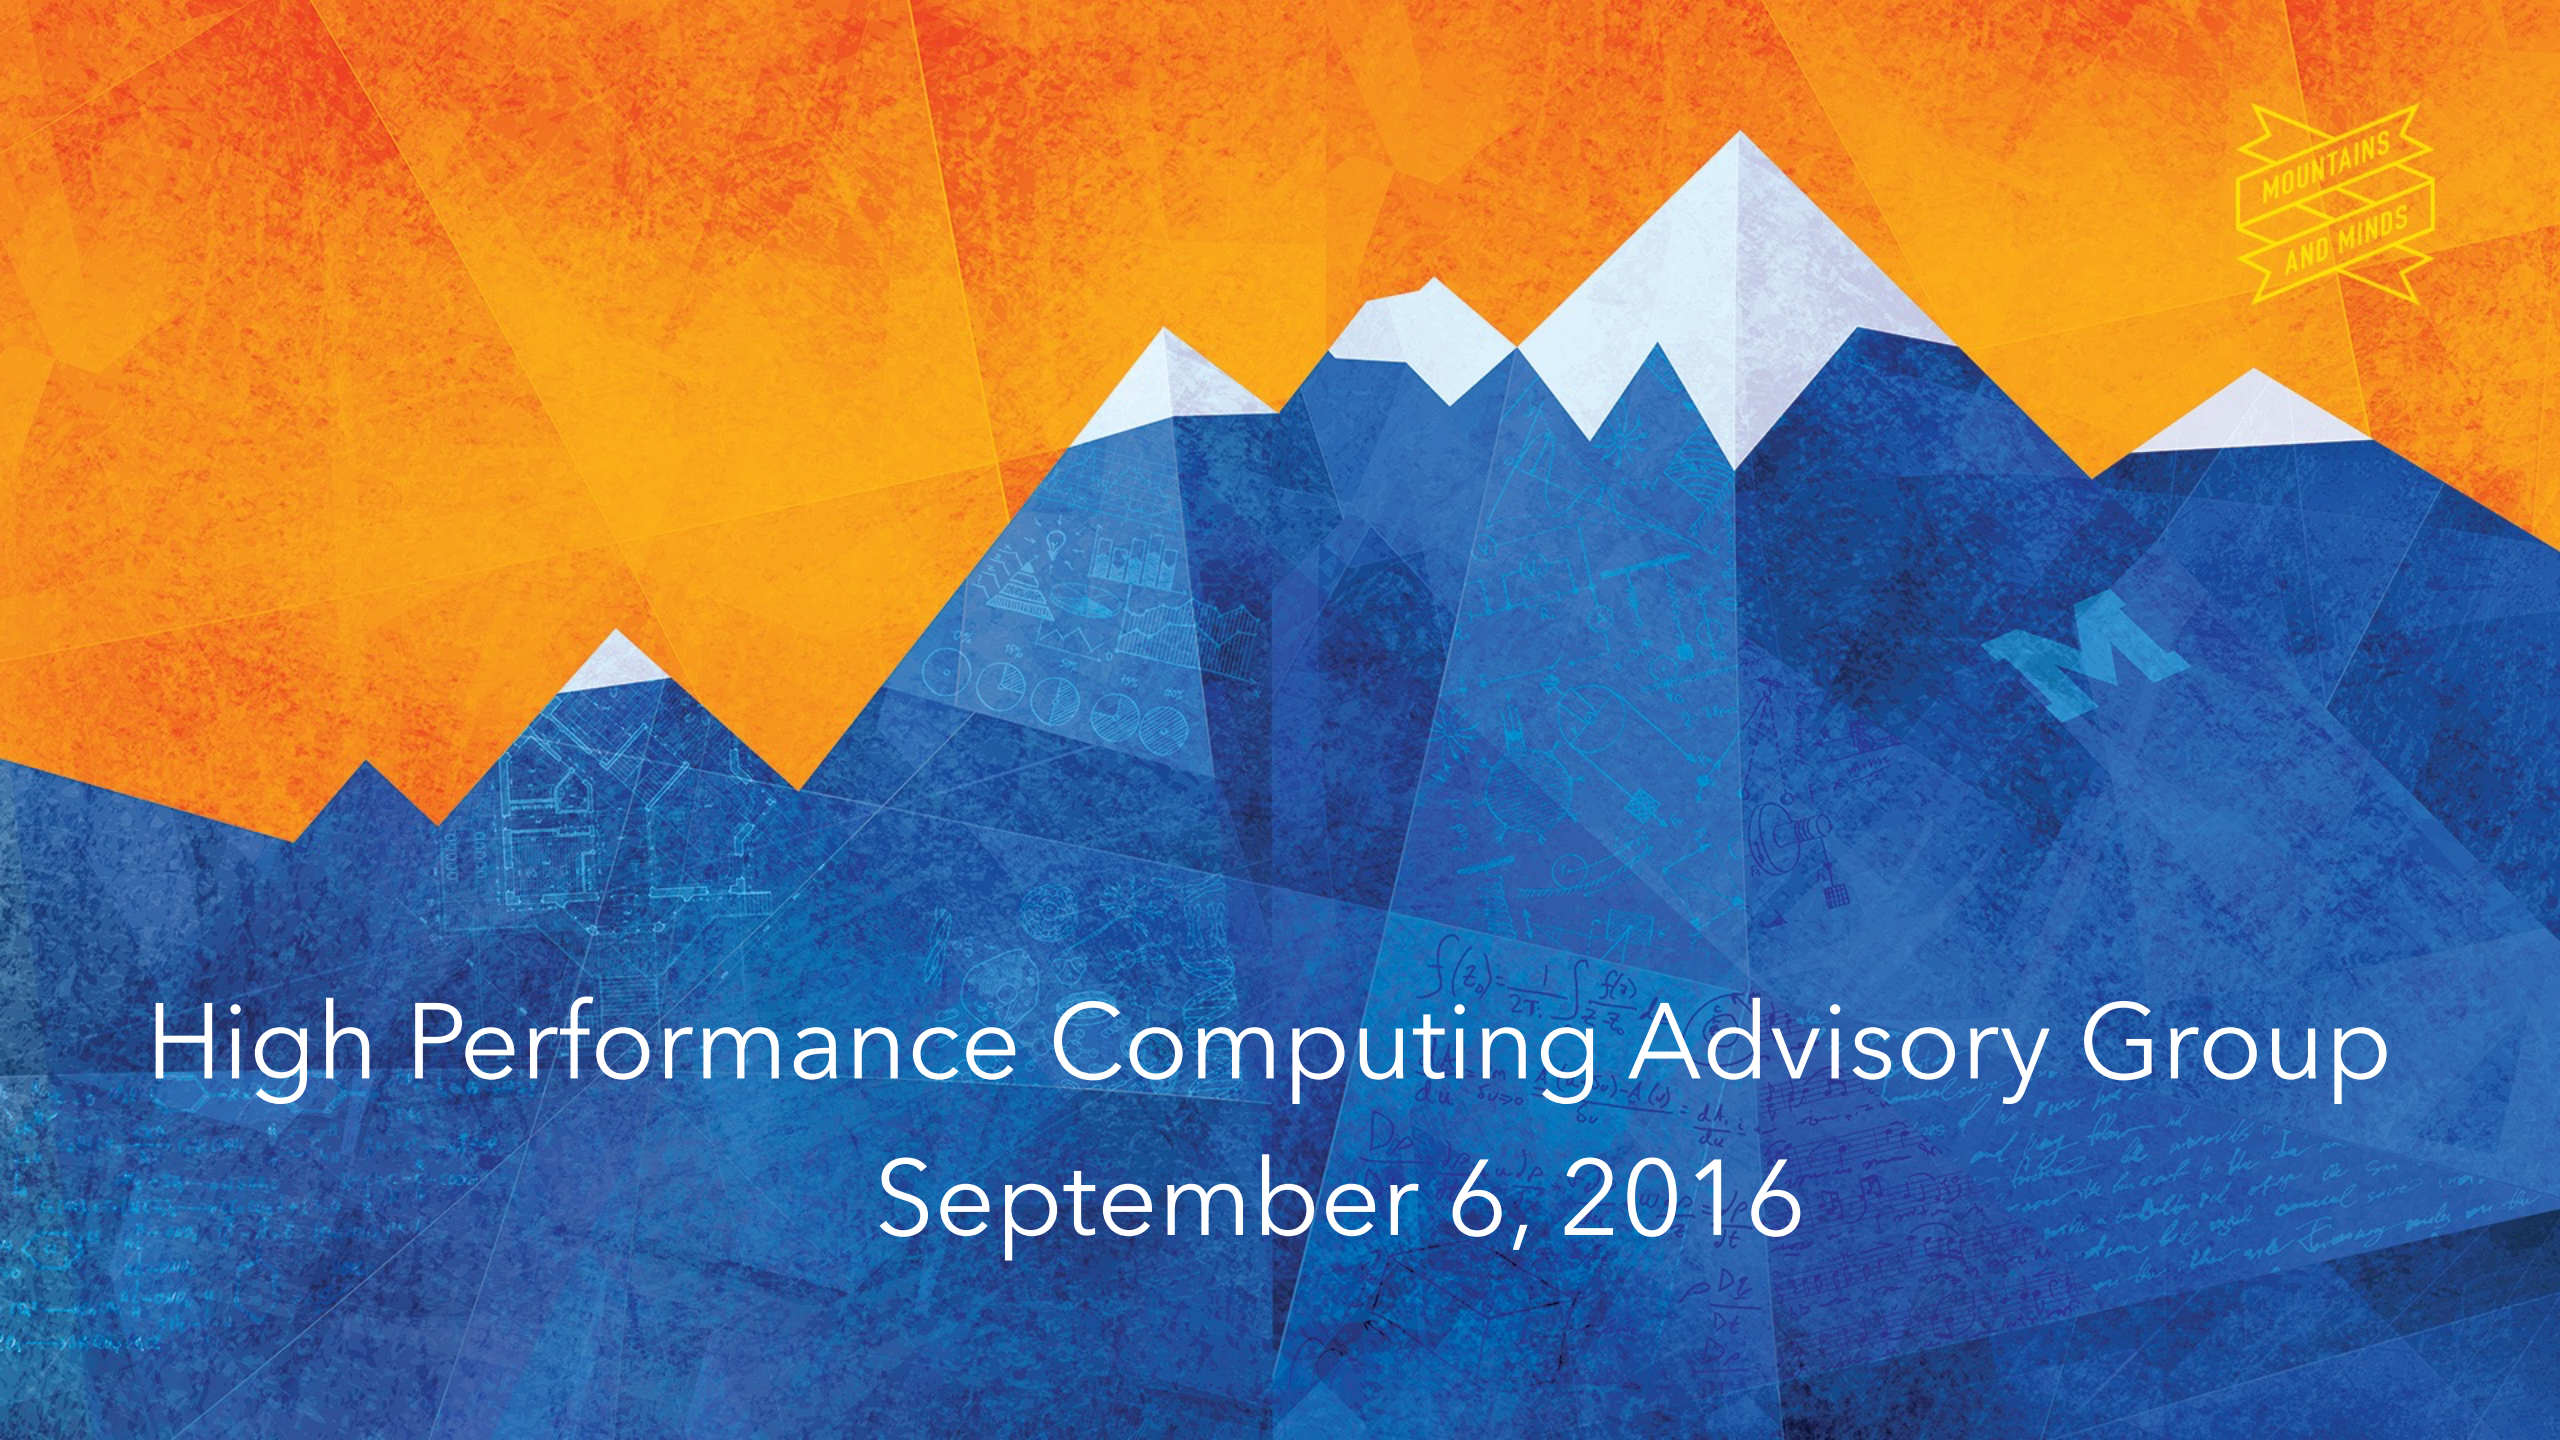

High Performance Computing Advisory Group

September 6, 2016

Hyalite Expansion

16 new Hyalite Nodes were installed and provisioned in June

- We had an issue with one of the nodes (Compute50) that was resolved by replacing a bad hard drive during the maintenance window last week.

