

Supercomputer application integral characteristics analysis for the whole queued job collection of large-scale HPC systems

D.A. Nikitenko, Vl.V. Voevodin, A.M. Teplov, S.A. Zhumatiy, Vad.V. Voevodin,
K.S. Stefanov, P.A. Shvets

Research Computing Center M.V. Lomonosov Moscow State University

Efficient use and high output of any supercomputer depends on a great number of factors. The problem of controlling granted resource utilization is one of those, and becomes especially noticeable in conditions of concurrent work of many user projects. It is important to provide users with detailed information on peculiarities of their executed jobs. At the same time it is important to provide project managers with detailed information on resource utilization by project members by giving access to the detailed job analysis. Unfortunately, such information is rarely available. This gap should be eliminated with our proposed approach to supercomputer application integral characteristics analysis for the whole queued job collection of large-scale HPC systems based on system monitoring data management and study, building integral job characteristics, revealing job categories and single job run peculiarities.

Keywords: supercomputer, efficiency, system monitoring, job categories, integral job characteristics, queued job collection, job queue, resource utilization control.

1. Introduction

Securing efficient resource utilization of HPC systems is one of the most important and challenging tasks at present trends of rapid growth of scales and capabilities of modern supercomputers [1, 2]. There is a variety of approaches that are aimed at analysis of efficient utilization of certain HPC system components or systems as a whole. Some of them are based on system monitoring data analysis [3,4]. This type of approaches sets especially strict requirements on monitoring system implementation and configuration [5], as well as for the means of data storage and access. At the same time these approaches possess a number of fundamental advantages.

First, the analyzed data reflects physical, real levels of HPC system components and appropriate resource utilization.

Second, filtering system monitoring data obtained from known set of components and period of time allows binding this data to certain jobs. Thus, allowing analyzing resource utilization history and trends by certain applications, users, projects, partitions, and so on.

Third, typically it is possible to configure monitoring systems obtaining data from the whole system in such a way that it induces acceptable overhead. This allows collecting data with a rougher granulation, when possible, but still sufficient for basic analysis of resource utilization by any and every job. To have more detailed information on certain job, of course, used monitoring system should likely support data acquisition rate reconfiguration on-the-fly for specified sensor sets and sources. If not (of course, it is much less efficient way), most monitoring systems can be started in a higher granularity mode to record certain job activity and restarted in a normal mode afterwards. There are other options available, for example, data aggregation implementation that is precise for first, say, 30 seconds of job execution (to study short jobs) that is later switches to rough mode (for longer jobs). Anyhow, there are a number of techniques available to study certain application behavior.

The existing methods and techniques that base on system monitoring data analysis, allow both analysis of dynamic characteristics of certain application runs and peculiarities of resource utilization within system partitions and systems as a whole [6]. With a project-oriented workflow, when a number of users run jobs as a part of one applied research, it is very useful to let administrator and system manager have a clear view of resource utilization distribution in the workgroup to have a possibility to influence permissions or workflow inside the workgroup to meet the granted resources limitation [7].

Nevertheless there is still need for specialized tools and techniques to analyze available system monitoring data. In a point of fact, the one is needed as a valuable additional tool to the set of implemented approaches in every-day practice of MSU Supercomputer Center [8-11] — a tool for job queue analysis based on system monitoring that would allow revealing job categories, job grouping by some criteria, starting from belonging to user or project domain and other resource manager specific characteristics, to categories by levels and peculiarities of HPC system resource utilization or its combinations.

As a basic technique for such grouping implementing tagging system seems to be an adequate option — assigning special tags to a job description as soon as each tag description criteria is met by job characteristics. Tagging principles are widely successfully used for categorizing and search purposes managing huge collections of data in Internet: news, videos, photos, notes, and so forth, that is quite close to the challenge that is being tackled.

2. Job categories and tagging

The combined analysis of system monitoring data and resource manager log data, as was already mentioned, allows binding raw system monitoring data to certain jobs. This provides means to analyze job dynamics as far as data granularity allows. To analyze the average rate of application resource utilization every dynamic characteristic can serve basis for the calculation of minimum, maximum, average and median values. These types of values are often named integral job characteristics.

When one takes a look at the whole scope of executed jobs for analysis of job queue structure, application run sequences, jobs comparative analysis and even searching for outstanding single job behavior it becomes obvious that it would be very useful to have tools that provide means for revealing job categories based on various criteria.

