

4th HPC I/O in the Data Center Workshop



Image under free license (CC0)



Limitless Storage
Limitless Possibilities

<https://hps.vi4io.org>

Julian M. Kunkel, Jay Lofstead

Sponsors of the Workshop



The workshop is powered by:



esiwace

CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER
AND CLIMATE IN EUROPE



The Virtual Institute for I/O

EU funded Project: ESiWACE



The Centre of Excellence in Simulation of Weather and Climate in Europe

- Representing the European community for
 - ▶ climate modelling and numerical weather simulation
- Goals in respect to HPC environments:
 - ▶ Improve efficiency and productivity
 - ▶ Supporting the end-to-end workflow of global Earth system modelling
 - ▶ Establish demonstrator simulations that run at highest affordable resolution
- Funding via the European Union's Horizon 2020 program (grant #675191)

<http://esiwace.eu>



esiwace

CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER
AND CLIMATE IN EUROPE

The Virtual Institute for I/O



Goals of the Virtual Institute for I/O

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

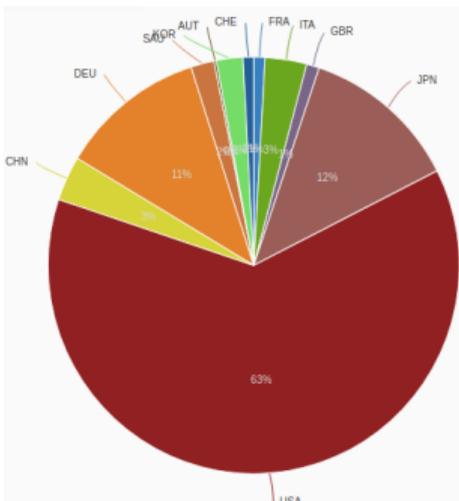


<https://www.vi4io.org>

CDCL Storage List 2018

Features

- Table view with selectable columns
- Flexible metrics and aggregation



Capacity grouped by country



2018

#	site.institution	site.storage.system.net	site.supercomputer.compute	site.supercomputer.memory
		capacity	peak	capacity
		in PB	in PFLOPS	in TB
1	Oak Ridge National Laboratory	250.04	220.64	3511.66
2	Los Alamos National Laboratory	72.83	11.06	2110.00
3	German Climate Computing Center	52.00	3.89	663.60
4	Lawrence Livermore National Laboratory	48.85	20.10	1500.00
5	RIKEN Advanced Institute for Computational Science	38.77	10.62	1250.00
6	National Center for Atmospheric Research	37.00	5.33	202.75
7	National Energy Research Scientific Computing Center	30.00	4.90	224.30
8	National Center for Supercomputing Applications	27.80	13.40	1649.27
9	Global Scientific Information and Computing Center	25.84	17.89	275.98
10	Joint Center for Advanced HPC	24.10	24.91	919.29
11	Cineca	23.71	12.93	455.17
12	Argonne National Laboratory	21.32	10.00	768.00
13	Forschungszentrum Jülich	20.30	6.25	454.15
14	Japan Agency for Marine-Earth Science and Technology	19.62	1.31	320.00
15	Korea Meteorological Administration	19.27	2.90	0.00
16	National Supercomputing Center in Wuxi	17.76	125.00	1310.00
17	Maryland Advanced Research Computing Center	17.00	0.87	82.67
18	King Abdullah University of Science and Technology	16.96	7.20	790.00
19	Air Force Research Laboratory	15.54	5.61	447.00
20	Leibniz Supercomputing Centre	15.00	3.56	194.00
21	National Supercomputing Center in Guangzhou	14.40	99.00	1286.00
22	National Aeronautics and Space Administration	14.21	4.97	664.00
23	Texas Advanced Computing Center	12.43	9.60	270.00
24	Engineer Research and Development Center - US Army Corps	10.66	4.57	441.60
25	Sandia National Laboratories	9.93	0.50	22.10
26	Karlsruhe Institute of Technology (KIT)	9.57	1.61	222.00
27	High-Performance Computing Centre Stuttgart	8.88	7.40	964.00
28	Total Exploration Production	8.17	6.71	54.00
29	Swiss National Supercomputing Centre	7.73	25.32	521.00
30	Eni S.p.A.	6.66	4.60	0.00
31	Nagoya University	5.33	3.20	82.00
32	PGS	5.33	5.37	584.00
33	European Centre for Medium-Range Weather Forecasts	5.33	4.25	0.00
34	Army Research Laboratory DxD Supercomputing Resource	4.99	3.70	424.00
35	University of Edinburgh	3.91	2.55	0.00
36	Pacific Northwest National Laboratory	2.40	3.40	184.00
37	Navy DoD Supercomputer Resource Center	2.11	2.05	0.00
38	Vienna Scientific Cluster	1.81	0.66	42.18
39	Center for Scientific Computing	0.75	0.51	77.57

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 and monitoring
- Issues during production and potential solutions



Workshop results



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

Morning Agenda: Shared Between WOPSSS and IODC



9:00 *Welcome*

9:15 **Keynote**

10:00 **Research paper session**

11:00 *Coffee break*

11:30 **Talks**

12:30 Discussion round (hot topic)

13:00 *Lunch + End*

14:00 *HPC-IODC Workshop and WOPSSS workshop in separate rooms*

Afternoon Agenda: HPC-IODC



14:00 *Welcome*

14:02 **Expert talk session (30 min each)**

15:00 **Research paper session (20 min each)**

16:00 *Coffee break*

16:30 **Expert talk session (30 min each)**

17:30 Discussion round (hot topic)

18:00 *Farewell*