

---

# **dune Documentation**

**J r mie Dimino**

**Aug 15, 2018**



<b>1</b>	<b>Quickstart</b>	<b>1</b>
1.1	Building a hello world program . . . . .	1
1.2	Building a hello world program using Lwt . . . . .	1
1.3	Building a hello world program using Core and Jane Street PPXs . . . . .	2
1.4	Defining a library using Lwt and ocaml-re . . . . .	2
1.5	Setting the OCaml compilation flags globally . . . . .	2
1.6	Using cppo . . . . .	3
1.6.1	Using the .cppo.ml style like the ocamlbuild plugin . . . . .	3
1.7	Defining a library with C stubs . . . . .	3
1.8	Defining a library with C stubs using pkg-config . . . . .	3
1.9	Using a custom code generator . . . . .	4
1.10	Defining tests . . . . .	5
1.11	Building a custom toplevel . . . . .	5
<b>2</b>	<b>Overview</b>	<b>7</b>
<b>3</b>	<b>Terminology</b>	<b>9</b>
<b>4</b>	<b>Project Layout and Metadata Specification</b>	<b>11</b>
4.1	Metadata format . . . . .	11
4.1.1	Comments . . . . .	11
4.1.2	Atoms . . . . .	12
4.1.3	Strings . . . . .	12
4.1.4	End of line strings . . . . .	12
4.1.5	Lists . . . . .	13
4.1.6	Variables . . . . .	13
4.2	dune-project files . . . . .	13
4.2.1	name . . . . .	13
4.2.2	version . . . . .	14
4.3	<package>.opam files . . . . .	14
4.3.1	Scopes . . . . .	14
4.3.2	Package version . . . . .	14
4.3.3	Odig conventions . . . . .	15
4.4	jbuild-ignore (deprecated) . . . . .	15
<b>5</b>	<b>dune files</b>	<b>17</b>
5.1	Stanzas . . . . .	17

5.1.1	jbuild_version	17
5.1.2	library	17
5.1.3	executable	20
5.1.4	executables	22
5.1.5	rule	22
5.1.6	ocamllex	23
5.1.7	ocamlyacc	24
5.1.8	menhir	24
5.1.9	alias	24
5.1.10	install	25
5.1.11	copy_files	26
5.1.12	include	26
5.1.13	tests	27
5.1.14	test	27
5.1.15	env	27
5.1.16	ignored_subdirs	27
5.1.17	include_subdirs	27
5.2	Common items	28
5.2.1	Ordered set language	28
5.2.2	Boolean Language	29
5.2.3	Variables expansion	29
5.2.4	Library dependencies	31
5.2.5	Preprocessing specification	32
5.2.6	Dependency specification	33
5.2.7	OCaml flags	35
5.2.8	js_of_ocaml	35
5.2.9	User actions	35
5.2.10	Locks	37
5.2.11	Diffing and promotion	37
5.3	OCaml syntax	38
<b>6</b>	<b>Writing and running tests</b>	<b>39</b>
6.1	Running tests	39
6.2	Inline tests	39
6.2.1	Inline expectation tests	40
6.2.2	Specifying inline test dependencies	41
6.2.3	Passing special arguments to the test runner	41
6.2.4	Using additional libraries in the test runner	42
6.2.5	Defining your own inline test backend	42
6.3	Custom tests	43
6.3.1	Diffing the result	43
<b>7</b>	<b>Generating Documentation</b>	<b>45</b>
7.1	Prerequisites	45
7.2	Writing Documentation	45
7.3	Building Documentation	45
7.4	Documentation Stanza	46
7.4.1	Examples	46
<b>8</b>	<b>Usage</b>	<b>47</b>
8.1	Finding the root	47
8.1.1	dune-workspace	47
8.1.2	Current directory	48
8.1.3	Forcing the root (for scripts)	48

8.2	Interpretation of targets	48
8.2.1	Resolution	48
8.2.2	Aliases	48
8.2.3	Default alias	49
8.3	Finding external libraries	49
8.4	Running tests	49
8.5	Launching the Toplevel (REPL)	50
8.5.1	Requirements & Limitations	50
8.6	Restricting the set of packages	50
8.7	Invocation from opam	50
8.8	Tests	51
8.9	Installation	51
8.9.1	Destination	51
8.10	Workspace configuration	51
8.10.1	dune-workspace	52
8.11	Distributing Projects	53
8.12	Watermarking	53
8.13	dune subst	53
8.14	Custom Build Directory	54
<b>9</b>	<b>Advanced topics</b>	<b>55</b>
9.1	META file generation	55
9.2	Findlib integration and limitations	55
9.3	Cross Compilation	56
9.3.1	How does it work?	57
9.4	Classical ppx	57
<b>10</b>	<b>Configurator</b>	<b>59</b>
10.1	Usage	59
10.2	Upgrading from the old Configurator	60
<b>11</b>	<b>Menhir</b>	<b>63</b>
<b>12</b>	<b>js_of_ocaml</b>	<b>65</b>
12.1	Compiling to JS	65
12.2	js_of_ocaml field	65
12.3	Separate Compilation	66
<b>13</b>	<b>FAQ</b>	<b>67</b>
13.1	Why do many dune projects contain a Makefile?	67
13.2	How to add a configure step to a dune project?	67
13.3	Can I use topkg with dune?	67
13.4	How do I publish my packages with dune?	67
13.5	here can I find some examples of projects using dune?	68
<b>14</b>	<b>Known Issues</b>	<b>69</b>
14.1	mli only modules	69
<b>15</b>	<b>Migration</b>	<b>71</b>
15.1	Timeline	71
15.1.1	July 2018: release of Dune 1.0.0	71
15.1.2	January 2019: deprecation of Jbuilder	71
15.1.3	July 2019: support for Jbuilder is dropped	72
15.1.4	January 2020: the jbuilder binary goes away	72
15.1.5	Distant future	72

15.2	Check List	72
15.2.1	New configuration files	72
15.2.2	dune-project files	72
15.2.3	dune files	73
15.2.4	dune-workspace	73
15.2.5	Variable Syntax	73
15.2.6	(files_recursively_in ..) is removed	73
15.2.7	Escape Sequences	73
15.2.8	Comments Syntax	73
15.2.9	Renamed Variables	73
15.2.10	Removed Variables	74
15.2.11	# JBUILDER_GEN renamed	74

This document gives simple usage examples of dune. You can also look at [examples](#) for complete examples of projects using dune.

## 1.1 Building a hello world program

In a directory of your choice, write this dune file:

```
;; This declares the hello_world executable implemented by hello_world.ml
(executable
 (name hello_world))
```

This `hello_world.ml` file:

```
print_endline "Hello, world!"
```

And build it with:

```
dune build hello_world.exe
```

The executable will be built as `_build/default/hello_world.exe`. Note that native code executables will have the `.exe` extension on all platforms (including non-Windows systems). The executable can be run with: `$ dune exec ./hello_world.exe`

## 1.2 Building a hello world program using Lwt

In a directory of your choice, write this dune file:

```
(executable
 (name hello_world)
 (libraries lwt.unix))
```

This `hello_world.ml` file:

```
Lwt_main.run (Lwt_io.printf "Hello, world!\n")
```

And build it with:

```
dune build hello_world.exe
```

The executable will be built as `_build/default/hello_world.exe`

### 1.3 Building a hello world program using Core and Jane Street PPXs

Write this dune file:

```
(executable
 (name hello_world)
 (libraries core)
 (preprocess (pps ppx_jane)))
```

This `hello_world.ml` file:

```
open Core

let () =
  Sexp.to_string_hum [%sexp ([3;4;5] : int list)]
  |> print_endline
```

And build it with:

```
dune build hello_world.exe
```

The executable will be built as `_build/default/hello_world.exe`

### 1.4 Defining a library using Lwt and ocaml-re

Write this dune file:

```
(library
 (name      mylib)
 (public_name mylib)
 (libraries re lwt))
```

The library will be composed of all the modules in the same directory. Outside of the library, module `Foo` will be accessible as `Mylib.Foo`, unless you write an explicit `mylib.ml` file.

You can then use this library in any other directory by adding `mylib` to the `(libraries ...)` field.

### 1.5 Setting the OCaml compilation flags globally

Write this dune file at the root of your project:



```
(env
 (dev
  (flags (:standard -w +42)))
 (release
  (flags (:standard -O3))))
```

*dev* and *release* correspond to build profiles. The build profile can be selected from the command line with *-profile foo* or from a *dune-workspace* file by writing:

```
(profile foo)
```

## 1.6 Using cppo

Add this field to your `library` or `executable` stanzas:

```
(preprocess (action (run %{bin:cppo} -V OCAML:%{ocaml_version} %{input-file})))
```

Additionally, if you are include a `config.h` file, you need to declare the dependency to this file via:

```
(preprocessor_deps config.h)
```

### 1.6.1 Using the `.cppo.ml` style like the `ocamlbuild` plugin

Write this in your `dune` file:

```
(rule
 (targets foo.ml)
 (deps (:first-dep foo.cppo.ml) <other files that foo.ml includes>)
 (action (run %{bin:cppo} %{first-dep} -o %{targets})))
```

## 1.7 Defining a library with C stubs

Assuming you have a file called `mystubs.c`, that you need to pass `-I/blah/include` to compile it and `-lblah` at link time, write this `dune` file:

```
(library
 (name          mylib)
 (public_name   mylib)
 (libraries     re lwt)
 (c_names       mystubs)
 (c_flags       (-I/blah/include))
 (c_library_flags (-lblah)))
```

## 1.8 Defining a library with C stubs using `pkg-config`

Same context as before, but using `pkg-config` to query the compilation and link flags. Write this `dune` file:

```
(library
  (name          mylib)
  (public_name  mylib)
  (libraries    re lwt)
  (c_names      mystubs)
  (c_flags      (:include c_flags.sexp))
  (c_library_flags (:include c_library_flags.sexp)))

(rule
  (targets c_flags.sexp c_library_flags.sexp)
  (deps    (:discover config/discover.exe))
  (action  (run %{discover} -ocamlc %{OCAMLC})))
```

Then create a config subdirectory and write this dune file:

```
(executable
  (name discover)
  (libraries base stdio configurator))
```

as well as this discover.ml file:

```
open Base
open Stdio
module C = Configurator

let write_sexp fn sexp =
  Out_channel.write_all fn ~data:(Sexp.to_string sexp)

let () =
  C.main ~name:"mylib" (fun c ->
    let default : C.Pkg_config.package_conf =
      { libs = ["-lblah"]
        ; cflags = []
        }
    in
    let conf =
      match C.Pkg_config.get c with
      | None -> default
      | Some pc ->
        Option.value (C.Pkg_config.query pc ~package:"blah") ~default
    in

    write_sexp "c_flags.sexp" (sexp_of_list sexp_of_string conf.cflags);
    write_sexp "c_library_flags.sexp" (sexp_of_list sexp_of_string conf.libs))
```

## 1.9 Using a custom code generator

To generate a file `foo.ml` using a program from another directory:

```
(rule
  (targets foo.ml)
  (deps    (:gen ../generator/gen.exe))
  (action  (run %{gen} -o %{targets})))
```

## 1.10 Defining tests

Write this in your dune file:

```
(test (name my_test_program))
```

And run the tests with:

```
dune runtest
```

It will run the test program (the main module is `my_test_program.ml`) and error if it exits with a nonzero code.

In addition, if a `my_test_program.expected` file exists, it will be compared to the standard output of the test program and the differences will be displayed. It is possible to replace the `.expected` file with the last output using:

```
dune promote
```

## 1.11 Building a custom toplevel

A toplevel is simply an executable calling `Topmain.main ()` and linked with the compiler libraries and `-linkall`. Moreover, currently toplevels can only be built in bytecode.

As a result, write this in your dune file:

```
(executable
  (name      mytoplevel)
  (libraries compiler-libs.toplevel mylib)
  (link_flags (-linkall))
  (modes     byte))
```

And write this in `mytoplevel.ml`

```
let () = Topmain.main ()
```



Dune is a build system for OCaml and Reason. It is not intended as a completely generic build system that is able to build any given project in any language. On the contrary, it makes lots of choices in order to encourage a consistent development style.

This scheme is inspired from the one used inside Jane Street and adapted to the opam world. It has matured over a long time and is used daily by hundred of developers, which means that it is highly tested and productive.

When using dune, you give very little and high-level information to the build system, which in turn takes care of all the low-level details, from the compilation of your libraries, executables and documentation, to the installation, setting up of tests, setting up of the development tools such as merlin, etc.

In addition to the normal features one would expect from a build system for OCaml, dune provides a few additional ones that detach it from the crowd:

- you never need to tell dune where things such as libraries are. Dune will always discover them automatically. In particular, this means that when you want to re-organize your project you need to do no more than rename your directories, dune will do the rest
- things always work the same whether your dependencies are local or installed on the system. In particular, this means that you can always drop in the source for a dependency of your project in your working copy and dune will start using it immediately. This makes dune a great choice for multi-project development
- cross-platform: as long as your code is portable, dune will be able to cross-compile it (note that dune is designed internally to make this easy but the actual support is not implemented yet)
- release directly from any revision: dune needs no setup stage. To release your project, you can simply point to a specific tag. You can of course add some release steps if you want to, but it is not necessary

The first section of this document defines some terms used in the rest of this manual. The second section specifies the dune metadata format and the third one describes how to use the `dune` command.



