Integrating HPC Resources and Other Scientific Instruments with Workflows

Sean Wilkinson

wilkinsonsr@ornl.gov

Workflow and Ecosystem Services Group

Oak Ridge Leadership Computing Facility

ORNL is managed by UT-Battelle LLC for the US Department of Energy

Programmatic HPC Access

- How to treat HPC resources as scientific instruments?
  - Efforts like DOE’s IRI and and ORNL’s INTERSECT
- Intriguing question: “**Programmatic HPC Access *for whom*?**”
  - Traditionally, humans
  - Increasingly, humans and machines together (e.g. GitHub Copilot)
  - Eventually, machines → autonomous workflows
- How do we get there?

What are computational workflows?

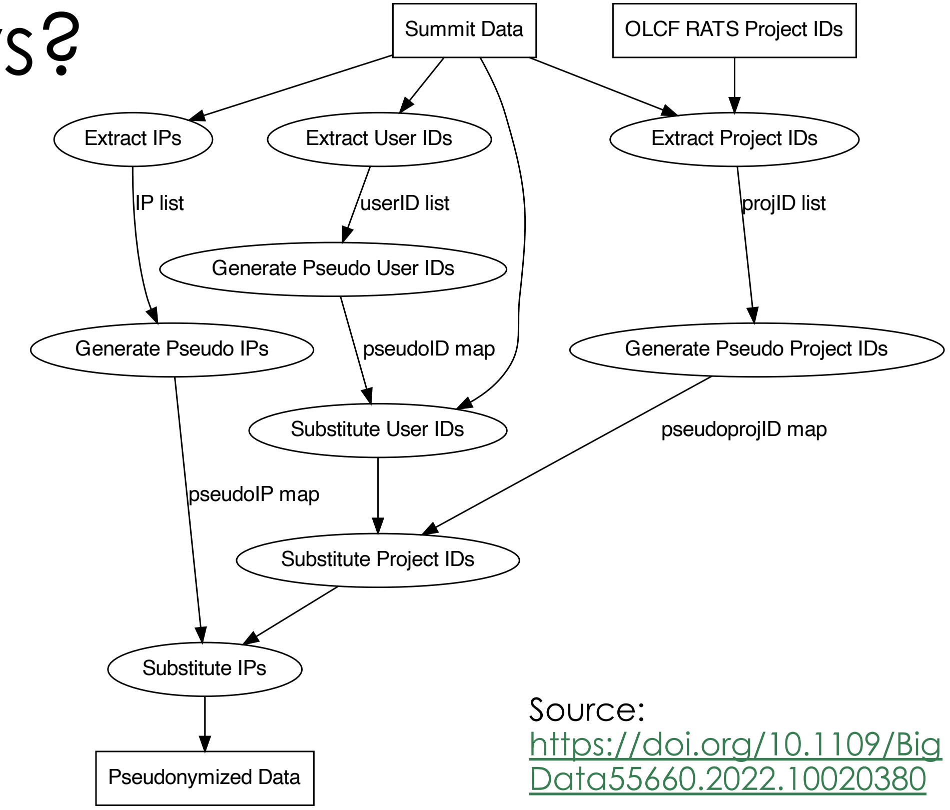

Source: https://doi.org/10.1109/BigData55660.2022.10020380

- This is a deceptively hard question!
- Different from "business workflows"
- A few choices:
  - Programs of programs
  - Directed acyclic graphs (DAGs) of tasks
  - Compositions of transformations for data
  - Programs that use workflow management systems (!)
- Today's definition: **"multi-step computational processes for working with data".**

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

Easier question: *Why* computational workflows?

Manage complexity

- Automation **avoids human errors**
- Workflows enable you to use the **right tool for each task** instead of writing one big, ugly program
- Workflows enable automatic **expert optimizations**
  - Allocating and scheduling resources
  - Error handling and task restarting

Improve scientific practice

- Fully automating an experiment enables
  - you to **repeat** the experiment easily with different or same conditions
  - others to **reproduce** your results from the exact same programs
  - others to **reuse** part or all of your computational methods for their own experiments so they don't have to re-invent the wheel

Toward autonomous workflows…

- Workflows are about automation.
- Autonomous workflows are a higher-level automation.
- Must complexity increase?
- **The same things that enhance human-machine interactions also enable machine-machine interactions.**

