`raco`: Racket Command-Line Tools

Version 9.4.0.3

October 29, 2026

The `raco` program supports various Racket tasks from a command line. The first argument to `raco` is always a specific command name. For example, `raco make` starts a command to compile a Racket source module to bytecode format.

The set of commands available through `raco` is extensible. Use `raco help` to get a complete list of available commands for your installation. This manual covers the commands that are available in a typical Racket installation.

Contents

1 `raco make`: Compiling Source to Bytecode **7**

1.1 Running `raco make` . . . 7

1.2 Bytecode Files . . . 7

1.3 Dependency Files . . . 9

1.4 API for Making Bytecode . . . 9

1.5 API for Parallel Builds . . . 18

1.6 Compilation Manager Hook for Syntax Transformers . . . 20

1.7 API for Simple Bytecode Creation . . . 20

1.8 API for Bytecode Paths . . . 21

1.9 Compiling to Raw Bytecode . . . 22

1.10 API for Raw Compilation . . . 23

1.10.1 Bytecode Compilation . . . 23

1.10.2 Recognizing Module Suffixes . . . 26

1.10.3 Loading Compiler Support . . . 28

1.10.4 Options for the Compiler . . . 28

1.10.5 The Compiler as a Unit . . . 29

1.11 API for Reading Compilation Dependencies . . . 29

2 `raco exe`: Creating Stand-Alone Executables **31**

2.1 API for Creating Executables . . . 35

2.1.1 Executable Creation Signature . . . 43

2.1.2 Executable Creation Unit . . . 43

2.1.3 Finding the Racket Executable . . . 43

2.2 Installation-Specific Launchers . . . 44

2.2.1 Creating Launchers 44
2.2.2 Launcher Path and Platform Conventions 48
2.2.3 Launcher Configuration 52
2.2.4 Launcher Creation Signature 54
2.2.5 Launcher Creation Unit 54
2.3 Mac OS Dynamic Library Paths 54

3 `raco distribute`: Sharing Stand-Alone Executables 56
3.1 API for Distributing Executables 57
3.2 API for Bundling Distributions 58

4 `raco planet`: Automatic Package Distribution 59

5 `raco pkg`: Package Management 60

6 `raco setup`: Installation Management 61
6.1 Running `raco setup` 61
6.2 Installing `".plt"` Archives 66
6.3 Controlling `raco setup` with `"info.rkt"` Files 66
6.3.1 Document Descriptions in `"info.rkt"` Files 71
6.3.2 Document Setup Adjustments 76
6.3.3 Documentation Language Families 77
6.3.4 Rendering Documentation Listings 78
6.4 `"info.rkt"` File Format 81
6.5 Package Dependency Checking 82
6.5.1 Declaring Build-Time Dependencies 83
6.5.2 How Dependency Checking Works 83

6.6 API for Setup 84
6.6.1 `raco setup` Unit 86
6.6.2 Options Unit 87
6.6.3 Options Signature 87
6.6.4 Setup Start Module 92
6.7 API for Installing `".plt"` Archives 92
6.7.1 Non-GUI Installer 92
6.8 API for Finding Installation Directories 94
6.9 API for Reading `"info.rkt"` Files 104
6.10 API for Relative Paths 106
6.10.1 Representing Collection-Based Paths 107
6.10.2 Representing Paths Relative to `"collects"` 107
6.10.3 Representing Paths Relative to the Documentation 108
6.10.4 Displaying Paths Relative to a Common Root 108
6.11 API for Collection Names 110
6.12 API for Collection Searches 110
6.13 API for Platform Specifications 111
6.14 API for Cross-Platform Configuration 112
6.15 API for Cross-References for Installed Manuals 114
6.16 API for Materializing User-Specific Documentation 115
6.17 API for Staging Documentation Installs 116
6.18 Layered Installations 117
6.19 Tethered Installations 118

7 `raco decompile`: Decompiling Bytecode **120**

7.1 Racket CS Decompilation 120

7.2 Racket BC Decompilation 121
7.3 API for Decompiling 122
7.4 API for Parsing Bytecode 123
7.5 API for Marshaling Bytecode 124
7.6 Bytecode Representation 124
7.6.1 Prefix 125
7.6.2 Forms and Inline Variants 128
7.6.3 Expressions 129
7.7 Machine-Independent Linklets 136

8 `raco demod`: Demodularizing Programs **138**
8.1 Demodularizing Libraries 139
8.2 Language for Demodularizing 140

9 `raco link`: Library Collection Links **143**
9.1 API for Collection Links 144

10 `raco pack`: Packing Library Collections **146**
10.1 Format of `".plt"` Archives 148
10.2 API for Packing 150

11 `raco unpack`: Unpacking Library Collections **155**
11.1 Unpacking API 155

12 `raco ctool`: Working with C Code **158**
12.1 Compiling and Linking C Extensions 158
12.1.1 API for 3m Transformation 158
12.2 Embedding Modules via C 159

13 `raco test`: Run tests **161**

13.1 Test Configuration by Submodule 164

13.2 Test Configuration by `"info.rkt"` 164

13.3 Responsible-Party and Varying-Output Logging 166

13.4 Logging Test Results 166

14 `raco docs`: Documentation Search **168**

15 `raco expand`: Macro Expansion **169**

16 `raco read`: Reading and Pretty-Printing **170**

17 `raco scribble`: Building Documentation **171**

18 Adding a `raco` Command **172**

18.1 Command Argument Parsing 173

18.2 Accessing `raco` Commands 174

19 Installation Configuration and Search Paths **175**

1 `raco make`: Compiling Source to Bytecode

The `raco make` command accept filenames for Racket modules to be compiled to bytecode format. Modules are re-compiled only if the source Racket file is newer than the bytecode file and has a different SHA-1 hash, or if any imported module is recompiled or has a different SHA-1 hash for its compiled form plus dependencies.

1.1 Running `raco make`

The `raco make` command accepts a few flags:

- `-l` ⟨*path*⟩ — Compiles ⟨*path*⟩ interpreted as a collection-based module path, as for `require`.
- `-j` ⟨*n*⟩ — Compiles argument modules in parallel, using up to ⟨*n*⟩ parallel tasks.
- `--disable-inline` — Disables function inlining while compiling (but does not re-compile files that are already up-to-date). This flag is often useful to simplify generated code before decompiling, and it corresponds to setting `compile-context-preservation-enabled` to `#t`.
- `--disable-constant` — Disables inference of definitions within a module as constant (but does not re-compile files that are already up-to-date). The value associated with a non-constant definition is never inlined or constant-propagated, either within its own module or an importing module. This flag corresponds to setting `compile-enforce-module-constants` to `#f`.
- `--no-deps` — Compiles a non-module file (i.e., one that is run via `load` instead of `require`). See §1.9 “Compiling to Raw Bytecode” for more information.
- `-p` ⟨*file*⟩ or `--prefix` ⟨*file*⟩ — For use with `--no-deps`; see §1.9 “Compiling to Raw Bytecode”.
- `-no-prim` — For use with `--no-deps`; see §1.9 “Compiling to Raw Bytecode”.
- `-v` — Verbose mode, which shows which files are compiled.
- `--vv` — Very verbose mode, which implies `-v` and also shows every dependency that is checked.

1.2 Bytecode Files

A file "⟨*name*⟩.⟨*ext*⟩" is compiled to bytecode that is saved as "compiled/⟨*name*⟩_⟨*ext*⟩.zo" relative to the file. As a result, the bytecode file is normally

used automatically when "⟨*name*⟩.⟨*ext*⟩" is required as a module, since the underlying load/use-compiled operation detects such a bytecode file.

For example, in a directory that contains the following files:

- "a.rkt":

```
#lang racket
(require "b.rkt" "c.rkt")
(+ b c)
```

- "b.rkt":

```
#lang racket
(provide b)
(define b 1)
```

- "c.rkt":

```
#lang racket
(provide c)
(define c 1)
```

then

```
raco make a.rkt
```

triggers the creation of "compiled/a_rkt.zo", "compiled/b_rkt.zo", and "compiled/c_rkt.zo". A subsequent

```
racket a.rkt
```

loads bytecode from the generated ".zo" files, paying attention to the ".rkt" sources only to confirm that each ".zo" file has a later timestamp (unless the PLT_COMPILED_FILE_CHECK environment variable is set to exists, in which case the compiled file is used without a timestamp check).

In contrast,

```
raco make b.rkt c.rkt
```

would create only "compiled/b_rkt.zo" and "compiled/c_rkt.zo", since neither "b.rkt" nor "c.rkt" imports "a.rkt".

1.3 Dependency Files

In addition to a bytecode file, `raco make` creates a file `"compiled/`⟨*name*⟩`_`⟨*ext*⟩`.dep"` that records dependencies of the compiled module on other module files and the source file's SHA-1 hash. Using this dependency information, a re-compilation request via `raco make` can consult both the source file's timestamp/hash and the timestamps/hashes for the bytecode of imported modules. Furthermore, imported modules are themselves compiled as necessary, including updating the bytecode and dependency files for the imported modules, transitively.

Continuing the `raco make a.rkt` example from the previous section, the `raco make` command creates `"compiled/a_rkt.dep"`, `"compiled/b_rkt.dep"`, and `"compiled/c_rkt.dep"` at the same time as the `".zo"` files. The `"compiled/a_rkt.dep"` file records the dependency of `"a.rkt"` on `"b.rkt"`, `"c.rkt"` and the `racket` library. If the `"b.rkt"` file is modified (so that its SHA-1 hash changes), then running

```
raco make a.rkt
```

again rebuilds `"compiled/a_rkt.zo"` and `"compiled/b_rkt.zo"`.

For module files that are within library collections, `raco setup` uses the same `".zo"` and `".dep"` conventions and files as `raco make`, so the two tools can be used together.

As long as the `PLT_COMPILED_FILE_CHECK` environment variable is not set or is set to `modify`, then `raco make` updates the timestamp on a compiled bytecode file if it is older than the source, even if the file does not need to be recompiled.

1.4 API for Making Bytecode

`(require compiler/cm)` package: `base`

The `compiler/cm` module implements the compilation and dependency management used by `raco make` and `raco setup`.

```
(make-compilation-manager-load/use-compiled-handler
 [delete-zos-when-rkt-file-does-not-exist?
  #:security-guard security-guard])
→ (path? (or/c symbol? #f) . -> . any)
 delete-zos-when-rkt-file-does-not-exist? : any/c = #f
 security-guard : (or/c security-guard? #f) = #f
```

Returns a procedure suitable as a value for the `current-load/use-compiled` parameter. The returned procedure passes its arguments on to the `current-load/use-compiled`

procedure that is installed when `make-compilation-manager-load/use-compiled-handler` is called, but first it automatically compiles a source file to a `".zo"` file if

- the file is expected to contain a module (i.e., the second argument to the handler is a symbol);
- the value of each of `(current-eval)`, `(current-load)`, and `(namespace-module-registry (current-namespace))` is the same as when `make-compilation-manager-load/use-compiled-handler` was called;
- the value of `use-compiled-file-paths` contains the first path that was present when `make-compilation-manager-load/use-compiled-handler` was called;
- the value of `current-load/use-compiled` is the result of this procedure; and
- one of the following holds:
  - the source file is newer than the `".zo"` file in the first sub-directory listed in `use-compiled-file-paths` (at the time that `make-compilation-manager-load/use-compiled-handler` was called), and either no `".dep"` file exists or it records a source-file SHA-1 hash that differs from the current version and source-file SHA-1 hash;
  - no `".dep"` file exists next to the `".zo"` file;
  - the version recorded in the `".dep"` file does not match the result of `(version)`;
  - the target machine recorded in the `".dep"` file does not match the result of `(current-compile-target-machine)`;
  - the source hash recorded in the `".dep"` file does not match the current source hash;
  - one of the files listed in the `".dep"` file has a `".zo"` timestamp newer than the target `".zo"` and `use-compiled-file-check` is set to `'modify-seconds`;
  - the combined hashes of the dependencies recorded in the `".dep"` file does not match the combined hash recorded in the `".dep"` file.

If SHA-1 hashes override a timestamp-based decision to recompile the file, then the target `".zo"` file's timestamp is updated to the current time, unless the `use-compiled-file-check` parameter is not set to `'modify-seconds`.

After the handler procedure compiles a `".zo"` file, it creates a corresponding `".dep"` file that lists the current version and the identification of every file that is directly `required` by the module in the compiled file. Additional dependencies can be installed during compilation via `compiler/cm-accomplice`. The `".dep"` file also records the SHA-1 hash of the module's source, and it records a combined SHA-1 hash of all of the dependencies that includes their recursive dependencies. If a bytecode file is generated by recompiling a bytecode file that was formerly compiled as machine-independent, then the `".dep"` file also records the SHA-1 hash of the machine-independent form, since the recompiled module's behavior should be exactly the same.

The special combination of (cross-installation?) or (current-multi-compile-any) as #t, (current-compile-target-machine) as #f, and (current-compiled-file-roots) having two or more elements triggers a special compilation mode. Bytecode specific to the running Racket is written to the directory determined by the first element of (current-compiled-file-roots). Bytecode specific to either the cross-compilation target for (cross-installation?) or machine-independent format if (current-multi-compile-any) is written to the directory determined by the second element of (current-compiled-file-roots). By configuring (current-compiled-file-roots) so that the first element is outside a build tree and the second element is inside the build tree, cross-compilation can create a build tree suitable for the target machine while building and loading bytecode (for macro expansion, etc.) that is usable on the current machine. This mode works correctly for a build directory that starts with only source code and machine-independent bytecode.

The handler caches timestamps when it checks ".dep" files, and the cache is maintained across calls to the same handler. The cache is not consulted to compare the immediate source file to its ".zo" file, which means that the caching behavior is consistent with the caching of the default module name resolver (see current-module-name-resolver).

If use-compiled-file-paths contains an empty list when make-compilation-manager-load/use-compiled-handler is called, then an exn:fail:contract exception is raised.

If the *delete-zos-when-rkt-file-does-not-exist?* argument is a true value, then the returned handler will delete ".zo" files when there is no corresponding original source file.

If the *security-guard* argument is supplied, it is used when creating ".zo" files, ".dep" files, and "compiled/" directories, and when it adjusts the timestamps for existing files. If it is #f, then the security guard in the current-security-guard when the files are created is used (not the security guard at the point make-compilation-manager-load/use-compiled-handler is called).

The continuation of the compilation of a module is marked with a managed-compiled-context-key and the module's source path.

Do not install the result of make-compilation-manager-load/use-compiled-handler when the current namespace contains already-loaded versions of modules that may need to be recompiled—unless the already-loaded modules are never referenced by not-yet-loaded modules. References to already-loaded modules may produce compiled files with inconsistent timestamps and/or ".dep" files with incorrect information.

The handler logs messages to the topic 'compiler/cm at the level 'info. These messages are instances of a compile-event prefab structure:

```
(struct compile-event (timestamp path type) #:prefab)
```

The timestamp field is the time at which the event occurred in milliseconds since the epoch.

The path field is the path of a file being compiled for which the event is about. The type field is a symbol which describes the action the event corresponds to. The currently logged values are 'locking, 'start-compile, 'finish-compile, and 'already-done.

Changed in version 6.1.1.8 of package base: Added identification of the compilation context via managed-compiled-context-key.
Changed in version 6.6.0.3: added check on a source's SHA1 hash to complement the timestamp check, where the latter can be disabled via use-compile-file-check.

```
(managed-compile-zo  file
                     [read-src-syntax
                      #:security-guard security-guard]) → void?
  file : path-string?
  read-src-syntax : (any/c input-port? . -> . syntax?)
                  = read-syntax
  security-guard : (or/c security-guard? #f) = #f
```

Compiles the given module source file to a ".zo", installing a compilation-manager handler while the file is compiled (so that required modules are also compiled), and creating a ".dep" file to record the timestamps of immediate files used to compile the source (i.e., files required in the source).

Compilation is triggered by loading a module into the current namespace, so if a module that is a dependency of file has already been loaded into the current namespace, then that module will not necessarily be (re-)compiled. The handler used to trigger compilation is created with make-compilation-manager-load/use-compiled-handler, so all the rules and constraints there apply.

If file is compiled from source, then read-src-syntax is used in the same way as read-syntax to read the source module. The normal read-syntax is used for any required files, however.

If security-guard is not #f, then the provided security guard is used when creating the "compiled/" directories, ".dep" and ".zo" files, and when it adjusts the timestamps of existing files. If it is #f, then the security guard in the current-security-guard when the files are created is used (not the security guard at the point managed-compile-zo is called).

While compiling file, the error-display-handler parameter is set to (make-compilation-context-error-display-handler (error-display-handler)), so that errors from uncaught exceptions will report the compilation context.

Changed in version 6.1.1.8 of package base: Added error-display-handler configuration.

```
managed-compiled-context-key : any/c
```

A key used as a continuation mark key by make-compilation-manager-load/use-

compiled-handler for the continuation of a module compilation. The associated value is a path to the module's source.

Added in version 6.1.1.8 of package base.

```
(make-compilation-context-error-display-handler orig-handlers)
 → (string? any/c . -> . void?)
  orig-handlers : (string? any/c . -> . void?)
```

Produces a handler suitable for use as an error-display-handler value, given an existing such value. The generated handler shows information about the compilation context when the handler's second argument is an exception whose continuation marks include managed-compiled-context-key keys.

Added in version 6.1.1.8 of package base.

```
(trust-existing-zos) → boolean?
(trust-existing-zos trust?) → void?
  trust? : any/c
```

A parameter that is intended for use by raco setup when installing with pre-built ".zo" files. It causes a compilation-manager load/use-compiled handler to "touch" out-of-date ".zo" files instead of re-compiling from source.

```
(make-caching-managed-compile-zo
 [read-src-syntax
  #:security-guard security-guard])
 → (path-string? . -> . void?)
  read-src-syntax : (any/c input-port? . -> . syntax?)
                  = read-syntax
  security-guard : (or/c security-guard? #f) = #f
```

Returns a procedure that behaves like managed-compile-zo (providing the same *read-src-syntax* each time), but a cache of timestamp information is preserved across calls to the procedure.

A handler to support compilation is created with make-compilation-manager-load/use-compiled-handler each time the result of make-caching-managed-compile-zo is called, so the current namespace and other parameter values are relevant at that time, not when make-caching-managed-compile-zo is called.

```
(manager-compile-notify-handler) → (path? . -> . any)
(manager-compile-notify-handler notify) → void?
  notify : (path? . -> . any)
```

A parameter for a procedure of one argument that is called whenever a compilation starts. The argument to the procedure is the file's path.

```
(manager-trace-handler) → (string? . -> . any)
(manager-trace-handler notify) → void?
  notify : (string? . -> . any)
```

A parameter for a procedure of one argument that is called to report compilation-manager actions, such as checking a file. The argument to the procedure is a string.

The default value of the parameter logs the argument, along with `current-inexact-milliseconds`, to a logger named `'compiler/cm` at the `'debug` level.

```
(manager-skip-file-handler)
 → (-> path? (or/c (cons/c number? promise?) #f))
(manager-skip-file-handler proc) → void?
  proc : (-> path? (or/c (cons/c number? promise?) #f))
```

A parameter whose value is called for each file that is loaded and needs recompilation. If the procedure returns a pair, then the file is skipped (i.e., not compiled); the number in the pair is used as the timestamp for the file's bytecode, and the promise may be `force`d to obtain a string that is used as hash of the compiled file plus its dependencies. If the procedure returns `#f`, then the file is compiled as usual. The default is `(lambda (x) #f)`.

```
(current-path->mode)
 → (or/c #f (-> path? (and/c path? relative-path?)))
(current-path->mode path->mode) → void?
  path->mode : (or/c #f (-> path? (and/c path? relative-path?)))
 = #f
```

Used by `make-compilation-manager-load/use-compiled-handler` and `make-caching-managed-compile-zo` to override `use-compiled-file-paths` for deciding where to write compiled `".zo"` files. If it is `#f`, then the first element of `use-compiled-file-paths` is used. If it isn't `#f`, then it is called with the original source file's location and its result is treated the same as if it had been the first element of `use-compiled-file-paths`.

Note that this parameter is not used by `current-load/use-compiled`. So if the parameter causes `".zo"` files to be placed in different directories, then the correct `".zo"` file must still be communicated via `use-compiled-file-paths`, and one way to do that is to override `current-load/use-compiled` to delete `".zo"` files that would cause the wrong one to be chosen right before they are loaded.

Added in version 6.4.0.14 of package `base`.

```
(file-stamp-in-collection p)
 → (or/c (cons/c number? promise?) #f)
  p : path?
```

Calls `file-stamp-in-paths` with *p* and `(current-library-collection-paths)`.

```
(file-stamp-in-paths p paths)
 → (or/c (cons/c number? promise?) #f)
  p : path?
  paths : (listof path?)
```

Returns the file-modification date and `delayed` hash of *p* or its bytecode form (i.e., `".zo"` file), whichever exists and is newer, if *p* is an extension of any path in *paths* (i.e., exists in the directory, a subdirectory, etc.). Otherwise, the result is `#f`.

This function is intended for use with `manager-skip-file-handler`.

```
(get-file-sha1 p) → (or/c string? #f)
  p : path?
```

Computes a SHA-1 hash for the file *p*; the result is `#f` if *p* cannot be opened.

```
(get-compiled-file-sha1 p) → (or/c string? #f)
  p : path?
```

Computes a SHA-1 hash for the bytecode file *p*, appending any dependency-describing hash available from a `".dep"` file when available (i.e., the suffix on *p* is replaced by `".dep"` to locate dependency information). The result is `#f` if *p* cannot be opened.

```
(with-compile-output p proc) → any
  p : path-string?
  proc : ([port input-port?] [tmp-path path?]  . -> . any)
```

A wrapper on `call-with-atomic-output-file` that passes along any security guard put in place by `make-compilation-manager-load/use-compiled-handler`, etc.

```
(parallel-lock-client)
 → (or/c #f
         (->i ([command (or/c 'lock 'unlock)]
                [file bytes?])
              [res (command) (if (eq? command 'lock)
                                 boolean?
                                 void?)]))
(parallel-lock-client proc) → void?
  proc : (or/c #f
               (->i ([command (or/c 'lock 'unlock)]
                      [file bytes?])
                    [res (command) (if (eq? command 'lock)
                                       boolean?
                                       void?)]))
```

Creates a parallel compilation lock client, which is used by the result of make-compilation-manager-load/use-compiled-handler to prevent compilation races between parallel builders.

When *proc* is #f (the default), no checking for parallel compilation is done (and thus multiple threads or places running compilations via make-compilation-manager-load/use-compiled-handler will potentially corrupt each other's ".zo" files).

When *proc* is a function, its first argument is a command 'lock pr 'unlock, which indicates whether the caller wants to lock or unlock a target *zo-path*, and the second argument is the target *zo-path* (expressed as a byte string).

When *proc* returns #t for a 'lock command, the current builder has obtained the lock for *zo-path*. Once compilation of *zo-path* is complete, the builder process must release the lock by calling *proc* 'unlock with the exact same *zo-path*.

When *proc* returns #f for a 'lock command, another parallel builder obtained the lock first and has already compiled the target. The parallel builder should continue without compiling *zo-path*. (In this case, make-compilation-manager-load/use-compiled-handler's result will not call *proc* with 'unlock.)

Example:

```
> (let* ([lc (parallel-lock-client)]
         [zo-name  #"collects/racket/compiled/draw_rkt.zo"]
         [locked? (and lc (lc 'lock zo-name))]
         [ok-to-compile? (or (not lc) locked?)])
    (dynamic-wind
      (lambda () (void))
      (lambda ()
        (when ok-to-compile?
          (printf "Do compile here ...\n")))
      (lambda ()
        (when locked?
          (lc 'unlock zo-name)))))
Do compile here ...
```

```
(compile-lock->parallel-lock-client  pc
                                     [cust
                                      current-shutdown-evt])
 → (-> (or/c 'lock 'unlock) bytes? boolean?)
  pc : place-channel?
  cust : (or/c #f custodian?) = #f
  current-shutdown-evt : (-> evt?) = (lambda () never-evt)
```

Returns a function that follows the parallel-lock-client protocol by communicating over *pc*, where *pc* must be a result of make-compile-lock.

This communication protocol implementation is not kill-safe when *cust* is #f. Making the protocol kill-safe requires a sufficiently powerful custodian (i.e., one that is not subject to termination unless all of the participants in the compilation are also terminated) supplied as *cust*. The given custodian is used to create a thread that monitors the threads that are perform the compilation. If one of the threads is terminated, the presence of the custodian lets another one continue. (The custodian is also used to create a thread that manages a thread-safe table.)

Just checking for thread termination is not always sufficient to release a lock, because a thread created with thread/suspend-to-kill is merely suspending by removing its ability to run. The *current-shutdown-evt* argument returns an synchronizable event that the monitor thread waits on at the same time as it waits for a thread to terminate. If the event becomes ready, then the monitor releases a lock the same as if the thread was terminated. For example, *current-shutdown-evt* might return a custodian box to detect a custodian shutdown.

Changed in version 8.1.0.7 of package base: Added the *current-shutdown-evt* argument.

```
(make-compile-lock) → place-channel?
```

Creates a place-channel that can be used with compile-lock->parallel-lock-client to avoid concurrent compilations of the same Racket source files in multiple places.

```
(install-module-hashes! bstr [start end]) → void?
  bstr : bytes?
  start : exact-nonnegative-integer? = 0
  end : exact-nonnegative-integer? = (bytes-length bstr)
```

Adjusts the bytecode representation in *bstr* (from bytes *start* to *end*) to install a hash code, including any submodules within the region. The existing representation should have zero bytes in place of each hash string, which is what write produces for a compiled form.

Added in version 6.3 of package base.

```
(current-multi-compile-any) → boolean?
(current-multi-compile-any on?) → void?
  on? : any/c
```

A parameter that enables compilation of both current-machine bytecode and machine-independent bytecode by a handler created with make-compilation-manager-load/use-compiled-handler.

Added in version 8.1.0.2 of package base.

1.5 API for Parallel Builds

```
(require setup/parallel-build)
```
package: base

The setup/parallel-build library provides the parallel-compilation functionality of raco setup and raco make.

Both parallel-compile-files and parallel-compile log messages to the topic 'setup/parallel-build at the level 'info. These messages are instances of a parallel-compile-event prefab structure:

```
(struct parallel-compile-event (worker event) #:prefab)
```

The worker field is the index of the worker that the created the event. The event field is a compile-event as documented in make-compilation-manager-load/use-compiled-handler.

```
(parallel-compile-files list-of-files
                        [#:worker-count worker-count
                         #:use-places? use-places?
                         #:handler handler])
 → (or/c void? #f)
  list-of-files : (listof path-string?)
  worker-count : exact-positive-integer? = (processor-count)
  use-places? : any/c = #t
  handler : (->i ([worker-id exact-integer?] = void
                  [handler-type symbol?]
                  [path path-string?]
                  [msg string?]
                  [out string?]
                  [err string?])
                 void?)
```

The parallel-compile-files utility function is used by raco make to compile a list of paths in parallel. The optional #:worker-count argument specifies the number of compile workers to spawn during parallel compilation. The compile workers are implemented as Racket places if *use-places?* is true, otherwise the compile workers are implemented as separate Racket processes. The callback, *handler*, is called with the symbol 'done as the *handler-type* argument for each successfully compiled file, 'output when a successful compilation produces stdout/stderr output, 'error when a compilation error has occurred, or 'fatal-error when an unrecoverable error occurs. The other arguments give more information for each status update. The return value is (void) if it was successful, or #f if there was an error.

```
(parallel-compile-files
  source-files
```

```
#:worker-count 4
#:handler
(lambda (type work msg out err)
  (match type
    ['done (when (verbose) (printf " Made ~a\n" work))]
    ['output (printf " Output from: ~a\n~a~a" work out err)]
    [_ (printf " Error compiling ~a\n~a\n~a~a"
               work
               msg
               out
               err)]))
```

Changed in version 7.0.0.19 of package `base`: Added the `#:use-places?` argument.

```
(parallel-compile  worker-count
                   setup-fprintf
                   append-error
                   collects-tree
                  [#:use-places? use-places?]) → (void)
  worker-count : non-negative-integer?
  setup-fprintf : (->i ([stage string?] [format string?])
                       ()
                       #:rest (listof any/c) void)
  append-error : (->i ([cc cc?]
                       [prefix string?]
                       [exn (or/c exn? (cons/c string? string?) #f)]
                       [out string?]
                       [err string?]
                       [message string?])
                      void?)
  collects-tree : (listof any/c)
  use-places? : any/c = #t
```

The `parallel-compile` function is used by `raco setup` to compile collections in parallel. The *worker-count* argument specifies the number of compilation workers to spawn during parallel compilation. The *use-places?* argument specified whether to use places, otherwise separate processes are used. The *setup-fprintf* and *append-error* functions communicate intermediate compilation results and errors. The *collects-tree* argument is a compound data structure containing an in-memory tree representation of the collects directory.

When the *exn* argument to *append-error* is a pair of strings, the first string is a long form of the error message, and the second string is a short form (omitting evaluation context information, for example).

Changed in version 6.1.1.8 of package `base`: Changed *append-error* to allow a pair of error strings.

Changed in version 7.0.0.19: Added the #:use-places? argument.

1.6 Compilation Manager Hook for Syntax Transformers

```
(require compiler/cm-accomplice)      package: base
```

```
(register-external-file  file
                         [#:indirect? indirect?]) → void?
  file : (and path? complete-path?)
  indirect? : any/c = #f
```

Logs a message (see log-message) to the current logger at level 'info with the topic 'cm-accomplice. The message data is a file-dependency prefab structure type with two fields; the first field's value is *file* and the second field's value is #f (to indicate a non-module dependency). If the *indirect?* argument is true, the data is more specifically an instance of a file-dependency/options prefab structure type that is a subtype of file-dependency with one extra field: a hash table mapping 'indirect to #t.

A compilation manager implemented by compiler/cm looks for such messages to register an external dependency. In response, the compilation manager records (in a ".dep" file) the path as contributing to the implementation of the module currently being compiled. Afterward, if the registered file is modified, the compilation manager will know to recompile the module. An indirect dependency has no effect on recompilation, but it can signal to other tools, such as a package-dependency checker, that the dependency is indirect (and should not imply a direct package dependency).

The include macro, for example, calls this procedure with the path of an included file as it expands an include form.

```
(register-external-module  file
                           [#:indirect? indirect?]) → void?
  file : (and path? complete-path?)
  indirect? : any/c = #f
```

Like register-external-file, but logs a message with a file-dependency prefab structure type whose second field is #t.

A compilation manager implemented by compiler/cm recognizes the message to register a dependency on a module (which implies a dependency on all of that module's dependencies, etc.).

1.7 API for Simple Bytecode Creation

```
(require compiler/compile-file)      package: base
```

```
(compile-file src [dest filter]) → path?
  src : path-string?
  dest : path-string?
       = (let-values ([(base name dir?) (split-path src)])
           (build-path base "compiled"
                       (path-add-suffix name #".zo")))
  filter : (any/c . -> . any/c) = values
```

Compiles the Racket file *src* and saves the compiled code to *dest*. If *dest* is not provided and the `"compiled"` subdirectory does not already exist, the subdirectory is created. The result of `compile-file` is the destination file's path.

If the *filter* procedure is provided, it is applied to each source expression, and the result is compiled.

Beware that `compile-file` uses the current reader parameterization to read *src*. Typically, `compile-file` should be called from a thunk passed to `with-module-reading-parameterization` so that the source program is parsed in a consistent way and allowing `#lang`.

Each expression in *src* is compiled independently. If *src* does not contain a single `module` expression, then earlier expressions can affect the compilation of later expressions when *src* is loaded directly. An appropriate *filter* can make compilation behave like evaluation, but the problem is also solved (as much as possible) by the `compile-zos` procedure.

See also `managed-compile-zo`.

1.8 API for Bytecode Paths

```
(require compiler/compilation-path)
```
package: `base`

Added in version 6.0.1.10 of package `base`.

```
(get-compilation-dir+name  path
                           [#:modes modes
                            #:roots roots
                            #:default-root default-root])
 → path? path?
  path : path-string?
  modes : (non-empty-listof (and/c path-string? relative-path?))
        = (use-compiled-file-paths)
  roots : (non-empty-listof (or/c path-string? 'same))
        = (current-compiled-file-roots)
  default-root : (or/c path-string? 'same) = (car roots)
```

Determines the directory that holds the bytecode form of *path* plus the base name of *path*.

The directory is determined by checking *roots* in order, and for each element of *roots* checking *modes* in order. The first such directory that contains a file whose name matches *path* with `".zo"` added (in the sense of path-add-suffix) is reported as the return directory path. If no such file is found, the result corresponds to the first element of *modes* combined with *default-root*.

Changed in version 7.1.0.9 of package base: Added the `#:default-root` argument.

```
(get-compilation-dir  path
                      [#:modes modes
                       #:roots roots
                       #:default-root default-root]) → path?
  path : path-string?
  modes : (non-empty-listof (and/c path-string? relative-path?))
        = (use-compiled-file-paths)
  roots : (non-empty-listof (or/c path-string? 'same))
        = (current-compiled-file-roots)
  default-root : (or/c path-string? 'same) = (car roots)
```

The same as get-compilation-dir+name, but returning only the first result.

Changed in version 7.1.0.9 of package base: Added the `#:default-root` argument.

```
(get-compilation-bytecode-file  path
                                [#:modes modes
                                 #:roots roots
                                 #:default-root default-root])
 → path?
  path : path-string?
  modes : (non-empty-listof (and/c path-string? relative-path?))
        = (use-compiled-file-paths)
  roots : (non-empty-listof (or/c path-string? 'same))
        = (current-compiled-file-roots)
  default-root : (or/c path-string? 'same) = (car roots)
```

The same as get-compilation-dir+name, but combines the results and adds a `".zo"` suffix to arrive at a bytecode file path.

Changed in version 7.1.0.9 of package base: Added the `#:default-root` argument.

1.9 Compiling to Raw Bytecode

The `--no-deps` mode for `raco make` is an improverished form of the compilation, because it does not track import dependencies. It does, however, support compilation of non-module

source in a namespace that initially imports `scheme`.

Outside of a module, top-level `define-syntaxes`, `module`, `#%require`, `define-values-for-syntax`, and `begin` expressions are handled specially by `raco make --no-deps`: the compile-time portion of the expression is evaluated, because it might affect later expressions.

For example, when compiling the file containing

```
(require racket/class)
(define f (class object% (super-new)))
```

the `class` form from the `racket/class` library must be bound in the compilation namespace at compile time. Thus, the `require` expression is both compiled (to appear in the output code) and evaluated (for further computation).

Many definition forms expand to `define-syntaxes`. For example, `define-signature` expands to `define-syntaxes`. In `--no-deps` mode, `raco make --no-deps` detects `define-syntaxes` and other expressions after expansion, so top-level `define-signature` expressions affect the compilation of later expressions, as a programmer would expect.

In contrast, a `load` or `eval` expression in a source file is compiled—but *not evaluated!*—as the source file is compiled. Even if the `load` expression loads syntax or signature definitions, these will not be loaded as the file is compiled. The same is true of application expressions that affect the reader, such as `(read-case-sensitive #t)`. The `-p` or `--prefix` flag for `raco make` takes a file and loads it before compiling the source files specified on the command line.

By default, the namespace for compilation is initialized by a `require` of `scheme`. If the `--no-prim` flag is specified, the namespace is instead initialized with `namespace-require/copy`, which allows mutation and redefinition of all initial bindings (other than syntactic forms, in the case of mutation).

In general, a better solution is to put all code to compile into a module and use `raco make` in its default mode.

1.10 API for Raw Compilation

`(require compiler/compiler)` package: `base`

The `compiler/compiler` library provides the functionality of `raco make` for compilation to bytecode, but through a Racket API.

1.10.1 Bytecode Compilation

```
((compile-zos  expr
               [#:module? module?
                #:verbose? verbose?])
  racket-files
  dest-dir)                              → void?
  expr : any/c
  module? : any/c = #f
  verbose? : any/c = #f
  racket-files : (listof path-string?)
  dest-dir : (or/c path-string? #f 'auto)
```

Supplying just *expr* returns a compiler that is initialized with the expression *expr*, as described below.

The compiler takes a list of Racket files and compiles each of them to bytecode, placing the resulting bytecode in a `".zo"` file within the directory specified by *dest-dir*. If *dest-dir* is `#f`, each bytecode result is placed in the same directory as its source file. If *dest-dir* is `'auto`, each bytecode file is placed in a `"compiled"` subdirectory relative to the source; the directory is created if necessary.

If *expr* is anything other than `#f`, then a namespace is created for compiling the files that are supplied later, and *expr* is evaluated to initialize the created namespace. For example, *expr* might load a set of macros. In addition, the expansion-time part of each expression later compiled is evaluated in the namespace before being compiled, so that the effects are visible when compiling later expressions.

If *expr* is `#f`, then no compilation namespace is created (the current namespace is used), and expressions in the files are assumed to compile independently (so there's no need to evaluate the expansion-time part of an expression to compile).

Typically, *expr* is `#f` for compiling `module` files, and it is `(void)` for compiling files with top-level definitions and expressions.

If *module?* is `#t`, then the given files are read and compiled as modules (so there is no dependency on the current namespace's top-level environment).

If *verbose?* is `#t`, the output file for each given file is reported through the current output port.

```
(compile-collection-zos
 collection ...+
 [#:skip-path skip-path
  #:skip-paths skip-paths
  #:skip-doc-sources? skip-docs?
  #:managed-compile-zo managed-compile-zo])
→ void?
 collection : string?
```

```
skip-path : (or/c path-string? #f) = #f
skip-paths : (listof path-string?) = null
skip-docs? : any/c = #f
managed-compile-zo : (path-string? . -> . void?)
                   = (make-caching-managed-compile-zo)
```

Compiles the specified collection's files to `".zo"` files by using *managed-compile-zo* on each source file. The `".zo"` files are placed into the collection's `"compiled"` directory.

