

openQA Infrastructure - action #44078

coordination # 64746 (Blocked): [saga][epic] Scale up: Efficient handling of large storage to be able to run current tests efficiently but keep big archives of old results

coordination # 102710 (Blocked): [epic] Improve our backup

Implement proper backups for o3 size:M

2018-11-20 22:16 - okurz

Status:	In Progress	Start date:	2018-11-20
Priority:	Low	Due date:	
Assignee:	mkittler	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	Ready		
Description			
Motivation			
We should find a backup space for o3, e.g. 4TB for test data/assets			
Acceptance criteria			
<ul style="list-style-type: none">• AC1: Test assets are backed up and can be restored• AC2: Test results are backed up and can be restored• AC3: Screenshots are backed up and can be restored			
Suggestions			
<ul style="list-style-type: none">• Use storage.qa.suse.de• We already have database backups elsewhere• Consider snapshots (if available) or rsync			
Out of scope			
<ul style="list-style-type: none">• /etc -> already covered by https://gitlab.suse.de/qa-sle/backup-server-salt• database -> already covered by https://gitlab.suse.de/qa-sle/backup-server-salt• List of installed packages -> #102713			
Further details			
<ul style="list-style-type: none">• Current backup is explained in https://progress.opensuse.org/projects/openqav3/wiki#Backup• See for reference the existing salt recipes in https://gitlab.suse.de/qa-sle/backup-server-salt			
Related issues:			
Related to openQA Infrastructure - action #69577: Handle installation of the ...		Resolved	2020-08-04
Related to openQA Infrastructure - action #88546: Make use of the new "Storag...		Blocked	
Related to openQA Infrastructure - action #96269: Define what a "complete OSD...		Blocked	2021-07-29
Copied to openQA Infrastructure - action #94015: proper backup for osd		New	
Copied to openQA Infrastructure - action #102713: Implement proper backups fo...		Feedback	

History

#2 - 2019-01-21 10:51 - okurz

As o3 is in a special network I recommend rsnapshot on backup.qa.suse.de syncing the data from o3.

#3 - 2019-01-31 08:49 - okurz

Or could we just add the host to <https://gitlab.suse.de/buildops/infra/tree/master/backup> ?

#4 - 2019-02-26 06:57 - okurz

- Status changed from New to In Progress

- Assignee set to okurz

Talked with tbro. 4TB is a problem. Anything in the range of up to 100GB is no problem to have, using space on old netapp, on request. However, I think then we can just as easily go with backup.qa which is easier for us given that we have full control of that machine and can trigger backups from there as suggested in [#44078#note-2](#)

#5 - 2019-03-01 08:39 - okurz

- Status changed from In Progress to Workable

- Assignee deleted (okurz)

Unassigning again as preparation for longer absence. I should really not leave tickets assigned to me dangling "In Progress" :)

#7 - 2019-03-18 15:27 - okurz

Automatic backup for the o3 webui host introduced with <https://gitlab.suse.de/okurz/backup-server-salt/tree/master/rsnapshot> covering so far /etc and the SQL database dumps. As next steps I recommend to save a bit more if feasible from /var/lib/openqa as well as from the workers which are transactional-servers. Maybe as well just /etc plus an autoyast profile.

#8 - 2020-07-28 09:47 - okurz

- Priority changed from Normal to Low

- Target version set to Ready

#9 - 2020-10-18 12:41 - okurz

- Related to action #69577: Handle installation of the new "Storage Server" added

#10 - 2020-10-18 12:42 - okurz

- Status changed from Workable to Blocked

- Assignee set to okurz

waiting for [#69577](#) or #41918 first

#11 - 2021-06-15 09:48 - okurz

- Copied to action #94015: proper backup for osd added

#12 - 2021-06-28 11:35 - okurz

- Status changed from Blocked to New

- Assignee deleted (okurz)

storage.qa.suse.de in place and usable

#13 - 2021-07-14 09:42 - cdywan

- Subject changed from proper backup for o3 to Implement proper backups for o3 size:M

- Description updated

- Status changed from New to Workable

#14 - 2021-07-29 12:03 - okurz

- Related to action #88546: Make use of the new "Storage Server", e.g. complete OSD backup added

#15 - 2021-07-29 12:04 - okurz

- Related to action #96269: Define what a "complete OSD backup" should or can include added

