

openQA Infrastructure - action #90755

aarch64 openQA jobs on osd not properly processed since 2021-04-05

2021-04-07 08:05 - okurz

| | | | |
|------------------------|----------|------------------------|------------|
| Status: | Resolved | Start date: | 2021-04-07 |
| Priority: | High | Due date: | 2021-04-22 |
| Assignee: | mkittler | % Done: | 0% |
| Category: | | Estimated time: | 0.00 hour |
| Target version: | Ready | | |

Description

Observation

As visible on <https://monitor.qa.suse.de/d/nRDab3Jiz/openqa-jobs-test?viewPanel=12&orgId=1&from=1617573600000&to=1617832799000> job queues on aarch64 within the osd-infrastructure do not seem to go down. This points to a problem that jobs are not properly processed at all for aarch64.

Acceptance criteria

- **AC1:** aarch64 jobs are processed within usual time frames for aarch64
- **AC2:** We have an alert or other mean preventing a similar situation in the future going unnoticed

History

#1 - 2021-04-07 08:10 - okurz

hm, looking at what workers are doing from <https://openqa.suse.de/admin/workers.html> maybe the problem of "something has stuck" has actually been resolved by the last osd-deployment as it seems now aarch64 jobs are picked up and the worker instances on openqaworker-arm-1 through openqaworker-arm-2 are working ok. Still, the problem should be evaluated, e.g. from worker logs what the machines had been doing for the past two days

#2 - 2021-04-07 08:13 - mkittler

- Assignee set to mkittler

- Priority changed from Urgent to High

I've just did the same and can confirm that it doesn't look broken at this point. I can look into worker logs and their job history.

How did you notice the problem?

#3 - 2021-04-07 08:38 - mkittler

It doesn't look too bad actually:

```
openqa=# select id, t_created, t_started, t_finished, (select concat(host, ':', instance) from workers where id = assigned_worker_id) as worker from jobs where ARCH = 'aarch64' order by t_created desc limit 100;
```

| id | t_created | t_started | t_finished | worker |
|---------|---------------------|---------------------|---------------------|-----------------------|
| 5785787 | 2021-04-07 08:09:36 | | | |
| 5785786 | 2021-04-07 08:09:34 | | | |
| 5785785 | 2021-04-07 08:09:33 | | | |
| 5785784 | 2021-04-07 08:09:28 | | 2021-04-07 08:09:47 | |
| 5785754 | 2021-04-07 07:53:22 | | | |
| 5785734 | 2021-04-07 07:47:13 | | | |
| 5785732 | 2021-04-07 07:45:42 | | | |
| 5785731 | 2021-04-07 07:45:07 | | | |
| 5785729 | 2021-04-07 07:44:29 | | | |
| 5785726 | 2021-04-07 07:43:29 | | | |
| 5785709 | 2021-04-07 07:17:00 | 2021-04-07 07:23:56 | | openqaworker-arm-2:17 |
| 5785627 | 2021-04-07 05:51:16 | 2021-04-07 06:23:55 | 2021-04-07 07:50:21 | openqaworker-arm-1:10 |
| 5785626 | 2021-04-07 05:51:16 | 2021-04-07 06:23:37 | 2021-04-07 07:13:55 | openqaworker-arm-2:17 |
| 5785625 | 2021-04-07 05:51:14 | | | |
| 5785624 | 2021-04-07 05:51:14 | 2021-04-07 06:13:09 | 2021-04-07 08:04:20 | openqaworker-arm-1:8 |
| 5785620 | 2021-04-07 05:51:11 | 2021-04-07 08:06:37 | 2021-04-07 08:14:20 | openqaworker-arm-1:8 |
| 5785618 | 2021-04-07 05:51:10 | 2021-04-07 08:01:35 | | openqaworker-arm-2:7 |