By default, all files with the extension `".rkt"`, `".ss"`, or `".scm"` in a collection are compiled, as are all such files within subdirectories; the set of such suffixes is extensible globally as described in `get-module-suffixes`, and `compile-collection-zos` recognizes suffixes from the `'libs` group. However, any file or directory whose path starts with *skip-path* or an element of *skip-paths* is skipped. ("Starts with" means that the simplified complete path *p*'s byte-string form after `(simplify-path p #f)` starts with the byte-string form of `(simplify-path skip-path #f)`; not that each *skip-path* should normally be a complete path.)

The collection compiler reads the collection's `"info.rkt"` file (see §6.4 "`"info.rkt"` File Format") to obtain further instructions for compiling the collection. The following fields are used:

- `name` : The name of the collection as a string, used only for status and error reporting.
- `compile-omit-paths` : Either a list of paths and regexp values or `'all`. In a list, a path is treated as a file that should not be compiled or a directory whose files should not be compiled and whose `"info.rkt"` files should be ignored by `raco setup`; the paths are relative to the collection (i.e., directory containing the `"info.rkt"` file) and can refer to files and directories in subcollections that are that are represented by subdirectories. A regexp in the list is matched against file and directory paths relative to the collection (so, for example, start a regexp with `^` to match only paths in the immediate collection and not in subcollections) to exclude those files and directories from compilation and `raco setup`. The value `'all` is equivalent to specifying all files and directories in the collection (to effectively ignore the collection for compilation). Automatically omitted files and directories are `"compiled"`, `"doc"`, and those whose names start with `.`.

  Files that are required by other files are always compiled in the process of compiling the requiring file—even when the required file is listed with this field or when the field's value is `'all`.
- `compile-omit-files` : A list of filenames (without directory paths) that are not compiled, in addition to the contents of `compile-omit-paths`. Do not use this field; it is for backward compatibility.
- `scribblings` : A list of lists, each of which starts with a path for documentation source. See §6.3 "Controlling `raco setup` with `"info.rkt"` Files" for more infor-

mation. The sources (and the files that they require) are compiled in the same way as other module files, unless *skip-docs?* is a true value.

- `compile-include-files` : A list of filenames (without directory paths) to be compiled, in addition to files that are compiled based on the file's extension, being in `scribblings`, or being `required` by other compiled files.

- `module-suffixes` and `doc-module-suffixes` : Used indirectly via `get-module-suffixes`.

Changed in version 6.3 of package `base`: Added support for `compile-include-files`.
Changed in version 7.8.0.8: Changed "starts with" for *skip-path* to include an exact match.
Changed in version 8.1.0.5: Added support for regexps in `compile-omit-paths`.

```
(compile-directory-zos
 path
 info
 [#:verbose verbose?
  #:skip-path skip-path
  #:skip-paths skip-paths
  #:skip-doc-sources? skip-docs?
  #:managed-compile-zo managed-compile-zo])
→ void?
 path : path-string?
 info : procedure?
 verbose? : any/c = #f
 skip-path : (or/c path-string? #f) = #f
 skip-paths : (listof path-string?) = null
 skip-docs? : any/c = #f
 managed-compile-zo : (path-string? . -> . void?)
                    = (make-caching-managed-compile-zo)
```

Like `compile-collection-zos`, but compiles the given directory rather than a collection. The *info* function behaves like the result of `get-info` to supply `"info.rkt"` fields, instead of using an `"info.rkt"` file (if any) in the directory.

Changed in version 7.8.0.8 of package `base`: Changed *info* handling to use *info* for `'compile-omit-paths`, ignoring any `"info.rkt"` files in parent and child directories.

1.10.2 Recognizing Module Suffixes

```
(require compiler/module-suffix)
```
package: `base`

The `compiler/module-suffix` library provides functions for recognizing file suffixes that correspond to Racket modules for the purposes of compiling files in a directory, running

tests for files in a directory, and so on. The set of suffixes always includes `".rkt"`, `".ss"`, and `".scm"`, but it can be extended globally by `"info.rkt"` configurations in collections.

Added in version 6.3 of package `base`.

```
(get-module-suffixes [#:group group
                      #:mode mode
                      #:namespace namespace]) → (listof bytes?)
  group : (or/c 'all 'libs 'docs) = 'all
  mode : (or/c 'preferred 'all-available 'no-planet 'no-user)
       = 'preferred
  namespace : (or/c #f namespace?) = #f
```

Inspects `"info.rkt"` files (see §6.4 “`"info.rkt"` File Format”) of installed collections to produce a list of file suffixes that should be recognized as Racket modules. Each suffix is reported as a byte string that does not include the `.` that precedes a suffix.

The *mode* and *namespace* arguments are propagated to `find-relevant-directories` to determine which collection directories might configure the set of suffixes. Consequently, suffix registrations are found reliably only if `raco setup` (or package installations or updates that trigger `raco setup`) is run.

The *group* argument determines whether the result includes all registered suffixes, only those that are registered as general library suffixes, or only those that are registered as documentation suffixes. The set of general-library suffixes always includes `".rkt"`, `".ss"`, and `".scm"`. The set of documentation suffixes always includes `".scrbl"`.

The following fields in an `"info.rkt"` file extend the set of suffixes:

- `module-suffixes` : A list of byte strings that correspond to general-library module suffixes (without the `.` that must appear before the suffix). Non-lists or non-byte-string elements of the list are ignored.
- `doc-module-suffixes` : A list of byte strings as for `module-suffixes`, but for documentation modules.

```
(get-module-suffix-regexp [#:group group
                           #:mode mode
                           #:namespace namespace]) → byte-regexp?
  group : (or/c 'all 'libs 'docs) = 'all
  mode : (or/c 'preferred 'all-available 'no-planet 'no-user)
       = 'preferred
  namespace : (or/c #f namespace?) = #f
```

Returns a regexp value that matches paths ending with a suffix as reported by `get-module-suffixes`. The pattern includes a subpatterns for the suffix without its leading `.`

1.10.3 Loading Compiler Support

The compiler unit loads certain tools on demand via `dynamic-require` and `get-info`. If the namespace used during compilation is different from the namespace used to load the compiler, or if other load-related parameters are set, then the following parameter can be used to restore settings for `dynamic-require`.

```
(current-compiler-dynamic-require-wrapper)
 → ((-> any) . -> . any)
(current-compiler-dynamic-require-wrapper proc) → void?
  proc : ((-> any) . -> . any)
```

A parameter whose value is a procedure that takes a thunk to apply. The default wrapper sets the current namespace (via `parameterize`) before calling the thunk, using the namespace in which the `compiler/compiler` library was originally instantiated.

1.10.4 Options for the Compiler

`(require compiler/option)` package: `base`

The `compiler/option` module provides options (in the form of parameters) that control the compiler's behaviors.

More options are defined by the `dynext/compile` and `dynext/link` libraries, which control the actual C compiler and linker that are used for compilation via C.

```
(somewhat-verbose) → boolean?
(somewhat-verbose on?) → void?
  on? : any/c
```

A `#t` value for the parameter causes the compiler to print the files that it compiles and produces. The default is `#f`.

```
(verbose) → boolean?
(verbose on?) → void?
  on? : any/c
```

A `#t` value for the parameter causes the compiler to print verbose messages about its operations. The default is `#f`.

```
(compile-subcollections) → boolean?
(compile-subcollections on?) → void?
  on? : any/c
```

A parameter that specifies whether sub-collections are compiled by `compile-collection-zos`. The default is `#t`.

1.10.5 The Compiler as a Unit

Signatures

```
(require compiler/sig)
```
package: compiler-lib

```
compiler^ : signature
```

Includes all of the names exported by compiler/compiler.

```
compiler:option^ : signature
```

Includes all of the names exported by compiler/option.

Main Compiler Unit

```
(require compiler/compiler-unit)
```
package: compiler-lib

```
compiler@ : unit?
```

Provides the exports of compiler/compiler in unit form, where C-compiler operations are imports to the unit, although they are not used.

The unit imports compiler:option^, dynext:compile^, dynext:link^, and dynext:file^. It exports compiler^.

Options Unit

```
(require compiler/option-unit)
```
package: compiler-lib

```
compiler:option@ : unit?
```

Provides the exports of compiler/option in unit form. It imports no signatures, and exports compiler:option^.

1.11 API for Reading Compilation Dependencies

```
(require compiler/depend)
```
package: base

The compiler/depend module provides a function to inspect and traverse the dependency information generated by raco make, raco setup, or compiler/cm.

Added in version 6.90.0.13 of package base.

```
(module-recorded-dependencies module-file)
 → (listof (and path? (complete-path? path?)))
  module-file : path?
```

Given a *module-file* for a file that has been compiled with `raco make`, `raco setup`, or `compiler/cm`, returns a list of dependencies for *module-file* by reading and traversing dependency-information files left behind by compilation.

2 `raco exe`: Creating Stand-Alone Executables

To achieve a faster startup time, instead of trying `raco exe`, use a smaller base language—such as `#lang racket/base` instead of `#lang racket`. Also, ensure that bytecode files are compiled by using `raco make`. For further improvements, try using `raco demod`.

Compiled code produced by `raco make` relies on Racket executables to provide run-time support to the compiled code. However, `raco exe` can package code together with its run-time support to form an executable, and `raco distribute` can package the executable into a distribution that works on other machines. Running an executable produced by `raco exe` will not improve performance over `raco make`.

The `raco exe` command embeds a module, from source or byte code, into a copy of the `racket` executable. (On Unix, the embedding executable is actually a copy of a wrapper executable.) The created executable invokes the embedded module on startup. The `--gui` flag causes the program to be embedded in a copy of the `gracket` executable. If the embedded module refers to other modules via `require`, then the other modules are also included in the embedding executable.

For example, the command

```
raco exe --gui hello.rkt
```

produces either `"hello.exe"` (Windows), `"hello.app"` (Mac OS), or `"hello"` (Unix), which runs the same as running the `"hello.rkt"` module in `gracket`.

Library modules or other files that are referenced dynamically—through `eval`, `load`, or `dynamic-require`—are not automatically embedded into the created executable. Such modules can be explicitly included using the `++lib` flag to `raco exe`. Alternately, use `define-runtime-path` to embed references to the run-time files in the executable; the files are then copied and packaged together with the executable when creating a distribution (as described in §3 "`raco distribute`: Sharing Stand-Alone Executables"). A submodule is included if its enclosing module is included and the submodule contains a sub-submodule named `declare-preserve-for-embedding` (where the implementation of the sub-submodule is ignored).

Language reader modules that are used only via `#lang` are also not automatically embedded. To support dynamic use of `#lang` with a language specification, supply the `++lang` flag to `raco exe`. The argument after `++lang` can be a language name, but more generally it can be text to appear just after `#lang`. For example, `at-exp racket/base` makes sense as an argument to `++lang` to allow `at-exp` combined with `racket/base` as a language for dynamically loaded modules.

Modules that are implemented directly by extensions—i.e., extensions that are automatically loaded from `(build-path "compiled" "native" (system-library-subpath))` to satisfy a `require`—are treated like other run-time files: a generated executable uses them from their original location, and they are copied and packaged together when creating a distribution.

When a module is embedded in an executable, it gets a symbolic name instead of its original

filesystem-based name. The module-name resolver is configured in the embedding executable to map collection-based module paths to the embedded symbolic name, but no such mapping is created for filesystem paths. By default, a module's symbolic name is generated in an unspecified but deterministic way where the name starts with `#%embedded:`, except that the main module is prefixed with `#%mzc:`. The relative lack of specification for module names can be a problem for language constructs that are sensitive to module names, such as serialization. To take more control over a module's symbolic name, use the `++named-lib` or `++named-file` argument to specify a prefix that is appended before the module's base name to generate a symbolic name.

The `raco exe` command works only with module-based programs. The `compiler/embed` library provides a more general interface to the embedding mechanism.

A stand-alone executable is "stand-alone" in the sense that you can run it without starting `racket`, `gracket`, or DrRacket. However, the executable may depend on Racket shared libraries and possibly other run-time files declared via `define-runtime-path`. Using `--embed-dlls` on Windows or `--orig-exe` on Unix may produce an executable that is more stand-alone than otherwise. Options used when building Racket itself affect the degree to which executables are stand-alone. In any case, the executable can be packaged with support libraries to create a self-contained distribution using `raco distribute`, as described in §3 "`raco distribute`: Sharing Stand-Alone Executables".

Then standard distribution uses options that make executables as stand-alone as possible. For a Unix build, configuring with `--enable-shared` makes executables less stand-alone. For a Mac OS build, configuring without `--enable-embedfw` makes non-GUI executables less stand-alone.

The `raco exe` command accepts the following command-line flags:

- `-o` ⟨*file*⟩ — create the executable as ⟨*file*⟩, adding a suffix to ⟨*file*⟩ as appropriate for the platform and executable type. On Mac OS in `--gui` mode, ⟨*file*⟩ is actually a bundle directory, but it appears as a file within Finder.

- `--gui` — create a graphical executable based on `gracket` instead of `racket`.

- `-l` or `--launcher` — create a launcher (see §2.2 "Installation-Specific Launchers"), instead of a stand-alone executable. Flags such as `--config-path`, `--collects-path`, and `--lib` have no effect on launchers. Beware that the default command-line flags to build into the launcher prevent access to packages that are installed in user scope; use `--exf -U` to enable access to user-scope packages from the launcher.

- `--embed-dlls` — On Windows, for a stand-alone executable, copies any needed DLLs into the executable. Embedding DLLs makes the resulting executable truly stand-alone if it does not depend on other external files. Not all DLLs work with embedding, and limitations are mostly related to thread-local storage and resources, but all DLLs within the main Racket distribution work with `--embed-dlls`.

- `--config-path` ⟨*path*⟩ — set ⟨*path*⟩ within the executable as the path to the configuration directory; if the path is relative, it will be treated as relative to the executable. The default path is `"etc"`, with the expectation that no such directory will exist at run time.

- `--collects-path` ⟨*path*⟩ — set ⟨*path*⟩ within the executable as the path to the main collection directory; if the path is relative, it will be treated as relative to the executable. The default is to have no path, which means that the `current-library-collection-paths` and `current-library-collection-links` parameters are initialized as `null` when the executable starts. Beware that various other directories are located relative to the main collection directory by default (see §19 "Installation Configuration and Search Paths"), so that installing ⟨*path*⟩ may allow other directories to be found—intentional or not.

- `--collects-dest` ⟨*path*⟩ — write modules to be included with the executable into ⟨*path*⟩ (relative to the current directory), instead of embedded within the executable. The `--collects-dest` flag normally makes sense only in combination with `--collects-path`. This mode currently does not prune unreferenced submodules (and it pulls along any dependencies of submodules).

- `--ico` ⟨*.ico-path*⟩ — on Windows, set the icons for the generated executable to ones extracted from ⟨*.ico-path*⟩; see `create-embedding-executable`'s use of the `'ico` auxiliary association for more information about expected icon sizes and transformations.

- `--icns` ⟨*.icns-path*⟩ — on Mac OS, set the icons for the generated executable to be the content of ⟨*.icns-path*⟩.

- `--orig-exe` — on Unix, generate an executable based on the original `racket` or `gracket` executable, instead of a wrapper executable that redirects to the original. If the original executable is statically linked to the Racket runtime library, then the resulting executable is similarly stand-alone. Beware that if the original executable links to Racket as a shared library, however, then `raco distribute` cannot work with executables that are created with `--orig-exe` (because the wrapper executable normally takes care of finding the shared libraries when the executable is distributed to a different machine).

- `--cs` — generate an executable based on the CS implementation of Racket, which is the default unless running a `raco exe` that is based on the BC implementation.

- `--3m` — generate an executable based on the 3m variant of Racket, which is the default only when running a `raco exe` that is based on the 3m variant of the BC implementation.

- `--cgc` — generate an executable based on the CGC variant of Racket, which is the default only when running a `raco exe` that is based on the CGC variant of the BC implementation.

- `++aux` ⟨*file*⟩ — attach information to the executable based on ⟨*file*⟩'s suffix; see `extract-aux-from-path` for a list of recognized suffixes and meanings, and see `create-embedding-executable`'s use of auxiliary association for more specific information about how each kind of file is used.

- `++lib` ⟨*module-path*⟩ — include ⟨*module-path*⟩ in the executable, even if it is not referenced by the main program, so that it is available via `dynamic-require`.

- `++lang` ⟨*lang*⟩ — include modules needed to load modules starting `#lang` ⟨*lang*⟩ dynamically. The ⟨*lang*⟩ does not have to be a plain language or module name; it might be a more general text sequence, such as `at-exp racket/base` to support language constructors like `at-exp`. The initial `require` for a `module` read as ⟨*lang*⟩ must be available though the language reader's `get-info` function and the `'module-language` key; languages implemented with `syntax/module-reader` support that key automatically.

- `++named-lib` ⟨*prefix*⟩ ⟨*module-path*⟩ — like `++lib`, but the embedded module's symbolic name is specified to be ⟨*prefix*⟩ appended before the library file's base name. Specifying a module's symbolic name can be useful with language constructs that depend reflexively on a module name, such as a serialization format (where a module name is record so that a function can be found later for deserialization).

- `++named-file` ⟨*prefix*⟩ ⟨*file-path*⟩ — include ⟨*file-path*⟩ in the executable, even if it is not referenced by the main program, and use ⟨*prefix*⟩ before the file's base name as the embedded module's symbolic name. Since the embedded module's symbolic name is predictable, the module might be accessed at run time via `dynamic-require`. A predictable module name can also help with serialized data in the same way as `++named-lib`.

- `++exf` ⟨*flag*⟩ — provide the ⟨*flag*⟩ command-line argument on startup to the embedded `racket` or `gracket`.

- `--exf` ⟨*flag*⟩ — remove ⟨*flag*⟩ from the command-line arguments to be provided on startup to the embedded `racket` or `gracket`.

- `--exf-clear` — remove all command-line arguments to be provided on startup to the embedded `racket` or `gracket`.

- `--exf-show` — show (without changing) the command-line arguments to be provided on startup to the embedded `racket` or `gracket`.

- `-v` — report progress verbosely.

- `--vv` — report progress more verbosely than `-v`.

Changed in version 6.3.0.11: Added support for `declare-preserve-for-embedding`.

Changed in version 6.90.0.23: Added `--embed-dlls`.

Changed in version 7.0.0.17: Added `++lang`.

Changed in version 7.3.0.6: Added `++named-lib` and `++named-file`, and changed generation of symbolic names for embedded modules to make it deterministic.

2.1 API for Creating Executables

`(require compiler/embed)` package: `base`

The `compiler/embed` library provides a function to embed Racket code into a copy of Racket or GRacket, thus creating a stand-alone Racket executable. To package the executable into a distribution that is independent of your Racket installation, use `assemble-distribution` from `compiler/distribute`.

Embedding walks the module dependency graph to find all modules needed by some initial set of top-level modules, compiling them if needed, and combining them into a "module bundle." In addition to the module code, the bundle extends the module name resolver, so that modules can be `required` with their original names, and they will be retrieved from the bundle instead of the filesystem.

The `create-embedding-executable` function combines the bundle with an executable (Racket or GRacket). The `write-module-bundle` function prints the bundle to the current output port, instead; this stream can be `load`ed directly by a running program, as long as the `read-accept-compiled` parameter is true.

```
(create-embedding-executable
 dest
 #:modules mod-list
 [#:early-literal-expressions early-literal-sexps
 #:configure-via-first-module? config-via-first?
 #:literal-files literal-files
 #:literal-expression literal-sexp
 #:literal-expressions literal-sexps
 #:cmdline cmdline
 #:gracket? gracket?
 #:mred? mred?
 #:variant variant
 #:aux aux
 #:collects-path collects-path
 #:collects-dest collects-dest
 #:launcher? launcher?
 #:verbose? verbose?
 #:expand-namespace expand-namespace
 #:compiler compile-proc
 #:src-filter src-filter
 #:on-extension ext-proc
 #:get-extra-imports extras-proc])
→ void?
 dest : path-string?
```

```
mod-list : (listof (or/c (list/c (or/c symbol? #f #t)
                                 (or/c module-path? path?))
                         (list/c (or/c symbol? #f #t)
                                 (or/c module-path? path?)
                                 (listof symbol?))))
early-literal-sexps : list? = null
config-via-first? : any/c = #f
literal-files : (listof path-string?) = null
literal-sexp : any/c = #f
literal-sexps : list? = (if literal-sexp
                            (list literal-sexp)
                            null)
cmdline : (listof string?) = null
gracket? : any/c = #f
mred? : any/c = #f
variant : (or/c 'cgc '3m 'cs) = (system-type 'gc)
aux : (listof (cons/c symbol? any/c)) = null
collects-path : (or/c #f                    = #f
                      path-string?
                      (listof path-string?))
collects-dest : (or/c #f path-string?) = #f
launcher? : any/c = #f
verbose? : any/c = #f
expand-namespace : namespace? = (current-namespace)
compile-proc : (any/c . -> . compiled-expression?)
             = (lambda (e)
                 (parameterize ([current-namespace
                                 expand-namespace])
                   (compile e)))
src-filter : (path? . -> . any) = (lambda (p) #t)
ext-proc : (or/c #f (path-string? boolean? . -> . any)) = #f
extras-proc : (path? compiled-module-expression?
                . -> . (listof module-path?))
            = (lambda (p m) null)
```

Copies the Racket (if *gracket?* and *mred?* are `#f`) or GRacket (otherwise) binary, embedding code into the copied executable to be loaded on startup. On Unix, the binary is actually a wrapper executable that `exec`s the original; see also the `'original-exe?` tag for *aux*.

The embedding executable is written to *dest*, which is overwritten if it exists already (as a file or directory).

The embedded code consists of module declarations followed by additional (arbitrary) code. When a module is embedded, every module that it imports is also embedded. Library modules are embedded so that they are accessible via their `lib` paths in the initial namespace.

The `#:modules` argument *`mod-list`* designates modules to be embedded, as described below. The `#:early-literal-expressions`, `#:literal-files`, and `#:literal-expressions` arguments specify literal code to be copied into the executable: each element of *`early-literal-sexps`* is copied in order, then the content of each file in *`literal-files`* in order (with no intervening spaces), and then each element of *`literal-sexps`*. The *`literal-files`* files or *`early-literal-sexps`* or *`literal-sexps`* lists can contain compiled bytecode, and it's possible that the content of the *`literal-files`* files only parse when concatenated; the files and expression are not compiled or inspected in any way during the embedding process. Beware that the initial namespace contains no bindings; use compiled expressions to bootstrap the namespace. The `#:literal-expression` (singular) argument is for backward compatibility.

If the `#:configure-via-first-module?` argument is specified as true, then the language of the first module in *`mod-list`* is used to configure the run-time environment before the expressions added by `#:literal-files` and `#:literal-expressions` are evaluated, but after the expressions of `#:early-literal-expressions`. See also §18.1.5 "Language Run-Time Configuration".

The `#:cmdline` argument *`cmdline`* contains command-line strings that are prefixed onto any actual command-line arguments that are provided to the embedding executable. A command-line argument that evaluates an expression or loads a file will be executed after the embedded code is loaded.

Each element of the `#:modules` argument *`mod-list`* is a two- or three-item list, where the first item is a prefix for the module name, and the second item is a module path datum (that's in the format understood by the default module name resolver), and the third is a list of submodule names to be included if they are available. The prefix can be a symbol, `#f` to indicate no prefix, or `#t` to indicate an auto-generated prefix. For example,

```
'((#f "m.rkt"))
```

embeds the module `m` from the file `"m.rkt"`, without prefixing the name of the module; the `literal-sexpr` argument to go with the above might be `'(require m)`. When submodules are available and included, the submodule is given a name by symbol-appending the `write` form of the submodule path to the enclosing module's name.

When an embedded module is not listed in the `#:modules` argument or not given a prefix there, a symbolic name for the embedded module is generated automatically. The names are generated in a deterministic but unspecified way, so that they are not conveniently accessible. The generated names may depend on the path of the first element of *`mod-list`*. Modules that were included via a collection-based path remain accessible at run time through their collection-based paths (via a module name resolver that is installed for the embedding executable).

Modules are normally compiled before they are embedded into the target executable; see also `#:compiler` and `#:src-filter` below. When a module declares run-time paths via

`define-runtime-path`, the generated executable records the path (for use both by immediate execution and for creating a distribution that contains the executable).

If *collects-dest* is a path instead of `#f`, then instead of embedding collection-based modules into the executable, the modules (in compiled form, only) are copied into collections in the *collects-dest* directory.

The optional `#:aux` argument is an association list for platform-specific options (i.e., it is a list of pairs where the first element of the pair is a key symbol and the second element is the value for that key). See also `build-aux-from-path`. The currently supported keys are as follows:

- `'icns` (Mac OS) : An icon file path (suffix `".icns"`) to use for the executable's desktop icon.
- `'ico` (Windows) : An icon file path (suffix `".ico"`) to use for the executable's desktop icon.

  Changed in version 6.3 of package `base`: All icons in the executable are replaced with icons from the file, instead of setting only certain sizes and depths.
- `'creator` (Mac OS) : Provides a 4-character string to use as the application signature.
- `'file-types` (Mac OS) : Provides a list of association lists, one for each type of file handled by the application; each association is a two-element list, where the first (key) element is a string recognized by Finder, and the second element is a plist value (see `xml/plist`). See `"drracket.filetypes"` in the `"drracket"` collection for an example.
- `'uti-exports` (Mac OS) : Provides a list of association lists, one for each Uniform Type Identifier (UTI) exported by the executable; each association is a two-element list, where the first (key) element is a string recognized in a UTI declaration, and the second element is a plist value (see `xml/plist`). See `"drracket.utiexports"` in the `"drracket"` collection for an example.
- `'resource-files` (Mac OS) : extra files to copy into the `"Resources"` directory of the generated executable.
- `'config-dir` : A string/path to a directory that contains configuration information, such as `"config.rtkd"` (see §19 "Installation Configuration and Search Paths"). If no value is supplied, the path is left as-is and converted to absolute form as needed. If `#f` is supplied, the path is left as-is (in potentially relative form). Note that if *collects-path* is provided as an empty list, then the configuration-directory path is not used by Racket's start up process (in contrast to a normal Racket start-up, where the configuration directory is consulted for information about collection link files).
- `'framework-root` (Mac OS) : A string to prefix the executable's path to the Racket and GRacket frameworks (including a separating slash); note that when the prefix start `"@executable_path/"` works for a Racket-based application, the corresponding

prefix start for a GRacket-based application is `"@executable_path/../../../"`; if `#f` is supplied, the executable's framework path is left as-is, otherwise the original executable's path to a framework is converted to an absolute path if it was relative.

- `'dll-dir` (Windows) : A string/path to a directory that contains Racket DLLs needed by the executable, such as `"racket`⟨*version*⟩`.dll"`, or a boolean; a path can be relative to the executable; if `#f` is supplied, the path is left as-is; if `#t` is supplied, the path is dropped (so that the DLLs must be in the system directory or the user's `PATH`); if no value is supplied the original executable's path to DLLs is converted to an absolute path if it was relative.

- `'embed-dlls?` (Windows) : A boolean indicating whether to copy DLLs into the executable, where the default value is `#f`. Embedded DLLs are instantiated by an internal linking step that bypasses some operating system facilities, so it will not work for all Windows DLLs, but typical DLLs will work as embedded.

- `'subsystem` (Windows) : A symbol, either `'console` for a console application or `'windows` for a consoleless application; the default is `'console` for a Racket-based application and `'windows` for a GRacket-based application; see also `'single-instance?`, below.

- `'single-instance?` (Windows) : A boolean for GRacket-based apps; the default is `#t`, which means that the app looks for instances of itself on startup and merely brings the other instance to the front; `#f` means that multiple instances are expected.

- `'forget-exe?` (Unix, Windows, Mac OS) : A boolean; `#t` for a launcher (see *launcher?* below) does not preserve the original executable name for `(find-system-path 'exec-file)`; one consequence is that library collections will be found relative to the launcher instead of the original executable.

- `'original-exe?` (Unix) : A boolean; `#t` means that the embedding uses the original Racket or GRacket executable, instead of a wrapper binary that `exec`s the original; the default is `#f`.

- `'relative?` (Unix, Windows, Mac OS) : A boolean; `#t` means that, to the degree that the generated executable must refer to another, it can use a relative path (so the executables can be moved together, but not separately), and it implies `#f` for `'config-dir`, `'framework-dir`, and `'dll-dir`, unless those are explicitly provided; a `#f` value (the default) means that absolute paths should be used (so the generated executable can be moved).

- `'wm-class` (Unix) : A string; used as the default `WM_CLASS` program class for the program's windows.

If the `#:collects-path` argument is `#f`, then the created executable maintains its built-in (relative) path to the main `"collects"` directory—which will be the result of `(find-system-path 'collects-dir)` when the executable is run—plus a potential list of other

directories for finding library collections—which are used to initialize the current-library-collection-paths list in combination with the PLTCOLLECTS environment variable. Otherwise, the argument specifies a replacement; it must be either a path, string, or list of paths and strings. In the last case, the first path or string specifies the main collection directory, and the rest are additional directories for the collection search path (placed, in order, after the user-specific "collects" directory, but before the main "collects" directory; then the search list is combined with PLTCOLLECTS, if it is defined). If the list is empty, then (find-system-path 'collects-dir) will return the directory of the executable, but current-library-collection-paths is initialized to an empty list, and use-collection-link-paths is set to false to disable the use of collection links files.

If the #:launcher? argument is #t, then *mod-list* should be null, *literal-files* should be null, and *literal-sexp* should be #f. The embedding executable is created in such a way that (find-system-path 'exec-file) produces the source Racket or GRacket path instead of the embedding executable (but the result of (find-system-path 'run-file) is still the embedding executable), unless 'forget-exe? is associated to a true value in *aux*.

The #:variant argument indicates which variant of the original binary to use for embedding. The default is (system-type 'gc); see also current-launcher-variant.

The #:compiler argument is used to compile the source of modules to be included in the executable (when a compiled form is not already available). It should accept a single argument that is a syntax object for a module form. The default procedure uses compile parameterized to set the current namespace to *expand-namespace*.

The #:expand-namespace argument selects a namespace for expanding extra modules (and for compiling using the default *compile-proc*). Extra-module expansion is needed to detect run-time path declarations in included modules, so that the path resolutions can be directed to the current locations (and, ultimately, redirected to copies in a distribution).

The #:src-filter *src-filter* argument takes a path and returns true if the corresponding file source should be included in the embedding executable in source form (instead of compiled form), #f otherwise. The default returns #f for all paths. Beware that the current output port may be redirected to the result executable when the filter procedure is called. Each path given to *src-filter* corresponds to the actual file name (e.g., ".ss"/".rkt" conversions have been applied as needed to refer to the existing file).

If the #:on-extension argument is a procedure, the procedure is called when the traversal of module dependencies arrives at an extension (i.e., a DLL or shared object). The default, #f, causes a reference to a single-module extension (in its current location) to be embedded into the executable. The procedure is called with two arguments: a path for the extension, and a #f (for historical reasons).

The #:get-extra-imports *extras-proc* argument takes a source pathname and compiled module for each module to be included in the executable. It returns a list of quoted module paths (absolute, as opposed to relative to the module) for extra modules to be in-

cluded in the executable in addition to the modules that the source module `requires`. For example, these modules might correspond to reader extensions needed to parse a module that will be included as source, as long as the reader is referenced through an absolute module path. Each path given to *extras-proc* corresponds to the actual file name (e.g., `".ss"/".rkt"` conversions have been applied as needed to refer to the existing file).

Changed in version 6.90.0.23 of package `base`: Added `embed-dlls?` as an `#:aux` key.
Changed in version 7.3.0.6: Changed generation of symbolic names for embedded modules to make it deterministic.

```
(make-embedding-executable  dest
                            mred?
                            verbose?
                            mod-list
                            literal-files
                            literal-sexp
                            cmdline
                           [aux
                            launcher?
                            variant
                            collects-path]) → void?
  dest : path-string?
  mred? : any/c
  verbose? : any/c
  mod-list : (listof (or/c (list/c (or/c symbol? #f #t)
                                   (or/c module-path? path?))
                           (list/c (or/c symbol? #f #t)
                                   (or/c module-path? path?)
                                   (listof symbol?))))
  literal-files : (listof path-string?)
  literal-sexp : any/c
  cmdline : (listof string?)
  aux : (listof (cons/c symbol? any/c)) = null
  launcher? : any/c = #f
  variant : (or/c 'cgc '3m 'cs) = (system-type 'gc)
  collects-path : (or/c #f                      = #f
                        path-string?
                        (listof path-string?))
```

Old (keywordless) interface to `create-embedding-executable`.

```
(write-module-bundle verbose?
                     mod-list
                     literal-files
                     literal-sexp) → void?
  verbose? : any/c
```

```
  mod-list : (listof (or/c (list/c (or/c symbol? #f #t)
                                   (or/c module-path? path?))
                           (list/c (or/c symbol? #f #t)
                                   (or/c module-path? path?)
                                   (listof symbol?))))
  literal-files : (listof path-string?)
  literal-sexp : any/c
```

Like make-embedding-executable, but the module bundle is written to the current output port instead of being embedded into an executable. The output of this function can be read to load and instantiate *mod-list* and its dependencies, adjust the module name resolver to find the newly loaded modules, evaluate the forms included from *literal-files*, and finally evaluate literal-sexpr. The read-accept-compiled parameter must be true to read the stream.

```
(embedding-executable-is-directory? mred?) → boolean
  mred? : any/c
```

Indicates whether Racket/GRacket executables for the current platform correspond to directories from the user's perspective. The result is currently #f for all platforms.

```
(embedding-executable-is-actually-directory? mred?) → boolean?
  mred? : any/c
```

Indicates whether Racket/GRacket executables for the current platform actually correspond to directories. The result is #t on Mac OS when *mred?* is #t, #f otherwise.

```
(embedding-executable-put-file-extension+style+filters mred?)
 → (or/c string? #f)
    (listof (or/c 'packages 'enter-packages))
    (listof (list/c string? string?))
  mred? : any/c
```

Returns three values suitable for use as the extension, style, and filters arguments to put-file, respectively.

If Racket/GRacket launchers for the current platform were directories from the user's perspective, the style result is suitable for use with get-directory, and the extension result may be a string indicating a required extension for the directory name.

```
(embedding-executable-add-suffix path
                                 mred?) → path-string?
  path : path-string?
  mred? : any/c
```

Adds a suitable executable suffix, if it's not present already.

Changed in version 8.1.0.7 of package `base`: Changed to actually add a suffix, instead of replacing an existing suffix.

2.1.1 Executable Creation Signature

```
(require compiler/embed-sig)      package: compiler-lib
```

```
compiler:embed^ : signature
```

Includes the identifiers provided by `compiler/embed`.

2.1.2 Executable Creation Unit

```
(require compiler/embed-unit)      package: compiler-lib
```

```
compiler:embed@ : unit?
```

A unit that imports nothing and exports `compiler:embed^`.

2.1.3 Finding the Racket Executable

```
(require compiler/find-exe)      package: base
```

```
(find-exe [#:cross? cross?
           #:untethered? untethered?
           gracket?
           variant])                    → path?
  cross? : any/c = #f
  untethered? : any/c = #f
  gracket? : any/c = #f
  variant : (or/c 'cgc '3m 'cs) = (if cross?
                                       (cross-system-type 'gc)
                                       (system-type 'gc))
```

Finds the path to the `racket` or `gracket` (when *gracket?* is true) executable.

If *cross?* is true, the executable is found for the target platform in cross-installation mode.

If *untethered?* is true, then the original executable is found, instead of an executable that is tethered to a configuration or addon directory via `(find-addon-tethered-console-bin-dir)` and related functions.

Changed in version 6.2.0.5 of package `base`: Added the `#:untethered?` argument.
Changed in version 6.3: Added the `#:cross?` argument.

2.2 Installation-Specific Launchers

A *launcher* is similar to a stand-alone executable, but a launcher is usually smaller and can be created more quickly, because it depends permanently on the local Racket installation and the program's sources. In the case of Unix, a launcher is simply a shell script that runs `racket` or `gracket`. Launchers *cannot* be packaged into a distribution using `raco distribute`. The `raco exe` command creates a launcher when the `-l` or `--launcher` flag is specified.

```
(require launcher/launcher)
```
package: `base`

The `launcher/launcher` library provides functions for creating launchers.

2.2.1 Creating Launchers

```
(make-gracket-launcher args
                       dest
                      [aux
                       #:tether-mode tether-mode]) → void?
  args : (listof string?)
  dest : path-string?
  aux : (listof (cons/c symbol? any/c)) = null
  tether-mode : (or/c 'addon 'config #f) = 'addon
```

Creates the launcher *dest*, which starts GRacket with the command-line arguments specified as strings in *args*. Extra arguments passed to the launcher at run-time are appended (modulo special Unix/X flag handling, as described below) to this list and passed on to GRacket. If *dest* exists already, as either a file or directory, it is replaced.

The optional *aux* argument is an association list for platform-specific options (i.e., it is a list of pairs where the first element of the pair is a key symbol and the second element is the value for that key). See also `build-aux-from-path`. See `create-embedding-executable` for a list that applies to both stand-alone executables and launchers on Windows and Mac OS GRacket; the following additional associations apply to launchers:

- `'independent?` (Windows) — a boolean; `#t` creates an old-style launcher that works with any Racket or GRacket binary, like `raco.exe`. No other *aux* associations are used for an old-style launcher.
- `'exe-name` (Mac OS, `'script-3m`, `'script-cgc` or `'script-cs` variant) — provides the base name for a `'3m`-/`'cgc`-/`'cs`-variant launcher, which the script will call

ignoring *args*. If this name is not provided, the script will go through the GRacket executable as usual.

- `'exe-is-gracket` (when `'exe-name` is used) — indicates that `'exe-name` refers to the GRacket executable, which is potentially in a `"lib"` subdirectory instead of with other GUI applications.

- `'relative?` (all platforms) — a boolean, where `#t` means that the generated launcher should find the base GRacket executable through a relative path.

