

Workshop on Storage Challenges in the UK

3rd SIG-IO-UK Workshop



Image under free license (CC0)



Limitless Storage
Limitless Possibilities

<https://hps.vi4io.org>

Julian M. Kunkel

2019-03-06

- 1 SIG IO UK
- 2 Workshop
- 3 Next Generation Interfaces: Status Update

The Special Interest Group I/O in the UK



Goals

- Strengthen networking the community in the UK:
 - ▶ System, middleware, ..., researchers and developers
 - ▶ Users with I/O-intensive workflows
 - ▶ Data centers
 - ▶ Storage and independent software vendors
- To bring forward RD&E:
 - ▶ Research
 - ▶ Product development
 - ▶ Funding opportunities

Communication

- Bi-annual meetings
- Mailing list for opportunities

Outline



- 1 SIG IO UK
- 2 Workshop
- 3 Next Generation Interfaces: Status Update

Sponsoring



The workshop is funded by
Data Direct Networks



The workshop is powered by



esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER
AND CLIMATE IN EUROPE



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

- Provide a platform for I/O enthusiasts for exchanging information
- Foster training and collaboration in the field of high-performance I/O
- Track and encourage the deployment of large storage systems by hosting information about high-performance storage systems

<https://www.vi4io.org>



Structure of Today's Agenda



10:00 4 talks

12:00 networking lunch with proper food thanks to DDN!

13:00 3 talks

14:30 Coffee break

15:00 4 talks

17:00 Adjourn

Next Generation Interfaces: Community Approach



Goal a new data-driven compute/storage and NGI

- Workflows and metadata as first-class citizens
- Storage and compute for heterogenous environments
- Smart software (and hardware) instead of manual
- Improving over time (self-aware/learning)
- Standardized interfaces beyond POSIX



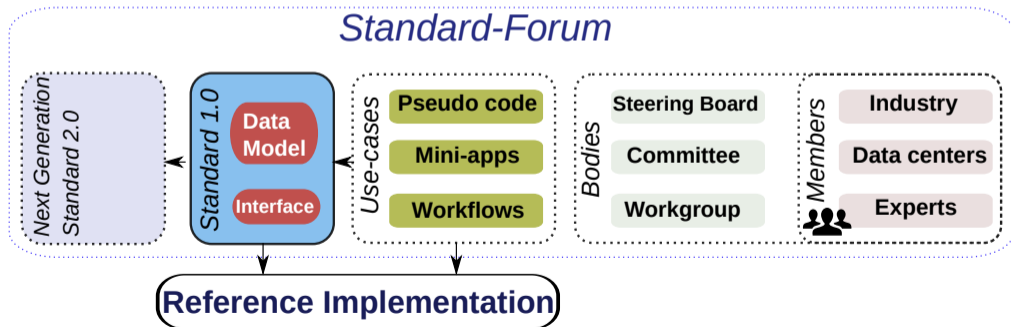
Why do we need a new domain-independent API?

- Other domains have similar issues: Harness RD&E effort across domains
- It is a hard problem approached by countless approaches
- Existing approaches address only a subset of the problems

Community-Driven Development of Data Model & API



- Establishing a Forum (similarly to MPI)
- Model targets High-Performance Computing and data-intensive compute
- Open board: encourage community collaboration



Project Plan



Approach

- Development of a Whitepaper
 - ▶ Limitations of the current state of practice
 - ▶ A vision for scientific computing in data centers in 2025+
 - ▶ Vision for next generation interfaces
 - ▶ Selected use-cases
 - ▶ Community approach
- Get more community on board

Status

- Submitted a BoF for ISC'19
- Coarse draft written by Julian in Dec/2018
 - ▶ Contributions from individuals
 - ▶ Have to refine the paper; likely to have different papers
 - Simplify issues for companies that (already) agreed to join consortium