OSG Virtual Organizations Status and Outlook

1. ALICE – Jeff Porter
2. ATLAS – Rob Gardner
3. Belle II – Leo Piilonen
4. CDF – Rick Snider
5. CMS – Frank Würthwein
6. DZero – Joel Snow
7. Engage – John McGee
8. Fermilab – Steven Timm
9. HCC – David Swanson
10. SBGrid – Piotr Sliz

ALICE & ALICE-USA

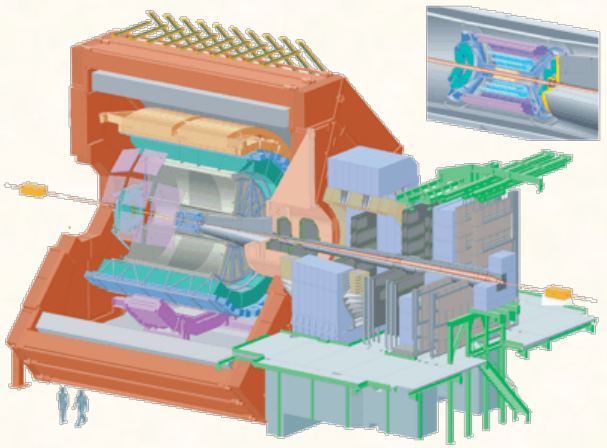

- ALICE Experiment
  - Dedicated heavy-ion experiment at the LHC

    Study physics of strongly interacting matter
    at extremely high energy densities
  - 1000 physicists, 33 countries
- ALICE-USA Collaboration
  - Formed to manage construction & operation of ALICE EMCal
  - ~50 physicists at 11 participating US institutions
- ALICE-USA Computing Project
  - Build a facility to meet ALICE-USA computing obligations to provide share of computing resources for ALICE data analysis and simulations, enabling US ALICE Scientist to pursue research goals
  - Proposal developed in 2008/2009 & project approved in 2010:
    - Locate facility at two sites: LLNL/LC and LBNL/NERSC/PDSF
    - LBNL as host lab, Ron Soltz & Jeff Cunningham (LLNL), Jeff Porter & Keith Beattie (LBNL)

ALICE-USA Facilities: LLNL/LC & NERSC/PDSF

- LLNL/LC Computing: Institutional-HPC supporting LLNL Science & Engineering projects
  - Cost effective procurement & operations model
    - large scalable unit purchases
    - In-house standard software → OS (CHAOS) & batch (SLURM)
  - Pursuing external collaborations (e.g. Green Data Oasis)
  - ALICE Resources
    - 800 cores, 650TB Disk space
    - Fixed resources for 3 year life cycle

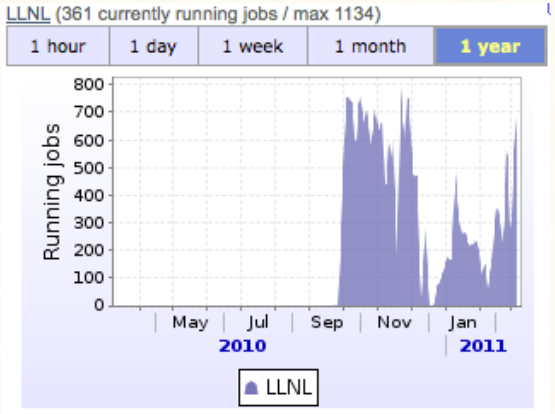

- NERSC: DOE Office of Science Flagship facility for High Performance Scientific Computing
  - Supports DOE sponsored research by competitive allocations
    - Two Cray Systems: Hopper & Franklin: ~200k cores
    - Special Clusters: Visualization, Analytics, HENP (PDSF)
    - Evaluation Systems: GPU and Cloud Clusters
    - PB-scale Global File System (GPFS), Tape Storage (HPSS ~40 PB)
    - Data Transfer, Science Gateways & Grid Services
  - PDSF: HENP funded cluster operated by NERSC
    - STAR Tier 1, ALICE Tier 2
    - ATLAS, IceCube, Daya Bay,
  - ALICE Resources
    - 320 cores, 400TB disk space , PB-scale tape storage allocation in NERSC/HPSS
    - Annual growth plan of 400 cores, 500TB disk

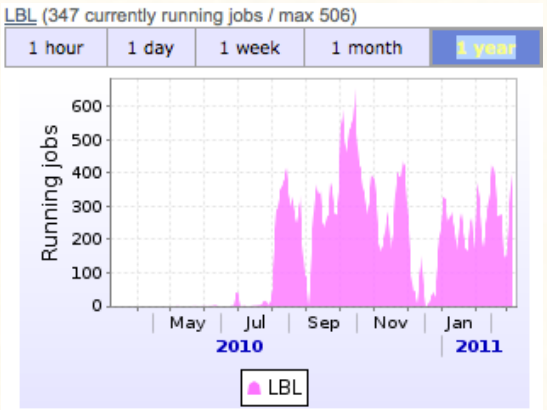

ALICE Grid Computing Model

- ALICE Grid framework in production for since 2004
- AliEn, "Alice Environment"
  - VO Box at each site manages
    - Job submissions (CEs or LRMA directly)
    - Proxy renewal
    - Software deployment (AliRoot, Geant, …)
    - Monitoring
  - Central Task Queue at CERN
  - JobAgents (pilot jobs) submitted to WNs
    - Evaluates local resources
    - pull jobs from the task queue (for 1 user)
- Data Management
  - AliEn FileCatalog
  - Xrootd-based SEs with global redirector
- Grid monitoring with MonALISA
  - Job/site status: ~30k concurrent jobs
  - SE capacities & availability: ~10PB
  - Dynamic network topology mapping

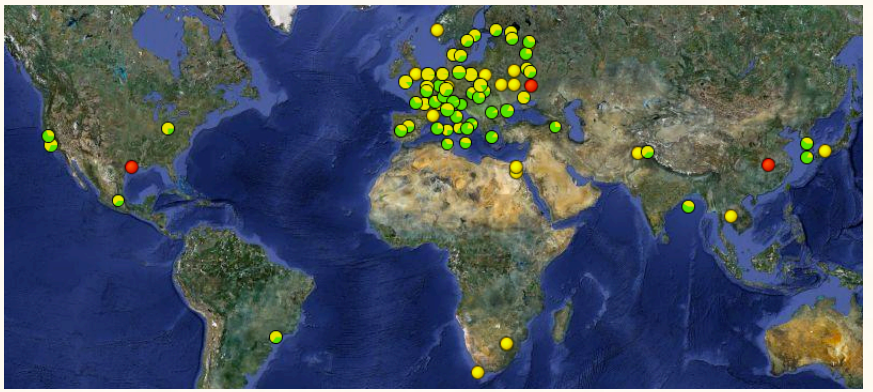

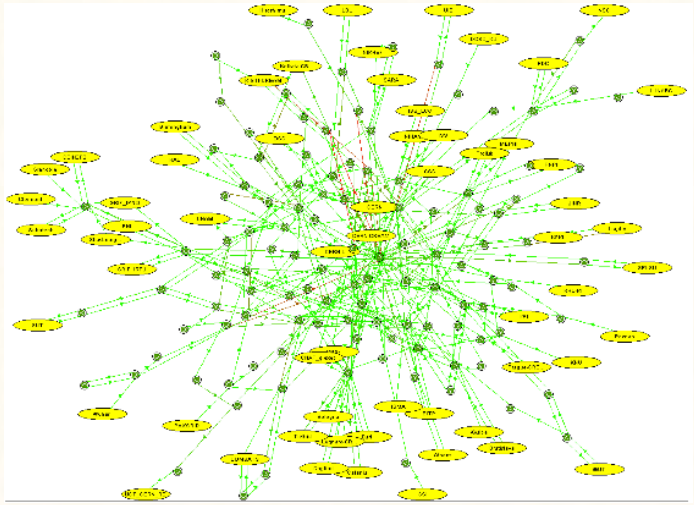

network topology map

ALICE-USA & OSG: 2011-2012

- ALICE-USA Project Plan: leverage OSG capabilities
  - OSG RA participation
  - Resource Accounting Reports to WLCG
    - Very good response from OSG support
      - Thanks to John W., Brian B., Chander, ...
    - WLCG-registration waiting final MOU signatures
  - Disposition of unused resources
    - PDSF is already an OSG resource
    - LLNL currently hidden OSG resource, full support for limited VO set is planned this year

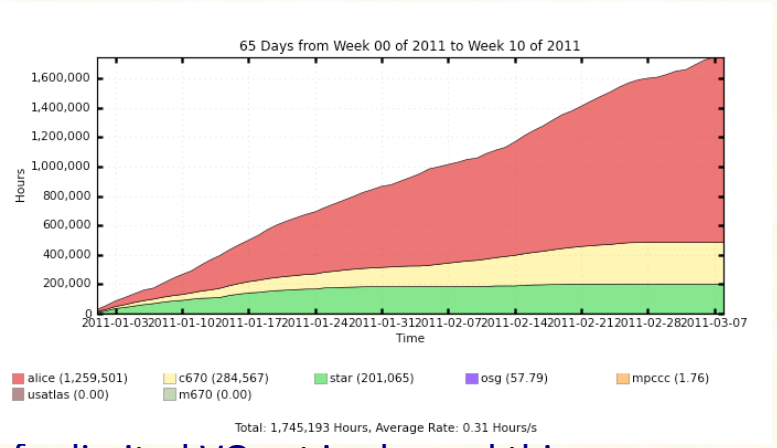

- Opportunistic use of OSG resources
  - ALICE requirements
    - 2-4GB/core, 10GB local scratch, WN outgoing network connectivity
    - RHEL5/SLC5 or newer
  - Need to relax current on-site VO-box requirement
    - Will deploy a test ALICE VO-box at LBNL to submit to remote OSG resources
    - Software distribution into $OSG_APPS via modified AliEn tools

ATLAS Computing on OSG

Rob Gardner
VO Session, OSG All Hands Meeting
Boston March 9, 2011

- 37 countries, 174 institutions, >3000 physicists, >1000 grad students
- US: 39 universities, 4 DOE labs, 21 states involved, > 700 physicists, >200 students

Forefront Particle Physics

Origin of Mass (Higgs)
Supersymmetry
What is dark matter?
Do extra dimensions exist?

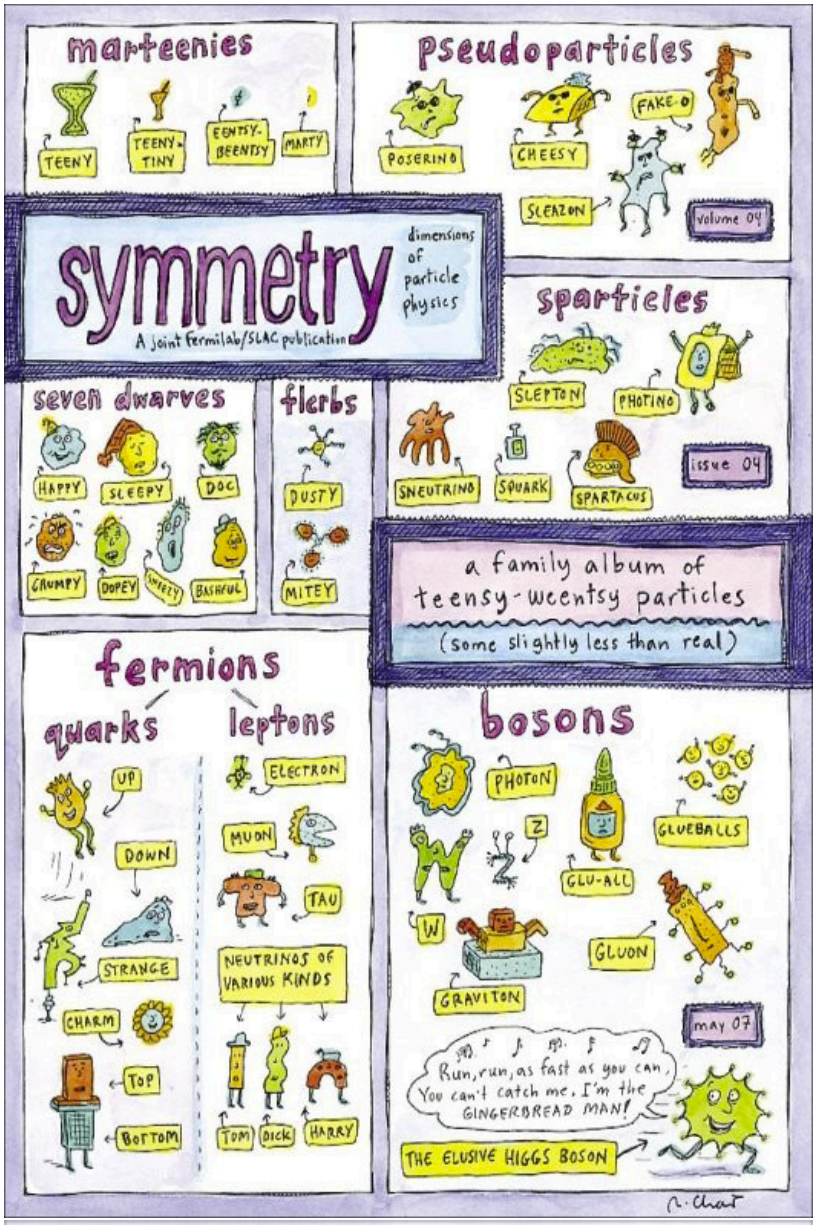

Methodology

- Two modes: central production (simulation, reprocessing, pileup) and user analysis
- Both use the (pilot-based) Panda system which submits to "queues" at sites
- On OSG: Tier 1 at BNL, 5 (federated) Tier 2 centers (9 sites) - (17k slots, 15 PB disk)
- 2010 model: pre-placement of input datasets for analysis

Outlook 2011-2012

- Expect constant production and analysis for the 2011-12 LHC run
- Tier 2 centers - have been operating within the US "cloud" - now going multi-Cloud
- Data pre-placement to (global) caching model
  - ★ BIG change! 27 PB in old model reduced to 9 PB new model)
- ~26 Tier 3 sites will be ramping up
- Collaborating with OSG, CMS towards Xrootd federated data layer across all ATLAS storage facilities

Belle II

Leo Piilonen, Virginia Tech

piilonen@vt.edu

on behalf of the
Belle II Computing Group

http://belle2.kek.jp

OSG All Hands Meeting March 2011, Boston

The Team

- ✓ Computing team led by Thomas Kuhr and Takanori Hara [原隆宣]

Kuhr

Kuhr *Fifield*

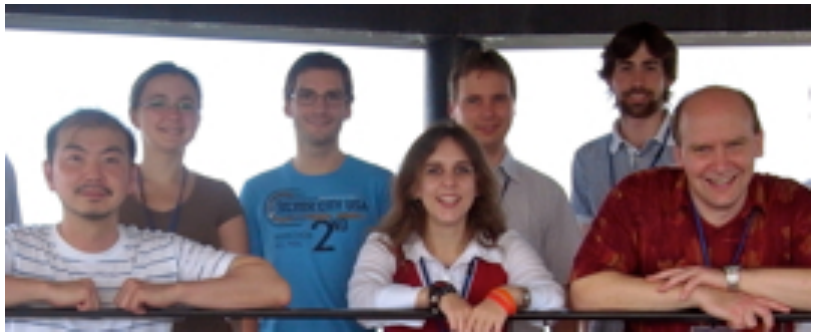

Hara *Sevior*

Hara

- ✓ Distributed Computing and Data Management based in Australia (Tom Fifield, Martin Sevior) and Korea (Kihyeon Cho, JungHyun Kim, Soonwook Hwang, Sunhil Ahn, Taesang Huh)
- ✓ ... with support from China (Yan Liang Han [韩艳良]), Japan (Go Iwai [岩井剛]), Poland (Rafał Grzymkowski, Miłosz Zydbał), Slovenia (Marko Bračko) and all the people on the next slide
- ✓ Thanks to Ruth Pordes (OSG VO creation), Brian Bockelman and Burt Holzman (EGI interoperability, getting things running), and Douglas Olsen (OSG security).
- ✓ Thanks to the DIRAC team, without whom this would not be possible.

