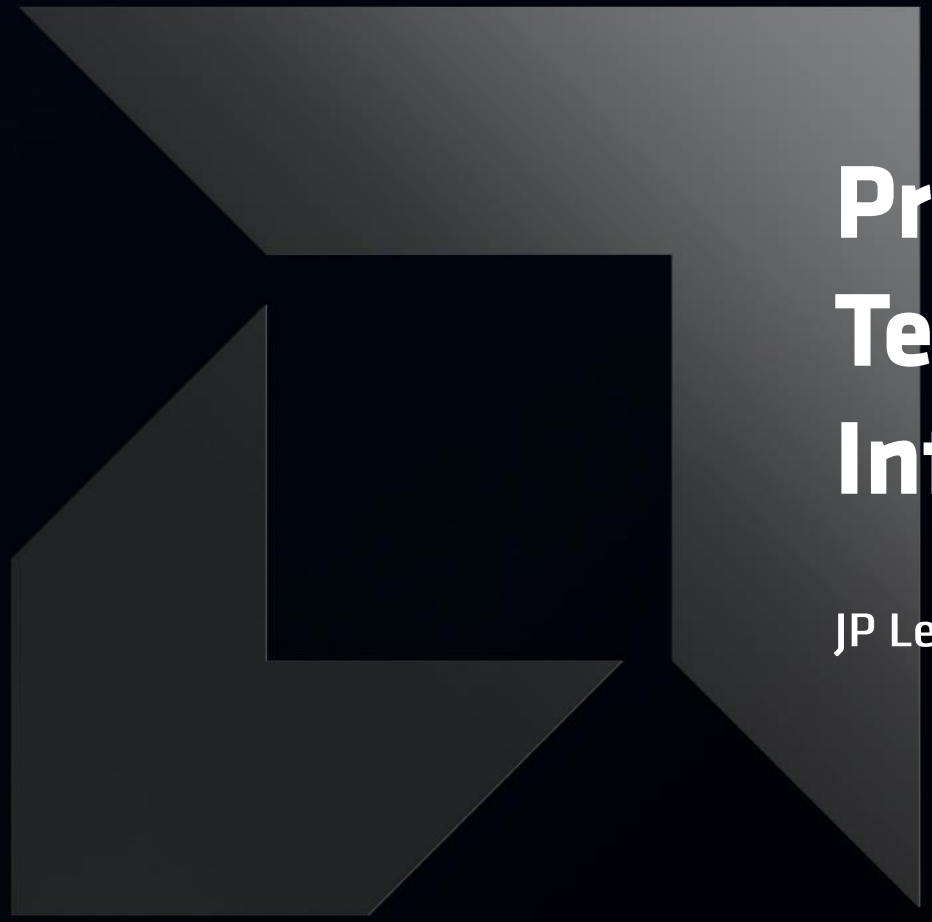

Programming is fun; Testing is needed; Infra is ...