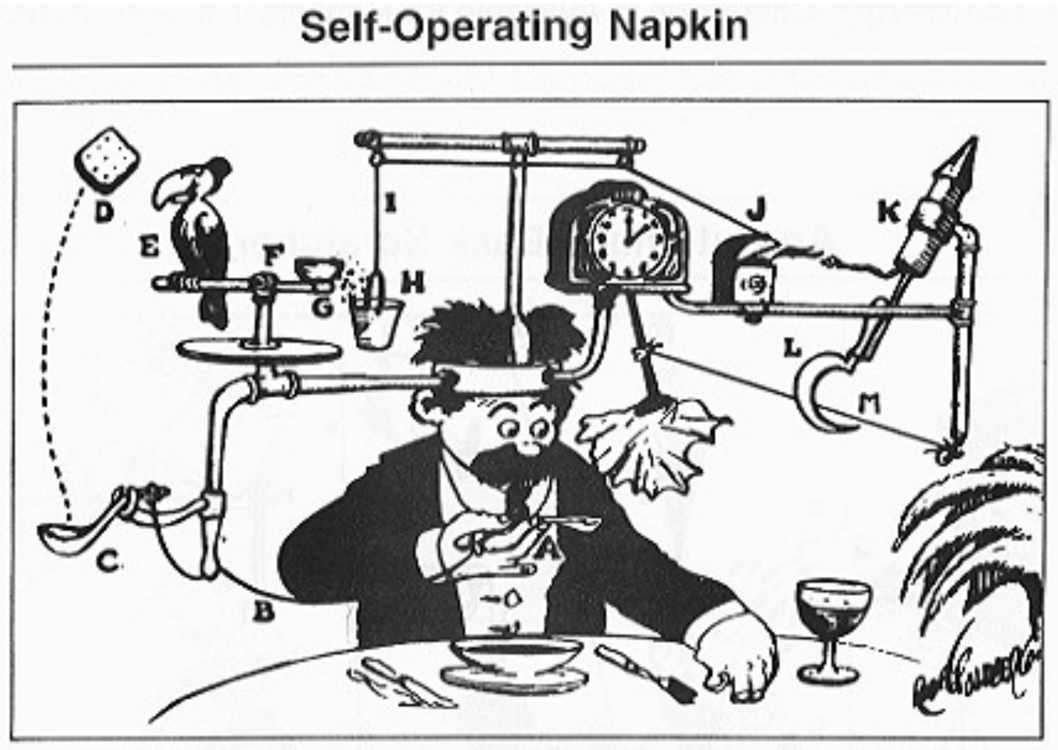

Source: https://w.wiki/9PX6

Surprise! It’s the FAIR Principles!

> FAIR data are essential to the future of human–machine collaboration and autonomous machine-to-machine communication.
>
> - Sansone *et al.* 2023

Source: https://doi.org/10.1142/9789811265679_0033

The FAIR Guiding Principles for Scientific Data

Published in 2016 by Mark Wilkinson *et al.* Listed 4 main emphases for scientific data management and stewardship:

- Findable (4 principles)
- Accessible (2 princ + 2 sub)
- Interoperable (3 principles)
- Reusable (1 princ + 3 sub)

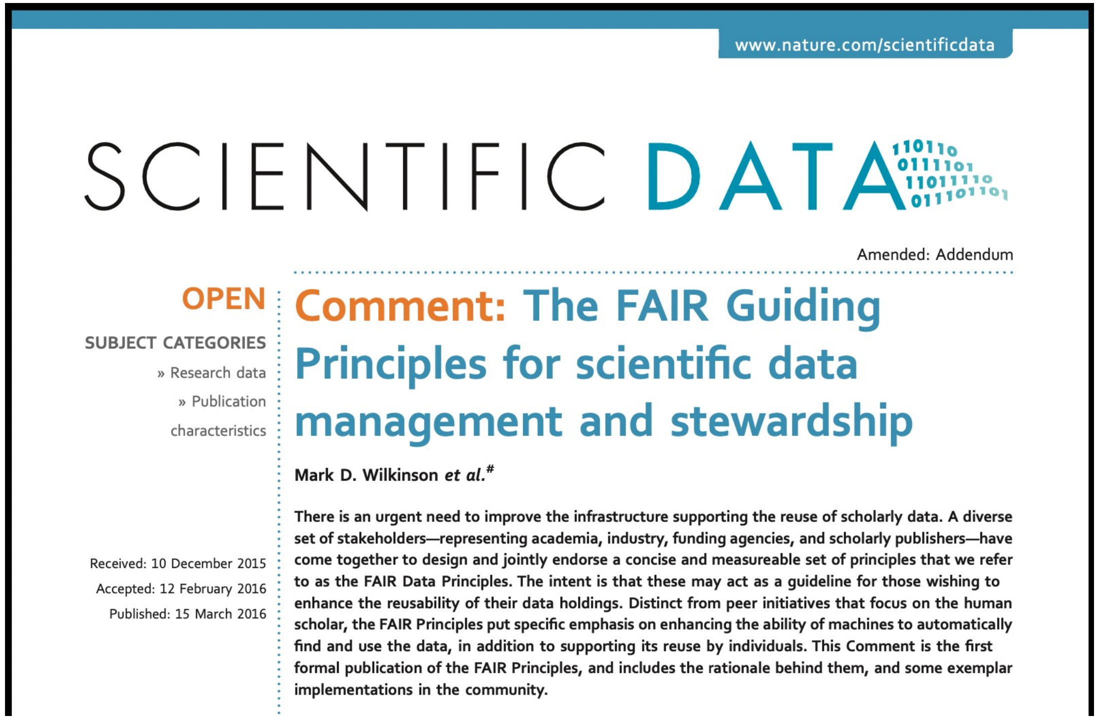
www.nature.com/scientificdata

SCIENTIFIC DATA

Amended: Addendum

OPEN

SUBJECT CATEGORIES
» Research data
» Publication characteristics

Received: 10 December 2015
Accepted: 12 February 2016
Published: 15 March 2016

Comment: The FAIR Guiding Principles for scientific data management and stewardship

Mark D. Wilkinson *et al.*#

There is an urgent need to improve the infrastructure supporting the reuse of scholarly data. A diverse set of stakeholders—representing academia, industry, funding agencies, and scholarly publishers—have come together to design and jointly endorse a concise and measureable set of principles that we refer to as the FAIR Data Principles. The intent is that these may act as a guideline for those wishing to enhance the reusability of their data holdings. Distinct from peer initiatives that focus on the human scholar, the FAIR Principles put specific emphasis on enhancing the ability of machines to automatically find and use the data, in addition to supporting its reuse by individuals. This Comment is the first formal publication of the FAIR Principles, and includes the rationale behind them, and some exemplar implementations in the community.

The FAIR Principles strongly emphasize metadata to enable machine actionability.

Directly from the abstract of the original FAIR paper:

> Distinct from peer initiatives that focus on the human scholar, the FAIR Principles put specific emphasis on enhancing the ability of machines to automatically find and use the data, in addition to supporting its reuse by individuals.
>
> – Wilkinson *et al.* 2016

Source: https://doi.org/10.1038/sdata.2016.18

The FAIR Principles for Research Software (FAIR4RS)

Published in 2022 by Michelle Barker *et al.*

The FAIR Principles are applied slightly differently because software is

- executable and
- made of components.

