

NuComp News

Lee Lueking

September 14, 2011

Outline

- ▶ Maintenance day stuff
- ▶ IF Computing Resource Summary
- ▶ A look at FY12 Requests

Maintenance Day: Tomorrow 9/15

- ▶ The following systems will be rebooted for new kernel:
 - ▶ GPCF024, DBWEB3
MNVNEARLINE3, MNVNEARLINE4
MBDATA01, MBDB02
MINOS-MYSQL2, MINOS-SAM02
MIPPDBSRV01, MIPPDBSRV02
NUDATA03, NUDATA04
- ▶ FermiGrid expects to do a VDT upgrade on all three GP Grid gatekeepers on the Sep. 15 downtime. This is the last feature addition release of vdt 2.0.0 although there may be further patches. Expected downtime 1.5 hours. Running jobs won't be disturbed due to this work. New job submissions during the downtime will not start until after the work is completed.

Maintenance Day Next Week, and Next Month

- ▶ Thursday 9/22, 8–10AM. Transitioning to use the new BlueArc storage space. This will affect the following areas:
 - /minerva/data
 - /mu2e/app
 - /minerva/app
 - /nusoft/app
 - /lbne/app
 - /microboone/app
 - /gm2/app
- ▶ If these areas are used during the scheduled downtime, a stale mount may result. In some cases rebooting the affected machine(s) may be required. Jobs running on GP grid cluster are also affected and it is best to not have jobs running that employ any of these areas during this period. Condor queues will be drained Wed. night.
- ▶ Next Month 10/20: IF/GPCF/MINOS/MB reboots for new kernel

SLF4 Demise, and SLF6 Readiness

- ▶ SLF4 will no longer be supported after Feb. 2, 2012
- ▶ SLF6 is close to being ready for full deployment.
- ▶ (more details from Jason)

IFront Computing Resource Summary

» Central Storage
GPCF
GP Grid Cluster

Central Disk (BlueArc) Allocations

	2009	2010	2011 (replace)	2012 request
MiniBooNE	✓80	30	✓25 (+13)	0
MINOS	✓90	✓30	✓30 (✓30+30)	✓30(+50)*
MINERvA	✓10	✓20	✓40+✓70*	40
Argoneut	✓20	0	0	0
NOvA	✓4	✓26	✓30	100
MicroBooNE	✓2	0	✓5	30
LBNE (all)	0	✓2	✓10+✓10	0
Mu2e	0	✓2	✓2	0
g minus 2	0	0	✓10	0

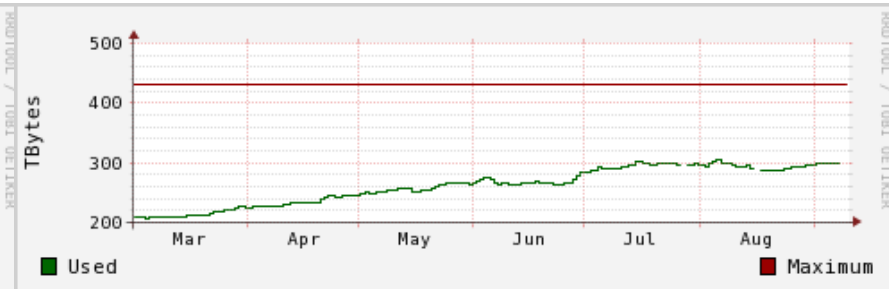
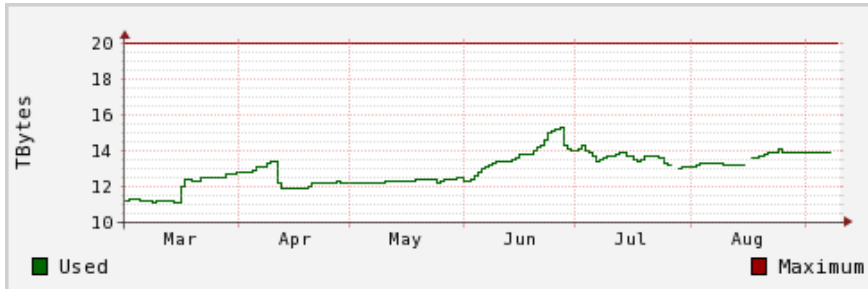
✓ satisfied
 ✓ in-hand
 * FY12
 buy-
 ahead

Disk Usage Trends (1 / 3)

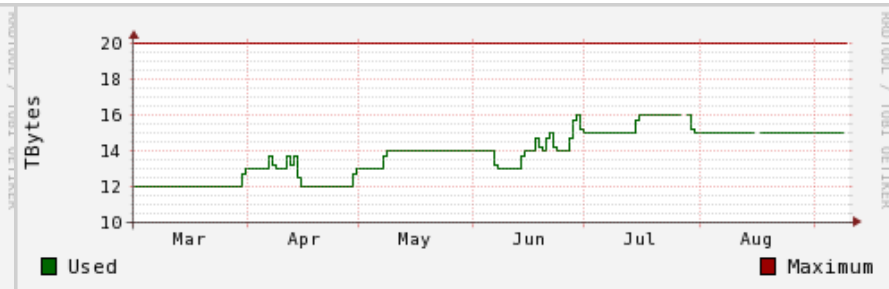
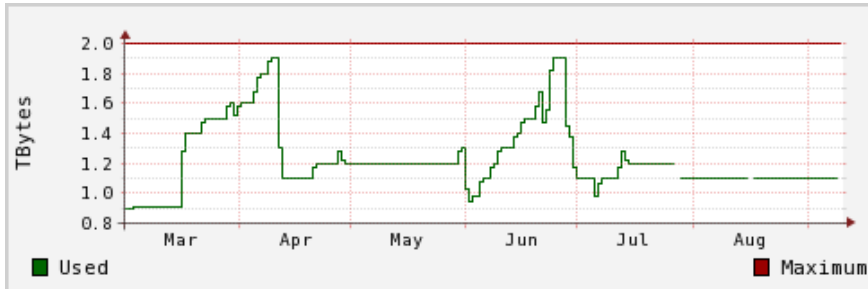
/app

