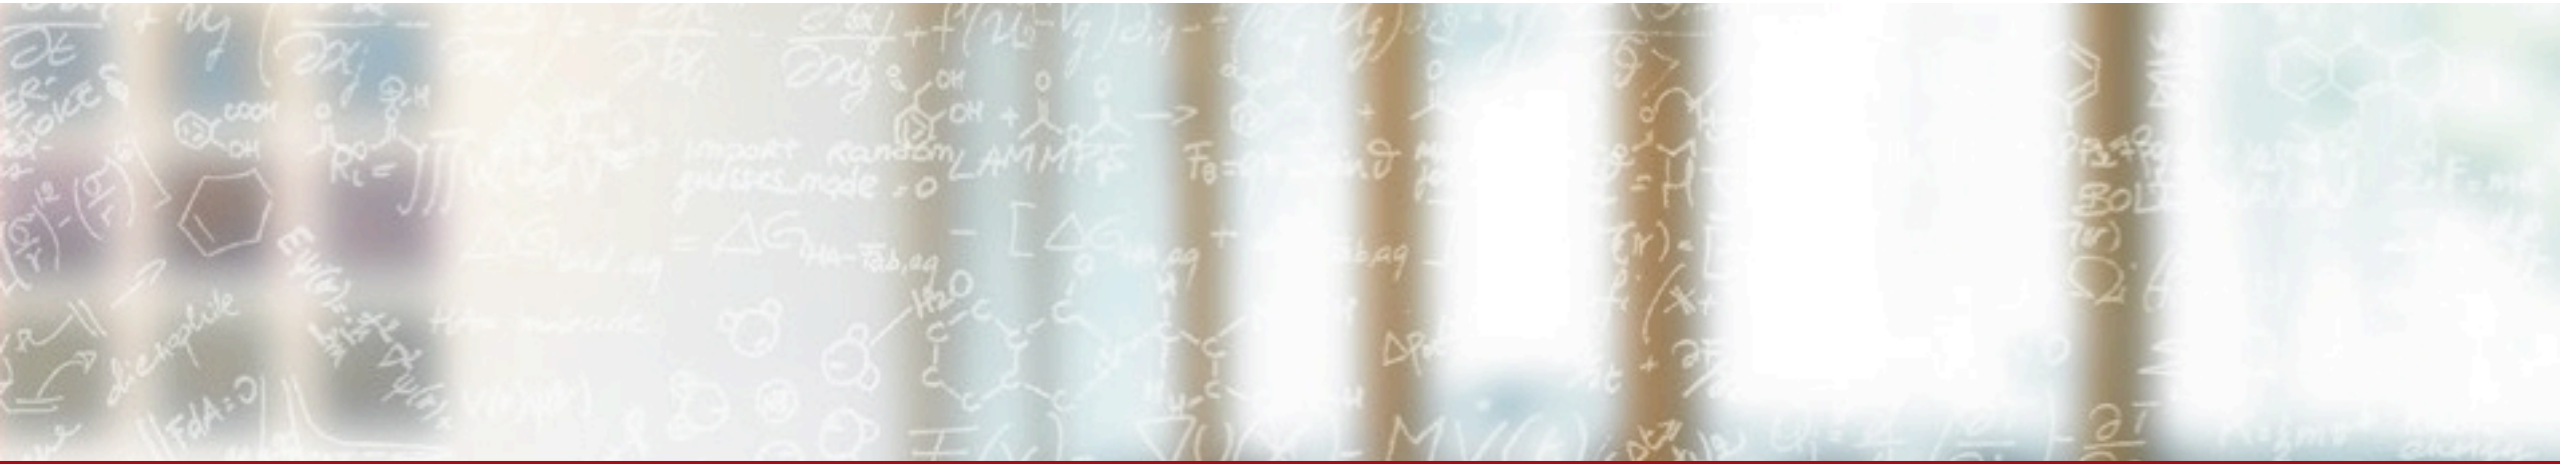




CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich



ReFrame: A Regression Testing and Continuous Integration Framework for HPC systems

