

FAST AND RESILIENT INTEGRATION TESTING

CONTINUOUS LIFECYCLE 2015

   Dr. Thomas Schank

  Max F. Albrecht

<http://drtom.ch/talks/2015/CL>



Version 1.1.0

This work is licensed under a [Creative Commons Attribution-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nd/4.0/).

Digitales Areal ZHdK

Ideen, Projekte, Werke – künstlerisch und wissenschaftlich: Das Medienarchiv der Künste ist die Plattform der ZHdK zum gemeinschaftlichen Arbeiten mit Medien und Teilen von Inhalten.

ZHdK-Login

Externe

Alle Funktionen nutzen und auf mehr Inhalte zugreifen.

Anmelden

Erkunden

Suche

Support

ZHdK-Katalog [Alle anzeigen →](#)

MADEK TEAM & US

 ZURICH UNIVERSITY OF THE ARTS



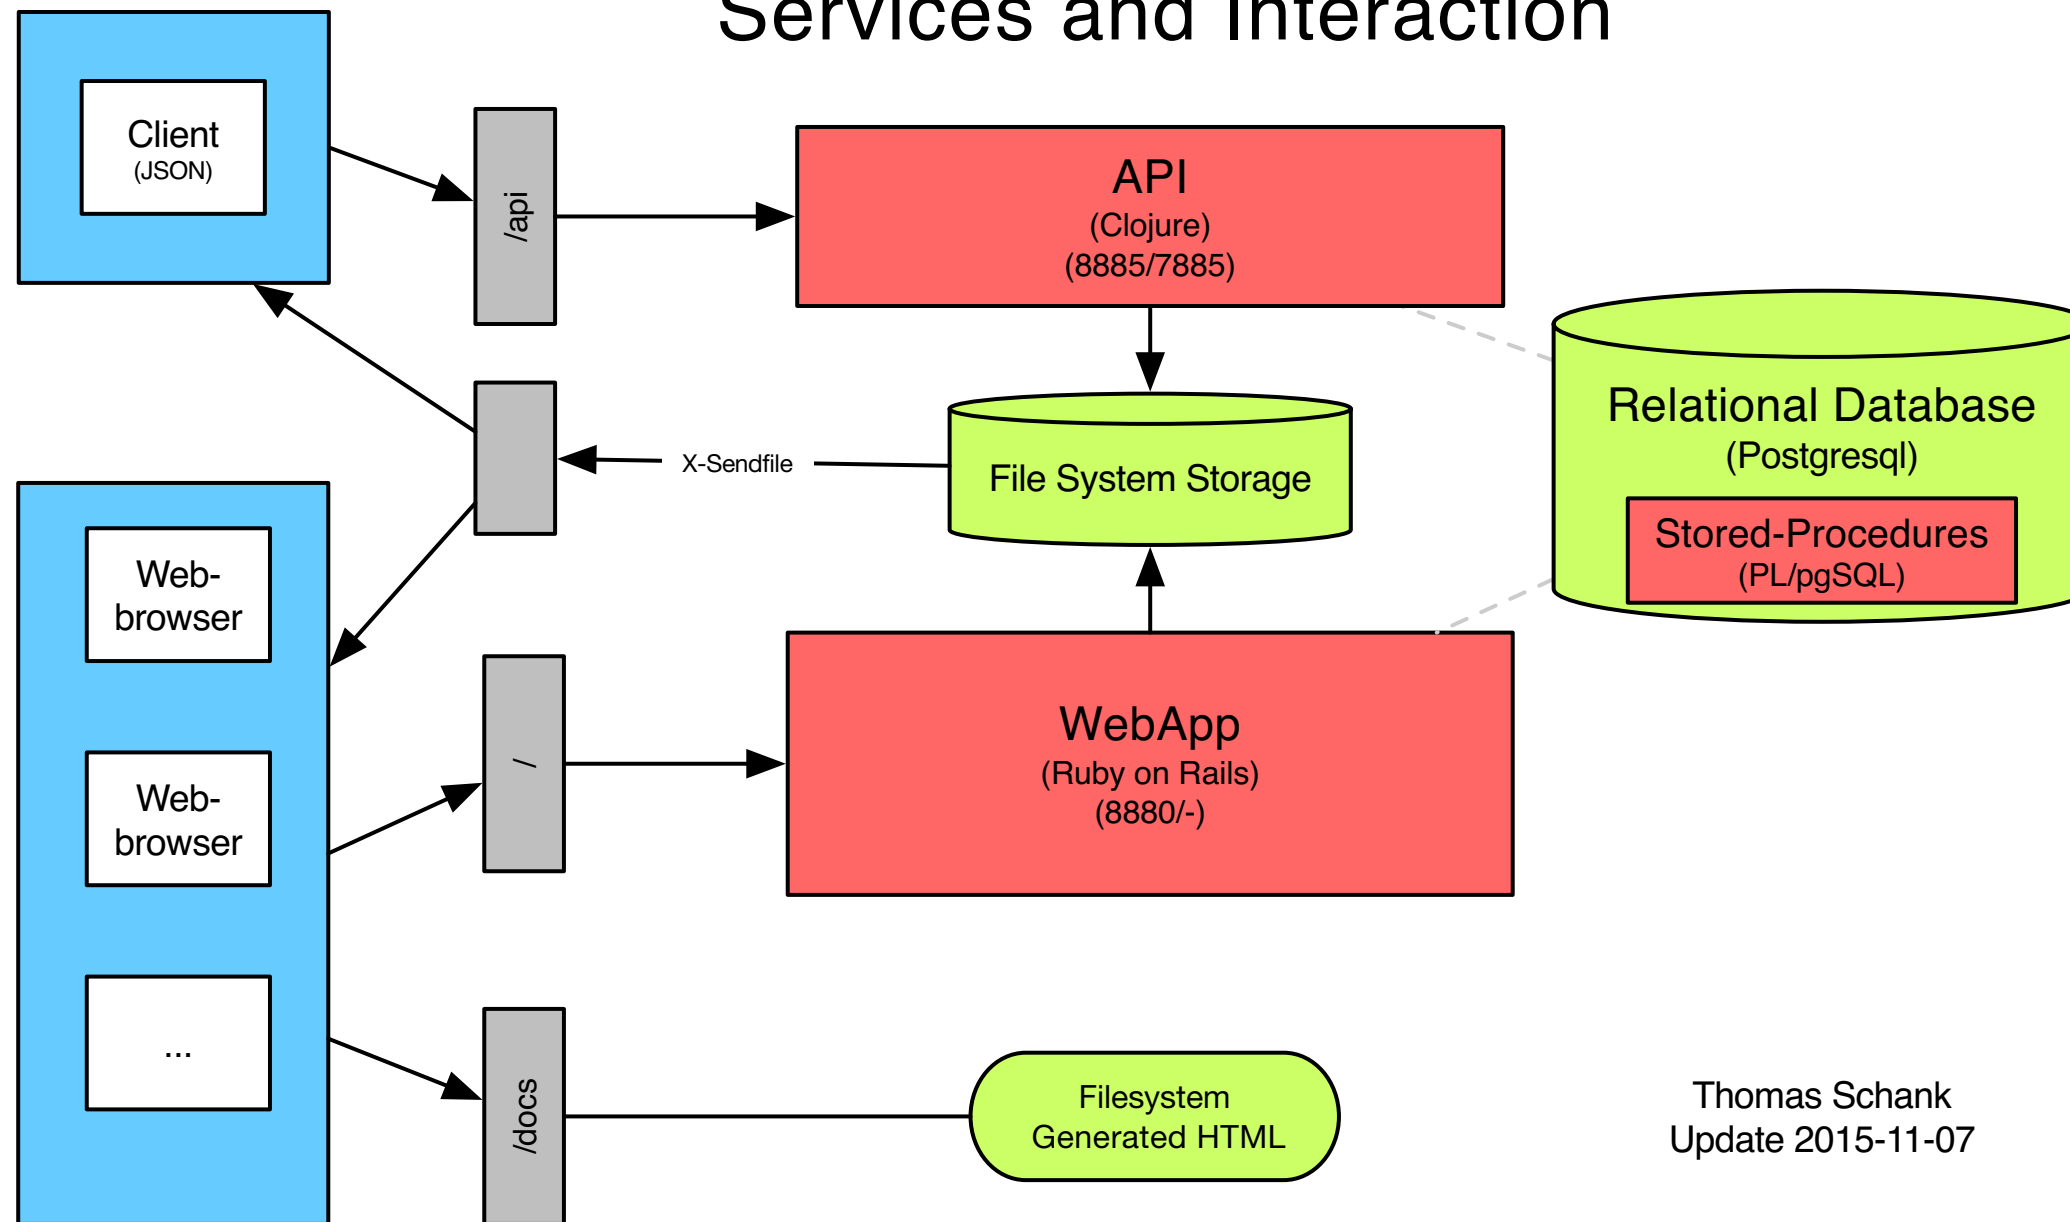
- Thomas: Software-Architect, Developer, CI-Infrastructure
- Max: Frontend Software-Engineer, Meta-Data Concepts

