

openQA Infrastructure - action #96551

coordination # 96447 (Blocked): [epic] Failed systemd services and job age alerts

Persistent records of systemd journal size:S

2021-08-04 09:11 - cdywan

Status: Feedback	Start date: 2021-10-22
Priority: Low	Due date:
Assignee: okurz	% Done: 0%
Category:	Estimated time: 0.00 hour
Target version: Ready	

Description

Observation

When alerts are triggered there's a brief window where the journal can be checked but we don't seem to retain this long/ at all.

Acceptance criteria

- **AC1:** systemd journal is kept for 1 week

Suggestion

- Look up the man page of journal
- Add config to salt / <https://gitlab.suse.de/openqa/salt-states-openqa/>

Related issues:

Related to openQA Project - action #101379: Reduce amount of unhelpful log me...	Resolved	2021-10-22
Copied to openQA Infrastructure - action #96552: Persistent records of I/O us...	Workable	2021-08-04

History

#1 - 2021-08-04 09:16 - okurz

- Status changed from Workable to Feedback
- Assignee set to okurz

https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/535

#2 - 2021-08-04 09:19 - cdywan

- Parent task set to #96447

#3 - 2021-08-04 09:25 - okurz

- Due date set to 2021-08-18

#4 - 2021-08-04 09:27 - cdywan

- Copied to action #96552: Persistent records of I/O usage by process size:M added

