CMS

Compact Muon Solenoid

Status Report to WLCG-MB
June 10, 2008

- **CSA status and results**
- **CCRC status and results**

Matthias Kasemann

May 08: CSA08 and CCRC08

CSA08 is part of CCRC08-phase-2

- **CSA goals are those of a CMS-specific exercise strongly tight to the commissioning&physics schedule**
  - **Exercise first 3 months of data taking**
    - completely driven by requirements of CMS Commissioning and Physics

- **CCRC goals are a multi-VO computing scale test**
  - **CCRC activities to augmented CSA loads with additional computing tests**
  - **ensure all systems are being stress tested end-to-end to prove metric**

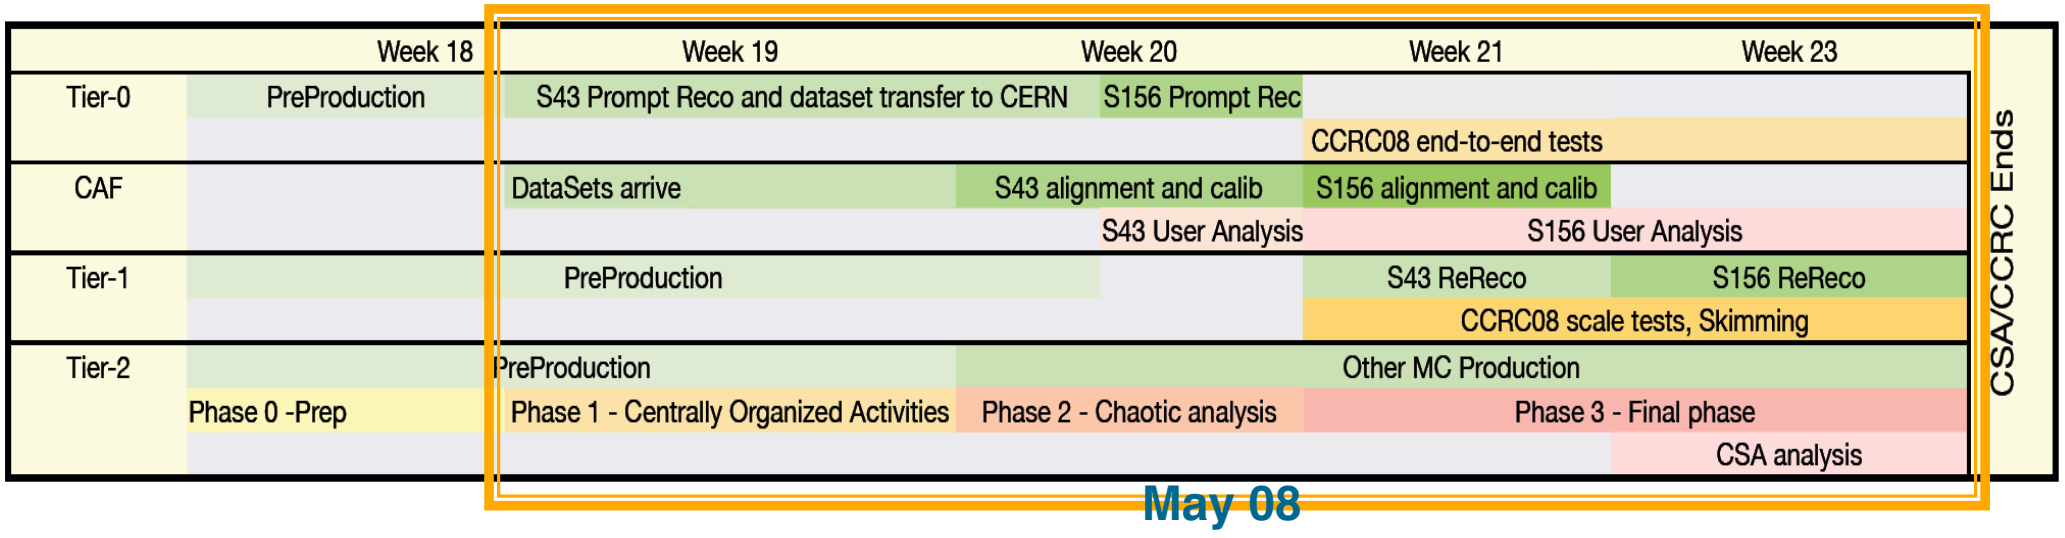

CSA08/CCRC08 Workflows

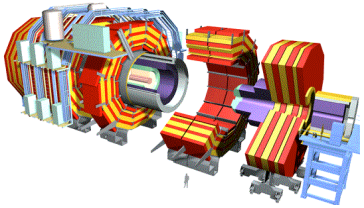

Prompt Reconstruction

TIER-0

CASTOR

CAF

Calibration
Express-Stream
Analysis

600MB/s

50-500MB/s

TIER-1

TIER-1

TIER-1

Re-Reco
Skims

50-500MB/s

~20MB/s

Simulation
Analysis

TIE

TIE

TIE

TIER-2

CSA08: Simulating Data

- **Production of MC samples with CMSSW 2.0**

- **Focus on 2 scenarios:**
  - **"S43": 43 x 43 bunches, L~2x1030, 1 pb-1, 120M events**
  - **"S156": 156 x 156 bunches, L~2x1031, 10 pb-1, 120M events**
    - Samples have big overlap: total production of ~150 M events
  - **Conditions:**
    No pile-up; assume a complete detector; zero suppression
  - **Generator streams**
    - physics samples (minbias, jets, leptons) + technical samples cosmics
    - Assume Storage Manager bandwidth can sustain 300 MB/s
    - Primary Datasets defined by generator

CSA08: Alignment & Calibration

- **Situation similar to real-data startup**
  - **full complexity of almost 20 concurrent alignment & calibration workflows (with interdependencies)**
    - AlCaReco streams with realistic skim filters
    - partially with event content reduction at HLT
  - **all workflows executed on the CAF**
  - **seven teams from DPGs, more than 30 persons directly involved**
  - **regular use of CMS centre (part of work done from remote)**
  - **daily run meetings**
  - **database operations**
  - **constants validation & sign-off**
- **Teams were very successful in running their workflows**
  - **constants for re-reconstruction were provided on time**

Physics

- **Four analysis teams focusing on startup topics**

For results see backup slides

CSA08: Preproduction

- Production of samples started with delay of 11 days
- Events produced at T0, T1 and T2
- Up to 14000 CPU-slots used
- 150M events produced in ~2 weeks

- CSA08 processing started on May 5 (as planed)

- mostly back on schedule 1 week into the schedule
- later completely back on schedule

