Improving Science Gateway Monitoring, Caveats and Goals

Jeanette Sperhac
Center for Computational Research
University at Buffalo
Buffalo, NY 14203
Email: jsperhac@buffalo.edu

Joseph P. White
Center for Computational Research
University at Buffalo
Buffalo, NY 14203
Email: jpwhite4@buffalo.edu

Ben Plessinger
Center for Computational Research
University at Buffalo
Buffalo, NY 14203
Email: bpless@buffalo.edu

Robert L. DeLeon
Center for Computational Research
University at Buffalo
Buffalo, NY 14203
Email: rldeleon@buffalo.edu

Matthew D. Jones
Center for Computational Research
University at Buffalo
Buffalo, NY 14203
Email: jonesm@buffalo.edu

Thomas R. Furlani
Roswell Park
Comprehensive Cancer Center
Buffalo, NY 14203
Email: Thomas.Furlani@roswellpark.org

***Abstract*—The U.S. National Science Foundation (NSF) has invested heavily in research computing, including funding the eXtreme Digital (XD) program, to make computing infrastructure available to researchers, and sponsoring the Science Gateways Community Institute (SGCI), to support the integration of supercomputers with science gateways. XD Metrics on Demand (XDMoD), an NSF-funded tool that collects and presents detailed data about computing resources, was developed to manage just such varied computational resources. XDMoD metrics incorporate accounting and performance data for computational jobs, and include resources consumed (computation, memory, disk, network, etc.), wait times, and quality of service. In recent years, XDMoD has been enhanced to provide monitoring for cloud computing and storage, and metrics such as power consumption and GPU support. As the use of science gateways becomes more prominent in research computing, the influence of these user-friendly portals grows. New features under development now in XDMoD will serve the gateways community. In this short paper, we outline our plans to introduce new gateways monitoring functionality that integrates with existing monitoring and metrics.**

***Index Terms*—science gateways, high-performance computing, monitoring**

I. Introduction

XDMoD is a comprehensive tool designed to facilitate the management of computational resources. To do this, XDMoD gathers data from many different data sources. This includes accounting information about HPC jobs from resource managers, performance data from cluster monitoring software, allocation data from resource allocation databases, and virtual machine (VM) provisioning data from cloud computing platforms. Examples of XDMoD metrics include accounting and performance data for computational jobs such as resource usage, wait times, power consumption, [1] and quality of service. XDMoD can provide information on individual jobs as well as data aggregated over an ensemble of jobs. Its web interface offers charting, exploration, and reporting of these metrics, for user-selected time ranges, and across all resources; refer to Figure 1.

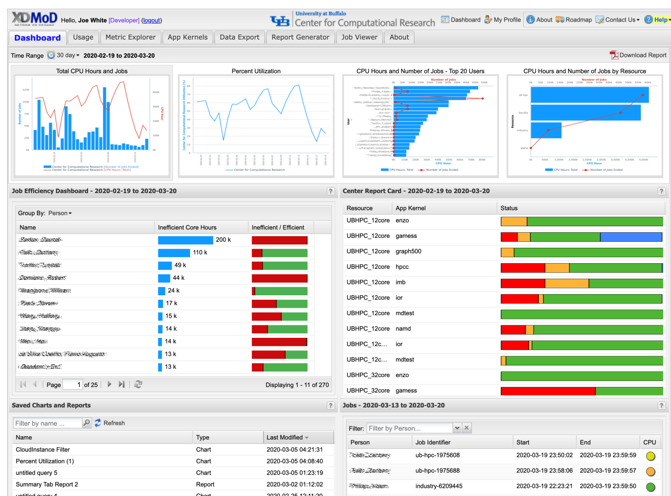

Fig. 1. The XDMoD web interface enables users to plot and explore numerous metrics. This summary page displays the efficiency of recent jobs, the health of the computing center, cluster status, current usage charts, and so forth, tailored to the role of the logged-in user.

