Fermi Large Area Telescope:

Collaboration Computing Needs

Richard Dubois
Stanford Linear Accelerator Center
richard@slac.stanford.edu

It Works!

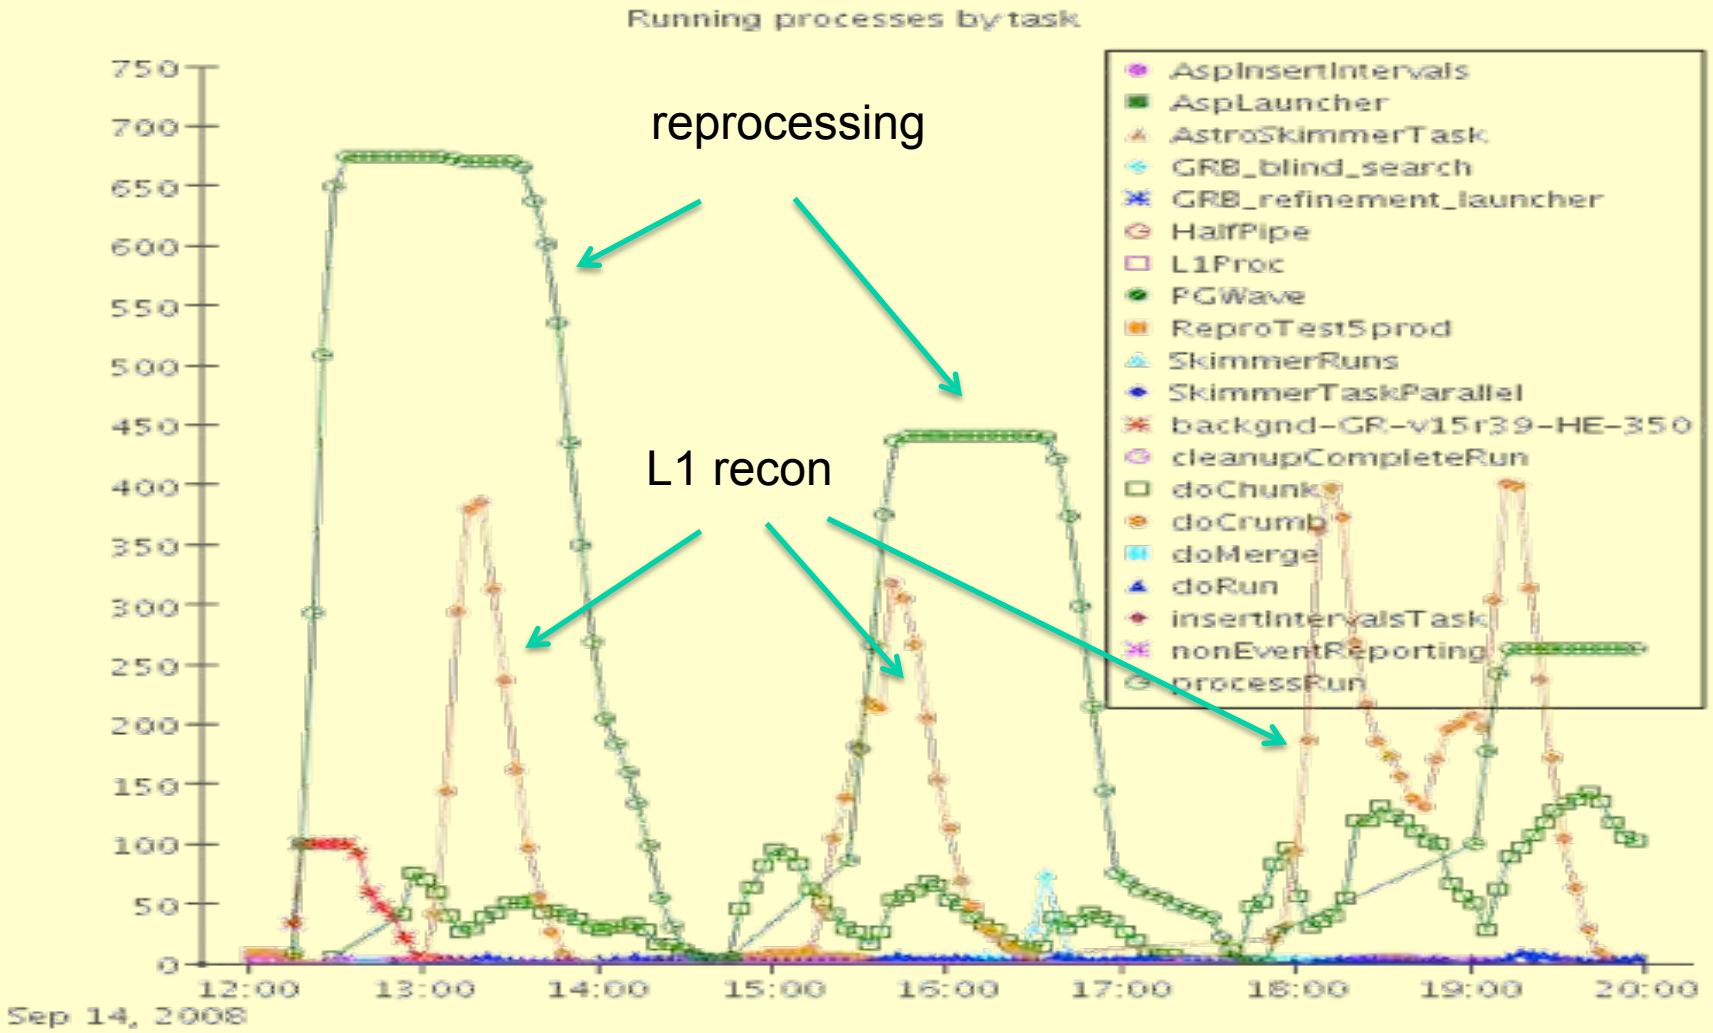
Running processes by task
reprocessing
L1 recon
AspInsertIntervals
AspLauncher
AstroSkimmerTask
GRB_blind_search
GRB_refinement_launcher
HalfPipe
L1Proc
PGWave
ReproTest5prod
SkimmerRuns
SkimmerTaskParallel
backgnd-GR-v15r39-HE-350
cleanupCompleteRun
doChunk
doCrumb
doMerge
doRun
insertIntervalsTask
nonEventReporting
processRun
750
700
650
600
550
500
450
400
350
300
250
200
150
100
50
0
12:00
13:00
14:00
15:00
16:00
17:00
18:00
19:00
20:00
Sep 14, 2008

