

Plan for shutdown and 2009 data-taking

Kors Bos

Plans for 2008/2009

- September + October 2008: cosmic ray data taking (see next slide)
- November + December 2008: preparations re-processing and analysis
- Early 2009: reprocessing 2008 cosmic ray data (reduction to ~20%) and data analysis
- *No official statement from CERN management but we need to make a plan*
- So we assume for 2009:
 - ▣ ~1/3 of the time, starting in January: Functional Tests, Rre-Processing, MC Prod., Analysis Challenges, ...
 - ▣ ~2/3 of the time: cosmic ray and collisions data taking
- *Until we are told differently*

Data Type	June 24 th - October 28 th , 2008			January 1 st - June 24 th , 2008		
	Files	Events	Tot. Size [MB]	Files	Events	Tot. Size [MB]
RAW	693,509	464,362,139	1,104,305,943	87,796	40,059,042	136,324,739
CALRAW	17,028	110,463,488	17,558,049	130	n/a	1,886
ESD	651,109	434,868,635	240,578,525	104,645	56,379,218	14,017,253
ESD_FILTERED	474,716	n/a	7,795,534	70,284	n/a	242,337
CBNT	651,109	434,868,635	404,906,622	104,645	56,379,218	21,286,696
AOD	2,572	79,046,922	1,894,468	n/a	n/a	n/a
TAG	2,572	79,046,922	4,407	n/a	n/a	n/a
TAG_COMM	4,623	325,676,278	27,558	1,914	17,289,650	388
HIST	5,378	n/a	183,594	2278	n/a	77,183
NTUP_PIXELCALIB	354	n/a	21,974	n/a	n/a	n/a
NTUP_MUONCALIB	5,022	n/a	2,285,878	n/a	n/a	n/a

Glossary:

- RAW: physics and debug streams from DAQ
- CALRAW: calibration streams from DAQ
- ESD_FILTERED: ESDs of events with at least one ID track
- TAG_COMM: commissioning TAGs (produced in reconstruction step)
- TAG: physics TAGs (produced in AOD merging step)
- HIST: merged DQM histograms
- NTUP_*CALIB: merged n-tuples for Pixel and Muon detector calibration purposes

RAW event size = 2.4 MB
because of many samplings
nominal = 1.6 MB

Data Type	June 24 th - October 28 th , 2008			January 1 st - June 24 th , 2008		
	Files	Events	Tot. Size [MB]	Files	Events	Tot. Size [MB]
RAW	693,509	464,362,139	1,104,305,943	87,796	40,059,042	136,324,739
CALRAW	17,028	110,463,488	17,558,049	130	n/a	1,886
ESD	651,109	434,868,635	240,578,525	104,645	56,379,218	14,017,253
ESD_FILTERED	474,716	n/a	7,795,534	70,284	n/a	242,337
CBNT	651,109	434,868,635	404,906,622	104,645	56,379,218	21,286,696
AOD	2,572	79,046,922	1,894,468	n/a	n/a	n/a
TAG	2,572	79,046,922	4,407	n/a	n/a	n/a
TAG_COMM	4,623	325,676,278	27,558	1,914	17,289,650	388
HIST	5,378	n/a	183,594	2278	n/a	77,183
NTUP_PIXELCALIB	354	n/a	21,974	n/a	n/a	n/a
NTUP_MUONCALIB	5,022	n/a	2,285,878	n/a	n/a	n/a

Glossary:

- RAW: physics and debug streams from DAQ
- CALRAW: calibration streams from DAQ
- ESD_FILTERED: ESDs of events with at least one ID track
- TAG_COMM: commissioning TAGs (produced in reconstruction step)
- TAG: physics TAGs (produced in AOD merging step)
- HIST: merged DQM histograms
- NTUP_*CALIB: merged n-tuples for Pixel and Muon detector calibration purposes

ESD event size = 0.6 MB
nominal 1 MB

Data Type	June 24 th - October 28 th , 2008			January 1 st - June 24 th , 2008		
	Files	Events	Tot. Size [MB]	Files	Events	Tot. Size [MB]
RAW	693,509	464,362,139	1,104,305,943	87,796	40,059,042	136,324,739
CALRAW	17,028	110,463,488	17,558,049	130	n/a	1,886
ESD	651,109	434,868,635	240,578,525	104,645	56,379,218	14,017,253
ESD_FILTERED	474,716	n/a	7,795,534	70,284	n/a	242,337
CBNT	651,109	434,868,635	404,906,622	104,645	56,379,218	21,286,696
AOD	2,572	79,046,922	1,894,468	n/a	n/a	n/a
TAG	2,572	79,046,922	4,407	n/a	n/a	n/a
TAG_COMM	4,623	325,676,278	27,558	1,914	17,289,650	388
HIST	5,378	n/a	183,594	2278	n/a	77,183
NTUP_PIXELCALIB	354	n/a	21,974	n/a	n/a	n/a
NTUP_MUONCALIB	5,022	n/a	2,285,878	n/a	n/a	n/a

Glossary:

- RAW: physics and debug streams from DAQ
- CALRAW: calibration streams from DAQ
- ESD_FILTERED: ESDs of events with at least one ID track
- TAG_COMM: commissioning TAGs (produced in reconstruction step)
- TAG: physics TAGs (produced in AOD merging step)
- HIST: merged DQM histograms
- NTUP_*CALIB: merged n-tuples for Pixel and Muon detector calibration purposes

