

Web Applications in Racket

Version 9.3.0.9

Jay McCarthy

September 30, 2026

This manual describes the Racket libraries for building Web applications.

§1 “Running Web Servlets” describes how to run the servlets you’ve written.

§2 “Stateful Servlets” and §3 “Stateless Servlets” describe two ways to write Web applications. §2 “Stateful Servlets” use the entire Racket language, but their continuations are stored in the Web server’s memory. §3 “Stateless Servlets” use a slightly restricted Racket language, but their continuations can be stored by the Web client or on a Web server’s disk. If you can, you want to use §3 “Stateless Servlets” for the improved scalability.

The §4 “HTTP: Hypertext Transfer Protocol” section describes the common library functions for manipulating HTTP requests and creating HTTP responses. In particular, this section covers cookies, authentication, and request bindings.

The final five sections (§5 “URL-Based Dispatch”, §6 “Formlets: Functional Form Abstraction”, §7 “Templates: Separation of View”, §8 “Page: Short-hand for Common Patterns”, and §9 “Testing Servlets”) cover utility libraries that ease the creation of typical Web applications.

This manual closes with a frequently asked questions section: §10 “Troubleshooting and Tips”.

1 Running Web Servlets

There are a number of ways to run Web servlets.

1.1 Instant Servlets

```
#lang web-server/insta      package: web-server-doc
```

The fastest way to get a servlet running in the Web server is to use the "Insta" language in DrRacket. Enter the following into DrRacket:

```
#lang web-server/insta

(define (start req)
  (response/xexpr
    `(html (head (title "Hello world!"))
          (body (p "Hey out there!")))))
```

And press Run. A Web browser will open up showing your new servlet. This servlet will only be accessible from your local machine.

Behind the scenes, DrRacket has used `serve/servlet` to start a new server that uses your `start` function as the servlet. You are given the entire `web-server/servlet` API.

The following API is provided to customize the server instance:

```
(no-web-browser) → void
```

Calling this will instruct DrRacket to *not* start a Web browser when you press Run.

```
(static-files-path path) → void
  path : path-string?
```

This instructs the Web server to serve static files, such as stylesheets and images, from `path`.

If you want more control over specific parameters, keep reading about `web-server/servlet-env`.

1.2 Simple Single Servlet Servers

```
(require web-server/servlet-env)      package: web-server-lib
```

The Web Server provides a way to quickly configure and start a servlet with more customizability than `web-server/insta` provides. This is provided by the `web-server/servlet-env` module.

1.2.1 Examples

Here is a simple example of its use:

```
#lang racket
(require web-server/servlet
         web-server/servlet-env)

(define (start req)
  (response/xexpr
   `(html (head (title "Hello world!"))
          (body (p "Hey out there!")))))

(serve/servlet start)
```

Unlike the `web-server/insta` language, `start` is not a special identifier, so we could just as well have written the example as:

```
#lang racket
(require web-server/servlet
         web-server/servlet-env)

(define (my-app req)
  (response/xexpr
   `(html (head (title "Hello world!"))
          (body (p "Hey out there!")))))

(serve/servlet my-app)
```

Let's look at some of the customizations `serve/servlet` allows.

Suppose you'd like to change the port to something else, change the last line to:

```
(serve/servlet my-app
 #:port 8080)
```

Suppose you want to accept connections from external machines:

```
(serve/servlet my-app
 #:listen-ip #f)
```

By default the URL for your servlet is "http://localhost:8000/servlets/standalone.rkt", suppose you wanted it to be "http://localhost:8000/hello.rkt":

```
(serve/servlet my-app
  #:servlet-path "/hello.rkt")
```

Suppose you wanted it to capture top-level requests:

```
(serve/servlet my-app
  #:servlet-regexp #rx"")
```

Or, perhaps just some nice top-level name:

```
(serve/servlet my-app
  #:servlet-path "/main")
```

Suppose you wanted to use a style-sheet ("style.css") found on your Desktop ("/Users/jay/Desktop/"):

```
(serve/servlet my-app
  #:extra-files-paths
  (list
    (build-path "/Users/jay/Desktop")))
```

These files are served *in addition* to those from the `#:server-root-path` "htdocs" directory. You may pass any number of extra paths.

If you want to use `serve/servlet` in a start up script for a Web application, and don't want a browser opened or the DrRacket banner printed, then you can write:

```
(serve/servlet my-app
  #:command-line? #t)
```

Stateless Servlets

Suppose you would like to start a server for a stateless Web servlet "servlet.rkt" that provides `start`:

```
#lang racket
(require "servlet.rkt"
  web-server/servlet-env)

(serve/servlet start #:stateless? #t)
```

You can also put the call to `serve/servlet` in the `web-server` module directly:

```
#lang web-server
(require web-server/servlet-env)

(define (start req)
  (start
    (send/suspend
      (lambda (k-url)
        (response/xexpr
          `(html (body (a ([href ,k-url]) "Hello world!"))))))))

(serve/servlet start #:stateless? #t)
```

Like always, you don't even need to save the file.

1.2.2 Full API

```

(serve/servlet
  start
  [#:command-line? command-line?
   #:connection-close? connection-close?
   #:launch-browser? launch-browser?
   #:quit? quit?
   #:banner? banner?
   #:listen-ip listen-ip
   #:port port
   #:max-waiting max-waiting
   #:safety-limits safety-limits
   #:servlet-path servlet-path
   #:servlet-regexp servlet-regexp
   #:stateless? stateless?
   #:stuffer stuffer
   #:manager manager
   #:servlet-namespace servlet-namespace
   #:server-root-path server-root-path
   #:extra-files-paths extra-files-paths
   #:servlets-root servlets-root
   #:servlet-current-directory servlet-current-directory
   #:file-not-found-responder file-not-found-responder
   #:servlet-loading-responder responders-servlet-loading
   #:servlet-responder responders-servlet
   #:mime-types-path mime-types-path
   #:ssl? ssl?
   #:ssl-cert ssl-cert
   #:ssl-key ssl-key
   #:log-file log-file
   #:log-format log-format])
→ any
start : (request? . -> . can-be-response?)
command-line? : boolean? = #f
connection-close? : boolean? = #f
launch-browser? : boolean? = (not command-line?)
quit? : boolean? = (not command-line?)
banner? : boolean? = (not command-line?)
listen-ip : (or/c #f string?) = "127.0.0.1"
port : listen-port-number? = 8000
max-waiting : exact-nonnegative-integer? = 511
safety-limits : safety-limits?
               = (make-safety-limits #:max-waiting max-waiting)
servlet-path : string? = "/servlets/standalone.rkt"
servlet-regexp : regexp?
               = (regexp
                  (format "~a$" (regexp-quote servlet-path)))

```

```

stateless? : boolean? = #f
stuffer : (stuffer/c serializable? bytes?) = default-stuffer
manager : manager?
          = (make-threshold-LRU-manager #f (* 128 1024 1024))
servlet-namespace : (listof module-path?) = empty
server-root-path : path-string?
                  = (collection-path "web-server" "default-web-root")
extra-files-paths : (listof path-string?)
                  = (list (build-path server-root-path "htdocs"))
servlets-root : path-string?
               = (build-path server-root-path "htdocs")
servlet-current-directory : path-string? = servlets-root
file-not-found-responder : (request? . -> . can-be-response?)
                          = (gen-file-not-found-responder
                             (build-path
                               server-root-path
                               "conf"
                               "not-found.html"))
responders-servlet-loading : (url? any/c . -> . can-be-response?)
                            = servlet-loading-responder
responders-servlet : (url? any/c . -> . can-be-response?)
                    = servlet-error-responder
mime-types-path : path-string? = ....
ssl? : boolean? = #f
ssl-cert : (or/c #f path-string?)
           = (and ssl? (build-path server-root-path "server-cert.pem"))
ssl-key : (or/c #f path-string?)
          = (and ssl? (build-path server-root-path "private-key.pem"))
log-file : (or/c #f path-string? output-port?) = #f
log-format : (or/c log-format/c format-reqresp/c)
             = 'apache-default

```

This sets up and starts a fairly default server instance.

`start` is loaded as a servlet and responds to requests that match `servlet-regexp`. The current directory of servlet execution is `servlet-current-directory`.

If `launch-browser?` is true, then a web browser is opened to "http://localhost:<port><servlet-path>". `servlet-path` has no other purpose, if `servlet-regexp` is provided.

If `quit?` is true, then the URL `"/quit"` ends the server.

If `stateless?` is true, then the servlet is run as a stateless

```
#lang web-server
```

module and *stuffer* is used as the stuffer.

serve/servlet is simpler interface over *serve/launch/wait*, *dispatch/servlet*, and a few of the standard §2 “Dispatchers”. Some options, like *port* and *safety-limits* are transparently passed to *serve/launch/wait*. Some advanced customization requires using these underlying pieces of the *web-server* directly. However, many simpler customizations do not, which the rest of this section describes.

The server listens on *listen-ip* and port *port*. If *listen-ip* is *#f*, then the server accepts connections to all of the listening machine’s addresses. Otherwise, the server accepts connections only at the interface(s) associated with the given string. For example, providing “127.0.0.1” (the default) as *listen-ip* creates a server that accepts only connections to “127.0.0.1” (the loopback interface) from the local machine.

If *ssl-cert* and *ssl-key* are not false, then the server runs in HTTPS mode with *ssl-cert* and *ssl-key* as the certificates and private keys.

The servlet is loaded with *manager* as its continuation manager. (The default manager limits the amount of memory to 64 MB and deals with memory pressure as discussed in the *make-threshold-LRU-manager* documentation.)

The server files are rooted at *server-root-path* (which is the “default-web-root” directory in the *web-server* collection by default.) File paths, in addition to the “htdocs” directory under *server-root-path* may be provided with *extra-files-paths*. These paths are checked first, in the order they appear in the list.

Other servlets are served from *servlets-root*. The modules specified by *servlet-namespace* are shared between servlets found in *servlets-root* and the current namespace (and therefore the *start* procedure.)

If a file cannot be found, *file-not-found-responder* is used to generate an error response. If a servlet fails to load, *responders-servlet-loading* is used. If a servlet errors during its operation, *responders-servlet* is used.

If *banner?* is true, then an informative banner is printed. You may want to use this when running from the command line, in which case the *command-line?* option controls similar options.

MIME types are looked up at *mime-types-path*. By default the “mime.types” file in the *server-root-path* is used, but if that file does not exist, then the file that ships with the Web Server is used instead. Of course, if a path is given, then it overrides this behavior.

If *log-file* is given, then it used to log requests using *log-format* as the format. Allowable formats are those allowed by *log-format->format*. If *log-format* is a function, it is used directly to render the log entry. If *log-file* is a filepath, the given file is opened and written in a different thread. If *log-file* is an *output-port?*, logs are written directly to the port.

If `connection-close?` is `#t`, then every connection is closed after one request. Otherwise, the client decides based on what HTTP version it uses.

Changed in version 1.3 of package `web-server-lib`: Added support for providing `log-file` as an output port.

Changed in version 1.6: Added the `safety-limits` argument as with `serve/launch/wait`: see compatibility note.

1.3 Command-line Tools

One command-line utility is provided with the Web Server:

```
plt-web-server [-f <file-name> -p <port> -a <ip-address> --ssl]
```

The optional `file-name` argument specifies the path to a `configuration-table` S-expression (see `configuration-table->sexpr` for the syntax documentation.) If this is not provided, the default configuration shipped with the server is used. The optional `port` and `ip-address` arguments override the corresponding portions of the `configuration-table`. If the `ssl` option is provided, then the server uses HTTPS with "server-cert.pem" and "private-key.pem" in the current directory, with 443 as the default port. (See the `openssl` module for details on the SSL implementation.)

The `configuration-table` is given to `configuration-table->web-config@` and used to construct a `web-config^` unit, and is linked with the `web-server@` unit. The resulting unit is invoked, and the server runs until the process is killed.

2 Stateful Servlets

```
(require web-server/servlet)      package: web-server-lib
```

2.1 Example

A stateful servlet should provide the following exports:

```
interface-version : (one-of/c 'v2)
```

This indicates that the servlet is a version two servlet.

```
manager : manager?
```

The manager for the continuations of this servlet. See §2.5 “Continuation Managers” for options.

```
(start initial-request) → can-be-response?  
initial-request : request?
```

This function is called when an instance of this servlet is started. The argument is the HTTP request that initiated the instance.

An example version 2 module:

```
#lang racket  
(require web-server/http  
         web-server/managers/none)  
(provide interface-version manager start)  
  
(define interface-version 'v2)  
(define manager  
  (create-none-manager  
    (lambda (req)  
      (response/xexpr  
        `(html (head (title "No Continuations Here!"))  
              (body (h1 "No Continuations Here!"))))))))  
(define (start req)  
  (response/xexpr  
    `(html (head (title "Hello World!"))  
          (body (h1 "Hi Mom!")))))
```

These servlets have an extensive API available to them: `net/url`, `web-server/http`, `web-server/http/bindings`, `web-server/servlet/servlet-structs`, `web-server/servlet/web`, `web-server/servlet/web-cells`, and `web-server/dispatch`. Some of these are documented in the subsections that follow.

2.2 Responses

```
(require web-server/servlet/servlet-structs)
package: web-server-lib
```

Servlets communicate to the Web Server by returning HTTP responses. In order to accommodate lightweight programs (and backwards compatibility), the Web Server provides an indirection from application-specific response formats and the internal HTTP response format, `response`.

```
(can-be-response? x) → boolean?
  x : any/c
(any->response x) → (or/c false/c response?)
  x : any/c
(set-any->response! new-any->response) → void
  new-any->response : (-> any/c (or/c false/c response?))
```

`any->response` coerces any value into a response or returns `#f` if coercion is not possible. `any->response` guarantees that any `response?` input must always be returned exactly (i.e. `eq?`.) The default always returns `#f`, signifying that no coercion is possible.

`can-be-response?` returns `#t` if `x` is a response or can be turned into a response by calling `any->response`.

Users of `any->response` should protect themselves by using `can-be-response?` as a contract. If they do so, they can safely ignore the `#f` return case of `any->response`.

`set-any->response!` replaces the global `any->response` with the supplied argument. This function should return the same value for `eq?` inputs to ensure that `can-be-response?` is any accurate predicate. Similarly, this function should be cheap to call multiple times on the same input, since it will be used in contract checking as well as coercion before transmission. You may want to use a weak `eq?`-based hash-table to cache the results for this purpose. (See `make-weak-hasheq`.)

2.3 Web Interaction

```
(require web-server/servlet/web) package: web-server-lib
```

The `web-server/servlet/web` library provides the primary functions of interest for the servlet developer.

```
(send/back response) → void?
  response : can-be-response?
```

Sends *response* to the client. No continuation is captured, so the servlet is done.

Example:

```
(send/back
  (response/xexpr
    `(html
      (body
        (h1 "The sum is: "
          ,(+ first-number
              second-number))))))
```

```
(send/suspend make-response) → request?
make-response : (string? . -> . can-be-response?)
```

Captures the current continuation, stores it with (*current-servlet-continuation-expiration-handler*) as the expiration handler, and binds it to a URL. *make-response* is called with this URL and is expected to generate a *can-be-response?*, which is sent to the client. If the continuation URL is invoked, the captured continuation is invoked and the request is returned from this call to *send/suspend*.

Example:

```
(send/suspend
  (lambda (k-url)
    (response/xexpr
      `(html (head (title "Enter a number"))
        (body
          (form ([action ,k-url])
            "Enter a number: "
            (input ([name "number"])))
            (input ([type "submit"]))))))))))
```

When this form is submitted by the browser, the request will be sent to the URL generated by *send/suspend*. Thus, the request will be “returned” from *send/suspend* to the continuation of this call.

```
(send/suspend/url make-response) → request?
make-response : (url? . -> . can-be-response?)
```

Like *send/suspend* but with a URL struct.

```
(send/suspend/dispatch make-response) → any
make-response : (((request? . -> . any) . -> . string?) . -> . can-be-response?)
```

Calls *make-response* with a function (often named *embed/url*) that, when called with a procedure from *request?* to *any/c* will generate a URL, that when invoked will call the function with the *request?* object and return the result to the caller of *send/suspend/dispatch*. Therefore, if you pass *embed/url* the identity function, *send/suspend/dispatch* devolves into *send/suspend*:

```
(define (send/suspend response-generator)
  (send/suspend/dispatch
    (lambda (embed/url)
      (response-generator (embed/url (lambda (x) x))))))
```

Use *send/suspend/dispatch* when there are multiple ‘logical’ continuations of a page. For example, we could either add to a number or subtract from it:

```
(define (count-dot-com i)
  (count-dot-com
    (send/suspend/dispatch
      (lambda (embed/url)
        (response/xexpr
          `(html
            (head (title "Count!"))
            (body
              (h2 (a ([href
                        ,(embed/url
                          (lambda (req)
                            (sub1 i)))]
                        "-"))
                (h1 ,(number->string i))
                (h2 (a ([href
                          ,(embed/url
                            (lambda (req)
                              (add1 i)))]
                          "+"))))))))))))
```

Notice that in this example the result of the handlers are returned to the continuation of *send/suspend/dispatch*. However, it is very common that the return value of *send/suspend/dispatch* is irrelevant in your application and you may think of it as “embedding” value-less callbacks. Here is the same example in this style:

```
(define (count-dot-com i)
  (send/suspend/dispatch
    (lambda (embed/url)
      (response/xexpr
        `(html
          (head (title "Count!"))
```

```

(body
  (h2 (a ([href
            ,(embed/url
              (lambda (req)
                (count-dot-com (sub1 i))))))])
    "-"))
(h1 ,(number->string i))
(h2 (a ([href
            ,(embed/url
              (lambda (req)
                (count-dot-com (add1 i))))))])
    "+")))))))

```

```

(send/suspend/url/dispatch make-response) → any
make-response : (((request? . -> . any) . -> . url?) . -> . can-be-response?)

```

Like `send/suspend/dispatch`, but with a URL struct.

```

(send/forward make-response) → request?
make-response : (string? . -> . can-be-response?)

```

Calls `clear-continuation-table!`, then `send/suspend`.

Use this if the user can logically go ‘forward’ in your application, but cannot go backward.

```

(send/finish response) → void?
response : can-be-response?

```

Calls `clear-continuation-table!`, then `send/back`.

Use this if the user is truly ‘done’ with your application. For example, it may be used to display the post-logout page:

```

(send/finish
 (response/xexpr
  `(html (head (title "Logged out"))
    (body (p "Thank you for using the services "
      "of the Add Two Numbers, Inc.))))))

```

```

(redirect/get [#:headers hs]) → request?
hs : (listof header?) = empty

```

Calls `send/suspend` with `redirect-to`, passing `hs` as the headers and `see-other` as the redirection status.

This implements the Post-Redirect-Get pattern. Use this to prevent the Refresh button from duplicating effects, such as adding items to a database.