MADEK - MEDIENARCHIV DER KÜNSTE

ARCHITECTURE & TECHNOLOGIES

- Ruby on Rails, Clojure
- React with progressive enhancement
- 3-tier web-application
- towards micro-services
- deployment via Ansible to private cloud

Madek 3.0 Architecture Services and Interaction



Thomas Schank
Update 2015-11-07

MADEK TESTING

"specification by example"

→ integration testing

→ components interaction

1. THE PROBLEM



MADEK PROJECT 2012

many new features, many new tests

- testing time 1 1/2 - 2 hours, increasing
- more and more failing tests: **false negatives**
- 1/8 builds pass



TRY TO IMPROVE TESTS

- very time and resources consuming
- improvement for some time
- new features and new tests made efforts futile

MANUAL RETRYING

automated tests, local retries

automatic → semi automatic testing

2. COMPREHENSION



PROBABILITY OF A FALSE NEGATIVE FOR A WHOLE TEST-SUITE

	Expression	Example
probability false negative single test	p_f	3%
probability "success"	$p_s = 1 - p_f$	0.97
number of tests	n	100
probability "success" whole suite	$P_s = p_s^n = (1 - p_f)^n$	$\approx 5\%$

→ only one out of 20 will pass as it should

"succes" = true positive

WHY RETRYING WORKS SO WELL

let k number of independent retries per test

$$P_s(n) = (1 - p_f)^n \Rightarrow P'_s(n, k) = (1 - p_f^k)^n$$


Expected successful outcome for $n = 100$ and $p_f = 0.03$

k	P'_s
1	5%
2	91%
3	99.7%

2. COMPREHENSION - CONCLUSION

- more tests → exponential increase of likelihood for false negatives
- compensate by **retrying** single tests just a few times
 - retrying is not an anti-pattern
 - it can be a necessity

