Progress on User Analysis on Tier-2s

Ian Fisk

September 25, 2008

User Analysis on Tier-2s

The CAF at CERN is a very valuable resources

- It will have access to the really prompt reconstruction and calibration samples the quickest
- It's useful for low latency analysis and some other very high profile tasks involving data promptly
- Unfortunately it's small

There are many Tier-2s and many of them are big

- A nominal Tier-2 is 1MSI2k and 200TB of disk
  - Half is devoted to simulation the other half to analysis
  - We already have individual Tier-2s with at least half the batch slots for analysis that the CAF has
- With the exception of the limited number of tasks that can only be done on the CAF users are going to find more resources at Tier-2s

What are we doing to make them more efficient to use

In CMS Jobs go to Data

How is the Storage managed?

Storage at Tier-2 centers is broken into 6 pieces

- Transient and unmanaged to more persistent and centrally managed

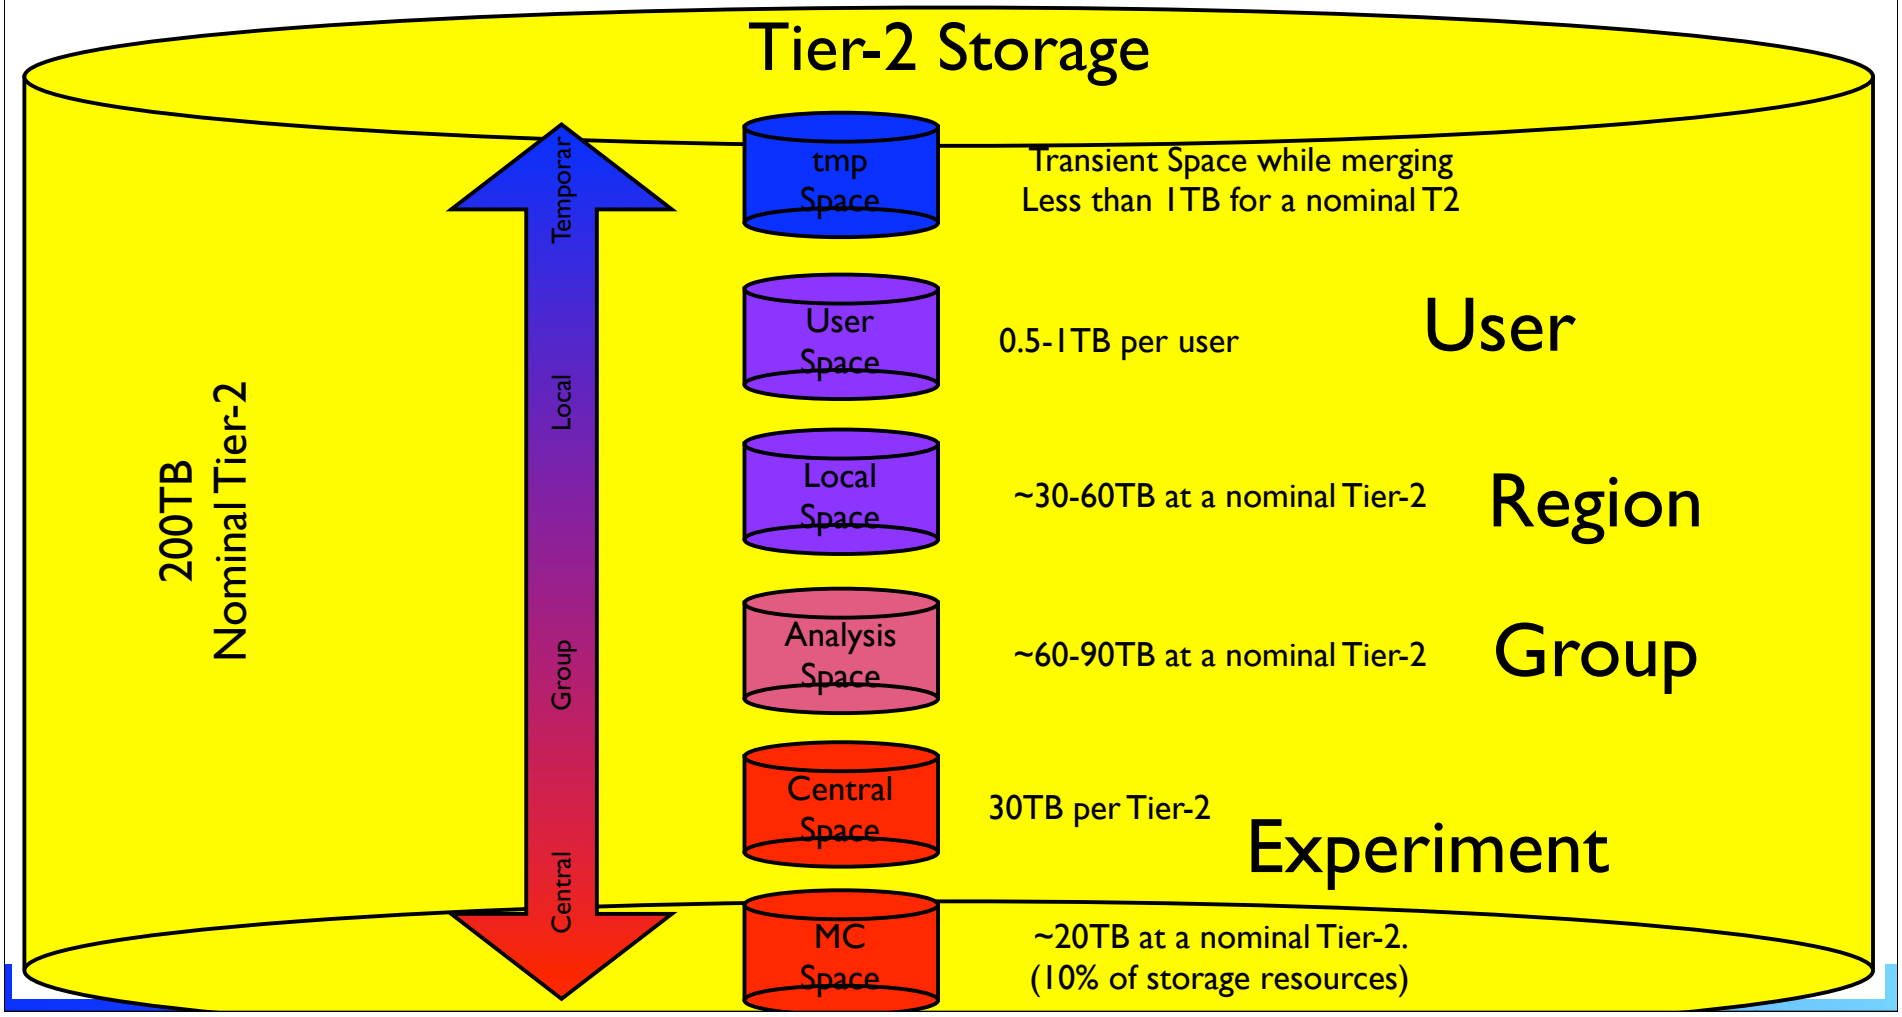

Who Controls the Storage?

All numbers are for a nominal Tier-2

Central Space 30TB

- Intended for RECO samples of Primary Datasets.
  - In 2008 we had expected to be able to store 2 copies of MC and data sample using the identified T2 space

Physics Group Space 60-90TB

- Assigned to 1-3 physics groups. Space allocated by physics data manager. The site data manager still approves the request, but only to ensure the group is below quota

Local Storage Space 30TB-60TB

- Controlled by the local storage manager. Intended to benefit the geographically associated community

User Space 0.5-1TB per person in the geographically associated community

- controlled by individuals

Motivations for User Space at Tier-2s

We need to give users a predictable space to write and Grid accessible storage

- People need places to write to that are not Castor at CERN
  - CERN Castor writes to tape
    - This uses tape resources, which we need for real data
    - User Files are often small, which is lowering the average file size on tape and the efficiency of the tape system impacts the ability to access data
    - Need to support users on disk resources at Tier-2s

The concept of keeping it on the local Tier-2 was to divide the problem

- At a nominal Tier-2 40 users are supported
- User Space is assigned at the Tier-2 geographically associated with the institution
- Keeping it by institution and local users provides us with better chance for efficient support and management.

