HPC I/O in the Data Center Workshop (HPC-IODC)

Julian M. Kunkel, Jay Lofstead, Jean-Thomas Acquaviva

[1] University of Reading
Whiteknights
Reading RG6 6AY, UK
j.m.kunkel@reading.ac.uk
[2] Center for Computing Research
Sandia National Laboratories
Albuquerque, USA
gflofst@sandia.gov
[3] Data Direct Networks, France
jtacquaviva@ddn.com

1 Introduction

Many public and privately funded data centers host supercomputers for running large scale simulations and analyzing experimental and observational data. These supercomputers run usually tightly coupled parallel applications that require hardware components that deliver the best performance. In contrast, commercial data centers, such as Facebook and Google, execute loosely coupled workloads with a broad assumption of regular failures. The dimension of the data centers is enormous. A 2013 article summarizes commercial data centers' dimensions [4]. It estimates, for example, that Facebook hosts around 100 PB of storage, and Google and Microsoft manage around 1 million servers each – although the hardware is split among several physical data centers – a modus operandi not suitable for HPC centers. With the increasing importance of using machine learning to reveal underlying patterns in data, the data storage rates are accelerating to feed these additional use cases. Combining traditional modeling and simulation with ML workloads yields both a write and read-intensive workload for a single workflow.

Management of the huge amount of data is vital for the effective use of the contained information. However, with limited budgets, it is a daunting task for data center operators, especially as the design and storage system required hardware depends heavily on the executed workloads. A co-factor of the increasing difficulty is the increase in complexity of the storage hierarchy with the adoption of SSD and memory class storage technology. The US Department of Energy recognizes the importance of data management, listing it among the top 10 research challenges for Exascale [3].

There are several initiatives, consortia and special tracks in conferences that target RD&E audiences. Examples are the Storage Networking Industry Association (SNIA) for enterprises, the Big Data and Extreme-Scale Computing

(BDEC) initiative[1], the Parallel Data Systems Workshop (PDSW) and the HEC FSIO workshop [1].

There are many I/O workloads studies and performance analysis reports for parallel I/O available. Additionally, many surveys of enterprise technology usage include predictions of analysis for future storage technology and the storage market [2]. However, the analysis conducted for HPC typically focuses on applications and not on the data center perspective. Information about data center operational aspects is usually described in file system-specific user groups and meetings or described partially in research papers as part of the evaluation environment.

In the HPC IODC workshop, we bring together I/O experts from data centers and application workflows to share current practices for scientific workflows, issues, and obstacles for both hardware and the software stack, and RD&E to overcome these issues.

Due to the COVID-19 crisis, the ISC conference changed to a digital edition. We preserved the nature of the workshop and organized it as a virtual full-day meeting on the regular workshop day with minimal changes to the agenda.

2 Organization of the Workshop

The workshop was organized by

- Julian Kunkel (*University of Reading, UK*)
- Jay Lofstead (*Sandia National Labs, USA*)
- Jean-Thomas Acquaviva (*DDN*)

The workshop is supported by the Centre of Excellence in Simulation of Weather and Climate in Europe (ESiWACE), the Virtual Institute for I/O (VI4IO)[2] and the Journal of High-Performance Storage (JHPS)[3].
The workshop covered three tracks:

- **Research paper presentations** – authors needed to submit a paper regarding relevant state-of-the-practice or research for I/O in the datacenter.
- **Talks from I/O experts** – authors needed to submit a rough outline for the talk related to the operational aspects of the data center.
- A moderated **discussion** to identify key issues and potential solutions in the community.

The CFP has been issued at the beginning of January. Important deadlines were:

- Submission deadline: 2020-04-15 AoE
- Author notification: 2020-05-03
- Workshop: 2020-06-25
- Camera-ready papers: 2020-07-25

From all submissions, the program committee selected three talks from I/O experts and eleven research papers for presentation during the workshop.