3. IMPLEMENTATION



Projekt MAdeK_AT__next___AGGREGATOR



Arbeitsbereich



Letzte Änderungen


Vorgelagerte Projekte

- [MAdeK AT next feature admin media sets](#)
- [MAdeK AT next feature admin people](#)
- [MAdeK AT next feature admin previews](#)
- [MAdeK AT next feature admin specify links and logo](#)
- [MAdeK AT next feature admin users](#)
- [MAdeK AT next feature apply to all](#)
- [MAdeK AT next feature batch edit permissions](#)
- [MAdeK AT next feature browse](#)
- [MAdeK AT next feature configurable contexts in lists](#)
- [MAdeK AT next feature context](#)
- [MAdeK AT next feature dashboard page](#)
- [MAdeK AT next feature delete](#)
- [MAdeK AT next feature explore page](#)
- [MAdeK AT next feature filter media resource types](#)
- [MAdeK AT next feature filter panel](#)
- [MAdeK AT next feature filter set](#)
- [MAdeK AT next feature graph visualization](#)
- [MAdeK AT next feature quest user](#)
- [MAdeK AT next feature import](#)
- [MAdeK AT next feature import 01](#)
- [MAdeK AT next feature import 02](#)
- [MAdeK AT next feature import encoding](#)
- [MAdeK AT next feature import metadata](#)
- [MAdeK AT next feature import via dropbox](#)
- [MAdeK AT next feature inheritance of contexts](#)
- [MAdeK AT next feature login](#)
- [MAdeK AT next feature metadata](#)
- [MAdeK AT next feature permissions](#)
- [MAdeK AT next feature permissions 01](#)
- [MAdeK AT next feature permissions 02](#)
- [MAdeK AT next feature preview](#)
- [MAdeK AT next feature search](#)
- [MAdeK AT next feature showing media resources](#)
- [MAdeK AT next feature uberadmin](#)
- [MAdeK AT next feature welcome page](#)
- [MAdeK AT next feature workgroups](#)
- [MAdeK AT next spec admin](#)
- [MAdeK AT next spec controllers](#)
- [MAdeK AT next spec controllers 02](#)
- [MAdeK AT next spec misc](#)
- [MAdeK AT next spec models](#)
- [MAdeK AT next spec requests](#)

JENKINS

- fall 2012
- build creates other builds via the Jenkins API
- last build aggregates
- solved false negative problem (partly)
- testing time: 15 - 25 minutes

→ it worked



- frequent code pushes interfere
 - "REST-like style API" → not much like REST
 - considerable effort and maintenance
- **Jenkins and "CI-X" just aren't made for this**


CIDER-CI

homegrown solution, started in spring 2013

- inherent support for **retries** and **parallelization**
- test **reproducibility**
- tight **integration** with **source code**
- **manage services** while testing
- support everything from **testing** to **deployment**

ready to use in fall 2013, **never looked back**

4. CONCEPTS IN CONTEXT

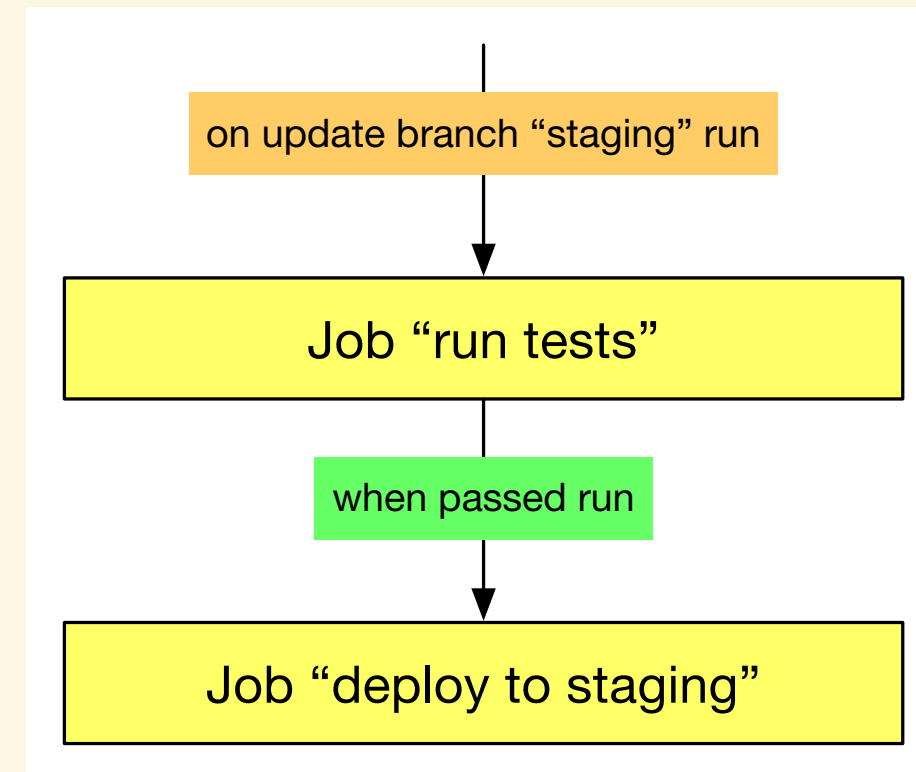


JOBS

EXAMPLES

- run test-suite
- perform static code checks
- build
- deploy

jobs can be **triggered** and can **depend on each other**



PROJECT CONFIGURATION

`cider-ci.yml` file in the project

```
jobs:
  deploy_test:
    name: Deploy to test

    depends-on:
      - type: job
        job: integration-tests
        states: [passed]

    run-on:
      - type: branch
        include-match: ^master$


# specify tasks etc
```

The source is the truth.

configuration: reproducible, reviews, audits ???

