

## Debian Support - Issue #6920

### Pulp 3 - pulp-deb : Issue synchronizing bullseye-security repo

06/08/2020 10:32 AM - swisscom

<b>Status:</b>	CLOSED - WORKSFORME	<b>Start date:</b>	
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>		<b>Estimated time:</b>	0:00 hour
<b>Category:</b>			
<b>Sprint/Milestone:</b>			
<b>Severity:</b>	2. Medium	<b>Groomed:</b>	No
<b>Version - Debian:</b>	master	<b>Sprint Candidate:</b>	No
<b>Platform Release:</b>		<b>Tags:</b>	Katello
<b>Target Release - Debian:</b>		<b>Sprint:</b>	
<b>OS:</b>	CentOS 7	<b>Quarter:</b>	
<b>Triaged:</b>	No		

#### Description

Dear support team,

I just discovered that the sync task fails when bullseye-security distribution is included in a remote, like the following :

```
{
  "architectures": "amd64",
  "ca_cert": null,
  "client_cert": null,
  "client_key": null,
  "components": null,
  "distributions": "bullseye-security",
  "download_concurrency": 20,
  "gpgkey": null,
  "name": "debian-security",
  "policy": "on_demand",
  "proxy_url": null,
  "pulp_created": "2020-06-08T07:33:26.302753Z",
  "pulp_href": "/pulp/api/v3/remotes/deb/apt/d394132f-413f-416b-b5aa-64e5e6883179/",
  "pulp_last_updated": "2020-06-08T07:56:41.103430Z",
  "sync_installer": false,
  "sync_sources": false,
  "sync_udebs": false,
  "tls_validation": true,
  "url": "http://security.debian.org/debian-security/"
},
```

The error is the following :

```
File \"/usr/local/lib/pulp/lib64/python3.6/site-packages/rq/worker.py", line 886, in perform_job
\n    rv = job.perform()\n File \"/usr/local/lib/pulp/lib64/python3.6/site-packages/rq/job.py"
, line 664, in perform\n    self._result = self._execute()\n File \"/
/usr/local/lib/pulp/lib64/python3.6/site-packages/rq/job.py", line 670, in _execute\n
    return self.func(*self.args, **self.kwargs)\n File \"/
/usr/local/lib/pulp/lib64/python3.6/site-packages/pulp_deb/app/tasks/synchronizing.py"
, line 106, in synchronize\n
    DebDeclarativeVersion(first_stage, repository, mirror=mirror).create()\n File \"/
/usr/local/lib/pulp/lib64/python3.6/site-packages/pulpcore/plugin/stages/declarative_version.py"
, line 149, in create\n    loop.run_until_complete(pipeline)\n File \"/
/usr/lib64/python3.6/asyncio/base_events.py", line 484, in run_until_complete\n
    return future.result()\n File \"/
/usr/local/lib/pulp/lib64/python3.6/site-packages/pulpcore/plugin/stages/api.py"
, line 228, in create_pipeline\n    await asyncio.gather(*futures)\n File \"/>
```

```
/usr/local/lib/pulp/lib64/python3.6/site-packages/pulpcore/plugin/stages/api.py\"
, line 43, in __call__\n    await self.run()\n File \"
/usr/local/lib/pulp/lib64/python3.6/site-packages/pulpcore/plugin/stages/content_stages.py\"
, line 120, in run\n    ContentArtifact.objects.bulk_get_or_create(content_artifact_bulk)\n File
\"/usr/local/lib/pulp/lib64/python3.6/site-packages/pulpcore/app/models/content.py\"
, line 55, in bulk_get_or_create\n    objs[i] = objs[i].__class__.objects.get(objs[i].q())\n
File \"/usr/local/lib/pulp/lib64/python3.6/site-packages/django/db/models/manager.py\"
, line 82, in manager_method\n    return getattr(self.get_queryset(), name)(*args, **kwargs)\n
File \"/usr/local/lib/pulp/lib64/python3.6/site-packages/django/db/models/query.py\"
, line 408, in get\n    self.model._meta.object_name\n
```

If I patch the remote and include some components like "main contrib", the sync task do not fail anymore but the resulting publication's (dists/bullseye-security) contain only the Release file, not the "main" and "contrib" directories.

