

Innovations from the early user phase on the Jetstream Research Cloud

Richard Knepper, Jeremy Fischer, Craig Stewart, David Hancock and Matthew Link
Pervasive Technology Institute
Indiana University, Bloomington, Indiana 47408
Email: rich@iu.edu

Abstract—We describe the Jetstream cyberinfrastructure for research, a purpose-built system with the goal of supporting “long-tail” research by providing a flexible infrastructure that can provide a set cloud services tuned for research applications, whether they be traditional HPC applications, science gateways, or desktop applications. Jetstream offers a library of virtual machines and allows the user to create their own virtual machines in order to provide an open cloud for science that allows both on-demand and persistent instances. The system is currently in early-user mode and a number of users at partner institutions are already creating and using images in the system. This paper details some of the early work being done with the system to create high performance clusters in an on-demand fashion to support scientific work directly as well as serve as capability backend to scientific gateways such as CyVerse and Galaxy.

I. INTRODUCTION

THE JETSTREAM SYSTEM is designed to provide a production cloud resource in support of general science and engineering activities in the eXtreme Digital (XD) ecosystem. While the United States National Science Foundation (NSF) has funded a number of high performance computing (HPC) as well as high throughput computing (HTC) resources, but there remains a significant population of researchers who have computational and data analysis needs that neither HPC nor HTC resources are fit[1]. The NSF has noted [4]the benefits of increasing diversity in the range of cyberinfrastructure (CI) resources available to researchers. Other efforts on the part of the NSF to provide more support for research in the “long tail of science” include the Comet[2] system and the Wrangler data storage and data analytics system[3].

The Jetstream resource and planned activities are well-described in [1]. The current system has been implemented at Indiana University (IU) and Texas Advanced Computing Center (TACC) and is currently in early user mode. In this paper we review the Jetstream architecture and software environment and detail early user experiences with the system, including innovative use of Apache Mesos software for building on-demand cluster resources, using the Atmosphere software suite to serve as a capability service for science gateways, and a “desktop mode” for scientific computing in the field or at sites with limited resources. We discuss some of the challenges of managing a multi-zoned research cloud system in a seamless fashion. We conclude the paper with a discussion of future activities with the service.

II. JETSTREAM OVERVIEW

The Jetstream system is designed to provide general purpose cloud resources for research in a configurable fashion. The system provides on-demand and persistent virtual systems that support a wide range of scientific software in the form of configurable environments.

A. Service Functions

The Jetstream system as designed supports multiple modalities of use in support of scientific research that are currently not provided within the broader cyberinfrastructure ecosystem. These include: self-serve academic cloud services, based on virtual machines images provided by the user or selected from a library; persistent virtual machine systems which support the delivery of science gateways such as the Galaxy [5] life science research gateway; data movement via the Globus Connect and authentication via Globus Auth [6]; facilities for publishing and sharing virtual machine images via IU’s persistent digital repository, IUScholarWorks, accessible via Digital Object Identifier (DOI)[7]; and provide virtual desktop services to institutions with limited resources.

Within the cyberinfrastructure ecosystem, the Open Science Grid, the Extreme Science and Engineering Discovery Environment (XSEDE), and other projects provide a broad range of resources for computational support of research: high-performance computing at large scale, high-memory resources, big data, and visualization systems, but to date no system provides a highly configurable virtualized environment with the capabilities described above. The Jetstream system is designed to provide resources for the “long tail” of science [8], who frequently need more access to interactive computational resources than they have locally available, and provide interactive access to a handful of systems as needed in an on-demand fashion, rather than forcing them to work with an allocation system, as in XSEDE, or with a virtual organization, as in the Open Science Grid.

B. Hardware Configuration

The Jetstream system hardware architecture follows commercial cloud offerings in terms of uptime and availability. Two production systems, one at IU and one at TACC, provide a 100Gbps-linked, distributed cluster infrastructure. A third development and testing system resides at the University of Arizona. This configuration offers “zoning” between the two production centers, similar to that offered in Amazon EC2

