

The `build2` Package Manager

Copyright © 2014-2019 Code Synthesis Ltd

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

Revision 0.10, March 2019

This revision of the document describes the `build2` package manager 0.10.x series.

Table of Contents

Preface	1
1 Package Name	1
2 Package Version	1
3 Manifests	5
3.1 Manifest Format	5
3.2 Package Manifest	9
3.2.1 name	10
3.2.2 version	10
3.2.3 project	10
3.2.4 priority	10
3.2.5 summary	11
3.2.6 license	11
3.2.7 tags	12
3.2.8 description	12
3.2.9 changes	12
3.2.10 url	13
3.2.11 doc-url	13
3.2.12 src-url	13
3.2.13 package-url	13
3.2.14 email	14
3.2.15 package-email	14
3.2.16 build-email	14
3.2.17 build-warning-email	14
3.2.18 build-error-email	14
3.2.19 depends	14
3.2.20 requires	17
3.2.21 builds	18
3.2.22 build-{include,exclude}	20
3.3 Package List Manifest for pkg Repositories	20
3.3.1 sha256sum (list manifest)	21
3.3.2 location (package manifest)	21
3.3.3 sha256sum (package manifest)	21
3.4 Package List Manifest for dir Repositories	22
3.4.1 location	22
3.4.2 fragment	22
3.5 Repository Manifest	22
3.5.1 location	23
3.5.2 type	23
3.5.3 role	24
3.5.4 trust	24

3.5.5 url	24
3.5.6 email	25
3.5.7 summary	25
3.5.8 description	25
3.5.9 certificate	26
3.5.10 fragment	26
3.6 Repository List Manifest	26
3.7 Signature Manifest for pkg Repositories	27
3.7.1 sha256sum	27
3.7.2 signature	28

Preface

This document describes `bpkg`, the `build2` package dependency manager. For the package manager command line interface refer to the **`bpkg(1)`** man pages.

1 Package Name

The `bpkg` package name can contain ASCII alphabetic characters (`[a-zA-Z]`), digits (`[0-9]`), underscores (`_`), plus/minus (`+/-`), and dots/periods (`.`). The name must be at least two characters long with the following additional restrictions:

1. It must start with an alphabetic character.
2. It must end with an alphabetic, digit, or plus character.
3. It must not be any of the following illegal names:

```
build
con prn aux nul
com1 com2 com3 com4 com5 com6 com7 com8 com9
lpt1 lpt2 lpt3 lpt4 lpt5 lpt6 lpt7 lpt8 lpt9
```

The use of the plus (`+`) character in package names is discouraged. Pluses are used in URL encoding which makes specifying packages that contain pluses in URLs cumbersome.

The use of the dot (`.`) character in package names is discouraged except for distinguishing the implementations of the same functionality for different languages. For example, `libfoo` and `libfoo.bash`.

Package name comparison is case-insensitive but the original case must be preserved for display, in file names, etc. The reason for case-insensitive comparison is Windows file names.

If the package is a library then it is strongly recommended that you start its package name with the `lib` prefix, for example, `libfoo`. Some package repositories may make this a requirement as part of their submission policy.

If a package (normally a library) supports usage of multiple major versions in the same project, then it is recommended to append the major version number to the package name starting from version `2.0.0`, for example, `libfoo` (before `2.0.0`), `libfoo2`, `libfoo3` (`3.Y.Z`), etc.

2 Package Version

The `bpkg` package version format tries to balance the need of accommodating existing software versions on one hand and providing a reasonably straightforward comparison semantics on another. For some background on this problem see **`deb-version(1)`** and the Semantic

Versioning specification.

Note also that if you are starting a new project that will use the `build2` toolchain, then it is strongly recommended that you use the *standard versioning* scheme which is a more strictly defined subset of semantic versioning and that allows automation of many version management tasks. See `version` Module for details.

The `bpkg` package version has the following form:

```
[+<epoch>-]<upstream>[-<prerelease>][+<revision>][#<iteration>]
```

The *epoch* part should be an integer. It can be used to change to a new versioning scheme that would be incompatible with the old one. If not specified, then *epoch* defaults to 1 except for a stub version (see below) in which case it defaults to 0. The explicit zero *epoch* can be used if the current versioning scheme (for example, date-based) is known to be temporary.

The *upstream* part is the upstream software version that this package is based on. It can only contain alpha-numeric characters and `.`. The `.` character is used to separate the version into *components*.

The *prerelease* part is the upstream software pre-release marker, for example, alpha, beta, candidate, etc. Its format is the same as for *upstream* except for two special values: the absent *prerelease* (for example, `1.2.3`) signifies the maximum or final release while the empty *prerelease* (for example, `1.2.3-`) signifies the minimum or earliest possible release. The minimum release is intended to be used for version constraints (for example, `libfoo < 1.2.3-`) rather than actual releases.

The *revision* part should be an integer. It is used to version package releases that are based on the same upstream versions. If not specified, then *revision* defaults to 0.

The *iteration* part is an integer. It is used internally by `bpkg` to automatically version modifications to the packaging information (specifically, to package manifest and lockfile) in *external packages* that have the same upstream version and revision. As a result, the *iteration* cannot not be specified by the user and is only shown in the `bpkg` output (for example, by `pkg-status` command) in order to distinguish between package iterations with otherwise identical versions. Note also that *iteration* is relative to the `bpkg` configuration. Or, in other words, it is an iteration number of a package as observed by a specific configuration. As a result, two configuration can "see" the same package state as two different iterations.

Package iterations are used to support package development during which requiring the developer to manually increment the version or revision after each modification would be impractical. This mechanism is similar to the automatic commit versioning provided by the *standard version* except that it is limited to the packaging information but works for uncommitted changes.

