

Osmocom.org Servers - Bug #4083

git.osmocom.org: error: garbage at end of loose object

07/01/2019 08:34 AM - osmith

Status: New	Start date: 07/01/2019
Priority: Low	Due date:
Assignee: sysmocom	% Done: 90%
Category:	
Target version:	
Spec Reference:	
Description	
Right now, I can't clone libosmo-sccp from https://git.osmocom.org/libosmo-sccp . It fails with:	
<pre>\$ git -C /home/user/code/osmo-dev/src/osmo-ci/_temp/repos clone https://git.osmocom.org/libosmo-sccp Cloning into 'libosmo-sccp'... error: garbage at end of loose object 'a70b4599f63b50c4d654eb9e1b1a099a7c8cdc36' fatal: loose object a70b4599f63b50c4d654eb9e1b1a099a7c8cdc36 (stored in /home/user/code/osmo-dev/src/osmo-ci/_temp/repos/libosmo-sccp/.git/objects/a7/0b4599f63b50c4d654eb9e1b1a099a7c8cdc36) is corrupt</pre>	
I have noticed this multiple times with various repositories in the past week.	
A few minutes later, it is working again. So maybe this is related to some sync job?	
Related issues:	
Related to Osmocom.org Servers - Bug #4774: occasional 504 error when accessi...	In Progress 10/03/2020

History

#1 - 07/22/2019 06:10 AM - osmith

This is still happening, e.g. here:

https://jenkins.osmocom.org/jenkins/job/Osmocom_OBS_latest/664/console

#2 - 07/22/2019 07:30 AM - laforge

There is a "synch job" and that's the gerrit replication job that pushes/mirrors gerrit.osmocom.org to git.osmocom.org. We're using the built-in gerrit replication service for this, which in turn AFAIK just does a git+ssh push, which shouldn't introduce any races. After all, it happens quite frequently that somebody clones a public git repo while somebody else is pushing something.

However, pulling/cloning over https means there's a nginx reverse proxy between the server and the client, and there may be some problem introduced here.

It would be useful to have exact timestamps to check if there's anything in the log. So maybe it's worth to modify related scripts to print a timestamp before doing the clone?

The interesting bit from our config is that the http+https access to git.osmocom.org are actually redirected to cgit.osmocom.org

To be honest, I never really looked too deep into how git over http works, given that in the past I've only used it with the native git protocol and with ssh. For a long time AFAIR it was even discouraged to use http* with it.

#3 - 07/22/2019 11:25 AM - osmith

Thanks for looking into it.

It would be useful to have exact timestamps to check if there's anything in the log. So maybe it's worth to modify related scripts to print a timestamp before doing the clone?

<https://gerrit.osmocom.org/c/osmo-ci/+/14894>

#4 - 07/29/2019 07:05 AM - osmith

[laforge](#), here is one set of timestamps:

```
+ date '+%Y-%m-%d %H:%M:%S'
2019-07-29 07:01:59
+ git clone https://git.osmocom.org/libosmo-sccp
Cloning into 'libosmo-sccp'...
error: garbage at end of loose object '70751fce24f989d7334cba180e34ca65840ed0f9'
fatal: loose object 70751fce24f989d7334cba180e34ca65840ed0f9 (stored in /home/osmocom-build/jenkins/workspace/Osmocom_OBS_nightly/nightly-3g_QnWQVB/repo/libosmo-sccp/.git/objects/70/751fce24f989d7334cba180e34ca65840ed0f9) is corrupt
+ date '+%Y-%m-%d %H:%M:%S'
2019-07-29 07:02:10
+ exit 1
```

https://jenkins.osmocom.org/jenkins/job/Osmocom_OBS_nightly/173/console

#5 - 08/29/2019 06:31 AM - osmith

```
+ date '+%Y-%m-%d %H:%M:%S'
2019-08-29 01:26:58
+ git clone https://git.osmocom.org/osmo-gsm-manuals
Cloning into 'osmo-gsm-manuals'...
error: garbage at end of loose object 'f880d707aa636080def299450b13901904614449'
fatal: loose object f880d707aa636080def299450b13901904614449 (stored in /home/osmocom-build/jenkins/workspace/Osmocom_OBS_nightly/nightly-3g_6IdnDB/repo/osmo-gsm-manuals/.git/objects/f8/80d707aa636080def299450b13901904614449) is corrupt
+ date '+%Y-%m-%d %H:%M:%S'
2019-08-29 01:27:11
```

