

Debian Support - Issue #8756

Repeated syncs with mirror=false lead to duplicate metadata files within repo versions (causing the sync to fail)

05/14/2021 02:52 PM - knzivid

Status: NEW	Start date:
Priority: Normal	Due date:
Assignee:	Estimated time: 0:00 hour
Category:	
Sprint/Milestone: Katello	
Severity: 3. High	Groomed: No
Version - Debian:	Sprint Candidate: No
Platform Release:	Tags:
Target Release - Debian:	Sprint:
OS:	Quarter:
Triaged: No	

Description

I hit this one during a sync

```
pulp_1 | 12:47:10 Traceback (most recent call last):
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 124, in validate_version_paths
pulp_1 |       validate_file_paths(paths)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/app/files.py", line 134, in validate_file_paths
pulp_1 |       raise ValueError(_("Path is duplicated: {path}").format(path=path))
pulp_1 | ValueError: Path is duplicated: dists/xenial-updates/universe/i18n/Translation-en.gz
pulp_1 |
pulp_1 | During handling of the above exception, another exception occurred:
pulp_1 |
pulp_1 | Traceback (most recent call last):
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/rq/worker.py", line 1008, in perform_job
pulp_1 |     rv = job.perform()
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/rq/job.py", line 706, in perform
pulp_1 |     self._result = self._execute()
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/rq/job.py", line 729, in _execute
pulp_1 |     result = self.func(*self.args, **self.kwargs)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulp_deb/app/tasks/synchronizing.py", line 122, in synchronize
pulp_1 |     DebDeclarativeVersion(first_stage, repository, mirror=mirror).create()
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/stages/declarative_version.py", line 149, in create
pulp_1 |     loop.run_until_complete(pipeline)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/app/models/repository.py", line 912, in __exit__
pulp_1 |     repository.finalize_new_version(self)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulp_deb/app/models/repository.py", line 53, in finalize_new_version
pulp_1 |     validate_repo_version(new_version)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 139, in validate_repo_version
pulp_1 |     validate_version_paths(version)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 126, in validate_version_paths
pulp_1 |     raise ValueError(_("Cannot create repository version. {err}").format(err=e))
pulp_1 | ValueError: Cannot create repository version. Path is duplicated: dists/xenial-updates/universe/i18n/Translation-en.gz.
pulp_1 | Traceback (most recent call last):
```

```
pulp_1 | File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 124, in validate_version_paths
pulp_1 |     validate_file_paths(paths)
pulp_1 | File "/usr/local/lib/python3.6/site-packages/pulpcore/app/files.py", line 134, in validate_file_paths
pulp_1 |     raise ValueError(_("Path is duplicated: {path}").format(path=path))
pulp_1 | ValueError: Path is duplicated: dists/xenial-updates/universe/il8n/Translation-en.gz
pulp_1 |
pulp_1 | During handling of the above exception, another exception occurred:
pulp_1 |
pulp_1 | Traceback (most recent call last):
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/rq/worker.py", line 1008, in perform_job
pulp_1 |     rv = job.perform()
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/rq/job.py", line 706, in perform
pulp_1 |     self._result = self._execute()
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/rq/job.py", line 729, in _execute
pulp_1 |     result = self.func(*self.args, **self.kwargs)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulp_deb/app/tasks/synchronizing.py", line 122, in synchronize
pulp_1 |     DebDeclarativeVersion(first_stage, repository, mirror=mirror).create()
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/stages/declarative_version.py", line 149, in create
pulp_1 |     loop.run_until_complete(pipeline)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/app/models/repository.py", line 912, in __exit__
pulp_1 |     repository.finalize_new_version(self)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulp_deb/app/models/repository.py", line 53, in finalize_new_version
pulp_1 |     validate_repo_version(new_version)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 139, in validate_repo_version
pulp_1 |     validate_version_paths(version)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 126, in validate_version_paths
pulp_1 |     raise ValueError(_("Cannot create repository version. {err}").format(err=e))
pulp_1 | ValueError: Cannot create repository version. Path is duplicated: dists/xenial-updates/universe/il8n/Translation-en.gz.
```

Pulp versions

pulp-ansible	0.7.3
pulp-certguard	1.2.0
pulp-container	2.5.2
pulp-deb	2.11.1
pulp-file	1.7.0
pulp-maven	0.2.0
pulp-python	3.2.0
pulp-rpm	3.10.0
pulpcore	3.12.2

Related issues:

Is duplicate of Debian Support - Issue #8671: Sync with Main Debian Repo fail...