This functionality can be implemented by introducing special tags. Every tag is based on its own criteria, based on a single integral job characteristic or its combination, resource manager job-related information and any other available info from used data sources. For example tags can correspond to certain average rates of various resource utilization, job ownership, job duration, resource utilization specifics, special execution modes, detailed system monitoring data availability, and so forth.

The approach features means to make efficient grouping and filtration of whole job queue history collection by any improvised combination of specified tags.

Driven by experience of application efficiency and scalability study based on system monitoring data analysis, the authors propose introducing the following job categories on the first stage of implementation. Tag naming is designed to give self-explanatory tag description, nevertheless, every tag must have a detailed full-format description available.

2.1 System monitoring data based categories

CPU utilization

- Tag name: ***avg_CPU_user LOW***
Category: Low CPU user utilization.
Criteria: Average value of CPU_user doesn't exceed 20%.
- Tag name: ***avg_CPU_user HIGH***
Category: High CPU user utilization.
Criteria: Average value of CPU_user exceeds 40%.
- Tag name: ***avg_CPU_idle TOO HIGH***
Category: CPU is idle for a considerable time.
Criteria: Average CPU_idle value exceeds 25%.

Competition of processes for CPU cores

- Tag name: ***avg_LA LOW***
Category: User job is almost out of action, almost no utilization of CPU.
Criteria: Average Load Average is below 1.

- Tag name: ***avg_LA SINGLE CORE***
Category: Only one process per node is active as an average.
Criteria: Average Load Average is approximately 1.
- Tag name: ***avg_LA NORMAL***
Category: Optimal competition of processes.
Criteria: Average Load Average is approximately equal to the number of cores per node.
- Tag name: ***avg_LA HYPERTHREADED***
Category: Normal process competition with hyperthreading is on.
Criteria: Average Load Average value is approximately equal to the double number of CPU cores per node.

Floating point operations

- Tag name: ***avg_Flops HIGH***
Category: Intensive CPU floating point operations.
Criteria: Average value of floating point operations number exceeds 10% of theoretical CPU peak.

Interconnect activity

- Tag name: ***avg_IB_packages_num LOW***
Category: Low number of inter-node data transmissions.
Criteria: Average package send rate does not exceed 10^3 packages per second.
- Tag name: ***avg_IB_packages_size TOO LOW***
Category: Small size of packages.
Criteria: Average package send rate exceeds 10^3 packages per second while average data transmission rate is below 2 kilobytes per second.
- Tag name: ***avg_IB_speed HIGH***
Category: High data transmission intensity.
Criteria: Average data transmission rate is over 0,2 Gigabytes per second and up to 1 Gigabytes per second.
- Tag name: ***avg_IB_speed TOO HIGH***
Category: Very high data transmission intensity.
Criteria: Average data transmission rate is over 1 Gigabytes per second.

Memory utilization

- Tag name: ***avg_cache_L1/L3 TOO LOW***
Category: Very low efficiency of cache stack utilization.
Criteria: Ratio of the number of L1 misses to the number of L3 misses is below 5.
- Tag name: ***avg_cache_L1/L3 LOW***
Category: Reduced efficiency of cache stack utilization.
Criteria: Ratio of the number of L1 misses to the number of L3 misses is below 10.
- Tag name: ***avg_cache_L1/L3 HIGH***
Category: Good efficiency of cache stack utilization.
Criteria: Ratio of the number of L1 misses to the number of L3 misses exceeds 10.
- Tag name: ***avg_mem/cache_L1 LOW***
Category: Reduced efficiency of cache L1 utilization.
Criteria: Ratio of the number of total memory operations to the number of L1 misses does not exceed 15.
- Tag name: ***avg_memload HIGH***
Category: Intensive memory operations.
Criteria: Average number of memory operations exceeds 10^9 operations per second.

2.2 Resource manager based categories

Job execution status

- Tag name: ***job_status COMPLETED***
Category: Job is successfully finished.
- Tag name: ***job_status FAILED***
Category: Job is finished with an error in program.
- Tag name: ***job_status CANCELED***
Category: Job was cancelled by user.
- Tag name: ***job_status TIMEOUT***
Category: Job was cancelled by exceeding time limit.
- Tag name: ***job_status NODE_FAIL***
Category: Job is finished with system error.