| | | | | | | | | |
|---------|--|---------------------|--|---------------------|--|---------------------|--|-----------------------|
| 5785619 | | 2021-04-07 05:51:10 | | 2021-04-07 08:04:41 | | 2021-04-07 08:13:30 | | openqaworker-arm-1:2 |
| 5785617 | | 2021-04-07 05:51:09 | | 2021-04-07 07:55:21 | | | | openqaworker-arm-2:14 |
| 5785616 | | 2021-04-07 05:51:08 | | 2021-04-07 07:55:06 | | | | openqaworker-arm-1:5 |
| 5785615 | | 2021-04-07 05:51:07 | | 2021-04-07 07:52:03 | | | | openqaworker-arm-2:6 |
| 5785614 | | 2021-04-07 05:51:07 | | 2021-04-07 07:51:00 | | | | openqaworker-arm-1:7 |
| 5785613 | | 2021-04-07 05:51:06 | | 2021-04-07 07:45:05 | | | | openqaworker-arm-1:9 |
| 5785612 | | 2021-04-07 05:51:05 | | 2021-04-07 07:43:42 | | | | openqaworker-arm-2:18 |
| 5785611 | | 2021-04-07 05:51:05 | | 2021-04-07 07:42:50 | | | | openqaworker-arm-1:1 |
| 5785610 | | 2021-04-07 05:51:04 | | 2021-04-07 07:40:20 | | | | openqaworker-arm-1:3 |
| 5785609 | | 2021-04-07 05:51:03 | | 2021-04-07 07:37:22 | | | | openqaworker-arm-2:3 |
| 5785608 | | 2021-04-07 05:51:01 | | 2021-04-07 05:55:49 | | 2021-04-07 07:12:35 | | openqaworker-arm-2:16 |
| 5785606 | | 2021-04-07 05:51:00 | | 2021-04-07 05:55:43 | | 2021-04-07 07:12:46 | | openqaworker-arm-2:11 |
| 5785607 | | 2021-04-07 05:51:00 | | 2021-04-07 05:55:48 | | 2021-04-07 07:12:41 | | openqaworker-arm-2:9 |
| 5785605 | | 2021-04-07 05:50:59 | | 2021-04-07 05:55:50 | | 2021-04-07 07:13:08 | | openqaworker-arm-2:3 |
| 5785581 | | 2021-04-07 04:37:39 | | 2021-04-07 04:40:48 | | 2021-04-07 05:51:16 | | openqaworker-arm-2:18 |
| 5785580 | | 2021-04-07 04:37:39 | | 2021-04-07 04:40:50 | | 2021-04-07 05:51:15 | | openqaworker-arm-2:6 |
| 5785578 | | 2021-04-07 04:37:38 | | 2021-04-07 04:39:18 | | 2021-04-07 06:49:51 | | openqaworker-arm-1:4 |
| 5785579 | | 2021-04-07 04:37:38 | | 2021-04-07 04:39:17 | | 2021-04-07 06:38:20 | | openqaworker-arm-1:9 |
| 5785576 | | 2021-04-07 04:37:37 | | 2021-04-07 05:55:45 | | | | openqaworker-arm-2:8 |
| 5785577 | | 2021-04-07 04:37:37 | | 2021-04-07 05:55:47 | | | | openqaworker-arm-2:5 |
| 5785575 | | 2021-04-07 04:37:36 | | 2021-04-07 05:55:46 | | | | openqaworker-arm-2:19 |
| 5785574 | | 2021-04-07 04:37:35 | | 2021-04-07 05:55:44 | | | | openqaworker-arm-2:15 |
| 5785573 | | 2021-04-07 04:37:35 | | 2021-04-07 05:55:43 | | 2021-04-07 06:22:16 | | openqaworker-arm-2:17 |
| 5785572 | | 2021-04-07 04:37:35 | | 2021-04-07 07:34:19 | | 2021-04-07 07:43:28 | | openqaworker-arm-2:18 |
| 5785571 | | 2021-04-07 04:37:34 | | 2021-04-07 07:50:50 | | 2021-04-07 08:15:09 | | openqaworker-arm-1:10 |
| 5785570 | | 2021-04-07 04:37:33 | | 2021-04-07 07:50:57 | | 2021-04-07 08:05:50 | | openqaworker-arm-2:20 |
| 5785569 | | 2021-04-07 04:37:33 | | 2021-04-07 07:32:42 | | 2021-04-07 08:03:59 | | openqaworker-arm-1:2 |
| 5785568 | | 2021-04-07 04:37:32 | | 2021-04-07 06:39:57 | | 2021-04-07 06:43:57 | | openqaworker-arm-1:9 |
| 5785567 | | 2021-04-07 04:37:31 | | 2021-04-07 06:38:29 | | 2021-04-07 06:42:41 | | openqaworker-arm-1:5 |
| 5785566 | | 2021-04-07 04:37:30 | | | | 2021-04-07 05:51:14 | | |
| 5785565 | | 2021-04-07 04:37:29 | | 2021-04-07 04:40:51 | | 2021-04-07 05:51:14 | | openqaworker-arm-2:12 |
| 5785562 | | 2021-04-07 04:37:21 | | 2021-04-07 07:15:22 | | 2021-04-07 07:22:33 | | openqaworker-arm-2:17 |
| 5785564 | | 2021-04-07 04:37:21 | | | | | | |
| 5785563 | | 2021-04-07 04:37:21 | | | | | | |
| 5785561 | | 2021-04-07 04:37:20 | | 2021-04-07 07:13:44 | | 2021-04-07 07:27:34 | | openqaworker-arm-2:11 |
| 5785560 | | 2021-04-07 04:37:19 | | 2021-04-07 07:13:42 | | 2021-04-07 07:27:32 | | openqaworker-arm-2:9 |
| 5785559 | | 2021-04-07 04:37:18 | | 2021-04-07 07:13:46 | | 2021-04-07 07:36:58 | | openqaworker-arm-2:3 |
| 5785557 | | 2021-04-07 04:37:17 | | 2021-04-07 06:55:53 | | 2021-04-07 07:14:52 | | openqaworker-arm-1:5 |
| 5785558 | | 2021-04-07 04:37:17 | | 2021-04-07 07:07:42 | | 2021-04-07 08:08:47 | | openqaworker-arm-1:4 |