CRAB will have the ability to stage data to /store/user and stage from it

Tier-2 Analysis Workflow

Ideal Tier-2 Workflow

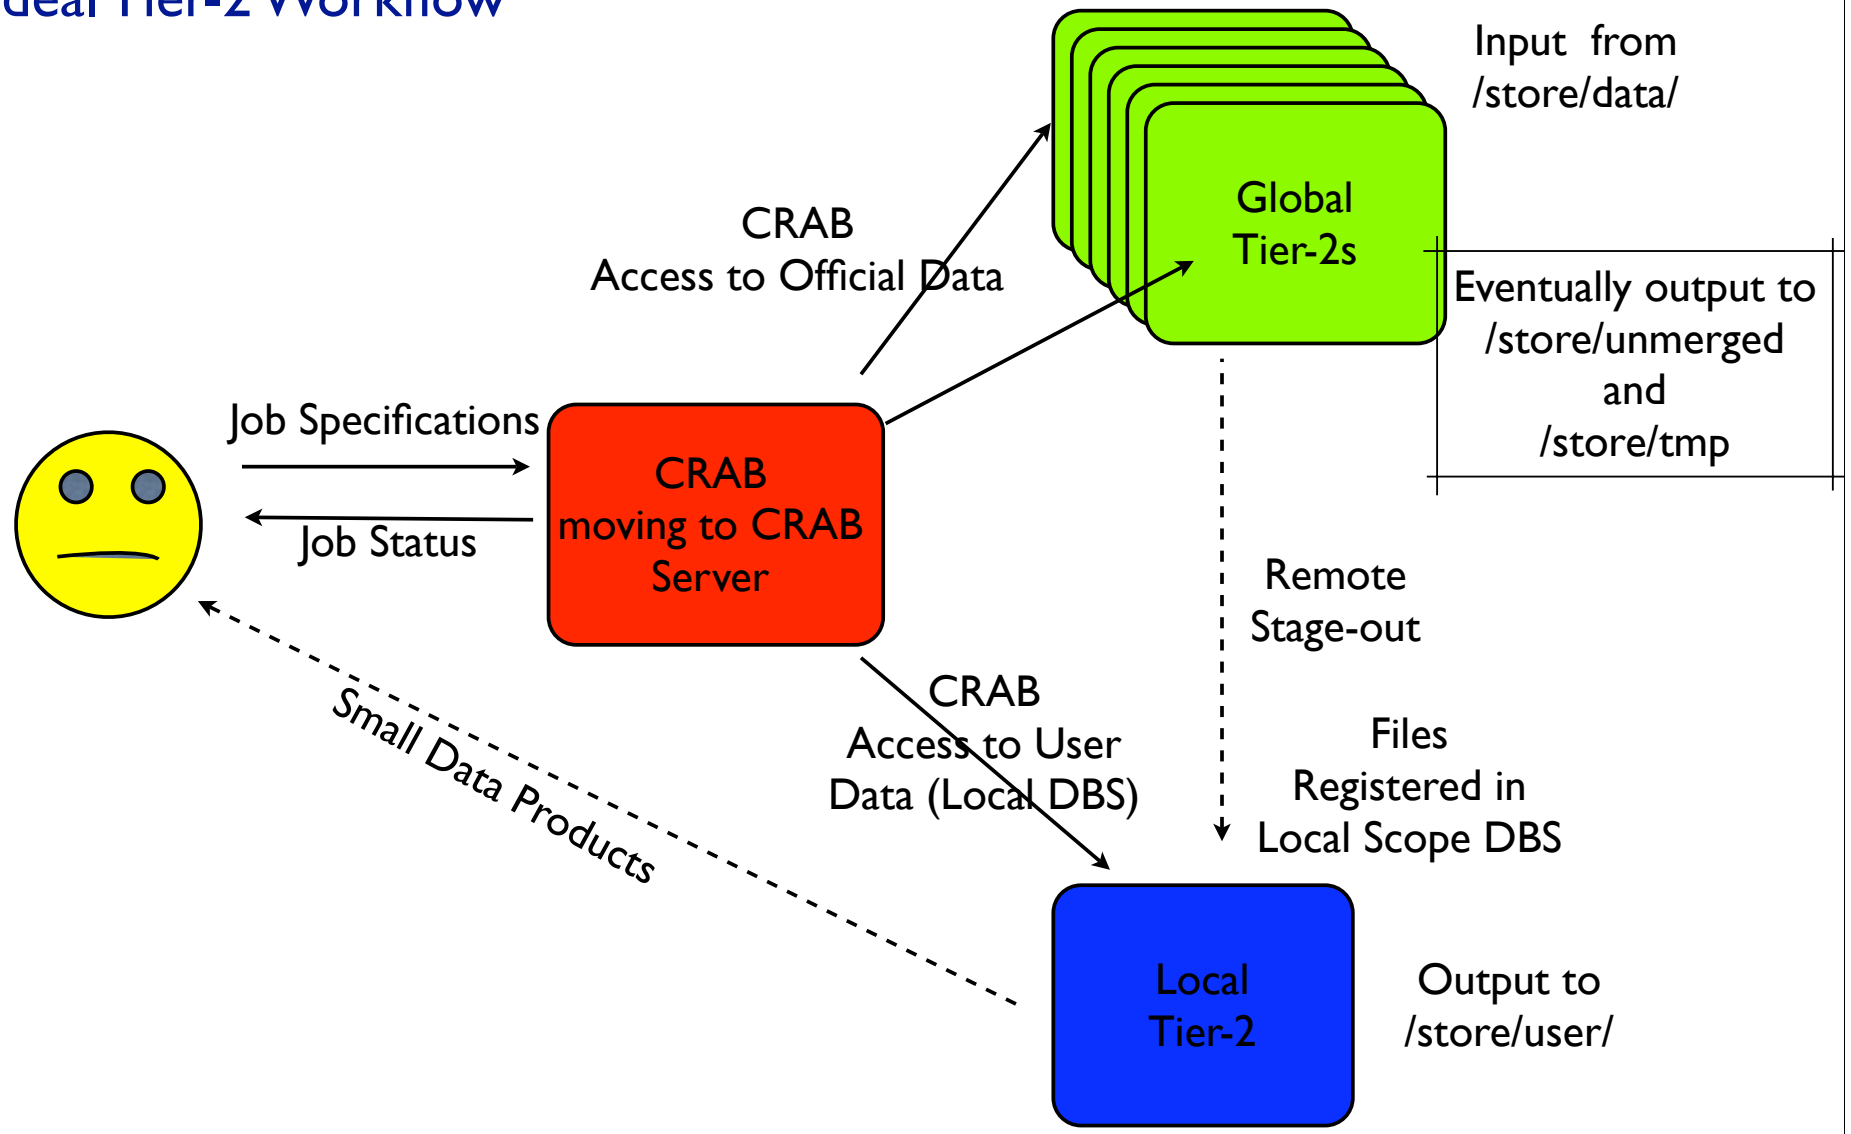

Why are we switching to CRAB Server?

Technical Improvements -

- CRAB server removes the limitations on the size of the input sandbox
- CRAB server allows the user to disconnect once the workflow is upload (Don't have to wait for many jobs to be submitted)
- Can provide better resubmission functionality

Support

- CRAB server gives CMS central points were jobs are submitted through. Support people will have access to logs and configurations which should improve our user support

Tier-2 Associations to Analysis Groups

After what seems like a very long time, we have the mappings of site to analysis groups

- The process and the mapping should be revisited one we have some operational experience with high energy data
- These mappings were chosen to optimize the types of physics expected early
  - Concentration on commissioning work (DPG, POG)
    - Each country was expected to have at least 40% DPG and POG
    - Concentration on standard model analysis
    - Leave resources for search groups, but the emphasis is on commissioning work

