

## Pulp - Issue #9213

### DNF gets wrong rpm download size, until package downloaded with curl, then works

08/04/2021 11:39 PM - optiz0r

<b>Status:</b>	CLOSED - CURRENTRELEASE	<b>Start date:</b>	
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	dkliban@redhat.com	<b>Estimated time:</b>	0:00 hour
<b>Category:</b>		<b>Groomed:</b>	No
<b>Sprint/Milestone:</b>	3.16.0	<b>Sprint Candidate:</b>	No
<b>Severity:</b>	3. High	<b>Tags:</b>	
<b>Version:</b>		<b>Sprint:</b>	Sprint 104
<b>Platform Release:</b>		<b>Quarter:</b>	
<b>OS:</b>			
<b>Triaged:</b>	Yes		

#### Description

I have setup Katello with [https://rpm.releases.hashicorp.com/RHEL/8/x86\\_64/stable/](https://rpm.releases.hashicorp.com/RHEL/8/x86_64/stable/) as an on-demand RPM repo. A content host is subscribed to this via the Library. Attempting to install any package results in an error, that the downloaded RPM size did not match the expected size. This is repeatable, persists across dnf clean all, and happens on multiple hosts.

Downloading the package with wget via the http repo url may result in an error the first time, but subsequent downloads will fetch the entire file successfully. As soon as the file has been downloaded successfully once via wget, dnf will then happily install the package. The behaviour is the same for all packages within the repo.

I have done multiple "Complete syncs", "Verify content checksum", and "Regenerate repository metadata" operations, none of which have had any impact.

Two examples, with some output snipped for brevity:

```
[root@alma-canary ~]# yum install waypoint -y
Stable (EL8)
          36 kB/s | 1.7 kB      00:00
Dependencies resolved.
=====
Package                Architecture      Version           Reposit
tory                   Size
=====
Installing:
waypoint                x86_64           0.4.2-1          sihnon
_hashicorp_stable-el8 46 M
Transaction Summary
=====
Install 1 Package

Total download size: 46 M
Installed size: 100 M
Downloading Packages:
[MIRROR] waypoint-0.4.2-1.x86_64.rpm: Interrupted by header callback: Server reports Content-Lengt
h: 47747976 but expected size is: 47746472
[MIRROR] waypoint-0.4.2-1.x86_64.rpm: Interrupted by header callback: Server reports Content-Lengt
h: 47747976 but expected size is: 47746472
[MIRROR] waypoint-0.4.2-1.x86_64.rpm: Interrupted by header callback: Server reports Content-Lengt
h: 47747976 but expected size is: 47746472
[MIRROR] waypoint-0.4.2-1.x86_64.rpm: Interrupted by header callback: Server reports Content-Lengt
h: 47747976 but expected size is: 47746472
[FAILED] waypoint-0.4.2-1.x86_64.rpm: No more mirrors to try - All mirrors were already tried with
```

out success

The downloaded packages were saved in cache until the next successful transaction.  
You can remove cached packages by executing 'yum clean packages'.  
Error: Error downloading packages:  
Cannot download Packages/w/waypoint-0.4.2-1.x86\_64.rpm: All mirrors were tried

```
[root@alma-canary ~]# yum install waypoint -y
Stable (EL8)
```

```
36 kB/s | 1.7 kB 00:00
```

Dependencies resolved.

```
=====
```

Package	Architecture	Version	Repository
waypoint	x86_64	0.4.2-1	sihnon

```
=====
```

Installing:

waypoint	x86_64	0.4.2-1	sihnon
_hashicorp_stable-el8		46 M	

Transaction Summary

```
=====
```

Install 1 Package

Total download size: 46 M

Installed size: 100 M

Downloading Packages:

```
[MIRROR] waypoint-0.4.2-1.x86_64.rpm: Interrupted by header callback: Server reports Content-Length: 47747976 but expected size is: 47746472
[MIRROR] waypoint-0.4.2-1.x86_64.rpm: Interrupted by header callback: Server reports Content-Length: 47747976 but expected size is: 47746472
[MIRROR] waypoint-0.4.2-1.x86_64.rpm: Interrupted by header callback: Server reports Content-Length: 47747976 but expected size is: 47746472
[MIRROR] waypoint-0.4.2-1.x86_64.rpm: Interrupted by header callback: Server reports Content-Length: 47747976 but expected size is: 47746472
[FAILED] waypoint-0.4.2-1.x86_64.rpm: No more mirrors to try - All mirrors were already tried with out success
```

