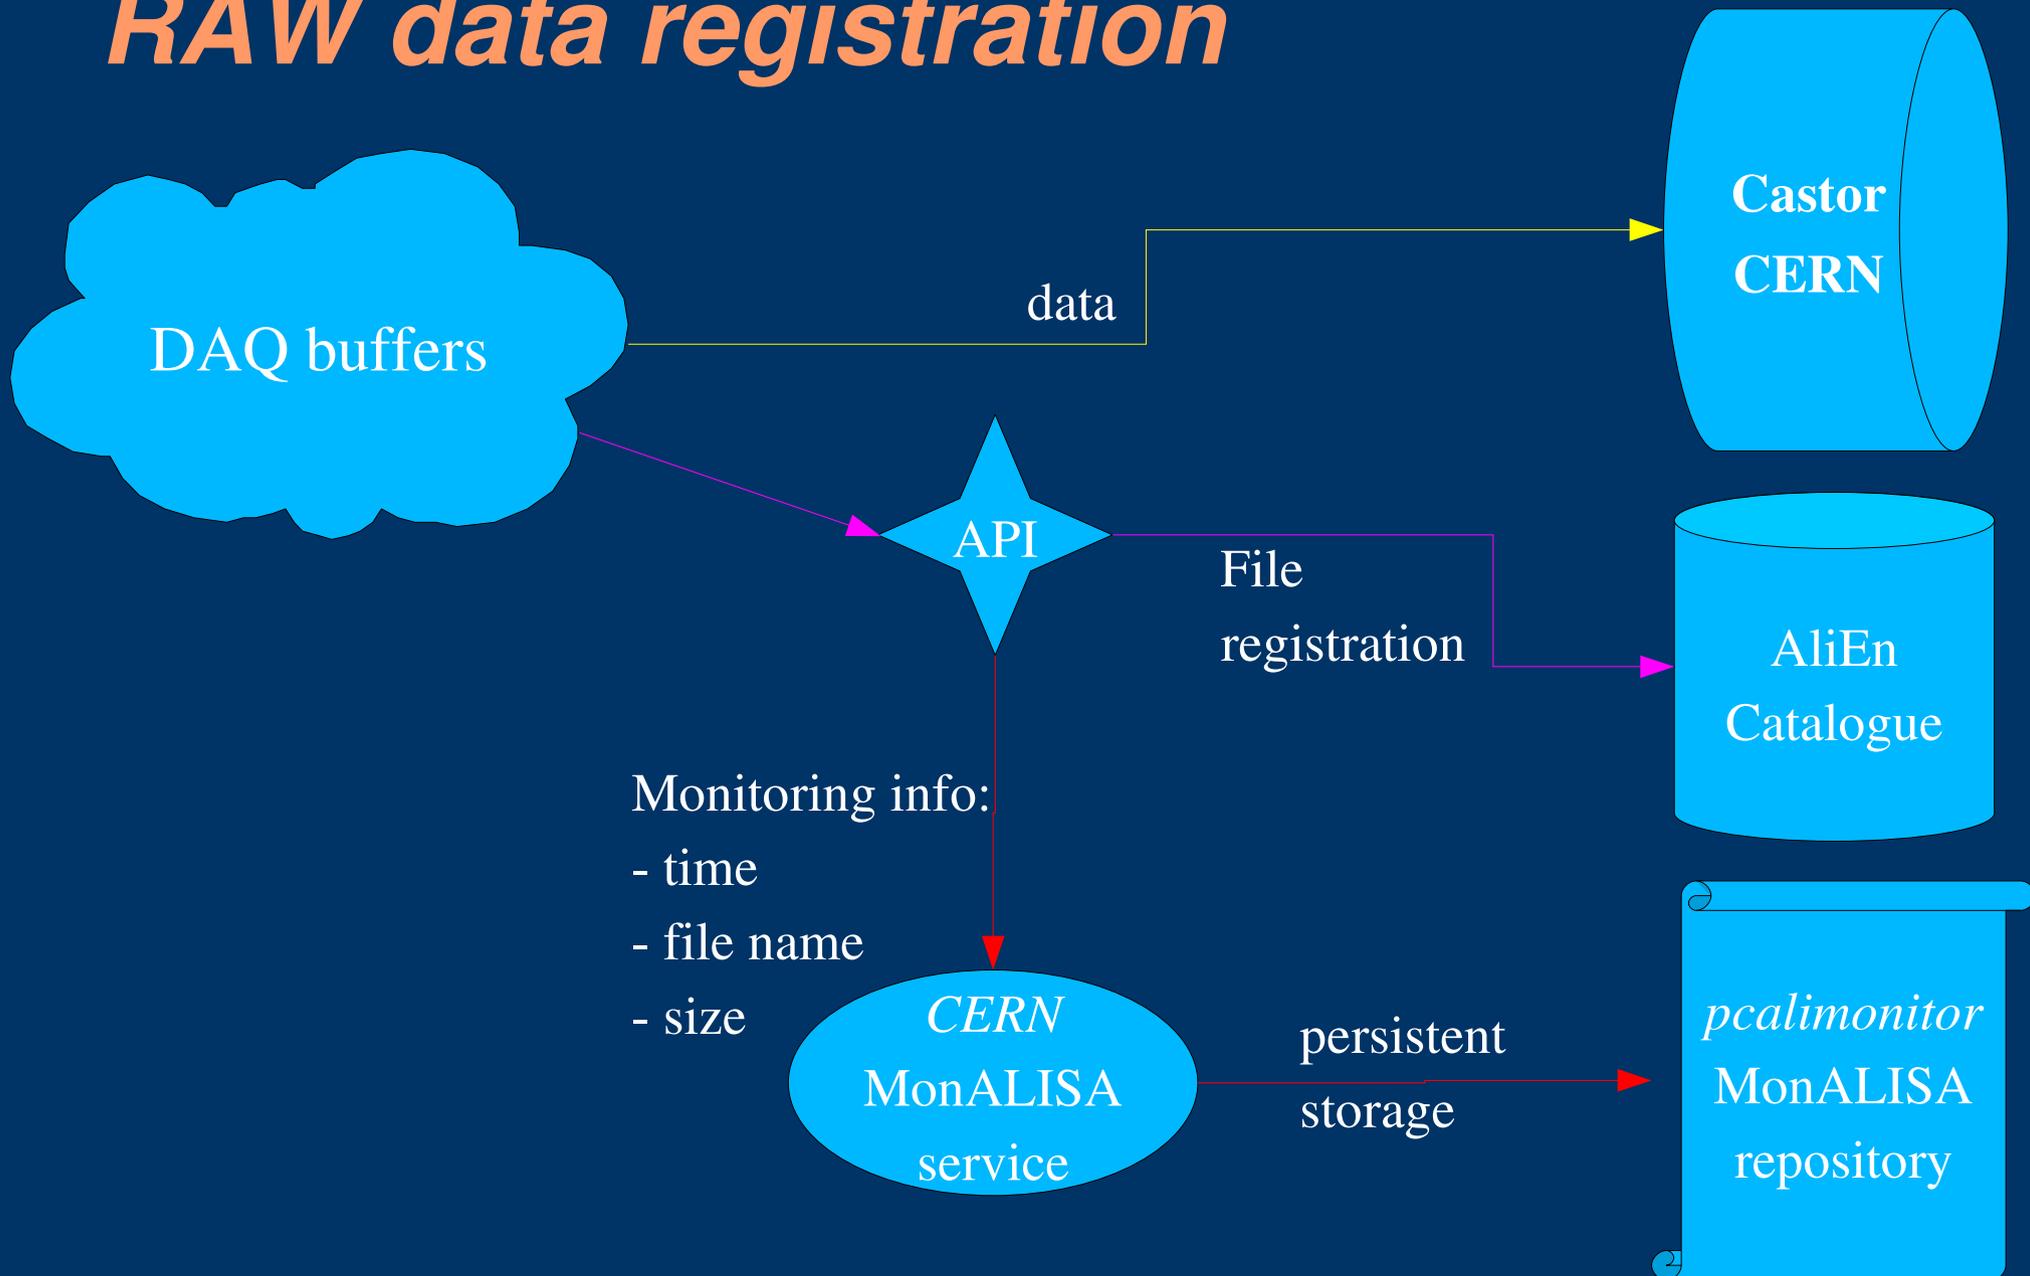


RAW data registration and replication

Costin Grigoras <grigoras@cern.ch>

ALICE Offline week, 04.2008

RAW data registration



Shifters' dashboard

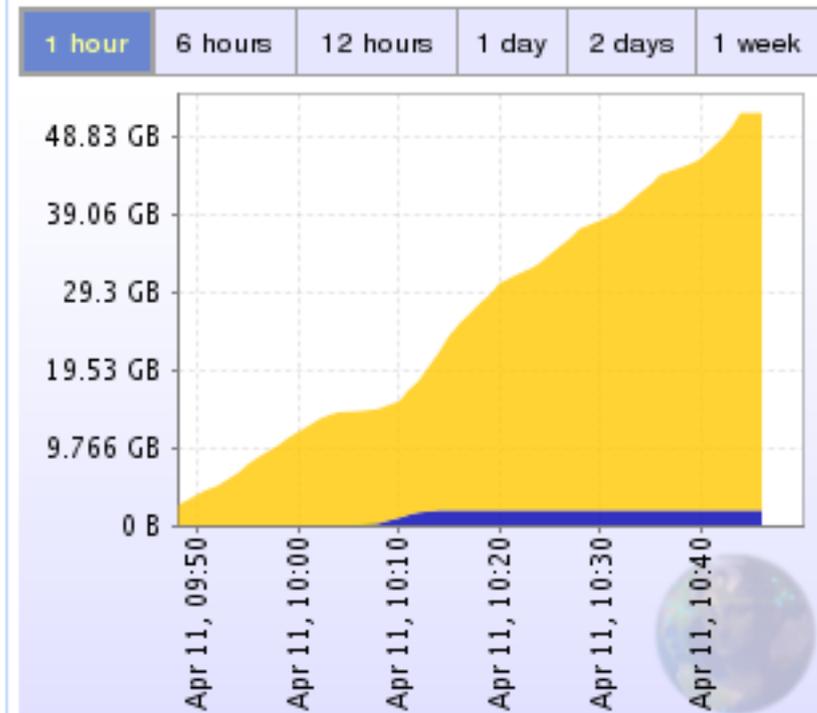
<http://pcalimonitor.cern.ch/dashboard/>

RAW data registration , last runs in the catalogue:

Run #	Partition	Chunks	Last time	Size
22226	TDSMtest	1431	today 10:44	1014 GB
26986	LHC08a_TPC	3	today 10:09	1.814 GB
26985	LHC08a_TPC	2	today 10:09	59.1 KB
26795	LHC08a_TPC	2	today 10:09	59.1 KB
26792	LHC08a_TPC	2	today 10:09	18.64 MB
22225	TDSMtest	36	yesterday 13:24	23.51 GB
22224	TDSMtest	18	yesterday 13:00	10.64 GB
22223	TDSMtest	9	yesterday 12:38	3.869 GB
26808	LHC08a_HMPID	2	yesterday 10:56	287.9 MB
26804	LHC08a_HMPID	2	yesterday 10:47	674.8 MB

[Full details »](#)

RAW data registration



[Full details »](#)

