


Advanced ports toolkit: near-perfect packing-list generation

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



Epita Research and
Development Laboratory

September 22, 2019



- llvm is quietly changing its software license for 9.0 onwards.
- it's going from BSD (perfect) to Apache V2 (waaaay less free)
- it's probably too late to do anything but at least you know.

- We're supposed to be BSD
- A big reason (besides technical excellence) is that we believe in the license
- We already got scammed big-time by the FSF with the GPLv3 toolchain
- And now, we are ready to accept Apache V2, which is also a compromise
- let's not be pushovers and be more vocal about licensing
- it's too late for llvm.

Binary packages

- we inherited `bsd.port.mk` from FreeBSD
- we converted to "staged installation" (we call that `fake`) 20 years ago
- went from `build` → `install` → `package` to `build` → `package` → `install`

Consequences

- Reproducible installation because we have every file
- Need some kind of MANIFEST for stuff that goes into packages
- ... generated with ad-hoc tools

make plist

The original `make_plist` predates staging areas:

- 150 lines of perl that looks under `/usr/local` for new files.
- Very simple and stupid.
- support for manual calls to `ldconfig`, `install-info`
- uses `mtree` to avoid recreating directory hierarchies
- gets much faster with staging area

pkg create

And `pkg_create` gets its plist from a sed script Horrible stuff like:

```
sed -e '/^!%%${_i}%%$$/d' -e '/^%%${_i}%%$$/r${PKGDIR}/PFRAG.${_i}'
```

The package tools were replaced with a perl ABI and perl tools

```
1  # This is the basic class, which is mostly abstract, except for
2  # create and register_with_factory.
3  # It does provide base methods for stuff under it, though.
4  package OpenBSD::PackingElement;
5
6  # concrete objects
7  package OpenBSD::PackingElement::Object;
8  our @ISA=qw(OpenBSD::PackingElement);
9
10 sub cwd
11 {
12     return ${$_[0]}->{cwd};
13 }
14
15 sub absolute_okay() { 0 }
```

More features II

```
16 sub compute_fullname
17 {
18     my ($self, $state) = @_;
19
20     $self->{cwd} = $state->{cwd};
21     $self->set_name(File::Spec->canonpath($self->name));
22     if ($self->name =~ m|^/|) {
23         unless ($self->absolute_okay) {
24             die "Absolute name forbidden: ", $self->name;
25         }
26     }
27 }
28
29 sub make_full
30 {
31     my ($self, $path) = @_;
32     if ($path !~ m|^/|o && $self->cwd ne '.') {
```

More features III

```
33         $path = $self->cwd."/".$path;
34         $path =~ s,^//,/,;
35     }
36     return $path;
37 }
38 [...]
39 # concrete objects with file-like behavior
40 package OpenBSD::PackingElement::FileObject;
41 our @ISA=qw(OpenBSD::PackingElement::Object);
42 # exec/unexec and friends
43 [...]
44 package OpenBSD::PackingElement::Action;
45 our @ISA=qw(OpenBSD::PackingElement::Object);
46 [...]
47 package OpenBSD::PackingElement::Lib;
48 our @ISA=qw(OpenBSD::PackingElement::FileBase);
49
```


More features IV

```
50 sub keyword() { "lib" }
51
52 sub mark_ldconfig_directory
53 {
54     my ($self, $state) = @_;
55     $state->ldconfig->mark_directory($self->fullname);
56 }
```

So make-plist was updated

It became able to read and generate more information... and grew to 300 lines.
Copying annotations like @mode, @owner, and file types like @lib.

Motivation

- Flavors are very simple, they're just a way to record compilation options into packages.
- This however results in more fragments... and variables.
- So `pkg_create` learnt to use them directly, exit the sed stuff

Enter FLAVORS and MULTI PACKAGES I

```
@comment $OpenBSD: PLIST,v 1.2 2019/05/18 15:00:41 espie Exp $
@option no-default-conflict
@option is-branch
@conflict nethack->=${BASEV}
@newgroup ${GAMEGRP}:806
@comment bin/license
@comment bin/nethack
bin/nethack-${V}
bin/recover-${V}
@mode 0775
@group ${GAMEGRP}
${GAMEDIR}/
@mode
${GAMEDIR}/license
@mode 2555
@bin ${GAMEDIR}/nethack
@bin ${GAMEDIR}/recover
@mode
${GAMEDIR}/nhdat
@mode 0664
@comment no checksum
${GAMEDIR}/logfile
${GAMEDIR}/perm
@comment no checksum
${GAMEDIR}/record
@mode 0664
@comment no checksum
```

