

The `build2` Build System

Copyright © 2014-2026 the build2 authors.

Permission is granted to copy, distribute and/or modify this document under the terms of the MIT License.

Revision 0.19, September 2026

This revision of the document describes the `build2` build system 0.19.x series.

Table of Contents

Preface	1
1 Introduction	1
1.1 Hello, World	2
1.2 Project Structure	7
1.3 Output Directories and Scopes	15
1.4 Operations	26
1.4.1 Configuring	27
1.4.2 Testing	30
1.4.3 Installing	35
1.4.4 Distributing	38
1.5 Target Importation	40
1.6 Library Exportation and Versioning	45
1.7 Subprojects and Amalgamations	50
1.8 Buildfile Language	54
1.8.1 Expansion and Quoting	56
1.8.2 Conditions (if-else)	60
1.8.3 Pattern Matching (switch)	62
1.8.4 Repetitions (for, while)	65
1.9 Implementing Unit Testing	66
1.10 Diagnostics and Debugging	69
2 Project Configuration	74
2.1 config Directive	78
2.2 Configuration Report	83
2.3 Configuration Propagation	85
3 Targets and Target Types	89
3.1 Target Types	89
3.1.1 target{}	91
3.1.2 alias{} and dir{}	92
3.1.3 fsdir{}	92
3.1.4 mtime_target{} and path_target{}	93
3.1.5 group{}	93
3.1.6 file{}	93
3.1.7 doc{}, legal{}, and man{}	93
3.1.8 exe{}	94
3.1.9 json{}	95
4 Variables	95
5 Functions	96
5.1 Builtin Functions	97
5.1.1 \$builtin.defined()	97
5.1.2 \$builtin.visibility()	97
5.1.3 \$builtin.type()	97
5.1.4 \$builtin.null()	97
5.1.5 \$builtin.empty()	97
5.1.6 \$builtin.first(), \$builtin.second()	97

5.1.7	\$builtin.quote()	98
5.1.8	\$builtin.getenv()	98
5.1.9	\$builtin.generate_uuid()	98
5.1.10	\$builtin.sha256sum()	98
5.1.11	\$builtin.xxh64sum()	98
5.2	String Functions	99
5.2.1	\$string.icasecmp()	99
5.2.2	\$string.contains()	99
5.2.3	\$string.starts_with()	99
5.2.4	\$string.ends_with()	99
5.2.5	\$string.compare()	100
5.2.6	\$string.replace()	100
5.2.7	\$string.trim()	101
5.2.8	\$string.lcase(), \$string.ucase()	101
5.2.9	\$string.size()	101
5.2.10	\$string.front()	101
5.2.11	\$string.back()	101
5.2.12	\$string.sort()	101
5.2.13	\$string.find()	102
5.2.14	\$string.find_index()	102
5.2.15	\$string.filter(), \$string.filter_out()	103
5.2.16	\$string.merge()	103
5.2.17	\$string.keys()	104
5.3	Integer Functions	104
5.3.1	\$integer.string()	104
5.3.2	\$integer.integer_sequence()	104
5.3.3	\$integer.size()	104
5.3.4	\$integer.front()	104
5.3.5	\$integer.back()	104
5.3.6	\$integer.sort()	105
5.3.7	\$integer.find()	105
5.3.8	\$integer.find_index()	105
5.4	Bool Functions	105
5.4.1	\$bool.string()	105
5.5	Path Functions	105
5.5.1	\$path.string()	105
5.5.2	\$path.posix_string()	105
5.5.3	\$path.representation()	106
5.5.4	\$path.posix_representation()	106
5.5.5	\$path.absolute()	106
5.5.6	\$path.simple()	106
5.5.7	\$path.sub_path()	106
5.5.8	\$path.super_path()	107
5.5.9	\$path.directory()	107
5.5.10	\$path.root_directory()	107
5.5.11	\$path.leaf()	107

5.5.12	\$path.relative()	107
5.5.13	\$path.base()	108
5.5.14	\$path.extension()	108
5.5.15	\$path.complete()	108
5.5.16	\$path.canonicalize()	108
5.5.17	\$path.normalize(), \$path.try_normalize()	108
5.5.18	\$path.actualize(), \$path.try_actualize()	109
5.5.19	\$path.size()	109
5.5.20	\$path.front()	109
5.5.21	\$path.back()	109
5.5.22	\$path.sort()	109
5.5.23	\$path.find()	110
5.5.24	\$path.find_index()	110
5.5.25	\$path.match()	110
5.6	Name Functions	111
5.6.1	\$name.name()	111
5.6.2	\$name.extension()	111
5.6.3	\$name.directory()	111
5.6.4	\$name.target_type()	111
5.6.5	\$name.project()	111
5.6.6	\$name.is_a()	111
5.6.7	\$name.filter(), \$name.filter_out()	112
5.6.8	\$name.size()	112
5.6.9	\$name.front()	112
5.6.10	\$name.back()	112
5.6.11	\$name.sort()	112
5.6.12	\$name.find()	112
5.6.13	\$name.find_index()	112
5.7	Target Functions	113
5.7.1	\$target.path()	113
5.7.2	\$target.process_path()	113
5.8	Regex Functions	113
5.8.1	\$regex.match()	114
5.8.2	\$regex.find_match()	114
5.8.3	\$regex.filter_match(), \$regex.filter_out_match()	114
5.8.4	\$regex.search()	114
5.8.5	\$regex.find_search()	115
5.8.6	\$regex.filter_search(), \$regex.filter_out_search()	115
5.8.7	\$regex.replace()	115
5.8.8	\$regex.replace_lines()	116
5.8.9	\$regex.split()	116
5.8.10	\$regex.merge()	117
5.8.11	\$regex.apply()	117
5.9	JSON Functions	117
5.9.1	\$json.value_type()	118
5.9.2	\$json.value_size()	118

5.9.3	\$json.member_name()	118
5.9.4	\$json.member_value()	119
5.9.5	\$json.object_names()	119
5.9.6	\$json.array_size()	119
5.9.7	\$json.array_front()	119
5.9.8	\$json.array_back()	119
5.9.9	\$json.array_find()	119
5.9.10	\$json.array_find_index()	119
5.9.11	\$json.load()	120
5.9.12	\$json.parse()	120
5.9.13	\$json.serialize()	120
5.9.14	\$json.size()	121
5.9.15	\$json.keys()	121
5.10	Process Functions	121
5.10.1	\$process.run()	121
5.10.2	\$process.run_regex()	121
5.10.3	\$process.search()	122
5.11	Filesystem Functions	122
5.11.1	\$filesystem.file_exists()	122
5.11.2	\$filesystem.directory_exists()	122
5.11.3	\$filesystem.path_search()	122
5.12	Project Name Functions	122
5.12.1	\$project_name.string()	123
5.12.2	\$project_name.base()	123
5.12.3	\$project_name.extension()	123
5.12.4	\$project_name.variable()	123
5.13	Process Path Functions	123
5.13.1	\$process_path.recall()	123
5.13.2	\$process_path.effect()	124
5.13.3	\$process_path.name()	124
5.13.4	\$process_path.checksum()	124
5.13.5	\$process_path.env_checksum()	124
5.14	Target Triplet Functions	124
5.14.1	\$target_triplet.string()	124
5.14.2	\$target_triplet.representation()	124
6	Directives	124
6.1	define	125
6.2	include	125
6.3	source	125
6.4	update	125
7	Attributes	127
8	Name Patterns	128
9	config Module	131
9.1	Hermetic Build Configurations	131
10	test Module	134
11	install Module	136

11.1 Relocatable Installation	137
11.2 Installation Filtering	138
12 version Module	139
13 bin Module	147
13.1 Binary Target Types	147
13.1.1 lib{}, liba{}, libs{}	147
13.1.2 libul{}, libue{}, libua{}, libus{}	147
13.1.3 obj{}, obje{}, obja{}, objs{}	148
13.1.4 bmi{}, bmie{}, bmia{}, bmis{}	148
13.1.5 hbmi{}, hbmie{}, hbmia{}, hbmis{}	148
13.1.6 def{}	149
14 cc Module	149
14.1 C-Common Configuration Variables	149
14.2 C-Common Target Types	152
14.2.1 pc{}, pca{}, pcs{}	152
14.3 Compilation Internal Scope	152
14.4 Automatic DLL Symbol Exporting	154
14.5 Compiler Predefined Macro Extraction	155
14.6 Importation of Installed Libraries	158
14.6.1 Rewriting Installed Libraries System Root (sysroot)	160
14.7 Compilation Database	161
14.8 GCC Compiler Toolchain	166
14.9 Clang Compiler Toolchain	166
14.9.1 Clang Targeting MSVC	167
14.10 MSVC Compiler Toolchain	168
15 c Module	169
15.1 C Configuration Variables	169
15.2 C Target Types	170
15.2.1 c{}, h{}	171
15.3 Objective-C Compilation	171
15.4 Assembler with C Preprocessor Compilation	171
15.5 C Compiler Predefined Macro Extraction	172
16 cxx Module	173
16.1 C++ Configuration Variables	173
16.2 C++ Target Types	174
16.2.1 cxx{}, hxx{}, ix{ }, txx{ }, mxx{ }	174
16.3 C++ Modules Support	174
16.3.1 Modules Introduction	175
16.3.2 Building Modules	181
16.3.3 Module Symbols Exporting	184
16.3.4 Modules Installation	186
16.3.5 Modules Design Guidelines	187
16.3.6 Modularizing Existing Code	193
16.4 Objective-C++ Compilation	194
16.5 C++ Compiler Predefined Macro Extraction	194
17 in Module	194

18 bash Module	197
19 Appendix A – JSON Dump Format	200

Preface

This document describes the `build2` build system. For the build system driver command line interface refer to the **b(1)** man pages. For other tools in the `build2` toolchain (package and project managers, etc) see the Documentation index.

1 Introduction

The `build2` build system is a native, cross-platform build system with a terse, mostly declarative description language, a conceptual model of build, and a uniform interface with consistent behavior across platforms and compilers.

Those familiar with `make` will see many similarities, though mostly conceptual rather than syntactic. This is not by accident since `build2` borrows the fundamental DAG-based build model from original `make` and many of its conceptual extensions from GNU `make`. We believe, paraphrasing a famous quote, that *those who do not understand make are condemned to reinvent it, poorly*. So our goal with `build2` was to reinvent `make` *well* while handling the demands and complexity of modern cross-platform software development.

Like `make`, `build2` is an "*honest*" build system without magic or black boxes. You can expect to understand what's going on underneath and be able to customize most of its behavior to suit your needs. This is not to say that it's not an *opinionated* build system and if you find yourself "fighting" some of its fundamental design choices, it would probably be wiser to look for alternatives.

We believe the importance and complexity of the problem warranted the design of a new purpose-built language and will hopefully justify the time it takes for you to master it. In the end we hope `build2` will make creating and maintaining build infrastructure for your projects a pleasant task.

Also note that `build2` is not specific to C/C++ or even to compiled languages; its build model is general enough to handle any DAG-based operations. See the `bash` module for a good example.

While the build system is part of a larger, well-integrated build toolchain that includes the package and project dependency managers, it does not depend on them and its standalone usage is the only subject of this manual.

We begin with a tutorial introduction that aims to show the essential elements of the build system on real examples but without getting into too much detail. Specifically, we want to quickly get to the point where we can build useful executable and library projects.

1.1 Hello, World

Let's start with the customary *"Hello, World"* example: a single source file from which we would like to build an executable:

```
$ tree hello/
hello/
|-- hello.cxx

$ cat hello/hello.cxx

#include <iostream>

int main ()
{
    std::cout << "Hello, World!" << std::endl;
}
```

While this very basic program hardly resembles what most software projects look like today, it is useful for introducing key build system concepts without getting overwhelmed. In this spirit we will also use the *build2 simple project* structure, which, similarly, should only be used for basic needs.

To turn our `hello/` directory into a simple project all we need to do is add a `buildfile`:

```
$ tree hello/
hello/
|-- hello.cxx
|-- buildfile

$ cat hello/buildfile

using cxx

exe{hello}: cxx{hello.cxx}
```

Let's start from the bottom: the second line is a *dependency declaration*. On the left hand side of `:` we have a *target*, the `hello` executable, and on the right hand side – a *prerequisite*, the `hello.cxx` source file. Those `exe` and `cxx` in `exe{...}` and `cxx{...}` are called *target types*. In fact, for clarity, target type names are always mentioned with trailing `{ }`, for example, "the `exe{ }` target type denotes an executable".

Notice that the dependency declaration does not specify *how* to build an executable from a C++ source file – this is the job of a *rule*. When the build system needs to update a target, it tries to *match* a suitable rule based on the types of the target and its prerequisites. The *build2* core has a number of predefined fundamental rules with the rest coming from *build system modules*. For example, the `cxx` module defines a number of rules for compiling C++ source code as well as linking executables and libraries.

It should now be easy to guess what the first line of our `buildfile` does: it loads the `cxx` module which defines the rules necessary to build our program (it also registers the `cxx{ }` target type).

Let's now try to build and run our program (b is the build system driver):

```
$ cd hello/ # Change to project root.

$ b
c++ cxx{hello} -> obje{hello}
ld exe{hello}

$ ls -l
buildfile
hello.cxx
hello
hello.d
hello.o
hello.o.d

$ ./hello
Hello, World!
```

Or, if we are on Windows and using Visual Studio:

```
> cd hello

> b
c++ cxx{hello} -> obje{hello}
ld exe{hello}

> dir /b
buildfile
hello.cxx
hello.exe
hello.exe.d
hello.exe.obj
hello.exe.obj.d

> .\hello.exe
Hello, World!
```

By default build2 uses the same C++ compiler it was built with and without passing any extra options, such as debug or optimization, target architecture, etc. To change these defaults we use *configuration variables*. For example, to specify a different C++ compiler we use `config.cxx`:

```
$ b config.cxx=clang++
```

For Visual Studio, build2 by default will use the latest available version and build for the `x86_64` target (`x64` in the Microsoft's terminology). You can, however, override these defaults by either running from a suitable Visual Studio development command prompt or by specifying an absolute path to `cl` that you wish to use. For example (notice the use of inner quotes):

```
> b "config.cxx='...\\VC\\Tools\\MSVC\\14.23.28105\\bin\\Hostx64\\x86\\cl'"
```

See MSVC Compiler Toolchain for details.

Similarly, for additional compile options, such as debug information or optimization level, there is `config.cxx.coptions`. For example:

```
$ b config.cxx=clang++ config.cxx.coptions=-g
```

These and other configuration variables will be discussed in more detail later. We will also learn how to make our configuration persistent so that we don't have to repeat such long command lines on every build system invocation.

Similar to `config.cxx`, there is also `config.c` for specifying the C compiler. Note, however, that if your project uses both C and C++, then you normally only need to specify one of them – `build2` will determine the other automatically.

Let's discuss a few points about the build output. Firstly, to reduce the noise, the commands being executed are by default shown abbreviated and with the same target type notation as we used in the `buildfile`. For example:

```
c++ cxx{hello} -> obje{hello}
ld exe{hello}
```

If, however, you would like to see the actual command lines, you can pass `-v` (to see even more, there is the `-V` as well as `--verbose` options; see **b (1)** for details). For example:

```
$ b -v
g++ -o hello.o -c hello.cxx
g++ -o hello hello.o
```

Most of the files produced by the build system should be self-explanatory: we have the object file (`hello.o`, `hello.obj`) and executable (`hello`, `hello.exe`). For each of them we also have the corresponding `.d` files which store the *auxiliary dependency information*, things like compile options, header dependencies, etc.

To remove the build system output we use the *clean operation* (if no operation is specified, the default is *update*):

```
$ b clean
rm exe{hello}
rm obje{hello}

$ ls -l
buildfile
hello.cxx
```

One of the main reasons behind the *target type* concept is the platform/compiler-specified variances in file names as illustrated by the above listings. In our `buildfile` we refer to the executable target as `exe{hello}`, not as `hello.exe` or `hello$EXT`. The actual file extension, if any, will be determined based on the compiler's target platform by the rule doing the linking. In this sense, target types are a platform-independent replacement of file extensions (though they do have other benefits, such as allowing non-file targets as well as being hierarchical; see Target Types for details).

Let's revisit the dependency declaration line from our `buildfile`:

```
exe{hello}: cxx{hello.cxx}
```

In light of target types replacing file extensions this looks tautological: why do we need to specify both the `cxx{ }` target type *and* the `.cxx` file extension? In fact, we don't have to if we specify the default file extension for the `cxx{ }` target type. Here is our updated `buildfile` in its entirety:

```
using cxx

cxx{*}: extension = cxx

exe{hello}: cxx{hello}
```

Let's unpack the new line. What we have here is a *target type/pattern-specific variable*. It only applies to targets of the `cxx{ }` type whose names match the `*` wildcard pattern. The `extension` variable name is reserved by the `build2` core for specifying target type extensions.

Let's see how all these pieces fit together. When the build system needs to update `exe{hello}`, it searches for a suitable rule. A rule from the `cxx` module matches since it knows how to build a target of type `exe{ }` from a prerequisite of type `cxx{ }`. When the matched rule is *applied*, it searches for a target for the `cxx{hello}` prerequisite. During this search, the `extension` variable is looked up and its value is used to end up with the `hello.cxx` file.

To resolve a rule match ambiguity or to override a default match `build2` uses *rule hints*. For example, if we wanted link a C executable using the C++ link rule:

```
[rule_hint=cxx] exe{hello}: c{hello}
```

Here is our new dependency declaration again:

```
exe{hello}: cxx{hello}
```

It has the canonical form: no extensions, only target types. Sometimes explicit extension specification is still necessary, for example, if your project uses multiple extensions for the same file type. But if unnecessary, it should be omitted for brevity.

If you prefer the `.cpp` file extension and your source file is called `hello.cpp`, then the only line in our `buildfile` that needs changing is the `extension` variable assignment:

```
cxx{*}: extension = cpp
```

Let's say our `hello` program got complicated enough to warrant moving some functionality into a separate source/header module (or a real C++ module). For example:

```
$ tree hello/
hello/
|-- hello.cxx
|-- utility.hxx
|-- utility.cxx
--- buildfile
```

This is what our updated `buildfile` could look like:

```
using cxx

hxx{*}: extension = hxx
cxx{*}: extension = cxx

exe{hello}: cxx{hello} hxx{utility} cxx{utility}
```

Nothing really new here: we've specified the default extension for the `hxx{ }` target type and listed the new header and source files as prerequisites. If you have experience with other build systems, then explicitly listing headers might seem strange to you. As will be discussed later, in `build2` we have to explicitly list all the prerequisites of a target that should end up in a source distribution of our project.

You don't have to list *all* headers that you include, only the ones belonging to your project. Like all modern C/C++ build systems, `build2` performs automatic header dependency extraction.

In real projects with a substantial number of source files, repeating target types and names will quickly become noisy. To tidy things up we can use *name generation*. Here are a few examples of dependency declarations equivalent to the above:

```
exe{hello}: cxx{hello utility} hxx{utility}
exe{hello}: cxx{hello} {hxx cxx}{utility}
```

The last form is probably the best choice if your project contains a large number of header/source pairs. Here is a more realistic example:

```
exe{hello}: {    cxx}{hello}          \
             {hxx  }{forward types}    \
             {hxx cxx}{format print utility}
```

Manually listing a prerequisite every time we add a new source file to our project is both tedious and error prone. Instead, we can automate our dependency declarations with *wildcard name patterns*. For example:

```
exe{hello}: {hxx cxx}{*}
```

Based on the previous discussion of default extensions, you can probably guess how this works: for each target type the value of the `extension` variable is added to the pattern and files matching the result become prerequisites. So, in our case, we will end up with files matching the `*.hxx` and `*.cxx` wildcard patterns.

In more complex projects it is often convenient to organize source code into subdirectories. To handle such projects we can use the recursive wildcard:

```
exe{hello}: {hxx cxx}{{**}}
```

Using wildcards is somewhat controversial. Patterns definitely make development more pleasant and less error prone: you don't need to update your `buildfile` every time you add, remove, or rename a source file and you won't forget to explicitly list headers, a mistake that is often only detected when trying to build a source distribution of a project. On the other hand, there is the possibility of including stray source files into your build without noticing. And, for more complex projects, name patterns can become fairly complex (see *Name Patterns* for details). Note also that on modern hardware the performance of wildcard searches hardly warrants a consideration.

In our experience, when combined with modern version control systems like `git(1)`, stray source files are rarely an issue and generally the benefits of wildcards outweigh their drawbacks. But, in the end, whether to use them or not is a personal choice and, as shown above, `build2` supports both approaches.

And that's about all there is to our `hello` example. To summarize, we've seen that to build a simple project we need a single `buildfile` which itself doesn't contain much more than a dependency declaration for what we want to build. But we've also mentioned that simple projects are only really meant for basics. So let's convert our `hello` example to the *standard project* structure which is what we will be using for most of our real development.

Simple projects have so many restrictions and limitations that they are hardly usable for anything but, well, *really* simple projects.

Specifically, such projects cannot be imported by other projects nor can they use build system modules that require bootstrapping. Notably, this includes the `dist` and `config` modules (the `test` and `install` modules are loaded implicitly). And without the `config` module there is no support for persistent configurations.

As a result, you should only use a simple project if you are happy to always build in the source directory and with the default build configuration or willing to specify the output directory and/or custom configuration on every invocation. In other words, expect an experience similar to a plain `Makefile`.

One notable example where simple projects are handy is a *glue buildfile* that "pulls" together several other projects, usually for convenience of development. See *Target Importation* for details.

1.2 Project Structure

A `build2 standard project` has the following overall layout:

```
hello/
|-- build/
|   |-- bootstrap.build
|   |-- root.build
|   ...
|-- ...
--- buildfile
```

Specifically, the project's root directory should contain the `build/` subdirectory as well as the root `buildfile`. The `build/` subdirectory contains project-wide build system information.

The **bdep-new(1)** command is an easy way to create the standard layout executable (`-t exe`) and library (`-t lib`) projects. To change the C++ file extensions to `.hpp/ .cpp`, pass `-l c++,cpp`. For example:

```
$ bdep new --no-init -l c++,cpp -t exe hello
```

It is also possible to use an alternative build file/directory naming scheme where every instance of the word *build* is replaced with *build2*, for example:

```
hello/
|-- build2/
|   |-- bootstrap.build2
|   |-- root.build2
|   ...
|-- ...
--- build2file
```

Note that the naming must be consistent within a project with all the filesystem entries either following *build* or *build2* scheme. In other words, we cannot call the directory `build2/` while still using `buildfile`.

The alternative naming scheme is primarily useful when adding `build2` support to an existing project along with other build systems. In this case, the fairly generic standard names might already be in use. For example, it is customary to have `build/` in `.gitignore`. Plus more specific naming will make it easier to identify files and directories as belonging to the `build2` support. For new projects as well as for existing projects that are switching exclusively to `build2` the standard naming scheme is recommended.

To create a project with the alternative naming using **bdep-new(1)** pass the `alt-naming` project type sub-option. For example:

```
$ bdep new -t exe,alt-naming ...
```

To support lazy loading of subprojects (discussed later), reading of the project's build information is split into two phases: bootstrapping and loading. During bootstrapping the project's `build/bootstrap.build` file is read. Then, when (and if) the project is loaded completely, its `build/root.build` file is read followed by the `buildfile` (normally from the project root but possibly from a subdirectory).

The `bootstrap.build` file is required. Let's see what it would look like for a typical project using our `hello` as an example:

```
project = hello

using version
using config
using test
using install
using dist
```

The first non-comment line in `bootstrap.build` should be the assignment of the project name to the `project` variable. After that, a typical `bootstrap.build` file loads a number of build system modules. While most modules can be loaded during the project load phase in `root.build`, certain modules have to be loaded early, while bootstrapping (for example, because they define new operations).

Let's examine briefly the modules loaded by our `bootstrap.build`: The `version` module helps with managing our project versioning. With this module we only maintain the version in a single place (the project's manifest file) and it is automatically made available in various convenient forms throughout our project (buildfiles, header files, etc). The `version` module also automates versioning of snapshots between releases.

The manifest file is what makes our build system project a *package*. It contains all the metadata that a user of a package might need to know: name, version, dependencies, etc., all in one place. However, even if you don't plan to package your project, it is a good idea to create a basic manifest if only to take advantage of the version management offered by the `version` module. So let's go ahead and add it next to our root buildfile:

```
$ tree hello/
hello/
|-- build/
|   |-- ...
|-- ...
|-- buildfile
--- manifest

$ cat hello/manifest
: 1
name: hello
version: 0.1.0
summary: hello C++ executable
```

The `config` module provides support for persistent configurations. While build configuration is a large topic that we will be discussing in more detail later, in a nutshell `build2` support for configuration is an integral part of the build system with the same mechanisms available to the build system core, modules, and your projects. However, without `config`, the configuration information is *transient*. That is, whatever configuration information was automatically discovered or that you have supplied on the command line is discarded after each build system invocation. With the `config` module, however, we can *configure* a project to make the configuration *persistent*. We will see an example of this shortly.

Next up are the `test`, `install`, and `dist` modules. As their names suggest, they provide support for testing, installation and preparation of source distributions. Specifically, the `test` module defines the `test` operation, the `install` module defines the `install` and `uninstall` operations, and the `dist` module defines the `dist` (meta-)operation. Again, we will try them out in a moment.

Moving on, the `root.build` file is optional though most projects will have it. This is the place where we define project's configuration variables (subject of Project Configuration), establish project-wide settings, as well as load build system modules that provide support for the languages/tools that we use. Here is what it could look like for our `hello` example:

```
cxx.std = latest

using cxx

hxx{*}: extension = hxx
cxx{*}: extension = cxx
```

As you can see, we've moved the loading of the `cxx` modules and setting of the default file extensions from the `root.buildfile` in our simple project to `root.build` when using the standard layout. We've also set the `cxx.std` variable to tell the `cxx` module to select the latest C++ standard available in any particular C++ compiler this project might be built with.

Selecting the C++ standard for our project is a messy issue. If we don't specify the standard explicitly with `cxx.std`, then the default standard in each compiler will be used, which, currently, can range from C++98 to C++14. So unless you carefully write your code to work with any standard, this is probably not a good idea.

Fixing the standard (for example, to `c++11`, `c++14`, etc) should work theoretically. In practice, however, compilers add support for new standards incrementally and many versions, while perfectly usable, are not feature-complete. As a result, a better practical strategy is to specify the set of minimum supported compiler versions rather than the C++ standard.

There is also the issue of using libraries that require a newer standard in old code. For example, headers from a library that relies on C++14 features will not compile when included in a project that is built as C++11. And, even if the headers compile (that is, C++14 features are only used in the implementation), strictly speaking, there is no guarantee that codebases compiled with different C++ standards are ABI compatible (in fact, some changes to the C++ language leave the implementations no choice but to break the ABI).

As result, our recommendation is to set the standard to `latest` and specify the minimum supported compilers and versions in your project's documentation (see `package manifest requires` value for one possible place). Practically, this should allow you to include and link any library, regardless of the C++ standard that it uses.

Let's now take a look at the `root.buildfile`:

```
./: {*/ -build/}
```

In plain English, this `buildfile` declares that building this directory (and, since it's the root of our project, building this entire project) means building all its subdirectories excluding `build/`. Let's now try to understand how this is actually achieved.

We already know this is a dependency declaration, `./` is the target, and what's after `:` are its prerequisites, which seem to be generated with some kind of a name pattern (the wildcard character in `*/` should be the giveaway). What's unusual about this declaration, however, is the lack of any target types plus that strange-looking `./`.

Let's start with the missing target types. In fact, the above `buildfile` can be rewritten as:

```
dir{.}: dir{* -build}
```

So the trailing slash (always forward, even on Windows) is a special shorthand notation for `dir{}`. As we will see shortly, it fits naturally with other uses of directories in `buildfiles` (for example, in scopes).

The `dir{}` target type is an *alias* (and, in fact, is derived from more general `alias{}`; see Target Types for details). Building it means building all its prerequisites.

If you are familiar with `make`, then you can probably see the similarity with the ubiquitous `all` pseudo-target. In `build2` we instead use directory names as more natural aliases for the "build everything in this directory" semantics.

Note also that `dir{}` is purely an alias and doesn't have anything to do with the filesystem. In particular, it does not create any directories. If you do want explicit directory creation (which should be rarely needed), use the `fsdir{}` target type instead.

The `./` target is a special *default target*. If we run the build system without specifying the target explicitly, then this target is built by default. Every `buildfile` has the `./` target. If we don't declare it explicitly, then its declaration is implied with the first target in the `buildfile` as its prerequisite. Recall our `buildfile` from the simple `hello` project:

```
exe{hello}: cxx{hello}
```

It is equivalent to:

```
./: exe{hello}
exe{hello}: cxx{hello}
```

If, however, we had several targets in the same directory that we wanted built by default, then we would need to explicitly list them as prerequisites of the default target. For example:

```
./: exe{hello}
exe{hello}: cxx{hello}

./: exe{goodby}
exe{goodby}: cxx{goodby}
```

While straightforward, this is somewhat inelegant in its repetitiveness. To tidy things up we can use *dependency declaration chains* that allow us to chain together several target-prerequisite declarations in a single line. For example:

```
./: exe{hello}: cxx{hello}

./: exe{goodby}: cxx{goodby}
```

With dependency chains a prerequisite of the preceding target becomes a target itself for the following prerequisites.

Let's get back to our root buildfile:

```
./: {*/ -build/}
```

The last unexplained bit is the {*/ -build/} name pattern. All it does is exclude build/ from the subdirectories to build. See Name Patterns for details.

Let's take a look at a slightly more realistic root buildfile:

```
./: {*/ -build/} doc{README.md LICENSE} manifest
```

Here we have the customary README.md and LICENSE files as well as the package manifest. Listing them as prerequisites achieves two things: they will be installed if/when our project is installed and, as mentioned earlier, they will be included into the project source distribution.

The README.md and LICENSE files use the doc{} target type. We could have used the generic file{} but using the more precise doc{} makes sure that they are installed into the appropriate documentation directory. The manifest file doesn't need an explicit target type since it has a fixed name (manifest{manifest} is valid but redundant).

Standard project infrastructure in place, where should we put our source code? While we could have everything in the root directory of our project, just like we did with the simple layout, it is recommended to instead place the source code into a subdirectory named the same as the project. For example:

```
hello/
|-- build/
|   |-- ...
|-- hello/
|   |-- hello.cxx
|   |-- buildfile
|-- buildfile
|-- manifest
-- README.md
```

There are several reasons for this layout: It implements the canonical inclusion scheme where each header is prefixed with its project name. It also has a predictable name where users can expect to find our project's source code. Finally, this layout prevents clutter in the project's root directory which usually contains various other files. See Canonical Project Structure for details.

Note, however, that this layout is not mandatory and `build2` is flexible enough to support various arrangements used in today's C and C++ projects. Furthermore, the **`bdep-new(1)`** command provides a number of customization options and chances are you will be able to create your preferred layout automatically. See `SOURCE LAYOUT` for more information and examples.

Note also that while we can name our header and source files however we like (but, again, see Canonical Project Structure for some sensible guidelines), C++ module interface files need to embed a sufficient amount of the module name suffix in their names to unambiguously resolve all the modules within a project. See `Building Modules` for details.

The source subdirectory `buildfile` is identical to that of the simple project minus the parts moved to `root.build`:

```
exe{hello}: {hxx cxx}{**}
```

Let's now build our project and see where the build system output ends up in this new layout:

```
$ cd hello/ # Change to project root.
$ b
c++ hello/cxx{hello} -> hello/obje{hello}
ld hello/exe{hello}

$ tree ./
./
|-- build/
|   |-- ...
|-- hello/
|   |-- hello.cxx
|   |-- hello
|   |-- hello.d
|   |-- hello.o
|   |-- hello.o.d
|   |-- buildfile
|-- buildfile
-- manifest

$ hello/hello
Hello, World!
```

If we don't specify a target to build (as in the example above), then `build2` will build the current directory or, more precisely, the default target in the `buildfile` in the current directory. We can also build a directory other than the current, for example:

```
$ b hello/
```

Note that the trailing slash is required. In fact, `hello/` in the above command line is a target and is equivalent to `dir{hello}`, just like in the `buildfiles`.

Or we can build a specific target:

```
$ b hello/exe{hello}
```

Naturally, nothing prevents us from building multiple targets or even projects in the same build system invocation. For example, if we had the `libhello` project next to our `hello/`, then we could build both at once:

```
$ ls -l
hello/
libhello/

$ b hello/ libhello/
```

Speaking of libraries, let's see what the standard project structure looks like for one, using `libhello` created by **bdep-new(1)** as an example:

```
$ bdep new --no-init -l c++ -t lib libhello

$ tree libhello/
libhello/
|-- build/
|   |-- bootstrap.build
|   |-- root.build
|   |-- export.build
|-- libhello/
|   |-- hello.hxx
|   |-- hello.cxx
|   |-- export.hxx
|   |-- version.hxx.in
|   |-- buildfile
|-- tests/
|   |-- ...
|-- buildfile
|-- manifest
--- README.md
```

The overall layout (`build/`, `libhello/` source subdirectory) as well as the contents of the root files (`bootstrap.build`, `root.build`, `root buildfile`) are exactly the same. There is, however, the new file `export.build` in `build/`, the new subdirectory `tests/`, and the contents of the project's source subdirectory `libhello/` look quite a bit different. We will examine all of these differences in the coming sections, as we learn more about the build system.

Again, this layout is not mandatory and **bdep-new(1)** can create a number of alternative library structures. For example, if you prefer the `include/src` split, try:

```
$ bdep new --no-init -l c++ -t lib,split libhello
```

See **SOURCE LAYOUT** for more examples.

The standard project structure is not type (executable, library, etc) or even language specific. In fact, the same project can contain multiple executables and/or libraries (for example, both `hello` and `libhello`). However, if you plan to package your projects, it is a good idea to keep them as separate build system projects (they can still reside in the same version control repository, though).

Speaking of projects, this term is unfortunately overloaded to mean two different things at different levels of software organization. At the bottom we have *build system projects* which, if packaged, become *packages*. And at the top, related packages are often grouped into what is also commonly referred to as *projects*. At this point both usages are probably too well established to look for alternatives.

And this completes the conversion of our simple `hello` project to the standard structure. Earlier, when examining `bootstrap.build`, we mentioned that modules loaded in this file usually provide additional operations. So we still need to discuss what exactly the term *build system operation* means and see how to use operations that are provided by the modules we have loaded. But before we do that, let's see how we can build our projects *out of source* tree and learn about another cornerstone `build2` concept: *scopes*.

1.3 Output Directories and Scopes

Two common requirements placed on modern build systems are the ability to build projects out of the source directory tree (referred to as just *out of source* vs *in source*) as well as isolation of `buildfiles` from each other when it comes to target and variable names. In `build2` these mechanisms are closely-related, integral parts of the build system.

This tight integration has advantages, like being always available and working well with other build system mechanisms, as well as disadvantages, like the inability to implement a completely different out of source arrangement and/or isolation model. In the end, if you find yourself "fighting" this aspect of `build2`, it will likely be easier to use a different build system than subvert it.

Let's start with an example of an out of source build for our `hello` project. To recap, this is what we have:

```
$ ls -l
hello/

$ tree hello/
hello/
|-- build/
|   |-- ...
|-- hello/
|   |-- ...
|-- buildfile
--- manifest
```

To start, let's build it in the `hello-out/` directory next to the project:

```
$ b hello/@hello-out/
mkdir fsdir{hello-out/}
mkdir hello-out/fsdir{hello/}
c++ hello/hello/cxx{hello} -> hello-out/hello/obje{hello}
ld hello-out/hello/exe{hello}

$ ls -l
hello/
hello-out/
```

```
$ tree hello-out/
hello-out/
|-- hello/
|   |-- hello
|   |-- hello.d
|   |-- hello.o
|   |-- hello.o.d
|   --
--
```

This definitely requires some explaining. Let's start from the bottom, with the `hello-out/` layout. It is *parallel* to the source directory. This mirrored side-by-side listing (of the relevant parts) should illustrate this clearly:

```
hello/          ~~> hello-out/
|-- hello/      ~~> |-- hello/
    |-- hello.cxx  ~~>    |-- hello.o
```

In fact, if we copy the contents of `hello-out/` over to `hello/`, we will end up with exactly the same result as in the in source build. And this is not accidental: an in source build is just a special case of an out of source build where the *out* directory is the same as *src*.

In `build2` this parallel structure of the out and src directories is a cornerstone design decision and is non-negotiable, so to speak. In particular, out cannot be inside src. And while we can stash the build system output (object files, executables, etc) into (potentially different) subdirectories, this is not recommended. As will be shown later, `build2` offers better mechanisms to achieve the same benefits (like reduced clutter, ability to run executables) but without the drawbacks (like name clashes).

Let's now examine how we invoked the build system to achieve this out of source build. Specifically, if we were building in source, our command line would have been:

```
$ b hello/
```

but for the out of source build, we have:

```
$ b hello/@hello-out/
```

In fact, that strange-looking construct, `hello/@hello-out/` is just a more elaborate target specification that explicitly spells out the target's src and out directories. Let's add an explicit target type to make it clearer:

```
$ b hello/@hello-out/dir{.}
```

What we have on the right of `@` is the target in the out directory and on the left – its src directory. In plain English, this command line says "build me the default target from `hello/` in the `hello-out/` directory".

As an example, if instead we wanted to build only the `hello` executable out of source, then the invocation would have looked like this:

```
$ b hello/hello/@hello-out/hello/exe{hello}
```

We could have also specified out for an in source build, but that's redundant:

```
$ b hello/@hello/
```

There is another example of this elaborate target specification that can be seen in the build diagnostics, for instance, when installing headers of a library (the `install` operation is discussed in the next section):

```
$ b install: libhello/@libhello-out/
...
install libhello/libhello/hxx{hello}@libhello-out/libhello/ ->
    /usr/local/include/
```

Notice, however, that now the target (`hxx{hello}`) is on the left of `@`, that is, in the `src` directory. It does, however, make sense if you think about it: our `hello.hxx` is a *source file*, in a sense that it is not built and it resides in the project's source directory. This is in contrast, for example, to the `exe{hello}` target which is the output of the build system and goes to the out directory. So in `build2` targets can be either in `src` or in `out` (there can also be *out of any project* targets, for example, installed files).

The elaborate target specification can also be used in `buildfiles`. We haven't encountered any so far because targets mentioned without explicit `src/out` default to `out` and, naturally, most of the targets we mention in `buildfiles` are things we want built. One situation where you may encounter an `src` target mentioned explicitly is when specifying its installability (discussed in the next section). For example, if our project includes the customary `INSTALL` file, it probably doesn't make sense to install it. However, since it is a source file, we have to use the elaborate target specification when disabling its installation:

```
doc{INSTALL}@./: install = false
```

Note also that only targets and not prerequisites have this notion of `src/out` directories. In a sense, prerequisites are relative to the target they are prerequisites of and are resolved to targets in a manner that is specific to their target types. For `file{}`-based prerequisites the corresponding target in `out` is first looked up and, if found, used. Otherwise, an existing file in `src` is searched for and, if found, the corresponding target (now in `src`) is used. In particular, this semantics gives preference to generated code over static.

More precisely, a prerequisite is relative to the scope (discussed below) in which the dependency is declared and not to the target that it is a prerequisite of. However, in most practical cases, this means the same thing.

And this pretty much covers out of source builds. Let's summarize the key points we have established so far: Every build has two parallel directory trees, `src` and `out`, with the in source build being just a special case where they are the same. Targets in a project can be either in the `src` or `out` directory though most of the time targets we mention in our `buildfiles` will be in `out`, which is the default. Prerequisites are relative to targets they are prerequisites of and `file{}`-based prerequisites are first searched for as declared targets in `out` and then as existing files in `src`.

Note also that we can have as many out of source builds as we want and we can place them anywhere we want (but not inside `src`), say, on a RAM-backed disk/filesystem. As an example, let's build our `hello` project with two different compilers:

```
$ b hello/@hello-gcc/      config.cxx=g++
$ b hello/@hello-clang/    config.cxx=clang++
```

In the next section we will see how to permanently configure our out of source builds so that we don't have to keep repeating these long command lines.

While technically you can have both in source and out of source builds at the same time, this is not recommended. While it may work for basic projects, as soon as you start using generated source code (which is fairly common in `build2`), it becomes difficult to predict where the compiler will pick generated headers. There is support for remapping mis-picked headers but this may not always work with older C/C++ compilers. Plus, as we will see in the next section, `build2` supports *forwarded configurations* which provide most of the benefits of an in source build but without the drawbacks.

Let's now turn to `buildfile` isolation. It is a common, well-established practice to organize complex software projects in directory hierarchies. One of the benefits of this organization is isolation: we can use the same, short file names in different subdirectories. In `build2` the project's directory tree is used as a basis for its *scope* hierarchy. In a sense, scopes are like C++ namespaces that automatically track the project's filesystem structure and use directories as their names. The following listing illustrates the parallel directory and scope hierarchies for our `hello` project. The `build/` subdirectory is special and does not have a corresponding scope.

[illegible]

Every `buildfile` is loaded in its corresponding scope, variables set in a `buildfile` are set in this scope and relative targets mentioned in a `buildfile` are relative to this scope's directory. Let's "load" the `buildfile` contents from our `hello` project to the above listing:

```
hello/
|
|-- buildfile
|
.-- hello/
    |
    .-- buildfile

hello/
{
    ./: {*/ -build/}

    hello/
    {
        exe{hello}: {hxx cxx}{**}
    }
}
```

In fact, to be absolutely precise, we should also add the contents of `bootstrap.build` and `root.build` to the project's root scope (module loading is omitted for brevity):

```
hello/                                     hello/
|                                         {
|-- build/                               {
|   |-- bootstrap.build      project = hello
|   |
|   |-- root.build           cxx.std = latest
|   |                       hxx{*}: extension = hxx
|   |                       cxx{*}: extension = cxx
|   |
|-- buildfile                 ./: {*/ -build/}
|-- hello/                   hello/
|   |                       {
|   |-- buildfile            exe{hello}: {hxx cxx}{**}
|   |                       }
|   |                       }
|   |                       }
```

The above scope structure is very similar to what you will see (besides a lot of other things) if you build with `--dump match`. With this option the build system driver dumps the build state after matching rules to targets (see [Diagnostics and Debugging](#) for more information). Here is an abbreviated output of building our `hello` with `--dump` (assuming an in source build in `/tmp/hello`):

```
$ b --dump match

/
{
  [target_triplet] build.host = x86_64-linux-gnu
  [string] build.host.class = linux
  [string] build.host.cpu = x86_64
  [string] build.host.system = linux-gnu

  /tmp/hello/
  {

    [dir_path] src_root = /tmp/hello/
    [dir_path] out_root = /tmp/hello/

    [dir_path] src_base = /tmp/hello/
    [dir_path] out_base = /tmp/hello/

    [project_name] project = hello
    [string] project.summary = hello executable
    [string] project.url = https://example.org/hello

    [string] version = 1.2.3
    [uint64] version.major = 1
    [uint64] version.minor = 2
    [uint64] version.patch = 3

    [string] cxx.std = latest

    [string] cxx.id = gcc
    [string] cxx.version = 8.1.0
    [uint64] cxx.version.major = 8
    [uint64] cxx.version.minor = 1
```

```

[uint64] cxx.version.patch = 0

[target_triplet] cxx.target = x86_64-w64-mingw32
[string] cxx.target.class = windows
[string] cxx.target.cpu = x86_64
[string] cxx.target.system = mingw32

hxx{*}: [string] extension = hxx
cxx{*}: [string] extension = cxx

hello/
{
  [dir_path] src_base = /tmp/hello/hello/
  [dir_path] out_base = /tmp/hello/hello/

  dir{./}: exe{hello}
  exe{hello.}: cxx{hello.cxx}
}

dir{./}: dir{hello/} manifest{manifest}
}

```

This is probably quite a bit more information than what you've expected to see so let's explain a couple of things. Firstly, it appears there is another scope outer to our project's root. In fact, `build2` extends scoping outside of projects with the root of the filesystem (denoted by the special `/`) being the *global scope*. This extension becomes useful when we try to build multiple unrelated projects or import one project into another. In this model all projects are part of a single scope hierarchy with the global scope at its root.

The global scope is read-only and contains a number of pre-defined *build-wide* variables such as the build system version, host platform (shown in the above listing), etc.

Next, inside the global scope, we see our project's root scope (`/tmp/hello/`). Besides the variables that we have set ourselves (like `project`), it also contains a number of variables set by the build system core (for example, `out_base`, `src_root`, etc) as well by build system modules (for example, `project.*` and `version.*` variables set by the `version` module and `cxx.*` variables set by the `cxx` module).

