Architecture and Resource @ ATLAS Western Tier 2

Wei Yang

ATLAS Western Tier 2 User Forum meeting
SLAC April 6-7 2009

What do we have, ATLAS specific

- **Xrootd Storage**
  - 275 usable TB on 10 Sun X4500 thumpers, format under ZFS and run Xrootd
    ZFS is a high performance, multiple parity software RAID, detects silent data corruption
  - Good for reading dominated data accessing. Optimized for analysis activities
  - See Andy Hanushevsky's talk about how Xrootd works

- **LSF batch system**
  - ~ 600 cores purchased by Western Tier 2, Redhat Linux 4, 64bit, 2GB/core
  - Neal Adam will talk about LSF batch system usage

- **Interactive user login machine**: atlint01.slac.stanford.edu
  - 8 core Intel Xeon @ 3Ghz, 8GB, Redhat Linux 4, 32bit
  - Xrootd File System (XrootdFS), for easy to use (and for its Posix FS interface)
  - DQ2 client tools (dq2-ls/get/put, etc.)

- **Special, temporary hardware resource needs can be arranged**
  - Out of SLAC's general computing resource pool
  - Relatively easy to setup for local users
  - Hard for users coming from Grid

What do we have, cont'd

- **Open Science Grid Middleware + Selected LHC Computing Grid (LCG) tools**
  - Computing Element, Storage Element, monitoring, accounting, client tools, etc.
  - LCG Utils
  - LCG File Catalog (LFC) service, LFC clients tools/Python API

- **ATLAS DDM (Distributed Data Management) site services and DQ2 client**
  - DQ2 Tier 2 site service
  - DQ2 client tools

- **Info on the Web**
  - WT2 web site: http://wt2.slac.stanford.edu
  - SLAC ATLAS experiment web page: http://www.slac.stanford.edu/exp/atlas
  - ATLAS HyperNews forums: https://espace.cern.ch/atlas-forums/default.aspx
    Look for "Grid Jobs at SLAC" and "Non-Grid jobs at SLAC"

- **Additional Resources**
  - ATLAS releases in SLAC AFS group space
  - SLAC ATLAS group NFS server

WT2 Architecture (user's view)

