



Fermi

Gamma-ray Space Telescope

Computing and Astroparticle Physics
2nd ASPERA workshop,
30-31 May 2011
Barcelona, Spain



Computing Model for Fermi Large Area Telescope

Johann Cohen-Tanugi




Université Montpellier 2 and IN2P3/CNRS

With great help from
Richard Dubois, Tony Johnson
SLAC

- **Fermi cheat sheet**
- **Data Processing Model : the LAT Pipeline**
 - **Design consideration & Pipeline description**
 - **Monitoring, Alerts, Automatic Science**
 - **Data catalog, access and selection tools**
- **Science Software Development:**
 - **GlastRelease : Event Reconstruction**
 - **ScienceTools : Science analysis “toolkit”**

The Fermi observatory


June 11, 2008



5 year nominal, possibly 10 year mission (No expandable)

• Sky survey is the nominal observation mode

The Fermi Large Area Telescope (LAT)



Anti-Coincidence Detector

- 4% R.L.
- 89 scintillating tiles
- efficiency (>0.9997) for MIPs

Tracking detector


- 16 tungsten foils (12x3%R.L., 4x18%R.L.)
- 18 pairs of silicon strip arrays
- 884736 strips (228 micron pitch)

Trigger

- Overall HW Trigger Rate ~few KHz
- Software Filters Reduce Rate
- Downlink: ~400-500 Hz
- Rate after Ground Cuts: ~few Hz

Calorimeter


- 8.5 radiation lengths
- 8 layers cesium iodide logs
- 1536 logs total (1200kg)




Unidentified
sources



Active Galactic Nuclei




Cosmic ray
acceleration

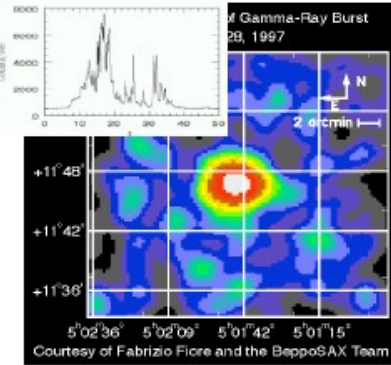


Solar flares


Fermi Science




Pulsars




Gamma Ray Bursts




Quantum Gravity?



Dark matter

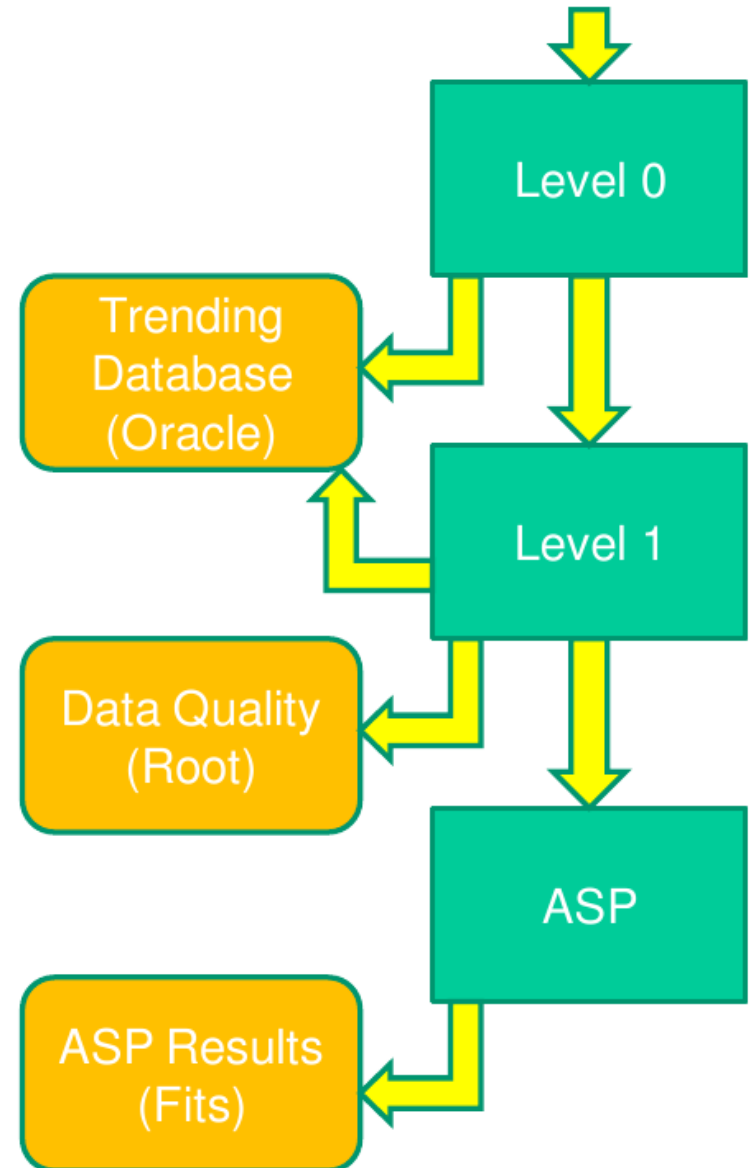


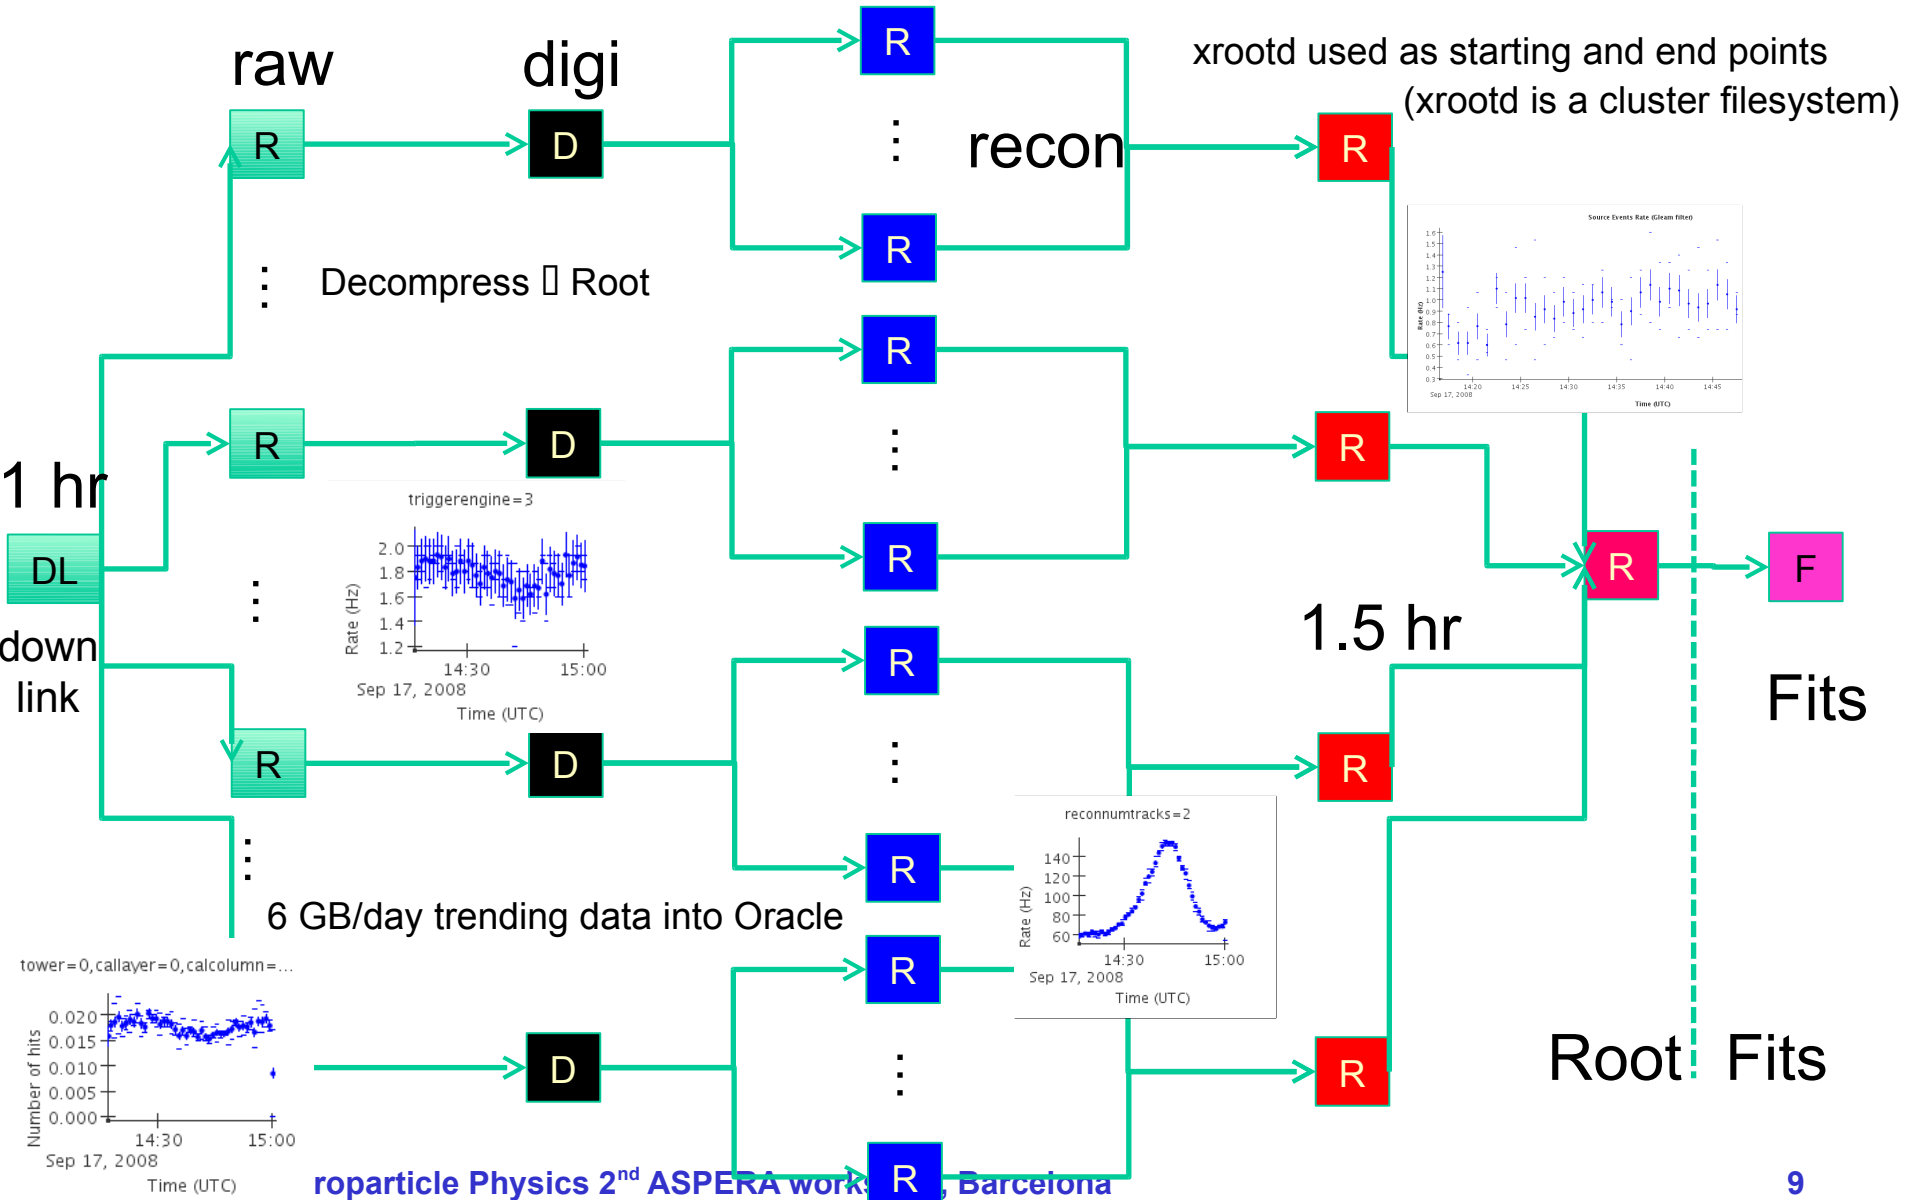
Fermi MISSION ELEMENTS



- **Generic pipeline processing to support simultaneous running**
 - **Automatic trigger of prompt processing & monitoring**
 - **Reprocessing and/or Simulations (should have no impact on above)**
- **Aim : full processing before next download (8/day → 3hrs goal)**
 - **Distributed computing taking advantage of the massive “embarrassing” parallelism of the problem (2000 jobs, 800 cores)**
 - **<0.01% jobs need manual intervention**
- **Scaling**
 - **10 years of operations foreseen : Many hundreds of terabytes of data**
 - **Hundreds of millions of datasets and processes**
- **Reliability**
 - **10’s of thousands of batch jobs per day**
 - **43k in a day is our record (approx 40k CPU-hrs)**
- **Ease of use within collaboration : web based frontend**


- **Downlink from Goddard Space Flight Center**
 - ~8 downloads per day
 - 15 GB total daily
- **Level 0 Processing**
 - Automatically launched as data arrives
 - Decode & repackage incoming data
 - Split science data from telemetry data
- **Level 1 Processing**
 - Full event reconstruction: 750 GB/day
 - Data Quality Monitoring
 - Transfer science summary files to Fermi Science Support Center - 200 MB/day
 - Immediately available to the public
- **ASP (Automated Science Processing)**
 - GRB and Flare detection
 - Spectral analysis (Routine Science Processing)
- **120,000 quantities continuously monitored**
 - Mixture of Oracle, Root, Fits data