CIDER-CI AND THE SOURCE CODE

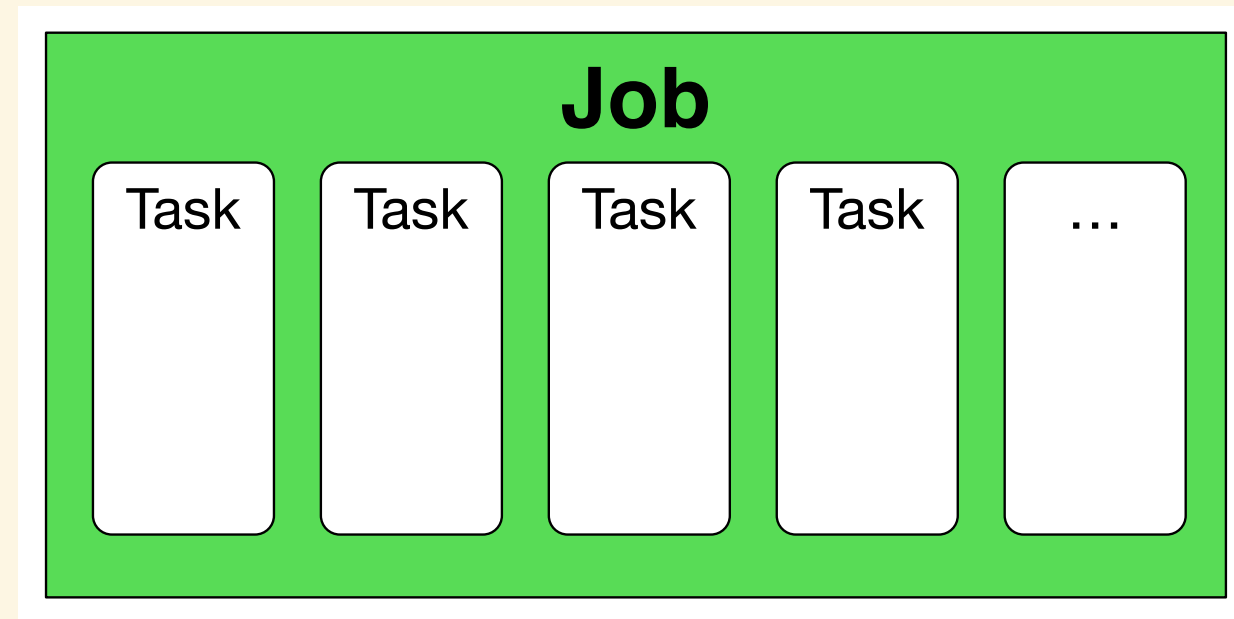
Cider-CI "knows" about commits, branches, submodules,...



`tree-id`: fingerprint of your source code

- **reproducibility**
- jobs can be **run at any time** (later)
- **binary search** for "bad" commits
- commit amends, squashing: **existing job remains valid**

