MetaCentrum

CESNET

PBSMon: A Web-Based Framework for Monitoring and Accounting the Czech National Grid and Cloud Infrastructure

Martin Kuba, Dalibor Klusáček
CESNET, Czech Republic

MetaCentrum e-Infrastructure

- **MetaCentrum** is the Czech national grid and cloud infrastructure
- operated by CESNET, an association of Czech public universities and the Czech Academy of Sciences
- CESNET is also the Czech NREN operator
- **MetaCentrum** is a rather heterogeneous infrastructure:
  - 544 computing machines
    - 29 clusters/SMP-machines of 9 owners, located in 7 cities
    - from 8-CPU cluster nodes up to 288-CPU SGI UltraViolet
    - all have CPUs, RAM, local HDDs
    - optionally GPUs, Infiniband, SSD disks, network scratch
  - 9 disk arrays with 1250 TiB capacity
  - 4 hierarchical storages (HSMs) with 14 PiB capacity

Complicated virtualisation

- computing jobs managed by 2 Torque (PBS) servers
- MetaCentrum physical machines are virtualised in 6 ways
  - no virtualisation (machines with GPUs or large RAM)
  - UltraViolet divided into hardware partitions
  - Xen managed by Magrathea (home-made)
    - 2 VMs, both managed by Torque
    - 1 VM in Torque, more VMs created on demand
  - OpenNebula hosts
    - with a VM managed by Torque
    - with general VMs
- OpenNebula hosts dynamically reassigned between Torque and general VMs depending on demand

PBSMon

- PBSMon visualises the complex infrastructure in a comprehensible way
- web framework showing
  - computing resources
    - physical machines
    - VMs
    - disk arrays
    - hierarchical data storages (HSMs)
  - Torque (PBS) state
    - jobs
    - queues
    - users
    - computing nodes (VMs and non-virtualized physical machines)
  - personalized information for users
    - user's jobs
    - used space and quotas on disk arrays and HSMs
    - qsub command refinement

CESNET (3232 CPU)

mandos.ics.muni.cz (896 CPU) - Cluster of 64-CPU SMP machines

mandos1 (64 CPU)	mandos2 (64 CPU)	mandos3 (64 CPU)	mandos4 (64 CPU)	mandos6 (64 CPU)	mandos7 (64 CPU)	mandos5 (64 CPU)
mandos8 (64 CPU)	mandos9 (64 CPU)	mandos10 (64 CPU)	mandos11 (64 CPU)	mandos12 (64 CPU)	mandos13 (64 CPU)	mandos14 (64 CPU)

doom.metacentrum.cz (480 CPU) - Cluster of machines with 2x GPU nVidia Tesla (Ostrava)

Cluster with GPGPU (General-purpose graphics processing unit). Each node contains 2x GPU nVidia Tesla K20 5GB (Kepler)

doom1 (16 CPU)	doom2 (16 CPU)	doom3 (16 CPU)	doom4 (16 CPU)	doom5 (16 CPU)	doom6 (16 CPU)	doom7 (16 CPU)
doom8 (16 CPU)	doom9 (16 CPU)	doom10 (16 CPU)	doom11 (16 CPU)	doom12 (16 CPU)	doom13 (16 CPU)	doom14 (16 CPU)
doom15 (16 CPU)	doom16 (16 CPU)	doom17 (16 CPU)	doom18 (16 CPU)	doom19 (16 CPU)	doom20 (16 CPU)	doom21 (16 CPU)
doom22 (16 CPU)	doom23 (16 CPU)	doom24 (16 CPU)	doom25 (16 CPU)	doom26 (16 CPU)	doom27 (16 CPU)	doom28 (16 CPU)
doom29 (16 CPU)	doom30 (16 CPU)					

gram.zcu.cz (160 CPU) - Cluster of machines with 4x GPU nVidia Tesla (Plzeň)

Cluster with GPGPU (General-purpose graphics processing unit). Každý uzel obsahuje 4x GPU nVidia Tesla M2090 6GB

gram1 (16 CPU)	gram2 (16 CPU)	gram3 (16 CPU)	gram4 (16 CPU)	gram5 (16 CPU)	gram6 (16 CPU)	gram7 (16 CPU)
gram8 (16 CPU)	gram9 (16 CPU)	gram10 (16 CPU)				

dukan.ics.muni.cz (240 CPU) - Cluster for cloud (Brno)

Cluster intended for cloud, i.e. running virtual machines when demanded

dukan1 (24 CPU)	dukan2 (24 CPU)	dukan3 (24 CPU)	dukan4 (24 CPU)	dukan5 (24 CPU)	dukan6 (24 CPU)	dukan7 (24 CPU)
dukan8 (24 CPU)	dukan9 (24 CPU)	dukan10 (24 CPU)				