#5 - 2021-08-15 20:38 - okurz

sudo journalctl --disk-usage says logs take up 2.9G but journalctl itself says

```
-- Logs begin at Sun 2021-07-18 03:30:00 CEST, end at Sun 2021-08-15 22:23:10 CEST. --
Jul 18 03:30:00 openqa sudo[25770]: pam_unix(sudo:session): session closed for user root
Jul 18 03:30:02 openqa systemd[25237]: Stopped target Default.
Jul 18 03:30:02 openqa systemd[25237]: Stopped target Basic System.
Jul 18 03:30:02 openqa systemd[25237]: Stopped target Timers.
Jul 18 03:30:02 openqa systemd[25237]: Stopped target Sockets.
Jul 18 03:30:02 openqa systemd[25237]: Stopped target Paths.
Jul 18 03:30:02 openqa systemd[25237]: Closed D-Bus User Message Bus Socket.
Jul 18 03:30:02 openqa systemd[25237]: Reached target Shutdown.
Jul 18 03:30:02 openqa systemd[25237]: Starting Exit the Session...
Jul 18 03:30:02 openqa systemd[25237]: Received SIGRTMIN+24 from PID 2783 (kill).
```

```

Jul 18 03:30:02 openqa systemd[25238]: pam_unix(systemd-user:session): session closed for user asmorodskiy
-- Reboot --
Aug 15 03:30:00 openqa systemd[29348]: Stopped target Default.
Aug 15 03:30:00 openqa systemd[29348]: Stopped target Basic System.
Aug 15 03:30:00 openqa systemd[29348]: Stopped target Sockets.
Aug 15 03:30:00 openqa systemd[29348]: Closed D-Bus User Message Bus Socket.
Aug 15 03:30:00 openqa systemd[29348]: Stopped target Timers.
Aug 15 03:30:00 openqa systemd[29348]: Stopped target Paths.
Aug 15 03:30:00 openqa systemd[29348]: Reached target Shutdown.
Aug 15 03:30:00 openqa systemd[29348]: Starting Exit the Session...
Aug 15 03:30:01 openqa systemd[29348]: Received SIGRTMIN+24 from PID 6482 (kill).
Aug 15 03:30:01 openqa systemd[29354]: pam_unix(systemd-user:session): session closed for user cdywan
Aug 15 03:30:01 openqa systemd[3598]: Stopped target Default.
Aug 15 03:30:01 openqa systemd[3598]: Stopped target Basic System.
Aug 15 03:30:01 openqa systemd[3598]: Stopped target Paths.
Aug 15 03:30:01 openqa systemd[3598]: Stopped target Sockets.
Aug 15 03:30:01 openqa systemd[3598]: Closed D-Bus User Message Bus Socket.
Aug 15 03:30:01 openqa systemd[3598]: Reached target Shutdown.
Aug 15 03:30:01 openqa systemd[3598]: Starting Exit the Session...
Aug 15 03:30:01 openqa systemd[3598]: Stopped target Timers.
Aug 15 03:30:01 openqa systemd[3598]: Received SIGRTMIN+24 from PID 6513 (kill).
Aug 15 03:30:01 openqa systemd[17836]: Stopped target Default.
Aug 15 03:30:01 openqa systemd[17836]: Stopped target Basic System.
Aug 15 03:30:01 openqa systemd[17836]: Stopped target Timers.
Aug 15 03:30:01 openqa systemd[17836]: Stopped target Sockets.
Aug 15 03:30:01 openqa systemd[17836]: Closed D-Bus User Message Bus Socket.
Aug 15 03:30:01 openqa systemd[17836]: Stopped target Paths.
Aug 15 03:30:01 openqa systemd[17836]: Reached target Shutdown.
Aug 15 03:30:01 openqa systemd[17836]: Starting Exit the Session...
Aug 15 03:30:01 openqa systemd[17836]: Received SIGRTMIN+24 from PID 6522 (kill).
-- Reboot --
Aug 15 18:00:01 openqa systemd[17066]: Reached target Timers.
Aug 15 18:00:01 openqa systemd[17066]: Starting D-Bus User Message Bus Socket.
Aug 15 18:00:01 openqa systemd[17066]: Reached target Paths.
Aug 15 18:00:01 openqa systemd[17066]: Listening on D-Bus User Message Bus Socket.
Aug 15 18:00:01 openqa systemd[17066]: Reached target Sockets.
Aug 15 18:00:01 openqa systemd[17066]: Reached target Basic System.
Aug 15 18:00:01 openqa systemd[17066]: Reached target Default.
Aug 15 18:00:01 openqa systemd[17066]: Startup finished in 137ms.
Aug 15 18:00:13 openqa systemd[17066]: Stopped target Default.
Aug 15 18:00:13 openqa systemd[17066]: Stopped target Basic System.
Aug 15 18:00:13 openqa systemd[17066]: Stopped target Timers.
Aug 15 18:00:13 openqa systemd[17066]: Stopped target Sockets.
Aug 15 18:00:13 openqa systemd[17066]: Closed D-Bus User Message Bus Socket.
Aug 15 18:00:13 openqa systemd[17066]: Stopped target Paths.
Aug 15 18:00:13 openqa systemd[17066]: Reached target Shutdown.
Aug 15 18:00:13 openqa systemd[17066]: Starting Exit the Session...
Aug 15 18:00:13 openqa systemd[17066]: Received SIGRTMIN+24 from PID 17847 (kill).
Aug 15 18:00:13 openqa systemd[17080]: pam_unix(systemd-user:session): session closed for user cvdupdate
Aug 15 18:03:27 openqa openqa-websockets-daemon[2219]: "HDDSIZEGB" => 40,
Aug 15 18:03:27 openqa openqa-websockets-daemon[2219]: "INCIDENT_ID" => 18273,
...

```

so there are some logs back to 2021-07-18 but then only since the last reboot. Also the entries look weird with the reboot at 03:30 as planned but then another line of "-- Reboot --" before "Aug 15 18:00:01" but uptime says that the machine is up since 18h:53m, that is 03:30 localtime.

Maybe the problem is openqa-webui which spams the logs. On o3 the logs go back way further and /var/log/openqa is configured to take the openQA logs, on osd it's not the case. Created https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/549

#6 - 2021-08-18 08:51 - okurz

merged the MR. Let's see if it helps to keep a longer system journal.

#7 - 2021-08-18 16:41 - okurz

- Due date changed from 2021-08-18 to 2021-08-20

/var/log/openqa has logs. The system journal still has a lot of information, e.g. from openqa-websockets-daemon, etc. Maybe I would need to restart these services as well? I will need to check on the next day if the journal persists now a bit longer. I don't think there is a better way than to simply wait hence bumping due-date.