[1] `http://www.exascale.org/bdec/`
[2] `http://vi4io.org`
[3] `https://jhps.vi4io.org/`

2.1 Programm Committee

- Thomas Boenisch *(High-performance Computing Center Stuttgart)*
- Suren Byna *(Lawrence Berkeley National Laboratory)*
- Matthew Curry *(Sandia National Laboratories)*
- Sandro Fiore *(CMCC)*
- Wolfgang Frings *(Juelich Supercomputing Centre)*
- Javier Garcia Blas *(Carlos III University)*
- Adrian Jackson *(The University of Edinburgh)*
- Ivo Jimenez *(University of California, Santa Cruz)*
- Anthony Kougkas *(Illinois Institute of Technology)*
- Glenn Lockwood *(Lawrence Berkeley National Laboratory)*
- Jay Lofstead *(Sandia National Laboratories)*
- Carlos Maltzahn *(University of California, Santa Cruz)*
- Suzanne McIntosh *(New York University)*
- Maria Perez *(Technical University of Madrid)*
- Robert Ross *(Argonne National Laboratory)*
- George S. Markomanolis *(Oak Ridge National Laboratory)*
- Feiyi Wang *(Oak Ridge National Laboratory)*
- Bing Xie *(Oak Ridge National Lab)*

3 Workshop Summary

Over the full-day program, about 40 attendees were constantly connected to the virtual session. More than 100 participants expressed their interest to be informed about the workshop slides and presentations. In spite of the workshop being held online squarely on British Time making attending from North America more difficult. This is in line with the in-person attendance at previous instances and included many North American attendees, including some from the American west coast.

We had a good mix of talks from I/O experts, data center relevant research, and two discussion sessions. A short summary of the presentations is given in the following. The slides and video recordings of the presentations are available on the workshop's webpage: `https://hps.vi4io.org/events/2020/iodc`.

3.1 Research Papers

We have shifted our peer review process to be more community building oriented. Last year, we changed the review process to shepherd all papers with a solid core, but potentially presentation flaws, to help develop them so that they are acceptable for publication. If a paper cannot be successfully be revised in time for the workshop, it will be rejected. We find this approach is better for building an interactive community.

Our goal was to create a more open, fully interactive process for quickly developing research papers into quality publishable results.

To support this activity, we were excited to explore the open review process of the Journal of High-Performance Storage for the research track – the papers were publicly shared using GoogleDoc in the JHPS incubator and open for comments. We allowed authors to submit extended papers that may be potentially accepted on JHPS while the 12-page shorter workshop submissions will be published in this edition of Springer's Lecture Notes on Computer Science together with the other ISC workshops as was done previously.

Unfortunately, due to COVID-19, we received notes from various authors that they would not be able to make the deadline, hence the number of submissions this year was lower than usual.

In the first research papers session, HPC-IODC was presenting three different talks sharing the same interest for I/O performance characterization and analysis. Each of these talks brings its own originality, either with the usage of none HPC API for high-performance storage, or taking into account network considerations for I/O, or machine learning techniques applied to IO analysis. The following presentation try to capture the presentation itself as some elements of the discussion triggered. In more details, the research presentation covered the following topics:

- **Characterizing I/O Optimization Effect Through Holistic Log Data Analysis of Parallel File Systems and Interconnects**
  Yuichi Tsujita, Yoshitaka Furutani, Hajime Hida, Keiji Yamamoto, Atsuya Uno
  In his presentation, Dr. Tsujita discussed the value of observing not only file system metrics but also networks aspect in order to understand the I/O behavior of an application. Regarding the result of these analyses, Dr. Tsujita notes the difference between observation and resolution: as a different toolbox is needed to optimize user-level I/O patterns. For instance, at the K computer, experts were providing hints or training to improve I/O specifically for the complex MPI-IO. It remains that most users preferred POSIX-I/O. Therefore, POSIX is now the highest priority issue. Discussing further the weight of user choices, the author notes that despite the availability of optimized MPI-IO libraries as dynamically loadable shared libraries, most of the users keep the default MPI-IO implementation. One motivation of this work is to raise awareness among the user community, and that access to profiling information could push users to have interests in advanced I/O optimization.
- **Investigating the Overhead of the REST Protocol to Reveal the Potential for Using Cloud Services for HPC Storage**
  Frank Gadban, Julian Kunkel, Thomas Ludwig
  This work investigated the relevance of cloud APIs to address performance-oriented tasks. The authors observe that a correctly configured *REST* service can deliver high performance and match MPI in terms of bandwidth, thus making *REST* a relevant alternative to MPI based IO tools. Bandwidth was not the only investigated metric and the authors monitor as well the CPU cost of operating data movement. They observed that MPI can lead to a poor