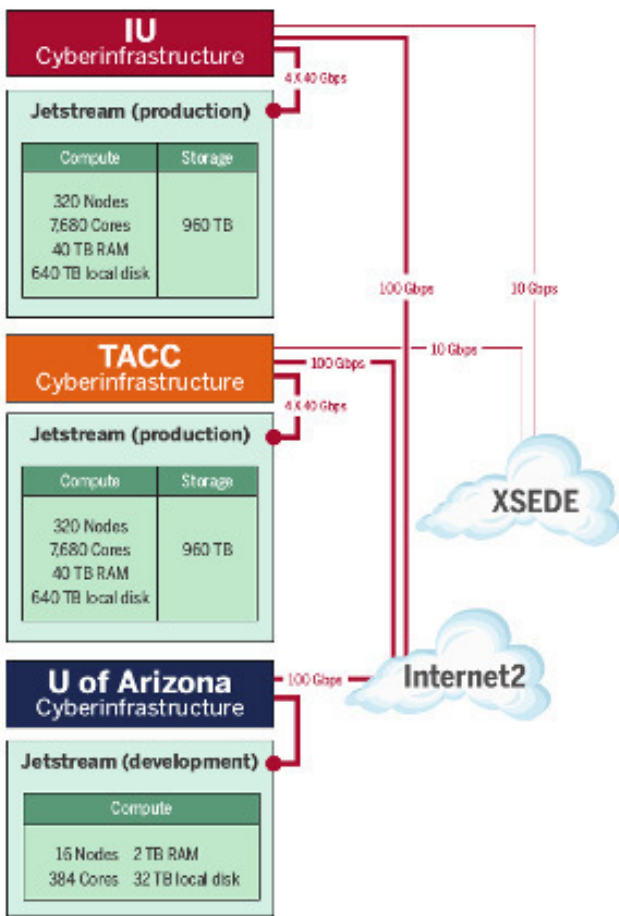


Fig. 1. Jetstream physical sites and network connectivity

resources and elsewhere. A schematic of the Jetstream system is presented in figure 1.

Jetstream provides multiple VM configurations, ranging from “Tiny”: 1 CPU, 2GB of memory, 20GB of storage, allowing as many as 46 concurrent virtual machine instances up to “XXL”: 44 CPUs, 120GB memory, 480GB storage, with one virtual instance. “Small”, “Medium”, “Large”, and “X-Large” configurations are also offered, with according sizes.

C. Software Architecture

Jetstream utilizes the Atmosphere[9] software stack for presenting a user interface, managing images, provisioning, monitoring and managing cloud infrastructure. Openstack provides host and virtual machine management, as well as virtual machine filesystem storage, with iRODS software providing replication between sites. Authentication is provided by Globus Auth, and user data transfer between the user’s desktop and virtual machine filesystems is provided by Globus Transfer. A diagram of the Atmosphere implementation on Jetstream is shown in figure 2.

Atmosphere offers a number of features essential to research computing in a cloud context. Identity management services, networking configuration, and security policies are

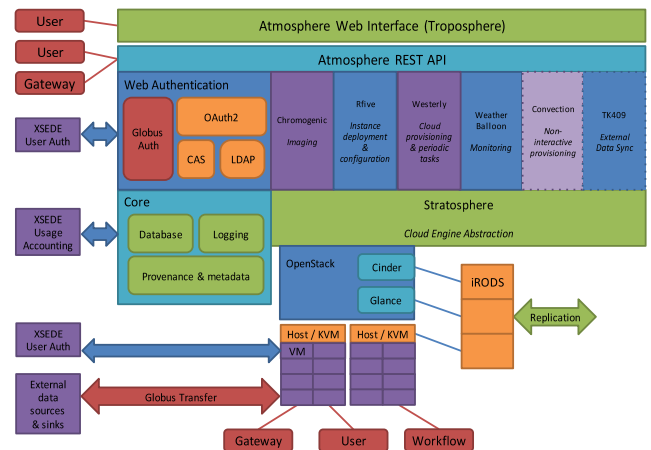


Fig. 2. Atmosphere Architecture in Jetstream

integrated with the software stack. The software provides complete functions for managing virtual machine instances throughout their lifecycle. Finally, data lifecycle management is simplified via the Atmosphere web portal and API. The web portal displays images that can be launched and worked with by the user, shared between users, and new images can be developed and uploaded. Once an image is launched it can be used interactively. A diagram of the Jetstream user interface is presented in figure 3.

III. MODALITIES OF USE

The flexibility of the Jetstream resource offers researchers a broad range of options for computational support of their research. During the initial implementation of Jetstream, the system was placed in “early-user” mode at the beginning of 2016 with the intention of providing the service to the Atmosphere development team at University of Arizona, as