However, there were lots of incompletes with the reason abandoned: associated worker openqaworker-arm-2:6 has not sent any status updates for too long, e.g. <https://openqa.suse.de/tests/5782816> (see select id, t_created, t_started, t_finished, state, result, reason, (select concat(host, ':', instance) from workers where id = assigned_worker_id) as worker from jobs where ARCH = 'aarch64' and result = 'incomplete' order by t_created desc limit 100; for full list).

#4 - 2021-04-07 09:10 - mkittler

The log of <https://openqa.suse.de/tests/5782816> looks like this on the worker-side:

```
Apr 07 06:05:27 openqaworker-arm-2 worker[14851]: [info] [pid:49749] sle-15-SP3-aarch64-172.3-textmode@aarch64
-virtio-multipath-lvm-encrypt.qcow2: Processing chunk 275/3383, avg. speed ~976.562 KiB/s
Apr 07 06:05:28 openqaworker-arm-2 worker[14851]: [info] [pid:49749] sle-15-SP3-aarch64-172.3-textmode@aarch64
-virtio-multipath-lvm-encrypt.qcow2: Processing chunk 276/3383, avg. speed ~976.562 KiB/s
Apr 07 06:05:29 openqaworker-arm-2 worker[14851]: [info] [pid:49749] sle-15-SP3-aarch64-172.3-textmode@aarch64
-virtio-multipath-lvm-encrypt.qcow2: Processing chunk 277/3383, avg. speed ~976.562 KiB/s
Apr 07 06:05:30 openqaworker-arm-2 worker[14851]: [info] [pid:49749] sle-15-SP3-aarch64-172.3-textmode@aarch64
-virtio-multipath-lvm-encrypt.qcow2: Processing chunk 278/3383, avg. speed ~976.562 KiB/s
Apr 07 06:05:31 openqaworker-arm-2 worker[14851]: [info] [pid:49749] sle-15-SP3-aarch64-172.3-textmode@aarch64
-virtio-multipath-lvm-encrypt.qcow2: Processing chunk 279/3383, avg. speed ~976.562 KiB/s
Apr 07 06:05:32 openqaworker-arm-2 worker[14851]: [info] [pid:49749] sle-15-SP3-aarch64-172.3-textmode@aarch64
-virtio-multipath-lvm-encrypt.qcow2: Processing chunk 280/3383, avg. speed ~976.562 KiB/s
Apr 07 06:05:33 openqaworker-arm-2 worker[14851]: [info] [pid:49749] sle-15-SP3-aarch64-172.3-textmode@aarch64
-virtio-multipath-lvm-encrypt.qcow2: Processing chunk 281/3383, avg. speed ~976.562 KiB/s
Apr 07 06:05:33 openqaworker-arm-2 worker[14851]: [debug] [pid:14851] REST-API call: POST http://openqa.suse.d
e/api/v1/jobs/5782816/status
Apr 07 06:05:34 openqaworker-arm-2 worker[14851]: [info] [pid:49749] sle-15-SP3-aarch64-172.3-textmode@aarch64
-virtio-multipath-lvm-encrypt.qcow2: Processing chunk 282/3383, avg. speed ~976.562 KiB/s
Apr 07 06:05:34 openqaworker-arm-2 worker[14851]: [debug] [pid:14851] Upload concluded (no current module)
Apr 07 06:05:34 openqaworker-arm-2 worker[14851]: [info] [pid:49749] sle-15-SP3-aarch64-172.3-textmode@aarch64