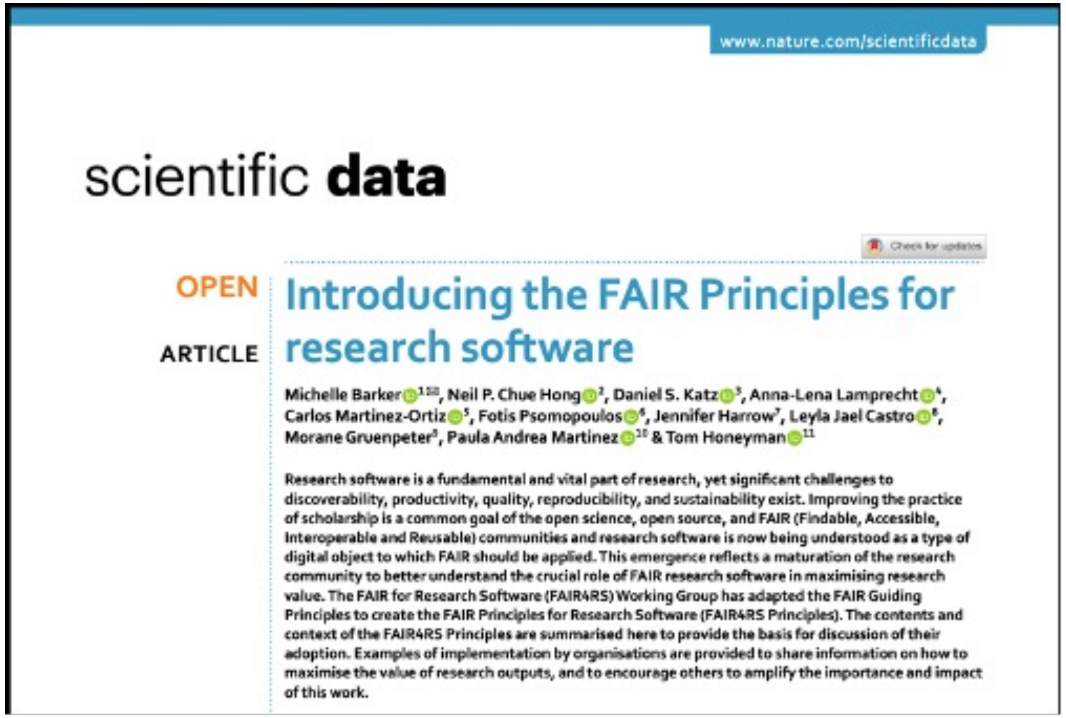
www.nature.com/scientificdata

scientific **data**

Check for updates

OPEN ARTICLE

Introducing the FAIR Principles for research software

Michelle Barker[1], Neil P. Chue Hong[2], Daniel S. Katz[3], Anna-Lena Lamprecht[4], Carlos Martinez-Ortiz[5], Fotis Psomopoulos[6], Jennifer Harrow[7], Leyla Jael Castro[8], Morane Gruenpeter[9], Paula Andrea Martinez[10] & Tom Honeyman[11]

Research software is a fundamental and vital part of research, yet significant challenges to discoverability, productivity, quality, reproducibility, and sustainability exist. Improving the practice of scholarship is a common goal of the open science, open source, and FAIR (Findable, Accessible, Interoperable and Reusable) communities and research software is now being understood as a type of digital object to which FAIR should be applied. This emergence reflects a maturation of the research community to better understand the crucial role of FAIR research software in maximising research value. The FAIR for Research Software (FAIR4RS) Working Group has adapted the FAIR Guiding Principles to create the FAIR Principles for Research Software (FAIR4RS Principles). The contents and context of the FAIR4RS Principles are summarised here to provide the basis for discussion of their adoption. Examples of implementation by organisations are provided to share information on how to maximise the value of research outputs, and to encourage others to amplify the importance and impact of this work.

FAIR4RS still emphasizes metadata strongly for the same reason – machine actionability.

The FAIR Principles for Computational Workflows

- To be published in 2024, hopefully, by members of the WCI FAIR Workflows working group. (Join us!)
- FAIR applies a little differently to workflows because they have attributes of **both data and software**.
  - FAIR can be applied to their components recursively.

FAIR Principles for workflows will definitely emphasize metadata!

So, what is machine actionability?

The capability of computational systems

- "to use services on data without human intervention"
  - https://doi.org/10.5334/dsj-2020-015

- "to find, access, interoperate, and reuse data with none or minimal human intervention"
  - https://www.go-fair.org/fair-principles/

So, what is machine actionability?

The capability of computational systems

- "**to use** services on data without human intervention"
  - https://doi.org/10.5334/dsj-2020-015

- "to find, access, interoperate, **and reuse data with** none or **minimal human intervention**"
  - https://www.go-fair.org/fair-principles/

Combined definition

Machine actionability (n.) –

The capability of computational systems to use and reuse data with minimal human intervention.

Yes, "use" and "reuse" are related, but in practice, they are not the same!

Combined definition

Machine actionability (n.) –

The capability of **computational systems** to use and reuse **data** with minimal **human intervention**.

So, the FAIR Principles aren't just for (meta)data nerds – they're for all of us.

FAIR applies at many levels of the ecosystem:

- Data – Wilkinson *et al.* 2016
- Software – Barker *et al.* 2022
- Hardware – https://zenodo.org/records/6506428 (2021-2022)
- AI / ML – https://doi.org/10.1038/s41597-023-02298-6 (2023)
- Workflows – in preparation by WCI working group (join us!)
- Facilities and instruments – see next slide!

Coming soon: FAIR Facilities and Instruments

Report just published in February 2024 for a workshop held in Boulder, Colorado in September 2023.

The focus is on the assignment of persistent identifiers (PIDs) to research facilities and instrumentation.