Job submission details

- Tag name: ***job_time_limit CUSTOM***
Category: Requested time limit is custom.
- Tag name: ***job_start_script CUSTOM***
Category: Job batch file is custom.
- Tag name: ***job_cores_requested FEW***
Category: Not all available CPU cores per node requested.
- Tag name: ***job_cores_requested SINGLE***
Category: Just a single CPU core per node requested.
- Tag name: ***job_MPI INTEL***
Category: MPI type used: Intel MPI.
- Tag name: ***job_MPI OpenMPI***
Category: MPI type used: OpenMPI.
- Tag name: ***job_nnodes SINGLE***
Category: Job used a single node.
- Tag name: ***job_nnodes FEW***
Category: Job used from 2 up to 8 nodes.
- Tag name: ***job_nnodes MANY***
Category: Job used 8 nodes and above.

System-dependent peculiarities and partition usage

Illustrated by the example of “Lomonosov” supercomputer partitions.

- Tag name: ***job_partition REGULAR4***
Category: Job allocated to REGULAR4 partition.
- Tag name: ***job_partition REGULAR6***
Category: Job allocated to REGULAR6 partition.
- Tag name: ***job_partition HDD4***
Category: Job allocated to HDD4 partition.
- Tag name: ***job_partition HDD6***
Category: Job allocated to HDD6 partition.

- Tag name: ***job_partition SMP***
Category: Job allocated to SMP partition.
- Tag name: ***job_partition GPU***
Category: Job allocated to GPU partition.
- Tag name: ***job_partition TEST***
Category: Job allocated to TEST partition.
- Tag name: ***job_partition GPUTEST***
Category: Job allocated to GPUTEST partition.
- Tag name: ***job_partition EXCEPT TEST***
Category: Job allocated to regular or high priority partition.
- Tag name: ***job_priority HIGH***
Category: Job allocated to partitions with a higher priority (в очередях reg4prio, gru_p, dedicated6)

Matching partition specifics

- Tag name: ***job_accell GPU***
Category: User application uses accelerators. Accelerator type: GPU.
- Tag name: ***job_accel GPU UNUSED***
Category: Job is run on GPU partition, but never uses GPUs.
- Tag name: ***job_disks UNUSED***
Category: Job is run on HDD-equipped partition, but never uses local I/O.
- Tag name: ***job_disks TOO LOW***
Category: Job is run on HDD-equipped partition, but I/O rate is very low.

2.3 Other categories

Beyond the tags that can be assigned automatically, it is possible to introduce manually-set tags. This is useful when the criteria is cannot be automatically determined. There are now different manual setting options available. First, most typical, selecting the one from known tags lost or introducing new one with human-read and formal description. Second, is pushing some tags like “higher system monitoring rate for the job” via the command line when submitting a job.

This applies for instance to general job description characterizing type of data processing as it is usually known a priori or determined in the course of job behavior study by a specialist:

job_behavior DATA MINING,
job_behavior MASTER-SLAVE,
job_behavior COMMUNICATION,
job_behavior ITERATIVE, etc.

In the same manner typical anomalies encountered during analysis course can be specified:

job_bug DEADLOCK,
job_bug DATA RACE, etc.

It is very useful to specify if a widely used algorithm implementation or software package is used. Just in case, this can provide a great contribution to scalability and algorithms-studying projects, like AlgoWiki [12]:

job_sw VASP,
job_sw FIREFLY,
job_sw GROMACS, etc.

If detailed reports on job efficiency analysis or issues is available, or specific standard report like JobDigest is available, it is useful to mark such a feature with another tag, for example:.

job_analized,
job_analized JobDigest.

3. Implementation details

We keep to the basis of building a tool that might be deployed at any supercomputer center with minimal efforts. We currently support Slurm[13], Cleo[14] resource managers and Ganglia [15], Collectd[16], Clustrx[17], DiMMon [5] (most promising) monitoring systems.

As for integral job characteristics derivation and tagging, PostgreSQL is used as data storage for coarsened system monitoring data and saved job information from resource managers. The saved job info is processed by JavaScript, jQuery with jQuery UI [18] and TagIt [19].

The tag can be assigned to a job only if it is already declared in tag description table. Such a table includes tag id, name, human-readable description, criteria (a specification of SQL request for automatic processing), comments, and a flag of availability that can be set only by administrator. Any user can suggest introducing a new tag, but it will be available only after administrator approval. Information on new tag author is saved in the comments attribute, added the user tag description suggestion and motivation.

All tags can be assigned in two ways: automatically and manually. Any tag set by mistake or error can be manually removed from a job.

Automatic mode

In this mode, the tags are automatically assigned:

- to all finished jobs according to SQL-based criteria regarding saved integral job characteristics data, information from resource manager and other available saved data;
- as a result of running a special script that processes whole saved job collection info.

In this mode a special attribute would indicate that the tag was set in package (automatic) mode of tag assignment.