ratio in terms of bandwidth/cycle. However, the discussion arises and the point was made that in HPC and MPI in particular, CPU cycles can be used to decrease latency for instance. Therefore in terms of methodology, a higher rate of idle cycles is not always positive, such metric would systematically penalize spin locking a classic usage of CPU resources to improve latency and time-critical performance. In the aspect of performance optimization, specifically for complex systems where performance can be capped by any of its components, reporting not raw performance but efficiency in respect of the network throughput and CPU utilization could an interesting way to extrapolate the results on the different architectures (e.g. HDR200 or faster CPUs).

– **Classifying Temporal Characteristics of Job I/O Patterns Using Machine Learning Techniques**
Eugen Betke, Julian Kunkel
In this talk, the authors present their results on the classification of HPC jobs based on collected information on their I/O pattern. I/O bottlenecks are a multi-factorial artifact due to a combination of I/O Patterns + IO Configuration System + I/O Workload. The processed data set covers an impressive 1.000.000 jobs. This work follows a previous and published study focused on the identification of IO-intensive jobs. However, their experience has shown that for large systems I/O intensive jobs could be numerous and the single job is not the right granularity to provide support and help to end-users. This has been the key motivation to pursue the research effort in order to classify and to cluster jobs. This is the purpose of the present paper. The resulting classification scheme is lightweight enough to be operated as soon as a job is over but can not be performed on-line (during job execution), so the method is near-line. The proposed methodology did not achieve the point where it can characterize I/O, for example, at the moment, it can not tell if jobs in this and these clusters can have a negative impact on the metadata server. This is work in progress, and at this stage understanding, I/O is still a manual task. It remains that classification is an important step toward the goal of automation. The benefit at the moment is the following: assuming a job is identified as slowing down of a file system (e.g massive meta-data operations), with a simple lookup it is possible to find the whole cluster of jobs that can potentially cause the same issue to the file system. During the discussion, a point was made on conceptual constraints induced by an approach based on Machine Learning and Deep Learning. The authors conclude their talk on a long term objective of a learning model representative of the whole complexity of HPC-IO systems.

– **A Reinforcement Learning Strategy to Tune Request Scheduling at the I/O Forwarding Layer**
Jean Luca Bez, Francieli Zanon Boito, Ramon Nou, Alberto Miranda, Toni Cortes, Philippe O. A. Navaux
I/O optimization techniques can improve performance for the access patterns they were designed to target, but they often decrease for others. Moreover, these techniques usually depend on the precise tune of their parameters,

which commonly falls back to the users. The authors propose an approach to tune parameters dynamically at runtime based on the I/O workload observed by the system. Our focusing is on the I/O forwarding layer as it is transparent to applications and file system independent. The approach uses a reinforcement learning technique to make the system capable of learning the best parameter value to each observed access pattern during its execution, eliminating the need for a complex and time-consuming training phase. The authors evaluated the proposal for the TWINS scheduling algorithm designed for the I/O forwarding layer seeking to reduce contention and coordinate accesses to the data servers. They demonstrate the approach can reach a precision of 88% on the parameter selection in the first hundreds of observations of an access pattern, achieving 99% of the optimal performance. After the talk, it was noted that a baseline to optimal performance is necessary for comparison reason, as a naive approach may already lead to good performance numbers.

