

Jetstream – performance, early experiences, and early results

Craig A Stewart
Indiana University Pervasive
Technology Institute (IUPTI)
stewart@iu.edu

Jeremy Fischer
IUPTI
jeremy@iu.edu

Nirav Merchant
University of Arizona
nirav@email.arizona.edu

Daniel C. Stanzione
University of Texas at Austin
dan@tacc.utexas.edu

David Y. Hancock
IUPTI
dyhancoc@iu.edu

Tim Cockerill
University of Texas at Austin
cockerill@tacc.utexas.edu

Therese Miller
IUPTI
millertm@iu.edu

James Taylor
Johns Hopkins University
james@taylorlab.org

Matthew Vaughn
University of Texas at Austin
vaughn@tacc.utexas.edu

Lee Liming
University of Chicago
lliming@uchicago.edu

John Michael Lowe
IUPTI
jomlowe@iu.edu

Edwin Skidmore
University of Arizona
edwin@cyverse.org

ABSTRACT

Jetstream is a first-of-a-kind system for the NSF - a distributed production cloud resource. The NSF awarded funds to create Jetstream in November 2014. Here we review the purpose for creating Jetstream, present the acceptance test results that define Jetstream's key characteristics, describe our experiences in standing up an OpenStack-based cloud environment, and share some of the early scientific results that have been obtained by researchers and students using this system. Jetstream offers unique capability within the XSEDE-supported US national cyberinfrastructure, delivering interactive virtual machines (VMs) via the Atmosphere interface developed by the University of Arizona. As a multi-region deployment that operates as a single integrated system, Jetstream is proving effective in supporting modes and disciplines of research traditionally underrepresented on larger XSEDE-supported clusters and supercomputers. Already, researchers in biology, network science, economics, earth science, and computer science have used Jetstream to perform research – much of it research in the “long tail of science.”

CCS Concepts

• **Computer systems organization** ~ **Cloud computing**
• **Software and its engineering** ~ **Virtual machines** • *Applied computing* ~ *Life and medical sciences* • *Applied computing* ~ *Physical sciences and engineering*

Keywords

Cloud computing; Virtual machines; OpenStack; Long tail of science; National Science Foundation.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the Owner/Author. Copyright is held by the owner/author(s).

XSEDE16, July 17-21, 2016,
ACM 978-1-4503-4755-6/16/07.
<http://dx.doi.org/10.1145/2949550.2949639>



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

1. INTRODUCTION

For decades, the National Science Foundation (NSF) has provided the national research community with resources collectively known as cyberinfrastructure (CI): “computing systems, data storage systems, advanced instruments and data repositories, visualization environments, and people, all linked together by software and high performance networks to improve research productivity and enable breakthroughs not otherwise possible” [1]. These CI resources have been provisioned through grant solicitations and grant awards to organizations that deliver NSF-funded resources to the research community on behalf of the NSF. One of the most recent major solicitations was NSF - Program Solicitation 14-536 [2], which begins as follows:

The intent of this solicitation is to request proposals from organizations willing to serve as Resource [Service] Providers within the NSF eXtreme Digital (XD) program. The current solicitation is intended to complement previous NSF investments in advanced computational infrastructure by exploring new and creative approaches to delivering innovative computational resources to an increasingly diverse community and portfolio of scientific research and education.

When we in the cyberinfrastructure community think of what it means to support research “not otherwise possible,” we often think of applications that require thousands of CPUs or large amounts of memory or very fast networks in order to do some fantastically large data analysis or simulation. But “not otherwise possible” can also mean something very different: Not otherwise possible because CI resources are provisioned, presented, and supported in a way that does not makes it possible for researchers to use them.

As a first step in planning to deliver “innovative computational resources to an increasingly diverse community,” we interviewed dozens of researchers, mostly working in areas underrepresented among existing users of resources in the NSF XD program. (We follow NSF usage in using “XD Program” to refer collectively to XSEDE and the NSF-funded resource providers that manage and deliver resources allocated and supported by XSEDE under NSF direction.) On the basis of these interviews we defined two canonical use cases for domain scientists, and a number of use cases based on mode of use. All are designed to meet the needs of researchers who can be described generally as working in the “long

tail of science” [3] – a label that applies to many researchers working in areas supported by the NSF but not strongly represented in XD program usage. These use cases are described briefly below and in greater detail in [4].

Two generalizations of domain-based use cases. The most common use case we heard was: A researcher wants straightforward access to specific, usually interactive, tools to analyze data delivered in a manner congruent with their normal operations and often driven by availability of new data. A related and also common use case we heard, generally from scientists who are developing new software, was: A tool producer develops new analysis routines and methods to address research bottlenecks, and needs to make said tool available to experimentalists without having them contend with technical complexities of operating system and software dependencies. *Proposed solution:* Develop an accessible platform where application creators can easily publish and share within a VM image, and end users can easily invoke runnable instances of these applications via virtualization.

Among the use cases distinguished by mode of access or mode of use of CI resources the following affect the largest number of users:

Enable analysis of public data sets at small schools with limited CI budgets (including MSIs). The ability to do computationally-intensive research at small schools with constrained budgets is often limited by local network capability and lack of local CI resources. *Proposed solution:* Provide a Jetstream VM image featuring a user-friendly virtual Linux desktop that executes on Jetstream with screen images delivered to tablet devices on cellular connections or to older PCs on slow networks.