JP Lehr

AMD

together we advance_

WHY TEST UPSTREAM LLVM

AMD ROCm

Support Upstream Developers

Guard Downstream ROCm™

LLVM TESTING LANDSCAPE

GitHub

- Used for GitHub actions
- Hosted on GitHub or self-hosted

Buildkite

- Used for pre-commit testing
- Hosted on GKE

Buildbot

- Used for post-commit testing
- Hosted everywhere

In this talk: Buildbot

TERMINOLOGY

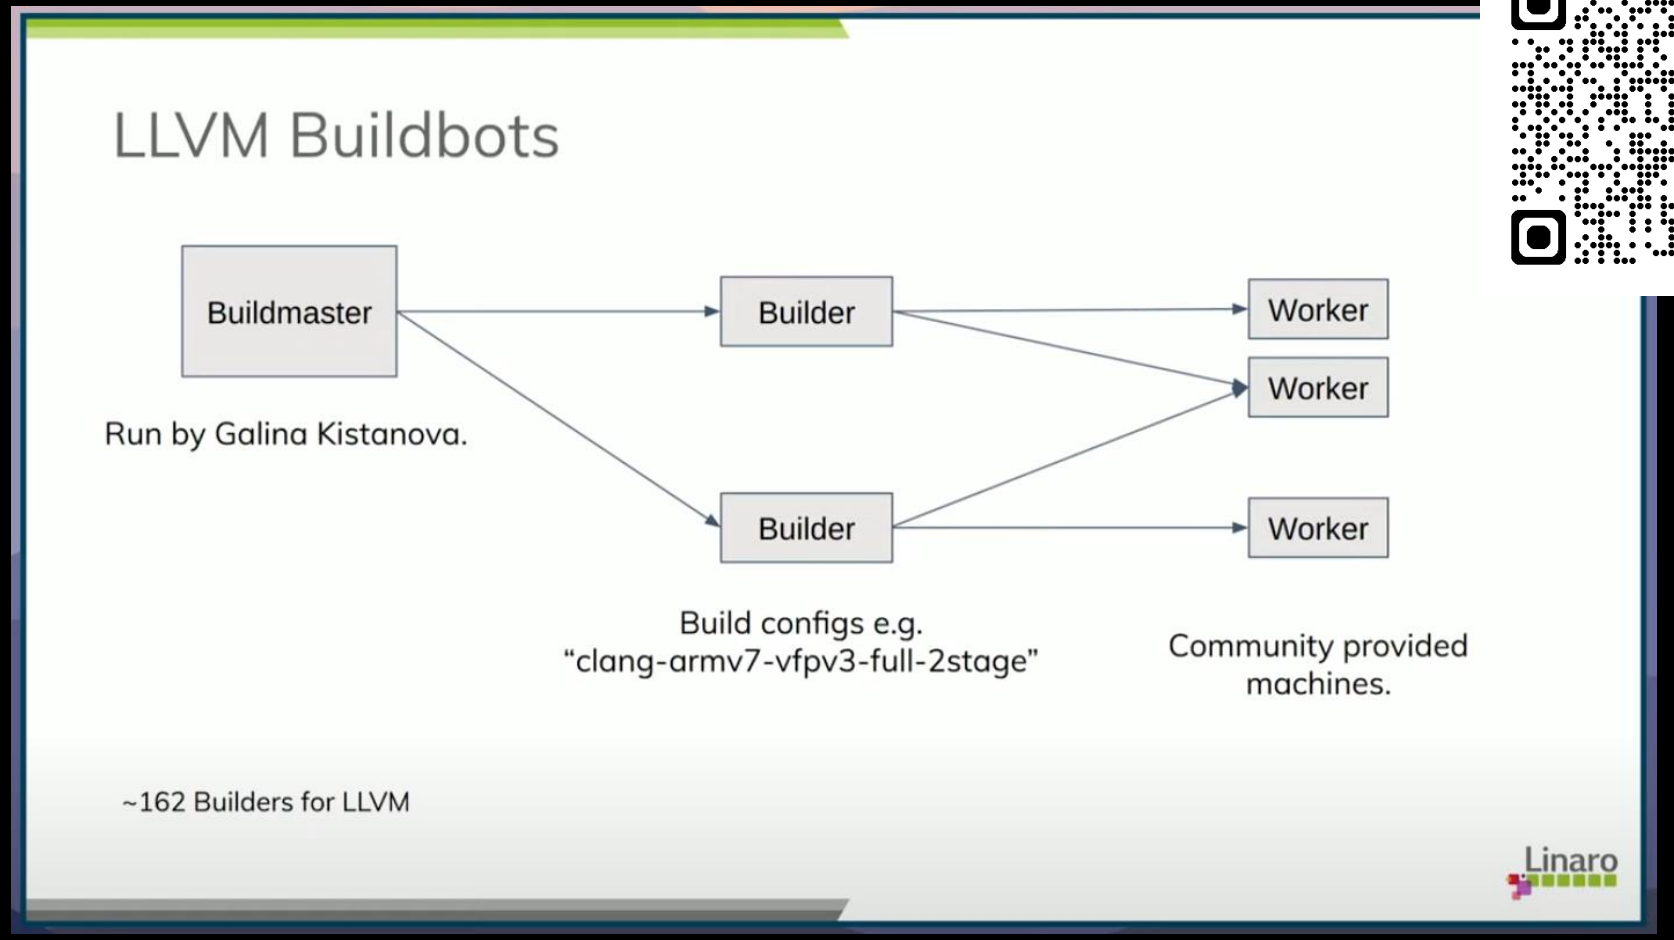

Talk by David Spickett: **2022 LLVM Dev Mtg: What does it take to run LLVM Buildbots?**

TERMINOLOGY

AnnotatedBuilder

- Build-your-own builder
- Script defines what the builder is doing
- Not tailored for any particular purpose
- Examples: HIP, libc, ...

OpenMPBuilder

- Off-the-shelve builder
- OpenMPBuilder defines what builder is doing
- Tailored to building OpenMP®
- Always adds the check-openmp target

```
{'name' : "amdgpu-offload-ubuntu-22-cmake-build-only",
'tags'  : ["openmp"],
'workernames' : ["rocm-docker-ubu-22"],
'builddir': "amdgpu-offload-ubuntu-22-cmake-build-only",
'factory' : AnnotatedBuilder.getAnnotatedBuildFactory(
                depends_on_projects=["llvm", "clang", "lld", "compiler-rt", "libcxx", "libcxxabi", "openmp", "offload", "libunwind"],
                script="amdgpu-offload-cmake.py",
                checkout_llvm_sources=True,
                script_interpreter=None
            )},
```

