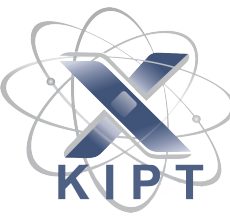




Support and development of KIPT Tier-2 center for processing of data from the Large Hadron Collider

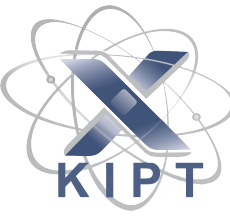
*O. Bunetskiy, K. Klimenko, L. Levchuk, S. Lukyanenko,
D. Soroka, P. Sorokin, A. Pristavka*

NSC Kharkov Institute of Physics and Technology, 61108, Kharkov, Ukraine



Plan

- History;
- The CMS Computing Model;
- T2_UA_KIPT structure;
- Computing facilities;
- Jobs statistic and monitoring;
- Mass Storage System;
- PhEDEx data transfers;
- Plans and prospects.

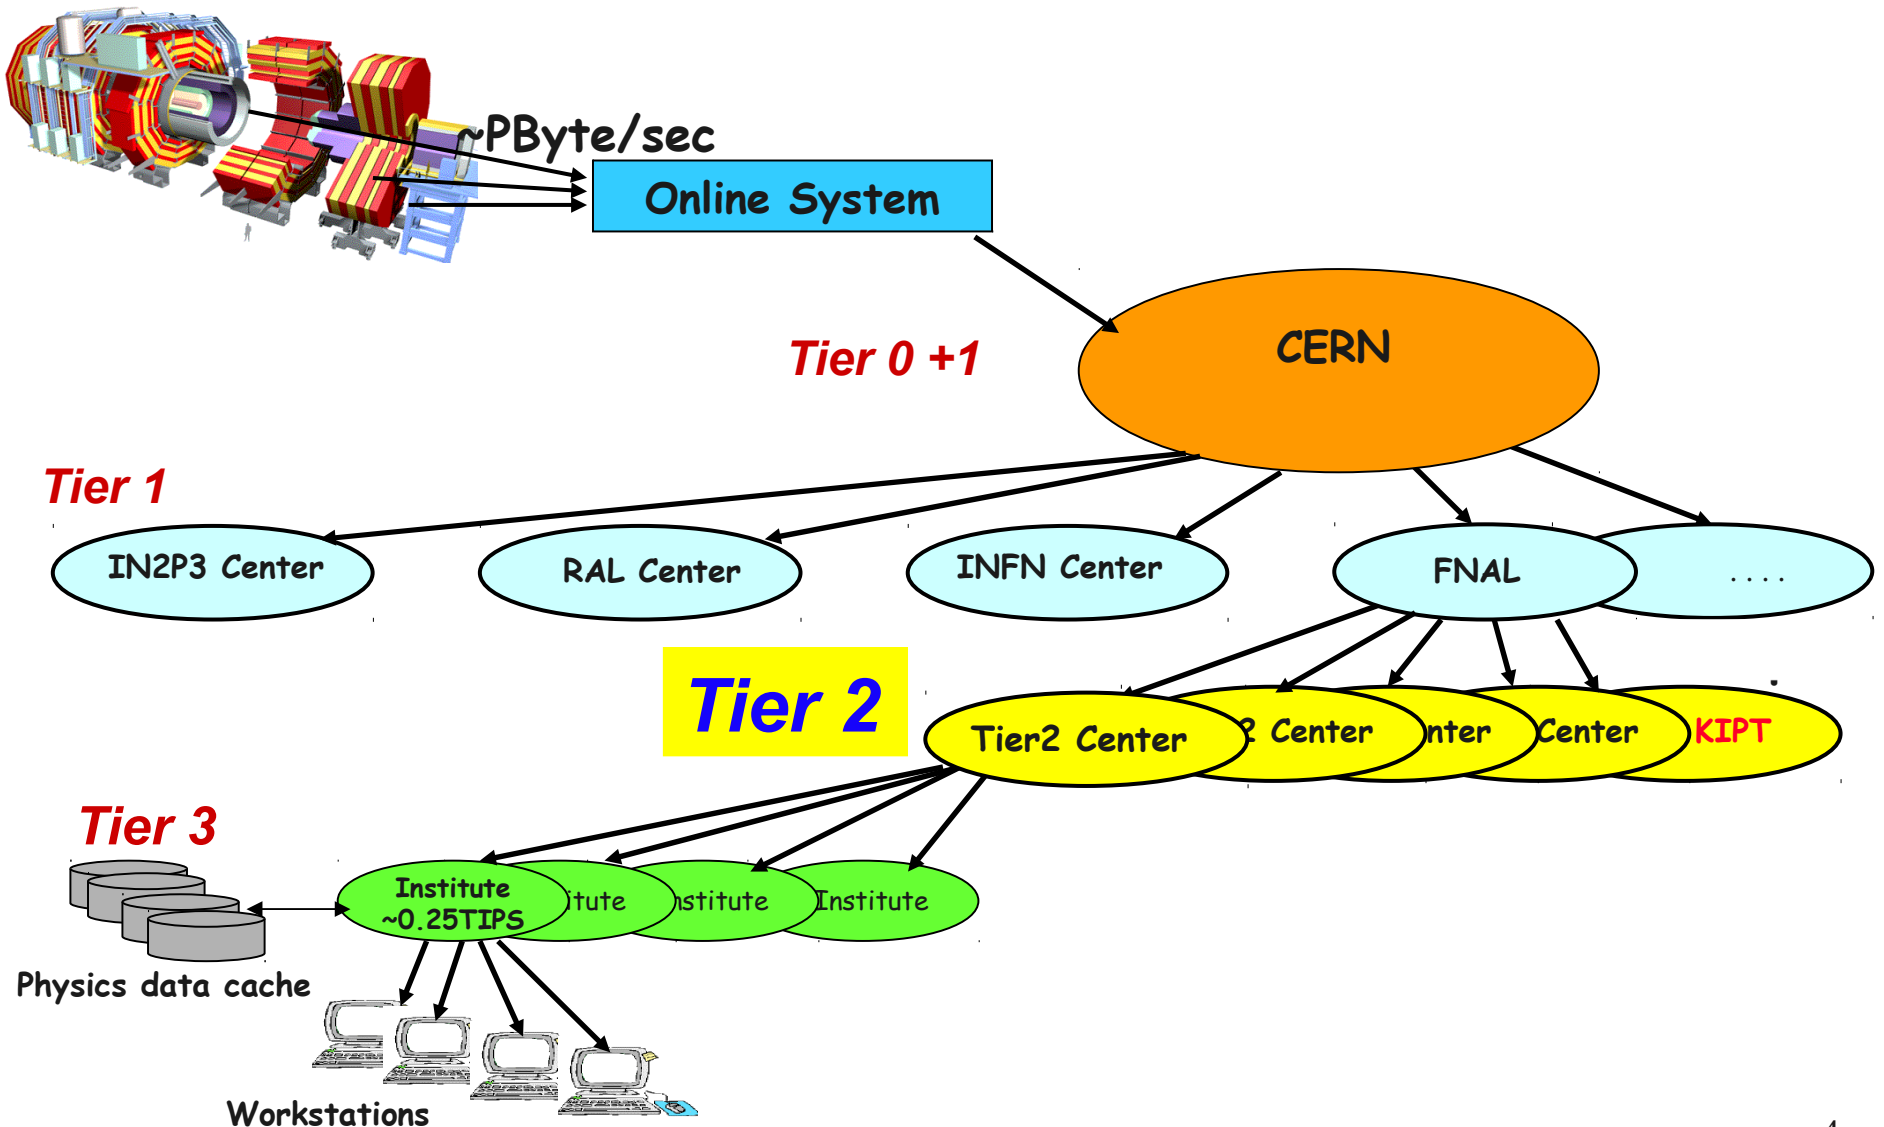
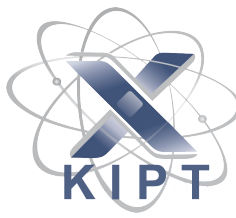


