

Troubleshooting



Albert Gil
Jason Booth


SchedMD
SLUG 2019

Copyright 2019 SchedMD
www.schedmd.com

Troubleshooting



The ability to fix a problem is as good as the tools and knowledge to fix it.



Copyright 2019 SchedMD
www.schedmd.com

Outline



- How can I get all the information?
 - Commands and Logs
- How to be sure that the ball/problem is actually in Slurm's court?
 - Setup checklist and best practices
- Why won't my job run?
 - Backfill
- How to get (the best) support?
 - Bugzilla

Commands

- Basic Info

- sinfo / squeue
- sprio / sshare
- smap

-o,--format → %letter
VS
-O,--Format → word

- Config

- scontrol show config
- sacctmgr show config

Check running config,
not only config files.
Implicit Configs.

- Details

- scontrol show [partition | reservation...]
- sacctmgr show [assoc | qos | event...]

- Stats & Diagnostic

- sdiag
- sacctmgr show stats
- sacctmgr show problem

- Automation & Notifications

- strigger
- HealthCheckProgram
- UnkillableStepProgram

- Autocorrection

- sacctmgr show RunawayJobs

- External Process

- scontrol [listpids <job> | pidinfo <pid>]
- taskset, stress...



Details - sacctmgr show [assoc | qos | ...]

- Information about Limits or QoS of Users and Accounts (ie, of Associations)

```
$ sacctmgr show assoc tree Format=Account,User,Share,MaxTRES,QoS
```

| Account | User | Share | MaxTRES | QoS |
|-------------|------|-------|---------|-------------|
| root | | 1 | | normal |
| root | root | 1 | | normal |
| development | | 1 | | normal,high |
| development | agil | 70 | | normal,high |
| development | bob | 70 | | normal,high |
| development | joe | 70 | | normal,high |
| development | sue | 70 | | normal,high |
| external | | 1 | | normal,low |
| external | jim | 30 | | normal,low |
| external | joe | 30 | | normal,low |

```
$ sacctmgr show qos Format=Name,Priority,UsageFactor,MaxTRES,MaxTRESPU%20
```

| Name | Priority | UsageFactor | MaxTRES | MaxTRESPU |
|--------|----------|-------------|----------------------|-----------|
| normal | 5 | 1.000000 | cpu=80,gres/gpu=10 | |
| high | 10 | 1.000000 | cpu=800,gres/gpu=100 | |
| low | 1 | 1.000000 | cpu=8,gres/gpu=1 | |

Think more in
Associations and Hierarchies
and less in
Users and Accounts

Double check with
`scontrol show assoc_mgr`

Do you really need *withassoc*?

Copyright 2019 SchedMD
www.schedmd.com



Stats & Diagnostic - sdiag

- How many
 - Jobs being examined
 - Jobs being started by the backfill and quick scheduler
 - RPC calls are being executed
- How long take to process
 - the scheduling operations
 - RPC calls
- Who is executing RPC calls



```
$ sdiag
Server thread count: 3
Agent queue size: 0
Agent count: 0
DBD Agent queue size: 0
```

```
Jobs submitted: 523
Jobs started: 523
Jobs completed: 501
Jobs canceled: 3
Jobs failed: 19
```

```
Main schedule statistics (microseconds):
Last cycle: 16
Max cycle: 53
Total cycles: 59
Mean cycle: 20
Mean depth cycle: 11
Cycles per minute: 1
Last queue length: 0
```

Check size of
queues and counts

```
Backfilling stats
Total backfilled jobs (since last slurm start): 0
Total backfilled jobs (since last stats cycle start): 0
Total cycles: 28
Last cycle when: Wed Dec 30 15:33:18 2018
Last cycle: 93
Max cycle: 5433
Last depth cycle: 0
Last depth cycle (try sched): 0
Last queue length: 0

Latency for 1000 calls to gettimeofday(): 15 microseconds

Remote Procedure Call statistics by message type
REQUEST_RESOURCE_ALLOCATION (4001) count:5 ave_time:1880
total_time:94042
REQUEST_JOB_READY (4019) count:5 ave_time:490
total_time:24520

Remote Procedure Call statistics by user
student (1002) count:32 ave_time:1405
total_time:44973
root ( 0) count:0 ave_time:0 total_time:0

Pending RPC statistics
No pending RPCs
```

