"LHC analysis in the times of ubiquitous Grids"

Or should this be

LHC Analysis in the era of the Ubiquitous Grid

Outline
- Random stuff you already know about the grid
- Stuff you already know about analysis

One or Many Grids?

A nuisance at present

Success would mean we don't care

- How many webs are there?
  - There's only one HTTP (well almost)

Ubiquitous grids implies

- An agreed set of standard protocols (XML's) for inter-enterprise collaborative computing
  - authentication, authorisation, accounting/logging
  - resource description, data access, workflow
- Security and trust
- Driven in large part by HEP developments/experience

Production Data Processing

Processing 10 PB is hard + time consuming

- Can't all do what we want when we want it
  - Resource limits will tend to impose a traditional model
    - Raw Data -> Processed data -> DST -> mini-DST -> ntuple
  - No requirement to do it in a historic way
    - Do it all once (asap)
    - "pointer" not copy … brings new power
    - Copies will be taken ….. Need to keep track of provenance
- We can build the system
  - Virtual, location independent, standard
- But coordinating resources of this scale is not trivial
  - **This is not a new problem…**

To Solve: the HENP "Data Problem"

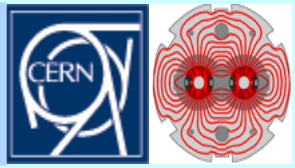

While the proposed future computing and data handling facilities are large by present-day standards,
They will not support FREE access, transport or reconstruction for more than a Minute portion of the data.

- ➔ Need for effective global strategies to handle and prioritise requests, based on both policies and marginal utility
- ➔ Strategies must be studied and prototyped, to ensure *Viability:* acceptable turnaround times; efficient resource utilization

Problems to be Explored; How To

- ➔ Meet the demands of hundreds of users who need transparent access to local and remote data, in disk caches and tape stores
- ➔ Prioritise thousands of requests from local and remote communities
- ➔ Ensure that the system is dimensioned "optimally", for the aggregate demand

Why do we need a Grid

Grid is next step of virtualisation and other trendy things

Virtualisation:

Grid is the next step…

- Read those bytes from that address
- Read the Nth entry from that file
- Return the Nth muon momentum from that file
- Read the momentum of the Nth entry from last weeks muon file
- Find last weeks muons and get me the momentum of the Nth muon
- Go forth and find me the most recent QA'd momentum of the Nth muon from last week

even Microsoft…

- AWT deprecated
- Programme to .NET
- Location independence
- Easy to move forwards -> distributed/internet computing

Service Based Architecture

Focus on what you want to do, not how (or where) you want to do it
Service based architecture based on loosely couple services.
Why ?

- Can always build a tightly coupled system on top
- But then have to deal with the loosely couple nature (failures)
- Whole system will (should) not come down with first crash

Virtualisation will allow us to scale/adapt the system for the long term
SBA will make scaleable and robustness more likely

What does the Grid give us?

Standard ways to use distributed resources

- Install the standard "stuff"
- Set the appropriate flags
- Central production can run the whole thing

Use systems even when there's no postdoc !

Stuff?

- AAA/resource access
  - Non expt specific
  - Non HEP specific
- "standard" sysadmin to do it

Need to be "opportunistic"

