



# **Field Notes Mark Two: Random Musings From Under a New Hat**

Tim Wickberg  
SchedMD

Slurm User Group Meeting 2018

# Field Notes - Overview



- Successor to last years talk - “Field Notes From The Frontlines of Slurm Support”
- <https://slurm.schedmd.com/SLUG17/FieldNotes.pdf>
- ... which serves in part as a best-practices guide.
  - I won't be repeating that content here; I'd suggest taking a look through it if you haven't seen it before.

# Field Notes - Overview



- Feel free to ask questions throughout.
  - Although I may ask to defer more involved discussions until later in the interest of time, and the attention of the rest of the audience.

# Field Notes - Overview



- “Best practices observed from three years behind SchedMD's customer support organization, alongside an impractical demonstration of Slurm's API capabilities, and musings on future features and direction for Slurm in the coming years.”
- TLDR: an ~hour of me rambling. : )



# Upgrading

- There is a specific sequence to use when moving between major Slurm releases.

# Upgrading



slurmdbd

>=

slurmctld

>=

slurmd

>=

commands



Must stay within 3 major releases.

E.g., {18.08, 17.11, 17.11, 17.02} is okay,  
but {18.08, 17.11, 17.02, 16.05} is not.

Within each major release, you can mix the  
maintenance release versions without issue.  
E.g., {18.08.0, 17.11.7, 17.11.9, 17.11.4} is okay.

# Upgrading

- RPMs do make this process difficult to do with the system live.
- While we ship and support the slurm.spec file, we do not actually recommend using RPMs to install Slurm.
- We suggest structuring installs in version-specific directories, and using symlinks and/or module files to manage versions.
  - This makes rolling upgrades much simpler.



# Example installation hierarchy

```
# ./configure --prefix=/apps/slurm/18.08.0/ --sysconfdir=/apps/slurm/etc/  
# ln -s /apps/slurm/18.08.0 /apps/slurm/dbd  
# ln -s /apps/slurm/18.08.0 /apps/slurm/ctld  
# ln -s /apps/slurm/18.08.0 /apps/slurm/d  
# ln -s /apps/slurm/18.08.0 /apps/slurm/current
```

Use the appropriate symlink in each service file, and add /apps/slurm/current symlink into \$PATH (through /etc/profile.d/ or a module file).

This makes a rolling upgrade much simpler, just move the symlink when ready to move that component forward onto the newer release.

# Upgrading



- Backing up the MySQL database used by slurmdbd is strongly encouraged when upgrading.
  - You should probably be doing this already as part of a regular backup strategy, but this would be a good time to make sure it works.
  - For larger databases, or more unusual systems, you may want to test the upgrading/conversion on a copy of the full production database on a separate machine.
  - Older MySQL versions (5.5 and before) have had problems with later conversion processes.

# Upgrading



- slurmdbd will automatically convert the MySQL schema.
  - This can take ~10-15 minutes or more depending on the size of the database.
  - If using the systemd service file, we do recommend handling the conversion directly on the command line, otherwise systemd may kill the conversion if it thinks it's not proceeding.
  - 17.11 and older slurmdbd processes did not write the PID until after the conversion, leading to systemd thinking the process was hung. This has changed in 18.08, and the PID is written out before conversion.

# Upgrading



- Taking a backup of StateSaveLocation is also recommended.
- Once a daemon has been upgraded, you cannot roll back to a prior major version without loss of data and/or your job queue.



# Node Addition and Removal



- Adding and removing nodes in Slurm is a sensitive operation.
  - This seems to cause problems for each site at least once early on.
- Certain internal datastructures are built off the node list at startup, and are used within the communication subsystems.
- Changing the Node definitions, and restarting only the `slurmctld`, will usually lead to communication errors as messages as misrouted internally.

# Node Addition and Removal



Safe procedure:

1. Stop slurmctld
2. Change configs
3. Restart all slurmd processes
4. Start slurmctld

# Node Addition and Removal



Less-Safe, but usually okay, procedure:

1. Change configs
2. Restart slurmctld
3. Restart all slurmd processes really quickly



# Node Addition and Removal



- We do have plans to make this less painful long-term.
- The new `cons_tres` plugin has split some of these datastructures apart, and will eventually let us change this.
- But refactoring to address this is blocked until `cons_res` is removed.



