

Web Server: PLT HTTP Server

Version 4.1.5

Jay McCarthy <jay at plt-scheme dot org>

March 21, 2009

The Web Server collection provides libraries that can be used to develop Web applications in Scheme.

Contents

1	Running the Web Server	7
1.1	Instant Servlets	7
1.1.1	Customization API	7
1.2	Simple Single Servlet Servers	7
1.3	Command-line Tools	13
1.4	Functional	13
2	Writing Servlets	16
2.1	Defining a Servlet	16
2.1.1	Version 1 Servlets	16
2.1.2	Version 2 Servlets	16
2.1.3	Stateless Servlets	17
2.2	APIs	19
2.2.1	Standard API	20
2.2.2	Stateless API	20
2.3	Common Contracts	20
2.4	HTTP	21
2.4.1	Requests	22
2.4.2	Bindings	24
2.4.3	Responses	25
2.4.4	Placing Cookies	27
2.4.5	Extracting Cookies	28
2.4.6	Redirect	29
2.4.7	Basic Authentication	30

2.4.8	Digest Authentication	31
2.5	Web Interaction	33
2.6	Stateless Web Interaction	37
2.6.1	Low Level	37
2.6.2	High Level	37
2.7	Web Cells	38
2.7.1	Stateless Web Cells	40
2.8	File Boxes	40
2.9	Stateless Web Parameters	41
2.10	URL-Based Dispatch	42
2.10.1	Using <code>web-server/dispatch</code>	42
2.10.2	API Reference	44
2.10.3	Built-in URL patterns	45
2.10.4	Extending <code>web-server/dispatch</code>	46
2.11	Formlets	47
2.11.1	Basic Formlet Usage	47
2.11.2	Syntactic Shorthand	48
2.11.3	Functional Usage	49
2.11.4	Predefined Formlets	51
2.11.5	Utilities	53
2.12	Templates	53
2.12.1	Static	54
2.12.2	Dynamic	54
2.12.3	Gotchas	56
2.12.4	HTTP Responses	58

2.12.5	API Details	58
2.12.6	Conversion Example	59
2.13	Continuation Managers	63
2.13.1	General	63
2.13.2	No Continuations	64
2.13.3	Timeouts	65
2.13.4	LRU	66
2.14	Stuffers	67
2.14.1	Basic Combinators	67
2.14.2	Serialization	68
2.14.3	Base64 Encoding	68
2.14.4	GZip Compression	69
2.14.5	Key/Value Storage	69
2.14.6	Hash-addressed Storage	70
2.14.7	HMAC-SHA1 Signing	70
2.14.8	Helpers	71
3	Extending the Web Server	73
3.1	Configuration	73
3.1.1	Configuration Table Structure	73
3.1.2	Configuration Table	75
3.1.3	Servlet Namespaces	77
3.1.4	Standard Responders	78
3.2	Setting Up Servlets	80
3.2.1	Internal Servlet Representation	81
3.3	Dispatchers	82

3.3.1	General	82
3.3.2	Mapping URLs to Paths	83
3.3.3	Sequencing	84
3.3.4	Timeouts	84
3.3.5	Lifting Procedures	84
3.3.6	Filtering Requests	85
3.3.7	Procedure Invocation upon Request	85
3.3.8	Logging	86
3.3.9	Password Protection	87
3.3.10	Virtual Hosts	88
3.3.11	Serving Files	88
3.3.12	Serving Servlets	89
3.3.13	Statistics	90
3.3.14	Limiting Requests	90
3.4	Web Config Unit	92
3.4.1	Configuration Signature	92
3.4.2	Configuration Units	93
3.5	Web Server Unit	93
3.5.1	Signature	93
3.5.2	Unit	94
3.6	Internal	95
3.6.1	Timers	95
3.6.2	Connection Manager	96
3.6.3	Dispatching Server	97
3.6.4	Serializable Closures	99

3.6.5	Cache Table	100
3.6.6	MIME Types	101
3.6.7	Serialization Utilities	101
3.6.8	URL Param	101
3.6.9	GZip	102
3.6.10	Miscellaneous Utilities	102
4	Troubleshooting and Tips	106
4.1	Why are my servlets not updating on the server when I change the code on disk?	106
4.2	What special considerations are there for security with the Web Server? . . .	106
4.3	How do I use Apache with the PLT Web Server?	106
4.4	IE ignores my CSS or behaves strange in other ways	107
4.5	Can the server create a PID file?	107
4.6	How do I set up the server to use HTTPS?	107
4.7	How do I limit the number of requests serviced at once by the Web Server? . . .	108
5	Acknowledgements	109
	Index	110

1 Running the Web Server

There are a number of ways to run the Web Server. They are given in order of simplest to most advanced.

1.1 Instant Servlets

```
#lang web-server/insta
```

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

```
#lang web-server/insta

(define (start request)
  '(html (head (title "Hello world!"))
        (body (p "Hey out there!"))))
```

And press Run. A Web browser will open up showing your new servlet.

Behind the scenes, DrScheme 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.

1.1.1 Customization API

```
(require web-server/insta/insta)
```

The following API is provided to customize the server instance:

```
(no-web-browser) → void
```

Calling this will instruct DrScheme 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 stylesheet and images, from `path`.

1.2 Simple Single Servlet Servers

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

The Web Server provides a way to quickly configure and start a server instance.

Here's a simple example:

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

(define (my-app request)
  '(html (head (title "Hello world!"))
        (body (p "Hey out there!"))))

(serve/servlet my-app)
```

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

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

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

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

Suppose you wanted it to capture top-level requests:

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

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. Notice that you may pass any number of extra paths.

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

```
#lang scheme
(require "servlet.ss"
         web-server/servlet-env)

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

Note: If you put the call to `serve/servlet` in the module like normal, strange things will happen because of the way the top-level interacts with continuations. (Read: Don't do it.)

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

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

```
(serve/servlet
 start
 [#:command-line? command-line?
 #:launch-browser? launch-browser?
 #:quit? quit?
 #:banner? banner?
 #:listen-ip listen-ip
 #:port port
 #:ssl? ssl?
 #: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
 #:mime-types-path mime-types-path
 #:log-file log-file
 #:log-format log-format])
→ void
start : (request? . -> . response/c)
command-line? : boolean? = #f
launch-browser? : boolean? = (not command-line?)
quit? : boolean? = (not command-line?)
```

```

banner? : boolean? = (not command-line?)
listen-ip : (or/c false/c string?) = "127.0.0.1"
port : number? = 8000
ssl? : boolean? = #f
servlet-path : string? = "/servlets/standalone.ss"
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 (* 1024 1024 64))
servlet-namespace : (listof module-path?) = empty
server-root-path : path-string? = default-server-root-path
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? . -> . response/c)
                          = (gen-file-not-found-responder
                             (build-path
                              server-root-path
                              "conf"
                              "not-found.html"))
mime-types-path : path-string? = ...
log-file : (or/c false/c path-string?) = #f
log-format : log-format/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>".

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.

Advanced users may need the following options:

The server listens on *listen-ip* and port *port*.

If *ssl?* is true, then the server runs in HTTPS mode with "*<server-root-path>/server-cert.pem*" and "*<server-root-path>/private-key.pem*" 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 modules specified by *servlet-namespace* are shared with other servlets.

The server files are rooted at *server-root-path* (which is defaultly the distribution root.) 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*.

If a file cannot be found, *file-not-found-responder* is used to generate an error response.

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*.

```
(dispatch/servlet
  start
  [#:regexp regexp
   #:stateless? stateless?
   #:stuffer stuffer
   #:manager manager
   #:namespace namespace
   #:current-directory servlet-current-directory])
→ dispatcher/c
start : (request? . -> . response/c)
regexp : regexp = #rx"
stateless? : boolean? = #f
stuffer : (stuffer/c serializable? bytes?) = default-stuffer
```

```

manager : manager?
        = (make-threshold-LRU-manager #f (* 1024 1024 64))
namespace : (listof module-path?) = empty
servlet-current-directory : path-string? = (current-directory)

```

`serve/servlet` starts a server and uses a particular dispatching sequence. For some applications, this nails down too much, but users are conflicted, because the interface is so convenient. For those users, `dispatch/servlet` does the hardest part of `serve/servlet` and constructs a dispatcher just for the `start` servlet.

The dispatcher responds to requests that match `regexp`. The current directory of servlet execution is `servlet-current-directory`.

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

```
#lang web-server
```

module and `stuffer` is used as the stuffer.

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 modules specified by `servlet-namespace` are shared with other servlets.

```

(serve/launch/wait make-dispatcher
  [#:launch-path launch-path
   #:banner? banner?
   #:listen-ip listen-ip
   #:port port
   #:ssl-keys ssl-keys]) → void
make-dispatcher : (semaphore? . -> . dispatcher/c)
launch-path : (or/c false/c string?) = #f
banner? : boolean? = #f
listen-ip : (or/c false/c string?) = "127.0.0.1"
port : number? = 8000
ssl-keys : (or/c false/c (cons/c path-string? path-string?))
          = #f

```

The other interesting part of `serve/servlet` is its ability to start up a server and immediately launch a browser at it. This is provided by `serve/launch/wait`.

It starts a server using the result of `make-dispatcher` as the dispatcher. `make-dispatcher` is supplied a semaphore that if posted, will cause the server to quit.

If `launch-path` is not false, then a browser is launched with that path appended to the URL to the server itself.

If `banner?` is true, then a banner is printed informing the user of the server's URL.

The server listens on `listen-ip` and port `port`.

If `ssl-keys` is not false, then the server runs in HTTPS mode with `(car ssl-keys)` and `(cdr ssl-keys)` as paths to the certificate and private key.

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 §3.1.2 “Configuration Table”). 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.

To run the web server with MrEd, use

```
mred -l- web-server/gui [-f <file-name> -p <port> -a <ip-address>]
```

1.4 Functional

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

“web-server.ss” provides a number of functions for easing embedding of the Web Server in other applications, or loading a custom dispatcher.