#8 - 2021-08-19 05:45 - okurz

- Due date changed from 2021-08-20 to 2021-09-03

- Priority changed from Normal to Low

Again there is not much left in the journal:

```
-- Logs begin at Sun 2021-07-18 03:30:00 CEST, end at Thu 2021-08-19 07:38:01 CEST. --
Jul 18 03:30:00 openqa sudo[25770]: pam_unix(sudo:session): session closed for user root
Jul 18 03:30:02 openqa systemd[25237]: Stopped target Default.
Jul 18 03:30:02 openqa systemd[25237]: Stopped target Basic System.
Jul 18 03:30:02 openqa systemd[25237]: Stopped target Timers.
Jul 18 03:30:02 openqa systemd[25237]: Stopped target Sockets.
Jul 18 03:30:02 openqa systemd[25237]: Stopped target Paths.
Jul 18 03:30:02 openqa systemd[25237]: Closed D-Bus User Message Bus Socket.
Jul 18 03:30:02 openqa systemd[25237]: Reached target Shutdown.
Jul 18 03:30:02 openqa systemd[25237]: Starting Exit the Session...
Jul 18 03:30:02 openqa systemd[25237]: Received SIGRTMIN+24 from PID 2783 (kill).
Jul 18 03:30:02 openqa systemd[25238]: pam_unix(systemd-user:session): session closed for user asmorodskyi
-- Reboot --
Aug 19 00:34:01 openqa systemd[1]: Started Session 11796 of user geekotest.
Aug 19 00:34:01 openqa cron[8641]: pam_unix(crond:session): session opened for user geekotest by (uid=0)
Aug 19 00:34:04 openqa systemd-journald[23984]: Suppressed 5660 messages from /system.slice/openqa-websockets.service
Aug 19 00:34:11 openqa CRON[8641]: pam_unix(crond:session): session closed for user geekotest
```

/var/log/openqa is right now 137M but now the problem seems to be mainly openqa-websockets. systemctl status systemd-journald shows

```
Aug 19 07:35:30 openqa systemd-journald[14984]: Suppressed 2269 messages from /system.slice/openqa-websockets.service
Aug 19 07:36:00 openqa systemd-journald[14984]: Suppressed 2717 messages from /system.slice/openqa-websockets.service
```

so I restarted openqa-websockets as well. Looks like now the messages end up in /var/log/openqa as well.

I am not sure if the settings controlled by salt are effective now so to be sure I additionally set the settings for now with vim /etc/systemd/journald.conf and did systemctl restart systemd-journald. Now systemctl status systemd-journald shows System journal (/var/log/journal/84f4f4f356b525388b60f0ae547597e0) is 2.4G, max 23.7G, 21.2G free. But I removed the settings again, only keeping in /etc/systemd/journald.conf.d/journal_size.conf and after systemctl restart systemd-journald the command systemctl status systemd-journald shows System journal (/var/log/journal/84f4f4f356b525388b60f0ae547597e0) is 2.4G, max 25.1G, 22.7G free. so that should be fine. Let's see how this changes over the next days. journalctl --disk-usage says Archived and active journals take up 2.4G in the file system.

As the original problem about I/O alerts seem to be less severe by now I will follow up with this with low prio.

#9 - 2021-08-26 12:30 - okurz

- Due date changed from 2021-09-03 to 2021-10-08

As of now the latest in systemctl status systemd-journald is still

```
Aug 20 11:59:21 openqa systemd-journald[8460]: System journal (/var/log/journal/84f4f4f356b525388b60f0ae547597e0) is 2.0G, max 22.5G, 20.4G free.
```

so looks good. journalctl --disk-usage reports 1.3G. journalctl goes back to Aug 19. I want to wait over the next automatic reboots if this still works, then revert the openQA log target and check again.