Fifth Annual Workshop on HPC User Support Tools – SC18

Vasileios Karakasis, CSCS

November 11, 2018

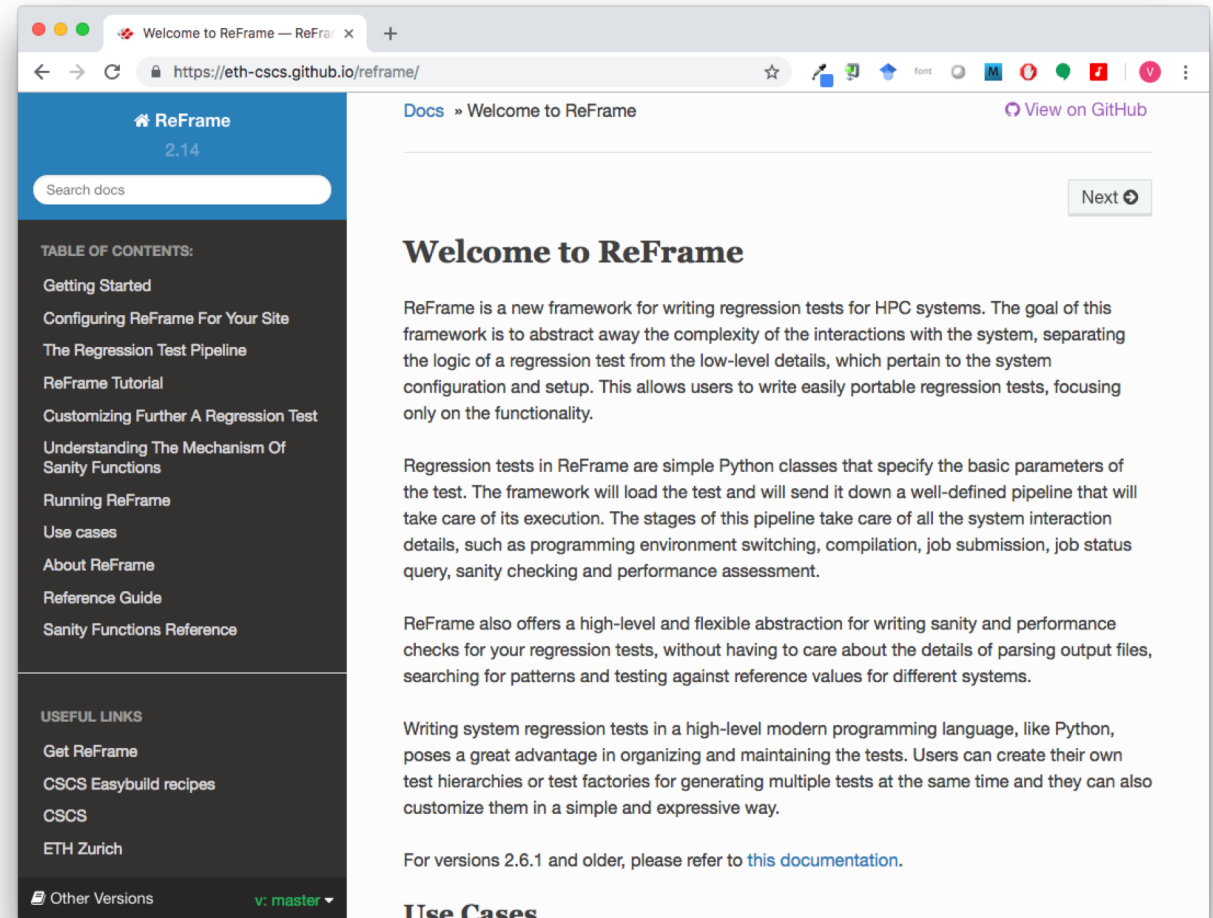
Table of Contents

1. Overview of ReFrame
2. Use Cases
3. Demo
4. Q & A

What is ReFrame?

A new regression testing framework that

- allows writing **portable HPC** regression tests in Python,
- **abstracts away** the system interaction details,
- lets users focus solely on the **logic** of their test.



The screenshot shows a web browser displaying the ReFrame documentation page. The page title is "Welcome to ReFrame" and the version is "2.14". The page content includes a "TABLE OF CONTENTS" section with links to "Getting Started", "Configuring ReFrame For Your Site", "The Regression Test Pipeline", "ReFrame Tutorial", "Customizing Further A Regression Test", "Understanding The Mechanism Of Sanity Functions", "Running ReFrame", "Use cases", "About ReFrame", "Reference Guide", and "Sanity Functions Reference". There is also a "USEFUL LINKS" section with links to "Get ReFrame", "CSCS Easybuild recipes", "CSCS", and "ETH Zurich". The main content area starts with "Welcome to ReFrame" and describes the framework's goal: "ReFrame is a new framework for writing regression tests for HPC systems. The goal of this framework is to abstract away the complexity of the interactions with the system, separating the logic of a regression test from the low-level details, which pertain to the system configuration and setup. This allows users to write easily portable regression tests, focusing only on the functionality." It then explains that regression tests are simple Python classes and that the framework handles execution details like environment switching and job submission. It also mentions that ReFrame offers a high-level abstraction for writing sanity and performance checks. The page ends with a note for older versions and a "Use Cases" section header.

<https://github.com/eth-cscs/reframe>

Design Goals

- Productivity
- Portability
- Speed and Ease of Use
- Robustness

Write once, test everywhere!

