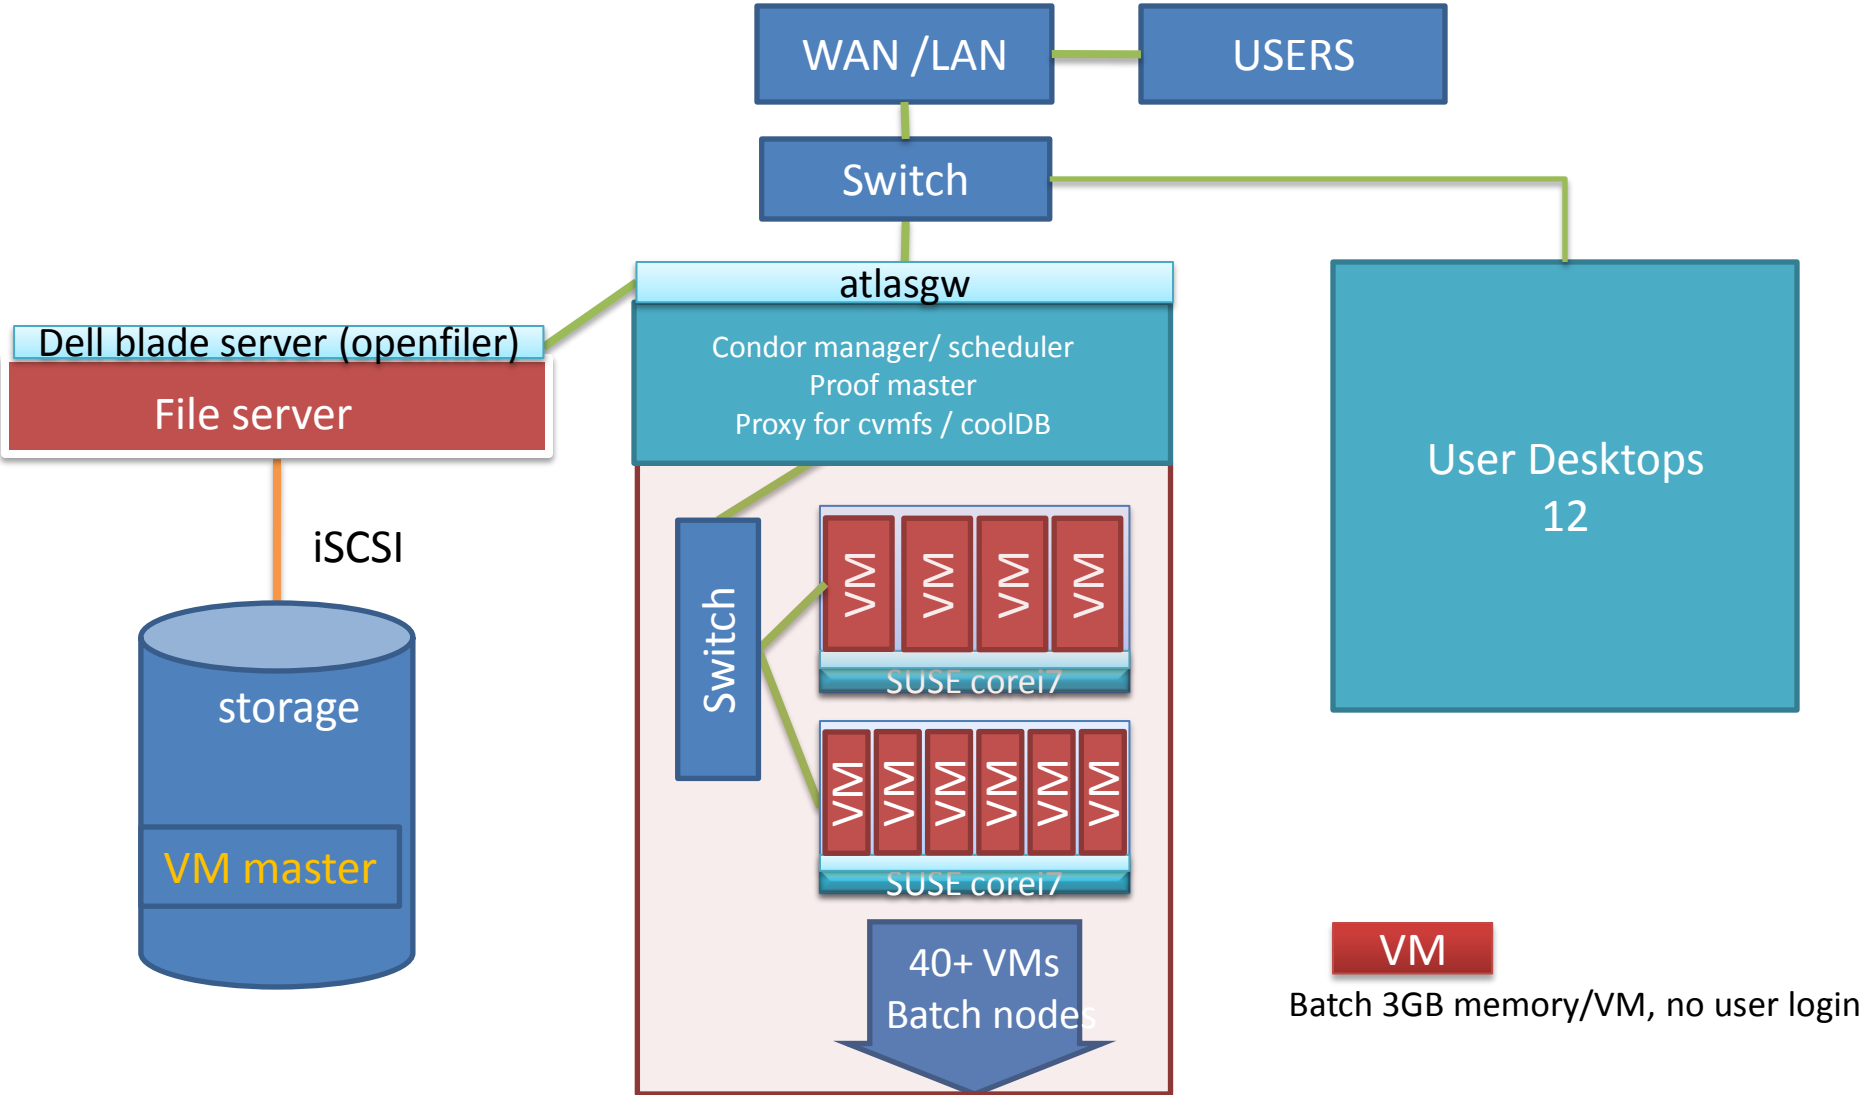


# Atlas Software Structure

- Complicated system maintained at CERN
  - Framework for Monte Carlo and real data (Athena)
    - MC data generation, simulation and reconstruction
    - Analysis of MC/real data
    - CPU bound
  - Conversion of data files (D3PD maker)
    - Raw data (AOD/ESD) → flat ntuples (rootuples)
    - Extreme I/O
  - Analysis of ntuples (root/proof)
    - I/O bound+CPU combination

Compare these programs in Real and Virtual Systems

# Atlas T3 Virtual System design



# Real/Virtual Comparison designs

## Dell PowerEdge R710

SL5  
8 core HT real Machine  
36GB memory total  
Athena locally installed  
Local File system  
Bonnie++ disk benchmarks

Read PerChr: 67955K/s  
Read Block: 270546K/s  
Random seeks: 428.9/sec


## Dell PowerEdge R710

SUSE 11.1 /Xen kernel  
16 cernvm (v1.6 Xen batch SL4)  
2GB memory per VM  
Athena with CVMFS  
Local File system  
Bonnie++ disk benchmarks(inside VM)