The scope for our project's source directory (`hello/`) should look familiar. We again have a few special variables (`out_base`, `src_base`). Notice also that the name patterns in prerequisites have been expanded to the actual files.

As you can probably guess from their names, the `src_*` and `out_*` variables track the association between scopes and `src/out` directories. They are maintained automatically by the build system core with the `src/out_base` pair set on each scope within the project and an additional `src/out_root` pair set on the project's root scope so that we can get the project's root directories from anywhere in the project. Note that directory paths in these variables are always absolute and normalized.

In the above example the corresponding `src/out` variable pairs have the same values because we were building in source. As an example, this is what the association will look like for an out of source build:

```

hello/  ~~>      hello-out/                                <~~  hello-out/
|
|      {
|      src_root = ../hello/
|      out_root = ../hello-out/
|
|      src_base = ../hello/
|      out_base = ../hello-out/
|
|-- hello/  ~~>      hello/                                <~~  |-- hello/
|      {
|      src_base = ../hello/hello/
|      out_base = ../hello-out/hello/
|      }
|      }

```

Now that we have some scopes and variables to play with, it's a good time to introduce variable expansion. To get the value stored in a variable we use `$` followed by the variable's name. The variable is first looked up in the current scope (that is, the scope in which the expansion was encountered) and, if not found, in the outer scopes all the way to the global scope.

To be precise, this is for the default *variable visibility*. Variables, however, can have more limited visibilities, such as *project*, *scope*, *target*, or *prerequisite*.

To illustrate the lookup semantics, let's add the following line to each buildfile in our hello project:

```

$ cd hello/ # Change to project root.

$ cat buildfile
...
info "src_base: $src_base"

$ cat hello/buildfile
...
info "src_base: $src_base"

```

And then build it:

```

$ b
buildfile:3:1: info: src_base: /tmp/hello/
hello/buildfile:8:1: info: src_base: /tmp/hello/hello/

```

In this case `src_base` is defined in each of the two scopes and we get their respective values. If, however, we change the above line to print `src_root` instead of `src_base`, we will get the same value from the root scope:

```

buildfile:3:1: info: src_root: /tmp/hello/
hello/buildfile:8:1: info: src_root: /tmp/hello/

```

In this section we've only scratched the surface when it comes to variables. In particular, variables and variable values in `build2` are optionally typed (those `[string]`, `[uint64]` we've seen in the build state dump). And in certain contexts the lookup semantics actually starts from the target, not from the scope (target-specific variables; there are also prerequisite-specific). These and other variable-related topics will be covered in subsequent sections.

One typical place to find `src/out_root` expansions is in the include search path options. For example, the source subdirectory `buildfile` generated by **`bdep-new(1)`** for an executable project actually looks like this (`poptions` stands for *preprocessor options*):

```
exe{hello}: {hxx cxx}{**}

cxx.poptions += "-I$out_root" "-I$src_root"
```

The strange-looking `+=` line is a *prepend* variable assignment. It adds the value on the right hand side to the beginning of the existing value. So, in the above example, the two header search paths will be added before any of the existing preprocessor options (and thus will be considered first).

There are also the *append* assignment, `+=`, which adds the value on the right hand side to the end of the existing value, as well as, of course, the normal or *replace* assignment, `=`, which replaces the existing value with the right hand side. One way to remember where the existing and new values end up in the `+=` and `+=` results is to imagine the new value taking the position of `=` and the existing value – of `+`.

The above `buildfile` allows us to include our headers using the project's name as a prefix, inline with the Canonical Project Structure guidelines. For example, if we added the `utility.hxx` header to our `hello` project, we would include it like this:

```
#include <iostream>

#include <hello/utility.hxx>

int main ()
{
...
}
```

Besides `poptions`, there are also `coptions` (compile options), `loptions` (link options), `aoptions` (archive options) and `libs` (extra libraries to link). If you are familiar with `make`, these are roughly equivalent to `CPPFLAGS`, `CFLAGS/CXXFLAGS`, `LDFLAGS`, `ARFLAGS`, and `LIBS/LDLIBS`, respectively. Here they are again in the tabular form:

<code>*.poptions</code>	preprocess	<code>CPPFLAGS</code>
<code>*.coptions</code>	compile	<code>CFLAGS/CXXFLAGS</code>
<code>*.loptions</code>	link	<code>LDFLAGS</code>
<code>*.aoptions</code>	archive	<code>ARFLAGS</code>
<code>*.libs</code>	extra libraries	<code>LIBS/LDLIBS</code>

More specifically, there are three sets of these variables: `cc.*` (stands for *C-common*) which applies to all C-like languages as well as `c.*` and `cxx.*` which only apply during the C and C++ compilation, respectively. We can use these variables in our buildfiles to adjust the compiler/linker behavior. For example:

```
if ($cc.class == 'gcc')
{
    cc.options += -fno-strict-aliasing # C and C++
    cxx.options += -fno-exceptions    # only C++
}

if ($c.target.class != 'windows')
    c.libs += -ldl # only C
```

Additionally, as we will see in [Configuring](#), there are also the `config.cc.*`, `config.c.*`, and `config.cxx.*` sets which are used by the users of our projects to provide external configuration. The initial values of the `cc.*`, `c.*`, and `cxx.*` variables are taken from the corresponding `config.*.*` values.

And, as we will learn in [Library Exportation](#), there are also the `cc.export.*`, `c.export.*`, and `cxx.export.*` sets that are used to specify options that should be exported to the users of our library.

If we adjust the `cc.*`, `c.*`, and `cxx.*` variables at the scope level, as in the above fragment, then the changes will apply when building every target in this scope (as well as in the nested scopes, if any). Usually this is what we want but sometimes we may need to pass additional options only when compiling certain source files or linking certain libraries or executables. For that we use the target-specific variable assignment. For example:

```
exe{hello}: {hxx cxx}{**}

obj{utility}: cxx.poptions += -DNDEBUG
exe{hello}: cxx.loptions += -static
```

Note that we set these variables on targets which they affect. In particular, those with a background in other build systems may, for example, erroneously expect that setting `poptions` on a library target will affect compilation of its prerequisites. For example, the following does not work:

```
exe{hello}: cxx.poptions += -DNDEBUG
```

The recommended way to achieve this behavior in `build2` is to organize your targets into subdirectories, in which case we can just set the variables on the scope. And if this is impossible or undesirable, then we can use target type/pattern-specific variables (if there is a common pattern) or simply list the affected targets explicitly. For example:

```
obj{*.test}: cxx.poptions += -DDEFINE_MAIN
obj{main utility}: cxx.poptions += -DNDEBUG
```

The first line covers compilation of source files that have the `.test` second-level extension (see Implementing Unit Testing for background) while the second simply lists the targets explicitly.

It is also possible to specify different options when producing different types of object files (`obje{}` – executable, `obja{}` – static library, or `objs{}` – shared library) or when linking different libraries (`liba{}` – static library or `libs{}` – shared library). See Library Exportation and Versioning for an example.

As mentioned above, each `buildfile` in a project is loaded into its corresponding scope. As a result, we rarely need to open scopes explicitly. In the few cases that we do, we use the following syntax:

```
<directory>/
{
    ...
}
```

If the scope directory is relative, then it is assumed to be relative to the current scope. As an exercise for understanding, let's reimplement our `hello` project as a single `buildfile`. That is, we move the contents of the source subdirectory `buildfile` into the root `buildfile`:

```
$ tree hello/
hello/
|-- build/
|   |-- ...
|-- hello/
|   |-- hello.cxx
|-- buildfile

$ cat hello/buildfile

./: hello/

hello/
{
    ./: exe{hello}: {hxx cxx}{**}
}
```

While this single `buildfile` setup is not recommended for new projects, it can be useful for non-intrusive conversion of existing projects to `build2`. One approach is to place the unmodified original project into a subdirectory (potentially automating this with a mechanism such as `git(1)` submodules) then adding the `build/` subdirectory and the root `buildfile` which explicitly opens scopes to define the build over the upstream project's subdirectory structure.

Seeing this merged `buildfile` may make you wonder what exactly caused the loading of the source subdirectory `buildfile` in our normal setup. In other words, when we build our `hello` from the project root, who loads `hello/buildfile` and why?

Actually, in the earlier days of `build2`, we had to explicitly load `buildfiles` that define targets we depend on with the `include` directive. In fact, we still can (and have to if we are depending on targets other than directories). For example:

```
./: hello/

include hello/buildfile
```

We can also omit `buildfile` for brevity and have just:

```
include hello/
```

This explicit inclusion, however, quickly becomes tiresome as the number of directories grows. It also makes using wildcard patterns for subdirectory prerequisites a lot less appealing.

To overcome this the `dir{ }` target type implements an interesting prerequisite to target resolution semantics: if there is no existing target with this name, a `buildfile` that (presumably) defines this target is automatically loaded from the corresponding directory. In fact, this mechanism goes a step further and, if the `buildfile` does not exist, then it assumes one with the following contents was implied:

```
./: */
```

That is, it simply builds all the subdirectories. This is especially handy when organizing related tests into directory hierarchies.

As mentioned above, this automatic inclusion is only triggered if the target we depend on is `dir{ }` and we still have to explicitly include the necessary `buildfiles` for other targets. One common example is a project consisting of a library and an executable that links it, each residing in a separate directory next to each other (as noted earlier, this is not recommended for projects that you plan to package). For example:

```
hello/
|-- build/
|   |-- ...
|-- hello/
|   |-- main.cxx
|   |-- buildfile
|-- libhello/
|   |-- hello.hxx
|   |-- hello.cxx
|   |-- buildfile
.-- buildfile
```

In this case the executable `buildfile` would look along these lines:

```
include ../libhello/ # Include lib{hello}.

exe{hello}: {hxx cxx}{**} ../libhello/lib{hello}
```

Note also that `buildfile` inclusion should only be used for accessing targets within the same project. For cross-project references we use Target Importation.

1.4 Operations

Modern build systems have to perform operations other than just building: cleaning the build output, running tests, installing/uninstalling the build results, preparing source distributions, and so on. And, if the build system has integrated configuration support, configuring the project would naturally belong to this list as well.

If you are familiar with `make`, you should recognize the parallel with the common `clean`, `test`, `install`, and `dist`, "operation" pseudo-targets.

In `build2` we have the concept of a *build system operation* performed on a target. The two pre-defined operations are `update` and `clean` with other operations provided by build system modules.

Operations to be performed and targets to perform them on are specified on the command line. As discussed earlier, `update` is the default operation and `./` in the current directory is the default target if no operation and/or target is specified explicitly. And, similar to targets, we can specify multiple operations (not necessarily on the same target) in a single build system invocation. The list of operations to perform and targets to perform them on is called a *build specification* or *buildspec* for short (see **b (1)** for details). Here are a few examples:

```
$ cd hello          # Change to project root.

$ b                # Update current directory.
$ b ./             # Same as above.
$ b update         # Same as above.
$ b update: ./     # Same as above.

$ b clean update   # Rebuild.

$ b clean: hello/      # Clean specific target.
$ b update: hello/exe{hello} # Update specific target

$ b update: libhello/ tests/ # Update two targets.
```

If you are running `build2` from PowerShell, then you will need to use quoting when updating specific targets, for example:

```
$ b update: 'hello/exe{hello}'
```

Let's revisit `build/bootstrap.build` from our `hello` project:

```
project = hello

using version
using config
using test
using install
using dist
```

Other than `version`, all the modules we load define new operations. Let's examine each of them starting with `config`.

1.4.1 Configuring

As mentioned briefly earlier, the `config` module provides support for persisting configurations by having us *configure* our projects. At first it may feel natural to call `configure` an operation. There is, however, a conceptual problem: we don't really configure a target. And, perhaps after some meditation, it should become clear that what we are really doing is configuring operations on targets. For example, configuring updating a C++ project might involve detecting and saving information about the C++ compiler while configuring installing it may require specifying the installation directory.

In other words, `configure` is an operation on operation on targets – a meta-operation. And so in `build2` we have the concept of a *build system meta-operation*. If not specified explicitly (as part of the `buildspec`), the default is `perform`, which is to simply perform the operation.

Back to `config`, this module provides two meta-operations: `configure` which saves the configuration of a project into the `build/config.build` file as well as `disfigure` which removes it.

While the common meaning of the word *disfigure* is somewhat different to what we make it mean in this context, we still prefer it over the commonly suggested alternative (*deconfigure*) for the symmetry of their Latin *con-* ("together") and *dis-* ("apart") prefixes.

Let's say for the in source build of our `hello` project we want to use Clang and enable debug information. Without persistence we would have to repeat this configuration on every build system invocation:

```
$ cd hello/ # Change to project root.
$ b config.cxx=clang++ config.cxx.coptions=-g
```

Instead, we can configure our project with this information once and from then on invoke the build system without any arguments:

```
$ b configure config.cxx=clang++ config.cxx.coptions=-g
$ tree ./
./
|-- build/
|   |-- ...
|   |-- config.build
|   ...
$ b
$ b clean
$ b
...
```

To remove the persistent configuration we use the `disfigure` meta-operation:

```
$ b disfigure
```

Let's again configure our project and take a look at `config.build`:

```
$ b configure config.cxx=clang++ config.cxx.coptions=-g
```

```
$ cat build/config.build
```

```
config.cxx = clang++
config.cxx.poptions = [null]
config.cxx.coptions = -g
config.cxx.loptions = [null]
config.cxx.aoptions = [null]
config.cxx.libs = [null]
...
```

As you can see, it's just a buildfile with a bunch of variable assignments. In particular, this means you can tweak your build configuration by modifying this file with your favorite editor. Or, alternatively, you can adjust the configuration by reconfiguring the project:

```
$ b configure config.cxx=g++
```

```
$ cat build/config.build
```

```
config.cxx = g++
config.cxx.poptions = [null]
config.cxx.coptions = -g
config.cxx.loptions = [null]
config.cxx.aoptions = [null]
config.cxx.libs = [null]
...
```

Any variable value specified on the command line overrides those specified in the build-files. As a result, `config.cxx` was updated while the value of `config.cxx.coptions` was preserved.

To revert a configuration variable to its default value, list its name in the special `config.config.disfigure` variable. For example:

```
$ b configure config.config.disfigure=config.cxx
```

Command line variable overrides are also handy to adjust the configuration for a single build system invocation. For example, let's say we want to quickly check that our project builds with optimization but without permanently changing the configuration:

```
$ b config.cxx.coptions=-O3 # Rebuild with -O3.
$ b                        # Rebuild with -g.
```

Besides the various `*.?options` variables, we can also specify the "compiler mode" options as part of the compiler executable in `config.c` and `config.cxx`. Such options cannot be modified by buildfiles and they will appear last on the command lines. For example:

```
$ b configure config.cxx="g++ -m32"
```

The compiler mode options are also the correct place to specify *system-like* header (`-I`) and library (`-L`, `/LIBPATH`) search paths. Where by system-like we mean common installation directories like `/usr/include` or `/usr/local/lib` which may contain older versions of the libraries we are trying to build and/or use. By specifying these paths as part of the mode options (as opposed to `config.*.poptions` and `config.*.loptions`) we make sure they will be considered last, similar to the compiler's build-in search paths. For example:

```
$ b configure config.cxx="g++ -L/opt/install"
```

If we would like to prevent subsequent changes to the environment from affecting our build configuration, we can make it *hermetic* (see Hermetic Build Configurations for details):

```
$ b configure config.config.hermetic=true ...
```

One prominent use of hermetic configurations is to preserve the build environment of the Visual Studio development command prompt. That is, hermetically configuring our project in a suitable Visual Studio command prompt makes us free to build it from any other prompt or shell, IDE, etc.

We can also configure out of source builds of our projects. In this case, besides `config.build`, `configure` also saves the location of the source directory so that we don't have to repeat that either. Remember, this is how we used to build our `hello` out of source:

```
$ b hello/@hello-gcc/    config.cxx=g++
$ b hello/@hello-clang/ config.cxx=clang++
```

And now we can do:

```
$ b configure: hello/@hello-gcc/    config.cxx=g++
$ b configure: hello/@hello-clang/ config.cxx=clang++

$ hello-clang/
hello-clang/
.-- build/
|  -- bootstrap/
|    .-- src-root.build
.-- config.build

$ b hello-gcc/
$ b hello-clang/
$ b hello-gcc/ hello-clang/
```

One major benefit of an in source build is the ability to run executables as well as examine build and test output (test results, generated source code, documentation, etc) without leaving the source directory. Unfortunately, we cannot have multiple in source builds and as was discussed earlier, mixing in and out of source builds is not recommended.

To overcome this limitation `build2` has a notion of *forwarded configurations*. As the name suggests, we can configure a project's source directory to forward to one of its out of source builds. Once done, whenever we run the build system from the source directory, it will automatically build in the corresponded forwarded output directory. Additionally, it will *backlink* (using symlinks or another suitable mechanism) certain "interesting" targets (`exe{ }`, `doc{ }`) to the source directory for easy access. As an example, let's configure our `hello/` source directory to forward to the `hello-gcc/` build:

```
$ b configure: hello/@hello-gcc/,forward

$ cd hello/ # Change to project root.
$ b
c++ hello/cxx{hello} -> ../hello-gcc/hello/obje{hello}
ld ../hello-gcc/hello/exe{hello}
ln ../hello-gcc/hello/exe{hello} -> hello/
```

Notice the last line in the above listing: it indicates that `exe{hello}` from the out directory was backlinked in our project's source subdirectory:

```
$ tree ./
./
|-- build/
|   |-- bootstrap/
|   |   |-- out-root.build
|   |   |-- ...
|-- hello/
|   |-- ...
|   |-- hello -> ../../hello-gcc/hello/hello*
|-- ...

$ ./hello/hello
Hello World!
```

By default only `exe{ }` and `doc{ }` targets are backlinked. This, however, can be customized with the `backlink` target-specific variable.

1.4.2 Testing

The next module we load in `bootstrap.build` is `test` which defines the `test` operation. As the name suggests, this module provides support for running tests.

There are two types of tests that we can run with the `test` module: simple and scripted.

A simple test is just an executable target with the `test` target-specific variable set to `true`. For example:

```
exe{hello}: test = true
```

A simple test is executed once and in its most basic form (typical for unit testing) doesn't take any inputs nor produce any output, indicating success via the zero exit status. If we test our `hello` project with the above addition to the `buildfile`, then we will see the following output:

```
$ b test
test hello/exe{hello}
Hello, World!
```

While the test passes (since it exited with zero status), we probably don't want to see that `Hello, World!` every time we run it (this can, however, be quite useful when running examples). More importantly, we don't really test its functionality and if tomorrow our `hello` starts swearing rather than greeting, the test will still pass.

Besides checking its exit status we can also supply some basic information to a simple test (more common for integration testing). Specifically, we can pass command line options (`test.options`) and arguments (`test.arguments`) as well as input (`test.stdin`, used to supply test's `stdin`) and output (`test.stdout`, used to compare to test's `stdout`).

Let's see how we can use this to fix our `hello` test by making sure our program prints the expected greeting. First, we need to add a file that will contain the expected output, let's call it `test.out`:

```
$ ls -l hello/
hello.cxx
test.out
buildfile

$ cat hello/test.out
Hello, World!
```

Next, we arrange for it to be compared to our test's `stdout`. Here is the new `hello/buildfile`:

```
exe{hello}: {hxx cxx}{**}
exe{hello}: file{test.out}: test.stdout = true
```

The last line looks new. What we have here is a *prerequisite-specific variable* assignment. By setting `test.stdout` for the `file{test.out}` prerequisite of target `exe{hello}` we mark it as expected `stdout` output of *this* target (theoretically, we could have marked it as `test.input` for another target). Notice also that we no longer need the `test` target-specific variable; it's unnecessary if one of the other `test.*` variables is specified.

Now, if we run our test, we won't see any output:

```
$ b test
test hello/exe{hello}
```

And if we try to change the greeting in `hello.cxx` but not in `test.out`, our test will fail printing the `diff(1)` comparison of the expected and actual output:

```

$ b test
c++ hello/cxx{hello} -> hello/obje{hello}
ld hello/exe{hello}
test hello/exe{hello}
--- test.out
+++ -
@@ -1 +1 @@
-Hello, World!
+Hi, World!
error: test hello/exe{hello} failed

```

Notice another interesting thing: we have modified `hello.cxx` to change the greeting and our test executable was automatically rebuilt before testing. This happened because the `test` operation performs `update` as its *pre-operation* on all the targets to be tested.

Let's make our `hello` program more flexible by accepting the name to greet on the command line:

```

#include <iostream>

int main (int argc, char* argv[])
{
    if (argc < 2)
    {
        std::cerr << "error: missing name" << std::endl;
        return 1;
    }

    std::cout << "Hello, " << argv[1] << '!' << std::endl;
}

```

We can exercise its successful execution path with a simple test fairly easily:

```

exe{hello}: test.arguments = 'World'
exe{hello}: file{test.out}: test.stdout = true

```

What if we also wanted to test its error handling? Since simple tests are single-run, this won't be easy. Even if we could overcome this, having expected output for each test in a separate file will quickly become untidy. And this is where script-based tests come in. Testscript is `build2`'s portable language for running tests. It vaguely resembles Bash and is optimized for concise test implementation and fast, parallel execution.

Just to give you an idea (see Testscript Introduction for a proper introduction), here is what testing our `hello` program with Testscript would look like:

```

$ ls -l hello/
hello.cxx
testscript
buildfile

$ cat hello/buildfile

exe{hello}: {hxx cxx}{**} testscript

```

And this is the contents of `hello/testscript`:

```
: basics
:
$* 'World' >'Hello, World!'

: missing-name
:
$* 2>>EOE != 0
error: missing name
EOE
```

A couple of key points: The `test.out` file is gone with all the test inputs and expected outputs incorporated into `testscript`. To test an executable with Testscript, all we have to do is list the corresponding `testscript` file as its prerequisite (and which, being a fixed name, doesn't need an explicit target type, similar to `manifest`).

To see Testscript in action, let's say we've made our program more forgiving by falling back to a default name if one wasn't specified:

```
#include <iostream>

int main (int argc, char* argv[])
{
    const char* n (argc > 1 ? argv[1] : "World");
    std::cout << "Hello, " << n << '!' << std::endl;
}
```

If we forget to adjust the `missing-name` test, then this is what we could expect to see when running the tests:

```
$ b test
c++ hello/cxx{hello} -> hello/obje{hello}
ld hello/exe{hello}
test hello/exe{hello} + hello/testscript{testscript}
hello/testscript:7:1: error: hello/hello exit code 0 == 0
    info: stdout: hello/test-hello/missing-name/stdout
```

Testscript-based integration testing is the default setup for executable (`-t exe`) projects created by **`bdep-new (1)`**. Here is the recap of the overall layout:

```
hello/
|-- build/
|   |-- ...
|-- hello/
|   |-- hello.cxx
|   |-- testscript
|   |-- buildfile
|-- buildfile
-- manifest
```

For libraries (`-t lib`), however, the integration testing setup is a bit different. Here are the relevant parts of the layout:

```

libhello/
|-- build/
|   |-- ...
|-- libhello/
|   |-- hello.hxx
|   |-- hello.cxx
|   |-- export.hxx
|   |-- version.hxx.in
|   |-- buildfile
|-- tests/
|   |-- build/
|   |   |-- bootstrap.build
|   |   |-- root.build
|   |-- basics/
|   |   |-- driver.cxx
|   |   |-- buildfile
|   |-- buildfile
|-- buildfile
--- manifest

```

Specifically, there is no `testscript` in `libhello/`, the project's source subdirectory. Instead, we have the `tests/` subdirectory which itself looks like a project: it contains the `build/` subdirectory with all the familiar files, etc. In fact, `tests` is a *subproject* of our `libhello` project.

While we will be examining `tests` in greater detail later, in a nutshell, the reason it is a subproject is to be able to test an installed version of our library. By default, when `tests` is built as part of its parent project (called *amalgamation*), the locally built `libhello` library will be automatically imported. However, we can also configure a build of `tests` out of its amalgamation, in which case we can import an installed version of `libhello`. We will learn how to do all that as well as the underlying concepts (*subproject/amalgamation*, *import*, etc) in the coming sections.

Inside `tests/` we have the `basics/` subdirectory which contains a simple test for our library's API. By default it doesn't use Testscript but if you want to, you can. You can also rename `basics/` to something more meaningful and add more tests next to it. For example, if we were creating an XML parsing and serialization library, then our `tests/` could have the following layout:

```

tests/
|-- build/
|   |-- ...
|-- parser/
|   |-- ...
|-- serializer/
|   |-- ...
--- buildfile

```

Nothing prevents us from having the `tests/` subdirectory for executable projects. And it can be just a subdirectory or a subproject, the same as for libraries. Making it a subproject makes sense if your program has complex installation, for example, if its execution requires configuration and/or data files that need to be found, etc. For simple programs, however, testing the executable before installing it is usually sufficient.

For a general discussion of functional/integration and unit testing refer to the Tests section in the toolchain introduction. For details on the unit test support implementation see Implementing Unit Testing.

1.4.3 Installing

The `install` module defines the `install` and `uninstall` operations. As the name suggests, this module provides support for project installation.

Installation in `build2` is modeled after UNIX-like operation systems though the installation directory layout is highly customizable. While `build2` projects can import `build2` libraries directly, installation is often a way to "export" them in a form usable by other build systems.

The root installation directory is specified with the `config.install.root` configuration variable. Let's install our `hello` program into `/tmp/install`:

```
$ cd hello/ # Change to project root.
$ b install config.install.root=/tmp/install/
```

And see what we've got (executables are marked with *):

```
$ tree /tmp/install/

/tmp/install/
|-- bin/
|   |-- *hello
|-- share/
|   |-- doc/
|       |-- hello/
|           |-- manifest
```

Similar to the `test` operation, `install` performs `update` as a pre-operation for targets that it installs.

We can also configure our project with the desired `config.install.*` values so that we don't have to repeat them on every `install/uninstall`. For example:

```
$ b configure config.install.root=/tmp/install/
$ b install
$ b uninstall
```

Now let's try the same for `libhello` (symbolic link targets are shown with `->` and actual static/shared library names may differ on your operating system):

```
$ rm -r /tmp/install

$ cd libhello/ # Change to project root.

$ b install config.install.root=/tmp/install/

$ tree /tmp/install/

/tmp/install/
```

```

|-- include/
|   |-- libhello/
|       |-- hello.hxx
|       |-- export.hxx
|       |-- version.hxx
|-- lib/
|   |-- pkgconfig/
|       |-- libhello.pc
|       |-- libhello.shared.pc
|       |-- libhello.static.pc
|   |-- libhello.a
|   |-- libhello.so -> libhello-0.1.so
|   |-- libhello-0.1.so
|-- share/
|   |-- doc/
|       |-- libhello/
|           |-- manifest

```

As you can see, the library headers go into the customary `include/` subdirectory while static and shared libraries (and their `pkg-config(1)` files) – into `lib/`. Using this installation we should be able to import this library from other build systems or even use it in a manual build:

```
$ g++ -I/tmp/install/include -L/tmp/install/lib greet.cxx -lhello
```

If we want to install into a system-wide location like `/usr` or `/usr/local`, then we most likely will need to specify the `sudo(1)` program:

```
$ b config.install.root=/usr/local/ config.install.sudo=sudo
```

In `build2` only actual `install/uninstall` commands are executed with `sudo(1)`. And while on the topic of sensible implementations, `uninstall` can be generally trusted to work reliably.

The default installability of a target as well as where it is installed is determined by its target type. For example, `exe{}` is by default installed into `bin/`, `doc{}` – into `share/doc/<project>/`, and `file{}` is not installed.

We can, however, override these defaults with the `install` target-specific variable. Its value should be either special `false` indicating that the target should not be installed or the directory to install the target to. As an example, here is what the root buildfile from our `libhello` project looks like:

```

./: {*/ -build/} manifest

tests/: install = false

```

The first line we have already seen and the purpose of the second line should now be clear: it makes sure we don't try to install anything in the `tests/` subdirectory.

If the value of the `install` variable is not `false`, then it is normally a relative path with the first path component being one of these names:

name	default	override
----	-----	-----
root		config.install.root
data_root	root/	config.install.data_root
exec_root	root/	config.install.exec_root
bin	exec_root/bin/	config.install.bin
sbin	exec_root/sbin/	config.install.sbin
lib	exec_root/lib/	config.install.lib
libexec	exec_root/libexec/<project>/	config.install.libexec
pkgconfig	lib/pkgconfig/	config.install.pkgconfig
etc	data_root/etc/	config.install.etc
include	data_root/include/	config.install.include
include_arch	include/	config.install.include_arch
share	data_root/share/	config.install.share
data	share/<project>/	config.install.data
buildfile	share/build2/export/<project>/	config.install.buildfile
doc	share/doc/<project>/	config.install.doc
legal	doc/	config.install.legal
man	share/man/	config.install.man
man<N>	man/man<N>/	config.install.man<N>

Let's see what's going on here: The default install directory tree is derived from the `config.install.root` value but the location of each node in this tree can be overridden by the user that installs our project using the corresponding `config.install.*` variables (see the `install` module documentation for details on their meaning). In our buildfiles, in turn, we use the node names instead of actual directories. As an example, here is a buildfile fragment from the source subdirectory of our `libhello` project:

```
hxx{*}:
{
  install      = include/libhello/
  install.subdirs = true
}
```

Here we set the installation location for headers to be the `libhello/` subdirectory of the `include` installation location. Assuming `config.install.root` is `/usr/`, the `install` module will perform the following steps to resolve this relative path to the actual, absolute installation directory:

```
include/libhello/
data_root/include/libhello/
root/include/libhello/
/usr/include/libhello/
```

In the above buildfile fragment we also see the use of the `install.subdirs` variable. Setting it to `true` instructs the `install` module to recreate subdirectories starting from this point in the project's directory hierarchy. For example, if our `libhello/` source subdirectory had the `details/` subdirectory with the `utility.hxx` header, then this header would have been installed as

```
.../include/libhello/details/utility.hxx.
```

By default the generated `pkg-config` files will contain `install.include` and `install.lib` directories as header (`-I`) and library (`-L`) search paths, respectively. However, these can be customized with the `{c,cxx}.pkgconfig.{include,lib}` variables. For example, sometimes we may need to install headers into a subdirectory of the include directory but include them without this subdirectory:

```
# Install headers into hello/libhello/ subdirectory of, say,
# /usr/include/ but include them as <libhello/*>.
#
hxx{*}:
{
    install      = include/hello/libhello/
    install.subdirs = true
}

lib{hello}: cxx.pkgconfig.include = include/hello/
```

1.4.4 Distributing

The last module that we load in our `bootstrap.build` is `dist` which provides support for the preparation of source distributions and defines the `dist` meta-operation. Similar to `configure`, `dist` is a meta-operation rather than an operation because, conceptually, we are preparing a distribution for performing operations (like `update`, `test`) on targets rather than targets themselves.

The preparation of a correct distribution requires that all the necessary project files (sources, documentation, etc) be listed as prerequisites in the project's `buildfiles`.

You may wonder why not just use the export support offered by many version control systems? The main reason is that in most real-world projects version control repositories contain a lot more than what needs to be distributed. In fact, it is not uncommon to host multiple build system projects/packages in a single repository. As a result, with this approach we seem to inevitably end up maintaining an exclusion list, which feels backwards: why specify all the things we don't want in a new list instead of making sure the already existing list of things that we do want is complete? Also, once we have the complete list, it can be put to good use by other tools, such as editors, IDEs, etc.

The preparation of a distribution also requires an out of source build. This allows the `dist` module to distinguish between source and output targets. By default, targets found in `src` are included into the distribution while those in `out` are excluded. However, we can customize this with the `dist` target-specific variable.

As an example, let's prepare a distribution of our `hello` project using the out of source build configured in `hello-out/`. We use `config.dist.root` to specify the directory to write the distribution to:

```
$ b dist: hello-out/ config.dist.root=/tmp/dist

$ ls -l /tmp/dist
hello-0.1.0/
```

```
$ tree /tmp/dist/hello-0.1.0/
/tmp/dist/hello-0.1.0/
|-- build/
|   |-- bootstrap.build
|   |-- root.build
|-- hello/
|   |-- hello.cxx
|   |-- testscript
|   |-- buildfile
|-- buildfile
--- manifest
```

As we can see, the distribution directory includes the project version (from the version variable which, in our case, is extracted from manifest by the version module). Inside the distribution directory we have our project's source files (but, for example, without any .gitignore files that we may have had in hello/).

We can also ask the dist module to package the distribution directory into one or more archives and generate their checksum files for us. For example:

```
$ b dist: hello-out/ \
  config.dist.root=/tmp/dist \
  config.dist.archives="tar.gz zip" \
  config.dist.checksums=sha256

$ ls -l /tmp/dist
hello-0.1.0/
hello-0.1.0.tar.gz
hello-0.1.0.tar.gz.sha256
hello-0.1.0.zip
hello-0.1.0.zip.sha256
```

We can also configure our project with the desired config.dist.* values so we don't have to repeat them every time. For example:

```
$ b configure: hello-out/ config.dist.root=/tmp/dist ...
$ b dist
```

Let's now take a look at an example of customizing what gets distributed. Most of the time you will be using this mechanism to include certain targets from out. Here is a fragment from the libhello source subdirectory buildfile:

```
hxx{version}: in{version} $src_root/manifest
```

Our library provides the version.hxx header that the users can include to obtain its version. This header is generated by the version module from the version.hxx.in template. In essence, the version module takes the version value from our manifest, splits it into various components (major, minor, patch, etc) and then preprocesses the in{} file substituting these values (see the version module documentation for details). The end result is an automatically maintained version header.

Usually there is no need to include this header into the distribution since it will be automatically generated if and when necessary. However, we can if we need to. For example, we could be porting an existing project and its users could be expecting the version header to be

shipped as part of the archive. Here is how we can achieve this:

```
hxx{version}: in{version} $src_root/manifest
{
    dist = true
    clean = ($src_root != $out_root)
}
```

Because this header is an output target, we have to explicitly request its distribution with `dist=true`. Notice that we have also disabled its cleaning for the in source build so that the `clean` operation results in a state identical to distributed.

1.5 Target Importation

Recall that if we need to depend on a target defined in another `buildfile` within our project, then we simply include the said `buildfile` and reference the target. For example, if our `hello` included both an executable and a library in separate subdirectories next to each other:

```
hello/
|-- build/
|   |-- ...
|-- hello/
|   |-- ...
|   |-- buildfile
|-- libhello/
|   |-- ...
|   |-- buildfile
```

Then our executable `buildfile` could look like this:

```
include ../libhello/ # Include lib{hello}.

exe{hello}: {hxx cxx}{**} ../libhello/lib{hello}
```

What if instead `libhello` were a separate project? The inclusion approach would no longer work for two reasons: we don't know the path to `libhello` (after all, it's an independent project and can reside anywhere) and we can't assume the path to the `lib{hello}` target within `libhello` (the project directory layout can change).

To depend on a target from a separate project we use *importation* instead of inclusion. This mechanism is also used to depend on targets that are not part of any project, for example, installed libraries.

The importing project's side is pretty simple. This is what the above `buildfile` will look like if `libhello` were a separate project:

```
import libs = libhello%lib{hello}

exe{hello}: {hxx cxx}{**} $libs
```

The `import` directive is a kind of variable assignment that resolves a *project-qualified* relative target (`libhello%lib{hello}`) to an unqualified absolute target and stores it in the variable (`libs` in our case). We can then expand the variable (`$libs`), normally in the dependency declaration, to get the imported target.

If we needed to import several libraries, then we simply repeat the `import` directive, usually accumulating the result in the same variable, for example:

```
import libs = libformat%lib{format}
import libs += libprint%lib{print}
import libs += libhello%lib{hello}
```

```
exe{hello}: {hxx cxx}{**} $libs
```

Let's now try to build our `hello` project that uses imported `libhello`:

```
$ b hello/
error: unable to import target libhello%lib{hello}
info: use config.import.libhello command line variable to specify
its project out_root
```

While that didn't work out well, it does make sense: the build system cannot know the location of `libhello` or which of its builds we want to use. Though it does helpfully suggest that we use `config.import.libhello` to specify its out directory (`out_root`). Let's point it to `libhello` source directory to use its in source build (`out_root == src_root`):

```
$ b hello/ config.import.libhello=libhello/
c++ libhello/libhello/cxx{hello} -> libhello/libhello/objs{hello}
ld libhello/libhello/libs{hello}
c++ hello/hello/cxx{hello} -> hello/hello/obje{hello}
ld hello/hello/exe{hello}
```

And it works. Naturally, the importation mechanism works the same for out of source builds and we can persist the `config.import.*` variables in the project's configuration. As an example, let's configure Clang builds of the two projects out of source:

```
$ b configure: libhello/@libhello-clang/ config.cxx=clang++
$ b configure: hello/@hello-clang/ config.cxx=clang++ \
  config.import.libhello=libhello-clang/

$ b hello-clang/
c++ libhello/libhello/cxx{hello} -> libhello-clang/libhello/objs{hello}
ld libhello-clang/libhello/libs{hello}
c++ hello/hello/cxx{hello} -> hello-clang/hello/obje{hello}
ld hello-clang/hello/exe{hello}
```

If the corresponding `config.import.*` variable is not specified, `import` searches for a project in a couple of other places. First, it looks in the list of subprojects starting from the importing project itself and then continuing with its outer amalgamations and their subprojects (see Subprojects and Amalgamations for details on this subject).

We've actually seen an example of this search step in action: the `tests` subproject in `libhello`. The test imports `libhello` which is automatically found as an amalgamation containing this subproject.

To skip searching in subprojects/amalgamations and proceed directly to the rule-specific search (described below), specify the `config.import.*` variable with an empty value. For example:

```
$ b configure: ... config.import.libhello=
```

If the project being imported cannot be located using any of these methods, then `import` falls back to the rule-specific search. That is, a rule that matches the target may provide support for importing certain target types based on rule-specific knowledge. Support for importing installed libraries by the C++ link rule is a good example of this. Internally, the `cxx` module extracts the compiler's library search paths (that is, paths that would be used to resolve `-lfoo`) and then the link rule uses them to search for installed libraries. This allows us to use the same `import` directive regardless of whether we import a library from a separate build, from a subproject, or from an installation directory.

Importation of an installed library will work even if it is not a `build2` project. Besides finding the library itself, the link rule will also try to locate its `pkg-config(1)` file and, if present, extract additional compile/link flags from it (see [Importation of Installed Libraries](#) for details). The link rule also automatically produces `pkg-config(1)` files for libraries that it installs.

A common problem with importing and using third-party C/C++ libraries is compiler warnings. Specifically, we are likely to include their headers into our project's source files which means we may see warnings in such headers (which we cannot always fix) mixed with warnings in our code (which we should normally be able to fix). See [Compilation Internal Scope](#) for a mechanism to deal with this problem.

Let's now examine the exporting side of the importation mechanism. While a project doesn't need to do anything special to be found by `import`, it does need to handle locating the exported target (or targets; there could be several) within the project as well as loading their `buildfiles`. And this is the job of an *export stub*, the `build/export.build` file that you might have noticed in the `libhello` project:

```
libhello
|-- build/
|   |-- export.build
|   ...
|   ...
```

Let's take a look inside:

```
$out_root/
{
    include libhello/
}

export $out_root/libhello/$import.target
```

An export stub is a special kind of `buildfile` that bridges from the importing project into exporting. It is loaded in a special temporary scope outside of any project, in a "no man's land" so to speak. The only variables set on the temporary scope are `src_root` and `out_root` of the project being imported as well as `import.target` containing the name of the target being imported (without project qualification; that is, `lib{hello}` in our example).

Typically, an export stub will open the scope of the exporting project, load the `buildfile` that defines the target being exported and finally "return" the absolute target name to the importing project using the `export` directive. And this is exactly what the export stub in our `libhello` does.

We now have all the pieces of the importation puzzle in place and you can probably see how they all fit together. To summarize, when the build system sees the `import` directive, it looks for a project with the specified name. If found, it creates a temporary scope, sets the `src/out_root` variables to point to the project and `import.target` – to the target name specified in the `import` directive. And then it load the project's export stub in this scope. Inside the export stub we switch to the project's root scope, load its `buildfile` and then use the `export` directive to return the exported target. Once the export stub is processed, the build system obtains the exported target and assigns it to the variable specified in the `import` directive.

Our export stub is quite "loose" in that it allows importing any target defined in the project's source subdirectory `buildfile`. While we found it to be a good balance between strictness and flexibility, if you would like to "tighten" your export stubs, you can. For example:

```
if ($import.target == lib{hello})
  export $out_root/libhello/$import.target
```

If no `export` directive is executed in an export stub then the build system assumes that the target is not exported by the project and issues appropriate diagnostics.

Let's revisit the executable `buildfile` with which we started this section. Recall that it is for an executable that depends on a library which resides in the same project:

```
include ../libhello/ # Include lib{hello}.

exe{hello}: {hxx cxx}{**} ../libhello/lib{hello}
```

If `lib{hello}` is exported by this project, then instead of manually including its `buildfile` we can use *project-local importation*:

```
import lib = lib{hello}

exe{hello}: {hxx cxx}{**} $lib
```

The main advantage of project-local importation over inclusion is the ability to move things around without having to adjust locations in multiple places (the only place we need to do it is the export stub). This advantage becomes noticeable in more complex projects with a large number of components.

An import is project-local if the target being imported has no project name. Note that the target must still be exported in the project's export stub. In other words, project-local importation use the same mechanism as the normal import.

Another special type of importation is *ad hoc importation*. It is triggered if the target being imported has no project name and is either absolute or is a relative directory (in which case it is interpreted as relative to the importing scope). Semantically this is similar a normal import but with the location of the project being imported hard-coded into the `buildfile`. While this would be a bad idea in most case, sometimes we may want to create a special *glue buildfile* that "pulls" together several projects, usually for convenience of development.

One typical case that calls for such a glue buildfile is a multi-package project. For example, we may have a `hello` project (in a more general sense, as in a version control repository) that contains the `libhello` library and `hello` executable packages (which are independent build system projects):

```
hello/
|-- .git/
|-- hello/
|   |-- build/
|   |   |-- ...
|   |-- hello/
|   |   |-- ...
|   |-- buildfile
|   |-- manifest
|-- libhello/
|   |-- build/
|   |   |-- ...
|   |-- libhello/
|   |   |-- ...
|   |-- buildfile
|   |-- manifest
```

Notice that the root of this repository is not a build system project and we cannot, for example, just run the build system driver without any arguments to update all the packages. Instead we have to list them explicitly:

```
$ b hello/ libhello/
```

And that's inconvenient. To overcome this shortcoming we can turn the repository root into a simple build system project by adding a glue `buildfile` that imports (using *ad hoc importation*) and builds all the packages:

```
import pkgs = */
./: $pkgs
```

Unlike other import types, *ad hoc importation* does not rely (or require) an export stub. Instead, it directly loads a `buildfile` that could plausibly declare the target being imported.

In the unlikely event of a project-local importation of a directory target, it will have to be spelled with an explicit `dir{ }` target type, for example:

```
import d = dir{tests/}
```

1.6 Library Exportation and Versioning

By now we have examined and explained every line of every buildfile in our `hello` executable project. There are, however, still a few lines to be covered in the source subdirectory buildfile in `libhello`. Here it is in its entirety:

```
intf_libs = # Interface dependencies.
impl_libs = # Implementation dependencies.

lib{hello}: {hxx ixh txx cxx}{** -version} hxx{version} \
    $impl_libs $intf_libs

hxx{version}: in{version} $src_root/manifest

# Build options.
#
cxx.poptions += "-I$out_root" "-I$src_root"

obja{*}: cxx.poptions += -DLIBHELLO_STATIC_BUILD
objb{*}: cxx.poptions += -DLIBHELLO_SHARED_BUILD

# Export options.
#
lib{hello}:
{
    cxx.export.poptions = "-I$out_root" "-I$src_root"
    cxx.export.libs = $intf_libs
}

liba{hello}: cxx.export.poptions += -DLIBHELLO_STATIC
libs{hello}: cxx.export.poptions += -DLIBHELLO_SHARED

# For pre-releases use the complete version to make sure they cannot
# be used in place of another pre-release or the final version. See
# the version module for details on the version.* variable values.
#
if $version.pre_release
    lib{hello}: bin.lib.version = "-$version.project_id"
else
    lib{hello}: bin.lib.version = "-$version.major.$version.minor"

# Install into the libhello/ subdirectory of, say, /usr/include/
# recreating subdirectories.
#
{hxx ixh txx}{*}:
{
    install          = include/libhello/
    install.subdirs = true
}
```

Let's start with all those `cxx.export.*` variables. It turns out that merely exporting a library target is not enough for the importers of the library to be able to use it. They also need to know where to find its headers, which other libraries to link, etc. This information is carried in a set of target-specific `cxx.export.*` variables that parallel the `cxx.*` set and that together with the library's prerequisites constitute the *library metadata protocol*. Every time a source file that depends on a library is compiled or a binary is linked, this information is automatically extracted by the compile and link rules from the library dependency chain, recursively. And when the library is installed, this information is carried over to its `pkg-config(1)` file.

