Red Hat

Faster, Cheaper, Leaner: Horizontally Scaling a CI Pipeline

Yorgos Saslis, Software Delivery Engineer

Michal Cichra, Principal Software Engineer

3scale BY RED HAT®

CI is a production workload

3scale BY RED HAT®

Red Hat

Customers?

YOU!

Yes, you!

Maintain Flow

CI should sustain flow. Not get in its way.

where are your manners?

Red Hat

yorgos
saslis

Community
Automation
OSS

Maintainability

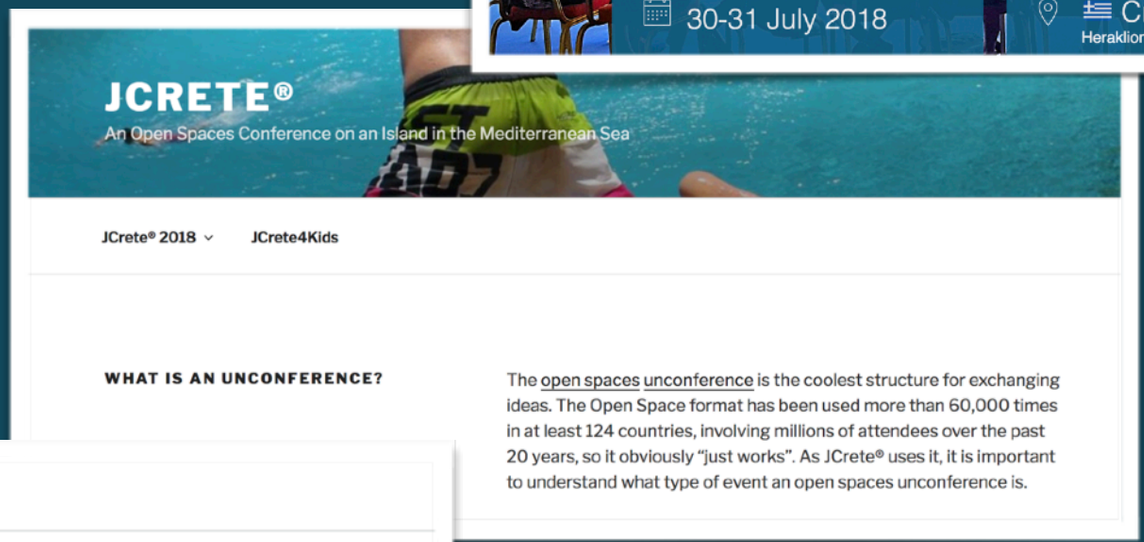

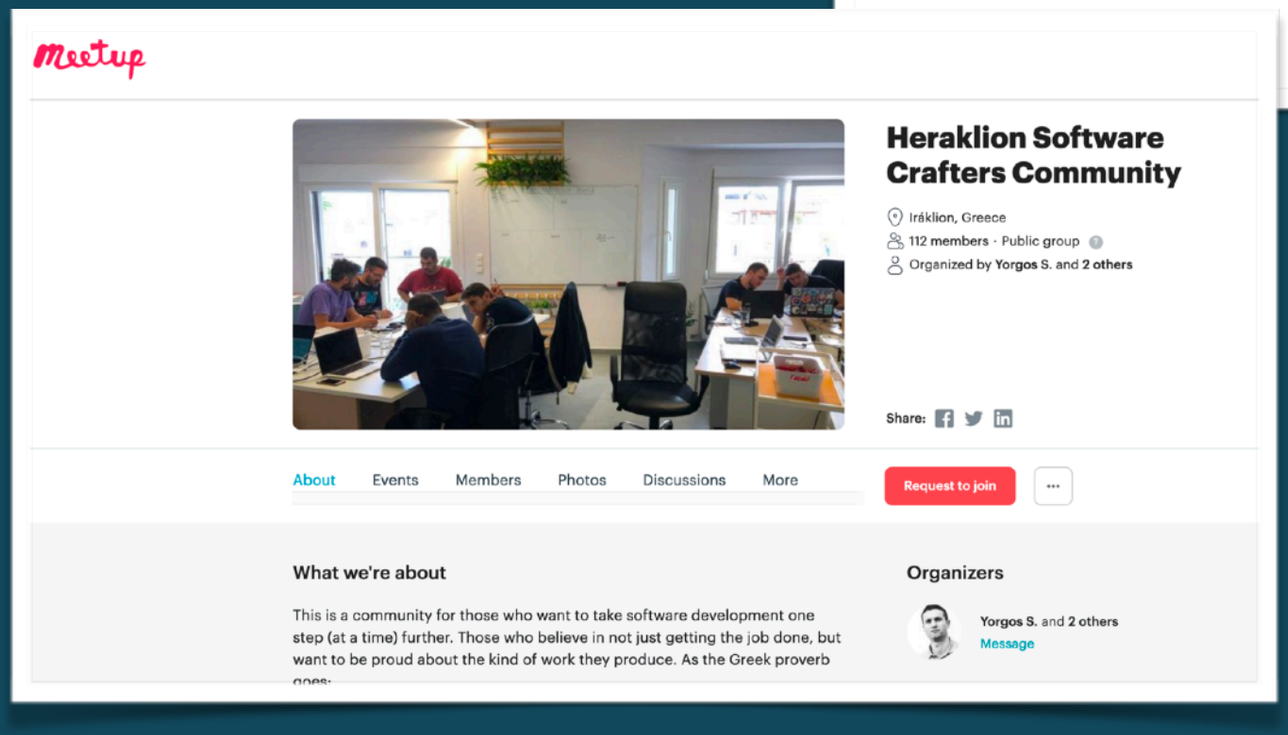

Red Hat

3scale
BY RED HAT®

Open Source API Management

Red Hat

a bit of history...

Red Hat

3scale Timeline

Important milestones

‘07 — 3scale founded

‘16 — 3scale acquired by Red Hat

‘18 — 3scale fully open source!

Red Hat

Open Source projects need CI

All projects need CI. OSS projects need it more!

Hmmm interesting project…

But I just need this extra feature!!

Maybe I can open a pull request…

But how will I know I didn't break anything with my PR ?

Aha!!

There are a bunch of checks on every PR that will protect me!

Contributing can be daunting

- Daunting task
  - especially for new contributors
- CI helps lower the barrier-to-entry

What does CI for a closed source project look like?

Single Jenkins Master

EC2 Cloud plugin for provisioning workers

Jenkins master provisioning automated through Makefiles + terraform

SCM Sync plugin used to persist jenkins configuration “as code”, in a github repository.

Job DSL for jenkins jobs in another github repository.

“HA” not so necessary…

Other Important Figures

Get the whole idea

5 person team

(main component)

10-20

builds per day

2-3

Open PRs per day

Jenkins Worker Nodes

Auto-scaling (both up and down to reduce costs when not used)