Enter FLAVORS and MULTI PACKAGES II

```

${GAMEDIR}/xlogfile
@mode 0774
${GAMEDIR}/save/
@mode
${GAMEDIR}/symbols
${GAMEDIR}/sysconf
@man man/man6/nethack-${V}.6
@man man/man6/recover-${V}.6
!%%no_x11%%

```

Motivation

- Compile once, package several times.
- No need to hack at makefiles/configure to just build part of some software.
- Create subpackages based on dependencies and/or sizes.
- ... but how do we sort files into the right subpackage

Enter FLAVORS and MULTI PACKAGES I

```
# $OpenBSD: Makefile,v 1.82 2019/07/12 20:47:14 sthen Exp $
```

```
COMMENT-main =           Japanese input method
COMMENT-dict =           dictionaries for Japanese Wnn
COMMENT-ko =             Korean input method
COMMENT-kodict =         dictionaries for Korean Wnn
```

```
[...]
```

```
DISTNAME =              Wnn4.2
```

```
CATEGORIES =            japanese
```

```
MULTI_PACKAGES =        -main -dict -zh -zhdic -ko -kodict -xwnmo -data
```

```
PKGNAME-main =          ja-Wnn-4.2
```

Enter FLAVORS and MULTI PACKAGES II

```
PKGNAME-dict =          ja-Wnndict-4.2
PKGNAME-ko =           ko-Wnn-4.2
PKGNAME-kodict =       ko-Wnndict-4.2
```

```
WNNDICBASE =          /var
WNNDICDIR =           ${WNNDICBASE}/dict/Wnn
WNNBASE =             ${LOCALBASE}
SUBST_VARS =          WNNBASE
```

```
PREFIX-dict =          ${WNNDICBASE}
PREFIX-kodict =        ${WNNDICBASE}
```

[...]

```
.include <bsd.port.mk>
```

So make-plist grew I

1100 lines in 2011! Horrible code:

```
#!/usr/bin/perl
# TODO
# - multi-packages with inter-dependencies still are not 100% correct with
# respect to common directories.
```

```
my %prefix;
my %plistname;
my %mtree;
my @subs;
my $baseprefix=$ENV{PREFIX};
my $shared_only;
my $make = $ENV{MAKE};
my $portsdir = $ENV{PORTSDIR};
```

So make-plist grew II

[...]

```
sub parse_arg
{
    my $_ = shift;
    if (m/^(DEPPATHS(-.*?))\=/) {
        $mtree{$1} = build_mtree($1, $');
        return;
    }
    if (m/\=/) {
        $subst->parse_option($_);
    }
    if (m/^\^PREFIX(\-.*?)\=(.*)\/?$/) {
        $prefix{$1} = $2;
    }
}
```

So make-plist grew III

```
    } elsif (m/^PLIST(\-.*?)\=/) {  
        $plistname{$1} = $';  
    }  
}
```

```
# Fragments are new PackingElement unique to make-plist and pkg_create,  
# to handle %%thingy%%.  
# (and so, make-plist will use a special PLIST reader)
```

```
# Method summary:  
# add_to_mtree: new directory in dependent package  
# add_to_haystack: put stuff so that it can be found on the FS  
# copy_extra: stuff that can't be easily deduced but should be copied  
# tag_along: set of items that associate themselves to this item
```

So make-plist grew IV

```
# (e.g., @exec, @unexec, @sample...)  
# clone_tags: copy tagged stuff over.  
# deduce_fragment: find fragment file name from %%stuff%%  
  
# note $plist->{nonempty}: set as soon as a plist holds anything  
# but a cvstag.  
  
[...]
```

Oh crap



Bad smells

- mismash of DESTDIR, PREFIX ...
- did not cope well with options
- ad-hoc code to make it work better
- each option took longer to implement



Yours is without a doubt the worst code I've ever run



But it runs



So around 2015, I decided to stop implementing new stuff
And started thinking about rewriting things
... People had started writing scripts to “help” `make_plist`

It took me a few years to get it into shape.

The basic structure was much easier, since I knew better.

But getting it to do correctly everything that the old tool did (incorrectly) took a long time!