The downloaded packages were saved in cache until the next successful transaction.  
You can remove cached packages by executing 'yum clean packages'.  
Error: Error downloading packages:  
Cannot download Packages/w/waypoint-0.4.2-1.x86\_64.rpm: All mirrors were tried

```
https://whitefall.example.com/pulp/content/sihnon/Library/custom/hashicorp/stable-el8/Packages/w/w
aypoint-0.4.2-1.x86_64.rpm
--2021-08-04 21:13:13-- https://whitefall.example.com/pulp/content/sihnon/Library/custom/hashicor
p/stable-el8/Packages/w/waypoint-0.4.2-1.x86_64.rpm
Resolving whitefall.example.com (whitefall.example.com)... 81.xx.xx.141
Connecting to whitefall.example.com (whitefall.example.com)|81.xx.xx.141|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 47747976 (46M) [binary/octet-stream]
Saving to: 'waypoint-0.4.2-1.x86_64.rpm'
```

```
waypoint-0.4.2-1.x86_64.rpm          99%[=====]
======> ] 45.53M 11.2MB/s  in 3.9s
```

2021-08-04 21:13:18 (11.8 MB/s) - Connection closed at byte 47746472. Retrying.

```
--2021-08-04 21:13:19-- (try: 2) https://whitefall.example.com/pulp/content/sihnon/Library/custo
m/hashicorp/stable-el8/Packages/w/waypoint-0.4.2-1.x86_64.rpm
Connecting to whitefall.example.com (whitefall.example.com)|81.xx.xx.141|:443... connected.
HTTP request sent, awaiting response... 416 Requested Range Not Satisfiable
```

The file is already fully retrieved; nothing to do.

```
[root@alma-canary ~]# wget --no-check-certificate https://whitefall.example.com/pulp/content/sihnon/Library/custom/hashicorp/stable-el8/Packages/w/waypoint-0.4.2-1.x86_64.rpm
--2021-08-04 21:13:35-- https://whitefall.example.com/pulp/content/sihnon/Library/custom/hashicorp/stable-el8/Packages/w/waypoint-0.4.2-1.x86_64.rpm
Resolving whitefall.example.com (whitefall.example.com)... 81.xx.xx.141
Connecting to whitefall.example.com (whitefall.example.com)|81.xx.xx.141|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 47746472 (46M) [application/x-rpm]
Saving to: 'waypoint-0.4.2-1.x86_64.rpm.1'
```

```
waypoint-0.4.2-1.x86_64.rpm.1      100%[=====]
=====>] 45.53M  107MB/s   in 0.4s
```

2021-08-04 21:13:36 (107 MB/s) - 'waypoint-0.4.2-1.x86\_64.rpm.1' saved [47746472/47746472]

```
[root@alma-canary ~]# yum install waypoint -y
Stable (EL8)
```

```
32 kB/s | 1.7 kB    00:00
```

Dependencies resolved.

Package	Architecture	Size	Version	Repository
<b>Installing:</b>				
waypoint	x86_64	46 M	0.4.2-1	sihnon
_hashicorp_stable-el8				

Transaction Summary

Install 1 Package

Total download size: 46 M

Installed size: 100 M

Downloading Packages:

```
waypoint-0.4.2-1.x86_64.rpm
63 MB/s | 46 MB    00:00
```

```
Total
63 MB/s | 46 MB    00:00
```

Running transaction check

Transaction check succeeded.

Running transaction test

Transaction test succeeded.

Running transaction

```
Preparing      :
Installing     : waypoint-0.4.2-1.x86_64 1/1
Verifying      : waypoint-0.4.2-1.x86_64 1/1
1/1
```

Installed products updated.

Installed:

```
waypoint-0.4.2-1.x86_64
```

Complete!

```
[root@alma-canary ~]# yum install boundary -y
Stable (EL8)
```

```
49 kB/s | 1.7 kB    00:00
```

Dependencies resolved.

Package	Architecture	Version	Repository
boundary	x86_64	0.4.0-1	sihnon

Transaction Summary

Install 1 Package

Total download size: 18 M

Installed size: 57 M

Downloading Packages:

```
[MIRROR] boundary-0.4.0-1.x86_64.rpm: Interrupted by header callback: Server reports Content-Length: 18515599 but expected size is: 18519432
[MIRROR] boundary-0.4.0-1.x86_64.rpm: Interrupted by header callback: Server reports Content-Length: 18515599 but expected size is: 18519432
[MIRROR] boundary-0.4.0-1.x86_64.rpm: Interrupted by header callback: Server reports Content-Length: 18515599 but expected size is: 18519432
[MIRROR] boundary-0.4.0-1.x86_64.rpm: Interrupted by header callback: Server reports Content-Length: 18515599 but expected size is: 18519432
[FAILED] boundary-0.4.0-1.x86_64.rpm: No more mirrors to try - All mirrors were already tried without success
```

The downloaded packages were saved in cache until the next successful transaction.

You can remove cached packages by executing 'yum clean packages'.

