ATLAS Computing Model: Resources and Network Requirements

Roger Jones

20th January 2005

SARA

ATLAS is not one experiment

Extra Dimensions

Higgs

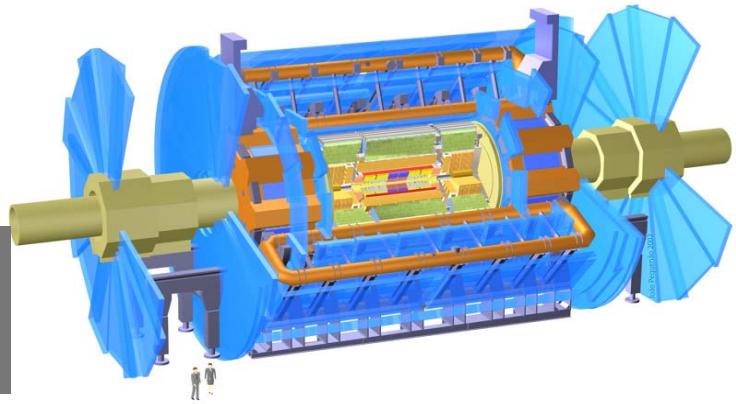

Heavy Ion Physics

SUSY

Electroweak

B physics

QCD

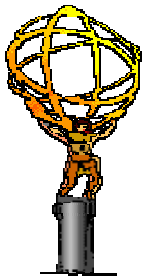

Computing Resources

- **Computing Model fairly well evolved**
  - Under external review
- **There are (and will remain for some time) many unknowns**
  - Calibration and alignment strategy is still evolving
  - Physics data access patterns MAY start to be exercised this Spring
    - Unlikely to know the real patterns until 2007/2008!
  - Still uncertainties on the event sizes
  - If there is a problem with resources, e.g. disk, the model will have to change
- **Lesson from the previous round of experiments at CERN (LEP, 1989-2000)**
  - Reviews in 1988 underestimated the computing requirements by an order of magnitude!

The System

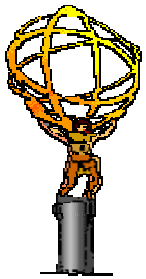

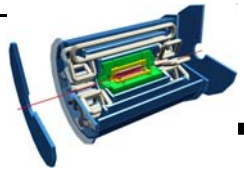

~Pb/sec

Event Builder

~100 Gb/sec

PC (2004) = ~1 kSpecInt2k

Event Filter
~7.5MSI2k

~3 Gb/sec raw

- Some data for calibration and monitoring to institutes
- Calibrations flow back

Tier 0

T0 ~5MSI2k

- **~5 PB/year**
- **No simulation**

~ 75MB/s/T1 raw for ATLAS

Tier 1

US Regional Centre

Dutch Regional Centre

French Regional Centre

UK Regional Centre (RAL)

- **~2MSI2k/T1**
- **~2 PB/year/T1**

≥622Mb/s links

10 Tier-1s

reprocess

house simulation

Group Analysis

Tier 2

Northern Tier ~200kSI2k

Tier2 Centre ~200kSI2k

Centre ~200kSI2k

Centre ~200kSI2k

- **~200 TB/year/T2**

≥622Mb/s links

Lancaster ~0.25TIPS

pool

anchester

Sheffield

Physics data cache

100 - 1000 Mb/s links

Desktop

Workstations

Each of ~30 Tier 2s have ~20 physicists (range) working on one or more channels

Each Tier 2 should have the full AOD, TAG & relevant Physics Group summary data

Tier 2 do bulk of simulation

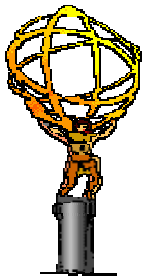

Processing

- **Tier-0:**
  - First pass processing on express/calibration physics stream
  - 24-48 hours later, process full physics data stream with reasonable calibrations

  → These imply large data movement from T0 to T1s
- **Tier-1:**
  - Reprocess 1-2 months after arrival with better calibrations
  - Reprocess all resident RAW at year end with improved calibration and software

  → These imply large data movement from T1 to T1 and T1 to T2

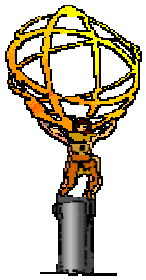

Analysis model

Analysis model broken into two components

- Scheduled central production of augmented AOD, tuples & TAG collections from ESD

→ Derived files moved to other T1s and to T2s

- Chaotic user analysis of augmented AOD streams, tuples, new selections etc and individual user simulation and CPU-bound tasks matching the official MC production

→ Modest job traffic between T2s

2008 data Tier 0 requirements

	Disk (TB)	Shelf.Tape (TB)
Raw	0	3040
ESD	0	1000
Buffer	127	0
Calibration	240	168
Total	**354**	**4208**

1 copy offsite to T1s

2 copies offsite to T1s

1 copy offsite

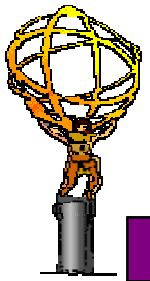

CERN Analysis Facility

	Disk (TB)	Tape (TB)
Raw	241	0
ESD (current)	229	0
ESD (previous)	0	18
AOD (current)	257	0
AOD (previous)	0	4
TAG (current)	3	0
TAG (previous)	0	2
MC ESD (current)	286	0
MC ESD (previous)	0	4
MC AOD (current)	57	0
MC AOD (previous)	0	40
MC Tag (current)	0.6	0
MC Tag (previous)	0	0.4
Calibration	240	168
User Data	303	212
Total	**1615**	**448**