#16 - 2021-08-27 10:07 - okurz

- Description updated

#17 - 2021-08-27 10:14 - okurz

- Description updated

#18 - 2021-08-27 10:14 - okurz

- Description updated

#19 - 2021-08-27 10:20 - cdywan

I'm taking it now. Discussing it in the weekly I feel like it might be easier to treat it as a spike and propose something, and then see if we're happy with that :-D

#20 - 2021-08-27 10:20 - cdywan

- Assignee set to cdywan

#21 - 2021-08-27 15:33 - cdywan

- AC4: List of installed packages is backed up for reference

```
$ rpm -qa
[...]
$ zypper in gnome-maps-40.1-1.2.x86_64
[...]
Loading repository data...
Reading installed packages...
'gnome-maps.x86_64 = 40.1-1.2' is already installed.
There is an update candidate 'gnome-maps-40.4-1.1.x86_64' for 'gnome-maps-40.1-1.2.x86_64', but it does not match the specified version, architecture, or repository.
$ rpm -qa > packages.txt; sudo zypper in -D $(sed 's/-[a-z].*//' packages.txt)
[...]
$ zypper -x se --type=package -i
[...]
$ pip3 install --user yq
[...]
$ zypper -x se --type=package -i > packages.xml
$ sudo zypper in -D $(~/local/bin/xq '.stream|"search-result"|"solvable-list"|"solvable|.|]|"@name"' packages.xml | sed 's/"//g')
```

Proof of concept using [xq](#) to process XML output from zypper seems to work for (re)storing the list of installed packages.

#22 - 2021-08-30 11:34 - okurz

I wasn't aware about "yq", sounds great! But why do you care about this? rpm -qa should suffice and we don't want to recover systems from this list, just have it ready as a reference. I see now that in https://progress.opensuse.org/journals/427051/diff?detail_id=404851 you added the "list of packages". As this is a rather short document I suggest to just include it in the existing rsnapshot config for backup.qa.suse.de, see <https://gitlab.suse.de/qa-sle/backup-server-salt/-/blob/master/rsnapshot/rsnapshot.conf#L27> for how we call commands to backup SQL. I suggest you exclude the part about "list of packages" into another ticket or just do it to get to the more demanding/interesting parts of screenshots+assets+results for which we should use storage.qa.suse.de as target

#23 - 2021-09-02 15:32 - cdywan

- Assignee deleted (cdywan)

okurz wrote:

I wasn't aware about "yq", sounds great! But why do you care about this? rpm -qa should suffice and we don't want to recover systems from this list, just have it ready as a reference.

Like I already said on the call, "proper backups" includes being able to recover from said backups in my book. And the other ACs mention that as well. I don't know what "a reference" means in this context. I'll unassign since it seems we're still unclear on the goal of the ticket.

#24 - 2021-11-09 14:49 - mkittler

- Assignee set to mkittler

- Target version deleted (Ready)

One simple option would be a rsnapshot setup like on <https://gitlab.suse.de/qa-sle/backup-server-salt/-/tree/master/rsnapshot>. It would be simple in the sense that we wouldn't need to invent something new and we wouldn't rely on specific filesystem features (other than hard links).

Another option would be creating a small script which runs on storage.qa.suse.de which would utilize btrfs: It would make a new btrfs subvolume/snapshot to store the previous state and invoke rsync to get in sync with o3. We'd run this script via cron or a systemd timer similar to rsnapshot. This shouldn't be hard to do although I'd needed to figure the required btrfs commands first.

#25 - 2021-11-09 15:21 - mkittler

- Target version set to Ready

#26 - 2021-11-09 16:24 - mkittler

- Target version deleted (Ready)

For played around with btrfs commands for 2nd option. I suppose the following commands would do it.

First time setup:

```
mkdir -p /storage/o3-backup/snapshots
btrfs subvolume create /storage/o3-backup/fs-tree
```

The actual backup and [snapshotting](#):

```
mkdir -p /storage/o3-backup/fs-tree/var/lib/openqa
rsync -aHP root@o3:/var/lib/openqa/ /storage/o3-backup/fs-tree/var/lib/openqa
btrfs subvolume snapshot -r "/storage/o3-backup/fs-tree" "/storage/o3-backup/snapshots/$(date --iso-8601=date)"
```

Cleanup of old snapshots:

```
cd /storage/o3-backup/snapshots
for snapshot in *; do test "$(( $(date +%s') - $(date --date="$snapshot" +%s') ))" -ge "$snapshot_retention_in_seconds" && btrfs subvolume delete "$snapshot"; done
```

Maybe I'm reinventing the wheel here but I suppose it doesn't need to be more complicated than this. Note that we cannot use btrfs send/receive because the fs on o3 isn't using btrfs.