Error: Error downloading packages:

Cannot download Packages/b/boundary-0.4.0-1.x86\_64.rpm: All mirrors were tried

```
[root@alma-canary ~]# yum install boundary -y
```

Stable (EL8)

```
6.8 kB/s | 1.7 kB 00:00
```

Dependencies resolved.

Package	Architecture	Version	Repository
boundary	x86_64	0.4.0-1	sihnon

Transaction Summary

Install 1 Package

Total download size: 18 M

Installed size: 57 M

Downloading Packages:

```
[MIRROR] boundary-0.4.0-1.x86_64.rpm: Interrupted by header callback: Server reports Content-Length: 18515599 but expected size is: 18519432
[MIRROR] boundary-0.4.0-1.x86_64.rpm: Interrupted by header callback: Server reports Content-Length: 18515599 but expected size is: 18519432
[MIRROR] boundary-0.4.0-1.x86_64.rpm: Interrupted by header callback: Server reports Content-Length: 18515599 but expected size is: 18519432
[MIRROR] boundary-0.4.0-1.x86_64.rpm: Interrupted by header callback: Server reports Content-Length: 18515599 but expected size is: 18519432
[FAILED] boundary-0.4.0-1.x86_64.rpm: No more mirrors to try - All mirrors were already tried without success
```

The downloaded packages were saved in cache until the next successful transaction.

You can remove cached packages by executing 'yum clean packages'.

Error: Error downloading packages:

```

Cannot download Packages/b/boundary-0.4.0-1.x86_64.rpm: All mirrors were tried
[root@alma-canary ~]# wget --no-check-certificate https://whitefall.example.com/pulp/content/sihnon/Library/custom/hashicorp/stable-el8/Packages/b/boundary-0.4.0-1.x86_64.rpm
--2021-08-04 21:15:36-- https://whitefall.example.com/pulp/content/sihnon/Library/custom/hashicorp/stable-el8/Packages/b/boundary-0.4.0-1.x86_64.rpm
Resolving whitefall.example.com(whitefall.example.com)... 81.xx.xx.141
Connecting to whitefall.example.com (whitefall.example.com)|81.xx.xx.141|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 18515599 (18M) [binary/octet-stream]
Saving to: 'boundary-0.4.0-1.x86_64.rpm'

```

```

boundary-0.4.0-1.x86_64.rpm      100%[=====]
=====>] 17.66M 12.1MB/s   in 1.5s

```

```

2021-08-04 21:15:37 (12.1 MB/s) - 'boundary-0.4.0-1.x86_64.rpm' saved [18515599/18515599]

```

```

[root@alma-canary ~]# yum install boundary -y
Stable (EL8)

```

```

6.8 kB/s | 1.7 kB    00:00

```

```

Dependencies resolved.

```

```

=====
Package                        Architecture      Version           Reposit
tory                            Size
=====
Installing:
boundary                        x86_64           0.4.0-1          sihnou
_hashicorp_stable-el8         18 M

```

```

Transaction Summary

```

```

Install 1 Package

```

```

Total download size: 18 M

```

```

Installed size: 57 M

```

```

Downloading Packages:

```

```

boundary-0.4.0-1.x86_64.rpm
59 MB/s | 18 MB    00:00

```

```

-----
Total
59 MB/s | 18 MB    00:00

```

```

Running transaction check

```

```

Transaction check succeeded.

```

```

Running transaction test

```

```

Transaction test succeeded.

```

```

Running transaction

```

```

Preparing      :
Installing     : boundary-0.4.0-1.x86_64 1/1
Verifying      : boundary-0.4.0-1.x86_64 1/1

```

```

Installed products updated.

```

```

Installed:

```

```

boundary-0.4.0-1.x86_64

```

```

Complete!

```

Attempting to download the file using chrome yields a "Failed - Network Error", until the file has been downloaded successfully via wget, and then the browser download completes successfully.

Another example, downloading vault on a PC with a different OS and more recent version of wget in quick succession yields different

sizes on the two downloads.

```
$ wget --no-check-certificate https://whitefall.example.com/pulp/content/sihnon/Library/custom/hashicorp/stable-el8/Packages/v/vault-1.8.0-1.x86_64.rpm
--2021-08-04 21:26:12-- https://whitefall.example.com/pulp/content/sihnon/Library/custom/hashicorp/stable-el8/Packages/v/vault-1.8.0-1.x86_64.rpm
HTTP request sent, awaiting response... 200 OK
Length: 52581251 (50M) [binary/octet-stream]
Saving to: 'vault-1.8.0-1.x86_64.rpm'
```

```
vault-1.8.0-1.x86_64.rpm          100%[=====]
=====>] 50.14M  11.3MB/s   in 4.3s
```

