The CLuEDØ Linux Cluster

Roger Moore
Michigan State University

Current Statistics

- Total of 182 machines with 215 CPUs ranging from 200-2000MHz
- 470 users from 48 institutes and growing
  - Majority of D0 now has a CLuEDØ account
- 5 geographic sub-clusters
  - Plum, White, Green, Scarlett, Mustard
- 20 system administrators
- 6 large disk servers all using IDE RAID arrays
  - 4x640GB + 1x800GB + 1x960GB = 4.24 TB
  - 4 more potentially on the way

RedHat Linux 7.1

- Currently running RedHat Linux 7.1
  - Updates regularly applied by autoRPM
  - Investigating upgrade to 7.3
- Several extra packages added
  - Open PBS batch queue system
  - Kerberized OpenSSH, XEmacs, Acroread etc. installed as standard
  - XFS journaling file system for RAID arrays
    - Investigating replacing with ext3 or JFS
- Support for most of latest hardware
  - Up to date Kernels (2.4.17)
  - WebCAM video conferencing via VRVS and via GNOME-Meeting (H323)

Disk Configuration

- AutoFS used to NFS mount disks
  - Improves NFS reliability since mounts timeout and are re-made as needed
- Several AutoFS volumes
  - /home: user & services home areas
  - /rooms: big IDE raid servers
  - /work: work (scratch) area, one per node
  - /clued0: local cluster utilities, programs etc
    - Includes OpenOffice, AutoRPM directories etc.
- NFS biggest cluster problem to date
  - Ripon or dØ2ka going down causes "Stale NFS handle" errors and kills batch jobs
  - Investigating OpenAFS as a replacement

PBS: Portable Batch System

- Batch queue system used on CLuEDØ
  - Open source and free
  - Commands similar/same as NQS and LSF
- Wrapper script 'cluesow' provided
  - Wraps the standard 'qsub' command
  - Sets users' institute group for resource tracking and job priority
  - Sets options so that memory limits are correctly enforced
- Supports interactive jobs so debugging runs don't interfere with batch system
  - "cluesow –I" gives new shell prompt

PBS Configuration

- Four queues available
  - SHORT: up to 2 hours CPU/4 hours wall
  - MEDIUM: up to 12 hours CPU/48 hours wall
  - LONG: up to 72 hours CPU/96 hours wall
  - GENERAL: sorts and pipes to other queues
- Users specify limits at submission time
  - Scheduler benefits correct estimates
- CPU shared according to institute's contribution to CLuEDØ
  - Implemented by scheduling priorities
- Jobs executed on fastest available node
  - CPU time allowed scaled to 1GHz P-III
    - i.e. 1 hour limit scales to 2 hours on 500MHz node

Batch Statistics

- PBS now in frequent use
  - More CPU power than DØmino
  - Regularly ~100 jobs in batch system
  - Limit is memory NOT CPU
    - Most machines have 256MB/CPU
    - Need 512MB to run DØ code, but…
    - …need >1GB to compile DØ code!
- Over the past two weeks:
  - Served 14,407 hours CPU (=46 CPUs for 24 hours/day)
  - Average job queue time 1 hour (max 35 hrs)
  - Job efficiency ~83% (fraction of the CPU requested which was used)

DØ Environment

- Full DØ development environment available on every workstation
  - Mounted from dØ2ka, always up to date thanks to Alan and Paul
  - Local packages mounted from CLuEDØ server to allow for customizations
- Home directories also mounted from dØ2ka
  - Dave Fagan does our backups
  - New scheme for combined dØmino and CLuEDØ home directories

SAM on CLuEDØ

- SAM installed but still not yet working…
- Current progress temporarily stalled
  - Tests could only be done on the CLuEDØ server since they require gigabit ethernet
  - …these tests overload and crash the server!
- PBS adapter hard coded with executable and queue names (doesn't use PBS API)
  - Work around for queues found
  - Special scripts needed to adapt CLI interface
- Still possible conflicts with PBS scheduling vs. SAM scheduling
  - Need SAM running in production use first to determine if this will be a problem

Remaining Problems

- Biggest problem on CLuEDØ so far is NFS
  - Linux NFS notoriously bad
  - Investigating AFS as a replacement
- Memory limits beginning to byte ☺
  - All new machines should come with 1+GB
    - Need AT LEAST 1GB to link DØ code
    - Need 512MB/CPU to run DØ code
  - In general please check the CLuEDØ hardware web page before buying a machine
    - http://www-clued0.fnal.gov/clued0/hardware/clued0_hardware.html
- Indiscriminate use of the batch system can take out cluster nodes…
  - 100 batch jobs attempting to read data files from a single node is NOT a good idea

