# HEP computing

DOE Site Visit

September 14, 2010

# Activities & Issues

- Development of a computational cosmology proposal
- Scientific Computing Applications group
- Status and issues for GEANT4
- Status and issues for SPIRES
- Update on scientific computing cost recovery model

# Computational Cosmology Collaboration Proposal

Risa Wechsler

*Sept 14, 2010*

# Computational Cosmology Collaboration Basic Motivation

- A large program of cosmological galaxy surveys is being supported by the DOE over the next decade, which will probe fundamental physics beyond the Standard Model, such as dark energy, dark matter, neutrino masses, and inflation.
- Extracting science from these surveys, as well as from laboratory and accelerator experiments, requires large, state-of-the-art cosmological simulations which make detailed, accurate predictions for structure formation on a wide range of scales.
- The computational challenge is large, and there is a need for a more coordinated national program that integrates hardware, software development, data curation & dissemination, and training initiatives.
- The HEP-supported DOE Labs have the infrastructure and expertise to host a cosmological computing collaboration to meet the above goals.

*Key People: Tom Abel, Risa Wechsler (SLAC); Scott Dodelson, Nick Gnedin (Fermilab); Uros Seljak (LBNL); Anze Slosar (BNL)*

# Computational Cosmology Collaboration Goals

- Carry out and organize computational cosmology research efforts within the participating labs in support of DOE-HEP science goals
- Engage observational and experimental physicists with the theoretical and computational cosmology community
- Develop, maintain, and support simulation analysis tools and two public state-of-the-art adaptive mesh refinement cosmological hydrodynamics codes for the community
- Create a vigorous training program for DOE laboratory staff and the cosmology community at large to carry out cosmological simulations and analyze the resulting data products.

# Computational Cosmology Collaboration Implementation

- Install and **maintain medium-scale computational facilities** enabling
  - development and optimization of cosmological codes for petascale machines
  - applied computations that require fast turnaround
  - interactive data analysis and visualization
  - analysis of the largest data sets produced at the leadership class facilities
- Deliver the **Cosmology Data Grid**
  - a data repository of curated numerical simulations using standardized data formats, derived data products, and observational data.
  - a user facility to share and disseminate simulation and observational data within and outside the collaboration.
- Support **three FTEs**, whose responsibilities will include
  - training and user support for the software developed by the collaboration
  - performance testing and supporting ongoing research to scale codes for leadership class computers
  - supporting the development of the cosmological data grid including curating the data products generated and organized by the collaboration.

# Specifics

- Example calculations needed for HEP science goals:
  - Large volume simulations for mock surveys
  - Calibration of the halo mass function and bias function
  - Calibration of baryonic effects for weak lensing
  - Determination of Ly-alpha flux 2-point statistics
  - Modeling the Cosmological 21cm signal
  - Calibration of the scale for Baryon Acoustic Oscillations
  - Improving Predictions for Dark Matter Indirect and Direct Detection
- Software support and development
  - Support two state-of-the-art hydrodynamical codes (Enzo & ART)
  - Develop and support a suite of simulation analysis tools (e.g. halo finding, ray tracing, merger trees, mock creation)
  - Bring codes to the petascale
- Data curation through the cosmology data grid
  - Currently a lot of duplication and no standardization, hard to make existing simulations public even when desired.
  - Coordinated effort with standard products

# Possible Expenditures for Hardware

- Specifics still under active discussion
  - Tradeoffs between building up at all labs and one large machine.
- Possible Strategy
  - Build up large machine at one site over two years
  - Develop capacity at all labs for local analysis and visualization
  - In third year focus on building IO at all labs to deal with data volumes from simulations produced.

| SLAC | Fermilab | LBL | BNL | Total |
|---|---|---|---|---|
| 1300 C/IO | 300 C | 200 C | 100 C | 1900 |
| 1500 C | 200 C/IO | 200 C/IO | 150 | 2050 |
| 700 C/IO | 600 C/IO | 400 C/IO | 150 | 1850 |
| 1800 C/IO | 150 C | 50 C | 50 | 2050 |
| 1150 C/IO | 400 C | 300 | 150 | 1850 |