CLOSED - CURRENTRELEASE

History

#1 - 05/14/2021 03:10 PM - knzivid

Updated pulp-deb to 2.12 and got the same

```
pulp_1 | 13:07:15 Traceback (most recent call last):
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 124, in validate_version_paths
pulp_1 |     validate_file_paths(paths)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/app/files.py", line 134, in validate_file_paths
pulp_1 |     raise ValueError(_("Path is duplicated: {path}").format(path=path))
pulp_1 | ValueError: Path is duplicated: dists/xenial-updates/universe/il8n/Translation-en.gz
pulp_1 |
pulp_1 | During handling of the above exception, another exception occurred:
```

```

pulp_1 |
pulp_1 | Traceback (most recent call last):
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/rq/worker.py", line 1008, in perform_job
pulp_1 |     rv = job.perform()
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/rq/job.py", line 706, in perform
pulp_1 |     self._result = self._execute()
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/rq/job.py", line 729, in _execute
pulp_1 |     result = self.func(*self.args, **self.kwargs)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulp_deb/app/tasks/synchronizing.py", line 122, in syn
nchronize
pulp_1 |     DebDeclarativeVersion(first_stage, repository, mirror=mirror).create()
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/stages/declarative_version.py", line
149, in create
pulp_1 |     loop.run_until_complete(pipeline)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/app/models/repository.py", line 912, in __ex
it__
pulp_1 |     repository.finalize_new_version(self)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulp_deb/app/models/repository.py", line 57, in final
ize_new_version
pulp_1 |     validate_repo_version(new_version)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 139, in
validate_repo_version
pulp_1 |     validate_version_paths(version)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 126, in
validate_version_paths
pulp_1 |     raise ValueError(_("Cannot create repository version. {err}").format(err=e))
pulp_1 | ValueError: Cannot create repository version. Path is duplicated: dists/xenial-updates/universe/i18n
/Translation-en.gz.
pulp_1 | Traceback (most recent call last):
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 124, in
validate_version_paths
pulp_1 |     validate_file_paths(paths)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/app/files.py", line 134, in validate_file_pa
ths
pulp_1 |     raise ValueError(_("Path is duplicated: {path}").format(path=path))
pulp_1 | ValueError: Path is duplicated: dists/xenial-updates/universe/i18n/Translation-en.gz
pulp_1 |
pulp_1 | During handling of the above exception, another exception occurred:
pulp_1 |
pulp_1 | Traceback (most recent call last):
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/rq/worker.py", line 1008, in perform_job
pulp_1 |     rv = job.perform()
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/rq/job.py", line 706, in perform
pulp_1 |     self._result = self._execute()
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/rq/job.py", line 729, in _execute
pulp_1 |     result = self.func(*self.args, **self.kwargs)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulp_deb/app/tasks/synchronizing.py", line 122, in syn
nchronize
pulp_1 |     DebDeclarativeVersion(first_stage, repository, mirror=mirror).create()
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/stages/declarative_version.py", line
149, in create
pulp_1 |     loop.run_until_complete(pipeline)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/app/models/repository.py", line 912, in __ex
it__
pulp_1 |     repository.finalize_new_version(self)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulp_deb/app/models/repository.py", line 57, in final
ize_new_version
pulp_1 |     validate_repo_version(new_version)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 139, in
validate_repo_version
pulp_1 |     validate_version_paths(version)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 126, in
validate_version_paths
pulp_1 |     raise ValueError(_("Cannot create repository version. {err}").format(err=e))
pulp_1 | ValueError: Cannot create repository version. Path is duplicated: dists/xenial-updates/universe/i18n
/Translation-en.gz.

```

#2 - 05/14/2021 03:31 PM - knzivid

I deleted and recreated both the repo and the remote. Then successive syncs worked fine.