- **package:** a package is a set of libraries, executables, ... that are built and installed as one by opam
- **project:** a project is a source tree, maybe containing one or more packages
- **root:** the root is the directory from where dune can build things. Dune knows how to build targets that are descendents of the root. Anything outside of the tree starting from the root is considered part of the **installed world**. How the root is determined is explained in *Finding the root*.
- **workspace:** the workspace is the subtree starting from the root. It can contain any number of projects that will be built simultaneously by dune
- **installed world:** anything outside of the workspace, that dune takes for granted and doesn't know how to build
- **installation:** this is the action of copying build artifacts or other files from the `<root>/_build` directory to the installed world
- **scope:** a scope determines where private items are visible. Private items include libraries or binaries that will not be installed. In dune, scopes are sub-trees rooted where at least one `<package>.opam` file is present. Moreover, scopes are exclusive. Typically, every project defines a single scope. See *Scopes* for more details
- **build context:** a build context is a subdirectory of the `<root>/_build` directory. It contains all the build artifacts of the workspace built against a specific configuration. Without specific configuration from the user, there is always a `default` build context, which corresponds to the environment in which dune is executed. Build contexts can be specified by writing a *dune-workspace* file
- **build context root:** the root of a build context named `foo` is `<root>/_build/<foo>`
- **alias: an alias is a build target that doesn't produce any file** and has configurable dependencies. Aliases are per-directory. However, on the command line, asking for an alias to be built in a given directory will trigger the construction of the alias in all children directories recursively. Dune defines the following standard aliases:
  - `default` which is the alias build by default when no targets are specified on the command line. See *Default alias* for details
  - `runtest` which runs user defined tests
  - `install` which depends on everything that should be installed

- `doc` which depends on the generated HTML documentation. See *Generating Documentation* for details.
- **environment:** in dune, each directory has an environment attached to it. The environment determines the default values of various parameters, such as the compilation flags. Inside a scope, each directory inherit the environment from its parent. At the root of every scope, a default environment is used. At any point, the environment can be altered using an *env* stanza.
- **build profile:** a global setting that influence various defaults. It can be set from the command line using `--profile <profile>` or from `dune-workspace` files. The following profiles are standard:
  - `release` which is the profile used for opam releases
  - `dev` which is the default profile when none is set explicitly, it has stricter warnings than the `release` one



---

## Project Layout and Metadata Specification

---

A typical dune project will have a `dune-project` and one or more `<package>.opam` file at toplevel as well as dune files wherever interesting things are: libraries, executables, tests, documents to install, etc. . .

It is recommended to organize your project so that you have exactly one library per directory. You can have several executables in the same directory, as long as they share the same build configuration. If you'd like to have multiple executables with different configurations in the same directory, you will have to make an explicit module list for every executable using `modules`.

The next sections describe the format of dune metadata files.

Note that the dune metadata format is versioned in order to ensure forward compatibility. There is currently only one version available, but to be future proof, you should still specify it in your dune files. If no version is specified, the latest one will be used.

### 4.1 Metadata format

All configuration files read by Dune are using a syntax similar to the one of S-expressions, which is very simple. The Dune language can represent three kinds of values: atoms, strings and lists. By combining these, it is possible to construct arbitrarily complex project descriptions.

A Dune configuration file is a sequence of atoms, strings or lists separated by spaces, newlines and comments. The other sections of this manual describe how each configuration file is interpreted. We describe below the syntax of the language.

#### 4.1.1 Comments

The Dune language only has end of line comments. End of line comments are introduced with a semicolon and span up to the end of the end of the current line. Everything from the semicolon to the end of the line is ignored. For instance:

```
; This is a comment
```

### 4.1.2 Atoms

An atom is a non-empty contiguous sequences of character other than special characters. Special characters are:

- spaces, horizontal tabs, newlines and form feed
- opening and closing parenthesis
- double quotes
- semicolons

For instance `hello` or `+` are valid atoms.

Note that backslashes inside atoms have no special meaning are always interpreted as plain backslashes characters.

### 4.1.3 Strings

A string is a sequence of characters surrounded by double quotes. A string represent the exact text between the double quotes, except for escape sequences. Escape sequence are introduced by the a backslash character. Dune recognizes and interprets the following escape sequences:

- `\n` to represent a newline character
- `\r` to represent a carriage return (character with ASCII code 13)
- `\b` to represent ASCII character 8
- `\t` to represent a horizontal tab
- `\NNN`, a backslash followed by three decimal characters to represent the character with ASCII code `NNN`
- `\xHH`, a backslash followed by two hexadecimal characters to represent the character with ASCII code `HH` in hexadecimal
- `\\`, a double backslash to represent a single backslash
- `\%{` to represent `%{` (see *Variables*)

Additionally, a backslash that comes just before the end of the line is used to skip the newline up to the next non-space character. For instance the following two strings represent the same text:

```
"abcdef"  
"abc\  
  def"
```

In most places where Dune expect a string, it will also accept an atom. As a result it possible to write most Dune configuration file using very few double quotes. This is very convenient in practice.

### 4.1.4 End of line strings

End of line strings are another way to write strings. The are a convenient way to write blocks of text inside a Dune file.

End of line strings are introduced by `"\|` or `"\>` and span up the end of the current line. If the next line starts as well by `"\|` or `"\>` it is the continuation of the same string. For readability, it is necessary that the text that follows the delimiter is either empty or starts with a space that is ignored.

For instance:

```
"\| this is a block
"\| of text
```

represent the same text as the string "this is a block\nof text".

Escape sequences are interpreted in text that follows "\|" but not in text that follows "\>". Both delimiters can be mixed inside the same block of text.

### 4.1.5 Lists

Lists are sequences of values enclosed by parentheses. For instance `(x y z)` is a list containing the three atoms `x`, `y` and `z`. Lists can be empty, for instance: `()`.

Lists can be nested, allowing to represent arbitrarily complex descriptions. For instance:

```
(html
 (head (title "Hello world!"))
 (body
  This is a simple example of using S-expressions))
```

### 4.1.6 Variables

Dune allows variables in a few places. Their interpretation often depend on the context in which they appear.

The syntax of variables is as follow:

```
%{var}
```

or, for more complex forms that take an argument:

```
%{fun:arg}
```

In order to write a plain `%{`, you need to write `\%{` in a string.

## 4.2 dune-project files

These files are used to mark the root of projects as well as define project-wide parameters. These files are required to have a `lang` which controls the names and contents of all configuration files read by Dune. The `lang` stanza looks like:

```
(lang dune 1.0)
```

Additionally, they can contains the following stanzas.

### 4.2.1 name

Sets the name of the project:

```
(name <name>)
```

### 4.2.2 version

Sets the version of the project:

```
(version <version>)
```

## 4.3 <package>.opam files

When a <package>.opam file is present, dune will know that the package named <package> exists. It will know how to construct a <package>.install file in the same directory to handle installation via [opam](#). Dune also defines the recursive `install` alias, which depends on all the buildable <package>.install files in the workspace. So for instance to build everything that is installable in a workspace, run at the root:

```
$ dune build @install
```

Declaring a package this way will allow you to add elements such as libraries, executables, documentation, ... to your package by declaring them in dune files.

Such elements can only be declared in the scope defined by the corresponding <package>.opam file. Typically, your <package>.opam files should be at the root of your project, since this is where `opam pin ...` will look for them.

Note that <package> must be non-empty, so in particular .opam files are ignored.

### 4.3.1 Scopes

Any directory containing at least one <package>.opam file defines a scope. This scope is the sub-tree starting from this directory, excluding any other scopes rooted in sub-directories.

Typically, any given project will define a single scope. Libraries and executables that are not meant to be installed will be visible inside this scope only.

Because scopes are exclusive, if you wish to include the dependencies of the project you are currently working on into your workspace, you may copy them in a `vendor` directory, or any other name of your choice. Dune will look for them there rather than in the installed world and there will be no overlap between the various scopes.

### 4.3.2 Package version

Note that dune will try to determine the version number of packages defined in the workspace. While dune itself makes no use of version numbers, it can be used by external tools such as [ocamlfind](#).

Dune determines the version of a package by trying the following methods in order:

- it looks in the <package>.opam file for a `version` variable
- it looks for a <package>.version file in the same directory and reads the first line
- it looks for the version specified in the `dune-project` if present
- it looks for a `version` file and reads the first line
- it looks for a `VERSION` file and reads the first line

<package>.version, version and VERSION files may be generated.

If the version can't be determined, dune just won't assign one.

### 4.3.3 Odig conventions

Dune follows the `odig` conventions and automatically installs any `README*`, `CHANGE*`, `HISTORY*` and `LICENSE*` files in the same directory as the `<package>.opam` file to a location where `odig` will find them.

Note that this includes files present in the source tree as well as generated files. So for instance a changelog generated by a user rule will be automatically installed as well.

## 4.4 `jbuild-ignore` (deprecated)

`jbuild-ignore` files are deprecated and replaced by `ignored_subdirs` stanzas in dune files.



dune files are the main part of dune. They are used to describe libraries, executables, tests, and everything dune needs to know about.

The syntax of dune files is described in *Metadata format* section.

## 5.1 Stanzas

dune files are composed of stanzas. For instance a typical dune looks like:

```
(library
 (name mylib)
 (libraries base lwt))

(rule
 (targets foo.ml)
 (deps    generator/gen.exe)
 (action  (run %{deps} -o %{targets})))
```

The following sections describe the available stanzas and their meaning.

### 5.1.1 jbuild\_version

Deprecated. This stanza is no longer used and will be removed in the future.

### 5.1.2 library

The `library` stanza must be used to describe OCaml libraries. The format of library stanzas is as follows:

```
library
  (name <library-name>
   <optional-fields>
  )
```

`<library-name>` is the real name of the library. It determines the names of the archive files generated for the library as well as the module name under which the library will be available, unless `(wrapped false)` is used (see below). It must be a valid OCaml module name but doesn't need to start with a uppercase letter.

For instance, the modules of a library named `foo` will be available as `FOO.XXX` outside of `foo` itself. It is however allowed to write an explicit `FOO` module, in which case this will be the interface of the library and you are free to expose only the modules you want.

Note that by default libraries and other things that consume OCaml/Reason modules only consume modules from the directory where the stanza appear. In order to declare a multi-directory library, you need to use the `include_subdirs` stanza.

`<optional-fields>` are:

- `(public_name <name>)` this is the name under which the library can be referred to as a dependency when it is not part of the current workspace, i.e. when it is installed. Without a `(public_name ...)` field, the library will not be installed by dune. The public name must start by the package name it is part of and optionally followed by a dot and anything else you want. The package name must be one of the packages that dune knows about, as determined by the *dune-project files*
- `(synopsis <string>)` should give a one-line description of the library. This is used by tools that list installed libraries
- `(modules <modules>)` specifies what modules are part of the library. By default dune will use all the `.ml/.re` files in the same directory as the dune file. This include ones that are present in the file system as well as ones generated by user rules. You can restrict this list by using a `(modules <modules>)` field. `<modules>` uses the *Ordered set language* where elements are module names and don't need to start with a uppercase letter. For instance to exclude module `Foo`: `(modules (:standard \ foo))`
- `(libraries <library-dependencies>)` is used to specify the dependencies of the library. See the section about *Library dependencies* for more details
- `(wrapped <boolean>)` specifies whether the modules of the library should be available only through the top-level library module, or should all be exposed at the top level. The default is `true` and it is highly recommended to keep it this way. Because OCaml top-level modules must all be unique when linking an executables, polluting the top-level namespace will make your library unusable with other libraries if there is a module name clash. This option is only intended for libraries that manually prefix all their modules by the library name and to ease porting of existing projects to dune
- `(preprocess <preprocess-spec>)` specifies how to preprocess files if needed. The default is `no_processing`. Other options are described in the *Preprocessing specification* section
- `(preprocessor_deps (<deps-conf list>))` specifies extra dependencies of the preprocessor, for instance if the preprocessor reads a generated file. The specification of dependencies is described in the *Dependency specification* section
- `(optional)`, if present it indicates that the library should only be built and installed if all the dependencies are available, either in the workspace or in the installed world. You can use this to provide extra features without adding hard dependencies to your project
- `(c_names (<names>))`, if your library has stubs, you must list the C files in this field, without the `.c` extension
- `(cxx_names (<names>))` is the same as `c_names` but for C++ stubs



- `(install_c_headers (<names>))`, if your library has public C header files that must be installed, you must list them in this field, without the `.h` extension
- `(modes <modes>)` modes which should be built by default. The most common use for this feature is to disable native compilation when writing libraries for the OCaml toplevel. The following modes are available: `byte`, `native` and `best`. `best` is `native` or `byte` when native compilation is not available
- `(no_dynlink)` is to disable dynamic linking of the library. This is for advanced use only, by default you shouldn't set this option
- `(kind <kind>)` is the kind of the library. The default is `normal`, other available choices are `ppx_rewriter` and `ppx_deriver` and must be set when the library is intended to be used as a ppx rewriter or a `[@@deriving ...]` plugin. The reason why `ppx_rewriter` and `ppx_deriver` are split is historical and hopefully we won't need two options soon
- `(ppx_runtime_libraries (<library-names>))` is for when the library is a ppx rewriter or a `[@@deriving ...]` plugin and has runtime dependencies. You need to specify these runtime dependencies here
- `(virtual_deps (<opam-packages>))`. Sometimes opam packages enable a specific feature only if another package is installed. This is for instance the case of `ctypes` which will only install `ctypes.foreign` if the dummy `ctypes-foreign` package is installed. You can specify such virtual dependencies here. You don't need to do so unless you use dune to synthesize the `depends` and `depopts` sections of your opam file
- `js_of_ocaml`. See the section about [js\\_of\\_ocaml](#)
- `flags`, `ocamlc_flags` and `ocamlopt_flags`. See the section about [OCaml flags](#)
- `(library_flags (<flags>))` is a list of flags that are passed as it to `ocamlc` and `ocamlopt` when building the library archive files. You can use this to specify `-linkall` for instance. `<flags>` is a list of strings supporting [Variables expansion](#)
- `(c_flags <flags>)` specifies the compilation flags for C stubs, using the [Ordered set language](#). This field supports `(:include ...)` forms
- `(cxx_flags <flags>)` is the same as `c_flags` but for C++ stubs
- `(c_library_flags <flags>)` specifies the flags to pass to the C compiler when constructing the library archive file for the C stubs. `<flags>` uses the [Ordered set language](#) and supports `(:include ...)` forms. When you are writing bindings for a C library named `bar`, you should typically write `-lbar` here, or whatever flags are necessary to link against this library
- `(self_build_stubs_archive <c-libname>)` indicates to dune that the library has stubs, but that the stubs are built manually. The aim of the field is to embed a library written in foreign language and/or building with another build system. It is not for casual uses, see the [re2 library](#) for an example of use
- `(modules_without_implementation <modules>)` specifies a list of modules that have only a `.mli` or `.rei` but no `.ml` or `.re` file. Such modules are usually referred as *mli only modules*. They are not officially supported by the OCaml compiler, however they are commonly used. Such modules must only define types. Since it is not reasonably possible for dune to check that this is the case, dune requires the user to explicitly list such modules to avoid surprises. `<modules>` must be a subset of the modules listed in the `(modules ...)` field.
- `(allow_overlapping_dependencies)` allows external dependencies to overlap with libraries that are present in the workspace
- `(no_keep_locs)` undocumented, it is a necessary hack until this is implemented: <https://github.com/ocaml/dune/issues/921>

Note that when binding C libraries, dune doesn't provide special support for tools such as `pkg-config`, however it integrates easily with [configurator](#) by using `(c_flags (:include ...))` and `(c_library_flags (:include ...))`.

