


Mesa Continuous Integration at Intel


Mark Janes
Clayton Craft



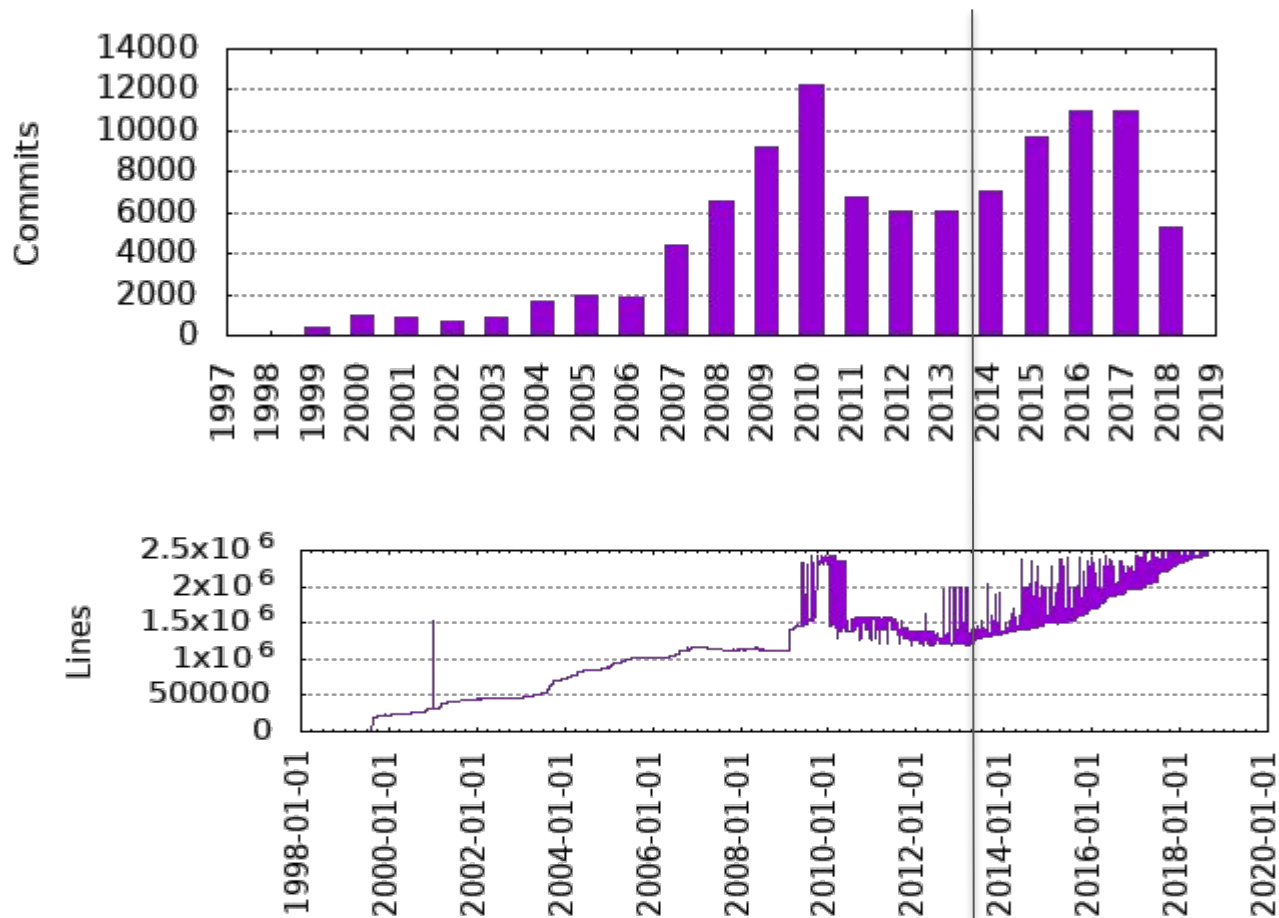
Why is continuous integration valuable for Mesa?

Summary of Mesa CI at Intel


- ~200 systems with full Intel hardware coverage going back to 2007
- dEQP, Piglit, VulkanCTS, Crucible, OpenGL CTS, OpenGL ES CTS
- Build-tests for Android and non-Intel platforms
- Millions of tests per run for every commit
- Target execution time of 30min.
- False positives ~0.0001%
- Generates performance trend lines for common benchmarks
- Open source CI implementation (https://gitlab.freedesktop.org/Mesa_CI)



