

# Slurm Version 15.08



Jacob Jenson  
jacob@schedmd.com

SchedMD LLC  
<http://www.schedmd.com>

# V15.08 – Road Map

- Burst buffers
- Expanded charging options
- Improved recovery from communication failures
- Add support for PMI Exascale (PMIx)
- rCUDA






# V15.08 – Burst buffers

- Data storage available for before, during and/or after job computation
- Use cases
  - Data file loaded prior to a job start
  - Persistent data store between jobs
  - Checkpoint
  - Improved performance when every compute node attempts to read the same shared object libraries or config file
  - Temporary job data
  - Post processing Read/Write Files Data Staged
    - Data analysis and visualization

# V15.08 – Charging options



- Historically charging has been based on CPU time.
- Change to a general system billing unit
  - Computed as a function of many different resources
  - CPUs, memory, power and GPUs



# V15.08 – Improved communication failure recovery

- Improve recovery time for communication failures when large numbers of nodes fail simultaneously.


## 14.11 vs 15.08



# V15.08 – PMI Exascale (PMIx)

- Support for PMIx to improve MPI scalability.
- The vision for Slurm is to extend workload management functionality to address Exascale requirements.

12 15 18  
101010  
TERA PETA EXASCALE & BEYOND



# V15.08 – rCUDA



UNIVERSITAT  
POLITÈCNICA  
DE VALÈNCIA


Sergio Iserte

Adrián Castelló

Rafael Mayo

Enrique S. Quintana-Ortí

Federico Silla, Jose Duato



UNIVERSITAT  
JAUME I


**rCUDA**  
*remote CUDA*

# V15.08 – rCUDA






# V15.08 – rCUDA



# V15.08 – rCUDA



# V15.08 – rCUDA

