

Group Space Management & Group Job Monitoring

Ueda I.
ATLAS S&C workshop
2011.04.05.

Group Space

Global quota vs Specific sites

Questions asked:

- A) which groups would like to utilize the global quota and automatic choice of destination
- B) which groups would like to keep spaces at specific sites, and if so, which sites
- C) which groups would like the mixture of the two

Tentative Results:

- A: 4, B: 2, C: 7

The DDM-dev has been consulted to implement C

- to satisfy all the groups (A + B + C)
- it is a big change in the system, and will take several months for implementation

Group space quota in 2011

We have asked the CREM to confirm the quota in 2011

- hopefully will be increased ~ next week

While waiting for the implementation of the “global quota”

- Should we move the group spaces from T2s to T1s for some groups who goes A and C?
 - ▶ specified T2 spaces will be kept
 - ▶ the total volume per group will be within the 2011 quota

Panda Dynamic Data Placement (PD2P)

- There is a plan to activate PD2P for group data
 - ▶ only for ‘Group Production’ data? or anything in the Group Spaces?
- the replicas will be created in T2_DATADISK
 - ▶ will not eat the group space quota

Group Job Monitoring

Definition of “groups” – Task Request



Task (data10_7TeV.00167607.physics_Egamma.merge.r1774_p327_p333_p491) Request Parameters

Input Attributes

Project
 Input dataset
 Transformation
 Transformation Version
 Transformation Cache
 Transformation Parameters
 Values
 Total Number Of Output Events
 First File Number in Input Dataset
 Total Input Files
 Number Of Events for Output File
 CPU per Event
 Memory Usage
 Grid Flavour
 E-mail

data10_7TeV
 data10_7TeV.00167607.physics_Egamma.merge.AOD.r1774_p327_p333
 Reco_trf.py
 16.0.3.8.3
 TopPhys
 --athenaopts,--ignoreerrors,--omitvalidation,DBRelease,autoConfiguratio
 NONE,True,ALL,15.1.2,everything,NONE,NONE,NONE,NONE,NONE,NON
 1684000
 1
 1684
 1000
 10
 3500
 panda@it
 Junji.Tojo@cern.ch

Output Attributes

Output Task Name :
 Output Datasets
 Priority
 SW Release :
 State
 Request Time :
 Comments
 Project Mode
 Modification Time :
 Modified by
 Request Id :
 Total_req_jobs :
 Total_done_jobs :
 Total_avail_jobs :
 Physics group :
 Queue :
 Ctag :
 Bug report :

data10_7TeV.00167607.physics_Egamma.merge.r1774_p327_p333_p491
 DAOD_ELLOOSE18
 650
 16.0.3
 running
 Mon Apr 4 17:30:32 2011
 group production SM WZ
 spacetoken=ATLASDATADISK;reprocessing=yes;lumiblock=yes;diskcount=6
 Never
 None
[292776](#)
 170
 6
 1684
 gr_SM
 default
 p491
 N/A

Total_avail_jobs :
 Physics group :
 Queue :
 Ctag :

1684
 gr_SM
 default
 p491

The group names are defined in the Task Requests

Definition of “groups” — pathena



How to make group defined outDS

```
$ pathena --official --voms atlas:/atlas/groupName/Role=production --outDS group.groupName.[otherFields].dataType.Version ...
```

where groupName for SUSY is phys-susy, for instance. See the document ATL-GEN-INT-2007-001 for dataset naming convention. Note that you need to have the production role for the group to produce group-defined datasets. If not, please request it in the ATLAS VO registration page. For example, to produce datasets for Higgs working group,

The group names are defined in VOMS and specified in the pathena command line

atlas VO Registration

/atlas/det-indet	production	Approved
/atlas/det-larg	production	Approved
/atlas/det-muon	production	Approved
/atlas/det-tile	production	Approved
/atlas/perf-egamma	production	Approved
/atlas/perf-flavtag	production	Approved
/atlas/perf-idtracking	production	Approved
/atlas/perf-jets	production	Approved
/atlas/perf-muons	production	Approved
/atlas/perf-tau	production	Approved
/atlas/phys-beauty	production	Approved
/atlas/phys-exotics	production	Approved
/atlas/phys-gener	production	Approved
/atlas/phys-hi	production	Approved
/atlas/phys-higgs	production	Approved
/atlas/phys-lumin	production	Approved
/atlas/phys-sm	production	Approved
/atlas/phys-susy	production	Approved
/atlas/phys-top	production	Approved
/atlas/trig-daq	production	Approved

Group activities on the Panda Monitor



Servers: **CERN:OK** **Logger:OK**

Active tasks: CA:6 CERN:2 DE:25 ES:9 FR:67 IT:17 ND:7 NL:18 TW:3 UK:14 US:30
 Bamboo [task brokerage](#), [job submissions](#), [status](#) over last 12 hours