Manual mode

Manual tag assignment is usually done by user, project manager or administrator in the following cases:

- as a result of certain job analysis (specifying algorithm implemented, etc.);
- as a result of specifying the tag via command line when submitting a job;
- any tag in a user-specific tag space (marking out important job runs as a part of the project, etc.).

In this mode an attribute addressing tag author is set, that also allows finding jobs, marked as a part of a certain project or by a certain user.

User-specific tag space consists of regular tags and custom user tags. Any manually assigned tags by a user are seen only in the scope of the project and system administrators. The members of other project see their own tag spaces and the general tag space is available for all of the users.

4. Use cases

Of course, real life use cases are very diverse. In this section we would like to share our experience of every-day usage of the proposed technique as a part of the developed tool approbation at Supercomputer Center of Moscow State University on a few examples just to give a general idea of it.

4.1 Revealing jobs, users and projects that practice inappropriate resource utilization

One of the problems of every-day practice of large-scale supercomputer center with a number of heterogeneous resources and considerable number of users concurring for the resources is a problem of unacceptable efficiency or inappropriate resource utilization. This is of a higher priority for specific limited resources, like compute nodes equipped with specialized accelerators, local disks, extra memory or other hardware and software that is critical for some applications and at the same time these nodes can still be used by applications that do not need that specific type of resources that the

nodes possess. Such nodes usually have a high potential for resource-demanding specific applications and for the large systems like “Lomonosov” are usually managed as a separate partition with a special queuing options to allow submitting jobs to the appropriate partition. This is vital for projects that perform computations only due to the advantages of such partitions, so by queuing to the desired partition user get a guarantee that their application would have all necessary resources at disposal.

Nevertheless, when analyzing the whole job collection for such partitions it appears that there are numerous job runs that do not use any partition facilities benefits. Of course, sometimes algorithm peculiarities can use resources with totally different intensity, but further analysis usually shows that the majority of suspicious jobs never use any benefit of such partitions. The reasons can be different, but usually it is a shorter wait time in a queue.

This can be seen on GPU partitions with user job runs that never use GPUs. A slightly different situation is seen on HDD-equipped nodes with absent or extreme low disk usage rate and finally, single-process application that don't benefit from multiple CPU cores can be seen almost on any partition regardless of hardware and software.

It is important to find the root cause of such applications behavior and as soon as the reason is found and changes by user or administrator are applied, the ratio of such jobs can be lowered that would immediately raise HPC system efficiency and overall throughput.

The most popular reasons are:

- Problems inside the application, program or algorithm. The user is sure that he needs resources, but in practice application doesn't utilize any or utilizes at extremely low rates.
- Problems of HPC system. The declared resources are not available on the nodes.
- Inappropriate job allocation. This can be both a mistake, and cheating for lower job waiting time.

Regardless of real reason, these job runs lead to a higher wait time for the jobs that really need specific resources.

The search for such jobs can be automated using integral job characteristics and some of introduced tags.

For example, to filter the jobs allocated to GPU partition with no usage of accelerator one can use tags ***job_partition GPU*** and ***job_accel GPU UNUSED*** at the same time. Next, one can cut off jobs allocated to the test partition as of no interest. The rest jobs that are assigned ***job_status COMPLETED*** tag probably do not need GPUs at all, as finished successfully with no registered GPU usage. At this point two options are available whether it is a mistake (user or system) or it was done by user intentionally, trying to reduce job wait time as wait time in specific partitions is sometimes less than in regular.

A very similar situation is seen for HDD-equipped nodes. Jobs that are tagged with ***job_partition HDD4***, ***job_partition HDD6***, ***job_disks UNUSED*** or ***job_disks TOO LOW*** can potentially be successfully executed at regular partitions. Note that there appears an option of very low resource utilization. This means that disk operations might be easily replaced with network file system operations with minimal additional overhead or even without it.

For those jobs that are tagged with ***job_nodes SINGLE*** and ***avg_LA SINGLE CORE*** or ***avg_LA LOW***, it is quite reasonable to inquire what for it was submitted to the supercomputer. Such jobs use a single node and a single core (or just few processes per node) and can potentially run well on a desktop. Unfortunately such jobs are met very often.

Users who submit types of jobs mentioned above must be contacted to figure out the reasons of the revealed facts of inappropriate and inefficient resource utilization. The problems found should be resolved. If cheating is met or it is proved that the executed jobs do not really need HPC resources, quotas for corresponding user accounts and projects can be reduced to the extent of blocking.