Stats & Diagnostic - sacctmgr show stats

- Provides detailed information about:
 - Rollups
 - How many Slurm DB RPC calls are being executed and by whom



```
$ sacctmgr show stats
Rollup statistics
  Hour      count:746   ave_time:307765 max_time:223576198   total_time:229593291
  Day       count:31    ave_time:2429   max_time:10972      total_time:75328
  Month     count:1     ave_time:32007  max_time:32007      total_time:32007

Remote Procedure Call statistics by message type
  DBD_CLUSTER_TRES      ( 1407) count:8948   ave_time:57837   total_time:517531274
  DBD_JOB_COMPLETE     ( 1424) count:5       ave_time:17972   total_time:89864
  DBD_FINI              ( 1401) count:5       ave_time:256     total_time:1284
  SLURM_PERSIST_INIT    ( 6500) count:4       ave_time:341     total_time:1367
  DBD_STEP_START       ( 1442) count:3       ave_time:4617   total_time:13852
  DBD_SEND_MULT_MSG    ( 1474) count:3       ave_time:1579   total_time:4738
  DBD_STEP_COMPLETE    ( 1441) count:3       ave_time:1252   total_time:3757
  DBD_SEND_MULT_JOB_START ( 1472) count:3       ave_time:3527   total_time:10581
  DBD_JOB_START        ( 1425) count:2       ave_time:1146   total_time:2292
  DBD_NODE_STATE       ( 1432) count:2       ave_time:2427   total_time:4854
  DBD_GET_USERS        ( 1415) count:1       ave_time:510    total_time:510
  DBD_GET ASSOCS       ( 1410) count:1       ave_time:1768   total_time:1768
  DBD_GET_RES          ( 1478) count:1       ave_time:274    total_time:274
  DBD_REGISTER_CTLD    ( 1434) count:1       ave_time:1065   total_time:1065
  DBD_GET_TRES         ( 1486) count:1       ave_time:357    total_time:357
  DBD_GET FEDERATIONS  ( 1494) count:1       ave_time:477    total_time:477
  DBD_GET_QOS          ( 1448) count:1       ave_time:208    total_time:208
  DBD_GET_STATS        ( 1489) count:1       ave_time:387    total_time:387
  DBD_GET_CONFIG       ( 1466) count:1       ave_time:99     total_time:99

Remote Procedure Call statistics by user
  agil      ( 1000) count:8987   ave_time:57601   total_time:517669008
```

Stats & Diagnostic - sacctmgr show problem

```
$ sacctmgr show problem
Cluster  Account  User                Problem
-----
                jim         User does not have a uid
                joe         User does not have a uid
```

Quick Poll?



Automation - strigger

- strigger manage event triggers, programs to be run on a variety of events
- Mostly used to notify sysadmins about failures, and quick containment

```
> strigger --set --primary_slurmctld_failure \  
    --program=/usr/sbin/primary_slurmctld_failure.bash  
  
> cat /usr/sbin/primary_slurmctld_failure.bash  
#!/bin/bash  
# Resubmit trigger  
strigger --set --primary_slurmctld_failure \  
    --program=/usr/sbin/primary_slurmctld_failure.bash  
# Send notification  
/bin/mail slurm_admin@site.com -s Primary_slurmctld_failure
```



Automation - strigger events

```
primary_slurmctld_failure
primary_slurmctld_resumed_control
primary_slurmctld_acct_buffer_full
backup_slurmctld_failure
backup_slurmctld_assumed_control
```

```
primary_slurmdbd_failure
primary_slurmdbd_resumed_operation
primary_database_failure
primary_database_resumed_operation
```

```
Down      (compute node enters DOWN state)
Drained   (compute node enters DRAINED state)
Fail      (compute node enters FAILING state)
```

Not a complete list.
See manpage of strigger for
complete listing



Underused?
Quick Poll?