2021-08-04 21:26:16 (11.8 MB/s) - 'vault-1.8.0-1.x86\_64.rpm' saved [52581251/52581251]

```
$ wget --no-check-certificate https://whitefall.example.com/pulp/content/sihnon/Library/custom/hashicorp/stable-el8/Packages/v/vault-1.8.0-1.x86_64.rpm
--2021-08-04 21:26:30-- https://whitefall.example.com/pulp/content/sihnon/Library/custom/hashicorp/stable-el8/Packages/v/vault-1.8.0-1.x86_64.rpm
HTTP request sent, awaiting response... 200 OK
Length: 52588104 (50M) [application/x-rpm]
Saving to: 'vault-1.8.0-1.x86_64.rpm.1'
```

```
vault-1.8.0-1.x86_64.rpm.1      100%[=====]
=====>] 50.15M  12.6MB/s   in 4.1s
```

2021-08-04 21:26:35 (12.3 MB/s) - 'vault-1.8.0-1.x86\_64.rpm.1' saved [52588104/52588104]

#### Versions:

```
[root@whitefall ~]# rpm -qa '*pulp*'
rubygem-pulp_rpm_client-3.13.3-1.el8.noarch
python3-pulp-file-1.8.1-1.el8.noarch
python3-pulp-rpm-3.14.0-1.el8.noarch
rubygem-pulp_container_client-2.7.0-1.el8.noarch
python3-pulp-cli-0.10.1-1.el8.noarch
rubygem-pulp_deb_client-2.13.0-1.el8.noarch
python3-pulp-container-2.7.0-1.el8.noarch
python3-pulpcore-3.14.3-1.el8.noarch
pulpcore-selinux-1.2.4-1.el8.x86_64
rubygem-pulp_file_client-1.8.1-1.el8.noarch
rubygem-pulp_ansible_client-0.8.0-1.el8.noarch
python3-pulp-certguard-1.4.0-1.el8.noarch
pulp-client-1.0-1.noarch
rubygem-pulp_certguard_client-1.4.0-1.el8.noarch
python3-pulp-ansible-0.8.0-1.el8.noarch
rubygem-pulpcore_client-3.14.1-1.el8.noarch
rubygem-pulp_python_client-3.4.0-1.el8.noarch
rubygem-smart_proxy_pulp-3.1.0-1.fm2_6.el8.noarch
python3-pulp-deb-2.13.0-1.el8.noarch

[root@whitefall ~]# rpm -qa '*katello*'
katello-4.2.0-0.1.master.el8.noarch
katello-debug-4.2.0-0.1.master.el8.noarch
rubygem-katello-4.2.0-0.4.pre.master.20210723175601git5f23a60.el8.noarch
katello-repos-4.2-0.2.nightly.el8.noarch
katello-server-ca-1.0-1.noarch
katello-ca-consumer-whitefall.example.com-1.0-1.noarch
katello-common-4.2.0-0.1.master.el8.noarch
katello-certs-tools-2.8.0-1.el8.noarch
foreman-installer-katello-3.0.0-0.1.develop.20210723110556gitc4bb060.el8.noarch
katello-selinux-4.0.2-1.el8.noarch
rubygem-hammer_cli_katello-1.1.2-1.el8.noarch
katello-client-bootstrap-1.7.6-1.el8.noarch
katello-default-ca-1.0-1.noarch
```

#### Related issues:

**Associated revisions****Revision 6d05548e - 09/01/2021 05:15 PM - dkliban@redhat.com**

Don't send Content-Encoding header with streamed responses

aioshttp automatically enflates gzipped responses. Pulp clients always receive uncompressed responses when requesting on\_demand content.

fixes: #9213

**History****#1 - 08/05/2021 05:32 PM - dalley**

- *Triaged changed from No to Yes*

- *Sprint set to Sprint 102*

**#2 - 08/12/2021 05:23 PM - rchan**

- *Sprint changed from Sprint 102 to Sprint 103*

**#3 - 08/22/2021 05:24 AM - dalley**

- *Status changed from NEW to ASSIGNED*

- *Assignee set to dalley*

**#4 - 08/22/2021 06:13 AM - dalley**

@optiz0r I have a vague hypothesis about what might be happening here, but I can't use the Hashicorp repos to test it. Could you use curl or httpie to print out the full HTTP response headers you see when you download the files from Hashicorp's server, and also from Pulp?

**#5 - 08/24/2021 07:26 PM - dalley**

- *Project changed from RPM Support to Pulp*

**#6 - 08/24/2021 09:52 PM - optiz0r**

Sorry, didn't see your last message, I don't appear to have email notifications working for plan.io despite it being configured to send notifications for any ticket I contribute to.