Similar to the `c.*` and `cc.*` sets discussed earlier, there are also `c.export.*` and `cc.export.*` sets.

Note, however, that there is no `*.export.options` since a library imposing compilation options on its consumers is bad practice (too coarse-grained, does not compose, etc). Instead, the recommended approach is to specify in the library documentation that it expects its consumers to use a certain compilation option. And if your library is unusable without exporting a compilation option and you are sure benefits outweigh the drawbacks, then you can specify it as part of `*.export.poptions` (it is still a good idea to prominently document this).

Here are the parts relevant to the library metadata protocol in the above buildfile:

```
intf_libs = # Interface dependencies.
impl_libs = # Implementation dependencies.

lib{hello}: ... $impl_libs $intf_libs

lib{hello}:
{
  cxx.export.poptions = "-I$out_root" "-I$src_root"
  cxx.export.libs = $intf_libs
}

liba{hello}: cxx.export.poptions += -DLIBHELLO_STATIC
libs{hello}: cxx.export.poptions += -DLIBHELLO_SHARED
```

As a first step we classify all our library dependencies into *interface dependencies* and *implementation dependencies*. A library is an interface dependency if it is referenced from our interface, for example, by including (importing) one of its headers (modules) from one of our (public) headers (modules) or if one of its functions is called from our inline or template functions. Otherwise, it is an implementation dependency.

To illustrate the distinction between interface and implementation dependencies, let's say we've reimplemented our `libhello` to use `libformat` to format the greeting and `libprint` to print it. Here is our new header (`hello.hxx`):

```
#include <libformat/format.hxx>

namespace hello
{
  void
```

```

say_hello_formatted (std::ostream&, const std::string& hello);

inline void
say_hello (std::ostream& o, const std::string& name)
{
    say_hello_formatted (o, format::format_hello ("Hello", name));
}
}

```

And this is the new source file (`hello.cxx`):

```

#include <libprint/print.hxx>

namespace hello
{
    void
    say_hello_formatted (ostream& o, const string& h)
    {
        print::print_hello (o, h);
    }
}

```

In this case, `libformat` is our interface dependency since we both include its header in our interface and call it from one of our inline functions. In contrast, `libprint` is only included and used in the source file and so we can safely treat it as an implementation dependency. The corresponding `import` directives in our `buildfile` will therefore look like this:

```

import intf_libs = libformat%lib{format}
import impl_libs = libprint%lib{print}

```

The preprocessor options (`poptions`) of an interface dependency must be made available to our library's users. The library itself should also be explicitly linked whenever our library is linked. All this is achieved by listing the interface dependencies in the `cxx.export.libs` variable:

```

lib{hello}:
{
    cxx.export.libs = $intf_libs
}

```

More precisely, the interface dependency should be explicitly linked if a user of our library may end up with a direct call to the dependency in one of their object files. Not linking such a library is called *underlinking* while linking a library unnecessarily (which can happen because we've included its header but are not actually calling any of its non-inline/template functions) is called *overlinking*. Underlinking is an error on some platforms while overlinking may slow down the process startup and/or waste its memory.

Note also that this only applies to shared libraries. In case of static libraries, both interface and implementation dependencies are always linked, recursively. Specifically, when linking a shared library, only libraries specified in its `*.export.libs` are linked. While when linking a static library, all its library prerequisites as well as those specified in `*.libs` are linked. Note that `*.export.libs` is not used when linking a static library since it is naturally assumed that all such libraries are also specified as library prerequisites or in `*.libs`.

The remaining lines in the library metadata fragment are:

```
lib{hello}:
{
  cxx.export.poptions = "-I$out_root" "-I$src_root"
}

liba{hello}: cxx.export.poptions += -DLIBHELLO_STATIC
libs{hello}: cxx.export.poptions += -DLIBHELLO_SHARED
```

The first line makes sure the users of our library can locate its headers by exporting the relevant `-I` options. The last two lines define the library type macros that are relied upon by the `export.hxx` header to properly setup symbol exporting.

The `liba{}` and `libs{}` target types correspond to the static and shared libraries, respectively. And `lib{}` is actually a target group that can contain one, the other, or both as its members.

Specifically, when we build a `lib{}` target, which members will be built is determined by the `config.bin.lib` variable with the `static`, `shared`, and `both` (default) possible values. So to only build a shared library we can run:

```
$ b config.bin.lib=shared
```

When it comes to linking `lib{}` prerequisites, which member is picked is controlled by the `config.bin.{exe,liba,libs}.lib` variables for the executable, static library, and shared library targets, respectively. Each contains a list of `shared` and `static` values that determine the linking preferences. For example, to build both shared and static libraries but to link executable to static libraries we can run:

```
$ b config.bin.lib=both config.bin.exe.lib=static
```

See the `bin` module documentation for more information.

Note also that we don't need to change anything in the above `buildfile` if our library is header-only. In `build2` this is handled dynamically and automatically based on the absence of source file prerequisites. In fact, the same library can be header-only on some platforms or in some configuration and "source-ful" in others.

In `build2` a header-only library (or a module interface-only library) is not a different kind of library compared to static/shared libraries but is rather a binary-less, or *binless* for short, static or shared library. So, theoretically, it is possible to have a library that has a binless static and a binary-ful (*binful*) shared variants. Note also that binless libraries can depend on binful libraries and are fully supported where the `pkg-config(1)` functionality is concerned.

One counter-intuitive aspect of having a binless library that depends on a system binful library, for example, `-lm`, is that you still have to specify the system library in both `*.export.libs` and `*.libs` because the latter is used when linking the static variant of the binless library. For example:

```
cxx.libs = -lm
lib{hello}: cxx.export.libs = -lm
```

If you are creating a new library with **bdep-new(1)** and are certain that it will always be binless and in all configurations, then you can produce a simplified buildfile by specifying the `binless` option, for example:

```
$ bdep new -l c++ -t lib,binless libheader-only
```

Let's now turn to the second subject of this section and the last unexplained bit in our buildfile: shared library versioning. Here is the relevant fragment:

```
if $version.pre_release
  lib{hello}: bin.lib.version = "-$version.project_id"
else
  lib{hello}: bin.lib.version = "-$version.major.$version.minor"
```

Shared library versioning is a murky, platform-specific area. Instead of trying to come up with a unified versioning scheme that few are likely to comprehend (similar to `autoconf`), `build2` provides a platform-independent versioning scheme as well as the ability to specify platform-specific versions in a native format.

The library version is specified with the `bin.lib.version` target-specific variable. Its value should be a sequence of @-pairs with the left hand side (key) being the platform name and the right hand side (value) being the version. An empty key (in which case @ can be omitted) signifies the platform-independent version (see the `bin` module documentation for the exact semantics). For example:

```
lib{hello}: bin.lib.version = -1.2 linux@3
```

While the interface for platform-specific versions is defined, their support is currently only implemented on Linux.

A platform-independent version is embedded as a suffix into the library name (and into its `soname` on relevant platforms) while platform-specific versions are handled according to the platform. Continuing with the above example, these would be the resulting shared library names on select platforms:

```
libhello.so.3      # Linux
libhello-1.2.dll   # Windows
libhello-1.2.dylib # Mac OS
```

With this background we can now explain what's going in our buildfile:

```
if $version.pre_release
  lib{hello}: bin.lib.version = "-$version.project_id"
else
  lib{hello}: bin.lib.version = "-$version.major.$version.minor"
```

Here we only use platform-independent library versioning. For releases we embed both major and minor version components assuming that patch releases are binary compatible. For pre-releases, however, we use the complete version to make sure it cannot be used in place of

another pre-release or the final version.

The `version.project_id` variable contains the project's (as opposed to package's), shortest "version id". See the `version` module documentation for details.

1.7 Subprojects and Amalgamations

In `build2` projects can contain other projects, recursively. In this arrangement the outer project is called an *amalgamation* and the inner – *subprojects*. In contrast to importation where we merely reference a project somewhere else, amalgamation is physical containment. It can be *strong* where the `src` directory of a subproject is within the amalgamating project or *weak* where only the `out` directory is contained.

There are several distinct use cases for amalgamations. We've already discussed the `tests/` subproject in `libhello`. To recap: traditionally, it is made a subproject rather than a subdirectory to support building it as a standalone project in order to test library installations.

As discussed in Target Importation, subprojects and amalgamations (as well as their subprojects, recursively) are automatically considered when resolving imports. As a result, amalgamation can be used to *bundle* dependencies to produce an external dependency-free distribution. For example, if our `hello` project imports `libhello`, then we could copy the `libhello` project into `hello`, for example:

```
$ tree hello/
hello/
|-- build/
|   |-- ...
|-- hello/
|   |-- hello.cxx
|   |-- ...
|-- libhello/
|   |-- build/
|       |-- ...
|       |-- libhello/
|           |-- hello.hxx
|           |-- hello.cxx
|           |-- ...
|       |-- tests/
|           |-- ...
|       |-- buildfile
|-- buildfile

$ b hello/
c++ hello/libhello/libhello/cxx{hello} ->
    hello/libhello/libhello/objs{hello}
ld hello/libhello/libhello/libs{hello}
c++ hello/hello/cxx{hello} -> hello/hello/obje{hello}
ld hello/hello/exe{hello}
```

Note, however, that while project bundling can be useful in certain cases, it does not scale as a general dependency management solution. For that, independent packaging and proper dependency management are the appropriate mechanisms.

By default `build2` looks for subprojects only in the root directory of a project. That is, every root subdirectory is examined to see if it itself is a project root. If you need to place a subproject somewhere else in your project's directory hierarchy, then you will need to specify its location (and of all other subprojects) explicitly with the `subprojects` variable in `bootstrap.build`. For example, if above we placed `libhello` into the `extras/` subdirectory of `hello`, then our `bootstrap.build` would need to start like this:

```
project = hello
subprojects = extras/libhello/
...
```

Note also that while importation of specific targets from subprojects is always performed, whether they are loaded and built as part of the overall project build is controlled using the standard subdirectories inclusion and dependency mechanisms. Continuing with the above example, if we adjust the root buildfile in `hello` to exclude the `extras/` subdirectory from the build:

```
./: {*/ -build/ -extras/}
```

Then while we can still import `libhello` from any buildfile in our project, the entire `libhello` (for example, its tests) will never be built as part of the `hello` build.

Similar to subprojects we can also explicitly specify the project's amalgamation with the `amalgamation` variable (again, in `bootstrap.build`). This is rarely necessary except if you want to prevent the project from being amalgamated, in which case you should set it to the empty value.

If either of these variables is not explicitly set, then they will contain the automatically discovered values.

Besides affecting importation, another central property of amalgamation is configuration inheritance. As an example, let's configure the above bundled `hello` project in its `src` directory:

```
$ b configure: hello/ config.cxx=clang++ config.cxx.coptions=-g

$ tree
hello/
|-- build/
|   |-- config.build
|   |-- ...
|-- libhello/
|   |-- build/
|   |   |-- config.build
|   |   |-- ...
|   |-- ...
|-- ...
```

As you can see, we now have the `config.build` files in both projects' `build/` subdirectories. If we examine the amalgamation's `config.build`, we will see the familiar picture:

```
$ cat hello/build/config.build
```

```
config.cxx = clang++
config.cxx.poptions = [null]
config.cxx.coptions = -g
config.cxx.loptions = [null]
config.cxx.aoptions = [null]
config.cxx.libs = [null]
```

```
...
```

The subproject's `config.build`, however, is pretty much empty:

```
$ cat hello/libhello/build/config.build
```

```
# Base configuration inherited from ../
```

As the comment suggests, the base configuration is inherited from the outer project. We can, however, override some values if we need to. For example (note that we are re-configuring the `libhello` subproject):

```
$ b configure: hello/libhello/ config.cxx.coptions=-O2
```

```
$ cat hello/libhello/build/config.build
```

```
# Base configuration inherited from ../
```

```
config.cxx.coptions = -O2
```

This configuration inheritance combined with import resolution is behind the most common use of amalgamations in `build2` – shared build configurations. Let's say we are developing multiple projects, for example, `hello` and `libhello` that it imports:

```
$ ls -l
hello/
libhello/
```

And we want to build them with several compilers, let's say GCC and Clang. As we have already seen in `Configuring`, we can configure several out of source builds for each compiler, for example:

```
$ b configure: libhello/@libhello-gcc/ config.cxx=g++
$ b configure: libhello/@libhello-clang/ config.cxx=clang++
```

```
$ b configure: hello/@hello-gcc/ \
               config.cxx=g++ \
               config.import.libhello=libhello-gcc/
```

```
$ b configure: hello/@hello-clang/ \
               config.cxx=clang++ \
               config.import.libhello=libhello-clang/
```

```
$ ls -l
hello/
hello-gcc/
hello-clang/
libhello/
libhello-gcc/
libhello-clang/
```

Needless to say, this is a lot of repetitive typing. Another problem is future changes to the configurations. If, for example, we need to adjust compile options in the GCC configuration, then we will have to (remember to) do it in both places.

You can probably sense where this is going: why not create two shared build configurations (that is, amalgamations), one for GCC and one for Clang, within each of which we build both of our projects (as their subprojects)? This is how we can do that:

```
$ b create: build-gcc/,cc config.cxx=g++
$ b create: build-clang/,cc config.cxx=clang++

$ b configure: libhello/@build-gcc/libhello/
$ b configure: hello/@build-gcc/hello/

$ b configure: libhello/@build-clang/libhello/
$ b configure: hello/@build-clang/hello/

$ ls -l
hello/
libhello/
build-gcc/
build-clang/
```

Let's explain what's going on here. First, we create two build configurations using the `create` meta-operation. These are real `build2` projects just tailored for housing other projects as subprojects. In `create`, after the directory name, we specify the list of modules to load in the project's `root.build`. In our case we specify `cc` which is a common module for C-based languages (see **b (1)** for details on `create` and its parameters).

When creating build configurations it is a good idea to get into the habit of using the `cc` module instead of `c` or `cxx` since with more complex dependency chains we may not know whether every project we build only uses C or C++. In fact, it is not uncommon for a C++ project to have C implementation details and even the other way around (yes, really, there are C libraries with C++ implementations).

Once the configurations are ready we simply configure our `libhello` and `hello` as subprojects in each of them. Note that now we neither need to specify `config.cxx`, because it will be inherited from the amalgamation, nor `config.import.*`, because the import will be automatically resolved to a subproject.

Now, to build a specific project in a particular configuration we simply build the corresponding subdirectory. We can also build the entire build configuration if we want to. For example:

```
$ b build-gcc/hello/
$ b build-clang/
```

In case you've already looked into **bpkg (1)** and/or **bdep (1)**, their build configurations are actually these same amalgamations (created underneath with the `create` meta-operation) and their packages are just subprojects. And with this understanding you are free to interact with them directly using the build system interface.

1.8 Buildfile Language

By now we should have a good overall sense of what writing `buildfiles` feels like. In this section we will examine the language in slightly more detail and with more precision.

Buildfile is primarily a declarative language with support for variables, pure functions, repetition (`for`, `while` loops), conditional inclusion/exclusion (`if-else`), and pattern matching (`switch`). At the lexical level, `buildfiles` are UTF-8 encoded text restricted to the Unicode graphic characters, tabs (`\t`), carriage returns (`\r`), and line feeds (`\n`).

Buildfile is a line-oriented language. That is, every construct ends at the end of the line unless escaped with line continuation (trailing `\`). For example:

```
exe{hello}: {hxx cxx}{**} \
    $libs
```

Some lines may start a *block* if followed by `{` on the next line. Such a block ends with a closing `}` on a separate line. Some types of blocks can nest. For example:

```
if ($cxx.target.class == 'windows')
{
    if ($cxx.target.system == 'ming32')
    {
        ...
    }
}
```

A comment starts with `#` and everything from this character and until the end of the line is ignored. A multi-line comment starts with `#\` on a separate line and ends with the same character sequence, again on a separate line. For example:

```
# Single line comment.

info 'Hello, World!' # Trailing comment.

#\
Multi-
line
comment.
#\
```

The three primary Buildfile constructs are dependency declaration, directive, and variable assignment. We've already used all three but let's see another example:

```
include ../libhello/                                # Directive.

exe{hello}: {hxx cxx}{**} ../libhello/lib{hello}    # Dependency.

cxx.poptions += -DNDEBUG                             # Variable.
```

There is also the scope opening (we've seen one in `export.build`) as well as target-specific and prerequisite-specific variable assignment blocks. The latter two are used to assign several entity-specific variables at once. For example:

```

details/                                # Scope.
{
    hxx{*}: install = false
}

lib{hello}:                             # Target-specific.
{
    cxx.export.poptions = "-I$src_root"
    cxx.export.libs = $intf_libs
}

exe{test}: file{test.roundtrip}:        # Prerequisite-specific.
{
    test.stdin  = true
    test.stdout = true
}

```

Variable assignment blocks can be combined with dependency declarations, for example:

```

h{config}: in{config}
{
    in.symbol = '@'
    in.mode = lax

    SYSTEM_NAME = $c.target.system
    SYSTEM_PROCESSOR = $c.target.cpu
}

```

In case of a dependency chain, if the chain ends with a colon (:), then the block applies to the last set of prerequisites. Otherwise, it applies to the last set of targets. For example:

```

./: exe{test}: cxx{main}
{
    test = true          # Applies to the exe{test} target.
}

./: exe{test}: libue{test}:
{
    bin.whole = false    # Applies to the libue{test} prerequisite.
}

```

All prerequisite-specific variables must be assigned at once as part of the dependency declaration since repeating the same dependency again duplicates the prerequisite rather than references the already existing one.

There is also the target type/pattern-specific variable assignment block, for example:

```

exe{*.test}:
{
    test = true
    install = false
}

```

See Variables for a more detailed discussion of variables.

Each `buildfile` is processed linearly with directives executed and variables expanded as they are encountered. However, certain variables, for example `cxx.poptions`, are also expanded by rules during execution in which case they will "see" the final value set in the `buildfile`.

Unlike GNU `make` (1), which has deferred (`=`) and immediate (`:=`) variable assignments, all assignments in `build2` are immediate. For example:

```
x = x
y = $x
x = X
info $y # Prints 'x', not 'X'.
```

1.8.1 Expansion and Quoting

While we've discussed variable expansion and lookup earlier, to recap, to get the variable's value we use `$` followed by its name. The variable name is first looked up in the current scope (that is, the scope in which the expansion was encountered) and, if not found, in the outer scopes, recursively.

There are two other kinds of expansions: function calls and evaluation contexts, or *eval contexts* for short. Let's start with the latter since function calls are built on top of eval contexts.

An eval context is essentially a fragment of a line with additional interpretations of certain characters to support value comparison, logical operators, and a few other constructs. Eval contexts begin with `(`, end with `)`, and can nest. Here are a few examples:

```
info ($src_root != $out_root)           # Prints true or false.
info ($src_root == $out_root ? 'in' : 'out') # Prints in or out.

macos = ($cxx.target.class == 'macos') # Assigns true or false.
linux = ($cxx.target.class == 'linux')  # Assigns true or false.

if ($macos || $linux) # Also eval context.
...
```

Below is the eval context grammar that shows supported operators and their precedence.

```
eval:      '(' (eval-comma | eval-qual)? ')'
eval-comma: eval-ternary '(' eval-ternary)*
eval-ternary: eval-or '(' '?' eval-ternary ':' eval-ternary)?
eval-or:    eval-and '(' '|' eval-and)*
eval-and:   eval-comp '(' '&&' eval-comp)*
eval-comp:  eval-value '(' ('==' | '!=' | '<' | '>' | '<=' | '>=') eval-value)*
eval-value: value-attributes? (<value> | eval | '!' eval-value)
eval-qual:  <name> ':' <name>

value-attributes: '[' <key-value-pairs> ']'
```

Note that `?:` (ternary operator) and `!` (logical not) are right-associative. Unlike C++, all the comparison operators have the same precedence. A qualified name cannot be combined with any other operator (including ternary) unless enclosed in parentheses. The `eval` option in the

`eval-value` production shall contain a single value only (no commas).

Additionally, the ``` (backtick) and `|` (bitwise or) tokens are reserved for future support of arithmetic evaluation contexts and evaluation pipelines, respectively.

A function call starts with `$` followed by its name and an eval context listing its arguments. Note that there is no space between the name and `(`. For example:

```
x =
y = Y

info $empty($x) # true
info $empty($y) # false

if $regex.match($y, '[A-Z]')
...

p = $src_base/foo.txt

info $path.leaf($src_base)           # foo.txt
info $path.directory($src_base)      # $src_base
info $path.base($path.leaf($src_base)) # foo
```

Note that the majority of functions in `build2` are *pure* in a sense that they do not alter the build state in any way (see Functions for details).

Functions in `build2` are currently defined either by the build system core or build system modules and are implemented in C++. In the future it will be possible to define custom functions in `buildfiles` (also in C++).

Variable and function names follow the C identifier rules. We can also group variables into namespaces and functions into families by combining multiple identifiers with `..`. These rules are used to determine the end of the variable name in expansions. If, however, a name is recognized as being longer than desired, then we can use the eval context to explicitly specify its boundaries. For example:

```
base = foo
name = $(base).txt
```

What is the structure of a variable value? Consider this assignment:

```
x = foo bar
```

The value of `x` could be a string, a list of two strings, or something else entirely. In `build2` the fundamental, untyped value is a *list of names*. A value can be typed to something else later but it always starts as a list of names. So in the above example we have a list of two names, `foo` and `bar`, the same as in this example (notice the extra spaces):

```
x = foo    bar
```

The motivation behind going with a list of names instead of a string or a list of strings is that at its core we are dealing with targets and their prerequisites and it would be natural to make the representation of their names (that is, the way we refer to them) the default. Consider the following two examples; it would be natural for them to mean the same thing:

```
exe{hello}: {hxx cxx}{**}

prereqs = {hxx cxx}{**}
exe{hello}: $prereqs
```

Note also that the name semantics was carefully tuned to be *reversible* to its syntactic representation for common non-name values, such as paths, command line options, etc., that are usually found in buildfiles.

To get to individual elements of a list, an expansion can be followed by a subscript. Note that subscripts are only recognize inside evaluation contexts and there should be no space between the expansion and [. For example:

```
x = foo bar

info ($x[0])                # foo
info ($regex.split('foo bar', ' ', '')[1]) # bar
```

Names are split into a list at whitespace boundaries with certain other characters treated as syntax rather than as part of the value. Here are a few examples:

```
x = $y           # expansion
x = (a == b)     # eval context
x = {foo bar}    # name generation
x = [null]       # attributes
x = name@value   # pairs
x = # start of a comment
```

The complete set of syntax characters is `$ ( ) { } @ # " '` plus space and tab, as well as `[ ]`, but only in certain contexts (see Attributes for details). If instead we need these characters to appear literally as part of the value, then we either have to *escape* or *quote* them.

Additionally, `*?[` will be treated as wildcards in name patterns (see Name Patterns for details). Note that this treatment can only be inhibited with quoting and not escaping.

While name patterns are recognized inside evaluation contexts, in certain cases the `?[` characters are treated as part of the ternary operator and value subscript, respectively. In such case, to be treat as wildcards rather than as syntax, these characters have to be escaped, for example:

```
x = (foo.\?xx)
y = ($foo\[123].txt)
```

To escape a special character, we prefix it with a backslash (`\`; to specify a literal backslash, double it). For example:

```
x = \$
y = C:\\Program\\ Files
```

Similar to UNIX shells, `build2` supports single (`' '`) and double (`" "`) quoting with roughly the same semantics. Specifically, expansions (variable, function call, and eval context) and escaping are performed inside double-quoted strings but not in single-quoted. Note also that quoted strings can span multiple lines with newlines treated literally (unless escaped in double-quoted strings). For example:

```
x = "(a != b)" # true
y = '(a != b)' # (a != b)

x = "C:\\Program Files"
y = 'C:\\Program Files'

t = 'line one
line two
line three'
```

Since quote characters are also part of the syntax, if you need to specify them literally in the value, then they will either have to be escaped or quoted. For example:

```
cxx.poptions += -DOUTPUT=' "debug"'
cxx.poptions += -DTARGET=\\ "$cxx.target\\"
```

An expansion can be one of two kinds: *spliced* or *concatenated*. In a spliced expansion the variable, function, or eval context is separated from other text with whitespaces. In this case, as the name suggests, the resulting list of names is spliced into the value. For example:

```
x = 'foo fox'
y = bar $x baz # Three names: 'bar' 'foo fox' 'baz'.
```

This is an important difference compared to the semantics of UNIX shells where the result of expansion is re-parsed. In particular, this is the reason why you won't see quoted expansions in `buildfiles` as often as in (well-written) shell scripts.

In a concatenated expansion the variable, function, or eval context are combined with unseparated text before and/or after the expansion. For example:

```
x = 'foo fox'
y = bar$(x)foz # Single name: 'barfoo foxbaz'
```

A concatenated expansion is typed unless it is quoted. In a typed concatenated expansion the parts are combined in a type-aware manner while in an untyped – literally, as string. To illustrate the difference, consider this `buildfile` fragment:

```
info $src_root/foo.txt
info "$src_root/foo.txt"
```

If we run it on a UNIX-like operating system, we will see two identical lines, along these lines:

```
/tmp/test/foo.txt
/tmp/test/foo.txt
```

However, if we run it on Windows (which uses backslashes as directory separators), we will see the output along these lines:

```
C:\test\foo.txt
C:\test/foo.txt
```

The typed concatenation resulted in a native directory separator because `dir_path` (the `src_root` type) did the right thing.

Not every typed concatenation is well defined and in certain situations we may need to force untyped concatenation with quoting. Options specifying header search paths (`-I`) are a typical case, for example:

```
cxx.poptions += "-I$out_root" "-I$src_root"
```

If we were to remove the quotes, we would see the following error:

```
buildfile:6:20: error: no typed concatenation of <untyped> to dir_path
      info: use quoting to force untyped concatenation
```

1.8.2 Conditions (**if-else**)

The `if` directive can be used to conditionally exclude buildfile fragments from being processed. The conditional fragment can be a single (separate) line or a block with the initial `if` optionally followed by a number of `elif` directives and a final `else`, which together form the `if-else` chain. An `if-else` block can contain nested `if-else` chains. For example:

```
if ($cxx.target.class == 'linux')
  info 'linux'
elif ($cxx.target.class == 'windows')
{
  if ($cxx.target.system == 'mingw32')
    info 'windows-mingw'
  elif ($cxx.target.system == 'win32-msvc')
    info 'windows-msvc'
  else
    info 'windows-other'
}
else
  info 'other'
```

The `if` and `elif` directive names must be followed by an expression that expands to a single, literal `true` or `false`. This can be a variable expansion, a function call, an `eval` context, or a literal value. For example:

```
if $version.pre_release
  ...

if $regex.match($x, '[A-Z]')
  ...
```

```

if ($cxx.target.class == 'linux')
    ...

if false
{
    # disabled fragment
}

x = X
if $x # Error, must expand to true or false.
    ...

```

There are also `if!` and `elif!` directives which negate the condition that follows (note that there is no space before `!`). For example:

```

if! $version.pre_release
    ...
elif! $regex.match($x, '[A-Z]')
    ...

```

Besides these general `if`-directives there is also a number of specialized shortcuts for checking whether a value is/is-not `null` or `empty`:

```

ifn ...      ~  if $null(...)
ife ...      ~  if $empty(...)

ifn! ...     ~  if! $null(...)
ife! ...     ~  if! $empty(...)

elifn ...    ~  elif $null(...)
elifn! ...   ~  elif $empty(...)

elifn! ...   ~  elif! $null(...)
elifn! ...   ~  elif! $empty(...)

```

For example, the following two constructs are equivalent:

```

if $null($foo)
    ...
elif! $empty($bar)
    ...

ifn $foo
    ...
elifn! $bar
    ...

```

Note that a `null` value is considered `empty`.

Note also that there is no notion of variable locality in `if-else` blocks and any value set inside is visible outside. For example:

```

if true
{
    x = X
}

info $x # Prints 'X'.

```

The `if-else` chains should not be used for conditional dependency declarations since this would violate the expectation that all of the project's source files are listed as prerequisites, irrespective of the configuration. Instead, use the special `include` prerequisite-specific variable to conditionally include prerequisites into the build. For example:

```

# Incorrect.
#
if ($cxx.target.class == 'linux')
    exe{hello}: cxx{hello-linux}
elif ($cxx.target.class == 'windows')
    exe{hello}: cxx{hello-win32}

# Correct.
#
exe{hello}: cxx{hello-linux}: include = ($cxx.target.class == 'linux')
exe{hello}: cxx{hello-win32}: include = ($cxx.target.class == 'windows')

```

1.8.3 Pattern Matching (**switch**)

The `switch` directive is similar to `if-else` in that it allows us to conditionally exclude buildfile fragments from being processed. The difference is in the way the conditions are structured: while in `if-else` we can do arbitrary tests, in `switch` we match one or more values against a set of patterns. For instance, this is how we can reimplement the first example from Conditionals (`if-else`) using `switch`:

```

switch $cxx.target.class, $cxx.target.system
{
    case 'linux'
        info 'linux'

    case 'windows', 'mingw32'
        info 'windows-mingw'

    case 'windows', 'win32-msvc'
        info 'windows-msvc'

    case 'windows'
        info 'windows-other'

    default
        info 'other'
}

```

Similar to `if-else`, the conditional fragment can be a single (separate) line or a block with a zero or more `case` lines/blocks optionally followed by `default`. A `case-default` block can contain nested `switch` directives (though it is often more convenient to use multiple values in a single `switch`, as shown above). For example:

```

switch $cxx.target.class
{
    ...
    case 'windows'
    {
        switch $cxx.target.system
        {
            case 'mingw32'
                info 'windows-mingw'

            case 'win32-msvc'
                info 'windows-msvc'

            default
                info 'windows-other'
        }
    }
    ...
}

```

All the case fragments are tried in the order specified with the first that matches evaluated and all the others ignored (that is, there is no explicit `break` nor the ability to fall through). If none of the case patterns matched and there is the `default` fragment, then it is evaluated. Multiple case lines can be specified for a single conditional fragment. For example:

```

switch $cxx.target.class, $cxx.id
{
    case 'windows', 'msvc'
    case 'windows', 'clang'
        info 'msvcrt'
}

```

The `switch` directive name must be followed by one or more *value expressions* separated with a comma (,). Similarly, the `case` directive name must be followed by one or more *pattern expressions* separated with a comma (,). These expressions can be variable expansions, function calls, eval contexts, or literal values.

If multiple values/patterns are specified, then all the case patterns must match in order for the corresponding fragment to be evaluated. However, if some trailing patterns are omitted, then they are considered as matching. For example:

```

switch $cxx.target.class, $cxx.target.system
{
    case 'windows', 'mingw32'
        info 'windows-mingw'

    case 'windows', 'win32-msvc'
        info 'windows-msvc'

    case 'windows'
        info 'windows-other'
}

```

The first pattern in the pattern expression can be optionally followed by one or more alternative patterns separated by a vertical bar (|). Only one of the alternatives need to match in order for the whole pattern expression to be considered as matching. For example:

```
switch $cxx.id
{
  case 'clang' | 'clang-apple'
    ...
}
```

The value in the value expression can be optionally followed by a colon (:) and a *match function*. If the match function is not specified, then equality is used by default. For example:

```
switch $cxx.target.cpu: regex.match
{
  case 'i[3-6]86'
    ...

  case 'x86_64'
    ...
}
```

The match function name can be optionally followed by additional values that are passed as the third argument to the match function. This is normally used to specify additional match flags, for example:

```
switch $cxx.target.cpu: regex.match icase
{
  ...
}
```

Other commonly used match functions are `regex.search()` (similar to `regex.match()` but searches for any match rather than matching the whole value), `path.match()` (match using shell wildcard patterns) and `string.icasecmp()` (match using equality but ignoring case). Additionally, any other function that takes the value as its first argument, the pattern as its second, and returns `bool` can be used as a match function.

Note that there is no special wildcard or match-anything pattern at the syntax level. In most common cases the desired semantics can be achieved with `default` and/or by omitting trailing patterns. If you do need it, then we recommend using `path.match()` and its `*` wildcard. For example:

```
switch $cxx.target.class: path.match, \
  $cxx.target.system: path.match, \
  $cxx.id: path.match
{
  case 'windows', '*', 'clang'
    ...
}
```

Note also that similar to `if-else`, there is no notion of variable locality in the `switch` and `case-default` blocks and any value set inside is visible outside. Additionally, the same considerations about conditional dependency declarations apply.

1.8.4 Repetitions (**for**, **while**)

The `for` directive can be used to repeat the same `buildfile` fragment multiple times, once for each element of a list. The fragment to repeat can be a single (separate) line or a block, which together form the `for` loop. A `for` block can contain nested `for` loops. For example:

```
for n: foo bar baz
{
    exe{$n}: cxx{$n}
}
```

The `for` directive name must be followed by the variable name (called *loop variable*) that on each iteration will be assigned the corresponding element, `:`, and an expression that expands to a potentially empty list of values. This can be a variable expansion, a function call, an eval context, or a literal list as in the above fragment. Here is a somewhat more realistic example that splits a space-separated environment variable value into names and then generates a dependency declaration for each of them:

```
for n: $regex.split($getenv(NAMES), ' ', '')
{
    exe{$n}: cxx{$n}
}
```

Note also that there is no notion of variable locality in `for` blocks and any value set inside is visible outside. At the end of the iteration the loop variable contains the value of the last element, if any. For example:

```
for x: x X
{
    y = Y
}

info $x # Prints 'X'.
info $y # Prints 'Y'.
```

Similarly, the `while` directive can be used to repeat the same `buildfile` fragment multiple times, while a certain condition is met. The fragment is repeated as long as the condition expression evaluates to the literal `true` value. The expression can be a variable expansion, a function call, an eval context, or a literal value. For example:

```
i = [uint64] 0
while ($i != 10)
{
    ...

    i += 1
}
```

There is also the `while!` variant which negates the condition (note that there is no space before `!`). For example:

```
done = [bool] false
while! $done
{
    ...

    if (...)
        done = true
}
```

The normal control flow of the `for` and `while` loops can be altered with the `continue` and `break` directives. The `continue` directive interrupts the current iteration of the nearest containing loop and starts the next iteration, subject to the normal `for` and `while` loop termination semantics. The `break` directive interrupts the current iteration and terminates the loop. For example:

```
for n: foo bar baz
{
    if (...)
        continue

    if (...)
        break

    ...
}
```

An example of a `while` loop:

```
i = [uint64] 0
while ($i != 10)
{
    if (...)
    {
        i += 1
        continue
    }

    if (...)
        break

    ...

    i += 1
}
```

1.9 Implementing Unit Testing

As an example of how many of these features fit together to implement more advanced functionality, let's examine a `buildfile` that provides support for unit testing. This support is added by the **`bdep-new(1)`** command if we specify the `unit-tests` option when creating executable (`-t exe,unit-tests`) or library (`-t lib,unit-tests`) projects. Here is the source subdirectory `buildfile` of an executable created with this option:

```

./: exe{hello}: libue{hello}: {hxx cxx}{** -**.test...}

# Unit tests.
#
exe{*.test}:
{
    test = true
    install = false
}

for t: cxx{**.test...}
{
    d = $directory($t)
    n = $name($t)...

    ./: $d/exe{$n}: $t $d/hxx{+$n} $d/testscript{+$n}
    $d/exe{$n}: libue{hello}: bin.whole = false
}

cxx.poptions += "-I$out_root" "-I$src_root"

```

The basic idea behind this unit testing arrangement is to keep unit tests next to the source code files that they test and automatically recognize and build them into test executables without having to manually list each in the `buildfile`. Specifically, if we have `hello.hxx` and `hello.cxx`, then to add a unit test for this module all we have to do is drop the `hello.test.cxx` source file next to them and it will be automatically picked up, built into an executable, and run during the `test` operation.

As an example, let's say we've renamed `hello.cxx` to `main.cxx` and factored the printing code into the `hello.hxx/hello.cxx` module that we would like to unit-test. Here is the new layout:

```

hello/
|-- build
|   |-- ...
|-- hello
|   |-- hello.cxx
|   |-- hello.hxx
|   |-- hello.test.cxx
|   |-- main.cxx
|   |-- buildfile
|-- ...

```

Let's examine how this support is implemented in our `buildfile`, line by line. Because now we link `hello.cxx` object code into multiple executables (unit tests and the `hello` program itself), we have to place it into a *utility library*. This is what the first line does (it has to explicitly list `exe{hello}` as a prerequisite of the default targets since we now have multiple targets that should be built by default):

```

./: exe{hello}: libue{hello}: {hxx cxx}{** -**.test...}

```

A utility library (**u** in `libue`) is a static library that is built for a specific type of a *primary target* (**e** in `libue` for executable). If we were building a utility library for a library then we would have used the `libul{}` target type instead. In fact, this would be the only difference in the above unit testing implementation if it were for a library project instead of an

executable:

```
./: lib{hello}: libul{hello}: {hxx cxx}{** -**.test...}

...

# Unit tests.
#
...

for t: cxx{**.test...}
{
    ...

    $d/exe{$n}: libul{hello}: bin.whole = false
}
```

Going back to the first three lines of the executable buildfile, notice that we had to exclude source files in the `*.test.cxx` form from the utility library. This makes sense since we don't want unit testing code (each with its own `main()`) to end up in the utility library.

The exclusion pattern, `-**.test...`, looks a bit cryptic. What we have here is a second-level extension (`.test`) which we use to classify our source files as belonging to unit tests. Because it is a second-level extension, we have to indicate this fact to the pattern matching machinery with the trailing triple dot (meaning "there are more extensions coming"). If we didn't do that, `.test` would have been treated as a first-level extension explicitly specified for our source files (see Target Types for details).

The next couple of lines set target type/pattern-specific variables to treat all unit test executables as tests that should not be installed:

```
exe{*.test}:
{
    test = true
    install = false
}
```

You may be wondering why we had to escape the second-level `.test` extension in the name pattern above but not here. The answer is that these are different kinds of patterns in different contexts. In particular, patterns in the target type/pattern-specific variables are only matched against target names without regard for extensions. See Name Patterns for details.

Then we have the `for`-loop that declares an executable target for each unit test source file. The list of these files is generated with a name pattern that is the inverse of what we've used for the utility library:

```

for t: cxx{**.test...}
{
  d = $directory($t)
  n = $name($t)...

  ./: $d/exe{$n}: $t $d/hxx{+$n} $d/testscript{+$n}
  $d/exe{$n}: libue{hello}: bin.whole = false
}

```

In the loop body we first split the test source file into the directory (remember, we can have sources, including tests, in subdirectories) and name (which contains the `.test` second-level extension and which we immediately escape with `. . .`). And then we use these components to declare a dependency for the corresponding unit test executable. There is nothing here that we haven't already seen except for using variable expansions instead of literal names.

By default utility libraries are linked in the "whole archive" mode where every object file from the static library ends up in the resulting executable or library. This behavior is what we want when linking the primary target but can normally be relaxed for unit tests to speed up linking. This is what the last line in the loop does using the `bin.whole` prerequisite-specific variable.

You can easily customize this and other aspects on a test-by-test basis by excluding the specific test(s) from the loop and then providing a custom implementation. For example:

```

for t: cxx{**.test... -special.test...}
{
  ...
}

./: exe{special.test...}: cxx{special.test...} libue{hello}

```

Note also that if you plan to link any of your unit tests in the whole archive mode, then you will also need to exclude the source file containing the primary executable's `main()` from the utility library. For example:

```

./: exe{hello}: cxx{main} libue{hello}
libue{hello}: {hxx cxx}{** -main -**.test...}

```

1.10 Diagnostics and Debugging

Sooner or later we will run into a situation where our `buildfiles` don't do what we expect them to. In this section we examine a number of techniques and mechanisms that can help us understand the cause of a misbehaving build.

To perform a build the build system goes through several phases. During the *load* phase the `buildfiles` are loaded and processed. The result of this phase is the in-memory *build state* that contains the scopes, targets, variables, etc., defined by the `buildfiles`. Next is the *match* phase during which rules are matched to the targets that need to be built, recursively. Finally, during the *execute* phase the matched rules are executed to perform the build.

The load phase is always serial and stops at the first error. In contrast, by default, both match and execute are parallel and continue in the presence of errors (similar to the "keep going" make mode). While beneficial in normal circumstances, during debugging this can lead to both interleaved output that is hard to correlate as well as extra noise from cascading errors. As a result, for debugging, it is usually helpful to run serially and stop at the first error, which can be achieved with the `--serial-stop|-s` option.

The match phase can be temporarily switched to either (serial) load or (parallel) execute. The former is used, for example, to load additional `buildfiles` during the `dir{}` prerequisite to target resolution, as described in Output Directories and Scopes. While the latter is used to update generated source code (such as headers) that is required to complete the match.

Debugging issues in each phase requires different techniques. Let's start with the load phase. As mentioned in Buildfile Language, `buildfiles` are processed linearly with directives executed and variables expanded as they are encountered. As we have already seen, to print a variable value we can use the `info` directive. For example:

```
x = X
info $x
```

This will print something along these lines:

```
buildfile:2:1: info: X
```

Or, if we want to clearly see where the value begins and ends (useful when investigating whitespace-related issues):

```
x = " X "
info "'$x'"
```

Which prints:

```
buildfile:2:1: info: ' X '
```

Besides the `info` directive, there are also `text`, which doesn't print the `info:` prefix, `warn`, which prints a warning, as well as `fail` which prints an error and causes the build system to exit with an error. Here is an example of using each:

```
text 'note: we are about to get an error'
warn 'the error is imminent'
fail 'this is the end'
info 'we will never get here'
```

This will produce the following output:

```
buildfile:1:1: note: we are about to get an error
buildfile:2:1: warning: the error is imminent
buildfile:3:1: error: this is the end
```

If you find yourself writing code like this:

```
if ($cxx.target.class == 'windows')
    fail 'Windows is not supported'
```

Then the `assert` directive is a more concise way to express the same:

```
assert ($cxx.target.class != 'windows') 'Windows is not supported'
```

The `assert` condition must be an expression that evaluates to `true` or `false`, similar to the `if` directive (see [Conditions \(if-else\)](#) for details). The description after the condition is optional and, similar to `if`, there is also the `assert!` variant, which fails if the condition is `true`.

All the diagnostics directives write to `stderr`. If instead we need to write something to `stdout` to, for example, send some information back to our caller, then we can use the `print` directive. For example, this will print the C++ compiler id and its target:

```
print "$cxx.id $cxx.target"
```

To query the value of a target-specific variable we use the qualified name syntax (the `eval-qual` production) of `eval` context, for example:

```
obj{main}: cxx.poptions += -DMAIN
info $(obj{main}: cxx.poptions)
```

There is no direct way to query the value of a prerequisite-specific variable since a prerequisite has no identity. Instead, we can use the `dump` directive discussed next to print the entire dependency declaration, including prerequisite-specific variables for each prerequisite.

While printing variable values is the most common mechanism for diagnosing `buildfile` issues, sometimes it is also helpful to examine targets and scopes. For that we use the `dump` directive.

Without any arguments, `dump` prints (to `stderr`) the contents of the scope it was encountered in and at that point of processing the `buildfile`. Its output includes variables, targets and their prerequisites, as well as nested scopes, recursively. As an example, let's print the source subdirectory scope of our `hello` executable project. Here is its `buildfile` with the `dump` directive at the end:

```
exe{hello}: {hxx cxx}{**}

cxx.poptions += "-I$out_root" "-I$src_root"

dump
```

This will produce the output along these lines:

```

buildfile:5:1: dump:
/tmp/hello/hello/
{
  [strings] cxx.poptions = -I/tmp/hello -I/tmp/hello
  [dir_path] out_base = /tmp/hello/hello/
  [dir_path] src_base = /tmp/hello/hello/

  buildfile{buildfile.}:

  exe{hello.?): cxx{hello.?)
}

```

The question marks (?) in the dependency declaration mean that the file extensions haven't been assigned yet, which happens during the match phase.

