

Pulp - Issue #1008

Package Missing from repo metadata

05/28/2015 01:57 AM - Ben.Stanley

Status:	CLOSED - WONTFIX	Start date:	
Priority:	High	Due date:	
Assignee:		Estimated time:	0:00 hour
Category:			
Sprint/Milestone:			
Severity:	3. High	Groomed:	No
Version:	2.6.1	Sprint Candidate:	No
Platform Release:		Tags:	Pulp 2
OS:	RHEL 7	Sprint:	
Triaged:	Yes	Quarter:	

Description

I have a problem with syncing

http://elrepo.org/linux/elrepo/el7/x86_64

It seems that the package
nvidia-x11-drv-304xx-304.125-1.el7.elrepo.x86_64.rpm
is not included in my yum metadata. However, it is published in my local
repo.

```
[root stig ]# yum list nvidia-x11-drv-304xx
Loaded plugins: langpacks, product-id, subscription-manager
Available Packages
nvidia-x11-drv-304xx.x86_64
304.123-2.el7.elrepo elrepo
```

Only the older version 304.123-2 of the package is found, despite the newer 304.125-1 being published in the repo.

When I change my /etc/yum.repos.d/elrepo.repo to point to upstream instead
of my local pulp server, yum finds the correct nvidia-x11-drv-304xx (after
yum clean metadata).

I have recently run a sync, but it has not fixed the problem. There is
nothing interesting in the log (no exceptions or errors).

I am running pulp 2.6.1 (package versions below).

I have found that I can fix this kind of problem by deleting the repo and
re-creating it, but I need to get to the bottom of this so that it stops
happening. I have complained about something like this previously in
pulp-2.5.1 at [0] . I created my elrepo in pulp-2.6.1, so it is not
possible to blame an older version of pulp this time.

I need to get to the bottom of this, and hopefully get this resolved
properly. Otherwise, I look like a goose when my local copy doesn't work.

I need to know that after I have completed a sync, that everything really
is OK, before I sync to my disconnected system. Performing a complete
internet sync takes a long time, and a disconnected sync is going to be
even worse. I want to prove that everything is right before I do the
disconnected sync.

Questions:

1) Is there a tool that I can use to compare the metadata of two repos and
determine equivalence (at least to the point of validating the sync)? Such
a tool would allow me to proactively detect such problems, and attempt to

correct them. It would also provide more test cases for you...

2) Is there some way of investigating why this particular file isn't appearing in the metadata?

```
===== rpm -qa | grep pulp =====  
python-pulp-common-2.6.1-1.el7.noarch  
pulp-rpm-admin-extensions-2.6.1-1.el7.noarch  
pulp-server-2.6.1-1.el7.noarch  
python-pulp-client-lib-2.6.1-1.el7.noarch  
python-pulp-puppet-common-2.6.1-1.el7.noarch  
pulp-puppet-plugins-2.6.1-1.el7.noarch  
pulp-admin-client-2.6.1-1.el7.noarch  
python-isodate-0.5.0-4.pulp.el7.noarch  
python-kombu-3.0.24-5.pulp.el7.noarch  
pulp-rpm-plugins-2.6.1-1.el7.noarch  
python-pulp-bindings-2.6.1-1.el7.noarch  
pulp-puppet-admin-extensions-2.6.1-1.el7.noarch  
python-pulp-rpm-common-2.6.1-1.el7.noarch  
pulp-selinux-2.6.1-1.el7.noarch  
=====
```

[0] <https://pulp.plan.io/issues/942>

History

#1 - 05/29/2015 04:43 PM - mhrivnak

- Priority changed from Normal to High
- Severity changed from 2. Medium to 3. High
- Triaged changed from No to Yes

#2 - 06/01/2015 10:37 AM - Ben.Stanley

I am preserving this repo in its defective state to preserve the evidence.

However, this is blocking other dependent tasks.

I would like to rename the faulty repo and create a new one to provide continuity of service, but there does not appear to be a command to do that.

As deleting and re-creating the repo has been observed to fix the problem, it seems likely that something in the repo state is faulty. I can run diagnostic commands, dump relevant portions of the database, and send relevant diagnostic information.

I have previously hacked MythTV stuff using phpMyAdmin. I may be able to use a similar tool to probe mongodb, but so far an attempt to install robomongo has failed. There are other tools I could try.

I would appreciate assistance with investigating this so that I can do what I need to do to get the repo synced (which will destroy the evidence).

#3 - 06/02/2015 10:32 PM - jortel@redhat.com

- Assignee set to jortel@redhat.com

#4 - 06/02/2015 11:31 PM - jortel@redhat.com

Hey Ben,

Let me be sure I have the correct steps to reproduce:

1. Copy (wget) the repository at: http://elrepo.org/linux/elrepo/el7/x86_64 to a local directory (mirror).
2. Remove nvidia-x11-drv-304xx-304.125-1.el7.elrepo.x86_64.rpm from the local mirror and run createrepo.
3. Created the Pulp repository with feed=mirror.
4. Sync the Pulp repository.
5. Add nvidia-x11-drv-304xx-304.125-1.el7.elrepo.x86_64.rpm to the local mirror and run createrepo.
6. Sync (and publish) the Pulp repository.
7. Observe that the pulp published YUM metadata still contains 304.123-2 and not 304.125-1.

Correct?

#5 - 06/03/2015 12:08 AM - Ben.Stanley

That isn't how the problem was created, but I think it will reproduce the

observation.

I suspect that doing this will bypass the defect that caused it.

It will be interesting to test whether a sync will fix this situation. I have not tested that yet.

I am also concerned that this problem may affect other packages as well. I want a way to test that automatically.

Ben.

#6 - 06/03/2015 04:22 AM - Ben.Stanley

I have tested something like [[<https://pulp.plan.io/issues/1008#note-4>]] as follows:

- 1) Use wget to copy data out of my local pulp repo into a local directory.
- 2) Delete the file nvidia-x11-drv-304xx-304.125-1.el7.elrepo.x86_64.rpm
- 3) createrepo
- 4) create a new pulp repo feeding from the local directory.
[root@bumblebee ~]# pulp-admin rpm repo create --repo-id=elrepo-7-x86_64-test1 --feed=file:///data/test-pulp/elrepo-7-x86_64-test1 --validate true --serve-https false --serve-http true --relative-url=elrepo/elrepo-7-x86_64-test1
- 5) sync the repo
[root@bumblebee ~]# pulp-admin rpm repo sync run --repo-id=elrepo-7-x86_64-test1
- 6) Configure a test machine to point to elrepo-7-x86_64-test1
- 7) Attempt to install kmod-nvidia-304xx on the test machine - note install fails due to missing nvidia-x11-drv-304xx = 304.125

```
-----
[root@stig ~]# yum clean metadata
Loaded plugins: langpacks
Cleaning repos: elrepo epel pulp-2-stable rhel-7-Server rhel-7-Server-extras
: rhel-7-Server-optional
30 metadata files removed
6 sqlite files removed
0 metadata files removed
[root@stig ~]# yum install $(nvidia-detect)
Loaded plugins: langpacks
elrepo | 2.1 kB 00:00
epel | 2.6 kB 00:00
pulp-2-stable | 2.1 kB 00:00
rhel-7-Server | 2.1 kB 00:00
rhel-7-Server-extras | 1.8 kB 00:00
rhel-7-Server-optional | 1.8 kB 00:00
(1/19): elrepo/group | 130 B 00:00
(2/19): elrepo/primary | 78 kB 00:00
(3/19): elrepo/updateinfo | 93 B 00:00
(4/19): epel/x86_64/group | 1.2 MB 00:00
(5/19): epel/x86_64/updateinfo | 376 kB 00:00
(6/19): epel/x86_64/primary | 2.3 MB 00:00
(7/19): epel/x86_64/pkgtags | 1.5 MB 00:00
(8/19): pulp-2-stable/x86_64/updateinfo | 93 B 00:00
(9/19): pulp-2-stable/x86_64/group | 4.2 kB 00:00
(10/19): pulp-2-stable/x86_64/primary | 12 kB 00:00
(11/19): rhel-7-Server/x86_64/group | 584 kB 00:00
(12/19): rhel-7-Server/x86_64/updateinfo | 826 kB 00:00
(13/19): rhel-7-Server/x86_64/primary | 6.2 MB 00:00
(14/19): rhel-7-Server-extras/x86_64/group | 130 B 00:00
(15/19): rhel-7-Server-extras/x86_64/updateinfo | 16 kB 00:00
(16/19): rhel-7-Server-extras/x86_64/primary | 25 kB 00:00
(17/19): rhel-7-Server-optional/x86_64/group | 130 B 00:00
(18/19): rhel-7-Server-optional/x86_64/updateinfo | 716 kB 00:00
(19/19): rhel-7-Server-optional/x86_64/primary | 1.0 MB 00:00
elrepo 146/146
epel 8310/8310
pulp-2-stable 56/56
rhel-7-Server 6990/6990
rhel-7-Server-extras 96/96
rhel-7-Server-optional 5808/5808
Resolving Dependencies
--> Running transaction check
--> Package kmod-nvidia-304xx.x86_64 0:304.125-2.el7.elrepo will be installed
--> Processing Dependency: nvidia-x11-drv-304xx = 304.125 for package: kmod-nvidia-304xx-304.125-2.el7.elrepo.x86_64
--> Finished Dependency Resolution
Error: Package: kmod-nvidia-304xx-304.125-2.el7.elrepo.x86_64 (elrepo)
Requires: nvidia-x11-drv-304xx = 304.125
Available: nvidia-x11-drv-304xx-304.123-2.el7.elrepo.x86_64 (elrepo)
```

```
nvidia-x11-drv-304xx = 304.123-2.el7.elrepo
You could try using --skip-broken to work around the problem
You could try running: rpm -Va --nofiles --nodigest
[root@stig ~]#
```