Preproduction status 13.5.08

S43 + S156	Requested Events	Produced Events	% of requested Events produced	Produced Events transferred to CERN	% of produced Events transferred to CERN	% of requested Events transferred to CERN
Minbias	25,000,000	27,633,750	110.54	27,633,750	100.00	110.54
JetET20	4,000,000	4,140,800	103.52	4,140,800	100.00	103.52
JetET30	4,000,000	4,132,800	103.32	4,132,800	100.00	103.32
JetET50	4,000,000	4,122,400	103.06	4,012,600	97.34	100.32
JetET80	4,000,000	3,441,800	86.05	2,699,000	78.42	67.48
JetET110	4,000,000	4,484,000	112.10	3,983,000	88.83	99.58
MuonPT5	10,000,000	12,063,785	120.64	9,857,517	81.71	98.58
MuonCosmicBON	10,000,000	12,005,000	120.05	12,005,000	100.00	120.05
MuonCosmicBOFF	10,000,000	11,820,000	118.20	10,491,000	88.76	104.91
TkCosmicBON	10,000,000	10,755,882	107.56	10,755,882	100.00	107.56
TkCosmicBOFF	10,000,000	11,887,943	118.88	11,887,943	100.00	118.88
MuonBeamHalo	10,000,000	11,960,000	119.60	11,960,000	100.00	119.60
TkBeamHalo	200,000	213,196	106.60	213,196	100.00	106.60
TrackerLaser	5,000	5,110	102.20	5,110	100.00	102.20
HCALNZS	6,000,000	6,124,000	102.07	6,124,000	100.00	102.07
HCALNZS V2	6,000,000	4,800,000	80.00	2,394,000	49.88	39.90
HCALIsoTracks30 S43	5,000,000	6,762,510	135.25	5,219,768	77.19	104.40
	122,205,000	**136,352,976.00**	**111.58**	**127,515,366.00**	**93.52**	**104.35**

S156	Requested Events	Produced Events	% of requested Events produced	Produced Events transferred to CERN	% of produced Events transferred to CERN	% of requested Events transferred to CERN
JetET150	4,000,000	3,400,000	85.00	1,968,000	57.88	49.20
GammaJets	2,000,000	161,873	8.09	161,873	100.00	8.09
MuonPT11	10,000,000	9,533,402	95.33	7,993,090	83.84	79.93
JPsi	750,000	795,369	106.05	795,369	100.00	106.05
Upsilon	250,000	260,165	104.07	260,165	100.00	104.07
Zee	13,000	16,708	128.52	16,708	100.00	128.52
Zmumu	13,000	16,862	129.71	16,862	100.00	129.71
Wenu	119,000	108,402	91.09	108,402	100.00	91.09
Wmunu	119,000	266,859	224.25	266,859	100.00	224.25
HCALIsoTracks50 S156	5,000,000	4,055,853	81.12	3,713,544	91.56	74.27
	22,264,000	**18,615,493**	**83.61**	**15,300,872**	**82.19**	**68.72**

CSA08: Data Processing steps

- prompt S43 reco at T0
- S43 rereco on T1 level
- prompt S156 reco at T0
- ALCARECO week
- S156 rereco

	Tier-0	Tier-1	CAF	User
Week 19	S43 reco			
Week 20	S156 reco	S43 re-reco 1	S43 AlCa exercise	CAF S43 analysis
Week 21			S156 AlCa exercise	
Week 22		S43/156 re-reco 2/1		
Week 23				Tier-2 analysis

C.Paus: iCSA08 Workflow Schedule

Status:

- Wed. 14.5. Finished S43 reco
- Thu. 15.5. Started S156 reco
- Mon. 19.5. Finished S156 reco
- Mon. 19.5. Started re-reco with new ALCO constants
- Mon. 26.5. Started S156 rereco

Overall: production ~ on schedule

Plans after the challenge

- Cleanup PreCSA08 pre-production from tape at T1 sites and CERN
- Cleanup of all (tape) samples at CERN
- No central cleanup of CAF
- All samples are stored (custodial) at T1
- Provide opportunity to T1 sites to cleanup non-custodial samples

CAF usage in CSA08

- **T0/CAF resources:**
  - **T0: 2330 cores = 5.0 MSI2k**
  - **CAF: 700 cores = 1.4 MSI2k**
  - **Public resources ~ 1.0 MSI2k**

  CMS request: 7.7 MSI2k

  - disk will ramp up over next weeks

CAF mid-May08	648 slots + 48 “cmsexpress”	535 TB
CAF Pledge 08	1200 slots	1600 TB
T0 May08	2328 slots	350 TB
T0 Plegde 2008	3000 slots	400TB

- **Coordinated use cases (ALCA, physics, DPG)**
  - **283 users registered, 81 where active in May**

cmscaf Disc Pool Utilization

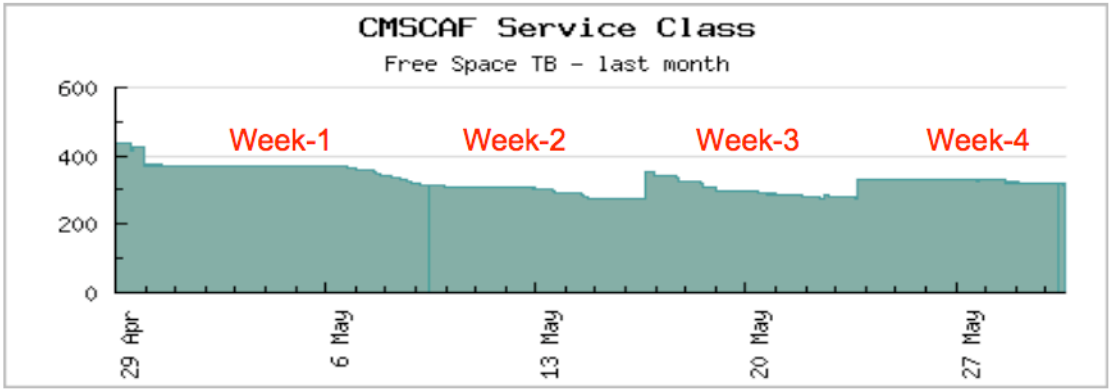

Currently free CAF-disk: 316 TB

Tier-0 Workflows: sustain reco at 150-300 Hz

- great performance for CSA and Cruzet processing/reprocessing
  - reco as such seems fine, job submission and WM systems holds up

- end-to-end test: schedule depends on experts available
  - "re-play" of CRUZET data, through re-packer and prompt reco, and moving data out to Tier-1s
  - assess latency from "taking data" to prompt processing to primary datasets available at Tier-1s

- also still open: test of 300MB/sec from P5 for 1 week
  - can and should be scheduled during June, preferably before WLCG mtg.... June 13/14

Data Transfers: T0—>T1s: fast enough?

