



Compute Resources and Infrastructure

*An overview of UC Davis' existing and
planned setup*

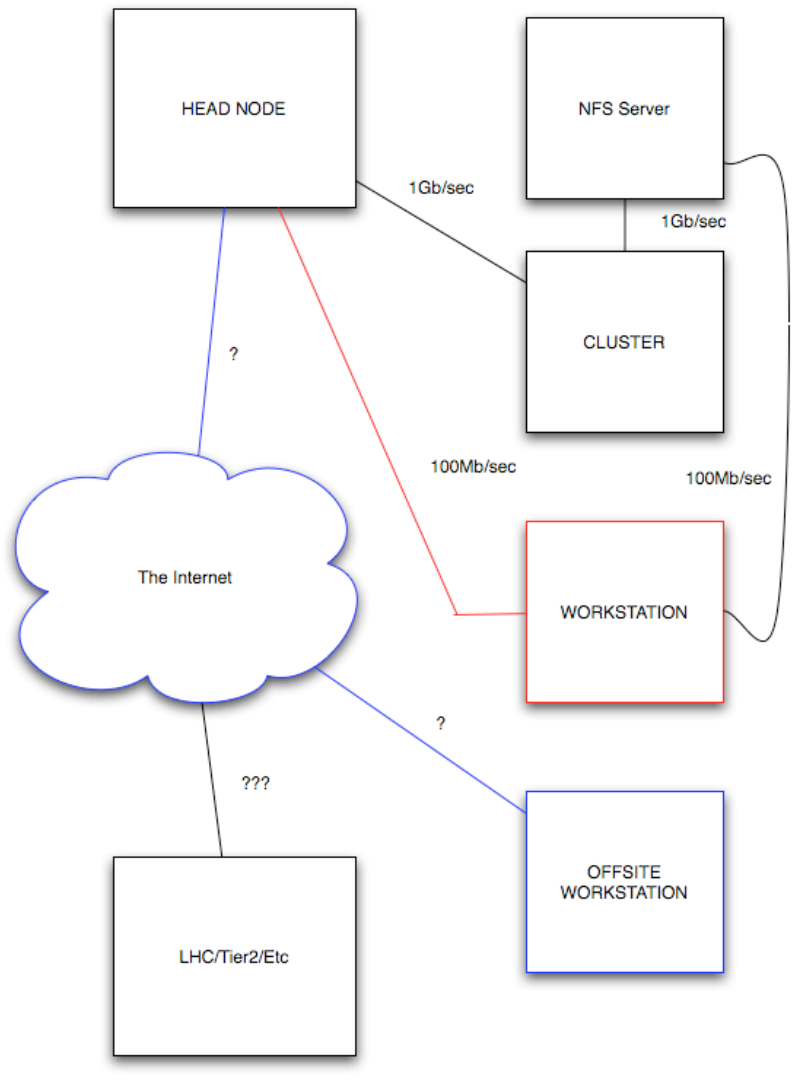
Michael Squires

Existing Setup



- ✧ UC Davis HEP currently utilizes a 14 node (two processors per node) ROCKS cluster
- ✧ A workstation connects to the head node via the WAN or the LAN (10/100) and then utilizes condor or cluster-fork to submit jobs for the compute nodes to work on
- ✧ The cluster is serviced by dedicated NFS shares that allow for uniform software presentation and data access

Existing Setup- A more detailed view



- ⌘ To the left is a block diagram of several possible connection paths for a given workflow
- ⌘ We exclude the type of physics analysis being performed from the graph, but it does make a difference in terms of data rate required
- ⌘ We see that in every case network bandwidth is a huge factor in over all performance

Existing Setup: Issues

- ⌘ Complex software setup
- ⌘ Poor speed between workstations and cluster. On the order of 1MB/sec
- ⌘ Poor speed between our institution and other institutions. On the order of .5MB/sec
- ⌘ Both are in spite of performance tuning.
- ⌘ On the bright side most operations are compute bound outside of the data set transfer itself :-)

Planned Setup

- ⌘ A staged roll out of a large ~ 100 node cluster
- ⌘ Start with a 10 node miniature version and continually add to it as resources allow
- ⌘ Software architecture and network connectivity are identical with the only difference being that the machine is located physically in the campus data center
- ⌘ The challenges and issues are the same in nature but we require more bandwidth for LHC data set transfers