```
(serve
  #:dispatch dispatch
  [#:tcp@ tcp@
  #:port port
  #:listen-ip listen-ip
  #:max-waiting max-waiting
  #:initial-connection-timeout initial-connection-timeout])
```

```

→ (-> void)
dispatch : dispatcher/c
tcp@ : (unit/c (import) (export tcp^)) = raw:tcp@
port : integer? = 80
listen-ip : (or/c string? false/c) = #f
max-waiting : integer? = 40
initial-connection-timeout : integer? = 60

```

Constructs an appropriate `dispatch-server-config^`, invokes the `dispatch-server@`, and calls its `serve` function.

The `#:tcp@` keyword is provided for building an SSL server. See §4.6 “How do I set up the server to use HTTPS?”.

Here’s an example of a simple web server that serves files from a given path:

```

(define (start-file-server base)
  (serve
   #:dispatch
   (files:make
    #:url->path (make-url->path base)
    #:path->mime-type
    (lambda (path)
      #"application/octet-stream")))
   #:port 8080))

```

```

(serve/ports
 #:dispatch dispatch
 [#:tcp@ tcp@
 #:ports ports
 #:listen-ip listen-ip
 #:max-waiting max-waiting
 #:initial-connection-timeout initial-connection-timeout])
→ (-> void)
dispatch : dispatcher/c
tcp@ : (unit/c (import) (export tcp^)) = raw:tcp@
ports : (listof integer?) = (list 80)
listen-ip : (or/c string? false/c) = #f
max-waiting : integer? = 40
initial-connection-timeout : integer? = 60

```

Calls `serve` multiple times, once for each port, and returns a function that shuts down all of the server instances.

```

(serve/ips+ports
 #:dispatch dispatch
 [#:tcp@ tcp@
 #:ips+ports ips+ports
 #:max-waiting max-waiting
 #:initial-connection-timeout initial-connection-timeout])
→ (-> void)
dispatch : dispatcher/c
tcp@ : (unit/c (import) (export tcp^)) = raw:tcp@
ips+ports : (listof (cons/c (or/c string? false/c) (listof integer?)))
           = (list (cons #f (list 80)))
max-waiting : integer? = 40
initial-connection-timeout : integer? = 60

```

Calls `serve/ports` multiple times, once for each `ip`, and returns a function that shuts down all of the server instances.

```

(serve/web-config@ config@ [#:tcp@ tcp@]) → (-> void)
config@ : (unit/c (import) (export web-config^))
tcp@ : (unit/c (import) (export tcp^)) = raw:tcp@

```

Starts the Web Server with the settings defined by the given `web-config^` unit.

It is very useful to combine this with `configuration-table->web-config@` and `configuration-table-sexpr->web-config@`:

```

(serve/web-config@
 (configuration-table->web-config@
  default-configuration-table-path))

```

```

(do-not-return) → void

```

This function does not return. If you are writing a script to load the Web Server you are likely to want to call this functions at the end of your script.

2 Writing Servlets

2.1 Defining a Servlet

A *servlet* is a module with particular exports. There three kinds of servlets.

2.1.1 Version 1 Servlets

```
interface-version : (one-of/c 'v1)
```

This indicates that the servlet is a version one servlet.

```
timeout : integer?
```

This number is used as the `continuation-timeout` argument to a timeout-based continuation manager used for this servlet. (See §2.13.3 “Timeouts”.) (i.e., you do not have a choice of the manager for this servlet and will be given a timeout-based manager.)

```
(start initial-request) → response/c  
  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 1 module:

```
#lang scheme  
  
(define interface-version 'v1)  
(define timeout (* 60 60 24))  
(define (start req)  
  '(html (head (title "Hello World!"))  
        (body (h1 "Hi Mom!"))))
```

These servlets should use the `web-server/servlet` API.

2.1.2 Version 2 Servlets

```
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 §2.13 “Continuation Managers” for options.

`(start initial-request) → response/c`
`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 scheme
(require web-server/managers/none)

(define interface-version 'v2)
(define manager
  (create-none-manager
   (lambda (req)
     '(html (head (title "No Continuations Here!"))
            (body (h1 "No Continuations Here!"))))))
(define (start req)
  '(html (head (title "Hello World!"))
         (body (h1 "Hi Mom!"))))
```

These servlets should use the `web-server/servlet` API.

2.1.3 Stateless Servlets

`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`.

```
(start initial-request) → response/c
  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
(define interface-version 'stateless)
(define stuffer
  (stuffer-chain
   serialize-stuffer
   (md5-stuffer (build-path (find-system-path 'home-dir) ".urls"))))
(define (start req)
  `(html (body (h2 "Look ma, no state!"))))
```

The `web-server` language automatically provides the `web-server/lang/lang-api` API.

Usage Considerations

```
#lang web-server
```

A 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. (`"lang/elim-letrec.ss"`)
- The program is converted into ANF (Administrative Normal Form), making all continuations explicit. (`"lang/anormal.ss"`)
- All continuations (and other continuations marks) are recorded in the continuation marks of the expression they are the continuation of. (`"lang/elim-callcc.ss"`)
- All calls to external modules are identified and marked. (`"lang/elim-callcc.ss"`)
- All uses of `call/cc` are removed and replaced with equivalent gathering of the continuations through the continuation-marks. (`"lang/elim-callcc.ss"`)
- The program is defunctionalized with a serializable data-structure for each anonymous lambda. (`"lang/defun.ss"`)

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. Thus, your servlet has no

cost to the server other than execution. This is very attractive if you've used Scheme servlets and had memory problems.

This process IS defined on all of PLT Scheme and occurs AFTER macro-expansion, so you are free to use all interesting features of PLT Scheme. 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 PLT Scheme. However, it is theoretically a benefit. The main implications would be due to optimizations MzScheme attempts to perform that will no longer apply. Ideally, your program should be optimized first.

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 a bug!

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 capturable—must not be in the context of another module. For example, the following will not work:

```
(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.

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 before the server is restarted or the memory is garbage collected.

This process is derived from the paper *Continuations from Generalized Stack Inspection*. We thank Greg Pettyjohn for his initial implementation of this algorithm.

2.2 APIs

There are two API sets provided by the Web Server. One is for standard servlets, the other is for stateless servlets.

2.2.1 Standard API

(require web-server/servlet)

This API provides:

- net/url,
- web-server/servlet/web-cells,
- web-server/dispatch,
- web-server/http/bindings,
- web-server/http,
- web-server/servlet/servlet-structs, and
- web-server/servlet/web.

2.2.2 Stateless API

(require web-server/lang/lang-api)

This API provides:

- net/url,
- web-server/dispatch,
- web-server/http,
- web-server/stuffers,
- web-server/lang/abort-resume,
- web-server/lang/web,
- web-server/lang/web-cells,
- web-server/lang/web-param, and
- web-server/lang/file-box.

2.3 Common Contracts

```
(require web-server/servlet/servlet-structs)
```

"servlet/servlet-structs.ss" provides a number of contracts for use in servlets.

`k-url?` : `contract?`

Equivalent to `string?`.

Example: `"http://localhost:8080/servlets;1*1*20131636/examples/add.ss"`

`response-generator/c` : `contract?`

Equivalent to `(-> k-url? response/c)`.

Example:

```
(lambda (k-url)
  '(html
    (body
      (a ([href ,k-url])
         "Click Me to Invoke the Continuation!")))))
```

`expiration-handler/c` : `contract?`

Equivalent to `(or/c false/c (-> request? response/c))`.

Example:

```
(lambda (req)
  '(html (head (title "Expired"))
        (body (h1 "Expired")
              (p "This URL has expired. "
                 "Please return to the home page.")))))
```

`embed/url/c` : `contract?`

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

This is what `send/suspend/dispatch` gives to its function argument.

2.4 HTTP

```
(require web-server/http)
```

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

2.4.1 Requests

```
(require web-server/http/request-structs)
```

```
(struct header (field value))  
  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))  
  id : bytes?
```

Represents a binding of `id`.

```
(struct (binding:form binding) (value))  
  value : bytes?
```

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

```
(struct (binding:file binding) (filename headers content))  
  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.)

```
(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`.

```
(struct request (method
                  uri
                  headers/raw
                  bindings/raw
                  post-data/raw
                  host-ip
                  host-port
                  client-ip))
method : bytes?
uri : url?
headers/raw : (listof header?)
bindings/raw : (listof binding?)
post-data/raw : (or/c false/c 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.