Level 1 Processing 3-Ring Circus

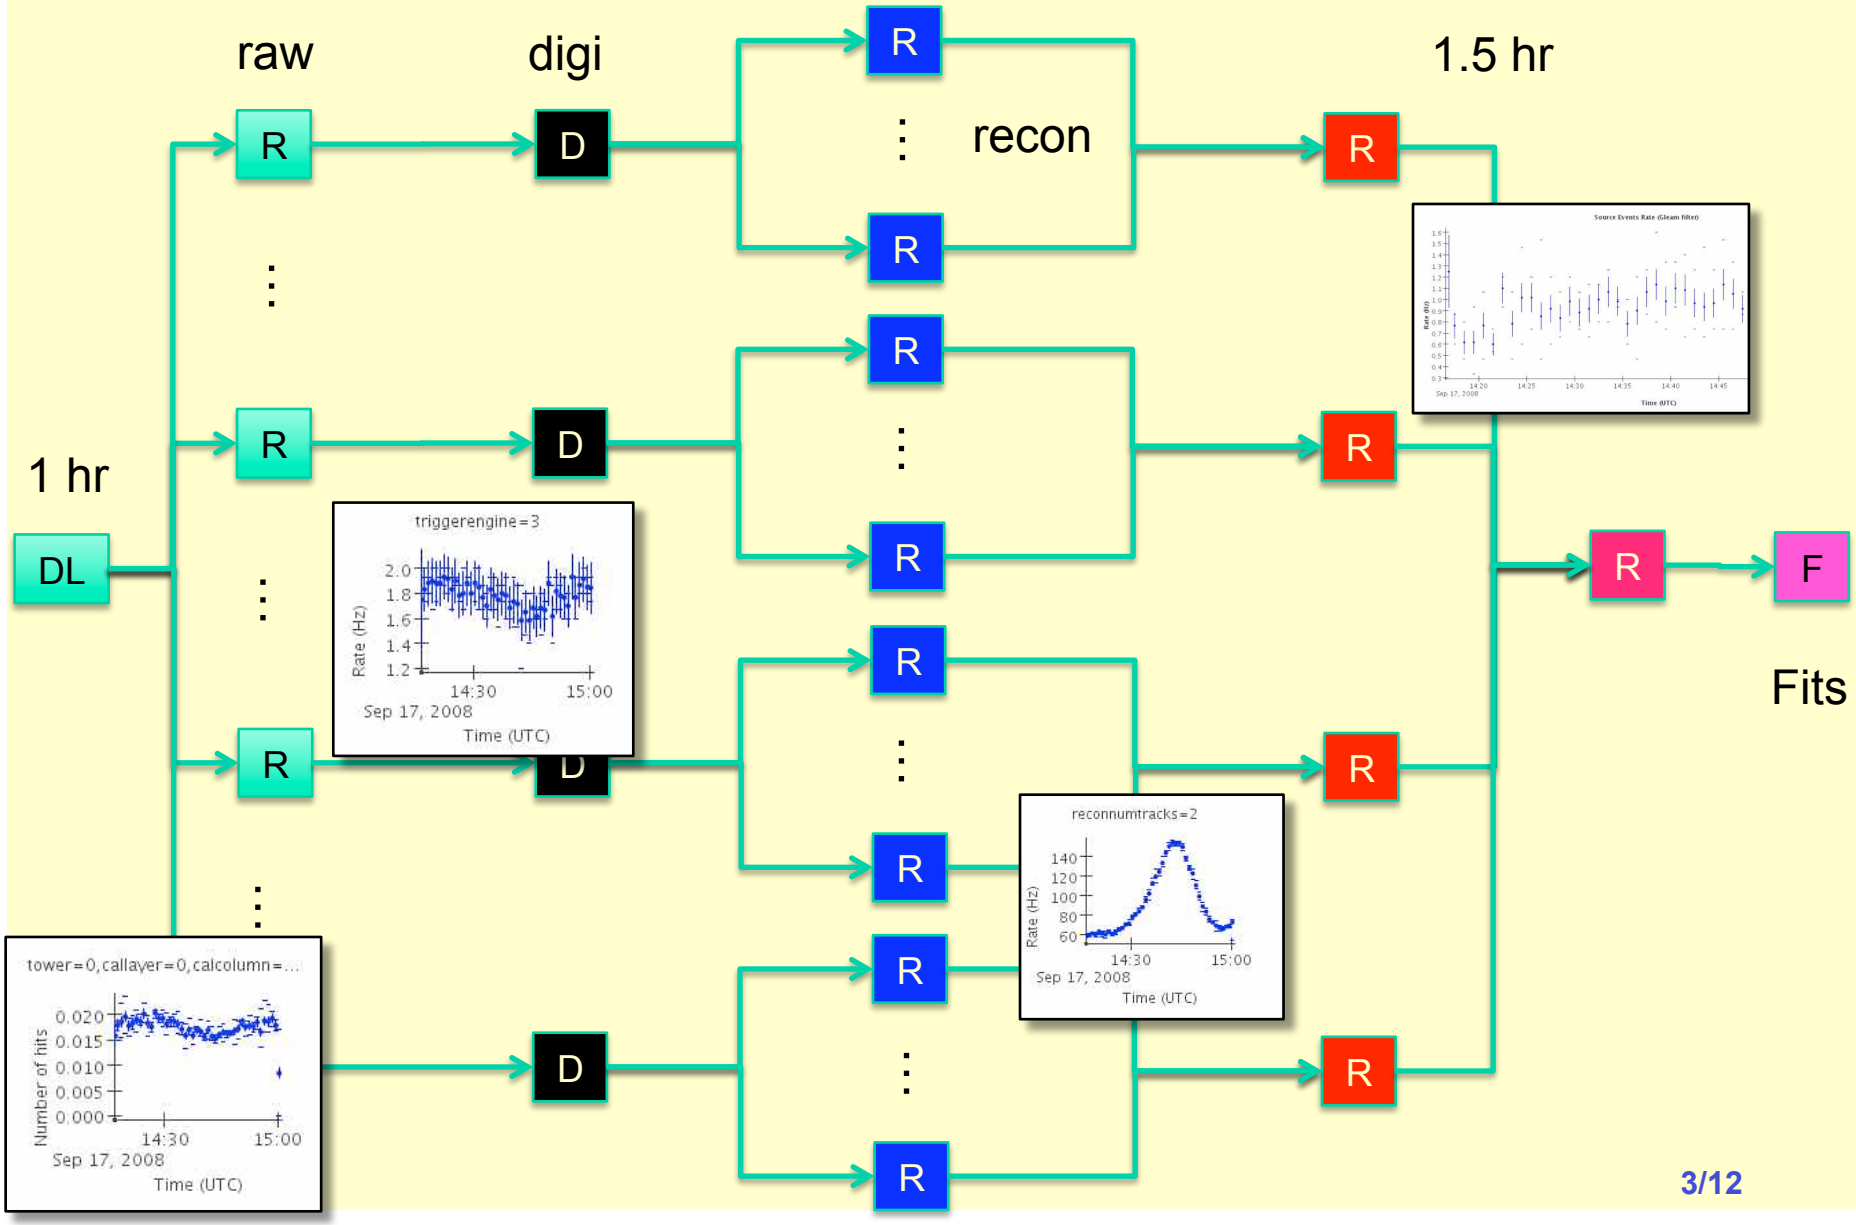

Current State

- At Launch
  - 800 cores + 150 TB disk + 250 TB tape available
  - 64 TB additional disk just delivered last week
  - Downlink rate ~30-40% above pre-launch bandwidth spec
- Level 1 processing + ASP
  - ~300 cores to turn around downlink promptly
  - ASP not yet noticeable
  - L1 will exhaust disk in 27 weeks
- Simulations have been mostly restricted to validating code builds
  - Will start to ramp up with ghost tracks/Pass7 development
  - Encourage MC away from SLAC
- Routine Science processing needs not scoped out yet
  - Resources for the Science Groups to do automated analyses on timeframes intermediate between ASP and Catalogue.
  - Expecting a need for 200-300 cores
- Reprocessings – the big gorilla
  - A year's full reprocessing would take (for example: new calibrations) roughly 3-4 weeks with 1000 cores
  - Expect to call on SLAC's general queues to help in such cases.
  - Need to model future needs once "10% solution" understood