# PPA Scientific Computing Applications

Anders Borgland & Tony Johnson

Richard Dubois

*SCA Department*

*Sept 14, 2010*

# SCA in a Nutshell

- Created in 2010 in a merger of most of the Fermi offline team and the Computing Division Scientific Applications group
  - Staff from Fermi, GEANT4, xrootd, ATLAS and BABAR
- Mission
  - Coordinate offline computing in PPA amongst its projects
  - Be the face of offline computing to the Lab and CD
  - Leverage expertise to support multiple PPA and Lab projects
  - Provide advice to PPA Management on computing issues
  - Work to improve the scientific computing environment at the Lab
  - Reach out to collaborators outside SLAC

# Current Activities

- Support of Fermi, EXO & CDMS data/MC handling and some offline infrastructure
- Direct support of BABAR and ATLAS projects with embedded staff fully devoted to these programs
- Computing coordinators for Fermi, SuperCDMS and EXO.
- GEANT4 [discussed separately]
- xrootd cluster file system
  - Heavily used by ATLAS, BABAR and Fermi at SLAC
  - Growing use by LHC outside the Lab
  - Understanding multi-tier storage
- Ongoing direct support of LCLS offline infrastructure
- Working with the Lab to set the path for Scientific Computing

# Upcoming Activities

- Continue support of Fermi ISOC data handling
- Ramp up in SuperCDMS and EXO
  - EXO expecting to take data this year (Nov)!
  - New Oracle server to be deployed for these two projects
- Looking now at how to work with the LSST Camera group
  - Interfaces defined; opportunity for up to 3 FTEs
- Ramp up CTA as the SLAC side of the project firms up
- GEANT4 [discussed separately]
- Expand xrootd to support global clusters, plus ongoing ATLAS & Fermi operations

# Financial and Manpower

- 23 FTEs operated as a Professional Center
  - Extract travel & M&S from 3% Center overhead
- $4M in labor costs
- $50k travel
- $10k M&S

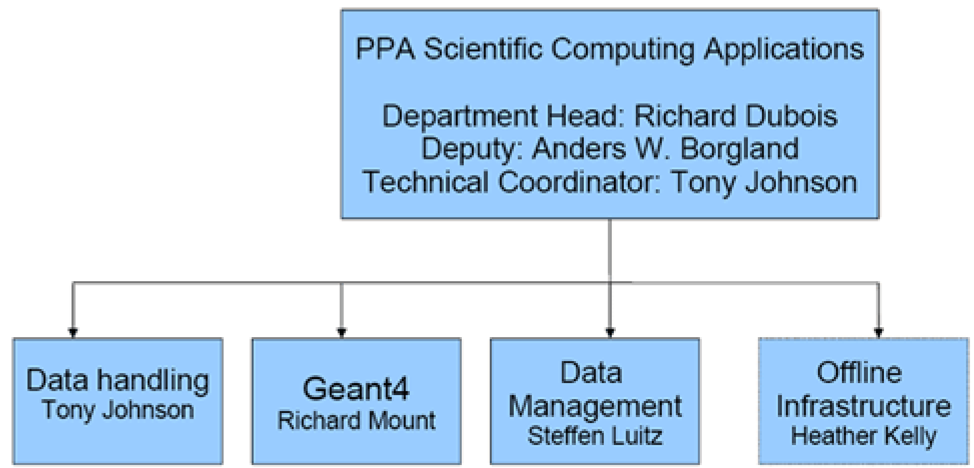


5 (6) 5 7 2

# Issues and Risks

- Data Handling team is understaffed to take on multiple projects
  - Losing key pipeline developer
    - Need to replace him ASAP
  - Need two hires to shore up current effort
    - Budget uncertainty has put one on hold
- GEANT4 funding is insufficient to support the group
- xrootd team needs to cement its relationship to LHC
- Project transitions will be challenging
  - Rates for migration from old to new are often unequal
  - Having no headroom makes it worse

# GEANT4

Richard Mount

# Geant4 at SLAC

- HEP-founded object-oriented toolkit for simulating the interactions of particles with matter:
  - Supports complex geometries
  - Wide range of physics processes for all particle types
- SLAC Geant4 Team
  - Includes the founding architect of Geant4
  - Includes the Spokesperson, leader of Hadronics, leader of Visualization
  - Expertise in kernel architecture, hadronic physics, nuclear physics, and visualization
  - Leads support for US HEP and outreach to US space and medicine
  - Supports use of Geant4 by SLAC projects
    - Using funding from those projects
    - Goal is each team member works 50% on support of SLAC science