... more of the team ...

APROC:
Chen Yi Chien
ShuTing Liao
Eric Yen
Jhen-Wei Huang
Horng-Liang Shih
WeiJen Chang
Felix Lee
Syue-Yi Liaw
Tz Ke Wu
Albert Uang

IJS:
Marko Bracko
Andrej Filipcic
Borut Kersevan
Dejan Lesjak
Jan Jona Javorsek
Jernej Porenta

CESNET:
Jan Kundrat
Jan Svec
Jiri Chudoba
Milos Mulac
Miroslav Ruda
Tomas Kouba

CYFRONET:
Andrzej Ozieblo
Lukasz Flis
Maciej Pawlik
Marek Magrys
Patryk Lason
Wojciech Ziajka

KEK:
Go Iwai
Manabu Matsui
Takashi Sasaki
Yoshiyuki Watase

KIT:
Andreas Heiss
Andreas Motzke
Artem Trunov
Bruno Hoeft
Christopher Jung
Dimitri Nilsen
Doris Ressmann
Foued Jrad
Holger Marten
Ingrid Schaeffner
Jos van Wezel
Manfred Alef
Marian Zvada
Silke Halstenberg
Ursula Epting
Xavier Mol

Melbourne:
Tom Fifield
Tim Dyce

KISTI:
Cristophe Bonnaud
Beob Kyun Kim
Jae-Hyuk Kwak
Jonghu Lee
Soonwook Hwang
Sunil Ahn

UNL:
Brian Bockelman
Carl Lundstedt
Garhan Attebury

VPI:
Miles Gentry
Roger Link

Fermilab:
Burt Holzman

Versi:
Ulrich Felzmann

Belle II and SuperKEKB
at KEK in Japan

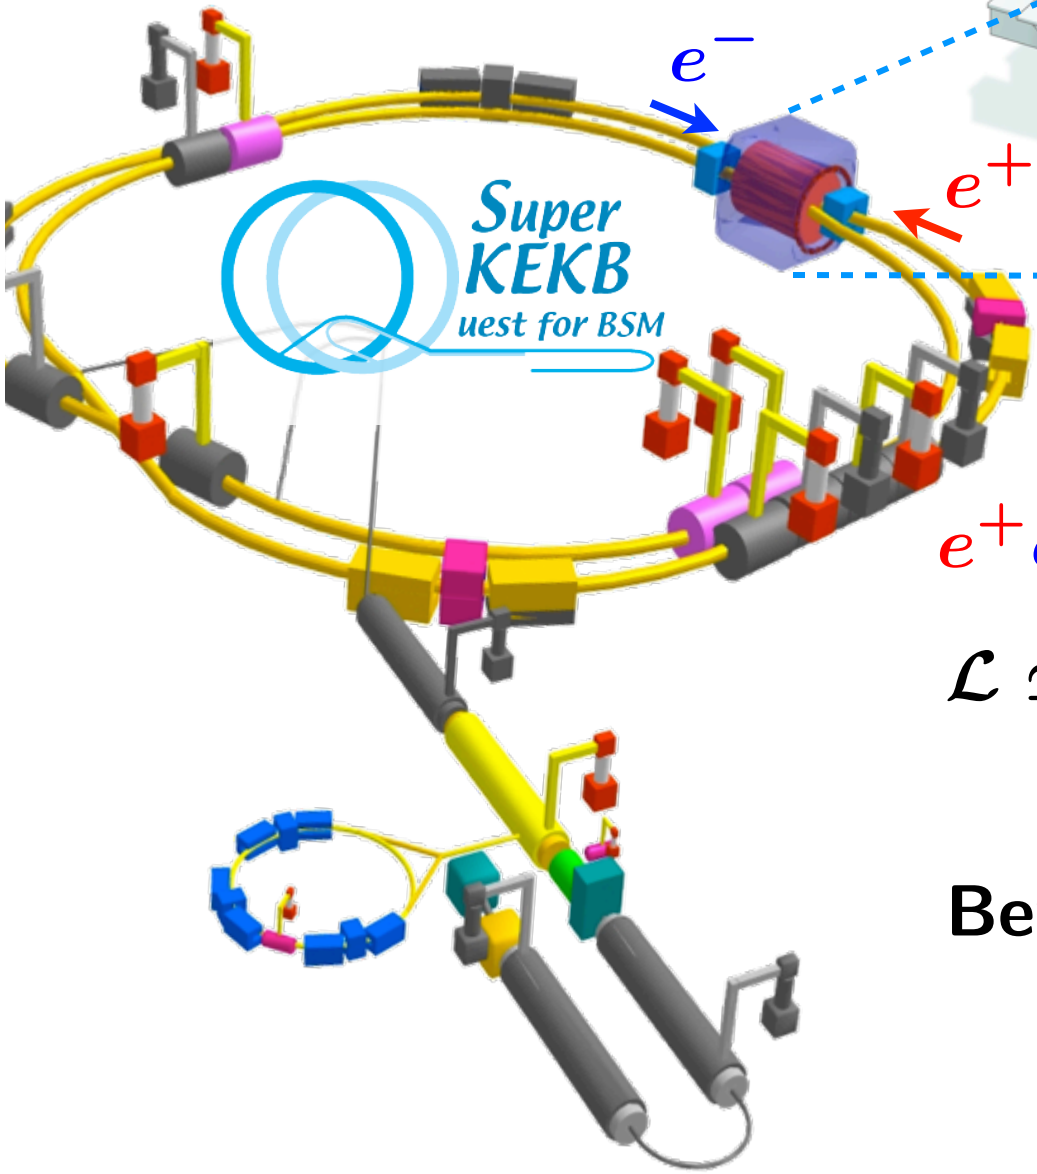

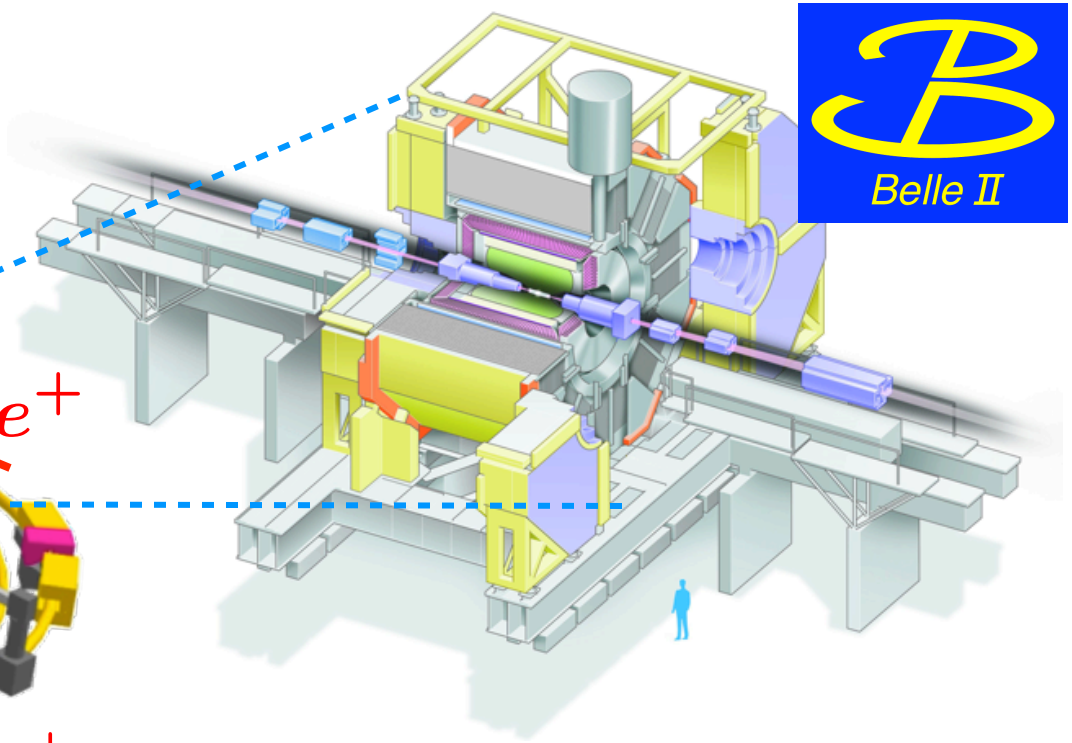

e^+e^- **collisions at** $E_{\text{cm}} = M_{\Upsilon(4S)}$

$\mathcal{L} \simeq 8 \times 10^{35}\,\text{cm}^{-2}\text{s}^{-1}$

$\rightarrow \int \mathcal{L} dt \approx 50\,\text{ab}^{-1}$ **by 2021**

Beyond the Standard Model :

$B \rightarrow \tau\nu,\ B \rightarrow K\pi,$

$B \rightarrow X_s\gamma,\ B \rightarrow K^{(*)}\ell\ell,\ \ldots$

rare charm and τ **decays**

Belle II Collaboration: 13 countries/regions, 60 institutions, ~350 collaborators

Belle II's U.S. grid site will be at
Pacific Northwest National Laboratory

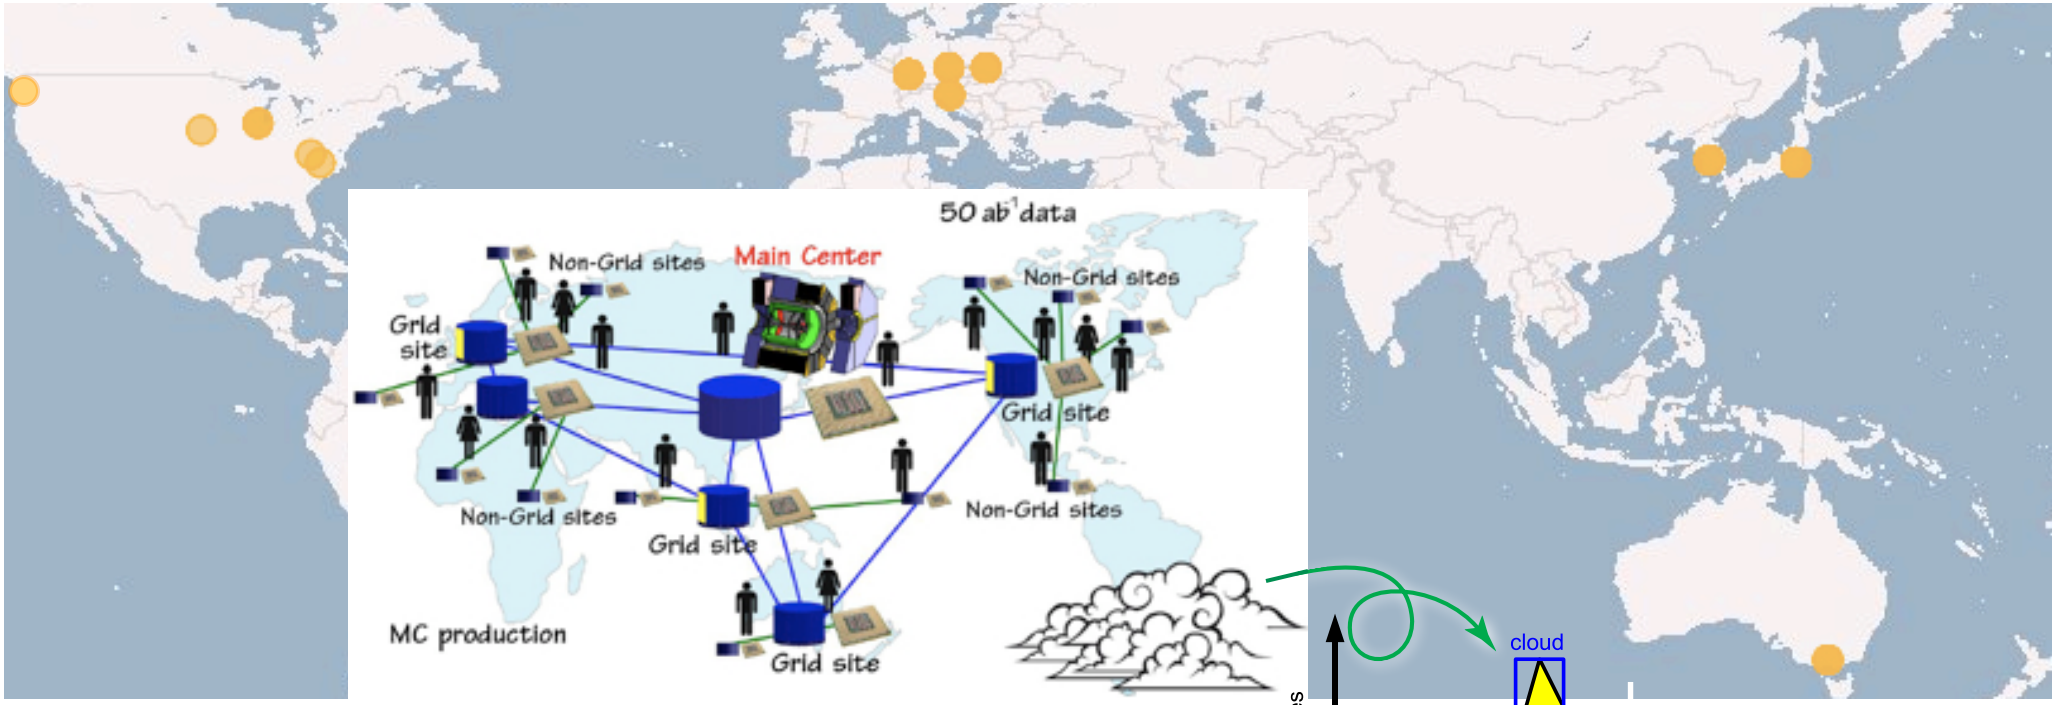

- ✓ EGI–OSG interoperable VO: 13 countries
- ✓ Jobs submitted using DIRAC (of LHCb fame)
  - ❖ Input defined by Metadata queries (powered by AMGA)
  - ❖ Works on the idea of a “Project”
    - a collection of jobs
    - different data per job
    - same analysis code for each job
- ✓ Data source at KEK; distributed on demand
  - ❖ ~50 PB raw data and ~300 PB simulated/skimmed data by 2016
- ✓ MC production ramps up in 2012 in preparation for data in 2014