#10 - 2021-09-02 20:59 - okurz

osd was rebooted 4 days ago. journal size is reported to take 1.2G, systemctl status systemd-journald says "Warning: Journal has been rotated since unit was started. Log output is incomplete or unavailable.", not that helpful. journalctl says

```
-- Logs begin at Sun 2021-07-18 03:30:00 CEST, end at Thu 2021-09-02 22:24:00 CEST. --
Jul 18 03:30:00 openqa sudo[25770]: pam_unix(sudo:session): session closed for user root
Jul 18 03:30:02 openqa systemd[25237]: Stopped target Default.
Jul 18 03:30:02 openqa systemd[25237]: Stopped target Basic System.
Jul 18 03:30:02 openqa systemd[25237]: Stopped target Timers.
Jul 18 03:30:02 openqa systemd[25237]: Stopped target Sockets.
Jul 18 03:30:02 openqa systemd[25237]: Stopped target Paths.
Jul 18 03:30:02 openqa systemd[25237]: Closed D-Bus User Message Bus Socket.
Jul 18 03:30:02 openqa systemd[25237]: Reached target Shutdown.
Jul 18 03:30:02 openqa systemd[25237]: Starting Exit the Session...
Jul 18 03:30:02 openqa systemd[25237]: Received SIGRTMIN+24 from PID 2783 (kill).
Jul 18 03:30:02 openqa systemd[25238]: pam_unix(systemd-user:session): session closed for user asmorodskyi
-- Reboot --
Aug 31 09:39:57 openqa systemd[6377]: Reached target Timers.
...
```

so there are still some entries from July 18 but then only after August 31 though the last reboot was August 29? Restarted "systemd-journald" again and status says "System journal (/var/log/journal/84f4f4f356b525388b60f0ae547597e0) is 1.2G, max 21.1G, 19.8G free"

So it does not look like the 20G would really be used.

