Slurm Version 15.08

Morris Jette, Danny Auble (SchedMD)
Yiannis Georgiou (Bull)

Slurm User Group 2015

Version 15.08

- Version 15.08.0 released on August 31
- Massive changes from version 14.11
  - Diff file >250,000 lines

Trackable Resources (TRES)

- Tracks utilization of memory, GRES, burst buffer, license, and any other configurable resources in the accounting database
- Any TRES can be used as a factor in computing a job's billing weight as used in calculating its resource utilization
- Any TRES can be used as a factor on calculating a job's priority
- Separate presentation with more details

Per-Partition QOS

- Each partition can now have an associated QOS
- Each partition now has all of the limits available in a QOS
- Separate presentation with more details

Burst Buffers

- Burst buffers are a cluster-wide high-performance file system
- Slurm can allocate resources, stage-in files before a job starts, stage-out files after a job completes and otherwise manage burst buffer resources
- Separate presentation with more details

Network Topologies

- Optimized resource allocations for
  - SGI Hypercube (work by SGI)
  - Dragonfly

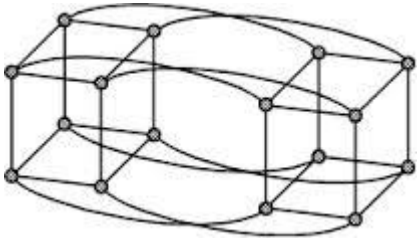

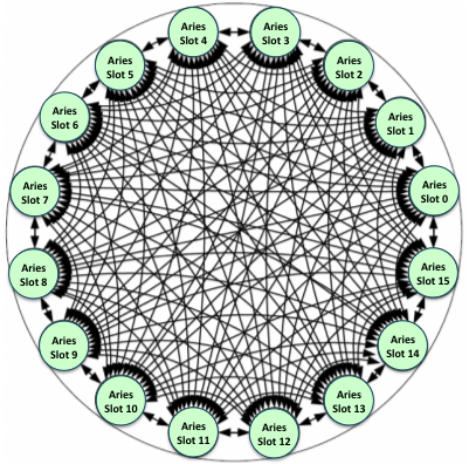

Advanced Reservations

- New flag “replace”
  - As resources are allocated, replace them with idle resources to the extent possible
  - Maintains constant size of available resources
- Replace flag "License_only" with flag "Any_Nodes".
  - Used to indicate the advanced reservation resources (licenses and/or burst buffers) can be used with any compute nodes

Job Preemption

- Permit “PreemptMode=suspend,gang” and “PreemptType=qos” and to be used together
  - A high-priority QOS job will now oversubscribe resources and gang schedule, but only if there are insufficient resources for the job to be started without preemption

Single User Per Node

- Compute nodes can be allocated to multiple jobs, but restricted to a single user
  - New job option “--exclusive=user”
  - New partition configuration parameter “ExclusiveUser=yes”
  - Only accounts for resources allocated to jobs, idle resources currently not changed to the user who has been allocated the compute node

Elasticsearch Job Records

- New job completion plugin records a job's details into Elasticsearch database
- Many Elasticsearch tools available for analysis
- (Work by Alejandro Sanchez Graells, Barcelona Supercomputing Center, Universitat Politecnica de Catalunya)

