ATLAS DC2

Status
LCG-GDB meeting
15 June 2004

Gilbert Poulard
for ATLAS DC; Grid and Operations teams

DC2: goals

- The goal includes:
  - Full use of Geant4; POOL; LCG applications
  - Pile-up and digitization in Athena
  - Deployment of the complete Event Data Model and the Detector Description
  - Simulation of full ATLAS and 2004 combined Testbeam
  - Test the calibration and alignment procedures
  - **Use widely the GRID middleware and tools**
  - **Large scale physics analysis**
  - **Computing model studies (document end 2004)**
  - **Run as much as possible of the production on Grids**
  - **Demonstrate use of multiple grids**

Task Flow for DC2 data

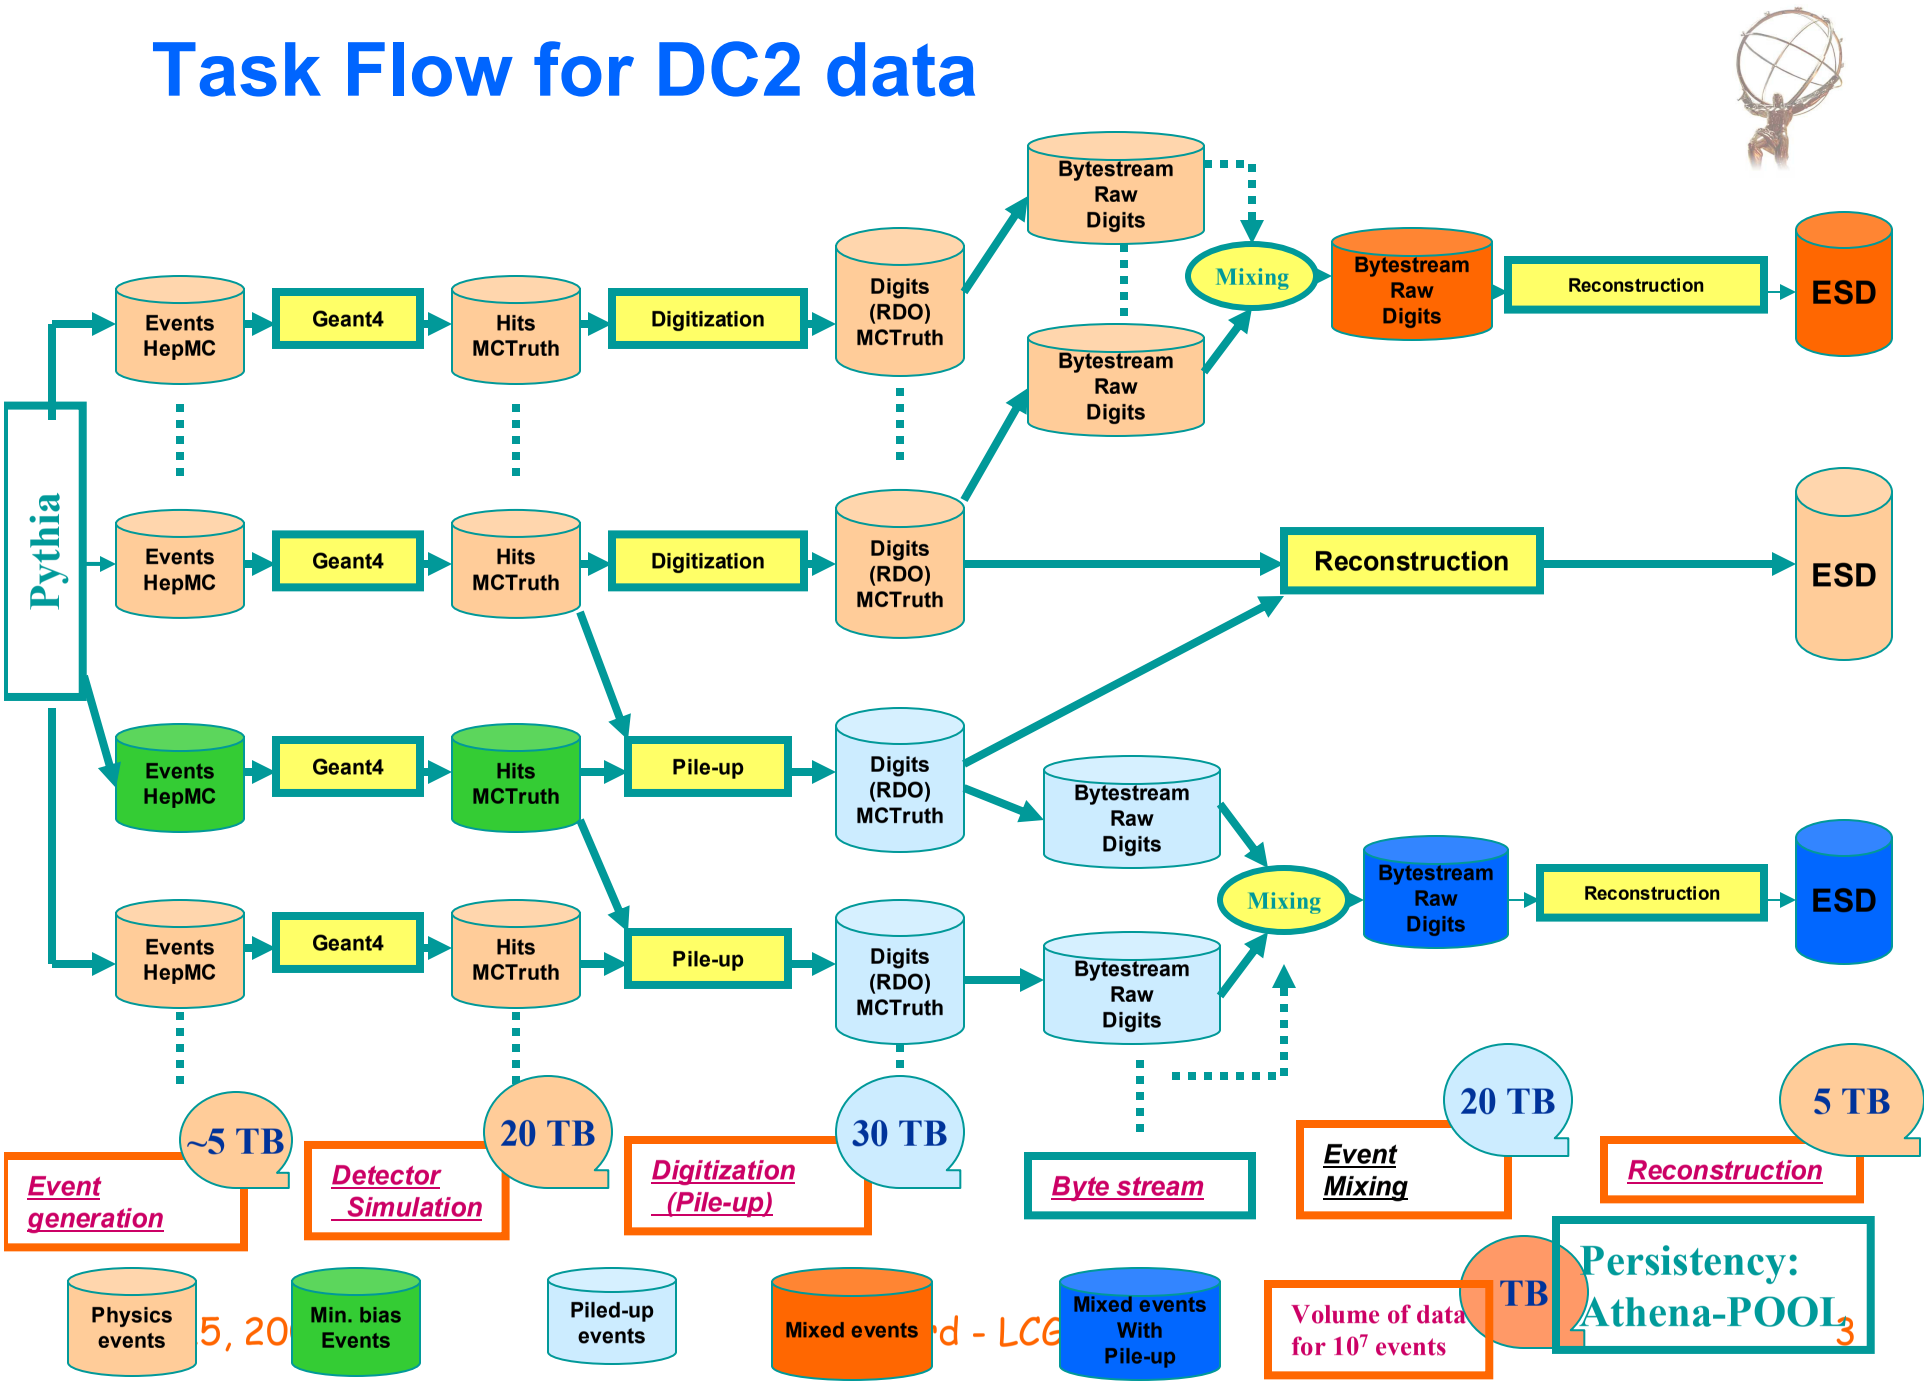

DC2 operation

- Consider DC2 as a three-part operation:
  - part I: production of simulated data (June-July 2004)
    - needs Geant4, digitization and pile-up in Athena, POOL persistency
    - "minimal" reconstruction just to validate simulation suite
    - will run **"preferably" on "Grid"**
  - part II: test of Tier-0 operation (August 2004)
    - needs full reconstruction software following RTF report design, definition of AODs and TAGs
    - (calibration/alignment and) reconstruction will run on Tier-0 prototype as if data were coming from the online system (at 10% of the rate)