Phoronix: making more whitespace commits



CI Builds vs Commits since 8/2014, Developers with > 30 commits



Reverts and Fixes tags

```
6521d4a659b911bb86d979564de03665616a671e
Author:      Samuel Pitoiset <samuel.pitoiset@gmail.com>
Commit:     Samuel Pitoiset <samuel.pitoiset@gmail.com>
```

```
Revert "radv: Optimize rebinding the same descriptor set"
```

```
This introduces random GPU hangs on Vega, at least on Linux.
```

```
This reverts commit
```

```
02a43edf186cb9998741ba765cb948bb238a122d.
```

```
02a43edf186cb9998741ba765cb948bb238a122d
```

```
Author:      Bas Nieuwenhuizen <bas@basnieuwenhuizen.nl>
```

```
radv: Optimize rebinding the same descriptor set.
```

```
This makes it cheaper to just change the dynamic offsets with
the same descriptor sets.
```

```
Suggested-by: Philip Rebohle <philip.rebohle@tu-dortmund.de>
```

```
Reviewed-by: Samuel Pitoiset <samuel.pitoiset@gmail.com>
```

```
c75a4e5b465261e982ea31ef875325a3cc30e79d
Author:      Dylan Baker <dylan@pnwbakers.com>
```

```
meson: Check for actual LLVM required versions
```

```
Currently we always check for 3.9.0, which is pretty
safe since everything except radv works with >= 3.9 and
3.9 is pretty old at this point. However, radv actually
requires 4.0, and there is a patch for radeonsi to do
the same.
```

```
Fixes: 673dda833076 ("meson: build "radv" vulkan driver
for radeon hardware")
```

```
Signed-off-by: Dylan Baker <dylan.c.baker@intel.com>
```

```
Reviewed-by: Bas Nieuwenhuizen <bas@basnieuwenhuizen.nl>
```

```
673dda8330769309a319d3e7f24a029cd72a1caf
```

```
Author:      Dylan Baker <dylan@pnwbakers.com>
```

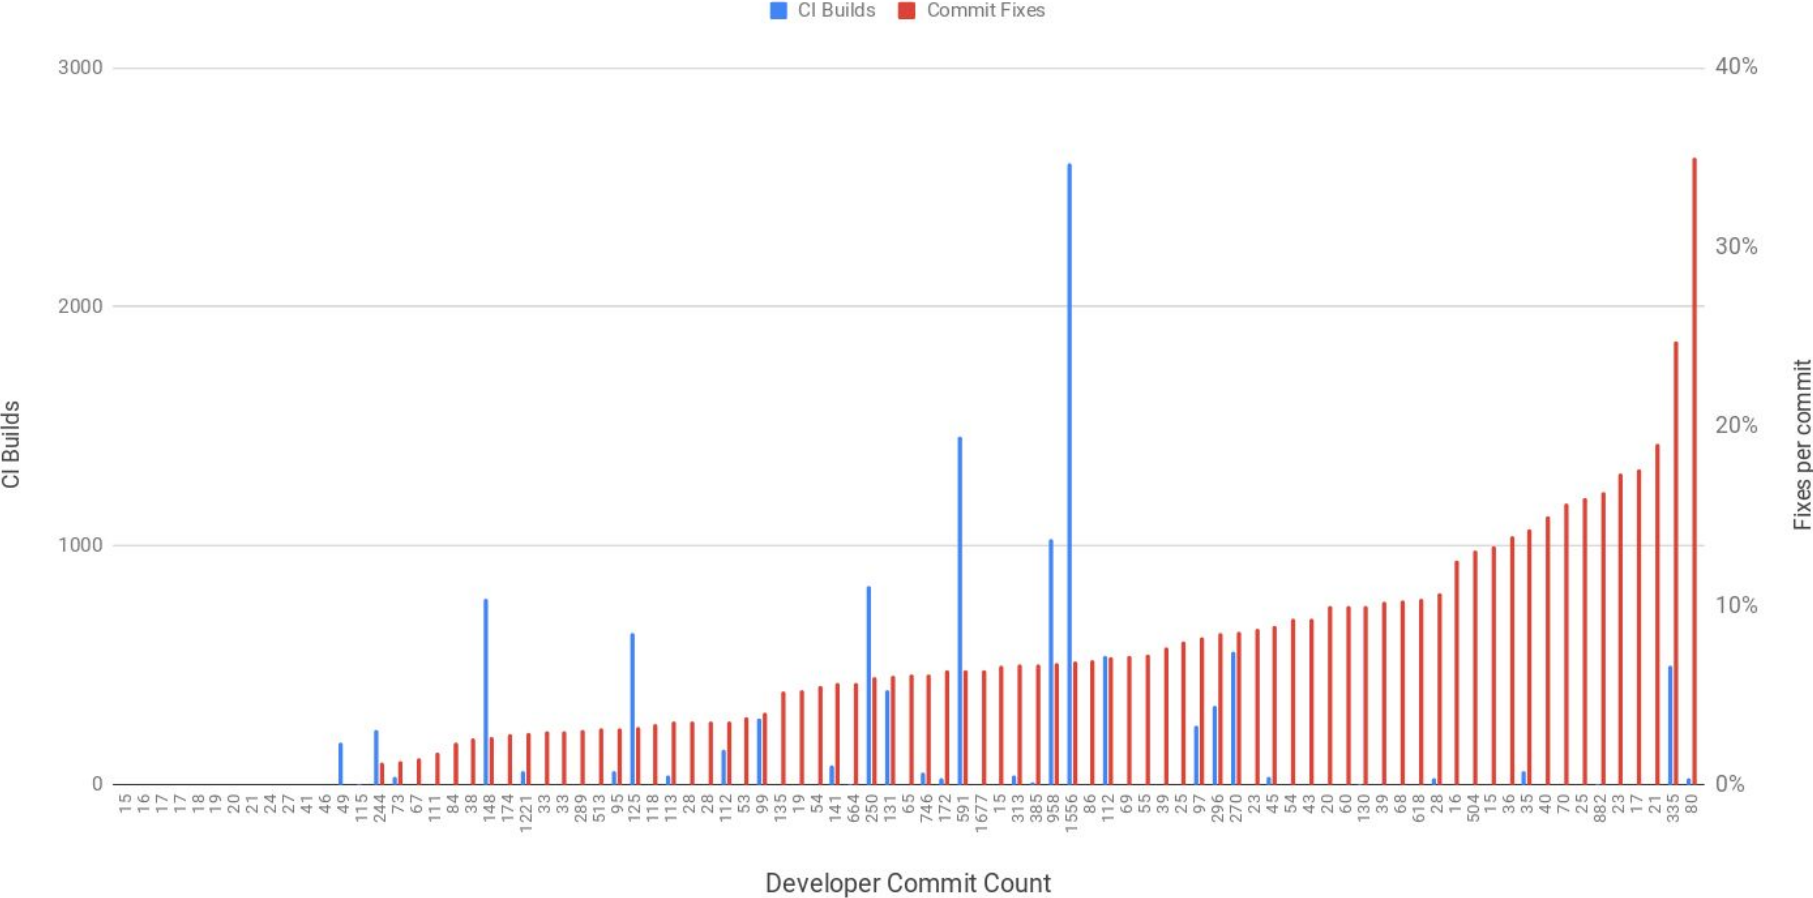
```
meson: build "radv" vulkan driver for radeon hardware
```

```
This builds, installs, and has been tested on a r290x
(Hawaii) with the Vulkan CTS. It dies horribly in a fire at
the same point for the meson build as the autotools build.
```

