

openQA Infrastructure - action #89821

alert: PROBLEM Service Alert: openqa.suse.de/fs_/srv is WARNING (flaky, partial recovery with OK messages)

2021-03-10 08:41 - okurz

Status:	Resolved	Start date:	2021-03-10
Priority:	Normal	Due date:	
Assignee:	mkittler	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	Ready		

Description

Observation

Multiple alert email reports:
Notification: PROBLEM
Host: openqa.suse.de
State: WARNING
Date/Time: Tue Mar 9 13:17:18 UTC 2021
Info: WARN - 80.1% used (64.06 of 79.99 GB), trend: +573.77 MB / 24 hours

Service: fs_/srv

See Online: https://thruk.suse.de/thruk/cgi-bin/extinfo.cgi?type=2&host=openqa.suse.de&service=fs_%2Fsrv

Acceptance criteria

- **AC1:** /srv on osd has enough free space
- **AC2:** alert is handled
- **AC3:** icinga alert is only triggering if internal grafana alert is not handled or not effective

Suggestions

- Follow the above thruk link to understand the monitoring data
- Crosscheck alert limit "80%" with the limit we have in grafana
- Make sure the grafana limit is smaller
- Ensure there is enough space, e.g. ask EngInfra to increase or cleanup

History

#1 - 2021-03-10 08:41 - okurz

- Tags set to alert, thruk, icinga, srv, grafana, osd, postgres, storage, space

#2 - 2021-03-10 09:25 - mkittler

Looks like the relevant device is /dev/vdb so this is not about assets/results:

```
martchus@openqa:~> df -h / /srv/ /var/lib/openqa/share /var/lib/openqa/testresults
Dateisystem Größe Benutzt Verf. Verw% Eingehängt auf
/dev/vdal 20G 8,1G 11G 44% /
/dev/vdb 80G 61G 20G 77% /srv
/dev/vdc 7,0T 5,6T 1,5T 80% /var/lib/openqa/share
/dev/vdd 5,5T 4,6T 1021G 82% /var/lib/openqa
```

In fact, it might be about the PostgreSQL database:

```
martchus@openqa:~> mount | grep /dev/vdb
/dev/vdb on /srv type xfs (rw,noatime,nodiratime,attr2,inode64,logbufs=8,logbsize=256k,noquota)
/dev/vdb on /var/lib/pgsql type xfs (rw,noatime,nodiratime,attr2,inode64,logbufs=8,logbsize=256k,noquota)
```

#3 - 2021-03-10 11:04 - okurz

Yes, /srv is the filesystem where we store mostly the database but a bit of other data as well. Sorry if that was not clear. I have mentioned "postgres" only as a tag, not in the text of the description

#4 - 2021-03-10 15:20 - mkittler

It is really mostly the database:

```
--- /srv
 54,0 GiB [#####] /PSQL10
  6,2 GiB [#          ] /log

  1,3 GiB [          ] homes.img
 23,1 MiB [          ] /salt
 10,0 MiB [          ] /pillar
```

We've already shared some figures within the chat and the indexes using most of the disk space. Here are the commands I've used to check this out: <https://github.com/Martchus/openQA-helper#show-postgresql-table-sizes>

Not sure whether we can easily improve/optimize this. We also likely can't just drop most of the indexes because they are actually there for a reason.

So I'd follow the last suggestion for now and asked infra to increase the size of /dev/vdb, e.g. to 100 GiB (so far we have 80 GiB).

#5 - 2021-03-10 18:16 - okurz

So I'd follow the last suggestion for now and asked infra to increase the size of /dev/vdb, e.g. to 100 GiB (so far we have 80 GiB).

Good. I assume you did so with ticket. Can you please reference the ticket? And did you include osd-admins@suse.de in CC?

#6 - 2021-03-11 09:24 - mkittler

- Assignee set to mkittler

I haven't done anything because I wanted to wait for at least one reply within the team. I'll do it now then.

#7 - 2021-03-12 10:54 - mkittler

- Status changed from Workable to Blocked

<https://infra.nue.suse.com/SelfService/Display.html?id=186825>

#8 - 2021-03-12 15:49 - mkittler

- Status changed from Blocked to Feedback

- Priority changed from Urgent to Normal

```
martchus@openqa:~> df -h /dev/vdb
Dateisystem  Größe Benutzt Verf. Verw% Eingehängt auf
/dev/vdb     80G   62G   19G   78% /srv
martchus@openqa:~> sudo xfs_growfs /var/lib/pgsql
meta-data=/dev/vdb          isize=256    agcount=16, agsize=1310720 blks
           =                sectsz=4096  attr=2,   projid32bit=1
           =                crc=0      finobt=0  spinodes=0 rmapbt=0
           =                reflink=0
data       =                bsize=4096 blocks=20971520, imaxpct=25
           =                sunit=0    swidth=0 blks
naming     =version 2      bsize=4096 ascii-ci=0 ftype=0
log        =Intern      bsize=4096 blocks=2560, version=2
           =                sectsz=4096 sunit=1 blks, lazy-count=1
realtime   =keine      extsz=4096 blocks=0, rtextents=0
Datenblöcke von 20971520 auf 26214400 geändert.
martchus@openqa:~> df -h /dev/vdb
Dateisystem  Größe Benutzt Verf. Verw% Eingehängt auf
/dev/vdb     100G   62G   39G   62% /srv
```

They gave us more space so I guess we're good for now (AC1 and AC2).

About AC3: Should I change the threshold in our monitoring from 90 % to 80 % or should I asked Infra to change the threshold in their monitoring from 80 % to 90 %. Or we just keep the 2 alerts differently so the Infra alert serves as initial alert and our own alert as a last reminder.

#9 - 2021-03-12 16:06 - okurz

About AC3: Should I change the threshold in our monitoring from 90 % to 80 % or should I asked Infra to change the threshold in their monitoring from 80 % to 90 %. Or we just keep the 2 alerts differently so the Infra alert serves as initial alert and our own alert as a last reminder.

I suggest to keep the 2 alerts but the other way around: Our alert first on 80% and the EngInfra alert on 90%, i.e. ask 'em to bump to 90%.

#10 - 2021-03-12 16:35 - mkittler

I've asked them: <https://infra.nue.suse.com/SelfService/Display.html?id=186868>

#11 - 2021-03-29 13:10 - okurz

- *Status changed from Feedback to Blocked*

blocked by <https://infra.nue.suse.com/SelfService/Display.html?id=186868>

#12 - 2021-04-13 08:13 - okurz

- *Status changed from Blocked to Resolved*

Confirmed with Daniel Rodríguez,

https://thruk.suse.de/thruk/cgi-bin/extinfo.cgi?type=2&host=openqa.suse.de&service=fs_/srv#pnp_th2/1618211527/1618301527/0 looks good