### 5.1.3 executable

The executable stanza must be used to describe an executable. The format of executable stanzas is as follows:

```
(executable
  (name <name>)
  <optional-fields>
)
```

<name> is a module name that contains the main entry point of the executable. There can be additional modules in the current directory, you only need to specify the entry point. Given an executable stanza with (name <name>), dune will know how to build <name>.exe, <name>.bc and <name>.bc.js. <name>.exe is a native code executable, <name>.bc is a bytecode executable which requires `ocamlrun` to run and <name>.bc.js is a JavaScript generated using `js_of_ocaml`.

Note that in case native compilation is not available, <name>.exe will in fact be a custom byte-code executable. Custom in the sense of `ocamlc -custom`, meaning that it is a native executable that embeds the `ocamlrun` virtual machine as well as the byte code. As such you can always rely on <name>.exe being available. Moreover, it is usually preferable to use <name>.exe in custom rules or when calling the executable by hand. This is because running a byte-code executable often requires loading shared libraries that are locally built, and so requires additional setup such as setting specific environment variables and dune doesn't do at the moment.

Native compilation is considered not available when there is no `ocamlopt` binary at the same place as where `ocamlc` was found.

Executables can also be linked as object or shared object files. See *linking modes* for more information.

<optional-fields> are:

- (public\_name <public-name>) specifies that the executable should be installed under that name. It is the same as adding the following stanza to your dune file:

```
(install
  (section bin)
  (files (<name>.exe as <public-name>)))
```

- (package <package>) if there is a (public\_name ...) field, this specifies the package the executables are part of
- (libraries <library-dependencies>) specifies the library dependencies. See the section about *Library dependencies* for more details
- (link\_flags <flags>) specifies additional flags to pass to the linker. This field supports (:include ...) forms
- (link\_deps (<deps-conf list>)) specifies the dependencies used only by the linker, for example when using a version script. See the *Dependency specification* section for more details.
- (modules <modules>) specifies which modules in the current directory dune should consider when building this executable. Modules not listed here will be ignored and cannot be used inside the executable described by the current stanza. It is interpreted in the same way as the (modules ...) field of *library*
- (modes (<modes>)) sets the *linking modes*. The default is (byte exe)
- (preprocess <preprocess-spec>) is the same as the (preprocess ...) field of *library*
- (preprocessor\_deps (<deps-conf list>)) is the same as the (preprocessor\_deps ...) field of *library*
- js\_of\_ocaml. See the section about *js\_of\_ocaml*
- **flags, ocamlc\_flags and ocamlopt\_flags.** See the section about specifying *OCaml flags*

- (`modules_without_implementation <modules>`) is the same as the corresponding field of *library*
- (`allow_overlapping_dependencies`) is the same as the corresponding field of *library*

## Linking modes

The `modes` field allows to select what linking modes should be used to link executables. Each mode is a pair (`<compilation-mode> <binary-kind>`) where `<compilation-mode>` describes whether the byte code or native code backend of the OCaml compiler should be used and `<binary-kind>` describes what kind of file should be produced.

`<compilation-mode>` must be `byte`, `native` or `best`, where `best` is `native` with a fallback to `byte-code` when native compilation is not available.

`<binary-kind>` is one of:

- `exe` for normal executables
- `object` for producing static object files that can be manually linked into C applications
- `shared_object` for producing object files that can be dynamically loaded into an application. This mode can be used to write a plugin in OCaml for a non-OCaml application.

For instance the following executables stanza will produce byte code executables and native shared objects:

```
(executables
  ((names (a b c))
   (modes ((byte exe) (native shared_object)))))
```

Additionally, you can use the following short-hands:

- `exe` for `(best exe)`
- `object` for `(best object)`
- `shared_object` for `(best shared_object)`
- `byte` for `(byte exe)`
- `native` for `(native exe)`

For instance the following modes fields are all equivalent:

```
(modes (exe object shared_object))
(modes ((best exe)
        (best object)
        (best shared_object)))
```

The extensions for the various linking modes are chosen as follows:

compilation mode	binary kind	extensions
byte	exe	.bc and .bc.js
native/best	exe	.exe
byte	object	.bc%{ext_obj}
native/best	object	.exe%{ext_obj}
byte	shared_object	.bc%{ext_dll}
native/best	shared_object	%{ext_dll}

Where `%{ext_obj}` and `%{ext_dll}` are the extensions for object and shared object files. Their value depends on the OS, for instance on Unix `%{ext_obj}` is usually `.o` and `%{ext_dll}` is usually `.so` while on Windows `%{ext_obj}` is `.obj` and `%{ext_dll}` is `.dll`.

Note that when `(byte exe)` is specified but neither `(best exe)` nor `(native exe)` are specified, Jbuild still knows how to build an executable with the extension `.exe`. In such case, the `.exe` version is the same as the `.bc` one except that it is linked with the `-custom` option of the compiler. You should always use the `.exe` rather than the `.bc` inside build rules.

### 5.1.4 executables

The `executables` stanza is the same as the `executable` stanza, except that it is used to describe several executables sharing the same configuration.

It shares the same fields as the `executable` stanza, except that instead of `(name ...)` and `(public_name ...)` you must use:

- `(names <names>)` where `<names>` is a list of entry point names. As for `executable` you only need to specify the modules containing the entry point of each executable
- `(public_names <names>)` describes under what name each executable should be installed. The list of names must be of the same length as the list in the `(names ...)` field. Moreover you can use `-` for executables that shouldn't be installed

### 5.1.5 rule

The `rule` stanza is used to create custom user rules. It tells dune how to generate a specific set of files from a specific set of dependencies.

The syntax is as follows:

```
(rule
  (targets <filenames>)
  (action <action>)
  <optional-fields>)
```

`<filenames>` is a list of file names. Note that currently dune only support user rules with targets in the current directory.

`<action>` is the action to run to produce the targets from the dependencies. See the [User actions](#) section for more details.

`<optional-fields>` are:

- `(deps <deps-conf list>)` to specify the dependencies of the rule. See the [Dependency specification](#) section for more details.
- `(mode <mode>)` to specify how to handle the targets, see [modes](#) for details
- `(fallback)` is deprecated and is the same as `(mode fallback)`
- `(locks (<lock-names>))` specify that the action must be run while holding the following locks. See the [Locks](#) section for more details.

Note that contrary to makefiles or other build systems, user rules currently don't support patterns, such as a rule to produce `%.y` from `%.x` for any given `%`. This might be supported in the future.

## modes

By default, the target of a rule must not exist in the source tree and dune will error out when this is the case.

However, it is possible to change this behavior using the `mode` field. The following modes are available:

- `standard`, this is the standard mode
- `fallback`, in this mode, when the targets are already present in the source tree, dune will ignore the rule. It is an error if only a subset of the targets are present in the tree. The common use of fallback rules is to generate default configuration files that may be generated by a configure script.
- `promote`, in this mode, the files in the source tree will be ignored. Once the rule has been executed, the targets will be copied back to the source tree
- `promote-until-clean` is the same as `promote` except than `dune clean` will remove the promoted files from the source tree

There are two use cases for promote rules. The first one is when the generated code is easier to review than the generator, so it's easier to commit the generated code and review it. The second is to cut down dependencies during releases: by passing `--ignore-promoted-rules` to dune, rules will (mode `promote`) will be ignored and the source files will be used instead. The `-p/--for-release-of-packages` flag implies `--ignore-promote-rules`.

## inferred rules

When using the action DSL (see *User actions*), it is most of the time obvious what are the dependencies and targets.

For instance:

```
(rule
 (targets b)
 (deps    a)
 (action  (copy %{deps} %{targets})))
```

In this example it is obvious by inspecting the action what the dependencies and targets are. When this is the case you can use the following shorter syntax, where dune infers dependencies and targets for you:

```
(rule <action>)
```

For instance:

```
(rule (copy a b))
```

Note that in dune, targets must always be known statically. Especially, this mean that dune must be able to statically determine all targets. For instance, this `(rule ...)` stanza is rejected by dune:

```
(rule (copy a b.%{read:file}))
```

### 5.1.6 ocamllex

`(ocamllex <names>)` is essentially a shorthand for:

```
(rule
 (targets <name>.ml)
 (deps    <name>.mll))
```

(continues on next page)

(continued from previous page)

```
(action (chdir %{workspace_root}
         (run %{bin:ocamllex} -q -o %{targets} %{deps}))))
```

To use a different rule mode, use the long form:

```
(ocamllex
 (modules <names>)
 (mode <mode>))
```

### 5.1.7 ocaml yacc

(ocaml yacc <names>) is essentially a shorthand for:

```
(rule
 (targets <name>.ml <name>.mli)
 (deps <name>.mly)
 (action (chdir %{workspace_root}
          (run %{bin:ocaml yacc} %{deps}))))
```

To use a different rule mode, use the long form:

```
(ocaml yacc
 (modules <names>)
 (mode <mode>))
```

### 5.1.8 menhir

A `menhir` stanza is available to support the *menhir* parser generator. See the *Menhir* section for details.

### 5.1.9 alias

The `alias` stanza lets you add dependencies to an alias, or specify an action to run to construct the alias.

The syntax is as follows:

```
(alias
 (name <alias-name>)
 (deps <deps-conf list>)
 <optional-fields>)
```

<name> is an alias name such as `runtest`.

<deps-conf list> specifies the dependencies of the alias. See the *Dependency specification* section for more details.

<optional-fields> are:

- <action>, an action to run when constructing the alias. See the *User actions* section for more details.
- (package <name>) indicates that this alias stanza is part of package <name> and should be filtered out if <name> is filtered out from the command line, either with `--only-packages <pkg>` or `-p <pkg>`
- (locks (<lock-names>)) specify that the action must be run while holding the following locks. See the *Locks* section for more details.

- `(enabled_if <blang expression>)` specifies the boolean condition that must be true for the tests to run. The condition is specified using the *blang*, and the field allows for *variables* to appear in the expressions.

The typical use of the `alias` stanza is to define tests:

```
(alias
 (name   runtest)
 (action (run %{exe:my-test-program.exe} blah)))
```

See the section about *Running tests* for details.

Note that if your project contains several packages and you run test the tests from the opam file using a `build-test` field, then all your `runtest` alias stanzas should have a `(package ...)` field in order to partition the set of tests.

### 5.1.10 install

The `install` stanza is what lets you describe what dune should install, either when running `dune install` or through opam.

Libraries and executables don't need an `install` stanza to be installed, just a `public_name` field. Everything else needs an `install` stanza.

The syntax is as follows:

```
(install
 (section <section>)
 (files  <filenames>)
 <optional-fields>)
```

`<section>` is the installation section, as described in the opam manual. The following sections are available:

- `lib`
- `lib_root`
- `libexec`
- `libexec_root`
- `bin`
- `sbin`
- `toplevel`
- `share`
- `share_root`
- `etc`
- `doc`
- `stublibs`
- `man`
- `misc`

`<files>` is the list of files to install. Each element in the list must be either a literal filename or a S-expression of the form:

```
(<filename> as <destination>)
```

where `<destination>` describe how the file will be installed. For instance, to install a file `mylib.el` as `emacs/site-lisp/mylib.el` in the `share_root` section:

```
(install
 (section share_root)
 (files (mylib.el as emacs/site-lisp/mylib.el)))
```

`<optional-fields>` are:

- `(package <name>)`. If there are no ambiguities, you can omit this field. Otherwise you need it to specify which package these files are part of. The package is not ambiguous when the first parent directory to contain a `<package>.opam` file contains exactly one `<package>.opam` file

### Handling of the .exe extension on Windows

Under Microsoft Windows, executables must be suffixed with `.exe`. Dune tries to make sure that executables are always installed with this extension on Windows.

More precisely, when installing a file via an `(install ...)` stanza, if the source file has extension `.exe` or `.bc`, then dune implicitly adds the `.exe` extension to the destination, if not already present.

#### 5.1.11 copy\_files

The `copy_files` and `copy_files#` stanzas allow to specify that files from another directory could be copied if needed to the current directory.

The syntax is as follows:

```
(copy_files <glob>)
```

`<glob>` represents the set of files to copy, see the [glob](#) for details.

The difference between `copy_files` and `copy_files#` is the same as the difference between the `copy` and `copy#` action. See the [User actions](#) section for more details.