- Do the results agree ?
- Shouldn't need to ask the question

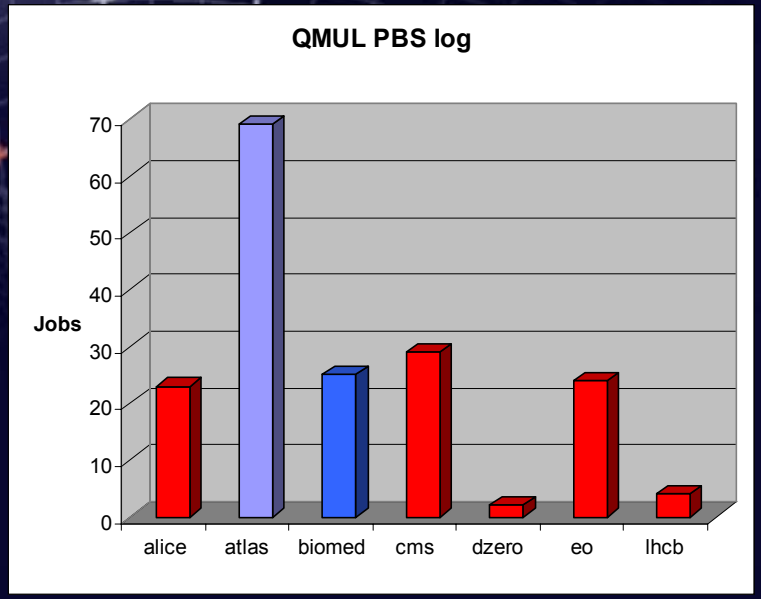

climate*prediction*.net

- Have created extremely powerful and distributed climate modelling facility by running model simulation on home computers (cf. SETI@home)
- Launch ensemble of coupled simulations of 1950-2000 and compare with observations.
- Run on to 2050 under a range of natural and anthropogenic forcing scenarios.
- Investigates sensitivity of climate system to increasing CO_2 with range of parameter values
- Have collaborated with other universities and industry to build system

- "Screensaver" requires
- 10 CPU days on a 1.4GHz P4 (min is 800MHz machine)
- >128MB memory
- 600MB disk space allocated to the programme

NATURAL ENVIRONMENT RESEARCH COUNCIL

climate*prediction*.net results

- Already largest climate model ensemble ever (by factor of >200)
- >45,000 users, >15,000 complete model runs, >1,000,000 model years in ~3 months (this is equivalent to 1.5 Earth Simulators)

- Global outreach (participants in all 7 continents, inc. Antarctica!)
- Generated much interest in schools (coolkidsforacoolclimate.com)