Science gateway administrators and end users can use XDMoD to view detailed accounting and performance metrics about HPC jobs. However, while XDMoD provides detailed information about jobs from the point of view of the HPC resource, the current version of the tool does not incorporate comprehensive gateway-specific accounting information. The XDMoD team is actively integrating gateway job metadata into the tool to enhance the experience of gateway end users,

This work is supported by the National Science Foundation under grant numbers OCI 1025159 (Technology Audit Service), OCI 1203560 (SUPReMM), and ACI 1445806 (XD Net Metrics Service for High Performance Computing).

Presented at Gateways 2020, Online, USA, October 12–23, 2020. https://osf.io/meetings/gateways2020/

gateway administrators, and resource administrators alike.

In this paper, we describe prior work on HPC and research computing monitoring (Section II); provide an overview of science gateways and their monitoring (Section III); explain the rationale for the present work, including current caveats (Section IV); describe the proposed features of expanded science gateways monitoring (Section V); and mention potential future work (Section VI).

II. Prior Work

The XDMoD tool is designed to facilitate the management of high-performance computing (HPC) systems and associated cloud computing resources, storage resources, and more. For ten years, the evolving XDMoD [2] tool has been used to monitor NSF's XSEDE portfolio; it is also available in an open-source version for monitoring of any computing resource, public or private [3]. The tool offers support for job accounting metrics, as well as collection and charting of job-level performance data [4] and quality-of-service metrics known as application kernels [5], [6]. XDMoD has a modular architecture, and supports federation of multiple instances [7]. In addition to assisting with the management of computing resources, this tool has also been used to perform major workload analyses [8], [9], and studies of node sharing on HPC clusters [10] and HPC usage patterns [11]. Additionally, XDMoD already offers numerous features useful to the Gateways community; their use has been discussed elsewhere [12].

The monitoring of HPC systems is a topic of great interest in research computing circles. Many tools exist for monitoring job performance and tracking job-level activity on HPC clusters including Ganglia [13], Lightweight Distributed Metric Service (LDMS) [14], [15], Performance Co-Pilot (PCP) [16], TACC_stats [17], and XALT [18]. These tools work primarily by tracking hardware and OS performance counters on the compute nodes. XDMoD is able to ingest data from such performance monitoring tools. Performance data from PCP and TACC_Stats are actively supported by the XDMoD development team, while LDMS data have also been used as a data source for XDMoD [8]. The XDMoD tool combines data originating in different organizations and administrative domains. Grid platforms such as Globus Toolkit [19], Open Science Grid [20], and BOINC [21], among others, have confronted this challenge since the mid 1990s. Overall, one of XDMoD's strengths lies in its modular design, which allows integration of additional data sources and configuration for use on different systems, and provides extensible views tailored to a range of user roles. As architectures and practices evolve, XDMoD will adapt to track new metrics and offer new analysis tools.

III. Overview

Science gateways are not usually computational resources in themselves, but rather portals to these resources. They typically offer a web-based user interface that permits a specific research community to access a set of curated applications tailored to its needs. Science gateways provide a straightforward way for users, especially those in disciplines that are not traditionally heavily computational, to access and easily utilize high performance and cloud computing [22]. As non-traditional fields turn more to high performance computing (HPC), the use of gateways for submission of HPC jobs has increased substantially as a fraction of all research computing. A recent workload analysis of XSEDE indicated that science gateways usage increased fivefold between 2011 and 2017, and continues to grow [23]. To illustrate this, Figure 2 shows the growth in XSEDE job submissions by all gateways between 2015 and 2019. Blue bars show the annual counts of gateways jobs submitted to XSEDE resources. Red circles show the annual count of single-core jobs submitted to these resources over the same timeframe. The plot indicates that gateways submitted 5 times more jobs to these resources, and nearly 70 times more single-core jobs, in 2019 than in 2015.