History

- 2001** – start of work (5 PIII nodes)
- 2002** – start of CMS MC production at KIPT
- 2005** – registration in WLCG as **Kharkov-KIPT-LCG2**
- 2008** – registration in CMS SiteDB as **T2_UA_KIPT**
- 2009** – commissioning of **T2_UA_KIPT** in CMS Grid infrastructure
- 2010** – association of **T2_UA_KIPT** with EWK PAG

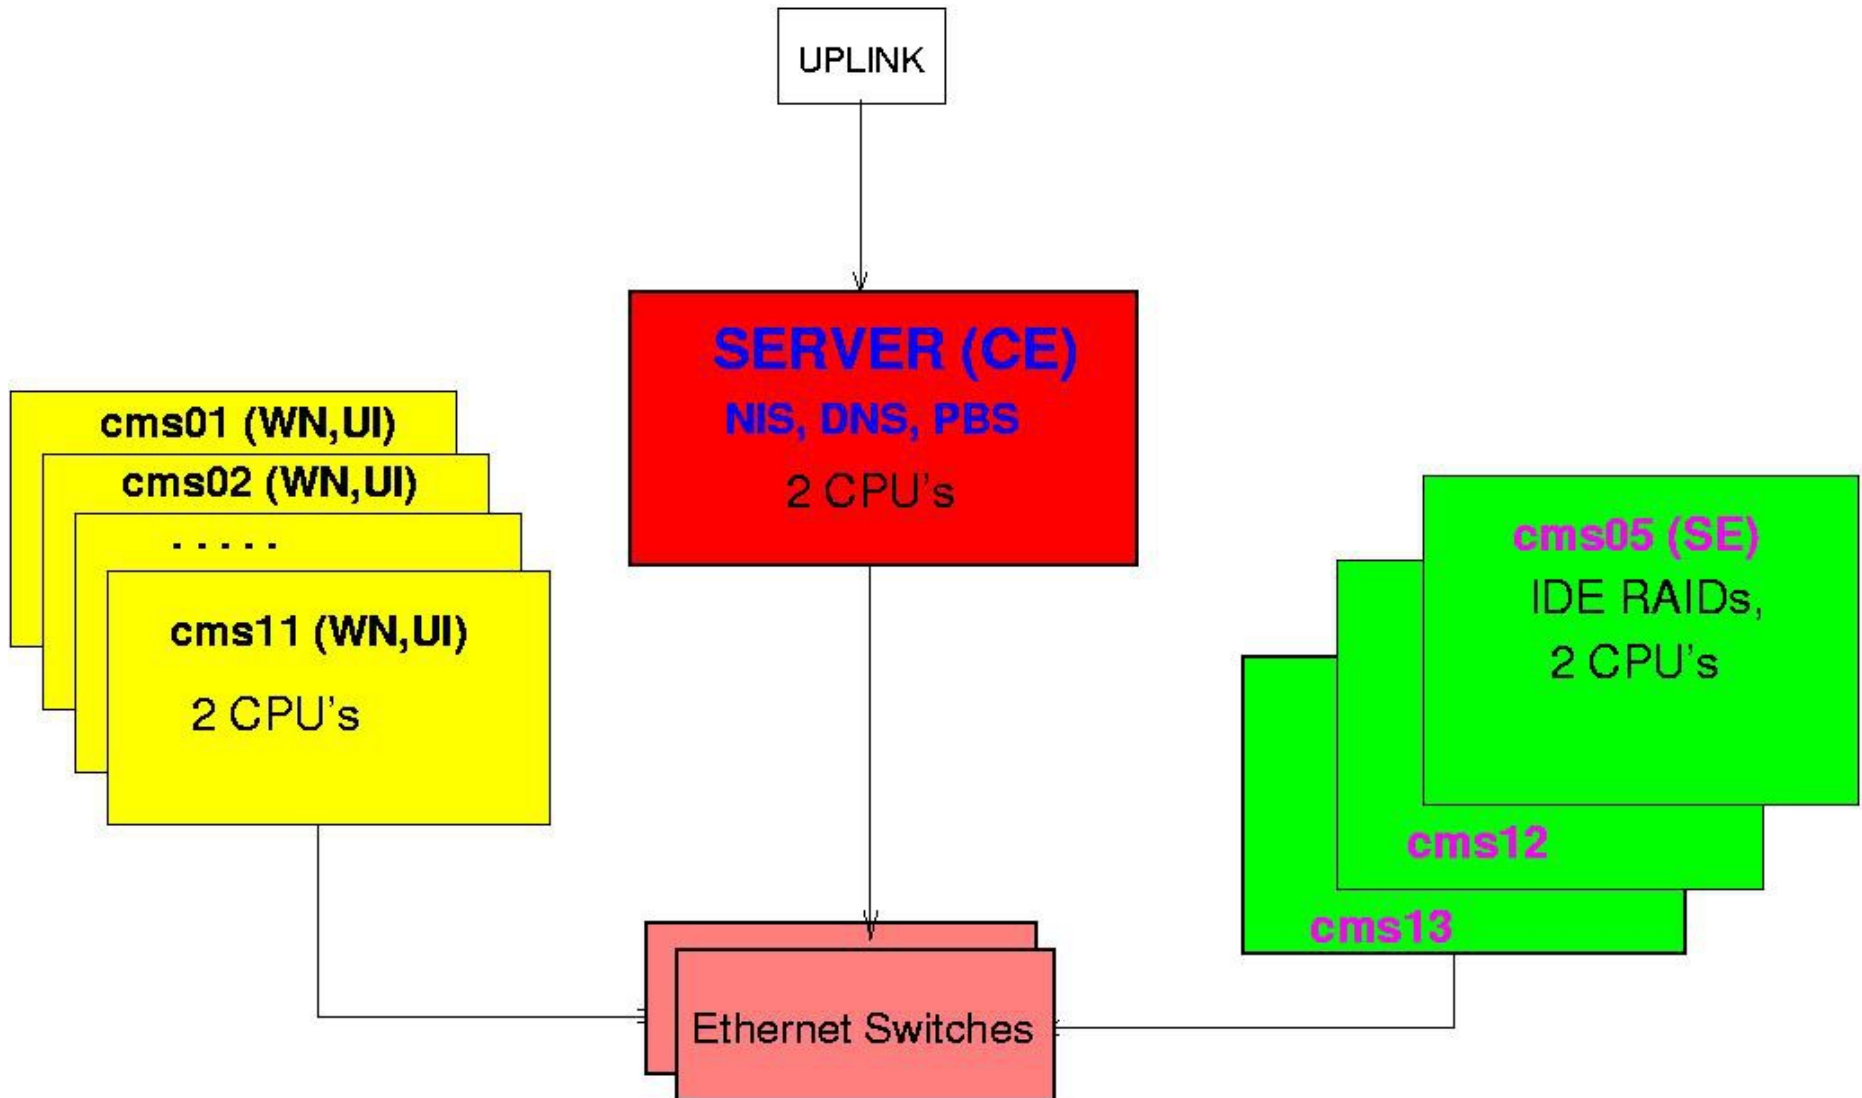
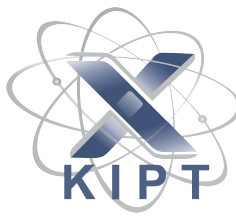