# Development and Patch Submission



- We love to see external submissions, and do try to provide timely feedback regardless of who submitted.
  - You do not need to be a SchedMD customer to submit patches.
- While we at SchedMD handle the bulk of development, some major subsystems have originated externally, and we love that aspect of open-source.
  - E.g., `pam_slurm_adopt` was conceived by BYU.

# Development and Patch Submission



- Please read CONTRIBUTING.md for style guidelines and notes on submission and review.
- Patches that add new functionality should be of general interest and usable by other sites.
  - We generally do not land features only usable by a specific site.
- New functionality is only permitted in the next release; please make sure patches are based on the master branch.

# Development and Patch Submission



- Bug fixes should be based on the most recent stable release branch. Right now that's slurm-18.08.
- Please be patient with the review process; we do thoroughly review all patches, and this can be a significant investment of time and effort.
- Minor style corrections we will generally address; more significant issues we'll bring back to you to discuss and correct.

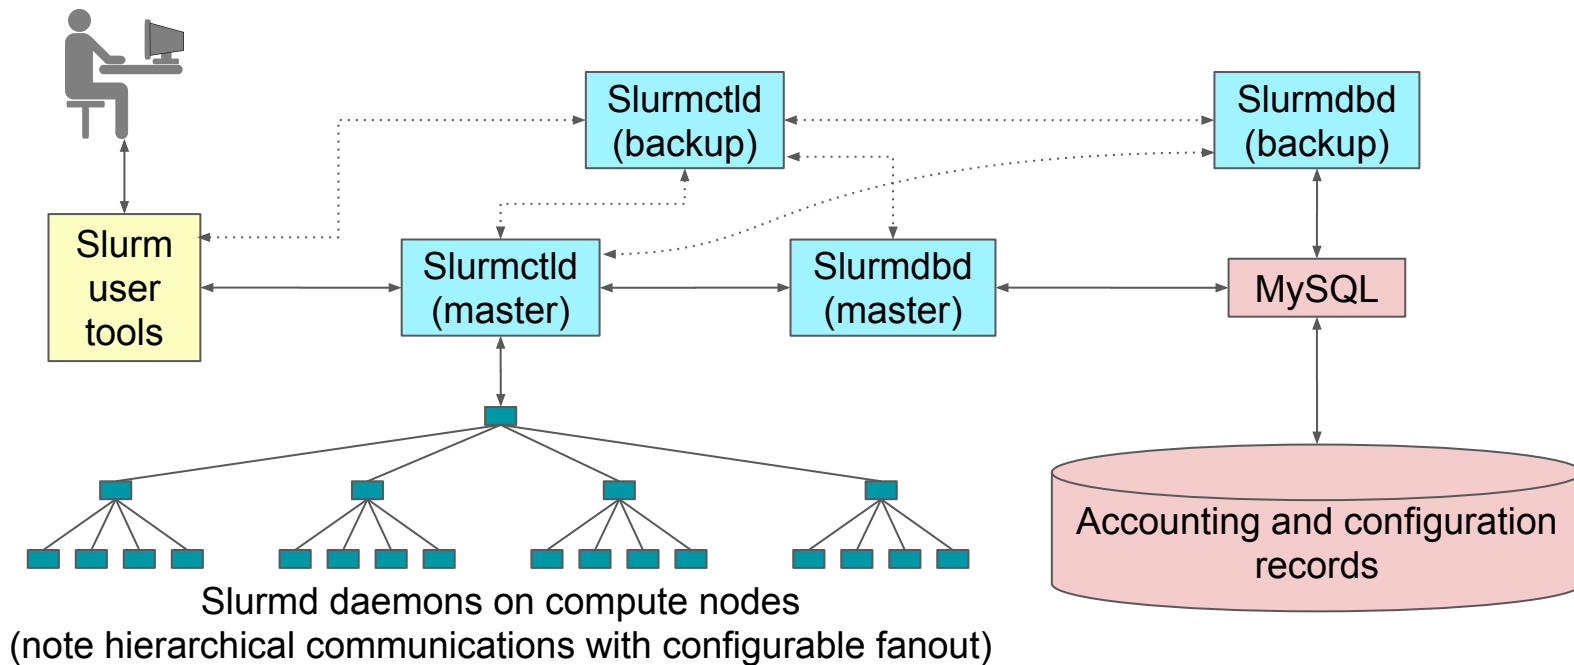
# Development and Patch Submission



- Documentation is critical, and must land alongside any new features.
  - Testsuite additions are also appreciated.
- When submitting, please attach patches to a new bug in our Bugzilla instance.
  - <https://bugs.schedmd.com>
  - Set the Severity to “C - Contributions”
  - ‘git format -am’ strongly preferred for patches.