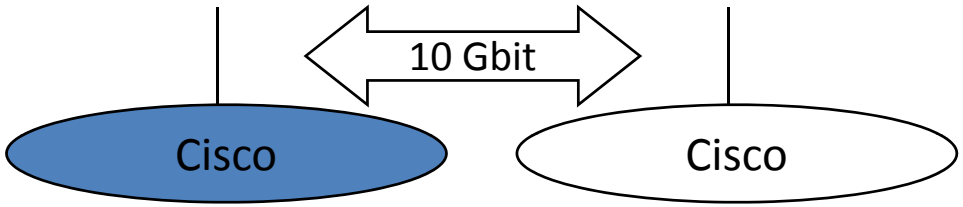

OSG middleware/proxies/ATLAS DQ2 service

User Interactive machine: XrootdFS, dq2 client

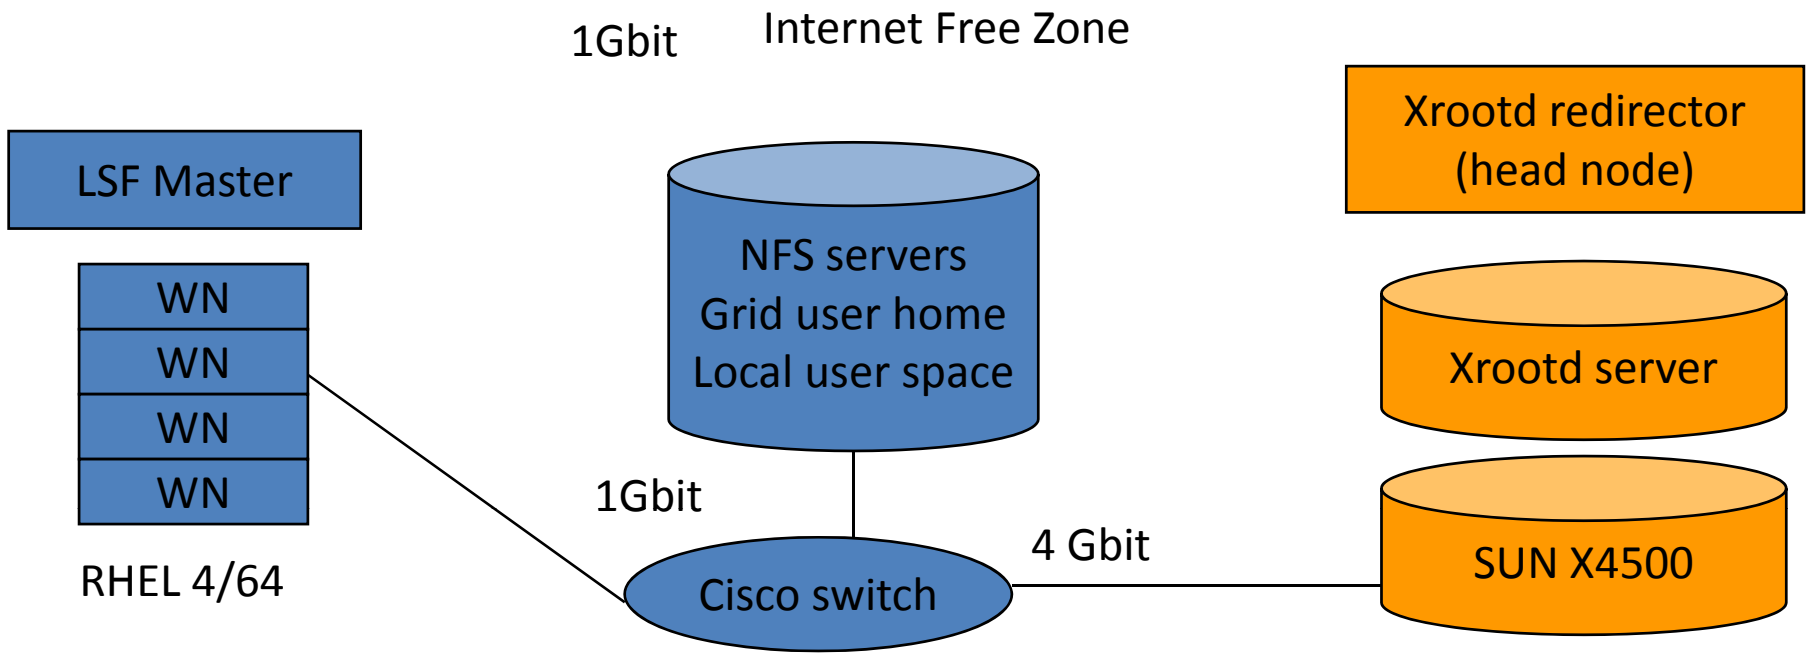

~ 600 cores, 275 TB usable

ATLAS DQ2 sites hosted by SLAC

- **Each ATLAS Tier 1 / 2 site has multiply DQ2 sites**
  ATLAS uses DQ2 site name/types (Space Token) for data categorization
  - SLACXRD_DATADISK, _MCDISK ⇔ "real" data and MC data, long term (months)
  - SLACXRD_USERDISK, _GROUPDISK ⇔ short term, mainly used by Panda user jobs
  - SLACXRD_PRODDISK ⇔ Panda production jobs, very short term
  - SLACXRD ⇔ Initial SLAC DQ2 site. Not longer used by ATLAS central DDM
- **Data/Datasets in SLACXRD_* sites are available to users**
  - Treat DQ2 datasets in these sites as readonly.
  - Storing in Xrootd space. **Deleting may create disaster**
  - SLACXRD_PRODDISK: data come and go, don't depend on them
- **Use dq2 client tools (dq2-ls) to check dataset availability**
  - what datasets are available at SLACXRD_* sites, and location in storage
  - which sites have an interesting dataset
- **Request datasets to be moved to SLACXRD_***
  - http://panda.cern.ch:25880/server/pandamon/query?mode=reqsubs0
  - e-mail your request to HyperNews. We will try to help/discuss possibilities

Interactive Machine, what can you do

- **atlint01.slac.stanford.edu ⇔ Current interactive machine**
- **Run interactive job**
  - ✧ Debug Athena jobs
  - ✧ Run ROOT analysis:

    $ root –b –l

    root [0] TXNetFile f(“root://atl-xrdr//atlas/xrootd/file.root”)
- **Submit batch jobs**
- **Organize your own Xrootd data via XrootdFS**
  - ✧ Xrootd space are mounted under Unix file system tree, like NFS disks

    /xrootd/atlas/atlasdatadisk

    /xrootd/atlas/atlasmcdisk

    …
  - ✧ Support most unix commands, “cd”, “ls”, “rm”, “cp”, “cat”, “find”, etc.
  - ✧ Transfer individual data file between CERN and SLAC using bbcp
- **Use DQ2 client tools**

  dq2-ls, dq2-get, dq2-put

SLAC Computing in General

- **Obtain a SLAC computing account**
  - SLAC computing is centrally managed by Scientific Computing and Computing Services
  - For ATLAS users, pls follow instructions at http://wt2.slac.stanford.edu for an account
- **AFS**
  - Home directories of SLAC unix computing accounts are located in AFS.
  - AFS access permissions are controlled by ACL (access control list) and AFS token
    fs listacl $HOME/dir
    make sure you have an AFS token: /usr/local/bin/qtoken, /usr/local/bin/kinit
  - AFS is good for long term storage of documents, code.
  - AFS is NOT GOOD for data. Use NFS or Xrootd instead.
  - AFS space unit is "ASF volume", mounted any where under /afs/slac.stanford.edu
    fs listquota $HOME/dir
  - To request increasing AFS volume size, or request a new AFS volume:
    http://www.slac.stanford.edu/comp/unix/afs-req.html
- **Unix interactive machines**
  For general usage. rhel4-32.slac.stanford.edu, rhel4-64, noric
- **unix-admin@slac.stanford.edu for any Unix question**

SLAC Computing in General, cont'd

- **Grid computing:**
  - After registered person certificate with ATLAS, you will be assigned a SLAC Grid account
    - This is NOT your SLAC local Unix account
    - Currently only dq2-get and dq2-put use this Grid account
    - Not activated automatically, e-mail unix-admin if can't use dq2-get/put with SLACXRD_*
  - FYI, OSG CE is osgserv01.slac.stanford.edu, SE (SRM) is osgserv04.slac.stanford.edu

- **Windows account, E-mail**

  SLAC recommends Microsoft Exchange e-mail. See:
  http://www2.slac.stanford.edu/comp/messaging/

- **Computing security**

  http://www2.slac.stanford.edu/computing/security/

- **General help : http://www2.slac.stanford.edu/comp/helptrak/**