

# Tier 3 Task Force

Gustaaf Brooijmans, Columbia,  
Jim Cochran, Iowa State University,  
Michael Ernst, Brookhaven National Laboratory,  
Amir Farbin, University of Texas at Arlington,  
Bruce Melado, University of Wisconsin,  
Mark Neubauer, University of Illinois,  
Flera Rizatdinova, Oklahoma State University,  
Paul Tipton, Yale University, Gordon Watts, University of Washington,  
Chip Brock, Michigan State University  
adopted Kyle Cranmer, NYU, Akira Shibata, and Marco Mambelli

# charge, 1

## 1. Use Cases.

Typical workflows for physicists analyzing ATLAS data from their home institutions should be enumerated. This needs to be inclusive, but not in excruciating detail.

It should be defined from within the ATLAS computing/analysis models, the existing sets of T2 centers, and their expected evolutions.

If there are particular requirements in early running, related to detector commissioning and/or special low-luminosity considerations, this should be noted.

If particular ATLAS institutions have subsystem responsibilities not covered by the existing T1/2 deployment, this should be noted.

Is the previous whitepaper relevant?

## 2. Characterizations of generic T3 configurations.

Some T3's may be very significant because of special infrastructure availabilities and some T3's maybe relatively modest.

Is there only 1 kind of T3 center, or are their possible functional distinctions which might characterize roles for some T3's that might not be necessary for others?

Description of "classes" of T3 centers, if relevant, should be made.

Support needs and suggestions for possible support models should be considered.

# charge, 2

## 3. Funding.

This is not part of the US ATLAS Operations budget, so funding must come out of the institutes through core funding or local sources. We would like to make it easier for institutes to secure funding for ATLAS computing--this can only happen if it fits in the DOE and NSF budgets ( precedent: the amount of funding groups got for computing equipment in Tevatron experiments) and it must fit in the overall US ATLAS model.

For the latter, we have to make the case that the existing T1/2 centers are not enough.

Perhaps a recommendation can be justified for an estimated \$ amount needed for a viable Tier 3 cluster -- something like  $X + n*Y$  \$'s where  $n$  = number of active physicists.

# activities

A few meetings & lots of small group discussions

Fixated on the roles of T2s and the degree to which they may or may not be saturated.

- Some DØ and CDF concern, based on experiences.
- A simulation underway...with early results suggesting that MC production will be a significant load for T2s - assuming 10% of total dataset full-simulation
- lots of questions, though: who/where/when: D1->D2? D2->D3? fast MC's?

# Three Surveys:

Other ATLAS national, T3 plans

Other LHC experiments' T3 plans

A survey of current capabilities with mixed results

- a number of institutions did not respond to 3 requests:

*Hampton, Howard, Penn, NIU, OSU, BU, UM, U-Mass, Yale, New Mexico, Albany*

# Two essential agreements:

Partly based on a personal bias about Academic HEP:

- We should seek to promote active, visible, *tangible*, on-campus computing

Anticipate 4 classes of T3 centers, based on distinct criteria

- a sense that some T3's could be brought into service for some production tasks if necessary

A support model

Table 12: Comparison of the 1997 Computing Plans for the DØ experiment looked at from 2006 [1].

	1997 projections	2006 actual
Peak (average) data rate (Hz)	50 (20)	100(35)
Events collected	600M/year	1500M/year
Raw Data Size (kB/event)	250	250
Reconstructed Data size(kB/event)	100	80
User format (kB/event)	1	40
Tape Storage	280 TB/year	1.6 PB on tape
Tape reads/writes (weekly)		30 TB/7TB
Analysis/cache disk	7 TB/year	220 TB
Reconstruction time (GHz-s/event)	2.0	50
User analysis times (GHz-s/event)	?	1
User analysis weekly reads	?	3B events
Primary reconstruction farm size (THz)	0.6	2.4 THz
Central analysis farm size (GHz)	0.6	2.2 THz
Remote resources (GHz)	?	~ 2.5THz

Table 13: The DØ experiment “tiered” computing clusters and the closest ATLAS analogs.

	reconstruction farm	CAB cluster	CLuED0 cluster
DØ	400 nodes	1252 nodes, 2 clusters	500 nodes
	dedicated batch	~ 5000 job slots batch	interactive & batch
ATLAS	~ T1?	~ T2's?	~ T3's?

# Flexible and Nimble



# “Kinds of T3’s”

“T3gs”: a center with full **g**rid **s**ervices

- likely a significant center with infrastructure in place
- production-capable

“T3g”: a cluster with **g**rid connectivity

- maybe a desktop cluster
- enough storage to support a full D3PD analysis year’s worth of data, 5ishTB?
- DQ2 client, but maybe sharing a DQ2 site services relationship with T2

*in order to support data subscriptions*

“T3w”: individual, personal **w**orkstations

- RootTuple analyses, grid submission

“T3af”: within the confines of a an **a**nalysis **f**acility

- like BAF or an ANL AF

# support model

An ~FTE, “peripatetic” support person

- U.S. ATLAS supported
- will travel
- will maintain up to date documentation
- will sponsor workshops for site administrators

We are characterizing “shopping lists”:

- of hardware examples, servers required, software required for each T3 flavor

# where are we

## writing

draft among just a handful of people at the moment  
trying to be comprehensive about numbers, current policies, etc.

## calculating

use-case examples of T2 loads and T3gs, T3g, and T3w capabilities

**Surely this will end!**