# Cluster Architecture - Typical Linux Cluster



Copyright 2018 SchedMD LLC

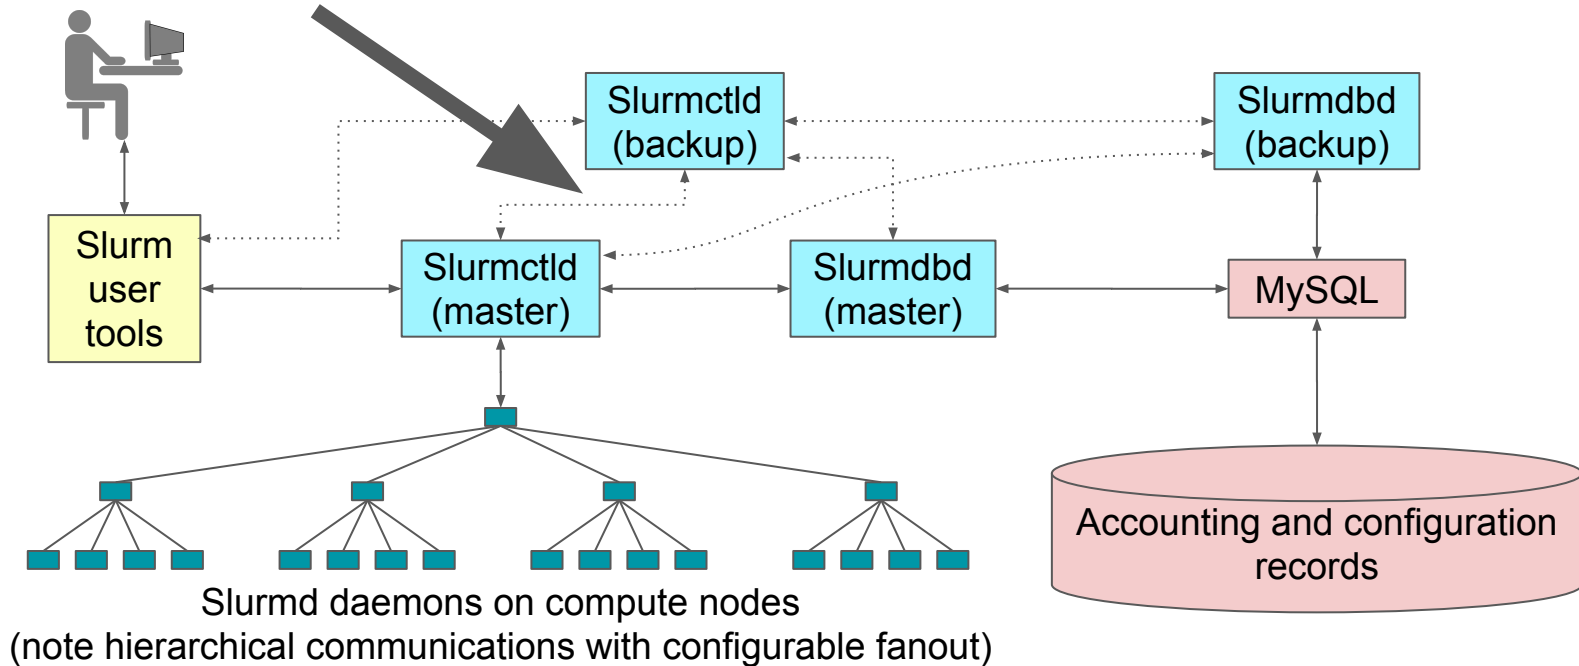
[www.schedmd.com](http://www.schedmd.com)

USC Sept. 17-19, 2018



# Cluster Architecture - Typical Linux Cluster

StateSaveLocation is hiding in here, and establishes your failure domain for slurmctld.