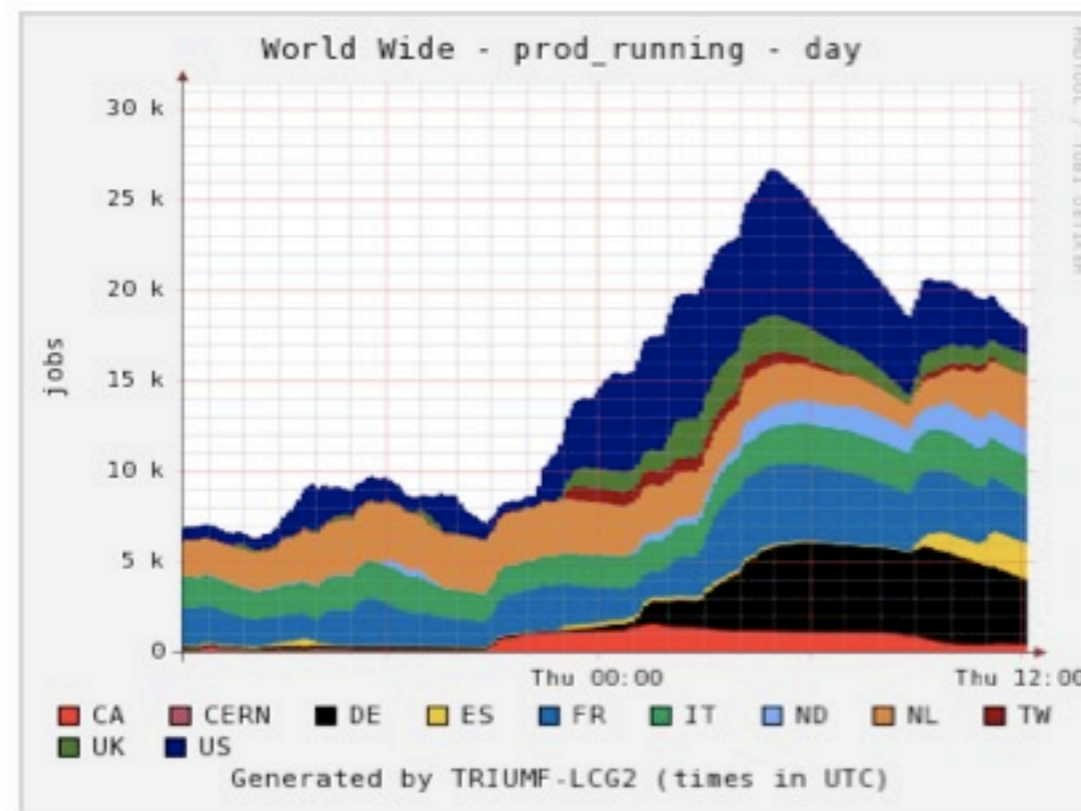
Jobs updated >12 hrs ago: **activated:** [1193](#) **running:** [none](#)
 Jobs updated >36 hrs ago: **transferring:** [none](#)

[Summary plots by cloud](#)

[Cloud efficiency history](#)

Ganglia World Wide Summary

More [plots](#): jobs in last



Panda Monitor shows the on-going activities

- per processing type
- per group
- per user

Production job summary at 03-31 12:13, last 12 hours (Details: [errors](#), [nodes](#))

~~Processing types: evgen(82) merge(2155) pile(263) prod_test(10) reco(1419) reprocessing(6835) simul(48679) validation(2345)~~

~~Working groups: BPhysics(2) Exotics(52) Generators(80) HeavyIon(176) Higgs(21337) Physics(8970) Reprocessing(5421) SM(2) Susy(1) Top(20449) Validation(2206) gr_Detector(16) gr_SM(2800) gr_Susy(22) gr_Top(5) heavyion(20) higgs(52) validation(177)~~

~~Users: Jinji.Tojo(28) areinsch(1327) borut.kersevan(51117) c.gwenlan(80) douglas(10) katarzp(22) leonardo.carminati(10) menke(16) r.henderson(2327) strandbe(5411) yhj(1440)~~

Panda production monitoring page shows group production activities as “gr_Group” as defined in the task request

Panda analysis monitoring page shows pathena jobs for groups with VOMS group names

Analysis job summary at 03-31 12:10, last 12 hours (Details: [errors](#), [nodes](#)) [pathena analysis queue status](#)