The Belle II Computing Model

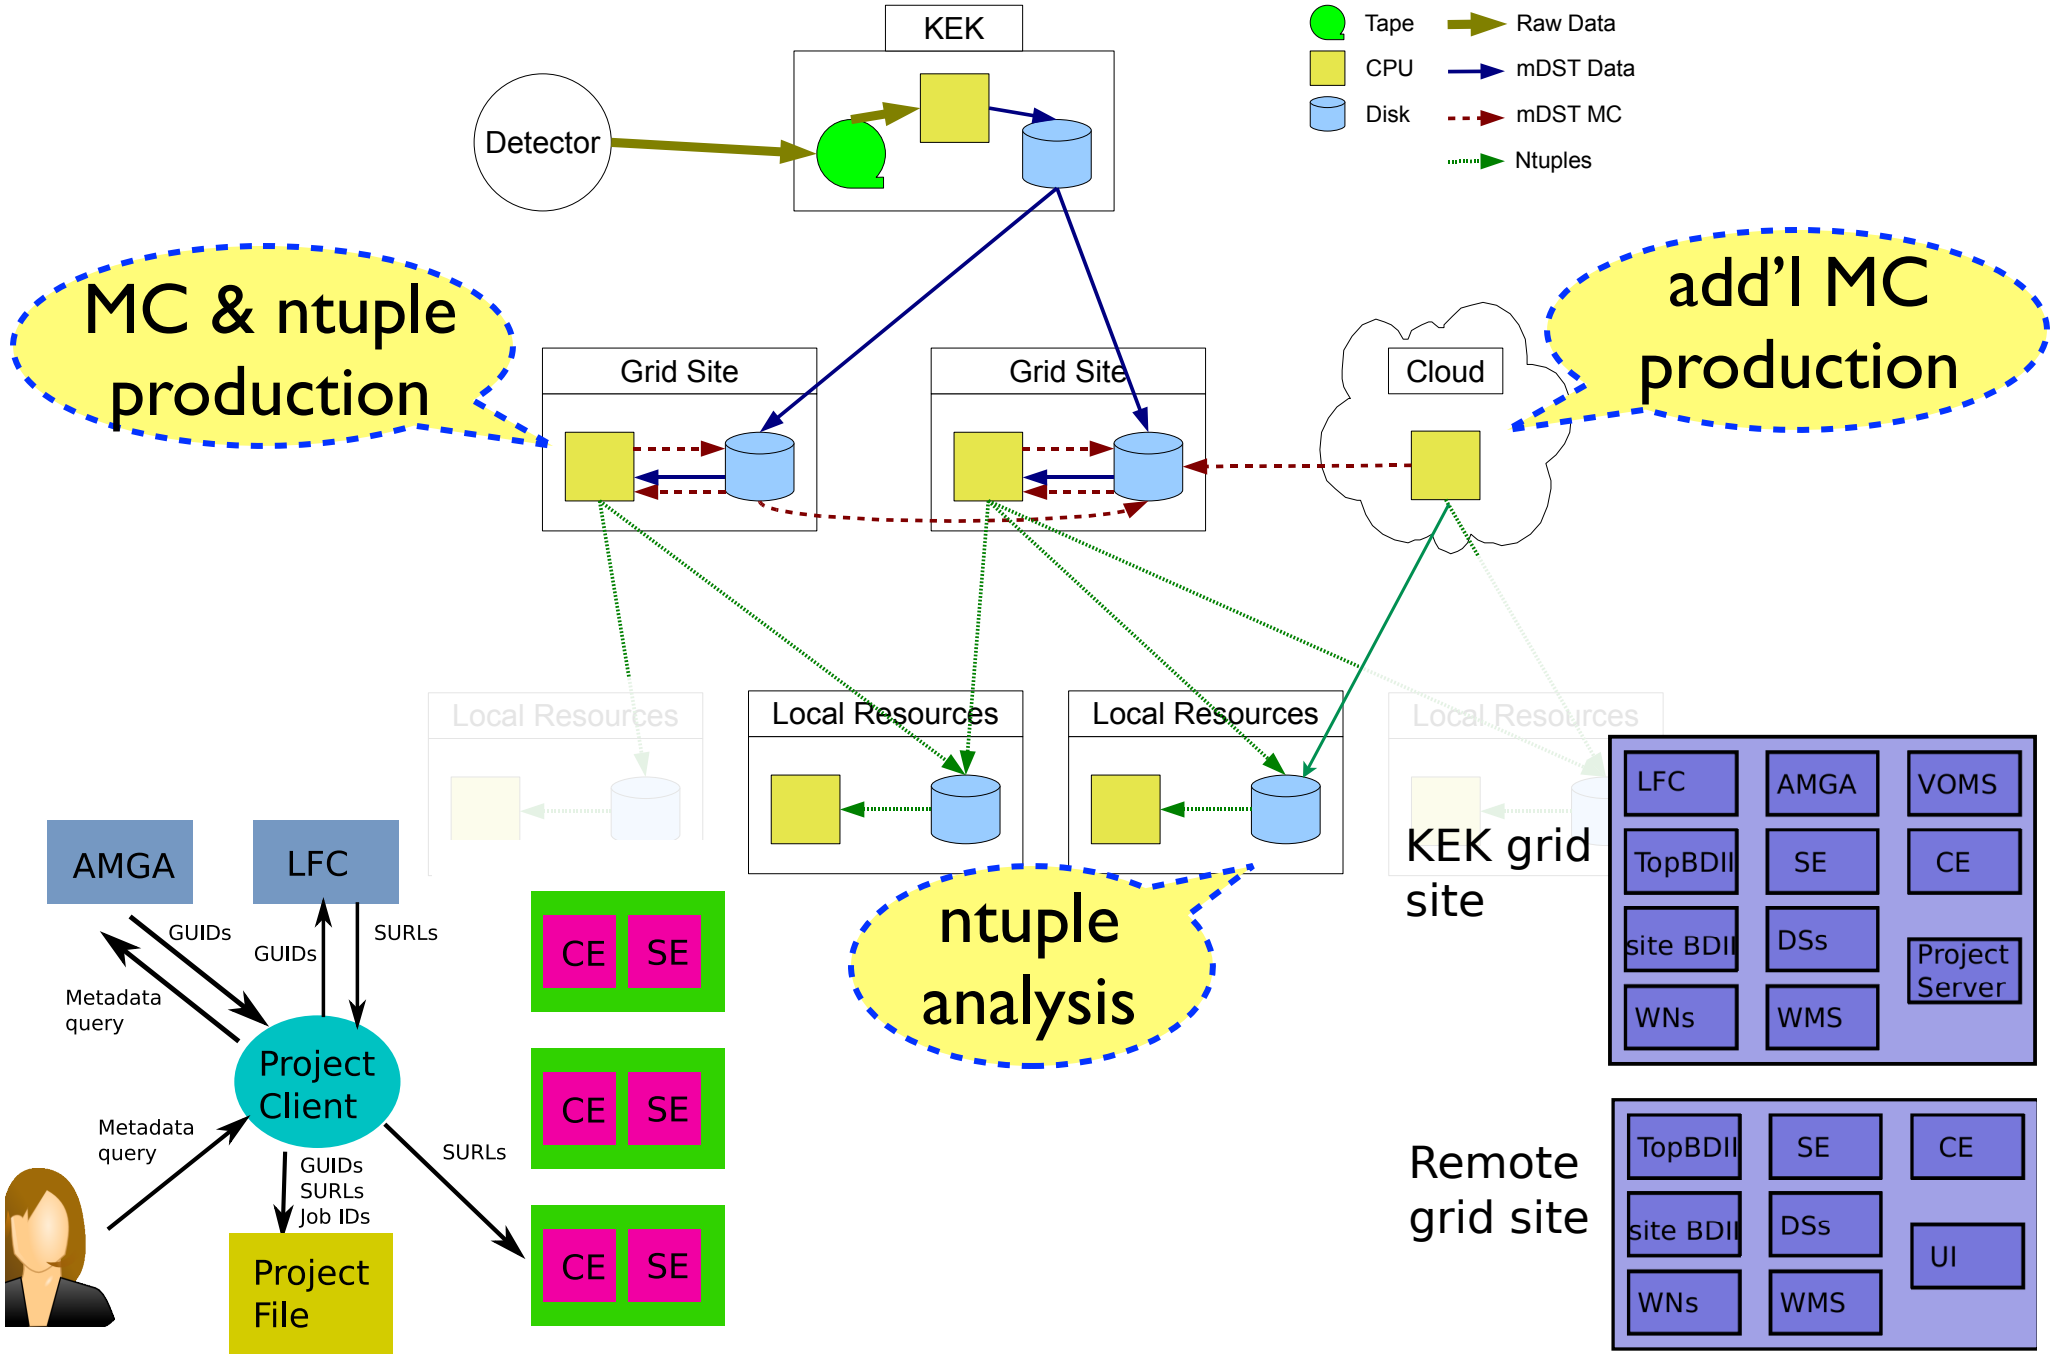

CDF VO status and outlook

Rick Snider
for the CDF Offline Group

OSG All Hands Meeting
Harvard Medical School
March 9, 2011

The CDF experiment

- Fermilab Tevatron Collider
  - Study the elementary structure of matter
    - Measure properties of know particles, forces
    - Search for new phenomena
  - Collides 1 TeV beams of protons and antiprotons head-on at center of CDF experiment