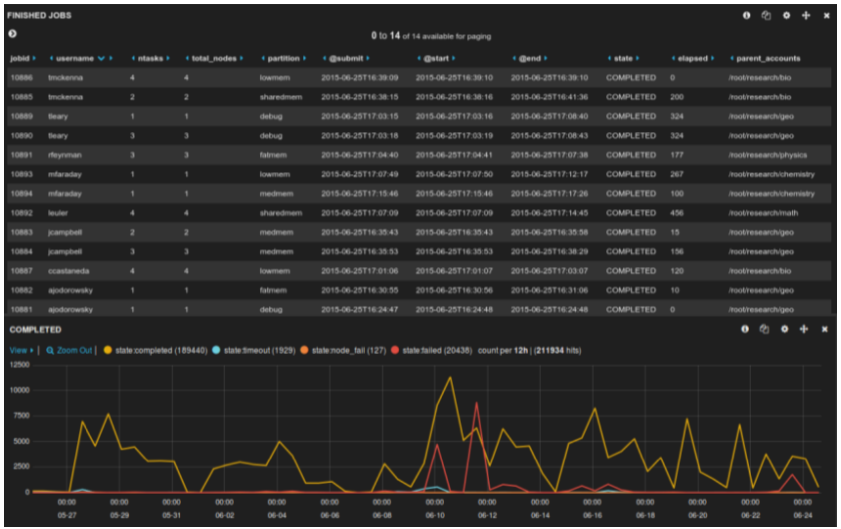

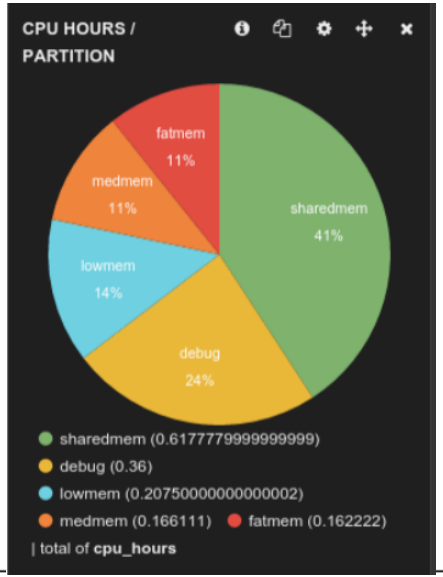

New Job Options

- sbast command can operate on nodes associated with either a job step (new) or an entire job
- Job “--mail” options now apply to a job array as a whole rather than each task of the job array
- OR'ed job dependencies
  - For example: “--depend=afterok:123?afternotok:124”
- Job “--kill-on-invalid-dep” to control behavior if the job's dependency can never be satisfied

New Environment Variables

- New job array environment variables
  - For example: --array=1-99:2
    - SLURM_ARRAY_TASK_MIN=1
    - SLURM_ARRAY_TASK_MAX=99
    - SLURM_ARRAY_TASK_STEP=2
- Other new environment variables
  - SLURM_JOB_ACCOUNT
  - SLURM_JOB_QOS
  - SLURM_JOB_RESERVATION

Task Layout Options (1 of 3)

- Job “--dist” option has “pack” and “nopack” flags
  - “pack” fills nodes, resulting in uneven distribution
  - “nopack” distributes tasks evenly across nodes (default)