New Cluster Overview:

- 60 Nodes (Xeon, 36 Sandy Bridge, 24 Haswell)
- 16 cores per Node for a total of 960 cores (1920 HT)
- 4 GB Ram per core
- 620 TB of Lustre scratch storage
- 10 GbE fabric w/ RDMA

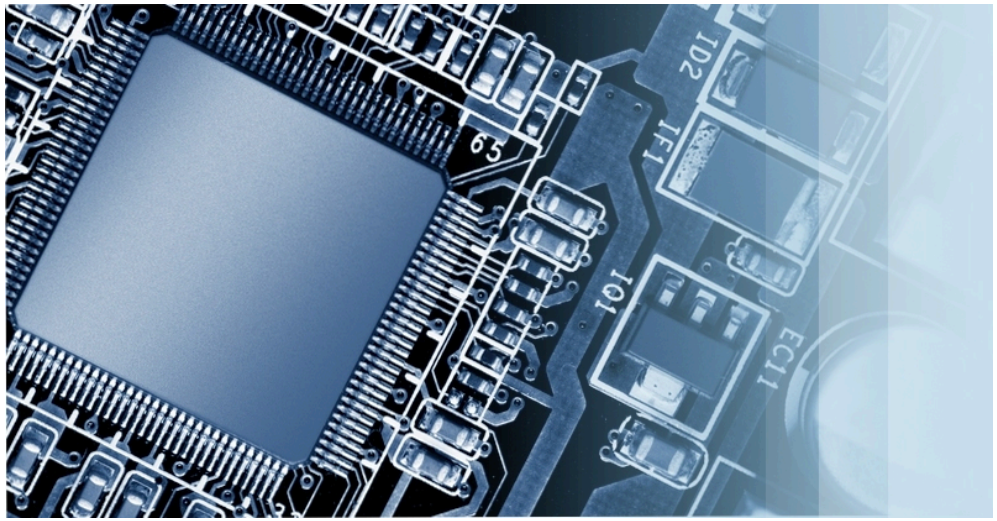

MONTANA STATE UNIVERSITY

2016-09 Maintenance

- Lustre updated to 2.5.42
- 10GbE Network Drivers Updated
- Migrated IPMI to new network
- RobinHood installed and initialized
- RDMA installed
  - <64byte latency reduced from 8.5 to 1.9 microseconds
  - Similar to QDR infiniband

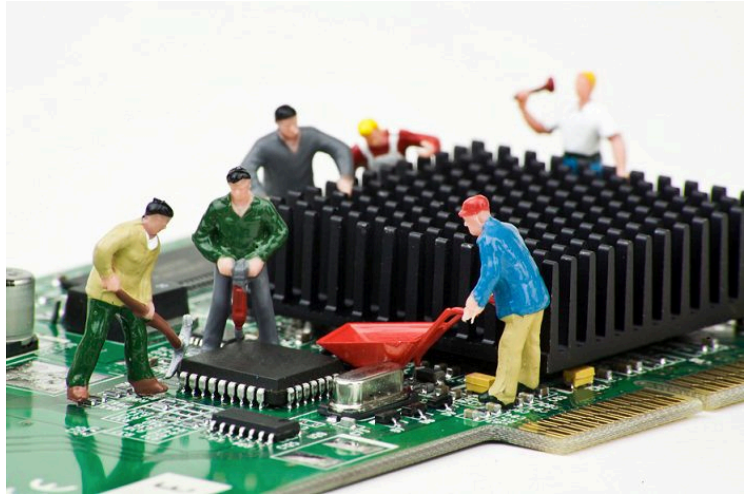

MONTANA STATE UNIVERSITY

XDMoD Stats – 2016 UPDATE

- XDMoD Stats
  - Active Users: 35
  - Total Jobs: 85,409
  - Total CPU Hours: 1,687,511
  - Average Job: 19.86 hours
  - Average Wait Time: 20 hours

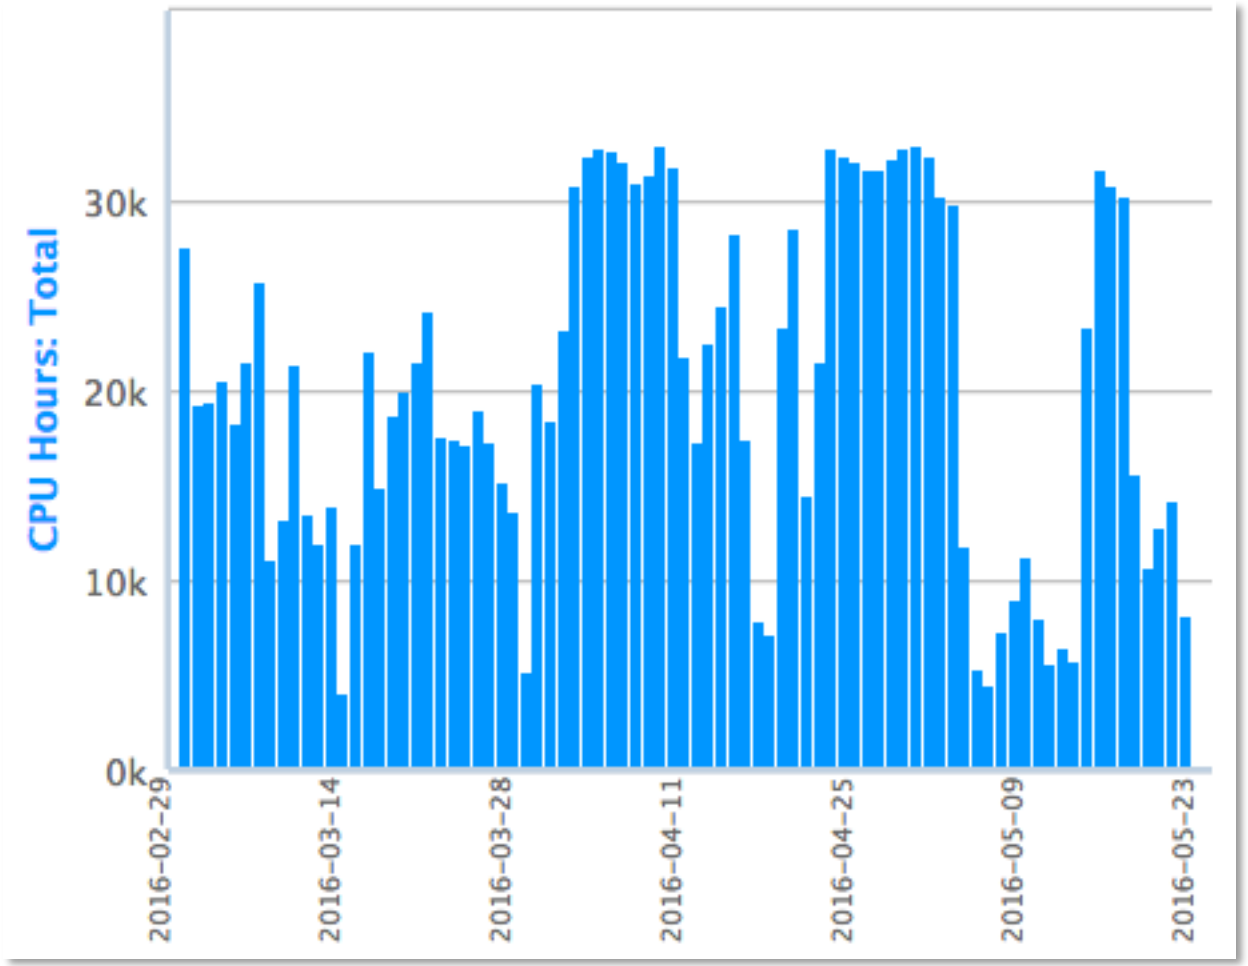

MONTANA STATE UNIVERSITY

Who has used the most time?

//Jobs by User/Wall Hours: Total

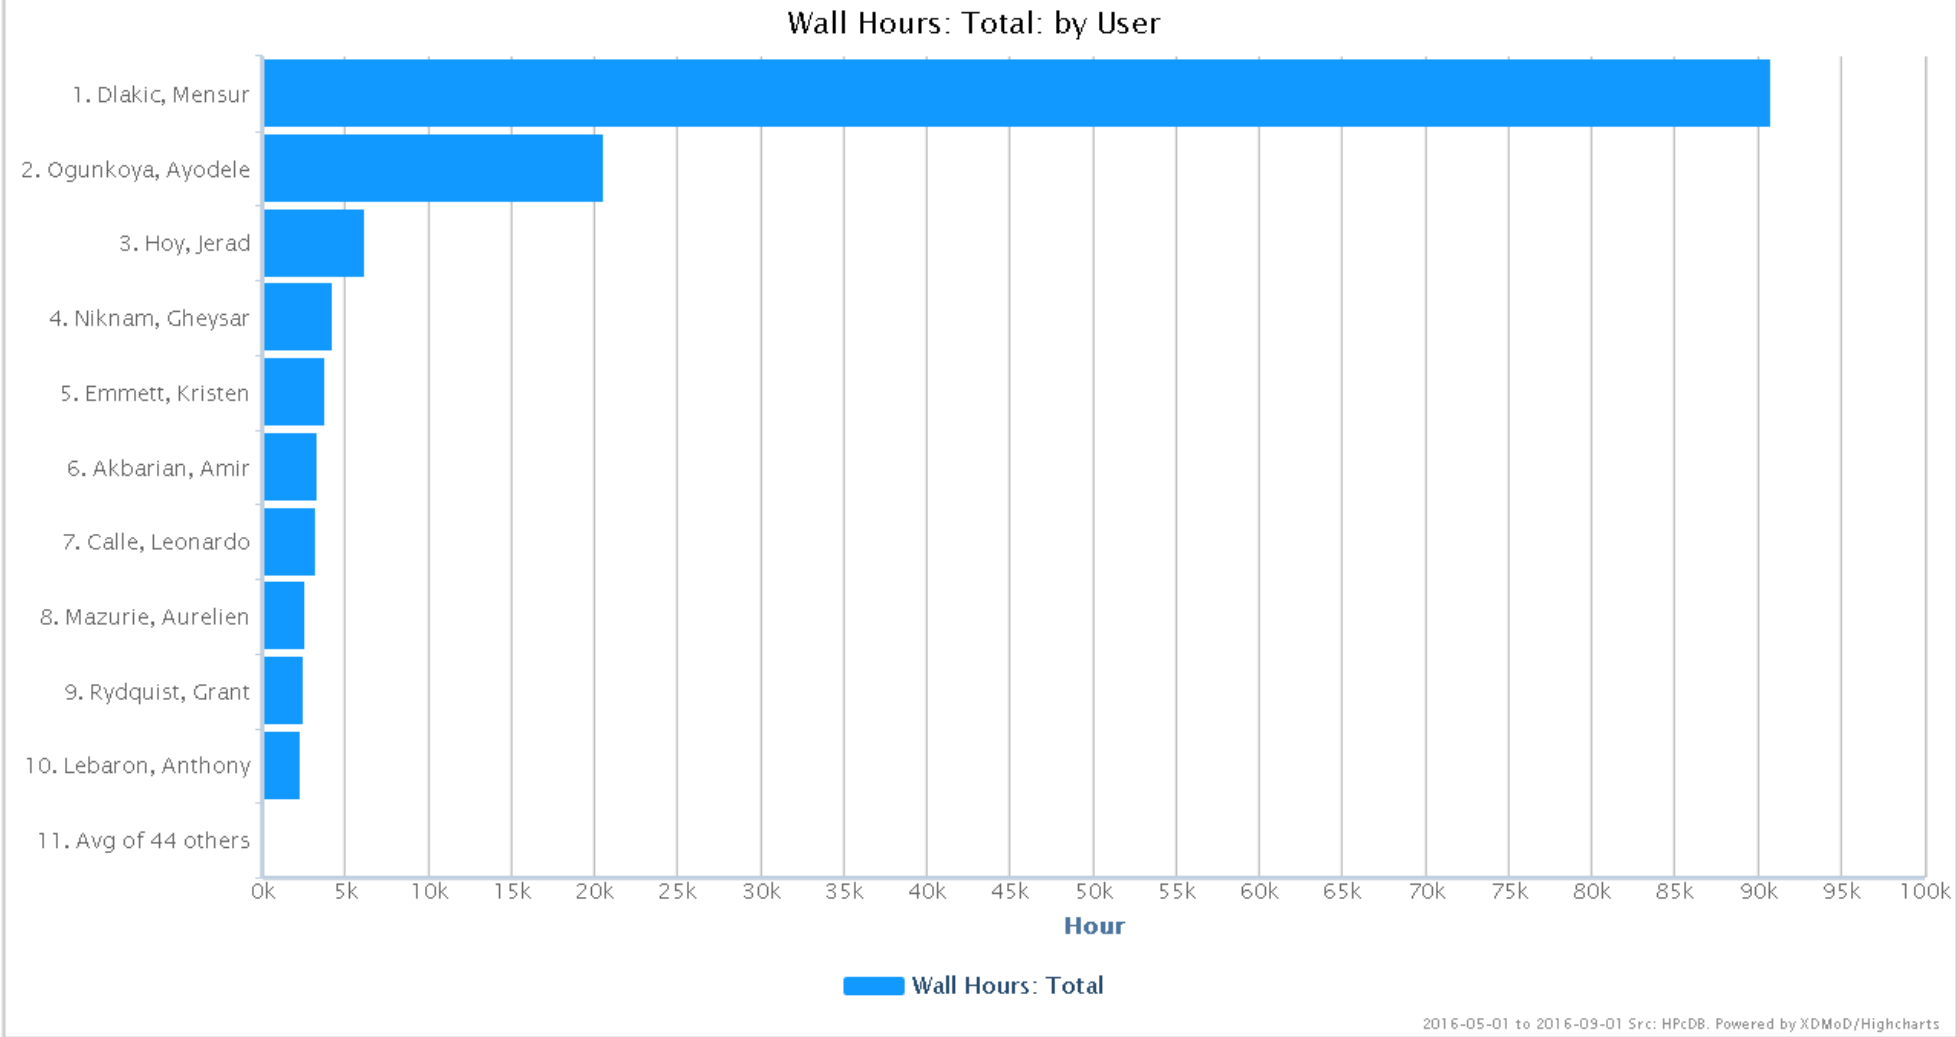

Description

- **User:** A person who is on a PIs allocation, hence able run jobs on resources.
- **Wall Hours: Total**: The total time, in hours, Information Technology Services Research Cyberinfrastructure jobs took to execute.
  Wall Time: Wall time is defined as the linear time between start and end time of execution for a particular job.