8) Switch the feed on elrepo-7-x86_64-test1 to come from an upstream source.

```
[root@bumblebee ~]# pulp-admin rpm repo update --repo-id=elrepo-7-x86_64-test1 --feed=http://mirror.aarnet.edu.au/pub/elrepo/elrepo/el7/x86\_64
--validate true --serve-https false --serve-http true --relative-url=elrepo/elrepo-7-x86_64-test1
```

9) Sync elrepo-7-x86_64-test1

```
[root@bumblebee ~]# pulp-admin rpm repo sync run --repo-id=elrepo-7-x86_64-test1
The sync completes successfully.
```

10) Attempt to install kmod-nvidia-304xx on the test machine - install would succeed.

```
[root@stig ~]# yum clean metadata
```

Loaded plugins: langpacks

Cleaning repos: elrepo epel pulp-2-stable rhel-7-Server rhel-7-Server-extras

: rhel-7-Server-optional

36 metadata files removed

7 sqlite files removed

0 metadata files removed

```
[root@stig ~]# yum install $(nvidia-detect)
```

Loaded plugins: langpacks

elrepo | 2.1 kB 00:00

epel | 2.6 kB 00:00

pulp-2-stable | 2.1 kB 00:00

rhel-7-Server | 2.1 kB 00:00

rhel-7-Server-extras | 1.8 kB 00:00

rhel-7-Server-optional | 1.8 kB 00:00

(1/19): elrepo/group | 130 B 00:00

(2/19): elrepo/primary | 78 kB 00:00

(3/19): elrepo/updateinfo | 93 B 00:00

(4/19): epel/x86_64/group | 1.2 MB 00:00

(5/19): epel/x86_64/updateinfo | 376 kB 00:00

(6/19): epel/x86_64/primary | 2.3 MB 00:00

(7/19): epel/x86_64/pkgtags | 1.5 MB 00:00

(8/19): pulp-2-stable/x86_64/group | 4.2 kB 00:00

(9/19): pulp-2-stable/x86_64/updateinfo | 93 B 00:00

(10/19): pulp-2-stable/x86_64/primary | 12 kB 00:00

(11/19): rhel-7-Server/x86_64/group | 584 kB 00:00

(12/19): rhel-7-Server/x86_64/updateinfo | 826 kB 00:00

(13/19): rhel-7-Server/x86_64/primary | 6.2 MB 00:00

(14/19): rhel-7-Server-extras/x86_64/group | 130 B 00:00

(15/19): rhel-7-Server-extras/x86_64/updateinfo | 16 kB 00:00

(16/19): rhel-7-Server-extras/x86_64/primary | 25 kB 00:00

(17/19): rhel-7-Server-optional/x86_64/group | 130 B 00:00

(18/19): rhel-7-Server-optional/x86_64/updateinfo | 716 kB 00:00

(19/19): rhel-7-Server-optional/x86_64/primary | 1.0 MB 00:00

elrepo 147/147

epel 8310/8310

pulp-2-stable 56/56

rhel-7-Server 6990/6990

rhel-7-Server-extras 96/96

rhel-7-Server-optional 5808/5808

Resolving Dependencies

--> Running transaction check

---> Package kmod-nvidia-304xx.x86_64 0:304.125-2.el7.elrepo will be installed

--> Processing Dependency: nvidia-x11-drv-304xx = 304.125 for package: kmod-nvidia-304xx-304.125-2.el7.elrepo.x86_64

--> Running transaction check

---> Package nvidia-x11-drv-304xx.x86_64 0:304.125-1.el7.elrepo will be installed

--> Finished Dependency Resolution

Dependencies Resolved

```
=====
Package Arch Version Repository Size
=====
```

Installing:

kmod-nvidia-304xx x86_64 304.125-2.el7.elrepo elrepo 4.3 M

Installing for dependencies:

nvidia-x11-drv-304xx x86_64 304.125-1.el7.elrepo elrepo 25 M

```
Transaction Summary
=====
```

Install 1 Package (+1 Dependent package)

Total download size: 29 M

Installed size: 127 M

Is this ok [y/d/N]:

11) sync the original repo, which uses the same feed:

```
[root@bumblebee ~]# pulp-admin rpm repo sync run --repo-id=elrepo-7-x86_64-elrepo
```

The sync completes successfully.

12) Configure the test machine to point to elrepo-7-x86_64-elrepo

13) Attempt to install kmod-nvidia-304xx on the test machine

```
[root@stig ~]# yum clean metadata
```

Loaded plugins: langpacks

Cleaning repos: elrepo epel pulp-2-stable rhel-7-Server rhel-7-Server-extras

: rhel-7-Server-optional

36 metadata files removed

7 sqlite files removed

0 metadata files removed

```
[root@stig ~]# yum install $(nvidia-detect)
```

Loaded plugins: langpacks

elrepo | 2.1 kB 00:00

epel | 2.6 kB 00:00

pulp-2-stable | 2.1 kB 00:00

rhel-7-Server | 2.1 kB 00:00

rhel-7-Server-extras | 1.8 kB 00:00

rhel-7-Server-optional | 1.8 kB 00:00

(1/19): elrepo/x86_64/group | 130 B 00:00

(2/19): elrepo/x86_64/primary | 106 kB 00:00

(3/19): elrepo/x86_64/updateinfo | 93 B 00:00

(4/19): epel/x86_64/group | 1.2 MB 00:00

(5/19): epel/x86_64/updateinfo | 376 kB 00:00

(6/19): epel/x86_64/primary | 2.3 MB 00:00

(7/19): epel/x86_64/pkgtags | 1.5 MB 00:00

(8/19): pulp-2-stable/x86_64/group | 4.2 kB 00:00

(9/19): pulp-2-stable/x86_64/updateinfo | 93 B 00:00

(10/19): pulp-2-stable/x86_64/primary | 12 kB 00:00

(11/19): rhel-7-Server/x86_64/group | 584 kB 00:00

(12/19): rhel-7-Server/x86_64/updateinfo | 826 kB 00:00

(13/19): rhel-7-Server/x86_64/primary | 6.2 MB 00:00

(14/19): rhel-7-Server-extras/x86_64/group | 130 B 00:00

(15/19): rhel-7-Server-extras/x86_64/updateinfo | 16 kB 00:00

(16/19): rhel-7-Server-extras/x86_64/primary | 25 kB 00:00

(17/19): rhel-7-Server-optional/x86_64/group | 130 B 00:00

(18/19): rhel-7-Server-optional/x86_64/updateinfo | 716 kB 00:00

(19/19): rhel-7-Server-optional/x86_64/primary | 1.0 MB 00:00

epel 8310/8310

pulp-2-stable 56/56

rhel-7-Server 6990/6990

rhel-7-Server-extras 96/96

rhel-7-Server-optional 5808/5808

Resolving Dependencies

--> Running transaction check

---> Package kmod-nvidia-304xx.x86_64 0:304.125-2.el7.elrepo will be installed

--> Processing Dependency: nvidia-x11-drv-304xx = 304.125 for package: kmod-nvidia-304xx-304.125-2.el7.elrepo.x86_64

--> Finished Dependency Resolution

Error: Package: kmod-nvidia-304xx-304.125-2.el7.elrepo.x86_64 (elrepo)

Requires: nvidia-x11-drv-304xx = 304.125

Available: nvidia-x11-drv-304xx-304.123-2.el7.elrepo.x86_64 (elrepo)

nvidia-x11-drv-304xx = 304.123-2.el7.elrepo

You could try using --skip-broken to work around the problem

You could try running: rpm -Va --nofiles --nodigest

```
[root@stig ~]#
```

In summary:

- 1) This method of creating a repo missing one file, and then syncing it to an upstream repo, does not reproduce the problem.
- 2) The original repo still has the problem, and can't be corrected by running a sync.

It seems that there is something intrinsic to the repo that is broken, rather than anything in the update process.

