JobDigest – Detailed System Monitoring-Based Supercomputer Application Behavior Analysis

Dmitry Nikitenko, Alexander Antonov, Pavel Shvets, Sergey Sobolev, Konstantin Stefanov, Vadim Voevodin, Vladimir Voevodin, Sergey Zhumatiy

Research Computing Center
Lomonosov Moscow State University

26.09.2017 Moscow, Russia

Moscow State HPC Facilities

MSU HPC Center
2017 Q3

- **"Lomonosov-2"**
  - 2,9 PF peak
  - 42688 cores
- **"Lomonosov"**
  - 1,7 PF peak
  - 82468 cores
- Some 10x TF systems
- Advanced Infrastructure

"Strela" – USSR 1st mass prod.

"Setun" – ternary logic

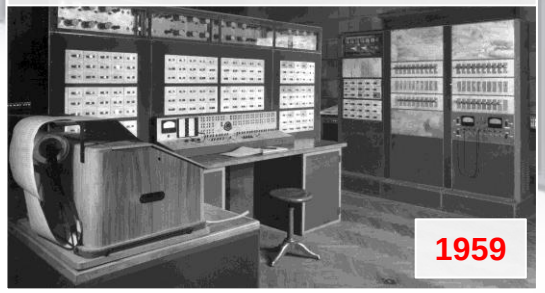

"BESM-6"

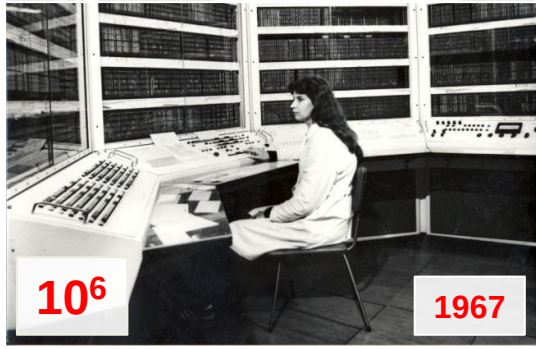

"Lomonosov-2"

"Lomonosov"

"Chebyshev"

Scientific projects of MSU HPC Center

The Largest HPC Center in Russia

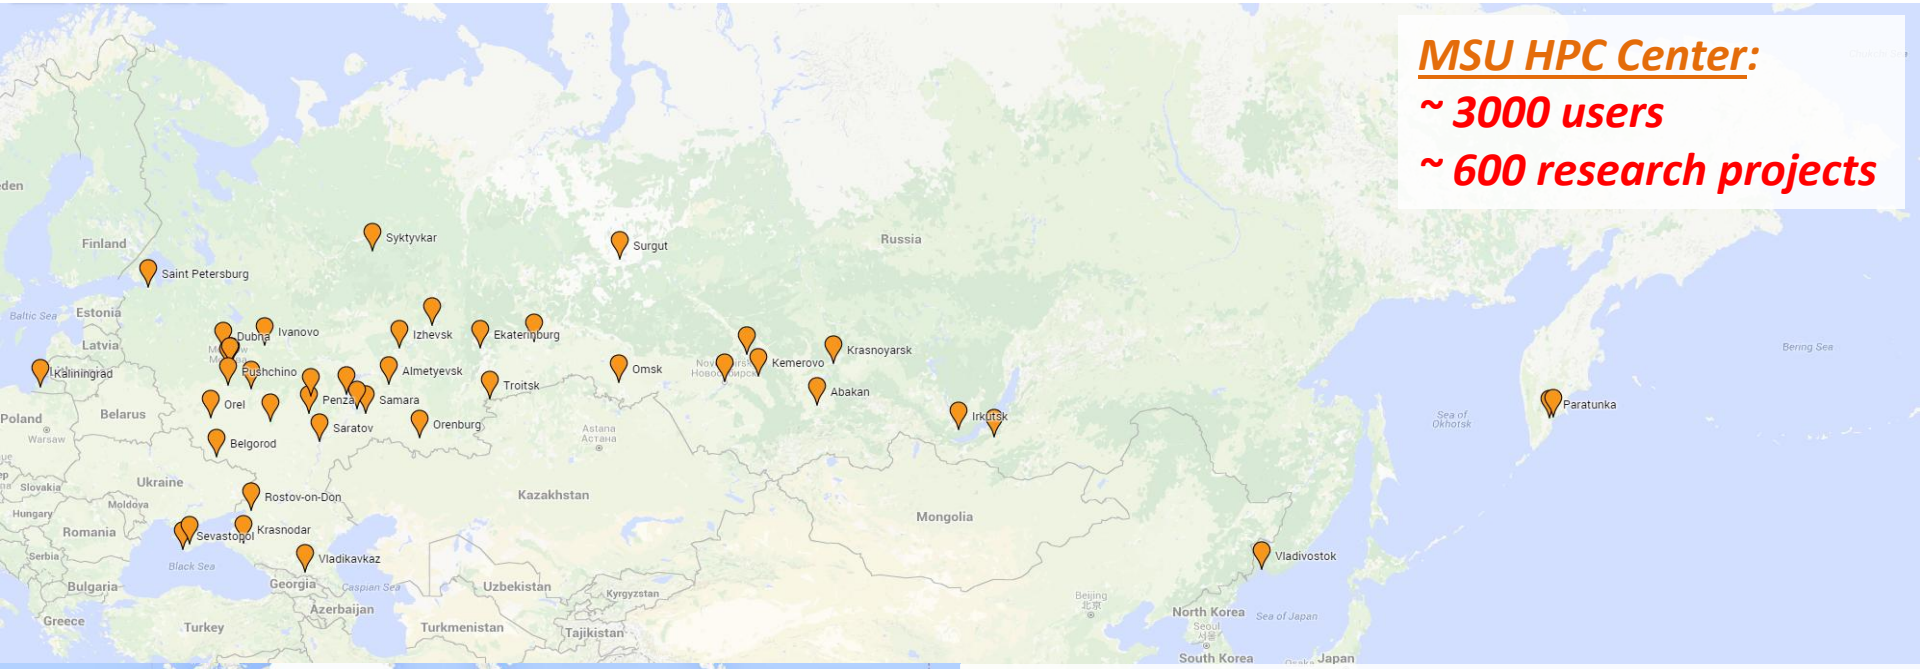

MSU HPC Center:
~ 3000 users
~ 600 research projects

- ***Scientific groups from 350+ organizations and institutions allover Russia (~50 cities)***

- ***Collaboration of leading teams and specialists all over the world:***
~150 collaborative projects with international scientific groups from over 90 locations worldwide

User Applications of MSU HPC Center

The Largest HPC Center in Russia

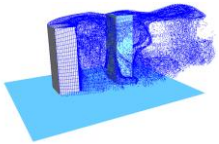
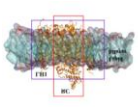

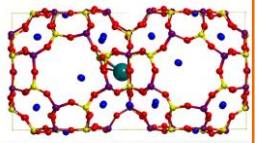
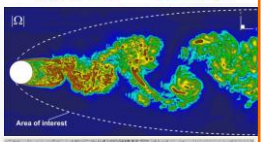
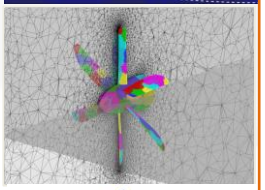
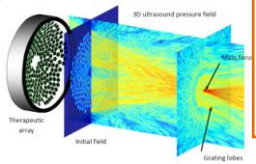

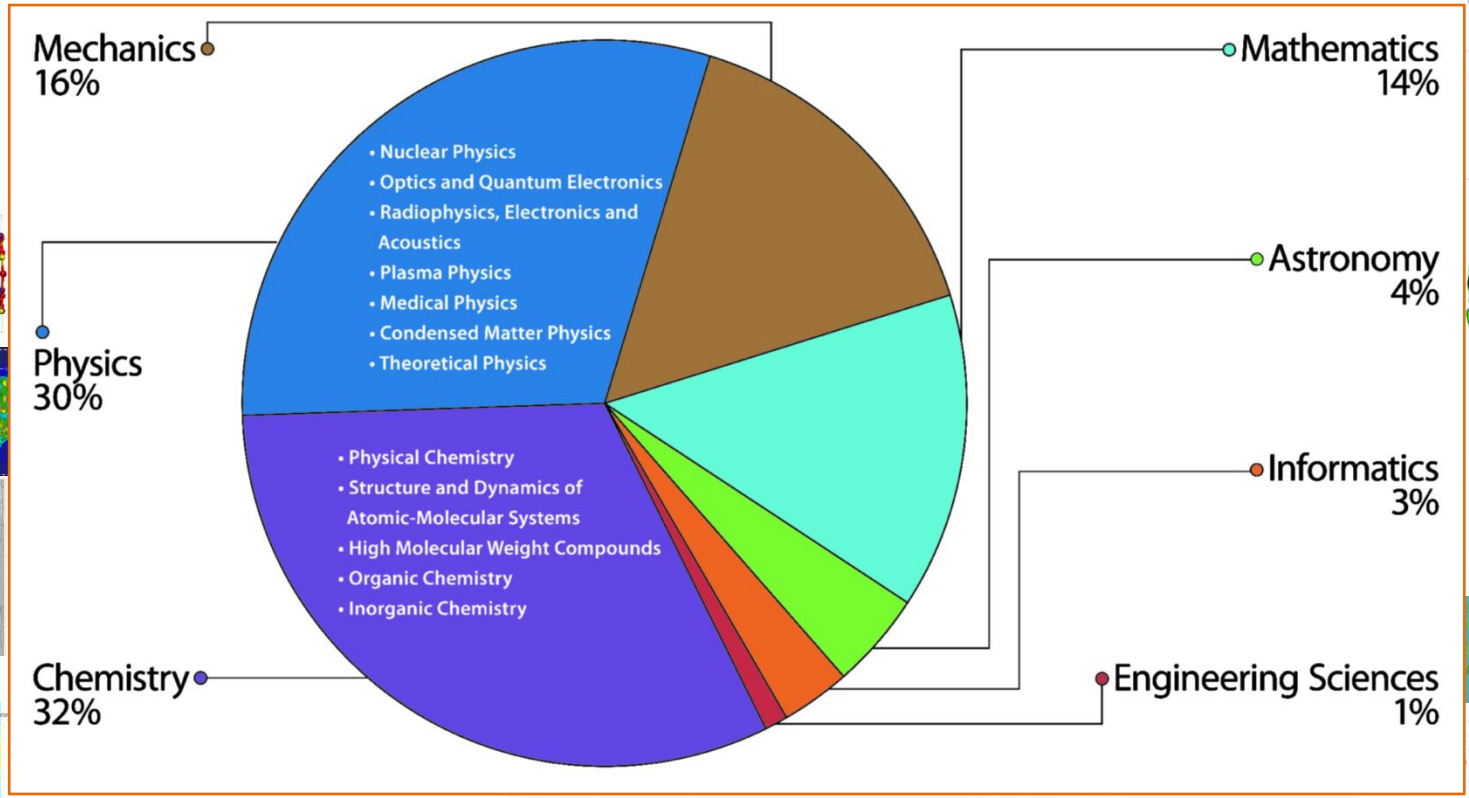

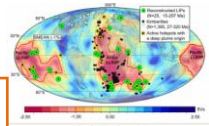
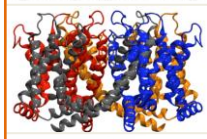
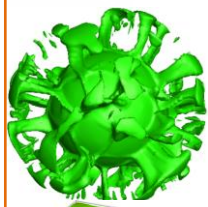
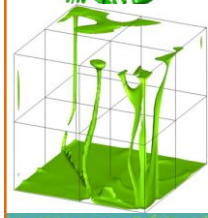
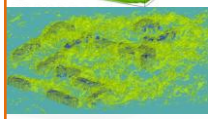
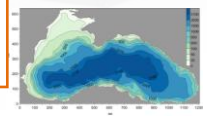

TOTAL DIVERSITY:

- *Programming Languages & Technologies*
- *Algorithms & Implementations*
- *Hardware & Software Envolved*