Curl downloading from the upstream location:

```
# F=consul-1.6.0-1.x86_64.rpm; curl -v -o $F https://rpm.releases.hashicorp.com/RHEL/8/x86_64/stable/$F && ls
-l $F
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload   Total   Spent    Left   Speed
  0     0     0     0     0     0     0     0  --:--:-- --:--:-- --:--:--    0*   Trying 2a04:4e42:200::561...
* TCP_NODELAY set
* connect to 2a04:4e42:200::561 port 443 failed: Permission denied
*   Trying 2a04:4e42:600::561...
* TCP_NODELAY set
* connect to 2a04:4e42:600::561 port 443 failed: Permission denied
*   Trying 2a04:4e42::561...
* TCP_NODELAY set
* connect to 2a04:4e42::561 port 443 failed: Permission denied
*   Trying 2a04:4e42:400::561...
* TCP_NODELAY set
* connect to 2a04:4e42:400::561 port 443 failed: Permission denied
*   Trying 151.101.66.49...
* TCP_NODELAY set
* Connected to rpm.releases.hashicorp.com (151.101.66.49) port 443 (#0)
* ALPN, offering h2
* ALPN, offering http/1.1
* successfully set certificate verify locations:
*   CAfile: /etc/pki/tls/certs/ca-bundle.crt
   CAspath: none
} [5 bytes data]
* TLSv1.3 (OUT), TLS handshake, Client hello (1):
} [512 bytes data]
* TLSv1.3 (IN), TLS handshake, Server hello (2):
{ [106 bytes data]
* TLSv1.2 (IN), TLS handshake, Certificate (11):
{ [2862 bytes data]
```

```

* TLSv1.2 (IN), TLS handshake, Server key exchange (12):
{ [300 bytes data]
* TLSv1.2 (IN), TLS handshake, Server finished (14):
{ [4 bytes data]
* TLSv1.2 (OUT), TLS handshake, Client key exchange (16):
} [37 bytes data]
* TLSv1.2 (OUT), TLS change cipher, Change cipher spec (1):
} [1 bytes data]
* TLSv1.2 (OUT), TLS handshake, Finished (20):
} [16 bytes data]
* TLSv1.2 (IN), TLS handshake, Finished (20):
{ [16 bytes data]
* SSL connection using TLSv1.2 / ECDHE-RSA-AES128-GCM-SHA256
* ALPN, server accepted to use h2
* Server certificate:
* subject: CN=*.releases.hashicorp.com
* start date: May  3 19:03:11 2021 GMT
* expire date: Jun  4 19:03:10 2022 GMT
* subjectAltName: host "rpm.releases.hashicorp.com" matched cert's "*.releases.hashicorp.com"
* issuer: C=BE; O=GlobalSign nv-sa; CN=GlobalSign Atlas R3 DV TLS CA 2020
* SSL certificate verify ok.
* Using HTTP2, server supports multi-use
* Connection state changed (HTTP/2 confirmed)
* Copying HTTP/2 data in stream buffer to connection buffer after upgrade: len=0
} [5 bytes data]
* Using Stream ID: 1 (easy handle 0x55f1ae8bf4c0)
} [5 bytes data]
> GET /RHEL/8/x86_64/stable/consul-1.6.0-1.x86_64.rpm HTTP/2
> Host: rpm.releases.hashicorp.com
> User-Agent: curl/7.61.1
> Accept: */*
>
{ [5 bytes data]
* Connection state changed (MAX_CONCURRENT_STREAMS == 100)!
} [5 bytes data]
 0    0    0    0    0    0    0    0    0  --:--:--  --:--:--  --:--:--    0< HTTP/2 200
< x-amz-id-2: 3ErbRwV9Kh8H0AAriY55ZM7VFGyDFz/fa+OkXO0byMOQFEpGhTzONAKUO1qruA832ZkgK0rVFh4=
< last-modified: Mon, 20 Jul 2020 06:18:57 GMT
< etag: "a8aeda452065918e68173c3f0fbf7fe5"
< content-type: binary/octet-stream
< server: AmazonS3
< accept-ranges: bytes
< age: 1
< date: Tue, 24 Aug 2021 19:39:15 GMT
< via: 1.1 varnish
< x-served-by: cache-lon4259-LON
< x-cache: MISS
< x-cache-hits: 0
< vary: Accept-Encoding
< content-length: 31755252
<
{ [1088 bytes data]
100 30.2M 100 30.2M  0    0 9913k  0 0:00:03 0:00:03 --:--:-- 9910k
* Connection #0 to host rpm.releases.hashicorp.com left intact

```

```
-rw-r--r--. 1 root root 31755252 Aug 24 19:39 consul-1.6.0-1.x86_64.rpm
```