#7 - 06/03/2015 11:33 PM - jortel@redhat.com

Tried steps in #4 with expected results. After initial sync, only 304.123-2 existed in the pulp repo and was published. After adding 304.125-1 to the upstream; createrepo; then re-running pulp repo sync, both existed in the pulp repo and were published.

```
[jortel@localhost pulp]$ pulp-admin rpm repo content rpm --str-eq="name=nvidia-x11-drv-304xx" --repo-id=issue1
```

008

Arch: x86_64
Buildhost: Build64R7.elrepo.org
Checksum: 3cba49b9cc256bd2cd3a2809473fc2cb49cd7f9d9c9f1e5c7d962e40b64cbaf5
Checksumtype: sha256
Description: This package provides the proprietary NVIDIA 304xx OpenGL X11 display driver files.
Epoch: 0
Filename: nvidia-x11-drv-304xx-304.123-2.el7.elrepo.x86_64.rpm
License: Distributable
Name: nvidia-x11-drv-304xx
Provides: config(nvidia-x11-drv-304xx) = 304.123-2.el7.elrepo-0,
libGL.so.1() (64bit), libOpenCL.so.1() (64bit),
libXvMCNVIDIA_dynamic.so.1() (64bit), libcuda.so.1() (64bit),
libglx.so.304.123() (64bit), libnvcuvid.so.1() (64bit),
libnvidia-cfg.so.1() (64bit),
libnvidia-compiler.so.304.123() (64bit),
libnvidia-eglcore.so.304.123() (64bit), libnvidia-ml.so.1() (64bit),
libnvidia-openssl.so.1() (64bit), libnvidia-tls.so.304.123() (64bit),
libvdpau_nvidia.so.1() (64bit), nvidia-x11-drv-304xx =
304.123-2.el7.elrepo-0, nvidia-x11-drv-304xx(x86-64) =
304.123-2.el7.elrepo-0
Release: 2.el7.elrepo
Requires: /bin/sh, /bin/sh, /sbin/ldconfig, dracut, grubby, grubby,
libX11.so.6() (64bit), libXext.so.6() (64bit), libXv.so.1() (64bit),
libXvMC.so.1() (64bit), libatk-1.0.so.0() (64bit),
libc.so.6(GLIBC_2.2.5) (64bit), libdl.so.2() (64bit),
libdl.so.2(GLIBC_2.2.5) (64bit), libgcc_s.so.1() (64bit),
libgcc_s.so.1(GCC_3.0) (64bit), libgcc_s.so.1(GCC_3.3) (64bit),
libgdk-x11-2.0.so.0() (64bit), libgdk_pixbuf-2.0.so.0() (64bit),
libglib-2.0.so.0() (64bit), libglx.so() (64bit),
libgmodule-2.0.so.0() (64bit), libgobject-2.0.so.0() (64bit),
libgtk-x11-2.0.so.0() (64bit), libm.so.6() (64bit),
libm.so.6(GLIBC_2.2.5) (64bit), libpango-1.0.so.0() (64bit),
libpthread.so.0() (64bit), libpthread.so.0(GLIBC_2.2.5) (64bit),
librt.so.1() (64bit), librt.so.1(GLIBC_2.2.5) (64bit),
libz.so.1() (64bit), nvidia-304xx-kmod = 304.123-0,
nvidia-304xx-kmod = 304.123-0, perl, xorg-x11-server-Xorg <=
1.16.99-0
Vendor: None
Version: 304.123

Arch: x86_64
Buildhost: Build64R7.elrepo.org
Checksum: 172ccc73efcb58fd1099406465bb6a2089ab6e0f396ccddc2a52c9b9d587b9db
Checksumtype: sha256
Description: This package provides the proprietary NVIDIA 304xx OpenGL X11 display driver files.
Epoch: 0
Filename: nvidia-x11-drv-304xx-304.125-1.el7.elrepo.x86_64.rpm
License: Distributable
Name: nvidia-x11-drv-304xx
Provides: config(nvidia-x11-drv-304xx) = 304.125-1.el7.elrepo-0,
libGL.so.1() (64bit), libOpenCL.so.1() (64bit),
libXvMCNVIDIA_dynamic.so.1() (64bit), libcuda.so.1() (64bit),
libglx.so.304.125() (64bit), libnvcuvid.so.1() (64bit),
libnvidia-cfg.so.1() (64bit),
libnvidia-compiler.so.304.125() (64bit),
libnvidia-eglcore.so.304.125() (64bit), libnvidia-ml.so.1() (64bit),
libnvidia-openssl.so.1() (64bit), libnvidia-tls.so.304.125() (64bit),
libvdpau_nvidia.so.1() (64bit), nvidia-x11-drv-304xx =
304.125-1.el7.elrepo-0, nvidia-x11-drv-304xx(x86-64) =
304.125-1.el7.elrepo-0
Release: 1.el7.elrepo
Requires: /bin/sh, /bin/sh, /sbin/ldconfig, dracut, grubby, grubby,
libX11.so.6() (64bit), libXext.so.6() (64bit), libXv.so.1() (64bit),
libXvMC.so.1() (64bit), libatk-1.0.so.0() (64bit),
libc.so.6(GLIBC_2.2.5) (64bit), libdl.so.2() (64bit),
libdl.so.2(GLIBC_2.2.5) (64bit), libgcc_s.so.1() (64bit),
libgcc_s.so.1(GCC_3.0) (64bit), libgcc_s.so.1(GCC_3.3) (64bit),
libgdk-x11-2.0.so.0() (64bit), libgdk_pixbuf-2.0.so.0() (64bit),
libglib-2.0.so.0() (64bit), libglx.so() (64bit),
libgmodule-2.0.so.0() (64bit), libgobject-2.0.so.0() (64bit),
libgtk-x11-2.0.so.0() (64bit), libm.so.6() (64bit),
libm.so.6(GLIBC_2.2.5) (64bit), libpango-1.0.so.0() (64bit),

```
libpthread.so.0() (64bit), libpthread.so.0 (GLIBC_2.2.5) (64bit),
librt.so.1() (64bit), librt.so.1 (GLIBC_2.2.5) (64bit),
libz.so.1() (64bit), nvidia-304xx-kmod = 304.125-0,
nvidia-304xx-kmod = 304.125-0, perl, xorg-x11-server-Xorg <=
1.17.99-0
Vendor: None
Version: 304.125
```

```
<package type="rpm">
  <name>nvidia-x11-drv-304xx</name>
  <arch>x86_64</arch>
  <version epoch="0" rel="2.el7.elrepo" ver="304.123" />
  <checksum pkgid="YES" type="sha256">3cba49b9cc256bd2cd3a2809473fc2cb49cd7f9d9c9f1e5c7d962e40b64cbaf5</checksum>
  <summary>NVIDIA 304xx OpenGL X11 display driver files</summary>
  <description>This package provides the proprietary NVIDIA 304xx OpenGL X11 display driver files.</description>
  <packager>Philip J Perry <mailto:phil@elrepo.org></packager>
  <url>http://www.nvidia.com/</url>
  <time build="1412440457" file="1412440625" />
  <size archive="111576428" installed="111558206" package="26114416" />
  <location href="nvidia-x11-drv-304xx-304.123-2.el7.elrepo.x86_64.rpm"/>
  <package type="rpm">
    <name>nvidia-x11-drv-304xx</name>
    <arch>x86_64</arch>
    <version epoch="0" rel="1.el7.elrepo" ver="304.125" />
    <checksum pkgid="YES" type="sha256">172ccc73efcb58fd1099406465bb6a2089ab6e0f396ccddc2a52c9b9d587b9db</checksum>
    <summary>NVIDIA 304xx OpenGL X11 display driver files</summary>
    <description>This package provides the proprietary NVIDIA 304xx OpenGL X11 display driver files.</description>
    <packager>Philip J Perry <mailto:phil@elrepo.org></packager>
    <url>http://www.nvidia.com/</url>
    <time build="1419085844" file="1419085958" />
    <size archive="113184780" installed="113166559" package="26201740" />
    <location href="nvidia-x11-drv-304xx-304.125-1.el7.elrepo.x86_64.rpm"/>
```

Ben,
Can you attach a tarball of the repodata/ under /var/lib/pulp/published/... for the broken repository?

#8 - 06/05/2015 03:04 AM - Ben.Stanley

- File `elrepo-el7-x86_64-repdata-20150605.bz` added

Please find a tarball of the repodata attached.