84TB registered in Feb-Mar run

Run details

<http://pcalimonitor.cern.ch/DAQ/>

RAW Data Registration, Transferring and Processing

Run#	Partition	Chunks	Total size	First seen	Last seen	Transfer status	Processing status	Staging status	
25744-257	- All -			- All -		- All -	- All -	- All -	»
25755	LHC08a_HMPID	6	4.219 GB	07 Mar 2008 10:42	12 Mar 2008 14:48	ALICE::NDGF::DCACHE		Stage	<input type="checkbox"/>
25744	LHC08a	2503	2.226 TB	07 Mar 2008 09:51	07 Mar 2008 22:47	ALICE::NDGF::DCACHE	2416 / 2503 (96%)	Stage	<input checked="" type="checkbox"/>
25766	LHC08a	6	2.037 GB	07 Mar 2008 15:12	07 Mar 2008 15:16	ALICE::GSI::SE	3 / 6 (50%)	Stage	<input type="checkbox"/>
25758	LHC08a	4	57.65 MB	07 Mar 2008 15:16	07 Mar 2008 15:16	ALICE::NDGF::DCACHE	2 / 4 (50%)	Stage	<input checked="" type="checkbox"/>
25764	LHC08a	6	691.4 MB	07 Mar 2008 14:20	07 Mar 2008 14:22	ALICE::NDGF::DCACHE	3 / 6 (50%)	Stage	<input checked="" type="checkbox"/>
25763	LHC08a	6	13.05 MB	07 Mar 2008 14:03	07 Mar 2008 14:05	ALICE::NDGF::DCACHE		Stage	<input type="checkbox"/>
25757	LHC08a	4	554.5 KB	07 Mar 2008 10:40	07 Mar 2008 13:55	ALICE::NDGF::DCACHE	1 / 4 (25%)	Stage	<input checked="" type="checkbox"/>
25765	LHC08a_TO	2	444 MB	07 Mar 2008 11:52	07 Mar 2008 11:52	ALICE::NDGF::DCACHE		Stage	<input type="checkbox"/>
25762	LHC08a_TO	2	358.9 MB	07 Mar 2008 11:28	07 Mar 2008 11:28	ALICE::NDGF::DCACHE		Stage	<input type="checkbox"/>
25753	LHC08a_HMPID	2	874.3 MB	07 Mar 2008 10:41	07 Mar 2008 10:42	ALICE::NDGF::DCACHE		Stage	<input type="checkbox"/>
25749	LHC08a_HMPID	8	5.857 GB	07 Mar 2008 10:13	07 Mar 2008 10:41	ALICE::NDGF::DCACHE		Stage	<input type="checkbox"/>
25754	LHC08a	4	1.718 MB	07 Mar 2008 10:34	07 Mar 2008 10:39	ALICE::NDGF::DCACHE	2 / 4 (50%)	Stage	<input checked="" type="checkbox"/>
25752	LHC08a	4	315.4 KB	07 Mar 2008 10:32	07 Mar 2008 10:32	ALICE::NDGF::DCACHE		Stage	<input type="checkbox"/>
25748	LHC08a_HMPID	2	1.254 MB	07 Mar 2008 10:13	07 Mar 2008 10:13	ALICE::NDGF::DCACHE		Stage	<input type="checkbox"/>
14 runs		2559 files	2.24 TB			0 0 14	8 6 0	14 0 0	

Flag run 25744 for analysis

LTM (Lightweight Transfer Manager)