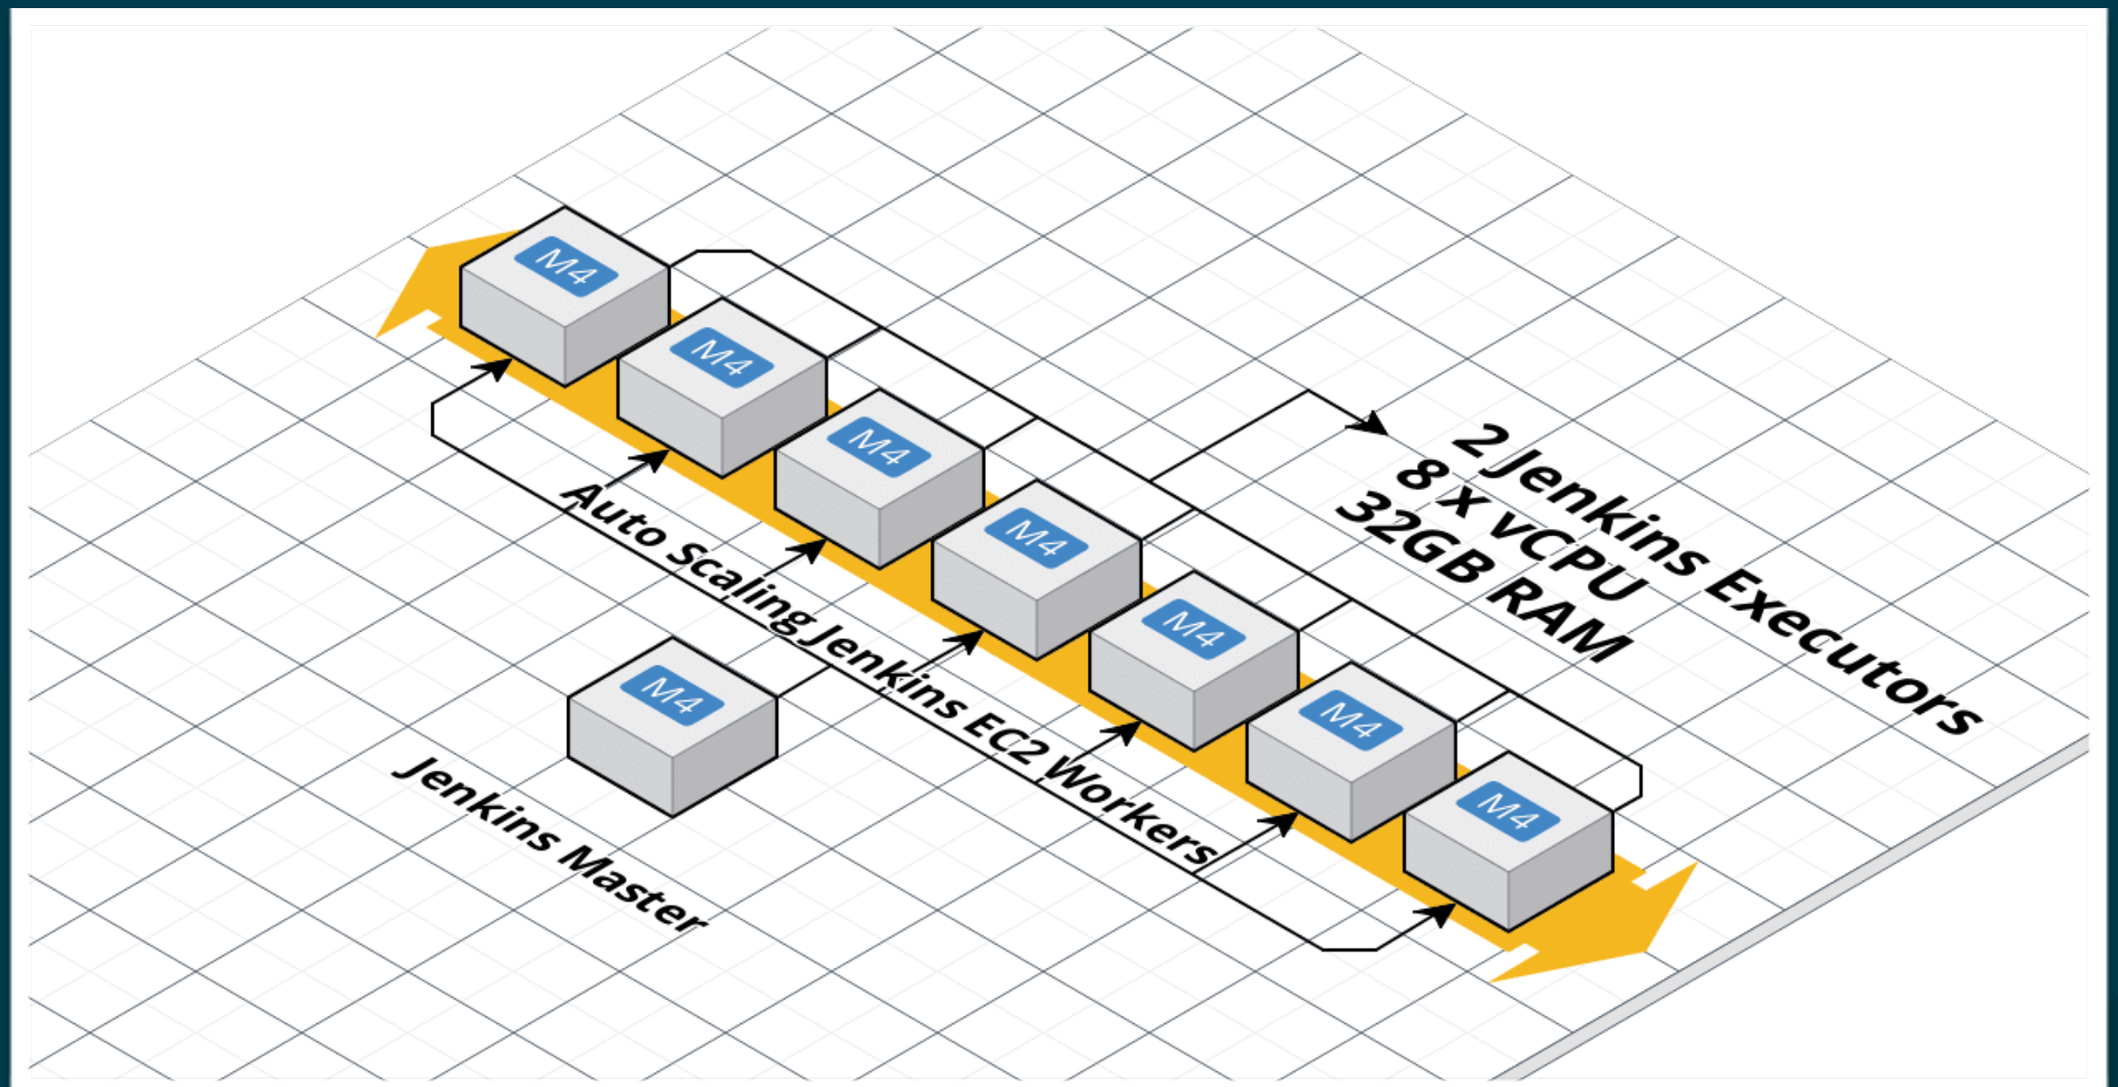

Red Hat

Build Time

For “warm” build

~15 minutes

~11 hours
CPU time

45 vCPUs
90GB RAM

How do we fit 11 hours into… 15 minutes?

Bending Space-Time

NOT

(yet…)

Test suite Parallelization

Homegrown parallelization

15 tasks
manually split

6 executors
for one build

4 languages
to understand
(Groovy, Ruby, Shell, Make)

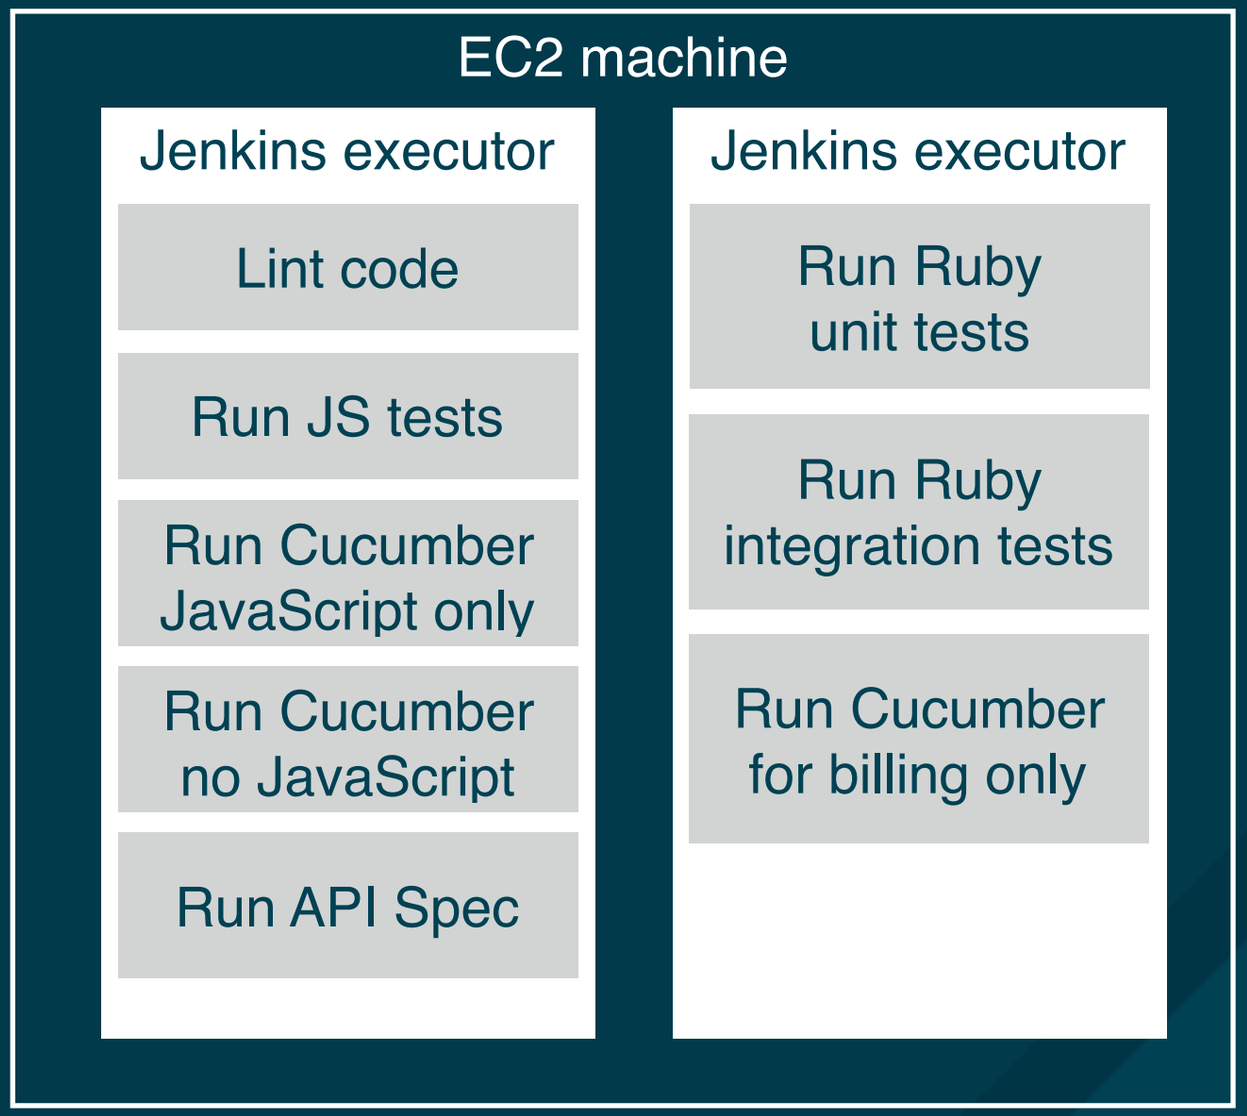

Red Hat

Test Parallelisation

Red Hat

Parallelising Test Execution

Separate test phases

vs.

tests within phase

Separate Test Phases

Usually depending on:

- how long tests take
- what environment they run in (how expensive)

```
@Fast
@Test
void myFastTest() {
    // ...
}
```

Separate Test Phases

“Fast” (seconds)

Commit-phase (5-10min)

Nightlies (hours)

Parts of same test phase, in parallel

Multiple processes responsible for each running some part of the test suite, aggregating results at the end

Group I

Group II

Group III

Group IV

…but how to group, such that they all end at the same time?

