ATLAS

LHCC, CERN, September 21 2009

Kors Bos

Content

- Not much about operations
  - A lot was covered last LHCC when we reported on STEP
- Data protection and dCache
- DB access
- Software Performance improvements
- Thermo dynamical Data Model

Data Protection on dCache

- Long standing requirement
- Dangers
  - (un-) intended deletion of data ← did not happen (yet)
  - (un-) intended access to tapes ← happened few times already
- dCache 1.9.4 out since end July supports ACL's
- dCache 1.9.5 "Golden Release" due end October
- All sites are still on 1.9.3.x
- Lyon will upgrade 28 Sept. (announced 72 hrs downtime)
- Difficult balance between "silver" and "gold"

DB components @CERN

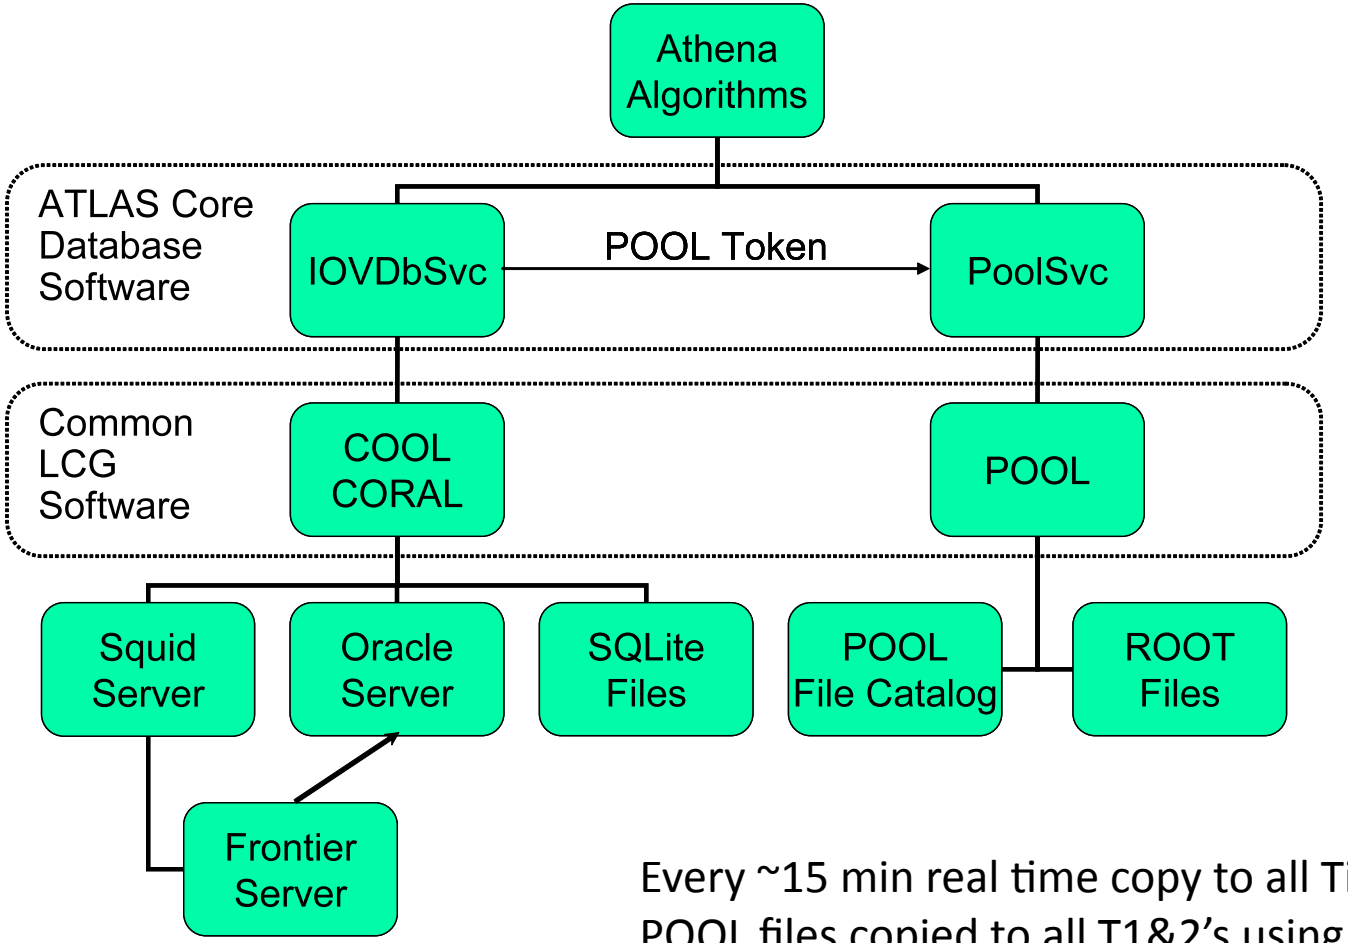

Every ~15 min real time copy to all Tier-1's
POOL files copied to all T1&2's using DDM

DB access use cases and technologies

- Major use cases:
  - Simulation production
    - Tier-1s, Tier-2s, Tier-3s
  - Tier-0 processing
    - Tier-0
  - Reprocessing
    - Tier-1s and some Tier-2s
  - Calibration and alignments
    - Tier-1s, CAF, calibrations centres
  - Group and user analysis
    - Everywhere
- Available technologies:
  - Direct access to Oracle databases
    - Good for jobs running at CERN and Tier-1s but possible overload problems
  - (Conditions) DB release
    - Best for production tasks needing fixed conditions
  - FroNTier/Squid (access to Oracle databases through web caches)
    - FroNTier cache in front of Oracle server helps with load problems and local (Squid) cache solves latency problems for jobs running at Tier-2/3s
  - DB-on-demand (SQLite extraction of DB data)
    - Useful to run on "disconnected laptops"

Improving DB access for analysis

1) Augment HammerCloud tests with jobs requiring DB access
   a) Test FroNTier/Squid access vs direct Oracle access
   b) Collect information on load levels of FroNTier, Squid and Oracle servers
2) Define the baseline parameters of FroNTier and Squid servers
3) Make sure that POOL conditions files are replicated to Tier-2s and then evaluate the need for further replicas at Tier-1 sites
4) Evaluate the global system efficiency and the necessary evolution of hardware and software to cope with the 2009-2010 data-taking period
5) Discuss and implement a way to set the environment on each site so as to point to the nearest Squid and the local POOL file catalogue

Software Performance

- Reconstruction
  - fits within 2 GB now: gained -200 MB !
  - Detailed studies of memory consumption per domain
  - 64 bit code needed for upgrade studies but still difficult to run
- Simulation
  - Reduced the event size of HITS by 40%
  - Memory footprint ~600 MB → may use 64 bit code
  - Simulation time reduced by almost factor 2