(Metadata again…)


NCAR | National Center for
UCAR | Atmospheric Research

FAIR Facilities and Instruments Workshop #1 Report: Exploring Persistent Identifier Needs, Barriers and Incentives

Andrew Johnson
Renaine Julian
Matt Mayernik
Claudius Mundoma
Matthew Murray
Aditya Ranganath
Greg Stossmeister

https://doi.org/10.5065/zgsx-2d06

Machine actionability strongly depends on metadata!

Humans and machines are

- similar because
  - without proper metadata, they're just guessing.
  - with proper metadata, they can make informed decisions.
- different because
  - machines don't complain that recording metadata is tedious.

Terminal — 67×30

~ — vi metadata.json


```
{
  "@context": "https://w3id.org/ro/crate/1.1/context",
  "@graph": [
    {
      "@id": "https://example.com/research-project",
      "@type": "Dataset",
      "name": "Multi-omics Study on Cancer Subtypes",
      "description": "This RO-Crate describes a multi-omics study a
      "datePublished": "2023-10-01",
      "dateModified": "2024-02-15",
      "creator": [
        {
          "@id": "https://orcid.org/0000-0001-9876-5432",
          "name": "Dr. Alice Johnson",
          "@type": "Person"
        },
        {
          "@id": "https://orcid.org/0000-0002-1234-5678",
          "name": "Dr. Robert Smith",
          "@type": "Person"
        }
      ],
      "contributor": [
        {
          "@id": "https://orcid.org/0000-0003-4567-8901",
          "name": "Dr. Emily White",
          "@type": "Person"
        },
        {
:
```

Examples from OLCF's IRI Effort

INTEGRATED RESEARCH INFRASTRUCTURE

ACE TESTBED

- Sandbox for diverse computing and data resource evaluation
- Drives HPC technology development for OLCF/DOE
- Open, flexible environment for testing emerging technologies

FLEXIBLE QUEUEING

- Preemptable workflow testing on ACE Testbed
- Prioritize IRI jobs for faster initiation

FEDERATED IDENTITY

Exploring IAM solution used in DOE headquarters and labs (OneID and PingFed)

DATA STREAMING TO HPC

OLCF system for bidirectional data streaming with external experiments

CONTAINERIZATION

Container solutions for HPC applications and workflows

INTEGRATION WITH INTERSECT

Enabling INTERSECT use cases on ACE Testbed

FACILITY API

Shared API library for accelerated development and uniform functionality across facility services