One of those rare moments in life when...

WE CAN FIX SOME TECH DEBT!

Red Hat

Problem 1: Build Queues

Almost never empty.

(during working hours)

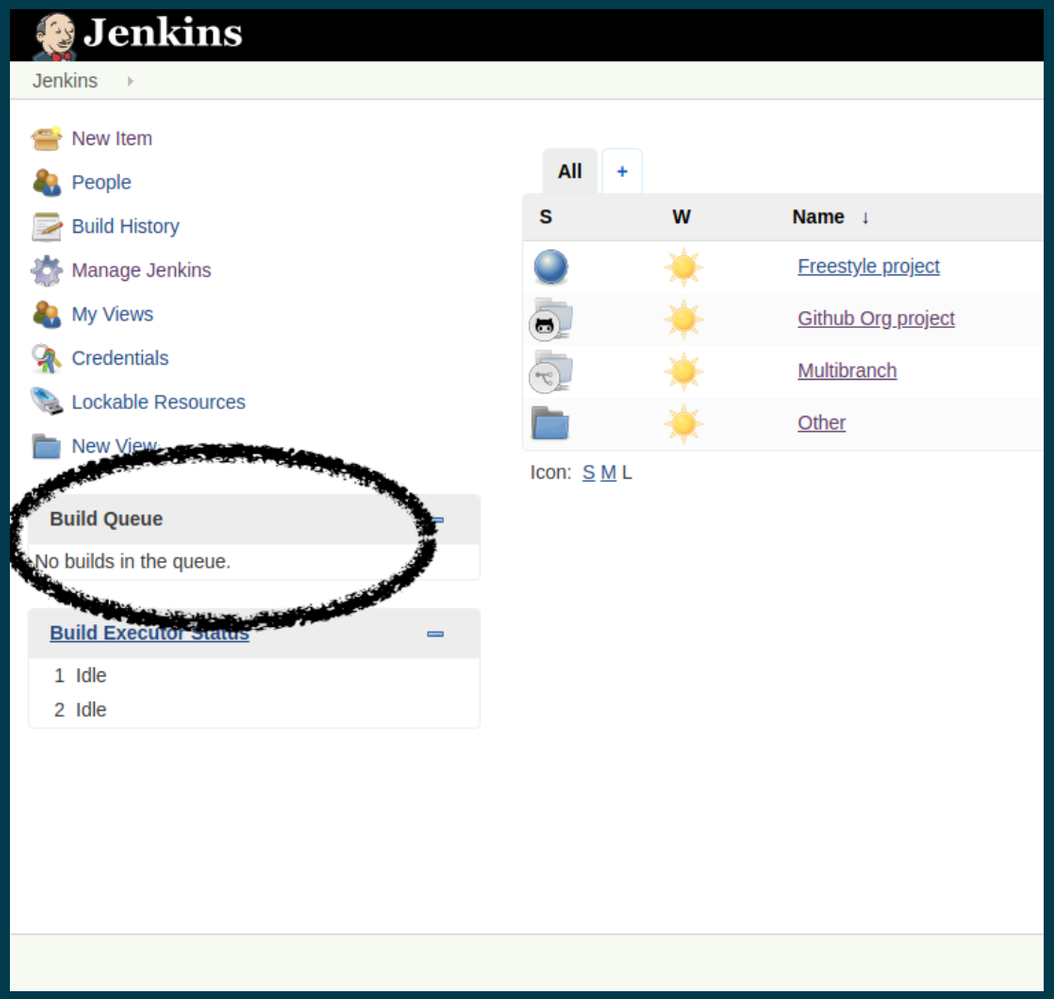

Jenkins AWS Plugin did spin up new nodes, but:

- new worker nodes took ~5 minutes just to be provisioned (EC2 + user-data)
- max 7 EC2 instances (4xlarge)
- one build took up several EC2 instances
- Jenkins EC2 cloud plugin scaled up by one at a time
- Typical for cold builds to take > 30 mins

Problem 2: Random test failures

False positives

At least **ONE** failure per day,
not related to actual changes made.

Overcome by always rerunning
FULL pipeline on failure.

2-3 runs necessary for
build to pass some times.

BAD for team
confidence in test suite.

MORE delays…

Problem 3: Jenkins maintenance

Devs are expensive. Devs rely on CI. Therefore, CI is a prod system.

Hosting own CI is like hosting any other production system.

You need to maintain it, test before making changes to it and ensure it is up and running.

Any degradation of the service can block the whole team including production deploys.

Preparing staging environment for verifying any Jenkins core or plugin updates can cost a lot of time.

It felt like security updates happen almost weekly.

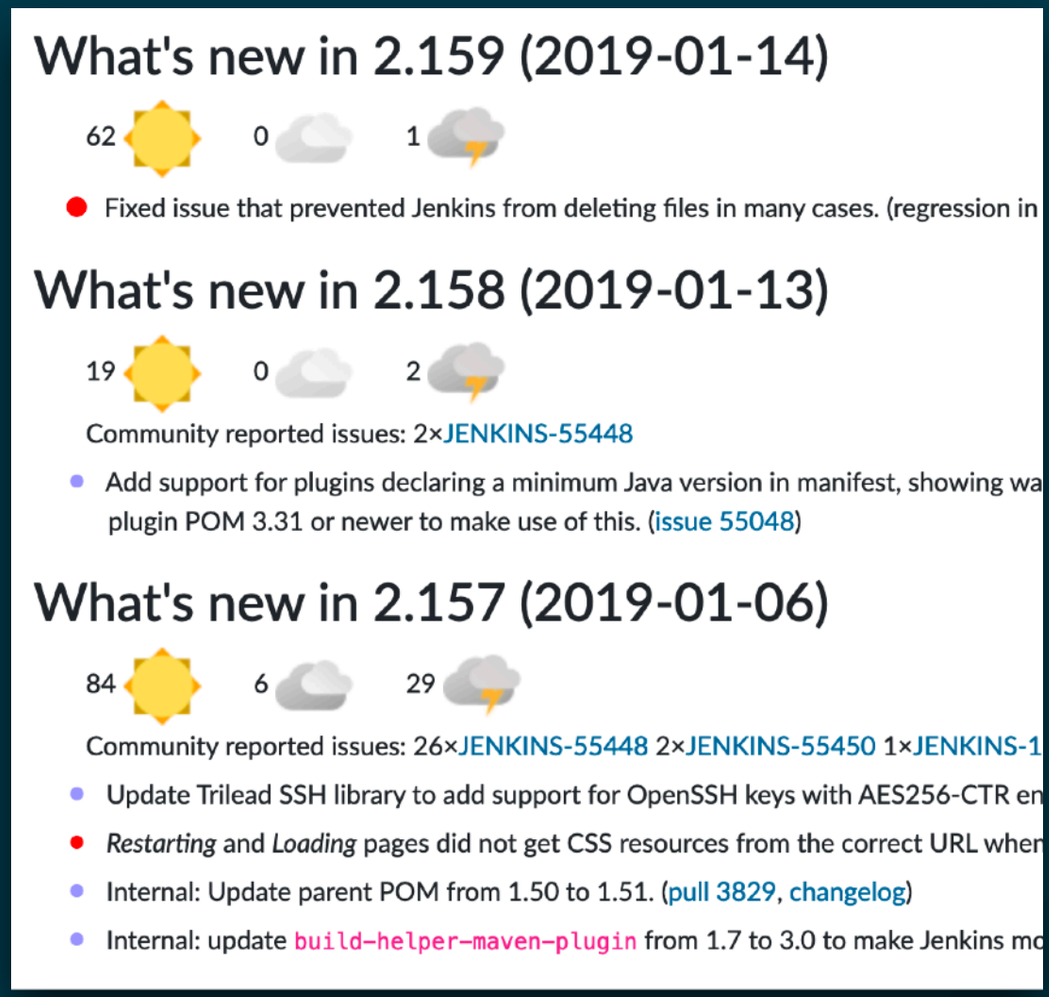

Problem 4: AWS Costs

Growing concern, especially as team was expected to grow

~2.5K EUR / month (just for AWS)

Total Costs =
AWS Costs +
Maintenance costs +
Dev team slow-down

Choosing our CI

Red Hat

Concerns

Our shopping list

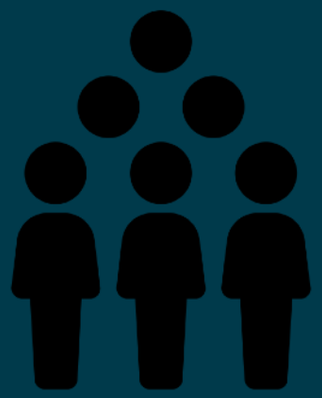

Created by Alice Design
from Noun Project

Publicly accessible build information

external contributors should be able to see if their build failed and why!