Instead of printing the entire scope, we can also print individual targets by specifying one or more target names in `dump`. To make things more interesting, let's convert our `hello` project to use a utility library, similar to the unit testing setup (Implementing Unit Testing). We will also link to the `dl` library to see an example of a target-specific variable being dumped:

```

exe{hello}: libue{hello}: bin.whole = false
exe{hello}: cxx.libs += -ldl
libue{hello}: {hxx cxx}{**}

dump exe{hello}

```

The output will look along these lines:

```

buildfile:5:1: dump:
/tmp/hello/hello/exe{hello.?):
{
  [strings] cxx.libs = -ldl
}
/tmp/hello/hello/exe{hello.?): /tmp/hello/hello/:libue{hello.?):
{
  [bool] bin.whole = false
}

```

The output of `dump` might look familiar: in [Output Directories and Scopes](#) we've used the `--dump` option to print the entire build state, which looks pretty similar. In fact, the `dump` directive uses the same mechanism but allows us to print individual scopes and targets from within a `buildfile`.

There is, however, an important difference to keep in mind: `dump` prints the state of a target or scope at the point in the `buildfile` load phase where it was executed. In contrast, the `--dump` option can be used to print the state after the load phase (`--dump load`) and/or after the match phase (`--dump match`). In particular, the after match printout reflects the changes to the build state made by the matching rules, which may include entering of additional dependencies, setting of additional variables, resolution of prerequisites to targets, assignment of file extensions, etc. As a result, while the `dump` directive should be sufficient in most cases, sometimes you may need to use the `--dump` option to examine the build state just before rule execution.

It is possible to limit the output of `--dump` to specific scopes and/or targets with the `--dump-scope` and `--dump-target` options.

Let's now move from state to behavior. As we already know, to see the underlying commands executed by the build system we use the `-v` options (which is equivalent to `--verbose 2`). Note, however, that these are *logical* rather than actual commands. You can still run them and they should produce the desired result, but in reality the build system may have achieved the same result in a different way. To see the actual commands we use the `-V` option instead (equivalent to `--verbose 3`). Let's see the difference in an example. Here is what building our `hello` executable with `-v` might look like:

```
$ b -s -v
g++ -o hello.o -c hello.cxx
g++ -o hello hello.o
```

And here is the same build with `-V`:

```
$ b -s -V
g++ -MD -E -fdirectives-only -MF hello.o.t -o hello.o.ii hello.cxx
g++ -E -fpreprocessed -fdirectives-only hello.o.ii
g++ -o hello.o -c -fdirectives-only hello.o.ii
g++ -o hello hello.o
```

From the second listing we can see that in reality `build2` first partially preprocessed `hello.cxx` while extracting its header dependency information. It then preprocessed it fully, which is used to extract module dependency information, calculate the checksum for ignorable change detection, etc. When it comes to producing `hello.o`, the build system compiled the partially preprocessed output rather than the original `hello.cxx`. The end result, however, is the same as in the first listing.

Verbosity level 3 (`-V`) also triggers printing of the build system module configuration information. Here is what we would see for the `cxx` module:

```
cxx hello@/tmp/hello/
  cxx      g++@/usr/bin/g++
  id       gcc
  version  7.2.0 (Ubuntu 7.2.0-1ubuntu1~16.04)
  major    7
  minor    2
  patch    0
  build     (Ubuntu 7.2.0-1ubuntu1~16.04)
  signature gcc version 7.2.0 (Ubuntu 7.2.0-1ubuntu1~16.04)
  checksum 09b3b59d337eb9a760dd028fa0df585b307e6a49c2bfa00b3[...]
  target   x86_64-linux-gnu
  runtime  libgcc
  stdlib   libstdc++
  c stdlib  glibc
...
```

Verbosity levels higher than 3 enable build system tracing. In particular, level 4 is useful for understanding why a rule doesn't match a target or if it does, why it determined the target to be out of date. For example, assuming we have an up-to-date build of our `hello`, let's change a compile option:

```
$ b -s --verbose 4
info: /tmp/hello/dir{hello/} is up to date

$ b -s --verbose 4 config.cxx.poptions+==DNDEBUG
trace: cxx::compile_rule::apply: options mismatch forcing update
of /tmp/hello/hello/obje{hello.o}
...
```

Higher verbosity levels result in more and more tracing statements being printed. These include `buildfile` loading and parsing, prerequisite to target resolution, as well as build system module and rule-specific logic.

While the tracing statements can be helpful in understanding what is happening, they don't make it easy to see why things are happening a certain way. In particular, one question that is often encountered during build troubleshooting is which dependency chain causes matching or execution of a particular target. These questions can be answered with the help of the `--trace-match` and `--trace-execute` options. For example, if we want to understand what causes the update of `obje{hello}` in the `hello` project above:

```
$ b -s --trace-execute 'obje{hello}'
info: updating hello/obje{hello}
  info: using rule cxx.compile
  info: while updating hello/libue{hello}
  info: while updating hello/exe{hello}
  info: while updating dir{hello/}
  info: while updating dir{./}
```

Another useful diagnostics option is `--mtime-check`. When specified, the build system performs a number of file modification time sanity checks that can be helpful in diagnosing spurious rebuilds.

If neither state dumps nor behavior analysis are sufficient to understand the problem, there is always an option of running the build system under a C++ debugger in order to better understand what's going on. This can be particularly productive for debugging complex rules.

Finally, to help with diagnosing the build system performance issues, there is the `--stat` option. It causes `build2` to print various execution statistics which can be useful for pin-pointing the bottlenecks. There are also a number of options for tuning the build system's performance, such as, the number of jobs to perform in parallel, the stack size, queue depths, etc. See the **b(1)** man pages for details.

2 Project Configuration

As discussed in the introduction (specifically, Project Structure) support for build configurations is an integral part of `build2` with the same mechanism used by the build system core (for example, for project importation via the `config.import.*` variables), by the build system modules (for example, for supplying compile options such as `config.cxx.poptions`), as well as by our projects to provide any project-specific configurability. Project configuration is the topic of this chapter.

The `build2` build system currently provides no support for `autoconf`-style probing of the build environment in order to automatically discover available libraries, functions, features, etc.

The main reason for omitting this support is the fundamental ambiguity and the resulting brittleness of such probing due to the reliance on compiler, linker, or test execution failures. Specifically, in many such tests it is impossible for a build system to distinguish between a missing feature, a broken test, and a misconfigured build environment. This leads to requiring a user intervention in the best case and to a silently misconfigured build in the worst. Other issues with this approach include portability, speed (compiling and linking takes time), as well as limited applicability during cross-compilation (specifically, inability to run tests).

As a result, we recommend using *expectation-based* configuration where your project assumes a feature to be available if certain conditions are met. Examples of such conditions at the source code level include feature test macros, platform macros, runtime library macros, compiler macros, etc., with the build system modules exposing some of the same information via variables to allow making similar decisions in `buildfiles`. The standard pre-installed `autoconf` build system module provides emulation of GNU `autoconf` using this approach.

Another alternative is to automatically adapt to missing features using more advanced techniques such as C++ SFINAE. And in situations where none of this is possible, we recommend delegating the decision to the user via a configuration value. Our experience with `build2` as well as those of other large cross-platform projects such as Boost show that this is a viable strategy.

Having said that, `build2` does provide the ability to extract configuration information from the environment (`$getenv()` function) or other tools (`$process.run*()` family of functions). Note, however, that for this to work reliably there should be no ambiguity between the "no configuration available" case (if such a case is possible) and the "something went wrong" case. We show a realistic example of this in Configuration Report where we extract the GCC plugin directory while dealing with the possibility of it being configured without plugin support.

Before we delve into the technical details, let's discuss the overall need for project configurability. While it may seem that making one's project more user-configurable is always a good idea, there are costs: by having a choice we increase the complexity and open the door for potential incompatibility. Specifically, we may end up with two projects in the same build needing a shared dependency with incompatible configurations.

While some languages, such as Rust, support having multiple differently-configured projects in the same build, this is not something that is done often in C/C++. This ability is also not without its drawbacks, most notably code bloat.

As a result, our recommendation is to strive for simplicity and avoid user configurability whenever possible. For example, there is a common desire to make certain functionality optional in order not to make the user pay for things they don't need. This, however, is often

better addressed either by always providing the optional functionality if it's fairly small or by factoring it into a separate project if it's substantial. If a configuration value is to be provided, it should have a sensible default with a bias for simplicity and compatibility rather than the optimal result. For example, in the optional functionality case, the default should probably be to provide it.

As discussed in the introduction, the central part of the build configuration functionality are the *configuration variables*. One of the key features that make them special is support for automatic persistence in the `build/config.build` file provided by the `config` module (see [Configuring](#) for details).

Another mechanism that can be used for project configuration is environment variables. While not recommended, sometimes it may be forced on us by external factors. In such cases, environment variables that affect the build result should be reported with the `config.environment` directive as discussed in [Hermetic Build Configurations](#).

The following example, based on the `libhello` project from the introduction, gives an overview of the project configuration functionality with the remainder of the chapter providing the detailed explanation of all the parts shown as well as the alternative approaches.

```
libhello/
|-- build/
|   |-- root.build
|   |-- ...
|-- libhello/
|   |-- hello.cxx
|   |-- buildfile
|   |-- ...
|-- ...

# build/root.build

config [string] config.libhello.greeting ?= 'Hello'

# libhello/buildfile

cxx.poptions += "-DLIBHELLO_GREETING=\"${config.libhello.greeting}\""

// libhello/hello.cxx

void say_hello (ostream& o, const string& n)
{
    o << LIBHELLO_GREETING " ", " << n << '!' << endl;
}

$ b configure config.libhello.greeting=Hi -v
config libhello@/tmp/libhello/
    greeting    Hi

$ cat build/config.build
config.libhello.greeting = Hi

$ b -v
g++ ... -DLIBHELLO_GREETING="Hi" ...
```

By (enforced) convention, configuration variables start with `config.`, for example, `config.import.libhello`. In case of a build system module, the second component in its configuration variables should be the module name, for example, `config.cxx`, `config.cxx.options`. Similarly, project-specific configuration variables should have the project name as their second component, for example, `config.libhello.greeting`.

More precisely, a project configuration variable must match the `config[.*].<project>.*` pattern where additional components may be present after `config.` in case of subprojects. Overall, the recommendation is to use hierarchical names, such as `config.libcurl.tests.remote` for subprojects, similar to build system submodules.

If a build system module for a tool (such as a source code generator) and the tool itself share a name, then they may need to coordinate their configuration variable names in order to avoid clashes. Note also that when importing an executable target in the `<project>%exe{<project>}` form, the `config.<project>` variable is treated as an alias for `config.import.<project>.<project>.exe`.

For an imported buildfile, `<project>` may refer to either the importing project or the project from which the said buildfile was imported.

The build system core reserves `build` as any intermediate component and `import` as the second component in configuration variables.

A variable in the `config.<project>.develop` form has pre-defined semantics: it allows a project to distinguish between *development* and *consumption* builds. While normally there is no distinction between these two modes, sometimes a project may need to provide additional functionality during development. For example, a source code generator which uses its own generated code in its implementation may need to provide a bootstrap step from the pre-generated code. Normally, such a step is only needed during development.

While some communities, such as Rust, believe that building and running tests is only done during development, we believe it's reasonable for an end-user to want to run tests for all their dependencies. As a result, we strongly discourage restricting tests to the development mode only. Tests are an integral part of the project and should always be available.

If used, the `config.<project>.develop` variable should be explicitly defined by the project with the `bool` type and the `false` default value. For example:

```
# build/root.build

config [bool] config.libhello.develop ?= false
```

If the `config.<project>.develop` variable is specified by the user of the project but the project does not define it (that is, the project does not distinguish between development and consumption), then this variable is silently ignored. By default **bdep-init (1)** configures projects being initialized for development. This can be overridden with explicit

```
config.<project>.develop=false.
```

2.1 config Directive

To define a project configuration variable we add the `config` directive into the project's `build/root.build` file (see Project Structure). For example:

```
config [bool]    config.libhello.fancy    ?= false
config [string]  config.libhello.greeting ?= 'Hello'
```

The irony does not escape us: these configuration variables are exactly of the kind that we advocate against. However, finding a reasonable example of build-time configurability in a *"Hello, World!"* library is not easy. In fact, it probably shouldn't have any. So, for this chapter, do as we say, not as we do.

Similar to `import` (see Target Importation), the `config` directive is a special kind of variable assignment. Let's examine all its parts in turn.

First comes the optional list of variable attributes inside `[ ]`. The only attribute that we have in the above example is the variable type, `bool` and `string`, respectively. It is generally a good idea to assign static types to configuration variables because their values will be specified by the users of our project and the more automatic validation we provide the better (see Variables for the list of available types). For example, this is what will happen if we misspell the value of the `fancy` variable:

```
$ b configure config.libhello.fancy=fals
error: invalid bool value 'fals' in variable config.libhello.fancy
```

After the attribute list we have the variable name. The `config` directive will validate that it matches the `config[.*].<project>.*` pattern (with one exception discussed in Configuration Report).

Finally, after the variable name comes the optional default value. Note that unlike normal variables, the default value assignment (`?=`) is the only valid form of assignment in the `config` directive.

The semantics of the `config` directive is as follows: First an overridable variable is entered with the specified name, type (if any), and global visibility. Then, if the variable is undefined and the default value is specified, it is assigned the default value. After this, if the variable is defined (either as user-defined or default), it is marked for persistence. Finally, a defined variable is also marked for reporting as discussed in Configuration Report. Note that if the variable is user-defined, then the default value is not evaluated.

Note also that if the configuration value is not specified by the user and you haven't provided the default, the variable will be undefined, not `null`, and, as a result, omitted from the persistent configuration (`build/config.build` file). In fact, unlike other variables, project configuration variables are by default not *nullable*. For example:

```
$ b configure config.libhello.fancy=[null]
error: null value in non-nullable variable config.libhello.fancy
```

There are two ways to make `null` a valid value of a project configuration variable. Firstly, if the default value is `null`, then naturally the variable is assumed nullable. This is traditionally used for *optional* configuration values. For example:

```
config [string] config.libhello.fallback_name ?= [null]
```

If we need a nullable configuration variable but with a non-`null` default value (or no default value at all), then we have to use the `null` variable attribute. For example:

```
config [string, null] config.libhello.fallback_name ?= "World"
```

A common approach for representing an C/C++ enum-like value is to use `string` as a type and pattern matching for validation. In fact, validation and propagation can often be combined. For example, if our library needed to use a database for some reason, we could handle it like this:

```
config [string] config.libhello.database ?= [null]

using cxx

switch $config.libhello.database
{
  case [null]
  {
    # No database in use.
  }
  case 'sqlite'
  {
    cxx.poptions += -DLIBHELLO_WITH_SQLITE
  }
  case 'pgsql'
  {
    cxx.poptions += -DLIBHELLO_WITH_PGSQL
  }
  default
  {
    fail "invalid config.libhello.database value \
'$config.libhello.database' "
  }
}
```

While it is generally a good idea to provide a sensible default for all your configuration variables, if you need to force the user to specify its value explicitly, this can be achieved with an extra check. For example:

```
config [string] config.libhello.database

if! $defined(config.libhello.database)
  fail 'config.libhello.database must be specified'
```

A configuration variable without a default value is omitted from `config.build` unless the value is specified by the user. This semantics is useful for values that are normally derived from other configuration values but could also be specified by the user. If the value is derived, then we don't want it saved in `config.build` since that would prevent it from being re-derived if the configuration values it is based on are changed. For example:

```
config [strings] config.hello.database

assert ($size($config.hello.database) > 0) \
    'database must be specified with config.hello.database'

config [bool, config.report.variable=multi] config.hello.multi_database

multi = ($defined(config.hello.multi_database) \
    ? $config.hello.multi_database \
    : $size(config.hello.database) > 1)

assert ($multi || $size(config.hello.database) == 1) \
    'one database can be specified if config.hello.multi_database=false'
```

If computing the default value is expensive or requires elaborate logic, then the handling of a configuration variable can be broken down into two steps along these lines:

```
config [string] config.libhello.greeting

if! $defined(config.libhello.greeting)
{
    greeting = ... # Calculate default value.

    if ($greeting == [null])
        fail "unable to calculate default greeting, specify manually \
with config.libhello.greeting"

    config config.libhello.greeting ?= $greeting
}
```

Other than assigning the default value via the `config` directive, configuration variables should not be modified by the project's `buildfiles`. Instead, if further processing of the configuration value is necessary, we should assign the configuration value to a different, non-`config.*`, variable and modify that. The two situations where this is commonly required are post-processing of configuration values to be more suitable for use in `buildfiles` as well as further customization of configuration values. Let's see examples of both.

To illustrate the first situation, let's say we need to translate the database identifiers specified by the user:

```
config [string] config.libhello.database ?= [null]

switch $config.libhello.database
{
    case [null]
        database = [null]

    case 'sqlite'
        database = 'SQLITE'
```

```

case 'pgsql'
    database = 'PGSQL'

case 'mysql'
case 'mariadb'
    database = 'MYSQL'

default
    fail "...
}
}

using cxx

if ($database != [null])
    cxx.poptions += "-DLIBHELLO_WITH_${database}"

```

For the second situation, the typical pattern looks like this:

```

config [strings] config.libhello.options

options = # Overridable options go here.
options += $config.libhello.options
options += # Non-overridable options go here.

```

That is, assuming that the subsequently specified options (for example, command line options) override any previously specified, we first set default `buildfile` options that are allowed to be overridden by options from the configuration value, then append such options, if any, and finish off by appending `buildfile` options that should always be in effect.

As a concrete example of this approach, let's say we want to make the compiler warning level of our project configurable (likely a bad idea; also ignores compiler differences):

```

config [strings] config.libhello.woptions

woptions = -Wall -Wextra
woptions += $config.libhello.woptions
woptions += -Werror

using cxx

cxx.coptions += $woptions

```

With this arrangement, the users of our project can customize the warning level but cannot disable the treatment of warnings as errors. For example:

```

$ b -v config.libhello.woptions=-Wno-extra
g++ ... -Wall -Wextra -Wno-extra -Werror ...

```

If you do not plan to package your project, then the above rules are the only constraints you have. However, if your project is also a package, then other projects that use it as a dependency may have preferences and requirements regarding its configuration. And it becomes the job of the package manager (`bpkg`) to negotiate a suitable configuration between all the dependents of your project (see [Dependency Configuration Negotiation](#) for details). This can be a difficult problem to solve optimally in a reasonable time and to help the package manager come up with the best configuration quickly you should follow the below additional rules and

recommendations for configuration of packages (but which are also generally good ideas):

1. Prefer `bool` configuration variables. For example, if your project supports a fixed number of backends, then provide a `bool` variable to enable each rather than a single variable that lists all the backends to be enabled.
2. Avoid project configuration variable dependencies, that is, where the default value of one variable depends on the value of another. But if you do need such a dependency, make sure it is expressed using the original `config.<project>.*` variables rather than any intermediate/computed values. For example:

```
# Enable Y only if X is enabled.
#
config [bool] config.hello.x ?= false
config [bool] config.hello.y ?= $config.libhello.x
```

3. Do not make project configuration variables conditional. In other words, the set of configuration variables and their types should be a static property of the project. If you do need to make a certain configuration variable "unavailable" or "disabled" if certain conditions are met (for example, on a certain platform or based on the value of another configuration variable), then express this with a default value and/or a check. For example:

```
windows = ($cxx.target.class == 'windows')

# Y should only be enabled if X is enabled and we are not on
# Windows.
#
config [bool] config.hello.x ?= false
config [bool] config.hello.y ?= ($config.hello.x && !$windows)

if $config.libhello.y
{
    assert $config.hello.x "Y can only be enabled if X is enabled"
    assert (!$windows)    "Y cannot be enabled on Windows"
}
```

Additionally, if you wish to factor some `config` directives into a separate file (for example, if you have a large number of them or you would like to share them with subprojects) and source it from your `build/root.build`, then it is recommended that you place this file into the `build/config/` subdirectory, where the package manager expects to find such files (see Package Build System Skeleton for background). For example:

```
# root.build
#
...

source $src_root/build/config/common.build
```

If you would prefer to keep such a file in a different location (for example, because it contains things other than `config` directives), then you will need to manually list it in your package's manifest file, see the `build-file` value for details.

Another effect of the `config` directive is to print the configuration variable in the project's configuration report. This functionality is discussed in the following section. While we have already seen some examples of how to propagate the configuration values to our source code, Configuration Propagation discusses this topic in more detail.

2.2 Configuration Report

One of the effects of the `config` directive is to mark a defined configuration variable for reporting. The project configuration report is printed automatically at a sufficiently high verbosity level along with the build system module configuration. For example (some of the `cxx` module configuration is omitted for brevity):

```
$ b config.libhello.greeting=Hey -v
cxx libhello@/tmp/libhello/
  cxx      g++@/usr/bin/g++
  id       gcc
  version  9.1.0
  ...
config libhello@/tmp/libhello/
  fancy    false
  greeting Hey
```

The configuration report is printed immediately after loading the project's `build/root.build` file. It is always printed at verbosity level 3 (`-V`) or higher. It is also printed at verbosity level 2 (`-v`) if any of the reported configuration variables have a *new* value. A value is considered new if it was set to default or was overridden on the command line.

The project configuration report header (the first line) starts with the special `config` module name (the `config` module itself does not have a report) followed by the project name and its `out_root` path. After the header come configuration variables with the `config[.*].<project>` prefix removed. The configuration report for each variable can be customized using a number of `config.report*` attributes as discussed next.

The `config.report` attribute controls whether the variable is included into the report and, if so, the format to print its value in. For example, this is how we can exclude a variable from the report:

```
config [bool, config.report=false] config.libhello.selftest ?= false
```

While we would normally want to report all our configuration variables, if some of them are internal and not meant to be used by the users of our project, it probably makes sense to exclude them.

The only currently supported alternative printing format is `multiline` which prints a list value one element per line. Other printing formats may be supported in the future. For example:

```

config [dir_paths, config.report=multiline] config.libhello.search_dirs

$ b config.libhello.search_dirs="/etc/default /etc" -v
config libhello@/tmp/libhello/
  search_dirs
    /etc/default/
    /etc/

```

The `config.report` attribute can also be used to include a `non-config.*` variable into a report. This is primarily useful for configuration values that are always discovered automatically but that are still useful to report for troubleshooting. Here is a realistic example:

```

using cxx

# Determine the GCC plugin directory.
#
if ($cxx.id == 'gcc')
{
  plugin_dir = [dir_path] $process.run($cxx.path -print-file-name=plugin)

  # If plugin support is disabled, then -print-file-name will print
  # the name we have passed (the real plugin directory will always
  # be absolute).
  #
  if ("plugin_dir" == plugin)
    fail "$recall($cxx.path) does not support plugins"

  config [config.report] plugin_dir
}

```

This is the only situation where a variable that does not match the `config[.*].<project>.*` pattern is allowed in the `config` directive. Note also that a value of such a variable is never considered new.

Note that this mechanism should not be used to report configuration values that require post-processing because of the loss of the new value status (unless you are reporting both the original and post-processed values). Instead, use the `config.report.variable` attribute to specify an alternative variable for the report. For example:

```

config [strings, config.report.variable=woptions] \
  config.libhello.woptions

woptions = -Wall -Wextra
woptions += $config.libhello.woptions
woptions += -Werror

$ b config.libhello.woptions=-Wno-extra -v
config libhello@/tmp/libhello/
  woptions -Wall -Wextra -Wno-extra -Werror

```

The `config.report.module` attribute can be used to override the reporting module name, that is, `config` in the `config libhello@/tmp/libhello/` line above. It is primarily useful in imported buildfiles that wish to report `non-config.*` variables under their own name. For example:

```

config [string] config.rtos.board

# Load the board description and report key information such as the
# capability revoker.
#
...
revoker = ...

config [config.report.module=rtos] revoker

$ b config.rtos.board=ibex-safe-simulator -v
rtos hello@/tmp/hello/
  board      ibex-safe-simulator
  revoker    hardware

```

2.3 Configuration Propagation

Using configuration values in our buildfiles is straightforward: they are like any other buildfile variables and we can access them directly. For example, this is how we could provide optional functionality in our library by conditionally including certain source files: See Conditions (`if-else`) for why we should not use `if` to implement this.

```

# build/root.build

config [strings] config.libhello.io ?= true

# libhello/buildfile

lib{hello}: {hxx ixx txx cxx}{** -version -hello-io} hxx{version}
lib{hello}: {hxx cxx}{hello-io}: include = $config.libhello.io

```

On the other hand, it is often required to propagate the configuration information to our source code. In fact, we have already seen one way to do it: we can pass this information via C/C++ preprocessor macros defined on the compiler's command line. For example:

```

# build/root.build

config [bool]    config.libhello.fancy    ?= false
config [string]  config.libhello.greeting ?= 'Hello'

# libhello/buildfile

if $config.libhello.fancy
  cxx.poptions += -DLIBHELLO_FANCY

cxx.poptions += "-DLIBHELLO_GREETING=\"${config.libhello.greeting}\""

// libhello/hello.cxx

void say_hello (ostream& o, const string& n)
{
#ifdef LIBHELLO_FANCY
  // TODO: something fancy.
#else
  o << LIBHELLO_GREETING " , " << n << '!' << endl;
#endif
}

```

We can even use the same approach to export certain configuration information to our library's users (see Library Exportation and Versioning for details):

```
# libhello/buildfile

# Export options.
#
if $config.libhello.fancy
    lib{hello}: cxx.export.poptions += -DLIBHELLO_FANCY
```

This mechanism is simple and works well across compilers so there is no reason not to use it when the number of configuration values passed and their size are small. However, it can quickly get unwieldy as these numbers grow. For such cases, it may make sense to save this information into a separate auto-generated source file with the help of the `in` module, similar to how we do it for the version header.

The often-used approach is to generate a header file and include it into source files that need access to the configuration information. Historically, this was a C header full of macros called `config.h`. However, for C++ projects, there is no reason not to make it a C++ header and, if desired, to use modern C++ features instead of macros. Which is what we will do here.

As an example of this approach, let's convert the above command line-based implementation to use the configuration header. We will continue using macros as a start (or in case this is a C project) and try more modern techniques later. The `build/root.build` file is unchanged except for loading the `in` module:

```
# build/root.build

config [bool]    config.libhello.fancy    ?= false
config [string]  config.libhello.greeting ?= 'Hello'

using in
```

The `libhello/config.hxx.in` file is new:

```
// libhello/config.hxx.in

#pragma once

#define LIBHELLO_FANCY    $config.libhello.fancy$
#define LIBHELLO_GREETING "$config.libhello.greeting$"
```

As you can see, we can reference our configuration variables directly in the `config.hxx.in` substitutions (see the `in` module documentation for details on how this works).

With this setup, the way to export configuration information to our library's users is to make the configuration header public and install it, similar to how we do it for the version header.

The rest is changed as follows:

```
# libhello/buildfile

lib{hello}: {hxx ixx txx cxx}{** -version -config} hxx{version config}

hxx{config}: in{config}
{
    install = false
}

// libhello/hello.cxx

#include <libhello/config.hxx>

void say_hello (ostream& o, const string& n)
{
    #if LIBHELLO_FANCY
        // TODO: something fancy.
    #else
        o << LIBHELLO_GREETING ", " << n << '!' << endl;
    #endif
}
```

Notice that we had to replace `#ifdef LIBHELLO_FANCY` with `#if LIBHELLO_FANCY`. If you want to continue using `#ifdef`, then you will need to make the necessary arrangements yourself (the `in` module is a generic preprocessor and does not provide any special treatment for `#define`). For example:

```
#define LIBHELLO_FANCY $config.libhello.fancy$
#if !LIBHELLO_FANCY
#    undef LIBHELLO_FANCY
#endif
```

Now that the macro-based version is working, let's see how we can take advantage of modern C++ features to hopefully improve on some of their drawbacks. As a first step, we can replace the `LIBHELLO_FANCY` macro with a compile-time constant and use `if constexpr` instead of `#ifdef` in our implementation:

```
// libhello/config.hxx.in

namespace hello
{
    inline constexpr bool fancy = $config.libhello.fancy$;
}

// libhello/hello.cxx

#include <libhello/config.hxx>

void say_hello (ostream& o, const string& n)
{
    if constexpr (fancy)
    {
        // TODO: something fancy.
    }
    else
        o << LIBHELLO_GREETING ", " << n << '!' << endl;
}
```

Note that with `if constexpr` the branch not taken must still be valid, parsable code. This is both one of the main benefits of using it instead of `#if` (the code we are not using is still guaranteed to be syntactically correct) as well as its main drawback (it cannot be used, for example, for platform-specific code without extra efforts, such as providing shims for missing declarations, etc).

Next, we can do the same for `LIBHELLO_GREETING`:

```
// libhello/config.hxx.in

namespace hello
{
    inline constexpr char greeting[] = "$config.libhello.greeting$";
}

// libhello/hello.cxx

#include <libhello/config.hxx>

void say_hello (ostream& o, const string& n)
{
    if constexpr (fancy)
    {
        // TODO: something fancy.
    }
    else
        o << greeting << ", " << n << '!' << endl;
}
```

Note that for `greeting` we can achieve the same result without using inline variables or `constexpr` and which would be usable in older C++ and even C. All we have to do is add the `config.cxx.in` source file next to our header with the definition of the `greeting` variable. For example:

```
// libhello/config.hxx.in

namespace hello
{
    extern const char greeting[];
}

// libhello/config.cxx.in

#include <libhello/config.hxx>

namespace hello
{
    const char greeting[] = "$config.libhello.greeting$";
}
```

```
# libhello/buildfile

lib{hello}: {hxx ixx txx cxx}{** -config} {hxx cxx}{config}

hxx{config}: in{config}
{
    install = false
}

cxx{config}: in{config}
```

As this illustrates, the `in` module can produce as many auto-generated source files as we need. For example, we could use this to split the configuration header into two, one public and installed while the other private.

3 Targets and Target Types

This chapter is a work in progress and is incomplete.

3.1 Target Types


A target type is part of a target's identity. The core idea behind the concept of target types is to abstract away from file extensions which can vary from project to project (for example, C++ source files extensions) or from platform to platform (for example, executable file extensions). It also allows us to have non-file-based targets.

Target types form a *base-derived* inheritance tree. The root of this tree is the abstract `target{}` type. The `build2` core defines a number of standard target types, such as `file{}`, `doc{}`, and `exe{}`. Build system modules can define additional target types that are based on the standard ones (or on types defined by other modules). For example, the `c` module that provides the C compilation support defines the `h{}` and `c{}` target types. Finally, `buildfiles` can derive project-local target types using the `define` directive.

If a target type represents a file type with a well-established extension, then by convention such an extension is used as the target type name. For example, the C language header and source files use the `.h` and `.c` extensions and the target types are called `h{}` and `c{}`.

Speaking of conventions, as you may have noticed, when mentioning a target type we customarily add `{}` after its name. We found that this helps with comprehension since target type names are often short (you can also search for `<type>{` to narrow it down to target types). In a way this is a similar approach to adding `()` after a function name except here we use `{}`, which mimics target type usage in target names, for example `c{hello}` for `hello.c`.

The following listing shows the hierarchy of the standard target types defined by the `build2` core (the abstract target types are marked with `*`) while the following sections describe each standard target type in detail. For target types defined by a module refer to the respective module documentation.



While target types replace (potentially variable) extensions, there still needs to be a mechanism for specifying them since in most cases targets have to be mapped to files. There are several ways this can be achieved.

If a target type represents a file type with a well-established extension, then such an extension is normally used by default and we don't need to take any extra steps. For example the `h{ }` and `c{ }` target types for C header and source files default to the `.h` and `.c` extensions, respectively, and if our project follows this convention, then we can simply write:

```
exe{utility}: c{utility} h{utility}
```

And `c{utility}` will be mapped to `utility.c` and `h{utility}` – to `utility.h`.

There are two variants of this default extension case: fixed extension and customizable extension. A target type may choose to fix the default extension if it's a bad idea to deviate from the default extension. A good example of such a target is `man1{ }`, which fixes the default extension to be `.1`. More commonly, however, a target will have a default extension but will allow customizing it with the `extension` variable.

A good example where extension customization is often required are the `hxx{ }` and `cxx{ }` target types for C++ header and source files, which default to the `.hxx` and `.cxx` extensions, respectively. If our project uses other extensions, for example, `.hpp` and `.cpp`, then we can adjust the defaults (typically done in `root.build`, after loading the `cxx` module):

```
hxx{*}: extension = hpp
cxx{*}: extension = cpp
```

Then we can write:

```
exe{utility}: cxx{utility} hxx{utility}
```

And `cxx{utility}` will be mapped to `utility.cpp` and `hxx{utility}` – to `utility.hpp`.

What about `exe{utility}`, where does its extension come from? This is an example of a target type with an extension that varies from platform to platform. In such cases the extension is expected to be assigned by the rule that matches the target. In the above example, the link rule from the `cxx` module that matches updating `exe{utility}` will assign a suitable

extension based on the target platform of the C++ compiler that it was instructed to use.

Finally, it is always possible to specify the file extension explicitly as part of the target name. For example:

```
exe{utility}: cxx{utility.cc} hxx{utility.hh}
```

This is normally only needed if the default extension is not appropriate or if the target type does not have a default extension, as is the case, for example, for the `file{}` and `doc{}` target types. This mechanism can also be used to override the automatically derived extension. For example:

```
exe{($cxx.target.class == 'windows' ? utility.com : utility)}: ...
```

If you need to specify a name that does not have an extension, then end it with a single dot. For example, for a header `utility` you would write `hxx{utility.}`. If you need to specify a name with an actual trailing dot, then escape it with a double dot, for example, `hxx{utility..}`.

More generally, anywhere in a name, a double dot can be used to specify a dot that should not be considered the extension separator while a triple dot – which should. For example, in `obja{foo.a.o}` the extension is `.o` and if instead we wanted `.a.o` to be considered the extension, then we could rewrite it either as `obja{foo.a..o}` or as `obja{foo...a.o}`. The trailing triple dots signify the "unspecified (default) extension", for example, `cxx{foo.test...}` for `foo.test.cxx`. Note that an even number of dots is always treated as a sequence of double dots and the odd number of dots other than one or three is illegal.

To derive a new target type in a `buildfile` we use the `define` directive. Such target types are project-local, meaning they cannot be exported to other projects. Typically this is used to provide a more meaningful name to a set of files and also avoid having to specify their extensions explicitly. Compare:

```
./: doc{README.md PACKAGE-README.md INSTALL.md}
```

To:

```
define md: doc
doc{*}: extension = md

./: md{README PACKAGE-README INSTALL}
```

3.1.1 target{ }

The `target{ }` target type is a root of the target type hierarchy. It is abstract and is not commonly used directly, except perhaps in patterns (target type/pattern-specific variable, pattern rules).

3.1.2 alias{} and dir{}

The `alias{}` target type is used for non-file-based targets that serve as aliases for their prerequisite.

Alias targets in `build2` are roughly equivalent to phony targets in `make`.

For example:

```
alias{tests}: exe{test1 test2 test3}
```

```
$ b test: alias{tests}
```

An `alias{}` target can also serve as an "action" if supplied with an ad hoc recipe (or matched by an ad hoc pattern rule). For example:

```
alias{strip}: exe{hello}
{{
    diag strip $<
    strip $path($<)
}}
```

The `dir{}` target type is a special kind of alias that represents a directory. Building it means building everything inside the directory. See Project Structure for background.

A target without a type that ends with a directory separator (/) is automatically treated as `dir{}`. For example, the following two lines are equivalent:

```
./: exe{test1 test2}
dir{./}: exe{test1 test2}
```

Omitting the target type in such situations is customary.

3.1.3 fsdir{}

The `fsdir{}` target type represents a filesystem directory. Unlike `dir{}` above, it is not an alias and listing an `fsdir{}` directory as a prerequisite of a target will cause that directory to be created on `update` and removed on `clean`.

While we usually don't need to list explicit `fsdir{}` prerequisites for our targets, one situation where this is necessary is when the target resides in a subdirectory that does not correspond to an existing source directory. A typical example of this situation is placing object files into subdirectories. Compare:

```
obj{foo}: c{foo}
sub/obj{bar}: c{bar} fsdir{sub/}
```

3.1.4 mtime_target{} and path_target{}

The `mtime_target{}` target type represents a target that uses modification times to determine if it is out of date. The `path_target{}` target type represents a target that has a corresponding filesystem entry. It is derived from `mtime_target{}` and uses the modification time of that filesystem entry to determine if the target is out of date.

Both of these target types are abstract and are not commonly used directly, except perhaps in patterns (target type/pattern-specific variable, pattern rules).

3.1.5 group{ }

The `group{ }` target type represents a user-defined explicit target group, that is, a target that has multiple member targets that are all built together with a single recipe.

Normally this target type is not used to declare targets or prerequisites but rather as a base of a derived group. If desired, such a derived group can be marked with an attribute as "see-through", meaning that when the group is listed as a prerequisite of a target, the matching rule "sees" its members, rather than the group itself. For example:

```
define [see_through] thrift_cxx: group
```

3.1.6 file{ }

The `file{ }` target type represents a generic file. This target type is used as a base for most of the file-based targets and can also be used to declare targets and prerequisites when there are no more specific target types.

A target or prerequisite without a target type is automatically treated as `file{ }`. However, omitting a target type in such situations is not customary.

The `file{ }` target type has no default extension and one cannot be assigned with the `extension` variable. As a result, if a `file{ }` target has an extension, then it must be specified explicitly as part of the target name. For example:

```
./: file{example.conf}
```

3.1.7 doc{ }, legal{ }, and man{ }

The `doc{ }` target type represents a generic documentation file. It has semantics similar to `file{ }` (from which it derives): it can be used as a base or declare targets/prerequisites and there is no default extension. One notable difference, however, is that `doc{ }` targets are by default installed into the `doc/` installation location (see `install` Module). For example:

```
./: doc{README.md ChangeLog.txt}
```

The `legal{ }` target type is derived from `doc{ }` and represents a legal documentation file, such as a license, copyright notice, authorship information, etc. The main purpose of having a separate target type like this is to help with installing licensing-related files into a different location. To this effect, `legal{ }` targets are installed into the `legal/` installation location, which by default is the same as `doc/` but can be customized. For example:

```
./: legal{COPYRIGHT LICENSE AUTHORS.md}
```

The `man{ }` target type is derived from `doc{ }` and represents a manual page. This target type requires an explicit extension specification and is installed into the `man/` installation location

If you are using the `man{ }` target type directly (instead of one of `man<N>{ }` described below), for example, to install a localized version of a man page, then you will likely need to adjust the installation location on the per target basis.

The `man<N>{ }` target types (where `<N>` is an integer between 1 and 9) are derived from `man{ }` and represent manual pages in the respective sections. These target types have fixed default extensions `.<N>` (but an explicit extension can still be specified, for example `man1{foo.lp}`) and are installed into the `man<N>/` installation locations. For example:

```
./: man1{foo}
```

3.1.8 exe{ }

The `exe{ }` target type represents an executable file. Executables in `build2` appear in two distinct but sometimes overlapping contexts: We can build an executable target, for example from C source files. Or we can list an executable target as a prerequisite in order to execute it as part of a recipe. And sometimes this can be the same executable target. For example, one project may build an executable target that is a source code generator and another project may import this executable target and use it in its recipes in order to generate some source code.

To support this semantics the `exe{ }` target type has a peculiar default extension logic. Specifically, if the `exe{ }` target is "output", then the extension is expected to be assigned by the matching rule according to the target platform for which this executable is built. But if it does not, then we fall back to no extension (for example, a script). If, however, the `exe{ }` target is "input" (that is, it's listed as a prerequisite and there is no corresponding "output" target), then the extension of the host platform is used as the default.

In all these cases the extension can also be specified explicitly. This, for example, would be necessary if the executable were a batch file:

```
h{generate}: exe{generate.bat}
{{
    diag $< -> $>
    $< -o $path($>)
}}
```

Here, without the explicit extension, the `.exe` extension would have been used by default.

3.1.9 `json{ }`

The `json{ }` target type represents a JSON text file. It is derived from `file{ }` and has the `.json` default extension.

4 Variables

This chapter is a work in progress and is incomplete.

Variable names/components that start with underscore (`_`), variables in the `build`, `import`, and `export` namespaces, as well as the `build` intermediate component are reserved by the build system core. For example:

```
_x = 1          # error
x._y = 1        # error
build.x = 1     # error
x.build.y = 1  # error
```

The following variable/value types can currently be used in `buildfiles`:

```
bool

int64
int64s

uint64
uint64s

string
strings
string_set
string_map

path
paths
dir_path
dir_paths

json
json_array
json_object
json_set
json_map

name
names
name_pair

cmdline
project_name
target_triplet
```

Note that while expansions in the target and prerequisite-specific assignments happen in the corresponding target and prerequisite contexts, respectively, for type/pattern-specific assignments they happen in the scope context. Plus, a type/pattern-specific prepend/append is applied at the time of expansion for the actual target. For example:

```
x = s

file{foo}:                # target
{
  x += t      # s t
  y = $x y    # s t y
}

file{foo}: file{bar}      # prerequisite
{
  x += p      # x t p
  y = $x y    # x t p y
}

file{b*}:                # type/pattern
{
  x += w      # <append w>
  y = $x w    # <assign s w>
}

x = S

info $(file{bar}: x) # S w
info $(file{bar}: y) # s w
```

5 Functions

This chapter is a work in progress and is incomplete.

Functions in `build2` are organized into families, such as the `$string.*()` family for manipulating strings or `$regex.*()` for working with regular expressions. Most functions are pure and those that are not, such as `$builtin.getenv()`, are explicitly documented as such.

Some functions, such as from the `$regex.*()` family, can only be called fully qualified with their family name. For example:

```
if $regex.match($name, '(.)-(.)')
...
```

While other functions can be called without explicit qualification. For example:

```
path = $getenv('PATH')
```

There are also functions that can be called unqualified only for certain types of arguments (this fact will be reflected in their synopsis and/or documentation). Note, however, that every function can always be called qualified.

5.1 Builtin Functions

The `$builtin.*()` function family contains fundamental `build2` functions.

5.1.1 `$builtin.defined()`

```
$defined(<variable>)
```

Return true if the specified variable is defined in the calling scope or any outer scopes.

Note that this function is not pure.

5.1.2 `$builtin.visibility()`

```
$visibility(<variable>)
```

Return variable visibility if it is known and `null` otherwise.

Possible visibility value are:

```
global  -- all outer scopes
project -- this project (no outer projects)
scope   -- this scope (no outer scopes)
target  -- target and target type/pattern-specific
prereq  -- prerequisite-specific
```

Note that this function is not pure.

5.1.3 `$builtin.type()`

```
$type(<value>)
```

Return the type name of the value or empty string if untyped.

5.1.4 `$builtin.null()`

```
$null(<value>)
```

Return true if the value is `null`.

5.1.5 `$builtin.empty()`

```
$empty(<value>)
```

Return true if the value is empty.

5.1.6 `$builtin.first()`, `$builtin.second()`

```
$first(<value> [, <not_pair>])
$second(<value> [, <not_pair>])
```

Return the first or the second half of a pair, respectively. If a value is not a pair, then return `null` unless the `not_pair` argument is `true`, in which case return the non-pair value.

If multiple pairs are specified, then return the list of first/second halves. If an element is not a pair, then omit it from the resulting list unless the `not_pair` argument is `true`, in which case add the non-pair element to the list.

5.1.7 `$builtin.quote()`

```
$quote(<value> [, <escape>])
```

Quote the value returning its string representation. If `escape` is `true`, then also escape (with a backslash) the quote characters being added (this is useful if the result will be re-parsed, for example as a script command line).

5.1.8 `$builtin.getenv()`

```
$getenv(<untyped>)
```

Get the value of the environment variable. Return `null` if the environment variable is not set.

Note that if the build result can be affected by the variable being queried, then it should be reported with the `config.environment` directive.

Note that this function is not pure.

5.1.9 `$builtin.generate_uuid()`

```
$generate_uuid()
```

Generate a UUID and return its string representation in the form:

```
33f5de8d-1a29-4b1d-ba28-09e3e55c9d7a
```

Note that this function is not pure.

5.1.10 `$builtin.sha256sum()`

```
$sha256sum(<untyped>)
```

Return SHA256 checksum of the passed string.

See also `$xxh64sum()` for a faster but non-cryptographic alternative.

5.1.11 `$builtin.xxh64sum()`

```
$xxh64sum(<untyped>)
```

Return XXH64 checksum of the passed string.

See also `$sha256sum()` for a slower but cryptographic alternative.

5.2 String Functions

5.2.1 `$string.icasecmp()`

```
$string.icasecmp(<untyped>, <untyped>)
$casecmp(<string>, <string>)
```

Compare ASCII strings ignoring case and returning the boolean value.

5.2.2 `$string.contains()`

```
$string.contains(<untyped>, <untyped> [, <flags>])
$contains(<string>, <string> [, <flags>])
```

Check if the string (first argument) contains the given substring (second argument). The substring must not be empty.