Automation - trigger events (for users)

- This mechanism can also be used to notify user of job-specific events
 - Node failure in a job's allocation
 - Job nearing time limit

```
> strigger --set --time --offset=600 --jobid=123 \  
    --program=/home/student/time_msg.bash  
  
> cat /home/student/time_msg.bash  
#!/bin/bash  
# Send notification when job within 600 seconds of time limit  
/bin/mail student@site.com -s "Check job time limit"
```



Automation - HealthChecks



Because a healthy node is more than just a healthy slurmd:

- `HealthCheckProgram`
 - Fully qualified pathname of a script to execute as user root periodically on all compute nodes that are not in the `NOT_RESPONDING` state.
- `HealthCheckInterval`
 - The interval in seconds between executions of `HealthCheckProgram`.
- `HealthCheckNodeState`
 - Identify what node states should execute the `HealthCheckProgram`.

Integration with [LBNL Node Health Check](#)

Automation - Unkillable

Well, Slurm also needs to be able to kill jobs/processes:

- `UnkillableStepProgram`
 - If the processes in a job step are determined to be unkillable for a period of time specified by the `UnkillableStepTimeout` variable, the program specified by `UnkillableStepProgram` will be executed.
- `UnkillableStepTimeout`
 - The length of time, in seconds, that Slurm will wait before deciding that processes in a job step are unkillable (after they have been signaled with SIGKILL) and execute `UnkillableStepProgram`.

Last quick poll?

Autocorrection

```
$ sacctmgr show runaway
```

```
NOTE: Runaway jobs are jobs that don't exist in the controller but have a start time and no end time in the database
```

| ID | Name | Partition | Cluster | State | TimeSubmit | TimeStart | TimeEnd |
|-------|------------|-----------|---------|---------|---------------------|-----------|---------|
| 17614 | allocation | | cluster | PENDING | 2019-08-15T11:15:11 | Unknown | Unknown |

```
Would you like to fix these runaway jobs?
```

```
(This will set the end time for each job to the latest out of the start, eligible, or submit times, and set the state to completed.  
Once corrected, this will trigger the rollup to reroll usage from before the earliest submit time of all the runaway jobs.)
```

```
(You have 30 seconds to decide)
```

```
(N/y): y
```

Runway Jobs mess with
your accounting.
Stay clean.

There are other
autocorrection
commands, but you
shouldn't use them
unless proper support
advise you to do it.




External Process - scontrol [listpids | pidinfo]

```
NAME
    scontrol - Used view and modify Slurm configuration and state.

SYNOPSIS
    scontrol [OPTIONS...] [COMMAND...]

OPTIONS
    pidinfo proc_id
        Print the Slurm job id and scheduled termination time corresponding to the supplied process id, proc_id, on the current node. This will work only with processes on node on which scontrol is run, and only for those processes spawned by Slurm and their descendants.

    listpids [job_id[.step_id]] [NodeName]
        Print a listing of the process IDs in a job step (if JOBID.STEPID is provided), or all of the job steps in a job (if job_id is provided), or all of the job steps in all of the jobs on the local node (if job_id is not provided or job_id is "*"). This will work only with processes on the node on which scontrol is run, and only for those processes spawned by Slurm and their descendants. Note that some Slurm configurations (ProctrackType value of pgid) are unable to identify all processes associated with a job or job step.
```



```
$ sbatch -n 4 -w c6 --wrap "srun sleep 300"
Submitted batch job 395152

$ ssh c6

$ scontrol listpids 395152
PID      JOBID    STEPID  LOCALID  GLOBALID
21462    395152  0       0        0
21463    395152  0       1        1
21464    395152  0       2        2
21465    395152  0       3        3
21446    395152  batch  0        0
21448    395152  batch  -        -
21449    395152  batch  -        -

$ scontrol pidinfo 21464
Slurm job id 395152 ends at Fri Sep 04 15:50:47 2020
slurm_get_rem_time is 31535906
```

