

Experiences in Building a User Portal for Expanse Supercomputer

Scott Sakai

San Diego Supercomputer Center,
University of California, San Diego
ssakaidmishin@ucsd.edu

Mahidhar Tatineni

San Diego Supercomputer Center,
University of California, San Diego
mtatineni@ucsd.edu

Christopher Irving

San Diego Supercomputer Center,
University of California, San Diego
cirving@ucsd.edu

Dmitry Mishin

San Diego Supercomputer Center,
University of California, San Diego

Martin Kandes

San Diego Supercomputer Center,
University of California, San Diego
mkandes@ucsd.edu

Shawn Strande

San Diego Supercomputer Center,
University of California, San Diego
sstrande@ucsd.edu

Subhashini Sivagnanam

San Diego Supercomputer Center,
University of California, San Diego
ssivagnanam@ucsd.edu

Mary Thomas

San Diego Supercomputer Center,
University of California, San Diego
mpthomas@ucsd.edu

Michael Norman

San Diego Supercomputer Center,
University of California, San Diego
mlnorman@ucsd.edu

ABSTRACT

A User Portal is being developed for NSF-funded Expanse supercomputer. The Expanse portal is based on the NSF-funded Open OnDemand HPC portal platform which has gained widespread adoption at HPC centers. The portal will provide a gateway for launching interactive applications such as MATLAB, RStudio, and an integrated web-based environment for file management and job submission. This paper discusses the early experience in deploying the portal and the customizations that were made to accommodate the requirements of the Expanse user community.

CCS CONCEPTS

• **Computing methodologies** → Modeling and simulation; Simulation support systems; Simulation environments.

KEYWORDS

User portal, HPC, Gateway, Cyberinfrastructure

ACM Reference Format:

Scott Sakai, Mahidhar Tatineni, Christopher Irving, Dmitry Mishin, Martin Kandes, Shawn Strande, Subhashini Sivagnanam, Mary Thomas, and Michael Norman. 2021. Experiences in Building a User Portal for Expanse Supercomputer. In *Practice and Experience in Advanced Research Computing (PEARC '21)*, July 18–22, 2021, Boston, MA, USA. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3437359.3465590>



This work is licensed under a Creative Commons Attribution International 4.0 License.

PEARC '21, July 18–22, 2021, Boston, MA, USA
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8292-2/21/07.
<https://doi.org/10.1145/3437359.3465590>

1 INTRODUCTION

In recent years, the need for accessing, using, monitoring, and managing HPC resources through a web interface has been constantly growing, partly owing to the increase in popularity of science gateways [1], and more generally, the democratization of supercomputing that has come in response to new approaches to computing. A user portal serves as a single interface providing access to distributed resources and services for a community of users, researchers, and students interested in using the resource. For decades, multiple science gateways such as CIPRES [2], Neuroscience Gateway [3], SEAGrid [4], I-TASSER [5] have provided and customized such interfaces to access high-performance computing resources, meeting the needs of a specific scientific community. This paper describes the early experiences in building the Expanse user portal that is designed to serve as the main web-based interface to the National Science Foundation (NSF) funded Expanse supercomputer [6] at San Diego Supercomputer Center (SDSC). The Expanse portal offers an easy, and unified interface that enables users of the Expanse system to manage data and cluster operations and provide interactive access to popular scientific applications. The Expanse portal will also provide information about allocation usage and resource utilization to the Expanse user community. The paper is organized as follows. Section 2 describes the desired features of the user portal and Open OnDemand that was used to build the portal; Section 3 describes the architecture and customization of the portal; and the current status and future work are mentioned in Section 4 of this short paper.

2 DESIRED FEATURES FOR EXPANSE USER PORTAL

In December 2020, SDSC launched Expanse to support “longtail” of computational science and engineering research with a theme of “computing without boundaries” [7]. Expanse consists of 729 compute nodes powered by two 64-core AMD EPYC 7742 processors and 56 GPU nodes containing four NVIDIA V100s with the entire system delivering 5.16 peak petaflops. The Expanse cluster

is managed using the Bright Computing HPC Cluster management system and uses the SLURM workload manager for job scheduling. Expanse will support thousands of users of batch-oriented and science gateway computing. It will provide new capabilities that will enable research increasingly dependent upon heterogeneous and distributed resources.

In response to a growing and more diverse user community, the Expanse user portal is being developed to provide a single point of service for many of the system features and to provide user information. The desired features of the portal include having a browser-based interface that allows users of the Expanse system to

- Review and monitor allocations/account information
- View detailed information regarding user's compute jobs
- Perform file and job management actions
- Set up reservations of resources for development work
- Allow interactive access for visualization tasks
- Manage resources for complex workflows
- Access a dynamic knowledgebase
- Access on-demand training materials

Providing a self-serving portal that can orchestrate services such as interactive visualization will enable users to explore and analyze their data with ease via a wide variety of visualization tools such as using batch computation, server-client interactive computation (e.g., MATLAB), and interactive web computation (e.g., Jupyter notebook) without the need to install any external software.

