

NuComp News

(and FY11 budget analysis)

Lee Lueking

March 16, 2011

Outline

- ▶ New stuff and planned stuff
- ▶ Maintenance day stuff
- ▶ Miscellaneous
- ▶ Usage Eval and FY11 Requests review
 - BlueArc Disk
 - GPCF VM & Batch
 - FermiGrid, sp. GP Cluster

IF web area

- ▶ BlueArc disk area is being served by FNAL web server
 - Disk: /nusoft/app/web/htdocs/<group>
 - Admin: owner:group = nusoft:nusoft
 - GrpAdmin (proposed):
owner:group=<group>web:<group>
 - [url:http://nusoft.fnal.gov/<group>](http://nusoft.fnal.gov/<group>) (ex. nova, lbne)
- ▶ This is intended for automated/monitoring information. It is not a replacement for experiment's afs web area.

Group /home areas for if/gpcf

- ▶ Group areas have used “local” /home/<group> area with .k5login
- ▶ This is awkward; there are lots of group accounts now.
- ▶ Inventory and plan:
<https://cdcvs.fnal.gov/redmine/attachments/1770/gpcf-home-accounts20110211.pdf>
- ▶ move to /<exp>/app/home/<group>
 - REX will set up areas, and move obvious stuff (.k5login).
 - Exp’s/groups need to help/double check.
 - FEF will change to new homes
 - Should be transparent, will coordinate with each exp.
 - Existing /home areas will remain for a while.

Maintenance Day: Tomorrow 3/17

- ▶ Reboot GPCF/IF/MINOS/MB nodes for kernel updates
 - ▶ https://fefweb.fnal.gov/mediawiki/index.php/EAG,_GP_Grid,_IF_Experiments_-_17_Mar_2011
- ▶ Upgrade GlideinWMS on GPSN01, batch unavailable.
- ▶ IFORA PRD Oracle v10 to v11 upgrades:
MINOS, MINERvA, NOvA
- ▶ Nova & Argoneut BlueArc volume migrations to new back-end arrays:
 - blue3:/nova/app blue3:/nova/data
 - blue3:/argoneut/data blue3:/argoneut/app
- ▶ MINOS specific work
 - Bluearc /minos/data remount globally to eliminate data1,data2 paths
 - Removal of /minos/scratch compatibility link
- ▶ FermiGrid
 - rebooting the head nodes of the GP Grid cluster to do kernel updates., also changing automount maps to account for the bluearc changes.

Security Audit Next Week

- ▶ Next week the security auditors will be onsite.
 - Please be extra vigilant!
 - Do not have un/pw on display
 - Use password locked screen savers,
 - Lock you office
 - Don't let people tailgate into secure areas.
 - Ask to see badge ID when it doubt
 - Use common sense
- ▶ Spread the word

Next NuComp

- ▶ April 13 or 20, 2011 (?)

Usage Stats and Evaluation of FY11 Budget

- ▶ Where are we?
- ▶ How should we modify requests if budget:
 - Cut by 10%?
 - Cut by 30%?
- ▶ Disk and GP Cluster CPU were purchased near the end of FY10 to get a leg up on FY11.
- ▶ FY requests and FY budgets are offset in time.
- ▶ Cost of equipment might be different (hopefully lower) than estimated.


Central Disk (BlueArc)

	2009	2010	2011 (replace)	2012 request
MiniBooNE	✓80	30	✓15+✓10 (+13)	0
MINOS	✓90	✓30	✓30 (✓30+30)	30
MINERvA	✓10	✓20	✓20+✓15	30
Argoneut	✓20	0	0	0
NOvA	✓4	✓26	✓30	100
MicroBooNE	✓2	0	✓5	30
LBNE (all)	0	✓2	✓20	0
Mu2e	0	✓2	✓2	0
g minus 2	0	0	?20	0