**Collaboration with NERSC and ALCF

MULTI-TENANCY

Implement multi-tenancy in HPC for efficient IRI use case support

PATHFINDER PROJECTS

ESGF, JGI, LCLS, GRETA, CNMS

Slide credit: Rafael Ferreira da Silva

Testbed Network Connections

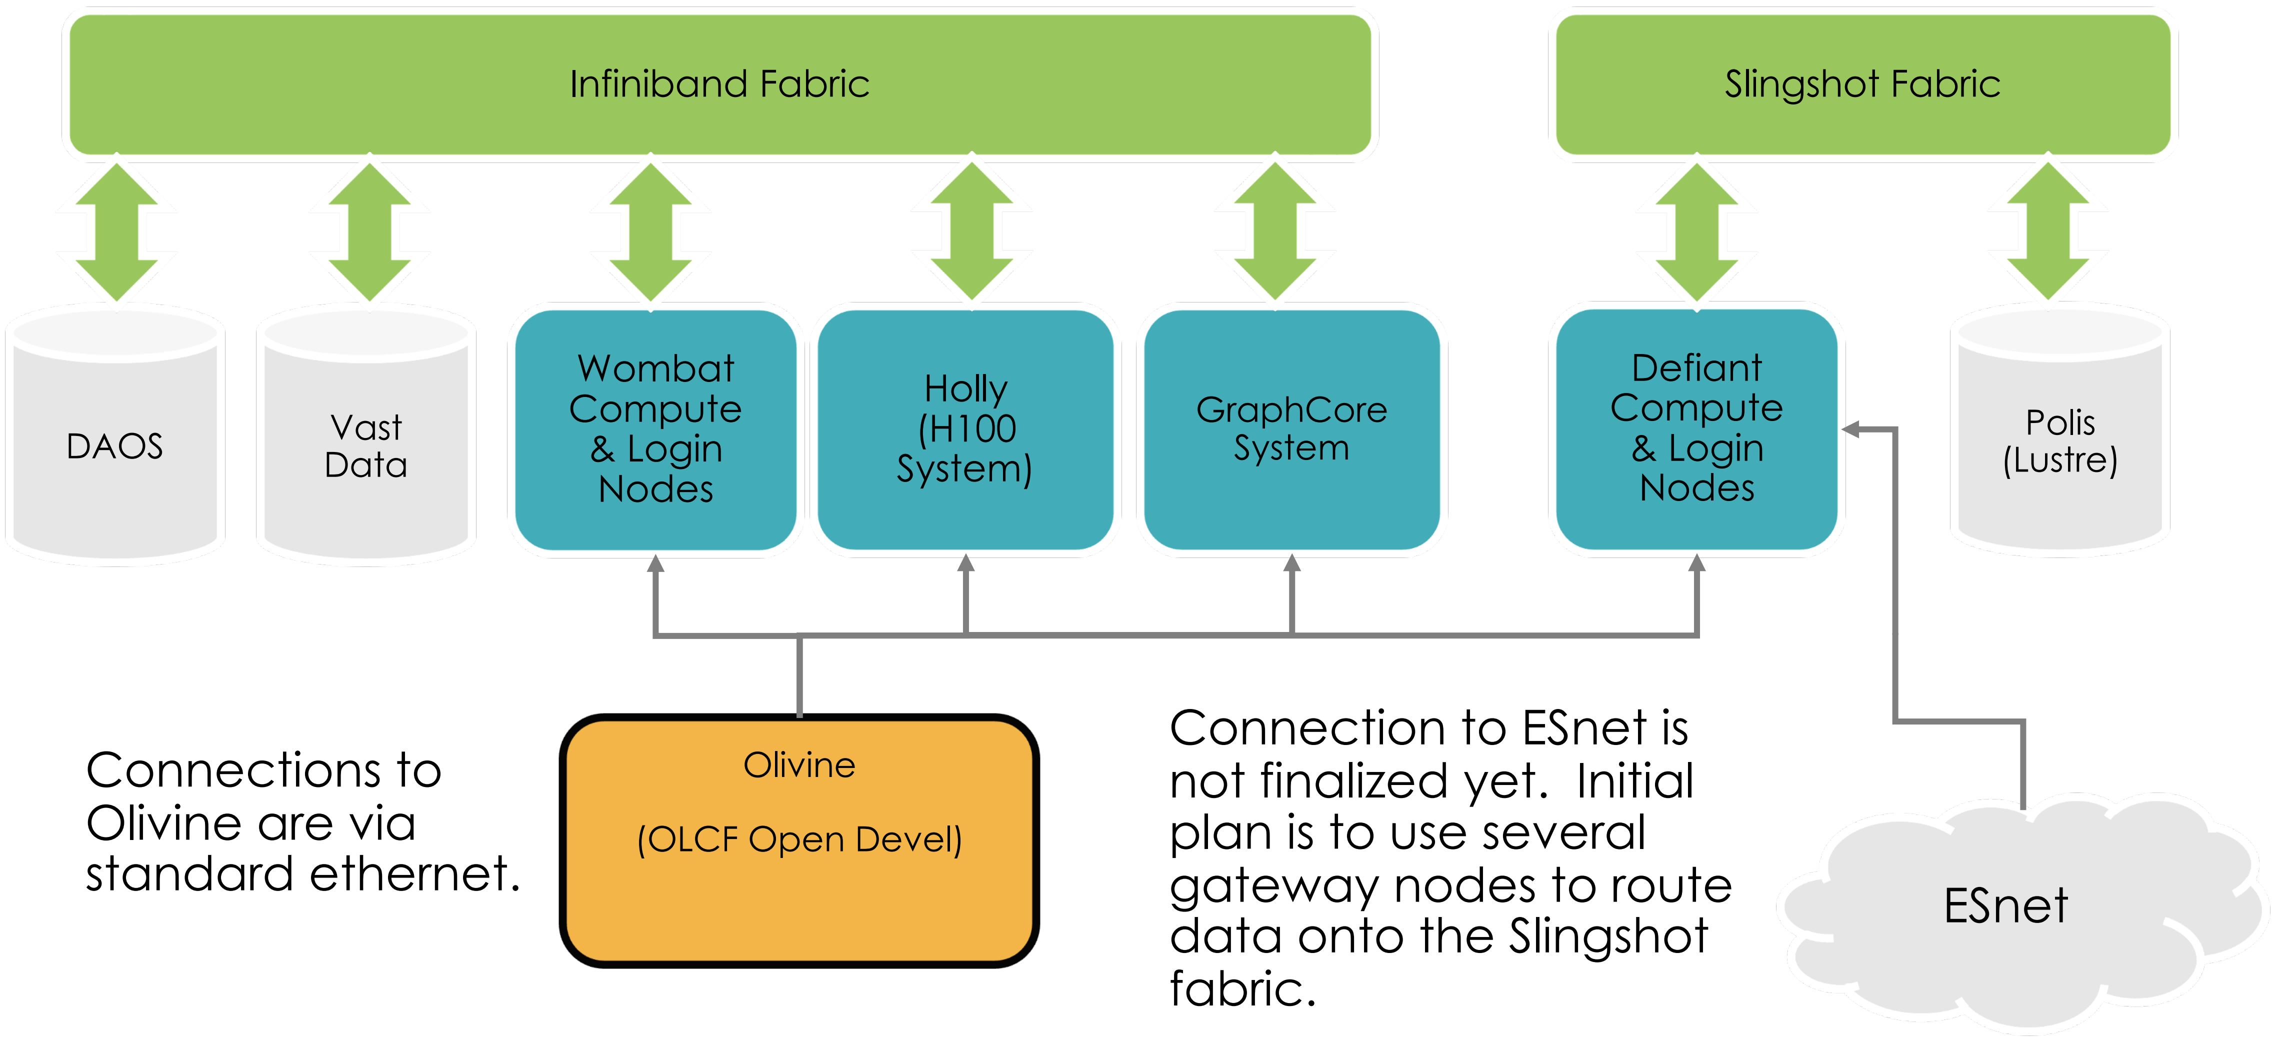

Connections to Olivine are via standard ethernet.

Connection to ESnet is not finalized yet. Initial plan is to use several gateway nodes to route data onto the Slingshot fabric.

Slide credit: Ross Miller

ACE/IRI Testbed Resources (as of January 2024)

- Compute:
  - Defiant: 36 nodes. AMD EPYC + 4 MI100 GPUs per node. Slingshot networking
  - Wombat: 14 nodes. Various AArch64 + GPU configurations. IB networking
    - 8 new Grace/Hopper nodes due to arrive this month
  - Graphcore: Specialized AI appliance from GraphCore. 16 "IPUs"
  - Holly: 8 H100 GPUs in a single server.
  - **Coming soon**: Kubernetes cluster (named Olivine), 16 x86 compute nodes, direct connection to ESnet, NDR IB
- Storage:
  - 8 servers for testing DAOS. ~30TB flash per server
    - 4 servers installed and working. 4 more racked and awaiting software install.
    - On the IB fabric. Will be usable by Wombat, Holly & Graphcore
  - Storage appliance from VastData. ~600TB usable space. All flash.
    - On the IB fabric. Mounted by Wombat, Holly & Graphcore
  - Lustre filesystem (Polis). HPE ClusterStor. ~1.2PB usable space. Mix of flash and spinning disk
    - On the slingshot fabric. Mounted by Defiant.

