


# Scaling Molecular Software Testing with Self-Hosted GitHub Actions Runners

Ethan Holz, David W.H. Swenson

Ecosystem Infrastructure, Open Molecular Software Foundation, Davis, California 95616



## Introduction

At the OMSF, our projects are primarily hosted on GitHub and make heavy use of the GitHub Actions platform for CI. While this has worked for much of our history, we lacked an **affordable and configurable option** for testing our GPU code paths. Thus, we turned to self-hosted runners to handle this case. As such, **we needed a solution that is flexible, affordable, and built for our community.**

## Requirements

- Needs a GPU available to run GPU code paths
- **Ephemeral** environment
- Easily auditable by our community
- Uses **existing GitHub Actions** workflows
- Well tested
- **Cloud-agnostic**

## Easily Extensible

- Uses a abstract bases classes to utilize inheritance for adding new machine providers
- Built with **cloud-provider SDKs** in mind
- If your provider has a REST API for provisioning a machine, this action can be built to use it
- **MIT licensed** for ease of adoption

## Current Solutions

### `machulav/ec2-github-runner`

- ✓ Uses existing GitHub Actions syntax
- ✓ Used in existing community context
- ✗ Only works on AWS

### `MonolithProjects/ansible-github_actions_runner`

- ✓ Uses Ansible
- ✗ Requires a way to provision infrastructure outside of the pipeline

### `actions/actions-runner-controller`

- ✓ Developed by GitHub
- ✓ Uses a Kubernetes cluster to manage infrastructure
- ✗ Operates on the assumption that you have long-standing infrastructure