And from Katello/pulp the first time:

```

# F=consul-1.6.0-1.x86_64.rpm; curl -v -o $F https://whitefall.jellybean.sihnon.net/pulp/content/sihnon/Library/custom/hashicorp/stable-el8/$F && ls -l $F
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload  Total   Spent    Left     Speed
 0     0    0     0    0     0     0     0  --:--:--  --:--:--  --:--:--    0*   Trying 2001:8b0:3b3:0:250:56ff:fe9a:9f35...
* TCP_NODELAY set
* connect to 2001:8b0:3b3:0:250:56ff:fe9a:9f35 port 443 failed: Permission denied
*   Trying 81.187.154.141...
* TCP_NODELAY set
* Connected to whitefall.jellybean.sihnon.net (81.187.154.141) port 443 (#0)
* ALPN, offering h2
* ALPN, offering http/1.1
* successfully set certificate verify locations:
*   CAfile: /etc/pki/tls/certs/ca-bundle.crt

```



```

CApath: none
} [5 bytes data]
* TLSv1.3 (OUT), TLS handshake, Client hello (1):
} [512 bytes data]
* TLSv1.3 (IN), TLS handshake, Server hello (2):
{ [122 bytes data]
* TLSv1.3 (IN), TLS handshake, [no content] (0):
{ [1 bytes data]
* TLSv1.3 (IN), TLS handshake, Encrypted Extensions (8):
{ [25 bytes data]
* TLSv1.3 (IN), TLS handshake, [no content] (0):
{ [1 bytes data]
* TLSv1.3 (IN), TLS handshake, Request CERT (13):
{ [197 bytes data]
* TLSv1.3 (IN), TLS handshake, [no content] (0):
{ [1 bytes data]
* TLSv1.3 (IN), TLS handshake, Certificate (11):
{ [3697 bytes data]
* TLSv1.3 (IN), TLS handshake, [no content] (0):
{ [1 bytes data]
* TLSv1.3 (IN), TLS handshake, CERT verify (15):
{ [520 bytes data]
* TLSv1.3 (IN), TLS handshake, [no content] (0):
{ [1 bytes data]
* TLSv1.3 (IN), TLS handshake, Finished (20):
{ [52 bytes data]
* TLSv1.3 (OUT), TLS change cipher, Change cipher spec (1):
} [1 bytes data]
* TLSv1.3 (OUT), TLS handshake, [no content] (0):
} [1 bytes data]
* TLSv1.3 (OUT), TLS handshake, Certificate (11):
} [8 bytes data]
* TLSv1.3 (OUT), TLS handshake, [no content] (0):
} [1 bytes data]
* TLSv1.3 (OUT), TLS handshake, Finished (20):
} [52 bytes data]
* SSL connection using TLSv1.3 / TLS_AES_256_GCM_SHA384
* ALPN, server accepted to use http/1.1
* Server certificate:
*  subject: C=US; ST=North Carolina; O=Katello; OU=SomeOrgUnit; CN=whitefall.jellybean.sihnon.net
*  start date: Jul  4 19:43:32 2021 GMT
*  expire date: Jan 17 19:43:32 2038 GMT
*  subjectAltName: host "whitefall.jellybean.sihnon.net" matched cert's "whitefall.jellybean.sihnon.net"
*  issuer: C=US; ST=North Carolina; L=Raleigh; O=Katello; OU=SomeOrgUnit; CN=whitefall.jellybean.sihnon.net
*  SSL certificate verify ok.
} [5 bytes data]
* TLSv1.3 (OUT), TLS app data, [no content] (0):
} [1 bytes data]
> GET /pulp/content/sihnon/Library/custom/hashicorp/stable-el8/consul-1.6.0-1.x86_64.rpm HTTP/1.1
> Host: whitefall.jellybean.sihnon.net
> User-Agent: curl/7.61.1
> Accept: */*
>
{ [5 bytes data]
* TLSv1.3 (IN), TLS handshake, [no content] (0):
{ [1 bytes data]
* TLSv1.3 (IN), TLS handshake, Newsession Ticket (4):
{ [297 bytes data]
* TLSv1.3 (IN), TLS handshake, [no content] (0):
{ [1 bytes data]
* TLSv1.3 (IN), TLS handshake, Newsession Ticket (4):
{ [297 bytes data]
0 0 0 0 0 0 0 0 0 --:--:-- 0:00:02 --:--:-- 0* TLSv1.3 (IN), TLS app data, [n
o content] (0):
{ [1 bytes data]
< HTTP/1.1 200 OK
< Date: Tue, 24 Aug 2021 19:40:19 GMT
< Server: AmazonS3
< Content-Type: binary/octet-stream
< Content-Length: 31752344
< x-amz-id-2: 6Ajm/lq2FhzmZJPDGTfUL1CUDAHymUU3D0TUfrSvi4c62aENMt foahDxVCEdRy3jUd3F73d2f6A=
< Last-Modified: Mon, 20 Jul 2020 06:18:57 GMT
< Etag: "a8aeda452065918e68173c3f0fbf7fe5"
< Content-Encoding: gzip
< Accept-Ranges: bytes