Evolution of memory consumption

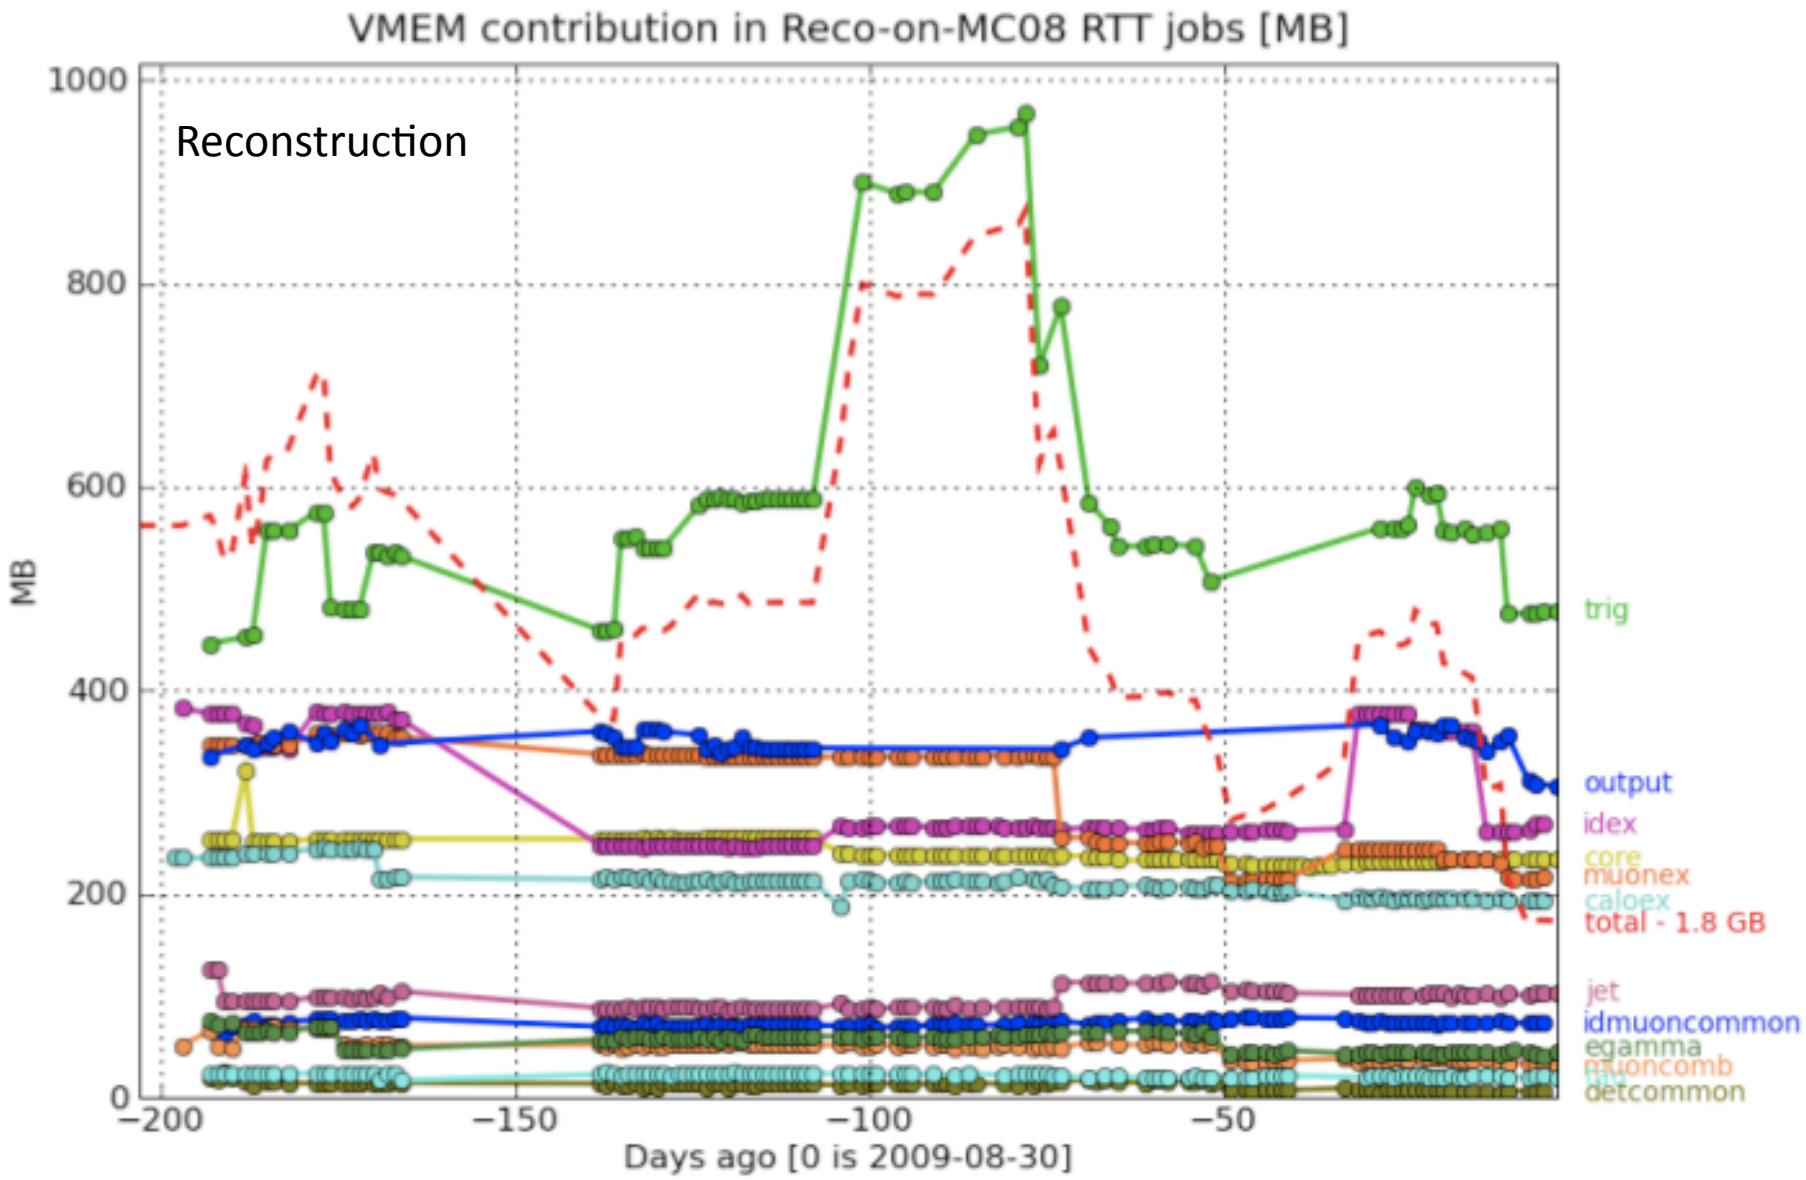

Memory usage per domain

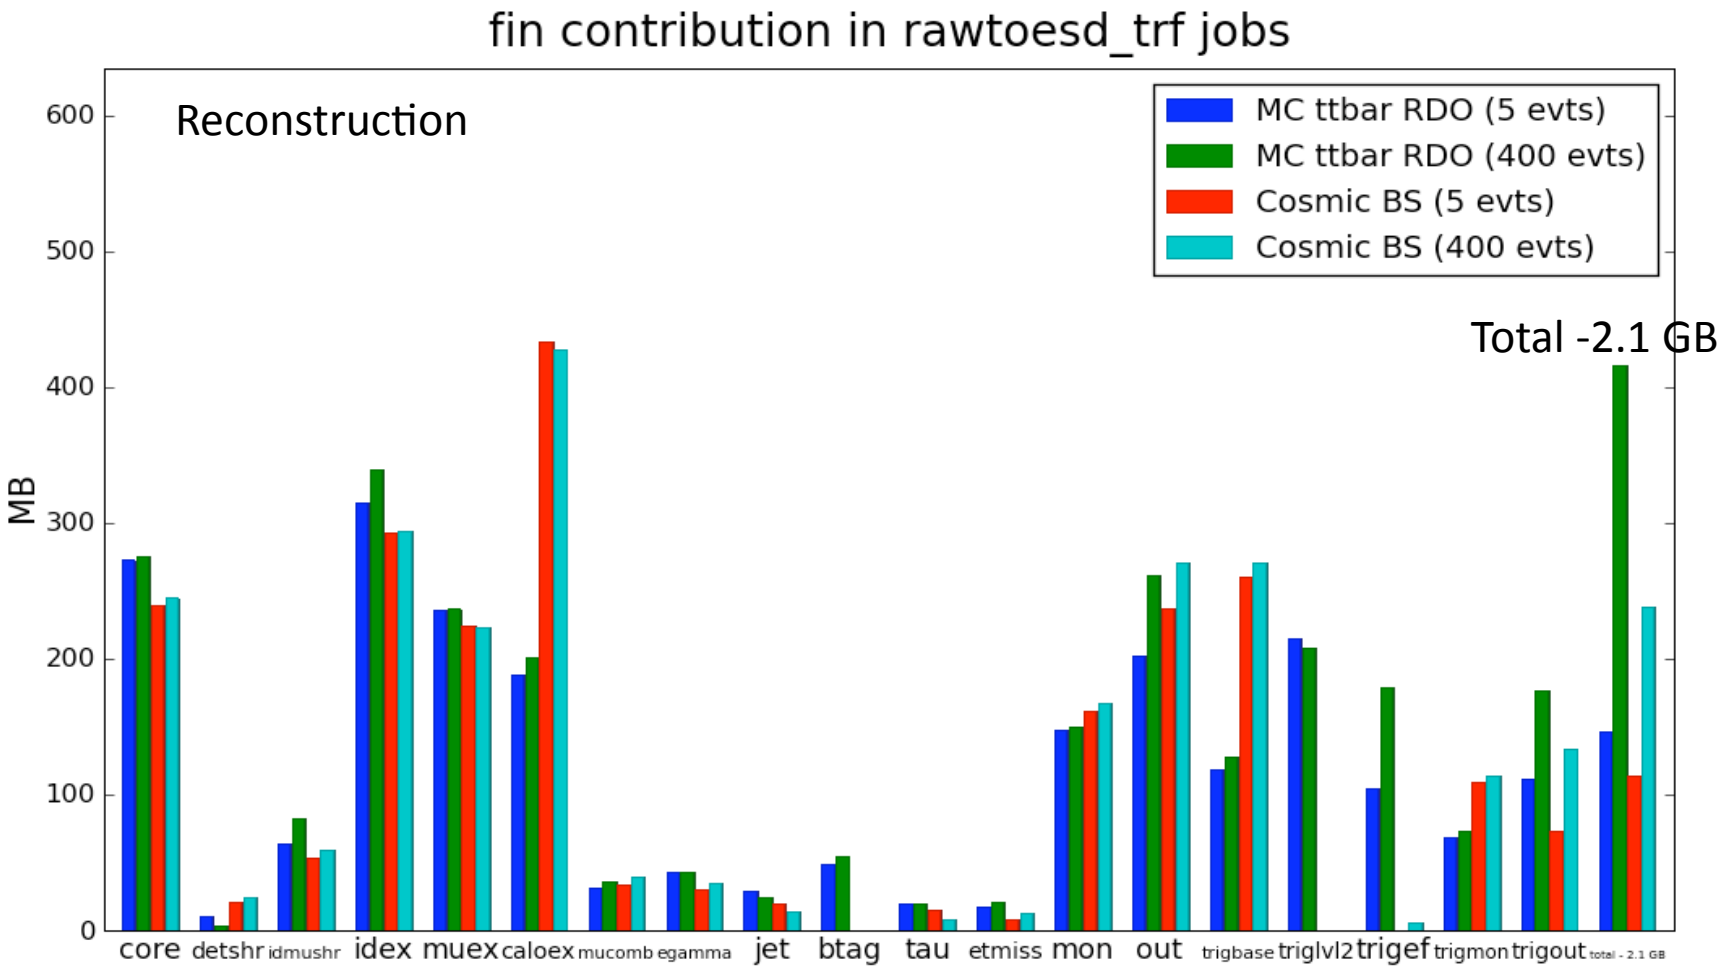

Per domain memory usage 15.4.X.Y

Simulation: HITS event size

Container	14.2.25	15.3.0
MBTS Hits	1.36	1.53
Track Records - MuonEntry	2.57	3.69
MDT Hits	3.82	1.59
Lar Hits HEC	19.8	18.5
Tile Hits	26.0	26.4
Track Records - CaloEntry	26.9	38.5
Lar Hits FCAL	39.1	35.5
Pixel Hits	74.4	67.2
Lar Hits EMB	85.4	86.1
Lar Hits EMEC	93.5	89.9
Truth Event	96.3	82.4
Lar Calibration Hits	235	221
SCT Hits	282	237
TRT Hits	**834**	**196**
Total [kb/Event]	**1821**	**1106**

<50 ttbar events>

Hit container sizes in kB per event

Thanks TRT!!