Builds from 3scale team as fast as possible

(willing to pay for that)

Builds from forks should be possible

but not billed on Red Hat (abuse cases in the past)

Upstream CI options

We need to give contributors an easy way to run the test suite

Red Hat

OR

Public CI

OR

Hybrid

Red Hat

It’s all about Developer Experience...

Red Hat

Public Build Info - Smooth DX

No account needed to access build information.

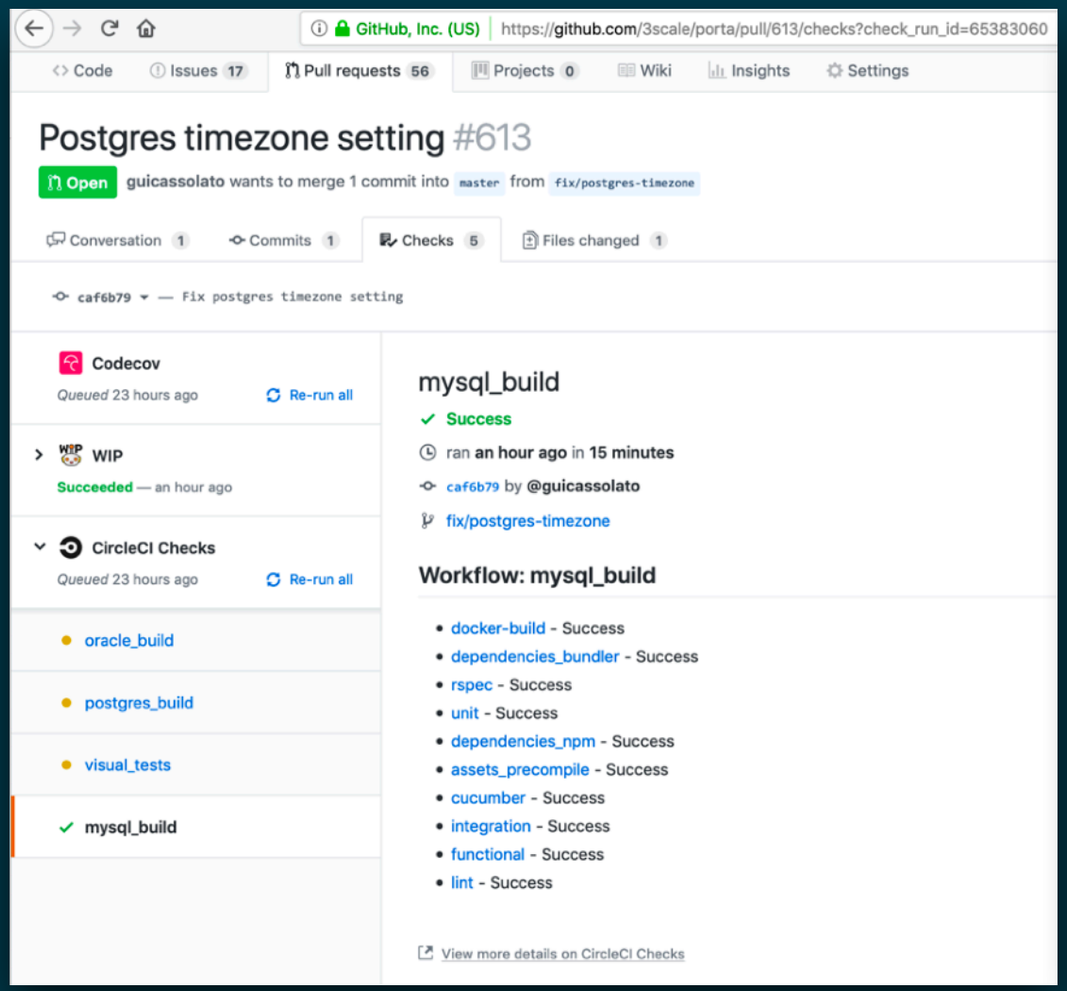

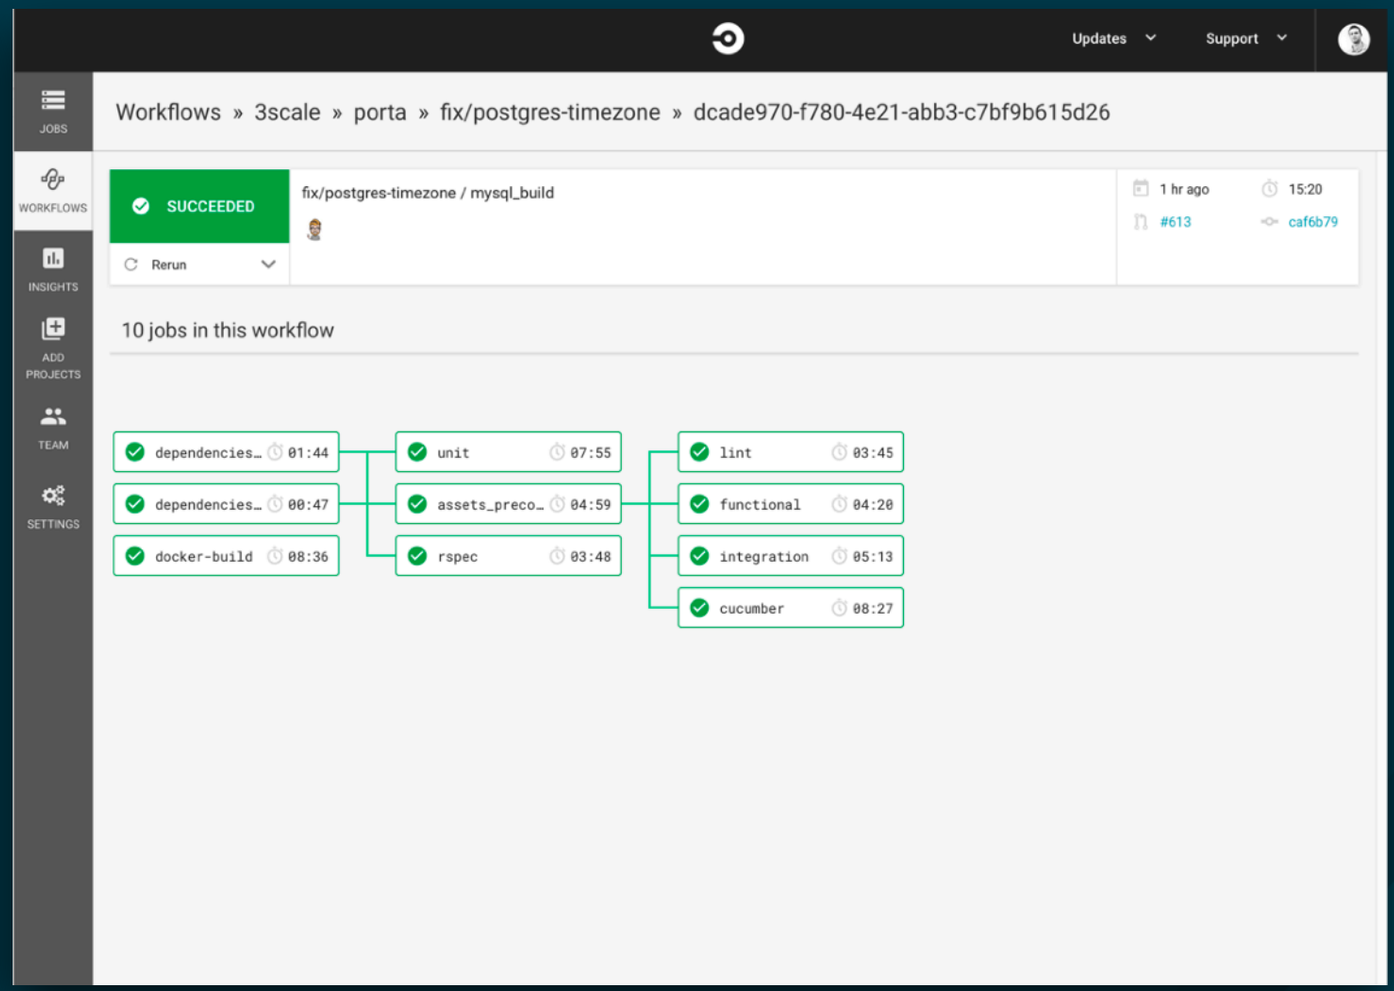

Accessible right from the GitHub pull request, to dive into detail

Red Hat

No more maintaining CI server!!