Maintained in llvm-zorg repository

CURRENT BUILDBOT FLEET

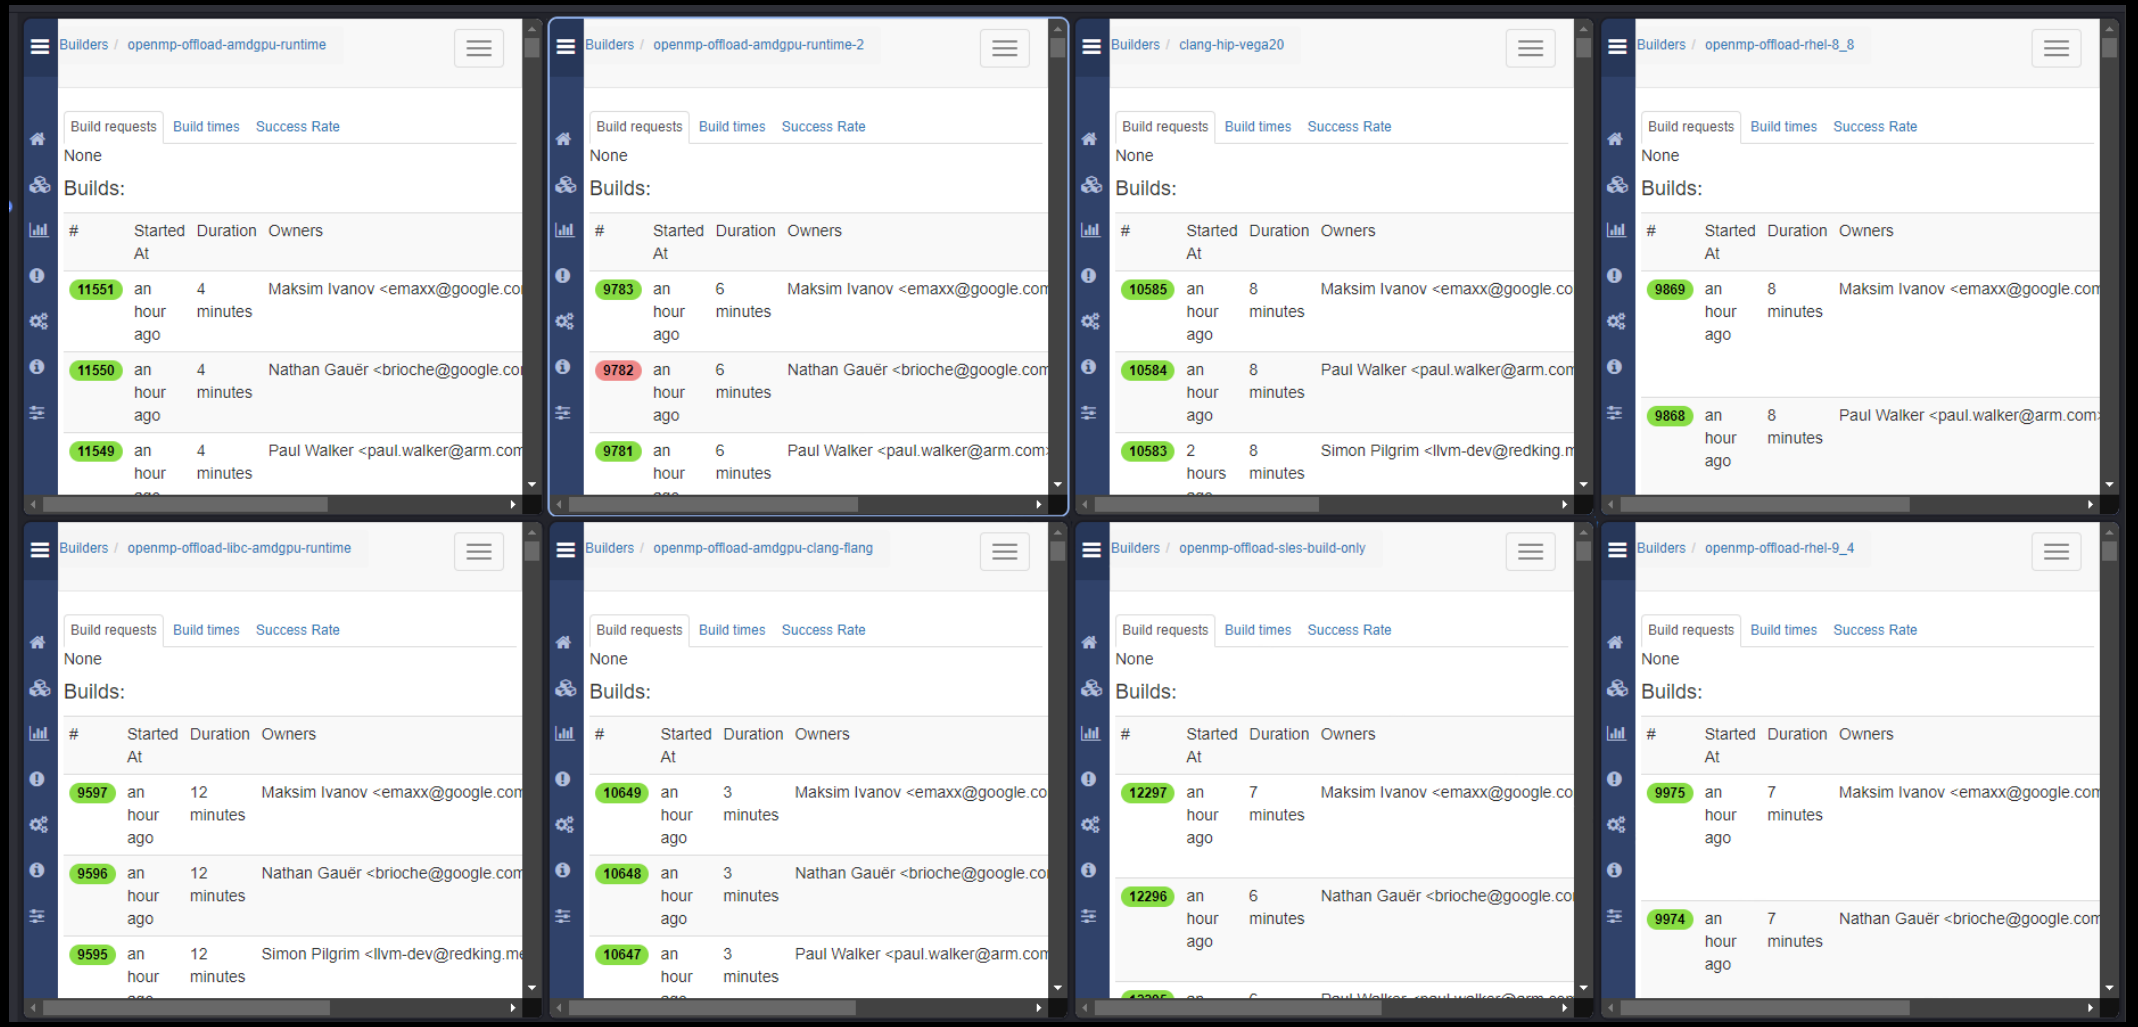

HISTORY

The inherited time

The intermediate time

The current time

The (bright) future

HISTORY

The inherited time

The intermediate time

The current time

The (bright) future

THE INHERITED TIME

- Three machines: all running bare metal
- Three machines: all w/ different usernames and different systems
- Partly unclear (and undocumented) setup steps
- Existing documentation only available internally

- OpenMP® builder uses OpenMPBuilder
- HIP builder uses AnnotatedBuilder

- Changes to llvm-zorg automatically pushed to staging area every 2 hours
- … unless there was a problem