- `'install-mode` (Windows, Unix) — either `'main`, `'user`, `'config-tethered`, or `'addon-tethered`, indicates that the launcher is being installed to an installation-wide place, a user-specific place, an installation-wide place that embeds the configuration path, or a specific place that embeds an addon-directory path; the install mode, in turn, determines whether and where to record `'start-menu`, `'extension-registry`, and/or `'desktop` information.

- `'start-menu` (Windows) — a boolean or real number; `#t` indicates that the launcher should be in the Start menu by an installer that includes the launcher. A number value is treated like `#t`, but also requests that the installer automatically start the application, where the number determines a precedence relative to other launchers that may request starting. A `'start-menu` value is used only when `'install-mode` is also specified.

- `'extension-register` (Windows) — a list of document types for file-extension registrations to be performed by an installer. Each document type is described by a list of six items:

  – a human-readable string describing the document type, such as `"Racket Document"`;

  – a string to use as a key for the document type, such as `"Racket.Document"`;

  – a list of strings, where each string is a file extension without the dot, such as `'("rkt" "rktl" "rktd")`;

  – a path to a file that supplies the icon, such as `"doc.ico"`;

  – a string to represent the command line to handle a document with a matching extension, such as `"\"%1\""`, where the string will be prefixed with a path to the launcher, and where `%1` will be replaced with the document path

  An `'extension-registry` value is used only when `'install-mode` is also specified.

- `'desktop` (Unix) — a string containing the content of a `".desktop"` file for the launcher, where `Exec` and `Icon` entries are added automatically. If an `Exec` entry exists in the string, and if its value starts with a non-empty sequence of alpha-numeric ASCII characters followed by a space, then the space and remainder of the value is appended to the automatically generated value. The `".desktop"` file is written to the directory produced by `(find-apps-dir)` or `(find-user-apps-dir)`. A `'desktop` value is used only when `'install-mode` is also specified.

- `'png` (Unix) : An icon file path (suffix `".png"`) to be referenced by a `".desktop"` file (if any); a `'png` value takes precedence over a `'ico` value, but neither is used unless a `'desktop` value is also present.

- `'ico` (Unix, in addition to more general Windows use) : An icon file path (suffix `".ico"`) that is used in the same way as `'png` if no `'png` value is available.

For Unix/X, the script created by `make-mred-launcher` detects and handles X Windows flags specially when they appear as the initial arguments to the script. Instead of appending these arguments to the end of *args*, they are spliced in after any X Windows flags already listed in *args*. The remaining arguments (i.e., all script flags and arguments after the last X Windows flag or argument) are then appended after the spliced *args*.

The *tether-mode* argument indicates how much to preserve the current installation's tethering to a configuration directory and/or addon directory based on `(find-addon-tether-console-bin-dir)` and `(find-config-tether-console-bin-dir)`. The `'addon` mode allows full tethering, the `'config` mode allows only configuration-directory tethering, and the `#f` mode disables tethering.

Changed in version 6.5.0.2 of package `base`: Added the `#:tether-mode` argument.

```
(make-racket-launcher args dest [aux]) → void?
  args : (listof string?)
  dest : path-string?
  aux : (listof (cons/c symbol? any/c)) = null
```

Like `make-gracket-launcher`, but for starting Racket. On Mac OS, the `'exe-name` *aux* association is ignored.

```
(make-gracket-program-launcher file
                               collection
                               dest)      → void?
  file : string?
  collection : string?
  dest : path-string?
```

Calls `make-gracket-launcher` with arguments that start the GRacket program implemented by *file* in *collection*: `(list "-l-" (string-append collection "/" file))`. The *aux* argument to `make-gracket-launcher` is generated by stripping the suffix (if any) from *file*, adding it to the path of *collection*, and passing the result to `build-aux-from-path`.

```
(make-racket-program-launcher file
                              collection
                              dest)      → void?
  file : string?
```

```
  collection : string?
  dest : path-string?
```

Like make-gracket-program-launcher, but for make-racket-launcher.

```
(install-gracket-program-launcher file
                                  collection
                                  name)      → void?
  file : string?
  collection : string?
  name : string?
```

Same as

```
(make-gracket-program-launcher
 file collection
 (gracket-program-launcher-path name))
```

```
(install-racket-program-launcher file
                                 collection
                                 name)      → void?
  file : string?
  collection : string?
  name : string?
```

Same as

```
(make-racket-program-launcher
 file collection
 (racket-program-launcher-path name))
```

```
(make-mred-launcher args dest [aux]) → void?
  args : (listof string?)
  dest : path-string?
  aux : (listof (cons/c symbol? any/c)) = null
(make-mred-program-launcher file
                            collection
                            dest)      → void?
  file : string?
  collection : string?
  dest : path-string?
(install-mred-program-launcher file
                               collection
                               name)      → void?
  file : string?
  collection : string?
  name : string?
```

Backward-compatible version of make-gracket-launcher, etc., that adds "-I" "scheme/gui/init" to the start of the command-line arguments.

```
(make-mzscheme-launcher args dest [aux]) → void?
  args : (listof string?)
  dest : path-string?
  aux : (listof (cons/c symbol? any/c)) = null
(make-mzscheme-program-launcher file
                                collection
                                dest)      → void?
  file : string?
  collection : string?
  dest : path-string?
(install-mzscheme-program-launcher file
                                   collection
                                   name)      → void?
  file : string?
  collection : string?
  name : string?
```

Backward-compatible version of make-racket-launcher, etc., that adds "-I" "scheme/init" to the start of the command-line arguments.

2.2.2 Launcher Path and Platform Conventions

```
(gracket-program-launcher-path  name
                               [#:user? user?
                                #:tethered? tethered?
                                #:console? console?]) → path?
  name : string?
  user? : any/c = #f
  tethered? : any/c = #f
  console? : any/c = #f
```

Returns a pathname for an executable called something like *name* in

- the Racket installation — when *user?* is #f and *tethered?* is #f;
- the user's Racket executable directory — when *user?* is #t and *tethered?* is #f;
- an additional executable directory for executables tethered to a particular configuration directory — when *user?* is #f and *tethered?* is #t; or
- an additional executable directory for executables tethered to a particular addon and configuration directory — when *user?* is #t and *tethered?* is #t.

For Windows, the ".exe" suffix is automatically appended to *name*. For Unix, *name* is changed to lowercase, whitespace is changed to -, and the path includes the "bin" subdirectory of the Racket installation. For Mac OS, the ".app" suffix is appended to *name*.

If *console?* is true, then the path is in the console executable directory, such as the one reported by (find-console-bin-dir), instead of the GUI executable directory, such as the one reported by (find-gui-bin-dir).

Changed in version 6.5.0.2 of package base: Added the #:tethered? argument.
Changed in version 6.8.0.2: Added the #:console? argument.

```
(racket-program-launcher-path  name
                               [#:user? user?
                                #:tethered? tethered?
                                #:console? console?]) → path?
  name : string?
  user? : any/c = #f
  tethered? : any/c = #f
  console? : any/c = #f
```

Returns the same path as (gracket-program-launcher-path *name* #:user? *user?* #:tethered *tethered?* #:console? *console?*).

Changed in version 6.5.0.2 of package base: Added the #:tethered? argument.
Changed in version 6.8.0.2: Added the #:console? argument.

```
(gracket-launcher-is-directory?) → boolean?
```

Returns #t if GRacket launchers for the current platform are directories from the user's perspective. For all currently supported platforms, the result is #f.

```
(racket-launcher-is-directory?) → boolean?
```

Like gracket-launcher-is-directory?, but for Racket launchers.

```
(gracket-launcher-is-actually-directory?) → boolean?
```

Returns #t if GRacket launchers for the current platform are implemented as directories from the filesystem's perspective. The result is #t for Mac OS, #f for all other platforms.

```
(racket-launcher-is-actually-directory?) → boolean?
```

Like gracket-launcher-is-actually-directory?, but for Racket launchers. The result is #f for all platforms.

```
(gracket-launcher-add-suffix path-string?) → path?
  path-string? : path
```

Returns a path with a suitable executable suffix added, if it's not present already.

```
(racket-launcher-add-suffix path-string?) → path?
  path-string? : path
```

Like `gracket-launcher-add-suffix`, but for Racket launchers.

```
(gracket-launcher-put-file-extension+style+filters)
 → (or/c string? #f)
    (listof (or/c 'packages 'enter-packages))
    (listof (list/c string? string?))
```

Returns three values suitable for use as the `extension`, `style`, and `filters` arguments to `put-file`, respectively.

If GRacket launchers for the current platform were directories from the user's perspective, the `style` result is suitable for use with `get-directory`, and the `extension` result may be a string indicating a required extension for the directory name.

```
(racket-launcher-put-file-extension+style+filters)
 → (or/c string? #f)
    (listof (or/c 'packages 'enter-packages))
    (listof (list/c string? string?))
```

Like `gracket-launcher-get-file-extension+style+filters`, but for Racket launchers.

```
(mred-program-launcher-path  name
                             [#:user? user?
                              #:tethered? tethered?]) → path?
  name : string?
  user? : any/c = #f
  tethered? : any/c = #f
(mred-launcher-is-directory?) → boolean?
(mred-launcher-is-actually-directory?) → boolean?
(mred-launcher-add-suffix path-string?) → path?
  path-string? : path
(mred-launcher-put-file-extension+style+filters)
 → (or/c string? #f)
    (listof (or/c 'packages 'enter-packages))
    (listof (list/c string? string?))
```

Backward-compatible aliases for `gracket-program-launcher-path`, etc.

Changed in version 6.5.0.2 of package `base`: Added the `#:tethered?` argument.

```
(mzscheme-program-launcher-path  name
                                 [#:user? user?
                                  #:tethered? tethered?]) → path?
  name : string?
  user? : any/c = #f
  tethered? : any/c = #f
(mzscheme-launcher-is-directory?) → boolean?
(mzscheme-launcher-is-actually-directory?) → boolean?
(mzscheme-launcher-add-suffix path-string?) → path?
  path-string? : path
(mzscheme-launcher-put-file-extension+style+filters)
 → (or/c string? #f)
    (listof (or/c 'packages 'enter-packages))
    (listof (list/c string? string?))
```

Backward-compatible aliases for `racket-program-launcher-path`, etc.

Changed in version 6.5.0.2 of package `base`: Added the `#:tethered?` argument.

```
(installed-executable-path->desktop-path exec-path
                                         user?
                                         tethered?)
 → (or/c (and/c path? complete-path?) #f)
  exec-path : path-string?
  user? : any/c
  tethered? : any/c
```

Returns a path for a `".desktop"` file to describe the installed executable at *exec-path*. Only the filename part of *exec-path* is used. The *user?* argument should be true if *exec-path* is installed in a user-specific location (in which case the result path will also be user-specific). The *tethered?* argument should be true for a tethered install. The result can be `#f` only when *tethered?* is true and `find-addon-tethered-apps-dir` (when *user?* is true) or `find-config-tethered-apps-dir` (when *user?* is `#f`) returns `#f`.

Changed in version 8.3.0.11 of package `base`: Added the *tethered?* argument.

```
(installed-desktop-path->icon-path desktop-path
                                   user?
                                   suffix)
 → (and/c path? complete-path?)
  desktop-path : path-string?
  user? : any/c
  suffix : bytes?
```

Returns a path for an icon file to be referenced by the "desktop" file at *desktop-path*. Only the filename part of *desktop-path* is used. The *user?* argument should be true if *desktop-path* is installed in a user-specific location (in which case the result path will also be user-specific). The *suffix* argument provides the icon-file suffix, normally either #"png" or #"ico".

2.2.3 Launcher Configuration

```
(gracket-launcher-up-to-date? dest aux) → boolean?
  dest : path-string?
  aux : (listof (cons/c symbol? any/c))
```

Returns #t if the GRacket launcher *dest* does not need to be updated, assuming that *dest* is a launcher and its arguments have not changed.

```
(racket-launcher-up-to-date? dest aux) → boolean?
  dest : path-string?
  aux : (listof (cons/c symbol? any/c))
```

Analogous to gracket-launcher-up-to-date?, but for a Racket launcher.

```
(build-aux-from-path path) → (listof (cons/c symbol? any/c))
  path : path-string?
```

Creates an association list suitable for use with make-gracket-launcher or create-embedding-executable. It builds associations by adding to *path* suffixes, such as ".icns", checking whether such a file exists, and calling extract-aux-from-path if so. The results from all recognized suffixes are appended together.

```
(extract-aux-from-path path) → (listof (cons/c symbol? any/c))
  path : path-string?
```

Creates an association list suitable for use with make-gracket-launcher or create-embedding-executable. It builds associations by recognizing the suffix of *path*, where the recognized suffixes are as follows:

- ".icns" → 'icns file for use on Mac OS
- ".ico" → 'ico file for use on Windows or Unix
- ".png" → 'png file for use on Unix
- ".lch" → 'independent? as #t (the file content is ignored) for use on Windows

- `".creator"` → `'creator` as the initial four characters in the file for use on Mac OS
- `".filetypes"` → `'file-types` as `read` content (a single S-expression), and `'resource-files` as a list constructed by finding `"CFBundleTypeIconFile"` entries in `'file-types` (and filtering duplicates); for use on Mac OS
- `".utiexports"` → `'uti-exports` as `read` content (a single S-expression); for use on Mac OS
- `".wmclass"` → `'wm-class` as the literal content, removing a trailing newline if any; for use on Unix
- `".desktop"` → `'desktop` as the literal content; for use on Unix
- `".startmenu"` → `'start-menu` as the file content if it `read`s as a real number, `#t` otherwise, for use on Windows
- `".extreg"` → `'extension-register` as `read` content (a single S-expression), but with relative (to the `".extreg"` file) paths converted to absolute paths; for use on Windows

```
(current-launcher-variant) → symbol?
(current-launcher-variant variant) → void?
  variant : symbol?
```

A parameter that indicates a variant of Racket or GRacket to use for launcher creation and for generating launcher names. The default is the result of `(system-type 'gc)`. On Unix and Windows, the possibilities are `'cgc`, `'3m`, and `'cs`. On Mac OS, the `'script-cgc`, `'script-3m`, and `'script-cs` variants are also available for GRacket launchers.

```
(available-gracket-variants) → (listof symbol?)
```

Returns a list of symbols corresponding to available variants of GRacket in the current Racket installation. The list normally includes at least one of `'3m`, `'cgc`, or `'cs`— whichever is the result of `(system-type 'gc)`—and may include the others, as well as `'script-3m`, `'script-cgc`, and/or `'script-cs` on Mac OS.

```
(available-racket-variants) → (listof symbol?)
```

Returns a list of symbols corresponding to available variants of Racket in the current Racket installation. The list normally includes at least one of `'3m`, `'cgc`, or `'cs`—whichever is the result of `(system-type 'gc)`—and may include the others.

```
(mred-launcher-up-to-date? dest aux) → boolean?
  dest : path-string?
  aux : (listof (cons/c symbol? any/c))
```

```
(mzscheme-launcher-up-to-date? dest aux) → boolean?
  dest : path-string?
  aux : (listof (cons/c symbol? any/c))
(available-mred-variants) → (listof symbol?)
(available-mzscheme-variants) → (listof symbol?)
```

Backward-compatible aliases for `gracket-launcher-up-to-date?`, etc.

2.2.4 Launcher Creation Signature

```
(require launcher/launcher-sig)
```
package: `compiler-lib`

```
launcher^ : signature
```

Includes the identifiers provided by `launcher/launcher`.

2.2.5 Launcher Creation Unit

```
(require launcher/launcher-unit)
```
package: `compiler-lib`

```
launcher@ : unit?
```

A unit that imports nothing and exports `launcher^`.

2.3 Mac OS Dynamic Library Paths

```
(require compiler/exe-dylib-path)
```
package: `base`

The `compiler/exe-dylib-path` library provides functions for reading and adjusting dynamic-library references in a Mac OS executable.

Added in version 6.3 of package `base`.

```
(find-matching-library-path exe-path
                            library-str) → (or/c #f string?)
  exe-path : path-string?
  library-str : string?
```

Searches dynamic-linking information in *exe-path* for a library reference whose name includes *library-str* and returns the executable's path to the library for the first match. If no match is found, the result is `#f`.

```
(update-matching-library-path exe-path
                              library-str
                              library-path-str) → void?
  exe-path : path-string?
  library-str : string?
  library-path-str : string?
```

Searches dynamic-linking information in *exe-path* for each library reference whose name includes *library-str* and replaces the executable's path to that library with *library-path-str*.

A single match is expected, and the update assumes enough space for the new path, perhaps because the executable is linked with `-headerpad_max_install_names`.

3 `raco distribute`: Sharing Stand-Alone Executables

The `raco distribute` command combines a stand-alone executable created by `raco exe` with all of the shared libraries that are needed to run it, along with any run-time files declared via `define-runtime-path`. The resulting package can be moved to other machines that run the same operating system.

On Windows and Mac OS, native libraries tend to be included with the output of `raco distribute`. On Unix platforms, native libraries tend not to be included, so system libraries will be used on the host machine. The difference is whether a Racket installation itself includes bundled native libraries or relies on system-installed libraries. Adding a symbolic link in Racket's `"lib"` directory to a system-installed library causes that library to be included with a distribution directory created by `raco distribute`; see also `define-runtime-path`.

After the `raco distribute` command, supply a directory to contain the combined files for a distribution. Each command-line argument is an executable to include in the distribution, so multiple executables can be packaged together. For example, on Windows,

```
raco distribute greetings hello.exe goodbye.exe
```

creates a directory `"greetings"` (if the directory doesn't exist already), and it copies the executables `"hello.exe"` and `"goodbye.exe"` into `"greetings"`. It also creates a `"lib"` sub-directory in `"greetings"` if needed to contain DLLs, and in that case it adjusts the copied `"hello.exe"` and `"goodbye.exe"` to use the DLLs in `"lib"`.

The number of needed support files depends in part on the way that executables for a distribution are created. Supplying `--embed-dlls` or `--orig-exe` to `raco exe` reduces the need for support files, but at the expense of making the distribution larger if it contains multiple executables.

The layout of files within a distribution directory is platform-specific:

- On Windows, executables are put directly into the distribution directory, and DLLs and other run-time files go into a `"lib"` sub-directory.
- On Mac OS, GUI executables go into the distribution directory, other executables go into a `"bin"` subdirectory, and frameworks (i.e., shared libraries) go into a `"lib"` sub-directory along with other run-time files. As a special case, if the distribution has a single `--gui-exe` executable, then the `"lib"` directory is hidden inside the application bundle.
- On Unix, executables go into a `"bin"` subdirectory, shared libraries (if any) go into a `"lib"` subdirectory along with other run-time files, and wrapped executables are placed into a `"lib/plt"` subdirectory with version-specific names. This layout is consistent with Unix installation conventions; the version-specific names for shared libraries and wrapped executables means that distributions can be safely unpacked into a standard place on target machines without colliding with an existing Racket installation or other executables created by `raco exe`.

A distribution also has a `"collects"` directory that is used as the main library collection directory for the packaged executables. By default, the directory is empty. Use the `++collects-copy` flag of `raco distribute` to supply a directory whose content is copied

into the distribution's "collects" directory. The ++collects-copy flag can be used multiple times to supply multiple directories.

When multiple executables are distributed together, then separately creating the executables with raco exe can generate multiple copies of collection-based libraries that are used by multiple executables. To share the library code, instead, specify a target directory for library copies using the --collects-dest flag with raco exe, and specify the same directory for each executable (so that the set of libraries used by all executables are pooled together). Finally, when packaging the distribution with raco distribute, use the ++collects-copy flag to include the copied libraries in the distribution.

3.1 API for Distributing Executables

(require compiler/distribute) package: base

The compiler/distribute library provides a function to perform the same work as raco distribute.

```
(assemble-distribution dest-dir
                       exec-files
                      [#:executables? executables?
                       #:relative-base relative-base
                       #:collects-path path
                       #:copy-collects dirs])        → void?
  dest-dir : path-string?
  exec-files : (listof path-string?)
  executables? : any/c = #t
  relative-base : (or/c path-string? #f) = #f
  path : (or/c #f (and/c path-string? relative-path?)) = #f
  dirs : (listof path-string?) = null
```

Copies the executables in *exec-files* to the directory *dest-dir*, along with DLLs, frameworks, shared libraries, and/or runtime files that the executables need to run a different machine. If *executables?* is #f, then the *exec-files* are treated as plain data files, instead of executables, and they are modified in-place.

The arrangement of the executables and support files in *dest-dir* depends on the platform. In general, assemble-distribution tries to do the Right Thing, but a non-#f value for *relative-base* specifies a path for reaching the assembled content relative to the executable at run time. When *executables?* is #f, then the default access path is *dest-dir*, with its relativeness preserved.

If a #:collects-path argument is given, it overrides the default location of the main "collects" directory for the packaged executables. It should be relative to the *dest-dir* directory (typically inside it).

The content of each directory in the `#:copy-collects` argument is copied into the main `"collects"` directory for the packaged executables.

Changed in version 6.3 of package `base`: Added the `#:executables?` and `#:relative-base` arguments.

3.2 API for Bundling Distributions

```
(require compiler/bundle-dist)
```
package: `compiler-lib`

The `compiler/bundle-dist` library provides a function to pack a directory (usually assembled by `assemble-distribution`) into a distribution file. On Windows, the result is a `".zip"` archive; on Mac OS, it's a `".dmg"` disk image; on Unix, it's a `".tgz"` archive.

```
(bundle-directory dist-file dir [for-exe?]) → void?
  dist-file : file-path?
  dir : file-path?
  for-exe? : any/c = #f
```

Packages *dir* into *dist-file*. If *dist-file* has no extension, a file extension is added automatically (using the first result of `bundle-put-file-extension+style+filters`).

The created archive contains a directory with the same name as *dir*—except on Mac OS when *for-exe?* is true and *dir* contains a single a single file or directory, in which case the created disk image contains just the file or directory. The default for *for-exe?* is `#f`.

Archive creation fails if *dist-file* exists.

```
(bundle-put-file-extension+style+filters)
 → (or/c string? #f)
    (listof (or/c 'packages 'enter-packages))
    (listof (list/c string? string?))
```

Returns three values suitable for use as the `extension`, `style`, and `filters` arguments to `put-file`, respectively to select a distribution-file name.

4 `raco planet`: Automatic Package Distribution

See *PLaneT: Automatic Package Distribution* for information on the `raco planet` command, which is used for managing packages that can be automatically downloaded and installed from the PLaneT server.

5 `raco pkg`: Package Management

See *Package Management in Racket* for information on the `raco pkg` command, which is used for managing external code packages.

6 `raco setup`: Installation Management

The `raco setup` command builds bytecode, documentation, executables, and metadata indexes for all installed collections.

The collections that are built by `raco setup` can be part of the original Racket distribution, installed via the package manager (see *Package Management in Racket*), installed via PLaneT (see *PLaneT: Automatic Package Distribution*), linked via `raco link`, in a directory that is listed in the `PLTCOLLECTS` environment variable, or placed into one of the default collection directories.

The `raco setup` tool itself does not directly support the installation of collections, except through the now-discouraged `-A` flag (see §6.2 "Installing `".plt"` Archives"). The `raco setup` command is used by installation tools such as the package manager or PLaneT. Programmers who modify installed collections may find it useful to run `raco setup` as an alternative to un-installing and re-installing a set of collections.

6.1 Running `raco setup`

With no command-line arguments, `raco setup` finds all of the current collections—see §18.2 "Libraries and Collections"—and compiles libraries in each collection. (Directories that are named `".git"` or `".svn"` are not treated as collections.)

To restrict `raco setup` to a set of collections, provide the collection names as arguments. For example, `raco setup scribblings/raco` would only compile and render the documentation for `raco`, which is implemented in a `"scribblings/raco"` collection.

An optional `"info.rkt"` within the collection can indicate specifically how the collection's files are to be compiled and other actions to take in setting up a collection, such as creating executables or building documentation. See §6.3 "Controlling `raco setup` with `"info.rkt"` Files" for more information.

The `raco setup` command accepts the following command-line flags:

- Constraining to specified collections or PLaneT packages:
  - `--only` — restrict setup to specified collections and PLaneT packages, even if none are specified. This mode is the default if any collection is specified as a command-line argument or through the `-l`, `--pkgs`, or `-P` flag.
  - `-l` ⟨*collection*⟩ ... — constrain setup actions to the specified ⟨*collection*⟩s (i.e., the same as providing ⟨*collections*⟩s without a flag, but with no possibility that a ⟨*collection*⟩ is interpreted as a flag).
  - `--pkgs` ⟨*pkg*⟩ ... — constrain setup actions to collections that are within (or partially within) the named ⟨*pkg*⟩s.

  - `-P` ⟨*owner*⟩ ⟨*package-name*⟩ ⟨*maj*⟩ ⟨*min*⟩ — constrain setup actions to the specified PLaneT package, in addition to any other specified PLaneT packages or collections.
  - `--doc-index` — build collections that implement documentation indexes (when documentation building is enabled), in addition to specified collections.
  - `--tidy` — remove metadata cache information and documentation for non-existent collections or documentation to clean up after removal, even when setup actions are otherwise confined to specified collections. Although tidying is not confined to specified collections, it can be constrained with `--avoid-main` or `--no-user`.

- Constraining to specific tasks:

  - `--clean` or `-c` — delete existing `".zo"` files, thus ensuring a clean build from the source files. The exact set of deleted files can be controlled by `"info.rkt"`; see `clean` for more information. Unless `--no-info-domain` or `-d` is also specified, the `"info.rkt"` cache is cleared. Unless `--no-docs` or `-D` is also specified, the documentation-index database is reset.
  - `--fast-clean` — like `--clean`, but without forcing a bootstrap of `raco setup` from source (which means that `--fast-clean` cannot clean corruption that affects `raco setup` itself).
  - `--no-zo` or `-n` — refrain from compiling source files to `".zo"` files.
  - `--trust-zos` — fix timestamps on `".zo"` files on the assumption that they are already up-to-date (unless the `PLT_COMPILED_FILE_CHECK` environment variable is set to `exists`, in which case timestamps are ignored).
  - `--recompile-only` — disallow recompilation of modules from source, imposing the constraint that each `".zo"` file is up-to-date, needs only a timestamp adjustment, or can be recompiled from an existing `".zo"` in machine-independent format (when compiling to a machine-dependent format).
  - `--recompile-cache` ⟨*dir*⟩ — cache module recompilations (from machine-independent format to machine-dependent format) in ⟨*dir*⟩.
  - `--sync-docs-only` — synchronize or move documentation into place to "build" it, but do not run or render documentation sources.
  - `--no-launcher` or `-x` — refrain from creating executables or installing `man` pages (as specified in `"info.rkt"`; see §6.3 "Controlling `raco setup` with `"info.rkt"` Files").
  - `--no-foreign-libs` or `-F` — refrain from installing foreign libraries (as specified in `"info.rkt"`; see §6.3 "Controlling `raco setup` with `"info.rkt"` Files").
  - `--only-foreign-libs` — disable actions other than installing foreign libraries; equivalent to `-nxiIdD`, except that `--only-foreign-libs` doesn't reject (redundant) specification of those individual flags.

- `--no-install` or `-i` — refrain from running pre-install actions (as specified in `"info.rkt"` files; see §6.3 “Controlling `raco setup` with `"info.rkt"` Files”).
- `--no-post-install` or `-I` — refrain from running post-install actions (as specified in `"info.rkt"` files; see §6.3 “Controlling `raco setup` with `"info.rkt"` Files”).
- `--no-info-domain` or `-d` — refrain from building a cache of metadata information from `"info.rkt"` files. This cache is needed by other tools. For example, `raco` itself uses the cache to locate plug-in tools.
- `--no-docs` or `-D` — refrain from building documentation.
- `--only-extra-docs` — disable actions other than rendering for `--doc-pdf`, `--doc-xelatex`, `--doc-typst-pdf`, or `--doc-markdown`.
- `--doc-pdf` ⟨*dir*⟩ — in addition to building HTML documentation, render documentation to PDF and place files in ⟨*dir*⟩ via Latex rendering and `pdflatex`.
- `--doc-xelatex` ⟨*dir*⟩ — like `--doc-pdf`, but using `xelatex`.
- `--doc-typst-pdf` ⟨*dir*⟩ — like `--doc-pdf`, but via Typst rendering and `typst compile`.
- `--doc-markdown` ⟨*dir*⟩ — in addition to building HTML documentation, render documentation to Markdown (using Scribble’s Markdown backend) and place files in ⟨*dir*⟩.
- `--no-pkg-deps` or `-K` — refrain from checking whether dependencies among libraries are properly reflected by package-level dependency declarations, whether modules are declared by multiple packages, and whether package version dependencies are satisfied. See §6.5 “Package Dependency Checking” for more information.
- `--check-pkg-deps` — checks package dependencies (unless explicitly disabled) even when specific collections are provided to `raco setup`, and even for packages that have no dependency declarations. See §6.5 “Package Dependency Checking” for more information.
- `--fix-pkg-deps` — attempt to correct dependency mismatches by adjusting package `"info.rkt"` files (which makes sense only for packages that are installed as links). See §6.5 “Package Dependency Checking” for more information.
- `--unused-pkg-deps` — attempt to report dependencies that are declared but are unused. Beware that some package dependencies may be intentionally unused (e.g., declared to force installation of other packages as a convenience), and beware that package dependencies may be reported as unused only because compilation of relevant modules has been suppressed. See §6.5 “Package Dependency Checking” for more information.

- Constraining user versus installation setup:

- `--no-user` or `-U` — refrain from any user-specific (as opposed to installation-specific) setup actions.
- `--no-planet` — refrain from any setup actions for PLaneT actions; this flag is implied by `--no-user`.
- `--avoid-main` — refrain from any setup actions that affect the installation, as opposed to user-specific actions.
- `--force-user-docs` — when building documentation, create a user-specific documentation entry point even if it has the same content as the main installation.

• Selecting parallelism and other build modes:

- `--jobs` ⟨*n*⟩, `--workers` ⟨*n*⟩, or `-j` ⟨*n*⟩ — use up to ⟨*n*⟩ parallel processes. By default, `raco setup` uses `(processor-count)` jobs, which typically uses all of the machine's processing cores.
- `--places` — use Racket places for parallel jobs; this mode is the default if Racket places run in parallel.
- `--processes` — use separate processes for parallel jobs; this mode is the default if Racket places cannot run in parallel.
- `--quiet` or `-q` — suppress status and progress output, including verbose output that is otherwise requested. This flags is implied if the `PLT_SETUP_QUIET` environment variable is set (to anything).
- `--verbose` or `-v` — more verbose output about `raco setup` actions.
- `--make-verbose` or `-m` — more verbose output about dependency checks.
- `--compiler-verbose` or `-r` — even more verbose output about dependency checks and compilation.
- `--mode` ⟨*mode*⟩ — use a `".zo"` compiler other than the default compiler, and put the resulting `".zo"` files in a subdirectory (of the usual place) named by ⟨*mode*⟩. The compiler is obtained by using ⟨*mode*⟩ as a collection name, finding a `"zo-compile.rkt"` module in that collection, and extracting its `zo-compile` export. The `zo-compile` export should be a function like `compile`; see the `"errortrace"` collection for an example.
- `--fail-fast` — attempt to break as soon as any error is discovered.
- `--error-out` ⟨*file*⟩ — handle survivable errors by writing ⟨*file*⟩ and exiting as successful, which facilitates chaining multiple `raco setup` invocations in combination with `--error-in`. If there are no errors and ⟨*file*⟩ already exists, it is deleted.
- `--error-in` ⟨*file*⟩ — treat the existence of ⟨*file*⟩ as a "errors were reported by a previous process" error. Typically, ⟨*file*⟩ is created by previous `raco setup` run using `--error-out`. A file for `--error-in` is detected before creating a file via `--error-out`, so the same file can be used to chain a sequence of `raco setup` steps.

  - `--pause` or `-p` — pause for user input if any errors are reported (so that a user has time to inspect output that might otherwise disappear when the `raco setup` process ends).

- Unpacking `".plt"` archives:

  - `-A` ⟨*archive*⟩ `...` — Install each ⟨*archive*⟩; see §6.2 “Installing `".plt"` Archives”.
  - `--force` — for use with `-A`, treat version mismatches for archives as mere warnings.
  - `--all-users` or `-a` — for use with `-A`, install archive into the installation instead of a user-specific location.

- Bootstrapping:

  - `--boot` ⟨*module-file*⟩ ⟨*build-dir*⟩ — For use by directly running `setup` instead of through `raco setup`, loads ⟨*module-file*⟩ in the same way that `raco setup` normally loads itself, auto-detecting the need to start from sources and rebuild the compiled files—even for the compilation manager itself. The ⟨*build-dir*⟩ path is installed as the only path in `current-compiled-file-roots`, so all compiled files go there.
  - `--chain` ⟨*module-file*⟩ ⟨*build-dir*⟩ — Like `--boot`, but adds ⟨*build-dir*⟩ to the start of `current-compiled-file-roots` instead of replacing the current value, which means that libraries already built in the normal location (including the compilation manager itself) will be used instead of rebuilt. This mode makes sense for cross-compilation.

When building `racket`, flags can be provided to `raco setup` as run by `make install` by setting the `PLT_SETUP_OPTIONS` makefile variable. For example, the following command line uses a single process to build collections during an install:

```
make install PLT_SETUP_OPTIONS="-j 1"
```

Running `raco setup` is sensitive to the `PLT_COMPILED_FILE_CHECK` environment variable in the same way as `raco make`. Specifically, if `PLT_COMPILED_FILE_CHECK` is set to `exists`, then `raco make` does not attempt to update a compiled file’s timestamp if the file is not recompiled.

Some additional environment variables are useful for performance debugging:

- `PLT_SETUP_DMS_ARGS` triggers a call to `dump-memory-stats` after each collection is compiled, where the environment variable’s value is parsed with `read` to obtain a list of arguments to `dump-memory-stats`.
- `PLT_SETUP_LIMIT_CACHE` (set to anything) avoids caching compiled-file information across different collections, which is useful to reduce noise when looking for memory leaks.

- `PLT_SETUP_NO_FORCE_GC` (set to anything) suppresses a call to `collect-garbage` that is issued by default for non-parallel builds after each collection is compiled and after each document is run or rendered.
- `PLT_SETUP_SHOW_TIMESTAMPS` (set to anything) appends the current process time after `@` for each status message printed by `raco setup`.

Changed in version 6.1: Added the `--pkgs`, `--check-pkg-deps`, and `--fail-fast` flags.
Changed in version 6.1.1: Added the `--force-user-docs` flag.
Changed in version 6.1.1.6: Added the `--only-foreign-libs` flag.
Changed in version 6.6.0.3: Added support for `PLT_COMPILED_FILE_CHECK`.
Changed in version 7.0.0.19: Added `--places` and `--processes`.
Changed in version 7.2.0.7: Added `--error-in` and `--error-out`.
Changed in version 7.2.0.8: Added `--recompile-only`.
Changed in version 7.9.0.3: Added `PLT_SETUP_NO_FORCE_GC`, `PLT_SETUP_SHOW_TIMESTAMPS`, and `--sync-docs-only`.
Changed in version 8.17.0.2: Added the `--recompile-cache` flag.
Changed in version 9.2.0.4: Added the `--doc-markdown` and `--only-extra-docs` flags.
Changed in version 9.3.0.1: Added the `--doc-xelatex` and `--doc-typst-pdf` flags.
Changed in version 9.3.0.10: Added the `--quiet`/`-q` flag and `PLT_SETUP_QUIET` environment variable.

6.2 Installing `".plt"` Archives

A `".plt"` file is a platform-independent distribution archive for software based on Racket. A typical `".plt"` file can be installed as a package using `raco pkg` (see *Package Management in Racket*), in which case `raco pkg` supplies facilities for uninstalling the package and managing dependencies.

An older approach is to supply a `".plt"` file to `raco setup` with the `-A` flag; the files contained in the `".plt"` archive are unpacked (according to specifications embedded in the `".plt"` file) and only collections specified by the `".plt"` file are compiled and setup. Archives processed in this way can include arbitrary code that is executed at install time, in addition to any actions triggered by the normal collection-setup part of `raco setup`.

Finally, the `raco unpack` (see §11 "`raco unpack`: Unpacking Library Collections") command can list the content of a `".plt"` archive or unpack the archive without installing it as a package or collection.

6.3 Controlling `raco setup` with `"info.rkt"` Files

To compile a collection's files to bytecode, `raco setup` uses the `compile-collection-zos` procedure. That procedure, in turn, consults the collection's `"info.rkt"` file, if it exists, for specific instructions on compiling the collection. See `compile-collection-zos`

for more information on the fields of `"info.rkt"` that it uses, and see §6.4 "`"info.rkt"` File Format" for information on the format of an `"info.rkt"` file.

Additional fields are used by the Racket package manager and are documented in §4 "Package Metadata". The `raco test` command also recognizes additional fields, which are documented in §13.2 "Test Configuration by `"info.rkt"`".

Optional `"info.rkt"` fields trigger additional actions by `raco setup`:

- `scribblings` : `(listof (cons/c string? list?))` — A list of documents to build. Each document in the list is itself represented as a list, where each document's list starts with a string that is a collection-relative path to the document's source file. A document name (which is derived from the source module's name by default) is intended to be globally unique in the same way as a package or module name. See §6.3.1 "Document Descriptions in `"info.rkt"` Files" for more information about a `scribblings` value. Before a document is rendered by `raco setup`, the document's main `part` is adjusted in several ways; see §6.3.2 "Document Setup Adjustments".

- `release-note-files` : `(listof (cons/c string? (cons/c string? list?)))` — A list of release-notes text files to link from the main documentation pages. Each note is itself represented as a list, and the list can specify auxiliary notes that are grouped with the main note.

  A `release-note-files` entry must be a value that can be generated from an expression matching the following *entry* grammar:

```
      entry = (list note ...)

       note = (list label-string note-path)
            | (list label-string note-path order-integer)
            | (list label-string note-path order-integer
                    (list sub-note ...))

   sub-note = (list label-string note-path)
```

  The *order-integer* is used to order notes and defaults to `0`.