/data

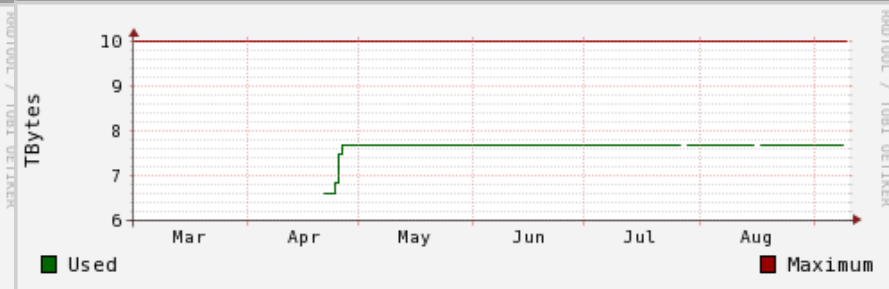
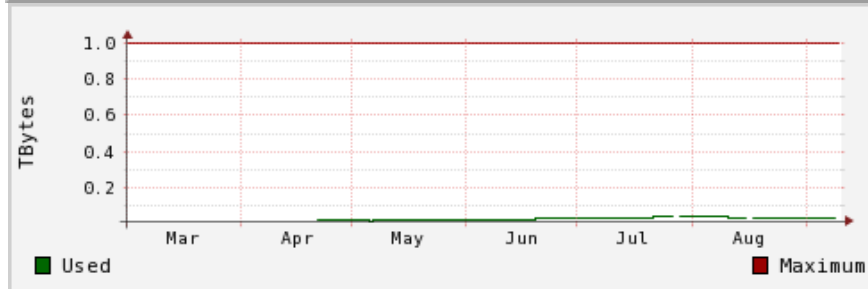
All



/argoneut



/gm2

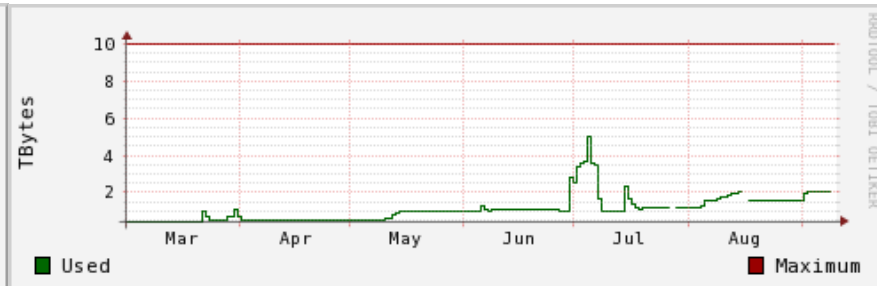
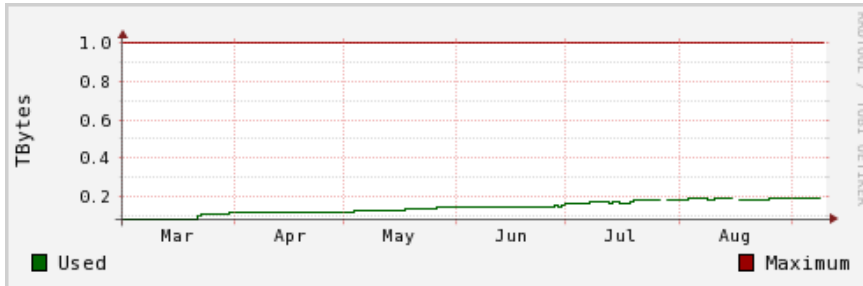


Disk Usage Trends (2 / 3)