Real data traffic internal to CERN

All MC from offsite T2

To/from T1/T2s

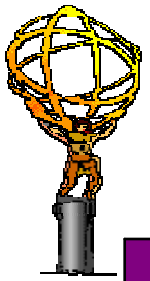

2008 data Combined Tier 1 Requirements

	Disk (TB)	Tape (TB)
Raw	430	3040
ESD (current)	2570	900
ESD (previous)	1290	900
AOD	2830	360
TAG	30	0
Calibration	2400	0
MC RAW	0	800
MC ESD (current)	570	200
MC ESD (previous)	290	200
AOD Simulation	630	80
Tag Simulation	10	0
Group User Data	1260	0
Total	**12300**	**6480**

From CERN

20 copies to CERN, other T1s and T2s

From CERN

All MC data from T2s

Copies to and from other T1s and to T2s

2008 Data Combined Tier-2

	Disk (TB)
Raw	43.4
General ESD (curr.)	385.7
General ESD (prev..)	0.0
AOD	2571.4
TAG	77.1
RAW Sim	0.0
ESD Sim (curr.)	171.4
ESD Sim (prev.)	0.0
AOD Sim	571.4
Tag Sim	17.1
User Group	1257.1
User Data	1815.3
Total	**6910.1**

From T1 or CERN

From T1

From Ts – probably local

From T1
Local

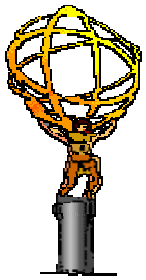

Important points:

- **Storage of Simulation**
  - Assumed to be at T1s
  - Need partnerships to plan networking
  - Must have fail-over to other sites
- **Simulation fraction is an important tunable parameter in T2 numbers!**
  - Increased simulation increases the (Tier 1) storage requirement

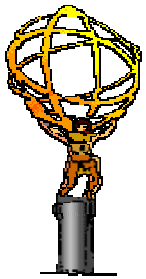

Networking – CERN to T1s

- EF⇔T0 maximum 320MB/s (450MB/s with headroom)
- The ATLAS Tier 1s will be: CCIN2P3-Lyon, RAL, NIKHEF, F2K-Karlsruhe, ASCC, BNL, PIC, NorduGrid, CNAF, TRIUMF
  - They vary in size!
- Traffic from T0 to average Tier-1 is ~75MB/s raw
- With proposed headroom, efficiency and recovery factors for service challenges this is 3.5Gbps/sec
- Most ATLAS T1s are shared with other experiments, so aggregate bandwidth & contention larger

Data tyoe	Inbound from CERN (MB/s)	Outbound to CERN (MB/s)
RAW	30.4	
ESD Versions	20	1.41
AOD versions	18	0.28
TAG Versions	0.18	0
Group Derived Physics Datasets		0.81
Total CERN/AverageTier-1	68.58	2.51

Networking – T1 to T1

- **Significant traffic of ESD and AOD from reprocessing between T1s**
  - 52MB/sec raw
  - ~2.5Gbps after usual factors
- ➔ **Tier-2 to Tier-1 networking requirements far more uncertain**
  - ➔ without job traffic ~18.5MB/s for 'average' T2
  - ➔ ~900Mbps required?
- ➔ **Limited but unquantified need for Tier-2 to CERN connections for calibration**

Source	Inbound from other Tier-1s (MB/s)	Outbound to other Tier-1s (MB/s)
ESD Versions	12	14
AOD versions	20.8	2.08
TAG Versions	0.21	0.02
Group DPD	0	4
Total Tier-1 to Tier-1	33	19

Networking and Tier—2s

- Tier-2 to Tier-1 networking requirements far more uncertain
  - without job traffic ~15.6MB/s for 'average' T2
  - ~750Mbps required?
- Limited but unquantified need for Tier-2 to CERN connections for calibration

Source	Inbound (MB/s)	Outbound (MB/s)
RAW	0.9	2.7
ESD Versions	0.8	1.1
AOD versions	5.5	0.2
TAG Versions	0.2	0.0
Group DPD	4.2	0.0
Total for average Tier-2	11.6	4.0

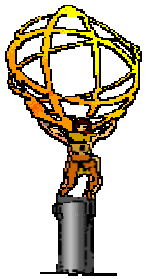

Reality check

- **The requirement is not completely matched by the current pledges**
  - **The table presents Atlas's estimate of the lower bound on the Tier-1 resources available**
  - **Recent indications suggest the CPU shortage will be met, but the storage remains a problem**

Snapshot of 2008 Tier-1 status

Summary Tier1s	Split 2008	ATLAS
CPU (kSI2K)	Offered	18300
	Required	26500
	Balance	-31%
Disk (Tbytes)	Offered	5100
	Required	15500
	Balance	-67%
Tape (Pbytes)	Offered	9.9
	Required	10.1
	Balance	-2%