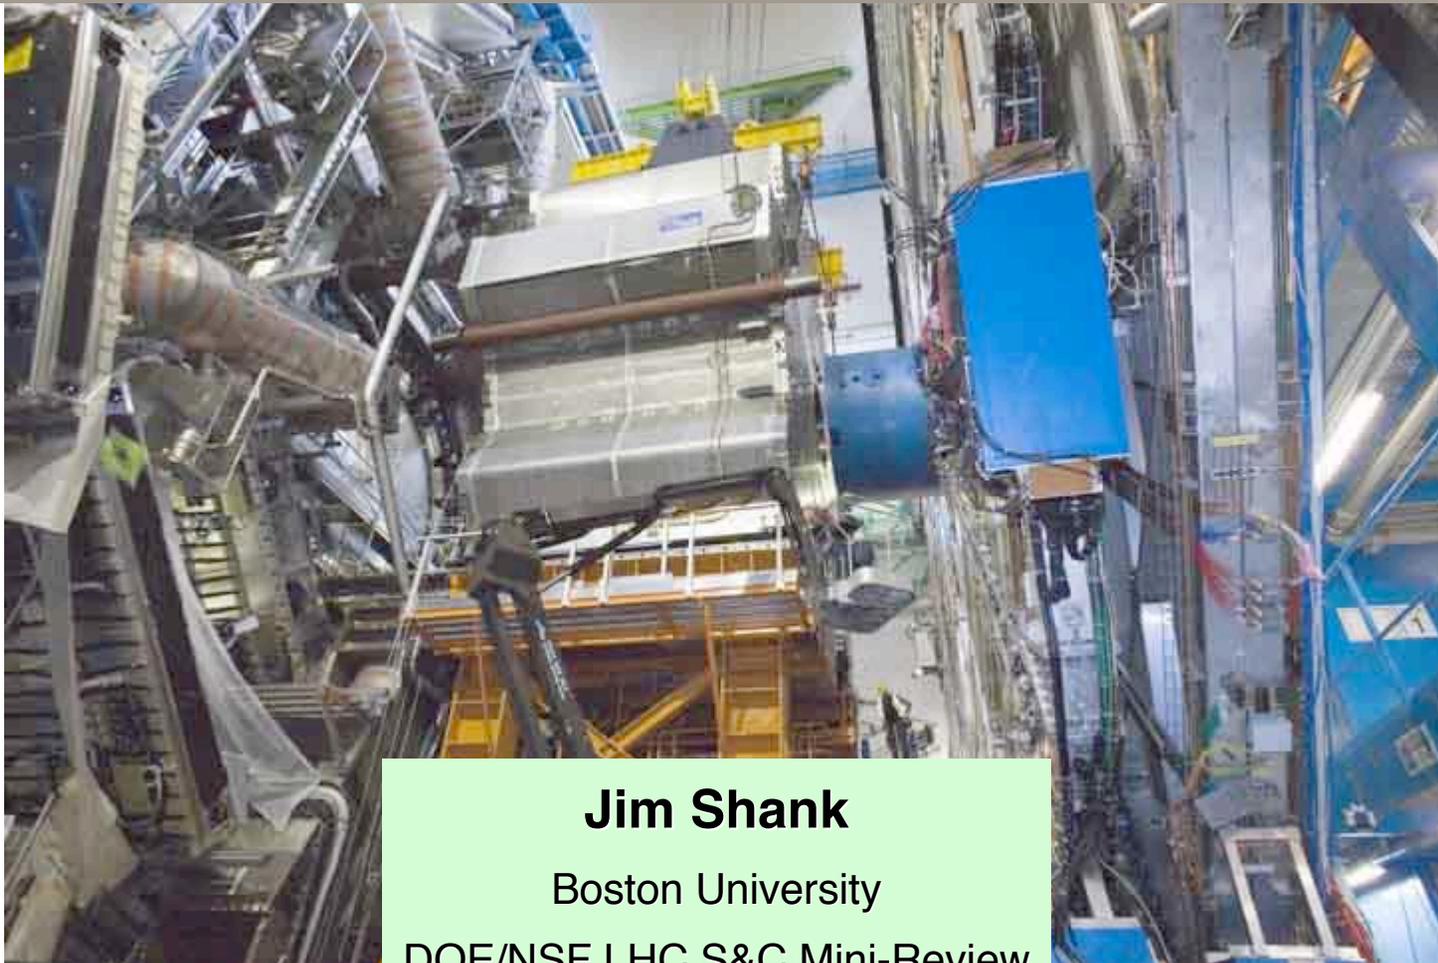


# U.S. ATLAS Computing and Physics Support



**Jim Shank**

Boston University

DOE/NSF LHC S&C Mini-Review

URA

9 Aug., 2007

# Overview



- Selected issues from the questions for this review
  - ATLAS Distributed Data Management, schedules, the ATLAS Computing Model, ...
- Cybersecurity issues
- Tier 3 centers
- Computing System Commissioning/Final Dress Rehearsal
  - Monte Carlo production on going
- Production and Distributed Analysis (PanDA)
- Physics Analysis Support
- Collaboratory Tools
- Current Funding Profile

# Issues from the Pre-Review Questions (1)



- Distributed Data Management (DDM)
  - Review report is available
    - [https://twiki.cern.ch/twiki/pub/Atlas/DDMReviewNovember06/ATLAS\\_DDMReview\\_Feb07\\_report\\_v1.11-1.pdf](https://twiki.cern.ch/twiki/pub/Atlas/DDMReviewNovember06/ATLAS_DDMReview_Feb07_report_v1.11-1.pdf)
  - Follow-up review was held June 27, 2007
    - <https://twiki.cern.ch/twiki/bin/view/Atlas/DDMReviewFollowUpJun07>
  - Major DDM software upgrade was deployed ~2 months ago
    - Did not go smoothly, services still not back to the level before the upgrade
  - DDM remains a big concern

# Issues from the Pre-Review Questions (2)



- Schedule for major tests of ATLAS Computing
  - Combined Cosmics Test
    - Series of milestones (M3-M5) for major cosmics test
      - M3 completed earlier in the summer
      - M4 to start end of August
      - M5 October
    - Each has larger parts of the detector included and tests full chain of data movement/data processing/physics analysis
    - Generating a lot of interest in the detector and physics communities
  - Full Dress Rehearsal
    - Complementary to the cosmic tests and use a realistic mix of simulated physics samples
      - First test Oct 2007
      - Second early 2008
        - Tests will take place for 2 weeks and simulate real running activities complete with full shifts

# Issues from the Pre-Review Questions (2)



- The ATLAS Computing Model
  - Deals with data sizes, CPU/Disk resource needs

| Type & Release | Computing TDR | Baseline 2008 | 1 1 . 0 . X<br>2005/2006 | 1 2 . 0 . X<br>2006/2007 | 1 3 . 0 . X<br>2007/2008 |
|----------------|---------------|---------------|--------------------------|--------------------------|--------------------------|
| RAW            | 1.6/2.0MB     | 1.6/2.0MB     | N/A                      | ~1.6MB                   | N/A                      |
| ESD            | 0.5MB         | 1.0/1.2MB     | 0.8/1.0MB                | 2.0/2.4MB                | 1.0/1.2MB                |
| AOD            | 100kB         | 100/140kB     | 80/160kB                 | 270/350kB                | 120/180kB                |
| TAG            | 1kB           | 1kB           | ~1kB                     | ~1kB                     | N/A                      |

Raw Data/Monte Carlo

- Stepped-up effort to revise the Computing Model estimates in time for the Oct. RRB meetings
  - Latest LHC schedule
  - Measured data sizes with latest software release
- U.S. will look at the Tier2 shortfall in 2010/2011 after this revision

# Facilities



- What plan has U.S. ATLAS put in place with BNL to mitigate effects of cyber security incidents that might effect the availability of the BNL Tier-1 facility?
  - Present Situation if Tier-1 services are not available
    - Primarily affected are PanDA based MC Production and User Analysis
      - Currently running jobs can finish within about a day
      - A new Pilot coming in later can pick up the results
      - No new jobs can be initiated
    - Other Services
      - Content of LXR (indexing service), HyperNews, wiki will be replicated to servers at other US sites
        - Access to existing content will be maintained (Providing static information)
  - Near term Solution
    - Services will be replicated at CERN (in progress)
      - Provides enough redundancy in the system to set up a fail-over
        - New Jobs can be initiated using services at CERN (or elsewhere in the U.S.)
        - Both instances can be used for load sharing