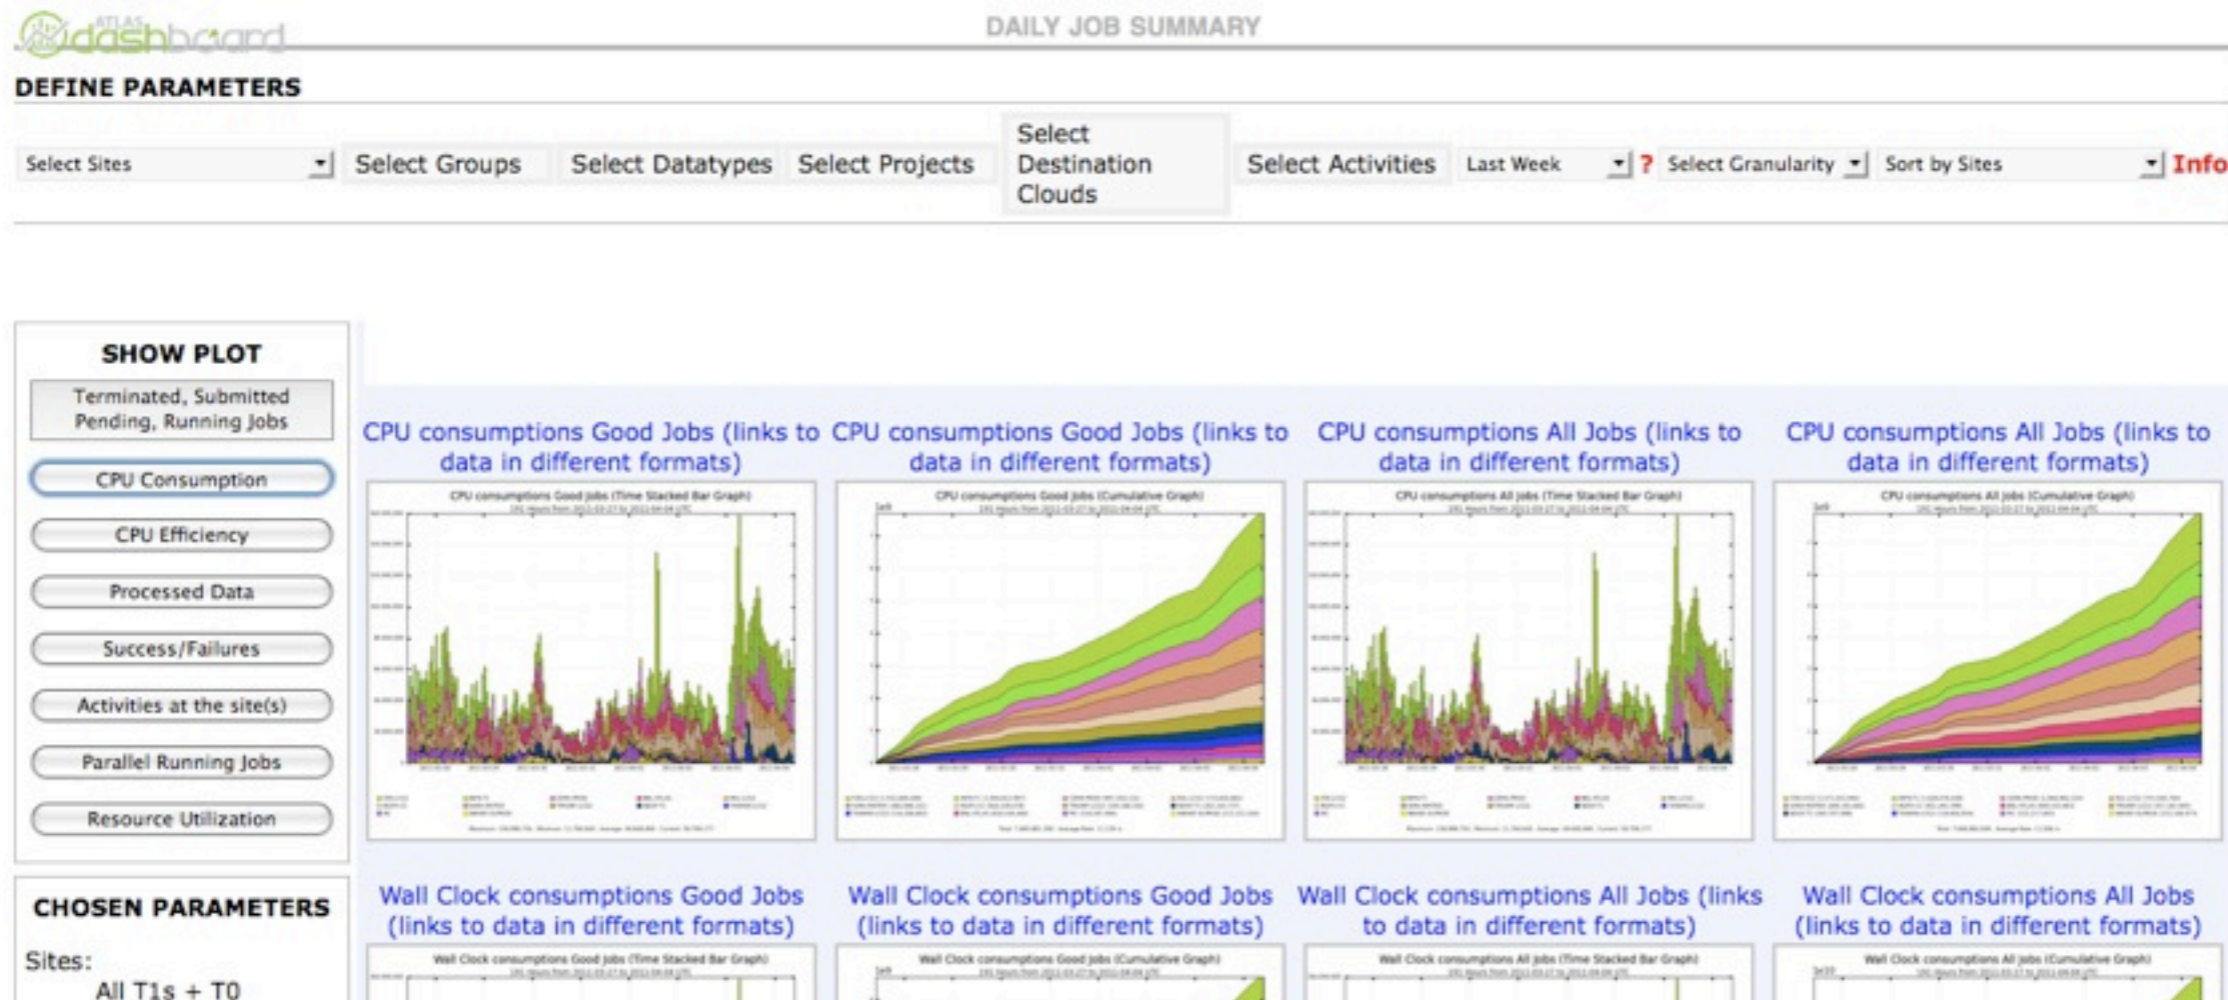
~~Processing types: ITB_Integration(320) ganga(3042) gangarobot(3189) hammercloud(5230) pathena(32839) prun(65590) usermerge(1)~~