PDF only

THE INHERITED TIME

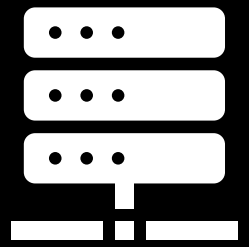

clang-hip-vega20

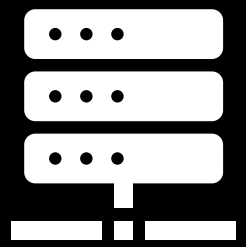

openmp-offload-amdgpu-runtime

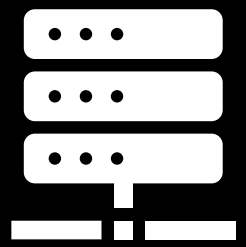

openmp-offload-amdgpu-runtime-experimental

THE INHERITED TIME

llvm-zorg

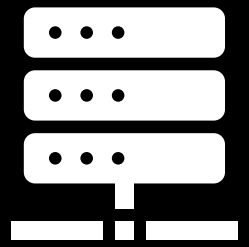

clang-hip-vega20

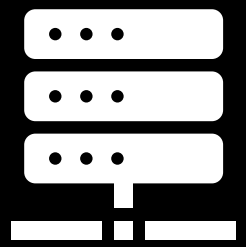

openmp-offload-amdgpu-runtime

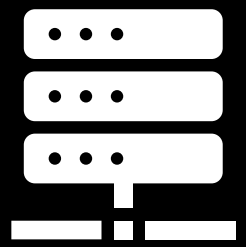

openmp-offload-amdgpu-runtime-experimental

THE INHERITED TIME

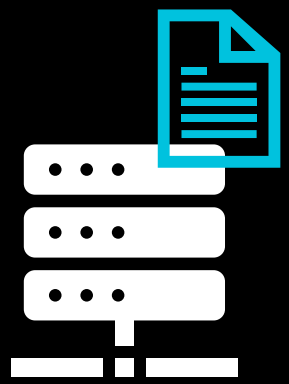

clang-hip-vega20

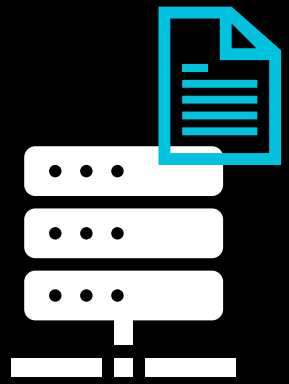

openmp-offload-amdgpu-runtime

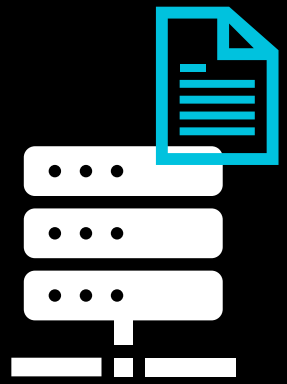

openmp-offload-amdgpu-runtime-experimental

THE INHERITED TIME

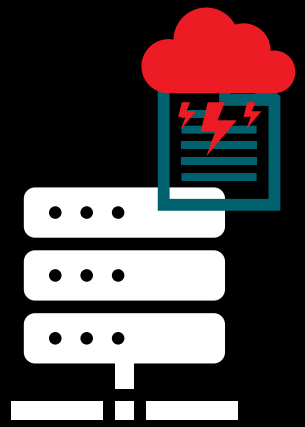

clang-hip-vega20

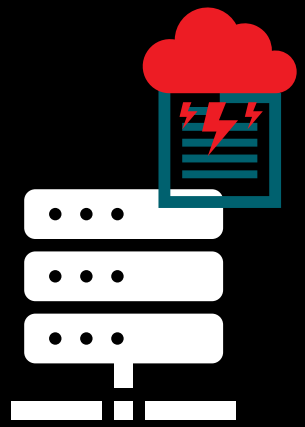

openmp-offload-amdgpu-runtime

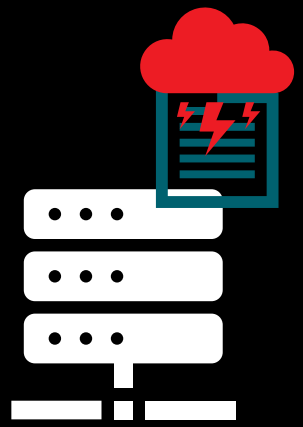

openmp-offload-amdgpu-runtime-experimental

THE INHERITED TIME

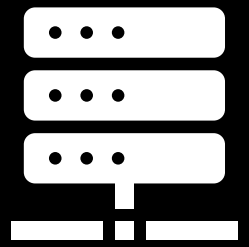

clang-hip-vega20

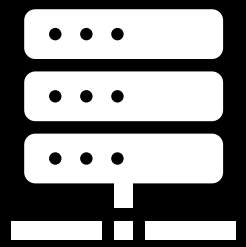

openmp-offload-amdgpu-runtime

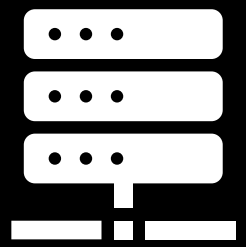

openmp-offload-amdgpu-runtime-experimental

HISTORY

The inherited time

The intermediate time

The current time

The (bright) future

HISTORY

The inherited time

The intermediate time

The current time

The (bright) future

THE INTERMEDIATE TIME

- Additional test-targets for one of two OpenMP® builders to cover libc on AMDGPU
- Patches to OpenMPBuilder to allow additional lit arguments
- Added containers for RHEL 8, RHEL 9 and SLES 15 builders
- Setting up new builders manual process (documentation only internally avail)

