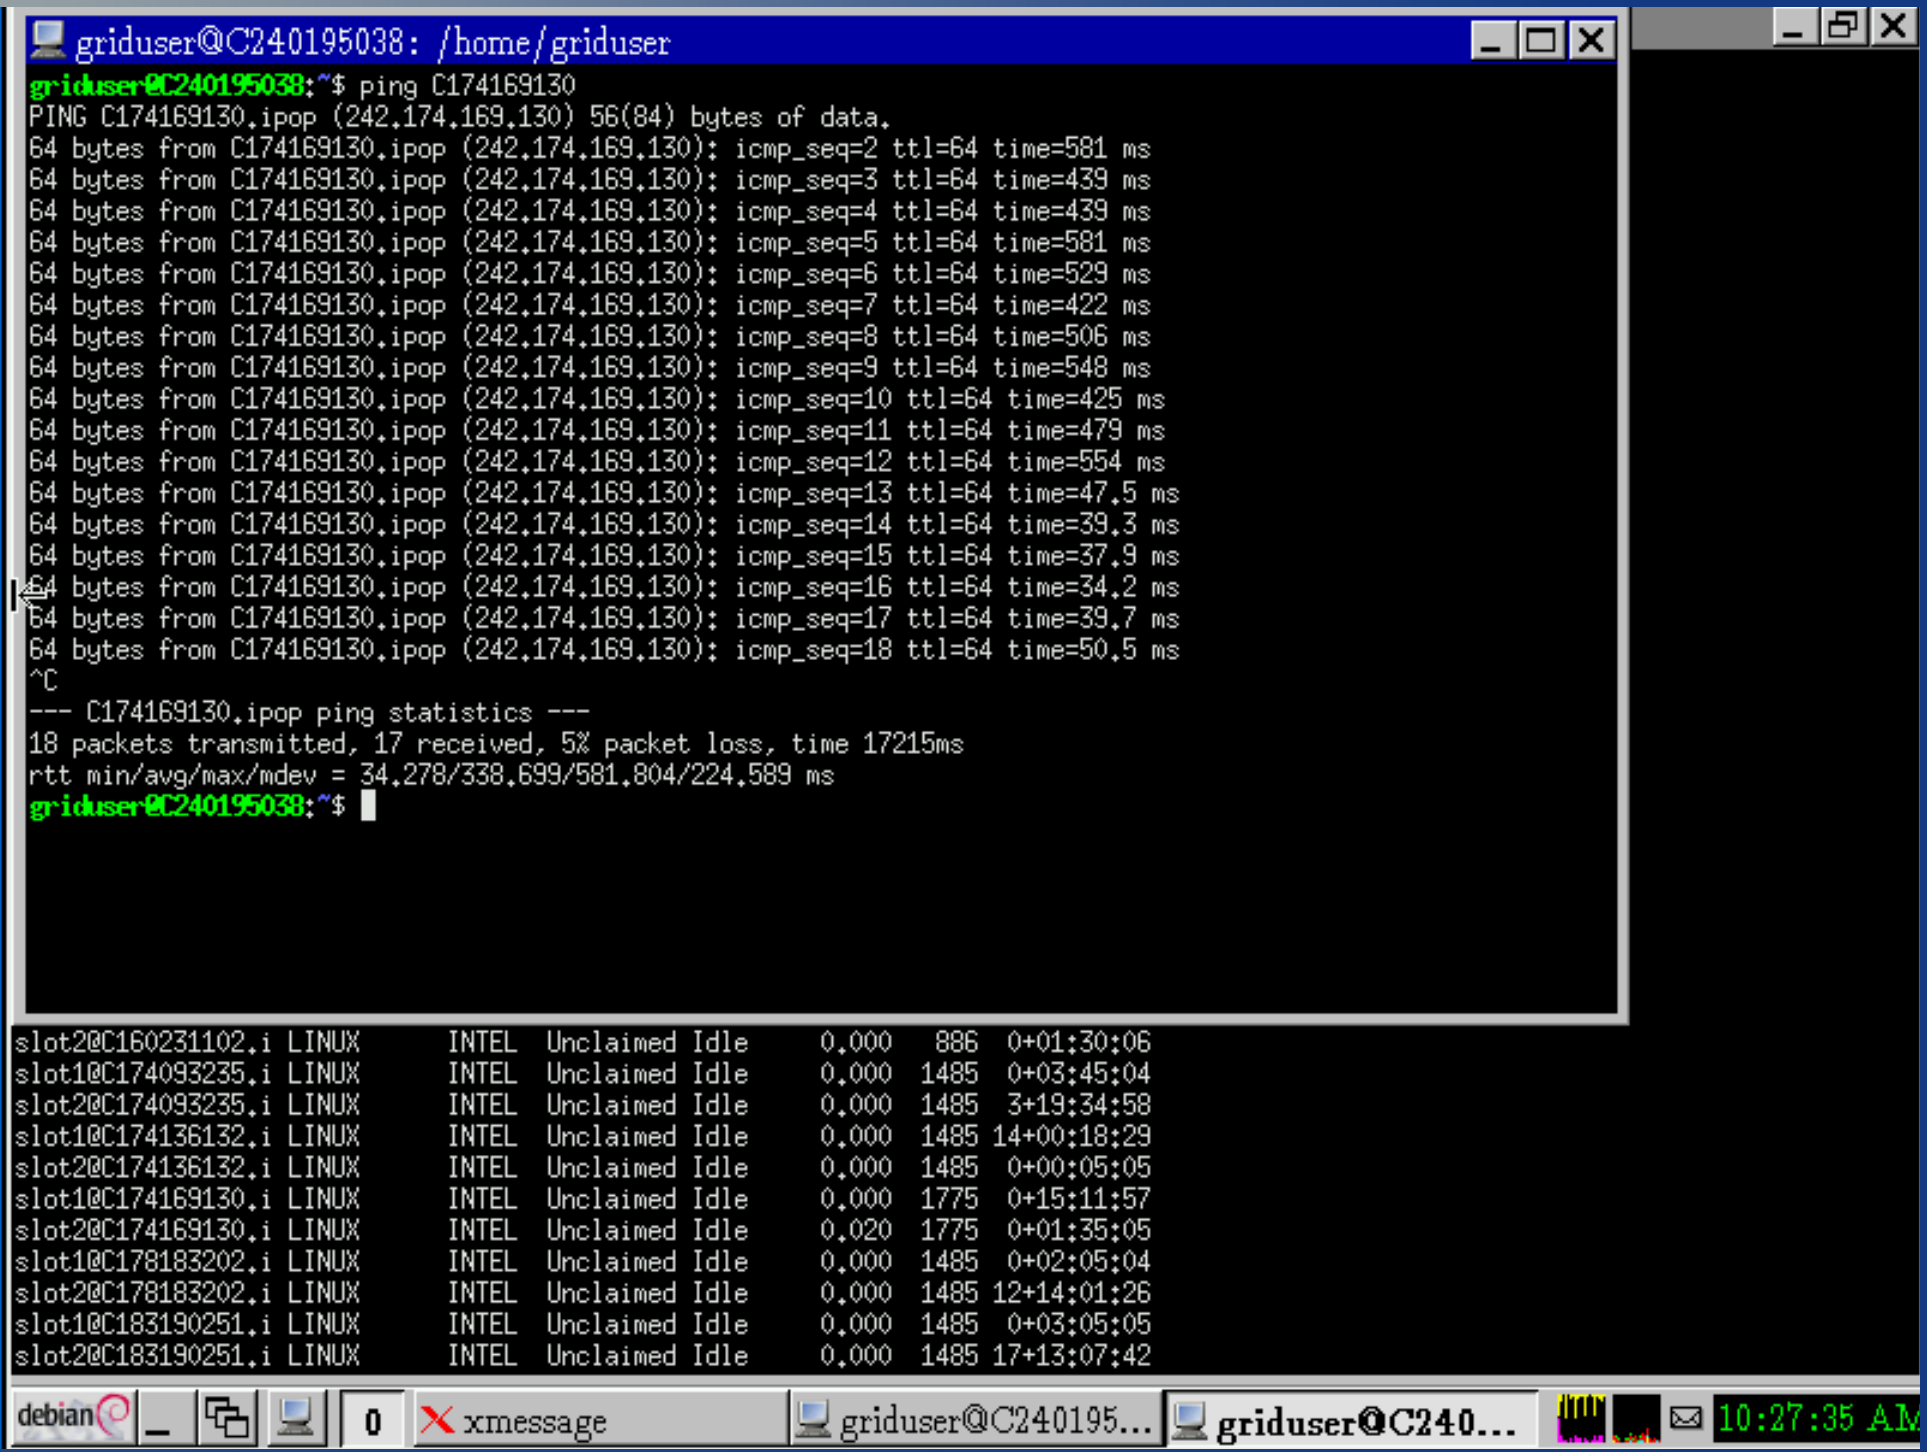

UF UNIVERSITY of FLORIDA ACIS

Where We Are Now

- CondorView
- Tutorials
- Hadoop Appliance (Uses Condor to
- Organize)
- Archer
  - Computer Architecture research/education
  - Hundreds of cores distributed over 6 universities in the US over 3 years
- Over 100,000 CPU Hours used since October