#### 5.1.12 include

The `include` stanza allows to include the contents of another file into the current dune file. Currently, the included file cannot be generated and must be present in the source tree. This feature is intended to be used in conjunction with promotion, when parts of a dune file are to be generated.

For instance:

```
(include dune.inc)

(rule (with-stdout-to dune.inc.gen (run ./gen-dune.exe)))

(alias
 (name runtest)
 (action (diff dune.inc dune.inc.gen)))
```

With this dune file, running dune as follow will replace the `dune.inc` file in the source tree by the generated one:

```
$ dune build @runtest --auto-promote
```



### 5.1.13 tests

The `tests` stanza allows one to easily define multiple tests. For example we can define two tests at once with:

```
(tests
 (names mytest expect_test)
 <optional fields>)
```

This will define an executable named `mytest.exe` that will be executed as part of the `runtest` alias. If the directory also contains an `expect_test.expected` file, then `expect_test` will be used to define an expect test. That is, the test will be executed and its output will be compared to `expect_test.expected`.

The optional fields that are supported are a subset of the alias and executables fields. In particular, all fields except for `public_names` are supported from the `executables` stanza. Alias fields apart from `name` and `action` are allowed.

### 5.1.14 test

The `test` stanza is the singular form of `tests`. The only difference is that it's of the form:

```
(test
 (name foo)
 <optional fields>)
```

where the `name` field is singular. The same optional fields are supported.

### 5.1.15 env

The `env` stanza allows to modify the environment. The syntax is as follow:

```
(env
 (<profile1> <settings1>)
 (<profile2> <settings2>)
 ...
 (<profilen> <settingsn>))
```

The first form (`<profile> <settings>`) that correspond to the selected build profile will be used to modify the environment in this directory. You can use `_` to match any build profile.

Currently `<settings>` can be any OCaml flags field, see *OCaml flags* for more details.

### 5.1.16 ignored\_subdirs

The `ignored_subdirs` stanza allows to tell Dune to ignore one or more sub-directories. The syntax is as follow:

```
(ignored_subdirs (<sub-dir1> <sub-dir2> ...))
```

A directory that is ignored will not be eagerly scanned by Dune. Any `dune` or other special files in it won't be interpreted either and will be treated as raw data. It is however possible to depend on files inside ignored sub-directories.

### 5.1.17 include\_subdirs

The `include_subdirs` stanza is used to control how dune considers sub-directories of the current directory. The syntax is as follow:

```
(include_subdirs <mode>)
```

Where `<mode>` maybe be one of:

- `no`, the default
- `unqualified`

When the `include_subdirs` stanza is not present or `<mode>` is `no`, dune considers sub-directories as independent. When `<mode>` is `unqualified`, dune will assume that the sub-directories of the current directory are part of the same group of directories. In particular, dune will scan all these directories at once when looking for OCaml/Reason files. This allows you to split a library between several directories. `unqualified` means that modules in sub-directories are seen as if they were all in the same directory. In particular, you cannot have two modules with the same name in two different directories. It is planned to add a `qualified` mode in the future.

Note that sub-directories are included recursively, however the recursion will stop when encountering a sub-directory that contains another `include_subdirs` stanza. Additionally, it is not allowed for a sub-directory of a directory with `(include_subdirs <x>)` where `<x>` is not `no` to contain one of the following stanzas:

- `library`
- `executable(s)`
- `test(s)`

## 5.2 Common items

### 5.2.1 Ordered set language

A few fields takes as argument an ordered set and can be specified using a small DSL.

This DSL is interpreted by dune into an ordered set of strings using the following rules:

- `:standard` denotes the standard value of the field when it is absent
- an atom not starting with a `:` is a singleton containing only this atom
- a list of sets is the concatenation of its inner sets
- `(<sets1> \ <sets2>)` is the set composed of elements of `<sets1>` that do not appear in `<sets2>`

In addition, some fields support the inclusion of an external file using the syntax `(:include <filename>)`. This is useful for instance when you need to run a script to figure out some compilation flags. `<filename>` is expected to contain a single S-expression and cannot contain `(:include ...)` forms.

Note that inside an ordered set, the first element of a list cannot be an atom except if it starts with `-` or `:`. The reason for this is that we are planning to add simple programmatic features in the futures so that one may write:

```
(flags (if (>= %{ocaml_version}} 4.06) ...))
```

This restriction will allow to add this feature without introducing a breaking changes. If you want to write a list where the first element doesn't start by `-`, you can simply quote it: `("x" y z)`.

Most fields using the ordered set language also support *Variables expansion*. Variables are expanded after the set language is interpreted.

## 5.2.2 Boolean Language

The boolean language allows the user to define simple boolean expressions that dune can evaluate. Here's a semi formal specification of the language:

```
op := '=' | '<' | '>' | '<>' | '>=' | '<='

expr := (and <expr>+)
       | (or <expr>+)
       | (<op> <template> <template>)
       | <template>
```

After an expression is evaluated, it must be exactly the string `true` or `false` to be considered as a boolean. Any other value will be treated as an error.

Here's a simple example of a condition that expresses running on OSX and having an flambda compiler with the help of variable expansion:

```
(and ${ocamlc-config:flambda} (= ${ocamlc-config:system} macosx))
```

## 5.2.3 Variables expansion

Some fields can contains variables of the form `${var}` that are expanded by dune.

Dune supports the following variables:

- **project\_root** is the root of the current project. It is typically the toplevel directory of your project and as long as you have a `dune-project` file there, `project_root` is independent of the workspace configuration
- `workspace_root` is the root of the current workspace. Note that the value of `workspace_root` is not constant and depends on whether your project is vendored or not
- `CC` is the C compiler command line (list made of the compiler name followed by its flags) that was used to compile OCaml in the current build context
- `CXX` is the C++ compiler command line being used in the current build context
- `ocaml_bin` is the path where `ocamlc` lives
- `ocaml` is the `ocaml` binary
- `ocamlc` is the `ocamlc` binary
- `ocamlopt` is the `ocamlopt` binary
- `ocaml_version` is the version of the compiler used in the current build context
- `ocaml_where` is the output of `ocamlc -where`
- `arch_sixtyfour` is `true` if using a compiler targeting a 64 bit architecture and `false` otherwise
- `null` is `/dev/null` on Unix or `nul` on Windows
- `ext_obj`, `ext_asm`, `ext_lib`, `ext_dll` and `ext_exe` are the file extension used for various artifacts
- `ocaml-config:v` for every variable `v` in the output of `ocamlc -config`. Note that dune processes the output of `ocamlc -config` in order to make it a bit more stable across versions, so the exact set of variables accessible this way might not be exactly the same as what you can see in the output of `ocamlc -config`. In particular, variables added in new versions of OCaml needs to be registered in dune before they can be used
- `profile` the profile selected via `--profile`

In addition, (`action ...`) fields support the following special variables:

- `targets` expands to the list of target
- `deps` expands to the list of dependencies
- `^` expands to the list of dependencies, separated by spaces
- `dep:<path>` expands to `<path>` (and adds `<path>` as a dependency of the action)
- `exe:<path>` is the same as `<path>`, except when cross-compiling, in which case it will expand to `<path>` from the host build context
- `bin:<program>` expands to a path to `program`. If `program` is installed by a package in the workspace (see *install* stanzas), the locally built binary will be used, otherwise it will be searched in the `PATH` of the current build context. Note that `(run %{bin:program} ...)` and `(run program ...)` behave in the same way. `%{bin:...}` is only necessary when you are using `(bash ...)` or `(system ...)`
- `lib:<public-library-name>:<file>` expands to a path to file `<file>` of library `<public-library-name>`. If `<public-library-name>` is available in the current workspace, the local file will be used, otherwise the one from the installed world will be used
- `libexec:<public-library-name>:<file>` is the same as `lib:...`  except when cross-compiling, in which case it will expand to the file from the host build context
- `lib-available:<library-name>` expands to `true` or `false` depending on whether the library is available or not. A library is available iff at least one of the following condition holds:
  - it is part the installed worlds
  - it is available locally and is not optional
  - it is available locally and all its library dependencies are available
- `version:<package>` expands to the version of the given package. Note that this is only supported for packages that are being defined in the current scope
- `read:<path>` expands to the contents of the given file
- `read-lines:<path>` expands to the list of lines in the given file
- `read-strings:<path>` expands to the list of lines in the given file, unescaped using OCaml lexical convention

The `%<kind>:...}` forms are what allows you to write custom rules that work transparently whether things are installed or not.

Note that aliases are ignored by `%{deps}`

The intent of this last form is to reliably read a list of strings generated by an OCaml program via:

```
List.iter (fun s -> print_string (String.escaped s)) l
```

### 1. Expansion of lists

Forms that expands to list of items, such as `%{cc}`, `%{deps}`, `%{targets}` or `%{read-lines:...}`, are suitable to be used in, say, `(run <prog> <arguments>)`. For instance in:

```
(run foo %{deps})
```

if there are two dependencies `a` and `b`, the produced command will be equivalent to the shell command:

```
$ foo "a" "b"
```

If you want the two dependencies to be passed as a single argument, you have to quote the variable as in:

```
(run foo "%{deps}")
```

which is equivalent to the following shell command:

```
$ foo "a b"
```

(the items of the list are concatenated with space). Note that, since `%{deps}` is a list of items, the first one may be used as a program name, for instance:

```
(rule
 (targets result.txt)
 (deps    foo.exe (glob_files *.txt))
 (action  (run %{deps})))
```

Here is another example:

```
(rule
 (targets foo.exe)
 (deps    foo.c)
 (action  (run %{cc} -o %{targets} %{deps} -lfoolib)))
```

## 5.2.4 Library dependencies

Dependencies on libraries are specified using `(libraries ...)` fields in `library` and `executables` stanzas.

For libraries defined in the current scope, you can use either the real name or the public name. For libraries that are part of the installed world, or for libraries that are part of the current workspace but in another scope, you need to use the public name. For instance: `(libraries base re)`.

When resolving libraries, libraries that are part of the workspace are always preferred to ones that are part of the installed world.

### Alternative dependencies

In addition to direct dependencies you can specify alternative dependencies. This is described in the [Alternative dependencies](#) section

It is sometimes the case that one wants to not depend on a specific library, but instead on whatever is already installed. For instance to use a different backend depending on the target.

Dune allows this by using a `(select ... from ...)` form inside the list of library dependencies.

Select forms are specified as follows:

```
(select <target-filename> from
 (<literals> -> <filename>)
 (<literals> -> <filename>)
 ...)
```

`<literals>` are lists of literals, where each literal is one of:

- `<library-name>`, which will evaluate to true if `<library-name>` is available, either in the workspace or in the installed world
- `!<library-name>`, which will evaluate to true if `<library-name>` is not available in the workspace or in the installed world

When evaluating a select form, dune will create `<target-filename>` by copying the file given by the first (`<literals> -> <filename>`) case where all the literals evaluate to true. It is an error if none of the clauses are selectable. You can add a fallback by adding a clause of the form (`-> <file>`) at the end of the list.

### 5.2.5 Preprocessing specification

Dune accepts three kinds of preprocessing:

- `no_preprocessing`, meaning that files are given as it to the compiler, this is the default
- `(action <action>)` to preprocess files using the given action
- `(pps <ppx-rewriters-and-flags>)` to preprocess files using the given list of ppx rewriters
- **`(staged_pps <ppx-rewriters-and-flags>)` is similar to `“(pps ...)”` but behave slightly differently and is needed for certain ppx rewriters (see below for details)**

Dune normally assumes that the compilation pipeline is sequenced as follow:

- code generation (including preprocessing)
- dependency analysis
- compilation

Dune uses this fact to optimize the pipeline and in particular share the result of code generation and preprocessing between the dependency analysis and compilation phases. However, some specific code generators or preprocessors require feedback from the compilation phase. As a result they must be applied in stages as follows:

- first stage of code generation
- dependency analysis
- second step of code generation in parallel with compilation

This is the case for ppx rewriters using the OCaml typer for instance. When using such ppx rewriters, you must use `staged_pps` instead of `pps` in order to force Dune to use the second pipeline, which is slower but necessary in this case.

#### Preprocessing with actions

`<action>` uses the same DSL as described in the *User actions* section, and for the same reason given in that section, it will be executed from the root of the current build context. It is expected to be an action that reads the file given as only dependency named `input-file` and outputs the preprocessed file on its standard output.

More precisely, `(preprocess (action <action>))` acts as if you had setup a rule for every file of the form:

```
(rule
  (targets file.pp.ml)
  (deps    file.ml)
  (action  (with-stdout-to %{targets}
            (chdir %{workspace_root} <action>))))
```

The equivalent of a `-pp <command>` option passed to the OCaml compiler is `(system "<command> %{input-file")`.

## Preprocessing with ppx rewriters

`<ppx-rewriters-and-flags>` is expected to be a sequence where each element is either a command line flag if starting with a `-` or the name of a library. If you want to pass command line flags that do not start with a `-`, you can separate library names from flags using `--`. So for instance from the following `preprocess` field:

```
(preprocess (pps ppx1 -foo ppx2 -- -bar 42))
```

The list of libraries will be `ppx1` and `ppx2` and the command line arguments will be: `-foo -bar 42`.

Libraries listed here should be libraries implementing an OCaml AST rewriter and registering themselves using the `ocaml-migrate-parsetree.driver` API.

Dune will build a single executable by linking all these libraries and their dependencies. Note that it is important that all these libraries are linked with `-linkall`. Dune automatically uses `-linkall` when the `(kind ...)` field is set to `ppx_rewriter` or `ppx_deriver`.

