

Distributed and Asynchronous Management of Science Workflows: Data and Computation Across DOE User Facilities

Sean R. Wilkinson (wilkinsonsr@ornl.gov), Joshua S. Brown, Ketan C. Maheshwari, Olga A. Kuchar, Rafael Ferreira da Silva
Data Lifecycle and Scalable Workflows Group, Oak Ridge Leadership Computing Facility, National Center for Computational Sciences, Oak Ridge National Laboratory

DOE User Facilities as an Ecosystem

- The United States Department of Energy sponsors 28 user facilities: 4 in ASCR, 12 in BES, 3 in BER, 2 in FES, 3 in HEP, and 4 in NP.
- These facilities should share resources for data and computation to accelerate scientific research, but each facility is different [1].
- Ideally, DOE's domain science facilities would connect to its data and computing facilities seamlessly as a single "ecosystem".

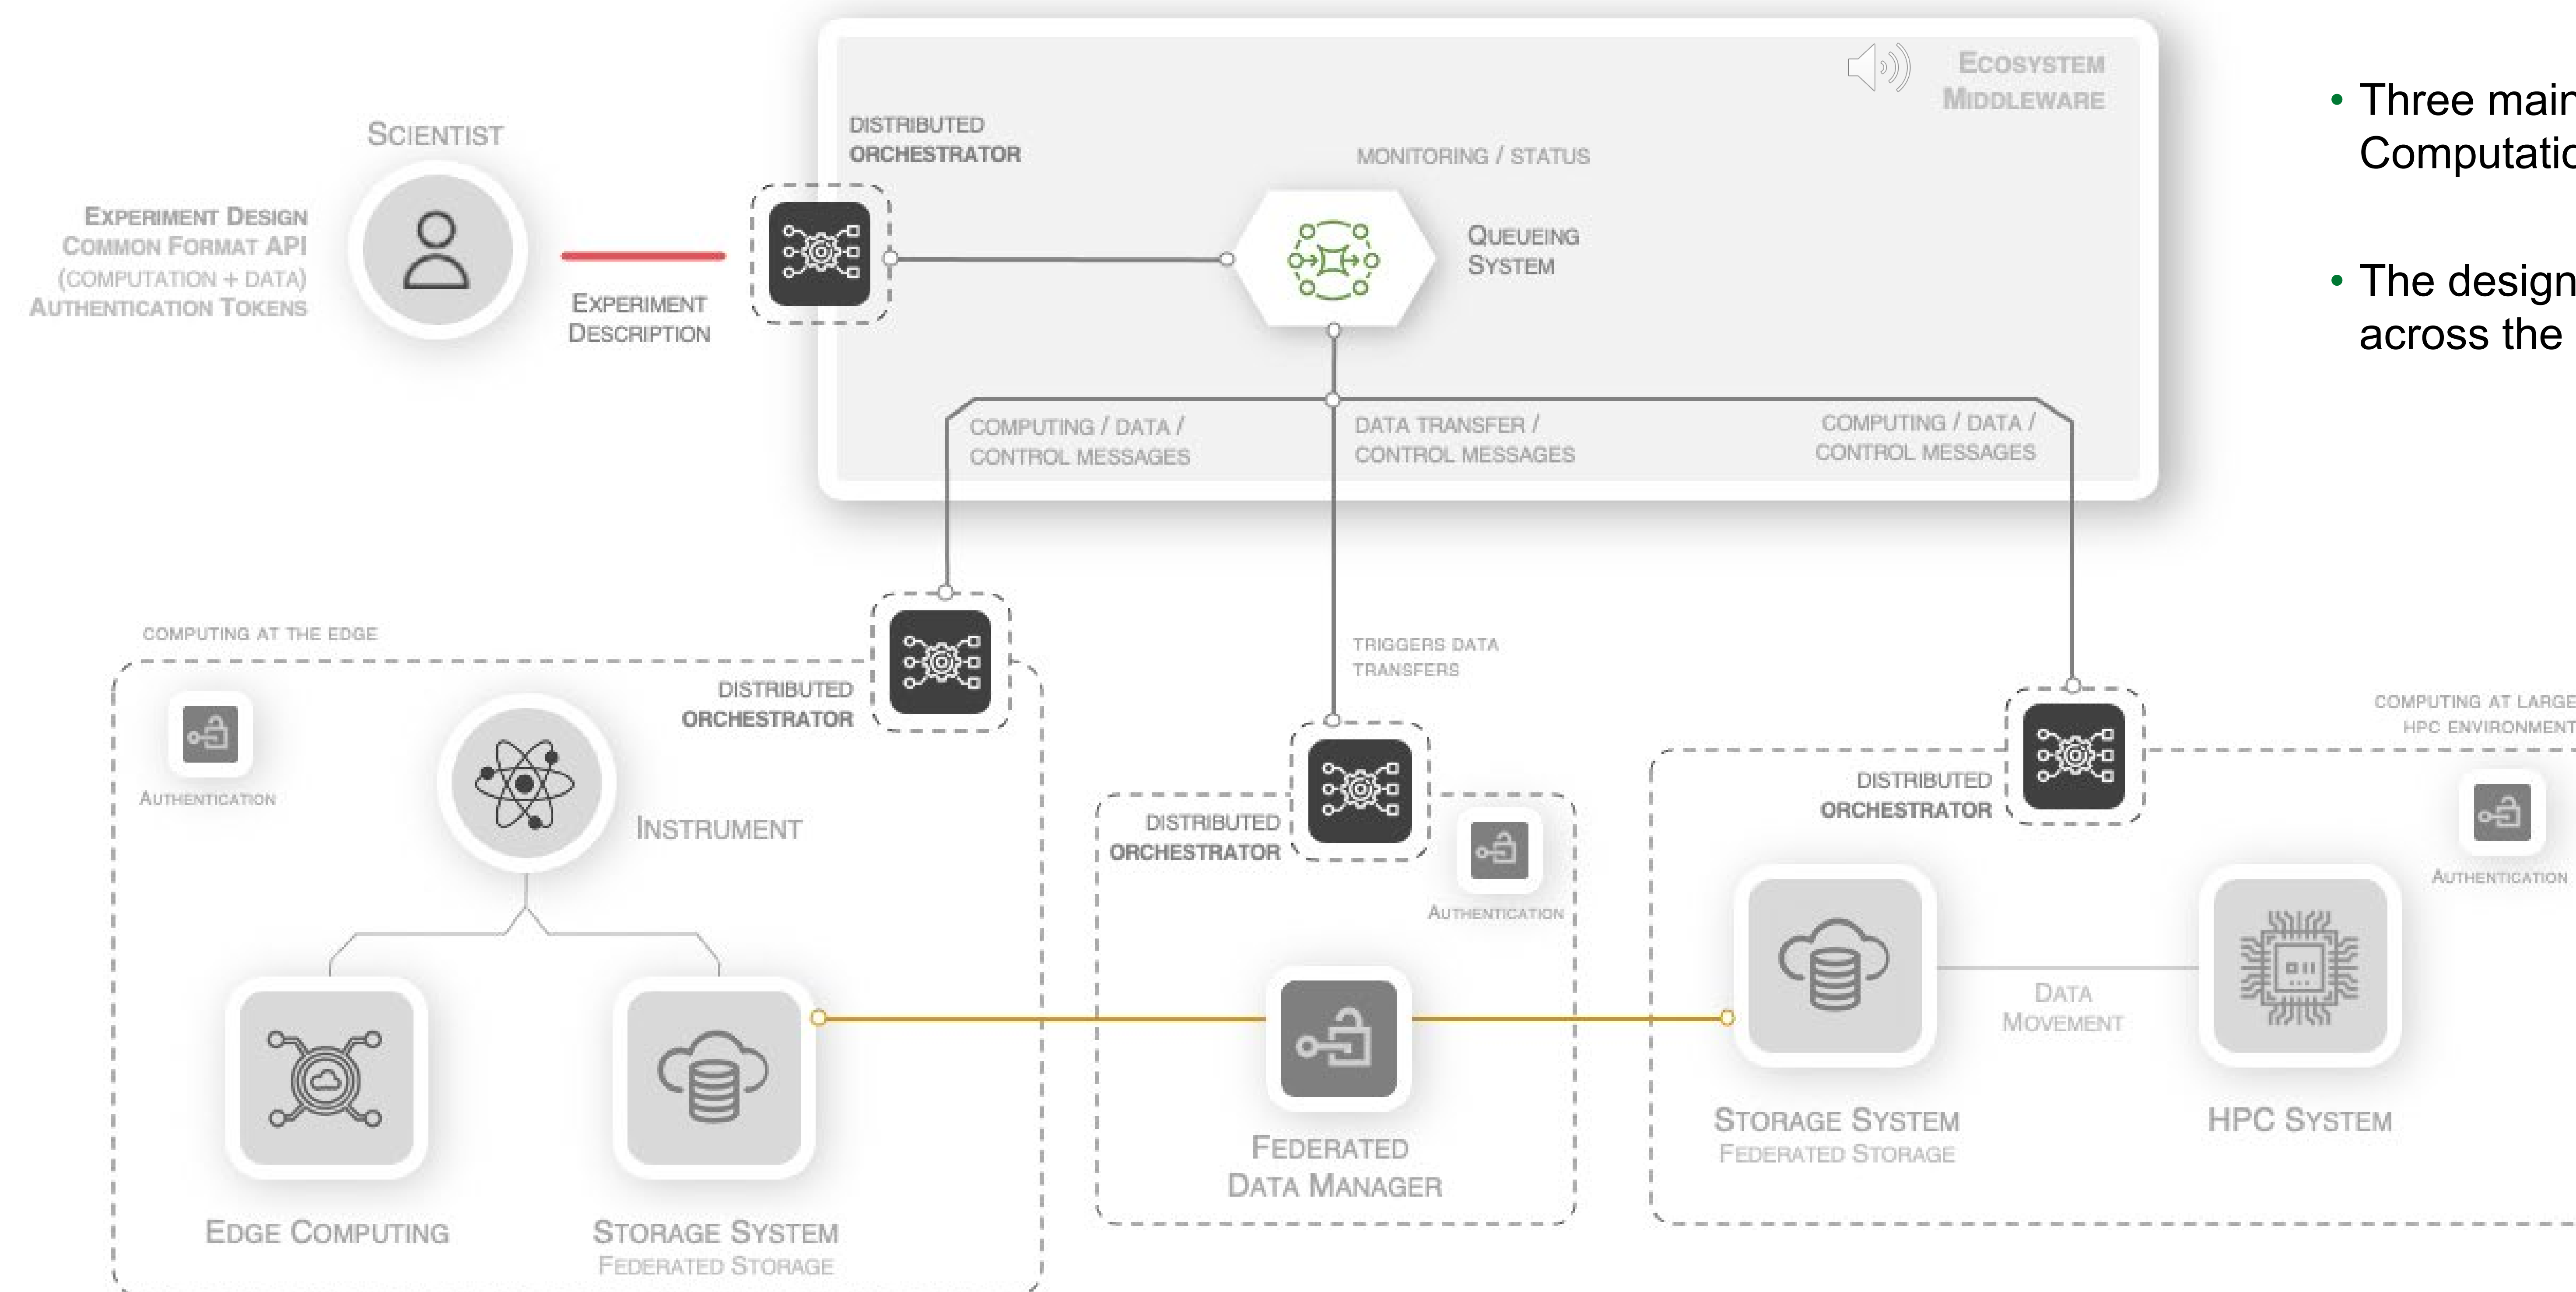
Cross-Facility Workflow Efforts