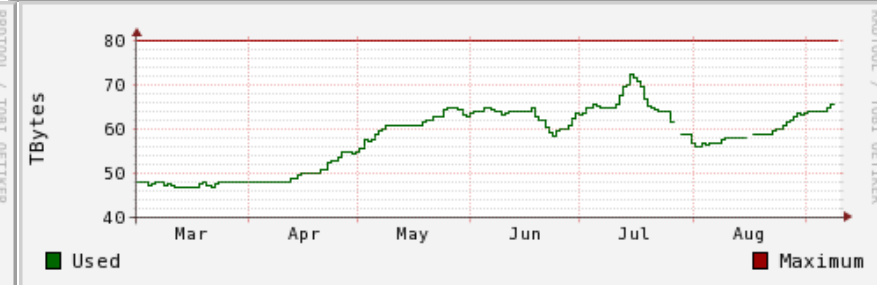
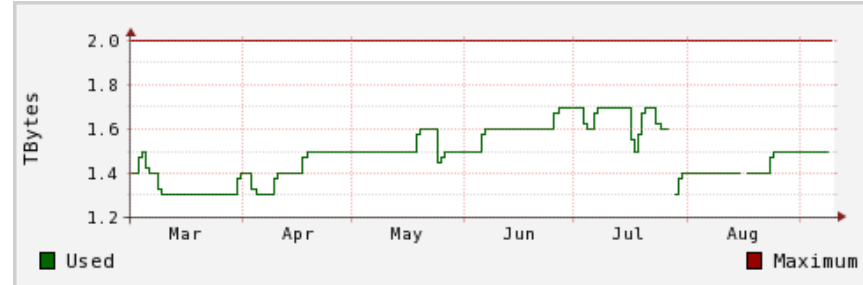
/app

/data

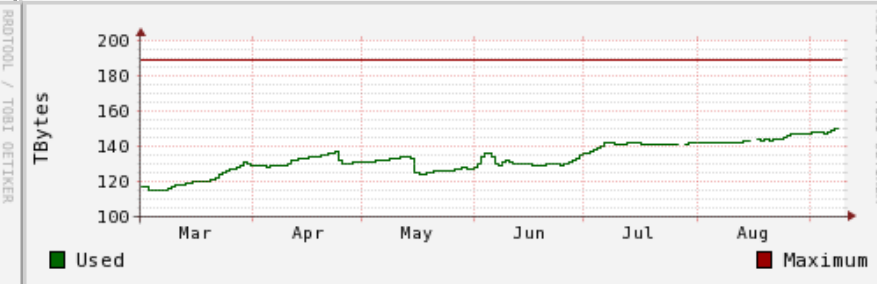
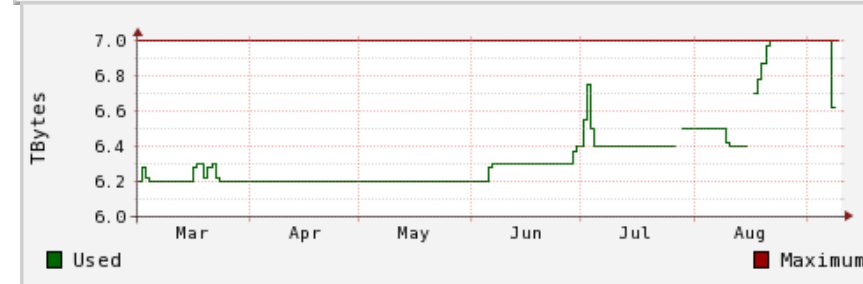
/lbne



/minerva



/minos

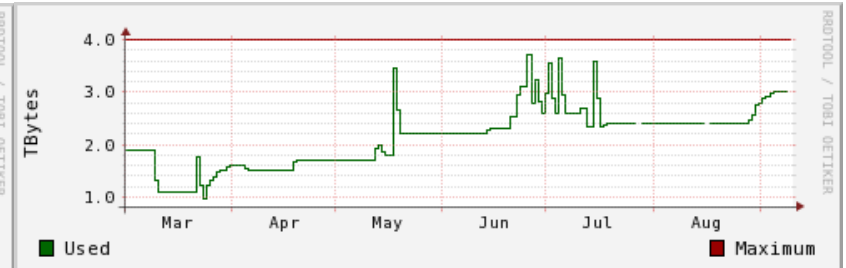
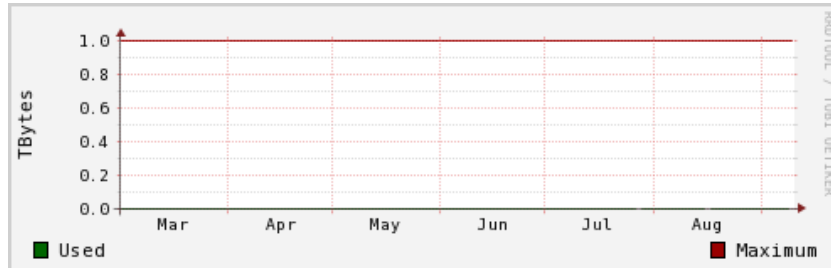


Disk Usage Trends (3 / 3)

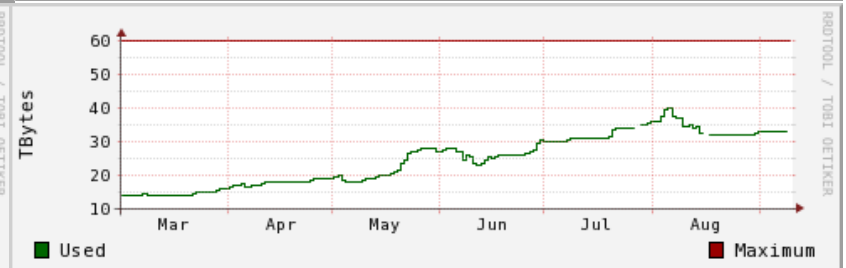
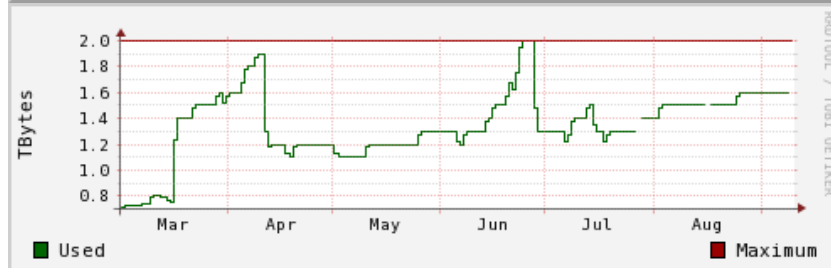
/app

