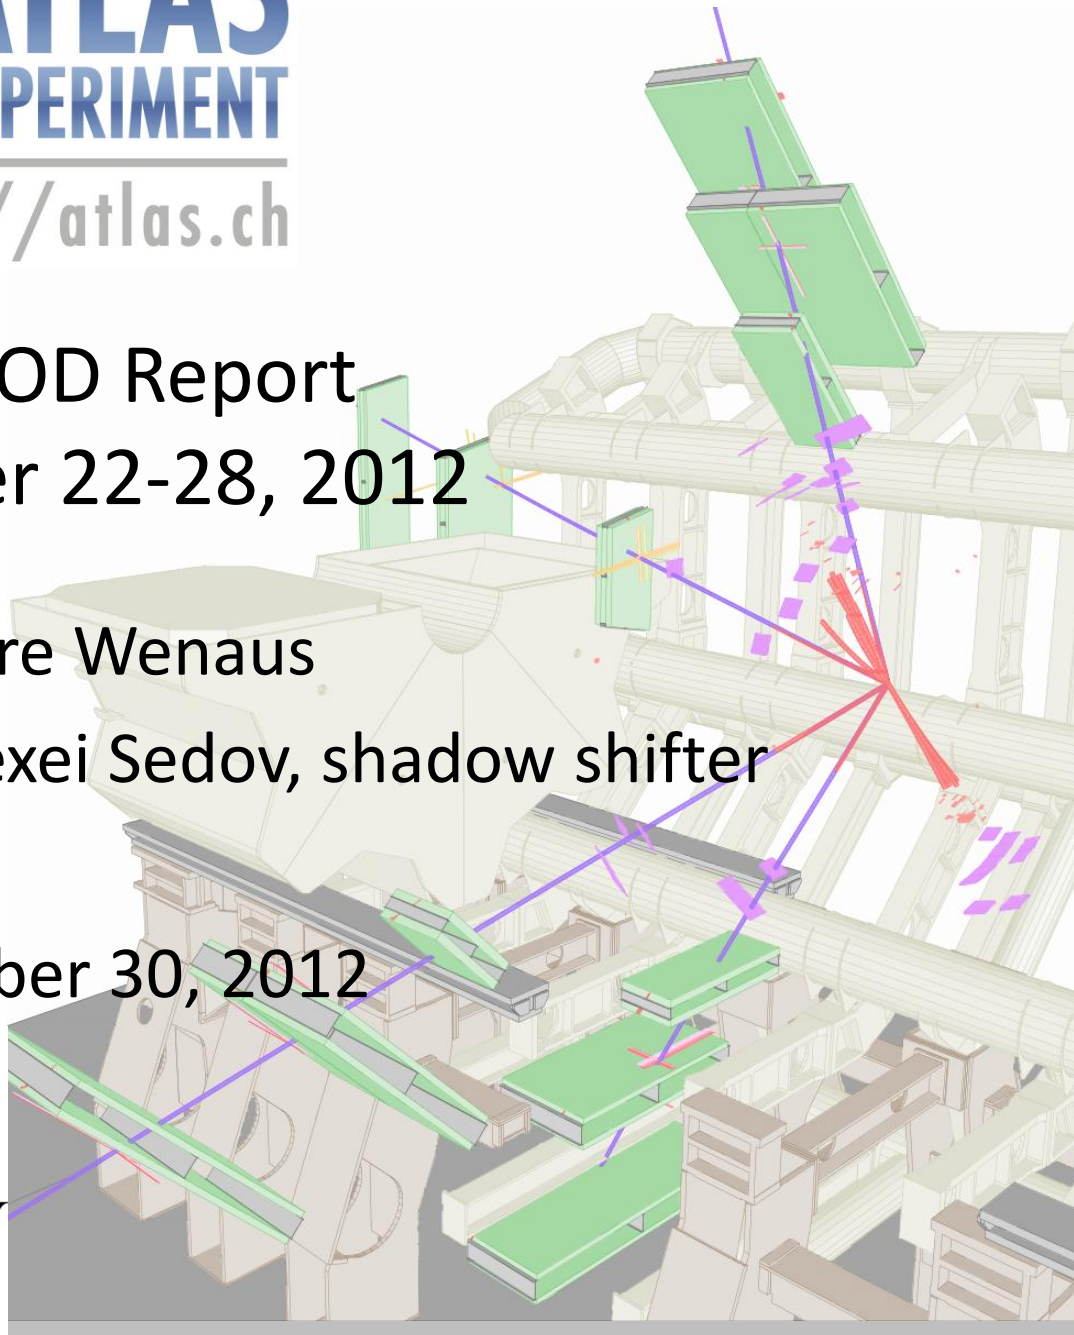

AMOD Report
October 22-28, 2012

Torre Wenaus

With thanks to Alexei Sedov, shadow shifter

October 30, 2012

BROOKHAVEN NATIONAL LABORATORY

Activities

- Datataking (not too many long fills)
- Tail end of express reprocessing
- ~2.3M production jobs (group, MC, validation and reprocessing)
- ~2.7M analysis jobs
- ~550 analysis users