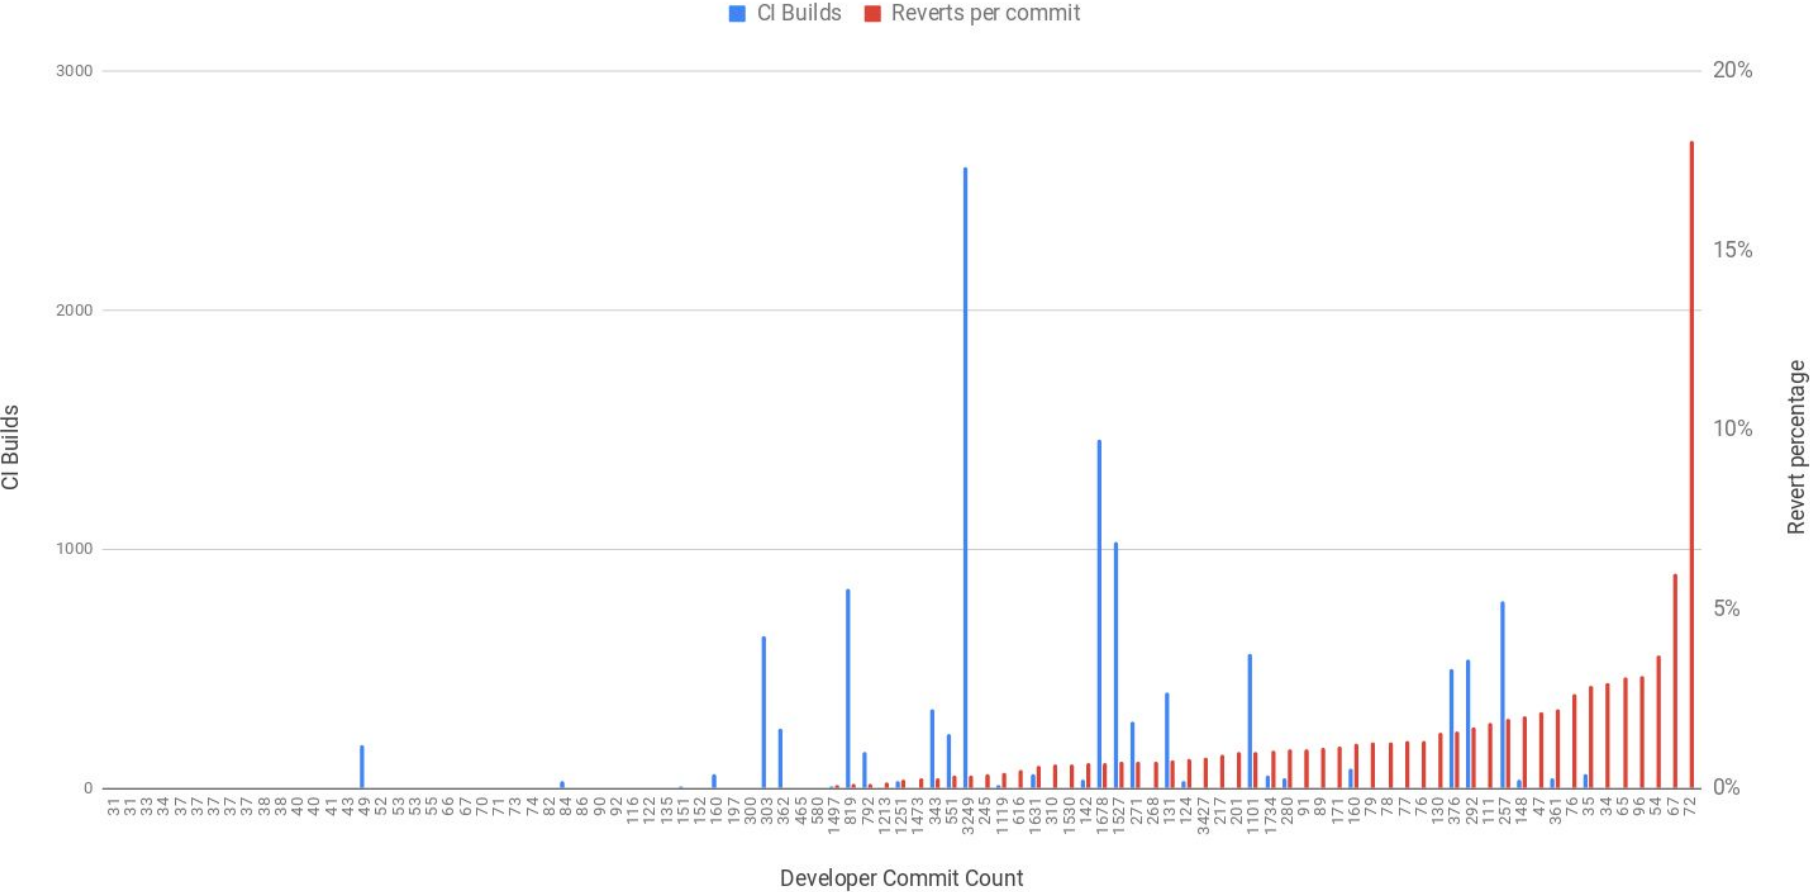
```
Signed-off-by: Dylan Baker <dylanx.c.baker@intel.com>
```

```
Reviewed-by: Bas Nieuwenhuizen <bas@basnieuwenhuizen.nl>
```