✓ satisfied
✓ in hand


BlueArc Data Areas

BlueArc /<exp>/data



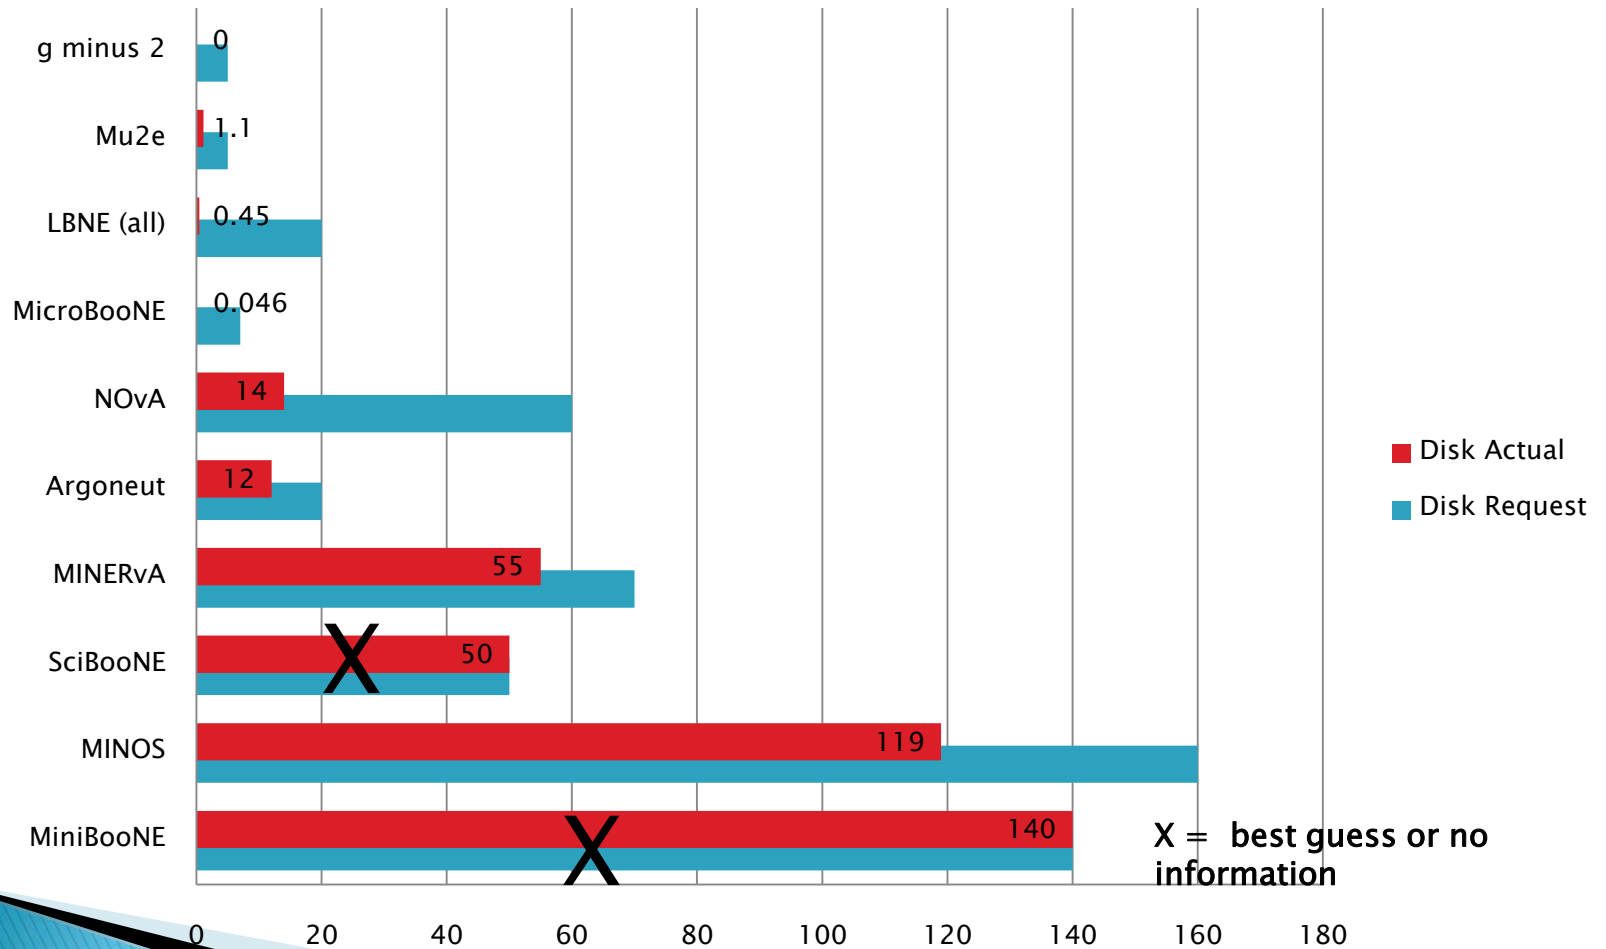
BlueArc App Areas

BlueArc /<exp>/app



BlueArc Disk usage vs. FY11 Request

DISK Usage: Actual and Requested


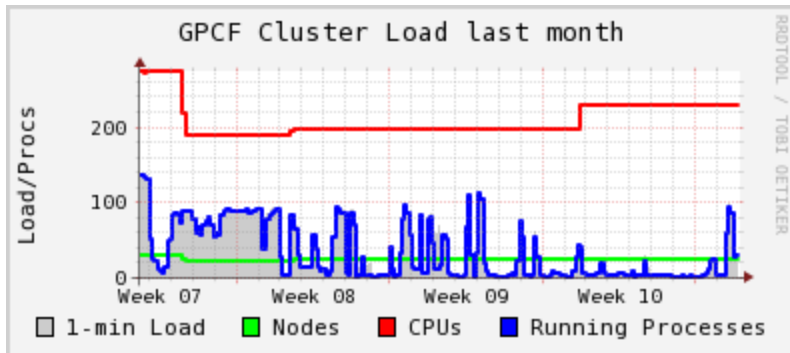




GPCF Status

	Int. Login status	Comments	Local Batch
ArgoNeuT	1 VM ✓		Jobs submitted through gpsn01 Currently 84 job slots: 7 worker nodes with 12 slots on each (limited by memory, plan to purchase). Four additional nodes to be added soon: FKA gpcf026,30,31,32.
LBNE	2 VMs ✓	Gpcf026 to gpwn pool	
MicroBooNE	1 VM ✓		
MINERvA	If01-if04 2 VM's ✓	If05 to be returned soon. (slf5?)	
MiniBooNE	(not assigned)		
MINOS	Minos50-54 ✓ 1VM (for testing)		
Mu2e	1 VM ✓ + 1 VM (for testing)		
NOvA	5 VM's ✓ + 2 x gpcf nodes temporarily	Gpcf028+29 until VM I/O issues resolved	


Ganglia Stats

- ▶ Linked from https://cdcv.s.fnal.gov/redmine/projects/ifront/wiki/Monitoring_Stats