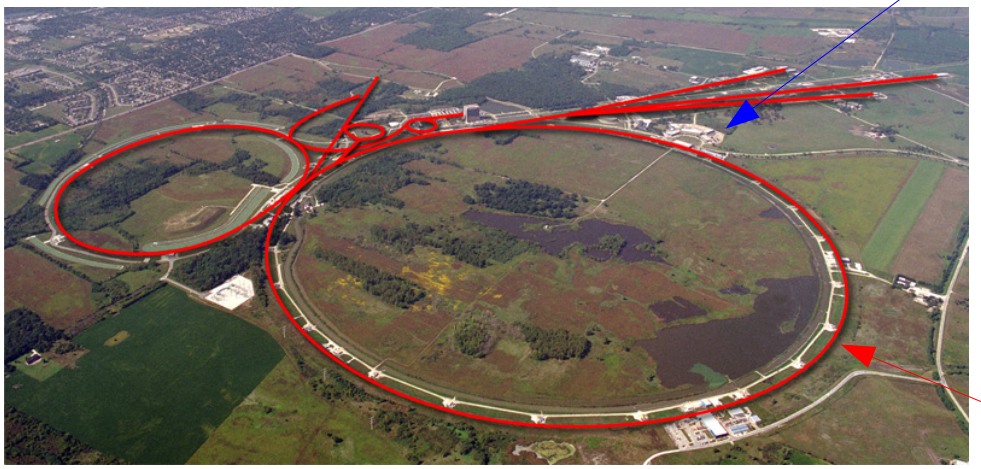

Tevatron collider

Known particles, forces

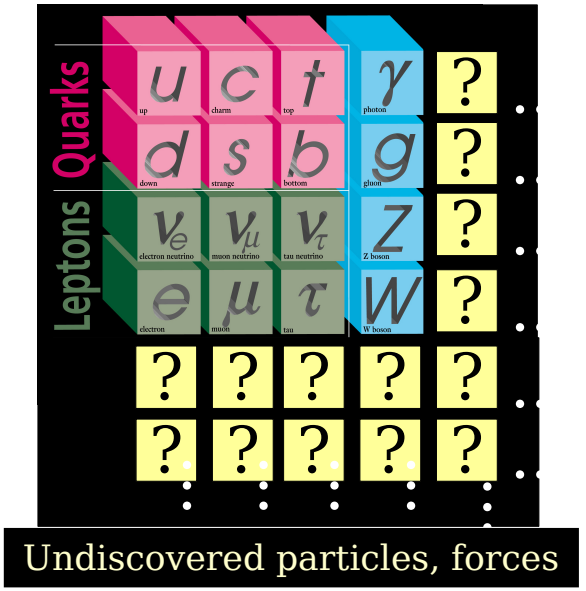

The CDF Collaboration

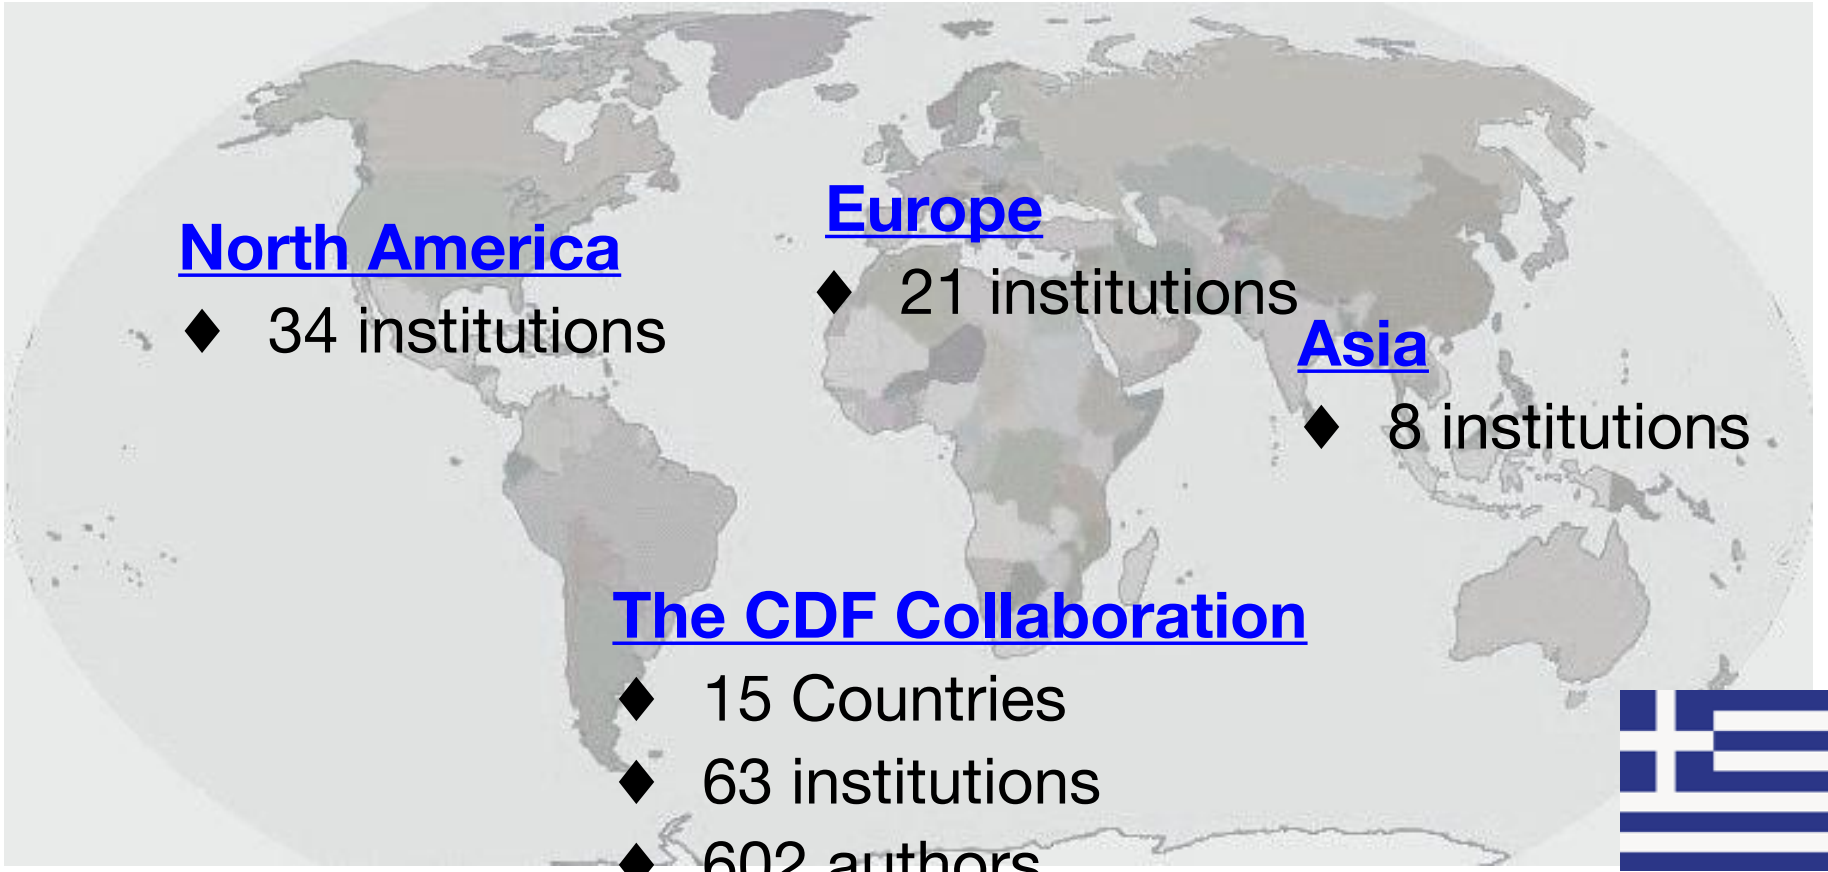

North America
- 34 institutions

Europe
- 21 institutions

Asia
- 8 institutions

The CDF Collaboration
- 15 Countries
- 63 institutions
- 602 authors

Collider Detector at Fermilab experiment

- CDF detector
  - Detector surrounds one of the collision points
    - Collisions every 400 ns
  - Measures critical properties of each “event”
    - Where collision products go, energy and momentum, type of particle, whether some escape without being detected
  - 750k channels of electronics
    - About 150 kB of data per event
  - Record data at 20 MB/s
  - Log data 100+ hrs per week, 40+ weeks per year since 2002
    - 13 x 10^9 events, 1.9 PB of raw data to date

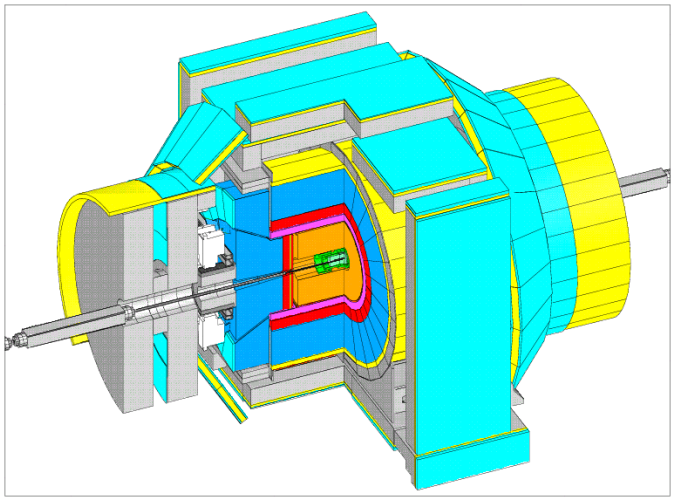

CDF computing model

- Data analysis
  - Process raw data (1.3 times) to reconstruct particle trajectories, energy, momenta, particle ID, etc.
  - Process the results to refine the reconstruction, produce reduced datasets
  - Occasional perform large-scale re-processing of previous two steps
  - Monte Carlo simulation data produced in parallel following the same processing chain
  - Individual physicists:
    - Further process, analyze the reduced datasets, MC results
    - Run toy MC / pseudo-experiments / matrix element calculations

CDF computing model

- Data analysis
  - Process raw data (1.3 times) to reconstruct particle trajectories, energy, momenta, particle ID, etc.
  - Process the results to refine the reconstruction, produce reduced datasets
  - Occasional perform large-scale re-processing of previous two steps
  - Monte Carlo simulation data produced in parallel following the same processing chain
  - Individual physicists:
    - Further process, analyze the reduced datasets, MC results
    - Run toy MC / pseudo-experiments / matrix element calculations

Process on "CDF" OSG resources

Running jobs

5.5k

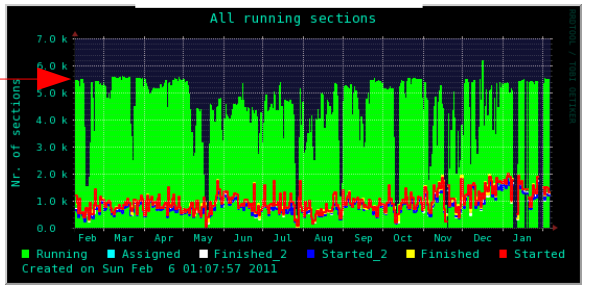

CDF computing model

- Data analysis
  - Process raw data (1.3 times) to reconstruct particle trajectories, energy, momenta, particle ID, etc.
  - Process the results to refine the reconstruction, produce reduced datasets
  - Occasional perform large-scale re-processing of previous two steps
  - Monte Carlo simulation data produced in parallel following the same processing chain
  - Individual physicists:
    - Further process, analyze the reduced datasets, MC results
    - Run toy MC / pseudo-experiments / matrix element calculations

Running jobs by site

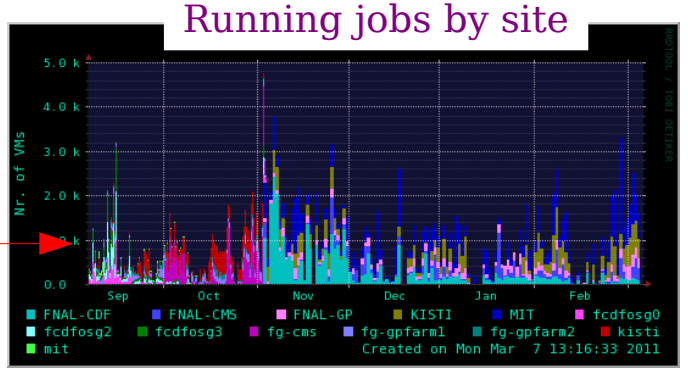

1k

Process primarily off-site and opportunistically on Fermigrid

CDF computing model

- Data analysis
  - Process raw data (1.3 times) to reconstruct particle trajectories, energy, momenta, particle ID, etc.
  - Process the results to refine the reconstruction, produce reduced datasets
  - Occasional perform large-scale re-processing of previous two steps
  - Monte Carlo simulation data produced in parallel following the same processing chain
  - Individual physicists:
    - Further process, analyze the reduced datasets, MC results
    - Run toy MC / pseudo-experiments / matrix element calculations

CDF site + Fermigrid + CDF collaborating sites*

* New within the past year.

Outlook for 2012

- OSG has been critical for success of CDF computing
  - Depend upon software infrastructure, operations support,...
  - Run ends in 2011, but large-scale computing will continue for several years
- A major task: re-processing
  - Several large re-processing tasks started last year, some still pending
  - Expect to re-process significant fraction of dataset after run ends
  - Re-processing model has been to tax general users
    - Costly to users who often have pressing short-term physics goals
  - Investigated running reconstruction on remote CDF + non-CDF resources
    - Have proof in principle that this works at some scale
    - Will consider using OSG more broadly for some re-processing

Compact Muon Solenoid

- 2000 physicists across 200 institutions in 40 countries, give or take a few.
- In 2010, we accumulated 35/pb of data.
  => Led to ~ 30 publications so far.
- For 2012, we expect up to 2000/pb of data.
- 300Hz rate of 200kB events at 30% duty cycle of the accelerator = 6e11 kB/year ~ PetaByte of user analysis data/year
  - Not counting RAW, RECO, nor simulation.
  - Not counting same data processed using multiple different software releases.
- A typical analysis processes more than 10% of the events.

Science Objectives of the LHC

- We study the fundamental constituents of matter and their interactions.
  - The Origin of Mass => Higgs Mechanism
  - Matter and Energy content of the Universe
    - Only ~4% is made of ordinary matter.
    - ~29% is “dark matter” and ~67% is “dark energy”, neither of which we have an explanation for.
  - Nature of space-time
    - Are there more than 4 dimensions in nature as “desired” by some theories ?

Science Objectives of the LHC

- We study the fundamental constituents of matter and their interactions.
  - The Origin of Mass => Higgs
  - Matter and Energy content of the Universe
    - Only ~4% is made of ordinary matter.
    - ~29% is "dark matter" and ~67% is "dark energy", neither of which we have an explanation for.
  - Nature of space-time
    - Are there more than 4 dimensions in nature as "desired" by theories ?

Guaranteed

Well Motivated

Speculative

CMS Outlook

- 2011: full year of data taking
- 2012: full year of data taking
- 2013/14: LHC accelerator shutdown
- Beyond that, fkw can’t read the tea leaves.

The DØ Collaboration

DØ is an international collaboration of 487 physicists from 20 nations who have designed, built and operate the DØ detector at the Tevatron and perform data analysis

Institutions: 80 total, 36 US, 44 non-US

The DØ Experiment

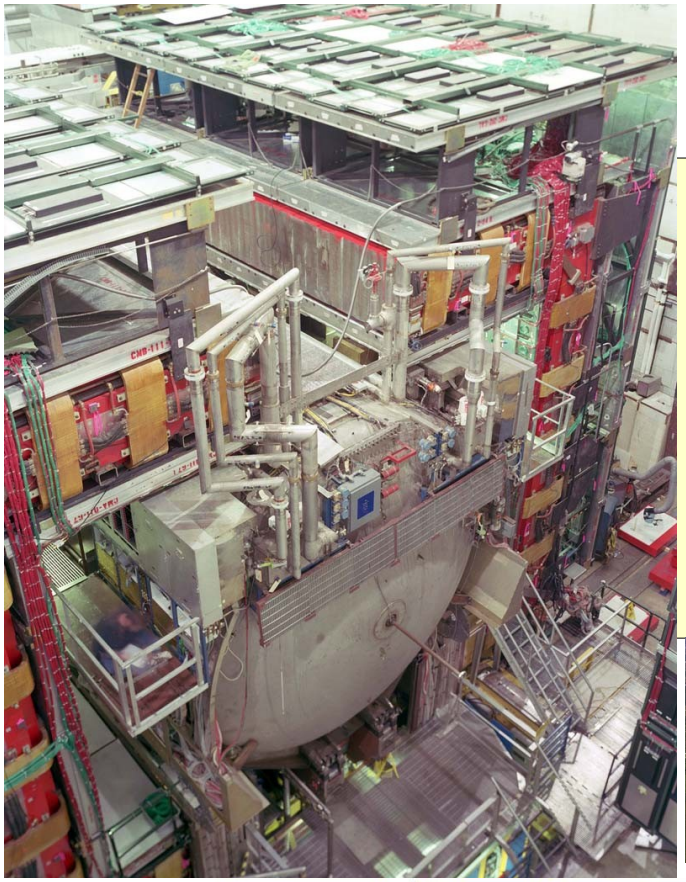

Studies fundamental physical constituents and their interactions by colliding 1 TeV protons with 1 TeV anti-protons

Topics of inquiry include:

Bottom quark physics
Higgs boson physics
Quantum Chromodynamics
Top quark physics
Electroweak physics
New phenomena beyond SM