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))]]))
```

2.4.2 Bindings

```
(require web-server/http/bindings)
```

These functions, while convenient, could introduce subtle bugs into your application. Examples: that they are case-insensitive could introduce a bug; 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, which are left as is. 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))))
```

2.4.3 Responses

```
(require web-server/http/response-structs)
```

```
(struct response/basic (code message seconds mime headers))
code : number?
message : bytes?
seconds : number?
mime : bytes?
headers : (listof header?)
```

A basic HTTP response containing no body. `code` is the response code, `message` the message, `seconds` the generation time, `mime` the MIME type of the file, and `extras` are the extra headers, in addition to those produced by the server.

Example:

```
(make-response/basic
 301 #"Moved Permanently"
 (current-seconds) TEXT/HTML-MIME-TYPE
 (list (make-header #"Location"
                  #"http://www.plt-scheme.org/downloads")))
```

```
(struct (response/full response/basic) (body))
body : (listof bytes?)
```

As with `response/basic`, except with `body` as the response body.

Example:

```
(make-response/full
 301 #"Moved Permanently"
```

```

(current-seconds) TEXT/HTML-MIME-TYPE
(list (make-header #"Location"
                  #"http://www.plt-scheme.org/downloads"))
(list #"<html><body><p>"
      #"Please go to <a href=\"\"
      #"http://www.plt-scheme.org/downloads"
      #"\">here</a> instead."
      #"</p></body></html>"))

```

```

(struct (response/incremental response/basic) (generator))
generator : (((() () #:rest (listof bytes?) . ->* . any) . -> . any)

```

As with `response/basic`, except with `generator` as a function that is called to generate the response body, by being given an `output-response` function that outputs the content it is called with.

Here is a short example:

```

(make-response/incremental
 200 #"OK" (current-seconds)
 #"application/octet-stream"
 (list (make-header #"Content-Disposition"
                   #"attachment; filename=\"file\""))
 (lambda (send/bytes)
   (send/bytes #"Some content")
   (send/bytes)
   (send/bytes #"Even" #"more" #"content!")
   (send/bytes #"Now we're done")))

```

`response/c` : contract?

Equivalent to `(or/c response/basic? (cons/c bytes? (listof (or/c string? bytes?))) xexpr/c)`.

```

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

xexpr : xexpr/c
code : number? = 200
message : bytes? = #"Okay"
seconds : number? = (current-seconds)

```

```
mime-type : bytes? = TEXT/HTML-MIME-TYPE
headers : (listof header?) = empty
```

Equivalent to

```
(make-response/full
  code message seconds mime-type headers
  (list (string->bytes/utf-8 (xexpr->string xexpr))))
```

```
(normalize-response close? response)
→ (or/c response/full? response/incremental?)
close? : boolean?
response : response/c
```

Coerces *response* into a full response, filling in additional details where appropriate.

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

Equivalent to `"text/html; 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**.

2.4.4 Placing Cookies

```
(require web-server/http/cookie)
```

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
   #:secure? secure?]) → cookie?
name : string?
value : string?
comment : (or/c false/c string?) = #f
domain : (or/c false/c valid-domain?) = #f
max-age : (or/c false/c exact-nonnegative-integer?) = #f
path : (or/c false/c string?) = #f
secure? : (or/c false/c boolean?) = #f
```

Constructs a cookie with the appropriate fields.

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

Constructs a header that sets the cookie.

```
(xexpr-response/cookies cookies xexpr) → response/full?  
cookies : (listof cookie?)  
xexpr : xexpr/c
```

Constructs a response using *xexpr* that sets all the cookies in *cookies*.

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)  
    (xexpr-response/cookies  
      (list time-cookie id-cookie)  
      '(html (head (title "Cookie Example"))  
              (body (h1 "You're cookie'd!"))))))
```

Warning: When using cookies, make sure you follow the advice of the MIT Cookie Eaters, or you will be susceptible to dangerous attacks.

2.4.5 Extracting Cookies

```
(require web-server/http/cookie-parse)
```

```
(struct client-cookie (name value domain path))  
name : string?
```

```
value : string?
domain : (or/c false/c valid-domain?)
path : (or/c false/c string?)
```

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.

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)
  '(html (head (title "Hello!"))
        (body
         (h1 "Hello "
            ,who))))
```

2.4.6 Redirect

```
(require web-server/http/redirect)
```

```
(redirect-to uri
  [perm/temp
   #:headers headers]) → response/c
uri : non-empty-string/c
perm/temp : redirection-status? = temporarily
```

```
headers : (listof header?) = (list)
```

Generates an HTTP response that redirects the browser to *uri*, while including the *headers* in the response.

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

```
(redirection-status? v) → boolean?  
v : any/c
```

Determines if *v* is one of the following values.

```
permanently : redirection-status?
```

A `redirection-status?` for permanent redirections.

```
temporarily : redirection-status?
```

A `redirection-status?` for temporary redirections.

```
see-other : redirection-status?
```

A `redirection-status?` for "see-other" redirections.

2.4.7 Basic Authentication

```
(require web-server/http/basic-auth)
```

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)
     '(html (head (title "Basic Auth Test"))
            (body (h1 "User: " ,(bytes->string/utf-8 user))
                  (h1 "Pass: " ,(bytes->string/utf-8 pass)))))]
    [else
     (make-response/basic
      401 #"Unauthorized" (current-seconds) TEXT/HTML-MIME-TYPE
      (list
       (make-basic-auth-header
        (format "Basic Auth Test: ~a" (gensym))))))]))
```

2.4.8 Digest Authentication

```
(require web-server/http/digest-auth)
```

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 scheme/pretty)

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

(define (start req)
  (match (request->digest-credentials req)
    [#f
     (make-response/basic
      401 #"Unauthorized" (current-seconds) TEXT/HTML-MIME-TYPE
      (list (make-digest-auth-header
              (format "Digest Auth Test: ~a" (gensym))
              private-key opaque)))]
    [alist
     (define check
       (make-check-digest-credentials
```



```

      (password->digest-HA1 (lambda (username realm) "pass"))))
(define pass?
  (check "GET" alist))
'(html (head (title "Digest Auth Test"))
      (body
        (h1 ,(if pass? "Pass!" "No Pass!"))
        (pre ,(pretty-format alist))))))

```

2.5 Web Interaction

```
(require web-server/servlet/web)
```

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

```
(send/back response) → void?
  response : response/c
```

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

Example:

```

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

```

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

Captures the current continuation, stores it with `exp` as the expiration handler, and binds it to a URL. *make-response* is called with this URL and is expected to generate a `response/c`, 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)
  '(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? . -> . response/c)

```

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

```

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

```

Calls `make-response` with a function (`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)
      `(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)
      '(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?) . -> . response/c)

```

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

```

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

```

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 : response/c

```

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

`(redirect/get)` → `request?`

Calls `send/suspend` with `redirect-to`.

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

`(redirect/get/forget)` → `request?`

Calls `send/forward` with `redirect-to`.

`current-servlet-continuation-expiration-handler` : `(parameter/c expiration-handler/c)`

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

Example:

```
(parameterize
  ([current-servlet-continuation-expiration-handler
   (lambda (req)
     '(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 : (response/c . -> . 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.

2.6 Stateless Web Interaction

2.6.1 Low Level

```
(require web-server/lang/abort-resume)
```

```
(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.

2.6.2 High Level

```
(require web-server/lang/web)
```

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

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/c)
```

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

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

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/c)
```

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) → request?  
make-response : (((request? . -> . any) . -> . string?) . -> . response/c)
```

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

```
(redirect/get) → request?
```

See `web-server/servlet/web`.

2.7 Web Cells

```
(require web-server/servlet/web-cells)
```

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)  
      '(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)
  (let/cc k
    (let loop ()
      (k
       (lambda (embed/url)
         '(div (h3 ,(number->string (web-cell-ref a-counter)))
              (a ([href
                   ,(embed/url
                     (lambda _
                       ; A new frame has been created
                       (define last (web-cell-ref a-counter))
                       ; We can inspect the value at the parent
                       (web-cell-shadow a-counter (add1 last))
                       ; The new frame has been modified
                       (loop))))])
                "+"))))))))

```

2.7.1 Stateless Web Cells

```
(require web-server/lang/web-cells)
```

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`.

2.8 File Boxes

```
(require web-server/lang/file-box)
```


As mentioned earlier, it is dangerous to rely on the store in Web Language servlets, due to the deployment scenarios available to them. "lang/file-box.ss" 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.

2.9 Stateless Web Parameters

```
(require web-server/lang/web-param)
```

It is not easy to use parameterize in the Web Language. "lang/web-param.ss" 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.

2.10 URL-Based Dispatch

```
(require web-server/dispatch)
```

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.

2.10.1 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
                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"

```

2.10.2 API Reference

```

(dispatch-rules
 [dispatch-pattern dispatch-fun]
 ...
 [else else-fun])
(dispatch-rules
 [dispatch-pattern dispatch-fun]
 ...)

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

```

Returns two values: the first is a dispatching function with the contract `(-> request? response/c)` that calls the appropriate `dispatch-fun` based on the first `dispatch-pattern` that matches the request's URL; 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.

```

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

```

```

(dispatch-case
 [dispatch-pattern dispatch-fun]
 ...
 [else else-fun])
(dispatch-case
 [dispatch-pattern dispatch-fun]
 ...)

```

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

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

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

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

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

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

Calls `serve/servlet` with appropriate arguments so that every request is handled by `dispatch`.

2.10.3 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.

2.10.4 Extending `web-server/dispatch`

```
(require web-server/dispatch/extend)
```

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.

```
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

```

2.11 Formlets

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

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*.

2.11.1 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))
```

The rest of the manual gives the details of formlet usage and extension.