Facilitate reproducible data analyses. Enable reproducibility of data analyses and published research (as discussed in [5, 6]). *Proposed solutions:* Enable researchers to easily publish a VM containing their analysis tools – and, in the case of published research, the input data, scripts, and output data in a VM image identifiable and discoverable via a DOI (Digital Object Identifier).

Enhance ease of science gateway deployment. Science Gateways provide web-accessible implementations of particular analyses and scientific workflows. While straightforward to use, they can be labor intensive to create. Extensive server-side programming and integration of workflow orchestration tools are often required to make a science gateway work. More than a dozen research groups currently operate XSEDE science gateways, and nationally even more wish to do so [7]. *Proposed solution:* Provide a gateway builder’s toolkit, including VMs with commonly used workflow engines installed and ready to configure with XSEDE tools, coupled with a platform that supports persistently hosting web services.

1.1 Meeting the needs described in use cases

The use cases identified clearly called for a solution that can generally be described as “provide a handful of CPU cores to an end user now, whenever now is, interactively.” Even in terms of supporting science gateway deployment, the current challenges seem greater in provisioning the interactive front of gateways than the sometimes massive supercomputers that constitute behind-the-scenes resources. A cloud-based solution was obvious – not necessarily because the use cases span the set of existing functionality of cloud systems, but because cloud systems provide a straightforward way to deliver and manage VMs using open source software.

Jetstream is designed to deliver the services and programming models needed by researchers working in the “long tail of science.”

It is designed to deliver services in a way that is and is perceived to be easy, accessible, and valuable. In particular, Jetstream:

- Offers “self-serve” cloud services, enabling researchers or students to select a pre-existing VM image or to create a new virtual environment for personalized research computing
- Hosts persistent Science Gateways
- Enables data movement, storage, and dissemination
 - Jetstream supports data transfer with Globus Connect.
 - Users are able to store VMs in the Indiana University digital repository, IUScholarWorks, and make them discoverable with a Digital Object Identifier (DOI).
- Provides virtual desktop services to tablet devices, increasing CI access for users at resource-limited institutions (e.g., small schools, schools in EPSCoR states, and Minority Serving Institutions).

To foster adoption by scientists working largely in the long tail of science, who were not previously users of XD program resources, we took into account current understandings that suggest technology adoption is driven by: performance expectancy (perceived value), effort expectancy (perceived ease of use), social influence, and facilitating conditions (including knowledge of a technology and the belief that end users will find it accessible) [8]. We thus picked “deliver VMs interactively” as the primary function that Jetstream provides. The user interface and familiar software environments were selected primarily to address user views on the issues of perceived value and perceived ease of use and thus facilitate adoption of Jetstream. The Jetstream user interface is based on Atmosphere, designed by the University of Arizona to be intuitive, combined with powerful cloud service management and orchestration capabilities. Atmosphere was developed for the iPlant project [9] and is thus widely known in the biology research community. OpenStack cloud software [10] is the most widely used open source cloud environment, and thus presumably the most reliable in terms of long-term availability. An open source cloud environment seemed also a prudent choice given the budget limits specified in NSF solicitation 14-536.

Since the initial publications of our plans for Jetstream [4], we have sharpened our mission and vision statements as follows. Our vision for Jetstream is straightforward: Jetstream will be a managed science and engineering cloud – a cloud managed and operated in order to support open science and engineering research in the US. Jetstream as a system, and the Jetstream team as a management and support group, will complement existing NSF-funded CI resources supported by XSEDE (the eXtreme Science and Engineering Discovery Environment). In particular, Jetstream aims to provide resources for interactive use any time a handful of processor cores are needed, and for large-scale computational use during “non-peak” hours. Our objective is that Jetstream be known for the distinctive research results and training outcomes it has enabled.

Our basic plan to meet these needs is a system that interactively delivers VMs to end users, providing a modest number of processor cores and modest computational power, and supporting several modes of access and enhanced reproducibility of analyses. Our belief was that we could implement a system that over the course of its life would provide resources for several thousand users.

Here, we present new information on Jetstream’s capabilities, demonstrated as part of fulfilling the criteria for formal acceptance of the system by the NSF. We also discuss early achievements made with Jetstream in its first months of availability to the national research community.

2. SYSTEM IMPLEMENTATION

Jetstream consists of three components: Jetstream-IU, Jetstream-TACC, and Jetstream-AZ. We designed Jetstream to deliver availability and reliability in a way that would mimic commercial cloud experiences and be manageable on a budget of less than \$8M. Jetstream-AZ is housed at the University of Arizona and serves as a test and development environment. The two production components of Jetstream are located in two highly secure and robust data centers in two different geographical regions of the US. They are Jetstream-IU (operated by the Indiana University Pervasive Technology Institute and located on the IU campus in Bloomington, IN) and Jetstream-TACC (operated by the Texas Advanced Computing Center on the campuses of the University of Texas at Austin). These systems are identical, and each consists of:







- Compute nodes: 320 Dell M630 blades with a total of 640 CPUs, 15,360 processor cores, 258 TFLOPS peak processing capability, and 40 TB RAM. Each blade contains two Intel “Haswell” E5-2680v3 (12-core) 2.5 GHz CPUs for single-node peak performance of 806 GFLOPS with 128 GB RAM.
- Storage nodes: 20 Dell R730 servers, with a total of 40 CPUs, 960 processing cores, 1.2 TB RAM, 16 TB local storage, 960 TB of storage, and a peak processing capability of 16.1 TFLOPS.
- Management nodes: 7 Dell R630 servers, with a total of 14 CPUs, 168 processing cores, 448 GB RAM, 5.6 TB local storage, and a peak processing capability of 5.6 TF

The system is based on Dell PowerEdge servers with a 10/40 Gb/s Fat-Tree Ethernet network, oversubscribed 2:1 ratio [16].