Cluster hdc.cerit-sc.cz

Data from PBS server arien: Oct 24, 2014 3:57:07 PM
Data from PBS server wagap: Oct 24, 2014 3:57:08 PM
Data from PBS cache : Oct 24, 2014 3:57:05 PM
Data from OpenNebula : Oct 24, 2014 3:56:37 PM
Displayed: Oct 24, 2014 3:58:06 PM

cluster of 16-CPU machines, hosts virtual machines zapat (Jihlava)

A High Density cluster of CERIT-SC

Cluster hdc.cerit-sc.cz contains 112 nodes, each of the nodes has the following hardware specification:

CPU	2x 8-core Intel E5-2670 2.6GHz
RAM	128 GB
disk	2x 600 GB 15k
net	1x Infiniband 40 Gbit/s, 2x Ethernet 1 Gbit/s
comment	SPECfp2006 471 (i.e. cca 29 per core)
owner	CERIT-SC/MU

hdc1 (16 CPU)	hdc2 (16 CPU)	hdc3 (16 CPU)	hdc4 (16 CPU)	hdc5 (16 CPU)	hdc6 (16 CPU)	hdc7 (16 CPU)
hdc8 (16 CPU)	hdc9 (16 CPU)	hdc10 (16 CPU)	hdc11 (16 CPU)	hdc12 (16 CPU)	hdc13 (16 CPU)	hdc14 (16 CPU)
hdc15 (16 CPU)	hdc16 (16 CPU)	hdc17 (16 CPU)	hdc18 (16 CPU)	hdc19 (16 CPU)	hdc20 (16 CPU)	hdc21 (16 CPU)
hdc22 (16 CPU)	hdc23 (16 CPU)	hdc24 (16 CPU)	hdc25 (16 CPU)	hdc26 (16 CPU)	hdc27 (16 CPU)	hdc28 (16 CPU)
hdc29 (16 CPU)	hdc30 (16 CPU)	hdc31 (16 CPU)	hdc32 (16 CPU)	hdc33 (16 CPU)	hdc34 (16 CPU)	hdc35 (16 CPU)
hdc36 (16 CPU)	hdc37 (16 CPU)	hdc38 (16 CPU)	hdc39 (16 CPU)	hdc40 (16 CPU)	hdc41 (16 CPU)	hdc42 (16 CPU)

About MetaCentrum VO
Current affairs
Documentation and services
Getting an account
My account
Current state
Personal view
Physical machines
Virtual machines
Job queues
Jobs
Jobs queued
Users
Machine properties
List of hardware
Cloud
Statistics
User Support
Seminars
Portal map
Search

The physical machine has PBS node **minos12.zcu.cz**.

The physical machine has 2 virtual machines:.

- **minos12-1.zcu.cz**
- **minos12-2.zcu.cz**

Physical machine minos12.zcu.cz

This machine is dedicated for creation of new virtual machines on demand.

Virtual machine minos12-1.zcu.cz

name	minos12-1.zcu.cz
state	job-exclusive
state in PBS	job-exclusive
state in Magrathea	running
node type	virtual
properties	q_2w, debian7, xen, q_2h, q_2w_plus, q_1w, q_4d, q_2d, q_1d, q_4h, infiniband, q_privileged, q_long, hyperthreading, data-kky, q_backfill, home_nympha, home_minos, home_konos, cl_minos, nfs4, q_short, q_normal, quadcore, nodecpus12, xeon, x86, x86_64, em64t, linux, plzen
space in /scratch	483.8 GiB

reserved CPUs	100%	12 / 12
reserved memory	62%	12gb / 19gb
reserved HDD	0%	110gb / 495gb
used HDD	2%	11gb / 495gb

job	user	scratch	CPU	name	state	start	max till	queue
6892659.arien.ics.muni.cz	**matunovp**	local 60gb	8 CPU	trDR1b4f	R - running	23.10. 10:22	22.11. 9:23	**long**
6905712.arien.ics.muni.cz	**vandrush**	local 50gb	4 CPU	12000_gaff.inp.bsh	R - running	24.10. 14:30	25.10. 14:30	**q_1d**

The last job should end before 22.11.2014 09:23.

Accessible thru these queues on the PBS server **arien**:

queue	Priority	time limits	required property	jobs

Documentation and services
Getting an account
My account
Current state
Personal view
Physical machines
Virtual machines
Job queues
Jobs
Jobs queued
Users
Machine properties
List of hardware
Cloud
Statistics
User Support
Seminars
Portal map
Search
RSS
CESNET
IPv6

Storages

Scratch

Each computing node has fast local storage, its size is displayed in detail of each node.

Use scratch **during computation only**. Data are deleted 14 days after their job ended.

