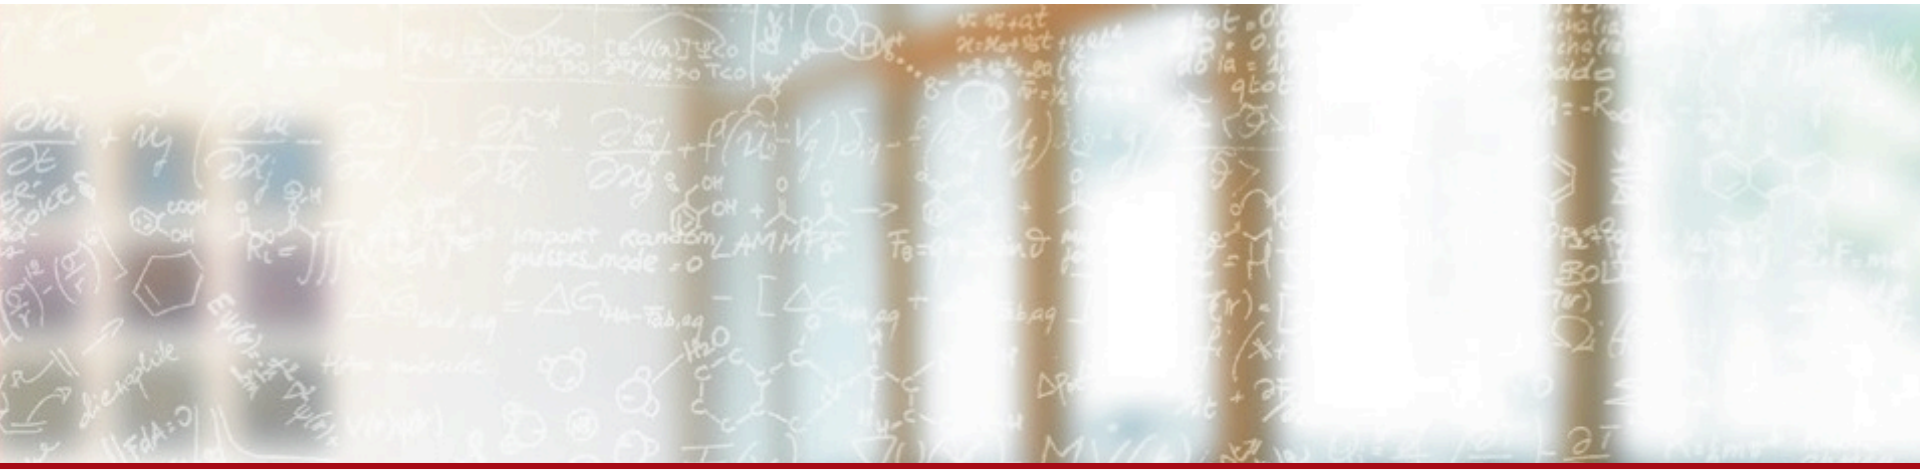

EasyBuild @ CSCS: Site Update

6th EasyBuild User Meeting
Jan 25th – 29th 2021, Online Event

Luca Marsella

Scientific Computing Support, Swiss National Supercomputing Center (CSCS)

Outline

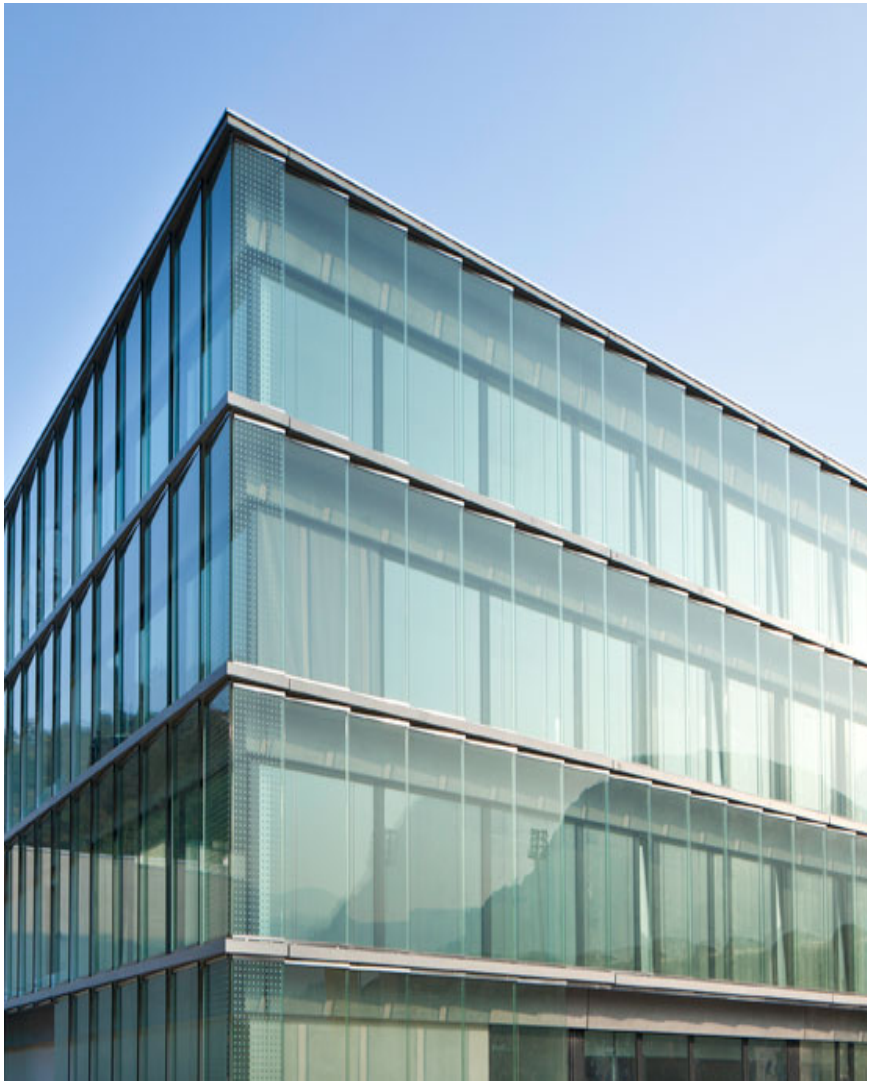

- **EasyBuild timeline @ CSCS**
- Overview of CSCS HPC systems
  - MeteoSwiss systems
  - Piz Daint
  - Alps
- EasyBuild for CSCS Users
  - Custom User builds
- Easybuild with Jenkins
  - Jenkins pipelines
    - ProductionEB
    - ReBuildEB
    - TestingEB
    - UpdateEB
- Final Remarks