Down 40% from MC08 to MC09

Evolution of memory consumption

- Well within 1 GB now
- Could run in 64 bit mode (factor 1.5 in size) but ~20% faster

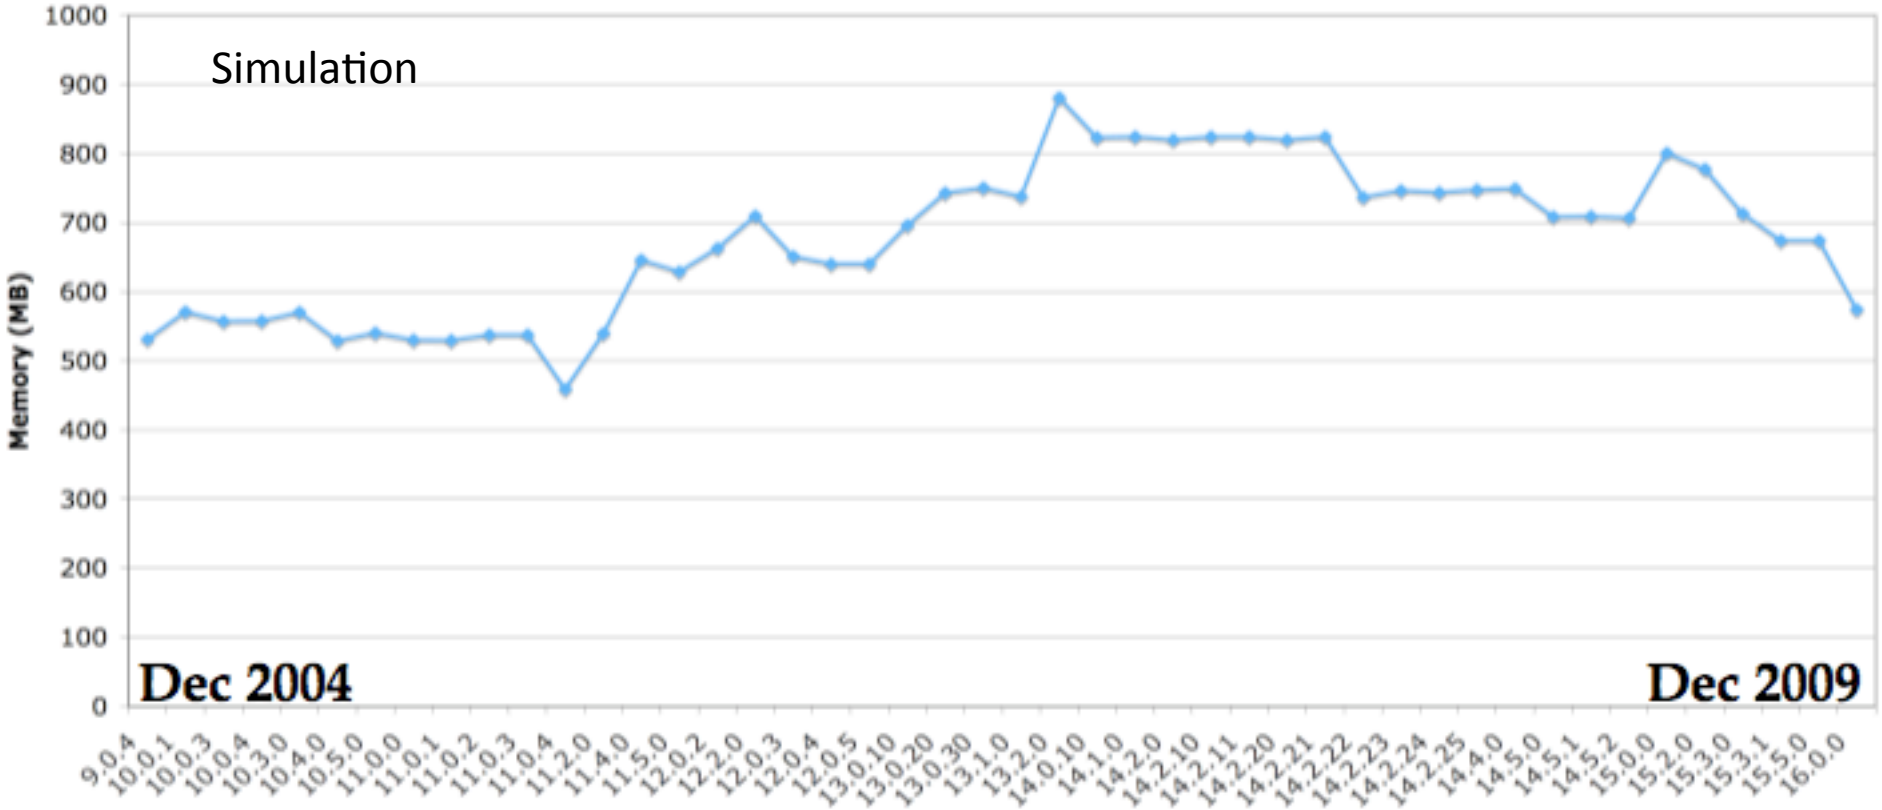

Simulation CPU performance

- Still ~10 minutes per event
- Gain mostly in stepping in magnetic field

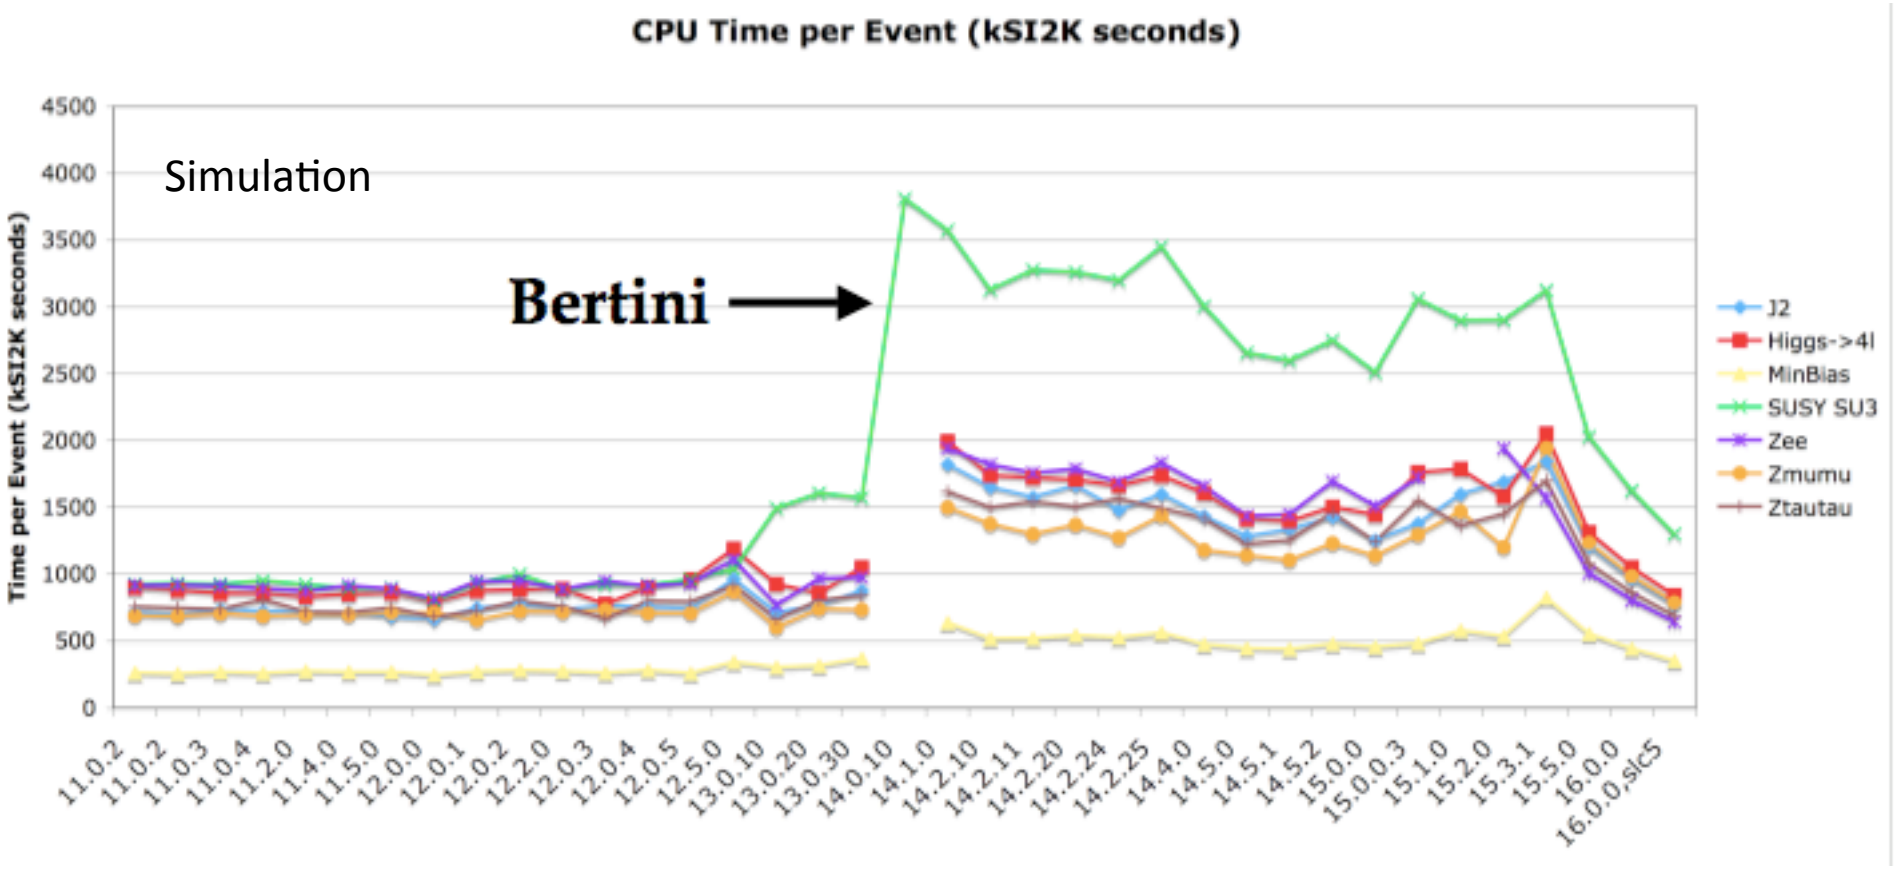

Thermodynamic data model

- ATLAS pre-stages data generously to optimize analysis resources
- Can save disk space by measuring “temperature” of data
- Hot data gets used a lot
- Not used data cools down with time
- If below critical temperature: reduce the number of copies
- Now done by hand but has to become more automatic
- Use the DDM Central Catalog

Count file transactions

- File access by PanDA, Ganga, dq2-get
- Not subscriptions nor local access
- since May 2008, analysis since May 2009
- ~33 transactions/second