    - output (ESD+AOD) will be **distributed to Tier-1s in real time** for analysis
  - part III: test of distributed analysis on the Grid (Sept.-Oct. 2004)
    - access to event and non-event data from anywhere in the world both in organized and chaotic ways
  - in parallel: run distributed reconstruction on simulated data (from RODs)

DC2: where are we?

- DC2 Phase I
  - Part 1: event generation
    - Release 8.0.1 (end April) for Pythia generation (70% of data)
      - tested, validated, distributed
      - test production started 2 weeks ago
        - a few minor bugs fixed since
      - real production started this week with current release 8.0.4
  - Part 2: Geant4 simulation
    - Release 8.0.2 (mid May) reverted to Geant4 6.0 (with MS from 5.2)
      - tested, validated, distributed: MAJOR BUG FOUND!
        - TileCal rotated by 180 degrees around vertical axis
    - Release 8.0.4 (early June) was supposed to be used
      - New problem in endcap TRT just discovered
      - Release 8.0.5 being prepared
  - Part 3: pile-up and digitization
    - Release 8.0.5
      - currently under test (performance optimization)
      - production later in June

DC2: where are we?

- DC2 Phase I
  - Part 4: Data transfer
    - RDOs (or ByteStream "raw data") to be sent to CERN
      - ~30 TB in 4 weeks
  - Part 5: Event mixing
    - Read many input files
    - Mix the physics channels (in "adhoc" proportion)
    - If done from RDOs create BysteStream data (raw data)
    - Release 8.0.5

DC2: where are we?

- DC2 Phase II
  - Reconstruction
    - Reconstruction from ByteStream
      - Creates ESD and AOD
  - In parallel distributes ESD and AOD to Tier-1s in real time
  - Release 9.0.x

- DC2 Phase III
  - Calibration and Reprocessing
  - Test of Distributed Analysis on the Grid

Production scenario

	Input		Output		Comments
Event generation	none		Generated events	< 2 GB files	
G4 simulation	Generated Events	"part of" < 2 GB files	Hits + MCTruth	< 2 GB files	Job duration limited to 24h! ~ 2000 jobs/day ~ 500 GB/day ~ 5 MB/s
Detector response	Hits + MCTruth (Generated events)	1 file	Digits +MCTruth	RDO (or BS)	No MCTruth if BS
Pile-up	Hits "signal" +MCTruth Hits "min.b"	1 file Several 10 files	Digits +MCTruth	RDO (or BS)	~ 2000 jobs/day Input: ~ 10 GB/job ~ 30 TB/day ~ 350 MB/s
Byte-stream	"pile-up" data RDO	1 (or few) files	BS		Still some work
Events mixing	RDO or BS	Several files	BS		"
Reconstruction	RDO or BS		ESD		
AOD production	ESD		AOD		Streaming?