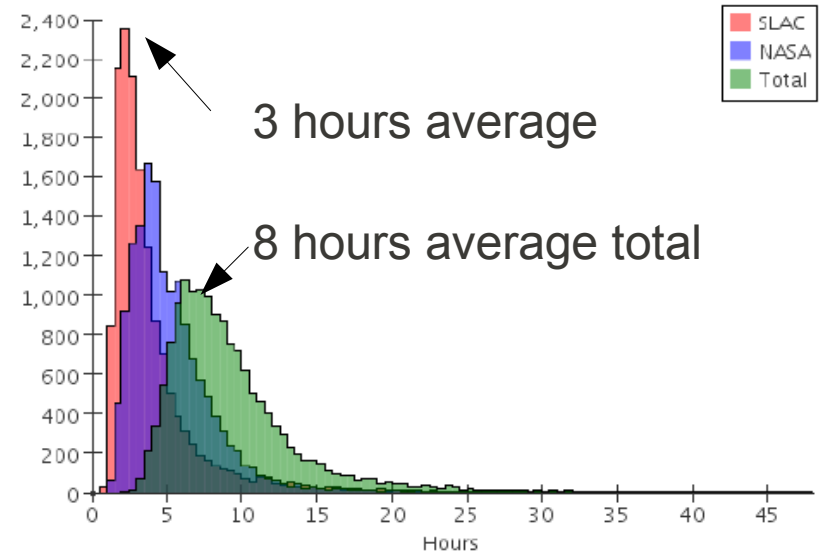
Pipeline Performance and Reliability

Data processing elapsed time per run vs MET



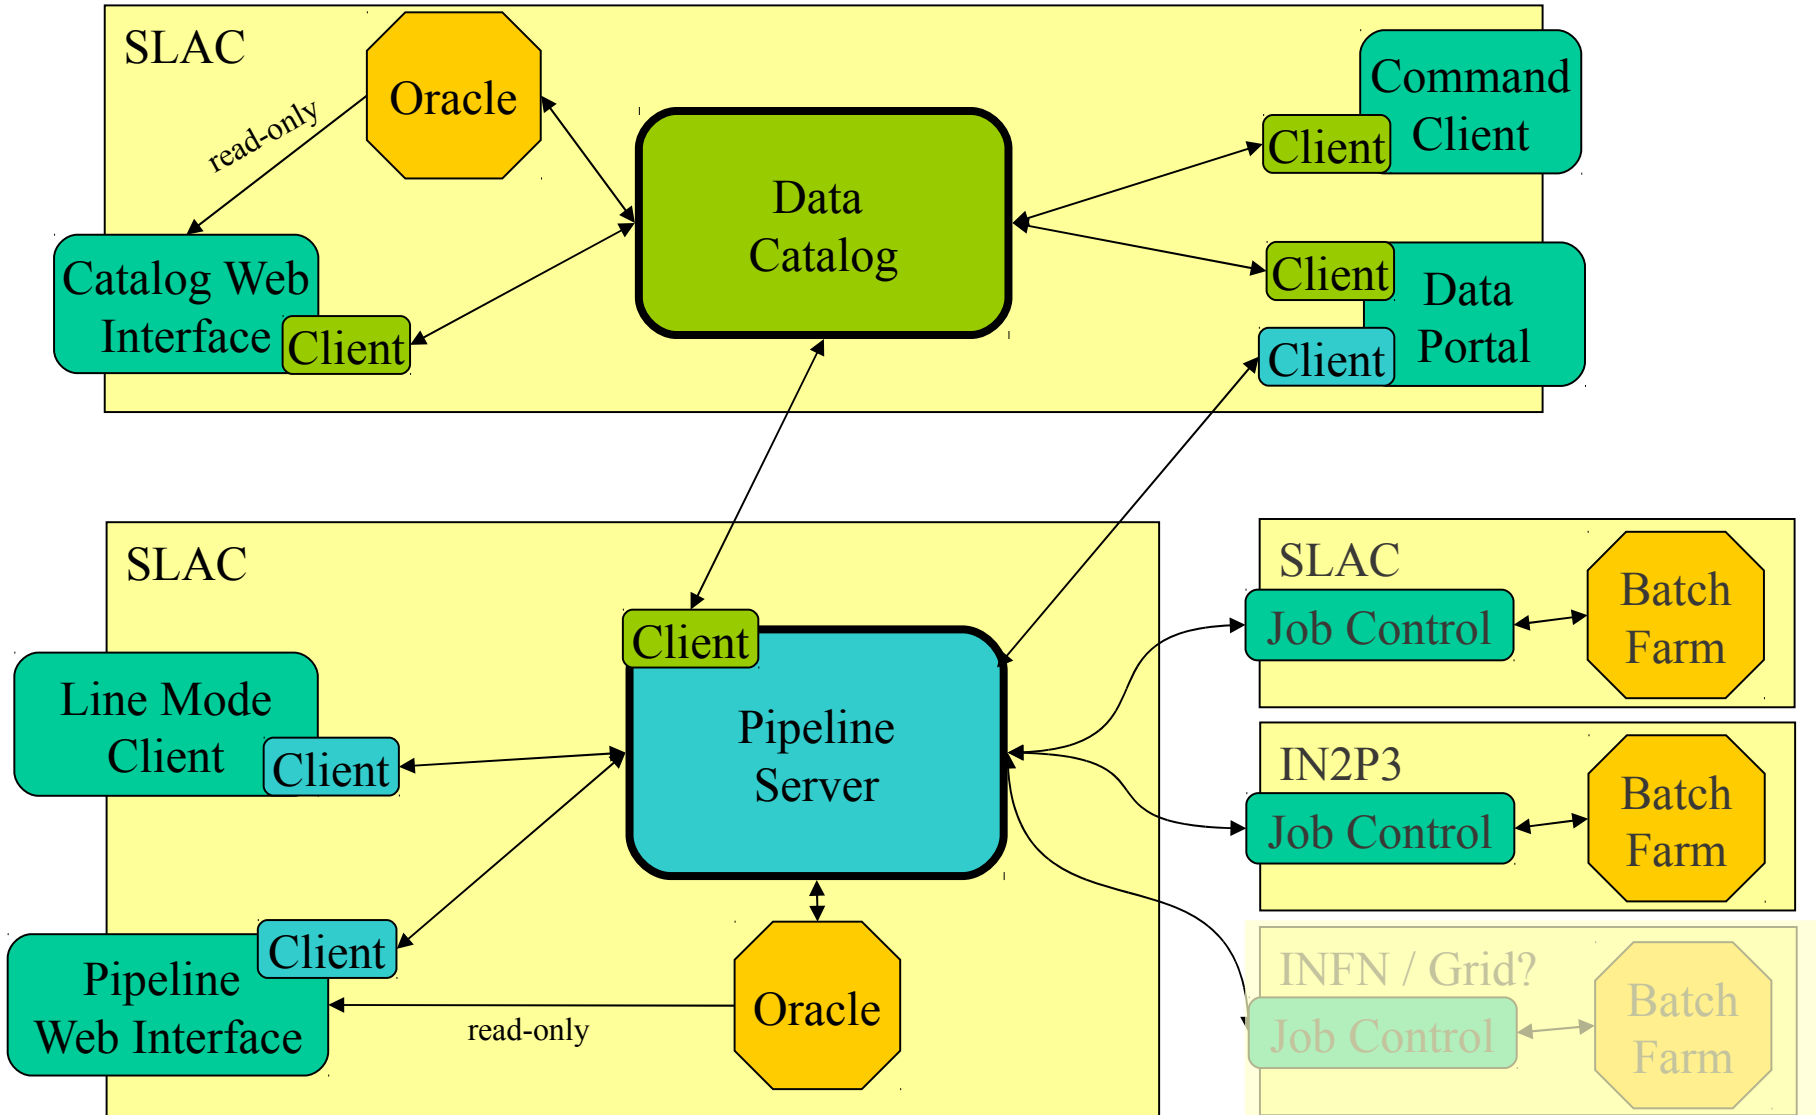
Elapsed time between data being recorded on satellite and arriving at SLAC (blue), and between arriving at SLAC and being totally processed (red), and total elapsed time (green). Almost all data is fully processed <24 hours after being recorded.

Data processing elapsed time per run



- **Pipeline, Web Servers implemented in Java**
 - **Apache tomcat web server**
 - **Redundant servers for reliability, scalability**
 - **Java Server Pages (JSP) for web pages**
 - **Extensive use of 3rd party and custom tag libraries**
 - **DisplayTag for tabular data**
 - **AIDA tld for dynamic plot generation**
 - **JMX for monitoring, control**
 - **JaSIG CAS single-signon for user authentication**
- **Extensive use of database for storing state, history**
 - **Oracle (10g, 11)**
 - **Java Stored Procedures for performance**
 - **GridControl for performance, tuning, monitoring**

Pipeline Implementation



- **SLAC compute farm (Fermi allocation)**
 - **1600 CPUs in batch farm running LSF (+ peak loads of >3000)**
 - **1100 TB disk - xroot – 32 TB raid 10 Sun thumpers**
 - **1500 TB tape in robotic silo (HPSS)**
- **Lyon compute farm (CCIN2P3, France)**
 - **1200 CPUs in batch farm running BQS**
 - **few TB disk allocated for all Fermi uses (transfer generated files to SLAC)**
 - **seamlessly used by pipeline from SLAC**
- **Plans for using Italian GRID**
- **Extensive use of relational DB**
 - **Two ‘Niagara’ redundant Oracle servers each with 1.5 TB raid disk**
 - **Science monitoring**
 - **LAT+Spacecraft ‘housekeeping’ trending**
 - **Bookkeeping (dataCatalog, processing config, etc.)**
 - **Science data (e.g., GRB catalog)**
 - **Resource trending**

Use open source tools to watch resources and services

Ganglia for usage – trends for loads, I/O

- **Oracle, xrootd, afs, nfs, linux servers, ...**

Nagios for function – aliveness etc

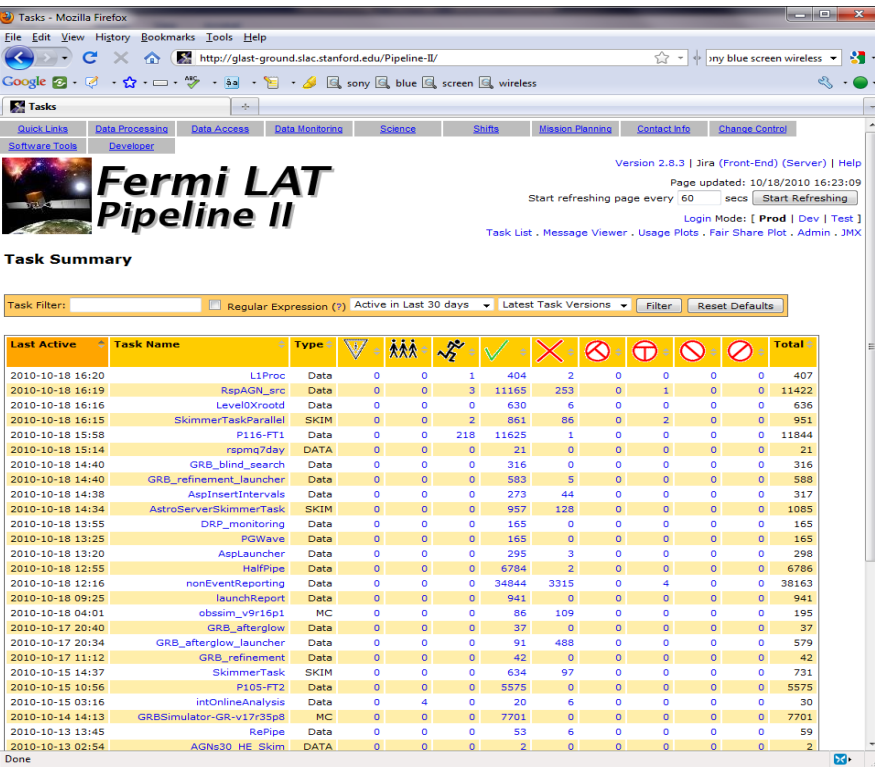
- **Many standard functions: ping, database checks, ...**
- **Customisable with plugins to check anything**
- **email, pager alerts with warning levels and history**

Our own trending system to track performance & resource usage

- **Be generous with info stored in the database!**



- “Duty Scientists” monitoring data quality daily
- All of the data processing and data quality monitoring can be done from the web

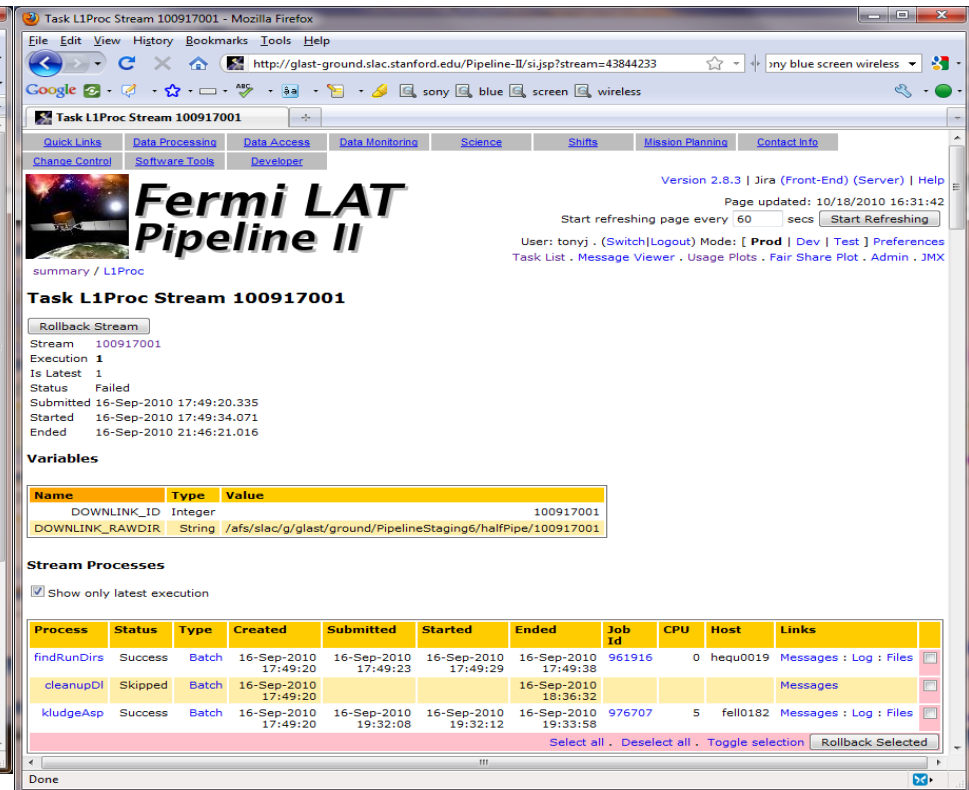


Version 2.8.3 | Jira (Front-End) (Server) | Help
Page updated: 10/18/2010 16:23:09
Start refreshing page every 60 secs [Start Refreshing](#)
Login Mode: [**Prod** | Dev | Test]
[Task List](#) . [Message Viewer](#) . [Usage Plots](#) . [Fair Share Plot](#) . [Admin](#) . [JMX](#)

Task Summary

Task Filters: Regular Expression (?) Active in Last 30 days Latest Task Versions Filter