EasyBuild timeline @ CSCS

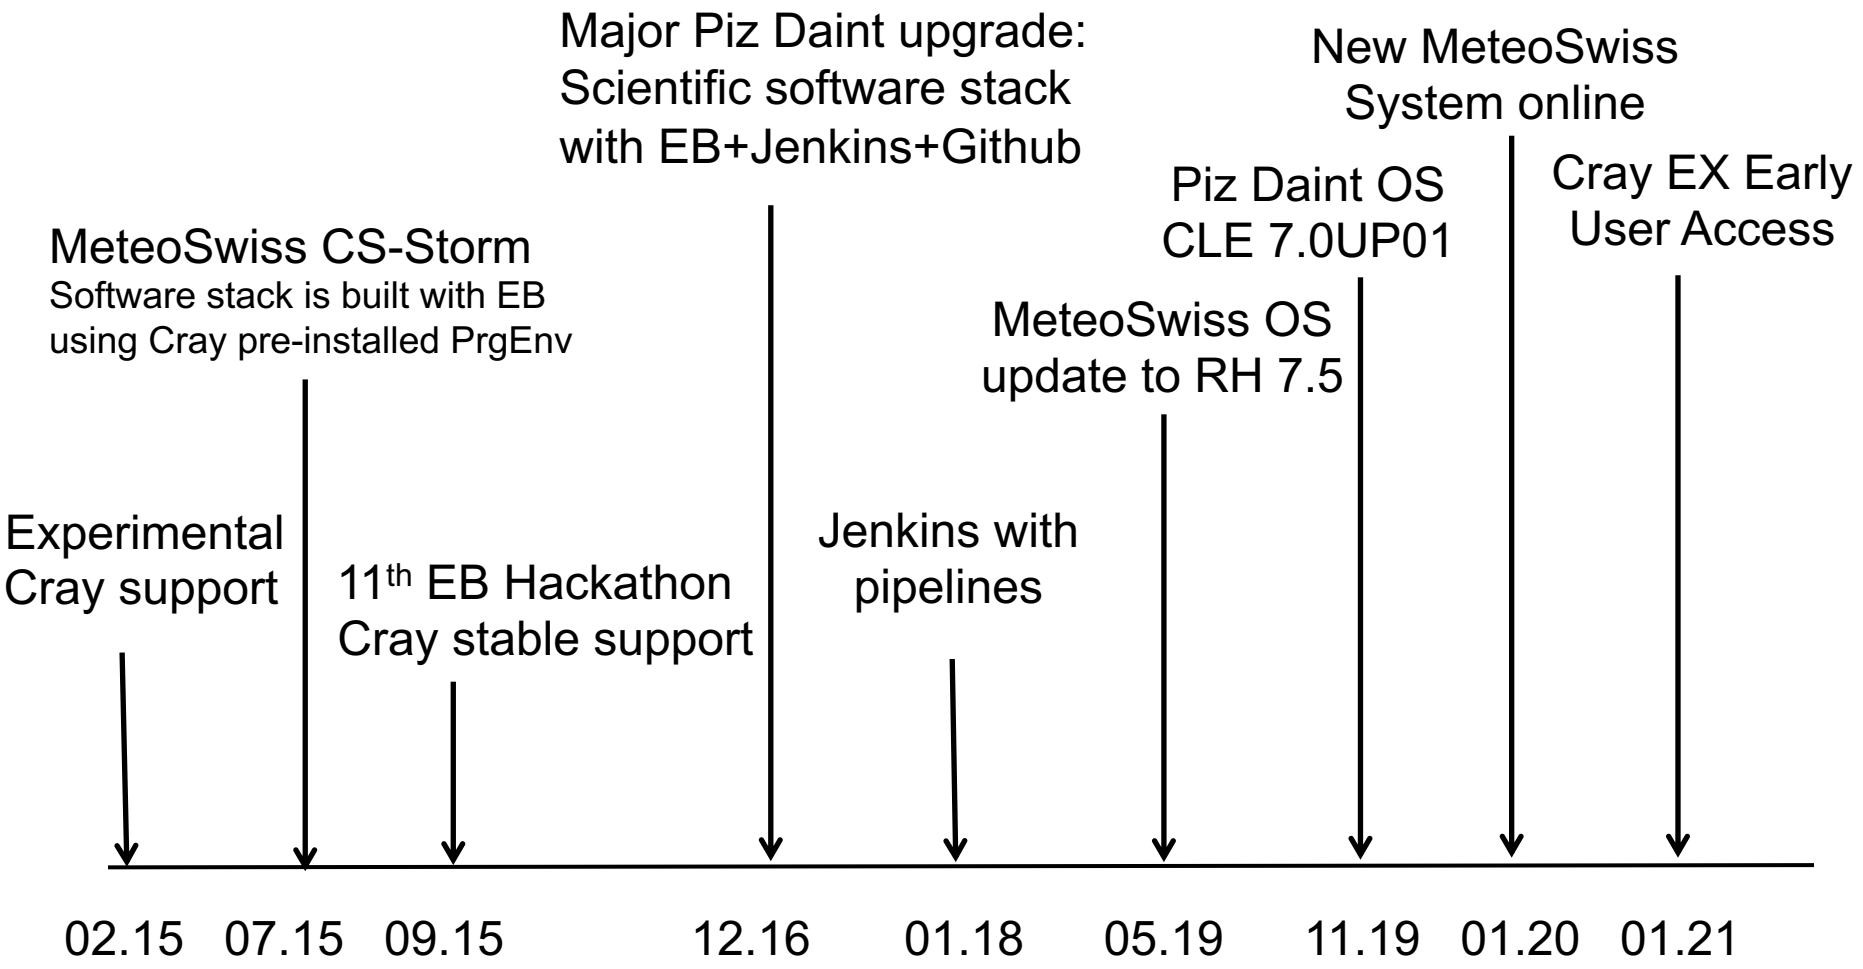

Outline

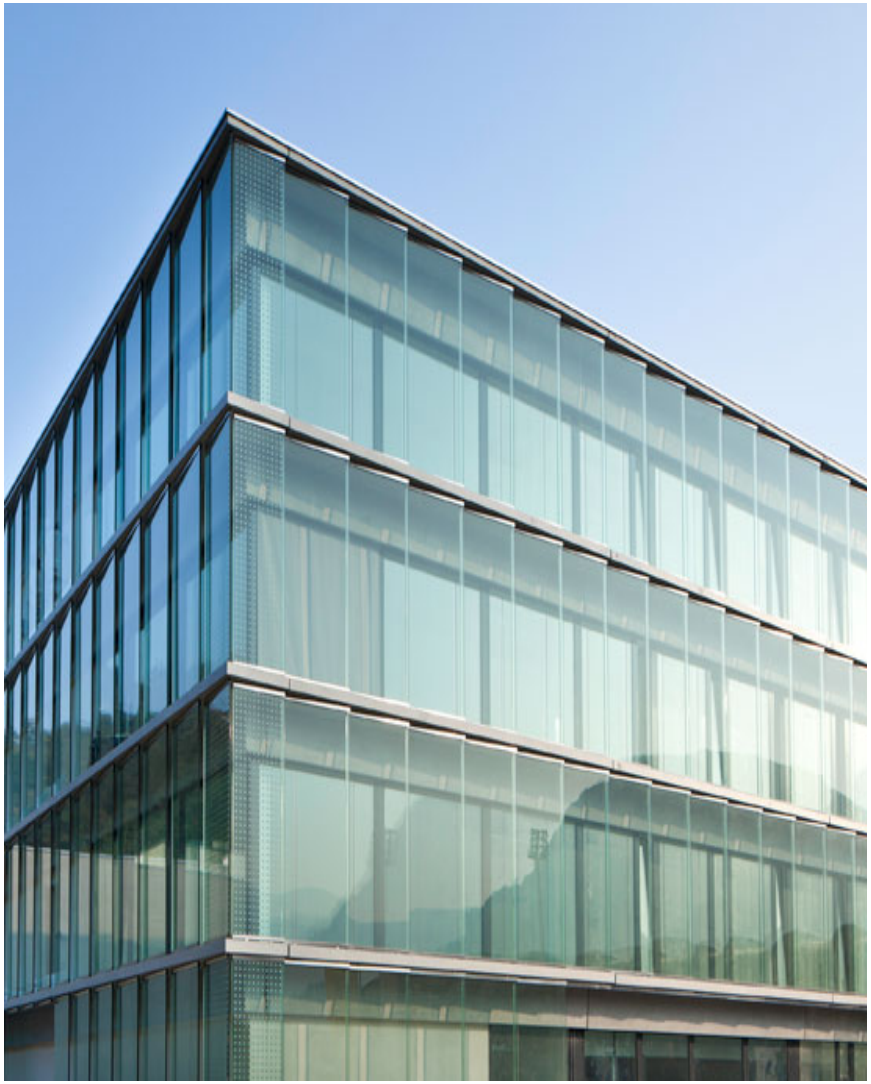

- EasyBuild timeline @ CSCS
- **Overview of CSCS HPC systems**
  - **MeteoSwiss systems**
  - **Piz Daint**
  - **Alps**
- EasyBuild for CSCS Users
  - Custom User builds
- Easybuild with Jenkins
  - Jenkins pipelines
    - Production EB
    - ReBuild EB
    - Testing EB
    - Update EB
- Final Remarks

Overview of CSCS HPC systems

System	Scope	Accelerators / node	Architecture
Piz Daint	CSCS User Lab	1 GPU	Cray XC50 / XC40 P100,Haswell/Broadwell
Alps	Early User Access	CPU only	Cray EX AMD Rome
Arolla / Tsa	MeteoSwiss	8 GPU	V100, Intel SkyLake

MeteoSwiss Production System

Arolla and Tsa are the names of the new systems:

- Compute nodes ft. Intel Skylake and Tesla V100
- EB software stack available since 01.20

Module names are lowercase, with few exceptions:

- EasyBuild-custom (CSCS EasyBuild modulefile)
- PrgEnv-gnu
- PrgEnv-pgi

Meta-modules provide a hierarchical environment:

E.g.: PrgEnv-pgi/19.9 unfolds the following modules

- hdf5/1.10.5-pgi-19.9-gcc-8.3.0
- netcdf-c++/4.3.0-pgi-19.9-gcc-8.3.0
- netcdf-fortran/4.4.5-pgi-19.9-gcc-8.3.0
- netcdf/4.7.0-pgi-19.9-gcc-8.3.0
- openmpi/4.0.2-pgi-19.9-gcc-8.3.0-cuda-10.1
  pgi/19.9-gcc-8.3.0

Alps

Alps (Eiger) is a HPE/Cray EX Supercomputing system open to selected users

Compute Nodes (CN) feature two sockets with one AMD EPYC™ 7742 64-Core processor per socket, that interface with the high-speed HPE Slingshot interconnect

Current EB software stack features:
- Custom toolchains to address the new Cray CPE (cpeCCE, cpeGNU)
- Modules versions of Cray PE are pinned to avoid using collections