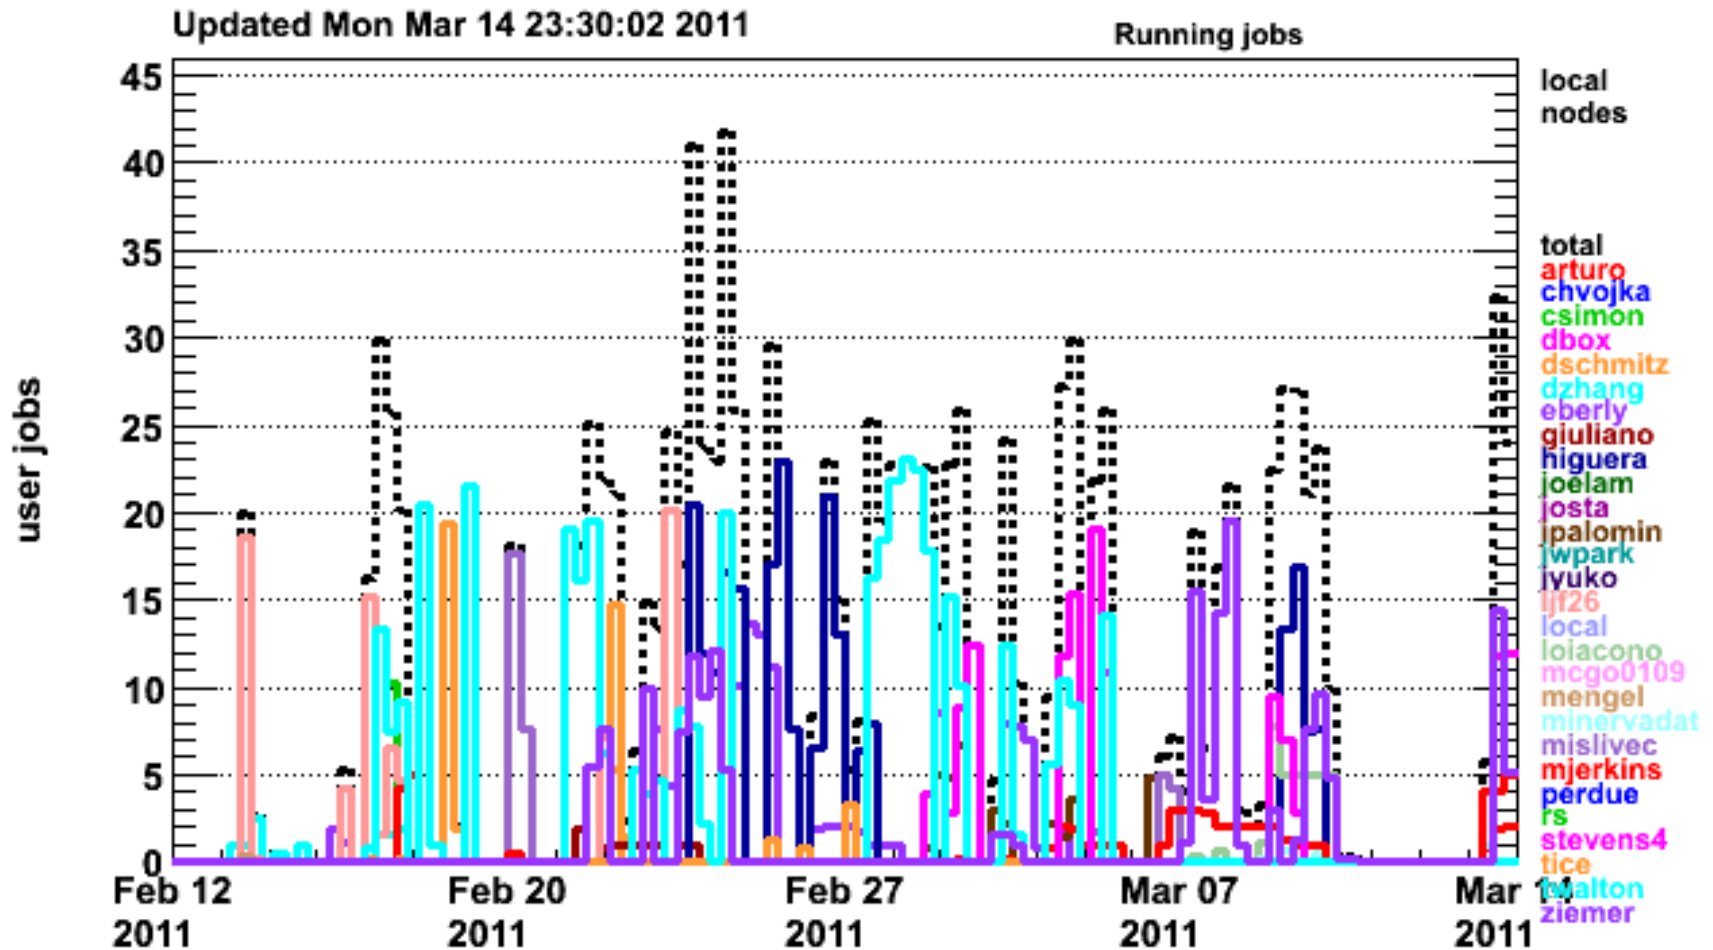
Batch Jobs gpsn01 Cluster

- ▶ http://gpsn01.fnal.gov/condor_monitor/index_month_grid.html



Batch Jobs on if01 Cluster


http://if01.fnal.gov:8080/condor_monitoring/index_month_local.html



24 total users

GPCF/IF/MINOS usage vs. FY11 Request


Int./Local Batch: Actual vs. Request



Includes rough estimates of gpsn01 batch usage

FermiGrid GP Cluster Last Month

fnpc4x2 - fermilab V0 Usage by group - Last Month




	Maximum	Average	Minimum	LastVal
total slots	3299	3276	0	3295
fermilab	2888	1200	0	65
fnal_acc	600	92	0	0
fnal_thy	499	2	0	0
fnalgrid	1010	146	0	56
astro	0	0	0	0
cdms	381	12	0	0
hypercp	0	0	0	0
ktev	0	0	0	0
minerva	520	155	0	1
miniboone	0	0	0	0
minos	2110	646	0	8
mipp	144	2	0	0
mu2e	500	115	0	0
nova	525	11	0	0
numi	92	0	0	0
patriot	10	0	0	0
storage	398	12	0	0
test	0	0	0	0
uboone	0	0	0	0

Data for fnpc4x2 between 06-Feb-2011 and 10-Mar-2011
 Plot generated on 10-Mar-2011 08:27:18

FermiGrid GP Cluster Last Quarter

fnpc4x2 - fermilab V0 Usage by group - Last Quarter




■ total slots
 ■ fnal_acc
 ■ fnal_thy
 ■ fnalgrid
 ■ astro
 ■ cdms
 ■ hypercp
 ■ ktev
 ■ minerva
 ■ miniboone