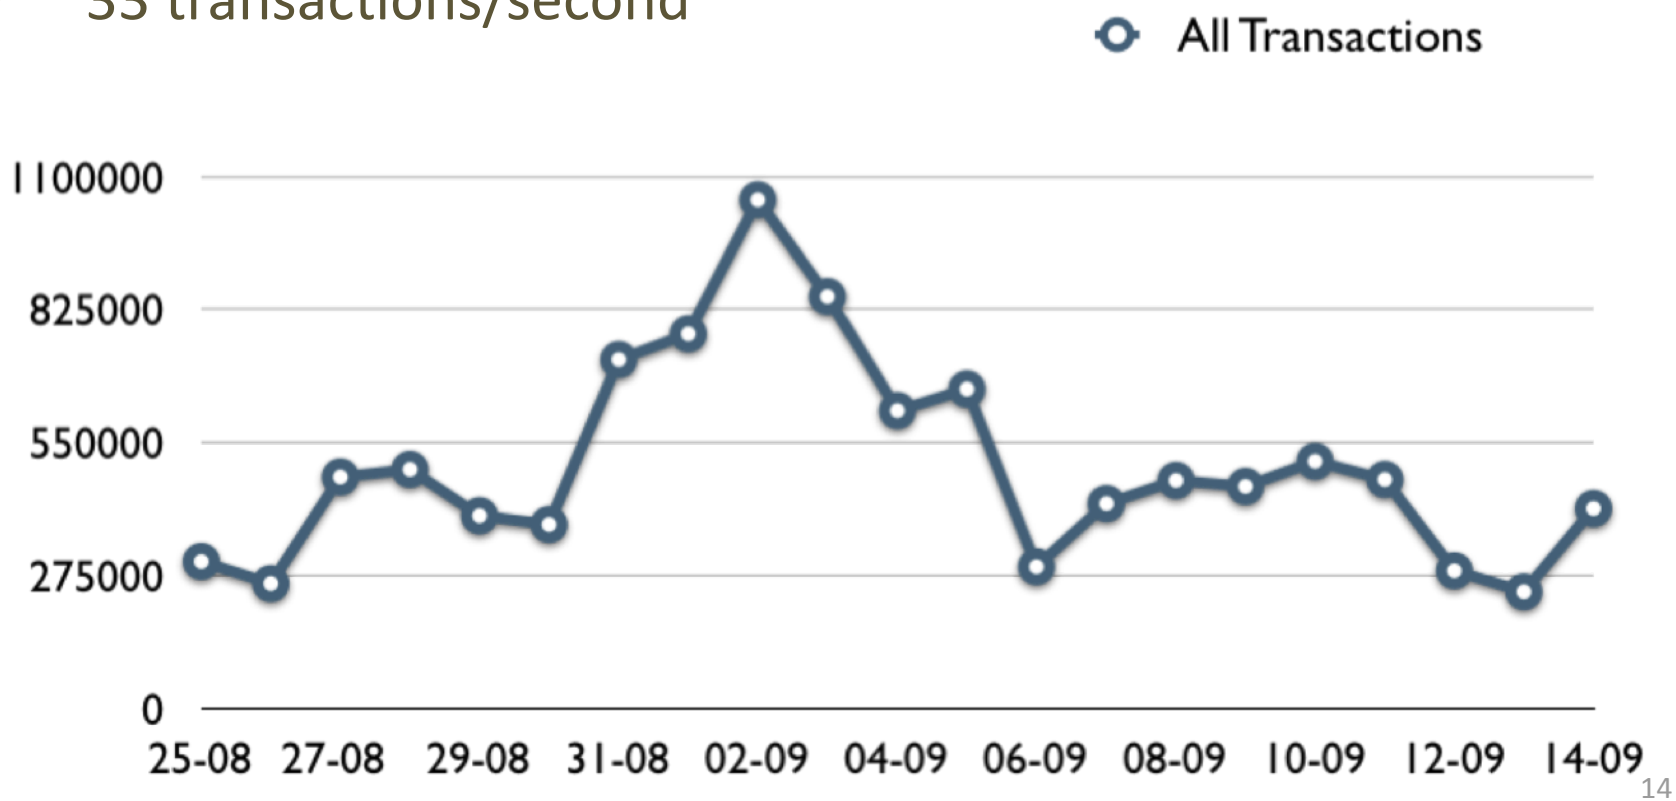

Screenshot 1: user interface

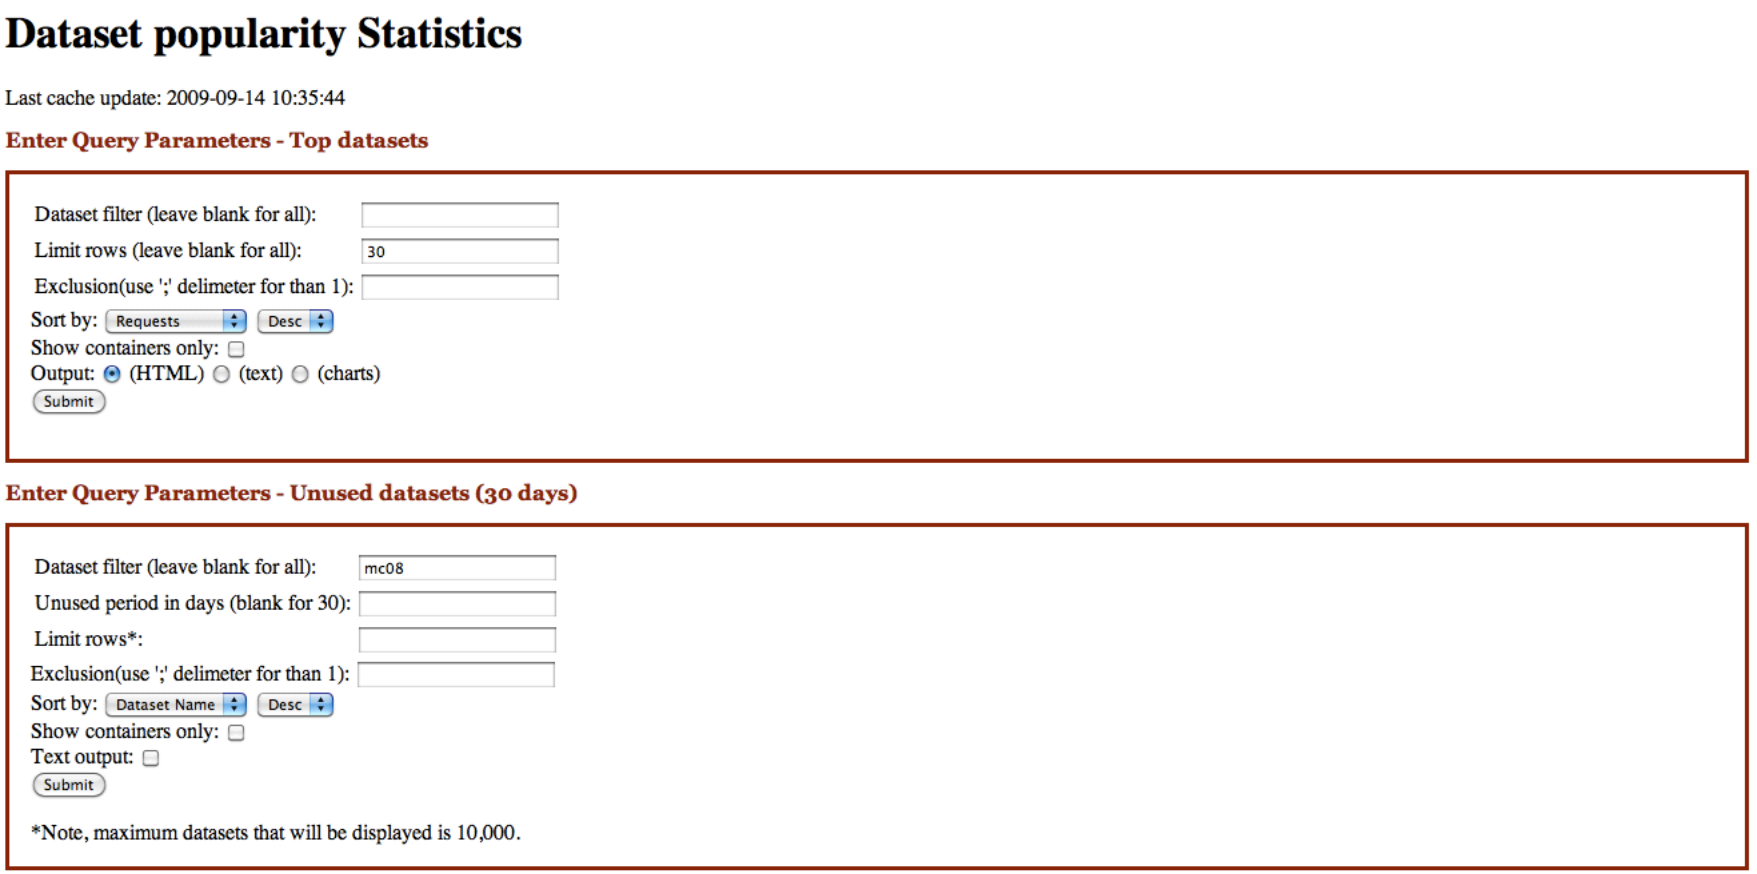

Screenshot 2: most used files

Dataset popularity (30 days)