- Lmod supported as of 20.10 (modules based on Lua: Version 8.3.1)

Scientific software available in the software stack:
- CDO/1.9.9-cpeGNU-20.10
- CP2K/7.1-cpeGNU-20.10
- GROMACS/2020.4-cpeGNU-20.10
- GSL/2.6-cpeCCE-20.10
- LAMMPS/29Oct20-cpeGNU-20.10
- NAMD/2.14-cpeGNU-20.10
- QuantumESPRESSO/6.6-cpeGNU-20.10
- VASP/6.1.0-cpeGNU-20.10

Piz Daint

Model	Cray XC50 / XC40
XC50 node	Intel® Xeon® E5-2690 v3 (Haswell) @ 2.60GHz (12 cores, 64GB RAM) and NVIDIA® Tesla® P100 16GB
XC40 node	Intel® Xeon® E5-2695 v4 (Broadwell) @ 2.10GHz (18 cores, 64/128 GB RAM)
Login node	Intel® Xeon® CPU E5-2650 v3 @ 2.30GHz (10 cores, 256 GB RAM)
Interconnect	Aries routing and communications ASIC, and Dragonfly network topology
Scratch	8.8 PB (Lustre / Sonexion 3000)

Flagship production system with hybrid nodes for the User Lab at CSCS:

- EB software stack in production since 11.16
- Successfully updated OS to CLE 7.0UP02 in 11.20
- Automated update of Easyconfig files in production for new Cray PEs

Scientific Software built with EasyBuild on Piz Daint

Amber/20-3-0-CrayIntel-20.08-cuda
Boost/1.70.0-CrayGNU-20.08-python3
CDO/1.9.5-CrayGNU-20.08
CP2K/7.1-CrayGNU-20.08-cuda
CPMD/4.1-CrayIntel-20.08
GREASY/19.03-cscs-CrayGNU-20.08
GROMACS/2020.3-CrayGNU-20.08-cuda
GSL/2.5-CrayGNU-20.08
HPX/1.4.0-CrayGNU-20.08-cuda
Horovod/0.19.1-CrayGNU-20.08-tf-2.2.0
Julia/1.5.0-CrayGNU-20.08-cuda
JuliaExtensions/1.5.0-CrayGNU-20.08
LAMMPS/03Mar20-CrayGNU-20.08-cuda
NAMD/2.13-CrayIntel-20.08-cuda
NCL/6.4.0
NCO/4.8.1-CrayGNU-20.08
PLUMED/2.5.1-CrayGNU-20.08
ParaView/5.8.1-CrayGNU-20.08-EGL-python3
QuantumESPRESSO/6.5a1-CrayPGI-20.08-cuda
QuantumESPRESSO-SIRIUS/6.5-rc4-CrayGNU-20.08-cuda
VASP/6.1.0-CrayIntel-20.08-cuda
VTK/9.0.0-CrayGNU-20.08-EGL-python3