Last Active	Task Name	Type	0	1	2	3	4	5	6	7	8	9	10	Total
2010-10-18 16:20	L1Proc	Data	0	0	1	404	2	0	0	0	0	0	0	407
2010-10-18 16:19	RspAGN_src	Data	0	0	3	11165	253	0	1	0	0	0	0	11422
2010-10-18 16:16	Level0Xrootd	Data	0	0	0	630	6	0	0	0	0	0	0	636
2010-10-18 16:15	SkimmerTaskParallel	SKIM	0	0	2	861	86	0	2	0	0	0	0	951
2010-10-18 15:58	P116-FT1	Data	0	0	218	11625	1	0	0	0	0	0	0	11844
2010-10-18 15:14	rspmq7day	DATA	0	0	0	21	0	0	0	0	0	0	0	21
2010-10-18 14:40	GRB_blind_search	Data	0	0	0	316	0	0	0	0	0	0	0	316
2010-10-18 14:40	GRB_refinement_launcher	Data	0	0	0	593	5	0	0	0	0	0	0	598
2010-10-18 14:38	AspInsertIntervals	Data	0	0	0	273	44	0	0	0	0	0	0	317
2010-10-18 14:34	AstroServerSkimmerTask	SKIM	0	0	0	957	128	0	0	0	0	0	0	1085
2010-10-18 13:55	DRP_monitoring	Data	0	0	0	165	0	0	0	0	0	0	0	165
2010-10-18 13:25	PGIWave	Data	0	0	0	165	0	0	0	0	0	0	0	165
2010-10-18 13:20	AspLauncher	Data	0	0	0	295	3	0	0	0	0	0	0	298
2010-10-18 12:55	HalfPipe	Data	0	0	0	6784	2	0	0	0	0	0	0	6786
2010-10-18 12:16	nonEventReporting	Data	0	0	0	34844	3315	0	4	0	0	0	0	38163
2010-10-18 09:25	launchReport	Data	0	0	0	941	0	0	0	0	0	0	0	941
2010-10-18 04:01	obsim_v9r16p1	MC	0	0	0	86	109	0	0	0	0	0	0	195
2010-10-17 20:40	GRB_afterglow	Data	0	0	0	37	0	0	0	0	0	0	0	37
2010-10-17 20:34	GRB_afterglow_launcher	Data	0	0	0	91	488	0	0	0	0	0	0	579
2010-10-17 11:12	GRB_refinement	SKIM	0	0	0	42	0	0	0	0	0	0	0	42
2010-10-15 14:37	SkimmerTask	SKIM	0	0	0	634	97	0	0	0	0	0	0	731
2010-10-15 10:56	P105-FT2	Data	0	0	0	5575	0	0	0	0	0	0	0	5575
2010-10-15 03:16	intOnlineAnalysis	Data	0	4	0	20	0	0	0	0	0	0	0	20
2010-10-14 14:13	GRBSimulator-GR-v1735p8	MC	0	0	0	7701	0	0	0	0	0	0	0	7701
2010-10-13 13:45	RePipe	Data	0	0	0	53	6	0	0	0	0	0	0	59
2010-10-13 02:54	AGNe30 HE_Skim	DATA	0	0	0	2	0	0	0	0	0	0	0	2



Version 2.8.3 | Jira (Front-End) (Server) | Help
Page updated: 10/18/2010 16:31:42
Start refreshing page every 60 secs [Start Refreshing](#)
User: tonjy . (Switch[Logout] Mode: [**Prod** | Dev | Test] Preferences
[Task List](#) . [Message Viewer](#) . [Usage Plots](#) . [Fair Share Plot](#) . [Admin](#) . [JMX](#)

Task L1Proc Stream 100917001

[Rollback Stream](#)

Stream 100917001
Execution 1
Is Latest 1
Status Failed
Submitted 16-Sep-2010 17:49:20.335
Started 16-Sep-2010 17:49:34.071
Ended 16-Sep-2010 21:46:21.016

Variables

Name	Type	Value
DOWNLINK_ID	Integer	100917001
DOWNLINK_RAWDIR	String	/afs/slac/g/glast/ground/PipelineStaging5/halfPipe/100917001

Stream Processes

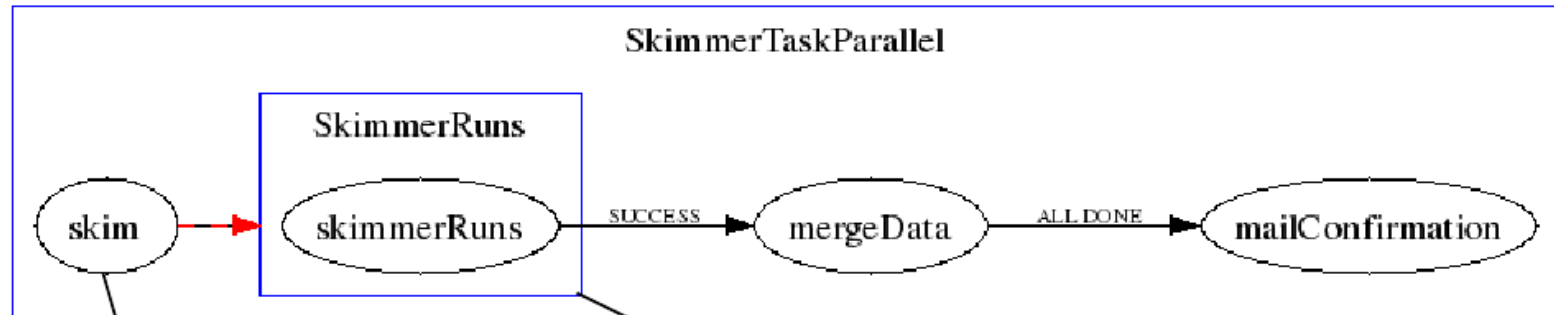
Show only latest execution

Process	Status	Type	Created	Submitted	Started	Ended	Job Id	CPU	Host	Links
findRunDirs	Success	Batch	16-Sep-2010 17:49:20	16-Sep-2010 17:49:23	16-Sep-2010 17:49:29	16-Sep-2010 17:49:38	961916	0	hequ0019	Messages : Log : Files
cleanupDI	Skipped	Batch	16-Sep-2010 17:49:20			16-Sep-2010 18:36:32				Messages
kludgeAsp	Success	Batch	16-Sep-2010 17:49:20	16-Sep-2010 19:32:08	16-Sep-2010 19:32:12	16-Sep-2010 19:33:58	976707	5	fell0182	Messages : Log : Files

Select all . Deselect all . Toggle selection [Rollback Selected](#)

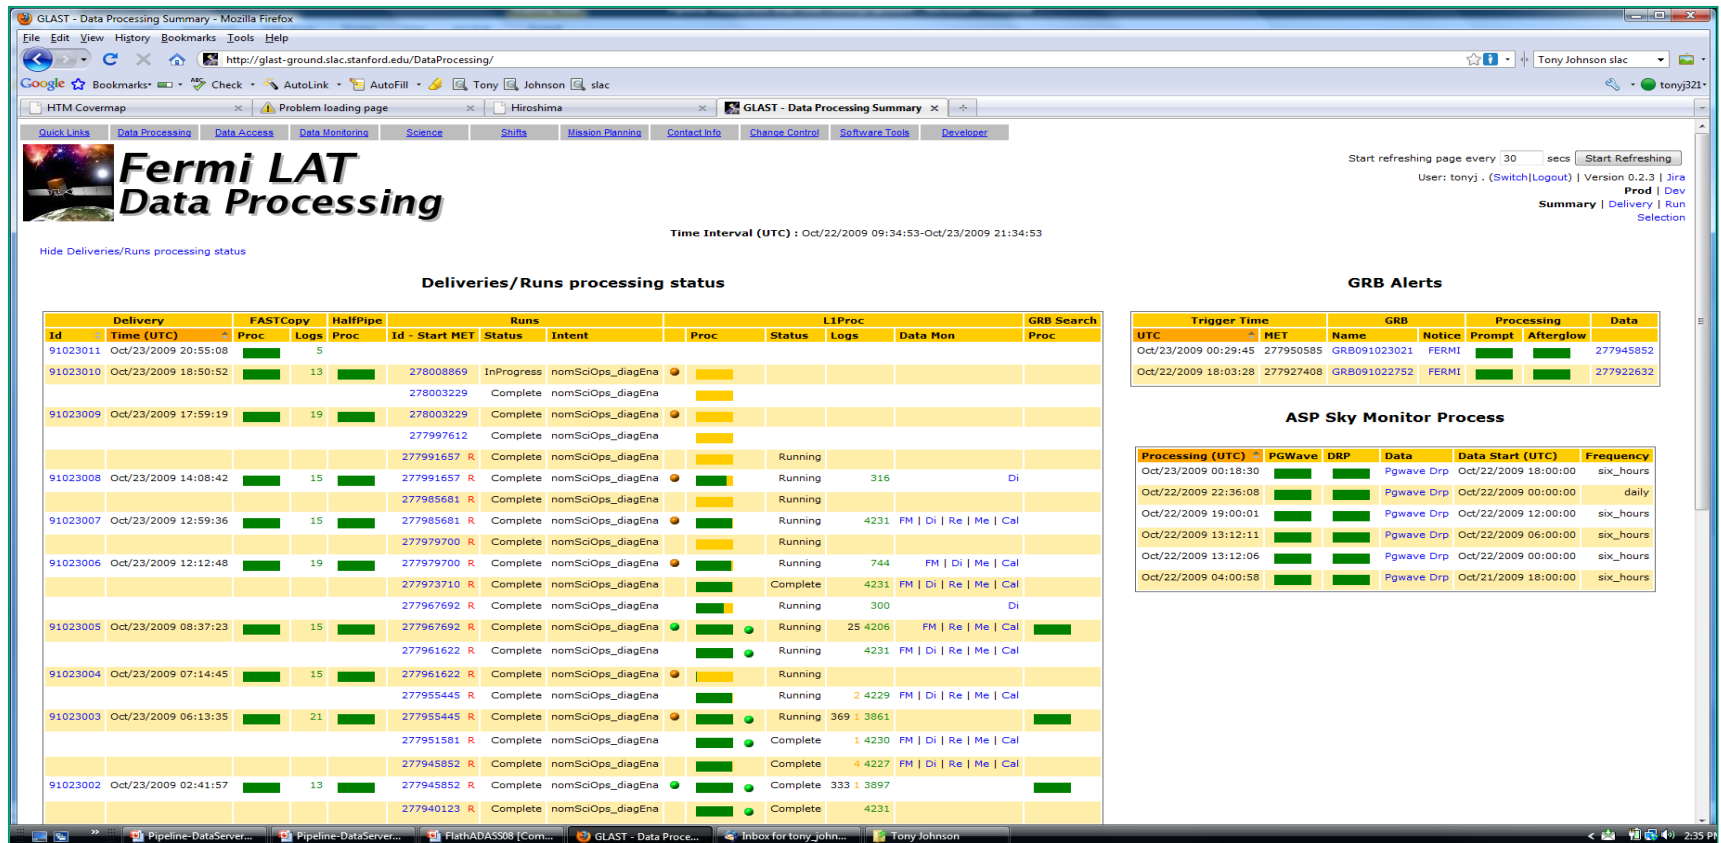
- Pipeline web interface allows
 - Many views of data processing, down to log files of individual jobs
 - Job submission (but normally done from command line)
 - If jobs do fail they can be “rolled back” directly from the web interface

Tasks, Subtasks, Processes, Streams, specified in user written XML



```
<process name="skim">
  <script>
    <![CDATA[
from java.util import HashMap
start = 0
chunk_size = DP_FILE_LINES/DP_SUBTASKS + 1
for i in range(DP_SUBTASKS):
  vars = HashMap()
  vars.put("DP_START", start)
  end = min( start + chunk_size, DP_FILE_LINES )
  vars.put("DP_END", end)
  if end>start:
    pipeline.createSubstream("SkimmerRuns",i,vars)
  start = end
]]>
  </script>
  <createsSubtasks>
    <subtask>SkimmerRuns</subtask>
  </createsSubtasks>
</process>
```

```
<task name="SkimmerRuns" type="Data" version="1.0">
  <process name="skimmerRuns">
    <variables>
      <var name="streamID">${format(pipeline.stream, "%04d")}</var>
    </variables>
    <job maxCPU="${DP_SKIMMER_MAXCPU}">
      cat ${DP_FILE_LIST} | head -${DP_END} | tail -${(DP_END-DP_START)} > partlist
      export SK_FILE_LIST_FILE=partlist
      export SK_OUT_DIR=${DP_OUT_DIR}/${streamID}
      export SK_ENFORCE_OUTPUT_FILES=false
      export SK_MAX_FILE_SIZE=0
      export SK_OUT_FILE_BODY=${DP_JOBNAME}-${streamID}
      mkdir -p ${SK_OUT_DIR}
      $SK_DIR/bin/skimmer
    </job>
  </process>
</task>
```



GLAST - Data Processing Summary - Mozilla Firefox

http://glast-ground.slac.stanford.edu/DataProcessing/

Quick Links: Data Processing, Data Access, Data Monitoring, Science, Shifts, Mission Planning, Contact Info, Change Control, Software Tools, Developer