Key Features

- Separation of system and prog. environment configuration from test's logic
- Support for cycling through prog. environments and system partitions
- Regression tests written in Python
 - Easy customization of tests
 - Flexibility in organizing the tests
- Support for sanity and performance tests
 - Allows complex and custom analysis of the output through an embedded mini-language for sanity and performance checking.
- Progress and result reports
- Performance logging with support for Graylog
- Clean internal APIs that allow the easy extension of the framework's functionality

More Features


- Multiple workload manager backends
 - SLURM
 - PBS/Torque
- Multiple parallel launcher backends
 - srun, mpirun, mpiexec etc.
- Multiple environment modules backends
 - Tmod, Tmod4, Lmod
- Build system backends
 - CMake, Autotools, Make
- Asynchronous execution of regression tests
- Complete documentation (tutorials, reference guide)
- ... and more (<https://github.com/eth-cscs/reframe>)

ReFrame's architecture

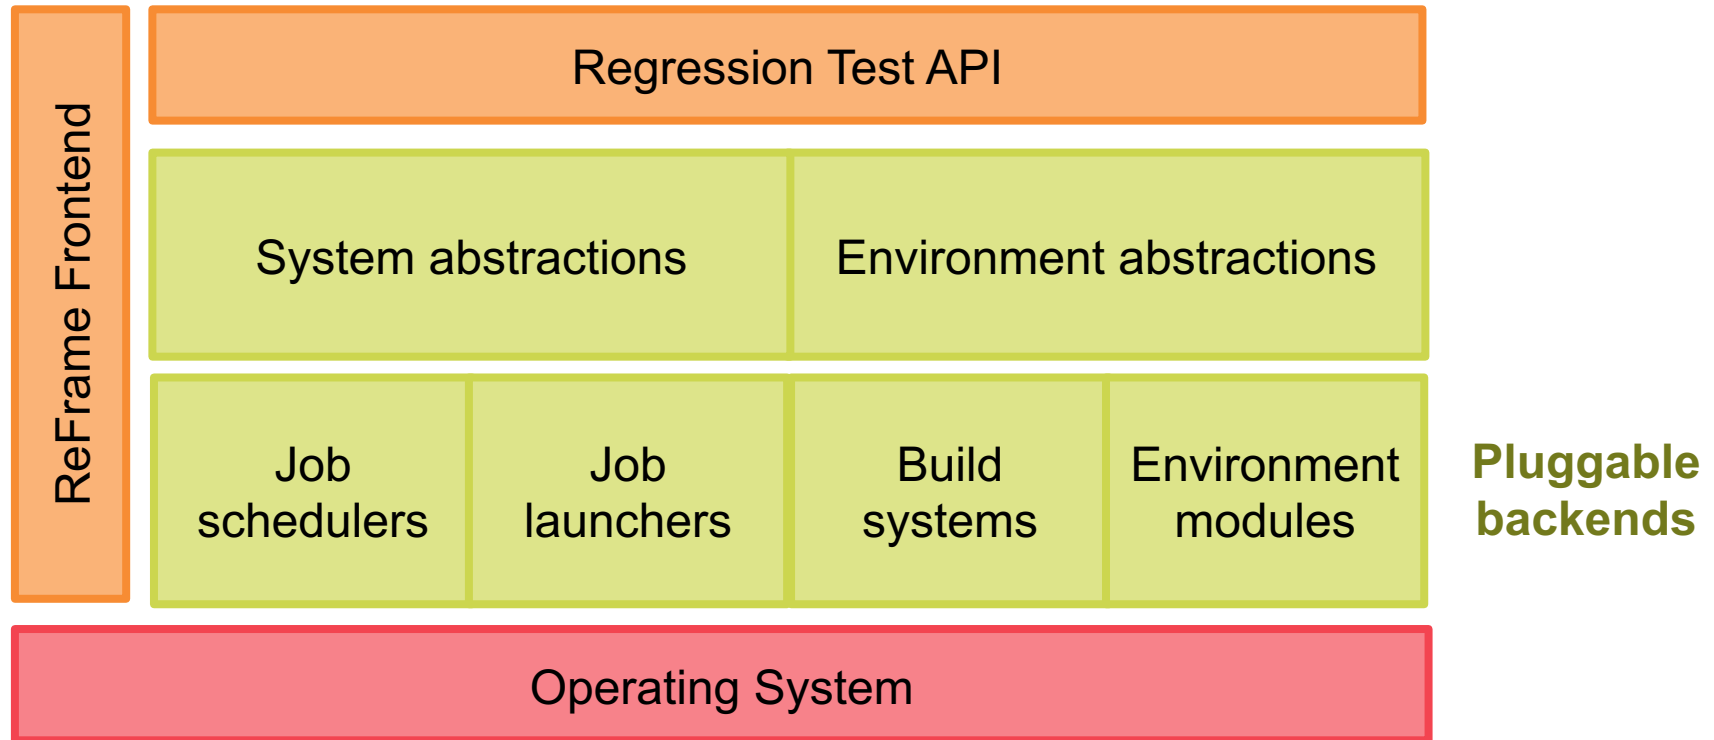


Developer of regression tests

```
@rfm.simple_test  
class MyTest(rfm.RegressionTest):
```



```
reframe -r
```



Writing a Regression Test in ReFrame

A regression test writer should not care about...