2.11.2 Syntactic Shorthand

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

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

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


Constructs a formlet with the specified *rendering* and the processing resulting in the *yields-expr* expression. The *rendering* form is a quasiquoted X-expression, with two special caveats:

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

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

2.11.3 Functional Usage

```
(require web-server/formlets/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 §2.11.4 “Predefined Formlets” for example low-level formlets using these combinators.

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

Equivalent to `(listof xexpr/c)`

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

Equivalent to `(-> integer? (values xexpr-forest/c (-> (listof binding?) (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.

```
(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 (any/c . -> . any/c))  
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 result 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`.

2.11.4 Predefined Formlets

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

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`.

```
(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?) = #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?) = #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.

```
(checkbox value checked? [#:attributes attrs])
```

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

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

```
(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 integer?)

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

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

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

2.11.5 Utilities

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

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

```
(send/formlet f [#:wrap wrapper]) → any/c
  f : (formlet/c any/c)
  wrapper : (xexpr/c . -> . response/c)
            = (lambda (form-xexpr)
                '(html (head (title "Form Entry"))
                      (body ,form-xexpr)))
```

Uses `send/suspend` to send `f`'s rendering (wrapped in a FORM tag 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 : embed/url/c
  f : (formlet/c any/c)
```

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

2.12 Templates

```
(require web-server/templates)
```

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 PLT-based Web applications.

2.12.1 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 Scheme 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>"
```

2.12.2 Dynamic

`include-template` gives the template access to the *complete lexical context* of the including program. This context can be accessed via the §3 “@-Reader” 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"))
```

Although all the examples here generate HTML, the template language and the §17 “Text Preprocessor” it is based on can be used to generate any text-based format: C, SQL, form emails, reports, etc.

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 Scheme used by templates: they can use macros, structs, continuation marks, threads, etc. However, Scheme values that are ultimately returned must be printable by the §17 “Text Preprocessor”. For example, consider the following outputs of the title line of different calls to `fast-template`:

- `(fast-template 'Templates)`
<head><title>Fastest Templates in the West!</title></head>
- `(fast-template 42)`
<head><title>Fastest 42 in the West!</title></head>
- `(fast-template (list "Noo" "dles"))`
<head><title>Fastest Noodles in the West!</title></head>

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

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

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

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

2.12.3 Gotchas

To obtain an `@` character in template output, you must escape the it, because it is the escape character of the §3 “@-Reader” 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 Scheme expression, which can be used for longer strings too.)

The §3 “@-Reader” will read Scheme 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 can subvert this by explicitly delimiting the identifier:

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

Another gotcha is that since the template is compiled into a Scheme 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>
</tr>
</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.

2.12.4 HTTP Responses

The quickest way to generate an HTTP response from a template is using the `list` response type:

```
(list #"text/html" (include-template "static.html"))
```

If you want more control then you can generate a `response/full` struct:

```
(make-response/full
 200 #"Okay"
 (current-seconds) TEXT/HTML-MIME-TYPE
 empty
 (list (include-template "static.html")))
```

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 create a `cdata` structure:

```
'(html ,(make-cdata #f #f (include-template "static.html")))
```

2.12.5 API Details

```
(include-template path)
```

Compiles the template at *path* using the §3 “@-Reader” syntax within the enclosing lexical context.

Example:

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

```
(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>
  }

```

Scheme Example:

```

(in c clients "<tr><td>" (car c) ", " (cdr c) "</td></tr>")

```

2.12.6 Conversion Example

Al Church has been maintaining a blog with PLT Scheme 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)
  `(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)
  '(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)))

```

Some of these problems go away by using here strings, as described in the documentation on §12.6.6 “Reading Strings”.

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.

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)
  (list TEXT/HTML-MIME-TYPE
        (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;
    }

    #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*.

2.13 Continuation Managers

Since Scheme 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.

2.13.1 General

```
(require web-server/managers/manager)
```

"managers/manager.ss" defines the manager interface. It is required by the users and implementers of managers.

```
(struct manager (create-instance
                 adjust-timeout!
                 clear-continuations!
                 continuation-store!
                 continuation-lookup))
create-instance : ((-> void) . -> . number?)
adjust-timeout! : (number? number? . -> . void)
clear-continuations! : (number? . -> . void)
continuation-store! : (number? any/c expiration-handler/c . -> . (list/c number? number?))
continuation-lookup : (number? number? number? . -> . any/c)
```

`create-instance` is called to initialize an 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.

```
(struct (exn:fail:servlet-manager:no-instance exn:fail) (expiration-handler
              expiration-handler : expiration-handler/c)
```

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
              expiration-handler : expiration-handler/c)
```

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

2.13.2 No Continuations

```
(require web-server/managers/none)
```


"managers/none.ss" defines a manager constructor:

```
(create-none-manager instance-expiration-handler) → manager?
  instance-expiration-handler : expiration-handler/c
```

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 §2.1.3 “Stateless Servlets”.)

2.13.3 Timeouts

```
(require web-server/managers/timeouts)
```

"managers/timeouts.ss" defines a manager constructor:

```
(create-timeout-manager instance-exp-handler
  instance-timeout
  continuation-timeout) → manager?
  instance-exp-handler : expiration-handler/c
  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 .

2.13.4 LRU

```
(require web-server/managers/lru)
```

"managers/lru.ss" 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 : expiration-handler/c
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 : expiration-handler/c
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.

2.14 Stuffers

```
(require web-server/stuffers)
```

The

```
#lang web-server
```

module 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.

2.14.1 Basic Combinators

```
(require web-server/stuffers/stuffer)
```

```
(struct stuffer (in out))  
  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 any/c any/c`)
`g` : (`stuffer any/c any/c`)
`f` : (`stuffer 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 any/c any/c`)
`f` : (`stuffer any/c any/c`)
`g` : (`stuffer any/c any/c`)

`stuffer-compose` with arguments swapped.

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

2.14.2 Serialization

(`require web-server/stuffers/serialize`)

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

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

2.14.3 Base64 Encoding

(`require web-server/stuffers/base64`)

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

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

Useful for getting URL-safe bytes.

2.14.4 GZip Compression

```
(require web-server/stuffers/gzip)
```

```
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.

2.14.5 Key/Value Storage

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

```
(require web-server/stuffers/store)
```

```
(struct store (write read))  
  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.

2.14.6 Hash-addressed Storage

(require web-server/stuffers/hash)

hash/c : contract?

Equivalent to (-> bytes? bytes?).

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

2.14.7 HMAC-SHA1 Signing

(require web-server/stuffers/hmac-sha1)

(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.

2.14.8 Helpers

```
(require web-server/lang/stuff-url)
```

```
(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"))
```


3 Extending the Web Server

3.1 Configuration

There are a number of libraries and utilities useful for configuring the Web Server .

3.1.1 Configuration Table Structure

```
(require web-server/configuration/configuration-table-structs)
```

"configuration/configuration-table-structs.ss" provides the following structures that represent a standard configuration (see §3.5 “Web Server Unit”) of the Web Server . The contracts on this structure influence the valid types of values in the configuration table S-expression file format described in §3.1.2 “Configuration Table”.

```
(struct configuration-table (port
                            max-waiting
                            initial-connection-timeout
                            default-host
                            virtual-hosts))

port : port-number?
max-waiting : natural-number/c
initial-connection-timeout : natural-number/c
default-host : host-table?
virtual-hosts : (listof (cons/c string? host-table?))
```

```
(struct host-table (indices log-format messages timeouts paths))
indices : (listof string?)
log-format : symbol?
messages : messages?
timeouts : timeouts?
paths : paths?
```

```
(struct host (indices
              log-format
              log-path
              passwords
              responders
              timeouts
              paths))
indices : (listof string?)
```

```
log-format : symbol?
log-path : (or/c false/c path-string?)
passwords : (or/c false/c path-string?)
responders : responders?
timeouts : timeouts?
paths : paths?
```

```
(struct responders (servlet
  servlet-loading
  authentication
  servlets-refreshed
  passwords-refreshed
  file-not-found
  protocol
  collect-garbage))
servlet : (url? any/c . -> . response/c)
servlet-loading : (url? any/c . -> . response/c)
authentication : (url? (cons/c symbol? string?) . -> . response/c)
servlets-refreshed : (-> response/c)
passwords-refreshed : (-> response/c)
file-not-found : (request? . -> . response/c)
protocol : (url? . -> . response/c)
collect-garbage : (-> response/c)
```

```
(struct messages (servlet
  authentication
  servlets-refreshed
  passwords-refreshed
  file-not-found
  protocol
  collect-garbage))
servlet : string?
authentication : string?
servlets-refreshed : string?
passwords-refreshed : string?
file-not-found : string?
protocol : string?
collect-garbage : string?
```

```
(struct timeouts (default-servlet
  password
  servlet-connection
  file-per-byte
  file-base))
```

```
default-servlet : number?
password : number?
servlet-connection : number?
file-per-byte : number?
file-base : number?
```

```
(struct paths (conf
  host-base
  log
  htdocs
  servlet
  mime-types
  passwords))
conf : (or/c false/c path-string?)
host-base : (or/c false/c path-string?)
log : (or/c false/c path-string?)
htdocs : (or/c false/c path-string?)
servlet : (or/c false/c path-string?)
mime-types : (or/c false/c path-string?)
passwords : (or/c false/c path-string?)
```

3.1.2 Configuration Table

```
(require web-server/configuration/configuration-table)
```

"configuration/configuration-table.ss" provides functions for reading, writing, parsing, and printing `configuration-table` structures.