RSP – next Growth Industry

- Leverage Pipeline to support Science Groups' needs
  - Automated spectral fitting and light curve creation
  - Pulsar implemented already
  - Binaries, GRB, AGN in progress
- Should be flexible to span times between ASP and catalogue
  - Run weekly in most cases
  - Also handle reprocessed data

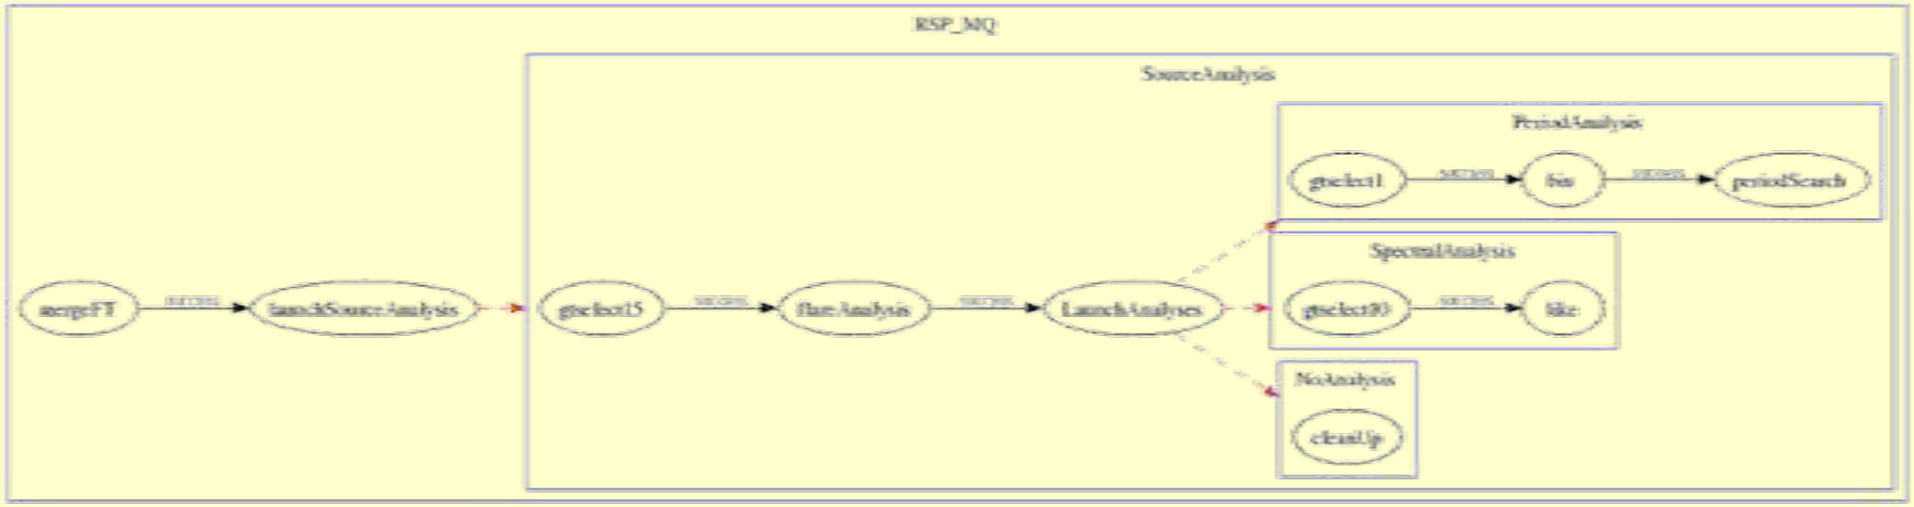

Leveraging the Pipeline: USP

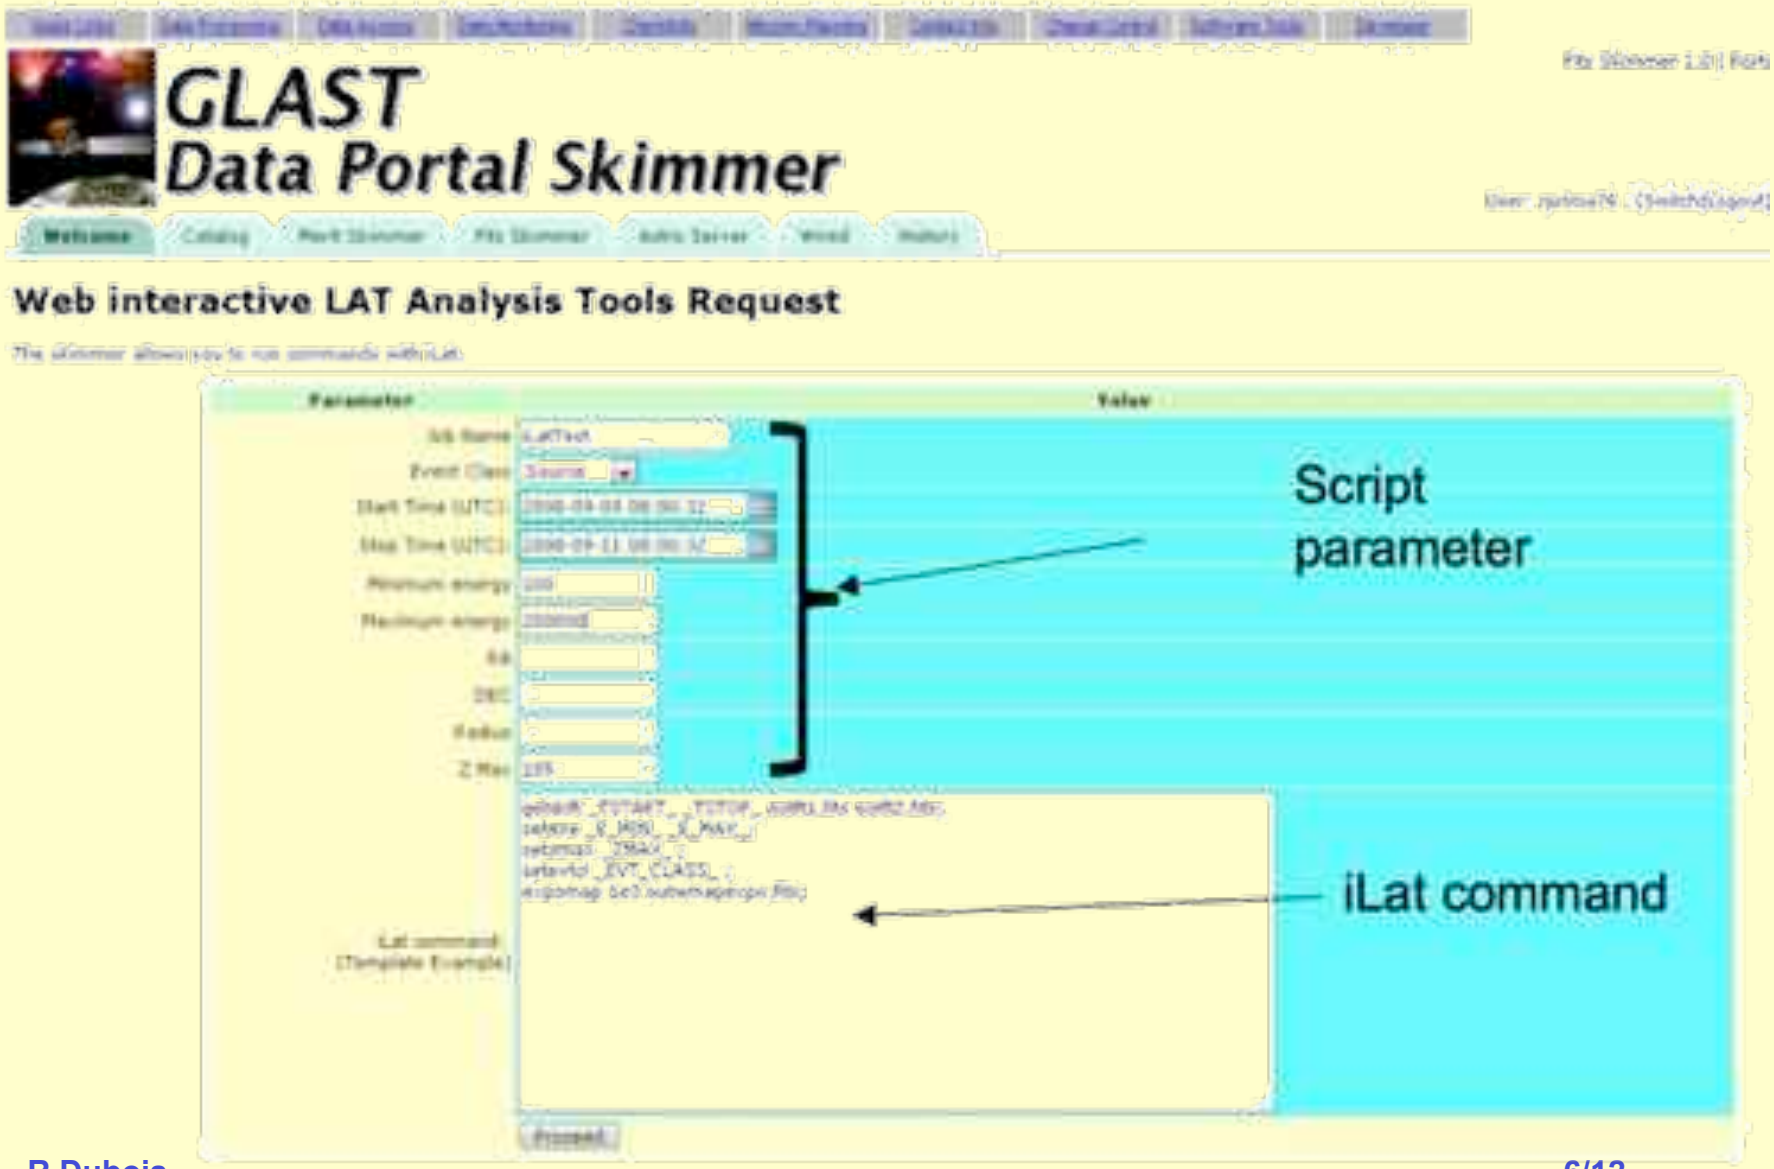

Needs for 2009

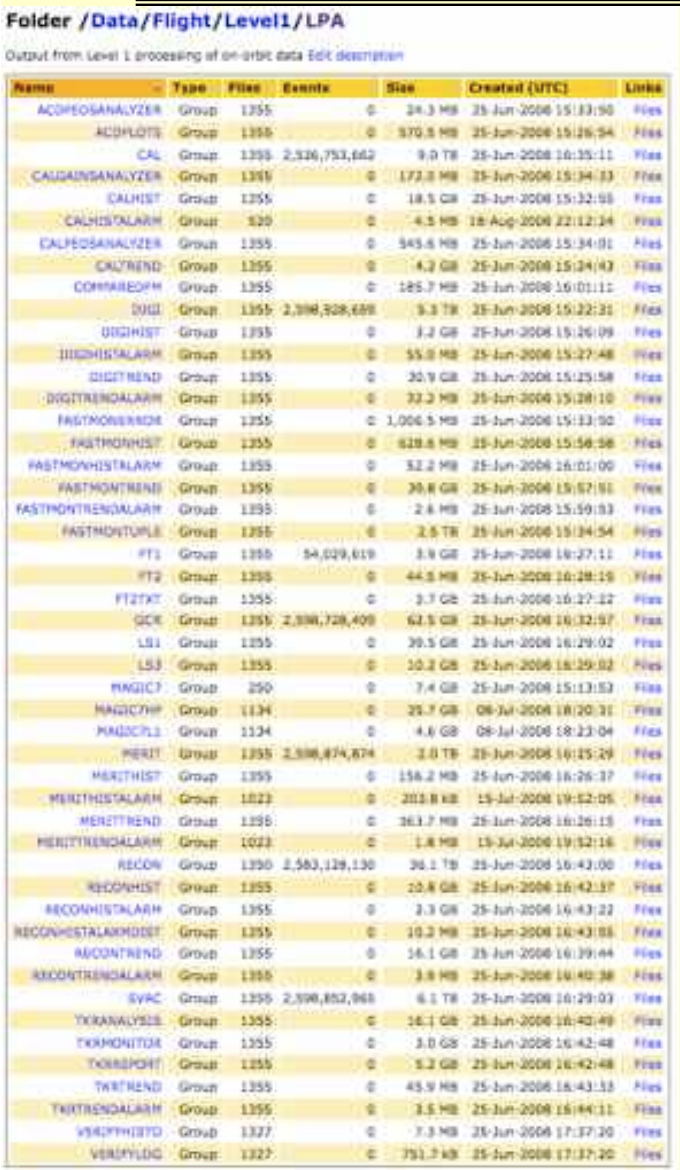

Folder /Data/Flight/Level1/LPA

Output from Level 1 processing of on-orbit data Edit description

Name	Type	Files	Events	Size	Created (UTC)	Links
ACDPEDSANALYZER	Group	1355	0	24.3 MB	25-Jun-2008 15:33:50	Files
ACDPLOTS	Group	1355	0	570.5 MB	25-Jun-2008 15:26:54	Files
CAL	Group	1355	2,526,753,662	9.0 TB	25-Jun-2008 16:35:11	Files
CALGAINSANALYZER	Group	1355	0	172.0 MB	25-Jun-2008 15:34:33	Files
CALHIST	Group	1355	0	18.5 GB	25-Jun-2008 15:32:55	Files
CALHISTALARM	Group	520	0	4.5 MB	18-Aug-2008 22:12:34	Files
CALPEDSANALYZER	Group	1355	0	545.6 MB	25-Jun-2008 15:34:01	Files