HPC.MSU.RU

Computing Facilities of MSU

Diversity of parallel programming technologies in use

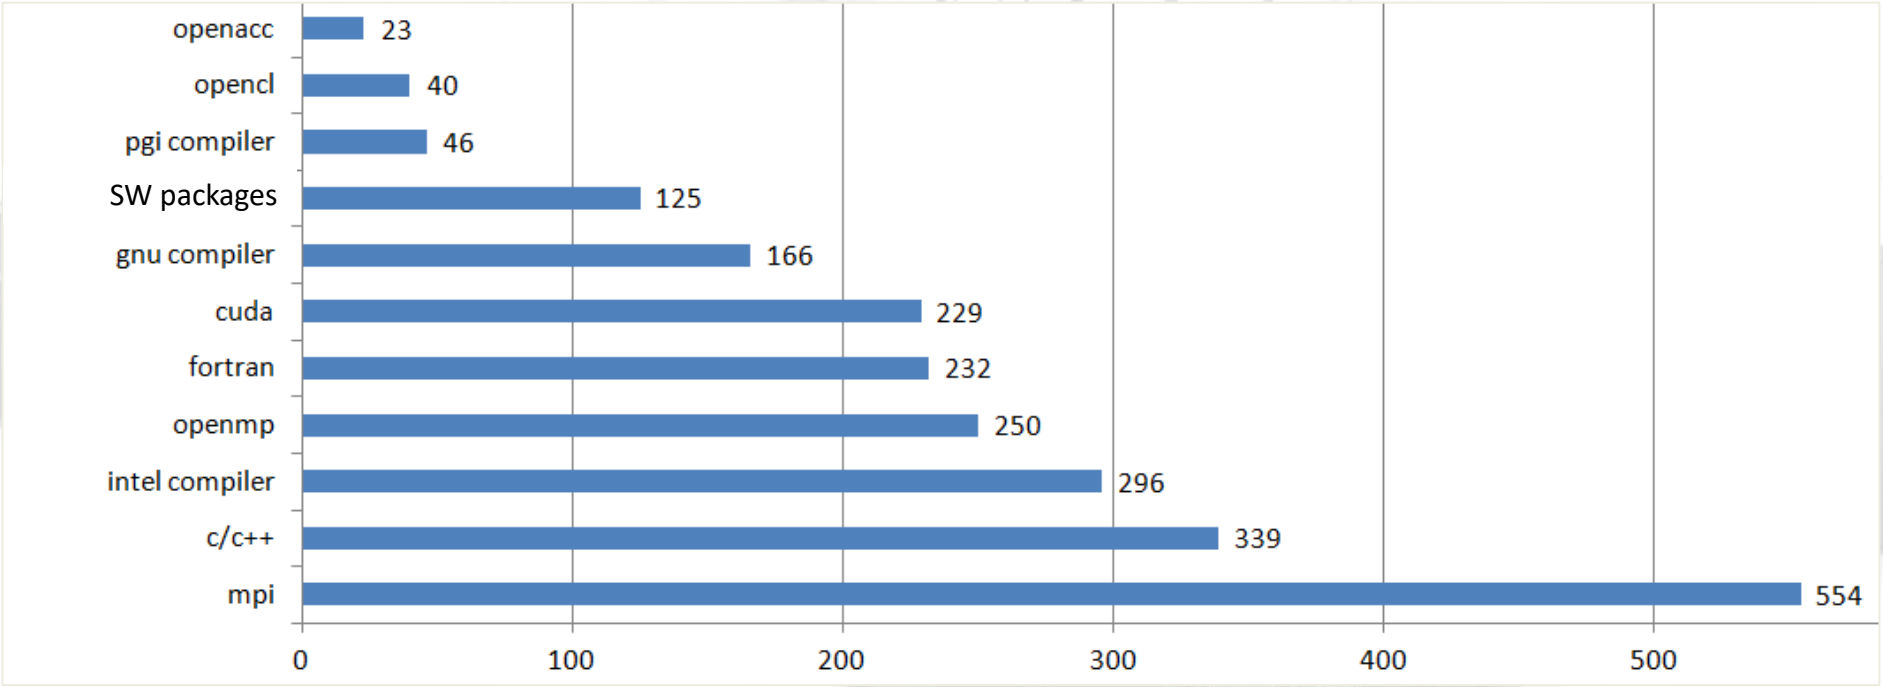

Computing Facilities of MSU

Diversity of software in use

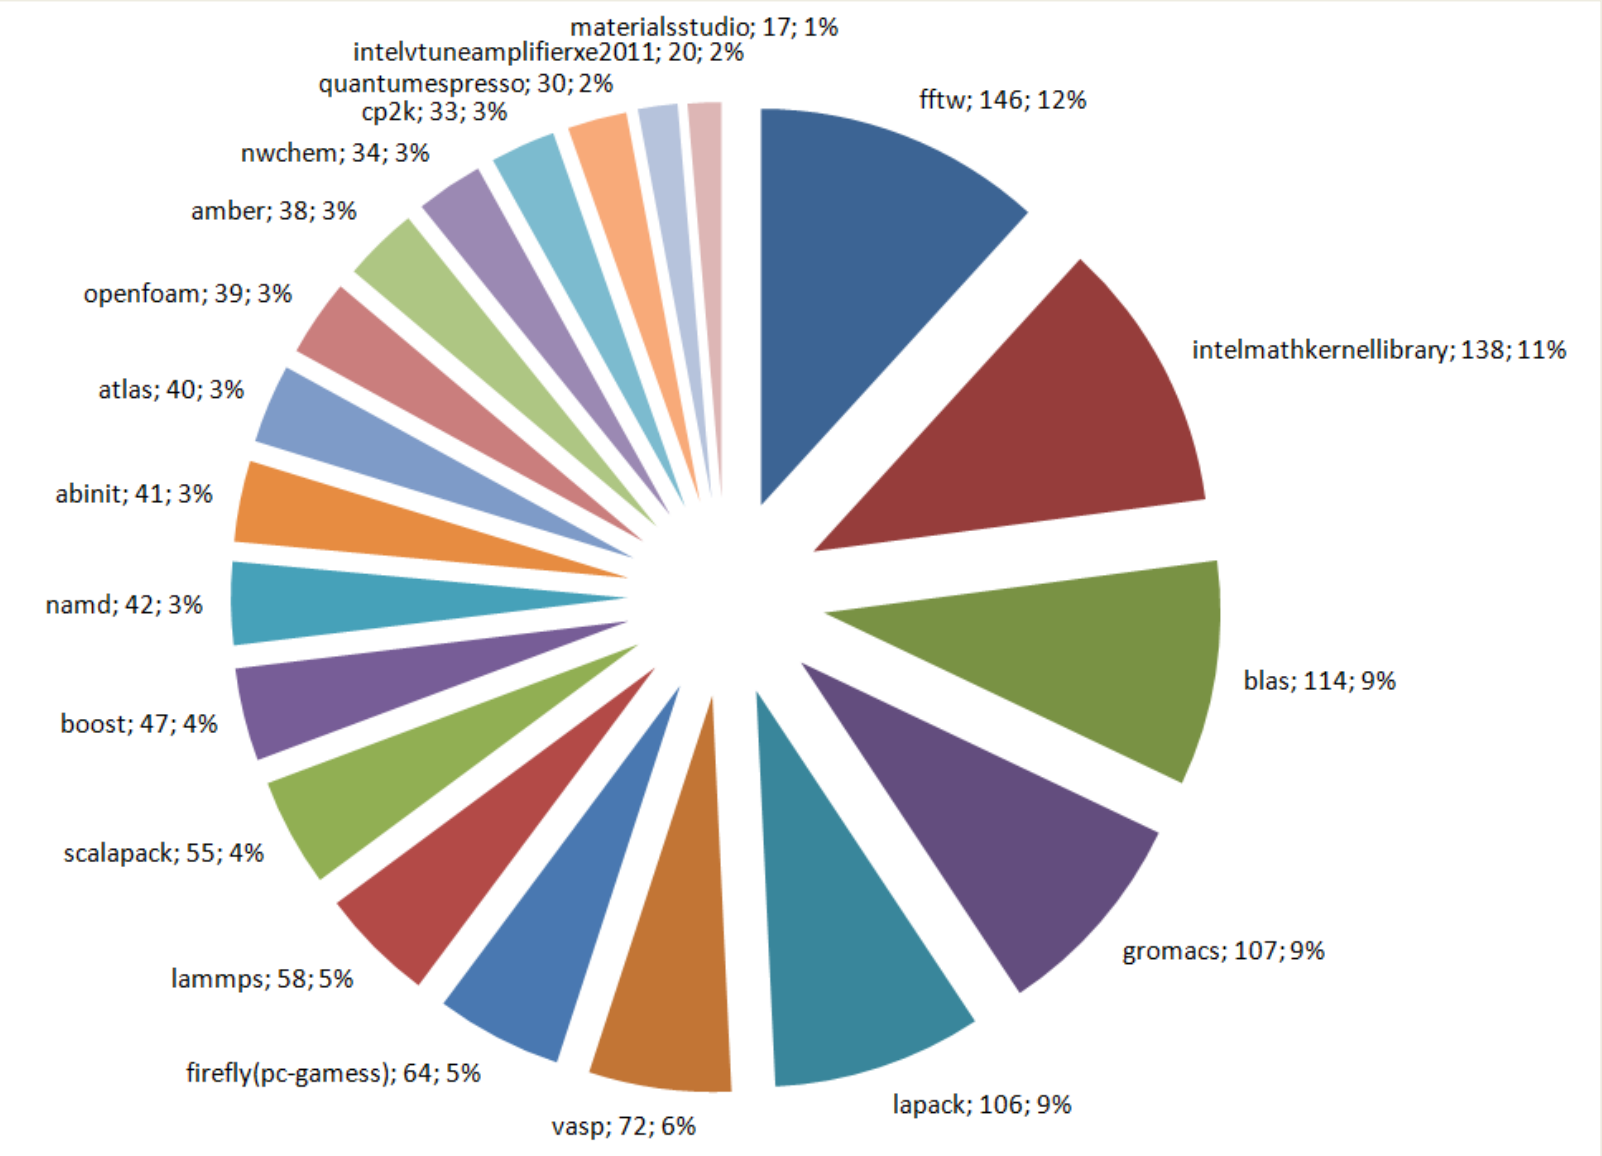

Operability and productivity

Monitoring and automated reactions needed

Idle system: expensive
Inefficient load: expensive
Late decisions: expensive
dangerous

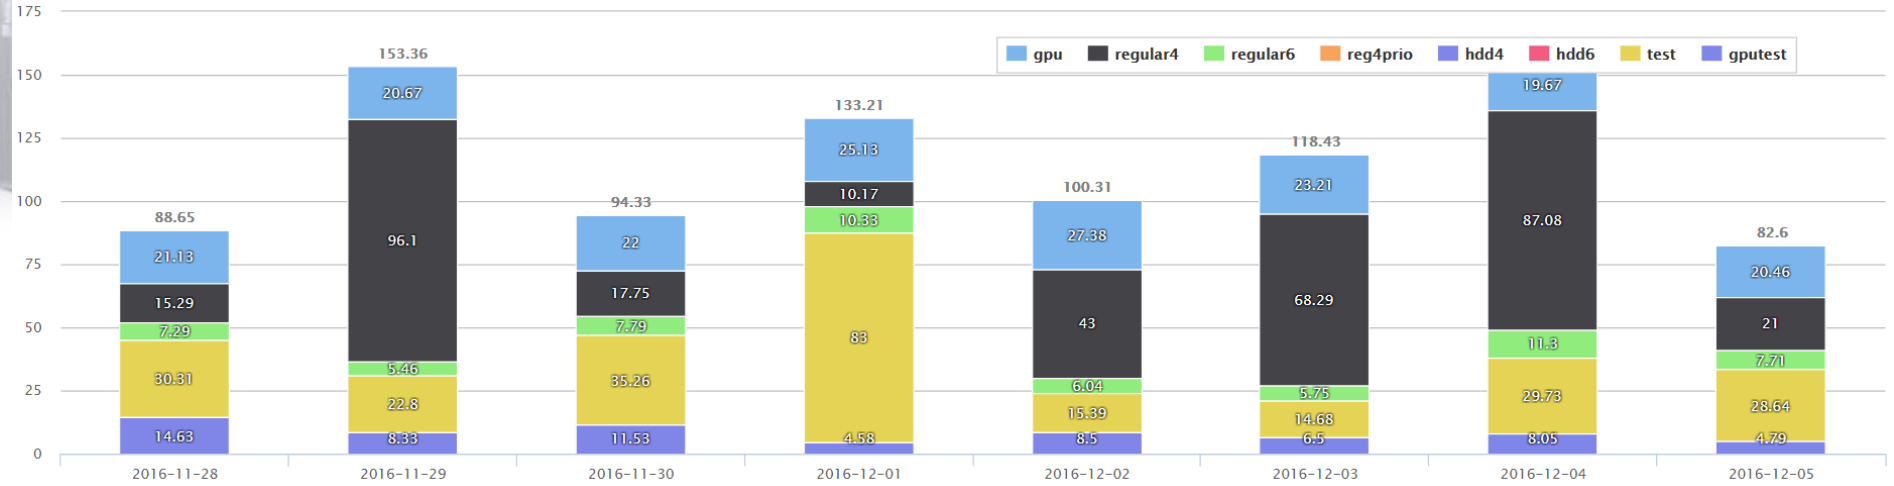

Recent “Lomonosov” stats:
100-150 simultaneously run jobs, over 1000 jobs per day

Operability and productivity

Everything plays role and contributes. Total monitoring.

It is imperative to permanently keep track on all components that influence efficiency of large-scale system output

- **Computing HW:** nodes, CPUs, memory stack, disks and storage hierarchy, networks, etc.
- **Infrastructure HW:** much more fault-tolerance critical. It includes cooling system: chillers, heat exchangers, air conditioners; piping, pumps; a set of components of the power system in conjunction with an uninterruptible power supply; fire safety systems and smoke removal; access control; etc.
- **Whole SW stack:** OS, package and license usage rates and peculiarities, etc.
- **Job scheduling** and queuing details from various points of view.
- **Job execution details:** dynamics and peculiarities of resource utilization by any and every executed user job.
- **Users.** Peculiarities of resource utilization, using available SW, cheating, etc.

Total Monitoring

Scopes of analysis

Time period:

- **Past**
  post-mortem analysis: performance analysis, event processing efficiency
- **Present**
  quick and immediate reactions, present state screening
- **Future**
  prediction and planning

Role-specific viewpoints and interests:

- **User**
  - Regular User
  - Project Manager
- **HPC Center**
  - Administrator
  - Expert / Supervisor / Reviewer
  - System Holder

Total Monitoring

Data Sources

- **Recent events** and HW/SW state streaming data for performance analysis & quick reactions

  (monitoring systems: DiMMon, Collectd, Nagios, Zabbix, etc.)
  - Compute nodes
  - Infrastructure

- **Historic coarse-grained data**
  (DataBases, CSV, JSON, logs, etc.)
  - Resource Managers (SLURM, CLEO, etc.)
  - Resilience Systems
  - User/Account Management Systems