The CMS Computing Model



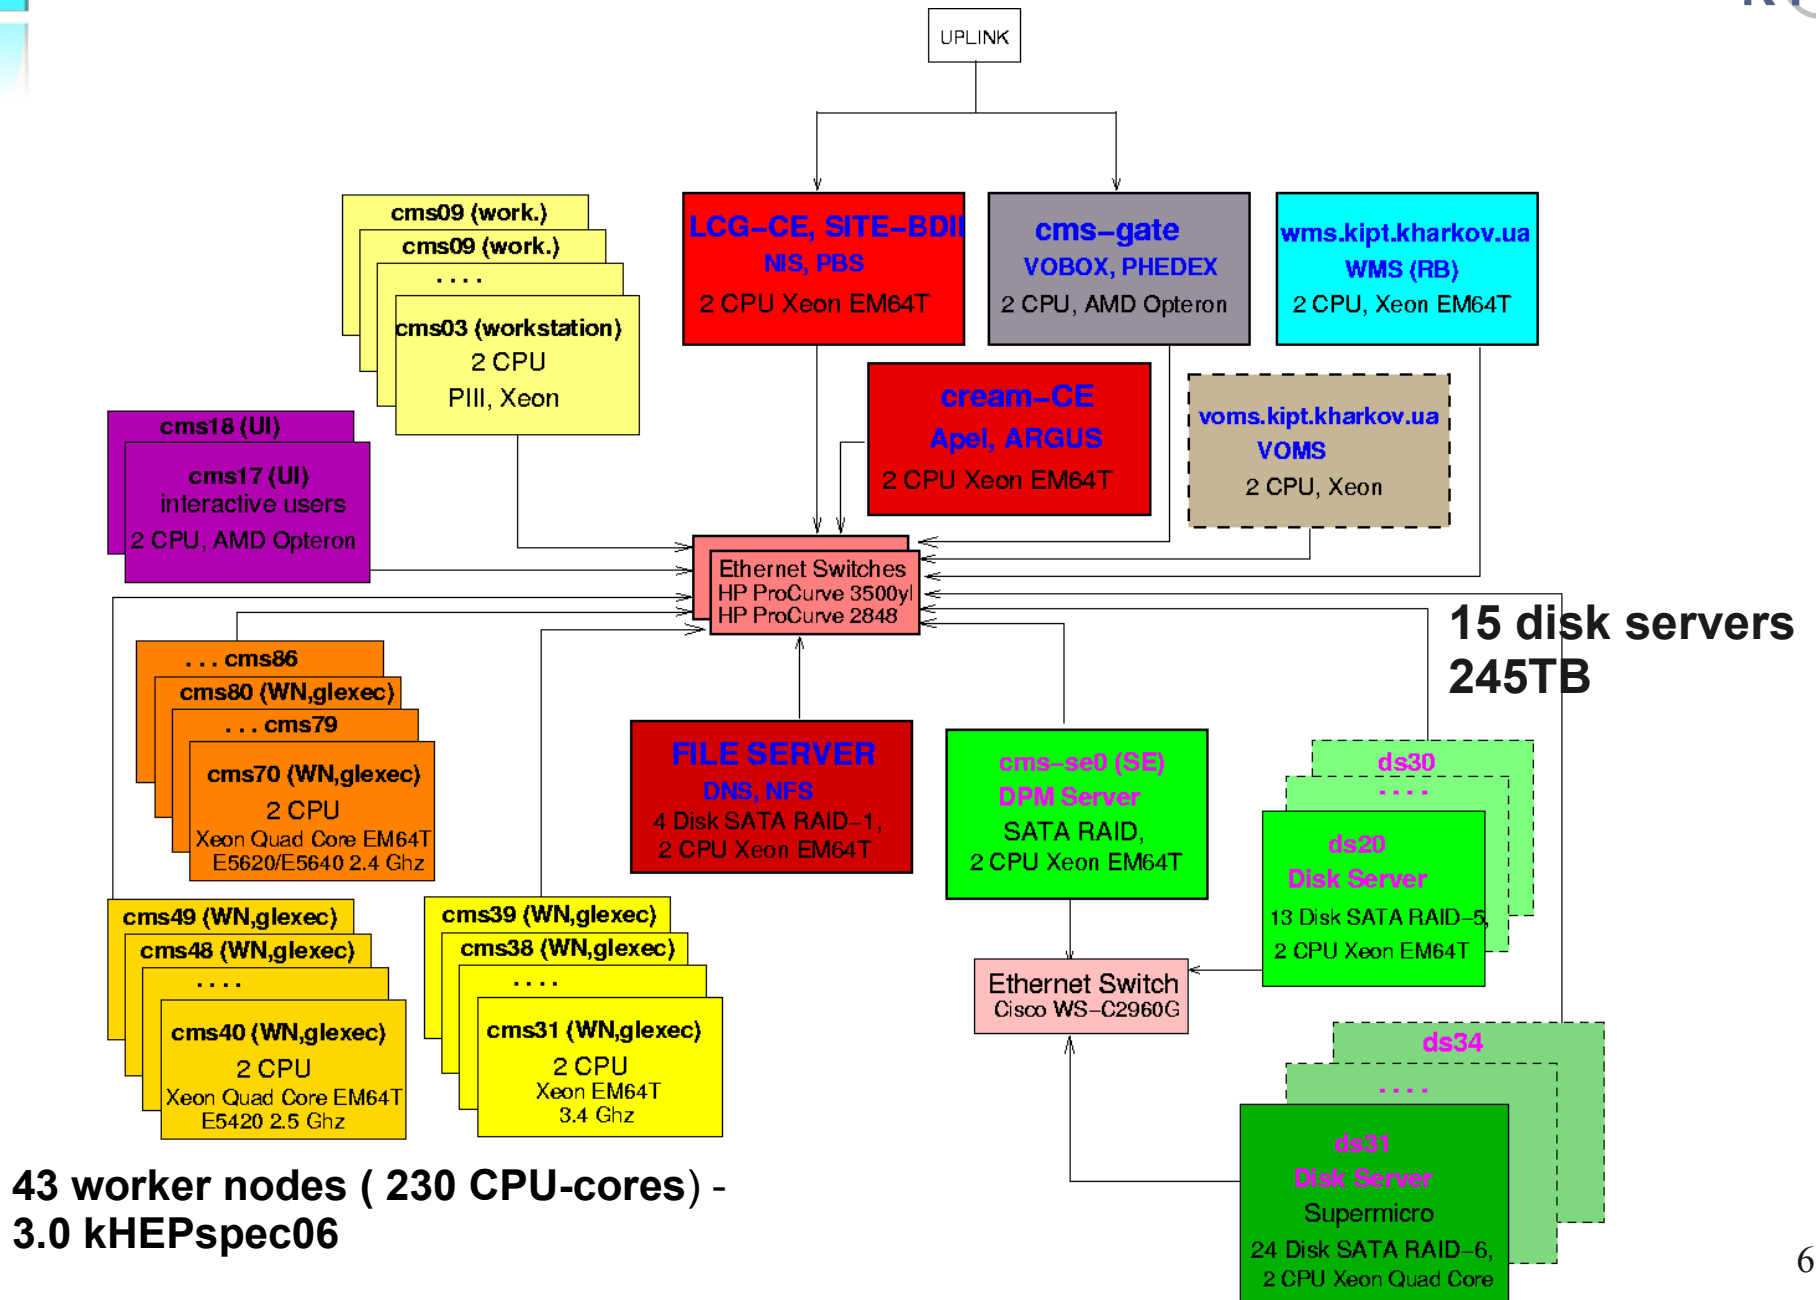
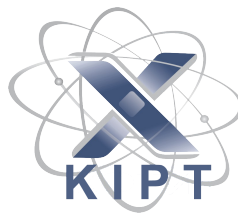


