Computer Science Department

HPS

University of Reading

The Virtual Institute of I/O

Limitless Storage
Limitless Possibilities

https://hps.vi4io.org

Julian M. Kunkel

Supercomputing 2018

2018-11-14

LIMITLESS **POTENTIAL** | LIMITLESS **OPPORTUNITIES** | LIMITLESS **IMPACT**

Outline

University of Reading

Introduction

University of Reading

Goals of the Virtual Institute for I/O

- 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

Introduction

University of Reading

Philosophical cornerstones of the institute

- Treat every member and participant equally
- Allow free participation without any membership fee inclusive to all
- Be independent of vendors and research facilities

Open Organization

University of Reading

- The organization uses a wiki as central hub
  - Everybody (registered users) can edit the content
  - Major changes should be discussed (see below)
  - The wiki uses tag clouds to link between similar entities
- Supported by mailing lists
  - Call-for-papers
  - Announce list for relevant information
  - Contribute list to discuss and steer organizational issues
- Major changes should be discussed on the contribute mailing list
- Members can vote for changes

Everybody is welcome to participate

Wiki Content

University of Reading

- Groups involved in high-performance storage
  Overview of research groups (evtl. companies involved in research)
  - Product development the group is involved in
  - Research projects (with links to their source)
  - Tags for layers, products and knowledge
- Tools: *Overview of relevant tools with small descriptions*
  - Types of tools: analysis, benchmarking, I/O middleware
  - Tags for layers and features
- Data Comprehensive Center List (CDCL) / High-Performance storage list
  Characteristics of data center systems
  - Editable and owned by the community
- Internal section
  Provides templates and describes rules for editing the page

Comprehensive Data Center List (CDCL)

The CDCL contains system characteristics for sites, supercomputer and storage

System Model

- The system model has been refined since ISC
  - Now 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, ...

CDCL Storage List 2018

Features

- Table view with selectable columns
- Flexible metrics and aggregation

#	site.institution	site.storage system.net capacity	site.supercomputer.compute peak	site.supercomputer.memory capacity
		in PiB	*in PFLOPS*	*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
37	Pacific Northwest National Laboratory	2.40	3.40	184.00
38	Navy DoD Supercomputer Resource Center	2.11	2.05	0.00
39	Vienna Scientific Cluster	1.81	0.68	42.18
40	Center for Scientific Computing	0.75	0.51	77.57
41	University of Bristol	0.44	0.38	43.01

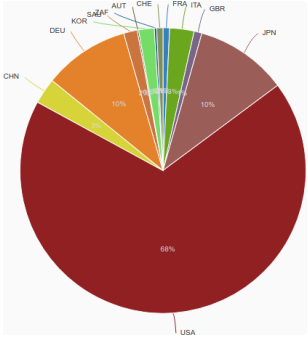

Capacity grouped by country

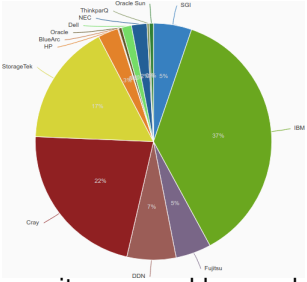

Capacity grouped by vendor

Status

University of Reading

CDCL: What's New

- Form to create stubs for pages: `https://www.vi4io.org/cdcl-add`
  - This simplifies the addition of new systems significantly
- New sites:
  - National Energy Research Scientific Computing Center (#2): **Comprehensive**
  - ORNL, University of Bristol, Cambridge, Queen Mary U of London

Started: Web page data center representation

- Provide a Javascript for embedding into any data center web page
  - ⇒ Towards a **standardized presentation** of systems!
  - Allowing the site to describe and visualize their system
  - Hosted by the site directly
  - Allowing a simple export into VI4IO data center list
- Polish presentation of site's information

Roadmap for 2019

University of Reading

Supported community activities

- Official roadmaps for the community benchmarks IOR/MDtest
  - Various contributions to IOR/MDtest happened
- Achieve the standardized presentation of systems
  - Help is welcome: Requirements? Testers?
- Link IO-500 results with CDCL
- New training page linking resources for learning high-performance storage
- Supporting a forum for **Next-Generation Storage Interfaces**

Community Development of Next Generation Interfaces

University of Reading

Towards developing a new I/O stack API considering:

- Cover storage and data-flow computation together
- Utilizing heterogeneous storage and compute landscapes
  - Scheduler optimizes plans; beyond tiering; liquid computing
- Smart hardware and software components
  - Self-aware system instead of unconscious
  - Improving over time (self-learning, hardware upgrades)
- User metadata, ILM, and workflows as first-class citizens

Indeed many research prototypes address subproblems

- But not all aspects together
- Competing approaches; the standardization does not compete!

Pursued Community Strategy

University of Reading

- The **standardization** of a high-level *data model & interface*
  - Lifting semantic access to a new level (e.g. NetCDF + X)
  - Targeting data intensive and HPC workloads
  - To have a future: must be beneficial for Cloud/Big Data + Desktop, too
- Supporting a reference implementation of a **smart runtime system**
  - Implementing key features
  - Build on top of the available solutions!
- Demonstration of benefits on socially relevant data-intense apps

Development of the Data Model and API

University of Reading

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

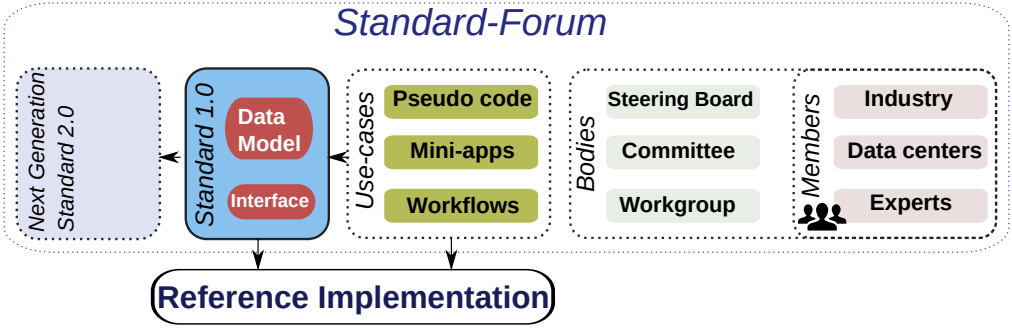

Summary

University of Reading

- The Virtual Institute for I/O is a community hub
  - Open to everybody and free to join
- It contains information about
  - Tools, benchmarks
  - Research groups
  - Standardization efforts
- It hosts the Comprehensive Data Center List (CDCL)
  - Covers many metrics and allows flexible visualization
  - Will track metrics across years
  - Can be updated by members
- Towards **standardization** of system presentation and APIs
  - We are preparing a white-paper for NGI; join the mailinglist now to participate!
- ***We need you to participate!***