/data

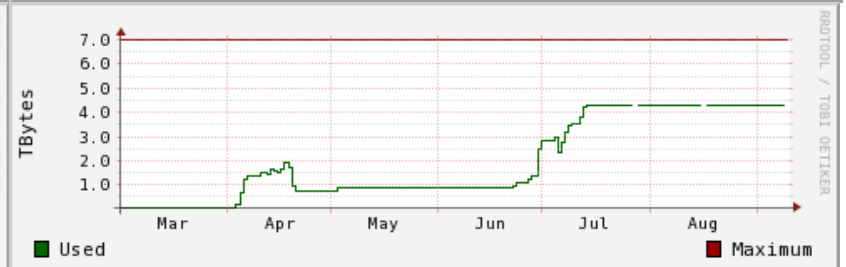
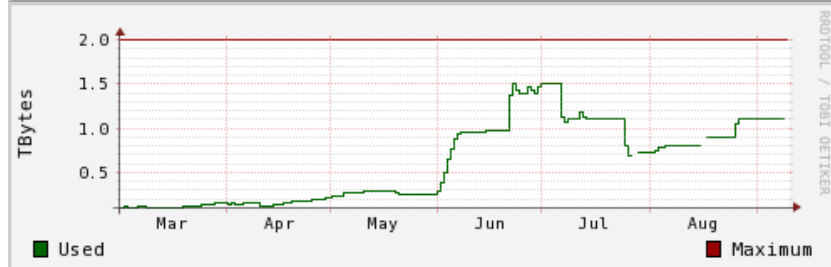
/mu2e



