

Open OnDemand as a platform for Virtual Learning in Higher Education

Robert E. Settlage¹(0000-0002-1354-7609), Alan Chalker²(0000-0002-5475-8779), Jeff Ohrstrom², Eric Franz²(0000-0002-9662-412X), Doug Johnson²(0000-0002-4331-8508), and David Hudak²(0000-0002-9043-0850)

¹ Virginia Tech, Blacksburg VA 24061, USA,
rsettlag@vt.edu

² Ohio Supercomputer Center, Columbus OH 43212, USA

Abstract. The need for distance and virtual learning platforms has been emphasized by the COVID-19 pandemic. With the closure of campuses in Spring of 2020 and many classes moving to on-line only in Fall of 2020, platforms for facilitating computationally oriented curriculum have had to be quickly adopted. Open OnDemand offers a familiar web-based portal to computational resources such as high-performance computing and cloud. Through OnDemand’s customizable dashboard, students can be offered an interface tailored to the course schedule giving them a “just what I need” view. Advantages to instructors include a web accessible, platform agnostic interface leading to less time troubleshooting local student platforms and more time for discussion of the core course curriculum, a fully customizable course page, access controls and more. Here we present Open OnDemand as a platform for developing, deploying and presenting software and course material to software-oriented classes as used at Ohio Supercomputer Center and Virginia Tech.

Keywords: Open OnDemand · Education · HPC · Virtual Learning · Distance Learning

1 Introduction: Virtual Learning Challenges and Opportunities

Education in the Arts and Sciences has been dominated by traditional teacher led in-classroom education in higher education institutions. Through the needs dictated by mitigating the COVID-19 global pandemic, i.e. social and physical distancing, classes have transitioned from in-class to on-line. This transition has highlighted gaps and opportunities in our approaches, physical infrastructure, and even our curriculum. This is a discussion of some of the issues in offering virtual learning labs and successes we have found in our approach of using Open OnDemand to ease the transition in classes involving computational/software resources. From a educators point of view, many challenges revolve around loss (or

dampening) of the interaction between people: instructor-student and student-student. While we can't remove these issues, we can positively affect the instructor, student and administrators experiences as we describe here.

2 Open OnDemand as a Virtual Learning Portal

Open OnDemand [1] provides familiar, site configurable, and extensible web-based access to high-performance computing (HPC) and cloud resources. The use of web based tools and web forms to request computing and software resources makes traditionally difficult to use computing environments fully accessible to even the most novice users. While providing familiar tools to use HPC / cloud resources, Open OnDemand also provides methods to customize the landing page, dashboard, to include site branding, announcements or messages of the day, etc to allow creation of fit for purpose portals to HPC and cloud. Additional customizations can include both the tools users have access to and the amount/type of access they have (see [2] for a discussion of the state of the project). In combination, institutions and instructors are able to create class based portals to HPC/cloud resources customized to instruction at the site as a whole and to individual classes.

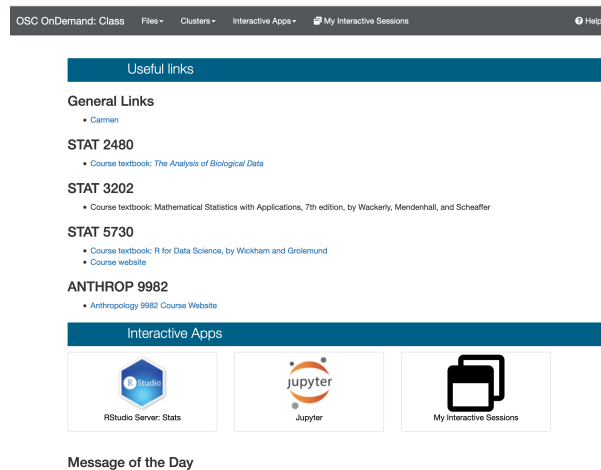


Fig. 1. Course oriented dashboard at OSC.

2.1 Center Branded Course Site

Customizing the branding and dashboard landing page are a core feature of Open OnDemand. Through this feature, sites are able to display announcements, messages of the day, system utilization metrics and more. An extension

of this feature is that sites can create instruction oriented Open OnDemand portals. Figure 1 shows the Virtual Learning Dashboard at the Ohio Supercomputer Center (OSC). Through this dashboard, students are able to access course materials and apps specific to their study.