T2_UA_KIPT structure in 2005





T2_UA_KIPT structure in 2012

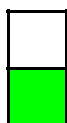




T2_UA_KIPT: status and development



Year	Number of Logical CPU	Productivity	Storage
2005	22 CPU 's	14 kSI2000	2.4 TB
2007	76 CPU's	96 kSI2000	57 TB
2009	42 CPU's (x86_64 only)	0.5 kHEPspec06	67 TB
2010	68 CPU's (x86_64 only)	1.4 kHEPspec06	90 TB
2011	230 CPU's (x86_64 only)	3.0 kHEPspec06	245 TB
2012	374 CPU's (x86_64 only)	5.5 kHEPcpec06	~350TB

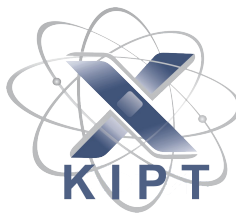


exists

Under configuration



Computing Facilities



Platform	Quantity	Test Result (HEPspec06)	Sum. Result (kHEPspec06)
Dual Intel Xeon EM64T 3.4 GHz	3	13	0.039
Dual CPU, Quad Core, Intel Xeon E5420 2.5 GHz	6	64	0.384
Dual CPU, Quad Core, Intel Xeon E5620 2.4 GHz	22	121	2.662
Dual CPU, Six Core, Intel Xeon E5-2620 2.0 GHz	12	207	2.484



exists

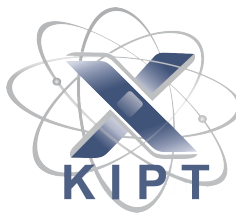
Under configuration

Already have **3.085 kHEPspec06**

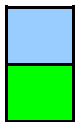
Expected **5.53 kHEPspec06**



Test of potential WN's



Platform	OS	Kernel	RAM	gcc	Test type	Result	Per 1 CPU
Dual Intel Xeon EM64T 3.4 GHz	SLC5	2.6.18-128.7.1.el5	4 GB	4.1.2	S2k6 all_cpp 64bit	13.8	6.9
Dual CPU, Quad Core, Intel Xeon E5420 2.5 GHz	SLC5	2.6.18-128.7.1.el5	16 GB	4.1.2	S2k6 all_cpp 64bit	63.8	8.0
Dual CPU, Quad Core, Intel Xeon E5620 2.4 GHz	SLC5	2.6.18-194.26.1.el5	16 GB	4.1.2	S2k6 all_cpp 64bit	121.5	15.2
Dual CPU, Intel Xeon E5 -2620 2.2 GHz	SLC5	2.6.18-128.7.1.el5	24 GB	4.1.2	S2k6 all_cpp 64bit	207.8	17.3