Fermi LAT Data Processing

Time Interval (UTC) : Oct/22/2009 09:34:53-Oct/23/2009 21:34:53

Hide Deliveries/Runs processing status

Deliveries/Runs processing status

Id	Delivery Time (UTC)	FASTCopy Proc	Logs	HalfPipe Proc	Id - Start MET	Status	Intent	Proc	Status	Logs	Data Mon	GRB Search
91023011	Oct/23/2009 20:55:08	Proc	5	Proc	278008869	InProgress	nomSciOps_diagEna	Proc	Running			
91023010	Oct/23/2009 18:50:52	Proc	13	Proc	278003229	Complete	nomSciOps_diagEna	Proc	Running			
91023009	Oct/23/2009 17:59:19	Proc	19	Proc	278003229	Complete	nomSciOps_diagEna	Proc	Running			
					277997612	Complete	nomSciOps_diagEna	Proc	Running			
					277991657	R	Complete	nomSciOps_diagEna	Proc	Running		
91023008	Oct/23/2009 14:08:42	Proc	15	Proc	277991657	R	Complete	nomSciOps_diagEna	Proc	Running	316	Di
					277985681	R	Complete	nomSciOps_diagEna	Proc	Running		
91023007	Oct/23/2009 12:59:36	Proc	15	Proc	277985681	R	Complete	nomSciOps_diagEna	Proc	Running	4231	FM Di Re Me Cal
					277979700	R	Complete	nomSciOps_diagEna	Proc	Running		
91023006	Oct/23/2009 12:12:48	Proc	19	Proc	277979700	R	Complete	nomSciOps_diagEna	Proc	Running	744	FM Di Me Cal
					277973710	R	Complete	nomSciOps_diagEna	Proc	Complete	4231	FM Di Re Me Cal
					277967692	R	Complete	nomSciOps_diagEna	Proc	Running	300	Di
91023005	Oct/23/2009 08:37:23	Proc	15	Proc	277967692	R	Complete	nomSciOps_diagEna	Proc	Running	25 4206	FM Re Me Cal
					277961622	R	Complete	nomSciOps_diagEna	Proc	Running	4231	FM Di Re Me Cal
91023004	Oct/23/2009 07:14:45	Proc	15	Proc	277961622	R	Complete	nomSciOps_diagEna	Proc	Running		
					277955445	R	Complete	nomSciOps_diagEna	Proc	Running	2 4229	FM Di Re Me Cal
91023003	Oct/23/2009 06:13:35	Proc	21	Proc	277955445	R	Complete	nomSciOps_diagEna	Proc	Running	369 1 3861	
					277951581	R	Complete	nomSciOps_diagEna	Proc	Complete	1 4230	FM Di Re Me Cal
					277945852	R	Complete	nomSciOps_diagEna	Proc	Complete	4 4227	FM Di Re Me Cal
91023002	Oct/23/2009 02:41:57	Proc	13	Proc	277945852	R	Complete	nomSciOps_diagEna	Proc	Complete	333 1 3897	
					277940123	R	Complete	nomSciOps_diagEna	Proc	Complete	4231	

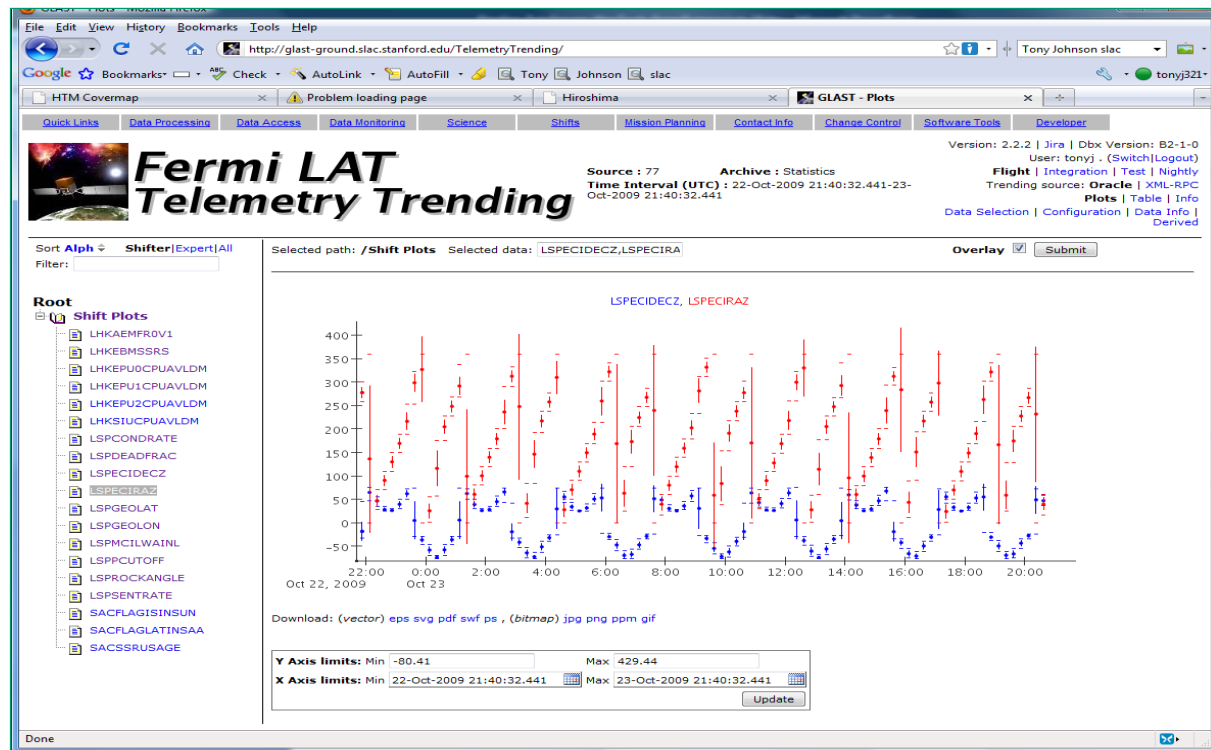
GRB Alerts

Trigger Time	MET	Name	GRB	Notice	Processing	Afterglow	Data
Oct/23/2009 00:29:45	277950585	GRB091023021	FERMI	Proc	Proc		277945852
Oct/22/2009 18:03:28	277927408	GRB091022752	FERMI	Proc	Proc		277922632

ASP Sky Monitor Process

Processing (UTC)	PGWave	DRP	Data	Data Start (UTC)	Frequency
Oct/23/2009 00:18:30	Proc	Proc	Pgwave Drp	Oct/22/2009 18:00:00	six_hours
Oct/22/2009 22:36:08	Proc	Proc	Pgwave Drp	Oct/22/2009 00:00:00	daily
Oct/22/2009 19:00:01	Proc	Proc	Pgwave Drp	Oct/22/2009 12:00:00	six_hours
Oct/22/2009 13:12:11	Proc	Proc	Pgwave Drp	Oct/22/2009 06:00:00	six_hours
Oct/22/2009 13:12:06	Proc	Proc	Pgwave Drp	Oct/22/2009 00:00:00	six_hours
Oct/22/2009 04:00:58	Proc	Proc	Pgwave Drp	Oct/21/2009 18:00:00	six_hours

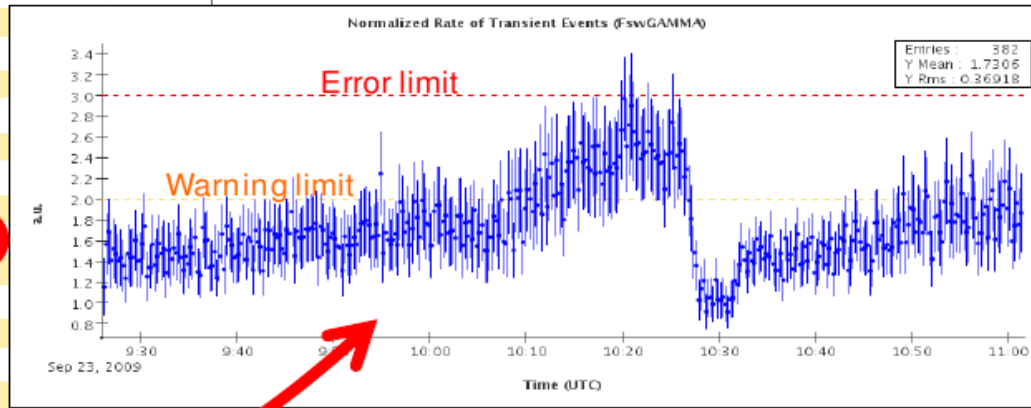
- **Web interface allows**
 - Quick overview of data processing
 - Flags runs requiring further attention
 - Allows “drill-down” to isolate/identify problems



- Web interface allows
 - Dynamic selection of time period
 - Dynamic overlay of quantities
 - Customized tree to draw attention to important plots
 - Can be customized for individuals or groups
- Cross trending of housekeeping and level 1 data

Alarms for run 275390766

Mode	Type	Error	Warning	Undefined	Clean
acdPedsAnalyzer	Hist	0	0	0	14
calGainsAnalyzer	Hist	0	0	0	18
CalPed	Hist	0	0	0	2
calPedsAnalyzer	Hist	0	0	0	
Digi	Hist	0	0	0	
Digi	Trend	0	0	0	
FastMon	Hist	0	0	0	
FastMon	Trend	0	0	0	
fastMonError	Trend	0	0	0	
Merit	Hist	0	0	0	
Merit	Trend	0	1	0	
Recon	Hist	0	0	0	
Recon	Trend	0	0	0	
TkrMon	Trend	0	0	0	
verifyFlt2Error	Trend	0	0	0	
verifyLog	Trend	0	0	0	



WARNING Status

Severity	Mode	Type	Variable Name	Algorithm	Value	Limits	Details
5	Merit	Trend	OutF_NormRateTransientEvts	values	2.97 +- 0.81	[-1.0E10 -1.0E10 --- 2.0 3.0]	View

- Automated alarms are used to alert duty scientists to anomalies
- Use fixed limits and reference histograms
- Many quantities are highly orbit dependent, so particle fluxes, geomagnetic variables must be taken into account
 - 20 different alarm algorithms

Fermi data is immediately available to the public

- **Via Fermi Science Support Center**
- **ISOC supports collaboration data servers which provide**

Full access to data via web based “data catalog”

- **Access to public files plus extended event formats**
- **Search based on arbitrary “meta-data” associated with datasets**

Web based event display for looking at detailed reconstruction of individual events

Web based data selection tools

- **With support for producing Root and Fits files**

Data Catalog Web Interface

Run Min Max Status: ALL
 MET Start Stop
 Filter Clear

Folder /Data/Flight/Level1/LPA Group FT1

7,310 items found, displaying 1 to 500.
 [First/Prev] 1, 2, 3, 4, 5, 6, 7, 8 [Next/Last]

Name	Type	Format	Run Min	Run Max	MET Start	MET Stop	Events	Size	Status	Created (UTC)
r0277985681	FT1	ft	277985681	277985681	277985683.905165	277990271.085179	30,504	2.7 MB	OK	23-Oct-2009 19:56:21
r0277979700	FT1	ft	277979700	277979700	277979702.903274	277984145.083137	23,534	2.1 MB	OK	23-Oct-2009 19:19:55
r0277973710	FT1	ft	277973710	277973710	277973712.964956	277977906.089333	30,101	2.7 MB	OK	23-Oct-2009 18:26:11
r0277967692	FT1	ft	277967692	277967692	277967694.903479	277971873.08714	15,808	1.4 MB	OK	23-Oct-2009 15:28:46
r0277961622	FT1	ft	277961622	277961622	277961624.903399	277969984.084222	28,896	2.6 MB	OK	23-Oct-2009 15:31:20
r0277955445	FT1	ft	277955445	277955445	277955447.916756	277960098.085045	41,867	3.7 MB	OK	23-Oct-2009 17:47:14
r0277951581	FT1	ft	277951581	277951581	277951583.905027	277954232.085327	23,772	2.1 MB	OK	23-Oct-2009 10:34:10
r0277948852	FT1	ft	277948852	277948852	277948854.903315	277951571.085071	64,889	5.7 MB	OK	23-Oct-2009 14:09:05
r0277940123	FT1	ft	277940123	277940123	277940125.911704	277945842.086144	48,507	4.3 MB	OK	23-Oct-2009 13:54:41
r0277934394	FT1	ft	277934394	277934394	277934396.906468	277940113.085254	60,327	5.3 MB	OK	23-Oct-2009 09:33:40
r0277928665	FT1	ft	277928665	277928665	277928667.906623	277934384.085057	47,486	4.2 MB	OK	23-Oct-2009 09:02:42
r0277922632	FT1	ft	277922632	277922632	277922634.903501	277928655.099030	35,059	4.7 MB	OK	23-Oct-2009 05:28:09
r0277917385	FT1	ft	277917385	277917385	277917387.905128	277922501.086144	39,463	3.5 MB	OK	23-Oct-2009 05:28:09
r0277911633	FT1	ft	277911633	277911633	277911635.907119	277916406.085126	37,451	3.3 MB	OK	23-Oct-2009 05:28:09

Fermi LAT Data Portal Catalog

Catalog version 1.9 | Site | Portal Version 3.1 | Site

User: tomy (Switch/Logout) Config: OnOrbit

Model: [Read | Dev | Test]

View: [Tree | Data Types | File Formats | Messages | Admin | Problems]

Buttons: Welcome, Catalog, Mark Skimmer, File Skimmer, Astro Server, Wired, History

Folder /Data/Flight/Level1/LPA Group FT1

FT1 files from level 1 processing of on-orbit data. Edit description

Created (UTC): 25-Jun-2008 16:27:11

Run Min: 238084237
 Run Max: 277985681
 Files: 7310 (Errors 38)
 Events: 227,354,343
 Size: 16.9 GB
 Data Type: FT1

Meta-data

Name	Value	Type
astroDB-LEOScience	true	STRING
astroDB-Level1	true	STRING
FT1Skim	Level 1 LPA data	STRING
L1_Pc_public_v1	true	STRING
nKeyData	20	NUMBER

Buttons: Edit meta-data

Folder /Data/Flight/Level1/LPA Group FT1

Dataset r0277967692 version 0

Standard Data

Name	Value
Created (UTC):	23-Oct-2009 15:28:46
Run Min:	277967692
Run Max:	277967692
Events:	15,808
Size:	1.4 MB
Format:	ft
Type:	FT1
Source:	PIPELINE
Task:	doRun
Links	Download History

Meta-data

Name	Value	Type
L1_Pc_public_v1	true	STRING
nDownlink	91023095	NUMBER
nMetStart	277967694.903479	NUMBER
nMetStop	277971873.08714	NUMBER
nHookKey	2557	NUMBER
nRun	277967692	NUMBER
cCreator	L1Proc-1.79	STRING
cDataSource	LPA	STRING
sIntent	nomSOOps_diagEna	STRING

Location

Site	Status	Checked (UTC)	Location
SLAC	OK	23-Oct-2009 16:01:39	/afs/fermi/g/glast/u20/FT1-20copies/glast/Data/Flight/Level 1/LPA/prod/1.79/ft1/glt_ph_r0277967692_v000.ft
SLAC_XROOT	OK	23-Oct-2009 15:29:44	root://glast-01.slac.stanford.edu/glast/data/Flight/Level 1/LPA/prod/1.79/ft1/glt_ph_r0277967692_v000.ft

Drill down to get more details

GLAST Download Manager

File Edit Help

Buttons: Add..., Remove, Details, Clean Up, Start downloading, Pause downloading

File Name	File Size	Status
r0236084237_ft1.ft	3.8 MB	Done
r0236090205_ft1.ft	2.9 MB	Downloading
r0236096298_ft1.ft	513.1 kB	Queued
r0236102471_ft1.ft	478.1 kB	Queued
r0236108455_ft1.ft	689.1 kB	Queued
r0236121367_ft1.ft	191.2 kB	Queued
r0236126697_ft1.ft	174.4 kB	Queued
r0236135175_ft1.ft	12.4 MB	Queued

Progress

Total Download Size: 16.9 GB
 Downloaded So Far: 5.5 MB
 Download Rate: 270.2 kB/sec
 Time Remaining: 18:11:49

Download manager, reliable download of multiple files

“Astro” Server Web Interface

“duplicates” for use internal to the collaboration the NASA FSSC official data server

The P6_public_v1 event sample currently contains 190,185,596 events covering the time period 2008-08-04 15:43:36 UTC (239,557,418 MET) to 2009-10-22 11:23:53 UTC (277,903,436 MET) .