New Search
Number of records: 4746

Dataset Name	Version	Sites	Locals	Users	Requests	L-Requests	Replicas	File No	File Sizes
mc08.108863.Hijing_beamgas.simul.HITS.e4_s470_tid024853	1	13	13	5	156359	156359	12	400	37616190967
mc08.108852.BeamHaloInputs.simul.HITS.e4_e348_s473_tid025027	1	12	12	5	75259	75259	11	400	1009217177
mc08.105200.T1_McAtNlo_Jimmy.recon.AOD.e357_a68_tid028060	1	6	1	3	34107	30	6	1997	124091608567
mc08.106071.PythiaZmumuJet_Ptcut.recon.AOD.e352_s462_r541_tid026367	1	6	5	4	19459	19458	5	712	30391194298
mc08.106050.PythiaZee_1Lepton.recon.AOD.e347_s462_r635_tid047180	1	3	1	4	17908	17887	4	994	39384985300
mc08.105012.J3_pythia_jetjet.recon.ESD.e344_s479_d153_r643_tid064886	1	3	2	3	15931	15881	2	944	683010700951
mc08.105807.JF35_pythia_jet_filter.evgen.EVNT.e449_tid0771480	1	13	13	1	14263	14263	1	10000	107277454534
mc08.105200.T1_McAtNlo_Jimmy.recon.AOD.e357_a68_tid039300	1	3	1	3	13775	2	5	7957	494140698957
mc08.105200.T1_McAtNlo_Jimmy.recon.AOD.e357_a68_tid039301	1	3	1	5	9955	2	5	9953	618266968259
mc08.106059.PythiaB_ccmu15X.recon.AOD.e401_a84_tid070734	2	1	1	1	9266	9266	1	9966	189495345088
mc08.105001.pythia_minbias.digit.RDO.e375_s462_d150_tid038924	1	4	3	3	8657	8655	3	19992	1008141778541
mc08.105010.J1_pythia_jetjet.recon.AOD.e344_s475_r586_tid029094	1	3	2	3	8074	8073	22	1599	60174470719
mc08.108857.SherpaW3jetstoenuQCD.recon.AOD.e418_a84_tid068737	1	1	None	1	7525	None	0	0	None
mc08.105802.JF17_pythia_jet_filter.recon.ESD.e347_s462_d153_r643_tid064843	1	4	2	2	7392	7356	2	999	676863477488
mc08.105012.J3_pythia_jetjet.recon.ESD.e344_s479_d153_r643_tid064887	1	1	1	1	7326	7326	2	565	408376636328
mc08.106038.PythiaB_bbe5X.evgen.log.e401_tid060065	1	1	None	1	7069	None	1	5008	10945401399
mc08.106021.PythiaWmunu_1Lepton.recon.AOD.e352_s462_r541_tid028751	1	4	4	3	7008	7008	7	984	36176086677
mc08.105802.JF17_pythia_jet_filter.evgen.EVNT.e347_tid023941	1	71	70	3	6642	6636	30	799	76660799621
mc08.105807.JF35_pythia_jet_filter.recon.AOD.e359_a84_tid065858	1	1	1	1	5781	5781	1	16000	628037930582
mc08.105807.JF35_pythia_jet_filter.evgen.EVNT.e449_tid0771481	1	14	13	2	5415	5414	1	10000	107271925578
mc08.106050.PythiaZee_1Lepton.recon.AOD.e347_s462_r541_tid028675	1	3	2	3	5405	1049	14	997	38294496603
mc08.109440.PythiaB_bb_ll.recon.AOD.e432_s495_r635_tid076825	1	4	3	3	5095	5094	1	1996	118843905182
mc08.106039.PythiaB_bbmu5X.evgen.log.e401_tid066131	1	1	None	2	5008	None	1	5006	11537068144
mc08.105001.pythia_minbias.evgen.EVNT.e375_tid029950	2	25	25	1	4973	4973	6	599	31436977654
mc08.108086.PythiaPhotonJet_JetFilter_Nj2Et17.merge.AOD.e372_s462_d153_r643_t53_tid065835	1	12	12	2	4851	4851	32	100	300563820373

Screenshot 3: least used files

Unused datasets (30 days)

New Search

Dataset Name	Replicas	Last Read
mc08.007085.singlepart_gamma_E500.recon.AOD.e339_s439_r570_tid028033	23	2009-02-09 14:28:04
mc08.007083.singlepart_gamma_E75.recon.AOD.e339_s439_r570_tid028024	22	2009-03-09 18:54:30
mc08.007082.singlepart_gamma_E25.recon.AOD.e339_s439_r570_tid029055	22	2009-08-11 16:41:25.938860
mc08.007076.singlepart_e_E1000.recon.AOD.e339_s439_r546_tid026443	22	2008-11-13 02:03:19
mc08.007076.singlepart_e_E1000.recon.AOD.e339_s439_r570_tid028012	22	2009-08-08 13:04:52.082562
mc08.007061.singlepart_e_E100.recon.AOD.e339_s439_r546_tid026438	22	2008-11-13 06:40:33
mc08.007072.singlepart_e_E25.recon.AOD.e339_s439_r570_tid027994	22	2009-08-11 16:48:35.309884
mc08.007072.singlepart_e_E25.recon.AOD.e339_s439_r570_tid027995	22	2009-08-11 16:39:48.167155
mc08.007063.singlepart_gamma_E100.recon.AOD.e339_s439_r570_tid028237	22	2009-07-08 22:04:07.110989
mc08.007074.singlepart_e_E200.recon.AOD.e339_s439_r546_tid026440	22	2008-11-14 04:25:17
mc08.007233.singlepart_mu100.recon.AOD.e342_s483_r562_tid027659	21	2009-08-11 16:48:45.578301
mc08.007086.singlepart_gamma_E1000.recon.AOD.e339_s439_r570_tid028036	21	2009-03-11 12:09:03
mc08.007085.singlepart_gamma_E500.recon.AOD.e339_s439_r546_tid026457	21	2008-11-13 16:19:24
mc08.007082.singlepart_gamma_E25.recon.AOD.e339_s439_r570_tid028018	21	2009-03-09 02:33:22
mc08.007071.singlepart_e_E10.recon.AOD.e339_s439_r570_tid027990	21	2009-08-11 15:49:31.190728
mc08.007075.singlepart_e_E500.recon.AOD.e339_s439_r546_tid026441	21	2008-11-14 14:06:38
mc08.007075.singlepart_e_E500.recon.AOD.e339_s439_r546_tid026442	21	2008-11-13 09:14:50
mc08.007075.singlepart_e_E500.recon.AOD.e339_s439_r570_tid028009	21	2009-08-07 19:03:06.997947
mc08.007072.singlepart_e_E25.recon.AOD.e339_s439_r570_tid027993	21	2009-08-11 16:00:53.835096
mc08.007076.singlepart_e_E1000.recon.AOD.e339_s439_r546_tid026444	21	2008-11-13 00:35:35
mc08.007086.singlepart_gamma_E1000.recon.AOD.e339_s439_r546_tid026459	20	2008-11-14 00:51:55
mc08.007086.singlepart_gamma_E1000.recon.AOD.e339_s439_r546_tid026460	20	2008-11-14 12:14:30
mc08.007085.singlepart_gamma_E500.recon.AOD.e339_s439_r546_tid026458	20	2008-11-13 02:25:26
mc08.007084.singlepart_gamma_E200.recon.AOD.e339_s439_r546_tid026455	20	2008-11-13 05:42:04