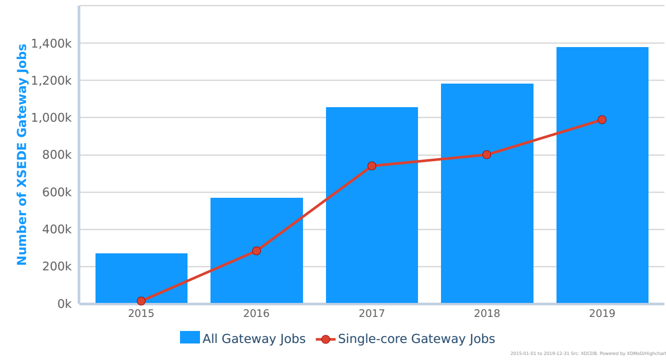

Fig. 2. This XDMoD plot displays the annual count of gateways jobs (blue bars) run on XSEDE resources between 2015 and 2019. Also shown are the annual count of single-core gateway jobs (red circles), indicating that the rapid growth in gateways jobs was due to single-core submissions.

Science gateways help make computing accessible to end user communities, but involve complexity under the hood. Typically, gateways run computing jobs on external computing resources. In order to do so, gateways must stage and transfer data, queue and run remote computations, then collect, transfer, and represent the results for the gateway end user to collect. Tracking and associating the different parts of this process, which can take place on computational resources located in multiple computing centers, time zones, or regions, makes gateway usage tracking a challenge. However, given the substantial recent gains in gateway utilization, a complete understanding of the usage of flagship and national compute resources should also incorporate usage by gateways.

High-level monitoring of gateways is important for ensuring access to gateways and the underlying computing resources, planning for future acquisition, and other concerns. Some support already exists for monitoring gateways. In particular, our NSF-funded project, XDMoD, provides limited support for monitoring gateways-initiated jobs that are run on HPC resources. In a previous presentation we described how the existing features of XDMoD enable examination of gateways jobs, including such metrics as counts, sizes, durations, and resource consumption of gateway jobs submitted to different computing resources. The application also supports, for XSEDE resources only, limited reporting on counts of new and unique gateway end users [12].

However, access to the gateways metrics may not be obvious to the casual XDMoD user, and gateways jobs can be difficult to locate using the search and reporting functions. Gateway end users may not have accounts on XDMoD, and thus lack access to the data describing their jobs. Furthermore, XDMoD user accounts are not linked to gateway end user names, meaning that gateways end users lack access to advanced metrics, such as performance metrics describing their individual jobs. With this presentation we describe XDMoD development in progress now that will put gateways data, monitoring, and metrics on a more equal footing with those from other computational jobs.

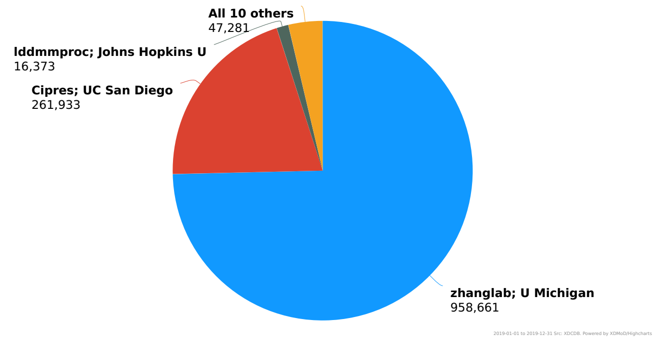

Fig. 3. This XDMoD plot displays the count of gateways jobs run on the SDSC Comet resource in 2019, by gateway name. Gateways shown are Zhanglab (blue, 958,661 jobs), Cipres (red, 261,933 jobs), and Lddmmproc, or Computational Anatomy gateway (grey, 16,373 jobs). All other gateways combined submitted 47,281 jobs to the resource.

