

Alexander Grund
Center for Information Services and High Performance Computing (ZIH)

EasyBuild at ZIH - TU Dresden

5th EasyBuild User Meeting, Barcelona
31st January, 2020

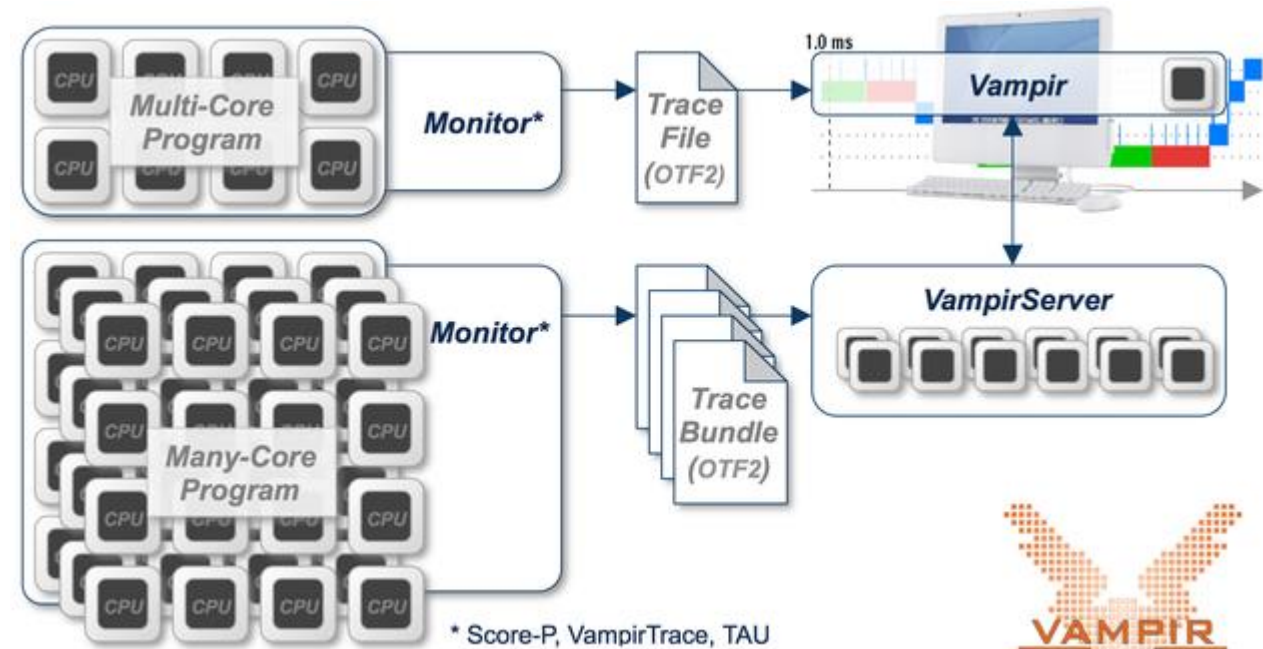
TU (University of Technology) Dresden

- Since 1828/1961
- Ca. 32 000 students, 8 300 employees
- 18 faculties, 121 disciplines
- Focus on Biomedicine, Bioengineering, Materials sciences, Information technology, Microelectronics, Energy and Environment
- Since 2012 “University of Excellence” with Clusters
 - Center for Advancing Electronics Dresden (cfaed)
 - Physics of Life
 - Complexity and Topology in Quantum Materials
 - ...



Center for Information Services and High Performance Computing (ZIH)

- Central scientific institution
- Member of the Gauß-Allianz
- Tasks:
 - Communication infrastructure
 - Operation of central IT infrastructure
 - Support for other departments
- Own Research and development work



HPC Cluster

Main cluster: Taurus

- ~47 000 processors
- 270 Sandy Bridge nodes (gone)
- ~1600 Haswell nodes
 - 64 - 256 GB RAM
 - 128 GB local SSD
- 44 GPU nodes (2x K20), Sandy Bridge
- 64 GPU nodes (4x K80)
- 32 Power9 nodes (6x V100)
- Some consumer GPU nodes (GTX 1080Ti)




HPC Cluster

Special Islands

- SMP nodes
 - 2 nodes @ 1 TB RAM
 - 4 nodes @ 2 TB RAM

- Intel Xeon Phi
 - 32 nodes w/ Knights Landing

- HPC Data Analytics
 - 90 nodes
 - 2 PB NVMe flash memory



HPC Cluster

Upcoming:

- 192 nodes with AMD EPYC 7702 processors
 - 2 / node → 128 cores / node
- > 750 TFlops

“Superdome Flex”

- 32 Intel Cascade Lake 8276 → 896 cores
- 48 TB RAM in unified address space
- 400 TB NVMe memory



Software environment

Previously:

- RHEL 6
- EnvironmentModules 3.2.1 (TCL/C Modules)
- Custom build software & modules

Since 10/2016:


- Lmod as alternative (opt-in)
- Separate EasyBuild environment: modenv/eb vs modenv/classic

Since 08/2017:

- Lmod as default
- New software build with EasyBuild (exceptions to modenv/classic)

Since 06/2018:

- Bull SCS5 w/ RHEL 7
- Clean new modenv/scs5 w/ EasyBuild modules (currently ~1300)
- Flat module structure due to user habits
- Some custom EasyConfigs, no custom EasyBlocks



Module environment

Problem:

Heterogeneous cluster with multiple CPU architectures (westmere, sandy bridge, haswell, ...)

Solution:

- Shared module folder
 - Separate folders for software (binaries etc.)
 - /sw/installed is symlinked to /sw/haswell or /sw/sandy on every node
-
- Most optimized builds w/ custom build settings possible
 - Transparent for users
 - Regression checks for “hardware dependent” bugs
 - Symlinks may lead to trouble with build systems → Better use bind mounts

Now mixture again:

- modenv/scs5 with Symlinks
- modenv/ml as separate module environment for Power9

User experience

- SLURM
 - 1 Architecture on every partition
- Lustre
 - Different storages (from SSD to tape)
 - Workspaces → Expire after given period
 - Some NFS
- Lmod

- Extensive wiki
 - Instructions for using SLURM
 - Examples for using certain software
 - List of installed modules

- Ticket system
 - Issues with system
 - Software requests
 - General help



Software support

- For long time only 1 person allowed to install software
- Some exceptions for particular software, e.g. Vampir

- Now 2-3 people supporting the software stack

- Mostly official EasyConfigs used
- Some custom ones required (minor changes, different version)
- Some user-provided EasyConfigs
- No EasyBlocks written so far
- Some software still installed manually (LS-DYNA, STAR-CCM+)

Software support – EasyBuild extensions

Helper script for installation: eb_zih

- Determine architecture and set up install targets (/sw/{haswell,sandy,ml})
- Load EasyBuild module and settings
- Set up group write permissions
- Add --group-writable-installdir to eb command

Separate script:

- Copies module files to user-visible location
- Check for module availability on all architectures

- EB Hooks for CMAKE_PREFIX_PATH in module files
- Shared folder for custom ECs, migration to git repo planned

Conclusion

- EasyBuild greatly simplifies software installation
- Community EasyConfigs save a lot of time
- Toolchain concept of EasyBuild makes cutting-edge software installation hard
 - Lot's of EasyConfigs to create on toolchain update
 - --try-* options nice but doesn't handle dependencies well
 - Trainee project for a webapp in work to help
- Mid term goal: Hierarchical module structure
 - Already in use in CLARA cluster from DLR