# Facilities (cont'd)



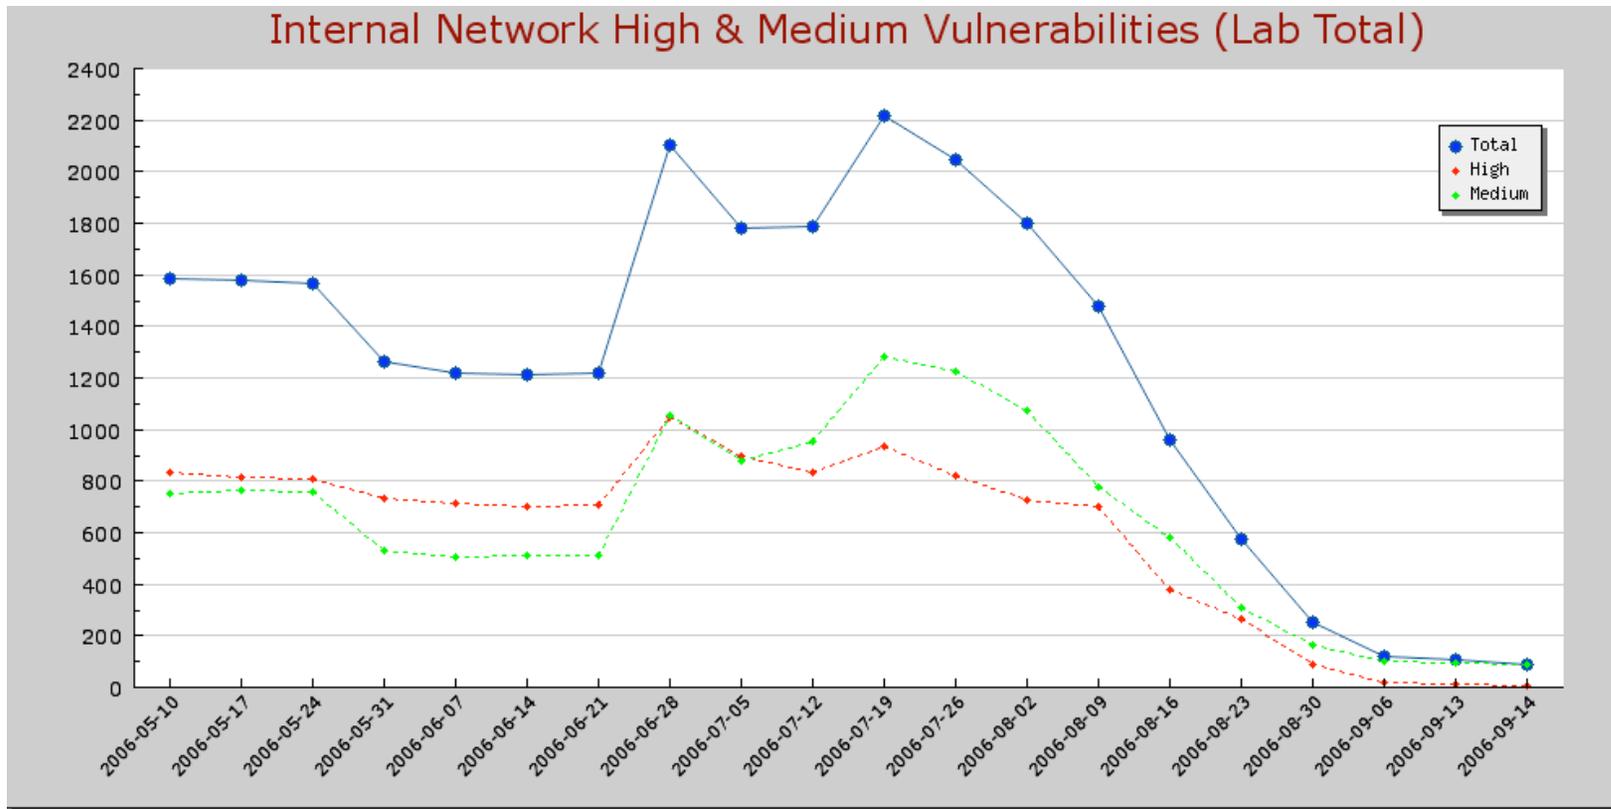
- What plan has U.S. ATLAS put in place with BNL to mitigate effects of cyber security incidents that might effect the availability of the BNL Tier-1 facility?
  - Measures to protect the Tier-1 facility from the affects of a BNL-wide network stand-down
    - Tier-1 center especially protected as an enclave in BNL LAN
      - Firewall between the Facility and the rest of the BNL Campus LAN (besides the Firewall between the public Internet and the Facilities)
      - Facility resources are isolated from Incidents on Campus LAN
    - o Protection is realized by using true Firewalls
      - Traffic between ATLAS and RHIC Computing, and the BNL Campus LAN has to pass the Firewall
    - Decouple actions (i.e. network stand-down) necessary to protect the Campus LAN from the Facility
      - Leave the Tier-1 Center connected while the rest of BNL may have to go off-line
      - Document in preparation, to be signed by Lab Director, will be submitted to DOE for approval
- As a general remark affecting all sites we are concerned as to how Cyber Security is handled as part of DOE's TMR (Technical Management Requirement) implementation process
  - Can easily cause conflict with the scientific mission of the respective Research Program(s)

# Facilities (cont'd)



- What plan has U.S. ATLAS put in place with BNL to mitigate effects of cyber security incidents that might effect the availability of the BNL Tier-1 facility?
  - Significant Cyber Security improvements have been made at BNL – Changes have resulted in
    - Fewer Incidents (graph on next slide)
    - Better Configuration Management
    - Early Detection of potential Vulnerabilities
      - Vigorous network vulnerability scanning program
      - If vulnerabilities are not addressed device is quarantined
    - Improved Detection and Response to potential Intrusions
  - Security Management Act of 2002 (FISMA) requires BNL to perform Certification and Accreditation of their Information Systems
    - Involves Threat and Risk Assessment and Security Plans
    - Science Application International Corporation (SAIC) performed an independent security test and evaluation of BNL's security control in Jan 2007
      - Recommended that DOE Site Office grant BNL an Authority to Operate (ATO)
      - BNL now has an ATO until January 2010
      - BNL was the first DOE Office of Science lab to successfully complete the entire Certification and Accreditation process

# CyberSecurity Vulnerabilities at BNL



# ATLAS MC Production

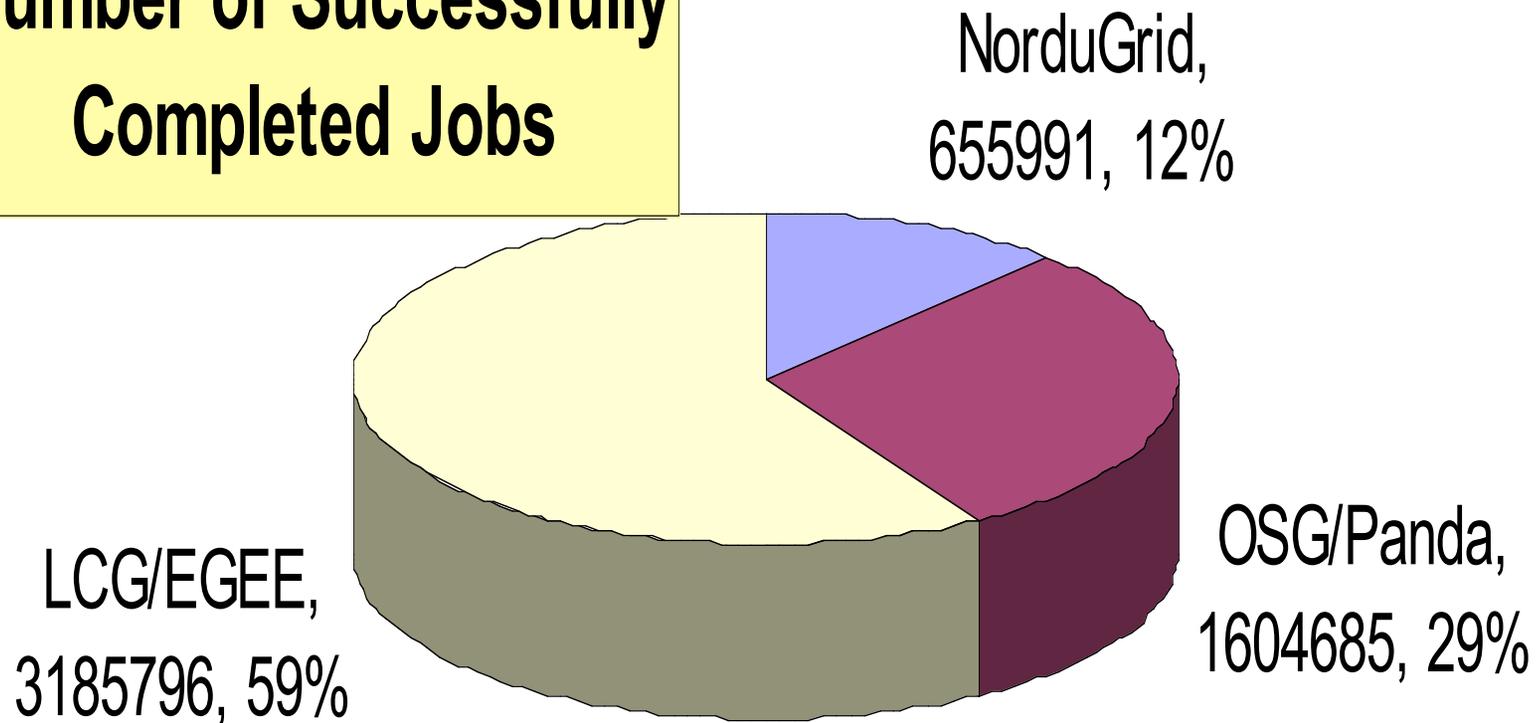


- Computing System Commissioning (CSC) in 2006-2007
- Massive MC samples (approaching  $10^8$ ) produced for software validation, physics studies, calibration and commissioning
- Many hundreds of different physics processes fully simulated with Geant 4 – largest such exercise ever
- Over 8200 different tasks were successfully completed on the grid (each task is a collection of 10-10,000 similar jobs)
- More than 10k CPU's participated in this exercise (average usage about 3-5k CPU's/day in 2007)
- Over 600 TBytes of data (including some replicas) produced and stored