Changed in version 1.3 of package `web-server-lib`: Use `see-other` instead of `temporarily`.

```
(redirect/get/forget [#:headers hs]) → request?
  hs : (listof header?) = empty
```

Like `redirect/get`, but using `send/forward` instead of `send/suspend`.

Changed in version 1.3 of package `web-server-lib`: Use `see-other` instead of `temporarily`, as with `redirect/get`.

```
current-servlet-continuation-expiration-handler
: (parameter/c (or/c false/c
                  (request? . -> . can-be-response?)))
```

Holds the expiration handler to be used when a continuation captured in this context is expired, then looked up.

Example:

```
(parameterize
  ([current-servlet-continuation-expiration-handler
    (lambda (req)
      (response/xexpr
        `(html (head (title "Custom Expiration!")))))]])
  (send/suspend
    ....))
```

```
(clear-continuation-table!) → void?
```

Calls the servlet's manager's `clear-continuation-table!` function. Normally, this deletes all the previously captured continuations.

```
(with-errors-to-browser send/finish-or-back
                          thunk) → any
  send/finish-or-back : (can-be-response? . -> . request?)
  thunk : (-> any)
```

Calls *thunk* with an exception handler that generates an HTML error page and calls *send/finish-or-back*.

Example:

```
(with-errors-to-browser
  send/back
  (lambda ()
    (/ 1 (get-number (request-number)))))
```

```
(adjust-timeout! t) → void?
t : number?
```

Calls the servlet's manager's *adjust-timeout!* function.

Warning: This is deprecated and will be removed in a future release.

```
(continuation-url? u)
→ (or/c false/c (list/c number? number? number?))
u : url?
```

Checks if *u* is a URL that refers to a continuation, if so returns the instance id, continuation id, and nonce.

```
servlet-prompt : continuation-prompt-tag?
```

The tag used for Web interaction continuation capture.

2.4 Web Cells

```
(require web-server/servlet/web-cells)
package: web-server-lib
```

The *web-server/servlet/web-cells* library provides the interface to Web cells.

A Web cell is a kind of state defined relative to the *frame tree*. The frame-tree is a mirror of the user's browsing session. Every time a continuation is invoked, a new frame (called the *current frame*) is created as a child of the current frame when the continuation was captured.

You should use Web cells if you want an effect to be encapsulated in all interactions linked from (in a transitive sense) the HTTP response being generated. For more information on their semantics, consult the paper "Interaction-Safe State for the Web".

```
(web-cell? v) → boolean?  
v : any/c
```

Determines if *v* is a web-cell.

```
(make-web-cell v) → web-cell?  
v : any/c
```

Creates a web-cell with a default value of *v*.

```
(web-cell-ref wc) → any/c  
wc : web-cell?
```

Looks up the value of *wc* found in the nearest frame.

```
(web-cell-shadow wc v) → void  
wc : web-cell?  
v : any/c
```

Binds *wc* to *v* in the current frame, shadowing any other bindings to *wc* in the current frame.

Below is an extended example that demonstrates how Web cells allow the creation of reusable Web abstractions without requiring global transformations of the program into continuation or store passing style.

```
#lang web-server/insta  
  
(define (start initial-request)  
  (define counter1 (make-counter))  
  (define counter2 (make-counter))  
  (define include1 (include-counter counter1))  
  (define include2 (include-counter counter2))  
  (send/suspend/dispatch  
    (lambda (embed/url)  
      (response/xexpr  
        `(html  
          (body (h2 "Double Counters")  
                (div (h3 "First")  
                      ,(include1 embed/url))  
                (div (h3 "Second")  
                      ,(include2 embed/url))))))))))
```

```

(define (make-counter)
  (make-web-cell 0))

(define (include-counter a-counter)
  (call-with-current-continuation
    (λ (k)
      (let loop ()
        (k
          (lambda (embed/url)
            `(div (h3 ,(number->string (web-cell-ref a-counter)))
              (a ([href
                    ,(embed/url
                      (lambda _

                        (define last (web-cell-ref a-counter))

                        (web-cell-shadow a-counter (add1 last))

                        (loop))))])
              "+"))))))
    servlet-prompt))

```

2.5 Continuation Managers

Since Racket servlets store their continuations on the server, they take up memory on the server. Furthermore, garbage collection can not be used to free this memory, because there are roots outside the system: users' browsers, bookmarks, brains, and notebooks. Therefore, some other strategy must be used if memory usage is to be controlled. This functionality is pluggable through the *manager* interface.

2.5.1 General

```

(require web-server/managers/manager)
package: web-server-lib

```

This module defines the manager interface. It is required by the users and implementors of managers.

```
(struct manager (create-instance
  adjust-timeout!
  clear-continuations!
  continuation-store!
  continuation-lookup
  continuation-peek)
  #:extra-constructor-name make-manager)
create-instance : ((-> void) . -> . number?)
adjust-timeout! : (number? number? . -> . void)
clear-continuations! : (number? . -> . void)
continuation-store! : (number? any/c
  (or/c false/c
    (request? . -> . can-be-response?))
  . -> . (list/c number? number?))
continuation-lookup : (number? number? number? . -> . any/c)
continuation-peek : (number? number? number? . -> . any/c)
```

`create-instance` is called to initialize a instance, to hold the continuations of one servlet session. It is passed a function to call when the instance is expired. It runs the id of the instance.

`adjust-timeout!` is a to-be-deprecated function that takes an instance-id and a number. It is specific to the timeout-based manager and will be removed.

`clear-continuations!` expires all the continuations of an instance.

`continuation-store!` is given an instance-id, a continuation value, and a function to include in the exception thrown if the continuation is looked up and has been expired. The two numbers returned are a continuation-id and a nonce.

`continuation-lookup` finds the continuation value associated with the instance-id, continuation-id, and nonce triple it is given.

`continuation-peek` is identical to `continuation-lookup` except that its use must not affect the resource management policy decisions on the instance or continuation accessed. It is intended to be used by debuggers and benchmarks.

```
(struct exn:fail:servlet-manager:no-instance exn:fail
  (expiration-handler)
  #:extra-constructor-name
  make-exn:fail:servlet-manager:no-instance)
expiration-handler : (or/c false/c
  (request? . -> . can-be-response?))
```

This exception should be thrown by a manager when an instance is looked up that does not exist.

```
(struct exn:fail:servlet-manager:no-continuation exn:fail
      (expiration-handler)
      #:extra-constructor-name
      make-exn:fail:servlet-manager:no-continuation)
expiration-handler : (or/c false/c
                    (request? . -> . can-be-response?))
```

This exception should be thrown by a manager when a continuation is looked up that does not exist.

2.5.2 No Continuations

```
(require web-server/managers/none)      package: web-server-lib
```

This module defines a manager constructor:

```
(create-none-manager instance-expiration-handler) → manager?
instance-expiration-handler : (or/c false/c
                              (request? . -> . can-be-response?))
```

This manager does not actually store any continuation or instance data. You could use it if you know your servlet does not use the continuation capturing functions and want the server to not allocate meta-data structures for each instance.

If you *do* use a continuation capturing function, the continuation is simply not stored. If the URL is visited, the *instance-expiration-handler* is called with the request.

If you are considering using this manager, also consider using the Web Language. (See §3 “Stateless Servlets”.)

2.5.3 Timeouts

```
(require web-server/managers/timeouts)
package: web-server-lib
```

This module defines a manager constructor:

```
(create-timeout-manager instance-exp-handler
                       instance-timeout
                       continuation-timeout) → manager?
instance-exp-handler : (or/c false/c
                      (request? . -> . can-be-response?))
```

```

instance-timeout : number?
continuation-timeout : number?

```

Instances managed by this manager will be expired *instance-timeout* seconds after the last time it is accessed. If an expired instance is looked up, the `exn:fail:servlet-manager:no-instance` exception is thrown with *instance-exp-handler* as the expiration handler.

Continuations managed by this manager will be expired *continuation-timeout* seconds after the last time it is accessed. If an expired continuation is looked up, the `exn:fail:servlet-manager:no-continuation` exception is thrown with *instance-exp-handler* as the expiration handler, if no expiration-handler was passed to *continuation-store!*.

adjust-timeout! corresponds to *reset-timer!* on the timer responsible for the servlet instance.

This manager has been found to be... problematic... in large-scale deployments of the Web Server .

2.5.4 LRU

```

(require web-server/managers/lru)      package: web-server-lib

```

This module defines a manager constructor:

```

(create-LRU-manager instance-expiration-handler
                    check-interval
                    collect-interval
                    collect?
                    [#:initial-count initial-count
                     #:inform-p inform-p]) → manager?
instance-expiration-handler : (or/c false/c
                                   (request? . -> . can-be-response?))

check-interval : integer?
collect-interval : integer?
collect? : (-> boolean?)
initial-count : integer? = 1
inform-p : (integer? . -> . void) = (lambda _ (void))

```

Instances managed by this manager will be expired if there are no continuations associated with them, after the instance is unlocked. If an expired instance is looked up, the `exn:fail:servlet-manager:no-instance` exception is thrown with *instance-exp-handler* as the expiration handler.

Continuations managed by this manager are given a "Life Count" of *initial-count* initially. If an expired continuation is looked up, the `exn:fail:servlet-manager:no-continuation` exception is thrown with *instance-exp-handler* as the expiration handler, if no expiration-handler was passed to *continuation-store!*.

Every *check-interval* seconds *collect?* is called to determine if the collection routine should be run. Every *collect-interval* seconds the collection routine is run.

Every time the collection routine runs, the "Life Count" of every continuation is decremented by 1. If a continuation's count reaches 0, it is expired. The *inform-p* function is called if any continuations are expired, with the number of continuations expired.

The recommended usage of this manager is codified as the following function:

```
(make-threshold-LRU-manager instance-expiration-handler
                            memory-threshold)
→ manager?
instance-expiration-handler : (or/c false/c
                                (request? . -> . can-be-response?))
memory-threshold : number?
```

This creates an LRU manager with the following behavior: The memory limit is set to *memory-threshold* bytes. Continuations start with 24 life points. Life points are deducted at the rate of one every 10 minutes, or one every 5 seconds when the memory limit is exceeded. Hence the maximum life time for a continuation is 4 hours, and the minimum is 2 minutes.

If the load on the server spikes—as indicated by memory usage—the server will quickly expire continuations, until the memory is back under control. If the load stays low, it will still efficiently expire old continuations.

3 Stateless Servlets

```
#lang web-server/base      package: web-server-doc
#lang web-server
```

3.1 Example

A stateless servlet should provide the following exports:

```
interface-version : (one-of/c 'stateless)
```

This indicates that the servlet is a stateless servlet.

```
stuffer : (stuffer/c serializable? bytes?)
```

This is the stuffer that will be used for the servlet.

If it is not provided, it defaults to `default-stuffer`.

```
manager : manager?
```

This is the manager that will be used for the servlet.

If it is not provided, it defaults to `(create-none-manager #f)`.

```
(start initial-request) → response?
initial-request : request?
```

This function is called when an instance of this servlet is started. The argument is the HTTP request that initiated the instance.

An example `'stateless` servlet module:

```
#lang web-server
(require web-server/http)
(provide interface-version stuffer start)
(define interface-version 'stateless)
(define stuffer
  (stuffer-chain
    serialize-stuffer
    (md5-stuffer (build-path (find-system-path 'home-dir) ".urls"))))
(define (start req)
  (response/xexpr
    `(html (body (h2 "Look ma, no state!")))))
```

The `web-server/base` language exports all of the functions and syntax from `racket/base` and nothing else.

The `web-server` language exports all of the functions and syntax from the following libraries: `racket`, `net/url`, `web-server/http`, `web-server/http/bindings`, `web-server/lang/abort-resume`, `web-server/lang/web`, `web-server/lang/native`, `web-server/lang/web-param`, `web-server/lang/web-cells`, `web-server/lang/file-box`, `web-server/lang/soft`, `web-server/dispatch`, and `web-server/stuffers`. Some of these are documented in the subsections that follow.

3.2 Usage Considerations

A stateless servlet has the following process performed on it automatically:

- All uses of `letrec` are removed and replaced with equivalent uses of `let` and imperative features.
- The program is converted into ANF (A-Normal Form), making all continuations explicit.
- All continuations and continuation marks are recorded in the continuation marks of the expression they are the continuation of.
- All calls to external modules are identified and marked.
- All uses of `call/cc` are removed and replaced with equivalent gathering of the continuations through the continuation marks installed earlier.
- The program is defunctionalized with a serializable data-structure for each `lambda`.

This process allows the continuations captured by your servlet to be serialized. This means they may be stored on the client's browser or the server's disk.

This means your servlet has no cost to the server other than execution. This is very attractive if you've used Racket servlets and had memory problems.

This means your server can restart in the middle of a long running Web interaction without the URLs that have been shared with the client expiring. This is very attractive if you've used Racket servlets and had session timeout problems.

This process is defined on all of Racket and occurs after macro-expansion, so you are free to use all interesting features of Racket. However, there are some considerations you must make.

First, this process drastically changes the structure of your program. It will create an immense number of lambdas and structures your program did not normally contain. The performance implication of this has not been studied with Racket.

Second, the defunctionalization process is sensitive to the syntactic structure of your program. Therefore, if you change your program in a trivial way, for example, changing a constant, then all serialized continuations will be obsolete and will error when deserialization is attempted. This is a feature, not an error! It is a small price to pay for protection from the sorts of errors that would occur if your program were changed in a meaningful way. If you use the `default-stuffer` or `web-server/stuffers/hash`, then whenever you change your servlet's code, you can safely delete all saved continuations, because they won't be used any longer.

Third, the values in the lexical scope of your continuations must be serializable for the continuations itself to be serializable. This means that you must use `define-serializable-struct` rather than `define-struct`, and take care to use modules that do the same. Similarly, you may not use `parameterize`, because parameterizations are not serializable.

Fourth, and related, this process only runs on your code, not on the code you require. Thus, your continuations—to be serializable—must not be in the context of another module. For example, the following will fail with an "unsafe context" exception:

```
(define requests
  (map (lambda (rg) (send/suspend/url rg))
       response-generators))
```

because `map` is not transformed by the process. However, if you defined your own `map` function, there would be no problem. Another solution is to store the `map` part of the continuation on the server with `serial->native` and `native->serial`:

```
(define requests
  (serial->native
   (map (lambda (rg) (native->serial (send/suspend/url rg)))
        response-generators)))
```

Fifth, the store is **not** serialized. If you rely on the store you will be taking huge risks. You will be assuming that the serialized continuation is invoked on the same server before the server is restarted or the memory is garbage collected.

This process is derived from the papers *Continuations from Generalized Stack Inspection* by Pettyjohn et al. in 2005, *Automatically RESTful Web Applications, Or Marking Modular Serializable Continuations* by Jay McCarthy in 2009, and *The Two-State Solution: Native and Serializable Continuations Accord* by Jay McCarthy in 2010. We thank Greg Pettyjohn for his initial implementation of this algorithm.

3.3 Serializable Continuations

```
(require web-server/lang/abort-resume)
package: web-server-lib
```

The main purpose of the stateless language is to provide serializable continuations to your servlet.

```
(call-with-serializable-current-continuation response-generator)
→ any
response-generator : (continuation? . -> . any)
```

Captures the current continuation in a serializable way and calls *response-generator* with it, returning the result.

This potentially uses resources of the current servlet's *manager* if *serial->native* and *native->serial* were used to capture an untransformable context.

```
(serial->native expr)
```

serial->native informs the serializing runtime that *expr* is potentially a call to an untransformed context. This sets up the necessary information for *native->serial* to signal to *call-with-serializable-current-continuation* to capture the native (and thus unserializable) section of the context and store it on the server.

```
(native->serial expr)
```

native->serial informs the serializing runtime that *expr* marks first expression after returning from an untransformed context. This captures the untransformed context such that *call-with-serializable-current-continuation* can store it on the server and reference it from serializable continuations.

For example,

```
(build-list
  3
  (lambda (i)
    (call-with-serializable-current-continuation
      (lambda (k) (serialize k))))))
```

will fail at runtime because *build-list* is not transformed. However,

```
(serial->native
```

```
(build-list
 3
 (lambda (i)
  (native->serial
   (call-with-serializable-current-continuation
    (lambda (k) (serialize k)))))))
```

will succeed and `k` will reference a cell in the current servlet's `manager` that stores the part of the continuation in `build-list`.

3.4 Native Interfaces

```
(require web-server/lang/native)      package: web-server-lib
```

It is sometimes inconvenient to use `serial->native` and `native->serial` throughout your program. This module provides a macro for creating wrappers.

```
(define-native (native arg-spec ...) original)

arg-spec : ho
arg-spec : _
```

Builds an interface around `original` named `native` such that calls to `native` are wrapped in `serial->native` and all arguments marked with `ho` in `arg-spec` are assumed to be procedures and are wrapped in `native->serial`.

For example,

```
(define-native (build-list/native _ ho) build-list)
```

is equivalent to

```
(define (build-list/native fst snd)
  (serial->native
   (build-list
    fst
    (lambda args
      (native->serial
       (apply snd args)))))))
```

3.5 Stateless Web Interaction

```
(require web-server/lang/web)      package: web-server-lib
```

```
(send/suspend/url response-generator) → request?
response-generator : (url? . -> . response?)
```

Captures the current continuation. Serializes it and stuffs it into a URL. Calls *response-generator* with this URL and delivers the response to the client. If the URL is invoked the request is returned to this continuation.

```
(send/suspend response-generator) → request?
response-generator : (string? . -> . response?)
```

Like `send/suspend/url` but with a string URL representation.

```
(send/suspend/hidden response-generator) → request?
response-generator : (url? xexpr/c . -> . response?)
```

Captures the current continuation. Serializes it and stuffs it into a hidden INPUT form element. Calls *response-generator* with this URL and form field and delivers the response to the client. If the URL is invoked with form data containing the hidden form, the request is returned to this continuation.

```
(send/suspend/url/dispatch make-response) → any
make-response : (((request? . -> . any) . -> . url?) . -> . response?)
```

Calls *make-response* with a function that, when called with a procedure from *request?* to *any/c* will generate a URL, that when invoked will call the function with the *request?* object and return the result to the caller of `send/suspend/dispatch`.

```
(send/suspend/dispatch make-response) → any
make-response : (((request? . -> . any) . -> . string?) . -> . response?)
```

Like `send/suspend/url/dispatch` but with a string URL representation.

```
(redirect/get [#:headers hs]) → request?
hs : (listof header?) = empty
```

See `web-server/servlet/web`.

Changed in version 1.3 of package `web-server-lib`: Added *hs* argument and changed to use *see-other* instead of *temporarily*.

3.6 Stateless Web Cells

```
(require web-server/lang/web-cells)
package: web-server-lib
```

The `web-server/lang/web-cells` library provides the same API as `web-server/servlet/web-cells`, but in a way compatible with the Web Language. The one difference is that `make-web-cell` is syntax, rather than a function.

```
(web-cell? v) → boolean?
  v : any/c
(make-web-cell default-expr)
(web-cell-ref wc) → any/c
  wc : web-cell?
(web-cell-shadow wc v) → void
  wc : web-cell?
  v : any/c
```

See `web-server/servlet/web-cells`.