CALTREND	Group	1355	0	4.2 GB	25-Jun-2008 15:24:43	Files
COMPAREOFM	Group	1355	0	185.7 MB	25-Jun-2008 16:01:11	Files
DIGI	Group	1355	2,398,928,689	5.3 TB	25-Jun-2008 15:22:31	Files
DIGIHIST	Group	1355	0	3.2 GB	25-Jun-2008 15:26:09	Files
DIGIHISTALARM	Group	1355	0	55.0 MB	25-Jun-2008 15:27:48	Files
DIGITREND	Group	1355	0	30.9 GB	25-Jun-2008 15:25:58	Files
DIGITRENDALARM	Group	1355	0	32.2 MB	25-Jun-2008 15:28:10	Files
FASTMONRUNMON	Group	1355	0	1,006.5 MB	25-Jun-2008 15:33:50	Files
FASTMONHIST	Group	1355	0	628.6 MB	25-Jun-2008 15:56:56	Files
FASTMONHISTALARM	Group	1355	0	52.2 MB	25-Jun-2008 16:01:00	Files
FASTMONTREND	Group	1355	0	30.8 GB	25-Jun-2008 15:57:51	Files
FASTMONTRENDALARM	Group	1355	0	2.6 MB	25-Jun-2008 15:59:53	Files
FASTMONTURPLE	Group	1355	0	2.5 TB	25-Jun-2008 15:34:54	Files
FT1	Group	1355	54,029,619	3.9 GB	25-Jun-2008 16:27:11	Files
FT2	Group	1355	0	44.5 MB	25-Jun-2008 16:28:15	Files
FT2TXT	Group	1355	0	3.7 GB	25-Jun-2008 16:27:22	Files
GCR	Group	1355	2,398,728,409	62.5 GB	25-Jun-2008 16:32:57	Files
LS1	Group	1355	0	39.5 GB	25-Jun-2008 16:29:02	Files
