Enhancing Scientific Onboarding in CIRRUS: Cloud-Native Infrastructure for Earth System Scientists

Shilin Chhabra

Mentors: Nick Cote, Kevin Hrpcek, Varsha Banda

***S**ummer **I**nternships in **P**arallel **C**omputational **S**cience (**SIParCS**)*

July 29, 2025

This material is based upon work supported by the NSF National Center for Atmospheric Research, which is a major facility sponsored by the U.S. National Science Foundation under Cooperative Agreement No. 1852977.

Why CIRRUS?

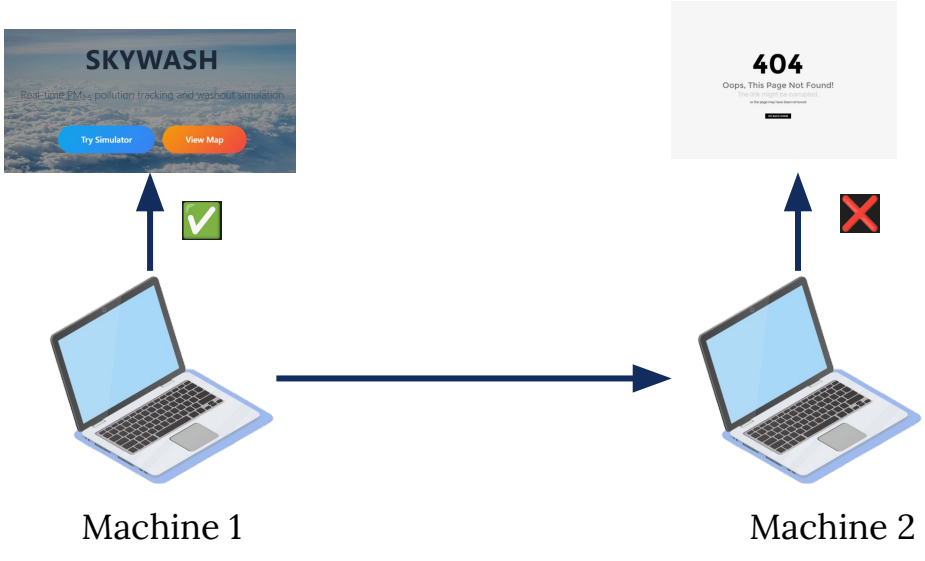

Face roadblocks:

- Library/version conflicts
- Missing GPU/compiler drivers
- OS mismatch
- Manual, unreliable installs

Hours lost rebuilding environments & debugging cryptic errors

Wish: a drop-in container, unchanged, reproducible, shareable

Why CIRRUS?

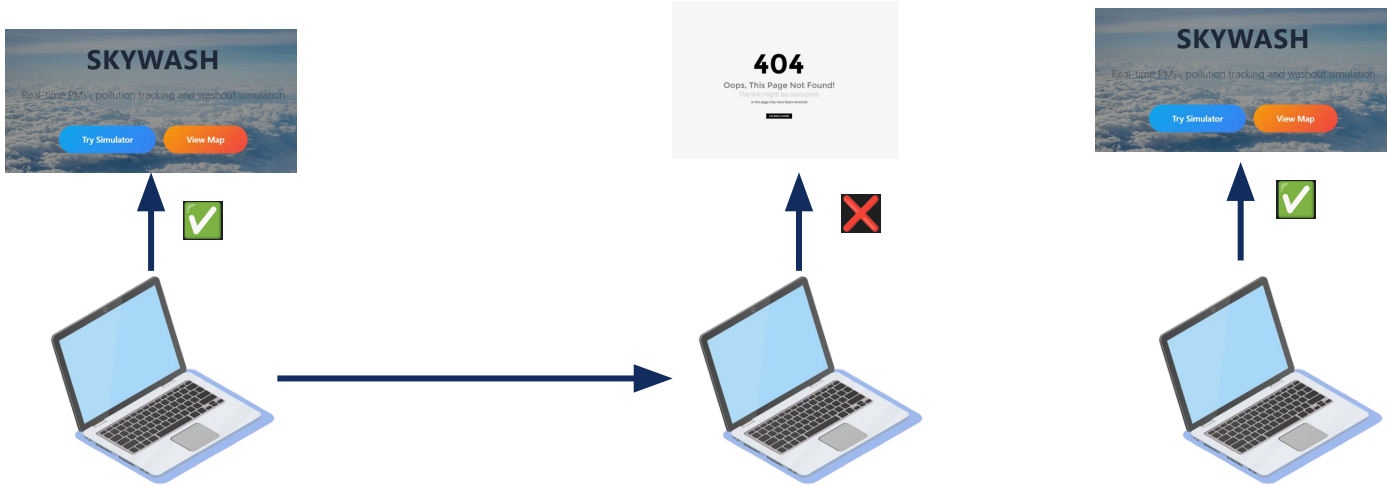

One container that runs anywhere

Access | Share | Reproduce | Scale

Core Components of CIRRUS

Researcher does…

1) Push code to *GitHub*

(one time only)