e-Science

NATURAL ENVIRONMENT RESEARCH COUNCIL

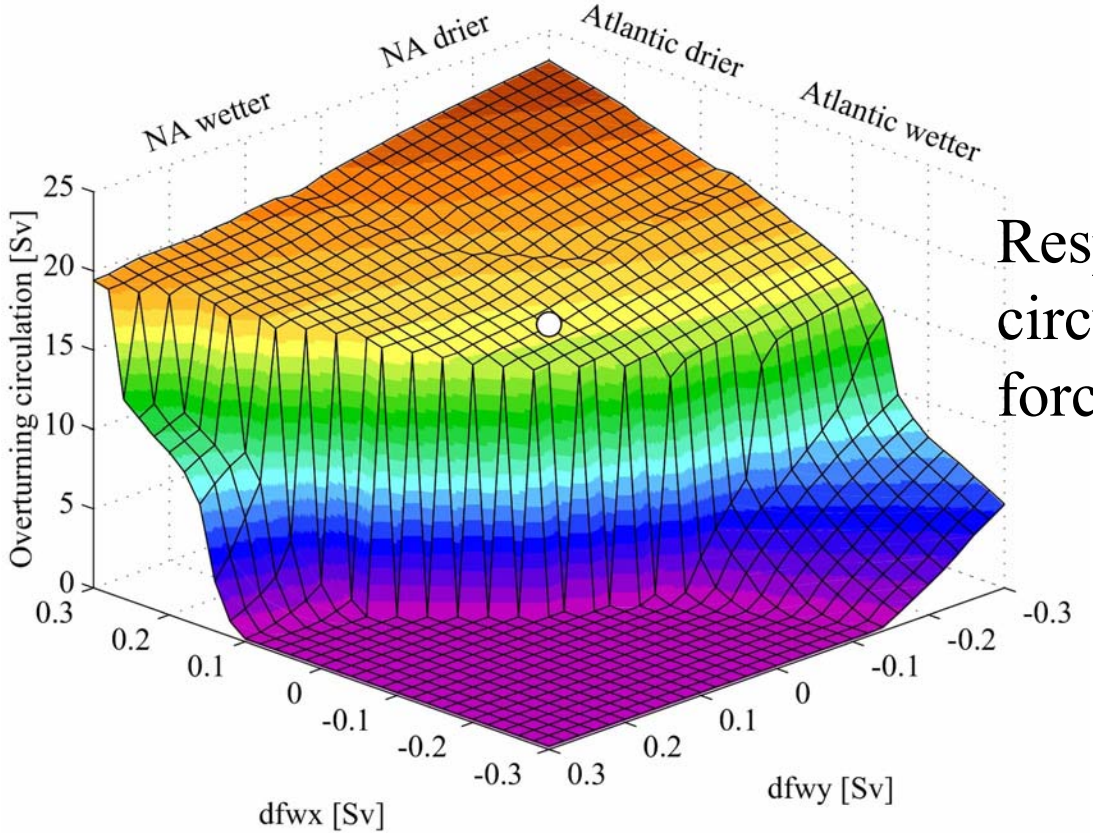

Response of Atlantic circulation to freshwater forcing

LHC Analysis…

Experience of large collaborations (LEP, B-factories, Collider)

- Only a small fraction of physicists know (in detail) how the data gets processed
- Even fewer care (except when it impacts their physics)

Physicists stand on each others shoulders

"If I have seen further, it is by standing on the shoulders of giants." Isaac Newton

"Mathematicians stand on each other's shoulders." - Gauss

"Mathematicians stand on each other's shoulders while computer scientists stand on each other's toes." - Richard Hamming

"It has been said that physicists stand on one another's shoulders. If this is the case, then programmers stand on one another's toes, and software engineers dig each other's graves." - Unknown

- Gi'me all the events with high P_t J/Ψ events, with a good D^0, and 5 identified kaons (and I only want the good stuff) …
- And I want it on my lap-top

The grid? What about the grid?

BaBar

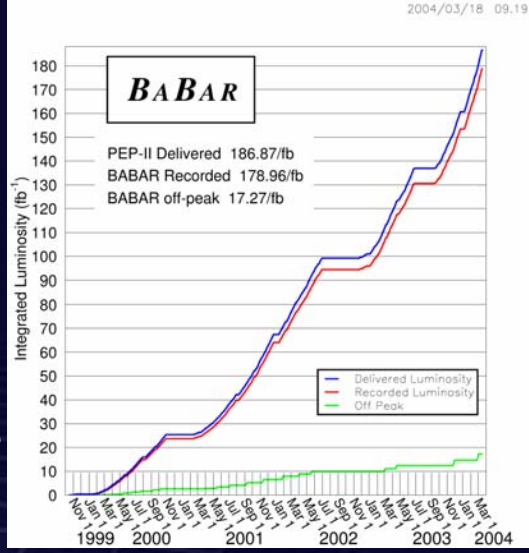

Most analysis uses "ntuples" (rootuples)

- Analysis organised in physics process based groups
  - Groups define the "tags" and composites produced
  - An "expert" produces the "ntuples" for the group
  - Group activity focused around the ntuple

Lots of reasons for the focus on ntuples

- Big reason – analysing 1PB data is hard !
  - Don't want to do it very often, leverage the work of each other
  - Book-keeping

Other experts defined the kaons and the pions

- Do everything once
- BaBar started off pretty well (k/μ/π/p…)
- Slow to standardise/include composites
- No real attempt/manpower to develop integrated analysis

Big step forward is "merging" of "ntuples" and "dst"

CDF/D0

CDF

- Analysis in physics oriented groups
- Most analysis on pinned staged data on Central Analysis Facility
  - Distributed analysis means setting up new CAF's
- Users write out their own samples/ntuple to "take home"
- High statistics analyses (e.g. B) already on analysis group ntuples

D0

- Commitment to distributed production computing – Peter Maettig
  - SAM
  - DST Reprocessing requires no database access
- Analysis based on DSTs + "Thumbnail"s (root files/trees ?)
  - Disk resident – but users don't know (SAM)
  - Confidence in contents from initial emphasis on "object-id"
- Organise groups around physics processes

What I want to do

Just before I leave for the airport

- Skim 200GB sHiggs sample to my laptop
  - Recipe -> tool(s) -> a button

On the (transatlantic) flight [or during the summary talk]