Proctrack Issues
(eg CG)
use to be related to
pstree / cgroups

External Process - taskset


```
NAME
    taskset - set or retrieve a process's CPU affinity

SYNOPSIS
    taskset [options] mask command [argument...]
    taskset [options] -p [mask] pid

OPTIONS
    -c, --cpu-list
        Interpret mask as numerical list of processors
        instead of a bitmask. Numbers are separated by
        commas and may include ranges. For example: 0,5,8-11.

    -p, --pid
        Operate on an existing PID and do not launch a new
        task.
```

Small “whereami” apps help a lot troubleshooting bindings, MPI, GPU ID...



```
$ sbatch --array=0-7 --ntasks-per-core=1 --wrap \
"srun bash -c 'printenv SLURMD_NODENAME; taskset -cp \${$}'"
Submitted batch job 520
$ tail slurm-520_*
==> slurm-520_0.out <==
c1
pid 25834's current affinity list: 0,2
==> slurm-520_1.out <==
c1
pid 25857's current affinity list: 1,3
==> slurm-520_2.out <==
c2
pid 25917's current affinity list: 0,2
==> slurm-520_3.out <==
c2
pid 25955's current affinity list: 1,3
==> slurm-520_4.out <==
c3
pid 25936's current affinity list: 0,2
==> slurm-520_5.out <==
c3
pid 25942's current affinity list: 1,3
==> slurm-520_6.out <==
c4
pid 25962's current affinity list: 0,2
==> slurm-520_7.out <==
c4
pid 25975's current affinity list: 1,3
```


External Process - stress(-ng)

```
NAME
  stress-ng - a tool to load and stress a computer system

SYNOPSIS
  stress-ng [OPTION [ARG]] ...

DESCRIPTION
  stress-ng will stress test a computer system in various selectable
  ways. It was designed to exercise various physical subsystems of a com-
  puter as well as the various operating system kernel interfaces.
  stress-ng also has a wide range of CPU specific stress tests that exer-
  cise floating point, integer, bit manipulation and control flow.
```




Copyright 2019 SchedMD
www.schedmd.com

Stressing jobs to
stress clusters
is stressing..
but
stress tests
save lives!

Prevent is better
than cure

Ensure config
works as expected
or
Replicate an issue

Logs



Logs are not an easy thing to deal with, but are (and will be) an important tool.



Logs - Two ways to increase logs

- Modular

- Specific subsystems
- DebugFlags (slurmctld and slurmdbd)

Be more modular
than verbose

- Levels

- More/Less verbosity
- slurmctld
 - Slurmctld(Syslog) Debug
 - SlurmSchedLogLevel
- slurmdbd
 - DebugLevel(Syslog)
- slurmd
 - Slurmd(Syslog) Debug

Logs - List of Flags and Levels

| | |
|---------------------|--|
| Backfill | Backfill scheduling |
| BurstBuffer | BurstBuffer state |
| CPU_Bind | Binding of tasks to CPUs |
| CpuFrequency | CPU governor and frequency management |
| Energy | Energy use |
| ExtSensors | External sensors |
| Gres | Generic Resource allocation and use |
| Gang | Gang scheduling |
| License | License scheduling |
| Priority | Job priority |
| Power | System power management |
| Reservations | Advanced reservations |
| SelectType | Resource selection |
| Steps | Job step activities |
| Switch | Network resources (e.g. the switch plugin) |
| Triggers | Event triggers |

Partial list.

See manpage of `slurm.conf` for complete listing

| | |
|----------------|--|
| quiet | Nothing |
| fatal | Only fatal errors |
| error | Errors and fatal errors |
| info | Errors and general messages |
| verbose | Errors and verbose messages |
| debug | Errors, verbose and debug messages |
| debug2 | Errors, verbose and more debug messages |
| debug3 | Errors, verbose and even more debug messages |
| debug4 | Errors, verbose and even more debug messages |
| debug5 | Errors, verbose and even more debug messages |