	Extra target 850 [MB/s]	Target 600 [MB/s]	Acceptable 425 [MB/s]
SITE	CMS nominal rate	WLCG nominal rate	CCRC/phase-1 extra target rate
ASGC	76	54	32
CERN T1	0	0	0
CNAF	92	65	39
FNAL	238	168	113
FZK	114	80	59
IN2P3	134	95	78
PIC	81	57	42
RAL	116	82	62

T0-T1 transfer targets

Tier-0 to Tier-1 transfer rates during CCRC

day	Total	ASGC	CNAF	FNAL	FZK	IN2P3	PIC	RAL
May 10, 2008	496	152	0	59	55	86	65	79
May 11, 2008	663	163	26	104	0	73	109	188
May 12, 2008	454	95	20	119	0	52	73	95
May 13, 2008	545	151	18	168	14	59	80	55
May 14, 2008	592	85	32	203	71	91	75	35
May 15, 2008	750	181	77	164	120	87	78	43
May 16, 2008	885	240	104	193	77	110	105	56
May 17, 2008	1077	183	123	302	63	170	136	100
May 18, 2008	974	152	119	183	34	143	178	165
May 19, 2008	690	161	104	71	49	81	148	76
May 20, 2008	506	154	51	94	17	77	82	31
May 21, 2008	522	118	88	126	48	50	86	6
May 22, 2008	707	177	160	133	66	22	64	85
May 23, 2008	1017	122	176	319	98	122	66	114
May 24, 2008	773	104	192	58	49	152	116	102
May 25, 2008	956	242	208	0	1	137	142	226
May 26, 2008	947	248	108	0	88	169	164	170
May 27, 2008	752	151	42	11	127	164	112	145
May 28, 2008	811	115	106	171	76	143	91	109
May 29, 2008	728	97	86	168	33	109	100	135
May 30, 2008	309	27	39	107	20	16	36	64

- main metric: 2/4 weeks with >=3 days above nominal rate (600MB/sec)
  - not a test of bandwidth, but of reliability and robustness—> watch for red patches!!!
- Successfully met target for: Total rate, ASGC, CNAF, IN2P3, PIC, RAL
  - FZK, FNAL had difficulties to succeed, to be analyzed
  - many instances of 3-days-red: (all but ASGC and PIC)

- 4 instances of 4-days-red: in data taking times this could be BAD
  - had this been real data, we would have been stuck at the T0 overrunning buffers

Data Transfers T1—T1: replicate re-reco results/AODs

- **Main use case: distribute the AOD after re-reco within 4 days!**
  - chose total sample size of 28.6TB
    - was a size scaled to 3-days at nominal rates
    - is not a unreasonable sample size: ~250M AOD events
- **All sites passed the metric!!**
  - **including sending/receiving the 14TB FNAL sample!**
  - **shown: time between first and last transfer to be ready**
    - e.g. ASGC: transfers took between 0.5 and 2.5 days
- **Want to repeat "other way":**
  - **get re-RECO and AOD from T1 back to CERN-CAF**

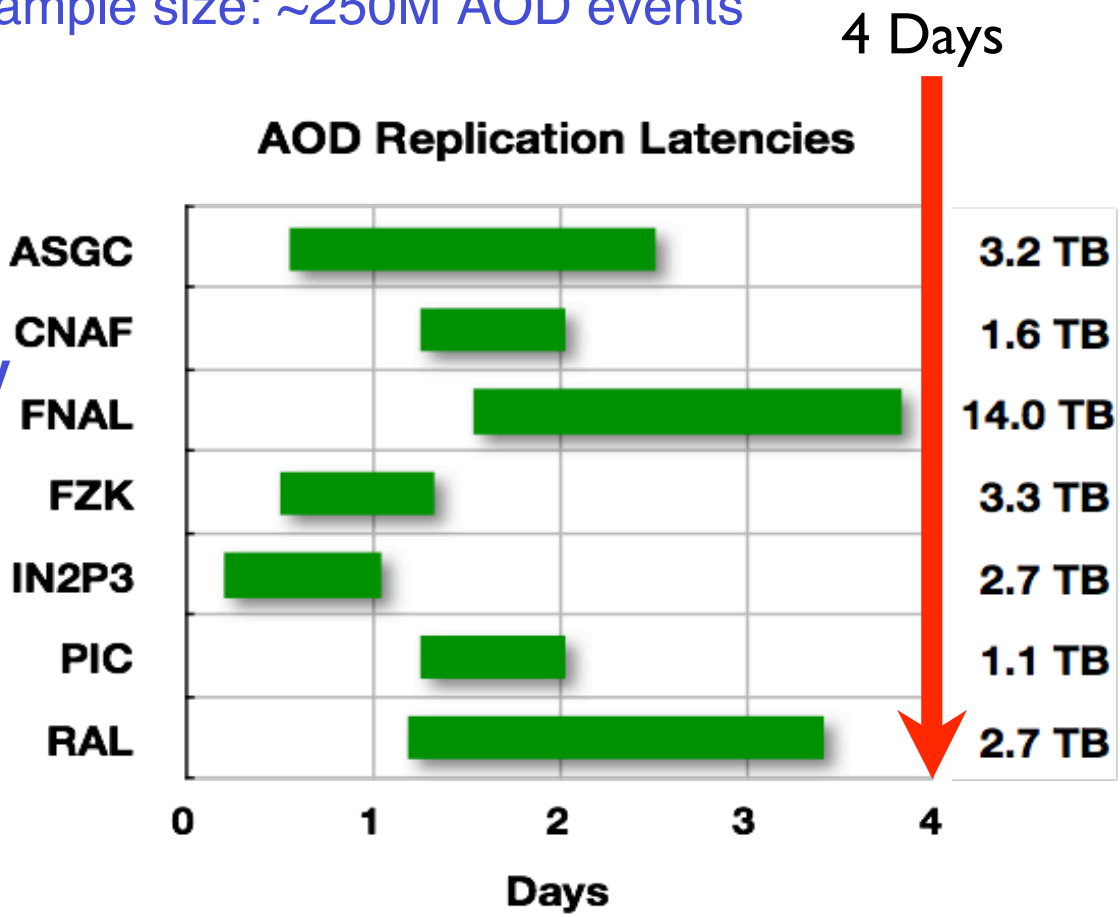

T1 - T2 transfer tests

- **Goal: exercise the full matrix of regional and non-regional transfers in Prod.**
  - **Regional and non-regional T2s enter in the same way - with different target rates though.**
  - **Started May 15., 2 cycles planed**
- **Transfer metric: Latency, participation, rate**

- **Status 178/193 Links tested, ~90% links tested**

T1-->Region avg. rate/target	ASGC	CERN	CNAF	FNAL	FZK	IN2P3	PIC	RAL
ASGC	2	-1	5	2	3	2	6	2
CERN	1	4	4	3	2	4	4	3
CNAF	1	-1	5	4	3	-1	2	1
FNAL	0.2	-1	1	2	2	2	1	2
FZK	1	-1	2	1	4	2	0.9	1
IN2P3	1	-1	2	3	3	3	2	-1
PIC	1	-2	5	2	3	3	2	4
RAL	-1	-1	2	1	1	1	2	1

OK | ACCEPTABLE | FAILED | No links

World-wide CMS Job submission:
scale test of CMS WM systems

- we routinely submit 100,000 jobs per day to all tiers in the system
- peaks of 200k/day seem to work just fine