Amber/20-3-0-CrayIntel-20.08
Boost/1.70.0-CrayGNU-20.08-python3
CDO/1.9.5-CrayGNU-20.08
CP2K/7.1-CrayGNU-20.08
CPMD/4.1-CrayIntel-20.08
GREASY/19.03-cscs-CrayGNU-20.08
GROMACS/2020.3-CrayGNU-20.08
GSL/2.5-CrayGNU-20.08
HPX/1.4.0-CrayGNU-20.08
Julia/1.5.0-CrayGNU-20.08
JuliaExtensions/1.5.0-CrayGNU-20.08
LAMMPS/03Mar20-CrayGNU-20.08
NAMD/2.13-CrayIntel-20.08
NCL/6.4.0
NCO/4.8.1-CrayGNU-20.08
PLUMED/2.5.1-CrayGNU-20.08
ParaView/5.8.1-CrayGNU-20.08-OSMesa-python3
QuantumESPRESSO/6.5-CrayIntel-20.08
Spark/2.3.1-CrayGNU-20.08-Hadoop-2.7
VASP/6.1.0-CrayIntel-20.08
Visit/3.1.2-CrayGNU-20.08

Outline

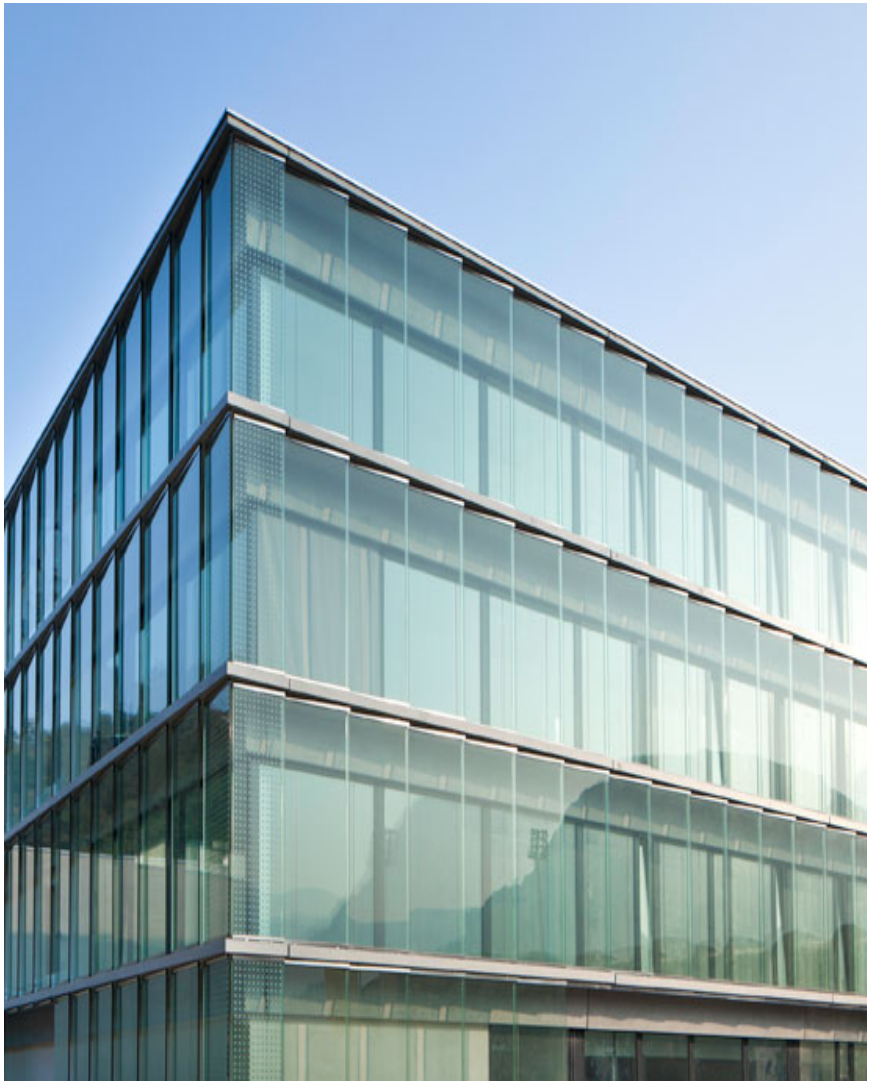

- EasyBuild timeline @ CSCS
- Overview of CSCS HPC systems
  - MeteoSwiss systems
  - Piz Daint
  - Alps
- **EasyBuild for CSCS Users**
  - Custom User builds
- Easybuild with Jenkins
  - Jenkins pipelines
    - Production EB
    - ReBuild EB
    - Testing EB
    - Update EB
- Final Remarks

EasyBuild for CSCS Users

- User are given EasyBuild recipes when requesting software
  - Instead of error-prone manual steps on how to build and run
- EasyBuild documentation on the CSCS User Portal