Regardless of this I see that vsftpd plays a big part in filling the logs, e.g. entry blocks like

```
Sep 02 00:18:46 openqa vsftpd[18289]: [ftp] OK DOWNLOAD: Client "::ffff:10.160.0.245", "/SLE-15-SP4-Module-Bas
esystem-POOL-aarch64-Build31.2-Media1/aarch64/dhcp-client-4.3.6.P1-6.11.1.aarch64.rpm", 747476 bytes, 34388.14
Kbyte/sec
Sep 02 00:18:46 openqa vsftpd[18289]: [ftp] OK DOWNLOAD: Client "::ffff:10.160.0.245", "/SLE-15-SP4-Module-Bas
esystem-POOL-aarch64-Build31.2-Media1/aarch64/dracut-049.1+suse.203.g8ee14a90-3.35.1.aarch64.rpm", 442592 byte
s, 54669.71Kbyte/sec
Sep 02 00:18:46 openqa vsftpd[18289]: [ftp] OK DOWNLOAD: Client "::ffff:10.160.0.245", "/SLE-15-SP4-Module-Bas
esystem-POOL-aarch64-Build31.2-Media1/aarch64/util-linux-systemd-2.36.2-2.1.aarch64.rpm", 166636 bytes, 74991.
00Kbyte/sec
Sep 02 00:18:46 openqa vsftpd[18289]: [ftp] OK DOWNLOAD: Client "::ffff:10.160.0.245", "/SLE-15-SP4-Module-Bas
esystem-POOL-aarch64-Build31.2-Media1/aarch64/postfix-3.5.9-5.3.3.aarch64.rpm", 1164512 bytes, 76068.14Kbyte/s
ec
Sep 02 00:18:46 openqa vsftpd[18289]: [ftp] OK DOWNLOAD: Client "::ffff:10.160.0.245", "/SLE-15-SP4-Module-Bas
esystem-POOL-aarch64-Build31.2-Media1/aarch64/device-mapper-1.02.163-8.33.1.aarch64.rpm", 142104 bytes, 30634.
31Kbyte/sec
Sep 02 00:18:46 openqa vsftpd[18289]: [ftp] OK DOWNLOAD: Client "::ffff:10.160.0.245", "/SLE-15-SP4-Module-Bas
esystem-POOL-aarch64-Build31.2-Media1/aarch64/mdadm-4.1-24.3.1.aarch64.rpm", 473004 bytes, 24733.24Kbyte/sec
Sep 02 00:18:46 openqa vsftpd[18289]: [ftp] OK DOWNLOAD: Client "::ffff:10.160.0.245", "/SLE-15-SP4-Module-Bas
esystem-POOL-aarch64-Build31.2-Media1/aarch64/liblvm2cmd2_03-2.03.05-8.33.1.aarch64.rpm", 740500 bytes, 47000.
16Kbyte/sec
Sep 02 00:18:46 openqa vsftpd[18289]: [ftp] OK DOWNLOAD: Client "::ffff:10.160.0.245", "/SLE-15-SP4-Module-Bas
esystem-POOL-aarch64-Build31.2-Media1/aarch64/kpartx-0.8.6+10+suse.4771137-1.1.aarch64.rpm", 98888 bytes, 4977
8.51Kbyte/sec
Sep 02 00:18:47 openqa vsftpd[18289]: [ftp] OK DOWNLOAD: Client "::ffff:10.160.0.245", "/SLE-15-SP4-Module-Bas
esystem-POOL-aarch64-Build31.2-Media1/aarch64/lvm2-2.03.05-8.33.1.aarch64.rpm", 1227116 bytes, 32950.82Kbyte/s
ec
Sep 02 00:18:47 openqa kernel: vsftpd: page allocation failure: order:5, mode:0x40dc0(GFP_KERNEL|__GFP_COMP|__
GFP_ZERO), nodemask=(null),cpuset=/,mems_allowed=0
Sep 02 00:18:47 openqa kernel: CPU: 4 PID: 19545 Comm: vsftpd Not tainted 5.3.18-lp152.87-default #1 openSUSE
Leap 15.2
Sep 02 00:18:47 openqa kernel: Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS Bochs 01/01/2011
Sep 02 00:18:47 openqa kernel: Call Trace:
Sep 02 00:18:47 openqa kernel: dump_stack+0x66/0x8b
Sep 02 00:18:47 openqa kernel: warn_alloc+0x10b/0x190
Sep 02 00:18:47 openqa kernel: ? __alloc_pages_direct_compact+0x155/0x160
Sep 02 00:18:47 openqa kernel: __alloc_pages_slowpath+0xcc3/0xd00
Sep 02 00:18:47 openqa kernel: __alloc_pages_nodemask+0x2c4/0x300
Sep 02 00:18:47 openqa kernel: kmalloc_order+0x18/0x70
Sep 02 00:18:47 openqa kernel: kmalloc_order_trace+0x1d/0xa0
Sep 02 00:18:47 openqa kernel: __kmalloc+0x25f/0x2c0
Sep 02 00:18:47 openqa kernel: ? _cond_resched+0x15/0x40
Sep 02 00:18:47 openqa kernel: nfsd_reply_cache_init+0xb2/0x140 [nfsd]
Sep 02 00:18:47 openqa kernel: nfsd_init_net+0x5d/0x100 [nfsd]
Sep 02 00:18:47 openqa kernel: ops_init+0x40/0x110
Sep 02 00:18:47 openqa kernel: setup_net+0xd2/0x210
Sep 02 00:18:47 openqa kernel: copy_net_ns+0xf5/0x230
Sep 02 00:18:47 openqa kernel: create_new_namespaces+0x124/0x2a0
Sep 02 00:18:47 openqa kernel: copy_namespaces+0x7e/0xd0
Sep 02 00:18:47 openqa kernel: copy_process+0xc4b/0x1b10
Sep 02 00:18:47 openqa kernel: ? __alloc_file+0x8d/0xf0
Sep 02 00:18:47 openqa kernel: _do_fork+0x74/0x3b0
Sep 02 00:18:47 openqa kernel: ? do_sigaction+0x1c9/0x250
Sep 02 00:18:47 openqa kernel: __x64_sys_clone+0x8f/0xc0
Sep 02 00:18:47 openqa kernel: do_syscall_64+0x65/0x1f0
Sep 02 00:18:47 openqa kernel: entry_SYSCALL_64_after_hwframe+0x44/0xa9
Sep 02 00:18:47 openqa kernel: RIP: 0033:0x7f01a3f55759
```

showing both a lot of download status messages as well as kernel reports about "vsftpd: page allocation failure" and kernel stack traces. Couldn't find anything useful for the kernel stack trace in a web search.