3.7 File Boxes

```
(require web-server/lang/file-box) package: web-server-lib
```

As mentioned earlier, it is dangerous to rely on the store in Web Language servlets, due to the deployment scenarios available to them. This module provides a simple API to replace boxes in a safe way.

```
(file-box? v) → boolean?
  v : any/c
```

Checks if `v` is a file-box.

```
(file-box p v) → file-box?
  p : path-string?
  v : serializable?
```

Creates a file-box that is stored at `p`, with the default contents of `v`.

```
(file-unbox fb) → serializable?
  fb : file-box?
```

Returns the value inside *fb*

```
(file-box-set? fb) → boolean?  
fb : file-box?
```

Returns `#t` if *fb* contains a value.

```
(file-box-set! fb v) → void  
fb : file-box?  
v : serializable?
```

Saves *v* in the file represented by *fb*.

Warning: If you plan on using a load-balancer, make sure your file-boxes are on a shared medium.

3.8 Stateless Web Parameters

```
(require web-server/lang/web-param)  
package: web-server-lib
```

It is not easy to use `parameterize` in the Web Language. This module provides (roughly) the same functionality in a way that is serializable. Like other serializable things in the Web Language, they are sensitive to source code modification.

```
(make-web-parameter default)
```

Expands to the definition of a web-parameter with *default* as the default value. A web-parameter is a procedure that, when called with zero arguments, returns *default* or the last value web-parameterized in the dynamic context of the call.

```
(web-parameter? v) → boolean?  
v : any/c
```

Checks if *v* appears to be a web-parameter.

```
(web-parameterize ([web-parameter-expr value-expr] ...) expr ...)
```

Runs `(begin expr ...)` such that the web-parameters that the *web-parameter-exprs* evaluate to are bound to the *value-exprs*. From the perspective of the *value-exprs*, this is like `let`.

3.9 Soft State

```
(require web-server/lang/soft)      package: web-server-lib
```

Sometimes you want to reference a large data-structure from a stateless program without the data-structure being serialized and increasing the size of the serialization. This module provides support for this scenario.

```
(soft-state? v) → boolean?  
  v : any/c
```

Determines if *v* is a soft state record.

```
(make-soft-state thunk) → soft-state?  
  thunk : (-> any/c)
```

Creates a piece of soft state that is computed by *thunk*. This value is serializable.

```
(soft-state-ref ss) → any/c  
  ss : soft-state?
```

Extracts the value associated with *ss*. If the value is not available (perhaps because of garbage collection, deserialization in an uninitialized process, etc), then the thunk associated with *ss* is invoked and the value is cached.

```
(soft-state expr ...)
```

Equivalent to `(make-soft-state (lambda () expr ...))`.

Here's an example servlet that uses soft state:

```
#lang web-server

(provide interface-version start)
(define interface-version 'stateless)

(define softie
  (soft-state
    (printf "Doing a long computation...\n")
    (sleep 1)))

(define (start req)
```

```

(soft-state-ref softie)
(sprintf "Done\n")
(start
  (send/suspend
    (lambda (k-url)
      (response/xexpr
        `(html (body (a ([href ,k-url]) "Done"))))))))

```

When this is run and the link is clicked a few times, the output is:

```

$ plt-web-server -p 8080
Doing a long computation...
Done
Done
Done
Done

```

If the server is restarted or the hostname in the URL is changed to a different host with the same code, and the URL is clicked:

```

^Cuser break
$ plt-web-server -p 8080
Doing a long computation...
Done

```

3.10 Stuffers

```

(require web-server/stuffers)      package: web-server-lib

```

The `web-server` language provides serializable continuations. The serialization functionality is abstracted into *stuffers* that control how it operates. You can supply your own (built with these functions) when you write a stateless servlet.

3.10.1 Basic Combinators

```

(require web-server/stuffers/stuffer)
package: web-server-lib

```

```

(struct stuffer (in out)
  #:extra-constructor-name make-stuffer)
in : (any/c . -> . any/c)

```

```
out : (any/c . -> . any/c)
```

A stuffer is essentially an invertible function captured in this structure. The following should hold:

```
(out (in x)) = x  
(in (out x)) = x
```

```
(stuffer/c dom rng) → contract?  
  dom : any/c  
  rng : any/c
```

Constructs a contract for a stuffer where `in` has the contract `(-> dom rng)` and `out` has the contract `(-> rng dom)`.

```
id-stuffer : (stuffer/c any/c any/c)
```

The identity stuffer.

```
(stuffer-compose g f) → (stuffer/c any/c any/c)  
  g : (stuffer/c any/c any/c)  
  f : (stuffer/c any/c any/c)
```

Composes `f` and `g`, i.e., applies `f` then `g` for `in` and `g` then `f` for `out`.

```
(stuffer-sequence f g) → (stuffer/c any/c any/c)  
  f : (stuffer/c any/c any/c)  
  g : (stuffer/c any/c any/c)
```

`stuffer-compose` with arguments swapped.

```
(stuffer-if c f) → (stuffer/c bytes? bytes?)  
  c : (bytes? . -> . boolean?)  
  f : (stuffer/c bytes? bytes?)
```

Creates a stuffer that stuffs with `f` if `c` is true on the input to `in`. Similarly, applies `f` during `out` if it was applied during `in` (which is recorded by prepending a byte.)

```
(stuffer-chain x ...) → stuffer?  
  x : (or/c stuffer? (bytes? . -> . boolean?))
```

Applies `stuffer-sequence` and `stuffer-if` to successive tails of `x`.

3.10.2 Serialization

```
(require web-server/stuffers/serialize)  
package: web-server-lib
```

| `serialize-stuffer` : (`stuffer/c serializable? bytes?`)

A stuffer that uses `serialize` and `write/bytes` and `deserialize` and `read/bytes`.

3.10.3 Base64 Encoding

```
(require web-server/stuffers/base64)  
package: web-server-lib
```

| `base64-stuffer` : (`stuffer/c bytes? bytes?`)

A stuffer that uses `base64-encode` and `base64-decode`.

Useful for getting URL-safe bytes.

3.10.4 GZip Compression

```
(require web-server/stuffers/gzip) package: web-server-lib
```

| `gzip-stuffer` : (`stuffer/c bytes? bytes?`)

A stuffer that uses `gzip/bytes` and `gunzip/bytes`.

Warning: You should compose this with `base64-stuffer` to get URL-safe bytes.

3.10.5 Key/Value Storage

The `web-server/stuffers/hash` stuffers rely on a key/value store.

```
(require web-server/stuffers/store)  
package: web-server-lib
```

```
(struct store (write read)
  #:extra-constructor-name make-store)
write : (bytes? bytes? . -> . void)
read : (bytes? . -> . bytes?)
```

The following should hold:

```
(begin (write k v) (read k)) = v
```

```
(dir-store root) → store?
root : path-string?
```

A store that stores key `key`'s value in a file located at

```
(build-path
 root
 (bytes->string/utf-8 key))
```

It should be easy to use this interface to create store for databases like SQLite, CouchDB, or BerkeleyDB.

3.10.6 Hash-addressed Storage

```
(require web-server/stuffers/hash)      package: web-server-lib
```

```
hash-fun/c : contract?
```

Equivalent to `(-> bytes? bytes?)`.

```
(hash-stuffer H store) → (stuffer/c bytes? bytes?)
H : hash-fun/c
store : store?
```

A content-addressed storage stuffer that stores input bytes, `input`, in `store` with the key `(H input)` and returns the key. Similarly, on `out` the original bytes are looked up.

```
(md5-stuffer root) → (stuffer/c bytes? bytes?)
root : path-string?
```

Equivalent to `(hash-stuffer md5 (dir-store root))`

3.10.7 HMAC-SHA1 Signing

```
(require web-server/stuffers/hmac-sha1)
package: web-server-lib
```

```
(HMAC-SHA1 kb db) → bytes?
  kb : bytes?
  db : bytes?
```

Performs a HMAC-SHA1 calculation on *db* using *kb* as the key. The result is guaranteed to be 20 bytes. (You could curry this to use it with [hash-stuffer](#), but there is little value in doing so over [md5](#).)

```
(HMAC-SHA1-stuffer kb) → (stuffer/c bytes? bytes?)
  kb : bytes?
```

A stuffer that signs input using [HMAC-SHA1](#) with *kb* as the key. The result of the stuffer is the hash prepended to the input data. When the stuffer is run in reverse, it checks if the first 20 bytes are the correct has for the rest of the data.

Warning: You should compose this with [base64-stuffer](#) to get URL-safe bytes.

Warning: Without explicit provision, it is possible for users to modify the continuations they are sent through the other stuffers. This stuffer allows the servlet to certify that stuffed data was truly generated by the servlet. Therefore, you **should** use this if you are not using the [hash-stuffers](#).

Warning: This stuffer does **not** encrypt the data in anyway, so users can still observe the stuffed values.

3.10.8 Helpers

```
(require web-server/lang/stuff-url)
package: web-server-lib
```

```
(is-url-too-big? v) → boolean?
  v : bytes?
```

Determines if stuffing *v* into the current servlet's URL would result in a URL that is too big for Internet Explorer. (IE only supports URLs up to 2048 characters.)

```
(make-default-stuffer root) → (stuffer/c serializable? bytes?)  
  root : path-string?
```

Constructs a stuffer that serializes, then if the URL is too big, compresses (and base64-encodes), if the URL is still too big then it stores it in an MD5-indexed database rooted at *root*.

Equivalent to:

```
(stuffer-chain  
  serialize-stuffer  
  is-url-too-big?  
  (stuffer-chain  
    gzip-stuffer  
    base64-stuffer)  
  is-url-too-big?  
  (md5-stuffer root))
```

```
default-stuffer : (stuffer/c serializable? bytes?)
```

Equivalent to:

```
(make-default-stuffer  
  (build-path  
    (find-system-path 'home-dir)  
    ".urls"))
```

4 HTTP: Hypertext Transfer Protocol

```
(require web-server/http)      package: web-server-lib
```

The Web Server implements many HTTP libraries that are provided by this module.

4.1 Requests

```
(require web-server/http/request-structs)
package: web-server-lib
```

```
(struct header (field value)
  #:extra-constructor-name make-header)
field : bytes?
value : bytes?
```

Represents a header of `field` to `value`.

```
(headers-assq id heads) → (or/c false/c header?)
id : bytes?
heads : (listof header?)
```

Returns the header with a field equal to `id` from `heads` or `#f`.

```
(headers-assq* id heads) → (or/c false/c header?)
id : bytes?
heads : (listof header?)
```

Returns the header with a field case-insensitively equal to `id` from `heads` or `#f`.

You almost **always** want to use this, rather than `headers-assq` because Web browsers may send headers with arbitrary casing.

```
(struct binding (id)
  #:extra-constructor-name make-binding)
id : bytes?
```

Represents a binding of `id`.

```
(struct binding:form binding (value)
  #:extra-constructor-name make-binding:form)
value : bytes?
```

Represents a form binding of `id` to `value`.

```
(struct binding:file binding (filename headers content)
  #:extra-constructor-name make-binding:file)
filename : bytes?
headers : (listof header?)
content : bytes?
```

Represents the uploading of the file `filename` with the id `id` and the content `content`, where `headers` are the additional headers from the MIME envelope the file was in. For example, the `#"Content-Type"` header may be included by some browsers.

See also `binding:file/port-in`, an alternative interface to file uploads that can be significantly more memory efficient.

Changed in version 1.6 of package `web-server-lib`: Extended to support a port-based representation: see `binding:file/port-in`.

```
(binding:file/port-in binding) → input-port?
binding : binding:file/port?
(binding:file/port? v) → boolean?
v : any/c
(make-binding:file/port id
  filename
  headers
  content-in) → binding:file/port?

id : bytes?
filename : bytes?
headers : (listof header?)
content-in : input-port?
```

The web server can avoid storing uploaded files in memory. In particular, a safety limits value can instruct this library to offload files to disk if they are larger than some threshold. Even for file uploads that are not written to disk, the web server initially places the content in an input port, rather than a byte-string, so that storage need not be retained after the content has been read.

The port-based interface is exposed to programmers, and it can be significantly more memory efficient than the byte-string-based interface. However, the port-based interface is state-

ful: programmers who use it take responsibility for managing the state of the input port. Read on for details.

To maintain compatability, the port-based interface uses a private, opaque subtype of `binding:file`. Instances of this extended type are recognized by the predicate `binding:file/port?` and are created using `make-binding:file/port`, which is like `make-binding:file`, but takes the file content as an input port rather than a byte string. Only `binding:file` instances recognized by `binding:file/port?` support `binding:file/port-in`. The web server uses `make-binding:file/port` when reading `request` structures, which is the primary way most programs encounter `binding:file` instances: however, deserialized instances (see below) and instances constructed manually using `make-binding:file` do not support the port-based API.

It is important to be aware of how `binding:file-content` works with port-based instances. The first time `binding:file-content` is called on a port-based instance `v`, it consumes the port's remaining content as with `(port->bytes (binding:file/port-in v))`, memoizes the result for future calls to `binding:file-content`, and closes the input port. This behavior means that:

- A given byte of input may be either stored in the `binding:file-content` field or read directly by from the input port, but never both; and
- If the input port has already been closed directly when `binding:file-content` is called for the first time, `binding:file-content` will raise an exception.

Accessing the `binding:file-content` field indirectly, such as by using `match`, has the same behavior as calling `binding:file-content` directly. In particular, calling `serialize` on a `binding:file` instance implicitly uses `binding:file-content`, and deserialized instances are effectively constructed using `make-binding:file`.

Added in version 1.6 of package `web-server-lib`.

```
(bindings-assq id binds) → (or/c false/c binding?)
  id : bytes?
  binds : (listof binding?)
```

Returns the binding with an id equal to `id` from `binds` or `#f`.

```
(bindings-assq-all id binds) → (listof binding?)
  id : bytes?
  binds : (listof binding?)
```

Like `bindings-assq`, but returns a list of all bindings matching `id`.

```
(struct request (method
                  uri
                  headers/raw
                  bindings/raw-promise
                  post-data/raw
                  host-ip
                  host-port
                  client-ip)
  #:extra-constructor-name make-request)
method : bytes?
uri : url?
headers/raw : (listof header?)
bindings/raw-promise : (promise/c (listof binding?))
post-data/raw : (or/c #f bytes?)
host-ip : string?
host-port : number?
client-ip : string?
```

An HTTP `method` request to `uri` from `client-ip` to the server at `host-ip:host-port` with `headers/raw` headers, `bindings/raw` GET and POST queries and `post-data/raw` POST data.

You are **unlikely to need to construct** a request struct.

Changed in version 1.6 of package `web-server-lib`: Fixed to answer `#f` to `serializable?`, as all `request` instances contain non-serializable pieces.

```
(request-bindings/raw r) → (listof binding?)
r : request?
```

Forces `(request-bindings/raw-promise r)`.

Here is an example typical of what you will find in many applications:

```
(define (get-number req)
  (match
    (bindings-assq
      #"number"
      (request-bindings/raw req))
    [(? binding:form? b)
     (string->number
      (bytes->string/utf-8
       (binding:form-value b)))]
    [_
     (get-number (request-number))]))
```

4.2 Bindings

```
(require web-server/http/bindings)      package: web-server-lib
```

These functions, while convenient, could introduce subtle errors into your application. Examples: that they are case-insensitive could introduce an error; if the data submitted is not in UTF-8 format, then the conversion to a string will fail; if an attacker submits a form field as if it were a file, when it is not, then the `request-bindings` will hold a `bytes?` object and your program will error; and, for file uploads you lose the filename. **Therefore, we recommend against their use, but they are provided for compatibility with old code.**

```
(request-bindings req)
→ (listof (or/c (cons/c symbol? string?)
                (cons/c symbol? bytes?)))
req : request?
```

Translates the `request-bindings/raw` of `req` by interpreting `bytes?` as `string?`s, except in the case of `binding:file` bindings, whose contents are returned as bytes. Ids are then translated into lowercase symbols.

```
(request-headers req) → (listof (cons/c symbol? string?))
req : request?
```

Translates the `request-headers/raw` of `req` by interpreting `bytes?` as `string?`s. Ids are then translated into lowercase symbols.

```
(extract-binding/single id binds) → string?
id : symbol?
binds : (listof (cons/c symbol? string?))
```

Returns the single binding associated with `id` in the a-list `binds` if there is exactly one binding. Otherwise raises `exn:fail`.

```
(extract-bindings id binds) → (listof string?)
id : symbol?
binds : (listof (cons/c symbol? string?))
```

Returns a list of all the bindings of `id` in the a-list `binds`.

```
(exists-binding? id binds) → boolean?
id : symbol?
binds : (listof (cons/c symbol? string?))
```

Returns `#t` if `binds` contains a binding for `id`. Otherwise, `#f`.

Here is an example typical of what you will find in many applications:

```
(define (get-number req)
  (string->number
    (extract-binding/single
      'number
      (request-bindings req))))
```

4.3 Responses

```
(require web-server/http/response-structs)
package: web-server-lib
```

```
(struct response (code message seconds mime headers output))
  code : response-code/c
  message : bytes?
  seconds : real?
  mime : (or/c #f bytes?)
  headers : (listof header?)
  output : (output-port? . -> . any)
response-code/c : flat-contract? = (integer-in 100 999)
```

An HTTP response where `output` produces the body by writing to the output port. `code` is the response code, `message` the message, `seconds` the generation time, `mime` the MIME type of the file, and `headers` are the headers.

If `headers` does not include `Date`, `Last-Modified`, or `Server` headers, then the server will automatically add them, where `Date` is based on `current-seconds`, `Last-Modified` is based on `seconds`, and `Server` is `Racket`.

If `headers` does not include `Content-Type` and `mime` is not `#f`, then `mime` is added as a `Content-Type` header.

The server will always replace your `Connection` header if it needs to ensure the connection will be closed. (Typically with an HTTP/1.0 client.)

The server will always put headers it generates before those in the `response` structure and guarantees that the headers supplied appear in the output in the order given. (This is relevant if multiple occurrences of the same header have a different interpretation by the client, such as with `Set-Cookie`.)

Examples:

```

(response
  301 #"OK"
  (current-seconds) TEXT/HTML-MIME-TYPE
  empty
  (λ (op) (write-bytes #"<html><body>Hello,
World!</body></html>" op)))
(response
  301 #"Moved Permanently"
  (current-seconds) TEXT/HTML-MIME-TYPE
  (list (make-header #"Location"
                    #"http://racket-lang.org/download"))
  (λ (op) (write-bytes #"Moved" op)))
(response
  304 #"Not Modified"
  (current-seconds) #f
  (list (make-header #"Location"
                    #"http://racket-lang.org/download"))
  void)

```

Changed in version 1.2 of package `web-server-lib`: Contract on `output` weakened to allow `any` as the result (instead of demanding `void?`).