Slide credit: Ross Miller

Examples from OLCF's IRI Effort

ACE TESTBED

- Sandbox for diverse computing and data resource evaluation
- Drives HPC technology development for OLCF/DOE
- Open, flexible environment for testing emerging technologies

FLEXIBLE QUEUEING

- Preemptable workflow testing on ACE Testbed
- Prioritize IRI jobs for faster initiation

FEDERATED IDENTITY

Exploring IAM solution used in DOE headquarters and labs (OneID and PingFed)

DATA STREAMING TO HPC

OLCF system for bidirectional data streaming with external experiments

CONTAINERIZATION

Container solutions for HPC applications and workflows

INTEGRATION WITH INTERSECT

Enabling INTERSECT use cases on ACE Testbed

FACILITY API

Shared API library for accelerated development and uniform functionality across facility services

**Collaboration with NERSC and ALCF

MULTI-TENANCY

Implement multi-tenancy in HPC for efficient IRI use case support

PATHFINDER PROJECTS

ESGF, JGI, LCLS, GRETA, CNMS

Slide credit: Rafael Ferreira da Silva

Extending Resource Interactions Through Facility APIs

Why?

- APIs, or application program interfaces, are vital in enabling software-to-software interaction
- Underlying systems and decades of HPC experience are still crucial - the API is another avenue of interaction for users
- APIs allow facilities to expose resources to externally developed applications and data portals through a more robust interface
- Facility APIs will be foundational to the advancement of orchestrating and automating workflows across laboratories

Benefit

- Higher-level portals will be able to more flexibly interact with facility provided systems
- The Facility API will provide a streamlined development experience to users who desire to develop and provide orchestration and workflow systems to their community
- The Facility API will enable a clean software interface for new types of integrations and enable users to drive workflows in advanced ways
- Focuses on providing a central service that is capable of enforcing policy, logging, authenticating, authorizing, and securing all incoming API requests to a range of services.

OLCF TEAM:

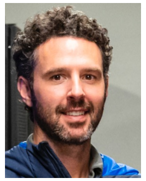

Ryan Prout

Paul Bryant

A.J. Ruckman

Tyler Skluzacek

Rafael Ferreira da Silva

Lawrence Sorrillo

Slide credit: Ryan Prout

Facility API – Streamlining Integrations

Slide credit: Ryan Prout

Earth Systems Grid Federation

Atmospheric Radiation Measurement Facility

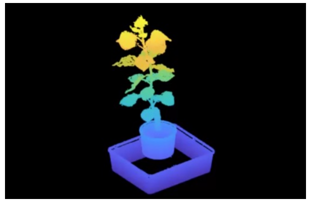

Advanced Plant Phenotyping Laboratory

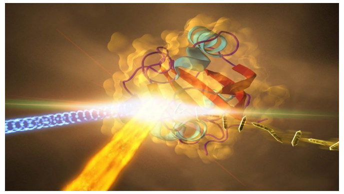

Linac Coherent Light Source-II

Gamma-Ray Energy Tracking Array

Users

External

OLCF

- Data Services
- Workflow and Data Streaming Orchestration
- Software Deployment and Integration
- Facility-provided portals

Facility API

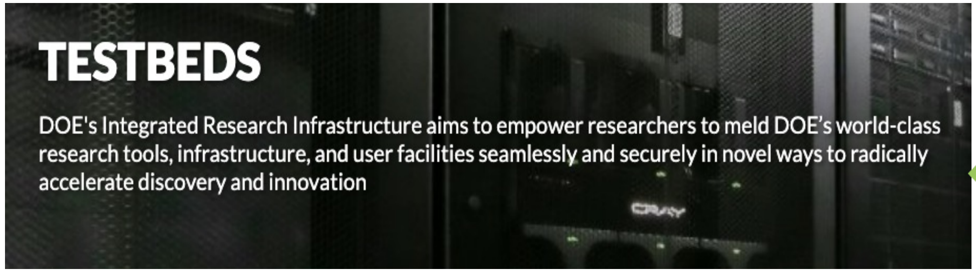

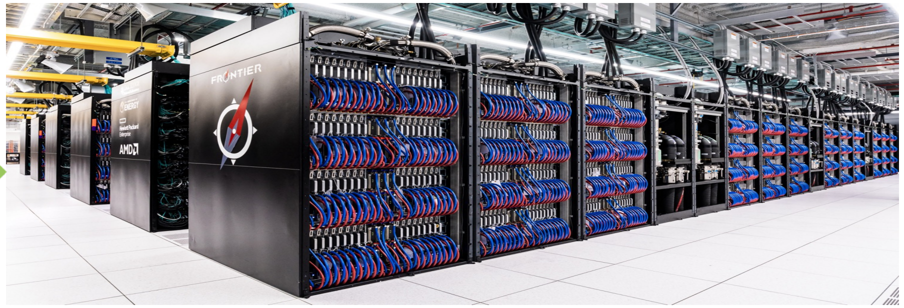

Examples from OLCF's IRI Effort

INTEGRATED RESEARCH INFRASTRUCTURE

ACE TESTBED

- Sandbox for diverse computing and data resource evaluation
- Drives HPC technology development for OLCF/DOE
- Open, flexible environment for testing emerging technologies

DATA STREAMING TO HPC

OLCF system for bidirectional data streaming with external experiments

FACILITY API

Shared API library for accelerated development and uniform functionality across facility services

**Collaboration with NERSC and ALCF

MULTI-TENANCY

Implement multi-tenancy in HPC for efficient IRI use case support

FLEXIBLE QUEUEING

- Preemptable workflow testing on ACE Testbed
- Prioritize IRI jobs for faster initiation

FEDERATED IDENTITY

Exploring IAM solution used in DOE headquarters and labs (OneID and PingFed)

