



## CircleCI RSpec Suite Improvements: Q3 2024

Presented by Konstantin Gredeskoul, Staff Software Engineer, Infrastructure



## CircleCI RSpec Suite Improvements: Q3 2024

### What this is about

We use CircleCI to run an automated test suite, which ensures critical parts of the system work with every change we push (if any test fails, then we know we have something that needs fixing).

### Problem:

- Running the test suite on a single dev box takes **over 2 hours & 35 minutes**

real 155m 4.954s

- So we parallelize into 40 containers to speed up the process to ~ 18 minutes on average back in June 2024.



Filters

Everyone's Pipelines ▾





































academia-app ▾

Select a Branch ▾

Greater than 91 days ▾

🔍 ▾

Auto-expand ☐

Project	Status	Workflow	Trigger Event	Start	Duration	Actions
academia-app 135527	▶ <span>✓ Success</span>	build_and_ test	 master a58df97 [KG-105] Switch shadow paper mentions algorithm (#35764)	3mo ago	15m 13s <span>↑ 61%</span>	   ...
academia-app 135525	▶ <span>✓ Success</span>	build_and_ test	 fake_door_2024 cda9422 Pass feature to frontend	3mo ago	15m 45s <span>↑ 67%</span>	   ...
academia-app 135524	▶ <span>✓ Success</span>	build_and_ test	 toby/SRE-2521-rm-ff-in-EmailOpen- firehose_batch_frequency 79f788e Remove feature flag for EmailOpen.firehose_batch_frequency	3mo ago	21m 38s <span>↑ 129%</span>	   ...
academia-app 135523	▶ <span>✓ Success</span>	build_and_ test	 soc-2176 31fa655 Generate content scores along with candidates <span>RERUN FROM START</span>	3mo ago	16m 57s <span>↑ 80%</span>	   ...
	▶ <span>✗ Failed</span>	build_and_ test	 soc-2176 31fa655 Generate content scores along with candidates	3mo ago	13m 33s	   ...
academia-app 135521	▶ <span>✓ Success</span>	build_and_ test	 patbl/harmonize-rubocop-configs 8d2f4a0 .rubocop.yml: Enable Academia/YAMLLoad cop <span>RERUN FROM START</span>	3mo ago	17m 55s <span>↑ 90%</span>	   ...
	▶ <span>✗ Failed</span>	build_and_ test	 patbl/harmonize-rubocop-configs 8d2f4a0 .rubocop.yml: Enable Academia/YAMLLoad cop	3mo ago	36m 46s	   ...
academia-app 135520	▶ <span>✓ Success</span>	build_and_ test	 patbl/harmonize-rubocop-configs a672398 .rubocop.yml: Enable Academia/YAMLLoad cop	3mo ago	14m 41s <span>↑ 56%</span>	   ...
academia-app 135519	▶ <span>✓ Success</span>	build_and_ test	 master 691f8d2 Fix send_batch_to_firehose test (#35770)	3mo ago	19m 45s <span>↑ 109%</span>	   ...



## CircleCI RSpec Suite Improvements

### The Problem

- That's 18 minutes of waiting to know whether your feature broke anything.
- Multiplied by the number of engineers, and by the number of times a day each engineer pushes a change to a PR, and that's a lot of wasted minutes.



## CircleCI RSpec Suite Improvements

### What makes test suite slow?

- A sheer number of specs (tests) in the test suite
  - NOTE: Speeding up individual tests was outside the scope of the effort
- Installing and validating large number of dependencies
- Number of concurrent CircleCI builds running at a given time
- General busyness of the cloud hardware shared with other CircleCI customers






























## CircleCI RSpec Suite Improvements

### Solutions: 1. KnasackPro

- In August, we integrated our test suite with a 3rd party service called **Knapsack Pro** which remotely and more efficiently coordinates which container runs which test, ensuring they all finish around the same time.
- **Outcome:** the average test suite time went down from 12-14 **minutes** on average to about 7-8 mins (approx. a **50% improvement**)
  - NOTE: all improvement % are approximate ball-park figures,



Project	Status	Workflow	Trigger Event	Start	Duration	Actions
academia-app 140384	✓ Success	build_and_ test	 master b31fbab [soc-2274] job for filling posts into extended social OpenSearch index (#36672)	1mo ago	7m 0s ↓ 26%	 
academia-app 140383	✓ Success	build_and_ test	 master 14d9a1e Add department on Editor Cards for microsite	1mo ago	7m 10s ↓ 24%	 
academia-app 140379	✓ Success	build_and_ test	 master 27ac3ff Hotfix - set variant properly in transient redis setex (#36676)	1mo ago	6m 59s ↓ 26%	 
academia-app 140378	✓ Success	build_and_ test	 fix_empty_mention_title 3935112 fix empty titles being displayed on mentions	1mo ago	6m 47s ↓ 28%	 
academia-app 140377	✓ Success	build_and_ test	 soc-2274-extended-social-opensearch c442aac refactor into abstract class	1mo ago	6m 46s ↓ 28%	 
academia-app 140376	✓ Success	build_and_ test	 isaac--fix-unintended-keypass-fire 2be7a7e Refactor sentry full error code and handle running when sentry is off	1mo ago	6m 47s ↓ 28%	 
academia-app 140375	✓ Success	build_and_ test	 tiger--publishing-feed a8e2ad9 Redeploy rotating feed test with sandbox and improved fallback feed	1mo ago	6m 40s ↓ 29%	 
academia-app 140373	✓ Success	build_and_ test	 bree/PRI-88-hotfix-redis-set 10a463d Hotfix - set variant properly in transient redis setex	1mo ago	6m 43s ↓ 29%	 
academia-app 140372	✓ Success	build_and_ test	 bee--editor-card-add-department 7f853ae Add department on Editor Cards for microsite	1mo ago	6m 45s ↓ 28%	 