Jetstream-AZ was ordered in spring 2015 and passed its acceptance tests in June 2015 [11]. The Jetstream production components were ordered from Dell in summer 2015 and delivered in October. The production systems passed basic system functionality tests at the beginning of November 2015. Jetstream-AZ was available for a few very, very friendly users as of that month. Cloud-based workflows running on Jetstream were demonstrated at SC15. Jetstream-IU also passed its hardware performance tests in November 2015, and Jetstream-TACC passed its tests in February 2016. Table 1 provides a summary of the hardware performance criteria and achieved results. One result merits some explanation. The file system tests provide different results for Jetstream-TACC and Jetstream-IU. IU and TACC are using slightly different versions of Ceph, and Jetstream-IU disks are configured in RAID 0 pairs leading to better large block performance. Jetstream-TACC is faster with 1 MB block size, while Jetstream-IU is faster with 8 MB block size. The complete acceptance review report submitted to the NSF is available online [12].

The NSF specified several criteria that relate to basic integration of a CI resource with XSEDE, all of which were fulfilled while the hardware performance tests were being completed: system available for allocation via XSEDE; resource accepted into the XSEDE Service Provider Forum; accounts created via receipt of packets from XSEDE AMIE system; system connected to XSEDE network. Jetstream became available to the national research community in “Early Operations” mode on February 10, 2016 (this means that the system is available to allocated users in an “as is state,” and use is not charged against allocations). On May 27, 2016 the NSF granted Indiana University authority to pay our vendor partner, Dell Inc., for the Jetstream system hardware.

Table 1. Jetstream hardware performance functionality test results. In the outcome column, G indicates that a metric is being met successfully.

Test	Success criteria	Achieved values	Outcome
Single-Node Performance Tests			
High-Performance Linpack (HPL): Single node Linpack performance (for a problem size using at least half of the memory in a node)	VM performance 80% or more of that achieved in Linux OS	<ul style="list-style-type: none"> • Achieved in Linux OS: Jetstream-IU 697 GFLOPS; Jetstream-TACC 701 GFLOPS • Achieved inside VM: 678 GFLOPS • Performance in VM: 97% of that in OS 	
STREAM: Single node OpenMP threaded performance (aggregate across the node)	65 GB/s	Jetstream-IU: 100 GB/s Jetstream-TACC: 113 GB/s	
10 Gigabit Ethernet Bandwidth link performance: At least 1 GB/s for large-message point-to-point	1 GB/s	1.1 GB/s on the Indiana cluster and 1.2 GB/s on the TACC cluster	
File System and Storage Benchmarks			
200 MB/s data transfer rate for data reads, 100 MB/s writes from within a virtual machine to the block storage (system totals). Initial tests done with dd. Subsequent tests done with IOR.	200 MB/s read 100 MB/s write	<ul style="list-style-type: none"> • Jetstream-IU: 273 MB/s read, 170 MB/s write • Jetstream-TACC: 247 MB/s read, 251 MB/s write 	
Use nttcp to analyze network bandwidth and packet loss.**	N/A	• Jetstream – both clouds: 1.2 GB/s, 0 packet loss	Not part of acceptance test
Use IOR to analyze network bandwidth and I/O to the storage systems. **	N/A	<ul style="list-style-type: none"> • Jetstream-IU: 263 MB/s read, 155 MB/s write • Jetstream-TACC: 229 MB/s read, 235 MB/s write 	Not part of acceptance test
System Reliability & Capacity			
Uptime for a period of 14 days	95% uptime for 14 days	100% uptime as an integrated resource during all of April	
VMs simultaneously operating	320 per location minimum	<ul style="list-style-type: none"> • Jetstream-IU: 998 • Jetstream-TACC: 832 	

** This test was not part of the initial acceptance test set; it was suggested by the review panel appointed by the NSF to review Jetstream.

3. JETSTREAM: SIMULTANEOUSLY FIRST-OF-A-KIND PILOT SYSTEM AND PRODUCTION SYSTEM

The NSF has funded three previous experimental grids and cloud systems – FutureGrid, CloudLab, and Chameleon [13-15]. As previously stated, Jetstream is a first-of-a-kind acquisition and implementation for the NSF within the NSF-funded national CI. Because of that and the a modest (< \$8M) budget, the system is in a sense a pilot experiment of a production cloud system. There was a calculated leap of faith for all involved, including Dell, in proposing Jetstream as an integrated, functioning cloud environment. Thus the NSF could not verify that Jetstream had been implemented as described in the proposal until it passed functionality tests, ultimately demonstrating that it was a geographically distributed, integrated cloud environment. For Dell this meant delivering a system that functions according to standard system delivery; “it works” was not sufficient to ensure payment by the NSF. Dell had to depend on the initial implementation team – IUPTI, TACC, the University of Arizona, the University of Chicago Computation Institute, and Johns Hopkins University – to correctly and successfully implement system software over which Dell had neither influence nor control.