LS3	Group	1355	0	10.2 GB	25-Jun-2008 16:29:02	Files
MAGIC7	Group	250	0	7.4 GB	25-Jun-2008 15:13:53	Files
MAGIC7HP	Group	1134	0	25.7 GB	08-Jul-2008 18:20:31	Files
MAGIC7L1	Group	1134	0	4.6 GB	08-Jul-2008 18:23:04	Files
MERIT	Group	1355	2,398,874,874	2.0 TB	25-Jun-2008 16:25:29	Files
MERITHIST	Group	1355	0	156.2 MB	25-Jun-2008 16:26:37	Files
MERITHISTALARM	Group	1023	0	203.8 kB	15-Jul-2008 19:52:05	Files
MERITTREND	Group	1355	0	363.7 MB	25-Jun-2008 16:26:15	Files
MERITTRENDALARM	Group	1023	0	1.8 MB	15-Jul-2008 19:52:16	Files
RECON	Group	1350	2,583,129,130	36.1 TB	25-Jun-2008 16:43:00	Files
RECONHIST	Group	1355	0	10.8 GB	25-Jun-2008 16:42:37	Files
RECONHISTALARM	Group	1355	0	3.3 GB	25-Jun-2008 16:43:22	Files
RECONHISTALARMDIST	Group	1355	0	10.2 MB	25-Jun-2008 16:43:05	Files
RECONTREND	Group	1355	0	16.1 GB	25-Jun-2008 16:39:44	Files
RECONTRENDALARM	Group	1355	0	3.8 MB	25-Jun-2008 16:40:38	Files
SVAC	Group	1355	2,398,852,965	6.1 TB	25-Jun-2008 16:29:03	Files
TKRANALYSIS	Group	1355	0	16.1 GB	25-Jun-2008 16:40:49	Files
TKRMONITOR	Group	1355	0	3.0 GB	25-Jun-2008 16:42:48	Files
