

MirrorManager and CentOS Stream 9

Adrian Reber

2022, February 03

About myself

- ▶ Running a mirror server since 1998
- ▶ Involved in Fedora (fedora.us) for almost 20 years
- ▶ First MirrorManager patches in 2008
- ▶ Main point of contact for MirrorManager since 2016
- ▶ Red Hat Core Kernel Team
- ▶ Working on checkpoint and restore (CRIU)

Agenda

History

`dnf install`

Fedora setup

Statistics

Recent changes

Known Problems

MirrorManager in use since 2007

Fedora and EPEL

RPM Fusion

Rocky Linux

CentOS Stream 9

What happens during `dnf install`

Download metalink

`https://mirrors.centos.org/metalink?repo=centos-baseos-9-stream&arch=x86_64`

```
1 <files>
2   <file name="repomd.xml">
3     <mm0:timestamp>1642386319</mm0:timestamp>
4     <size>4843</size>
5     <verification>
6       <hash type="md5">304051e1daed2d...bfff15ba781c</hash>
7       <hash type="sha1">be5753534d8df1a8...224bc4a892ad</hash>
8       <hash type="sha256">8e3f7dc1a2f9640b15686db...b1a0b8</hash>
9       <hash type="sha512">3d2b429dc31310f41135220...d31c0520b5d</hash>
10    </verification>
11    <mm0:alternates>
12      <mm0:alternate>
13        <mm0:timestamp>1642298581</mm0:timestamp>
14        <size>4844</size>
15        <verification>
16          <hash type="md5">6101a2d6f2e652...0aed60036</hash>
17          <hash type="sha1">3badc0625b7a47b2...72b93992ed6b</hash>
18          <hash type="sha256">392544cbd0a6e88a0...c66c128e65712</hash>
19          <hash type="sha512">78c79ec7aaff6b38b...1c9a2559c89f8d6</hash>
20        </verification>
```

```
1 <resources maxconnections="1">
2   <url protocol="https" type="https" location="DE" preference="100">https://
   mirror.netzwerke.de/centos-stream/9-stream/BaseOS/x86_64/os/repdata/
   repomd.xml</url>
3   <url protocol="https" type="https" location="DE" preference="99">https://
   mirror1.hs-esslingen.de/pub/Mirrors/centos-stream/9-stream/BaseOS/x86_64/
   os/repdata/repomd.xml</url>
4   <url protocol="https" type="https" location="DE" preference="98">https://ftp.
   plusline.net/centos-stream/9-stream/BaseOS/x86_64/os/repdata/repomd.xml</
   url>
5   <url protocol="https" type="https" location="DK" preference="97">https://
   mirror.netsite.dk/centos-stream/9-stream/BaseOS/x86_64/os/repdata/repomd.
   xml</url>
6   <url protocol="https" type="https" location="NL" preference="96">https://
   centos-stream.mirror.wearetriple.com/9-stream/BaseOS/x86_64/os/repdata/
   repomd.xml</url>
7   <url protocol="https" type="https" location="CH" preference="95">https://
   linuxsoft.cern.ch/centos-stream/9-stream/BaseOS/x86_64/os/repdata/repomd.
   xml</url>
```

Download `repomd.xml`

Download repository files

Download actual RPM

Mirror selection process

Take IP from client (DNF)

Get mirrors from same netblock

Private mirrors

Get mirrors from same ASN or neighboring ASN

Get nearby Internet2 mirrors (GEANT in Europe)

Get mirrors from same country (GeoIP)

Get mirrors from same continent

Get global mirrors

Randomize and order by prefix size

Weighted shuffle by bandwidth

Fedora's setup

Proxies (25)

Frontends (3)

Backend

Crawlers (2)

Database

MirrorManager2 since 2015

Proxies on Fedora

We needed Python 3 support

Python 3 support (2017)

Backend Python 2

Proxies Python 3

Switched to Protocol Buffers

protobuf

Introducing Rust

Starting with the Proxies

Introducing Rust

Step two

(protobuf generation)

Introducing Rust

Step three


(scan primary mirror)

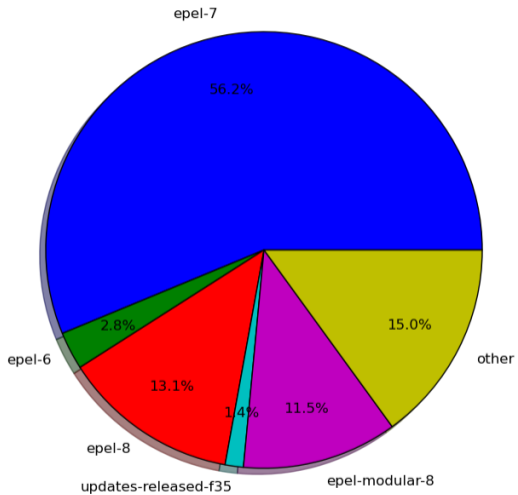
Statistics

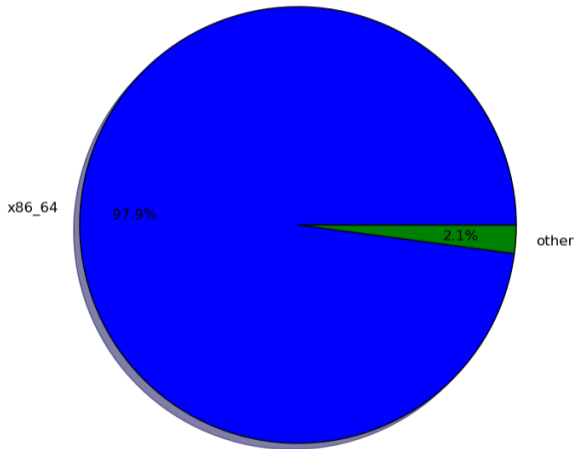
Over 1200 mirrors



Around 68 million requests per day







Known Problems

baseos

appstream

crb

Primary mirror scanning

The End