CLuB: CLuEdØ Backend

- The design goals of CLuB are:
  - Execute mid-sized analysis jobs ~2TB
  - Allow institutes to contribute resources
    - Modular design
  - Integrate with the CLuEDØ cluster
    - Accept batch jobs from CLuEDØ nodes
    - Provide CLuB/CLuEDØ shared disk and SAM servers
  - Provide for easy remote access
- Design based closely on CLuEDØ but some differences
  - Single location to allow fast, local networking
  - More uniform hardware (don't choose your own)

CLuB Design

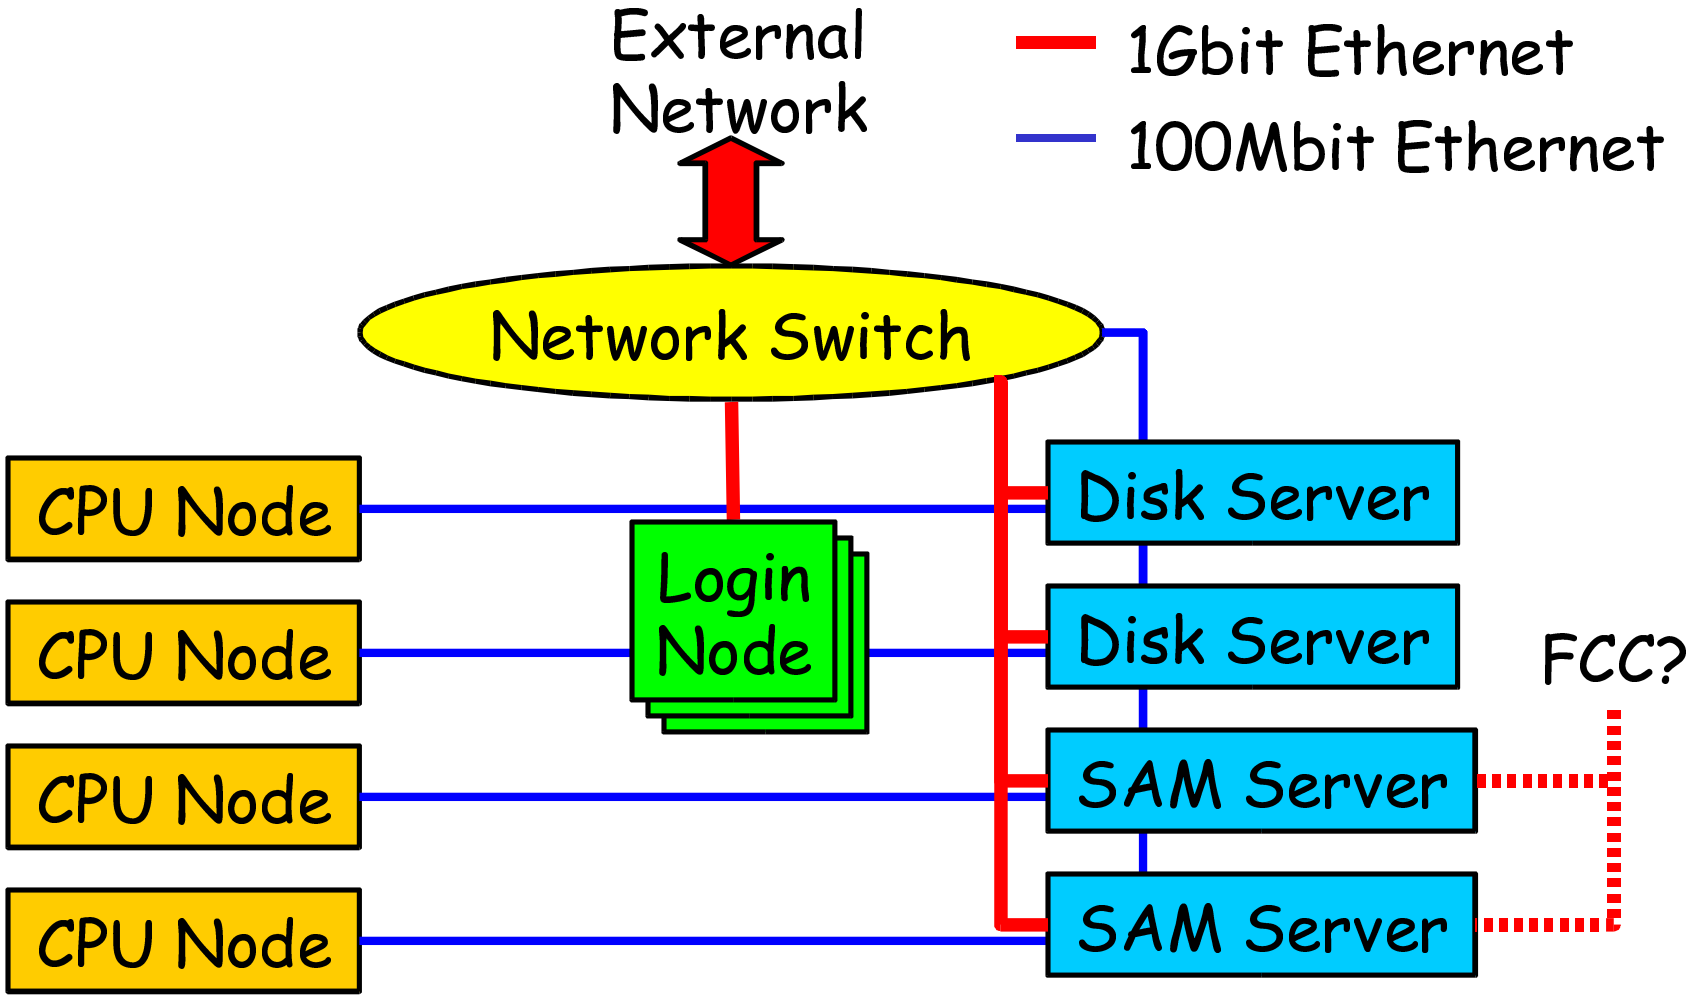
External Network
1Gbit Ethernet
100Mbit Ethernet
Network Switch
CPU Node
CPU Node
CPU Node
CPU Node
Login Node
Disk Server
Disk Server
SAM Server
SAM Server
FCC?