Fixes per commit since 2017, Developers with > 15 commits



Reverts per commit since 9/2014, Developers with > 30 commits



Tracking test status

ig65 CI tracks all test status changes in configuration files.

Known issues are filtered from the results, to make new regressions obvious in results:

- Developer pushes a broken commit.
- CI regresses at least one test.
- ig65 CI staff investigate, write FDO bug, where resolution is tracked.
- Regressed tests are filtered to “skip” status by CI config.
- Developer fixes bug.
- Test status changes, creating unexpected result in CI.
- ig65 CI staff investigate.
- CI configs updated to reflect fix.

CI Config file: SandyBridge Piglit results

[expected-failures]

```
piglit.shaders.arb_texture_gather-miplevels
piglit.shaders.point-vertex-id
...
piglit.shaders.glsl-deriv-varyings = piglit cd62eff8e5
Piglit.spec.ext_texture_compression_s3tc = piglit_d05448d06f
```

[expected-crashes]

```
piglit.fast_color_clear.fcc-front-buffer = mesa 880573e7
piglit.spec.egl_1_4.egl-copy-buffers = piglit 85e3b32b32
```

[fixed-tests]

```
piglit.spec.glsl-es-1_00.linker.fface-invariant = mesa 9b5c0c520
piglit.spec.glsl-es-1_00.linker.fcoord-invariant = mesa 9b5c0c520
```

Filtering results based on blamed commit

```
[expected-failures]  
piglit.spec.foo
```

```
[expected-crashes]
```

```
[fixed-tests]
```

```
e1623da8185 * idr test bug fix 1
```



```
49e4248a93a * i965/nir: export  
nir_optimize
```




Filtering results based on blamed commit

When test status changes, i965 CI staff triage results, close FDO bugs, and update CI configuration.

```
[expected-failures]
iglit.spec.foo
[expected-crashes]

[fixed-tests]
iglit.spec.foo
```