Disk arrays

directory	usage - used free	size
/storage/brno2/home	88%	110 TiB
/storage/brno3-cerit/home	75%	374 TiB
/storage/budejovice1/home	64%	44 TiB
/storage/jihlava1-cerit/home	75%	374 TiB
/storage/ostrava1/home	26%	88 TiB
/storage/plzen1/home	75%	44 TiB
/storage/plzen3-kky/home	77%	70 TiB
/storage/praha1/home	61%	51 TiB
/storage/praha2-elixir/home	25%	95 TiB
Total storage space in MetaCentrum		**1250 TiB**
Total used		852 TiB
Total free		398 TiB

Disk arrays consist of many connected disks. A file is stored on multiple disks, moreover its data are stored redundantly. When compared with a single disk, disk arra has higher read and write performance, and is more robust against data loss.

Use disk arrays **for storing data between jobs**.

Hierarchical storages

directory	usage - used free	size
/storage/plzen2-archive	43%	8034 TiB
/storage/jihlava2-archive	6%	3799 TiB
/storage/brno5-archive	11%	1587 TiB
/storage/brno4-cerit-hsm	26%	512 TiB
Total storage space in MetaCentrum		**13932 TiB**
Total used		3951 TiB
Total free		9981 TiB

GPU

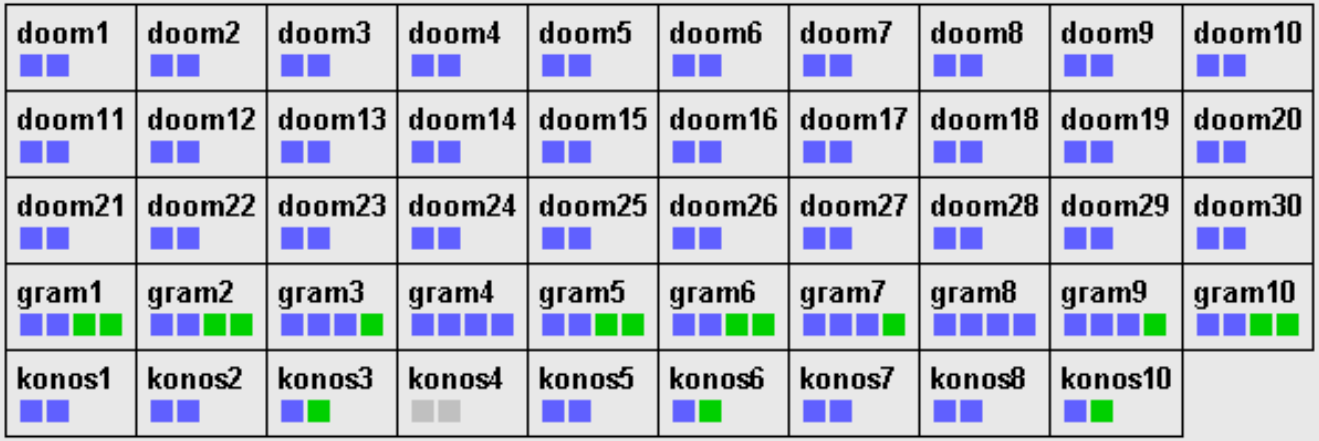
doom1
doom2
doom3
doom4
doom5
doom6
doom7
doom8
doom9
doom10
doom11
doom12
doom13
doom14
doom15
doom16
doom17
doom18
doom19
doom20
doom21
doom22
doom23
doom24
doom25
doom26
doom27
doom28
doom29
doom30
gram1
gram2
gram3
gram4
gram5
gram6
gram7
gram8
gram9
gram10
konos1
konos2
konos3
konos4
konos5
konos6
konos7
konos8
konos10

History

- started as a web frontend for PBSPro using hyperlinks among objects
- colors inspired by “xpbsmon” tool for X-window
- migrated from PBSPro to Torque
- extended from 1 server to multiple servers
- accommodated virtualization and mapping of VMs to phys. machines
- added charts for last 24 hours
- extended to display OpenNebula state
- enhanced to compute machine load not only from CPUs, but also RAM, GPUs, HDD and SSD usage
- support for multiple fairshares
- support for plan-based scheduler

About MetaCentrum VO
Current affairs
Documentation and services
Getting an account
My account
Current state
Personal view
Physical machines
Virtual machines
Job queues
Jobs
Jobs queued
Users
Machine properties
List of hardware
Cloud
Statistics
User Support
Seminars
Portal map
Search

Job 6852987.arien.ics.muni.cz

Data from PBS server arien: Oct 20, 2014 1:04:33 PM
Data from PBS server wagap: Oct 20, 2014 1:04:34 PM
Data from PBS cache : Oct 20, 2014 1:04:31 PM
Data from OpenNebula : Oct 20, 2014 1:00:40 PM
Displayed: Oct 20, 2014 1:05:53 PM