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

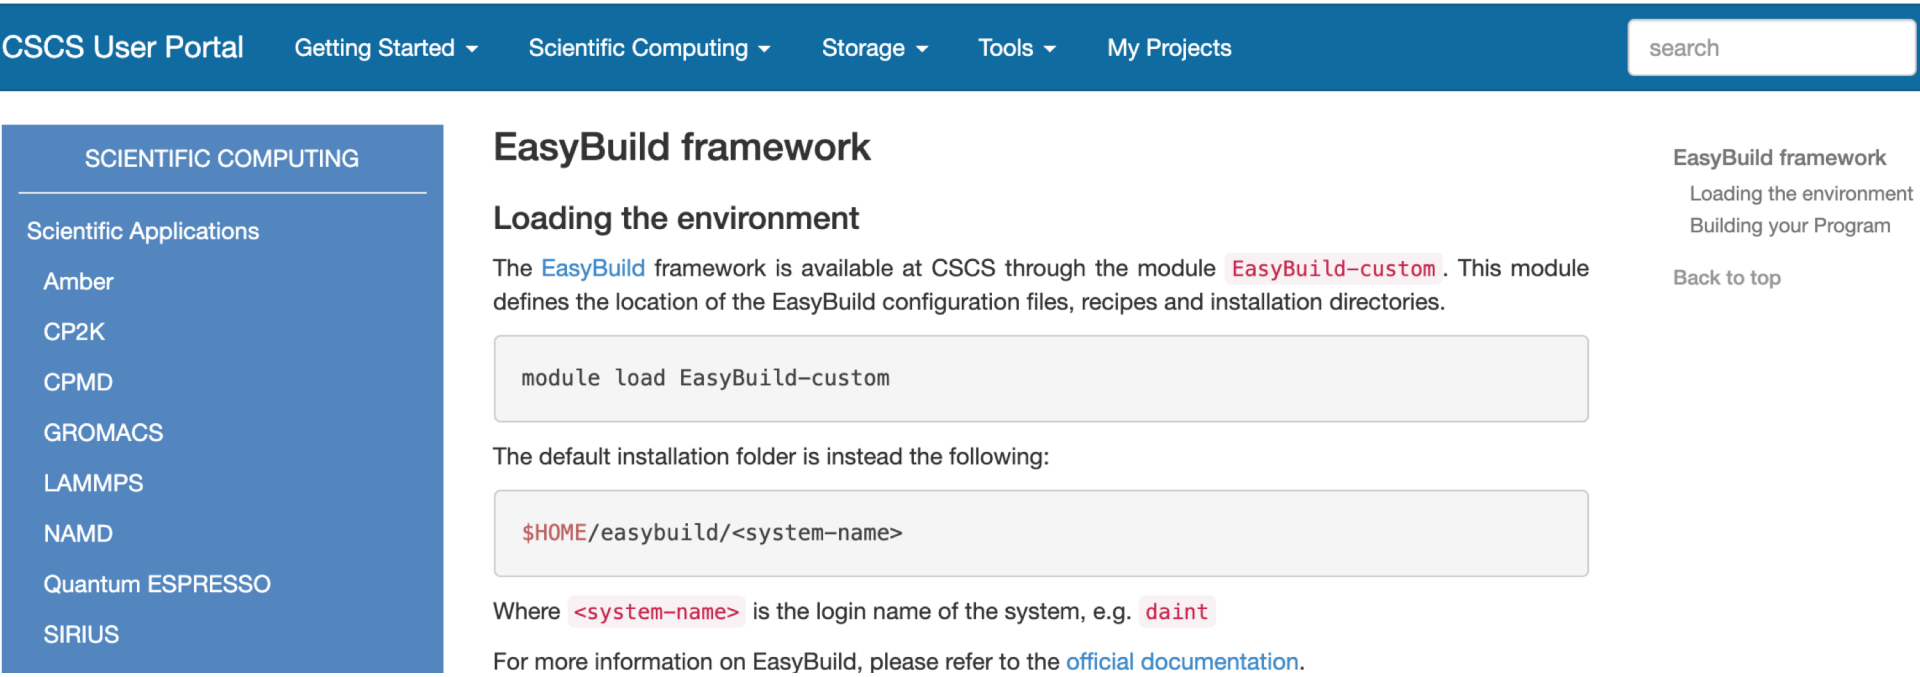

Custom User builds

- Users can extend or customize the CSCS EasyBuild recipes cloning the CSCS project production from GitHub
  - git clone https://github.com/eth-cscs/production.git
- Export the corresponding EasyBuild environment variable:
  - **EB_CUSTOM_REPOSITORY=/<path>/production/easybuild**
    module load EasyBuild-custom/cscs
- The modulefile **EasyBuild-custom/cscs** will add CSCS production Easyconfigs to the local robot path for search

Outline

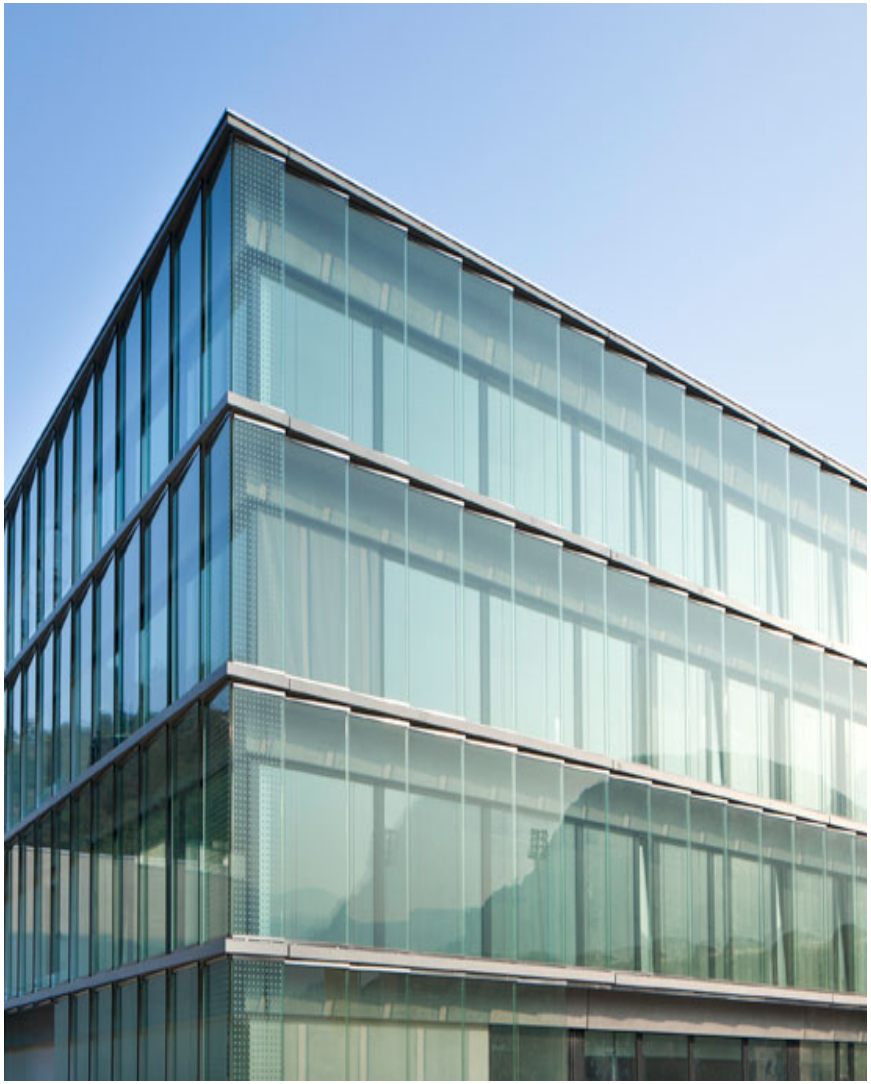

- EasyBuild timeline @ CSCS
- Overview of CSCS HPC systems
  - MeteoSwiss systems
  - Piz Daint
  - Alps
- EasyBuild for CSCS Users
  - Custom User builds
- **Easybuild with Jenkins**
  - **Jenkins pipelines**
    - **Production EB**
    - **ReBuild EB**
    - **Testing EB**
    - **Update EB**
- Final Remarks