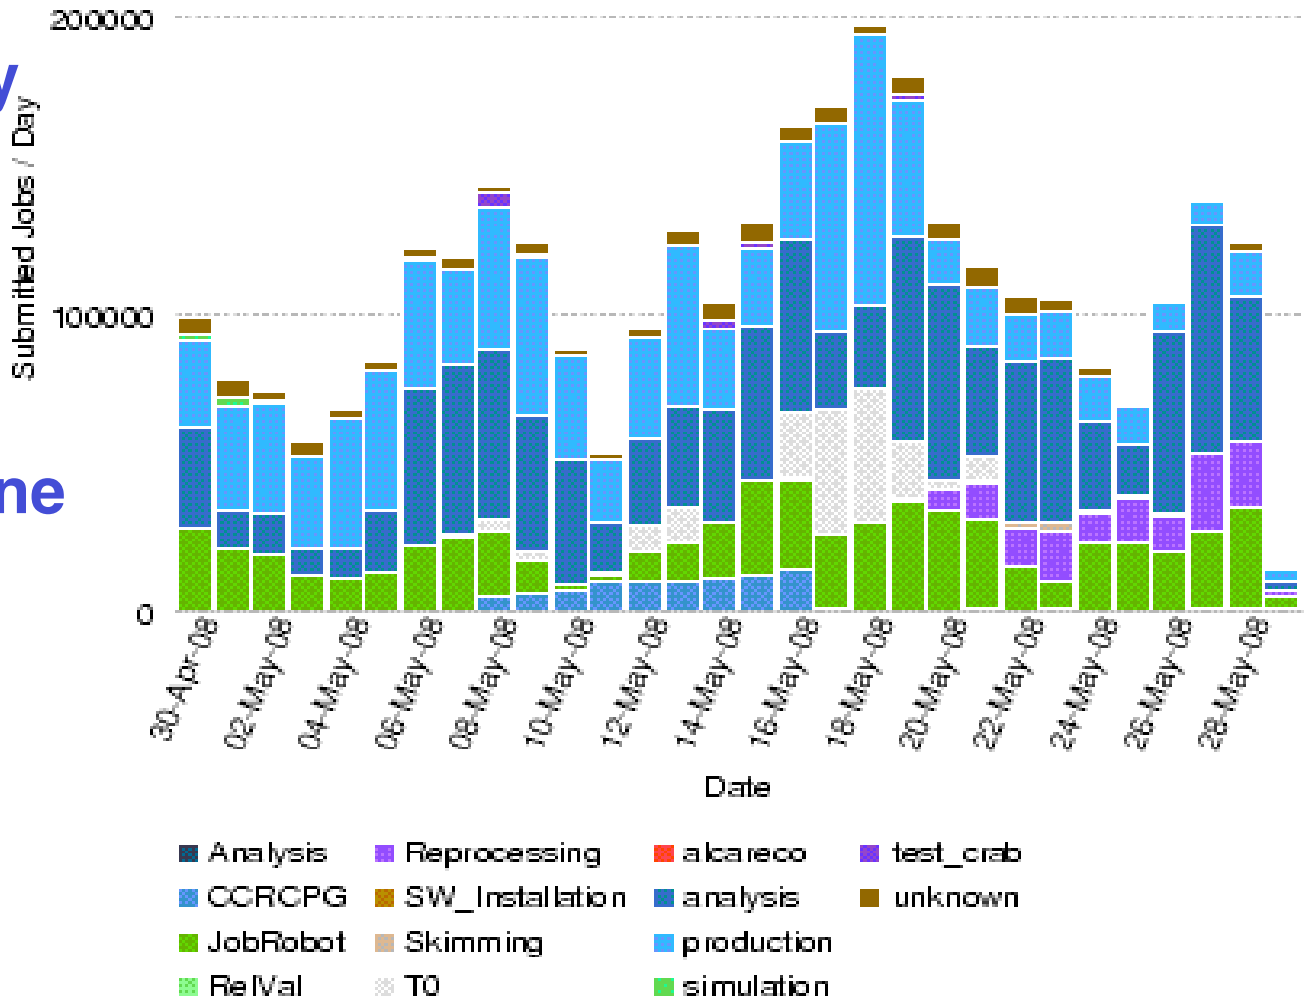

Can CMS routinely submit the required numbers of Jobs/Day? —YES!