The following flags are supported:

```
icase - compare ignoring case
once  - check if the substring occurs exactly once
```

See also `$string.starts_with()`, `$string.ends_with()`, `$regex.search()`, `$string.compare()`.

5.2.3 `$string.starts_with()`

```
$string.starts_with(<untyped>, <untyped> [, <flags>])
$starts_with(<string>, <string> [, <flags>])
```

Check if the string (first argument) begins with the given prefix (second argument). The prefix must not be empty.

The following flags are supported:

```
icase - compare ignoring case
```

See also `$string.contains()` and `$string.compare()`.

5.2.4 `$string.ends_with()`

```
$string.ends_with(<untyped>, <untyped> [, <flags>])
$ends_with(<string>, <string> [, <flags>])
```

Check if the string (first argument) ends with the given suffix (second argument). The suffix must not be empty.

The following flags are supported:

`icase` - compare ignoring case

See also `$string.contains()` and `$string.compare()`.

5.2.5 \$string.compare()

```
$string.compare(<untyped>, <untyped> [, <flags>])
$compare(<string>, <string> [, <flags>])
```

Compare two strings according to flags.

If no flags other than `icase` are specified, then compare strings lexicographically and return 0 if the passed strings are equivalent, -1 if the first string is less than the second one, and 1 if the first string is greater than the second one.

If any of the `contains`, `contains_once`, `starts_with`, or `ends_with` flags are specified, then check if the string (first argument) contains the sub-string (second argument) according to the flags combination. Return 0 if the sub-string is contained as requested and non-0 otherwise. The sub-string must not be empty.

The following flags are supported:

`icase` - compare ignoring case

`contains` - check if string contains sub-string

`contains_once` - check if sub-string occurs in string exactly once

`starts_with` - check if string begins with sub-string

`ends_with` - check if string ends with sub-string

See also `$string.starts_with()`, `$string.ends_with()`, `$string.contains()`.

5.2.6 \$string.replace()

```
$string.replace(<untyped>, <from>, <to> [, <flags>])
$replace(<string>, <from>, <to> [, <flags>])
```

Replace occurrences of substring *from* with *to* in a string. The *from* substring must not be empty.

The following flags are supported:

`icase` - compare ignoring case

`first_only` - only replace the first match

`last_only` - only replace the last match

If both `first_only` and `last_only` flags are specified, then *from* is replaced only if it occurs in the string once.

See also `$regex.replace()`.

5.2.7 \$string.trim()

```
$string.trim(<untyped>)
$trim(<string>)
```

Trim leading and trailing whitespaces in a string.

5.2.8 \$string.lcase(), \$string.ucase()

```
$string.lcase(<untyped>)
$string.ucase(<untyped>)
$lcase(<string>)
$ucase(<string>)
```

Convert ASCII string into lower/upper case.

5.2.9 \$string.size()

```
$size(<strings>)
$size(<string-set>)
$size(<string-map>)
$size(<string>)
```

First three forms: return the number of elements in the sequence.

Fourth form: return the number of characters (bytes) in the string.

5.2.10 \$string.front()

```
$front(<strings>)
```

Return the first string in the sequence.

5.2.11 \$string.back()

```
$back(<strings>)
```

Return the last string in the sequence.

5.2.12 \$string.sort()

```
$sort(<strings> [, <flags>])
```

Sort strings in ascending order.

The following flags are supported:

icase - sort ignoring case

dedup - in addition to sorting also remove duplicates

5.2.13 \$string.find()

`$find(<strings>, <string> [, <flags>])`

Return true if for any of the elements in the string sequence the `$compare(element, string, flags)` function call returns 0.

The following flags are supported:

icase - compare ignoring case

contains - check if string contains sub-string

contains_once - check if sub-string occurs in string exactly once

starts_with - check if string begins with sub-string

ends_with - check if string ends with sub-string

See also `$regex.find_match()`, `$regex.find_search()`, `$string.compare()`.

5.2.14 \$string.find_index()

`$find_index(<strings>, <string> [, <flags>])`

Return the index of the first element in the string sequence for which the `$compare(element, string, flags)` function call returns 0 or `$size(strings)` if no such element is found.

The following flags are supported:

icase - compare ignoring case

contains - check if string contains sub-string

contains_once - check if sub-string occurs in string exactly once

starts_with - check if string begins with sub-string

ends_with - check if string ends with sub-string

See also `$string.compare()`.

5.2.15 \$string.filter(), \$string.filter_out()

```
$filter(<strings>, <string> [, <flags>])
$filter_out(<strings>, <string> [, <flags>])
```

Return elements of a string sequence for which the `$compare(element, string, flags)` function call returns 0 (filter) or non-0 (filter_out).

The following flags are supported:

```
icase           - compare ignoring case
contains        - check if string contains sub-string
contains_once   - check if sub-string occurs in string exactly once
starts_with     - check if string begins with sub-string
ends_with       - check if string ends with sub-string
```

See also `$regex.filter_match()`, `$regex.filter_out_match()`, `$regex.filter_search()`, `$regex.filter_out_search()`, `$string.compare()`.

5.2.16 \$string.merge()

```
$string.merge(<untyped> [, <delim> [, <flags>]])
$merge(<strings> [, <delim> [, <flags>]])
```

Merge elements of a string sequence into a single string, optionally separating each subsequent element with the specified delimiter. For example:

```
info $string.merge(1 2 3, ', ') # prints 1, 2, 3
```

The following flags are supported:

```
trim           - trim elements (see $string.trim())
omit_empty     - omit empty elements
separate_first - add the delimited before the first element
separate_last  - add the delimited after the last element
```

Note that if there are no elements to merge, then `separate_first` and `separate_last` have no effect and an empty string is returned.

See also `$regex.merge()`.

5.2.17 `$string.keys()`

```
$keys(<string-map>)
```

Return the list of keys in a string map.

Note that the result is sorted in ascending order.

5.3 Integer Functions

5.3.1 `$integer.string()`

```
$string(<int64> [, <base> [, <width>]])
$string(<uint64> [, <base> [, <width>]])
```

Convert an integer to a string, optionally with the desired base and width. For example:

```
x = [uint64] 0x0000ffff

c.poptions += "-DOFFSET=$x"           # -DOFFSET=65535
c.poptions += "-DOFFSET=$string($x, 16)" # -DOFFSET=0xffff
c.poptions += "-DOFFSET=$string($x, 16, 8)" # -DOFFSET=0x0000ffff
```

Note that the minus sign is not counted for the width.

5.3.2 `$integer.integer_sequence()`

```
$integer_sequence(<begin>, <end> [, <step>])
```

Return the list of uint64 integers starting from *begin* (including) to *end* (excluding) with the specified *step* or 1 if unspecified. If *begin* is greater than *end*, empty list is returned.

5.3.3 `$integer.size()`

```
$size(<ints>)
```

Return the number of elements in the sequence.

5.3.4 `$integer.front()`

```
$front(<ints>)
```

Return the first integer in the sequence.

5.3.5 `$integer.back()`

```
$back(<ints>)
```

Return the last integer in the sequence.

5.3.6 `$integer.sort()`

```
$sort(<ints> [, <flags>])
```

Sort integers in ascending order.

The following flags are supported:

```
dedup - in addition to sorting also remove duplicates
```

5.3.7 `$integer.find()`

```
$find(<ints>, <int>)
```

Return true if the integer sequence contains the specified integer.

5.3.8 `$integer.find_index()`

```
$find_index(<ints>, <int>)
```

Return the index of the first element in the integer sequence that is equal to the specified integer or `$size(ints)` if none is found.

5.4 Bool Functions

5.4.1 `$bool.string()`

```
$string(<bool>)
```

Convert a boolean value to a string literal `true` or `false`.

5.5 Path Functions

The `$path.*()` function family contains function that manipulating filesystem paths.

5.5.1 `$path.string()`

```
$string(<paths>)
```

Return the traditional string representation of a path (or a list of string representations for a list of paths). In particular, for directory paths, the traditional representation does not include the trailing directory separator (except for the POSIX root directory). See `$representation()` below for the precise string representation.

5.5.2 `$path.posix_string()`

```
$posix_string(<paths>)
$path.posix_string(<untyped>)
```

Return the traditional string representation of a path (or a list of string representations for a list of paths) using the POSIX directory separators (forward slashes).

5.5.3 `$path.representation()`

```
$representation(<paths>)
```

Return the precise string representation of a path (or a list of string representations for a list of paths). In particular, for directory paths, the precise representation includes the trailing directory separator. See `$string()` above for the traditional string representation.

5.5.4 `$path.posix_representation()`

```
$posix_representation(<paths>)
$path.posix_representation(<untyped>)
```

Return the precise string representation of a path (or a list of string representations for a list of paths) using the POSIX directory separators (forward slashes).

5.5.5 `$path.absolute()`

```
$absolute(<path>)
$path.absolute(<untyped>)
```

Return true if the path is absolute and false otherwise.

5.5.6 `$path.simple()`

```
$simple(<path>)
$path.simple(<untyped>)
```

Return true if the path is simple, that is, has no directory component, and false otherwise.

Note that on POSIX `/foo` is not a simple path (it is `foo` in the root directory) while `/` is (it is the root directory).

5.5.7 `$path.sub_path()`

```
$sub_path(<path>, <path>)
$path.sub_path(<untyped>, <untyped>)
```

Return true if the path specified as the first argument is a sub-path of the one specified as the second argument (in other words, the second argument is a prefix of the first) and false otherwise. Both paths are expected to be normalized. Note that this function returns true if the paths are equal. Empty path is considered a prefix of any path.

5.5.8 \$path.super_path()

```
$super_path(<path>, <path>)
$path.super_path(<untyped>, <untyped>)
```

Return true if the path specified as the first argument is a super-path of the one specified as the second argument (in other words, the second argument is a suffix of the first) and false otherwise. Both paths are expected to be normalized. Note that this function returns true if the paths are equal. Empty path is considered a suffix of any path.

5.5.9 \$path.directory()

```
$directory(<paths>)
$path.directory(<untyped>)
```

Return the directory part of a path (or a list of directory parts for a list of paths) or an empty path if there is no directory. A directory of a root directory is an empty path.

5.5.10 \$path.root_directory()

```
$root_directory(<paths>)
$path.root_directory(<untyped>)
```

Return the root directory of a path (or a list of root directories for a list of paths) or an empty path if the specified path is not absolute.

5.5.11 \$path.leaf()

```
$leaf(<paths>)
$path.leaf(<untyped>)
$leaf(<paths>, <dir-path>)
$path.leaf(<untyped>, <dir-path>)
```

First form (one argument): return the last component of a path (or a list of last components for a list of paths).

Second form (two arguments): return a path without the specified directory part (or a list of paths without the directory part for a list of paths). Return an empty path if the paths are the same. Issue diagnostics and fail if the directory is not a prefix of the path. Note: expects both paths to be normalized.

5.5.12 \$path.relative()

```
$relative(<paths>, <dir-path>)
$path.relative(<untyped>, <dir-path>)
```

Return the path relative to the specified directory that is equivalent to the specified path (or a list of relative paths for a list of specified paths). Issue diagnostics and fail if a relative path cannot be derived (for example, paths are on different drives on Windows).

Note: to check if a path is relative, use `$path.absolute()`.

5.5.13 `$path.base()`

```
$base(<paths>)
$path.base(<untyped>)
```

Return the base part (without the extension) of a path (or a list of base parts for a list of paths).

5.5.14 `$path.extension()`

```
$extension(<path>)
$path.extension(<untyped>)
```

Return the extension part (without the dot) of a path or empty string if there is no extension.

5.5.15 `$path.complete()`

```
$complete(<paths>)
$path.complete(<untyped>)
```

Complete the path (or list of paths) by prepending the current working directory unless the path is already absolute.

5.5.16 `$path.canonicalize()`

```
$canonicalize(<paths>)
$path.canonicalize(<untyped>)
```

Canonicalize the path (or list of paths) by converting all the directory separators to the canonical form for the host platform. Note that multiple directory separators are not collapsed.

5.5.17 `$path.normalize()`, `$path.try_normalize()`

```
$normalize(<paths>)
$path.normalize(<untyped>)
$try_normalize(<path>)
$path.try_normalize(<untyped>)
```

Normalize the path (or list of paths) by collapsing the `.` and `..` components if possible, collapsing multiple directory separators, and converting all the directory separators to the canonical form for the host platform.

If the resulting path would be invalid, the `$normalize()` version issues diagnostics and fails while the `$try_normalize()` version returns `null`. Note that `$try_normalize()` only accepts a single path.

5.5.18 \$path.actualize(), \$path.try_actualize()

```
$actualize(<paths>)
$path.actualize(<untyped>)
$try_actualize(<path>)
$path.try_actualize(<untyped>)
```

Actualize the path (or list of paths) by first normalizing it and then for host platforms with case-insensitive filesystems obtaining the actual spelling of the path.

Only an absolute path can be actualized. If a path component does not exist, then its (and all subsequent) spelling is unchanged. Note that this is a potentially expensive operation.

If the resulting path would be invalid or in case of filesystem errors (other than non-existent component), the `$actualize()` version issues diagnostics and fails while the `$try_actualize()` version returns null. Note that `$try_actualize()` only accepts a single path.

Note that this function is not pure.

5.5.19 \$path.size()

```
$size(<paths>)
$size(<path>)
```

First form: return the number of elements in the paths sequence.

Second form: return the number of characters (bytes) in the path. Note that for `dir_path` the result does not include the trailing directory separator (except for the POSIX root directory).

5.5.20 \$path.front()

```
$front(<paths>)
```

Return the first path in the sequence.

5.5.21 \$path.back()

```
$back(<paths>)
```

Return the last path in the sequence.

5.5.22 \$path.sort()

```
$sort(<paths> [, <flags>])
```

Sort paths in ascending order. Note that on host platforms with a case-insensitive filesystem the order is case-insensitive.

The following flags are supported:

`dedup` - in addition to sorting also remove duplicates

5.5.23 \$path.find()

`$find(<paths>, <path>)`

Return true if the paths sequence contains the specified path. Note that on host platforms with a case-insensitive filesystem the comparison is case-insensitive.

5.5.24 \$path.find_index()

`$find_index(<paths>, <path>)`

Return the index of the first element in the paths sequence that is equal to the specified path or `$size(paths)` if none is found. Note that on host platforms with a case-insensitive filesystem the comparison is case-insensitive.

5.5.25 \$path.match()

`$path.match(<entry>, <pattern> [, <start-dir>])`

Match a filesystem entry name against a name pattern (both are strings), or a filesystem entry path against a path pattern. For the latter case the start directory may also be required (see below). The pattern is a shell-like wildcard pattern. The semantics of the *pattern* and *entry* arguments is determined according to the following rules:

1. The arguments must be of the string or path types, or be untyped.
2. If one of the arguments is typed, then the other one must be of the same type or be untyped. In the later case, an untyped argument is converted to the type of the other argument.
3. If both arguments are untyped and the start directory is specified, then the arguments are converted to the path type.
4. If both arguments are untyped and the start directory is not specified, then, if one of the arguments is syntactically a path (the value contains a directory separator), then they are converted to the path type, otherwise -- to the string type (match as names).

If pattern and entry paths are both either absolute or relative and not empty, and the first pattern component is not a self-matching wildcard (doesn't contain `***`), then the start directory is not required, and is ignored if specified. Otherwise, the start directory must be specified and be an absolute path.

5.6 Name Functions

The `$name.*()` function family contains function that operate on target and prerequisite names. See also the `$target.*()` function family for functions that operate on actual targets.

5.6.1 `$name.name()`

```
$name(<names>)
```

Return the name of a target (or a list of names for a list of targets).

5.6.2 `$name.extension()`

```
$extension(<name>)
```

Return the extension of a target.

Note that this function returns `null` if the extension is unspecified (default) and empty string if it's specified as no extension.

5.6.3 `$name.directory()`

```
$directory(<names>)
```

Return the directory of a target (or a list of directories for a list of targets).

5.6.4 `$name.target_type()`

```
$target_type(<names>)
```

Return the target type name of a target (or a list of target type names for a list of targets).

5.6.5 `$name.project()`

```
$project(<name>)
```

Return the project of a target or `null` if not project-qualified.

5.6.6 `$name.is_a()`

```
$is_a(<name>, <target-type>)
```

Return true if the *name*'s target type is-a *target-type*. Note that this is a dynamic type check that takes into account target type inheritance.

5.6.7 `$name.filter()`, `$name.filter_out()`

```
$filter(<names>, <target-types>)  
$filter_out(<names>, <target-types>)
```

Return names with target types which are-a (`filter`) or not are-a (`filter_out`) one of *target-types*. See `$is_a()` for background.

5.6.8 `$name.size()`

```
$size(<names>)
```

Return the number of elements in the sequence.

5.6.9 `$name.front()`

```
$front(<names>)
```

Return the first element in the sequence.

5.6.10 `$name.back()`

```
$back(<names>)
```

Return the last element in the sequence.

5.6.11 `$name.sort()`

```
$sort(<names> [, <flags>])
```

Sort names in ascending order.

The following flags are supported:

`dedup` - in addition to sorting also remove duplicates

5.6.12 `$name.find()`

```
$find(<names>, <name>)
```

Return true if the name sequence contains the specified name.

5.6.13 `$name.find_index()`

```
$find_index(<names>, <name>)
```

Return the index of the first element in the name sequence that is equal to the specified name or `$size(names)` if none is found.

5.7 Target Functions

The `$target.*()` function family contains function that operate on targets. See also the `$name.*()` function family for functions that operate on target (and prerequisite) names.

5.7.1 `$target.path()`

```
$path(<names>)
```

Return the path of a target (or a list of paths for a list of targets). The path must be assigned, which normally happens during match. As a result, this function is normally called from a recipe, but can also be called from a buildfile provided the target has been updated during load.

Note that while this function is technically not pure, we don't mark it as such since it can only be called (normally from a recipe) after the target has been matched, meaning that this target is a prerequisite and therefore this impurity has been accounted for.

5.7.2 `$target.process_path()`

```
$process_path(<name>)
```

Return the process path of an executable target.

Note that while this function is not technically pure, we don't mark it as such for the same reasons as for `$path()` above.

5.8 Regex Functions

The `$regex.*()` function family contains function that provide comprehensive regular expression matching and substitution facilities. The supported regular expression flavor is ECMAScript, more precisely, ECMA-262-based C++11 regular expressions. Note that the `match_not_null` flag is in effect unless the string being matched is empty.

In the `$regex.*()` functions the substitution escape sequences in the format string (the *fmt* argument) are extended with a subset of the Perl escape sequences: `\n`, `\u`, `\l`, `\U`, `\L`, `\E`, `\1` ... `\9`, and `\\`. Note that the standard ECMAScript escape sequences (`$1`, `$2`, `$&`, etc) are still supported.

Note that functions from the `$regex.*()` family can only be called fully qualified with their family name. For example:

```
if $regex.match($name, '(.)-(.)')
...

```

5.8.1 \$regex.match()

```
$regex.match(<val>, <pat> [, <flags>])
```

Match a value of an arbitrary type against the regular expression. Convert the value to string prior to matching. Return the boolean value unless `return_subs` flag is specified (see below), in which case return names (or null if no match).

The following flags are supported:

```

icase          - match ignoring case

return_subs    - return names (rather than boolean), that contain
                  sub-strings that match the marked sub-expressions
                  and null if no match

```

5.8.2 \$regex.find_match()

```
$regex.find_match(<vals>, <pat> [, <flags>])
```

Match list elements against the regular expression and return true if the match is found. Convert the elements to strings prior to matching.

The following flags are supported:

```
icase - match ignoring case
```

5.8.3 \$regex.filter_match(), \$regex.filter_out_match()

```

$regex.filter_match(<vals>, <pat> [, <flags>])
$regex.filter_out_match(<vals>, <pat> [, <flags>])

```

Return elements of a list that match (`filter`) or do not match (`filter_out`) the regular expression. Convert the elements to strings prior to matching.

The following flags are supported:

```
icase - match ignoring case
```

5.8.4 \$regex.search()

```
$regex.search(<val>, <pat> [, <flags>])
```

Determine if there is a match between the regular expression and some part of a value of an arbitrary type. Convert the value to string prior to searching. Return the boolean value unless `return_match` or `return_subs` flag is specified (see below) in which case return names (null if no match).

The following flags are supported:

```

icase           - match ignoring case

return_match - return names (rather than boolean), that contain a
                  sub-string that matches the whole regular expression
                  and null if no match

return_subs  - return names (rather than boolean), that contain
                  sub-strings that match the marked sub-expressions
                  and null if no match

```

If both `return_match` and `return_subs` flags are specified then the sub-string that matches the whole regular expression comes first.

See also `$string.contains()`, `$string.starts_with()`, `$string.ends_with()`.

5.8.5 \$regex.find_search()

```
$regex.find_search(<vals>, <pat> [, <flags>])
```

Determine if there is a match between the regular expression and some part of any of the list elements and return true if the match is found. Convert the elements to strings prior to matching.

The following flags are supported:

```
icase - match ignoring case
```

5.8.6 \$regex.filter_search(), \$regex.filter_out_search()

```

$regex.filter_search(<vals>, <pat> [, <flags>])
$regex.filter_out_search(<vals>, <pat> [, <flags>])

```

Return elements of a list for which there is a match (`filter`) or no match (`filter_out`) between the regular expression and some part of the element. Convert the elements to strings prior to matching.

The following flags are supported:

```
icase - match ignoring case
```

5.8.7 \$regex.replace()

```
$regex.replace(<val>, <pat>, <fmt> [, <flags>])
```

Replace matched parts in a value of an arbitrary type, using the format string. Convert the value to string prior to matching. The result value is always untyped, regardless of the argument type.

The following flags are supported:

```

icase                - match ignoring case

format_first_only    - only replace the first match

format_no_copy       - do not copy unmatched value parts into the
                      result

```

If both `format_first_only` and `format_no_copy` flags are specified then the result will only contain the replacement of the first match.

See also `$string.replace()`.

5.8.8 \$regex.replace_lines()

```
$regex.replace_lines(<val>, <pat>, <fmt> [, <flags>])
```

Convert the value to string, parse it into lines and for each line apply the `$regex.replace()` function with the specified pattern, format, and flags. If the format argument is null, omit the "all-null" replacements for the matched lines from the result. Return unmatched lines and line replacements as a name list unless `return_lines` flag is specified (see below), in which case return a single multi-line simple name value.

The following flags are supported in addition to the `$regex.replace()` function's flags:

```

return_lines - return the simple name (rather than a name list)
               containing the unmatched lines and line replacements
               separated with newlines.

```

Note that if `format_no_copy` is specified, unmatched lines are not copied either.

5.8.9 \$regex.split()

```
$regex.split(<val>, <pat>, <fmt> [, <flags>])
```

Split a value of an arbitrary type into a list of unmatched value parts and replacements of the matched parts, omitting empty ones (unless the `format_copy_empty` flag is specified). Convert the value to string prior to matching.

The following flags are supported:

```

icase                - match ignoring case

format_no_copy       - do not copy unmatched value parts into the
                      result

format_copy_empty    - copy empty elements into the result

```

5.8.10 `$regex.merge()`

```
$regex.merge(<vals>, <pat>, <fmt> [, <delim> [, <flags>]])
```

Replace matched parts in a list of elements using the regex format string. Convert the elements to strings prior to matching. The result value is untyped and contains concatenation of transformed non-empty elements (unless the `format_copy_empty` flag is specified) optionally separated with a delimiter.

The following flags are supported:

```

icase                - match ignoring case

format_first_only - only replace the first match

format_no_copy     - do not copy unmatched value parts into the
                      result

format_copy_empty - copy empty elements into the result

```

If both `format_first_only` and `format_no_copy` flags are specified then the result will be a concatenation of only the first match replacements.

5.8.11 `$regex.apply()`

```
$regex.apply(<vals>, <pat>, <fmt> [, <flags>])
```

Replace matched parts of each element in a list using the regex format string. Convert the elements to strings prior to matching. Return a list of transformed elements, omitting the empty ones (unless the `format_copy_empty` flag is specified).

The following flags are supported:

```

icase                - match ignoring case

format_first_only - only replace the first match

format_no_copy     - do not copy unmatched value parts into the
                      result

format_copy_empty - copy empty elements into the result

```

If both `format_first_only` and `format_no_copy` flags are specified then the result elements will only contain the replacement of the first match.

5.9 JSON Functions

The `$json.*()` function family contains functions that operate on the JSON types, `json`, `json_array`, and `json_object`, which represent JSON, JSON5, or JSON5E value. For example:

5.9.1 \$json.value_type()

```
j = [json] '{
  one   : 1,
  two   : "abc",
  three : {x:1, y:-1}
}'

# Alternative construction syntax for the same value.
#
#j = [json] one@1 two@abc three@([json] x@1 y@-1)

for m: $j
{
  n = $member_name($m)
  v = $member_value($m)

  info $n $value_type($v) $v
}
```

Note also that conversions of JSON values to other types are performed with explicit type casts. For example:

```
j = [json] '[1, 2, 3]'
a = [uint64s] $j
l = $regex.merge([strings] $j, '.*', '$&\n')
```

5.9.1 \$json.value_type()

`$value_type(<json> [, <distinguish_numbers>])`

Return the type of a JSON value: null, boolean, number, string, array, or object. If the *distinguish_numbers* argument is true, then instead of number return signed number, unsigned number, hexadecimal signed number, or hexadecimal unsigned number. Note that the hexadecimal numbers are only available in JSON5 and JSON5E.

5.9.2 \$json.value_size()

`$value_size(<json>)`

Return the size of a JSON value.

The size of a null value is 0. The sizes of simple values (boolean, number, and string) is 1. The size of array and object values is the number of elements and members, respectively.

Note that the size of a string JSON value is not the length of the string. To get the length call `$string.size()` instead by casting the JSON value to the string value type.

5.9.3 \$json.member_name()

`$member_name(<json-member>)`

Return the name of a JSON object member.

5.9.4 \$json.member_value()

```
$member_value(<json-member>)
```

Return the value of a JSON object member.

5.9.5 \$json.object_names()

```
$object_names(<json-object>)
```

Return the list of names in the JSON object. If the JSON `null` is passed instead, assume it is a missing object and return an empty list.

5.9.6 \$json.array_size()

```
$array_size(<json-array>)
```

Return the number of elements in the JSON array. If the JSON `null` value is passed instead, assume it is a missing array and return 0.

5.9.7 \$json.array_front()

```
$array_front(<json-array>)
```

Return the first element in the JSON array. If the JSON `null` value is passed instead, assume it is a missing array and treat it as empty.

5.9.8 \$json.array_back()

```
$array_back(<json-array>)
```

Return the last element in the JSON array. If the JSON `null` value is passed instead, assume it is a missing array and treat it as empty.

5.9.9 \$json.array_find()

```
$array_find(<json-array>, <json>)
```

Return true if the JSON array contains the specified JSON value. If the JSON `null` value is passed instead, assume it is a missing array and return `false`.

5.9.10 \$json.array_find_index()

```
$array_find_index(<json-array>, <json>)
```

Return the index of the first element in the JSON array that is equal to the specified JSON value or `$array_size(json-array)` if none is found. If the JSON `null` value is passed instead, assume it is a missing array and return 0.

5.9.11 \$json.load()

```
$json.load(<path> [, <flags>])
```

Parse the contents of the specified file as JSON (default), JSON5, or JSON5E input text and return the result as a value of the `json` type.

The following flags are supported:

```
json    - parse as JSON input text (default)
```

```
json5   - parse as JSON5 input text
```

```
json5e  - parse as JSON5E input text
```

See also `$json.parse()`.

Note that this function is not pure.

5.9.12 \$json.parse()

```
$json.parse(<text> [, <flags>])
```

Parse the specified JSON (default), JSON5, or JSON5E input text and return the result as a value of the `json` type.

The following flags are supported:

```
json    - parse as JSON input text (default)
```

```
json5   - parse as JSON5 input text
```

```
json5e  - parse as JSON5E input text
```

See also `$json.load()` and `$json.serialize()`.

5.9.13 \$json.serialize()

```
$serialize(<json> [, <indentation> [, <flags>]])
```

Serialize the specified JSON value and return the resulting JSON (default), JSON5, or JSON5E output text.

The optional *indentation* argument specifies the number of indentation spaces that should be used for pretty-printing. If 0 is passed, then no pretty-printing is performed. The default is 2 spaces.

The following flags are supported:

```
json    - serialize as JSON output text (default)
```

```
json5   - serialize as JSON5 output text
```

```
json5e  - serialize as JSON5E output text
```

Note that the exact semantics of JSON5 and JSON5E serialization is implementation-defined in that certain constructs may still be serialized using their JSON syntax.

See also `$json.parse()`.

5.9.14 `$json.size()`

```
$size(<json-set>)
$size(<json-map>)
```

Return the number of elements in the sequence.

5.9.15 `$json.keys()`

```
$keys(<json-map>)
```

Return the list of keys in a json map as a json array.

Note that the result is sorted in ascending order.

5.10 Process Functions

5.10.1 `$process.run()`

```
$process.run(<prog>[ <args>...])
```

Run builtin or external program and return trimmed `stdout` output as a single value. If you need to split the output into lines, see the `$process.run_regex()` function. If you need to split the output in other ways (words, etc), see the `$regex.*()` function family.

Note that if the result of executing the program can be affected by environment variables and this result can in turn affect the build result, then such variables should be reported with the `config.environment` directive.

Note that this function is not pure and can only be called during the load phase.

5.10.2 `$process.run_regex()`

```
$process.run_regex(<prog>[ <args>...], <pat> [, <fmt>])
```

Run builtin or external program and return `stdout` output lines matched and optionally processed with a regular expression. The result is a list of values, one per line.

Each line of `stdout` (including the customary trailing blank) is matched (as a whole) against *pat* and, if successful, returned, optionally processed with *fmt*, as an element of a list. See the `$regex.*()` function family for details on regular expressions and format strings.

Note that if the result of executing the program can be affected by environment variables and this result can in turn affect the build result, then such variables should be reported with the `config.environment` directive.

Note that this function is not pure and can only be called during the load phase.

5.10.3 `$process.search()`

```
$process.search(<prog>)
```

Return the effective path of an executable, that is, the absolute path to the executable that will also include any omitted extensions, etc. Return `null` if the executable is not found.

Note that this function is not pure.

5.11 Filesystem Functions

5.11.1 `$filesystem.file_exists()`

```
$file_exists(<path>)
```

Return true if a filesystem entry at the specified path exists and is a regular file (or is a symlink to a regular file) and false otherwise.

Note that this function is not pure.

5.11.2 `$filesystem.directory_exists()`

```
$directory_exists(<path>)
```

Return true if a filesystem entry at the specified path exists and is a directory (or is a symlink to a directory) and false otherwise.

Note that this function is not pure.

5.11.3 `$filesystem.path_search()`

```
$path_search(<pattern> [, <start-dir>])
```

Return filesystem paths that match the shell-like wildcard pattern. If the pattern is an absolute path, then the start directory is ignored (if present). Otherwise, the start directory must be specified and be absolute, except for Shellsript. For Shellsript, if the start directory is not specified, then the current working directory is assumed, and if the relative start directory is specified, then the current working directory is used as a base.

Note that this function is not pure.

5.12 Project Name Functions

The `$project_name.*()` function family contains function that operate on the `project_name` type.

5.12.1 `$project_name.string()`

```
$string(<project-name>)
```

Return the string representation of a project name. See also the `$variable()` function below.

5.12.2 `$project_name.base()`

```
$base(<project-name> [, <extension>])
```

Return the base part (without the extension) of a project name.

If *extension* is specified, then only remove that extension. Note that *extension* should not include the dot and the comparison is always case-insensitive.

5.12.3 `$project_name.extension()`

```
$extension(<project-name>)
```

Return the extension part (without the dot) of a project name or empty string if there is no extension.

5.12.4 `$project_name.variable()`

```
$variable(<project-name>)
```

Return the string representation of a project name that is sanitized to be usable as a variable name. Specifically, `.`, `-`, and `+` are replaced with `_`.

5.13 Process Path Functions

The `$process_path.*()` function family contains function that operate on the `process_path` type and its extended `process_path_ex` variant. These types describe a path to an executable that, if necessary, has been found in `PATH`, completed with an extension, etc. The `process_path_ex` variant includes additional metadata, such as the stable process name for diagnostics and the executable checksum for change tracking.

5.13.1 `$process_path.recall()`

```
$recall(<process-path>)
```

Return the recall path of an executable, that is, a path that is not necessarily absolute but which nevertheless can be used to re-run the executable in the current environment. This path, for example, could be used in diagnostics when printing the failing command line.

5.13.2 `$process_path.effect()`

```
$effect(<process-path>)
```

Return the effective path of an executable, that is, the absolute path to the executable that will also include any omitted extensions, etc.

5.13.3 `$process_path.name()`

```
$name(<process-path-ex>)
```

Return the stable process name for diagnostics.

5.13.4 `$process_path.checksum()`

```
$checksum(<process-path-ex>)
```

Return the executable checksum for change tracking.

5.13.5 `$process_path.env_checksum()`

```
$env_checksum(<process-path-ex>)
```

Return the environment checksum for change tracking.

5.14 Target Triplet Functions

The `$target_triplet.*()` function family contains function that operate on the `target_triplet` type that represents the ubiquitous *cpu-vendor-os* target platform triplet.

5.14.1 `$target_triplet.string()`

```
$string(<target-triplet>)
```

Return the canonical (that is, without the unknown vendor component) target triplet string.

5.14.2 `$target_triplet.representation()`

```
$representation(<target-triplet>)
```

Return the complete target triplet string that always contains the vendor component.

6 Directives

This chapter is a work in progress and is incomplete.

6.1 define

```
define <derived>: <base>
```

Define a new target type `<derived>` by inheriting from existing target type `<base>`. See Target Types for details.

6.2 include

```
include <file>
include <directory>
```

Load the specified file (the first form) or `buildfile` in the specified directory (the second form). In both cases the file is loaded in the scope corresponding to its directory. Subsequent inclusions of the same file are automatically ignored. See also `source`.

6.3 source

```
source <file>
```

Load the specified file in the current scope as if its contents were copied and pasted in place of the `source` directive. Note that subsequent sourcing of the same file in the same scope are not automatically ignored. See also `include`.

6.4 update

```
update <target>...
```

Update the specified targets during load before continuing evaluating the `buildfile`. Updating a target during load is primarily useful when the information it contains is required in the `buildfile` itself. For example, we may need to know the target architecture byte-order in order to decide which source files must be included into the build. And, at least in case of the C/C++ compilation, the only reliable source of this information are the compiler macros. To extract this information from the compiler and make it available during the `buildfile` evaluation, we can generate a `buildfile` fragment during load and then source it into the main `buildfile`. For example (see C Compiler Predefined Macro Extraction for the actual extraction):

```
./: buildfile{byte-order} # Make sure it gets cleaned.

buildfile{byte-order}:
{{
    diag gen $>
    echo 'little_endian = true' >$path($>)
}}

update buildfile{byte-order}

source $path(buildfile{byte-order})
```

```
./: exe{hello}: cxx{hello}
exe{hello}: cxx{hello-big}:    include = (!$little_endian)
exe{hello}: cxx{hello-little}: include = $little_endian
```

Once a target is updated during load, its path can be queried with the `$path()` function (or `$process_path()`, for executables). The resulting path is normally used to load the information contained in the target into the `buildfile`, for example, using the `source` directive as in the above example, using the `$json.load()` function if the target is a JSON file, or using the `run` directive or `$process.run*()` function if the target is either an executable or the information it contains can be extracted using builtin commands such as `cat` or `sed`.

Updating targets during load should only be used as a last resort because such updates happen serially and block further `buildfile` evaluation until completed. Even during the incremental build where the target in question is already up-to-date, this check is performed serially during load whereas it would be performed in parallel with other targets if updated normally.

More precisely, each `update` directive is processed serially as the `buildfile` is evaluated. However, targets specified in each directive (and their prerequisites) are updated in parallel. As result, if you need to update several targets during load, it is beneficial to do it with a single `update` directive if possible.

As a result, if you do need to update certain targets during load, try to make the update (and the up-to-date check) as fast as possible by limiting the amount of work done during load to the absolute minimum, including limiting the number of prerequisites to only what's necessary. If a certain prerequisite, for example, generated `config.h`, is also used for other purposes and thus contains more information than what's needed during load, consider making a smaller version (or even a static version, if possible) specifically for update during load.

Another reason to avoid update during load unless absolutely necessary is the counter-intuitive behavior during operations other than update. Since updating such targets is required to continue evaluating the `buildfile`, this is performed regardless of the operation requested by the user, including, for example, during `clean`. As a result, the user may see targets being updated in unexpected situations.

While there is nothing we can do about `clean`, we can limit update during load to the perform meta-operation only. This will be especially helpful for the `configure` meta-operation since seeing update commands during configuration will be surprising, especially if they are coming from a third-party dependency.

Limiting update during load to `perform` is safe to do if the accurate information is only required during `perform` and we are able to provide a suitable fallback for other meta-operations. For example:

```

./: buildfile{byte-order} # Make sure it gets cleaned.

buildfile{byte-order}:
{{
    diag gen $>
    echo 'little_endian = true' >$path($>)
}}

if ($build.meta_operation == 'perform')
{
    update buildfile{byte-order}
    source $path(buildfile{byte-order})
}
else
{
    little_endian = true
}

./: exe{hello}: cxx{hello}
exe{hello}: cxx{hello-big}:    include = (!$little_endian)
exe{hello}: cxx{hello-little}: include = $little_endian

```

One scenario where limiting update during load to perform would not work is if, for example, we were using the extracted information to set default values of the configuration variables.

A target being updated during load should be defined in the `buildfile` containing the `update` directive and this `buildfile` should be standalone, that is, it should be possible to load it on its own. Note also that everything pertaining to updating the target (prerequisites, options, recipes/rules, etc) should be in effect before the `update` directive since the target will be updated without evaluating the rest of the `buildfile`.

The update during load mechanism also has a number of restrictions and limitations:

- It cannot be used in `bootstrap.build` (but can be used in `root.build`).
- Targets (and their prerequisites) that are updated during load should not be aliases (`alias{}`, including `dir{}`).
- Post hoc prerequisites will not yet be updated after the `update` directive.

7 Attributes

This chapter is a work in progress and is incomplete.

The only currently recognized target attribute is `rule_hint` which specifies the rule hint. Rule hints can be used to resolve ambiguity when multiple rules match the same target as well as to override an unambiguous match. For example, the following rule hint makes sure our executable is linked with the C++ compiler even though it only has C sources:

```
[rule_hint=cxx] exe{hello}: c{hello}
```

8 Name Patterns

For convenience, in certain contexts, names can be generated with shell-like wildcard patterns. A name is a *name pattern* if its value contains one or more unquoted wildcard characters or character sequences. For example:

```
./: */                # All (immediate) subdirectories
exe{hello}: {hxx cxx}{**} # All C++ header/source files.
pattern = '*.txt'      # Literal '*.txt'.
```

Pattern-based name generation is not performed in certain contexts. Specifically, it is not performed in target names where it is interpreted as a pattern for target type/pattern-specific variable assignments. For example.

```
s = *.txt             # Variable assignment (performed).
./: cxx{*}            # Prerequisite names (performed).
cxx{*}: dist = false  # Target pattern (not performed).
```

In contexts where it is performed, it can be inhibited with quoting, for example:

```
pat = 'foo*bar'
./: cxx{'foo*bar'}
```

The following wildcards are recognized:

```
*      - match any number of characters (including zero)
?      - match any single character
[...] - match a character with a bracket expression
```

Currently only literal character and range bracket expressions are supported. Specifically, no character or equivalence classes, etc., are supported nor the special characters backslash-escaping. See the "Pattern Matching Notation" section in the POSIX "Shell Command Language" specification for details.

Note that some wildcard characters may have special meaning in certain contexts. For instance, [at the beginning of a value will be interpreted as the start of the attribute list while ? and [in the eval context are part of the ternary operator and value subscript, respectively. In such cases the character will need to be escaped in order to be treated as a wildcard, for example:

```
x = \[1-9]-foo.txt
y = (foo.\?xx)
z = ($foo\[123].txt)
```

If a pattern ends with a directory separator, then it only matches directories. Otherwise, it only matches files. Matches that start with a dot (.) are automatically ignored unless the pattern itself also starts with this character.

In addition to the above wildcards, ** and *** are recognized as wildcard sequences. If a pattern contains **, then it is matched just like * but in all the subdirectories, recursively, but excluding directories that contain the .buildignore file. The *** wildcard behaves like

****** but also matches the start directory itself. For example:

```
exe{hello}: cxx{**} # All C++ source files recursively.
```

A group-enclosed (`{}`) pattern value may be followed by inclusion/exclusion patterns/matches. A subsequent value is treated as an inclusion or exclusion if it starts with a literal, unquoted plus (+) or minus (-) sign, respectively. In this case the remaining group values, if any, must all be inclusions or exclusions. If the second value doesn't start with a plus or minus, then all the group values are considered independent with leading pluses and minuses not having any special meaning. For regularity as well as to allow patterns without wildcards, the first pattern can also start with the plus sign. For example:

```
exe{hello}: cxx{f* -foo}           # Exclude foo if exists.
exe{hello}: cxx{f* +bar}           # Include bar if exists.
exe{hello}: cxx{f* -fo?}           # Exclude foo and fox if exist.
exe{hello}: cxx{f* +b* -foo -bar}  # Exclude foo and bar if exist.
exe{hello}: cxx{+f* +b* -foo -bar} # Same as above.
exe{hello}: cxx{+foo}              # Pattern without wildcards.
exe{hello}: cxx{f* b* -z*}         # Names matching three patterns.
```

Inclusions and exclusions are applied in the order specified and only to the result produced up to that point. The order of names in the result is unspecified. However, it is guaranteed not to contain duplicates. The first pattern and the following inclusions/exclusions must be consistent with regards to the type of filesystem entry they match. That is, they should all match either files or directories. For example:

```
exe{hello}: cxx{f* -foo +*oo} # Exclusion has no effect.
exe{hello}: cxx{f* +*oo}      # Ok, no duplicates.
./: {*/ -build}              # Error: exclusion not a directory.
```

As a more realistic example, let's say we want to exclude source files that reside in the `test/` directories (and their subdirectories) anywhere in the tree. This can be achieved with the following pattern:

```
exe{hello}: cxx{** -***/test/**}
```

Similarly, if we wanted to exclude all source files that have the `-test` suffix:

```
exe{hello}: cxx{** -**test}
```

In contrast, the following pattern only excludes such files from the top directory:

```
exe{hello}: cxx{** -*test}
```

If many inclusions or exclusions need to be specified, then an inclusion/exclusion group can be used. For example:

```
exe{hello}: cxx{f* -{foo bar}}
exe{hello}: cxx{+{f* b*} -{foo bar}}
```

This is particularly useful if you would like to list the names to include or exclude in a variable. For example, this is how we can exclude certain files from compilation but still include them as ordinary file prerequisites (so that they are still included into the source distribution):

```
exc = foo.cxx bar.cxx
exe{hello}: cxx{+{f* b*} -{$exc}} file{$exc}
```

If we want to specify our pattern in a variable, then we have to use the explicit inclusion syntax, for example:

```
pat = 'f*'
exe{hello}: cxx{+$pat} # Pattern match.
exe{hello}: cxx{$pat} # Literal 'f*'.

pat = '+f*'
exe{hello}: cxx{$pat} # Literal '+f*'.

inc = 'f*' 'b*'
exc = 'f*o' 'b*r'
exe{hello}: cxx{+{$inc} -{$exc}}
```

One common situation that calls for exclusions is auto-generated source code. Let's say we have auto-generated command line parser in `options.hxx` and `options.cxx`. Because of the in/out of source builds, our name pattern may or may not find these files. Note, however, that we cannot just include them as non-pattern prerequisites. We also have to exclude them from the pattern match since otherwise we may end up with duplicate prerequisites. As a result, this is how we have to handle this case provided we want to continue using patterns to find other, non-generated source files:

```
exe{hello}: {hxx cxx}{* -options} {hxx cxx}{options}
```

If all our auto-generated source files have a common prefix or suffix, then we can exclude them wholesale with a pattern. For example, if all our generated files end with the '-options' suffix:

```
exe{hello}: {hxx cxx}{** -**-options} {hxx cxx}{foo-options bar-options}
```

If the name pattern includes an absolute directory, then the pattern match is performed in that directory and the generated names include absolute directories as well. Otherwise, the pattern match is performed in the *pattern base* directory. In buildfiles this is `src_base` while on the command line – the current working directory. In this case the generated names are relative to the base directory. For example, assuming we have the `foo.cxx` and `b/bar.cxx` source files:

```
exe{hello}: $src_base/cxx{**} # $src_base/cxx{foo} $src_base/b/cxx{bar}
exe{hello}: cxx{**} # cxx{foo} b/cxx{bar}
```

Pattern matching as well as inclusion/exclusion logic is target type-specific. If the name pattern does not contain a type, then the `dir{}` type is assumed if the pattern ends with a directory separator and `file{}` otherwise.