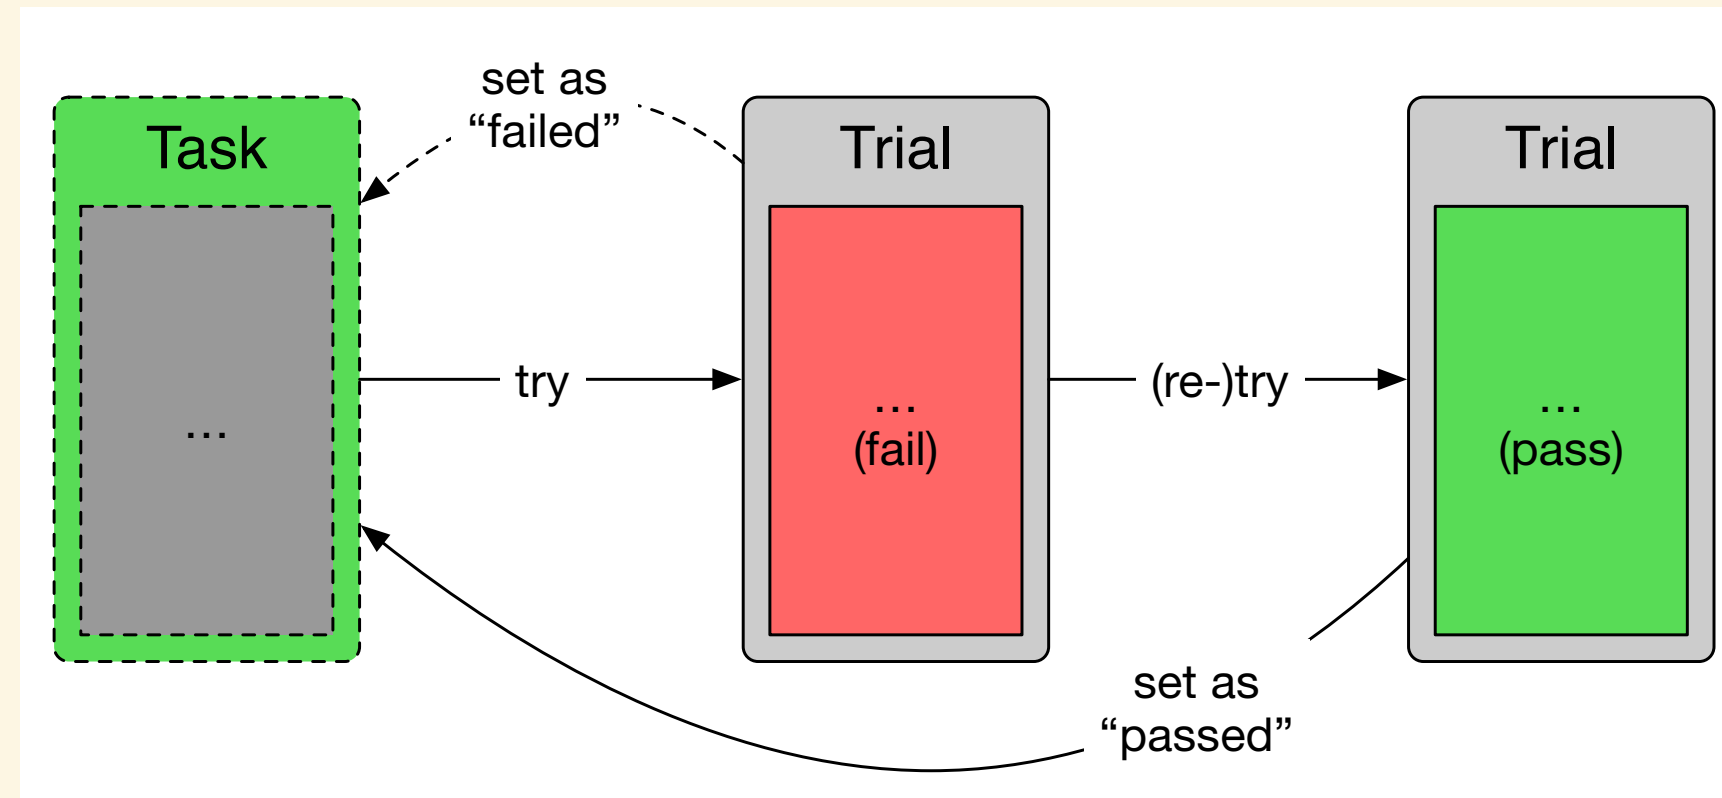
JOBS & TASKS



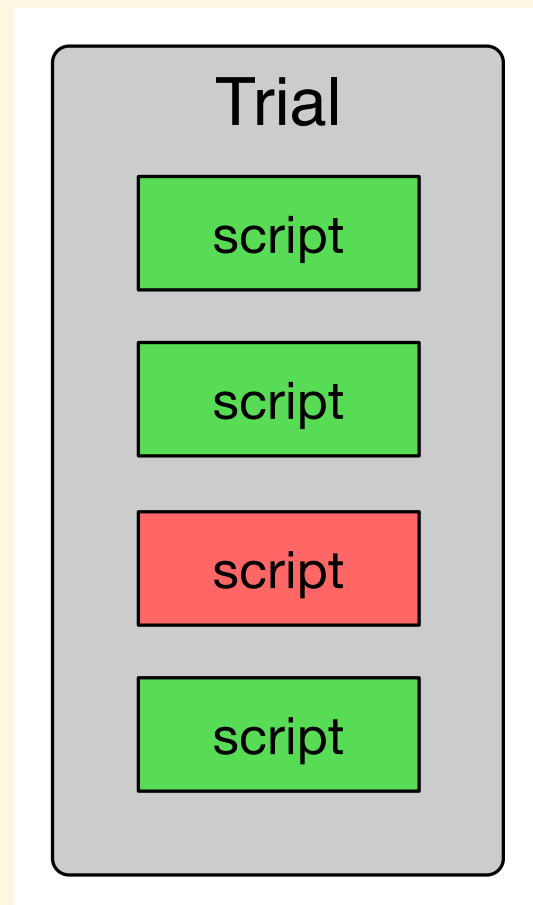
job: container and state aggregate for tasks

