

SHIP-HATS Runner Variants

SHIP-HATS Team: Jason Kwan | Leon Leow

GOVTECH
SINGAPORE





STACK MEETUP FOR SGTS: THE ONE ABOUT AI

SPEAKERS & PANNELISTS



WANG ZHEN
Advisor for Moonshot AI &
Founder of JAI



KEVIN AUBRY
CTO, Palo IT



WINSTON TEO
Director / Head of Practice
GovTech



ANI ADHIKARY
Software Engineer (II)
GovTech



VICTOR ONG
Lead Data Scientist
GovTech



SHEIKH MOHD SHAHEED
Senior Solutions Consultant
GovTech

WHO SHOULD ATTEND?

AI Practitioners, Leaders &
Techies with interest in AI,
Software Developers, and
Engineers

Levelling: **All levels**



IN-PERSON MEETUP

MON, 26 MAY 2025
3:00PM – 5:30PM
GOVTECH HIVE

GovTech Hive Open Space (Level 8)
@1 Fusionopolis View,
#08-01 Sandcrawler, 138577

Sign Up: <https://go.gov.sg/sgts-meetup-signup>

Date: 26 May 2025

Time: 3:00 PM to 5:30 PM

Format: In-person

Programme Details

Scaling AI in software delivery

Two Gaping Gaps in Post-AI Software Engineering

Tiny Team x AI: Our Prototype to Production Journey

Data Exploitation with Coding Assistants

Leveraging Agentic AI Tools with Lean Design Thinking

Panel Discussion

Networking & Refreshments



<https://go.gov.sg/sgts-meetup-signup>

Q & A

<https://pigeonhole.at/SHIPHATSRUNNERS>



Agenda Slide

01

Context & Why are introducing Runner Variants

An overview of the context: From community feedback, to our north star and how runners play a part.

02

Our Approach to introducing runner variants

Upcoming variants, SaaS vs Self-hosted, will this result in additional cost?

03

New Runner Variants and how you can take advantage of this

New Runners, Specs, Fair-use

04

Q & A

Share your questions

Why are we introducing SHIP-HATS Runner Variants?

Community Feedback

“Slowness” in Runners

- Insufficient runner specification leading to long build times (> 1 hour)
- Low availability of runners, leading to long queue times (>10 min)
- Large number of jobs and pipeline execution during peak periods (before noon and 5pm)

More Runner Variety

- ARM64 and macOS support

Our North Star

Improve Agency Deployment Frequency

- Remove blockers to frequent deployment
- Speed up builds and tests, reducing time
- Handles complex jobs without failing or timing out
- Enables more frequent and reliable deployments

Reaping Economies of Scale


Encourage more Shared Runner usage

- Reduce toil of self-hosting runners
- Focus on application development

Optimizing Runner Performance in FY24

Category	Changes
Performance Upgrade	GitLab Infrastructure upgrade
	Scaling of runner count during peak periods
	Upgraded Storage Volume and Performance (GP2 -> GP3)
Scaled Resource Management	Auto-scaling of Docker Runner resources based on demand
Increased Observability	Runner performance tracking via StackOps (Elastic)
	Runner Insights Dashboard: SHIP-HATS Docs (Available for all)
Best Practices & Self-help	SHIP-HATS Docs

Why are we introducing SHIP-HATS Runner Variants?



In our approach, we aim to eliminate or reduce bottleneck from onboarding to first deployment and subsequent deployments.

Ultimately, we want to **reduce the time** it takes for project teams using SHIP-HATS to reach a stage where they have the capability to deploy on demand.

Our Approach to Runner variants

Our Aim

As-is:

One-size fits all runner specs, optimized to provide affordable and secure performance while catering to a broad base. Users will self-host if not

To-be:

Multiple Runner types of various specifications and scaling fit for different use-cases

Approach

Build versus Buy:

Hosting by SHIP-HATS vs GitLab

Cost Management:

Can we realistically absorb the costs of hosting runners?

What this means for Users?

More use-case coverage with better performance, reducing need to self-host

Reduced bottlenecks (technical or resource) to deployment

Runner Variant Roadmap

FY25 Q1

Shared Runners (Docker)

Shared Runners (Cstack)

Shared Runners (Windows)

New!

Shared Runners (Docker Med / Large)

Higher Specifications to meet
performance demands

What is in the pipeline FY25

Support for more OS and
deployment environments

Shared Runners (ARM64)

Shared Runners (macOS)

Shared Runners+ (for GCC+) **

Under Consideration

DinD Capability + Privilege **

Higher Specifications for other
runners

** Undergoing security / technical evaluation

Runner Variants & Best Practices

How can you maximize the value of runner variants?

