

# Experience and Benchmarking with Virtual Machines

W. Fernando, H. Kagan  
Ohio State University

US ATLAS Tier 3 Meeting  
ANL Oct 29-30, 2009

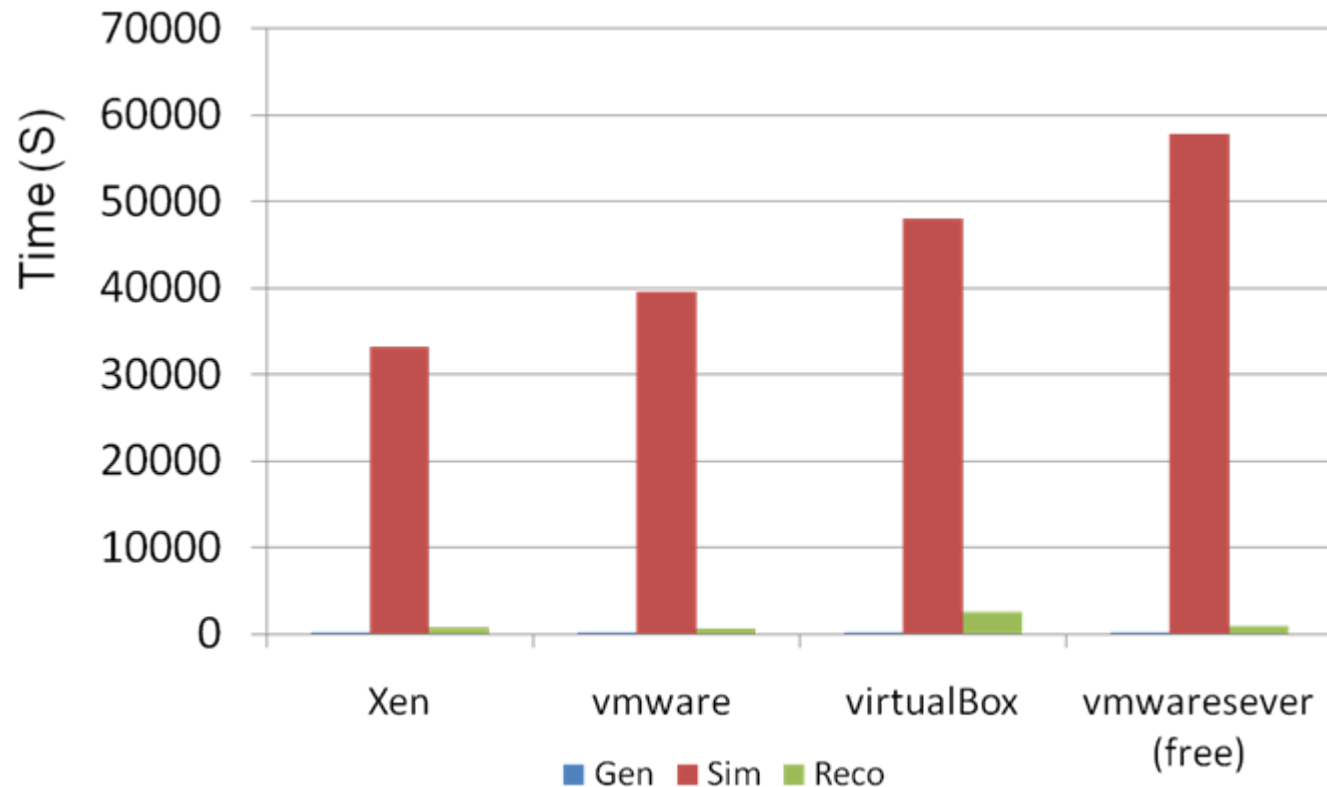
# Virtual Environments

- ATLAS software is difficult to use with new hardware, difficult to maintain
- One solution is to use new hardware with your favorite new operating system but running an “old” virtual machine on it – does this work?
- Virtual Machine Systems considered:
  - Xen
  - VMware
  - SunVirtualBox
  - KVM