 ■ minos
 ■ mipp
 ■ mu2e
 ■ nova
 ■ numi
 ■ patriot
 ■ storage
 ■ test
 ■ uboone
 ■ fermilab

	Maximum	Average	Minimum	LastVal
total slots	3299	3002	611	3295
fermilab	2803	817	0	83
fnal_acc	600	46	0	0
fnal_thy	498	1	0	0
fnalgrid	1008	52	0	58
astro	0	0	0	0
cdms	553	49	0	0
hypercp	0	0	0	0
ktev	0	0	0	0
minerva	500	107	0	1
miniboone	0	0	0	0
minos	2000	459	0	7
mipp	105	1	0	0
mu2e	500	96	0	17
nova	525	4	0	1
numi	120	1	0	0
patriot	20	0	0	0
storage	250	5	0	0
test	0	0	0	0
uboone	0	0	0	0

Data for fnpc4x2 between 06-Dec-2010 and 10-Mar-2011
 Plot generated on 10-Mar-2011 08:27:21

GP Cluster Usage vs. FY11 Request


GRID: Actual vs. Request




Used last month's accounting for "Actual"

FY12 Requests (very uncertain)

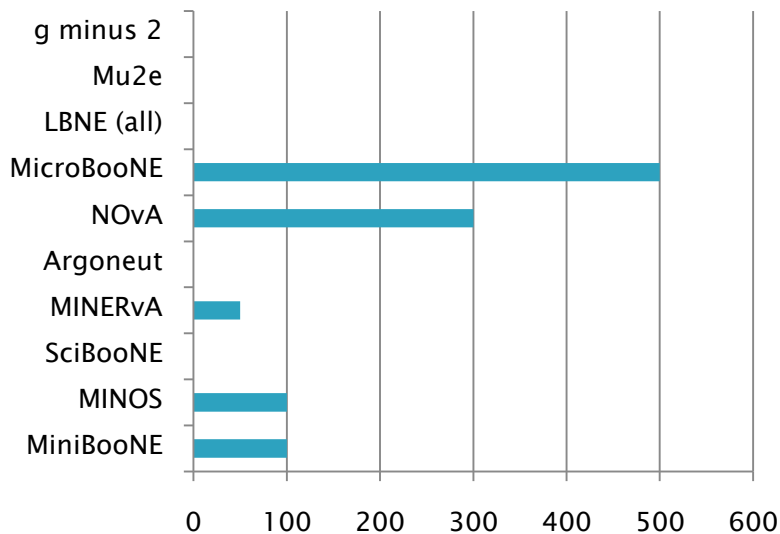
FY12 Disk Requests




FY12 Int/batch Requests



FY12 Tape Requests



FY12 GRID Requests



Other Thoughts and Conclusions

- ▶ Replacement hardware for EOL is also in budget.
- ▶ 10% cut scenario
 - Reduce grid replacements
 - Reduce BlueArc some
- ▶ 30% cut scenario, 10% cuts plus
 - Reduce added grid nodes by 25%
 - Reduce BlueArc by 20% (might be done with cheaper alternative).

finish