PDF only

THE INTERMEDIATE TIME

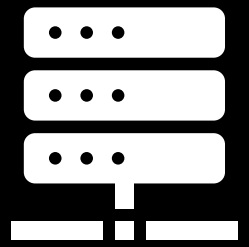

clang-hip-vega20

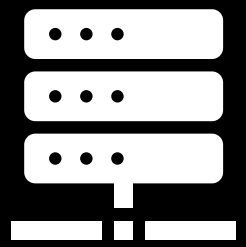

openmp-offload-amdgpu-runtime

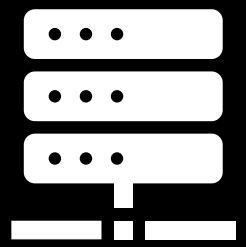

openmp-offload-amdgpu-runtime-experimental

THE INTERMEDIATE TIME

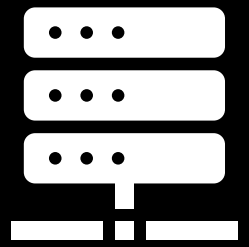

clang-hip-vega20

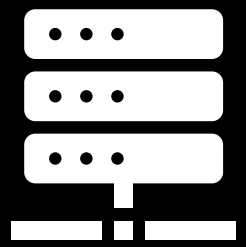

openmp-offload-amdgpu-runtime

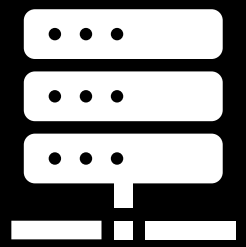

openmp-offload-libc-amdgpu-runtime

THE INTERMEDIATE TIME

openmp-offload-sles-build-only

openmp-offload-rhel-8_8

openmp-offload-rhel-9_4

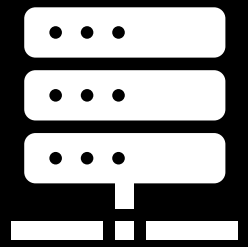

clang-hip-vega20

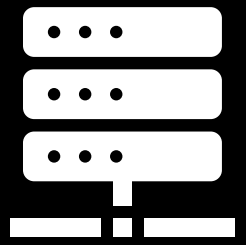

openmp-offload-amdgpu-runtime

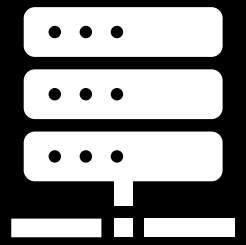

openmp-offload-amdgpu-clang-flang

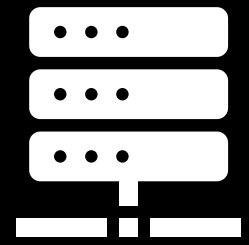

openmp-offload-libc-amdgpu-runtime

THE INTERMEDIATE TIME

openmp-offload-sles-build-only

openmp-offload-rhel-9_4

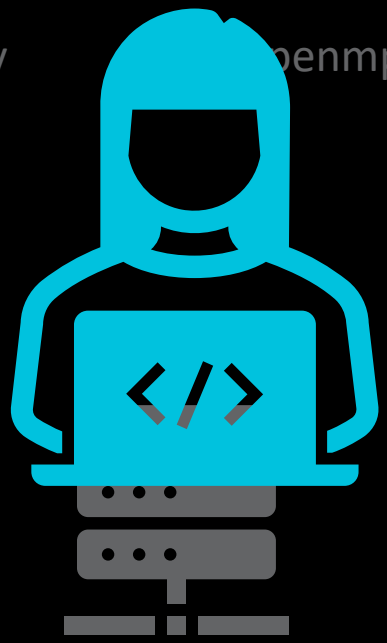

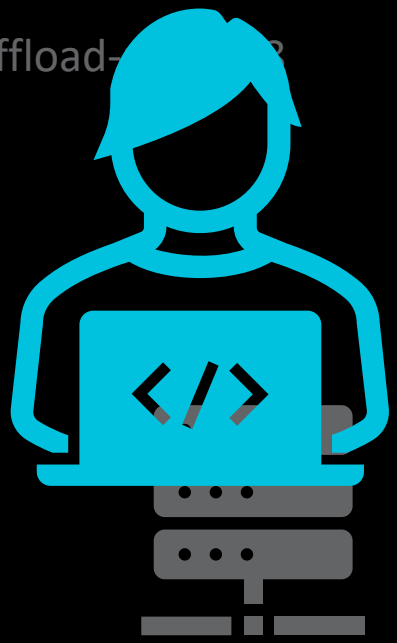

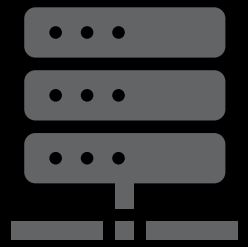

clang-hip-vega20

openmp-offload-amdgpu-runtime

openmp-offload-amdgpu-clang-flang

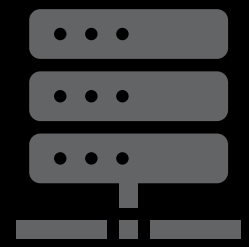

openmp-offload-libc-amdgpu-runtime

HISTORY

The inherited time

The intermediate time

The current time

The (bright) future

HISTORY

The inherited time

The intermediate time

The current time

The (bright) future

THE CURRENT TIME

- Still mix of bare metal and containerized builders
- Struggling with flakiness of some OpenMP®/GPU tests
- Most actual (not flaky) problems are during build

What it should be

- All builders containerized
- All machine setup and builder deployment automated via ansible
- Publicly available Dockerfiles for all builders
- Rely on CMake cache files in-tree for build configuration
- All builders based on AnnotatedBuilder