TKRREPORT	Group	1355	0	5.2 GB	25-Jun-2008 16:42:48	Files
TKRTREND	Group	1355	0	45.9 MB	25-Jun-2008 16:43:33	Files
TKRTRENDALARM	Group	1355	0	3.5 MB	25-Jun-2008 16:44:11	Files
VERIFYHISTO	Group	1327	0	7.3 MB	25-Jun-2008 17:37:20	Files
VERIFYLOG	Group	1327	0	751.7 kB	25-Jun-2008 17:37:20	Files

Estimates are for ONE copy of the data

We will have to get smarter by the time of the first major reprocessing

Request 250 TB and assume we will figure out how to live in it in time

Still not plateaued: buy 400 more cores for data/MC load

We now estimate ~300 cores to reconstruct a downlink in 1-2 hours

Total data ~ 750 GB/day => ~275 TB

MC ~ 50 TB per yr of simulation (WAG)

WAG at 2010

- We must have dealt with issues of removing backgrounds in '08-09
- Assume a 20% solution accomplished: data still piles up!
  - 3x reduction in overall volume
    - We don't toss 80% of everything
    - 92 TB/yr of 2010 data
    - 123 TB previous years
    - 50 TB MC
    - 357 TB total
- No additional batch cores
- We still need to learn the usage patterns for reprocessing