2.2 Extensibility

Open OnDemand has a plug-in based architecture allowing for creation of *custom* applications. These custom applications are, at their core, form based HPC/cloud job submission requests. These can be as simple as a form to start RStudio or Jupyter, or more complex to start a compute pipeline. For our purposes in Virtual Learning Labs, we are creating forms to start instances of software specific to a class with job characteristics suitable to that class. As an example, see Figure 2. For this introductory bio statistics class at Virginia Tech (VT), we have predetermined what a student will receive in terms of memory and cores, limited the length of an interactive session, and only exposed the form fields necessary to start a session. This fine control allows for more flexible instruction. For sessions where parallel computing may be used, we could change the resource request to allow for multi-core functions. For sessions where memory will be a component, again, we can size the backend resources as appropriate. Note, we could have exposed these fields in the web form or we could have simplified the submission form further limiting the account and partition students have access to, but chose to leave these as topics for discussion of computing on large heterogeneous systems as a teaser. These are choices instructors can make to suit the needs of the class.

The screenshot displays the OnDemand Virtual Learning interface. On the left, a sidebar lists various applications: Dragonstooth, DNAmonitor, Jupyter Notebook, MATLAB, QGIS, Remote Desktop, Rstudio, Rstudio for STAT3615 (highlighted), Stata, and VisIT. The main panel shows the configuration for 'Rstudio for STAT3615'. It includes a description: 'This app will launch a Rstudio Server in a container on one node.' Below this, there are sections for 'Account' (with a text input field containing 'lrjohn-STAT3615-sp20'), 'Partition' (with a dropdown menu set to 'normal_q'), and 'Number of hours (min-1, max-4) -- after this time, your job will be killed' (with a numeric input field set to '1'). A checkbox for 'I would like to receive an email when the session starts' is present and unchecked. A blue 'Launch' button is at the bottom. A footnote states: '* All Rstudio for STAT3615 session data is generated and stored under the user's home directory in the corresponding [data root directory](#).'

Fig. 2. Cluster specific OnDemand Interactive Apps highlighting the presence of a class app and the simplicity a job request form can take.

2.3 App Access Control

"A key feature of OnDemand over a traditional web service is the per-user web server model where all the web applications are run as the Linux user" more fully described in [3]. In contrast to many web services, OnDemand leverages the Linux kernel for security and accounting which means server side security models and access control are honored. Operationally, access to specific apps is controlled through normal Linux group membership because the web service is running as the user. Individual apps are deployed on the server within a folder specific to that app. Permissions can be applied to that folder to limit access via group memberships. Users only see apps they have permissions to access because the web service accessing files within the various app folders using the user permissions. This creates an easy environment to maintain, all access is controlled through group memberships in the normal way! In Figure 2 above, the menu item for the class app is only visible to members of the stat3615 group.

3 Open OnDemand Use Cases

Instructional use of Open OnDemand had started prior to COVID-19. However, COVID-19 has stressed the need for stable, consistent and scalable infrastructure for teaching. Below we highlight impacts at Virginia Tech and across Ohio via OSC.

3.1 Intro to HPC Computing

The Advanced Research Computing group at Virginia Tech is often asked to give an introductory "Intro to HPC" lecture in classes such as statistics, math, and engineering. Before Open OnDemand, a typical "Intro to University HPC Resources" lecture would be consumed by platform driven connection issues. "From a Mac, use the terminal, SSH to ..., from a Windows machine, download and install an SSH client (e.g. Putty, MobaXterm, etc), SSH to Now, if you are configured for 2FA with your phone, have the Duo app ready, and go" For a class of approximately 40 students, getting the students logged in to the cluster was a frustrating exercise often consuming the majority of a 1 hour lecture. Open OnDemand changes the discussion to: "using a browser, navigate to *ondemand.<site>.edu*, log in as you would any other university resource, now click the pulldown for cluster access". In most cases, students are logged into the cluster in 5 minutes leaving the majority of the hour to discuss topics of interest to the class subject.

3.2 Enabling Instruction Across Ohio

The Ohio Supercomputer Center is a statewide resource that provides super computing services and computational science expertise to Ohio university researchers as well as Ohio industries. With a goal of leading science and engineering research efforts, education and training potential users is an important

component of the OSC mission. Students across Ohio can gain exposure to HPC through a variety of opportunities including summer workshops and research events and now class instruction via Open OnDemand. Figure 3 and Table 1 show the reach of the OSC Open OnDemand education of students across Ohio with nearly 1900 students across 14 universities.