Changed in version 1.3: Added `response-code/c` and made the contracts on `code` and `seconds` stronger (rather than accepting `number?`).

```

(response/full code
               message
               seconds
               mime
               headers
               body) → response?
code : response-code/c
message : (or/c #f bytes?)
seconds : real?
mime : (or/c #f bytes?)
headers : (listof header?)
body : (listof bytes?)

```

A constructor for responses where `body` is the response body.

Example:

```

(response/full
  301 #"Moved Permanently"
  (current-seconds) TEXT/HTML-MIME-TYPE
  (list (make-header #"Location"

```

```

                                #"http://racket-lang.org/download"))
(list #"<html><body><p>"
      #"Please go to <a href=\"\"
      #"http://racket-lang.org/download"
      #"\">here</a> instead."
      #"</p></body></html>"))

```

If *message* is not supplied or is *#f*, a status message will be inferred based on *code*. Status messages will be inferred based on RFCs 7231 (“Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content”) and 7235 (“Hypertext Transfer Protocol (HTTP/1.1): Authentication”). These are the following:

Code Message

100	Continue
101	Switching Protocols
200	OK
201	Created
202	Accepted
203	Non-Authoritative Information
204	No Content
205	Reset Content
206	Partial Content
300	Multiple Choices
301	Moved Permanently
302	Found
303	See Other
305	Use Proxy
307	Temporary Redirect
308	Permanent Redirect
400	Bad Request
401	Unauthorized
402	Payment Required
403	Forbidden
404	Not Found
405	Method Not Allowed
406	Not Acceptable
407	Proxy Authentication Required
408	Request Timeout
409	Conflict
410	Gone
411	Length Required
413	Payload Too Large
414	URI Too Long
415	Unsupported Media Type
417	Expectation Failed
426	Upgrade Required

500 Internal Server Error
 501 Not Implemented
 502 Bad Gateway
 503 Service Unavailable
 504 Gateway Timeout
 505 HTTP Version Not Supported

Changed in version 1.3 of package `web-server-lib`: Updated contracts on `code` and `seconds` as with `response`.

Changed in version 1.4 of package `web-server-lib`: Contract on `message` relaxed to allow both `#f` and a `bytes?`, with a default of `#f`. Previously, `bytes?` was required, and had a default of `#"Okay"`.

```

(response/output output
  [#:code code
   #:message message
   #:seconds seconds
   #:mime-type mime-type
   #:headers headers]) → response?

output : (-> output-port? any)
code : number? = 200
message : (or/c false/c bytes?) = #f
seconds : number? = (current-seconds)
mime-type : (or/c bytes? #f) = TEXT/HTML-MIME-TYPE
headers : (listof header?) = '()
  
```

Equivalent to

```
(response code message seconds mime-type headers output)
```

, with the understanding that if `message` is missing, it will be inferred from `code` using the association between status codes and messages found in RFCs 7231 and 7235. See the documentation for `response/full` for the table of built-in status codes.

Changed in version 1.2 of package `web-server-lib`: Contract on `output` weakened to allow `any` as the result (instead of demanding `void?`).

Changed in version 1.3: Updated contracts on `code` and `seconds` as with `response`.

Changed in version 1.4: Contract on `message` relaxed to allow both `#f` and a `bytes?`, with a default of `#f`. Previously, `bytes?` was required, and had a default of `#"Okay"`.

```
TEXT/HTML-MIME-TYPE : bytes?
```

Equivalent to `#"text/html; charset=utf-8"`.

`APPLICATION/JSON-MIME-TYPE : bytes?`

Equivalent to `"application/json; charset=utf-8"`.

Warning: If you include a Content-Length header in a response that is inaccurate, there **will be an error** in transmission that the server **will not catch**.

4.4 Placing Cookies

`(require web-server/http/cookie)` package: web-server-lib

This module provides functions to create cookies and responses that set them.

```
(make-cookie name
             value
             [#:comment comment
              #:domain domain
              #:max-age max-age
              #:path path
              #:expires expires
              #:secure? secure?
              #:http-only? http-only?
              #:extension extension]) → cookie?

name : cookie-name?
value : cookie-value?
comment : any/c = #f
domain : (or/c domain-value? #f) = #f
max-age : (or/c (and/c integer? positive?) #f) = #f
path : (or/c path/extension-value? #f) = #f
expires : (or/c date? string? #f) = #f
secure? : any/c = #f
http-only? : any/c = #f
extension : (or/c path/extension-value? #f) = #f
```

Constructs a cookie with the appropriate fields.

This is a wrapper around `make-cookie` from `net/cookies/server` for backwards compatibility. The `comment` argument is ignored. If `expires` is given as a string, it should match RFC 7231, Section 7.1.1.2, in which case it will be converted to a `date?` value. If conversion fails, an `exn:fail:contract?` is raised.

Changed in version 1.3 of package `web-server-lib`: Added support for RFC 6265 via `net/cookies/server`. Enforce stronger contracts on string-valued arguments. Allow `expires` to be a `date?` and allow `secure` to be

`any/c` (rather than `boolean?`). Forbid 0 for `max-age`. Support `http-only?` and `extension` arguments. Ignore `comment`.

```
(cookie->header c) → header?  
c : cookie?
```

Constructs a header that sets the cookie.

Examples:

```
(define time-cookie  
  (make-cookie "time" (number->string (current-seconds))))  
(define id-cookie  
  (make-cookie "id" "joseph" #:secure? #t))  
  
(redirect-to  
  "http://localhost/logged-in"  
  see-other  
  #:headers  
  (map cookie->header  
    (list time-cookie id-cookie)))  
  
(send/suspend  
  (lambda (k-url)  
    (response/xexpr  
      #:cookies (list time-cookie id-cookie)  
      `(html (head (title "Cookie Example"))  
              (body (h1 "You're cookie'd!"))))))
```

4.5 Authenticated Cookies

```
(require web-server/http/id-cookie)  
package: web-server-lib
```

Cookies are useful for storing information of user's browsers and particularly useful for storing identifying information for authentication, sessions, etc. However, there are inherent difficulties when using cookies as authenticators, because cookie data is fully controlled by the user, and thus cannot be trusted.

This module provides functions for creating and verifying authenticated cookies that are intrinsically timestamped. It is based on the algorithm proposed by the MIT Cookie Eaters: if you store the data `data` at time `authored-seconds`, then the user will receive `digest&authored-seconds&data`, where `digest` is an HMAC-SHA1 digest of `authored-seconds` and `data`, using an arbitrary secret key. When you receive a cookie,

it will reverify this digest and check that the cookie's `authored-seconds` is not after a timeout period, and only then return the cookie data to the program.

The interface represents the secret key as a byte string. **For security, this should be created using cryptographic-quality randomness.** A convenient purely Racket-based option is `make-secret-salt/file`, which is implemented using `crypto-random-bytes`. You can also generate random bytes using something like OpenSSL or `/dev/random`: this FAQ lists a few options.

```
(make-id-cookie name
                value
                #:key secret-salt
                [#:path path
                 #:expires expires
                 #:max-age max-age
                 #:domain domain
                 #:secure? secure?
                 #:http-only? http-only?
                 #:extension extension]) → cookie?

name : cookie-name?
value : cookie-value?
secret-salt : bytes?
path : (or/c path/extension-value? #f) = #f
expires : (or/c date? #f) = #f
max-age : (or/c (and/c integer? positive?) #f) = #f
domain : (or/c domain-value? #f) = #f
secure? : any/c = #f
http-only? : any/c = #f
extension : (or/c path/extension-value? #f) = #f

(make-id-cookie name
                secret-salt
                value
                [#:path path
                 #:expires expires
                 #:max-age max-age
                 #:domain domain
                 #:secure? secure?
                 #:http-only? http-only?
                 #:extension extension]) → cookie?

name : cookie-name?
secret-salt : bytes?
value : cookie-value?
path : (or/c path/extension-value? #f) = #f
expires : (or/c date? #f) = #f
max-age : (or/c (and/c integer? positive?) #f) = #f
domain : (or/c domain-value? #f) = #f
```

```

secure? : any/c = #f
http-only? : any/c = #t
extension : (or/c path/extension-value? #f) = #f

```

Generates an authenticated cookie named *name* containing *value*, signed with *secret-salt*.

The calling conventions allow *secret-salt* to be given either as a keyword argument (mirroring the style of *make-cookie*) or a by-position argument (for compatibility with older versions of this library).

The other arguments are passed to *make-cookie*; however, note that the default value for *http-only?* is *#t*. Users will also likely want to set *secure?* to *#t* when using HTTPS.

Changed in version 1.3 of package *web-server-lib*: Added support for RFC 6265 as with *make-cookie*, including adding the optional arguments *expires*, *max-age*, *domain*, *secure*, *extension*, and *http-only?* (which is *#true* by default). Allowed *secret-salt* to be given with the keyword *#:key* instead of by position.

Changed in version 1.6: Changed to accept any *cookie-name?* or *cookie-value?* (rather than only strings) for the *name* and *value* arguments, respectively, for consistency with *make-cookie*. Fixed a bug that had incorrectly truncated cookie signatures: note that previous versions of this library will not recognize cookies created by the fixed *make-id-cookie* as validly signed, and vice versa.

```

(request-id-cookie request
  #:name name
  #:key secret-salt
  [#:timeout timeout
   #:shelf-life shelf-life])
→ (or/c #f (and/c string? cookie-value?))
request : request?
name : cookie-name?
secret-salt : bytes?
timeout : real? = +inf.0
shelf-life : real? = +inf.0
(request-id-cookie name
  secret-salt
  request
  [#:timeout timeout
   #:shelf-life shelf-life])
→ (or/c #f (and/c string? cookie-value?))
name : cookie-name?
secret-salt : bytes?
request : request?
timeout : real? = +inf.0
shelf-life : real? = +inf.0

```

Extracts the first authenticated cookie named *name* that was previously signed with *secret-salt* from *request*, with the allowable age of the cookie is controlled by *shelf-life* and *timeout* as with *valid-id-cookie?*.

If no valid cookie is available, returns *#f*.

Changed in version 1.3 of package *web-server-lib*: Added *shelf-life* argument and support for giving *name* and *secret-salt* by keyword instead of by position. Added support for RFC 6265 as with *make-cookie*.

Changed in version 1.6: Changed *name* argument to accept any *cookie-name?* as with *make-id-cookie*. Corrected the documented contract for the *timeout* argument. Fixed a bug that had incorrectly truncated cookie signatures: note that the fixed *request-id-cookie* will reject cookies created by previous versions of this library, and vice versa.

```
(valid-id-cookie? cookie
  #:name name
  #:key secret-salt
  [#:timeout timeout
   #:shelf-life shelf-life])
→ (or/c #f (and/c string? cookie-value?))
cookie : any/c
name : cookie-name?
secret-salt : bytes?
timeout : real? = +inf.0
shelf-life : real? = +inf.0
```

Recognizes authenticated cookies named *name* that were previously signed with *secret-salt*. Values satisfying either *cookie?* or *client-cookie?* can be recognized.

The *shelf-life* specifies the maximum age of the cookie in seconds. Cookies created more than *shelf-life* seconds ago will not be considered valid. The default value, *+inf.0*, permits all properly named and signed cookies.

Counterintuitively, the *timeout* argument requires that the cookie have been created *before* a certain moment in time: in other words, it requires that the cookie be *older* than a certain age. This is not usually what you want to restrict. Specifically, *valid-id-cookie?* tests that $(\leq \text{authored } \text{timeout})$, where *authored* is the value returned by *(current-seconds)* when the cookie was created. The default value, *+inf.0*, permits all properly named and signed cookies.

Added in version 1.3 of package *web-server-lib*.

Changed in version 1.6: Changed *name* argument to accept any *cookie-name?* as with *make-id-cookie*. Corrected the documented contract for the *timeout* argument. Fixed a bug that had incorrectly truncated cookie signatures: note that the fixed *valid-id-cookie?* will reject cookies created by previous versions of this library, and vice versa.

```
(logout-id-cookie name
  [#:path path
   #:domain domain]) → cookie?
name : cookie-name?
path : (or/c #f string?) = #f
domain : (or/c #f domain-value? #f) = #f
```

Generates a cookie named *name* that is not validly authenticated and expires in the past.

This will cause non-malicious browsers to overwrite a previously set cookie. If you use authenticated cookies for login information, you could send this to cause a “logout.” However, malicious browsers do not need to respect such an overwrite. Therefore, this is not an effective way to implement timeouts or protect users on public (i.e. possibly compromised) computers. The only way to securely logout on the compromised computer is to have server-side state keeping track of which cookies (sessions, etc.) are invalid. Depending on your application, it may be better to track live sessions or dead sessions, or never set cookies to begin with and just use (stateful) continuations, which you can revoke with [send/finish](#).

Changed in version 1.3 of package `web-server-lib`: Added support for RFC 6265 as with [make-cookie](#), including adding the *domain* argument.

Changed in version 1.6: Fixed to accept any [cookie-name?](#) for the *name* argument, as was previously documented.

```
(make-secret-salt/file secret-salt-path) → bytes?
secret-salt-path : path-string?
```

Extracts the bytes from *secret-salt-path*. If *secret-salt-path* does not exist, then it is created and initialized with 128 cryptographic-quality random bytes from [crypto-random-bytes](#).

Changed in version 1.3 of package `web-server-lib`: Changed to use cryptographic-quality randomness to initialize *secret-salt-path*.

4.6 Extracting Cookies

```
(require web-server/http/cookie-parse)
package: web-server-lib
```

```
(struct client-cookie (name value domain path)
  #:extra-constructor-name make-client-cookie)
name : (and/c string? cookie-name?)
value : (and/c string? cookie-value?)
domain : (or/c #f domain-value?)
```

```
path : (or/c #f path/extension-value?)
```

While server cookies are represented with `cookie?`s, cookies that come from the client are represented with a `client-cookie` structure.

```
(request-cookies req) → (listof client-cookie?)  
req : request?
```

Extracts the cookies from `req`'s headers.

Changed in version 1.3 of package `web-server-lib`: Added support for RFC 6265 via `net/cookies/common`.

Examples:

```
(define (start req)  
  (define cookies (request-cookies req))  
  (define id-cookie  
    (findf (lambda (c)  
            (string=? "id" (client-cookie-name c)))  
           cookies))  
  (if id-cookie  
      (hello (client-cookie-value id-cookie))  
      (redirect-to  
        (url->string (request-uri req))  
        see-other  
        #:headers  
        (list  
         (cookie->header (make-cookie "id" "joseph"))))))  
  
(define (hello who)  
  (response/xexpr  
    `(html (head (title "Hello!"))  
          (body  
            (h1 "Hello "  
              ,who))))))
```

4.7 Redirect

```
(require web-server/http/redirect)      package: web-server-lib
```

```
(redirect-to uri [status #:headers headers]) → response?  
uri : non-empty-string?  
status : redirection-status? = temporarily  
headers : (listof header?) = '()
```

```
(redirection-status? v) → boolean?
  v : any/c
temporarily : redirection-status?
temporarily/same-method : redirection-status?
see-other : redirection-status?
permanently : redirection-status?
```

The function `redirect-to` generates an HTTP response that redirects the browser to `uri`, while including the `headers` in the response. The `status` argument is a *redirection status* value, which determines the specific type of HTTP redirect to be used.

The default redirection status, `temporarily`, is preserved for backwards compatibility: new code should usually use either `temporarily/same-method` or `see-other`, instead. The `temporarily` redirection status corresponds to `302 Found`. Unfortunately, browsers have not implemented this status consistently for methods other than `GET` and (in practice, with all but some very old browsers) `POST`.

The `temporarily/same-method` redirection status uses `307 Temporary Redirect`. This redirects the browser to `uri` using the same HTTP method as the original request.

The `see-other` redirection status corresponds to `303 See Other`. It is most often used to implement the *Post-Redirect-Get* pattern: as a response to a request using `POST` or another HTTP method with side-effects, it causes the browser to perform a `GET` or `HEAD` request for `uri`, which gives a response to the original `POST` request. This prevents the Back and Refresh buttons from duplicating effects, such as making a purchase or adding items to a database. The web server provides `redirect/get` for added convenience with Post-Redirect-Get.

The `permanently` redirection status uses the HTTP status `301 Moved Permanently`. It is like `temporarily`, except that, as the name suggests, it signifies that the move is permanent and that search engines, for example, should use `uri` instead of the URI of the original request. Unfortunately, `permanently` is also like `temporarily` in that browsers have implemented it inconsistently for methods other than `GET` and `HEAD`: in particular, RFC 7231 permits that, “for historical reasons, a user agent **may** change the request method from `POST` to `GET` for the subsequent request. When it is important to ensure that the request to `uri` use the same method, there are some possible alternatives:

- RFC 7231 suggests using `307 Temporary Redirect`, i.e. `temporarily/same-method`. This has the disadvantage that search engines and others won’t update references to the old URI.
- RFC 7538 specifies a new HTTP status, `308 Permanent Redirect`, which forbids changing the request method, analogously to `307 Temporary Redirect`. However, the RFC also highlights some important deployment considerations for this status. In particular, older browsers—including, as of this writing, some that remain in relatively common use—do not understand this status and will fall back to the semantics of

300 Multiple Choices, which is often undesirable.

- The application can note the method of the original request and use `permanently` for `GET` and `HEAD` requests or one of the other alternatives for other methods.

Example: `(redirect-to "http://www.add-three-numbers.com" permanently)`

Changed in version 1.3 of package `web-server-lib`: Added `temporarily/same-method`.

4.8 Basic Authentication

```
(require web-server/http/basic-auth)
package: web-server-lib
```

An implementation of HTTP Basic Authentication.

```
(make-basic-auth-header realm) → header?
realm : string?
```

Returns a header that instructs the Web browser to request a username and password from the client using Basic authentication with `realm` as the realm.

```
(request->basic-credentials req)
→ (or/c false/c (cons/c bytes? bytes?))
req : request?
```

Returns a pair of the username and password from the authentication header in `req` if they are present, or `#f`.

Example:

```
#lang web-server/insta

(define (start req)
  (match (request->basic-credentials req)
    [(cons user pass)
     (response/xexpr
      `(html (head (title "Basic Auth Test"))
              (body (h1 "User: " ,(bytes->string/utf-8 user))
                    (h1 "Pass: " ,(bytes->string/utf-8 pass))))))]
    [else
     (response
      401 #"Unauthorized" (current-seconds) TEXT/HTML-MIME-TYPE)])
```

```
(list
  (make-basic-auth-header
    (format "Basic Auth Test: ~a" (gensym))))
void]]))
```

4.9 Digest Authentication

```
(require web-server/http/digest-auth)
package: web-server-lib
```

An implementation of HTTP Digest Authentication.

```
(make-digest-auth-header realm
                        private-key
                        opaque) → header?

realm : string?
private-key : string?
opaque : string?
```

Returns a header that instructs the Web browser to request a username and password from the client using Digest authentication with *realm* as the realm, *private-key* as the server's contribution to the nonce, and *opaque* as the opaque data passed through the client.

```
(request->digest-credentials req)
→ (or/c false/c (listof (cons/c symbol? string?)))
req : request?
```

Returns the Digest credentials from *req* (if they appear) as an association list.

```
username*realm->password/c : contract?
```

Used to look up the password for a user in a realm.

