

RPM Support - Test #4727

rpm_rsync_distributor is broken for yum_repo_metadata_file units

04/23/2019 09:41 PM - bherring

Status: CLOSED - COMPLETE	
Priority: High	
Assignee: kersom	
Category:	
Sprint/Milestone:	
Version:	Sprint:
Platform Release:	Quarter:
Tags: Pulp 2	
Description	
rpm_rsync_distributor doesn't work at all for units of type yum_repo_metadata_file. It will publish broken symlinks to the remote host.	
This is due to the following interaction between yum_distributor and rpm_rsync_distributor:	
First, in yum_distributor, the metadata file is symlinked from repo publish dir to storage path:	
<pre>class PublishMetadataStep(platform_steps.UnitModelPluginStep): ... def process_main(self, item=None): """ Copy the metadata file into place and add it to the repomd file. :param item: The unit to process :type item: pulp.server.db.model.ContentUnit """ unit = item # Copy the file to the location on disk where the published repo is built publish_location_relative_path = os.path.join(self.get_working_dir(), REPO_DATA_DIR_NAME) metadata_file_name = os.path.basename(unit._storage_path) link_path = os.path.join(publish_location_relative_path, metadata_file_name) plugin_misc.create_symlink(unit._storage_path, link_path) # Add the proper relative reference to the metadata file to repomd self.parent.repomd_file_context.\ add_metadata_file_metadata(unit.data_type, link_path)</pre>	
That results in a yum repo with a symlink as in following example:	
<pre>\$ ls -l /var/lib/pulp/published/yum/master/yum_distributor/rhel-ha-for-rhel-7-for-system-z-rpms__7 Server__s390x/1552949091.82/repodata/ total 200 -rw-r--r--. 1 apache apache 2087 Mar 18 22:44 29e7003fe4c25a3291e80c16006ca0edc315c8b89f930f8be9 f435bb10c1000b-updateinfo.xml.gz -rw-r--r--. 1 apache apache 160494 Mar 18 22:44 792d01f805670173c58c8c55b1b960576de6b72215bd61d845 b077af254d7e14-other.xml.gz -rw-r--r--. 1 apache apache 16643 Mar 18 22:44 79f6997620d82b1085fb9dc22499b81d1b4ccb8a6113372039 273d0dbab5be9f-filelists.xml.gz lrwxrwxrwx. 1 apache apache 153 Mar 18 22:44 85771b84-06dd-4e4d-9e88-1b1d5616b3cc -> /var/lib/p ulp/content/units/yum_repo_metadata_file/6b/74066bca70043c0cec6f1c43bc973a4bb84aa9ff6579757b8b0a9b e1bcf81b/85771b84-06dd-4e4d-9e88-1b1d5616b3cc -rw-r--r--. 1 apache apache 124 Mar 18 22:44 a27718cc28ec6d71432e0ef3e6da544b7f9d93f6bb7d0a55aa cd592d03144b70-comps.xml</pre>	

```
-rw-r--r--. 1 apache apache 3575 Mar 18 22:44 a57cc25570ba03b0c7fc82ad4502fad14fb20b404f065801563bc58ea8340eda-primary.xml.gz
-rw-r--r--. 1 apache apache 2405 Mar 18 22:44 repomd.xml
```

Then, rpm_rsync_distributor performs rsync of repodata directory **with links=True**, meaning rsync --links, meaning "copy symlinks as symlinks":

```
self.add_child(RSyncPublishStep("Rsync step (repodata)",
                                repodata_file_list,
                                "%s/" % os.path.join(master_dir, 'repodata'),
                                "%s/" % os.path.join(remote_repo_path, 'repodata'),
                                exclude=[".*", "repodata.old"],
                                config=config, links=True,
                                delete=self.config.get("delete")))
```

This indeed rsyncs the symlink onto the remote host exactly as it is:

```
$ ls -l /mnt/redhat/cdn/content/dist/rhel/system-z/7/7Server/s390x/highavailability/os/repodata |
grep 85771b84-06dd-4e4d-9e88-1b1d5616b3cc
lrwxrwxrwx. 1 vagrant vagrant 153 Mar 18 22:44 85771b84-06dd-4e4d-9e88-1b1d5616b3cc -> /var/lib
/pulp/content/units/yum_repo_metadata_file/6b/74066bca70043c0cec6f1c43bc973a4bb84aa9ff6579757b8b0a
9be1bcf81b/85771b84-06dd-4e4d-9e88-1b1d5616b3cc
```

That's a broken symlink to /var/lib/pulp, which is not expected to exist on the remote. Hence the file is not able to be accessed and the published repo is incomplete.

Steps to reproduce

- Create a repo with yum_distributor, rpm_rsync_distributor
- Add any yum_repo_metadata_file to the repo
- Publish the repo with yum_distributor and then rpm_rsync_distributor

Expected result

All yum_repo_metadata_file in the repo are accessible in the repo published to remote host via rpm_rsync_distributor

Actual result

All yum_repo_metadata_file in published repo are broken symlinks.

Related issues:

Copied from RPM Support - Issue #4550: rpm_rsync_distributor is broken for yu...

CLOSED - CURRENTRELEASE

History

#1 - 04/23/2019 09:41 PM - bherring

- Copied from Issue #4550: rpm_rsync_distributor is broken for yum_repo_metadata_file units added

#2 - 05/07/2019 02:44 PM - bherring

- Priority changed from Normal to High

#3 - 05/17/2019 05:58 PM - kersom

- Status changed from NEW to ASSIGNED

- Assignee set to kersom

#4 - 05/31/2019 10:08 PM - kersom

- Status changed from ASSIGNED to CLOSED - COMPLETE