DC2 resources (based on release 8.0.3)

Process	No. of events	Time duration	CPU power	Volume of data	At CERN	Off site	
		months	kSI2k	TB	TB	TB	
Simulation	10^7	1	2000*	20	4	16	Phase I (June -July)
RDO	10^7	1	200	20	4	16	
Pile-up (*) Digitization	10^7	1	600	35 (?)	35 (?)	~30(?)	
Event mixing & Byte-stream	10^7	1	(small)	20	20	0	
Total Phase I	10^7	1	2800	~100	~60	~60	
Reconstruction Tier-0	10^7	0.5	600	5	5	10 ?	Phase II (mid-August)
Reconstruction Tier-1	10^7	2	600	5	0	5	
Total	10^7			100	63 (39?)	71	

DC2 resources (based on release 8.0.3)

Process	Time duration	Input			Output		
		Event size	I/O rate		Event size	I/O rate	
	months	MB	MB/s	GB/day	MB	MB/s	GB/day
Simulation	1				1.9	9	800
Pile-up Digitization	1	1.9 (65 m.b.)	350	30000	3.5(?)	17	1500
Reconstruction	0.5						

Process	Time duration	Volume of data (Bytestream)	Bandwith
	Weeks	TB	MB/s
Data Transfer	4	30	~12