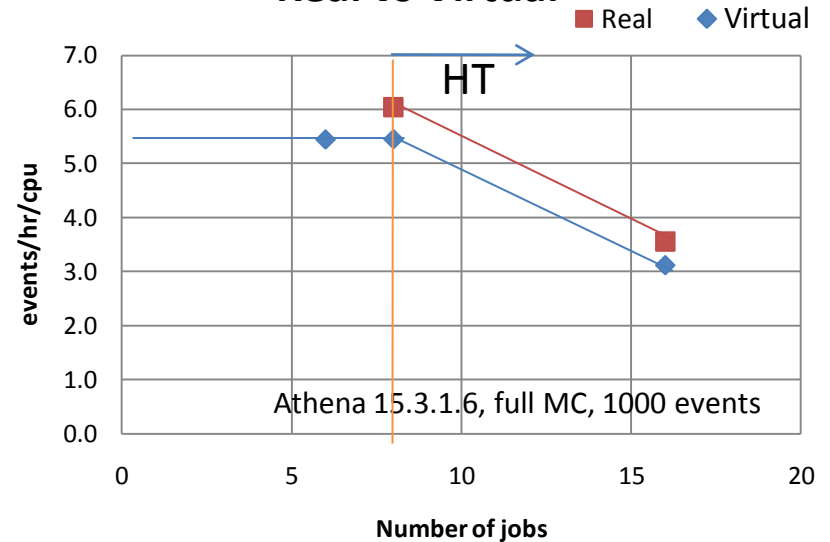
Read PerChr: 55828K/s  
Read Block: 252955K/s  
Random seeks: 410/sec