Equivalent to `(-> string? string? string?)`.

```
username*realm->digest-HA1/c : contract?
```

Used to compute the user's secret hash.

Equivalent to `(-> string? string? bytes?)`.

```
(password->digest-HA1 lookup-password)
→ username*realm->digest-HA1/c
lookup-password : username*realm->password/c
```

Uses *lookup-password* to find the password, then computes the secret hash of it.

```
(make-check-digest-credentials lookup-HA1)
→ (string? (listof (cons/c symbol? string?))) . -> . boolean?)
lookup-HA1 : username*realm->digest-HA1/c
```

Constructs a function that checks whether particular Digest credentials (the second argument of the returned function) are correct given the HTTP method provided as the first argument and the secret hash computed by *lookup-HA1*.

This will result in an exception if the Digest credentials are missing portions.

Example:

```
#lang web-server/insta
(require racket/pretty)

(define private-key "private-key")
(define opaque "opaque")

(define (start req)
  (match (request->digest-credentials req)
    [#f
     (response
      401 #"Unauthorized" (current-seconds) TEXT/HTML-MIME-TYPE
      (list (make-digest-auth-header
              (format "Digest Auth Test: ~a" (gensym))
              private-key opaque))
      void)]
    [alist
     (define check
       (make-check-digest-credentials
        (password->digest-HA1 (lambda (username realm) "pass"))))
     (define pass?
       (check "GET" alist))
     (response/xexpr
      `(html (head (title "Digest Auth Test"))
              (body
               (h1 ,(if pass? "Pass!" "No Pass!"))
               (pre ,(pretty-format alist))))))]))
```

4.10 X-expression Support

```
(require web-server/http/xexpr)      package: web-server-lib

(response/xexpr xexpr
  [#:code code
   #:message message
   #:seconds seconds
   #:mime-type mime-type
   #:headers headers
   #:cookies cookies
   #:preamble preamble]) → response?

xexpr : xexpr/c
code : response-code/c = 200
message : (or/c #f bytes?) = #f
seconds : real? = (current-seconds)
mime-type : (or/c #f bytes?) = TEXT/HTML-MIME-TYPE
headers : (listof header?) = empty
cookies : (listof cookie?) = empty
preamble : bytes? = #""
```

Equivalent to

```
(response/full
  code message seconds mime-type
  (append headers (map cookie->header cookies)))
(list preamble (string->bytes/utf-8 (xexpr->string xexpr))))
```

This is a viable function to pass to `set-any->response!`.

See the documentation for `response/full` to see how `#f` is handled for `message`.

Changed in version 1.3 of package `web-server-lib`: Updated contracts on `code` and `seconds` as with `response`.

Changed in version 1.4: Contract on `message` relaxed to allow both `#f` and `bytes?`, with a default of `#f`. Previously, `bytes?` was required, and had a default of `#"Okay"`.

4.11 Empty Responses

```
(require web-server/http/empty)      package: web-server-lib
```

```
(response/empty [#:code code
                #:message message
                #:cookies cookies
                #:seconds seconds
                #:headers headers]) → response?
code : number? = 204
message : (or/c false/c bytes?) = #f
cookies : (listof cookie?) = '()
seconds : number? = (current-seconds)
headers : (listof header?) = '()
```

Generates a response with an empty body. The usual Content-Type header will be absent, unless passed in via *headers*. Equivalent to

```
(response code message seconds #f headers (λ (o) (write-
bytes #"" o)))
```

, with the understanding that if *message* is missing (or *#f*), it will be inferred from *code* using the association between status codes and messages found in RFCs 7231 and 7235. See the documentation for *response/full* for the table of built-in status codes.

Added in version 1.6 of package *web-server-lib*.

Changed in version 1.7: Make default response code 204 rather than 200.

Changed in version 1.7: Ensure a Content-Length header is present in the response, with value 0.

4.12 JSON Support

```
(require web-server/http/json)      package: web-server-lib
```

JSON is a widely used data format for the web. Racket's JSON library meets the web server with *response/jsexpr*, which is for JSON what *response/xexpr* is for XML.

```
(response/jsexpr jsexpr
                [#:code code
                #:message message
                #:seconds seconds
                #:mime-type mime-type
                #:headers headers
                #:cookies cookies]) → response?
jsexpr : jsexpr?
code : response-code/c = 200
message : (or/c #f bytes?) = #f
seconds : real? = (current-seconds)
mime-type : (or/c #f bytes?) = APPLICATION/JSON-MIME-TYPE
```

```
headers : (listof header?) = empty
cookies : (listof cookie?) = empty
```

Equivalent to

```
(response/full
  code message seconds mime-type
  (append headers (map cookie->header cookies))
  (list (jsexpr->bytes jsexpr)))
```

See the documentation for [response/full](#) to see how *message*, if *#f*, is turned into a *bytes?*.

Added in version 1.5 of package `web-server-lib`.

5 URL-Based Dispatch

```
(require web-server/dispatch)    package: web-server-lib
```

The library allows the creation of two-way mappings between permanent URLs and request-handling procedures.

This library was
inspired by the
(planet
untyped/dispatch)
package.

5.1 Using web-server/dispatch

Suppose you are writing a blog application and want pretty URLs for different views of the site. You would define some URL dispatching rules as follows:

```
> (define-values (blog-dispatch blog-url)
  (dispatch-rules
    [("") list-posts]
    [("posts" (string-arg)) review-post]
    [("archive" (integer-arg) (integer-arg)) review-archive]
    [else list-posts]))
```

And define your request handlers as follows:

```
> (define (list-posts req) `(list-posts))
> (define (review-post req p) `(review-post ,p))
> (define (review-archive req y m) `(review-archive ,y ,m))
```

Now when a request is sent to your application, it will be directed to the appropriate handler:

```
> (define (url->request u)
  (make-request #"GET" (string->url u) empty
    (delay empty) #f "1.2.3.4" 80 "4.3.2.1"))
> (blog-dispatch
  (url->request "http://www.chrlsnchrg.com"))
'(list-posts)
> (blog-dispatch
  (url->request "http://www.chrlsnchrg.com/"))
'(list-posts)
> (blog-dispatch
  (url->request
    "http://www.chrlsnchrg.com/posts/Extracurricular-Activity"))
'(review-post "Extracurricular-Activity")
> (blog-dispatch
  (url->request "http://www.chrlsnchrg.com/archive/1984/10"))
'(review-archive 1984 10)
```

```
> (blog-dispatch
  (url->request "http://www.chrlsnchrg.com/contact"))
'(list-posts)
```

You can also generate these pretty URLs from procedure calls:

```
> (blog-url list-posts)
"/"
> (blog-url review-post "Another-Saturday-Night")
"/posts/Another-Saturday-Night"
> (blog-url review-archive 1984 11)
"/archive/1984/11"
```

After mastering the world of blogging software, you decide to put the ubiquitous Add-Two-Numbers.com out of business with Sum.com:

```
> (define-values (sum-dispatch sum-url)
  (dispatch-rules
    [((integer-arg) ...) sum]
    [else (lambda (req) (sum req empty))]))
> (define (sum req is)
  (apply + is))
> (sum-dispatch (url->request "http://www.sum.com/"))
0
> (sum-dispatch (url->request "http://www.sum.com/2"))
2
> (sum-dispatch (url->request "http://www.sum.com/2/3/4"))
9
> (sum-dispatch (url->request "http://www.sum.com/5/10/15/20"))
50
> (sum-url sum empty)
"/"
> (sum-url sum (list 1))
"/1"
> (sum-url sum (list 2 3 5 7))
"/2/3/5/7"
```

When you use `web-server/dispatch` with `serve/servlet`, you almost always want to use the `#:servlet-regexp` argument with the value `"` to capture all top-level requests. However, make sure you don't include an `else` in your rules if you are also serving static files, or else the filesystem server will never see the requests.

5.2 API Reference

```

(dispatch-rules
 dispatch-clause ...
 maybe-else-clause)

dispatch-clause = [dispatch-pattern maybe-method dispatch-fun]

dispatch-pattern = ()
                  | (string . dispatch-pattern)
                  | (bidi-match-expander ... . dispatch-pattern)
                  | (bidi-match-expander . dispatch-pattern)

maybe-method =
               | #:method method

method = pat

maybe-else-clause =
                 | [else else-fun]

else-fun : (request? . -> . any)
dispatch-fun : (request? any/c ... . -> . any)

```

Returns two values: the first is a dispatching function with the contract `(-> request? any)` that calls the appropriate `dispatch-fun` based on the first `dispatch-pattern` that matches the request's URL (and method), the second is a URL-generating function with the contract `(-> procedure? any/c ... string?)` that generates a URL using `dispatch-pattern` for the `dispatch-fun` given as its first argument.

If `else-fun` is left out, one is provided that calls `(next-dispatcher)` to signal to the Web Server that this dispatcher does not apply.

The `method` syntax is used in a match expression to match the `request-method` part of the incoming request object. However, since HTTP allows methods to use any case, the byte string from `request-method` is normalized to a lower-case string. Thus, valid patterns are things like: `"get"`, `"post"`, `"head"`, (or `"get"` `"post"`), etc.

If `method` is left out, it assumed to apply to requests without methods and GET methods.

```

(dispatch-rules+applies
 dispatch-clause ...
 maybe-else-clause)

```

Like `dispatch-rules`, `except` returns a third value with the contract `(-> request? boolean?)` that returns `#t` if the dispatching rules apply to the request and `#f` otherwise.

```
(dispatch-case
  dispatch-clause ...
  maybe-else-clause)
```

Returns a dispatching function as described by `dispatch-rules`.

```
(dispatch-url
  [dispatch-pattern dispatch-fun]
  ...)

dispatch-fun : (request? any/c ... . -> . any)
```

Returns a URL-generating function as described by `dispatch-rules`.

```
(serve/dispatch dispatch) → void
dispatch : (request? . -> . can-be-response?)
```

Calls `serve/servlet` with a `#:servlet-regex` argument (`#rx"`) so that every request is handled by `dispatch`.

5.3 Imperative Dispatch Containers

`dispatch-rules` is purely functional. This presents a more declarative interface, but inhibits some programming and modularity patterns. *Containers* provide an imperative overlay atop `dispatch-rules`.

```
(container? x) → boolean?
x : any/c
```

Identifies containers.

```
(define-container container-id (dispatch-id url-id))
```

Defines `container-id` as a container as well as `dispatch-id` as its dispatching function and `url-id` as its URL lookup function.

```
(dispatch-rules! container-expr [dispatch-pattern dispatch-fun] ...)
```

Like `dispatch-rules`, but imperatively adds the patterns to the container specified by `container-expr`. The new rules are consulted *before* any rules already in the container.

5.4 Built-in URL patterns

`web-server/dispatch` builds in a few useful URL component patterns.

`(number-arg)`

A bi-directional match expander that parses a `number?` from the URL and generates a URL with a number's encoding as a string.

`(integer-arg)`

A bi-directional match expander that parses a `integer?` from the URL and generates a URL with an integer's encoding as a string.

`(real-arg)`

A bi-directional match expander that parses a `real?` from the URL and generates a URL with a real's encoding as a string.

`(string-arg)`

A bi-directional match expander that parses a `string?` from the URL and generates a URL containing the string.

`(symbol-arg)`

A bi-directional match expander that parses a `symbol?` from the URL and generates a URL with a symbol's encoding as a string.

5.5 Extending `web-server/dispatch`

```
(require web-server/dispatch/extend)
package: web-server-lib
```

You can create new URL component patterns by defining bi-directional match expanders.

`(define-bidi-match-expander id in-xform out-xform)`

Binds *id* to a *bi-directional match expander* where *in-xform* is a match expander (defined by `define-match-expander`) that is used when parsing URLs and *out-xform* is one used when generating URLs.

Both *in-xform* and *out-xform* should use the syntax `(xform arg ... id)` where the *args* are specific to *id* and compatible with both *in-xform* and *out-xform*. *id* will typically be provided automatically by `dispatch-rules`.

`| bidi-match-going-in?`

A syntax parameter used by bi-directional match expanders to determine if a URL is being parsed or generated.

When defining new patterns, you may find it useful to use these helper functions:

`| (define-coercion-match-expander id test? coerce)`

Binds *id* to a match expander that expands `(id x)` to `(? test? (app coerce x))` (i.e., uses *test?* to determine if the pattern matches and *coerce* to transform the binding.)

`| (make-coerce-safe? coerce) → (any/c . -> . boolean?)`
`coerce : (any/c . -> . any/c)`

Returns a function that returns `#t` if *coerce* would not throw an exception or return `#f` on its input.

Examples:

```
> (define string->number? (make-coerce-safe? string->number))
> (string->number? "1")
#t
> (string->number? "1.2")
#t
> (string->number? "+inf.0")
#t
> (string->number? "one")
#f
```

6 Formlets: Functional Form Abstraction

```
(require web-server/formlets)      package: web-server-lib
```

The Web Server provides a kind of Web form abstraction called a formlet.

Formlets originate in the work of the Links research group in their paper *The Essence of Form Abstraction*.

6.1 Basic Formlet Usage

Suppose we want to create an abstraction of entering a date in an HTML form. The following *formlet* captures this idea:

```
(define date-formlet
  (formlet
    (div "Month:" ,{input-int . => . month}
         "Day:" ,{input-int . => . day}))
    (list month day)))
```

The first part of the formlet syntax is the template of an X-expression that is the rendering of the formlet. It can contain elements like `(=> formlet name)` where *formlet* is a formlet expression and *name* is an identifier bound in the second part of the formlet syntax.

This formlet is displayed (with `formlet-display`) as the following X-expression forest (list):

```
(list
  '(div "Month:" (input ([name "input_0"])))
    "Day:" (input ([name "input_1"]))))
```

`date-formlet` not only captures the rendering of the form, but also the request processing logic. If we send it an HTTP request with bindings for `"input_0"` to `"10"` and `"input_1"` to `"3"`, with `formlet-process`, then it returns:

```
(list 10 3)
```

which is the second part of the formlet syntax, where `month` has been replaced with the integer represented by the `"input_0"` and `day` has been replaced with the integer represented by the `"input_1"`.

The real power of formlet is that they can be embedded within one another. For instance, suppose we want to combine two date forms to capture a travel itinerary. The following formlet does the job:

```
(define travel-formlet
```

```
(formlet
  (div
    "Name:" ,{input-string . => . name}
    (div
      "Arrive:" ,{date-formlet . => . arrive}
      "Depart:" ,{date-formlet . => . depart})
    (list name arrive depart))))
```

(Notice that `date-formlet` is embedded twice.) This is rendered as:

```
(list
  '(div
    "Name:"
    (input ([name "input_0"])))
    (div
      "Arrive:"
      (div "Month:" (input ([name "input_1"])))
        "Day:" (input ([name "input_2"])))
      "Depart:"
      (div "Month:" (input ([name "input_3"])))
        "Day:" (input ([name "input_4"]))))))
```

Observe that `formlet-display` has automatically generated unique names for each input element. When we pass bindings for these names to `formlet-process`, the following list is returned:

```
(list "Jay"
  (list 10 3)
  (list 10 6))
```

In all these examples, we used the `input-int` and `input-string` formlets. Any value with the formlet contract can be used in these positions. For example, `(to-string (required (text-input)))` could be used as well. The rest of the manual gives the details of formlet usage, extension, and existing formlet combinators.

6.2 Static Syntactic Shorthand

```
(require web-server/formlets/syntax)
package: web-server-lib
```

Most users will want to use the syntactic shorthand for creating formlets.

```
| (formlet rendering-xexpr yields-expr)
```

Constructs a formlet with the specified *rendering-expr* and the processing result is the evaluation of the *yields-expr* expression. The *rendering-expr* form is a quasiquoted syntactic X-expression, with three special caveats:

`{=> formlet-expr name}` embeds the formlet given by *formlet-expr*; the result of processing this formlet is available in the *yields-expr* as *name*.

`{=> formlet-expr (values name ...)}` embeds the formlet given by *formlet-expr*; the results of processing this formlet is available in the *yields-expr* as *name*

`(#%# xexpr ...)` renders an X-expression forest.

These forms *may not* appear nested inside `unquote` or `unquote-splicing`. For example, this is illegal:

```
(formlet (div ,@(for/list ([i (in-range 10)])
  `(p ,((text-input) . => . name))))
  name)
```

| `#%#`

Only allowed inside `formlet` and `formlet*`.

6.3 Dynamic Syntactic Shorthand

```
(require web-server/formlets/dyn-syntax)
package: web-server-lib
```

The formlet syntax is too restrictive for some applications because it forces the *rendering* to be *syntactically* an X-expression. You may discover you want to use a more "dynamic" shorthand.

| `(formlet* rendering-expr yields-expr)`

Constructs a formlet where *rendering-expr* is evaluated (with caveats) to construct the rendering and the processing result is the evaluation of the *yields-expr* expression. The *rendering-expr* should evaluate to an "X-expression" that may embed the results of the following forms that only have meaning within `formlet*`:

`{=>* formlet-expr name}` embeds the formlet given by *formlet-expr*; the result of processing this formlet is available in the *yields-expr* as *name*.

`{=>* formlet-expr (values name ...)}` embeds the formlet given by *formlet-expr*; the results of processing this formlet is available in the *yields-expr* as *name*

(`#/%# xexpr-expr . . .`) renders an X-expression forest.

Each of these forms evaluates to an opaque value that *rendering-expr* may not manipulate in any way, but if it is returned to `formlet*` as part of an "X-expression" it will be rendered and the formlets processing stages will be executed, etc.

Because these forms *may* appear anywhere in *rendering-expr*, they may be duplicated. Therefore, the formlet may render (and be processed) multiple times. Thus, in *yields-expr* the formlet result names are bound to lists of results rather than single results as in `formlet`. The result list is ordered according to the order of the formlets in the result of *rendering-expr*. For example, in

```
(formlet* `(div ,@(for/list ([i (in-range 1 10)])
                          `(p ,(number->string i)
                             ,((text-input) . =>* . name))))
  name)
```

`name` is bound to a list of strings, not a single string, where the first element is the string that was inputted next to the string `1` on the Web page.

In this example, it is clear that this is the desired behavior. However, sometimes the value of a formlet's result may be surprising. For example, in

```
(formlet* `(div (p ,((text-input) . =>* . name)))
  name)
```

`name` is bound to a list of strings, because `formlet*` cannot syntactically determine if the formlet whose result is bound to `name` is used many times.

`=>*`

Only allowed inside `formlet*`.

