

Deploying File System Performance Metrics Through XSEDE Metrics on Demand (XDMD)



Bailey Kleespies
SIParCS Intern

July 31, 2019



Project:

- Deploying XDMoD and SUPReMM to help CISL monitor CPU, memory, and file system metrics
 - Utilization, quality of service, job level performance, etc.

Purpose:

- To make adjustments to the installation of XDMoD on NCAR systems

Goal:

- To be able to deliver an accurate weekly report to management

- XDMoD
 - XSEDE Metrics on Demand
- SUPReMM
 - Job Performance Data



Version 8.1

Tools

- Ganglia on the nodes of Cheyenne
- PBS scheduler account log
- VM service provided by EIO (Enterprise Engineering Organization)



What We Did

1. Request and Prepare Server Resource
2. Set up Open XDMoD Instance
3. Set up SUPReMM Modules and Implement NCAR Variation
4. Manage Data Flow my cronjob and Generate Report from XDMoD Portal
5. Set up Public Access to XDMoD Portal
6. Set up Special Access to XDMoD Portal via port 8080

Public Mode



Hello, [Sign In](#) to view personalized information.

[About](#) [Roadmap](#) [Contact Us](#) [Help](#)


Summary Usage About

Duration: Previous month Start: 2019-06-01 End: 2019-06-30 Refresh


Quick Filters

Activity		Jobs	CPU Time (h)		Wait Time (h)		Wall Time (h)		Processors	
Users:	Pls:	Total:	Total:	Avg (Per Job):	Avg (Per Job):	Total:	Avg (Per Job):	Max:	Avg (Per Job):	
624	320	441,347	41,560,528.5	93.45	0.17	78,672.1	0.18	285,984	108	

Total CPU Hours By Resource (Top 10)



Total CPU Hours by Job Size



User Mode


Metrics and Options

Title:

Legend:

Font Size:

- ### Jobs Summary
- CPU Hours: Per Job
 - CPU Hours: Total
 - Job Size: Max
 - Job Size: Min
 - Job Size: Normalized
 - Job Size: Per Job
 - Job Size: Weighted By CPU Hours
 - Node Hours: Per Job
 - Node Hours: Total
 - Number of Jobs Ended
 - Number of Jobs Running
 - Number of Jobs Started
 - Number of Jobs Submitted
 - Number of PIs: Active
 - Number of Resources: Active
 - Number of Users: Active



Description

- The National Center for Atmospheric Research:** Summarizes jobs reported to the The National Center for Atmospheric Research central database (excludes non-The National Center for Atmospheric Research usage of the resource).
- CPU Hours: Per Job:** The average CPU hours (number of CPU cores x wall time hours) per The National Center for Atmospheric Research job.
For each job, the CPU usage is aggregated. For example, if a job used 1000 CPUs for one minute, it would be aggregated as 1000 CPU minutes or 16.67 CPU hours.
- CPU Hours: Total:** The total CPU hours (number of CPU cores x wall time hours) used by The National Center for Atmospheric Research jobs.
For each job, the CPU usage is aggregated. For example, if a job used 1000 CPUs for one minute, it would be aggregated as 1000 CPU minutes or 16.67 CPU hours.