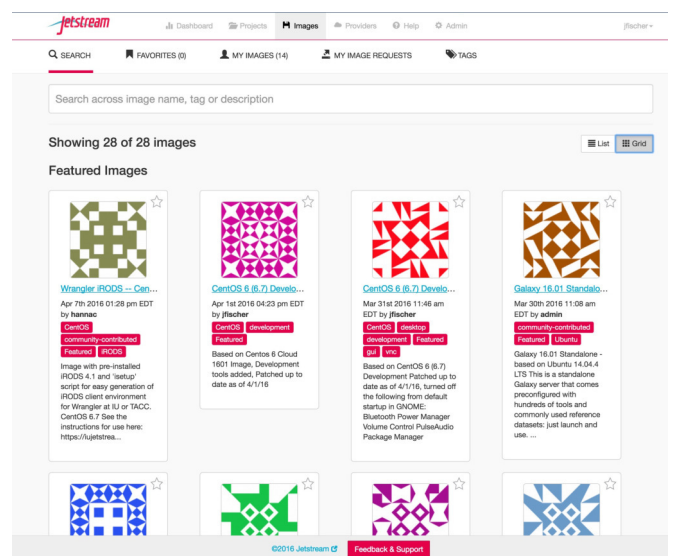


Fig. 3. Jetstream User Interface

well as other scientific users with some familiarity with cloud research in order to test the capabilities of Jetstream and determine what modalities of use might be possible and which ones users may prefer. In this section we describe some of the modalities of use established with these early system users.

A. Research Cloud Capabilities

As a research cloud resource, the Jetstream system must support individual usage of the system directly via interactive sessions with running virtual machine images, but also must support utility computing, particularly in support of science gateway services. In this scenario, users are able to choose workflows and data via a web-based science gateway and workflows are executed in a capability environment via middleware utilities. One of the common software frameworks for supporting science gateway computing is Apache Airavata [10]. In a similar fashion, utility computing resources can be utilized by providing high throughput computing capabilities. Three utility computing modalities for Jetstream have been established in the time since the system was introduced: support for CyVerse; the Galaxy life science gateway; and the ATLAS high energy physics project.

1) *CyVerse*: One of the initial project goals of the Jetstream system as proposed to the NSF was to improve the availability of compute resources to the CyVerse project¹ (previously known as the iPlant Collaboration). CyVerse was created in order to support life sciences research and improve access to existing cyberinfrastructure, and CyVerse created the Atmosphere cloud service which powers Jetstream. For the CyVerse project, Jetstream provides a set of on-demand system image toolkits, which can be instantiated and run to complete analyses, and then archived or repeated, with the image file available via DOI. Developers at Arizona and IU have created a number of system image toolkits for specific types of research. These toolkits include:

- 1) a general life science toolkit
- 2) an R toolkit
- 3) an astronomy toolkit
- 4) a data transfer toolkit, with iRODS interface to the Wrangler system at Texas Advanced Computing Center
- 5) a phylogenetics toolkit

The ability for Jetstream project users to create, save, and share toolkit system images provides a flexible means of collaborating on multi-researcher projects.

2) *Galaxy*: The Galaxy life science gateway service provides a comprehensive platform for genomic research supporting advanced data management capabilities with an intent to support reproducibility of analyses[5]. The Galaxy web service can utilize multiple different types of resources for analysis, and researchers can either make use of the main Galaxy portal at <http://usegalaxy.org> or they can set up their own Galaxy servers providing for specific community or lab requirements. Galaxy users at IU have created two means of supporting Galaxy with Jetstream. In the first, the user is able to instantiate a persistent virtual machine image and run a local Galaxy gateway on Jetstream, either completing analyses within the

same virtual machine or incorporating other virtual machines in Jetstream to provide additional computational capability. Jetstream VM's are configured as resources for the Galaxy Cloud cyberinfrastructure management tool. Jetstream VM's are with images ready to receive jobs and create a local data caches. The main Galaxy system distributes jobs to VMs via slurm workload manager with pulsar, Galaxy's remote execution service. At this time, about 4 months after Jetstream was initially opened to early use, 7,299 Galaxy main jobs have been run on Jetstream, completing work requests from 758 distinct users.

In the second path, Jetstream system images with Galaxy server deployed are launched within Jetstream and either use their own local resources or submit jobs to Galaxy Cloud. This latter on-demand service provides an easily created and archived Galaxy instance for short to medium term analyses that can be archived in the IU Scholarworks system and retrieved via DOI for the purposes of replication or further analysis.

3) *ATLAS*: The ATLAS experiment at the Large Hadron Collider utilizes considerable computational resources via the Open Science Grid, many at participating sites. Open Science Grid jobs are well-prepared to take advantage of cloud resources, as Open Science Grid software is able to provide "glide-in" capabilities which allow for jobs to be distributed to resources via a factory. The Jetstream resource is able to support virtual machine images which are a simple base operating system with the software to accept glide-in submissions. ATLAS experiment users can submit jobs to their virtual organization Condor scheduler, which will start the virtual machines on Jetstream via the Atmosphere API. The jobs will be submitted to the virtual machine as glide-in jobs and managed by Open Science Grid monitoring and scheduling resources.