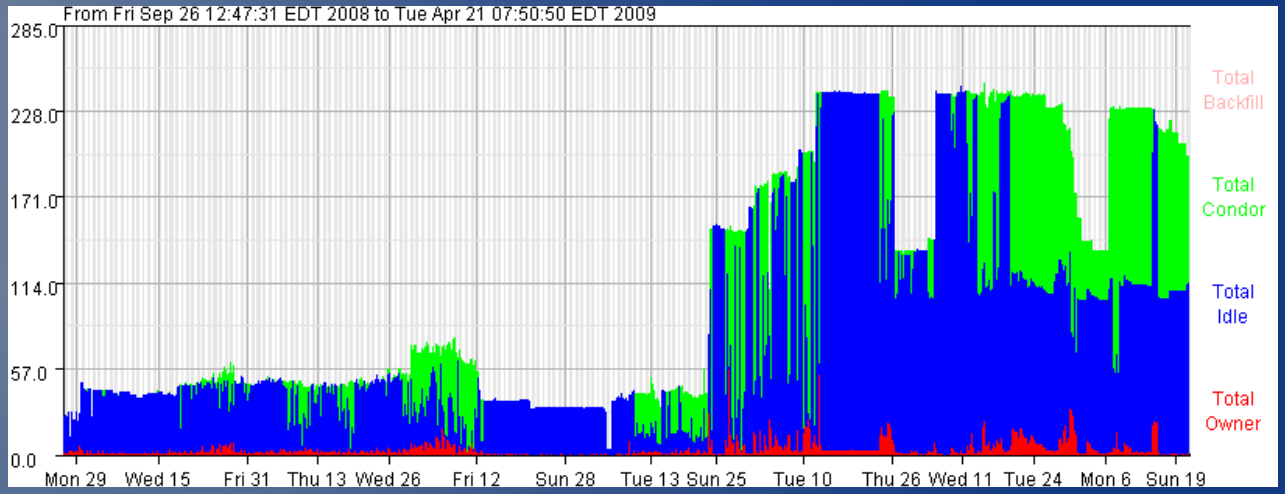

UF UNIVERSITY of FLORIDA ACIS

Projects Using IPOP

- Purdue – BoilerGrid
- Clemson – Campus Grids
- FCCN in Portugal – Hadoop-based Web Indexing - GaPPa

Questions

- Our Projects
  - Grid Appliance, Archer – grid-appliance.org
  - IPOP – ipop-project.org
- Acknowledgements
  - ACIS P2P Group
  - Condor Group – Ben Burnett and Alain Roy
  - Effort sponsored by the NSF under grants OCI-0721867 and CNS-0751112.
  - Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

UF UNIVERSITY of FLORIDA ACIS

Simics w/ NFS

- Simics requires license server – runs on condor master – Time sync issue :(
- Can use LARGE (many GB) input files
- Expensive to transfer over the wide area!
- NFS allows block level access, only transfer what you need, and data cache, so if a job comes again, we reduce our transfer overhead
- From http://www.grid-appliance.org/wiki/index.php/Archer:Simics

User Feedback

The best thing we like about Archer is its computing power and easy accessibility and usability. We can easily access 50 or more Intel Xeon cores remotely for our computing needs to accomplish something we cannot do before. For example, in our recent data prefetching and cache management studies, we used more than 10000 simulation hours on Archer.

Besides the huge computational resources, Archer is unbelievably easy to use. All the needed software is packaged in a virtual machine, including the IPOP, which provides communication among the grid. It is all transparent to users. Users just need to download the Grid Appliance package, and is ready to go. Archer even provides a couple of popular CPU simulators by default, like Simics, SimpleScalar and PTLsim, etc. Our group uses Simics frequently and glad to see it is available on Archer. Besides, Archer employs Condor to manage all the tasks and resources, which makes it easy to deploy/monitor the tasks and need not worry about the resources. With the newly added feature of NFS interface, we can do more in a customized way. It allows mounting a local virtual disk to the grid, and sharing user-specific files, i.e., large Simics checkpoints. All the grid nodes can access the files shared in the NFS file system. This feature helps build our own simulation environments efficiently.

- Jih-Kwon Peir