For the `dir{}` target type the trailing directory separator is added to the pattern and all the inclusion/exclusion patterns/matches that do not already end with one. Then the filesystem search is performed for matching directories. For example:

```
./: dir{* -build} # Search for */, exclude build/.
```

For the `file{}` and `file{}`-based target types the default extension (if any) is added to the pattern and all the inclusion/exclusion patterns/matches that do not already contain an extension. Then the filesystem search is performed for matching files.

For example, the `cxx{}` target type obtains the default extension from the `extension` variable (see Target Types for background). Assuming we have the following line in our `root.build`:

```
cxx{*}: extension = cxx
```

And the following in our `buildfile`:

```
exe{hello}: {cxx}{* -foo -bar.cxx}
```

The pattern match will first search for all the files matching the `*.cxx` pattern in `src_base` and then exclude `foo.cxx` and `bar.cxx` from the result. Note also that target type-specific decorations are removed from the result. So in the above example if the pattern match produces `baz.cxx`, then the prerequisite name is `cxx{baz}`, not `cxx{baz.cxx}`.

If the name generation cannot be performed because the base directory is unknown, target type is unknown, or the target type is not directory or file-based, then the name pattern is returned as is (that is, as an ordinary name). Project-qualified names are never considered to be patterns.

9 config Module

This chapter is a work in progress and is incomplete.

9.1 Hermetic Build Configurations

Hermetic build configurations save environment variables that affect the project along with other project configuration in the `build/config.build` file. These saved environment variables are then used instead of the current environment when performing operations on the project, thus making sure the project "sees" exactly the same environment as during configuration.

While currently hermetic configurations only deal with the environment, in the future this functionality may be extended to also support disallowing changes to external resources (compilers, system headers and libraries, etc).

To create a hermetic configuration we use the `config.config.hermetic` configuration variable. For example:

```
$ b configure config.config.hermetic=true
```

Hermetic configurations are not the default because they are not without drawbacks. Firstly, a hermetic configuration may break if the saved environment becomes incompatible with the rest of the system. For example, you may re-install an external program (say, a compiler) into a different location and update your `PATH` to match the new setup. However, a hermetic configuration will "see" the first change but not the second.

Another issue is the commands printed during a hermetic build: they are executed in the saved environment which may not match the environment in which the build system was invoked. As a result, we cannot easily re-execute such commands, which is often handy during build troubleshooting.

It is also important to keep in mind that a non-hermetic build configuration does not break or produce incorrect results if the environment changes. Instead, changes to the environment are detected and affected targets are automatically rebuilt.

The two use-cases where hermetic configurations are especially useful are when we need to save an environment which is not generally available (for example, an environment of a Visual Studio development command prompt) or when our build results need to exactly match the specific configuration (for example, because parts of the overall result have already been built and installed, as is the case with build system modules).

If we now examine `config.build`, we will see something along these lines:

```
$ cat build/config.build

config.config.hermetic = true
config.config.environment = CPATH CPLUS_INCLUDE_PATH PATH=...
```

Hermetic configuration support is built on top of the low-level `config.config.environment` configuration variable which allows us to specify custom environment variables and their values. Specifically, it contains a list of environment variable "sets" (*name=value*) and "unsets" (*name*). For example:

```
$ b configure \
  config.config.environment="PATH=/bin:/usr/bin LD_LIBRARY_PATH"
```

Specifying `config.config.hermetic=true` simply instructs the `config` module to collect and save in `config.config.environment` environment variables that affect the project. These include:

- built-in variables (such as `PATH` and `LD_LIBRARY_PATH` or equivalent),
- variables that affect external programs as reported by build system modules (such as `CPLUS_INCLUDE_PATH` reported by the `cxx` module) or by imported programs via metadata,
- variables reported by the project itself with the `config.environment` directive

(discussed below).

Reconfiguring a hermetic configuration preserves the saved environment unless *re-hermetization* is explicitly requested with the `config.config.hermetic.reload` configuration variable. For example:

```
$ b configure config.config.hermetic.reload=true
```

Note that `config.config.hermetic.reload` is transient and is not stored in `config.build`. In other words, there is no way to create a hermetic configuration that is re-hermetized by default during reconfiguration.

To *de-hermetize* a hermetic build configuration, reconfigure it with `config.config.hermetic=false`.

The `config.config.hermetic` variable has essentially a tri-state value: `true` means keep hermetized (save the environment in `config.config.environment`), `false` means keep de-hermetized (clear `config.config.environment`) and `null` or `undefined` means don't touch `config.config.environment`.

We can adjust the set of environment variables saved in a hermetic configuration using the `config.config.hermetic.environment` configuration variable. It contains a list of inclusions (*name*) and exclusions (*name@false*) which are applied to the final set of environment variables that affect the project. For example:

```
LC_ALL=C b configure \
  config.config.hermetic=true \
  config.config.hermetic.environment="LC_ALL PATH@false"
```

Typically, the set of environment variables that affect the project is discovered automatically. Specifically, modules that we use (such as `cxx`) are expected to report the environment variables that affect the programs they invoke (such as the C++ compiler). Similarly, programs that we import in our `buildfiles` (for example to use in ad hoc recipes) are expected to report environment variables that affect them as part of their metadata.

However, there are situations where we need to report an environment variable manually. These include calling the `$getenv()` function from a `buildfile` or invoking a program (either in an ad hoc recipe, the `run` directive, or the `$run*()` function family) that either does not provide the metadata or does not report the environment as part of it. In such cases we should report the environment variable manually using the `config.environment` directive. For example:

```
config.environment USE_FOO

foo = $getenv(USE_FOO)

if ($foo != [null])
  cxx.poptions += "-DUSE_FOO=$foo"
```

Additionally, if invoking a program in an ad hoc recipe that either does not provide the meta-data or does not report the environment as part of it, then we additionally should track the changes to the relevant environment variables manually using the `depdb env` builtin. For example:

```
import! foo = foo%exe{foo} # Uses FOO and BAR environment variables.

config.environment FOO BAR

file{output}: file{input} $foo
{{
  diag foo $>
  depdb env FOO BAR
  $foo $path($<[0]) >$path($>)
}}
```

Normally, we would want to report variables that affect the build result rather than build byproducts (for example, diagnostics). This is, for example, the reason why locale-related environment variables are not saved by default. Also, sometime environment variables only affect certain modes of a program. If such modes are not used, then there is no need to report the corresponding variables.

10 test Module

This chapter is a work in progress and is incomplete.

The targets to be tested as well as the tests/groups from testscripts to be run can be narrowed down using the `config.test` variable. While this value is normally specified as a command line override (for example, to quickly re-run a previously failed test), it can also be persisted in `config.build` in order to create a configuration that will only run a subset of tests by default. For example:

```
$ b test config.test=foo/exe{driver} # Only test foo/exe{driver} target.
$ b test config.test=bar/baz          # Only run bar/baz testscript test.
```

The `config.test` variable contains a list of @-separated pairs with the left hand side being the target and the right hand side being the testscript id path. Either can be omitted (along with @). If the value contains a target type or ends with a directory separator, then it is treated as a target name. Otherwise – an id path. The targets are resolved relative to the root scope where the `config.test` value is set. For example:

```
$ b test config.test=foo/exe{driver}@bar
```

To specify multiple id paths for the same target we can use the pair generation syntax:

```
$ b test config.test=foo/exe{driver}@{bar baz}
```

If no targets are specified (only id paths), then all the targets are tested (with the testscript tests to be run limited to the specified id paths). If no id paths are specified (only targets), then all the testscript tests are run (with the targets to be tested limited to the specified targets). An id path without a target applies to all the targets being considered.

A directory target without an explicit target type (for example, `foo/`) is treated specially. It enables all the tests at and under its directory. This special treatment can be inhibited by specifying the target type explicitly (for example, `dir{foo/}`).

The test execution time can be limited using the `config.test.timeout` variable. Its value has the `<operation-timeout>/<test-timeout>` form where the timeouts are specified in seconds and either of them (but not both) can be omitted. The left hand side sets the timeout for the whole `test` operation and the right hand side – for individual tests. The zero value clears the previously set timeout. For example:

```
$ b test config.test.timeout=20    # Test operation.
$ b test config.test.timeout=20/5 # Test operation and individual tests.
$ b test config.test.timeout=/5    # Individual tests.
```

The test timeout can be specified on multiple nested root scopes. For example, we can specify a greater timeout for the entire build configuration and lesser ones for individual projects. The tests must complete before the nearest of the enclosing scope timeouts. Failed that, the timed out tests are terminated forcibly causing the entire `test` operation to fail. See also the `timeout` builtin for specifying timeouts from within the tests and test groups.

The programs being tested can be executed via a *runner program* by specifying the `config.test.runner` variable. Its value has the `<path> [<options>]` form. For example:

```
$ b test config.test.runner="valgrind -q"
```

When the runner program is specified, commands of simple and Testscript tests are automatically adjusted so that the runner program is executed instead, with the test command passed to it as arguments. For ad hoc test recipes, the runner program has to be handled explicitly. Specifically, if `config.test.runner` is specified, the `test.runner.path` and `test.runner.options` variables contain the runner program path and options, respectively, and are set to `null` otherwise. These variables can be used by ad hoc recipes to detect the presence of the runner program and, if so, arrange appropriate execution of desired commands. For example:

```
exe{hello}:
% test
{{
  diag test $>

  cmd = ($test.runner.path == [null] \
    ? $> \
    : $test.runner.path $test.runner.options $path($>))

  $cmd 'World' >>>'Hello, World!'
}}
```

11 install Module

This chapter is a work in progress and is incomplete.

The `install` module provides support for installing and uninstalling projects.

As briefly discussed in the Installing section of the Introduction, the `install` module defines the following standard installation locations:

name	default	config.install.* (c.i.*) override
----	-----	-----
root		c.i.root
data_root	root/	c.i.data_root
exec_root	root/	c.i.exec_root
bin	exec_root/bin/	c.i.bin
sbin	exec_root/sbin/	c.i.sbin
lib	exec_root/lib/<private>/	c.i.lib
libexec	exec_root/libexec/<private>/<project>/	c.i.libexec
pkgconfig	lib/pkgconfig/	c.i.pkgconfig
etc	data_root/etc/	c.i.etc
include	data_root/include/<private>/	c.i.include
include_arch	include/	c.i.include_arch
share	data_root/share/	c.i.share
data	share/<private>/<project>/	c.i.data
buildfile	share/build2/export/<project>/	c.i.buildfile
doc	share/doc/<private>/<project>/	c.i.doc
legal	doc/	c.i.legal
man	share/man/	c.i.man
man<N>	man/man<N>/	c.i.man<N>

The `include_arch` location is meant for architecture-specific files, such as configuration headers. By default it's the same as `include` but can be configured by the user to a different value (for example, `/usr/include/x86_64-linux-gnu/`) for platforms that support multiple architectures from the same installation location. This is how one would normally use it from a buildfile:

```
# The configuration header may contain target architecture-specific
# information so install it into include_arch/ instead of include/.
#
h{*}:      install = include/libhello/
h{config}: install = include_arch/libhello/
```

The `buildfile` location is meant for exported buildfiles that can be imported by other projects. If a project contains any `**build` buildfiles in its `build/export/` directory (or `**build2` and `build2/export/` in the alternative naming scheme), then they are automatically installed into this location (recreating subdirectories).

The `<project>`, `<version>`, and `<private>` substitutions in these `config.install.*` values are replaced with the project name, version, and private subdirectory, respectively. If either is empty, then the corresponding directory component is ignored.

The optional private installation subdirectory (`<private>`) mechanism can be used to hide the implementation details of a project. This is primarily useful when installing an executable that depends on a bunch of libraries into a shared location, such as `/usr/local/`. By hiding the libraries in the private subdirectory we can make sure that they will not interfere with anything that is already installed into such a shared location by the user and that any further such installations won't interfere with our executable.

The private installation subdirectory is specified with the `config.install.private` variable. Its value must be a relative directory and may include multiple components. For example:

```
$ b install \
  config.install.root=/usr/local/ \
  config.install.private=hello/
```

If you are relying on your system's dynamic linker defaults to automatically find shared libraries that are installed with your executable, then adding the private installation subdirectory will most definitely cause this to stop working. The recommended way to resolve this problem is to use *rpath*, for example:

```
$ b install \
  config.install.root=/usr/local/ \
  config.install.private=hello/ \
  config.bin.rpath=/usr/local/lib/hello/
```

11.1 Relocatable Installation

A relocatable installation can be moved to a directory other than its original installation location. Note that the installation should be moved as a whole preserving the directory structure under its root (`config.install.root`). To request a relocatable installation, set the `config.install.relocatable` variable to `true`. For example:

```
$ b install \
  config.install.root=/tmp/install \
  config.install.relocatable=true
```

A relocatable installation is achieved by using paths relative to one filesystem entry within the installation to locate another. Some examples include:

- Paths specified in `config.bin.rpath` are made relative using the `$ORIGIN` (Linux, BSD) or `@loader_path` (Mac OS) mechanisms.
- Paths in the generated `pkg-config` files are made relative to the `${pcfiledir}` built-in variable.
- Paths in the generated installation manifest (`config.install.manifest`) are made relative to the location of the manifest file.

While these common aspects are handled automatically, if a project relies on knowing its installation location, then it will most likely need to add manual support for relocatable installations.

As an example, consider an executable that supports loading plugins and requires the plugin installation directory to be embedded into the executable during the build. The common way to support relocatable installations for such cases is to embed a path relative to the executable and complete it at runtime, normally by resolving the executable's path and using its directory as a base.

If you would like to always use the relative path, regardless of whether the installation is relocatable or not, then you can obtain the library installation directory relative to the executable installation directory like this:

```
plugin_dir = $install.resolve($install.lib, $install.bin)
```

Alternatively, if you would like to continue using absolute paths for non-relocatable installations, then you can use something like this:

```
plugin_dir = $install.resolve( \
    $install.lib,                \
    ($install.relocatable ? $install.bin : [dir_path] ))
```

Finally, if you are unable to support relocatable installations, the correct way to handle this is to assert this fact in `root.build` of your project, for example:

```
assert (!$install.relocatable) 'relocatable installation not supported'
```

11.2 Installation Filtering

While project authors determine what gets installed at the `buildfile` level, the users of the project can further filter the installation using the `config.install.filter` variable.

The value of this variable is a list of key-value pairs that specify the filesystem entries to include or exclude from the installation. For example, the following filters will omit installing headers and static libraries (notice the quoting of the wildcard).

```
$ b install config.install.filter='include/@false "*.a"@false'
```

The key in each pair is a file or directory path or a path wildcard pattern. If a key is relative and contains a directory component or is a directory, then it is treated relative to the corresponding `config.install.*` location. Otherwise (simple path, normally a pattern), it is matched against the leaf of any path. Note that if an absolute path is specified, it should be without the `config.install.chroot` prefix.

The value in each pair is either `true` (include) or `false` (exclude). The filters are evaluated in the order specified and the first match that is found determines the outcome. If no match is found, the default is to include. For a directory, while `false` means exclude all the sub-paths inside this directory, `true` does not mean that all the sub-paths will be included wholesale.

Rather, the matched component of the sub-path is treated as included with the rest of the components matched against the following sub-filters. For example:

```
$ b install config.install.filter='
  include/x86_64-linux-gnu/@true
  include/x86_64-linux-gnu/details/@false
  include/@false'
```

The `true` or `false` value may be followed by comma and the `symlink` modifier to only apply to symlink filesystem entries. For example:

```
$ b config.install.filter=' "*.so"@false,symlink'
```

A filter can be negated by specifying `!` as the first pair. For example:

```
$ b install config.install.filter='! include/@false "*.a"@false'
```

Note that the filtering mechanism only affects what gets physically copied to the installation directory without affecting what gets built for install or the view of what gets installed at the buildfile level. For example, given the `include/@false *.a@false` filters, static libraries will still be built (unless arranged not to with `config.bin.lib`) and the `pkg-config` files will still end up with `-I` options pointing to the header installation directory. Note also that this mechanism applies to both `install` and `uninstall` operations.

If you are familiar with the Debian or Fedora packaging, this mechanism is somewhat similar to (and can be used for a similar purpose as) the Debian's `.install` files and Fedora's `%files` spec file sections, which are used to split the installation into multiple binary packages.

As another example, the following filters will omit all the development-related files (headers, `pkg-config` files, static libraries, and shared library symlinks; assuming the platform uses the `.a/.so` extensions for the libraries):

```
$ b install config.install.filter='
  include/@false
  pkgconfig/@false
  "lib/*.a"@false
  "lib/*.so"@false,symlink'
```

12 version Module

A project can use any version format as long as it meets the package version requirements. The toolchain also provides additional functionality for managing projects that conform to the `build2 standard version` format. If you are starting a new project that uses `build2`, you are strongly encouraged to use this versioning scheme. It is based on much thought and, often painful, experience. If you decide not to follow this advice, you are essentially on your own where version management is concerned.

The standard `build2` project version conforms to Semantic Versioning and has the following form:

```
<major>.<minor>.<patch>[-<prerelease>]
```

For example:

```
1.2.3
1.2.3-a.1
1.2.3-b.2
```

The `build2` package version that uses the standard project version will then have the following form (*epoch* is the versioning scheme version and *revision* is the package revision):

```
[+<epoch>-]<major>.<minor>.<patch>[-<prerelease>][+<revision>]
```

For example:

```
1.2.3
1.2.3+1
+2-1.2.3-a.1+2
```

The *major*, *minor*, and *patch* should be numeric values between 0 and 99999 and all three cannot be zero at the same time. For initial development it is recommended to use 0 for *major*, start with version 0.1.0, and change to 1.0.0 once things stabilize.

In the context of C and C++ (or other compiled languages), you should increment *patch* when making binary-compatible changes, *minor* when making source-compatible changes, and *major* when making breaking changes. While the binary compatibility must be set in stone, the source compatibility rules can sometimes be bent. For example, you may decide to make a breaking change in a rarely used interface as part of a minor release (though this is probably still a bad idea if your library is widely depended upon). Note also that in the context of C++ deciding whether a change is binary-compatible is a non-trivial task. There are resources that list the rules but no automated tooling yet. If unsure, increment *minor*.

If present, the *prerelease* component signifies a pre-release. Two types of pre-releases are supported by the standard versioning scheme: *final* and *snapshot* (non-pre-release versions are naturally always final). For final pre-releases the *prerelease* component has the following form:

```
(a|b).<num>
```

For example:

```
1.2.3-a.1
1.2.3-b.2
```

The letter 'a' signifies an alpha release and 'b' – beta. The alpha/beta numbers (*num*) should be between 1 and 499.

Note that there is no support for release candidates. Instead, it is recommended that you use later-stage beta releases for this purpose (and, if you wish, call them "release candidates" in announcements, etc).

What version should be used during development? The common approach is to increment to the next version and use that until the release. This has one major drawback: if we publish intermediate snapshots (for example, for testing) they will all be indistinguishable both between each other and, even worse, from the final release. One way to remedy this is to increment the pre-release number before each publication. However, unless automated, this will be burdensome and error-prone. Also, there is a real possibility of running out of version numbers if, for example, we do continuous integration by publishing and testing each commit.

To address this, the standard versioning scheme supports *snapshot pre-releases* with the *prerel* component having the following extended form:

```
(a|b).<num>.<snapsn>[.<snapid>]
```

For example:

```
1.2.3-a.1.20180319215815.26efe301f4a7
```

In essence, a snapshot pre-release is after the previous final release but before the next (a.1 and, perhaps, a.2 in the above example) and is uniquely identified by the snapshot sequence number (*snapsn*) and optional snapshot id (*snapid*).

The *num* component has the same semantics as in the final pre-releases except that it can be 0. The *snapsn* component should be either the special value 'z' or a numeric, non-zero value that increases for each subsequent snapshot. It must not be longer than 16 decimal digits. The *snapid* component, if present, should be an alpha-numeric value that uniquely identifies the snapshot. It is not required for version comparison (*snapsn* should be sufficient) and is included for reference. It must not be longer than 16 characters.

Where do the snapshot number and id come from? Normally from the version control system. For example, for *git*, *snapsn* is the commit date in the *YYYYMMDDhhmmss* form and UTC timezone and *snapid* is a 12-character abbreviated commit id. As discussed below, the *build2 version* module extracts and manages all this information automatically (but the use of *git* commit dates is not without limitations; see below for details).

The special 'z' *snapsn* value identifies the *latest* or *uncommitted* snapshot. It is chosen to be greater than any other possible *snapsn* value and its use is discussed further below.

As an illustration of this approach, let's examine how versions change during the lifetime of a project:

```
0.1.0-a.0.z      # development after a.0
0.1.0-a.1        # pre-release
0.1.0-a.1.z      # development after a.1
0.1.0-a.2        # pre-release
0.1.0-a.2.z      # development after a.2
0.1.0-b.1        # pre-release
```

```

0.1.0-b.1.z      # development after b.1
0.1.0            # release
0.1.1-b.0.z      # development after b.0 (bugfix)
0.2.0-a.0.z      # development after a.0
0.1.1            # release (bugfix)
1.0.0            # release (jumped straight to 1.0.0)
...

```

As shown in the above example, there is nothing wrong with "jumping" to a further version (for example, from alpha to beta, or from beta to release, or even from alpha to release). We cannot, however, jump backwards (for example, from beta back to alpha). As a result, a sensible strategy is to start with `a.0` since it can always be upgraded (but not downgraded) at a later stage.

When it comes to the version control systems, the recommended workflow is as follows: The change to the final version should be the last commit in the (pre-)release. It is also a good idea to tag this commit with the project version. A commit immediately after that should change the version to a snapshot, "opening" the repository for development.

The project version without the snapshot part can be represented as a 64-bit decimal value comparable as integers (for example, in preprocessor directives). The integer representation has the following form:

```
AAAAABBBBBCCCCDDDE
```

```

AAAAA - major
BBBBB - minor
CCCCC - patch
DDD    - alpha / beta (DDD + 500)
E      - final (0) / snapshot (1)

```

If the *DDDE* value is not zero, then it signifies a pre-release. In this case one is subtracted from the *AAAAABBBBBCCCCC* value. An alpha number is stored in *DDD* as is while beta – incremented by 500. If *E* is 1, then this is a snapshot after *DDD*.

For example:

```

                                AAAAABBBBBCCCCDDDE
0.1.0                          0000000001000000000
0.1.2                          0000000001000020000
1.2.3                          0000100002000030000
2.2.0-a.1                      0000200001999990010
3.0.0-b.2                      0000299999999995020
2.2.0-a.1.z                    0000200001999990011

```

A project that uses standard versioning can rely on the `build2 version` module to simplify and automate version managements. The `version` module has two primary functions: eliminate the need to change the version anywhere except in the project's manifest file and automatically extract and propagate the snapshot information (sequence number and id).

The `version` module must be loaded in the project's `bootstrap.build`. While being loaded, it reads the project's manifest and extracts its version (which must be in the standard form).

Another function of the `version` module is to check the `build2` version requirement that is customarily specified in the manifest. This check is also the reason why we normally load the `version` module in subprojects, such as `tests/`, which don't have their own manifest file. In this case the `version` module loads the amalgamating project's manifest.

The extracted version is parsed and presented as the following build system variables (which can be used in the buildfiles):

```
[string] version                # +2-1.2.3-b.4.1234567.deadbeef+3

[string] version.project        # 1.2.3-b.4.1234567.deadbeef
[uint64] version.project_number # 0000100002000025041
[string] version.project_id     # 1.2.3-b.4.deadbeef

[bool]   version.stub          # false (true for 0[+<revision>])

[uint64] version.epoch         # 2

[uint64] version.major         # 1
[uint64] version.minor         # 2
[uint64] version.patch         # 3

[bool]   version.alpha         # false
[bool]   version.beta          # true
[bool]   version.pre_release   # true
[string] version.pre_release_string # b.4
[uint64] version.pre_release_number # 4

[bool]   version.snapshot      # true
[uint64] version.snapshot_sn    # 1234567
[string] version.snapshot_id    # deadbeef
[string] version.snapshot_string # 1234567.deadbeef
[bool]   version.snapshot_committed # true

[uint64] version.revision       # 3
```

As a convenience, the `version` module also extracts the `summary` and `url` manifest values and sets them as the following build system variables (this additional information is used, for example, when generating the `pkg-config` files):

```
[string] project.summary
[string] project.url
```

If the version is the latest snapshot (that is, it's in the `.z` form), then the `version` module extracts the snapshot information from the version control system used by the project. Currently only `git` is supported with the following semantics.

If the project's source directory (`src_root`) is clean (that is, it does not have any changed or untracked files), then the `HEAD` commit date and id are used as the snapshot number and id, respectively.

Otherwise (that is, the project is between commits), the `HEAD` commit date is incremented by one second and is used as the snapshot number with no id. While we can work with such uncommitted snapshots locally, we should not distribute or publish them since they are indistinguishable from each other.

Finally, if the project does not have HEAD (that is, the project has no commits yet), the special 19700101000000 (UNIX epoch) commit date is used.

The use of `git` commit dates for snapshot ordering has its limitations: they have one second resolution which means it is possible to create two commits with the same date (but not the same commit id and thus snapshot id). We also need all the committers to have a reasonably accurate clock. Note, however, that in case of a commit date clash/ordering issue, we still end up with distinct versions (because of the commit id) – they are just not ordered correctly. As a result, we feel that the risks are justified when the only alternative is manual version management (which is always an option, nevertheless).

When we prepare a source distribution of a snapshot, the `version` module automatically adjusts the package name to include the snapshot information as well as patches the manifest file in the distribution with the snapshot number and id (that is, replacing `.z` in the version value with the actual snapshot information). The result is a package that is specific to this commit.

Besides extracting the version information and making it available as individual components, the `version` module also provides rules for installing the manifest file as well as automatically generating version headers (or other similar version-based files).

By default the project's manifest file is installed as documentation, just like other `doc{ }` targets (thus replacing the `version` file customarily shipped in the project root directory). The manifest installation rule in the `version` module in addition patches the installed manifest file with the actual snapshot number and id, just like during the preparation of distributions.

The version header rule is based on the `in` module rule and can be used to preprocess a template file with version information. While it is usually used to generate C/C++ version headers (thus the name), it can really generate any kind of files.

The rule matches a file-based target that has the corresponding `in` prerequisite and also depends on the project's manifest file. As an example, let's assume we want to auto-generate a header called `version.hxx` for our `libhello` library. To accomplish this we add the `version.hxx.in` template as well as something along these lines to our buildfile:

```
lib{hello}: {hxx cxx}{** -version} hxx{version}

hxx{version}: in{version} $src_root/file{manifest}
```

The header rule is a line-based preprocessor that substitutes fragments enclosed between (and including) a pair of dollar signs (\$) with \$\$ being the escape sequence (see the `in` module for details). As an example, let's assume our `version.hxx.in` contains the following lines:

```
#ifndef LIBHELLO_VERSION

#define LIBHELLO_VERSION      $libhello.version.project_number$ULL
#define LIBHELLO_VERSION_STR "$libhello.version.project$"

#endif
```

If our `libhello` is at version `1.2.3`, then the generated `version.hxx` will look like this:

```
#ifndef LIBHELLO_VERSION

#define LIBHELLO_VERSION      100002000030000ULL
#define LIBHELLO_VERSION_STR "1.2.3"

#endif
```

The first component after the opening `$` should be either the name of the project itself (like `libhello` above) or a name of one of its dependencies as listed in the manifest. If it is the project itself, then the rest can refer to one of the `version.*` variables that we discussed earlier (in reality it can be any variable visible from the project's root scope).

If the name refers to one of the dependencies (that is, projects listed with `depends:` in the manifest), then the following special substitutions are recognized:

```
$<name>.version$           - textual version constraint
$<name>.condition(<VERSION>[,<SNAPSHOT>])$ - numeric satisfaction condition
$<name>.check(<VERSION>[,<SNAPSHOT>])$    - numeric satisfaction check
```

Here *VERSION* is the version number macro and the optional *SNAPSHOT* is the snapshot number macro. The snapshot is only required if you plan to include snapshot information in your dependency constraints.

As an example, let's assume our `libhello` depends on `libprint` which is reflected with the following line in our manifest:

```
depends: libprint >= 2.3.4
```

We also assume that `libprint` provides its version information in the `libprint/version.hxx` header and uses analogous-named macros. Here is how we can add a version check to our `version.hxx.in`:

```
#ifndef LIBHELLO_VERSION

#define LIBHELLO_VERSION      $libhello.version.project_number$ULL
#define LIBHELLO_VERSION_STR "$libhello.version.project$"

#include <libprint/version.hxx>

$libprint.check(LIBPRINT_VERSION)$

#endif
```

After the substitution our `version.hxx` header will look like this:

```
#ifndef LIBHELLO_VERSION

#define LIBHELLO_VERSION      100002000030000ULL
#define LIBHELLO_VERSION_STR "1.2.3"

#include <libprint/version.hxx>
```

```
#ifdef LIBPRINT_VERSION
# if !(LIBPRINT_VERSION >= 200003000040000ULL)
#   error incompatible libprint version, libprint >= 2.3.4 is required
# endif
#endif

#endif
```

The `version` and `condition` substitutions are the building blocks of the check substitution. For example, here is how we can implement a check with a customized error message:

```
#if !($libprint.condition(LIBPRINT_VERSION)$)
# error bad libprint, need libprint $libprint.version$
#endif
```

The `version` module also treats one dependency in a special way: if you specify the required version of the build system in your manifest, then the module will automatically check it for you. For example, if we have the following line in our manifest:

```
depends: * build2 >= 0.5.0
```

And someone tries to build our project with `build2 0.4.0`, then they will see an error like this:

```
build/bootstrap.build:3:1: error: incompatible build2 version
    info: running 0.4.0
    info: required 0.5.0
```

What version constraints should be used when depending on another project? We start with a simple case where we depend on a release. Let's say `libprint 2.3.0` added a feature that we need in our `libhello`. If `libprint` follows the source/binary compatibility guidelines discussed above, then any `2.X.Y` version should work provided `X >= 3`. And this how we can specify it in the manifest:

```
depends: libprint ^2.3.0
```

Let's say we are now working on `libhello 2.0.0` and would like to start using features from `libprint 3.0.0`. However, currently, only pre-releases of `3.0.0` are available. If you would like to add a dependency on a pre-release (most likely from your own pre-release), then the recommendation is to only allow a specific version, essentially "expiring" the combination as soon as newer versions become available. For example:

```
version: 2.0.0-b.1
depends: libprint == 3.0.0-b.2
```

Finally, let's assume we are feeling adventurous and would like to test development snapshots of `libprint` (most likely from our own snapshots). In this case the recommendation is to only allow a snapshot range for a specific pre-release with the understanding and a warning that no compatibility between snapshot versions is guaranteed. For example:


```
version: 2.0.0-b.1.z
depends: libprint [3.0.0-b.2.1 3.0.0-b.3)
```

13 bin Module

This chapter is a work in progress and is incomplete.

13.1 Binary Target Types

The following listing shows the hierarchy of the target types defined by the `bin` module while the following sections describe each target type in detail (`target{}` and `file{}` are standard target types defined by the `build2` core; see Target Types for details).



13.1.1 `lib{}`, `liba{}`, `libs{}`

The `liba{}` and `libs{}` target types represent static (archive) and shared libraries, respectively.

The `lib{}` target type is a group with the `liba{}` and/or `libs{}` members. A rule that encounters a `lib{}` prerequisite may pick a member appropriate for the target being built or it may build all the members according to the `bin.lib` variable. See Library Exportation and Versioning for background.

The `lib*{}` file extensions are normally automatically assigned by the matching rules based on the target platform.

13.1.2 `libul{}`, `libue{}`, `libua{}`, `libus{}`

The `libu*{}` target types represent utility libraries. Utility libraries are static libraries with object files appropriate for linking an executable (`libue{}`), static library (`libua{}`), or shared library (`libus{}`). Where possible, utility libraries are built in the "thin archive" mode.

The `libul{}` target type is a group with the `libua{}` and/or `libus{}` members. A rule that encounters a `libul{}` prerequisite picks a member appropriate for the target being built.

The `libu*{}` file extensions are normally automatically assigned by the matching rules based on the target platform.

13.1.3 obj{ }, obje{ }, obja{ }, objs{ }

The `obj*{ }` target types represent object files appropriate for linking an executable (`obje{ }`), static library (`obja{ }`), or shared library (`objs{ }`).

In `build2` we use distinct object files for the three types of binaries (executable, static library, and shared library). The distinction between static and shared libraries is made to accommodate build differences such as the need for position-independent code (`-fPIC`) in shared libraries. While in most cases the same object file can be used for executables and static libraries, they are kept separate for consistency and generality.

The `obj{ }` target type is a group with the `obje{ }`, and/or `obja{ }`, and/or `objs{ }` members. A rule that encounters an `obj{ }` prerequisite picks a member appropriate for the target being built.

The `obj*{ }` file extensions are normally automatically assigned by the matching rules based on the target platform.

13.1.4 bmi{ }, bmie{ }, bmia{ }, bmis{ }

The `bmi*{ }` target types represent binary module interfaces (BMI) for C++20 named modules appropriate for linking an executable (`bmie{ }`), static library (`bmia{ }`), or shared library (`bmis{ }`).

The `bmi{ }` target type is a group with the `bmie{ }`, and/or `bmia{ }`, and/or `bmis{ }` members. A rule that encounters an `bmi{ }` prerequisite picks a member appropriate for the target being built.

The `bmi*{ }` file extensions are normally automatically assigned by the matching rules based on the target platform.

13.1.5 hbmi{ }, hbmie{ }, hbmia{ }, hbmis{ }

The `hbmi*{ }` target types represent binary module interfaces (BMI) for C++20 header units appropriate for linking an executable (`hbmie{ }`), static library (`hbmia{ }`), or shared library (`hbmis{ }`).

The `hbmi{ }` target type is a group with the `hbmie{ }`, and/or `hbmia{ }`, and/or `hbmis{ }` members. A rule that encounters an `hbmi{ }` prerequisite picks a member appropriate for the target being built.

The `hbmi*{ }` file extensions are normally automatically assigned by the matching rules based on the target platform.

13.1.6 `def{ }`

The `def{ }` target type represents Windows module definition files and has the fixed default extension `.def`.

14 cc Module

This chapter is a work in progress and is incomplete.

This chapter describes the `cc` build system module which provides the common compilation and linking support for C-family languages.

14.1 C-Common Configuration Variables

```

config.c
config.cxx
  cc.id

  cc.target
  cc.target.cpu
  cc.target.vendor
  cc.target.system
  cc.target.version
  cc.target.class

config.cc.poptions
  cc.poptions

config.cc.coptions
  cc.coptions

config.cc.loptions
  cc.loptions

config.cc.aoptions
  cc.aoptions

config.cc.libs
  cc.libs

config.cc.internal.scope
  cc.internal.scope

config.cc.reprocess
  cc.reprocess

config.cc.pkgconfig.sysroot

config.cc.compiledb
config.cc.compiledb.name
config.cc.compiledb.filter
config.cc.compiledb.filter.input
config.cc.compiledb.filter.output

```

Note that the compiler mode options are "cross-hinted" between `config.c` and `config.cxx` meaning that if we specify one but not the other, mode options, if any, will be added to the absent. This may or may not be the desired behavior, for example:

```
# Ok: config.c="gcc -m32"
$ b config.cxx="g++ -m32"

# Not OK: config.c="clang -stdlib=libc++"
$ b config.cxx="clang++ -stdlib=libc++"
```

In addition to the `obj{}/bmi{}` targets, the compile option (`*.poptions` and `*.coptions`) can be specified on the `exe{}/lib{}` targets. We refer to such options as "binary-specific compile options". They are propagated to `obj{}/bmi{}` targets that are synthesized for source prerequisites of the binary. Note that this propagation does not apply to explicit (non-synthesized) `obj{}/bmi{}` prerequisites. For example:

```
exe{foo}: cxx{foo} obj{common}
{
    cc.poptions += -DFOO
}

exe{bar}: cxx{bar} obj{common}
{
    cc.poptions += -DBAR
}

obj{common}: cxx{common}
{
    cc.poptions += -DCOMMON
}
```

In this example, `cxx{foo}` will be compiled with `-DFOO`, `cxx{bar}` – with `-DBAR`, and `cxx{common}` – with `-DCOMMON`.

Note that if a source prerequisite is shared between several binaries, then the values of the propagated compile options (or their absence) must match. For instance, the following variant of the above example would result in an error because `cxx{common}` would have contradictory `cc.poptions` values:

```
exe{foo}: cxx{foo common}
{
    cc.poptions += -DFOO
}

exe{bar}: cxx{bar common}
{
    cc.poptions += -DBAR
}
```

As another example, here is how we can rewrite a typical C++ library `buildfile` (which requires different macros to distinguish between shared/static builds) using this mechanism, in this case, to build two libraries in the same scope:

```

./: lib{foo}: {hxx cxx}{*-foo}
./: lib{bar}: {hxx cxx}{*-bar}

cxx.poptions += "-I$out_root" "-I$src_root"

lib{foo}:
{
  cxx.poptions += -DFOO
  cxx.export.poptions = "-I$out_root" "-I$src_root"
}

liba{foo}:
{
  cxx.poptions += -DLIBFOO_STATIC_BUILD
  cxx.export.poptions += -DLIBFOO_STATIC
}

libs{foo}:
{
  cxx.poptions += -DLIBFOO_SHARED_BUILD
  cxx.export.poptions += -DLIBFOO_SHARED
}

lib{bar}:
{
  cxx.poptions += -DBAR
  cxx.export.poptions = "-I$out_root" "-I$src_root"
}

liba{bar}:
{
  cxx.poptions += -DLIBBAR_STATIC_BUILD
  cxx.export.poptions += -DLIBBAR_STATIC
}

libs{bar}:
{
  cxx.poptions += -DLIBBAR_SHARED_BUILD
  cxx.export.poptions += -DLIBBAR_SHARED
}

```

The exact semantics of this mechanism is as-if the binary-specific compile options were set on the synthesized `obj{ }/bmi{ }` targets at the end of the `buildfile`.

One nuance to keep in mind is that target type/pattern-specific assign/append/prepend specified for `obj{ }/bmi{ }` will not be in effect for options specified on `lib{ }/exe{ }`. For example:

```

cc.poptions += -DSCOPE
obj{*}: cc.poptions += -DTARGET
exe{foo}: cc.poptions += -DFOO

```

Here the effective `cc.poptions` for `exe{foo}` prerequisites will be `-DSCOPE -DFOO` since `exe{foo}` does not match the `obj{ * }` pattern. As a result, if using this mechanism, remember to include the binary target types in `obj{ }` patterns. For example:

```
{obj exe}{*}: cxx.poptions += -DTARGET
```

14.2 C-Common Target Types

The following listing shows the hierarchy of the target types defined by the `cc` module while the following sections describe each target type in detail (`file{}` is a standard target type defined by the `build2` core; see Target Types for details). Every `cc`-based module (such as `c` and `cxx`) will have these common target types defined in addition to the language-specific ones.

```

.--file--
|         |
h         pc
         |
         pca
         pcs

```

While the `h{}` target type represents a C header file, there is hardly a C-family compilation without a C header inclusion. As a result, this target type is defined by all `cc`-based modules.

For the description of the `h{}` target type refer to `c{}`, `h{}` in the `C` module documentation.

14.2.1 `pc{}`, `pca{}`, `pcs{}`

The `pc*{}` target types represent `pkg-config` files. The `pc{}` target type represents the common file and has the fixed default extension `.pc`. The `pca{}` and `pcs{}` target types represent the static and shared files and have the fixed default extensions `.static.pc` and `.shared.pc`, respectively. See Importation of Installed Libraries for background.

14.3 Compilation Internal Scope

While this section uses the `cxx` module and C++ compilation as an example, the same functionality is available for C compilation – simply replace `cxx` with `c` in the module and variable names.

The `cxx` module has a notion of a project's internal scope. During compilation of a project's C/C++ translation units a header search path (`-I`) exported by a library that is outside of the internal scope is considered external and, if supported by the compiler, the corresponding `-I` option is translated to an appropriate "external header search path" option (`-isystem` for GCC/Clang, `/external:I` for MSVC 16.10 and later). In particular, this suppresses compiler warnings in such external headers (`/external:W0` is automatically added unless a custom `/external:Wn` is specified).

While the aim of this functionality is to control warnings in external libraries, the underlying mechanisms currently provided by compilers have limitations and undesirable side effects. In particular, `-isystem` paths are searched after `-I` so translating `-I` to `-isystem` alters the search order. This should normally be harmless when using a development build of a library but may result in a change of semantics for installed libraries. Also, marking the search path

as system has additional (to warning suppression) effects, see System Headers in the GCC documentation for details. On the MSVC side, `/external:W0` currently does not suppress some warnings (refer to the MSVC documentation for details).

Another issue is warnings in template instantiations. Each such warning could be either due to a (potential) issue in the template itself or due to the template arguments we are instantiating it with. By default, all such warnings are suppressed and there is currently no way to change this with GCC/Clang `-isystem`. While MSVC provides `/external:templates-`, it cannot be applied on the library by library basis, only globally for the entire compilation. See MSVC `/external:templates-` documentation for more background on this issue.

In the future this functionality will be extended to side-building BMIs for external module interfaces and header units.

The internal scope can be specified by the project with the `cxx.internal.scope` variable and overridden by the user with the `config.cxx.internal.scope` variable. Note that `cxx.internal.scope` must be specified before loading the `cxx` module (`cxx.config`, more precisely) and after which it contains the effective value (see below). For example:

```
# root.build

cxx.std = latest
cxx.internal.scope = current

using cxx
```

Valid values for `cxx.internal.scope` are:

```
current  -- current root scope (where variable is assigned)
base     -- target's base scope
root     -- target's root scope
bundle   -- target's bundle amalgamation
strong   -- target's strong amalgamation
weak     -- target's weak amalgamation
global   -- global scope (everything is internal)
```

Valid values for `config.cxx.internal.scope` are the same except for `current`.

Note also that there are `[config.]cc.internal.scope` variables that can be used to specify the internal scope for all the `cc`-based modules.

The project's effective internal scope is chosen based on the following priority list:

1. `config.cxx.internal.scope`
2. `config.cc.internal.scope`
3. effective scope from bundle amalgamation
4. `cxx.internal.scope`
5. `cc.internal.scope`

In particular, item #3 allows an amalgamation that bundles a project to override its internal scope.

If no `*.internal.scope` is specified by the project, user, or bundle, then this functionality is disabled and all libraries are treated as internal regardless of their location.

While it may seem natural to have this enabled by default, the limitations and side effects of the underlying mechanisms as well as cases where it would be undesirable (such as in separate `*-tests` projects, see below) all suggest that explicit opt-in is probably the correct choice.

The recommended value for a typical project is `current`, meaning that only headers inside the project will be considered internal. The `tests` subproject, if present, will inherit its value from the project (which acts as a bundle amalgamation), unless it is being built out of source (for example, to test an installed library).

A project can also whitelist specific libraries using the `cxx.internal.libs` variable. If a library target name (that is, the name inside `lib{}`) matches any of the wildcard patterns listed in this variable, then the library is considered internal regardless of its location. For example (notice that the pattern is quoted):

```
# root.build

cxx.std = latest
cxx.internal.scope = current
cxx.internal.libs = foo 'bar-*'

using cxx
```

Note that this variable should also be set before loading the `cxx` module and there is the common `cc.internal.libs` equivalent. However, there are no `config.*` versions nor the override by the bundle amalgamation semantics.

Typically you would want to whitelist libraries that are developed together but reside in separate build system projects. In particular, a separate `*-tests` project for a library should whitelist the library being tested if the internal scope functionality is in use. Another reason to whitelist is to catch warnings in instantiations of templates that belong to a library that is otherwise warning-free (see the `MSVC /external:templates-` option for background).

Note also that if multiple libraries are installed into the same location (or otherwise share the same header search paths, for example, as a family of libraries), then the whitelist may not be effective.