Total monitoring of HPC systems

All-round and Holistic Approach practiced at MSU

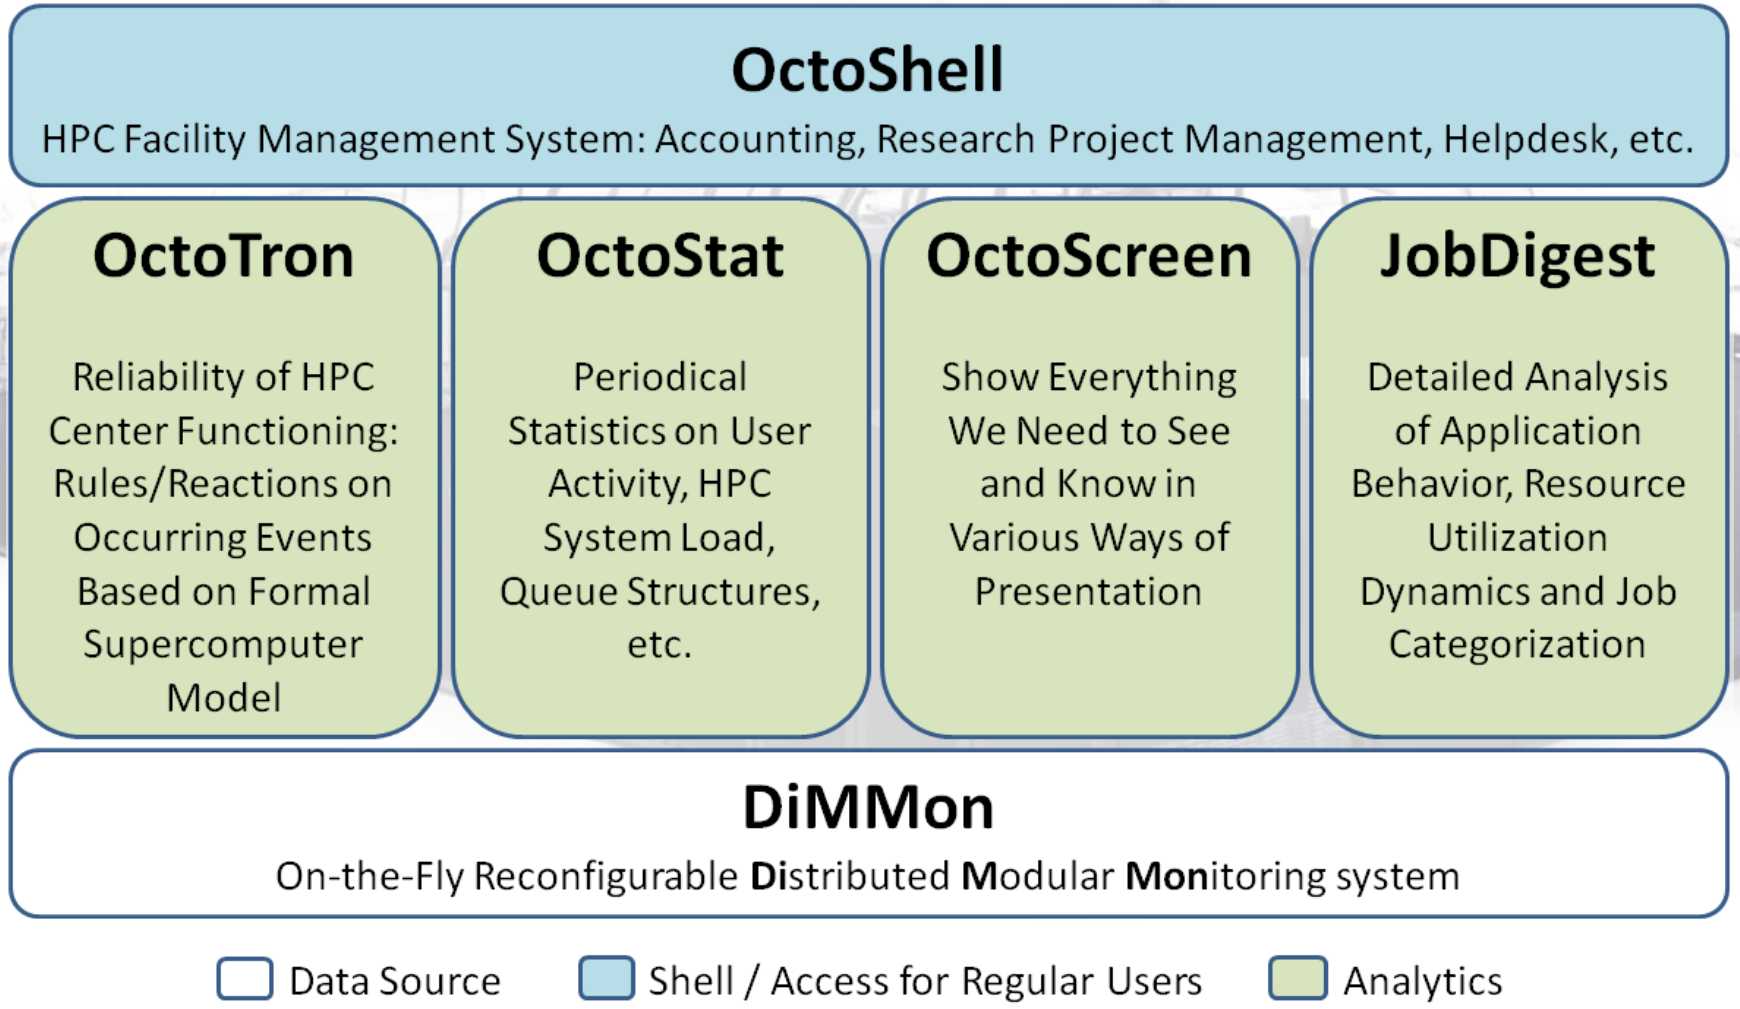

Approach Design Principles

Key feature - availability of sufficient info on every job

- **General information on every finished job must be available**, including data from resource manager and average rate of resource utilization (integral job characteristics) obtained from the monitoring system.
- The used monitoring system (sensor set, polling frequency and saved data coarsening) must be configured to grant **availability of job profiles for all finished jobs** right after execution with no resource-intensive post mortem operations.
- There must be means for **job marking and categorization** based on job characteristics in manual and automatic modes.
- Job information access restrictions must meet workflow regulations, supporting **various scopes of analysis**: regular user, research project manager, system administrator, etc.
- Flexible configuration, supporting **diverse data sources**.

General Job Characteristics

Key feature - availability of sufficient info on every job

- **General Job Information**
  - general job queue structure analysis (OctoStat)
  - stats on user/workgroup resource utilization (OctoStat, OctoShell)
  - based on resource manager data / logs
- **Dynamical Job Characteristics**
  - represent the dynamics of resource utilization during application execution
  - application execution resource utilization dynamics analysis, finding abnormal program behavior (JobDigest)
  - based on resource manager data + system monitoring data
- **Integral Job Characteristics**
  - represent average resource utilization during application execution
  - Resource utilization analysis for job collections (user/workgroup, partition, set of nodes, etc.)
  - based on resource manager data + system monitoring data

System monitoring:

- CPU user
- Flops
- LoadAverage
- Interconnect usage
- L1 replacements
- LLC misses
- Memory access rate
- GPU load
- etc.

Resource Manager:

- ID
- Owner
- Status (finished, cancelled, etc.)
- Timestamps (queued, started, finished)
- Command line (to parse for details)
- Partition
- Node set allocated
- Allocated number of cores
- Execution time and CPUh

Approach Design Principles

General workflow based on system monitoring data analysis

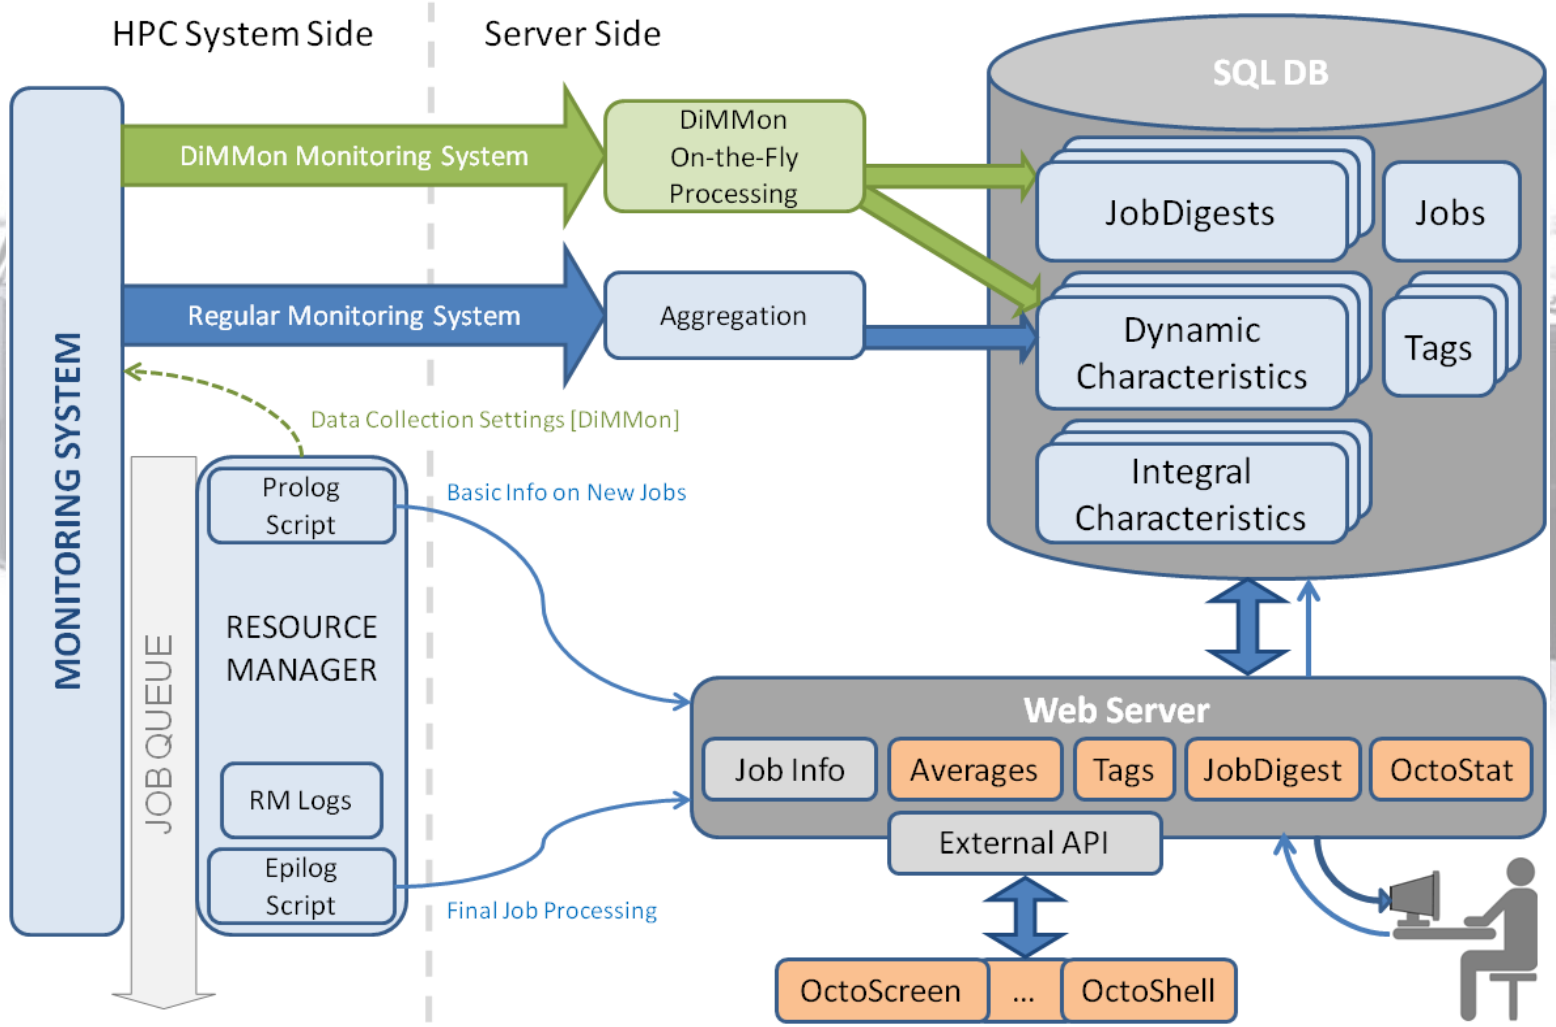

Job and Queue Structure Analysis Hierarchy

General information on execution history for all jobs

Integral job characteristics, job categories and collections, tags

Dynamic job characteristics for any and every job (JobDigest)

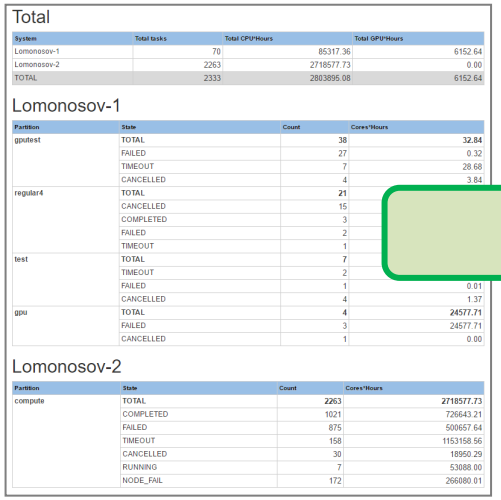

Total

System	Total tasks	Total CPU*Hours	Total GPU*Hours
Lomonosov-1	70	85317.36	6152.64
Lomonosov-2	2263	2718577.73	0.00
TOTAL	2333	2803895.08	6152.64

Lomonosov-1

Partition	State	Count	Cores*Hours
gputest	TOTAL	38	32.84
	FAILED	27	0.32
	TIMEOUT	7	28.68
	CANCELLED	4	3.84
regular4	TOTAL	21	
	CANCELLED	15	
	COMPLETED	3	
	FAILED	2	
	TIMEOUT	1	
test	TOTAL	7	
	TIMEOUT	2	
	FAILED	1	0.01
	CANCELLED	4	1.37
gpu	TOTAL	4	24577.71
	FAILED	3	24577.71
	CANCELLED	1	0.00

Lomonosov-2

Partition	State	Count	Cores*Hours
compute	TOTAL	2263	2718577.73
	COMPLETED	1021	726643.21
	FAILED	875	500657.64
	TIMEOUT	158	1153158.56
	CANCELLED	30	18950.29
	RUNNING	7	53088.00
	NODE_FAIL	172	266080.01

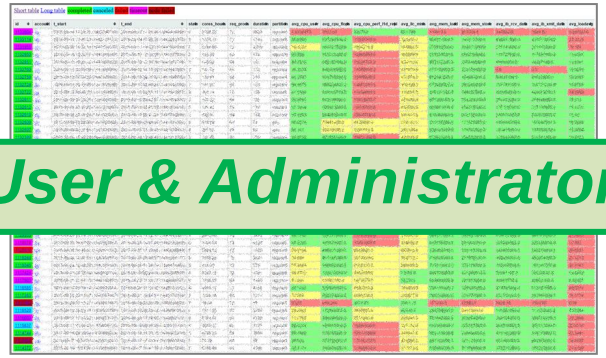

User & Administrator

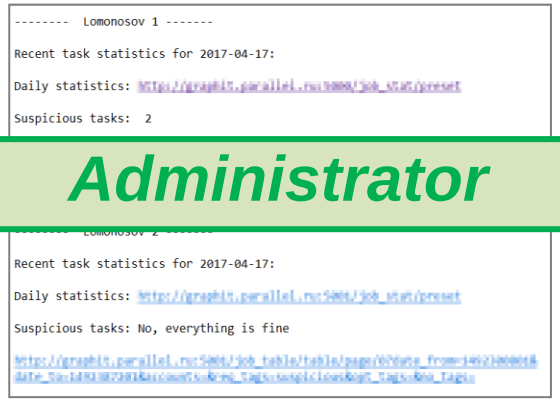

Administrator

Job History and Summary

Available for every user in personal account web page

Ability groups:

- **Regular user**
  - personal accounts only
- **Project manager**
  - all accounts of the workgroup
- **Expert**
- **Administrator**

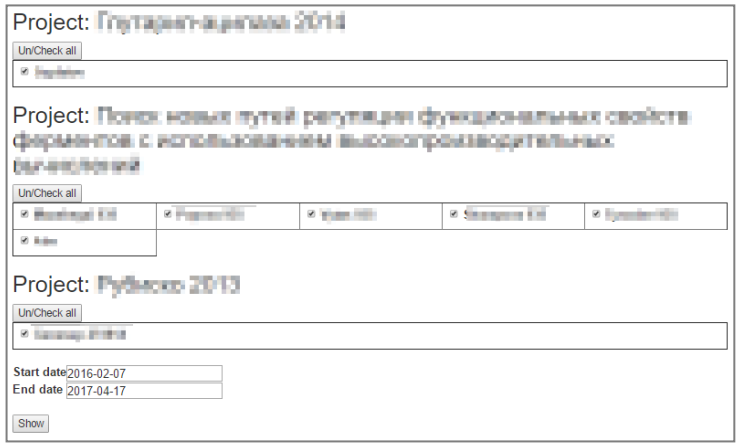

Total

System	Total tasks	Total CPU*Hours	Total GPU*Hours
Lomonosov-1	70	85317.36	6152.64
Lomonosov-2	2263	2718577.73	0.00
TOTAL	2333	2803895.08	6152.64

Lomonosov-1

Partition	State	Count	Cores*Hours
gputest	**TOTAL**	**38**	**32.84**
	FAILED	27	0.32
	TIMEOUT	7	28.68
	CANCELLED	4	3.84
regular4	**TOTAL**	**21**	**60697.23**
	CANCELLED	15	35631.08
	COMPLETED	3	7559.38
	FAILED	2	6582.61
	TIMEOUT	1	10924.16
test	**TOTAL**	**7**	**9.59**
	TIMEOUT	2	8.21
	FAILED	1	0.01
	CANCELLED	4	1.37
gpu	**TOTAL**	**4**	**24577.71**
	FAILED	3	24577.71
	CANCELLED	1	0.00

Lomonosov-2

Partition	State	Count	Cores*Hours
compute	**TOTAL**	**2263**	**2718577.73**
	COMPLETED	1021	726643.21
	FAILED	875	500657.64
	TIMEOUT	158	1153158.56
	CANCELLED	30	18950.29
	RUNNING	7	53088.00
	NODE_FAIL	172	266080.01

Workgroup Job Collection Details

Research project jobs summary

TIME PERIOD	2016-01-01 18:47:05 -- 2016-12-09 18:47:05
LOGINS	shvets vadim coctic serg
TOTAL CPU TIME (CPUh)	805.79
USER JOBS	178
COMPLETED	167
CANCELLED	4
TIMEOUT	0
FAILED	7
NODE_FAIL	0

"Lomonosov-2": CPU load, GPU load, LA, IB MPI, IB FS

id	account	t_start	t_end	state	cores_hours	num_cores	duration	partition	cpu_user	gpu_load	loadavg	ib_rcv_data_mpi	ib_rcv_data_fs
122141	shvets	2016-12-09 17:48:13	2016-12-09 18:33:51	COMPLETED	42.59	56	45.63	compute	46.26	35.73	13.96	1030631127.20	1030724468.14
122062	vadim	2016-12-09 15:50:32	2016-12-09 17:50:51	COMPLETING	28.07	14	120.32	compute	0.04	0.00	0.03	0.00	276.98
122064	vadim	2016-12-09 15:50:33	2016-12-09 17:50:51	COMPLETING	28.07	14	120.30	compute	0.11	0.00	0.02	0.00	279.83
122084	shvets	2016-12-09 16:58:45	2016-12-09 17:38:18	COMPLETING	9.23	14	39.55	compute	0.04	0.00	0.00	0.00	187.13
122049	shvets	2016-12-09 15:45:11	2016-12-09 16:23:57	COMPLETED	36.18	56	38.77	compute	45.87	25.99	13.98	939884640.36	925061313.08
122072	vadim	2016-12-09 15:58:46	2016-12-09 16:09:24	COMPLETED	2.48	14	10.63	compute	18.24	0.00	13.98	0.00	485.15
122073	vadim	2016-12-09 15:58:47	2016-12-09 16:08:56	COMPLETED	2.37	14	10.15	compute	19.45	0.00	16.03	0.00	6328.39
122070	vadim	2016-12-09 15:52:10	2016-12-09 16:05:58	COMPLETED	3.22	14	13.80	compute	11.80	99.00	8.56	0.00	457.38
122067	vadim	2016-12-09 15:58:45	2016-12-09 16:05:51	COMPLETED	3.31	28	7.10	compute	38.35	0.00	15.57	152445207.65	133954687.73
122071	vadim	2016-12-09 15:55:01	2016-12-09 16:05:33	COMPLETED	2.46	14	10.53	compute	19.16	0.00	12.30	0.00	494.26
122063	vadim	2016-12-09 15:52:05	2016-12-09 15:59:56	COMPLETED	7.33	56	7.85	compute	47.72	0.00	17.33	187262462.26	187675724.51
122066	vadim	2016-12-09 15:52:06	2016-12-09 15:59:54	COMPLETED	3.64	28	7.80	compute	24.66	99.12	11.44	164043411.16	166637443.66
122069	vadim	2016-12-09 15:52:08	2016-12-09 15:59:20	COMPLETED	3.36	28	7.20	compute	45.34	99.00	19.73	152720283.82	147813432.90
122068	vadim	2016-12-09 15:52:07	2016-12-09 15:59:14	COMPLETED	3.32	28	7.12	compute	39.08	99.29	18.35	226869299.78	229201610.44
122065	vadim	2016-12-09 15:50:34	2016-12-09 15:58:37	COMPLETED	7.51	56	8.05	compute	50.04	0.00	21.64	167065425.74	167085160.13

Highlighting statuses (left column) and average resource utilization (right colored block)

Job Details

Information on a job collection

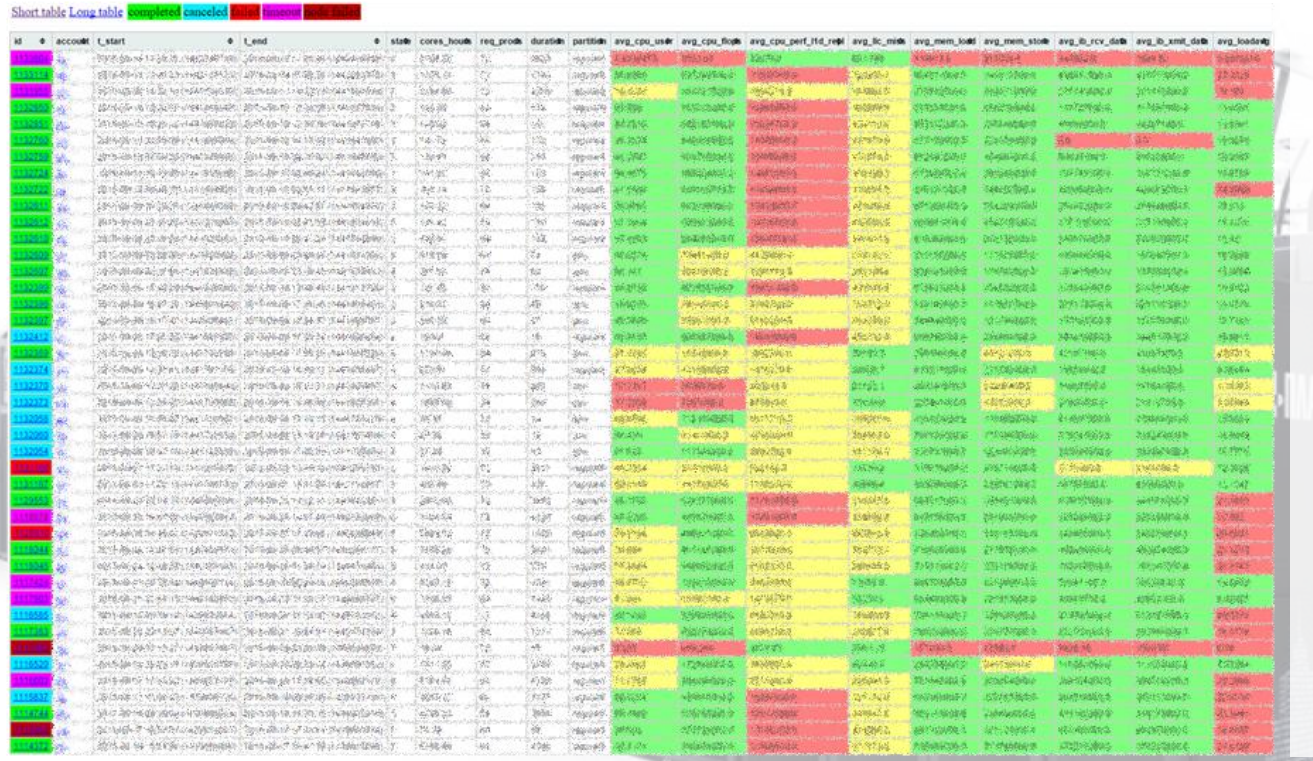

“Lomonosov” system average resource utilization rate (left to right):

- CPU user
- Flops
- L1 replacements
- LLC misses
- Memory read&write
- IB send&receive
- LoadAverage

Highlighting statuses (left column) and average resource utilization (right colored block)

Categories of similar jobs by combination of criteria and tagging:

- Formal criteria (mostly thresholds)
- Machine learning

Information on a job collection

auto_low_gpu_load + partition_GPU

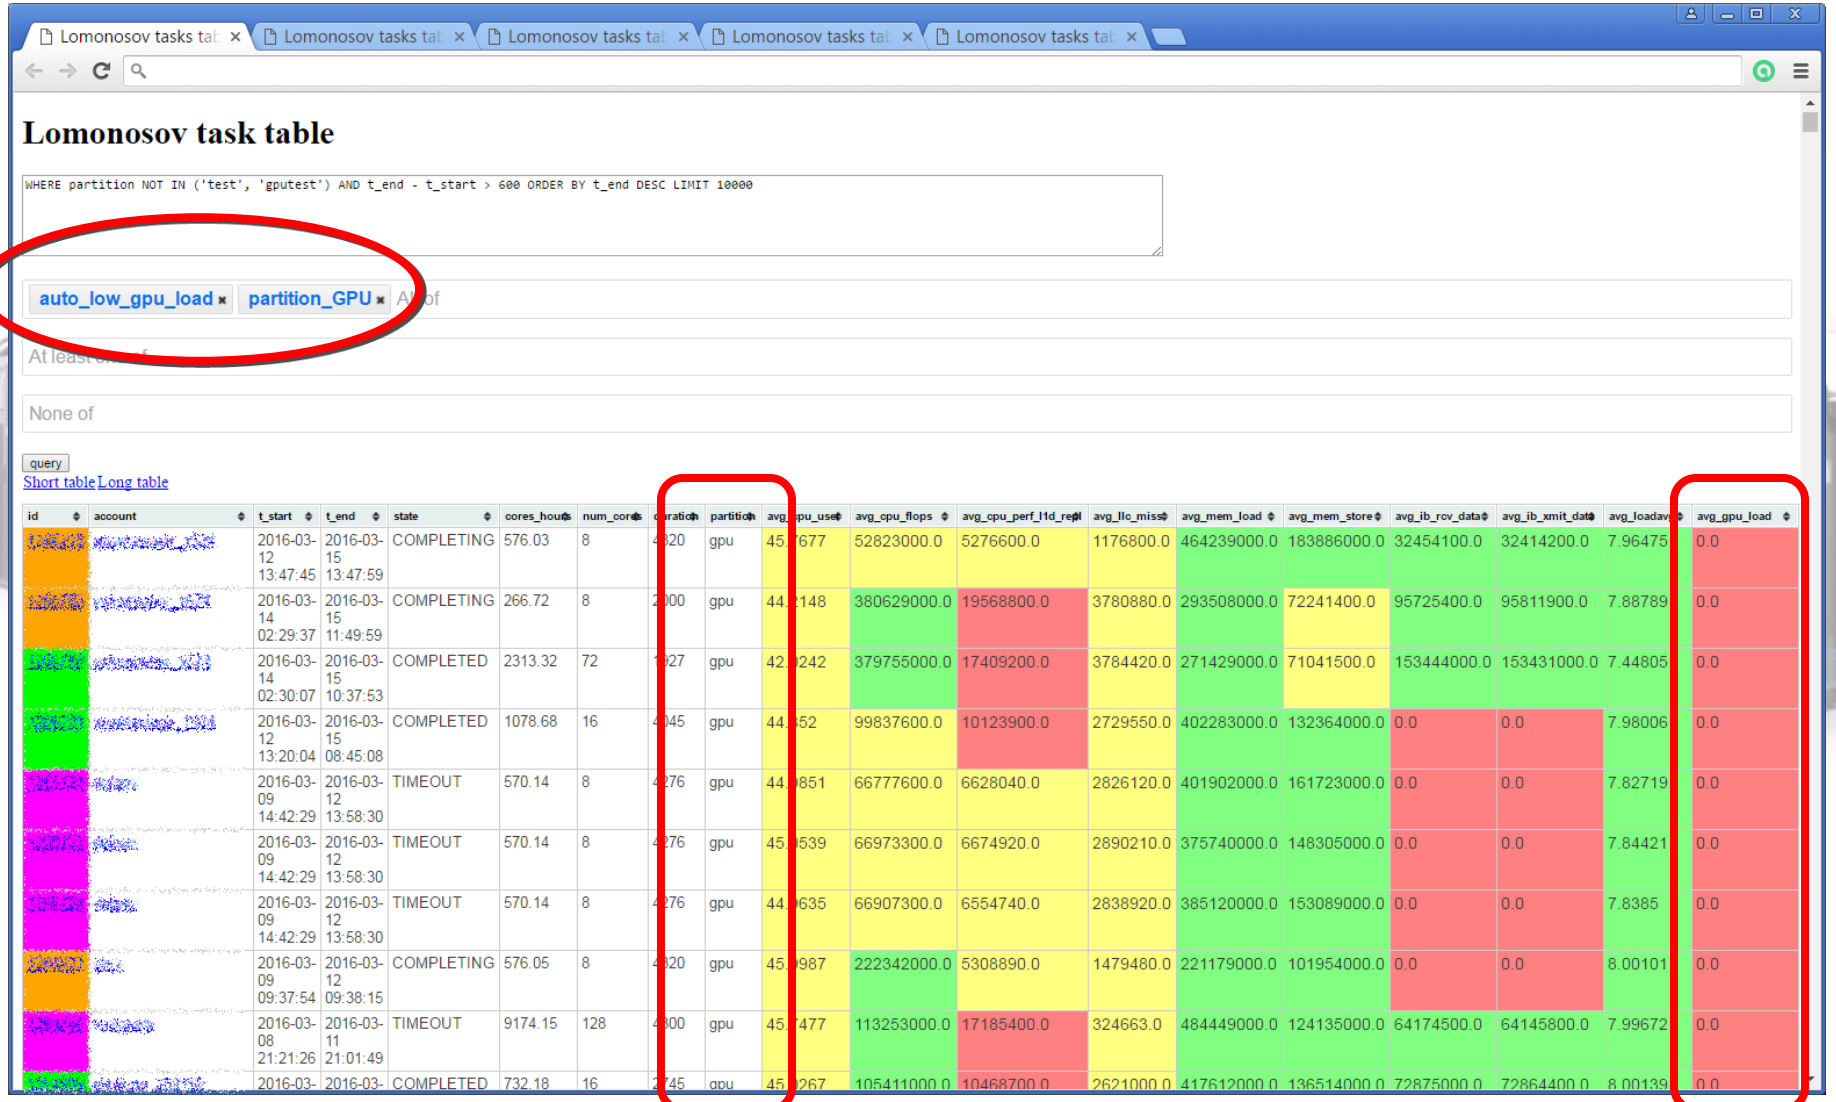

Information on a job collection

Single core jobs not in test partitions

auto_avg_LA_SINGLE_CORE × Add tags to filter the table

query

Short table Long table

single node **CPU_user** **LoadAvg**

id	account	t_start	t_end	state	cores_hours	num_cores	duration	partition	avg_cpu_user	avg_cpu_flops	avg_cpu_perf_l1d_repl	avg_llc_miss	avg_mem_load	avg_mem_store	avg_ib_rcv_data	avg_ib_xmit_data	avg_loadavg
1238308	[illegible]	2016-02-04 06:25:13	2016-02-04 07:51:49	COMPLETED	11.55	8	86	regular4	6.52	8979330.0	183216.0	118759.0	37346000.0	22979700.0	1732.93	1967.65	1.00467
1236734	[illegible]	2016-02-03 05:54:57	2016-02-03 07:58:22	FAILED	16.46	8	123	regular4	6.07174	28813500.0	1132060.0	101059.0	30475200.0	13338400.0	1719400.0	4526380.0	0.994348
1237308	[illegible]	2016-02-03 00:59:36	2016-02-03 04:24:47	COMPLETED	27.36	8	205	hdd4	0.0923077	1498.76	30833.7	1890.64	594213.0	245048.0	80744.6	79506.6	0.994103
1237180	[illegible]	2016-02-02 22:01:55	2016-02-02 22:22:17	COMPLETED	2.72	8	20	regular4	10.27	8689710.0	166478.0	124506.0	37970300.0	23214700.0	2142.9	2049.93	1.0
1237093	[illegible]	2016-02-02 18:53:35	2016-02-02 19:54:02	TIMEOUT	8.06	8	60	regular4	7.03	8864480.0	153599.0	96245.3	37148900.0	22867400.0	0.0	0.0	1.001
1237055	[illegible]	2016-02-02 16:58:39	2016-02-02 17:33:38	FAILED	4.66	8	34	regular4	6.12	9650980.0	56315.4	20220.9	36315500.0	22835900.0	0.0	0.0	0.99
1235368	[illegible]	2016-01-28 14:51:53	2016-01-31 14:52:04	TIMEOUT	576.02	8	4320	regular4	6.19784	7540280.0	287333.0	32770.1	8116410.0	4075520.0	50941.7	2798470.0	0.994907
1236081	[illegible]	2016-01-31 02:07:46	2016-01-31 05:33:30	COMPLETED	27.43	8	205	regular4	0.0489744	25378.4	177844.0	46266.7	19580800.0	12530700.0	54673600.0	844684.0	0.900513
1230559	[illegible]	2016-01-24 08:28:31	2016-01-27 08:28:33	COMPLETING	576.00	8	4320						[illegible]20100.0	5824880.0	2890.48	2911.17	0.98058
1230224	[illegible]	2016-01-23 09:14:31	2016-01-25 05:30:34	COMPLETED	354.14	8	[illegible]						[illegible]63900.0	12936200.0	5563.94	5926180.0	1.00758
1230057	[illegible]	2016-01-23 03:22:49	2016-01-24 18:48:44	COMPLETING	315.46	8	2365	regular4	6.08654	311.465	6483.29	409.6	114895000.0	25839200.0	1704.98	1907.76	1.01085
1230597	[illegible]	2016-01-24 13:02:31	2016-01-24 13:28:52	FAILED	3.51	8	26	hdd4	0.253333	116354.0	2007430.0	539.793	51266300.0	32041400.0	3184.72	2307.22	0.986667
1230097	[illegible]	2016-01-23 06:34:42	2016-01-23 11:09:05	COMPLETING	36.58	8	274	regular4	0.0424528	310.48	8636.37	420.936	39045500.0	6020430.0	1760.95	1911.44	1.00528
1230221	[illegible]	2016-01-23 09:14:31	2016-01-23 09:46:39	FAILED	4.28	8	32	regular4	5.418	22515000.0	1447870.0	83310.4	33935100.0	16352300.0	7107.31	149630.0	1.002
1230145	[illegible]	2016-01-22 21:52:12	2016-01-23 02:52:32	TIMEOUT	40.04	8	300	regular4	0.825517	113100.0	1877160.0	7024.57	49070900.0	30596800.0	46279200.0	133681.0	1.01707
1230198	[illegible]	2016-01-23 00:16:42	2016-01-23 00:53:46	COMPLETED	4.94	8	37	hdd4	0.904	51709.4	2285050.0	3309.51	42122700.0	24885900.0	25130.5	2207980.0	1.048
1230194	[illegible]	2016-01-23 00:07:12	2016-01-23 00:49:11	COMPLETED	5.60	8	41	hdd4	0.923333	49520.7	2304560.0	4253.29	42744700.0	25238900.0	23307.5	2071330.0	1.03
1229822	[illegible]	2016-01-22 13:03:29	2016-01-22 14:43:45	TIMEOUT	13.37	8	100	regular4	0.639444	118970.0	1941090.0	5203.37	50921000.0	31777900.0	54200400.0	107747.0	1.00778
1229238	[illegible]	2016-01-21 12:44:15	2016-01-21 15:25:46	FAILED	21.54	8	161	regular4	0.462903	377364.0	1954350.0	1676.62	49992900.0	31221100.0	3537700.0	32939.6	1.01516

Up to 3 days long!

Information on a job collection

auto_avg_CPU_idle TOO HIGH (CPU idle > 25%)