```
default-configuration-table-path : path?
```

The default configuration table S-expression file.

```
configuration-table-sexpr? : (any . -> . boolean?)
```

Equivalent to `list?`.

```
(sexpr->configuration-table sexpr) → configuration-table?
sexpr : configuration-table-sexpr?
```

This function converts a `configuration-table` from an S-expression.

```
(configuration-table->sexpr ctable)
→ configuration-table-sexpr?
  ctable : configuration-table?
```

This function converts a `configuration-table` to an S-expression.

```
`((port ,integer?)
  (max-waiting ,integer?)
  (initial-connection-timeout ,integer?)
  (default-host-table
   ,host-table-sexpr?)
  (virtual-host-table
   (list ,symbol? ,host-table-sexpr?)
   ...))
```

where a `host-table-sexpr` is:

```
`(host-table
  (default-indices ,string? ...)
  (log-format ,symbol?)
  (messages
   (servlet-message ,path-string?)
   (authentication-message ,path-string?)
   (servlets-refreshed ,path-string?)
   (passwords-refreshed ,path-string?)
   (file-not-found-message ,path-string?)
   (protocol-message ,path-string?)
   (collect-garbage ,path-string?))
  (timeouts
   (default-servlet-timeout ,integer?)
   (password-connection-timeout ,integer?)
   (servlet-connection-timeout ,integer?)
   (file-per-byte-connection-timeout ,integer?)
   (file-base-connection-timeout ,integer))
  (paths
   (configuration-root ,path-string?)
   (host-root ,path-string?)
   (log-file-path ,path-string?)
   (file-root ,path-string?)
   (servlet-root ,path-string?)
   (mime-types ,path-string?)
   (password-authentication ,path-string?)))
```

In this syntax, the `'messages` paths are relative to the `'configuration-root` directory. All the paths in `'paths` are relative to `'host-root` (other than `'host-root` obviously.)

Allowable `'log-formats` are those accepted by `log-format->format`.

Note: You almost always want to leave everything in the `'paths` section the default except the `'host-root`.

```
(read-configuration-table path) → configuration-table?  
  path : path-string?
```

This function reads a `configuration-table` from `path`.

```
(write-configuration-table ctable path) → void  
  ctable : configuration-table?  
  path : path-string?
```

This function writes a `configuration-table` to `path`.

3.1.3 Servlet Namespaces

```
(require web-server/configuration/namespace)
```

"`configuration/namespace.ss`" provides a function to help create the `make-servlet-namespace` procedure needed by the `make` functions of "`dispatchers/dispatch-servlets.ss`" and "`dispatchers/dispatch-lang.ss`".

```
make-servlet-namespace/c : contract?
```

Equivalent to

```
(->* ()  
  (#:additional-specs (listof module-path?))  
  namespace?)
```

.

```
(make-make-servlet-namespace #:to-be-copied-module-specs to-be-copied-module-specs)  
→ make-servlet-namespace/c  
  to-be-copied-module-specs : (listof module-path?)
```

This function creates a function that when called will construct a new `namespace` that has all the modules from `to-be-copied-module-specs` and `additional-specs`, as well as `mzscheme` and `mred`, provided they are already attached to the `(current-namespace)` of the call-site.

Example:

```
(make-make-servlet-namespace
 #:to-be-copied-module-specs '((lib "database.ss" "my-module")))
```

Why this is useful

A different namespace is needed for each servlet, so that if servlet A and servlet B both use a stateful module C, they will be isolated from one another. We see the Web Server as an operating system for servlets, so we inherit the isolation requirement on operating systems.

However, there are some modules which must be shared. If they were not, then structures cannot be passed from the Web Server to the servlets, due to a subtlety in the way MzScheme implements structures.

Since, on occasion, a user will actually want servlets A and B to interact through module C. A custom `make-servlet-namespace` can be created, through this procedure, that attaches module C to all servlet namespaces. Through other means (see §3.3 “Dispatchers”) different sets of servlets can share different sets of modules.

3.1.4 Standard Responders

```
(require web-server/configuration/responders)
```

"`configuration/responders.ss`" provides some functions that help constructing HTTP responders. These functions are used by the default dispatcher constructor (see §3.5 “Web Server Unit”) to turn the paths given in the `configuration-table` into responders for the associated circumstance.

```
(file-response http-code
              short-version
              text-file
              header ...) → response/c
http-code : natural-number/c
short-version : string?
text-file : string?
header : header?
```

Generates a `response/full` with the given `http-code` and `short-version` as the corresponding fields; with the content of the `text-file` as the body; and, with the `headers` as, you guessed it, headers.

This does not cause redirects to a well-known URL, such as "`conf/not-found.html`",

but rather use the contents of "not-found.html" (for example) as its contents. Therefore, any relative URLs in *text-file* are relative to whatever URL *file-response* is used to respond to. Thus, you should probably use absolute URLs in these files.

```
(servlet-loading-responder url exn) → response/c
  url : url?
  exn : exn?
```

Gives *exn* to the *current-error-handler* and response with a stack trace and a "Servlet didn't load" message.

```
(gen-servlet-not-found file) → ((url url?) . -> . response/c)
  file : path-string?
```

Returns a function that generates a standard "Servlet not found." error with content from *file*.

```
(servlet-error-responder url exn) → response/c
  url : url?
  exn : exn?
```

Gives *exn* to the *current-error-handler* and response with a stack trace and a "Servlet error" message.

```
(gen-servlet-responder file)
→ ((url url?) (exn any/c) . -> . response/c)
  file : path-string?
```

Prints the *exn* to standard output and responds with a "Servlet error." message with content from *file*.

```
(gen-servlets-refreshed file) → (-> response/c)
  file : path-string?
```

Returns a function that generates a standard "Servlet cache refreshed." message with content from *file*.

```
(gen-passwords-refreshed file) → (-> response/c)
  file : path-string?
```

Returns a function that generates a standard "Passwords refreshed." message with content from *file*.

```
(gen-authentication-responder file)
→ ((url url?) (header header?) . -> . response/c)
   file : path-string?
```

Returns a function that generates an authentication failure error with content from *file* and *header* as the HTTP header.

```
(gen-protocol-responder file) → ((url url?) . -> . response/c)
   file : path-string?
```

Returns a function that generates a "Malformed request" error with content from *file*.

```
(gen-file-not-found-responder file)
→ ((req request?) . -> . response/c)
   file : path-string?
```

Returns a function that generates a standard "File not found" error with content from *file*.

```
(gen-collect-garbage-responder file) → (-> response/c)
   file : path-string?
```

Returns a function that generates a standard "Garbage collection run" message with content from *file*.

3.2 Setting Up Servlets

```
(require web-server/servlet/setup)
```

This module is used internally to build and load servlets. It may be useful to those who are trying to extend the server.

```
(make-v1.servlet directory timeout start) → servlet?
   directory : path-string?
   timeout   : integer?
   start     : (request? . -> . response/c)
```

Creates a version 1 servlet that uses *directory* as its current directory, a timeout manager with a *timeout* timeout, and *start* as the request handler.

```
(make-v2.servlet directory manager start) → servlet?
   directory : path-string?
```



```
manager : manager?
start : (request? . -> . response/c)
```

Creates a version 2 servlet that uses *directory* as its current directory, a *manager* as the continuation manager, and *start* as the request handler.

```
(make-stateless.servlet directory start) → servlet?
  directory : path-string?
  start : (request? . -> . response/c)
```

Creates a stateless web-server servlet that uses *directory* as its current directory and *start* as the request handler.

```
default-module-specs : (listof module-path?)
```

The modules that the Web Server needs to share with all servlets.

```
path->servlet/c : contract?
```

Equivalent to `(-> path? servlet?)`.

```
(make-default-path->servlet
 [#:make-servlet-namespace make-servlet-namespace
  #:timeouts-default-servlet timeouts-default-servlet])
→ path->servlet/c
  make-servlet-namespace : make-servlet-namespace/c
                        = (make-make-servlet-namespace)
  timeouts-default-servlet : integer? = 30
```

Constructs a procedure that loads a servlet from the path in a namespace created with *make-servlet-namespace*, using a timeout manager with *timeouts-default-servlet* as the default timeout (if no manager is given.)

3.2.1 Internal Servlet Representation

```
(require web-server/private/servlet)
```

```
(struct servlet (custodian namespace manager directory handler)
 #:mutable)
  custodian : custodian?
  namespace : namespace?
  manager : manager?
```

```
directory : path-string?  
handler : (request? . -> . response/c)
```

Instances of this structure hold the necessary parts of a servlet: the `custodian` responsible for the servlet's resources, the `namespace` the servlet is executed within, the `manager` responsible for the servlet's continuations, the current `directory` of the servlet, and the `handler` for all requests to the servlet.

3.3 Dispatchers

The Web Server is really just a particular configuration of a dispatching server. There are a number of dispatchers that are defined to support the Web Server . Other dispatching servers, or variants of the Web Server , may find these useful. In particular, if you want a peculiar processing pipeline for your Web Server installation, this documentation will be useful.

3.3.1 General

```
(require web-server/dispatchers/dispatch)
```

"dispatchers/dispatch.ss" provides a few functions for dispatchers in general.

