

# **Special Hardware Needs of ATLAS Tier 2**

**Wei Yang**

Embarrassingly parallel jobs, up to now

Devoted significant effort to meet the storage requirement

## **What do we need to achieve in storage:**

### **Need to handle random read dominated IO**

IO from user analysis jobs peak at ~200 bytes per read

### **Must be reliable**

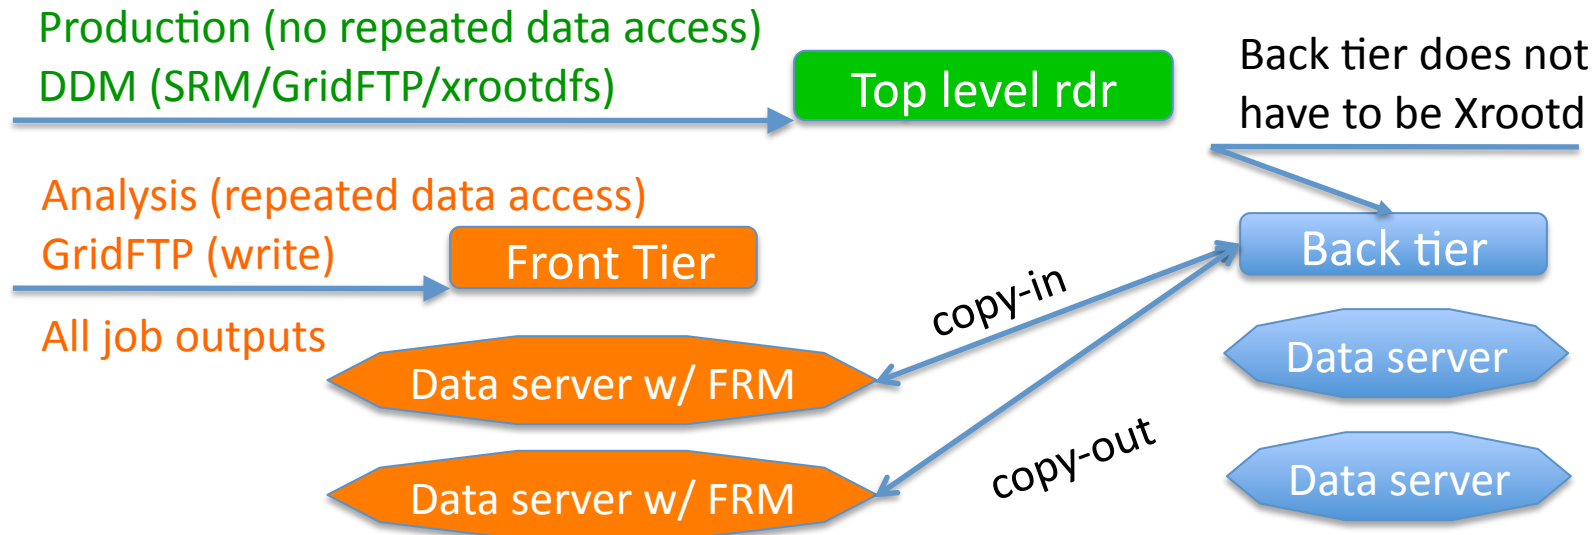
Monthly site availability and reliability report to ATLAS and World Wide LHC Computing Grid

### **Must be cheap**

- ATLAS determines CPU and storage pledges Tier 2s
- Tier 2s are competing on cost vs. effectiveness
- Dell provides cheap low end storage, at very low density

# Two-Tier Xrootd Disk Storage

1TB =  $10^{12}$  Bytes




Front tier: moves data in and out of the back tier storage


- 650 TB useable in 20 Thumper/Thors,
- 1 raw TB/spindle, 7200 rpm SATA
- provide 5000 iops/box without hurting WAN data transferring

Back tier:


- 1080TB useable in 42 MD1000s attached to 5 Dell R710s
- 560TB more ordered
- 2 raw TB/spindle, 5400 rpm SATA



Data in and out  
via GridFTP



Front Tier



Back Tier

## **Future storage for ATLAS Tier 2**

Actively looking for high density storage at optimum price

Can we utilize SSD to improve random IO performance

Federated Xrootd storage system

- R&D
- Federate multiple ATLAS sites' storage across the WAN
- Free users from data transferring and data management tasks
- Using Xrootd protocol but not limited to Xrootd storages
- Will change the data accessing methods and pattern from users