PDF only

THE CURRENT TIME

openmp-offload-sles-build-only

openmp-offload-rhel-8_8

openmp-offload-rhel-9_4

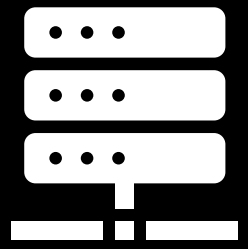

clang-hip-vega20

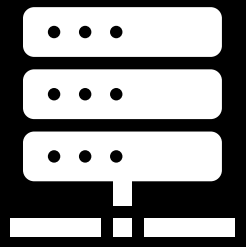

openmp-offload-amdgpu-runtime

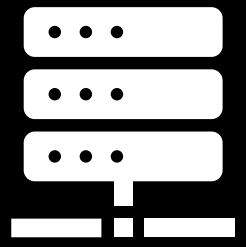

openmp-offload-amdgpu-clang-flang

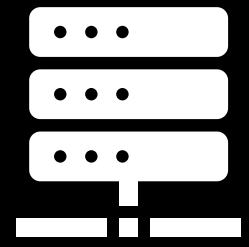

openmp-offload-libc-amdgpu-runtime

THE CURRENT TIME

```
amdkcl.ko:
Running module version sanity check.
 - Original module
   - This kernel never originally had a module by this name
 - Installation
   - Installing to /lib/modules/5.15.0-130-generic/updates/dkms/

amd-sched.ko:
Running module version sanity check.
 - Original module
   - This kernel never originally had a module by this name
 - Installation
   - Installing to /lib/modules/5.15.0-130-generic/updates/dkms/

amddrm_ttm_helper.ko:
Running module version sanity check.
 - Original module
   - This kernel never originally had a module by this name
 - Installation
   - Installing to /lib/modules/5.15.0-130-generic/updates/dkms/

amddrm_bu
Running m
 - Origin
   - This
 - Instal
   - Inst

amdxcp.ko
Running m
 - Origin
   - This
 - Instal
   - Inst

depmod...
update-in
W: Possib
W: Possib
W: Possib
W: Possib
W: Possib
W: Possib
W: Possib
W: Possib
```

```
W: Possible missing firmware /lib/firmware/ast_dp501_fw.bin for module ast
Scanning processes...

Scanning candidates...

Scanning processor microcode...

Scanning linux images...


Restarting services...
 /etc/needrestart/restart.d/systemd-manager
 systemctl restart cron.service irqbalance.service packagekit.service polkit.service rsyslog.service ssh.service systemd-journald.service systemd-networkd.service systemd-res
olved.service systemd-timesyncd.service systemd-udevd.service udisks2.service upower.service
janplehr@r15:~/git/apps/openmp-ci/buildbots$ ./connect_bot -d rack-178
ssh: connect to host [redacted] port 22: Connection timed out
janplehr@r15:~/git/apps/openmp-ci/buildbots$ ./connect_bot -d rack-178
ssh: connect to host [redacted] port 22: Connection timed out
janplehr@r15:~/git/apps/openmp-ci/buildbots$ ./connect_bot -d rack-178
ssh: connect to host [redacted] port 22: Connection timed out
janplehr@r15:~/git/apps/openmp-ci/buildbots$ ./connect_bot -d rack-178
ssh: connect to host [redacted] port 22: Connection timed out
janplehr@r15:~/git/apps/openmp-ci/buildbots$ ./connect_bot -d rack-178
ssh: connect to host [redacted] port 22: Connection timed out
```