# CSC Production – Job Breakdown



**Number of Successfully Completed Jobs**

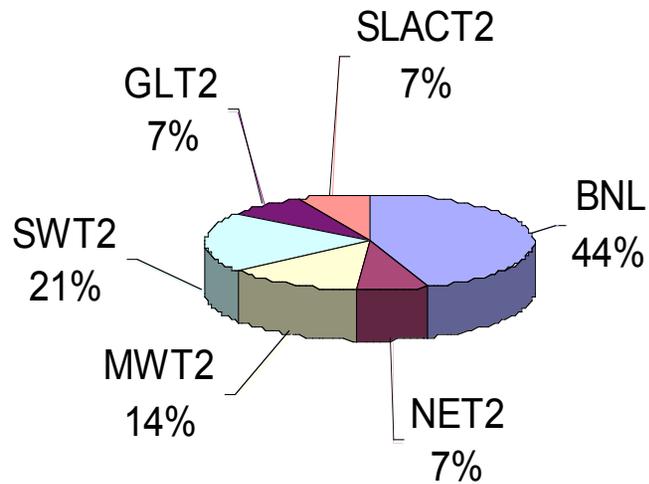


**Total Waltime usage (successful jobs): 3,700 CPU years, since 2006!**

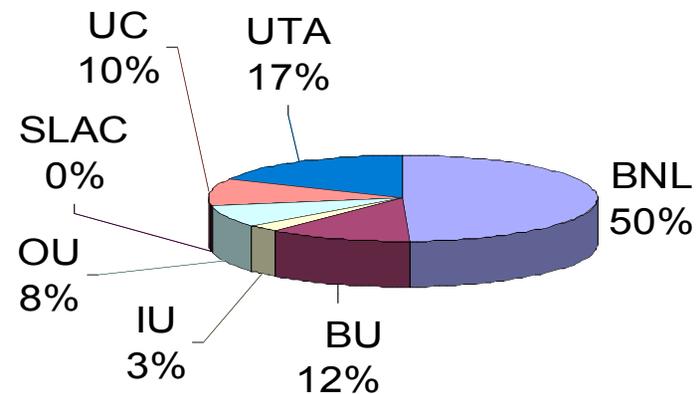
# U.S. ATLAS Production Tier 1 and Tier 2s



### CSC Production - Jobs finished in 2007



### Jobs Finished in 2006



# PanDA Monitor



Panda Production Operations Dashboard

http://gridui02.usatlas.bnl.gov:25880/server/pandamon/query?dash=prod

Getting Started Latest Headlines lotw Google Panda group USATLAS PanDA Monitor US Atlas Reporting C... French to English CERN market: Mess...

Edit in Contribute Post to Blog

Atlas Collaboration Panda Production Operations D... Screen Capture, Screen Shots, ...

**Configuration** Dashboards: [Production](#) [DDM](#) [AutoPilot](#) [Sites & Grids](#) [Analysis](#) [Physics data](#) [Usage & Quotas](#) [Plots](#) [ProdDash](#) [DDMDash](#)

0 min old [Update](#) Not logged in. [List users](#)

**Panda monitor** **Panda Production Operations Dashboard**

[Quick guide](#), [twiki](#) Panda shift [guide](#) [calendar](#) [mailing list](#)

**User info**

**Jobs** - [search](#)  
Recent [running](#),  
[activated](#), [waiting](#),  
[assigned](#), [defined](#),  
[finished](#), [failed](#) jobs  
Select [analysis](#),  
[production](#), [test](#) jobs  
**Quick search**  
Job  
Dataset  
Task  
File

**Summaries**  
Blocks:  days  
Errors:  days  
Nodes:  days  
[Daily usage](#)

**Tasks** - [search](#)  
[Generic Task Req](#)  
[EvGen Task Req](#)  
[CTBsim Task Req](#)  
[Task list](#)  
[Task browser](#)

**Datasets** - [search](#)  
[Dataset browser](#)  
[New datasets](#)  
[Aborted MC datasets](#)  
[Panda subscriptions](#)  
[All subscriptions](#)

**Datasets Distribution**  
[DDM Req](#)  
[Req list](#)  
[AODs](#)  
[RDOs](#)  
[Conditions DS](#)  
[DB Releases](#)

**Servers:** [Panda:OK](#) [Panda-dev:OK](#) [Logger:OK](#) [DQ2:offline](#)

**Tasks assigned to OSG**

**Jobs updated >12 hrs ago:** [activated:451](#) [running:none](#)  
**Jobs updated >36 hrs ago:** [transferring:203](#)

**Space available at sites:**

| Site                            | GB    | As of       |
|---------------------------------|-------|-------------|
| <a href="#">AGLT2</a>           | 1033  | 08-04 09:05 |
| <a href="#">BU_ATLAS_Tier2</a>  | 17181 | 08-04 05:40 |
| <a href="#">BU_ATLAS_Tier2o</a> | 17138 | 08-04 09:07 |
| <a href="#">IU_ATLAS_Tier2</a>  | 51223 | 08-04 09:05 |
| <a href="#">MWT2_IU</a>         | 51223 | 08-04 08:57 |
| <a href="#">MWT2_UC</a>         | 89270 | 08-04 08:07 |
| <a href="#">OU_OCHEP_SWT2</a>   | 153   | 08-04 09:06 |
| <a href="#">SLACXRD</a>         | 6232  | 08-04 09:06 |
| <a href="#">UC_ATLAS_MWT2</a>   | 89274 | 08-04 07:10 |
| <a href="#">UTA-DPCC</a>        | 125   | 08-04 08:51 |
| <a href="#">UTA_SWT2</a>        | 5587  | 08-04 09:02 |

**Pilot job requests per hour, last 3 hours**

| Production Analysis  | Count |
|----------------------|-------|
| AGLT2                | 24    |
| ANALY_BNL_ATLAS_1    | 28    |
| ANALY_LONG_BNL_ATLAS | 5     |
| BNL_ATLAS_1          | 246   |
| BU_ATLAS_Tier2o      | 68    |
| IU_ATLAS_Tier2       | 2     |
| MWT2_IU              | 15    |
| MWT2_UC              | 20    |
| OU_OCHEP_SWT2        | 19    |
| OU_OSCER_ATLAS       | 35    |
| SLACXRD              | 39    |
| UTA-DPCC             | 1     |
| UTA_SWT2             | 31    |

**Production job summary, last 12 hours** (Details: [errors](#), [nodes](#))

| Site            | Nodes | Jobs  | Latest      | defined | assigned | waiting | activated | running | holding | transferring | finished | failed | tot  | trf | other |
|-----------------|-------|-------|-------------|---------|----------|---------|-----------|---------|---------|--------------|----------|--------|------|-----|-------|
| All             | 1535  | 16131 | 08-04 09:07 | 16      | 380      | 2       | 3298      | 3071    | 1461    | 2826         | 4904     | 173    | 3%   | 0%  | 3%    |
| AGLT2           | 33    | 1468  | 08-04 09:07 | 0       | 132      | 0       | 180       | 373     | 2       | 417          | 294      | 70     | 19%  | 0%  | 19%   |
| BNL_ATLAS_1     | 338   | 4799  | 08-04 09:07 | 0       | 0        | 0       | 1189      | 1005    | 31      | 0            | 2549     | 25     | 1%   | 0%  | 1%    |
| BNL_ATLAS_2     | 0     | 0     |             | 0       | 0        | 0       | 0         | 0       | 0       | 0            | 0        | 0      |      |     |       |
| BU_ATLAS_Tier2  | 26    | 280   | 08-04 09:06 | 0       | 0        | 0       | 107       | 61      | 0       | 1            | 108      | 3      | 3%   | 1%  | 2%    |
| BU_ATLAS_Tier2o | 82    | 1562  | 08-04 09:07 | 0       | 0        | 0       | 481       | 324     | 20      | 346          | 388      | 3      | 1%   | 0%  | 1%    |
| IU_ATLAS_Tier2  | 49    | 600   | 08-04 09:06 | 0       | 0        | 0       | 63        | 33      | 431     | 67           | 0        | 6      | 100% | 0%  | 100%  |
| MWT2_IU         | 40    | 1517  | 08-04 09:06 | 0       | 0        | 0       | 120       | 108     | 786     | 495          | 0        | 8      | 100% | 0%  | 100%  |
| MWT2_UC         | 35    | 850   | 08-04 09:07 | 0       | 0        | 0       | 204       | 135     | 136     | 182          | 188      | 5      | 3%   | 0%  | 3%    |

Done

# Tier 3 Centers



- **Summary from whitepaper:**
  - Some local compute resources, beyond Tier-1 and Tier-2, are required to do physics analysis in ATLAS.
  - These resources are termed Tier-3 and could be as small as a modern desktop computer on each physicist's desk, or as large as a Linux farm, perhaps operated as part of a shared facility from an institution's own resources.
  - Resources outside of the U.S. ATLAS Research Program are sometimes available for Tier-3 centers. A small amount of HEP Core Program money can sometimes leverage a large amount of other funding for Tier-3 centers. Decisions on when it is useful to spend Core money in this way will have to be considered on a case by case basis.
  - Support for Tier-3 centers can be accommodated in the U.S. Research Program provided the Tier-3 centers are part of the Open Science Grid and that they provide access to those resources with appropriate priority settings to US ATLAS via the Virtual Organization authentication, authorization and accounting infrastructure.
- Recent T2/T3 meeting: <http://indico.cern.ch/conferenceDisplay.py?confId=15523>
  - On going work to assess resource needs for Physics Analysis
    - See the very interesting talk of Amir Farbin at this meeting
      - Following 4 slides from Amir's talk
- **ATLAS has just formed a T3 task force -- M. Ernst is a US member. S. Gowdy chairs**

# Recap

- Different types of Analysis Activity

- *Algorithmic analysis.*

Tier 1/2

- Re-reconstruction/re-calibration, selection/overlap removal, combinatorics, kinematic-fits, observable calculation.
- Multistep reduction of data into DPD: Skimming, Thinning, Slimming, Reducing

- *Interactive Analysis.* Start with DPD or AOD

Tier 3

- Prototyping above steps
- Making plots, performing studies
- *Statistical Analysis*
- Fits, toy MCs, significance calculation, etc...

# Estimating Analysis Requirements

- Analysis activity is unpredictable... some illustrative examples: (possible stages in one analysis)
  - Step 0: Re-reconstruction/re-calibration
    - Read in 75% of AOD/event, write out 50% (DPD).
    - CPU intensive: 250 ms/event.
  - Step 1: Algorithmic Analysis (selections, overlap removal, combinatorics, ...)
    - Read in 25% of AOD/event, write out 10% (DPD).
    - CPU: 20 ms/event.
  - Step 2: Interactive Analysis
    - Read in <1% of AOD/event. No output.
    - CPU: Effectively 0.
- Earlier steps done infrequently... later steps done very frequently.