Basic info

job	CPU	reserved mem	used mem	name	user	CPU time	Wall time	state	queue
6852987.arien.ics.muni.cz	48	96gb	34gb	re_p1lt_df_40.r	nagnaat	133:59:31	02:47:54	R - running	ncbr_medium

resources	nodes=6:ppn=8:cl_perian
time created	Monday, October 20, 2014 10:13:47 AM
time eligible	Monday, October 20, 2014 10:13:47 AM
start	Monday, October 20, 2014 10:14:00 AM
max till	Saturday, October 25, 2014 10:14:04 AM
time last changed state	Monday, October 20, 2014 10:16:40 AM
comment	Job successfully started at Mon Oct 20 10:14:01 2014
working directory	$HOME
SCRATCHDIR	/scratch/nagnaat/job_6852987.arien.ics.muni.cz
exec host/cpu	**perian23-2/7** **perian23-2/6 perian23-2/5 perian23-2/4 perian23-2/3 perian23-2/2 perian23-2/1** **perian23-2/0 perian24-2/7 perian24-2/6 perian24-2/5 perian24-2/4 perian24-2/3** **perian24-2/2 perian24-2/1 perian24-2/0 perian25-2/7 perian25-2/6 perian25-2/5** **perian25-2/4 perian25-2/3 perian25-2/2 perian25-2/1 perian25-2/0 perian27-2/7** **perian27-2/6 perian27-2/5 perian27-2/4 perian27-2/3 perian27-2/2 perian27-2/1** **perian27-2/0 perian28-2/7 perian28-2/6 perian28-2/5 perian28-2/4 perian28-2/3** **perian28-2/2 perian28-2/1 perian28-2/0 perian29-2/7 perian29-2/6 perian29-2/5** **perian29-2/4 perian29-2/3 perian29-2/2 perian29-2/1 perian29-2/0**
variables	PBS_O_QUEUE=ncbr_medium PBS_O_HOST=META INF_JOB_NAME=re_p1lt_df_40.run INF_JOB_NAME_SUFFIX= INF_JOB_PATH=/storage/brno2/home/nagnaat/Ch_14-3-3/REMD_test/df_p1lt/dr_40 INF_JOB_MACHINE=skirit.ics.muni.cz INF_JOB_KEY=6b450745-2ca8-e4cf-3b45-195978c2a3db ABS_ROOT=/software/ncbr/softrepo/common/abs/2.0.5466/x86_64/single

Server wagap.cerit-sc.cz - Environment of center CERIT-SC

queue		Priority	time limits	required property	jobs					max CPUs per user	fairshare
					queued	running /max	completed	total	max jobs per user		
p2ptest@wagap.cerit-sc.cz		99	0 - 0		0	0 /	0	1			
ops@wagap.cerit-sc.cz		99	0 - 0		0	0 / 1000	0	7	4000		
maintenance@wagap.cerit-sc.cz		99	0 - 0		0	0 / 1000	0	0	4000		
reserved@wagap.cerit-sc.cz		90	0 - 0		0	0 / 1000	0	0	4000		
eli@wagap.cerit-sc.cz		60	0 - 744:00:00	eli	0	0 / 1000	0	0	4000		
q_1w@wagap.cerit-sc.cz		50	96:00:01 - 168:00:00	q_1w	38	85 / 1000	54	177	1200		
q_2w@wagap.cerit-sc.cz		50	168:00:01 - 336:00:00	q_2w	27	229 / 1000	130	391	800		
q_4d@wagap.cerit-sc.cz		50	48:00:01 - 96:00:00	q_4d	1	34 / 2000	191	226	4000		
q_2h@wagap.cerit-sc.cz		50	0 - 02:00:00	q_2h	0	0 / 1000	11	13	4000		
q_2d@wagap.cerit-sc.cz		50	24:00:01 - 48:00:00	q_2d	35	26 / 1000	179	240	2000		
q_1d@wagap.cerit-sc.cz		50	04:00:01 - 24:00:00	q_1d	26	112 / 1000	181	319	3200		
q_4h@wagap.cerit-sc.cz		50	02:00:01 - 04:00:00	q_4h	0	0 / 1000	0	0	4000		
q_2w_plus@wagap.cerit-sc.cz		50	336:00:01 - 1488:00:00	q_2w_plus	4	33 / 1000	3	40	400		
uv@wagap.cerit-sc.cz		30	00:00:01 - 96:00:00	uv	0	1 / 10000	16	17	8000		
backfill@wagap.cerit-sc.cz		20	00:00:01 - 24:00:00	q_backfill	0	0 / 10000	0	0	8000		
default@wagap.cerit-sc.cz		0	0 - 0		0	0 /	0	1			