https://jenkins.osmocom.org/jenkins/job/Osmocom_OBS_nightly/204/console

#6 - 10/21/2019 08:03 AM - osmith

```
+ date '+%Y-%m-%d %H:%M:%S'
2019-10-21 01:26:55
+ git clone https://git.osmocom.org/osmo-gsm-manuals
Cloning into 'osmo-gsm-manuals'...
error: garbage at end of loose object '51c7d508314156e6677972975c8a12a9224ccbe4'
fatal: loose object 51c7d508314156e6677972975c8a12a9224ccbe4 (stored in /home/osmocom-build/jenkins/workspace/Osmocom_OBS_nightly/nightly-3g_7tHS0x/repo/osmo-gsm-manuals/.git/objects/51/c7d508314156e6677972975c8a12a9224ccbe4) is corrupt
+ date '+%Y-%m-%d %H:%M:%S'
2019-10-21 01:27:07
```

https://jenkins.osmocom.org/jenkins/job/Osmocom_OBS_nightly/259/console

#7 - 10/21/2019 10:28 AM - laforge

I cannot reproduce it here. From which build slave[s] is this happening? Maybe it relates to a specific git client version?

#8 - 10/21/2019 12:15 PM - osmith

laforge wrote:

I cannot reproduce it here. From which build slave[s] is this happening? Maybe it relates to a specific git client version?

All three errors linked above happened on build2-deb9build-ansible. I've also experienced this problem on my local machine (also running Debian 9, currently git 2.11.0 installed), but very rarely. Maybe it is related to the git version.

When I try to pull / clone again at a later point (maybe a few minutes, I did not watch the time), it works again.

It would be useful to have exact timestamps to check if there's anything in the log. So

maybe it's worth to modify related scripts to print a timestamp before doing the clone?

The log output above contains timestamps. Is there anything useful in the logs?

All in all, this problem appears very rarely so it does not seem worth spending much time on it.

#9 - 12/12/2019 10:13 AM - osmith

It seems that cloning from <https://gerrit.osmocom.org> is not affected. How about I change the OBS jobs (these are the ones where I notice this problem every now and then), so it clones from that URL instead?

Of course we could also use the git protocol, but considering that we are building source packages that OBS will use to build the binary packages for the official repository, HTTPS is the better choice imho.

```
user@gsmdevel /tmp/test % date
Thu Dec 12 11:07:08 CET 2019
user@gsmdevel /tmp/test % git clone https://git.osmocom.org/osmo-gsm-manuals
Cloning into 'osmo-gsm-manuals'...
error: garbage at end of loose object 'f69689dc20498ac729baea89f1a5ced9f51715ce'
fatal: loose object f69689dc20498ac729baea89f1a5ced9f51715ce (stored in /tmp/test/osmo-gsm-manuals/.git/object
s/f6/9689dc20498ac729baea89f1a5ced9f51715ce) is corrupt
user@gsmdevel /tmp/test % git clone https://gerrit.osmocom.org/osmo-gsm-manuals
Cloning into 'osmo-gsm-manuals'...
remote: Counting objects: 150, done
remote: Finding sources: 100% (148/148)
remote: Total 2894 (delta 74), reused 2860 (delta 74)
Receiving objects: 100% (2894/2894), 1.31 MiB | 0 bytes/s, done.
Resolving deltas: 100% (1708/1708), done.
user@gsmdevel /tmp/test % rm -rf osmo-gsm-manuals
user@gsmdevel /tmp/test % git clone https://git.osmocom.org/osmo-gsm-manuals
Cloning into 'osmo-gsm-manuals'...
error: garbage at end of loose object 'f69689dc20498ac729baea89f1a5ced9f51715ce'
fatal: loose object f69689dc20498ac729baea89f1a5ced9f51715ce (stored in /tmp/test/osmo-gsm-manuals/.git/object
s/f6/9689dc20498ac729baea89f1a5ced9f51715ce) is corrupt
user@gsmdevel /tmp/test % date
Thu Dec 12 11:08:31 CET 2019
```

EDIT: cloning this repository from git.osmocom.org worked again at Thu Dec 12 11:15:47 CET 2019.

#10 - 12/12/2019 10:14 AM - osmith

Or alternatively, set up cloning from ssh + git for jenkins?

#11 - 12/12/2019 10:50 AM - laforge

On Thu, Dec 12, 2019 at 10:13:59AM +0000, osmith [REDMINE] wrote:

It seems that cloning from <https://gerrit.osmocom.org> is not affected. How about I change the OBS jobs (these are the ones where I notice this problem every now and then), so it clones from that URL instead?

good idea, works for me.

#12 - 12/12/2019 01:06 PM - osmith

- Status changed from New to In Progress

- % Done changed from 0 to 90

<https://gerrit.osmocom.org/c/osmo-ci/+/16561>

#13 - 10/03/2020 02:07 PM - laforge

- Related to Bug #4774: occasional 504 error when accessing git.osmocom.org added

#14 - 10/03/2020 02:08 PM - laforge

- Status changed from In Progress to New

- Priority changed from Normal to Low