openQA Project - coordination #101048

[epic] Investigate and fix higher instability of openqaworker-arm-4/5 vs. arm-1/2/3

2021-10-15 14:39 - okurz

Status:	Blocked	**Start date:**	2021-10-15
Priority:	Normal	**Due date:**	2022-01-25
Assignee:	okurz	**% Done:**	100%
Category:	Concrete Bugs	**Estimated time:**	0.00 hour
Target version:	Ready		
Difficulty:			

Description

Observation

According to https://monitor.qa.suse.de/d/nRDab3Jiz/openqa-jobs-test?viewPanel=27&orgId=1&from=now-30d&to=now (sort by "avg" in the table on the right-hand side) openqaworker-arm-4/5 have a fail-ratio of 33-36% vs. openqaworker-arm-1/2/3 with a fail-ratio of 15-17%

Acceptance criteria

- **AC1:** openqaworker-arm-4/5 have a fail-ratio less or equal to arm-1/2/3

Suggestions

- Confirm if typing issues cause the failures (look for timeouts, observe additional or missing characters in typed commands)
- Upgrade arm3 to Leap 15.3 and compare failure rate -> #101265 => Leap 15.3 behaves similar as Leap 15.2
- Consider switching to kernel-stable or kernel-head -> #101271 => "kernel-default" from Kernel:stable behaves same as openSUSE:Leap:15.3 one
- ~~Consider downgrading kernel to what's used in 15.2~~ -> same upstream version is running on most
- Bring back arm 4 and 5 after verifying stability
- Run typing.pm from os-autoinst as test in production -> #101262

Subtasks:

action # 101262: Document running os-autoinst full-stack.t on OSD workers size:M	**Resolved**
action # 101265: Upgrade arm3 to Leap 15.3 and compare failure rate size:M	**Resolved**
action # 101271: Try Kernel:stable on arm4+arm5 and compare failure rate size:M	**Resolved**
action # 104304: Crosscheck results of https://github.com/os-autoinst/os-autoinst#verif...	**Resolved**

Related issues:

Related to openQA Project - action #101030: Typing problems on aarch64	**Resolved**	**2021-10-15**
Copied to openQA Project - action #101265: Upgrade arm3 to Leap 15.3 and comp...	**Resolved**	**2021-10-15**

History

#1 - 2021-10-16 12:36 - okurz

- Related to action #101030: Typing problems on aarch64 added

#2 - 2021-10-18 11:37 - nicksinger

- File sysctl_diff.html added

First investigation shows that we run leap15.2 on the "old" workers and 15.3 on the new ones. Kernel-version seems to be quite different between different workers:
arm-1: 5.8.3-1.gbad027a-default
arm-2: 5.7.12-1.g9c98feb-default
arm-3: 5.3.18-lp152.95-default
arm-4: 5.3.18-59.27-default
arm-5: 5.3.18-59.27-default

given that arm3 has at least a similar kernel version I'd exclude kernel for now.
Kernel cmdline is managed by salt and therefore the same on all 5 machines.

I also tried to diff the sysctl's currently set in the system. Due to different kernels this is a quite tedious task and I didn't see much which could make a difference here. Attaching the diff as html.

#3 - 2021-10-18 11:44 - nicksinger

despite being way less utilized arm4 sees load-spikes up to 75 and is around 25 quite constantly according to: https://monitor.qa.suse.de/d/WDopenqaworker-arm-4/worker-dashboard-openqaworker-arm-4?viewPanel=54694&orgId=1&from=1633952613651&to=1634557413651 - this could hint to IO performance issues.

#4 - 2021-10-18 11:51 - nicksinger

Disk IO is 10x less compared to arm3 on arm4: https://monitor.qa.suse.de/d/WDopenqaworker-arm-4/worker-dashboard-openqaworker-arm-4?viewPanel=13782&orgId=1&from=1633953007558&to=1634557807558 (expected, less load) but IO response times seem to be ~50% worse according to https://monitor.qa.suse.de/d/WDopenqaworker-arm-4/worker-dashboard-openqaworker-arm-4?viewPanel=56720&orgId=1&from=1633953007558&to=1634557807558

#5 - 2021-10-18 11:55 - nicksinger

network IO seems fine. Higher packet drops can be observed on arm4 & arm5 (exactly the same pattern, so hinting to the switch) but IMHO this shouldn't cause such a performance-hit