# Activities of SLAC Geant4 Group – FY 2010

| Geant4 Core | Other Group Activities |
| --- | --- |
| • Makoto Asai Elected Spokesperson | • ATLAS full simulation realized a 30% speedup (mainly from FY 2009 work) |
| • Geant4 Architectural Review | • Cleanup of ATLAS muon simulation – volume clashes etc. (mainly FY 2009) |
| • Bertini Cascade rewritten for physics and speed improvements | • ATLAS cavern background simulation (Geant4 geometry + Fluka physics) |
| • Development and maintenance of other hadronic and nuclear models | • CDMS simulation – especially germanium detector and phonons! |
| • Event biasing and scoring | • Ongoing support for BaBar and FGST simulation code |
| • Users' Workshops and outreach to new science areas | • Support for EXO, ILC and accelerator studies |
| • Extending/improving visualization | • NIH-funded proton therapy simulation (TOPAS) |
| | • Other externally funded space/medical projects (NASA(Vanderbilt), Varian …) |

# Activities of SLAC Geant4 Group – FY 2011

| Geant4 Core | Other Group Activities |
|---|---|
| • Spokesperson and other leadership roles | • ATLAS – limited consultancy |
| • Develop plans for vital architectural improvements | • ATLAS cavern background simulation – maturing/decreasing effort |
| • Address (mainly hadronic) issues revealed by increasing LHC data | • CDMS simulation – especially germanium detector and phonons! |
| • Upgrading Bertini and radioactive decay | • Ongoing support for BaBar, FGST and EXO simulation code |
| • High-precision neutron code maintenance and improvement | • NIH-funded proton therapy simulation (TOPAS) |
| • Event biasing and scoring | |
| • Users' Workshops and outreach to new science areas | |

# Geant4 - Issues

- Emerging needs of the LHC are the principal driver for HEP funding – they include:
  - Improving the precision and speed of hadronic shower modeling
  - Improving the precision and speed of electromagnetic modeling
  - Improving the precision, robustness and speed of transportation
  - Systematic validation, with input from existing sources plus LHC data, leading to usable estimates of precision
  - Efficient exploitation of multicore (and perhaps later GPU) hardware
  - Ensuring maintainability for 20 years
- Architectural revisions will be needed to address the multicore and maintainability issues
- Estimating the likely return on effort investment for each area of need is difficult – plans and efforts should be reviewed annually

# Geant4 – Specific SLAC Concerns

- The funding needed exceeds that planned by SLAC by ~$100k
  - i.e. by approximately the additional cost of Makoto Asai's spokesperson role
- The SLAC plan excludes urgently needed work on electromagnetic physics validation and improvement
  - Essential to the US HEP program.

Travis Brooks

Scientific Information Services

*Sept 13, 2010*

# SPIRES/INSPIRE Goals

- SPIRES:
  - Continue to provide quality service to HEP community during software migration – without compromising development effort
- INSPIRE:
  - Reproduce SPIRES' functionality on a new platform
    - Faster
    - More Agile
  - Revitalize the HEP information infrastructure:
    - Getting authors right
    - Beyond papers
    - Google-like searching in the HEP research world
  - Build new partnerships

# SPIRES/INSPIRE FY10

- http://inspirebeta.net
  - live April 2010
  - Speed
  - Speed
  - Speed
- Oct 2010 inspirebeta.net will become “recommended site”
  - Also will have – beyond SPIRES:
    - Plots extracted from arXiv
    - Fulltext search of 3 years of arXiv
    - Authors disambiguated
    - LHC notes

# SPIRES/INSPIRE FY11

- Turn off SPIRES
  - Finish back-end curation tools
  - Leverage partners (old and new) in workflow
- Improved Interactions
  - User Feedback Made easy
  - Citation Analysis
- New Features
  - Claim Your Papers
  - Fulltext search in journals
  - Community Knowledge

# SPIRES/INSPIRE FY11

- Partnership Building
  - Solidify Central Partnership
    - CERN DESY Fermilab SLAC
    - Conversations with KEK
  - Build infrastructure, leveraging others in the community
    - arXiv.org and NASA-ADS exchange of software and data
    - Data preservation with ICFA DPHEP
    - ORCID and Publishers
    - PDG: Beyond Interlinking

# SPIRES/INSPIRE Financials

- US effort: 8FTE (6 SLAC / 2 Fermilab)
    - SLAC has management (1) + development (2) + admin/curation(3)
    - Fermilab has admin/curation(2)
  - Labor is only significant cost
    - Travel due to global collaboration
- Proposal to DOE to fund SLAC effort
  - Expand and rebalance admin/curation between Fermilab and SLAC
  - Keep Management/Development at SLAC where expertise is