Don't be above error

Don't be below verbose if not troubleshooting

Logs - How to increase/decrease logs

- Changing Logging

- All daemons and commands
 - Command Line: `-vvv`
- All daemons
 - Edit `.conf` + reconfigure
- `scontrol setdebug LEVEL`
- `scontrol schedloglevel LEVEL`
- `scontrol setdebugflags [+|-] FLAG`

Increase logs, test
and reduce logs.

```
Increase logging verbosity
> scontrol setdebug debug2
> scontrol setdebugflags +backfill
Wait for event...
Restore original logging levels
> scontrol setdebug info
> scontrol setdebugflags -backfill
```

Use scontrol

Logs - Logrotate

- Don't restart, reconfigure or SIGHUP slurmctld, slurmdbd or slurmd to drop and reopen logfiles.



Use SIGUSR2
for logrotate!

How to be sure that the ball is in Slurm's court?

Slurm relies in a base setup.

Although it works on degraded setups, a healthy setup also helps troubleshooting.



Copyright 2019 SchedMD
www.schedmd.com

How to be sure that the ball is in Slurm's court?

- Network issues

- DNS
- Firewalls
- NTP (munge)

RPC plays
a key role

- Filesystem

- StateSaveLocation
- Unkillable



How to be sure that the ball is in Slurm's court?

- RAM usage on DB
 - Huge queries
 - Constant queries
 - Caching is not leaking

sdiag & stats
to detect DoS!

MaxQueryTimeRange
Archive & Purge



How to be sure that the ball is in Slurm's court?

- **Reliable constraints:** `cgroups`
 - `ConstrainDevices`
 - `ConstrainCores`
 - `ConstrainRAMSpace`
 - `ConstrainSwapSpace`
 - `ConstrainKmemSpace`

`JobAcctGatherParams/OverMemoryKill`
is limited to `JobAcctGatherFreq`

Demonized process in jobs
escape from `linuxproc`

`CUDA_VISIBLE_DEVICES`
is NOT a constraint

Avoid multiple `cgroups`
controllers in the same mount



How to be sure that the ball is in Slurm's court?

- Submit and Custom Plugins

No Comment

Have you tried [cli_filter](#)?



Why won't my job run?



Copyright 2019 SchedMD
www.schedmd.com

Why won't my job run?


- To understand this let's look at what happens in:
 - Scheduling / Backfill



Copyright 2019 SchedMD
www.schedmd.com

Scheduling & Backfill

- This is not an all inclusive training on scheduling & backfill rather an introduction to help you, as the admin, troubleshoot issues.



Copyright 2019 SchedMD
www.schedmd.com


Main scheduler (Backfill - Where it fits in)



- Slurm is designed to perform a quick and simple scheduling attempt at frequent intervals
 - At each job submission
 - At job completion on each of it's allocated nodes
 - At configuration changes
- Slower and more comprehensive scheduling attempts performed less frequently

Backfill

- The Backfill scheduler will start lower priority jobs if by doing so does not delay the expected start time of any higher priority job (configurable exceptions)
- Setting accurate and reasonable run times is required for backfill to start lower priority jobs.



Backfill

- Making sure you have correctly tuned your SchedulerParameters is key when trying to push your sites workflow through optimally.
- Parameters such as:



Backfill

- Making sure you have correctly tuned your SchedulerParameters is key when trying to push your sites workflow through optimally.
- Parameters such as:
 - `bf_window`
 - The number of minutes into the future to look when considering jobs to schedule.
 - A value at least as long as the highest allowed time limit is generally advisable to prevent job starvation.



Backfill


- Making sure you have correctly tuned your SchedulerParameters is key when trying to push your sites workflow through optimally.
- Parameters such as:
 - `bf_window`
 - `bf_continue`
 - Setting this option will cause the backfill scheduler to continue processing pending jobs from its original job list after releasing locks even if job or node state changes.



Backfill

- Making sure you have correctly tuned your SchedulerParameters is key when trying to push your sites workflow through optimally.
- Parameters such as:
 - bf_window
 - bf_continue
 - bf_resolution
 - The number of seconds in the resolution of data maintained about when jobs begin and end.