Version `+0-0-` (least possible version) is reserved and specifying it explicitly is illegal. Explicitly specifying this version does not make much sense since `libfoo < +0-0-` is always false and `libfoo > +0-0-` is always true. In the implementation this value is used as a special empty version.

Version 0 (with a potential revision, for example, `0+1`, `0+2`) is used to signify a *stub package*. A stub is a package that does not contain source code and can only be "obtained" from other sources, for example, a system package manager. Note that at some point a stub may be converted into a full-fledged package at which point it will be assigned a "real" version. It is assumed that this version will always be greater than the stub version.

When displaying the package version or when using the version to derive the file name, the default *epoch* value as well as zero *revision* and *iteration* values are omitted (even if they were explicitly specified, for instance, in the package manifest). For example, `+1-1.2.3+0` will be used as `libfoo-1.2.3`.

This versioning scheme and the choice of delimiter characters (`.-+`) is meant to align with semantic versioning.

Some examples of versions:

```
0+1
+0-20180112
1.2.3
1.2.3-a1
1.2.3-b2
1.2.3-rc1
1.2.3-alpha1
1.2.3-alpha.1
1.2.3-beta.1
1.2.3+1
+2-1.2.3
+2-1.2.3-alpha.1+3
+2.2.3#1
1.2.3+1#1
+2-1.2.3+1#2
```

The version sorting order is *epoch*, *upstream*, *prerelease*, *revision*, and finally, *iteration*. The *upstream* and *prerelease* parts are compared from left to right, one component at a time, as described next.

To compare two components, first the component types are determined. A component that only consists of digits is an integer. Otherwise, it is a string. If both components are integers, then they are compared as integers. Otherwise, they are compared lexicographically and case-insensitively. The reason for case-insensitive comparison is Windows file names.

A non-existent component is considered 0 if the other component is an integer and an empty string if the other component is a string. For example, in 1.2 vs 1.2.0, the third component in the first version is 0 and the two versions are therefore equal. As a special exception to this rule, an absent *prerelease* part is always greater than any non-absent part. And thus making the final release always older than any pre-release.

This algorithm gives correct results for most commonly-used versioning schemes, for example:

```
1.2.3 < 12.2
1.alpha < 1.beta
20151128 < 20151228
2015.11.28 < 2015.12.28
```

One notable versioning scheme where this approach gives an incorrect result is hex numbers (consider A vs 1A). The simplest work around is to convert such numbers to decimal. Alternatively, one can fix the width of the hex number and pad all the values with leading zeros, for example: 00A vs 01A.

It is also possible to convert the *upstream* and *prerelease* parts into a *canonical representation* that will produce the correct comparison result when always compared lexicographically and as a whole. This can be useful, for example, when storing versions in the database which would otherwise require a custom collation implementation to obtain the correct sort order.

To convert one of these parts to its canonical representation, all its string components are converted to the lower case while all its integer components are padded with leading zeros to the fixed length of 16 characters, with all trailing zero-only components removed. Note that this places an implementation limit on the length of integer components which should be checked by the implementation when converting to the canonical representation. The 16 characters limit was chosen to still be able to represent (with some spare) components in the *YYYYMMDDhhmmss* form while not (visually) bloating the database too much. As a special case, the absent *prerelease* part is represented as ~. Since the ASCII code for ~ is greater than any other character that could appear in *prerelease*, such a string will always be greater than any other representation. The empty *prerelease* part is represented as an empty string.

Note that because it is not possible to perform a reverse conversion without the possibility of loss (consider 01.AA.BB), the original parts may also have to be stored, for example, for display, to derive package archive names, etc.

In quite a few contexts the implementation needs to ignore the *revision* and/or *iteration* parts. For example, this is needed to implement the semantics of newer revisions/iterations of packages replacing their old ones since we do not keep multiple revisions/iterations of the same upstream version in the same repository. As a result, in the package object model, we have a version key as just {*epoch*, *upstream*, *prerelease*} but also store the package revision and iteration so that it can be shown to the user, etc.

3 Manifests

This chapter describes the general manifest file format as well as the concrete manifests used by `bpkg`.

Currently, three manifests are defined: package manifest, repository manifest, and signature manifest. The former two manifests can also be combined into a list of manifests to form the list of available packages and the description of a repository, respectively.

3.1 Manifest Format

The manifest format is a UTF-8 encoded text containing a list of name-value pairs in the form:

```
<name>: <value>
```

For example:

```
name: libfoo
version: 1.2.3
```

The name can contain any characters except `:` and whitespaces. Newline terminates the pair unless escaped with `\` (see below). Leading and trailing whitespaces before and after name and value are ignored except in the multi-line mode (see below).

If, the first non-whitespace character on the line is `#`, then the rest of the line is treated as a comment and ignored except if the preceding newline was escaped or in the multi-line mode (see below). For example:

```
# This is a comment.
short: This is #not a comment
long: Also \
#not a comment
```

The first name-value pair in the manifest file should always have an empty name. The value of this special pair is the manifest format version. The version value shall use the default (that is, non-multi-line) mode and shall not use any escape sequences. Currently it should be `1`, for example:

```
: 1
name: libfoo
version: 1.2.3
```

Any new name that is added without incrementing the version must be optional so that it can be safely ignored by older implementations.