2) Save *Harbor* robot token as repo secrets

3) Merge deploy YAML / click “Upgrade”

4) Open dashboard

CIRRUS does …

1) Builds the *Docker* image

2) Pushes image *Harbor* registry

3) *Harbor* runs a vulnerability scan; blocks critical CVEs

4) *Kubernetes* + *Argo CD* rolls out the new image and auto-scales pods

5) *OpenBao* Secret Manager injects runtime secrets

Left Column ≈ 2 min of human effort | Right column = full DevOps pipeline on autopilot

Barriers to Adoption

Lengthy narrative **documentation**

Disconnect between the **working example** and **reference docs**

Extensive **security workflow** yet unclear, leaving newcomers **uncertain**

≈ 4 hours for a newcomer to reach a first successful deploy (not a smooth onboarding process)

Why onboarding matters:

Quick, clear onboarding turns interest into adoption, cuts support time, and lets scientists start producing results sooner.

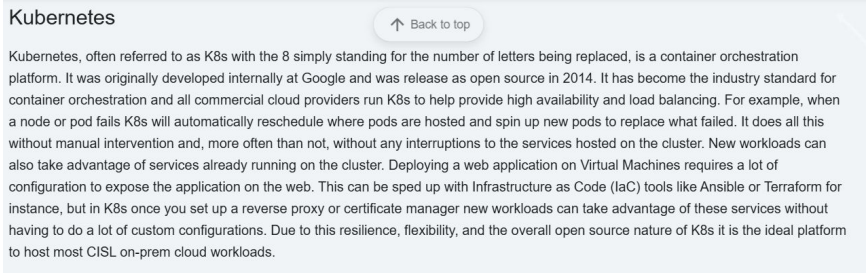

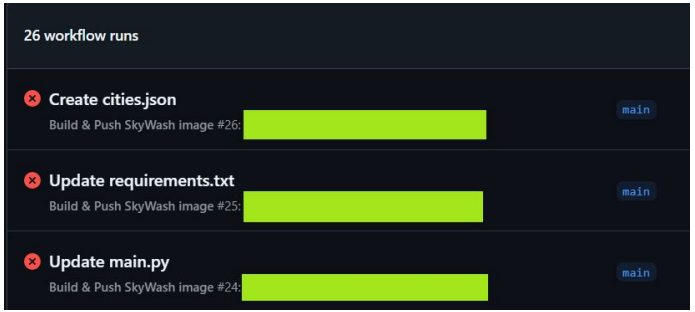

Approach to Improving Onboarding

Aim: *Craft the docs to be easy for a first-time user as well as detailed and coherent enough for an advanced user to make the most of the platform.*

Re-audited CIRRUS docs from a first-time user's view

Filed **600+ improvement issues** and merged suggestions weekly with mentor reviews

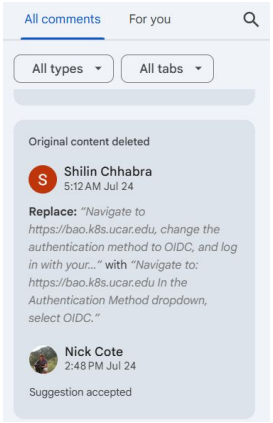

Built & deployed **SkyWash**, exercising containers, CI/CD, Kubernetes, Harbor scans, and OpenBao secrets

Outcome: *Estimated onboarding time* ***slashed from ~4 h to ~1 h***

Skywash Live Demo

Purpose: interactive tool spotlighting air-quality challenges in South & East Asia

Wash-Out Simulator: cloud-seeding model (FastAPI) forecasts post-rain $PM_{2.5}$ / AQI

Visuals: React + Leaflet map & Chart.js bar chart

Deployed end-to-end on CIRRUS
(GitHub → Harbor → K8s → Ingress)