IV. Rationale

Gateways are frequently used to implement scientific workflows, which execute a series of computational or data manipulation steps. Workflows may consist of automated software pipelines, or of loosely associated tasks executed one after the other. Workflows typically entail the use of widely distributed computing resources ranging from storage to compute nodes. These may enable the processing and analysis of data from large-scale scientific experiments. Workflows can help to organize, streamline, and document the scientific process, and can be shared along with datasets, contributing to reproducibility, research transparency, and dissemination of research results and methods. Gateways assist in these ends by providing locations where scientific communities stage and perform workflow computations and share datasets. [24]

Furthermore, web-based science gateways are frequently examples of workflows. An example: On the I-TASSER (Zhanglab) Gateway [25], the process begins when a user accesses the Zhanglab webserver to submit input data, which is validated and then copied over a network for further processing on a compute cluster. Next, the dataset is prepared and submitted with an HPC job to the XSEDE Comet resource. Once the HPC job is complete, the results are copied back to the gateway, where the final workflow output is prepared. Finally, the user is sent an informational email, and can then access the webserver and download the results. Each of these steps consists of multiple computational and data processing tasks. Each must be linked back to the submitted workflow job and to the originating gateway.

XDMoD is instrumented to collect metrics on many aspects of high-performance computing, but it does not comprehensively collect data on science gateways. One of the obstacles is that any individual gateway may perform many disparate tasks on a variety of resources (or contain workflows that do the same). Each such task must be tracked and associated with the others that form the gateway submission; when job metadata and origin data are missing, these associations cannot be made.

For example, numerous jobs submitted by XSEDE gateways lack metadata that could associate them to their gateway end users. In the case of monitoring bare metal cloud computing usage, tracing the resources consumed and the details of the job executed can be challenging, since cloud instances are not typically instrumented for such data collection. Resource usage by gateways on computational infrastructure such as webservers and databases is also difficult to acquire. However, even the gateway jobs data from traditional HPC resources are often acquired asynchronously, or associated with an unexpected resource username, so that the association to a gateways job can be missed.

Another obstacle has to do with user membership. Science gateways are useful portals to HPC and other computing resources, but gateway end users do not map one-to-one with usernames active on those resources. Gateway usage on these HPC resources can be traced in part because gateways tend to use a dedicated proxy user for submitting all their jobs to computing resources. That single system username is then used to run all jobs for a given gateway. When it comes time for gateway end users to review the metrics describing their jobs, they may have difficulty identifying the exact jobs they submitted, and even gaining access to the job details. The acquisition and linking of job metadata indicating which jobs are tied to gateway usage, and which gateway end users and applications are involved, is necessary to connect these dots. We are presently undertaking to improve these metadata pipelines, correct existing associations with gateways users, resources, and jobs, accommodate asynchronous data arrivals, and thus improve the gateways data we collect and the metrics we offer.

V. Proposed Features

XDMoD already provides limited support for monitoring gateways-initiated jobs run on HPC resources. This information can be found in the application's Jobs realm, where gateway job counts, job sizes, and resource utilization by gateway can be monitored. An example is shown in Figure 3, which shows the total number of jobs run by science gateways on the SDSC Comet resource in all of 2019. For the XSEDE resource portfolio, additional gateway metrics can be found in the Custom Queries, which allow scrutiny of the counts of new and unique gateway end users, among others. Refer to our previous presentation for further details about these features [12].

However, there are substantial gaps in XDMoD's metrics offerings for gateways. For example, filtering and searching by gateway end user is not yet possible, and additionally, gateway end users lack access to data about the performance of their own individual jobs. Errors in our metadata processing pipeline have contributed to miscounts among the Custom Queries for gateways. Furthermore, access to many of the gateways-focused metrics is not obvious, and the gateway metrics in the Custom Queries do not offer the same flexibility as those in other parts of the system. Therefore, we are now developing a set of gateways-specific metrics in a prominent location within XDMoD, and revisiting our processing of gateways metadata. This will enable us to correctly monitor gateways' use of remote resources (cloud, HPC) in an easily accessible way.

We are working on several enhancements that address these shortcomings directly. Expanding the gateways metadata collected by XDMoD is central to this work. For instance, we are undertaking improved tracking of jobs that gateways submit to HPC resources; introducing new gateway-focused metrics, such as CPU Hours by Gateway; and offering the ability to search and monitor jobs by gateway end user. As a part of this work, existing gateways data will be reinterpreted and re-aggregated to provide correct historical totals.