The special empty name pair can also be used to separate multiple manifests. In this case the version may be omitted in the subsequent manifests, for example:

```
: 1
name: libfoo
version: 1.2.3
:
name: libbar
version: 2.3.4
```

To disable treating of a newline as a name-value pair terminator we can escape it with `\`. Note that `\` is only treated as an escape sequence when followed by a newline and both are simply removed from the stream (as opposed to being replaced with a space). To enter a literal `\` at the end of the value, use the `\\` sequence. For example:

```
description: Long text that doesn't fit into one line \
so it is continued on the next line.

windows-path: C:\foo\bar\\
```

Notice that in the final example only the last `\` needs special handling since it is the only one that is followed by a newline.

One may notice that in this newline escaping scheme a line consisting of just `\` followed by a newline has no use, except, perhaps, for visual presentation of, arguably, dubious value. For example, this representation:

```
description: First line. \
\
Second line.
```

Is semantically equivalent to:

```
description: First line. Second line.
```

As a result, such a sequence is "overloaded" to provide more useful functionality in two ways: Firstly, if `:` after the name is immediately followed (ignoring whitespaces) by `\` and a newline, then it signals the start of the multi-line mode. In this mode all subsequent newlines and `#` are treated as ordinary characters rather than value terminators or comments until a line consisting of just `\` and a newline (the multi-line mode terminator). For example:

```
description:\
First paragraph.
#
Second paragraph.
\
```

Expressed as a C-string, the value in the above example is:

```
"First paragraph.\n#\nSecond paragraph."
```

If we didn't expect to ever need to specify a name with an empty value, then an empty value could have turned on the multi-line mode, for example:

```
description:
First paragraph.
#
Second paragraph.
\
```

There are two reasons we don't do this: we don't want to close the door on empty values and we want a more explicit "introductor" for the multi-line mode since it is quite different compared to the simple mode.

Note that in the multi-line mode we can still use newline escaping to split long lines, for example:

```
description:\
First paragraph that doesn't fit into one line \
so it is continued on the next line.
Second paragraph.
\
```

In the simple (that is, non-multi-line) mode, the sole `\` and newline sequence is overloaded to mean a newline. So the previous example can also be represented like this:

```
description: First paragraph that doesn't fit into one \
line so it is continued on the next line.\
\
Second paragraph.
```

Note that the multi-line mode can be used to capture a value with leading and/or trailing whitespaces, for example:

```
description:\
  test
\
```

The C-string representing this value is:

```
"  test\n"
```

EOF can be used instead of a newline to terminate both simple and multi-line values. For example the following representation results in the same value as in the previous example.

```
description:\
  test

<EOF>
```

By convention, names are all in lower case and multi-word names are separated with `-`. Note that names are case-sensitive.

Also by convention, the following name suffixes are used to denote common types of values:

```
-file
-url
-email
```

For example:

```
description: Inline description
description-file: README
package-url: http://www.example.com
package-email: john@example.com
```

Other common name suffixes (such as `-feed`) could be added later.

Generally, unless there is a good reason not to, we keep values lower-case (for example, requires values such as `c++11` or `linux`). An example where we use upper/mixed case would be `license`; it seems unlikely `gplv2` would be better than `GPLv2`.

A number of name-value pairs described below allow for the value proper to be optionally followed by `;` and a comment. Such comments serve as additional documentation for the user and should be full sentence(s), that is start with a capital letter and end with a period. Note that unlike `#`-style comments which are ignored, these comments are considered to be part of the value. For example:

```
email: foo-users@example.com ; Public mailing list.
```

It is recommended that you keep comments short, single-sentence. Note that non-comment semi-colons in such values have to be escaped with a backslash, for example:

```
url: http://git.example.com/?p=foo\;a=tree
```

In the manifest specifications described below optional components are enclosed in square brackets (`[]`). If the name is enclosed in `[]` then the name-value pair is optional, otherwise – required. For example:

```
name: <name>
license: <licenses> [; <comment>]
[description]: <text>
```

In the above example `name` is required, `license` has an optional component (`comment`), and `description` is optional.

In certain situations (for example, shell scripts) it can be easier to parse the binary manifest representation. The binary representation does not include comments and consists of a sequence of name-value pairs in the following form:

```
<name>:<value>\0
```

That is, the name and the value are separated by a colon and each pair (including the last) is terminated with the NUL character. Note that there can be no leading or trailing whitespace characters around the name and any whitespaces after the colon and before the NUL terminator are part of the value. Finally, the manifest format versions are always explicit (that is, not empty) in binary manifest lists.

3.2 Package Manifest

The package manifest (the `manifest` file found in the package's root directory) describes a `bpkg` package. The manifest synopsis is presented next followed by the detailed description of each value in subsequent sections.

The subset of the values up to and including `license` constitute the package manifest header. Note that the header is a valid package manifest since all the other values are optional. There is also no requirement for the header values to appear first or to be in a specific order. In particular, in a full package manifest they can be interleaved with non-header values.

```
name: <name>
version: <version>
[project]: <name>
[priority]: <priority> [; <comment>]
summary: <text>
license: <licenses> [; <comment>]

[tags]: <tags>
[description]: <text>
[description-file]: <path> [; <comment>]
[changes]: <text>
[changes-file]: <path> [; <comment>]

[url]: <url> [; <comment>]
[doc-url]: <url> [; <comment>]
[src-url]: <url> [; <comment>]
[package-url]: <url> [; <comment>]

[email]: <email> [; <comment>]
[package-email]: <email> [; <comment>]
[build-email]: <email> [; <comment>]
[build-warning-email]: <email> [; <comment>]
[build-error-email]: <email> [; <comment>]
```

3.2.1 name

```
[depends]: [?][*] <alternatives> [; <comment>]
[requires]: [?] [<alternatives>] [; <comment>]

[builds]: <class-expr> [; <comment>]
[build-include]: <config>[/<target>] [; <comment>]
[build-exclude]: <config>[/<target>] [; <comment>]
```

3.2.1 name

name: <name>

The package name. See Package Name for the package name format description. Note that the name case is preserved for display, in file names, etc.

