Google

An Update on LLVM Premerge Testing

Core Developer Software Foundations

aidengrossman@google.com

April 2025

The New Beta System

Beta Premerge System Status

- Both old and new systems are running in parallel
- But the new system marks all jobs as passing
- Sorting out a few stability problems
- Launch expected in Q2

Beta System Features

- Autoscaling means more machines at peak load (workday afternoons Pacific time), meaning better latency and typically <20 minutes of queue time.
- Analytics
- Oncall team support during working hours
- Dedicated engineering staff for maintenance and improvements

Metrics

- Displays key metrics that have significantly impacted premerge functionality in the past.
- Metrics is coupled with alerting to notify an on-call rotation when things go awry.

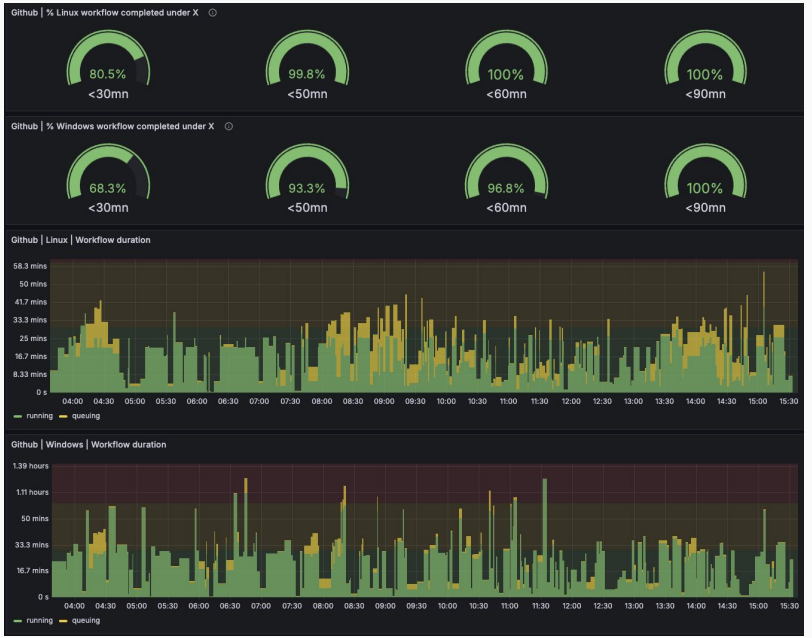

Beta System Features/Improvements

Done:

- Computing what projects to test is now much simpler and unit tested.

Exploring:

- Improving node/container setup times.
- Full reproduction instructions using containers.
- Faster toolchain on Windows?

Beta System Launch Q2

- Launch in this case means marking the new system premerge tests as authoritative, meaning a failing test will report as failed to GitHub
- At this time we will also turn down the old presubmit infrastructure

Testing Resources Prioritization

- There’s always a trade-off between test coverage and premerge latency
- More test coverage means more premerge latency
- Google will do its best to deliver a high-reliability and high-performance system
- But beyond that, these tradeoffs are best made by the community
- We will defer coverage decisions to the Infrastructure Area Leads

Credits

From within Google:

- Aiden Grossman
- Caroline Tice
- Guillaume Chatelet
- Lucile Rose Nihlen
- Nathan Gauër

And the following OSS contributors:

- David Spickett
- Tom Stellard

Questions?