2.1 Open OnDemand

The Expanse portal is built using NSF-funded Open OnDemand platform [8, 12]. Open OnDemand (OOD) is an open-source HPC portal platform developed by the Ohio Supercomputing Center. The motivation for using the OOD platform includes:

- OOD platform has been adopted by multiple universities and centers and has a record for providing an easy way for administrators to integrate web access to their HPC/campus cluster resources.
- Using the OOD platform leverages prior NSF investment, and avoids building yet another user portal from scratch.
- OOD platform has many of the desired features that the Expanse team is interested in providing for the users of Expanse system such as
 - Providing an integrated browser-based environment
 - Graphical environment for file and job management
 - Terminal access without the need to install external clients such as Putty
 - Link and support third party credentials (e.g., XSEDE credentials)
 - Interactive access to applications
 - Ability to add new features such as knowledge base

The OOD platform was customized and modified to meet the needs of the Expanse user community, with modifications to the architecture to adhere to SDSC's security and system requirements; details are presented in the following section.

3 EXPANSE PORTAL ARCHITECTURE AND CUSTOMIZATIONS

The original conceptual architecture of the OOD platform is shown in Figure 1. The modified architecture for the Expanse portal is shown in Figure 2. Our major concerns were: 1) the dedicated hardware required as a part of the Expanse cluster to host the OOD platform; and 2) the requirement to place complete trust in the OOD platform and its dependencies. To address the hardware concern, we acquired separate server resources at SDSC for the current deployment of OOD for Expanse, with a design that allowed for future operation of OOD on public cloud resources if needed. To address security, our custom modifications allow the OOD platform to operate like a science gateway, which obtains and uses short-lived credentials for active users. In this mode of operation, it does not receive additional trust, nor does it introduce new risks that we have not already accepted. Described below are the new features of OOD implemented for the Expanse User Portal, these being implemented via custom modifications of the OOD source code [9] to meet the security and system requirements of using the Expanse cluster.

Batch system interaction with GSISSH from a dedicated host:

We evaluated OOD platform software and found that the interaction with the batch system was through standard command-line programs. This allowed us to run OOD on an external host similar to how science gateways connect to Expanse from an external host. The code was configured to wrap the command line interactions with a script to use GSISSH to connect to Expanse. With the addition of SSH connection multiplexing, frequent one-off GSISSH logins to the cluster were reduced to a single per-user GSISSH session.

Accessing user directories using SSHFS:

To allow the Expanse portal to display users' home and Lustre scratch directories, the OOD code was modified to mount the user's home/scratch directory using SSHFS [10]. The mount is active only as long as the user is active.

Supporting XSEDE authentication:

OOD relies upon Apache httpd for authentication and supports mod_auth_openidc which is an authentication/authorization module for the Apache 2.x HTTP server that authenticates users with an OpenID Connect (OIDC) provider. XSEDE runs a compatible OAuth2 endpoint as an OIDC provider. To allow users of Expanse portal to authenticate to our portal with their XSEDE credential, the OOD LUA code was modified to invoke our script, *setup_user.sh*, which obtains a MyProxy certificate for the user, and stores it in the OOD environment. We also created a user-mapping script, *user_map*, to convert XSEDE identities to Expanse account names, leveraging a process already in place to map XSEDE Kerberos principals to Expanse accounts.

Modification to the OOD process:

The following steps capture the modified OOD workflow process starting from when a user logs in to the Expanse portal:

- 1) Apache web server (using mod_auth_openidc) redirects the user to XSEDE's OAuth2 endpoint.
- 2) Apache (mod_auth_openidc) accepts the returned identity claim and stores the identity token.




Figure 1: Open OnDemand standard architecture places OOD inside cluster trust boundary.




Figure 2: Open OnDemand standard architecture places OOD inside cluster trust boundary.

- 3) Apache (OOD LUA + setup_user.sh) uses the identity token to retrieve the GSI certificate from XSEDE's MyProxy CA and stores the certificate and key in a location accessible to the GSISsh process.
- 4) Apache (OOD LUA + user_map) maps the user's XSEDE identity from OIDC to the username on Expanse and starts the per-user NGINX process (PUN).
- 5) PUN (our modification to the script) calls SSHFS to mount the users' home directory and scratch directory from the Expanse cluster's login node.

- 6) SSHFS calls GSISsh and uses the credentials that were previously stored.
- 7) GSISsh establishes a control channel with the cluster and subsequent SSH calls use this channel.
- 8) PUN applications work with job-related activity calling wrapped commands, which use SSH to run the commands on the cluster's login node.

Implementation of customized features:

The Expanse portal is deployed at <https://portal.expanse.sdsc.edu>. Access to the Expanse portal requires users to have a valid

XSEDE allocation on Expanse resource. Since Expanse supports only Rstudio client and allows a multitiered approach to running Jupyter notebooks [11], customizations were made to the OOD code to accommodate running of Rstudio client and Jupyter notebook as allowed on Expanse.

