



## Status on GPU support

EESSI Community Meeting @ Amsterdam

15 Sept 2022

Alan O'Cais (CECAM), Michael Hübner (U Bonn / HPC.NRW)

# Design challenges for GPU support in EESSI

- Dealing with CUDA drivers
- Installing CUDA compatibility drivers
- Installing CUDA itself
- Testing that CUDA software is working
- Making CUDA software visible *only* when everything is working





# Dealing with CUDA drivers

DRV: Regular driver release branch every 3 months

NFB: New feature branch

Production Branch: Driver branch supported for 1 year

LTSB: Long Term Support Branch supported for 3 years




- CUDA compatibility libraries offer a way out
- Ability to use CUDA compatibility libraries depend on driver version
  - If compat library was released after driver EOL then it most likely won't work
  - But... can only really know by checking, need to iterate until we get a working combination
  - Should probably recommend that people stick to latest LTSB driver (470)

# Installing CUDA Compatibility Libraries




- CUDA compatibility libraries are released per distro
  - Libraries themselves are identical (at least in versions checked)
  - Can just focus on RHEL8 RPMs
    - Largest number of versions of CUDA compatibility libraries
    - Support likely for quite a few years
    - Need p7zip in the EESSI stack to unpack them
- Place CUDA compatibility libraries in a place where *our* linker will automatically find them
  - Tweak the default linker search paths of the Gentoo Prefix linker
  - Choose a special location common to all EESSI stacks as we will always want the latest usable version of these libraries
  - Will only work when *our* linker is used!

# Installing CUDA itself




- CUDA installation is fat
  - Need ~15GB available to do the installation
- Most likely we will never be allowed to ship CUDA
  - US export regulations would likely prohibit this
  - What about just runtime libraries? (Doesn't solve driver issue though)
- Need for EESSI to be able to **see** the CUDA installation
  - Create an EasyBuild hook so that CUDA installation is located in a **fixed** alternate path
    - Use `#{EESSI_SOFTWARE_PATH/versions/host_injections}`
      - Only relevant for the software, **not** the module
    - Tweak Lmod to refuse to load CUDA module unless this path exists
      - Error with a message that explains the EESSI GPU installation process

# Testing that CUDA software is working



- Only way to really check is to compile and run something
  - Can use `deviceQuery` for this
    - Currently this compiled during the installation process
  - CUDA samples no longer shipped with CUDA
    - Can we ship (compiled) CUDA samples with EESSI?
- Allow 5 iterations of CUDA compatibility libraries before we give up
- Could extend checks easily once we start shipping software

# Making CUDA software visible only when everything works




- Use an EasyBuild hook to add an Lmod property to anything that depends on CUDA
  - Using existing Lmod GPU property, but this would need to change with AMD GPU support
- Use Lmod hook to hide modules with that property unless certain condition(s) are met
  - Exact condition(s) are WIP (perhaps existence of specific file in CUDA/compat installation?)
- Leave CUDA itself visible so people *are* tempted to try and load it
  - Module load error will detail how they can actually enable CUDA support and lead them to any associated documentation

# Current status



- Most of the effort is in PR from Michael Hübner
  - <https://github.com/EESSI/software-layer/pull/172>
- Works for use cases investigated
  - Tested on “live” system and within build containers
  - LAMMPS, GROMACS, CUDA samples compiled and run on live system
  - There are some limitations to CUDA compat libraries, yet to see if these will be encountered by actual software we install
- Lmod hook requires very recent Lmod
  - Releasing a new EESSI version would also allow simplifying some things





**Paper (open access):** <https://doi.org/10.1002/spe.3075>

Website: <https://www.eessi-hpc.org>

**Join our mailing list & Slack channel**

<https://www.eessi-hpc.org/join>

Documentation: <https://eessi.github.io/docs>

GitHub: <https://github.com/eessi>

Twitter: [@eessi\\_hpc](https://twitter.com/eessi_hpc)

[youtube.com/channel/UCKLS5X7\\_oMWhUrAZuzSwBxQ](https://youtube.com/channel/UCKLS5X7_oMWhUrAZuzSwBxQ)

Monthly online meetings (first Thursday, 2pm CEST)