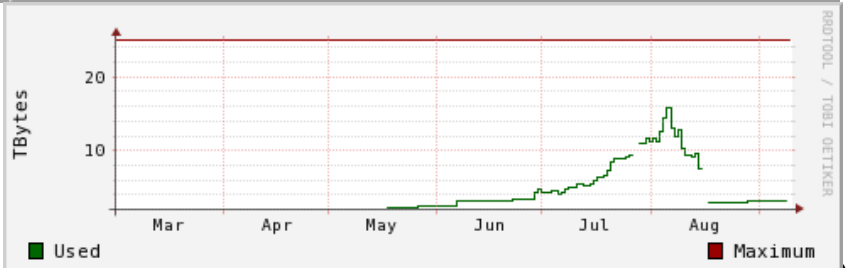
/nova



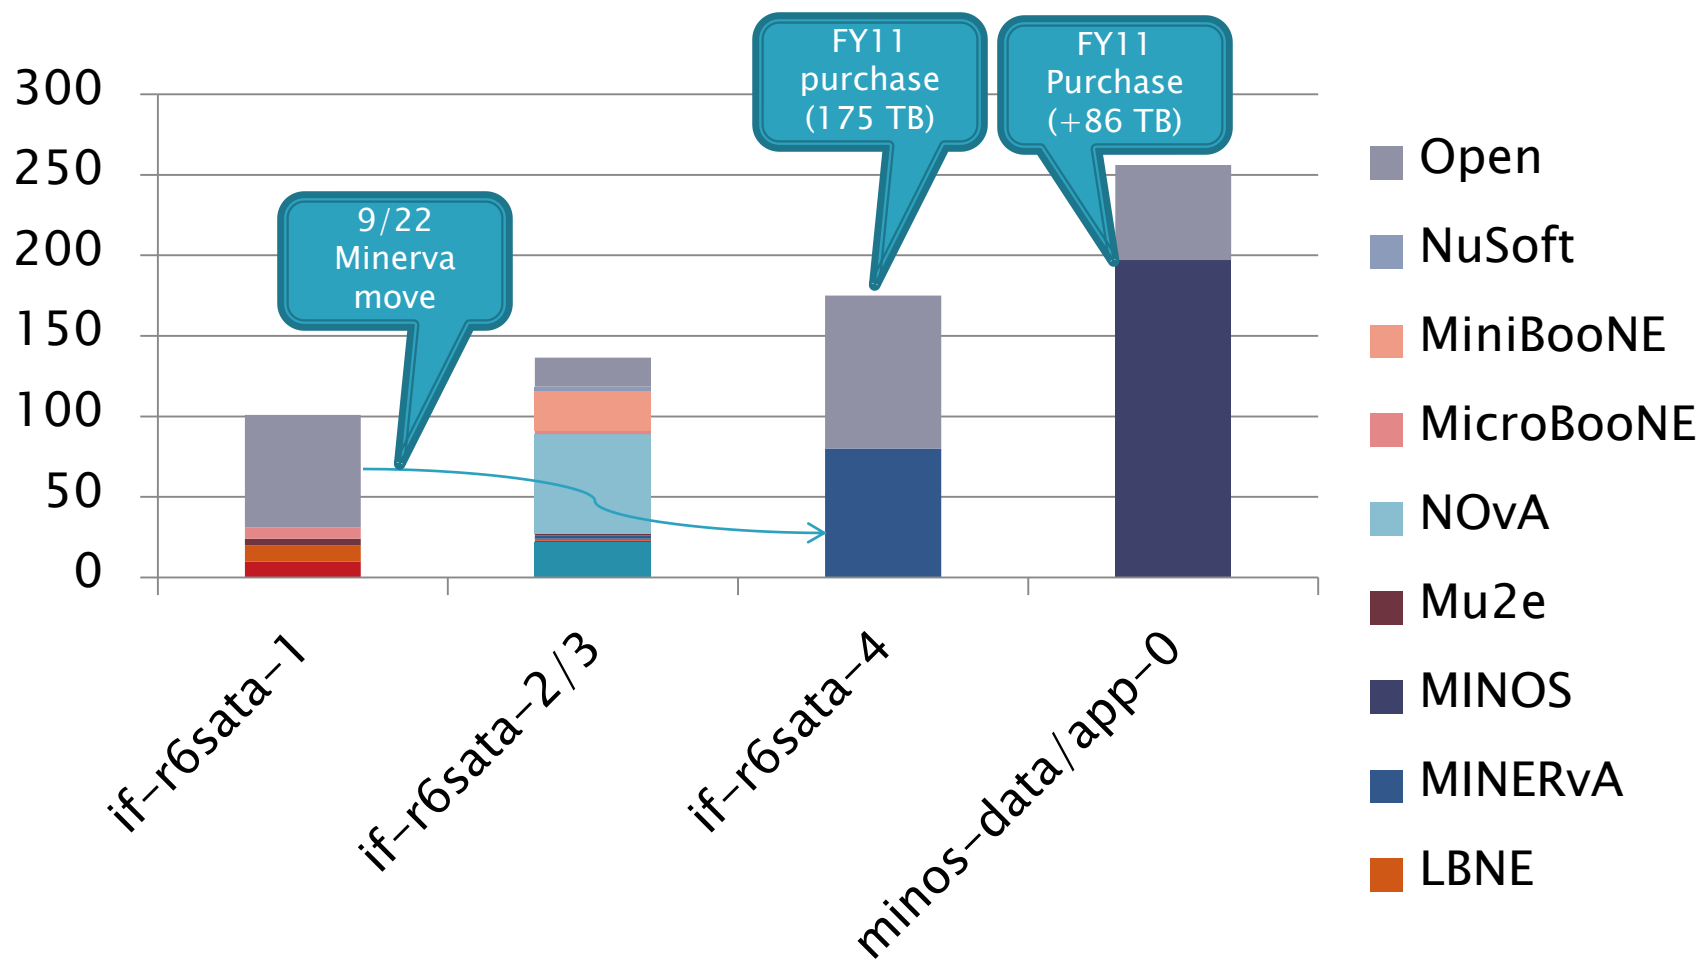
/microboone



/miniboone



Experiment Assignment to Disk Pools



Includes *app* + *data* areas

GPCF Status

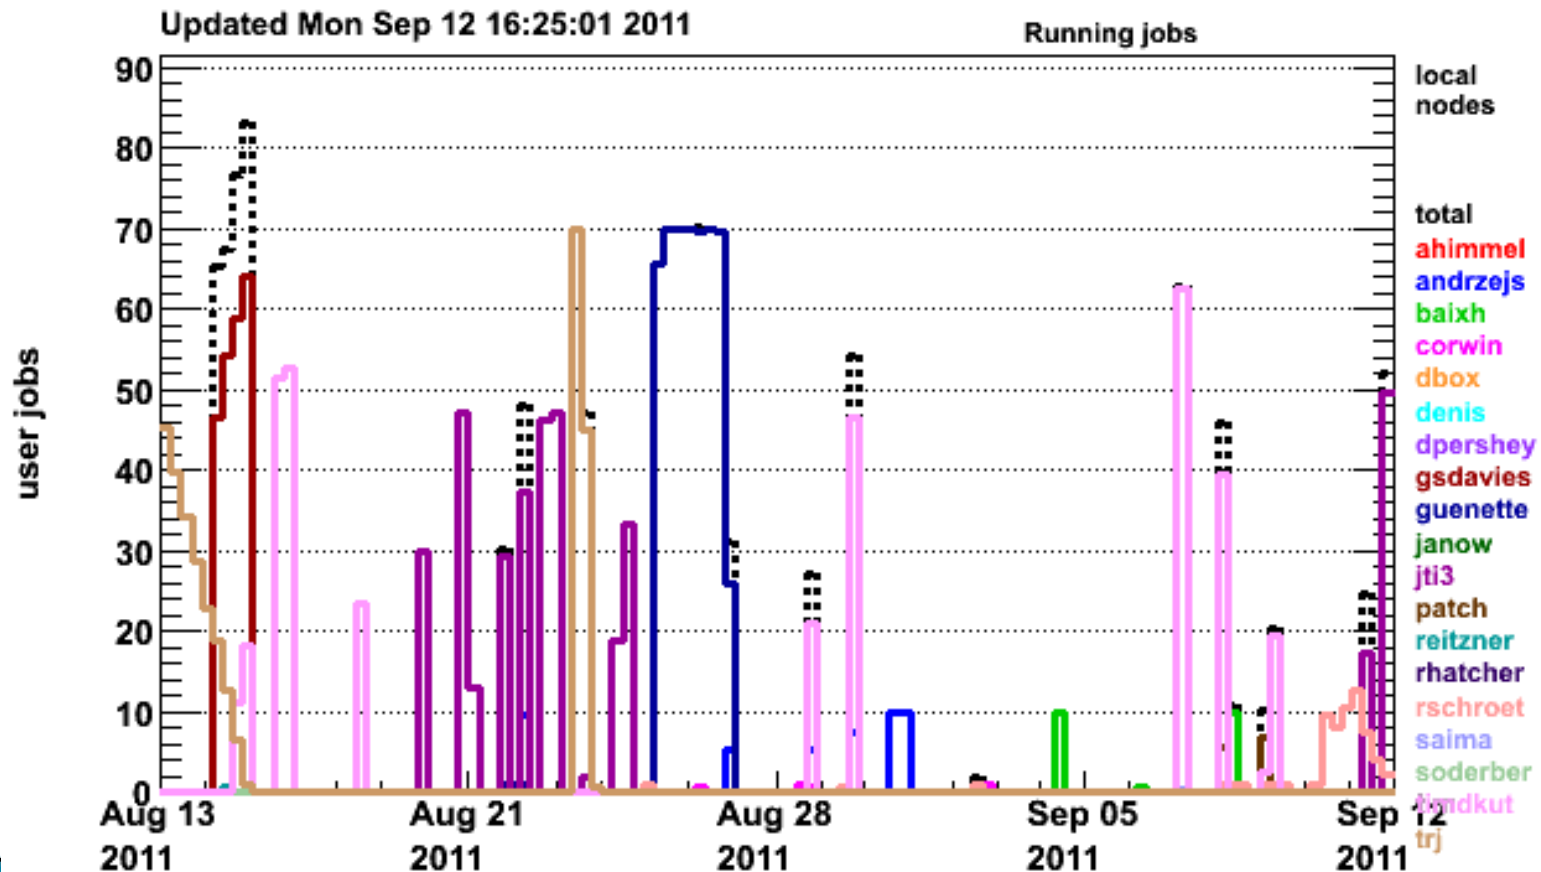
	Int. Login status	Comments	Local Batch
ArgoNeuT	1 VM		Jobs submitted through gpsn01
gm2	1 VM		
LBNE	1 VMs		
MicroBooNE	2 VM		Currently 7 worker nodes: 48 x 2GB slots 18 x 4GB slots
MINERvA	If01-if05 3 VM		
MiniBooNE	(not assigned)		
MINOS	Minos50-54 1VM (for testing)		On deck: gpscf026,30,31,32.
Mu2e	2 VM		
NOvA	5 VM 2 x gpscf nodes		