## Per module preprocessing specification

By default a preprocessing specification will apply to all modules in the library/set of executables. It is possible to select the preprocessing on a module-by-module basis by using the following syntax:

```
(preprocess (per_module
  (<spec1> (<module-list1>))
  (<spec2> (<module-list2>))
  ...))
```

Where `<spec1>`, `<spec2>`, ... are preprocessing specifications and `<module-list1>`, `<module-list2>`, ... are list of module names.

For instance:

```
(preprocess (per_module
  (((action (run ./pp.sh X=1 %{input-file})) (foo bar)))
  (((action (run ./pp.sh X=2 %{input-file})) (baz)))))
```

## 5.2.6 Dependency specification

Dependencies in dune files can be specified using one of the following syntax:

- `(:name <dependencies>)` will bind the the list of dependencies to the name variable. This variable will be available as `%{name}` in actions.
- `(file <filename>)` or simply `<filename>`: depend on this file
- `(alias <alias-name>)`: depend on the construction of this alias, for instance: `(alias src/runtest)`
- `(alias_rec <alias-name>)`: depend on the construction of this alias recursively in all children directories wherever it is defined. For instance: `(alias_rec src/runtest)` might depend on `(alias src/runtest)`, `(alias src/foo/bar/runtest)`,...
- `(glob_files <glob>)`: depend on all files matched by `<glob>`, see the *glob* for details
- `(source_tree <dir>)`: depend on all source files in the subtree with root `<dir>`

- `(universe)`: depend on everything in the universe. This is for cases where dependencies are too hard to specify. Note that dune will not be able to cache the result of actions that depend on the universe. In any case, this is only for dependencies in the installed world, you must still specify all dependencies that come from the workspace.
- `(package <pkg>)` depend on all files installed by `<package>`, as well as on the transitive package dependencies of `<package>`. This can be used to test a command against the files that will be installed

In all these cases, the argument supports *Variables expansion*.

### Named Dependencies

dune allows a user to organize dependency lists by naming them. The user is allowed to assign a group of dependencies a name that can later be referred to in actions (like the `%{deps}` and `%{targets}` built in variables).

One instance where this is useful is for naming globs. Here's an example of an imaginary bundle command:

```
(rule
 (targets archive.tar)
 (deps
  index.html
  (:css (glob_files *.css))
  (:js foo.js bar.js)
  (:img (glob_files *.png) (glob_files *.jpg)))
 (action
  (run %{bin:bundle} index.html -css %{css} -js %{js} -img %{img} -o %{targets})))
```

Note that such named dependency list can also include unnamed dependencies (like `index.html` in the example above). Also, such user defined names will shadow built in variables. So `(:workspace_root x)` will shadow the built in `%{workspace_root}` variable.

### Glob

You can use globs to declare dependencies on a set of files. Note that globs will match files that exist in the source tree as well as buildable targets, so for instance you can depend on `*.cmi`.

Currently dune only support globbing files in a single directory. And in particular the glob is interpreted as follows:

- anything before the last `/` is taken as a literal path
- anything after the last `/`, or everything if the glob contains no `/`, is interpreted using the glob syntax

The glob syntax is interpreted as follows:

- `\<char>` matches exactly `<char>`, even if it is a special character (`*`, `?`, ...)
- `*` matches any sequence of characters, except if it comes first in which case it matches any character that is not `.` followed by anything
- `**` matches any character that is not `.` followed by anything, except if it comes first in which case it matches anything
- `?` matches any single character
- `[<set>]` matches any character that is part of `<set>`
- `[!<set>]` matches any character that is not part of `<set>`
- `{<glob1>, <glob2>, ..., <globn>}` matches any string that is matched by one of `<glob1>`, `<glob2>`, ...



## 5.2.7 OCaml flags

In `library`, `executable`, `executables` and `env` stanzas, you can specify OCaml compilation flags using the following fields:

- `(flags <flags>)` to specify flags passed to both `ocamlc` and `ocamlopt`
- `(ocamlc_flags <flags>)` to specify flags passed to `ocamlc` only
- `(ocamlopt_flags <flags>)` to specify flags passed to `ocamlopt` only

For all these fields, `<flags>` is specified in the *Ordered set language*. These fields all support `(:include ...)` forms.

The default value for `(flags ...)` is taken from the environment, as a result it is recommended to write `(flags ...)` fields as follows:

```
(flags (:standard <my options>))
```

## 5.2.8 js\_of\_ocaml

A `js_of_ocaml` field exists in `executable` and `libraries` stanzas that allows one to customize options relevant to `jsoc`.

## 5.2.9 User actions

`(action ...)` fields describe user actions.

User actions are always run from the same subdirectory of the current build context as the dune file they are defined in. So for instance an action defined in `src/foo/dune` will be run from `_build/<context>/src/foo`.

The argument of `(action ...)` fields is a small DSL that is interpreted by dune directly and doesn't require an external shell. All atoms in the DSL support *Variables expansion*. Moreover, you don't need to specify dependencies explicitly for the special `%{<kind>:...}` forms, these are recognized and automatically handled by dune.

The DSL is currently quite limited, so if you want to do something complicated it is recommended to write a small OCaml program and use the DSL to invoke it. You can use `shexp` to write portable scripts or `configurator` for configuration related tasks.

The following constructions are available:

- `(run <prog> <args>)` to execute a program. `<prog>` is resolved locally if it is available in the current workspace, otherwise it is resolved using the `PATH`
- `(chdir <dir> <DSL>)` to change the current directory
- `(setenv <var> <value> <DSL>)` to set an environment variable
- `(with-<outputs>-to <file> <DSL>)` to redirect the output to a file, where `<outputs>` is one of: `stdout`, `stderr` or `outputs` (for both `stdout` and `stderr`)
- `(ignore-<outputs> <DSL>)` to ignore the output, where `<outputs>` is one of: `stdout`, `stderr` or `outputs`
- `(progn <DSL>...)` to execute several commands in sequence
- `(echo <string>)` to output a string on `stdout`
- `(write-file <file> <string>)` writes `<string>` to `<file>`
- `(cat <file>)` to print the contents of a file to `stdout`

- `(copy <src> <dst>)` to copy a file
- `(copy# <src> <dst>)` to copy a file and add a line directive at the beginning
- `(system <cmd>)` to execute a command using the system shell: `sh` on Unix and `cmd` on Windows
- `(bash <cmd>)` to execute a command using `/bin/bash`. This is obviously not very portable
- `(diff <file1> <file2>)` is similar to `(run diff <file1> <file2>)` but is better and allows promotion. See *Diffing and promotion* for more details
- `(diff? <file1> <file2>)` is the same as `(diff <file1> <file2>)` except that it is ignored when `<file1>` or `<file2>` doesn't exist
- `(cmp <file1> <file2>)` is similar to `(run cmp <file1> <file2>)` but allows promotion. See *Diffing and promotion* for more details

As mentioned `copy#` inserts a line directive at the beginning of the destination file. More precisely, it inserts the following line:

```
# 1 "<source file name>"
```

Most languages recognize such lines and update their current location, in order to report errors in the original file rather than the copy. This is important as the copy exists only under the `_build` directory and in order for editors to jump to errors when parsing the output of the build system, errors must point to files that exist in the source tree. In the beta versions of dune, `copy#` was called `copy-and-add-line-directive`. However, most of the time one wants this behavior rather than a bare copy, so it was renamed to something shorter.

Note: expansion of the special `%{<kind>:...}` is done relative to the current working directory of the part of the DSL being executed. So for instance if you have this action in a `src/foo/dune`:

```
(action (chdir ../../.. (echo %{path:dune})))
```

Then `%{path:dune}` will expand to `src/foo/dune`. When you run various tools, they often use the filename given on the command line in error messages. As a result, if you execute the command from the original directory, it will only see the basename.

To understand why this is important, let's consider this dune file living in `src/foo`:

```
(rule
 (targets blah.ml)
 (deps    blah.mll)
 (action  (run ocamllex -o %{targets} %{deps})))
```

Here the command that will be executed is:

```
ocamllex -o blah.ml blah.mll
```

And it will be executed in `_build/<context>/src/foo`. As a result, if there is an error in the generated `blah.ml` file it will be reported as:

```
File "blah.ml", line 42, characters 5-10:
Error: ...
```

Which can be a problem as your editor might think that `blah.ml` is at the root of your project. What you should write instead is:

```
(rule
 (targets blah.ml)
```

(continues on next page)

(continued from previous page)

```
(deps   blah.ml1)
(action (chdir %{workspace_root} (run ocamllex -o %{targets} %{deps}))))
```

## 5.2.10 Locks

Given two rules that are independent, dune will assume that their associated actions can be run concurrently. Two rules are considered independent if none of them depend on the other, either directly or through a chain of dependencies. This basic assumption allows to parallelize the build.

However, it is sometimes the case that two independent rules cannot be executed concurrently. For instance this can happen for more complicated tests. In order to prevent dune from running the actions at the same time, you can specify that both actions take the same lock:

```
(alias
 (name  runtest)
 (deps  foo)
 (locks m)
 (action (run test.exe %{deps})))

(alias
 (name  runtest)
 (deps  bar)
 (locks m)
 (action (run test.exe %{deps})))
```

Dune will make sure that the executions of `test.exe foo` and `test.exe bar` are serialized.

Although they don't live in the filesystem, lock names are interpreted as file names. So for instance `(with-lock m ...)` in `src/dune` and `(with-lock ../src/m)` in `test/dune` refer to the same lock.

Note also that locks are per build context. So if your workspace has two build contexts setup, the same rule might still be executed concurrently between the two build contexts. If you want a lock that is global to all build contexts, simply use an absolute filename:

```
(alias
 (name  runtest)
 (deps  foo)
 (locks /tcp-port/1042)
 (action (run test.exe %{deps})))
```

## 5.2.11 Diffing and promotion

`(diff <file1> <file2>)` is very similar to `(run diff <file1> <file2>)`. In particular it behaves in the same way:

- when `<file1>` and `<file2>` are equal, it does nothing
- when they are not, the differences are shown and the action fails

However, it is different for the following reason:

- the exact command used to diff files can be configured via the `--diff-command` command line argument. Note that it is only called when the files are not byte equals
- by default, it will use `patdiff` if it is installed. `patdiff` is a better diffing program. You can install it via `opam` with:

```
$ opam install patdiff
```

- on Windows, both `(diff a b)` and `(diff? a b)` normalize the end of lines before comparing the files
- since `(diff a b)` is a builtin action, dune knows that `a` and `b` are needed and so you don't need to specify them explicitly as dependencies
- you can use `(diff? a b)` after a command that might or might not produce `b`. For cases where commands optionally produce a *corrected* file
- it allows promotion. See below

Note that `(cmp a b)` does no end of lines normalization and doesn't print a diff when the files differ. `cmp` is meant to be used with binary files.

### Promotion

Whenever an action `(diff <file1> <file2>)` or `(diff? <file1> <file2>)` fails because the two files are different, dune allows you to promote `<file2>` as `<file1>` if `<file1>` is a source file and `<file2>` is a generated file.

More precisely, let's consider the following dune file:

```
(rule
  (with-stdout-to data.out (run ./test.exe)))

(alias
  (name   runtest)
  (action (diff data.expected data.out)))
```

Where `data.expected` is a file committed in the source repository. You can use the following workflow to update your test:

- update the code of your test
- run `dune runtest`. The diff action will fail and a diff will be printed
- check the diff to make sure it is what you expect
- run `dune promote`. This will copy the generated `data.out` file to `data.expected` directly in the source tree

You can also use `dune runtest --auto-promote` which will automatically do the promotion.

## 5.3 OCaml syntax

If a dune file starts with `(* -*- tuareg -*- *)`, then it is interpreted as an OCaml script that generates the dune file as described in the rest of this section. The code in the script will have access to a `Jbuild_plugin` module containing details about the build context it is executed in.

The OCaml syntax gives you an escape hatch for when the S-expression syntax is not enough. It is not clear whether the OCaml syntax will be supported in the long term as it doesn't work well with incremental builds. It is possible that it will be replaced by just an `include` stanza where one can include a generated file.

Consequently **you must not** build complex systems based on it.

---

## Writing and running tests

---

Dune tries to streamline the testing story as much as possible, so that you can focus on the tests themselves and not bother with setting up with various test frameworks.

In this section, we will explain the workflow to deal with tests in dune. In particular we will see how to run the testsuite of a project, how to describe your tests to dune and how to promote tests result as expectation.

We distinguish two kinds of tests: inline tests and custom tests. Inline tests are usually written directly inside the ml files of a library. They are the easiest to work with and usually requires nothing more than writing `(inline_tests)` inside your library stanza. Custom tests consist on executing an executable and sometimes do something afterwards, such as diffing its output.

### 6.1 Running tests

Whatever the tests of a project are, the usual way to run tests with dune is to call `dune runtest` from the shell. This will run all the tests defined in the current directory and any sub-directory recursively. You can also pass a directory argument to run the tests from a sub-tree. For instance `dune runtest test` will only run the tests from the `test` directory and any sub-directory of `test` recursively.

Note that in any case, `dune runtest` is simply a short-hand for building the `runtest` alias, so you can always ask dune to run the tests in conjunction with other targets by passing `@runtest` to `dune build`. For instance:

```
$ dune build @install @runtest
$ dune build @install @test/runtest
```

### 6.2 Inline tests

There are several inline tests framework available for OCaml, such as `ppx_inline_test` and `qtest`. We will use `ppx_inline_test` as an example as at the time of writing this document it has the necessary setup to be used with dune out of the box.

