PPA Scientific Computing Applications

Richard Dubois
SCA Department
Dec 10, 2010

Outline

- Plan for Lab-wide scientific computing
- Paying for computing at SLAC
- Scientific computing in PPA
  - SCA – Scientific Computing Applications department
    - Organisation & mission
    - Sub-group activities and goals
    - Issues and Risks

What is Scientific Computing?

- CPU &/or data intensive computations on high performance hardware
  - Compute and storage clusters and databases
- Data acquisition
- Scientific algorithm development & deployment
- Collaboration tools
  - Web, wiki, forum

Sketch of the Lab Scientific Computing Plan

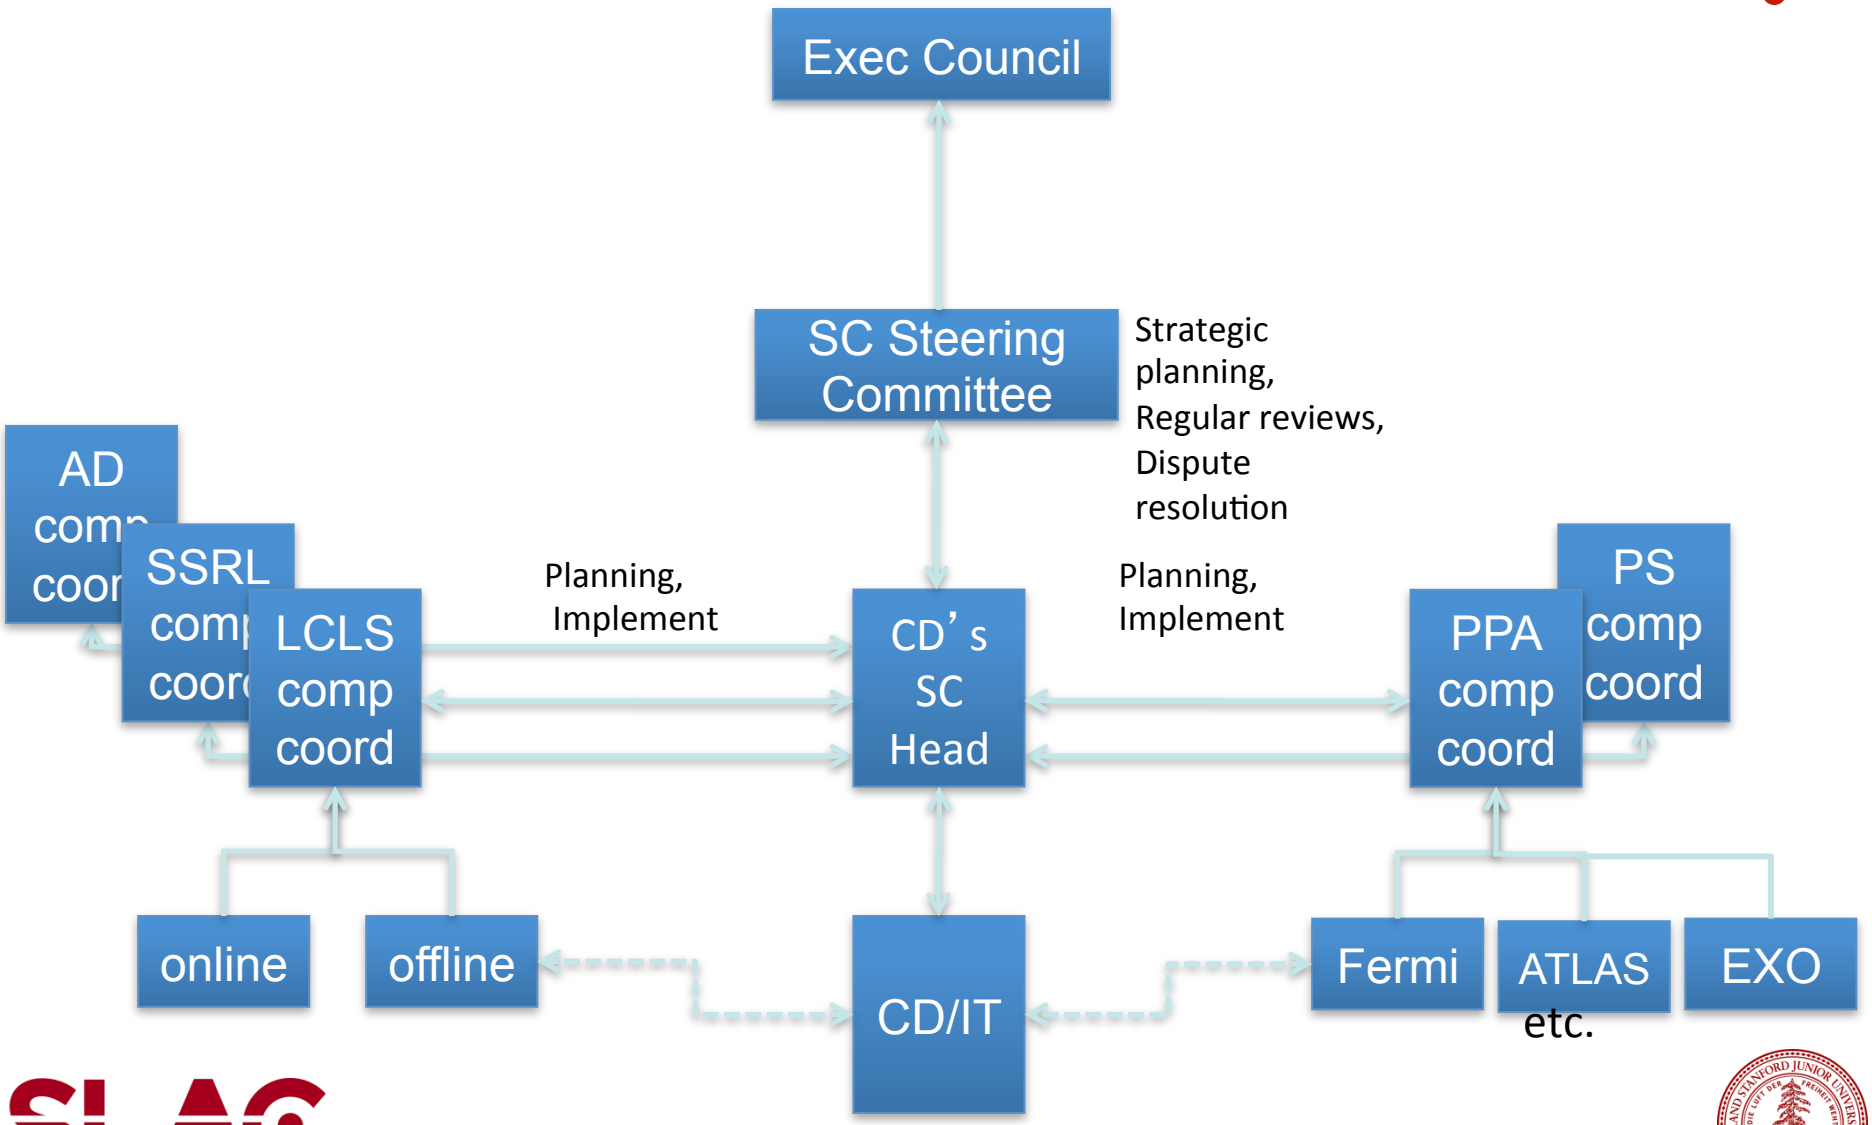

Scientific Computing Plan Labwide

- What is the relationship between organisations
  - CD
    - Head of Scientific Computing works with computing coordinators of Science directorates to plan resource volume and technology
    - Develop overall technology and resource roadmap
    - Hardware acquisition, installation and support performed by CD personnel. Manpower levels result from semi-annual (TBD) roadmaps by Head of SC. Recharge model must smoothly account for manpower needs.
    - Project teams defined in CD to work with Science Directorate client projects
    - Head of SC identifies technologies emerging from the Directorates as suitable for transition to CD
    - Work with Science Directorates to identify the needed services and technology development areas by which CD provides value for money
  - SC Steering Committee
    - Report to ALDs
    - Provide SC strategic advice
    - Review regular reports from CD Head of SC
    - Take ownership of the scientific computing funding model
    - Attempt to resolve disputes