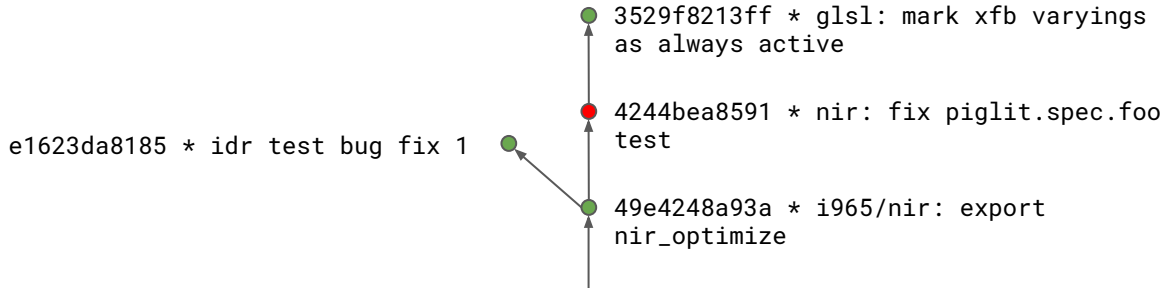


Filtering results based on blamed commit

Updated CI configuration files will report success for subsequent CI builds.

Test failure patterns are hardware specific. Each platform needs a separate configuration file. Some test suites require separate configuration for 32-bit builds.

```
[expected-failures]
[expected-crashes]
[fixed-tests]
piglit.spec.foo
```



Filtering results based on blamed commit

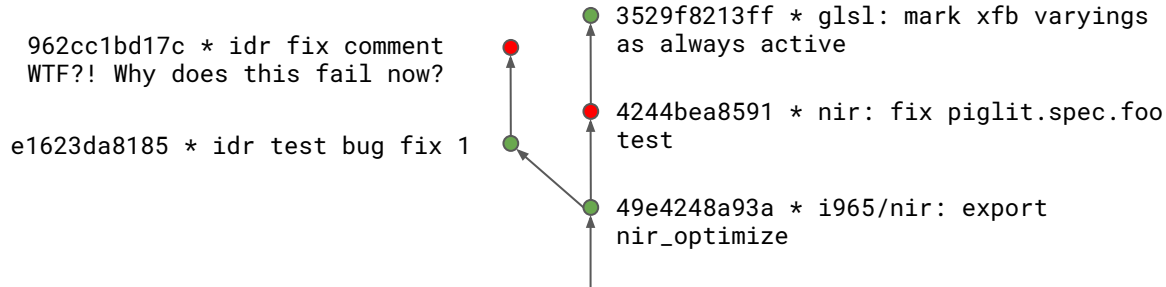
Branches will report spurious test status changes as CI tracks progress in the master branch.

Ian's branch does not contain the fix for piglit.spec.foo. Ian's test run will fail that test. The failure state does NOT match CI expectations.

```
[expected-failures]
```

```
[expected-crashes]
```

```
[fixed-tests]  
piglit.spec.foo
```



Filtering results based on blamed commit


Mesa CI records the blamed commit for every test status change.

For every unexpected test result, Mesa checks to see if the target branch contains the commit blamed by the CI config.

lan's branch does not contain 4244bea8591, so CI comprehends that the test status ought to be wrong for the branch.

```
[expected-failures]
[expected-crashes]
[fixed-tests]
piglit.spec.foo = mesa 4244bea8591
```

```
majanes@giraffe:~/src/mesa$ git branch -a --contains 4244bea8591
* master
* remotes/curro/wip/test
majanes@giraffe:~/src/mesa$
```



Filtering results based on blamed commit

Over time, CI configuration allows testing of releases for Mesa stable branches and old test suites.

Automated tests are fixed on a daily basis. Over time, this represents thousands of test results. Testing a stable point release is non-trivial.

Typically, CI systems fork the entire CI to test a stable branch. This is incompatible with hardware updates and other changes that affect all branches.