This gateway support will be added to both the XSEDE and open source versions of XDMoD. The XSEDE version of XDMoD presently pulls accounting data from XSEDE's accounting database. This already includes metadata provided by various gateways and cross referenced to the corresponding HPC jobs. The open source XDMoD version must obtain gateway accounting information directly from the gateways. Since it is impractical for XDMoD to directly support the many hundreds of different gateway implementations, the approach we will take is to define a standard set of information and a standard data interchange format. We plan to provide reference implementations for some common gateway software, and will provide support for gateway software developers, enabling them to produce gateway accounting information in the standard format. The gateway accounting logs must provide sufficient information to allow XDMoD to correctly associate the gateway-specific metadata with the corresponding HPC job. The job metadata we require from gateways, much as XSEDE requires from its member gateways, is as follows [26]:

- gateway proxy user's username
- gateway end user's username
- local gateway job id
- resource manager job id
- numeric id of gateway end user
- name of submitted script or executable
- cluster or resource name
- submission time and timezone

Gateway managers, when associated to their gateway as a Principal Investigator (PI), can already use XDMoD to monitor submitted gateway jobs, even down to performance details. However, this fact is largely hidden by the user interface. Simple enhancements to the user interface are planned that will present gateway-oriented XDMoD users with a collection of the relevant metrics as well as summaries of recent activity. Refer to Figure 1, which displays a user dashboard tailored to a HPC center staff role. Such a dashboard for a gateway manager could include usage metric charts pertinent to the gateway; information about the efficiency of recent gateway jobs; resource utilization by top gateway end user consumers; application usage; and the current status of the computing resources used by the gateway.

VI. Future Work

Several desirable features for gateways will not be implemented in the current revision of XDMoD. For example, enabling gateway end users access to the performance data for their own jobs will be added in the future.

On-gateway metric support is another. Though substantial gateways usage is submitted to off-gateway resources for processing, gateways also host computational tools that are run directly on the gateway. Most gateway platorms (HUBzero, Apache Airavata, Tapis, etc.) accumulate usage statistics for on-gateway user activity. [27]–[29] Centralized tracking of this on-gateway usage within XDMoD would involve the ingestion and standardization of on-gateway statistics, including user counts, CPU hours, job counts, individual tool, script, and executable use, and so forth.

Additional metrics could further enhance our picture of the infrastructure used and resources consumed by gateways. In particular, network usage incurred by transferring large data sets; better characterization of gateways workflows; on-gateway resource, tool, and computing usage; and resource usage on non-HPC computational infrastructure, such as web-servers, cloud resources, and databases; would further describe gateways job submissions.

VII. Conclusion

XDMoD is a tool that mines log files from computational resources to provide a multitude of job metrics describing resource usage and performance, including number of jobs, CPU hours, wall times, wait times, power consumption, FLOPS and job sizes. These metrics may be examined for a single job or aggregated over any desired time frame, in order to monitor and report collective job activity on a resource. XDMoD supports functions such as filtering, grouping and drill-down for these data, and reporting capabilities that include charting, data export, and custom report generation.

In this paper we have discussed the existing XDMoD support for gateways and indicated some areas where improvements are required. We have outlined enhancements to the XDMoD data acquisition pipeline, user interface, and metrics, and indicated the direction that future development will take to support improved monitoring for the science gateway community.

Acknowledgment

The authors would like to thank the entire XDMoD development team, including Cynthia Cornelius, Greg Dean, Steven M. Gallo, Jeffrey T. Palmer, Ryan Rathsam, Nikolay Simakov, and former members Rudra Chakraborty, Ryan Gentner, Amin Ghadersohi, Martins Innus, and Thomas Yearke.

This work was sponsored by the National Science Foundation (NSF) under award NSF ACI 1445806 for the XD Metrics Service (XMS).