Every first-of-a-kind computer system is a bit of a technical and social experiment: “Can we make it run, and will people who have said they wanted it actually use it if we do?” In this case, the challenge came down to the dual role of a system intended to deliver production cloud-based services and yet be first-of-a-kind. Much of the challenge surrounded defining tests to demonstrate that Jetstream functioned as an integrated cloud environment.

3.1 Cloud functionality tests

Working with the NSF, and developing a program execution plan (PEP) that was peer reviewed by outside experts, we arrived at the following tests of integrated cloud functionality of the system. The function we committed to deliver was academic self-serve cloud services - provide "self-serve" academic cloud services, enabling researchers or students to select a VM image from a published library, or alternatively to create or customize their own virtual environment for discipline- or task-specific personalized research computing. Authentication to this “self-serve” environment will be via Globus. *As agreed by NSF and the Jetstream team, proof of functionality consisted of success in accomplishing the following tasks on Jetstream:*

- An authorized and knowledgeable user will be able to authenticate to the Jetstream user interface (which uses Globus as the mechanism for verification of credentials).
- After so doing, an authorized and knowledgeable user will be able to launch a virtual machine from a menu of pre-packaged VMs on the production hardware located in Indiana or Texas.
- After so doing, an authorized and knowledgeable user will be able to quiesce a VM image running on production hardware in Indiana or Texas, move it from one production system to another, and reactivate said VM.
- An authorized and knowledgeable user will be able to create and access persistent cloud storage on the Indiana or Texas production hardware.
- An authorized and knowledgeable user will be able to modify a preexisting VM image and manually store that VM image to one of the production locations within Jetstream.



Jetstream now reliably does all of these things. During the NSF review of Jetstream, we exhibited the ability to pass these tests in

live demonstrations. The final version of the acceptance report submitted to the NSF [12] includes screenshots of the critical steps.

3.2 Gateway functionality tests

The Jetstream PEP also included functionality tests for supporting science gateways, developed by agreement between the NSF and Jetstream. The function we committed to deliver was science gateway support - Jetstream will support persistent science gateways, including the capability of hosting persistent science gateways within a VM when the nature of the gateway is consistent with VM operation. Galaxy will be one of the initial science gateways supported. *Proof of functionality: Described in Table 2 below.*

Table 2. Jetstream Gateway functionality and performance tests.

Test	Success criteria	Key test metric result achieved	Outcome
Deliver Galaxy gateway	Gateway executes a known workflow in \leq 125% of the time required	Normalized for clock speed, a known workflow completes in 80% of the time to run on Stampede	
Deliver one other gateway functioning in XSEDE environments	Gateway will operate correctly and be available within 2% of overall system availability in 14-day test period	SEAGrid operated continuously for 14 days	

3.3 Data movement, storage & dissemination

The Jetstream PEP also included functionality tests for supporting data movement, storage, and dissemination. The function we committed to deliver was data movement, storage, and dissemination. Jetstream will support data transfer with Globus Connect; users will be able to store VMs in the Indiana University persistent digital repository, IUScholarWorks (scholarworks.iu.edu), and obtain a Digital Object Identifier (DOI) associated with the VM stored. *Proof of functionality:*

- An authorized and knowledgeable user will be able to select a file to which they have rights on a system outside Jetstream, and move and save that file to storage on Jetstream (with the condition that the file size is within the storage quota set for their use on Jetstream).
- An authorized and knowledgeable user will be able to select a file to which they have rights on Jetstream, and move and save that file to storage on a system for which that user has rights and access from open public networks (with the condition that the file size is within the storage quota set for their use on Jetstream).
- An authorized and knowledgeable user will be able to successfully save a VM previously stored to disk storage on Jetstream in a format supported by DSpace, upload that file to scholarworks.iu.edu, and submit that document for publication via IUScholarWorks via online forms. Subsequent to that, the VM will appear in IUScholarWorks and the user will receive a DOI identifier for that object.

These tests were passed successfully in live demonstrations for the NSF review panel and documented via screenshots [12].

4. DEMONSTRATIONS OF UTILITY

4.1 Use and allocations

By the end of April, 327 different people had used Jetstream. Of these, 159 were “end-user researchers or students” and 168 staff. Jetstream SUs (XSEDE Service Units) are based on virtual CPU

hours with varying VM sizes. One SU is one wall clock hour for one vCPU and associated RAM and storage. So far, a total of 53 allocations have been awarded. Table 3 shows a breakdown of allocations and SUs by area of science as of the end of May 2016.

The early operations phase demonstrated success with critical NSF criteria for XD program CI resources – that the system be 90% allocatable and 90% allocated. That something is 90% allocatable is trivial to accomplish. All that is required is that we inform the XSEDE Resource Allocation Service that the system is available for allocation. 90% *allocated* is another matter. In a 30-day month, Jetstream produces 1.67 million SUs at 96% uptime (the required level of uptime specified by the NSF solicitation 14-536). As of the writing of this report, existing Jetstream allocations will consume 90% of the monthly SUs produced for allocation by XSEDE for more than six months after the system goes into full production. Once usage starts being charged against allocations, new requests are expected on an ongoing basis sufficient to have the system 90% allocated for the foreseeable future.

Table 3. Distribution of allocations by discipline for Jetstream and for all other systems supported and allocated via XSEDE.

