

Analyzing IO Usage Patterns of User Jobs to Improve Overall HPC System Efficiency

Syed Sadat Nazrul*, Cherie Huang*, Mahidhar Tatineni,
Nicole Wolter, Dimitry Mishin, Trevor Cooper and Amit Majumdar
San Diego Supercomputer Center
University of California San Diego
* students at the time of project

SCEC2018, Delhi, Dec 13-14, 2018

Comet

“HPC for the long tail of science”




iPhone panorama photograph of 1 of 2 server rows


Comet: System Characteristics

- **Total peak flops ~2.1 PF**
- **Dell primary integrator**
 - Intel Haswell processors w/ AVX2
 - Mellanox FDR InfiniBand
- **1,944 standard compute nodes (46,656 cores)**
 - Dual CPUs, each 12-core, 2.5 GHz
 - 128 GB DDR4 2133 MHz DRAM
 - 2*160GB GB SSDs (local disk)
- **72 GPU nodes**
 - 36 nodes same as standard nodes plus Two NVIDIA K80 cards, each with dual Kepler3 GPUs
 - 36 nodes with 2 14-core Intel Broadwell CPUs plus 4 NVIDIA P100 GPUs
- **4 large-memory nodes**
 - 1.5 TB DDR4 1866 MHz DRAM
 - Four Haswell processors/node
 - 64 cores/node
- **Hybrid fat-tree topology**
 - FDR (56 Gbps) InfiniBand
 - Rack-level (72 nodes, 1,728 cores) full bisection bandwidth
 - 4:1 oversubscription cross-rack
- **Performance Storage (Aeon)**
 - 7.6 PB, 200 GB/s; Lustre
 - Scratch & Persistent Storage segments
- **Durable Storage (Aeon)**
 - 6 PB, 100 GB/s; Lustre
 - Automatic backups of critical data
- **Home directory storage**
- **Gateway hosting nodes**
- **Virtual image repository**
- **100 Gbps external connectivity to Internet2 & ESNet**