14.4 Automatic DLL Symbol Exporting

The `bin.def` module (automatically loaded by the `c` and `cxx` modules for the `*-win32-msvc` targets) provides a rule for generating symbol-exporting `.def` files. This allows automatically exporting all symbols for all the Windows targets/compilers using the following arrangement (showing for `cxx` in this example):

```
lib{foo}: libul{foo}: {hxx cxx}{**} ...

libs{foo}: def{foo}: include = ($cxx.target.system == 'win32-msvc')
def{foo}: libul{foo}

if ($cxx.target.system == 'mingw32')
  cxx.loptions += -Wl,--export-all-symbols
```

That is, we use the `.def` file approach for MSVC (including when building with Clang) and the built-in support (`--export-all-symbols`) for MinGW.

You will likely also want to add the generated `.def` file (or the blanket `*.def`) to your `.gitignore` file.

Note that it is also possible to use the `.def` file approach for MinGW. In this case we need to explicitly load the `bin.def` module (which should be done after loading `c` or `cxx`) and can use the following arrangement:

```
# root.build

using cxx

if ($cxx.target.class == 'windows')
  using bin.def

lib{foo}: libul{foo}: {hxx cxx}{**} ...

libs{foo}: def{foo}: include = ($cxx.target.class == 'windows')
def{foo}: libul{foo}
```

Note also that this only deals with exporting of the symbols from a DLL. In order to work, code that uses such a DLL should be able to import the symbols without explicit `__declspec(dllimport)` declarations. This works thanks to the symbol auto-importing support in Windows linkers. Note, however, that auto-importing only works for functions and not for global variables.

14.5 Compiler Predefined Macro Extraction

The `cc`-based modules provide the `*.predefs` submodule which can be loaded in order to register a rule that extracts predefined compiler macros. Note that the `*.predefs` module must be loaded after the respective main module and the rule will only match with an explicit rule hint. Typical usage:

```
# root.build
#
using c
using c.predefs

# buildfile
#
[rule_hint=c.predefs] h{predefs}:
```

The `predefs` rule has two modes: the so-called "pure predefs", where we preprocess an empty translation unit with the resulting set of macros consisting of only what is pre-defined by the compiler, and "custom predefs" where we preprocess a custom input header with the resulting set of macros including what is defined by such a header and headers that it includes.

The mode is determined by the presence or absence of a prerequisite of a header type. For example:

```
[rule_hint=c.predefs] h{predefs}:          # Pure.
[rule_hint=c.predefs] h{predefs}: h{config} # Custom.
```

Note that the explicit rule hint is required in both modes.

If the custom input header is only used to extract the predefined macros during the build, then you will want to make sure it does not get installed.

The `autoconf` build system module can be used to generate the custom input header for a number of common checks.

The `predefs` rule can produce its output in three forms: a header file containing a number of `#define` directive, a JSON file containing an object with each macro recorded as its member, and a `buildfile` with each macro recorded as a variable assignment. For example:

```
/* config.h */
#define FOO 1
#define BAR 123ULL
#define BAZ 0xFFFF
#undef BIZ

c.poptions += -DFOX

c.predefs.poptions = true # Include *.poptions.

[rule_hint=c.predefs] h{predefs}:          h{config}
[rule_hint=c.predefs] json{predefs}:       h{config}
[rule_hint=c.predefs] buildfile{predefs}: h{config}
{
    c.predefs.macros = FOO BAR BAZ BIZ FOX
}
```

The resulting `predefs.h` header would have the following contents:

```
#define FOO 1
#define BAR 123ULL
#define BAZ 0xffff
#define FOX 1
...
```

The resulting `predefs.json` file would have the following contents:

```
{
  "FOO": 1,
  "BAR": 123,
  "BAZ": 65535,
  "FOX": 1,
  ...
}
```

The resulting `predefs.build` file would have the following contents:

```
FOO = [uint64] 1
BAR = [uint64] 123
BAZ = [uint64] 0xffff
FOX = [uint64] 1
BIZ = [null]
```

The `buildfile` output of the `predefs` rule is typically used with the `update` during load functionality (see `update` directive) in order to communicate macro values to the `buildfiles`.

Note that macro values do not undergo recursive macro-expansion. Instead, we get the value as defined, which could be another macro. This could be a problem for the JSON and `buildfile` output since the extracted value may not be directly usable. The `__BYTE_ORDER__` macro provides a good illustration of this problem:

```
#define __ORDER_BIG_ENDIAN__ 4321
#define __ORDER_LITTLE_ENDIAN__ 1234
#define __BYTE_ORDER__ __ORDER_LITTLE_ENDIAN__
```

As a result, you may need to implement custom evaluation logic in your header in order to communicate the actual value of a macro. For example:

```
#if __BYTE_ORDER__ == __ORDER_LITTLE_ENDIAN__
#  define BYTE_ORDER_LITTLE_ENDIAN true
#elif __BYTE_ORDER__ == __ORDER_BIG_ENDIAN__
#  define BYTE_ORDER_LITTLE_ENDIAN false
#else
#  error unexpected byte order
#endif
```

In the future the `predefs` rule may perform simple macro expansions and expression evaluations automatically.

A number of variables control the behavior of the `predefs` rule:

```
[bool]      *.predefs.poptions
[string]    *.predefs.default
[string_map] *.predefs.macros
```

The `*.predefs.poptions` variable controls whether `*.poptions` are included on the compiler command line, in which case any macro definitions they may contain will end up in the output. It is `false` by default for pure `predefs` and is required if we are preprocessing a custom header (since command line macros may affect its contents).

The `*.predefs.default` variable specifies the default macro value to use in the JSON and buildfile output for macros that are not defined to any value (that is, just `#define FOO`). If not specified, then 1 is used (which is what macros specified on the command line as `-DFOO` end up being defined to by the compilers).

The `*.predefs.macros` variable specifies the macros to extract for the JSON and buildfile output. Additionally, optional mapping to member/variable name can be specified as the second half of a pair for each macro. For example:

```
c.predefs.macros = FOO BAR __SIZEOF_SIZE_T__@SIZEOF_SIZE_T
```

Note that for the buildfile output specifying `*.predefs.macros` is mandatory (since undefined macros need to be explicitly set to `null`).

Finally note also that the MSVC compiler only supports the predefined macro extraction starting from Visual Studio 2019 (16.0; `cl.exe` version 19.20). If support for earlier versions is required, then you will need to provide a fallback implementation appropriate for your project. For example:

```
[rule_hint=c.predefs] h{predefs}:
% update
if ($c.id == 'msvc' && \
    ($c.version.major < 19 || \
    ($c.version.major == 19 && $c.version.minor < 20)))
{{
    diag c-predefs $>

    cat <<EOF >$path($>)
    #define _WIN32
    EOF
}}
```

Similarly, custom predefs extraction is only supported in Clang version 12 and later due to bugs in earlier versions related to producing both macros and header dependency information into the same stream.

14.6 Importation of Installed Libraries

As discussed in Target Importation, searching for installed C/C++ libraries is seamlessly integrated into the general target importation mechanism. This section provides more details on the installed library search semantics and `pkg-config` integration. These details can be particularly useful when dealing with libraries that were not built with `build2` and which often use idiosyncratic `pkg-config` file names.

The `cc`-based modules use the common installed library search implementation with the following semantics. To illustrate the finer points, we assume the following import:

```
import libs = libbar%lib{Xfoo}
```

1. First, the ordered list of library search directories is obtained by combining two lists: the lists of the compiler's system library search directories (extracted, for example, with `-print-search-dirs` GCC/Clang options) and the list of user library search directories (specified, for example, with the `-L` options in `*.loptions`).

The key property of this combined list is that it matches the search semantics that would be used by the compiler to find libraries specified with the `-l` option during linking.

2. Given the list obtained in the previous step, a library binary (shared and/or static library) is searched for in the correct order and using the target platform-appropriate library prefix and extension (for example, `lib` prefix and the `.so/.a` extensions if targeting Linux).

For example (continuing with the above import and assuming Linux), each directory will be checked for the presence of `libXfoo.so` and `libXfoo.a` (where the `Xfoo` stem is the imported target name).

If only a shared or static binary is found in a given directory, no further directories are checked for the missing variant. Instead, the missing variant is assumed to be unavailable.

If neither a shared nor static library is found in a given directory, then it is also checked for the presence of the corresponding `pkg-config` file as in the following step. If such a file is found, then the library is assumed to be *binless* (header-only, etc).

3. If a static and/or shared library is found (or if looking for a binless library), the corresponding `pkg-config` subdirectory (normally just `pkgconfig/`) is searched for the library's `.pc` file.

More precisely, we first look for the `.static.pc` file for a static library and for the `.shared.pc` file for a shared library falling back to the common `.pc` if they don't exist.

It is often required to use different options for consuming static and shared libraries. While there is the `Libs.private` and `Cflags.private` mechanism in `pkg-config`, its semantics is to append options to `Libs` and `Cflags` rather than to provide alternative options. And often the required semantics is to provide different options for static and shared libraries, such as to provide a macro which indicates whether linking static or shared in order to setup symbol exporting.

As a result, in `build2` we produce separate `.pc` files for static and shared libraries in addition to the "best effort" common `.pc` file for compatibility with other build systems. Similarly, when consuming a library we first look for the `.static.pc` and `.shared.pc` files falling back to the common `.pc` if they are not available.

To deal with idiosyncrasies in `pkg-config` file names, the following base names are tried in order, where *name* is the imported target name (`Xfoo` in the above import), *proj* is the imported project name (`libbar` in the above import), and *ext* is one of the

above-mentioned `pkg-config` extensions (`static.pc`, `shared.pc`, or `pc`). The concrete name tried for the above import is shown in parenthesis as an example.

1. `libname.ext(libXfoo.pc)`
2. `name.ext(Xfoo.pc)`
3. `lowercase libname.ext(libxfoo.pc)`
4. `lowercase name.ext(xfoo.pc)`
5. `proj.ext(libbar.pc)`; this test is omitted if not project-qualified)

In particular, the last try (for `proj.ext`) serves as an escape hatch for cases where the `.pc` file name does not have anything to do with the names of library binaries. The canonical example of this is `zlib` which names its library binaries `libz.so/libz.a` while its `.pc` file – `zlib.pc`. To be able to import `zlib` that was not built with `build2`, we have to use the following import:

```
import libs = zlib%lib{z}
```

Note also that these complex rules (which are unfortunately necessary to deal with the lack of any consistency in `.pc` file naming) can sometimes produce surprising interactions. For example, it may appear that a clearly incorrect import nevertheless appears to somehow work, as in the following example:

```
import libs = zlib%lib{znonsense}
```

What happens here is that while no library binary is found, `zlib.pc` is found and as a result the library ends up being considered binless with the `-lz` (that is found in the `Libs` value of `zlib.pc`) treated as a prerequisite library, resolved using the above algorithm, and linked. In other words, in this case we end up with a binless library `lib{znonsense}` that depends on `lib{z}` instead of a single `lib{z}` library.

14.6.1 Rewriting Installed Libraries System Root (sysroot)

Sometimes the installed libraries are moved to a different location after the installation. This is especially common in embedded development where the code is normally cross-compiled and the libraries for the target platform are placed into a host directory, called system root or *sysroot*, that doesn't match where these libraries were originally installed to. For example, the libraries might have been installed into `/usr/` but on the host machine they may reside in `/opt/target/usr/`. In this example, `/opt/target/` is the *sysroot*.

While such relocations usually do not affect the library headers or binaries, they do break the `pkg-config`'s `.pc` files which often contain `-I` and `-L` options with absolute paths. Continue with the above example, a `.pc` file as originally installed may contain `-I/usr/include` and `-L/usr/lib` while now, that the libraries have been relocated to `/opt/target/`, they somehow need to be adjusted to `-I/opt/target/usr/include` and `-L/opt/target/usr/lib`.

While it is possible (and perhaps correct) to accomplish this by fixing the `.pc` files to match the new location, it is not always possible or easy. As a result, `build2` provides a mechanism for automatically adjusting the system root in the `-I` and `-L` options extracted from `.pc` files.

This functionality is roughly equivalent to that provided with the `PKG_CONFIG_SYSROOT_DIR` environment variable by the `pkg-config` utility.

Specifically, the `config.cc.pkgconfig.sysroot` variable can be used to specify an alternative system root. When specified, all absolute paths in the `-I` and `-L` options that are not already in this directory will be rewritten to start with this `sysroot`.

Note that this mechanism is a workaround rather than a proper solution since it is limited to the `-I` and `-L` options. In particular, it does not handle any other options that may contain absolute paths nor `pkg-config` variables that may be queried.

As a result, it should only be used for dealing with issues in third-party `.pc` files that do not handle relocation (for example, using the `${pcfiledir}` built-in `pkg-config` variable). In particular, for `build2`-generated `.pc` files a relocatable installation should be used instead.

14.7 Compilation Database

The `cc`-based modules provide support for generating and maintaining the JSON Compilation Database which can be used by other tools (static analyzers, language servers, IDEs, etc) to understand how a codebase is compiled. "Maintaining" in the previous sentence means that if new source files get added to the project or old ones removed, or if any compilation options change, then the corresponding entries in the compilation database will be automatically updated when you update your project. This helps maintain the database in sync with the project state.

Note that the removal of old database entries only happens during the `update` operation and only for entries that reside in subdirectories being updated.

The generation of compilation databases and their configuration are controlled with a number of `config.cc.compiledb.*` variables. The `config.cc.compiledb` variable provides a simplified interface that enables the generation of one database per project with the resulting database containing entries for all the source and object files. The rest of the variables provide a more flexible interface that allows you to generate multiple databases in different locations as well as filter the entries that end up in each database.

Let's start with the simplified interface as provided by `config.cc.compiledb`. The value of this configuration variable is a single *name* or a *name* and *path* pair in the `name[@path]` form.

The *name* part is the compilation database name that can be used to refer to it in filters (see below). If *path* is absent or is (syntactically) a directory, then *name* is also used to derive the compilation database file by appending the `.json` extension to it.

If *path* is absent, then the compilation database is placed into the top-level amalgamation that loads any `cc`-based module. Otherwise, the database is placed into the specified location.

The special `- name` is interpreted as an instruction to dump the database to `stdout`.

Let's see some examples of using `config.cc.compiledb` to handle a few common scenarios. Here we will use **bdep (1)** to create amalgamations (configurations) and configure (initialize) one or more projects. We will assume we have `hello` and `libhello` as if created like this:

```
$ bdep new -t exe hello
$ bdep new -t lib libhello
```

The most common scenario is likely having a compilation database per project:

```
$ cd libhello
$ bdep config create ../build-gcc @gcc cc config.cxx=g++
$ bdep init @gcc config.cc.compiledb=libhello
$ cd ..

$ cd hello
$ bdep config add ../build-gcc @gcc
$ bdep init @gcc config.cc.compiledb=hello
$ cd ..

$ b hello/ libhello/
```

Or if you prefer to create/add configuration as part of `init` (notice the `--` separator):

```
$ bdep init -C ../build-gcc @gcc cc config.cxx=g++ -- \
  config.cc.compiledb=libhello

$ bdep init -A ../build-gcc @gcc config.cc.compiledb=hello
```

After the update (the last command), we will have `hello.json` and `libhello.json` in `build-gcc/` which contain the compilation command lines for each project.

Only source files that are compiled end up being added to the compilation database.

To illustrate this point, let's assume our `hello` project imports and links `libhello`. And instead of updating both as in the above example, we will first update only `hello`:

```
$ b hello/
```

In this case `libhello.json` will still be generated but it will only contain a subset of the expected entries – only those that were caused to be compiled by `hello`. The missing entries can be added by updating `libhello`:

```
$ b libhello/
```

In the above setup it feels natural to call each database after the project and place them into the output directory. However, some consumers, such as IDEs and LSP servers, may not handle this setup well. Specifically, they may only recognize the canonical `compile_commands.json` file as the compilation database, opening all other files as generic JSON. They may also assume the directory where this file resides to be the project source directory root. To accommodate these assumptions we can instead place each database into the project's source directory and call it `compile_commands.json`:

```
$ cd libhello
$ bdep init @gcc config.cc.compiledb=libhello@./compile_commands.json

$ cd hello
$ bdep init @gcc config.cc.compiledb=hello@./compile_commands.json
```

To facilitate this use-case, `config.cc.compiledb` supports another shortcut: if we specify just *name* and it contains a directory component, then it is interpreted as *path* rather than *name*. In this case *name* is taken to be the name of the last directory component in *path* (which would typically be a project or package name). And if *path* is a directory, then the database file name is taken to be `compile_commands.json`. Or, in other words, the following:

```
config.cc.compiledb=.../<dir>/
```

Is equivalent to:

```
config.cc.compiledb=<dir>@.../<dir>/compile_commands.json
```

This shortcut allows us to simplify the above `init` commands to read:

```
$ cd libhello
$ bdep init @gcc config.cc.compiledb=./

$ cd hello
$ bdep init @gcc config.cc.compiledb=./
```

Note also that in this case it will be your responsibility to remove the database files if and when necessary. **bdep-new(1)** adds `compile_commands.json` to `.gitignore` it generates.

If instead of having a separate database for each project we wanted to place all the entries into a single database (and in the output directory), then the relevant commands would change as follows:

```
$ bdep init @gcc config.cc.compiledb=compiledb

$ bdep init @gcc config.cc.compiledb=compiledb
```

This would give us a single `build-gcc/compiledb.json` that contains the compilation command lines for both projects.

In the above example only `hello` and `libhello` will end up in the database, but not any of their dependencies. What if we wanted entries for everything in `build-gcc`? In this case, we should enable the compilation database for the entire configuration rather than for individual projects:

```
$ bdep config create ../build-gcc @gcc cc \
  config.cxx=g++ \
  config.cc.compiledb=compiledb
$ bdep init @gcc

$ bdep config add ../build-gcc @gcc
$ bdep init @gcc
```

If multiple linked configurations are involved, then we would often want projects initialized in different configurations share the compilation database. The representative scenario here is a tool, such as a source code generator, which is initialized in the host configuration, and its runtime library plus tests/examples, which are initialized in the target configuration. Let's assume that in our example `hello` is the tool and `libhello` is the runtime library and both are part of the same project. This is how we can arrange for them to share the compilation database:

```
$ bdep config create @host ../host-gcc --type host cc config.cxx=g++
$ bdep config create @target ../build-gcc cc config.cxx=g++

$ bdep init @host -d hello config.cc.compiledb=hello@../build-gcc/
$ bdep init @target -d libhello config.cc.compiledb=hello

$ bdep update @host @target
```

With this setup the `hello.json` database in `build-gcc/` will contain entries for both `hello` and `libhello`.

If instead of configuring and maintaining the compilation database in a file you want to dump it somewhere once, the recommended approach is to write it to `stdout`. For example:

```
$ b -n hello/ libhello/ config.cc.compiledb=- >/tmp/compiledb.json
```

Note that writing to `stdout` forces recompilation of all the targets that would be updated in order to make sure their entries end up in the database. If you don't want the actual recompilation, then you can use the dry run mode (`-n` option above).

If your projects are spread across multiple linked configurations and you would like to get compilation command lines for all of them, then use the global override for `config.cc.compiledb`:

```
$ b '!config.cc.compiledb=-' ...
```

As mentioned earlier, the entries that will end up in such a database are determined by what gets updated.

Let's now turn to the rest of the `config.cc.compiledb.*` configuration variables that provide a lower-level but more flexible interface. The following listing shows their synopsis:

```
config.cc.compiledb.name           = <name>[@<path>]...
config.cc.compiledb.filter         = [<name>@]<bool>...
config.cc.compiledb.filter.input   = [<name>@]<target-type>...
config.cc.compiledb.filter.output  = [<name>@]<target-type>...
```

The `config.cc.compiledb.name` variable specifies the name and location of one or more compilation databases. The semantics of the `name[@path]` pair is the same as in `config.cc.compiledb` discussed above, except that if `path` is absent, then the database is placed into the project being configured rather than into the top-level amalgamation.

Also, unlike `config.cc.compiledb`, this variable does not automatically enable writing to the specified databases. Instead, this is the job of `config.cc.compiledb.filter`. Splitting this logic into two steps allows us to configure the database name/location in one place, typically an outer amalgamation, and then enable writing to it in other places, typically specific subprojects.

The `config.cc.compiledb.filter.{input,output}` variables allow us to filter the entries that end up in the databases based on the input (`c{}`, `cxx{}`, etc) and output (`obja{}`, `objs{}`, etc) target types.

Note that in all three `.filter` variables the values are examined in the reverse order and the first entry that matches determines the outcome. Entries without `name` apply to all databases and the target types are matched taking into account inheritance (so `target{}` will match any type) and groups (so `obj{}` will match any `obj[eas]{}`). If no target type filter (input or output) is specified, then no corresponding target filtering is performed.

The `config.cc.compiledb=<name>` semantics can be expressed as the following set of lower-level variables:

```
config.cc.compiledb.name           = <name>@../path/to/amalgamation/
config.cc.compiledb.filter         += <name>@true
config.cc.compiledb.filter.input   += <name>@target
config.cc.compiledb.filter.output  += <name>@target
```

The last three assignments only apply if the corresponding variable is not set to a custom value for this project.

Let's look at a few examples of using these lower-level configuration variables. The common use for the output target filtering is getting rid of `obja{}` or `objs{}` entries in libraries. Unless configured otherwise, when we build a library we end up with both static and shared variants. And this means that each source file for the library is compiled twice, once to produce `obja{}` that goes to the static library and once -- `objs{}`. And that, in turn, means that we will end up with two compilation database entries for each such source file. If we don't want that for some reason (for instance, because the consumer of the database does not handle this well), then we can filter one of them out. For example, below is how we can initialize `libhello` to achieve this (notice that we also include `obje{}` to keep object files for executables, such as tests):

```
$ bdep init @gcc \
  config.cc.compiledb=libhello \
  config.cc.compiledb.filter.output='obje objs'
```

As an example of the input target type filtering, below is how we can keep entries only for the C and C++ source files, filtering out everything else (assembler, Objective-C/C++), for instance, because the consumer of our database does not recognize them:

```
$ bdep init @gcc \
  config.cc.compiledb=libhello \
  config.cc.compiledb.filter.input='c cxx'
```

As an example of a more advanced configuration, consider a compilation database for a project that use C++ modules. To know how such a project is compiled we not only need to know how its own source files are compiled, but also how to compile all the module interfaces that it consumes, including from other projects, transitively. One way to set this up would be to enable writing entries of the `bmi{ }` output target type to any database in the amalgamation:

```
$ bdep config create ../build-gcc @gcc cc \
  config.cxx=g++ \
  config.cc.compiledb.filter=true \
  config.cc.compiledb.filter.output=bmi \

$ bdep init @gcc config.cc.compiledb=libhello

$ bdep init @gcc config.cc.compiledb=hello
```

With this setup `libhello.json` and `hello.json` will contain module interface entries from all the dependencies.

When debugging complex compilation database setups it can be helpful to increase diagnostics verbosity to level 6 in order to get a trace of filtering decisions (the relevant lines will contain the `compiledb` keyword).

14.8 GCC Compiler Toolchain

The GCC compiler id is `gcc`.

14.9 Clang Compiler Toolchain

The vanilla Clang compiler id is `clang` (including when targeting the MSVC runtime), Apple Clang compiler id is `clang-apple`, and Clang's `cl` compatibility driver (`clang-cl`) id is `msvc-clang`.

14.9.1 Clang Targeting MSVC

There are two common ways to obtain Clang on Windows: bundled with the MSVC installation or as a separate installation. If you are using the separate installation, then the Clang compiler is most likely already in the `PATH` environment variable. Otherwise, if you are using Clang that is bundled with MSVC, the `cc` module will attempt various search strategies described below. Note, however, that in both cases once the Clang compiler binary is located, the mode (32 or 64-bit) and the rest of the environment (locations of binary utilities as well as the system headers and libraries) are obtained by querying Clang.

Normally, if Clang is invoked from one of the Visual Studio command prompts, then it will use the corresponding Visual Studio version and environment (it is, however, still up to you to match the mode with the `-m32/-m64` options, if necessary). Otherwise, Clang will try to locate the latest version of Visual Studio and Platform SDK and use that (in this case it matches the environment to the `-m32/-m64` options). Refer to Clang documentation for details.

If you specify the compiler as just `config.c=clang` or `config.cxx=clang++` and it is found in the `PATH` environment variable or if you specify it as an absolute path, then the `cc` module will use that.

Otherwise, if you are building from one of the Visual Studio development command prompts, the `cc` module will look for the corresponding bundled Clang (`%VCINSTALLDIR%\Tools\Llvm\bin`).

Finally, the `cc` module will attempt to locate the latest installed version of Visual Studio and look for a bundled Clang in there.

The default mode (32 or 64-bit) depends on the Clang configuration and can be overridden with the `-m32/-m64` options. For example:

```
> b "config.cxx=clang++ -m64"
```

The default MSVC runtime selected by the `cc` module is multi-threaded shared (the `/MD` option in `cl`). Unfortunately, the Clang driver does not yet provide anything equivalent to the `cl /M*` options (see Clang bug #33273) and selection of an alternative runtime has to be performed manually:

```
> rem /MD - multi-threaded shared (default)
> rem
> b "config.cxx=clang++ -nostdlib -D_MT -D_DLL" ^
    config.cc.libs=/DEFAULTLIB:msvcrt

> rem /MDd - multi-threaded debug shared
> rem
> b "config.cxx=clang++ -nostdlib -D_MT -D_DLL -D_DEBUG" ^
    config.cc.libs=/DEFAULTLIB:msvcrt_d

> rem /MT - multi-threaded static
> rem
> b "config.cxx=clang++ -nostdlib -D_MT" ^
```

```

config.cc.libs=/DEFAULTLIB:libcmt

> rem /MTd - multi-threaded debug static
> rem
> b "config.cxx=clang++ -nostdlib -D_MT -D_DEBUG" ^
    config.cc.libs=/DEFAULTLIB:libcmt

```

By default the MSVC's binary utilities (`link` and `lib`) are used when compiling with Clang. It is, however, possible to use LLVM's versions instead, for example:

```

> b config.cxx=clang++ ^
    config.bin.ld=lld-link ^
    config.bin.ar=llvm-lib

```

In particular, one benefit of using `llvm-lib` is support for thin archives which, if available, is automatically enabled for utility libraries.

While there is basic support for Clang's `cl` compatibility driver (`clang-cl`), its use is not recommended. This driver is a very thin wrapper over the standard Clang interface that does not always recreate the `cl`'s semantics exactly. Specifically, its diagnostics in the `/showIncludes` mode does not match that of `cl` in the presence of missing headers. As a result, `clang-cl`'s use, if any, should be limited to projects that do not have auto-generated headers.

If you need to link with other projects that use `clang-cl`, then the recommended approach is to discover any additional `cl` options passed by `clang-cl` by comparing the `-v` output of a test compilation with `clang-cl` and `clang/clang++` and then passing them explicitly to `clang/clang++`, potentially prefixed with `-Xclang`. For example:

```
b "config.cxx=clang++ -Xclang -fms-volatile ..."
```

Relevant additional options that are passed by `clang-cl` at the time of this writing:

```

-fno-strict-aliasing
-fstack-protector-strong
-Xclang -fms-volatile
-ffunction-sections

```

14.10 MSVC Compiler Toolchain

The Microsoft VC (MSVC) compiler id is `msvc`.

There are several ways to specify the desired MSVC compiler and mode (32 or 64-bit) as well as the corresponding environment (locations of binary utilities as well as the system headers and libraries).

Unlike other compilers, MSVC compiler (`cl`) binaries are target-specific, that is, there are no `-m32/-m64` options nor something like the `/MACHINE` option available in `link`.

If the compiler is specified as just `cl` in `config.{c,cxx}` and it is found in the `PATH` environment variable, then the `cc` module assumes the build is performed from one of the Visual Studio development command prompts and expects the environment (the `PATH`, `INCLUDE`, and `LIB` environment variables) to already be setup.

If, however, `cl` is not found in `PATH`, then the `cc` module will attempt to locate the latest installed version of Visual Studio and Platform SDK and use that in the 64-bit mode.

Finally, if the compiler is specified as an absolute path to `cl`, then the `cc` module will attempt to locate the corresponding Visual Studio installation as well as the latest Platform SDK and use that in the mode corresponding to the specified `cl` executable. Note that to specify an absolute path to `cl` (which most likely contains spaces) we have to use two levels of quoting:

```
> b "config.cxx='...\\VC\\Tools\\MSVC\\14.23.28105\\bin\\Hostx64\\x86\\cl'"
```

The latter two methods are only available for Visual Studio 15 (2017) and later and for earlier versions the development command prompt must be used.

The default MSVC runtime selected by the `cc` module is multi-threaded shared (the `/MD cl` option). An alternative runtime can be selected by passing one of the `cl /M*` options, for example:

```
> b "config.cxx=cl /MT"
```

15 c Module

This chapter is a work in progress and is incomplete.

This chapter describes the `c` build system module which provides the C compilation and linking support. Most of its functionality, however, is provided by the `cc` module, a common implementation for the C-family languages.

15.1 C Configuration Variables

The following listing summarizes the `c` module configuration variables as well as the corresponding module-specific variables that are derived from their values. See also C-Common Configuration Variables.

```
config.c
  c.path
  c.mode

config.c.id
  c.id
  c.id.type
  c.id.variant
  c.class

config.c.version
  c.version
  c.version.major
```

```

c.version.minor
c.version.patch
c.version.build

config.c.target
  c.target
  c.target.cpu
  c.target.vendor
  c.target.system
  c.target.version
  c.target.class

config.c.std
  c.std

config.c.poptions
  c.poptions

config.c.coptions
  c.coptions

config.c.loptions
  c.loptions

config.c.aoptions
  c.aoptions

config.c.libs
  c.libs

config.c.internal.scope
  c.internal.scope

```

15.2 C Target Types

The following listing shows the hierarchy of the target types defined by the `c` module while the following sections describe each target type in detail (`file{}` is a standard target type defined by the `build2` core; see Target Types for details). See also C-Common Target Types for target types defined by all the `cc`-based modules.

```

.--file--
|   |   |
c   m   S
h

```

The `m{}` target type represents an Objective-C source file, see Objective-C Compilation for details.

The `S{}` target type represents an Assembler with C Preprocessor file, see Assembler with C Preprocessor Compilation for details.

15.2.1 `c{}`, `h{}`

The `c{}` and `h{}` target types represent C source and header files. They have the default extensions `.c` and `.h`, respectively, which can be customized with the `extension` variable.

15.3 Objective-C Compilation

The `c` module provides the `c.objc` submodule which can be loaded in order to register the `m{}` target type and enable Objective-C compilation in the `C` compile rule. Note that `c.objc` must be loaded after the `c` module and while the `m{}` target type is registered unconditionally, compilation is only enabled if the C compiler supports Objective-C for the target platform. Typical usage:

```
# root.build
#
using c
using c.objc

# buildfile
#
lib{hello}: {h c}{*}
lib{hello}: m{*}: include = ($c.target.class == 'macos')
```

Note also that while there is support for linking Objective-C executables and libraries, this is done using the C compiler driver and no attempt is made to automatically link any necessary Objective-C runtime library (such as `-lobjc`).

15.4 Assembler with C Preprocessor Compilation

The `c` module provides the `c.as-cpp` submodule which can be loaded in order to register the `S{}` target type and enable Assembler with C Preprocessor compilation in the `C` compile rule. Note that `c.as-cpp` must be loaded after the `c` module and while the `S{}` target type is registered unconditionally, compilation is only enabled if the C compiler supports Assembler with C Preprocessor compilation. Typical usage:

```
# root.build
#
using c
using c.as-cpp

# buildfile
#
exe{hello}: {h c}{*} -hello.c

# Use C implementation as a fallback if no assembler.
#
assembler = ($c.class == 'gcc' && $c.target.cpu == 'x86_64')

exe{hello}: S{hello}: include = $assembler
exe{hello}: c{hello}: include = (!$assembler)
```

```

/* hello.S
 */
#ifdef HELLO_RESULT
# define HELLO_RESULT 0
#endif

text

.global hello
hello:
    /* ... */
    movq $HELLO_RESULT, %rax
    ret

#ifdef __ELF__
.section .note.GNU-stack, "", @progbits
#endif

```

The default file extension for the `S{ }` target type is `.S` (capital) but that can be customized using the standard mechanisms. For example:

```

# root.build
#
using c
using c.as-cpp

h{*}: extension = h
c{*}: extension = c
S{*}: extension = sx

```

Note that `*.coptions` are passed to the C compiler when compiling Assembler with C Preprocessor files because compile options may cause additional preprocessor macros to be defined. Plus, some of them (such as `-g`) are passed (potentially translated) to the underlying assembler. To pass additional options when compiling Assembler files use `c.poptions` and `c.coptions`. For example (continuing with the previous example):

```

if $assembler
{
    obj{hello}:
    {
        c.poptions += -DHELLO_RESULT=1
        c.coptions += -Wa,--no-pad-sections
    }
}

```

15.5 C Compiler Predefined Macro Extraction

The `c` module provides the `c.predefs` submodule which can be loaded in order to register a rule that generates a C header with predefined compiler macros. Note that the `c.predefs` module must be loaded after the `c` module and the rule will only match with an explicit rule hint. Typical usage:

```

# root.build
#
using c
using c.predefs

```

```
# buildfile
#
[rule_hint=c.predefs] h{predefs}:
```

See Compiler Predefined Macro Extraction for details.

16 cxx Module

This chapter is a work in progress and is incomplete.

This chapter describes the `cxx` build system module which provides the C++ compilation and linking support. Most of its functionality, however, is provided by the `cc` module, a common implementation for the C-family languages.

16.1 C++ Configuration Variables

The following listing summarizes the `cxx` module configuration variables as well as the corresponding module-specific variables that are derived from their values. See also C-Common Configuration Variables.

```
config.cxx
  cxx.path
  cxx.mode

config.cxx.id
  cxx.id
  cxx.id.type
  cxx.id.variant
  cxx.class

config.cxx.version
  cxx.version
  cxx.version.major
  cxx.version.minor
  cxx.version.patch
  cxx.version.build

config.cxx.target
  cxx.target
  cxx.target.cpu
  cxx.target.vendor
  cxx.target.system
  cxx.target.version
  cxx.target.class

config.cxx.std
  cxx.std

config.cxx.poptions
  cxx.poptions

config.cxx.coptions
  cxx.coptions

config.cxx.loptions
  cxx.loptions
```

```

config.cxx.aoptions
  cxx.aoptions

config.cxx.libs
  cxx.libs

config.cxx.internal.scope
  cxx.internal.scope

config.cxx.translate_include
  cxx.translate_include

```

16.2 C++ Target Types

The following listing shows the hierarchy of the target types defined by the `cxx` module while the following sections describe each target type in detail (`file{}` is a standard target type defined by the `build2` core; see Target Types for details). See also C-Common Target Types for target types defined by all the `cc`-based modules.

```

      .--file--.
      |         |
    cxx         mm
    hxx
    ixx
    txx
    mxx

```

The `mm{}` target type represents an Objective-C++ source file, see Objective-C++ Compilation for details.

16.2.1 `cxx{}`, `hxx{}`, `ixx{}`, `txx{}`, `mxx{}`

The `cxx{}`, `hxx{}`, `ixx{}`, `txx{}`, and `mxx{}` target types represent C++ source, header, inline, template, and module interface files. They have the default extensions `.cxx`, `.hxx`, `.ixx`, `.txx`, and `.mxx`, respectively, which can be customized with the `extension` variable. For example (normally done in `root.build`):

```

using cxx

cxx{*}: extension = cpp
hxx{*}: extension = hpp
mxx{*}: extension = cppm

```

16.3 C++ Modules Support

This section describes the build system support for C++ modules.

16.3.1 Modules Introduction

The goal of this section is to provide a practical introduction to C++ Modules and to establish key concepts and terminology. You can skip directly to Building Modules if you are already familiar with this topic.

A pre-modules C++ program or library consists of one or more *translation units* which are customarily referred to as C++ source files. Translation units are compiled to *object files* which are then linked together to form a program or library.

Let's also recap the difference between an *external name* and a *symbol*: External names refer to language entities, for example classes, functions, and so on. The *external* qualifier means they are visible across translation units.

Symbols are derived from external names for use inside object files. They are the cross-referencing mechanism for linking a program from multiple, separately-compiled translation units. Not all external names end up becoming symbols and symbols are often *decorated* with additional information, for example, a namespace. We often talk about a symbol having to be satisfied by linking an object file or a library that provides it. Similarly, duplicate symbol issues may arise if more than one object file or library provides the same symbol.

What is a C++ module? It is hard to give a single but intuitive answer to this question. So we will try to answer it from three different perspectives: that of a module consumer, a module producer, and a build system that tries to make those two play nice. But we can make one thing clear at the outset: modules are a *language-level* not a preprocessor-level mechanism; it is `import`, not `#import`.

One may also wonder why C++ modules, what are the benefits? Modules offer isolation, both from preprocessor macros and other modules' symbols. Unlike headers, modules require explicit exportation of entities that will be visible to the consumers. In this sense they are a *physical design mechanism* that forces us to think how we structure our code. Modules promise significant build speedups since importing a module, unlike including a header, should be essentially free. Modules are also the first step to not needing the preprocessor in most translation units. Finally, modules have a chance of bringing to mainstream reliable and easy to setup distributed C++ compilation, since with modules build systems can make sure compilers on the local and remote hosts are provided with identical inputs.

To refer to a module we use a *module name*, a sequence of dot-separated identifiers, for example `hello.core`. While the specification does not assign any hierarchical semantics to this sequence, it is customary to refer to `hello.core` as a submodule of `hello`. We discuss submodules and provide the module naming guidelines below.

From a consumer's perspective, a module is a collection of external names, called *module interface*, that become *visible* once the module is imported:

```
import hello.core;
```

What exactly does *visible* mean? To quote the standard: *An import-declaration makes exported declarations [...] visible to name lookup in the current translation unit, in the same namespaces and contexts [...]. [Note: The entities are not redeclared in the translation unit containing the module import declaration. -- end note]* One intuitive way to think about this visibility is *as if* there were only a single translation unit for the entire program that contained all the modules as well as all their consumers. In such a translation unit all the names would be visible to everyone in exactly the same way and no entity would be redeclared.

This visibility semantics suggests that modules are not a name scoping mechanism and are orthogonal to namespaces. Specifically, a module can export names from any number of namespaces, including the global namespace. While the module name and its namespace names need not be related, it usually makes sense to have a parallel naming scheme, as discussed below. Finally, the `import` declaration does not imply any additional visibility for names declared inside namespaces. Specifically, to access such names we must continue using the existing mechanisms, such as qualification or using declaration/directive. For example:

```
import hello.core;           // Exports hello::say().

say ();                     // Error.
hello::say ();              // Ok.

using namespace hello;
say ();                     // Ok.
```

Note also that from the consumer's perspective a module does not provide any symbols, only C++ entity names. If we use names from a module, then we may have to satisfy the corresponding symbols using the usual mechanisms: link an object file or a library that provides them. In this respect, modules are similar to headers and as with headers, module's use is not limited to libraries; they make perfect sense when structuring programs. Furthermore, a library may also have private or implementation modules that are not meant to be imported by the library's consumers.

The producer perspective on modules is predictably more complex. In pre-modules C++ we only had one kind of translation unit (or source file). With modules there are three kinds: *module interface unit*, *module implementation unit*, and the original kind which we will call a *non-module translation unit*.

There are two additional modular translation units: module interface partition and module implementation partition. While partitions are supported, they are not covered in this introduction. A link to a complete example that uses both types of partitions will be given in the next section.

From the producer's perspective, a module is a collection of module translation units: one interface unit and zero or more implementation units. A simple module may consist of just the interface unit that includes implementations of all its functions (not necessarily inline). A more complex module may span multiple implementation units.

A translation unit is a module interface unit if it contains an *exporting module declaration*:

```
export module hello;
```

A translation unit is a module implementation unit if it contains a *non-exporting module declaration*:

```
module hello;
```

While module interface units may use the same file extension as normal source files, we recommend that a different extension be used to distinguish them as such, similar to header files. While the compiler vendors suggest various (and predictably different) extensions, our recommendation is `.mxx` for the `.hxx/.cxx` source file naming and `.mpp` for `.hpp/.cpp`. And if you are using some other naming scheme, then perhaps now is a good opportunity to switch to one of the above. Continuing using the source file extension for module implementation units appears reasonable and that's what we recommend.

A modular translation unit (that is, either module interface or implementation) that does not start with one of the above module declarations must then start with the module introducer:

```
module;

...

export module hello;
```

The fragment from the module introducer and until the module declaration is called the *global module fragment*. Any name declared in the global module fragment belongs to the *global module*, an implied module containing "old" or non-modular declarations that don't belong to any named module.

A module declaration (exporting or non-exporting) starts a *module purview* that extends until the end of the module translation unit. Any name declared in a module's purview *belongs* to the said module. For example:

```
module;                                // Start of global module fragment.

#include <cassert>                       // Not in purview.

export module hello;                   // Start of purview.

import std;                           // In purview.

void say_hello (const std::string&);   // In purview.
```

A name that belongs to a module is *invisible* to the module's consumers unless it is *exported*. A name can be declared exported only in a module interface unit, only in the module's purview, and there are several syntactic ways to accomplish this. We can start the declaration with the `export` specifier, for example:

```
export module hello;

export enum class volume {quiet, normal, loud};

export void say_hello (const char*, volume);
```

Alternatively, we can enclose one or more declarations into an *exported group*, for example:

```
export module hello;

export
{
    enum class volume {quiet, normal, loud};

    void say_hello (const char*, volume);
}
```

Finally, if a namespace definition is declared exported, then every name in its body is exported, for example:

```
export module hello;

export namespace hello
{
    enum class volume {quiet, normal, loud};

    void say_hello (const char*, volume);
}

namespace hello
{
    void impl (const char*, volume); // Not exported.
}
```

Up until now we've only been talking about names belonging to a module. What about the corresponding symbols? All the major C++ compilers have chosen to implement the so-called strong ownership model, where for both exported and non-exported names, the corresponding symbols are decorated with the module name. As a result, they cannot clash with symbols for identical names from other named modules or the global module.

What about the preprocessor? Modules do not export preprocessor macros, only C++ names. A macro defined in the module interface unit cannot affect the module's consumers. And macros defined by the module's consumers cannot affect the module interface they are importing. In other words, module producers and consumers are isolated from each other where the preprocessor is concerned. For example, consider this module interface:

```
export module hello;

#ifdef SMALL
#define HELLO
export void say_hello (const char*);
#endif
```

And its consumer:

```
// module consumer
//
#define SMALL          // No effect.
import hello;

#ifdef HELLO           // Not defined.
...
#endif
```

This is not to say that the preprocessor cannot be used by either the module interface or its consumer, it just that macros don't "leak" through the module interface. One practical consequence of this model is the insignificance of the importation order.

If a module imports another module in its purview, the imported module's names are not made automatically visible to the consumers of the importing module. This is unlike headers and can be surprising. Consider this module interface as an example:

```
export module hello;

import std;

export std::string formal_hello (const std::string&);
```

And its consumer:

```
import hello;

int
main ()
{
    std::string s (format_hello ("World"));
}
```

This example will result in a compile error and the diagnostics may confusingly indicate that there is no member `string` in namespace `std`. But with the understanding of the difference between `import` and `#include` the reason should be clear: while the module interface "sees" `std::string` (because it imported its module), we (the consumer) do not (since we did not). So the fix is to explicitly import `std`:

```
import std;
import hello;

int
main ()
{
    std::string s (format_hello ("World"));
}
```

A module, however, can choose to re-export a module it imports. In this case, all the names from the imported module will also be visible to the importing module's consumers. For example, with this change to the module interface the first version of our consumer will compile without errors (note that whether this is a good design choice is debatable, as discussed below):

```
export module hello;

export import std;

export std::string formal_hello (const std::string&);
```

One way to think of a re-export is *as if* an import of a module also "injects" all the imports the said module re-exports, recursively. That's essentially how most compilers implement it.

Module re-export is the mechanism for assembling bigger modules out of submodules. As an example, let's say we had the `hello.core`, `hello.basic`, and `hello.extra` modules. To make life easier for users that want to import all of them we can create the `hello` module that re-exports the three:

```
export module hello;

export
{
    import hello.core;
    import hello.basic;
    import hello.extra;
}
```

Besides starting a module purview, a non-exporting module declaration in the implementation unit makes (non-internal linkage) names declared or made visible (via import) in the module purview of an interface unit also visible in the module purview of the implementation unit. In this sense a non-exporting module declaration acts as a special import. The following example illustrates this point:

```
module;

import hello.impl;           // Not visible (exports impl()).

#include <string.h>           // Not visible (declares strlen()).

export module hello.extra;   // Start of module purview (interface).

import hello.core;          // Visible (exports core()).

void extra ();               // Visible.

static void extra2 ();       // Not visible (internal linkage).
```

And this is the implementation unit:

```
module hello.extra;          // Start of module purview (implementation).

void
f ()
{
    impl ();                 // Error.
    strlen ("");             // Error.
    core ();                  // Ok.
    extra ();                 // Ok.
    extra2 ();                // Error.
}
```

In particular, this means that while the relative order of imports is not significant, the placement of imports in the module interface unit relative to the module declaration can be.