- `racket-launcher-names` : `(listof string?)` — A list of executable names to be generated in the installation's executable directory to run Racket-based programs implemented by the collection. A parallel list of library names must be provided by `racket-launcher-libraries` or `racket-launcher-flags`.

  For each name, a launching executable is set up using `make-racket-launcher`. The arguments are `-l-` and ⟨*colls*⟩`/.../`⟨*file*⟩, where ⟨*file*⟩ is the file named by `racket-launcher-libraries` and ⟨*colls*⟩`/...` are the collections (and subcollections) of the `"info.rkt"` file.

  In addition,

```
(build-aux-from-path
 (build-path (collection-path ⟨colls⟩ ...) ⟨suffixless-file⟩))
```

is provided for the optional `aux` argument (for icons, etc.) to `make-racket-launcher`, where ⟨*suffixless-file*⟩ is ⟨*file*⟩ without its suffix.

If `racket-launcher-flags` is provided, it is used as a list of command-line arguments passed to `racket` instead of the above default, allowing arbitrary command-line arguments. If `racket-launcher-flags` is specified together with `racket-launcher-libraries`, then the flags will override the libraries, but the libraries can still be used to specify a name for `build-aux-from-path` (to find related information like icon files etc).

- `racket-launcher-libraries` : `(listof path-string?)` — A list of library names in parallel to `racket-launcher-names`.
- `racket-launcher-flags` : `(listof string?)` — A list of command-line flag lists, in parallel to `racket-launcher-names`.
- `mzscheme-launcher-names`, `mzscheme-launcher-libraries`, and `mzscheme-launcher-flags` — Backward-compatible variant of `racket-launcher-names`, etc.
- `gracket-launcher-names` : `(listof string?)` — Like `racket-launcher-names`, but for GRacket-based executables. The launcher-name list is treated in parallel to `gracket-launcher-libraries` and `gracket-launcher-flags`.
- `gracket-launcher-libraries` : `(listof path-string?)` — A list of library names in parallel to `gracket-launcher-names`.
- `gracket-launcher-flags` : `(listof string?)` — A list of command-line flag lists, in parallel to `gracket-launcher-names`.
- `mred-launcher-names`, `mred-launcher-libraries`, and `mred-launcher-flags` — Backward-compatible variant of `gracket-launcher-names`, etc.
- `copy-foreign-libs` : `(listof (and/c path-string? relative-path?))` — Files to copy into a directory where foreign libraries are found by `ffi-lib`. If `install-platform` is defined, then the files are copied only if the current platform matches the definition.

  On Mac OS, when a Mach-O file is copied, if the copied file includes a library reference that starts `@loader_path/`, and if the referenced library exists in a different location among the paths listed by `(get-lib-search-dirs)`, then the library reference is updated to an absolute path.

  On Unix, when an ELF file is copied, if the copied file includes an RPATH setting of `$ORIGIN` and the file is being installed to a user-specific location, then the file's RPATH is adjusted to `$ORIGIN:` followed by the path to the main installation's library directory as reported by `(find-lib-dir)`.

  On Windows, deleting a previously installed foreign library may be complicated by a lock on the file, if it is in use. To compensate, `raco setup` deletes a foreign-library file by first renaming the file to have the prefix `"raco-setup-delete-"`; it then attempts to delete the renamed file and merely issues a warning on a failure to delete the

renamed file. Meanwhile, in modes where `raco setup` removes uninstalled libraries, it attempts to delete any file in the foreign-library directory whose name starts with `"raco-setup-delete-"` (in an attempt to clean up after previous failures).

- `move-foreign-libs` : `(listof (and/c path-string? relative-path?))` — Like `copy-foreign-libs`, but the original file is removed after it is copied (which makes sense for precompiled packages).

- `copy-shared-files` : `(listof (and/c path-string? relative-path?))` — Files to copy into a directory where shared files are found. If `install-platform` is defined, then the files are copied only if the current platform matches the definition.

  On Windows, uninstalled files are deleted in the same way as for `copy-foreign-libs`, and the name prefix `"raco-setup-delete-"` is similarly special.

- `move-shared-files` : `(listof (and/c path-string? relative-path?))` — Like `copy-shared-files`, but the original file is removed after it is copied (which makes sense for precompiled packages).

- `copy-man-pages` : `(listof (and/c path-string? relative-path? filename-extension))` — Files to copy into a `man` directory. The file suffix determines its category; for example, `.1` should be used for a `man` page describing an executable.

  On Windows, uninstalled files are deleted in the same way as for `copy-foreign-libs`, and the name prefix `"raco-setup-delete-"` is similarly special.

- `move-man-pages` : `(listof (and/c path-string? relative-path? filename-extension))` — Like `copy-man-pages`, but the original file is removed after it is copied (which makes sense for precompiled packages).

- `install-platform` : `platform-spec?` If this specification matches the current platform, the foreign libraries associated with this package are copied or moved into useful locations. See `copy-foreign-libs`, `move-foreign-libs`, `copy-shared-files`, and `move-shared-files`. Also see `matching-platform?` for information on the way that the specification is compared to `(system-type)` and `(system-library-subpath #f)`.

- `install-collection` : `path-string?` — A library module relative to the collection that provides `installer`. The `installer` procedure must accept one, two, three, or four arguments:

  - The first argument is a directory path to the parent of the Racket installation's `"collects"` directory.
  - The second argument, if accepted, is a path to the collection's own directory.
  - The third argument, if accepted, is a boolean indicating whether the collection is installed as user-specific (`#t`) or installation-wide (`#f`).

  - The fourth argument, if accepted, is a boolean indicating whether the collection is installed as installation-wide and should nevertheless avoid modifying the installation; an `installer` procedure that does not accept this argument is never called when the argument would be `#t`. An installer that does accept this argument is called with `#t` to that it can perform user-specific work, even though the collection is installed installation-wide.

- `pre-install-collection` : `path-string?` — Like `install-collection`, except that the corresponding installer procedures are called *before* the normal `".zo"` build, instead of after. The provided procedure is `pre-installer`, so it can be provided by the same file that provides an `installer` procedure.

- `post-install-collection` : `path-string?` — Like `install-collection` for a procedure that is called right after the `install-collection` procedure is executed. The `--no-install` flag can be provided to `raco setup` to disable `install-collection` and `pre-install-collection`, but not `post-install-collection`. The `post-install-collection` function is therefore expected to perform operations that are always needed, even after an installation that contains precompiled files. The provided procedure is `post-installer`, so it can be provided by the same file that provides an `installer` procedure.

- `assume-virtual-sources` : `any/c` — A true value indicates that bytecode files without a corresponding source file should not be removed from `"compiled"` directories, and no files should not be removed when the `--clean` or `-c` flag is passed to `raco setup`.

- `clean` : `(listof path-string?)` — A list of pathnames to be deleted when the `--clean` or `-c` flag is passed to `raco setup`. The pathnames must be relative to the collection. If any path names a directory, each of the files in the directory are deleted, but none of the subdirectories of the directory are checked. If the path names a file, the file is deleted. The default, if this flag is not specified, is to delete all files in the `"compiled"` subdirectory, and all of the files in the platform-specific subdirectory of the compiled directory for the current platform.

  Just as compiling `".zo"` files will compile each module used by a compiled module, deleting a module's compiled image will delete the `".zo"` of each module that is used by the module. More specifically, used modules are determined when deleting a `".dep"` file, which would have been created to accompany a `".zo"` file when the `".zo"` was built by `raco setup` or `raco make` (see §1.3 "Dependency Files"). If the `".dep"` file indicates another module, that module's `".zo"` is deleted only if it also has an accompanying `".dep"` file. In that case, the `".dep"` file is deleted, and additional used modules are deleted based on the used module's `".dep"` file, etc. Supplying a specific list of collections to `raco setup` disables this dependency-based deletion of compiled files.

- `compile-omit-paths`, `compile-omit-files`, and `compile-include-files` — Used indirectly via `compile-collection-zos`.

- `module-suffixes` and `doc-module-suffixes` — Used indirectly via `get-module-suffixes`.

- `language-family` — A list of hash tables, where each hash table describes a language family. See §6.3.3 “Documentation Language Families” for information about the content of ach hash table.

  Added in version 9.0.0.11.

- `main-doc-index` — A collection name (in the sense of `collection-name?`) or a list of collection names to be added to a `raco setup` request when `--doc-index` is specified without `--avoid-main`.

  Added in version 9.0.0.11.

- `user-doc-index` — A collection name (in the sense of `collection-name?`) or a list of collection names to be added to a `raco setup` request when `--doc-index` is specified without `--no-user`.

  Added in version 9.0.0.11.

6.3.1 Document Descriptions in `"info.rkt"` Files

A `scribblings` entry is introduced in §6.3 “Controlling `raco setup` with `"info.rkt"` Files” as having the shape `(listof (cons/c string? list?))`, but it more precisely must be a value that can be generated from an expression matching the following *entry* grammar:

```
          entry = (list doc ...)

            doc = (list src-string)
                | (list src-string flags)
                | (list src-string flags category)
                | (list src-string flags category name)
                | (list src-string flags category name out-k)
                | (list src-string flags category name out-k order-n)

          flags = (list mode-symbol ...)

       category = (list category-name)
                | (list category-name sort-number)
                | (list category-name sort-number lang-fam)

  category-name = symbol
                | string
                | (box string)

       lang-fam = (list string ...)
```

```
name = string
     | #f
```

A document entry *doc* must have at least a *src-string*, and it optionally continues with information on how to build the document. If a document's list contains a second item, *flags*, it must be a list of mode symbols (described below). If a document's list contains a third item, *category*, it must be a list that categorizes the document (described further below). If a document's list contains a fourth item, *name*, it is a name to use for the generated documentation, instead of defaulting to the source file's name (sans extension), where `#f` means to use the default; a non-`#f` value for *name* must fit the grammar of a collection-name element as checked by `collection-name-element?`. If a document's list contains a fifth item, *out-k*, it is used a hint for the number of files to use for the document's cross-reference information; see below. If a document's list contains a fourth item, *order-n*, it is used a hint for the order of rendering; see below.

Each mode symbol in *flags* can be one of the following, where only `'multi-page` is commonly used:

- `'multi-page` : Generates multi-page HTML output, instead of the default single-page format.
- `'main-doc` : Indicates that the generated documentation should be written into the main installation directory, instead of to a user-specific directory. This mode is the default for a collection that is itself located in the main installation.
- `'user-doc` : Indicates that the generated documentation should be written a user-specific directory. This mode is the default for a collection that is not itself located in the main installation.
- `'depends-all` : Indicates that the document should be rebuilt if any other document is rebuilt—except for documents that have the `'no-depend-on` flag.
- `'depends-all-main` : Indicates that the document should be rebuilt if any other document is rebuilt that is installed into the main installation—except for documents that have the `'no-depend-on` flag.
- `'depends-all-user` : Indicates that the document should be rebuilt if any other document is rebuilt that is installed into the user's space—except for documents that have the `'no-depend-on` flag.
- `'depend-family` : Indicates that the document should be rebuilt if the set of registered language families changes. This flag normally should be combined with `'depends-all`, `'depends-all-main`, or `'depends-all-user`; if it is combined with `'depends-all-main`, then it depends only on language families in the main installation.
- `'always-run` : Build the document every time that `raco setup` is run, even if none of its dependencies change.

- `'no-depend-on` : Removes the document for consideration for other dependencies. Furthermore, references from the document to other documents are always direct, instead of potentially indirect (i.e., resolved at document-viewing time and potentially redirected to a remote site).
- `'main-doc-root` : Designates the root document for the main installation. The document that currently has this mode should be the only one with the mode.
- `'user-doc-root` : Designates the root document for the user-specific documentation directory. The document that currently has this mode should be the only one with the mode.
- `'keep-style` : Leave the document's style as-is, instead of imposing the document style for manuals.
- `'no-search` : Build the document without a search box.
- `'every-main-layer` : With `'main-doc`, indicates that the document should be rendered separately at every installation layer (see §6.18 "Layered Installations").

The *category* list specifies how to show the document in the root table of contents and, for the *lang-fam* part, how to classify the documentation's content for searching. This information can be extended or overridden through a `'doc-properties` table within the `tag-prefix` of the document's main `part`, but we first consider *category* on its own:

- A *category* list must start with a *category-name*, which determines where the manual appears in a document listing such as the root documentation page. A category is a symbol, string, or a boxed string. If it is a string or a boxed string, then the string is the category label on the root page (when the document's language families include the language family used for the listing, which is `"Racket"` for the root documentation page). If it is a symbol, then it should be one of the following categories listed below:
  - `'getting-started` : High-level, introductory documentation, typeset at the same level as other category titles.
  - `'core` : A core reference or library for a language family as may be specified with a *lang-fam*.
  - `'racket-core` : A core reference or library for Racket. This category normally should be used only by specific packages in the main Racket distribution. When rendering a listing for a language family other than `"Racket"`, these documents appear after `'library` instead of after `'core`.
  - `'teaching` : Documentation for a teaching language or library. Documents in this category appear after `'language` if `'racket-core` is moved to later.
  - `'language` : Documentation for a prominent programming language. If `'racket-core` is moved to later, documents in this category appear immediately after `'racket-core` and before `'teaching`.

- `'tool` : Documentation for an executable.
- `'gui-library` : Documentation for GUI and graphics libraries.
- `'net-library` : Documentation for networking libraries.
- `'parsing-library` : Documentation for parsing libraries.
- `'tool-library` : Documentation for programming-tool libraries (i.e., not important enough for the more prominent `'tool` category).
- `'interop` : Documentation for interoperability tools and libraries.
- `'drracket-plugin` : Documentation for DrRacket Plugins.
- All string and boxed-string categories as ordered by `string<=?` appear at this point relative to other categories.
- `'library` : Documentation for miscellaneous libraries.
- All documents whose language families do not include the current language family appear at this point, at least for most categories. Documents are ordered by `string<=?` on the first family name; within a language family, they are ordered as in a documentation listing for that language family. A document whose category is `'language`, `'teaching`, `'experimental`, `'legacy`, or `'racket-core` is always listed independent of its language family, however.

  Unless a document's category is a boxed string, the label used for the category in this section is prefixed by the first family name in the document's families. A boxed string avoid this prefixing.
- `'legacy` : Documentation for deprecated libraries, languages, and tools.
- `'experimental` : Documentation for an experimental language or library.
- `'other` : Other documentation.
- `'omit` : Documentation that should not be listed on the root page or indexed for searching.
- `'omit-start` : Documentation that should not be listed on the root page but should be indexed for searching.

If the *category* list is not given, or if the category symbol is unrecognized, the documentation is added to the (`'library`) category.

- If the category list has a second element, *sort-number*, it must be a real number that designates the manual's sorting position with the category; manuals with the same sorting position are ordered alphabetically. For a pair of manuals with sorting numbers *n* and *m*, the groups for the manuals are separated by space if `(truncate (/ n 10))` and `(truncate (/ m 10))` are different.

- If the category list has a third element, *lang-fam*, then it must be a list of strings, where each string names a language family. The default for *lang-fam* is `(list "Racket")`. This language family list is used to organize a listing of all documentation, and is also used for index entries that are extracted from the document and used for searching. For index entries, the document, a part within the document, or an individual index entry may specify its own language family, and *lang-fam* provides only a default for entries that do not otherwise specify a language family.

- If a document's main `part` has a `tag-prefix` hash table that maps `'doc-properties` to another hash table, the inner hash table can override and generalize the *`category`* list:
  - If `'language-family` is mapped to a list of strings, it provides a replacement for *`lang-fam`*.
  - If `'category` is mapped to a hash table *`cat-ht`*, it is used to get a *`category`* replacement specific to a language family. If *`cat-ht`* maps the listing's language family name to a list, that list is used of *`category`*. Otherwise, if *`cat-ht`* maps `'default` to a list, that list is used instead of *`category`*. In either case, the replacement list cannot contain a *`lang-fam`* component; a `'language-family` mapping (as described in the previous bullet) is the only way to replace a *`lang-fam`* component.

The *`out-k`* specification is a hint on whether to break the document's cross-reference information into multiple parts, which can reduce the time and memory use for resolving a cross-reference into the document. It must be a positive, exact integer, and the default is `1`.

The *`order-n`* specification is a hint for ordering document builds, since documentation references can be mutually recursive. The order hint can be any real number. A value of `-10` or less disables running the document in parallel to other documents. The main Racket reference is given a value of `-11`, the search page is given a value of `10`, and the default is `0`.

A directory for pre-rendered documentation is computed from the source file name by starting with the directory of the `"info.rkt"` file, adding `"doc"`, and then using the document name (which is usually the source file's name without a suffix); if such a directory exists and does not have a `"synced.rktd"` file, then it is treated as pre-rendered documentation and moved into place, in which case the documentation source file need not be present. Moving documentation into place may require no movement at all, depending on the way that the enclosing collection is installed, but movement includes adding a `"synced.rktd"` file to represent the installation.

The destination for a rendered document depends on whether the enclosing collection is part of a Racket installation or installed as a package in user package scope. When the document is in user scope, then it is rendered within the package in the same location as for pre-rendered documentation. An exception is when the documentation is declared in `scribblings` with `'depends-all`, `'depends-all-main`, or `'depends-all-user` and without `'every-main-layer`; in that case, it is rendered in a more central location (and not included in a pre-rendered form) as part of the strategy described in §6.3.4 "Rendering Documentation Listings".

If the a document's main `part`'s has a `tag-prefix` as a hash table, if that hash table has `'doc-properties` mapped to another hash table, and if the inner table maps `'supplant` to a string, then `raco setup` copies a rendered `"index.html"` to a sibling directory name by the `'supplant` string. This step, which is performed at the end of the `raco setup` document-rendering phase, is intended to support document listing as described in §6.3.4

“Rendering Documentation Listings”.

Changed in version 6.4: Allow a category to be a string instead of a symbol.
Changed in version 8.9.0.6: Add the 'drracket-plugin category symbol.
Changed in version 8.14.0.5: Added optional *lang-fam* within *category*.
Changed in version 9.0.0.11: Added support for 'depend-family and for 'doc-properties in a document's main part and for boxed-string category names.

6.3.2 Document Setup Adjustments

Before a document is rendered by `raco setup`, its main part is adjusted in several ways:

- The tag-prefix field of the part is adjusted to have the named document's module path as its 'tag-prefix, which means that other documents can refer to the rendered document via secref or other-doc.
- A '(part "top") tag is added to the part's tag field if it is not present already.
- A document-version style property is added using (version) if no document-version property is present already.
- A body-id style property is added with "doc-racket-lang-org" if no body-id property is present already.
- A document-source style property is added with the document's module path.
- A 'show-language-family style property is added.
- A default language family is determined as *lang-fam* from *category* in a scribblings entry or (if not present) the value of a 'default-language-family key in the part's tag-prefix as a hash table (perhaps originally supplied to title). That list of strings, if either, is added as 'language-family to a new table that is paired with 'index-extras (if any) already in the table. That way, *category* or the main part of a document can supply the default language family for all index entries generated from the document.
- When the part's tag-prefix is a hash table with 'doc-properties mapped to a hash table value, the value is recorded for cross references using the tag `(doc-properties "top") combined with the document's module path. This addition allows a 'doc-properties table to configure the document's listing in more general ways than a *category* specification within scribblings as described in §6.3.1 “Document Descriptions in "info.rkt" Files”.

The document's rendering may be further adjusted at the renderer level (see §6.4 “Renderers”), including configuration at the level of CSS or Latex.

6.3.3 Documentation Language Families

```
(require setup/language-family)
```
package: base

Added in version 9.0.0.11 of package base.

A *language family* is a classification used in documentation that affects the way that search results are shown and filtered, and it also affects the way that documentation is categorized and shown in a listing of all documentation. A language family is not merely a module-based language, but instead stands for a set of languages that share a module-naming convention. As a rule of thumb, a language family is distinct enough that it might have its own downloadable distribution.

Language families are declared and used in several ways and places:

- A Racket installation is configured with a default language family, where the default is "Racket"; see get-main-language-family. This default is used by the main documentation listing, and it is used by the documentation-search page as the default language family.
- A language-family definition in an "info.rkt" file in a collection can declare a language family. This declaration is used in rendered documentation to list language families for navigation, and it is used by raco docs to map a language family name to an entry point into documentation and configuration for navigation.

  A language-family definition's value is a list of hash tables, where each table can have the following keys:

  - 'family: The language-family name as a string—technically optional, but effectively required.
  - 'describe-doc: A module path for the source of a document that describes the language family or its representative language. If 'describe-doc is not mapped, then 'doc (if mapped) is used for the description document, or else no description link is provided for the language family.
  - 'start-doc: A module path for the starting document for the language family or its representative language. If 'start-doc is not mapped, then 'doc (if mapped) is used for the starting document. If neither 'start-doc nor 'doc is mapped, but 'family-root is present, then 'family-root determines the starting page. If none of those keys are mapped, a default starting page is used.
  - 'doc: A module path for a document used as 'describe-doc and/or 'start-doc when those are not mapped.
  - 'order: A real number that orders the family relative to other families (higher is earlier in the list, the Racket language is 100, and the default is 0).
  - 'family-root: The name of the document (if any) that should be considered the starting listing for the language family, so that "top" and "up" navigations arrive at this document. It must be a document with the 'main-doc style and also

a 'user-doc plus 'supplant in the 'doc-properties table in tag-prefix for the document's main part.

- Each regular document rendered via raco setup declares a list of language families. The document is considered in each of those families for the purposes of generating a listing, prioritizing search results, or filtering search results. A document's language family can be declared by the *lang-fam* part of the document's scribblings entry in an "info.rkt" file (see §6.3.1 "Document Descriptions in "info.rkt" Files"), or it can be declared with the document source via the document's main part, and specifically within the tag-prefix of the part. The default is normally (list "Racket"), but the analog of title for Rhombus's Scribble dialect injects (list "Rhombus") by default, for example. Individual index entries in a document, which correspond to different results that can be shown by a search, can specify their own language families, so a document that bridges languages can declare different language families for different parts of the document; per-entry information is in a index-desc, where a function like index picks up configuration via part context.

- A module-based language used via #lang can specify a language family for navigating from a programming environment to documentation. See 'documentation-language-family in DrRacket's documentation for more information.

Navigation and searching for HTML can be adapted to a language family at viewing time (as opposed to rendering time) via query parameters:

- fam: A language family, normally a capitalized name like Racket or Rhombus. This query parameter adjust the prioritization and printing of search results, but it does not filter those results.

- famroot: A document name, normally a case-folded name like rhombus, representing an entry point for a language family. This query parameter adjust the document to be reached by following "up" or "top" links while navigating documents, and it is used only when a fam query is also present.

```
(get-language-families [#:user? user?
                        #:namespace namespace]) → (listof hash?)
  user? : any/c = #t
  namespace : (or/c #f namespace?) = #f
```

Finds language-family declarations in "info.rkt" files and returns the hash tables for well-formed entries. The *user?* argument determines whether user-scope language families are included as well as installation-scope declarations. The *namespace* argument is passed along to get-info/full.

6.3.4 Rendering Documentation Listings

```
(require scribblings/main/contents)
```
package: racket-index

Added in version 9.0.0.11 of package `racket-index`.

```
(build-contents
 [#:user? user?
  #:supplant supplant
  #:style style
  #:main-language-family language-family
  #:title-content title-content
  #:self-path self-path
  #:bug-url bug-url
  #:default-category default-category
  #:doc-properties doc-properties
  #:default-language-family default-language-family
  #:version doc-version
  #:date doc-date])
→ pre-part?
 user? : any/c = #f
 supplant : (or/c #f string?) = #f
 style : (or/c style? #f string? symbol? (listof symbol?)) = #f
 language-family : string? = (get-main-language-family)
 title-content : content?
               = (list language-family
                       (element (style #f '(aux)) " Documentation"))
 self-path : (or/c #f path-string?) = #f
 bug-url : (or/c #f string?) = #f
 default-category : list? = '(language)
 doc-properties : hash? = (make-default-doc-properties
                           main-family
                           default-category
                           supplant)
 default-language-family : (or/c #f (nonempty-listof string?))
                         = #f
 doc-version : (or/c #f string?) = #f
 doc-date : (or/c #f string?) = #f
```

Creates the content of a document that lists all installed documentation from the perspective of `main-family`.

The *`language-family`* argument selects the language family used to render the document (i.e., compared to language families of listed documents), while the *`default-language-family`* argument specifies the language families that the document listing declares for itself.

Use this function as follows:

- Create a subcollection (say, `"my-language/scribblings/main"`) with an `"info.rkt"` file and a sub-collectiopn of that one (say, `"my-`

`language/scribblings/main/user`") with its own `"info.rkt"` file. Each layer will have one document.

- In the outer subcollection, create a document (say, `"my-language.scrbl"`) that renders on the assumption that it's in installation scope (i.e., the *user?* argument to `build-contents` should be `#f`). This is the document that `secref` might reasonably refer to, in case that's useful. Make sure that the document is listed in `scribblings` for the subcollection's `"info.rkt"`, and include the `'depends-all` or `'depends-all-main` flag and the `'no-depend-on` flag.

- In the nested collection, create another document (say, `"user-my-language.scrbl"`) that also has the flags `'user-doc`, `'depends-all-user`, and `'no-depend-on` in the nested subcollection's `"info.rkt"` as `scribblings`. Use the category `'(omit)`. The document should render as in user scope (i.e., the *user?* argument to `build-contents` should be `#t`). If the containing package is in installation scope, this document will be rendered only when there are some other user-scope documents installed. This documentation generally should *not* be the target of a cross-reference, because it won't always get rendered.

- In the latter document (perhaps `"user-my-language.scrbl"`), include the `'supplant` key in a `'doc-properties` table in `tag-prefix` for the document's main `part`. The value of `'supplant` should be the name of the destination for the former document (perhaps `"my-language"`). A `raco setup` that builds documentation will arrange for the `"index.html"` of the supplanting document to be added or to overrwite `"index.html"` for the supplanted document. That way, when user-scope documentation exists at all, there will be a user-scope document using the name of the main document, whether or not the main document would render into user scope. It's also important that a `'depends-all` flag puts a document in the user-scope root documentation directory, instead of keeping it in the installed package's directory.

- In the outer subcollection's `"info.rkt"`, include a `main-doc-index` entry with the subcollection's own name. That entry cases the document to be rerendered when packages are installed. Putting the document in a fresh subcollection minimize the work that this rebuild triggers.

- In the nested subcollection's `"info.rkt"`, include a `user-doc-index` entry with the subcollection's own name.

```
(make-default-doc-properties
 [#:language-family language-family
  #:default-category default-category
  #:supplant supplant])
→ hash?
 language-family : string? = (get-main-language-family)
 default-category : list? = '(language)
 supplant : (or/c #f string?) = #f
```

Constructs a hash table suitable for the `#:doc-properties` argument of `build-contents`.

6.4 "info.rkt" File Format

```
#lang info      package: base
 #lang setup/infotab
```

In each collection, a special module file "info.rkt" provides general information about a collection for use by various tools. For example, an "info.rkt" file specifies how to build the documentation for a collection, and it lists plug-in tools for DrRacket or commands for raco that the collection provides.

The fields specified in an "info.rkt" file are documented in §4 "Package Metadata" for packages and in §6.3 "Controlling raco setup with "info.rkt" Files" for collections.

Although an "info.rkt" file contains a module declaration, the declaration has a highly constrained form. It must match the following grammar of *info-module*:

```
    info-module = (module info info-mod-path
                    decl
                    ...)

  info-mod-path = info
                | setup/infotab
                | (lib "info/main.rkt")
                | (lib "setup/infotab.ss")
                | (lib "setup/infotab.rkt")
                | (lib "main.rkt" "info")
                | (lib "infotab.rkt" "setup")
                | (lib "infotab.ss" "setup")

         decl = (define id info-expr)

    info-expr = (quote datum)
              | (quasiquote datum)
              | (if info-expr info-expr info-expr)
              | (info-primitive info-expr ...)
              | id
              | string
              | number
              | boolean

info-primitive = cons
               | car
               | cdr
               | list
               | list*
               | reverse
               | append
               | equal?
               | string-append
```

```
| make-immutable-hash
| hash
| hash-set
| hash-set*
| hash-remove
| hash-clear
| hash-update
| path->string
| build-path
| collection-path
| system-library-subpath
| getenv
```

For example, the following declaration could be the `"info.rkt"` library of the `"games"` collection. It contains definitions for three info tags, `name`, `gracket-launcher-libraries`, and `gracket-launcher-names`.

```
#lang info
(define name "Games")
(define gracket-launcher-libraries '("main.rkt"))
(define gracket-launcher-names     '("PLT Games"))
```

As illustrated in this example, an `"info.rkt"` file can use `#lang` notation, but only with the `info` (or `setup/infotab`) language.

Although `getenv` is allowed in an `info` module, the `get-info` function loads the module with an environment that prunes any variable not listed in the `PLT_INFO_ALLOW_VARS` environment variable, which holds a list of `;`-separated variable names. By default, the set of allowed environment variables is empty.

See also `get-info` from `setup/getinfo`.

Changed in version 6.5.0.2 of package `base`: Added `if`, `equal?`, and `getenv`.

6.5 Package Dependency Checking

When `raco setup` is run with no arguments, after building all collections and documentation, `raco setup` checks package dependencies. Specifically, it inspects compiled files and documentation to check that references across package boundaries are reflected by dependency declarations in each package-level `"info.rkt"` file (see §4 “Package Metadata”).

Unless `--check-pkg-deps` is specified, dependency checking is disabled if any collection is specified for `raco setup`.

Dependency checking in `raco setup` is intended as an aid to package developers to help them declare dependencies correctly. The `raco setup` process itself does not depend on package dependency declarations. Similarly, a package with a missing dependency declaration may install successfully for other users, as long as they happen to have the dependencies

installed already. A missing dependency creates trouble for others who install a package without having the dependency installed already.

Practically every package depends on the "base" package, which includes the collections that are in a minimal variant of Racket. Declaring a dependency on "base" may seem unnecessary, since its collections are always installed. In a future version of Racket, however, the minimal collections may change, and the new set of minimal collections will then have a package name, such as "base2". Declaring a dependency on "base" ensures forward compatibility, and raco setup complains if the declaration is missing.

To accommodate the early stages of package development, missing dependencies are not treated as an error for a package that has no dependency declarations.

6.5.1 Declaring Build-Time Dependencies

A build-time dependency is one that is not present in a package if it is converted to a binary package (see §5 "Source, Binary, and Built Packages"). For example, "tests" and "scribblings" directories are stripped away in a binary package by default, so cross-package references from directories with those names are treated as build dependencies. Similarly, test and doc submodules are stripped away, so references within those submodules create build dependencies.

Build-time-only dependencies can be listed as build-deps instead of deps in a package's "info.rkt" file. Dependencies listed in deps, meanwhile, are treated as both run-time and build-time dependencies. The advantage of using build-deps, instead of listing all dependencies in deps, is that a binary version of the package can install with fewer dependencies.

6.5.2 How Dependency Checking Works

Dependency checking uses ".zo" files, associated ".dep" files (see §1.3 "Dependency Files"), and the documentation index. Dynamic references, such as through dynamic-require, are not visible to the dependency checker; only dependencies via require, define-runtime-module-path-index, and other forms that cooperate with raco make are visible for dependency checking.

Dependency checking is sensitive to whether a dependency is needed only as a build-time dependency. If raco setup detects that a missing dependency could be added as a build-time dependency, it will suggest the addition, but raco setup will not suggest converting a normal dependency to a build-time dependency (since every normal dependency counts as a build-time dependency, too).

6.6 API for Setup

```
(require setup/setup)
```
package: base

```
(setup [#:file file
        #:collections collections
        #:pkgs pkgs
        #:planet-specs planet-specs
        #:make-user? make-user?
        #:avoid-main? avoid-main?
        #:make-docs? make-docs?
        #:make-doc-index? make-doc-index?
        #:force-user-docs? force-user-docs?
        #:check-pkg-deps? check-pkg-deps?
        #:fix-pkg-deps? fix-pkg-deps?
        #:unused-pkg-deps? unused-pkg-deps?
        #:clean? clean?
        #:tidy? tidy?
        #:recompile-only? recompile-only?
        #:recompile-cache recompile-cache
        #:jobs jobs
        #:quiet? quiet
        #:fail-fast? fail-fast?
        #:get-target-dir get-target-dir])   → boolean?
  file : (or/c #f path-string?) = #f
  collections : (or/c #f (listof (listof path-string?))) = #f
  pkgs : (or/c #f (listof string?)) = #f
  planet-specs : (or/c #f
                       (listof (list/c string?
                                       string?
                                       exact-nonnegative-integer?
                                       exact-nonnegative-integer?)))
               = #f
  make-user? : any/c = #t
  avoid-main? : any/c = #f
  make-docs? : any/c = #t
  make-doc-index? : any/c = #f
  force-user-docs? : any/c = #f
  check-pkg-deps? : any/c = #f
  fix-pkg-deps? : any/c = #f
  unused-pkg-deps? : any/c = #f
  clean? : any/c = #f
  tidy? : any/c = #f
  recompile-only? : any/c = #f
  recompile-cache : (or/c path-string? #f) = #f
  jobs : exact-nonnegative-integer? = #f
```

```
quiet : any/c = #f
fail-fast? : any/c = #f
get-target-dir : (or/c #f (-> path-string?)) = #f
```

Runs `raco setup` with various options:

- *file* — if not `#f`, installs *file* as a `".plt"` archive.
- *collections* — if not `#f`, constrains setup to the named collections (along with *pkgs* and *planet-specs*, if any)
- *pkgs* — if not `#f`, constrains setup to the named packages (along with *collections* and *planet-specs*, if any)
- planet-spec — if not `#f`, constrains setup to the named PLaneT packages (along with *collections* and *pkgs*, if any)
- *make-user?* — if `#f`, disables any user-specific setup actions
- *avoid-main?* — if true, avoids setup actions that affect the main installation, as opposed to user directories
- *make-docs?* — if `#f`, disables any documentation-specific setup actions
- *make-doc-index?* — if true, builds documentation index collections in addition to *collections*, assuming that documentation is built
- *force-user-docs?* — if true, then when building documentation, creates a user-specific documentation entry point even if it has the same content as the installation
- *check-pkg-deps?* — if true, enables package-dependency checking even when *collections*, *pkgs*, or *planet-specs* is non-`#f`.
- *fix-pkg-deps?* — if true, implies *check-pkg-deps?* and attempts to automatically correct discovered package-dependency problems
- *unused-pkg-deps?* — if true, implies *check-pkg-deps?* and also reports dependencies that appear to be unused
- *clean?* — if true, enables cleaning mode instead of setup mode
- *tidy?* — if true, enables global tidying of documentation and metadata indexes even when *collections* or *planet-specs* is non-`#f`
- *recompile-only?* — if true, disallows compilation from source, allowing only timestamp adjustments and recompilation from machine-independent form
- *recompile-cache* — if not `#f`, a directory to cache recompilations from machine-independent form to machine-dependent form

- *jobs* — if not #f, determines the maximum number of parallel tasks used for setup
- quiet? — if true, suppresses status and progress output
- *fail-fast?* — if true, breaks the current thread as soon as an error is discovered
- *get-target-dir* — if not #f, treated as a value for current-target-directory-getter

The result is #t if raco setup completes without error, #f otherwise.

Instead of using PLT_COMPILED_FILE_CHECK, setup is sensitive to the use-compiled-file-check parameter.

Changed in version 6.1 of package base: Added the *fail-fast?* argument.
Changed in version 6.1.1: Added the *force-user-docs?* argument.
Changed in version 7.2.0.7: Added the *check-pkg-deps?*, *fix-pkg-deps?* , and *unused-pkg-deps?* arguments.
Changed in version 7.2.0.8: Added the *recompile-only?* argument.
Changed in version 8.17.0.2: Added the *recompile-cache* argument.
Changed in version 9.3.0.10: Added the quiet? argument.

6.6.1 raco setup Unit

```
(require setup/setup-unit)
```
package: compiler-lib

The setup/setup-unit library provides raco setup in unit form. The associated setup/option-sig and setup/option-unit libraries provides the interface for setting options for the run of raco setup.

For example, to unpack a single ".plt" archive "x.plt", set the archives parameter to (list "x.plt") and leave specific-collections as null.

Link the options and setup units so that your option-setting code is initialized between them, e.g.:

```
(compound-unit
  ...
  (link ...
    [((OPTIONS : setup-option^)) setup:option@]
    [() my-init-options@ OPTIONS]
    [() setup@ OPTIONS ...])
  ...)
```

```
setup@ : unit?
```

Imports

- `setup-option^`
- `compiler^`
- `compiler:option^`
- `launcher^`
- `dynext:file^`

and exports nothing. Invoking `setup@` starts the setup process.

6.6.2 Options Unit

```
(require setup/option-unit)
```
package: `compiler-lib`

```
setup:option@ : unit?
```

Imports nothing and exports `setup-option^`.

6.6.3 Options Signature

```
(require setup/option-sig)
```
package: `compiler-lib`

```
setup-option^ : signature
```

Provides parameters used to control `raco setup` in unit form.

```
(setup-program-name) → string?
(setup-program-name name) → void?
  name : string?
```

The prefix used when printing status messages. The default is `"raco setup"`.

```
(setup-compiled-file-paths)
 → (or/c #f (listof (and/c path? relative-path?)))
(setup-compiled-file-paths paths) → void?
  paths : (or/c #f (listof (and/c path? relative-path?)))
```

If not #f, supplies a value like the one for use-compiled-file-paths to control operations such as cleaning, where use-compiled-file-paths may have been set to null to avoid loading bytecode.