#9 - 06/05/2015 10:44 PM - jortel@redhat.com

Inspection of the primary.xml referenced in the repomd.xml confirms 123-2 is included but NOT 125-1. However, the files contained in repodata/ is a little puzzling. As far as I know, Pulp 2.6 generates a new repodata/ directory on each publish so not sure how this repodata/ contains ~4 of each type of metadata file.

```
-rw-r--r-- 1 jortel jortel      93 Jan 17 06:36 012aed642ab0c19fafc32ea33fa7fe8776ae51dc7c4ab7c2fbefb09804e057
34-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel      93 Jan  7 09:39 0304eae93e58677c154a8d073e585bafd92cfe215513e6f5ee271fa1c63374
eb-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel    47755 Dec 26 05:51 1dfe537fc4c4062bd16bfbfd25a18227e6f78e7e01b4f820d0b2b68346ec8ab
42-primary.xml.gz
-rw-r--r-- 1 jortel jortel      93 May 18 12:33 33ebed9fd4a1362048df5149be1ba385bddaebd36ff6ad4203dd95bb634679
bb-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel      93 May 26 05:00 353764a52be47da60e38f397842fab2a26cf0fb2d1afc56a228bbb39efa113
e4-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel      93 Dec 19 19:50 371a4c6886e1a018d8e4b9304b3b6acf97abb9b96a2de402df2e0701a8b46d
10-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel 1193194 Jun  2 21:15 398644749f62717e7d047ee7d808f86d0a17d41c75ee909a2e734e400b32a3
19-primary.xml
-rw-r--r-- 1 jortel jortel      93 Feb  8 23:58 3bc091baf522f75f9b0f32dbe8574ee80b6970ca599918ca3f5090548883a7
2b-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel     130 Jun  2 21:15 3e800f8eb4bfc89f8e320cda3dd709a8e0bef5f4884e3b622df63b17249e82
42-comps.xml
-rw-r--r-- 1 jortel jortel      93 Dec 26 05:51 419863f046e4cbdb06ecd76170b234a3146ba9d6245bead1161ddad1359d25
b9-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel     7545 Dec 29 03:04 464de9c9cf686b86c8029e11c3f5b4a406cc4ea9011d380839fec128f41ca4
f9-other.xml.gz
-rw-r--r-- 1 jortel jortel    10552 Dec 29 03:04 4982a47dbf096309d358187dd23e04ce3f0aa05123fc3e838ea8db64a2d2be
```

```

69-filelists.xml.gz
-rw-r--r-- 1 jortel jortel      93 Jun  2 21:15 4cd0d88e46f50d021615bb36d5f4a3ac725917bdfc4f4af7561e4e8fd74187
39-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel      93 Jan  9 19:13 6211fc6935128575b93ca5e833a634b55c991ecbe4208991ba8032eb9e49d8
76-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel      93 May 27 03:27 67dd51fc369ce7101cb6dbc517fc2f6970039e91c919ca56384e53b6cc9e65
2b-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel      93 Dec 29 03:04 71cce97e3a542c481f50c1e9deefec56186b1cc3a4f1c1ad69d5062045d4cd
94-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel    42113 Dec 19 19:50 78015285d920e47024c2434db8d1fb8fac5df0f6531f19403f8ea16881aa01
85-primary.xml.gz
-rw-r--r-- 1 jortel jortel      93 Jan  2 05:16 78e2e4ee70997f276d04c8b3e62d9afd8cedfd6c005e619e658a886d502455
dc-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel      93 Feb 15 00:28 88ad766a05c364b46c932c22ebcf82f224967a742d2d45c1719b859e61a7de
54-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel      93 Jan 25 10:38 8eb66818cbc38d32c620f4c9e84206cc87d62ecd4c90a016417d33c390726
08-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel      93 Feb 22 14:56 94f514bed10104c944c6087c2fe4c9ac555ed728a5af8800e29e15bfeeca95
a9-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel      93 Jun  1 22:37 9bf9c8c55d4079b5358bf907a75552f2607c3ccdb6d2f60e5a58820093a17d
c2-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel    20762 Jun  2 21:15 a0dbaa1821ef4b7ec09dd6ee3f64ca55708308c1804626eb84b896f7d285d1
fb-filelists.xml.gz
-rw-r--r-- 1 jortel jortel    11658 Dec 26 05:51 ac2ed9e1b64ea179daeb0b27ce54d95bd4cd5fc4d960a82bc1ee20579f4dac
bf-filelists.xml.gz
-rw-r--r-- 1 jortel jortel      93 May 19 10:21 b08cc51c53632a485ca3b6416e98073128b46e27376bf12cd80e81e5b2ead5
88-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel    10552 Dec 19 19:50 b38e613606b0e3e04134ba105b3ea822b7844c2f12d562629b5cf9910bb98a
3f-filelists.xml.gz
-rw-r--r-- 1 jortel jortel      7545 Dec 19 19:50 bb41f13a038a72c6eda57eb7409b276736726ab797c8bdd342cc2aada094d2
e4-other.xml.gz
-rw-r--r-- 1 jortel jortel      8319 Dec 26 05:51 c62015885deea245da5959455bcd6449075eded4c70fe00f43ae6a1efb0a61
4b-other.xml.gz
-rw-r--r-- 1 jortel jortel      93 Jun  1 01:28 cb104b0157a9044957a5cd2913adcc5099daceeb799422b1cb96998e095889
da-updateinfo.xml.gz
-rw-r--r-- 1 jortel jortel    42114 Dec 29 03:04 ec60964e2563a3d4ac3144ed59cb589a50f873b10dc5cbc7dc634594b62ed1
5c-primary.xml.gz
-rw-r--r-- 1 jortel jortel    14250 Jun  2 21:15 ecb7b4787a8fb21a8eecbccf27b3f8e0b9704b058ab6b24501bd598f9e90c1
9e-other.xml.gz

```

#10 - 06/05/2015 11:29 PM - Ben.Stanley

I observe this accumulation of stuff in the repodata dir in all my repos.

This is published by http, not https, if that makes a difference.

This sounds like a new issue.

Regarding the original problem, is there a critical piece of code that I could instrument before running another test? Since pulp is all written in python, I could insert a few logging statements, run a test, and send you the log (or a relevant extract).

I am still running 2.6.1.

Ben.

On 6 June 2015 6:44:23 AM Pulp <dropbox+pulp+c71e@plan.io> wrote:

#11 - 06/08/2015 05:53 PM - jortel@redhat.com

Let's try the following:

1. Change pulp logging to DEBUG level and restart Pulp services; sync the repo and attach logged messages. It's my understanding that the yum importer logs a lot of information at DEBUG.
2. Using pulp-admin, do a repo search like I did in [#7](#). I'd like to confirm that the unit for the missing package (version) has not been created. Knowing this will help determine if it's an importer issue and not a publisher issue.

If we still need more information, We can suggest additional logging statements to add.

#12 - 06/09/2015 03:39 AM - Ben.Stanley

- File `pulp_20150609_DEBUG_sync_test.log` added

```
[root@bumblebee pulp]# less /etc/pulp/server.conf
[server]
```


1. server_name: server_hostname
2. key_url: /pulp/gpg
3. ks_url: /pulp/ks
4. default_login: admin
5. default_password: admin
6. debugging_mode: false
7. log_level: INFO
log_level: DEBUG

pulp-stop
pulp-start

Then I had to kill about a thousand tasks that started up...

```
[root@bumblebee pulp]# date  
Tue Jun 9 10:57:46 AEST 2015  
[root@bumblebee pulp]# pulp-admin rpm repo sync run --repo-id=elrepo-7-x86_64-elrepo
```

Synchronizing Repository [elrepo-7-x86_64-elrepo]

This command may be exited via ctrl+c without affecting the request.

Downloading metadata...
[-]
... completed

Downloading repository content...
[=====] 100%
RPMs: 0/0 items
Delta RPMs: 0/0 items
... completed

Downloading distribution files...
[=====] 100%
Distributions: 0/0 items
... completed

Importing errata...
[-]
... completed

Importing package groups/categories...
[-]
... completed

Task Succeeded

Copying files
[-]
... completed

Initializing repo metadata
[-]
... completed

Publishing Distribution files
[-]
... completed

Publishing RPMs
[-]
... completed

Publishing Delta RPMs
... skipped

Publishing Errata
[-]
... completed

Publishing Comps file
[-]
... completed

Publishing Metadata.

[-]
... completed

Closing repo metadata

[-]
... completed

Generating sqlite files

... skipped

Publishing files to web

[-]
... completed

Writing Listings File

[-]
... completed

Task Succeeded

[root@bumblebee pulp]# date
Tue Jun 9 10:58:43 AEST 2015

/var/log/pulp.log

Jun 9 10:53:33 bumblebee pulp: qpid.messaging:DEBUG: poll(3 fds, timeout=None)
Jun 9 11:00:07 bumblebee pulp: qpid.messaging:DEBUG: poll() returned 0 fds

It seems that the log file is missing something...

/var/log/messages

Jun 9 10:53:33 bumblebee rsyslogd-2177: imjournal: begin to drop messages due to rate-limiting
Jun 9 11:00:07 bumblebee rsyslogd-2177: imjournal: 39210 messages lost due to rate-limiting

So now I edit my /etc/rsyslog.conf file and put in these lines:

1. Try to eliminate
2. rsyslogd-2177: imjournal: begin to drop messages due to rate-limiting
 [SystemLogRateLimitInterval 0
 [SystemLogRateLimitBurst 50

And re-start rsyslogd:

[root@bumblebee pulp]# systemctl restart rsyslog

And re-try the test:

[root@bumblebee pulp]# date
Tue Jun 9 11:10:52 AEST 2015
[root@bumblebee pulp]# pulp-admin rpm repo sync run --repo-id=elrepo-7-x86_64-elrepo

Synchronizing Repository [elrepo-7-x86_64-elrepo]

This command may be exited via ctrl+c without affecting the request.