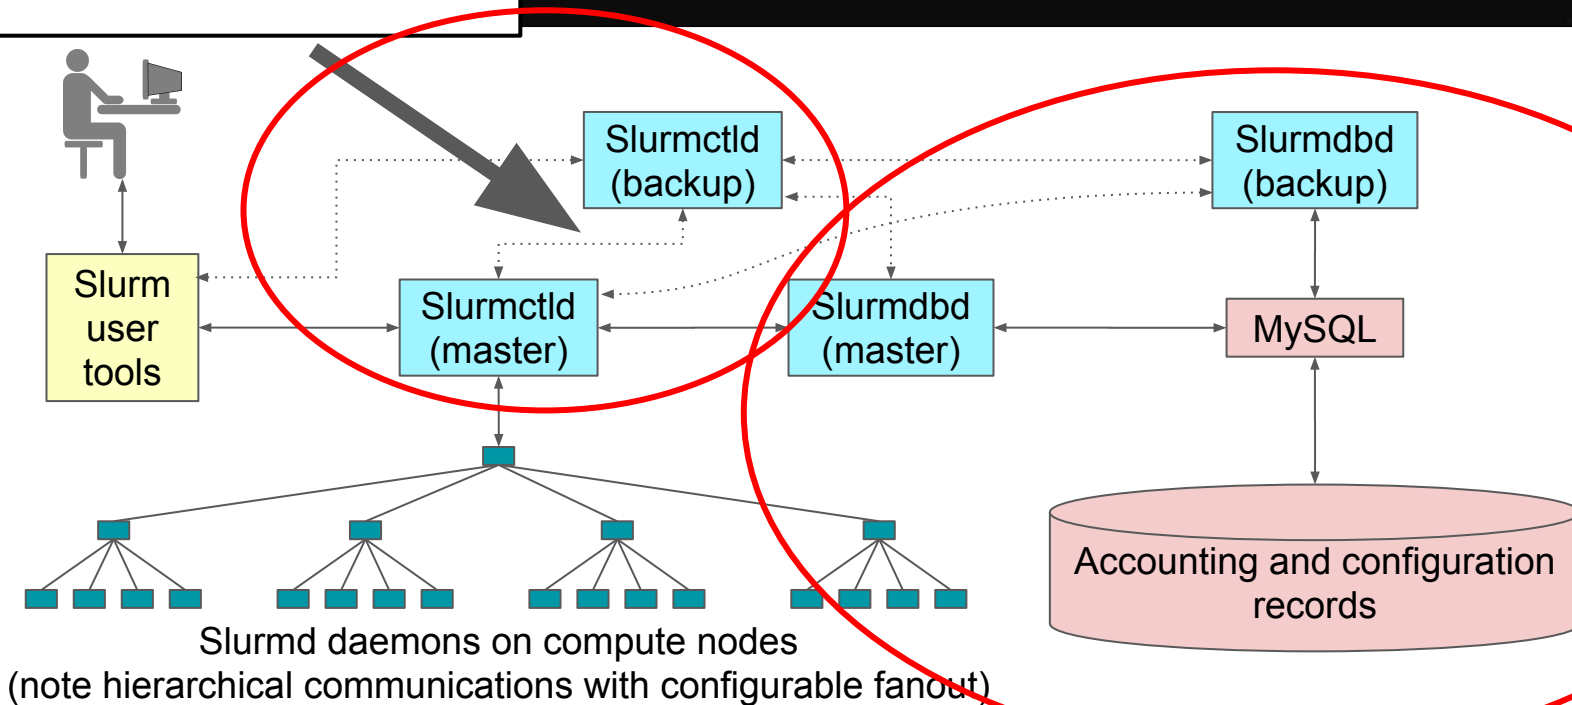
Copyright 2018 SchedMD LLC

[www.schedmd.com](http://www.schedmd.com)

USC Sept. 17-19, 2018

# Cluster Architecture - Typical Linux Cluster

StateSaveLocation is hiding in here, and establishes your failure domain for slurmctld.