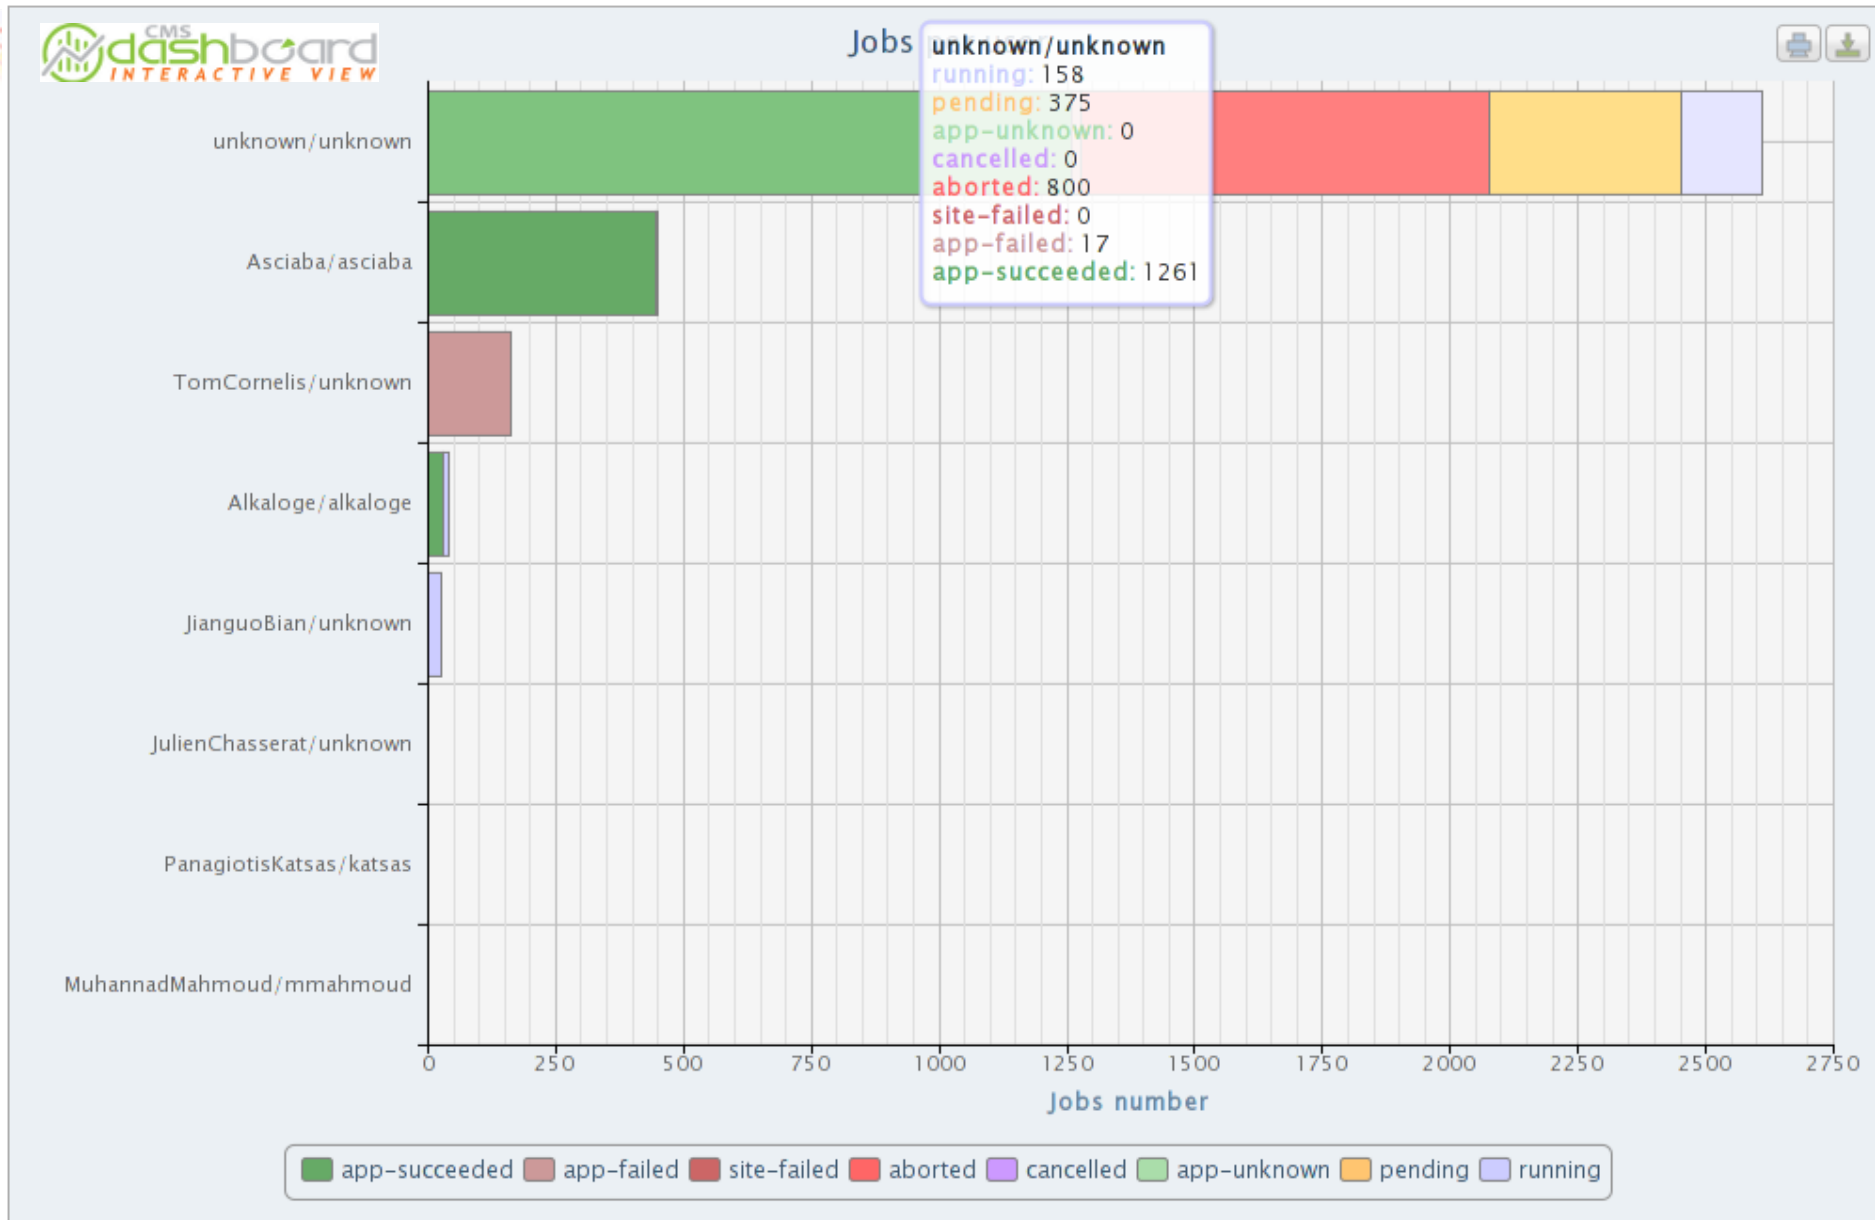
exists

Under configuration

* all_cpp — 4 floating point tests(444.namd, 447.dealll, 450.soplex, 453.povray) and 3 integer (471.omnetpp, 473.astar, 483.xalancbmk)



