



Set up a cluster with SLURM

Institute of Computer Science

High-Performance Storage

David Nelles

Supervisor: Vanessa End

February 23, 2023




SLURM

- Simple Linux Utility for Resource Management
- Workload manager
- Installation
 - Package manager
 - Compile on system






Workload manager




Workload manager






SLURM Setup

Minimal setup



SLURM Setup


Minimal setup






SLURM Setup

Minimal setup




Multiple control nodes common






Installing SLURM





Installing SLURM





Install Prerequisites

Compiling

- C-Compiler (GCC)
- Perl
- Python3



Install Prerequisites

Compiling

- C-Compiler (GCC)
- Perl
- Python3

Execution

- Munge
- NFS



Setup Munge (Authentication service)

- 1 Install via package manager
- 2 Create munge user
- 3 Generate mung key and distribute on all nodes
- 4 Enable and start daemon



Why we need NFS

Controller nodes

- SLURM config file
- Database config file

Compute nodes

- SLURM config file



Why we need NFS

Controller nodes

- SLURM config file
- Database config file

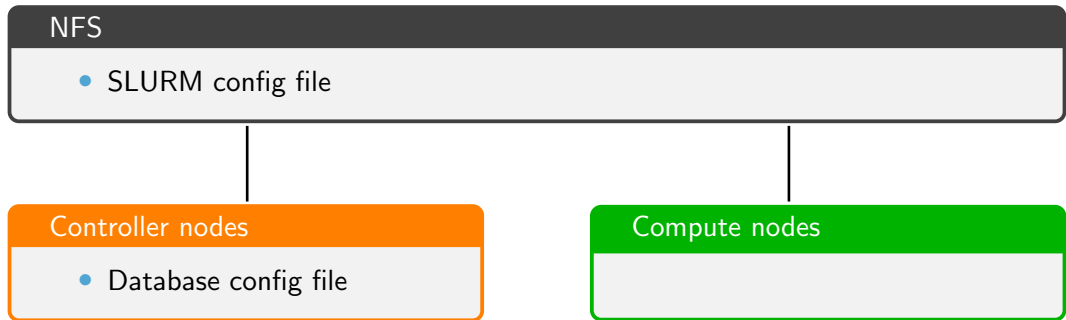
Compute nodes

- SLURM config file

Must be the same configuration

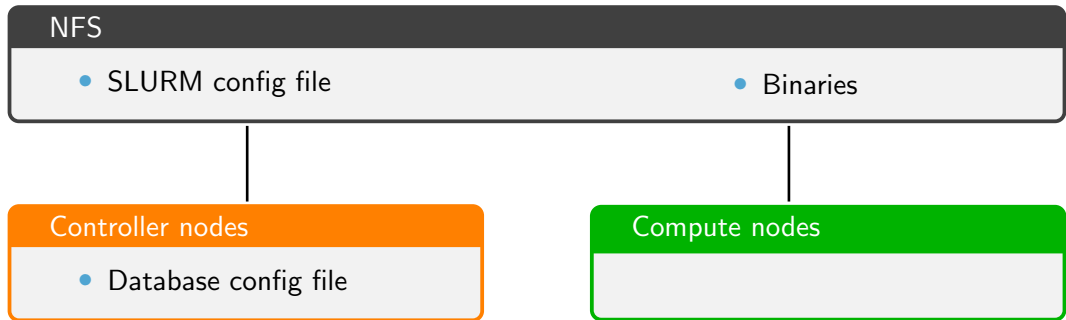


Why we need NFS






Why we need NFS






NFS-Structure





NFS-Structure


- Configuration files





NFS-Structure


- Configuration files
- Log files





NFS-Structure


- Configuration files
- Log files
- Soft link to *version*





NFS-Structure


- Configuration files
- Log files
- Soft link to *version*
- Containing a specific version of SLURM





NFS-Structure


- Configuration files
- Log files
- Soft link to *version*
- Containing a specific version of SLURM
- Source code





NFS-Structure

- Configuration files
- Log files
- Soft link to *version*
- Containing a specific version of SLURM
- Source code
- Installation (Binaries)





Setup NFS

- 1 Create file system structure
- 2 Install NFS from package manager
- 3 Enable via systemd
- 4 Configure in `/etc/exports` (Host)
- 5 Configure `fstab` (Clients)



Setup NFS

- 1 Create file system structure
- 2 Install NFS from package manager
- 3 Enable via systemd
- 4 Configure in `/etc/exports` (Host)
- 5 Configure `fstab` (Clients)