~~Working groups: calib-muon(580) det-indet(451) det-muon(1002) perf-jets(84) perf-tau(1452) phys-higgs(3) phys-sm(5740)~~

~~Special Handling: localpool(301)~~

New Monitoring available

<http://dashb-atlas-job-prototype.cern.ch/dashboard/request.py/dailysummary>



Daily

Sort by:
Activities

Report a bug or a suggestion

Still under development

- please report when you see something wrong
- or if you have suggestions



Group production monitoring

S. Jézéquel (LAPP)



- ▶ **New version of 'Historical view' monitoring for ATLAS Grid CPU activity**
<http://dashb-atlas-job-prototype.cern.ch/dashboard/request.py/dailysummary>
- ▶ **Now usefull for group production monitoring**
- ▶ **Still under validation but many checks already done**

- ▶ **Informations collected since 3 weeks**

- ▶ **Detailed presentation in 'Monitoring session' (in parallel to this session)**



DEFINE PARAMETERS

Select Sites | Select Groups | Select Datatypes | Select Projects | Select Destination Clouds | Select Activities | Last Week | ? Select Granularity | Sort by Sites | Info

SHOW PLOT

- Terminated, Submitted Pending, Running Jobs
- CPU Consumption
- CPU Efficiency
- Processed Data
- Success/Failures
- Activities at the site(s)
- Parallel Running Jobs
- Resource Utilization

CHOSEN PARAMETERS

Sites: All T1s + T0

Groups: all

Datatypes: all

DAILY JOB SUMMARY

24h, 48h, week, month, custom | Hourly, Daily, Weekly, Monthly

All Projects

- data09_900GeV
- data10_7TeV
- data10_7TeV*physics_JetTauEtmis
- data10_cos
- data10_cvalid
- data10_hi
- data11_2p76TeV
- data11_7TeV
- data11_calib

Done | Cancel

Sort by Sites

- Sort by Sites
- Sort by Activities
- Sort by Generic Activities
- Sort by Clouds
- Sort by Datatypes
- Sort by Groups
- Sort by Projects
- Sort by DestinationClouds

