

Partition QOS



Danny Auble
SchedMD LLC

Slurm User Group Meeting 2015

Outline

- Reasoning and Previous Limitations
- How To
- Benefits

Reasoning and Previous Limitations




- Want for limits of a QOS on a Partition Level
- ≤ 14.11 “floating partition” not really adequate
- This QOS would be the QOS the job used

How To

- Any QOS can be used as a Partition QOS
- Pick QOS and add QOS=\$name to the PartitionName line
- Restart or reconfigure your slurmctld
- The Partition QOS will override the job's QOS.
 - If the opposite is desired you need to have the job's QOS have the 'OverPartQOS' flag which will reverse the order of precedence.


Order of Precedence

The Partition QOS will override the job's QOS.



Order of Precedence

Use 'OverPartQOS' flag on the Job QOS for it to override the Partition QOS.



Floating Partition Example



- Define QOS with limit GrpCPU (Not GrpNodes) set appropriately
- Define partition will **ALL** nodes in the system with the Partition QOS set to the new QOS
- This partition will only have access to a group of cpus now on any of the nodes
- Helpful for debug like partitions with short run times
- Makes it so you don't have nodes sitting idle when other jobs could run if they weren't carved out in a different partition

Partition vs Job QOS



- Every Job will have a QOS if enforcing them
- Not every Partition has to have a QOS
- Non-applicable PartitionQOS features
 - GraceTime
 - UsageFactor
 - UsageThreshold
 - Preemption
 - All flags except 'DenyOnLimit'