~67 TF supercomputer in a rack

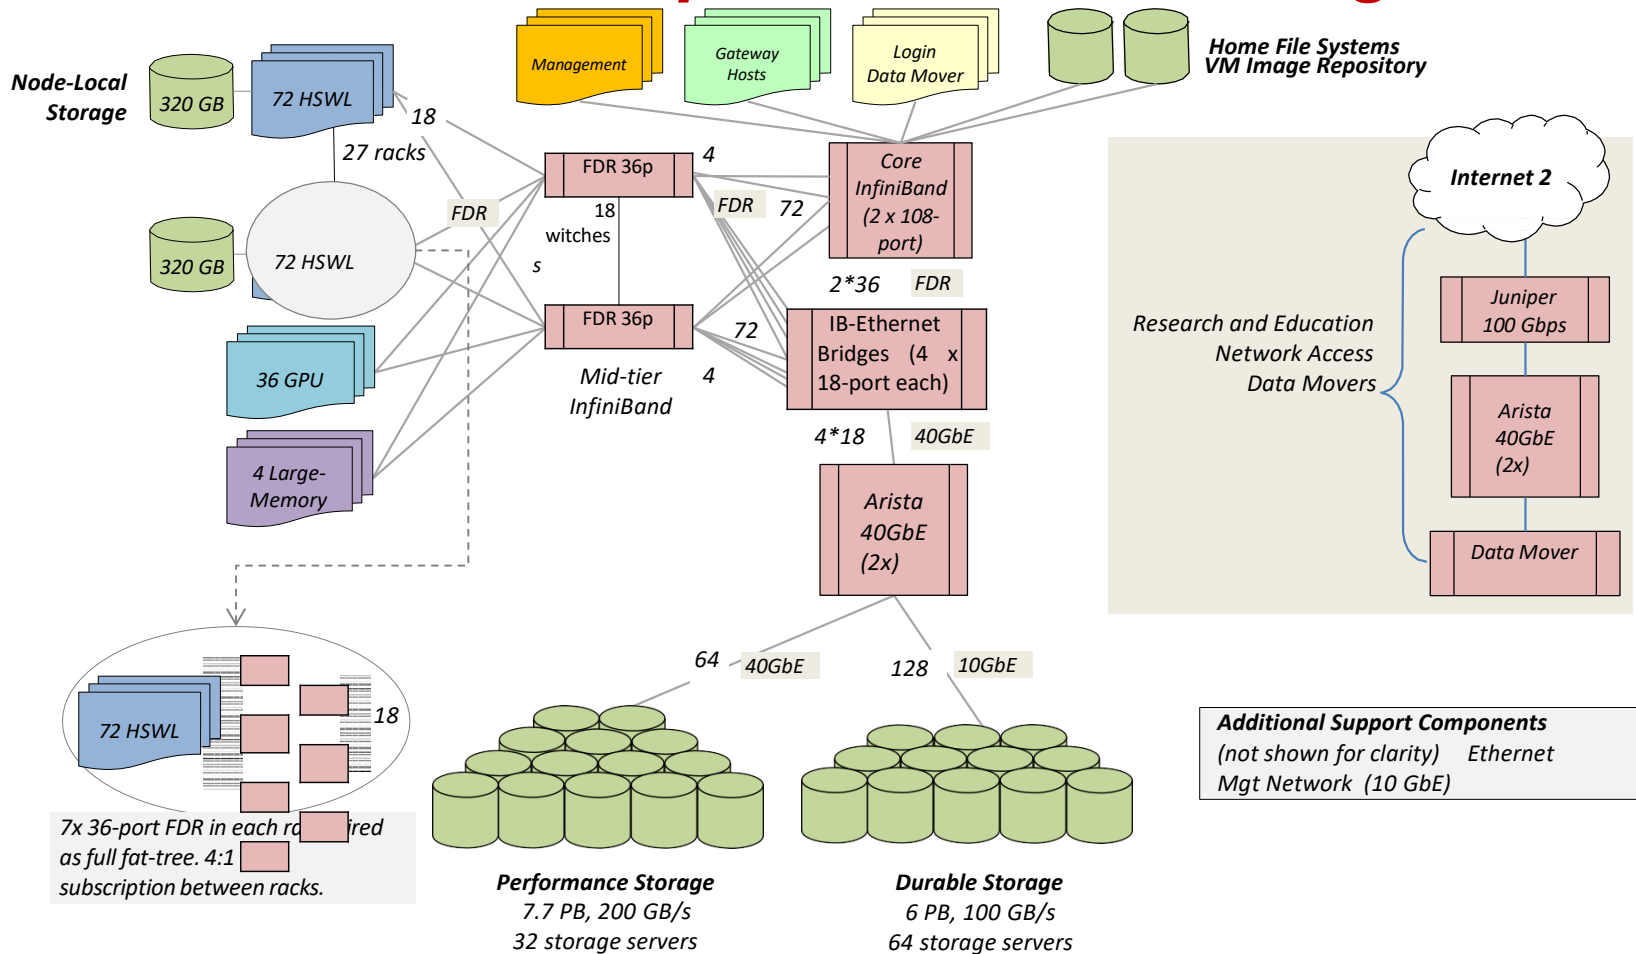


And 27 single-rack supercomputers



Comet Network Architecture

InfiniBand compute, Ethernet Storage



Comet: Filesystems

- **Lustre filesystems – Good for scalable large block I/O**
 - Accessible from all compute and GPU nodes.
 - */oasis/scratch/comet* - 2.5PB, peak performance: 100GB/s. Good location for storing large scale scratch data during a job.
 - */oasis/projects/nsf* - 2.5PB, peak performance: 100 GB/s. Long term storage.
 - **Not good for lots of small files or small block I/O.**
- **SSD filesystems**
 - */scratch* local to each native compute node – 210GB on regular compute nodes, 285GB on GPU, large memory nodes, 1.4TB on selected compute nodes.
 - SSD location is good for writing small files and temporary scratch files. Purged at the end of a job.
- **Home directories (*/home/\$USER*)**
 - Source trees, binaries, and small input files.
 - **Not good for large scale I/O.**

Motivation

- **Currently HPC systems monitor/collect lots of data**
 - Network traffic, file system traffic (I/O), CPU utilization etc.
 - Analyzing users' job data can provide insight into static and dynamic loads on
 - File system
 - Network
 - Processors
- **How to analyze data, observe patterns, use those for improved system operation**
- **Analysis of I/O usage patterns of users' jobs**
 - Insight into which jobs to schedule together or not
 - System admins perform I/O work coordinating with specific user jobs etc.