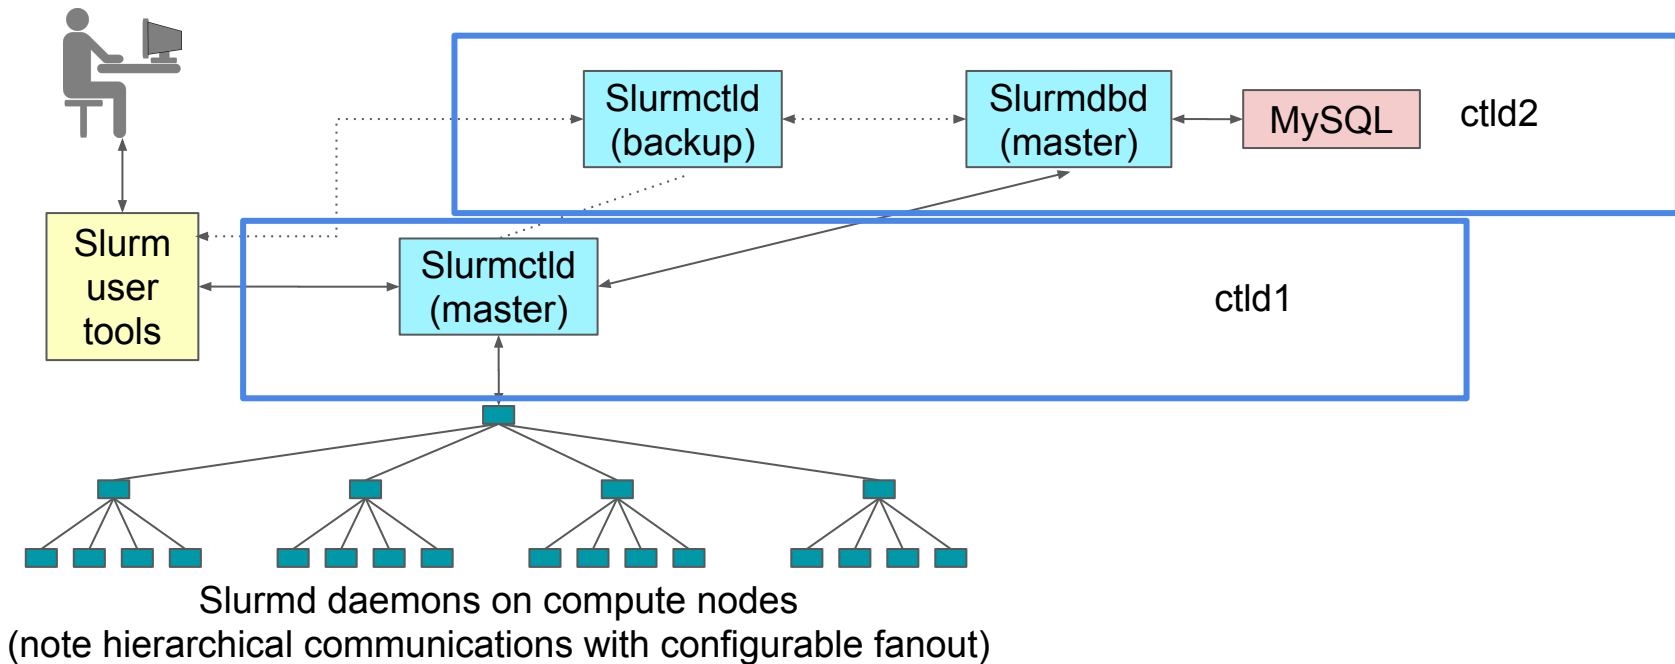


Copyright 2018 SchedMD LLC

[www.schedmd.com](http://www.schedmd.com)

USC Sept. 17-19, 2018

# My Preferred Deployment Pattern



Copyright 2018 SchedMD LLC

[www.schedmd.com](http://www.schedmd.com)

USC Sept. 17-19, 2018

# Justification

- Separating slurmctld and slurmdbd in normal production use is recommended.
- Master/backup slurmctld is common, and - as long as the performance for StateSaveLocation is kept high - not that difficult to implement.