EasyBuild with Jenkins

- Jenkins service for Continuous Integration
  - Deploy software packages on the systems in production
  - Test new Easyconfig files submitted by staff and users
  - Check regressions of Easyconfigs listed in production
  - Update production recipes in view of system upgrades
- Jenkins project names run with EasyBuild
  - ProductionEB builds the Easyconfigs once they are in production
  - TestingEB is triggered when a new pull request appears on Github
  - ReBuildEB runs EasyBuild from scratch to ensure reproducibility
  - UpdateEB runs EasyBuild to update recipes for a new Cray PE
- Jenkins projects defined by **Pipelines**
  - Enhanced flexibility of the actions performed by Jenkins
  - **Jenkinsfile** script of each project is version controlled
  - The CI can run in parallel optimizing the available resources

CSCS EB production repository on GitHub

- **How to submit a pull request:**
  - Add the EasyBuild configuration files to a **new branch** in your **fork**
  - The pull request must include **all the required dependencies**
- Create **a pull request** following these policies:
  - the title **must match a supported system**, otherwise the build will fail
  - The system names **have to be enclosed in square brackets**
  - if the title matches **WIP** ("Work In Progress"), then the test build will be aborted immediately, as work in progress is not supposed to be tested
  - Dom and Piz Daint can test both software stacks **-gpu** and **-mc** at once:
    - if the title matches only ${system}-gpu or ${system}-mc, only that is used:
      - [dom-gpu] NAMD will build using -gpu, [dom-mc] NAMD will use -mc
    - if the title matches both or none, then both will be used, one after another:
      - [dom] NAMD will build using both -gpu and -mc in a loop
      - [dom-gpu,dom-mc] NAMD will do the same
- The CSCS Jenkins project **TestingEB** will test the build of new EasyBuild recipes. The Jenkins pipeline script is available at https://github.com/eth-cscs/production/tree/master/jenkins

ProductionEB Pipeline

ProductionEB < 1600 | Pipeline | Changes | Tests | Artifacts | Logout

Branch: –
Commit: –
15m 50s
a day ago
Changes by Guilherme Peretti-Pezzi
Started by an SCM change

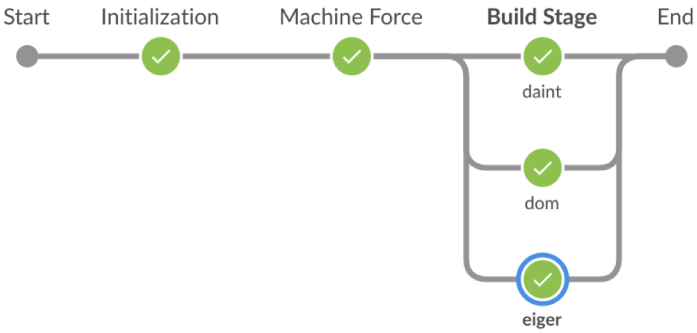

Build Stage / eiger - 15m 43s

Check out from version control	3s
echo $PWD — Shell Script	<1s
#!/bin/bash -l echo $SCRATCH — Shell Script	<1s
#!/bin/bash -l echo $XDG_RUNTIME_DIR/build — Shell Script	<1s
#!/bin/bash -l echo /apps/eiger/UES/jenkins/1.3.2/20.10 — Shell Script	<1s
Shell Script	6m 52s