User Mode

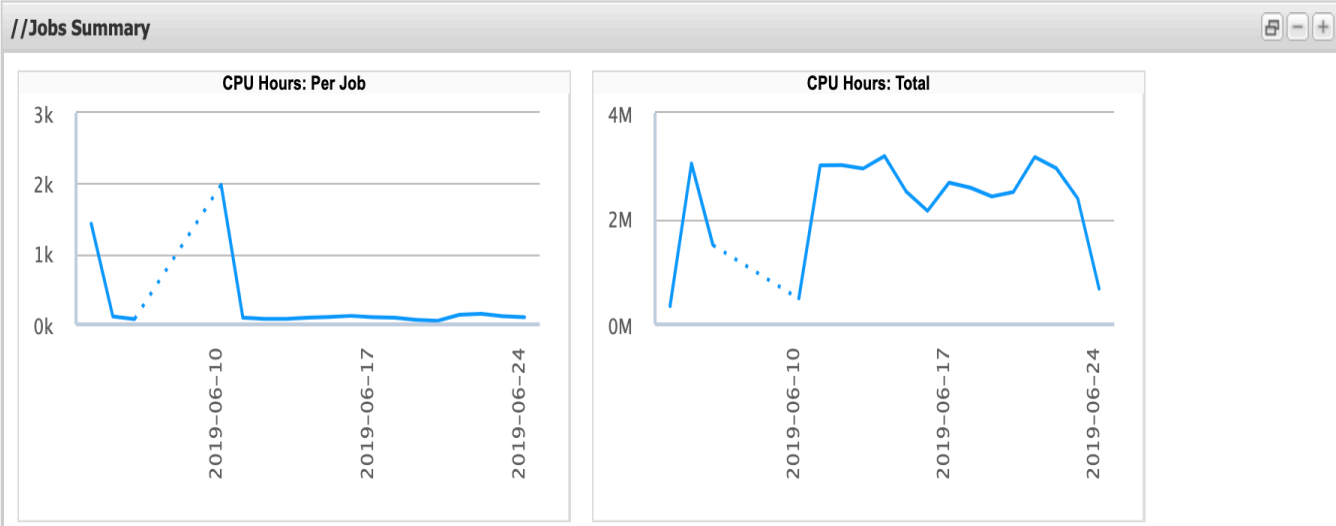
Metrics and Options

Title:

Legend:

Font Size:

- Jobs Summary
- CPU Hours: Per Job
- CPU Hours: Total
- Job Size: Max
- Job Size: Min
- Job Size: Normalized
- Job Size: Per Job
- Job Size: Weighted By CPU Hours
- Node Hours: Per Job
- Node Hours: Total
- Number of Jobs Ended
- Number of Jobs Running
- Number of Jobs Started
- Number of Jobs Submitted
- Number of PIs: Active
- Number of Resources: Active
- Number of Users: Active



Description

- The National Center for Atmospheric Research:** Summarizes jobs reported to the The National Center for Atmospheric Research central database (excludes non-The National Center for Atmospheric Research usage of the resource).
- CPU Hours: Per Job:** The average CPU hours (number of CPU cores x wall time hours) per The National Center for Atmospheric Research job.
For each job, the CPU usage is aggregated. For example, if a job used 1000 CPUs for one minute, it would be aggregated as 1000 CPU minutes or 16.67 CPU hours.
- CPU Hours: Total:** The total CPU hours (number of CPU cores x wall time hours) used by The National Center for Atmospheric Research jobs.
For each job, the CPU usage is aggregated. For example, if a job used 1000 CPUs for one minute, it would be aggregated as 1000 CPU minutes or 16.67 CPU hours.

User Mode

XDMoD Hello, Broday Walker ([logout](#)) Dashboard My Profile About Roadmap Contact Us Help

Summary **Usage** Metric Explorer Report Generator Job Viewer About

Duration: Previous month Start: 2019-06-01 End: 2019-06-30 Refresh Filter Display Top 10 Export Print Available For Report

Quick Filters

Metrics and Options

Title:

Legend:

Font Size:

Jobs Summary

- CPU Hours: Per Job
- CPU Hours: Total**
- Job Size: Max
- Job Size: Min
- Job Size: Normalized
- Job Size: Per Job
- Job Size: Weighted By CPU Hours
- Node Hours: Per Job
- Node Hours: Total
- Number of Jobs Ended
- Number of Jobs Running
- Number of Jobs Started
- Number of Jobs Submitted
- Number of PIs: Active
- Number of Resources: Active
- Number of Users: Active
- The National Center for Atmospheric Research Utilization