#3 - 05/25/2021 08:07 AM - knzivid

I also hit this for Release files

```

pulp_1 | 06:03:39 Traceback (most recent call last):
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 124, in
validate_version_paths
pulp_1 |     validate_file_paths(paths)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/app/files.py", line 134, in validate_file_pa
ths
pulp_1 |     raise ValueError(_("Path is duplicated: {path}").format(path=path))
pulp_1 | ValueError: Path is duplicated: dists/xenial-updates/Release
pulp_1 |
pulp_1 | During handling of the above exception, another exception occurred:
pulp_1 |
pulp_1 | Traceback (most recent call last):
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/rq/worker.py", line 1008, in perform_job
pulp_1 |     rv = job.perform()
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/rq/job.py", line 706, in perform
pulp_1 |     self._result = self._execute()
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/rq/job.py", line 729, in _execute
pulp_1 |     result = self.func(*self.args, **self.kwargs)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulp_deb/app/tasks/synchronizing.py", line 122, in sy
nchronize
pulp_1 |     DebDeclarativeVersion(first_stage, repository, mirror=mirror).create()
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/stages/declarative_version.py", line
149, in create
pulp_1 |     loop.run_until_complete(pipeline)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/app/models/repository.py", line 912, in __ex
it__
pulp_1 |     repository.finalize_new_version(self)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulp_deb/app/models/repository.py", line 57, in final
ize_new_version
pulp_1 |     validate_repo_version(new_version)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 139, in
validate_repo_version
pulp_1 |     validate_version_paths(version)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 126, in
validate_version_paths
pulp_1 |     raise ValueError(_("Cannot create repository version. {err}").format(err=e))
pulp_1 | ValueError: Cannot create repository version. Path is duplicated: dists/xenial-updates/Release.
pulp_1 | Traceback (most recent call last):
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 124, in
validate_version_paths
pulp_1 |     validate_file_paths(paths)
pulp_1 |   File "/usr/local/lib/python3.6/site-packages/pulpcore/app/files.py", line 134, in validate_file_pa
ths
pulp_1 |     raise ValueError(_("Path is duplicated: {path}").format(path=path))
pulp_1 | ValueError: Path is duplicated: dists/xenial-updates/Release

```

#4 - 05/25/2021 12:14 PM - knzivid

All of the following files are duplicated

```

'dists/xenial-updates/InRelease'
'dists/xenial-updates/Release'
'dists/xenial-updates/Release.gpg'
'dists/xenial-updates/main/binary-amd64/Packages'
'dists/xenial-updates/main/binary-amd64/Packages.gz'
'dists/xenial-updates/main/binary-amd64/Packages.xz'
'dists/xenial-updates/main/binary-amd64/Release'
'dists/xenial-updates/universe/binary-amd64/Packages'
'dists/xenial-updates/universe/binary-amd64/Packages.gz'
'dists/xenial-updates/universe/binary-amd64/Packages.xz'
'dists/xenial-updates/universe/binary-amd64/Release'

```

I listed these by finding non-unique items in `ContentArtifact.objects.filter(content__pk__in=version.content).values_list("relative_path", flat=True)` in `validate_version_paths` in `repo_version_utils.py`. How did the repository end up with duplicate contents?

#5 - 05/25/2021 01:24 PM - quba42

The duplicate path for translation files problem will be fixed with 2.12.1 (to be released soon).

I have not yet encountered the duplicate path problem with respect to the Release file. This should be retested with version 2.12.1. For now I am marking this issue as a duplicate of <https://pulp.plan.io/issues/8671>. But if the error persists with version 2.12.1 then this is a separate issue. It may be worth cleaning up orphans before retesting with 2.12.1 as well.

#6 - 05/25/2021 01:24 PM - quba42

- Is duplicate of Issue #8671: Sync with Main Debian Repo fails on a translation file added

#7 - 05/25/2021 01:32 PM - knzivid

When I run sync with RepositorySyncURL(remote=remote_pulp_href, mirror=True), it works!

By default, mirror=False and that fails. When I am syncing from a remote, does it matter if mirror=True? Packages and Release files will be overwritten anyway.

If True, synchronization will remove all content that is not present in the remote repository. If False, sync will be additive only.

#8 - 05/25/2021 02:50 PM - quba42

That sounds like pulp_deb does not implement any reasonable mirror=False logic for metadata files. (Retaining multiple sets of metadata files in a single repository version does not make any sense.)

In that case you have stumbled on another bug, as well as the workaround (use mirror=True).

#9 - 05/26/2021 10:59 AM - quba42

Note 2.12.2 has been released which should take care of at least one of the errors mentioned in this ticket.