# Software Tests of Virtual Machines

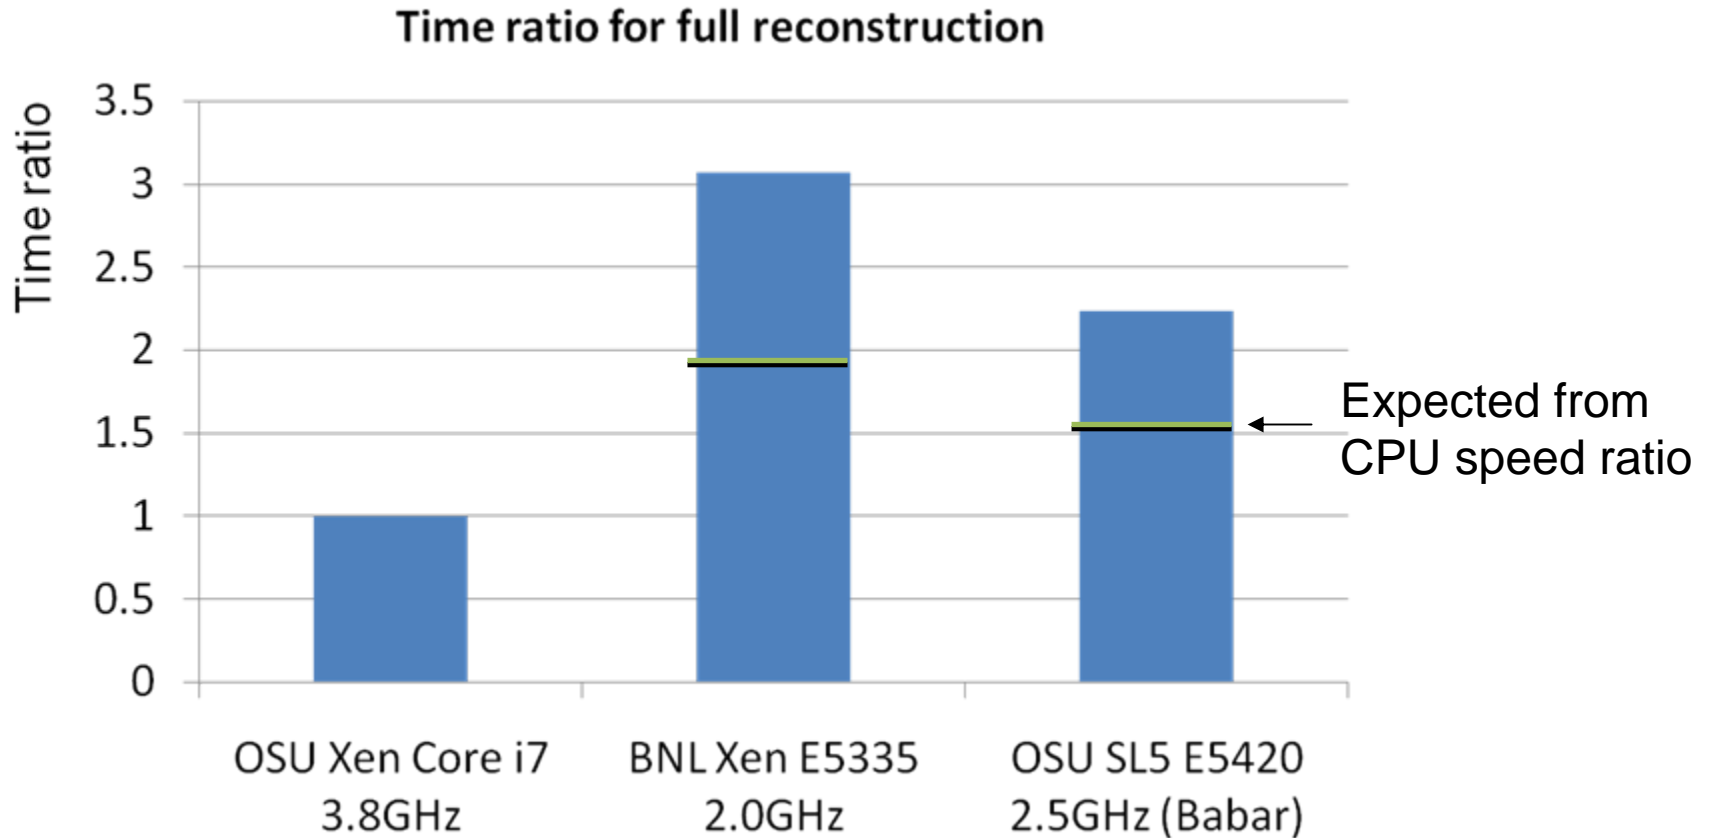
- ATLAS Software used:
  - ATHENA (generation, simulation, reconstruction)
  - PROOF /XROOTD (analyze 1M events)
    - Physics analyzed  $Z \rightarrow e\mu$  (goal is 100K /sec)
- Write scripts which are system independent
  - ATHENA 14, 15 - measure speed
  - PROOF – measure how many nodes necessary
- Compare various Tier2/3 facilities for speed and functionality