- Try out that new idea
- (Complain that the satellite link [or wireless] isn't working)

When I get off the plane

- Refit tracks for my 10 selected events (needs some raw data)
- Join the analysis group access grid meeting and show plots
- Call up and display the full events for the gold plated CP violating sHiggs events I found (tool)
- Show I used the latest calibration and processing (tool)
- Add my sHiggs events to the "DST"

Book e-ticket to Stockholm 

Other Grid stuff

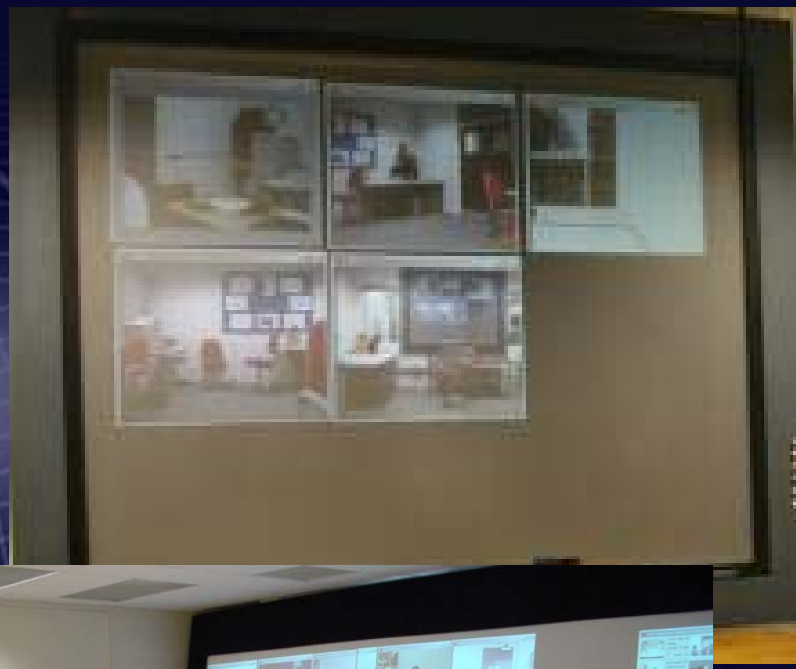

Access grid
(+ desk top integration)

New developments

Fall into 2 types (not black and white)

- The "invisible" stuff
  - Users never need to know, it just happens
    - New network router, new standard c library, bigger tapes in the robot, new grid s/w
- The new opportunity stuff
  - "I know this looks harder (ie is different) but it will allow us to cross correlate everything when we write the tools"
    - Click on the histogram and see the calibration constants
    - Return all the raw data for my se
  - YOU NEED TO WRITE THE TOO
    - If the tools don't do the job, the
    - Someone will write "new" ones
    - NEED USERS + TRAINING

ARDA/Data Challenges
Extremely important here
- only 3 years to go

Grids for 2007

- Standard and coherent access
  - Transparent virtual data storage, resource access, single sign on…
- Analysis will only be really distributed if the tools exist
  - Calibrate the data, Track the data, Verify the data, use the data

Successful LHC Collaborations

Will

- Develop systems with one eye on the future
- Coordinate their data processing
- Not simply assume that it will sort itself out
- Allow users to work where ever
- Facilitate lightweight "take home" event samples
- Keep track of calibration/processing/provenance
- Provide simple access to the power of the underlying computing system
- Plan for extensible analysis systems
- Get organised

Will not

- Ignore reality/priorities/human nature
- Stop users taking home event samples
- Expect all collaborators to understand/care about the computing system
- Undervalue their computing (personnel) ☺

Conclusion
LHC Analysis in 2010

All Data processed in real time

- Full detector calibration, track fitting and tagging
- Events passing standard criteria also have D^0, D^*, B, H reconstructed

PLUS

- The system is person-centric ("me"-centric)
- The things I need to do are easy to initiate
  - Defined by analysis group organisation
- I can synchronise my laptop anywhere ~~at the press of a button~~

M

Some tuples are HUGE (several TB)

- Actually several copies stored over many different locations