Daily jobs statistic





Local jobs monitoring



```
alexander@cms42:~/JOBS - Терминал
2012-07-10_18:41:42.2      PBS server: cms.kipt.kharkov.ua
JobNameList  RUNNING  QUEUED  WAITING  EXITING  TRANSITING  HELD  AllUserJobs
pilcms01      2         0         1         0         0         0         3
pilcms04      0         0         1         0         0         0         1
pilcms05      0         0         1         0         0         0         1
pilcms06      1         0         1         0         0         0         2
pilcms07      0         0         1         0         0         0         1
pilcms08      2         0         1         0         0         0         3
pilcms09      1         0         0         0         0         0         1
pilcms10      3         0         0         0         0         0         3
prdcms92     124        1         62         0         0         0       187
JobsState    133        1         68         0         0         0       202
```

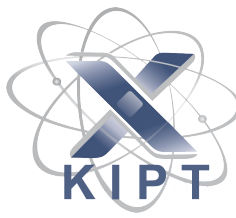
```
alexander@cms81:~/JOBS - Терминал
2012-07-10_18:41:43.8      PBS server: cream.kipt.kharkov.ua
JobNameList  RUNNING  QUEUED  WAITING  EXITING  TRANSITING  HELD  AllUserJobs
cms063        1         0         0         0         0         0         1
cms072       29         0         0         0         0         0       29
cms110       13         0         0         0         0         0       13
cms151        2         0         0         0         0         0         2
cms171       11         18         0         0         0         0       29
JobsState    56        18         0         0         0         0       74
```

```
cms34: 2 1 0 0 0 0 3
cms37: 8 0 0 0 0 0 8
cms41: 8 0 0 0 0 0 8
cms42: 8 0 0 0 0 0 8
cms43: 8 0 0 0 0 0 8
cms44: 8 0 0 0 0 0 8
cms45: 8 0 0 0 0 0 8
cms46: 8 0 0 0 0 0 8
cms48: 7 0 0 0 0 0 7
cms70: 7 0 0 0 0 0 7
cms71: 8 0 0 0 0 0 8
cms72: 7 0 0 0 0 0 7
cms73: 8 0 0 0 0 0 8
cms74: 0 0 68 0 0 0 68
cms75: 8 0 0 0 0 0 8
cms76: 8 0 0 0 0 0 8
cms77: 7 0 0 0 0 0 7
```

```
cms80: 8 0 0 0 0 0 8
cms81: 8 0 0 0 0 0 8
cms82: 8 0 0 0 0 0 8
cms83: 8 0 0 0 0 0 8
cms84: 8 0 0 0 0 0 8
cms85: 8 0 0 0 0 0 8
cms86: 8 0 0 0 0 0 8
```



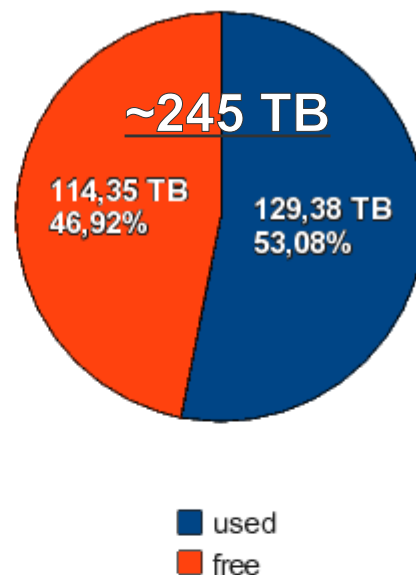
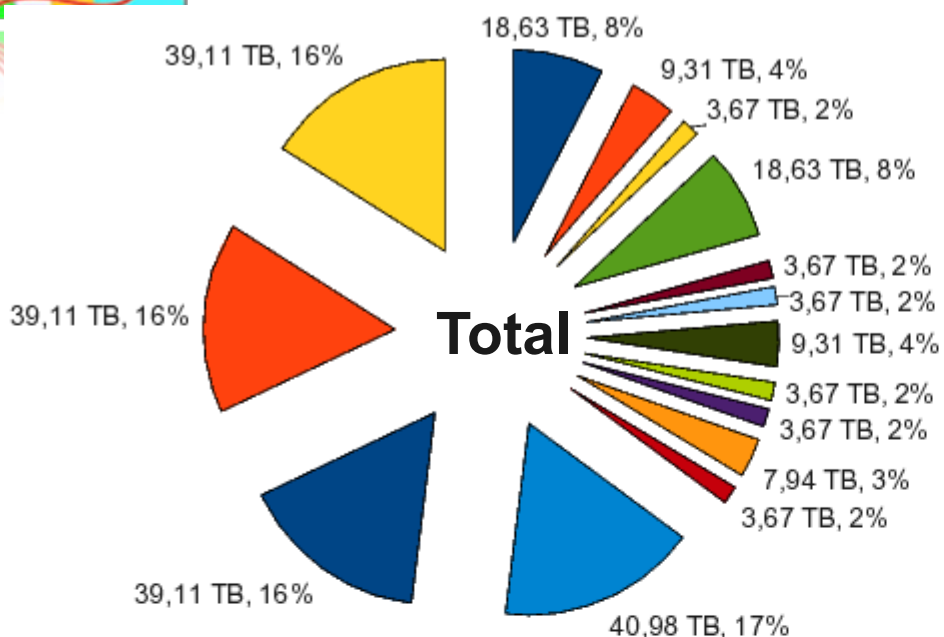
Configutration of disk servers