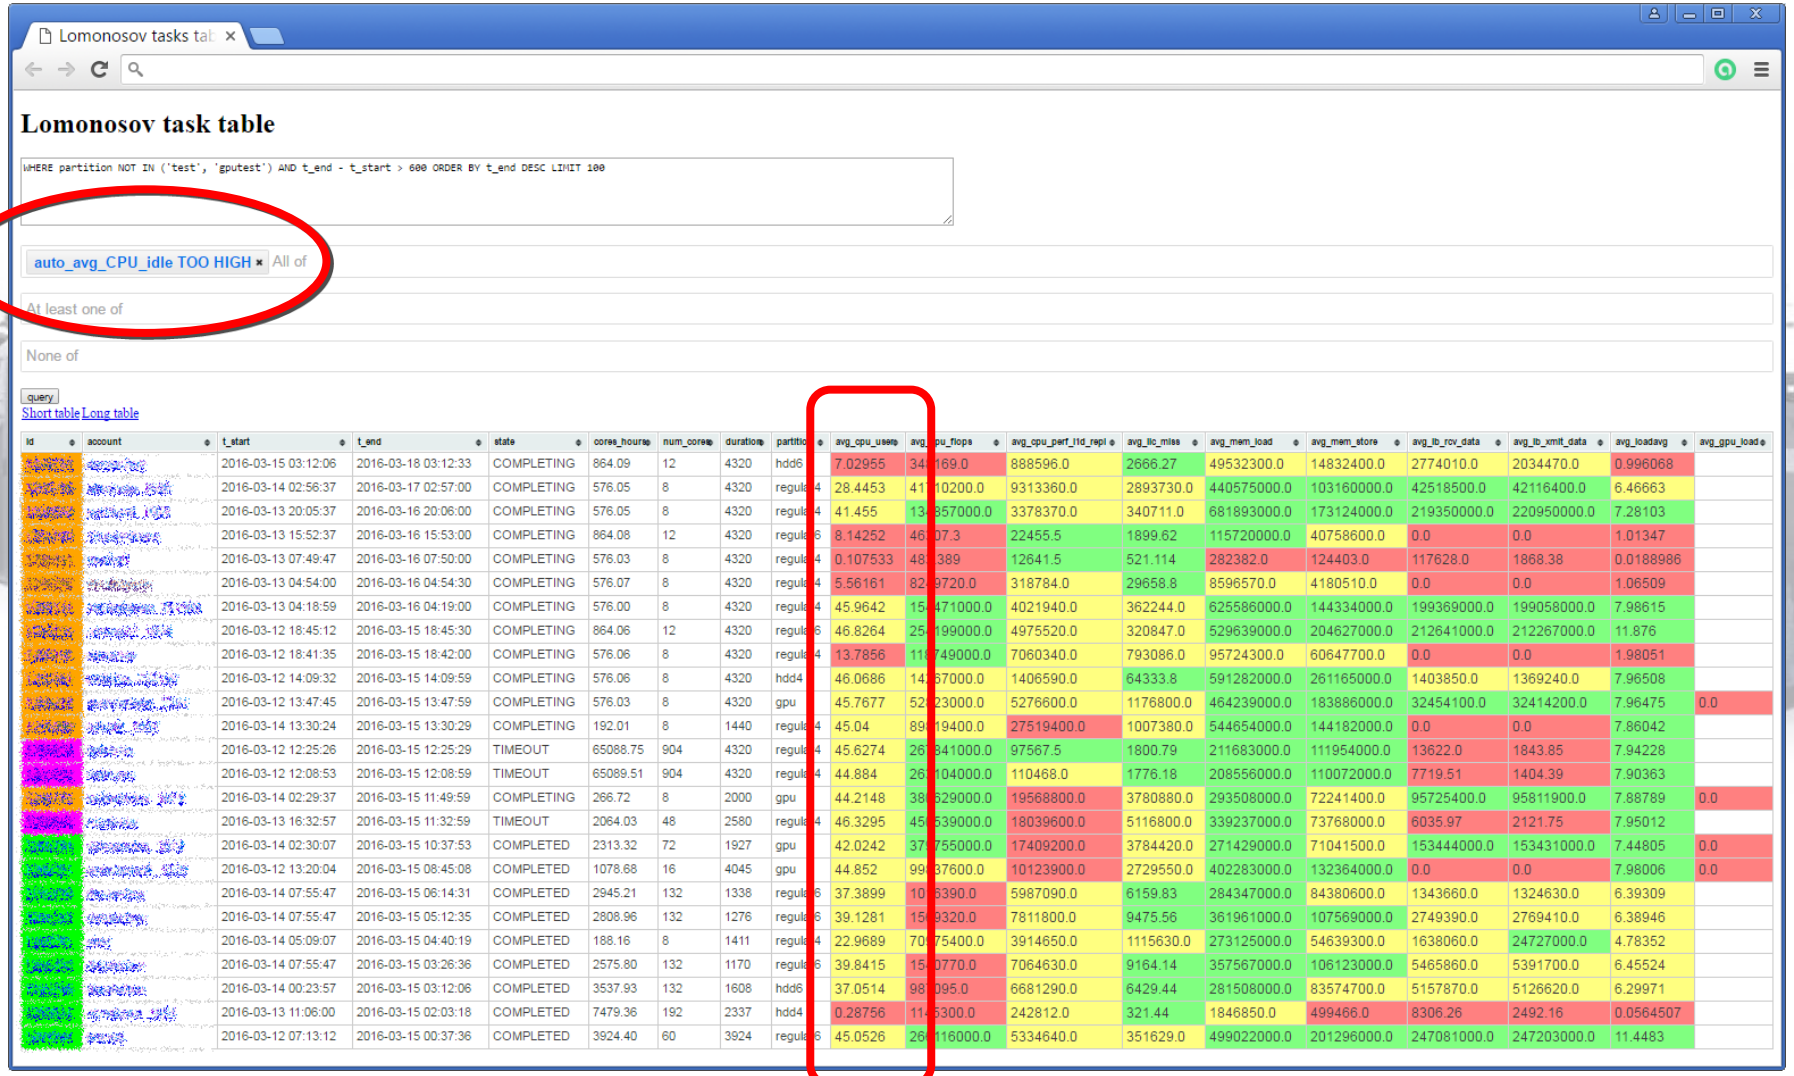

id	account	t_start	t_end	state	cores_hours	num_cores	duration	partition	avg_cpu_usem	avg_cpu_flops	avg_cpu_perf_l1d_repl	avg_llc_miss	avg_mem_load	avg_mem_store	avg_lb_rcv_data	avg_lb_xmit_data	avg_loadavg	avg_gpu_load
[illegible]	[illegible]	2016-03-15 03:12:06	2016-03-18 03:12:33	COMPLETING	864.09	12	4320	hdd6	7.02955	34[illegible]169.0	888596.0	2666.27	49532300.0	14832400.0	2774010.0	2034470.0	0.996068	
[illegible]	[illegible]	2016-03-14 02:56:37	2016-03-17 02:57:00	COMPLETING	576.05	8	4320	regula4	28.4453	41[illegible]10200.0	9313360.0	2893730.0	440575000.0	103160000.0	42518500.0	42116400.0	6.46663	
[illegible]	[illegible]	2016-03-13 20:05:37	2016-03-16 20:06:00	COMPLETING	576.05	8	4320	regula4	41.455	13[illegible]857000.0	3378370.0	340711.0	681893000.0	173124000.0	219350000.0	220950000.0	7.28103	
[illegible]	[illegible]	2016-03-13 15:52:37	2016-03-16 15:53:00	COMPLETING	864.08	12	4320	regula6	8.14252	46[illegible]07.3	22455.5	1899.62	115720000.0	40758600.0	0.0	0.0	1.01347	
[illegible]	[illegible]	2016-03-13 07:49:47	2016-03-16 07:50:00	COMPLETING	576.03	8	4320	regula4	0.107533	48[illegible]389	12641.5	521.114	282382.0	124403.0	117628.0	1868.38	0.0188986	
[illegible]	[illegible]	2016-03-13 04:54:00	2016-03-16 04:54:30	COMPLETING	576.07	8	4320	regula4	5.56161	82[illegible]9720.0	318784.0	29658.8	8596570.0	4180510.0	0.0	0.0	1.06509	
[illegible]	[illegible]	2016-03-13 04:18:59	2016-03-16 04:19:00	COMPLETING	576.00	8	4320	regula4	45.9642	15[illegible]471000.0	4021940.0	362244.0	625586000.0	144334000.0	199369000.0	199058000.0	7.98615	
[illegible]	[illegible]	2016-03-12 18:45:12	2016-03-15 18:45:30	COMPLETING	864.06	12	4320	regula6	46.8264	25[illegible]199000.0	4975520.0	320847.0	529639000.0	204627000.0	212641000.0	212267000.0	11.876	
[illegible]	[illegible]	2016-03-12 18:41:35	2016-03-15 18:42:00	COMPLETING	576.06	8	4320	regula4	13.7856	11[illegible]749000.0	7060340.0	793086.0	95724300.0	60647700.0	0.0	0.0	1.98051	
[illegible]	[illegible]	2016-03-12 14:09:32	2016-03-15 14:09:59	COMPLETING	576.06	8	4320	hdd4	46.0686	14[illegible]67000.0	1406590.0	64333.8	591282000.0	261165000.0	1403850.0	1369240.0	7.96508	
[illegible]	[illegible]	2016-03-12 13:47:45	2016-03-15 13:47:59	COMPLETING	576.03	8	4320	gpu	45.7677	52[illegible]23000.0	5276600.0	1176800.0	464239000.0	183886000.0	32454100.0	32414200.0	7.96475	0.0
[illegible]	[illegible]	2016-03-14 13:30:24	2016-03-15 13:30:29	COMPLETING	192.01	8	1440	regula4	45.04	89[illegible]19400.0	27519400.0	1007380.0	544654000.0	144182000.0	0.0	0.0	7.86042	
[illegible]	[illegible]	2016-03-12 12:25:26	2016-03-15 12:25:29	TIMEOUT	65088.75	904	4320	regula4	45.6274	26[illegible]841000.0	97567.5	1800.79	211683000.0	111954000.0	13622.0	1843.85	7.94228	
[illegible]	[illegible]	2016-03-12 12:08:53	2016-03-15 12:08:59	TIMEOUT	65089.51	904	4320	regula4	44.884	26[illegible]104000.0	110468.0	1776.18	208556000.0	110072000.0	7719.51	1404.39	7.90363	
[illegible]	[illegible]	2016-03-14 02:29:37	2016-03-15 11:49:59	COMPLETING	266.72	8	2000	gpu	44.2148	38[illegible]629000.0	19568800.0	3780880.0	293508000.0	72241400.0	95725400.0	95811900.0	7.88789	0.0
[illegible]	[illegible]	2016-03-13 16:32:57	2016-03-15 11:32:59	TIMEOUT	2064.03	48	2580	regula4	46.3295	45[illegible]539000.0	18039600.0	5116800.0	339237000.0	73768000.0	6035.97	2121.75	7.95012	
[illegible]	[illegible]	2016-03-14 02:30:07	2016-03-15 10:37:53	COMPLETED	2313.32	72	1927	gpu	42.0242	37[illegible]755000.0	17409200.0	3784420.0	271429000.0	71041500.0	153444000.0	153431000.0	7.44805	0.0
[illegible]	[illegible]	2016-03-12 13:20:04	2016-03-15 08:45:08	COMPLETED	1078.68	16	4045	gpu	44.852	99[illegible]37600.0	10123900.0	2729550.0	402283000.0	132364000.0	0.0	0.0	7.98006	0.0
[illegible]	[illegible]	2016-03-14 07:55:47	2016-03-15 06:14:31	COMPLETED	2945.21	132	1338	regula6	37.3899	10[illegible]6390.0	5987090.0	6159.83	284347000.0	84380600.0	1343660.0	1324630.0	6.39309	
[illegible]	[illegible]	2016-03-14 07:55:47	2016-03-15 05:12:35	COMPLETED	2808.96	132	1276	regula6	39.1281	15[illegible]9320.0	7811800.0	9475.56	361961000.0	107569000.0	2749390.0	2769410.0	6.38946	
[illegible]	[illegible]	2016-03-14 05:09:07	2016-03-15 04:40:19	COMPLETED	188.16	8	1411	regula4	22.9689	70[illegible]75400.0	3914650.0	1115630.0	273125000.0	54639300.0	1638060.0	24727000.0	4.78352	
[illegible]	[illegible]	2016-03-14 07:55:47	2016-03-15 03:26:36	COMPLETED	2575.80	132	1170	regula6	39.8415	15[illegible]0770.0	7064630.0	9164.14	357567000.0	106123000.0	5465860.0	5391700.0	6.45524	
[illegible]	[illegible]	2016-03-14 00:23:57	2016-03-15 03:12:06	COMPLETED	3537.93	132	1608	hdd6	37.0514	98[illegible]095.0	6681290.0	6429.44	281508000.0	83574700.0	5157870.0	5126620.0	6.29971	
[illegible]	[illegible]	2016-03-13 11:06:00	2016-03-15 02:03:18	COMPLETED	7479.36	192	2337	hdd4	0.28756	11[illegible]5300.0	242812.0	321.44	1846850.0	499466.0	8306.26	2492.16	0.0564507	
[illegible]	[illegible]	2016-03-12 07:13:12	2016-03-15 00:37:36	COMPLETED	3924.40	60	3924	regula6	45.0526	26[illegible]116000.0	5334640.0	351629.0	499022000.0	201296000.0	247081000.0	247203000.0	11.4483	

Information on a job collection

auto_avg_CPU_idle TOO HIGH (CPU idle > 25%)