Server wagap.cerit-sc.cz - Environment of center CERIT-SC

job	CPU	resources	reserved mem	user	planned start	planned nodes	waiting reason
592104.wagap.cerit-sc.cz	48	1:ppn=48:brno:cl_mandos	64gb	**kakanek**	2/22/70 12:07 PM		Never Running: This jobs requirements will never be satisfied under the current grid configuration.
595411.wagap.cerit-sc.cz	32	1:ppn=32:brno:cl_mandos	64gb	**kakanek**	2/22/70 12:07 PM		Never Running: This jobs requirements will never be satisfied under the current grid configuration.
595412.wagap.cerit-sc.cz	32	1:ppn=32:brno:cl_mandos	64gb	**kakanek**	2/22/70 12:07 PM		Never Running: This jobs requirements will never be satisfied under the current grid configuration.
596099.wagap.cerit-sc.cz	4	1:ppn=4:x86_64:linux:cl_zewura9	30gb	**gkesan**	2/22/70 12:07 PM		Never Running: This jobs requirements will never be satisfied under the current grid configuration.
596100.wagap.cerit-sc.cz	4	1:ppn=4:x86_64:linux:cl_zewura9	30gb	**gkesan**	2/22/70 12:07 PM		Never Running: This jobs requirements will never be satisfied under the current grid configuration.
577201.wagap.cerit-sc.cz	1	1:brno:cl_zegox#excl	3gb	**kesslr**	10/24/14 4:23 PM	**zegox5.cerit-sc.cz**	
577204.wagap.cerit-sc.cz	1	1:brno:cl_zegox#excl	3gb	**kesslr**	10/24/14 4:23 PM	**zegox2.cerit-sc.cz**	
595982.wagap.cerit-sc.cz	16	1:ppn=16	16gb	**kulhanek**	10/24/14 5:05 PM	**zapat12.cerit-sc.cz**	
595956.wagap.cerit-sc.cz	16	1:ppn=16	205gb	**tomsvo**	10/24/14 6:13 PM	**zewura8.cerit-sc.cz**	
595164.wagap.cerit-sc.cz	8	1:ppn=8	24gb	**kubikuk**	10/24/14 6:25 PM	**zapat44.cerit-sc.cz**	
595165.wagap.cerit-sc.cz	8	1:ppn=8	24gb	**kubikuk**	10/24/14 6:28 PM	**zapat31.cerit-sc.cz**	
595981.wagap.cerit-sc.cz	10	1:ppn=10:cl_zewura	10gb	**slamavl**	10/24/14 6:28 PM	**zewura17.cerit-sc.cz**	
595989.wagap.cerit-sc.cz	10	1:ppn=10:cl_zewura	10gb	**slamavl**	10/24/14 6:28 PM	**zewura18.cerit-sc.cz**	
595589.wagap.cerit-sc.cz	16	1:ppn=16	40gb	**mnovak**	10/24/14 6:31 PM	**zapat29.cerit-sc.cz**	
595590.wagap.cerit-sc.cz	16	1:ppn=16	40gb	**mnovak**	10/24/14 7:17 PM	**zapat49.cerit-sc.cz**	
595986.wagap.cerit-sc.cz	16	1:ppn=16	16gb	**kulhanek**	10/24/14 7:17 PM	**zapat71.cerit-sc.cz**	
577205.wagap.cerit-sc.cz	1	1:brno:cl_zegox#excl	3gb	**kesslr**	10/24/14 7:45 PM	**zegox7.cerit-sc.cz**	
596051.wagap.cerit-sc.cz	16	1:ppn=16	24gb	**kulhanek**	10/24/14 7:46 PM	**zewura10.cerit-sc.cz**	
596079.wagap.cerit-sc.cz	16	1:ppn=16:cl_zapat	115gb	**kulhanek**	10/25/14 1:18 PM	**zapat14.cerit-sc.cz**	
596060.wagap.cerit-sc.cz	8	1:ppn=8:x86_64:linux:cl_zigur	50gb	**gkesan**	10/25/14 2:32 PM	**zigur1.cerit-sc.cz**	
592906.wagap.cerit-sc.cz	16	1:ppn=16	32gb	**straka**	10/25/14 9:53 PM	**zewura7.cerit-sc.cz**	
593786.wagap.cerit-sc.cz	16	1:ppn=16	117gb	**vicha**	10/25/14 9:53 PM	**zewura7.cerit-sc.cz**	
595574.wagap.cerit-sc.cz	16	1:ppn=16:cl_zapat	40gb	**ailar**	10/26/14 6:03 AM	**zapat68.cerit-sc.cz**	
595591.wagap.cerit-sc.cz	16	1:ppn=16:cl_zapat	40gb	**ailar**	10/26/14 7:13 AM	**zapat53.cerit-sc.cz**	
595595.wagap.cerit-sc.cz	16	1:ppn=16:cl_zapat	40gb	**ailar**	10/26/14 7:26 AM	**zapat82.cerit-sc.cz**	
596052.wagap.cerit-sc.cz	16	1:ppn=16	24gb	**kulhanek**	10/26/14 9:17 AM	**zewura13.cerit-sc.cz**	
596072.wagap.cerit-sc.cz	16	1:ppn=16	24gb	**kulhanek**	10/26/14 10:11 AM	**zewura11.cerit-sc.cz**	
596073.wagap.cerit-sc.cz	16	1:ppn=16	24gb	**kulhanek**	10/26/14 10:11 AM	**zewura11.cerit-sc.cz**	
596074.wagap.cerit-sc.cz	16	1:ppn=16	24gb	**kulhanek**	10/26/14 10:11 AM	**zewura11.cerit-sc.cz**	
596075.wagap.cerit-sc.cz	16	1:ppn=16	24gb	**kulhanek**	10/26/14 11:30 AM	**zewura11.cerit-sc.cz**	
595067.wagap.cerit-sc.cz	16	1:ppn=16:cl_zapat	20gb	**stolba**	10/26/14 12:26 PM	**zapat30.cerit-sc.cz**	
595051.wagap.cerit-sc.cz	16	1:ppn=16:cl_zapat	20gb	**stolba**	10/26/14 4:05 PM	**zapat12.cerit-sc.cz**	
595052.wagap.cerit-sc.cz	16	1:ppn=16:cl_zapat	20gb	**stolba**	10/26/14 4:08 PM	**zapat62.cerit-sc.cz**	

