

To cache or not to cache making pkg_add faster

Marc Espie <espie@openbsd.org>, <espie@lse.epita.fr>



June 3, 2022

This is just a "lightning talk" to give you some of the recent elements. There is more to the story, which will get expanded upon in this year's summer week, hopefully

`pkg_add` is a tool specific to OpenBSD. Historically, we do "just in time" updates.

- open new package and peek at meta information
- decide whether we want to update
- if so, extract the new package, then delete the old one

The meta information

We got structured information (packing-lists) that looks like this:

```
1 @pkgpath x11/dbus
2 @newgroup _dbus:572
3 @newuser _dbus:572:_dbus::dbus user:/nonexistent:/sbin/nologin
4 @extra ${SYSCONFDIR}/machine-id
5 @rcscript ${RCDIR}/messagebus
6 @bin bin/dbus-cleanup-sockets
7 @bin bin/dbus-daemon
8 @bin bin/dbus-launch
9 @bin bin/dbus-monitor
10 @bin bin/dbus-run-session
11 @bin bin/dbus-send
12 @bin bin/dbus-test-tool
13 [... more files]
```

This is the source information for the dbus package, telling us it requires some user/groups, has a service start-up script, and contains a bunch of files

The whole story I

After going through `pkg_create`, the full *packing-list* looks more like

```
1 @name dbus-1.14.0v0
2 @version 8
3 @comment pkgpath=x11/dbus,-main ftp=yes
4 @arch amd64
5 +DESC
6 @sha TYbBC2o07XX0XqnQ0FU6qikEuiN+fqoN2azXrJA9jJg=
7 @size 448
8 @pkgpath x11/dbus
9 @wantlib X11.18.0
10 @wantlib c.96.1
11 @wantlib execinfo.3.0
12 @wantlib expat.14.0
13 @wantlib pthread.26.1
14 @wantlib xcb.4.1
15 @newgroup _dbus:572
16 @newuser _dbus:572:_dbus::dbus user:/nonexistent:/sbin/nologin
17 @cwd /usr/local
```

The whole story II

```
18 @extra /etc/machine-id
19 @rcscript /etc/rc.d/messagebus
20 @sha G8InGF0+lEOiPUMpXqicxP01KEkofH0guRhxV9sMXHk=
21 @size 172
22 @ts 1653570364
23 @bin bin/dbus-cleanup-sockets
24 @sha lew9j03YckJ1VnMPtypbKh1k1eedAXgwCvYU3hE44jU=
25 @size 13318
26 @ts 1653570364
27 [... more files]
```

- a *packing-list* is a structured object which has constructors. Most often it starts as

```
my $plist = OpenBSD::PackingList->from_file("filename");
```
- objects (*packing elements*) can be added to it using the right method:

```
OpenBSD::PackingElement::Wantlib->add($plist, $w);
```
- some complex objects can have a multiline representation, like files:
 - name
 - extra modes
 - checksum
 - timestamp
 - ownership

OO properties

- there's a whole hierarchy of objects: anything file-system related is a FileObject, annotations are Meta, anything depend-related is a Depend
- objects are emitted in a specific order: first all the meta information, then the actual objects (in order)
- most operations happen as visitors on the packing-list
- there are specialized scanners that take advantage of the text structure of the packing-list to avoid reading it all

```
1  sub DependOnly
2  {
3      my ($fh, $cont) = @_;
4      while (<$fh>) {
5          if (m/^\@(?:libset|depend|wantlib|define-tag)\b/o) {
6              &$cont($_);
7              # XXX optimization
8          } elsif (m/^\@(?:newgroup|newuser|cwd)\b/o) {
9              last;
10         }
11     }
```


In order to decide whether to update a package

- we look up all packages that have the same name with a different version number
- we open every package to see whether it's a valid candidate
- we filter the ones we don't want
- pathological case: autoconf. We have a branch for each version, which means 17 packages to consider.
- ... and we decide to update