-virtio-multipath-lvm-encrypt.qcow2: Processing chunk 283/3383, avg. speed ~976.562 KiB/s
-- Reboot --
Apr 07 06:37:40 openqaworker-arm-2 systemd[1]: Starting openQA Worker #1...
Apr 07 06:37:40 openqaworker-arm-2 systemd[1]: Started openQA Worker #1.
Apr 07 06:37:44 openqaworker-arm-2 worker[3086]: [info] [pid:3086] worker 1:
Apr 07 06:37:44 openqaworker-arm-2 worker[3086]: - config file: /etc/openqa/workers.ini
```

```

Apr 07 06:37:44 openqaworker-arm-2 worker[3086]: - worker hostname:      openqaworker-arm-2
Apr 07 06:37:44 openqaworker-arm-2 worker[3086]: - isotovideo version:      23
Apr 07 06:37:44 openqaworker-arm-2 worker[3086]: - websocket API version:  1
Apr 07 06:37:44 openqaworker-arm-2 worker[3086]: - web UI hosts:          openqa.suse.de
Apr 07 06:37:44 openqaworker-arm-2 worker[3086]: - class:                qemu_aarch64,tap,openqaworker-arm-2
Apr 07 06:37:44 openqaworker-arm-2 worker[3086]: - no cleanup:          no
Apr 07 06:37:44 openqaworker-arm-2 worker[3086]: - pool directory:      /var/lib/openqa/pool/1
Apr 07 06:37:44 openqaworker-arm-2 worker[3086]: [error] [pid:3086] Worker cache not available: No workers active in the cache service
Apr 07 06:37:44 openqaworker-arm-2 worker[3086]: [info] [pid:3086] CACHE: caching is enabled, setting up /var/lib/openqa/cache/openqa.suse.de
Apr 07 06:37:44 openqaworker-arm-2 worker[3086]: [info] [pid:3086] Project dir for host openqa.suse.de is /var/lib/openqa/share
Apr 07 06:37:44 openqaworker-arm-2 worker[3086]: [info] [pid:3086] Registering with openQA openqa.suse.de
Apr 07 06:37:45 openqaworker-arm-2 worker[3086]: [warn] [pid:3086] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Apr 07 06:37:55 openqaworker-arm-2 worker[3086]: [info] [pid:3086] Registering with openQA openqa.suse.de
...
Apr 07 06:39:55 openqaworker-arm-2 worker[3086]: [warn] [pid:3086] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Apr 07 06:40:05 openqaworker-arm-2 worker[3086]: [info] [pid:3086] Registering with openQA openqa.suse.de
Apr 07 06:40:05 openqaworker-arm-2 worker[3086]: [info] [pid:3086] Establishing ws connection via ws://openqa.suse.de/api/v1/ws/886
Apr 07 06:40:05 openqaworker-arm-2 worker[3086]: [info] [pid:3086] Registered and connected via websockets with openQA host openqa.suse.de and worker ID 886
Apr 07 06:40:43 openqaworker-arm-2 worker[3086]: [debug] [pid:3086] Accepting job 5782972 from openqa.suse.de.
Apr 07 06:40:43 openqaworker-arm-2 worker[3086]: [debug] [pid:3086] Setting job 5782972 from openqa.suse.de up