Job and CPU counts

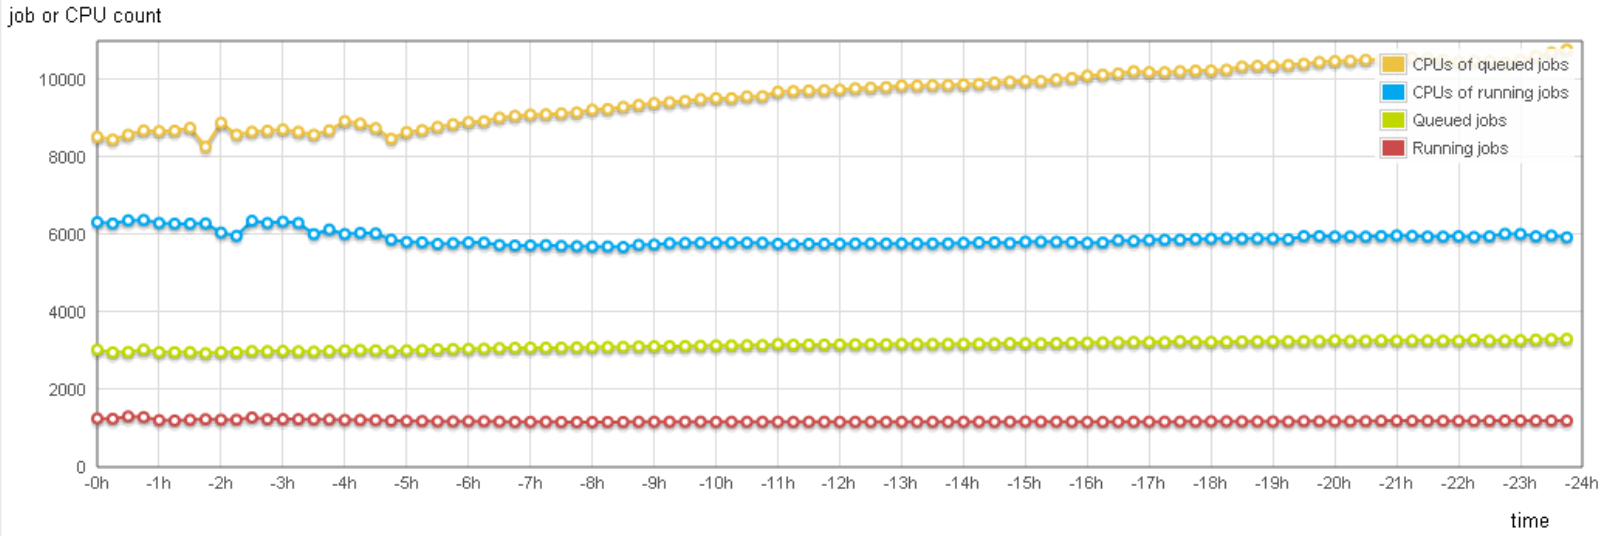

The line chart shows the numbers of running and wainting jobs and their CPUs, in 15 minute intervals during the last day. The numbers are **total counts** of jobs and CPUs.