Remember: it is a production system!

Rerun from failed stage

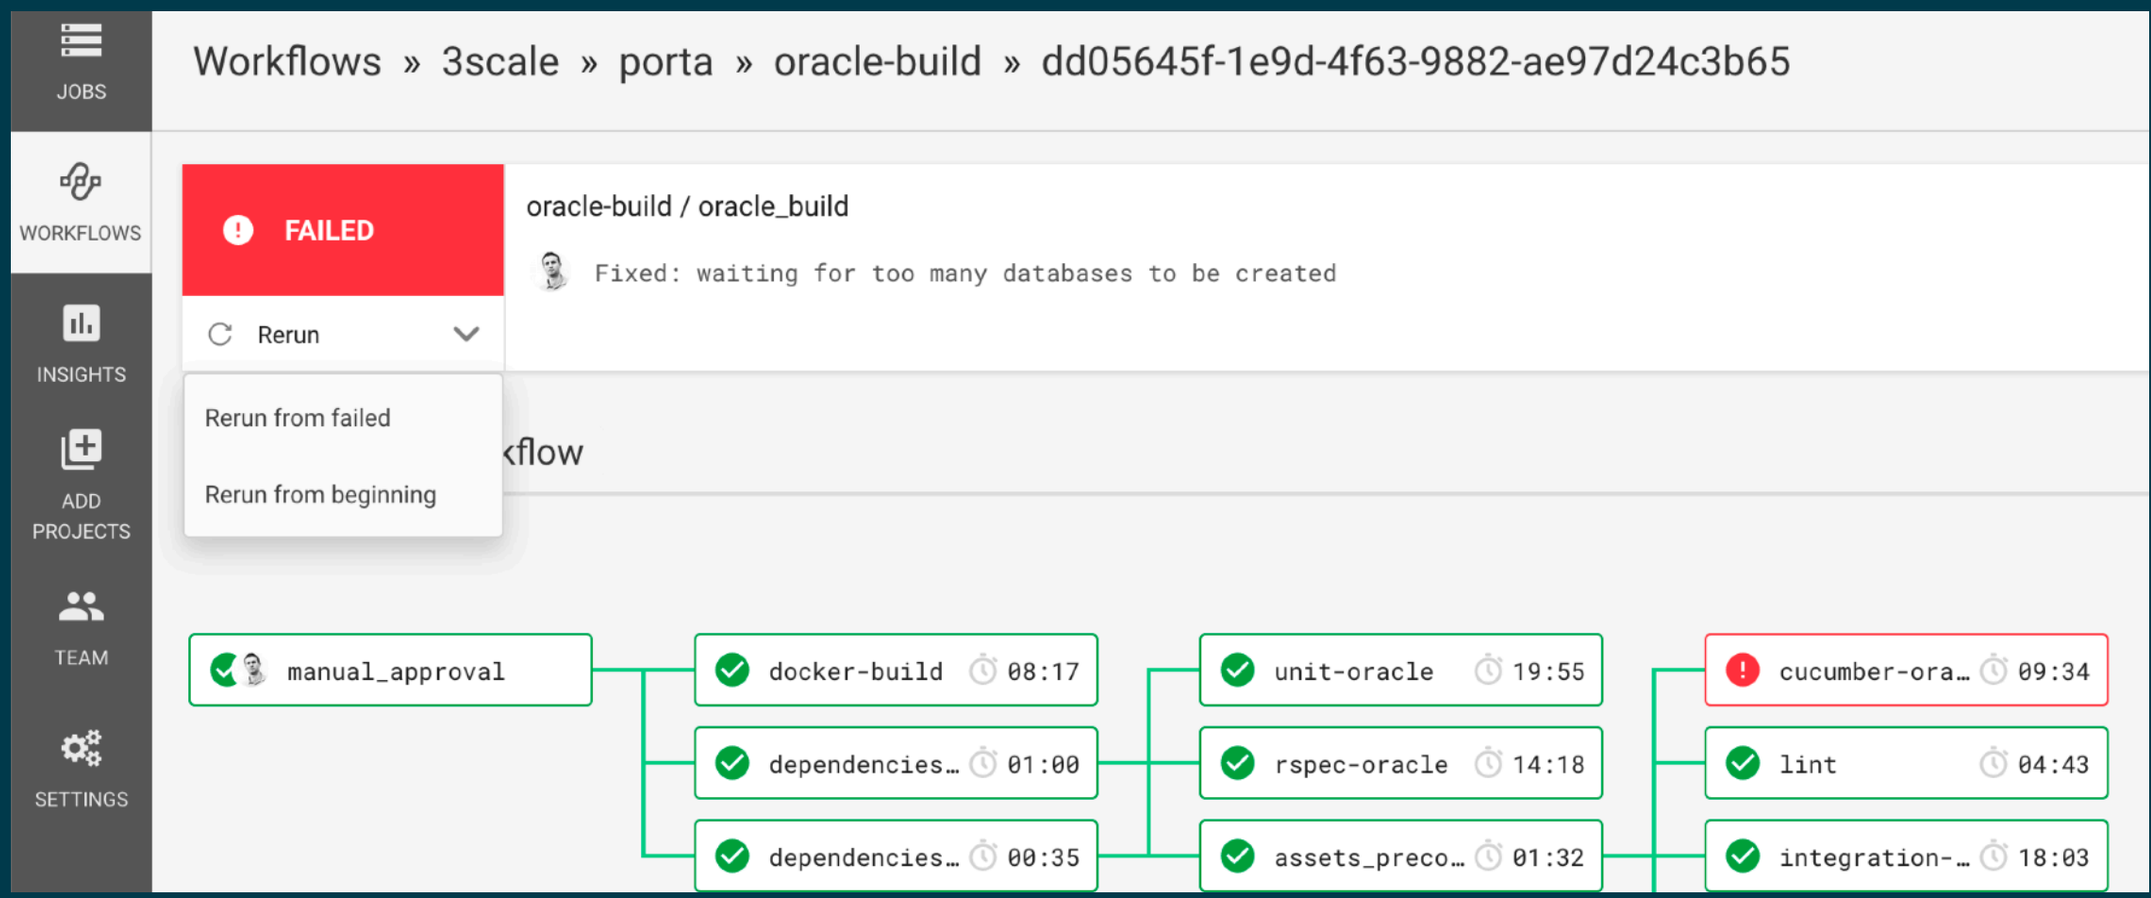

Pipeline only starts from segment that failed.

No waiting around, no billing for re-running same segments.

Debug CI failures

SSH to container that is running builds (allows us to get builds passing much faster)