3.2.2 version

version: <version>

The package version. See Package Version for the version format description. Note that the version case is preserved for display, in file names, etc.

3.2.3 project

[project]: <name>

The project this package belongs to. The project name has the same restrictions as the package name (see Package Name for details) and its case is preserved for display, in directory names, etc. If unspecified, then the project name is assumed to be the same as the package name.

Projects are used to group related packages together in order to help with organization and discovery in repositories. For example, packages `hello`, `libhello`, and `libhello2` could all belong to project `hello`. By convention, projects of library packages are named without the `lib` prefix.

3.2.4 priority

[priority]: <priority> [; <comment>]

<priority> = `security` | `high` | `medium` | `low`

The release priority (optional). As a guideline, use `security` for security fixes, `high` for critical bug fixes, `medium` for important bug fixes, and `low` for minor fixes and/or feature releases. If not specified, `low` is assumed.

3.2.5 summary

summary: <text>

The short description of the package.

3.2.6 license

license: <licenses> [; <comment>]

<licenses> = <license> [, <license>]*

The package lincense. The format is a comma-separated list of license names under which this package is distributed. This list has the *AND* semantics, that is, the user must comply with all the licenses listed. To capture alternative licensing options use multiple `license` values, for example:

```
license: LGPLv2, MIT
license: BSD
```

In the above example, the package can be used either under the BSD license or both LGPLv2 and MIT.

For complex licensing schemes it is recommended to add comments as an aid to the user, for example:

```
license: LGPLv2, MIT ; If linking with GNU TLS.
license: BSD          ; If linking with OpenSSL.
```

To assist automated processing, the following pre-defined values should be used for the common licenses:

```
MIT
GPLv2
GPLv3
LGPLv2
LGPLv3
ASLv1      ; Apache License 1.0.
ASLv1.1    ; Apache License 1.1.
ASLv2      ; Apache License 2.0.
BSD        ; BSD 3-clause.
BSD-Original ; BSD original.

public domain
available source ; Not free software/open source.
proprietary
```

An example of automated processing would be filtering for non-copyleft licensed packages.

3.2.7 tags

```
[tags]: <tags>
```

```
<tags> = <tag> [, <tag>]*
```

The package tags (optional). The format is a comma-separated list of words that describe this package.

Originally, support for multi-word tags were considered, however this quickly degenerated into quite a complex functionality. In particular, in such a case we should probably treat multi-word tags as if the single word as well as shorter, multi-word derivations were also mentioned. That is, "xml pull parser" should be equivalent to "xml pull parser, xml pull, xml parser, pull parser, xml, pull, parser".

On the user side, if exact phrase matches are favored, then this serves as an encouragement to come up with ever more elaborate multi-word tags to "cover" as much ground as possible. For example, "streaming c++ xml pull parser".

As a result, we will start simple and only allow single-word tags.

3.2.8 description

```
[description]: <text>
```

```
[description-file]: <path> [; <comment>]
```

The detailed description of the package. It can be provided either inline as a text fragment or by referring to a file within a package (e.g., README), but not both.

In the web interface (`brep`) the description is displayed as pre-formatted plain text. We could use the comment to "hint" at the file format.

3.2.9 changes

```
[changes]: <text>
```

```
[changes-file]: <path> [; <comment>]
```

The description of changes in the release.

The tricky aspect is what happens if the upstream release stays the same (and has, say, a NEWS file to which we point) but we need to make another package release, for example, to apply a critical patch.

Multiple changes values can be present which are all concatenated in the order specified, that is, the first value is considered to be the most recent (similar to ChangeLog and NEWS files). For example:


```
changes: 1.2.3-2: applied upstream patch for critical bug bar
changes: 1.2.3-1: applied upstream patch for critical bug foo
changes-file: NEWS
```

Or:

```
changes:\
1.2.3-2
  - applied upstream patch for critical bug bar
  - regenerated documentation

1.2.3-1
  - applied upstream patch for critical bug foo
\
changes-file: NEWS
```

In the web interface (brep) the changes are displayed as pre-formatted plain text, similar to the package description.

3.2.10 url

```
[url]: <url> [; <comment>]
```

The project home page URL.

3.2.11 doc-url

```
[doc-url]: <url> [; <comment>]
```

The project documentation URL.

3.2.12 src-url

```
[src-url]: <url> [; <comment>]
```

The project source repository URL.

3.2.13 package-url

```
[package-url]: <url> [; <comment>]
```

The package home page URL. If not specified, then assumed to be the same as `url`. It only makes sense to specify this value if the project and package are maintained separately.

3.2.14 email

```
[email]: <email> [; <comment>]
```

The project email address. For example, a support mailing list.

3.2.15 package-email

```
[package-email]: <email> [; <comment>]
```

The package email address. If not specified, then assumed to be the same as `email`. It only makes sense to specify this value if the project and package are maintained separately.

3.2.16 build-email

```
[build-email]: <email> [; <comment>]
```

The build notification email address. It is used to send build result notifications by automated build bots. If none of the `build-*email` values are specified, then it is assumed to be the same as `package-email`. If it is specified but empty, then no build result notifications for this package are sent by email.

3.2.17 build-warning-email

```
[build-warning-email]: <email> [; <comment>]
```

The build warning notification email address. Unlike `build-email`, only build warning and error notifications are sent to this email.

3.2.18 build-error-email

```
[build-error-email]: <email> [; <comment>]
```

The build error notification email address. Unlike `build-email`, only build error notifications are sent to this email.