Try it yourself: QR code to live site

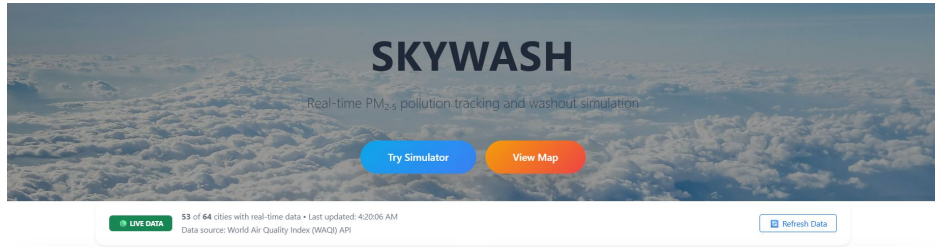

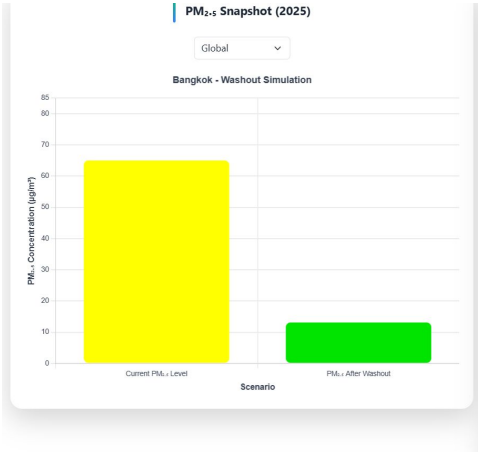

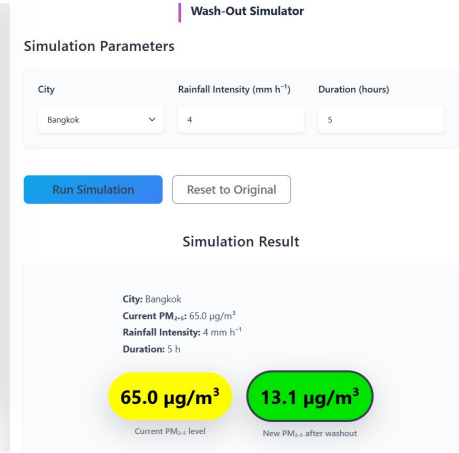

Skywash Deployment Pipeline

Build multi-stage Docker image
Push image the Harbor registry
Trigger Github Actions workflow
Deploy via Kubernetes *(Deployment + Service + Ingress)*
Inject runtime secret with OpenBao
Read/Write data in object storage

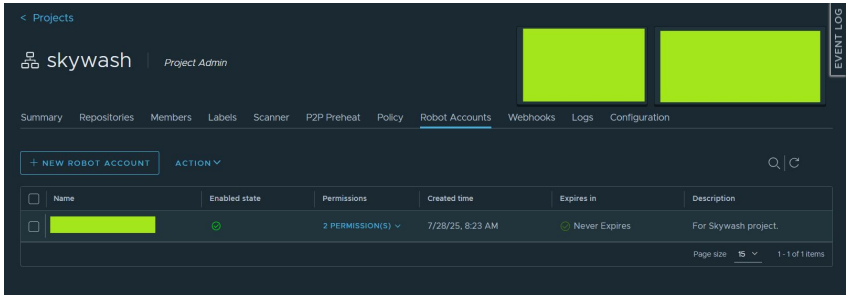

GitHub Actions

Looking Ahead

AI Doc Chatbot for instant, context-aware help

Self-Service Portal to create namespaces, storage, and CI runners in one click

Advanced Tutorials covering GPU, Dask, and Binder workflows

Usage & Cost Dashboards so PIs can track resources in real time

Conclusion

Docs are the gateway:
refined guides move CIRRUS from pilot curiosity to production-ready platform.

Proven path:
Skywash + updates cut estimated first deploy time
4 h → 1 h,

Impact:
faster onboarding drives adoption and guarantees reproducible science on CIRRUS.

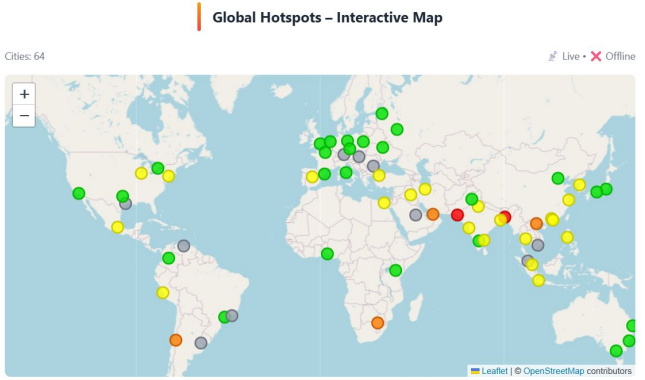

CIRRUS
NSF NCAR Cloud Infrastructure for Remote Research, Universities and Scientists

ACKNOWLEDGMENTS

Thank you to my mentors:

Nick Cote, Kevin Hrpcek, Varsha Banda

Thank you to the internship organizers:

Virginia Do, Jessica Wang, Prashansa Agrawal

Thank you to the organizations that support us:

This material is based upon work supported by the U.S. National Science Foundation National Center for Atmospheric Research, which is a major facility sponsored by the U.S. National Science Foundation under Cooperative Agreement No. 1755088. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U.S. National Science Foundation.

Questions?

chhabra7@wisc.edu