`ppx_inline_test` allows to write tests directly inside ml files as follow:

```
let rec fact n = if n = 1 then 1 else n * fact (n - 1)

let%test _ = fact 5 = 120
```

The file has to be preprocessed with the `ppx_inline_test` ppx rewriter, so for instance the `jbuild` file might look like this:

```
(library
 (name foo)
 (preprocess (pps ppx_inline_test)))
```

In order to instruct dune that our library contains inline tests, all we have to do is add an `inline_tests` field:

```
(library
 (name foo)
 (inline_tests)
 (preprocess (pps ppx_inline_test)))
```

We can now build and execute this test by running `dune runtest`. For instance, if we make the test fail by replacing 120 by 0 we get:

```
$ dune runtest
[...]
File "src/fact.ml", line 3, characters 0-25: <<(fact 5) = 0>> is false.

FAILED 1 / 1 tests
```

Note that in this case Jbuild knew how to build and run the tests without any special configuration. This is because `ppx_inline_test` defines an inline tests backend and it is used by the library. Some other frameworks, such as `qtest` don't have any special library or ppx rewriter. To use such a framework, you must tell dune about it since it cannot guess it. You can do that by adding a `backend` field:

```
(library
 (name foo)
 (inline_tests (backend qtest)))
```

### 6.2.1 Inline expectation tests

Inline expectation tests are a special case of inline tests where you write a bit of OCaml code that prints something followed by what you expect this code to print. For instance, using `ppx_expect`:

```
let%expect_test _ =
  print_endline "Hello, world!";
  [%expect{|
    Hello, world!
  |}]
```

The test procedure consist of executing the OCaml code and replacing the contents of the `[%expect]` extension point by the real output. You then get a new file that you can compare to the original source file. Expectation tests are a neat way to write tests as the following test elements are clearly identified:

- the code of the test
- the test expectation
- the test outcome

You can have a look at [this blog post](#) to find out more about expectation tests. To dune, the workflow for expectation tests is always as follows:

- write the test with some empty expect nodes in it
- run the tests
- check the suggested correction and promote it as the original source file if you are happy with it

Dune makes this workflow very easy, simply add `ppx_expect` to your list of `ppx` rewriters as follow:

```
(library
 (name foo)
 (inline_tests)
 (preprocess (pps ppx_expect)))
```

Then calling `dune runtest` will run these tests and in case of mismatch dune will print a diff of the original source file and the suggested correction. For instance:

```
$ dune runtest
[...]
-src/fact.ml
+src/fact.ml.corrected
File "src/fact.ml", line 5, characters 0-1:
let rec fact n = if n = 1 then 1 else n * fact (n - 1)

let%expect_test _ =
  print_int (fact 5);
- [%expect]
+ [%expect{| 120 |}]
```

In order to accept the correction, simply run:

```
$ dune promote
```

You can also make dune automatically accept the correction after running the tests by typing:

```
$ dune runtest --auto-promote
```

Finally, some editor integration is possible to make the editor do the promotion and make the workflow even smoother.

## 6.2.2 Specifying inline test dependencies

If your tests are reading files, you must say it to dune by adding a `deps` field to the `inline_tests` field. The argument of this `deps` field follows the usual *Dependency specification*. For instance:

```
(library
 (name foo)
 (inline_tests (deps data.txt))
 (preprocess (pps ppx_expect)))
```

## 6.2.3 Passing special arguments to the test runner

Under the hood, a test executable is built by dune. Depending on the backend used this runner might take useful command line arguments. You can specify such flags by using a `flags` field, such as:

```
(library
 (name foo)
 (inline_tests (flags (-foo bar)))
 (preprocess (pps ppx_expect)))
```

The argument of the `flags` field follows the *Ordered set language*.

## 6.2.4 Using additional libraries in the test runner

When tests are not part of the library code, it is possible that tests require additional libraries than the library being tested. This is the case with `qtest` as tests are written in comments. You can specify such libraries using a `libraries` field, such as:

```
(library
 (name foo)
 (inline_tests (backend qtest)
               (libraries bar)))
```

## 6.2.5 Defining your own inline test backend

If you are writing a test framework, or for specific cases, you might want to define your own inline tests backend. If your framework is naturally implemented by a library or ppx rewriter that the user must use when they want to write tests, then you should define this library has a backend. Otherwise simply create an empty library with the name you want to give for your backend.

In order to define a library as an inline tests backend, simply add an `inline_tests.backend` field to the library stanza. An inline tests backend is specified by three parameters:

1. How to create the test runner
2. How to build the test runner
3. How to run the test runner

These three parameters can be specified inside the `inline_tests.backend` field, which accepts the following fields:

```
(generate_runner <action>
 (runner_libraries (<ocaml-libraries>))
 (flags <flags>)
 (extends (<backends>))
```

For instance:

`<action>` follows the *User actions* specification. It describe an action that should be executed in the directory of libraries using this backend for their tests. It is expected that the action produces some OCaml code on its standard output. This code will constitute the test runner. The action can use the following additional variables:

- `#{library-name}` which is the name of the library being tested
- `#{impl-files}` which is the list of implementation files in the library, i.e. all the `.ml` and `.re` files
- `#{intf-files}` which is the list of interface files in the library, i.e. all the `.mli` and `.rei` files

The `runner_libraries` field specifies what OCaml libraries the test runner uses. For instance, if the `generate_runner` actions generates something like `My_test_framework.runtests ()`, the you should probably put `my_test_framework` in the `runner_libraries` field.



If your test runner needs specific flags, you should pass them in the `flags` field. You can use the `${library-name}` variable in this field.

Finally, a backend can be an extension of another backend. In this case you must specify by in the `extends` field. For instance, `ppx_expect` is an extension of `ppx_inline_test`. It is possible to use a backend with several extensions in a library, however there must be exactly one *root backend*, i.e. exactly one backend that is not an extension of another one.

When using a backend with extensions, the various fields are simply concatenated. The order in which they are concatenated is unspecified, however if a backend `b` extends of a backend `a`, then `a` will always come before `b`.

## Example of backend

In this example, we put tests in comments of the form:

```
(*TEST: assert (fact 5 = 120) *)
```

The backend for such a framework looks like this:

```
(library
  (name simple_tests)
  (inline_tests.backend
    (generate_runner (run sed "s/(\\*TEST:\\(.+\\)\\*)/let () = \\1;;/" %{impl-files}))
  ))
```

Now all you have to do is write `(inline_tests ((backend simple_tests)))` wherever you want to write such tests. Note that this is only an example, we do not recommend using `sed` in your build as this would cause portability problems.

## 6.3 Custom tests

We said in *Running tests* that to run tests dune simply builds the `runtest` alias. As a result, to define custom tests, you simply need to add an action to this alias in any directory. For instance if you have a binary `tests.exe` that you want to run as part of running your testsuite, simply add this to a `jbuild` file:

```
(alias
  (name runtest)
  (action (run ./tests.exe)))
```

Hence to define an a test a pair of alias and executable stanzas are required. To simplify this common pattern, dune provides a *tests* stanza to define multiple tests and their aliases at once:

```
(tests (names test1 test2))
```

### 6.3.1 Diffing the result

It is often the case that we want to compare the output of a test to some expected one. For that, dune offers the `diff` command, which in essence is the same as running the `diff` tool, except that it is more integrated in dune and especially with the `promote` command. For instance let's consider this test:

```
(rule
  (with-stdout-to tests.output (run ./tests.exe)))

(alias
  (name runtest)
  (action (diff tests.expected test.output)))
```

After having run `tests.exe` and dumping its output to `tests.output`, `dune` will compare the latter to `tests.expected`. In case of mismatch, `dune` will print a diff and then the `dune promote` command can be used to copy over the generated `test.output` file to `tests.expected` in the source tree.

Alternatively, the `tests` also supports this style of tests.

```
(tests (names tests))
```

Where `dune` expects a `tests.expected` file to exist to infer that this is an expect tests.

This provides a nice way of dealing with the usual *write code*, *run*, *promote* cycle of testing. For instance:

```
$ dune runtest
[...]
-tests.expected
+tests.output
File "tests.expected", line 1, characters 0-1:
-Hello, world!
+Good bye!
$ dune promote
Promoting _build/default/tests.output to tests.expected.
```

Note that if available, the diffing is done using the `patdiff` tool, which displays nicer looking diffs than the standard `diff` tool. You can change that by passing `--diff-command CMD` to `dune`.

---

## Generating Documentation

---

### 7.1 Prerequisites

Documentation in dune is done courtesy of the `odoc` tool. Therefore, to generate documentation in dune, you will need to install this tool. This should likely be done with `opam`:

```
$ opam install odoc
```

### 7.2 Writing Documentation

Documentation comments will be automatically extracted from your OCaml source files following the syntax described in the section `Text formatting of the OCaml manual`.

Additional documentation pages may be attached to a package can be attached using the *Documentation Stanza*.

### 7.3 Building Documentation

Building the documentation using the `@doc` alias. Hence, all that is required to generate documentation for your project is building this alias:

```
$ dune build @doc
```

An index page containing links to all the `opam` packages in your project can be found in:

```
$ open _build/default/_doc/_html/index.html
```

Documentation for private libraries may also be built with:

```
$ dune build @doc-private
```

But this libraries will not be in the main html listing above, since they do not belong to any particular package. But the generated html will still be found in `_build/default/_doc/_html/<library>`.

## 7.4 Documentation Stanza

Documentation pages will be automatically generated for from `.ml` and `.mli` files that include `ocaml doc` fragments. Additional manual pages may be attached to packages using the `documentation` stanza. These `.mld` files must contain text in the same syntax as `ocaml doc` comments.

```
(documentation (<optional-fields>))
```

Where `<optional-fields>` are:

- `(package <name>)` the package this documentation should be attached to. If this absent, `dune` will try to infer it based on the location of the stanza.
- `(mld_files <arg>)` where `<arg>` field follows the *Ordered set language*. This is a set of extension-less, `mld` file base names that are attached to the package. Where `:standard` refers to all the `.mld` files in the stanza's directory.

The `index.mld` file (specified as `index` in `mld_files`) is treated specially by `dune`. This will be the file used to generate the entry page for the package. This is the page that will be linked from the main package listing. If you omit writing an `index.mld`, `dune` will generate one with the entry modules for your package. But this generated will not be installed.

All `mld` files attached to a package will be included in the generated `.install` file for that package, and hence will be installed by `opam`.

### 7.4.1 Examples

This stanza use attach all the `.mld` files in the current directory in a project with a single package.

```
(documentation ())
```

This stanza will attach three `mld` files to package `foo`. The `mld` files should be named `foo.mld`, `bar.mld`, and `baz.mld`

```
(documentation
  ((package foo)
   (mld_files (foo bar baz))))
```

This stanza will attach all `mld` files excluding `wip.mld` in the current directory to the inferred package:

```
(documentation
  ((mld_files (:standard \ wip))))
```

This section describe usage of dune from the shell.

## 8.1 Finding the root

### 8.1.1 dune-workspace

The root of the current workspace is determined by looking up a `dune-workspace` or `dune-project` file in the current directory and parent directories.

`dune` prints out the root when starting if it is not the current directory:

```
$ dune runtest
Entering directory '/home/jdimino/code/dune'
...
```

More precisely, it will choose the outermost ancestor directory containing a `dune-workspace` file as root. For instance if you are in `/home/me/code/myproject/src`, then `dune` will look for all these files in order:

- `/dune-workspace`
- `/home/dune-workspace`
- `/home/me/dune-workspace`
- `/home/me/code/dune-workspace`
- `/home/me/code/myproject/dune-workspace`
- `/home/me/code/myproject/src/dune-workspace`

The first entry to match in this list will determine the root. In practice this means that if you nest your workspaces, `dune` will always use the outermost one.

In addition to determining the root, `dune` will read this file as to setup the configuration of the workspace unless the `--workspace` command line option is used. See the section [Workspace configuration](#) for the syntax of this file.

## 8.1.2 Current directory

If the previous rule doesn't apply, i.e. no ancestor directory has a file named `dune-workspace`, then the current directory will be used as root.

## 8.1.3 Forcing the root (for scripts)

You can pass the `--root` option to `dune` to select the root explicitly. This option is intended for scripts to disable the automatic lookup.

Note that when using the `--root` option, targets given on the command line will be interpreted relative to the given root, not relative to the current directory as this is normally the case.

## 8.2 Interpretation of targets

This section describes how `dune` interprets the targets given on the command line. When no targets are specified, `dune` builds the `default` alias, see *Default alias* for more details.

### 8.2.1 Resolution

All targets that `dune` knows how to build live in the `_build` directory. Although, some are sometimes copied to the source tree for the need of external tools. These includes:

- `.merlin` files
- `<package>.install` files

As a result, if you want to ask `dune` to produce a particular `.exe` file you would have to type:

```
$ dune build _build/default/bin/prog.exe
```

However, for convenience when a target on the command line doesn't start with `_build`, `dune` will expand it to the corresponding target in all the build contexts where it knows how to build it. When using `--verbose`, It prints out the actual set of targets when starting:

```
$ dune build bin/prog.exe --verbose
...
Actual targets:
- _build/default/bin/prog.exe
- _build/4.03.0/bin/prog.exe
- _build/4.04.0/bin/prog.exe
```

### 8.2.2 Aliases

Targets starting with a `@` are interpreted as aliases. For instance `@src/runtest` means the alias `runtest` in all descendant of `src` in all build contexts where it is defined. If you want to refer to a target starting with a `@`, simply write: `./@foo`.

To build and run the tests for a particular build context, use `@_build/default/runtest` instead.

So for instance:

- `dune build @_build/foo/runtest` will run the tests only for the `foo` build context

- `dune build @runtest` will run the tests for all build contexts

You can also build an alias non-recursively by using `@@` instead of `@`. For instance to run tests only from the current directory:

```
dune build @@runtest
```

### 8.2.3 Default alias

When no targets are given to `dune build`, it builds the special default alias. Effectively `dune build` is equivalent to:

```
dune build @@default
```

When a directory doesn't explicitly define what the default alias means via an *alias* stanza, the following implicit definition is assumed:

```
(alias
 (name default)
 (deps (alias_rec install)))
```

Which means that by default `dune build` will build everything that is installable.