3.2.19 depends

```
[depends]: [?][*] <alternatives> [; <comment>]
```

```
<alternatives> := <dependency> [ ' | ' <dependency> ] *
<dependency>  := <name> [<constraint>]
<constraint>   := <comparison> | <shortcut> | <range>
<comparison>  := ('==' | '>' | '<' | '>=' | '<=') <version>
<shortcut>    := ('^' | '~') <version>
<range>       := ('(' | '[') <version> <version> (')' | ']')
```

The prerequisite packages. If the `depends` value start with `*`, then it is a *build-time* prerequisite. Otherwise it is *run-time*.

Most of the build-time prerequisites are expected to be tools such as code generator, so you can think of `*` as the executable mark printed by `ls`. An important difference between the two kind of dependencies is that in case of cross-compilation a build-time dependency must be built for the build machine, not the target.

Two special build-time dependency names are recognized and checked in an ad hoc manner: `build2` (the `build2` build system) and `bpkg` (the `build2` package manager). This allows us to specify the required build system and package manager versions, for example:

```
depends: * build2 >= 0.6.0
depends: * bpkg >= 0.6.0
```

Each `depends` value can specify multiple packages with the *OR* semantics. While multiple `depends` values are used to specify multiple packages with the *AND* semantics. A value that starts with `?` is a conditional prerequisite. Whether such a prerequisite will be in effect can only be determined at the package configuration time. It is recommended that you provide a comment for each conditional prerequisite as an aid to the user. For example:

```
depends: libz
depends: libfoo ~1.2.0 ; Only works with libfoo 1.2.*.
depends: libgnutls >= 1.2.3 | libopenssl >= 2.3.4
depends: ? libboost-regex >= 1.52.0 ; Only if no C++11 <regex>.
depends: ? libqtcore >= 5.0.0 ; Only if GUI is enabled.
```

It is recommended that you specify unconditional dependencies first with simple (no alternatives) dependencies leading each set.

The optional version constraint can be specified using comparison operators, shortcut (to range) operators, and ranges.

The shortcut operators can only be used with standard versions (a semantic version without the pre-release part is a standard version). They are equivalent to the following ranges. The `X.Y.Z-` version signifies the earliest pre-release in the `X.Y.Z` series; see [Package Version](#) for details.

```
~X.Y.Z  [X.Y.Z  X.Y+1.0-)
^X.Y.Z  [X.Y.Z  X+1.0.0-)  if X >  0
^0.Y.Z  [0.Y.Z  0.Y+1.0-)  if X == 0
```

That is, the tilde (`~`) constraint allows upgrades to any further patch version while the caret (`^`) constraint – also to any further minor version.

Zero major version component is customarily used during early development where the minor version effectively becomes major. As a result, the tilde constraint has special semantics for this case.

Note that the shortcut operators can only be used with the complete, three-component versions (X.Y.Z with the optional pre-release part per the standard version). Specifically, there is no support for special $\wedge X.Y$ or $\sim X$ semantics offered by some package manager – if desired, such functionality can be easily achieved with ranges. Also, the 0.0.Z version is not considered special except as having zero major component for the tilde semantics discussed above.

Note also that pre-releases do not required any special considerations when used with the shortcut operators. For example, if package `libfoo` is usable starting with the second beta of the 2.0.0 release, then our constraint could be expressed as:

```
depends: libfoo ^2.0.0-b.2
```

Internally shortcuts and comparisons can be represented as ranges (that is, $[\nu, \nu]$ for `==`, (ν, inf) for `>`, etc). However, for display and serialization such representations should be converted back to simple operators. While it is possible that the original manifest specified equality or shortcuts as full ranges, it is acceptable to display/serialize them as simpler operators.

Instead of a specific version, the constraint can be specified in terms of the dependent package's version (that is, its `version` value) using the special `$` value. A `depends` value that contains `$` is called incomplete. This mechanism is primarily useful when developing related packages that should track each other's versions exactly or closely. For example:

```
name: sqlite3
version: 3.18.2
depends: libsqlite3 == $
```

In comparison operators and ranges the `$` value is replaced with the dependent version ignoring the revision. For shortcut operators, the dependent version must be a standard version and the following additional processing is applied depending on whether the version is a release, final pre-release, or a snapshot pre-release.

1. For a release we set the min version patch to zero. For \wedge we also set the minor version to zero, unless the major version is zero (reduces to \sim). The max version is set according to the standard shortcut logic. For example, $\sim \$$ is completed as follows:

```
1.2.0 -> [1.2.0 1.3.0-)
1.2.1 -> [1.2.0 1.3.0-)
1.2.2 -> [1.2.0 1.3.0-)
```

And $\wedge \$$ is completed as follows:

```
1.0.0 -> [1.0.0 2.0.0-)
1.1.1 -> [1.0.0 2.0.0-)
```

- For a final pre-release the key observation is that if the patch component for `~` or minor and patch components for `^` are not zero, then that means there has been a compatible release and we treat this case the same as release, ignoring the pre-release part. If, however, it/they are zero, then that means there may yet be no final release and we have to start from the first alpha. For example, for the `~$` case:

```
1.2.0-a.1 -> [1.2.0-a.1 1.3.0-)
1.2.0-b.2 -> [1.2.0-a.1 1.3.0-)
1.2.1-a.1 -> [1.2.0      1.3.0-)
1.2.2-b.2 -> [1.2.0      1.3.0-)
```

And for the `^$` case:

```
1.0.0-a.1 -> [1.0.0-a.1 2.0.0-)
1.0.0-b.2 -> [1.0.0-a.1 2.0.0-)
1.0.1-a.1 -> [1.0.0      2.0.0-)
1.1.0-b.2 -> [1.0.0      2.0.0-)
```

- For a snapshot pre-release we distinguish two cases: a patch snapshot (the patch component is not zero) and a major/minor snapshot (the patch component is zero). For the patch snapshot case we assume that it is (most likely) developed independently of the dependency and we treat it the same as the final pre-release case. For example, if the dependent version is `1.2.1-a.0.nnn`, the dependency could be `1.2.0` or `1.2.2` (or somewhere in-between).

