Orion: Discovery Environment for HPC Research and Bridging XSEDE Resources

Semir Sarajlic*, Nieranjan Edirisinghe, Yuriy Lukinov, Michael Walters, Brock Davis, Gregori Faroux
Research Solutions, Georgia State University

34 Broad Street
Atlanta, GA 30303

{ssarajlic1, nieranjan, ylukinov, mwalters12, bdavis64, gfaroux} @gsu.edu

ABSTRACT

We present a case study on how Georgia State University (GSU) has grown its active High Performance Computing (HPC) research community by 80% in 2015 over previous year, and how GSU is projected to double its active HPC research community for 2016 over 2015. In October 2015, GSU launched an institutional HPC resource, Orion, which provides batch and interactive compute environment. Currently, Orion supports both the traditional and non-traditional research communities on our campus as well as our affiliates from Qatar University, University of Toronto, and Georgia Tech. At GSU, Research Solutions' HPC facilitators are responsible for facilitating the HPC research, which is done in a form of providing technical support in developing pipelines and automating job submission process for various applications that researchers need for their research. This approach has resulted in nearly 80% growth in our active HPC users from 2014 to 2015, and currently we are tracking at doubling our active HPC user community in 2016. XSEDE remains a backbone of our ambitious goals, as we rely on XSEDE for providing us the necessary resources for select users whose research quickly exceeds our local infrastructure.

CCS Concepts

• **General and reference~Metrics** • *General and reference~Design* • **Applied computing~Education** • Computer systems organization~Heterogeneous (hybrid) systems

Keywords

High Performance Computing; metrics; automation; training

1. INTRODUCTION

The advancements in data sharing and procurement have led to an explosion in the volume of data worldwide that rapidly grows across all domains. This explosion of data is driven by automated instruments [9, 10], social media [2, 6], and scientific research in

fields such as astrophysics, bioinformatics, and geoinformatics [7, 12, 15]. The exponential growth in data [3, 5] presents a multifaceted challenge to researchers due to its volume, variety, velocity, and veracity of the data, which adds complexity to the problem by orders of magnitude that creates a challenge to traditional researchers and a greater challenge to non-traditional researchers.

Using the knowledge gained from the XSEDE Campus Champion and Advanced Cyberinfrastructure-Research, Education Facilitator (ACI-REF), and Southern Partnership in Advanced Networking (SPAN) programs, GSU launched its new institutional HPC resource, Orion. The design of Orion enables service provision to both traditional and non-traditional researchers, which enables rapid growth in our active HPC research community, which we define as any user with an active project during a year. In the following sections, Orion's design and our approach in working with traditional and non-traditional researchers will be discussed along with the results of our efforts.

2. DESIGN OF ORION

2.1 Hardware

Orion (Figure 1) offers roughly 8 teraflops (based on LINPACK benchmark) of capacity across 9 compute and GPU nodes (7 compute nodes and 2 accelerator based nodes) with 87 TB of NFS storage. This system is built from hardware procured through IBM, DELL and Silicon Mechanics. The specific hardware is as follows:

- IBM x3850 compute nodes: 4 socket, 10 core Sandy Bridge processor, 512 GB RAM, 2 TB disk
- Silicon Mechanics 8027R compute node: 4 socket, 10 core Ivy Bridge Processor, 512 GB RAM, 4 TB disk
- R720 GPU server: 2 socket, 10 core Ivy Bridge Processor, 128 GB RAM, 2 TB disk, 2x NVIDIA Tesla K40m
- R720 GPU server: 2 socket, 10 core Ivy Bridge Processor, 128 GB RAM, 2 TB disk, 2x Intel Phi 7120

2.2 Management, Scheduling, and Application Stack

Orion is our first system to utilize two level user account provisioning. First level, user login with campus ID via LDAP, which provides users private home directory space of 200 GB. Second level is project accounts, which users log into from their home directory (i.e., campus ID) by initiating a command, 'myprojects', which will provide them a list of projects they have access to and enable them to select the project they wish to work on and submit jobs once logged into project (Figure 2). Login to the project from home/LDAP account is enabled by RSA keys. Project accounts provides users with 1 TB of work space, and multiple users who are working on the same project may have

* Corresponding Author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

XSEDE16, July 17-21, 2016, Miami, FL, USA

© 2016 ACM. ISBN 978-1-4503-4755-6/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2949550.2952770>

Figure 1. Design of Orion with LSF scheduler as the workload manager. Note, all user files and applications reside on a network file system (NFS) consisting of 87 TB of space. Open XDMoD [11] and Ganglia [8] are used for tracking utilization and metrics of the system.

access to the same project account on Orion, which enables collaboration within the system and reduces duplication of data on the system. Additionally, researchers may have multiple project accounts on Orion, which allows compartmentalization. Project accounts provisioning is our innovative approach at provisioning accounts on HPC systems that enables maintaining metrics for HPC usage based on projects rather than individual users. In addition, we are able to measure the growth of our research portfolio based on the HPC research. This information is used in guiding our strategic initiatives in supporting our researchers.