//Jobs Summary/CPU Hours: Total

CPU Hours: Total

Date	CPU Hours: Total (M)
2019-06-04	~0.5
2019-06-10	~0.5
2019-06-12	~3.0
2019-06-14	~3.2
2019-06-16	~2.5
2019-06-18	~2.8
2019-06-20	~2.5
2019-06-22	~3.2
2019-06-24	~0.8

Drilldown to:


- CISL
- CSG
- Job Size
- Job Wait Time
- Job Wall Time
- NCAR
- Node Count
- PI
- Queue
- Resource
- Resource Type
- System Username
- User

Description

- The National Center for Atmospheric Research:** Summarizes jobs reported to the The National Center for Atmospheric Research central database.
- CPU Hours: Total:** (number of CPU cores x wall time hours) used by The National Center for Atmospheric Research jobs. For each job, the CPU usage is aggregated. For example, if a job used 1000 CPUs for one minute, it would be aggregated as 1000 CPU minutes or 16.67 CPU hours.

The National Center for Atmospheric Research
2019-06-01 to 2019-06-30 Src: HPCDB. Powered by XDMoD/Highcharts

User Mode

 Hello, **Broday Walker** ([logout](#)) [Dashboard](#) [My Profile](#) [About](#) [Roadmap](#) [Contact Us](#) [Help](#)

Summary **Usage** Metric Explorer Report Generator Job Viewer About

Duration: Previous month Start: 2019-06-01 End: 2019-06-30 Refresh Filter Display Top 10 Export Print Available For Report

Quick Filters

Metrics and Options

Title:

Legend:

Font Size:

Jobs Summary

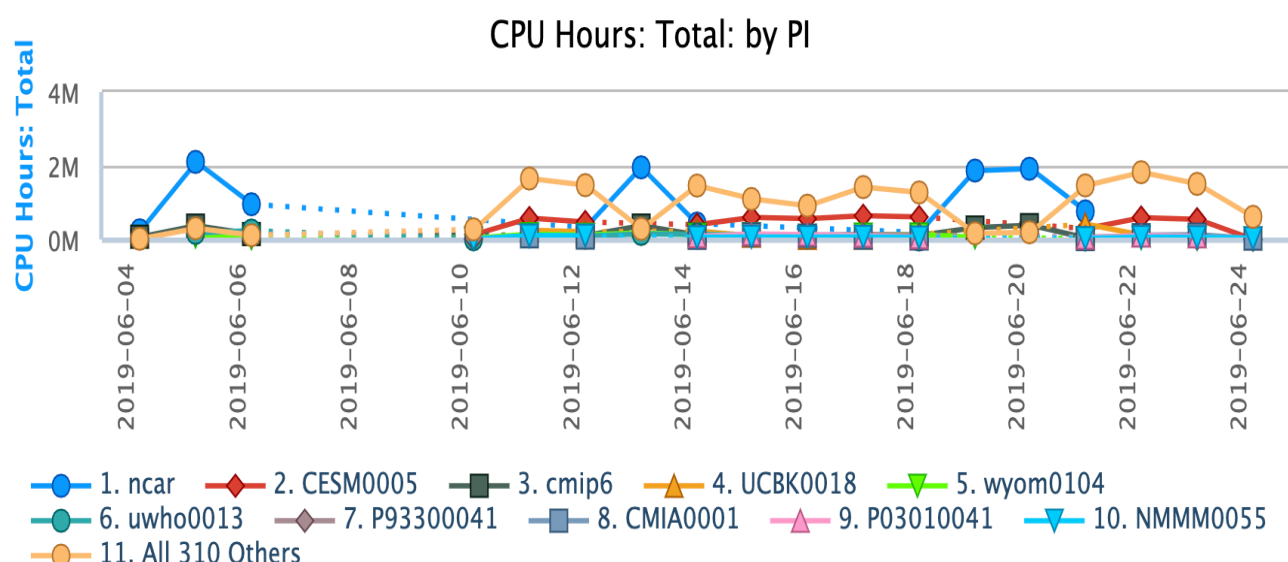
- CPU Hours: Per Job
- CPU Hours: Total
- by PI**