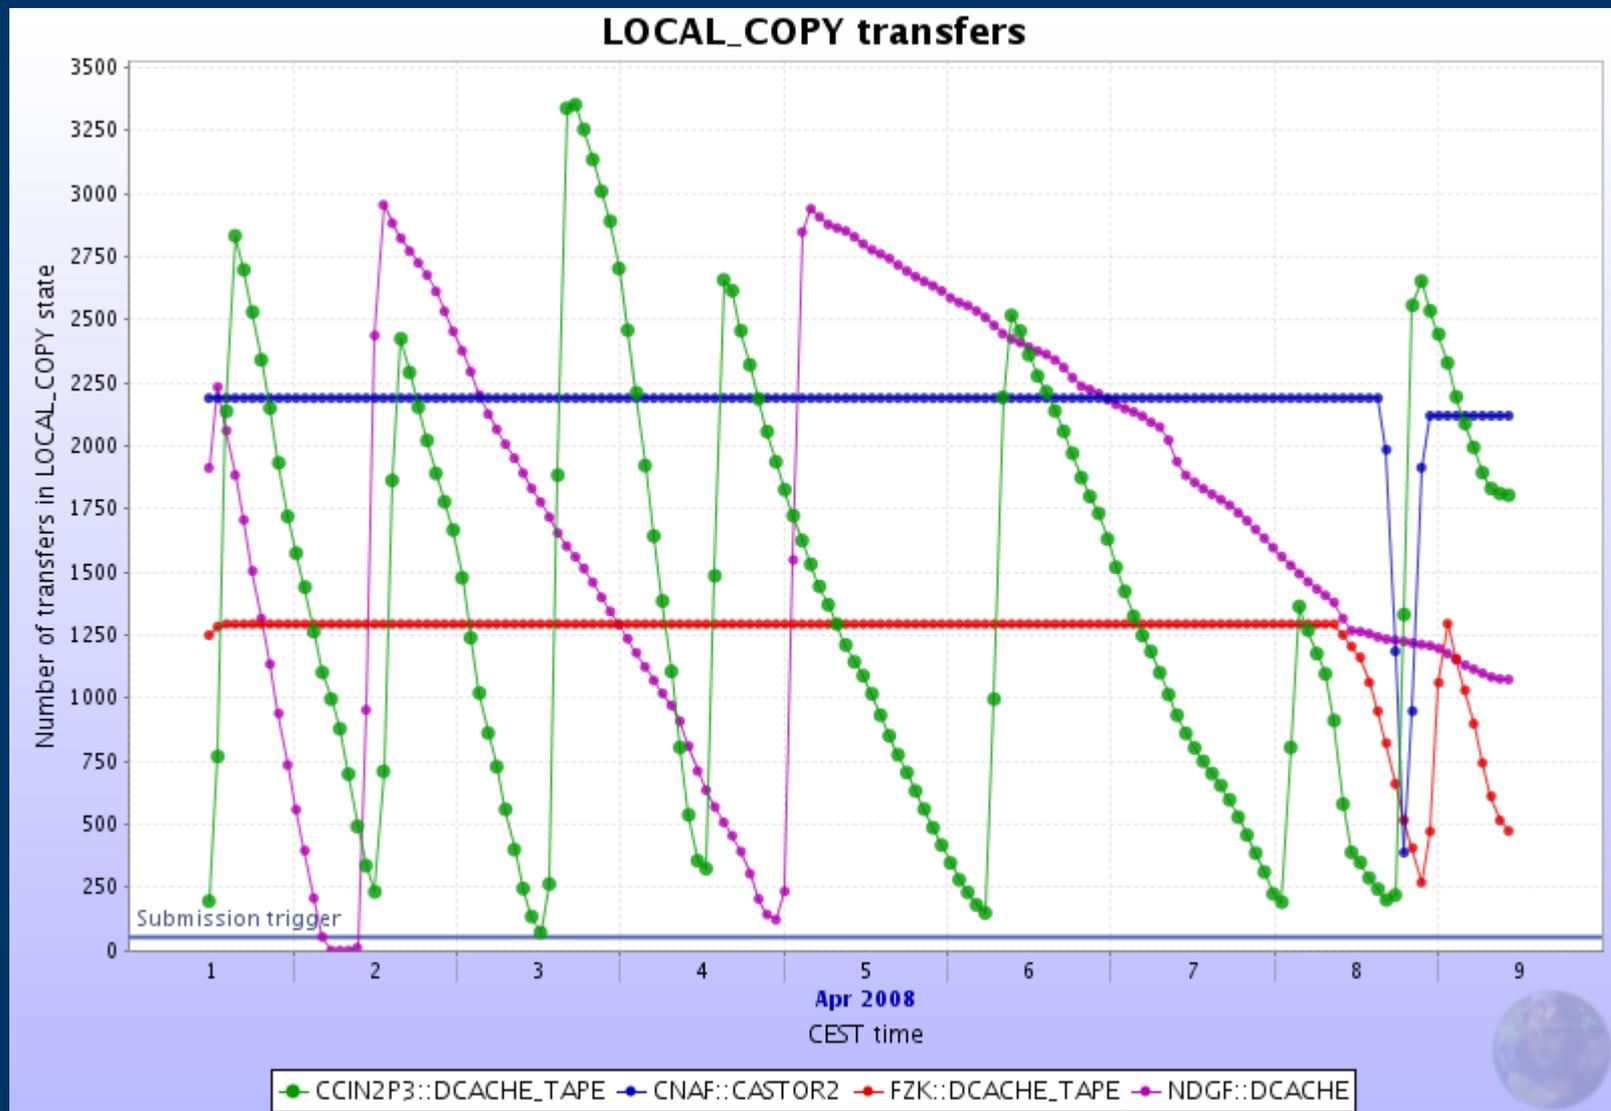
- ML repository watches the transfer queues to each of the defined targets and, when needed, resubmits incomplete runs until completion or submits newly registered runs to the SE that triggered the action