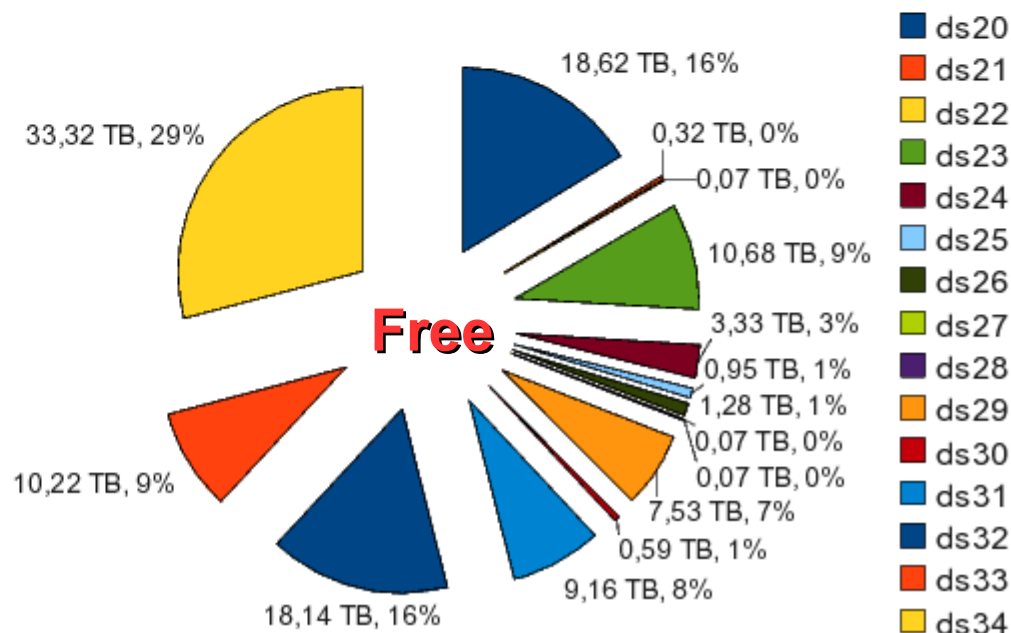
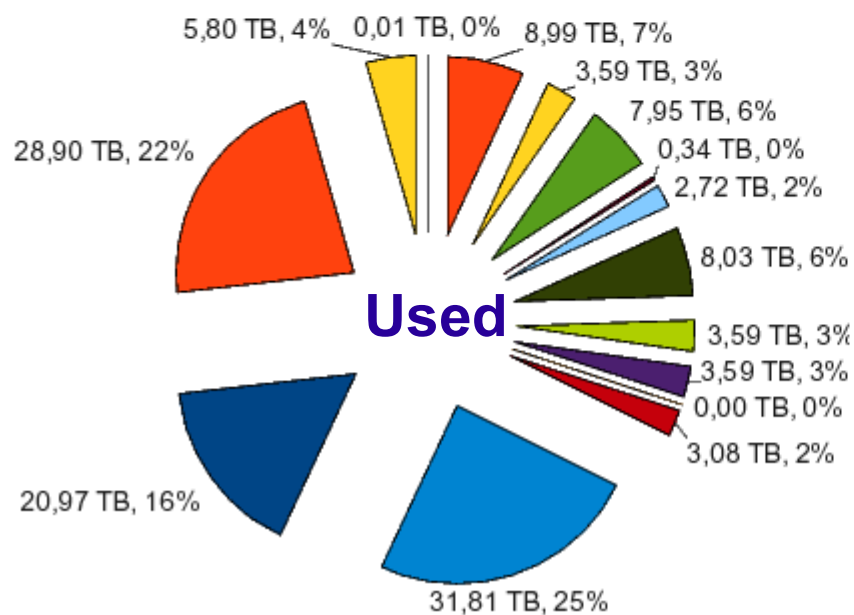
- **3.7 TB DS's (3ware 9550SX-16ML)**
 - System(400GB), RAID5 (400GB x 11), SPARE(400GB)
- **7.9 TB DS (3ware 9500S-12)**
 - RAID1 (2000GB x 2), RAID5 (2000GB x 5), SPARE (2000GB)
- **9.3 TB DS's (3ware 9550SXU-16ML)**
 - System(1000GB), RAID5(1000GB x 11), SPARE(1000GB)
- **18.6 TB DS's (3ware 9550SX-16ML)**
 - SINGLE(400GB), RAID5 (2000GB x 11), SPARE(2000GB)
- **39.1 TB DS's (3ware 9690SA-8I)**
 - SINGLE(2000GB), **RAID6** (2000GB x 23)
- **40.9 TB DS's (3ware 9690SA-4I4E)**
 - System(2000GB), **RAID6** (2000GB x **24**)



Mass Storage System

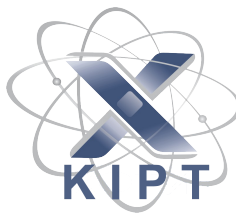


Capacity (TB)	Number of nodes
3.667	6
7.936	1
9.313	2
18.625	2
39.114	3
40.976	1
~100	1



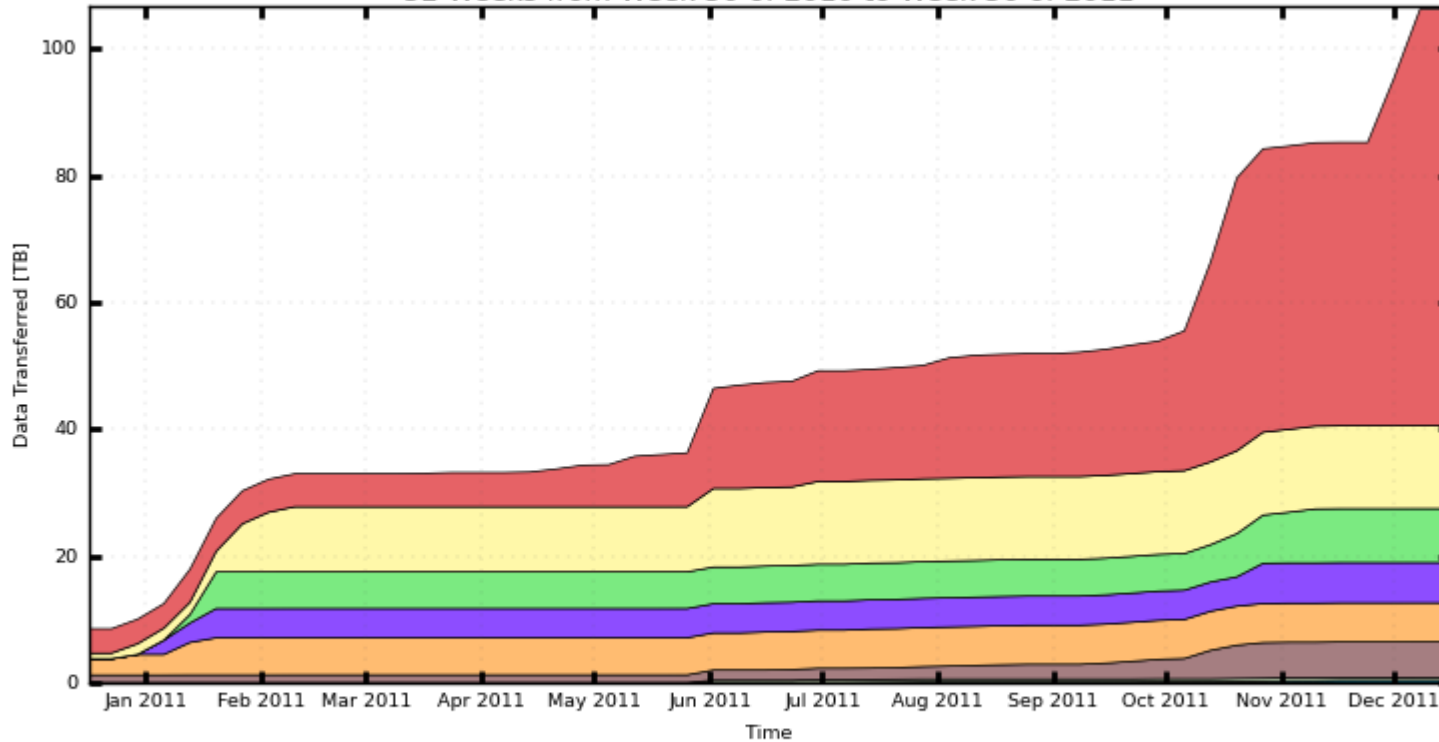


CMS data transfers to T2_UA_KIPT in 2011



CMS PhEDEx - Cumulative Transfer Volume

52 Weeks from Week 50 of 2010 to Week 50 of 2011



T1_US_FNAL_Buffer T1_UK_RAL_Buffer T1_IT_CNAF_Buffer T1_FR_CCIN2P3_Buffer T1_DE_KIT_Buffer
T1_ES_PIC_Buffer T1_TW_ASGC_Buffer T1_CH_CERN_Buffer