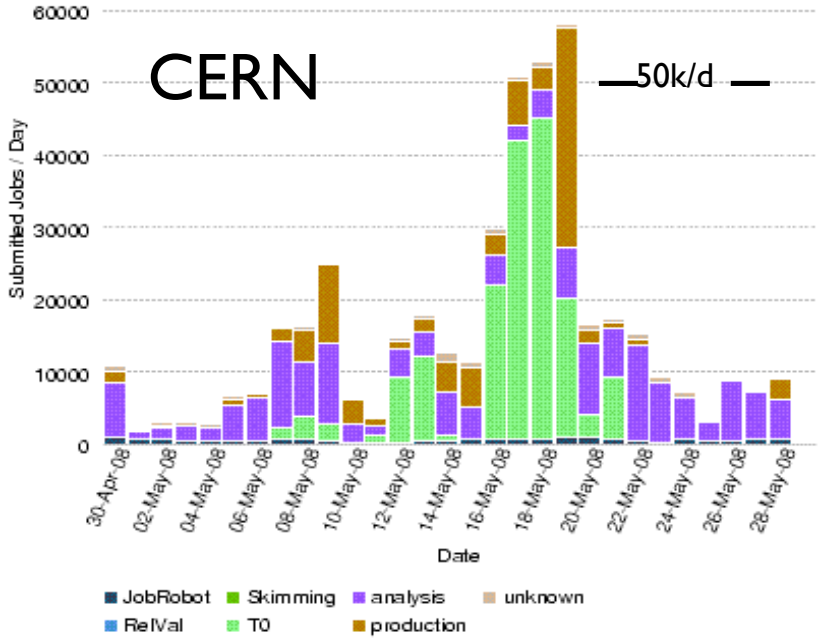

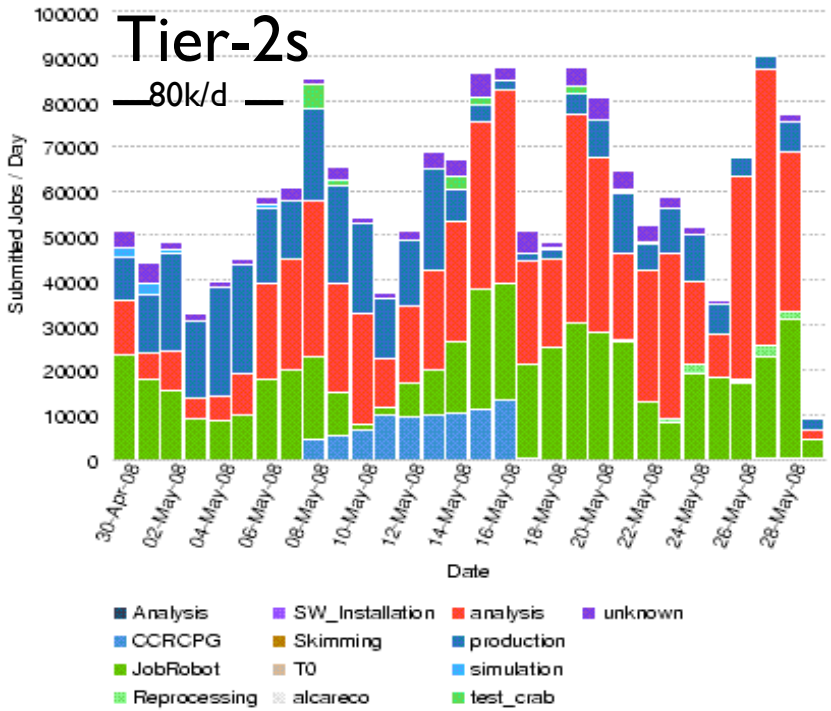

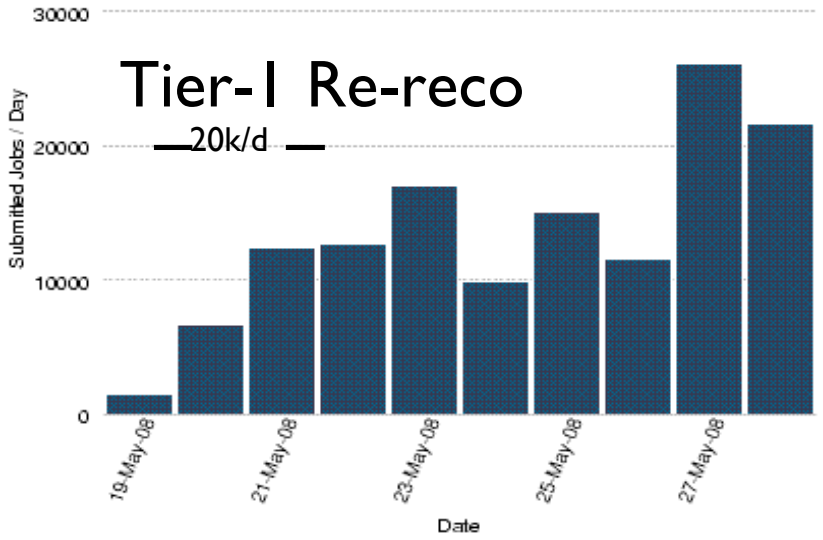

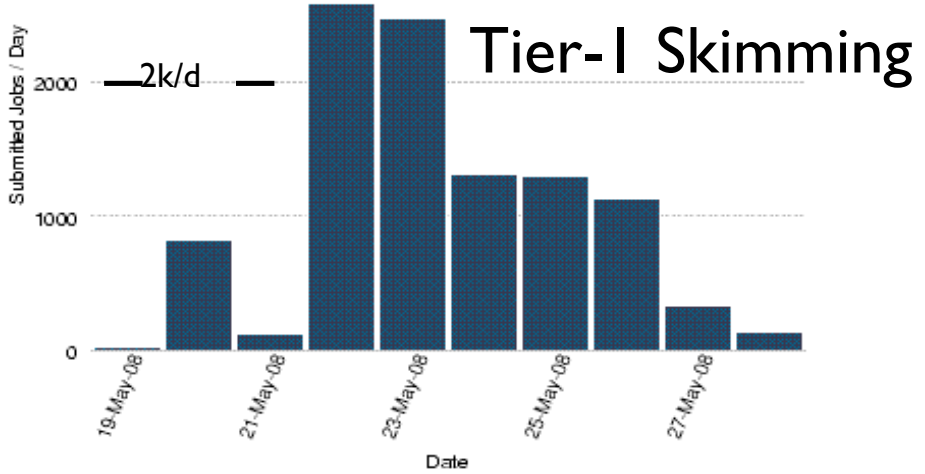

Re-reco and Skimming at Tier-1 centers

- **CSA re-processing great success**
  - **high throughput, and great performance of Tier-1 sites**
    - 127M and 106M events within ~5 days each
      (NB: 1M events is 1h of data@300Hz)
    - nominally 6300 slots, but could “opportunistically” use more

- **CCRC Skimming of RECO and AOD data sets**
  - **Skimming exercise is useful, uncovered interesting “features”**
  - **more skimming exercises (FEVT…) this week under “controlled” cond.**

- **Custodial storage and organizing CMS file catalogue (TFC) and Tape Families**
  - **The tools work and provide the required flexibility**

Tier-2 Analysis: ready to sustain massive CMS Analysis?

- **Exercise proceeding — possibly into this week**
  - **chaotic submission exercise**
    - Driven by many, also T2 managers

Chaotic submission stats

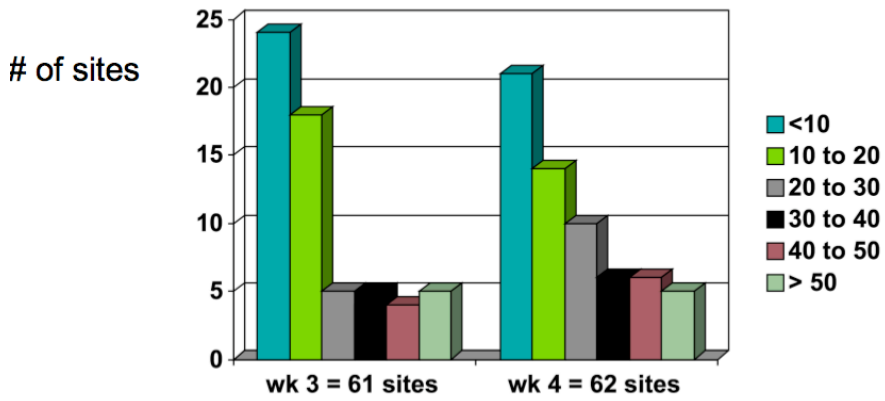

Distribution of number of people who submitted to a site in week 3 and week 4 so far.
Caveats: This includes T1,2,3, and only half of wk 4.

  - **"stop-watch" exercise**
    - Mainly driven by 3-4 individuals

- **Last week ramped up T2 MC production again to perform tests in "realistic" scenario of MC load at T2s**
  - **Results to be summarized…**

CCRC08 T2 analysis exercise

Centrally Organized Activities

- **Measure success rate for "typical" physics analysis jobs**

Performance exercise:

- **Exercise 1: modest IO**
  - failure rate ~0.3 to 0.6%
- **Exercise 2: nominal IO**
  - failure rate ~9%
    - 90% of the failures are in just 3 sites where the SE failed
- **Exercise 3: local stage-out** (May 26th 10:00 - ongoing)
  - Looks promising…

"Physics Group" Exercise Success Statistics

Real analysis like job: running for 3-4hours reading data and with remote stageout of O(20MB)

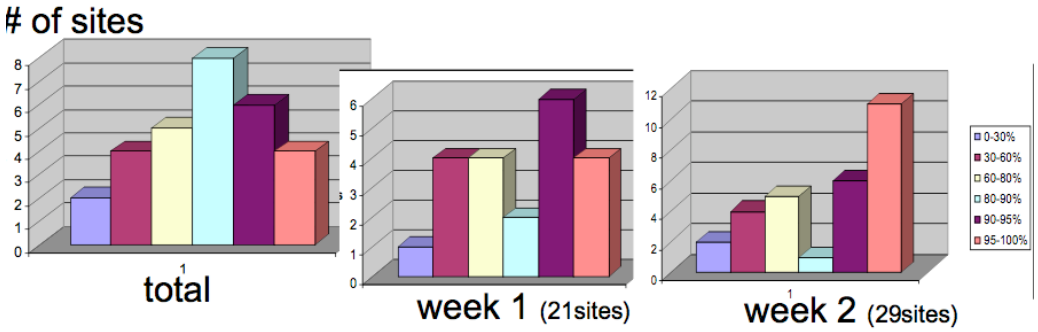

T2 success rate	
>95%	18
90-95%	4
80-90%	0
70-80%	2
60-70%	2
50-60%	0
<50%	2
total sites	28

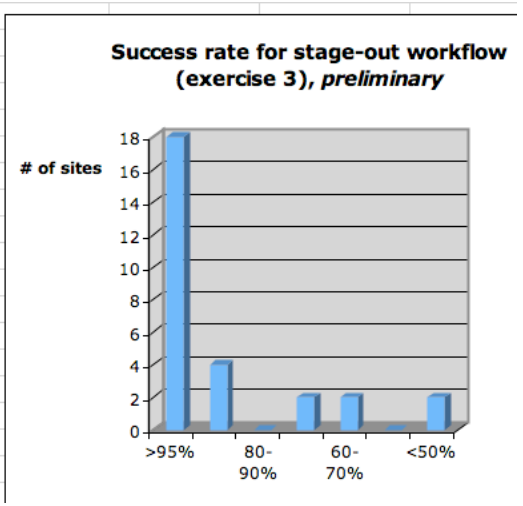

CCRC: Common challenge with other experiments?