See: <https://listman.redhat.com/archives/pulp-list/2021-May/msg00053.html>

#10 - 06/02/2021 11:16 AM - simon_t

```
{
  "component": "core",
  "version": "3.13.0"
},
{
  "component": "file",
  "version": "1.7.0"
},
{
  "component": "deb",
  "version": "2.12.0"
}
```

It looks like we are hitting the duplicate Release file issue too. Syncs (mirror=false) are failing at the same point across multiple attempts (5) with the following error:

```
{
  "child_tasks": [],
  "created_resources": [],
  "error": {
    "description": "Cannot create repository version. Path is duplicated: dists/bullseye/Release.",
    "traceback": " File \"/usr/local/lib/pulp/lib/python3.6/site-packages/rq/worker.py", line 1013, in p
perform_job\n rv = job.perform()\n File \"/usr/local/lib/pulp/lib/python3.6/site-packages/rq/job.py", line
709, in perform\n self._result = self._execute()\n File \"/usr/local/lib/pulp/lib/python3.6/site-packages
/rq/job.py", line 732, in _execute\n result = self.func(*self.args, **self.kwargs)\n File \"/usr/local/li
b/pulp/lib/python3.6/site-packages/pulp_deb/app/tasks/synchronizing.py", line 122, in synchronize\n DebDec
larativeVersion(first_stage, repository, mirror=mirror).create()\n File \"/usr/local/lib/pulp/lib/python3.6/s
ite-packages/pulpcore/plugin/stages/declarative_version.py", line 149, in create\n loop.run_until_complete
(pipeline)\n File \"/usr/local/lib/pulp/lib/python3.6/site-packages/pulpcore/app/models/repository.py", line
956, in __exit__\n repository.finalize_new_version(self)\n File \"/usr/local/lib/pulp/lib/python3.6/site-
packages/pulp_deb/app/models/repository.py", line 57, in finalize_new_version\n validate_repo_version(new_
version)\n File \"/usr/local/lib/pulp/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", li
ne 139, in validate_repo_version\n validate_version_paths(version)\n File \"/usr/local/lib/pulp/lib/python
3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 126, in validate_version_paths\n raise Valu
eError(_("\nCannot create repository version. {err}.\n").format(err=e))\n"
  },
  "finished_at": "2021-06-02T07:56:26.918462Z",
  "logging_cid": "b66faefac6b241e095a918176071e3d6",
  "name": "pulp_deb.app.tasks.synchronizing.synchronize",
  "parent_task": null,
  "progress_reports": [
    {
      "code": "sync.downloading.artifacts",
      "done": 3,
      "message": "Downloading Artifacts",
      "state": "completed",
      "suffix": null,
    }
  ]
}
```

```

    "total": null
  },
  {
    "code": "update.release_file",
    "done": 1,
    "message": "Update ReleaseFile units",
    "state": "completed",
    "suffix": null,
    "total": null
  },
  {
    "code": "update.packageindex",
    "done": 3,
    "message": "Update PackageIndex units",
    "state": "completed",
    "suffix": null,
    "total": null
  },
  {
    "code": "associating.content",
    "done": 1992,
    "message": "Associating Content",
    "state": "completed",
    "suffix": null,
    "total": null
  }
],
"pulp_created": "2021-06-02T07:11:37.144496Z",
"pulp_href": "/pulp/api/v3/tasks/7d376c27-17b4-48bb-8186-bc1646e10378/",
"reserved_resources_record": [
  "/pulp/api/v3/remotes/deb/apt/3f95927d-b39c-4bda-a24e-f850d73f2095/",
  "/pulp/api/v3/repositories/deb/apt/dbd1bbf7-a5ef-46b6-922a-0ddf62b92ebc/"
],
"started_at": "2021-06-02T07:45:50.499516Z",
"state": "failed",
"task_group": null,
"worker": "/pulp/api/v3/workers/788be55c-4399-4578-98b4-0cf268a8dd18/"
}