Design Details

- Jobs submitted via PBS
  - Read/write data via SAM or local disk servers
- External access to job output
  - RCP/SCP/FTP or AFS? for remote sites
    - Automate? Still need a cache in case network down when job completes
    - Long term: OpenAFS or replacement
  - NFS(AFS?) for CLuEDØ (RCP etc. still allowed)
- Login nodes provide remote access
  - Also remote access for CLuEDØ

SAM

- All data input handled by SAM
  - Read from Enstore or DØmino cache
  - Output can be written into SAM but not required
- Only SAM servers have access to external data sources
- Each CPU node has local cache
  - Use RCP to distribute data between caches
  - SAM takes into account local cache contents when assigning files from a project to a batch job

Hardware

- CPU nodes
  - Fastest, reasonable dual processors available
    - Currently dual 1.6GHz Athlon MP
  - 1-2GB memory (DØ code needs a lot!)
  - 80GB disk: almost all SAM cache
  - Cost: ~$2.5k
- Disk/SAM servers
  - Same base as CPU nodes
  - 3ware Escalade IDE RAID controller
    - 8x160GB IDE disks = 1.28TB disk
  - Cost: ~$5.5k
- Racks with remotely switchable power
  - Power cycle nodes remotely

Procurement

- Initial cluster seed is being bought by DØ
  - Blades for network switch (gigabit for servers)
  - Racks, power, terminal server (for serial console)
  - One 1TB disk server, ~8-10 CPU nodes
- Maintain detailed machine specifications with a single vendor (currently ASA Computers)
  - Institutes wishing to contribute should contact Dugan, Dave (Evans), or myself
- Timescale for CLuB implementation
  - Seed arrives at FNAL ~mid-July
  - Setup, install etc. by September

System Maintenance

- Administer CLuB as a CLuEDØ subcluster
  - Approx. same setup which existing CLuEDØ administrators already know
  - More machines but similar hardware and most without interactive users
  - No 24 hour support but could add remote sysadmins in, for example, Europe
- We would welcome any new sysadmins
  - Need working knowledge of Linux/UNIX
    - But you don't need to be an expert!
  - Fringe benefit: learn how to change job priorities on the batch system! ☺

Resource Allocation

- CLuB will be a central DØ resource
  - Administered by ORB
- Follow standard DØ practice
  - 50% goes to the institute
  - 50% goes to a physics group of their choice
- Apply this to both disk and CPU
- Share CPU by batch system scheduler
  - Means unused CPU can be used by other groups
  - Over several days allocation will work out (if the system is fully loaded)

Requests

- CLuEDØ needs help with system administration
  - Good response to last request in DØNEWS
    - Thank you!
  - Currently 20 x 0.2 volunteer FTEs
    - good for support, not for 'big' jobs
  - DØ NT cluster similar (smaller?) in size to CLuEDØ and has 2-3 fulltime FNAL personnel
  - Some support for CLuEDØ + CLuB would be nice…1-2 FTEs requested in computing plan
- CLuB will need institute supplied hardware
  - Disk and CPU servers

Conclusion

- CLuEDØ is reaching maturity
  - Growth slowed to <1 machine/week
  - Currently the largest source of analysis CPU for DØ (but I/O and memory limits)
- Institutes should start to think about adding disk resources to CLuB, not CLuEDØ
  - Better utilization of network
- Help with administration needed
  - FNAL-supplied professionals
  - Better incentives for volunteer administrators to encourage an increase in FTE fraction
    - Fixed length "tour of duty"
    - Shift exemption/credits?