Let us take a look at one of real-life examples. Figure 1 illustrates the filtered job list allocated to regular partitions with automatically ***avg_LA_SINGLE_CORE*** tag assigned. It is clearly seen that the jobs have a low LoadAverage close to 1 as filtered by the tag, at the same time having very low CPU_user. Note, that it is not a test partition and all jobs are run on a single node, grabbing 8 cores on regular4 and hdd4 partitions!

id	account	T_start	T_end	state	cores_host	mem_cores	duration	partition	avg_cpu_user	avg_cpu_flops	avg_cpu_perf_fld_rps	avg_ic_miss	avg_mem_load	avg_mem_store	avg_rb_rcv_data	avg_rb_xmt_data	avg_loadavg
1230001	user@host	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
1230002	user@host	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
1230003	user@host	2016-02-03 03:59:28	2016-02-03 04:24:47	COMPLETED	27.36	8	205	hdd4	0.092307	1498.76	30833.7	1890.64	504213.0	245048.0	80744.6	79506.6	0.994103
1230004	user@host	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
1230005	user@host	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
1230006	user@host	2016-02-02 16:58:39	2016-02-02 17:33:38	FAILED	4.66	8	34	regular4	6.12	9850980.0	56315.4	20220.9	36315500.0	22835900.0	0.0	0.0	0.99
1230007	user@host	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
1230008	user@host	2016-01-31 02:07:46	2016-01-31 05:33:30	COMPLETED	27.43	8	2487	hdd5	0.048974	25378.4	177844.0	46266.7	19580800.0	12530700.0	54673600.0	844684.0	0.900515
1230009	user@host	2016-01-24 08:28:31	2016-01-27 08:28:33	COMPLETING	576.00	8	4320	regular4	10.6563	36771700.0	323722.0	876.039	25520100.0	5824880.0	2890.48	2911.17	0.98058
1230010	user@host	2016-01-23 09:14:31	2016-01-25 05:30:34	COMPLETED	354.14	8	2656	regular4	6.25704	29631700.0	1163360.0	110499.0	29963900.0	12936200.0	5563.94	5626180.0	1.00758
1230011	user@host	2016-01-23 03:22:49	2016-01-24 19:46:44	COMPLETING	315.46	8	2365	regular4	6.08654	311.465	6483.29	409.6	114895000.0	25839200.0	1704.98	1907.76	1.01085
1230012	user@host	2016-01-24 13:02:31	2016-01-24 13:28:52	FAILED	3.51	8	26	hdd4	0.253333	110354.0	2007430.0	539.793	51296300.0	32041400.0	3184.72	2307.22	0.98666
1230013	user@host	2016-01-23 06:34:42	2016-01-23 11:09:05	COMPLETING	36.58	8	274	regular4	0.042452	310.48	8636.37	420.936	39045500.0	6020430.0	1760.95	1911.44	1.00528
1230014	user@host	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
1230015	user@host	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	30568800.0	46279200.0	133681.0	1.01707
1230016	user@host	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
1230017	user@host	2016-01-23 00:07:12	2016-01-23 00:49:11	COMPLETED	5.60	8	41	hdd4	0.823333	49520.7	2304560.0	4253.29	42744700.0	25238900.0	23307.5	2071330.0	1.03
1230018	user@host	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	50821900.0	31777900.0	54200400.0	107747.0	1.00778
1230019	user@host	2016-01-21 12:44:15	2016-01-21 15:25:49	FAILED	21.54	8	161	regular4	0.462903	377364.0	1954350.0	1676.62	49962900.0	31221100.0	3537700.0	32939.6	1.01516

Fig 1. Filtered single-process jobs found in real job queue in various regular partitions

A close look at the longest job owner that was cancelled by timeout illustrates that the user always runs such single-node, even single-process jobs regardless of partitions (Figure 2).