- for a long time, the network was slow, bandwidth-wise, so opening lots of files was not a big issue
- actually properly closing was an issue with ftp: premature closing requires full telnet support (with "attention" commands")
- and so I had to fix ftp-proxy back in the day
- ... but recently, latency is more of an issue, most people have lots of bandwidth, and so does our current setup

We got a CDN

- ftp is dead, long live http(s).
- establishing connections might be a bit slow
- the cdn first gives you a redirect which means two connections
- we've parsed the redirect from the start, to make sure an update connects to exactly one mirror
- bandwidth is not an issue, latency is

https ? not such a good idea

- http connection establishment is $1 \frac{1}{2}$ RTT
- https is $2 \frac{1}{2}$ RTT at best !
- we did implement session resumption (with fun results)
- ... so updates is slow, because we establish lots of connections
- also, signatures

Signatures ?

```
1  untrusted comment: verify with openbsd-69-pkg.pub
2  RWSG2ib5ZXSfQTrcxxj+A9b6oeFI/OiJVB49nvIs+UPIull+Mk/BclTXRuG4a+XbnyoiZffDILfP58BNelK0yMjZNE
3  date=2021-02-26T23:06:32Z
4  key=/etc/signify/openbsd-69-pkg.sec
5  algorithm=SHA512/256
6  blocksize=65536
7
8  9d61ddfc76218e7c3745bd942a29725ff1bc651f64af27a450da33a73f292d69
9  8621c7932e29c838783177287fc5779186c854b35eaa541e787979f78288c2a6
10 d895cc173cb9058341bbcbe6abe3c018b915eb9218fd65c31f490f9af9c11041
11 9895735d7a109e497ef3f616f35938ae4d6e66f851f038ba50aa2a69808ef53a
12 ce23313490656aaeda9b21aa137a7e70fb268db9372cafeefe860e3fb98c4dfb
13 d34eedc74d714c7a5702b386d36ee422d614d0239cf45e3ae417dd5cd6a09f6f
14 55330726f9221f239c76d4809463ebc251a634360f7098cff98931f8948b7669
15 e84f66e180f1be0c5ef057ea2c4bc74106791b6b794e2de74dc56a9968fa8410
16 e2e4283c81ace8474a32dfc6e43fa3515f02e9bdc93daa86d84875cf9d4ac72d
17 aa1588b1ca21bf13dc132fd12e485cf0edebc787ee53a4cf6df6aa8d5e5e5611
18 9f6723f0419bc16b0a1230407ab3e25015dda27793c424bc50a6ace4f7de4a2e
```

We'd like to store the update info somewhere but

- we don't have any db tools in the base system
- we need to generate and grab it securely from the cdn

- But we have `locate` in the base system.
- it's been designed to store efficiently "similar" strings (by sorting and compressing according to prefix)
- already used for `pkglocatedb`
- this stores each path in packages prefixed by the `pkgname/path` location

```
1  nausicaa$ pkglocate /usr/local/bin/vim
2  graphviz-2.42.3p0:math/graphviz,-main:/usr/local/bin/vimdot
3  vim-8.2.5036-gtk3-lua:editors/vim,-main,gtk3,lua:/usr/local/bin/vim
4  vim-8.2.5036-gtk3-lua:editors/vim,-main,gtk3,lua:/usr/local/bin/vimdiff
5  vim-8.2.5036-gtk3-lua:editors/vim,-main,gtk3,lua:/usr/local/bin/vimtutor
6  vim-8.2.5036-gtk3-perl-python3-ruby:editors/vim,-main,gtk3,perl,python3,ruby:/usr/local/bin/vim
7  vim-8.2.5036-gtk3-perl-python3-ruby:editors/vim,-main,gtk3,perl,python3,ruby:/usr/local/bin/vim
8  vim-8.2.5036-gtk3-perl-python3-ruby:editors/vim,-main,gtk3,perl,python3,ruby:/usr/local/bin/vim
9  vim-8.2.5036-gtk3-python3:editors/vim,-main,gtk3,python3:/usr/local/bin/vim
10 [ ... ]
```


- It is very efficient: 300MB compress to 23MB
- It is fast
- It is in the base system