```

Re-trying the sync with "mirror=true" results in it completing at the same point as the previous failures.

```

{
  "child_tasks": [],
  "created_resources": [
    "/pulp/api/v3/repositories/deb/apt/dbd1bbf7-a5ef-46b6-922a-0ddf62b92ebc/versions/6/"
  ],
  "error": null,
  "finished_at": "2021-06-02T08:12:27.377670Z",
  "logging_cid": "5056776354914a0787a8465b271acea8",
  "name": "pulp_deb.app.tasks.synchronizing.synchronize",
  "parent_task": null,
  "progress_reports": [
    {
      "code": "sync.downloading.artifacts",
      "done": 3,
      "message": "Downloading Artifacts",
      "state": "completed",
      "suffix": null,
      "total": null
    },
    {
      "code": "update.release_file",
      "done": 1,
      "message": "Update ReleaseFile units",
      "state": "completed",
      "suffix": null,
      "total": null
    },
    {
      "code": "update.packageindex",
      "done": 3,
      "message": "Update PackageIndex units",
      "state": "completed",
      "suffix": null,
      "total": null
    }
  ]
}

```

```

    },
    {
      "code": "unassociating.content",
      "done": 2018,
      "message": "Un-Associating Content",
      "state": "completed",
      "suffix": null,
      "total": null
    },
    {
      "code": "associating.content",
      "done": 1992,
      "message": "Associating Content",
      "state": "completed",
      "suffix": null,
      "total": null
    }
  ],
  "pulp_created": "2021-06-02T08:01:38.201191Z",
  "pulp_href": "/pulp/api/v3/tasks/aaec8b6a-b442-483f-a359-a857cc4dba26/",
  "reserved_resources_record": [
    "/pulp/api/v3/remotes/deb/apt/3f95927d-b39c-4bda-a24e-f850d73f2095/",
    "/pulp/api/v3/repositories/deb/apt/dbdlbbf7-a5ef-46b6-922a-0ddf62b92ebc/"
  ],
  "started_at": "2021-06-02T08:01:38.353696Z",
  "state": "completed",
  "task_group": null,
  "worker": "/pulp/api/v3/workers/760fb8d8-8cbd-4f2d-852f-2253cfe30773/"
}

```

This leaves the repository in an incomplete state and attempts to access it via a distribution result in this error on the client:

```

admin@debian-demo:~$ sudo apt-get update
Ign:1 https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye InRelease
Err:2 https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye Release
 404 Not Found [IP: 192.168.1.1 443]
Reading package lists... Done
E: The repository 'https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye Release' does not have a Release file.
N: Updating from such a repository can't be done securely, and is therefore disabled by default.
N: See apt-secure(8) manpage for repository creation and user configuration details.

Jun 02 08:20:44 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:20:44 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/InRelease HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
Jun 02 08:20:45 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:20:44 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/Release HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"

```

And this on the Pulp host:

```

Jun 02 08:20:44 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:20:44 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/InRelease HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
Jun 02 08:20:45 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:20:44 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/Release HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"

```

Setting "[trusted=true]" for the repo in the clients sources.list doesn't help:

```

admin@debian-demo:~$ sudo apt-get update
Ign:1 https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye InRelease
Ign:2 https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye Release
Ign:3 https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye/main amd64 Packages
Ign:4 https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye/main all Packages
Ign:3 https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye/main amd64 Packages
Ign:4 https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye/main all Packages
Ign:3 https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye/main amd64 Packages
Ign:4 https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye/main all Packages
Ign:3 https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye/main amd64 Packages
Ign:4 https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye/main all Packages
Ign:3 https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye/main amd64 Packages
Ign:4 https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye/main all Packages
Ign:3 https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye/main amd64 Packages
Err:3 https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye/main amd64 Packages
 404 Not Found [IP: 192.168.1.1 443]
Ign:4 https://pulphost.domain.internal/pulp/content/debian_bullseye bullseye/main all Packages
Reading package lists... Done