#6 - 2021-10-21 09:14 - okurz

- Priority changed from High to Urgent

#7 - 2021-10-21 09:27 - cdywan

- Description updated

#8 - 2021-10-21 09:37 - cdywan

- Tracker changed from action to coordination

- Subject changed from Investigate higher instability of openqaworker-arm-4/5 vs. arm-1/2/3 to [epic] Investigate higher instability of openqaworker-arm-4/5 vs. arm-1/2/3

- Description updated

#9 - 2021-10-21 09:48 - cdywan

- Copied to action #101265: Upgrade arm3 to Leap 15.3 and compare failure rate size:M added

#10 - 2021-10-21 09:50 - okurz

- Description updated

#11 - 2021-10-21 09:54 - okurz

- Description updated

#12 - 2021-10-21 10:01 - okurz

- Status changed from New to Blocked

- Assignee set to okurz

blocked by subtasks

#13 - 2021-11-14 14:50 - okurz

- Description updated

#14 - 2021-11-14 14:53 - okurz

- Description updated

#15 - 2021-11-14 14:54 - okurz

- Status changed from Blocked to New

- Assignee deleted (okurz)

I updated the epic with the results from #101265 and #101271 . We can now continue defining more hypotheses to follow-up with.

#16 - 2021-11-16 10:10 - okurz

- Subject changed from [epic] Investigate higher instability of openqaworker-arm-4/5 vs. arm-1/2/3 to [epic] Investigate and fix higher instability of openqaworker-arm-4/5 vs. arm-1/2/3

#17 - 2021-11-17 14:08 - kraih

There's still a lot of failed jobs from the #101271 stress test that should be searched for patterns. Maybe that will give some hints for where to look with followup investigations.

#18 - 2021-11-18 18:49 - okurz

- Status changed from New to Workable

#20 - 2021-12-08 14:11 - mkittler

The workers arm-4/5 went offline on 05.12.2021. IPMI still responds so I invoked a power cycle. However, they both workers didn't boot successfully. They've got both stuck in the early boot:

```
Loading Linux 5.15.5-lp153.2.g83fc974-default ...
Loading initial ramdisk ...
EFI stub: Booting Linux Kernel...
EFI stub: EFI_RNG_PROTOCOL unavailable
EFI stub: ERROR: FIRMWARE BUG: kernel image not aligned on 64k boundary
EFI stub: ERROR: FIRMWARE BUG: Image BSS overlaps adjacent EFI memory region
EFI stub: Using DTB from configuration table
EFI stub: Exiting boot services...
INFO:    Node: 0 :: REP: 0x0, REP-FAIL: 0x0, MBIST: 0x0, MBIST-FAIL: 0x803c3c
INFO:    Node: 1 :: REP: 0x0, REP-FAIL: 0x0, MBIST: 0x0, MBIST-FAIL: 0x803c3c
[    0.000000][    T0] Booting Linux on physical CPU 0x0000000000 [0x431f0af2]
[    0.000000][    T0] Linux version 5.15.5-lp153.2.g83fc974-default (geeko@buildhost) (gcc (SUSE Linux) 11.2.
1 20210816 [revision 056e324ce46a7924b5cf10f61010cf9dd2ca10e9], GNU ld (GNU Binutils; SUSE Linux Enterprise 15
) 2.37.20211103-7.26) #1 SMP Thu Nov 25 09:36:40 UTC 2021 (83fc974)
[    0.000000][    T0] efi: EFI v2.70 by American Megatrends
[    0.000000][    T0] efi: ESRT=0xf9515018 SMBIOS=0xfe390000 SMBIOS 3.0=0xfe380000 ACPI 2.0=0xfd8d0000 MOKvar
=0xf7bd7000 MEMRESERVE=0xf4801798
[    0.000000][    T0] esrt: Reserving ESRT space from 0x00000000f9515018 to 0x00000000f9515050.
[    0.000000][    T0] ACPI: Early table checksum verification disabled
…
[   39.152724][    T1] pci_bus 0000:80: resource 4 [mem 0x60000000-0x7fffffff window]
[   39.160282][    T1] pci_bus 0000:80: resource 5 [mem 0x14000000000-0x17fffffffff window]
[   39.168366][    T1] pci_bus 0000:91: resource 1 [mem 0x60000000-0x600fffff]
[   39.214998][    T1] iommu: Default domain type: Passthrough
[   39.220723][    T1] pci 0000:0d:00.0: vgaarb: VGA device added: decodes=io+mem,owns=none,locks=none
[   39.229763][    T1] pci 0000:0d:00.0: vgaarb: bridge control possible
[   39.236201][    T1] pci 0000:0d:00.0: vgaarb: setting as boot device (VGA legacy resources not available)
[   39.245755][    T1] vgaarb: loaded
[   39.249503][    T1] SCSI subsystem initialized
[   39.254128][    T1] pps_core: LinuxPPS API ver. 1 registered
[   39.259785][    T1] pps_core: Software ver. 5.3.6 - Copyright 2005-2007 Rodolfo Giometti <giometti@linux.it
>
[   39.269652][    T1] PTP clock support registered
[   39.274282][    T1] EDAC MC: Ver: 3.0.0
[   39.278314][    T1] Registered efivars operations
[   39.287410][    T1] NetLabel: Initializing
[   39.291497][    T1] NetLabel:  domain hash size = 128
[   39.296541][    T1] NetLabel:  protocols = UNLABELED CIPSOv4 CALIPSO
[   39.302898][    T1] NetLabel:  unlabeled traffic allowed by default
<no further log messages>
```