5-Apr-11

ATLAS Software & Computing Workshop

4

Nomenclature for group monitoring



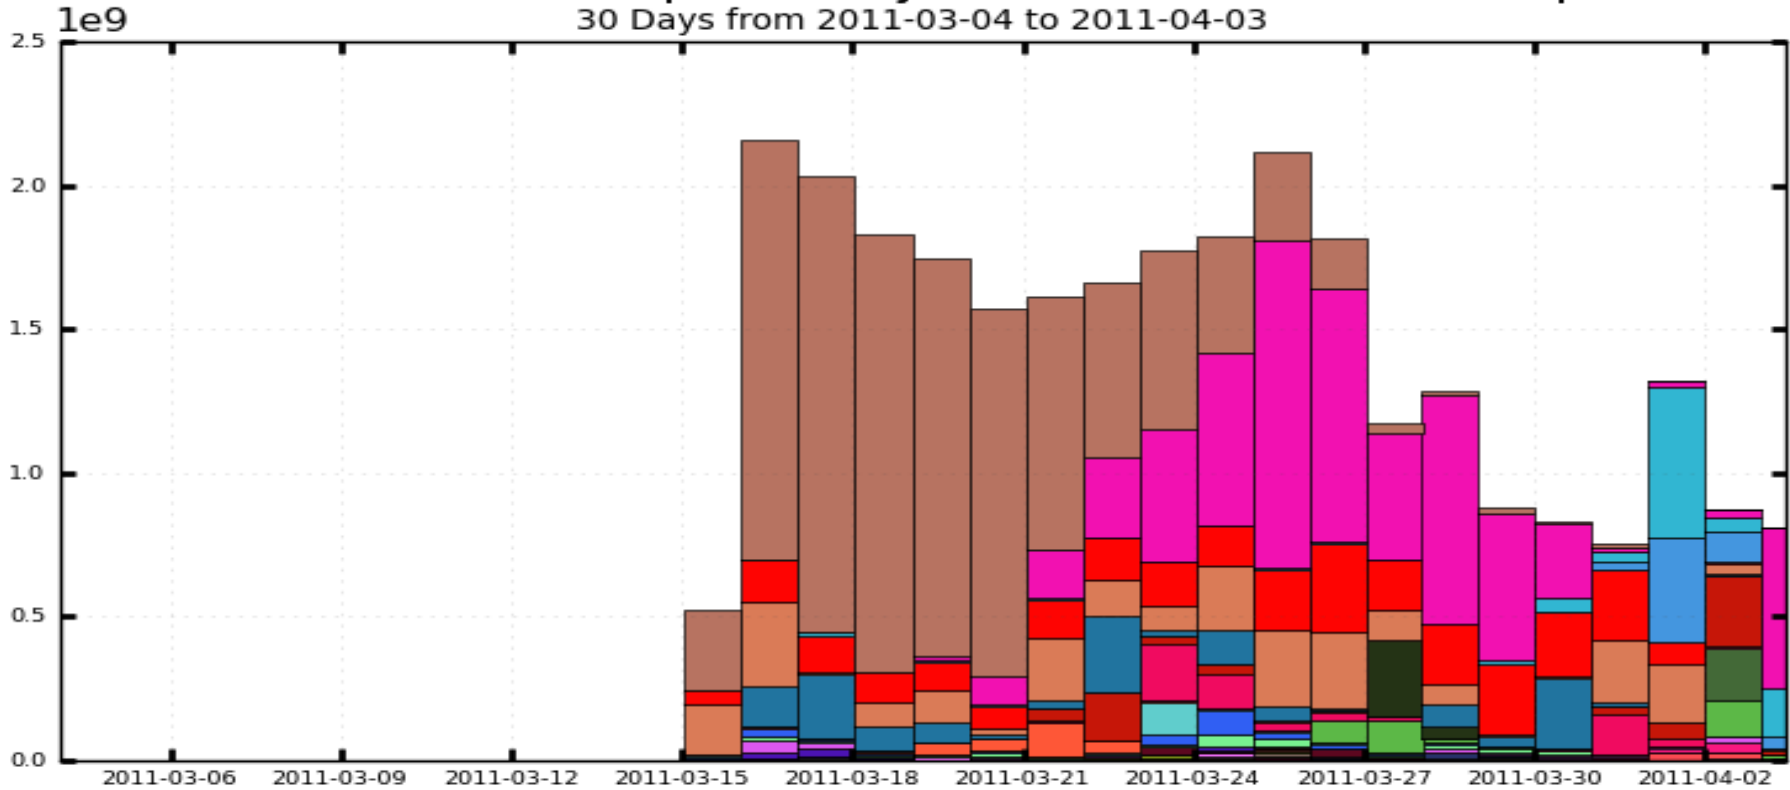
- ▶ Naming convention for group activity
 - ▶ Nomenclature defined in Panda
- ▶ Related to voms group:phys-*,perf-*,...
 - ▶ Pathena jobs with group priority
- ▶ Starts with gr_
 - ▶ Central production for groups :
- ▶ Localuser/user : pathena/ganga jobs without specific priority
- ▶ Remaining part :
 - ▶ Validation, reprocessing, Physics, ... : Central production



