

openQA Infrastructure - action #80538

flaky and misleading alerts about "openQA minion workers alert" as well as "Minion Jobs alert"

2020-11-27 18:09 - okurz

Status:	Resolved	Start date:	2020-11-27
Priority:	Normal	Due date:	2021-04-14
Assignee:	mkittler	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	Ready		

Description

Observation

On 2020-11-26 and 2020-11-27 there were the following alerts and automatic ok-messages:

- 2020-11-26 2045Z: [Alerting] openQA minion workers alert
- 2020-11-26 2046Z: [OK] openQA minion workers alert
- 2020-11-26 2106Z: [OK] openqaworker8: Minion Jobs alert
- 2020-11-27 0140Z: [Alerting] QA-Power8-4-kvm: Minion Jobs alert
- 2020-11-27 0255Z: [OK] QA-Power8-4-kvm: Minion Jobs alert

I assume no one did anything about these alerts to go to "OK" again.

the alert "openQA minion workers alert" has the description:

```
Minion workers down. Check systemd services on the openQA host
```

and the alert "\$worker: Minion Jobs alert" has the description:

```
to remove all failed jobs on the machine `` /usr/share/openqa/script/openqa-workercache eval 'my $jobs = app->minion->jobs({states => ["failed"]}); while (my $job = $jobs->next) { $job->remove }' ``
```

Problems

- meaning of alert is unclear
- alerts should be stable, i.e. not flaky
- description does not explain enough so that team members would understand what needs to be done and what the meaning is

Acceptance criteria

- **AC1:** Multiple team members can confirm that updated name and/or description explains them what to do
- **AC2:** Allow at least one other team member to handle the alert based on documentation
- **AC3:** The alerts are stable and only persist when there is a problem

Suggestions

- Look at the [o3 minion dashboard](#)
- Locate the alert on [grafana](#)
- Understand what the data source for the alert is and how "minion jobs" are involved
- Research what is the expected state for "minion jobs"
- Extend the description to explain, optionally reference this ticket
- Consider updating the name of the alert as well
- Stabilize alert

Related issues:

Related to openQA Project - action #89560: Add alert for blocked gitlab accou...	Workable	2021-03-05
Has duplicate openQA Infrastructure - action #78061: [Alerting] openQA minion...	Rejected	2020-11-16

History

#1 - 2021-01-21 12:06 - cdywan

- Description updated

- Priority changed from High to Normal

I lowered the *Priority* since I've not seen it at all since. Determining where the actual definition is is part of the ticket, so if there is something that was deleted or disabled I wouldn't know. But I added the obvious places to check.

#2 - 2021-01-22 12:27 - okurz

- Description updated

We discussed that what you wrote as "acceptance criteria" are suggestions which we already have. I changed it back to be real criteria but kept your addition. The project <https://gitlab.suse.de/openqa/grafana-webhook-actions> does not contain any related alert definitions.

#3 - 2021-02-01 10:42 - cdywan

[Alerting] openQA minion workers alert

Minion workers down. Check systemd services on the openQA host

Metric name
Sum

Value
0.984

Saw this alert, followed by an [OK]. I don't know what 0.984 tells me here. But it seems like it's still flaky since I didn't see a ticket or chat conversations about someone having fixed problems. Please correct me if I'm wrong here.

#4 - 2021-02-04 10:52 - cdywan

Metric name
Sum

Value
0.995

As seen today, also followed by an OK shortly after, like the previous one. On a side note the email didn't render properly in TB either.

#5 - 2021-03-04 13:02 - cdywan

Can anyone remind me why [#78061](#) is not a duplicate of this?

#6 - 2021-03-04 22:18 - okurz

- Has duplicate action [#78061](#): [Alerting] openQA minion workers alert - alert turned "OK" again after 20 minutes and we don't know what was wrong added

#7 - 2021-03-05 11:44 - cdywan

Discussion during the weekly:

- cache service down? possibly the gru service down?
- Hypothesis, influxdb
- <https://stats.openqa-monitor.qa.suse.de/d/WebuiDb/webui-summary?orgId=1&editPanel=17&tab=alert>
- [mkittler](#) claims to have an idea and volunteered to pick it up
- <https://github.com/os-autoinst/openQA/blob/master/lib/OpenQA/CacheService/Controller/Influxdb.pm>
- <https://github.com/os-autoinst/openQA/blob/master/lib/OpenQA/WebAPI/Controller/Admin/Influxdb.pm>

#8 - 2021-03-05 12:26 - okurz

- Related to action [#89560](#): Add alert for blocked gitlab account when users are unable to save/commit needles added

#9 - 2021-03-30 10:43 - mkittler

- Status changed from Workable to In Progress

- Assignee set to mkittler

#10 - 2021-03-30 11:37 - mkittler

SR: https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/469

Not sure what else to do (see SR description).

#11 - 2021-03-31 04:15 - openqa_review

- Due date set to 2021-04-14

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

#12 - 2021-03-31 14:55 - mkittler

- Status changed from *In Progress* to *Feedback*

I'm waiting for feedback from the team.

#13 - 2021-04-01 10:29 - okurz

- Status changed from *Feedback* to *Resolved*

MR was merged and as effective on <https://monitor.qa.suse.de/d/WebuiDb/webui-summary?orgId=1&editPanel=19&tab=alert> . ilausuch has also confirmed that the message is clear to him.