Discipline or area of interest	# of Jetstream allocations	SUs allocated on Jetstream	% of SUs allocated on Jetstream	% of all SUs allocated on other XSEDE-supported systems
Astronomy	2	300,000	3.5%	10.24%
Biological sciences	11	2,500,000	29.15%	4.05%
Chemistry	1	17,250	.20%	9.99%
Computational Science	5	700,000	8.16%	2.00%
Computer Science	6	800,000	9.33%	1.34%
Earth and Geo Sciences	6	405,120	4.72%	3.61%
Engineering	1	50,000	.58%	3.12%
Molecular science	8	775,000	9.04%	25.85%
Ocean Science	2	150,000	1.74%	1.56%
Physics	1	2,119,920	24.72%	28.93%
Visualization	1	150,000	1.74%	.81%
XSEDE staff	1	250,000	2.91%	.21%
Other / campus champions	8	350,000	4.19%	1.68%
Disciplines not represented on Jetstream	-	-	-	6.59%
Total		8,577,290	100%	100%

Table 3 demonstrates early success in one of the goals stated by the NSF in solicitation 14-536: Increase the diversity of users and uses of resources of the XD program. Relative to typical allocations on large clusters supported by XSEDE, Jetstream allocations show much more interest on the part of biologists (working in areas other than molecular biosciences) in particular than XSEDE as a whole. Principal Investigators with allocations represent 23 states and the District of Columbia. There are allocations to PIs in 10 EPSCoR states. There are also allocations to PIs at two Minority Serving Institutions.

4.2 Technology adoption, community support

In the 1990s, Indiana University pioneered a “leveraged support model” in which a central Information Technology (IT) organization leverages a suite of self-serve online help resources and community partners to deliver more support to a large user Today, new tools and social media allow more sophisticated interactions between support communities, end users, communities of practice (CoPs), and Virtual Organizations (VOs), but the principle of delivering support by leveraging partner effort remains valuable. Our implementation and support strategy for Jetstream is based on collaboration with formal and informal groups of professional CI experts, expert users, instructors, students, and new users as well as the staff of XSEDE. In this regard we have explicitly planned to focus on the factors of social influence and facilitating conditions (including knowledge of a technology and the belief that end users will find it accessible) as described in [8]. That the day-in, day-out support of Jetstream will fall largely on our partner CoPs and VOs means that social influences and facilitating conditions will speed adoption of Jetstream by researchers, educators, and students who are not currently users of XSEDE-funded CI resources. These VO partners are listed in Table 4. Most will not receive direct funding from IU as part of the Jetstream project. This is a strong indication of their beliefs the utility of Jetstream.

Table 4. Partners in deployment and support of Jetstream. * Indicates that a partner is funded to deliver support services through a subcontract of the NSF award to operate Jetstream

Discipline or mode of use	Lead partners
Biology	iPlant, University of Arizona*, Galaxy, Johns Hopkins University*, National Genome Analysis Support Center
Earth Science/Polar Science	National Snow and Ice Data Center Network (RCN)
Field station research	University of Arizona
Network Science	IU Network Institute
Observational astronomy	WIYN Consortium
Social sciences	Odum Institute, University of North Carolina
Campus bridging	XSEDE, Cornell, IU
Under-resourced schools	University of Hawaii, University of Arkansas Pine Bluff, University of Texas San Antonio*
Use of proprietary software	Mathworks
Reproducible data analyses	University of Chicago Computation Institute
Enhanced science gateway deployment	University of California San Diego (Supercomputing Center), XSEDE
Visualization and analysis	IU, University of Texas

Indeed, already during the friendly user phase and early operations phase VOs and XSEDE Campus Champions have assisted with disseminating help information about Jetstream. With regards to assistance from XSEDE staff, we are pragmatic in our expectations given the funding limitations expected in the future. We expect significant expert help from the allocations staff and processes of XSEDE. In the early days of Jetstream operations, this may be XSEDE staff’s most critical contribution to Jetstream operations. Because we expect XSEDE’s training capabilities to be limited in scope, we have built the creation of Cornell Virtual Workshops into the Management and Operations budget for Jetstream. Over time,

we expect vital assistance from XSEDE Extended Collaborative Support Services with in-depth performance tuning and science gateway development. However, we expect little to no practical help from XSEDE in resolving day in, day out user problems other than authentication and accounting issues; there simply is no XSEDE budget for that moving forward.

4.3 New science results

Perhaps the strongest sign of Jetstream’s ability are analyses performed by people not affiliated with the project, who have already produced useful incremental results that will accelerate the submission of scientific technical reports to peer-reviewed journals. Early results involve several priority research areas from our initial proposal – genomics and field biology, psychology, computer and computational science. So far use of Jetstream has helped provide:

- New insights into the evolution of endemic fish species in the US Southwest, evolution of poisonous snakes, and the evolution of plants in South Africa
- Support for development of new methods for brain mapping
- Storage facilities to enable replication of scientific analyses








The system has also been used for undergraduate education in classes at IU and as a tool for support of doctoral student research at SUNY Binghamton. Jetstream’s utility in aiding research breakthroughs represents two underlying facts: There were researchers in the US who did not have access to computational resources suitable for completing their research data analyses – some of which were computationally intense; and the user-friendliness of Jetstream was such that these researchers could start from scratch and produce publication-ready data analyses in a matter of 10 weeks or less. Going back to the definition of cyberinfrastructure as enabling research results “not otherwise possible” we see that Jetstream is indeed enabling research discoveries that were not otherwise possible – in ways not envisaged by IU when it codified this as part of the definition of cyberinfrastructure!