MONTANA STATE UNIVERSITY

Who is running long (time) jobs?

//Jobs by User/Wall Hours: Per Job

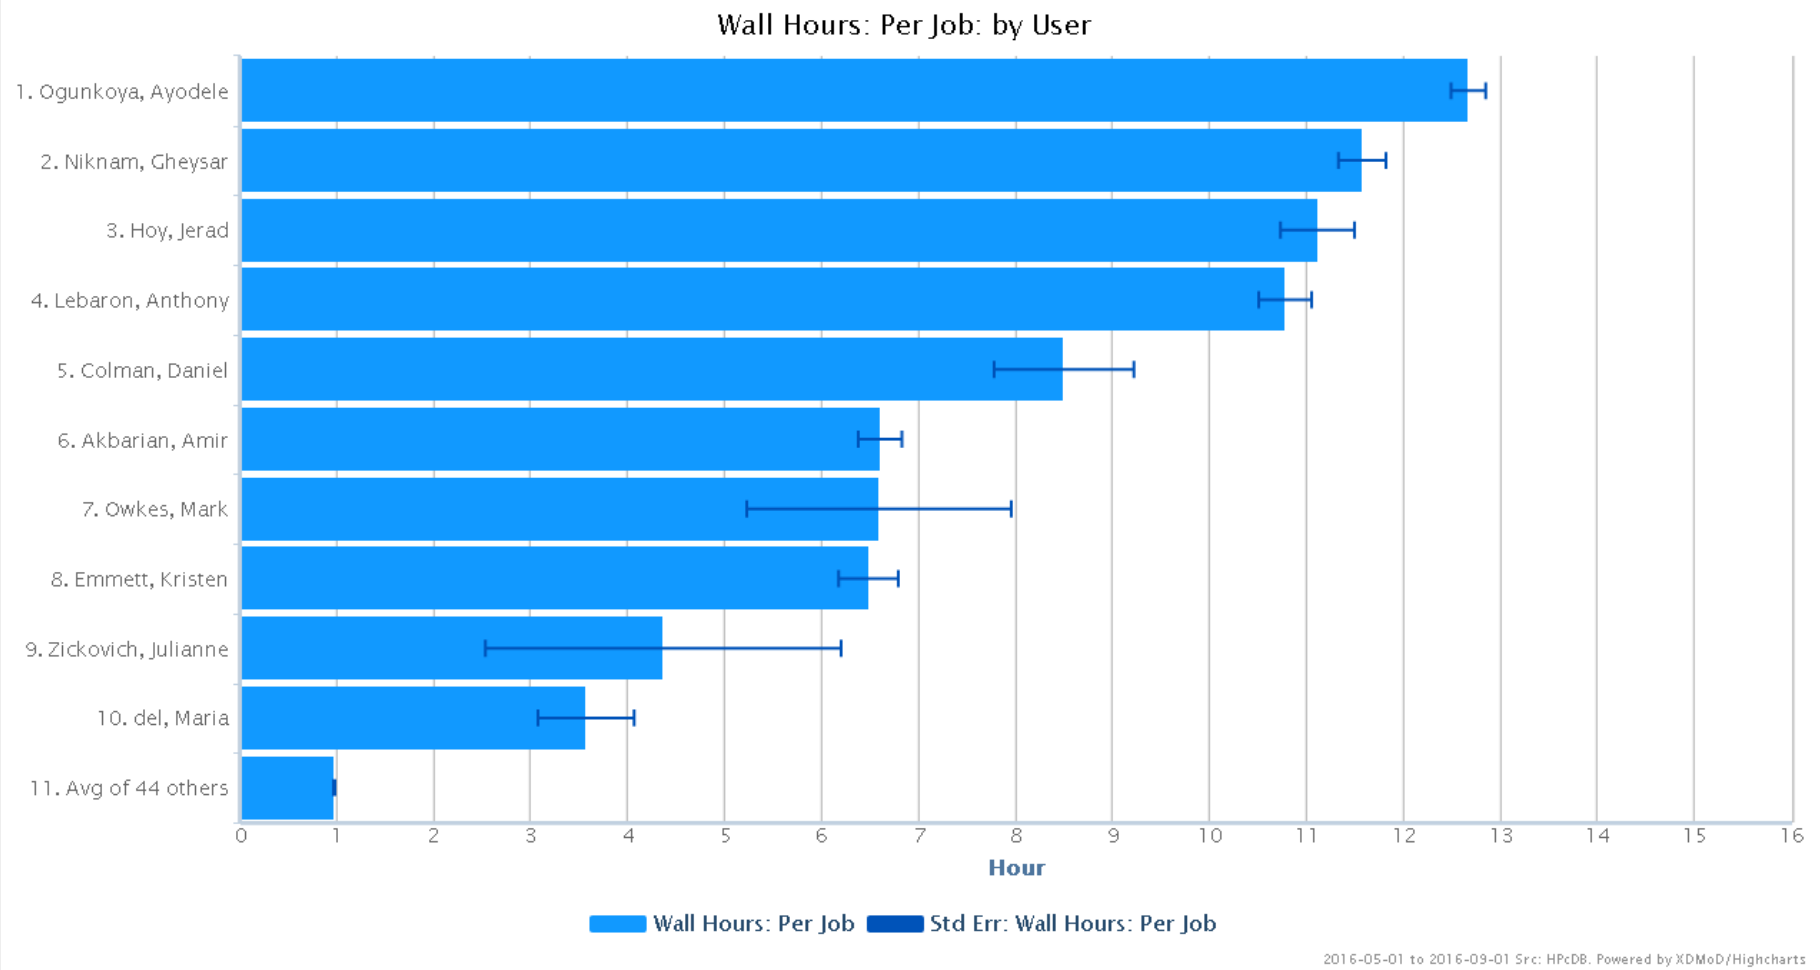

Description

- **User:** A person who is on a PIs allocation, hence able run jobs on resources.
- **Wall Hours: Per Job**: The average time, in hours, a Information Technology Services Research Cyberinfrastructure job takes to execute.
  Wall Time: Wall time is defined as the linear time between start and end time of execution for a particular job.

MONTANA STATE UNIVERSITY

Who is waiting the most?

//Jobs by User/Wait Hours: Total

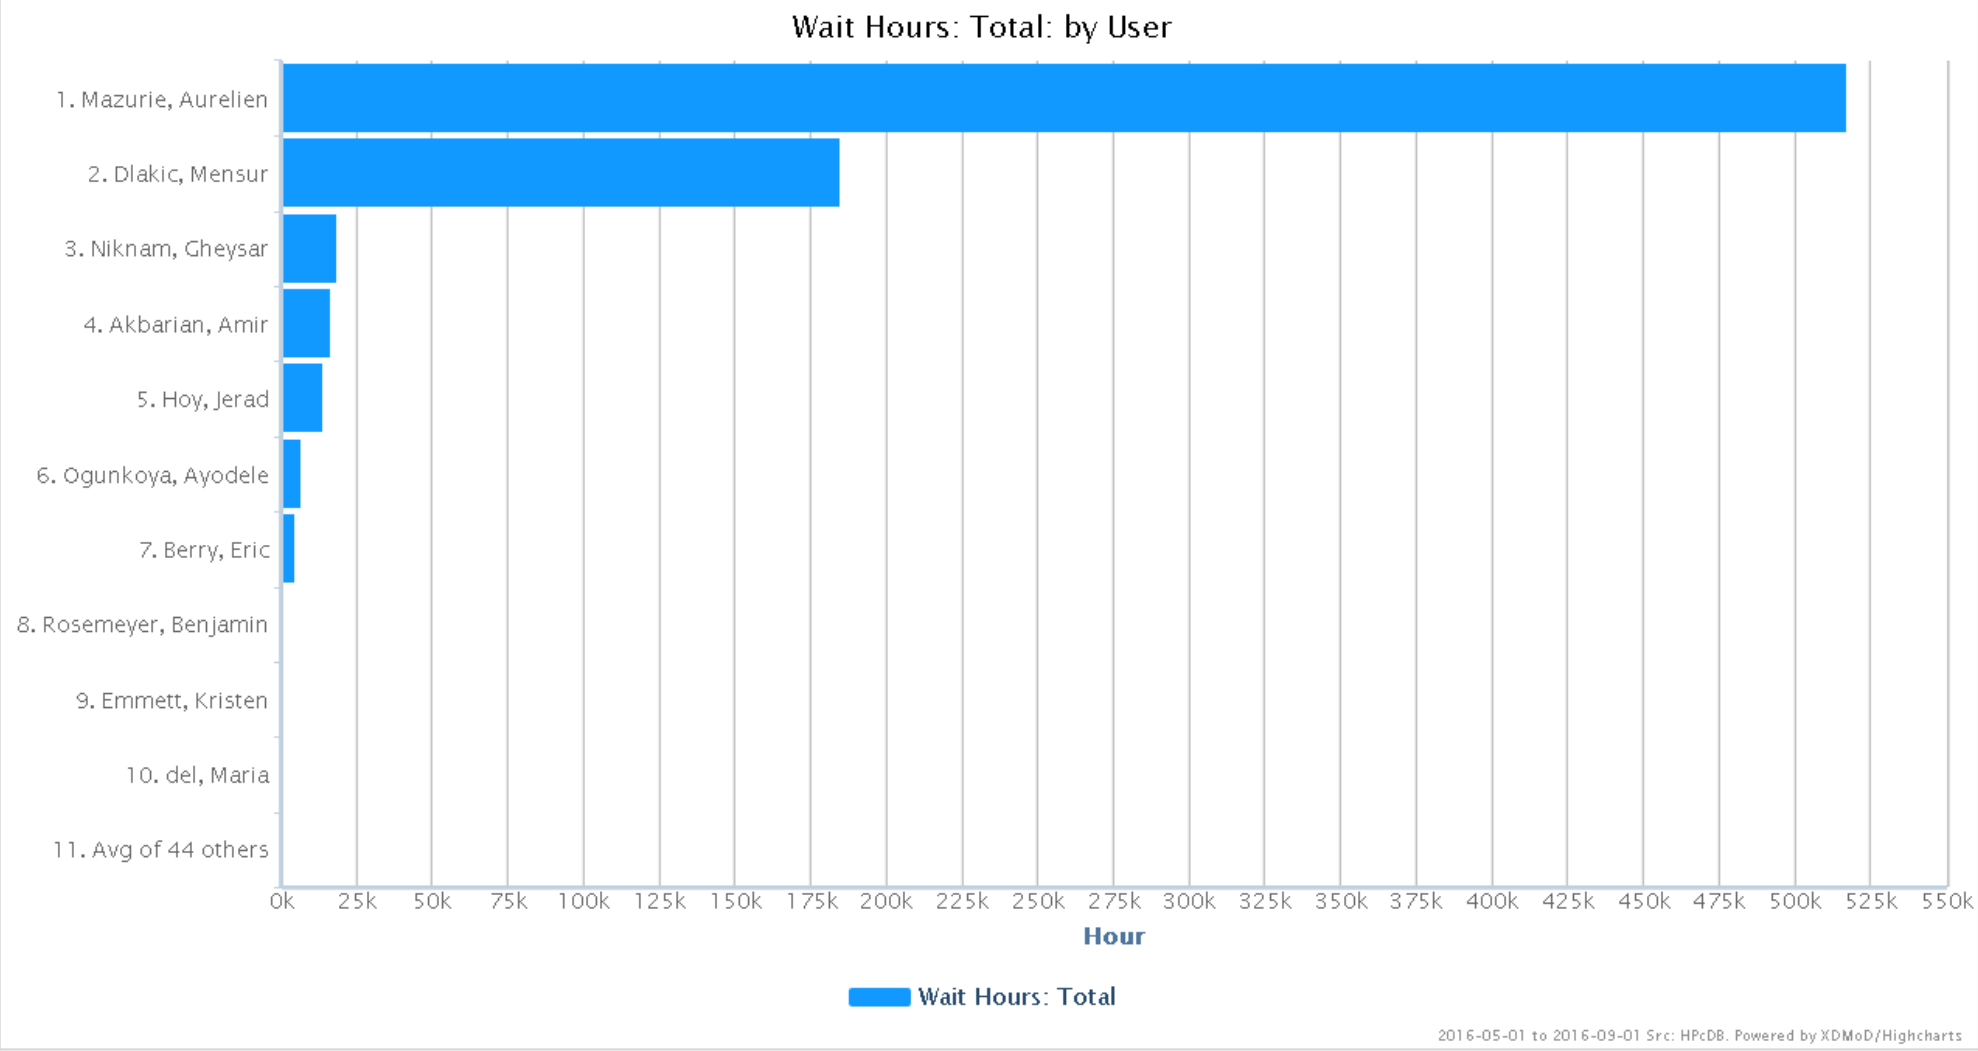

Description

- **User:** A person who is on a PIs allocation, hence able run jobs on resources.
- **Wait Hours: Total**: The total time, in hours, Information Technology Services Research Cyberinfrastructure jobs waited before execution on their designated resource.
  Wait Time: Wait time is defined as the linear time between submission of a job by a user until it begins to execute.