Status

- Persis to commission the Steering Committee in early January
  - Defining initial tasks, communications paths etc
  - Work has started
    - Organise scientific computing workshop in late winter
    - Introduce communications tools
    - Review of CD methods, manpower, licences and contracts
    - Start scientific computing seminar series
    - Leading to technology roadmap and improved community environment
- Search is on for Head of Scientific Computing in CD
  - Offer imminent
  - We will charge ahead in the interim with Randy Melen and Imre Kabai (from CD) filling in and see how things work…

What is needed for a Scientific Computing Professional Center?

- Capacity
  - The ability to support high volume, intensive computing for a variety of independent customers
  - Thousands of cores, many PB disk and robotic tape
- Basic Services
  - Needed by most researchers as basic needs
    - HPC compute cores, storage & databases
- Expertise
  - Deep knowledge of the high performance systems to run them professionally
  - Provide the range of services commonly need by an eclectic collection of customers
  - Advise customers how to craft their resources and tune them for optimal efficiency
- Attractiveness
  - Entice customers to SLAC to use these facilities
  - The Center needs to know more than the customers about compute technologies and it must be reasonably priced.

Paying 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
  - ALDs showed tremendous support for SciComp and willingness to pay for it!
  - Only 5.5 FTEs shared in cost recovery, remainder built into indirects
  - SLAC directly charging customers only 20¢ on the $

Model Costs per Project/Directorate

More indirects:

- shared services
- M&S
- 5.5 FTE – min cluster support
- about $800k “lost revenue to Ops from no OH on indireects

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: LCLS pays for most of AD

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
  - Computing Coordinators from all PPA projects are involved
- Mission
  - Coordinate scientific computing in PPA amongst its projects
  - Be the face of scientific 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

PPA SCA Department Goals

- Critical Outcome (long term):
  - Active and successful PPA and Laboratory wide Scientific Computing program
- Strategic (2-5 yrs):
  - Make SCA the focal point of Scientific Computing in PPA
  - Support of projects external to PPA like LCLS
  - Dedicated funding for Scientific Computing: base effort + project funding
- Tactical (0-2 yrs):
  - Develop departmental identity
  - Understand skills and explore career paths for members
  - Support implementation of the Computing Recharge Model

PPA SCA Data Handling Goals - I

- Critical Outcome:
  - Well supported key software infrastructure components ready for adoption by experiments at SLAC and other universities and laboratories
- Strategic:
  - Invest in new technologies for future experiments like Cloud and Interactive Web Applications
  - Continue to encourage software reuse between experiments at SLAC
  - Build collaborations with the Computing Department, LCLS and other laboratories
- Tactical:
  - Continued support of Fermi ISOC data handling
  - Continued support of SiD
  - Instrumental in smooth ramp up of EXO data processing
  - Ramp up support for SuperCDMS and LSST CCS
  - Define support for CTA and LCLS

PPA SCA Data Handling Goals - II

- Critical Outcome:
  - xrootd to be a major player in HEP data access and distribution
- Strategic:
  - Leverage MoU for increased external collaboration
  - Obtain stable funding and manpower
- Tactical:
  - Ongoing support of ATLAS and BaBar.
  - Continue xrootd support of Fermi and ATLAS
  - Expand xrootd to support global clusters
  - Xrootd storage view simplification
  - Xrootd configuration simplification

Geant4 at SLAC

- Part of international team: 3 Labs with primary roles – SLAC, CERN and FNAL
- 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