RAW data replication targets				
SE name	Runs	Files	Size	Options
ALICE::CCIN2P3::DCACHE	59	208	55.2 GB	Edit Enable
ALICE::CCIN2P3::DCACHE_TAPE	670	28,939	23.97 TB	Edit Disable
ALICE::CNAF::CASTOR2	155	5,899	4.886 TB	Edit Disable
ALICE::FZK::DCACHE_TAPE	532	18,532	15.24 TB	Edit Disable
ALICE::FZK::LOG	68	203	63.04 GB	Edit Enable
ALICE::GSI::SE	133	2,935	693.3 GB	Edit Enable
ALICE::NDGF::DCACHE	129	13,292	11.48 TB	Edit Disable
ALICE::SARA::DCACHE	202	1,418	752.8 GB	Edit Enable
ALICE::SARA::DCACHE_TAPE	290	6,695	5.444 TB	Edit Disable
5 active, 4 disabled	2,238	78,121	62.55 TB	Add SE

LTM Settings	
LTM Status:	ACTIVE
(Re)submission trigger:	50
Stop after submitting:	3000

[Edit LTM Settings](#)

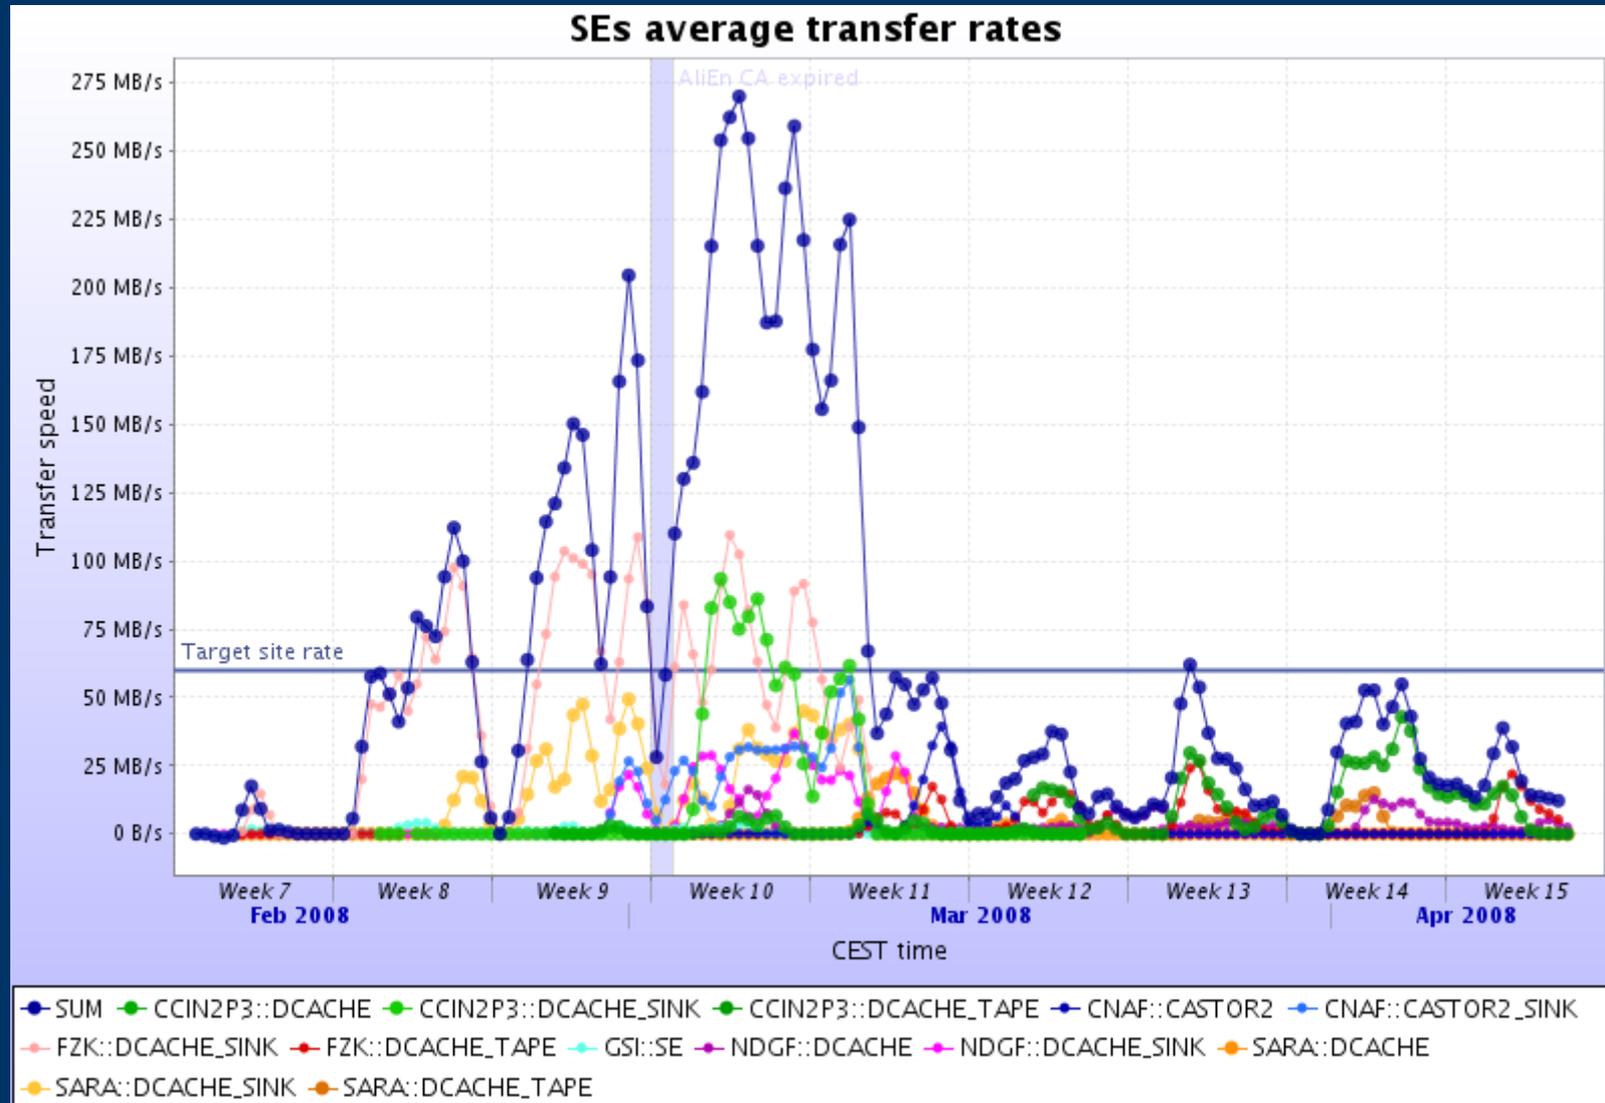
LTM in action



Data transfers

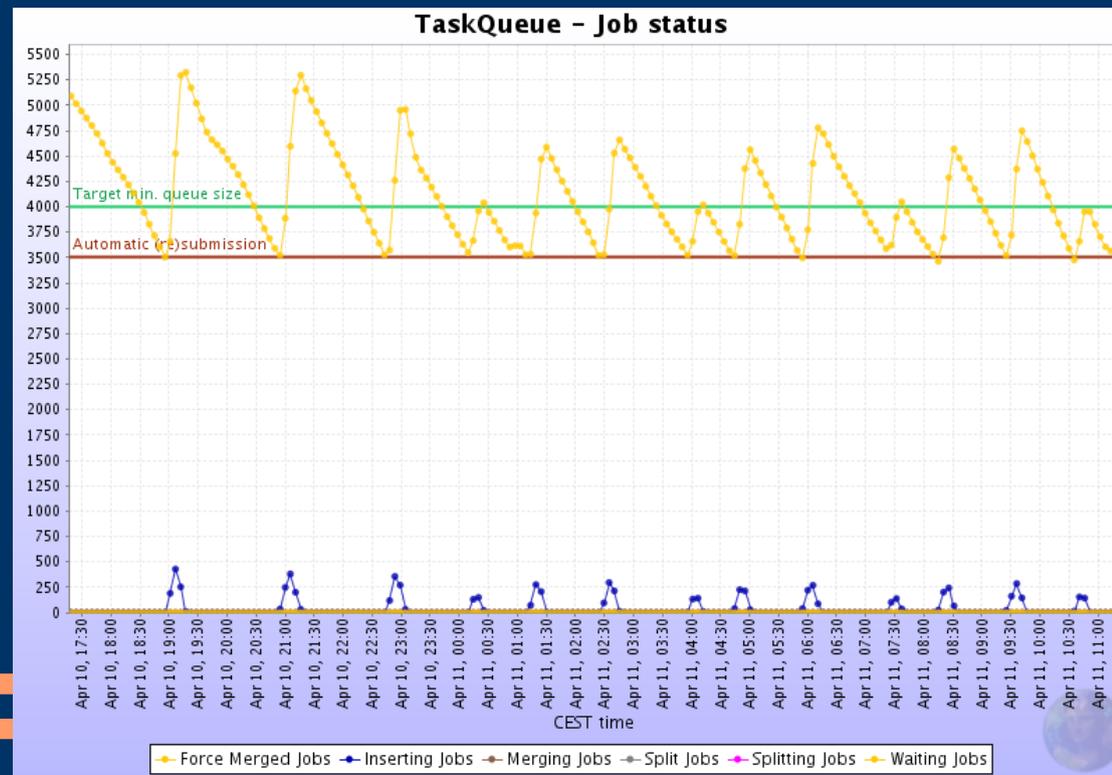
- LTM also took care of GSI's requests for particular runs:
<http://wiki.gsi.de/cgi-bin/view/Grid/TransferRequestsToGsi>
 - We transferred ~300TB to T1s + GSI during the last 2 months at an overall average speed of 60MB/s.
 - The difference in size is due to dummy data sent to “sink” SEs, in the attempt to fill all the resources during the last exercise.
-
-

Data transfers



LPM (Lightweight Production Manager)

- Just like LTM, keeping an eye on the job queue it resubmits error jobs or submits new jobs when necessary
- It treats separately *aliproduct* and *alidaq*
 - *aliproduct* jobs run numbers are just incremental counters
 - *alidaq* jobs are executed on the manually selected runs (See Marco's presentation after the break)



Staging

- Flagged runs are scheduled to be staged, using Andreas' Castor mass staging code.
 - After the initial staging request, all the flagged runs are periodically (once per week) re-requested, just to keep them on disk in case the files were not touched in the mean time.
-
-

Plans

- Enlarging the base of T2 storages, FTD at NIHAM and Legnaro are being debugged, Prague is next on the list.
 - T2-T2 transfers, and we started with Subatech-Legnaro.
 - Fully implementing the transfers on demand to T2 sites, LTM needs a few improvements + a user interface where the requests would be made.
-
-