```

So the worker simply crashed but was automatically recovered in minutes. So I suppose these incompletes are unfortunately expected considering the arm workers are generally unstable.

If I scroll down further within the log I see one more reboot:

```

Apr 07 07:20:33 openqaworker-arm-2 worker[3086]: [debug] [pid:3086] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5782972/status
Apr 07 07:20:33 openqaworker-arm-2 worker[3086]: [debug] [pid:3086] Upload concluded (at wait_children)
-- Reboot --
Apr 07 07:52:32 openqaworker-arm-2 systemd[1]: Starting openQA Worker #1...
Apr 07 07:52:32 openqaworker-arm-2 systemd[1]: Started openQA Worker #1.

```

When looking at the full system log there's nothing interesting to see shortly before the reboot.

The same happened on openqaworker-arm-1 as well, e.g. <https://openqa.suse.de/tests/5782884>.

```

Apr 07 02:44:36 openqaworker-arm-1 worker[35652]: [debug] [pid:35652] Upload concluded (at force_scheduled_tasks)
Apr 07 02:44:36 openqaworker-arm-1 worker[32730]: [debug] [pid:32730] Upload concluded (at yast2_migration)
Apr 07 02:44:36 openqaworker-arm-1 worker[27610]: [debug] [pid:45867] Uploading artefact patch_sle-173.png as 5d00a4bb879836a77c918cbcd36032eb
-- Reboot --
Apr 07 03:15:52 openqaworker-arm-1 kernel: Booting Linux on physical CPU 0x000000000 [0x431f0a11]
Apr 07 03:15:52 openqaworker-arm-1 kernel: Linux version 5.8.3-1.gbad027a-default (geeko@buildhost) (gcc (SUSE Linux) 10.2.1 20200805 [revision dda1e9d08434def88ed86557d08b23251332c5aa], GNU ld (GNU Binutils; openSUSE Tu mbleweed) 2.34.0.20200325-1) #1 SMP Sat Aug 22 06:31:07 UTC 2020 (bad027a)
Apr 07 03:15:52 openqaworker-arm-1 kernel: Machine model: Cavium ThunderX CN88XX board

```

Another crash the day before:

```

Apr 06 23:11:23 openqaworker-arm-1 worker[27300]: [debug] [pid:27300] Upload concluded (at slurm_master_backup)
Apr 06 23:11:23 openqaworker-arm-1 worker[28372]: [debug] [pid:28372] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5778496/status
Apr 06 23:11:23 openqaworker-arm-1 worker[28372]: [debug] [pid:28372] Upload concluded (at slurm_master)
Apr 06 23:11:23 openqaworker-arm-1 worker[26700]: [debug] [pid:26700] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5778493/status
-- Reboot --
Apr 06 23:41:52 openqaworker-arm-1 kernel: Booting Linux on physical CPU 0x000000000 [0x431f0a11]
Apr 06 23:41:52 openqaworker-arm-1 kernel: Linux version 5.8.3-1.gbad027a-default (geeko@buildhost) (gcc (SUSE Linux) 10.2.1 20200805 [revision dda1e9d08434def88ed86557d08b23251332c5aa], GNU ld (GNU Binutils; openSUSE Tu mbleweed) 2.34.0.20200325-1) #1 SMP Sat Aug 22 06:31:07 UTC 2020 (bad027a)
Apr 06 23:41:52 openqaworker-arm-1 kernel: Machine model: Cavium ThunderX CN88XX board
Apr 06 23:41:52 openqaworker-arm-1 kernel: efi: EFI v2.40 by American Megatrends

```

So I would assume we just see the same instability of the workers as usual and the automatic recovery works. However, maybe the frequency of the crashes has increased.