Computing Resource Projections

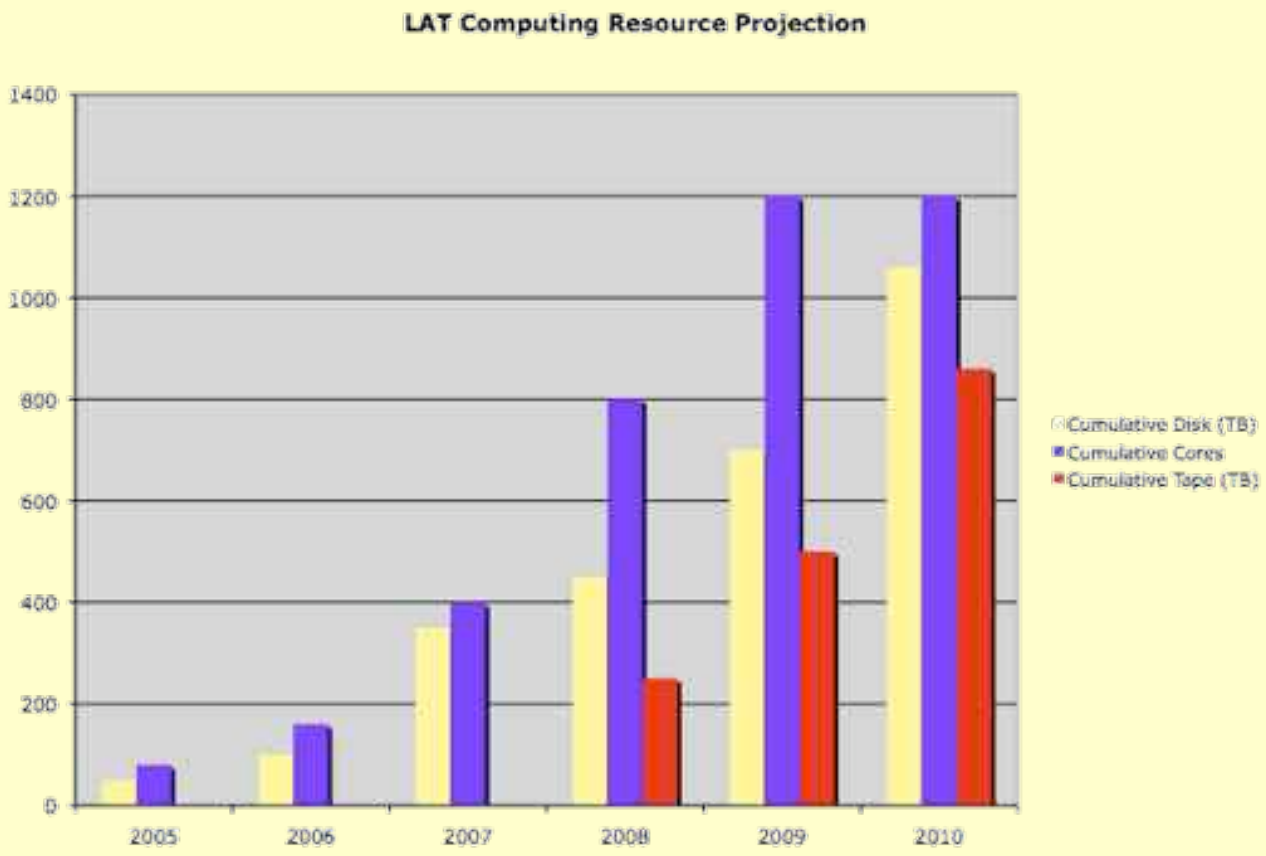

Need for/in '09:
- 250 TB disk: $225k
- 250 TB tape: $100k
- 400 cores: $200k
- Sundry servers: $50k

Total: $575k

Current pricing: $0.90k/TB disk; $0.40k/TB tape; $0.5k/core

The LAT Computing "GRID"

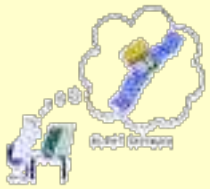

- SLAC
  - Pipeline2 configured originally to run at SLAC
  - 800 LAT cores now in batch farm
  - Plan to add 50% this coming year
- Lyon
  - P2 ported to Lyon
    - Stress testing operation completed
    - Plan to transfer all files to SLAC to minimise disk commitment at Lyon
    - Have demonstrated 500 cores
  - Ready for production use
    - Agreement with Lyon for Fermi capacity now needed
- Italy
  - Constrained to use the "real" GRID
  - 100 cores in '08 – asking for 200 more in '09
  - Real work needed to interface to Pipeline

Miscellaneous

- **Bought 2 new Oracle servers in 2008**
  - Doing the trick
  - Success of science trending data is stressing data volume
    - Need to "sparsify" data as it ages
    - May need to add disk to the servers in 2009

- **Expect to buy 4-5 standalone servers with remainder of '08 money**
  - Beef up astroserver
  - Allow isolation of more key servers

- **Upgrades to Event Display and Code Development gui**
  - One time cost of $25k - propose using 2008 OCF under-run
  - Upgrade these tools by software house (authors have moved on to this outfit)
  - Take over ma ion

Abandoned!

Take-away Messages

- **Our handling of disk volume still not understood**
- **Will need to understand reprocessing loads**

Computing Model not settled yet!

- **Growth areas**
  - **Increasing Automation and Reliability of L1/ASP**
  - **RSP**
  - **Enabling science analysis**