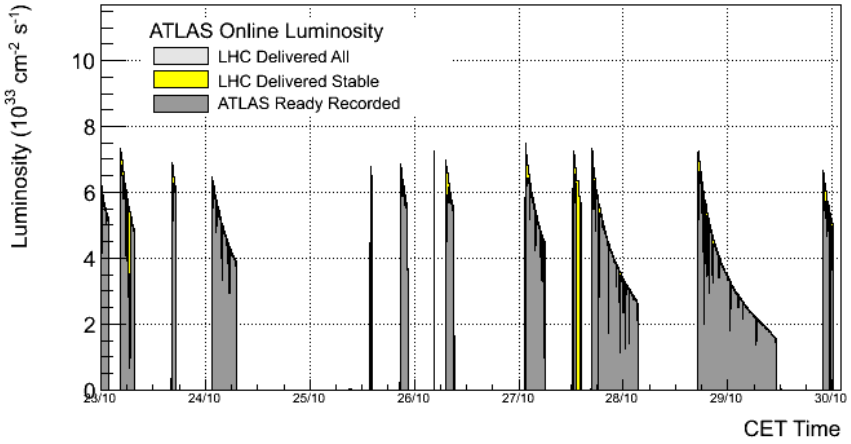

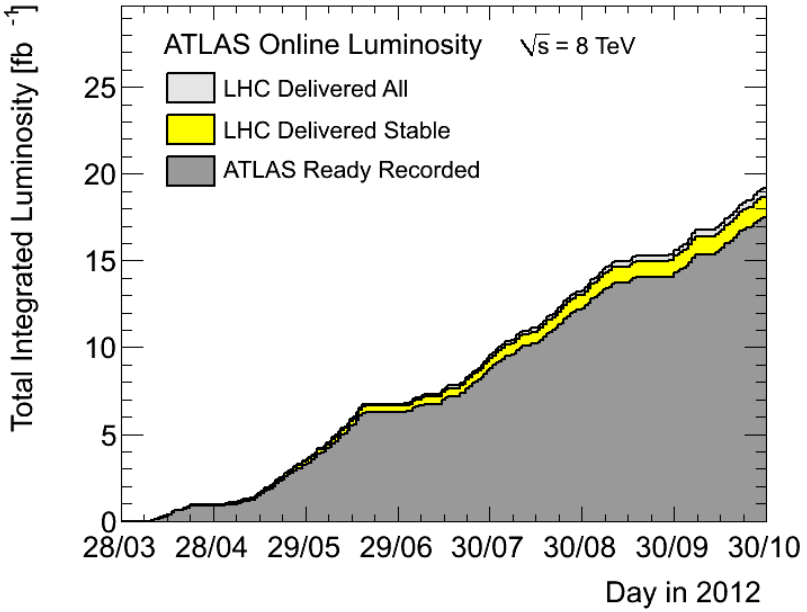

Production & Analysis

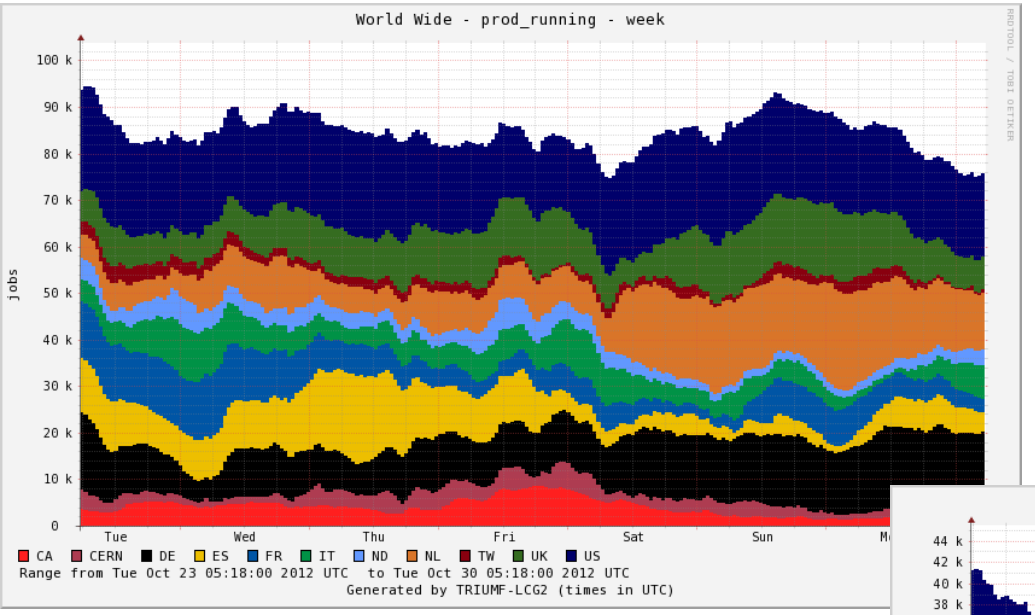

Sustained activity, production and analysis

Constant analysis workload, not spiky
24k min – 41k max

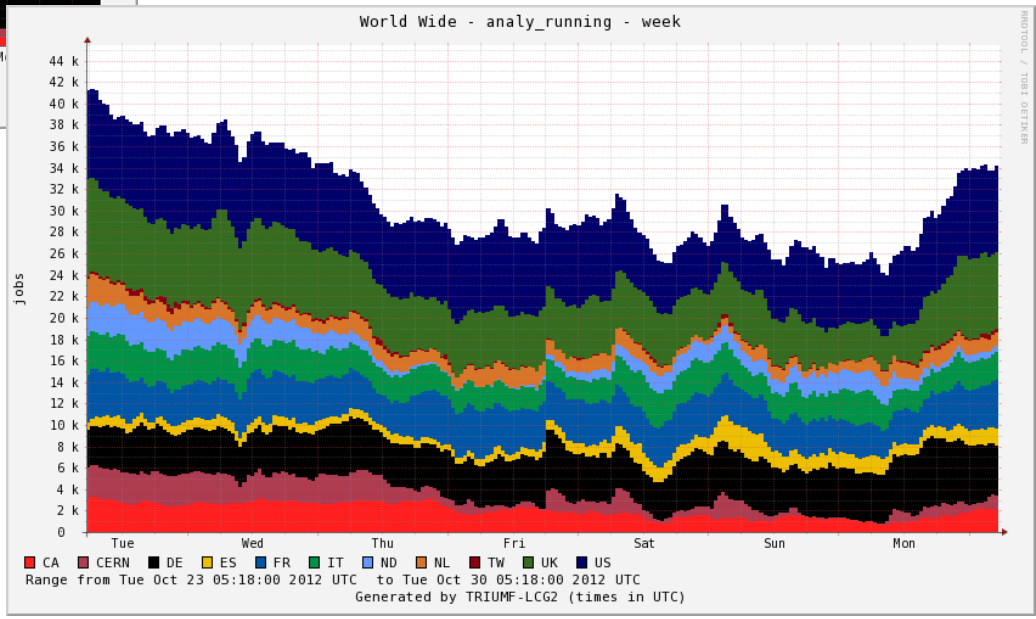

Data transfer - source

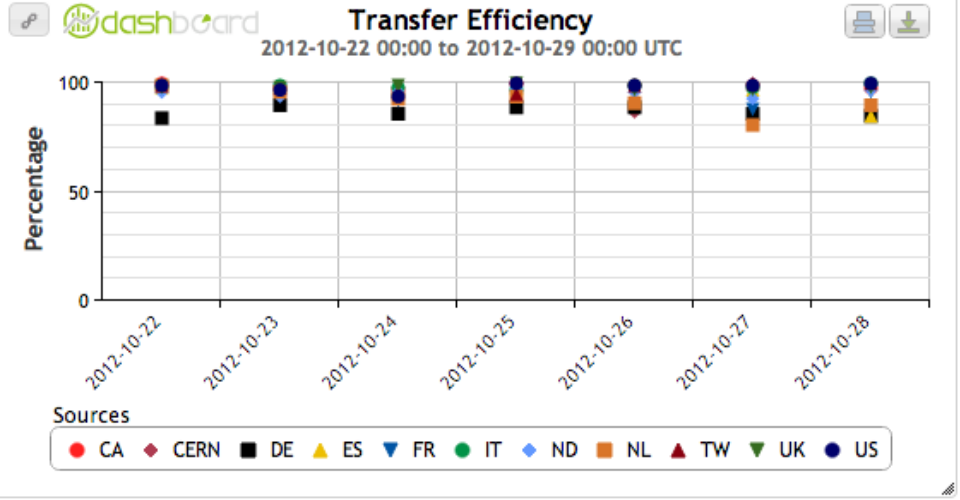

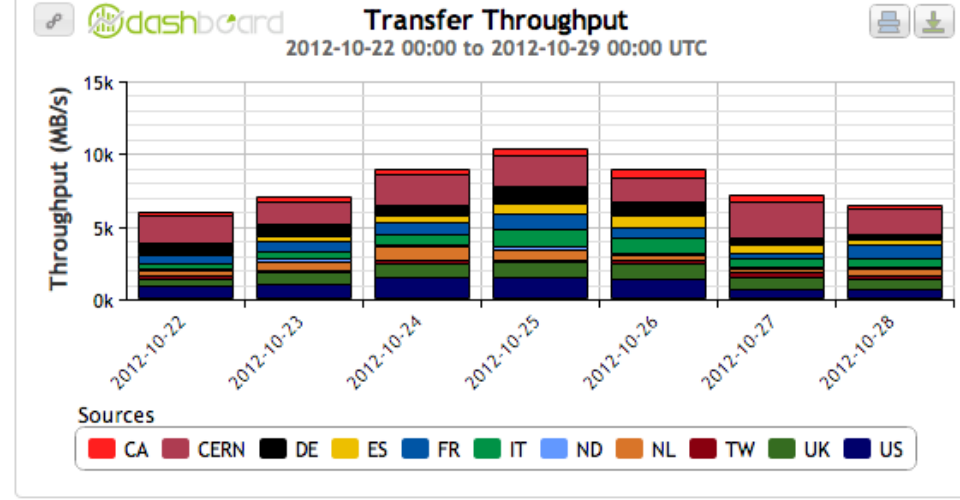

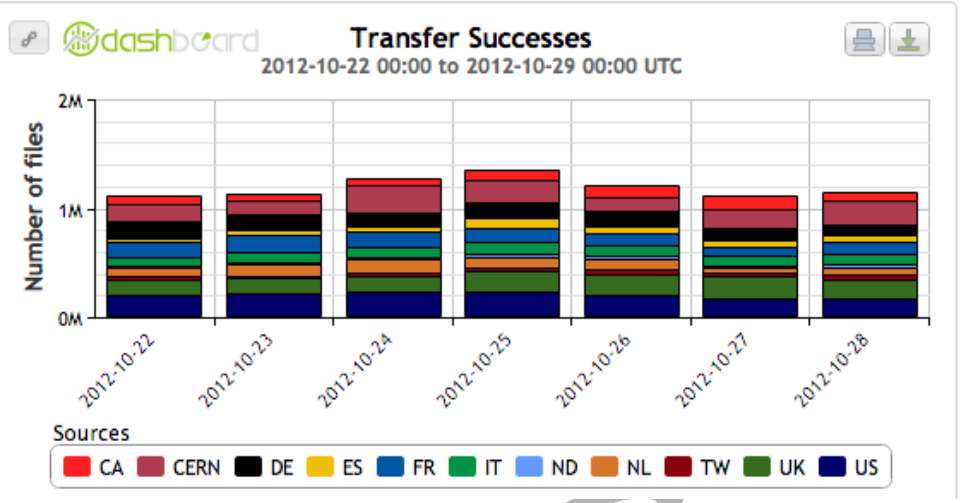

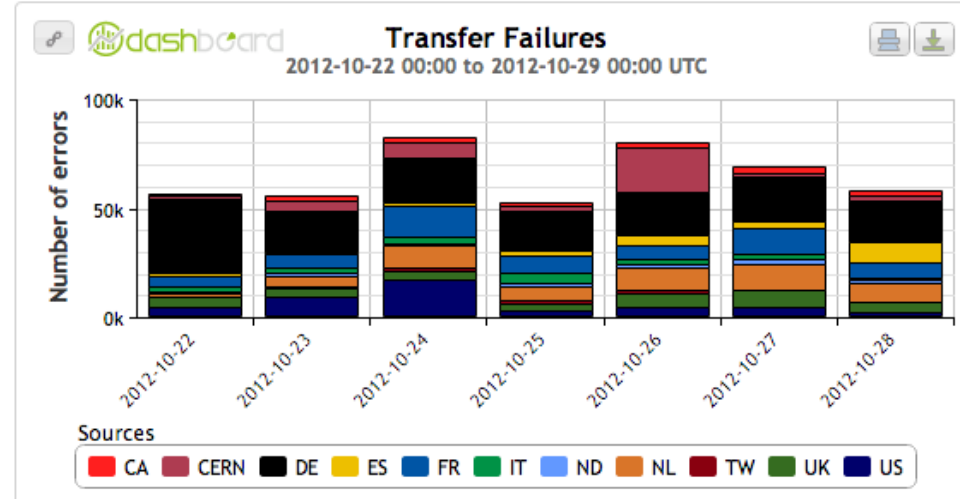

Data transfer - destination

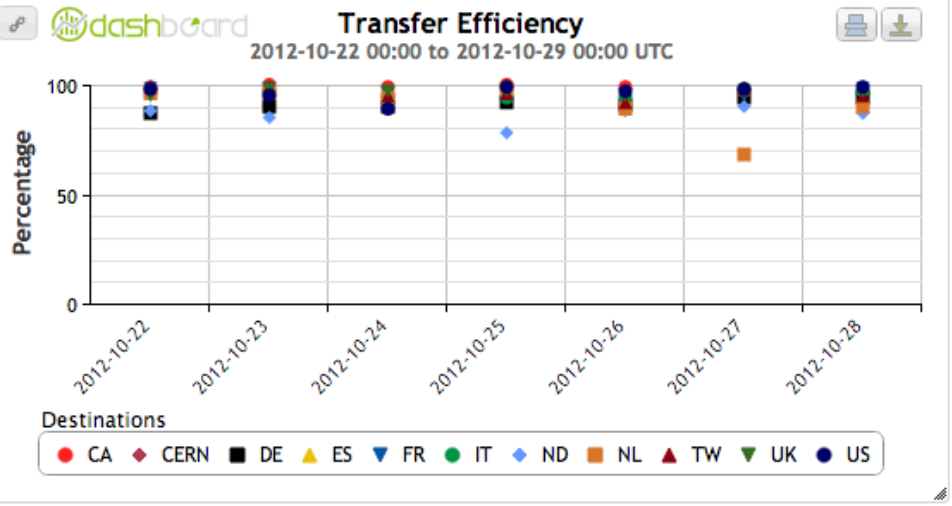

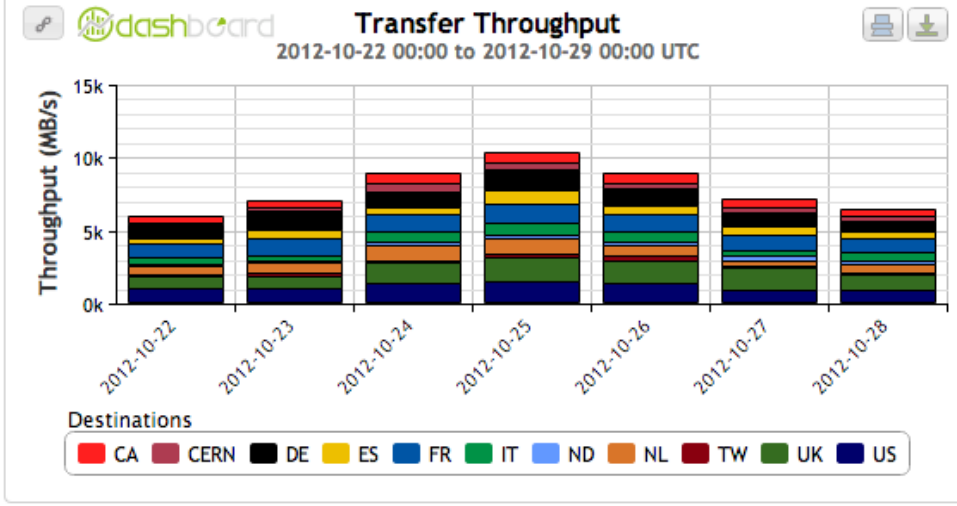