References

[1] J. White, M. Innus, R. DeLeon, M. Jones, and T. Furlani, "Monitoring and Analysis of Power Consumption on HPC clusters using XDMoD," in *Proceedings of the Practice and Experience in Advanced Research Computing*, ser. PEARC '20. New York, NY, USA: Association for Computing Machinery, 2020, awarded Best Paper. [Online]. Available: https://doi.org/10.1145/3311790.3396624

[2] J. C. Browne, R. L. DeLeon, A. K. Patra, W. L. Barth, J. Hammond, M. D. Jones, T. R. Furlani, B. I. Schneider, S. M. Gallo, A. Ghadersohi, R. J. Gentner, J. T. Palmer, N. Simakov, M. Innus, A. E. Bruno, J. P. White, C. D. Cornelius, T. Yearke, K. Marcus, G. von Laszewski, and F. Wang, "Comprehensive, open-source resource usage measurement and analysis for HPC systems," *Concurrency and Computation: Practice and Experience*, vol. 26, no. 13, pp. 2191–2209, 2014, cPE-14-0027.R1. [Online]. Available: https://dx.doi.org/10.1002/cpe.3245

[3] J. T. Palmer, S. M. Gallo, T. R. Furlani, M. D. Jones, R. L. DeLeon, J. P. White, N. Simakov, A. K. Patra, J. M. Sperhac, T. Yearke, R. Rathsam, M. Innus, C. D. Cornelius, J. C. Browne, W. L. Barth, and R. T. Evans, "Open XDMoD: A tool for the comprehensive management of high-performance computing resources," *Computing in Science and Engineering*, vol. 17, no. 4, pp. 52–62, 2015. [Online]. Available: https://dx.doi.org/10.1109/MCSE.2015.68

[4] Center for Computational Research, University at Buffalo, "XDMoD Performance (SUPReMM) Module," http://supremm.xdmod.org/, 2018.

[5] N. A. Simakov, J. P. White, R. L. DeLeon, A. Ghadersohi, T. R. Furlani, M. D. Jones, S. M. Gallo, and A. K. Patra, "Application kernels: HPC resources performance monitoring and variance analysis," *Concurrency and Computation: Practice and Experience*, vol. 27, no. 17, pp. 5238–5260, 2015, cPE-14-0402.R1. [Online]. Available: https://dx.doi.org/10.1002/cpe.3564

[6] T. R. Furlani, M. D. Jones, S. M. Gallo, A. E. Bruno, C.-D. Lu, A. Ghadersohi, R. J. Gentner, A. K. Patra, R. L. DeLeon, G. von Laszewski, L. Wang, and A. Zimmerman, "Performance Metrics and Auditing Framework Using Applications Kernels for High Performance Computer Systems," *Concurrency and Computation: Practice and Experience*, vol. 25, no. 7, pp. 918–931, 2013. [Online]. Available: https://dx.doi.org/10.1002/cpe.2871

[7] J. Sperhac, B. Plessinger, J. Palmer, R. Chakraborty, G. Dean, M. Innus, R. Rathsam, N. Simakov, J. White, T. Furlani, S. Gallo, R. L. DeLeon, M. D. Jones, C. Cornelius, and A. K. Patra, "Federating XDMoD to monitor affiliated computing resources," in *Proceedings of the 2018 IEEE International Conference on Cluster Computing (IEEE-Cluster)*, IEEE. Piscataway, NJ, USA: IEEE Press, 2018, pp. 548–557, doi:10.1109/CLUSTER.2018.00074.

[8] J. P. White, M. Innus, M. D. Jones, R. L. DeLeon, N. Simakov, J. T. Palmer, S. M. Gallo, T. R. Furlani, M. Showerman, R. Brunner, A. Kot, G. Bauer, B. Bode, J. Enos, and W. Kramer, "Challenges of Workload Analysis on Large HPC Systems: A Case Study on NCSA Blue Waters," in *Proceedings of the Practice and Experience in Advanced Research Computing 2017 on Sustainability, Success and Impact*, ser. PEARC17. New York, NY, USA: ACM, 2017, pp. 6:1–6:8. [Online]. Available: http://doi.acm.org/10.1145/3093338.3093348

[9] T. R. Furlani, B. I. Schneider, M. D. Jones, J. Towns, D. L. Hart, S. M. Gallo, R. L. DeLeon, C. Lu, A. Ghadersohi, R. J. Gentner, A. K. Patra, G. Laszewski, F. Wang, J. T. Palmer, and N. Simakov, "Using XDMoD to facilitate XSEDE operations, planning and analysis," in *Proceedings of the Conference on Extreme Science and Engineering Discovery Environment: Gateway to Discovery (XSEDE '13)*. ACM, 2013, p. 8.