Downloading metadata...

[/]
... completed

Downloading repository content...

[=====] 100%
RPMs: 0/0 items
Delta RPMs: 0/0 items

... completed

Downloading distribution files...

[=====] 100%
Distributions: 0/0 items
... completed

Importing errata...

[-]
... completed

Importing package groups/categories...

[-]
... completed

Task Succeeded

Copying files

[-]
... completed

Initializing repo metadata

[-]
... completed

Publishing Distribution files

[-]
... completed

Publishing RPMs

[-]
... completed

Publishing Delta RPMs

... skipped

Publishing Errata

[-]
... completed

Publishing Comps file

[-]
... completed

Publishing Metadata.

[-]
... completed

Closing repo metadata

[-]
... completed

Generating sqlite files

... skipped

Publishing files to web

[-]
... completed

Writing Listings File

[-]
... completed

Task Succeeded

```
[root@bumblebee pulp]# date
Tue Jun 9 11:11:20 AEST 2015
[root@bumblebee pulp]#
```

And the log file is huge. See attached file pulp_20150609_DEBUG_sync_test.log .

The interesting stuff is extracted below:

```
Jun 9 11:11:06 bumblebee pulp: celery.pool:DEBUG: TaskPool: Apply <function _fast_trace_task at 0x26d7a28>
(args:('pulp.server.async.tasks._queue_reserved_task', '56dc34a9-8ebf-4cda-9b5c-d6112144d5cb', ['pulp.server.managers.repo.sync.sync',
'd9504b8e-55cd-441e-9b7c-f6d914602f97', u'repository:elrepo-7-x86_64-elrepo', (), {'tags': [u'pulp:repository:elrepo-7-x86_64-elrepo',
'pulp:action:sync'], 'kwargs': {'sync_config_override': None, 'repo_id': u'elrepo-7-x86_64-elrepo'}}), {}, {'utc': True, u'is_eager': False, 'chord': None,
u'group': None, 'args': ['pulp.server.managers.repo.sync.sync', 'd9504b8e-55cd-441e-9b7c-f6d914602f97', u'repository:elrepo-7-x86_64-elrepo', (),
{'tags': [u'pulp:repository:elrepo-7-x86_64-elrepo', 'pulp:action:sync'], 'kwargs': {'sync_config_override': None, 'repo_id': u'elrepo-7-x86_64-elrepo'}}]),
'retries': 0, u'delivery_info': {'priority': 0, u'redelivered': None, u'routing_key': u'resource_manager', u'exchange': u'resource_manager'}, 'expires':
None, u'hostname': 'resource_manager@bumblebee', 'task': 'pulp.server.async.tasks._queue_reserved_task', 'callbacks': None,... kwargs: {})
```

Jun 9 11:11:07 bumblebee pulp: celery.pool:DEBUG: TaskPool: Apply <function _fast_trace_task at 0x2c6ba28>
(args:('pulp.server.managers.repo.sync.sync', 'd9504b8e-55cd-441e-9b7c-f6d914602f97', [], {'sync_config_override': None, 'repo_id':
u'elrepo-7-x86_64-elrepo'}, {'utc': True, u'is_eager': False, 'chord': None, u'group': None, 'args': [], 'retries': 0, u'delivery_info': {u'priority': 0,
u'redelivered': None, u'routing_key': u'reserved_resource_worker-2@bumblebee', u'exchange': u'C.dq'}, 'expires': None, u'hostname':
'reserved_resource_worker-2@bumblebee', 'task': 'pulp.server.managers.repo.sync.sync', 'callbacks': None, u'correlation_id':
u'd9504b8e-55cd-441e-9b7c-f6d914602f97', 'errbacks': None, 'timelimit': (None, None), 'taskset': None, 'kwargs': {'sync_config_override': None,
'repo_id': u'elrepo-7-x86_64-elrepo'}, 'eta': None, u'reply_to': u'67d22352-8c81-3467-94dc-49ccca78e1d8', 'id':
'd9504b8e-55cd-441e-9b7c-f6d914602f97', u'headers': {}}) kwargs: {})

Jun 9 11:11:07 bumblebee pulp: pulp_rpm.plugins.importers.yum.sync:INFO: Downloading metadata from
[http://mirror.aarnet.edu.au/pub/elrepo/el7/x86_64/](http://mirror.aarnet.edu.au/pub/elrepo/elrepo/el7/x86_64/).

Jun 9 11:11:07 bumblebee pulp: requests.packages.urllib3.connectionpool:DEBUG: "GET
http://mirror.aarnet.edu.au/pub/elrepo/el7/x86_64/repodata/repomd.xml HTTP/1.1" 200 None

Jun 9 11:11:08 bumblebee pulp: pulp_rpm.plugins.importers.yum.sync:INFO: Parsing metadata.

Jun 9 11:11:08 bumblebee pulp: pulp_rpm.plugins.importers.yum.sync:INFO: Downloading metadata files.

Jun 9 11:11:08 bumblebee pulp: nectar.downloaders.threaded:DEBUG: starting workers

Jun 9 11:11:08 bumblebee pulp: requests.packages.urllib3.connectionpool:INFO: Starting new HTTP connection (1): oreo.cs.adi-limited.com

Jun 9 11:11:08 bumblebee pulp: requests.packages.urllib3.connectionpool:INFO: Starting new HTTP connection (1): oreo.cs.adi-limited.com

Jun 9 11:11:08 bumblebee pulp: requests.packages.urllib3.connectionpool:INFO: Starting new HTTP connection (1): oreo.cs.adi-limited.com

Jun 9 11:11:08 bumblebee pulp: requests.packages.urllib3.connectionpool:DEBUG: "GET
http://mirror.aarnet.edu.au/pub/elrepo/el7/x86_64/repodata/8957bc085d17f2e0a9afda6569a584caf7d416313b97a6cb7c685facf93053-filelist_s.xml.gz HTTP/1.1" 200 22335

Jun 9 11:11:08 bumblebee pulp: requests.packages.urllib3.connectionpool:DEBUG: "GET
http://mirror.aarnet.edu.au/pub/elrepo/el7/x86_64/repodata/cd94be54953804ce25991cc669c26e2f4fc57fc4e60e99214bc412b83a927fb9-other_xml.gz HTTP/1.1" 200 14628

Jun 9 11:11:08 bumblebee pulp: requests.packages.urllib3.connectionpool:DEBUG: "GET
http://mirror.aarnet.edu.au/pub/elrepo/el7/x86_64/repodata/1292c8each6d11c0b643ea2863d2d866492028984f41e1446546248e30f8d1e1-pri_mary.xml.gz HTTP/1.1" 200 125780

Jun 9 11:11:08 bumblebee pulp: pulp_rpm.plugins.importers.yum.sync:INFO: Parsing metadata.

Jun 9 11:11:08 bumblebee pulp: pulp_rpm.plugins.importers.yum.sync:INFO: Downloading metadata files.

Jun 9 11:11:08 bumblebee pulp: nectar.downloaders.threaded:DEBUG: starting workers

Jun 9 11:11:08 bumblebee pulp: requests.packages.urllib3.connectionpool:INFO: Starting new HTTP connection (1): oreo.cs.adi-limited.com

Jun 9 11:11:08 bumblebee pulp: requests.packages.urllib3.connectionpool:INFO: Starting new HTTP connection (1): oreo.cs.adi-limited.com

Jun 9 11:11:08 bumblebee pulp: requests.packages.urllib3.connectionpool:INFO: Starting new HTTP connection (1): oreo.cs.adi-limited.com

Jun 9 11:11:08 bumblebee pulp: requests.packages.urllib3.connectionpool:DEBUG: "GET
http://mirror.aarnet.edu.au/pub/elrepo/el7/x86_64/repodata/8957bc085d17f2e0a9afda6569a584caf7d416313b97a6cb7c685facf93053-filelist_s.xml.gz HTTP/1.1" 200 22335

Jun 9 11:11:08 bumblebee pulp: requests.packages.urllib3.connectionpool:DEBUG: "GET
http://mirror.aarnet.edu.au/pub/elrepo/el7/x86_64/repodata/cd94be54953804ce25991cc669c26e2f4fc57fc4e60e99214bc412b83a927fb9-other_xml.gz HTTP/1.1" 200 14628

Jun 9 11:11:08 bumblebee pulp: requests.packages.urllib3.connectionpool:DEBUG: "GET
http://mirror.aarnet.edu.au/pub/elrepo/el7/x86_64/repodata/1292c8each6d11c0b643ea2863d2d866492028984f41e1446546248e30f8d1e1-pri_mary.xml.gz HTTP/1.1" 200 125780

Jun 9 11:11:13 bumblebee pulp: requests.packages.urllib3.connectionpool:DEBUG: "GET
http://mirror.aarnet.edu.au/pub/elrepo/el7/x86_64/treeinfo HTTP/1.1" 404 323