5. OPERATIONAL METRICS

In April we collected the metrics specified by the NSF for all production CI systems in the XD program and additional metrics agreed upon between the NSF and the Jetstream team. The results are shown below in Table 5. The Jetstream team and the NSF had to determine how to interpret and apply some of the metrics to a cloud system. Availability and job completion success are metrics that go back many years in NSF cyberinfrastructure funding. IU and NSF added a statistic for “capacity availability” because uptime is not a meaningful measure of availability for a cloud environment. In general, a cluster or a supercomputer tends to be in operation or not; it’s usually a binary condition. A distributed cloud system could be “up,” in the sense of “at least partially available to accept logins,” while still operating at well less than full capacity. Therefore, total capacity availability of at least 95% was added as a metric. The “job completion success” metric took considerably more work to interpret and express for a cloud computing environment. That has for years been defined by the NSF as the percentage of jobs that complete without having to be resubmitted as a result of a failure in the system hardware or software. There is no obvious analog within a cloud system for a “job.” However, in a more general sense, this metric measures the ability of the hardware and software to properly execute a series of commands. Within Jetstream, there are three types of VM images: those that are private to an individual user; those that are created by an individual and made available for use to the Jetstream user community; and what we refer to as “featured” VMs – those

certified to work properly and produce correct results (in terms of executing software within that VM). For Jetstream, we decided meeting this metric meant that 95% of the launches of featured VMs were executed correctly – to the point of a VM reporting its status to the Atmosphere interface and orchestration system as “in operation” after being launched. Evaluation results are presented in Table 5.

Table 5. Operational metrics for April 2015. G indicates a metric is being met. Y indicates 60% of target, a metric for focused attention and effort. * Indicates a metric required by the NSF for all major CI acquisitions ** Indicates data collected between 15 and 30 April * Annual goals are marked Green when the rate of growth per month is at or above the level required to achieve the annual goal**

Metric	Goal	Achieved	Outcome
System availability (uptime)*	95%	100%	
Capacity availability (% of system capacity available)*	95%	100%	
Job completion success*	96%	97.7%	
Number of distinct users	1,000	327***	
Use - average number of VMs active	320	290; peak of 1217**	
CPU % utilization	6%	Mean of 4.2%; peak 20.3% **	
VM images published with a DOI via IUScholarWorks	10/year	6 in one quarter***	

Perhaps the most challenging metric to transpose into a cloud setting – and the one metric that remains open as a topic of discussion between the Jetstream team and the NSF – is the concept of the busyness of the system. This has to do with both the extent to which the system is fully occupied doing as many computing jobs as possible, and the extent to which CPUs are well exploited.

For a cloud environment, there are inherent challenges in measuring the extent to which a system is fully occupied. We commonly hear discussion of cluster and supercomputer operations at more than 80% or 90% capacity. These statistics are really measures of the extent to which queues are scheduled with jobs. They do not measure CPU utilization – even with very sophisticated codes running with highly tuned batch management, systems can operate with CPU utilization percentages (percent of CPU’s total mathematical capacity) in the teens. In a sense, a cluster or supercomputer operates with a subscription rate of one “job” per CPU. Cloud systems deliver part of their value because of the ability to oversubscribe, assigning multiple virtual CPUs (vCPUs) in different VMs to one physical CPU. So in a cloud environment, “% busy” is a function of the computational intensity of the job being run, the efficiency of the code in exploiting the CPU, the extent to which oversubscription is allowed, and the amount of time a researcher sits pondering results of one interactive command before typing another. For a particular level of activity on a cloud system, the % busy statistic may be arbitrarily manipulated by changing the permitted level of oversubscription.

The challenges of any sort of “% busy” statistic in a cloud environment brought about a discussion of utilization as measured by the hardware counters on the CPUs. We were challenged by the general lack of published figures for traditional clusters (CPU utilization as a percent of the total core / floating point unit

utilization possible – not whether a CPU was reserved). The shortcomings of the Linpack benchmark, on which the Top500 list is based, are well known. A particular problem is that the Linpack benchmark rewards investment in CPUs, which are generally not as well used by other applications as by Linpack, while Linpack has less dependency on memory and network capabilities. Dr. Jack Dongarra, creator of the benchmark, wrote in a blog that “We have reached a point where designing a system for good Linpack performance can actually lead to design choices that are wrong for the real application mix” [18]. Open publication of monthly CPU utilization for major supercomputers and clusters would provide cloud providers like Jetstream with better data to help interpret cloud CPU utilization. This might also lead to better overall system design and put greater focus on software efficiency at the per-node level. It would also provide data to help the CI community address an important challenge common to HPC and cloud computing applications: node level performance.

Our approach so far has been to optimize for the productivity of individual Jetstream users. In order to provide a consistent user experience, we have in the early operations phase of Jetstream kept the maximum level of oversubscription at 2 vCPUs, with 1 physical CPU and an average of 320 VMs active simultaneously. We have operated the system with a combined human and script-generated load of 1,217 VMs. We have demonstrated an ability to drive CPU utilization to over 20% – a level that would be respectable for any cluster or supercomputer workload. Still, in early operations we have come relatively close to achieving our operational benchmark goals for operations (as indicated by the “Yellow” outcome indicator showing we are within 60% of achieving goals).