```

```

< Age: 2
< Via: 1.1 varnish
< X-Served-By: cache-lcy19228-LCY
< X-Cache: MISS
< X-Cache-Hits: 0
< Vary: Accept-Encoding
< Via: 1.1 whitefall.jellybean.sihnon.net
<
  0 30.2M  0  0  0  0  0  0  0  --:--:--  0:00:04  --:--:--  0{ [5 bytes data]
* TLSv1.3 (IN), TLS app data, [no content] (0):
{ [1 bytes data]
* TLSv1.3 (IN), TLS app data, [no content] (0):
{ [1 bytes data]
...
* TLSv1.3 (IN), TLS app data, [no content] (0):
{ [1 bytes data]
* TLSv1.3 (IN), TLS app data, [no content] (0):
{ [1 bytes data]
100 30.2M 100 30.2M  0  0 4518k  0 0:00:06 0:00:06 --:--:-- 4518k
* Connection #0 to host whitefall.jellybean.sihnon.net left intact

```

```
-rw-r--r--. 1 root root 31752344 Aug 24 19:40 consul-1.6.0-1.x86_64.rpm
```

And a second time:

```

# F=consul-1.6.0-1.x86_64.rpm; curl -v -o $F https://whitefall.jellybean.sihnon.net/pulp/content/sihnon/Library/custom/hashicorp/stable-el8/$F && ls -l $F
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
           %             0         0     0    0  --:--:--  --:--:--  --:--:--  0*   Trying 2001:8b0:3b3:0:250:56ff:fe9a:9f35...
ff:fe9a:9f35...
* TCP_NODELAY set
* connect to 2001:8b0:3b3:0:250:56ff:fe9a:9f35 port 443 failed: Permission denied
*   Trying 81.187.154.141...
* TCP_NODELAY set
* Connected to whitefall.jellybean.sihnon.net (81.187.154.141) port 443 (#0)
* ALPN, offering h2
* ALPN, offering http/1.1
* successfully set certificate verify locations:
*   CAfile: /etc/pki/tls/certs/ca-bundle.crt
  CApath: none
} [5 bytes data]
* TLSv1.3 (OUT), TLS handshake, Client hello (1):
} [512 bytes data]
* TLSv1.3 (IN), TLS handshake, Server hello (2):
{ [122 bytes data]
* TLSv1.3 (IN), TLS handshake, [no content] (0):
{ [1 bytes data]
* TLSv1.3 (IN), TLS handshake, Encrypted Extensions (8):
{ [25 bytes data]
* TLSv1.3 (IN), TLS handshake, [no content] (0):
{ [1 bytes data]
* TLSv1.3 (IN), TLS handshake, Request CERT (13):
{ [197 bytes data]
* TLSv1.3 (IN), TLS handshake, [no content] (0):
{ [1 bytes data]
* TLSv1.3 (IN), TLS handshake, Certificate (11):
{ [3697 bytes data]
* TLSv1.3 (IN), TLS handshake, [no content] (0):
{ [1 bytes data]
* TLSv1.3 (IN), TLS handshake, CERT verify (15):
{ [520 bytes data]
* TLSv1.3 (IN), TLS handshake, [no content] (0):
{ [1 bytes data]
* TLSv1.3 (IN), TLS handshake, Finished (20):
{ [52 bytes data]
* TLSv1.3 (OUT), TLS change cipher, Change cipher spec (1):
} [1 bytes data]
* TLSv1.3 (OUT), TLS handshake, [no content] (0):
} [1 bytes data]
* TLSv1.3 (OUT), TLS handshake, Certificate (11):
} [8 bytes data]
* TLSv1.3 (OUT), TLS handshake, [no content] (0):
} [1 bytes data]