[10] J. P. White, R. L. DeLeon, T. R. Furlani, S. M. Gallo, M. D. Jones, A. Ghadersohi, C. D. Cornelius, A. K. Patra, J. C. Browne, W. L. Barth, and J. Hammond, "An analysis of node sharing on HPC clusters using XDMoD/TACC_Stats," in *Proceedings of the 2014 Annual Conference on Extreme Science and Engineering Discovery Environment*, ser. XSEDE '14. New York, NY, USA: ACM, 2014, pp. 31:1–31:8. [Online]. Available: https://dx.doi.org/10.1145/2616498.2616533

[11] R. L. DeLeon, T. R. Furlani, S. M. Gallo, J. P. White, M. D. Jones, A. K. Patra, M. Innus, T. Yearke, J. T. Palmer, J. M. Sperhac, R. Rathsam, N. A. Simakov, G. von Laszewski, and F. Wang, "TAS view of XSEDE users and usage," in *Proceedings of the 2015 XSEDE Conference: Scientific Advancements Enabled by Enhanced Cyberinfrastructure*, ser. XSEDE '15. New York, NY, USA: ACM, 2015, pp. 21:1—-21:8. [Online]. Available: https://dx.doi.org/10.1145/2792745.2792766

[12] J. M. Sperhac, R. L. DeLeon, T. R. Furlani, S. M. Gallo, M. Innus, M. D. Jones, J. T. Palmer, A. Patra, B. D. Plessinger, R. Rathsam, N. Simakov, J. P. White, R. Chakraborty, and G. Dean, "Managing computational gateway resources with XDMoD," *Future Generation Computer Systems*, vol. 98, pp. 154 – 166, 2019. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0167739X18310732

[13] M. L. Massie, B. N. Chun, and D. E. Culler, "The ganglia distributed monitoring system: design, implementation, and experience," *Parallel Computing*, vol. 30, no. 7, pp. 817 – 840, 2004. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0167819104000535

[14] A. Agelastos, B. Allan, J. Brandt, P. Cassella, J. Enos, J. Fullop, A. Gentile, S. Monk, N. Naksinehaboon, J. Ogden, M. Rajan, M. Showerman, J. Stevenson, N. Taerat, and T. Tucker, "The lightweight distributed metric service: A scalable infrastructure for continuous monitoring of large scale computing systems and applications," in *SC14: International Conference for High Performance Computing, Networking, Storage and Analysis*, Nov 2014, pp. 154–165.

[15] Sandia National Laboratories, Albuquerque, NM, "OVIS (LDMS)," https://github.com/ovis-hpc/ovis/wiki, accessed 16 July 2018.

[16] "Performance Co-Pilot," https://pcp.io/, accessed 16 July 2018.

[17] T. Evans, W. L. Barth, J. C. Browne, R. L. DeLeon, T. R. Furlani, S. M. Gallo, M. D. Jones, and A. K. Patra, "Comprehensive resource use monitoring for HPC systems with TACC_Stats," in *Proceedings of the First International Workshop on HPC User Support Tools*, ser. HUST '14. Piscataway, NJ, USA: IEEE Press, 2014, pp. 13–21. [Online]. Available: https://dx.doi.org/10.1109/HUST.2014.7

[18] K. Agrawal, M. R. Fahey, R. McLay, and D. James, "User environment tracking and problem detection with XALT," in *Proceedings of the First International Workshop on HPC User Support Tools*, ser. HUST '14. Piscataway, NJ, USA: IEEE Press, 2014, pp. 32–40. [Online]. Available: https://dx.doi.org/10.1109/HUST.2014.6