The National Center for Atmospheric Research: The Nation

- Job Size: Max
- Job Size: Min
- Job Size: Normalized
- Job Size: Per Job
- Job Size: Weighted By CPU Hours
- Node Hours: Per Job
- Node Hours: Total
- Number of Jobs Ended
- Number of Jobs Running
- Number of Jobs Started
- Number of Jobs Submitted
- Number of PIs: Active
- Number of Resources: Active

//Jobs Summary/CPU Hours: Total/by PI

CPU Hours: Total: by PI



Date	1. ncar	2. CESM0005	3. cmip6	4. UCBK0018	5. wyom0104	6. uwho0013	7. P93300041	8. CMIA0001	9. P03010041	10. NMMM0055	11. All 310 Others
2019-06-04	0.5M	0.2M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M
2019-06-06	2.2M	0.2M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M
2019-06-10	0.5M	0.2M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M
2019-06-12	0.5M	0.2M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	1.8M
2019-06-14	2.2M	0.2M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	1.8M
2019-06-16	0.5M	0.2M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	1.2M
2019-06-18	0.5M	0.2M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	1.5M
2019-06-20	2.2M	0.2M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.5M
2019-06-22	0.5M	0.2M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	1.8M
2019-06-24	0.5M	0.2M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.1M	0.8M

2019-06-01 to 2019-06-30 Src: HPCDB. Powered by XDMoD/Highcharts

Description

- PI:** The principal investigator of a project has a valid allocation, which can be used by him/her or the members of the project to run jobs on.
- CPU Hours: Total:** The total CPU hours (number of CPU cores x wall time hours) used by The National Center for Atmospheric Research jobs. For each job, the CPU usage is aggregated. For example, if a job used 1000 CPUs for one minute, it would be aggregated as 1000 CPU minutes or 16.67 CPU hours.

User Mode

Duration: Previous month Start: 2019-06-01 End: 2019-06-30 Refresh Export Print Available For Report

Metric Catalog Load Chart New Chart Save As Save Undo Redo Delete Add Data Data Quick Filters Add Filter Filters

- Jobs
- SUPREMM**
- Avg CPU %: Idle: weighted by core-hour
- Avg CPU %: System: weighted by core-hour
- Avg CPU %: User: weighted by core-hour
- Avg GPU0 usage: weighted by node-hour (GPU %)
- Avg: CPI: Per Core weighted by core-hour
- Avg: CPLD: Per Core weighted by core-hour
- Avg: CPU User CV: weighted by core-hour
- Avg: CPU User Imbalance: weighted by core-hour (%)
- Avg: FLOPS: Per Core weighted by core-hour (ops/s)
- Avg: InfiniBand rate: Per Node weighted by node-hour (bytes/s)
- Avg: Memory Bandwidth: Per Core weighted by core-hour (bytes/s)
- Avg: Memory: Per Core weighted by core-hour (bytes)
- Avg: Total Memory: Per Core weighted by core-hour (bytes)
- Avg: block sda read ops rate: Per Node weighted by node-hour (ops/s)
- Avg: block sda read rate: Per Node weighted by node-hour (bytes/s)
- Avg: block sda write ops rate: Per Node weighted by node-hour (ops/s)
- Avg: block sda write rate: Per Node weighted by node-hour (bytes/s)
- Avg: eth0 receive rate: Per Node weighted by node-hour (bytes/s)
- Avg: eth0 transmit rate: Per Node weighted by node-hour (bytes/s)
- Avg: gpfs receive rate: Per Node weighted by node-hour (bytes/s)
- Avg: gpfs transmit rate: Per Node weighted by node-hour (bytes/s)
- Avg: ib0 receive rate: Per Node weighted by node-hour (bytes/s)
- Avg: ib0 transmit rate: Per Node weighted by node-hour (bytes/s)
- Avg: isilon receive rate: Per Node weighted by node-hour (bytes/s)

No data is available for viewing

Please refer to the instructions below:

1 Click on **New Chart** to get started and select either a **Timeseries** or **Aggregate** Chart, Then select how you want the data to be displayed. (i.e. Line, Bar...)