For the major/minor snapshot we assume that all the packages are developed in the lockstep and have the same `X.Y.0` version. In this case we make the range start from the earliest possible version in this "snapshot series" and end before the final pre-release. For example (in this case `~` and `^` are treated the same):

```
1.2.0-a.0.nnn -> [1.2.0-a.0.1 1.2.0-a.1)
2.0.0-b.2.nnn -> [2.0.0-b.2.1 2.0.0-b.3)
```

3.2.20 requires

```
[requires]: [?] [<alternatives>] [; <comment>]
```

```
<alternatives> := <requirement> [ ' | ' <requirement> ]*
<requirement> := <id> | <dependency>
```

The package requirements (other than other packages). Such requirements are normally checked during package configuration by the build system and the only purpose of capturing them in the manifest is for documentation. Similar to `depends`, a value that starts with `?` is a conditional requirement. For example:

3.2.21 builds

```
requires: linux | windows | macosx
requires: c++11
requires: ? ; VC 15 or later if targeting Windows.
requires: ? ; libc++ if using Clang on Mac OS.
```

Notice that in the last two cases the id is omitted altogether with only the comment specifying the requirement.

Note that `requires` should also be used to specify dependencies on external libraries, that is, the ones not packaged or not in the repository. In this case it may make sense to also specify the version constraint. For example:

```
requires: zlib >= 1.2.0 ; Most systems already have it or get from zlib.net.
```

It is recommended that you specify unconditional requirements first with simple (no alternatives) requirements leading each set.

To assist automated processing, the following pre-defined ids should be used for the common requirements:

```
c++98
c++03
c++11
c++14
c++17
c++20
c++23
```

```
posix
linux
macos
freebsd
windows
```

```
gcc[_X.Y.Z] ; For example: gcc_6, gcc_4.9, gcc_5.0.0
clang[_X.Y] ; For example: clang_6, clang_3.4, clang_3.4.1
msvc[_NU] ; For example: msvc_14, msvc_15u3
```

3.2.21 builds

```
[builds]: [<class-uset> ':' ] [<class-expr>] [; <comment>]
```

```
<class-uset> := <class-name> [ <class-name>]*
<class-expr> := <class-term> [ <class-term>]*
<class-term> := ('+' | '-' | '&') ['!'] (<class-name> | '(' <class-expr> ')')
```

The package build configurations. They specify the build configuration classes the package should or should not be built for by automated build bots. For example:

```
builds: -windows
```

Build configurations can belong to multiple classes with their names and semantics varying between different build bot deployments. However, the pre-defined `none`, `default`, and `all` classes are always provided. If no `builds` value is specified in the package manifest, then the `default` class is assumed.

A build configuration class can also derive from another class in which case configurations that belong to the derived class are treated as also belonging to the base class (or classes, recursively). See the Build Configurations page of the build bot deployment for the list of available build configurations and their classes.

The `builds` value consists of an optional underlying class set (`<class-uset>`) followed by a class set expression (`<class-expr>`). The underlying set is a space-separated list of class names that define the set of build configurations to consider. If not specified, then all the configurations belonging to the `default` class are assumed. The class set expression can then be used to exclude certain configurations from this initial set.

The class expression is a space-separated list of terms that are evaluated from left to right. The first character of each term determines whether the build configuration that belong to its set are added to (+), subtracted from (-), or intersected with (&) the current set. If the second character in the term is !, then its set of configuration is inverted against the underlying set. The term itself can be either the class name or a parenthesized expression. Some examples:

```
builds: none                ; None.
builds: all                 ; All.
builds: default legacy     ; Default and legacy.
builds: -windows           ; Default except Windows.
builds: all : -windows     ; All except Windows.
builds: all : &gcc         ; All with GCC only.
builds: all : &gcc-8+      ; All with GCC 8 and up only.
builds: gcc : -optimized   ; GCC without optimization.
builds: gcc : &( +linux +macos ) ; GCC on Linux or Mac OS.
```

Notice that the colon and parentheses must be separated with spaces from both preceding and following terms.

Multiple `builds` values are evaluated in the order specified and as if they were all part of a single expression. Only the first value may specify the underlying set. The main reason for having multiple values is to provide individual reasons (as the `builds` value comments) for different parts of the expression. For example:

```
builds: default experimental ; Only modern compilers are supported.
builds: -gcc                 ; GCC is not supported.
builds: -clang               ; Clang is not supported.
```

The `builds` value comments are used by the web interface (`brep`) to display the reason for the build configuration exclusion.

After evaluating all the `builds` values, the final configuration set can be further fine-tuned using the `build-{include,exclude}` patterns.