CONTAINERIZATION

Container solutions for HPC applications and workflows

INTEGRATION WITH INTERSECT

Enabling INTERSECT use cases on ACE Testbed

PATHFINDER PROJECTS

ESGF, JGI, LCLS, GRETA, CNMS

Slide credit: Rafael Ferreira da Silva

Data streaming to HPC ("memory to memory")

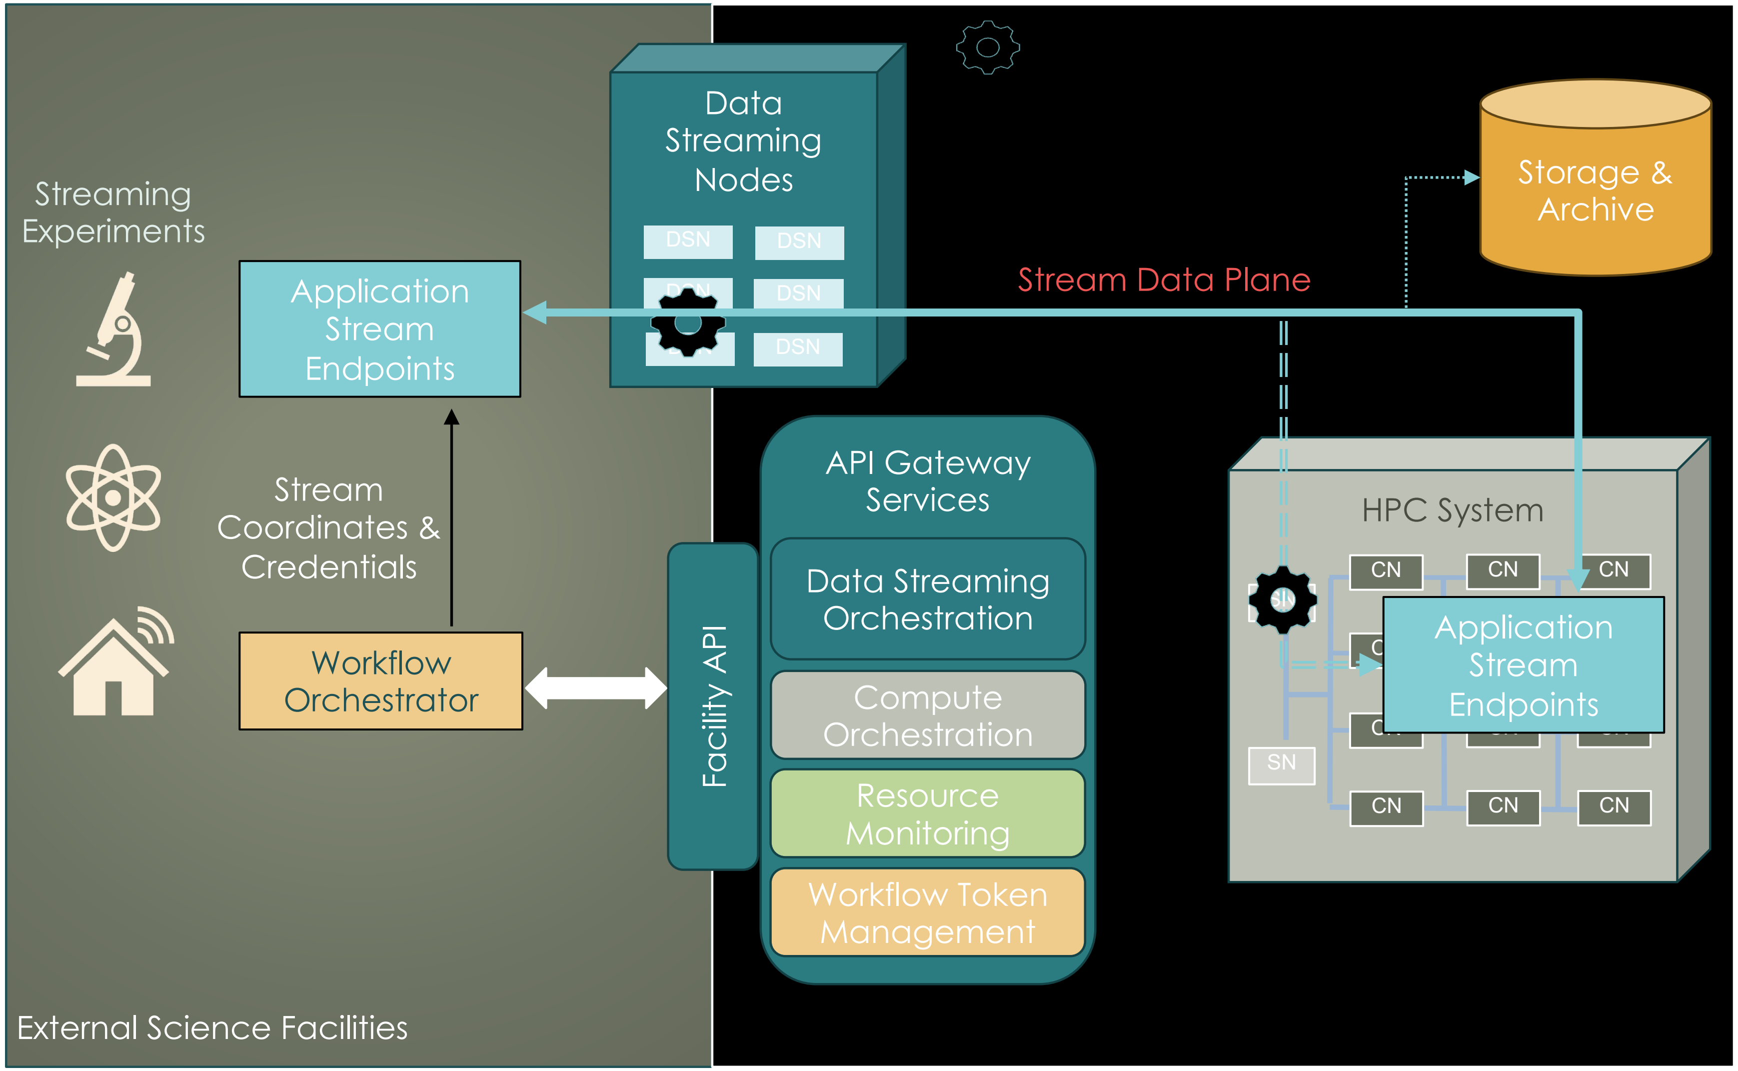

Slide credit: Mike Brim

Workflow orchestrator is still under construction…

Source:
https://w.wiki/9QdL

WorkflowHub: a FAIR workflow registry

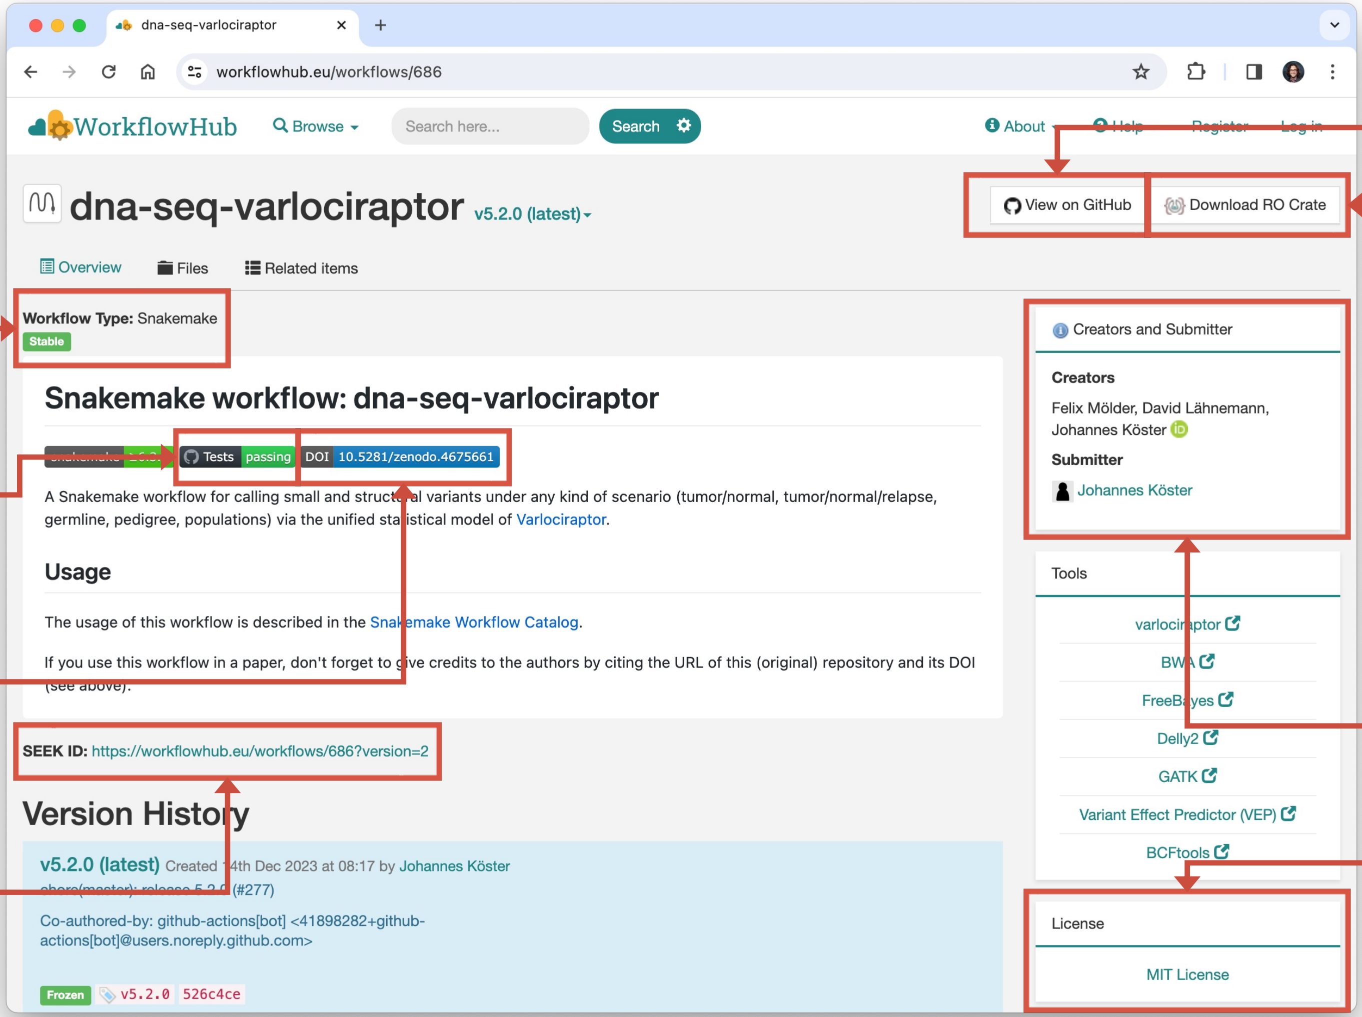

Workflow management system

Unit test badge

DOI to permanent repository

SEEK ID for use with SEEK API

Source code in version-controlled repository

RO-Crate download ("everything")

Contact information

License information

Summary

- The FAIR Principles are not just for (meta)data.
- To enable full programmatic HPC access for current humans and future machines, we need machine actionability at the
  - data level,
  - software level,
  - hardware level,
  - AI/ML level,
  - workflow level, and
  - facility and instrumentation level.

Acknowledgments

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

The End

- To contact me
  - wilkinsonsr@ornl.gov
- To join the WCI FAIR Computational Workflows working group
  - https://workflows.community/groups/fair/
- Thank you!