Orion utilizes IBM's Load Sharing Facility (LSF) scheduler for workload management. Within our scheduler, we created three user group types: sponsored research (SP), non-sponsored/academic research (AR), and owners (NI). This allows us to assign different priority for different projects belonging to each of these three groups. For example, SP has higher priority than AR and it can preempt AR job, and NI has the highest priority on select nodes that were purchased by researchers and NI may preempt any running job on owners' host. Jobs are submitted from project accounts only, which enables us to track metrics of system usage by parsing the LSF log data (i.e., *lsb.acct*) using Open XDMoD¹ [11]. Our design and implementation of management and scheduling on Orion creates a model for condominium approach to the HPC cluster that provides a solution to departments and researchers wanting to have dedicated resources for their research without the responsibility of maintaining the system.

Advanced scientific software for large scale computation and visualization is available on Orion. Our software packages include commercial applications (e.g., MATLAB®, Gaussian, STATA®, SAS®, Amber, etc.) and open source applications (e.g., R, VisIt, ParaView, NAMD, etc.). All our applications are centrally installed and managed, which offers tight control of the system and provides

a level of abstraction for our researchers. Modulefiles are used for implementing necessary environment variables for various applications. We offer automated job submission scripts (i.e., wrappers) that simplify the complexity of using HPC resources and enable users to conduct work on Orion in both interactive and batch modes [1]. In the next section, we will talk about our approach in facilitating HPC research at GSU.

3. APPROACH

Our approach in facilitating HPC research is as follows:

- Simplify access to HPC resources
- Provide automated job submission scripts for specific applications
- Provide user training workshops
- Work with researchers in scaling their research to XSEDE resources

Figure 2. A schematic of a two level account provisioning on Orion. In this example, a user campusID has three project that he/she may access via 'myprojects' command, then once in project they may launch jobs using 'bsub'/wrapper.

¹ GSU HPC metrics: Open XDMoD – <http://xdmod.rs.gsu.edu>

Figure 3. Distribution of average job size on Orion.
Data is presented using Open XDMoD [11].

3.1 Simplified access to HPC resources

We simplified the access process by developing a web application² using C# and RESTful web services, which is hosted in Azure. Our researchers are asked a few simple questions, and an object is created that provisions a project account on Orion. This provides researchers with rapid access to Orion, and researchers wanting to use common applications (e.g., MATLAB®, STATA®, R, Gaussian...) are able to submit jobs using our automated job submission scripts [1] that are already in place and information about them is provided to the researchers via our HPC knowledge base³.

3.2 Automation of job submission scripts

We created wrappers for automating job submission scripts for major applications that are used frequently such as: MATLAB®, R, STATA®, Gaussian 09, SAS®, and Trinity.

For example, 'qR' is the command we use to submit R batch jobs on our system. This command takes 3 input arguments:

qR [R file] [number of processors] [time requirement]

R '.r' file is the first argument, the second argument is the number of cores, and last argument is time limit. Then, 'qR' script determines correct queue, writes LSF job submission script, and writes prolog and epilog files necessary to perform file transfers from remote /home/rshome directories to local work directories /runjobs, and submits the jobs. This allows users to bypass the requirement to learn additional steps to use HPC resources, and it allows system administrators to manage the system efficiently. For example, using these wrappers we were able to set CPU/Memory affinity and use /dev/shm/ space without the user involvement.

We created wrappers to address both HPC and HTC needs of our researchers as we have equal amount of researchers who run parallel and serial jobs as shown in Figure 3. Our wrappers and pipelines that we developed have attracted researchers from Georgia Tech, Qatar University, and University of Toronto to our system, Orion, because our wrappers simplify the HPC usage.

Figure 4. Orion's utilization by researcher's projects.
Note, RS10085 accounting nearly 50% and RS10071 accounting nearly 25% of Orion's utilization. Data is presented using Open XDMoD [11].