PPA SCA GEANT4 Goals

- Critical Outcome:
  - Provide the simulation toolkit to the world and US HEP program that is required in maximizing the precision of physics measurements and sensitivity to new physics.
- Strategic:
  - Leverage collaboration with HEP colleagues and with other disciplines, both national and international, to maximize the value of GEANT4 to US science and commerce.
- Tactical:
  - Achieve and maintain leadership roles in the GEANT4 Collaboration.
  - Ensure adequate and reliable funding for the SLAC GEANT4 group.
  - Lead the effort to plan and execute architectural changes to GEANT4 to ensure efficiency on new computing architectures and longevity on and beyond LHC.
  - Expand the SLAC role into GEANT4 Electromagnetic simulation to ensure that this code meets US needs for precision, reliability and performance
  - Take stewardship of the effort to quantify the validity of GEANT4 by systematic comparison with relevant exisiting and emerging experiemental data.
  - Engage with SLAC's non-HEP science programs in exploring the applicability of GEANT4, especially for accelerator-based science such as the LCLS experiments

Picture Gallery

Workflow engine

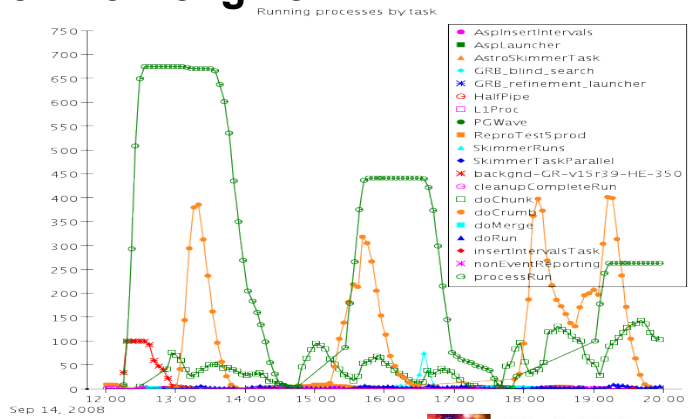

GLAST Data Processing

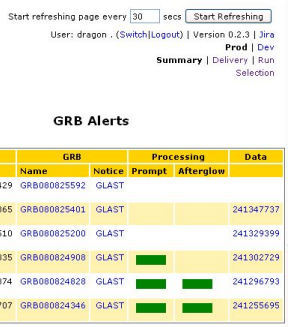

Possible US ATLAS/CMS Production Cluster

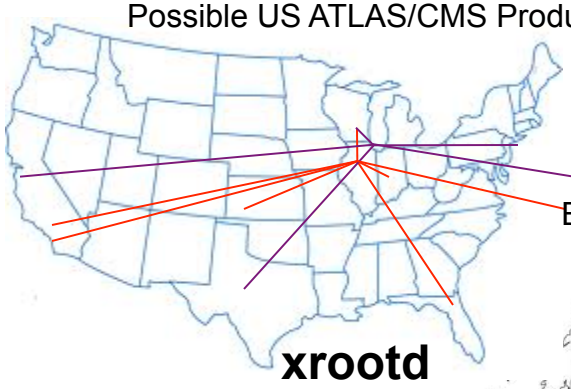

ALICE Production Cluster

SIMULATION SuperCDMS/G4

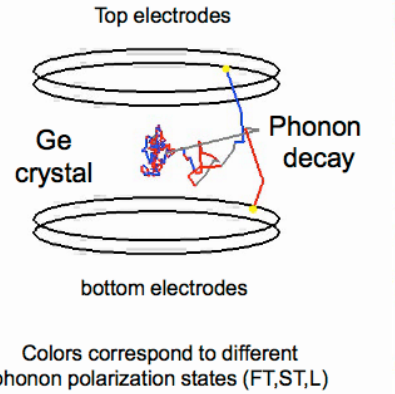

Colors correspond to different phonon polarization states (FT,ST,L)