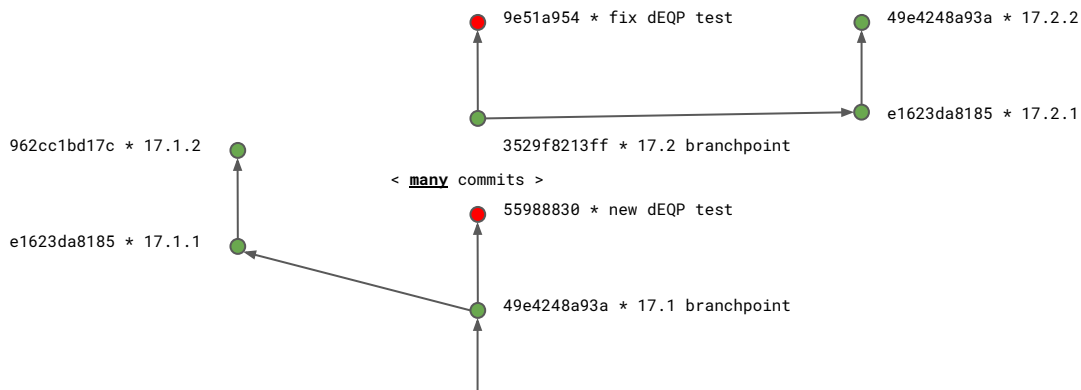
```
[expected-failures]
dEQP-GLES3.functional.shaders.preprocessor.builtin.line_expression_fragment = deqp 55988830
dEQP-GLES3.functional.shaders.preprocessor.builtin.line_expression_vertex = deqp 55988830

[expected-crashes]

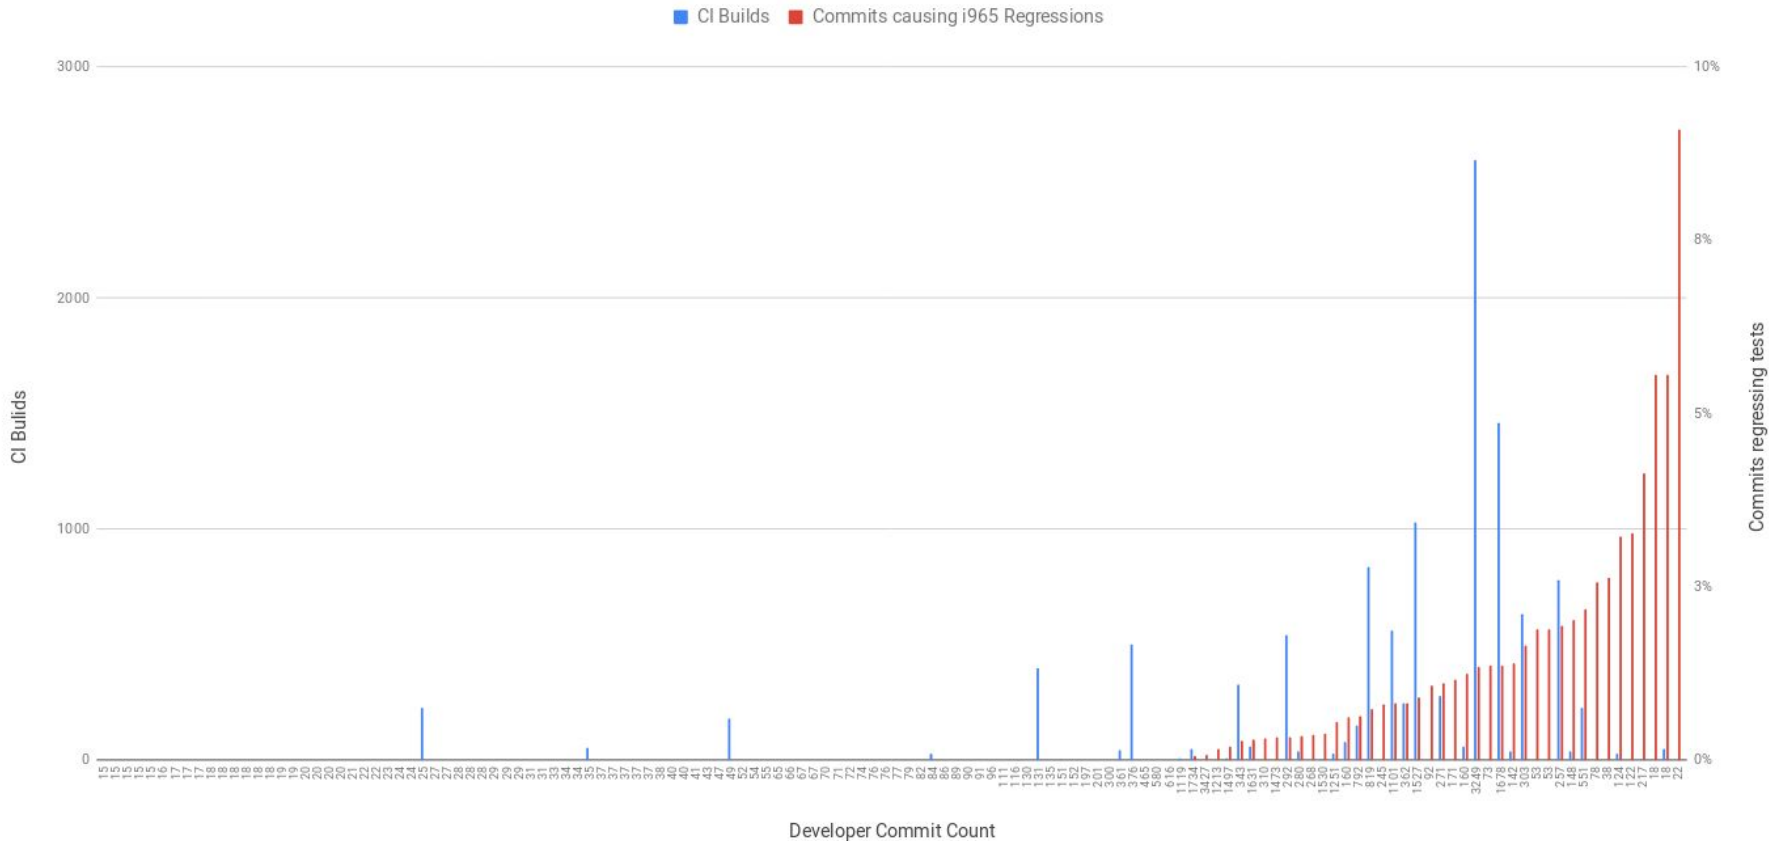
[fixed-tests]
dEQP-GLES3.functional.state_query.integers.stencil_value_mask_getfloat = mesa 37d63b50
dEQP-GLES3.functional.state_query.integers.stencil_back_value_mask_getfloat = mesa 37d63b50
dEQP-EGL.functional.color_clears.multi_thread.gles1_gles2.rgba8888_window = deqp 89c3844c
dEQP-EGL.functional.color_clears.multi_context.gles1_gles2.rgb888_pbuffer = deqp 89c3844c
dEQP-GLES3.functional.state_query.integers.stencil_value_mask_separate_both_getfloat = mesa 37d63b50
dEQP-GLES3.functional.state_query.integers.stencil_back_value_mask_separate_both_getfloat = mesa 37d63b50

<many more fixed tests ... >


EQP-EGL.functional.create_context.no_config = mesa 5e2909e7
dEQP-GLES31.functional.debug.negative_coverage.log.tessellation.single_tessellation_stage = mesa e6e8475b
dEQP-GLES3.functional.negative_api.texture.teximage3d = deqp 9e51a954
```