B. Innovative use of Jetstream as research cloud

Early users have also made inroads in optimizing the capability offered by configurable systems for research. Using the Apache Mesos cloud manager system, early users have demonstrated the allocation of virtual system hosts on Jetstream to be managed by Mesos. Mesos provides information on cpu and memory available, and supports a number of frameworks to aid in the management of scaling tools. In this instance, Marathon was used to long-running services in Docker. Within Mesos, users were able to start instances running Docker containers with Unidata IDV, a tool for analyzing and visualizing geoscience data. This early demonstration usage showed that Jetstream could be used as the service resource for a user-defined cluster management framework, instantiating resources and completing work within the temporary cluster, and then releasing those resources to the research cloud and archiving system image information for later use.

IV. OPERATING A RESEARCH CLOUD

A. Functionality tests and operational metrics

In order to establish functionality of the Jetstream system as proposed, a number of functional tests of the openstack environment, atmosphere API, and application software were conducted in order to demonstrate the system's viability as

¹<http://www.cyverse.org>

a distributed cloud platform for research activities. A large part of the functional activity tests of Jetstream focus on the set of tasks a user can manage for themselves in order to make use of the architecture, demonstrating a “self-service” system that allows the bulk of activities to take place on the end user’s initiative, without requiring the intervention of IT staff. The base set of use cases to articulate this are (assuming that a user has authorization via an allocation and has sufficient understanding of the system):

- Cloud functionality. A sufficiently authorized and knowledgeable user can:
 - authenticate to the Jetstream web interface
 - launch a virtual machine from a library of images on either cluster location
 - quiesce an image running on one cluster, move it to another cluster, and reactivate it
 - create and access a permanent cloud data storage space
 - modify a pre-existing image and store the changed image to the image library on either cluster
- Data management functionality. A sufficiently authorized and knowledgeable user can:
 - move a file from another system into the Jetstream system
 - select a file from Jetstream and move it to another system
 - save a VM image and upload it to IUScholarworks and receive a DOI for the object

The Jetstream system has demonstrated the ability of a user to perform all of these use cases as part of early user mode and acceptance testing for the NSF. Additionally the system has as part of testing activities demonstrated the ability to provide gateway services. These tests consisted of demonstrating that a virtual machine could be run in a continuous fashion, providing access to computational workflows via a standard web interface. In order to test this, the Galaxy and SEAgrid science gateway software were implemented on Jetstream virtual machines with the requirement of running workflows in a comparable amount of time as XSEDE resources. The results gateway services tests are detailed in table I.

TABLE I. GATEWAY FUNCTIONALITY TESTS OF THE JETSTREAM SYSTEM

Test	Success Criteria	Result achieved
Deliver Galaxy gateway	Workflow execution within $\leq 125\%$ of time required	Workflow completes in 80% of time to run on Stampede
Deliver additional gateway in XSEDE	Correct operation and availability within 2% of overall system availability over 14-day test period	SEAgrid gateway operated continuously for 14 days

Furthermore, the system had a number of operational goals to meet as part of early user stages. These goals include the availability of the system, system capacity, job completion, number of users, number of active VM’s, CPU utilization, and number of images published to IUScholarworks. The system met all but two of these goals during early user phase - number

of distinct users and active VM’s. The outcome of these tests is show in table II.

TABLE II. OPERATIONAL TESTS OF THE JETSTREAM SYSTEM

Metric	Goal	Achieved
System availability (uptime)	95%	100%
% Capacity available	95%	100%
Job completion success	96%	97.7%
Number of distinct users	1,000	327
Use - average number of active VMs	320	290 (1217 peak)
CPU % utilization	6%	4.2% (20.3% peak)
VM images published to IUScholarworks	10/year	6/quarter

B. Challenges of managing the Jetstream system

Along with this early-user utilization of Jetstream, a number of challenges for supporting the research cloud framework have been identified that remain to be addressed. Some of these issues result from the fact that Jetstream incorporates an installation at IU and one at TACC. While this provides a “zoning” function that allows for resources to be available in different location and retain functionality in the event of a large-scale network outage or system-wide maintenance, it also means that these two systems must be kept synchronized in order to ensure coherence for users across both systems. User accounts, managed by LDAP and Globus OAuth, must be synchronized in both places, which has been a manageable process with standard automation tools. Data sync across the clusters represents a more difficult challenge. In order for Jetstream to present a seamless interface, no matter what zone a user accesses, the image library must be synchronized across both cluster, which means that system images up to the XXL size (480GB allocated storage) may need to be copied across clusters. In addition, ensuring that user identification numbers (UIDs) and permissions are the same across both systems. Furthermore, Jetstream engineers are still working on a reliable means of transporting a system across the two clusters, so that a running virtual machine image can be quiesced, relocated, and restarted on a different cluster. This requires both fast data transfer between clusters as well as robust scripts for managing ownership and permissions in both places to ensure that the data can be accessed by the right people.