This work - preliminary

- Looked at I/O traffic of users' job on Comet for three months – early phase of Comet: June – November 2015
- Analyze data and extract information
 - Monitor system operation
 - Improve system operation
- Aggregate I/O usage pattern of users' jobs
 - On NFS, Lustre and node-local SSDs
- Data science applied to tie I/O usage pattern to users' particular codes

Data Analysis


- Data collected using TACC Stats (still being collected continuously)
- ~700,000 jobs that ran during the time period, and is around 500 GB in size
 - Collects user job's I/O stats on file systems every 10 min interval
- Looked at Compute and GPU queue (not shared queue for first pass)
- Data can be quickly extracted as inputs for learning algorithms – NFS, Lustre, node local SSD I/O data



- Ran controlled IOR for validating the data processing pipeline

Scatter plot

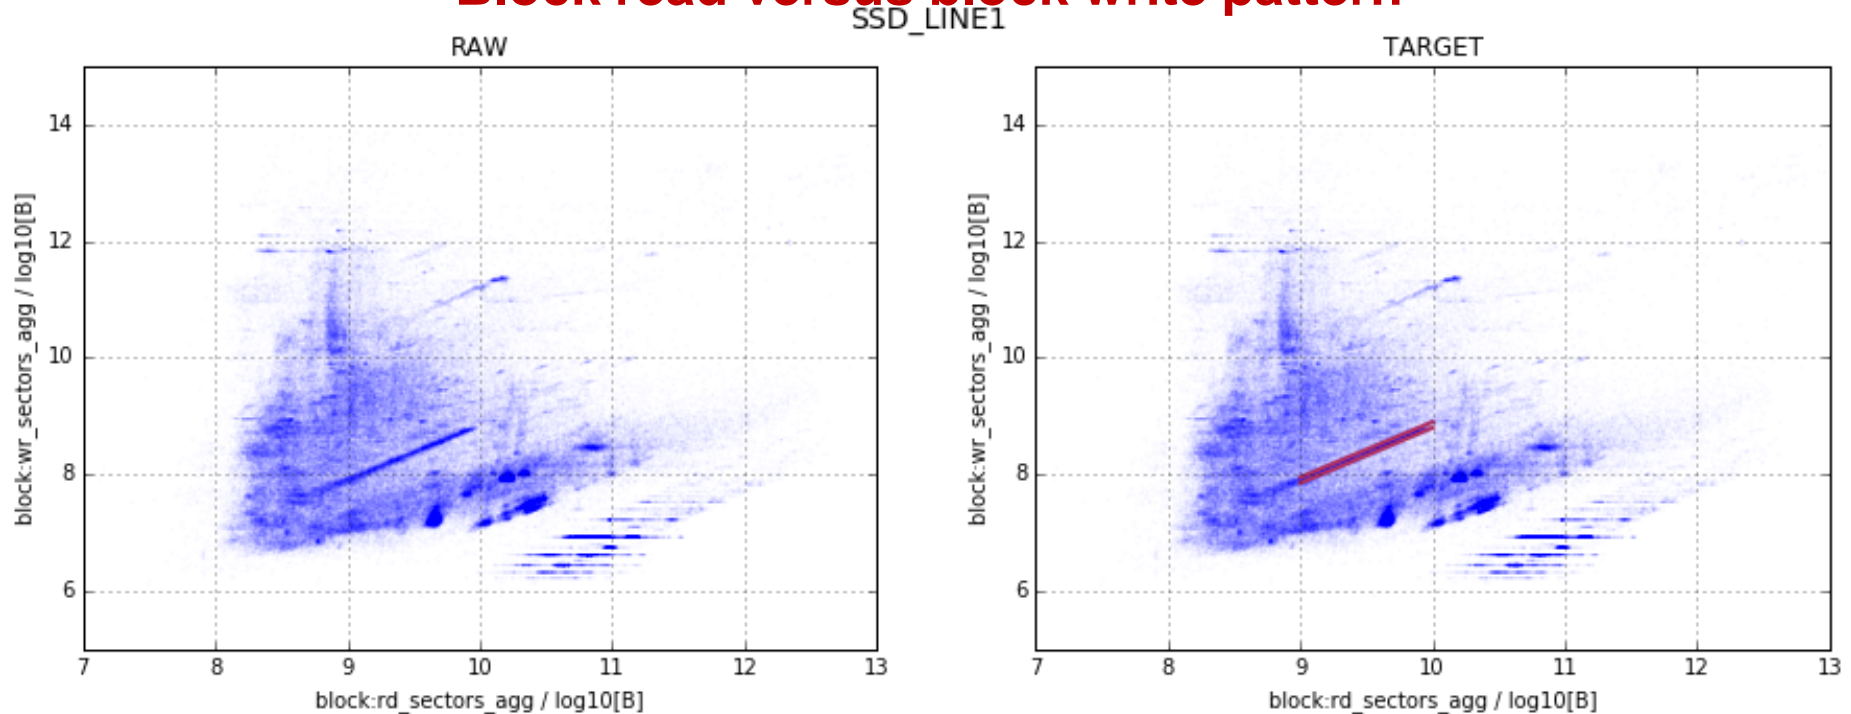
No Clustering
Units: Log10[B]
Low Alpha



- *scatter matrix from Scikit-learn*
- Block refers to SSD
- llite refers to Lustre
- Analyzed the linear patterns
- Tried to tie to apps