Number of events selected: 369311

Parameter	Value
Job Name	%u-%t-%n Arbitrary name: %u=user name, %t=job type, %n=unique id
Event Sample	P6_public_v1 Event selection help
Energy Range	Min: <input type="text"/> Max: <input type="text"/> MeV (Leave blank for no limit)
Time Range	Min: <input type="text"/> Max: <input type="text"/> Mission elapsed time (MET) (Leave blank for no limit)
Position	RA: <input type="text"/> DEC: <input type="text"/> degrees (Leave blank for full sky) or astronomical object: <input type="text"/> using NED overrides ra, dec above help
Radius	<input type="text"/> degrees
Event Class	Diffuse
Output (FT2 Files)	<input checked="" type="checkbox"/> 30 second (fits) <input type="checkbox"/> 1 second (fits)
Output (Event Data)	<input checked="" type="checkbox"/> FT1 (fits) <input type="checkbox"/> LS1 (fits) <input type="checkbox"/> Merit (root) <input type="checkbox"/> Event-List (text)
Debug Mode	False
User Comment	LS I 61+303
Expert Options	help

Parameter	Value
Job Name	%u-%t-%n
Event Source	P6_public_v1
Minimum energy	
Maximum energy	
Minimum MET	
Maximum MET	
RA	40.1
DEC	61.225
Galactic Object	
Radius	10.0
Event Class	Diffuse
Output (FT2 Files)	30-second
Output (Event Data)	FT1
Debug	false
User Comment	LS I 61+303
Expert Options	

Astro job submitted

Your job tonyj-AstroServer-00040 has been submitted.

Your data will be available for download from <ftp://ftp-glast.slac.stanford.edu/glast.u27/DataServer/1256243366055>

You will be sent an e-mail at tonyj@slac.stanford.edu when your job has completed.

You can monitor your job's progress using the [Pipeline](#)

Note: Clicking on the Status column will take you to the pipeline task that ran the job. Clicking on the Job column will allow you to rerun this task, or a similar one. Clicking on the Output Directory column will take you to the output.

Submit Time	Job	User	Task Type	Status	Output Directory	User Comment
22-Oct-2009 13:31	tonyj-AstroServer-00040	tonyj	AstroServer	Success	View dir	LS I 61+303
22-Oct-2009 00:01	andrig_ana_4	andrig	AstroServer	Success	View dir	
22-Oct-2009 07:30	andrig_ana_3	andrig	AstroServer	Success	View dir	
22-Oct-2009 07:22	andrig_ana_3	andrig	AstroServer	Success	View dir	
22-Oct-2009 00:36	Aug60_Oct06_2006_10953p0755	bjarb	AstroServer	Success	View dir	
21-Oct-2009 13:46	Abdo-FT2-Oct21	abdo	AstroServer	Success	View dir	
21-Oct-2009 09:59	Abdo-AISky_6	abdo	SimpleSkimmer	Failed	View dir	
21-Oct-2009 08:12	andrig_ana_2	andrig	AstroServer	Success	View dir	
21-Oct-2009 08:12	andrig_ana_1	andrig	AstroServer	Success	View dir	
21-Oct-2009 01:01	bernd-AstroServer-0004	bernd	AstroServer	Success	View dir	
20-Oct-2009 14:32	bernd-SimpleSkimmer-0002	bernd	SimpleSkimmer	Success	View dir	Pass7.2 Diffuse - Aug-Sept
20-Oct-2009 14:18	uchiyama-1915-13mm	uchiyama	Astro	Failed	View dir	
20-Oct-2009 14:03	bernd-SimpleSkimmer-0001	bernd	SimpleSkimmer	Success	View dir	Pass7.2 Diffuse - Aug-Sept - Per Pass7.2 validation
20-Oct-2009 08:37	guillermo-AstroServer-0003	guillermo	AstroServer	Success	View dir	
20-Oct-2009 08:36	guillermo-AstroServer-0002	guillermo	AstroServer	Success	View dir	

[Up to higher level directory](#)

Name

- tonyj-AstroServer-00040-README.html
- tonyj-AstroServer-00040-ft1.fits
- tonyj-AstroServer-00040-ft2-30s.fits

Size Last Modified

4 KB	10/22/2009 8:31:00 PM
33351 KB	10/22/2009 8:39:00 PM
144206 KB	10/22/2009 8:46:00 PM

Failure rates have to be very low when running hundreds of thousands of jobs!

- **use local disks as much as possible, staging input and output files back to networked storage.**
- **nfs has been fragile when asked to handle hundreds of simultaneous connections**

Defined pre-emptive queues on 900 cores for prompt processing : if all jobs slots busy, someone else's jobs are suspended

Auto-rollback of failed jobs by pipeline : Biggest failure rate is transient errors

Auto-restart sick batch machines : Sick machines are sinks for failing jobs since they quickly return to an available state

Control access to the database : We limit access to the pipeline server for bulk processing

Be vigilant and give yourself time to react :

- **nagios with warning levels on disk space to give us heads up on imminent disasters**
- **ganglia to watch loads**

Test ahead of time! We used full scale operations simulations and data challenges to stress the systems

- The LAT collaboration develops and manages 2 software
 - ***GlastRelease*** : the workhorse underneath the pipeline (simulation, digitization, reconstruction, rejection & classification)
 - ***ScienceTools*** : the analysis toolkit to do science with the LAT data
- GlastRelease is internal to the collaboration
- ScienceTools is a deliverable as per NASA mandate (remember that Fermi is an 'observatory' : the data are public, and so are the means to analyze them)
- Both software share codes, but their requirements are different
 - **Different data products**
 - All detailed/structured files to be in ROOT
 - High level analysis only in FITS
 - **Different code structures**

Package	Description	Provider
ACD, CAL, TKR Recon	Data reconstruction	LAT
ACD, CAL, TKR Sim	Instrument sim	LAT
GEANT4	Particle transport sim	G4 worldwide collaboration
xml	Parameters	World standard
Root 5.26.00	C++ object I/O	HEP standard
Gaudi	Code skeleton	CERN standard
doxygen	Code doc tool	World standard
Visual C++/gnu	Development envs	World standards
CMT (now Scons)	Code mgmt tool	HEP standard
ViewCvs	cvs web viewer	World standard
cvs	File version mgmt	World standard

Output of a telescope (for “event” data) :

- **Location on sky, time, energy, quality**
- **Very simple output tuple : implemented in Root and FITS**

NASA mandates that all its space missions use FITS as a data format (In use for 25+ years)

Format fully documented and files self documenting

- **File headers are an integral part of the format**
- **Interface library supplied for popular languages**

And that the data/software be made public (1 year hiatus negotiated)

Funds a Science Support Center to interface to the public :

- **Instrument teams not asked to do this (10+ FTEs for Fermi)**
- **Charter is to support the Fermi data “forever”**

The analysis model is atomic operations with files used as the messengers

- **Linked together with scripts (called pipelines)**
- **Designed for sequential simple processing**
 - **Not made for (parallel) batch processing**
 - **Very easy to step on the parameter files**
 - **No structures/objects in the files : Images or tables**

FITS headers are nice

- **Intrinsic to the files : all tools respect the header**
- **Carries processing history etc etc**

LAT collaboration mitigation : ScienceTools to python!

Divide code into packages (small is beautiful) :

- **CVS for versioning (with web access) : obsolete but legacy**
- **CMT (now scons) for code management (cross-platform mandatory)**

Extensive use of electronic communications : Web conferencing (evo), Instant Messaging (icq)

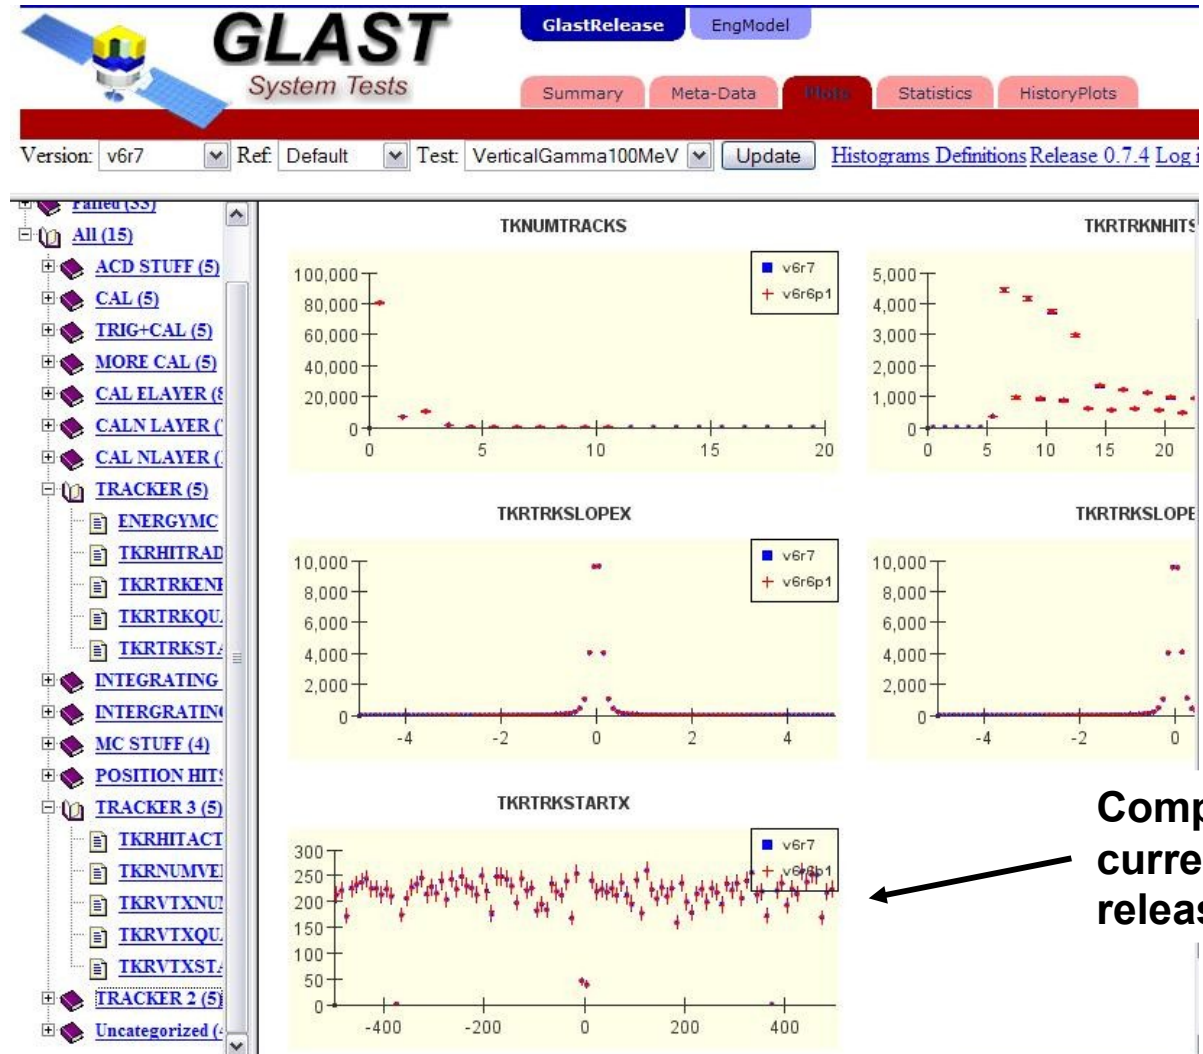
Software Development on Windows or linux; compute cycles on linux

Documentation and coding reviews enforce coding rules

Most important : “Continuous integration”

- **Eliminate surprises for incoming code releases**
- **Build code when packages are tagged; alert owners to failures in build or running of unit tests. Results tracked in database.**
- **Developing comprehensive system tests in multiple source configurations. Track results in database; web viewable.**

System Tests




Comparison of current to previous release.




Conclusions

- **Fermi data pipeline, data catalog and monitoring tools in production use for 4 years**
 - **Have proved very reliable for data processing**
 - **Web based tools : monitoring load distributed world wide**
- **Design decision : avoid tight coupling to specific experiment**
 - **EXO, CDMS, being evaluated for use by JWST**
- **Support submission of pipeline jobs to more systems**
 - **Currently support LSF (SLAC), BQS (Lyon), Condor**
 - **Adding support for Grid Engine, EGE Grid**
- **Increase interactivity of web apps (AJAX, Web 2.0, GWT)**




Fermi LAT
Pipeline
processing
at IN2P3



More slides.....