# Running Time/Resources

- ATLAS will record 200 Hz of data, regardless of luminosity →  $10^9$  event/year.
- CM Assumption 700 Analyzers: 12 tier 2 CPU/person for analysis at any give time.
- Not unusual for an analysis to start with 50% of the data.
- Assuming perfect software/hardware (10 MB/s read in = ROOT limit).

|        | Laptop<br>1 CPU | Tier 3<br>25 CPUs | Tier 2<br>10 Persons<br>100 CPUs | Tier 2<br>100 Persons<br>1000 CPUs |                                 |          |
|--------|-----------------|-------------------|----------------------------------|------------------------------------|---------------------------------|----------|
| Step 0 | 1 Hour          | 0.0001%           | 0.0035%                          | 0.0140%                            | Working<br>group on<br>Tier 2   |          |
|        | Overnight       | 0.0017%           | 0.0419%                          | 0.1678%                            |                                 |          |
|        | 1 Week          | 0.0235%           | 0.5872%                          | 2.3487%                            |                                 | 23.4874% |
|        | 1 Month         | 0.1007%           | 2.5165%                          | 10.0660%                           |                                 | All      |
| Step 1 | 1 Hour          | 0.0016%           | 0.0400%                          | 0.1600%                            | Analysis<br>group on<br>Tier 2  |          |
|        | Overnight       | 0.0192%           | 0.4800%                          | 1.9200%                            |                                 | 19.2000% |
|        | 1 Week          | 0.2688%           | 6.7200%                          | 26.8800%                           |                                 | All      |
|        | 1 Month         | 1.1520%           | 28.8000%                         | All                                |                                 | All      |
| Step 2 | 1 Hour          | 0.3600%           | 9.0000%                          | 36.0000%                           | Single<br>Analyzer<br>on Tier 3 |          |
|        | Overnight       | 4.3200%           | All                              | All                                |                                 | All      |
|        | 1 Week          | 60.4800%          | All                              | All                                |                                 | All      |
|        | 1 Month         | All               | All                              | All                                |                                 | All      |

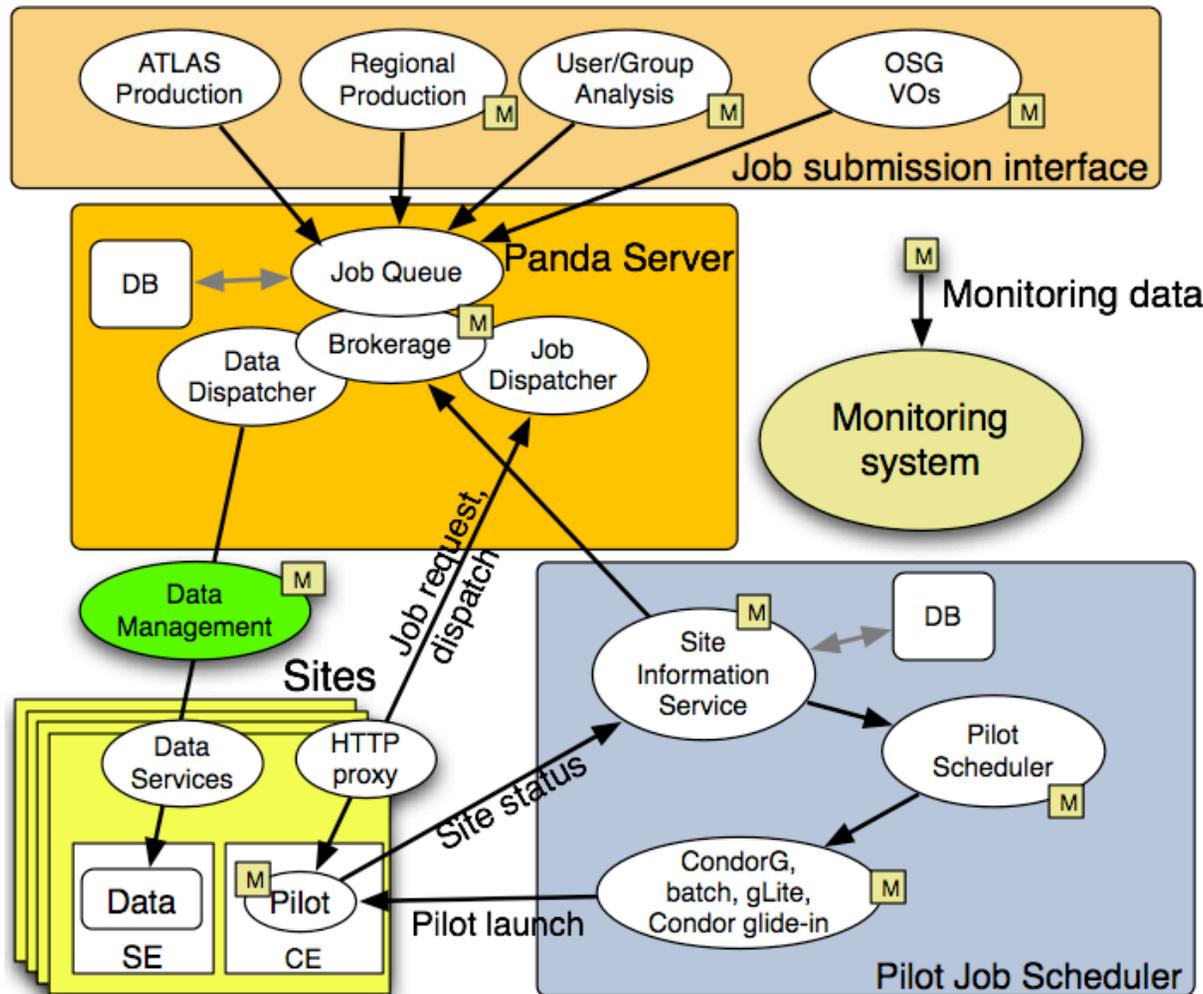
# Tier 3 CPU

- Note that your Tier 3 is the most likely place for your daily interaction with ATLAS data.
  - Every day you will work on your Tier 3... (develop, analyze, etc...)
  - But you will likely use Tier 2 CPU periodically... (run over lots of data)
- Scale of Tier 3 is mostly driven by cost and management capability.
- Tier 2s provide 12 cores/person for analysis at any given time... aggregate cores by working cooperatively (and working asynchronously).
- But Tier 3's are personal.
- Seems “logical” that a Tier 3 provides more CPU per person than at Tier 2.
- In the table I assumed 25 cores per simultaneously active person... less/more means you wait longer/less. This is 3 8-core, \$4000K machines.
- This means over night, you can just barely make plots (step 2) on 1 year's worth of data. (With PROOF, for example).

# PanDA Basics



Workload management system for Production ANd Distributed Analysis  
 Panda team @ BNL, UT Arlington, U Chicago



- Launched 8/05 to achieve scalable data-driven WMS
  - Prototype 9/05
  - Production 12/05
- OSG program 9/06
  - VO-neutral, Condor++
- Integrated with data mgmt
- Pilot-based 'CPU harvesting'
- Analysis as well as production
- Automation, monitoring, low operations manpower
- *Insulate* users (end- and VO-) from grid complexity, problems
  - Lower entry threshold
- Cautious in its dependencies
  - Proven components