MONTANA STATE UNIVERSITY

Who waits the most per job?

//Jobs by User/Wait Hours: Per Job

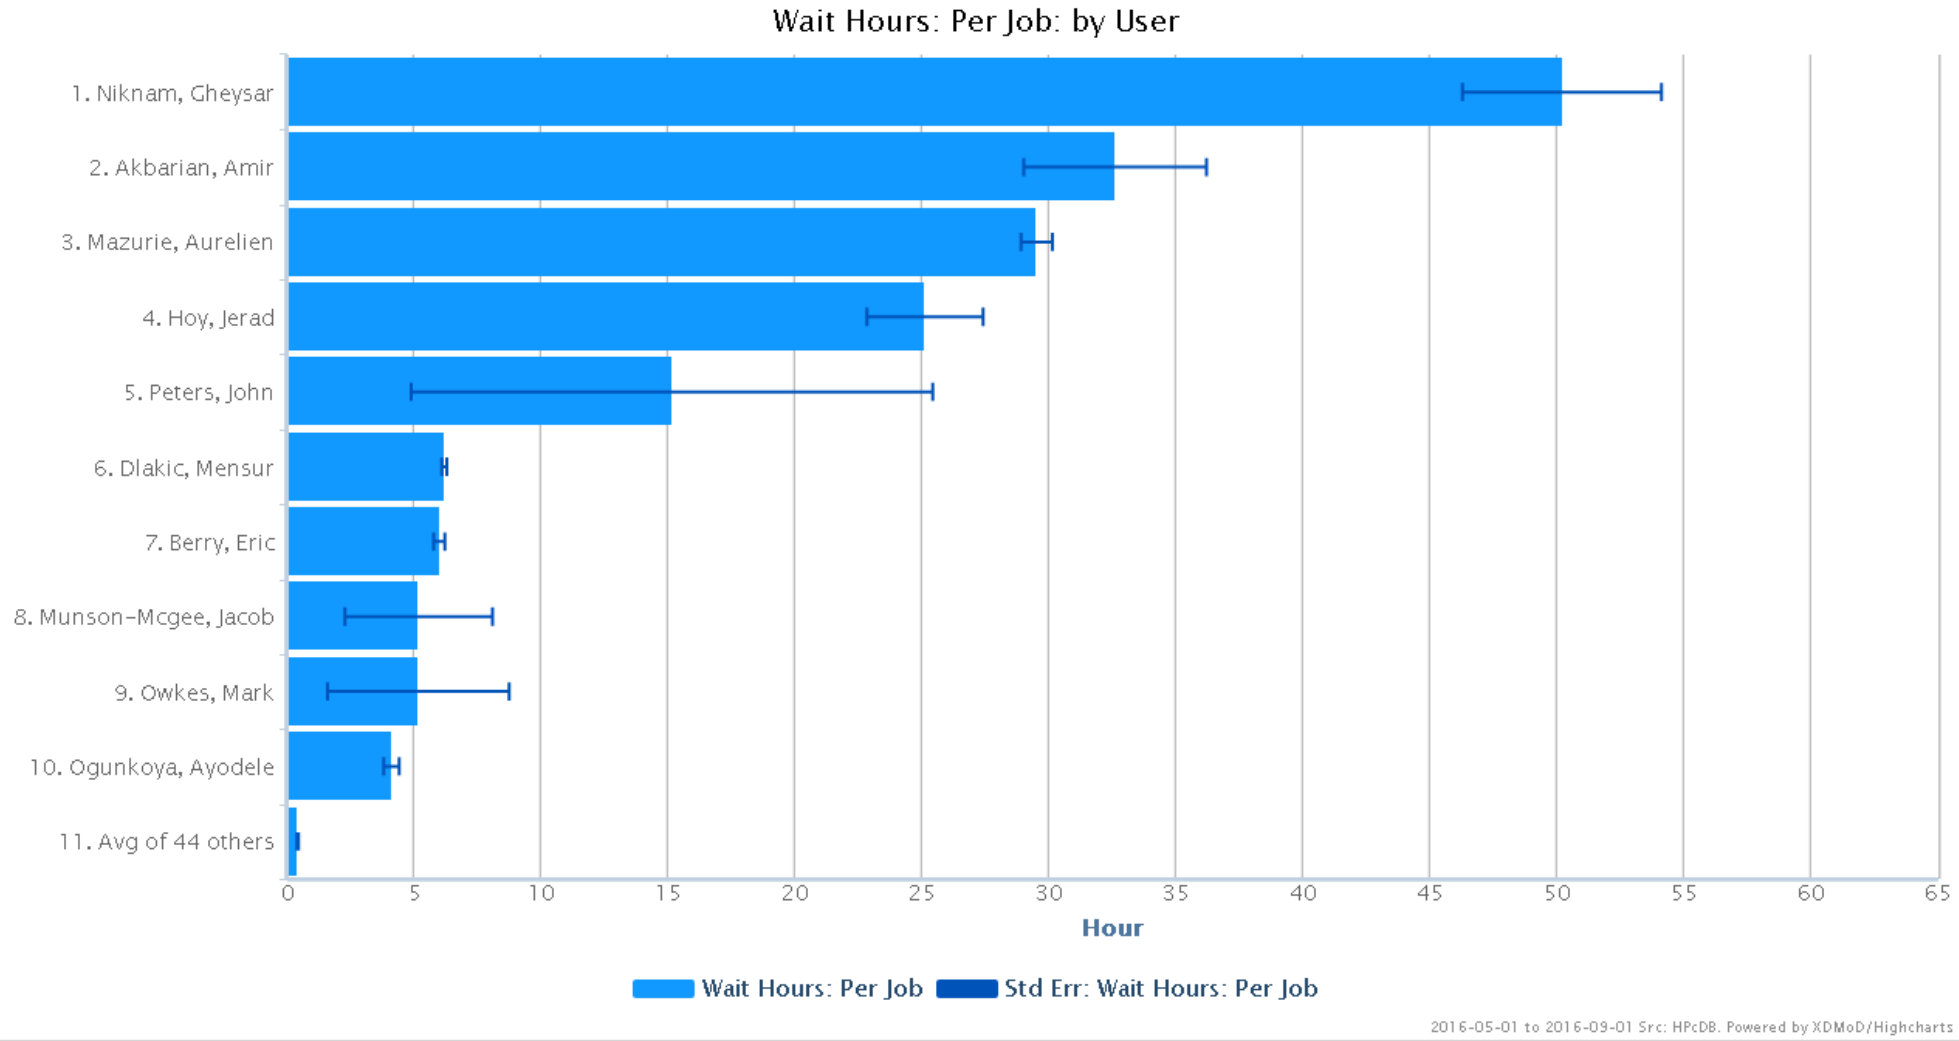

Description

- **User:** A person who is on a PIs allocation, hence able run jobs on resources.
- **Wait Hours: Per Job**: The average time, in hours, a Information Technology Services Research Cyberinfrastructure job waits before execution on the designated resource.
  Wait Time: Wait time is defined as the linear time between submission of a job by a user until it begins to execute.

MONTANA STATE UNIVERSITY

Waiting hours vs job length:

//Jobs by Job Wall Time/Wait Hours: Per Job

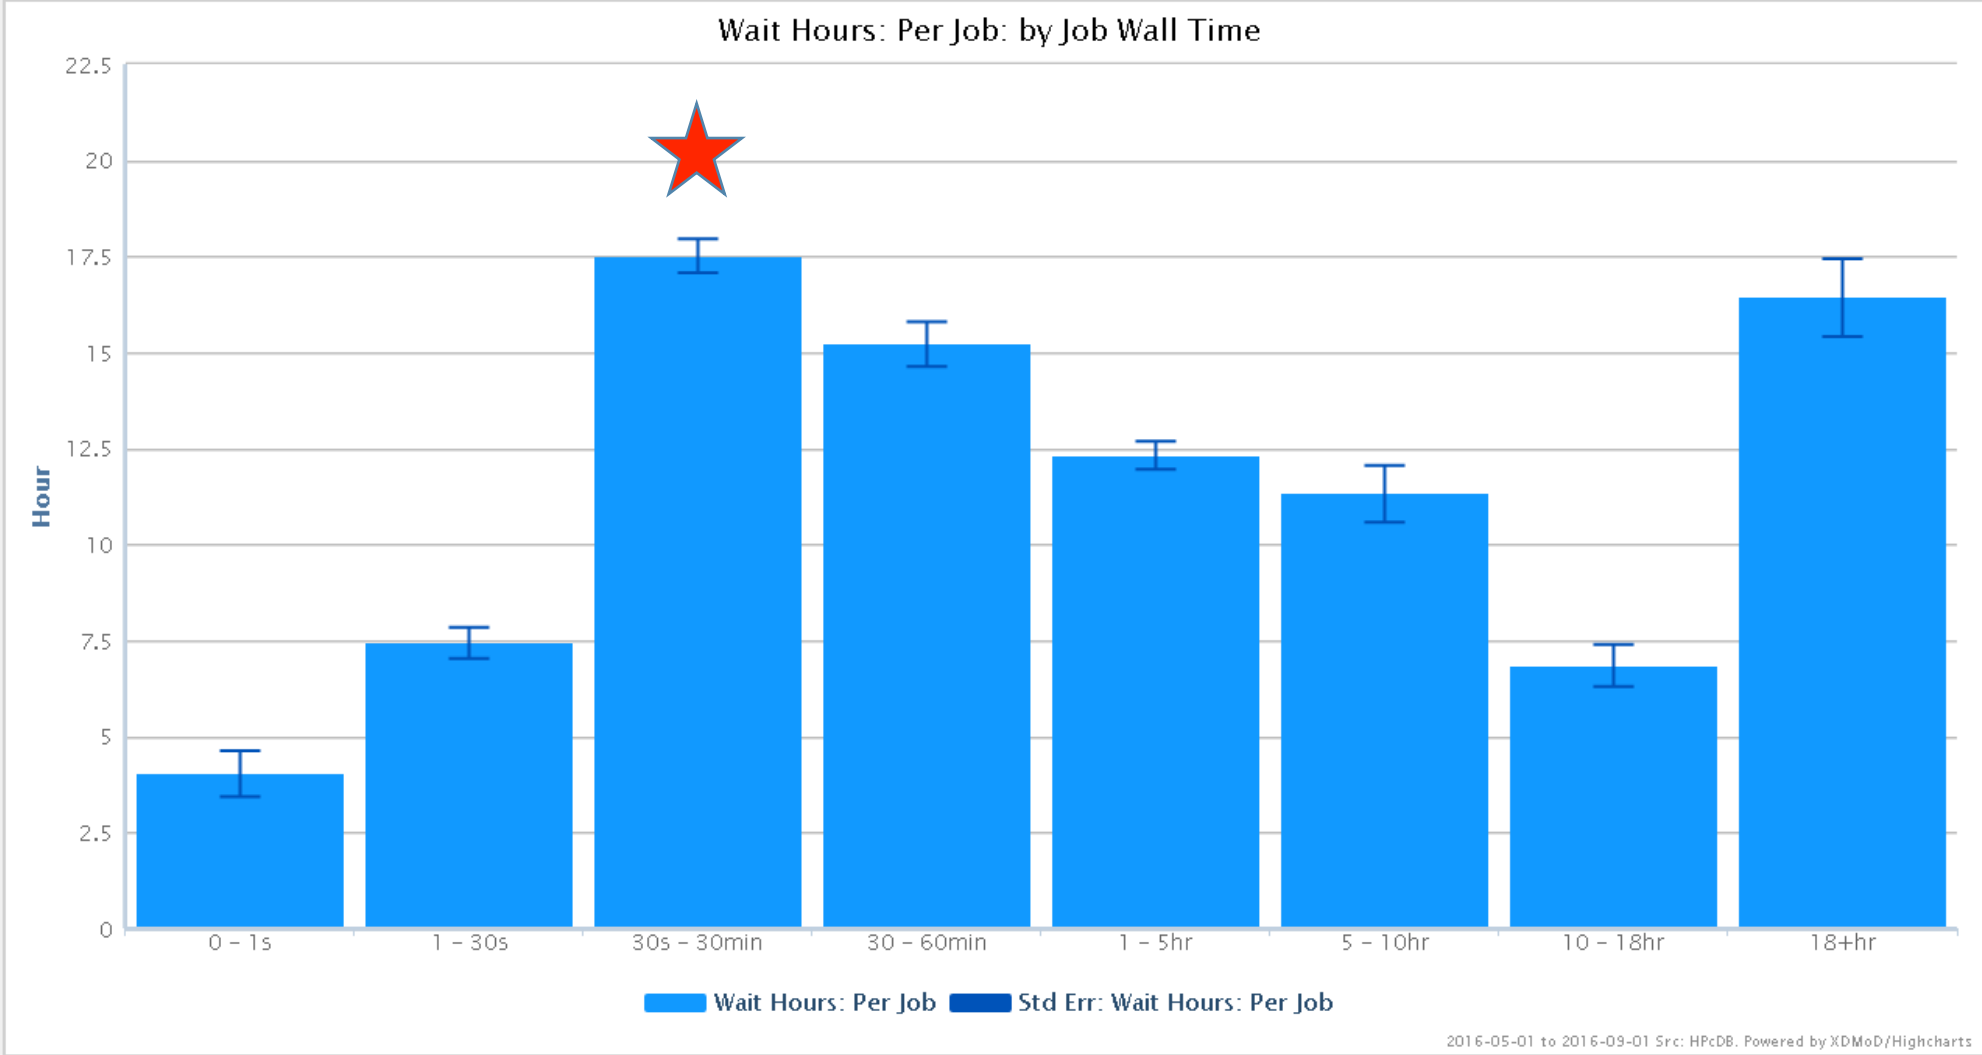

Description

- **Job Wall Time:** A categorization of jobs into discrete groups based on the total linear time each job took to execute.
- **Wait Hours: Per Job**: The average time, in hours, a Information Technology Services Research Cyberinfrastructure job waits before execution on the designated resource.
  Wait Time: Wait time is defined as the linear time between submission of a job by a user until it begins to execute.