I don't see anything we can improve since the automatic recovery is working and all jobs I looked at have been automatically restarted. I suppose we have to live with the fact that aarch64 jobs are piling up until we get more workers (which are not that unstable).

#5 - 2021-04-07 09:20 - okurz

mkittler wrote:

How did you notice the problem?

sorry, forgot about that part. I was informed in direct chat by user mgrifalconi

#6 - 2021-04-07 09:25 - okurz

ok, but if both openqaworker-arm-1 and openqaworker-arm-2 became more unstable at the same time then it's likely software updates that caused this. I suggest to try the following ideas:

- Run stress tests with <https://software.opensuse.org/package/stress-ng>
- Collect serial logs from from this machines over a screen session on a persistent machine, e.g. root@osd
- Look into better solutions for log aggregation from IPMI serial ports so that we do not need to do manual screen session setups in the future
- Crosscheck the logs for "unexpected reboots" vs. the package updates that had been recently applied, potentially revert package installations to an earlier stage for crosschecking

#7 - 2021-04-07 11:00 - mkittler

They've been unstable before so somehow I double it is a software issue (at least no new software issue). I suppose that [#89815](#) is unresolved makes it also worse.

Considering that jobs have already been processed slowly I'd refrain from running additional stress tests for now.

I've been starting IPMI SOL sessions within screen sessions on root@osd.

#8 - 2021-04-08 04:16 - openqa_review

- Due date set to 2021-04-22

Setting due date based on mean cycle time of SUSE QE Tools

#9 - 2021-04-08 07:43 - mkittler

Nothing to see on the serial consoles so far. Looks like the workers haven't crashed again. The job queue also went down.

#10 - 2021-04-08 11:53 - mkittler

- Status changed from Workable to Feedback

There was a reboot on openqaworker-arm-1 (most likely triggered by <https://gitlab.suse.de/openqa/grafana-webhook-actions/-/jobs/383701>).

This is the SOL output before:

```
openqa:/home/martchus # ipmitool -I lanplus -C 3 -H openqaworker-arm-1-ipmi.suse.de -U ADMIN -P ADMIN sol activate
[SQL Session operational. Use ~? for help]

openqaworker-arm-1 login: [46533.178250] nvme nvme0: I/O 44 QID 23 timeout, aborting
[46533.184118] nvme nvme0: Abort status: 0x0
[50156.378279] nvme nvme0: I/O 303 QID 25 timeout, aborting
[50156.384531] nvme nvme0: Abort status: 0x0
[80762.384375] perf: interrupt took too long (3154 > 3127), lowering kernel.perf_event_max_sample_rate to 6340
0
[85318.804084] nvme nvme0: I/O 360 QID 7 timeout, aborting
[85318.809759] nvme nvme0: Abort status: 0x0
```

I suppose everything which comes after is already after the power cycle:

```
Cavium SOC
Locking L2 cache
PASS: CRC32 verification
Transferring to thread♦UART: Baut rate configuration: 115200
Read CPLD data a
Set CPLD data a
Check CPLD data a
=====
Cavium THUNDERX Boot Stub
```

=====
BDK Version: 04.01.13c, 2.14.3, Branch: thunder-sdk-2.1.0-build5, Built: Fri Sep 23 08:23:26 UTC 2016

Board Model: gbt-mt30
Board name: MT30-GS1-00
System name: R120-T32-00
Board Revision: 1.x
Board SKU: 04 SATA
SKU version: 1.1

Node: 0 (Fixed)
Chip: Oxal Pass 2.1
SKU: CN8890-2000BG2601-CP-Y-G
L2: 16384 KB
RCLK: 2000 Mhz
SCLK: 800 Mhz
Boot: SPI24(5)
VRM: Disabled
Trust: Disabled, Non-secure Boot
CCPI: Disabled