# SPIRES/INSPIRE Risks

- Software rot
  - Lack of infrastructure investment for last ~15 years
  - Fixed the symptom with INSPIRE migration
    - Prevent recurrence as well
- Global balance
  - Maintain U.S. expertise and equal partnership with E.U. labs.
    - 4 out of 6 INSPIRE "directors" in E.U.
- Funding model
  - Community needs consistent service
  - Project funding doesn't fit service model

# Scientific Computing Cost Recovery Model

# FY2010 operations for scientific computing

- Completed exercise of bottoms up re-evaluation of services and costs
  - Catalog of services and required labor used to build complete model
  - Validated against LBL and BNL operations
- Need ~22 FTEs minimum to operate the current installation
  - 7 FTEs ($1.1M) charged to Scientific Computing in FY10 more appropriately assigned to indirects
  - ~4 FTEs to maintain shared services (including Head of SciComp)
  - 11 FTEs to operate current size of facility (2200 compute servers; 230 fileservers)
    - ~5 needed minimally to maintain expertise, 6 set by scale of facility
  - M&S ($840k) for licenses, maintenance
- Planned model
  - Only 5.5 FTEs shared in cost recovery, remainder built into indirects

# Examples of Costs

| | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Lustre | LCLS/LCLS | 1 | Cluster | | 1 | | 12 | | | 1 | 12 | 24 | 0 | 0 | 0 |
| | LCLS/LCLS | 10 | Server | | 0.8 | | 0.4 | | | 0 | 0 | 0 | 0 | 0.5 | 1 |
| | LCLS/LCLS | 100 | Drive | | 5 | | 2 | | | 0.2 | 0.4 | 0.8 | 0 | 0 | 0 |
| | LCLS/LCLS | 10 | Backup (TB/Week) | | | 1 | 1.6 | | | 0 | 0.8 | 1.6 | 0 | 0 | 0 |
| | LCLS/LCLS | 1 | Rack | | 4 | | 0 | | | | | | | | |
| SUM | | | | | | | 16.0 | | | | | | | | |
| xrootd | PPA/FERMI | 1 | Cluster | | 1 | | 4 | | | 1 | 4 | 8 | 0 | 0 | 0 |
| | PPA/FERMI | 10 | Server | | 2.7 | | 1.35 | | | 0 | 0 | 0 | 0.1 | 0.5 | 1 |
| | PPA/FERMI | 100 | Drive | | 11.8 | | 4.704 | | | 0.2 | 0.4 | 0.8 | 0 | 0 | 0 |
| | PPA/FERMI | 10 | HPSS Storage (TB/Week) | | 0.7 | | 0.21 | | | 0 | 0.3 | 0.5 | 0 | 0 | 0 |
| | PPA/FERMI | 1 | Rack | | 4 | | 0 | | | | | | | | |
| SUM | | | | | | | 10.3 | | | | | | | | |
| Fermi (share of general queues) | PPA/FERMI | 1 | Cluster | | 1 | | 4 | | | | | | 0.1 | 4 | 0 |
| | PPA/FERMI | 100 | Server | | 1.72 | | 5.16 | | | | | | 0.3 | 3 | 5 |
| | PPA/FERMI | 1 | Rack | | 5 | | 0 | | | | | | 0 | 0 | 0 |
| SUM | | | | | | | 9.2 | | | | | | | | |
| Pulse - Martinez | PS/Pulse | 1 | Cluster | 1 | | | 0.1 | | | | | | 0.1 | 4 | 0 |
| | PS/Pulse | 100 | Server | 1.4 | | | 0.35 | | | | | | 0.3 | 3 | 5 |
| | PS/Pulse | 1 | Rack | 5 | | | 0 | | | | | | 0 | 0 | 0 |
| SUM | | | | | | | 0.5 | | | | | | | | |

**LCLS Lustre Storage Cluster**
**Silver/Gold level, 500 drives, 8 servers, 10TB backup up per week, 4 racks: 16 hours/week = $137K/year**

**Fermi xrootd Storage Cluster**
**Silver level,1180 drives, 27 servers, 4 racks: 10.3hours/week=$88K/year**
**$3k/ "cluster" & $2000/server**

**Fermi Compute Cluster**
**Silver level, 172 servers: 9.2 hours/week=$80K**
**$34k/cluster &**
**$260/additional server**