Added in version 1.7 of package compiler-lib.

```
(verbose) → boolean?
(verbose on?) → void?
  on? : any/c
```

If on, prints messages from make to stderr. The default is #f.

```
(make-verbose) → boolean?
(make-verbose on?) → void?
  on? : any/c
```

If on, verbose make. The default is #f.

```
(compiler-verbose) → boolean?
(compiler-verbose on?) → void?
  on? : any/c
```

If on, verbose compiler. The default is #f.

```
(clean) → boolean?
(clean on?) → void?
  on? : any/c
```

If on, delete ".zo" and ".so"/".dll"/".dylib" files in the specified collections. The default is #f.

```
(compile-mode) → (or/c path? #f)
(compile-mode path) → void?
  path : (or/c path? #f)
```

If a path is given, use a ".zo" compiler other than plain compile, and build to (build-path "compiled" (compile-mode)). The default is #f.

```
(make-zo) → boolean?
(make-zo on?) → void?
  on? : any/c
```

If on, compile ".zo". The default is #t.

```
(make-info-domain) → boolean?
(make-info-domain on?) → void?
  on? : any/c
```

If on, update "info-domain/compiled/cache.rkt" for each collection path. The default is #t.

```
(make-launchers) → boolean?
(make-launchers on?) → void?
  on? : any/c
```

If on, make collection "info.rkt"-specified launchers and man pages. The default is #t.

```
(make-foreign-lib) → boolean?
(make-foreign-lib on?) → void?
  on? : any/c
```

If on, install collection "info.rkt"-specified libraries. The default is #t.

```
(make-docs) → boolean?
(make-docs on?) → void?
  on? : any/c
```

If on, build documentation. The default is #t.

```
(make-user) → boolean?
(make-user on?) → void?
  on? : any/c
```

If on, build the user-specific collection tree. The default is #t.

```
(make-planet) → boolean?
(make-planet on?) → void?
  on? : any/c
```

If on, build the planet cache. The default is #t.

```
(avoid-main-installation) → boolean?
(avoid-main-installation on?) → void?
  on? : any/c
```

If on, avoid building bytecode in the main installation tree when building other bytecode (e.g., in a user-specific collection). The default is #f.

```
(make-tidy) → boolean?
(make-tidy on?) → void?
  on? : any/c
```

If on, remove metadata cache information and documentation for non-existent collections (to clean up after removal) even when specific-collections or specific-planet-dirs is non-'() or make-only is true. The default is #f.

```
(call-install) → boolean?
(call-install on?) → void?
  on? : any/c
```

If on, call collection `"info.rkt"`-specified setup code. The default is `#t`.

```
(call-post-install) → boolean?
(call-post-install on?) → void?
  on? : any/c
```

If on, call collection `"info.rkt"`-specified post-install code. The default is `#t`.

```
(pause-on-errors) → boolean?
(pause-on-errors on?) → void?
  on? : any/c
```

If on, in the event of an error, prints a summary error and waits for `stdin` input before terminating. The default is `#f`.

```
(parallel-workers) → exact-nonnegative-integer?
(parallel-workers num) → void?
  num : exact-nonnegative-integer?
```

Determines the number of places to use for compiling bytecode and for building the documentation. The default is `(min (processor-count) 8)`.

```
(fail-fast) → boolean?
(fail-fast on?) → void?
  on? : any/c
```

If on, breaks the original thread as soon as an error is discovered. The default is `#f`.

Added in version 1.2 of package `compiler-lib`.

```
(force-unpacks) → boolean?
(force-unpacks on?) → void?
  on? : any/c
```

If on, ignore version and already-installed errors when unpacking a `".plt"` archive. The default is `#f`.

```
(specific-collections) → (listof (listof path-string?))
(specific-collections colls) → void?
  colls : (listof (listof path-string?))
```

A list of collections to set up; the empty list means set-up all collections if the archives list and `specific-planet-dirs` is also `'()`. The default is `'()`.

```
(specific-planet-dirs)
 → (listof (list/c string?
                   string?
                   exact-nonnegative-integer?
                   exact-nonnegative-integer?))
(specific-planet-dirs dir) → void?
  dir : (listof (list/c string?
                        string?
                        exact-nonnegative-integer?
                        exact-nonnegative-integer?))
```

A list of planet package version specs to set up; the empty list means to set-up all planet collections if the archives list and `specific-collections` is also `'()`. The default is `'()`.

```
(make-only) → boolean?
(make-only on?) → void?
  on? : any/c
```

If true, set up no collections if `specific-collections` and `specific-planet-dirs` are both `'()`.

```
(archives) → (listof path-string?)
(archives arch) → void?
  arch : (listof path-string?)
```

A list of `".plt"` archives to unpack; any collections specified by the archives are set-up in addition to the collections listed in specific-collections. The default is `null`.

```
(archive-implies-reindex) → boolean?
(archive-implies-reindex on?) → void?
  on? : any/c
```

If on, when `archives` has a non-empty list of packages, if any documentation is built, then suitable documentation start pages, search pages, and master index pages are rebuilt. The default is `#t`.

```
(current-target-directory-getter) → (-> path-string?)
(current-target-directory-getter thunk) → void?
  thunk : (-> path-string?)
```

A thunk that returns the target directory for unpacking a relative `".plt"` archive; when unpacking an archive, either this or the procedure in `current-target-plt-directory-getter` will be called. The default is `current-directory`.

```
(current-target-plt-directory-getter)
 → (path-string?
    path-string?
    (listof path-string?) . -> . path-string?)
(current-target-plt-directory-getter proc) → void?
  proc : (path-string?
          path-string?
          (listof path-string?) . -> . path-string?)
```

A procedure that takes a preferred path, a path to the parent of the main `"collects"` directory, and a list of path choices; it returns a path for a "plt-relative" install; when unpacking an archive, either this or the procedure in `current-target-directory-getter` will be called, and in the former case, this procedure may be called multiple times. The default is `(lambda (preferred main-parent-dir choices) preferred)`.

6.6.4 Setup Start Module

`(require setup)` package: `base`

The `setup` library implements `raco setup`, including the part that bootstraps `raco setup` if its own implementation needs to be compiled.

When running `setup` via `racket`, supply the `-N raco` to ensure that command-line arguments are parsed the same way as for `raco setup`, as opposed to a legacy command-line mode.

6.7 API for Installing `".plt"` Archives

The `setup/plt-single-installer` module provides a function for installing a single `".plt"` file.

6.7.1 Non-GUI Installer

`(require setup/plt-single-installer)` package: `base`

```
(run-single-installer
  file
  get-dir-proc
 [#:show-beginning-of-file? show-beginning-of-file?])
 → void?
```

```
  file : path-string?
  get-dir-proc : (-> (or/c path-string? #f))
  show-beginning-of-file? : any/c = #f
```

Creates a separate thread and namespace, runs the installer in that thread with the new namespace, and returns when the thread completes or dies. It also creates a custodian (see §14.7 "Custodians") to manage the created thread, sets the exit handler for the thread to shut down the custodian, and explicitly shuts down the custodian when the created thread terminates or dies.

The *get-dir-proc* procedure is called if the installer needs a target directory for installation, and a #f result means that the user canceled the installation. Typically, *get-dir-proc* is current-directory.

If *show-beginning-of-file?* is a true value and the installation fails, then run-single-installer prints the first 1,000 characters of the file (in an attempt to help debug the cause of failures).

```
(install-planet-package file directory spec) → void?
  file : path-string?
  directory : path-string?
  spec : (list/c string? string?
                 (listof string?)
                 exact-nonnegative-integer?
                 exact-nonnegative-integer?)
```

Similar to run-single-installer, but runs the setup process to install the archive *file* into *directory* as the PLaneT package described by *spec*. The user-specific documentation index is not rebuilt, so reindex-user-documentation should be run after a set of PLaneT packages are installed.

```
(reindex-user-documentation) → void?
```

Similar to run-single-installer, but runs only the part of the setup process that rebuilds the user-specific documentation start page, search page, and master index.

```
(clean-planet-package directory spec) → void?
  directory : path-string?
  spec : (list/c string? string?
                 (listof string?)
                 exact-nonnegative-integer?
                 exact-nonnegative-integer?)
```

Undoes the work of install-planet-package. The user-specific documentation index is not rebuilt, so reindex-user-documentation should be run after a set of PLaneT packages are removed.

6.8 API for Finding Installation Directories

```
(require setup/dirs)
```
package: base

The setup/dirs library provides several procedures for locating installation directories. Many of these paths can be configured through the configuration directory (see §19 “Installation Configuration and Search Paths”).

In cross-platform build mode (see §6.14 “API for Cross-Platform Configuration”), the functions provided by setup/dirs generally report target-system paths, instead of current-system paths. The exceptions are get-lib-search-dirs and find-dll-dir, which report current-system paths while get-cross-lib-search-dirs and find-cross-dll-dir report target-system paths.

```
(find-collects-dir) → (or/c path? #f)
```

Returns a path to the installation’s main "collects" directory, or #f if none can be found. A #f result is likely only in a stand-alone executable that is distributed without libraries.

```
(find-user-collects-dir) → path?
```

Returns a path to the user-specific "collects" directory; the directory indicated by the returned path may or may not exist.

The user-specific path depends on at least (find-system-path 'addon-dir) and (get-installation-name).

```
(get-collects-search-dirs) → (listof path?)
```

Returns the same result as (current-library-collection-paths), which means that this result is not sensitive to the value of the use-user-specific-search-paths parameter.

```
(get-main-collects-search-dirs) → (listof path?)
```

Returns a list of paths to installation "collects" directories, normally including the result of find-collects-dir. These directories are normally included in the result of (current-library-collection-paths), but a PLTCOLLECTS setting or change to the parameter may cause them to be omitted. Any other path in (current-library-collection-paths) is treated as user-specific. The directories indicated by the returned paths may or may not exist.

The main-collections search path can be configured via 'collects-search-dirs in "config.rktd" (see §19 “Installation Configuration and Search Paths”).

```
(find-config-dir) → (or/c path? #f)
```

Returns a path to the installation's `"etc"` directory, which contains configuration and package information—including configuration of some of the other directories (see §19 "Installation Configuration and Search Paths"). A `#f` result indicates that no configuration directory is available.

```
(find-links-file) → (or/c path? #f)
```

Returns a path to the installation's collection links file. The file indicated by the returned path may or may not exist. A `#f` result indicates that no links file is available.

See also `'links-file` in §19 "Installation Configuration and Search Paths".

```
(find-user-links-file [vers]) → path?
  vers : string? = (get-installation-name)
```

Returns a path to the user's collection links file. The file indicated by the returned path may or may not exist.

The user-specific path depends on at least `(find-system-path 'addon-dir)` and *vers*.

```
(get-links-search-files) → (listof path?)
```

Returns a list of paths to installation collection links files to search in order. (Normally, the result includes the result of `(find-links-file)`, which is where new installation-wide links are installed by `raco link` or `links`.) The result of `find-user-links-file` is *not* added to the returned list. The files indicated by the returned paths may or may not exist.

See also `'links-search-files` in §19 "Installation Configuration and Search Paths".

```
(find-pkgs-dir) → (or/c path? #f)
```

Returns a path to the directory containing packages with installation scope; the directory indicated by the returned path may or may not exist. A `#f` result indicates that no package-installation directory is available.

See also `'pkgs-dir` in §19 "Installation Configuration and Search Paths".

```
(find-user-pkgs-dir [vers]) → path?
  vers : string? = (get-installation-name)
```

Returns a path to the directory containing packages with user-specific scope for installation name *vers*; the directory indicated by the returned path may or may not exist.

The user-specific path depends on at least `(find-system-path 'addon-dir)` and *vers*.

```
(get-pkgs-search-dirs) → (listof path?)
```

Returns a list of paths to the directories containing packages in installation scope. (Normally, the result includes the result of (find-pkgs-dir), which is where new packages are installed by raco pkg install.) The result of find-user-pkgs-dir is *not* added to the returned list. The directories indicated by the returned paths may or may not exist.

See also 'pkgs-search-dirs in §19 "Installation Configuration and Search Paths".

(find-doc-dir) → (or/c path? #f)

Returns a path to the installation's "doc" directory. The result is #f if no such directory is available.

See also 'doc-dir in §19 "Installation Configuration and Search Paths".

(find-user-doc-dir) → path?

Returns a path to a user-specific "doc" directory. The directory indicated by the returned path may or may not exist.

The user-specific path depends on at least (find-system-path 'addon-dir) and (get-installation-name).

(get-doc-search-dirs) → (listof path?)

Returns a list of paths to search for documentation, not including documentation stored in individual collections. Unless it is configured otherwise, the result includes any non-#f result of (find-doc-dir) and (find-user-doc-dir)—but the latter is included only if the value of the use-user-specific-search-paths parameter is #t.

See also 'doc-search-dirs in §19 "Installation Configuration and Search Paths".

(get-doc-extra-search-dirs) → (listof path?)

Like get-doc-search-dirs, but refrains from adding (find-doc-dir) and (find-user-doc-dir) to the underlying 'doc-search-dirs configuration.

Added in version 8.1.0.6 of package base.

(find-lib-dir) → (or/c path? #f)

Returns a path to the installation's "lib" directory, which contains libraries and other build information. The result is #f if no such directory is available.

See also 'lib-dir in §19 "Installation Configuration and Search Paths".

(find-user-lib-dir) → path?

Returns a path to a user-specific `"lib"` directory; the directory indicated by the returned path may or may not exist.

The user-specific path depends on at least `(find-system-path 'addon-dir)` and `(get-installation-name)`.

```
(get-lib-search-dirs) → (listof path?)
```

Returns a list of paths to search for foreign libraries.

Unless it is configured otherwise, and except in cross-platform build mode, the result includes any non-`#f` result of `(find-lib-dir)` and `(find-user-lib-dir)`—but the latter is included only if the value of the `use-user-specific-search-paths` parameter is `#t`.

In cross-platform build mode (see §6.14 "API for Cross-Platform Configuration"), `get-lib-search-dirs` reports a result suitable for the current system, instead of the target system. See also `get-cross-lib-search-dirs`.

See also `'lib-search-dirs` in §19 "Installation Configuration and Search Paths".

Changed in version 6.1.1.4 of package `base`: Dropped `(find-dll-dir)` from the set of paths to explicitly include in the default.
Changed in version 6.9.0.1: Changed behavior in cross-platform build mode.

```
(get-cross-lib-search-dirs) → (listof path?)
```

Like `get-lib-search-dirs`, but in cross-platform build mode, reports directories for the target system (including any non-`#f` result of `(find-lib-dir)`, etc.) instead of the current system.

Added in version 6.9.0.1 of package `base`.

```
(get-cross-lib-extra-search-dirs) → (listof path?)
```

Like `get-cross-lib-search-dirs`, but refrains from adding `(find-lib-dir)` and `(find-user-lib-dir)` to the underlying `'lib-search-dirs` configuration.

Added in version 8.1.0.6 of package `base`.

```
(find-dll-dir) → (or/c path? #f)
```

Returns a path to the directory that contains DLLs for use with the current executable (e.g., `"libracket.dll"` on Windows). The result is `#f` if no such directory is available, or if no specific directory is available (i.e., other than the platform's normal search path).

In cross-platform build mode (see §6.14 "API for Cross-Platform Configuration"), `find-dll-dir` reports a result suitable for the current system, instead of the target system. See also `find-cross-dll-dir`.

Changed in version 6.9.0.1 of package base: Changed behavior in cross-platform build mode.

(find-cross-dll-dir) → (or/c path? #f)

Like find-dll-dir, but in cross-platform build mode, reports a directory for the target system instead of the current system.

Added in version 6.9.0.1 of package base.

(find-share-dir) → (or/c path? #f)

Returns a path to the installation's "share" directory, which contains installed packages and other platform-independent files. The result is #f if no such directory is available.

See also 'share-dir in §19 "Installation Configuration and Search Paths".

(find-user-share-dir) → path?

Returns a path to a user-specific "share" directory; the directory indicated by the returned path may or may not exist.

The user-specific path depends on at least (find-system-path 'addon-dir) and (get-installation-name).

(get-share-search-dirs) → (listof path?)

Returns a list of paths to search for files that are normally in a "share" directory.

Unless it is configured otherwise, the result includes any non-#f result of (find-share-dir) and (find-user-share-dir)—but the latter is included only if the value of the use-user-specific-search-paths parameter is #t.

See also 'share-search-dirs in §19 "Installation Configuration and Search Paths".

Added in version 8.1.0.6 of package base.

(get-share-extra-search-dirs) → (listof path?)

Like get-share-search-dirs, but refrains from adding (find-share-dir) and (find-user-share-dir) to the underlying 'share-search-dirs configuration.

Added in version 8.1.0.6 of package base.

(find-include-dir) → (or/c path? #f)

Returns a path to the installation's "include" directory, which contains ".h" files for building Racket extensions and embedding programs. The result is #f if no such directory is available.

See also 'include-dir in §19 “Installation Configuration and Search Paths”.

(find-user-include-dir) → path?

Returns a path to a user-specific "include" directory; the directory indicated by the returned path may or may not exist.

The user-specific path depends on at least (find-system-path 'addon-dir) and (get-installation-name).

(get-include-search-dirs) → (listof path?)

Returns a list of paths to search for ".h" files. Unless it is configured otherwise, the result includes any non-#f result of (find-include-dir) and (find-user-include-dir)—but the latter is included only if the value of the use-user-specific-search-paths parameter is #t.

See also 'include-search-dirs in §19 “Installation Configuration and Search Paths”.

(find-console-bin-dir) → (or/c path? #f)

Returns a path to the installation’s executable directory, where the stand-alone Racket executable resides. The result is #f if no such directory is available.

See also 'bin-dir in §19 “Installation Configuration and Search Paths”.

(find-gui-bin-dir) → (or/c path? #f)

Returns a path to the installation’s executable directory, where the stand-alone GRacket executable resides. The result is #f if no such directory is available.

See also 'gui-bin-dir in §19 “Installation Configuration and Search Paths”.

(find-user-console-bin-dir) → path?

Returns a path to the user’s executable directory; the directory indicated by the returned path may or may not exist.

The user-specific path depends on at least (find-system-path 'addon-dir) and (get-installation-name).

(find-user-gui-bin-dir) → path?

Returns a path to the user’s executable directory for graphical programs; the directory indicated by the returned path may or may not exist.

The user-specific path depends on at least (find-system-path 'addon-dir) and (get-installation-name).

```
(get-console-bin-search-dirs) → (listof path?)
```

Analogous to `get-share-search-dirs`, but for paths to search for executables such as `racket`.

See also `'bin-search-dirs` in §19 “Installation Configuration and Search Paths”.

Added in version 8.1.0.6 of package `base`.

```
(get-console-bin-extra-search-dirs) → (listof path?)
```

Analogous to `get-share-extra-search-dirs` for the underlying `'bin-search-dirs` configuration.

Added in version 8.1.0.6 of package `base`.

```
(get-gui-bin-search-dirs) → (listof path?)
```

Analogous to `get-share-search-dirs`, but for paths to search for executables such as `gracket`.

See also `'gui-bin-search-dirs` in §19 “Installation Configuration and Search Paths”.

Added in version 8.1.0.6 of package `base`.

```
(get-gui-bin-extra-search-dirs) → (listof path?)
```

Analogous to `get-share-extra-search-dirs` for the underlying `'gui-bin-search-dirs` configuration.

Added in version 8.1.0.6 of package `base`.

```
(find-apps-dir) → (or/c path? #f)
```

Returns a path to the installation’s directory `".desktop"` files (for Unix). The result is `#f` if no such directory exists.

See also `'apps-dir` in §19 “Installation Configuration and Search Paths”.

```
(find-user-apps-dir) → path?
```

Returns a path to the user’s directory for `".desktop"` files (for Unix); the directory indicated by the returned path may or may not exist.

The user-specific path depends on at least `(find-system-path 'addon-dir)` and `(get-installation-name)`.

```
(find-man-dir) → (or/c path? #f)
```

Returns a path to the installation's man-page directory. The result is `#f` if no such directory exists. See also `'man-dir` in §19 "Installation Configuration and Search Paths".

```
(find-user-man-dir) → path?
```

Returns a path to the user's man-page directory; the directory indicated by the returned path may or may not exist.

The user-specific path depends on at least `(find-system-path 'addon-dir)` and `(get-installation-name)`.

```
(get-man-search-dirs) → (listof path?)
```

Analogous to `get-share-search-dirs`, but for paths to search for man pages.

See also `'man-search-dirs` in §19 "Installation Configuration and Search Paths".

Added in version 8.1.0.6 of package `base`.

```
(get-man-extra-search-dirs) → (listof path?)
```

Analogous to `get-share-extra-search-dirs` for the underlying `'man-search-dirs` configuration.

Added in version 8.1.0.6 of package `base`.

```
(get-info-domain-root) → (or/c #false path?)
```

Returns `#f` or a path to be used as a prefix to redirect the paths used for recording and finding `"info.rkt"` information via `find-relevant-directories`.

Added in version 8.10.0.4 of package `base`.

```
(get-doc-search-url) → string?
```

Returns a string that is used by the documentation system, augmented with a version and search-key query, for remote documentation links.

See also `'doc-search-url` in §19 "Installation Configuration and Search Paths".

```
(get-doc-open-url) → (or/c string? #f)
```

Returns `#f` or a string for a root URL to be used as an alternative to opening a local file for documentation. A non-`#f` configuration means that DrRacket, for example, performs keyword searches for documentation via the specified URL instead of from locally installed documentation.

See also `'doc-open-url` in §19 “Installation Configuration and Search Paths”.

Added in version 6.0.1.6 of package `base`.

```
(get-installation-name config) → string?
  config : (read-installation-configuration-table)
```

Returns the current installation’s name, which is often `(version)`, but an installation name can be set through a combination of a `'installation-name` value in *config* plus a user-specific directory state if `(use-user-specific-search-paths)` is `#t`.

A user-specific result depends on whether a directory `"other-version"` exists within `(find-system-path 'addon-dir)`. If that directory exists, and it no directory with the installation’s configured name exists, then `"other-version"` is used as the installation name. So, by creating the `"other-version"` directory, a user can opt into sharing of packages and collections across installations/versions, while opting out for a specific installation/version by creating a directory with that installation’s name.

Changed in version 8.4.0.3 of package `base`: Added the *config* argument and support for a user-specific installation name.

```
(get-build-stamp) → (or/c #f string?)
```

Returns a string that identifies an installation build, which can be used to augment the Racket version number to more specifically identify the build. An empty string is normally produced for a release build. The result is `#f` if no build stamp is available.

See also `'build-stamp` in §19 “Installation Configuration and Search Paths”.

```
(get-main-language-family) → string?
```

Returns a string that names the installation’s main language family. The default is `"Racket"`.

See also `'main-language-family` in §19 “Installation Configuration and Search Paths”.

Added in version 8.14.0.5 of package `base`.

```
(get-base-documentation-packages) → (listof string?)
(get-distribution-documentation-packages) → (listof string?)
```

Returns a list of package names that represent a distribution’s base-language documentation and all of the documentation that is part of the distribution, respectively. These lists are

used to classify and sort documentation search results. If a package is part of the base documentation, that classification takes precedence over distribution documentation.

See also 'base-documentation-packages and 'distribution-documentation-packages in §19 “Installation Configuration and Search Paths”.

Added in version 8.14.0.5 of package base.

```
(get-absolute-installation?) → boolean?
```

Returns #t if this installation uses absolute path names for executable and library references, #f otherwise.

```
(find-addon-tethered-console-bin-dir) → (or/c #f path?)
(find-addon-tethered-gui-bin-dir) → (or/c #f path?)
(find-addon-tethered-apps-dir) → (or/c #f path?)
```

Returns a path to a user-specific directory to hold an extra copy of each installed executable and ".desktop" file (for Unix), where the extra copy is created by raco setup and tethered to a particular result for (find-system-path 'addon-dir) and (find-config-dir).

Unlike other directories, which are configured via "config.rktd" in the (find-config-dir) directory (see §19 “Installation Configuration and Search Paths”), these paths are configured via 'addon-tethered-console-bin-dir, 'addon-tethered-gui-bin-dir, and 'addon-tethered-apps-dir entries in "config.rktd" in (build-path (find-system-path 'addon-dir) "etc"). If no configuration is present, the result from the corresponding function, find-addon-tethered-console-bin-dir, find-addon-tethered-gui-bin-dir, or find-addon-tethered-apps-dir, is #f instead of a path.

See §6.19 “Tethered Installations” for more information.

Added in version 6.5.0.2 of package base.
Changed in version 8.3.0.11: Added find-addon-tethered-apps-dir.

]

```
(find-config-tethered-console-bin-dir) → (or/c #f path?)
(find-config-tethered-gui-bin-dir) → (or/c #f path?)
(find-config-tethered-apps-dir) → (or/c #f path?)
```

Similar to find-addon-tethered-console-bin-dir, find-addon-tethered-gui-bin-dir, and find-addon-tethered-apps-dir, but configured via "config.rktd" in the (find-config-dir) directory (see §19 “Installation Configuration and Search Paths”) and triggers executables that are tethered only to a particular value of (find-config-dir).

See §6.19 “Tethered Installations” for more information.

Added in version 6.5.0.2 of package `base`.
Changed in version 8.3.0.11: Added `find-addon-tethered-apps-dir`.

6.9 API for Reading `"info.rkt"` Files

`(require setup/getinfo)` package: `base`

The `setup/getinfo` library provides functions for accessing fields in `"info.rkt"` files. The file format for `"info.rkt"` files is documented in §6.4 “`"info.rkt"` File Format”.

```
(get-info  collection-names
           [#:namespace namespace
            #:bootstrap? bootstrap?])
 → (or/c
     ((symbol?) ((-> any)) . ->* . any)
     #f)
  collection-names : (listof string?)
  namespace : (or/c namespace? #f) = #f
  bootstrap? : any/c = #f
```

Accepts a list of strings naming a collection or sub-collection, and calls `get-info/full` with the full path corresponding to the named collection and the *namespace* argument.

```
(get-info/full  path
                [#:namespace namespace
                 #:bootstrap? bootstrap?])
 → (or/c
     ((symbol?) ((-> any)) . ->* . any)
     #f)
  path : path-string?
  namespace : (or/c namespace? #f) = #f
  bootstrap? : any/c = #f
```

Accepts a path to a directory. If it finds either a well-formed `"info.rkt"` file or an `"info.ss"` file (with preference for the `"info.rkt"` file), it returns an info procedure that accepts either one or two arguments. The first argument to the info procedure is always a symbolic name, and the result is the value of the name in the `"info.rkt"` file, if the name is defined. The optional second argument, *thunk*, is a procedure that takes no arguments to be called when the name is not defined; the result of the info procedure is the result of the *thunk* in that case. If the name is not defined and no *thunk* is provided, then an exception is raised.

The `get-info/full` function returns `#f` if there is no `"info.rkt"` (or `"info.ss"`) file in the directory. If there is a `"info.rkt"` (or `"info.ss"`) file that has the wrong shape

(i.e., not a module using `info` or `setup/infotab`), or if the `"info.rkt"` file fails to load, then an exception is raised. If the `"info.rkt"` file loaded, `get-info/full` returns the info procedure. If the `"info.rkt"` file does not exist, then `get-info/full` does the same checks for the `"info.ss"` file, either raising an exception or returning the info procedure from the `"info.ss"` file.

The `"info.rkt"` (or `"info.ss"`) module is loaded into *namespace* if it is not `#f`, or a private, weakly-held namespace otherwise.

If *bootstrap?* is true, then `use-compiled-file-paths` is set to `'()` while reading `"info.rkt"` (or `"info.ss"`), in case an existing compiled file is broken. Furthermore, the `info` and `setup/infotab` modules are attached to *namespace* from the namespace of `get-info/full` before attempting to load `"info.rkt"` (or `"info.ss"`).

As the module is loaded, the environment variable set is pruned to contain only environment variables that are listed in the `PLT_INFO_ALLOW_VARS` environment variable, which contains a `;`-separated list of names. By default, the list of allowed variable names is empty.

Changed in version 6.5.0.2 of package `base`: Added environment-variable pruning and `PLT_INFO_ALLOW_VARS` support.

```
(find-relevant-directories syms [mode]) → (listof path?)
  syms : (listof symbol?)
  mode : (or/c 'preferred 'all-available 'no-planet 'no-user)
       = 'preferred
```

Returns a list of paths identifying collections and installed PLaneT packages whose `"info.rkt"` file defines one or more of the given symbols. The result is based on a cache that is computed by `raco setup`.

Note that the cache may be out of date by the time you call `get-info/full`, so do not assume that every returned directory's `"info.rkt"` file will supply one of the requested symbols.

The result is in a canonical order (sorted lexicographically by directory name), and the paths it returns are suitable for providing to `get-info/full`.

If *mode* is specified, it must be either `'preferred` (the default), `'all-available`, `'no-planet`, or `'no-user`. If *mode* is `'all-available`, `find-relevant-directories` returns all installed directories whose info files contain the specified symbols—for instance, all versions of all installed PLaneT packages will be searched if `'all-available` is specified. If *mode* is `'preferred`, then only a subset of "preferred" packages will be searched: only the directory containing the most recent version of any PLaneT package will be returned. If *mode* is `'no-planet`, then PLaneT packages are not included in the search. If *mode* is `'no-user`, then only installation-wide directories are searched, which means omitting PLaneT package directories.

Regardless of *mode*, note that find-relevant-directories will not consider package-level "info.rkt" files for multi-collection packages, since those files are not part of any collection or PLaneT package. In contrast, a single-collection package's "info.rkt" file is part of a collection, and thus will be considered.

Collection links from the installation-wide collection links file or packages with installation scope are cached with the installation's main "lib" directory, and links from the user-specific collection links file and packages are cached with the user-specific directory (build-path (find-system-path 'addon-dir) "collects") for all-version cases, and in (build-path (find-system-path 'addon-dir) (version) "collects") for version-specific cases. These cache paths can be redirected by an 'info-domain-root entry in "config.rktd" (see §19 "Installation Configuration and Search Paths").

```
(find-relevant-directory-records syms key)
 → (listof directory-record?)
  syms : (listof symbol?)
  key : (or/c 'preferred 'all-available 'no-planet 'no-user)
```

Like find-relevant-directories, but returns directory-record structs instead of path?s.

```
(struct directory-record (maj min spec path syms)
    #:extra-constructor-name make-directory-record)
  maj : integer?
  min : integer?
  spec : any/c
  path : path?
  syms : (listof symbol?)
```

A struct that records information about a collection or a PLaneT package that has been installed. Collections will have the major version being 1 and the minor version being 0. The spec field is a quoted module spec; the path field is where the info.rkt file for this collection or PLaneT package exists on the filesystem; the syms field holds the identifiers defined in that file.

```
(reset-relevant-directories-state!) → void?
```

Resets the cache used by find-relevant-directories.

6.10 API for Relative Paths

The Racket installation tree can usually be moved around the filesystem. To support this, care must be taken to avoid absolute paths. The following two APIs cover two aspects of this: a way to convert a path to a value that is relative to the "collects" tree, and a way to display such paths (e.g., in error messages).

6.10.1 Representing Collection-Based Paths

```
(require setup/collects)
```
package: base

```
(path->collects-relative path
                         [#:cache cache])
 → (or/c path-string?
         (cons/c 'collects
                 (cons/c bytes? (non-empty-listof bytes?))))
  path : path-string?
  cache : (or/c #f (and/c hash? (not/c immutable?))) = #f
```

Checks whether *path* (normalized by `path->complete-path` and `simplify-path` with `#f` as its second argument) matches the result of `collection-file-path`. If so, the result is a list starting with `'collects` and containing the relevant path elements as byte strings. If not, the path is returned as-is.

The *cache* argument is used with `path->pkg`, if needed.

```
(collects-relative->path rel) → path-string?
  rel : (or/c path-string?
              (cons/c 'collects
                      (cons/c bytes? (non-empty-listof bytes?))))
```

The inverse of `path->collects-relative`: if *rel* is a pair that starts with `'collects`, then it is converted back to a path using `collection-file-path`.

```
(path->module-path path [#:cache cache])
 → (or/c path-string? normalized-lib-module-path?)
  path : path-string?
  cache : (or/c #f (and/c hash? (not/c immutable?))) = #f
```

Like `path->collects-relative`, but the result is either *path* or a normalized (in the sense of `collapse-module-path`) module path.

6.10.2 Representing Paths Relative to `"collects"`

```
(require setup/main-collects)
```
package: base

```
(path->main-collects-relative path)
 → (or/c path? (cons/c 'collects (non-empty-listof bytes?)))
  path : (or/c bytes? path-string?)
```

Checks whether *path* has a prefix that matches the prefix to the main "collects" directory as determined by (find-collects-dir). If so, the result is a list starting with 'collects and containing the remaining path elements as byte strings. If not, the path is returned as-is.

The *path* argument should be a complete path. Applying simplify-path before path->main-collects-relative is usually a good idea.

For historical reasons, *path* can be a byte string, which is converted to a path using bytes->path.

See also collects-relative->path.

```
(main-collects-relative->path rel) → path?
  rel : (or/c bytes?
              path-string?
              (cons/c 'collects (non-empty-listof bytes?)))
```

The inverse of path->main-collects-relative: if *rel* is a pair that starts with 'collects, then it is converted back to a path relative to (find-collects-dir).

6.10.3 Representing Paths Relative to the Documentation

```
(require setup/main-doc)
```
package: base

```
(path->main-doc-relative path)
 → (or/c path? (cons/c 'doc (non-empty-listof bytes?)))
  path : (or/c bytes? path-string?)
```

Like path->main-collects-relative, except that it checks for a prefix relative to (find-doc-dir) and returns a list starting with 'doc if so.

```
(main-doc-relative->path rel) → path>
  rel : (or/c bytes?
              path-string?
              (cons/c 'doc (non-empty-listof bytes?)))
```

Like path->main-collects-relative, except it is the inverse of path->main-doc-relative.

6.10.4 Displaying Paths Relative to a Common Root

```
(require setup/path-to-relative)
```
package: base

```
(path->relative-string/library  path
                                [default
                                 #:cache cache]) → any/c
  path : path-string?
  default : (or/c (-> path-string? any/c) any/c)
          = (lambda (x) (if (path? x) (path->string x) x))
  cache : (or/c #f (and/c hash? (not/c immutable?))) = #f
```

Produces a string suitable for display in error messages. If the path is an absolute one that is inside a package, the result is a string that begins with `"<pkgs>/"`. If the path is an absolute one that is inside the `"collects"` tree, the result is a string that begins with `"<collects>/"`. Similarly, a path in the user-specific collects results in a prefix of `"<user-collects>/"`, a PLaneT path results in `"<planet>/"`, and a path into documentation results in `"<doc>/"` or `"<user-doc>/"`.

If *cache* is not `#f`, it is used as a cache argument for `path->pkg` to speed up detection and conversion of package paths.

If the path is not absolute, or if it is not in any of these, it is returned as-is (converted to a string if needed). If *default* is given, it specifies the return value instead: it can be a procedure that is applied onto the path to get the result, or the result itself.

Note that this function can return a non-string only if *default* is given and it does not return a string.

```
(path->relative-string/setup  path
                              [default
                               #:cache cache]) → any/c
  path : path-string?
  default : (or/c (-> path-string? any/c) any/c)
          = (lambda (x) (if (path? x) (path->string x) x))
  cache : (or/c #f (and/c hash? (not/c immutable?))) = #f
```

The same as `path->relative-string/library`, for backward compatibility.

```
(make-path->relative-string dirs [default])
 → (path-string? any/c . -> . any)
  dirs : (listof (cons (-> path?) string?))
  default : (or/c (-> path-string? any/c) any/c)
          = (lambda (x) (if (path? x) (path->string x) x))
```

This function produces functions like `path->relative-string/library` and `path->relative-string/setup`.

The *dirs* argument determines the prefix substitutions. It must be an association list mapping a path-producing thunk to a prefix string for paths in the specified path.

default determines the default for the resulting function (which can always be overridden by an additional argument to this function).

6.11 API for Collection Names

```
(require setup/collection-name)
```
package: base

```
(collection-name? v) → boolean?
  v : any/c
```

Returns #t if *v* is a string that is syntactically valid as a collection name, which means that it is one or more /-separated strings for which collection-name-element? returns true.

```
(collection-name-element? v) → boolean?
  v : any/c
```

Returns #t if *v* is a string that is syntactically valid as a top-level collection name or as a part of a collection name, which means that it is non-empty and contains only ASCII letters, ASCII digits, -, +, _, and %, where a % is allowed only when followed by two lowercase hexadecimal digits, and the digits must form a number that is not the ASCII value of a letter, digit, -, +, or _.

6.12 API for Collection Searches

```
(require setup/collection-search)
```
package: base

Added in version 6.3 of package base.

```
(collection-search mod-path
                   [#:init result
                    #:combine combine
                    #:break? break?
                    #:all-possible-roots? all-possible-roots?])
 → any/c
  mod-path : normalized-lib-module-path?
  result : any/c = #f
  combine : (any/c (and/c path? complete-path?) . -> . any/c)
          = (lambda (r v) v)
  break? : (any/c . -> . any/c) = (lambda (r) #f)
  all-possible-roots? : any/c = #f
```

Generalizes collection-file-path to support folding over all possible locations of a collection-based file in the current configuration. Unlike collection-file-path,

`collection-search` takes the file to location in module-path form, but always as a `'lib` path.

Each possible path for the file (not counting a `".ss"` to/from `".rkt"` conversion) is provided as a second argument to the *combine* function, where the first argument is the current result, and the value produced by *combine* becomes the new result. The `#:init` argument provides the initial result.

The *break?* function short-circuits a search based on the current value. For example, it could be used to short-circuit a search after a suitable path is found.

If *all-possible-roots?* is `#f`, then *combine* is called only on paths within `"collects"`-like directories (for the current configuration) where at least a matching collection directory exists.

```
(normalized-lib-module-path? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a module path (in the sense of `module-path?`) of the form `'(lib str)` where *str* contains at least one slash. The `collapse-module-path` function produces such module paths for collection-based module references.

6.13 API for Platform Specifications

```
(require setup/matching-platform)
```
package: `base`

Added in version 6.0.1.13 of package `base`.