Regular, daily operational meetings organized by WLCG (J.Shiers)

- Attended by experiments, T0, T1 and T2 sites
- Experiments status and plans presented (for CMS: D.Bonacorsi)
- Mostly operational issues discussed

Very effective to solve operational issues

Little "cross-experiment' coordination and scheduling

- Each experiment had full schedule of activities
- Some correlated tests planned at sites (collect information now)

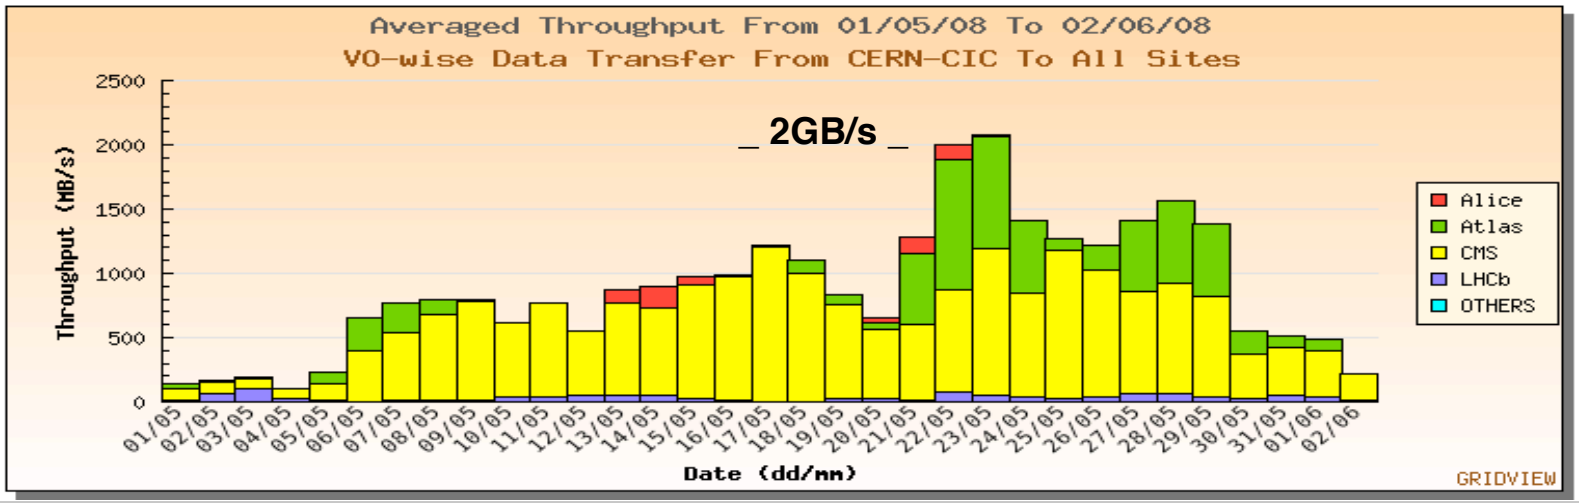

CSA08 Summary

- CSA08 has been very successful
  - production, calibration & alignment completed
  - physics analysis still ongoing, no show stopper
- Calibration and alignment in CSA08 worked very well
  - S43 & S156 exercises completed on time by all sub-detectors
  - All required constants uploaded to the production database
  - Re-reconstruction could proceed on schedule
- Various problems were encountered but were quickly overcome and the overall time schedule for CSA08 did not slip
- Organizational challenges were mastered
  - complexity of a large number of workflows
  - inter-dependences between workflows
  - management of database conditions

CCRC summary

- For CMS CCRC08 was very successful
  - Demonstrated all key use case performances of T0, CAF, T1, T2 infrastructure
- Demonstrated simultaneity with major DPG/ALCA/Physics activities and, and the same time, stress tested the computing infrastructure with real and artificial load

- Some tests will happen next weeks
  - T0 end-to-end test
  - Some coordinated tests with ATLAS at T1 centers
  - Data deletion exercise
  - Review T0 - T1 performance, re-test where required…

- Lessons learned are being compiled now
  - Review at computing management meetings, June 18-20
  - WLCG CCRC workshop scheduled for June 12/13

Backup slides

CCRC Goals Metric

Service	2008 Goal	CSA07 Goal	Status	CSA06 Goal	Status 2006
Tier-0 Reco Rate	150Hz - 300Hz	100Hz	Achieved Bursts	50Hz	Achieved
Network Transfers between T0-T1	600MB/s	300MB/s	Achieved Bursts	150MB/s	Achieved (6/7 cont.)
Network Transfers between T1-T2	50-500 MB/s	20-200 MB/s	Achieved Most Sites	10-100 MB/s	Achieved (15 sites)
Network Transfers T1-T1	100MB/s	50MB/s	Only Opportunistic	NA	Not Attempted
Job Submission to Tier-1s	50k jobs/d	25k jobs/d	Achieved	12k jobs/d	3k jobs/d
Job Submissions to Tier-2s	150k jobs/d	75k jobs/d	20k jobs	48k jobs/d	Achieved
MC Simulation	1.5 10^9 events/year	50M per month	Achieved	NA	Not Attempted

ECAL Calibration: Φ-Symmetry

- Use ϕ-symmetry of event activity to intercalibrate rings of crystals in ϕ:
  - Step 1: Accumulation of energy sums from 20M minbias events
  - Step 2: Derivation of constants from energy sums
  - Step 3: Validation like in "real life operation" (blind to MC truth)

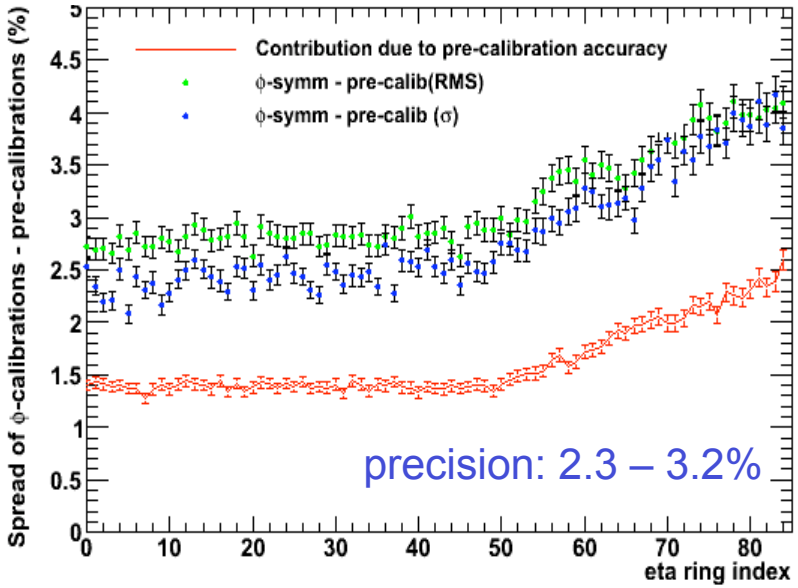

Stefano Argiro,
Tommaso Tabarelli di Fatis

- As expected, cannot improve on pre-calibrations for barrel
- Result slightly worse than seen in CSA07 due to use of a higher ET threshold (more robust, but need more statistics), and increased tracker material budget