id	account	T_start	T_end	state	cores_host	mem_cores	duration	partition	avg_cpu_user	avg_cpu_flops	avg_cpu_perf_fld_rps	avg_ic_miss	avg_mem_load	avg_mem_store	avg_rb_rcv_data	avg_rb_xmt_data	avg_loadavg
1230020	user@host	2016-02-01 11:38:57	2016-02-04 11:40:05	TIMEOUT	894.23	12	4321	hdd5	4.15912	18580900.0	617467.0	48928.4	19568700.0	8900490.0	9570.56	5600530.0	1.01378
1230021	user@host	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
1230022	user@host	2016-01-29 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
1230023	user@host	2016-01-28 05:26:13	2016-01-31 05:26:26	TIMEOUT	864.04	12	4320	regular4	8.77655	35173200.0	1282790.0	111023.0	36844200.0	16344900.0	9099.59	5280840.0	0.929849
1230024	user@host	2016-01-28 05:26:13	2016-01-31 05:26:26	TIMEOUT	864.04	12	4320	regular4	4.186	17402500.0	642886.0	57124.4	18246300.0	8096580.0	8317.03	5365910.0	0.943202
1230025	user@host	2016-01-28 03:02:12	2016-01-31 03:02:18	TIMEOUT	864.02	12	4320	hdd5	3.8758	16885000.0	512477.0	39170.0	17881700.0	8301980.0	8482.06	4987560.0	0.951705
1230026	user@host	2016-01-28 02:58:55	2016-01-31 02:59:18	TIMEOUT	864.08	12	4320	hdd5	3.84146	17714600.0	596533.0	44937.1	17912600.0	8211870.0	7297.53	4735980.0	0.939977
1230027	user@host	2016-01-28 05:26:13	2016-01-29 14:47:57	COMPLETED	400.35	12	2001	regular4	3.92857	14801900.0	549469.0	36831.0	16487800.0	7705300.0	8528.29	4514320.0	0.879649
1230028	user@host	2016-01-29 14:51:53	2016-01-29 21:38:14	COMPLETED	54.18	8	406	regular4	0.232025	101889.0	1674760.0	426.425	44143200.0	27594000.0	233985.0	2114.24	0.966203
1230029	user@host	2016-01-23 11:07:01	2016-01-26 11:07:01	COMPLETING	864.00	12	4320	regular4	3.12961	11653100.0	430194.0	33640.2	12861800.0	5889370.0	5031.85	3782510.0	0.88348
1230030	user@host	2016-01-23 11:07:01	2016-01-26 11:07:01	COMPLETING	864.00	12	4320	regular4	2.83212	12887500.0	464008.0	40734.5	13325100.0	5830350.0	4762.23	3930580.0	0.880048
1230031	user@host	2016-01-23 11:07:01	2016-01-26 11:07:01	COMPLETING	864.00	12	4320	regular4	3.04609	12865600.0	490842.0	44162.1	13322300.0	5808760.0	4718.03	3975610.0	1.66074
1230032	user@host	2016-01-23 09:14:31	2016-01-26 09:15:01	TIMEOUT	576.07	8	4320	regular4	4.33655	5432760.0	204065.0	23269.9	5665850.0	2800020.0	3349.31	2095440.0	0.721378
1230033	user@host	2016-01-23 01:13:21	2016-01-26 01:13:31	COMPLETING	864.03	12	4320	hdd5	3.39625	15536900.0	556017.0	46465.0	15968200.0	7040100.0	5270.36	4771740.0	0.81913
1230034	user@host	2016-01-23 09:14:31	2016-01-25 05:30:34	COMPLETED	354.14	8	2656	regular4	6.25704	29631700.0	1163360.0	110499.0	29963900.0	12936200.0	5563.94	5626180.0	1.00758
1230035	user@host	2016-01-23 09:14:31	2016-01-24 19:09:57	COMPLETED	497.56	12	2487	hdd5	4.19259	17882200.0	671697.0	52291.5	19600100.0	8950310.0	8002.2	6291390.0	1.01723
1230036	user@host	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
1230037	user@host	2016-01-14 22:18:39	2016-01-17 22:18:58	TIMEOUT	864.06	12	4320	regular4	4.63891	20412900.0	729143.0	68806.3	20612200.0	9038480.0	5362.19	5557470.0	1.01106
1230038	user@host	2016-01-14 10:33:03	2016-01-17 10:33:28	TIMEOUT	864.08	12	4320	hdd5	4.14311	20556500.0	736184.0	69529.3	20302000.0	8827940.0	5982.59	5083250.0	1.00708
1230039	user@host	2016-01-14 10:29:03	2016-01-17 10:29:28	TIMEOUT	576.06	8	4320	regular4	6.14744	28706300.0	1051800.0	105584.0	29685900.0	12895800.0	6108.46	5543890.0	1.00313
1230040	user@host	2016-01-14 10:27:03	2016-01-17 10:27:28	TIMEOUT	864.08	12	4320	regular4	4.20599	19444700.0	725574.0	64734.5	19696400.0	8532810.0	6016.08	5398490.0	1.01332
1230041	user@host	2016-01-14 10:21:03	2016-01-17 10:21:28	TIMEOUT	864.08	12	4320	hdd5	4.12558	19613900.0	723350.0	65947.4	20077000.0	8794960.0	6225.52	5569600.0	1.01158
1230042	user@host	2016-01-14 10:19:03	2016-01-17 10:19:28	TIMEOUT	576.06	8	4320	regular4	6.31773	26716300.0	903650.0	85406.0	29256300.0	13303000.0	8462.67	5527590.0	1.00814
1230043	user@host	2016-01-14 10:17:03	2016-01-15 10:17:28	TIMEOUT	576.06	8	4320	regular4	6.80788	28269100.0	1019160.0	96569.8	29877400.0	13262000.0	6019.81	4728780.0	1.00182
1230044	user@host	2016-01-14 10:35:03	2016-01-15 21:11:48	CANCELLED	415.35	12	2076	regular4	4.18375	19663500.0	697625.0	65144.6	20088600.0	8652860.0	8551.98	5572640.0	1.02867
1230045	user@host	2016-01-14 10:14:03	2016-01-14 12:55:14	CANCELLED	32.24	12	161	regular4	4.10677	19690100.0	715612.0	64964.6	19935300.0	8569990.0	5196.9	5150180.0	1.00065
1230046	user@host	2016-01-11 10:36:07	2016-01-14 09:32:41	FAILED	567.54	8	4256	regular4	4.31972	22211700.0	761182.0	76557.1	21713200.0	9393660.0	3464.24	3535160.0	0.898968
1230047	user@host	2016-01-13 00:13:04	2016-01-14 09:32:40	FAILED	399.92	12	1999	hdd5	1.44431</								