THE CURRENT TIME

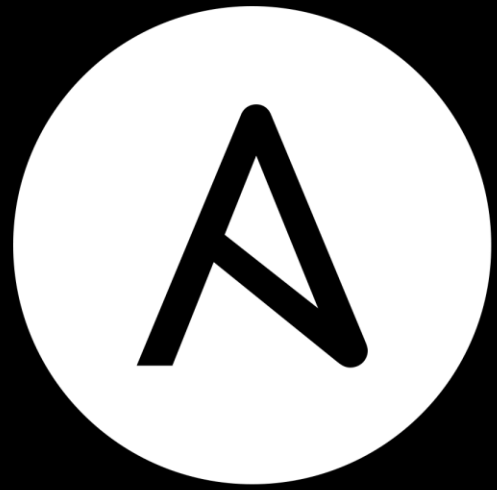

THE CURRENT TIME

openmp-offload-sles-build-only

openmp-offload-rhel-8_8

openmp-offload-rhel-9_4

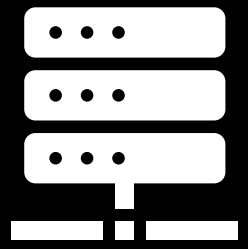

clang-hip-vega20

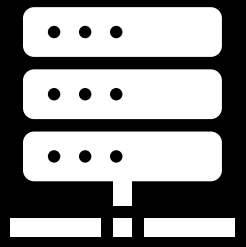

openmp-offload-amdgpu-runtime

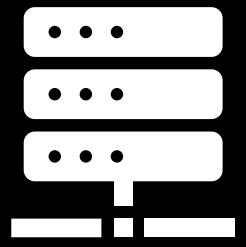

openmp-offload-amdgpu-clang-flang

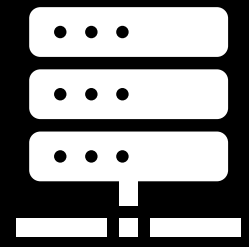

openmp-offload-libc-amdgpu-runtime

THE CURRENT TIME

openmp-offload-sles-build-only

openmp-offload-rhel-8_8

openmp-offload-rhel-9_4

clang-hip-vega20

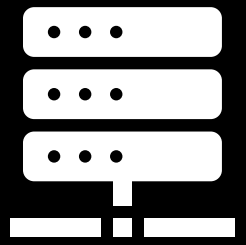

openmp-offload-amdgpu-runtime

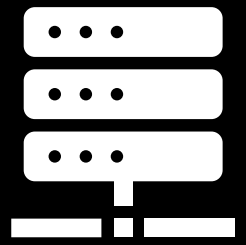

openmp-offload-amdgpu-clang-flang

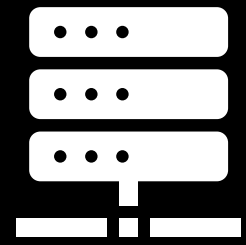

openmp-offload-libc-amdgpu-runtime

HISTORY

The inherited time

The intermediate time

The current time

The (bright) future

HISTORY

The inherited time

The intermediate time

The current time

The (bright) future

THE (BRIGHT) FUTURE

Ubuntu-22-fast

RHEL-8-fast

RHEL-9-fast

SLES-15-fast

Post Commit

Ubuntu-22-slow

RHEL-8-slow

RHEL-9-slow

SLES-15-slow

THE (BRIGHT) FUTURE

kubernetes

THE (BRIGHT) FUTURE

Pre Commit Build

kubernetes

THE (BRIGHT) FUTURE

Pre Commit Build

Pre Commit Test on GPU

kubernetes

THE BRIGHT FUTURE

Add pre-commit testing

- Add self-hosted runners to GitHub for pre-commit building and on-GPU testing
- Based on the same container as post-commit buildbots
- Plan is to use the same in-tree CMake config as in post-commit buildbots

Expand test coverage

- Add more compile targets (e.g., HIP programs or ROCm™ libraries) to post-commit testing
- Likely going to have fast and slow bots in post-commit

PDF only

LESSONS LEARNED/BEST PRACTICES

- Running a single (inherited) buildbot is easy enough and well documented
- Testing a buildbot locally can now be easily done (Thank you!)
- Don’t forget to enable mail-send on build-fail in buildbot/osuosl/master/config/status.py

Builders should be easily reproducible locally

- Environment (e.g., via containers)
- Build config (e.g., via CMake cache file)

Running a fleet of buildbot is more complex. Because you are managing a fleet, not because of buildbot.

- Deployment should be automated: serves as documentation and eases re-deployment

LESSONS LEARNED/BEST PRACTICES

- Running a single (inherited) buildbot is easy enough and well documented
- Testing a buildbot locally can now be easily done (Thank you!)
- Don’t forget to enable mail-send on build-fail in buildbot/osuosl/master/config/status.py

Builders should be easily reproducible locally

- Environment (e.g., via containers)
- Build config (e.g., via CMake cache file)

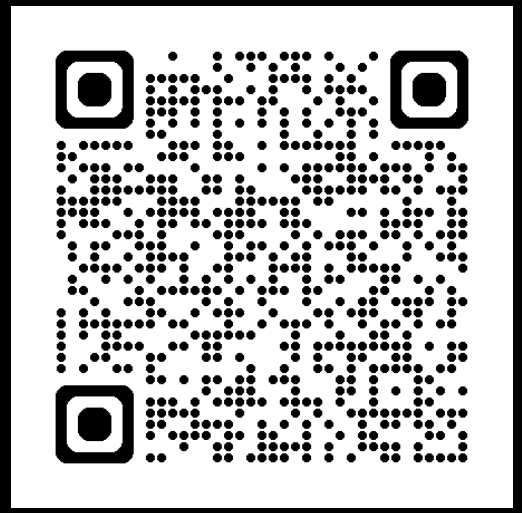

Running a fleet of buildbot is more complex. Because you are managing a fleet, not because of buildbot.

- Deployment should be automated: serves as documentation and eases re-deployment

LESSONS LEARNED/BEST PRACTICES

- Running a single (inherited) buildbot is easy enough and well documented
- Testing a buildbot locally can now be easily done (Thank you!)
- Don’t forget to enable mail-send on build-fail in buildbot/osuosl/master/config/status.py

Builders should be easily reproducible locally

- Environment (e.g., via containers)
- Build config (e.g., via CMake cache file)

Running a fleet of buildbot is more complex. Because you are managing a fleet, not because of buildbot.

- Deployment should be automated: serves as documentation and eases re-deployment

We are hiring.

Disclaimer

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

THIS INFORMATION IS PROVIDED ‘AS IS.” AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.


AMD