Thomas Kress and the Tier-2 Liaisons (Giuseppe Bagliesi and Ken Bloom) as well as the Physics leadership and the Computing Resource Board were instrumental in getting us this far

Country Mappings

	T2_AT	T2_BE	T2_BR	T2_DE	T2_CH	T2_CN	T2_EE	T2_ES	T2_FI	T2_FR	T2_IT	T2_KR	T2_PT	T2_RU	T2_UK	T2_US
FWD phys				1												1
QCD				1						1						2
Higgs								1		1	1					1
EWK								1		1	1				1	1
SUSY	1			1							1				1	1
Top		1		1				1		1						1
Exotica										1				1	1	1
B Physics					1	1			1							1
Heavy Ions														1		0
egamma										1	1				1	2
Jets/MissET				1					1			1		1		1
Muons								1			1			1		2
B-Tagging	1		1							1						1
Tracker				1						1	1					1
Tau / Pflow							1			1	1					1
Trigger DPG								1							1	1
Reserve																2
Unallocated		?											1			1
Current Resources	0	1	1	3	0	0	1	5	2	8	5	1	0	1	4	15
Fall Resources (*)	2	1	1	6	1	1	1	5	2	9	7	1	1	4	5	21
POGs/DPGs	1	0	1	3	0	0	1	2	1	4	4	1	0	2	2	10
POG fraction	0.5	0	1	0.5	0	0	1	0.4	0.5	0.44	0.6	1		0.5	0.4	0.48

By Group (1/2)

Forward - T2_US_Wisconsin (T2_DE_DESY at start of data taking)

QCD - T2_DE_DESY, T2_FR_CCIN2P3, T2_US_Caltech, (T2_US_MIT)

Higgs - T2_ES_IFCA, T2_FR_GRIF, (T2_IT_Roma), T2_US_MIT

EWK - T2_ES_CIEMAT, T2_FR_CCIN2P3, T2_IT_Legnaro, T2_UK_London_Brunel, T2_US_UCSD

SUSY - (T2_AT_Vienna), (T2_DE_RWTH), T2_IT_Bari, T2_UK_London_IC, T2_US_Florida

Top - T2_BE_IIHE, (T2_DE_DESY), T2_ES_IFCA, T2_FR_IPHC, T2_US_UCSD

Exotica - T2_FR_GRIF, (T2_RU), T2_UK_SGrid_RALPP, T2_US_Purdue

By Group (2/2)

B-Physics - (T2_CH_CSCS), (T2_CN_Beijing), T2_FI_HIP, T2_US_MIT

Heavy Ion- T2_RU

E-gamma - T2_FR_GRIF, T2_IT_Roma, T2_UK_London_IC, T2_US_Caltech, (T2_US_UCSD)

Jets/MET HCAL - T2_DE_DESY, T2_FI_HIP, T2_KR_KNU, T2_US_Purdue (T2_RU)

Muon - T2_ES_CIEMAT, (T2_IT_Legnaro) (T2_RU), T2_US_Purdue, (T2_US_Florida)

B-Tagging - (T2_AT_Vienna), T2_BR_UERJ, T2_FR_IPHC, T2_US_Nebraska

Tracker - T2_DE_RWTH, (T2_FR_CCIN2P3), T2_IT_PISA, T2_US_Nebraska

Tau/PFlow - T2_EE_Estonia, T2_FR_CCIN2P3, T2_IT_PISA, T2_US_Florida

Trigger - T2_ES_CIEMAT, (T2_UK_London), T2_US_Wisconsin

Outlook

The Tier-2 association to analysis group took longer than we hoped

- I hope in the next few months we can use the associations with simulation to exercise the system and train people on the task of data management
  - The latency will be lower and the data sets at the Tier-2s will better reflect the needs of the group, if they are controlled by those closest to the work

The transition to /store/user is becoming automated and it should give users a consistent place to store data products

- Large output can be accessed with CRAB
- Small output can be pulled back

The CRAB server should improve functionality and support

- We will be working on a smooth transition through the fall.