TestingEB Pipeline

Github Pull Request

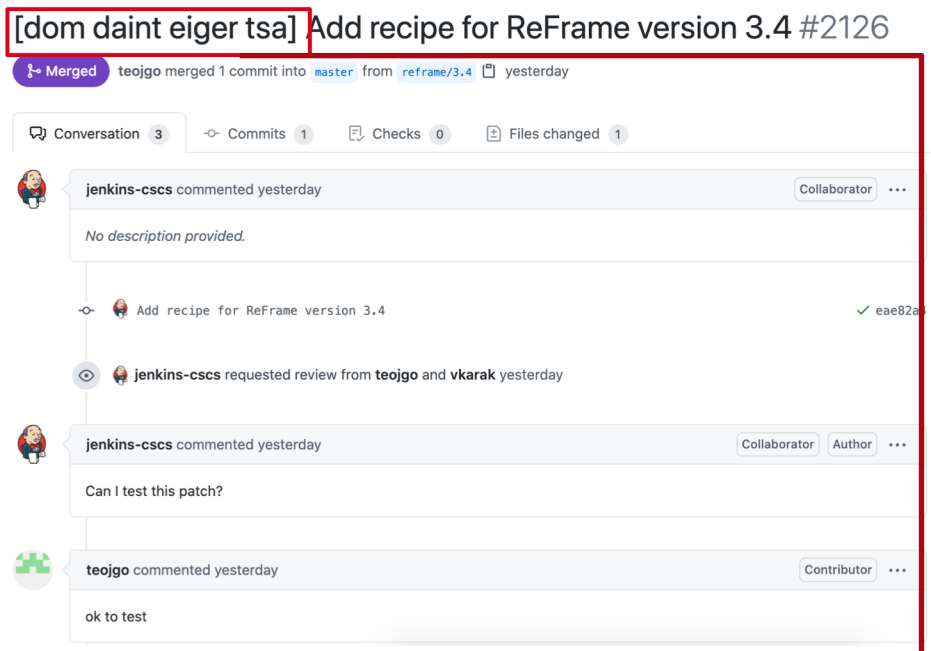

Pipeline triggered on Jenkins

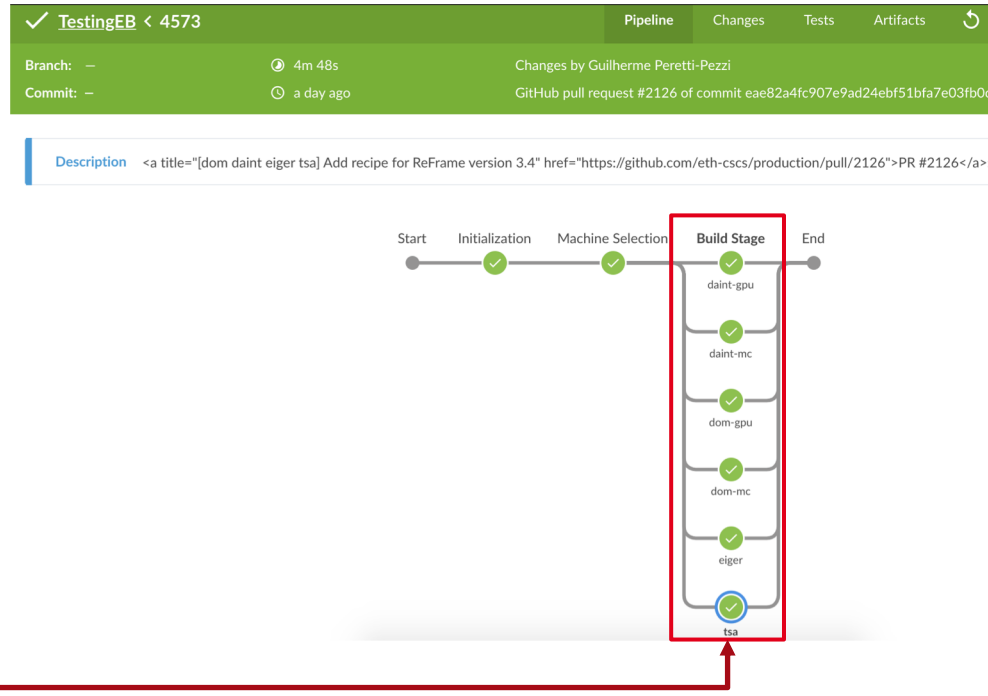

ReBuildEB Pipeline

- ReBuilds easyconfigs in parallel downloading from scratch
- Atlassian Jira integration: automated issue creation for failure

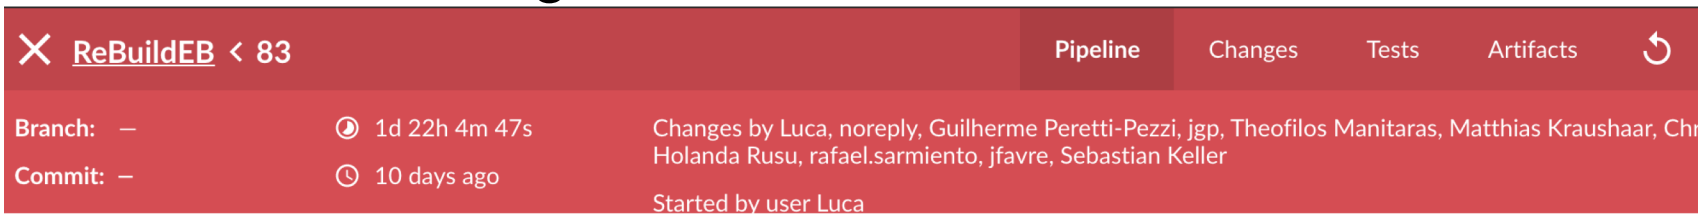

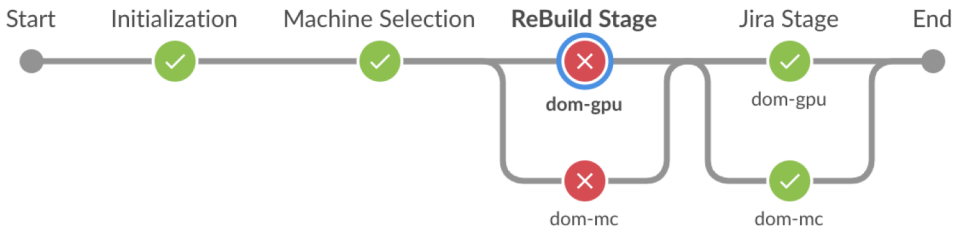

UpdateEB Pipeline

- Automated GitHub PR for successful updates (GitHub stage)

UpdateEB < 49 >

Pipeline | Changes | Tests | Artifacts

Branch: –
Commit: –

19h 15m 24s
11 days ago

No changes
Started by user Luca

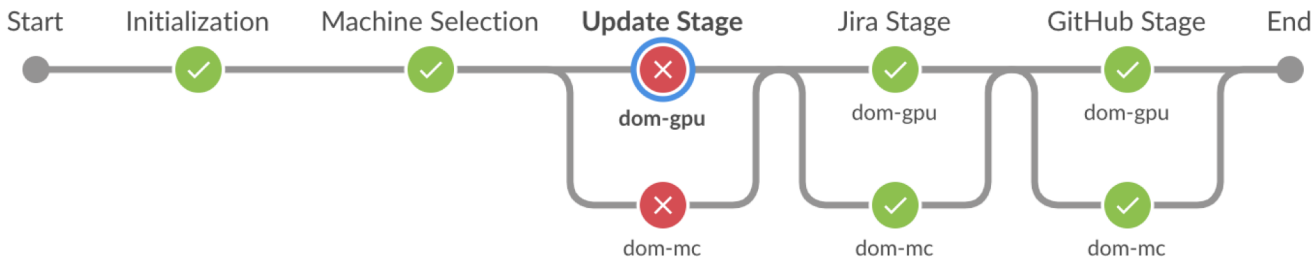

Update Stage / dom-gpu - 18h 48m 23s

- ✓ Check out from version control
- ✓ List of unuse paths: /apps/dom/UES/jenkins/7.0.UP02/gpu/easybuild/tools/modules/all /apps/dom/UES/jenkins/7.0.UP02/gpu/easybuild/modules/all
- ✗ Shell Script