```
(platform-spec? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a symbol, string, or regexp value (in the sense of `regexp?`), `#f` otherwise.

```
(matching-platform? spec
                    [#:cross? cross?
                     #:system-type sys-type
                     #:system-library-subpath sys-lib-subpath])
 → boolean?
  spec : platform-spec?
  cross? : any/c = #f
  sys-type : (or/c #f symbol?) = (if cross?
                                     (cross-system-type)
                                     (system-type))
  sys-lib-subpath : (or/c #f path-for-some-system?)
                  = (if cross?
                        (cross-system-library-subpath #f)
                        (system-library-subpath #f))
```

Reports whether *spec* matches *sys-type* or *sys-lib-subpath*, where #f values for the latter are replaced with the default values.

If *spec* is a symbol, then the result is #t if *sys-type* is the same symbol, #f otherwise.

If *spec* is a string, then the result is #t if (path->string *sys-lib-subpath*) is the same string, #f otherwise.

If *spec* is a regexp value, then the result is #t if the regexp matches (path->string *sys-lib-subpath*), #f otherwise.

Changed in version 6.3 of package base: Added #:cross? argument and changed the contract on *sys-lib-subpath* to accept path-for-some-system? instead of just path?.

6.14 API for Cross-Platform Configuration

See the raco cross documentation for information about raco cross, a tool that is provided by the "raco-cross" package as a convenient interface to cross-compilation for Racket. The underlying API documented here supports raco cross and other tools.

(require setup/cross-system) package: base

The setup/cross-system library provides functions for querying the system properties of a destination platform, which can be different than the current platform in cross-installation modes.

A Racket installation includes a "system.rktd" file in the directory reported by (find-lib-dir). When the information in that file does not match the running Racket's information, then the setup/cross-system module infers that Racket is being run in cross-installation mode.

For example, if an in-place Racket BC installation for a different platform resides at ⟨*cross-dir*⟩, then running Racket BC as

```
racket -C -G ⟨cross-dir⟩/etc -X ⟨cross-dir⟩/collects -l- raco pkg
```

runs raco pkg using the current platform's racket executable, but using the collections and other configuration information of ⟨*cross-dir*⟩, as well as modifying the packages of ⟨*cross-dir*⟩. That can work as long as no platform-specific libraries need to run to perform the requested raco pkg action (e.g., when installing built packages), or as long as the current platform's installation already includes those libraries.

For Racket CS, cross compilation is more complicated, because Racket CS ".zo" files are platform-specific:

- A target installation ⟨*cross-dir*⟩ is needed that includes cross-compilation support for

the host platform as plug-in within the installation's "⟨*cross-dir*⟩/lib" directory. That installation might be created by compiling from source on the host platform. Only Racket CS can use a CS cross-compilation plug-in.

When running `racket` in cross mode, use the `--cross-compiler` flag to specify the target machine and path to the "⟨*cross-dir*⟩/lib" directory.

- A flag combination `-MCR` with argument "⟨*absolute-zo-dir*⟩:" is needed to enable ".zo" file creation for both the host platform (which uses the directory before a `:`) and the target platform (which uses the normal compiled-file subdirectory when the path after the `:` is empty).

  The ⟨*absolute-zo-dir*⟩ can be any absolute path. It generally should be populated by running `raco setup` in cross mode before commands like `raco pkg`.

For example, the `raco pkg` example for Racket CS is

```
racket --cross-compiler ⟨target-machine⟩ ⟨cross-dir⟩/lib \
  -MCR ⟨absolute-zo-dir⟩: \
  -G ⟨cross-dir⟩/etc -X ⟨cross-dir⟩/collects -l- raco pkg
```

The ⟨*target-machine*⟩ provided to `--cross-compiler` should be the same as the `target-machine` entry in "⟨*cross-dir*⟩/lib/systemd.rktd".

The `-C` flag is shorthand for `--cross`, `-M` is short for `--compile-any`, `-R` is short for `--compiled`, `-G` is short for `--config`, `-X` is short for `--collects`, and `-MCR` is short for `-M -C -R`.

Added in version 6.3 of package `base`.

```
(cross-system-type [mode])
 → (or/c symbol? string? bytes? exact-positive-integer? vector?)
  mode : (or/c 'os 'os* 'arch 'word 'so-find 'platform = 'os
               'gc 'vm 'link 'machine 'target-machine
               'so-suffix 'so-mode 'fs-change)
```

Like `system-type`, but for the target platform instead of the current platform in cross-installation mode. When not in cross-installation mode, the results are the same as for `system-type`.

See also `'cross` mode for `system-type`.

```
(cross-system-library-subpath [mode]) → path-for-some-system?
  mode : (or/c 'cgc '3m 'cs #f) = (system-type 'gc)
```

Like `system-library-subpath`, but for the target platform instead of the current platform in cross-installation mode. When not in cross-installation mode, the results are the same as for `system-library-subpath`.

In cross-installation mode, the target platform may have a different path convention than the current platform, so the result is `path-for-some-system?` instead of `path?`.

```
(cross-installation?) → boolean?
```

Returns `#t` if cross-installation mode has been detected, `#f` otherwise.

6.15 API for Cross-References for Installed Manuals

```
(require setup/xref)      package: racket-index
```

```
(load-collections-xref [on-load]) → xref?
  on-load : (-> any/c) = (lambda () (void))
```

Either creates and caches or returns a cached cross-reference record created with `make-collections-xref`. The *on-load* function is called only when a previously cached record is not returned.

```
(make-collections-xref
 [#:no-user? no-user?
  #:no-main? no-main?
  #:doc-db db-path
  #:quiet-fail? quiet-fail?
  #:register-shutdown! register-shutdown!])
→ xref?
 no-user? : any/c = #f
 no-main? : any/c = #f
 db-path : (or/c #f path?) = #f
 quiet-fail? : any/c = #f
 register-shutdown! : ((-> any) . -> . any) = void
```

Like `load-xref`, but automatically finds all cross-reference files for manuals that have been installed with `raco setup`. The resulting cross-reference record takes advantage of a cross-reference database *db-path*, when support is available, to delay the loading of cross-reference details until needed.

Cross-reference information is skipped when it is installed in the main installation or in a user-specific location, respectively, if *no-main?* or *no-user?* is `#t`.

If *quiet-fail?* is true, then errors are suppressed while loading cross-reference information.

The *register-shutdown!* callback may be called to register a function that closes database connections when the result of `make-collections-xref` is no longer needed. If *register-shutdown!* is not supplied or if a function sent to *register-shutdown!* is never called, database connections will be closed only though a custodian.

```
(get-rendered-doc-directories  no-user?
                               no-main?
                              [#:keep-omit? keep-omit?])
 → (listof path?)
  no-user? : any/c
  no-main? : any/c
  keep-omit? : any/c = #f
```

Returns a list of directories for all documentation for all installed collections, omitting documentation that is installed in the main installation or in a user-specific location, respectively, if *no-main?* or *no-user?* is `#t`.

If *keep-omit?* is true, then the result includes documentation that has the `'omit` category.

Changed in version 1.5 of package `racket-index`: Added the `#:keep-omit?` argument.

```
(get-current-doc-state) → doc-state?
```

Records the time stamps of files that are touched whenever the documentation is changed.

Added in version 1.2 of package `racket-index`.

```
(doc-state-changed? doc-state) → boolean?
  doc-state : doc-state?
```

Returns `#t` when the time stamps of the files in *doc-state* changed (or new files appeared) and `#f` otherwise.

If the result is `#t`, then the documentation in this installation of Racket has changed and otherwise it hasn't.

Added in version 1.2 of package `racket-index`.

```
(doc-state? v) → boolean?
  v : any/c
```

A predicate to recognize the result of `get-current-doc-state`.

Added in version 1.2 of package `racket-index`.

6.16 API for Materializing User-Specific Documentation

```
(require setup/materialize-user-docs)
```

package: racket-index

Added in version 1.1 of package racket-index.

```
(materialize-user-docs
 [on-setup
  #:skip-user-doc-check? skip-user-doc-check?])
→ void?
 on-setup : ((-> boolean?) . -> . any)
          = (lambda (setup) (setup))
 skip-user-doc-check? : any/c = #f
```

Checks whether a user-specific documentation entry point already exists in (find-user-doc-dir), and if not, runs raco setup in a mode that will create the entry point (to have the same content as the installation's documentation entry point.) If *skip-user-doc-check?* is not #f, then skips the check for the user-specific documentation entry point.

The run of raco setup is packaged in a thunk that is provided to *on-setup*, which can adjust the current output and error ports as appropriate and check the thunk's result for success.

The *on-setup* argument is not called if the documentation entry point already exists in (find-user-doc-dir).

Changed in version 1.1 of package racket-index: Added the *skip-user-doc-check?* argument.

6.17 API for Staging Documentation Installs

```
(require setup/doc-to-destdir)
```
package: racket-index

The setup/doc-to-destdir module provides support for staging rendered documentation to better match the filesystem layout of a package installation.

When the setup/doc-to-destdir module is run directly, such as with racket -l, it expects command-line arguments to pass along to move-rendered-docs-to-destdir. Provide the --help flag for more information.

Added in version 9.2.0.6 of package racket-index.

```
(move-rendered-docs-to-destdir dir
                               doc-destdir) → void?
 dir : path-string?
 doc-destdir : path-string?
```

Scans *dir* for collections (typically within package implementations) that specify documentation through a "info.rkt" module, and where the specified documentation has been ren-

dered to a local "doc" subdirectory. The rendered documentation is moved out of "doc" and into *doc-destdir*. This movement imitates a raco setup step of moving documentation that is already rendered within a collection into a centralized directory for installation-scoped packages.

The move-rendered-docs-to-destdir function is intended for use in combination with raco pkg install --destdir, where *dir* is the same directory as provided after --destdir, and *doc-destdir* is an accompanying directory to be moved into place as rendered documentation instead of package implementation.

6.18 Layered Installations

A typical Racket configuration includes two layers: an *installation* layer and a *user* layer. The intent is that the *installation* layer is read-only to all users of a system, while the *user* layer allows each individual user to install additional packages that extend the *installation* layer. The *installation* layer is intended not only to be read-only, but to not change after users start installing in their own layers.

In an environment where Racket itself is under development, the *installation* layer will change. In that setting, if the *user* layer is used at all, care must be taken to not create conflicts for the user layer when modifying the installation layer—or else the user layer must be repaired on occasion.

By default, raco setup updates both layers whenever it is run; if a user does not have write permission the installation, raco setup with no arguments is all but certain to report permission errors. The actions of raco setup can be constrained to the *user* layer by supplying the --avoid-main argument, or raco setup can be constrained to the *installation* layer by using the --no-user argument. When raco pkg performs setup actions, it effectively supplies one of the other of those based on the package's scope (and raco pkg refuses to operate on both scopes/layers at once).

The *user* layer is always both user- and version-specific. More precisely, it is specific to the user and an installation's name, where the installation's name is typically its version number. However, the name of an installation can be changed through the 'installation setting in "config.rktd" (see §19 "Installation Configuration and Search Paths"). Setting an installation name changes the directory where packages and executables reside within (find-system-path 'addon-dir). The result of (find-system-path 'addon-dir) itself can be changed through 'addon-dir in "config.rktd".

The *installation* and *user* configuration layers can be generalized to multiple layers by setting search paths in "config.rktd". These search paths essentially treat the layer closest to *user* as the *installation* layer that might be adjusted by raco setup and raco pkg, but search paths can chain to an existing (unchanging) implementation in much the same way that *user* chains to *installation*. To build a new layer, create new "config.rktd" that is like the underlying layer's "config.rktd", but

- each of `'lib-dir`, `'share-dir`, `'links-file`, `'pkgs-dir`, `'bin-dir`, `'gui-bin-dir`, `'apps-dir`, `'doc-dir`, and `'man-dir` is a new directory or file; and
- the corresponding search lists `'lib-search-dirs`, `'share-search-dirs`, `'links-search-files`, `'pkgs-search-dirs`, `'bin-search-dirs`, `'gui-bin-search-dirs`, (no `'apps-dir` search needed), `'doc-search-dirs`, and `'man-search-dirs` each add the old directory or file to the search list just after `#f`; note that the default for each search list is `(list #f)`.

There is no argument to `raco setup` that is analogous to `--avoid-main` to avoid modifying nested layer; instead, nested layers are expected to be fully set up so that `raco setup` need not change them. When `raco setup` would otherwise install an executable into the directory configured as `'bin-dir`, it consults the `'bin-search-dirs` list to check whether the executable is already installed in one of those directories, and if so, it will refrain from creating a copy in the new layer. The same search-list check also applies to native libraries, shared files, and man pages, but with the additional check that the file to install matches the one that is already installed.

The default path to `"config.rktd"` is hardwired within a `racket` executable. In some cases, it can make sense for the innermost layer's configuration to point to another layer, perhaps because the filesystem provides an indirection. For example, on Unix, a Racket installation in `"/usr"` might reasonably configure the *installation* layer's directories to be in `"/usr/local"` with `"/usr"` directories included in the search lists.

To use `racket` with a new `"config.rktd"`, you can supply the `--config` or `--G` flag to `racket` or set the `PLTCONFIGDIR` environment variable to point to the directory containing `"config.rktd"`. Alternatively, you can create a tethered layer that creates replacement executables like `racket` that are hardwired to the layer's configuration directory.

6.19 Tethered Installations

A *tethered* installation of Racket is a layer (see §6.18 "Layered Installations") that includes a wrapper executable for every executable across the installation's layers. Each wrapper executable points back to the new layer's `"config.rktd"` (see §19 "Installation Configuration and Search Paths") without the use of a `PLTCONFIGDIR` environment variable or `--config` flag. In other words, a tethered installation provides executables such as `racket`, `raco`, and `drracket` that are tied to the layer. Tethering thus helps to create a layer of installation that behaves in a more self-contained way, but with minimal duplication of the underlying layers.

Tethering works at either a *user* or *installation* layer:

- A *user* layer with tethering is represented by a fresh directory ⟨*addon-dir*⟩ and a `"`⟨*addon-dir*⟩`/etc/config.rktd"` file that maps `'addon-tethered-console-bin-dir` to ⟨*tethered-bin-dir*⟩, `'addon-tethered-gui-bin-dir` to ⟨*tethered-gui-*

bin-dir⟩, and (optionally) `'addon-tethered-apps-dir` to ⟨*tethered-apps-dir*⟩. Initialize the tethered layer with

```
racket -A ⟨addon-dir⟩ -l- raco setup --avoid-main
```

- An *installation* layer with tethering is like a one without tethering (see §6.18 “Layered Installations”), but where the layer’s `"`⟨*layer-dir*⟩`/etc/config.rktd"` file maps `'config-tethered-console-bin-dir` to ⟨*tethered-bin-dir*⟩, `'config-tethered-gui-bin-dir` to ⟨*tethered-gui-bin-dir*⟩, and (optionally) `'config-tethered-apps-dir` to ⟨*tethered-apps-dir*⟩. The `'bin-dir` and `'gui-bin-dir` configurations can point to the same directories, but executables are not specifically created there by `raco setup`. Initialize the tethered layer with

  ```
  racket -G ⟨layer-dir⟩/etc -l- raco setup
  ```

In either case, initialization creates tethered executables in the directories ⟨*tethered-bin-dir*⟩ and ⟨*tethered-gui-bin-dir*⟩, writing `".desktop"` files (for Unix) in ⟨*tethered-apps-dir*⟩ (if specified). Thereafter, tethered executables like ⟨*tethered-bin-dir*⟩`/racket` and ⟨*tethered-bin-dir*⟩`/raco` can be used to work with the tethered layer.

7 `raco decompile`: Decompiling Bytecode

The `raco decompile` command takes the path of a bytecode file (which usually has the file extension `".zo"`) or a source file with an associated bytecode file (usually created with `raco make`) and converts the bytecode file's content back to an approximation of Racket code. When the "bytecode" file contains machine code, as for the CS variant of Racket, then it cannot be converted back to an approximation of Racket, but installing the `"disassemble"` package may enable disassembly of the machine code. Decompilation is mostly useful for checking the compiler's transformation and optimization of the source program.

The `raco decompile` command accepts the following command-line flags:

- `--force` — skip modification-date comparison on the given file's path and an associated `".zo"` file (if any)
- `-n` ⟨*n*⟩ or `--columns` ⟨*n*⟩ — format output for a display with ⟨*n*⟩ columns
- `--linklet` — decompile only as far as linklets, instead of decoding linklets to approximate Racket `module` forms
- `--no-disassemble` — show machine code as-is in a byte string, instead of attempting to disassemble
- `--no-syntax` — avoid decompiling syntax-object literals
- `--partial-fasl` — preserve more of the original structure of the bytecode file, instead of focusing on procedure bodies

Changed in version 1.8: Added `--no-disassemble`.
Changed in version 1.9: Added `--partial-fasl`.
Changed in version 1.17: Added `--no-syntax`.

7.1 Racket CS Decompilation

Decompilation of Racket CS bytecode mostly shows the structure of a module around machine-code implementations of procedures.

- A `#%machine-code` form corresponds to machine code that is not disassembled, where the machine code is in a byte string.
- A `#%assembly-code` form corresponds to disassembled machine code, where the assembly code is shown as a sequence of strings.
- A `#%interpret` form corresponds to a compiled form of a large procedure, where only smaller nested procedures are compiled to machine code.

7.2 Racket BC Decompilation

Racket BC bytecode has a structure that is close enough to Racket's core language that it can more often be converted to an approximation of Racket code. To the degree that it can be converted back, many forms in the decompiled code have the same meanings as always, such as `module`, `define`, and `lambda`. Other forms and transformations are specific to the rendering of bytecode, and they reflect a specific execution model:

- Top-level variables, variables defined within the module, and variables imported from other modules are prefixed with `_`, which helps expose the difference between uses of local variables versus other variables. Variables imported from other modules, moreover, have a suffix starting with `@` that indicates the source module. Finally, imported variables with constantness have a midfix: `:c` to indicate constant shape across all instantiations, `:f` to indicate a fixed value after initialization, `:p` to indicate a procedure, `:P` to indicate a procedure that preserves continuation marks on return, `:t` to indicate a structure type, `:mk` to indicate a structure constructor, `:?` to indicate a structure predicate, `:ref` to indicate a structure accessor, or `:set!` to indicate a structure mutator.

  Non-local variables are always accessed indirectly though an implicit `#%globals` or `#%modvars` variable that resides on the value stack (which otherwise contains local variables). Variable accesses are further wrapped with `#%checked` when the compiler cannot prove that the variable will be defined before the access.

  Uses of core primitives are shown without a leading `_`, and they are never wrapped with `#%checked`.

- Local-variable access may be wrapped with `#%sfs-clear`, which indicates that the variable-stack location holding the variable will be cleared to prevent the variable's value from being retained by the garbage collector. Variables whose name starts with `unused` are never actually stored on the stack, and so they never have `#%sfs-clear` annotations. (The bytecode compiler normally eliminates such bindings, but sometimes it cannot, either because it cannot prove that the right-hand side produces the right number of values, or the discovery that the variable is unused happens too late with the compiler.)

  Mutable variables are converted to explicitly boxed values using `#%box`, `#%unbox`, and `#%set-boxes!` (which works on multiple boxes at once). A `set!-rec-values` operation constructs mutually-recursive closures and simultaneously updates the corresponding variable-stack locations that bind the closures. A `set!`, `set!-values`, or `set!-rec-values` form is always used on a local variable before it is captured by a closure; that ordering reflects how closures capture values in variable-stack locations, as opposed to stack locations.

- In a `lambda` form, if the procedure produced by the `lambda` has a name (accessible via `object-name`) and/or source-location information, then it is shown as a quoted constant at the start of the procedure's body. Afterward, if the `lambda` form captures any bindings from its context, those bindings are also shown in a quoted constant.

Neither constant corresponds to a computation when the closure is called, though the list of captured bindings corresponds to a closure allocation when the `lambda` form itself is evaluated.

A `lambda` form that closes over no bindings is wrapped with `#%closed` plus an identifier that is bound to the closure. The binding's scope covers the entire decompiled output, and it may be referenced directly in other parts of the program; the binding corresponds to a constant closure value that is shared, and it may even contain cyclic references to itself or other constant closures.

- A form `(#%apply-values proc expr)` is equivalent to `(call-with-values (lambda () expr) proc)`, but the run-time system avoids allocating a closure for *expr*. Similarly, a `#%call-with-immediate-continuation-mark` call is equivalent to a `call-with-immediate-continuation-mark` call, but avoiding a closure allocation.

- A `define-values` form may have `(begin '%%inline-variant%% expr1 expr2)` for its expression, in which case *expr2* is the normal result, but *expr1* may be inlined for calls to the definition from other modules. Definitions of functions without an `'%%inline-variant%%` are never inlined across modules.

- Function arguments and local bindings that are known to have a particular type have names that embed the known type. For example, an argument might have a name that starts `argflonum` or a local binding might have a name that starts `flonum` to indicate a flonum value.

7.3 API for Decompiling

```
(require compiler/decompile)      package: compiler-lib
```

```
(decompile top
           [#:to-linklets? to-linklets?
            #:skip-syntax-literals? skip-syntax-literals?])
 → any/c
  top : (or/c linkl-directory? linkl-bundle? linkl?
              linklet? faslable-correlated-linklet?)
  to-linklets? : any/c = #f
  skip-syntax-literals? : any/c = #f
```

Consumes the result of parsing bytecode and returns an S-expression (as described above) that represents the compiled code.

If *to-linklets?* is true, then the result S-expression shows raw `linklet` forms within *top* instead of reconstructing a `module` form.

If *skip-syntax-literals?* is true, then the result S-expression omits the decompilation of syntax-object literals.

Changed in version 1.17 of package compiler-lib: Added the #:skip-syntax-literals? argument.

7.4 API for Parsing Bytecode

(require compiler/zo-parse) package: zo-lib

The compiler/zo-parse module re-exports compiler/zo-structs in addition to zo-parse.

```
(zo-parse [in]) → (or/c linkl-directory? linkl-bundle?)
  in : input-port? = (current-input-port)
```

Parses a port (typically the result of opening a ".zo" file) containing bytecode. Beware that the structure types used to represent the bytecode are subject to frequent changes across Racket versions.

The parsed bytecode is returned in a linkl-directory or linkl-bundle structure—the latter only for the compilation of a module that contains no submodules.

Beyond the linklet bundle or directory structure, the result of zo-parse contains linklets that depend on the machine for which the bytecode is compiled:

- For a machine-independent bytecode file, a linklet is represented as a faslable-correlated-linklet.
- For a Racket CS bytecode file, a linklet is opaque, because it is primarily machine code, but decompile can extract some information and potentially disassemble the machine code.
- For Racket BC bytecode, the bytecode can be parsed into structures as described below.

The rest of this section is specific to BC bytecode.

Within a linklet, the bytecode representation of an expression is closer to an S-expression than a traditional, flat control string. For example, an if form is represented by a branch structure that has three fields: a test expression, a "then" expression, and an "else" expression. Similarly, a function call is represented by an application structure that has a list of argument expressions.

Storage for local variables or intermediate values (such as the arguments for a function call) is explicitly specified in terms of a stack. For example, execution of an application structure reserves space on the stack for each argument result. Similarly, when a let-one structure (for a simple let) is executed, the value obtained by evaluating the right-hand side expression is pushed onto the stack, and then the body is evaluated. Local variables are

always accessed as offsets from the current stack position. When a function is called, its arguments are passed on the stack. A closure is created by transferring values from the stack to a flat closure record, and when a closure is applied, the saved values are restored on the stack (though possibly in a different order and likely in a more compact layout than when they were captured).

When a sub-expression produces a value, then the stack pointer is restored to its location from before evaluating the sub-expression. For example, evaluating the right-hand size for a `let-one` structure may temporarily push values onto the stack, but the stack is restored to its pre-`let-one` position before pushing the resulting value and continuing with the body. In addition, a tail call resets the stack pointer to the position that follows the enclosing function's arguments, and then the tail call continues by pushing onto the stack the arguments for the tail-called function.

Values for global and module-level variables are not put directly on the stack, but instead stored in "buckets," and an array of accessible buckets is kept on the stack. When a closure body needs to access a global variable, the closure captures and later restores the bucket array in the same way that it captured and restores a local variable. Mutable local variables are boxed similarly to global variables, but individual boxes are referenced from the stack and closures.

7.5 API for Marshaling Bytecode

```
(require compiler/zo-marshal)
```
package: zo-lib

```
(zo-marshal-to top out) → void?
  top : (or/c linkl-directory? linkl-bundle?)
  out : output-port?
```

Consumes a representation of bytecode and writes it to *out*.

```
(zo-marshal top) → bytes?
  top : (or/c linkl-directory? linkl-bundle?)
```

Consumes a representation of bytecode and generates a byte string for the marshaled bytecode.

7.6 Bytecode Representation

```
(require compiler/zo-structs)
```
package: zo-lib

The `compiler/zo-structs` library defines the bytecode structures that are produced by `zo-parse` and consumed by `decompile` and `zo-marshal`.

Warning: The `compiler/zo-structs` library exposes internals of the Racket bytecode abstraction. Unlike other Racket libraries, `compiler/zo-structs` is subject to incompatible changes across Racket versions.

```
(struct zo ()
    #:extra-constructor-name make-zo
    #:prefab)
```

A supertype for all forms that can appear in compiled code.

7.6.1 Prefix

```
(struct linkl-directory zo (table)
    #:extra-constructor-name make-linkl-directory
    #:prefab)
  table : (hash/c (listof symbol?) linkl-bundle?)
(struct linkl-bundle zo (table)
    #:extra-constructor-name make-linkl-bundle
    #:prefab)
  table : (hash/c (or/c symbol? fixnum?) (or linkl? any/c))
```

Wraps compiled code.

Module and top-level compilation produce one or more linklets that represent independent evaluation in a specific phase. Even a single top-level expression or a module with only run-time code will generate multiple linklets to implement metadata and syntax data. A module with no submodules is represented directly by a `linkl-bundle`, while any other compiled form is represented by a `linkl-directory`.

A linklet bundle maps an integer to a linklet representing forms to evaluate at the integer-indicated phase. Symbols are mapped to metadata, such as a module's name as compiled or a linklet implementing literal syntax objects. A linklet directory normally maps `'()` to the main linklet bundle for a module or a single top-level form; for a linklet directory that corresponds to a sequence of top-level forms, however, there is no "main" linklet bundle, and symbol forms of integers are used to order the linkets.

For a module with submodules, the linklet directory maps submodule paths (as lists of symbols) to linklet bundles for the corresponding submodules.

An individual linklet is represented as a `linkl` only if the source bytecode file was for Racket BC. A CS bytecode linklet will be represented by an opaque linklet (in the sense of `linklet?` from `racket/linklet`). A machine-independent linklet is represented as a `faslable-correlated-linklet` structure.

```
(struct linkl zo (name
                  importss
                  import-shapess
                  exports
                  internals
                  lifts
                  source-names
                  body
                  max-let-depth
                  need-instance-access?)
   #:extra-constructor-name make-linkl
   #:prefab)
  name : symbol?
  importss : (listof (listof symbol?))
  import-shapess : (listof (listof  (or/c #f 'constant 'fixed
                                          function-shape?
                                          struct-shape?)))
  exports : (listof symbol?)
  internals : (listof (or/c symbol? #f))
  lifts : (listof symbol?)
  source-names : (hash/c symbol? symbol?)
  body : (listof (or/c form? any/c))
  max-let-depth : exact-nonnegative-integer?
  need-instance-access? : boolean?
```

Represents a linklet, which corresponds to a module body or a top-level sequence at a single phase.

The `name` of a linklet is for debugging purposes, similar to the inferred name of a `lambda` form.

The `importss` list of lists describes the linklet's imports. Each of the elements of the out list corresponds to an import source, and each element of an inner list is the symbolic name of an export from that source. The `import-shapess` list is in parallel to `imports`; it reflects optimization assumptions by the compiler that are used by the bytecode validator and checked when the linklet is instantiated.

The `exports` list describes the linklet's defined names that are exported. The `internals` list describes additional definitions within the linket, but they are not accessible from the outside of a linklet or one of its instances; a `#f` can appear in place of an unreferenced internal definition that has been removed. The `lifts` list is an extension of `internals` for procedures that are lifted by the compiler; these procedures have certain properties that can be checked by the bytecode validator.

Each symbol in `exports`, `internals`, and `lifts` must be distinct from any other symbol

in those lists. The source-names table maps symbols in exports, internals, and lifts to other symbols, potentially not distinct, that correspond to original source names for the definition. The source-names table is used only for debugging.

When a linklet is instantiated, variables corresponding to the flattening of the lists importss, exports, internals, and lifts are placed in an array (in that order) for access via toplevel references. The initial slot is reserved for a variable-like reference that strongly retains a connection to an instance of its enclosing linklet.

The bodys list is the executable content of the linklet. The value of the last element in bodys may be returned when the linklet is instantiated, depending on the way that it's instantiated.

The max-let-depth field indicates the maximum size of the stack that will be created by any body.

The need-instance-access? boolean indicates whether the linklet contains a toplevel for position 0. A #t is allowed (but suboptimal) if not such reference is present in the linklet body.

```
(struct function-shape (arity preserves-marks?)
    #:extra-constructor-name make-function-shape
    #:prefab)
  arity : procedure-arity?
  preserves-marks? : boolean?
```

Represents the shape of an expected import, which should be a function having the arity specified by arity. The preserves-marks? field is true if calling the function is expected to leave continuation marks unchanged by the time it returns.

```
(struct struct-shape ()
    #:extra-constructor-name make-struct-shape
    #:prefab)
(struct struct-type-shape struct-shape (field-count authentic?)
    #:extra-constructor-name make-struct-type-shape
    #:prefab)
  field-count : exact-nonnegative-integer?
  authentic? : boolean?
(struct constructor-shape struct-shape (arity)
    #:extra-constructor-name make-constructor-shape
    #:prefab)
  arity : exact-nonnegative-integer?
(struct predicate-shape struct-shape (authentic?)
    #:extra-constructor-name make-predicate-shape
    #:prefab)
  authentic? : boolean?
```

```
(struct accessor-shape struct-shape (field-count authentic?)
    #:extra-constructor-name make-accessor-shape
    #:prefab)
  field-count : exact-nonnegative-integer?
  authentic? : boolean?
(struct mutator-shape struct-shape (field-count authentic?)
    #:extra-constructor-name make-mutator-shape
    #:prefab)
  field-count : exact-nonnegative-integer?
  authentic? : boolean?
(struct struct-type-property-shape struct-shape (has-guard?)
    #:extra-constructor-name make-struct-type-property-shape
    #:prefab)
  has-guard? : boolean?
(struct property-predicate-shape struct-shape ()
    #:extra-constructor-name make-property-predicate-shape
    #:prefab)
(struct property-accessor-shape struct-shape ()
    #:extra-constructor-name make-property-accessor-shape
    #:prefab)
(struct struct-other-shape struct-shape ()
    #:extra-constructor-name make-struct-other-shape
    #:prefab)
```

Represents the shape of an expected import as a structure-type binding, constructor, etc.

7.6.2 Forms and Inline Variants

```
(struct form zo ()
    #:extra-constructor-name make-form
    #:prefab)
```

A supertype for all forms that can appear in a linklet body (including `expr`s), except for literals that are represented as themselves.

```
(struct def-values form (ids rhs)
    #:extra-constructor-name make-def-values
    #:prefab)
  ids : (listof toplevel?)
  rhs : (or/c expr? seq? inline-variant? any/c)
```

Represents a `define-values` form. Each element of `ids` references a defined variable in the enclosing linklet.

After `rhs` is evaluated, the stack is restored to its depth from before evaluating `rhs`.

```
(struct inline-variant zo (direct inline)
    #:extra-constructor-name make-inline-variant
    #:prefab)
  direct : expr?
  inline : expr?
```

Represents a function that is bound by `define-values`, where the function has two variants. The first variant is used for normal calls to the function. The second may be used for cross-module inlining of the function.

7.6.3 Expressions

```
(struct expr form ()
    #:extra-constructor-name make-expr
    #:prefab)
```

A supertype for all expression forms that can appear in compiled code, except for literals that are represented as themselves.

```
(struct lam expr (name
                  flags
                  num-params
                  param-types
                  rest?
                  closure-map
                  closure-types
                  toplevel-map
                  max-let-depth
                  body)
    #:extra-constructor-name make-lam
    #:prefab)
  name : (or/c symbol? vector?)
  flags : (listof (or/c 'preserves-marks 'is-method 'single-result
                        'only-rest-arg-not-used 'sfs-clear-rest-args))
  num-params : exact-nonnegative-integer?
  param-types : (listof (or/c 'val 'ref 'flonum 'fixnum 'extflonum))
  rest? : boolean?
  closure-map : (vectorof exact-nonnegative-integer?)
  closure-types : (listof (or/c 'val/ref 'flonum 'fixnum 'extflonum))
  toplevel-map : (or/c #f (set/c exact-nonnegative-integer?))
  max-let-depth : exact-nonnegative-integer?
  body : (or/c expr? seq? any/c)
```

Represents a `lambda` form. The `name` field is a name for debugging purposes. The `num-params` field indicates the number of arguments accepted by the procedure, not counting a rest argument; the `rest?` field indicates whether extra arguments are accepted and collected into a "rest" variable. The `param-types` list contains `num-params` symbols indicating the type of each argument, either `'val` for a normal argument, `'ref` for a boxed argument (representing a mutable local variable), `'flonum` for a flonum argument, or `'extflonum` for an extflonum argument.

The `closure-map` field is a vector of stack positions that are captured when evaluating the `lambda` form to create a closure. The `closure-types` field provides a corresponding list of types, but no distinction is made between normal values and boxed values; also, this information is redundant, since it can be inferred by the bindings referenced though `closure-map`.

When a closure captures top-level or module-level variables or refers to a syntax-object constant, the variables and constants are represented in the closure by capturing a prefix (in the sense of `prefix`). The `toplevel-map` field indicates which top-level variables (i.e., linklet imports and definitions) are actually used by the closure (so that variables in a prefix can be pruned by the run-time system if they become unused) and whether any syntax objects are used (so that the syntax objects as a group can be similarly pruned). A `#f` value indicates either that no prefix is captured or all variables and syntax objects in the prefix should be considered used. Otherwise, numbers in the set indicate which variables and lifted variables are used. Variables are numbered consecutively by position in the prefix starting from `0`, but the number equal to the number of non-lifted variables corresponds to syntax objects (i.e., the number is include if any syntax-object constant is used). Lifted variables are numbered immediately afterward—which means that, if the prefix contains any syntax objects, lifted-variable numbers are shifted down relative to a `toplevel` by the number of syntax object in the prefix (which makes the `toplevel-map` set more compact).

When the function is called, the rest-argument list (if any) is pushed onto the stack, then the normal arguments in reverse order, then the closure-captured values in reverse order. Thus, when `body` is run, the first value on the stack is the first value captured by the `closure-map` array, and so on.

The `max-let-depth` field indicates the maximum stack depth created by `body` plus the arguments and closure-captured values pushed onto the stack. The `body` field is the expression for the closure's body.

Changed in version 6.1.1.8 of package `zo-lib`: Added a number to `toplevel-map` to indicate whether any syntax object is used, shifting numbers for lifted variables up by one if any syntax object is in the prefix.

```
(struct closure expr (code gen-id)
    #:extra-constructor-name make-closure
    #:prefab)
  code : lam?
  gen-id : symbol?
```

A `lambda` form with an empty closure, which is a procedure constant. The procedure constant can appear multiple times in the graph of expressions for bytecode, and the `code` field can be a cycle for a recursive constant procedure; the `gen-id` is different for each such constant.

```
(struct case-lam expr (name clauses)
    #:extra-constructor-name make-case-lam
    #:prefab)
  name : (or/c symbol? vector?)
  clauses : (listof lam?)
```

Represents a `case-lambda` form as a combination of `lambda` forms that are tried (in order) based on the number of arguments given.

```
(struct let-one expr (rhs body type unused?)
    #:extra-constructor-name make-let-one
    #:prefab)
  rhs : (or/c expr? seq? any/c)
  body : (or/c expr? seq? any/c)
  type : (or/c #f 'flonum 'fixnum 'extflonum)
  unused? : boolean?
```

Pushes an uninitialized slot onto the stack, evaluates `rhs` and puts its value into the slot, and then runs `body`. If `type` is not `#f`, then `rhs` must produce a value of the corresponding type, and the slot must be accessed by `localref`s that expect the type. If `unused?` is `#t`, then the slot must not be used, and the value of `rhs` is not actually pushed onto the stack (but `rhs` is constrained to produce a single value).

After `rhs` is evaluated, the stack is restored to its depth from before evaluating `rhs`. Note that the new slot is created before evaluating `rhs`.

```
(struct let-void expr (count boxes? body)
    #:extra-constructor-name make-let-void
    #:prefab)
  count : exact-nonnegative-integer?
  boxes? : boolean?
  body : (or/c expr? seq? any/c)
```

Pushes `count` uninitialized slots onto the stack and then runs `body`. If `boxes?` is `#t`, then the slots are filled with boxes that contain `#<undefined>`.

```
(struct install-value expr (count pos boxes? rhs body)
    #:extra-constructor-name make-install-value
    #:prefab)
  count : exact-nonnegative-integer?
  pos : exact-nonnegative-integer?
```

```
  boxes? : boolean?
  rhs : (or/c expr? seq? any/c)
  body : (or/c expr? seq? any/c)
```

Runs rhs to obtain count results, and installs them into existing slots on the stack in order, skipping the first pos stack positions. If boxes? is #t, then the values are put into existing boxes in the stack slots.