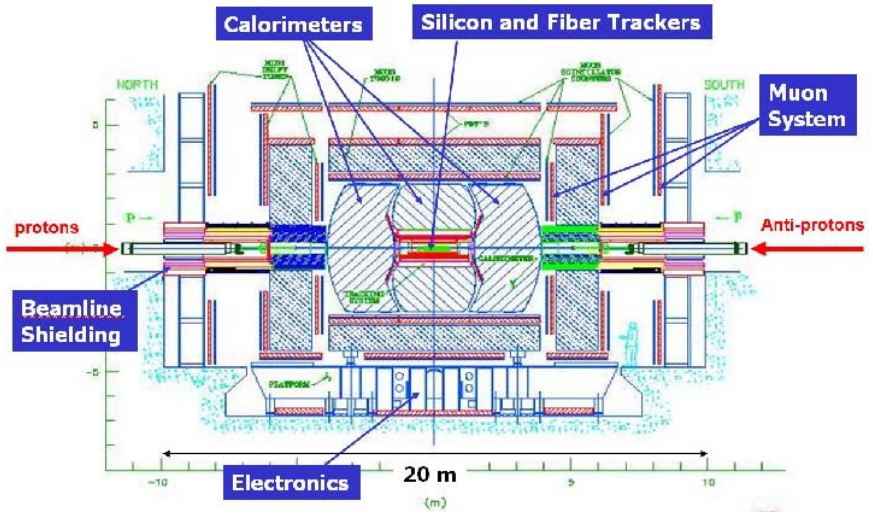

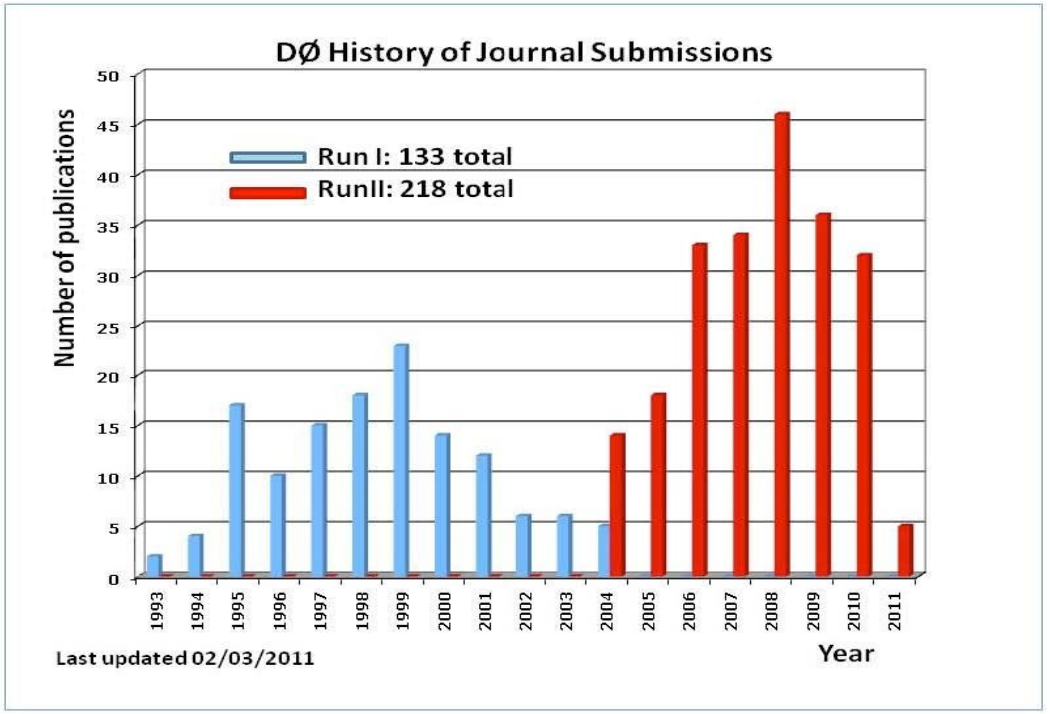

DØ OSG Use

Monte Carlo Production

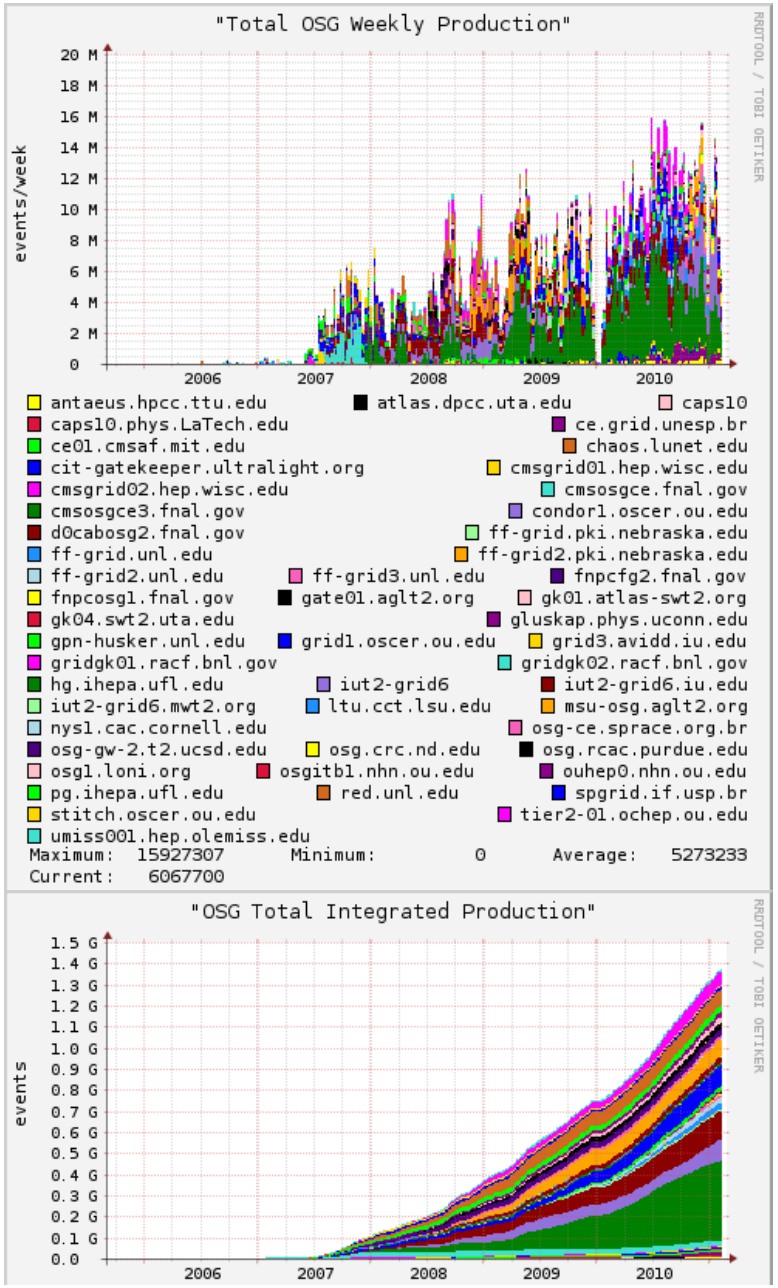

Data Production

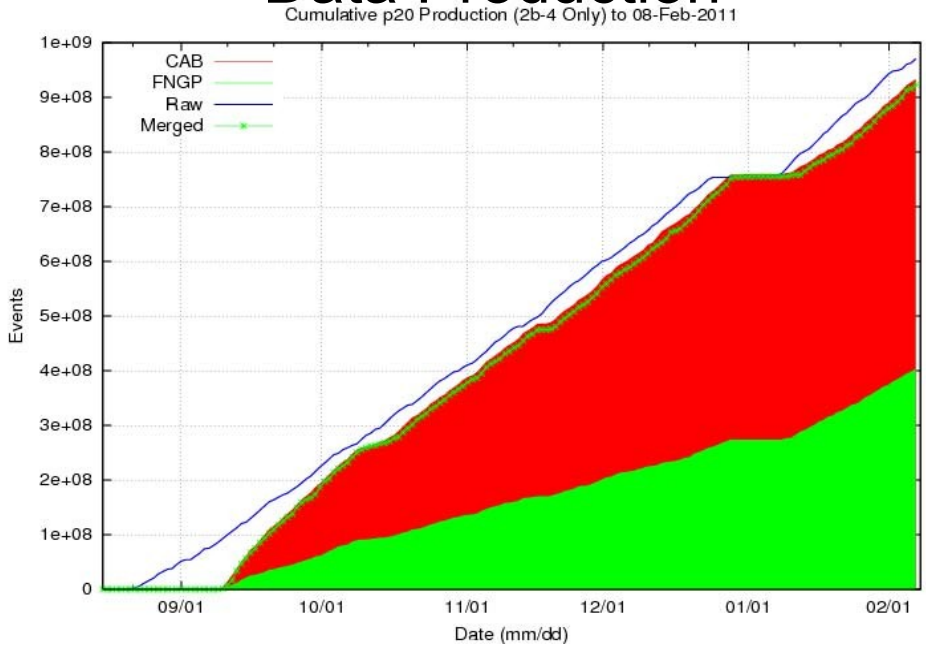

- **DØ** integrates legacy grid and data handling systems with the OSG. Job submission is all done at FNAL.
- **D**ata production only uses a few CE's located at FNAL with dedicated workers.
- **M**C jobs are submitted from a work queue by a daemon to a scheduler. MC is almost completely opportunistic and currently uses 30 CE's & 8 SE's.
- **T**here are also some CPU intensive analysis jobs done on the OSG.

1.4 billion MC events represents about 5000 CPU-years

DØ 2012 Outlook

- Tevatron program ends in September 2011
  - Data production use of the OSG will end
  - MC production and analysis will continue for several years
- DØ use of the OSG in 2012 is crucial to provide adequate and timely simulated data for physics analyses that produce publications.
- The DØ collaboration thanks the DOE, NSF, and the various international funding agencies that support this work

Joel Snow, Langston University/FNAL, on behalf of the DØ collaboration

OSG Engage VO
Status and Outlook

John McGee, Steve Cox, Mats Rynge

OSG AHM 2011

OSG Engage

- Infrastructure, Services, and human expertise to help new users and communities ramp up to national CI
- A broad set of science domains have been engaged over the years
  - Biochemistry, genomics, mechanical engineering, mathematics, Systems biology, coastal modeling, electrical engineering, etc
  - ‘09/10 project year: 10 publications citing OSG usage due to Engage effort
- Detailed independent analysis of Engage program conducted by Howison et al of CMU via VOSS award

Classes of Engage Users

- Individual researcher
  - many examples
- PI/Lab with steady stream of students
  - Steffen Bass at Duke University, theoretical physics
- Systems Integration
  - Andreas Prlic, RCSB Protein Data Bank

Cumulative Hours Spent on Jobs By Facility

1469 Days from Week 08 of 2007 to Week 10 of 2011

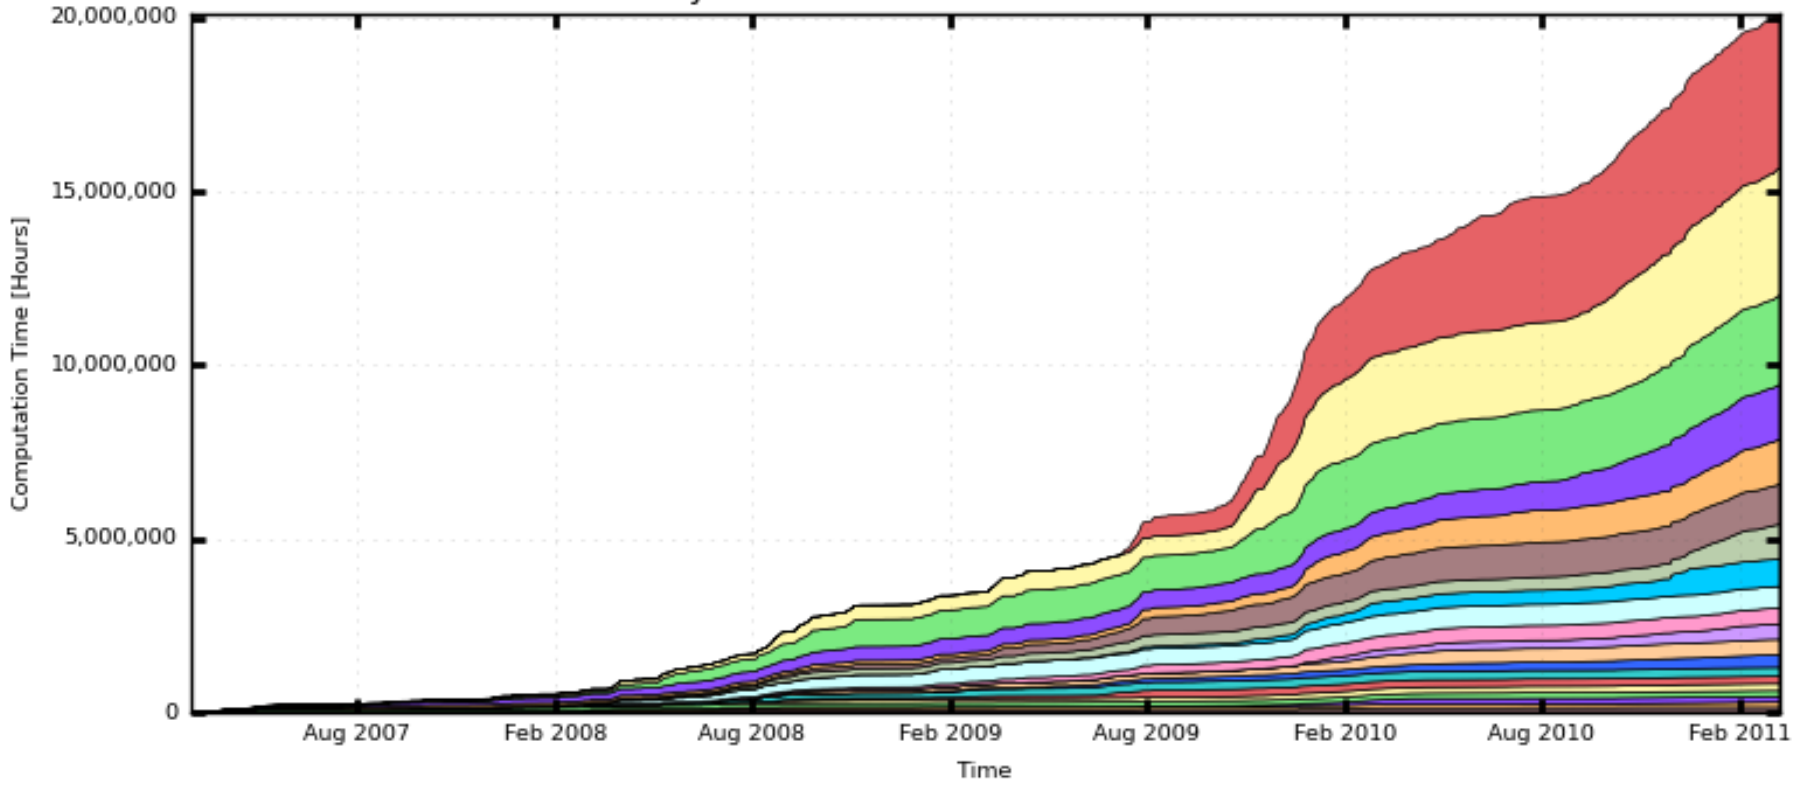

- Firefly (4,462,412)
- USCMS-FNAL-WC1 (3,672,708)
- FNAL_FERMIGRID (2,576,694)
- Other (1,564,796)
- Nebraska (1,292,597)
- UFlorida-HPC (1,156,412)
- UCSDT2 (996,956)
- Clemson-Palmetto (787,443)
- MIT_CMS (623,853)
- Purdue-RCAC (459,315)
- prairiefire (440,683)
- FNAL_GPGRID_1 (439,132)
- RENCI-Engagement (367,954)
- NYSGRID_CORNELL_NYS1 (256,288)
- SMU_PHY (210,671)
- UCR-HEP (204,133)
- TTU-ANTAEUS (165,097)
- IU_OSG (164,287)
- BNL-ATLAS (148,358)
- CIT_CMS_T2 (139,889)