In addition, we partner with departments on projects that help and enable researchers from the department to conduct HPC research. For example, we helped develop application pipelines for Neuroscience Institute that enabled their researchers to conduct transcriptome assembly and annotation [13, 14]. Our partnership with Neuroscience Institute had extended further with them purchasing a node that was added to Orion. This allowed the Neuroscience researchers and their teams to conduct their research without interruptions (i.e., preemptions), and it enabled GSU researchers to utilize free cycles when Neuroscience researchers were not using their node.

3.3 User Training

GSU is an XSEDE Level 3 Service Provider, and we host monthly XSEDE workshops for our HPC research community in addition to having custom HPC workshops we provide to our researchers. Our workshops range from Introduction to Linux and Workload Management in HPC Environment to application specific workshops such as R and Spark. Most of the time our workshops are used as an opportunity to identify new researchers, whom we can help. This enables us to build a relationship with researchers, and in some cases, we conduct research together.

Figure 5. GSU's HPC community of active HPC researchers. Note, 2016 metrics are as of April 15 2016, and our projected total number of users is expected to reach 220, which would represent a 100% increase over 2015.

² GSU HPC Portal – <https://hpc.gsu.edu>

³ GSU HPC Knowledge Base (for job submission scripts click on applications link) – <https://help.rs.gsu.edu/display/PD/Orion>

Figure 6. Representation of GSU HPC research community according to their department/discipline as of April 15, 2016. 29 departments/disciplines are represented in the above diagram.

3.4 Transition to XSEDE resources

After our researchers have their codes/scripts production ready, and if they require more resources than we can supply, then we work with the researchers in transitioning their workload to national resources that are part of XSEDE (e.g., Stampede, OSG, SuperMIC) [16]. For example, on one of the projects, we helped researchers conduct their molecular dynamic research and migrate their workload onto XSEDE, which can be read about in the following publication [4]. Figure 4 shows Orion's utilization by researchers' projects. Project RS10071 that accounted over 50% of our system utilization (early 2016) was successfully migrated to OSG. Now, our current effort is in transitioning RS10085 project's workload onto Stampede via XSEDE [16]. In migrating RS10085's workload, we implemented similar wrapper concepts on Stampede in a form of modified *.bashrc* file, and writing a wrapper for HTC Launcher and R.

Our ability to tap into XSEDE resources has enabled us to provide an end-to-end solution for our researchers, as we work with them in their initial development and then scale their workload to national resources. At GSU, Research Solutions' HPC Facilitators focus their efforts on working with researchers in enabling them to use Orion and XSEDE resources along with helping researchers improve their scripts/codes that enable researchers to utilize HPC resources optimally. This approach has been critical at growing our active HPC research community as indicated in Figures 5 and 6 that will be discussed further in the Results section.

4. RESULTS

Design and implementation of Orion management, scheduling and application stack with wrappers has enabled us to provide a user friendly system with a collaborative environment that enables a PI

of a project to share the data within a project amongst his/her colleagues and GRAs working on the same project. Additionally, our approach at working with researchers has allowed us to grow our active HPC research community substantially as illustrated in Figure 5, in which we were able to increase our active HPC research community by approximately 80% in 2015 over previous year, and we are currently tracking at doubling our active HPC research community in 2016. The data presented in Figure 5 is generated based on the active project accounts during/within the year. If the project is not active (i.e., no jobs are submitted within a year), then we do not count those users. Additionally, Figure 6 lists the departments that our researchers are affiliated with, currently, we have researchers from 29 departments. Furthermore, we were able to track the HPC usage among our sponsored projects (SP) on Orion, which allowed us to measure the impact of HPC research on our university's research portfolio. As of April 2016, HPC research that is facilitated on Orion accounted for 3.2% of GSU's yearly research portfolio.

5. CONCLUSION AND FUTURE WORK

This case study presents our design of Orion that features a collaborative environment enabled by projects, and wrappers for variety of applications for interactive and batch compute environment. The design of Orion simplifies usage of HPC, which non-traditional researchers have found very convenient to use, and this has enabled rapid growth of HPC research community at GSU. In addition, we present our approach on working with researchers and scaling them out to XSEDE resources, which enables us to provide an end-to-end solution to our researchers from development stage to large scale production stage of their research. Our approach in simplifying HPC usage has yielded significant growth in our active HPC research community.

In order to maintain our momentum we will bring a training system online to our HPC research community that would enable hands on training and allow us to expand our training efforts across the university. As a university initiative, we are teaming with the department's faculty (e.g., Geosciences, Neuroscience, etc.) to incorporate a session(s) in their current computational/data analytical courses in which we would be able to introduce the students to Big Data analytical methods. This initiative will act as a seed for fostering a community of Big Data researchers at GSU.