```

E: Failed to fetch https://pulphost.domain.internal/pulp/content/debian_bullseye/dists/bullseye/main/binary-amd64/Packages 404 Not Found [IP: 192.168.1.1 443]
E: Some index files failed to download. They have been ignored, or old ones used instead.

```
Jun 02 08:25:19 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:25:19 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/InRelease HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
Jun 02 08:25:19 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:25:19 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/Release HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
Jun 02 08:25:19 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:25:19 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/main/binary-amd64/Packages.xz HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
Jun 02 08:25:19 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:25:19 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/main/binary-all/Packages.xz HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
Jun 02 08:25:19 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:25:19 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/main/binary-amd64/Packages.bz2 HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
Jun 02 08:25:19 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:25:19 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/main/binary-all/Packages.bz2 HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
Jun 02 08:25:19 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:25:19 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/main/binary-amd64/Packages.lzma HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
Jun 02 08:25:19 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:25:19 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/main/binary-all/Packages.lzma HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
Jun 02 08:25:19 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:25:19 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/main/binary-amd64/Packages.gz HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
Jun 02 08:25:19 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:25:19 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/main/binary-all/Packages.gz HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
Jun 02 08:25:19 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:25:19 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/main/binary-amd64/Packages.lz4 HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
Jun 02 08:25:19 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:25:19 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/main/binary-all/Packages.lz4 HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
Jun 02 08:25:19 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:25:19 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/main/binary-amd64/Packages.zst HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
Jun 02 08:25:19 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:25:19 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/main/binary-all/Packages.zst HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
Jun 02 08:25:19 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:25:19 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/main/binary-amd64/Packages HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
Jun 02 08:25:19 pulphost.domain.internal gunicorn[22814]: [02/Jun/2021:08:25:19 +0000] "GET /pulp/content/debian_bullseye/dists/bullseye/main/binary-all/Packages HTTP/1.0" 404 178 "-" "Debian APT-HTTP/1.3 (2.2.3)"
```

#11 - 06/02/2021 01:56 PM - quba42

- Subject changed from *ValueError: Path is duplicated: dists/xenial-updates/universe/i18n/Translation-en.gz* to *Syncs with mirror=false lead to duplicate metadata files within repo versions (causing the sync to fail)*

- *Sprint/Milestone set to Katello*

#12 - 06/02/2021 02:03 PM - quba42

- Subject changed from *Syncs with mirror=false lead to duplicate metadata files within repo versions (causing the sync to fail)* to *Repeated syncs with mirror=false lead to duplicate metadata files within repo versions (causing the sync to fail)*

#13 - 06/23/2021 10:30 AM - quba42

I have now analyzed what needs doing: `mirror=false` works by omitting the `ContentUnassociation` stage when creating a new `DeclarativeVersion`. Since it never makes sense to retain altered metadata files, we need to add a stage to the sync which mimics the `ContentUnassociation` stage but just for the metadata content types.

(I think that should work).

#14 - 08/19/2021 06:55 AM - simon_t

This issue also affects the package indicies. An example from our last sync:

Repo: `"/pulp/api/v3/repositories/deb/apt/4ab4b277-886b-4f32-85f4-eed4ccd0bc6e/"`

```
"error": {
```

```
  "description": "Cannot create repository version. Path is duplicated: dists/bullseye/contrib/binary-amd64/Packages.gz.",
```

```
  "traceback": " File \"/usr/local/lib/pulp/lib/python3.6/site-packages/pulpcore/tasking/pulpcore_worker.py", line 272, in _perform_task\n    result = func(*args, **kwargs)\n File \"/usr/local/lib/pulp/lib/python3.6/site-packages/pulpcore/tasking/synchronizing.py", line 124, in synchronize\n    DebDeclarativeVersion(\n    first_stage, repository, mirror=mirror).create()\n File \"/usr/local/lib/pulp/lib/python3.6/site-packages/pulpcore/plugin/stages/declarative_version.py", line 151, in create\n    loop.run_until_complete(pipeline)\n File \"/usr/local/lib/pulp/lib/python3.6/site-packages/pulpcore/app/models/repository.py", line 963, in __exit_\n    repository.finalize_new_version(self)\n File \"/usr/local/lib/pulp/lib/python3.6/site-packages/pulp_de
```



```
b/app/models/repository.py", line 57, in finalize_new_version\n    validate_repo_version(new_version)\n File  
  \"/usr/local/lib/pulp/lib/python3.6/site-packages/pulpcore/plugin/repo_version_utils.py", line 139, in valid  
ate_repo_version\n    validate_version_paths(version)\n File \"/usr/local/lib/pulp/lib/python3.6/site-package  
s/pulpcore/plugin/repo_version_utils.py", line 126, in validate_version_paths\n    raise ValueError(_("\nCanno  
t create repository version. {err}.\n").format(err=e))\n"
```

```
}
```