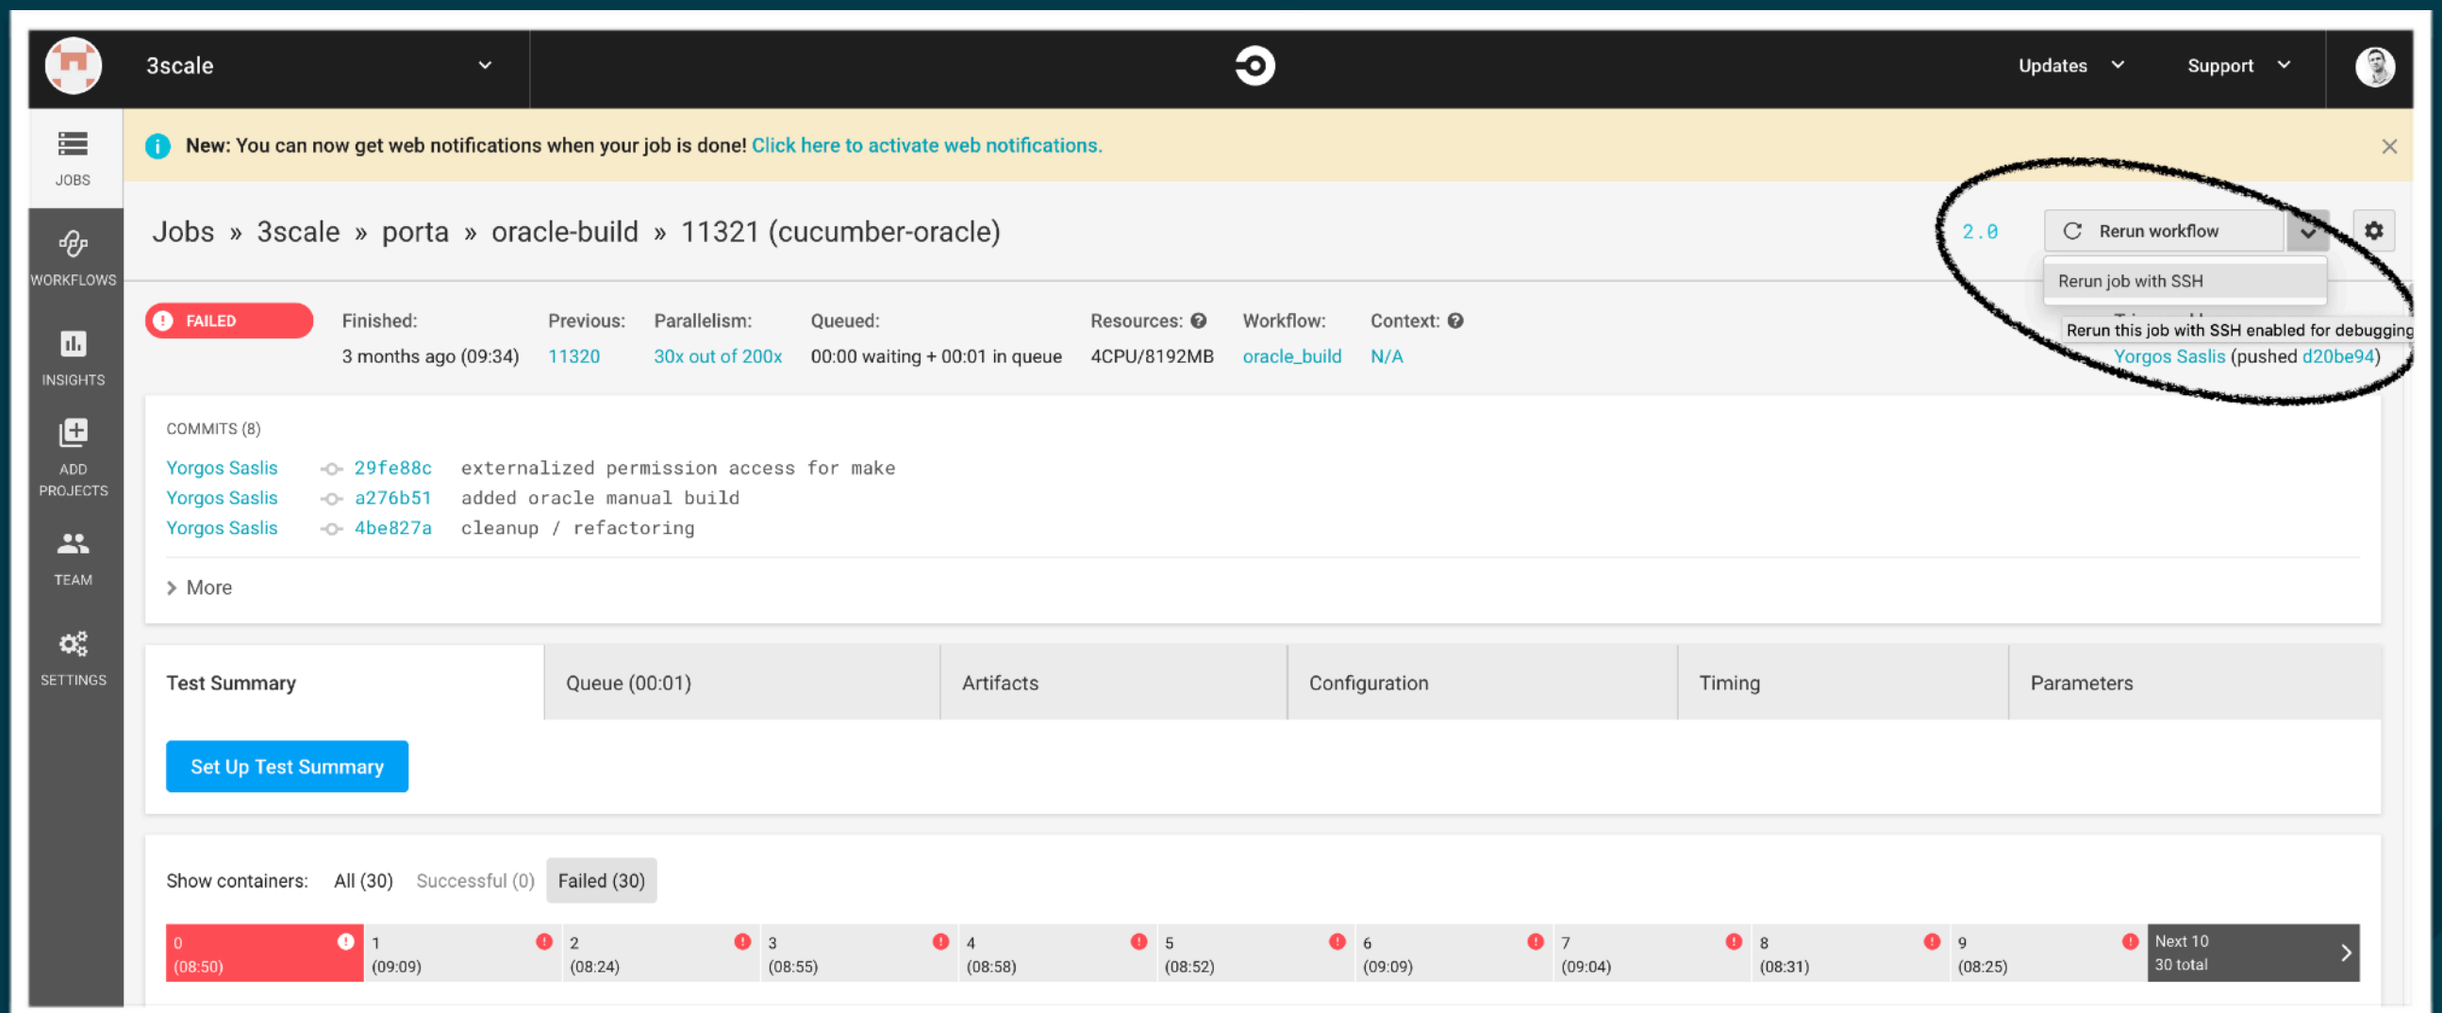

Bring up the environment to debug the failing build in just a couple of mins

...and price!

Red Hat

Price

It is cheaper because of better resource usage.

Using a fleet of short lived containers is better than VMs

2.5K EUR

vs

1.2K EUR

Let’s get back to parallelisation...

How do we parallelize our tests?

If we have to run
${numberOfTests = 1022}
tests,
how do we split them
into
${numberOfContainers = 40}
containers?

How do we parallelize our tests?

- Alphabetical
- Statically grouped
- maven phases
- JUnit Categories
- filesystem directories
- …

Wouldn’t it be great if we knew how long each test takes?