V. CONCLUSION

We have described some of the early successes and challenges faced by the Jetstream system in its first few months of operation. Jetstream demonstrates efficacy as a research cloud resource, in contrast to projects such as FutureGrid, which are largely used for the exploration of cloud technology and management software. Jetstream projects are capable of supporting recognized research workflows, carried out by researchers, without the aid of grad students or computer science consultants.

The next steps for work on Jetstream are the configuration of images and automation via the OpenStack API to create both persistent and dynamic resources for workflows initiated in the broader cyberinfrastructure (such as through the CyVerse portal or another existing portal) or initiated from gateways running within Jetstream itself. Further work on desktop-like access to Jetstream images will also benefit the Jetstream user base who need short-term interactive resources that can be suspended and revived as necessary.

ACKNOWLEDGMENTS

Implementation of Jetstream is supported by NSF award 1445604. The Indiana University Pervasive Technology Institute was established with support of funding from the Lilly Endowment, Inc., and provided support for this research. CyVerse is supported by NSF award DBI-0735191 and DBI-1265383. Any opinions expressed here are those of the authors and do not necessarily reflect the views of the National Science Foundation, or the Lilly Endowment.

REFERENCES

- [1] Stewart, Craig A., et al. "Jetstream: A self-provisioned, scalable science and engineering cloud environment." Proceedings of the 2015 XSEDE Conference: Scientific Advancements Enabled by Enhanced Cyberinfrastructure. 2015. ACM: St. Louis, MO, USA. p. 29-37.
- [2] Moore, R.L., C. Baru, D. Baxter, G. Fox, A. Majumdar, P. Papadopoulos, W. Pfeiffer, R.S. Sinkovits, S. Strande, M. Tatineni, R.P. Wagner, N. Wilkins-Diehr, M.L. Norman. Gateways to Discovery: Cyberinfrastructure for the Long Tail of Science. In Proceedings of the 2014 Annual Conference on Extreme Science and Engineering Discovery Environment. 2014. ACM: Atlanta, GA, USA. p. 1-8.
- [3] Texas Advanced Computing Center. Wrangler. <https://www.tacc.utexas.edu/-/wrangler-data-intensive-system-opens-to-scientists>
- [4] Committee on Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science in 2017-2020. 2014. Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science and Engineering in 2017- 2020: Interim Report. Washington, DC. The National Academies Press. 48 pp. <http://www.nap.edu/catalog/18972/future-directions-for-nsf-advanced-computing-infrastructure-to-support-us-science-and-engineering-in-2017-2020>
- [5] Goecks, J, A. Nekrutenko, J. Taylor, and The Galaxy Team. 2010. Galaxy: a comprehensive approach for supporting accessible, reproducible, and transparent computational research in the life sciences. *Genome Biol.* 11(8):R86. <http://genomebiology.com/2010/11/8/R86>
- [6] Foster I, Kesselman C. Globus: A metacomputing infrastructure toolkit. *International Journal of High Performance Computing Applications.* 1997 Jun 1;11(2):115-28.
- [7] Bobay J. Institutional Repositories: Why Go There?. *Indiana Libraries.* 2008 Jan 1;27(1):7-9
- [8] Heidorn, P.B. Shedding Light on the Dark Data in the Long Tail of Science. *Library Trends.* 2008. 57(2), p. 280-299. http://muse.jhu.edu/journals/library_trends/v057/57.2.heidorn.html
- [9] Skidmore, E., S.-j. Kim, S. Kuchimanchi, S. Singaram, N. Merchant, and D. Stanzione. iPlant Atmosphere: A Gateway to Cloud Infrastructure for the Plant Sciences. In Proceedings of the 2011 ACM workshop on Gateway computing environments. 2011, ACM: Seattle, Washington, USA. p. 59-64. <http://dl.acm.org/citation.cfm?id=2110495>
- [10] Marru S, Gunathilake L, Herath C, Tangchaisin P, Pierce M, Mattmann C, Singh R, Gunarathne T, Chinthaka E, Gardler R, Slominski A. Apache airavata: a framework for distributed applications and computational workflows. In Proceedings of the 2011 ACM workshop on Gateway computing environments 2011 Nov 18 (pp. 21-28). ACM.