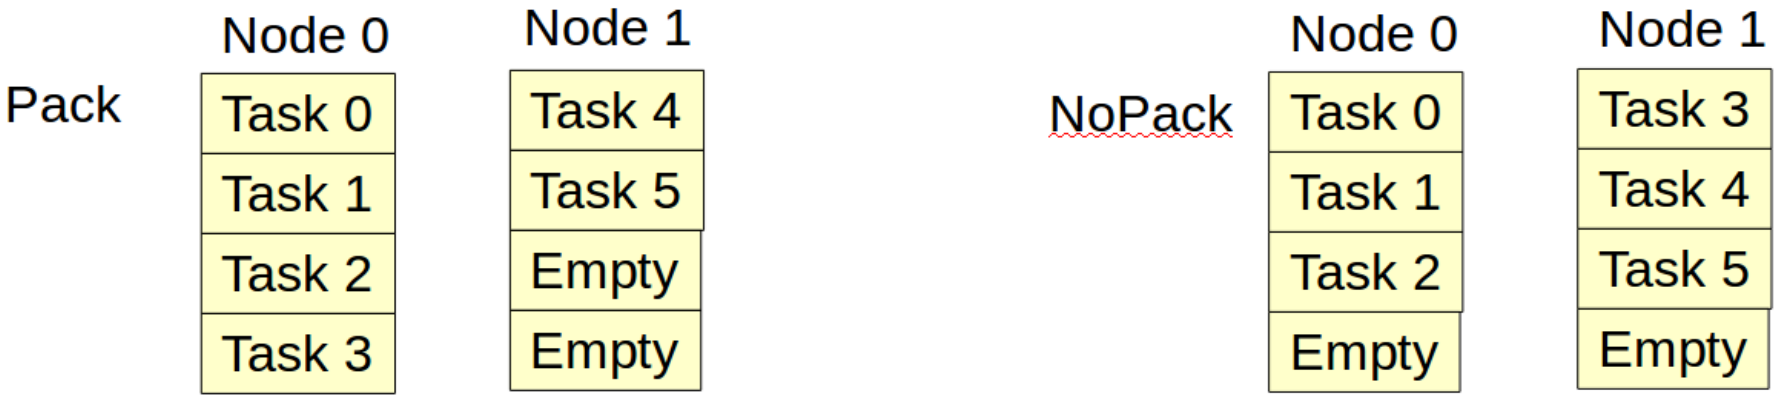

Task Layout Options (2 of 3)

- Job “--thread-spec” option can reserve a hyper-thread for system use rather than an entire core (the “--core-spec” option)
- Job “--accel-bind” option binds tasks to nearest GPU and NIC

Task Layout Options (3 of 3)

- To better support the Xeon Phi processor architecture and OpenMP
  - New parameter to control the distribution of allocated threads across cores for binding to tasks
- Allows greater control over the placement of jobs
- The new syntax for the --distribution option is as follows. The new parameters are shown in bold

```
-m, --distribution=*|block|cyclic|arbitrary|plane=<options> [:*|block|
cyclic|fcyclic [:*|block|cyclic|fcyclic]][,Pack|NoPack]
```

Accounting – Profiling
Support of multiple energy sensors

- Support for one sensor per node (until 14.11)

```
$ ipmi-sensors
62 | Power          | Current     | 175.80    | W          | 'OK'
```

- Support for multiple sensors per node (from 15.08)

```
$ipmi-sensors
85 | CPU0 Pwr          | Power Supply        | 10.00    | W    | 'OK'
86 | CPU1 Pwr          | Power Supply        | 6.00     | W    | 'OK'
87 | CPU0 DIM01 Pwr    | Power Supply        | 2.00     | W    | 'OK'
88 | CPU0 DIM23 Pwr    | Power Supply        | 0.00     | W    | 'OK'
89 | CPU1 DIM01 Pwr    | Power Supply        | 1.00     | W    | 'OK'
90 | CPU1 DIM23 Pwr    | Power Supply        | 0.00     | W    | 'OK'
91 | Blade Pwr         | Power Supply        | 112.00   | W    | 'OK'
```

Accounting – Profiling
Support of multiple energy sensors

- Configuration slurm.conf (No Change)

```
$cat slurm.conf
….
AcctGatherEnergyType=acct_gather_energy/ipmi
```

- Configuration acct_gather.conf (Changed - Attention)
  - Node required, other options custom and optional

```
$cat slurm.conf
….
#Old version :
EnergyIPMIPowerSensor=62

#New version (details for hdf5) :
EnergyIPMIPowerSensors=Node=91;CPU0=85,87,88;CPU1=86,89,90
```

Accounting – Profiling
Support of multiple energy sensors, Usage

```
$ scontrol show n=mo38 | grep ConsumedJoules
   CurrentWatts=105 LowestJoules=105 ConsumedJoules=17877

$ srun –acctg-freq=energy=3 –profile=energy -n ..... ./app

$ sstat – j 15 -o JobID,ConsumedEnergy,ConsumedEnergyRaw
     JobID ConsumedEnergy ConsumedEnergyRaw
------------ -------------- -----------------
15.0                2.10K       2103.000000

$ sacct – j 15 -o JobID,ConsumedEnergy,ConsumedEnergyRaw
       JobID ConsumedEnergy ConsumedEnergyRaw
------------ -------------- -----------------
15                  9.26K       9257.000000

$ sh5util -j 15
# And any application using hdf5.
```

Accounting – Profiling
Support of multiple energy sensors, Usage

```
[root@nd25 slurm]# cat energy_38.csv

ElapsedTime; EpochTime; NodePower; socket0Power; socket1Power
0; 1441786898; 75; 14; 18
1; 1441786899; 126; 82; 54
2; 1441786900; 171; 84; 54
3; 1441786901; 186; 83; 53
4; 1441786902; 189; 88; 55
```