#27 - 2021-11-09 16:24 - mkittler

- Target version set to Ready

#28 - 2021-11-09 17:54 - okurz

mkittler wrote:

Maybe I'm reinventing the wheel here

I am sure you are :) Your approach sounds nice but I am sure there is something we overlook and sooner or later we want a true backup application. I suggest to stay with rsnapshot but if you want something more fancy maybe <https://github.com/borgbackup/borg> sounds better to you? Regarding btrfs snapshots I think they are great for snapshotting system drives (and taking backups from these snapshots) but not for storing the backups in snapshots.

#29 - 2021-11-10 17:58 - mkittler

- Status changed from Workable to In Progress

I've checked available backup tools and we've discussed them in the chat. For the sake of simplicity we decided to stick with rsnapshot (which we're already using on the backup.qa.suse.de). We couldn't benefit much from the features of borg and restic anyways and btrbk isn't an option as our source file system doesn't use btrfs.

Draft SR for configuration changes: https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/612

I've also already started a backup manually and it is still running right now. Files will be stored under /storage/rsnapshot (which is a btrfs subvolume and has compression enabled), e.g. /storage/rsnapshot/alpha.0/openqa.opensuse.org/root/var/lib/openqa/....

#30 - 2021-11-11 10:13 - mkittler

Unfortunately the backup failed with rsync error code 12. It had copied 181 GiB before. Maybe the connection was interrupted in the middle. I'll try again.

I'm also currently checking the size of /var/lib/openqa on o3 to have some value to compare the backup against (to check for completeness). This shows that simply walking through this huge filesystem tree takes a while. I suppose we'll have to consider this when thinking about the frequency we'd like to run rsnapshot.

#31 - 2021-11-11 15:25 - mkittler

The backup I've started this morning is still running (in a tmux session of user martchus). Let's see how long it is going to take.

By the way, we're dealing with roughly 6 TiB of data:

```
martchus@ariel:~> df -h /var/lib/openqa /var/lib/openqa/share
Dateisystem      Größe Benutzt Verf. Verw% Eingehängt auf
/dev/vdb1        5,0T   3,5T  1,6T   69% /var/lib/openqa
/dev/mapper/vg0-assets 4,0T   3,2T  869G   79% /var/lib/openqa/share
```

(Not everything on /dev/vdb1 is under /var/lib/openqa, though.)

#32 - 2021-11-12 14:41 - mkittler

The cleanup I started yesterday is still running. It is still filling up /storage, we're currently at:

```
martchus@storage:~> sudo btrfs filesystem df /storage
Data, RAID1: total=486.00GiB, used=485.83GiB
System, RAID1: total=8.00MiB, used=96.00KiB
Metadata, RAID1: total=21.00GiB, used=20.30GiB
GlobalReserve, single: total=512.00MiB, used=0.00B
```

#33 - 2021-11-15 11:03 - mkittler

The backup is still ongoing:

```
martchus@storage:~> sudo btrfs filesystem df /storage
Data, RAID1: total=1.27TiB, used=1.27TiB
System, RAID1: total=8.00MiB, used=208.00KiB
Metadata, RAID1: total=52.00GiB, used=50.73GiB
GlobalReserve, single: total=512.00MiB, used=0.00B
martchus@storage:~> sudo btrfs filesystem df /storage
Data, RAID1: total=1.28TiB, used=1.28TiB
System, RAID1: total=8.00MiB, used=208.00KiB
Metadata, RAID1: total=52.00GiB, used=51.26GiB
GlobalReserve, single: total=512.00MiB, used=0.00B
```

#34 - 2021-11-18 10:12 - mkittler

The backup is still ongoing:

```
martchus@storage:~> sudo btrfs filesystem df /storage
Data, RAID1: total=7.82TiB, used=7.82TiB
System, RAID1: total=8.00MiB, used=1.09MiB
Metadata, RAID1: total=101.00GiB, used=99.36GiB
GlobalReserve, single: total=512.00MiB, used=0.00B
```

I doubt that we can practically fulfill all the acceptance criteria. This is just getting too big and slow.

#35 - 2021-11-19 07:48 - okurz

Agreed. Also the alert "storage: partitions usage (%) alert" triggered. I paused the alert now. Please ensure to re-enable it when the situation was resolved.