After rhs is evaluated, the stack is restored to its depth from before evaluating rhs.

```
(struct let-rec expr (procs body)
    #:extra-constructor-name make-let-rec
    #:prefab)
  procs : (listof lam?)
  body : (or/c expr? seq? any/c)
```

Represents a letrec form with lambda bindings. It allocates a closure shell for each lambda form in procs, installs each onto the stack in previously allocated slots in reverse order (so that the closure shell for the last element of procs is installed at stack position 0), fills out each shell's closure (where each closure normally references some other just-created closures, which is possible because the shells have been installed on the stack), and then evaluates body.

```
(struct boxenv expr (pos body)
    #:extra-constructor-name make-boxenv
    #:prefab)
  pos : exact-nonnegative-integer?
  body : (or/c expr? seq? any/c)
```

Skips pos elements of the stack, setting the slot afterward to a new box containing the slot's old value, and then runs body. This form appears when a lambda argument is mutated using set! within its body; calling the function initially pushes the value directly on the stack, and this form boxes the value so that it can be mutated later.

```
(struct localref expr (unbox? pos clear? other-clears? type)
    #:extra-constructor-name make-localref
    #:prefab)
  unbox? : boolean?
  pos : exact-nonnegative-integer?
  clear? : boolean?
  other-clears? : boolean?
  type : (or/c #f 'flonum 'fixnum 'extflonum)
```

Represents a local-variable reference; it accesses the value in the stack slot after the first pos slots. If unbox? is #t, the stack slot contains a box, and a value is extracted from the box.

If clear? is #t, then after the value is obtained, the stack slot is cleared (to avoid retaining a reference that can prevent reclamation of the value as garbage). If other-clears? is #t, then some later reference to the same stack slot may clear after reading. If type is not #f, the slot is known to hold a specific type of value.

```
(struct toplevel expr (depth pos const? ready?)
    #:extra-constructor-name make-toplevel
    #:prefab)
  depth : exact-nonnegative-integer?
  pos : exact-nonnegative-integer?
  const? : boolean?
  ready? : boolean?
```

Represents a reference to an imported or defined variable within a linklet. The depth field indicates the number of stack slots to skip to reach the prefix array, and pos is the offset into the array.

When the toplevel is an expression, if both const? and ready? are #t, then the variable definitely will be defined, its value stays constant, and the constant is effectively the same for every module instantiation. If only const? is #t, then the value is constant, but it may vary across instantiations. If only ready? is #t, then the variable definitely will be defined, but its value may change. If const? and ready? are both #f, then a check is needed to determine whether the variable is defined.

When the toplevel is the left-hand side for def-values, then const? is #f. If ready? is #t, the variable is marked as immutable after it is defined.

```
(struct application expr (rator rands)
    #:extra-constructor-name make-application
    #:prefab)
  rator : (or/c expr? seq? any/c)
  rands : (listof (or/c expr? seq? any/c))
```

Represents a function call. The rator field is the expression for the function, and rands are the argument expressions. Before any of the expressions are evaluated, (length rands) uninitialized stack slots are created (to be used as temporary space).

```
(struct branch expr (test then else)
    #:extra-constructor-name make-branch
    #:prefab)
  test : (or/c expr? seq? any/c)
  then : (or/c expr? seq? any/c)
  else : (or/c expr? seq? any/c)
```

Represents an if form.

After test is evaluated, the stack is restored to its depth from before evaluating test.

```
(struct with-cont-mark expr (key val body)
    #:extra-constructor-name make-with-cont-mark
    #:prefab)
  key : (or/c expr? seq? any/c)
  val : (or/c expr? seq? any/c)
  body : (or/c expr? seq? any/c)
```

Represents a `with-continuation-mark` expression.

After each of `key` and `val` is evaluated, the stack is restored to its depth from before evaluating `key` or `val`.

```
(struct seq expr (forms)
    #:extra-constructor-name make-seq
    #:prefab)
  forms : (listof (or/c expr? any/c))
```

Represents a `begin` form.

After each form in `forms` is evaluated, the stack is restored to its depth from before evaluating the form.

```
(struct beg0 expr (seq)
    #:extra-constructor-name make-beg0
    #:prefab)
  seq : (listof (or/c expr? seq? any/c))
```

Represents a `begin0` expression.

After each expression in `seq` is evaluated, the stack is restored to its depth from before evaluating the expression.

Unlike the `begin0` source form, the first expression in `seq` is never in tail position, even if it is the only expression in the list.

```
(struct varref expr (toplevel dummy constant? from-unsafe?)
    #:extra-constructor-name make-varref
    #:prefab)
  toplevel : (or/c toplevel? #t #f symbol?)
  dummy : (or/c toplevel? #f)
  constant? : boolean?
  from-unsafe? : boolean?
```

Represents a `#%variable-reference` form. The `toplevel` field is `#t` if the original reference was to a constant local binding, `#f` if the variable reference is for `(#%variable-reference)` and does not refer to a specific variable, or a symbol if the variable reference

refers to a primitive variable. The `dummy` field accesses a variable bucket that strongly references its namespace (as opposed to a normal variable bucket, which only weakly references its namespace); it can be `#f`.

The value of `constant?` is true when the `toplevel` field is not `#t` but the referenced variable is known to be constant. The value of `from-unsafe?` is true when the module that created the reference was compiled in unsafe mode.

```
(struct assign expr (id rhs undef-ok?)
    #:extra-constructor-name make-assign
    #:prefab)
  id : toplevel?
  rhs : (or/c expr? seq? any/c)
  undef-ok? : boolean?
```

Represents a `set!` expression that assigns to a top-level or module-level variable. (Assignments to local variables are represented by `install-value` expressions.) If `undef-ok?` is true, the assignment to `id` succeeds even if `id` was not previously defined (see also `compile-allow-set!-undefined`).

After `rhs` is evaluated, the stack is restored to its depth from before evaluating `rhs`.

```
(struct apply-values expr (proc args-expr)
    #:extra-constructor-name make-apply-values
    #:prefab)
  proc : (or/c expr? seq? any/c)
  args-expr : (or/c expr? seq? any/c)
```

Represents `(call-with-values (lambda () args-expr) proc)`, which is handled specially by the run-time system.

```
(struct with-immed-mark expr (key def-val body)
    #:extra-constructor-name make-with-immed-mark
    #:prefab)
  key : (or/c expr? seq? any/c)
  def-val : (or/c expr? seq? any/c)
  body : (or/c expr? seq? any/c)
```

Represents a `(call-with-immediate-continuation-mark key (lambda (arg) body) val)` expression that is handled specially by the run-time system to avoid a closure allocation. One initialized slot is pushed onto the stack after `expr` and `val` are evaluated and before `body` is evaluated.

After each of `key` and `val` is evaluated, the stack is restored to its depth from before evaluating `key` or `val`.

```
(struct primval expr (id)
    #:extra-constructor-name make-primval
    #:prefab)
  id : exact-nonnegative-integer?
```

Represents a direct reference to a variable imported from the run-time kernel.

7.7 Machine-Independent Linklets

```
(require compiler/faslable-correlated)
```
package: zo-lib

> **Warning:** The compiler/faslable-correlated library exposes internals of the Racket bytecode abstraction. Unlike other Racket libraries, compiler/faslable-correlated is subject to incompatible changes across Racket versions.

Added in version 1.3 of package zo-lib.

```
(struct faslable-correlated-linklet (expr name)
    #:extra-constructor-name make-faslable-correlated-linklet
    #:prefab)
  expr : any/c
  name : symbol?
```

A faslable-correlated-linklet structure represents a linklet that has been "compiled" to machine-independent form, which just contains an S-expression representing the linklet form. The S-expression is enriched with source-location information by wrapping some nested S-expressions with faslable-correlated structures.

Since faslable-correlated-linklet is a prefab structure type, the contracts documented above for its fields are not enforced.

```
(struct faslable-correlated (e
                             source
                             position
                             line
                             column
                             span
                             props)
    #:extra-constructor-name make-faslable-correlated
    #:prefab)
  e : any/c
  source : any/c
```

```
  position : exact-positive-integer?
  line : exact-positive-integer?
  column : exact-nonnegative-integer?
  span : exact-nonnegative-integer?
  props : (hash/c symbol? any/c)
```

Wraps an S-expression `e` to give it a source location. The S-expression `e` may contain nested `faslable-correlated` structures, but nesting is expected only within pairs.

Since `faslable-correlated` is a prefab structure type, the contracts documented above for its fields are not enforced.

```
(strip-correlated e) → any/c
  e : any/c
```

Recurs through `e` to strip away any `faslable-correlated` structures that are reachable through pairs. The given `e` must not contain any cycles that are reachable through pairs.

8 `raco demod`: Demodularizing Programs

The `raco demodularize` command (usually used with the shorthand `raco demod`) takes a Racket module and flattens its dependencies into a single compiled module, potentially with submodules. A file "⟨*name*⟩`.rkt`" is demodularized into "⟨*name*⟩`_rkt_merged.zo`".

See `compiler/demod` for an alternative way to use the demodularizer. Using `#lang compiler/demod` can cooperate with tools like `raco make` and `raco setup`, which is especially important for library modules (as opposed to end-user programs).

In its default configuration, `raco demod` supports flattening a module that represents an end-user program, so it discards all syntax and compile-time support in the module and its dependencies. Submodules are preserved, but their syntax and compile-time support are similarly discarded. The demodularized "`.zo`" file can be run by passing it as an argument to the `racket` command-line program, or it can be turned into an executable with `raco exe`.

Supply the `-s` or `--syntax` flag to preserve syntax and compile-time components of the module, so that it can be `required` the same as the original module. In that case, modules whose instances need to be shared with other libraries should be omitted from the demodularization using `-x` or `--exclude-library`. For example, `-x racket/base` is normally needed.

A large single module generated by the demodularizer is compiled as if `(#%declare #:unlimited-compile)` is specified, so the value of the `PLT_CS_COMPILE_LIMIT` environment variable does not limit compilation of the module.

The `raco demod` command accepts these flags:

- `-o` ⟨*file*⟩ — writes the flattened module to ⟨*file*⟩ instead of "⟨*name*⟩`_`⟨*ext*⟩`_merged.zo`" for an input file "⟨*name*⟩`.`⟨*ext*⟩".
- `-x` ⟨*module-path*⟩ or `--exclude-library` ⟨*module-path*⟩ — excludes the module in ⟨*module-path*⟩ from flattening, as well as all of its dependencies. An error is reported if ⟨*module-path*⟩ is not a dependency of the input module and has no submodules that are dependencies.
- `-e` ⟨*path*⟩ or `--exclude-module` ⟨*path*⟩ — excludes the module in relative-file ⟨*path*⟩ from flattening, as well as all of its dependencies. An error is reported if ⟨*path*⟩ is not a dependency of the input module and has no submodules that are dependencies. For backward compatibility, `--exclude-modules` is an alias for `--exclude-module`.
- `-s` or `--syntax` — preserve syntax objects and phase levels greater than the run-time phase in the flattened result. Otherwise, only the run-time phase is preserved, and unused (or merely exported) definitions are pruned, since they cannot be referenced through syntax.

- `-M` or `--compile-any` — flattens the module to machine-independent form, instead of recompiling the flattened module to the current platform and Racket virtual machine; the output generated with `-M` loads more slowly than a machine-specific form, but `raco decompile` can show the flattened module in a format that is closer to source. See also `--dump-mi`.

- `-r` or `--recompile` — (re)compiles the module to machine-dependent form after flattening; this mode is the default.

- `--work` ⟨*dir*⟩ — uses ⟨*dir*⟩ to cache compiled modules in an intermediate form for flattening; using `--work` with the same ⟨*dir*⟩ for multiple uses of `raco demod` can greatly speed up demodularization, and since the cache is based on `raco make`, it works even with different input files or when modules to be flattened have changed since the last use of the cache.

- `-g` or `--prune-definitions` — increases pruning of definitions that are unreferenced on the unsound assumption that the right-hand side of a definition has no side effect. When syntax is preserved, a definition can be pruned as long as no syntax literal includes an identifier that is bound to the definition. Since these assumptions are unchecked, conversion may not preserve the behavior of the input module. For backward compatibility, `--garbage-collect` is an alias for `--prune-definitions`.

- `--dump` ⟨*file*⟩ — writes an S-expression representation of the module's content to ⟨*file*⟩, which can be helpful for understanding the content that is in the compiled flatten module.

- `--dump-mi` ⟨*file*⟩ — writes a machine-independent form of the flattened module to ⟨*file*⟩, the same as `-M` would write, but useful when `-M` is not used.

In addition to preserving submodules or of the source module, demodularization may introduce new submodules to hold portions of the flattening. The introduced submodules have names `demod-pane-` followed by an integer.

Changed in version 1.10 of package `compiler-lib`: Added `-M`/`--compile-any`, `--work`, and support for Racket CS.

Changed in version 1.15: Added `-x`/`--exclude-library`, `-s`/`--syntax`, `--dump`, `--dump-mi`, `--prune-definitions` (as a new name for `--garbage-collect`), and preservation of submodules.

Changed in version 1.16: Changed to reporting an error when a module named by `-x` or `-e` is not a dependency of the input module.

8.1 Demodularizing Libraries

Demodularization of a library module with `compiler/demod` can create a module whose meaning is different than the original, since transitive dependencies (that are not specified as excluded) are copied into the flattened module. That copying can break sharing as needed for generated structure types or bindings. As a specific example, separate copies of

racket/base will have distinct and incompatible implementations of keyword arguments for procedures.

To avoid problems, a good general strategy for flattening is

- put all modules to be flattened into an "private/amalgam" subcollection of, where modules within "private/amalgam" can freely refer to each other;
- create a module "private/amalgam-src.rkt" that requires modules from "private/amalgam" that need to be accessible from outside, where submodules in "private/amalgam-src.rkt" can provided different subsets of bindings from "private/amalgam";
- create a module "mine/private/amalgam.rkt" as

  ```
  #lang compiler/demod
  "amalgam-src.rkt"
  #:include (#:dir "amalgam")
  ```

  and
- from outside "private/amalgam", use only "private/amalgam.rkt", perhaps via public modules that reprovide from "private/amalgam.rkt".

8.2 Language for Demodularizing

#lang compiler/demod package: compiler-lib

A module using compiler/demod language compiles to a form that is the flattened (in the same sense as raco demod) version of a source module. See also §8.1 "Demodularizing Libraries".

A #lang compiler/demod module body starts with a *module-path* to flatten, it may be followed by options:

```
module-path
option
...
```

```
    option = mode
           | #:include (mod-spec ...)
           | #:exclude (mod-spec ...)
           | #:submodule-include (submod-spec ...)
           | #:submodule-exclude (submod-spec ...)
           | #:prune-definitions
           | #:dump file
           | #:dump-mi file
           | #:no-demod

      mode = #:exe
           | #:dynamic
           | #:static

  mod-spec = #:module module-path
           | #:dir dir-path
           | #:collect collect-name

submod-spec = identifier
           | (identifier ...)
```

The default *mode* is `#:dynamic`, which preserves syntax objects and compile-time support (like macros), but does not insist that all modules are copied into the flattened module. For example, if a module is referenced by a combination of submodules within *module-path* and no other module is reached by the same combination, then the benefit of copying the module into a submodule is limited. The `#:static` mode is like `#:dynamic`, but it ensures that all modules are included unless they are specified as excluded. The `#:exe` mode discards syntax and compile-time support, so it may be suitable for flattening a module that implements an end-user program.

When the `#:include` option is specified, then only modules covered by a *mod-spec* will be included in the flattened form; otherwise, all modules are candidates for inclusion. When the `#:exclude` option is specified, the modules covered by the *mod-spec*s are excluded, even if they would otherwise be included according to a `#:include` specification. In other words, `#:exclude` is applied after `#:include`. Each *mod-spec* must name a module by a filesystem or collection-based path, and it must not name a submodule; any submodule of the named module is implicitly included or excluded. If a *mod-spec* in the `#:include` or `#:exclude` list is not a dependency of `module-path` (and has no submodules that are dependencies), then an exception is raised.

The `#:submodule-include` and `#:submodule-exclude` specifications are analogous to `#:include` and `#:exclude`, but for submodules immediately with *module-path*. If *mode* is `#:exe`, then the list of inclusions defaults to `main` and `configure-runtime`, otherwise the default is to have no specific inclusions.

A *mod-spec* either indicates a specific module with `#:module` or it indicates all modules in

a given collection (and its subcollections) with `#:collect`. A *`collect-name`* is always a string with `/`-separated components.

If the `#:prune-definitions` option is specified, then unused definitions from the original module and its dependencies are more aggressively pruned, but unsoundly. When syntax is preserved for `#:dynamic` or `#:static` mode, then all definitions are normally preserved from the original module, because they might be reachable via `datum->syntax`; when `#:prune-definitions` is specified, a definition can be pruned if no syntax object literal includes an identifier bound to the definition. Meanwhile, in all modes including `#:exe`, a definition is normally preserved if its right-hand side might have a side effect, but `#:prune-definitions` allows pruning on the unchecked assumption that a definition has no side effect. Due to its unchecked assumptions, `#:prune-definitions` may not preserve the behavior of the input module.

As an example of where `#:prune-definitions` can go wrong, a module could export a macro that expands to a use of `syntax-parse`, and that use could include a `:` shorthand to combine a pattern variable and a syntax class (also defined in the module) as one identifier. The identifier would be split into variable and syntax-class components only when the macro is used, so the shorthand does not count as a literal that is bound to the syntax class. In that particular situation, use `~var` instead of the shorthand, and then the syntax class is referenced by its own identifier. Meanwhile, a macro that is not exported (directly or indirectly through another macro) can safely use the `:` shorthand, since its expansions are part of the module's implementation.

If the `#:no-demod` option is specified, then *`mod-spec`* is not flattened, after all. Instead, the new module `requires` and `reprovides` *`mod-spec`* and each of its submodules. This mode is always used when a `compiler/demod` module is expanded, since expansion must produce syntax instead of a compiled module. This mode also may be useful during for development to avoid longer compile times from flattening or to check whether copying of modules for flattening creates any trouble.

A flattened module using `compiler/demod` has a build dependency on the original module, so a tool like `raco make` or `raco setup` will trigger reflattening if the source module changes, but the flattened module does not have a run-time or expand-time dependency on the original module. Modules excluded from the flattening via `#:include` and `#:exclude` remain as run-time and expand-time dependencies of the flattened module. In the default `#:dynamic` mode, additional dependencies may be preserved for modules that cannot be usefully merged, but `#:static` or `#:exe` mode copies even those modules into new submodules.

Compilation and expansion of a `compiler/demod` module creates a `"compiled/ephemeral/demod"` subdirectory in the same directory as the module. That subdirectory that holds freshly compiled versions of all dependencies of the flattened module in a form that is suitable for demodularization. This extra compilation is managed using `compiler/cm`, so changes to dependencies can be handled incrementally, but still separate from normal compilation of the dependencies. Detecting that the compilation of the `compiler/demod` module is up-to-date does not depend on the `"compiled/ephemeral/demod"` subdirectory, so it can be safely discarded after compilation.

Added in version 1.15 of package `compiler-lib`.

Changed in version 1.16: Changed to raising an exception when a module listed in `#:include` or `#:iexclude` is not a dependency of `module-path`.

9 `raco link`: Library Collection Links

The `raco link` command inspects and modifies a collection links file to display, add, or remove mappings from collection names to filesystem directories.

Managing links directly is somewhat discouraged. Instead, use the package manager (see *Package Management in Racket*), which installs and manages links (i.e., it builds on `raco link`) in a way that more gracefully leads to sharing collections with others. Nevertheless, `raco link` is available for direct use.

For example, the command

```
raco link maze
```

installs a user-specific and version-specific link for the `"maze"` collection, mapping it to the `"maze"` subdirectory of the current directory. Supply multiple directory paths to create multiple links at once, especially with a command-shell wildcard:

```
raco link *
```

By default, the linked collection name is the same as each directory's name, but the collection name can be set separately for a single directory with the `--name` flag.

To remove the link created by the first example above, use

```
raco link --remove maze
```

or

```
raco link -r maze
```

Like link-adding mode, removing mode accepts multiple directory paths to remove multiple links, and all links that match any directory are removed. If `--name` is used with `--remove`, then only links matching both the collection name and directory are removed.

Full command-line options:

- `-l` or `--list` — Shows the current link table. If any other command-line arguments are provided that modify the link table, the table is shown after modifications. If no directory arguments are provided, and if none of `-u`, `--user`, `-i`, `--installation`, `-f`, or `--file` are specified, then the link table is shown for all user-specific and installation-wide collection links files.

- `-n` ⟨*name*⟩ or `--name` ⟨*name*⟩ — Sets the collection name for adding a single link or removing matching links. By default, the collection name for an added link is derived from the directory name. When the `-r` or `--remove` flag is also used, only links with a collection name matching ⟨*name*⟩ are removed, and if no directory arguments

are provided, all links with a match to ⟨*name*⟩ are removed. This flag is mutually exclusive with `-d` and `--root`.

- `-d` or `--root` — Treats each directory as a collection root that contains collection directories, instead of a directory for a specific collection. When the `-r` or `--remove` flag is also used, only collection-root links that match a directory are removed. This flag is mutually exclusive with `-n` and `--name`.

- `-D` or `--static-root` — Like `-d` or `--root`, but each directory is assumed to have a constant set of subdirectories (to improve the use of collection-search caches) as long as the links file itself does not change.

- `-x` ⟨*regexp*⟩ or `--version-regexp` ⟨*regexp*⟩ — Sets a version regexp that limits the link to use only by Racket versions (as reported by `version`) matching ⟨*regexp*⟩. This flag is normally used with `-u` or `--user` with installations that have different versions but the same installation name. When the `-r` or `--remove` flag is also used, only links with a version regexp matching ⟨*regexp*⟩ are removed.

- `-r` or `--remove` — Selects remove mode instead of add mode.

- `-u` or `--user` — Limits listing and removal of links to the user-specific collection links file and not the installation-wide collection links file. This flag is mutually exclusive with `-i`, `--installation`, `-f`, and `--file`.

- `-i` or `--installation` — Reads and writes links in installation-wide collection links file and not the user-specific collection links file. This flag is mutually exclusive with `-u`, `--user`, `-f`, and `--file`.

- `-f` ⟨*file*⟩ or `--file` ⟨*file*⟩ — Reads and writes links in ⟨*file*⟩ instead of the user-specific collection links file. This flag is mutually exclusive with `-u`, `--user`, `-s`, `--shared`, `-i`, and `--installation`.

- `-v` ⟨*vers*⟩ or `--version` ⟨*vers*⟩ — Selects ⟨*vers*⟩ as relevant installation name for operations on the user-specific collection links file.

- `--repair` — Enables repairs to the existing file content when the content is erroneous. The file is repaired by deleting individual links when possible.

9.1 API for Collection Links

`(require setup/link)` package: `base`

```
(links  dir
        ...
       [#:user? user?
        #:user-version user-version
        #:file file
        #:name name
        #:root? root?
        #:static-root? static-root?
        #:version-regexp version-regexp
        #:error error-proc
        #:remove? remove?
        #:show? show?
        #:repair? repair?
        #:with-path? with-path?])        → list?
  dir : path?
  user? : any/c = #t
  user-version : string? = (get-installation-name)
  file : (or/c path-string? #f) = #f
  name : (or/c string? #f) = #f
  root? : any/c = #f
  static-root? : any/c = #f
  version-regexp : (or/c regexp? #f) = #f
  error-proc : (symbol? string? any/c ... . -> . any) = error
  remove? : any/c = #f
  show? : any/c = #f
  repair? : any/c = #f
  with-path? : any/c = #f
```

A function version of the `raco link` command that always works on a single file—either *file* if it is a path string, the user–specific collection links file if *user?* is true, or the installation-wide collection links file otherwise. If *user?* is true, then *user-version* determines the relevant installation name (defaulting to the current installation's name).

The *static-root?* flag value is ignored unless *root?* is true and *remove?* is false, in which case each given *dir* is added as a static root if *static-root?* is true.

The *error-proc* argument is called to raise exceptions that would be fatal to the `raco link` command.

If *remove?* is true, the result is a list of entries that were removed from the file. If *remove?* is `#f` but *root?* is true, the result is a list of paths for collection roots. If *remove?* and *root?* are both `#f`, the result is a list for top-level collections that are mapped by *file* and that apply to the running version of Racket; the list is a list of strings for collection names if *with-path?* is `#f`, or it is a list of pairs of collection-name strings and complete paths if *with-path?* is true.

10 `raco pack`: Packing Library Collections

The `raco pack` command creates an archive of files and directories. Formerly, such archives were used directly to distribute library files to Racket users, but the package manager (see *Package Management in Racket*) is now the preferred mechanism for distribution.

A packed archive usually has the suffix `".plt"`. The `raco pkg` command recognizes a `".plt"` archive for installation as a package. The `raco setup` command (see §6 "`raco setup`: Installation Management") also supports `".plt"` unpacking and installation when using the `-A` flag, but such installations do not benefit from the more general management facilities of `raco pkg`, while the `raco unpack` command (see §11 "`raco unpack`: Unpacking Library Collections") unpacks an archive locally without attempting to install it. DrRacket recognizes the `".plt"` and currently treats such an archive in the same way as `raco setup -A`.

An archive contains the following elements:

- A set of files and directories to be unpacked, and flags indicating whether they are to be unpacked relative to the Racket add-ons directory (which is user-specific), the Racket installation directory, or a user-selected directory.

  The files and directories for an archive are provided on the command line to `raco pack`, either directly or in the form of collection names when the `--collect` flag is used.

  The `--at-plt` flag indicates that the files and directories should be unpacked relative to the user's add-ons directory, unless the user specifies the Racket installation directory when unpacking. The `--collection-plt` flag implies `--at-plt`. The `--all-users` flag overrides `--at-plt`, and it indicates that the files and directories should be unpacked relative to the Racket installation directory, always.

- A flag for each file indicating whether it overwrites an existing file when the archive is unpacked; the default is to leave the old file in place, but the `--replace` flag enables replacing for all files in the archive.

- A list of collections to be set-up (via `raco setup`) after the archive is unpacked; the `++setup` flag adds a collection name to the archive's list, but each collection for `--collection-plt` is added automatically.

- A name for the archive, which is reported to the user by the unpacking interface; the `--plt-name` flag sets the archive's name, but a default name is determined automatically when using `--collect`.

- A list of required collections (with associated version numbers) and a list of conflicting collections; the `raco pack` command always names the `"racket"` collection in the required list (using the collection's pack-time version), `raco pack` names each packed collection in the conflict list (so that a collection is not unpacked on top of a

different version of the same collection), and `raco pack` extracts other requirements and conflicts from the `"info.rkt"` files of collections when using `--collect`.

Specify individual directories and files for the archive when not using `--collect`. Each file and directory must be specified with a relative path. By default, if the archive is unpacked with DrRacket, the user will be prompted for a target directory, and if `raco setup` is used to unpack the archive, the files and directories will be unpacked relative to the current directory. If the `--at-plt` flag is provided, the files and directories will be unpacked relative to the user's Racket add-ons directory, instead. Finally, if the `--all-users` flag is provided, the files and directories will be unpacked relative to the Racket installation directory, instead.

Use the `--collect` flag to pack one or more collections; sub-collections can be designated by using a `/` as a path separator on all platforms. In this mode, `raco pack` automatically uses paths relative to the Racket installation or add-ons directory for the archived files, and the collections will be set-up after unpacking. In addition, `raco pack` consults each collection's `"info.rkt"` file, as described below, to determine the set of required and conflicting collections. Finally, `raco pack` consults the first collection's `"info.rkt"` file to obtain a default name for the archive. For example, the following command creates a `"sirmail.plt"` archive for distributing a `"sirmail"` collection:

```
raco pack --collect sirmail.plt sirmail
```

When packing collections, `raco pack` checks the following fields of each collection's `"info.rkt"` file (see §6.4 "`"info.rkt"` File Format"):

- `requires` — A list of the form `(list (list coll vers) ...)` where each *coll* is a non-empty list of relative-path strings, and each *vers* is a (possibly empty) list of exact integers. The indicated collections must be installed at unpacking time, with version sequences that match as much of the version sequence specified in the corresponding `vers`.

  A collection's version is indicated by a `version` field in its `"info.rkt"` file, and the default version is the empty list. The version sequence generalized major and minor version numbers. For example, version `'(2 5 4 7)` of a collection can be used when any of `'()`, `'(2)`, `'(2 5)`, `'(2 5 4)`, or `'(2 5 4 7)` is required.

- `conflicts` — A list of the form `(list coll ...)` where each *coll* is a non-empty list of relative-path strings. The indicated collections must *not* be installed at unpacking time.

For example, the `"info.rkt"` file in the `"sirmail"` collection might contain the following `info` declaration:

```
#lang info
(define name "SirMail")
(define mred-launcher-libraries (list "sirmail.rkt"))
```

```
(define mred-launcher-names (list "SirMail"))
(define requires (list (list "mred")))
```

Then, the "sirmail.plt" file (created by the command-line example above) will contain the name "SirMail." When the archive is unpacked, the unpacker will check that the "mred" collection is installed, and that "mred" has the same version as when "sirmail.plt" was created.

10.1 Format of ".plt" Archives

The extension ".plt" is not required for a distribution archive, but the ".plt"-extension convention helps users identify the purpose of a distribution file.

The raw format of a distribution file is described below. This format is uncompressed and sensitive to communication modes (text vs. binary), so the distribution format is derived from the raw format by first compressing the file using gzip, then encoding the gzipped file with the MIME base64 standard (which relies only the characters A-Z, a-z, 0-9, +, /, and =; all other characters are ignored when a base64-encoded file is decoded).

The raw format is

- PLT are the first three characters.
- An S-expression matching

```
(lambda (request failure)
  (case request
    [(name) name]
    [(unpacker) (quote mzscheme)]
    [(requires) (quote requires)]
    [(conflicts) (quote conflicts)]
    [(plt-relative?) plt-relative?]
    [(plt-home-relative?) plt-home-relative?]
    [(test-plt-dirs) test-dirs]
    [else (failure)]))
```

  where the *name*, *requires*, etc., meta-variables stand for S-expressions as follows:

  - *name* — a human-readable string describing the archive's contents. This name is used only for printing messages to the user during unpacking.
  - *requires* — a list of collections required to be installed before unpacking the archive, which associated versions; see the documentation of pack for details.
  - *conflicts* — a list of collections required *not* to be installed before unpacking the archive.

- *plt-relative?* — a boolean; if true, then the archive's content should be unpacked relative to the plt add-ons directory.
- *plt-home-relative?* — a boolean; if true and if 'plt-relative? is true, then the archive's content should be unpacked relative to the Racket installation.
- *test-plt-dirs* — #f or a 'paths where *paths* is a list of path strings; in the latter case, a true value of *plt-home-relative?* is cancelled if any of the directories in the list (relative to the Racket installation) is unwritable by the user.

The S-expression is extracted from the archive using read (and the result is *not* evaluated).

- An S-expression matching

```
(unit (import main-collects-parent-dir mzuntar)
      (export)
      (mzuntar void)
      (quote collections))
```

where *collections* is a list of collection paths (where each collection path is a list of strings); once the archive is unpacked, raco setup will compile and setup the specified collections.

The S-expression is extracted from the archive using read (and the result is *not* evaluated).

The archive continues with unpackables. Unpackables are extracted until the end-of-file is found (as indicated by an = in the base64-encoded input archive).

An *unpackable* is one of the following:

- The symbol 'dir followed by a list S-expression. The build-path procedure will be applied to the list to obtain a relative path for the directory (and the relative path is combined with the target directory path to get a complete path).

  The 'dir symbol and list are extracted from the archive using read (and the result is *not* evaluated).

- The symbol 'file, a list, a number, an asterisk, and the file data. The list specifies the file's relative path, just as for directories. The number indicates the size of the file to be unpacked in bytes. The asterisk indicates the start of the file data; the next n bytes are written to the file, where n is the specified size of the file.

  The symbol, list, and number are all extracted from the archive using read (and the result is *not* evaluated). After the number is read, input characters are discarded until an asterisk is found. The file data must follow this asterisk immediately.

- The symbol 'file-replace is treated like 'file, but if the file exists on disk already, the file in the archive replaces the file on disk.

10.2 API for Packing

```
(require setup/pack)
```
package: base

Although the `raco pack` command can be used to create most `".plt"` files, the `setup/pack` library provides a more general API for making `".plt"` archives.

```
(pack-collections-plt
 dest
 name
 collections
 [#:replace? replace?
  #:at-plt-home? at-home?
  #:test-plt-collects? test?
  #:extra-setup-collections collection-list
  #:file-filter filter-proc])
→ void?
 dest : path-string?
 name : string?
 collections : (listof (listof path-string?))
 replace? : boolean? = #f
 at-home? : boolean? = #f
 test? : boolean? = #t
 collection-list : (listof path-string?) = null
 filter-proc : (path-string? . -> . boolean?) = std-filter
```

Creates the `".plt"` file specified by the pathname *dest*, using the *name* as the name reported to `raco setup` as the archive's description.

The archive contains the collections listed in *collections*, which should be a list of collection paths; each collection path is, in turn, a list of relative-path strings.

If the `#:replace?` argument is `#f`, then attempting to unpack the archive will report an error when any of the collections exist already, otherwise unpacking the archive will overwrite an existing collection.

If the `#:at-plt-home?` argument is `#t`, then the archived collections will be installed into the Racket installation directory instead of the user's directory if the main `"collects"` directory is writable by the user. If the `#:test-plt-collects?` argument is `#f` (the default is `#t`) and the `#:at-plt-home?` argument is `#t`, then installation fails if the main `"collects"` directory is not writable.

The optional `#:extra-setup-collections` argument is a list of collection paths that are not included in the archive, but are set-up when the archive is unpacked.

The optional `#:file-filter` argument is the same as for `pack-plt`.

```
(pack-collections  dest
                   name
                   collections
                   replace?
                   extra-setup-collections
                  [filter
                   at-plt-home?])           → void?
  dest : path-string?
  name : string?
  collections : (listof (listof path-string?))
  replace? : boolean?
  extra-setup-collections : (listof path-string?)
  filter : (path-string? . -> . boolean?) = std-filter
  at-plt-home? : boolean? = #f
```

Old, keywordless variant of `pack-collections-plt` for backward compatibility.

```
(pack-plt  dest
           name
           paths
          [#:as-paths as-paths
           #:file-filter filter-proc
           #:encode? encode?
           #:file-mode file-mode-sym
           #:unpack-unit unpack-spec
           #:collections collection-list
           #:plt-relative? plt-relative?
           #:at-plt-home? at-plt-home?
           #:test-plt-dirs dirs
           #:requires mod-and-version-list
           #:conflicts mod-list])           → void?
  dest : path-string?
  name : string?
  paths : (listof path-string?)
  as-paths : (listof path-string?) = paths
  filter-proc : (path-string? . -> . boolean?) = std-filter
  encode? : boolean? = #t
  file-mode-sym : symbol? = 'file
  unpack-spec : any/c = #f
  collection-list : (listof path-string?) = null
  plt-relative? : any/c = #f
  at-plt-home? : any/c = #f
  dirs : (or/c (listof path-string?) #f) = #f
  mod-and-version-list : (listof (listof path-string?)    = null
                                 (listof exact-integer?))
```

mod-list : (listof (listof path-string?)) = null

Creates the ".plt" file specified by the pathname *dest*, using the string *name* as the name reported to raco setup as the archive's description. The *paths* argument must be a list of relative paths for directories and files; the contents of these files and directories will be packed into the archive. The optional *as-paths* list provides the path to be recorded in the archive for each element of *paths* (so that the unpacked paths can be different from the packed paths).

The #:file-filter procedure is called with the relative path of each candidate for packing. If it returns #f for some path, then that file or directory is omitted from the archive. If it returns 'file or 'file-replace for a file, the file is packed with that mode, rather than the default mode. The default is std-filter.

If the #:encode? argument is #f, then the output archive is in raw form, and still must be gzipped and mime-encoded (in that order). The default value is #t.

The #:file-mode argument must be 'file or 'file-replace, indicating the default mode for a file in the archive. The default is 'file.

The #:unpack-unit argument is usually #f. Otherwise, it must be an S-expression for the S-expression that describes unpacking; see §10.1 "Format of ".plt" Archives" more information about the unit. If the #:unpack-unit argument is #f, an appropriate S-expression is generated.

The #:collections argument is a list of collection paths to be compiled after the archive is unpacked. The default is the null.

If the #:plt-relative? argument is true (the default is #f), the archive's files and directories are to be unpacked relative to the user's add-ons directory or the Racket installation directories, depending on whether the #:at-plt-home? argument is true and whether directories specified by #:test-plt-dirs are writable by the user.

If the #:at-plt-home? argument is true (the default is #f), then #:plt-relative? must be true, and the archive is unpacked relative to the Racket installation directory. In that case, a relative path that starts with "collects" is mapped to the installation's main "collects" directory, and so on, for the following the initial directory names:

- "collects"
- "doc"
- "lib"
- "include"

If #:test-plt-dirs is a list, then #:at-plt-home? must be #t. In that case, when the archive is unpacked, if any of the relative directories in the #:test-plt-dirs list is

unwritable by the current user, then the archive is unpacked in the user's add-ons directory after all.

The `#:requires` argument should have the shape `(list (list coll-path version) ...)` where each *coll-path* is a non-empty list of relative-path strings, and each *version* is a (possibly empty) list of exact integers. The indicated collections must be installed at unpacking time, with version sequences that match as much of the version sequence specified in the corresponding *version*. A collection's version is indicated by the `version` field of its `"info.rkt"` file.