New Docker Runner Variants

Runner Size (runner tag)	vCPU	Memory	Storage	Idle Count (high = shorter queue time, low = possibly longer queue time when idle nodes are all taken)	Ideal For
Small (Default) (ship_docker)	2	8GB	30GB	High	Starting point for any jobs, quick tests, linting, basic builds
Medium (ship_docker_med)	4	16GB	50GB	Medium	Moderately heavy tasks, integration tests
Large (ship_docker_large)	8	32GB	100GB	Low	Resource-intensive builds, long-running jobs, large test suites

Best Practices

Optimize your pipelines before scaling up.

Keep jobs **small** and **modular**.

Use **caching**, **artifacts**, and **parallelization** to reduce time.

Review job logs and metrics to justify moving to larger runners.

Choosing the right runner:

1

Start with **Small** Runners (ship_docker)

- Use this for most jobs such as:
 - Unit tests
 - Code linting/formatting
 - Simple builds
- If your job completes in under 5 minutes, stick with the small runner.

2

If your job consistently takes longer than 5 minutes due to lack of CPU/memory resources

- Consider moving to a **Medium** runner.
- Useful when jobs start timing out or hitting memory/CPU bottlenecks.

3

Move to a **Large** runner only if absolutely necessary. (due to lack of CPU/memory resources)

- Recommended for:
 - Very large builds
 - End-to-end testing with high concurrency
 - Jobs requiring high memory/CPU (e.g. data processing)

Q & A

<https://pigeonhole.at/SHIPHATSRUNNERS>





STACK MEETUP FOR SGTS: THE ONE ABOUT AI

SPEAKERS & PANNELISTS



WANG ZHEN
Advisor for Moonshot AI &
Founder of JAI



KEVIN AUBRY
CTO, Palo IT



WINSTON TEO
Director / Head of Practice
GovTech



ANI ADHIKARY
Software Engineer (II)
GovTech



VICTOR ONG
Lead Data Scientist
GovTech



SHEIKH MOHD SHAHEED
Senior Solutions Consultant
GovTech

WHO SHOULD ATTEND?

AI Practitioners, Leaders &
Techies with interest in AI,
Software Developers, and
Engineers

Levelling: **All levels**



IN-PERSON MEETUP

MON, 26 MAY 2025
3:00PM – 5:30PM
GOVTECH HIVE

GovTech Hive Open Space (Level 8)
@1 Fusionopolis View,
#08-01 Sandcrawler, 138577

Sign Up: <https://go.gov.sg/sgts-meetup-signup>

Date: 26 May 2025

Time: 3:00 PM to 5:30 PM

Format: In-person

Programme Details

Scaling AI in software delivery

Two Gaping Gaps in Post-AI Software Engineering

Tiny Team x AI: Our Prototype to Production Journey

Data Exploitation with Coding Assistants

Leveraging Agentic AI Tools with Lean Design Thinking

Panel Discussion

Networking & Refreshments



<https://go.gov.sg/sgts-meetup-signup>



STACK MEETUP FOR SGTS: THE ONE ABOUT AI

SPEAKERS & PANNELISTS



WANG ZHEN
Advisor for Moonshot AI &
Founder of JAI



KEVIN AUBRY
CTO, Palo IT



WINSTON TEO
Director / Head of Practice
GovTech



ANI ADHIKARY
Software Engineer (II)
GovTech



VICTOR ONG
Lead Data Scientist
GovTech



SHEIKH MOHD SHAHEED
Senior Solutions Consultant
GovTech

WHO SHOULD ATTEND?

AI Practitioners, Leaders &
Techies with interest in AI,
Software Developers, and
Engineers

Levelling: **All levels**



IN-PERSON MEETUP

MON, 26 MAY 2025
3:00PM – 5:30PM
GOVTECH HIVE

GovTech Hive Open Space (Level 8)
@1 Fusionopolis View,
#08-01 Sandcrawler, 138577

Sign Up: <https://go.gov.sg/sgts-meetup-signup>

Date: 26 May 2025

Time: 3:00 PM to 5:30 PM

Format: In-person

Programme Details

Scaling AI in software delivery

Two Gaping Gaps in Post-AI Software Engineering

Tiny Team x AI: Our Prototype to Production Journey

Data Exploitation with Coding Assistants

Leveraging Agentic AI Tools with Lean Design Thinking

Panel Discussion

Networking & Refreshments



<https://go.gov.sg/sgts-meetup-signup>

Reach out to us



Give us Feedback

<https://go.gov.sg/sgts-fb>



Join our community

<https://go.gov.sg/ship-hats-tele>

GOVTECH
SINGAPORE