- **Data Systems at Scale in Climate and Weather: Activities in the ESiWACE Project**
  Julian Kunkel and the ESiWACE4 team.
  The ESiWACE project aims to enable global eddy-resolving weather and climate simulations on the upcoming (pre-)Exascale supercomputers. In this talk, a selection of efforts to mitigate the effects of the data deluge from such high-resolution simulations is introduced. In particular, the speaker described the advances in the Earth System Data Middleware (ESDM), which enables scalable data management and supports the inhomogeneous storage stack. ESDM provides a NetCDF compatible layer at a high-performance and portable-portable fashion. A selection of performance results was given and ongoing efforts for workflow support and active storage are discussed.
- **Phobos a scale-out object store implementing tape library support**
  Patrice Lucas, Philippe Deniel, Thomas Leibovici Phobos is an open-source scale-out distributed object store providing access to multiple backends from flash and hard drives to tape libraries. Very large datasets can be efficiently managed on inexpensive storage media without giving up performance, scalability or fault-tolerance. Phobos is designed to offer several data layouts, such as mirroring or erasure coding. IOs through tape drives are optimized by dedicated resource scheduling policies. Developed at CEA, Phobos is in production since 2016 to manage the France Genomique multi-petabyte dataset at TGCC.

3.2 Talks from Experts

The following talks from experts included some basic information about the site and typical application profiles but focus on information regarding I/O tools and strategies applied to mitigate the pressing issues. The first session of three experts talks are presenting aspect from production environments, the topics are rather diverse but the goal is shared: making a high-performance storage service to end-users available.

– **The ALICE data management pipeline**
Massimo Lamanna
ALICE is an experiment at CERN generating a large volume of data, in the range of 60PB. The data processing pipeline is complex starting with a very fast ingest phase in the range of 3.5 TB/sec. This phase is embarrassingly parallel and data can be compressed in the range x6. In the second phase, compressed data percolate through a compute-intensive process. The process refines and compresses further data, eventually leading to an additional factor x6 of compression. This later stage of compression is not completely lossless and is run on a large GPU farm. Previous experiments at CERN were relying on data replication as a protection scheme ALICE will be the first experiment using massively Erasure Code. The workload will be dominated by large files. Ceph would be a natural candidate for ALICE. In this respect, CERN has a dual policy toward Ceph: promotion and support. As an example, Ceph has been considered for OpenStack and Ceph is offered as a service in CERN alternatively to the homegrown file system EOS, but CERN tends to rely on EOS for extremely large configurations.EOS uses for some use cases Ceph as a back-end. The benefit of EOS is that CERN can shape the software to be exactly as needed and to satisfy specific requirements.
– **Accelerating your Application I/O with UnifyFS**
Kathryn Mohror
UnifyFS is an on-going effort initiated at Lawrence Livermore National Laboratory focuses on the development of a file system harnessing node-local storage resources. The motivation is the renewed importance of node-local storage with the emergence of Storage Class Memory. To limit the cost of synchronization and communication between peers of UnifyFS the system alleviates the POSIX constraints. Data are publicized when an explicit laminating process occurs. The lamination is quite similar to synchronization between the peer but the file is made immutable. Lamination is either implicit when a file is closed or explicit when *chmod()* is called. It should be noticed that lamination is about the content (file data) or not on the file path or metadata. The distance taken with the strict POSIX semantic is bringing performance improvements as illustrated by benchmark results or workload similar to checkpoint-restart.
– **How to recognise I/O bottlenecks and what to do about them**
Rosemary Francis
Rosemary Francis is the CEO of Ellexus company, a start-up specialized in I/O optimization and profiling. Ellexus develops two different products: Mistral and Breeze. Breeze allows an in-depth analysis of the I/O activity but generates a large amount of log information, in the range of *strace*. In order to integrate an I/O profiler within a job life-cycle, the second tool Mistral is more suitable. From her extensive operational knowledge of I/O management, Dr. Francis observes that AI applications are newer and cleaner: no legacy code or legacy libraries, consequently IO patterns tend to be much cleaner and saner. Furthermore, she observes that AI IO patterns are agnostic to parallelism, using a small size configuration it is easy to extrapolate

the IO behavior on a larger system. For instance, insights collected on a single desktop have allowed the speaker to detect and fix an x3 performance issue on the whole cluster. The root cause of the performance issue was a misconfiguration in the Linux settings.