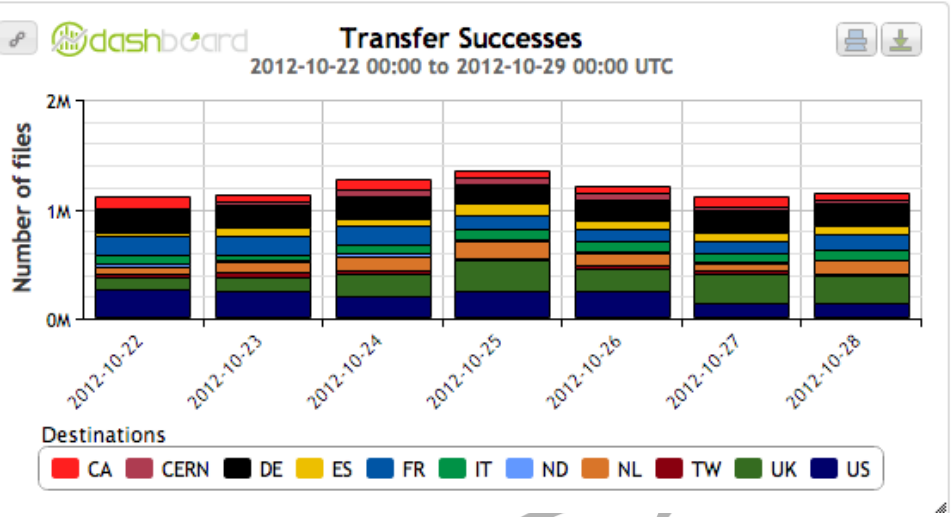

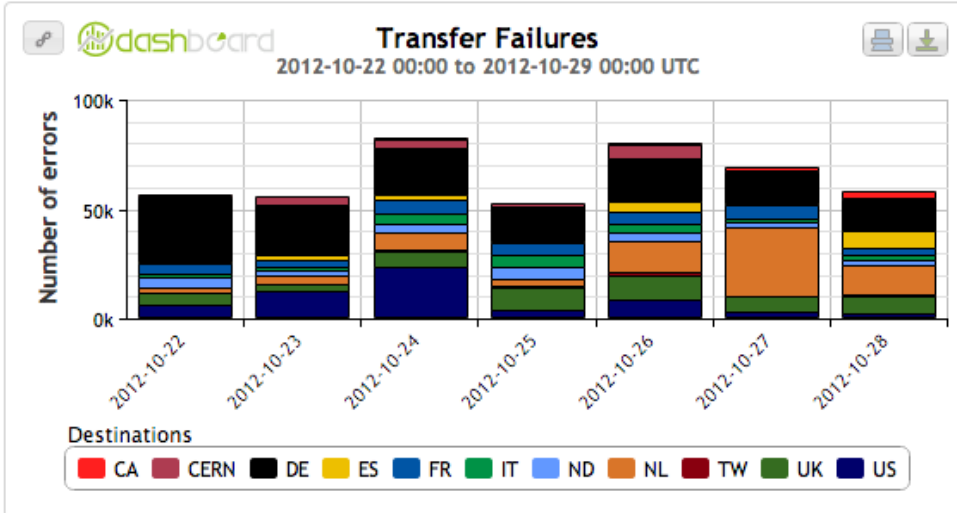

Data transfer - activity

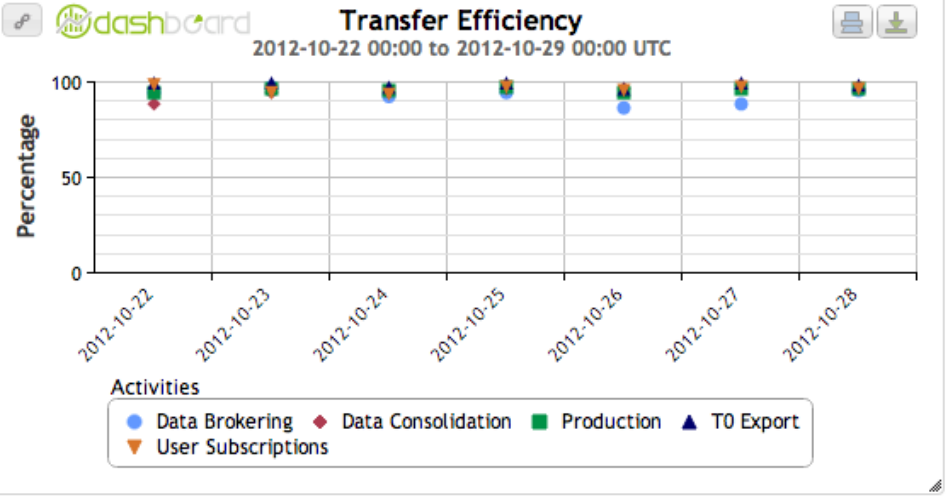

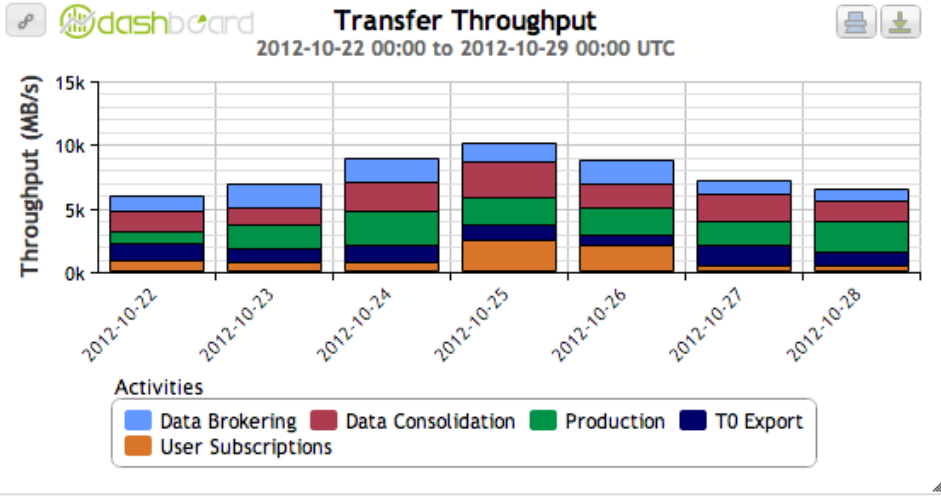

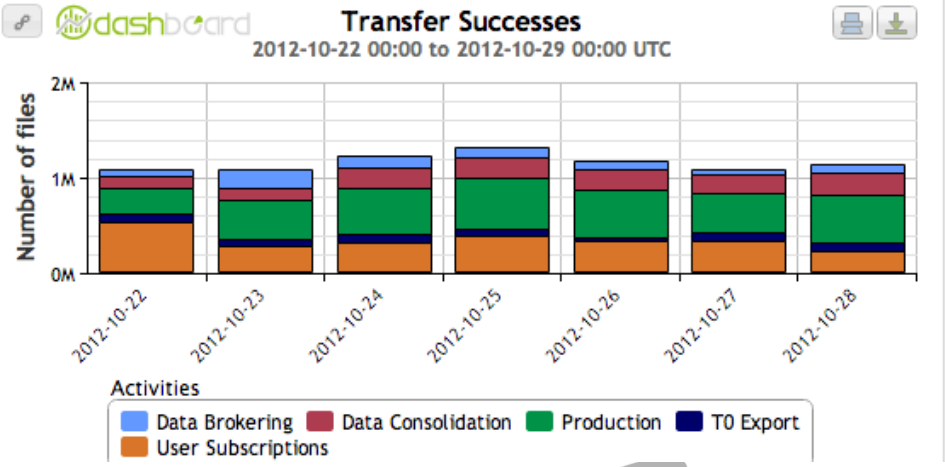

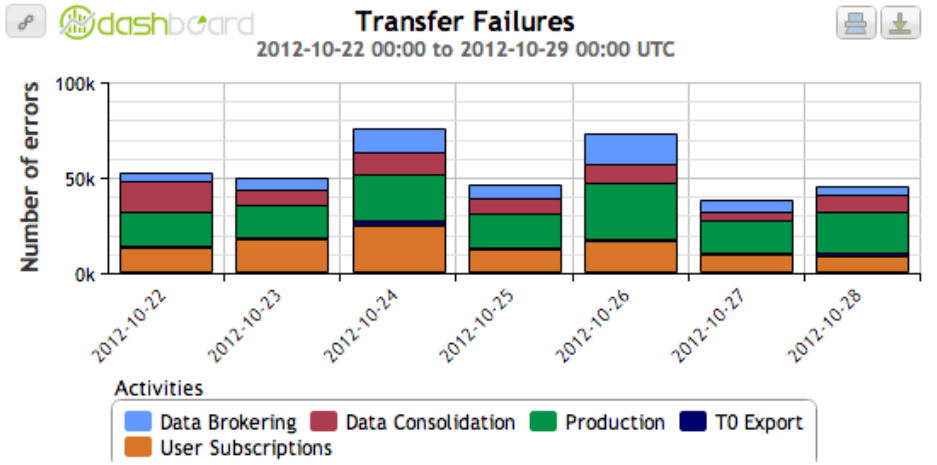

Express reprocessing finished

date	ALL	CA	CERN	DE	ES	FR	IT	ND	NL	TW	UK	US
12-10-14	10	0	0	0	0	0	0	0	10	0	0	0
12-10-15	8469	1420	0	83	81	1916	891	1	1266	1548	1124	139
12-10-16	106572	11829	461	4063	856	17962	15168	1348	12797	15521	13575	12992
12-10-17	56703	2928	0	5213	5121	5781	3375	7106	6382	4316	2764	13717
