

## openQA Infrastructure - action #73333

### Failed systemd services alert (workers) flaky

2020-10-14 07:07 - nicksinger

<b>Status:</b> Resolved	<b>Start date:</b> 2020-10-14
<b>Priority:</b> High	<b>Due date:</b> 2020-10-22
<b>Assignee:</b> okurz	<b>% Done:</b> 0%
<b>Category:</b>	<b>Estimated time:</b> 0.00 hour
<b>Target version:</b> Ready	

**Description**

#### Observation

In the last 12h we had quite some alerts for failing systemd services on ~~the worker~~ a host. Looking at <https://stats.openqa-monitor.qa.suse.de/d/KToPYLEWz/failed-systemd-services?orgId=1&panelId=6&fullscreen&edit&tab=alert&from=1602604933324&to=1602655964690> it seems like one service is repeatedly failing and recovering. The alert stated values for systemd\_failed.sum between 1.2 and 0.167 which I find kind of confusing and is a result how we sample the data

#### Expected result

- Alert does not fail on flaky [user@486.service](#) on staging-1.qa.suse.de

#### Suggestions

- check on staging-1.qa.suse.de why service "[user@486.service](#)" is failing, e.g. journalctl -u user@486.service, and fix that or prevent the alert, e.g. by disabling/masking telegraf on that host.

#### History

##### #1 - 2020-10-14 10:34 - okurz

- Target version set to Ready

##### #2 - 2020-10-17 20:04 - okurz

- Category set to Concrete Bugs

##### #3 - 2020-10-20 13:52 - nicksinger

the failing service in question is user@486.service on staging-1. We might still need to refine this alert to provide better insights

##### #4 - 2020-10-20 14:05 - okurz

- Tags changed from alert to alert, monitoring, systemd
- Project changed from openQA Project to openQA Infrastructure
- Description updated
- Category deleted (Concrete Bugs)
- Status changed from New to Workable
- Priority changed from Normal to Urgent

We should prevent the alert failing on that specific host to remove the urgency for now. Next step after that would be to improve the information we have available in the alert. However I would be ok to call this ticket done after we resolved the immediate issue except someone has a really good idea of how to improve the alert. We probably all agree that it could look better but wishful thinking also does not bring us further ;)

##### #5 - 2020-10-20 14:18 - okurz

- Due date set to 2020-10-22
- Status changed from Workable to Feedback
- Assignee set to okurz
- Priority changed from Urgent to High

created [https://gitlab.suse.de/openqa/salt-states-openqa/-/merge\\_requests/383](https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/383) to prevent the confusion due to the name of the alert and monitoring panel.

I assume that mkittler applied a complete "webui salt state" on that machine and that included even something like /etc/cron.d/SLES.CRON running some cron jobs which likely fail.

I removed

```
roles:  
- webui
```

from /etc/salt/grains, restarted salt-master service and did rm /etc/cron.d/SLES.CRON. Setting to feedback and waiting if the alert still triggers.

#### **#6 - 2020-10-20 17:58 - okurz**

I have another change prepared to cover the suggestions from nsinger to improve the output in case of alerts. But I can not push these changes to gitlab because my account seems to be blocked now, reported in <https://infra.nue.suse.com/Service/Display.html?id=178670>

#### **#7 - 2020-10-22 21:13 - okurz**

- *Status changed from Feedback to Resolved*

All necessary changes are in and no alert triggered again.