Does not include share of M&S and shared services

# Alternate Model Costs per Project/Directorate

| Groups | Total cost To Pay After Subsidy (K$) | Indirect by Lab budget |
|---|---|---|
| ACCLR/Various | 2.61 | 16.37 |
| ACCLR/MCC | 25.67 | 161.00 |
| ACCLR/ARD | 43.33 | 271.82 |
| ACCLR/KLYSTRON | 21.84 | 137.01 |
| SUM ACCLR | 93.45 | 0.00 |
| LCLS/LCLS | 143.73 | 901.59 |
| SUM LCLS | 143.73 | 1487.78 |
| PPA/KIPAC | 137.24 | 75.61 |
| PPA/FERMI | 207.82 | 114.49 |
| PPA/BABAR | 410.97 | 226.41 |
| PPA/ATLAS | 186.10 | 102.53 |
| PPA/University Groups | 4.12 | 2.27 |
| PPA/Various | 108.30 | 59.66 |
| PPA/SuperB | 25.68 | 14.15 |
| SUM PPA | 1080.23 | 595.11 |
| PS/CISC | 65.27 | 60.24 |
| PS/Pulse | 37.08 | 34.22 |
| PS/SIMES | 31.99 | 29.52 |
| SUM PS | 134.33 | 123.98 |
| SSRL/SDC | 43.33 | 192.65 |
| SSRL/Freia | 18.02 | 80.10 |
| SUM SSRL | 61.35 | 272.76 |
| SUM | 1513.09 | 2479.63 |

Note: ARD will be HEP/LCLS shared

# Financials

## FY11 Funding Request

| B&R | Description | FY 2009 Funding | FY 2010 Funding Programmatic | FY 2011 Requested' | FY 2012 Scenario A' Super B |
|---|---|---|---|---|---|
| KA1101021 | Atlas Research | $ 5,321 | $ 3,959 | $ 4,100 | $ 4,133 |
| KA1102054 | Atlas M&O ** | $ 22 | $ 1,193 | $ 909 | $ 909 |
| KA1102054 | Atlas Tier 2 Comp. | $ 600 | $ 600 | $ 600 | $ 1,150 |
| KA1202012 | BaBar Detector Operations | $ 4,157 | $ 4,194 | $ 2,641 | $ 4,042 |
| KA1202012 | Supplement for BaBar due to slower ramp down | | | | |
| KA1202022 | BaBar Equipment | $ 2,326 | $ 500 | $ 250 | $ 250 |
| KA1202012 | BaBar D&D | $ 3,871 | $ 4,085 | $ 4,009 | |
| | **Det Ops** | $ 10,354 | $ 8,779 | $ 6,900 | $ 4,292 |
| KA1202011 | PEP MMS/ D&D | $ 2,956 | $ 2,553 | $ 2,880 | $ 7,300 |
| | **B Factory ST** | **$ 13,310** | **$ 11,332** | **$ 9,780** | **$ 11,592** |
| | | $ 8,028 | | | |
| KA120102 | Electron Res (BaBar , Super B, SiD,..) | $ 8,319 | $ 8,322 | $ 6,321 | $ 4,328 |
| KA1301021 | Non Acc Res. - Cosmic | $ 18,945 | $ 16,011 | $ 18,781 | $ 20,734 |
| KA1301022 | Non Acc Res. - Intensity - EXO | | $ 3,199 | $ 3,519 | $ 2,521 |
| KA1301032 | LSST R&D** | $ 2,824 | $ 3,000 | TBD | TBD |
| KA1301032 | CDMS R&D | $ - | $ 300 | TBD | TBD |
| | CTA R&D | | | TBD | TBD |
| KA140102 | Theory | $ 7,795 | $ 7,795 | $ 7,989 | $ 8,149 |
| KA140105 | HEP Computing (SPIRES and G4) | | $ 1,458 | $ 858 | $ 2,695 |
| KA140105 | Supplement for SPIRES and SciDB | | | | |
| KA150302 | Det R&D | $ 3,194 | $ 3,178 | $ 3,398 | $ 3,784 |
| | ESTB | | $ 1,500 | | |
| KA150102 | Acc. Science* | $ 8,802 | $ 9,030 | $ 9,134 | $ 10,508 |
| KA1502011 | General Acc Development | $ 4,052 | $ 4,052 | $ 3,452 | $ 4,139 |
| KA1502021 | ILC R&D | $ 12,434 | $ 11,766 | $ 10,321 | $ 10,321 |
| KA1102053 | LARP | $ 1,728 | $ 2,281 | $ 1,845 | $ 1,845 |
| KA150102\ KA1202021 | FACET Operations \AIPs | $ 2,000 | $ 4,304 | $ 6,000 | $ 6,150 |
| | ST | $ 89,346 | $ 93,280 | $ 87,007 | $ 92,958 |
| | Computing Recharge *** | | $ 430 | $ 3,300 | $2.8 M Inc . Above |
| | FY 12 Supplement for Multi-program Financial Model | | | | ~$4.3M Supplement |
| | Core Programs | $ 56,428 | $ 57,004 | $ 57,552 | $ 60,991 |
| | ST w/o Comp Re-Charge, LSST, CTA, CDMS and ESTB | $ 86,522 | $ 88,480 | $ 87,007 | $ 92,958 |