Should I have done better ?

BOOrn again

- `make_plist` had an ad-hoc parser that did all of `pkg_create` work again.
- ... but `pkg_create` was now OO perl
- ... so the new `update_plist` could reuse `Pkg_Create.pm`
- synopsis would be `update_plist <options> -- pkg_create_args`
- and since *each* `pkg_create` synopsis was `pkg_create <options> pkg_name` it would work for multi-packages: just let the `Pkg_Create.pm` parser stop after the first “real” argument

To summarize

- One single `Pkg_Create.pm` parser
- that one has a `$state` that contains all the info
- for `update_plist`, we have a global `$state` with the `update_plist` options
- ... and one `$state` for each package in the multi-package set
- e.g., create a `Pkg_Create` object that just does this:

Parsing options I

```
1  sub process_next_subpackage
2  {
3      my $self = shift;
4      my $r = PlistReader->new;
5
6      my $s = PlistReader::State->new('update-plist', $self->{state});
7      $r->{state} = $s;
8      $s->handle_options;
9      $s->{opt}{q} = 1;
10     $r->{base_plists} = $s->{contents};
11     my $pkg = shift @ARGV;
12
13     $r->olist->set_pkgname($pkg);
14     $r->nlist->set_pkgname($pkg);
```

```
15     $self->{state}->say("Reading existing plist for #1", $pkg)
16         unless $self->{state}{quiet};
17     $r->read_all_fragments($s, $r->olist);
18     # add the cwd to new list as well!!!
19     OpenBSD::PackingElement::Cwd->add($r->nlist, $s->{prefix});
20     $r->add_extra_info($r->olist, $s);
21     push(@{$self->{lists}}, $r);
22 }
```

Checklist

- `parse_args` get the old plist(s)
- `known_objects` mark old files and directory locations
- `scan_fake_dir` just the old find with pattern recognition
- `handle_annotations` try to tie things to paths
- `copy_objects` actually create the new plists in correct order
- `add_missing_tags` global stuff needed for some files
- `add_delayed_objects` all about cwd
- `strip_dependency_directories`
- `adjust_final` whatever we didn't do

- Cool thing about OO, you can specialize things.
- There's an empty `annotate sub` in `Pkg_Create.pm`
- In `update_plist` it can match original lines with subst versions.

- there are 400 lines in `ReverseSubst.pm`
- basic algorithm is to sort a subpackage variable by reverse length, and try to reverse substitute the longest strings first...
- ... and there are 5 options related to specific usage
 - variables that only occur at start of path (`/etc`)
 - variables that only occur at end of paths (`.pyc`)
 - variables that shouldn't be added (only allowed if already there)
 - variables that expand to nothing
 - the python weird stuff

A typical python plist 1

```
1 @comment $OpenBSD: PLIST,v 1.1.1.1 2017/05/12 10:08:32 edd Exp $
2 lib/python${MODPY_VERSION}/site-packages/pyuv/
3 lib/python${MODPY_VERSION}/site-packages/pyuv/__init__.py
4 ${MODPY_COMMENT}lib/python${MODPY_VERSION}/site-packages/pyuv/${MODPY_PYCACHE}/
5 lib/python${MODPY_VERSION}/site-packages/pyuv/${MODPY_PYCACHE}__init__.${MODPY_PYC_MAGIC_TAG}pyc
6 lib/python${MODPY_VERSION}/site-packages/pyuv/${MODPY_PYCACHE}_version.${MODPY_PYC_MAGIC_TAG}pyc
7 lib/python${MODPY_VERSION}/site-packages/pyuv/_cpyuv.so
8 lib/python${MODPY_VERSION}/site-packages/pyuv/_version.py
```

- if several variables expand to the same value, we don't know, unless some variables are ignored. For instance, we prefer LOCALBASE to PREFIX
- there is a MODPY_COMMENT that expands to nothing in a python2 port and creates a python cache directory in python3.
- ... it is a comment because MODPY_CACHE vanishes, and we have duplicate dirs, and pkg_create does not like that.

```
1  # default backsubstitution and writing.
2  sub write_backsubst
3  {
4      my ($o, $file, $p) = @_ ;
5
6      if (defined (my $s = $o->{candidate_for_comment})) {
7          if ($p->{stash}{$s} > 1) {
8              $o->{prepared} =
9                  $p->subst->{maybe_comment}.$o->{prepared};
10         }
11     }
12     print {$file->fh} $o->{prepared}, "\n";
13 }
```

Basics

Really simple code, we just have a set of recognizers based on

- names
- contents

that create their own classes, which later get mapped to plist types.
... this can even be used for the subst handler (libraries...) and tags.