ATLAS Production system

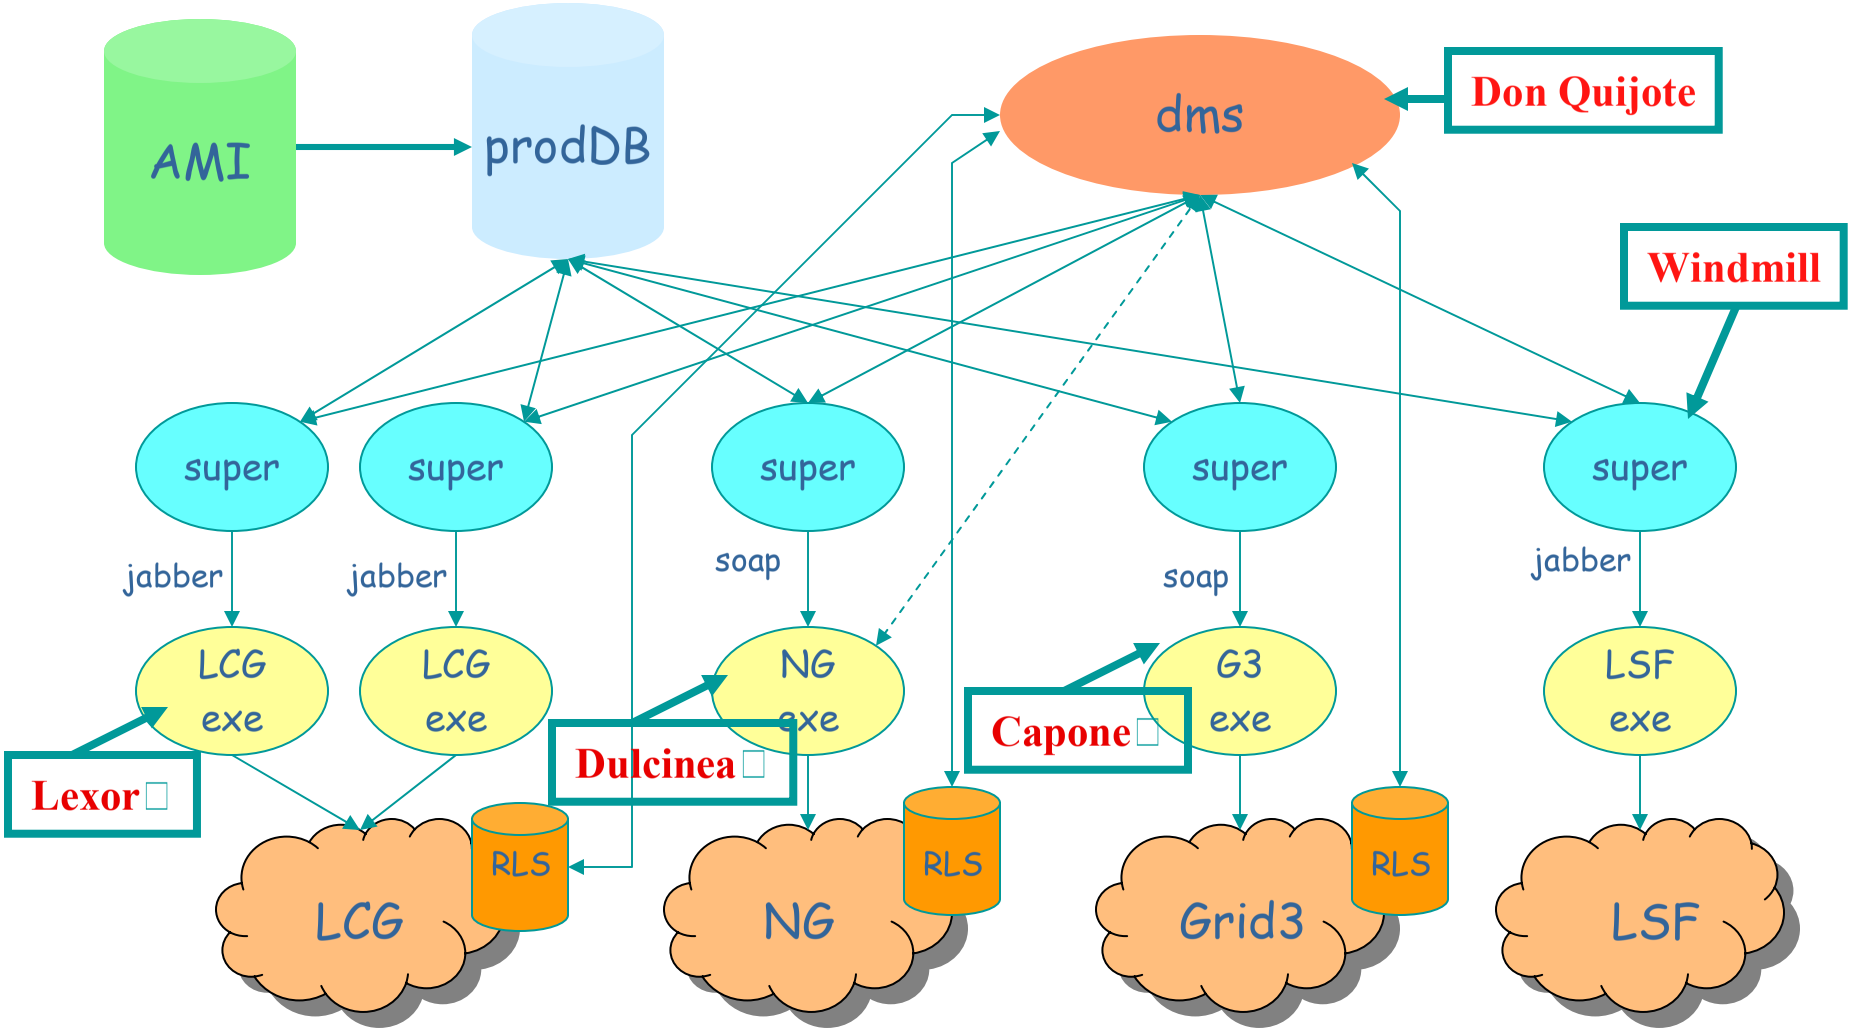

ATLAS Production System

- Components are there
  - Supervisor: Windmill
  - Executors: Capone (Grid3); Dulcinea (NG); Lexor (LCG); "Legacy systems"
  - Data Management System (DMS): Donquijote
  - Bookkeeping: AMI
- Needs
  - More testing
  - QA-QC and Robustness
  - To integrate all transformations (pile-up; event mixing; …)

Grid in DC2

- We are ready to use the 3 grid flavours
  - LCG-2, Grid3+ and NorduGrid
  - All 3 look "stable" (adiabatic evolution)
  - Since the production is just starting it's difficult to say more
- Newcomers:
  - Interface LCG to Grid Canada
    - UVic, NRC and Alberta accept LCG jobs via TRIUMF interface CE
      - ATLAS releases installed
      - Tests in DC2
      - Interests for this scheme London Tier-2