# One aside on StateSaveLocation



- Your maximum system throughput, and overall Slurm controller responsiveness under heavy load, will be governed by latency reading/writing from StateSaveLocation.
- In high-throughput (~200k+ jobs/day) environments, you may be much better off with a local NVMe drive in a single controller.
  - Especially if the alternative is an NFS mount shared with users that gets frequently hammered... you're more likely to see performance issues related to this than an outage from the controller system dying.

# Justification

- For slurmdbd, the critical element in the failure domain is MySQL, not slurmdbd. slurmdbd itself is stateless.
- slurmctld will cache accounting records (up to a limit) if slurmdbd is unavailable. This can be hours+ to days+ depending on your system without data loss.

# Justification

- IMNSHO, the additional complexity of a redundant MySQL deployment is more likely to cause an outage than it is to prevent one.
- So don't bother setting up a redundant slurmdbd, keep slurmdbd + MySQL local to a single server.

# Related Networking Notes



- Semi-frequently reported issue: changes made through sacctmgr aren't immediately showing up in slurmctld.
- But restarting slurmctld fixes this somehow?



# Related Networking Notes



- This is almost always an issue with a firewall in between slurmctld and slurmdbd.
- slurmdbd opens a new connection to slurmctld to push changes.
- If you've firewalled that off, the update will not be propagated.
- Restarting slurmctld forces it to open a connection to slurmdbd - the other direction - and pull down a full copy.

# Related Networking Notes



- slurmctld -> slurmdbd is the prevalent pattern, so this is easy to overlook, and everything else will work.
- The slurmdbd logs should be telling you this is an issue... but easy to miss this as well.



# Recommended Configuration



- My current recommended deployment has:
  - proctrack/cgroup, task/cgroup, jobacct\_gather/cgroup
  - cgroup.conf - ConstrainCores, ContrainDevices, ConstrainRAMSpace, ConstrainSwapSpace all enabled
  - PrologFlags=contain with pam\_slurm\_adopt setup appropriately
  - LaunchParameters=send\_gids
  - ReconfigFlags=KeepPartInfo

# LaunchParameters=send\_gids

- LDAP at scale sucks.

# LaunchParameters=send\_gids

- LDAP at scale sucks.
- How many people here have LDAP + sssd on your compute nodes?

# LaunchParameters=send\_gids

- LDAP at scale sucks.
- How many people here have LDAP + sssd on your compute nodes?
- And how many of you have noticed problems with this?
- Especially with missing extended group memberships?

# LaunchParameters=send\_gids



- This flag instructs the slurmctld to look up the extended group id list for a given user, and embed that securely as part of the launch credential.
- So the slurmd/slurmstepd has this provided, and does not need to look this up on the compute node on step launch.
- Which avoids the “thundering herd” problem at large scales when 100+ nodes all hit LDAP simultaneously.



# UNIX Groups Cause Problems



- Random trivia:
  - Partition settings include Allow/DenyAccount, Allow/DenyQOS, Allow/DenyUsers... but only AllowGroups.
  - This is intentional since we regularly see this missing group membership issue on production systems.
    - Which would defeat the access control.

# Demo #1



# nss\_slurm

- The slurmd/slurmstepd process on the compute node already has most of the important details about a given user.
- Do you need user or group info for people without active jobs on a node? I suspect the answer is “no” for most systems...
- So why even bother with LDAP/SSSD/syncing /etc/{passwd,group}?

# nss\_slurm

- Proof-of-concept; not ready, and Slurm will need to be modified to pass along additional user and group information it has not bother with in the past.
- Likely some version in 19.05, but we have not committed to it yet.
- Do people think this would be useful?
  - Are there things with it that I haven't thought of?
  - Bug #5773 will be tracking progress.

# Demo #1

- Things to highlight:
  - Cross-architecture Slurm support.
    - “blackhole” is amd64
    - “cashbox” is armv7l
  - Possible ways to extend Slurm’s realm of responsibility to cover other cluster concerns.

# Demo #2



- What if “pay-to-play” on your cluster had a physical manifestation?

# Demo #2



- Introducing the Prioramatic 5000™

# Demo #2



- Things to highlight:
  - Demo is showing two approaches to changing relative user priorities in an automatic way.
    - Setting a negative nice value administratively.
    - Setting a QOS with a higher priority that the user does not have access to already.
  - Slurm API use to query data and manipulate job settings
    - Along with Cgo bindings to libslurm.



# Questions?