# Backup

# Picture Gallery

## Workflow engine

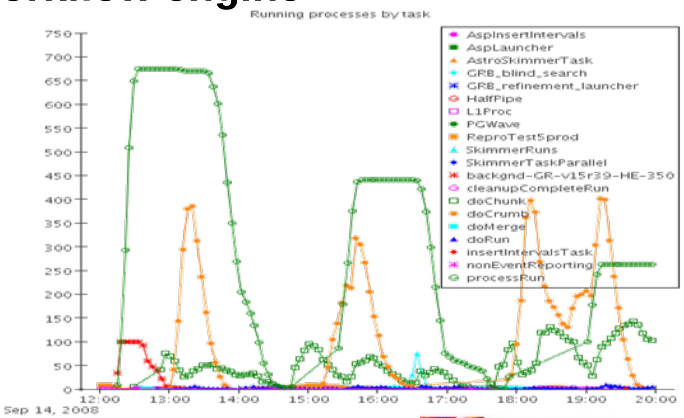


Possible US ATLAS/CMS Production Cluster

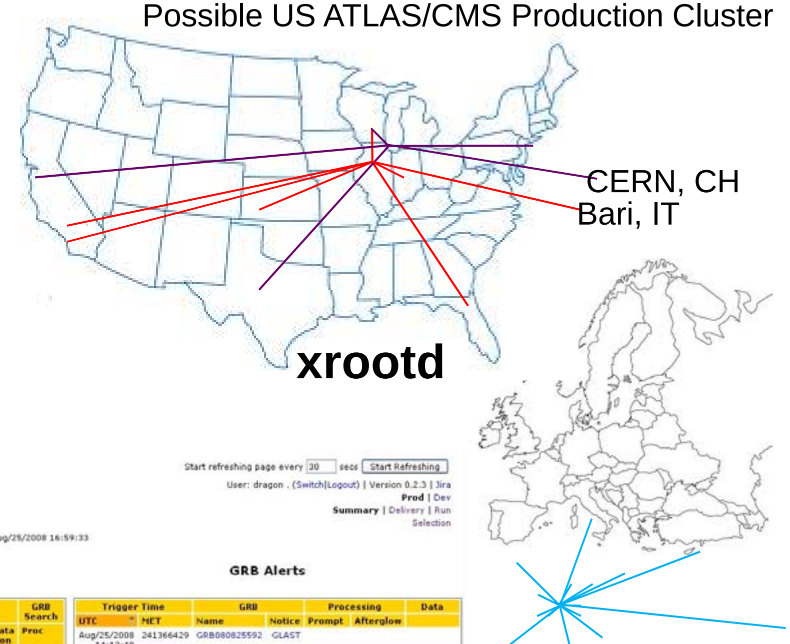


ALICE Production Cluster

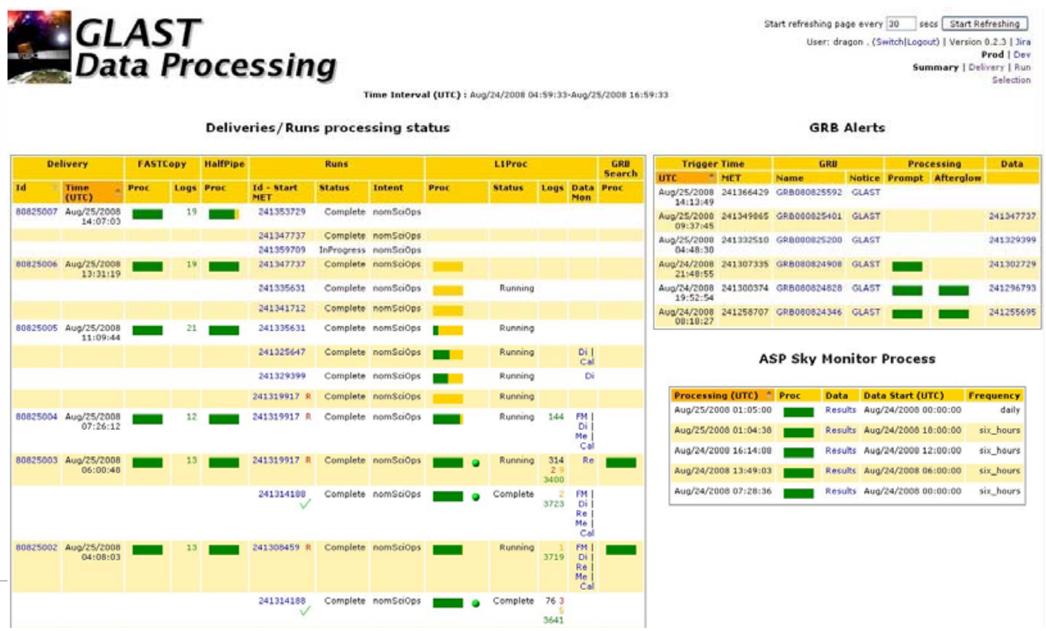


SIMULATION **SuperCDMS/G4**

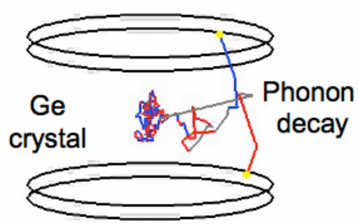


Colors correspond to different phonon polarization states (FT,ST,L)

# SLAC Geant4 Group – Core Activities FY 2011

| Activity | | SLAC efforts | Richard Mount | Makoto Asai | Dennis Wright | Tatsumi Koi | Joseph Perl | Mike Kelsey | Norman Graf | Administrative Support |
|---|---|---|---|---|---|---|---|---|---|---|
| Geant4 Spokesperson | | 40.0% | | 40% | | | | | | Included in Program Support |
| Hadonics | Leadership of G4 hadronics | 10.0% | | | 10.0% | | | | | |
| | RPG/LEP/HEP models | 10.0% | | | 10.0% | | | | | |
| | Bertini model | 47.0% | | | 10.0% | | | 37.0% | | |
| | CHIPS model | 3.0% | | | 3.0% | | | | | |
| | HP Neutron and its alternative models | 30.0% | | | | 30.0% | | | | |
| | QMD model | 7.0% | | | | 7.0% | | | | |
| | HadronicValidation | 4.0% | | | 4.0% | | | | | |