Screenshot 4: Dataset Popularity

- The most popular 30 data sets

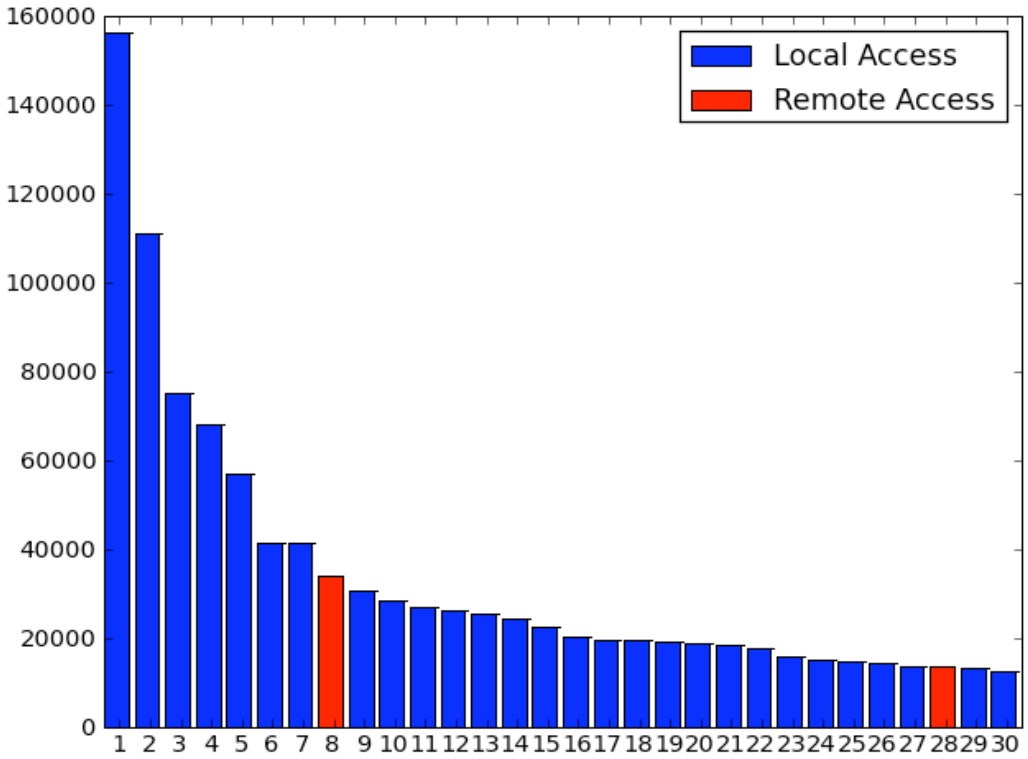

Screenshot 6: Dataset Popularity

- More statistics at log scale

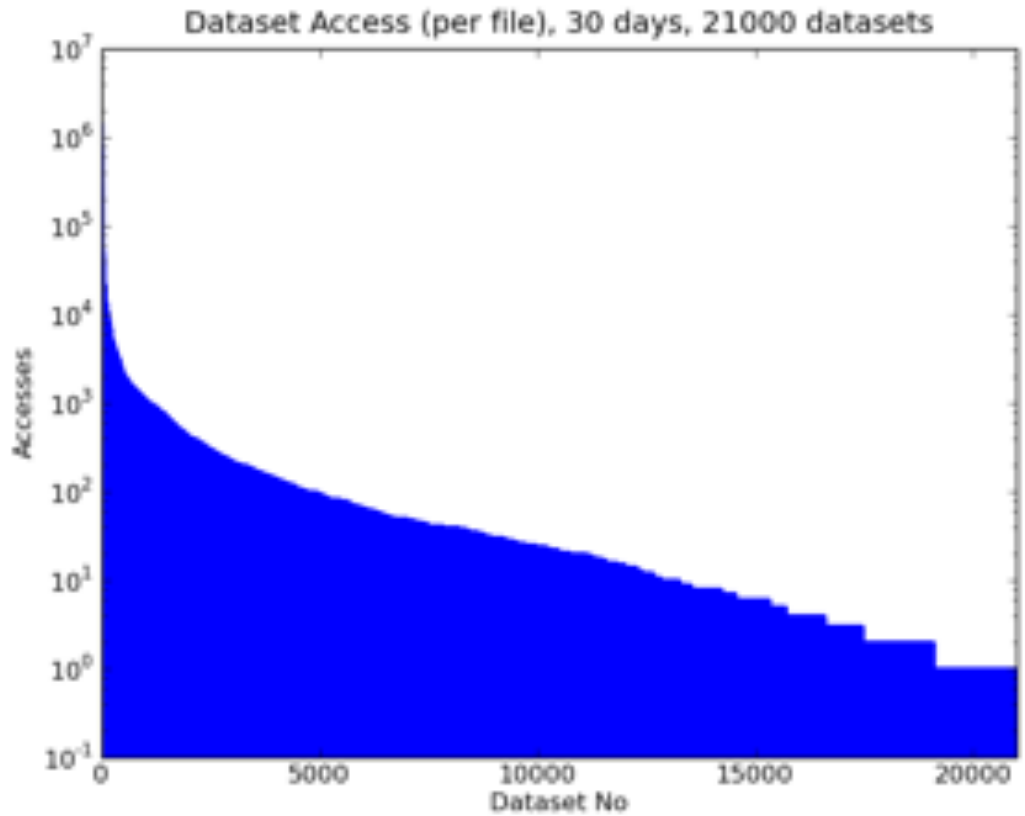

Finally

- Still many ongoing developments
  - Performance improvement, Thermo Model, ..
- Still many new versions
  - SL5, dCache, SS, FTS, …
- Still many more tests done
  - Pre-staging, data distribution challenge, HC analysis, …
- Most sites are in a reasonably stable state
  - Always something wrong somewhere
- Data taking starts on October 12
  - Starting with cosmics
- T0/1/2/3 Jamboree on October 13
  - http://indico.cern.ch/conferenceDisplay.py?confId=66012