## CircleCI RSpec Suite Improvements

### Solutions (September)

- In September we completed several additional projects aimed at reducing the time it takes to run our CircleCI test suite:
  - We switched from using **Bundle and Yarn cache** to our custom built Docker Images that already have the dependencies pre-installed
  - For feature branches the “**Datadog Intelligent Test Runner**” reduces the time by determining what tests it is safe NOT TO RUN, for a given change.
  - We simplified the step that sends a **webhook to a 3rd party service Coverall** (reduction from 1-2 minutes to 3-5 seconds)
  - We switched to using Ruby with YJIT and Jemalloc, and upgraded CI to 3.3





## Docker Image with Dependencies

### Solutions: 2. Bundle & Yarn Cache

- It is now possible to use the academia-docker-images repository to build and push the updated CircleCI Docker image with all of the dependencies pre-installed, using the packages defined by lock files on the Academia App's master branch at the time of the Docker build.
- This works on OS-X (using aws-vault) and on Linux (can't push from devbox)
- **Upside:** build is ~20% faster, now averaging ~8 minutes
- **Downside:** we must support a new project that rebuilds and re-pushes this image daily or weekly (NOTE: *currently in progress*).
  - In the meantime image updates are manual and happen about once a week.






























## Datadog Intelligent Test Runner

### Solutions: 3. Datadog Intelligent Test Runner

- This is also a 3rd party product that automatically tracks interdependencies within our code and tests, and is able to skip tests unaffected by the change.
- This is generally referred to as “test impact analysis”.
  - NOTE: We still run every test on **master** (i.e after the feature is merged)
- **Outcome:** test suite for feature branches now takes between 4 and 6 minutes on average (master is around ~ 6-8)
  - This saves us \$\$ on CircleCI! Additional bonus!
  - NOTE: Rohit will talk more about this next!

A

Project	Status	Workflow	Trigger Event	Start	Duration	Actions
academia-app 142786	✓ Success	build_and_ test	 master e7ceb4f Change table and field names to match Social Corpus DynamoDB tables (#37122)	28m ago	6m 52s ↓ 27%	 
academia-app 142784	✓ Success	build_and_ test	 SOC-2306--social-corpus-precomputations-fix 3cc3a45 Change table and field names to match existing tables	43m ago	6m 24s ↓ 32%	 
academia-app 142782	✓ Success	build_and_ test	 master 761c671 Add Dynamoid models for batch social recommendation calculations (#37063)	2h ago	7m 31s ↓ 20%	 
academia-app 142780	✓ Success	<a href="#">build_and_ test</a>	 SOC-2306--social-corpus-precomputations 6d5b854 Adjust defaults for AcademiaToSocialCorpusOrchestrator	2h ago	4m 45s ↓ 50%	 
academia-app 142779	✓ Success	build_and_ test	 patbl/SRE-2725--better-error-message-when- you-need-to-install-gems 675fe5b DBLab: Print error message sooner when you need to install gems	2h ago	5m 25s ↓ 43%	 
academia-app 142775	✓ Success	build_and_ test	 master 450b20e Merge pull request #37113 from academia-edu/ek/special-user-emilykuhn	2h ago	7m 24s ↓ 21%	 
academia-app 142774	✓ Success	build_and_ test	 patbl/SRE-2725--better-error-message-when- you-need-to-install-gems 247e2d0 DBLab: Print more-helpful error message when you need to install gems	2h ago	4m 35s ↓ 51%	 
academia-app 142773	✓ Success	build_and_ test	 PAC-94/dashboard_rollout 829cd3c fix the fixture	2h ago	5m 2s ↓ 47%	 
academia-app 142772	✓ Success	build_and_ test	 isaac--auto-cfp-search-by-hallucination 0d9ddd8 Switch academia CFP jobs to search	2h ago	7m 43s ↓ 18%	 



## Miscellaneous Changes

### Solutions (Misc)

- We are now ensuring that we use **jemalloc** (a better memory allocation library)
- We are ensuring that we use and enable **YJIT** (Ruby's just-in-time compiler)
- We switched to using **Ruby 3.3** on CI (production & dev are still on 3.2)



## Summary

### Outcome

- We were able to reduce the test suit runtime from **17-18 minutes to 4-6 minutes (branch), 6-8 minutes (master)**.
- That's a **61% decrease** (on master, and even more for branches)
- Who worked on it on the Infrastructure Team:
  - **Rohit** (implemented the Datadog ITR)
  - **Patrick** (sped up the Coverall webhook)
  - **Konstantin** (all other changes)
    - with kind assistance from Nick and Patrick (thanks!)

# A

## Next Steps

- Automate the Docker image with the dependencies built using Codebuild and Terraform so Circle always has fresh dependencies pre-installed
- Observe the test suite for a bit after all of the changes and see if anything is off
- Potential future projects:
  - Bring Datadog ITR in-house (to save \$\$ on Datadog), may not be worth it
  - Reduce the individual spec time
  - Test pushing/pulling the docker image to/from CircleCI instead of AWS
    - Could save up to additional 1-2 minutes
  - Test out BuildKite on our own infrastructure



**Thanks!**