Monitoring & Accounting

- At a very early stage in DC2
  - Needs more discussion within ATLAS
    - Metrics to be defined
    - Development of a coherent approach
  - Current efforts:
    - Job monitoring "around" the production database
      - Publish on the web, in real time, relevant data concerning the running of DC-2 and event production
      - SQL queries are submitted to the Prod DB hosted at CERN
      - Result is HTML formatted and web published
      - A first basic tool is already available as a prototype
    - On LCG: effort to verify the status of the Grid
      - Ø two main tasks: **site monitoring** and **job monitoring**
      - Ø **based** on **GridICE**, a tool deeply integrated with the current production Grid middleware
    - On Grid3: MonaLisa
    - On NG: NG monitoring

Savannah in DC2

Still at the level of a proposal:

assign bugs and requests directly to the responsible people for a service or tool

Categories:

- **ATLAS software**
- **ATLAS release package installation** (Alessandro De Salvo)
- **ATLAS production system**
- **AMI** (Solveig Albrand, Jerome Fulachier)
- **DC2 organization** (Armin Nairz, Nektarios Benekos)
- **GRID problems**
- **General** (a generic container for all other stuff)

ATLAS production

- Will be done as much as possible on Grid (All?)
  - Few production managers
  - Data stored on Tier1's
  - "Expression of Interests" to distribute the data in an "efficient" way
  - Keep the possibility to use "standard" batch facilities but using the same production system
  - Will use several "catalogs"; DMS will take care of them
  - Current plan:
    - 20% Grid3
    - 20% NorduGrid
    - 60% LCG-2 (10 "Tier1s")
    - To be adapted based on experience

Current Grid3 Status (3/1/04)
(http://www.ivdgl.org/grid2003)