→ parallelization

TASKS & TRIALS



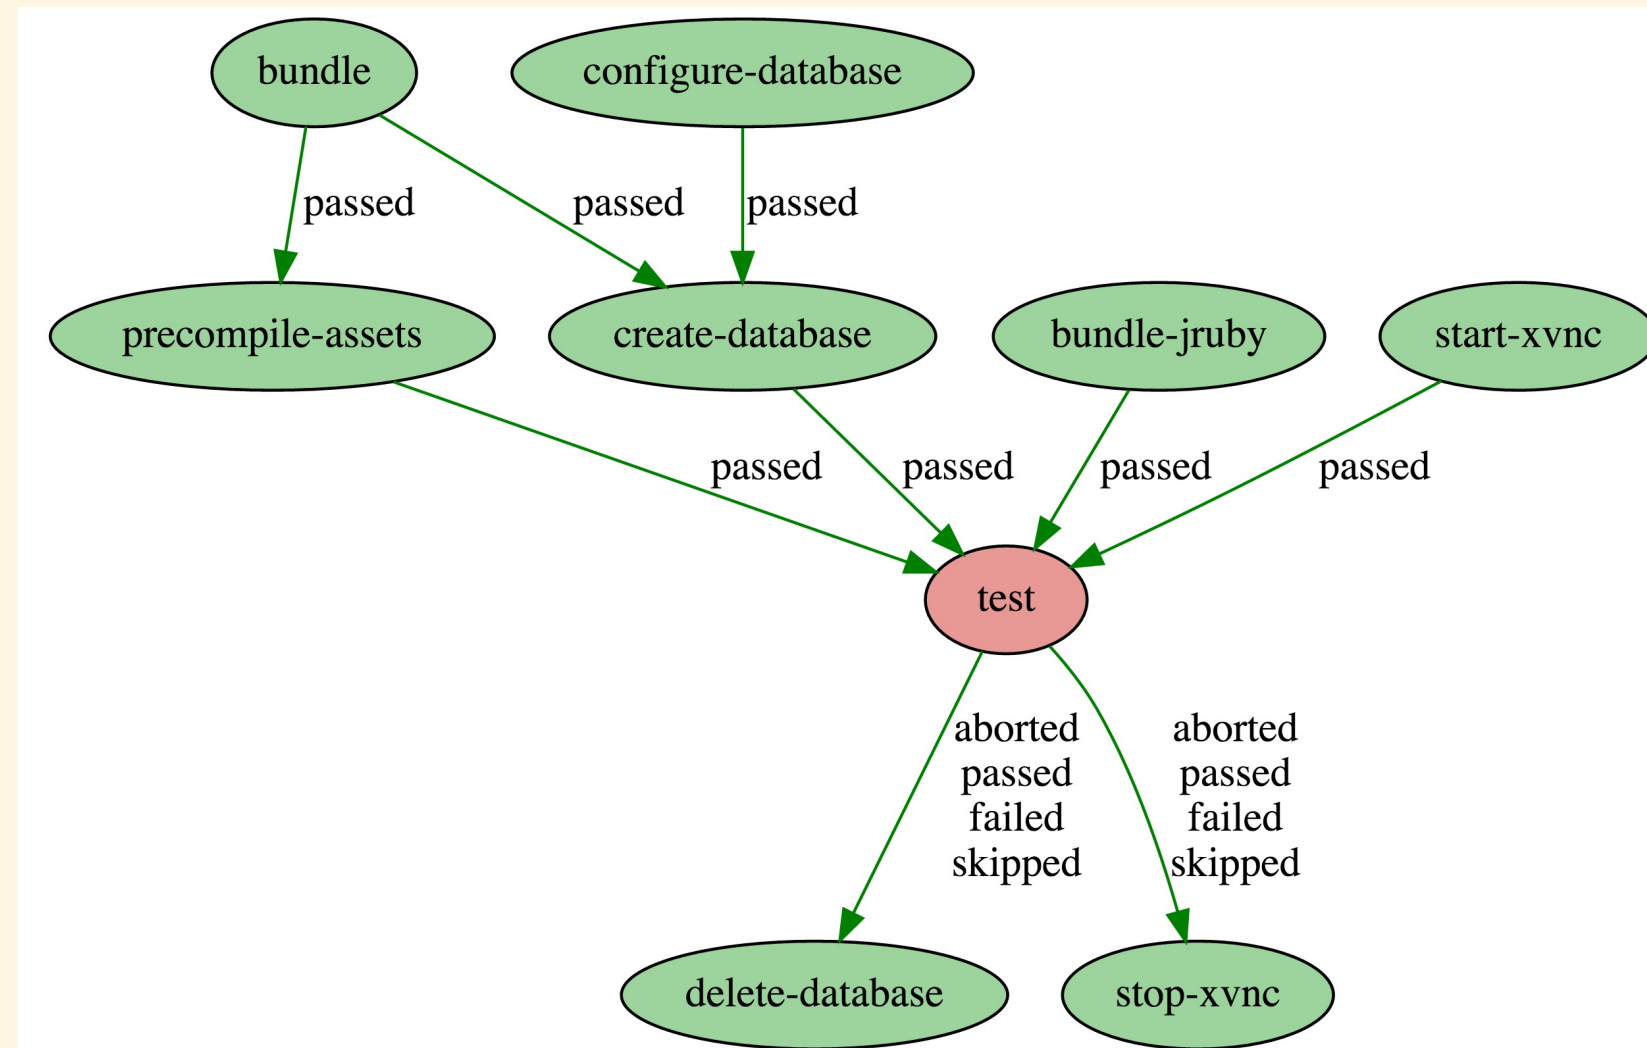
- blueprint
 - container and state aggregate for trials
- resilience



TRIAL & SCRIPTS


- actual unit of execution
- executed in the **same context**
- **depend** on each other