Global results(1)

Wall Clock consumptions All Jobs (Time Stacked Bar Graph)

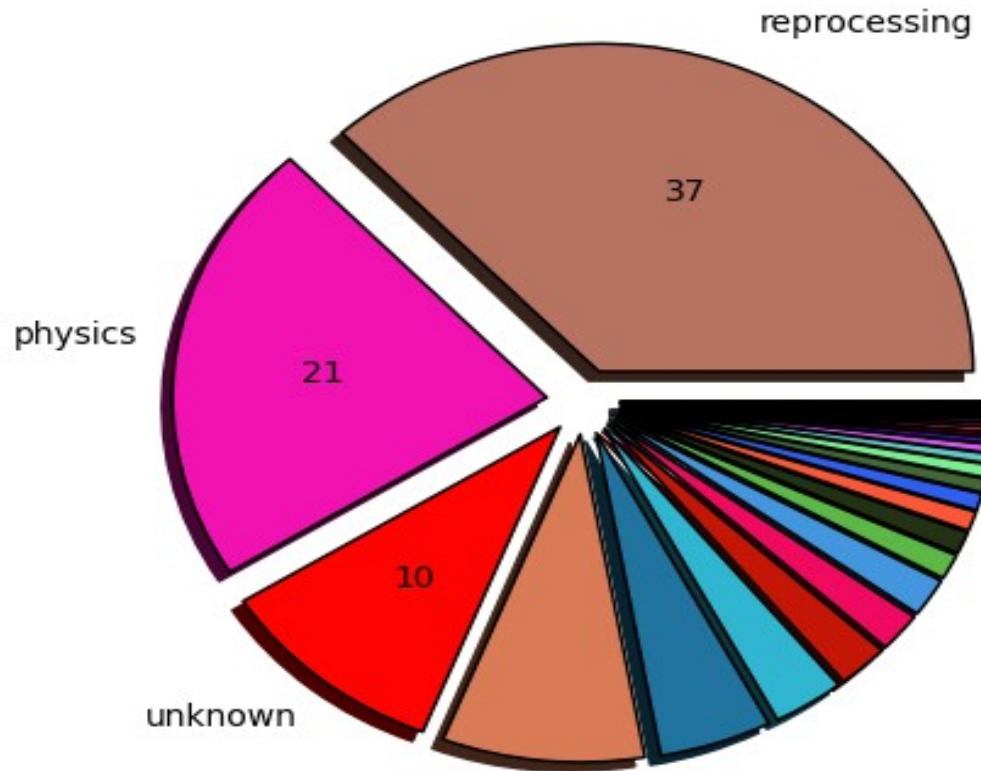
30 Days from 2011-03-04 to 2011-04-03



- | | | | | |
|--------------|-------------|------------|------------|------------|
| reprocessing | physics | gr_sm | user | unknown |
| top | validation | bphysics | heavyion | higgs |
| generators | sm | gr_susy | egamma | exotics |
| perf-jets | det-muon | gr_jetmet | perf-muons | susy |
| gr_detector | gr_top | phys-valid | phys-sm | gr_trigger |
| gr_bphysics | gr_btagging | perf-tau | det-indet | calib-muon |
| gr_higgs | jetetmiss | gr_egamma | localuser | gr_exotics |
| phys-higgs | | | | |

Maximum: 2,160,002,393 , Minimum: 0.00 , Average: 1,361,046,982 , Current: 806,609,916

Wall Clock consumptions All Jobs (Pie Chart in percentage) (Sum: 100.00)