- How to access system partitions and if there are any.
- How (programming) environments are switched.
- How the test's environment is actually set up.
- How a job script is generated and if it's needed at all.
- How a sanity/performance pattern is looked up in the output.

ReFrame allows you to focus on the logic of your test.

Writing a Regression Test in ReFrame

ReFrame tests are specially decorated classes

Valid systems and prog. environments

Compile and run setup

Sanity checking

Extract performance values from output

Reference values and performance thresholds


Tags for easy lookup

```
import reframe as rfm
import reframe.utility.sanity as sn

@rfm.simple_test
class Example7Test(rfm.RegressionTest):
    def __init__(self):
        super().__init__()
        self.descr = 'Matrix-vector multiplication (CUDA performance test)'
        self.valid_systems = ['daint:gpu']
        self.valid_prog_environs = ['PrgEnv-gnu', 'PrgEnv-cray', 'PrgEnv-pgi']
        self.sourcepath = 'example_matrix_vector_multiplication_cuda.cu'
        self.build_system = 'SingleSource'
        self.build_system.cxxflags = ['-O3']
        self.executable_opts = ['4096', '1000']
        self.modules = ['cudatoolkit']
        self.num_gpus_per_node = 1
        self.sanity_patterns = sn.assert_found(
            r'time for single matrix vector multiplication', self.stdout)
        self.perf_patterns = {
            'perf': sn.extractsingle(r'Performance:\s+(?P<Gflops>\S+) Gflop/s',
                                    self.stdout, 'Gflops', float)
        }
        self.reference = {
            'daint:gpu': {
                'perf': (50.0, -0.1, 0.1),
            }
        }
        self.maintainers = ['you-can-type-your-email-here']
        self.tags = {'tutorial'}
```

The Regression Test Pipeline / How ReFrame Executes Tests

A series of well defined phases that each regression test goes through



The Regression Test Pipeline / How ReFrame Executes Tests

- Tests may skip some pipeline stages
 - Compile-only tests
 - Run-only tests
- Users may define additional actions before or after every pipeline stage by overriding the corresponding methods of the regression test API.
 - E.g., override the setup stage for customizing the behavior of the test per programming environment and/or system partition.
- Frontend passes through three phases and drives the execution of the tests
 1. Regression test discovery and loading
 2. Regression test selection (by name, tag, prog. environment support etc.)
 3. Regression test listing or execution

Running ReFrame

```
reframe -C /path/to/config.py -c /path/to/checks -r
```

- ReFrame uses three directories when running:
 1. **Stage directory**: Stores temporarily all the resources (static and generated) of the tests
 - Source code, input files, generated build script, generated job script, output etc.
 - This directory is removed if the test finishes successfully.
 2. **Output directory**: Keeps important files from the run for later reference
 - Job and build scripts, outputs and any user-specified files.
 3. **Performance log directory**: Keeps performance logs for the performance tests
- ReFrame generates a summary report at the end with detailed failure information.

Running ReFrame (sample output)

```
[=====] Running 1 check(s)
[=====] Started on Fri Sep  7 15:32:50 2018

[-----] started processing Example7Test (Matrix-vector multiplication using CUDA)
[ RUN    ] Example7Test on daint:gpu using PrgEnv-cray
[      OK ] Example7Test on daint:gpu using PrgEnv-cray
[ RUN    ] Example7Test on daint:gpu using PrgEnv-gnu
[      OK ] Example7Test on daint:gpu using PrgEnv-gnu
[ RUN    ] Example7Test on daint:gpu using PrgEnv-pgi
[      OK ] Example7Test on daint:gpu using PrgEnv-pgi
[-----] finished processing Example7Test (Matrix-vector multiplication using CUDA)

[ PASSED ] Ran 3 test case(s) from 1 check(s) (0 failure(s))
[=====] Finished on Fri Sep  7 15:33:42 2018
```

Running ReFrame (sample failure)

```
[=====] Running 1 check(s)
[=====] Started on Fri Sep  7 16:40:12 2018

[-----] started processing Example7Test (Matrix-vector multiplication using CUDA)
[ RUN    ] Example7Test on daint:gpu using PrgEnv-gnu
[ FAIL   ] Example7Test on daint:gpu using PrgEnv-gnu
[-----] finished processing Example7Test (Matrix-vector multiplication using CUDA)

[ FAILED ] Ran 1 test case(s) from 1 check(s) (1 failure(s))
[=====] Finished on Fri Sep  7 16:40:22 2018
```

```
=====
SUMMARY OF FAILURES
```

```
-----
FAILURE INFO for Example7Test
```