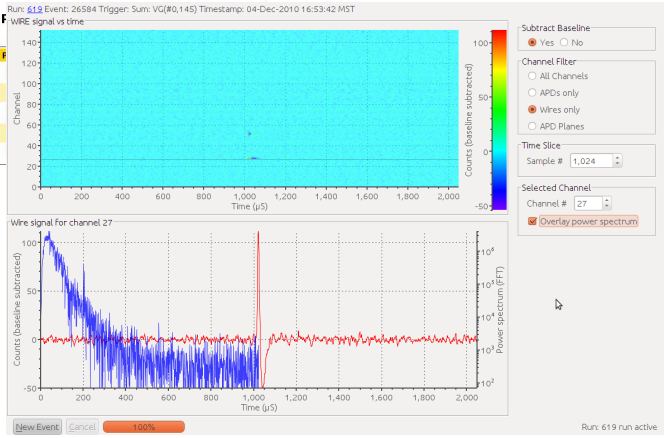

SLAC NATIONAL ACCELERATOR LABORATORY

PPA Advisory Committee: Dec 9-10, 2010

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
  - Spokesman, hadronics and visualisation leads
  - ATLAS cavern sim; CDMS phonons; NIH-funded proton therapy sims …..
- 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 taking data!
  - New Oracle server to be deployed for these two projects
- Started 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
  - Spokesman, hadronics and visualisation leads
  - ATLAS cavern sim; CDMS phonons; NIH-funded proton therapy sims …..
  - team discussions with LCLS experiments
- 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

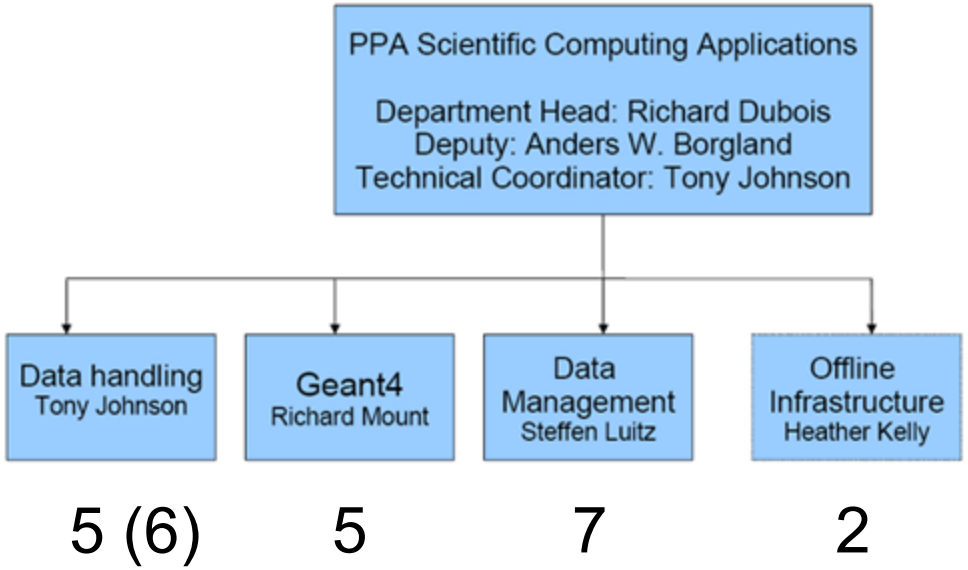

Issues and Risks - I

- Primary concern is funding
  - Ideally the setup would be core funding for base activities plus project funding to take care of specific projects and ebb and flow of effort
- Concept is much like Detector R&D
- OHEP in the past has discussed a computing line
  - There was a briefing on all the Labs' scientific computing in 2009
  - No action yet in creating such a line

Issues and Risks - II

- Data Handling team is understaffed to take on multiple projects
  - Lost key pipeline/Oracle developer to LCLS
    - Has been replaced; new person being trained
  - 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
- Keeping a diverse collection of people coherent

Bonus Slides

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
• Extending/improving visualization	• Support for EXO, ILC and accelerator studies
• Users' Workshops and outreach to new science areas	• 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 originally 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.