Jun 9 11:11:13 bumblebee pulp: nectar.downloaders.threaded:DEBUG: download failed: Download of
http://mirror.aarnet.edu.au/pub/elrepo/el7/x86_64/treeinfo failed with code 404: Not Found

Jun 9 11:11:14 bumblebee pulp: celery.pool:DEBUG: TaskPool: Apply <function _fast_trace_task at 0x26d7a28>
(args:('pulp.server.async.tasks._queue_reserved_task', 'af9d9e69-a79a-4dbe-a424-c95b1508bdcd', ['pulp.server.managers.repo.publish.publish',
'f84c98c2-9936-44f2-aa70-bf5edc0c16e6', u'repository:elrepo-7-x86_64-elrepo', (], {'tags': [u'pulp:repository:elrepo-7-x86_64-elrepo',
'pulp:action:publish'], 'kwargs': {'publish_config_override': None, 'repo_id': u'elrepo-7-x86_64-elrepo', 'distributor_id': u'yum_distributor'}}), {'utc':
True, u'is_eager': False, 'chord': None, u'group': None, 'args': ['pulp.server.managers.repo.publish.publish', 'f84c98c2-9936-44f2-aa70-bf5edc0c16e6',
u'repository:elrepo-7-x86_64-elrepo', (], {'tags': [u'pulp:repository:elrepo-7-x86_64-elrepo', 'pulp:action:publish'], 'kwargs': {'publish_config_override':
None, 'repo_id': u'elrepo-7-x86_64-elrepo', 'distributor_id': u'yum_distributor'}}), 'retries': 0, u'delivery_info': {u'priority': 0, u'redelivered': None,
u'routing_key': u'resource_manager', u'exchange': u'resource_manager'}, 'expires': None, u'hostname':... kwargs: {})

Jun 9 11:11:14 bumblebee pulp: celery.pool:DEBUG: TaskPool: Apply <function _fast_trace_task at 0x2c6ba28>
(args:('pulp.server.managers.repo.publish.publish', 'f84c98c2-9936-44f2-aa70-bf5edc0c16e6', [], {'publish_config_override': None, 'repo_id':
u'elrepo-7-x86_64-elrepo', 'distributor_id': u'yum_distributor'}, {'utc': True, u'is_eager': False, 'chord': None, u'group': None, 'args': [], 'retries': 0,
u'delivery_info': {u'priority': 0, u'redelivered': None, u'routing_key': u'reserved_resource_worker-2@bumblebee', u'exchange': u'C.dq'}, 'expires': None,
u'hostname': 'reserved_resource_worker-2@bumblebee', 'task': 'pulp.server.managers.repo.publish.publish', 'callbacks': None, u'correlation_id':
u'f84c98c2-9936-44f2-aa70-bf5edc0c16e6', 'errbacks': None, 'timelimit': (None, None), 'taskset': None, 'kwargs': {'publish_config_override': None,
'repo_id': u'elrepo-7-x86_64-elrepo', 'distributor_id': u'yum_distributor'}, 'eta': None, u'reply_to': u'67d22352-8c81-3467-94dc-49ccca78e1d8', 'id':
'f84c98c2-9936-44f2-aa70-bf5edc0c16e6', u'headers': {}}) kwargs: {})

Jun 9 11:11:14 bumblebee pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.distributor:DEBUG: Publishing yum repository:
elrepo-7-x86_64-elrepo

Jun 9 11:11:14 bumblebee pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.metadata:DEBUG: Opening metadata file:
/var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/repomd.xml

Jun 9 11:11:14 bumblebee pulp: qpid.messaging:DEBUG: SENT[2f64170]: Message(subject="task.started", properties={'qpid.subject': 'task.started'},
content={'body': <read-only buffer for 0x3053c30, size 4, offset 0 at 0x7f907c050df0>, "headers": {'hostname':
'reserved_resource_worker_2@bumblebee'}, 'content type': 'application/json', 'properties': {'body_encoding': 'base64', 'delivery_info': {'priority': 0,
'routing_key': u'task.started', 'exchange': 'celeryev', 'delivery_mode': 2, 'delivery_tag': 6848}, 'content encoding': 'utf-8'})

Jun 9 11:11:14 bumblebee pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.metadata:WARNING: Overwriting existing metadata file
[var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/repomd.xml]

Jun 9 11:11:14 bumblebee pulp: qpid.messaging:DEBUG: poll(1 fds, timeout=3000.0)

Jun 9 11:11:14 bumblebee pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.metadata:DEBUG: Opening metadata file handle for
[var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/repomd.xml]

Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.metadata:DEBUG: Writing XML header into metadata file:
/var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/repomd.xml
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.util.metadata_writer:DEBUG: Opening metadata file:
/var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/filelists.xml.gz
Jun 9 11:11:14 bumblebee-pulp: qpid.messaging:DEBUG: poll() returned 1 fds
Jun 9 11:11:14 bumblebee-pulp: qpid.messaging:DEBUG: poll(4 fds, timeout=None)
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.util.metadata_writer:DEBUG: Opening metadata file handle for
[var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/filelists.xml.gz]
Jun 9 11:11:14 bumblebee-pulp: qpid.messaging:DEBUG: poll() returned 1 fds
Jun 9 11:11:14 bumblebee-pulp: qpid.messaging:DEBUG: poll(1 fds, timeout=3000.0)
Jun 9 11:11:14 bumblebee-pulp: qpid.messaging:DEBUG: poll(3 fds, timeout=None)
Jun 9 11:11:14 bumblebee-pulp: qpid.messaging:DEBUG: poll(4 fds, timeout=None)
Jun 9 11:11:14 bumblebee-pulp: qpid.messaging:DEBUG: poll() returned 1 fds
Jun 9 11:11:14 bumblebee-pulp: qpid.messaging:DEBUG: poll(1 fds, timeout=1000.0)
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.util.metadata_writer:DEBUG: Opening metadata file:
/var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/other.xml.gz
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.util.metadata_writer:DEBUG: Opening metadata file handle for
[var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/other.xml.gz]
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.util.metadata_writer:DEBUG: Opening metadata file:
/var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/primary.xml.gz
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.util.metadata_writer:DEBUG: Opening metadata file handle for
[var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/primary.xml.gz]
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.util.metadata_writer:DEBUG: Closing metadata file:
/var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/filelists.xml.gz
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:DEBUG: (21867-76896) Writing repomd metadata:
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:DEBUG: (21867-76896) <data
type="filelists"><location href="repodata/f5599744abdb5e65ec311bf619bea7c76589af5bdba6c26b81dbfdeb7a298c62-filelists.xml.gz"
></location></filelists></data><timestamp>1433812274</timestamp><size>20762</size><checksum
type="sha256">f5599744abdb5e65ec311bf619bea7c76589af5bdba6c26b81dbfdeb7a298c62</checksum><open-size>208904</open-size><open-ch
ecksum type="sha256">e7610cf44bb4cddad92212a703ab41aa26ce25221f3f0bcc39d786ab5bbb0c6</open-checksum></data>
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.util.metadata_writer:DEBUG: Closing metadata file:
/var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/other.xml.gz
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:DEBUG: (21867-76896) Writing repomd metadata:
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:DEBUG: (21867-76896) <data
type="other"><location href="repodata/64f32d3a1922faed8943b9781a8e86093942238a1439acebfa88eb5b9b489255-other.xml.gz"
></location></other></data><timestamp>1433812274</timestamp><size>14250</size><checksum
type="sha256">64f32d3a1922faed8943b9781a8e86093942238a1439acebfa88eb5b9b489255</checksum><open-size>116978</open-size><open-c
hecksum type="sha256">ce455d017b92d05a396e94b25e64ea101ae151ffc05e41746f7582b512fa58e0</open-checksum></data>
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.util.metadata_writer:DEBUG: Closing metadata file:
/var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/primary.xml.gz
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:DEBUG: (21867-76896) Writing repomd metadata:
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:DEBUG: (21867-76896) <data
type="primary"><location href="repodata/43d3b4201ad4390e67f03a3f485075dbd8bfe076a91b84f88ee60b93f564a8ac-primary.xml.gz"
></location></primary></data><timestamp>1433812274</timestamp><size>108911</size><checksum
type="sha256">43d3b4201ad4390e67f03a3f485075dbd8bfe076a91b84f88ee60b93f564a8ac</checksum><open-size>1193194</open-size><open-c
hecksum type="sha256">233b1e1a5f08e026e804fb30e096743d1170ca887e0f2f7060be36156b8dbad8</open-checksum></data>
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.metadata:DEBUG: Opening metadata file:
/var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/updateinfo.xml.gz
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.metadata:DEBUG: Opening metadata file handle for
[var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/updateinfo.xml.gz]
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.metadata:DEBUG: Writing XML header into metadata file:
/var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/updateinfo.xml.gz
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.metadata:DEBUG: Closing metadata file:
/var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/updateinfo.xml.gz
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:DEBUG: (21867-76896) Writing repomd metadata:
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:DEBUG: (21867-76896) <data
type="updateinfo"><location href="repodata/2a4729573b542a011b89b7bbee6f1df1689c41ee6fec68054ec23f75950263e70-updateinfo.xml.gz"
></location></updateinfo></data><timestamp>1433812274</timestamp><size>93</size><checksum
type="sha256">2a4729573b542a011b89b7bbee6f1df1689c41ee6fec68054ec23f75950263e70</checksum><open-size>60</open-size><open-checks
um type="sha256">e12d66b45503fbedd6f6211054309902124b1e254cdc7634961025d5d7aaa1ed</open-checksum></data>
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.metadata:DEBUG: Opening metadata file:
/var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/comps.xml
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.metadata:DEBUG: Opening metadata file handle for
[var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/comps.xml]
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.metadata:DEBUG: Writing XML header into metadata file:
/var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/comps.xml
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.metadata:DEBUG: Closing metadata file:
/var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/comps.xml
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:DEBUG: (21867-76896) Writing repomd metadata:
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:DEBUG: (21867-76896) <data
type="group"><location href="repodata/3e800f8eb4bfc89f8c320cda3dd709a8e0bef5f4884e3b622df63b17249e8242-comps.xml"
></location></group></data><timestamp>1433812274</timestamp><size>130</size><checksum
type="sha256">3e800f8eb4bfc89f8c320cda3dd709a8e0bef5f4884e3b622df63b17249e8242</checksum></data>
Jun 9 11:11:14 bumblebee-pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.metadata:DEBUG: Closing metadata file:
/var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor/repodata/repomd.xml

```

Jun 9 11:11:14 bumblebee pulp: pulp.plugins.util.publish_step:DEBUG: Copying tree from
/var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributor to
/var/lib/pulp/published/yum/master/yum_distributor/elrepo-7-x86_64-elrepo/1433812274.53
Jun 9 11:11:14 bumblebee pulp: pulp.plugins.util.misc:DEBUG: Clearing out directory:
/var/lib/pulp/published/yum/master/yum_distributor/elrepo-7-x86_64-elrepo
Jun 9 11:11:14 bumblebee pulp: pulp.server.managers.repo.publish:DEBUG: publish succeeded for repo [elrepo-7-x86_64-elrepo] with distributor ID
[yum_distributor]
Jun 9 11:11:14 bumblebee pulp: pulp.server.async.tasks:DEBUG: Task successful : [f84c98c2-9936-44f2-aa70-bf5edc0c16e6]
Jun 9 11:11:14 bumblebee pulp: celery.pool:DEBUG: TaskPool: Apply <function _fast_trace_task at 0x2c6ba28>
(args:('pulp.server.async.tasks._release_resource', '28d325ec-c5c6-44a8-a000-4b7770e08108', ('f84c98c2-9936-44f2-aa70-bf5edc0c16e6',)), [], {'utc':
True, 'is_eager': False, 'chord': None, 'group': None, 'args': ('f84c98c2-9936-44f2-aa70-bf5edc0c16e6',), 'retries': 0, 'delivery_info': {'priority': 0,
'rdelivered': None, 'routing_key': 'reserved_resource_worker_2@bumblebee', 'exchange': 'u'C.dq'}, 'expires': None, 'hostname':
'reserved_resource_worker_2@bumblebee', 'task': 'pulp.server.async.tasks._release_resource', 'callbacks': None, 'correlation_id':
'u28d325ec-c5c6-44a8-a000-4b7770e08108', 'errbacks': None, 'timelimit': (None, None), 'taskset': None, 'kwargs': {}, 'eta': None, 'reply_to':
'u67d22352-8e81-3467-94dc-49ccca78e1d8', 'id': '28d325ec-c5c6-44a8-a000-4b7770e08108', 'headers': {}}) kwargs: {})