ECAL Calibration: Phi-Symmetry

- Systematic limit due to ϕ-inhomogeneity of tracker material and ECAL geometry.
- Investigate by increasing statistics through folding of crystals within η regions.
- Limit from observed to be about factor 2 higher than in CSA07.

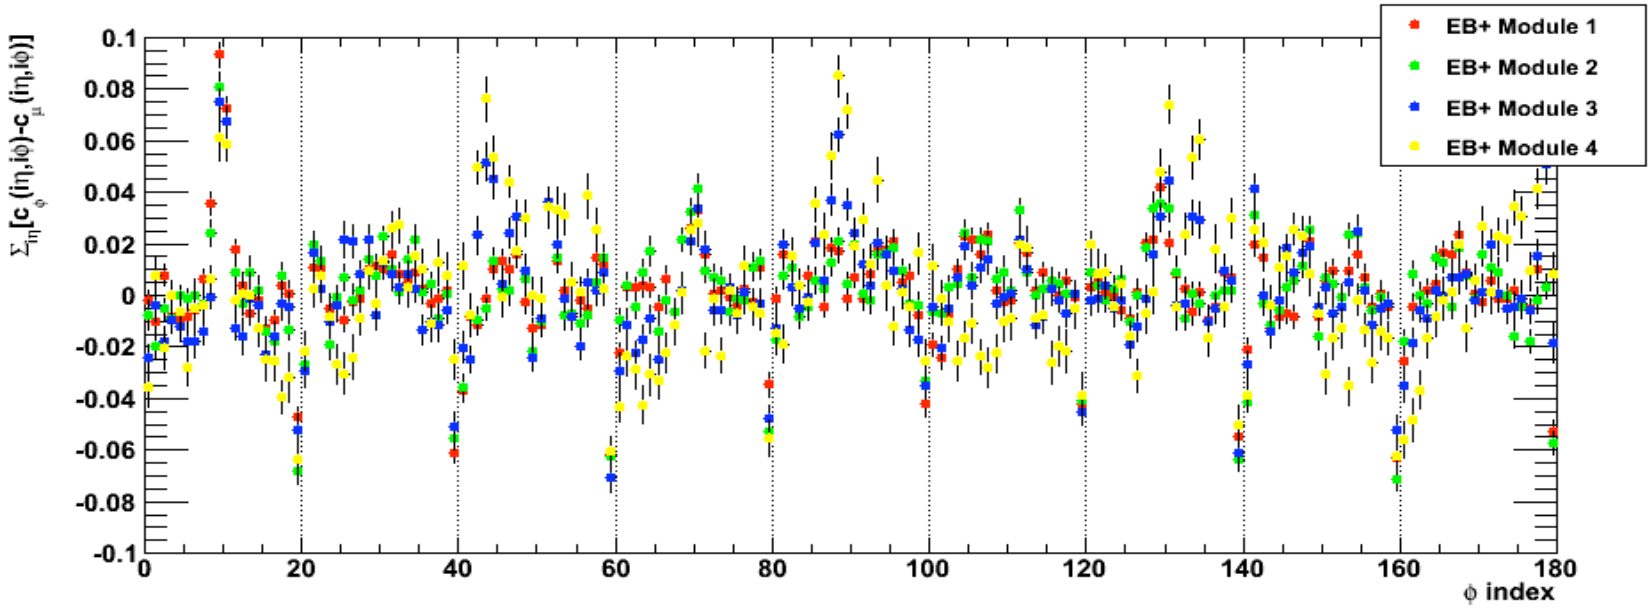

Mass Resolution $Z^0 \rightarrow e^+e^-$

One electron in barrel, one in endcaps S156 Both electrons in endcaps

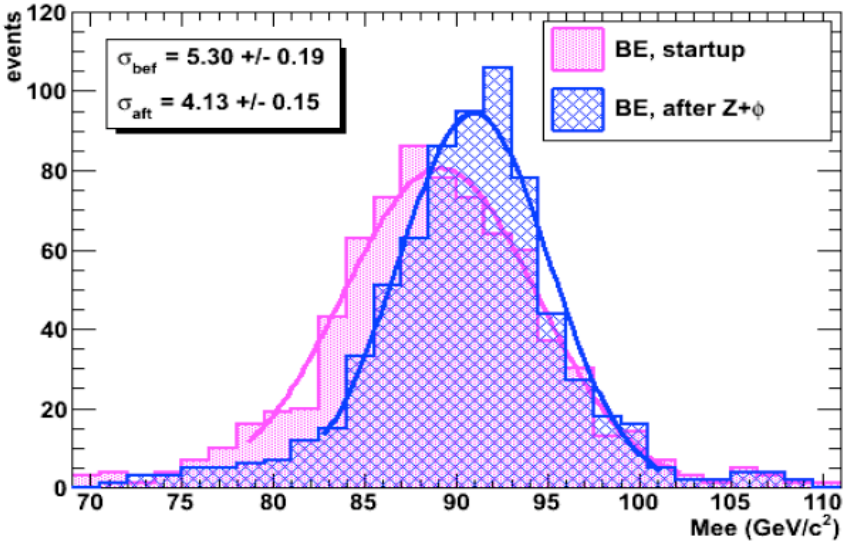

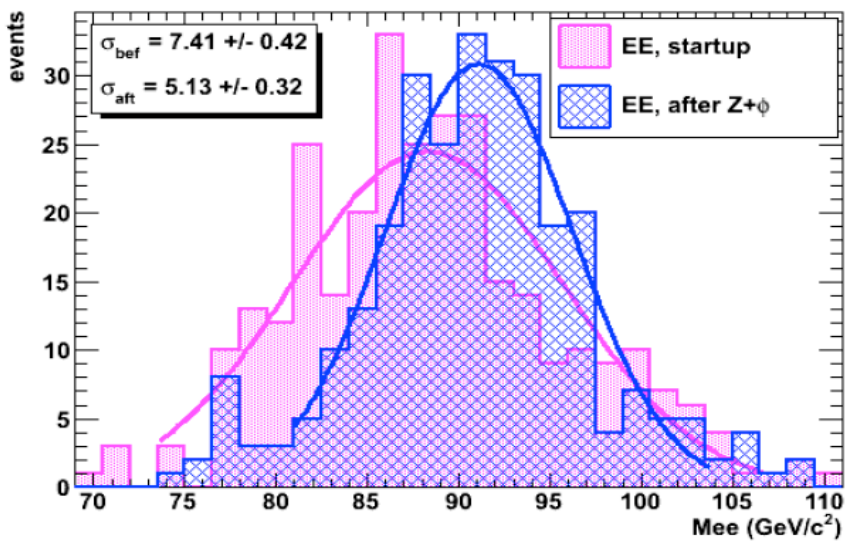

- **Due to the very good pre-calibration of the barrel, the relative improvement is more prominent with both electrons in the endcaps**
- **Achieved calibration precision in the endcaps is 7% (pre-calibration: 14%)**