- 28 sites, multi-VO
- shared resources
- ~2000 CPUs
- dynamic – roll in/out

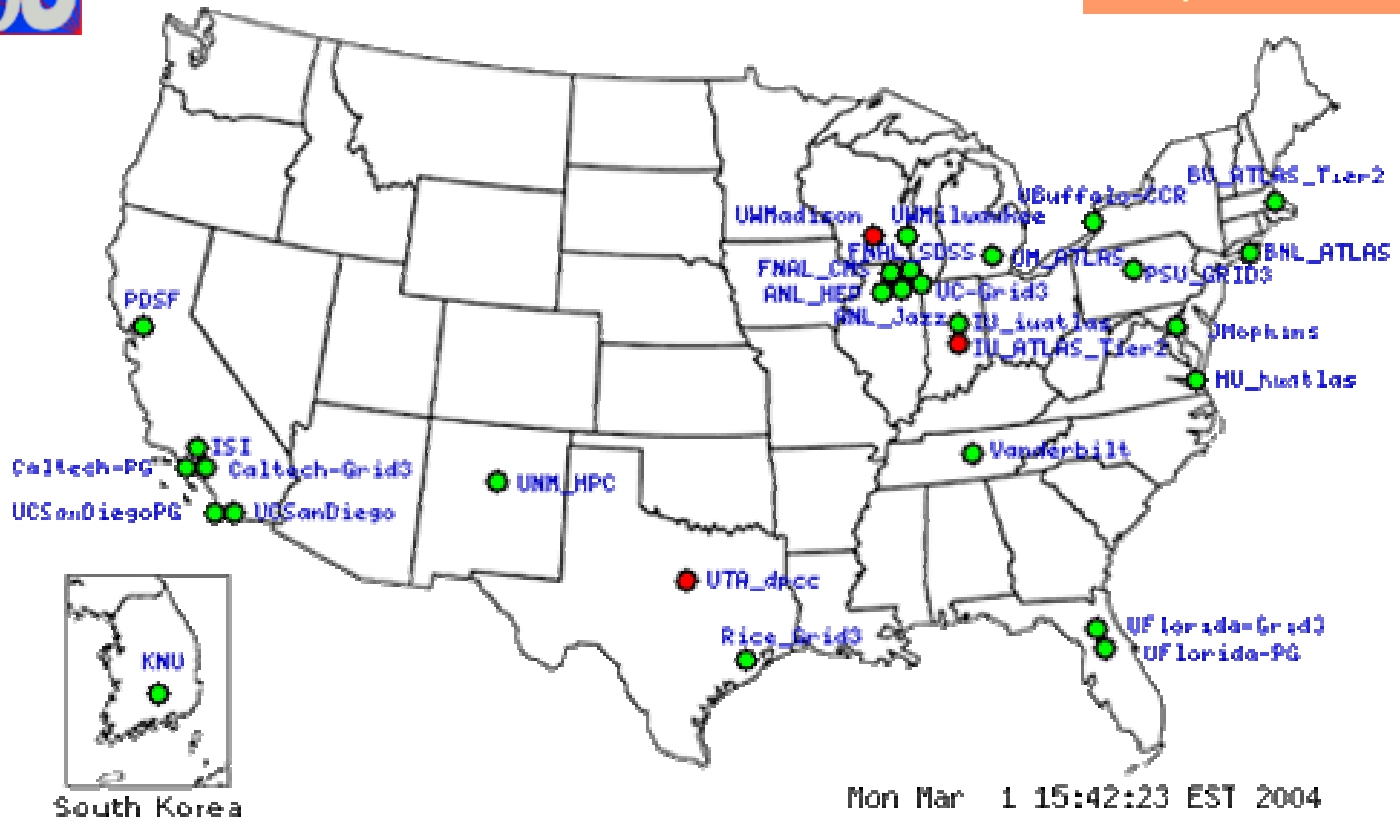

NorduGrid Resources: details

- NorduGrid middleware is deployed in:
  - Sweden (15 sites)
  - Denmark (10 sites)
  - Norway (3 sites)
  - Finland (3 sites)
  - Slovakia (1 site)
  - Estonia (1 site)
- Sites to join before/during DC2 (preliminary):
  - Norway (1-2 sites)
  - Russia (1-2 sites)
  - Estonia (1-2 sites)
  - Sweden (1-2 sites)
  - Finland (1 site)
  - Germany (1 site)
- Many of the resources will be available for ATLAS DC2 via the NorduGrid middleware
  - Nordic countries will coordinate their shares
  - For others, ATLAS representatives will negotiate the usage

LCG-2 today (May 14)

"Tiers" in DC2 (rough estimate)

Country	"Tier-1"	Sites	Grid	kSI2k
Australia			NG	12
Austria			LCG	7
Canada	TRIUMF	7	LCG	331
CERN	CERN	1	LCG	700
China				30
Czech Republic			LCG	25
France	CCIN2P3	1	LCG	~ 140
Germany	GridKa	3	LCG	90
Greece			LCG	10
Israel		2	LCG	23
Italy	CNAF	5	LCG	200
Japan	Tokyo	1	LCG	127
Netherlands	NIKHEF	1	LCG	75
NorduGrid	NG	~30	NG	380
Poland			LCG	80
Russia			LCG	~ 70
Slovakia			LCG	
Slovenia			NG	
Spain	PIC	4	LCG	50
Switzerland			LCG	18
Taiwan	ASTW	1	LCG	78
UK	RAL	8	LCG	~ 1000
US	BNL	28	Grid3/LCG	~ 1000
Total				~ 4500

Tiers in DC2

- Tier-1s will have to
  - Host simulated data produced by them or coming from Tier-2; plus ESD (& AOD) coming from Tier-0
  - Run reconstruction in parallel to Tier-0 exercise (~2 months)
    - This will include links to MCTruth
    - Produce and host ESD and AOD
  - Provide access to the ATLAS V.O. members

- Tier-2s
  - Run simulation (and other components if they wish to)
  - Copy (replicate) their data to Tier-1

After DC2: "continuous production"

- We have requests for
  - Single particles simulation (a lot)!
    - To be defined
      - The detector geometry (which layout?)
      - The luminosity if pile-up is required
      - Others? (eg. Cavern background)
  - Physics samples for the Physics workshop studies (June 2005)
    - DC2 uses ATLAS "Final Layout"
    - It is intended to move to "Initial Layout"
    - Assuming that the geometry description is ready by beginning of August we can foresee an intensive MC production starting ~mid-September
    - Initial thoughts:
      - ~ 50 Million Physics events; that means ~10 Million events per month from mid-September to February 2005
  - Production could be done either by the production team or by the Physics groups
    - The production system should be able to support both

Summary

- Major efforts on the past few months
  - Redesign of the ATLAS Event Data Model and Detector Description
  - Integration of the LCG components (G4; POOL; …)
  - Introduction of the Production System
    - Interfaced with 3 Grid flavours (and legacy systems)
- Delays in all activities have affected the schedule of DC2
  - Note that Combined Test Beam is ATLAS 1st priority
- DC2 is in front of us
  - Resources seem to be there
  - Production system should help
  - It's a challenge!