```

It seems to me that DEBUG is not high enough to see anything at the package level.

#13 - 06/10/2015 08:52 AM - Ben.Stanley

Reproduces on pulp 2.6.2 .

#14 - 06/10/2015 10:59 PM - jortel@redhat.com

Please add additional logging statements described in the following patches. Then, restart pulp services and re-sync the affected repository. Please attach the relevant portion of the log.

I suggest you reset the logging level back to INFO to reduce the noise level in the logs.

```

diff --git a/plugins/pulp_rpm/plugins/importers/yum/existing.py b/plugins/pulp_rpm/plugins/importers/yum/existing.py
index ceb5ac9..11a915c 100644
--- a/plugins/pulp_rpm/plugins/importers/yum/existing.py
+++ b/plugins/pulp_rpm/plugins/importers/yum/existing.py
@@ -91,6 +91,10 @@ def check_all_and_associate(wanted, sync_conduit):
     named tuples received as input were not found on the server.
     :rtype: set
     """
+
+ # debugging
+ for x in wanted:
+     _LOGGER.info('ISSUE-1008 - WANTED: %s', x)
+
     sorted_units = _sort_by_type(wanted)
     for unit_type, values in sorted_units.iteritems():
         model = models.TYPE_MAP[unit_type]
@@ -123,6 +127,10 @@ def check_all_and_associate(wanted, sync_conduit):
     # 1125388 - make sure we keep storage_path on the new unit model obj
     downloaded_unit.storage_path = unit.storage_path
+
+ # debug
+ _LOGGER.info('ISSUE-1008 - UNIT SAVED (not downloaded): %s', downloaded_unit)
+
     sync_conduit.save_unit(downloaded_unit)
+
     # Discard already downloaded unit from the return value.

```

```

diff --git a/plugins/pulp_rpm/plugins/importers/yum/listener.py b/plugins/pulp_rpm/plugins/importers/yum/listener.py
index ca5d6fe..216b19c 100644
--- a/plugins/pulp_rpm/plugins/importers/yum/listener.py
+++ b/plugins/pulp_rpm/plugins/importers/yum/listener.py
@@ -62,6 +62,9 @@ class ContentListener(DownloadEventListener):
     """
     model = report.data
+
+ # debug
+ _logger.info('ISSUE-1008 - DOWNLOADED: %s', model)
+
     try:
         self._verify_size(model, report)
         self._verify_checksum(model, report)
@@ -84,6 +87,10 @@ class ContentListener(DownloadEventListener):
     shutil.move(report.destination, unit.storage_path)

```

```

# save unit
self.sync_conduit.save_unit(unit)
+
+ # debug
+ _logger.info('ISSUE-1008 - UNIT SAVED (downloaded): %s', unit)
+
self.progress_report['content'].success(model)
self.sync_conduit.set_progress(self.progress_report)

diff --git a/plugins/pulp_rpm/plugins/importers/yum/sync.py b/plugins/pulp_rpm/plugins/importers/yum/sync.py
index 60af275..5259e96 100644
--- a/plugins/pulp_rpm/plugins/importers/yum/sync.py
+++ b/plugins/pulp_rpm/plugins/importers/yum/sync.py
@@ -384,6 +384,10 @@ class RepoSync(object):
# and associate them to the repo
to_download = existing.check_all_and_associate(not_found_in_the_repo,
self.sync_conduit)
+
+ # debugging
+ for x in to_download:
+ _logger.info('ISSUE-1008 - TO BE DOWNLOADED: %s', x)
+
count = len(to_download)
size = 0
for unit in to_download:

```

#15 - 06/11/2015 10:13 AM - Ben.Stanley

I performed these steps:

1. added the additional debugging statements mentioned in [#14](#)
2. reverted the log level to info
3. re-started pulp
4. date; pulp-admin rpm repo sync run --repo-id=elrepo-el7-x86_64-elrepo; date

I observed that none of the new logging statements were recorded in the log.

After chasing around in the repo sync code for a while, I determined that the unit in question was not downloaded because it was already downloaded and in the database. This was established by inspecting various stuff in sync.py RepoSync._decide_rpms_to_download().

I found the following facts regarding nvidia-x11-drv-304xx-304.125-1.el7.elrepo.x86_64.rpm :

1. It occurs in the wanted list: rpm(name='nvidia-x11-drv-304xx', epoch='0', version='304.125', release='1.el7.elrepo', arch='x86_64', checksumtype='sha256', checksum='172ccc73efcb58fd1099406465bb6a2089ab6e0f396ccddc2a52c9b9d587b9db')
2. It does NOT occur in not_found_in_the_repo returned by existing.check_repo() .
3. It occurs in the repo list of units inside existing.check_repo()

```

Jun 11 16:29:04 bumblebee pulp: pulp_rpm.plugins.importers.yum.existing:INFO: ISSUE-1008 - unit: Unit [key={u'name':
u'nvidia-x11-drv-304xx', u'checksum': u'3c3ba49b9c256bd2cd3a2809473fc2cb49cd7f9d9c9f1e5c7d962e40b64cbaf5', u'epoch': u'0', u'version':
u'304.123', u'release': u'2.el7.elrepo', u'arch': u'x86_64', u'checksumtype': u'sha256'}] [type=rpm] [id=130a8932-5007-4938-979f-b1a056062ac5]
Jun 11 16:29:04 bumblebee pulp: pulp_rpm.plugins.importers.yum.existing:INFO: ISSUE-1008 - unit: Unit [key={u'name':
u'nvidia-x11-drv-304xx', u'checksum': u'172ccc73efcb58fd1099406465bb6a2089ab6e0f396ccddc2a52c9b9d587b9db', u'epoch': u'0', u'version':
u'304.125', u'release': u'1.el7.elrepo', u'arch': u'x86_64', u'checksumtype': u'sha256'}] [type=rpm] [id=c75123e2-6fed-4699-b17b-f7a9ee008b22]

```

So it appears that it has been downloaded. What is going on in the publisher then? Or is the publishing being skipped too?