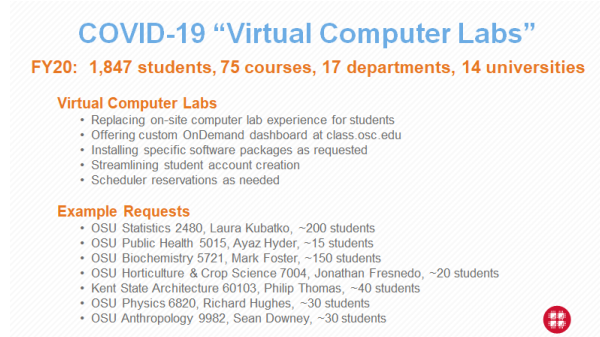


Fig. 3. Distribution of students accessing OSC resources via OnDemand across Ohio.

Table 1. Snapshot of student distribution. FY2020: 1,847 students, 75 courses, 17 departments, 14 universities

Institution	Students	Courses	Departments
Bowling Green State University	10	2	2
Denison University	3	1	1
Kent State University	23	6	2
Kenyon College	18	2	1
Miami University	33	2	2
Mount Union	8	2	1
Ohio State University	1,422	46	14
Ohio University	3	1	1
Stark State College	17	1	1
University of Cincinnati	285	8	3
University of Dayton	10	1	1
Urbana University	2	1	1
Wittenberg University	7	1	1
Xavier University	6	1	1

Examples of courses utilizing OSC’s OnDemand include: Kent State’s College of Architecture and Environmental Design ARCH 60103 provided students access to Maya; OSU Physics 6820 provided students a custom Python programming environment using Jupyter through OnDemand; College of Public Health PUBHTLH 5015 utilized a custom R environment through RStudio via OnDe-

mand; OSU Department of Horticulture and Crop Science's HCS 7004 Genome Analytics course utilized a custom R environment through RStudio via OnDemand. There are many other courses across many disciplines, including Business Analytics, Bio Chemistry, Statistics, and Anthropology.

4 Opportunities and Lessons Learned

While challenging, virtual learning offers opportunities. The opportunities span the gamut from cost savings, potentials to improve the student/instructor experience, modernize the instruction for methods more aligned with real world problems, and opportunities to scale the classroom to meet increases in demands for education. We highlight several of these below:

Remote Troubleshooting One of the hallmarks of instruction in classes involving software and computing are issues surrounding platform issues. As an example, consider RStudio/R for instruction in Statistics classes. Inevitably, immediately following the lecture demonstrating RStudio/R installation and use, student need for troubleshooting platform installation issues effectively monopolizes the instructors office hours for the next 1-2 weeks. Issues include everything from conflicts with previous installs of *LaTeX*, OS language problems (underlying OS was installed as Chinese version for instance), or simply poorly optimized installs missing links to BLAS or other libraries designed to improve performance. As an instructor, these issues are difficult to troubleshoot in person. Moving this to a remote troubleshooting session tests the ability of both student and instructor in both accurately describing the solution and succinctly describing instructions to fix the problem.

Using a web accessible University resource such as Open OnDemand backed by HPC or Cloud removes most, if not all platform issues. Plus or minus the browser, the software is the same for all students in the class. At Virginia Tech, on our HPC systems, instructors and students are equivalent in terms of access rights and permissions. This means troubleshooting a student issue is reduced to repeating the problem in an identical environment.

Cost Savings Moving instruction to virtual modes offers potential savings. Many Universities have computer laboratories. The labs are populated with physical computer workstations, often using some sort of virtualization and networked software licensing, hosting software in common courses such as STATA, R, Matlab, and more. To make sure there are enough workstations, the labs are slightly oversized, but there are still workstation availability bottlenecks at key points throughout a semester. Moving these workstations to a virtual lab giving students access to the software and computing resources through web based access obviates the need for monitors and even a physical lab. Further, depending on how the virtualization is done and what hardware backs the virtual workstations, there could be opportunities to share hardware that were not possible

when a physical keyboard was a limiting device. Frequently, the software use within an introductory course is very minimal from a computational standpoint suggesting high over-subscription rates are possible and use of the physical hardware can continue well past the depreciation period without degradation in user experience.[4] In many cases, both hardware and software can be shared across different uses, such as research and education.

Improve the Student Experience Many classes use scientific software to enhance the learning experience. However, the computer is more of a means to an end rather than an end in and of itself. For example, Statistics in Biology may have as the primary goal foundational topics around probability, testing and prediction. The computer is used to reinforce the learning objectives and begin discussions about data. Local platform issues are really a distraction to the learner. Removing these barriers, at least in introductory courses, removes the distraction and returns the focus of the student to the class topic. Importantly, virtualization should create a unified student experience.