MONTANA STATE UNIVERSITY

Waiting time (for jobs done in 0.5-30min)

//Jobs by Job Wall Time/Wait Hours: Per Job/by User

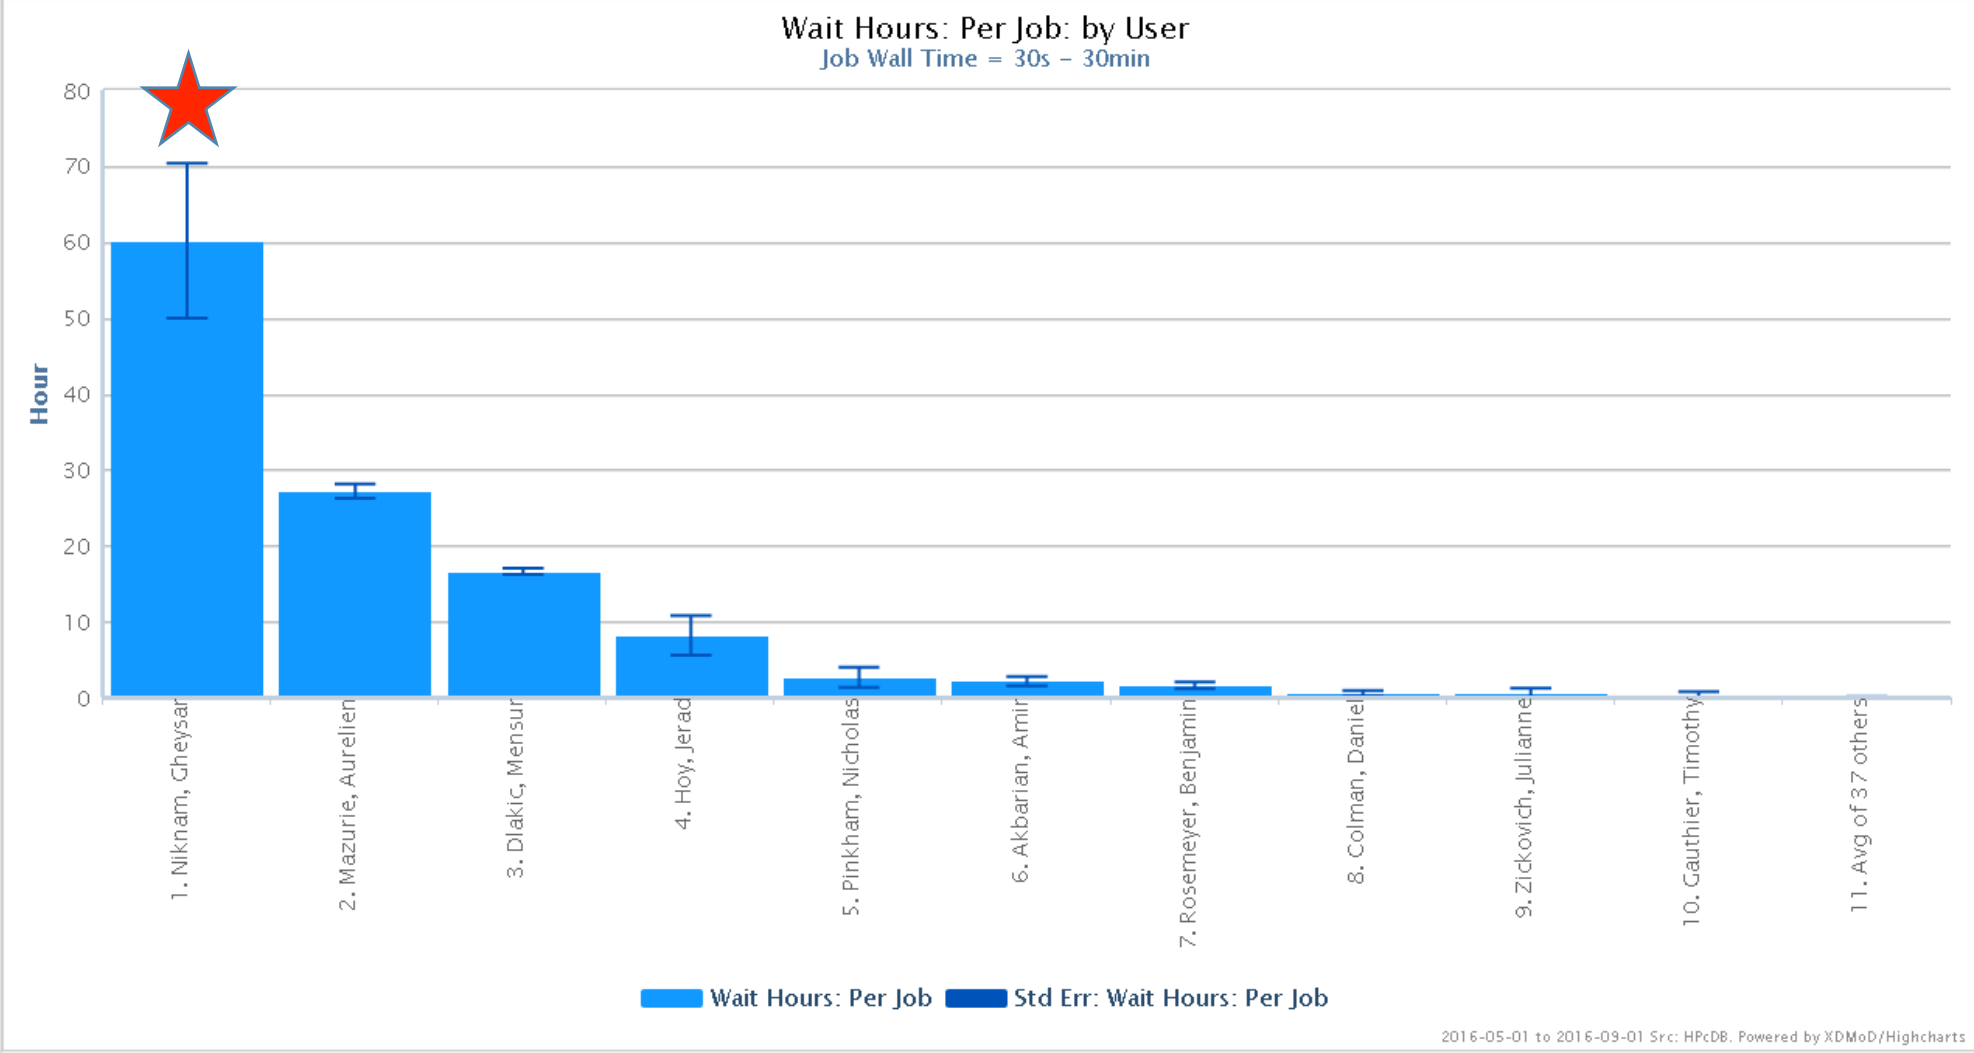

Description

- **User:** A person who is on a PIs allocation, hence able run jobs on resources.
- **Wait Hours: Per Job**: The average time, in hours, a Information Technology Services Research Cyberinfrastructure job waits before execution on the designated resource.
  Wait Time: Wait time is defined as the linear time between submission of a job by a user until it begins to execute.

MONTANA STATE UNIVERSITY

Example history of cpu-hours

//Jobs by User/CPU Hours: Per Job/by Job Size

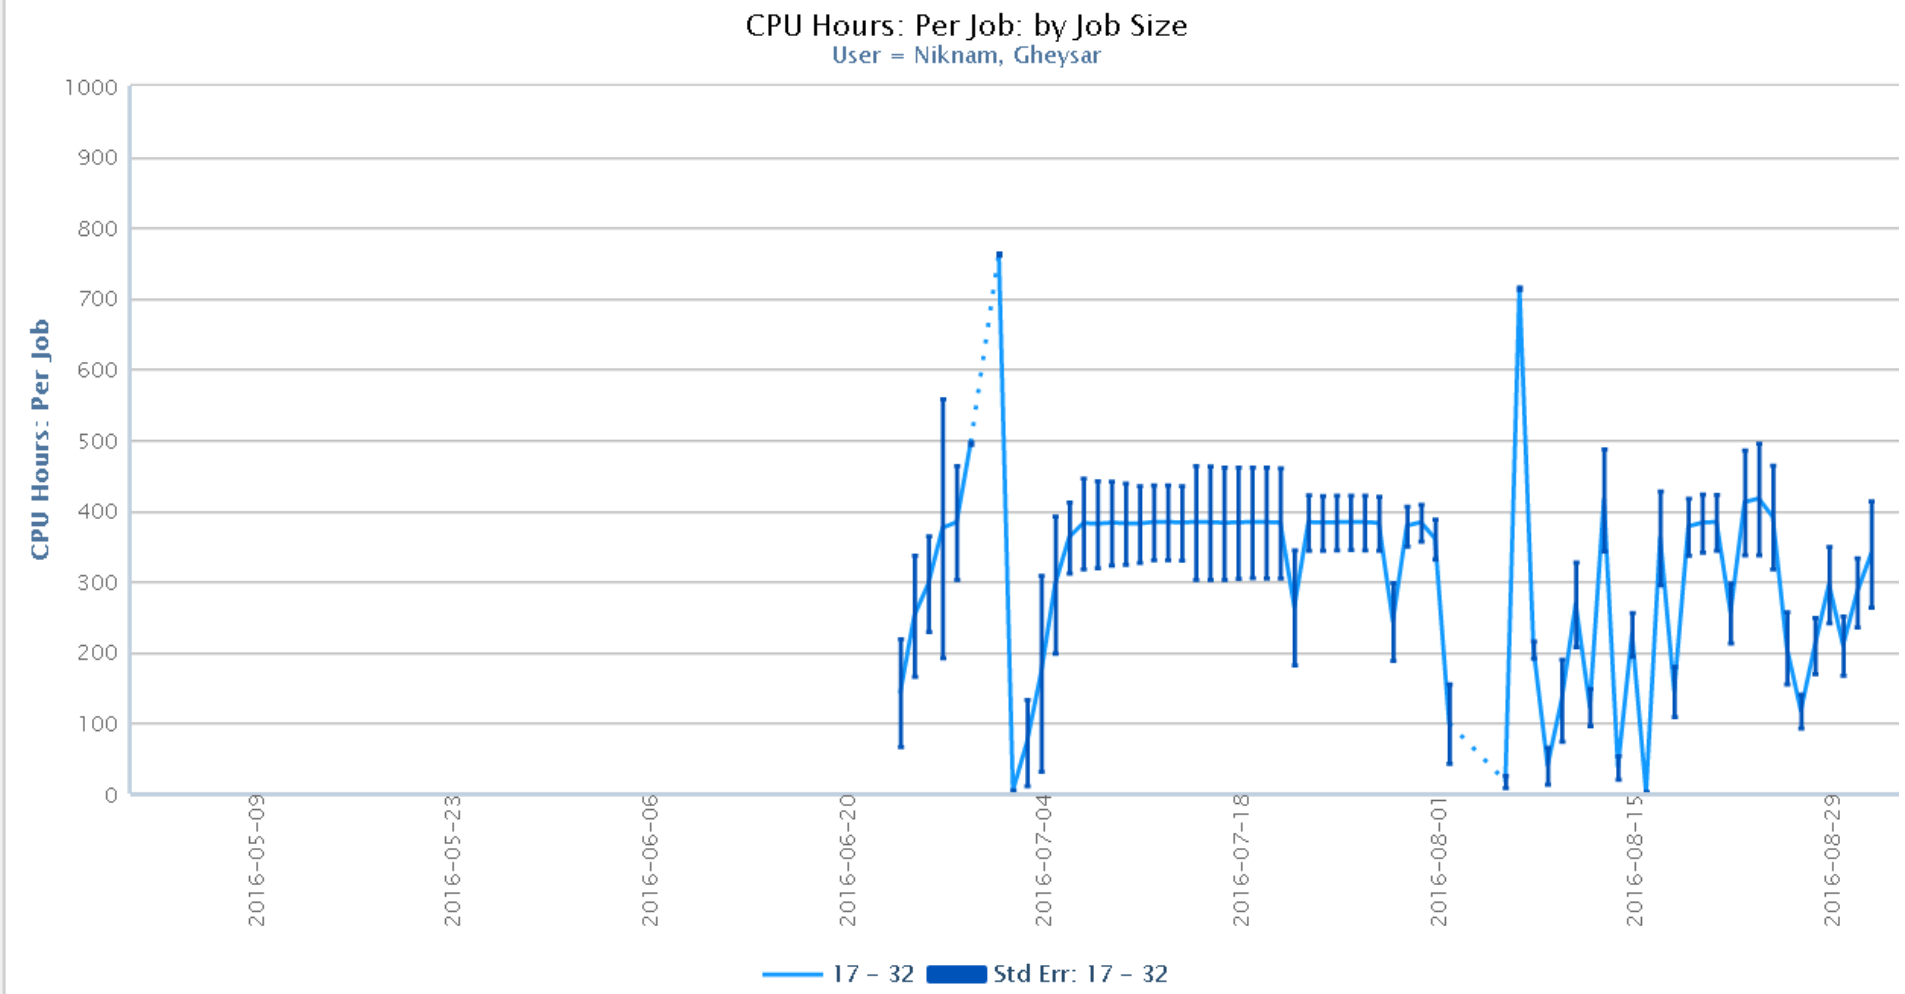

Description

- **Job Size:** A categorization of jobs into discrete groups based on the number of cores used by each job.
- **CPU Hours: Per Job**: The average CPU hours (number of CPU cores x wall time hours) per Information Technology Services Research Cyberinfrastructure 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.