12-10-18	30887	3383	0	3834	6158	1094	1169	256	3152	3197	1904	6740
12-10-19	4629	0	0	4438	6	1	1	5	8	14	19	137
12-10-20	2344	0	0	2343	0	0	0	0	0	0	1	0
12-10-21	775	0	0	768	0	0	3	0	0	2	2	0
12-10-22	1	0	0	0	0	0	0	0	0	0	1	0
12-10-23	30	0	0	0	0	0	0	0	0	0	30	0
12-10-24	5	0	0	0	0	0	0	0	0	0	5	0
total done	210425	19560	461	20742	12222	26754	20607	8716	23615	24598	19425	33725
total jobs	210449	19560	461	20742	12222	26764	20613	8716	23623	24598	19425	33725
% done	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
aborted	24	0	0	0	0	10	6	0	8	0	0	0
% aborted	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
jobs left	0	0	0	0	0	0	0	0	0	0	0	0
% left	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
running	0	0	0	0	0	0	0	0	0	0	0	0
days left	-	0	0	0	0	0	0	0	0	0	0	0

Ongoing issues

- Retired from ongoing issues in WLCG reporting:
  - mod_gridsite crashing due to access from hosts which don't have reverse DNS mapping. The fix is included in gridsite-1.7.22 and the panda server machines have installed that version. (GGUS:81757)
  - lfc.lfc_addreplica problem when using python 2.6.5. We cannot use python 2.6 in combination with LFC registrations in the pilot until this is fixed. The problem has been resolved and we are in the process of starting a large scale tests. GGUS:84716
- Ongoing:
  - FTS proxy expiration is observed at many FTS in transfers to several sites (GGUS:81844)
    - Ticket going back to May. Patch began deployment at beginning of October, leading to issue (next slide)
  - FZK-LCG2_MCTAPE errors staging from tape (more on this)

Central Services

- 10/19 Slow VOMRS updates to VOMS, ticketed for one user but found to be more general. GGUS:87597
  - Synchronization failing for all vomrs-voms instances at CERN
  - Solved 10/24: linked to an error made when renewing the certificates of CERN VOMRS
- 10/05 Problems reading VOMS attributes in FTS, “delegation ID can’t be generated” GGUS:86775
  - Problem arising from the mentioned patch for the old issue
  - Last week experts found the problem, developed a fix, has been under successful test since late last week
  - Deployment will be taken up this week

Central Services

- CERN-PROD: short-lived problem 10/26 Friday evening with transfer errors to EOS, promptly addressed: "hardware problem with the main headnode, and the namespace crashed. We are failing over to the second box." Turned out to be more software than hardware (we still have a failover box). GGUS:87850

