




Turning Kubernetes into a Developer Platform: Lessons from 300 Student Workspaces


Kubernetes-native dev workspaces – for development and training



**Isolated Workspaces.
Instantly. On Kubernetes.**


 Like GitHub Codespaces - self-hosted, and yours.

Architecture



Demo @app.k8shell.dev


1. Admin defines an onboarding rule for a student (GitHub sign-in)
2. Student signs in and provisions a workspace for a course
3. Student opens the workspace from the web terminal, SSH and VS Code and accesses the course materials.



Running in Production

300+

workspaces provisioned
every semester



Machine Learning Labs

CPU-bound

Python ML: training
models, e.g. classifiers



Systems, Network, Dev Labs

I/O-bound

Reverse proxies, network configs,
packages/binary installations

Two Concerns

Security


as few privileges as possible,
as many as the lab needs

Performance

hundreds of workspaces,
one shared cluster


Security: Privilege Levels

As little as possible, as much as the lab needs



Security: Worktrace


Lab containers run, escapes stop: the workspace is stopped at once



Plus: session recording, for audit, not prevention

Security: No Secrets in the Workspace

Credentials stay outside: a Git proxy, short-lived Kubernetes tokens




Performance: CPU


[Full report →](#)

Overcommitted 4.6×, 12-core limit per ML workspace

CPU limits against cores, per node



What capped an ML job when it was busy



588 / 128

cores of CPU limits on the allocatable cores

42 cores

the most ML workspaces used in one minute

2,800


CPU-hours of ML jobs over three months

Performance: Memory


[Full report →](#)

Never overcommitted, 12 GiB limit per ML workspace

Memory limits against memory, per node



Peak and typical memory of each workspace



45%

of its memory the busiest node
used at peak, all workloads

0

nodes ever under
memory pressure

10

out-of-memory kills,
in 5 of 38 workspaces

Join Us



Open Source

AGPL v3.0

<https://github.com/k8shell-io>

Extend it



<https://k8shell.io>

Connect with us



Thank You!

<https://k8shell.io>

Questions? Find me after the talk.