



Timeline-based I/O Behavior Assessment of Parallel Jobs

An Explorative Study on 10^6 Jobs

Analyzing Parallel I/O
November 20, 2019

Eugen Betke

Research Group
German Climate Computing Center

Julian Kunkel

Compute Science
University of Reading

Table Of Contents

- 1 Pre-processing Job Data
- 2 Job-IO-Metrics
- 3 Study on 10^6 Jobs
- 4 Summary



Mistral, the HPC system for Earth system research (HLRE-3)

<i>Peak performance</i>	<i>3.14 PetaFLOPS</i>
<i>Compute nodes</i>	<i>3,300</i>
<i>Compute cores</i>	<i>100,000</i>
<i>Memory</i>	<i>266 Terabytes</i>
<i>Storage (two file systems)</i>	<i>54 Petabytes</i>

- Goals: Finding jobs with
 - ▶ high I/O load, but inefficient data access
 - e.g., for application optimization
 - ▶ critical I/O load, that can degrade file system performance
 - e.g., for better job scheduling
- Approach:
 - ▶ Define simple job metrics
 - ▶ Use them for ranking and comparison of jobs

Table Of Contents

1 Pre-processing Job Data

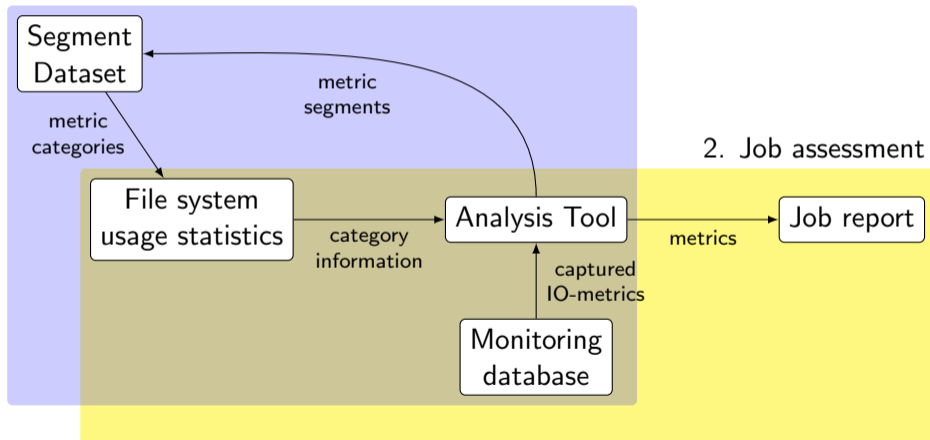
2 Job-IO-Metrics

3 Study on 10^6 Jobs


4 Summary

Analysis Workflow

1. Computing file system usage statistics



Segmentation and Scoring of Monitoring Data



Category	Criteria	MScore
LowIO	smaller than q_{99}	0
HighIO	between q_{99} and $q_{99.9}$	1
CriticalIO	larger than $q_{99.9}$	4

Categorization criteria and scores

1 Segmentation

- ▶ Segment size = 3 time points (in this example only)

2 Categorization

- ▶ Quantiles q_{99} and $q_{99.9}$ define thresholds

3 Scoring

- ▶ CriticalIO is at least 4x higher than HighIO

Score name	Definition
MScore	0,1 or 4
NScore	\sum MScore
JScore	\sum NScore

Segment scores

Table Of Contents

1 Pre-processing Job Data

2 Job-IO-Metrics

3 Study on 10^6 Jobs

4 Summary

Job-IO-Metrics [1/3]


Job-IO-Balance $\text{mean} \left(\left\{ \frac{\text{mean_score}(j)}{\text{max_score}(j)} \right\}_{j \in \text{IOJS}} \right)$

Job-IO-Utilization $\sum_{FS} \frac{\sum_{j \in \text{IOJS}} \text{max_score}(j)}{N}$

Job-IO-Problem-Time $\frac{\text{count}(\text{IOJS})}{\text{count}(\text{JS})}$

- FS: Filesystems
- JS: Job segments
- IOJS: IO-intensive job segments

Example



$$b_0 = \text{balance}(s_{j_0}) = 0.25$$

$$b_1 = \text{balance}(s_{j_5}) = 1$$

$$\text{Job-IO-Balance} = \text{mean}(\{b_0, b_1\}) = 0,625$$

Job-IO-Metrics [2/3]

Job-IO-Balance

$$\text{mean} \left(\left\{ \frac{\text{mean_score}(j)}{\text{max_score}(j)} \right\}_{j \in \text{IOJS}} \right)$$

Job-IO-Utilization


$$\sum_{FS} \frac{\sum_{j \in \text{IOJS}} \text{max_score}(j)}{N}$$

Job-IO-Problem-Time

$$\frac{\text{count}(\text{IOJS})}{\text{count}(\text{JS})}$$

- FS: Filesystems
- JS: Job segments
- IOJS: IO-intensive job segments

Example



$$\text{max}_0 = \text{max_score}(s_{j_0}) = 4$$

$$\text{max}_1 = \text{max_score}(s_{j_5}) = 1$$

$$U_{fs1} = \text{mean}(\{\text{max}_0, \text{max}_1\}) = 2.5$$

$$U_{fs2} = \text{mean}(\{\text{max}_0, \text{max}_1\}) = 2.5$$

$$\text{Job-IO-Utilization} = U_{fs1} + U_{fs2} = 5$$

Job-IO-Metrics [3/3]