Split by timings

Helps orchestrate your test workload, to run in parallel

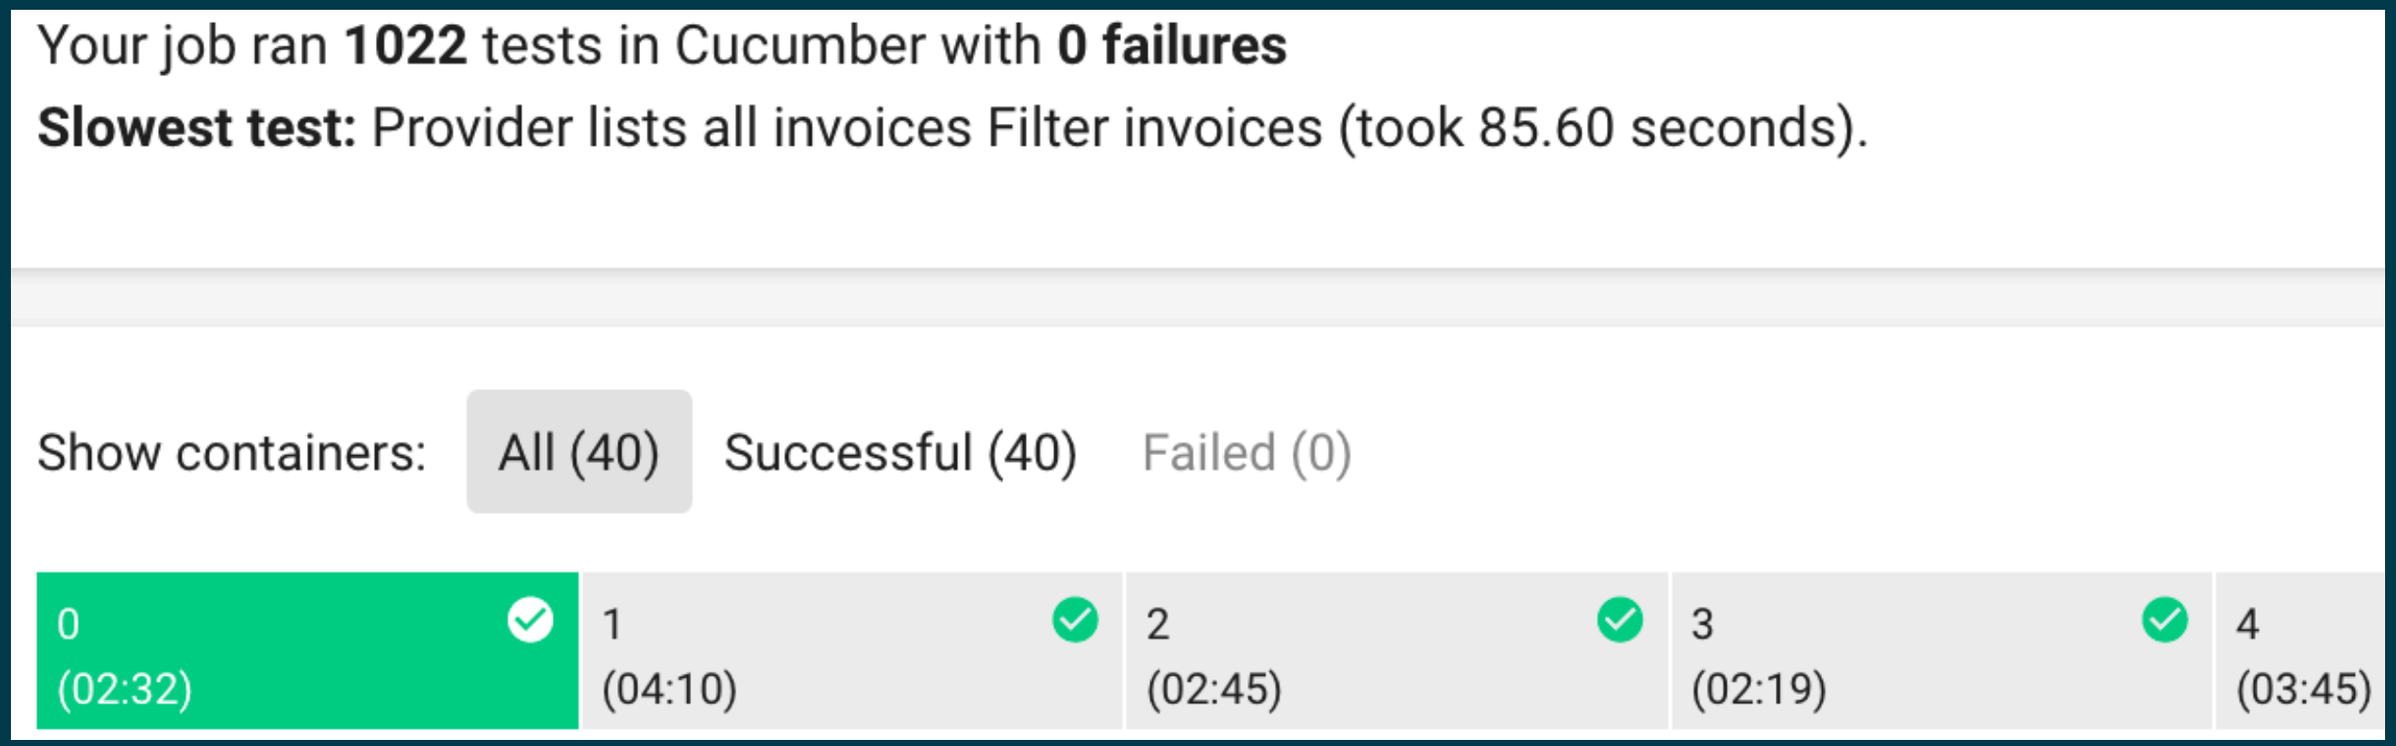

extract from `.circleci/config.yml` showing how cucumber tests are split

https://circleci.com/docs/2.0/parallelism-faster-jobs/#using-the-circleci-cli-to-split-tests

http://docs.shippable.com/ci/running-parallel-tests/

Split by timings - but how?

Helps orchestrate your test workload, to run in parallel

```
- run:
    name: Run cucumber tests
    concurrency: 40
    command: |
     bundle exec cucumber $(circleci tests glob “features/**/*.feature" \
             | circleci tests split –split-by=timings)
```