Ganglia Stats

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

Batch Jobs gpsn01 Cluster

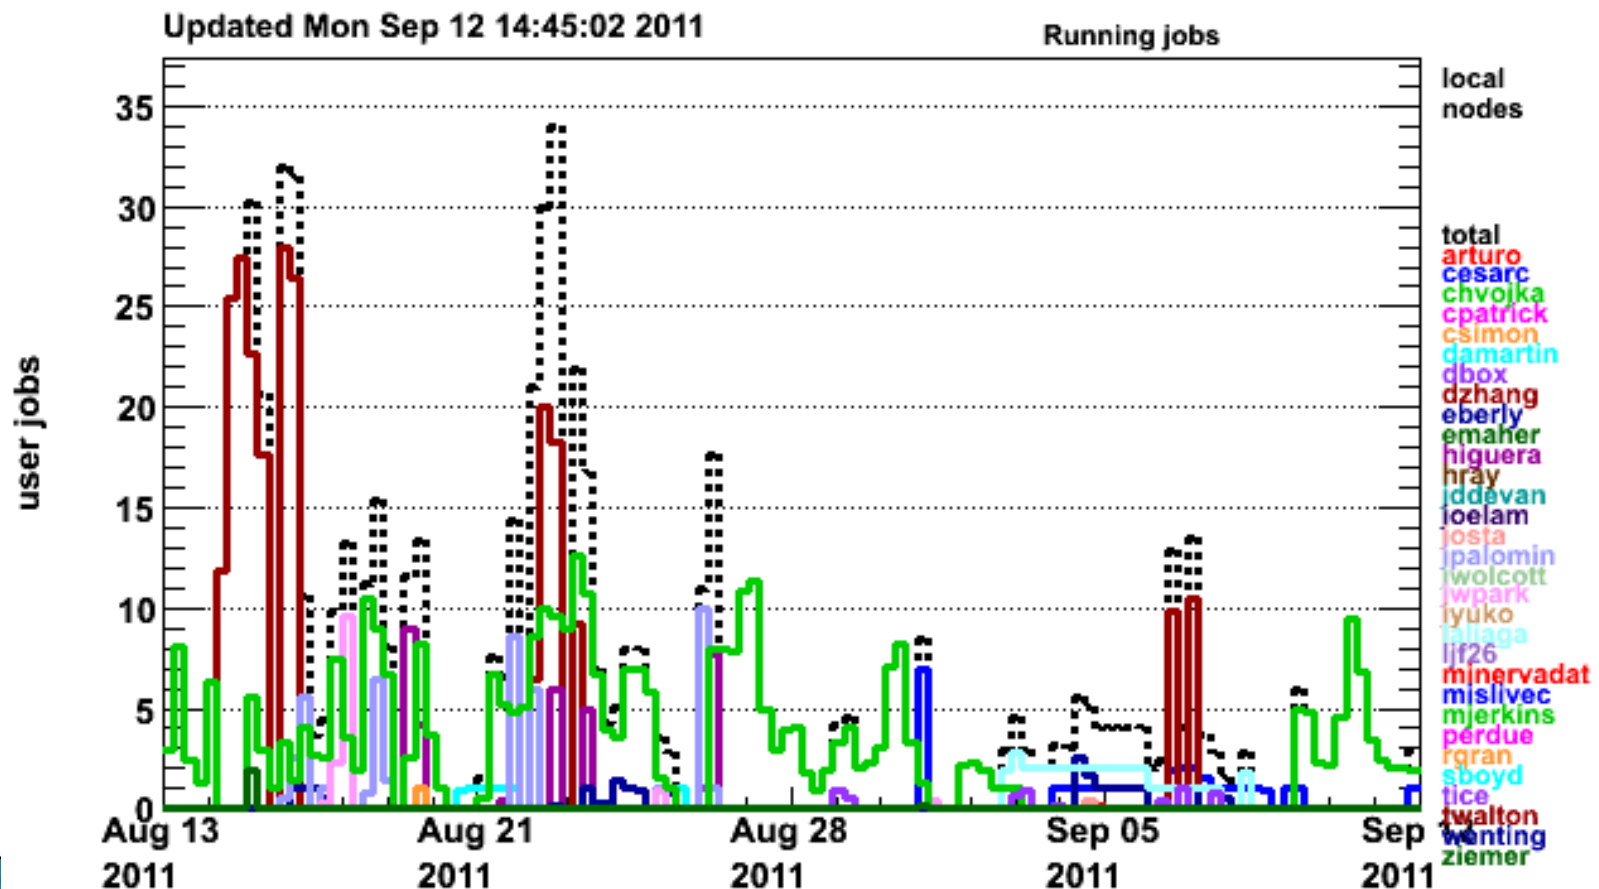
- ▶ http://gpsn01.fnal.gov/condor_monitor/index_month_grid.html



70 total slots

Batch Jobs on if01 Cluster

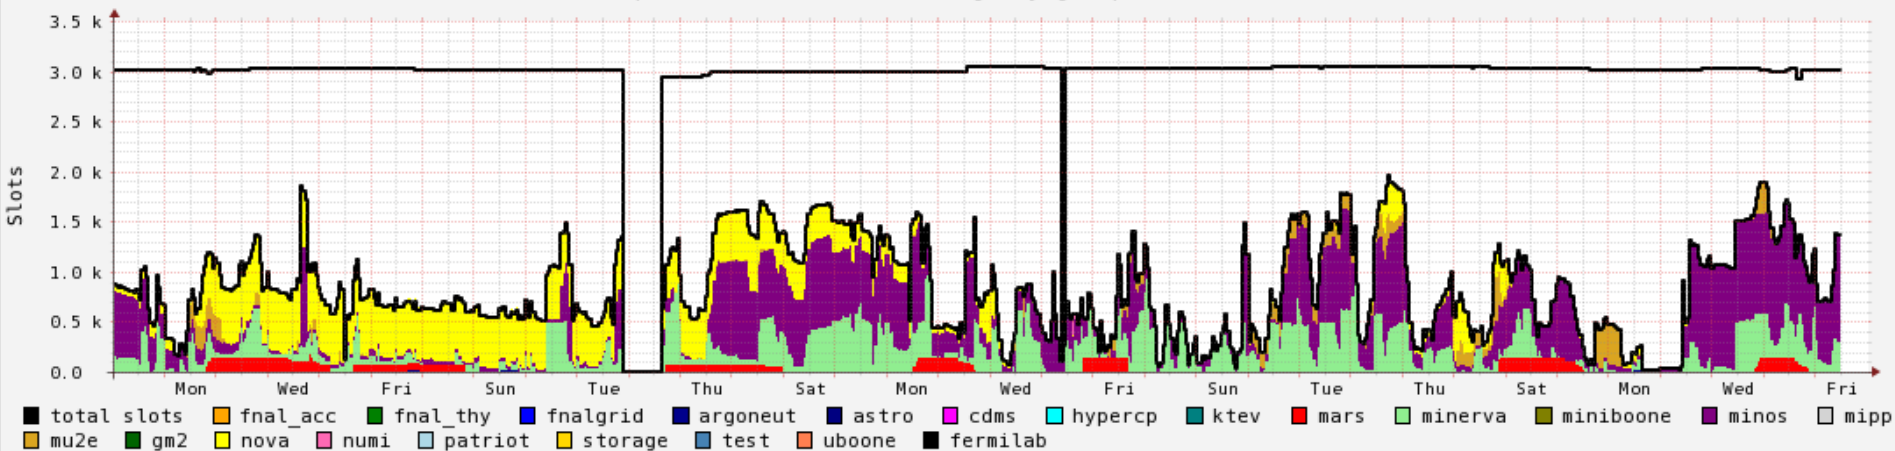
http://if01.fnal.gov:8080/condor_monitoring/index_month_local.html