So I guess Pulp has an issue to discover the components of bullseye-security remote

Any idea if the issue resides on Pulp side or on the remote itself ?

**Related issues:**

Related to Debian Support - Test #6051: Add tests for known cases where the d...	<b>CLOSED - COMPLETE</b>
Related to Debian Support - Issue #8028: Cannot sync Ubuntu repositories	<b>CLOSED - WORKSFORME</b>

**History**

**#1 - 06/17/2020 02:56 PM - quba42**

- Related to Test #6051: Add tests for known cases where the distribution is not the codename or suite added

**#2 - 06/17/2020 03:10 PM - quba42**

The remote looks fine to me.

It is also important to note, that Debian appears to have given up on their policy of publishing their security repos using an "unusual" structure starting with "bullseye".

This does NOT appear to be a codename != distribution issue!

I will need to investigate.

May I ask what version of pulp\_deb you are using?

**#3 - 06/19/2020 03:58 PM - swisscom**

Hi, I am using 2.3.0b1 and manually applied the fixes from following issues : 6907 6873 6876 6787

I have no clue if all these fixes have been published in a version installed with pip or not yet. If you'd like me to upgrade pulp\_deb or reinstall pulp server, let me know

**#4 - 07/27/2020 06:52 PM - quba42**

- Sprint/Milestone set to Katello

**#5 - 08/12/2020 02:00 PM - quba42**

As far as I can tell the main trouble with the bullseye-security repo is that it does not yet contain any packages.

(There is a possible second issue, in that this repository includes a weird circular symlink for backwards compatibility with the legacy dist folder structure).

The plugin is currently completely unable to handle valid repositories that do not contain any packages.

Syncing such repositories is a valid use case that needs to be fixed.

I have opened a draft PR for adding relevant test fixtures here: <https://github.com/pulp/pulp-fixtures/pull/187>

**#6 - 09/16/2020 02:10 PM - quba42**

Update: Syncing empty (no package) repositories works and has test coverage, but bullseye-security is still broken.

My only remaining theory is that the circular symlink is the root cause. However, it is still far from clear to me, why it should matter.

**#7 - 09/16/2020 02:11 PM - quba42**

- Tags Katello added

## #8 - 05/21/2021 05:58 PM - mbucher

So testing this with latest main-branch of pulp\_deb I got the following error on the first attempt. It seems to try to save the empty file artifact multiple times.