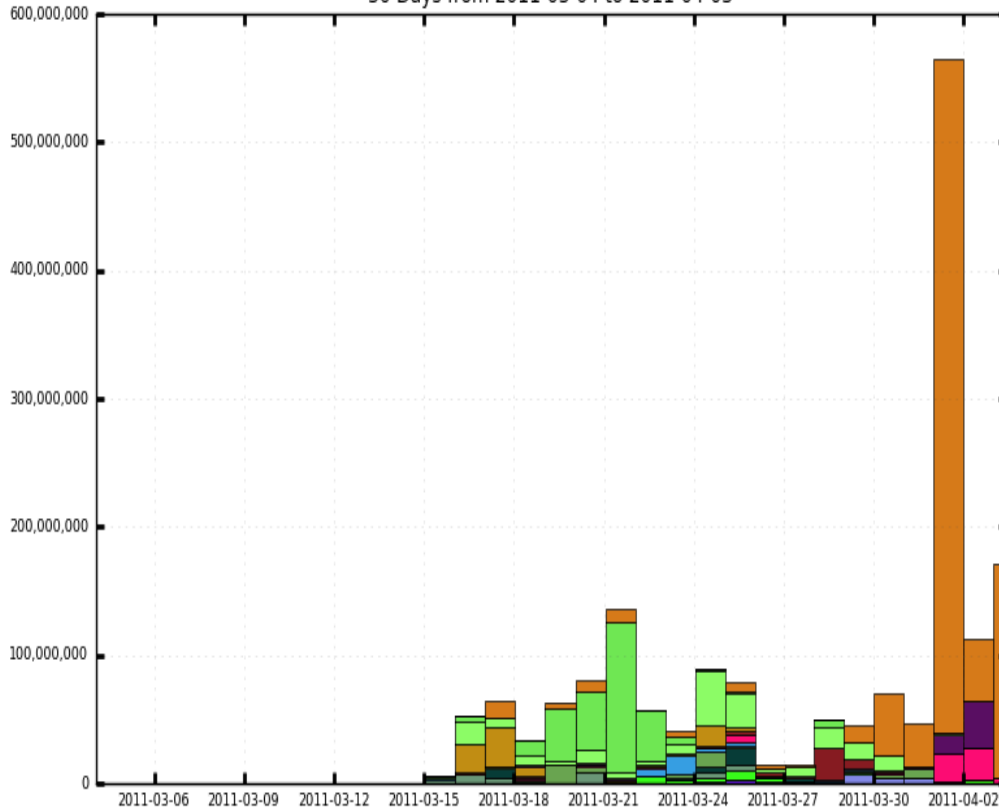


reprocessing (37.04)	physics (21.89)	unknown (10.10)	top (8.75)	validation (4.87)
gr_sm (3.12)	heavyion (2.22)	higgs (1.93)	user (1.89)	sm (1.24)
bphysics (1.09)	gr_susy (0.96)	exotics (0.82)	generators (0.67)	perf-jets (0.64)
egamma (0.47)	det-muon (0.42)	perf-muons (0.30)	susy (0.23)	gr_top (0.22)
gr_jetmet (0.19)	phys-sm (0.17)	gr_detector (0.16)	gr_bphysics (0.13)	gr_trigger (0.11)
gr_btagging (0.10)	phys-valid (0.10)	perf-tau (0.07)	calib-muon (0.05)	det-indet (0.03)
gr_egamma (0.01)	gr_higgs (0.01)	jetetmiss (0.01)	localuser (0.00)	gr_exotics (0.00)
phys-higgs (0.00)				





CPU consumptions Good Jobs (Time Stacked Bar Graph)

30 Days from 2011-03-04 to 2011-04-03



Maximum: 565,306,358 , Minimum: 0.00 , Average: 85,452,034 , Current: 170,936,923

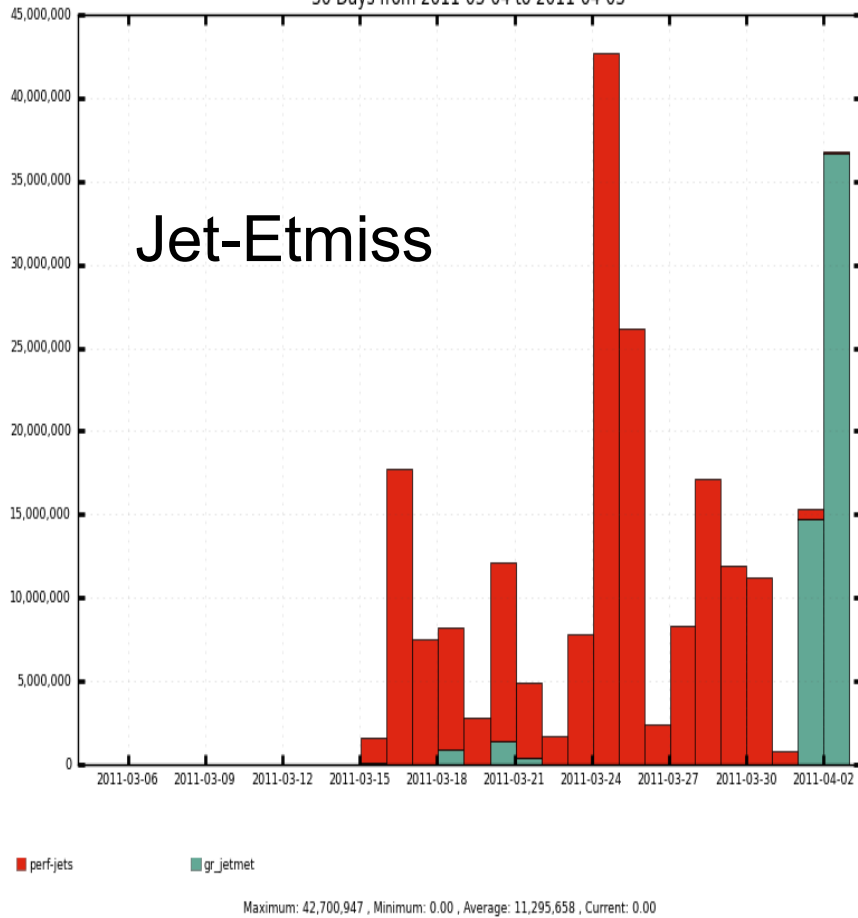
Wall Clock consumptions All Jobs (Pie Chart in percentage) (Sum: 99.00)