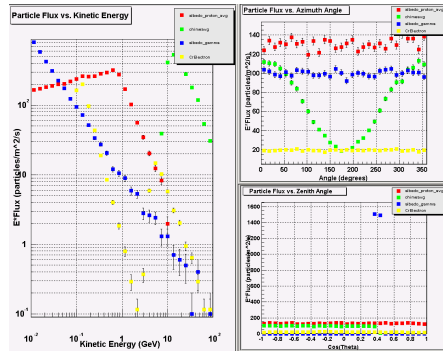
3 GeV gamma interaction




Source Fluxes

Particle Transport

Instrument data



“Raw” Data




3 GeV gamma recon

Recon

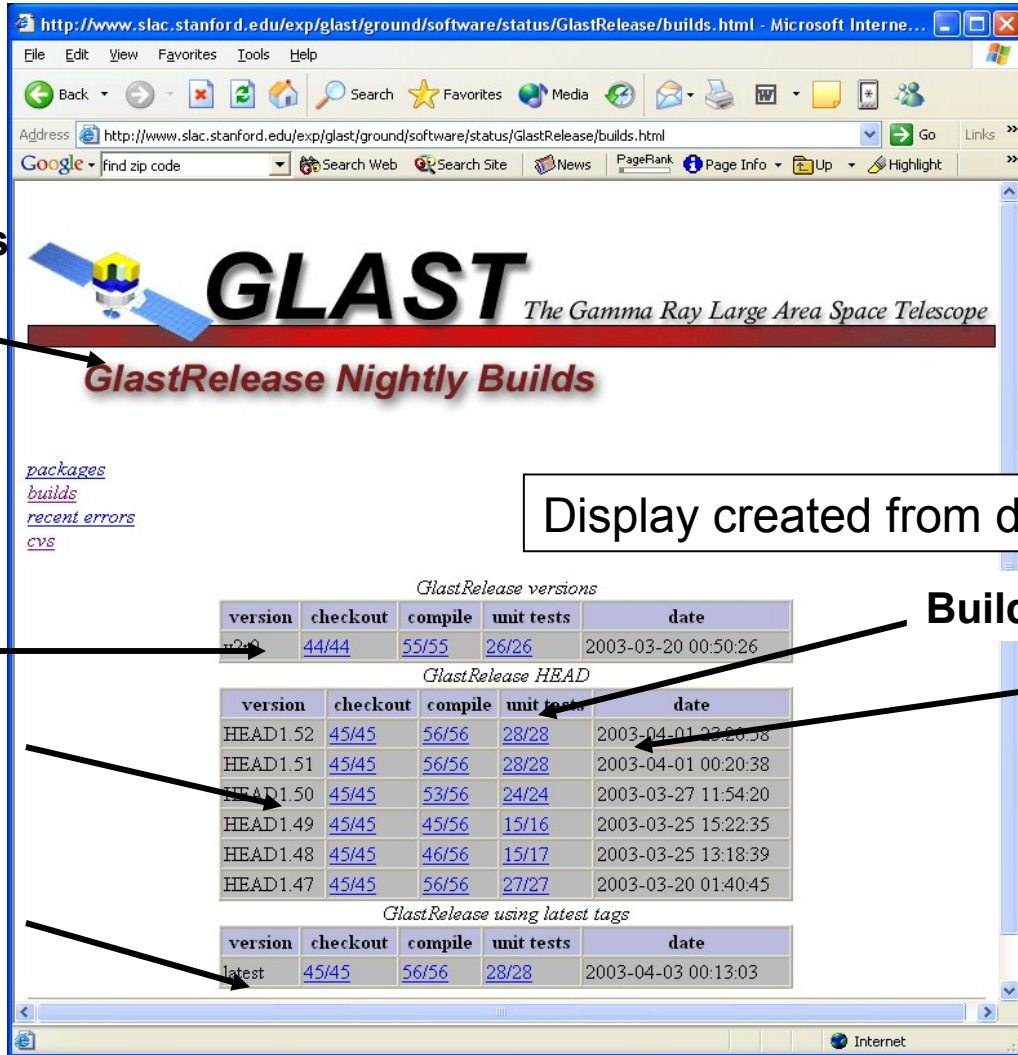
Geometry

Background Rejection
-
Particle ID

Full geometry in xml with C++ interface
G4 discovers instrument from the xml



CAL Detail



GLAST *The Gamma Ray Large Area Space Telescope*

GlastRelease Nightly Builds

[packages](#)
[builds](#)
[recent errors](#)
[cvs](#)

GlastRelease versions

version	checkout	compile	unit tests	date
v2.0	44/44	55/55	26/26	2003-03-20 00:50:26

GlastRelease HEAD

version	checkout	compile	unit tests	date
HEAD.1.52	45/45	56/56	28/28	2003-04-01 03:20:38
HEAD.1.51	45/45	56/56	28/28	2003-04-01 00:20:38
HEAD.1.50	45/45	53/56	24/24	2003-03-27 11:54:20
HEAD.1.49	45/45	45/56	15/16	2003-03-25 15:22:35
HEAD.1.48	45/45	46/56	15/17	2003-03-25 13:18:39
HEAD.1.47	45/45	56/56	27/27	2003-03-20 01:40:45

GlastRelease using latest tags

version	checkout	compile	unit tests	date
latest	45/45	56/56	28/28	2003-04-03 00:13:03

Performing builds for Science Tools also

Display created from database query

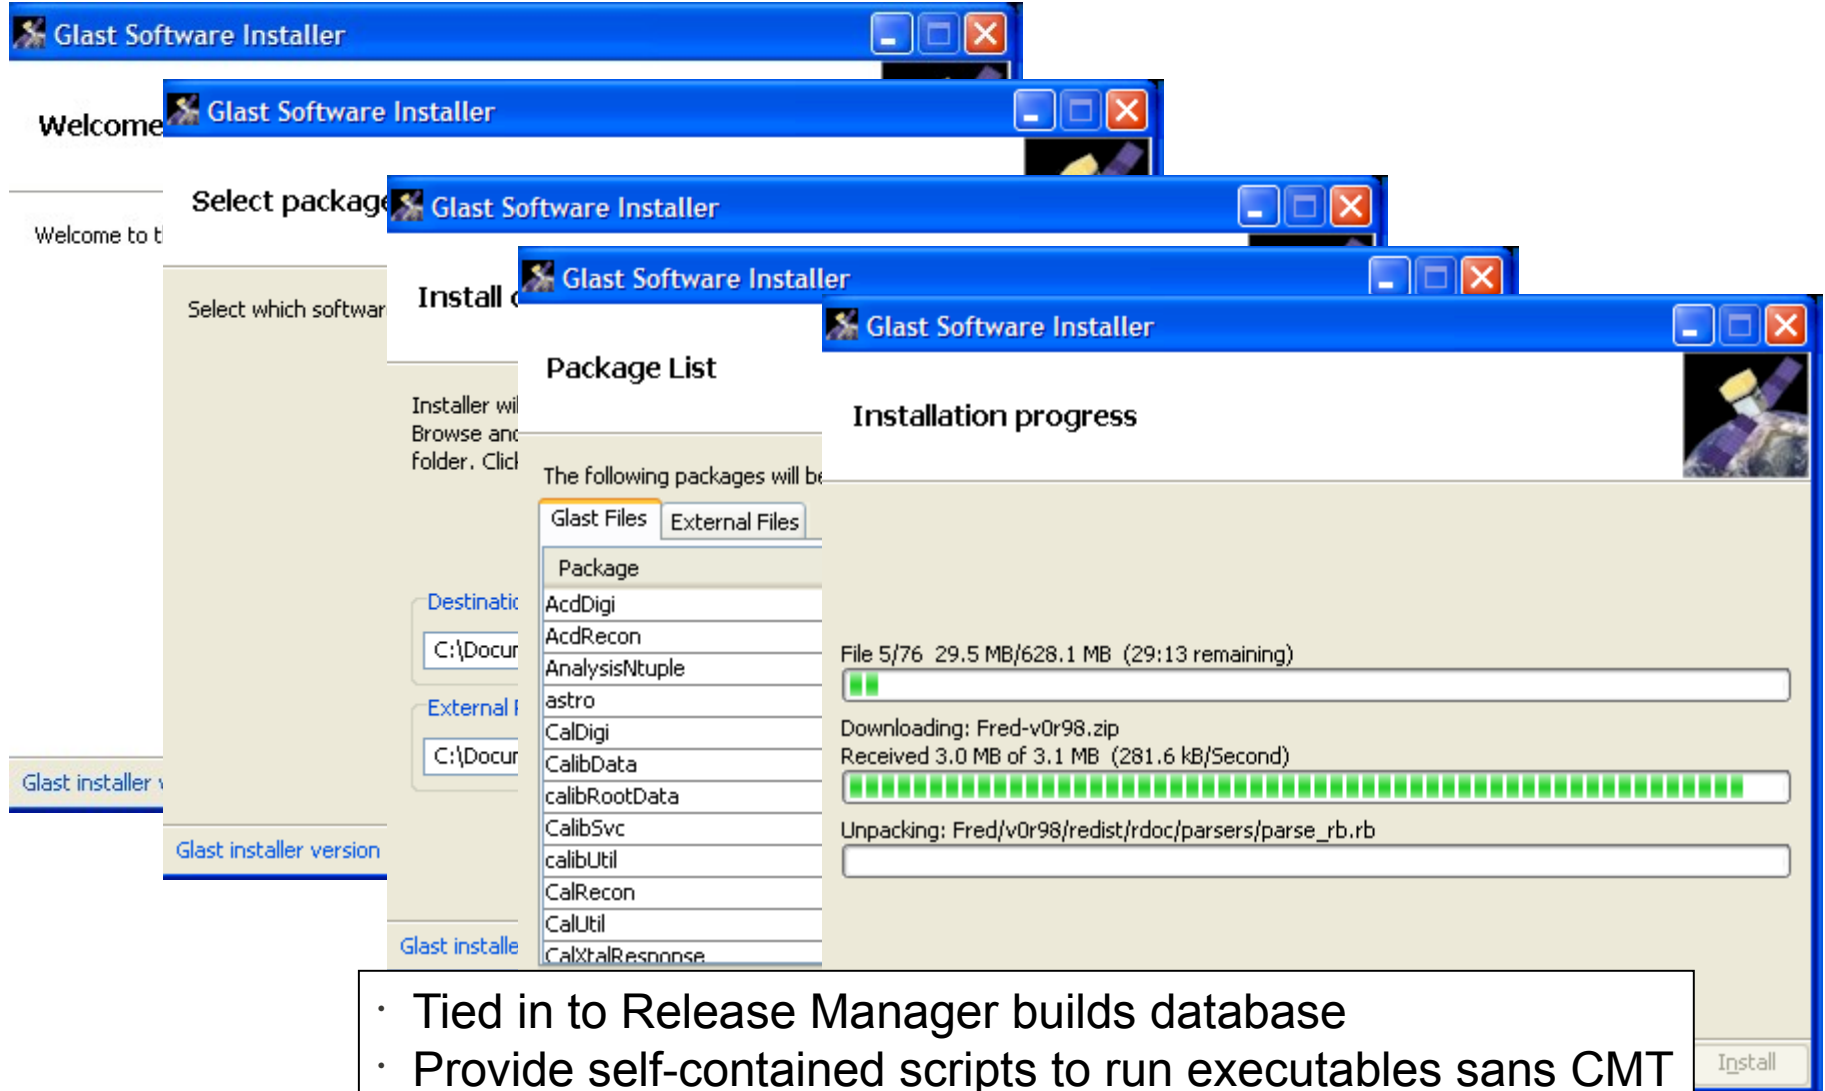
Build status

Unit test status

Past release

Release in progress

Future release



The screenshot shows the Glast Software Installer interface. The main window is titled "Glast Software Installer" and displays the "Package List" and "Installation progress" sections.

Package List

The following packages will be installed:

Package
AcqDigi
AcqRecon
AnalysisNtuple
astro
CalDigi
CalibData
calibRootData
CalibSvc
calibUtil
CalRecon
CalUtil
CalXtalResponse

Installation progress

File 5/76 29.5 MB/628.1 MB (29:13 remaining)

Downloading: Fred-v0r98.zip
Received 3.0 MB of 3.1 MB (281.6 kB/Second)

Unpacking: Fred/v0r98/redist/rdoc/parsers/parse_rb.rb

Buttons: "Glast installer", "Glast installer version", "Glast installer", "Install"

- Tied in to Release Manager builds database
- Provide self-contained scripts to run executables sans CMT

User JIRA web issues tracker

Commercial product

Handles bugs, features, improvements

Full user/group management

“roadmaps” for version evolution

Change Control Board: Data Processing Operations

Code used in pipeline – sim/recon; executive scripts; pipeline itself

Require documentation of all changes – preferably backed up by JIRA issues

Demonstration that fixes work; system tests on sim/recon

Using wiki tool to record actions

4-person board – adjudicated by email so far

Wiki

Confluence – Fermi now totally addicted to it

Commercial product (Atlassian – same parent as JIRA)

Simple web editing independent of user OS

Space management; same groups and users as JIRA

Ground software is amalgam of HEP instrument software and Astro FTOOLS

Adopt HEP's "Data Challenges" to create a series of end-to-end studies: create a progression of ever more demanding studies