- Some efforts aim to provide specialized APIs for data management across facilities, while others develop common interoperable standards for bridging workflow systems [2].
- One framework for executing cross-facility workflows used two DOE light sources and three DOE supercomputers simultaneously [3].
- Our work targets the development of an ecosystem framework.

An Ecosystem Framework for Cross-Facility Workflows

- Our architectural design (see figure) incorporates a centralized queueing system along with distributed services called "orchestrators" which can:
 - operate autonomously and asynchronously.
 - enact one or more services (e.g., to capture the scientist's experiment campaign description or manage a local storage or edge computing service).
 - publish tasks representing data, compute, and control operations to be executed by other orchestrators.
- Data task examples: data movement operations, as well as metadata and provenance capture.
- Computing task examples: execution of individual functions or complex large-scale computational workflows.
- Control task examples: service availability, system and task status

Centralized Queueing System and Distributed Orchestration System



- Three main workflow types: Science Campaign workflows, Computational workflows, and Data workflows.
- The design can be used even without federated identity management across the connected facilities.

References and Acknowledgments

1. Antypas *et al.* (2021). Enabling discovery data science through cross-facility workflows. doi:10.1109/BigData52589.2021.9671421
2. Enders *et al.* (2020). Cross-facility science with the Superfacility Project at LBNL. doi:10.1109/XLOOP51963.2020.00006
3. Salim *et al.* (2021). Toward Real-time Analysis of Experimental Science Workloads on Geographically Distributed Supercomputers. doi:10.48550/ARXIV.2105.06571

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.