```
dispatcher/c : contract?
```

Equivalent to `(-> connection? request? void)`.

```
(dispatcher-interface-version/c any) → boolean?  
any : any/c
```

Equivalent to `(symbols 'v1)`

```
(struct exn:dispatcher ())
```

An exception thrown to indicate that a dispatcher does not apply to a particular request.

```
(next-dispatcher) → void
```

Raises a `exn:dispatcher`

As the `dispatcher/c` contract suggests, a dispatcher is a function that takes a connection and request object and does something to them. Mostly likely it will generate some response and output it on the connection, but it may do something different. For example, it may

apply some test to the request object, perhaps checking for a valid source IP address, and error if the test is not passed, and call `next-dispatcher` otherwise.

Consider the following example dispatcher, that captures the essence of URL rewriting:

```
; (url? -> url?) dispatcher/c -> dispatcher/c
(lambda (rule inner)
  (lambda (conn req)
    ; Call the inner dispatcher...
    (inner conn
      ; with a new request object...
      (struct-copy request req
        ; with a new URL!
        [request-uri (rule (request-uri req))])))
```

3.3.2 Mapping URLs to Paths

```
(require web-server/dispatchers/filesystem-map)
```

"dispatchers/filesystem-map.ss" provides a means of mapping URLs to paths on the filesystem.

`url->path/c` : `contract?`

This contract is equivalent to `(->* (url?) (path? (listof path-element?)))`. The returned `path?` is the path on disk. The list is the list of path elements that correspond to the path of the URL.

```
(make-url->path base) → url->path/c
  base : path-string?
```

The `url->path/c` returned by this procedure considers the root URL to be `base`. It ensures that `".."`s in the URL do not escape the `base` and removes them silently otherwise.

```
(make-url->valid-path url->path) → url->path/c
  url->path : url->path/c
```

Runs the underlying `url->path`, but only returns if the path refers to a file that actually exists. If it does not, then the suffix elements of the URL are removed until a file is found. If this never occurs, then an error is thrown.

This is primarily useful for dispatchers that allow path information after the name of a service to be used for data, but where the service is represented by a file. The most prominent

example is obviously servlets.

```
(filter-url->path regex url->path) → url->path/c
  regex : regexp?
  url->path : url->path/c
```

Runs the underlying `url->path` but will only return if the path, when considered as a string, matches the `regex`. This is useful to disallow strange files, like GIFs, from being considered servlets when using the servlet dispatchers. It will return a `exn:fail:filesystem:exists?` exception if the path does not match.

3.3.3 Sequencing

```
(require web-server/dispatchers/dispatch-sequencer)
```

The `web-server/dispatchers/dispatch-sequencer` module defines a dispatcher constructor that invokes a sequence of dispatchers until one applies.

```
(make dispatcher ...) → dispatcher/c
  dispatcher : dispatcher/c
```

Invokes each `dispatcher`, invoking the next if the first calls `next-dispatcher`. If no `dispatcher` applies, then it calls `next-dispatcher` itself.

3.3.4 Timeouts

```
(require web-server/dispatchers/dispatch-timeout)
```

The `web-server/dispatchers/dispatch-timeout` module defines a dispatcher constructor that changes the timeout on the connection and calls the next dispatcher.

```
(make new-timeout) → dispatcher/c
  new-timeout : integer?
```

Changes the timeout on the connection with `adjust-connection-timeout!` called with `new-timeout`.

3.3.5 Lifting Procedures

```
(require web-server/dispatchers/dispatch-lift)
```

The `web-server/dispatchers/dispatch-lift` module defines a dispatcher constructor.

```
(make proc) → dispatcher/c
  proc : (request? . -> . response/c)
```

Constructs a dispatcher that calls `proc` on the request object, and outputs the response to the connection.

3.3.6 Filtering Requests

```
(require web-server/dispatchers/dispatch-filter)
```

The `web-server/dispatchers/dispatch-filter` module defines a dispatcher constructor that calls an underlying dispatcher with all requests that pass a predicate.

```
(make regex inner) → dispatcher/c
  regex : regexp?
  inner : dispatcher/c
```

Calls `inner` if the URL path of the request, converted to a string, matches `regex`. Otherwise, calls `next-dispatcher`.

3.3.7 Procedure Invocation upon Request

```
(require web-server/dispatchers/dispatch-pathprocedure)
```

The `web-server/dispatchers/dispatch-pathprocedure` module defines a dispatcher constructor for invoking a particular procedure when a request is given to a particular URL path.

```
(make path proc) → dispatcher/c
  path : string?
  proc : (request? . -> . response/c)
```

Checks if the request URL path as a string is equal to `path` and if so, calls `proc` for a response.

This is used in the standard Web Server pipeline to provide a URL that refreshes the password file, servlet cache, etc.

3.3.8 Logging

```
(require web-server/dispatchers/dispatch-log)
```

The `web-server/dispatchers/dispatch-log` module defines a dispatcher constructor for transparent logging of requests.

```
format-req/c : contract?
```

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

```
paren-format : format-req/c
```

Formats a request by:

```
(format
  "~s~n"
  (list 'from (request-client-ip req)
        'to (request-host-ip req)
        'for (url->string (request-uri req)) 'at
        (date->string
          (seconds->date (current-seconds)) #t)))
```

```
extended-format : format-req/c
```

Formats a request by:

```
(format
  "~s~n"
  '((client-ip ,(request-client-ip req))
    (host-ip ,(request-host-ip req))
    (referer
     ,(let ([R (headers-assq*
                #"Referer"
                (request-headers/raw req))])
        (if R
            (header-value R)
            #f))))
    (uri ,(url->string (request-uri req)))
    (time ,(current-seconds))))
```

```
apache-default-format : format-req/c
```

Formats a request like Apache's default.

```
log-format/c : contract?
```

Equivalent to `(symbols 'parenthesized-default 'extended 'apache-default)`.

```
(log-format->format id) → format-req/c  
  id : log-format/c
```

Maps `'parenthesized-default` to `paren-format`, `'extended` to `extended-format`, and `'apache-default` to `apache-default-format`.

```
(make [#:format format #:log-path log-path]) → dispatcher/c  
  format : format-req/c = paren-format  
  log-path : path-string? = "log"
```

Logs requests to `log-path` by using `format` to format the requests. Then invokes `next-dispatcher`.

3.3.9 Password Protection

```
(require web-server/dispatchers/dispatch-passwords)
```

The `web-server/dispatchers/dispatch-passwords` module defines a dispatcher constructor that performs HTTP Basic authentication filtering.

```
denied?/c : contract?
```

Equivalent to `(-> request? (or/c false/c string?))`. The return is the authentication realm as a string if the request is not authorized and `#f` if the request *is* authorized.

```
(make denied?  
  [#:authentication-responder authentication-responder])  
→ dispatcher/c  
  denied? : denied?/c  
  authentication-responder : (url? header? . -> . response/c)  
                             = (gen-authentication-responder "forbidden.html")
```

A dispatcher that checks if the request is denied based on `denied?`. If so, then `authentication-responder` is called with a `header` that requests credentials. If not, then `next-dispatcher` is invoked.

```
authorized?/c : contract?
```

Equivalent to `(-> string? (or/c false/c bytes?) (or/c false/c bytes?) (or/c false/c string?))`. The input is the URI as a string and the username and passwords as bytes. The return is the authentication realm as a string if the user is not authorized and `#f` if the request is authorized.

```
(make-basic-denied?/path password-file)
→ (-> void) authorized?/c
   password-file : path-string?
```

Creates an authorization procedure based on the given password file. The first returned value is a procedure that refreshes the password cache used by the authorization procedure.

`password-file` is parsed as:

```
(list ([domain : string?]
      [path : string?] ; This string is interpreted as a regex
      (list [user : symbol?]
            [pass : string?])
      ...))
...)
```

For example:

```
'(("secret stuff" "/secret(/.*)?" (bubba "bbq") (Billy "BoB")))
```

3.3.10 Virtual Hosts

```
(require web-server/dispatchers/dispatch-host)
```

The `web-server/dispatchers/dispatch-host` module defines a dispatcher constructor that calls a different dispatcher based upon the host requested.

```
(make lookup-dispatcher) → dispatcher/c
lookup-dispatcher : (symbol? . -> . dispatcher/c)
```

Extracts a host from the URL requested, or the Host HTTP header, calls `lookup-dispatcher` with the host, and invokes the returned dispatcher. If no host can be extracted, then `'none` is used.

3.3.11 Serving Files


```
(require web-server/dispatchers/dispatch-files)
```

The `web-server/dispatchers/dispatch-files` module allows files to be served. It defines a dispatcher construction procedure.

```
(make #:url->path url->path
      [#:path->mime-type path->mime-type
       #:indices indices]) → dispatcher/c
url->path : url->path/c
path->mime-type : (path? . -> . bytes?)
                 = (lambda (path) TEXT/HTML-MIME-TYPE)
indices : (listof string?) = (list "index.html" "index.htm")
```

Uses `url->path` to extract a path from the URL in the request object. If this path does not exist, then the dispatcher does not apply and `next-dispatcher` is invoked. If the path is a directory, then the `indices` are checked in order for an index file to serve. In that case, or in the case of a path that is a file already, `path->mime-type` is consulted for the MIME Type of the path. The file is then streamed out the connection object.

This dispatcher supports HTTP Range GET requests and HEAD requests.

3.3.12 Serving Servlets

```
(require web-server/dispatchers/dispatch-servlets)
```

The `web-server/dispatchers/dispatch-servlets` module defines a dispatcher constructor that runs servlets.