Improve the Instructor Experience While errors and other installation / configuration issues are often the bane of the first couple of weeks of a course using computational tools, there are other platform oddities that present hurdles to instructors. Consider Microsoft Office. Differences in the interface on Mac and PC versions cause constant frustrations during follow-along lectures. Menu items, tool bars etc are different even in the same tool! In some instances, platform differences are desirable, in others, they are frustrating obstacles.

Reduce Administrative Load In any computing oriented class, support is a key resource often delivered through a combination of instructor and system admin / computational scientist office hours. Open OnDemand reduces administrative support through pinning workflows to successful job submissions. Additionally, the additional users flowing through classes are administrated the same as normal users in terms of access controls, account allocations etc. In effect, students are normal HPC users with an intuitive interface. Through interface design, instructors are able to provide contextual help and guide students through use and problems where they would otherwise have been left with a blinking cursor.

Scaling Modern instruction in just about every discipline is impacted by the current data deluge. In fact, current estimates by the Visual Capitalist suggest we will generate 463 Exabytes of data everyday.[5]. To make sense of this data, we need to instruct students in modern data analysis methods. Open OnDemand facilitates the transition from novice to big data analyst by providing a common platform for both computing paradigms. The underlying infrastructure is the same, the interface can telescope based on user needs and access privileges.

Limitations of OnDemand for Classroom Support Providing a large icon launch interface for RStudio and Jupyter on the OSC course oriented dashboard, as shown in Fig. 1, required technical expertise in adding custom HTML to the localization files for the dashboard to include these links. OnDemand 2.0 plans to make this type of customized user or group tailored launch interface easier to provide.

While OnDemand provides a files app for uploading Jupyter notebooks from a student's laptop, professors asked if they could provide students the ability to save and load Jupyter notebooks from internet accessible storage environments such as OneDrive, Google Drive, or GitHub. Without these capabilities, students using tablets for the class may have a difficult time participating if requisite files are not uploaded before hand by the professor to a shared location on the file system. Furthermore, professors and TAs still need access to students completed assignments in the form of Jupyter notebooks in order to grade them. Students need to be able to submit completed assignments to their professor without inadvertently sharing those completed assignments with other students. Ideally, Students would be able to submit the completed assignment through the web interface, completely bypassing the need to first download the notebook to their device. Integrating OnDemand with GitHub Classroom or Canvas LMS may provide an optimal solution to these challenges.

5 The Future of Digital Learning

The future of learning is digital. As a society, nothing has more impact on our lives than the computer has now. To make sense of the unimaginable amount of data we are generating everyday, we have a need for computational resources. Educating our workforce in tools necessary to transform, analyse and extract knowledge from the information we are generating will require digital proficiency. Open OnDemand, as a portal tool set, enables students and instructors alike, provides a familiar framework for accessing computational power necessary for today's problems, and reduces administrative loads while allowing scaling for class demands and problem sets.

6 Acknowledgements

National Science Foundation of the United States Awards NSF SI2-SSE-1534949 and CSSI-Frameworks-1835725.

References

1. D. Hudak, D. Johnson, A. Chalker, J. Nicklas, E. Franz, T. Dockendorf, B. McMichael, Journal of Open Source Software (2018). DOI 10.21105/joss.00622
2. A. Chalker, E. Franz, M. Rodgers, T. Dockendorf, D. Johnson, D. Sajdak, J.P. White, B.D. Plessinger, M. Zia, S.M. Gallo, R.E. Settlage, D.E. Hudak, Concurrency and Computation: Practice and Experience (2020). DOI 10.1002/cpe.6114

3. D. Hudak, D. Johnson, E. Franz, T. Baer, T. Dockendorf, K. Cahill, (2018). DOI 10.6084/m9.figshare.7069691.v1. URL https://figshare.com/articles/Open_OnDemand_Access_Clusters_Gateways_and_Interactive_Apps/7069691
4. E. Franz, H.G. Ozer, T. Dockendorf, V.S. Gadepalli, A. Webb, A. Chalker, M. Pietrzak, M. Rodgers, D. Johnson, D.E. Hudak, Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (2019). DOI 10.1145/3332186.3332211
5. R. Ronconteur. Data (2019). URL <https://www.visualcapitalist.com/wp-content/uploads/2019/04/data-generated-each-day-wide.html>