Maybe making the vsftpd logs less noisy can help?: https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/571

#11 - 2021-09-23 21:16 - okurz

Taking look right now I see that the journal goes back only 3 days so again and still a problem. journalctl -u systemd-journald says

```
Sep 23 20:14:26 openqa systemd-journald[23151]: System journal (/var/log/journal/84f4f4f356b525388b60f0ae54759
7e0) is 1.2G, max 19.4G, 18.1G free.
```

so we do not use the available space at all. What's going on? I don't know. I suspected that maybe some other process using the same filesystem, e.g. postgresSQL, would use more space temporarily so that journalctl erroneously deletes too much but according to the man page of systemd-journald it does not behave like that and also looking at <https://monitor.qa.suse.de/d/WebuiDb/webui-summary?viewPanel=74&orgId=1&from=now-7d&to=now> I don't see that we would have peaks in bigger space usage.

There is rotation, the retention is configured to store max 80G. Hm, maybe the many user sessions cause the max files number in combination with size per rotated archive file to exceed?

https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/589

#12 - 2021-09-28 10:49 - okurz

no review was conducted :(merged now.

#13 - 2021-09-30 20:36 - okurz

All of the above still does not seem to work. Now the journal got truncated again to just 2 days worth of history. Trying to increase number of rotated files now: https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/598

Also temporarily I disabled logging of vsftpd to the system journal with `syslog_enable=NO` in `/etc/vsftpd.conf`

#14 - 2021-10-07 13:19 - cdywan

okurz wrote:

All of the above still does not seem to work. Now the journal got truncated again to just 2 days worth of history. Trying to increase number of rotated files now: https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/598

Also temporarily I disabled logging of vsftpd to the system journal with `syslog_enable=NO` in `/etc/vsftpd.conf`

```
> sudo journalctl --disk-usage
Archived and active journals take up 1.8G in the file system.
```

Also

```
> sudo journalctl
-- Logs begin at Tue 2021-09-28 01:30:09 CEST, end at Thu 2021-10-07 15:18:04 CEST. --
```

So it seems like you achieved something here?

#15 - 2021-10-08 06:53 - okurz

- Due date changed from 2021-10-08 to 2021-11-05

yes, that part seems to work. I would like to wait further until we see at least 2GiB of used space exceeded. Then we can slowly move back services from extra log files to systemd-journal

#16 - 2021-10-22 14:50 - kraih

- Related to action #101379: Reduce amount of unhelpful log messages at debug level added

#17 - 2021-11-16 12:25 - okurz

I reported https://bugzilla.opensuse.org/show_bug.cgi?id=1192748 about the kernel stack trace from vsftpd.

Next to that I see many lines like

```
Nov 16 12:57:01 openqa cron[29088]: pam_unix(cron:session): session opened for user geekotest by (uid=0)
Nov 16 12:57:01 openqa cron[29089]: pam_unix(cron:session): session opened for user geekotest by (uid=0)
Nov 16 12:57:02 openqa CRON[29088]: pam_unix(cron:session): session closed for user geekotest
Nov 16 12:57:09 openqa CRON[29089]: pam_unix(cron:session): session closed for user geekotest
```

a solution might be in

<https://languor.us/cron-pam-unix-cron-session-session-opened-closed-user-root-uid0>

but I have not followed up with this so far.

So currently journalctl goes back until 2021-10-11 so more than a month. That's good.

`systemctl restart systemd-journald` ; `systemctl status systemd-journald` says "System journal (/var/log/journal/84f4f4f356b525388b60f0ae547597e0) is 1.5G, max 46.7G, 45.1G free." so it's not using that much space. But I also wonder if it shouldn't use more space by now.

We previously had a daily logrotation for openQA logs, now we are currently using hourly log rotation with up to 20 archives stored meaning we are only storing openQA logs up to 21h, not longer.

Proposing to re-enable openQA logs in system journal with
https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/613