Linear Pattern

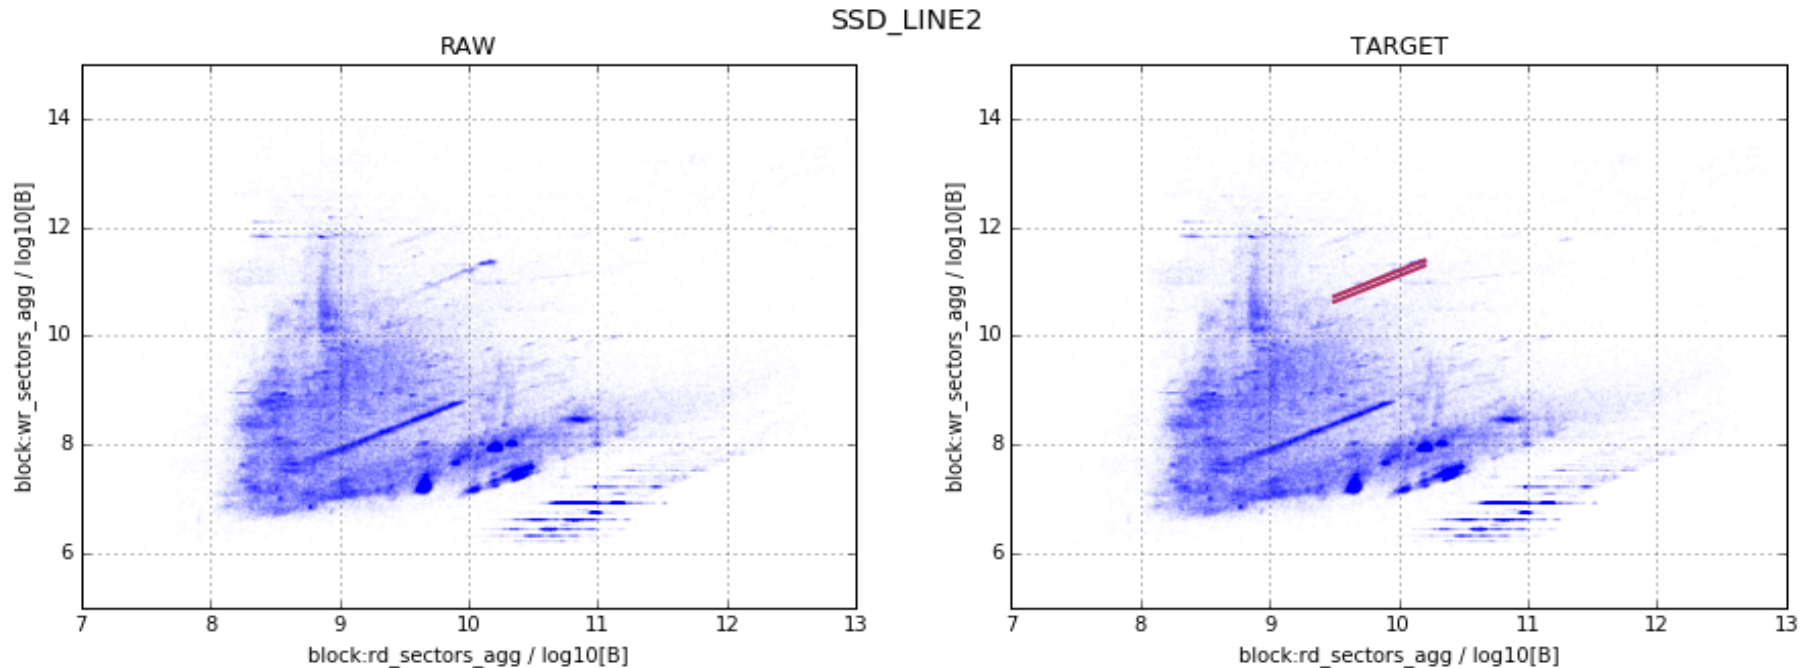
Block read versus block write pattern



- Linear patterns formed when analyzing aggregate write I/O and aggregate read I/O on the SSD
- Pertaining to all the jobs that are part of this pattern, we have seen that 1,877 (76%) jobs are Phylogenetics Gateway (CIPRES running RXML code) and Neuroscience Gateway (was mostly running spiking neuronal simulation) jobs
- We know that these jobs only produce I/O to NFS
- However they used OpenMPI for their MPI communication.
- This leads to runtime I/O activity (for example memory map information) in /tmp which is located on the SSDs

Linear Pattern

Block read versus block write pattern



- *Another linear pattern formed when analyzing aggregate write I/O and aggregate read I/O on the SSD*