## 8.3 Finding external libraries

When a library is not available in the workspace, `dune` will look it up in the installed world, and expect it to be already compiled.

It looks up external libraries using a specific list of search paths. A list of search paths is specific to a given build context and is determined as follow:

1. if the `ocamlfind` is present in the `PATH` of the context, use each line in the output of `ocamlfind printconf path` as a search path
2. otherwise, if `opam` is present in the `PATH`, use the outout of `opam config var lib`
3. otherwise, take the directory where `ocamlc` was found, and append `../lib` to it. For instance if `ocamlc` is found in `/usr/bin`, use `/usr/lib`

## 8.4 Running tests

There are two ways to run tests:

- `dune build @runtest`
- `dune runtest`

The two commands are equivalent. They will run all the tests defined in the current directory and its children recursively. You can also run the tests in a specific sub-directory and its children by using:

- `dune build @foo/bar/runtest`
- `dune runtest foo/bar`

## 8.5 Launching the Toplevel (REPL)

Dune supports launching a `utop` instance with locally defined libraries loaded.

```
$ dune utop <dir> -- <args>
```

Where `<dir>` is a directory containing a `dune` file defining all the libraries that will be loaded (using the `library` stanza). `<args>` will be passed as arguments to the `utop` command itself. For example, `dune utop lib --implicit-bindings` will start `utop` with the libraries defined in `lib` and implicit bindings for toplevel expressions.

### 8.5.1 Requirements & Limitations

- `utop` version `>= 2.0` is required for this to work.
- This subcommand only supports loading libraries. Executables aren't supported.
- Libraries that are dependencies of `utop` itself cannot be loaded. For example `Camomile`.
- Loading libraries that are defined in different directories into one `utop` instance isn't possible.

## 8.6 Restricting the set of packages

You can restrict the set of packages from your workspace that `dune` can see with the `--only-packages` option:

```
$ dune build --only-packages pkg1,pkg2,... @install
```

This option acts as if you went through all the `dune` files and commented out the stanzas referring to a package that is not in the list given to `dune`.

## 8.7 Invocation from opam

You should set the `build:` field of your `<package>.opam` file as follows:

```
build: [  
  ["dune" "subst"] {pinned}  
  ["dune" "build" "-p" name "-j" jobs]  
]
```

`-p pkg` is a shorthand for `--root . --only-packages pkg --profile release --default-target @install`. `-p` is the short version of `--for-release-of-packages`.

This has the following effects:

- it tells `dune` to build everything that is installable and to ignore packages other than `name` defined in your project
- it sets the root to prevent `dune` from looking it up
- it sets the build profile to `release`
- it uses whatever concurrency option `opam` provides
- it sets the default target to `@install` rather than `@@default`

Note that `name` and `jobs` are variables expanded by `opam`. `name` expands to the package name and `jobs` to the number of jobs available to build the package.



## 8.8 Tests

To setup the building and running of tests in opam, add this line to your `<package>.opam` file:

```
build-test: [["dune" "runtest" "-p" name "-j" jobs]]
```

## 8.9 Installation

Installing a package means copying the build artifacts from the build directory to the installed word.

When installing via opam, you don't need to worry about this step: dune generates a `<package>.install` file that opam will automatically read to handle installation.

However, when not using opam or doing local development, you can use dune to install the artifacts by hands. To do that, use the `install` command:

```
$ dune install [PACKAGE]...
```

without an argument, it will install all the packages available in the workspace. With a specific list of packages, it will only install these packages. If several build contexts are configured, the installation will be performed for all of them.

Note that `dune install` is a thin wrapper around the `opam-installer` tool, so you will need to install this tool in order to be able to use `dune install`.

### 8.9.1 Destination

The place where the build artifacts are copied, usually referred as **prefix**, is determined as follow for a given build context:

1. if an explicit `--prefix <path>` argument is passed, use this path
2. if opam is present in the `PATH` and is configured, use the output of `opam config var prefix`
3. otherwise, take the parent of the directory where `ocamlc` was found.

As an exception to this rule, library files might be copied to a different location. The reason for this is that they often need to be copied to a particular location for the various build system used in OCaml projects to find them and this location might be different from `<prefix>/lib` on some systems.

Historically, the location where to store OCaml library files was configured through `findlib` and the `ocamlfind` command line tool was used to both install these files and locate them. Many Linux distributions or other packaging systems are using this mechanism to setup where OCaml library files should be copied.

As a result, if none of `--libdir` and `--prefix` is passed to `dune install` and `ocamlfind` is present in the `PATH`, then library files will be copied to the directory reported by `ocamlfind printconf destdir`. This ensures that `dune install` can be used without opam. When using opam, `ocamlfind` is configured to point to the opam directory, so this rule makes no difference.

Note that `--prefix` and `--libdir` are only supported if a single build context is in use.

## 8.10 Workspace configuration

By default, a workspace has only one build context named `default` which correspond to the environment in which dune is run. You can define more contexts by writing a `dune-workspace` file.

You can point dune to an explicit `dune-workspace` file with the `--workspace` option. For instance it is good practice to write a `dune-workspace.dev` in your project with all the version of OCaml your projects support. This way developers can tests that the code builds with all version of OCaml by simply running:

```
$ dune build --workspace dune-workspace.dev @install @runtest
```

### 8.10.1 dune-workspace

The `dune-workspace` file uses the S-expression syntax. This is what a typical `dune-workspace` file looks like:

```
(lang dune 1.0)
(context (opam (switch 4.02.3)))
(context (opam (switch 4.03.0)))
(context (opam (switch 4.04.0)))
```

The rest of this section describe the stanzas available.

Note that an empty `dune-workspace` file is interpreted the same as one containing exactly:

```
(lang dune 1.0)
(context default)
```

This allows you to use an empty `dune-workspace` file to mark the root of your project.

#### profile

The build profile can be selected in the `dune-workspace` file by write a `(profile ...)` stanza. For instance:

```
(profile release)
```

Note that the command line option `--profile` has precedence over this stanza.

#### env

The `env` stanza can be used to set the base environment for all contexts in this workspace. This environment has the lowest precedence of all other `env` stanzas. The syntax for this stanza is the same dune's `env` stanza.

#### context

The `(context ...)` stanza declares a build context. The argument can be either `default` or `(default)` for the default build context or can be the description of an opam switch, as follows:

```
(context (opam (switch <opam-switch-name>)
              <optional-fields>))
```

`<optional-fields>` are:

- `(name <name>)` is the name of the subdirectory of `_build` where the artifacts for this build context will be stored
- `(root <opam-root>)` is the opam root. By default it will take the opam root defined by the environment in which dune is run which is usually `~/ .opam`
- `(merlin)` instructs dune to use this build context for merlin

- `(profile <profile>)` to set a different profile for a build context. This has precedence over the command line option `--profile`
- `(env <env>)` to set the environment for a particular context. This is of higher precedence than the toplevel `env` stanza in the workspace file. This field the same options as the `env` stanza.

Both `(default ...)` and `(opam ...)` accept a `targets` field in order to setup cross compilation. See *Cross Compilation* for more information.

Merlin reads compilation artifacts and it can only read the compilation artifacts of a single context. Usually, you should use the artifacts from the `default` context, and if you have the `(context default)` stanza in your `dune-workspace` file, that is the one dune will use.

For rare cases where this is not what you want, you can force dune to use a different build contexts for merlin by adding the field `(merlin)` to this context.

## 8.11 Distributing Projects

Dune provides support for building and installing your project. However it doesn't provide helpers for distributing it. It is recommended to use `dune-release` for this purpose.

The common defaults are that your projects include the following files:

- `README.md`
- `CHANGES.md`
- `LICENSE.md`

And that if your project contains several packages, then all the package names must be prefixed by the shortest one.

## 8.12 Watermarking

One of the feature `dune-release` provides is watermarking; it replaces various strings of the form `%%ID%%` in all files of your project before creating a release tarball or when the package is pinned by the user using `opam`.

This is especially interesting for the `VERSION` watermark, which gets replaced by the version obtained from the `vcs`. For instance if you are using `git`, `dune-release` invokes this command to find out the version:

```
$ git describe --always --dirty
1.0+beta9-79-g29e9b37
```

Projects using `dune` usually only need `dune-release` for creating and publishing releases. However they might still want to substitute the watermarks when the package is pinned by the user. To help with this, `dune` provides the `subst` sub-command.

## 8.13 dune subst

`dune subst` performs the same substitution `dune-release` does with the default configuration. i.e. calling `dune subst` at the root of your project will rewrite in place all the files in your project.

More precisely, it replaces all the following watermarks in source files:

- `NAME`, the name of the project
- `VERSION`, output of `git describe --always --dirty`

- `VERSION_NUM`, same as `VERSION` but with a potential leading `v` or `V` dropped
- `VCS_COMMIT_ID`, commit hash from the `vcs`
- `PKG_MAINTAINER`, contents of the `maintainer` field from the `opam` file
- `PKG_AUTHORS`, contents of the `authors` field from the `opam` file
- `PKG_HOMEPAGE`, contents of the `homepage` field from the `opam` file
- `PKG_ISSUES`, contents of the `issues` field from the `opam` file
- `PKG_DOC`, contents of the `doc` field from the `opam` file
- `PKG_LICENSE`, contents of the `license` field from the `opam` file
- `PKG_REPO`, contents of the `repo` field from the `opam` file

The name of the project is obtained by reading the `dune-project` file in the directory where `dune subst` is called. The `dune-project` file must exist and contain a valid `(name ...)` field.

Note that `dune subst` is meant to be called from the `opam` file and in particular behaves a bit different to other `dune` commands. In particular it doesn't try to detect the root of the workspace and must be called from the root of the project.

## 8.14 Custom Build Directory

By default `dune` places all build artifacts in the `_build` directory relative to the user's workspace. However, one can customize this directory by using the `--build-dir` flag or the `DUNE_BUILD_DIR` environment variable.

```
$ dune build --build-dir _build-foo

# this is equivalent to:
$ DUNE_BUILD_DIR=_build-foo dune build

# Absolute paths are also allowed
$ dune build --build-dir /tmp/build foo.exe
```

This section describes some details of dune for advanced users.

## 9.1 META file generation

Dune uses `META` files from the [findlib library manager](#) in order to interoperate with the rest of the world when installing libraries. It is able to generate them automatically. However, for the rare cases where you would need a specific `META` file, or to ease the transition of a project to dune, it is allowed to write/generate a specific one.

In order to do that, write or setup a rule to generate a `META.<package>.template` file in the same directory as the `<package>.opam` file. Dune will generate a `META.<package>` file from the `META.<package>.template` file by replacing lines of the form `# JBUILDER_GEN` by the contents of the `META` it would normally generate.

For instance if you want to extend the `META` file generated by dune you can write the following `META.foo.template` file:

```
# JBUILDER_GEN
blah = "..."
```

## 9.2 Findlib integration and limitations

Dune uses `META` files to support external libraries. However, it doesn't export the full power of findlib to the user, and especially it doesn't let the user specify *predicates*.

The reason for this limitation is that so far they haven't been needed, and adding full support for them would complicate things quite a lot. In particular, complex `META` files are often hand-written and the various features they offer are only available once the package is installed, which goes against the root ideas dune is built on.

In practice, dune interprets `META` files assuming the following set of predicates:

- `mt`: what this means is that using a library that can be used with or without threads with dune will force the threaded version

- `mt_posix`: forces the use of posix threads rather than VM threads. VM threads are deprecated and are likely to go away soon
- `ppx_driver`: when a library acts differently depending on whether it is linked as part of a driver or meant to add a `-ppx` argument to the compiler, choose the former behavior

### 9.3 Cross Compilation

Dune allows for cross compilation by defining build contexts with multiple targets. Targets are specified by adding a `targets` field to the definition of a build context.

`targets` takes a list of target name. It can be either:

- `native` which means using the native tools that can build binaries that run on the machine doing the build
- the name of an alternative toolchain

Note that at the moment, there is no official support for cross-compilation in OCaml. Dune supports the `opam-cross-x` repositories from the [ocaml-cross organization on github](#), such as:

- `opam-cross-windows`
- `opam-cross-android`
- `opam-cross-ios`

In particular:

- to build Windows binaries using `opam-cross-windows`, write `windows` in the list of targets
- to build Android binaries using `opam-cross-android`, write `android` in the list of targets
- to build IOS binaries using `opam-cross-ios`, write `ios` in the list of targets

For example, the following workspace file defines three different targets for the `default` build context:

```
(context (default (targets (native windows android))))
```

This configuration defines three build contexts:

- `default`
- `default.windows`
- `default.android`

Note that the `native` target is always implicitly added when not present. However, when implicitly added `dune build @install` will skip this context, i.e. `default` will only be used for building executables needed by the other contexts.

With such a setup, calling `dune build @install` will build all the packages three times.

Note that instead of writing a `dune-workspace` file, you can also use the `-x` command line option. Passing `-x foo` to `dune` without having a `dune-workspace` file is the same as writing the following `dune-workspace` file:

```
(context (default (targets (foo))))
```

If you have a `dune-workspace` and pass a `-x foo` option, `foo` will be added as target of all context stanzas.

### 9.3.1 How does it work?

In such a setup, binaries that need to be built and executed in the `default.windows` or `default.android` contexts as part of the build, will no longer be executed. Instead, all the binaries that will be executed will come from the `default` context. One consequence of this is that all preprocessing (ppx or otherwise) will be done using binaries built in the `default` context.