So let's reconsider the solutions. We can also achieve the acceptance criteria with the help of SUSE IT and the NetApp storage solution we already have. I suggest to ask Englnfra if and how we would be able to recover a) a complete filesystem content after a potential loss of data and b) selective data in case of partial loss, e.g. if by user error we loose some assets or results

For the current backup approach I suggest you exclude all non-fixed assets from the backup. Also exclude /assets/tests as all this content should be in git or is considered not important.

#36 - 2021-11-19 08:04 - okurz

- Parent task set to #102710

#37 - 2021-11-19 08:07 - okurz

- Copied to action #102713: Implement proper backups for o3 - list of installed packages added

#38 - 2021-11-19 08:09 - okurz

- Description updated

I split out the "save list of packages" as proposed and also as the last tasks conducted did not cover this -> [#102713](#)

#39 - 2021-11-19 10:12 - mkittler

I can exclude /var/lib/openqa/share/tests which would save us 5.4 GiB and /var/lib/openqa/share/factory/*/(fixed) which would save use 6 TiB. (By the way /var/lib/openqa/share/factory/*/(fixed) is only 3.17 GiB.)

I suppose I'd remove it already in the current backup.

#40 - 2021-11-19 12:47 - mkittler

I'm now running the backup with excludes. Because of a mistake in the config file I had to abort and retry the backup. Unfortunately now rsnapshot didn't include the `--link-dest` parameter on the second run:

```
storage:/storage # rsnapshot alpha
echo 22770 > /var/run/rsnapshot.pid
mv /storage/rsnapshot/alpha.1/ /storage/rsnapshot/_delete.22770/
mv /storage/rsnapshot/alpha.0/ /storage/rsnapshot/alpha.1/
mkdir -m 0755 -p /storage/rsnapshot/alpha.0/openqa.opensuse.org/
/usr/bin/rsync -a --delete --numeric-ids --relative --delete-excluded \
  --exclude=/var/lib/openqa/share/tests \
  +rsync_long_args=--exclude=/var/lib/openqa/share/factory/* \
  +rsync_long_args=--include=/var/lib/openqa/share/factory/*/fixed/** \
  --rsh=/usr/bin/ssh \
  --link-dest=/storage/rsnapshot/alpha.1/openqa.opensuse.org/root \
  root@openqa.opensuse.org:/var/lib/openqa \
  /storage/rsnapshot/alpha.0/openqa.opensuse.org/root
```

```
-----
rsnapshot encountered an error! The program was invoked with these options:
/usr/bin/rsnapshot alpha
-----
```

```
ERROR: /usr/bin/rsync syntax or usage error. Does this version of rsync support --link-dest?
```

```
touch /storage/rsnapshot/alpha.0/
rm -f /var/run/rsnapshot.pid
/usr/bin/rm -rf /storage/rsnapshot/_delete.22770
^CERROR: Warning! /usr/bin/rm failed.
ERROR: Error! rm_rf("/storage/rsnapshot/_delete.22770")
```

```
storage:/storage # rsnapshot alpha
echo 22883 > /var/run/rsnapshot.pid
mv /storage/rsnapshot/alpha.1/ /storage/rsnapshot/_delete.22883/
mv /storage/rsnapshot/alpha.0/ /storage/rsnapshot/alpha.1/
mkdir -m 0755 -p /storage/rsnapshot/alpha.0/openqa.opensuse.org/
/usr/bin/rsync -a --delete --numeric-ids --relative --delete-excluded \
  --exclude=/var/lib/openqa/share/tests \
  --exclude=/var/lib/openqa/share/factory/* \
  --include=/var/lib/openqa/share/factory/*/fixed/** --rsh=/usr/bin/ssh \
  root@openqa.opensuse.org:/var/lib/openqa \
  /storage/rsnapshot/alpha.0/openqa.opensuse.org/root
```

#41 - 2021-11-22 12:52 - mkittler

The storage server was rebooted so I'm not sure whether the backup has been fully concluded. I'll give it another manual try. This time it included `--link-dest` correctly.

I've also been creating an Infra ticket to clear up the questions: <https://sd.suse.com/servicedesk/customer/portal/1/SD-67510>

#42 - 2021-11-22 14:53 - mkittler

Trying again after running into `ERROR: /usr/bin/rsync returned 12 while processing root@openqa.opensuse.org:/var/lib/openqa`. If it happens more often (this is now the 2nd time, see [#44078#note-30](#)) I could try create a small rsync wrapper to implement retries.

#43 - 2021-11-26 12:56 - mkittler

I enabled now retries via the rsync config. The backup is still running after splitting it into multiple backup commands. It is still at the images directory (which is likely the directory which takes the longest).