How fine the lines are
on the graph paper



Backfill



- Making sure you have correctly tuned your SchedulerParameters is key when trying to push your sites workflow through optimally.
- Parameters such as:
 - bf_window
 - bf_continue
 - bf_resolution
- And many other tunable parameters based on your needs.
 - See the slurm.conf documentation

Backfill

- Ensuring that your scheduler is configured correctly can fix or improve job throughput and help solve job starvation issues.



Copyright 2019 SchedMD
www.schedmd.com

Why won't my job run? - continued



Copyright 2019 SchedMD
www.schedmd.com

Why won't my job run? - continued

- There are enough resources available.



Why won't my job run? - continued

- There are enough resources available.
- The cluster looks empty.



Why won't my job run? - continued

- There are enough resources available.
- The cluster looks empty.
- Do I have the right permissions to run?



Why won't my job run? - continued

- One common question we see is regarding a job that a user believes should run on a node that is already partially allocated.



Copyright 2019 SchedMD
www.schedmd.com

Why won't my job run? - continued

Example: CPU with 4 cores

gres.conf:

Name=gpu Type=p100 File=/dev/nvidia0 COREs=0,1

Name=gpu Type=gpu File=/dev/nvidia1 COREs=2,3

Job 123 is on cores 0,1

Job 124 needs the p100

Why won't my job run? - continued

- As soon as some reason is found why a job cannot be started, that is recorded in the job's "reason" field and the scheduler moves on to the next job

Some common reasons why jobs are pending:

| | |
|---------------------------------|--|
| Priority | Resources being reserved for higher priority job |
| Resources | Required resources are in use |
| Dependency | Job dependencies not yet satisfied |
| Reservation | Waiting for advanced reservation |
| AssociationJobLimit | User or account job limit reached |
| AssociationResourceLimit | User or account resource limit reached |
| AssociationTimeLimit | User or account time limit reached |
| QOSJobLimit | Quality Of Service (QOS) job limit reached |
| QOSResourceLimit | Quality Of Service (QOS) resource limit reached |
| QOSTimeLimit | Quality Of Service (QOS) time limit reached |

Why won't my job run? - continued

```
$ scontrol show job 21543
JobId=21543 JobName=wrap
UserId=agil(1000) GroupId=agil(1000) MCS_label=N/A
Priority=251879 Nice=0 Account=development QOS=normal
JobState=PENDING Reason=Dependency Dependency=afterok:21542
ReqQueue=1 Restarts=0 BatchFlag=1 Reboot=0
RunTime=00:00:00 TimeLimit=UNLIMITED
SubmitTime=2019-09-02T18:08:37 EligibleTime=2019-09-02T18:08:37
AccrueTime=2019-09-02T18:08:37
StartTime=Unknown EndTime=Unknown DeadlineTime=Unknown
SuspendTime=None SecsPreSuspend=0 LastSchedTime=2019-09-02T18:08:37
Partition=batch AllocNode:Sid=agil-work_station:22616
ReqNodeList=(null) ExcNodeList=(null)
NodeList=(null)
NumNodes=1 NumCPUs=1 NumTasks=1 CPUs/Task=1 ReqB:S:C:T=0:0:*:*
TRES=cpu=1,mem=512M,node=1,billing=1
Socks/Node=* NtasksPerN:B:S:C=0:0:*:* CoreSpec=*
MinCPUsNode=1 MinMemoryCPU=512M MinTmpDiskNode=0
Features=(null) DelayBoot=00:00:00
OverSubscribe=OK Contiguous=0 Licenses=(null) Network=(null)
Command=(null)
WorkDir=/home/agil/workspace/slurm/bugs/7468
StdErr=/home/agil/workspace/slurm/bugs/7468/slurm-21543.out
StdIn=/dev/null
StdOut=/home/agil/workspace/slurm/bugs/7468/slurm-21543.out
Power=
```