MONTANA STATE UNIVERSITY

Details on wait time, user example

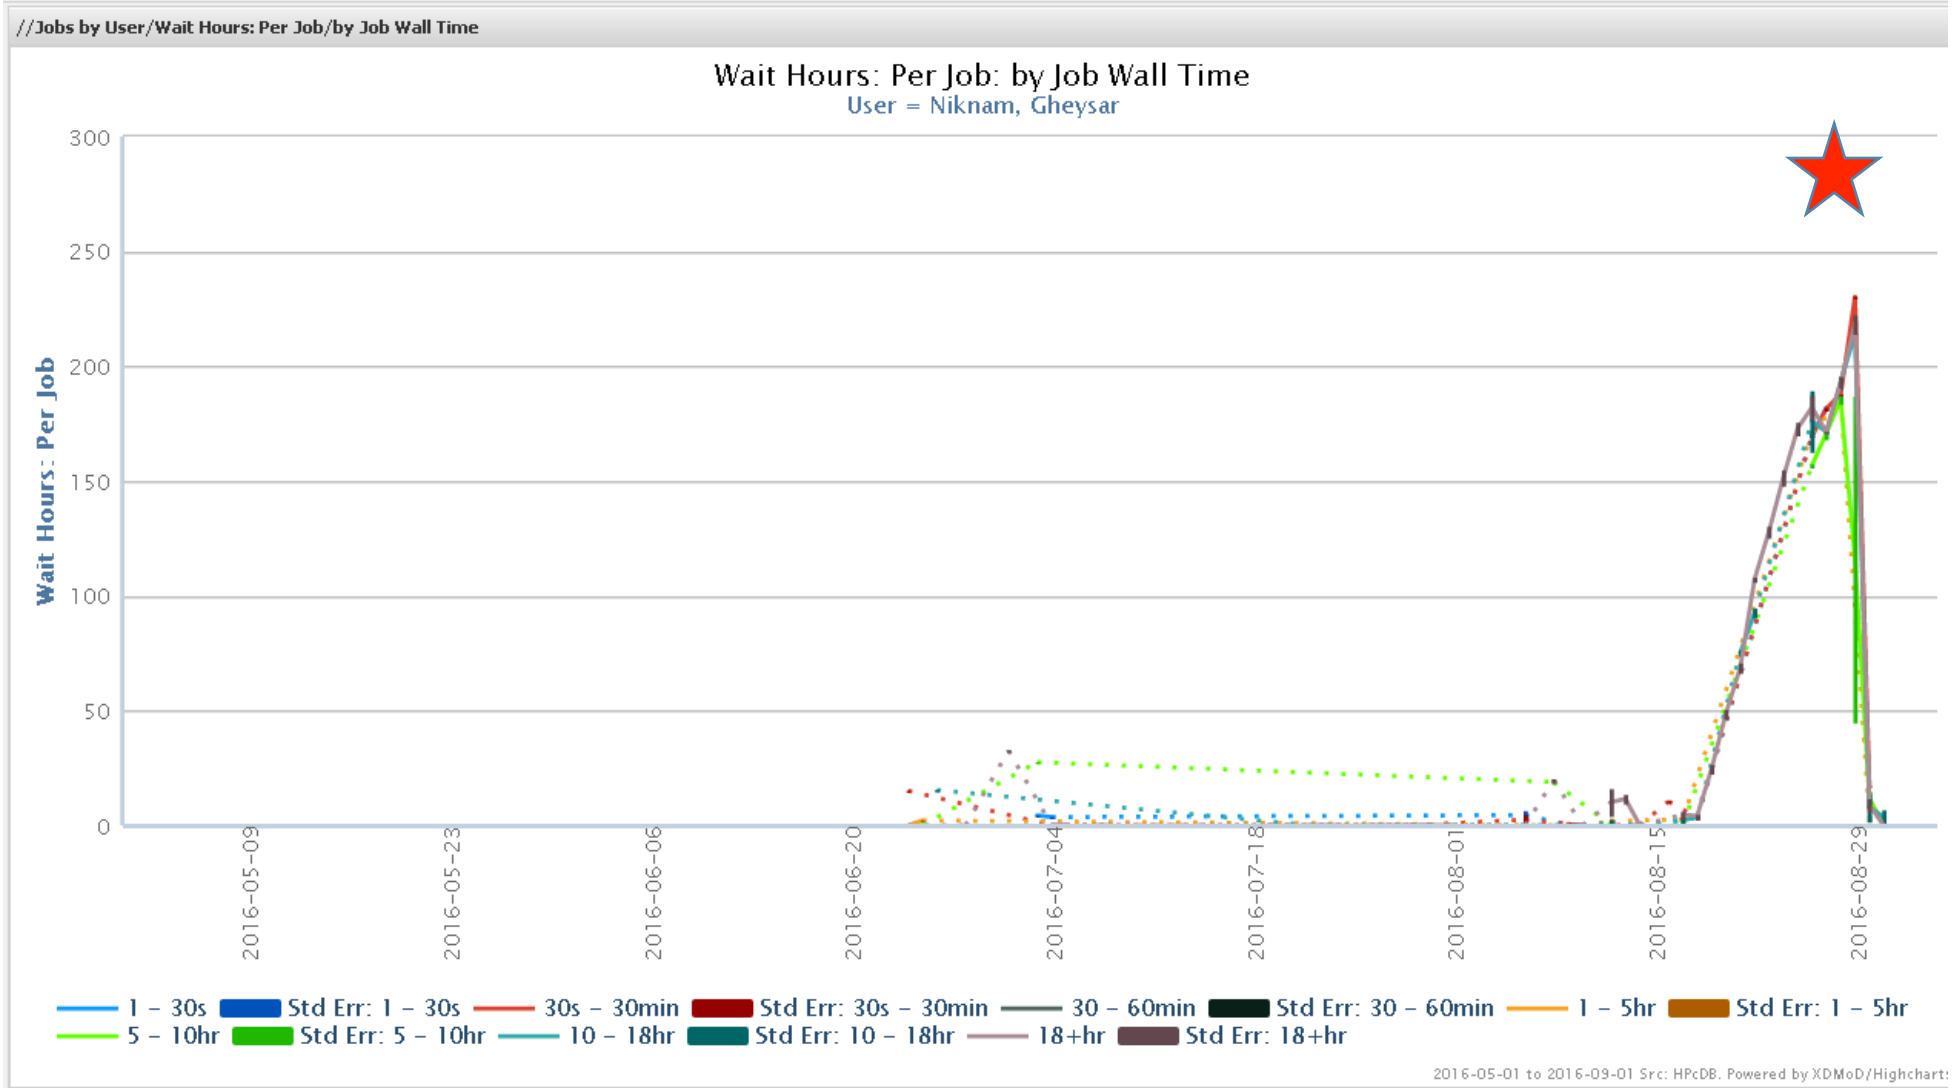

Description

- **Job Wall Time:** A categorization of jobs into discrete groups based on the total linear time each job took to execute.
- **Wait Hours: Per Job**: The average time, in hours, a Information Technology Services Research Cyberinfrastructure job waits before execution on the designated resource.
  Wait Time: Wait time is defined as the linear time between submission of a job by a user until it begins to execute.

MONTANA STATE UNIVERSITY

Waiting time, everyone, August:

//Jobs by Job Wall Time/Wait Hours: Per Job

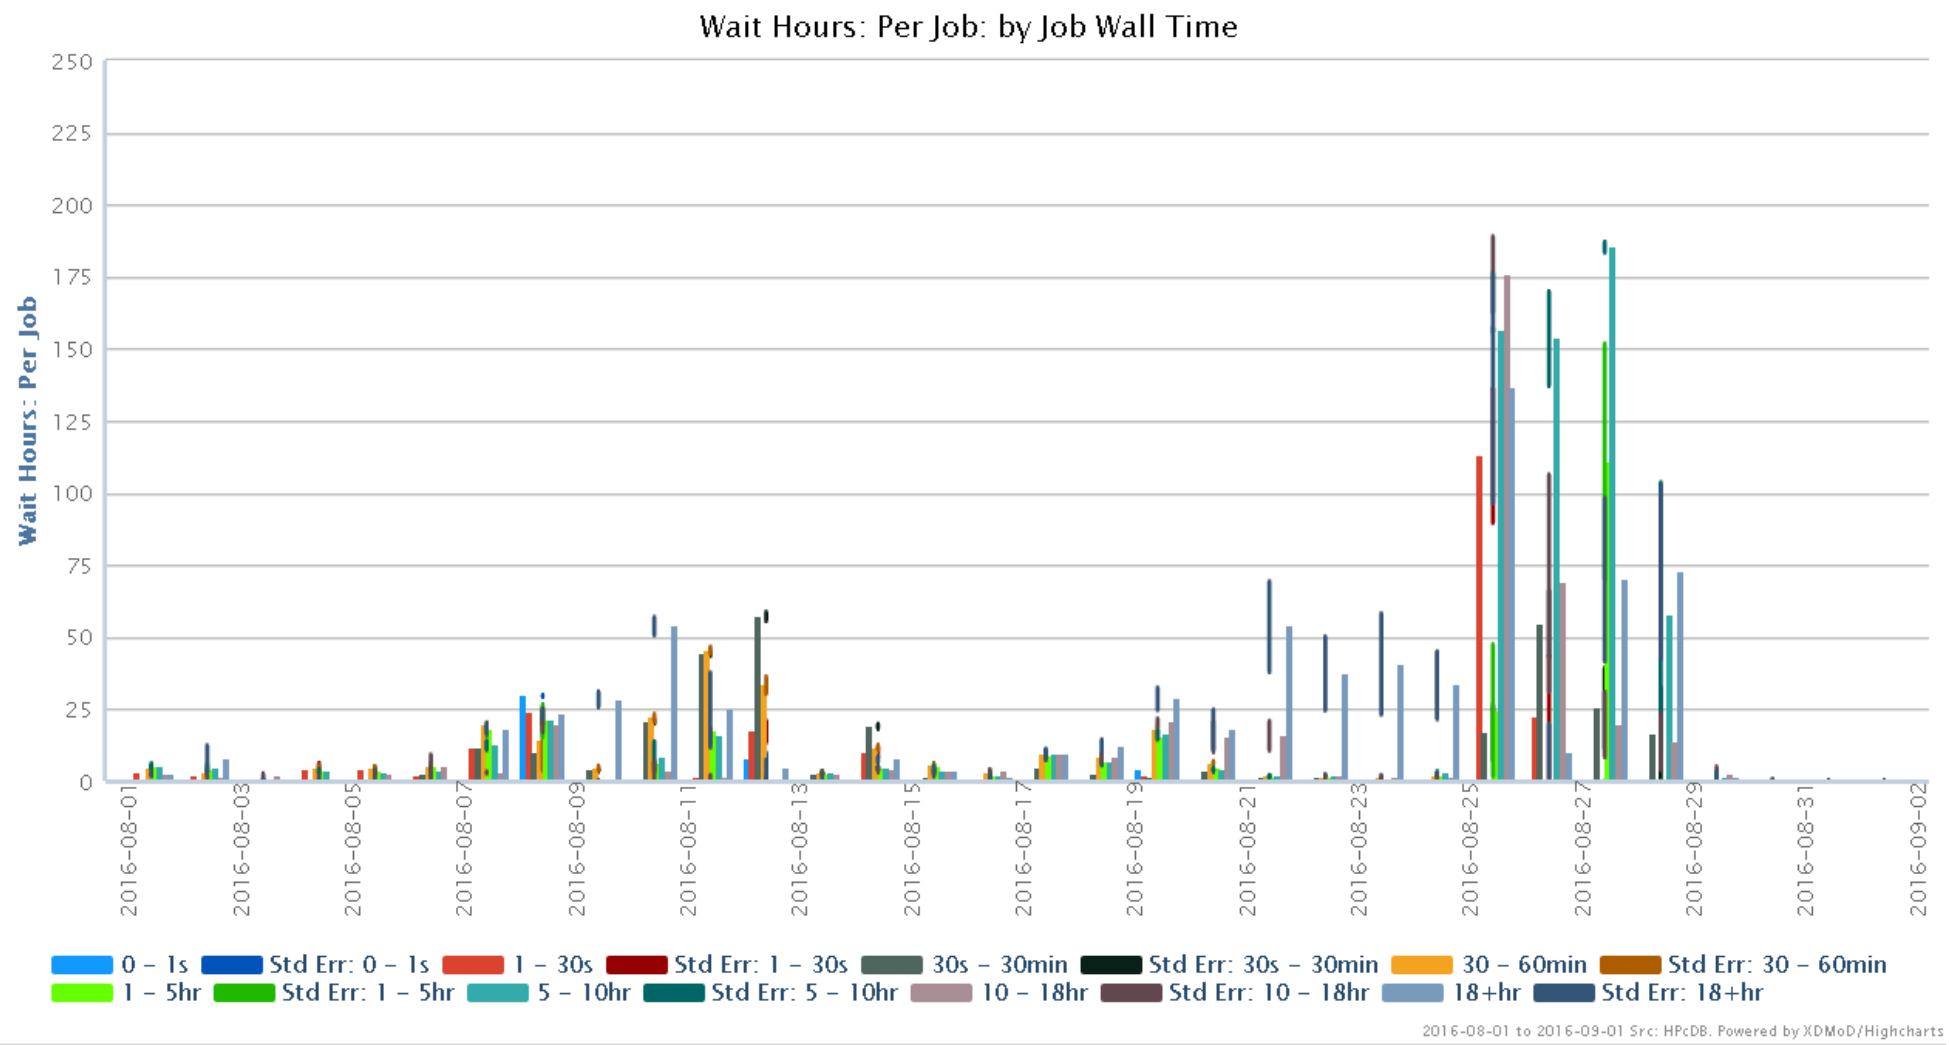

Description

- **Job Wall Time:** A categorization of jobs into discrete groups based on the total linear time each job took to execute.
- **Wait Hours: Per Job**: The average time, in hours, a Information Technology Services Research Cyberinfrastructure job waits before execution on the designated resource.
  Wait Time: Wait time is defined as the linear time between submission of a job by a user until it begins to execute.

MONTANA STATE UNIVERSITY

Waiting hours vs job length:

//Jobs by Job Wall Time/Wait Hours: Per Job

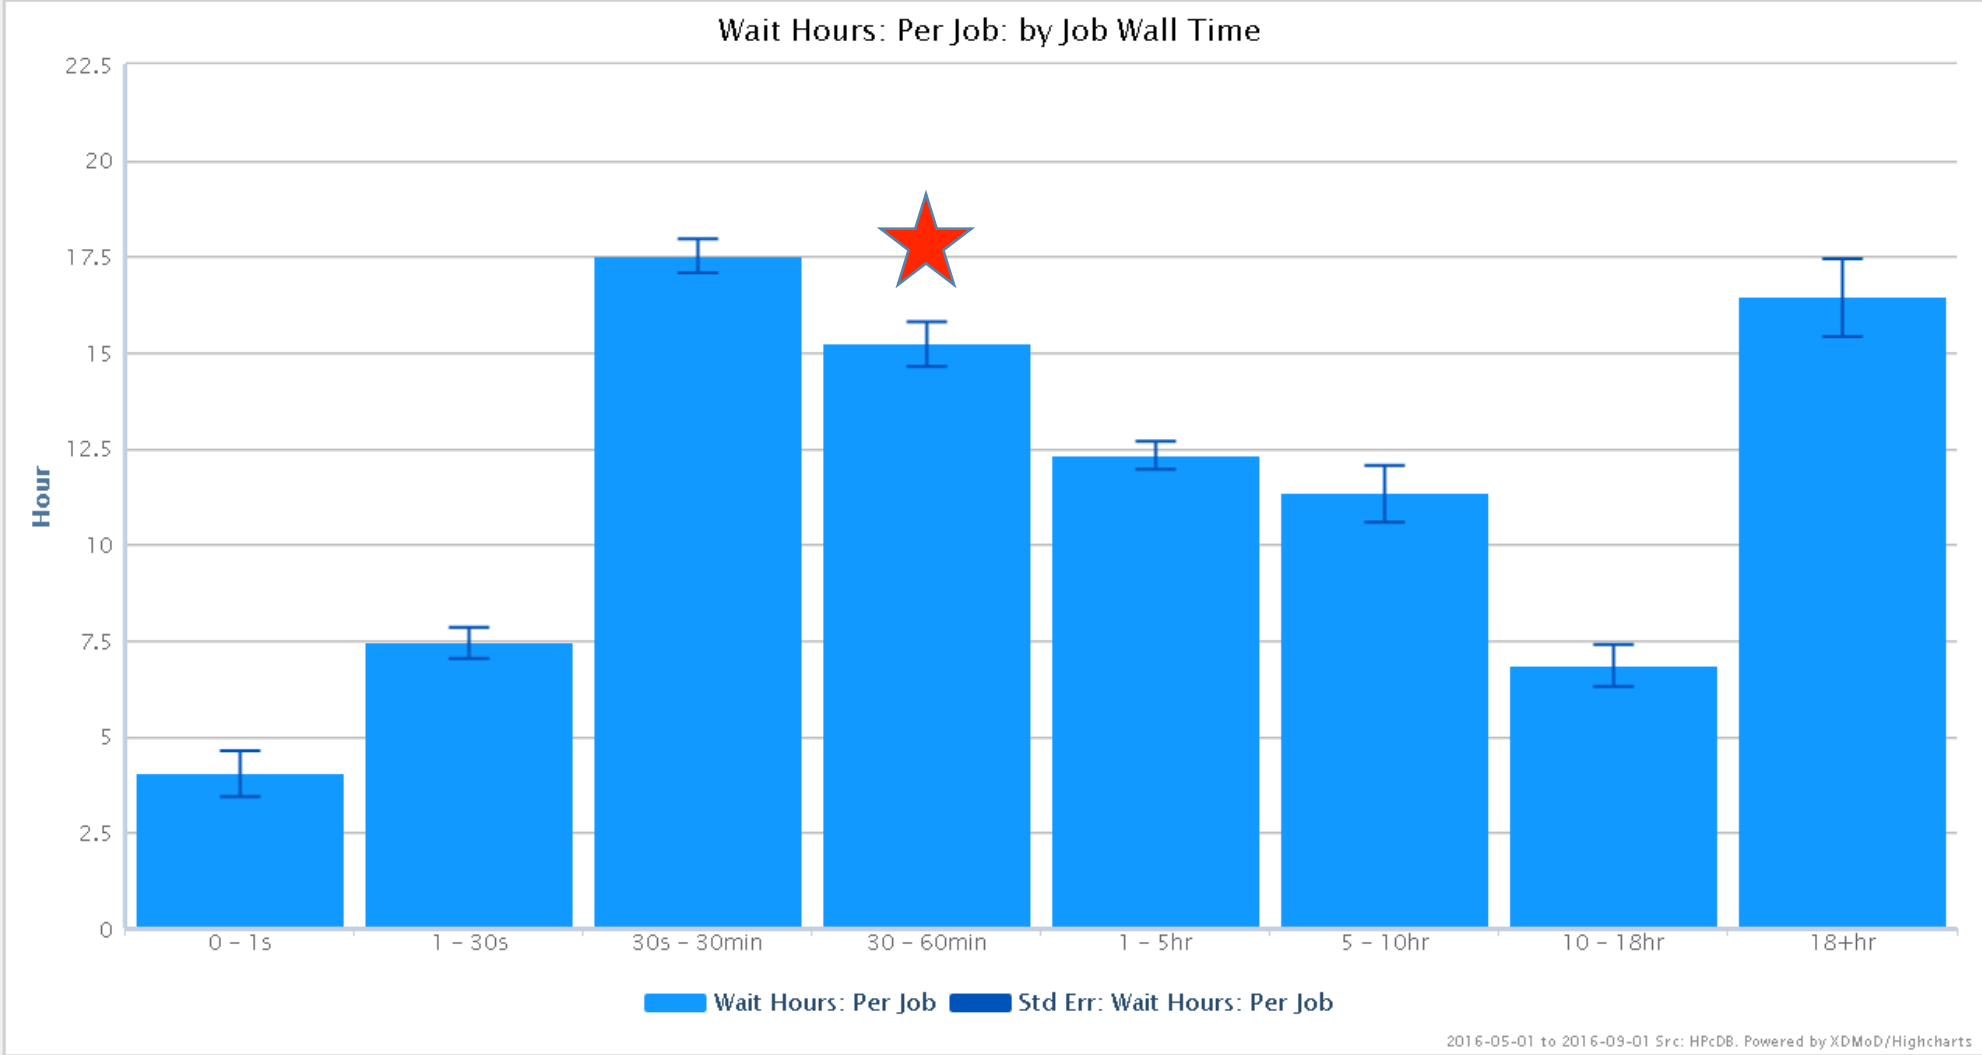

Description

- **Job Wall Time:** A categorization of jobs into discrete groups based on the total linear time each job took to execute.
- **Wait Hours: Per Job**: The average time, in hours, a Information Technology Services Research Cyberinfrastructure job waits before execution on the designated resource.
  Wait Time: Wait time is defined as the linear time between submission of a job by a user until it begins to execute.

MONTANA STATE UNIVERSITY

Why and who wait for brief jobs?