During the early operations phase, much of what we learned was that the system and its intended users operate largely as we had hoped and as we had written in our initial proposal. Scientists with data to analyze did so interactively, entering commands, looking at results, and entering more commands. We were completely surprised by one aspect of system use: the extent to which getting consistent, effective, high levels of utilization from the system can be managed through use of orchestration systems. Indeed, a significant portion of the CPU utilization to date has been generated by orchestration systems such as Mesos allowing interactive use and control of several VMs at once [19] and by science gateways such as SEAGrid [20]. We were surprised by the ease with which users of Jetstream, without any special system-level privileges, could manage the orchestration of large numbers of VMs within Jetstream from a single VM within Jetstream. We expect that orchestration systems will be a major factor in making good use of cyberinfrastructure systems while enabling large amounts of high-quality research.

6. WHY IT MAKES SENSE FOR THE NSF TO FUND A CLOUD SYSTEM

Several reports and many individuals have suggested that the NSF fund a cloud resource. Still, it’s worth asking at this point once again: Why not just depend upon the private sector? We believe the following factors, many of which are based on our early experience to date, suggest that it makes sense for the NSF to fund a resource such as Jetstream. It is, as we put it, a managed science and engineering research cloud; a cloud managed with science and engineering research as its first priority. Commercial clouds are not driven primarily by science needs. Why should the NSF invest in this? There are several reasons:

- Jetstream is without cost to the end user. The allocation process encourages scientists – particularly scientists in

domains that have not traditionally made deep use of advanced CI – to use it.

- The Jetstream team works from the user interaction layer of Atmosphere on down through OpenStack to network tuning, providing a degree of vertical integration, testing, and tuning that would be impossible to deliver on a commercial cloud.
- The Jetstream team tests, tunes, and certifies a set of VMs that we stand behind and guarantee to work. We work with Communities of Practice, VOs, and disciplinary groups to prioritize what we provide as a “featured VM.”
- We provide a clearinghouse of VMs contributed by scientists who want to make them available to the research community.
- We offer free storage for VMs and related products used on Jetstream as a digital object in a persistent digital archive (IUScholarWorks) with an associated DOI.

Lastly, and perhaps most importantly, there is inherent value in diversity, in terms of community stability. It is a well-known result from community ecology that diversity creates ecosystem community. This works in ecosystems of plants and animals, and in human-created ecosystems as well. Jetstream makes the community of cloud providers more diverse. Within the OpenStack community, it is one of the larger government supported clouds. Diversity is also better than monoculture over long time periods because it provides more sources from which to pick innovations.

There is also inherent value in Jetstream as a pilot for the NSF. System usage is metered by the capacity of the system they purchase as a capital investment. This makes it possible for the NSF to invest in a way that they can plan for what they otherwise might find hard to manage and budget in the form of payments to a commercial provider. As a pilot, it allows the NSF to make a significant but bounded commitment, learn from this pilot project, and on the basis of Jetstream make plans for the longer term future.

Over the next four or five years of production operations of Jetstream, we expect three primary sorts of benefits:

- Many important scientific discoveries made by communities of researchers that are new to the XD program and have not used XSEDE-supported resources before.
- Significant broader impacts stemming from use of Jetstream, ranging from workforce development to societally important outcomes of the work enabled.
- Significant knowledge for the NSF and for the open science community about what are, in practice, the costs and benefits of running a cloud system rather than buying services on commercial cloud systems.

7. CONCLUSIONS

Our key conclusion is that the system functions as proposed, and has thus far proved valuable to its initial users. A panel of experts appointed by the NSF reviewed the Jetstream implementation in April 2016, recommending that it transition into full production status. Jetstream has succeeded in providing high-end resources to a diverse research community. It provides valuable software tools for the communities identified as intended users of Jetstream. It supports multidisciplinary computational science and engineering, and early experiences suggest it will further the progress of the US open science and engineering community.

The NSF and the Jetstream team have worked together to set new precedents for measuring the effectiveness and operation of government funded cloud systems. The extensive scope of the acceptance review – spanning basic benchmarks to analysis of user experiences to operational metrics during early operations – has

helped ensure that the system implemented and presented to the NSF for acceptance is indeed the system we proposed in the introduction of our initial proposal (corrected for budget changes between proposal and award). We endorse this approach as one that will hold the cyberinfrastructure community highly accountable to itself, funding agencies, and our users – and as an approach that in the long run will increase the conformity of systems as delivered to systems as described in grant proposals to funding agencies.

One of the critical challenges going forward will be for the Jetstream team to demonstrate value in the form of return on investment for the NSF. For that reason we have created a DOI for a basic description of Jetstream and ask that researchers, educators, and students acknowledge Jetstream use by citing it [21] in all publications and products created via some use of Jetstream or Jetstream-related products (including use of VMs accessed from IUScholarWorks). This will allow the Jetstream team to easily discover products created with some contribution from Jetstream, so that we can document the value this system.

8. ACKNOWLEDGMENTS

This material is based upon work supported in part by the National Science Foundation under Award 1445604. We thank all the staff of the Jetstream partner organizations for making Jetstream such a success so far (especially George Turner, Bret Hammond, and Mike Packard). The Indiana University Pervasive Technology Institute and our partners have also supported Jetstream implementation. Any opinions expressed here are those of the authors and do not necessarily represent the opinions of any funding agencies. We thank reviewers of earlier versions of this report. Thanks to Dr. Winona Snapp-Childs for editorial contributions. We also wish to thank the staff who have worked with us during all phases of this project. We most particularly thank our primary vendor partner Dell and our primary vendor representative Jamie Stevens for their partnership in this adventure.