As a rule, most of jobs with high floating point intensity are tagged with ***avg_Flops HIGH***, ***avg_CPU_user HIGH*** tags and most efficient apps in terms of memory utilization are tagged ***avg_cache_L1/L3 HIGH***. Filtered by these criteria jobs are usually having good data locality, they are well-balanced and show high performance.

The deeper analysis of such jobs allows revealing optimal command line and compiler options for the variety of categories of standard applications and algorithm implementations. Once such a job is approved to be a well-optimized typical example of a SW package usage or algorithm implemented, a proper tag, corresponding to such a category can be set (like ***job_sw VASP***). This provides means for the comparative analysis of similar jobs. This can also serve as a good basis for the more detailed analysis of the whole job collection and revealing inefficient applications and users that use resources inefficiently.

4.3 Finding applications with special need for large amounts of memory

Many users of supercomputer complex run applications that are resource-demanding regarding amount of available memory per process. Such applications are usually effective enough, but are often scaled down in different ways to fit available memory, for example, reducing number of MPI processes per node and so on.

Such applications are usually run on a considerable number of nodes and LoadAvg values are below the number of CPU cores per node. These jobs are usually tagged with ***avg_memload HIGH***, and related to node and core usage tags ***avg_LA SINGLE CORE***, ***job_nnodes FEW*** or ***job_nnodes MANY***.

If such a job is found in the 6-cored CPU “Regular 6” partition (tagged with ***job_partition REGULAR6***), even changing allocation to the “Regular 4” partition can be an optimization choice leading to reducing the number of idle CPU cores per node by 4 cores ($2*6-2*4$).

Some of such applications can also benefit from moving to hybrid MPI+OpenMP or MPI+Cilk models. If such a model cannot be applied, some of the applications can be reallocated to SMP partition with much larger amounts of memory available.

4.4 Revealing categories of issues and inefficient behavior

The accumulation of statistics and knowledge on the problems of parallel applications is one of the most important components of the HPC center job collection analysis. The ability to add tags to the analyzed jobs related to the implementation issues found is a useful feature for this purpose as well as tags corresponding to non-efficient use of computing resources, hardware problems or other features found in course of job execution characteristics analysis.

When analyzing inefficient, abnormal application behavior of a single run or of a sequence of jobs, based on the certain software package, it is often needed to contact the user, the application owner who can provide additional information on the program details: algorithm implementation used, program architecture and structure, computing model and so on up to dependencies on input data and command line options. All this information should be recorded to aid further analysis of similar applications and categories.

If any application run is being analyzed it is useful to mark and tag the used system software details. This is true first of all regarding the math libraries used, compiler and compiler options, MPI type, etc. This provides the basis for the comparative analysis of similar jobs or sequences of jobs. If differences in behavior are found, one can continue deep study on the reason origin: user application reaction, system software configuration, etc.

All widely-met issues like data race, deadlocks and so forth can be marked by special tags (***job_bug DATA RACE***, ***job_bug DEADLOCK***, etc.). This can help in further analysis of other jobs. One can compare strange program behavior to the analyzed profiles marked as having specific issues. Once a similar behavior is found, it can be a key to resolving the problems of the originally analyzed job.