Press 'B' within 10 seconds for boot menu
Loading image file '/fatfs/init.bin'

Cavium SOC
Locking L2 cache
PASS: CRC32 verification
Transferring to thread scheduler
Using configuration from previous image
=====

Cavium Init
=====

BDK Version: 04.01.13c, 2.14.3, Branch: thunder-sdk-2.1.0-build5, Built: Fri Sep 23 08:23:26 UTC 2016

N0.LMC0 Configuration Completed: 65536 MB
N0.LMC1 Configuration Completed: 65536 MB
N0.LMC2 Configuration Completed: 65536 MB
N0.LMC3 Configuration Completed: 65536 MB
Node 0: DRAM: 262144 MB, 940 MHz, DDR4 RDIMM
Starting Test "Fast Scan" for [0x000005a0000:0x03ffffffff] using 1 core(s)
100.0% complete, testing [0x000005a0000:0x03ffffffff]0x02e6059fffff]
Read Node0 SKU
Read Node0 SKU from eeprom
Node0 SKU is correct.
GPIO: Node0 CP
N0.PCIe0: Link active, 1 lanes, speed gen1
N0.PCIe2: Link active, 4 lanes, speed gen3
N0.PCIe4: Link timeout, probably the slot is empty
Loading image file '/fatfs/cortina-app.bin'

Cavium SOC
PASS: CRC32 verification
Transferring to thread scheduler
Using configuration from previous image
CORTINA: Press 'M' within 2 seconds for Cortina menu

CORTINA: Running CORTINA PHY firmware updater...
CORTINA: Firmware update on date mismatch is configured in /rom/cortina.inf file.
CORTINA: Image file name is: /rom/cortina.img
CORTINA: Image file date is: 0x0507 0x2015 0x1530
CORTINA: ROM firmware date is: 0x0507 0x2015 0x1530
CORTINA: Firmware is up to date.

CORTINA: Loading next stage...
Loading image at /boot:0x400000

=====
BOARD MODEL = gbt-mt30
BMC BOOT TWSI bus=0x5, addr=0x20

```
BMC IPMI TWSI bus=0x4, addr=0x10
GPIO Shutdown pin IN = 0x25
GPIO Shutdown pin OUT = 0x24
=====
Mark memory region 0:: 0 to 300000 as secure (2)
Mark memory region 1:: 400000 to 4000000000 as non-secure (1)
#
GPIO(0) Node(0) init called config_base:848000030000 size:1000
table_size :96 bir:4
MSI-X vector base:803000f00000
GPIO(0)-NODE(0): Vector:122 address :801000000050 irq:33
GPIO(0)-NODE(0): Vector:123 address :801000000058 irq:33
Using TWSI func = 0x48
Using TWSI func = 0x49
Using TWSI func = 0x4a
Using TWSI func = 0x4b
Using TWSI func = 0x4c
Using TWSI func = 0x4d
TWSI1(5) Node(0) init called config_base:84800014d000 size:1000
table_size :1 bir:4
MSI-X vector base:87e0d5f00000
#####
#
#####
Booting trusted firmware boot loader stage 1
v0.3(release):
Built : 16:23:25, Sep 23 2016
Environment: FDT @: 0x20000
...
```