Currently, the deployed Expanse portal allows users to

- Access home and scratch directories using a graphical interface and perform file management activities such as copy/paste, upload small files, and edit files and scripts.
- Use of pre-filled templates that are provided for running batch jobs.
- Run interactive applications such as MATLAB, Rstudio and Jupyter notebook.
- Use an in-browser terminal without installing external clients such as Putty
- Run Expanse specific allocation usage and utilization monitoring tool (Expanse-client [7])

Code Contribution to the community:

Our custom scripts (setup_user.sh, user_map) and the Expanse portal code will be made available through the SDSC GitHub repository to other projects who are interested in using our deployment. We will also work with the OOD team to make the code available through their repository [9] for wider adoption.

4 CURRENT STATUS AND FUTURE WORK

4.1 Current Status

The Expanse portal went into production along with the Expanse system in Dec 2020. Expanse portal was demonstrated to users transitioning from SDSC's Comet to Expanse system as a part of a webinar in Oct 2020 and March 2021. Over 125 users have accessed the portal till March 2021.

We are utilizing a log parser to obtain and analyze Expanse portal web server statistics. The number of user visits per month is increasing and for the month of March we observed around 20-30 unique visits to the portal per day. These preliminary usage numbers demonstrate the interest in using the Expanse portal and usage is expected to grow as Expanse user base grows. However further analysis and monitoring of utilization of the Expanse portal are needed to identify details such as the number of failed jobs, and usage of data management activities which will help improve the portal functionalities.

4.2 Future Work

The initial version of Expanse portal implements a significant portion of features based on user requirements, but this is just a small set of the planned features. Other planned innovative features include incorporating a knowledge base which will use machine-learning techniques on user tickets to dynamically update content. The updates will be based on incoming tickets. An automated process to respond to commonly asked user questions will be incorporated to reduce the user support workload. Longer-term plans include providing event notification services, list and integrate data collections, and a consulting interface for Gateway administrators. We are closely following another research effort at SDSC to use

SSH certificates, which should open the cluster-detached mode of operation to environments that don't have GSISSH support.

ACKNOWLEDGMENTS

Expanse is supported by the National Science Foundation under award number 1928224. Additional support is provided by NSF award 1341698. The authors would like to acknowledge Trevor Petersen who contributed to the development of Expanse portal.

REFERENCES

- [1] Science Gateways. Retrieved from <https://sciencegateways.org/>
- [2] M Mark A. Miller, Wayne Pfeiffer, and Terri Schwartz. 2011. The CIPRES science gateway: a community resource for phylogenetic analyses. In *Proceedings of the 2011 TeraGrid Conference: Extreme Digital Discovery (TG '11)*. Association for Computing Machinery, New York, NY, USA, Article 41, 1–8. DOI:<https://doi.org/10.1145/2016741.2016785>
- [3] Sivagnanam, S., Majumdar, A., Yoshimoto, K., Astakhov, V., Bandrowski, A. E., Martone, M. E., & Carnevale, N. T. (2013). Introducing the Neuroscience Gateway. *IWSG*, 993.
- [4] Supun Nakandala, Sudhakar Pamidighantam, Shameera Yodage, Nipurn Doshi, Eroma Abeysinghe, Chathuri Peli Kankanamalage, Suresh Maruru, and Marlon Pierce. 2016. Anatomy of the SEAGrid Science Gateway. In *Proceedings of the XSEDE16 Conference on Diversity, Big Data, and Science at Scale (XSEDE16)*. Association for Computing Machinery, New York, NY, USA, Article 40, 1–8. DOI:<https://doi.org/10.1145/2949550.2949591>
- [5] Zheng, W., Zhang, C., Bell, E. W., & Zhang, Y. (2019). I-TASSER gateway: A protein structure and function prediction server powered by XSEDE. *Future generations computer systems : FGCS*, 99, 73–85. <https://doi.org/10.1016/j.future.2019.04.011>
- [6] Expanse homepage. Retrieved from <http://expanse.sdsc.edu>
- [7] Expanse userguide. Retrieved from https://www.sdsc.edu/support/user_guides/expanse.html
- [8] Hudak et al., (2018). Open OnDemand: A web-based client portal for HPC centers. *Journal of Open Source Software*, 3(25), 622, <https://doi.org/10.21105/joss.00622>
- [9] Open OnDemand GitHub. Retrieved from <https://github.com/OSC/ondemand>
- [10] SSHFS GitHub. Retrieved from <https://github.com/libfuse/sshfs>
- [11] RPS. Retrieved from <https://github.com/sdsc-hpc-training-org/reverse-proxy>
- [12] NSF award # 1534949. SI2-SSE: Open OnDemand: Transforming Computational Science through Omnidisciplinary Software Cyberinfrastructure