Total: 20,129,686 Hours, Average Rate: 0.16 Hours/s

Migrating to a GlideinWMS FE using SDSC Factory by April 2011

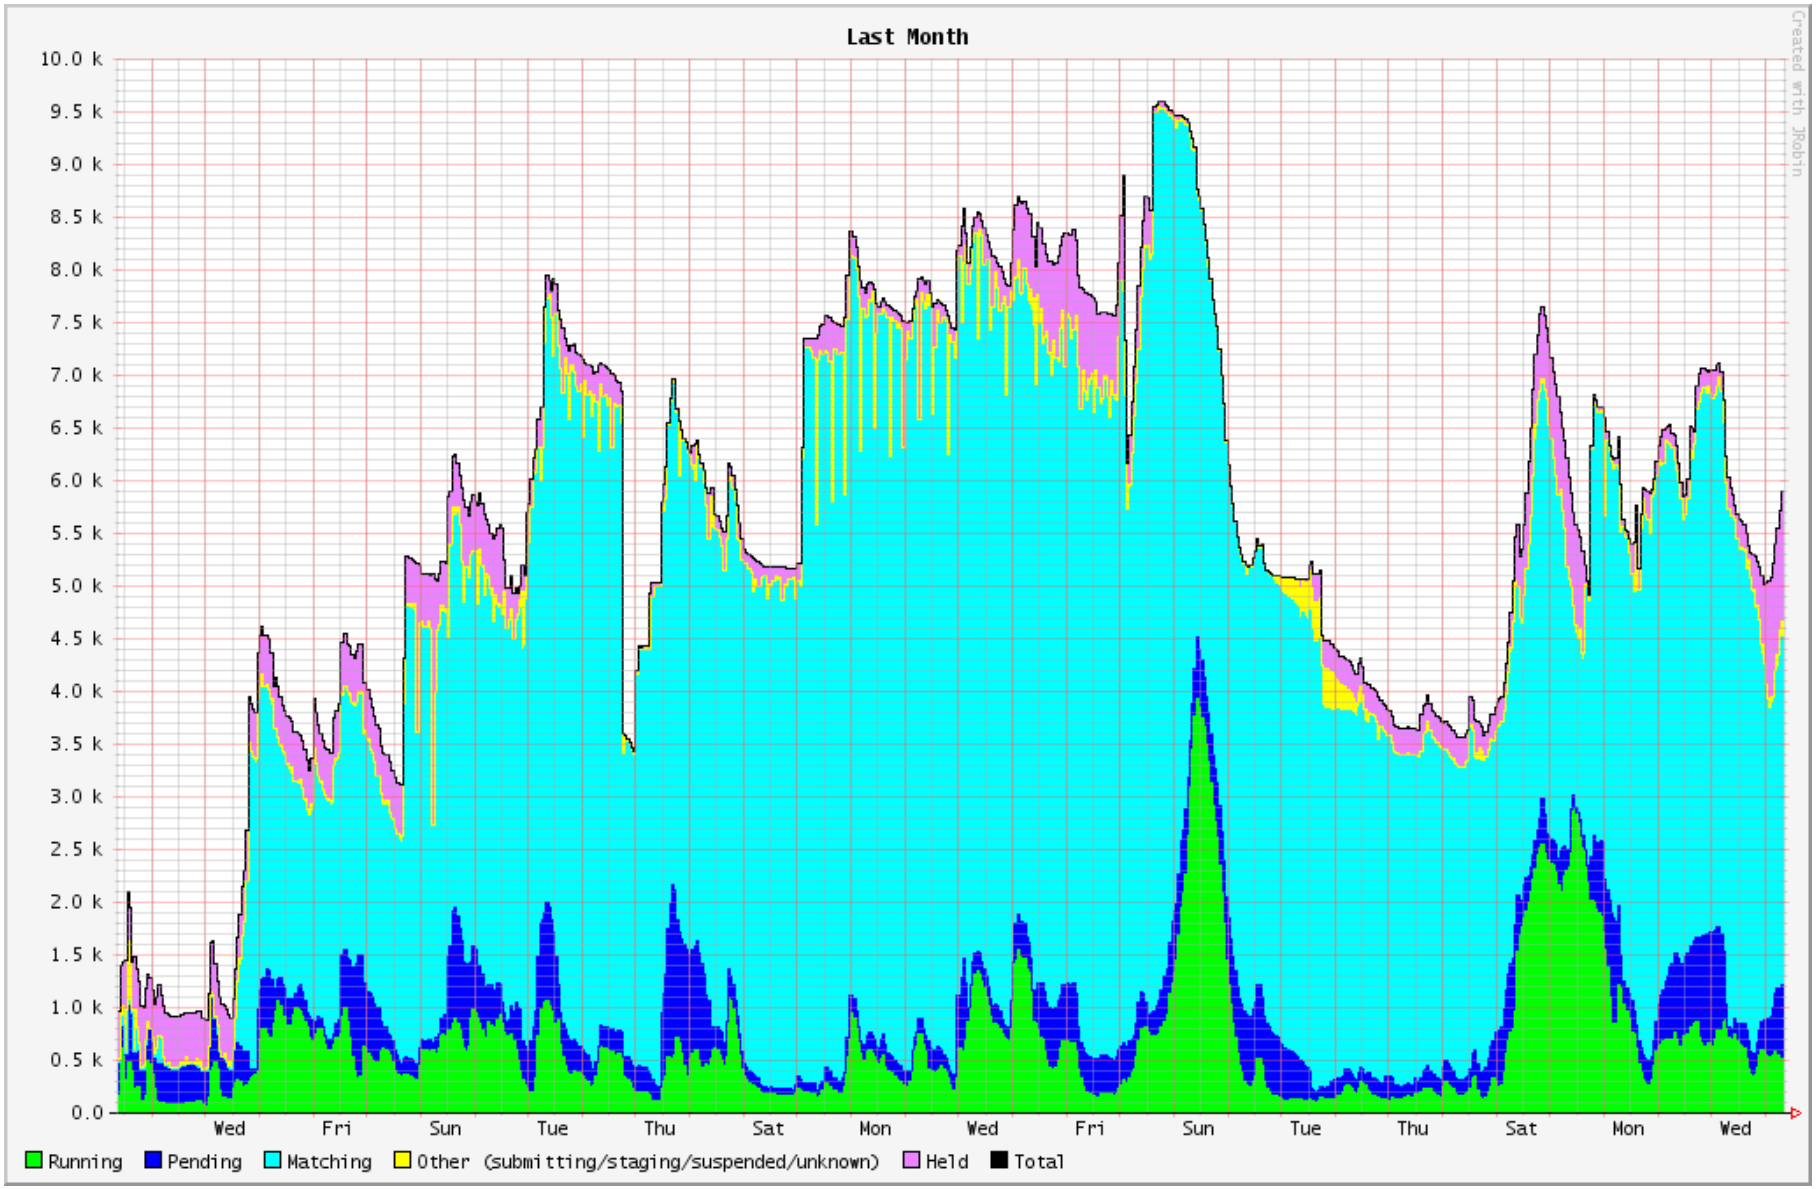

OSG Engage Outlook 1 of 2

- Satellite effort funded through 3/31/2011
- No cost extension available through 3/2012
- Continued support of the Engage infrastructure and submission mechanisms at RENCI
- VO management services may be migrated to the GOC
- It is not clear that funding agencies are interested in individual level researcher support, focus appears to be on large VO/science efforts and campus integration

OSG Engage Outlook 2 of 2

- A lot of interest in opportunistic HTPC from: environmental sciences, molecular dynamics, computational chemistry
- GPGPU
- Focus on PI groups and campus integration

Fermilab VO (Intensity Frontier)

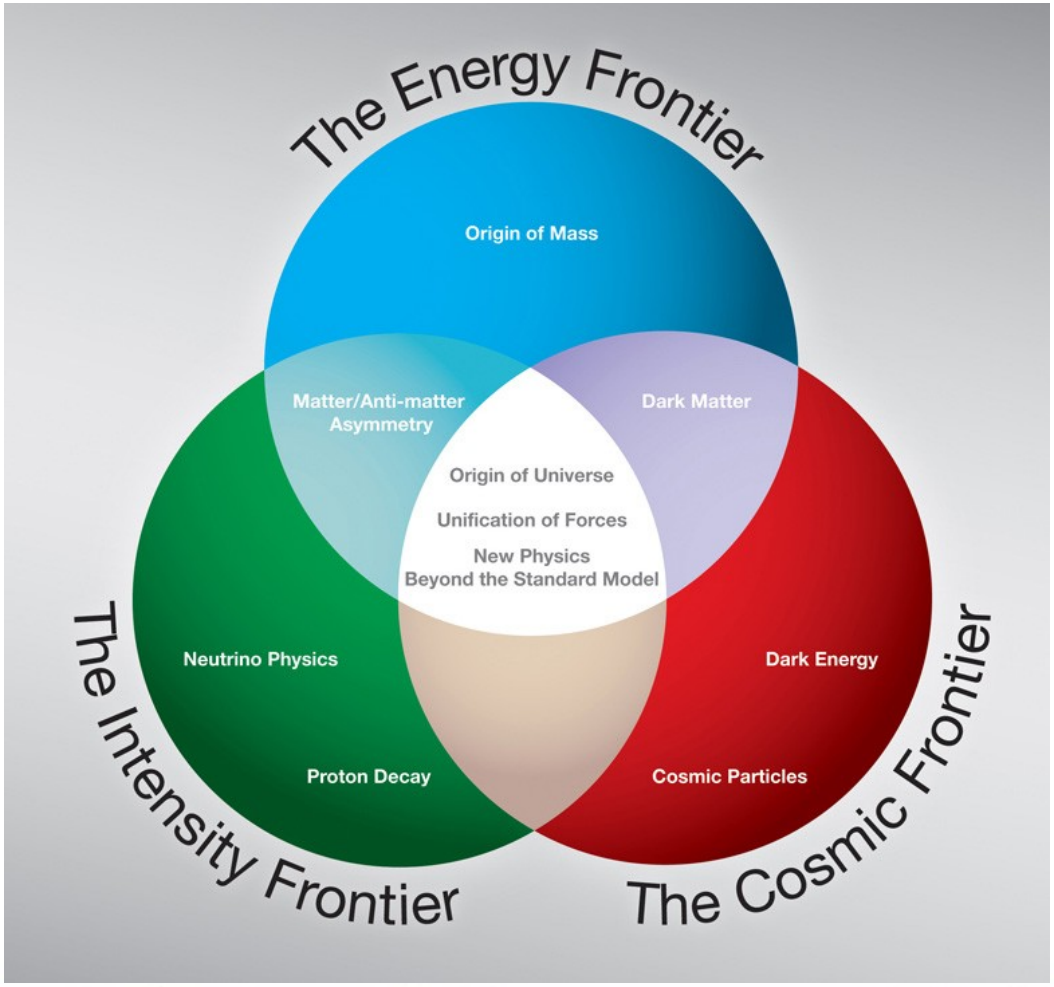

Fermilab VO projects

- All Fermilab employees and users part of Fermilab VO
- Only certain subgroups exported to OSG
- Intensity frontier experiments
  - Running: MINOS, MiniBooNE, MINERvA, NOvA, ArgoNeuT
  - Building/Planning: MicroBooNE, Mu2e, LBNE (soon to be own VO).
- Accelerator and beamline simulation
- Theoretical simulation
- Cosmic frontier—Cryogenic Dark Matter Search
- ALL publications of Intensity Frontier make use of OSG.