# Athena Performance - CPU Bound

## MC event rate vs CPU speed (Virtual)



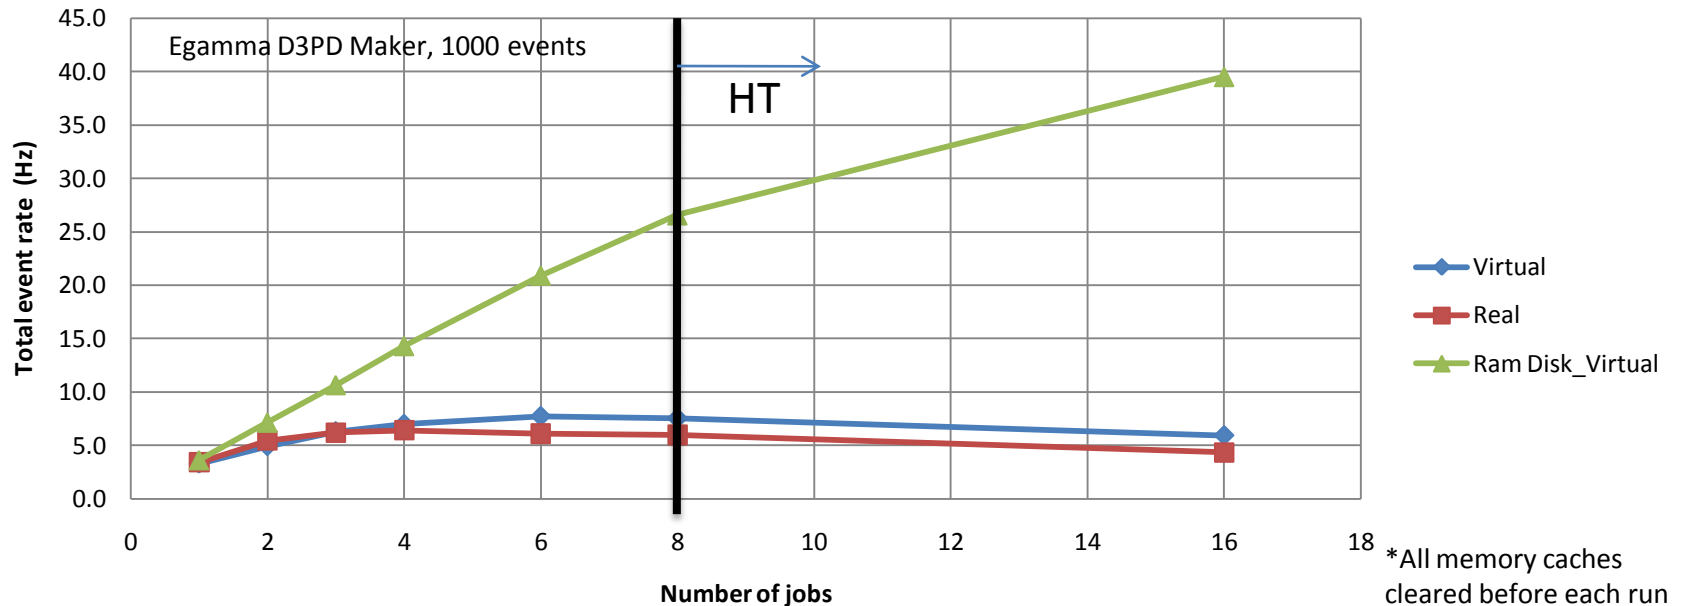
## Real vs Virtual



- For CPU bound jobs: 10% performance loss for virtual systems
- MC generation/simulation/reconstruction is CPU bound
- Performance scales with CPU speed

# D3PD Maker Performance - CPU+I/O Bound


## Real vs Virtual (same hardware same job)




- No I/O performance loss for virtual machines with ATLAS software
- Evidence for I/O performance (D3PD Maker) in virtual machines being better than in real machines (up to 20%) due to other limitations
- Local HD raid array with 250 MB/s read bandwidth (block) and 440 random seeks/sec can't provide enough I/O, Performance peaks at 6 jobs; performance loss noticeable at 3 jobs relative to Virtual RAM disk.

# iostat from real machines when 8 parallel D3PD jobs on real/virtual


## Rate of read requests



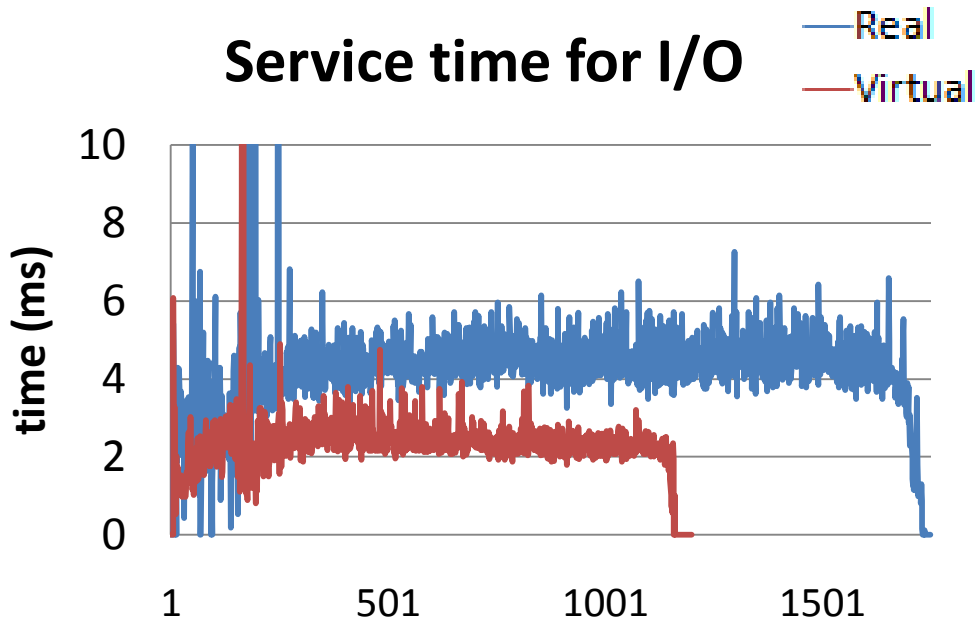
## I/O read speed (KB/s)



## I/O Wait time



## Service time for I/O



# Proof performance