- * System partition: daint:gpu
- * Environment: PrgEnv-gnu
- * Stage directory: /path/to/stage/daint/gpu/PrgEnv-gnu/Example7Test
- * Job type: batch job (id=823427)
- * Maintainers: ['you-can-type-your-email-here']
- * Failing phase: performance
- * Reason: sanity error: 50.363125 is beyond reference value 70.0 (l=63.0, u=77.0)

Running ReFrame (examining a failure)


- ReFrame executes each test case from a separate stage directory:
 - `/path/to/stage/<system>/<partition>/<testname>/<environ>`
- Auto-generated build script and compilation's standard output/error
 - `rfm_<testname>_build.sh`
 - `rfm_<testname>_build.out`
 - `rfm_<testname>_build.err`
- Auto-generated job script and execution's standard output/error
 - `rfm_<testname>_job.sh`
 - `rfm_<testname>_job.out`
 - `rfm_<testname>_job.err`

Running ReFrame (examining performance logs)

- `/path/to/reframe/prefix/perflogs/<testname>.log`
 - A single file named after the test's name is updated every time the test is run
 - Log record output is fully configurable

```
2018-09-07T15:32:59|reframe 2.14-dev2|Example7Test on daint:gpu using PrgEnv-cray|jobid=823394|perf=49.71432|ref=50.0 (l=-0.1, u=0.1)
2018-09-07T15:33:11|reframe 2.14-dev2|Example7Test on daint:gpu using PrgEnv-gnul|jobid=823395|perf=50.1609|ref=50.0 (l=-0.1, u=0.1)
2018-09-07T15:33:42|reframe 2.14-dev2|Example7Test on daint:gpu using PrgEnv-pgi|jobid=823396|perf=51.078648|ref=50.0 (l=-0.1, u=0.1)
2018-09-07T16:40:22|reframe 2.14-dev2|Example7Test on daint:gpu using PrgEnv-gnul|jobid=823427|perf=50.363125|ref=70.0 (l=-0.1, u=0.1)
```

- ReFrame can also send logs to a Graylog server, where you can plot them with web tools.





CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

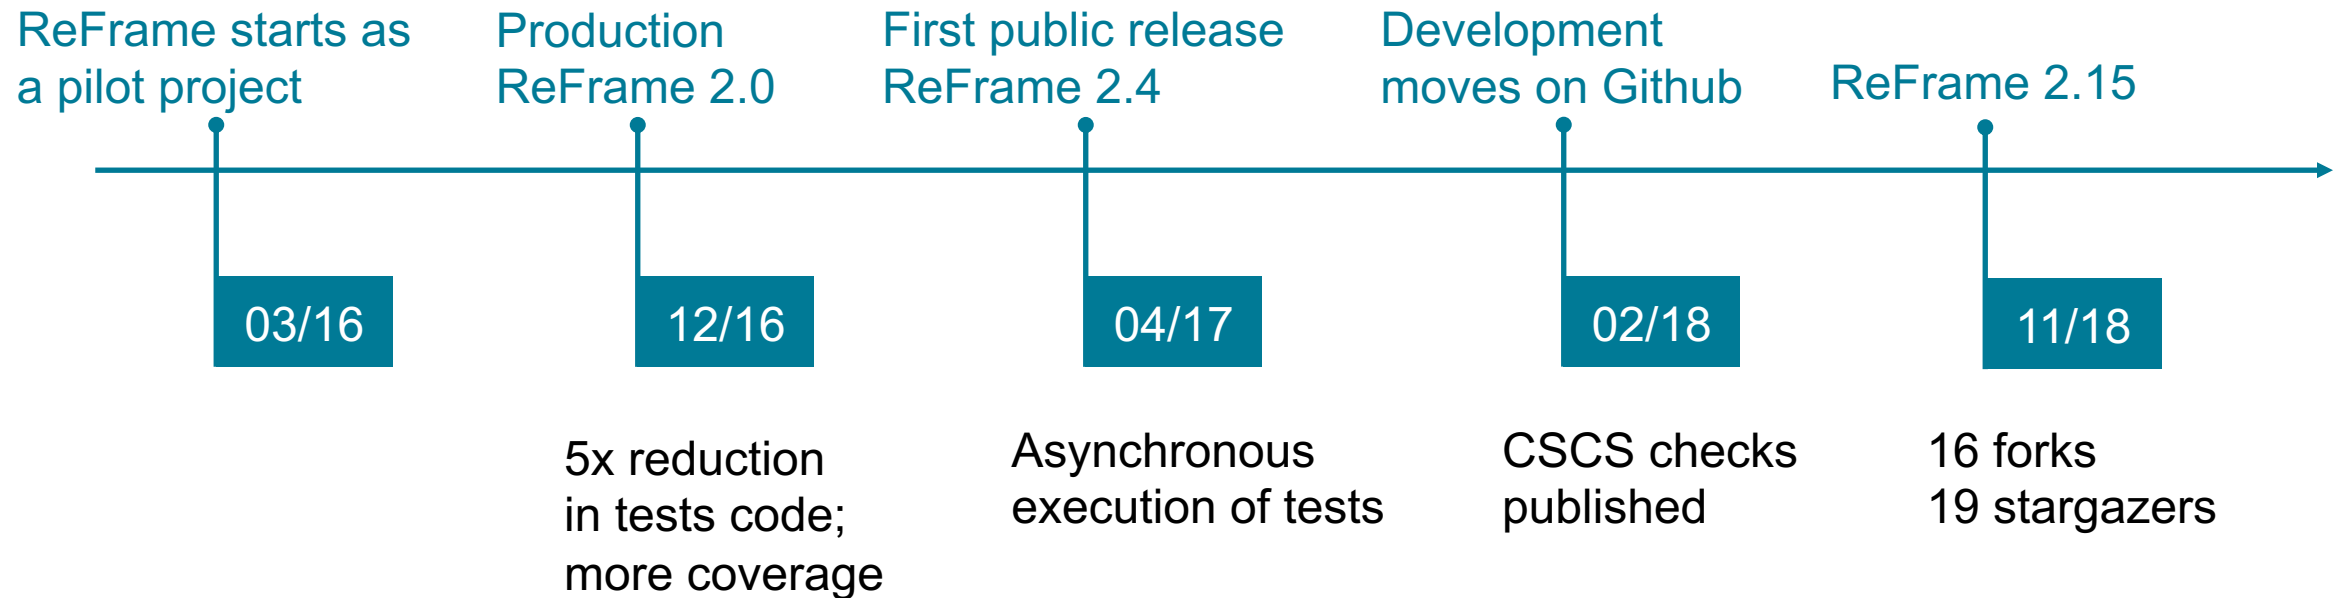
ETH zürich

Using ReFrame at CSCS

Background

- CSCS had a shell-script based regression suite
 - Tests very tightly coupled to system details
 - Lots of code replication across tests
 - 15K lines of test code
- Simple changes required significant team effort
 - Porting all tests to native Slurm took several weeks
- Fixing even simple bugs was a tedious task
 - Tens of regression test files had to be fixed