Total: 106.21 TB, Average Rate: 0.00 TB/s



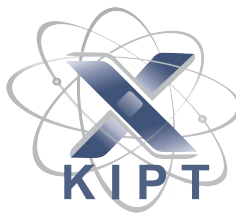
Storage consistency check: 9TB lost!

- At the beginning of 2012 two disks crashed on one of our disk servers (ds21, **RAID5** ,9TB), and the RAID controller is out of order on another one (ds32, RAID6, 39.1TB).
 - All data were successfully recovered at **ds32 (~25 TB)**;
 - RAID controller was replaced.
 - **9 TB** of data were, unfortunately, **LOST** at **ds21**;
 - All our attempts to recover these data failed.

Conclusion: 21208 files were lost at T2_UA_KIPT ?...



Storage consistency check: Fortune still with us!



- Storage consistency check campaign resulted in:
 - **of the 21208 files:**
 - 16350 are /store/unmerged — individual jobs output (should be cleaned-up every month). **Not real loss.**
 - 4582 are /store/mc — deprecated datasets, parts of the two deletion campaigns: Spring11SummerFall11. **Not real loss.**
 - 111 are /store/temp. **Not real loss.**
 - 10 are **LoadTestSource** files — **Important loss.**
 - 104 are /store/data — probably orphans. **Not real loss.**
 - **21** from /store/data — **Important loss. Were successfully re-transferred.**
 - 30 files are under /store/ops/ — not checked at all. **Not real loss.**

Conclusion: Only **31** files(~90GB) were **REALY** lost at T2_UA_KIPT;
Refuse usage of RAID 5 in future;
21 files were re-transferred to KIPT;
10 files(Load Test) were successfully re-created locally.



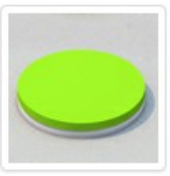
Plans and Prospects

- Increase the bandwidth of the outer channel up to 10 Gbit / sec;
- Upgrade the internal network infrastructure;
- Increase computing facility;
- Replace old disk servers with newer;
- Change hard disk drives on old disk servers to High capacity HDD (if applicable);
- Upgrade the Cooling system;
- Further support and maintenance of the System.



OVERVIEW CLUSTER KIPT Number of Jobs per year ZABBIX VIEW

- Overview Cluster
- Hardware Structure
- Software Content
- Network Bandwidth
- Computing Load
- Environment
- Certificates Validity
- Tickets
- Top Users



COMPUTING RESOURCE <small>(update)</small> 33 working nodes CPU/CORES/HEPSPEC 66/228/3114.32	STORAGE RESOURCE 13 storage units CAPACITY/FREE 209/79.9 TB	GRID JOBS 137 running TOTAL/QUEUE/WAITING 384 / 242 / 5
--	---	---

- Nodes Resource
- Storage Resource
- Jobs Payloads
- Accounting
- Temperature

The current review of Computing Resource, Storage Resource, Grid Jobs operating in Kharkov-KIPT-LCG2

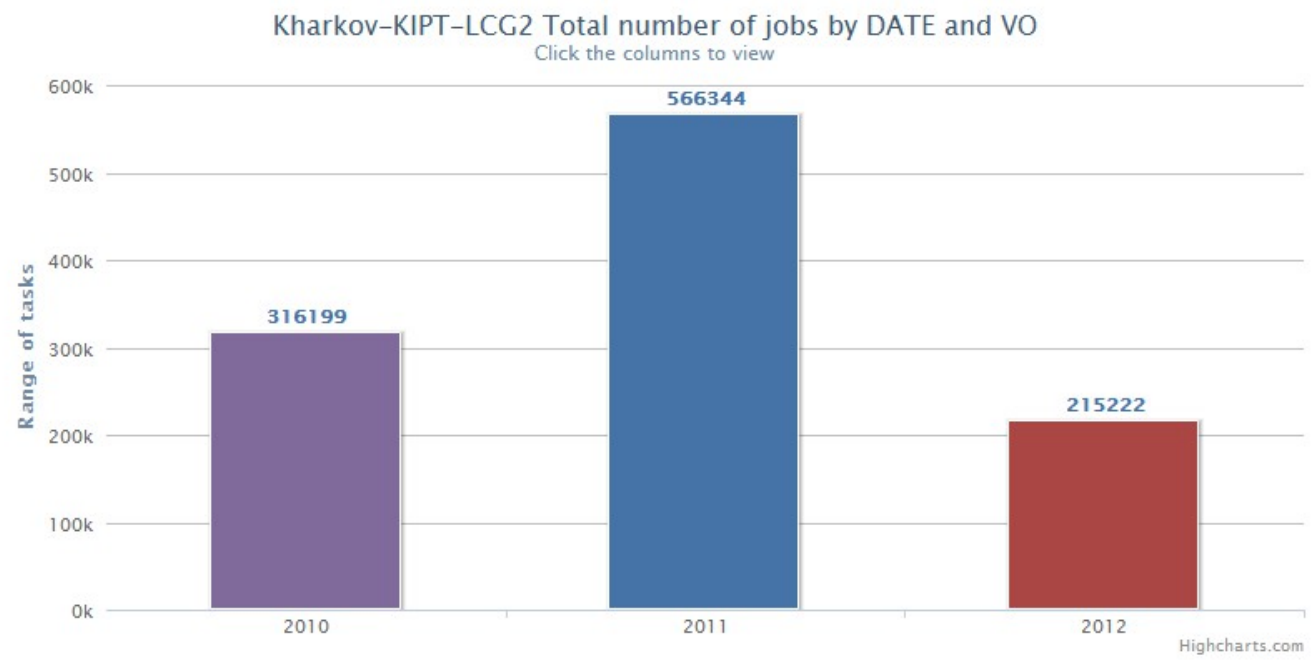


Chart shows the volume of completed tasks on the processing power T2-UA-KIPT cluster by years