| | New Models and Features | | | | | | | | | |
| Electromagnetics | | | | | | | | | | |
| Kernel | Leadership of G4 kernel | 10.0% | | 10.0% | | | | | | |
| | Event biasing | 20.0% | | 5.0% | | 5.0% | | 10.0% | | |
| | Scoring | 5.0% | | 5.0% | | | | | | |
| Geometry | | | | | | | | | | |
| Performance | Profiling for speed/memory improvement | | | | | | | | | |
| | Code Robustness and QA (reviews) | 10.0% | | 5.0% | 5.0% | | | | | |
| Documentation | | 10.0% | | | 5.0% | 5.0% | | | | |
| Visualization | | 10.0% | | | | | 10.0% | | | |
| Outreach to new science areas | | 31.0% | 2.0% | 5.0% | 3.0% | 3.0% | 15.0% | 3.0% | | |
| Project Management | | 13.0% | 3.0% | 10.0% | | | | | | |
| **Totals per person for FY 2011** | | **260.0%** | **5.0%** | **80.0%** | **50.0%** | **50.0%** | **25.0%** | **50.0%** | **0.0%** | |
| Total Staff Cost (incl benefits and overhead) | | 759 | | | | | | | | |
| M&S (desktops etc.) fully burdened $k | | 8 | | | | | | | | |
| Travel ($50k from 2010 experience + $10k for Spokesperson role) fully burdened | | 91 | | | | | | | | |
| **Total Cost $k** | | **858** | | | | | | | | |