3.2.22 `build-{include,exclude}`

```
[build-include]: <config>[/<target>] [; <comment>]
[build-exclude]: <config>[/<target>] [; <comment>]
```

The package build inclusions and exclusions. The `build-include` and `build-exclude` values further reduce the configuration set produced by evaluating the `builds` values. The *config* and *target* values are filesystem wildcard patterns which are matched against the build configuration names and target names (see the `bbot` documentation for details). In particular, the `*` wildcard matches zero or more characters within the name component while the `**` sequence matches across the components. Plus, wildcard-only pattern components match absent name components. For example:

```
build-exclude: windows**      # matches windows_10-msvc_15
build-exclude: macos*-gcc**    # matches macos_10.13-gcc_8.1-03
build-exclude: linux-gcc*-*    # matches linux-gcc_8.1 and linux-gcc_8.1-03
```

The exclusion and inclusion patterns are applied in the order specified with the first match determining whether the package will be built for this configuration and target. If none of the patterns match (or none we specified), then the package is built.

As an example, the following value will exclude 32-bit builds for the MSVC 14 compiler:

```
build-exclude: *-msvc_14**/i?86-* ; Linker crash.
```

As another example, the following pair of values will make sure that a package is only built on Linux:

```
build-include: linux**
build-exclude: ** ; Only supported on Linux.
```

Note that the comment of the matching exclusion is used by the web interface (`brep`) to display the reason for the build configuration exclusion.

3.3 Package List Manifest for pkg Repositories

The package list manifest (the `packages.manifest` file found in the **pkg** repository root directory) describes the list of packages available in the repository. First comes a manifest that describes the list itself (referred to as the list manifest). The list manifest synopsis is presented next:


```
sha256sum: <sum>
```

After the list manifest comes a (potentially empty) sequence of package manifests. These manifests shall not contain any `*-file` or incomplete `depends` values (such values should be converted to their inline versions or completed, respectively) but must contain the following additional (to package manifest) values:

```
location: <path>
sha256sum: <sum>
```

The detailed description of each value follows in the subsequent sections.

3.3.1 sha256sum (list manifest)

```
sha256sum: <sum>
```

The SHA256 checksum of the `repositories.manifest` file (described below) that corresponds to this repository. The *sum* value should be 64 characters long (that is, just the SHA256 value, no file name or any other markers), be calculated in the binary mode, and use lower-case letters.

This checksum is used to make sure that the `repositories.manifest` file that was fetched is the same as the one that was used to create the `packages.manifest` file. This also means that if `repositories.manifest` is modified in any way, then `packages.manifest` must be regenerated as well.

3.3.2 location (package manifest)

```
location: <path>
```

The path to the package archive file relative to the repository root. It should be in the POSIX representation.

if the repository keeps multiple versions of the package and places them all into the repository root directory, it can get untidy. With `location` we allow for sub-directories.

3.3.3 sha256sum (package manifest)

```
sha256sum: <sum>
```

The SHA256 checksum of the package archive file. The *sum* value should be 64 characters long (that is, just the SHA256 value, no file name or any other markers), be calculated in the binary mode, and use lower-case letters.

3.4 Package List Manifest for `dir` Repositories

The package list manifest (the `packages.manifest` file found in the `dir` repository root directory) describes the list of packages available in the repository. It is a (potentially empty) sequence of manifests with the following synopsis:

```
location: <path>
[fragment]: <string>
```

The detailed description of each value follows in the subsequent sections. The `fragment` value can only be present in a merged `packages.manifest` file for a multi-fragment repository.

As an example, if our repository contained the `src/` subdirectory that in turn contained the `libfoo` and `foo` packages, then the corresponding `packages.manifest` file could look like this:

```
: 1
location: src/libfoo/
:
location: src/foo/
```

3.4.1 `location`

```
location: <path>
```

The path to the package directory relative to the repository root. It should be in the POSIX representation.

3.4.2 `fragment`

```
[fragment]: <string>
```

The repository fragment id this package belongs to.

3.5 Repository Manifest

The repository manifest (only used as part of the repository manifest list described below) describes a `pkg`, `dir`, or `git` repository. The manifest synopsis is presented next followed by the detailed description of each value in subsequent sections.

```
[location]: <uri>
[type]: pkg|dir|git
[role]: base|prerequisite|complement
[trust]: <fingerprint>
[url]: <url>
[email]: <email> [; <comment>]
[summary]: <text>
[description]: <text>
[certificate]: <pem>
[fragment]: <string>
```

See also the Repository Chaining documentation for further information @@ TODO.

3.5.1 location

```
[location]: <uri>
```

The repository location. The location can only and must be omitted for the base repository. Since we got hold of its manifest, then we presumably already know the location of the base repository. If the location is a relative path, then it is treated as relative to the base repository location.

For the **git** repository type the relative location does not inherit the URL fragment from the base repository. Note also that the remote **git** repository locations normally have the **.git** extension that is stripped when a repository is cloned locally. To make the relative locations usable in both contexts, the **.git** extension should be ignored if the local prerequisite repository with the extension does not exist while the one without the extension does.

While POSIX systems normally only support POSIX paths (that is, forward slashes only), Windows is generally able to handle both slash types. As a result, it is recommended that POSIX paths are always used in the `location` values, except, perhaps, if the repository is explicitly Windows-only by, for example, having a location that is an absolute Windows path with the drive letter. The **bpkg** package manager will always try to represent the location as a POSIX path and only fallback to the native representation if that is not possible (for example, there is a drive letter in the path).

3.5.2 type

```
[type]: pkg|dir|git
```

The repository type. The type must be omitted for the base repository. If the type is omitted for a prerequisite/complement repository, then it is guessed from its `location` value as described in **bpkg-rep-add(1)**.

3.5.3 role

```
[role]: base|prerequisite|complement
```

The repository role. The `role` value can be omitted for the base repository only.