- *Pertaining to all the jobs that are part of this pattern, we have seen that 208 (82%) jobs have the same job name and from a particular project group*
- *Further investigation and discussion with the user showed that these I/O patterns were produced by Hadoop jobs*
- *On Comet, Hadoop is configured to use local SSD as the basis for its HDFS file system*
- *Hence, as expected, there is a significant amount of I/O to SSDs from these jobs*

Linear pattern

SSD read vs Lustre write; SSD read vs Lustre read

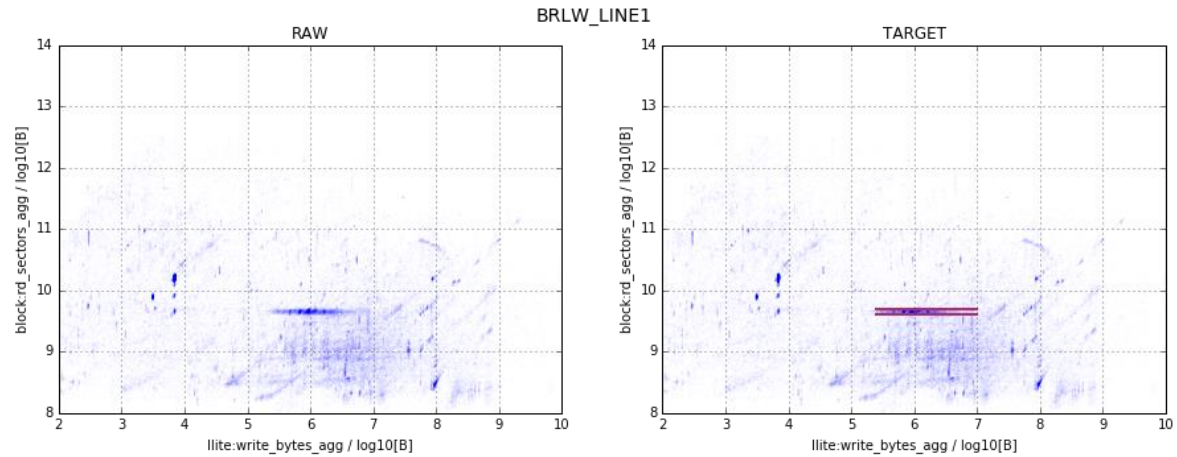


Fig. 6. Block read versus lustre write pattern (BRLW_LINE1).




Fig. 7. Block read versus lustre read pattern (BRLR_LINE1) – horizontal line.

