# Managing Computational Gateway Resources with XDMoD

12th Gateway Computing Environments Conference

Jeanette Sperhac, Robert DeLeon, Thomas Furlani, Steven Gallo, Martins Innus, Matthew D. Jones, Jeffrey T. Palmer, Abani Patra, Benjamin Plessinger, Ryan Rathsam, Nikolay Simakov, Joseph P. White, and Thomas Yearke 24 October 2017

Center for Computational Research, University at Buffalo

1. Introducing XDMoD

What can it do?

Open XDMoD

- 2. XDMoD for Gateways
- 3. Future Work

Gateways Realm

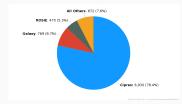
**Custom Queries** 

4. Conclusion

# Introducing XDMoD

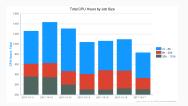
#### **XDMoD User Interface**





XDMoD collects, processes, aggregates, and verifies jobs and allocation data from supercomputers.

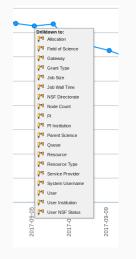
These data come from a wide variety of sources.



XDMoD users include anyone who uses or administers (or funds) high performance computing (HPC) resources, including:

- Gateway managers
- Gateway end users
- HPC center staff

Jobs	Performance	Accounting (XSEDE)		
CPU hours	Total memory	Allocation usage rate		
Job size	CPU usage	Number of allocations		
Node hours	Memory bandwidth	Number of institutions		
Number of jobs	I/O bandwidth	Service Units (XD SUs)		
Wall hours	Block read rate	Normalized Units (NUs)		
Wait hours	Block write rate	XSEDE utilization		



#### Describe the job or Describe the user that ran the job

Job	User		
Job Size (Core Count)	Username		
Wall Time	Institution		
Queue	Field of Science		
Node Count	PI		
(Computing) Resource	Grant Type		
	Allocation		

### **Exploring Metrics: Usage**

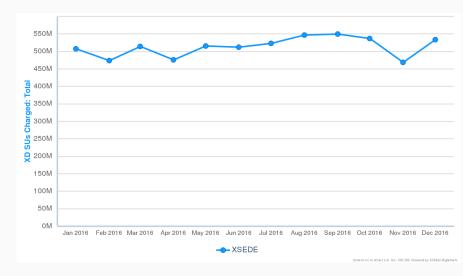
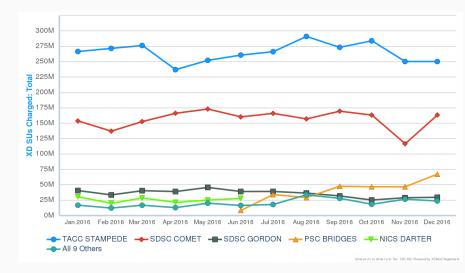


Figure 1: Total XD SUs charged, 2016

#### Exploring Metrics: Usage by Resource



#### Figure 2: Total XD SUs charged for top 5 Resources, 2016

XDMoD lets you customize plots:

- filter data
- overplot multiple datasets
- change scale
- change appearance

## Community (Gateway) Usage by Resource

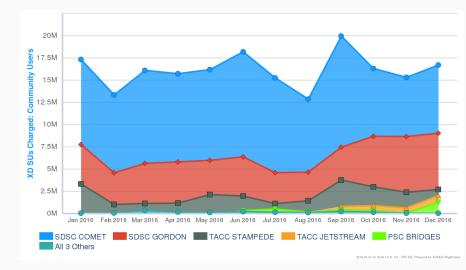


Figure 3: XD SUs charged by Community Users, 2016

- automatically generate reports
- export plots and datasets in multiple formats

## Export results

Month	Comet	Gordon	Stampede	Jetstream	Bridges	Other
1	9569751	4431442	3286014	0	0	36875
2	8766017	3555435	963777	0	0	53902
3	10459094	4512738	776036	0	0	347335
4	9925126	4636087	996356	0	0	170456
5	10209511	3854499	2011788	0	0	113293
6	11820967	4404377	1633338	0	122678	211118
7	10669700	3473191	606501	0	357075	149271
8	8215691	3225669	1256510	0	61122	107382
9	12499128	3696923	2968334	403487	155790	234507
10	7632032	5685713	2107839	573281	163999	145665
11	6661933	6268783	1756497	583061	4094	40196
12	7665562	6331866	686692	770874	1216712	28925

 Table 1: XSEDE SUs charged monthly by Community Users, 2016

XDMoD was developed to help manage NSF-funded XSEDE supercomputers.

The open-source version, which we run at our own center, can be installed on computational resources of any size.