# Different Virtual Environments on OSU Core i7 920 - 3.8GHz



Time for 100 full chain events using jobtransformations in Athena 15.3.1.4 production cache,  $Z \rightarrow ee$  105144 events using same script, same machine

# System Comparison



Time for 100 full chain events using jobtransformations in Athena 15.3.1.4 production cache,  $Z \rightarrow ee$  105144 events using same script

# Performance Summary

- Xen outperforms VMware and SunVirtualBox
- Few losses observed with Xen
- Performance gain observed above CPU speed ratio with high speed memory and disks
- Core i7's 3x faster than BNL user nodes  
2.3x faster than OSU Babar farm user node


To generate 100K signal events in 10 days requires 38 threads (5-6HT 4-core machines)  
- each 3.8GHz Core i7 (100 events/9 hrs)

# Performance Summary

- Proof /XROOTD tested only locally with 1M Zemu analysis (simple analysis) with D3PDs on Corei7 cluster  
104K events/sec with 28 threads


To analyze 100K events/sec requires  
28 threads (4HT 4-core machines)  
- each 3.8GHz Core i7

# Synthetic Benchmark Confirmation - Processor Speed






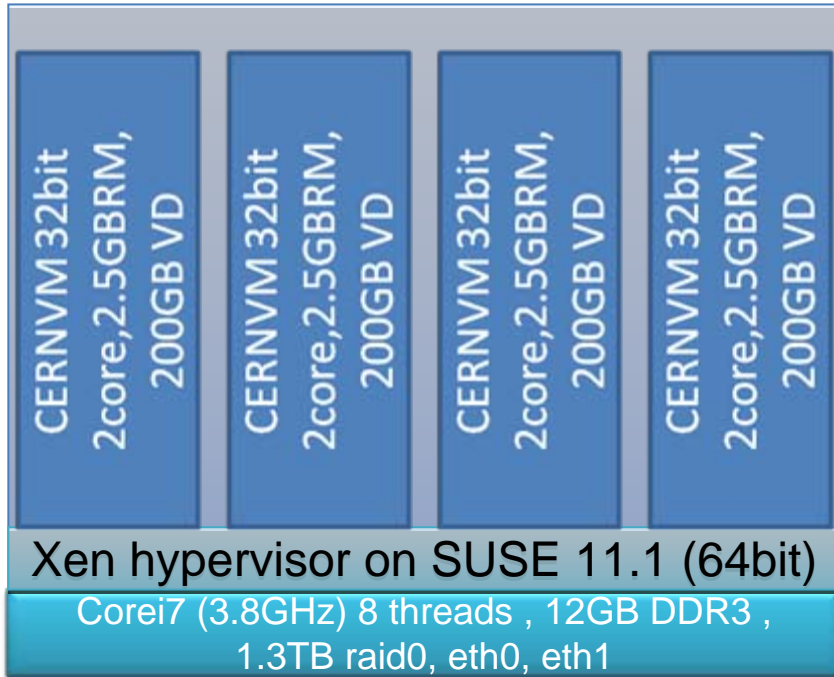
# Synthetic Benchmark Confirmation – Memory Speed



# Synthetic Benchmark Confirmation – Disk Read Speed



# Virtual Machine Setup - Details



## Advantages:

- Can have any host...like SUSE
  - No hardware driver issues
- Can run 32bit machines on 64bit hosts
  - 32bit machines < 4GB
  - Athena needs > 1.5GB/core
  - More problems with Athena in 64bit
- Different hosts same virtual machine
- Availability of cernvm with good support

# Software

## cernvm

- Athena (base, production, nightly)
- Eclipse Editor
- Event Display (VP1)
- **Condor**
- afs

## iSCSI

## Modified Tier3SW package (atlas Canada)


- Non-Athena Software

PBS to be replaced with Condor/Arcond

ROOT build from source with proof, xrootd and roofit

# Athena Versions Available (current)

- 14.2.25 (8,9,10,11)
- 14.5.0 (7)
- 14.5.1 (4,5,6)
- 14.5.2 (6,12)
- 15.0.0 (1,2,3,4)
- 15.0.1
- 15.1.0 (1,2,3,4,6,7)
- 15.2.0 (1)
- 15.3.0 (1,2)
- 15.3.1 (1,2,3,4,5,6,7,20)
- 15.4.0
- 15.4.1
- 15.5.0
- 15.5.1
- 14.1.23 (1,2,3,4)
- Nightly update (if needed)