It appears that the publisher is performing an incremental publish.

The publisher is not deleting previously published -filelists.xml.gz, -other.xml.gz, -primary.xml.gz, -updateinfo.xml.gz and -comps.xml files. They just accumulate.

```

Jun 11 17:11:26 bumblebee pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:INFO: (12383-60
672) ISSUE-1008 - Writing repomd metadata:
Jun 11 17:11:26 bumblebee pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:INFO: (12383-60
672) <data type="filelists"><location href="repodata/ae182b2b22533d4952ad23e3ee92c89ee69c1f09db1509dcfdc4f08d9
7867114-filelists.xml.gz" /><timestamp>1434006686</timestamp><size>20762</size><checksum type="sha256">ae182b2
b22533d4952ad23e3ee92c89ee69c1f09db1509dcfdc4f08d97867114</checksum><open-size>208904</open-size><open-checksu
m type="sha256">e7610cf44bb4cddadf92212a703ab41aa26ce25221f3f0bcc39d786ab5bbb0c6</open-checksum></data>
Jun 11 17:11:26 bumblebee pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:INFO: (12383-60
672) ISSUE-1008 - Writing repomd metadata:
Jun 11 17:11:26 bumblebee pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:INFO: (12383-60
672) <data type="other"><location href="repodata/6dda5ad6d8a4773e1e982842a9f1629f70ecbf7ce1779a09659fe89bbc226
5e9-other.xml.gz" /><timestamp>1434006686</timestamp><size>14250</size><checksum type="sha256">6dda5ad6d8a4773
e1e982842a9f1629f70ecbf7ce1779a09659fe89bbc2265e9</checksum><open-size>116978</open-size><open-checksum type="
sha256">ee455d017b92d05a396e94b25e64ea101ae151ffc05e41746f7582b512fa58c0</open-checksum></data>
Jun 11 17:11:26 bumblebee pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:INFO: (12383-60
672) ISSUE-1008 - Writing repomd metadata:

```

```
Jun 11 17:11:26 bumblebee pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:INFO: (12383-60
672) <data type="primary"><location href="repodata/b8744b3cc38cd5cd55e8adb48c97fcecbb08434598d39605de02f48c23
1bfe2-primary.xml.gz" /><timestamp>1434006686</timestamp><size>108911</size><checksum type="sha256">b8744b3cc3
8cd5cd55e8adb48c97fcecbb08434598d39605de02f48c231bfe2</checksum><open-size>1193194</open-size><open-checksum
type="sha256">233b1e1a5f08c926e804fb30c096743d1170ca887e0f2f7060bc36156b8dbad8</open-checksum></data>
Jun 11 17:11:26 bumblebee pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:INFO: (12383-60
672) ISSUE-1008 - Writing repomd metadata:
Jun 11 17:11:26 bumblebee pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:INFO: (12383-60
672) <data type="updateinfo"><location href="repodata/9f329c15b3270c4aeb6d362f08dd42d81d24de54121f1187edc8c575
88b044ec-updateinfo.xml.gz" /><timestamp>1434006686</timestamp><size>93</size><checksum type="sha256">9f329c15
b3270c4aeb6d362f08dd42d81d24de54121f1187edc8c57588b044ec</checksum><open-size>60</open-size><open-checksum typ
e="sha256">c12d66b45503fbedd6f6211054309902124b1e254cdc7634961025d5d7aaaled</open-checksum></data>
Jun 11 17:11:26 bumblebee pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:INFO: (12383-60
672) ISSUE-1008 - Writing repomd metadata:
Jun 11 17:11:26 bumblebee pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.repomd:INFO: (12383-60
672) <data type="group"><location href="repodata/3e800f8eb4bfc89f8e320cda3dd709a8e0bef5f4884e3b622df63b17249e8
242-comps.xml" /><timestamp>1434006686</timestamp><size>130</size><checksum type="sha256">3e800f8eb4bfc89f8e32
0cda3dd709a8e0bef5f4884e3b622df63b17249e8242</checksum></data>
```

The warning "Overwriting existing metadata file" is not useful.

```
Jun 11 17:11:26 bumblebee pulp: pulp.plugins.pulp_rpm.plugins.distributors.yum.metadata.metadata:WARNING: Over
writing existing metadata file [/var/lib/pulp/working/repos/elrepo-7-x86_64-elrepo/distributors/yum_distributo
r/repodata/repomd.xml]
```

publish.py PublishRpmStep.publish_unit() is not being called. Note that self.fast_forward is set to True by Publisher.__init__ .

The problem comes down to performing an incremental publish, which is decided in publish.py Publisher.__init__() at the line if os.path.exists(specific_master):

I forced a complete publish by deleting the specific_master /data/var_lib_pulp/published/yum/http/repos/elrepo/elrepo/el7/x86_64 and running a sync. This has now fixed the problem, and the problem rpm is now visible to the client.

How can this be prevented from happening again?

It seems that the repo was published incorrectly by some pulp version < 2.6.1. However, more recent versions of pulp were unable to fix the problem, because

1. they all performed an incremental publish
2. which fails to actually update anything
3. as it skips processing the RPMs

In summary, I have become very wary of skipping steps as an 'optimisation', because if there is any problem in the previously downloaded or published data, it is difficult to fix unless you know how to force the steps not to be skipped.

#16 - 06/11/2015 05:18 PM - bcourt

Hi Ben,

Was this Pulp installation ever running Pulp 2.5.0? There were some specific upgrade instructions for upgrading from 2.5.0 to a newer version (<https://pulp.readthedocs.org/en/latest/user-guide/release-notes/2.5.x.html#upgrade-instructions-for-2-5-0-2-5-1>).

More generally, incremental publishing does process the RPMs that have been added to the repo since the last publish completed successfully. It specifically skips all the RPMs that have already been published. If no RPMs have been added since the last successful publish then the step is skipped entirely. If an rpm repo publish fails for any reason it should not update the last_published timestamp on the repository. For a onetime update of all your repositories to force the next publish to not be incremental the simplest thing is to update all the yum distributors records in the mongo database directly to set last_publish to null, (db.repo_distributors.update({"distributor_type_id": "yum_distributor"},{\$set: {"last_update": null }},{multi: true}), and then republish all the repos normally. Alternately if you can tolerate the down time for the repos, just delete all the directories in /var/lib/pulp/published/yum/* and then republish your repos.

Given the behavior of the incremental publish, I'm concerned that you were able to get into a state where things were incorrect. Knowing the code I would like to understand how that could happen but there is not a good way to do that after the fact.

If you would like to disable incremental publishing there is not currently a way to do that. Please open a feature request RFE if you would like it. Thanks!

#17 - 06/11/2015 05:45 PM - Ben.Stanley

This pulp installation started at version 2.5, but the repo in question was created with pulp 2.6.1, so I don't think that https://bugzilla.redhat.com/show_bug.cgi?id=1171509 explains it.

I have had problems with large numbers of tasks on this server, and I periodically have to cancel tasks. It is possible that canceling a sync may have caused the error state.

The way to catch the cause of this problem is to add code to perform a deep validation that the published repo is equal to the source repo, after a sync. This will take extra time to perform, and should be optionally enabled, but will determine the conditions that produce the problem. Think of it as a paranoid debug mode with double checking built in.

On 12 June 2015 1:19:07 am Pulp <dropbox+pulp+c71e@plan.io> wrote:

#18 - 02/29/2016 05:28 PM - jortel@redhat.com

- Assignee deleted (*jortel@redhat.com*)

#19 - 04/12/2019 09:06 PM - bmbouter

- Status changed from NEW to CLOSED - WONTFIX

#20 - 04/12/2019 09:11 PM - bmbouter

Pulp 2 is approaching maintenance mode, and this Pulp 2 ticket is not being actively worked on. As such, it is being closed as WONTFIX. Pulp 2 is still accepting contributions though, so if you want to contribute a fix for this ticket, please reopen or comment on it. If you don't have permissions to reopen this ticket, or you want to discuss an issue, please reach out via the [developer mailing list](#).

#21 - 04/15/2019 10:48 PM - bmbouter

- Tags Pulp 2 added

Files

elrepo-el7-x86_64-repodata-20150605.bz	287 KB	06/05/2015	Ben.Stanley
pulp_20150609_DEBUG_sync_test.log	606 KB	06/09/2015	Ben.Stanley