Open Science Grid

# PanDA/pAthena Users



## Users

Users: 224  
Job count: 525842

| User                                   | Jobs  | Latest           | Sites used   | Job types run                               | Groups                        |
|--|-------|------------------|--|---|-------------------------------|
| <a href="#">akira shibata</a>          | 71210 | 2007-07-31 21:56 | ANALY_BNL_ATLAS_1 (1732) ANALY_CERN (8) ANALY_LONG_BNL_ATLAS (67796) BNL_ATLAS_1 (1674)  | panda (1762) user (69448)                   | all atlas                     |
| <a href="#">Peter Steinberg</a>        | 53366 | 2007-06-25 16:06 | ANALY_BNL_ATLAS_1 (45520) ANALY_LONG_BNL_ATLAS (7846)  | panda (82) user (53284)                     | all atlas<br>usatlas          |
| <a href="#">Eric Lancon</a>            | 44692 | 2007-07-18 08:20 | ANALY_LPC (8) ANALY_LONG_BNL_ATLAS (4) ANALY_LYON (34348) ANALY_BNL_ATLAS_1 (422) ANALY_TOKYO (32) ANALY_LONG_LYON (9822) ANALY_CPPM (8) ANALY_SACLAY (32) BNL_ATLAS_1 (16)  | panda (266) user (44426)                    | all atlas                     |
| <a href="#">Arthur Marques Moraes</a>  | 44296 | 2007-07-13 10:00 | ANALY_BNL_ATLAS_1 (462) ANALY_LONG_BNL_ATLAS (43834)   | panda (62) user (44234)                     | all atlas<br>usatlas          |
| <a href="#">Ana Damjanovic</a>         | 22056 | 2007-08-08 06:05 | TESTCHARMM (42) PROD_SLAC (52) UC_Teraport (24) MWT2_UC (26) AGLT2 (20) UTA-DPCC (16) BU_ATLAS_Tier2 (16) CHARMM (21830) UC_ATLAS_MWT2 (14) OU_OCHEP_SWT2 (16)   | test (22056)                                | all atlas<br>usatlas          |
| <a href="#">Ning Zhou</a>              | 19556 | 2007-04-23 08:11 | ANALY_BNL_ATLAS_1 (16508) ANALY_LONG_BNL_ATLAS (3048)  | panda (28) user (19528)                     | all atlas                     |
| <a href="#">TARRADE Fabien</a>         | 19020 | 2007-08-08 02:22 | ANALY_BNL_ATLAS_1 (16654) ANALY_LONG_BNL_ATLAS (2366)  | user (17916)<br>panda (1104)                | all atlas<br>usatlas          |
| <a href="#">Torre Wenaus</a>           | 17810 | 2007-08-07 11:16 | (14) BU_ATLAS_Tier2o (16) ANALY_BNL_ATLAS_1 (6) UTA_SWT2 (2) ANALY_LONG_LYON (42) ANALY_PIC (30) NULL (2) ANALY_CNAF (40) ANALY_LONG_UBC (38) MWT2_IU (4) PROD_SLAC (4) TPATHENA (3324) ANALY_UK (40) ANALY_TAIWAN (28) ANALY_SHEF (40) ANALY_TOKYO (30) | test (10486)<br>panda (2840)<br>user (4484) | admin all<br>atlas<br>usatlas |
| <a href="#">Sandrine Laplace</a>       | 15850 | 2007-05-30 07:30 | ANALY_BNL_ATLAS_1 (7098) ANALY_LYON (448) ANALY_LONG_LYON (8) ANALY_LONG_BNL_ATLAS (8296)  | panda (366) user (15484)                    | all atlas                     |
| <a href="#">Vikas Bansal</a>           | 14052 | 2007-07-04 16:05 | ANALY_BNL_ATLAS_1 (154) ANALY_LONG_BNL_ATLAS (13898)   | panda (42) user (14010)                     | all atlas<br>usatlas          |
| <a href="#">Theodota-Lagouri</a>       | 12826 | 2007-07-22 10:36 | ANALY_BNL_ATLAS_1 (10060) ANALY_LONG_BNL_ATLAS (2766)  | panda (678) user (12148)                    | all atlas                     |
| <a href="#">Victoria Rojo</a>          | 11746 | 2007-06-29 01:37 | ANALY_BNL_ATLAS_1 (10982) ANALY_LONG_BNL_ATLAS (764)   | panda (22) user (11724)                     | all atlas<br>usatlas          |
| <a href="#">Arbab K. Pal</a>           | 11506 | 2007-06-08 12:58 | UC_Teraport (202) ANALY_LONG_BNL_ATLAS (10300) BNL_ATLAS_2 (202) ANALY_BNL_ATLAS_1 (354) UTA-DPCC (22) UTA_SWT2 (224) UC_ATLAS_MWT2 (202)  | panda (252) user (11254)                    | all atlas<br>usatlas          |
| <a href="#">Kamal Benslama</a>         | 9972  | 2007-06-24 19:40 | ANALY_BNL_ATLAS_1 (1836) ANALY_LONG_BNL_ATLAS (8136)   | panda (16) user (9956)                      | all atlas<br>usatlas          |
| <a href="#">christina potter</a>       | 9392  | 2007-07-01 13:25 | ANALY_BNL_ATLAS_1 (3638) ANALY_LONG_BNL_ATLAS (5754)   | panda (312) user (9080)                     | all atlas                     |
| <a href="#">Kevin Black</a>            | 8282  | 2007-07-23 19:25 | ANALY_BNL_ATLAS_1 (3242) ANALY_LONG_BNL_ATLAS (5040)   | panda (258) user (8024)                     | all atlas<br>usatlas          |
| <a href="#">Jean-Francois Marchand</a> | 7218  | 2007-08-08 08:21 | ANALY_BNL_ATLAS_1 (5114) ANALY_LYON (812) ANALY_LONG_LYON (246) ANALY_LONG_BNL_ATLAS (1046)  | panda (190) user (7028)                     | all atlas                     |

Torre Wenaus, BNL



# PanDA on OSG, EGEE: AutoPilot



- Extensions to support broad OSG, EGEE deployment developed since Sep '06
  - Extends automation/monitoring into the pilot/scheduling subsystem
    - Keeps operational manpower low despite broader deployment
    - Rapid diagnostics of site, submission problems
  - Flexible use of 'tags' to dynamically define logical queue groupings for use by application communities
    - Queue content of the tag changed 'behind the scenes', either automatically (lcg-infosites) or manually (OSG), based on queue health
    - Insulates user from 'grid weather'; hit 'play' and forget
  - Centralized control, monitoring of multiple distributed pilot submit hosts for scalability, redundancy (BNL, Madison, CERN, Lyon)
    - Avoids Condor submission/monitoring scaling limits
  - Enables dynamically adjustable, feedback-driven pilot submit rate
- Operating stably on OSG+EGEE since fall '06; currently 255 gatekeepers, 360 queues, 281 with working pilots
  - OSG: 58 gatekeepers, 69 queues, 49 operational

# PanDA on EGEE/wLCG



- Panda for analysis (pathena) operational and in use in France
  - Lyon Tier 1 in use, several Tier 2s under test (Saclay, Marseilles, LAPP/Annecy, Clermont-Ferrand, Tokyo)
  - Thanks to intensive help, debugging and encouragement from our French colleagues, in particular Eric Lancon
- Panda/pathena deployed and under test at other Tier 1s
  - Italy, Germany, Spain, Holland, UK, Taiwan, Canada
- All using BNL Panda instance
  - No performance/latency issues so far at Panda level, but indications of limits at Condor level
    - Deploying more CondorG pilot submit hosts (CERN, Lyon)
- wLCG management expresses no objection to Panda, pilot jobs
  - Panda pilots do not waste CPU resources
    - They exit immediately if no work is available
    - Pilot rate dynamically adjustable based on workload

# PanDA, Condor Glide-ins, and OSG



- Use of glide-ins in Panda has been in the plan since Oct '05 meeting with Miron Livny et al @ Madison
- Actively pursued since Sep '06 when we gained manpower (a student) to work on it, support shared by ATLAS and OSG
- Initial priority is a new capability for Condor: schedd glide-ins to support site-level **pilot factory** to achieve better scalability, particularly for analysis
  - Moves pilot submission inside site perimeter to avoid GK GRAM bottleneck
  - Working directly with Condor team
  - Development complete, deployment in progress
- Just made a new OSG extensions hire at BNL which provides the manpower to proceed with startd glide-ins also
  - Re-implement Panda pilot using startd as basis of pilot
  - Use Igor Sfiligoi's glideinWMS as basis for startd glidein infrastructure
    - Well documented, code available, extensive security and monitoring features, welcomes collaboration
    - Objective: common glide-in infrastructure with CMS

# Security in PanDA



- Uses GSI based security for the server's LAMP software stack and its client communications (https)
- User ID, tracking, accounting, controls system is internal
  - Panda activity fully logged and accounted
    - Individual user activity (DN) recorded
- Will leverage Condor (startd glidein based pilots) to get glExec functionality (pilot ID = user ID) where needed
- Client<->Server validation, payload validation still to come
  - Expect to draw on CMS/FNAL work
- Data protection is responsibility of DDM system

# Near-term Plans



- In progress: Broaden LCG, OSG deployment for ATLAS analysis & production
  - Based on interest and local data availability
- In progress: Panda based production on opportunistic OSG sites, LCG sites (Canada WestGrid)
- Summer: Deploy schedd glide-in based pilot factory to key ATLAS analysis sites (BNL, UTA, ...)
- Summer: Extend Panda@LCG to ATLAS production, depending on ATLAS decisions/policies
- Summer/Fall: Integrate startd glide-ins as pilots
  - Leveraging CMS (Igor Sfiligoi) startd glidein factory
    - Planning a visit of Igor and Condor expert to BNL, late Aug
  - Selective deployment depending on requirements/performance (eg. glexec (user ID) support, multi-tasking pilot support (Condor VMs))

# PanDA Summary



- Panda performing very well for ATLAS production, analysis
  - Both as ATLAS production system component and as end-to-end system
  - Work on hardening, robustness, automation, monitoring has paid off
- Activity now is focusing on broadening deployment and usage, supporting scale-up, integrating middleware to extend functionality
  - ATLAS production/analysis across OSG and EGEE
  - Expand/improve OSG VO support, having learned from CHARMM
  - Condor extensions/integration in OSG program to support scale-up, extend pilot functionality (and simplify application-level code)
- Ready to provide stable and robust service for ATLAS when datataking starts
  - We're ready to start turning scalability knobs, but no operational need yet
  - Committed to making Panda the vehicle for effective ATLAS analysis throughout the US
- Demonstrated capability to support OSG VOs other than ATLAS
  - Ready, willing, and with the manpower resources to expand this
  - To provide low-threshold, low-maintenance WMS for OSG VOs
  - Will soon offer support for data handling and data-driven workflow which now is DIY

# Physics Analysis Support



- Keith Baker (Yale) has taken over from Stephane Willocq as Chair of the Analysis Support Group (effective March 2007).
  - Proactively engaging the whole U.S. community
    - Contacting EVERY institute to assess analysis support needs/issues
  - Reviving the Analysis Forum groups
    - Implementing recommendations of our Analysis Support review.
    - Making these groups effective means to facilitate full participation in ATLAS physics working groups.
  - Promoting the very successful series of Analysis Jamborees
    - Most recent is still on-going this week at BNL:
      - <http://www.usatlas.bnl.gov/twiki/bin/view/AtlasSoftware/BnlJamboreeAug2007>

# 1st Physics Workshop of the Americas



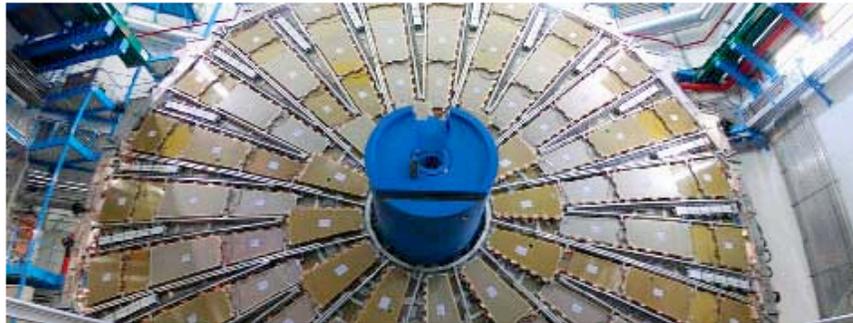
- Replaces the North American Physics Workshop series
- Now includes South Amer., Canada, U.S.