- Run by industrial and academic computing centers, and government labs, worldwide
- Source on github; RPMs available
- Supported by the XDMoD team
- Frequent new releases and features
- Thousands of downloads, hundreds of installs

## XDMoD for Gateways

- Gateways dimension
- Grant Type dimension: Science Gateways
- User dimension: Community (Gateway) users

### Performance data in aggregate

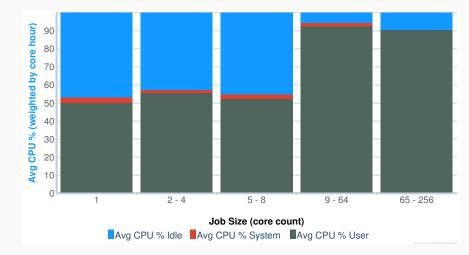


Figure 4: Average CPU percent by job size, Cipres gateway, 2016

#### Gateway Jobs on Comet

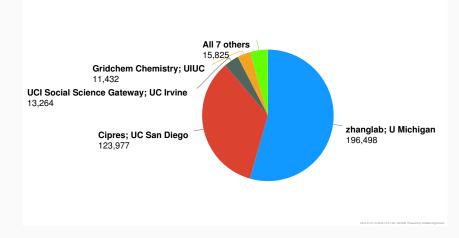


Figure 5: Count of XSEDE Gateway jobs submitted to Comet, 2016

#### **XSEDE Gateway CPU Hours**

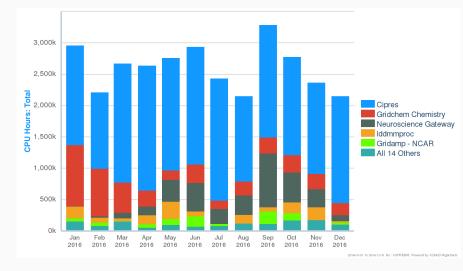
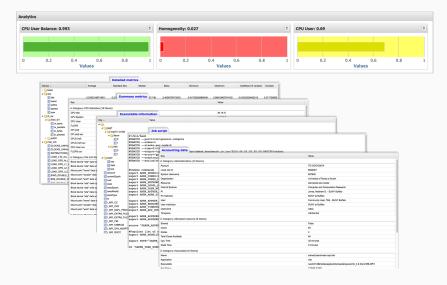


Figure 6: CPU hours used by top five XSEDE Gateways, 2016



#### Job Viewer

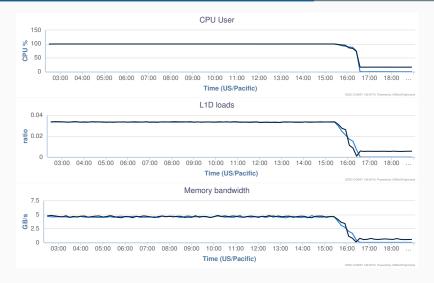


Figure 7: Job Viewer timeseries plots help diagnose problematic jobs

### Job Viewer

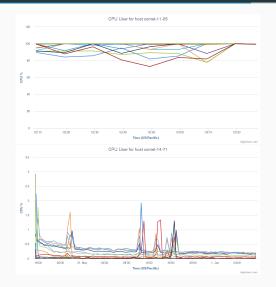


Figure 8: CPU User on efficient and inefficient jobs. Same executable!

**Future Work** 

#### A tale of two Gateways: Vhub and VIDIA



Figure 9: Command-line jobs started, CCR's VIDIA and Vhub gateways, 2017

Gateways will be first-class citizens, much as Resources are today. Features:

- Individual Gateway user IDs
- · Gateway manager access to job performance data
- $\cdot\,$  Granular access to gateway user and job information
- $\cdot$  Developed and prototyped on our own gateways^1  $\!$

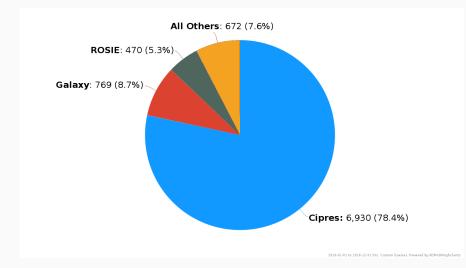
<sup>&</sup>lt;sup>1</sup>vhub.org and vidia.ccr.buffalo.edu

Thirteen gateway-specific metrics, by resource or gateway, include *counts* of:

- Unique gateway users
- Gateway jobs
- Gateway jobs with unknown users
- Gateway XD SUs vs. total XSEDE XD SUs

...available to authorized XSEDE users<sup>2</sup>

#### Gateways Custom Queries: Example



#### Figure 10: Unique XSEDE Gateway Users, 2016

XDMoD 7.1 for XSEDE is slated for release on 1 November 2017. Features include:

- Gateways custom queries
- Jetstream data
- Storage metrics (alpha version)
- Globus Authentication support

Conclusion

XDMoD helps the supercomputer community:

- assess their resource use;
- analyze individual job performance;
- inform code optimization and improvement;
- tune application performance;
- collect usage data for long-range analysis and planning.

Open XDMoD documentation, RPMs, and support:

#### http://open.xdmod.org/

XSEDE version:

https://xdmod.ccr.buffalo.edu/



Come see us at the Gateways 2017 poster session, 5:30 today.

#### NSF ACI 1025159 and ACI 1445806





#### Center for Computational Research SUNY University at Buffalo