4 Select a data series in order to edit how the data is displayed (refer to 4a) or click anywhere else on the chart to display the **Chart Options** dialog for editing global chart options including adding data and/or filters. (refer to 4b)

(a) **Data Series: XSEDE**

- Show
- Drilldown
- Metric
- Group By
- Compare To
- Display
- Stacking

(b) **Chart Options:**

- Aggregate
- Timeseries
- Add Data
- Add Filter
- Legend
- Font Size
- Log Scale Chart

Data Series 1 - 1 of 1 Page 1 of 1 Show Remainder (Disable Paging) Page Limit 10 Data Series

Description

User Mode

XDMoD Hello, Broday Walker ([logout](#)) Dashboard My Profile About Roadmap Contact Us Help

Summary Usage Metric Explorer **Report Generator** Job Viewer About

My Reports

Select New New Based On Edit Preview Send Now Download ← Delete

Name	Derived From	Schedule	Delivery Format	# Charts
TAS Report 1	Manual	Once	PDF	1 (1 per page)

1 report

Available Charts

Select Delete

Chart

- CPU Hours: Per Job**
2019-06-20 to 2019-06-20
User Defined
- CPU Hours: Total**
2019-06-20 to 2019-06-20
User Defined
- Total CPU Hours by PI**
2019-06-19 to 2019-07-12
User Defined
- untitled query 3**
2019-04-07 to 2019-04-09
User Defined


4 charts

User Mode

No job is selected for viewing

Please refer to the instructions below:

1 Click on **Search** in the top left to get started, If you know the **Resource & Job Number** use the **Quick Lookup** form (2a) for more options use **Advanced Search** (2b)

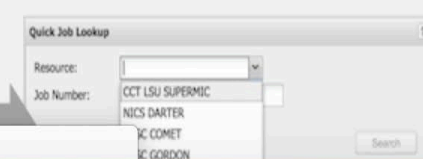


2b Within the **Advanced Search** form, select a timeframe and **Add** one or more filters. Click **Search** to run the search on the server.



Field	Operator	Value
Catastrophe Rank	=	0.1 - 0.5
Field of Science	=	Applied Mathematics

2a For a **Quick Job Lookup**, enter the **Resource** and **Job Number** and click **Search**



3 Select one or more **Jobs** and click **Save Results** to view data about the selected jobs.



Include	Job Id	Resource	Name
<input type="checkbox"/>	1245236	SDSC-COMET	Rumbell, Tim - Mt Sinai Sch of Med

User Mode

Search

Search Name:

Quick Job Lookup

Realm:

Resource:

Job Number:

Advanced Search

Start:

End:

Realm:

Filter: =

Field	Operator	Value
-------	----------	-------

Results

Include	Job Id	Resource	Name
---------	--------	----------	------

User Mode

- SUPREMM
 - search-2019-07-25-96
 - cheyenne-4954705
 - Accounting data**
 - Executable information
 - Summary metrics
 - Detailed metrics
 - search-2019-07-25-433

CPU User: 0.463 Homogeneity: N/A CPU User Balance: 0.79 Memory Headroom: N/A

Metric Missing: Not Available On The Metric Missing: Unknown Reason

Key	Value
Category: Administration (11 Items)	
Account	CESM0005
Local Job Id	4954705
Organization	The National Center for Atmospheric Research
Resource	cheyenne.ucar.edu
Hierarchy Bottom Level	Unknown
PI	CESM0005
PI Institution	The National Center for Atmospheric Research
User	cmip6
User Institution	The National Center for Atmospheric Research

Description

- Accounting data:** Shows information about the job that was obtained from the resource manager. This includes timing information such as the start and end time of the job as well as administrative information such as the user that submitted the job and the account that was charged.