Changes in job count

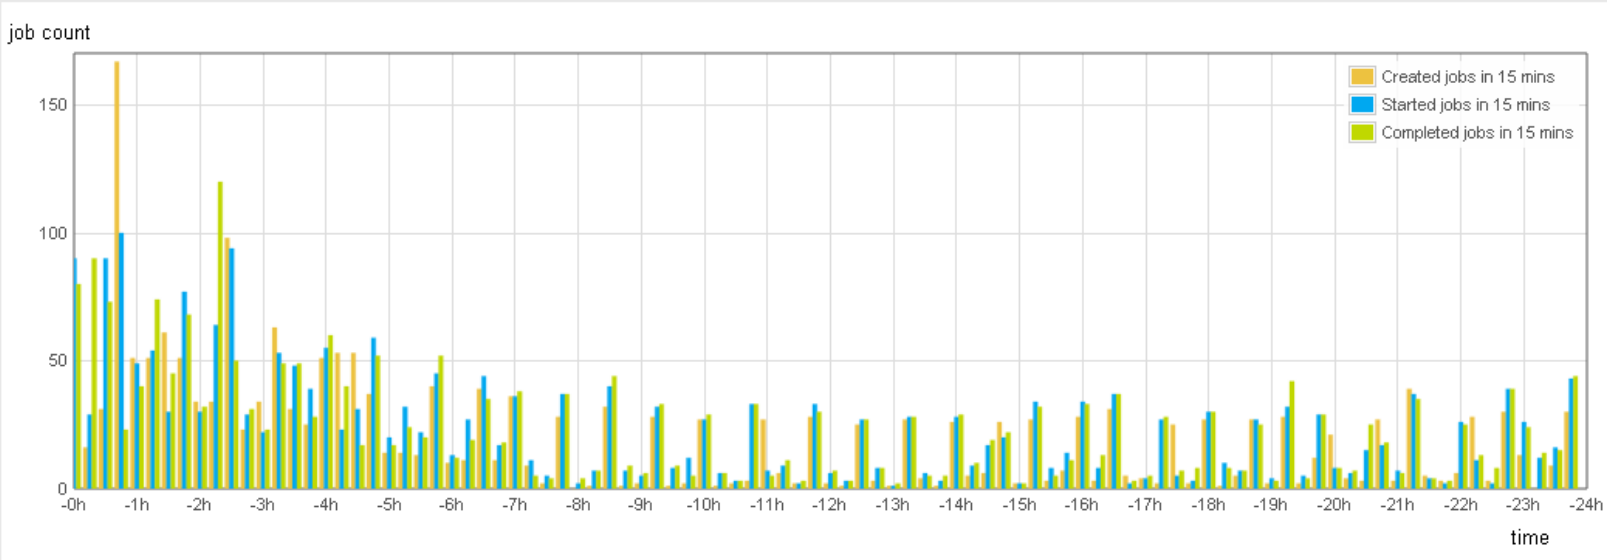

The column chart shows the numbers of created/started/completed jobs during 15 minute intervals during the last day. The numbers are **changes** in the counts of running and waiting jobs.

User personalisation

- user's jobs
- user quotas and used space
  - unifies disk arrays and HSMs
- qsub command refinement
  - user specifies resource requirements
  - tool simulates scheduler's algorithm for choosing nodes
  - shows Torque nodes matching the requirements

User uhlik

Data from
Data from
Data
Data

User **Filip Uhlik** from organization **Univerzita Karlova** belongs to research group **Katedra fyz. a makromolekulární chemie PFF UK (prof.Procházka)** and has a MetaCentrum account expiring on **Feb 1, 2015**. Computed **13,190 jobs** with the total CPU time of **23502 days**. Produced **19 publications** with acknowledgement to the MetaCentrum and **9 publications** with acknowledgement to the CERIT-SC.

year	number of jobs	CPUdays
2012	4772	6,448.1
2013	4443	8,991.2
2014	3975	8,062.6

Jobs in PBS

queue	job count					CPU count				
	total	queued	running	completed	other	total	queued	running	completed	other
q_1d@wagap.cerit-sc.cz	45	1	44	0	0	45	1	44	0	0
q_1w@wagap.cerit-sc.cz	1	0	1	0	0	16	0	16	0	0
total	**46**	**1**	**45**	**0**	**0**	**61**	**1**	**60**	**0**	**0**

job	CPU	reserved mem	used mem	name	CPU time	Wall time	state	host/cpu	
591496.wagap.cerit-sc.cz	16	32gb	9gb	ks/ks183full	1162:05:27	72:50:47	R - running	zewura1/79	q_
593482.wagap.cerit-sc.cz	1	2gb	12mb	100_1_20_0.3_1e-3_0_0.0	00:25:15	00:26:03	R - running	zewura13/42	q_
593483.wagap.cerit-sc.cz	1	2gb	12mb	100_1_20_0.3_1e-3_0_0.0634	00:26:06	00:26:36	R - running	zewura20/79	q_
593484.wagap.cerit-sc.cz	1	2gb	0b	100_1_20_0.3_1e-3_0_0.1			Q - queued		q_
593485.wagap.cerit-sc.cz	1	2gb	12mb	100_1_20_0.3_1e-4_0_0.0	00:29:00	00:29:01	R - running	zapat42/7	q_
593486.wagap.cerit-sc.cz	1	2gb	12mb	100_1_20_0.3_1e-4_0_0.0634	00:28:58	00:28:59	R - running	zapat42/10	q_
593487.wagap.cerit-sc.cz	1	2gb	12mb	100_1_20_0.3_1e-4_0_0.1	00:28:56	00:28:57	R - running	zapat42/11	q_