6.4 Functional Usage

```
(require web-server/formlets/lib)    package: web-server-lib
```

The syntactic shorthand abbreviates the construction of formlets with the following library. These combinators may be used directly to construct low-level formlets, such as those for new INPUT element types. Refer to §6.5 “Predefined Formlets” for example low-level formlets using these combinators.

`xexpr-forest/c` : `contract?`

Equivalent to `(listof xexpr/c)`

```
(formlet/c content ...) → contract?
  content : contract?
```

Equivalent to `(-> integer? (values xexpr-forest/c (-> (listof binding?) (values (coerce-contract 'formlet/c content) ...)) integer?)))`.

A formlet's internal representation is a function from an initial input number to an X-expression forest rendering, a processing function, and the next allowable input number.

(Actually, `formlet/c` is a macro which avoids using `dynamic->*` when the number of range contracts for the processing function is known at compile time.)

Changed in version 1.3 of package `web-server-lib`: Fixed support for multiple return values.

```
formlet*/c : contract?
```

Similar to the contracts created by `formlet/c`, but uses `any` to avoid checking the results (or even specifying the number of results) of the processing function.

Changed in version 1.3 of package `web-server-lib`: Fixed support for multiple return values.

```
(pure value) → (formlet/c any/c)
  value : any/c
```

Constructs a formlet that has no rendering and always returns `value` in the processing stage.

```
(cross f g) → (formlet/c any/c ...)
  f : (formlet/c procedure?)
  g : (formlet/c any/c ...)
```

Constructs a formlet with a rendering equal to the concatenation of the renderings of formlets `f` and `g`; a processing stage that applies `g`'s processing results to `f`'s processing result.

```
(cross* f g ...) → (formlet/c any/c)
  f : (formlet/c (() () #:rest (listof any/c) . ->* . any/c))
  g : (formlet/c any/c)
```

Equivalent to `cross` lifted to many arguments.

```
(xml-forest r) → (formlet/c procedure?)
  r : xexpr-forest/c
```

Constructs a formlet with the rendering *r* and the identity procedure as the processing step.

```
(xml r) → (formlet/c procedure?)
  r : xexpr/c
```

Equivalent to `(xml-forest (list r))`.

```
(text r) → (formlet/c procedure?)
  r : string?
```

Equivalent to `(xml r)`.

```
(tag-xexpr tag attrs inner) → (formlet/c any/c)
  tag : symbol?
  attrs : (listof (list/c symbol? string?))
  inner : (formlet/c any/c)
```

Constructs a formlet with the rendering `(list (list* tag attrs inner-rendering))` where *inner-rendering* is the rendering of *inner* and the processing stage identical to *inner*.

```
(formlet-display f) → xexpr-forest/c
  f : (formlet/c any/c)
```

Renders *f*.

```
(formlet-process f r) → any/c ...
  f : (formlet/c any/c ...)
  r : request?
```

Runs the processing stage of *f* on the bindings in *r*.

6.5 Predefined Formlets

```
(require web-server/formlets/input)
```

```
package: web-server-lib
```

These formlets are the main combinators for form input.

```
(make-input render) → (formlet/c (or/c false/c binding?))
  render : (string? . -> . xexpr/c)
```

This formlet is rendered with `render`, which is passed the input name, and results in the extracted `binding`.

```
(make-input* render) → (formlet/c (listof binding?))
  render : (string? . -> . xexpr/c)
```

This formlet is rendered with `render`, which is passed the input name, and results in all the `bindings` that use the name.

```
(input #:type type
      [#:value value
       #:size size
       #:max-length max-length
       #:read-only? read-only?
       #:attributes attrs])
→ (formlet/c (or/c false/c binding?))
  type : string?
  value : (or/c false/c bytes? string?) = #f
  size : (or/c false/c exact-nonnegative-integer?) = #f
  max-length : (or/c false/c exact-nonnegative-integer?) = #f
  read-only? : boolean? = #f
  attrs : (listof (list/c symbol? string?)) = empty
```

This formlet renders using an `INPUT` element with specified type and arguments.

```
(text-input [#:value value
            #:size size
            #:max-length max-length
            #:read-only? read-only?
            #:attributes attrs])
→ (formlet/c (or/c false/c binding?))
  value : (or/c false/c bytes? string?) = #f
  size : (or/c false/c exact-nonnegative-integer?) = #f
  max-length : (or/c false/c exact-nonnegative-integer?) = #f
  read-only? : boolean? = #f
  attrs : (listof (list/c symbol? string?)) = empty
```

This formlet renders using an INPUT element with the TEXT type and the attributes given in the arguments.

```
(password-input [#:value value
                 #:size size
                 #:max-length max-length
                 #:read-only? read-only?
                 #:attributes attrs])
→ (formlet/c (or/c false/c binding?))
value : (or/c false/c bytes? string?) = #f
size : (or/c false/c exact-nonnegative-integer?) = #f
max-length : (or/c false/c exact-nonnegative-integer?) = #f
read-only? : boolean? = #f
attrs : (listof (list/c symbol? string?)) = empty
```

This formlet renders using an INPUT element with the PASSWORD type and the attributes given in the arguments.

```
(textarea-input [#:value value
                 #:rows rows
                 #:cols cols
                 #:attributes attrs])
→ (formlet/c (or/c false/c binding?))
value : (or/c false/c bytes? string?) = #f
rows : (or/c false/c number?) = #f
cols : (or/c false/c number?) = #f
attrs : (listof (list/c symbol? string?)) = empty
```

This formlet renders using an TEXTAREA element with attributes given in the arguments.

```
(checkbox value checked? [#:attributes attrs])
→ (formlet/c (or/c false/c binding?))
value : (or/c bytes? string?)
checked? : boolean?
attrs : (listof (list/c symbol? string?)) = empty
```

This formlet renders using an INPUT element with the CHECKBOX type and the attributes given in the arguments.

```
(radio value checked? [#:attributes attrs])
→ (formlet/c (or/c false/c binding?))
value : (or/c bytes? string?)
checked? : boolean?
attrs : (listof (list/c symbol? string?)) = empty
```

This formlet renders using an INPUT element with the RADIO type and the attributes given in the arguments.

```
(radio-group l
  [#:attributes attrs
   #:checked? checked?
   #:display display
   #:wrap wrap]) → (formlet/c any/c)
l : sequence?
attrs : (any/c . -> . (listof (list/c symbol? string?)))
      = (λ (x) empty)
checked? : (any/c . -> . boolean?) = (λ (x) #f)
display : (any/c . -> . xexpr/c) = (λ (x) x)
wrap : (any/c any/c . -> . xexpr/c) = (λ (x y) (list x y))
```

This formlet renders using a sequence of INPUT elements of RADIO type where each element gets its attributes from *attrs* that share a single NAME. An element is checked if *checked?* returns *#t*. Elements are combined with the results of *display* into an X-expression specified in *wrap*. The result of processing this formlet is a single element of the sequence.

```
(checkbox-group l
  [#:attributes attrs
   #:checked? checked?
   #:display display
   #:wrap wrap])
→ (formlet/c (listof any/c))
l : sequence?
attrs : (any/c . -> . (listof (list/c symbol? string?)))
      = (λ (x) empty)
checked? : (any/c . -> . boolean?) = (λ (x) #f)
display : (any/c . -> . xexpr/c) = (λ (x) x)
wrap : (any/c any/c . -> . xexpr/c) = (λ (x y) (list x y))
```

This formlet renders using a sequence of INPUT elements of CHECKBOX type where each element gets its attributes from *attrs* that share a single NAME. An element is checked if *checked?* returns *#t*. Elements are combined with the results of *display* into an X-expression specified in *wrap*. The result of processing this formlet is a list of elements of the sequence.

```
(submit value [#:attributes attrs])
→ (formlet/c (or/c false/c binding?))
value : (or/c bytes? string?)
attrs : (listof (list/c symbol? string?)) = empty
```

This formlet renders using an INPUT element with the SUBMIT type and the attributes given in the arguments.

```
(reset value [#:attributes attrs])  
→ (formlet/c (or/c false/c binding?))  
  value : (or/c bytes? string?)  
  attrs : (listof (list/c symbol? string?)) = empty
```

This formlet renders using an INPUT element with the RESET type and the attributes given in the arguments.

```
(file-upload [#:attributes attrs])  
→ (formlet/c (or/c false/c binding?))  
  attrs : (listof (list/c symbol? string?)) = empty
```

This formlet renders using an INPUT element with the FILE type and the attributes given in the arguments.

```
(hidden value [#:attributes attrs])  
→ (formlet/c (or/c false/c binding?))  
  value : (or/c bytes? string?)  
  attrs : (listof (list/c symbol? string?)) = empty
```

This formlet renders using an INPUT element with HIDDEN type and the attributes given in the arguments.

```
(img alt  
  src  
  [#:height height  
    #:longdesc ldesc  
    #:usemap map  
    #:width width  
    #:attributes attrs]) → (formlet/c (or/c false/c binding?))  
alt : (or/c bytes? string?)  
src : (or/c bytes? string?)  
height : (or/c false/c exact-nonnegative-integer?) = #f  
ldesc : (or/c false/c bytes? string?) = #f  
map : (or/c false/c bytes? string?) = #f  
width : (or/c false/c exact-nonnegative-integer?) = #f  
attrs : (listof (list/c symbol? string?)) = empty
```

This formlet renders using an IMG element with the attributes given in the arguments.

```

(button type
      button-text
      [#:disabled disabled
       #:value value
       #:attributes attrs])
→ (formlet/c (or/c false/c binding?))
type : (or/c bytes? string?)
button-text : (or/c bytes? string?)
disabled : boolean? = #f
value : (or/c false/c bytes? string?) = #f
attrs : (listof (list/c symbol? string?)) = empty

```

This formlet renders using a `BUTTON` element with the attributes given in the arguments. `button-text` is the text that will appear on the button when rendered.

```

(multiselect-input l
  [#:attributes attrs
   #:multiple? multiple?
   #:selected? selected?
   #:display display]) → (formlet/c list?)

l : sequence?
attrs : (listof (list/c symbol? string?)) = empty
multiple? : boolean? = #t
selected? : (any/c . -> . boolean?) = (λ (x) #f)
display : (any/c . -> . xexpr/c) = (λ (x) x)

```

This formlet renders using an `SELECT` element with the attributes given with an `OPTION` for each element of the sequence. If `multiple?` is `#t`, then multiple options may be selected. An element is selected if `selected?` returns `#t`. Elements are displayed with `display`.

```

(select-input l
  [#:attributes attrs
   #:selected? selected?
   #:display display]) → (formlet/c any/c)

l : sequence?
attrs : (listof (list/c symbol? string?)) = empty
selected? : (any/c . -> . boolean?) = (λ (x) #f)
display : (any/c . -> . xexpr/c) = (λ (x) x)

```

This formlet renders using an `SELECT` element with the attributes given with an `OPTION` for each element of the sequence. An element is selected if `selected?` returns `#t`. Elements are displayed with `display`.

```
(required f) → (formlet/c bytes?)  
f : (formlet/c (or/c false/c binding?))
```

Constructs a formlet that extracts the `binding:form-value` from the binding produced by `f`, or errors.

```
(default def f) → (formlet/c bytes?)  
def : bytes?  
f : (formlet/c (or/c false/c binding?))
```

Constructs a formlet that extracts the `binding:form-value` from the binding produced by `f`, or returns `def`.

```
(to-string f) → (formlet/c string?)  
f : (formlet/c bytes?)
```

Converts `f`'s output to a string. Equivalent to `(cross (pure bytes->string/utf-8) f)`.

```
(to-number f) → (formlet/c number?)  
f : (formlet/c string?)
```

Converts `f`'s output to a number. Equivalent to `(cross (pure string->number) f)`.

```
(to-symbol f) → (formlet/c symbol?)  
f : (formlet/c string?)
```

Converts `f`'s output to a symbol. Equivalent to `(cross (pure string->symbol) f)`.

```
(to-boolean f) → (formlet/c boolean?)  
f : (formlet/c bytes?)
```

Converts `f`'s output to a boolean, if it is equal to `#"on"`.

```
input-string : (formlet/c string?)
```

Equivalent to `(to-string (required (text-input)))`.

```
input-int : (formlet/c number?)
```

Equivalent to `(to-number input-string)`. Note that, despite the name, the result **is not guaranteed to be an integer**.

Changed in version 1.3 of package `web-server-lib`: Weakened result contract to allow any number.

```
input-symbol : (formlet/c symbol?)
```

Equivalent to `(to-symbol input-string)`.

6.6 Utilities

```
(require web-server/formlets/servlet)
package: web-server-lib
```

A few utilities are provided for using formlets in Web applications.

```
(send/formlet f
  [#:method method
   #:wrap wrapper]) → any/c ...
f : (formlet/c any/c ...)
method : (or/c "GET" "POST" "get" "post") = "POST"
wrapper : (xexpr/c . -> . xexpr/c)
          = (lambda (form-xexpr)
              `(html (head (title "Form Entry"))
                     (body ,form-xexpr)))
```

Uses `send/suspend` and `response/xexpr` to send `f`'s rendering (wrapped in a FORM tag with method `method` whose action is the continuation URL (wrapped again by `wrapper`)) to the client. When the form is submitted, the request is passed to the processing stage of `f`.

```
(embed-formlet embed/url f) → xexpr/c
embed/url : ((request? . -> . any) . -> . string?)
f : (formlet/c any/c ...)
```

Like `send/formlet`, but for use with `send/suspend/dispatch`.

6.7 Formlets and Stateless Servlets

```
(require web-server/formlets/stateless)
```

```
(require web-server/formlets/unsafe)
package: web-server-lib
```

A few additional considerations apply when using formlets with stateless `#lang web-server` servlets.

First of all, continuations captured in your servlet cannot be serialized if they close over non-serializable data-structures. There are some generally-applicable ways to avoid having a data structure be part of the closure: for example, if you define all of your formlets as module-level variables, they will never be part of closures and will not need to be serialized. However, sometimes it can be useful to create formlets dynamically. To support this, all of the combinators documented above produce formlets that are serializable (as long as they contain only serializable values). This is not guaranteed to be true of third-party formlets.

Secondly, stateless `#lang web-server` servlets are based on different web interaction primitives than stateful servlets, so the version of `send/formlet` from `web-server/formlets` will not work. Instead, the library `web-server/formlets/stateless` provides the same API as `web-server/formlets`, but with a version of `send/formlet` for use in stateless servlets. (Using `web-server/formlets/stateless` also provides all of the bindings from `web-server/formlets/lib`, whereas `web-server/formlets` provides only some of them.) Alternatively, you can use the low-level `formlet-process` and `formlet-display` procedures directly.

Another issue concerns capturing continuations within the processing stage of a formlet. Recall that serializable continuations in `#lang web-server` can only be captured from within transformed contexts. The contract system is not transformed, so the contracts on this library prevent capturing continuations during the processing stage of formlets. In most cases, the best solution is simply to avoid using continuation-capturing operations during a formlet's processing stage. Instead, have the processing stage return a value, and interact with the user based on that value in code outside of the formlet. Alternatively, you can use generally-applicable approaches for capturing continuations from untransformed contexts, such as `web-server/lang/native`. However, if neither of those approaches are satisfactory, the library `web-server/formlets/unsafe` provides the same API as `web-server/formlets/stateless`, but *without enforcing contracts*. As the name implies, using `web-server/formlets/unsafe` may produce inscrutable error messages and other unpleasant effects of programming without contracts: you have been warned.

Changed in version 1.3 of package `web-server-lib`: Added `web-server/formlets/stateless` and `web-server/formlets/unsafe` and changed combinators from `web-server/formlets` to produce serializable formlets.

One potential pitfall for formlets and serialization is `pure`. Note that `(serialize (pure +))` will **fail**, because `+` is not serializable. To avoid this, you can write `(pure (λ args (apply + args)))` (in `#lang web-server`, where anonymous procedures are serializable, or using `web-server/lang/serial-lamb`

7 Templates: Separation of View

```
(require web-server/templates)    package: web-server-lib
```

The Web Server provides a powerful Web template system for separating the presentation logic of a Web application and enabling non-programmers to contribute to Racket-based Web applications.

Although all the examples here generate HTML, the template language and the §1 “Text Generation” it is based on can be used to generate any text-based format: C, SQL, form emails, reports, etc.

7.1 Static

Suppose we have a file "static.html" with the contents:

```
<html>
  <head><title>Fastest Templates in the West!</title></head>
  <body>
    <h1>Bang!</h1>
    <h2>Bang!</h2>
  </body>
</html>
```

If we write the following in our code:

```
(include-template "static.html")
```

Then the contents of "static.html" will be read *at compile time* and compiled into a racket program that returns the contents of "static.html" as a string:

```
"<html>\n  <head><title>Fastest Templates in the
West!</title></head>\n  <body>\n    <h1>Bang!</h1>\n    <h2>Bang!</h2>\n  </body>\n</html>"
```

7.2 Dynamic

include-template gives the template access to the *complete lexical context* of the including program. This context can be accessed via the §2 “@ Syntax” syntax. For example, if "simple.html" contains:

```
<html>
  <head><title>Fastest @thing in the West!</title></head>
  <body>
    <h1>Bang!</h1>
    <h2>Bang!</h2>
```

```
</body>
</html>
```

Then

```
(let ([thing "Templates"])
  (include-template "simple.html"))
```

evaluates to the same content as the static example.

There are no constraints on how the lexical context of the template is populated. For instance, you can build template abstractions by wrapping the inclusion of a template in a function:

```
(define (fast-template thing)
  (include-template "simple.html"))

(fast-template "Templates")
(fast-template "Noodles")
```

evaluates to two strings with the predictable contents:

```
<html>
  <head><title>Fastest Templates in the West!</title></head>
  <body>
    <h1>Bang!</h1>
    <h2>Bang!</h2>
  </body>
</html>
```

and

```
<html>
  <head><title>Fastest Noodles in the West!</title></head>
  <body>
    <h1>Bang!</h1>
    <h2>Bang!</h2>
  </body>
</html>
```

Furthermore, there are no constraints on the Racket used by templates: they can use macros, structs, continuation marks, threads, etc. However, Racket values that are ultimately returned must be printable by the §1 “Text Generation”. For example, consider the following outputs of the title line of different calls to `fast-template`:

- `(fast-template 'Templates)`

```
...<title>Fastest Templates in the West!</title>...
```

- `(fast-template 42)`

```
...<title>Fastest 42 in the West!</title>...
```

- `(fast-template (list "Noo" "dles"))`

```
...<title>Fastest Noodles in the West!</title>...
```

- `(fast-template (lambda () "Thunks"))`

```
...<title>Fastest Thunks in the West!</title>...
```

- `(fast-template (delay "Laziness"))`

```
...<title>Fastest Laziness in the West!</title>...
```

- `(fast-template (fast-template "Embedding"))`

```
...<title>Fastest ...<title>Fastest Embedding in the  
West!</title>... in the West!</title>...
```

7.3 Gotchas

7.3.1 @ Syntax: @ character, identifiers, and spaces

To obtain an @ character in template output, you must escape it since @ is the escape character of the §2 “@ Syntax”. For example, to obtain:

```
<head><title>Fastest @s in the West!</title></head>
```

You must write:

```
<head><title>Fastest @"@"s in the West!</title></head>
```

as your template: literal `@`s must be replaced with `@"`. (Note that the double-quotes are basically a Racket expression, which can be used for longer strings too.)