```
url->servlet/c : contract?
```

Equivalent to `(-> url? servlet?)`

```
(make-cached-url->servlet url->path
                          path->servlet)
→ (-> void) url->servlet/c
url->path : url->path/c
path->servlet : path->servlet/c
```

The first return value flushes the cache. The second is a procedure that uses `url->path` to resolve the URL to a path, then uses `path->servlet` to resolve that path to a servlet, caching the results in an internal table.

```

(make url->servlet
  [#:responders-servlet-loading responders-servlet-loading
   #:responders-servlet responders-servlet])
→ dispatcher/c
url->servlet : url->servlet/c
responders-servlet-loading : (url? exn? . -> . response/c)
                           = servlet-loading-responder
responders-servlet : (url? exn? . -> . response/c)
                    = servlet-error-responder

```

This dispatcher runs Scheme servlets, using `url->servlet` to resolve URLs to the underlying servlets. If servlets have errors loading, then `responders-servlet-loading` is used. Other errors are handled with `responders-servlet`. If a servlet raises calls `next-dispatcher`, then the signal is propagated by this dispatcher.

3.3.13 Statistics

```
(require web-server/dispatchers/dispatch-stat)
```

The `web-server/dispatchers/dispatch-stat` module provides services related to performance statistics.

```

(make-gc-thread time) → thread?
time : integer?

```

Starts a thread that calls `(collect-garbage)` every `time` seconds.

```
(make) → dispatcher/c
```

Returns a dispatcher that prints memory usage on every request.

3.3.14 Limiting Requests

```
(require web-server/dispatchers/limit)
```

The `web-server/dispatchers/limit` module provides a wrapper dispatcher that limits how many requests are serviced at once.

```

(make limit inner [#:over-limit over-limit]) → dispatcher/c
limit : number?
inner : dispatcher/c

```

```
over-limit : (symbols 'block 'kill-new 'kill-old) = 'block
```

Returns a dispatcher that defers to *inner* for work, but will forward a maximum of *limit* requests concurrently.

If there are no additional spaces inside the limit and a new request is received, the *over-limit* option determines what is done. The default (`'block`) causes the new request to block until an old request is finished being handled. If *over-limit* is `'kill-new`, then the new request handler is killed—a form of load-shedding. If *over-limit* is `'kill-old`, then the oldest request handler is killed—prioritizing new connections over old. (This setting is a little dangerous because requests might never finish if there is constant load.)

Consider this example:

```
#lang scheme

(require web-server/web-server
         web-server/http
         web-server/http/response
         (prefix-in limit: web-server/dispatchers/limit)
         (prefix-in filter: web-server/dispatchers/dispatch-filter)
         (prefix-in sequencer: web-server/dispatchers/dispatch-sequencer))

(serve #:dispatch
      (sequencer:make
       (filter:make
        #rx"/limited"
        (limit:make
         5
         (lambda (conn req)
           (output-response/method
            conn
            (make-response/full
             200 #"Okay"
             (current-seconds) TEXT/HTML-MIME-TYPE
             empty
             (list (string->bytes/utf-8
                   (format "hello world ~a"
                           (sort (build-list 100000 (lambda x (random 100)))
                                   <))))))
            (request-method req))))
       #:over-limit 'block))
      (lambda (conn req)
        (output-response/method
         conn
         (make-response/full 200 #"Okay"
                             (current-seconds) TEXT/HTML-MIME-TYPE
```

```

                                empty
                                (list #"<html><body>Unlimited</body></html>")
      (request-method req)))
      #:port 8080)

(do-not-return)

```

3.4 Web Config Unit

The Web Server offers a unit-based approach to configuring the server.

3.4.1 Configuration Signature

```
(require web-server/web-config-sig)
```

`web-config^` : signature

Provides contains the following identifiers.

`max-waiting` : integer?

Passed to `tcp-accept`.

`virtual-hosts` : (string? . -> . host?)

Contains the configuration of individual virtual hosts.

`initial-connection-timeout` : integer?

Specifies the initial timeout given to a connection.

`port` : port-number?

Specifies the port to serve HTTP on.

`listen-ip` : (or/c false/c string?)

Passed to `tcp-listen`.

`make-servlet-namespace` : make-servlet-namespace/c

Passed to `servlets:make` through `make-default-path->servlet`.

3.4.2 Configuration Units

```
(require web-server/web-config-unit)
```

```
(configuration-table->web-config@  
  path  
  [#:port port  
   #:listen-ip listen-ip  
   #:make-servlet-namespace make-servlet-namespace])  
→ (unit/c (import) (export web-config^))  
path : path-string?  
port : (or/c false/c port-number?) = #f  
listen-ip : (or/c false/c string?) = #f  
make-servlet-namespace : make-servlet-namespace/c  
                        = (make-make-servlet-namespace)
```

Reads the S-expression at *path* and calls `configuration-table-sexpr->web-config@` appropriately.

```
(configuration-table-sexpr->web-config@  
  sexpr  
  [#:web-server-root web-server-root  
   #:port port  
   #:listen-ip listen-ip  
   #:make-servlet-namespace make-servlet-namespace])  
→ (unit/c (import) (export web-config^))  
sexpr : list?  
web-server-root : path-string?  
                = (directory-part default-configuration-table-path)  
port : (or/c false/c port-number?) = #f  
listen-ip : (or/c false/c string?) = #f  
make-servlet-namespace : make-servlet-namespace/c  
                       = (make-make-servlet-namespace)
```

Parses *sexpr* as a configuration-table and constructs a `web-config^` unit.

3.5 Web Server Unit

The Web Server offers a unit-based approach to running the server.

3.5.1 Signature

```
(require web-server/web-server-sig)
```

`web-server^` : signature

```
(serve) → (-> void)
```

Runs the server and returns a procedure that shuts down the server.

```
(serve-ports ip op) → void
```

```
  ip : input-port?
```

```
  op : output-port?
```

Serves a single connection represented by the ports `ip` and `op`.

3.5.2 Unit

```
(require web-server/web-server-unit)
```

```
web-server@ : (unit/c (web-config^ tcp^)  
                  (web-server^))
```

Uses the `web-config^` to construct a `dispatcher/c` function that sets up one virtual host dispatcher, for each virtual host in the `web-config^`, that sequences the following operations:

- Logs the incoming request with the given format to the given file
- Performs HTTP Basic Authentication with the given password file
- Allows the `"/conf/refresh-passwords"` URL to refresh the password file.
- Allows the `"/conf/collect-garbage"` URL to call the garbage collector.
- Allows the `"/conf/refresh-servlets"` URL to refresh the servlets cache.
- Execute servlets in the mapping URLs to the given servlet root directory under htdocs.
- Serves files under the `"/"` URL in the given htdocs directory.

Using this `dispatcher/c`, it loads a dispatching server that provides `serve` and `serve-ports` functions that operate as expected.

3.6 Internal

The Web Server is a complicated piece of software and as a result, defines a number of interesting and independently useful sub-components. Some of these are documented here.

3.6.1 Timers

```
(require web-server/private/timer)
```

"private/timer.ss" provides a functionality for running procedures after a given amount of time, that may be extended.

```
(struct timer (evt expire-seconds action))  
  evt : evt?  
  expire-seconds : number?  
  action : (-> void)
```

evt is an `alarm-evt` that is ready at `expire-seconds`. `action` should be called when this `evt` is ready.

```
(start-timer-manager) → void
```

Handles the execution and management of timers.

```
(start-timer s action) → timer?  
  s : number?  
  action : (-> void)
```

Registers a timer that runs `action` after `s` seconds.

```
(reset-timer! t s) → void  
  t : timer?  
  s : number?
```

Changes `t` so that it will fire after `s` seconds.

```
(increment-timer! t s) → void  
  t : timer?  
  s : number?
```

Changes `t` so that it will fire after `s` seconds from when it does now.

```
(cancel-timer! t) → void
  t : timer?
```

Cancels the firing of `t` ever and frees resources used by `t`.

3.6.2 Connection Manager

```
(require web-server/private/connection-manager)
```

"private/connection-manager.ss" provides functionality for managing pairs of input and output ports. We have plans to allow a number of different strategies for doing this.

```
(struct connection (timer i-port o-port custodian close?))
  timer : timer?
  i-port : input-port?
  o-port : output-port?
  custodian : custodian?
  close? : boolean?
```

A connection is a pair of ports (`i-port` and `o-port`) that is ready to close after the current job if `close?` is `#t`. Resources associated with the connection should be allocated under `custodian`. The connection will last until `timer` triggers.

```
(start-connection-manager) → void
```

Runs the connection manager (now just the timer manager).

```
(new-connection timeout
  i-port
  o-port
  cust
  close?) → connection?

  timeout : number?
  i-port : input-port?
  o-port : output-port?
  cust : custodian?
  close? : boolean?
```

Constructs a connection with a timer with a trigger of `timeout` that calls `kill-connection!`.