Commits regressing i965 tests since 8/2014, developers with >15 commits



Tests regressed per commit since 8/2014, Developers with >15 commits



Mesa CI at Intel



Mesa CI at Intel




Mesa CI at Intel

High density storage



Summary of Mesa CI at Intel

High density storage



How are we improving it?

New public CI results site!

Features:

- Results for other jobs (eg master, kwg, jekstrand)
- Broken tests counts
- Revisions / transparency of sources
- Logs for broken components
- Test history
- Browse results by test suite
- It's fast!

Internal results site:	1 minute 24 seconds
New public results site:	0.62(ish) seconds

Time to load a CI build result page

<https://mesa-ci.01.org>

How are we improving it?

New public CI results site!

Jobs displayed by this server:

Jobs

Filter jobs:

Multiple terms to search for can be specified, separated with spaces. Example: ci_daily_mesa_vulkan

Name	Last Build	Last Failed Build	Last Failed Build Count
aplinetro	#38-mesa-eeaa0ff0613 (2018-09-12 13:46:27)	#38-mesa-eeaa0ff0613 (2018-09-12 13:46:27)	198
apureses	#58-mesa-a276011e188 (2018-09-18 10:29:17)	#58-mesa-a276011e188 (2018-09-18 10:29:17)	15
chadv	#4-mesa-ad22e6fc64 (2018-09-19 19:38:11)	#4-mesa-ad22e6fc64 (2018-09-19 19:38:11)	1
cmarco_el	#123-mesa-49c35239827-percheckin (2018-09-15 05:58:31)	#123-mesa-49c35239827-percheckin (2018-09-15 05:58:31)	1
cmarco_vulkan	#66-mesa-49c35239827-percheckin (2018-09-15 05:41:45)		
dcshaker	#251-piglit-05b9c02a8e-percheckin (2018-09-20 14:29:14)	#251-piglit-05b9c02a8e-percheckin (2018-09-20 14:29:14)	2
dcshaker_18_1	#118-mesa-a76c43c193 (2018-09-20 23:36:14)	#118-mesa-a76c43c193 (2018-09-20 23:36:14)	24
dcshaker_mesa	#107-mesa-a79d5d79d8d-percheckin (2018-09-14 17:34:04)	#107-mesa-a79d5d79d8d-percheckin (2018-09-14 17:34:04)	1
djdeath	#186-mesa-9ea2bced47-percheckin (2018-09-07 14:19:09)	#184-mesa-f1dc32a0bb2-percheckin (2018-09-07 11:29:55)	48
djdeath_vulkan	#129-mesa-47aac574d66-percheckin (2018-09-11 20:16:57)	#129-mesa-47aac574d66-percheckin (2018-09-11 20:16:57)	1
fröblich	#58-mesa-833ab6959 (2018-09-11 06:25:25)		
global_logic	#3-mesa-3812802fa5 (2018-09-21 16:13:04)	#3-mesa-3812802fa5 (2018-09-21 16:13:04)	6
ldr	#835-mesa-4624242829-percheckin (2018-09-11 00:05:25)	#835-mesa-4624242829-percheckin (2018-09-11 00:05:25)	2
ksral	#231-mesa-b6c5732228 (2018-09-13 10:53:54)		
jekstrand_el	#1251-mesa-bf526847511-percheckin (2018-09-21 14:41:50)	#1251-mesa-bf526847511-percheckin (2018-09-21 14:41:50)	2
jekstrand_vulkan	#1269-mesa-bf526847511-percheckin (2018-09-21 14:41:50)		292
jancaanova	#6-mesa-791710701d (2018-09-16 14:15:05)	#6-mesa-791710701d (2018-09-16 14:15:05)	175
mesa_ig_1_daily	#184-started on schedule (2018-09-21 10:26:59)	#184-started on schedule (2018-09-21 10:26:59)	72
mesa_ig_2	#11-mesa-6c355780da-percheckin (2018-09-21 07:44:33)	#11-mesa-6c355780da-percheckin (2018-09-21 07:44:33)	24
mesa_ig_2_staging	#25-mesa-5276435679-percheckin (2018-09-19 12:22:23)	#25-mesa-5276435679-percheckin (2018-09-19 12:22:23)	25
mesa_master	#14126-mesa-b56fba7885-percheckin (2018-09-21 17:10:00)	#14117-drm-4c31c3-percheckin (2018-09-20 18:49:23)	4
mesa_master_daily	#4429-started on schedule (2018-09-20 13:59:36)	#4429-started on schedule (2018-09-20 13:59:36)	6
schery	#354-mesa-eb0ff947559-percheckin (2018-09-20 22:07:44)	#354-mesa-eb0ff947559-percheckin (2018-09-20 22:07:44)	338
schery_vulkan	#424-mesa-513d63cb1a-percheckin (2018-09-21 00:08:30)		
sgarghne	#13-mesa-b1632237755-percheckin (2018-09-14 20:46:30)	#13-mesa-b1632237755-percheckin (2018-09-14 20:46:30)	1
tarreri	#960-mesa-e189717705 (2018-09-14 05:56:42)	#959-mesa-a7965c4fdaf5 (2018-09-08 05:21:15)	45
tpalli	#568-mesa-049faaf1707-percheckin (2018-09-20 11:34:05)	#568-mesa-049faaf1707-percheckin (2018-09-20 11:34:05)	73292
vulkancts	#9154-mesa-bc56f87885-percheckin (2018-09-21 17:09:59)	#9154-mesa-bc56f87885-percheckin (2018-09-21 17:09:59)	9