/opt/atlas/software  
(cvmfs)

AtlasLocalRootBase

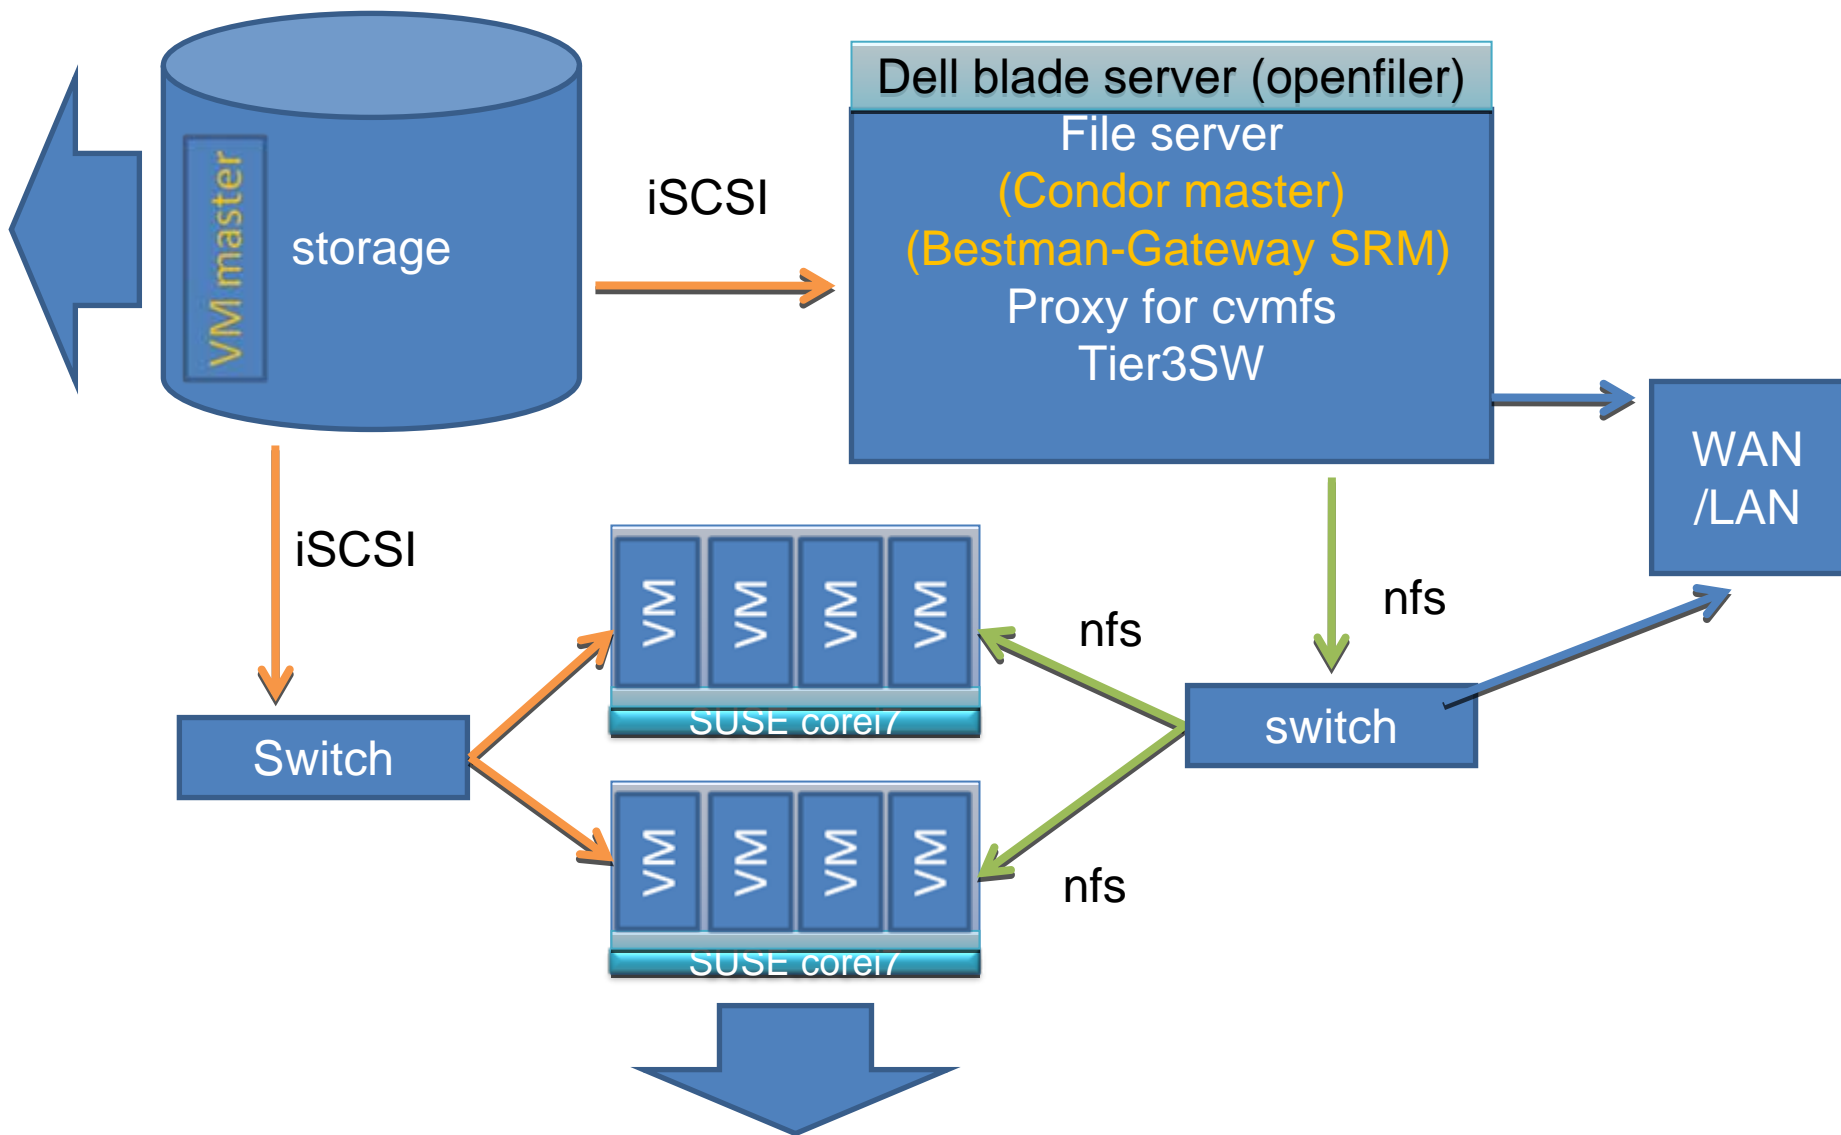
# Non Athena software

- gcc versions gcc346\_i686\_slc4, gcc432\_i686\_slc4, gcc432\_i686\_slc5
- Python 2.5.2, 2.5.3, 2.5.4, 2.6.1, 2.6.2, 2.6.3
- Atlantis
- DQ2Client 0.1.31
- Ganga 5.3.6
- gLite versions 3.1.38-0
- pacman-3.29
- ROOT 5.22.00-slc4-gcc3.4 , 5.24.00-slc4-gcc3.4
- PandaClient 0.1.95
- PBS (portable batch system) => replace with Condor (arcond)
- ROOT 5.24.02 with proof+XROOTD+roofit



Tier3SW  
package  
(atlas  
Canada)

# OSU Proposed System Design



ANL Tier 3 Meeting

# Summary

- Virtual Machines can perform all the Tier 3 required functions
- For ATLAS software Xen appears to be the best virtualization tool
- Careful selection of hardware can result in large performance gains