Job-IO-Balance $\text{mean} \left(\left\{ \frac{\text{mean_score}(j)}{\text{max_score}(j)} \right\}_{j \in \text{IOJS}} \right)$

Job-IO-Utilization $\sum_{FS} \frac{\sum_{j \in \text{IOJS}} \text{max_score}(j)}{N}$

Job-IO-Problem-Time $\frac{\text{count}(\text{IOJS})}{\text{count}(\text{JS})}$

- FS: Filesystems
- JS: Job segments
- IOJS: IO-intensive job segments

Example



$$N_{\text{IOJS}} = 2$$

$$N_{\text{JS}} = 6$$

$$\text{Job-IO-Problem-Time} = \frac{N_{\text{IOJS}}}{N_{\text{JS}}} \approx 0.33$$

Table Of Contents

1 Pre-processing Job Data


2 Job-IO-Metrics

3 Study on 10^6 Jobs

4 Summary

Monitoring Data of 10^6 jobs

- 1,000,000 jobs for time period of 99 days
 - ▶ from 2019-05-16 to 2019-08-23
- 323634 job data remain for analysis
 - ▶ Compute partition provides data for analysis
 - ▶ Only job with exit status COMPLETED are analysed



Slurm partitions

Captured Metrics

Metadata

(Source: `/proc/fs/lustre/llite/lustre*/*/stats`)

1	<code>md_read</code>	<code>getattr + getxattr + readdir + statfs + listxattr + open + close</code>
2	<code>md_mod</code>	<code>setattr + setxattr + mkdir + link + rename + symlink + rmdir</code>
3	<code>md_file_create</code>	<code>create</code>
4	<code>md_file_delete</code>	<code>unlink</code>
5	<code>md_other</code>	<code>truncate + mmap + ioctl + fsync + mknod</code>

Data

(Source: `/proc/fs/lustre/llite/lustre*/*/read_ahead_stats`)

6	<code>read_bytes</code>	
7	<code>read_calls</code>	
8	<code>write_bytes</code>	Application I/O requests to Lustre.
9	<code>write_calls</code>	
<hr/>		
10	<code>osc_read_bytes</code>	
11	<code>osc_read_calls</code>	Lustre object storage client requests to object storage.
12	<code>osc_write_bytes</code>	
13	<code>osc_write_calls</code>	


File System Usage Statistics

Metric		Limits		Number of occurrences		
Name	Unit	q99	q99.9	LowIO	HighIO	CriticalIO
md_file_create	Op/s	0.17	1.34	65,829K	622K	156K
md_file_delete	Op/s	0.00	0.41	65,824K	545K	172K
md_mod	Op/s	0.00	0.67	65,752K	642K	146K
md_other	Op/s	20.87	79.31	65,559K	763K	212K
md_read	Op/s	371.17	7084.16	65,281K	1,028K	225K
osc_read_bytes	MiB/s	1.98	93.58	17,317K	188K	30K
osc_read_calls	Op/s	5.65	32.23	17,215K	287K	33K
osc_write_bytes	MiB/s	8.17	64.64	16,935K	159K	26K
osc_write_calls	Op/s	2.77	17.37	16,926K	167K	27K
read_bytes	MiB/s	28.69	276.09	66,661K	865K	233K
read_calls	Op/s	348.91	1573.45	67,014K	360K	385K
write_bytes	MiB/s	9.84	80.10	61,938K	619K	155K
write_calls	Op/s	198.56	6149.64	61,860K	662K	174K

Job-IO-Metric Distributions


Job-IO-Balance

$$\text{mean} \left(\left\{ \frac{\text{mean_score}(j)}{\text{max_score}(j)} \right\}_{j \in \text{IOJS}} \right)$$




Job-IO-Utilization

$$\sum_{FS} \frac{\sum_{j \in \text{IOJS}} \text{max_score}(j)}{N}$$




Job-IO-Problem-Time

$$\frac{\text{count}(\text{IOJS})}{\text{count}(\text{JS})}$$




Jobs with high I/O-Intensity

$$\text{Job-IO-Intensity} = B \cdot PT \cdot U \cdot \text{total_nodes}$$




30 jobs ordered by IO-Intensity




Nodes: 100; B: 0.88; PT:1.0; U: 4.0

Jobs with high Problem-Score

$$\text{Problem-Score} = (1 - B) \cdot \text{PT} \cdot U$$



5000 jobs ordered by Problem-Size, and runtime $>$ 30min,



Example for high Problem-Size: Nodes: 4; B: 0.37;
PT: 1.0; U: 8.0

Table Of Contents

1 Pre-processing Job Data

2 Job-IO-Metrics

3 Study on 10^6 Jobs

4 Summary

Summary

- Applied methods
 - ▶ **Segmentation**: Preserves time line information
 - ▶ **Categorization**: Filters not significant I/O and make incompatible metrics compatible
 - ▶ **Scoring**: Allows mathematical computation
- Job-IO-Problem-Time, Job-IO-Balance and Job-IO-Utilization
 - ▶ Are **simple metrics**, that describe key properties of parallel jobs
- IO-Intensity and IO-Problem-Score
 - ▶ Are **penalty functions**, used for job ranking

Thank you for your attention!