```

```

* TLSv1.3 (OUT), TLS handshake, Finished (20):
} [52 bytes data]
* SSL connection using TLSv1.3 / TLS_AES_256_GCM_SHA384
* ALPN, server accepted to use http/1.1
* Server certificate:
*  subject: C=US; ST=North Carolina; O=Katello; OU=SomeOrgUnit; CN=whitefall.jellybean.sihnon.net
*  start date: Jul  4 19:43:32 2021 GMT
*  expire date: Jan 17 19:43:32 2038 GMT
*  subjectAltName: host "whitefall.jellybean.sihnon.net" matched cert's "whitefall.jellybean.sihnon.net"
*  issuer: C=US; ST=North Carolina; L=Raleigh; O=Katello; OU=SomeOrgUnit; CN=whitefall.jellybean.sihnon.net
*  SSL certificate verify ok.
} [5 bytes data]
* TLSv1.3 (OUT), TLS app data, [no content] (0):
} [1 bytes data]
> GET /pulp/content/sihnon/Library/custom/hashicorp/stable-el8/consul-1.6.0-1.x86_64.rpm HTTP/1.1
> Host: whitefall.jellybean.sihnon.net
> User-Agent: curl/7.61.1
> Accept: */*
>
{ [5 bytes data]
* TLSv1.3 (IN), TLS handshake, [no content] (0):
{ [1 bytes data]
* TLSv1.3 (IN), TLS handshake, Newsession Ticket (4):
{ [297 bytes data]
* TLSv1.3 (IN), TLS handshake, [no content] (0):
{ [1 bytes data]
* TLSv1.3 (IN), TLS handshake, Newsession Ticket (4):
{ [297 bytes data]
* TLSv1.3 (IN), TLS app data, [no content] (0):
{ [1 bytes data]
< HTTP/1.1 200 OK
< Date: Tue, 24 Aug 2021 19:42:14 GMT
< Server: Python/3.6 aiohttp/3.7.4
< Content-Type: application/x-rpm
< Last-Modified: Tue, 24 Aug 2021 19:40:23 GMT
< Content-Length: 31755252
< Accept-Ranges: bytes
< Via: 1.1 whitefall.jellybean.sihnon.net
<
{ [5 bytes data]
* TLSv1.3 (IN), TLS app data, [no content] (0):
{ [1 bytes data]
...
* TLSv1.3 (IN), TLS app data, [no content] (0):
{ [1 bytes data]
* TLSv1.3 (IN), TLS app data, [no content] (0):
{ [1 bytes data]
100 30.2M 100 30.2M 0 0 31.1M 0 ---:--:-- --:--:-- --:--:-- 31.0M
* Connection #0 to host whitefall.jellybean.sihnon.net left intact

```

```
-rw-r--r--. 1 root root 31755252 Aug 24 19:42 consul-1.6.0-1.x86_64.rpm
```

Once curl is retrieving the file with the correct size that matches upstream, the dnf install works. Curiously, repeated downloads via DNF doesn't seem to "kick" katello/pulp into serving up the right content. But using curl, typically the second download attempt does seem to trigger the correct download. I think I've seen it take more than two attempts with curl once in the past, but wasn't recording my actions that day.

(aside: katello/pulp appears to be serving up the download one byte at a time? I omitted many thousands of repeated lines showing 1 byte of data read from the two katello downloads)

#### #7 - 08/24/2021 10:47 PM - dalley

I think I uncovered at least part of this issue while investigating another issue. My theory is that a mistake is being made in one or both of these places:

- Some http client requests a file from Pulp
- Pulp sees this and requests the file from the remote server
  - ?? Pulp may provide an "Accept-Encoding" header that the original request did not ??
- The remote server starts returning data that is compressed and uses Content-Encoding headers to tell the client how to decompress it
  - ?? Pulp gets the data, decompresses it in order to save to artifact storage, and sends the (decompressed!!) bytes to the client while maintaining the Content-Encoding header ??
- The client gets data that wasn't gzip encoded when the headers claim that it was, or gets data that was gzip encoded when it wasn't expecting it
- The client throws a fit when it starts getting the incorrect bytes and/or Content-Length and cancels the whole request, so Pulp never completes downloading the file \*\*, and the cycle just repeats

And wget meanwhile powers through the error, gets Pulp to download the entire thing, at which point the header-mismatch issues go away.

\*\* I'm not actually sure why it doesn't do that, it probably should

**#8 - 08/27/2021 05:08 PM - rchan**

- Sprint changed from Sprint 103 to Sprint 104

**#9 - 08/27/2021 06:00 PM - dalley**

- Status changed from ASSIGNED to NEW

- Assignee deleted (dalley)

- Severity changed from 2. Medium to 3. High

**#10 - 08/31/2021 10:49 PM - dkliban@redhat.com**

- Status changed from NEW to ASSIGNED

- Assignee set to dkliban@redhat.com

**#11 - 08/31/2021 11:04 PM - dalley**

- Copied to Backport #9325: Backport #9213 "DNF gets wrong rpm download size, until package downloaded with curl, then works" to 3.14.z added

**#12 - 08/31/2021 11:10 PM - pulpbot**

- Status changed from ASSIGNED to POST

PR: <https://github.com/pulp/pulpcore/pull/1601>

**#13 - 09/01/2021 05:15 PM - dkliban@redhat.com**

- Status changed from POST to MODIFIED

Applied in changeset [pulpcore|6d05548e9c4bc41e56c34c9ab4d0aa73b7eb6b28](#).

**#14 - 09/01/2021 05:48 PM - dalley**

- Sprint/Milestone set to 3.16.0

**#15 - 10/05/2021 08:53 PM - pulpbot**

- Status changed from MODIFIED to CLOSED - CURRENTRELEASE