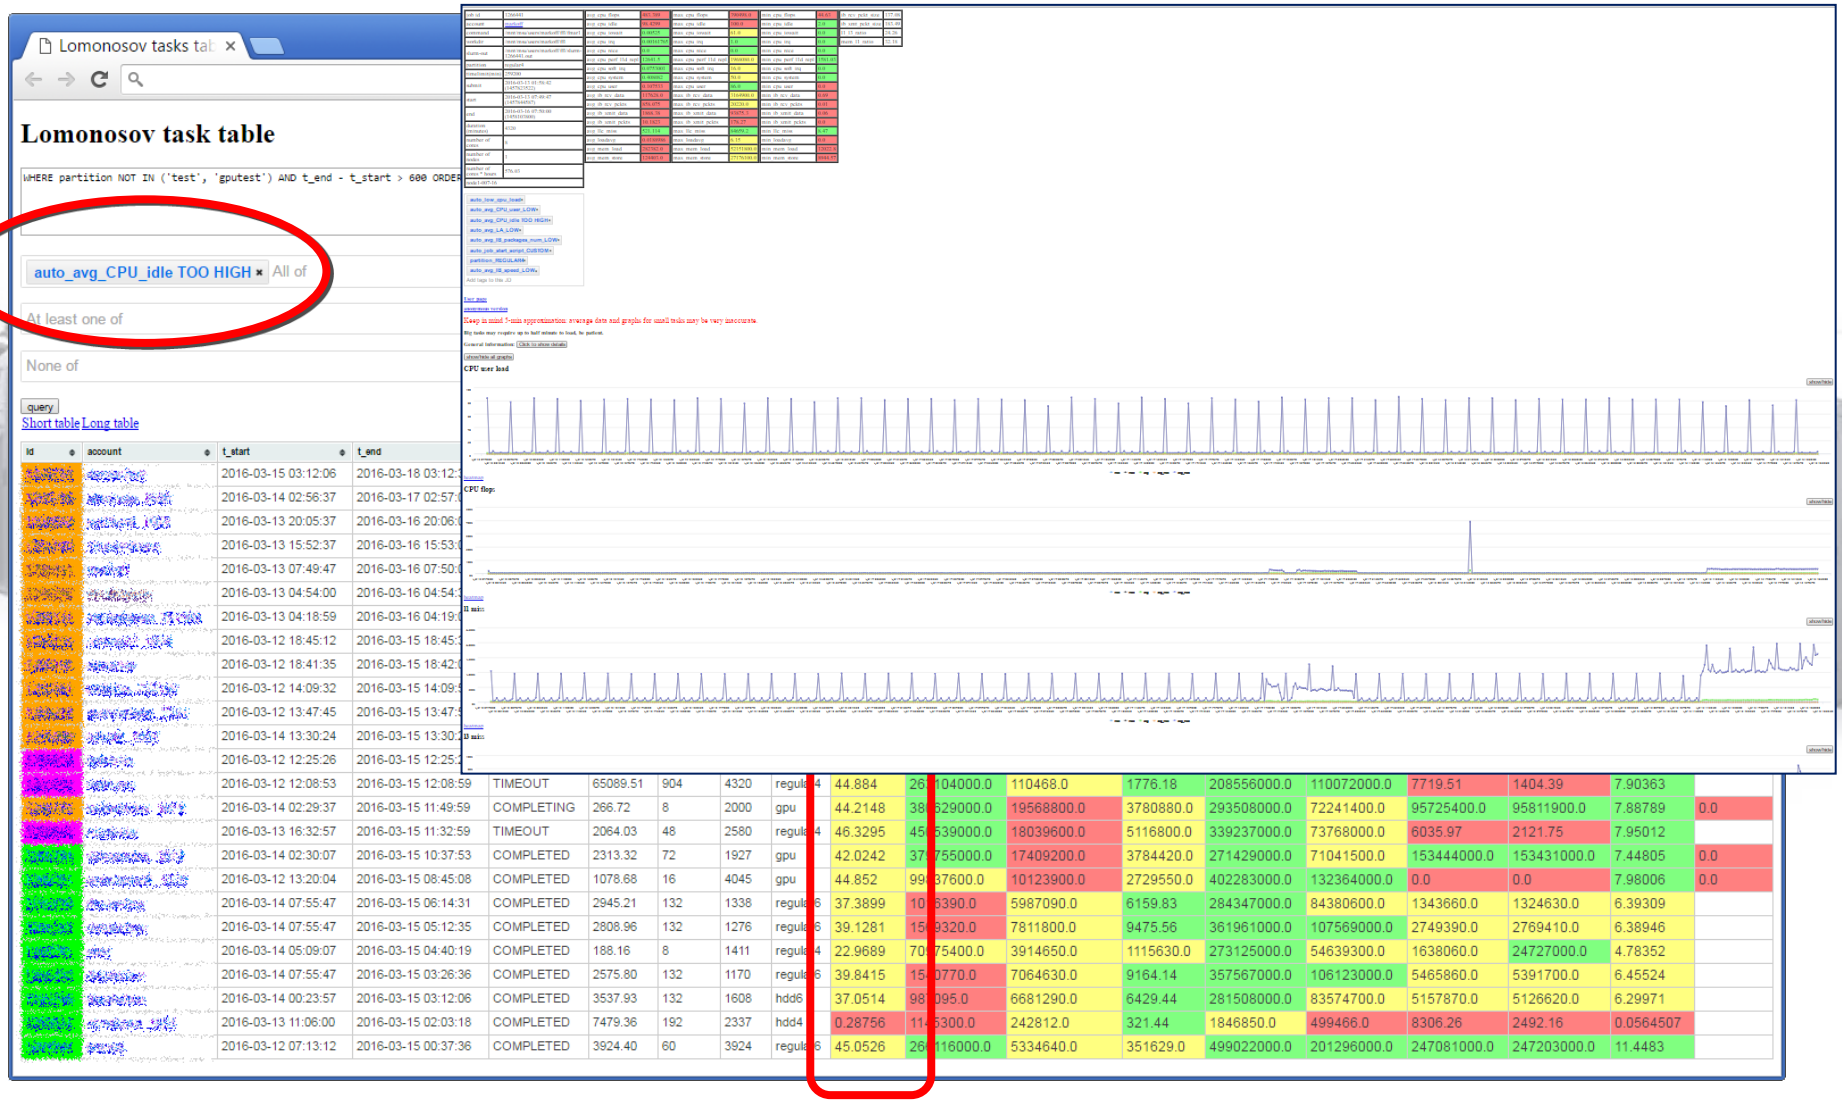

JobDigest

Analysis of job behavior and resource utilization dynamics

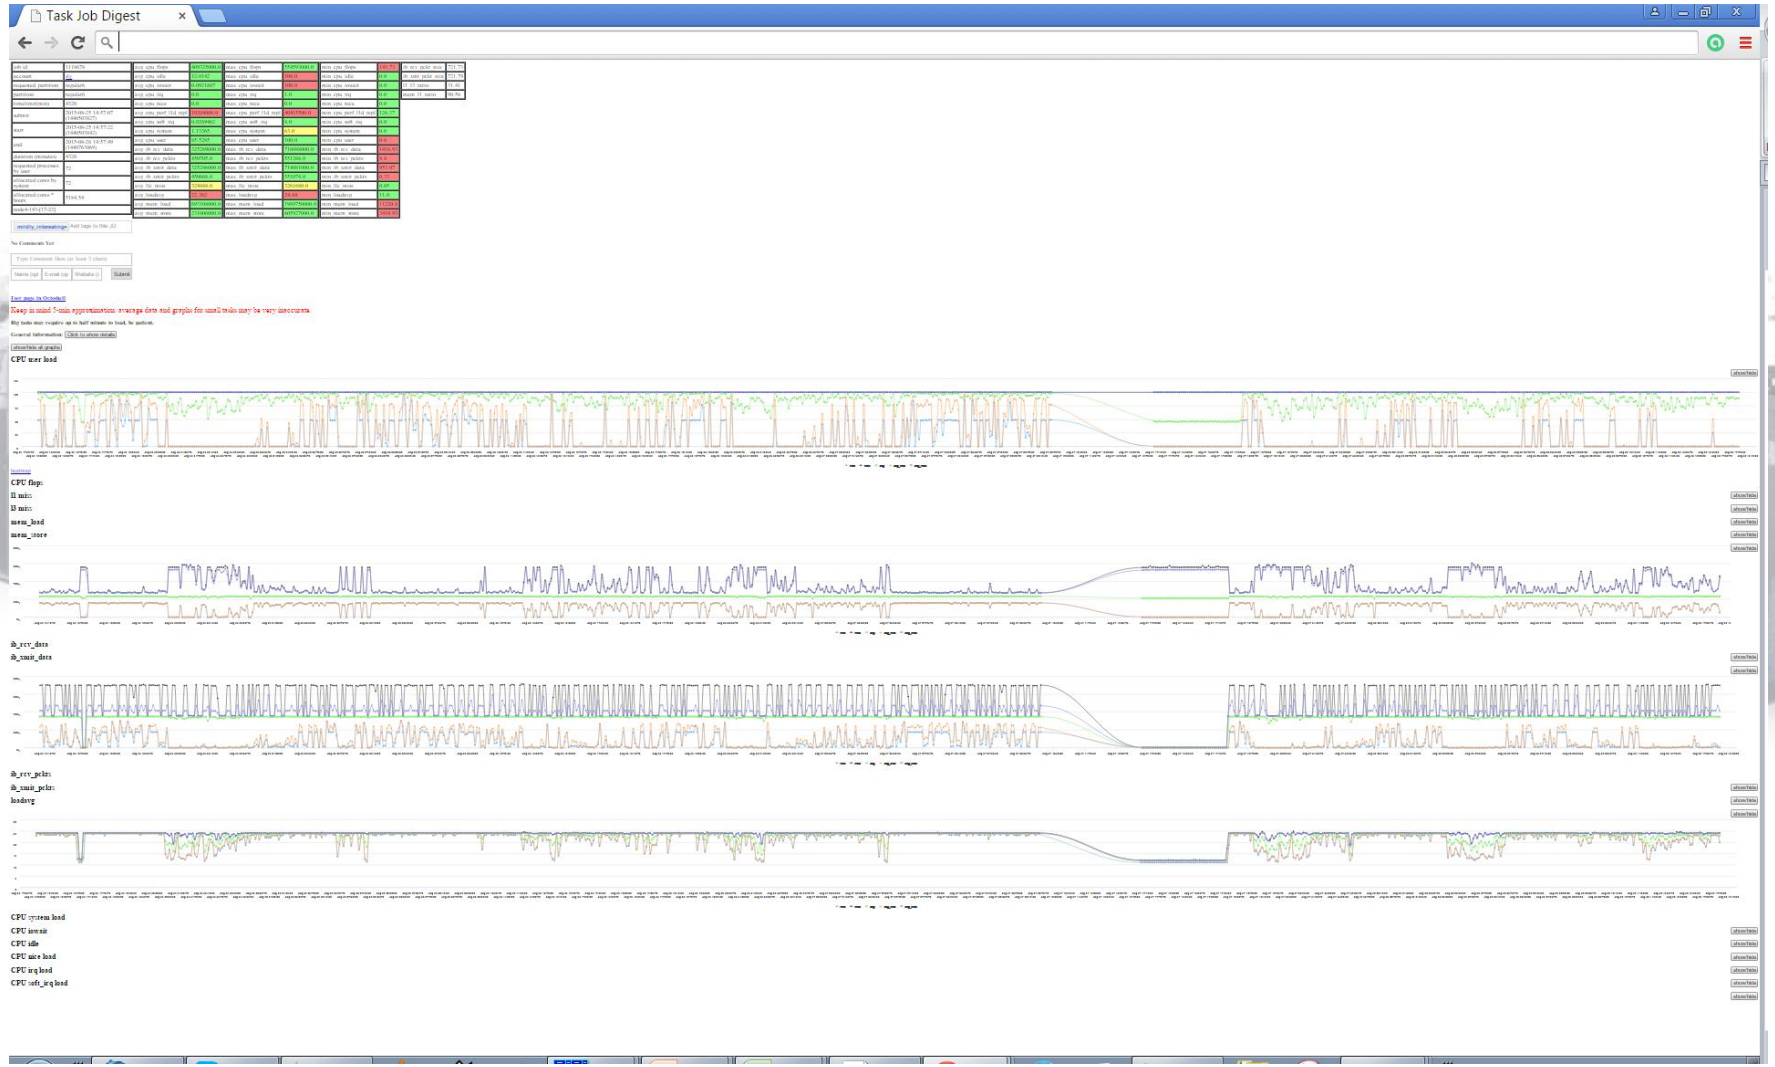

The JobDigest Report

JobDigest report blocks:

A – General job information,

B – Dynamic job characteristics,

C – Integral job characteristics,

D – Tags and job categories.

JobDigest

Dynamics of resource utilization execution

job id	1314965
account	
command	ffxp
workdir	unknown
slurm-out	unknown/slurm-1314965.out
partition	gpu
timelimit(min)	4320
submit	2016-07-09 01:18:40 (1468016320)
start	2016-07-12 22:45:07 (1468352707)
end	2016-07-15 22:46:10 (1468611970)
duration (minutes)	4321
number of cores	168
number of nodes	21
number of cores * hours	12098.94
node6-153-[02-03],node6-154-[06-07],node6-161-[03-05],node6-162-11,node6-164-16,node6-165-[01-03,15],node6-170-[07-12],node6-172-[12-13]	

auto_low_gpu_load × auto_avg_CPU_idle TOO HIGH ×
auto_avg_IB_packages_num_LOW ×
auto_job_start_script_CUSTOM × partition_GPU ×
status_TIMEOUT × auto_avg_IB_speed_LOW ×
Add tags to this JD

cpu_flops	240582000.0	cpu_flops	2086480000.0	cpu_flops	36.07
cpu_idle	47.5283	cpu_idle	100.0	cpu_idle	0.0
cpu_iowait	0.0178268	cpu_iowait	86.0	cpu_iowait	0.0
cpu_irq	0.000436416	cpu_irq	1.0	cpu_irq	0.0
cpu_nice	0.0	cpu_nice	0.0	cpu_nice	0.0
cpu_perf_l1d_repl	2661700.0	cpu_perf_l1d_repl	28549900.0	cpu_perf_l1d_repl	35.91
cpu_soft_irq	0.32768	cpu_soft_irq	6.0	cpu_soft_irq	0.0
cpu_system	1.21602	cpu_system	100.0	cpu_system	0.0
cpu_user	28.4004	cpu_user	100.0	cpu_user	0.0
gpu_load	0.0	gpu_load	0.0	gpu_load	0.0
gpu_mem_load	0.0	gpu_mem_load	0.0	gpu_mem_load	0.0
gpu_mem_usage	10658800.0	gpu_mem_usage	10858500.0	gpu_mem_usage	10334200.0
ib_rcv_data	4700.33	ib_rcv_data	16399600.0	ib_rcv_data	0.0
ib_rcv_pckts	15.3161	ib_rcv_pckts	608.17	ib_rcv_pckts	0.0
ib_xmit_data	1879.58	ib_xmit_data	129574.0	ib_xmit_data	0.0
ib_xmit_pckts	8.23421	ib_xmit_pckts	600.87	ib_xmit_pckts	0.0
llc_miss	102829.0	llc_miss	1668000.0	llc_miss	4.69
loadavg	4.8232	loadavg	14.88	loadavg	0.0
mem_load	377220000.0	mem_load	1639010000.0	mem_load	8287.4
mem_store	149185000.0	mem_store	634774000.0	mem_store	2581.3

Integral job characteristics

Resource manager data

JobDigest

Dynamics of resource utilization execution

job id	1314965
account	
command	
workdir	
slurm-out	
partition	
timelimit(min)	
submit	
start	
end	
duration (minutes)	
number of cores	
number of nodes	
number of cores * ho	
node6-153-[02-03],n [03-05],node6-162-1 [07-12],node6-172-[1	