The §2 “@ Syntax” will read Racket identifiers, so it does not terminate identifiers on punctuations or XML angle brackets. So,

```
<head><title>Fastest @thing in the @place!</title></head>
```

will complain that the identifier `place!</title></head>` is undefined. You might be tempted to subvert this by adding a space after the identifier:

```
<head><title>Fastest @thing in the @place !</title></head>
```

This will remove the error, but the generated HTML will not look like what you want. Sometimes this is not relevant, but sometimes it is. The safest thing to do is explicitly delimit the identifier with `|`s:

```
<head><title>Fastest @thing in the @|place|!</title></head>
```

If you intend to use templates a lot, you should familiarize yourself with the details of the §2 “@ Syntax”.

7.3.2 Iteration

Since the template is compiled into a Racket program, only its results will be printed. For example, suppose we have the template:

```
<table>
  @for([c clients]){
    <tr><td>@(car c), @(cdr c)</td></tr>
  }
</table>
```

If this is included in a lexical context with `clients` bound to

```
(list (cons "Young" "Brigham") (cons "Smith" "Joseph"))
```

then the template will be printed as:

```
<table>
</table>
```

because `for` does not return the value of the body. Suppose that we change the template to use `for/list` (which combines them into a list):

```
<table>
  @for/list([[c clients]]){
    <tr><td>@(car c), @(cdr c)</td></tr>
  }
</table>
```

Now the result is:

```
<table>
  </td></tr>
  </td></tr>
</table>
```

because only the final expression of the body of the `for/list` is included in the result. We can capture all the sub-expressions by using `list` in the body:

```
<table>
  @for/list([[c clients]]){
    @list{
      <tr><td>@(car c), @(cdr c)</td></tr>
    }
  }
</table>
```

Now the result is:

```
<table>
  <tr><td>Young, Brigham</td></tr>
  <tr><td>Smith, Joseph</td></tr>
</table>
```

The templating library provides a syntactic form to deal with this issue for you called `in`:

```
<table>
  @in[c clients]{
    <tr><td>@(car c), @(cdr c)</td></tr>
  }
</table>
```

Notice how it also avoids the absurd amount of punctuation on line two.

7.3.3 Escaping

Thanks to Michael W. for this section.

Because templates are useful for many things (scripts, CSS, HTML, etc), the Web Server does not assume that the template is for XML-like content. Therefore when templates are expanded, no XML escaping is done by default. Beware of *cross-site scripting* vulnerabilities! For example, suppose a servlet serves the following template where *some-variable* is an input string supplied by the client:

```
<html>
  <head><title>Fastest Templates in the West!</title></head>
  <body>
    @some-variable
  </body>
</html>
```

If the servlet contains something like the following:

```
(let ([some-variable (get-input-from-user)])
  (include-template "static.htm"))
```

There is nothing to prevent an attacker from entering `<script type="text/javascript">...</script>` to make the template expand into:

```
<html>
  <head><title>Fastest Templates in the West!</title></head>
  <body>
    <script type="text/javascript">...</script>
  </body>
</html>
```

Now the server will send the attacker's code to millions of innocent users. To keep this from happening when serving HTML, use the `xexpr->string` function from the `xml` module.

This can be done in the servlet:

```
(require xml)

(let ([some-variable (xexpr->string (get-input-from-user))])
  (include-template "static.htm"))
```

Alternatively, make the template responsible for its own escaping:

```
<html>
  <head><title>Fastest Templates in the West!</title></head>
  <body>
```

```

    @(xexpr->string some-variable)
  </body>
</html>

```

The improved version renders as:

```

<html>
  <head><title>Fastest Templates in the West!</title></head>
  <body>
    &lt;script type="text/javascript"&gt;...&lt;/script&gt;
  </body>
</html>

```

When writing templates, always remember to escape user-supplied input.

7.4 HTTP Responses

The quickest way to generate an HTTP response from a template is using a `response?` struct:

```

(response/output
  (λ (op) (display (include-template "static.html") op)))

```

Finally, if you want to include the contents of a template inside a larger X-expression :

```

`(html ,(include-template "static.html"))

```

will result in the literal string being included (and entity-escaped). If you actually want the template to be unescaped, then use `include-template/xml` to assert that the content is valid XML.

```

`(html ,(include-template/xml "static.html"))

```

7.5 API Details

```

(include-template path-spec)
(include-template #:command-char command-char path-spec)

```

Compiles the template at *path-spec* using the §2 “@ Syntax” syntax within the enclosing lexical context. The *path-spec* is the same format used by `include`. Use the *command-char* keyword to customize the escape character.

Examples:

```

(include-template "static.html")
(include-template #:command-char #\$ "dollar-static.html")

(include-template/xml path-spec)
(include-template/xml #:command-char command-char path-spec)

```

Like `include/template`, but expands to a `cdata` structure.

Added in version 1.3 of package `web-server-lib`.

```
(in x xs e ...)
```

Expands into

```

(for/list ([x xs])
  (begin/text e ...))

```

Template Example:

```

@in[c clients]{
  <tr><td>@(car c), @(cdr c)</td></tr>
}

```

Racket Example:

```
(in c clients "<tr><td>" (car c) ", " (cdr c) "</td></tr>")
```

7.6 Conversion Example

Al Church has been maintaining a blog with Racket for some years and would like to convert to `web-server/templates`.

The data-structures he uses are defined as:

```

(define-struct post (title body))

(define posts
  (list
    (make-post
      "(Y Y) Works: The Why of Y"
      "Why is Y, that is the question.")
    (make-post
      "Church and the States"
      "As you may know, I grew up in DC, not technically a
state.))))

```

Actually, Al Church-encodes these posts, but for explanatory reasons, we'll use structs.

He has divided his code into presentation functions and logic functions. We'll look at the presentation functions first.

The first presentation function defines the common layout of all pages.

```
(define (template section body)
  (response/xexpr
    `(html
      (head (title "Al's Church: " ,section))
      (body
        (h1 "Al's Church: " ,section)
        (div ([id "main"])
          ,@body))))))
```

One of the things to notice here is the unquote-splicing on the `body` argument. This indicates that the `body` is list of X-expressions. If he had accidentally used only unquote then there would be an error in converting the return value to an HTTP response.

```
(define (blog-posted title body k-url)
  `((h2 ,title)
    (p ,body)
    (h1 (a ([href ,k-url]) "Continue"))))
```

Here's an example of simple body that uses a list of X-expressions to show the newly posted blog entry, before continuing to redisplay the main page. Let's look at a more complicated body:

```
(define (blog-posts k-url)
  (append
    (apply append
      (for/list ([p posts])
        `((h2 ,(post-title p))
          (p ,(post-body p))))))
    `((h1 "New Post")
      (form ([action ,k-url])
        (input ([name "title"]))
        (input ([name "body"]))
        (input ([type "submit"]))))))
```

This function shows a number of common patterns that are required by X-expressions. First, `append` is used to combine different X-expression lists. Second, `apply append` is used to collapse and combine the results of a `for/list` where each iteration results in a list of X-expressions. We'll see that these patterns are unnecessary with templates. Another

annoying patterns shows up when Al tries to add CSS styling and some JavaScript from Google Analytics to all the pages of his blog. He changes the `template` function to:

```
(define (template section body)
  (response/xexpr
    `(html
      (head
        (title "Al's Church: " ,section)
        (style ([type "text/css"])
          "body {margin: 0px; padding: 10px;}"
          "#main {background: #dddddd;}"))
      (body
        (script
          ([type "text/javascript"])
          ,(make-cdata
            #f #f
            "var gaJsHost = ((\"https:\" =="
            "document.location.protocol)"
            "? \"https://ssl.\" : \"http://www.\");"
            "document.write(unescape(\"%3Cscript src='\" + gaJsHost"
            "+ \"google-analytics.com/ga.js' \""
            "type='text/javascript'%3E%3C/script%3E\"));"))
          (script
            ([type "text/javascript"])
            ,(make-cdata
              #f #f
              "var pageTracker = _gat._getTracker(\"UA-YYYYYYY-Y\");"
              "pageTracker._trackPageview();"))
            (h1 "Al's Church: " ,section)
            (div ([id "main"])
              ,@body))))))
```

The first thing we notice is that encoding CSS as a string is rather primitive. Encoding JavaScript with strings is even worse for two reasons: first, we are more likely to need to manually escape characters such as `"`; second, we need to use a CDATA object, because most JavaScript code uses characters that "need" to be escaped in XML, such as `&`, but most browsers will fail if these characters are entity-encoded. These are all problems that go away with templates.

Some of these problems go away by using here strings, as described in the documentation on §1.3.7 “Reading Strings”.

Before moving to templates, let’s look at the logic functions:

```
(define (extract-post req)
  (define binds
    (request-bindings req))
  (define title
    (extract-binding/single 'title binds))
```

```

(define body
  (extract-binding/single 'body binds))
(set! posts
  (list* (make-post title body)
    posts))
(send/suspend
  (lambda (k-url)
    (template "Posted" (blog-posted title body k-url))))
(display-posts))

(define (display-posts)
  (extract-post
    (send/suspend
      (lambda (k-url)
        (template "Posts" (blog-posts k-url))))))

(define (start req)
  (display-posts))

```

To use templates, we need only change `template`, `blog-posted`, and `blog-posts`:

```

(define (template section body)
  (response/full
    200 #"Okay"
    (current-seconds) TEXT/HTML-MIME-TYPE
    empty
    (list (string->bytes/utf-8 (include-template "blog.html")))))

(define (blog-posted title body k-url)
  (include-template "blog-posted.html"))

(define (blog-posts k-url)
  (include-template "blog-posts.html"))

```

Each of the templates are given below:

"blog.html":

```

<html>
<head>
  <title>Al's Church: @|section|</title>
  <style type="text/css">
    body {
      margin: 0px;
      padding: 10px;
    }
  </style>
</head>
<body>
  @|body|
</body>
</html>

```

```

    }

    #main {
        background: #dddddd;
    }
</style>
</head>
<body>
    <script type="text/javascript">
        var gaJsHost = (("https:" == document.location.protocol) ?
            "https://ssl." : "http://www.");
        document.write(unescape("%3Cscript src='" + gaJsHost +
            "google-analytics.com/ga.js'
            type='text/javascript'%3E%3C/script%3E"));
    </script>
    <script type="text/javascript">
        var pageTracker = _gat._getTracker("UA-YYYYYYY-Y");
        pageTracker._trackPageview();
    </script>

    <h1>Al's Church: @|section|</h1>
    <div id="main">
        @body
    </div>
</body>
</html>

```

Notice that this part of the presentation is much simpler, because the CSS and JavaScript can be included verbatim, without resorting to any special escape-escaping patterns. Similarly, since the `body` is represented as a string, there is no need to remember if splicing is necessary.

"blog-posted.html":

```

<h2>@|title|</h2>
<p>@|body|</p>

<h1><a href="@|k-url|">Continue</a></h1>

```

"blog-posts.html":

```

@in[p posts]{
    <h2>@(post-title p)</h2>
    <p>@(post-body p)</p>
}

```

```
<h1>New Post</h1>
<form action="@|k-url|">
  <input name="title" />
  <input name="body" />
  <input type="submit" />
</form>
```

Compare this template with the original presentation function: there is no need to worry about managing how lists are nested: the defaults *just work*.

8 Page: Short-hand for Common Patterns

```
(require web-server/page)      package: web-server-lib
```

The Web Server provides a simple utility library for building Web applications that consist mostly of `send/suspend/dispatch`-created pages and request handling.

Most Web applications rely heavily on `send/suspend/dispatch` and typically use the pattern:

```
(send/suspend/dispatch
  (λ (my-embed/url)
    .... (my-embed/url other-page) ....))
```

```
(page e ...)
```

The `page` macro automates this by expanding `(page e ...)` to a usage of `send/suspend/dispatch` where the syntax parameter `embed/url` is bound to the argument of `send/suspend/dispatch`.

```
embed/url
```

When used inside `page` syntactically, a rename transformer for the procedure embedding function; otherwise, a syntax error.

A simple example:

```
(page
  (response/xexpr
    `(html
      (body
        (a ([href
              ,(embed/url
                (λ (req)
                  "You clicked!"))])
            "Click me")))))
```

Similarly, many Web applications make use almost exclusively of functions that are arguments to `embed/url` and immediately invoke `send/suspend/dispatch`.

```
(lambda/page formals e ...)
(define/page (id . formals) e ...)
```

The `lambda/page` and `define/page` automate this by expanding to functions that accept a request as the first argument (followed by any arguments specified in *formals*) and immediately wrap their body in `page`. This functions also cooperate with `get-binding` by binding the request to the `current-request` parameter.

The binding interface of `web-server/http` is powerful, but subtle to use conveniently due to its protection against hostile clients.

```
(current-request) → request?
(current-request req) → void?
  req : request?
binding-id/c : contract?
binding-format/c : contract?
(get-binding id [req #:format format])
→ (or/c false/c string? bytes? binding?)
  id : binding-id/c
  req : request? = (current-request)
  format : binding-format/c = 'string
(get-bindings id [req #:format format])
→ (listof (or/c string? bytes? binding?))
  id : binding-id/c
  req : request? = (current-request)
  format : binding-format/c = 'string
```

The `get-binding(s)` interface attempts to resolve this by providing a powerful interface with convenient defaults.

`get-binding` extracts the first binding of a form input from a request, while `get-bindings` extracts them all.

They accept a form identifier (`id`) as either a byte string, a string, or a symbol. In each case, the user input is compared in a case-sensitive way with the form input.

They accept an optional request argument (`req`) that defaults to the value of the `current-request` parameter used by `lambda/page` and `define/page`.

Finally, they accept an optional keyword argument (`format`) that specifies the desired return format. The default, `'string`, produces a UTF-8 string (or `#f` if the byte string cannot be converted to UTF-8.) The `'bytes` format always produces the raw byte string. The `'file` format produces the file upload content (or `#f` if the form input was not an uploaded file.) The `'binding` format produces the binding object.

9 Testing Servlets

```
(require web-server/test)      package: web-server-lib
```

The Web Server provides a simple facility for writing tests for Web servlets and dispatchers.

The core functionality allows a request to be sent to the servlet and the response captured:

```
tester/c : contract?
```

This contract is equivalent to

```
(->* ()
  ((or/c string? url? request? false/c)
   (listof binding?)
   #:raw? boolean?
   #:headers? boolean?)
 (or/c bytes?
  xexpr?
  (cons/c bytes?
   (or/c bytes?
    xexpr?))))
```

It represents a function that accepts a request and returns the answer the servlet for that request. This interaction function has many possible calling patterns:

- No arguments: a call to the root URL path with no bindings.
- At least one argument: this may be a string, URL, or a request data structure.
- Two arguments: the first argument must be a string or a URL, but the second argument can specify the request bindings.
- The optional `#:raw?` keyword controls whether an X-expression or a byte string is returned as a result.
- The optional `#:headers?` keyword controls whether the headers are included in the return value as a byte string. When this is used, the two returns are returned in a cons.

```
(make-servlet-tester servlet) → tester/c
servlet : (-> request?
           can-be-response?)
```

This function accepts a servlet function and provides a tester function as described above. It is equivalent to `(make-dispatcher-tester (dispatch/servlet servlet))`, so if you need custom arguments to `dispatch/servlet`, use `make-dispatcher-tester`.

```
(make-dispatcher-tester d) → tester/c
d : dispatcher/c
```

This function accepts a dispatcher and provides a tester function as described above.

This facility is designed to be used in concert with a technique of extracting continuation URLs and relevant values; `xml/path` is one way to do this. Here is an extended example that tests an Add-Two-Numbers.com:

```
(define (test-add-two-numbers -s>)
  (define x (random 500))
  (define xs (string->bytes/utf-8 (number->string x)))
  (define y (random 500))
  (define ys (string->bytes/utf-8 (number->string y)))

  (define r0 (-s>))
  (define k0 (se-path* '(form #:action) r0))
  (define i0 (se-path* '(form input #:name) r0))
  (define r1
    (-s> (format "~a?~a=~a" k0 i0 xs)
         (list (make-binding:form (string->bytes/utf-8 i0) xs))))
  (define k1 (se-path* '(form #:action) r1))
  (define i1 (se-path* '(form input #:name) r1))
  (define r2
    (-s> (format "~a?~a=~a" k1 i1 ys)
         (list (make-binding:form (string->bytes/utf-8 i1) ys))))
  (define n (se-path* '(p) r2))
  (check-equal? n
    (format "The answer is ~a" (+ x y))))

(require
  (prefix-in ex:add1: web-server/default-web-
    root/htdocs/servlets/examples/add)
  (prefix-in ex:add2: web-server/default-web-
    root/htdocs/servlets/examples/add-v2))

(test-add-two-numbers
  (make-servlet-tester ex:add1:start))
(test-add-two-numbers
  (make-servlet-tester ex:add2:start))
```

10 Troubleshooting and Tips

10.1 Why is my servlet failing with a `can-be-response?` contract violation after updating Racket?

After 5.0.2, the Web Server had a backwards incompatible change that prevents X-expressions and lists of bytes from being directly returned from servlets. Please read "[PLTHOME/collects/web-server/compat/0/README](#)" to learn about porting your servlets forward. Don't worry. It's easy.

10.2 Why are my templates not updating on the server when I change the file on disk?

Templates are compiled into your application, so when you change them there is no connection between that change in the filesystem and the compiled bytecode that is already loaded in a running Web server process. For more discussion, see §10.4 “Why are my stateful servlets not updating on the server when I change the file on disk?”.

10.3 Why are templates compiled into programs?

Since templates can include arbitrary Racket code, macros, etc and refer to arbitrary identifiers, `include-template` is really just an obscured `require`.

10.4 Why are my stateful servlets not updating on the server when I change the file on disk?

If you are using `serve/servlet`, it starts a Web server that directly references a closure that has no connection to some file on the disk.

If you are using the command-line tool, or configuration file, then by default, the server uses `make-cached-url->servlet` to load servlets from the disk. As it loads them, they are cached and the disk is not referred to for future requests. This ensures that there is a single namespace for each servlet, so that different instances can share resources, such as database connections, and communicate through the store. The default configuration of the server (meaning the dispatcher sequence used when you load a configuration file) provides a special URL to localhost that will reset the cache: `"/conf/refresh-servlets"`.

If you want the server to reload your changed servlet code, then GET this URL and the server will reload the servlet on the next request. However, you may be surprised by what happens

on the next request. For more discussion, see §10.5 “After refreshing my stateful servlet, old captured continuations don’t change or old global effects are gone. Why?”.

