

3rd HPC I/O in the Data Center Workshop

Julian M. Kunkel, Jay Lofstead, Colin McMurtrie

kunkel@dkrz.de, gflfst@sandia.gov, cmurtrie@cscs.ch

DKRZ, Sandia National Lab, CSCS

06-22-2017





Agenda

9:00 *Welcome*

9:10 **Research paper session (10 min each)**

10:00 **Expert talk session 1 (20 min each)**

11:00 *Coffee break*

11:30 **Expert talk session 2 (15 min each)**

12:45 Discussion round (hot topics)

13:00 *Lunch + End*

14:00 *WOPSSS Workshop in the same room! Please stay.*

We stream the presentations on YouTube:

https://www.youtube.com/channel/UCcgmT-T19PS1pN_PETSb79Q

Sponsors of the Workshop

The workshop is powered by:

W·O·P·S·S·S
WORKSHOP ON PERFORMANCE AND SCALABILITY OF STORAGE SYSTEMS



The Virtual Institute for I/O

Goals of the Virtual Institute for I/O

- Provide a platform for I/O researchers and enthusiasts for exchanging information
- Foster training and international collaboration in the field of high-performance I/O
 - We support the community to establish conventions and standards
 - Example: We work on an **IO-500** benchmark!
- Track and encourage the deployment of large storage systems by hosting information about high-performance storage systems




<https://www.vi4io.org>

High-Performance Storage List

Features

- Model: Data centers, storage systems, supercomputers
- Tracks storage, compute, network, facility characteristics
- Community driven submission

storage capacity



#	Site		Supercomputer			Storage	
	Name	nationality	Name	compute_peak in PFLOPs	memory_capacity in TiB	Name	capacity ↑ in PiB
1	LANL	US	Trinity	11.08	1,919.03	Lustre	72.83
2	DKRZ	DE	Mistral	3.12	204.00	Lustre02 Lustre01 HPSS	52.00
3	LLNL	US	Sequoia	20.10	1,364.24	Grove	48.85
4	RIKEN	JP	K Computer	10.62	1,136.87	Lustre FEFS	39.77
5	NCAR	USA	Cheyenne	5.33	184.40	HPSS GPFS	37.00
6	NERSC	US	Cori Phase I	4.90	204.00	Lustre	30.00
7	ORNL	US	Titan	27.10	645.74	Spider 2	28.00
8	NCSA	US	Blue Waters	13.40	1,500.00	HPSS Lustre	26.40
9	JCAHPC	JP	Oakforest-PACS	24.91	836.09	Lustre Burst Buffer	24.10
10	CINECA	IT	Marconi A2 Fermi	12.93	413.97	GPFS GPFS	23.71
11	ANL	US	Mira	10.00	698.49	GPFS	21.32
12	JSC	DE	Juqueen	5.90	407.45	HPSS JUST	20.30
13	JAMSTEC	JP	Earth Simulator	1.31	291.04	Home Data Work Archive	19.62
14	KMA	KR	Miri	2.90	0.00	Lustre	19.27
15	NSCC	CN	TaihuLight	125.00	1,191.44	Sunway	17.76
16	AFRL	US	Thunder	5.61	406.54	Lustre	15.54
17	KAUST	SAU	Shaheen II	7.20	718.50	Lustre HPSS	15.28
18	LRZ	DE	SuperMUC Phase 2	3.58	176.44	GPFS	15.00
19	NASA	US	Pleiades	4.97	603.90	Lustre	14.21
20	NSCG	CN	Tianhe-2 Tianhe-1A	59.60	1,169.61	Tianhe-2 H2FS Tianhe-2 Lustre Lustre	14.18
21	TACC	US	Stampede	9.60	245.56	Lustre	12.43
22	ERDC DSRC	US	Topaz	4.57	401.63	Lustre	10.66
23	HLRS	DE	Hazel Hen	7.40	876.75	HPSS Lustre	8.88
24	TEP	FR	Pangea	6.71	49.11	Lustre	8.17
25	GSIC	JP	Tsukuba 2.5	5.76	67.67	Lustre	6.93
26	ENI	IT	HPC2	4.60	0.00	GPFS	6.66
27	PGS	US	Abel	5.37	531.14	Lustre	5.33
28	Nagoya University	JP	PRIMEHPC	3.20	83.67	Lustre	5.33
29	ECMWF	UK	Crazy XC40	4.25	0.00	HPSS Lustre	5.33
30	ARL	US	Excalibur	3.70	385.63	Lustre	4.09
31	EPCC	UK	Archer	2.55	0.00	Lustre	3.91
32	PNL	US	Cascade	3.40	167.35	Lustre	2.40
33	CSCS	CHE	Piz Daint	7.79	153.70	Lustre	2.22

Motivation for the Workshop

- I/O perspective of centers is often ignored
- Data centers aim to provide optimal service and performance

Providing a good storage strategy is challenging

- Though there are few HPC file systems: Lustre, GPFS, (BeeGFS)
 - Management of large volume/file numbers of data is difficult
 - Performance is often suboptimal: HDF5, NetCDF, small files
 - Shared storage and quality of service?
- Middleware to fix file system *issues* present in all file systems
 - PLFS, SIONlib, ADIOS, ...
 - Domain/Application-specific “solutions”, e.g. XIOS, CDI-PIO, ...
- Zoo of emerging storage approaches
 - Burst buffers, specialized storage for small files, ...
 - Alternative storage paradigms from BigData

Understanding Systems and Users

Knowing the behavior would allow to provide a better system

- A perfect understanding of usage and efficiency would allow for
 - selection of the right storage technology
 - gearing optimization effort towards mostly used I/O libraries
 - understanding the requirements for the procurement
 - optimizing the data center's efficiency as a whole
- But users often don't know their I/O patterns
- The I/O stack is challenging even for experts

Maybe I/O experts from data centers can make a difference

- From **individual** activity towards **community** effort and ultimately useful **conventions**

About the HPC-IODC Workshop

Goal: Bring together I/O experts from data centers

- Regardless of file system
- Foster information exchange
- Opportunity for networking

Topics of interest

- Scientific workload
- Usage characteristics (file, folders, scientific libraries)
- System perspective
- Architecture
- Performance aspects
- Monitoring
- Issues during production & potential solutions

Workshop results

- Presentations will be made available on our webpage
- Send the presentations ASAP to me, have to provide them to ISC folks on Friday!
- Research Papers are published in Springer LNCS
- We will write a preface and summarize the workshop results

Workshop Exascale I/O for Unstructured Grids

Topics

- Performance dealing with grids
- File formats
- Storage/data center perspective dealing with these types of data

Information

- When: 2017-09-25 / Sept. 25th + 26th
- Where: Hamburg, DKRZ
- <http://wr.informatik.uni-hamburg.de/events/2017/eiug>