– **Managing Decades of Scientific Data in Practice at NERSC**
Glenn Lockwood

The National Energy Research Scientific Computing Center (NERSC) has been operating since 1974 and has been storing and preserving user data continuously for over 45 years as a result. This has resulted in NERSC building significant expertise in how to store and manage user data for long periods of time–a decade or more–and the practical factors that must be considered when data must be retained for longer than the lifetime of the physical components of the data center, including the entire data center facility itself. As the relevance of HPC extends beyond modeling and simulation and the usable lifetime of data extends from months to years or decades, these best practices in long-term data stewardship are likely to become more important to more HPC facilities. To this end, the speaker presented some of the practical considerations, best practices, and lessons learned from managing the scientific data of NERSC's thousands of users over a period of four decades.

– **Portable Validations of Scientific Explorations with Container-native Workflows**
Ivo Jimenez

Researchers working in the computer, computational, or data science often find it difficult to reproduce experiments from artifacts like code, data, diagrams, and results which are left behind by previous researchers. The code developed on one machine often fails to run on other machines due to differences in hardware architecture, OS, software dependencies, among others. This is accompanied by the difficulty in understanding how artifacts are organized, as well as in using them in the correct order. Software container technology such as Docker, can solve most of the practical issues of portability, and in particular, container-native workflow engines can significantly aid experimenters in their work. In this talk, Popper was introduced, a container-native workflow engine that executes each step of a workflow in a separate dedicated container without assuming the presence of a Kubernetes cluster or any cloud-based Kubernetes service. With Popper, researchers can build and validate workflows easily in almost any environment of their choice including local machines, SLURM based HPC clusters, CI services, or Kubernetes based cloud computing environments. To exemplify the suitability of this workflow engine, three case studies with examples from Machine Learning and High-Performance Computing are turned into Popper workflows. It was also discussed how Popper can be used to aid in preparing artifacts associated with article submissions to conferences and journals, and in particular give an overview of the Journal of High-Performance Storage, a new eJournal that combines open reviews, living papers, digital reproducibility, and open access.

– **Tuning I/O Performance on Summit: HDF5 Write Use Case Study**
Xie Bing
The HDF5 I/O library is widely used in HPC across a variety of domain sciences for its simplicity, flexibility, and rich performance-tuning space. In this work, the authors addressed an observed HDF5 write performance issue on Summit at OLCF, which in particular is the poor write performance of HDF5 with the default configuration. To identify the performance issue, they developed an I/O benchmarking methodology to profile the HDF5 performance on Summit across scales, compute-node allocations, I/O configurations, and times. They developed a solution to the issue by altering the HDF5 alignment configuration which resulted in a 100x write performance improvement for the VPIC benchmark. The speaker expects the methodology and solution to be applicable to other platforms and technologies.

3.3 Discussion Sessions

The major distinguishing feature for this workshop compared to other venues is the discussion rounds. The opportunity for themed, open discussions about issues both pressing and relevant to the data center community facilitates sharing experiences, solutions, and problems.

Albeit the workshop was virtual, the discussion covered aspects around APIs, benchmarking, node-local IO vs. shared storage, setting defaults for users, storage, and code maintenance.

References

1. Bancroft, M., Bent, J., Felix, E., Grider, G., Nunez, J., Poole, S., Ross, R., Salmon, E., Ward, L.: Hec fsio 2008 workshop report. In: High End Computing Interagency Working Group (HECIWG), Sponsored File Systems and I/O Workshop HEC FSIO (2009)
2. IDC: Enterprise storage services survey. `http://www.idc.com/getdoc.jsp?containerId=254468`
3. Lucas, R., committee members: Top ten exascale research challenges. Online `http://science.energy.gov/~/media/ascr/ascac/pdf/meetings/20140210/Top10reportFEB14.pdf` (Feb 2014)
4. StorageServers Blog: Facts and stats of world's largest data centers. Online `https://storageservers.wordpress.com/2013/07/17/facts-and-stats-of-worlds-largest-data-centers/` (July 2013)