Two were performed before launch

Complete buy-in from the Science Working Groups

Deepest sky simulation at GeV energies at the time

Full scale operations simulation performed 3 months before launch

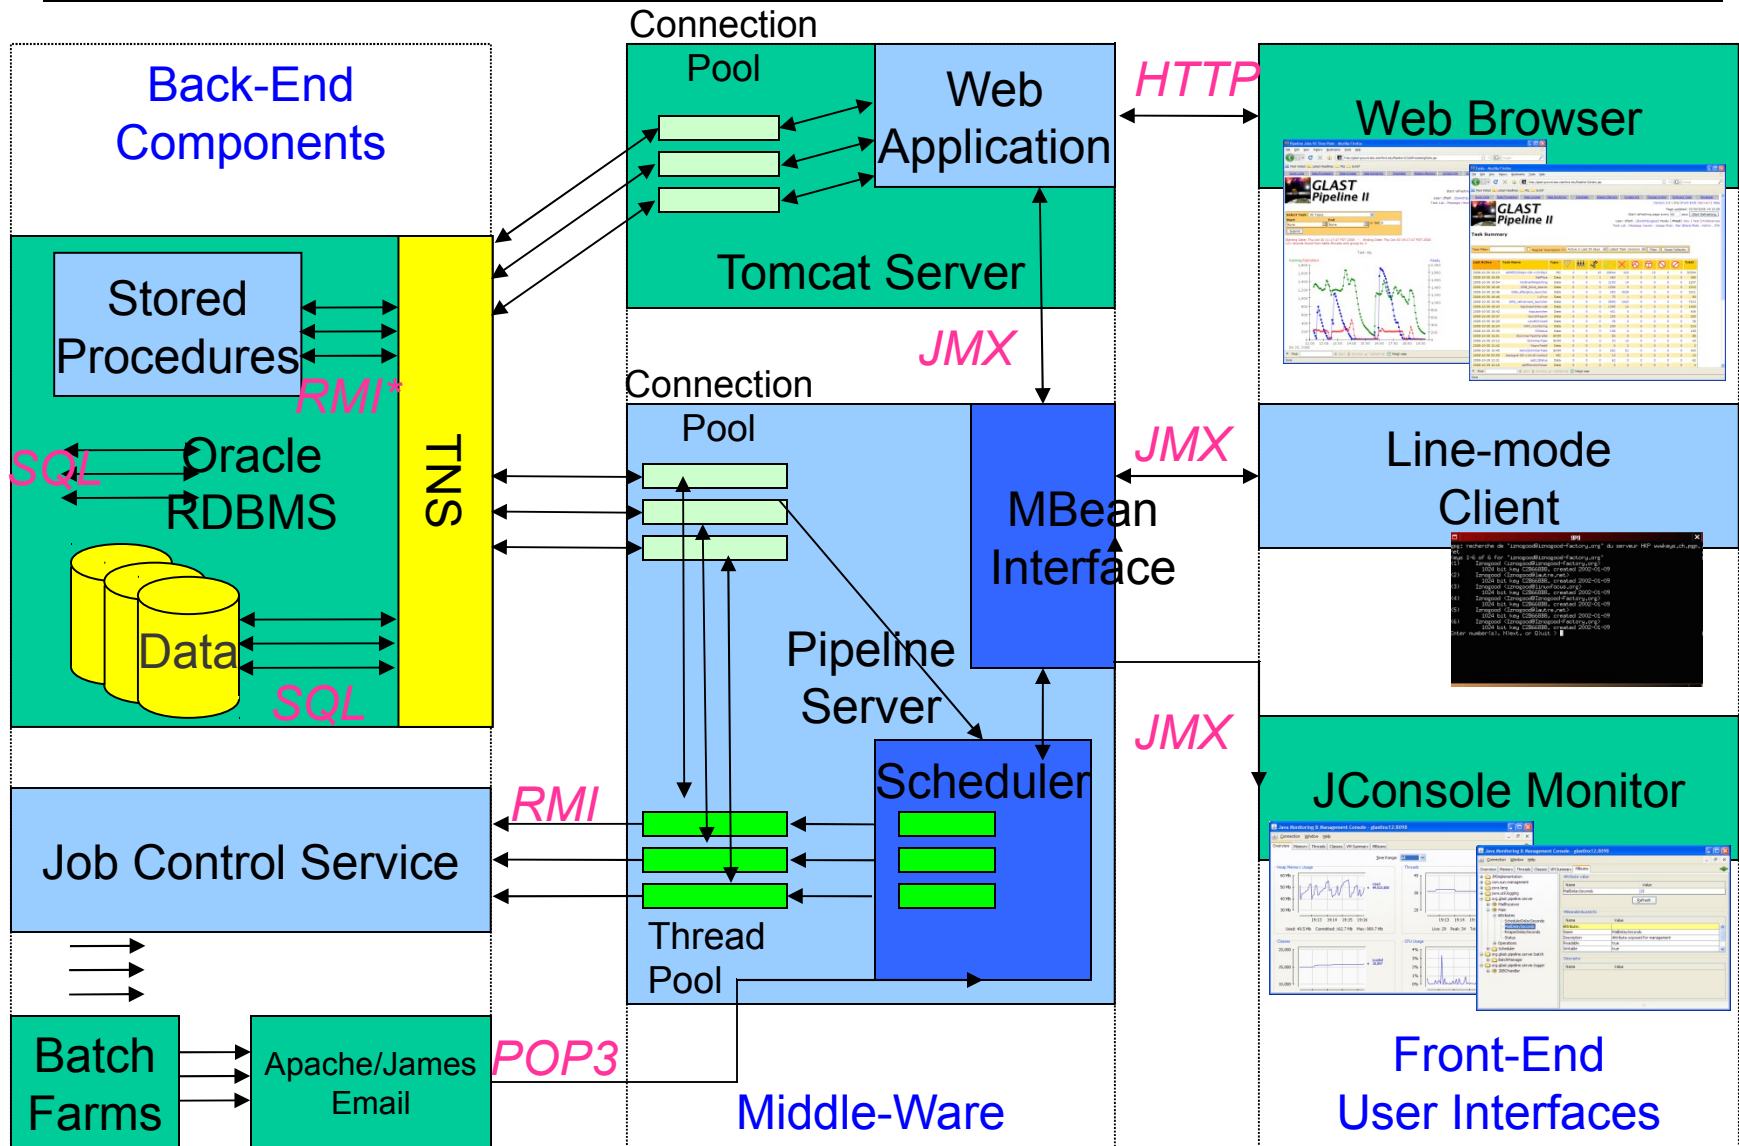
Simulated 1 full week of raw data and fed it to the operations center in scheduled "downlinks"

Ran data quality shifts etc etc


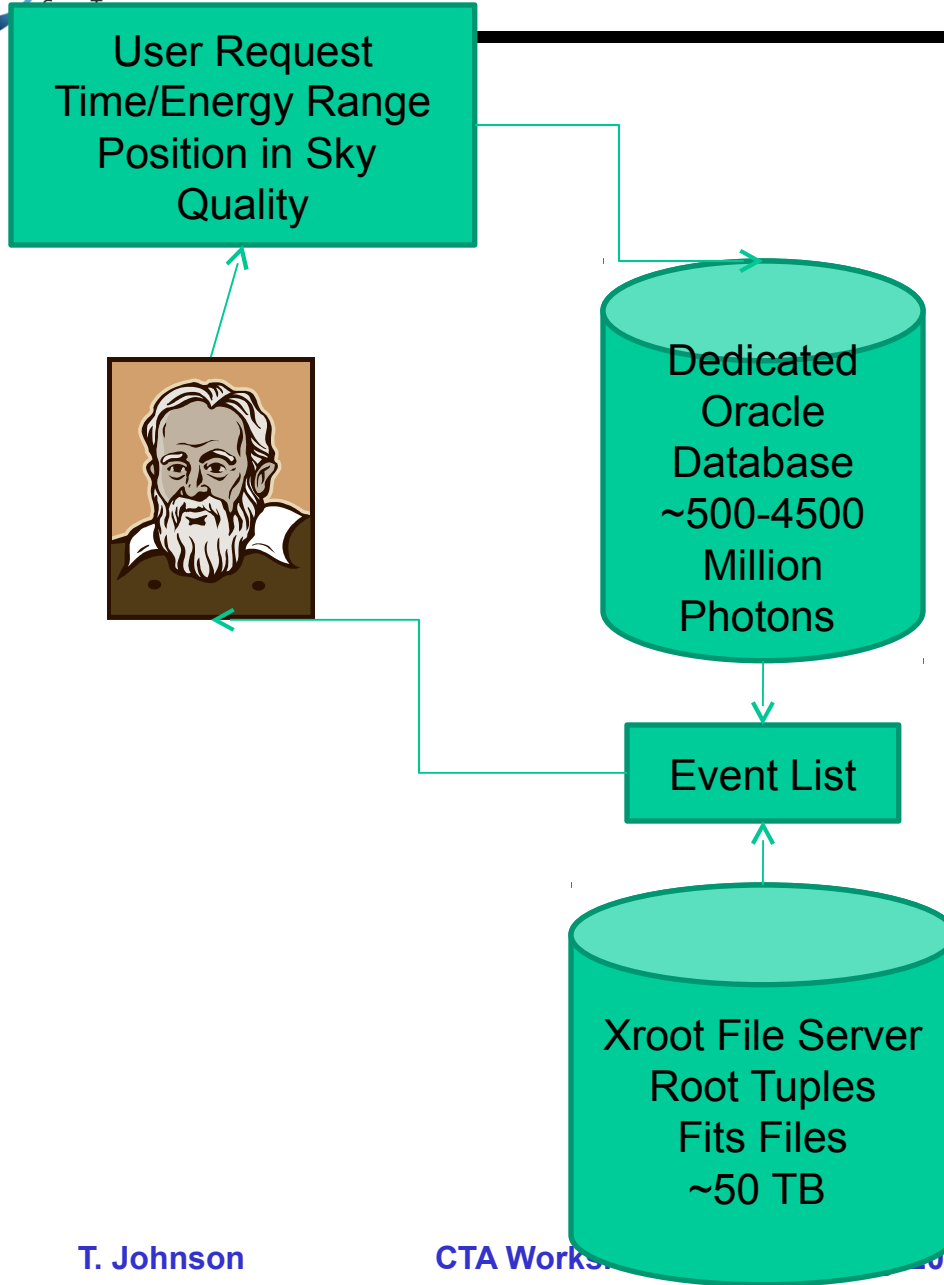
Uncovered many operational problems

- All solved in time for launch and a flawless startup.


Pipeline Implementation



"Astro" Server Implementation



Within the database events are indexed by time, energy and position using a hierarchical triangular mesh (HTM). Database partitions are used to split the data into 1 week time bins and 32 position bins within each time bin, each containing 1024 HTM regions (shown above). The use of HTM triangles makes it easy to identify which regions are entirely contained in the user request, and which are partially contained and require finer selection (below)



Database : Oracle

- **Java Stored Procedures for performance**
- **Scheduled Server-side Jobs for monitoring, stats-gathering**
- **Hierarchical queries**

Servers and Client Libraries (Pipeline, Data Catalog) : Java

- **Extensive use of threads, concurrency utilities for performance**

Jython interpreter for user scripts, **JMX MxBean** Interfaces for monitoring/comm.

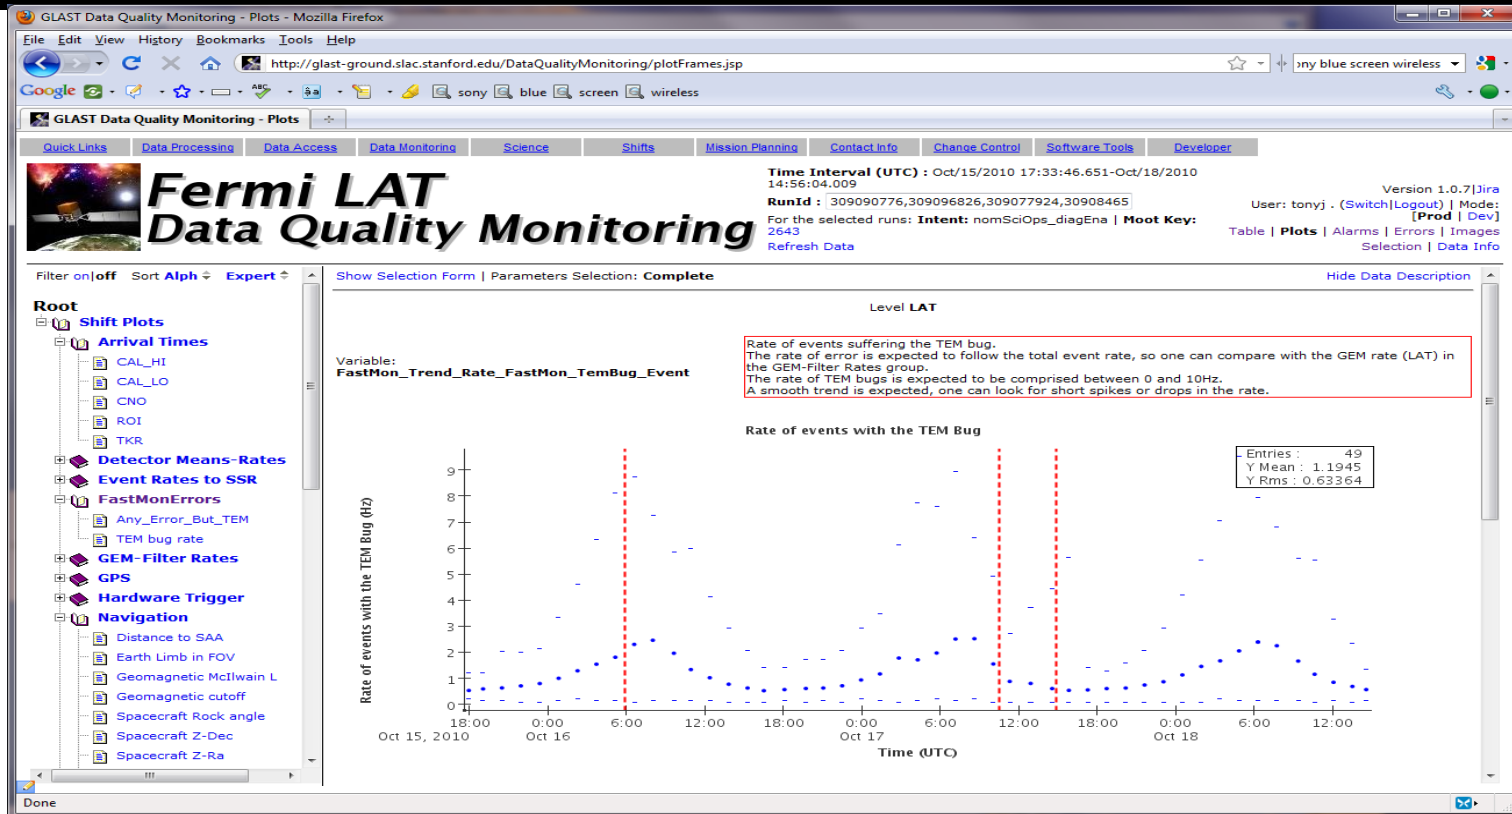
XML used for processing-task definitions

Pipeline : Apache/Tomcat servers for Web, **Apache/James Email** server for status notification

JSP for web pages

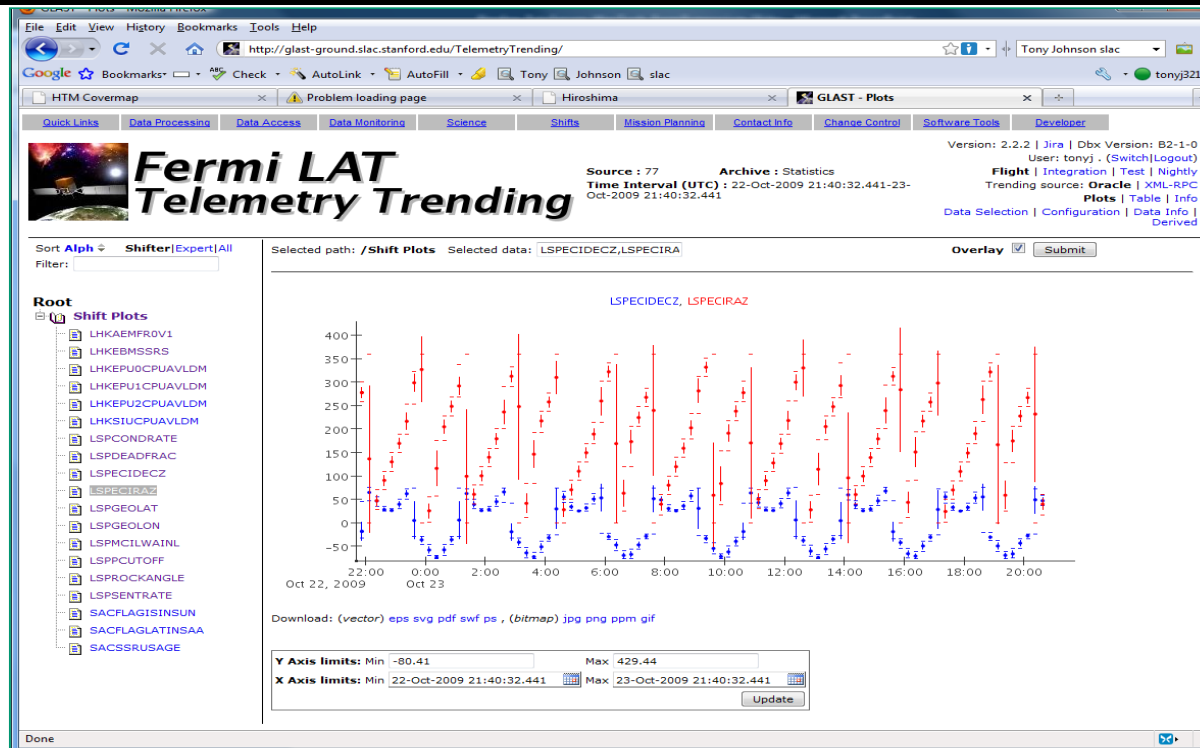
- **DisplayTag** for tabular data
- **AIDA tag libraries** for plotting
- **Custom tag libraries** expose Pipeline client methods