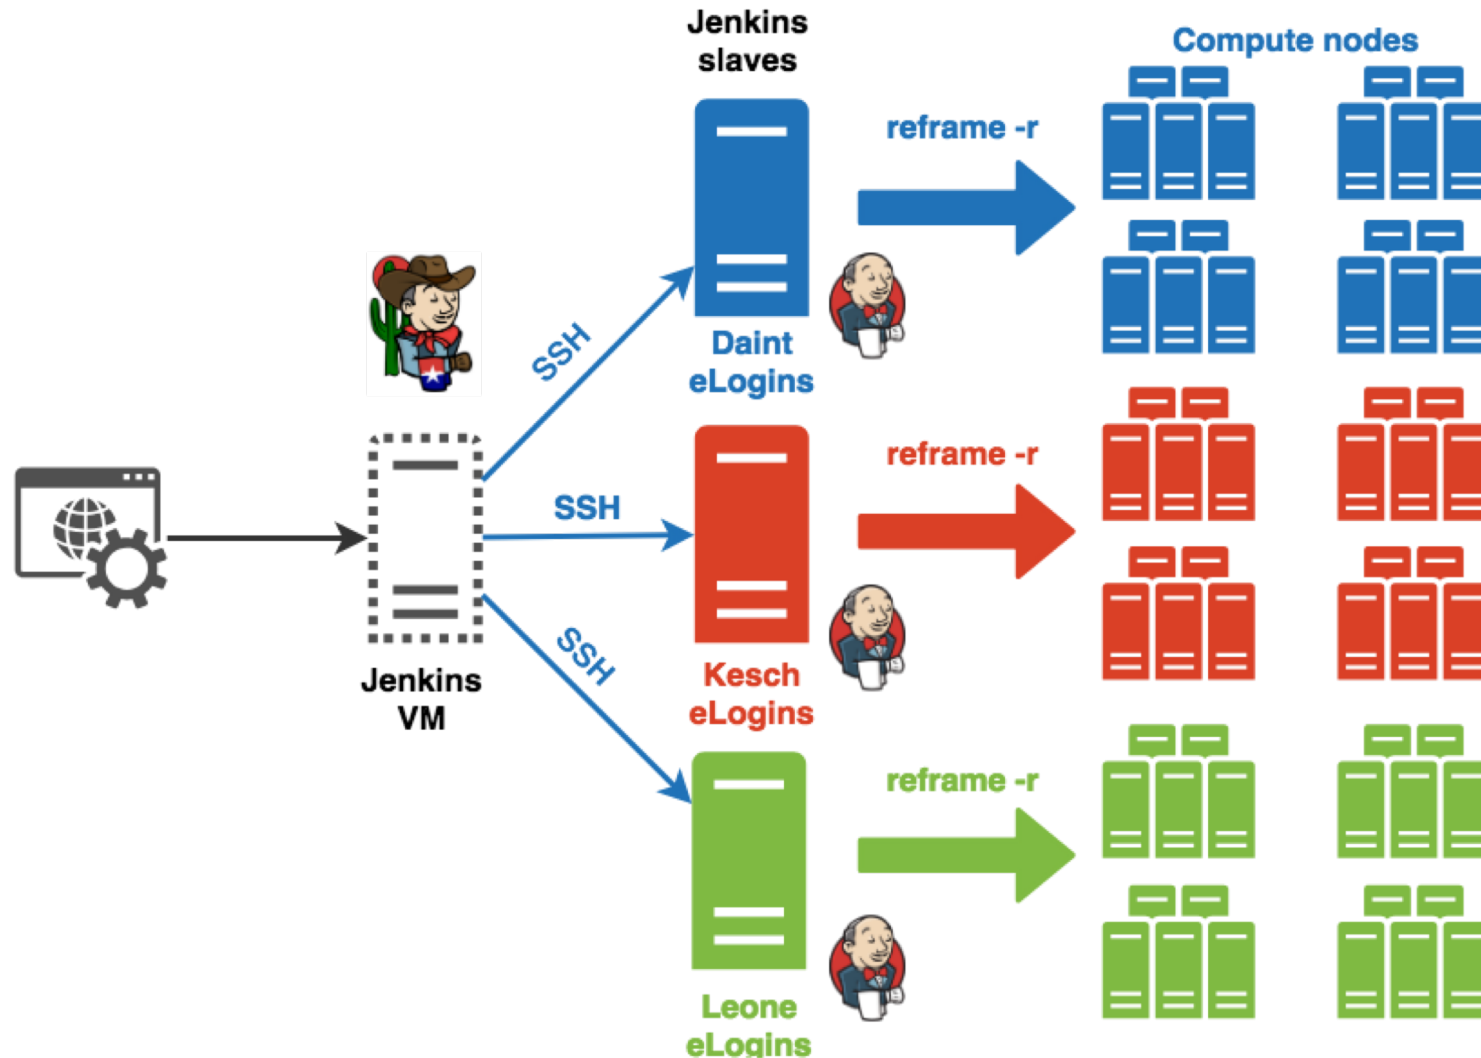
Timeline / ReFrame Evolution



ReFrame @ CSCS / Tests

- Used for continuously testing systems in production
 - Piz Daint: 179 tests
 - Piz Kesch: 75 tests
 - Leone: 45 tests
 - **Total: 241 different tests (reused across systems)**
- Three categories of tests
 1. Production (90min)
 - Applications, libraries, programming environments, profiling tools, debuggers, microbenchmarks
 - Sanity and performance
 - Run nightly by Jenkins
 2. Maintenance (10min)
 - Programming environment sanity and key user applications performance
 - Before/after maintenance sessions
 3. Diagnostics

ReFrame @ CSCS / Production set-up



ReFrame @ CSCS / Production set-up

The screenshot displays the Jenkins CI/CD pipeline for ReFrame at CSCS. It consists of three overlapping browser windows showing different views of the pipeline.

Left Window: Jenkins Dashboard

- URL: <https://jenkins.cscs.ch/view/Reframe%20Regression>
- Project: reframe-kesch-production-daily
- Build History Table:

Build #	Time	Status
#164	Nov 8, 2018 12:43 AM	Success
#163	Nov 7, 2018 12:43 AM	Failure
#162	Nov 6, 2018 12:43 AM	Success
#161	Nov 5, 2018 12:43 AM	Success
#160	Nov 4, 2018 12:43 AM	Success
#159	Nov 3, 2018 12:43 AM	Success
#158	Nov 2, 2018 12:43 AM	Success
#157	Nov 1, 2018 12:43 AM	Success
#156	Oct 31, 2018 12:43 AM	Success

Middle Window: Pipeline Configuration

- URL: <https://jenkins.cscs.ch/blue/organizations/jenkins/reframe-kesch-production-daily>
- Branch: — (43m 38s)
- Commit: — (3 days ago)
- Steps:

- ✓ > Check out from version control
- ✓ > Check out from version control
- ✓ > hello — Print Message
- ✓ > Shell Script
- ✓ > reframe.log — Archive the artifacts
- ✓ > exit 0 — Shell Script