About MetaCentrum VO
Current affairs
Documentation and services
Getting an account
My account
Personal information
Publications
Annual report and extension of accounts
Shell or quota change
Quotas overview
License
Report problem
Password change
My jobs
Current state
User Support
Seminars
Portal map
Search
RSS

Quotas on storages

Data timestamp 2014-10-24 16:36:20

Documentation for quotas is on the wiki page **Quotas**.

user	dir	space				files			
		used	soft quota	hard quota	grace	used	soft limit	hard limit	grace
makub	/storage/brno2/home/makub	681MB	17TB	18TB		3708	-	-	
makub	/storage/brno3-cerit/home/makub	12GB	1TB	3TB	none	37112	-	-	none
makub	/storage/brno5-archive/home/makub	0	85TB						
makub	/storage/budejovice1/home/makub	1MB	6TB	10TB		47	-	-	
makub	/storage/jihlava1-cerit/home/makub	1TB	14TB	17TB	none	150599	-	-	none
makub	/storage/jihlava2-archive/home/makub	0	5TB						
makub	/storage/ostrava1/home/makub	108kB	3TB	5TB		20	-	-	
makub	/storage/plzen1/home/makub	445MB	1TB	3TB	none	860	-	-	none
makub	/storage/plzen2-archive/home/makub	6TB	10TB						
makub	/storage/praha1/home/makub	9MB	3TB	5TB		194	-	-	

Personal view

Data from PBS server arien: Oct 24, 2014 4:42:40 PM
Data from PBS server wagap: Oct 24, 2014 4:42:41 PM
Data from PBS cache : Oct 24, 2014 4:42:38 PM
Data from OpenNebula : Oct 24, 2014 4:40:06 PM
Displayed: Oct 24, 2014 4:43:07 PM

This page shows a personal view of the PBS system for the user **makub**, i.e. queues and computing nodes accessible by the user.

Jobs of user "makub"

user	job count					CPU count				
	total	queued	running	completed	other	total	queued	running	completed	other
makub	0	0	0	0	0	0	0	0	0	0

list of jobs

personal view of storages.

personal view of computing machines.

Command qsub refining

qsub -l walltime= 0 w 1 d 0 h 0 m 0 s -q \
-l mem= 400 mb -l scratch= 400 mb : \
-l nodes= 1 :ppn= 11 :x86_64 :linux uloha.sh

For properties meaning see Machine properties.

Find machines mathing the resource specification

Result

Selection: qsub -l walltime=1d -l mem=400mb -l scratch=400mb -l nodes=1:ppn=11:x86_64:linux

OK

The requirement is 1 machine, and 5 such machines are free, out of 160 machines matching the requirements. The job may be started immediately. If it is not, do not panic, there may be a multi-CPU job in the queue ahead of jor job, and the planner may be collection free CPUs for it.

Machines available right now

doom6 (12 CPU, 49.2 GiB RAM, 1.7 TiB HDD)	**doom8** (12 CPU, 55.2 GiB RAM, 1.7 TiB HDD)	**doom28** (12 CPU, 56.1 GiB RAM, 1.7 TiB HDD)	**losgar2** (11 CPU, 1.8 GiB RAM, 1.0 TiB HDD)	**luna52** (11 CPU, 31.6 GiB RAM, 827.1 GiB HDD)

Implementation

- PBSMon is written in Java
- runs in Tomcat servlet container
- uses Stripes MVC framework, JSP + JSTL
- connects to Torque using C library with JNI calls
- reads files and makes HTTP calls to other sources of data
- integrates data from many sources
  - multiple Torque servers (jobs, queues, nodes)
  - /etc/group on Torque servers
  - list of phys. machines and users from Perun
  - fairshare and other data from pbs_cache
  - multiple OpenNebula servers
  - disk arrays' states from output of "df" command
  - HSMs' states from HSMs' accounting
  - users' past jobs from job accounting database
- generates HTML pages and JSON responses

Summary

- MetaCentrum has very complex infrastructure
- its state can be obtained only by integrating from many sources
- for resource owners provides state of physical resources
- for users and administrators provides state of jobs, queues etc.
- for users provides personalised information
- developed since 2003

Thank you for your attention