

RPM Support - Issue #9236

Pulp 3.14 - can't sync new repositories because of checksum validation failure

08/12/2021 04:28 PM - sbrock

Status:	CLOSED - NOTABUG	Start date:	
Priority:	Normal	Due date:	
Assignee:		Estimated time:	0:00 hour
Category:			
Sprint/Milestone:			
Severity:	2. Medium	Groomed:	No
Version:		Sprint Candidate:	No
Platform Release:		Tags:	
OS:		Sprint:	
Triaged:	No	Quarter:	

Description

Hi there! It seems I've discovered what may be a bug. I've proven this on two separate machines now. Both are fully up to date with Pulp 3.14 and are fresh installs. This seems to be related to issue [#9224](#) but I am not sure.

Syncing new repos fails with checksum validation failure on packages. Which packages it fails on seems random. Each time a new sync occurs, it fails on a different package. Syncing the repo over and over and over again eventually completes successfully, and going forward repo syncs fine. It does not seem to matter where the repository is, as I've tried different mirrors, http and https, and even my local cobbler install. It seems to often occur on the OS and AppStream repositories, and does not matter if its OS or Kickstart. I've got one on Alma's PowerTools repo as well. This also happens on CentOS, AlmaLinux and RockyLinux. And I've seen it on the EPEL repositories too.

Here's some of the error messages. I am also going to attach a couple of screenshots (with the org fudged out) Let me know what else you need from me, logs, configs, et al. THANKS!!

A file located at the url https://repo.almalinux.org/almalinux/8.4/PowerTools/x86_64/os/Packages/mingw64-headers-5.0.2-2.el8.noarch.rpm failed validation due to checksum.

A file located at the url https://repo.almalinux.org/almalinux/8.4/AppStream/x86_64/os/Packages/rust-std-static-1.52.1-1.module_el8.4.0+2520+0729bac9.x86_64.rpm failed validation due to checksum.

A file located at the url http://mirrors.rit.edu/centos/7.9.2009/os/x86_64/Packages/kdeartwork-wallpapers-4.10.5-4.el7.noarch.rpm failed validation due to checksum.

A file located at the url http://cobbler..org/cobbler/repo_mirror/CentOS_79-x86_64/Packages/thunderbird-68.10.0-1.el7.centos.x86_64.rpm failed validation due to checksum.

A file located at the url http://cobbler..org/cobbler/repo_mirror/EPEL7-x86_64/p/paraview-4.4.0-2.el7.x86_64.rpm failed validation due to checksum.

A file located at the url http://mirrors.rit.edu/rocky/8.4/BaseOS/x86_64/os/./././AppStream/x86_64/os/Packages/libkkc-data-0.2.7-12.el8.x86_64.rpm failed validation due to checksum.

A file located at the url http://mirrors.rit.edu/rocky/8.4/AppStream/x86_64/os/Packages/openblas-0.3.12-1.el8.i686.rpm failed validation due to checksum.

History

#1 - 08/17/2021 04:37 PM - dkliban@redhat.com

- Project changed from Pulp to RPM Support

#2 - 08/17/2021 04:55 PM - [dalley](#)

I think there is no relation to [#9224](#), the HTTP response headers returned by all of these servers appears to be correct

#3 - 08/18/2021 03:14 AM - dalley

So, this afternoon I tried syncing all of the repos listed above (apart from the cobbler ones) with pulpcore 3.14.4 and pulp_rpm 3.14.1 and was unable to reproduce the described errors with any of them, (un?)fortunately.

Syncing new repos fails with checksum validation failure on packages. Which packages it fails on seems random. Each time a new sync occurs, it fails on a different package. Syncing the repo over and over and over again eventually completes successfully. It does not seem to matter where the repository is, as I've tried different mirrors, http and https, and even my local cobbler install. It seems to often occur on the OS and AppStream repositories, and does not matter if its OS or Kickstart. Ive got one on Alma's PowerTools repo as well. This also happens on CentOS, AlmaLinux and RockyLinux. And I've seen it on the EPEL repositories too.

Given this, it's definitely not a problem with the actual repositories being synced, otherwise subsequent resyncs wouldn't help. And it's ?probably? not network issues, because that would get handled by TCP... At this point I would maybe start to suspect hardware issues (e.g. bad memory corrupting files as they are downloaded to /tmp), but...

I've proven this on two separate machines now.

The plot thickens. How much do these machines have in common? Are they different physical systems, different VMs on the same host, do they share a filesystem via NFS or some other means?

Could you provide the exact version numbers of the Pulp packages in use?

going forward repo syncs fine.

That part is expected, since subsequent syncs shouldn't re-download the files that are already downloaded. Retrying the download on checksum mismatch errors during the sync would help but that doesn't explain why you're seeing so many of them compared to any of our other users.

#4 - 08/18/2021 02:31 PM - sbrock

Thanks so much for looking at this!! It is an odd one, indeed. I was hoping not to be a unique case, but apparently that may be so :)

Here's a list of the pulp packages, grepping by pulp:

```
python3-pulp-container-2.7.1-1.el7.noarch
tfm-rubygem-pulp_rpm_client-3.13.3-1.el7.noarch
pulpcore-selinux-1.2.4-1.el7.x86_64
tfm-rubygem-pulp_ansible_client-0.8.0-1.el7.noarch
tfm-rubygem-pulp_container_client-2.7.0-1.el7.noarch
python3-pulpcore-3.14.3-1.el7.noarch
python3-pulp-file-1.8.1-1.el7.noarch
python3-pulp-rpm-3.14.0-1.el7.noarch
tfm-rubygem-pulp_file_client-1.8.1-1.el7.noarch
tfm-rubygem-pulp_deb_client-2.13.0-1.el7.noarch
pulp-client-1.0-1.noarch
tfm-rubygem-smart_proxy_pulp-3.0.0-1.fm2_5.el7.noarch
python3-pulp-ansible-0.9.0-1.el7.noarch
tfm-rubygem-pulpcore_client-3.14.1-1.el7.noarch
python3-pulp-certguard-1.4.0-1.el7.noarch
tfm-rubygem-pulp_certguard_client-1.4.0-1.el7.noarch
python3-pulp-deb-2.14.1-1.el7.noarch
```

So I verified that these two machines having problems are on separate VMWare hosts. Interestingly, the old katello instance is on the same host as one of the new instances. This host does not show this issue. It is the one we are retiring. It is pulp2 and we are unable to get it to upgrade to pulp3... plus its just really messy, easier to build anew.

As for storage, these are all using the same NFS share, sourced off an Isilon SAN. This is including the older Katello install, and our old Cobbler install. I don't think traffic saturation would be an issue here, since this is our old SAN, and most traffic has moved onto a Pure storage SAN. Maybe there's a bad disk? I could check that, but I doubt it would be an issue, as the Isilon would handle that. To rule that device out, I think I'll stop the services on the dev instance, attach a large vm disk, sync the data back and test again with a new product. I'll let you know how that goes.

#5 - 08/18/2021 06:53 PM - sbrock

Hi! So after attaching local storage and syncing the repositories of both Alma and CentOS, I discovered no errors. So something isn't correct on my Isilon side of things here. I suspect we can close this as Not A Bug! Thanks again for looking into it!

#6 - 08/18/2021 11:39 PM - dalley

- Status changed from NEW to CLOSED - NOTABUG

Thanks for your update! No problem at all, hardware issues are *the worst*

Files

Screenshot_20210812_102414.png	60.4 KB	08/12/2021	sbrock
Screenshot_20210812_102437.png	65.7 KB	08/12/2021	sbrock