- generate data with `pkgname:update-info-line`
- this should compress correctly
- where to put it to make this accessible
- I did a script that worked. Compression is okay (compresses 23M to 3M)

- I gave the script to my fellow builders and asked for pkgindex.tgz to be on the mirrors
- they did it for a while, but I got distracted
- and then they no longer did it
- right when I got motivated again

- it had to be on, all the time. Add glue at the end of dpb to generate it ?
- delivery system. Sign it specifically ? teach pkg_add how to read it ?
- scrape that, let's use quirks

- Quirks is the package that holds "Exceptions" to the rules (such as package renames, or packages that got dropped).
- First action of `pkg_add` ever is always to try to update quirks.
- So it's a natural location to drop update info

- so I got the script that builds the db into quirks
- told my friends to always regenerate quirks at the end
- and waited for the new package to show up

- (there was a small issue with "always-update" packages, let's avoid them)
- try to grab the updateinfo from the locate before going to the packages
- result **over twenty times speed-up**
- so worth making it work

- the db is linked to a given quirks, which means a given package repository.
- this is not a big issue because we got unique objects for repositories
- furthermore, quirks is an "always-update" package, so if we find we don't need to update it, it means the quirks we got contains update info for our packages
- we can actually put that in production !

- we run a separate locate for each updateinfo
- we can actually run a single locate upfront, because we got the list of pkgnames we want to handle

```
1 sub prime_update_info_cache
2 {
3     my ($self, $state, $setlist) = @_;
4
5     my $progress = $state->progress;
6     my $found = {};
7
8     my $pseudo_search = [$self];
9
10    for my $set (@{$setlist}) {
11        for my $h ($set->older, $set->hints) {
12            next if $h->{update_found};
13            my $name = $h->pkgname;
14            my $stem = OpenBSD::PackageName::splitstem($name);
```

```
15         next if $stem =~ m/^\.libs\d*\-/;
16         next if $stem =~ m/^\partial\-/;
17         $stem =~ s/\%.*//; # zap branch info
18         $stem =~ s/\-\-.*//; # and set flavors
19         $self->add_stem($stem);
20     }
21 }
22 my @list = sort keys %{$self->{stems}};
23 return if @list == 0;
24
25 my $total = scalar @list;
26 $progress->set_header(
27     $state->f("Reading update info for installed packages",
28         $total));
29 my $done = 0;
30 my $oldname = "";
31
32 open my $fh, "-|", $self->pipe_locate(map { "$_-[0-9]*" } @list)
```

```
33     or $state->fatal("Can't run locate: #1", $!);
34 while (<$fh>) {
35     if (m/^(.*?)\:(.*)/) {
36         my ($pkgname, $value) = ($1, $2);
37         $found->{OpenBSD::PackageName::splitstem($pkgname)} = 1;
38         $self->{raw_data}{$pkgname} //= '';
39         $self->{raw_data}{$pkgname} .= "$value\n";
40         if ($pkgname ne $oldname) {
41             $oldname = $pkgname;
42             $done++;
43         }
44         $progress->show($done, $total);
45     }
46 }
47 close($fh);
48 return unless $state->defines("CACHING_VERBOSE");
49 for my $k (@list) {
50     if (!defined $found->{$k}) {
```

```
51         $state->say("No cache entry for #1", $k);
52     }
53 }
54 }
```

- at first those were not handled at all
- it means a package that needs an update each time it changes
- after a few tries, I decided that storing a crypto hash would work
- so now it is `@option always-update <hash value>`
- and `pkg_create` generates it

Roughly ten lines in dpb:

```
1  if ($state->{all}) {
2      my $core = DPB::Core->get;
3      my $w = DPB::PkgPath->new('devel/quirks');
4      if ($state->{engine}{built_packages}) {
5          $state->grabber->clean_packages($core, $w->fullpkgpath);
6      }
7      my $subdirlist = {};
8      $w->add_to_subdirlist($subdirlist);
9      $state->grabber->grab_subdirs($core, $subdirlist, undef);
10     $state->engine->check_buildable;
11     $core->mark_ready;
12     main_loop();
13 }
```

that's all folks

Questions ? more fun details in the summer week