To clarify this with an example, let's assume that you have the following `src/dune` file:

```
(executable (name foo))
(rule (with-stdout-to blah (run ./foo.exe)))
```

When building `_build/default/src/blah`, dune will resolve `./foo.exe` to `_build/default/src/foo.exe` as expected. However, for `_build/default.windows/src/blah` dune will resolve `./foo.exe` to `_build/default/src/foo.exe`

Assuming that the right packages are installed or that your workspace has no external dependencies, dune will be able to cross-compile a given package without doing anything special.

Some packages might still have to be updated to support cross-compilation. For instance if the `foo.exe` program in the previous example was using `Sys.os_type`, it should instead take it as a command line argument:

```
(rule (with-stdout-to blah (run ./foo.exe -os-type ${os_type})))
```

## 9.4 Classical ppx

*classical ppx* refers to running ppx using the `-ppx` compiler option, which is composed using Findlib. Even though this is useful to run some (usually old) ppx's which don't support drivers, dune does not support preprocessing with ppx this way. but a workaround exists using the `ppxfind` tool.





# CHAPTER 10

---

## Configurator

---

Configurator is a small library designed to query features available on the system, in order to generate configuration for dune builds. Such generated configuration is usually in the form of command line flags, generated headers, stubs, but there are no limitations on this.

Configurator allows you to query for the following features:

- Variables defined in `ocamlc - config`,
- `pkg-config` flags for packages,
- Test features by compiling C code,
- Extract compile time information such as `#define` variables.

Configurator is designed to be cross compilation friendly and avoids `_running_` any compiled code to extract any of the information above.

Configurator started as an [independent library](#), but now lives in dune. You do not need to install anything to use configurator.

## 10.1 Usage

We'll describe configurator with a simple example. Everything else can be easily learned by studying [configurator's API](#).

To use configurator, we write an executable that will query the system using configurator's API and output a set of targets reflecting the results. For example:

```
module C = Configurator.V1

let clock_gettime_code = {}
#include <time.h>

int main()
```

(continues on next page)

(continued from previous page)

```

{
  struct timespec ts;
  clock_gettime(CLOCK_REALTIME, &ts);
  return 0;
}
|}

let () =
  C.main ~name:"foo" (fun c ->
    let has_clock_gettime = C.c_test c clock_gettime_code ~link_flags:["-lrt"] in

    C.C_define.gen_header_file c ~fname:"config.h"
    [ "HAS_CKOCK_GETTIME", Switch has_ckock_gettime ]);

```

Usually, the module above would be named `discover.ml`. The next step is to invoke it as an executable and tell dune about the targets that it produces:

```

(executable
  ((name discover)
   (libraries (dune.configurator))))

(rule
  ((targets (config.h))
   (action (run ./discover.exe))))

```

Another common pattern is to produce a flags file with configurator and then use this flag file using `:include`:

```

(library
  ((name mylib)
   (c_names (foo))
   (c_library_flags (:include (flags.sexp)))))

```

For this, generate the list of flags for your library — for example using `Configurator.V1.Pkg_config` — and then write them to a file, in the above example `flags.sexp`, with `Configurator.V1.write_flags "flags.sexp" flags`.

## 10.2 Upgrading from the old Configurator

The old configurator is the independent `configurator` opam package. It is deprecated and users are encouraged to migrate to dune's own configurator. The advantage of the transition include:

- No extra dependencies,
- No need to manually pass `-ocamlc` flag,
- New configurator is cross compilation compatible.

The following steps must be taken to transition from the old configurator:

- Mentions of the `configurator` opam package should be removed.
- The library name `configurator` should be changed `dune.configurator`.
- The `-ocamlc` flag in rules that run configurator scripts should be removed. This information is now passed automatically by dune.
- The new configurator API is versioned explicitly. The version that is compatible with old configurator is under the `V1` module. Hence, to transition one's code it's enough to add this module alias:

```
module Configurator = Configurator.V1
```



The basic form for defining `menhir` parsers (analogous to `ocamlyacc`) is:

```
(menhir
 (modules (<parser1> <parser2> ...)))
```

Modular parsers can be defined by adding a `merge_into` field. This correspond to the `--base` command line option of `menhir`. With this option, a single parser named `base_name` is generated.

```
(menhir
 ((merge_into <base_name>)
 (modules (<parser1> <parser2> ...)))
```

Extra flags can be passed to `menhir` using the `flags` flag:

```
(menhir
 ((flags (<option1> <option2> ...))
 (modules (<parser1> <parser2> ...)))
```



# CHAPTER 12

---

## js\_of\_ocaml

---

`js_of_ocaml` is a compiler from OCaml to JavaScript. The compiler works by translating OCaml bytecode to JS files. From here on, we'll abbreviate `js_of_ocaml` to `jsoc`. The compiler can be installed with `opam`:

```
$ opam install js_of_ocaml-compiler
```

## 12.1 Compiling to JS

Dune has full support building `jsoc` libraries and executables transparently. There's no need to customize or enable anything to compile ocaml libraries/executables to JS.

To build a JS executable, just define an executable as you would normally. Consider this example:

```
echo 'print_endline "hello from js"' > foo.ml
```

With the following dune file:

```
(executable (name foo))
```

And then request the `.js` target:

```
$ dune build ./foo.bc.js
$ node _build/default/foo.bc.js
hello from js
```

Similar targets are created for libraries, but we recommend sticking to the executable targets.

## 12.2 `js_of_ocaml` field

In library and executables stanzas, you can specify `js_of_ocaml` options using `(js_of_ocaml (<js_of_ocaml-options>))`.

`<js_of_ocaml-options>` are all optional:

- `(flags <flags>)` to specify flags passed to `js_of_ocaml`. This field supports `(:include ...)` forms
- `(javascript_files (<files-list>))` to specify `js_of_ocaml` JavaScript runtime files.

`<flags>` is specified in the *Ordered set language*.

The default value for `(flags ...)` depends on the selected build profile. The build profile `dev` (the default) will enable sourcemap and the pretty JavaScript output.

## 12.3 Separate Compilation

Dune supports two modes of compilation

- Direct compilation of a bytecode program to JavaScript. This mode allows `js_of_ocaml` to perform whole program deadcode elimination and whole program inlining.
- Separate compilation, where compilation units are compiled to JavaScript separately and then linked together. This mode is useful during development as it builds more quickly.

The separate compilation mode will be selected when the build profile is `dev`, which is the default. There is currently no other way to control this behaviour.



## 13.1 Why do many dune projects contain a Makefile?

Many dune projects contain a toplevel *Makefile*. It is often only there for convenience, for the following reasons:

1. there are many different build systems out there, all with a different CLI. If you have been hacking for a long time, the one true invocation you know is *make && make install*, possibly preceded by *./configure*
2. you often have a few common operations that are not part of the build and *make <blah>* is a good way to provide them
3. *make* is shorter to type than *dune build @install*

## 13.2 How to add a configure step to a dune project?

The [with-configure-step](#) example shows one way to do it which preserves composability; i.e. it doesn't require manually running *./configure* script when working on multiple projects at the same time.

## 13.3 Can I use topkg with dune?

It's possible using the [topkg-jbuilder](#) but it's not recommended. [dune-release](#) subsumes [topkg-jbuilder](#) and is specifically tailored to dune projects.

## 13.4 How do I publish my packages with dune?

Dune is just a build system and considers publishing outside of its scope. However, the [dune-release](#) project is specifically designed for releasing dune projects to opam. We recommend using tool for publishing dune packages.

## 13.5 here can I find some examples of projects using dune?

The [dune-universe](#) repository contains a snapshot of the latest versions of all opam packages depending on dune. It is therefore a useful reference to search through to find different approaches to constructing build rules.

### 14.1 mli only modules

These are supported, however using them might cause make it impossible for non-dune users to use your library. We tried to use them for some internal module generated by dune and it broke the build of projects not using dune:

<https://github.com/ocaml/dune/issues/567>

So, while they are supported, you should be careful where you use them. Using a *.ml* only module is still preferable.



Dune was initially called Jbuilder. Up to mid-2018, the package was still called *jbuilder* which only installed a *jbuilder* binary. This document explain how the migration to Dune will happen.

## 15.1 Timeline

The general idea is that the migration is gradual and existing Jbuilder projects don't need to be updated all at once. We encourage users to switch their development repositories and continue their usual release cycle. There is no need to re-release existing packages just to switch to Dune immediately.

The plan is as follows:

### 15.1.1 July 2018: release of Dune 1.0.0

First release of the opam package *dune*. The *jbuilder* package becomes a transitional package that depends on *dune*.

The *dune* package installs two binaries: *dune* and *jbuilder*. These two binaries are exactly the same and they work on both Jbuilder and Dune projects. Additionally they recognize both Jbuilder and Dune configuration files. The new Dune configuration files are described later in this document.

### 15.1.2 January 2019: deprecation of Jbuilder

At this point, the *jbuilder* binary emits a warning on every startup inviting users to switch to *dune*. When encountering *jbuild* or other Jbuilder configuration files, both binaries emit a warning. The rest is unchanged.

During this period, it makes sense for projects to do new releases just to switch to Dune if none of their existing releases is using Dune.

### 15.1.3 July 2019: support for Jbuilder is dropped

*jbuilder* is now a dummy executable that always exit with an error message on startup. *dune* no longer reads *jbuild* or other Jbuilder configuration files but still prints a warning when encountering them.

At this point, a conflict with newer versions of *dune* will be added to all opam packages that rely on the *jbuilder* binary or Jbuilder configuration files.

### 15.1.4 January 2020: the jbuilder binary goes away

The *dune* package no longer installs a *jbuilder* binary. The rest is unchanged.

### 15.1.5 Distant future

Once we are sure there are no more *jbuild* files out there, Dune will completely ignore *jbuild* and other Jbuilder configuration files.

## 15.2 Check List

This section is a concise list of migration tasks that will be required to transition from jbuilder to dune.

### 15.2.1 New configuration files

Until July 2019, *dune* will still read *jbuild* and other Jbuilder configuration files. There is no change in these files.

However, based on the experience acquired since the first release of Jbuilder, we made a few changes in the configuration files read by Dune. The most notable ones are the following:

- *jbuild* files are renamed simply *dune*
- projects now have a *dune-project* file at their root
- *jbuild-ignore* files are replaced by *ignored\_subdirs* stanzas in *dune* files
- *jbuild-workspace* are replaced by *dune-workspace* files
- *jbuild-workspace<suffix>* files no longer mean anything

Following are detailed explanation of the differences between the Jbuilder configuration files and the Dune ones.

### 15.2.2 dune-project files

These are a new kind of file. With Jbuilder, projects used to be identified by the presence of at least one *<package>.opam* file in a directory. This will still be supported until July 2019, however as Jbuilder evolved it became clear that we needed project files, so Dune introduces *dune-project* files to mark the root of projects.

Eventually, we are hoping that Dune will generate opam files. So users will only have to write a *dune-project* file.

The purpose of this file is to:

- delimit projects in larger workspaces
- set a few project-wide parameters, such as the name, the version of the Dune language in use or specification of extra features (plugins) used in the project

Eventually, for users who wish to do so it should be possible to centralize all the configuration of a project in this file.

### 15.2.3 dune files

These are the same as *jbuild* files.

### 15.2.4 dune-workspace

These are the same as *jbuild-workspace* files.

When looking for the root of the workspace, Jbuilder also looks for files whose name start with *jbuild-workspace*, such as *jbuild-workspace.in*. This rule will be kept until July 2019, however it is not preserved for *dune-workspace* files. I.e. a *dune-workspace.in* file means nothing.

This rule was only useful when we didn't have project files.

### 15.2.5 Variable Syntax

`#{foo}` and `$(foo)` are no longer valid variable syntax in dune files. Variables are defined as `%{foo}`. This change is done to simplify interoperability with bash commands which also use the `#{foo}` syntax.

### 15.2.6 (`files_recursively_in ..`) is removed

The `files_recursively_in` dependency specification is invalid in dune files. A *source\_tree* stanza has been introduced to reflect the actual function of this stanza.

### 15.2.7 Escape Sequences

Invalid escape sequences of the form `\x` where `x` is a character other than `[0-9]`, `x`, `n`, `r`, `t`, `b` are not allowed in dune files.

### 15.2.8 Comments Syntax

Block comments of the form `#| ... |#` and comments of the form `#;` are not supported in dune files.

### 15.2.9 Renamed Variables

All existing variables have been lowercased for consistency. Other variables have always been renamed. Refer to this table for details:

Jbuild	Dune
<code>#{@}</code>	<code>%{targets}</code>
<code>{^}</code>	<code>%{deps}</code>
<code>{path:file}</code>	<code>%{dep:file}</code>
<code>{SCOPE_ROOT}</code>	<code>%{project_root}</code>
<code>{ROOT}</code>	<code>%{workspace_root}</code>
<code>{findlib:..}</code>	<code>%{lib:..}</code>
<code>{CPP}</code>	<code>%{cpp}</code>
<code>{CC}</code>	<code>%{cc}</code>
<code>{CXX}</code>	<code>%{cxx}</code>
<code>{OCAML}</code>	<code>%{ocaml}</code>
<code>{OCAMLC}</code>	<code>%{ocamlc}</code>
<code>{OCAMLOPT}</code>	<code>%{ocamlopt}</code>
<code>{ARCH_SIXTYFOUR}</code>	<code>%{arch_sixtyfour}</code>
<code>{MAKE}</code>	<code>%{make}</code>

### 15.2.10 Removed Variables

`{path-no-dep:file}` and `{<}` have been removed.

A named dependency should be used instead of `{<}`. For instance the following jbuild file:

```
(alias
 (name  runtest)
 (deps  (input))
 (action (run ./test.exe %<))))
```

should be rewritten to the following dune file:

```
(alias
 (name  runtest)
 (deps  (:x input))
 (action (run ./test.exe %x))))
```

### 15.2.11 # JBUILDER\_GEN renamed

# DUNE\_GEN should be used instead of # JBUILDER\_GEN in META templates.