The final perspective that we consider is that of the build system. From its point of view the central piece of the module infrastructure is the *binary module interface* or BMI: a binary file that is produced by compiling the module interface unit and that is required when compiling any translation unit that imports this module as well as the module's implementation units.

Then, in a nutshell, the main functionality of a build system when it comes to modules support is figuring out the order in which all the translation units should be compiled and making sure that every compilation process is able to find the binary module interfaces it needs.

Predictably, the details are more complex. Compiling a module interface unit produces two outputs: the binary module interface and the object file. The latter contains object code for non-inline functions, global variables, etc., that the interface unit may define. This object file has to be linked when producing any binary (program or library) that uses this module.

Also, all the compilers currently implement module re-export as a shallow reference to the re-exported module name which means that their binary interfaces must be discoverable as well, recursively. In fact, currently, all the imports are handled like this, though a different implementation is at least plausible, if unlikely.

While the details vary between compilers, the contents of the binary module interface can range from a stream of preprocessed tokens to something fairly close to object code. As a result, binary interfaces can be sensitive to the compiler options and if the options used to produce the binary interface (for example, when building a library) are sufficiently different compared to the ones used when compiling the module consumers, the binary interface may be unusable. So while a build system should strive to reuse existing binary interfaces, it should also be prepared to compile its own versions "on the side".

This also suggests that binary module interfaces are not a distribution mechanism and should probably not be installed. Instead, we should install and distribute module interface sources and build systems should be prepared to compile them, again, on the side.

16.3.2 Building Modules

Compiler support for C++ modules is still experimental, incomplete, and often buggy. Also, in `build2`, the presence of modules changes the C++ compilation model in ways that would introduce unnecessary overheads for headers-only code. As a result, a project must explicitly enable modules using the `cxx.features.modules` boolean variable. This is what the relevant `root.build` fragment could look like for a modularized project:

```
cxx.std = latest
cxx.features.modules = true

using cxx

mxx{*}: extension = mxx
cxx{*}: extension = cxx
```

Note that you must explicitly enable modules in your project even if you are only importing other modules, including standard library modules (`std` or `std.compat`).

To support C++ modules the `cxx` build system module defines several additional target types. The `mxx{}` target is a module interface unit. As you can see from the above `root.build` fragment, in this project we are using the `.mxx` extension for our module interface files. While you can use the same extension as for `cxx{}` (source files), this is not recommended since some functionality, such as wildcard patterns, will become unusable.

The `bmi{}` group and its `bmie{}`, `bmia{}`, and `bmis{}` members are used to represent binary module interfaces targets. We normally do not need to mention them explicitly in our buildfiles except, perhaps, to specify additional, module interface-specific compile options.

To build a modularized executable or library we simply list the module interfaces as its prerequisites, just as we do for source files. As an example, let's build the `hello` program that we have started in the introduction (you can find the complete project in the `cxx20-modules-examples` repository under `hello-module`). Specifically, we assume our project contains the following files:

```
// file: hello.mxx (module interface)

export module hello;

import std;

export namespace hello
{
    void say_hello (const std::string_view& name);
}

// file: hello.cxx (module implementation)

module hello;

namespace hello
{
    void say_hello (const std::string_view& n)
    {
        std::cout << "Hello, " << n << '!' << std::endl;
    }
}

// file: main.cxx

import hello;

int
main ()
{
    hello::say_hello ("World");
}
```

To build a `hello` executable from these files we can write the following `buildfile`:

```
exe{hello}: cxx{main} {mxx cxx}{hello}
```

Or, if you prefer to use wildcard patterns:

```
exe{hello}: {mxx cxx}{*}
```

Module partitions, both interface and implementation, are compiled to BMIs and as a result must be listed as `mxx{}` prerequisites. See `hello-partition` in the `cxx20-modules-examples` repository for a complete example.

Alternatively, we can place the module into a library and then link the library to the executable (see `hello-library-module` in the `cxx20-modules-examples` repository):

```
exe{hello}: cxx{main} lib{hello}
lib{hello}: {mxx cxx}{hello}
```

Note that a library consisting of only module interface units is by default always `binful` (see `Library Exportation and Versioning` for background) since compiling a module interface always results in an object file, even if the module interface does not contain any non-inline/template functions or global variables. However, you can explicitly request for such a library to be treated as `binless`:

```
lib{hello}: mxx{hello}
{
    bin.binless = true
}
```

Note that if such a `binless` library has non-inline/template functions or global variables, then whether it can be used in all situations without causing duplicate symbols is platform-dependent.

As you might have surmised from this example, the modules support in `build2` automatically resolves imports to module interface units that are specified either as direct prerequisites or as prerequisites of library prerequisites.

To perform this resolution without a significant overhead, the implementation delays the extraction of the actual module name from module interface units (since not all available module interfaces are necessarily imported by all the translation units). Instead, the implementation tries to guess which interface unit implements each module being imported based on the interface file path. Or, more precisely, a two-step resolution process is performed: first a best match between the desired module name and the file path is sought and then the actual module name is extracted and the correctness of the initial guess is verified.

The practical implication of this implementation detail is that our module interface files must embed a portion of a module name, or, more precisely, a sufficient amount of "module name tail" to unambiguously resolve all the modules used in a project. Note that this guesswork is only performed for direct module interface prerequisites; for those that come from libraries the module names are known and are therefore matched exactly. And the guesses are always

verified before the actual compilation, so misguesses cannot go unnoticed.

As an example, let's assume our `hello` project had two modules: `hello.core` and `hello.extra`. While we could call our interface files `hello.core.mxx` and `hello.extra.mxx`, respectively, this doesn't look particularly good and may be contrary to the file naming scheme used in our project. To resolve this issue the match of module names to file names is made "fuzzy": it is case-insensitive, it treats all separators (dots, dashes, underscores, etc) as equal, and it treats a case change as an imaginary separator. As a result, the following naming schemes will all match the `hello.core` module name:

```
hello-core.mxx
hello_core.mxx
HelloCore.mxx
hello/core.mxx
```

We also don't have to embed the full module name. In our case, for example, it would be most natural to call the files `core.mxx` and `extra.mxx` since they are already in the project directory called `hello/`. This will work since our module names can still be guessed correctly and unambiguously.

If a guess turns out to be incorrect, the implementation issues diagnostics and exits with an error before attempting to build anything. To resolve this situation we can either adjust the interface file names or we can specify the module name explicitly with the `cxx.module_name` variable. The latter approach can be used with interface file names that have nothing in common with module names, for example:

```
mxx{foobar}@./: cxx.module_name = hello
```

Note also that the standard library modules (`std` and `std.compat`) are treated specially and are resolved in a compiler-specific manner.

When C++ modules are enabled and available, the build system makes sure the `__cpp_modules` feature test macro is defined. However, if the compiler version being used does not claim complete modules support, its value may not be 201907.

16.3.3 Module Symbols Exporting

When building a shared library, some platforms (notably Windows) require that we explicitly export symbols that must be accessible to the library consumers. If you don't need to support such platforms, you can thank your lucky stars and skip this section.

When using headers, the traditional way of achieving this is via an "export macro" that is used to mark exported APIs, for example:

```
LIBHELLO_EXPORT void say_hello (const string&);
```

This macro is then appropriately defined (often in a separate "export header") to export symbols when building the shared library and to import them when building the library's consumers (and to nothing when either building or consuming the static library).

The introduction of modules changes this in a number of ways, at least as implemented by MSVC and Clang. While we still have to explicitly mark exported symbols in our module interface unit, there is no need (and, in fact, no way) to do the same when said module is imported. Instead, the compiler automatically treats all such explicitly exported symbols (note: symbols, not names) as imported.

While the automatic importing may look like the same mechanism as what's used to support Automatic DLL Symbol Exporting, it appears not to be since it also works for global variables, not only functions. However, reportedly, it does appear to incur the same additional overhead as auto-importing, at least for functions.

One notable aspect of this new model is the locality of the export macro: it is only defined when compiling the module interface unit and is not visible to the consumers of the module. This is unlike headers where the macro has to have a unique per-library name (that `LIBHELLO_` prefix) because a header from one library can be included while building another library.

We can continue using the same export macro and header with modules and, in fact, that's the recommended approach if maintaining the dual, header/module arrangement for backwards compatibility. However, for modules-only codebases, we have an opportunity to improve the situation in two ways: we can use a single, keyword-like macro instead of a library-specific one and we can make the build system manage it for us thus getting rid of the export header.

To enable this functionality in `build2` we set the `cxx.features.symexport` boolean variable to `true` before loading the `cxx` module. For example:

```
cxx.std = latest
cxx.features.modules = true
cxx.features.symexport = true

using cxx

...
```

Once enabled, `build2` automatically defines the `__symexport` macro to the appropriate value depending on the platform and the type of library being built. As library authors, all we have to do is use it in appropriate places in our module interface units, for example:

```
export module hello;

import std;

export __symexport void say_hello (const std::string&);
```

You may be wondering why can't a module export automatically mean a symbol export? While you will normally want to export symbols of all your module-exported names, you may also need to do so for some non-module-exported ones. For example:

```
export module foo;

__symexport void f_impl ();

export __symexport inline void f ()
{
    f_impl ();
}
```

Furthermore, symbol exporting is a murky area with many limitations and pitfalls (such as auto-exporting of base classes). As a result, it would not be unreasonable to expect such an automatic module exporting to only further muddy the matter.

16.3.4 Modules Installation

As discussed in the introduction, binary module interfaces are not a distribution mechanism and installing module interface sources appears to be the preferred approach.

Module interface units are by default installed in the same location as headers (for example, `/usr/include`). However, instead of relying on a header-like search mechanism (`-I` paths, etc.), an explicit list of exported modules is provided for each library in its `.pc` (`pkg-config`) file.

Specifically, the library's `.pc` file contains the `cxx.modules` variable that lists all the exported C++ modules in the `<name>=<path>` form with `<name>` being the module's C++ name and `<path>` – the module interface file's absolute path. For example:

```
Name: libhello
Version: 1.0.0
Cflags:
Libs: -L/usr/lib -lhello
```

```
cxx.modules = hello.core=/usr/include/hello/core.mxx hello.extra=/usr/include/hello/extra.mxx
```

The `:` character in a module partition name is encoded as `..`. For example, for `hello:core` we would have:

```
cxx.modules = hello..core=/usr/...
```

Additional module properties are specified with variables in the `cxx.module_<property>.<name>` form, for example:

```
cxx.module_symexport.hello.core = true
cxx.module_preprocessed.hello.core = all
```

Currently, two properties are defined. The `symexport` property with the boolean value signals whether the module uses the `__symexport` support discussed above.

The `preprocessed` property indicates the degree of preprocessing the module unit requires and is used to optimize module compilation. Valid values are `none` (not preprocessed), `includes` (no `#include` directives in the source), `modules` (as above plus no module declarations depend on the preprocessor, for example, `#ifdef`, etc.), and `all` (the source is fully preprocessed). Note that for `all` the source may still contain comments and line contin-

uations.

16.3.5 Modules Design Guidelines

Modules are a physical design mechanism for structuring and organizing our code. Their explicit exportation semantics combined with the way they are built make many aspects of creating and consuming modules significantly different compared to headers. This section provides basic guidelines for designing modules. We start with the overall considerations such as module granularity and partitioning into translation units then continue with the structure of typical module interface and implementation units. The following section discusses practical approaches to modularizing existing code.

Unlike headers, the cost of importing modules should be negligible. As a result, it may be tempting to create "mega-modules", for example, one per library. After all, this is how the standard library is modularized with its `std` and `std.compat` modules.

There is, however, a significant drawback to this choice: every time we make a change, all consumers of such a mega-module will have to be recompiled, whether the change affects them or not. And the bigger the module the higher the chance that any given change does not (semantically) affect a large portion of the module's consumers. Note also that this is not an issue for the standard library modules since they are not expected to change often.

Another, more subtle, issue with mega-modules (which does affect the standard library) is the inability to re-export only specific interfaces, as will be discussed below.

The other extreme in choosing module granularity is a large number of "mini-modules". Their main drawback is the tediousness of importation by the consumers.

The sensible approach is then to create modules of conceptually-related and commonly-used entities possibly complemented with aggregate modules for ease of importation. This also happens to be generally good design.

As an example, let's consider a JSON library that provides support for both parsing and serialization. Since it is common for applications to only use one of the functionalities, it makes sense to provide the `json.parser` and `json.serializer` modules. Depending on the representation of JSON we use in our library, it will most likely have some shared types so it probably makes sense to have the `json.types` module that is re-exported by the parser and serializer modules. While it is not too tedious to import both `json.parser` and `json.serializer` if both are needed, for convenience we could also provide the `json` module that re-exports the two. Something along these lines:

```
// types.mxx

export module json.types;

export class json
{
    ...
};
```

```

// parser.mxx

export module json.parser;

export import json.types;

export json parse (...);

// serializer.mxx

export module json.serializer;

export import json.types;

export ... serialize (const json&);

// json.mxx

export module json;

export import json.types;
export import json.parser;
export import json.serializer;

```

Once we are past selecting an appropriate granularity for our modules, the next question is how to partition them into translation units. A module can consist of just the interface unit and, as discussed above, such a unit can contain anything an implementation unit can, including non-inline function definitions. Some may then view this as an opportunity to get rid of the header/source separation and have everything in a single file.

There are a number of drawbacks with this approach: Every time we change anything in the module interface unit, all its consumers have to be recompiled. If we keep everything in a single file, then every time we change the implementation we trigger recompilations that would have been avoided had the implementation been factored out into a separate unit. Note that a build system in cooperation with the compiler could theoretically avoid such unnecessary recompilations in certain cases: if the compiler produces identical binary interface files when the module interface is unchanged, then the build system could detect this and skip recompiling the module's consumers.

A related issue with single-file modules is the reduction in the build parallelization opportunities. If the implementation is part of the interface unit, then the build system cannot start compiling the module's consumers until both the interface and the implementation are compiled. On the other hand, had the implementation been split into a separate file, the build system could start compiling the module's consumers (as well as the implementation unit) as soon as the module interface is compiled.

Another issues with combining the interface with the implementation is the readability of the interface which could be significantly reduced if littered with implementation details. We could keep the interface separate by moving the implementation to the bottom of the interface file but then we might as well move it into a separate file and avoid the unnecessary recompilations or parallelization issues.

The sensible guideline is then to have a separate module implementation unit except perhaps for modules with a simple implementation that is mostly inline/template. Note that more complex modules may have several implementation units, however, based on our granularity guideline, those should be rare.

Once we start writing our first real module the immediate question that normally comes up is where to put `#include` directives and `import` declarations and in what order. To recap, a module unit, both interface and implementation, is split into two parts: before the module declaration, called the global module fragment, which obeys the usual or "old" translation unit rules and after the module declaration which is the module purview. Inside the module purview all declarations have their symbols invisible to any other module (including the global module). With this understanding, consider the following module interface:

```
export module hello;

#include <string>
```

Do you see the problem? We have included `<string>` in the module purview which means all its names (as well as all the names in any headers it might include, recursively) are now declared as having the `hello` module linkage. The result of doing this can range from silent code blot to strange-looking unresolved symbols.

The guideline this leads to should be clear: including a header in the module purview is almost always a bad idea. There are, however, a few types of headers that may make sense to include in the module purview. The first are headers that only define preprocessor macros, for example, configuration or export headers. There are also cases where we do want the included declarations to end up in the module purview. The most common example is inline/template function implementations that have been factored out into separate files for code organization reasons. As an example, consider the following module interface that uses an export header (which presumably sets up symbols exporting macros) as well as an inline file:

```
module;

#include <string>

export module hello;

#include <libhello/export.hxx>

export namespace hello
{
    ...
}

#include <libhello/hello.ixx>
```

A note on inline/template files: in header-based projects we could include additional headers in those files, for example, if the included declarations are only needed in the implementation. For the reasons just discussed, this does not work with modules and we have to move all the includes into the interface file, into the global module fragment. On the other hand, with modules, it is safe to use namespace-level using-directives (for example, `using names-`

pace `std;`) in inline/template files (and, with care, even in the interface file).

What about imports, where should we import other modules? Again, to recap, unlike a header inclusion, an `import` declaration only makes exported names visible without redeclaring them. As result, in module implementation units, it doesn't really matter where we place imports, in the module purview or the global module fragment. There are, however, two differences when it comes to module interface units: only imports in the purview are visible to implementation units and we can only re-export an imported module from the purview.

The guideline is then for interface units to import in the module purview unless there is a good reason not to make the import visible to the implementation units. And for implementation units to always import in the purview for simplicity. For example:

```
module;

#include <cassert>

export module hello;

import std;

#include <libhello/export.hxx>

export namespace hello
{
    ...
}

#include <libhello/hello.ixx>
```

By putting all these guidelines together we can then create a module interface unit template:

```
// Module interface unit.

module;                                // Start of global module fragment.

<header includes>

export module <name>;                  // Start of module purview.

<module imports>

<special header includes> // Configuration, export, etc.

<module interface>

<inline/template includes>
```

As well as the module implementation unit template:

```
// Module implementation unit.

module;                                // Start of global module fragment.

<header includes>

module <name>;                          // Start of module purview.

<extra module imports>                 // Only additional to interface.

<module implementation>
```

Let's now discuss module naming. Module names are in a separate "name plane" and do not collide with namespace, type, or function names. Also, as mentioned earlier, the standard does not assign a hierarchical meaning to module names though it is customary to assume module `hello.core` is a submodule of `hello` and, unless stated explicitly otherwise, importing the latter also imports the former.

It is important to choose good names for public modules (that is, modules packaged into libraries and used by a wide range of consumers) since changing them later can be costly. We have more leeway with naming private modules (that is, the ones used by programs or internal to libraries) though it's worth coming up with a consistent naming scheme here as well.

The general guideline is to start names of public modules with the library's namespace name followed by a name describing the module's functionality. In particular, if a module is dedicated to a single class (or, more generally, has a single primary entity), then it makes sense to use that name as the module name's last component.

As a concrete example, consider `libbutl` (the `build2` utility library): All its components are in the `butl` namespace so all its module names start with `butl`. One of its components is the `small_vector` class template which resides in its own module called `butl.small_vector`. Another component is a collection of string parsing utilities that are grouped into the `butl::string_parser` namespace with the corresponding module called `butl.string_parser`.

When is it a good idea to re-export a module? The two straightforward cases are when we are building an aggregate module out of submodules, for example, `json` out of `json.parser` and `json.serializer`, or when one module extends or supersedes another, for example, as `json.parser` extends `json.types`. It is also clear that there is no need to re-export a module that we only use in the implementation. The case when we use a module in our interface is, however, a lot less clear cut.

But before considering the last case in more detail, let's understand the issue with re-export. In other words, why not simply re-export any module we import in our interface? In essence, re-export implicitly injects another module import anywhere our module is imported. If we re-export `std` then consumers of our module will also automatically "see" all the names exported by `std`. They can then start using names from `std` without explicitly importing `std` and everything will compile until one day they no longer need to import our module or we no longer need to import `std`. In a sense, re-export becomes part of our interface and it is generally good design to keep interfaces minimal.

And so, at the outset, the guideline is then to only re-export the minimum necessary.

Let's now discuss a few concrete examples to get a sense of when re-export might or might not be appropriate. Unfortunately, there does not seem to be a hard and fast rule and instead one has to rely on their good sense of design.

To start, let's consider a simple module that uses `std::string` in its interface:

```
export module hello;

import std;

export namespace hello
{
    std::string format_hello (const std::string&);
}
```

Should we re-export `std` in this case? Most likely not. If consumers of our module want to refer to `std::string`, then it is natural to expect them to explicitly import the necessary module. In a sense, this is analogous to scoping: nobody expects to be able to use just `string` (without `std::`) because of using `namespace hello`;

So it seems that a mere usage of a name in an interface does not generally warrant a re-export. The fact that a consumer may not even use this part of our interface further supports this conclusion.

Let's now consider a more interesting case (inspired by real events):

```
export module small_vector;

import std;

template <typename T, std::size_t N>
export class small_vector: public std::vector<T, ...>
{
    ...
};
```

Here we have the `small_vector` container implemented in terms of `std::vector` by providing a custom allocator and with most of the functions derived as is. Consider now this innocent-looking consumer code:

```
import small_vector;

small_vector<int, 1> a, b;

if (a == b) // Error.
    ...
```

We don't reference `std::vector` directly so presumably we shouldn't need to import its module. However, the comparison won't compile: our `small_vector` implementation re-uses the comparison operators provided by `std::vector` (via implicit to-base conversion) but they aren't visible.

There is a palpable difference between the two cases: the first merely uses `std` interface while the second is *based on* and, in a sense, *extends* it which feels like a stronger relationship. Re-exporting `std` (or, better yet, `std::vector`, if it were available) seems less unreasonable.

Note also that there is no re-export of headers nor header inclusion visibility in the implementation units. Specifically, in the previous example, if the standard library is not modularized and we have to use it via headers, then the consumers of our `small_vector` will always have to explicitly include `<vector>`. This suggests that modularizing a codebase that still consumes substantial components (like the standard library) via headers can incur some development overhead compared to the old, headers-only approach.

16.3.6 Modularizing Existing Code

The aim of this section is to provide practical guidelines to modularizing existing codebases.

Predictably, a well modularized (in the general sense) set of headers makes conversion to C++ modules easier. Inclusion cycles will be particularly hard to deal with (C++ modules do not allow circular interface dependencies). Having a one-to-one header to module mapping will simplify this task. As a result, it may make sense to spend some time cleaning and re-organizing your headers prior to attempting modularization.

The recommended strategy for modularizing our own components is to identify and modularize inter-dependent sets of headers one at a time starting from the lower-level components. This way any newly modularized set will only depend on the already modularized ones. After converting each set we can switch its consumers to using imports keeping our entire project buildable and usable.

While ideally we would want to be able to modularize just a single component at a time, this does not seem to work in practice because we will have to continue consuming some of the components as headers. Since such headers can only be imported out of the module purview, it becomes hard to reason (both for us and often the compiler) what is imported/included and where. For example, it's not uncommon to end up importing the module in its implementation unit which is not something that all the compilers can handle gracefully.

If our module needs to "export" macros then the recommended approach is to simply provide an additional header that the consumer includes. While it might be tempting to also wrap the module import into this header, some may prefer to explicitly import the module and include the header, especially if the macros may not be needed by all consumers. This way we can also keep the header macro-only which means it can be included freely, in or out of module purviews.

16.4 Objective-C++ Compilation

The `cxx` module provides the `cxx.objcxx` submodule which can be loaded in order to register the `mm{ }` target type and enable Objective-C++ compilation in the C++ compile rule. Note that `cxx.objcxx` must be loaded after the `cxx` module and while the `mm{ }` target type is registered unconditionally, compilation is only enabled if the C++ compiler supports Objective-C++ for the target platform. Typical usage:

```
# root.build
#
using cxx
using cxx.objcxx

# buildfile
#
lib{hello}: {hxx cxx}{*}
lib{hello}: mm{*}: include = ($cxx.target.class == 'macos')
```

Note also that while there is support for linking Objective-C++ executables and libraries, this is done using the C++ compiler driver and no attempt is made to automatically link any necessary Objective-C runtime library (such as `-lobjc`).

16.5 C++ Compiler Predefined Macro Extraction

The `cxx` module provides the `cxx.predefs` submodule which can be loaded in order to register a rule that generates a C++ header with predefined compiler macros. Note that the `cxx.predefs` module must be loaded after the `cxx` module and the rule will only match with an explicit rule hint. Typical usage:

```
# root.build
#
using cxx
using cxx.predefs

# buildfile
#
[rule_hint=cxx.predefs] hxx{predefs}:
```

See Compiler Predefined Macro Extraction for details.

17 in Module

The `in` build system module provides support for `.in` (input) file preprocessing. Specifically, the `.in` file can contain a number of *substitutions* – build system variable names enclosed with the substitution symbol (`$` by default) – which are replaced with the corresponding variable values to produce the output file. For example:

```
# build/root.build

using in
```

```
// config.hxx.in

#define TARGET "$cxx.target$"

# buildfile

hxx{config}: in{config}
```

The `in` module defines the `in{ }` target type and implements the `in` build system rule.

While we can specify the `.in` extension explicitly, it is not necessary because the `in{ }` target type implements *target-dependent search* by taking into account the target it is a prerequisite of. In other words, the following dependency declarations produce the same result:

```
hxx{config}:      in{config}
hxx{config.hxx}: in{config}
hxx{config.hxx}: in{config.hxx.in}
```

By default the `in` rule uses `$` as the substitution symbol. This can be changed using the `in.symbol` variable. For example:

```
// data.cxx.in

const char data[] = "@data@";

# buildfile

cxx{data}: in{data}
{
    in.symbol = '@'
    data = 'Hello, World!'
}
```

Note that the substitution symbol must be a single character.

The default substitution mode is strict. In this mode every substitution symbol is expected to start a substitution with unresolved (to a variable value) names treated as errors. The double substitution symbol (for example, `$$`) serves as an escape sequence.

The substitution mode can be relaxed using the `in.mode` variable. Its valid values are `strict` (default) and `lax`. In the `lax` mode a pair of substitution symbols is only treated as a substitution if what's between them looks like a build system variable name (that is, it doesn't contain spaces, etc). Everything else, including unterminated substitution symbols, is copied as is. Note also that in this mode the double substitution symbol is not treated as an escape sequence.

The `lax` mode is mostly useful when trying to reuse existing `.in` files from other build systems, such as `autoconf`. Note, however, that the `lax` mode is still stricter than `autoconf`'s semantics which also leaves unresolved substitutions as is. For example:

```
# buildfile

h{config}: in{config} # config.h.in
{
    in.symbol = '@'
    in.mode = lax

    CMAKE_SYSTEM_NAME = $c.target.system
    CMAKE_SYSTEM_PROCESSOR = $c.target.cpu
}
```

The `in` rule tracks changes to the input file as well as the substituted variable values and automatically regenerates the output file if any were detected. Substituted variable values are looked up starting from the target-specific variables. Typed variable values are converted to string using the corresponding `builtin.string()` function overload before substitution.

While specifying substitution values as `buildfile` variables is usually natural, sometimes this may not be possible or convenient. Specifically, we may have substitution names that cannot be specified as `buildfile` variables, for example, because they start with an underscore (and are thus reserved) or because they refer to one of the predefined variables. Also, we may need to have different groups of substitution values for different cases, for example, for different platforms, and it would be convenient to pass such groups around as a single value.

To support these requirements the substitution values can alternatively be specified as key-value pairs in the `in.substitutions` variable. Note that entries in this substitution map take precedence over the `buildfile` variables. For example:

```
/* config.h.in */

#define _GNU_SOURCE    @_GNU_SOURCE@
#define _POSIX_SOURCE  @_POSIX_SOURCE@

# buildfile

h{config}: in{config}
{
    in.symbol = '@'
    in.mode = lax

    in.substitutions = _GNU_SOURCE@0 _POSIX_SOURCE@1
}
```

In the above example, the `@` characters in `in.symbol` and `in.substitutions` are unrelated.

Using an undefined variable in a substitution is an error. Using a `null` value in a substitution is also an error unless the fallback value is specified with the `in.null` variable. For example:

```
# buildfile

h{config}: in{config}
{
    in.null = '' # Substitute null values with empty string.
}
```

To specify a null value using the `in.substitutions` mechanism omit the value, for example:

```
in.substitutions = _GNU_SOURCE
```

A number of other build system modules, for example, `autoconf`, `version`, and `bash`, are based on the `in` module and provide extended functionality. The `in` preprocessing rule matches any `file{}`-based target that has the corresponding `in{}` prerequisite provided none of the extended rules match.

18 bash Module

The `bash` build system module provides modularization support for `bash` scripts. It is based on the `in` build system module and extends its preprocessing rule with support for *import substitutions* in the `@import <module>@` form. During preprocessing, such imports are replaced with suitable `source` builtin calls. For example:

```
# build/root.build

using bash

# hello/say-hello.bash

function say_hello ()
{
    echo "Hello, $1!"
}

#!/usr/bin/env bash

# hello/hello.in

@import hello/say-hello@

say_hello 'World'

# hello/buildfile

exe{hello}: in{hello} bash{say-hello}
```

By default the `bash` preprocessing rule uses the lax substitution mode and `@` as the substitution symbol but this can be overridden using the standard `in` module mechanisms.

In the above example, `say-hello.bash` is a *module*. By convention, `bash` modules have the `.bash` extension and we use the `bash{ }` target type (defined by the `bash` build system module) to refer to them in buildfiles.

The `say-hello.bash` module is *imported* by the `hello` script with the `@import hello/say-hello@` substitution. The *import path* (`hello/say-hello` in our case) is a path to the module file within the project. Its first component (`hello` in our case) must be both the project name and the top-level subdirectory within the project. The `.bash` module extension can be omitted. The constraint placed on the first component of the

import path is required to implement importation of installed modules, as discussed below.

During preprocessing, the import substitution will be replaced with a `source` builtin call and the import path resolved to one of the `bash{}` prerequisites from the script's dependency declaration. The actual module path used in `source` depends on whether the script is preprocessed for installation. If it's not (development build), then the absolute path to the module file is used. Otherwise, a path relative to the sourcing script's directory is derived. This allows installed scripts and their modules to be moved around.

The derivation of the sourcing script's directory works even if the script is executed via a symbolic link from another directory. Implementing this, however, requires `readlink(1)` with support for the `-f` option. One notable platform that does not provide such `readlink(1)` by default is Mac OS. The script, however, can provide a suitable implementation as a function. See the bash module tests for a sample implementation of such a function.

By default, bash modules are installed into a subdirectory of the `bin/` installation directory named as the project name plus the `.bash` extension. For instance, in the above example, the script will be installed as `bin/hello` and the module as `bin/hello.bash/say-hello.bash` with the script sourcing the module relative to the `bin/` directory. Note that currently it is assumed the script and all its modules are installed into the same `bin/` directory.

Naturally, modules can import other modules and modules can be packaged into *module libraries* and imported using the standard build system import mechanism. For example, we could factor the `say-hello.bash` module into a separate `libhello` project:

```
# build/export.build

$out_root/
{
    include libhello/
}

export $src_root/libhello/$import.target

# libhello/say-hello.bash

function hello_say_hello ()
{
    echo "Hello, $1!"
}
```

And then import it in a module of our `hello` project:

```
# hello/hello-world.bash.in

@import libhello/say-hello@

function hello_world ()
{
    hello_say_hello 'World'
}
```

```
#!/usr/bin/env bash

# hello/hello.in

@import hello/hello-world@

hello_world

# hello/buildfile

import mods = libhello%bash{say-hello}

exe{hello}:      in{hello}      bash{hello-world}
bash{hello-world}: in{hello-world} $mods
```

The bash preprocessing rule also supports importation of installed modules by searching in the PATH environment variable.

By convention, bash module libraries should use the `lib` name prefix, for example, `libhello`. If there is also a native library (that is, one written in C/C++) that provides the same functionality (or the bash library is a language binding for the said library), then it is customary to add the `.bash` extension to the bash library name, for example, `libhello.bash`. Note that in this case the top-level subdirectory within the project is expected to be called without the bash extension, for example, `libhello`.

Modules can be *private* or *public*. Private modules are implementation details of a specific project and are not expected to be imported from other projects. The `hello/hello-world.bash.in` module above is an example of a private module. Public modules are meant to be used by other projects and are normally packaged into libraries, like the `libhello/say-hello.bash` module above.

Public modules must take care to avoid name clashes. Since bash does not have a notion of namespaces, the recommended way is to prefix all module functions (and global variables, if any) with the library name (without the `lib` prefix), like in the `libhello/say-hello.bash` module above.

While using such decorated function names can be unwieldy, it is relatively easy to create wrappers with shorter names and use those instead. For example:

```
@import libhello/say-hello@

function say_hello () { hello_say_hello "$@"; }
```

A module should normally also prevent itself from being sourced multiple times. The recommended way to achieve this is to begin the module with a *source guard*. For example:

```
# libhello/say-hello.bash

if [ "$hello_say_hello" ]; then
    return 0
else
    hello_say_hello=true
fi

function hello_say_hello ()
{
    echo "Hello, $1!"
}
```

The `bash` preprocessing rule matches `exe{ }` targets that have the corresponding `in{ }` and one or more `bash{ }` prerequisites as well as `bash{ }` targets that have the corresponding `in{ }` prerequisite (if you need to preprocess a script that does not depend on any modules, you can use the `in` module's rule).

19 Appendix A – JSON Dump Format

This appendix describes the machine-readable, JSON-based build system state dump format that can be requested with the `--dump-format=json-v0.1` build system driver option (see **b (1)** for details).

The format is specified in terms of the serialized representation of C++ `struct` instances. See JSON OUTPUT for details on the overall properties of this format and the semantics of the `struct` serialization.

This format is currently unstable (thus the temporary `-v0.1` suffix) and may be changed in ways other than as described in JSON OUTPUT. In case of such changes the format version will be incremented to allow detecting incompatibilities but no support for older versions is guaranteed.

The build system state can be dumped after the load phase (`--dump=load`), once the build state has been loaded, and/or after the match phase (`--dump=match`), after rules have been matched to targets to execute the desired action. The JSON format differs depending on after which phase it is produced. After the load phase the format aims to describe the action-independent state, essentially as specified in the `buildfiles`. While after the match phase it aims to describe the state for executing the specified action, as determined by the rules that have been matched. The former state would be more appropriate, for example, for an IDE that tries to use `buildfiles` as project files. While the latter state could be used to determine the actual build graph for a certain action, for example, in order to infer which executable targets are considered tests by the `test` operation.

While it's possible to dump the build state as a byproduct of executing an action (for example, performing an update), it's often desirable to only dump the build state and do it as quickly as possible. For such cases the recommended option combinations are as follows (see the `--load-only` and `--match-only` documentation for details):

```
$ b --load-only --dump=load --dump-format=json-v0.1 .../dir/
$ b --match-only --dump=match --dump-format=json-v0.1 .../dir/
$ b --match-only --dump=match --dump-format=json-v0.1 .../dir/type{name}
```

Note that a match dump for a large project can produce a large amount of data, especially for the update operation (tens and even hundreds of megabytes is not uncommon). To reduce this size it is possible to limit the dump to specific scopes and/or targets with the `--dump-scope` and `--dump-target` options.

The complete dump (that is, not of a specific scope or target) is a tree of nested scope objects (see [Output Directories and Scopes](#) for background). The scope object has the serialized representation of the following C++ struct scope. It is the same for both load and match dumps except for the type of the `targets` member:

```
struct scope
{
    string          out_path;
    optional<string> src_path;

    vector<variable> variables; // Non-type/pattern scope variables.

    vector<scope> scopes; // Immediate children.

    vector<loaded_target|matched_target> targets;
};
```

For example (parts of the output are omitted for brevity):

The actual output is produced unindented to reduce the size.

```
$ cd /tmp
$ bdep new hello
$ cd hello
$ bdep new -C @gcc cc
$ b --load-only --dump=load --dump-format=json-v0.1
{
  "out_path": "",
  "variables": [ ... ],
  "scopes": [
    {
      "out_path": "/tmp/hello-gcc",
      "variables": [ ... ],
      "scopes": [
        {
          "out_path": "hello",
          "src_path": "/tmp/hello",
          "variables": [ ... ],
          "scopes": [
            {
              "out_path": "hello",
              "src_path": "/tmp/hello/hello",
              "variables": [ ... ],
              "targets": [ ... ]
            }
          ],
          "targets": [ ... ]
        }
      ],
      "targets": [ ... ]
    }
  ],
  "targets": [ ... ]
}
```

```

    ],
    "targets": [ ... ]
  }
]
}

```

The `out_path` member is relative to the parent scope. It is empty for the special global scope, which is the root of the tree. The `src_path` member is absent if it is the same as `out_path` (in source build or scope outside of project).

For the match dump, targets that have not been matched for the specified action are omitted.

In the load dump, the target object has the serialized representation of the following C++ struct `loaded_target`:

```

struct loaded_target
{
    string          name;    // Relative quoted/qualified name.
    string  display_name;    // Relative display name.
    string          type;    // Target type.
    optional<string> group;   // Absolute quoted/qualified group target.

    vector<variable> variables; // Target variables.

    vector<prerequisite> prerequisites;
};

```

For example (continuing with the previous `hello` setup):

```

{
  "out_path": "",
  "scopes": [
    {
      "out_path": "/tmp/hello-gcc",
      "scopes": [
        {
          "out_path": "hello",
          "src_path": "/tmp/hello",
          "scopes": [
            {
              "out_path": "hello",
              "src_path": "/tmp/hello/hello",
              "targets": [
                {
                  "name": "exe{hello}",
                  "display_name": "exe{hello}",
                  "type": "exe",
                  "prerequisites": [
                    {
                      "name": "cxx{hello}",
                      "type": "cxx"
                    },
                    {
                      "name": "testscript{testscript}",
                      "type": "testscript"
                    }
                  ]
                }
              ]
            }
          ]
        }
      ]
    }
  ]
}

```

```

    }
  ]
}
]
}
]
}

```

The target name member is the target name that is qualified with the extension (if applicable and known) and, if required, is quoted so that it can be passed back to the build system driver on the command line. The `display_name` member is unqualified and unquoted. Note that both the target name and `display_name` members are normally relative to the containing scope (if any).

The prerequisite object has the serialized representation of the following C++ struct `prerequisite`:

```

struct prerequisite
{
    string name; // Quoted/qualified name.
    string type;
    vector<variable> variables; // Prerequisite variables.
};

```

The prerequisite name member is normally relative to the containing scope.

In the match dump, the target object has the serialized representation of the following C++ struct `matched_target`:

```

struct matched_target
{
    string          name;
    string  display_name;
    string          type;
    optional<string> group;

    optional<path> path; // Absent if not path target, not assigned.

    vector<variable> variables;

    optional<operation_state> outer_operation; // null if not matched.
    operation_state          inner_operation; // null if not matched.
};

```

For example (outer scopes removed for brevity):

```

$ b --match-only --dump=match --dump-format=json-v0.1
{
  "out_path": "hello",
  "src_path": "/tmp/hello/hello",
  "targets": [
    {
      "name": "/tmp/hello/hello/cxx{hello.cxx}@./",
      "display_name": "/tmp/hello/hello/cxx{hello}@./",
      "type": "cxx",
      "path": "/tmp/hello/hello/hello.cxx",
      "inner_operation": {

```

```

        "rule": "build.file",
        "state": "unchanged"
    }
},
{
    "name": "obje{hello.o}",
    "display_name": "obje{hello}",
    "type": "obje",
    "group": "/tmp/hello-gcc/hello/hello/obj{hello}",
    "path": "/tmp/hello-gcc/hello/hello/hello.o",
    "inner_operation": {
        "rule": "cxx.compile",
        "prerequisite_targets": [
            {
                "name": "/tmp/hello/hello/cxx{hello.cxx}@./",
                "type": "cxx"
            },
            {
                "name": "/usr/include/c++/12/h{iostream.}",
                "type": "h"
            },
            ...
        ]
    }
},
{
    "name": "exe{hello.}",
    "display_name": "exe{hello}",
    "type": "exe",
    "path": "/tmp/hello-gcc/hello/hello/hello",
    "inner_operation": {
        "rule": "cxx.link",
        "prerequisite_targets": [
            {
                "name": "/tmp/hello-gcc/hello/hello/obje{hello.o}",
                "type": "obje"
            }
        ]
    }
}
]
}

```

The first four members in `matched_target` have the same semantics as in `loaded_target`.

The `outer_operation` member is only present if the action has an outer operation. For example, when performing `update-for-test`, `test` is the outer operation while `update` is the inner operation.

The operation state object has the serialized representation of the following C++ struct `operation_state`:

```

struct operation_state
{
    string rule; // null if direct recipe match.

    optional<string> state; // One of unchanged|changed|group.

    vector<variable> variables; // Rule variables.

    vector<prerequisite_target> prerequisite_targets;
};

```

The `rule` member is the matched rule name. The `state` member is the target state, if known after match. The `prerequisite_targets` array is a subset of prerequisites resolved to targets that are in effect for this action. The matched rule may add additional targets, for example, dynamically extracted additional dependencies, like `/usr/include/c++/12/h<iostream.>` in the above listing.

The prerequisite target object has the serialized representation of the following C++ struct `prerequisite_target`:

```

struct prerequisite_target
{
    string name; // Absolute quoted/qualified target name.
    string type;
    bool adhoc;
};

```

The `variables` array in the scope, target, prerequisite, and prerequisite target objects contains scope, target, prerequisite, and rule variables, respectively.

The variable object has the serialized representation of the following C++ struct `variable`:

```

struct variable
{
    string          name;
    optional<string> type;
    json_value      value; // null|boolean|number|string|object|array
};

```

For example:

```

{
    "out_path": "",
    "variables": [
        {
            "name": "build.show_progress",
            "type": "bool",
            "value": true
        },
        {
            "name": "build.verbosity",
            "type": "uint64",
            "value": 1
        },
        ...
    ],
}

```

```

"scopes": [
{
  "out_path": "/tmp/hello-gcc",
  "scopes": [
    {
      "out_path": "hello",
      "src_path": "/tmp/hello",
      "scopes": [
        {
          "out_path": "hello",
          "src_path": "/tmp/hello/hello",
          "variables": [
            {
              "name": "out_base",
              "type": "dir_path",
              "value": "/tmp/hello-gcc/hello/hello"
            },
            {
              "name": "src_base",
              "type": "dir_path",
              "value": "/tmp/hello/hello"
            },
            {
              "name": "cxx.poptions",
              "type": "strings",
              "value": [
                "-I/tmp/hello-gcc/hello",
                "-I/tmp/hello"
              ]
            },
            {
              "name": "libs",
              "value": "/tmp/hello-gcc/libhello/libhello/lib{hello}"
            }
          ]
        }
      ]
    }
  ]
}
]
}
]
}
]
}
}

```

The `type` member is absent if the variable value is untyped.

The `value` member contains the variable value in a suitable JSON representation. Specifically:

- `null` values are represented as JSON `null`.
- `bool` values are represented as JSON `boolean`.
- `int64` and `uint64` values are represented as JSON `number`.
- `string`, `path`, `dir_path` values are represented as JSON `string`.
- Untyped simple name values are represented as JSON `string`.
- Pairs of above values are represented as JSON objects with the `first` and `second` members corresponding to the pair elements.
- Untyped complex name values are serialized as target names and represented as JSON `string`.

- Containers of above values are represented as JSON arrays corresponding to the container elements.
- An empty value is represented as an empty JSON object if it's a typed pair, as an empty JSON array if it's a typed container or is untyped, and as an empty string otherwise.

One expected use-case for the match dump is to determine the set of targets for which a given action is applicable. For example, we may want to determine all the executables in a project that can be tested with the `test` operation in order to present this list to the user in an IDE plugin or some such. To further illuminate the problem, consider the following `buildfile` which declares a number of executable targets, some are tests and some are not:

```
exe{hello1}: ... testscript # Test because of testscript prerequisite.

exe{hello2}: test = true    # Test because of test=true.

exe{hello3}: ... testscript # Not a test because of test=false.
{
  test = false
}
```

As can be seen, trying to infer this information is not straightforward and doing so manually by examining prerequisites, variables, etc., while possible, will be complex and likely brittle. Instead, the recommended approach is to use the match dump and base the decision on the `state` target object member. Specifically, a rule which matched the target but determined that nothing needs to be done for this target, returns the special `noop` recipe. The `build2` core recognizes this situation and sets such target's state to `unchanged` during match. Here is what the match dump will look like for the above three executables:

```
$ b --match-only --dump=match --dump-format=json-v0.1 test
{
  "out_path": "hello",
  "src_path": "/tmp/hello/hello",
  "targets": [
    {
      "name": "exe{hello1.}",
      "display_name": "exe{hello1}",
      "type": "exe",
      "path": "/tmp/hello-gcc/hello/hello/hello1",
      "inner_operation": {
        "rule": "test"
      }
    },
    {
      "name": "exe{hello2.}",
      "display_name": "exe{hello2}",
      "type": "exe",
      "path": "/tmp/hello-gcc/hello/hello/hello2",
      "inner_operation": {
        "rule": "test"
      }
    },
    {
      "name": "exe{hello3}",
      "display_name": "exe{hello3}",
      "type": "exe",
      "inner_operation": {
```

```
        "rule": "test",  
        "state": "unchanged"  
    }  
]  
}
```