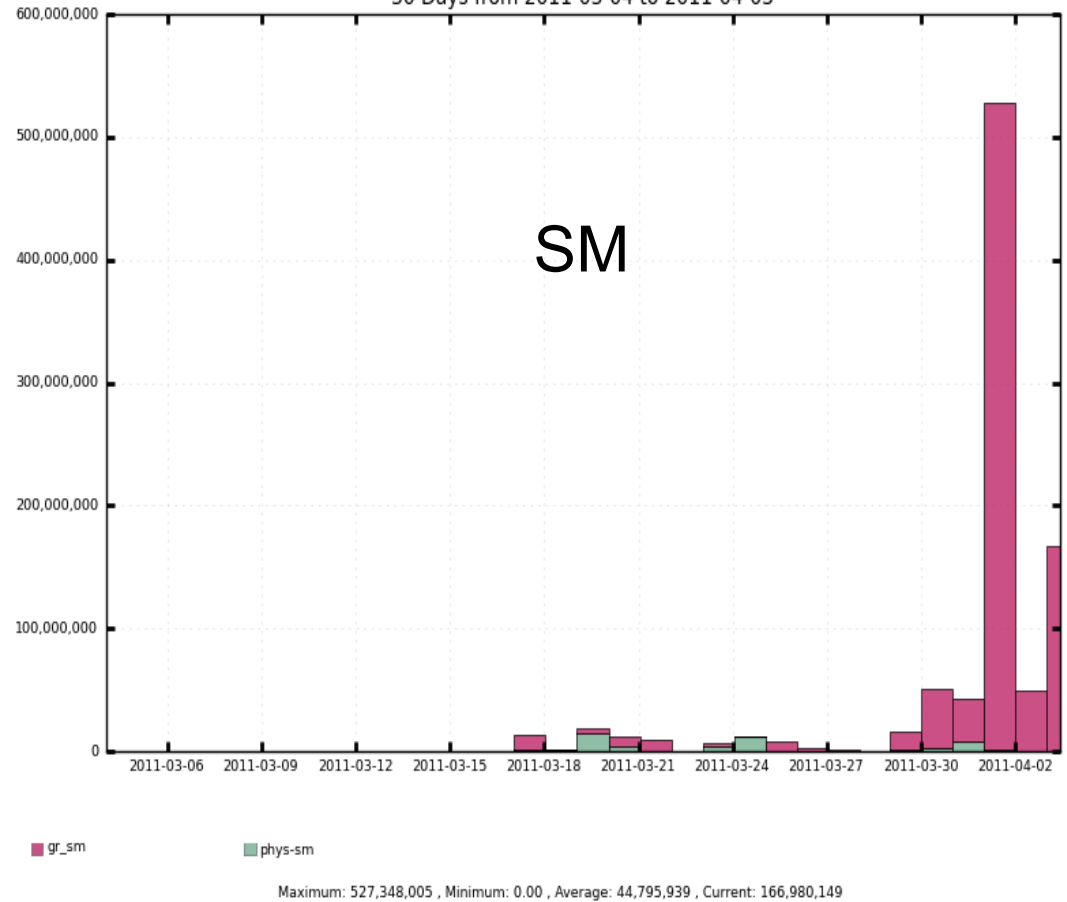
Specific groups



CPU consumptions Good Jobs (Time Stacked Bar Graph)
30 Days from 2011-03-04 to 2011-04-03




CPU consumptions Good Jobs (Time Stacked Bar Graph)
30 Days from 2011-03-04 to 2011-04-03



'Group Production' activities

CPU consumptions Good Jobs (Time Stacked Bar Graph)
21 Days from 2011-03-10 to 2011-03-31




gr_susy gr_sm gr_detector gr_btagging gr_trigger
gr_bphysics gr_top gr_jetmet gr_egamma gr_exotics
gr_higgs

Maximum: 113,145,912 , Minimum: 0.00 , Average: 27,713,585 , Current: 16,120,557

Pathena group jobs

CPU consumptions Good Jobs (Time Stacked Bar Graph)
21 Days from 2011-03-10 to 2011-03-31



Maximum: 126,455,667 , Minimum: 0.00 , Average: 27,272,833 , Current: 29,558,129