Results for vulkancts_daily/#1145-manually_started-daily/root:

i965 Mesa CI » vulkancts_daily » #1145-manually_started-daily

Revisions • Components • Subgroups • Search tests:

Revisions

iglit	9682dfa1bb
vulkancts	8301a9125
kc-cts	1af33eb
shaderc	30af9f9
cts	e7cc24097
glslang	0886a332
mesa_ci	686208cf
gmock	f703a2
mesa_jenkins	1d56ab66
sim-drm	9cc8875
waffle	1de0209
spirvheaders	d502e12
spirvtools	58e5ea8
glcts	d77c1f98
mesa_ci_internal	e50b019
drm	1e31cc49
mesa	b8b3517a95
gtest	1b077667
deep	82c07c4d5
crucible	09c0701

Components

Component	Arch	Shard	Start Time	End Time	Machine	Status	Artifact(s)
crucible	m64	0	2018-09-20 11:39:48	2018-09-20 11:38:33	builder-02	failure	[log]
vulkancts-test	m64	1/2	2018-09-20 11:42:07	2018-09-20 11:47:59	ctf-02	success	
shaderc	m64	0	2018-09-20 11:32:48	2018-09-20 11:33:59	builder-01	success	
vulkancts-test	m64	2/2	2018-09-20 11:44:01	2018-09-20 11:48:18	ctf-03	success	
vulkancts-test	m64	3/5	2018-09-20 11:42:02	2018-09-20 11:58:23	bsw-02	success	

<https://mesa-ci.01.org>

How are we improving it?

Future:

- Support new use cases for the results website:
 - Show the logs and the status of components during execution
 - Allow developers to trigger custom builds
 - Queue up a set of jobs, with names
 - Especially useful for piglit
 - Allow developers to do A/B comparisons of builds
- Collaborate with other GPU vendors or Distros.

Got ideas? Let us know!

- <https://www.pivotaltracker.com/n/projects/1471364>
- https://gitlab.freedesktop.org/Mesa_CI/mesa_ci_results

Backup

Caveats

Rebased branches (e.g. i915 kernel trees) break methodology for tracking regressions

- Blamed commit SHAs may no longer exist after a rebase

Can hide regressions at release time

- Regressions are accepted failures in CI, so i965 CI staff has to manually compare results to N-1 release to find regressions attributed to the release candidate

But... A/B testing is more costly and only shows data for a single delta!

Why is continuous integration valuable for Mesa?

Mesa CI at Intel is currently used for

- ... Developer (Intel and community) testing
- ... Mesa release verification
- ... Intel pre-silicon testing in simulation
- ... Performance testing
- ... Validation of upstream test suites (dEQP, Vulkancts, etc)