

# The Virtual Institute of I/O



**Limitless** Storage  
**Limitless** Possibilities

<https://hps.vi4io.org>

Julian M. Kunkel

IO500 Release SC 2020

2020-11-18

# Introduction



## Goals of the Virtual Institute for I/O

- Provide a platform for I/O enthusiasts for exchanging information
  - ▶ The organization uses a wiki as central hub
  - ▶ Supported by mailing lists
- Foster training and collaboration in the field of high-performance I/O
  - ▶ Collaboration with the HPC Certification Forum ([hpc-certification.org](http://hpc-certification.org))
- Track and encourage the deployment of large storage systems
  - ▶ Hosting information about high-performance storage systems
  - ▶ Supporting the IO-500 and the Data Center List

<https://www.vi4io.org>



# Outline



- 1 Overview
- 2 Comprehensive Data Center List (CDCL)
- 3 Roadmap for 2021

# Comprehensive Data Center List (CDCL)



The CDCL contains system characteristics for sites, supercomputer and storage

## System Model

- Based on an extensible JSON schema, optimized editor
- Supports now (all) logical components and subcomponents
- Characteristics and peak values
- Measured values \*-500


## Components with characteristics

- Site, supercomputer, online storage, tape archives
- Compute nodes, storage nodes, local storage, accelerators, ...
- Supporting: e.g., CPU type, memory available, costs, ...

# CDCL Storage List 2020: 56 Sites

## Features


- Table view with selectable columns
- Flexible metrics and aggregation



Capacity grouped by country

#	site.institution	site.storage.system.net capacity in PIB	site.supercomputer.compute peak in PFLOPS	site.supercomputer.memory capacity in TB
1	Oak Ridge National Laboratory	250.04	220.64	3511.66
2	National Energy Research Scientific Computing Center	197.65	37.71	857.03
3	Los Alamos National Laboratory	72.83	11.08	2110.00
4	German Climate Computing Center	52.00	3.69	683.60
5	Lawrence Livermore National Laboratory	48.85	20.10	1500.00
6	RIKEN Advanced Institute for Computational Science	39.77	10.62	1250.00
7	National Center for Atmospheric Research	37.00	5.33	202.75
8	National Center for Supercomputing Applications	27.60	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

46	Center for Scientific Computing	0.75	0.51	77.57
47	University of Bristol	0.44	0.38	43.01
48	Joint Institute for Nuclear Research	0.00	0.00	0.00
49	NVIDIA	0.00	0.00	15.00
50	Howard Hughes Medical Institute/Janelia Research Campus	0.00	0.00	0.00



Capacity grouped by vendor

# Relevant Features



## Webpage integration

- Provide a HTML stub for embedding into any data center web page
  - ⇒ Toward a **standardized presentation** of systems!
    - ▶ Allowing the site to describe and visualize their system
    - ▶ Can be hosted by a site directly
    - ▶ Allowing a simple export into VI4IO data center list
- Editor: <https://www.vi4io.org/hpsl/addpage>

## Simple usage for VI4IO hosted information

- Create wiki stubs for a data center/storage system:  
<https://www.vi4io.org/cdcl-add>
- Edit the page, save

# Editor

Press save to store any changes

Save changes Reset Cancel

site		
abbreviation	nca	
institution	National Center for Sup	
location		
webpage	https://bluewaters.nca	
nationality	USA	
annual staff costs		k\$ ▾
energy costs per kWh		\$ ▾
energy cost		
[supercomputer+] [network+] [storage system+] [building+]		
supercomputer Blue Waters		> -
network Gemini		> -
network External		> -
storage system Lustre		> -
name	Lustre	
vendor	Cray	
model	Sonexion	
type	file system ▾	
software	Lustre	
version		
installation date		
net capacity	26,40	PIB ▾
peak write	1100,00	GIB/s ▾
peak read	1100,00	GIB/s ▾
servers		[recalc]
energy cost nodes tape archive		
nodes Lustre		> -
name	Lustre	
count		
operating system		
compute peak		GFLOPS ▾ [recalc]
installation date		
cooling		
processor details memory accelerator local storage interconnect		

- Flexible schema (updated as needed)
  - ▶ Recursive component model
- Users may add (optional) components
- Compute information
  - ▶ e.g., two types of nodes, compute total memory
- Keep "model" data in database (to be extended)
- May also draw relation between components



# Outline



- 1 Overview
- 2 Comprehensive Data Center List (CDCL)
- 3 Roadmap for 2021**



# The Comprehensive Data Center List & The IO500



- Support IO500 by extending the current practice
  - ▶ Users can submit CDCL schema for an IO500 submission defining their system
- Provide storage-system specific schemas
  - ▶ Use language common to the storage/file system
  - ▶ For Lustre, allow to input details for OSTs, MDTs, +configuration flags
- Define rules and automatically compute comparable quantities
  - ▶ e.g., capacity, storage interface
- Capture production-relevant settings as well
- The basic tool is there...
  - ▶ We need your/vendor help to create most-appropriate schema

# The Journal of High-Performance Storage



New journal, tested approach/tools successfully in 2020

- LaTeX and GoogleDoc workflows
- Goal 2021: Scale up, involve more editors

## Features

- Open reviews, i.e., anyone can provide feedback
- Living papers, i.e., can improve over time
- Digital replicability (of analysis/experiments)
- Free open access

The logo for the Journal of High-Performance Storage (J∞HPS). The letter 'J' is grey, the infinity symbol '∞' is red with a white outline, and the letters 'HPS' are red with a white outline.

<https://jhps.vi4io.org>

Several papers are available for open review in the incubator

# Other Activities



## Supporting IO Benchmarking

- Evolving the IO benchmarks in HPC/IOR further
  - ▶ E.g., added md-workbench lately
  - ▶ Documentation ...
  - ▶ Testing using the IO500
  - ▶ Deploying performance regression with Jenkins
- Supporting users to run/discuss IO benchmarks

## Next Generation IO Interfaces

- Goal: create a community forum to discuss next generation API
- Must revitalize past approach...

Everyone is welcome to participate