

# GDB tests, CI & Buildbot BoF

Sergio Durigan Junior  
sergiodj@redhat.com

# License

- ▶ License: Creative Commons Attribution 4.0 International License (CC-BY-4.0)
- ▶ <https://creativecommons.org/licenses/by/4.0/>

## How was it?

- ▶ GDB Buildbot started in 2015 as a personal project.

## How was it?

- ▶ GDB Buildbot started in 2015 as a personal project.
- ▶ We just had 2 machines serving 4 Fedora x86\_64 workers at the time. And no *try builds*!

## How was it?

- ▶ GDB Buildbot started in 2015 as a personal project.
- ▶ We just had 2 machines serving 4 Fedora x86\_64 workers at the time. And no *try builds*!
- ▶ Initially it stored the test results in a git repository. This proved too inefficient over time...

## And now?

- ▶ The master runs in a dedicated VM at **OSCI** (**O**pen **S**ource **C**ommunity **I**nfrastructure).

## And now?

- ▶ The master runs in a dedicated VM at **OSCI** (**O**pen **S**ource **C**ommunity **I**nfrastructure).
- ▶ Most of our builders support *try builds*!

## And now?

- ▶ The master runs in a dedicated VM at **OSCI** (**O**pen **S**ource **C**ommunity **I**nfrastructure).
- ▶ Most of our builders support *try builds*!
- ▶ **14** workers (**11** machines):



## And now?

- ▶ The master runs in a dedicated VM at **OSCI** (**O**pen **S**ource **C**ommunity **I**nfrastructure).
- ▶ Most of our builders support *try builds*!
- ▶ **14** workers (**11** machines):
  - ▶ Sergio (Red Hat): **2** machines (Fedora x86\_64)
  - ▶ Alan Hayward (ARM): **2** machines (Ubuntu ARM 32 and 64)
  - ▶ Rainer Orth (CeBiTec.Uni-Bielefeld.DE): **2** machines (Solaris amd64 and sparcv9)
  - ▶ David Edelsohn: **3** machines (RHEL 7.1 s390x, AIX POWER8 and Debian Jessie s390x)
  - ▶ Edjunior Machado: **1** machine (CentOS 7 PPC64LE)
  - ▶ Mark Wielaard: **1** machine (Fedora s390x)

## And now?

- ▶ The master runs in a dedicated VM at **OSCI** (**O**pen **S**ource **C**ommunity **I**nfrastructure).
- ▶ Most of our builders support *try builds*!
- ▶ **14** workers (**11** machines):
  - ▶ Sergio (Red Hat): **2** machines (Fedora x86\_64)
  - ▶ Alan Hayward (ARM): **2** machines (Ubuntu ARM 32 and 64)
  - ▶ Rainer Orth (CeBiTec.Uni-Bielefeld.DE): **2** machines (Solaris amd64 and sparcv9)
  - ▶ David Edelsohn: **3** machines (RHEL 7.1 s390x, AIX POWER8 and Debian Jessie s390x)
  - ▶ Edjunior Machado: **1** machine (CentOS 7 PPC64LE)
  - ▶ Mark Wielaard: **1** machine (Fedora s390x)
- ▶ Test results are stored directly on-disk, and “garbage-collected” every week (tests older than 4 months are deleted).

How does it work?



# Notifications

- ▶ To *gdb-testers*: whenever we detect a possible regression in an upstream commit.

# Notifications

- ▶ To *gdb-testers*: whenever we detect a possible regression in an upstream commit.
- ▶ To the *author*: on *try builds*, or when his/her commit broke GDB.

# Notifications

- ▶ To *gdb-testers*: whenever we detect a possible regression in an upstream commit.
- ▶ To the *author*: on *try builds*, or when his/her commit broke GDB.
- ▶ To *gdb-patches*: when a commit breaks GDB.

# Notifications

- ▶ To *gdb-testers*: whenever we detect a possible regression in an upstream commit.
- ▶ To the *author*: on *try builds*, or when his/her commit broke GDB.
- ▶ To *gdb-patches*: when a commit breaks GDB.
- ▶ Breakage notifications are usually reliable. Regression notifications are not (just look at *gdb-testers*).

# Problems and challenges

- ▶ Racy testcases. Perhaps the most difficult/persistent problem?



# Problems and challenges

- ▶ Racy testcases. Perhaps the most difficult/persistent problem?
- ▶ Lots of test messages are non-unique. This makes it really hard to compare test results and find regressions.

# Problems and challenges

- ▶ Racy testcases. Perhaps the most difficult/persistent problem?
- ▶ Lots of test messages are non-unique. This makes it really hard to compare test results and find regressions.
- ▶ Better way to store and retrieve test results (current way is “enough” for what we need, but it can certainly be improved).

# Problems and challenges

- ▶ Racy testcases. Perhaps the most difficult/persistent problem?
- ▶ Lots of test messages are non-unique. This makes it really hard to compare test results and find regressions.
- ▶ Better way to store and retrieve test results (current way is “enough” for what we need, but it can certainly be improved).
- ▶