Interestingly, the second attempt succeeds.

```
pulp [cc8b664323d54993bb1e018edc73f432]: rq.worker:ERROR: Traceback (most recent call last):
  File "/usr/local/lib/pulp/lib64/python3.8/site-packages/django/db/backends/utils.py", line 84, in _execute
    return self.cursor.execute(sql, params)
psycopg2.errors.UniqueViolation: duplicate key value violates unique constraint "core_artifact_sha256_key"
DETAIL:  Key (sha256)=(e3b0c44298f1c149afb4c8996fb92427ae41e4649b934ca495991b7852b855) already exists.
The above exception was the direct cause of the following exception:
Traceback (most recent call last):
  File "/usr/local/lib/pulp/lib64/python3.8/site-packages/rq/worker.py", line 1013, in perform_job
    rv = job.perform()
  File "/usr/local/lib/pulp/lib64/python3.8/site-packages/rq/job.py", line 709, in perform
    self._result = self._execute()
  File "/usr/local/lib/pulp/lib64/python3.8/site-packages/rq/job.py", line 732, in _execute
    result = self.func(*self.args, **self.kwargs)
  File "/home/vagrant/devel/pulp_deb/pulp_deb/app/tasks/synchronizing.py", line 122, in synchronize
    DebDeclarativeVersion(first_stage, repository, mirror=mirror).create()
  File "/home/vagrant/devel/pulpcore/pulpcore/plugin/stages/declarative_version.py", line 149, in create
    loop.run_until_complete(pipeline)
  File "/usr/lib64/python3.8/asyncio/base_events.py", line 616, in run_until_complete
    return future.result()
  File "/home/vagrant/devel/pulpcore/pulpcore/plugin/stages/api.py", line 225, in create_pipeline
    await asyncio.gather(*futures)
  File "/home/vagrant/devel/pulpcore/pulpcore/plugin/stages/api.py", line 43, in __call__
    await self.run()
  File "/home/vagrant/devel/pulp_deb/pulp_deb/app/tasks/synchronizing.py", line 320, in run
    da.artifact.save()
  File "/home/vagrant/devel/pulpcore/pulpcore/app/models/content.py", line 126, in save
    super().save(*args, **kwargs)
  File "/usr/local/lib/pulp/lib64/python3.8/site-packages/django_lifecycle/mixins.py", line 134, in save
    save(*args, **kwargs)
  File "/usr/local/lib/pulp/lib64/python3.8/site-packages/django/db/models/base.py", line 743, in save
    self.save_base(using=using, force_insert=force_insert,
  File "/usr/local/lib/pulp/lib64/python3.8/site-packages/django/db/models/base.py", line 780, in save_base
    updated = self._save_table(
  File "/usr/local/lib/pulp/lib64/python3.8/site-packages/django/db/models/base.py", line 873, in _save_table
    result = self._do_insert(cls._base_manager, using, fields, update_pk, raw)
  File "/usr/local/lib/pulp/lib64/python3.8/site-packages/django/db/models/base.py", line 910, in _do_insert
    return manager._insert([self], fields=fields, return_id=update_pk,
  File "/usr/local/lib/pulp/lib64/python3.8/site-packages/django/db/models/manager.py", line 82, in manager_method
    return getattr(self.get_queryset(), name)(*args, **kwargs)
  File "/usr/local/lib/pulp/lib64/python3.8/site-packages/django/db/models/query.py", line 1186, in _insert
    return query.get_compiler(using=using).execute_sql(return_id)
  File "/usr/local/lib/pulp/lib64/python3.8/site-packages/django/db/models/sql/compiler.py", line 1377, in execute_sql
    cursor.execute(sql, params)
  File "/usr/local/lib/pulp/lib64/python3.8/site-packages/django/db/backends/utils.py", line 67, in execute
    return self._execute_with_wrappers(sql, params, many=False, executor=self._execute)
  File "/usr/local/lib/pulp/lib64/python3.8/site-packages/django/db/backends/utils.py", line 76, in _execute_with_wrappers
    return executor(sql, params, many, context)
  File "/usr/local/lib/pulp/lib64/python3.8/site-packages/django/db/backends/utils.py", line 84, in _execute
    return self.cursor.execute(sql, params)
  File "/usr/local/lib/pulp/lib64/python3.8/site-packages/django/db/utils.py", line 89, in __exit__
    raise dj_exc_value.with_traceback(traceback) from exc_value
  File "/usr/local/lib/pulp/lib64/python3.8/site-packages/django/db/backends/utils.py", line 84, in _execute
    return self.cursor.execute(sql, params)
django.db.utils.IntegrityError: duplicate key value violates unique constraint "core_artifact_sha256_key"
DETAIL:  Key (sha256)=(e3b0c44298f1c149afb4c8996fb92427ae41e4649b934ca495991b7852b855) already exists.
```

## #9 - 05/25/2021 01:18 PM - quba42

I was able to reproduce the same effect using the following remote:

```
REMOTE_HREF=$(http post "${BASE_ADDR}"/pulp/api/v3/remotes/deb/apt/ name="${ENTITIES_NAME}" distributions="buster/updates" url="http://ftp.de.debian.org/debian-security/" | jq -r '.pulp_href')
```

**#10 - 06/02/2021 03:39 PM - quba42**

- *Related to Issue #8028: Cannot sync Ubuntu repositories added*

**#11 - 11/04/2021 03:46 PM - quba42**

- *Status changed from NEW to CLOSED - WORKSFORME*

I can no longer reproduce this with latest pulp\_deb and pulpcore using either bullseye-security (now with packages) or bookworm-security (empty repo).

I am closing this issue "CLOSED - WORKSFORME". If someone finds a new reproducer, feel free to re-open or open a new issue!

**#12 - 11/04/2021 03:47 PM - quba42**

- *Sprint/Milestone deleted (Katello)*