We thank the NSF for helping get SA participation!

FIRST ATLAS PHYSICS WORKSHOP OF THE AMERICAS

SLAC

Home  
Registration  
Payment Options  
Agenda  
Participants  
Poster  
Accommodations  
General Information  
Visa Information  
Travel & Directions  
Contact  
Organizing Committee



## Welcome to the First ATLAS Physics Workshop of the Americas

**August 20 - 23, 2007**  
**Stanford Linear Accelerator Center**  
**Menlo Park, California**  
**Kavli Auditorium**

### REGISTRATION

Registration is required to participate in the workshop.

The registration fee is \$200.

Please register using our online registration form.

[Register >>](#)

### UPDATE

#### BANQUET

Banquet dinner on August 21st for \$50 per person is also available (includes wine and beer).

To pay for the banquet dinner now use our [online payment form](#).

**SLAC** Stanford Linear Accelerator Center, Menlo Park, CA  
Operated by Stanford University for the U.S. Dept. of Energy

[Page Contact](#)  
[- Top -](#)

Last update: July 17, 2007

*Collaborative Tools:*  
Current Activities for ATLAS & the  
LHC



**Steven Goldfarb**  
*ATLAS Overview Week*  
Glasgow, Scotland - 10 Jul 2007

# Remote Collaboration Task Force

## Remote Collaboration Task Force

- *Composition*
  - Chaired by Thomas Baron (IT/UDS)
  - Attended by IT Coordination, Developers, LHC Representatives
  - Occasional Attendance by Collaboration Finance Officers
- *Bi-Monthly Meetings*
  - Update from IT on Activities, Prototypes, Plans
  - Discussion of Priorities for Collaborations
  - Focus on Service-Level Agreements
- *Documentation*
  - <https://cern.ch/twiki/bin/view/RCTF/WebHome>
  - Includes conference room tutorials, guides, some planning, etc.

## ATLAS Participation

- *Attendance*
  - Roger Jones, Steven Goldfarb
- *Reporting*
  - Important Issues reported to CB Chair, CC: Spokesperson, Deputies, Finance
  - Relevant, Interesting News to [hn-atlas-collaborativeTools@cern.ch](mailto:hn-atlas-collaborativeTools@cern.ch)

# Conference Rooms

## General Strategy

- *Same or Similar Equipment When Possible*
  - Bulk Purchasing Power
  - Similar User Interface, Functionality
- *Central Management from IT*
  - Nightly Software Updates, Testing
  - On-Call Maintenance
    - *Operations Up To Users*



## Current Status

- *Prototypes for ATLAS, CMS in 40-R-B10 & 40-4-C01*
  - Tandberg Codex with PC, 4-way MCU
    - *Phone, VRVS, EVO, ECS, HERMES (SIP, H.323)*
  - Sound Optimization
    - *Installing ISDN line for phone level*
  - Simplifying Interface Based on User-Feedback
  - Status: <http://cern.ch/it-multimedia/Rooms>
  - Tutorials: <http://cern.ch/it-multimedia/tutorials>
- *Detailed Plans for Point 1*
  - Large (Expandable) Facility on 1st Floor of Control Room (SCX1)

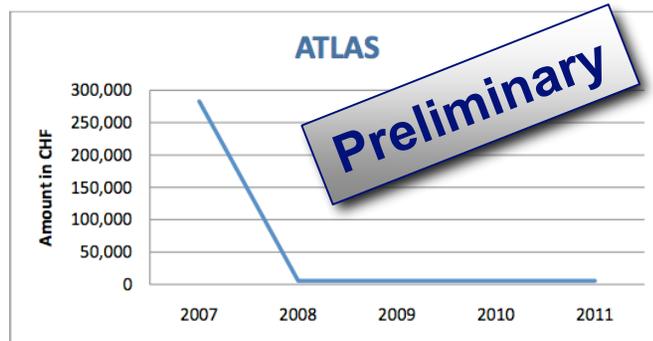
# Conference Rooms

## Next Five Years

- *Service-Level Agreement in Near-Final Form*
  - Defines Services Provided by CERN “Free” or “At Cost”
  - Experiments, Groups Select Services, Provide Funds
    - ATLAS, CMS, AB, DSU
    - ATLAS Provides 50-60 kCHF / Year Over 5 Years
  - IT Provides Installation & Maintenance Manpower (1 New Engineer!)
    - 40-SS-C01, 40-SS-D01, 40-R-C10, 40-R-D10, (40-4-D08)

### 3.1. ATLAS

ATLAS experiment has decided for a profile of maximum investment in the first year. The cost evolution graph in the case all rooms would be refurbished in 2007 is depicted below.



Due to manpower and resource constraints, IT-UDS-AVC can install or upgrade a maximum of 4 rooms per sponsor per year. The profile above will have to be adapted to these limitations. As mentioned in the beginning of this document, the service level agreed to satisfy this profile will be valid for the next 5 years (2007-2011).

*Essentially All Rooms for ATLAS by Next Year*

# Video Conferencing Systems

## EVO Replacing VRVS

<http://evo.caltech.edu>

▪ *Now in production!*

- WARNING: Some claims of problems with java version, some H.323 clients

## ECS (ESNet Collaboration System)

▪ *Recently Upgraded (3 Codian MCU's, new Tandberg Gatekeeper)*

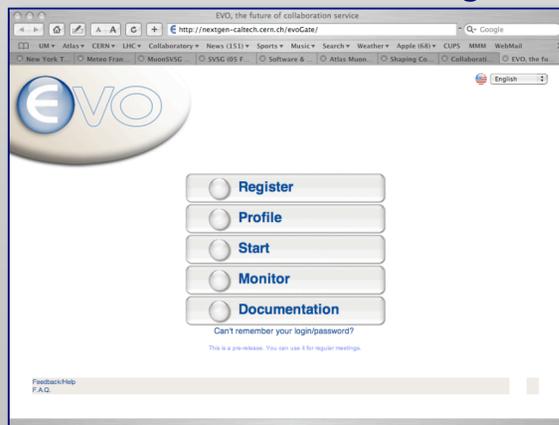
- Change your gatekeeper address: gk1.es.net (198.129.252.168)

## Integration

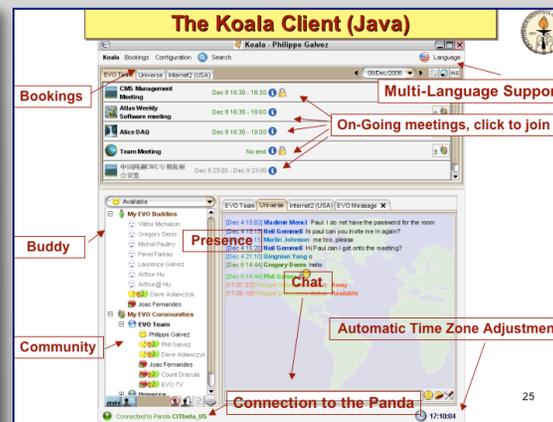
▪ *ECS and EVO allow bridging to phone, as does the hardware*

## Tutorials

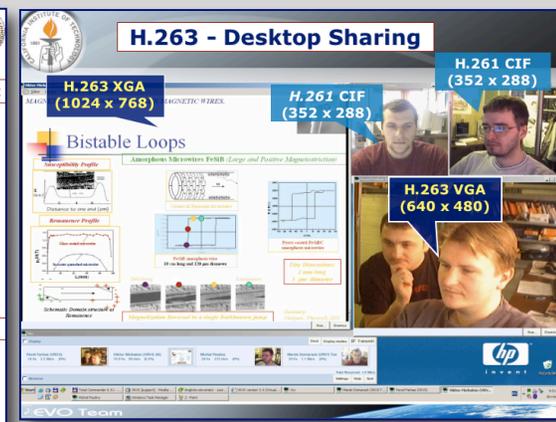
▪ *We will arrange tutorials on rooms and systems in coming months*



9 August 2007



S. Goldfarb ATLAS sw Week, Glasgow



Collab Tools - Slide 33

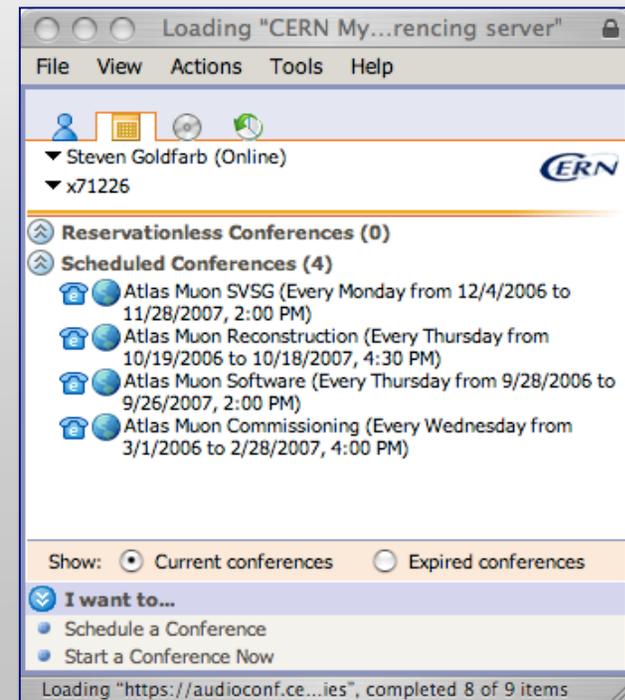
# Audio Conferencing System

Alcatel System at CERN (76000)

<http://edh.cern.ch/Document/AudioConf>

[http://cern.ch/audioconferencing/Audio\\_conference/audioconf-start-guide.html](http://cern.ch/audioconferencing/Audio_conference/audioconf-start-guide.html)