6. ACKNOWLEDGMENTS

This work was in part of Southern Partnership in Advanced Networking project that is supported by the National Science Foundation grant number ACI – 1440659. This work was supported by XSEDE Campus Champion Grant – GEO150002. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1053575.

7. REFERENCES

- [1] Edirisinhe, N., Sarajlic, S. HPC for everyone: Democratization of HPC resources for utilization by non-traditional researchers. XSEDE 2015 Conference, St. Louis, MO. July 26-30, 2015.
<https://conferences.xsede.org/xsede15/schedule/posters>
- [2] Fulgoni, G. and A. Lipsman, Digital Game Changers: How Social Media Will Help Usher in The Era of Mobile and Multi-Platform Campaign-Effectiveness Measurement. *Journal of Advertising Research*, 2014. 54(1): p. 11-16.
- [3] Greengard, S., Weathering a New Era of Big Data. *Communications of the ACM*, 2014. 57(9): p.12.
- [4] Ho, K. C., Hamelberg, D., (2016). Oscillatory Diffusion and Second-Order Cyclostationarity in Alanine Tripeptide from Molecular Dynamics Simulation. *Journal of Chemical Theory and Computation*, 12, 372–382
- [5] Howe, H., Costanzo, M., Fey, P., Gojobori, T., Hannick, L., Hide, W., Hill, D.P., Kania, R., Schaeffer, M., Pierre, S.S., Twigger, S., White, O., Rhee, S.Y., Big data: The future of biocuration. *Nature*, 2008. 455(7209): p. 47-50.
- [6] Kim, H.J., A. Pelaez, and E.R. Winston, Experiencing Big Data Analytics: Analyzing Social Media Data in Financial Sector as a Case Study. *Proceedings for the Northeast Region Decision Sciences Institute (NEDSI)*, 2013: p. 62-69.
- [7] Li, J., Meng L., Wang F. Z., Zhang W., Cai, Y., A map-reduce-enabled SOLAP cube for large scale remotely sensed data aggregation. *Computers & Geosciences*, 2014. 70: p. 110-119.
- [8] Massie, M., Li, B., Nicholes, B., Vuksan, V., Alexander, R., Buchbinder, J., Costa, F., Dean, A., Josephsen, D., Phaal, P., Pocock, D., 2012. *Monitoring with Ganglia* (1st ed.). O'Reilly Media, Inc..
- [9] Nelson, R., Collecting data, from sensors to systems. *EE: Evaluation Engineering*, 2014. 53(6): p.14-21.
- [10] Noor, A., Putting Big Data to Work. *Mechanical Engineering*, 2013. 135(10): p. 32-37.
- [11] Palmer, J. T., Gallo, S. M., Furlani, T. R., Jones, M. D., DeLeon, R. L., White, J. P., Simakov, N., Patra, A.K., Sperhac, J., Yearke, T., Rathsam, R., Innus, M., Cornelius, C. D., Browne, J.C., Barth, W.L., Evans, R.T., "Open XDMoD: A Tool for the Comprehensive Management of High-Performance Computing Resources", *Computing in Science & Engineering* 17.4 (2015):52-62, 2015 10.1109/MCSE.2015.68
- [12] Savage, N., Bioinformatics: Big data versus the big C. *Nature*, 2014. 509(7502): p. S66-S67.
- [13] Senatore, A., Edirisinhe, N., Katz P.S., (2105) Deep mRNA sequencing of the Tritonia diomedea brain transcriptome provides access to gene homologues for neuronal excitability, synaptic transmission and peptidergic signalling. *PLoS One*. 10(2): e0118321.Doi:10.1371/journal.pone.0118321.
- [14] Tamvacakis, A.N., Senatore, A., Katz P.S., (2015) Identification of genes related to learning and memory in the brain transcriptome of the mollusc, *Hermisenda crassicornis*. *Learning and Memory*. 22(12):617-21. PMID:26572652, <http://learnmem.cshlp.org/content/22/12/617>
- [15] The Sloan Digital Sky Survey: Mapping the Universe. 2014 [cited 2014; SDSS-IV:[Available from: <http://www.sdss.org/>]
- [16] Towns, J., Cockerill, T., Dahan, M., Foster, I., Gaither, K., Grimshaw, A., Hazlewood, V., Lathrop, S., Lifka, D., Peterson, G. D., Roskies, R., Scott J.R., Wilkins-Diehr, N., "XSEDE: Accelerating Scientific Discovery", *Computing in Science & Engineering*, vol.16, no. 5, pp. 62-74, Sept.-Oct. 2014, doi:10.1109/MCSE.2014.80