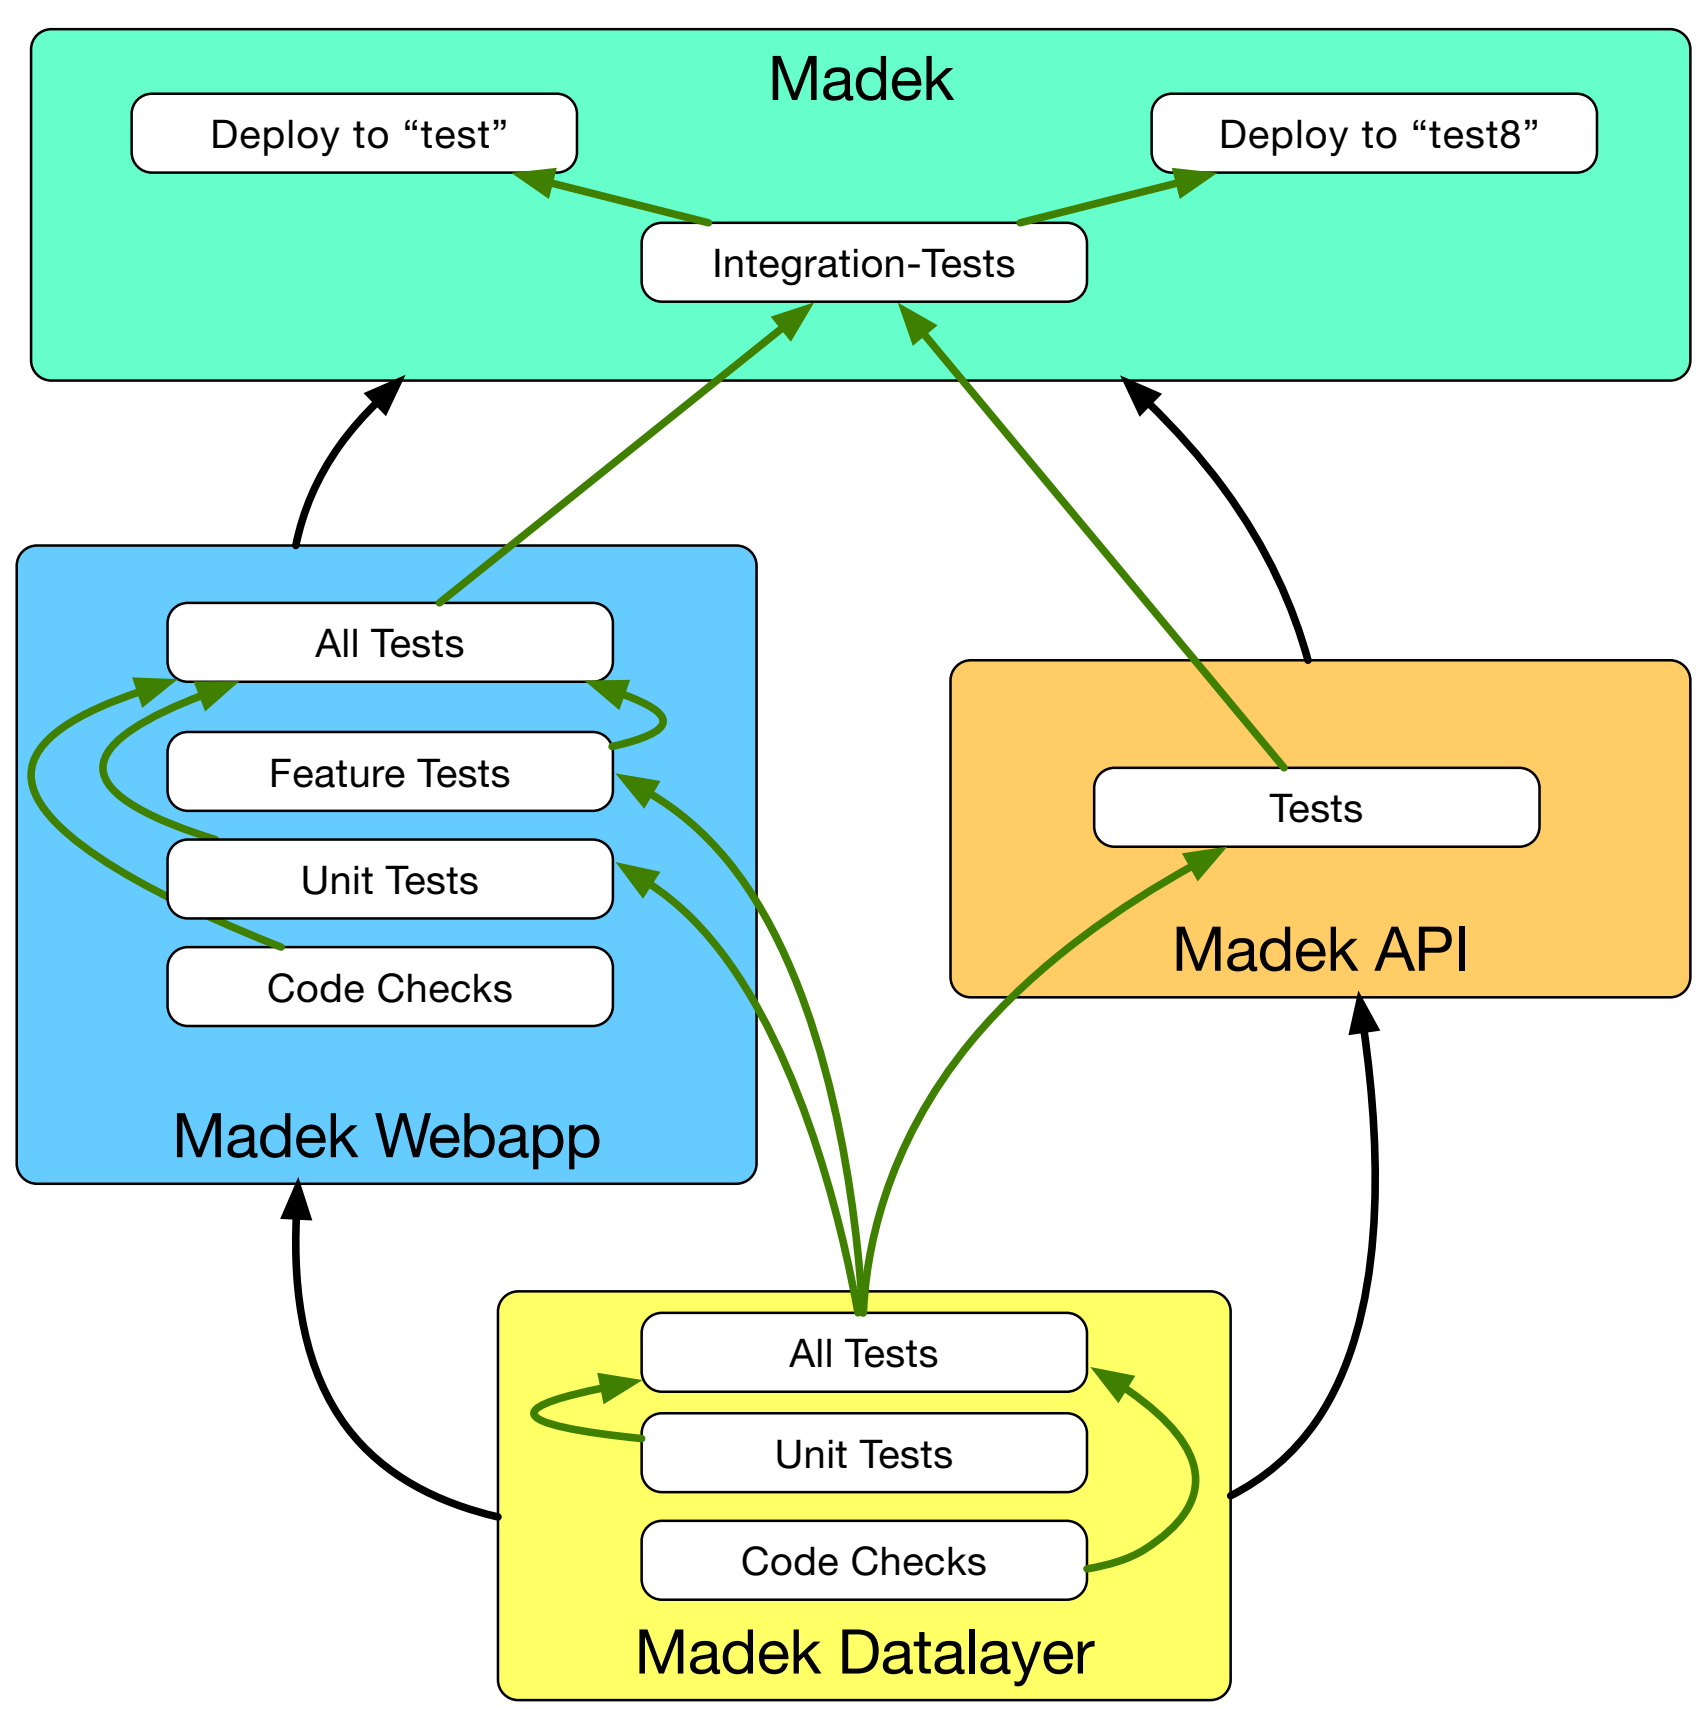
SCRIPT DEPENDENCIES



- traditional CI: one "build" \Leftrightarrow one script
- more modern: one main script + before and after "hooks"
- Cider-CI: **scripts with dependencies**

4. DEMOS





Cider-Cl Cardano 3.8.0-21 Workspace API Browser Documentation Administration drtom

Home / Workspace Next

Repository name 1, ... Branch name 1, ... Git-ref (tree-id, commit-id) Text search (subject, author, commite, ...) Reset Filter

Workflow and Job Dependencies Demo

4ffd50	324dda	2 minutes ago		Leihis	0 fs/procurement		Run ...
ecea56	30d1f5	8 minutes ago		Madek Webapp	0 ts_madek-v3		Run ...
2e2339	9d2750	41 minutes ago		Cider-Cl	0 tmp		Run ...
2cc8f5	dd7020	42 minutes ago		Cider-Cl User-Int...	0 tmp		Run ...
03e359	866c18	10 hours ago		Cider-Cl Integrat...	0 tmp		Run ...

```
thomas@nx-18122:~/...madek_v3/webapp$ git la -2
* 6509e57 6e28620 (HEAD -> ts_madek-v3) Remove `pending` declaration from JS Devtools test Thomas Schank, 21 seconds ago, committed 21 seconds ago
* ecea561 30d1f5e (github/ts_madek-v3) Perform fragment caching in dev mode when env var RAILS_CACHE is set Thomas Schank, 8 minutes ago, committed 8 minutes ago

thomas@nx-18122:~/...madek_v3/webapp$
```



Session: madek_v3.webapp 1 1 1:shell* 2:rails_server 3:rlwrap 4:vim#- 7:test 8:datalayer 09 Nov 08:31

ScreenFlow File Edit Mark Insert Font Actions View Window Help

Cider-Cl x Tom

ci.zhdk.ch/cider-ci/ui/workspace/jobs/9056ff32-3f14-4b4f-8ab6-92ea91935405

Cider-Cl **Cardano 3.8.0-21** Workspace API Browser Documentation Administration drtom

Home / Workspace / Job Analytics Specification

Job *Integration-Tests*

passed created 13 hours ago

This job runs the top level integration tests. It must be run from the master project.

Abort Retry & Resume Edit

4

Commit