- **Operator-Free 24h / 7d Service**
  - Web-Based Booking System (Requires registration for booking)
  - Call 76000, enter “Leader” or “Participant” code or Click for CERN callback
- **Status**
  - Now in Production
  - Anyone can register (no fee)
- **Resolved Issues**
  - Works with Skype (from outside)
  - Call-back reportedly fixed
- **Call-back fees**
  - Convener’s team account!
  - Looking into ATLAS-wide solution



# Computing Funding



- Overall Computing summary
- SW FY08 Breakdown
- T1 FY08
- FY08 Management Reserve Summary

# Overall Computing Needs



## US ATLAS Computing Needs Profile (AY k\$)

|                                     | FY07         | FY08         | FY09         | FY10         | FY11         |
|-------------------------------------|--------------|--------------|--------------|--------------|--------------|
| Research program target             | 15112        | 15260        | 17406        | 19006        | 17940        |
| <b>Current Computing Total</b>      | <b>15112</b> | <b>15260</b> | <b>17406</b> | <b>19006</b> | <b>17940</b> |
| Difference between Target-Total     | (0)          | 0            | 0            | (0)          | (0)          |
| <b>sw target</b>                    | <b>5268</b>  | <b>5179</b>  | <b>5641</b>  | <b>5835</b>  | <b>6067</b>  |
| sw mr                               | 0            | 624          | 483          | 501          | 523          |
| <b>Total sw</b>                     | <b>5268</b>  | <b>5803</b>  | <b>6124</b>  | <b>6336</b>  | <b>6590</b>  |
| <b>T1 target</b>                    | <b>6295</b>  | <b>6451</b>  | <b>8416</b>  | <b>9803</b>  | <b>8485</b>  |
| T1 mr                               | 0            | 1762         | 1397         | 1831         | 1251         |
| <b>Total T1</b>                     | <b>6295</b>  | <b>8213</b>  | <b>9813</b>  | <b>11634</b> | <b>9736</b>  |
| <b>DC/prod.</b>                     | <b>549</b>   | <b>630</b>   | <b>649</b>   | <b>668</b>   | <b>688</b>   |
| Operations Coordinator (MR)         |              | 250          | 260          | 270          | 281          |
| <b>T2</b>                           | <b>3000</b>  | <b>3000</b>  | <b>2700</b>  | <b>2700</b>  | <b>2700</b>  |
| T2 mr                               |              | 0            | 300          | 300          | 300          |
| <b>Total T2</b>                     | <b>3000</b>  | <b>3000</b>  | <b>3000</b>  | <b>3000</b>  | <b>3000</b>  |
| <b>Total Facilities (with MR)</b>   | <b>9844</b>  | <b>12093</b> | <b>13722</b> | <b>15573</b> | <b>13706</b> |
| <b>Total Fac. (no MR allocated)</b> | <b>9844</b>  | <b>10081</b> | <b>11765</b> | <b>13171</b> | <b>11873</b> |
| Total with no MR allocated          | 15112        | 15260        | 17406        | 19006        | 17940        |
| <b>Total with MR allocated</b>      | <b>15112</b> | <b>17896</b> | <b>19846</b> | <b>21909</b> | <b>20296</b> |

# WBS FTE Software Summary at Level 3



|            |                         | <b>FY07</b> | <b>FY08</b> | <b>FY09</b> | <b>FY10</b> | <b>FY11</b> |
|------------|-------------------------|-------------|-------------|-------------|-------------|-------------|
| 2.2.2      | Core Services           | 8           | 8           | 8           | 8           | 8           |
| 2.2.3      | Data Management         | 7           | 8           | 8           | 8           | 8           |
| 2.2.4      | Distributed Software    | 3.5         | 4           | 4           | 4           | 4           |
| 2.2.5      | Application Support     | 5           | 4.5         | 4.5         | 4.5         | 4.5         |
| 2.2.6      | Infrastructure Support  | 2           | 2.5         | 2.5         | 2.5         | 2.5         |
| 2.2.7      | Analysis Support Center | 2           | 4           | 4           | 4           | 4           |
| <b>2.2</b> | <b>TOTAL</b>            | <b>27.5</b> | <b>31</b>   | <b>31</b>   | <b>31</b>   | <b>31</b>   |

# BNL T1 FY08 estimates



## New FY08 Funding Plan

All in AY k\$

### Tier 1

|                              |      |      |
|------------------------------|------|------|
| Labor                        | 3855 | 2875 |
| Space + Power                | 356  | 356  |
| MST (travel, maintenance...) | 1220 | 1220 |
| Equipment RP \$              | 2327 | 2000 |
| Equipment MR \$              | 1435 | 1762 |
| Total Equipment              | 3762 | 3762 |
| <br>                         |      |      |
| Total Tier 1 RP \$           | 7758 | 6451 |
| Total Tier 1 MR \$           | 1435 | 1762 |
| Total Tier 1                 | 9193 | 8213 |

# FY08 Management Reserve Summary



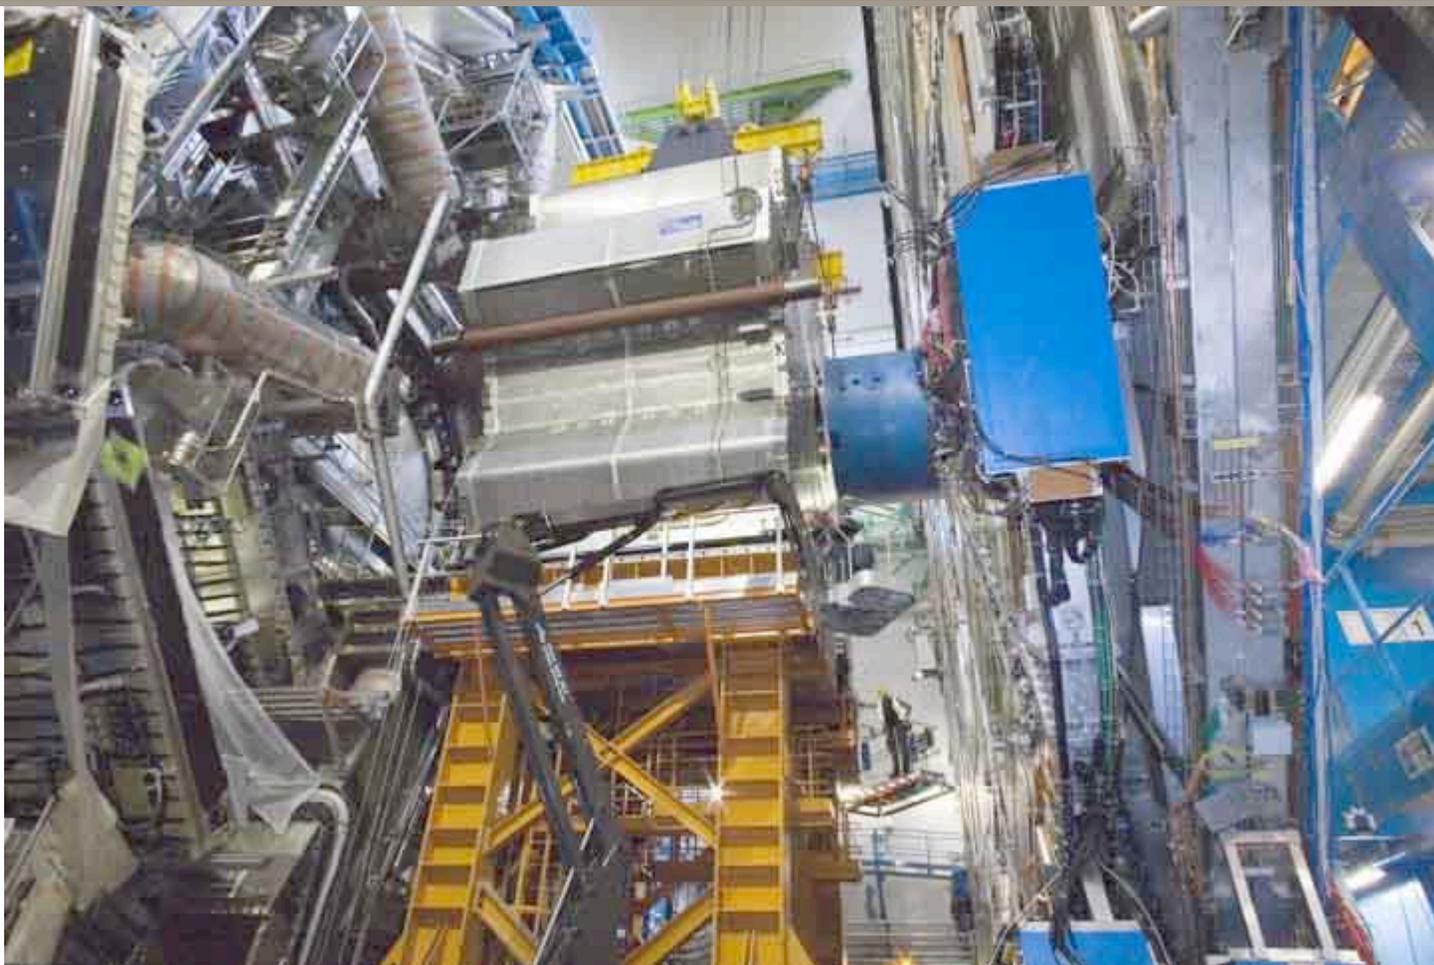
- **SW 624 k**
  - Application Support at BNL
  - Librarian Support at BNL
  - Analysis Support at LBL
  - Analysis Support at ANL
- **T1**
  - 1762 k in Equipment
    - The U.S. specific capacity
- **Facilities**
  - 250k for Operations Coordinator

# Conclusions



- The U.S. ATLAS Software and Facilities have achieved much success in FY07 and we are planning to be ready for the first beam collisions in 2008.
  - All Five Tier 2 Centers in Full Operation
  - Tier 1 under new leadership is doing great
  - The Software Group has a leadership role in ATLAS
- ATLAS test schedule leading to data taking mid-2008 is in place
  - Cosmic running, Full Dress Rehearsal
- ATLAS Distributed Data Management still has troubles
  - US management working actively with ATLAS management on solution
- PanDA very successful
  - Growing in popularity, being deployed at many sites worldwide
- Role and size of Tier 3 centers being defined
- U.S. ATLAS Physics Analysis Support, working now
  - Constantly being re-evaluated to ensure we are effectively getting U.S. physicists “plugged-in” to the overall ATLAS physics effort.