Outline

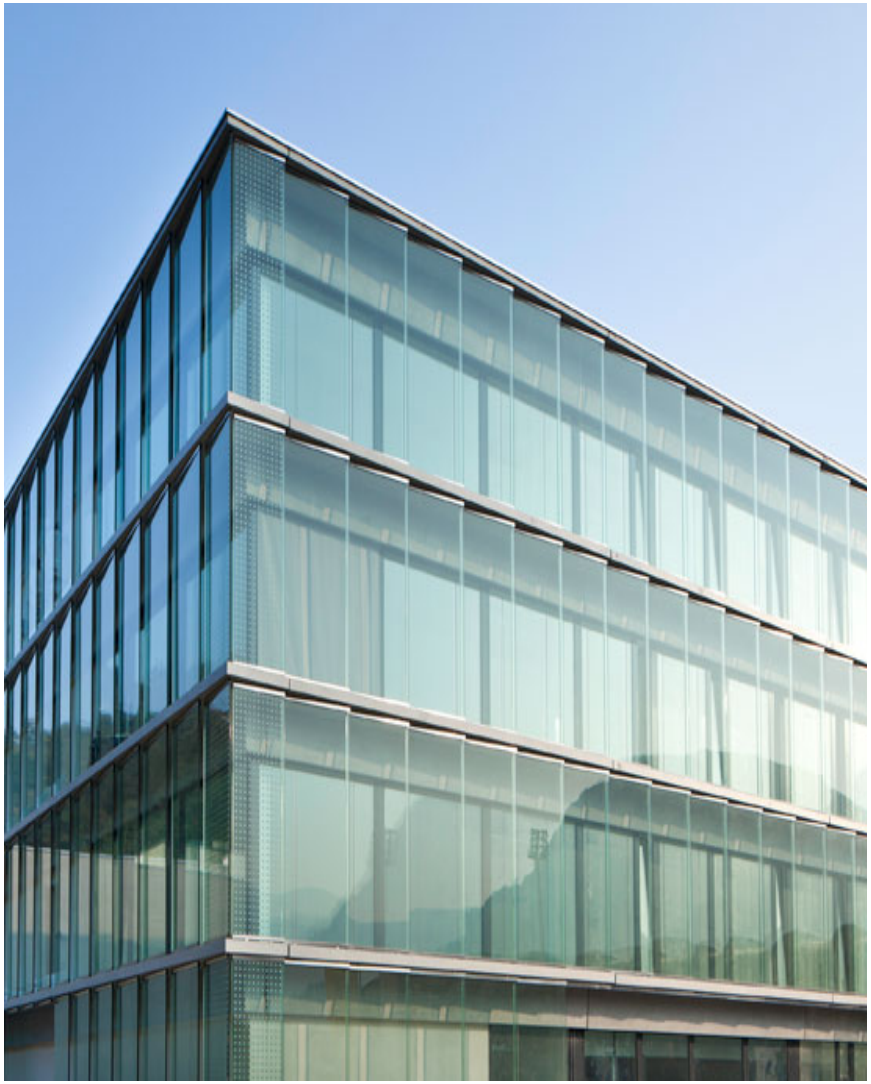

- EasyBuild timeline @ CSCS
- Overview of CSCS HPC systems
  - MeteoSwiss systems
  - Piz Daint
  - Alps
- EasyBuild for CSCS Users
- Easybuild with Jenkins
  - CSCS EB production repository
  - Jenkins pipelines
    - Production EB
    - ReBuild EB
    - Testing EB
    - Update EB
- Final Remarks

Final Remarks

- Moving HPC software stack to EB takes time
  - Learning curve
  - Resistance to change

- EasyConfigs vs. EasyBlocks
  - EasyBlocks
    - (+) Reusable: Great for well packaged & stable software
    - (-) Too much overhead for bleeding edge software
    - (-) Reproducibility: how to keep track of changes?
  - EasyConfigs
    - (-) Reuse by copy/paste (=> duplication)
    - **(+++) Self contained recipes**

Final Remarks

- Software Stack deployed on new Cray EX system
  - Custom toolchain cpeCCE and cpeGNU
  - Toolchains pinning Cray PE modules versions

- Automated Easyconfig updates for new Cray PEs
  - **New toolchain version** created for OS and Cray PE updates
  - Easyconfigs in production updated with **--try-toolchain**
  - **Integration** with Atlassian **Jira** for automated issue creation
  - Automated **GitHub PR** submitted for successful updates

- Work in progress
  - Automated update of software versions and dependencies
  - Replacement of current ProductionEB with UpdateEB pipeline
  - EasyBuild Jenkins projects with additional ReFrame step for check

Useful links for EasyBuild @ CSCS

- EasyBuild User Documentation at CSCS
  - https://user.cscs.ch/computing/compilation/easybuild

- CSCS EasyBuild repositories
  - List of production builds performed by Jenkins
    - https://github.com/eth-cscs/production/tree/master/jenkins-builds
  - Custom easyconfigs:
    - https://github.com/eth-cscs/production/tree/master/easybuild/easyconfigs
  - Custom easyblocks:
    - https://github.com/eth-cscs/production/tree/master/easybuild/easyblocks
  - CSCS repo mirrored under the EasyBuilders GitHub repository:
    - https://github.com/easybuilders/CSCS

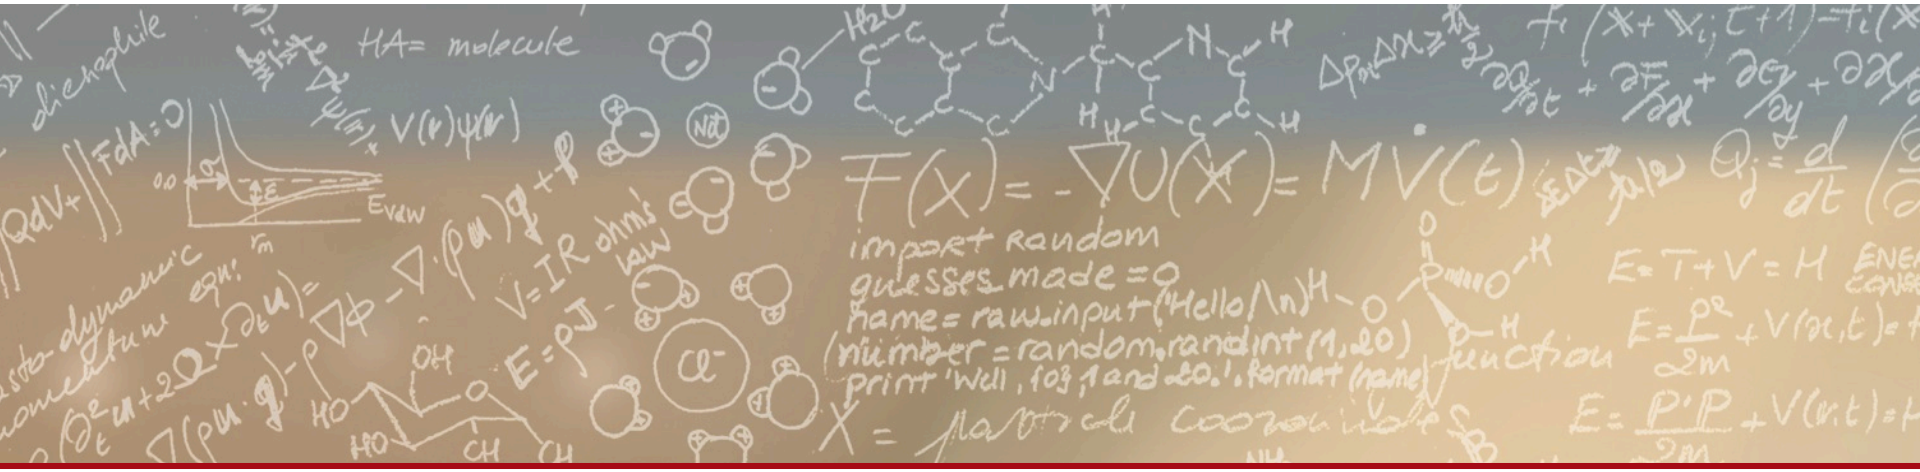

Thank you for your kind attention