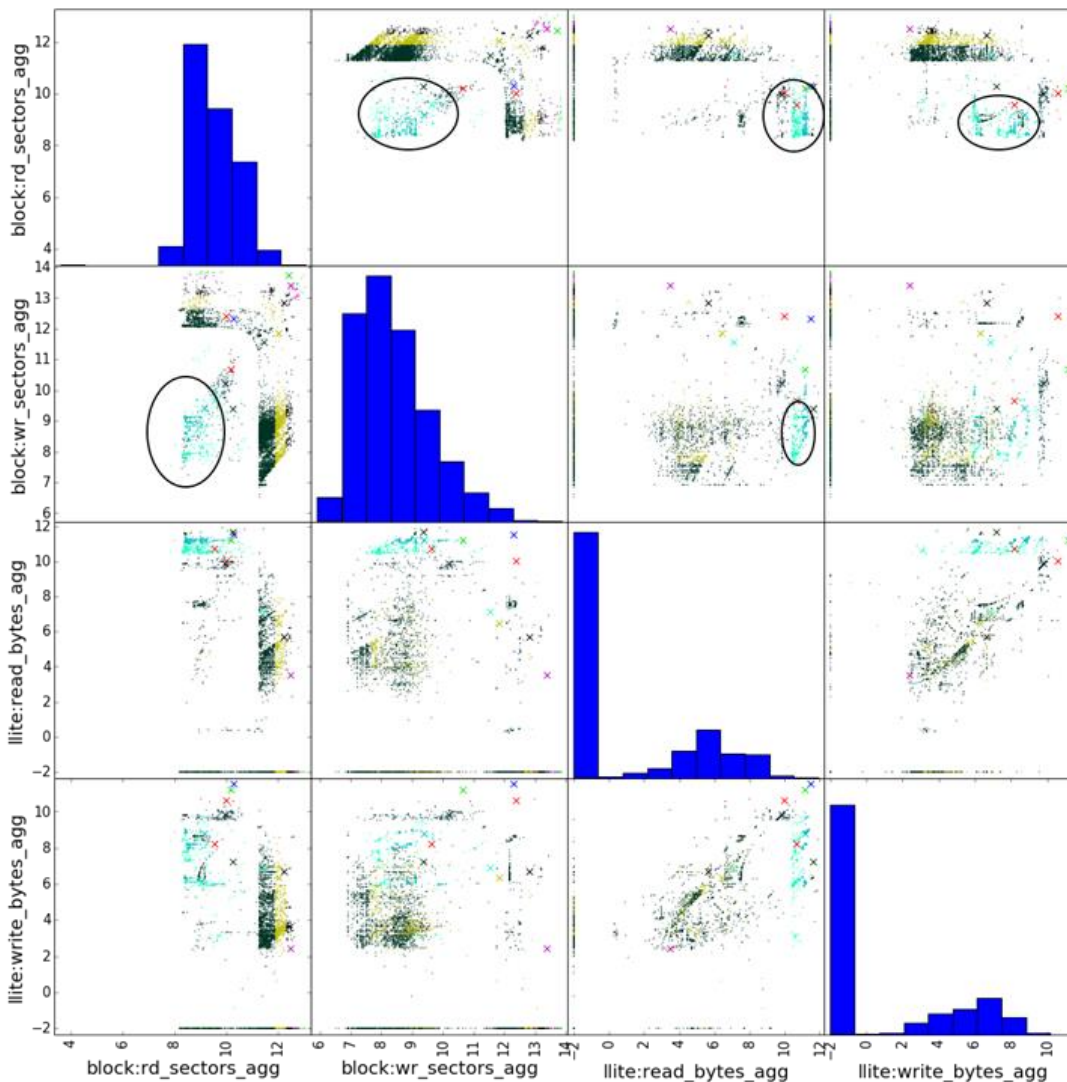
Linear pattern

SSD read vs Lustre write; SSD read vs Lustre read

- Horizontal linear patterns on SSD read I/O against Lustre Write I/O and Lustre Read I/O respectively
- Both show similar patterns.
- This indicates that they were both created by similar applications
- BRLW_LINE1 contains 232 (28%) VASP and CP2K jobs and 134 (16%) NAMD jobs
- We can say these applications require ~4 GB of read from the local SSD (this includes both scratch and system directories) and between 100 kB and 10 MB Lustre I/O (both read and write) to run the job

K-means analysis cluster center marks 'X' and cluster 10 encircled

k-means
Units: Log10[B]
Without Cluster 0 (97% of jobs)



K-means cluster analysis

- The teal colored cluster as shown in Figure, is characterized by low SSD read and SSD write (100 MB - 1 GB)
- However, this cluster shows very high Lustre read (>10 GB) and variable Lustre write (100 kB - 1 GB)
- At least 324 (89%) of these jobs had projects that indicate that these are astrophysics jobs

Summary

- **We did some other analysis such as using DBSCAN, longer (than 10 mins) time window for data etc.**
 - No distinct patterns
- **Presented work show we were able to analyze distinct patterns in the dataset caused by different applications**
- **We only looked at aggregate data**
 - In the future examine time series data - beginning, middle end of job
- **We can also analyze jobs separately based on parameters like run time of the job**

Acknowledgement: Partial funding from Engility for student research internship