Note that on openqaworker-arm-2 similar log messages appeared but the system did *not* crash:

```
openqaworker-arm-2 login: [27106.672640] nvme nvme0: I/O 89 QID 4 timeout, completion polled
[34975.968995] nvme nvme0: I/O 833 QID 4 timeout, aborting
[34975.976501] nvme nvme0: Abort status: 0x0
[49344.249699] nvme nvme0: I/O 770 QID 29 timeout, aborting
[49344.257731] nvme nvme0: Abort status: 0x0
[53478.057015] nvme nvme0: I/O 401 QID 24 timeout, completion polled
[57394.335646] nvme nvme0: I/O 184 QID 30 timeout, aborting
[57394.343307] nvme nvme0: Abort status: 0x0
[60914.954556] nvme nvme0: I/O 1010 QID 5 timeout, completion polled
[72155.818002] nvme nvme0: I/O 821 QID 27 timeout, completion polled
[76817.019783] nvme nvme0: I/O 204 QID 20 timeout, aborting
[76817.026788] nvme nvme0: Abort status: 0x0
[78558.610433] nvme nvme0: I/O 877 QID 21 timeout, aborting
[78558.617940] nvme nvme0: Abort status: 0x0
[81558.696410] nvme nvme0: I/O 885 QID 15 timeout, aborting
[81558.704016] nvme nvme0: Abort status: 0x0
[83895.136549] nvme nvme0: I/O 241 QID 28 timeout, aborting
[83895.143709] nvme nvme0: Abort status: 0x0
[86496.122580] nvme nvme0: I/O 1012 QID 13 timeout, aborting
[86496.129648] nvme nvme0: Abort status: 0x0
[87130.846554] BTRFS info (device sda3): qgroup scan completed (inconsistency flag cleared)
[87912.449656] nvme nvme0: I/O 65 QID 28 timeout, completion polled
[87946.687885] nvme nvme0: I/O 69 QID 12 timeout, aborting
[87946.694942] nvme nvme0: Abort status: 0x0
[90309.446352] nvme nvme0: I/O 938 QID 30 timeout, completion polled
[90542.794224] n[90542.801051] nvme nvme0: Abort status: 0x0
[91242.678125] nvme nvme0: I/O 17 QID 7 timeout, completion polled
[93534.880248] nvme nvme0: I/O 399 QID 1 timeout, completion polled
[98284.236349] nvme nvme0: I/O 632 QID 26 timeout, aborting
[98284.243783] nvme nvme0: Abort status: 0x0
[106580.691303] nvme nvme0: I/O 555 QID 10 timeout, completion polled
[106931.233390] nvme nvme0: I/O 246 QID 7 timeout, aborting
[106931.240622] nvme nvme0: Abort status: 0x0
```

So I assume that Abort status: 0x0 is not necessarily related to the crash. That means the serial log might not be very useful after all.

I've noticed that the timeout/delay of the recovery is quite long, e.g. Grafana triggers it only after 6 minutes and then the GitLab hook tries to ping for 600 seconds and then again for 600 seconds. So only after 26 minutes the power cycle is triggered. This roughly matches the timestamps in the logs from my previous comment. Maybe we should at least get rid of the first ping?

#11 - 2021-04-09 10:52 - mkittler

SR for getting rid of initial pings: https://gitlab.suse.de/openqa/grafana-webhook-actions/-/merge_requests/14

#12 - 2021-04-12 10:26 - okurz

MR merged. That should cover AC1. Ideas about AC2?

#13 - 2021-04-12 13:48 - mkittler

Maybe we could use <https://stats.openqa-monitor.qa.suse.de/d/7W06NBWGk/job-age?orgId=1> to trigger an alert? We could also add an alert if all ARM workers are offline at the same time (or if only one is online, now since openqaworker3 is back).

#14 - 2021-04-13 07:25 - okurz

mkittler wrote:

Maybe we could use <https://stats.openqa-monitor.qa.suse.de/d/7W06NBWGk/job-age?orgId=1> to trigger an alert?

Well, likely yes. We already have alerts for general job ages which are multiple days because some worker classes are providing only awful expectations and it's not something that we can fix. That is up to the testers.

We could also add an alert if all ARM workers are offline at the same time (or if only one is online, now since openqaworker3 is back).

You are again confusing openqaworker3 and openqaworker-arm-3 :D Hm, the idea could be to have an alert if any worker class is not covered by any machines. However any user can define an invalid worker class when triggering jobs so not sure. Maybe we can try again the idea I had and come up with something about calculating a ratio of "working/available" and raise an alert if not too many workers are working whereas more should be available. Maybe that's a topic we can bring up in a daily or weekly?

#15 - 2021-04-19 09:47 - okurz

- *Status changed from Feedback to Resolved*

Discussed with mkittler. We have "job age" and alerts although their thresholds are very high due to some cases that we simply can not control, e.g. the "virtualization" job group with very long job ages.