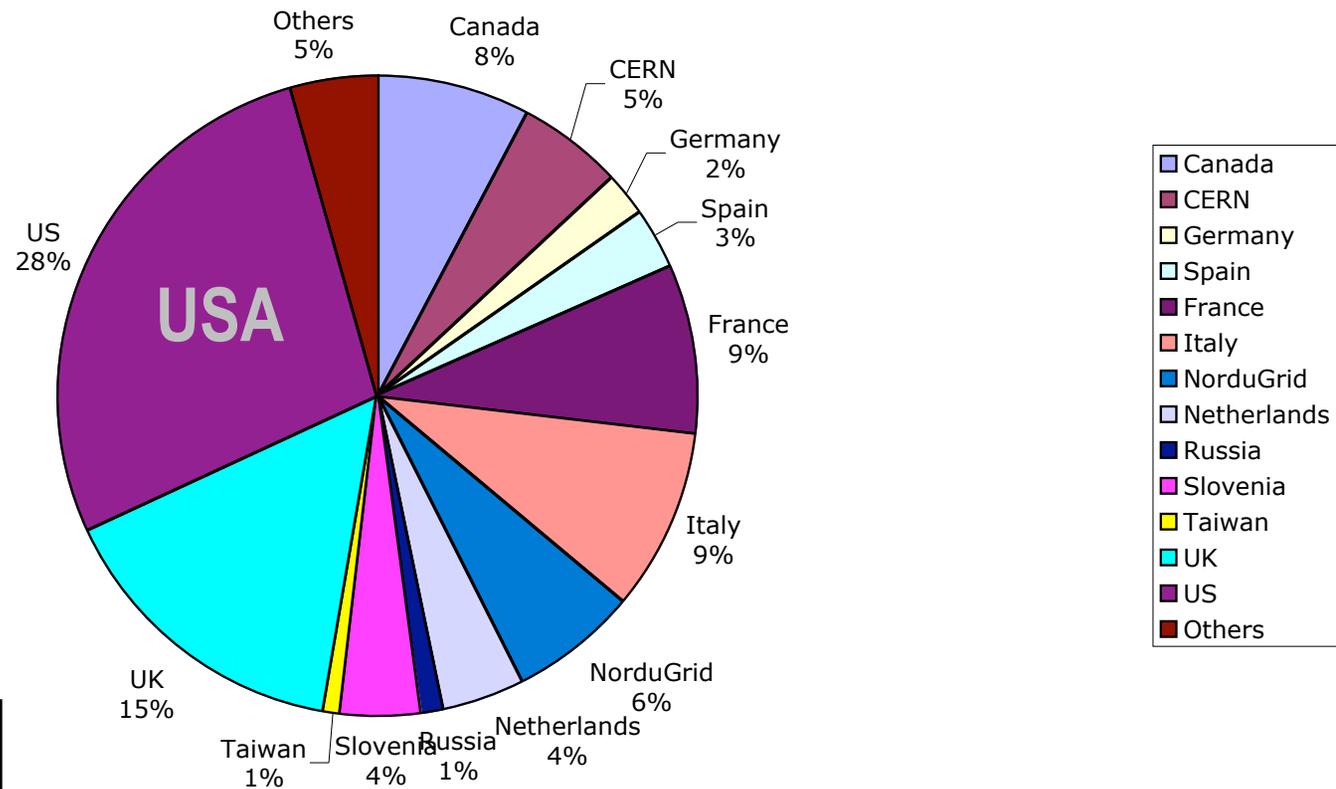
# Backup Slides



# Current CSC Production



## ATLAS Wall Time per country (Jan-Sep 2006)



## L3 WBS for Software & Analysis Support



|       |                         |                       |
|-------|-------------------------|-----------------------|
| 2.2.1 | Coordination            |                       |
| 2.2.2 | Core Services           | P. Calafiura (LBL)    |
| 2.2.3 | Data Management         | D. Malon (ANL)        |
| 2.2.4 | Distributed Software    | T. Wenaus (BNL)       |
| 2.2.5 | Application Software    | F. Luehring (Indiana) |
| 2.2.6 | Software Infrastructure | A. Undrus (BNL)       |
| 2.2.7 | Analysis Support        | K. Baker (Yale)       |

## 2.2.2 Core Services



|              |                        | FY06       | FY07     | FY08     | FY09     | FY10     |
|--------------|------------------------|------------|----------|----------|----------|----------|
| <b>2.2.2</b> | <b>Core Services</b>   | <b>6.5</b> | <b>8</b> | <b>8</b> | <b>7</b> | <b>7</b> |
| 2.2.2.1      | Framework              | 2          | 2        | 1.5      | 1.5      | 1.5      |
| 2.2.2.2      | EDM Infrastructure     | 1          | 1        | 1        | 1        | 0.5      |
| 2.2.2.3      | Detector Description   | 0.5        | 1        | 1        | 0.5      | 0.5      |
| 2.2.2.4      | Graphics               | 0.5        | 1        | 1        | 0.5      | 0.5      |
| 2.2.2.5      | Analysis Tools         | 1          | 1.5      | 1.5      | 1.5      | 1        |
| 2.2.2.6      | Grid Integration       | 0.5        | 0.5      | 0.5      | 0.5      | 0.5      |
| 2.2.2.7      | Core Service Usability | 1          | 1        | 1        | 1        | 1        |
| 2.2.2.8      | Framework Upgrades     | 0          | 0        | 0.5      | 0.5      | 1.5      |

•

## 2.2.3 Data Management



|              |                                 | <b>FY06</b> | <b>FY07</b> | <b>FY08</b> | <b>FY09</b> | <b>FY10</b> |
|--------------|---------------------------------|-------------|-------------|-------------|-------------|-------------|
| <b>2.2.3</b> | <b>Data Management</b>          | <b>6.5</b>  | <b>7</b>    | <b>7</b>    | <b>7</b>    | <b>7</b>    |
| 2.2.3.1      | Database services & Servers     | 1           | 1           | 1           | 1           | 1           |
| 2.2.3.2      | Common Data Mgmt Software       | 1.5         | 1.5         | 1.5         | 1.5         | 1.5         |
| 2.2.3.3      | Event Store                     | 2           | 2           | 2           | 1.25        | 1.25        |
| 2.2.3.4      | Non-Event Data Management       | 0           | 0.25        | 0           | 0.25        | 0.25        |
| 2.2.3.5      | Collections, Catalogs, Metadata | 1           | 1           | 1           | 1           | 1           |
| 2.2.3.6      | Distributed Data Management     | 1           | 1           | 1           | 1           | 1           |
| 2.2.3.7      | Data Access Support             | 0           | 0.25        | 0.5         | 1           | 1           |

•

## 2.2.4 Distributed Software (Personnel)



|              |                             | <b>FY06</b> | <b>FY07</b> | <b>FY08</b> | <b>FY09</b> | <b>FY10</b> |
|--------------|-----------------------------|-------------|-------------|-------------|-------------|-------------|
| <b>2.2.4</b> | <b>Distributed Software</b> | <b>3</b>    | <b>3.5</b>  | <b>3.5</b>  | <b>3.5</b>  | <b>3.5</b>  |
| 2.2.4.1      | Distributed Analysis        | 1           | 1           | 1.5         | 1           | 1           |
| 2.2.4.2      | Production System           | 2           | 2           | 1           | 1           | 1           |
| 2.2.4.3      | Production Support          | 0           | 0.5         | 1.5         | 1.5         | 1.5         |

•

## 2.2.5 Application Software



|              |                             | <b>FY06</b> | <b>FY07</b> | <b>FY08</b> | <b>FY09</b> | <b>FY10</b> |
|--------------|-----------------------------|-------------|-------------|-------------|-------------|-------------|
| <b>2.2.5</b> | <b>Application Software</b> | <b>4.25</b> | <b>5</b>    | <b>5.5</b>  | <b>5.5</b>  | <b>5.5</b>  |
| 2.2.5.1      | Generator Support           | 1           | 0.5         | 0           | 0           | 0           |
| 2.2.5.2      | Tracking Infrastructure     | 0.5         | 1           | 1           | 1           | 1           |
| 2.2.5.3      | Calorimeter Infrastructure  | 0.75        | 1           | 1           | 1           | 1           |
| 2.2.5.4      | Muon Infrastructure         | 1           | 1.5         | 1.5         | 1           | 1           |
| 2.2.5.5      | Monitoring Infrastructure   | 1           | 1           | 1           | 1           | 1           |
| 2.2.5.6      | Other Application Support   | 0           | 0           | 1           | 1.5         | 1.5         |

## 2.2.6 Infrastructure Support



|              |                               | FY06     | FY07     | FY08       | FY09       | FY10       |
|--------------|-------------------------------|----------|----------|------------|------------|------------|
| <b>2.2.6</b> | <b>Infrastructure Support</b> | <b>2</b> | <b>2</b> | <b>2.5</b> | <b>2.5</b> | <b>2.5</b> |
| 2.2.6.1      | Quality Assurance/Validation  | 1        | 1        | 1          | 1          | 1          |
| 2.2.6.2      | Librarian                     | 1        | 1        | 1.5        | 1.5        | 1.5        |

•