

Agenda

- ClueD0 Status
- Interfacing to d02ka
- Upgrade to RH7.1
- Splitting into Subclusters
- Future of SAM on ClueD0
- Handling Data (Not SAM)
- PBS (Our New Batch System)
- More Hardware (Who Pays?)
- A.O.B.

ClueDO Status

- 48 machines
 - 9 dual processor
 - 15 machines at or above 733Mhz (7 are dual)
- 21 institutes
- 115 users
- We are no longer a small group of rebels....we are the establishment, the man.

Interfacing to d02ka

- We have moved/created 69 users with home areas on d02ka
- We hope to also mount d0 code distributions from there (it will be sync'ed with the build machines)
- These steps ease the load on our sys-admins for system backup, server maintenance, etc.

Upgrade to RH7.1

- new machines are getting more difficult to install
- RH7.1 has 2.4 kernel, usb support, kerberos, etc.
- chance to "start clean"
- still link-compatible with d0 code...as far as we have tested
(d0trigsim)

Splitting into Subclusters

- We are currently one big cluster of 48 machines
 - We want to be composed of subclusters of ~15-25 machines with separate day-to-day administrators for each subcluster
 - What defines a subcluster? How do we do the split?
- Ideas (defining a subcluster):
 - subcluster is mainly defined by who administers it
 - each subcluster has a "server" which serves /usr/local so that a failure of one machine cannot take down all of clued0
 - servers also are slave servers for YP/LDAP
 - subclusters are principally defined by geography
 - start with 2 or 3 subclusters
 - need to identify sysadmins for these subclusters
 - aim to use CLuMP - MySQL database which stores all machines' configurations

Splitting (cont.)

- Ideas (how to split):
 - create first "new" server (ripon-clued0) running RH7.1
 - add clients one at a time to new server, reinstalling each client as we go
 - use DNS aliases to make future changes easier
 - server1-clued0, server2-clued0 etc. point to subcluster servers
 - clued0 becomes alias pointing to the login, WWW etc. server

Future of SAM on ClueDO

- currently have SAM development station on swale
- plan to move to production SAM station on "new" server.
Should we do this sooner?? We could convert swale and test.
- Need to have central disk data store (~100's GB) plus space on all machines (~5+GB)
 - Who will buy this? Split over several institutes...

Handling Data (not SAM)

- Several new "disk server" machines are order or exist
 - 640Gb Raid array
 - dual processor ~1Ghz
 - minimum 512Mb RAM
 - ~\$6000
- We need to have a coherent scheme
 - Do we ask/require people to share some fraction?
 - Do we use an every institute for itself scheme?
 - Should we have some scheme for smaller institutes to contribute disks rather than entire servers?
 - Really should have gigabit connections for these machines...need help from D0/fermilab...

PBS - Our New Batch?

- Free replacement for NQS from NASA
- Lots of nice features:
 - single, centralized server
 - more flexible scheduling policies available (even write your own for ultimate customization e.g. special queues for problem users! :-)
 - Dave Fagan is thinking about using it for d0mino backend PCs (it's a lot cheaper than LSF!)
- However we haven't made it work yet!
- How do we share queues fairly?
 - need to come up with a scheduling policy...suggestions welcome!

More Hardware? (*Who Pays?*)

- So far we fly by the seat of our pants....what do we do if we need something?
 - Should we have a wishlist for institutes?
 - Can major purchases be split over several institutes?
 - do we need a clued0 common fund to achieve this?
 - is this something we can/should expect D0 to provide? Dave Fagan has been very helpful in providing support so far (home dirs etc.)