The `#:conflicts` argument should have the shape `(list coll-path ...)` where each *coll-path* is a non-empty list of relative-path strings. The indicated collections must *not* be installed at unpacking time.

```
(pack  dest
       name
       paths
       collections
      [filter
       encode?
       file-mode
       unpack-unit
       plt-relative?
       requires
       conflicts
       at-plt-home?]) → void?
  dest : path-string?
  name : string?
  paths : (listof path-string?)
  collections : (listof path-string?)
  filter : (path-string? . -> . boolean?) = std-filter
  encode? : boolean? = #t
  file-mode : symbol? = 'file
  unpack-unit : any/c = #f
  plt-relative? : boolean? = #t
  requires : (listof (listof path-string?)    = null
                     (listof exact-integer?))
  conflicts : (listof (listof path-string?)) = null
  at-plt-home? : boolean? = #f
```

Old, keywordless variant of `pack-plt` for backward compatibility.

```
(std-filter p) → boolean?
  p : path-string?
```

Returns `#t` unless *p*, after stripping its directory path and converting to a byte string, matches

one of the following regular expressions: `^[.]git`, `^[.]svn$`, `^CVS$`, `^[.]cvsignore`, `^compiled$`, `^doc`, `~$`, `^#.*#$`, `^[.]#`, or `[.]plt$`.

```
(mztar  path
        [#:as-path as-path]
        output
        filter
        file-mode)          → void?
  path : path-string?
  as-path : path-string? = path
  output : output-port?
  filter : (path-string? . -> . boolean?)
  file-mode : (or/c 'file 'file-replace)
```

Called by `pack` to write one directory/file *path* to the output port *output* using the filter procedure *filter* (see `pack` for a description of *filter*). The *path* is recorded in the output as *as-path*, in case the unpacked path should be different from the original path. The *file-mode* argument specifies the default mode for packing a file, either `'file` or `'file-replace`.

11 `raco unpack`: Unpacking Library Collections

The `raco unpack` command unpacks a `".plt"` archive (see §10 "`raco pack`: Packing Library Collections") to the current directory without attempting to install any collections. Use `raco pkg` (see *Package Management in Racket*) to install a `".plt"` archive as a package, or use `raco setup -A` (see §6 "`raco setup`: Installation Management") to unpack and install collections from a `".plt"` archive.

Command-line flags:

- `-l` or `--list` — lists the content of the archive without unpacking it.
- `-c` or `--config` — shows the archive configuration before unpacking or listing the archive content.
- `-f` or `--force` — replace files that exist already; files that the archive says should be replaced will be replaced without this flag.

11.1 Unpacking API

```
(require setup/unpack)
```
package: `base`

The `setup/unpack` library provides raw support for unpacking a `".plt"` file.

```
(unpack archive
        [main-collects-parent-dir
         print-status
         get-target-directory
         force?
         get-target-plt-directory]) → void?
  archive : path-string?
  main-collects-parent-dir : path-string? = (current-directory)
  print-status : (string? . -> . any)
               = (lambda (x) (printf "~a\n" x))
  get-target-directory : (-> path-string?)
                       = (lambda () (current-directory))
  force? : any/c = #f
  get-target-plt-directory : (path-string?
                              path-string?
                              (listof path-string?)
                              . -> . path-string?)
                           = (lambda (preferred-dir main-dir options)
                               preferred-dir)
```

Unpacks *archive*.

The *main-collects-parent-dir* argument is passed along to *get-target-plt-directory*.

The *print-status* argument is used to report unpacking progress.

The *get-target-directory* argument is used to get the destination directory for unpacking an archive whose content is relative to an arbitrary directory.

If *force?* is true, then version and required-collection mismatches (comparing information in the archive to the current installation) are ignored.

The *get-target-plt-directory* function is called to select a target for installation for an archive whose is relative to the installation. The function should normally return one if its first two arguments; the third argument merely contains the first two, but has only one element if the first two are the same. If the archive does not request installation for all uses, then the first two arguments will be different, and the former will be a user-specific location, while the second will refer to the main installation.

```
(fold-plt-archive archive
                  on-config-fn
                  on-setup-unit
                  on-directory
                  on-file
                  initial-value) → any/c
  archive : path-string?
  on-config-fn : (any/c any/c . -> . any/c)
  on-setup-unit : (any/c input-port? any/c . -> . any/c)
  on-directory : ((or/c path-string?
                        (list/c (or/c 'collects 'doc 'lib 'include)
                                path-string?))
                  any/c
                  . -> . any/c)
  on-file : (or/c ((or/c path-string?
                         (list/c (or/c 'collects 'doc 'lib 'include)
                                 path-string?))
                   input-port?
                   any/c
                   . -> . any/c)
                  ((or/c path-string?
                         (list/c (or/c 'collects 'doc 'lib 'include)
                                 path-string?))
                   input-port?
                   (or/c 'file 'file-replace)
                   any/c
                   . -> . any/c))
  initial-value : any/c
```

Traverses the content of *archive*, which must be a ".plt" archive that is created with the default unpacking unit and configuration expression. The configuration expression is not evaluated, the unpacking unit is not invoked, and files are not unpacked to the filesystem. Instead, the information in the archive is reported back through on-config, *on-setup-unit*, *on-directory*, and *on-file*, each of which can build on an accumulated value that starts with *initial-value* and whose final value is returned.

The *on-config-fn* function is called once with an S-expression that represents a function to implement configuration information. The second argument to on-config is *initial-value*, and the function's result is passed on as the last argument to *on-setup-unit*.

The *on-setup-unit* function is called with the S-expression representation of the installation unit, an input port that points to the rest of the file, and the accumulated value. This input port is the same port that will be used in the rest of processing, so if *on-setup-unit* consumes any data from the port, then that data will not be consumed by the remaining functions. (This means that on-setup-unit can leave processing in an inconsistent state, which is not checked by anything, and therefore could cause an error.) The result of *on-setup-unit* becomes the new accumulated value.

For each directory that would be created by the archive when unpacking normally, *on-directory* is called with the directory path (described more below) and the accumulated value up to that point, and its result is the new accumulated value.

For each file that would be created by the archive when unpacking normally, *on-file* is called with the file path (described more below), an input port containing the contents of the file, an optional mode symbol indicating whether the file should be replaced, and the accumulated value up to that point; its result is the new accumulated value. The input port can be used or ignored, and parsing of the rest of the file continues the same either way. After *on-file* returns control, however, the input port is drained of its content.

A directory or file path can be a plain path, or it can be a list containing 'collects, 'doc, 'lib, or 'include and a relative path. The latter case corresponds to a directory or file relative to a target installation's collection directory (in the sense of find-collects-dir), documentation directory (in the sense of find-doc-dir), library directory (in the sense of find-lib-dir), or "include" directory (in the sense of find-include-dir).

12 `raco ctool`: Working with C Code

The `raco ctool` command works in various modes (as determined by command-line flags) to support various tasks involving C code.

12.1 Compiling and Linking C Extensions

A *dynamic extension* is a shared library (a.k.a. DLL) that extends Racket using the C API. An extension can be loaded explicitly via `load-extension`, or it can be loaded implicitly through `require` or `load/use-compiled` in place of a source *file* when the extension is located at

```
(build-path "compiled" "native" (system-library-subpath)
            (path-add-suffix file (system-type 'so-suffix)))
```

relative to *file*.

For information on writing extensions, see *Inside: Racket C API*.

`raco ctool` is provided by the `"cext-lib"` package.

Three `raco ctool` modes help for building extensions:

- `--cc` : Runs the host system's C compiler, automatically supplying flags to locate the Racket header files and to compile for inclusion in a shared library.
- `--ld` : Runs the host system's C linker, automatically supplying flags to locate and link to the Racket libraries and to generate a shared library.
- `--xform` : Transforms C code that is written without explicit GC-cooperation hooks to cooperate with Racket's 3m garbage collector; see §8 "Overview (BC)" in *Inside: Racket C API*.

Compilation and linking build on the `dynext/compile` and `dynext/link` libraries. The following `raco ctool` flags correspond to setting or accessing parameters for those libraries: `--tool`, `--compiler`, `--ccf`, `--ccf`, `--ccf-clear`, `--ccf-show`, `--linker`, `++ldf`, `--ldf`, `--ldf-clear`, `--ldf-show`, `++ldl`, `--ldl-show`, `++cppf`, `++cppf ++cppf-clear`, and `--cppf-show`.

The `--3m` flag specifies that the extension is to be loaded into the 3m variant of Racket. The `--cgc` flag specifies that the extension is to be used with the CGC. The default depends on `raco`: `--3m` if `raco` itself is running in 3m, `--cgc` if `raco` itself is running in CGC.

12.1.1 API for 3m Transformation

`(require compiler/xform)` package: `cext-lib`

```
(xform  quiet?
        input-file
        output-file
        include-dirs
       [#:keep-lines? keep-lines?]) → any/c
  quiet? : any/c
  input-file : path-string?
  output-file : path-string?
  include-dirs : (listof path-string?)
  keep-lines? : boolean? = #f
```

Transforms C code that is written without explicit GC-cooperation hooks to cooperate with Racket's 3m garbage collector; see §8 "Overview (BC)" in *Inside: Racket C API*.

The arguments are as for `compile-extension`; in addition *keep-lines?* can be `#t` to generate GCC-style annotations to connect the generated C code with the original source locations.

The file generated by `xform` can be compiled via `compile-extension`.

12.2 Embedding Modules via C

`raco ctool` is provided by the `"cext-lib"` package.

The `--c-mods` mode for `raco ctool` takes a set of Racket modules and generates a C source file that can be used as part of program that embeds the Racket runtime system. See §9 "Embedding into a Program (BC)" in *Inside: Racket C API* for an explanation of embedding programs. The `--mods` mode is similar, but it generates the raw bytes for the compiled module without encoding the bytes in C declarations.

The generated source or compiled file embeds the specified modules. Generated C source defines a `declare_modules` function that puts the module declarations into a namespace. Thus, using the output of `raco ctool --c-mods`, a program can embed Racket with a set of modules so that it does not need a `"collects"` directory to load modules at run time.

If the embedded modules refer to runtime files, the files can be gathered by supplying the `--runtime` argument to `raco ctool --c-mods`, specifying a directory ⟨*dir*⟩ to hold the files. Normally, ⟨*dir*⟩ is a relative path, and files are found at run time in ⟨*dir*⟩ relative to the executable, but a separate path (usually relative) for run time can be specified with `--runtime-access`.

Typically, `raco ctool --c-mods` is used with `++lib` to specify a collection-based module path. For example,

```
raco ctool --c-mods base.c ++lib racket/base
```

generates a "base.c" whose declare_modules function makes racket/base available for use via the scheme_namespace_require or scheme_dynamic_require functions within the embedding application.

When a module file is provided to raco ctool --c-mods, then declare_modules declares a module with the symbolic name of the module file. For example,

```
raco ctool --c-mods base.c hello.rkt
```

creates a declare_modules that defines the module 'hello, which could be required into the current namespace with (namespace-require ''hello) or similarly at the C level:

```
p = scheme_make_pair(scheme_intern_symbol("quote"),
                     scheme_make_pair(scheme_intern_symbol("hello"),
                                      scheme_make_null()));
scheme_namespace_require(p);
```

13 `raco test`: Run tests

The `raco test` command requires and runs the (by default) `test` submodule associated with each path given on the command line. Command-line flags can control which submodule is run, whether to run the main module if no submodule is found, and whether to run tests directly, in separate processes (the default), or in separate places. The current directory is set to a test file's directory before running the file.

When an argument path refers to a directory, `raco test` recursively discovers and runs all files within the directory that end in a module suffix (see `get-module-suffixes`, but the suffixes always include `".rkt"`, `".scrbl"`, `".ss"`, and `".scm"`) or have a (possibly empty) list of command-line arguments provided by `test-command-line-arguments` in an `"info.rkt"` file, or as directed by `test-include-paths` in an `"info.rkt"` file. At the same time, `raco test` omits files and directories within a directory as directed by `test-omit-paths` in an `"info.rkt"` file.

A test is counted as failing if it logs a failing test code via `test-log!`, causes Racket to exit with a non-zero exit code, or (when `-e` or `--check-stderr` is specified) if it produces output on the error port.

The `raco test` command accepts several flags:

- `-c` or `--collection` — Interprets the arguments as collections whose content should be tested (in the same way as directory content), and makes `--process` the default testing mode.
- `-p` or `--package` — Interprets the arguments as packages whose contents should be tested (in the same way as directory content). All package scopes are searched for the first, most specific package scope. This flag also makes `--process` the default testing mode.
- `-l` or `--lib` — Interprets the arguments as libraries that should be tested. Each argument ⟨*arg*⟩ is treated as a module path `(lib "@nonterm{arg}")`. The default testing mode is `--direct` if a single module is specified, `--process` if multiple modules are specified.
- `-m` or `--modules` — Not only interprets the arguments as paths (which is the default mode), but treats them the same as paths found in a directory, which means ignoring a file argument that does not have a module extension or is not enabled explicitly via `test-command-line-arguments` or `test-include-paths` in an `"info.rkt"` file; meanwhile, paths that are otherwise enabled can be disabled via `test-omit-paths` in an `"info.rkt"` file. The default testing mode is `--direct` if a single path is specified, `--process` if multiple paths are specified.
- `--drdr` — Configures defaults to imitate the DrDr continuous testing system: ignore non-modules, run tests in separate processes (unless `--thread` or `--direct` is specified) use as many jobs as available processors (unless `--jobs` is specified),

set the default timeout to 90 seconds (unless `--timeout` is specified), create a fresh `PLTUSERHOME` and `TMPDIR` for each test, count stderr output as a test failure, quiet program output, provide empty program input, and print a table of results.

- `-s` ⟨*name*⟩ or `--submodule` ⟨*name*⟩ — Requires the submodule ⟨*name*⟩ rather than `test`. Supply `-s` or `--submodule` to run multiple submodules, or combine multiple submodules with `--first-avail` to run the first available of the listed modules. Beware that if you use `-s` multiple times but supply a single module file as an argument, the default mode is still `--direct` (which likely means fewer fresh module instantiations than `--process` or `--place` mode).

- `-r` or `--run-if-absent` — Requires the top-level module of a file if a relevant submodule is not present. This is the default mode.

- `-x` or `--no-run-if-absent` — Ignores a file if the relevant submodule is not present.

- `--first-avail` — When multiple submodule names are provided with `-s` or `--submodule`, runs only the first available submodule.

- `--configure-runtime` — Run a `configure-runtime` submodule (if any) of each specified module before the module or a submodule is run. This mode is the default when only a single module is provided or when `--process` or `--place` mode is specified, unless a submodule name is provided via `-s` or `--submodule`.

- `--direct` — Runs each test in a thread, using a single namespace's module registry to load all tests. This mode is the default if a single file is specified. Multiple tests can interfere with each other and the overall test run by exiting, using unsafe operations that block (and thus prevent timeout), and so on.

- `--process` — Runs each test in a separate operating-system process. This mode is the default if multiple files are specified or if a directory, collection, or package is specified.

- `--place` — Runs each test in a place, instead of in an operating-system process.

- `-j` ⟨*n*⟩ or `--jobs` ⟨*n*⟩ — Runs up to ⟨*n*⟩ test files in parallel.

- `--timeout` ⟨*seconds*⟩ — Sets the default timeout (after which a test counts as failed) to ⟨*seconds*⟩. Use `+inf.0` to allow tests to run without limit but allow `timeout` sub-submodule configuration. If any test fails due to a timeout, the exit status of `raco test` is 2 (as opposed to 1 for only non-timeout failures or 0 for success). The default timeout corresponds to `+inf.0` if not specified via `--timeout` or `--drdr`.

- `--fresh-user` — When running tests in a separate process, creates a fresh directory and sets `PLTUSERHOME` and `TMPDIR`. The `PLTADDONDIR` environment variable is also set so that the add-on directory (which is where packages are installed, for example) does *not* change for each test process.

- `--empty-stdin` — Provide an empty stdin to each test program.

- `-Q` or `--quiet-program` — Suppresses output from each test program.
- `-e` or `--check-stderr` — Count any stderr output as a test failure.
- `--deps` — If considering arguments as packages, also check package dependencies.
- `++ignore-stderr` ⟨*pattern*⟩ — Don't count stderr output as a test failure if it matches ⟨*pattern*⟩. This flag can be used multiple times, and stderr output is treated as success as long as it matches any one ⟨*pattern*⟩.
- `--errortrace` — Dynamically loads `errortrace` before running the tests. Note that already-compiled files will not include the tracing information.
- `-y` or `--make` — Enable automatic generation and update of compiled `".zo"` files. Specifically, the result of `(make-compilation-manager-load/use-compiled-handler)` is installed as the value of `current-load/use-compiled` before module-loading actions.
- `-q` or `--quiet` — Suppresses output of progress information, responsible parties, and varying output (see §13.3 "Responsible-Party and Varying-Output Logging").
- `--heartbeat` — Periodically report that a test is still running after the test has been running at least 5 seconds.
- `--table` or `-t` — Print a summary table after all tests. If a test uses `rackunit`, or if a test at least uses `test-log!` from `raco/testing` to log successes and failures, the table reports test and failure counts based on the log.
- `++arg` ⟨*argument*⟩ — Adds ⟨*argument*⟩ to the list of arguments to the invoked test module, so that the invoked module sees ⟨*argument*⟩ in its `current-command-line-arguments`. These arguments are combined with any arguments specified in `"info.rkt"` by `test-command-line-arguments`.
- `++args` ⟨*arguments*⟩ — The same as `++arg`, but ⟨*arguments*⟩ is treated as a whitespace-delimited list of arguments to add. To specify multiple arguments using this flag within a typical shell, ⟨*arguments*⟩ must be enclosed in quotation marks.
- `--output` or `-o` ⟨*file*⟩ — Save all stdout and stderr output into ⟨*file*⟩. The target ⟨*file*⟩ will be overwritten if it exists already.

Changed in version 1.1 of package `compiler-lib`: Added `--heartbeat`.
Changed in version 1.4: Changed recognition of module suffixes to use `get-module-suffixes`, which implies recognizing `".ss"` and `".rkt"`.
Changed in version 1.5: Added `++ignore-stderr`.
Changed in version 1.6: Added `++arg` and `++args`.
Changed in version 1.8: Added `--output` and `-o`.
Changed in version 1.11: Added `--make`/`-y`.
Changed in version 1.12: Added `--errortrace`.

13.1 Test Configuration by Submodule

When `raco test` runs a test in a submodule, a `config` sub-submodule can provide additional configuration for running the test. The `config` sub-submodule should use the `info` module language to define the following identifiers:

- `timeout` — a real number in seconds to override the default timeout for the test, which applies only when timeouts are enabled.
- `responsible` — a string, symbol, or list of symbols and strings identifying a responsible party that should be notified when the test fails. See §13.3 "Responsible-Party and Varying-Output Logging".
- `lock-name` — a string that names a lock file that is used to serialize tests (i.e., tests that have the same lock name do not run concurrently). The lock file's location is determined by the `PLTLOCKDIR` environment variable or defaults to `(find-system-path 'temp-dir)`. The maximum time to wait on the lock file is determined by the `PLTLOCKTIME` environment variable or defaults to 4 hours.
- `ignore-stderr` — a string, byte string, or regexp value, as a pattern that causes error output to not be treated as a failure if the output matches the pattern.
- `random?` — if true, indicates that the test's output is expected to vary. See §13.3 "Responsible-Party and Varying-Output Logging".

In order to prevent evaluation of a file for testing purposes, it suffices to create a submodule that does not perform any tests and does not trigger the evaluation of the enclosing module. So, for instance, a file might look like this:

```
#lang racket

(/ 1 0)

; don't run this file for testing:
(module test racket/base)
```

Changed in version 1.5 of package `compiler-lib`: Added `ignore-stderr` support.

13.2 Test Configuration by `"info.rkt"`

Submodule-based test configuration is preferred (see §13.1 "Test Configuration by Submodule"). In particular, to prevent `raco test` from running a particular file, normally the file should contain a submodule that takes no action.

In some cases, however, adding a submodule is inconvenient or impossible (e.g., because the file will not always compile). Thus, `raco test` also consults any `"info.rkt"` file in the candidate test file's directory. In the case of a file within a collection, `"info.rkt"` files from any enclosing collection directories are also consulted for `test-omit-paths` and `test-include-paths`. Finally, for a file within a package, the package's `"info.rkt"` is consulted for `pkg-authors` to set the default responsible parties (see §13.3 "Responsible-Party and Varying-Output Logging") for all files in the package.

The following `"info.rkt"` fields are recognized:

- `test-omit-paths` — a list of path strings (relative to the enclosing directory) and regexp values (to omit all files within the enclosing directory matching the expression), or `'all` to omit all files within the enclosing directory. When a path string refers to a directory, all files within the directory are omitted.

- `test-include-paths` — a list of path strings (relative to the enclosing directory) and regexp values (to include all files within the enclosing directory matching the expression), or `'all` to include all files within the enclosing directory. When a path string refers to a directory, all files within the directory are included.

- `test-command-line-arguments` — a list of `(list` *`module-path-string`* `(list` *`argument-path-string`* `...))`, where `current-command-line-arguments` is set to a vector that contains the *`argument-path-string`* when running *`module-path-string`*.

- `test-timeouts` — a list of `(list` *`module-path-string real-number`*`)` to override the default timeout in seconds for *`module-path-string`*.

- `test-responsibles` — a list of `(list` *`module-path-string party`*`)` or `(list 'all` *`party`*`)` to override the default responsible party for *`module-path-string`* or all files within the directory (except as overridden), respectively. Each *`party`* is a string, symbol, or list of symbols and strings. See §13.3 "Responsible-Party and Varying-Output Logging".

- `test-lock-names` — a list of `(list` *`module-path-string lock-string`*`)` to declare a lock file name for *`module-path-string`*. See `lock-name` in §13.1 "Test Configuration by Submodule".

- `test-ignore-stderrs` — a list of `(list` *`module-path-string pattern`*`)` or `(list 'all` *`pattern`*`)` to declare patterns of standard error output that are allowed a non-failures for *`module-path-string`* or all files within the directory. Each *`pattern`* must be a string, byte string, or regexp value. See `ignore-stderr` in §13.1 "Test Configuration by Submodule".

- `test-randoms` — a list of path strings (relative to the enclosing directory) for modules whose output varies. See §13.3 "Responsible-Party and Varying-Output Logging".

- `module-suffixes` and `doc-module-suffixes` — Used indirectly via `get-module-suffixes`.

Changed in version 1.5 of package `compiler-lib`: Added `test-ignore-stderrs` support.

13.3 Responsible-Party and Varying-Output Logging

When a test has a declared responsible party, then the test's output is prefixed with a

`raco test:`⟨*which*⟩ `@(test-responsible '`⟨*responsible*⟩`)`

line, where ⟨*which*⟩ is a space followed by an exact non-negative number indicating a parallel task when parallelism is enabled (or empty otherwise), and ⟨*responsible*⟩ is a string, symbol, or list datum.

When a test's output (as written to stdout) is expected to vary across runs—aside from varying output that has the same form as produced by `time`—then it should be declared as varying. In that case, the test's output is prefixed with a

`raco test:`⟨*which*⟩ `@(test-random #t)`

line.

13.4 Logging Test Results

`(require raco/testing)` package: `compiler-lib`

This module provides a general purpose library for tracking test results and displaying a summary message. The command `raco test` uses this library to display test results. Therefore, any testing framework that wants to integrate with `raco test` should also use this library to log test results.

Added in version 1.13 of package `compiler-lib`.

```
(test-log! result) → void?
  result : any/c
```

Adds a test result to the running log. If *result* is false, then the test is considered a failure.

```
(test-report [#:display? display?
              #:exit? exit?])
 → (cons/c exact-nonnegative-integer?
           exact-nonnegative-integer?)
  display? : any/c = #f
```

```
  exit? : any/c = #f
```

Processes the running test log. The first integer is the failed tests, the second is the total tests. If *display?* is true, then a message is displayed. If there were failures, the message is printed on `(current-error-port)`. If *exit?* is true, then if there were failures, calls `(exit 1)`.

```
(test-log-enabled?) → boolean?
(test-log-enabled? enabled?) → void?
  enabled? : any/c
 = #t
```

When set to `#f`, `test-log!` is a no-op. This is useful to dynamically disable certain tests whose failures are expected and shouldn't be counted in the test log, such as when testing a custom check's failure behavior.

```
(current-test-invocation-directory) → (or/c #f path?)
(current-test-invocation-directory path) → void?
  path : (or/c #f path-string?)
 = #f
```

Contains the directory from which tests were invoked by, *e.g.*, `raco test`. This may differ from `current-directory` when the test runner changes directory before invoking a specific test file and should be set by test runners to reflect the directory from which they were originally invoked.

This should be used by test reports to display appropriate path names.

Added in version 1.14 of package `compiler-lib`.

14 `raco docs`: Documentation Search

The `raco docs` command searches the documentation for the given identifiers or search terms.

Command-line flags:

- `-f` ⟨*name*⟩ or `--family` ⟨*name*⟩ — Navigate documentation as language family ⟨*name*⟩, which may affect the starting point if no search terms are given, may affect the ordering of results for a search, and may affect "top" and "up" navigation.
- `-h` or `--help` — Show help information for this command.
- `--` — Do not treat remaining arguments as switches.

15 `raco expand`: Macro Expansion

The `raco expand` command macro-expands and pretty-prints the contents of the given source files. See also `expand`.

Command-line flags:

- `-n` ⟨*n*⟩ or `--columns` ⟨*n*⟩ — format output for a display with ⟨*n*⟩ columns
- `-h` or `--help` — show help information for this command
- `--` — do not treat remaining arguments as switches

16 `raco read`: Reading and Pretty-Printing

The `raco read` command `read`s and pretty-prints the contents of the given files. This command is useful for showing how a `#reader` or `#lang`-based reader extension converts input to an S-expression. It is also useful for pretty-printing a term that is already in S-expression form.

Command-line flags:

- `-n` ⟨*n*⟩ or `--columns` ⟨*n*⟩ — format output for a display with ⟨*n*⟩ columns
- `-h` or `--help` — show help information for this command
- `--` — do not treat remaining arguments as switches

Added in version 1.3 of package `compiler-lib`.

17 `raco scribble`: Building Documentation

See *Scribble: The Racket Documentation Tool* for information on the `raco scribble` command, which is used to run and render a Scribble document.

18 Adding a `raco` Command

The set of commands supported by `raco` can be extended by installed packages, PLaneT packages, and other collections. A command is added by defining `raco-commands` in the `"info.rkt"` library of a collection (see §6.4 "`"info.rkt"` File Format"), and then `raco setup` (as called directly or as part of a package or PLaneT installation) must index the `"info.rkt"` file.

The value bound to `raco-commands` must be a list of *command specifications*, where each specification is a list of four values:

```
(list command-string
      implementation-module-path
      description-string
      prominence)
```

The *`command-string`* is the command name. Any unambiguous prefix of a command name can be supplied to `raco` to invoke the command.

The *`implementation-module-path`* names the implementation though a module path (in the sense of `module-path?`). The module is loaded and invoked through `dynamic-require` to run the command. The module can access command-line arguments through the `current-command-line-arguments` parameter, which is adjusted before loading the command module to include only the arguments to the command. The `current-command-name` parameter is also set to the command name used to load the command. When `raco help` is used on a command, the command is launched with an initial `--help` argument in `current-command-line-arguments`.

The *`description-string`* is a short string used to describe the command in response to `raco help`. The description should not be capitalized or end with a period.

The *`prominence`* value should be a real number or `#f`. A `#f` value means that the command should not be included in the short list of "frequently used commands." A number indicates the relative prominence of the command; the `help` command has a value of `110`, and probably no command should be more prominent. The `pack` tool, which is currently ranked as the least-prominent of the frequently used commands, has a value of `10`.

As an example, the `"info.rkt"` of the `"compiler"` collection might contain the

```
(define raco-commands
  '(("make" compiler/commands/make "compile source to byte-
code" 100)
    ("decompile" compiler/commands/decompile "decompile byte-
code" #f)))
```

so that `make` is treated as a frequently used command, while `decompile` is available as an

infrequently used command.

18.1 Command Argument Parsing

```
(require raco/command-name)
```
package: base

The raco/command-name library provides functions to help a raco command identify itself to users.

```
(current-command-name) → (or/c string? #f)
(current-command-name name) → void?
  name : (or/c string? #f)
```

The name of the command currently being loaded via dynamic-require, or #f if raco is not loading any command.

A command implementation can use this parameter to determine whether it was invoked via raco or through some other means.

```
(short-program+command-name) → string?
```

Returns a string that identifies the current command. When current-command-name is a string, then the result is the short name of the raco executable followed by a space and the command name. Otherwise, it is the short name of the current executable, as determined by stripping the path from the result of (find-system-path 'run-file). In either case, on Windows, an ".exe" extension is removed from the executable name.

The result of this function is suitable for use with command-line. For example, the decompile tool parses command-line arguments with

```
(define source-files
  (command-line
   #:program (short-program+command-name)
   #:args source-or-bytecode-file
   source-or-bytecode-file))
```

so that raco decompile --help prints

```
  usage: raco decompile [ <option> ... ] [<source-or-bytecode-file>]
...

  <option> is one of

    --help, -h
```

```
      Show this help
    --
      Do not treat any remaining argument as a switch (at this
level)

  Multiple single-letter switches can be combined after
  one `-`. For example, `-h-` is the same as `-h --`.
```

```
(program+command-name) → string?
```

Like short-program+command-name, but the path (if any) is not stripped from the current executable's name.

18.2 Accessing raco Commands

```
(require raco/all-tools)
```
package: base

The raco/all-tools library collects the raco-commands specifications for installed packages, PLaneT packages, and other collections.

```
(all-tools)
 → (hash/c string? (list/c string? module-path? string? (or/c real? #f)))
```

Returns a hashtable with collection names as keys and command specifications as values. For example, the following program invokes raco make file.rkt:

```
(require raco/all-tools)

(define raco-make-spec (hash-ref (all-tools) "make"))

(parameterize ([current-command-line-arguments (vector "file.rkt")])
  (dynamic-require (second raco-make-spec) #f))
```

19 Installation Configuration and Search Paths

A *configuration directory* path is built into the Racket executable as selected at install time, and its location can be changed via the `PLTCONFIGDIR` directory or `--config`/`-G` command-line flag. Use `find-config-dir` to locate the configuration directory.

Modify the `"config.rktd"` file in the configuration directory to configure other directories as described below. Use the `setup/dirs` library (which combines information from the configuration files and other sources) to locate configured directories, instead of reading `"config.rktd"` directly. A `"config.rktd"` file can also appear in the directory `(build-path (find-system-path 'addon-dir) "etc")`, but it controls only the results of `find-addon-tethered-console-bin-dir` and `find-addon-tethered-gui-bin-dir`.

The path of the *main collection directory* is built into the Racket executable, and it can be changed via the `--collects`/`-X` flag, so it has no entry in `"config.rktd"`. Most paths that are specified in `"config.rktd"` have default values that are relative to the main collection directory. The paths of the configuration directory and main collection directory thus work together to determine a Racket configuration.

A `"config.rktd"` file in the configuration directory should contain a `read`able hash table with any of the following symbolic keys, where a relative path is relative to the main collection directory:

- `'installation-name` — a string for the installation name, which is used to determine user- and version-specific paths, such as the initial path produced by `find-library-collection-paths` and the location of packages that are installed in `user` package scope. The default is `(version)`.
- `'collects-search-dirs` — a list of paths, strings, byte strings, or `#f` representing the search path for collections. Each `#f` in the list, if any, is replaced with the main collection directory.
- `'lib-dir` — a path, string, or byte string for the *main library directory*. It defaults to a `"lib"` sibling directory of the main collection directory.
- `'lib-search-dirs` — a list of paths, strings, byte strings, or `#f` representing the search path for directories containing foreign libraries. Each `#f` in the list, if any, is replaced with the default search path, which is the user- and version-specific `"lib"` directory followed by the main library directory.
- `'dll-dir` — a path, string, or byte string for a directory containing shared libraries for the main executable. It defaults to the main library directory.
- `'share-dir` — a path, string, or byte string for the *main shared-file directory*, which normally includes installed packages. It defaults to a `"share"` sibling directory of the main collection directory.

- 'share-search-dirs — analogous to 'lib-search-dirs, where #f is replaced by the default search path, which is a user- and version-specific directory followed by a directory as potentially configured via 'share-dir.

  Added in version 8.1.0.6.

- 'links-file — a path, string, or byte string for the collection links file. It defaults to a "links.rktd" file in the main shared-file directory.

- 'links-search-files — like 'lib-search-dirs, but for collection links file. A #f is replaced by the default search path, which has the links file as potentially configured via 'links-file. A user- and version-specific links file is always added to the beginning of a search.

- 'pkgs-dir — a path, string, or byte string for packages that have installation package scope. It defaults to "pkgs" in the main shared-file directory.

- 'pkgs-search-dirs — similar to 'lib-search-dirs, but for packages in roughly installation package scope. More precisely, a #f value in the list is replaced with the directory specified by 'pkgs-dir, and that point in the search list corresponds to installation scope. Paths before or after a #f value in the list can be selected as a scopes to start searches at that path's point in the list. Directories listed in 'pkgs-search-dirs typically oblige a corresponding entry in 'links-search-files, where the corresponding entry is "links.rktd" within the directory.

  Changed in version 7.0.0.19: Adapt the package-search path in a general way for a directory scope.

- 'compiled-file-roots — a list of paths and 'same used to initialize current-compiled-file-roots. A path, which is relative or absolute, can be specified as a string or byte string that is converted to a path with string->path or bytes->path, respectively.

  Added in version 8.0.0.9.

- 'bin-dir — a path, string, or byte string for the installation's directory containing executables. It defaults to a "bin" sibling directory of the main collection directory.

- 'gui-bin-dir — a path, string, or byte string for the installation's directory containing GUI executables. It defaults to a the 'bin-dir value, if configured, or otherwise defaults in a platform-specific way: to the "bin" sibling directory of the main collection directory on Unix, and to the parent of the main collection directory on Windows and Mac OS.

  Added in version 6.8.0.2.

- 'bin-search-dirs — like 'lib-search-dirs, but for finding executables such as racket. A #f is replaced by the default search path, which is a user- and version-specific directory followed by the main console executable directory as potentially configured via 'bin-dir.

  Added in version 8.1.0.6.

- 'gui-bin-search-dirs — like 'bin-search-dirs, but for GUI executables, and defaults to the 'bin-search-dirs value.

  Added in version 8.1.0.6.

- 'apps-dir — a path, string, or byte string for the installation's directory for ".desktop" files. It defaults to a "applications" subdirectory of the main shared-file directory.

- 'man-dir — a path, string, or byte string for the installation's man-page directory. It defaults to a "man" sibling directory of the main shared-file directory.

- 'man-search-dirs — analogous to 'lib-search-dirs, where #f is replaced by the default search path, which is a user- and version-specific directory followed by a directory as potentially configured via 'man-dir.

  Added in version 8.1.0.6.

- 'doc-dir — a path, string, or byte string for the main documentation directory. The value defaults to a "doc" sibling directory of the main collection directory.

- 'doc-search-dirs — analogous to 'lib-search-dirs, where #f is replaced by the default search path, which is a user- and version-specific directory followed by a directory as potentially configured via 'doc-dir.

- 'doc-search-url — a URL string that is augmented with version and search-tag queries to form a remote documentation reference.

- 'doc-open-url — a URL string or #f; a string supplies a URL that is used instead of a local path to search and maybe open documentation pages (which normally makes sense only in an environment where opening a local HTML file does not work).

- 'include-dir — a path, string, or byte string for the main directory containing C header files. It defaults to an "include" sibling directory of the main collection directory.

- 'include-search-dirs — like doc-search-dirs, but for directories containing C header files.

- 'info-domain-root — a path, string, byte string, of #f; used as a prefix to redirect the paths used for recording and finding "info.rkt" information via find-relevant-directories. It defaults to #f, which uses paths as-is.

  Added in version 8.10.0.4.

- 'catalogs — a list of URL strings used as the search path for resolving package names. An #f in the list is replaced with the default search path. A string that does not start with alphabetic characters followed by :// is treated as a path, where a relative path is relative to the configuration directory.

- 'default-scope — either "user" or "installation", determining the default package scope for package-management operations.

- 'download-cache-dir — a path string used as the location for storing downloaded package archives. When not specified, packages are cached in a "download-cache" directory in the user's cache directory as reported by (find-system-path 'cache-dir).

- 'download-cache-max-files and 'download-cache-max-bytes — real numbers that determine limits on the download cache. When not specified, the cache is allowed to hold up to 1024 files that total up to 64 MB.

- 'build-stamp — a string that identifies a build, which can be used to augment the Racket version number to more specifically identify the build. An empty string is normally appropriate for a release build. The default "banner" also shows the build stamp when non-empty.

- Changed in version 8.11.1.7: Added build stamp to "banner".

- 'main-language-family — a string that names the main language family. The default is "Racket".

  Added in version 8.14.0.5.

- 'base-documentation-packages — a list of strings, each of which names a package. Any documentation provided by the package and its dependencies is considered part of the distribution's base language. This classification affects the way that documentation search results are sorted and reported. The default is '("racket-doc").

  Added in version 8.14.0.5.

- 'distribution-documentation-packages — like 'base-documentation-packages, but identifies a larger set of documentation that is considered part of the distribution beyond (but normally including) the base language. The default is '("main-distribution").

  Added in version 8.14.0.5.

- 'absolute-installation? — a boolean that is #t if the installation uses absolute path names, #f otherwise.

- 'cgc-suffix — a string used as the suffix (before the actual suffix, such as ".exe") for a "CGC" executable. Use Windows-style casing, and the string will be downcased as appropriate (e.g., for a Unix binary name). A #f value means that if the racket binary identifies itself as CGC, then the suffix is "", otherwise it is "CGC".

- '3m-suffix — analogous to 'cgc-suffix, but for 3m. A #f value means that if the racket binary identifies itself as 3m, then the suffix is "", otherwise it is "BC".

- 'cs-suffix — analogous to 'cgc-suffix, but for CS. A #f value means that if the racket binary identifies itself as CS, then the suffix is "", otherwise it is "CS".

- 'config-tethered-console-bin-dir and 'config-tethered-gui-bin-dir — a path for a directory to hold extra copies of executables that are tied to the configuration directory (as reported by find-config-dir) that is active at the time

the executables are created. See also §6.19 “Tethered Installations”, `find-config-tethered-console-bin-dir`, and `find-config-tethered-gui-bin-dir`.

- `'interactive-file` and `'gui-interactive-file` — a module path to the interactive module that runs when the REPL runs on startup, unless the `-q`/`--no-init-file` is provided. Defaults to `'racket/interactive` and `'racket/gui/interactive`.