Intensity Frontier Hardware Architecture

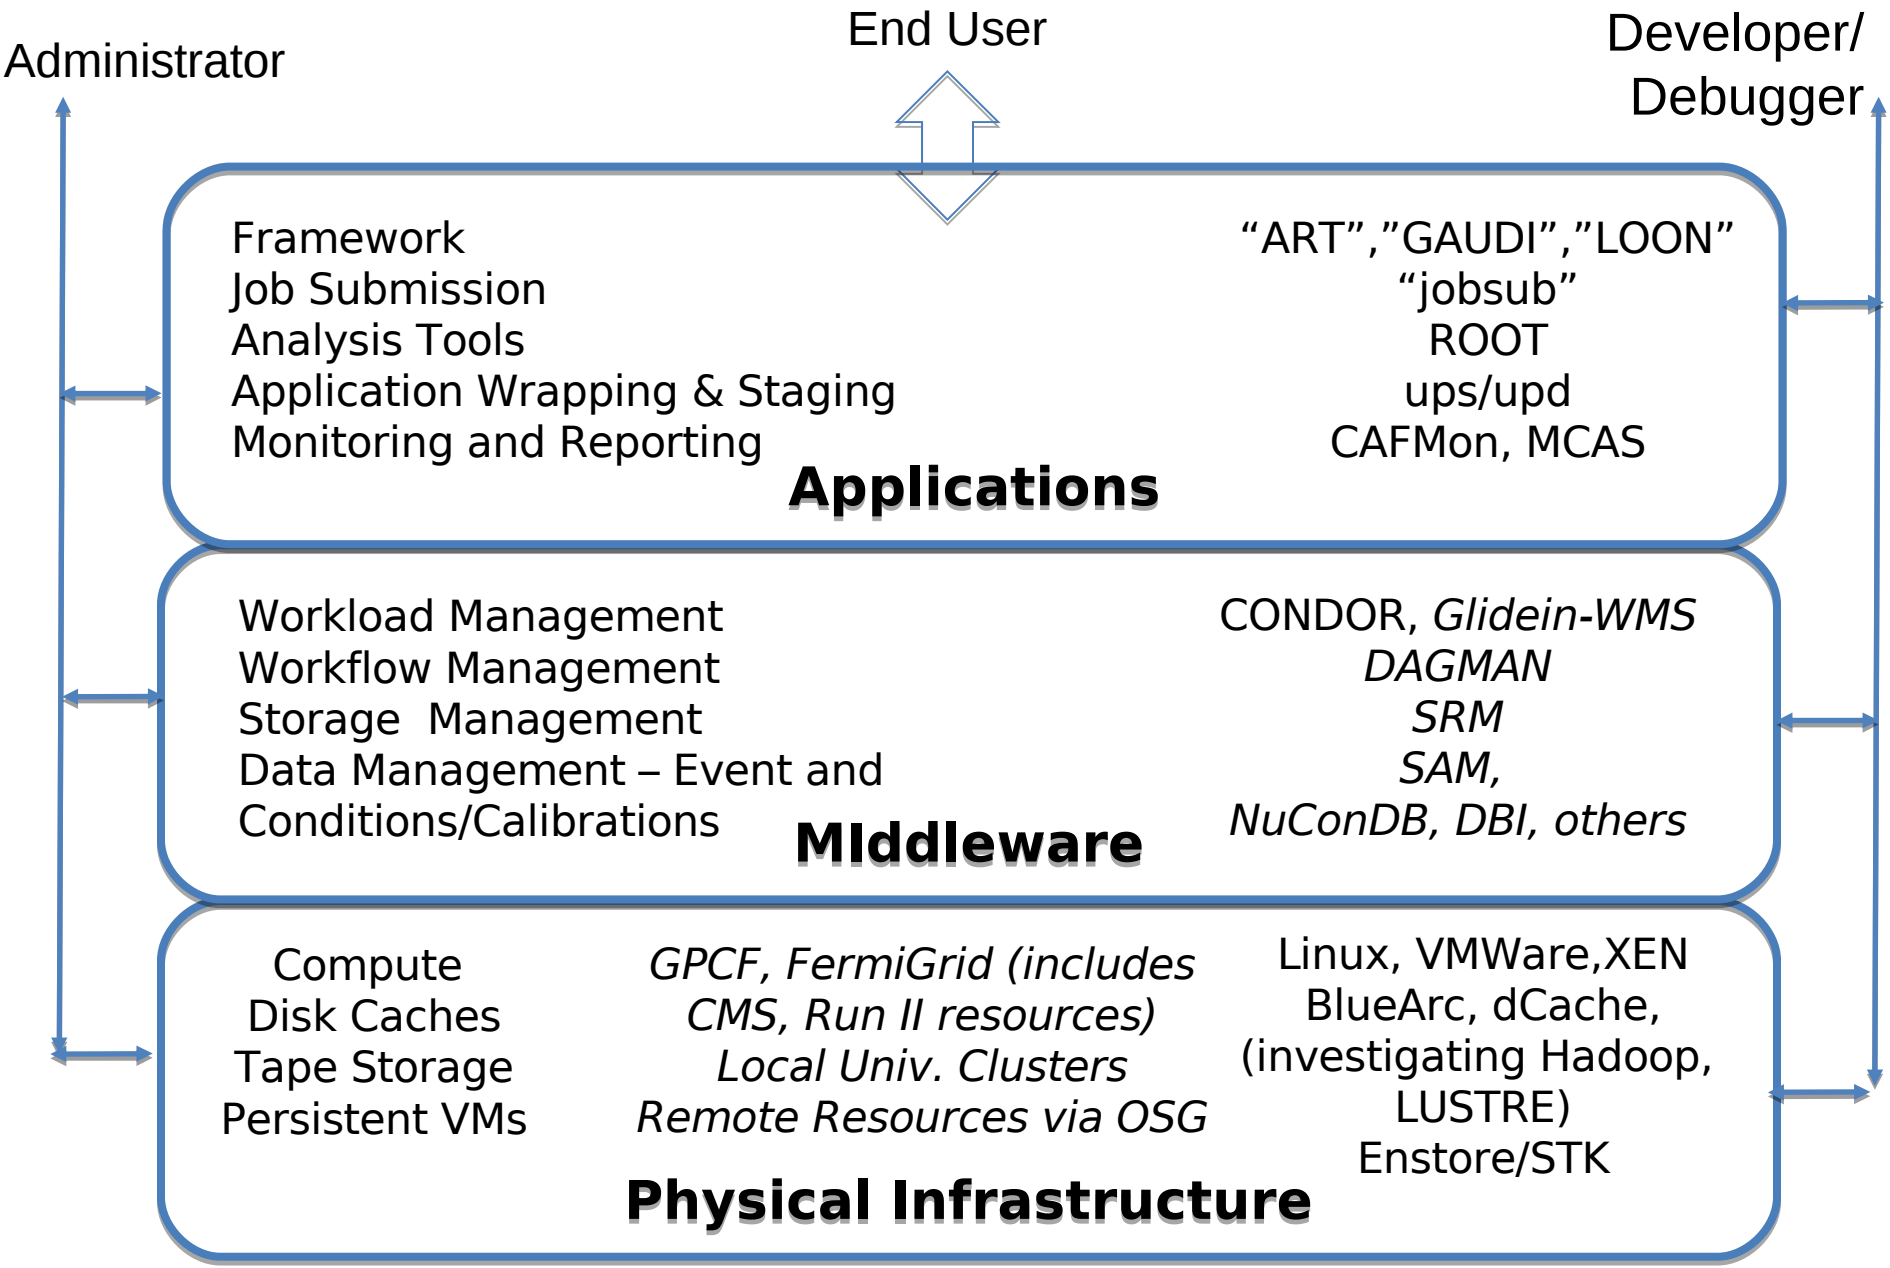

Intensity Frontier Hardware Architecture

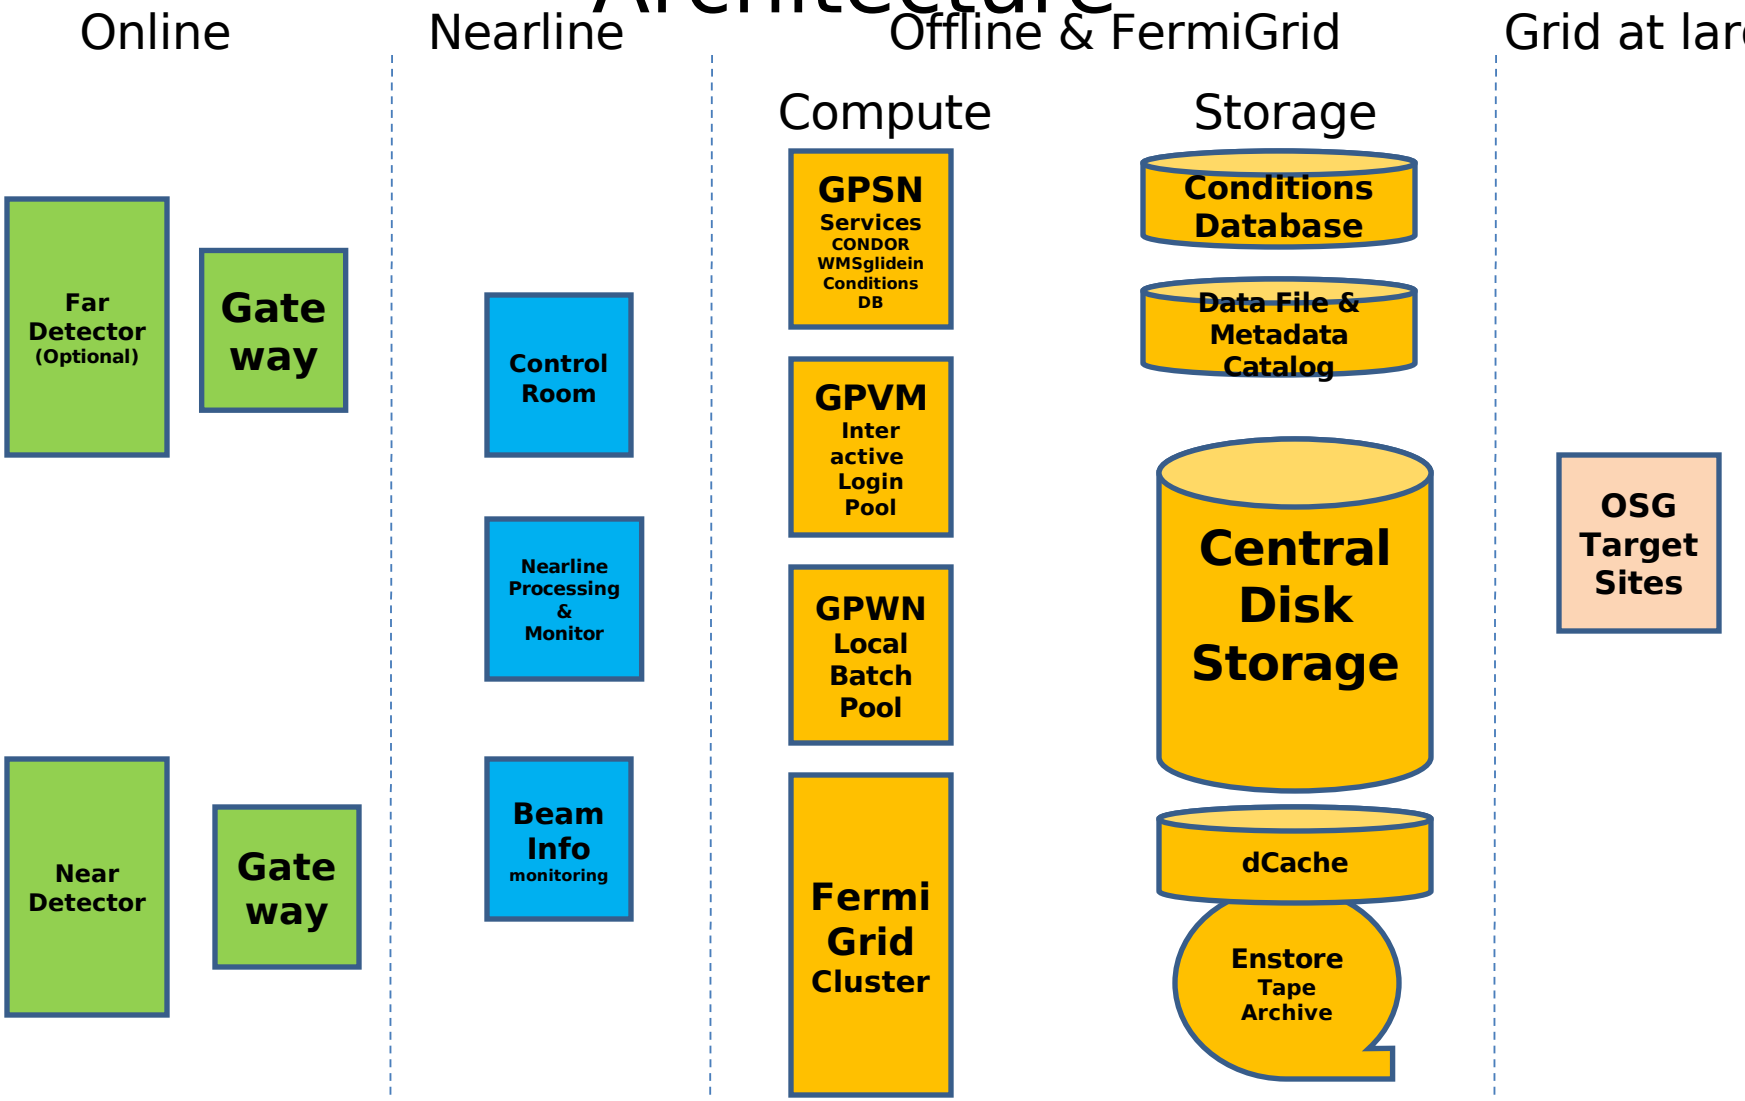

Mu2E High-statistics Background Simulation

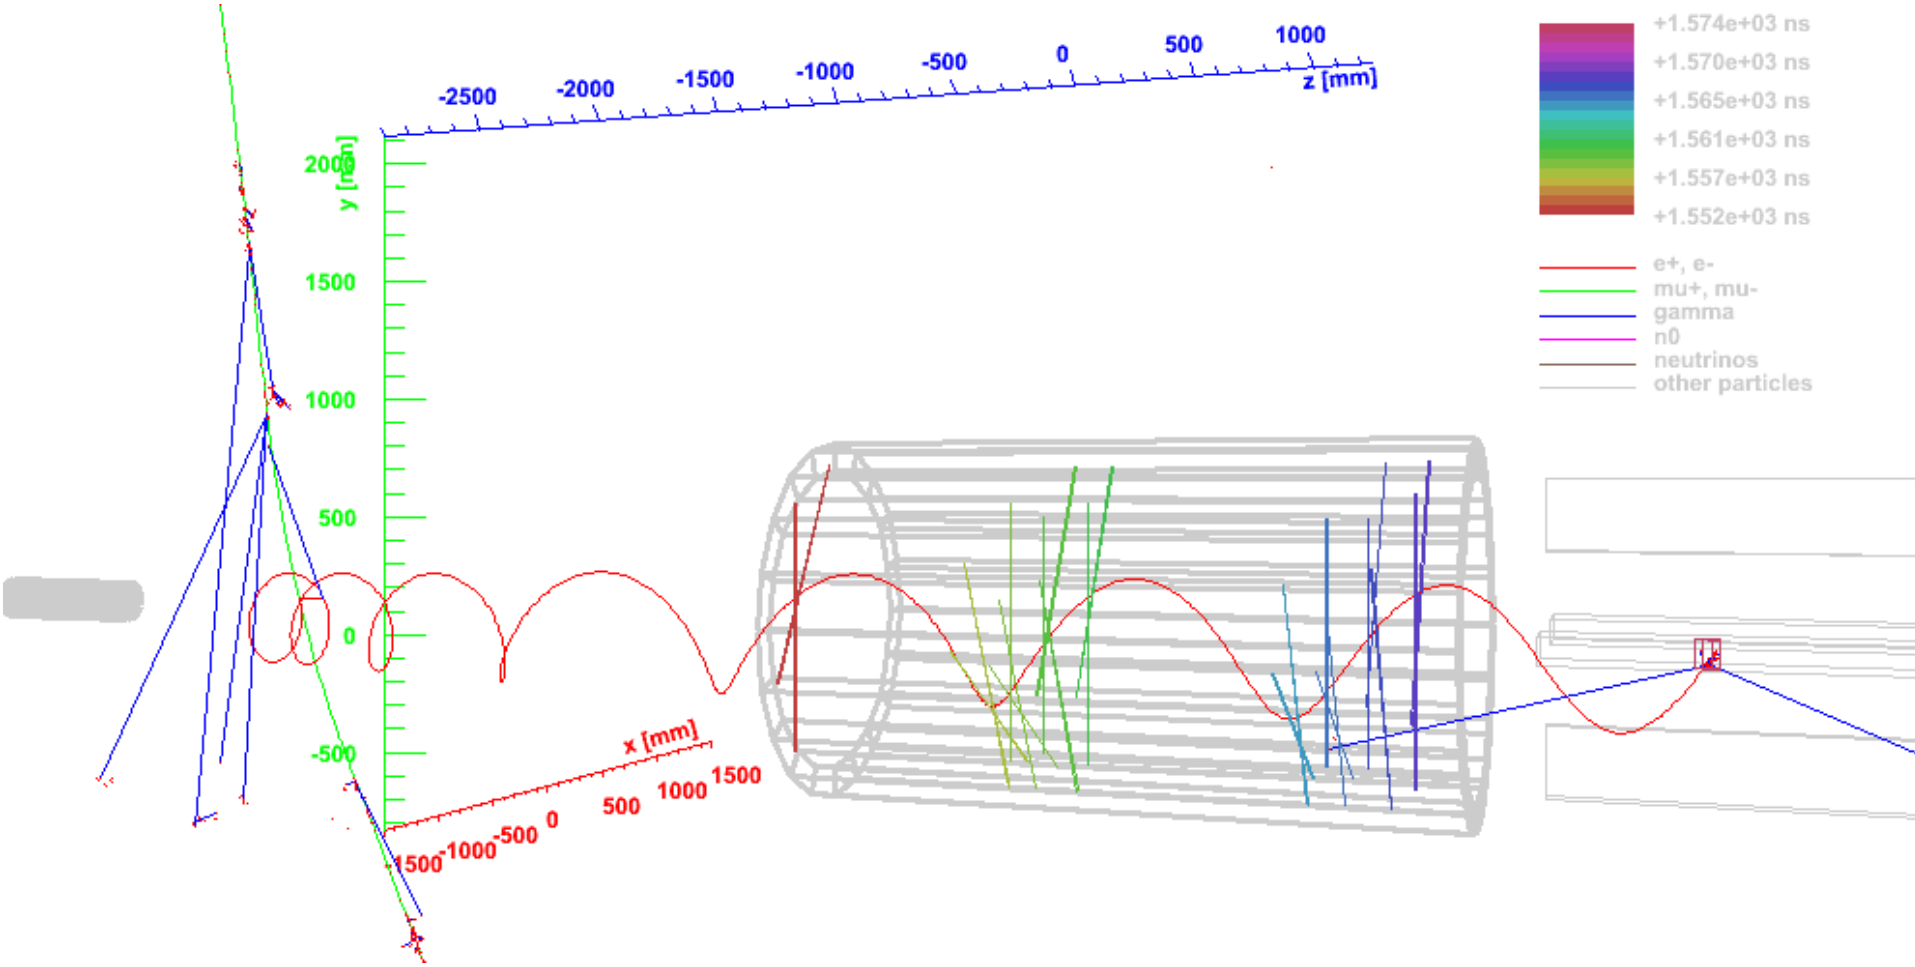

HCC TEAM

- CMS: Brian Bockelman, Carl Lundstedt, Garhan Attebury (Ken Bloom, Aaron Dominguez)
- OSG: Derek Weitzel, Adam Caprez, Ashu Guru
- HCC: Chris Cox, Ryan Lim, Tom Harvill

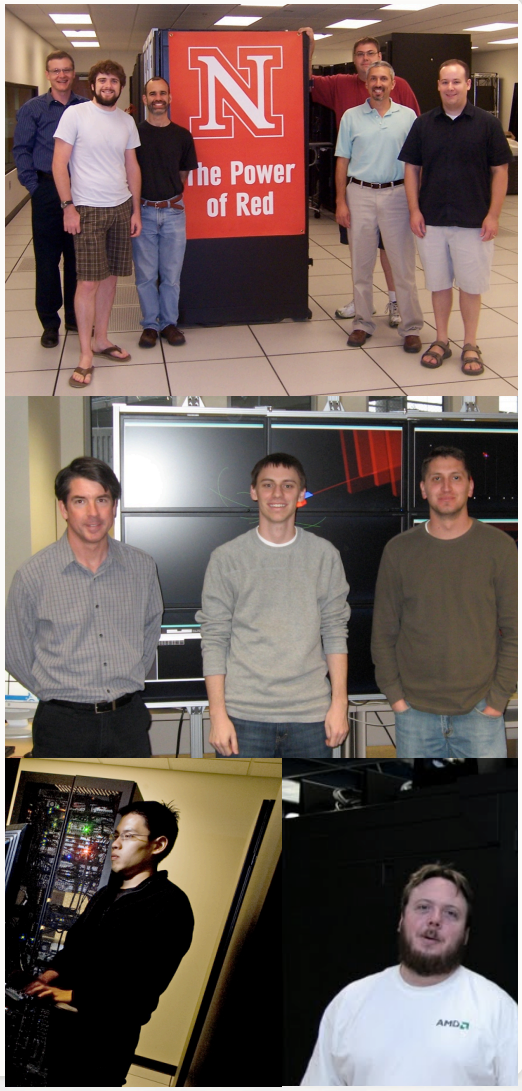

HCC SCIENCE

- CMS [8.8M cpu hr last yr] and other friends [30.2M]
  - Non-HCC engagement: TotalJobs: 472675 TotalWallTime: 1.2M]
- Dr. Bob Powers, UNL Chemistry (CPASS -Comparison of Protein Active Site Structures) [TotalJobs: 126880 TotalWallDuration: 68k]
- Derrick Stolee, UNL Mathematics (Graph Theory Existence Studies) [7.4M]
- Dr. Ding, UNMC (OMSSA - Nuclear Matrix Proteomics) [170,000 cpu hr one weekend]

HCC METHODOLOGY

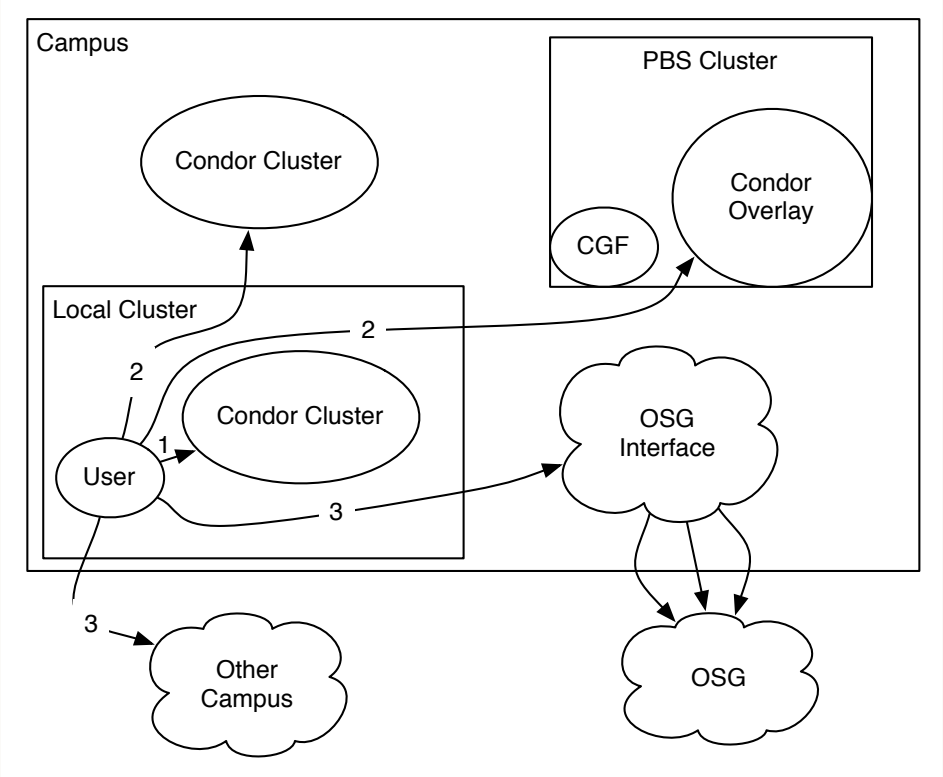

- Campus Grid approach

2011 - 2012

- Collaborative award with Puerto Rico (HubZero)
- Campus grids, HTPC, Xrootd
- DYNES
- 20 gbps between HCC sites in Omaha and Lincoln
- New hires (3)
- OSG AHM (!)

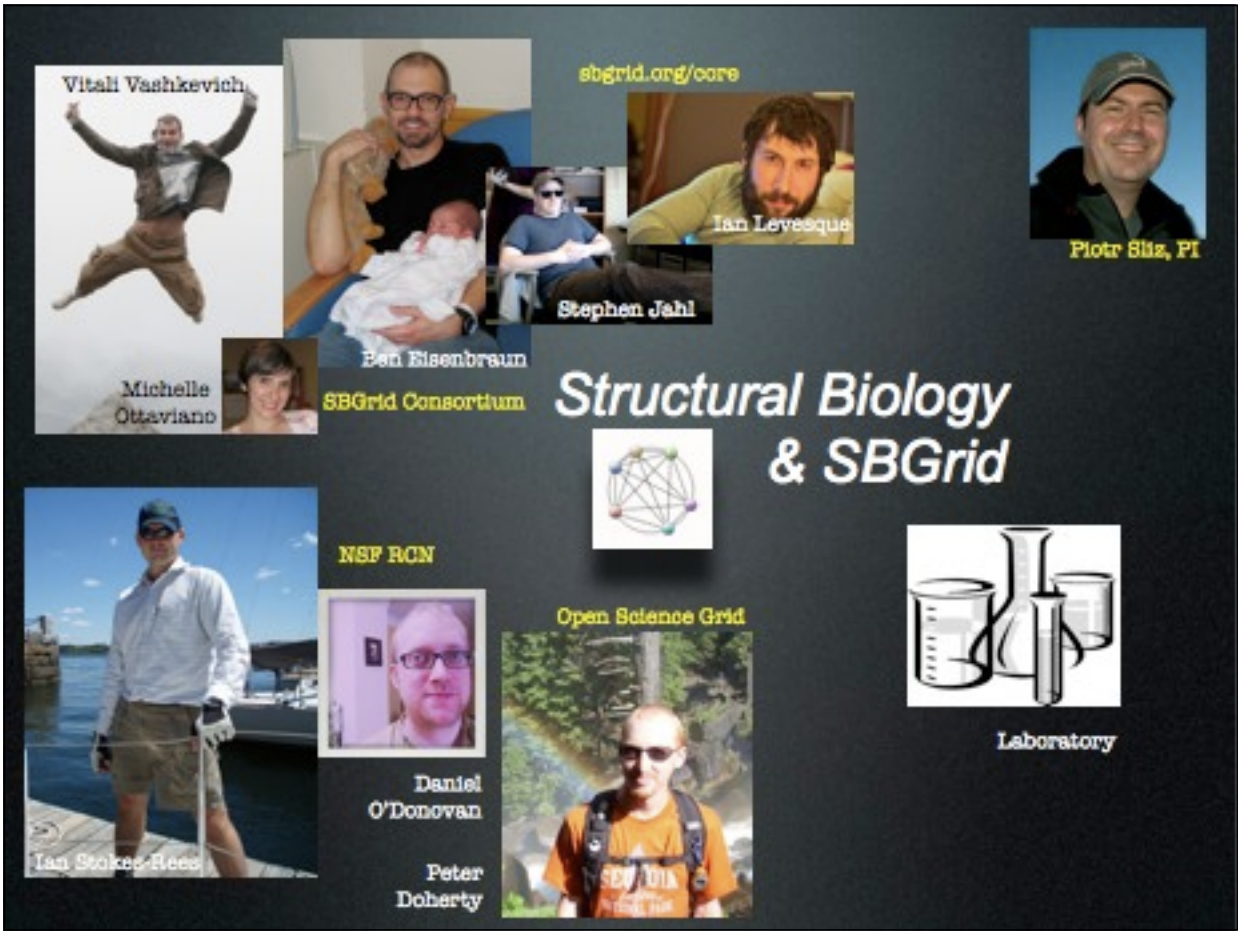

Science: Structural Biology

X-ray Crystallography: Wide Search Molecular Replacement