Right Window: Pipeline Execution Log

- URL: <https://jenkins.cscs.ch/blue/organizations/jenkins/reframe-kesch-production-daily>
- Log content (lines 972-1018):

```
972 [-----] started processing StreamTest (STREAM Benchmark)
973 [ RUN ] StreamTest on kesch:pn using PrgEnv-cray
974 [ RUN ] StreamTest on kesch:pn using PrgEnv-cray-nompi
975 [ HOLD ] StreamTest on kesch:pn using PrgEnv-cray-nompi
976 [ RUN ] StreamTest on kesch:pn using PrgEnv-pgi
977 [ HOLD ] StreamTest on kesch:pn using PrgEnv-pgi
978 [ RUN ] StreamTest on kesch:pn using PrgEnv-pgi-nompi
979 [ HOLD ] StreamTest on kesch:pn using PrgEnv-pgi-nompi
980 [ RUN ] StreamTest on kesch:pn using PrgEnv-gnu
981 [ HOLD ] StreamTest on kesch:pn using PrgEnv-gnu
982 [ RUN ] StreamTest on kesch:pn using PrgEnv-gnu-nompi
983 [ HOLD ] StreamTest on kesch:pn using PrgEnv-gnu-nompi
984 [ RUN ] StreamTest on kesch:cn using PrgEnv-cray
985 [ RUN ] StreamTest on kesch:cn using PrgEnv-cray-nompi
986 [ HOLD ] StreamTest on kesch:cn using PrgEnv-cray-nompi
987 [ RUN ] StreamTest on kesch:cn using PrgEnv-pgi
988 [ HOLD ] StreamTest on kesch:cn using PrgEnv-pgi
989 [ RUN ] StreamTest on kesch:cn using PrgEnv-pgi-nompi
990 [ HOLD ] StreamTest on kesch:cn using PrgEnv-pgi-nompi
991 [ RUN ] StreamTest on kesch:cn using PrgEnv-gnu
992 [ HOLD ] StreamTest on kesch:cn using PrgEnv-gnu
993 [ RUN ] StreamTest on kesch:cn using PrgEnv-gnu-nompi
994 [ HOLD ] StreamTest on kesch:cn using PrgEnv-gnu-nompi
995 [-----] finished processing StreamTest (STREAM Benchmark)
996
997 [-----] waiting for spawned checks to finish
998 [ FAIL ] StreamTest on kesch:pn using PrgEnv-cray
999 [ OK ] StreamTest on kesch:cn using PrgEnv-cray
1000 [ FAIL ] StreamTest on kesch:pn using PrgEnv-gnu-nompi
1001 [ OK ] StreamTest on kesch:cn using PrgEnv-gnu-nompi
1002 [ FAIL ] StreamTest on kesch:pn using PrgEnv-gnu
1003 [ OK ] StreamTest on kesch:cn using PrgEnv-gnu
1004 [ FAIL ] StreamTest on kesch:pn using PrgEnv-pgi-nompi
1005 [ OK ] StreamTest on kesch:cn using PrgEnv-pgi-nompi
1006 [ FAIL ] StreamTest on kesch:pn using PrgEnv-pgi
1007 [ OK ] StreamTest on kesch:cn using PrgEnv-pgi
1008 [ FAIL ] StreamTest on kesch:pn using PrgEnv-cray-nompi
1009 [ OK ] StreamTest on kesch:cn using PrgEnv-cray-nompi
1010 [-----] all spawned checks have finished
1011
1012 [ FAILED ] Ran 212 test case(s) from 73 check(s) (6 failure(s))
1013 [=====] Finished on Wed Nov 7 01:31:20 2018
1014
1015 SUMMARY OF RETRIES
1016
1017 * Test HelloWorldTestOpenMP_c_dynamic on kesch:cn using PrgEnv-cray was retried 1 time(s) and passed.
1018 * Test HelloWorldTestOpenMP_c_dynamic on kesch:cn using PrgEnv-cray-nompi was retried 1 time(s) and passed.
```

Conclusions and Future Directions

ReFrame is a powerful tool that allows you to continuously test an HPC environment without having to deal with the low-level system interaction details.

- High-level tests written in Python
 - Portability across HPC system platforms
 - Comprehensive reports and reproducible methods
-
- ReFrame is being actively developed with a regular release cycle.
 - Future directions
 - Test dependencies
 - Remote and asynchronous compilation of tests
 - Bug reports, feature requests, help @ <https://github.com/eth-cscs/reframe>

Acknowledgements

- Framework contributions
 - Andreas Jocksch
 - Matthias Kraushaar
 - Rafael Sarmiento
 - Samuel Omlin
 - Theofilos Manitaras
 - Victor Holanda
- Regression tests
 - SCS and OPS team



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich


ReFrame Demo ([link](#))



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich



Thank you for your attention.