cpu_flops 240582000.0 cpu_flops 2086480000.0 cpu_flops 36.07

auto_low_gpu_

auto_avg_IB_p

auto_job_start_

status_TIMEOU

Add tags to this J

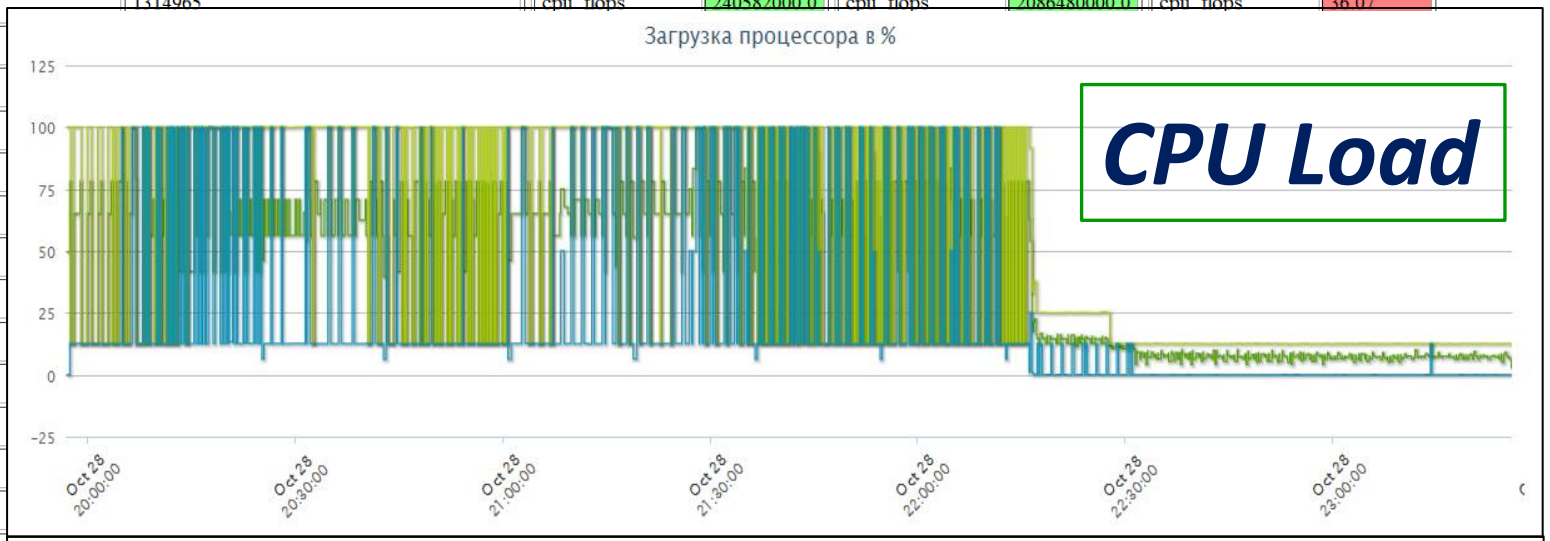

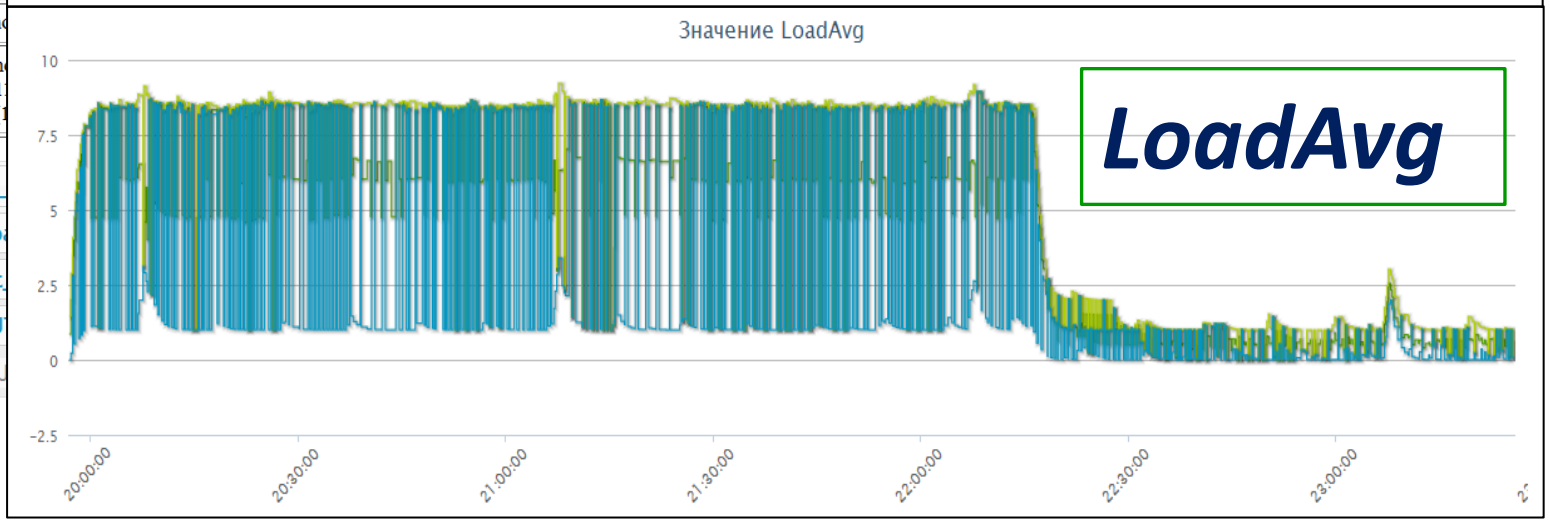

JobDigest

Dynamics of resource utilization execution

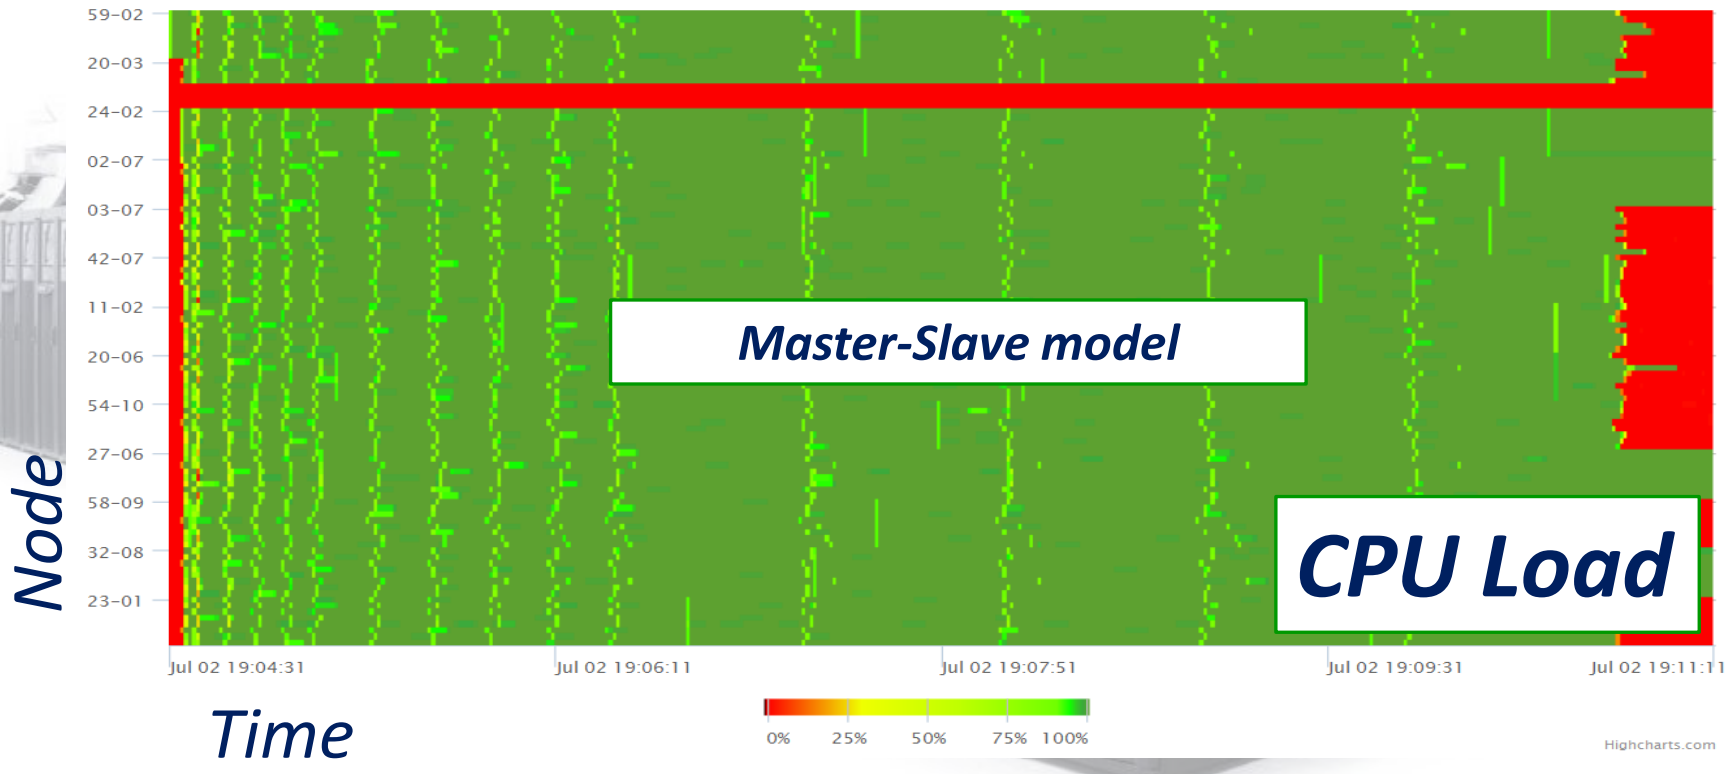

JobDigest

Dynamics of resource utilization execution

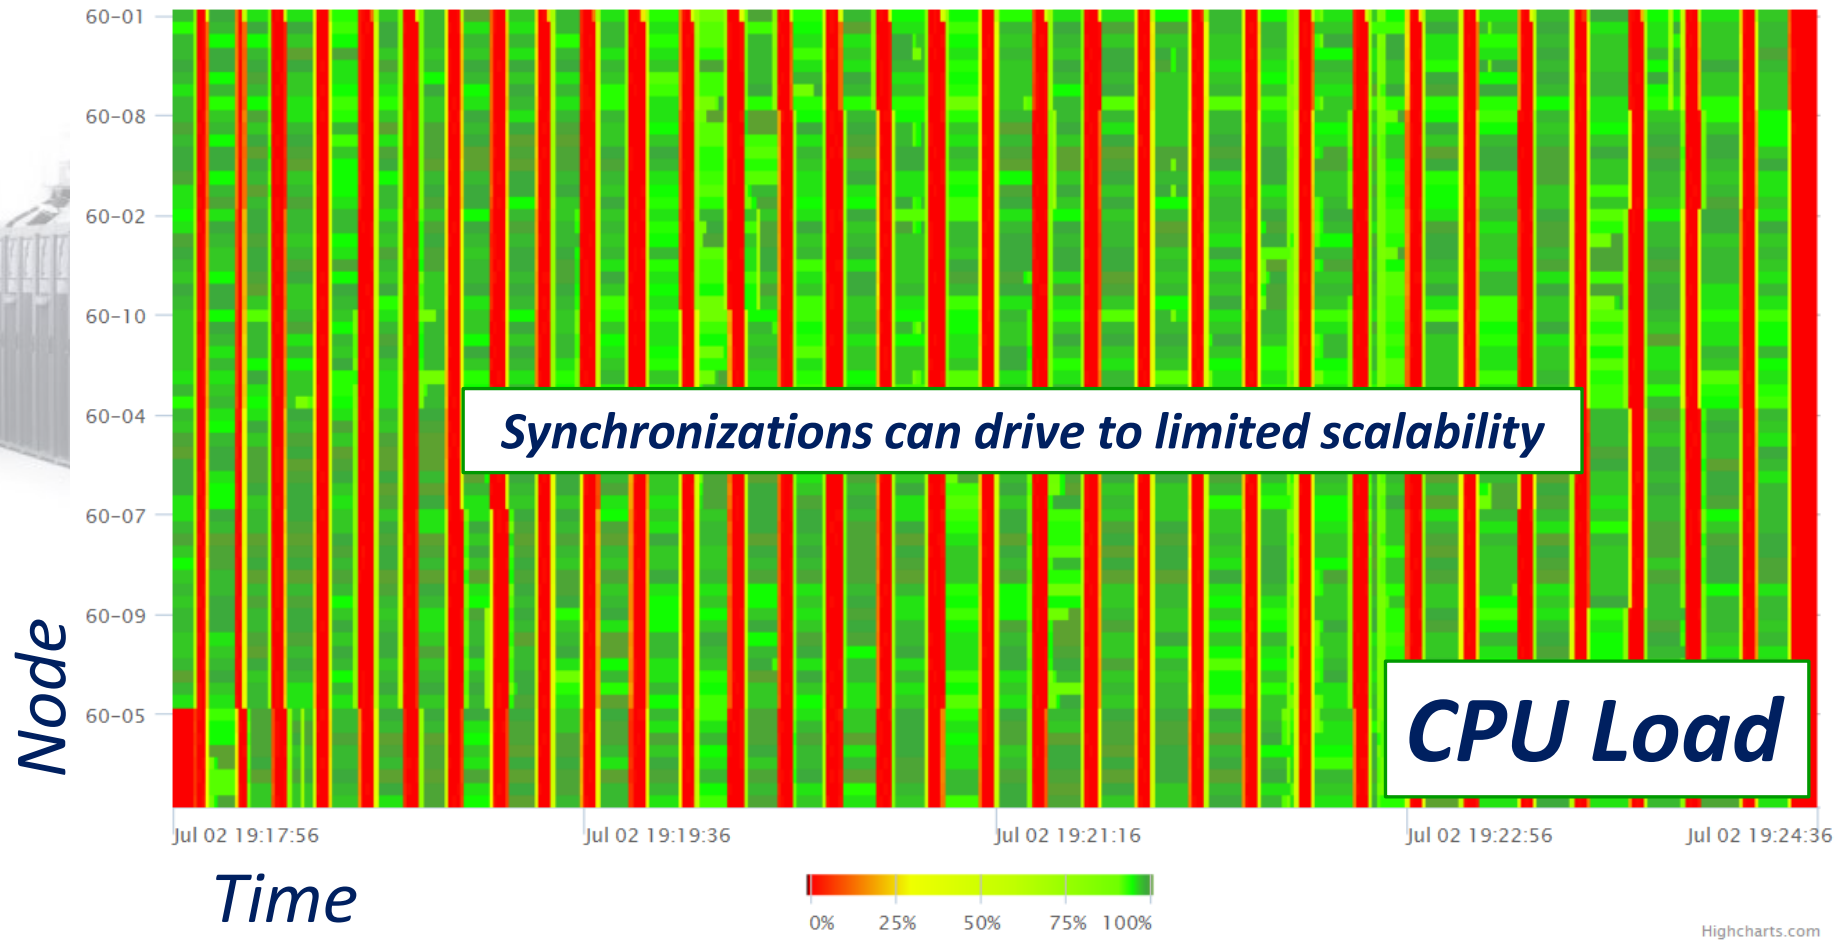

JobDigest

Dynamics of resource utilization execution

MSU/Uppsala application analysis

Joint project on apps analysis under STINT support

- Three different applications from Uppsala:
  1) “Chunks and Tasks” programming model.
     Evaluation of previously developed code at Uppsala.
  2) Parallel block preconditioned inner-outer solvers with application to Glacial Isostatic Adjustment (GIA).
     Extensive usage of third party libraries.
  3) The Fast Multipole Method (FMM).
     Self-written code at Uppsala.

- Methodology
  - Providing MSU’s Lomonosov-2 system as a testbed
  - Scalability and overall resource utilization analysis with the tools developed by MSU based on system monitoring
  - Further diagnostics with other tools (Scalasca, Valgrind, mpiP)

MSU/Uppsala application analysis

Parallel block preconditioned inner-outer solvers with application to Glacial Isostatic Adjustment (GIA)

1) JobDigest demonstrated low memory locality on the last phase of program execution – the solver
2) The solver demonstrates worse scalability among all application phases
3) Profiling of a single thread execution with Valgrind shows, that most time is spent in library call to epetra_dcrsmv_, which is internal function for sparse matrix - vector multiplication routine from Trilinos
4) The communication pattern shows that most interactions are seen between neighbor processes

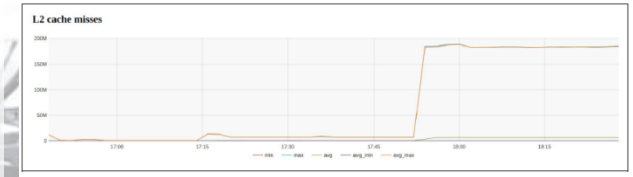

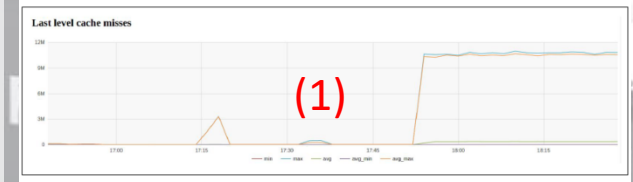

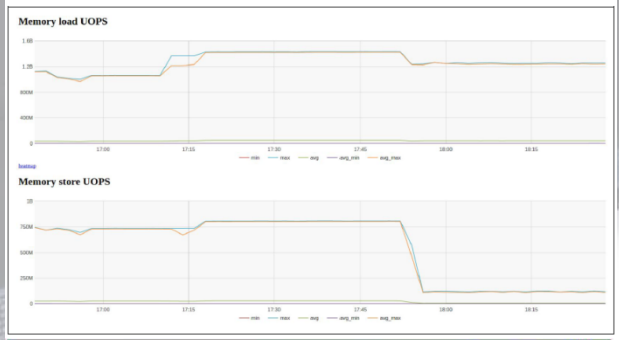

Processes	Assembling	Grid generator	Setup DOFs	Setup preconditioner	Solve	Total
1	479.00	64.80	253.00	27.40	962.00	1,790.00
2	245.00	36.50	285.00	16.20	548.00	1,130.00
4	124.00	19.10	224.00	8.53	302.00	678.00
8	61.80	9.93	103.00	4.74	208.00	387.00
16	31.10	5.39	51.60	2.87	105.00	197.00
32	16.20	3.93	27.20	2.44	77.40	128.00
64	9.34	2.57	16.10	2.81	57.70	89.70
128	4.83	2.03	12.70	3.12	62.90	86.50
256	3.27	1.95	11.30	7.27	96.90	123.00

(2)

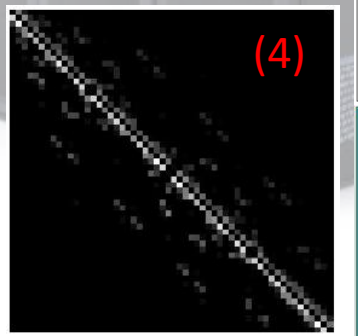

Transmitted data

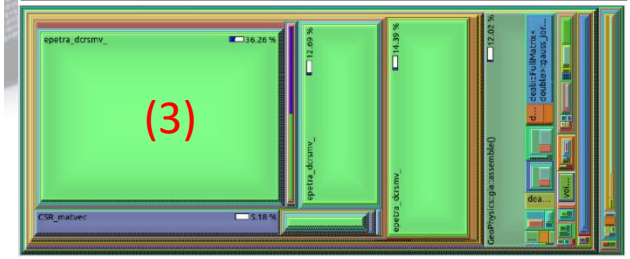

Taking care of resource utilization and productivity, take care of your users!


Russian Supercomputing Days 2017
International Supercomputing Conference
Moscow, Russia
Thank you!
Russian Supercomputing Days
2017
26.09.2017