ECAL Calibration: π^0

Vladimir Litvine, Marat Gataullin, Yong Yang

- **Full workflow tested**
  - **Samples used: Minbias, QCD Jets**
  - **Half of available statistics used to perform corrections as functions of η and ϕ, half to reconstruct mass peak and calibrate**
  - **Efficiency and S/B consistent with expectations**
- **CSA08 statistics insufficient to calibrate individual crystals:**
  - **Calibration relies on π^0 stripping at L1**
  - **CSA08 samples cover kinematic regions having high probability to pass HLT, whereas 70-80% of useful π^0s come from $p_T$hat 10-30 GeV.**
    - only 1 useful π^0 per 300 events
  - **Calibration exercise performed with 10x10 crystal folding. Constants not used for re-reco**
- **In real life, between 10-100 pb-1 the π^0 mass constraint is expected to improve the calibration considerably (→1%)**

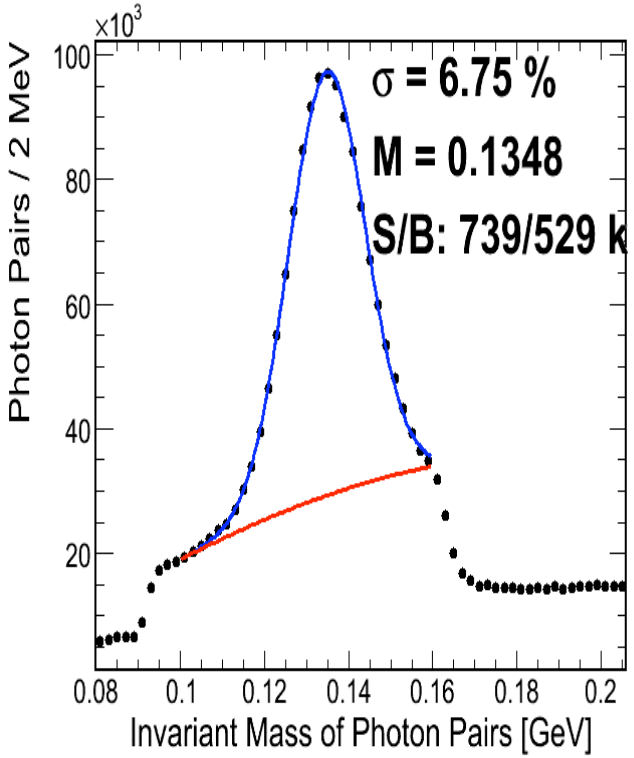

Tracker Alignment

E. Butz, M. Edelhoff, A. Gritsan, V. Karimäki, C.K. Lae, T. Lampén, S.Miglioranzi, A. Morelos Pineda, J. Olzem, M. Stoye, N. Tran, M. Weber, E. Widl, B.Wittmer

- **Methods used:**
  - **HIP in 2 variants:**
    - "2-step" = first align larger structures (rods, petals, ladders), then dets
    - "HIP survey" = align dets in 1 step, use survey constraint
  - **Kalman filter:**
    - align TOB first, then use as reference
  - **MillePede:**
    - module level, parameters u, (v,) w, γ (44244 free parameters)

Tracker Alignment (cont'd)

- χ^2 from track fit (after alignment) indicates alignment quality

- Millepede found to give best results
  - for S43, only minimum bias (6.6M) and muon (p_T>5 GeV) events were used
  - for S156, cosmics and more high-p_T muons were added

- Note: this estimator is rather insensitive to weak modes

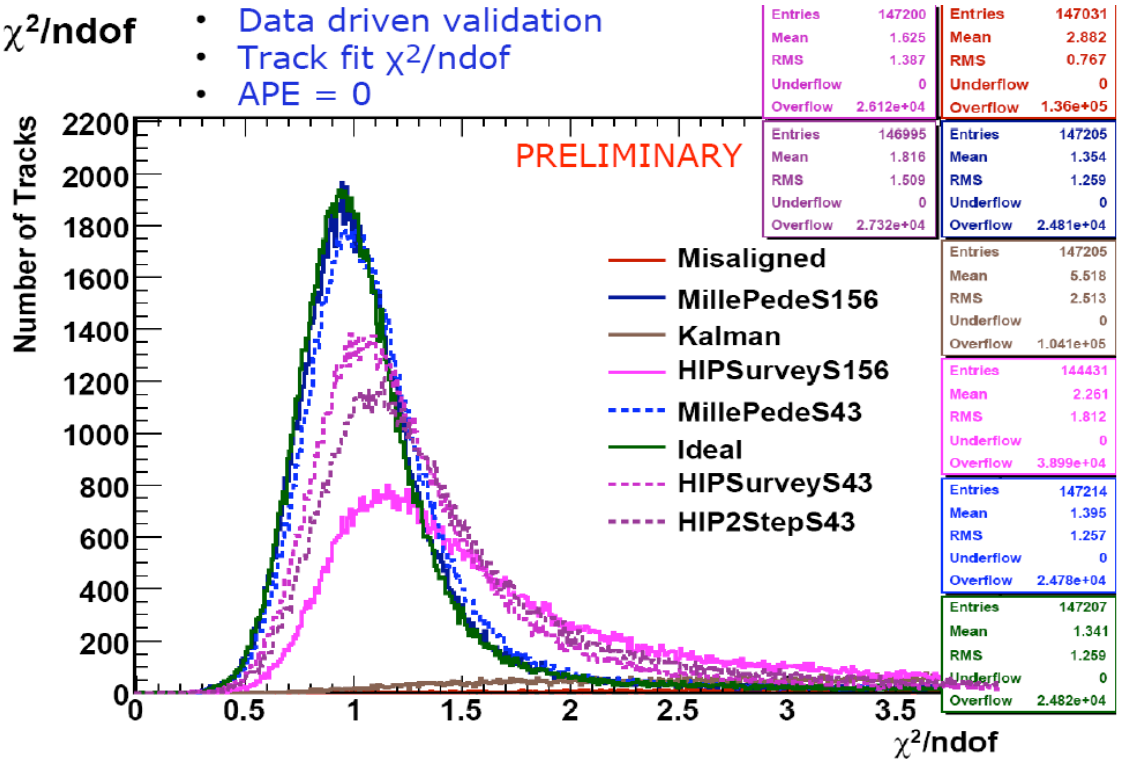

Tracker Alignment (cont'd)

- p_T resolution at high momentum is very sensitive to coordinate resolution & thus to alignment
  - also systematic effects (e.g. due to weak modes) can show here
- Visible improvement (Gaussian fits):

MillePede S43	3.0%
MillePede S156	2.2%
Ideal	1.7 %