The fingerprint

There is specific code for resolving symlinks, thanks to DESTDIR.

Recognizers have access to an extra stash (data) so they can leave useful info for later recognizers (for example, running objdump)

In theory, this could be located in individual files.

How do we sort everything

- I decided to sort on actual filenames...
- ... so that adding a variable would not change everything
- this does change manually generated lists
- there are annotations that “tag along” with the nearest file (@sample, @comment no checksum)

Order please I

```
1  @comment $OpenBSD: PLIST-main,v 1.2 2019/06/07 20:24:04 sthen Exp $
2  @option no-default-conflict
3  @option is-branch
4  @conflict php->=7.3,<7.4
5  @extraunexec rm -f ${SYSCONFDIR}/php-${PV}.sample/*
6  @extraunexec rm -f ${SYSCONFDIR}/php-fpm.d/*
7  @mode 1700
8  @owner www
9  @group www
10 @sample ${CHROOT_DIR}/tmp/
11 @mode
12 @owner
13 @group
14 @rcscript ${RCDIR}/php${SV}_fpm
15 bin/phar-${PV}
16 @bin bin/php-${PV}
17 bin/php-config-${PV}
18 bin/phpize-${PV}
19 lib/php-${PV}/
20 lib/php-${PV}/modules/
21 lib/php-${PV}/modules/opcache.so
22 @man man/man1/phar-${PV}.1
23 @man man/man1/php-${PV}.1
24 @man man/man1/php-config-${PV}.1
25 @man man/man1/phpize-${PV}.1
26 @man man/man8/php-fpm-${PV}.8
27 @bin sbin/php-fpm-${PV}
28 share/doc/pkg-readmes/${PKGSTEM}
```

Order please II

```
29 share/examples/php-${PV}/
30 @sample ${SYSCONFDIR}/php-${PV}.sample/
31 @sample ${SYSCONFDIR}/php-fpm.d/
32 share/examples/php-${PV}/opcache.ini
33 @sample ${SYSCONFDIR}/php-${PV}.sample/opcache.ini
34 share/examples/php-${PV}/php-fpm.conf
35 @sample ${SYSCONFDIR}/php-fpm.conf
36 share/examples/php-${PV}/php.ini-development
37 share/examples/php-${PV}/php.ini-production
38 @sample ${SYSCONFDIR}/php-${PV}.ini
39 share/php-${PV}/
40 share/php-${PV}/build/
41 share/php-${PV}/build/Makefile.global
42 share/php-${PV}/build/acinclude.m4
43 [...]
44 share/php-${PV}/include/sapi/cli/
45 share/php-${PV}/include/sapi/cli/cli.h
46 @sample ${SYSCONFDIR}/php-${PV}/
```

Sorting through mail

- comments are a bit tricky
- if they correspond to real files they're sorted with them
- some comments are not really comments (cvs tags, no checksum)...
- otherwise, they just tag along with the nearest item in the old plist
- this is the most frustrating part: it will move things a bit sometimes

State comes last

- There are @mode, @owner annotations
- ... it comes last
- might create back and forth
- but the file system “stays” sorted

- we now have @define-tag and tag for file-specific handling
- the file system scanner creates subclasses of file with the right type.
- they do not get specific annotations
- ... but update_plist visits and every item and notices special files.
- ... so we can add the @tag to the generated list

PRIVSEP

- `update_plist` starts its life as root, with two users to go to
- the file system scanner runs as `_pbuild`
- the plist writer runs as `user`

pkglocate

- File systems entries are run against `pkglocate` to check for collisions.
- If those collisions aren't registered, we warn.
- This is rather slow for large plists, so this runs in parallel by default.

- Thanks to yandex, I have big machines to test things on
- so I was able to verify that plists didn't change
- also a huge improvement for multi-packages
- ended up having per-port options (`UPDATE_PLIST_ARGS`, well mostly for python) and user options (`UPDATE_PLIST_OPTS`).

Side note

<https://github.com/marcespie/openbsd-backgrounds>



Thank you for your time

Questions ?