9. REFERENCES

1. Stewart, C.A., Simms, S., Plale, B., Link, M., Hancock, D., and Fox, G. *What is Cyberinfrastructure?* in *SIGUCCS 2010*. 2010. Norfolk, VA. <http://dx.doi.org/10.1145/1878335.1878347>
2. National Science Foundation. *NSF Solicitation 14-536*. 2014; Available from: <http://www.nsf.gov/pubs/2014/nsf14536/nsf14536.htm>.
3. 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
4. Stewart, C.A., Cockerill, T.M., Foster, I., Hancock, D., Merchant, N., Skidmore, E., Stanzione, D., Taylor, J., Tuecke, S., Turner, G., Vaughn, M., and Gaffney, N.I., Jetstream: a self-provisioned, scalable science and engineering cloud environment. 2015, In *Proceedings of the 2015 XSEDE Conference*. St. Louis, Missouri. ACM: 2792774. p. 1-8. <http://dx.doi.org/10.1145/2792745.2792774>
5. Stodden, V. and Miguez, S., *Best Practices for Computational Science: Software Infrastructure and Environments for Reproducible and Extensible Research*. Journal of Open Research Software, 2014. **2**(1): p. e21. <http://doi.org/10.5334/jors.ay>
6. Hurlin, C. and Perignon, C., *RunMyCode.org: A Research-Reproducibility Tool for Computational Sciences*, in *Implementing Reproducible Computational Research*, V. Stodden, F. Leisch, and R. and Peng, Editors. 2014, Chapman & Hall/CRC: Boca Raton.
7. Lawrence, K.A., Zentner, M., Wilkins-Diehr, N., Wernert, J.A., Pierce, M., Marru, S., and Michael, S., *Science gateways today and tomorrow: positive perspectives of nearly 5000 members of the research community*. Concurrency and Computation: Practice and Experience, 2015. **27**(16): p. 4252-4268. <http://dx.doi.org/10.1002/cpe.3526>
8. Venkatesh, V., Morris, M.G., Davis, F.D., Davis, G.B, *User Acceptance of Information Technology: Toward a Unified View*. MIS Quarterly, 2003. **27**(3): p. 425-478. [http://csdl-techreports.googlecode.com/svn/trunk/techreports/2005/05-06/doc/Venkatesh2003.pdf](http://csdl.techreports.googlecode.com/svn/trunk/techreports/2005/05-06/doc/Venkatesh2003.pdf)
9. iPlant Collaborative. *Home page*. Available from: <http://iplantcollaborative.org/>.
10. OpenStack. *Home page*. Available from: <http://www.openstack.org/>.
11. Hancock, D.Y., Link, M.R., and Stewart, C.A., *Acceptance Test for Jetstream Test Cluster – Jetstream-Arizona (JA) Dell PowerEdge Test and Development Cluster*. PTI Technical Report PTI-TR15-007, 2015. <http://hdl.handle.net/2022/20355>
12. Stewart, C.A., Foster, I., Merchant, N.C., Taylor, J., and Vaughn, M.W., *Jetstream – A Self-Provisioned, Scalable Science and Engineering Cloud Environment - NSF Acceptance Report*. PTI Technical Report, 2016. <http://hdl.handle.net/2022/20889>
13. Fox, G., Fortes, J., Grimshaw, A.S., Keahey, K., and Smith, W., *FutureGrid: An Experimental, High-Performance Grid Test-bed*. 2009, National Science Foundation: Indiana University.
14. *Chameleon: A configurable experimental environment for large-scale cloud research*. Available from: <https://www.chameleoncloud.org/>.
15. *Cloudlab*. Available from: <http://www.cloudlab.us/>.
16. Center for Computational Research - University at Buffalo. *Open XDMoD: Metrics on Demand*. Available from: <http://xdmod.sourceforge.net/>.
17. Voss, B.D., Jung-Gribble, Diane., Stewart, Craig A. *The Leveraged Support Model*. 1996. <http://scholarworks.iu.edu/dspace/handle/2022/14011>
18. Sieslack, N. *Jack Dongarra on TOP500: Past, Present, and Future*. 2015 [cited 2016 June 15]; Available from: <http://www.top500.org/news/jack-dongarra-on-top500-past-present-and-future/>.
19. DelValle, R., Rattihalli, G., Beltre, A., Madhusudhan, G., and Lewis, M., *Exploring the Design Space for Optimizations with Apache Aurora and Mesos*. accepted for publication in the IEEE International Conference on Cloud Computing (CLOUD), Applications Track, 2016.
20. Shen, N., Fan, Y., and Pamidighantam, S., *E-science infrastructures for molecular modeling and parametrization*. Journal of Computational Science, 2014. **5**(4): p. 576-589. <http://www.sciencedirect.com/science/article/pii/S1877750314000064>
21. Stewart, C., Stanzione, D., Cockerill, T., Skidmore, E., Fischer, J., Lowe, J.M., Hammond, B., Turner, G., Hancock, D., and Miller, T., *NSF High Performance Computing System Acquisition system description: Jetstream - a self-provisioned, scalable science and engineering cloud environment*. 2016. <http://doi.org/10.5967/K85Q4T1N>