`/etc/exports:`

```
# direcotry          client-ip  config    client-ip  config
/media/nfs          10.254.1.5(rw, sync)  10.254.1.11(rw, sync)
```

Setup NFS

- 1 Create file system structure
- 2 Install NFS from package manager
- 3 Enable via systemd
- 4 Configure in `/etc/exports` (Host)
- 5 Configure `fstab` (Clients)

`/etc/exports:`


```
# direcotry          client-ip  config    client-ip  config
  /media/nfs        10.254.1.5(rw, sync)  10.254.1.11(rw, sync)
```

`fstab:`

```
<server-ip>:<shared-directory> /mnt/nfs nfs4 defaults ,user ,exec 0 0
```




Download SLURM sources






Download SLURM sources



- Download from official website: (<https://www.schedmd.com/downloads.php>)
- Clone from GitHub: (<https://github.com/SchedMD/slurm.git>)




Configure Build






Configure Build



- Move sources to NFS
- Execute configure script



Configure Build




- Move sources to NFS
- Execute configure script

`prefix` Directory which should contain the installation.
`sysconfdir` Directory which contains the configuration files.



Build SLURM



- Use `make` to compile
- Use `make install` to install
- Always use superuser



Configure SLURM

Configuration file can be generated via the configurator.

Event logging

- SlurmctldDebug
- SlurmctldLogFile
- SlurmdDebug
- SlurmdLogFile

Control machines

- SlurmctldHost
- BackupController

Compute machines

- NodeName
- CPUs

Accounting

- JobAcctGatherType
- AccountingStorageType
- AccountingStorageHost
- AccountingStoragePort



Start SLURM

Use systemd for that

Controller nodes

- **slurmctld**
`sudo systemctl enable --now slurmctld`
- **slurmdbd**
`sudo systemctl enable --now slurmdbd`

Compute nodes

- **slurmd**
`sudo systemctl enable --now slurmd`



Challenges



Missing dependencies

Invalid MailProg configuration

Reason Missing mail client

Fix Install `mailx`

Cannot find cred/munge plugin

Reason Missing munge
development libraries

Further Could not find them on dnf

Fix Enable `powertools`
repository



Wrong locations

Systemd cannot find the service files

Expected by systemd

```
/usr/lib/systemd...  
/system/slurm*.service
```

Location of service files

```
/mnt/nfs/slurm/current...  
/source/etc/slurm*.service
```



Wrong locations

Systemd cannot find the service files

Expected by systemd

```
/usr/lib/systemd...  
/system/slurm*.service
```

Location of service files

```
/mnt/nfs/slurm/current...  
/source/etc/slurm*.service
```

Solution

Softlinks pointing to the service files



Wrong cgroup configuration

- Cgroup was not configured
- Create config file in etc directory

```
CgroupAutomount=yes  
ConstrainCores=no  
ConstrainRAMSpace=no
```



Accounting



Accounting prerequisites

Already fulfilled

- Munge

Need to be fulfilled

- MariaDB
- MariaDB development libraries



Accounting setup

Setup process

- 1 Make sure that `MariaDB-devel` is installed when compiling SLURM
- 2 Create a further configuration file
- 3 Clarify that SLURM config fits database config



Accounting setup

Setup process

- 1 Make sure that `MariaDB-devel` is installed when compiling SLURM
- 2 Create a further configuration file
- 3 Clarify that SLURM config fits database config

Important parameters

- `DdbHost`
- `StorageType`
- `StorageUser`
- `StoragePass`
- `StoragePort`



Recap

Basic SLURM setup

Challenges

Accounting setup



Recap

Basic SLURM setup

- Installing Prerequisites
- Compiling SLURM
- Configuring SLURM

Challenges

Accounting setup



Recap

Basic SLURM setup

- Installing Prerequisites
- Compiling SLURM
- Configuring SLURM

Challenges

- Missing dependencies
- Wrong systemd location
- Missing cgroup configuration

Accounting setup



Recap

Basic SLURM setup

- Installing Prerequisites
- Compiling SLURM
- Configuring SLURM

Challenges

- Missing dependencies
- Wrong systemd location
- Missing cgroup configuration

Accounting setup

- Dependencies
- Configuration



Set up a cluster with SLURM

Institute of Computer Science

High-Performance Storage

David Nelles

 david.nelles@stud.uni-goettingen.de

Supervisor: Vanessa End

February 23, 2023