# Backup Slides

# Hardware

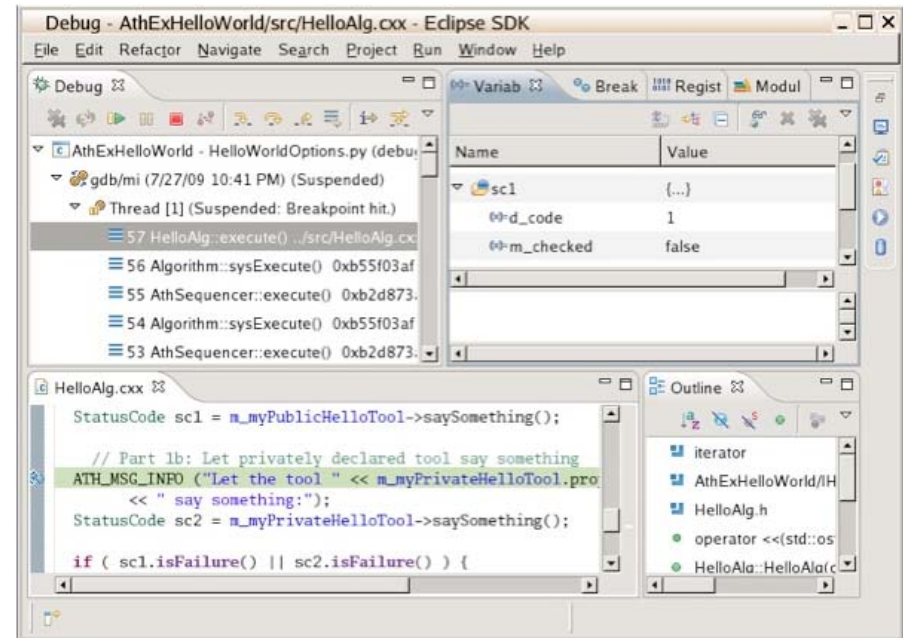
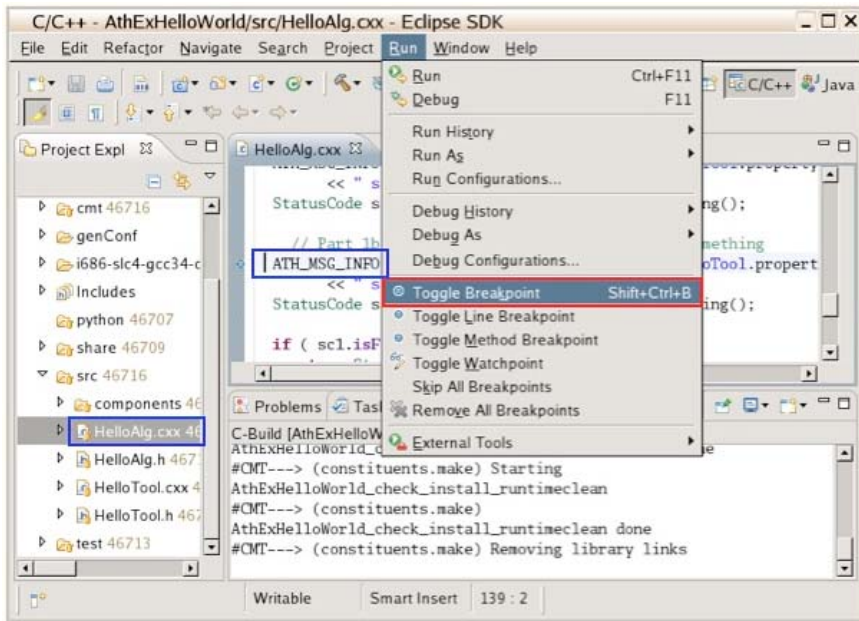
- A Dell blade server as a file server
- corei7 machines;
  - 8 threads, 12GB DDR3 RAM, 1.3T disks in raid0 array, dual ethernet
  - SUSE 11.1 / Xen
  - 2 tested, 3 ready to setup, more being built
- Hardware performance comparison
  - Script based – ATHENA, PROOF

# Cernvm

- Pre configured with cvmfs mounted athena(have many athena versions + caches +nightlies, update globally)
- Modified cernvm(xen,32bit) to use as worker nodes
  - 2 core, 2.5GB RAM, 5GB swap
  - Add virtual disk /data1 (200G, from local raid array)
  - 10TB through nfs as storage /n/atlasxx
  - LDAP
  - afs

# Eclipse Editor

- Preconfigured with cernvm
- <https://twiki.cern.ch/twiki/bin/view/Atlas/CernVMGuideEclipse>



# Manage T3SW update

- Athena

- add the base version /production cache to ALRB-athena-add.lis
- source parseConfigFiles.sh --pacmanOptions="-pretend-platform slc4 -allow trust-all-caches tar-overwrite" --mirror=<http://www.usatlas.bnl.gov/computing/cache/Atlas.mirror> --doRun (mirror ==nearest mirror)

- Other software

- Changed (add) the versions in the installSW.sh
- Source updateManageTier3SW.sh