//Jobs by Job Wall Time/Wait Hours: Per Job/by Queue

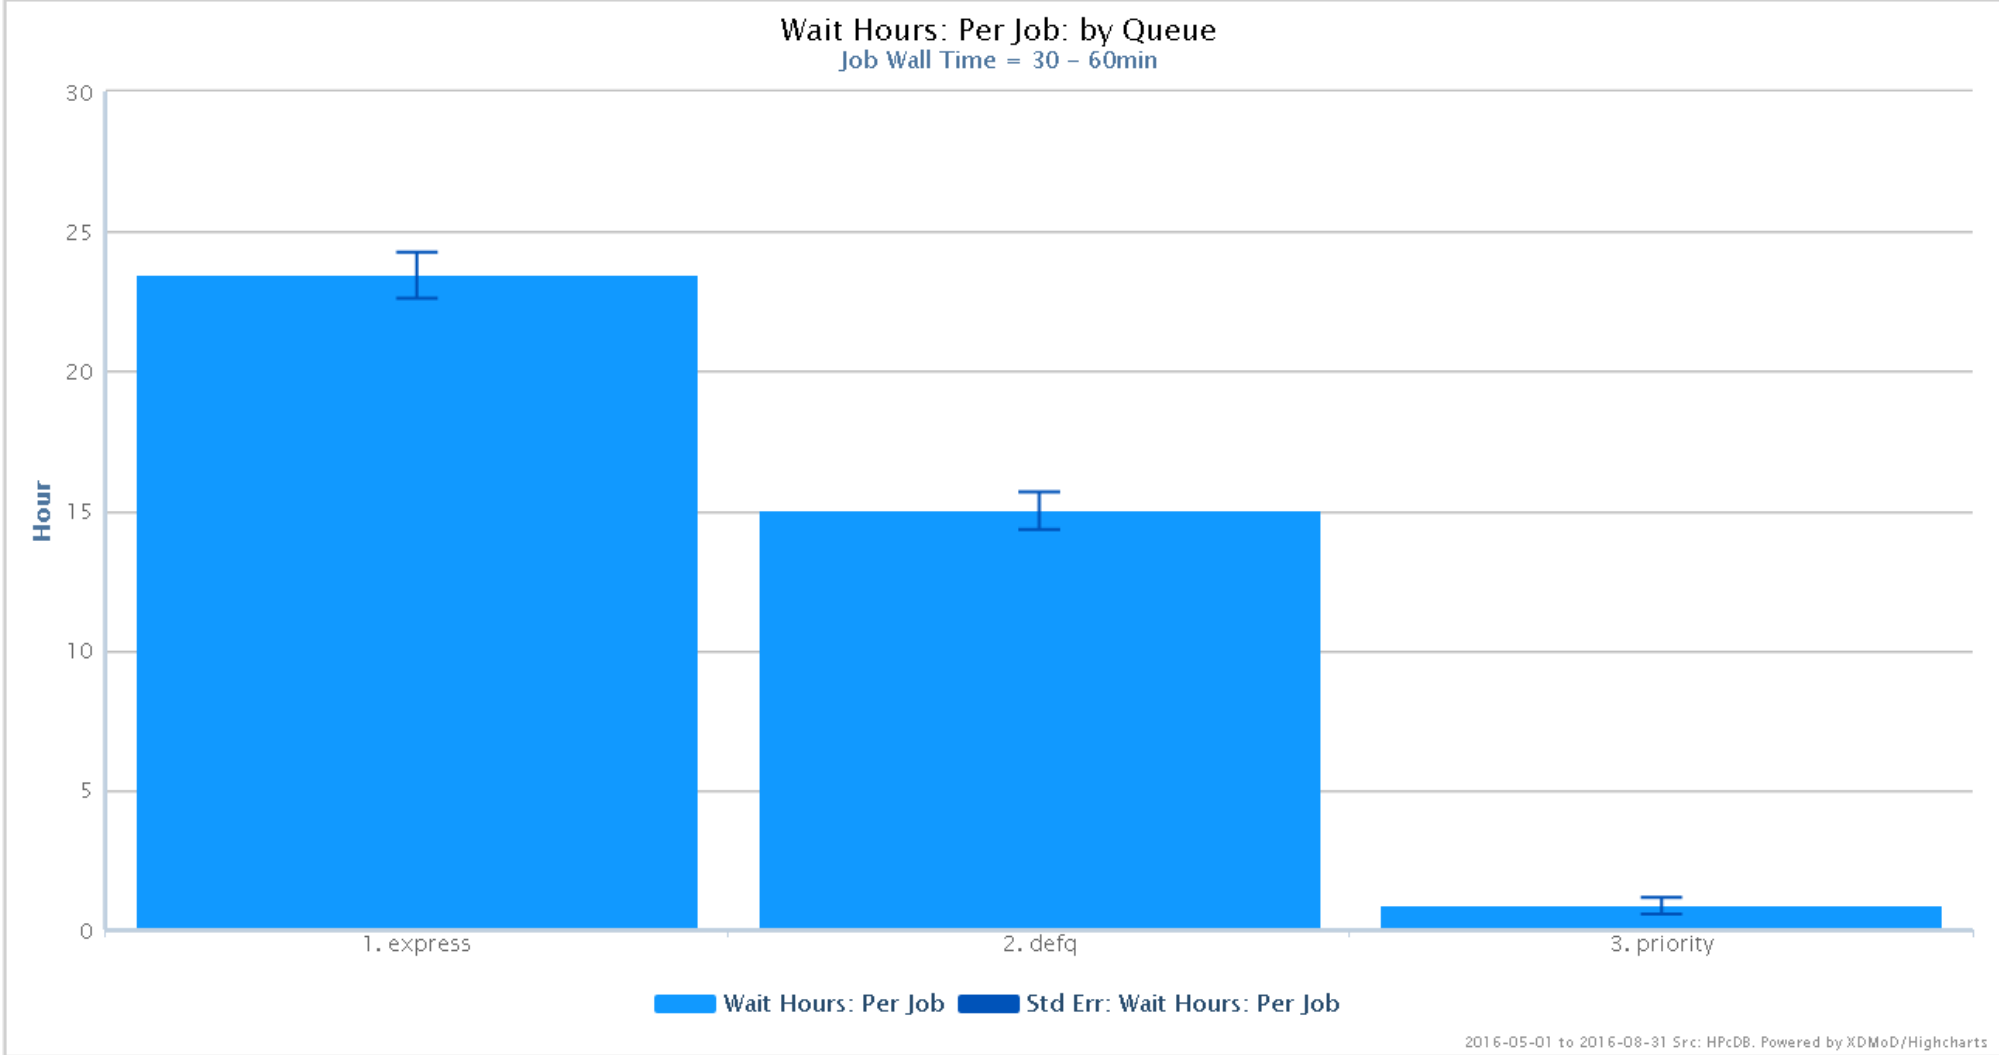

Description

- **Queue:** Queue pertains to the low level job queues on each resource.
- **Wait Hours: Per Job**: The average time, in hours, a Information Technology Services Research Cyberinfrastructure job waits before execution on the designated resource.
  Wait Time: Wait time is defined as the linear time between submission of a job by a user until it begins to execute.

MONTANA STATE UNIVERSITY

Why and who wait for brief jobs?

//Jobs by Job Wall Time/Wait Hours: Per Job/by User

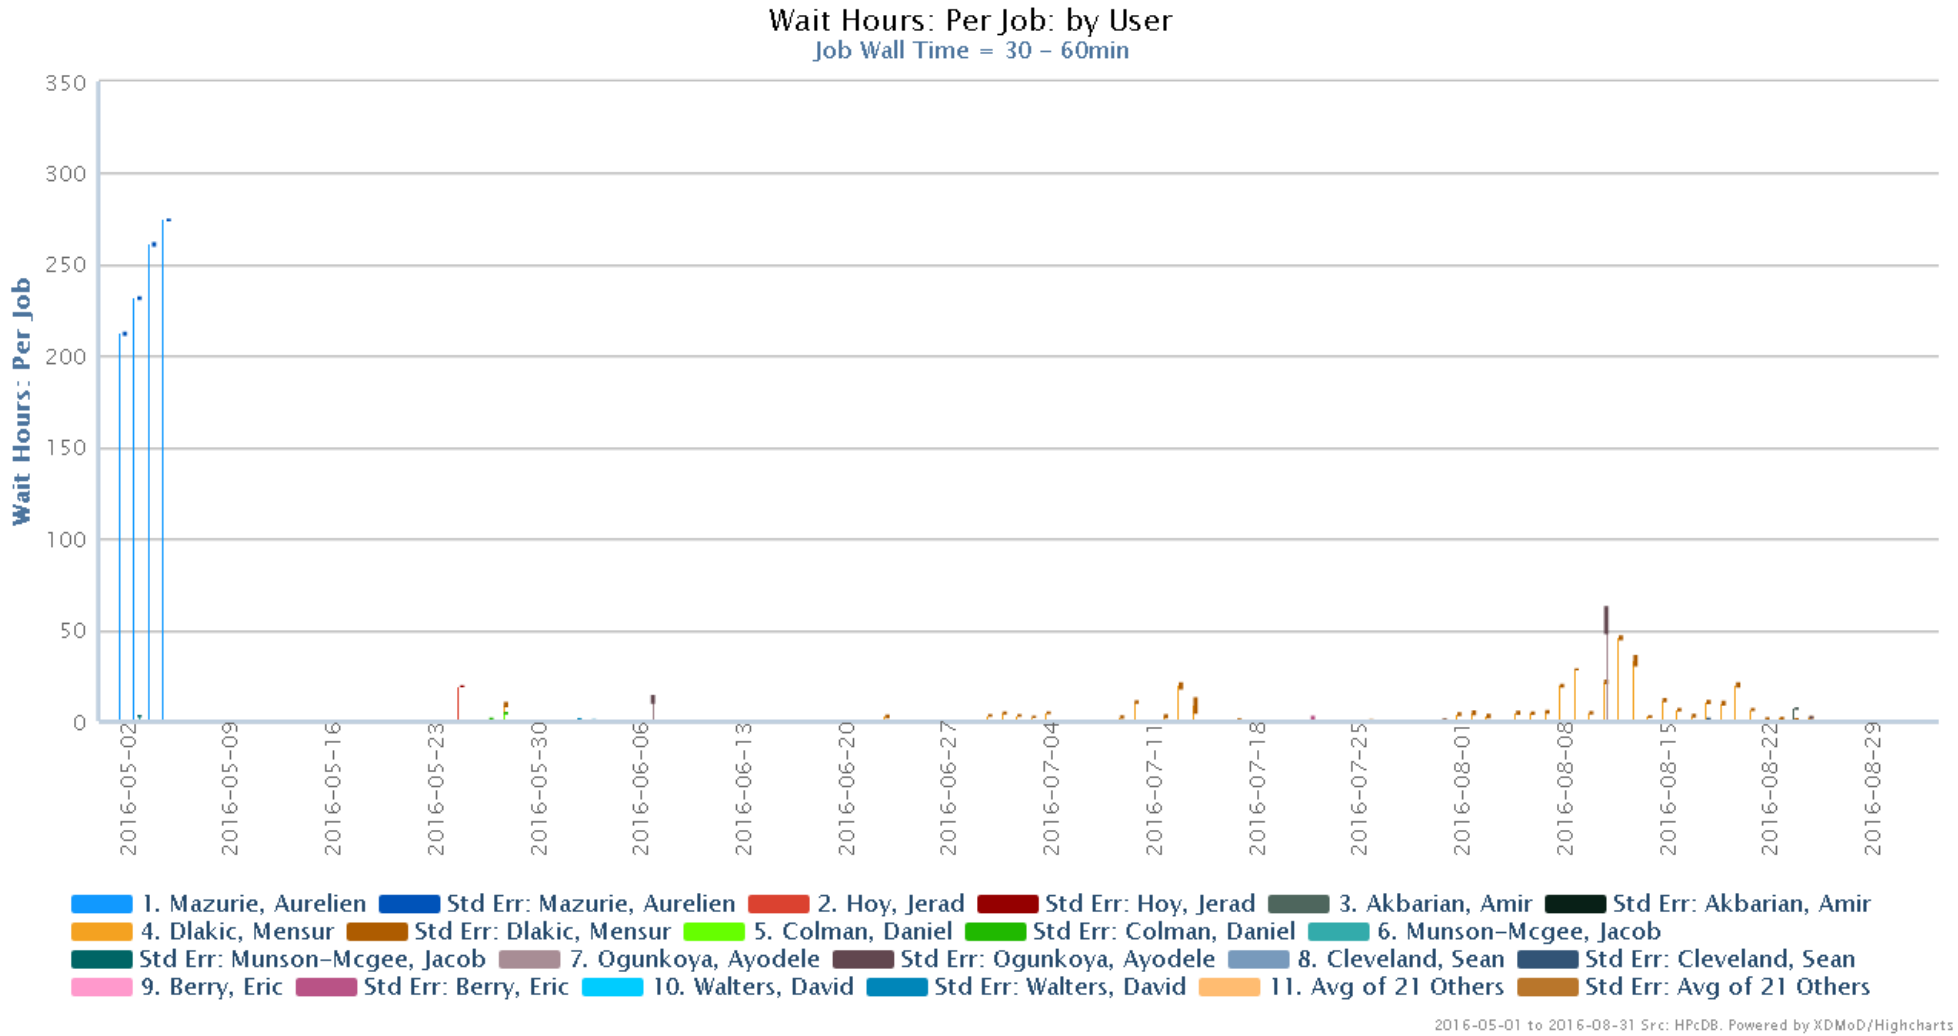

Description

- **User:** A person who is on a PIs allocation, hence able run jobs on resources.
- **Wait Hours: Per Job**: The average time, in hours, a Information Technology Services Research Cyberinfrastructure job waits before execution on the designated resource.
  Wait Time: Wait time is defined as the linear time between submission of a job by a user until it begins to execute.

MONTANA STATE UNIVERSITY

Matlab Update

- Matlab Total Academic Headcount License
  - Faculty, Staff, or Student
  - Any machine (Home or On-Campus)
- ITC Help Page about Installation
- Hyalite
  - Working on installation of R2016a
  - Will test Distributed Compute Server (32 worker)
- Matlab HPC Mentors Monthly Meeting

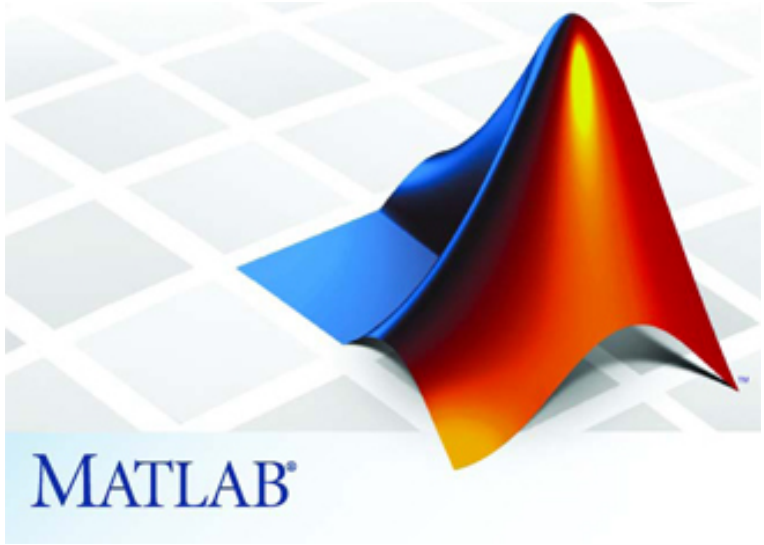

CHMY591 – Computational Chemistry

- Students: cap of 15, currently 4
- Usage plan:
  - Students start with own systems
  - Learn software, shared-user systems
  - Gradually move to cluster
  - By end of semester, submitting very long job
- Software
  - TINKER
  - MOPAC
  - DFTP+
  - Gaussian09
  - Tcl shell
- Estimates (rough Hilmer calculations)
  - Averaged: 25-50% of a single node's capacity, 24/7 for a semester
  - Heavily imbalanced: weighted towards end of semester
  - Very long jobs: up to week each