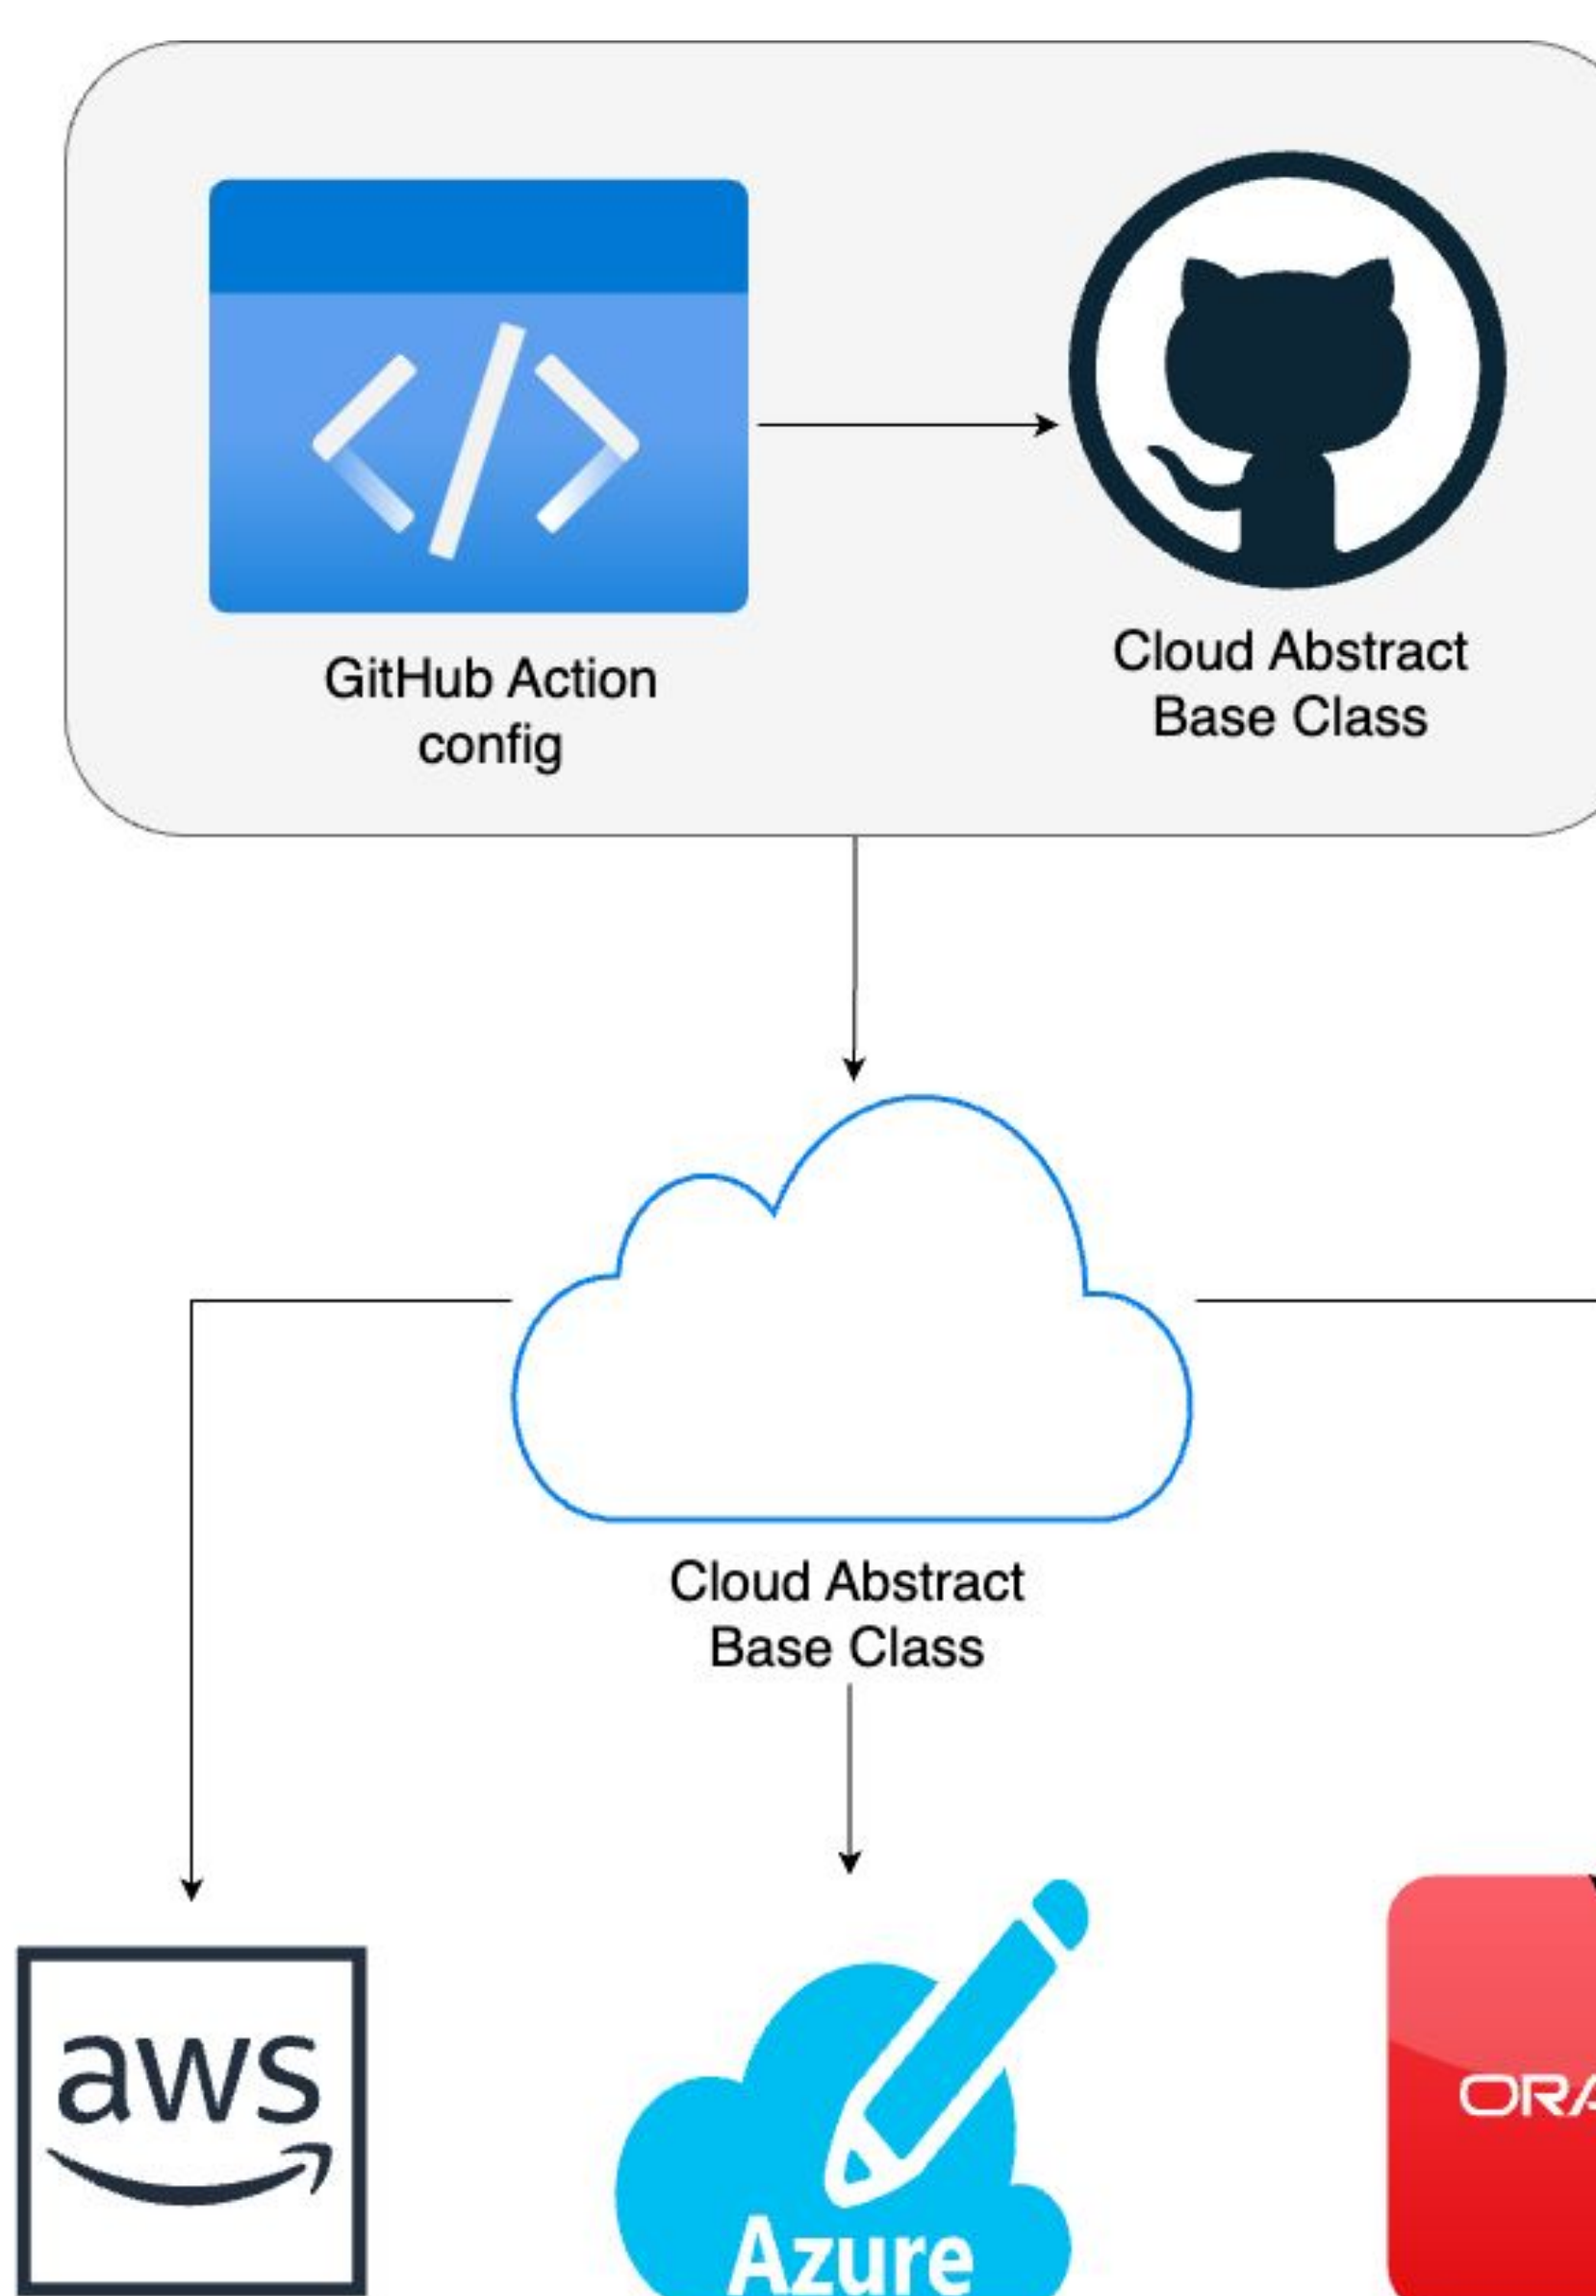
### `myoung34/docker-github-actions-runner`

- ✓ Uses Docker to isolate the runner daemon
- ✗ Disallows the use of Docker-based actions

### GitHub GPU Runners

- ✓ Turn-key solution for using GPU comput on the GitHub platform
- ✗ Only offers a Nvidia T4 GPUs
- ✗ Expensive (~\$4.20/hour)

## Architecture



## Built for the OMSF Community

- Uses **Python** *instead of* Javascript/Typescript due to community preference
- Encourages and requires unit tests for fearless contribution
- Extensible architecture for wherever people have compute
- Uses existing **GitHub Actions syntax** to ease adoption
- Extensible to support multiple cloud providers (but only AWS at this time)

## Where are we now?

- v0.3.0 has been released!
- 150+ commits and counting
- Basic architecture in place for contribution of new cloud providers
- Eliciting community feedback for new features and new providers

## Security

- GitHub **does not** recommend the use of self-hosted runners in public projects for workflows that run on PR due to the potential of running malicious code
- This led to a discussion on secrets management and permissions in our repos such as:
  - Restricting self-hosted runner workflows to not run from forks
  - Utilizing **OpenID Connect (OIDC)** for ephemeral credentials
    - Reduces runs outside of the repo
    - Provides a simpler configuration and setup

## Active Users



Find this project on GitHub!