I removed both workers from salt and paused the host-up alerts.

#21 - 2021-12-21 11:38 - kraih

I also gave power cycling arm-4 a try and for me it ended at a slightly different point:

```
...
[   27.159822][    T1] pci_bus 0000:0e: resource 1 [mem 0x43100000-0x432fffff]
[   27.196956][    T1] ARMH0011:00: ttyAMA0 at MMIO 0x402020000 root bus resource [mem 0x60000000-0x7fffffff w
indow]
[   38.675918][    T1] pci_bus 0000:80: root bus resource [mem 0x14000000000-0x17fffffffff window]
[   38.684606][    T1] pci_bus 0000:80: root bus resource [bus 80-ff]
[   38.690813][    T1] pci 0000:80:00.0: [177d:af00] type 00 class 0x060000
[   38.697639][    T1] pci 0000:80:01.0: [177d:af84] type 01 class 0x060400
[   38.704369][    T1] pci 0000:80:01.0: PME# supported from D0 D3hot D3cold
[   38.711273][    T1] pci 0000:80:02.0: [177d:af84] type 01 class 0x060400
```

```
[   38.718000][    T1] pci 0000:80:02.0: PME# supported from D0 D3hot D3cold
[   38.724911][    T1] pci 0000:80:03.0: [177d:af84] type 01 class 0x060400
[   38.731640][    T1] pci 0000:80:03.0: PME# supported from D0 D3hot D3cold
[   38.738539][    T1] pci 0000:80:04.0: [177d:af84] type 01 class 0x060400
[   38.745267][    T1] pci 0000:80:04.0: PME# supported from D0 D3hot D3cold
[   38.752161][    T1] pci 0000:80:05.0: [177d:af84] type 01 class 0x060400
[   38.758892][    T1] pcrom D0 D3hot D3cold
[   38.888346][    T1] pci 0000:80:0f.0: [14e4:9026] type 00 class 0x0c0330
[   38.895048][    T1] pci 0000:80:0f.0: reg 0x10: [mem 0x14000030000-0x1400003ffff 64bit pref]
[   38.903479][    T1] pci 0000:80:0f.0: reg 0x18: [mem 0x14000020000-0x1400002ffff 64bit pref]
[   38.911999][    T1] pci 0000:80:0f.1: [14e4:9026] type 00 class 0x0c0330
[   38.918695][    T1] pci 0000:80:0f.1: reg 0x10: [mem 0x14000010000-0x1400001ffff 64bit pref]
[   38.927129][    T1] pci 0000:80:0f.1: reg 0x18: [mem 0x14000000000-0x1400000ffff 64bit pref]
[   38.935693][    T1] acpiphp: Slot [1] registered
[   38.940380][    T1] acpiphp: Slot [1-1] registered
[   38.945240][    T1] acpiphp: Slot [1-2] registered
[   38.950096][    T1] acpiphp: Slot [1-3] registered
[   38.954941][    T1] pci 0000:91:00.0: [8086:0a54] type 00 class 0x010802
[   38.961645][    T1] pci 0000:91:00.0: reg 0x10: [mem 0x60000000-0x60003fff 64bit]
[   38.969148][    T1] pci 0000:91:00.0: reg 0x30: [mem 0xffff0000-0xffffffff pr
<no further log messages>
```

The machine does boot with the 5.14.14 kernel though. Upgrading to 5.15.10 did not work, gets stuck at the same point during boot.

#22 - 2021-12-21 13:48 - kraih

Downgraded both machines to the default Leap 15.3 kernel, so they are working again.

#23 - 2021-12-22 19:18 - okurz

- Status changed from Workable to Blocked

- Assignee set to okurz

Files

sysctl_diff.html	39.3 KB	2021-10-18	nicksinger