Workaround

```
~# http get $B/pulp/api/v3/content/deb/package_indices/?component="contrib"&architecture=amd64
```

```
{  
  "count": 3,  
  "next": null,  
  "previous": null,  
  "results": [  
    {  
      "architecture": "amd64",  
      "artifacts": {  
        "dists/bullseye/contrib/binary-amd64/Packages": "/pulp/api/v3/artifacts/772a6f11-e595-4ed8-b7c  
0-6be8d51d50a3/",  
        "dists/bullseye/contrib/binary-amd64/Packages.gz": "/pulp/api/v3/artifacts/fd4f3e0b-6f24-42de-  
a87d-b4b5c4da2aca/",  
        "dists/bullseye/contrib/binary-amd64/Packages.xz": "/pulp/api/v3/artifacts/3f0d0ffc-3faf-4163-  
9009-98f558fd2afa/",  
        "dists/bullseye/contrib/binary-amd64/Release": "/pulp/api/v3/artifacts/b16ebb49-97d7-43ee-8c06  
-0e4b83008ff9/"  
      },  
      "component": "contrib",  
      "pulp_created": "2021-08-18T01:30:49.608585Z",  
      "pulp_href": "/pulp/api/v3/content/deb/package_indices/1ea03ab7-dc35-46bc-9045-f8155b2473e2/",  
      "relative_path": "dists/bullseye/contrib/binary-amd64/Packages",  
      "release": "/pulp/api/v3/content/deb/release_files/3c08bddc-3303-4217-8f3f-cee36ccb28f9/"  
    },  
    {  
      "architecture": "all",  
      "artifacts": {  
        "dists/bullseye/contrib/binary-all/Packages": "/pulp/api/v3/artifacts/5ceed943-0210-473f-b767-  
43e4e98da99a/",  
        "dists/bullseye/contrib/binary-all/Packages.gz": "/pulp/api/v3/artifacts/ba0bc84b-858a-4a8c-98  
8f-52810b7a213f/",  
        "dists/bullseye/contrib/binary-all/Packages.xz": "/pulp/api/v3/artifacts/f8ac40ed-fd4d-459c-93  
86-e971df50e0e9/",  
        "dists/bullseye/contrib/binary-all/Release": "/pulp/api/v3/artifacts/503e29c1-4985-4f63-b974-1  
37b542c9ecb/"  
      },  
      "component": "contrib",
```

```

    "pulp_created": "2021-08-18T01:30:49.603219Z",
    "pulp_href": "/pulp/api/v3/content/deb/package_indices/ac504d9a-096c-4712-8c49-c80fc98a21bd/",
    "relative_path": "dists/bullseye/contrib/binary-all/Packages",
    "release": "/pulp/api/v3/content/deb/release_files/3c08bddc-3303-4217-8f3f-cee36ccb28f9/"
  },
  {
    "architecture": "amd64",
    "artifacts": {
      "dists/bullseye/contrib/binary-amd64/Packages": "/pulp/api/v3/artifacts/f6f34236-cc7e-41cc-b1c0-b6ce8ff8f47d/",
      "dists/bullseye/contrib/binary-amd64/Packages.gz": "/pulp/api/v3/artifacts/9efeb3fd-22ff-4131-b331-48c3870d5e7d/",
      "dists/bullseye/contrib/binary-amd64/Packages.xz": "/pulp/api/v3/artifacts/ee280975-75ba-415a-a4d2-0e23ce3c41c8/",
      "dists/bullseye/contrib/binary-amd64/Release": "/pulp/api/v3/artifacts/1bf430e9-bb9c-4574-aa6a-d5b07d47ba98/"
    },
    "component": "contrib",
    "pulp_created": "2021-06-25T05:44:18.270613Z",
    "pulp_href": "/pulp/api/v3/content/deb/package_indices/0433288a-f009-45c2-89dd-ee479bb38995/",
    "relative_path": "dists/bullseye/contrib/binary-amd64/Packages",
    "release": "/pulp/api/v3/content/deb/release_files/772dbd56-c977-4a59-a58e-9895bd7641a9/"
  }
]
}

```

```

~# http post $B/pulp/api/v3/repositories/deb/apt/4ab4b277-886b-4f32-85f4-eed4ccd0bc6e/modify/ remove_content_units='["/pulp/api/v3/content/deb/package_indices/0433288a-f009-45c2-89dd-ee479bb38995/"]'

```

The sync will now progress. If the same issue is hit on another file the process can be repeated to remove the relevant duplicated package indices

#15 - 08/19/2021 08:46 AM - quba42

Thanks for the additional information and workaround. Since indices are one of the metadata types, it is not entirely surprising that they are also affected.