extract from `.circleci/config.yml` showing how cucumber tests are split

2 levels of parallelism

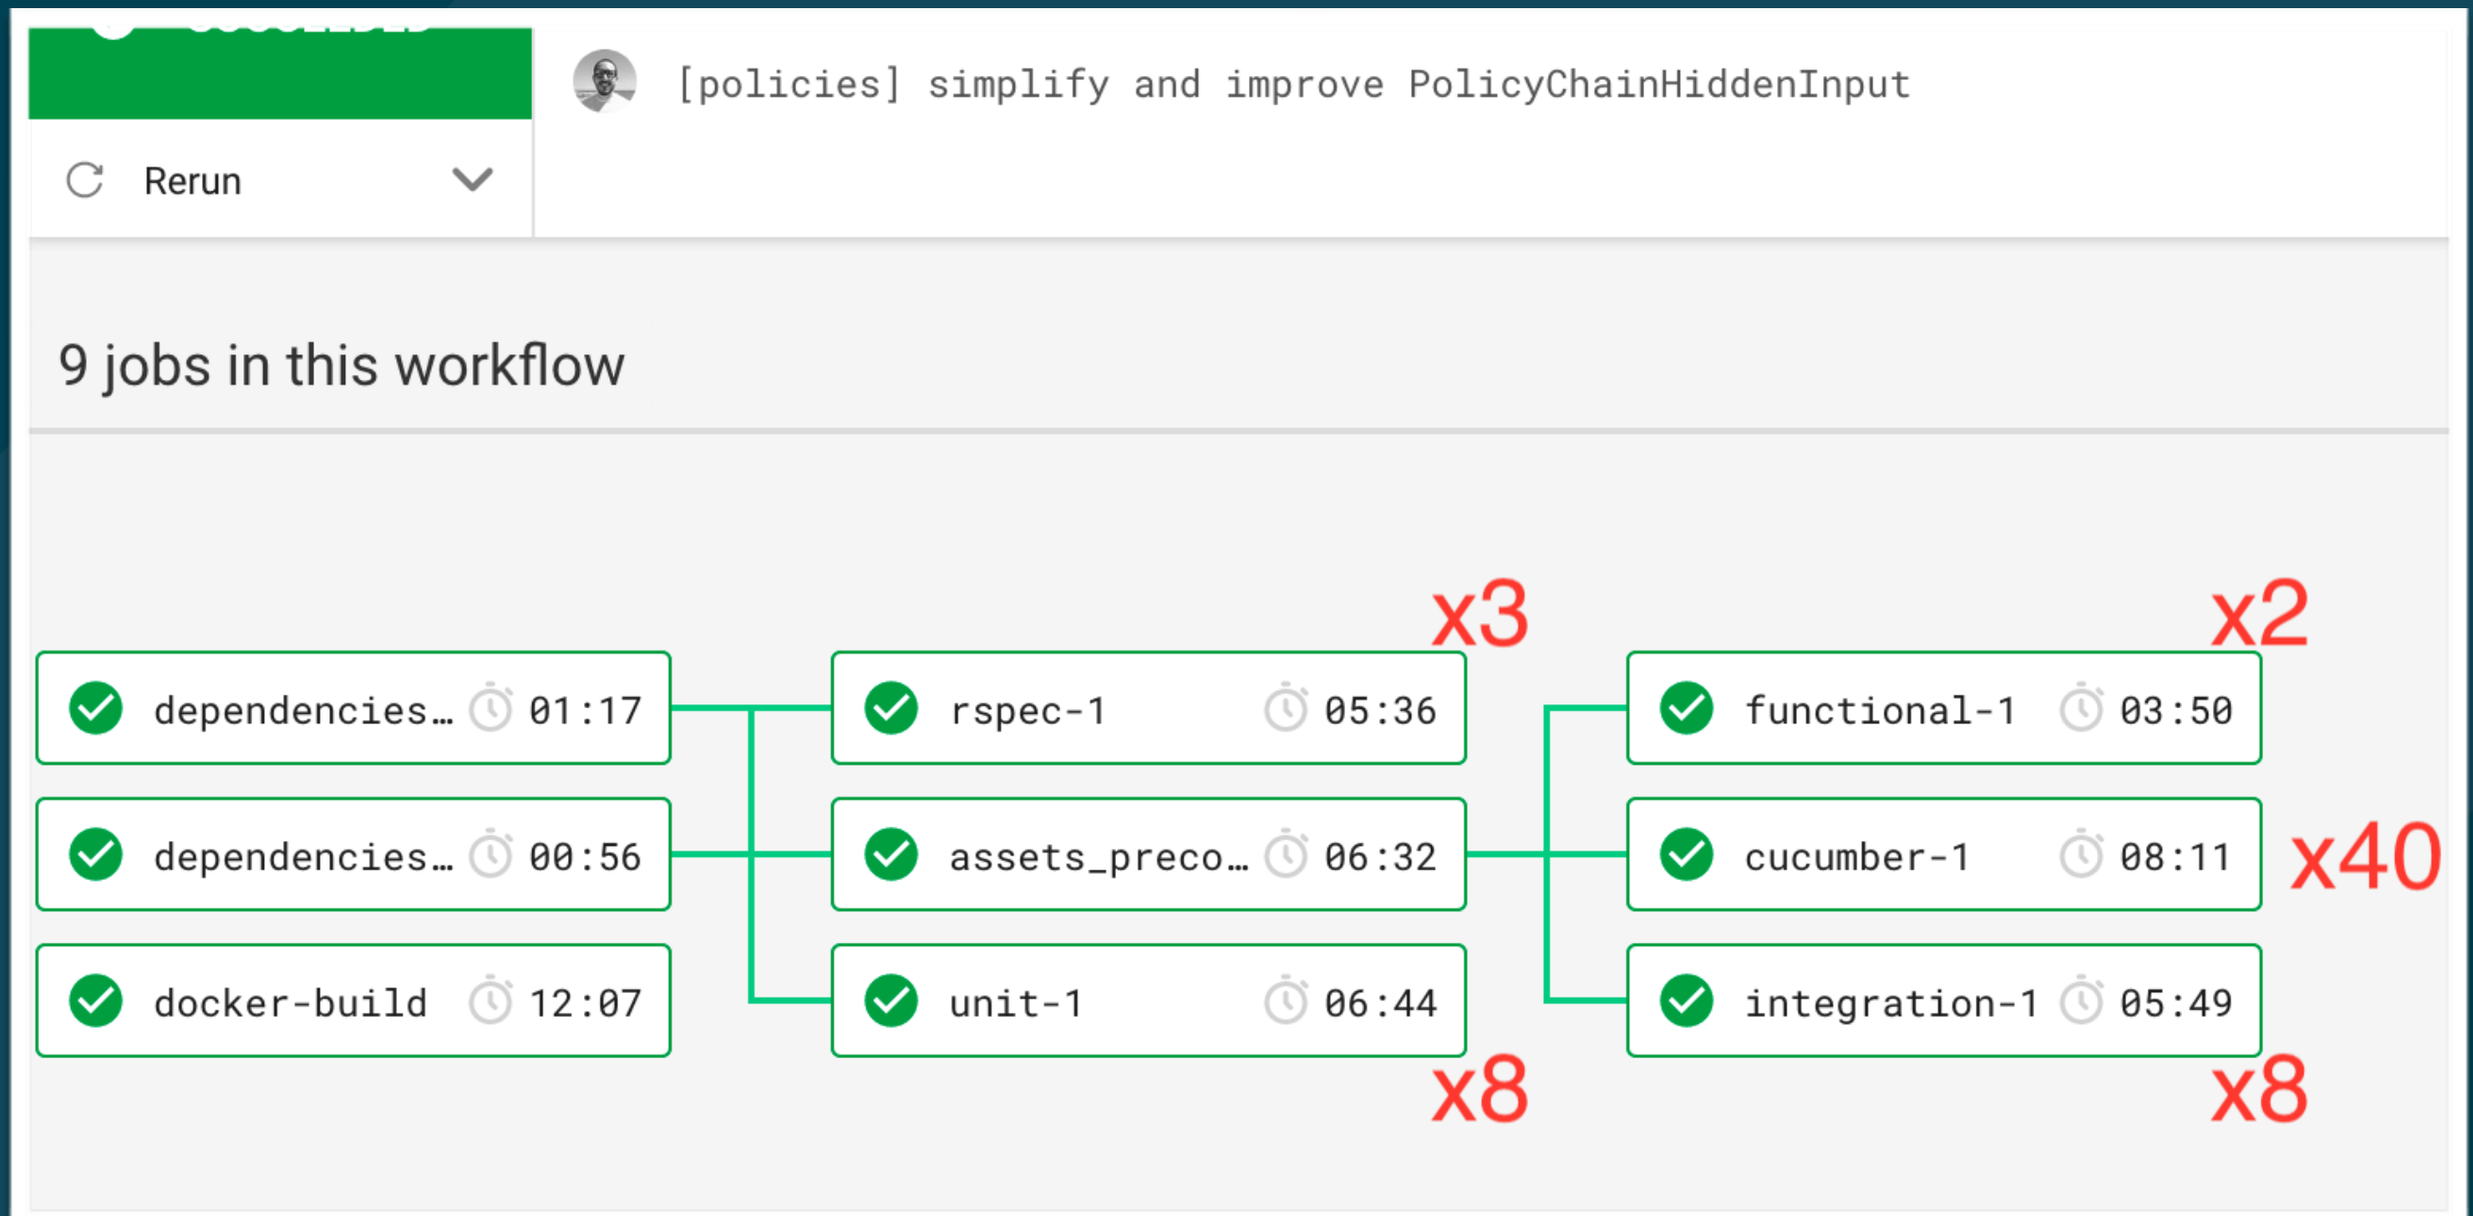

Red Hat

Tradeoffs

Nothing comes without sacrifice…

Costs

$$

Less configurable

than Jenkins

External

Dependency

Not

OSS

Not fully Open

Source Software

Enough about CI.

Let’s talk about tests!

Red Hat

flaky tests

Red Hat

Dirty State

If we rely on state for some tests, ensure it’s done properly.

LEFT-OVER STATE FROM PREVIOUS TESTS

Some tests that rely on bringing the System-Under-Test (SUT) into some “known” state - then running against that - don’t clean up after themselves properly.

BRINGING INTO KNOWN STATE ONLY COVERS SOME PARTS

E.g. if we rely on database for state, we didn’t restore a full database backup before every test (slow), rather we just modified some records in DB — but this does not ensure known state is what we expect it to be.

Reliance on other tests

Symptom: tests only pass if other tests have ran before them.

SomeFirstTest

SomeSecondTest

SomeThirdTest

Example: `SomeThirdTest` passes only when it happens to run after `SomeFirstTest` and `SomeSecondTest`

Tips how ensure test reliability

Discover randomly failing tests early

Randomize

Execute your tests in random order.

Verify you can rerun with the same seed.

Excercise

Run them 10 or 100 times a day if possible.

Not only on merge or pull requests.

Measure

Record test failures and times in machine readable format (JUnit, TAP, ...)

Steps to debug test order dependencies

The process we followed to identify problematic tests whenever a “random” failure occurred.

Reproduce

Run the batch of failing tests and reproduce the failure.

Bisect

Split the test batch in two.

Run only half of the tests.

Repeat

Go back to reproducing with just half of the tests.

Repeat until there are just two.

“The 13th factor: Tests”

Not enough focus on **test** codebase:

* parallelizable
* reliable
* independent of each other

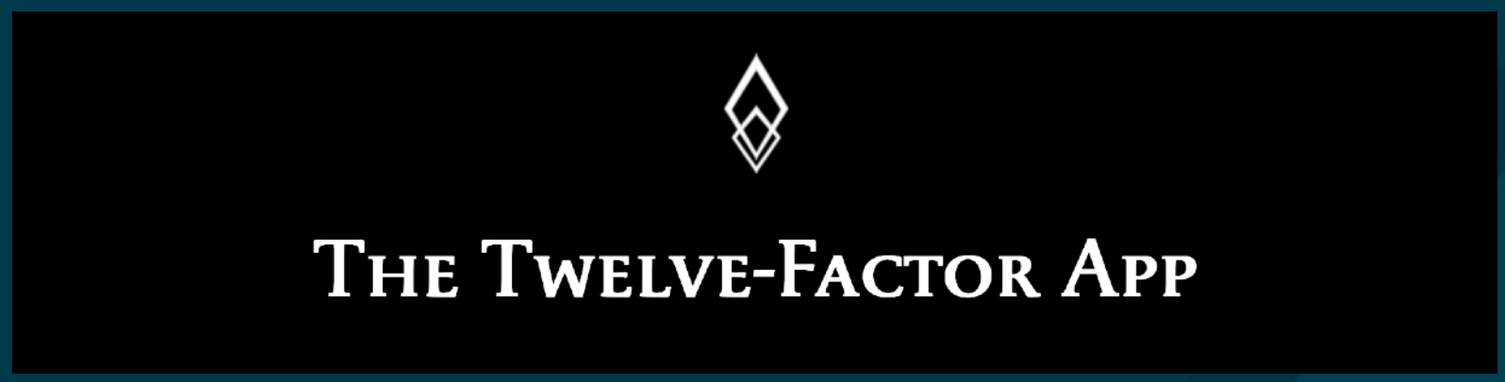

dependencies caching

Red Hat

External dependencies

Shave minutes off the build by avoiding to download from the internet

Use transitive dependency locking (Gemfile.lock, package-lock.json, Gopkg.lock, …)

* can be the same across builds
* no point running in “next” build if hasn’t changed from “previous” build
* use some cache

Try to use all CPU cores when installing dependencies.

External dependencies

Avoid reinstalling if they didn’t change since the last build

External dependencies

Don’t reinstall for each group. (don’t run `mvn clean verify` in each group…)

Group I

Group II

Group III

Group IV

Internal dependencies

Artifacts used inside the build

For example transpiled assets, bundling, optimizing images, etc.

The Future of CI/~~CD~~

Red Hat

CI != CD

- CI has very different needs than CD
  - Most deployments are usually simple
  - …compared to orchestrating the optimal, parallel execution of a test suite

- CI should only care about executing the tests, as fast as possible

- CI is a production workload with very predictable patterns
  - …unlike other production workloads

- CI should focus on Test, not on Pipeline

Next-gen CI

(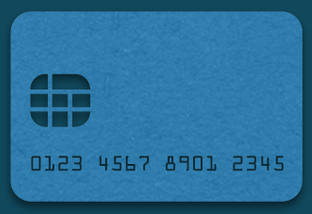 + Tests) =

Results

Red Hat

Dynamic test allocation

Optimising test suite parallelisation

Nodes pull more tests to run, when idle

Versus

Nodes get pushed a pre-allocated set of tests at start of test run

<3

your

tests!

Thanks for your attention!

Yorgos Saslis - @gsaslis

Michal Cichra - @mikz

github.com/3scale/porta

3scale BY RED HAT®

Red Hat