3.5.4 trust

```
[trust]: <fingerprint>
```

The repository fingerprint to trust. The `trust` value can only be specified for prerequisite and complement repositories and only for repository types that support authentication (currently only `pkg`). The *fingerprint* value should be an SHA256 repository fingerprint represented as 32 colon-separated hex digit pairs. The repository in question is only trusted for use as a prerequisite or complement of this repository. If it is also used by other repositories or is added to the configuration by the user, then such uses cases are authenticated independently.

3.5.5 url

```
[url]: <url>
```

The repository's web interface (`brep`) URL. It can only be specified for the base repository (the web interface URLs for prerequisite/complement repositories can be extracted from their respective manifests).

For example, given the following `url` value:

```
url: https://example.org/hello/
```

The package details page for `libfoo` located in this repository will be `https://example.org/hello/libfoo`.

The web interface URL can also be specified as relative to the repository location (the `location` value). In this case *url* should start with two path components each being either `.` or `...`. If the first component is `..`, then the `www`, `pkg` or `bpkg` domain component, if any, is removed from the `location` URL host, just like when deriving the repository name.

Similarly, if the second component is `..`, then the `pkg` or `bpkg` path component, if any, is removed from the `location` URL path, again, just like when deriving the repository name.

Finally, the version component is removed from the `location` URL path, the rest (after the two `./..` components) of the `url` value is appended to it, and the resulting path is normalized with all remaining `..` and `.` applied normally.

For examples, assuming repository location is:

```
https://pkg.example.org/test/pkg/1/hello/stable
```

The following listing shows some of the possible combinations (the <> marker is used to highlight the changes):

```
./..      -> https://pkg.example.org/test/pkg/hello/stable
../..     -> https://< >example.org/test/pkg/hello/stable
./...     -> https://pkg.example.org/test/< >hello/stable
../...    -> https://< >example.org/test/< >hello/stable
././...   -> https://pkg.example.org/test/pkg/hello< >
.././...  -> https://< >example.org/test< >
```

The rationale for the relative web interface URLs is to allow deployment of the same repository to slightly different configuration, for example, during development, testing, and public use. For instance, for development we may use the `https://example.org/pkg/` setup while in production it becomes `https://pkg.example.org/`. By specifying the web interface location as, say, `../..`, we can run the web interface at these respective locations using a single repository manifest.

3.5.6 email

```
[email]: <email> [; <comment>]
```

The repository email address. It must and can only be specified for the base repository. The email address is displayed by the web interface (`brep`) in the repository about page and could be used to contact the maintainers about issues with the repository.

3.5.7 summary

```
[summary]: <text>
```

The short description of the repository. It must and can only be specified for the base repository.

3.5.8 description

```
[description]: <text>
```

The detailed description of the repository. It can only be specified for the base repository.

In the web interface (`brep`) the description is formatted into one or more paragraphs using blank lines as paragraph separators. Specifically, it is not represented as `<pre>` so any kind of additional plain text formatting (for example, lists) will be lost and should not be used in the description.

3.5.9 certificate

[certificate]: <pem>

The X.509 certificate for the repository. It should be in the PEM format and can only be specified for the base repository. Currently only used for the **pkg** repository type.

The certificate should contain the CN and O components in the subject as well as the email: component in the subject alternative names. The CN component should start with name: and continue with the repository name prefix/wildcard (without trailing slash) that will be used to verify the repository name(s) that are authenticated with this certificate. See **bpkg-repository-signing(1)** for details.

If this value is present then the `packages.manifest` file must be signed with the corresponding private key and the signature saved in the `signature.manifest` file. See Signature Manifest for details.

3.5.10 fragment

[fragment]: <string>

The repository fragment id this repository belongs to.

3.6 Repository List Manifest

@@ TODO See the Repository Chaining document for more information on the terminology and semantics.

The repository list manifest (the `repositories.manifest` file found in the repository root directory) describes the repository. It is a sequence of repository manifests consisting of the base repository manifest (that is, the manifest for the repository that is being described) as well as manifests for its prerequisite and complement repositories. The individual repository manifests can appear in any order and the base repository manifest can be omitted.

The `fragment` values can only be present in a merged `repositories.manifest` file for a multi-fragment repository.

As an example, a repository manifest list for the `math/testing` repository could look like this:

```
# math/testing
#
: 1
email: math-pkg@example.org
summary: Math package repository
:
role: complement
location: ../stable
:
role: prerequisite
location: https://pkg.example.org/1/misc/testing
```

Here the first manifest describes the base repository itself, the second manifest – a complement repository, and the third manifest – a prerequisite repository. Note that the complement repository’s location is specified as a relative path. For example, if the base repository location were:

```
https://pkg.example.org/1/math/testing
```

Then the complement’s location would be:

```
https://pkg.example.org/1/math/stable
```

3.7 Signature Manifest for pkg Repositories

The signature manifest (the `signature.manifest` file found in the **pkg** repository root directory) contains the signature of the repository’s `packages.manifest` file. In order to detect the situation where the downloaded `signature.manifest` and `packages.manifest` files belong to different updates, the manifest contains both the checksum and the signature (which is the encrypted checksum). We cannot rely on just the signature since a mismatch could mean either a split update or tampering. The manifest synopsis is presented next followed by the detailed description of each value in subsequent sections.

```
sha256sum: <sum>
signature: <sig>
```

3.7.1 sha256sum

```
sha256sum: <sum>
```

The SHA256 checksum of the `packages.manifest` file. The *sum* value should be 64 characters long (that is, just the SHA256 value, no file name or any other markers), be calculated in the binary mode, and use lower-case letters.

3.7.2 signature

`signature: <sig>`

The signature of the `packages.manifest` file. It should be calculated by encrypting the above `sha256sum` value with the repository certificate's private key and then `base64`-encoding the result.