- Stokes-Rees I, Sliz P Protein Structure Determination by exhaustive search of Protein Data Bank derived databases. PNAS (2010).
- Stokes-Rees I, Sliz P. Compute and data management strategies for grid deployment of high throughput protein structure studies. 3rd IEEE workshop on Many-Task Computing on Grids and Supercomputers. (2010).

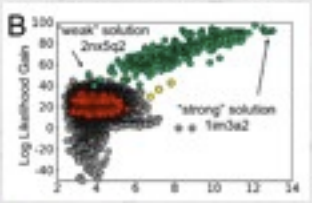

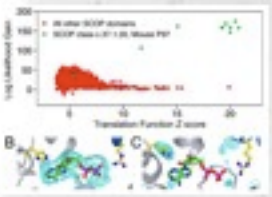

Molecular Dynamics:

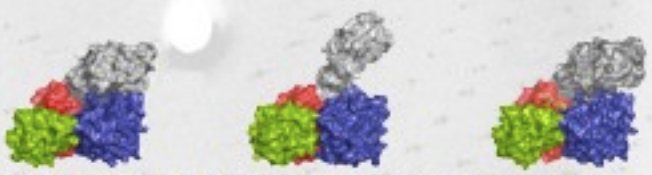

Lazarus MB, Nam Y, Jiang J, Sliz P, Walker S. Structure of human O-GlcNAc transferase and its complex with a peptide substrate. Nature (2011).

2011:

X-ray Crystallography: Low Resolution Refinement (DEN)

NMR: Fragment Approach

SBGrid - General Methodology

- New workflows to support experimental structure determination
- Reuse existing applications
- GlideinWMS for job submission (OSG-gfactory at UCSD)
- Web portal for WSMR workflow submission
- Extensive data analysis
- Typically 2000-7000 job slots
- Used 6 million hours in last 12 months
- 4th largest non-HEP VO

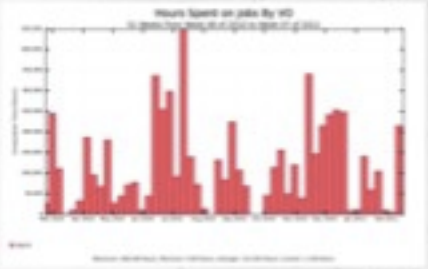

SBGrid 2011/12 Outlook

- 2-3 new workflows (NMR, X-ray crystallography)
- New resources
- Integrated user environment with LDAP, Kerberos, KX.509 and possibly SLCS
- Improved end-user data management facilities (exploring Globus Online)