Java Servlets : Serve GraphViz State diagrams



Web interface allows

- Show data from single run or aggregate set of runs
- View description of each plot, View/Print multiple plots
- Customized tree to draw attention to important plots
- Can be customized for individuals or groups

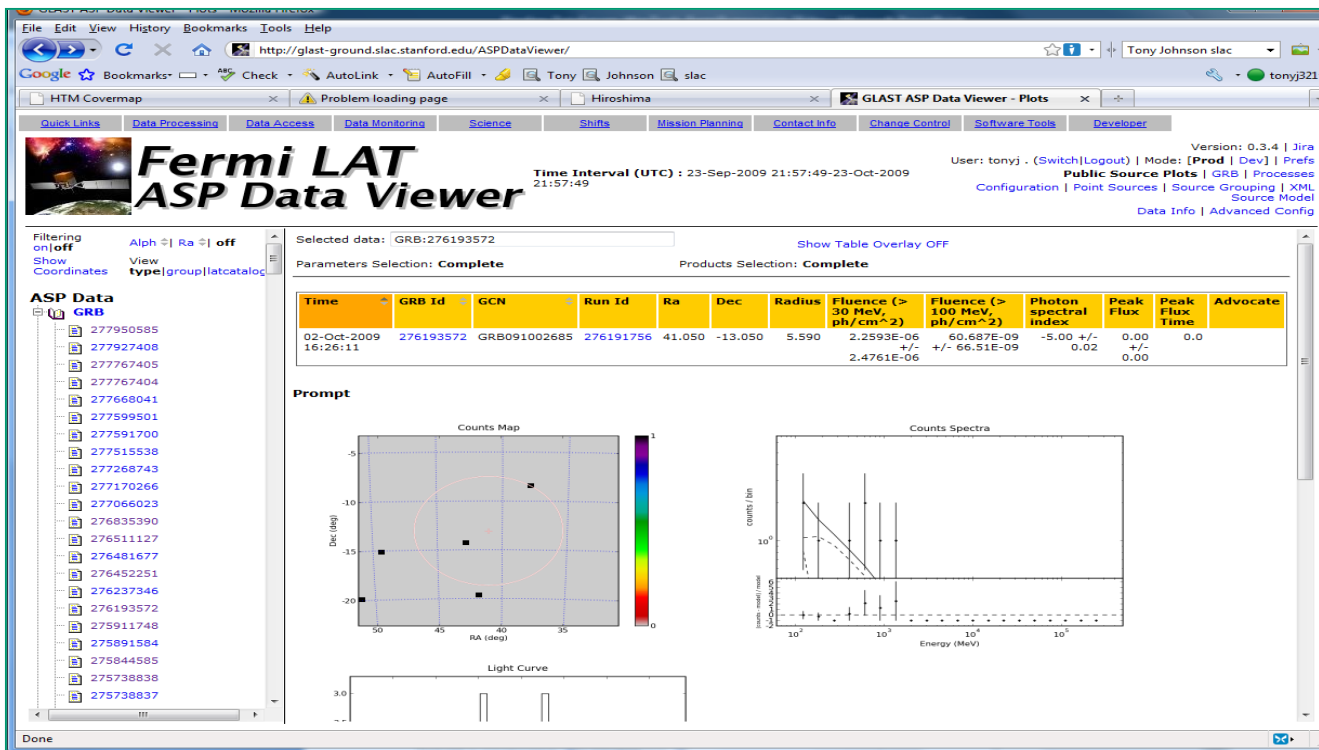


Web interface allows

- Dynamic selection of time period
- Dynamic overlay of quantities

Customized tree to draw attention to important plots

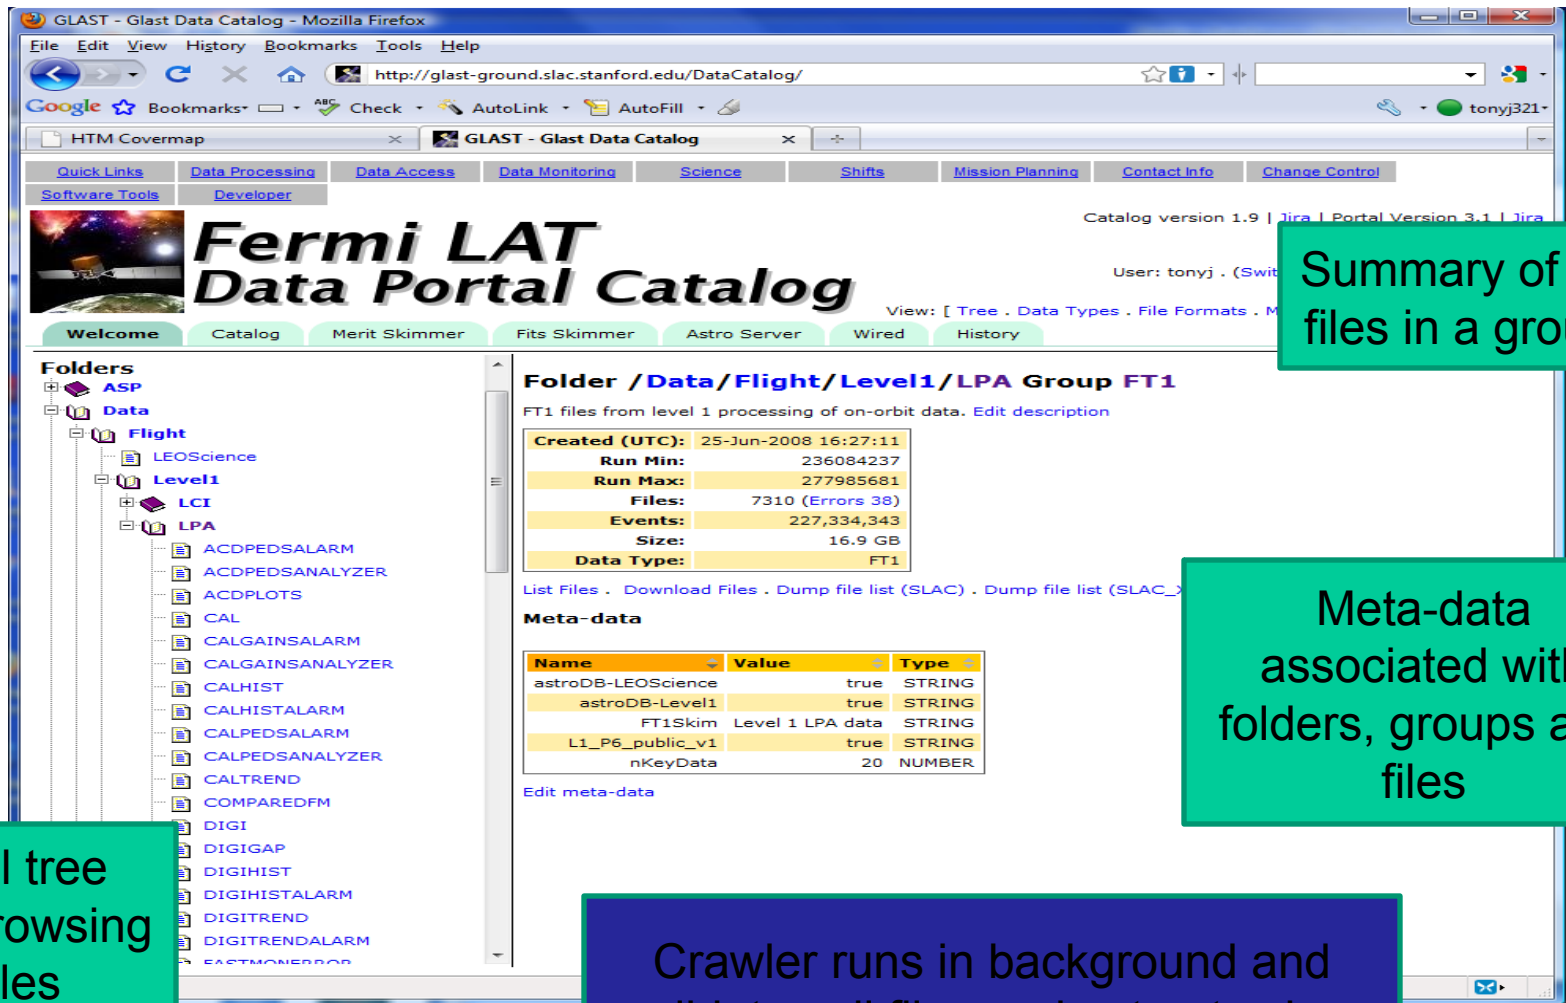
Cross trending of housekeeping and level 1 data



Used to rapidly detect Gamma Ray Bursts or other flaring events

Enabled timely notification of interesting events to external astrophysical community

Data Catalog Web Interface

GLAST - Glast Data Catalog - Mozilla Firefox

http://glast-ground.slac.stanford.edu/DataCatalog/

Quick Links | Data Processing | Data Access | Data Monitoring | Science | Shifts | Mission Planning | Contact Info | Change Control

Software Tools | Developer

Fermi LAT Data Portal Catalog

Catalog version 1.9 | Jira | Portal Version 3.1 | Jira

User: tonyj . (Switch)

View: [Tree . Data Types . File Formats . M]

Welcome | Catalog | Merit Skimmer | Fits Skimmer | Astro Server | Wired | History

Folders

- ASP
- Data
 - Flight
 - LEOScience
 - Level1
 - LCI
 - LPA
 - ACDPEDSALARM
 - ACDPEDSANALYZER
 - ACDPLOTS
 - CAL
 - CALGAINSALARM
 - CALGAINSANALYZER
 - CALHIST
 - CALHISTALARM
 - CALPEDSALARM
 - CALPEDSANALYZER
 - CALTREND
 - COMPAREDFM
 - DIGI
 - DIGIGAP
 - DIGIHIST
 - DIGIHISTALARM
 - DIGITREND
 - DIGITRENDALARM
 - EASTMONERROR

Folder /Data/Flight/Level1/LPA Group FT1

FT1 files from level 1 processing of on-orbit data. Edit description

Created (UTC):	25-Jun-2008 16:27:11
Run Min:	236084237
Run Max:	277985681
Files:	7310 (Errors 38)
Events:	227,334,343
Size:	16.9 GB
Data Type:	FT1

List Files . Download Files . Dump file list (SLAC) . Dump file list (SLAC)

Meta-data

Name	Value	Type
astroDB-LEOScience	true	STRING
astroDB-Level1	true	STRING
FT1Skim	Level 1 LPA data	STRING
L1_P6_public_v1	true	STRING
nKeyData	20	NUMBER

Edit meta-data

Summary of all files in a group

Meta-data associated with folders, groups and files

Logical tree allows browsing for files Folders and Groups

Crawler runs in background and validates all files and extracts size, #events, etc

Command line tools for direct or scripted interaction with middle-ware

Control Server

- Ping
- Restart
- Shutdown

Upload Task Definitions

Manage processing streams

- Create
- Delete
- Cancel
- Retry from failure point

Query processing history

Plus Interaction with Data Catalog

Provides all administrative functions available in a user-friendly interactive GUI

Interactive displays show active (and historical) processing

Filtering by Task, Process, Status(es), Stream-range, Date-range

Processing Statistics Plots

Provided by AIDA tag library

System throughput plots

- [Filterable by Task, Date-Range](#)

Individual process statistics plots


- [CPU time \(vs Wallclock\)](#)
- [Pending time](#)
- [By Batch Host-type](#)

Task diagrams generated by GraphViz and image-mapped to provide links to task element (Sub-tasks, processes) displays


Summary

recon

Wall Clock time (mins)




Pending time (mins)




PLOTS by BATCH NODES


Wall Clock time (mins)




Pending time (mins)




CPU time (mins)




CPU time/Wall Clock



CPU time (mins)



CPU time/Wall Clock



Task doChunk Process fastMonHist Stream 81031004.247117860.15155

Type Batch
 Status Running
 Stream [81031004.247117860.15155](#)
 CreateDate 30-Oct-2008 22:46:16.000
 SubmitDate 30-Oct-2008 23:14:54.482
 StartDate 30-Oct-2008 23:14:59.000
 EndDate
 CPU Used
 Memory Used
 Swap Used
 Execution Host boer0105
 Exit Code
 Working Dir [/nfs/farm/g/glast/u15/pipeline-II/prod/log/L1Proc/1.68/doRun/doChunk/fastMonHist/081xxxxxx/031xxx/004/247xxxxxx/117xxx/860/015xxx/155](#)
 Log File [/nfs/farm/g/glast/u15/pipeline-II/prod/log/L1Proc/1.68/doRun/doChunk/fastMonHist/081xxxxxx/031xxx/004/247xxxxxx/117xxx/860/015xxx/155/logFile.txt](#)
 Execution Number **1**
 Retry Number **0**
 Is Latest 1
 Batch Job ID [423937](#)

Links: [View Messages](#)

Variables

Nothing found to display.

- Show UpStream Process Instances
- Show Downstream Process Instances
- Show Created SubStreams

Upstream Process Instances

Wait Condition	Process	Status	Type	Created	Submitted	Started	Ended	Job Id	CPU	Host	Links
Success	fastMonTuple	Success	Batch	30-Oct-2008 22:46:16	30-Oct-2008 22:46:26	30-Oct-2008 22:46:33	30-Oct-2008 23:14:31	421514	1646	bali0244	Messages : Log : Files

Downstream Process Instances

Wait Condition	Process	Status	Type	Created	Submitted	Started	Ended	Job Id	CPU	Host	Links
Success	checkChunk	Waiting	Batch	30-Oct-2008 22:46:17							Messages
DONE	mergeFastMonHist	Waiting	Batch	30-Oct-2008 22:45:22							Messages