Reason Dependency
not satisfied

```
$ scontrol show job 26315
JobId=26315 JobName=wrap
UserId=jason(1000) GroupId=jason(1000) MCS_label=bio
Priority=25416 Nice=0 Account=bio QOS=normal
JobState=PENDING Reason=AssocMaxJobsLimit Dependency=(null)
ReqQueue=1 Restarts=0 BatchFlag=1 Reboot=0 ExitCode=0:0
RunTime=00:00:00 TimeLimit=00:00:00 TimeMin=N/A
SubmitTime=2019-09-12T13:50:39 EligibleTime=2019-09-12T13:50:39
AccrueTime=2019-09-12T13:50:39
StartTime=Unknown EndTime=Unknown DeadlineTime=Unknown
SuspendTime=None SecsPreSuspend=0 LastSchedTime=2019-09-12T13:50:40
Partition=batch AllocNode:Sid=jason-work_station:22616
ReqNodeList=(null) ExcNodeList=(null)
NodeList=(null)
NumNodes=1 NumCPUs=1 NumTasks=1 CPUs/Task=1 ReqB:S:C:T=0:0:*:*
TRES=cpu=1,node=1,billing=1
Socks/Node=* NtasksPerN:B:S:C=0:0:*:* CoreSpec=*
MinCPUsNode=1 MinMemoryNode=0 MinTmpDiskNode=0
Features=(null) DelayBoot=00:00:00
OverSubscribe=OK Contiguous=0 Licenses=(null) Network=(null)
Command=(null)
WorkDir=/home/jason/slurm/master
StdErr=/home/jason/slurm/master/slurm-26315.out
StdIn=/dev/null
StdOut=/home/jason/slurm/master/slurm-26315.out
Power=
```

Reason Max Job limit

Why won't my job run? - continued

- We continue to improve logging (for example)
 - Logs for BadConstraints have been improved see the following.
 - <https://github.com/SchedMD/slurm/commit/20c2b6151d6d10fee7e012555d30f9fd529b7dc3>



Copyright 2019 SchedMD
www.schedmd.com

Bugzilla@Mozilla - Main Page version 4.0.6+

[Home](#) | [New](#) | [Browse](#) | [Search](#) | [?] | [Reports](#) | [Requests](#) | [New Account](#) | [Log In](#) | [Forgot Password](#)

Welcome to Bugzilla



[Get Help](#)



[File a Bug](#)



[Search](#)



[Open a New Account](#)

[Quick Search help](#) | [Install the Quick Search plugin](#)

[Bugzilla User's Guide](#) | [Release Notes](#) | [Bugzilla Etiquette](#) | [Bug Writing Guidelines](#)

[Home](#) | [New](#) | [Browse](#) | [Search](#) | [?] | [Reports](#) | [Requests](#) | [New Account](#) | [Log In](#) | [Forgot Password](#)

[Privacy Policy](#)

Bugzilla - Consider the following

- Security
- Proper data / attachments
- Ticket visibility



Problem Reporting - Security Notice



- Bugzilla is by default: open to the internet
- Make sure to avoid uploading configs that contain passwords such as `slurmdbd.conf` (scrub the password).
- Attachments and comments can be marked as private by SchedMD

Problem Reporting - Proper data / attachments

- Report using bugzilla
 - <https://bugs.schedmd.com/>
- Identify Slurm version
- Provide configuration files
 - Essential to identify the plugins used
- Provide logs demonstrating the problems
- Make sure you compress large files



Practically perfect in every way!

Problem Reporting - Ticket Visibility

- When needed SchedMD can limit the visibility of comments, attachments and the entire bug if needed.



References

- <https://slurm.schedmd.com/troubleshoot.html>
- <https://slurm.schedmd.com/faq.html>
- <https://slurm.schedmd.com/news.html>
- https://slurm.schedmd.com/quickstart_admin.html (#upgrade)

- https://slurm.schedmd.com/SLUG18/field_notes2.pdf
- <https://slurm.schedmd.com/SLUG17/FieldNotes.pdf>

- <https://bugs.schedmd.com>
- <http://groups.google.com/group/slurm-users>

Pay attention to NOTES in manpages.

Troubleshooting



Questions?

Thanks

Copyright 2019 SchedMD
www.schedmd.com