5. Conclusion and future work

Close-future plans include implementation of the Octoshell [7] module for full project-oriented workflow support and authentication, thus securing accessibility of the proposed service for any user. We expect it ready by the middle of 2016. By that time we also plan to extend supported tag set and adjust criteria for existing tags if needed.

To sum up, a user-friendly, useful and effective technique for filtration, grouping and further analysis of the whole queued job collection of large-scale HPC systems based on system monitoring and resource manager data is proposed and implemented. The developed tool is evaluated in the everyday practice of the Supercomputer Center of Lomonosov Moscow State University, providing means for effective analysis for any and every user application run. The priceless collection of information on all finished jobs is already being enriched in a 24/7 mode for several month.

Acknowledgements

The work was funded in part by the Russian Found for Basic Research, grants №16–07–00972A, №13–07–00786A, and by the Ministry of Education and Science of the Russian Federation, Agreement No. 14.607.21.0006 (unique identifier RFMEFI60714X0006).

References

1. Top50 Supercomputers of Russia and CIS: [<http://top50.supercomputers.ru/>]. 15.02.2015.
2. Top500 Supercomputer sites: [<http://top500.org/>].15.02.2015.
3. Antonov A., Zhumatiy S., Nikitenko D., Stefanov K., Teplov A., Shvets P. Analysis of dynamic characteristics of job stream on supercomputer system Numerical Methods and Programming, 2013. Vol. 14, No. 2., P. 104-108.
4. Safonov A., Kostenetskiy P., Borodulin K., Melekhin F. A monitoring system for supercomputers of SUSU // Russian Supercomputing Days International Conference, Moscow, Russian Federation, 28-29 September, 2015, Proceedings. CEUR Workshop Proceedings, 2015. Vol. 1482, P. 662-666
5. Stefanov K. et al. Dynamically Reconfigurable Distributed Modular Monitoring System for Supercomputers (DiMMon) // Procedia Computer Science / Elsevier B.V., 2015. Vol. 66, P. 625–634.
6. Nikitenko D. Complex approach to performance analysis of supercomputer systems based on system monitoring data. Numerical Methods and Programming, 2014, Vol. 15, P. 85–97.
7. Voevodin V., Zhumatiy S., Nikitenko D. Octoshell: Large Supercomputer Complex Administration System // Russian Supercomputing Days International Conference, Moscow, Russian Federation, 28-29 September, 2015, Proceedings. CEUR Workshop Proceedings, 2015. Vol. 1482, P. 69-83.
8. Voevodin VI., Antonov A., Bryzgalov P., Nikitenko D., Zhumatiy S., Sobolev S., Stefanov K., Voevodin Vad. Practice of "Lomonosov" Supercomputer // Open systems, 2012. No. 7, P. 36–39.
9. Zhumatiy S., Nikitenko D. Approach to flexible supercomputers management // International supercomputing conference Scientific Services & Internet: all parallelism edges, Novorossiysk, Russian Federation, 23-28 September, 2013, Proceedings. MSU, 2013. P. 296–300.
10. Voevodin VI. Supercomputer situational screen // Open systems, 2014. No. 3, P. 36–39.
11. Shvets P. , Antonov A., Nikitenko D., Sobolev S., Stefanov K., Voevodin Vad., Voevodin V., Zhumatiy S. An Approach for Ensuring Reliable Functioning of a Supercomputer Based on a Formal Model. 11th Int. Conference on Parallel Processing and Applied Mathematics, Krakow, Poland, 6-9 September, 2015. Proceedings.

12. Voevodin V., Antonov A., Dongarra J. AlgoWiki: an Open Encyclopedia of Parallel Algorithmic Features // Supercomputing Frontiers and Innovations, 2015. Vol.2. c1. P. 4-18.
13. SLURM workload manager: [<http://slurm.schedmd.com/>]. 15.02.2015.
14. Cleo cluster batch system: [<http://sourceforge.net/projects/cleo-bs/>]. 15.02.2015.
15. Ganglia Monitoring System: [<http://ganglia.sourceforge.net/>]. 15.02.2015.
16. Collectd – The system statistics collection daemon: [<https://collectd.org/>]. 15.02.2015.
17. Clustrx: [<http://www.t-platforms.ru/products/software/clustrxproductfamily/clustrxwatch.html>]. 15.02.2015.
18. jQuery & jQuery UI: [<http://jqueryui.com/>]. 15.02.2015.
19. TagIt: [<http://aehlke.github.io/tag-it/>].15.02.2015.