- Accounting: Database aggregated Energy Consumption considering all sensors
- Profiling: HDF5 files separated sensors power values

Profiling
HDF5 Architecture Optimizations

- New more scalable and flexible architecture
  - AcctGatherProfile operates as a service
- Based upon the high level HDF5 API
  - Features such as data compression (TODO add new parameter)
- Update sh5util (kept backward compatibility)
  - Calculate statistics during merge and not during processing

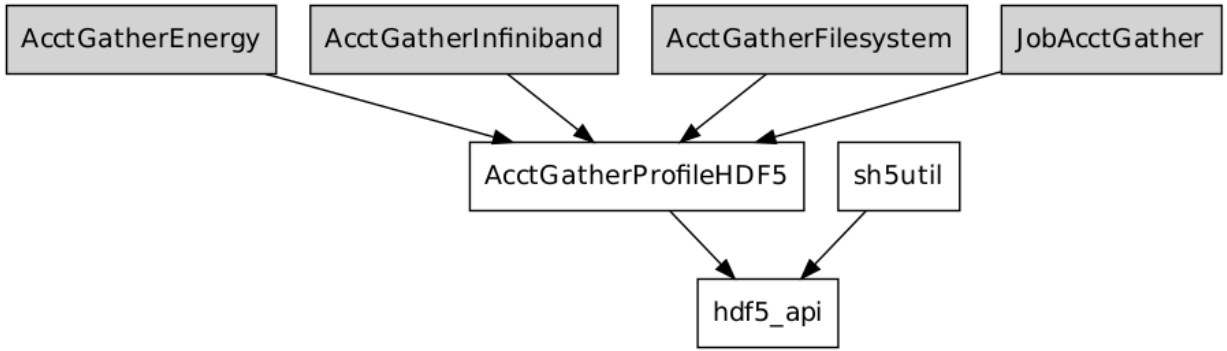

Profiling
HDF5 Architecture Optimizations

- Results: Profiling of a medium instance of HPLinpack upon 3 nodes (24min)
- Size of the profiling files:

Size (MiB)	Old	New
Each node	6,58	0,21
Consolidated	0,79	0,62

- Time to merge per node:

Time (sec)	Old	New
sh5util real time	1,36	0,083
sh5util user time	0,77	0,005

Support of Layouts Framework

- A generic representation of the arrangement of systems components associated to a particular aspect.
- Integrated within SLURM (15.08 - technical preview)
  - Attributes Inheritance features
  - API functions
  - Support of “scontrol update/show “
- Separate presentation with more details

Power Adaptive Scheduling

- Provide **centralized mechanism** to dynamically **adapt the instantaneous power consumption** of the whole platform
  - Reducing the number of usable resources or running them with lower power
- Based upon the layouts framework
- Separate presentations with more details

Power Management

- Job “--cpu-freq” option now supports minimum frequency (in addition to maximum frequency and governor) and supported for salloc and sbatch (for power adaptive scheduling)
- --cpu-freq =<p1[-p2[:p3]]>
  - p1 is current options or minimum frequency
  - optional p2 is maximum
  - optional p3 is scaling governor
- New configuration parameter “CpuFreqGovernors” identifies allowed governors

Message Aggregation

- Improve performance by aggregating epilog complete and node registration messages into a smaller number of composite messages,
- To reduce the number of incoming TCP connections to serve.
- To decrease the processing time of messages
- Based upon the reverse of the message forwarding/fanout mechanism
- Separate presentation with more details

WARNINGS (1 of 2)

- **The database schema has changed. Updating slurmdbd will take time. No records will be lost while upgrading, but the slurmdbd may not be responsive. It will not be possible to automatically revert the database for an earlier version of Slurm**
- In preparation for inter-cluster jobs, the MaxJobID has been reduced from 4,294,901,760 to 2,147,463,647. **Any job with an ID above 2,147,463,647 will be purged when upgrading**
- MVAPICH plugin now requires Munge for authentication
- Every plugin except SPANK must be built against the same version of Slurm (major and minor version number) to be loaded

WARNINGS (2 of 2)

- **HDF4 node-step file format changed** for improved performance
  - sh5util command has changed
  - Both file formats supported for next few releases

Questions?