10.5 After refreshing my stateful servlet, old captured continuations don’t change or old global effects are gone. Why?

Every load of your servlet is in a fresh namespace. When you refresh, a new namespace without the old effects is created. Old captured continuations refer to the original namespace and will never update. It is impossible, in general, to port a continuation from one namespace to another, because the code could be arbitrarily different.

10.6 How are stateless servlets different from stateful servlets vis a vis refreshing?

Continuations are serialized with a hash that ensures that any source code modifications makes all the old continuations incompatible for the same reason native continuations naturally are.

However, this hash only protects against changes in a single source file. Therefore if you modularize your application, then only continuations that refer to changed source files will be incompatible. For example, if you put all your templates in a single module, then it can change without invalidating old continuations.

10.7 What special considerations are there for security with the Web Server?

The biggest problem is that a naive usage of continuations will allow continuations to subvert authentication mechanisms. Typically, all that is necessary to execute a continuation is its URL. Thus, URLs must be as protected as the information in the continuation.

Consider if you link to a public site from a private continuation URL: the `Referrer` field in the new HTTP request will contain the private URL. Furthermore, if your HTTP traffic is in the clear, then these URLs can be easily poached.

One solution to this is to use a special cookie as an authenticator. This way, if a URL escapes, it will not be able to be used, unless the cookie is present. For advice about how to do this well, see *Dos and Don’ts of Client Authentication on the Web* from the MIT Cookie Eaters.

Note: It may be considered a great feature that URLs can be shared this way, because delegation is easily built into an application via URLs.

10.8 My browser displays my page strangely: my CSS is ignored, sections are missing, etc.

Most Web Server developers use X-expressions for representing the HTML of their page. However, X-expressions only represent XML and HTML is not exactly XML. This is a common source of problems.

For example, XML allows the "empty tag shorthand", e.g. `<img src='...' />`, on every tag, while HTML occasionally requires an end tag, e.g. `TEXTAREA`. Similarly, XML allows an end tag, e.g. `<img src='...'></img>`, on every tag, while HTML occasionally forbids an end tag, e.g. `IMG`. (Of course, browsers do not necessarily implement their HTML parsing as specified and may be more or less lenient towards XML-like HTML, so your test browser may not treat these forms as problematic.)

Since the Web Server uses `xml` to format X-expressions, it inherits `xml`'s default rendering behavior in general and its use of "empty tag shorthand" in particular. `xml`'s default is always use the shorthand with the tags from `html-empty-tags` and never otherwise. This list should contain the W3C's approved list. You can change it with the `empty-tag-shorthand` parameter.

You can also change your X-expression so that an end tag is forced. For example, `'(textarea [(name "text")])'` renders as `<textarea name="text" />`, while `'(textarea [(name "text")]) ""'` renders as `<textarea name="text"></textarea>`, because of the string content in the X-expression. In this case, the end tag will always be present regardless of the value of `empty-tag-shorthand`. It is not possible to force the other possibility; i.e., never include an end tag.

You may think the Web Server could do a better job advertising that the contents it serves is more like XML by default. Unfortunately, browser support for such advertisement is lacking. You can use `response/xexpr` to easily customize your application's MIME type and response headers.

Finally, you may find Web browser inspectors such as the Safari Inspector, Firebug, and the Google Chrome error console to be useful tools in identifying offending tags.

10.9 How do I use templates "dynamically"?

A common feature request is to include one template in another dynamically. It should hopefully be obvious that `include-template` can be included in a template to include a *static* sub-template. For example,

```
(include-template "posts.html")
```

may appear inside the `"blog.html"` template. But you will quickly find that `(include-`

template *expr*) will fail when *expr* is not syntactically a path, e.g.:

```
....  
(include-template (if logged-in?  
                      "user-info.html"  
                      "auth.html"))  
....
```

What is the solution? The templating system already allows you to parameterize templates so particular components come from the including scope. There is no reason those values can not be the results of other templates. In the previous example, suppose the includer was

```
(define (main-page logged-in?)  
  (include-template "site.html"))
```

We could change it to:

```
(define (main-page logged-in?)  
  (define user-content  
    (if logged-in?  
        (include-template "user-info.html")  
        (include-template "auth.html")))  
  (include-template "site.html"))
```

and "site.html" to:

```
....  
user-content  
....
```

This allows you to do the same thing but is safer and more efficient: safer because there is no way to include templates that are not named by the programmer and more efficient because all the templates are compiled (and optimized) with the rest of the code.

If you insist on dynamicism, there is always *eval*.

Index

"unsafe context", 25
#%#, 69
=>*, 70
@ Syntax: @ character, identifiers, and spaces, 83
[adjust-timeout!](#), 16
After refreshing my stateful servlet, old captured continuations don't change or old global effects are gone. Why?, 99
[any->response](#), 11
API Details, 87
API Reference, 62
[APPLICATION/JSON-MIME-TYPE](#), 47
Authenticated Cookies, 48
Base64 Encoding, 34
[base64-stuffer](#), 34
Basic Authentication, 55
Basic Combinators, 32
Basic Formlet Usage, 67
bi-directional match expander, 65
[bidi-match-going-in?](#), 66
[binding](#) (struct), 38
[binding-format/c](#), 95
[binding-id](#), 38
[binding-id/c](#), 95
[binding:file](#) (struct), 39
[binding:file-content](#), 39
[binding:file-filename](#), 39
[binding:file-headers](#), 39
[binding:file/port-in](#), 39
[binding:file/port?](#), 39
[binding:file?](#), 39
[binding:form](#) (struct), 39
[binding:form-value](#), 39
[binding:form?](#), 39
[binding?](#), 38
Bindings, 42
[bindings-assq](#), 40
[bindings-assq-all](#), 40
Built-in URL patterns, 65
[button](#), 77
[call-with-serializable-current-continuation](#), 26
[can-be-response?](#), 11
[checkbox](#), 74
[checkbox-group](#), 75
[clear-continuation-table!](#), 15
[client-cookie](#) (struct), 52
[client-cookie-domain](#), 52
[client-cookie-name](#), 52
[client-cookie-path](#), 52
[client-cookie-value](#), 52
[client-cookie?](#), 52
Command-line Tools, 9
[container?](#), 64
Containers, 64
Continuation Managers, 18
[continuation-url?](#), 16
Conversion Example, 88
[cookie->header](#), 48
[create-LRU-manager](#), 21
[create-none-manager](#), 20
[create-timeout-manager](#), 20
[cross](#), 71
[cross*](#), 71
[current-request](#), 95
[current-servlet-continuation-expiration-handler](#), 15
[default](#), 78
[default-stuffer](#), 37
define-bidi-match-expander, 65
define-coercion-match-expander, 66
define-container, 64
define-native, 27
define/page, 94
Digest Authentication, 56
[dir-store](#), 35
dispatch-case, 64
dispatch-rules, 63
dispatch-rules!, 64
dispatch-rules+applies, 63
dispatch-url, 64
Dynamic, 81

- Dynamic Syntactic Shorthand, 69
- [embed-formlet](#), 79
- [embed/url](#), 94
- Empty Responses, 58
- Escaping, 86
- Examples, 3
- [exists-binding?](#), 42
- [exn:fail:servlet-manager:no-continuation](#) (struct), 20
- [exn:fail:servlet-manager:no-continuation-expiration-handler](#), 20
- [exn:fail:servlet-manager:no-continuation?](#), 20
- [exn:fail:servlet-manager:no-instance](#) (struct), 19
- [exn:fail:servlet-manager:no-instance-expiration-handler](#), 19
- [exn:fail:servlet-manager:no-instance?](#), 19
- Extending `web-server/dispatch`, 65
- [extract-binding/single](#), 42
- [extract-bindings](#), 42
- Extracting Cookies, 52
- File Boxes, 29
- [file-box](#), 29
- [file-box-set!](#), 30
- [file-box-set?](#), 30
- [file-box?](#), 29
- [file-unbox](#), 29
- [file-upload](#), 76
- `formlet`, 68
- formlet*, 67
- `formlet*`, 69
- [formlet*/c](#), 71
- [formlet-display](#), 72
- [formlet-process](#), 72
- [formlet/c](#), 71
- Formlets and Stateless Servlets, 79
- Formlets: Functional Form Abstraction, 67
- Full API, 5
- Functional Usage, 70
- General, 18
- [get-binding](#), 95
- [get-bindings](#), 95
- Gotchas, 83
- GZip Compression, 34
- [gzip-stuffer](#), 34
- Hash-addressed Storage, 35
- [hash-fun/c](#), 35
- [hash-stuffer](#), 35
- [header](#) (struct), 38
- [header-field](#), 38
- [header-value](#), 38
- [header?](#), 38
- [headers-assq](#), 38
- [headers-assq*](#), 38
- Helpers, 36
- [hidden](#), 76
- [HMAC-SHA1](#), 36
- HMAC-SHA1 Signing, 36
- [HMAC-SHA1-stuffer](#), 36
- How are stateless servlets different from stateful servlets vis a vis refreshing?, 99
- How do I use templates “dynamically”?, 100
- HTTP Responses, 87
- HTTP: Hypertext Transfer Protocol, 38
- [id-stuffer](#), 33
- [img](#), 76
- Imperative Dispatch Containers, 64
- `in`, 88
- [include-template](#), 87
- [include-template/xml](#), 88
- [input](#), 73
- [input-int](#), 79
- [input-string](#), 78
- [input-symbol](#), 79
- Instant Servlets, 2
- `integer-arg`, 65
- [interface-version](#), 23
- [interface-version](#), 10
- [is-url-too-big?](#), 36
- Iteration, 84
- JSON Support, 59

- Key/Value Storage, 34
- lambda/page, 94
- logout-id-cookie, 52
- LRU, 21
- make-basic-auth-header, 55
- make-binding, 38
- make-binding:file, 39
- make-binding:file/port, 39
- make-binding:form, 39
- make-check-digest-credentials, 57
- make-client-cookie, 52
- make-coerce-safe?, 66
- make-cookie, 47
- make-default-stuffer, 37
- make-digest-auth-header, 56
- make-dispatcher-tester, 97
- make-exn:fail:servlet-manager:no-continuation, 20
- make-exn:fail:servlet-manager:no-instance, 19
- make-header, 38
- make-id-cookie, 49
- make-input, 73
- make-input*, 73
- make-manager, 19
- make-request, 41
- make-secret-salt/file, 52
- make-servlet-tester, 96
- make-soft-state, 31
- make-store, 35
- make-stuffer, 32
- make-threshold-LRU-manager, 22
- make-web-cell, 17
- make-web-cell, 29
- make-web-parameter, 30
- manager, 23
- manager (struct), 19
- manager, 10
- manager, 18
- manager-adjust-timeout!, 19
- manager-clear-continuations!, 19
- manager-continuation-lookup, 19
- manager-continuation-peek, 19
- manager-continuation-store!, 19
- manager-create-instance, 19
- manager?, 19
- md5-stuffer, 35
- multiselect-input, 77
- My browser displays my page strangely: my CSS is ignored, sections are missing, etc., 100
- Native Interfaces, 27
- native->serial, 26
- No Continuations, 20
- no-web-browser, 2
- number-arg, 65
- page, 94
- Page: Short-hand for Common Patterns, 94
- password->digest-HA1, 57
- password-input, 74
- permanently, 54
- Placing Cookies, 47
- plt-web-server, 9
- Post-Redirect-Get, 54
- Predefined Formlets, 72
- pure, 71
- radio, 74
- radio-group, 75
- real-arg, 65
- Redirect, 53
- redirect-to, 53
- redirect/get, 14
- redirect/get, 28
- redirect/get/forget, 15
- redirection status, 54
- redirection-status?, 54
- request (struct), 41
- request->basic-credentials, 55
- request->digest-credentials, 56
- request-bindings, 42
- request-bindings/raw, 41
- request-bindings/raw-promise, 41
- request-client-ip, 41
- request-cookies, 53

- [request-headers](#), 42
- [request-headers/raw](#), 41
- [request-host-ip](#), 41
- [request-host-port](#), 41
- [request-id-cookie](#), 50
- [request-method](#), 41
- [request-post-data/raw](#), 41
- [request-uri](#), 41
- [request?](#), 41
- Requests, 38
- [required](#), 78
- [reset](#), 76
- [response](#) (struct), 43
- [response-code](#), 43
- [response-code/c](#), 43
- [response-headers](#), 43
- [response-message](#), 43
- [response-mime](#), 43
- [response-output](#), 43
- [response-seconds](#), 43
- [response/empty](#), 59
- [response/full](#), 44
- [response/jsexpr](#), 59
- [response/output](#), 46
- [response/xexpr](#), 58
- [response?](#), 43
- Responses, 43
- Responses, 11
- Running Web Servlets, 2
- [see-other](#), 54
- [select-input](#), 77
- [send/back](#), 11
- [send/finish](#), 14
- [send/formlet](#), 79
- [send/forward](#), 14
- [send/suspend](#), 12
- [send/suspend](#), 28
- [send/suspend/dispatch](#), 12
- [send/suspend/dispatch](#), 28
- [send/suspend/hidden](#), 28
- [send/suspend/url](#), 28
- [send/suspend/url](#), 12
- [send/suspend/url/dispatch](#), 28
- [send/suspend/url/dispatch](#), 14
- [serial->native](#), 26
- Serializable Continuations, 26
- Serialization, 34
- [serialize-stuffer](#), 34
- [serve/dispatch](#), 64
- [serve/servlet](#), 6
- [servlet-prompt](#), 16
- [set-any->response!](#), 11
- Simple Single Servlet Servers, 2
- Soft State, 31
- [soft-state](#), 31
- [soft-state-ref](#), 31
- [soft-state?](#), 31
- [start](#), 10
- [start](#), 23
- Stateful Servlets, 10
- Stateless Servlets, 4
- Stateless Servlets, 23
- Stateless Web Cells, 29
- Stateless Web Interaction, 27
- Stateless Web Parameters, 30
- Static, 81
- Static Syntactic Shorthand, 68
- [static-files-path](#), 2
- [store](#) (struct), 35
- [store-read](#), 35
- [store-write](#), 35
- [store?](#), 35
- [string-arg](#), 65
- [struct:binding](#), 38
- [struct:binding:file](#), 39
- [struct:binding:form](#), 39
- [struct:client-cookie](#), 52
- [struct:exn:fail:servlet-manager:no-continuation](#), 20
- [struct:exn:fail:servlet-manager:no-instance](#), 19
- [struct:header](#), 38
- [struct:manager](#), 19
- [struct:request](#), 41

- [struct:response](#), 43
- [struct:store](#), 35
- [struct:stuffer](#), 32
- [stuffer](#) (struct), 32
- [stuffer](#), 23
- [stuffer-chain](#), 33
- [stuffer-compose](#), 33
- [stuffer-if](#), 33
- [stuffer-in](#), 32
- [stuffer-out](#), 32
- [stuffer-sequence](#), 33
- [stuffer/c](#), 33
- [stuffer?](#), 32
- Stuffers, 32
- stuffers*, 32
- [submit](#), 75
- [symbol-arg](#), 65
- [tag-xexpr](#), 72
- Templates: Separation of View, 81
- [temporarily](#), 54
- [temporarily/same-method](#), 54
- [tester/c](#), 96
- Testing Servlets, 96
- [text](#), 72
- [text-input](#), 73
- [TEXT/HTML-MIME-TYPE](#), 46
- [textarea-input](#), 74
- Timeouts, 20
- [to-boolean](#), 78
- [to-number](#), 78
- [to-string](#), 78
- [to-symbol](#), 78
- Troubleshooting and Tips, 98
- URL-Based Dispatch, 61
- Usage Considerations, 24
- [username*realm->digest-HA1/c](#), 56
- [username*realm->password/c](#), 56
- Using [web-server/dispatch](#), 61
- Utilities, 79
- [valid-id-cookie?](#), 51
- Web Applications in Racket, 1
- Web Cells, 16
- Web Interaction, 11
- [web-cell-ref](#), 29
- [web-cell-ref](#), 17
- [web-cell-shadow](#), 17
- [web-cell-shadow](#), 29
- [web-cell?](#), 17
- [web-cell?](#), 29
- [web-parameter?](#), 30
- [web-parameterize](#), 30
- [web-server](#), 23
- [web-server/base](#), 23
- [web-server/dispatch](#), 61
- [web-server/dispatch/extend](#), 65
- [web-server/formlets](#), 67
- [web-server/formlets/dyn-syntax](#), 69
- [web-server/formlets/input](#), 72
- [web-server/formlets/lib](#), 70
- [web-server/formlets/servlet](#), 79
- [web-server/formlets/stateless](#), 79
- [web-server/formlets/syntax](#), 68
- [web-server/formlets/unsafe](#), 79
- [web-server/http](#), 38
- [web-server/http/basic-auth](#), 55
- [web-server/http/bindings](#), 42
- [web-server/http/cookie](#), 47
- [web-server/http/cookie-parse](#), 52
- [web-server/http/digest-auth](#), 56
- [web-server/http/empty](#), 58
- [web-server/http/id-cookie](#), 48
- [web-server/http/json](#), 59
- [web-server/http/redirect](#), 53
- [web-server/http/request-structs](#), 38
- [web-server/http/response-structs](#), 43
- [web-server/http/xexpr](#), 58
- [web-server/insta](#), 2
- [web-server/lang/abort-resume](#), 26
- [web-server/lang/file-box](#), 29
- [web-server/lang/native](#), 27
- [web-server/lang/soft](#), 31
- [web-server/lang/stuff-url](#), 36
- [web-server/lang/web](#), 27

- [web-server/lang/web-cells](#), 29
- [web-server/lang/web-param](#), 30
- [web-server/managers/lru](#), 21
- [web-server/managers/manager](#), 18
- [web-server/managers/none](#), 20
- [web-server/managers/timeout](#)s, 20
- [web-server/page](#), 94
- [web-server/servlet](#), 10
- [web-server/servlet-env](#), 2
- [web-server/servlet/servlet-structs](#), 11
- [web-server/servlet/web](#), 11
- [web-server/servlet/web-cells](#), 16
- [web-server/stuffers](#), 32
- [web-server/stuffers/base64](#), 34
- [web-server/stuffers/gzip](#), 34
- [web-server/stuffers/hash](#), 35
- [web-server/stuffers/hmac-sha1](#), 36
- [web-server/stuffers/serialize](#), 34
- [web-server/stuffers/store](#), 34
- [web-server/stuffers/stuffer](#), 32
- [web-server/templates](#), 81
- [web-server/test](#), 96
- What special considerations are there for security with the Web Server?, 99
- Why are my stateful servlets not updating on the server when I change the file on disk?, 98
- Why are my templates not updating on the server when I change the file on disk?, 98
- Why are templates compiled into programs?, 98
- Why is my servlet failing with a [can-be-response?](#) contract violation after updating Racket?, 98
- [with-errors-to-browser](#), 15
- X-expression Support, 58
- [xexpr-forest/c](#), 70
- [xml](#), 72
- [xml-forest](#), 72