ADC

- MWT2 drained of production jobs at one point, which Rod found to be because MWT2 not getting jobs because #transferring is too high. It peaked at 10k twice in the last week. Brokerage starts when less than max(2*NRun,2000) jobs are transferring.
- Obscure thing for non-experts on brokerage to diagnose, so Tadashi added reporting of this condition to Bamboo logging visible in pandamon

— panda.mon.dev brokerage 10-22 18:35 WARNING MWT2 - too many transferring

Tier 1 Centers

- Transparent CE migrations to EMI CREAM CE at RAL and FZK, and they really were transparent, no problems
- Ditto for TRIUMF Frontier intervention, new server alias, transparent as advertised
- INFN-T1 excluded from Tier 0 export to DATADISK all week (since 10/15), insufficient space
- FZK tape service problems throughout the week. Out of T0 export all week; added back yesterday. GGUS:87510, 87526, 87630
- 10/24 FZK DATADISK urgently needed files unavailable, dCache pools went offline with a storage system problem, ~330TB unavailable (had been for a few days), recovering for the remainder of the week
  - List of lost files delivered yesterday. ~21k files, 17TB.
  - Still working on recovering remaining 34TB as of yesterday.

Tier 1 Centers

- 10/21 FT files apparently removed by ATLAS DDM at IN2P3? (We never did manage to have the 9am discussion on this we planned.) (cf. shifter question at end) GGUS:87635
- SARA-MATRIX: SRMV2STAGER errors, ticket from last week 10/18. Files appeared to copy to disk successfully but not reported as a successful transfer. GGUS:87531
  - "We're not sure what the 'Request clumping limit reached' means so we'll ask dCache support."
  - Asked for updates through the week. Want to be sure all is OK in advance of the start of bulk reprocessing. GGUS:87531
  - Reported yesterday: "The dCache developers think that there might be a bug involved here. We will keep you posted."

Tier 1 Centers

- 10/25 TAIWAN-LCG2 DATATAPE: File unavailable. Reported as due to a stuck tape mount. Resolved 10/29, files accessible again. GGUS:87796
- 10/18 IN2P3-CC: failure staging data from tape. Two files reported lost from a tape. Asked for definitive list of all lost files. 10/28 received confirmation that only the two files from the tape are lost, others moved to a new tape. GGUS:87529
- 10/29 IN2P3-CC: SRM failure early Mon morning, T0 export failure, raised to ALARM after 3hrs but addressed and resolved shortly thereafter by SRM restart. Briefly out of T0 export. GGUS:87872

Tier 1 Centers

- 10/14 IN2P3-CC: longstanding issue of AsyncWait error on source transfers, recurred, ticket updated. GGUS:87344
  - Explanation from site, repeating in part an earlier explanation:
    - If you have a source file which initially is in your permanent pool, then when the 1st FTS transfer request is executed, a disk to disk copy of the file is initiated. Normally the TURL for the "export" copy of the file is not returned until the transfer has fully finished. For slow/large files this is taking too long and the timeout on the FTS aborts the FTS request. However if the file is fully copied into the export pool, when the next request from the FTS to transfer the file is received, since the file is already in the export pool; the time to return the TURL for the transfer is short and so succeeds.
    - I notice form ATLAS monitoring that even when the number of failures is high with old settings, there was still a high through of successful transfers between in2p3 and RHUL. When reducing the channel settings, the failure rate has reduced, but so has the throughput rate.
  - (So is this internal staging transfer particularly slow at IN2P3 or something?)

Tier 1 Centers

- 10/28 Precautionary word from BNL that Hurricane Sandy likely to impact BNL operations [BNL closed Monday morning and is still closed]
- 10/29 BNL reported after review and contact with National Hurricane Center that decision taken to keep the Tier 1 operating, with the warning that the facility could go down abruptly in the event of persistent power loss. [So far so good, and the worst is past.]

Other

- Good question from Andrzej Olszewski, comp@p1:
  - Quite often some files used in functional tests are not available at some sites. Example is IL-TAU-HEP with GGUS ticket for this issue: GGUS:87849. I am not sure if in the answer to the question from site admin I should insist on declaring these (temporary I think) file as lost, or are they going to be resent to the site in the next tour of FT and thus we should just wait patiently?

Other

- A interesting comment from the Friday WLCG daily on difference between GGUS and SNOW tickets, in response to a question from CMS as to why GGUS tickets seem to be dealt with more promptly [my paraphrase]:
  - GGUS tickets: everyone responsible for a service that might be involved in an issue sees it right away.
  - SNOW tickets: when a GGUS ticket response filters down to the appropriate group, it is entered as a SNOW ticket and taken up. So there's an inherent delay in those extra steps.
  - There is also a perception that GGUS tickets are on average more important, so they stick out more and attract more attention.

Thanks

- Thanks to Alexei Sedov for his help as a shadow AMOD shifter and to all shifters and helpful experts!