CALRAW and CBNT
only at the Tier-0

Data Type	June 24 th - October 28 th , 2008			January 1 st - June 24 th , 2008		
	Files	Events	Tot. Size [MB]	Files	Events	Tot. Size [MB]
RAW	693,509	464,362,139	1,104,305,943	87,796	40,059,042	136,324,739
CALRAW	17,028	110,463,488	17,558,049	130	n/a	1,886
ESD	651,109	434,868,635	240,578,525	104,645	56,379,218	14,017,253
ESD_FILTERED	474,716	n/a	7,795,534	70,284	n/a	242,337
CBNT	651,109	434,868,635	404,906,622	104,645	56,379,218	21,286,696
AOD	2,572	79,046,922	1,894,468	n/a	n/a	n/a
TAG	2,572	79,046,922	4,407	n/a	n/a	n/a
TAG_COMM	4,623	325,676,278	27,558	1,914	17,289,650	388
HIST	5,378	n/a	183,594	2278	n/a	77,183
NTUP_PIXELCALIB	354	n/a	21,974	n/a	n/a	n/a
NTUP_MUONCALIB	5,022	n/a	2,285,878	n/a	n/a	n/a

Glossary:

- RAW: physics and debug streams from DAQ
- CALRAW: calibration streams from DAQ
- ESD_FILTERED: ESDs of events with at least one ID track
- TAG_COMM: commissioning TAGs (produced in reconstruction step)
- TAG: physics TAGs (produced in AOD merging step)
- HIST: merged DQM histograms
- NTUP_*CALIB: merged n-tuples for Pixel and Muon detector calibration purposes

AOD very small for cosmics
TAGs always small, but many

2008 Cosmic Ray Data Sample

- 460M events RAW and ESD at CERN
- and distributed to T1's to tape and disk
 - first: RAW to tape and ESD to disk
 - after half September: RAW to tape and disk and ESD to disk
- A Total of 1.35 PB to disk and 1.1 PB to tape in Tier-1's
- in 3.5 months → ~50 HZ
- nominal 200 HZ, but sometimes lower because of big event size
- 110M events CALRAW to CAF for calibration

- This may be a good guidance of what we will do in 2009 before collisions

2008 Cosmic Ray Data Reduction

- reduction to 20% of original volume
- not 100% sure that Reduced RAW will be saved
- but certainly Reduced ESD, as a matter of fact these are DPD's
- will still be 200 TB of Redu-RAW and 50 TB of Redu-ESD
- we would like to keep this on disk in 2009 for further analysis
- most RAW on disk, limited pre-staging needed

Cosmic Ray Data Re-Processing

- suitable runs and datasets will be defined for re-processing
- data is truly distributed, so all Tier-1's are needed
- we can at best replace 1 Tier-1 by CERN
- Processing times in Tier-1's may vary quite a bit
 - not similar CPU power at the Tier-1's
 - not optimized data distribution
- Conditions DB access was a problem, solved now (?)
- first tests planned in November/December

Re-Processing Bandwidth

- if we assume a typical 10% T1 has 1000 cores for re-processing
- the reconstruction of an event takes typically 10 sec on an average CPU (maybe a little less for these cosmics)
- an event is 2.4 MB (maybe less for collisions data)
- so we need a bandwidth of **240 MB/sec per 1000 CPU's**

- If we were to pre-stage those data from tape
- we would need to bring ~ 1 TB/hour on-line

DDM Functional Tests

- we get a VERY BIG one coming up
- next 2 weeks we accumulate a lot of data in the T1's (on DATADISK)
- in week 3 we then launch T1-T1 transfers all at once

this will test (a new vs of) DDM at an order of magnitude higher rates

- otherwise we shall always run the DDM FT
- to keep testing all sites
- and this is badly needed, unfortunately ☹

MC Production Functional Tests

- running at low rate now: 200 jobs/cloud*week
- need to increase the rate to become more significant
- need to be brokered indiscriminately so all T2 sites get jobs
- need more atomization
- and shifters to look at it
- we should always run them, at lower priority than real MC
- need to see if priorities work

Distributed Analysis Functional Tests

- real user analyses using GANGA ran in Italian (and German) cloud
 - AOD analysis in Milano, Bologna, Roma, Frascati and Napoli
 - 5 people submitting 50*5 jobs from UK, Germany, CERN and Italy
- should be expanded to all clouds
 - but needs people watching in (almost) all Tier-2's
 - should also include the US
- should also be done with Pathena
 - but then also in the EU clouds
- need to identify real user analysis use-cases
- need to be tested with Panda and WMS backend in EU clouds
- need to see if priorities work

Analysis Bandwidth

- if we assume a typical T2 has 400 cores for analysis
- an AOD analysis using GANGA reaches typically 15 Hz
 - on an average CPU, data on NFS mounted disk
- an AOD event is ~ 0.2 MB
- so we need a bandwidth of **1.2 GB/sec (10 Gb) per 400 CPU's**

- don't need to bring those AOD's on-line from tape
- don't know if DPD or NTUP analysis is faster
- have encountered LAN bandwidth problems in Italian T2's
- not using DB access yet

Organizational

- Post Mortem WS @CERN: January 19-23
 - T01/2/3 Jamboree on Thursday January 22
 - Shifters and other tutorials on Friday January 23
- T2/3 Jamborees organized per cloud
 - open to all people
 - central operations will attend
- Analysis workshops
 - combined with the physics coordination workshops (?)
 - need integration between US and EU clouds
 - and NDGF and Asia for that purpose