5a4326 546a67 13 hours ago Upgrade webapp to move pending declar... Madek

0	master	🟢	🟢	🟢	Run ...
0	ts_master				
1	mfa_master				

Tasks

All substr OR substr 12 Per page Reset Filter

🟢	spec/features/basic_spec.rb	🟢	Retry
🟢	spec/features/media-file_up-and-download_spec.rb	🟢	Retry
🟢	spec/features/traverse-webapp2api-via-session_spec.rb	🟢	Retry
🟢	spec/unit/datalayer-consistency_spec.rb	🟢	Retry

Scripts Dependency
and Task Retry Demo

5. ADDENDUM



MANAGING FALSE POSITIVES

Retrying randomized tests can hide problems!

"Generative Testing" e.g.

SOLUTION:

- reproducibility by initializing the pseudo random generator (we use the `tree_id` e.g.)
- statistics

GIT SCM AND GIT ONLY

- don't compromise
- can't support everything with reasonable effort

SECURITY & TRUST

- Cider-CI server itself **never runs any code** from projects
- **"blessed" executors** only accept trials for a particular project (repository)

MATCHING TRIALS TO EXECUTORS

- task specifies **required traits**, e.g: [bash, ruby-2.2]
- executors advertise **available traits**, e.g. [bash, maven, postgresql, ruby-2.1, ruby-2.2, ...]

Cider-CI will determine a suitable executor.

DEPLOYMENT

- Ansible
- Cider-CI deploy project, SCM managed, reproducible

CIDER-CI IS AN EXPERT SYSTEM

it is about making the hard possible, and not not about making the simple easy*

- for professionals
- no compromises
- steep learning curve
- high rewards

→ swiss army knife for devops

*see "Simple Made Easy" by Rich Hickey

CONCLUSION

- A false negative outcome becomes likely with an increasing number of tests.
- The problem must be solved by retrying single tests.
- Consider to build your own pipeline.
- Try Cider-CI, open source, installs with two commands:
<http://docs.cider-ci.info/introduction/quick-start/>



Thank You!