30 total slots

GP Grid Cluster Use; Last Month

fnpcl1x2 - fermilab V0 Usage by group - Last Month



	Maximum	Average	Minimum	LastVal
total slots	3042	2944.02	0	3010
fermilab	1955	826.12	0	1360
fnal_acc	0	0.00	0	0
fnal_thy	0	0.00	0	0
fnalgrid	2	0.00	0	0
argoneut	25	1.05	0	1
astro	0	0.00	0	0
cdms	0	0.00	0	0
hypercp	0	0.00	0	0
ktev	0	0.00	0	0
mars	174	36.68	0	0
minerva	932	203.04	0	297
miniboone	0	0.00	0	0
minos	1056	345.81	0	1056
mipp	0	0.00	0	0
mu2e	437	35.22	0	0
gm2	0	0.00	0	0
nova	727	204.31	0	7
numi	209	1.06	0	0
patriot	0	0.00	0	0
storage	0	0.00	0	0
test	0	0.00	0	0
uboone	0	0.00	0	0

	Maximum	Average	Minimum	LastVal
total slots	3042	2944.02	0	3010
fermilab	1955	826.12	0	1360
fnal_acc	0	0.00	0	0
fnal_thy	0	0.00	0	0
fnalgrid	2	0.00	0	0
argoneut	25	1.05	0	1
astro	0	0.00	0	0
cdms	0	0.00	0	0
hypercp	0	0.00	0	0
ktev	0	0.00	0	0
mars	174	36.68	0	0
minerva	932	203.04	0	297
miniboone	0	0.00	0	0
minos	1056	345.81	0	1056
mipp	0	0.00	0	0
mu2e	437	35.22	0	0
gm2	0	0.00	0	0
nova	727	204.31	0	7
numi	209	1.06	0	0
patriot	0	0.00	0	0
storage	0	0.00	0	0
test	0	0.00	0	0
uboone	0	0.00	0	0

Data for fnpcl1x2 between 07-Aug-2011 and 09-Sep-2011

Plot generated on 09-Sep-2011 10:27:03

FY12 Needs

- » The numbers
- Notes and thoughts

A Look at the FY12 Needs

Based on planning Spreadsheet

Experiment	Disk (TB)	Tape (TB)	Int. Login (4 core/VM)	Local Batch (cores)	GRID (cores)	Servers
MINOS	0* (30)	100	0	0	0	0* (3)
NOvA	100	300	5	20	300	0
MINERvA	40	50	5	20	200	0* (3)
Mu2e	5	0	0	0	30	0
LBNE	10	0	4	20	50	0
MicroBooNE	23	50	0	20	0	0
ArgoNeuT	0	0	0	0	0	0
MiniBooNE	15	100	0	0	0	0
Muon g - 2	0* (5)	0	1	5	0	0
(replace)	(?)			40	286	
TOTAL	193	600	15	125	866	0
TOTAL Cost	\$k	\$k	\$k	\$k	\$k	\$k

* Items pre-purchased in FY11 with the requested quantity in ().

FY12 Needs: Notes

- ▶ There was significant disk (150TB) purchased in FY11 that was a “pre-buy” for FY12.
 - Covers MINOS needs for FY12,13.
 - Provides Minerva an additional 70TB “cushion”.
- ▶ It appears that disk prices for FY12 may be significantly lower than FY11 (\$350/TB vs. \$800/TB). This may give us some flexibility:
 - Pre-buy portion of FY13 (FY13 total estimate is 320TB).
 - Provide additional storage for experiments.
 - Use disk archiving instead of tape.
- ▶ Assumptions for Local Batch cores,
 - We are retiring/repurposing IF nodes,
 - Doubling capacity of the Local Batch
- ▶ GP Grid resources released from Run II may be sufficient to meet most of our needs.

October NuComp



October NuComp

- ▶ October 12 (or 19)
 - (We could consider a new day/time if needed.)
- ▶ Topics
 - Condor DAGs– Submitting and running *Directed Acyclic Graph* jobs – Rick Snider
 - GPCF v 2.0 – plans for moving & reorganizing GPCF– Jason Allen
 - IF Beam Data – Status report on IF Beam Data Project – Igor Mandrichenko, Andrew Norman
 - gridFTP service – **Globus Online Integrated with glideinWMS** – Parag Mhashilkar, Dennis Box

finish

Short Term Plans



GPCF

- ▶ Replace existing worker nodes (7 nodes) with “farm” type nodes. Increase number of WN’s as needed.
- ▶ Increase number of VM host nodes with existing (7 worker + 4 unused) gpcf nodes.

Miscellaneous Servers

- ▶ General-IF-DB server plan
 - General purpose for conditions data, ECL, etc.
 - Three large, 16 core, 64GB Memory, 2TB data disk
 - Dev: Develop and test applications
 - Prod: Production service
 - Farm Replica: Replica for scaling to farm clients.
 - Both MySQL and PostgreSQL instances
 - Independent “clusters” for each experiment
- ▶ MINERvA: One event display machine, one VM server.