[19] I. Foster, "Globus Toolkit Version 4: Software for Service-Oriented Systems," in *Proceedings of the 2005 IFIP International Conference on Network and Parallel Computing*, ser. NPC'05. Berlin, Heidelberg: Springer-Verlag, 2005, p. 2–13. [Online]. Available: https://doi.org/10.1007/11577188_2

[20] M. Altunay, P. Avery, K. Blackburn, B. Bockelman, M. Ernst, D. Fraser, R. Quick, R. Gardner, S. Goasguen, T. Levshina, M. Livny, J. McGee, D. Olson, and R. Pordes, "A Science Driven Production Cyberinfrastructure—the Open Science Grid," *Journal of Grid Computing*, vol. 9, p. 201–218, 2011. [Online]. Available: https://doi.org/10.1007/s10723-010-9176-6

[21] D. P. Anderson, "BOINC: A System for Public-Resource Computing and Storage," in *Proceedings of the 5th IEEE/ACM International Workshop on Grid Computing*, ser. GRID '04. USA: IEEE Computer Society, 2004, p. 4–10. [Online]. Available: https://doi.org/10.1109/GRID.2004.14

[22] N. Wilkins-Diehr, "Measuring success: How science gateways define impact," Oct 2019. [Online]. Available: osf.io/tkzuy

[23] N. A. Simakov, J. P. White, R. L. DeLeon, S. M. Gallo, M. D. Jones, J. T. Palmer, B. D. Plessinger, and T. R. Furlani, "A Workload Analysis of NSF's Innovative HPC Resources Using

XDMoD," *CoRR*, vol. abs/1801.04306, 2018. [Online]. Available: http://arxiv.org/abs/1801.04306

[24] J. Sperhac, R. DeLeon, J. White, M. Jones, A. Bruno, R. Jones-Ivey, T. Furlani, J. Bard, and V. Chaudhary, "Towards Performant Workflows, Monitoring and Measuring," in *Proceedings of the 29th International Conference on Computer Communications and Networks (ICCCN 2020)*, ser. IEEE ICCCN '20. New York, NY, USA: IEEE Computer Society, 2020, invited talk.

[25] W. Zheng, C. Zhang, E. W. Bell, and Y. Zhang, "I-TASSER gateway: A protein structure and function prediction server powered by XSEDE," *Future Generation Computer Systems*, vol. 99, pp. 73 – 85, 2019. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0167739X18314705

[26] "XSEDE Gateways Metadata API," November 2019. [Online]. Available: https://xsede-xdcdb-api.xsede.org/api/

[27] M. Mclennan and R. Kennell, "HUBzero: A platform for dissemination and collaboration in computational science and engineering," *Computing in Science and Engineering*, vol. 12, no. 2, pp. 48–52, 2010.

[28] S. Marru, L. Gunathilake, C. Herath, P. Tangchaisin, M. Pierce, C. Mattmann, R. Singh, T. Gunarathne, E. Chinthaka, R. Gardler, A. Slominski, A. Douma, S. Perera, and S. Weerawarana, "Apache airavata: A framework for distributed applications and computational workflows," in *Proceedings of the 2011 ACM Workshop on Gateway Computing Environments*, ser. GCE '11. New York, NY, USA: Association for Computing Machinery, 2011, p. 21–28. [Online]. Available: https://doi.org/10.1145/2110486.2110490

[29] S. B. Cleveland, A. Jamthe, S. Padhy, J. Stubbs, M. Packard, J. Looney, S. Terry, R. Cardone, M. Dahan, and G. A. Jacobs, "Tapis API Development with Python: Best Practices In Scientific REST API Implementation: Experience Implementing a Distributed Stream API," in *Practice and Experience in Advanced Research Computing*, ser. PEARC '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 181–187. [Online]. Available: https://doi.org/10.1145/3311790.3396647