User Mode

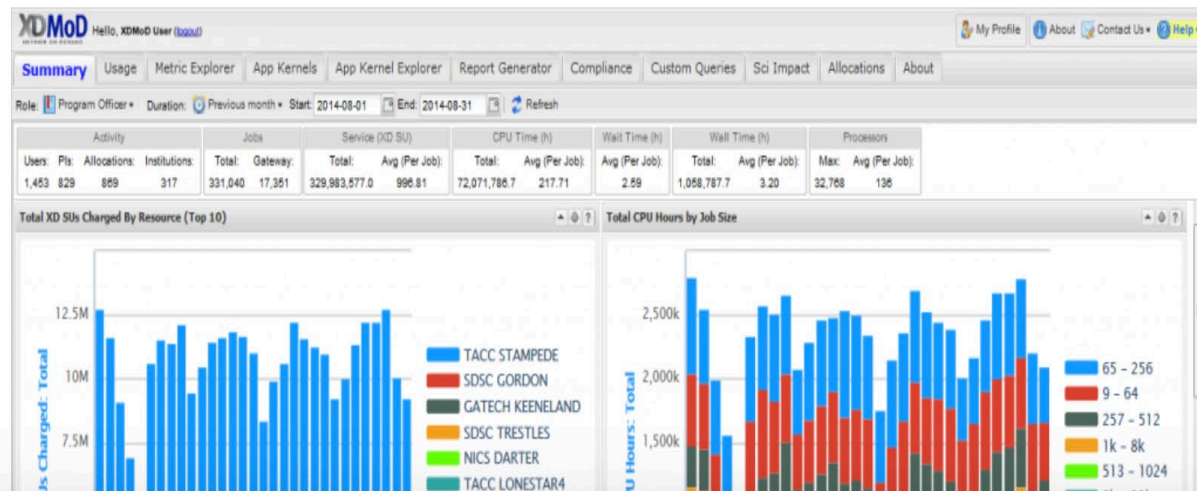


XDMoD: Comprehensive HPC System Management Tool

The University at Buffalo Center for Computational Research (CCR) has been at the forefront of the development of open source tools for use by national and campus level high performance computing (HPC) centers to help ensure their optimal operation as well as provide metrics to demonstrate the utility, service, competitive advantage, and return on investment that these centers provide.

The XDMoD (XD Metrics on Demand) tool provides HPC center personnel and senior leadership with the ability to easily obtain detailed operational metrics of HPC systems coupled with extensive analytical capability to optimize performance at the system and job level, ensure quality of service, and provide accurate data to guide system upgrades and acquisitions.

- XDMoD
- Open XDMoD
- SUPReMM
- Federated
- Roadmap
- Team
- Publications
- Presentations
- Links
- Release Notes



References

- Cheyenne. (n.d.). Retrieved from https://www2.cisl.ucar.edu/resources/computational_systems/cheyenne
- Cloud, Mobility, Networking & Security Solutions. (2019, July 26). Retrieved from [https:// www.vmware.com /](https://www.vmware.com/)
- Ganglia Monitoring System. (n.d.). Retrieved from [http:// ganglia.sourceforge.net/](http://ganglia.sourceforge.net/)
- XDMoD User Manual. (n.d.). Retrieved from https://xdmod.ccr.buffalo.edu/user_manual/
- XDMoD with SUPReMM set up procedure. (n.d.). Retrieved from [https://wiki.ucar.edu/display/csg/XDMoD with SUPReMM set up procedure](https://wiki.ucar.edu/display/csg/XDMoD+with+SUPReMM+set+up+procedure)

Acknowledgements

Many Thanks to:

- The University at Buffalo
- NCAR
- CSG
- EIO
- SSG
- SIParCS Program

A Special Thanks to:

- AJ Lauer
- Virginia Do
- Elliott Foust
- Blake Lewis
- Shiquan Su
- Mick Coady
- Broday Walker