```
(kill-connection! c) → void
```


`c : connection?`

Closes the ports associated with `c`, kills the timer, and shuts down the custodian.

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

Calls `reset-timer!` with the timer behind `c` with `t`.

3.6.3 Dispatching Server

The Web Server is just a configuration of a dispatching server. This dispatching server component is useful on its own.

Dispatching Server Signatures

```
(require web-server/private/dispatch-server-sig)
```

The `web-server/private/dispatch-server-sig` library provides two signatures.

`dispatch-server^` : signature

The `dispatch-server^` signature is an alias for `web-server^`.

```
(serve) → (-> void)
```

Runs the server and returns a procedure that shuts down the server.

```
(serve-ports ip op) → void
  ip : input-port?
  op : output-port?
```

Serves a single connection represented by the ports `ip` and `op`.

`dispatch-server-config^` : signature

```
port : port-number?
```

Specifies the port to serve on.

`listen-ip` : (or/c string? false/c)

Passed to `tcp-listen`.

`max-waiting` : integer?

Passed to `tcp-accept`.

`initial-connection-timeout` : integer?

Specifies the initial timeout given to a connection.

`(read-request c p port-addresses)` → any/c boolean?

`c` : connection?

`p` : port-number?

`port-addresses` : (input-port? . -> . (values string? string?))

Defines the way the server reads requests off connections to be passed to `dispatch`.

`dispatch` : (-> connection? any/c void)

How to handle requests.

Dispatching Server Unit

```
(require web-server/private/dispatch-server-unit)
```

The `web-server/private/dispatch-server-unit` module provides the unit that actually implements a dispatching server.

```
dispatch-server@ : (unit/c (import tcp^ dispatch-server-config^)  
                        (export dispatch-server^))
```

Runs the dispatching server config in a very basic way, except that it uses §3.6.2 “Connection Manager” to manage connections.

Threads and Custodians

The dispatching server runs in a dedicated thread. Every time a connection is initiated, a new thread is started to handle it. Connection threads are created inside a dedicated custodian that is a child of the server’s custodian. When the server is used to provide servlets, each servlet also receives a new custodian that is a child of the server’s custodian **not** the connection custodian.

3.6.4 Serializable Closures

```
(require web-server/private/closure)
```

The defunctionalization process of the Web Language (see §2.1.3 “Stateless Servlets”) requires an explicit representation of closures that is serializable. “private/closure.ss” is this representation. It provides:

```
(make-closure-definition-syntax tag
                               fvars
                               proc) → syntax?

tag : syntax?
fvars : (listof identifier?)
proc : syntax?
```

Outputs a syntax object that defines a serializable structure, with *tag* as the tag, that represents a closure over *fvars*, that acts a procedure and when invoked calls *proc*, which is assumed to be syntax of lambda or case-lambda.

```
(closure->deserialize-name c) → symbol?
c : closure?
```

Extracts the unique tag of a closure *c*.

These are difficult to use directly, so “private/define-closure.ss” defines a helper form:

Define Closure

```
(require web-server/private/define-closure)
```

```
(define-closure tag formals (free-vars ...) body)
```

Defines a closure, constructed with *make-tag* that accepts closure that returns *freevars* ..., that when invoked with *formals* executes *body*.

Here is an example:

```
#lang scheme
(require scheme/serialize)

(define-closure foo (a b) (x y)
  (+ (- a b)))
```

```

      (* x y)))

(define f12 (make-foo (lambda () (values 1 2))))
(serialize f12)
'((1) 1 (('page . foo:deserialize-info)) 0 () (0 1 2))
(f12 6 7)
1
(f12 9 1)
10

(define f45 (make-foo (lambda () (values 4 5))))
(serialize f45)
'((1) 1 (('page . foo:deserialize-info)) 0 () (0 4 5))
(f45 1 2)
19
(f45 8 8)
20

```

3.6.5 Cache Table

```
(require web-server/private/cache-table)
```

"private/cache-table.ss" provides a set of caching hash table functions.

```
(make-cache-table) → cache-table?
```

Constructs a cache-table.

```
(cache-table-lookup! ct id mk) → any/c
  ct : cache-table?
  id : symbol?
  mk : (-> any/c)
```

Looks up *id* in *ct*. If it is not present, then *mk* is called to construct the value and add it to *ct*.

```
(cache-table-clear! ct) → void?
  ct : cache-table?
```

Clears all entries in *ct*.

```
(cache-table? v) → boolean?
  v : any/c
```

Determines if `v` is a cache table.

3.6.6 MIME Types

```
(require web-server/private/mime-types)
```

"private/mime-types.ss" provides function for dealing with "mime.types" files.

```
(read-mime-types p) → (hash-table/c symbol? bytes?)  
  p : path-string?
```

Reads the "mime.types" file from `p` and constructs a hash table mapping extensions to MIME types.

```
(make-path->mime-type p) → (path? . -> . bytes?)  
  p : path-string?
```

Uses a `read-mime-types` with `p` and constructs a function from paths to their MIME type.

3.6.7 Serialization Utilities

```
(require web-server/private/mod-map)
```

The `scheme/serialize` library provides the functionality of serializing values. "private/mod-map.ss" compresses the serialized representation.

```
(compress-serial sv) → list?  
  sv : list?
```

Collapses multiple occurrences of the same module in the module map of the serialized representation, `sv`.

```
(decompress-serial csv) → list?  
  csv : list?
```

Expands multiple occurrences of the same module in the module map of the compressed serialized representation, `csv`.

3.6.8 URL Param

```
(require web-server/private/url-param)
```

The Web Server needs to encode information in URLs. If this data is stored in the query string, than it will be overridden by browsers that make GET requests to those URLs with more query data. So, it must be encoded in URL params. "private/url-param.ss" provides functions for helping with this process.

```
(insert-param u k v) → url?  
  u : url?  
  k : string?  
  v : string?
```

Associates *k* with *v* in the final URL param of *u*, overwriting any current binding for *k*.

```
(extract-param u k) → (or/c string? false/c)  
  u : url?  
  k : string?
```

Extracts the string associated with *k* in the final URL param of *u*, if there is one, returning *#f* otherwise.

3.6.9 GZip

```
(require web-server/private/gzip)
```

The Web Server provides a thin wrapper around `file/gzip` and `file/gunzip`.

```
(gzip/bytes ib) → bytes?  
  ib : bytes?
```

GZips *ib* and returns the result.

```
(gunzip/bytes ib) → bytes?  
  ib : bytes?
```

GUnzips *ib* and returns the result.

3.6.10 Miscellaneous Utilities

```
(require web-server/private/util)
```

There are a number of other miscellaneous utilities the Web Server needs. They are provided by "private/util.ss".

Contracts

`non-empty-string/c` : `contract?`

Contract for non-empty strings.

`port-number?` : `contract?`

Equivalent to `(between/c 1 65535)`.

`path-element?` : `contract?`

Equivalent to `(or/c path-string? (symbols 'up 'same))`.

Lists

`(list-prefix? l r) → boolean?`
`l` : `list?`
`r` : `list?`

True if `l` is a prefix of `r`.

URLs

`(url-replace-path proc u) → url?`
`proc` : `((listof path/param?) . -> . (listof path/param?))`
`u` : `url?`

Replaces the URL path of `u` with `proc` of the former path.

`(url-path->string url-path) → string?`
`url-path` : `(listof path/param?)`

Formats `url-path` as a string with `"/"` as a delimiter and no params.

Paths

```
(explode-path* p) → (listof path-element?)  
  p : path-string?
```

Like `normalize-path`, but does not resolve symlinks.

```
(path-without-base base p) → (listof path-element?)  
  base : path-string?  
  p : path-string?
```

Returns, as a list, the portion of `p` after `base`, assuming `base` is a prefix of `p`.

```
(directory-part p) → path?  
  p : path-string?
```

Returns the directory part of `p`, returning `(current-directory)` if it is relative.

```
(build-path-unless-absolute base p) → path?  
  base : path-string?  
  p : path-string?
```

Prepends `base` to `p`, unless `p` is absolute.

```
(strip-prefix-ups p) → (listof path-element?)  
  p : (listof path-element?)
```

Removes all the prefix `".."`s from `p`.

Exceptions

```
(network-error s fmt v ...) → void  
  s : symbol?  
  fmt : string?  
  v : any/c
```

Like `error`, but throws a `exn:fail:network`.

```
(exn->string exn) → string?  
  exn : (or/c exn? any/c)
```


Formats *exn* with (`error-display-handler`) as a string.

Strings

```
(lowercase-symbol! sb) → symbol?  
  sb : (or/c string? bytes?)
```

Returns *sb* as a lowercase symbol.

```
(read/string s) → serializable?  
  s : string?
```

reads a value from *s* and returns it.

```
(write/string v) → string?  
  v : serializable?
```

writes *v* to a string and returns it.

Bytes

```
(bytes-ci=? b1 b2) → boolean?  
  b1 : bytes?  
  b2 : bytes?
```

Compares two bytes case insensitively.

```
(read/bytes b) → serializable?  
  b : bytes?
```

reads a value from *b* and returns it.

```
(write/bytes v) → bytes?  
  v : serializable?
```

writes *v* to a bytes and returns it.

4 Troubleshooting and Tips

4.1 Why are my servlets not updating on the server when I change the code on disk?

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.

4.2 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.

4.3 How do I use Apache with the PLT Web Server?

You may want to put Apache in front of your PLT Web Server application. Apache can rewrite and proxy requests for a private (or public) PLT Web Server:

```
RewriteRule ^(.*)$ http://localhost:8080/$1 [P]
```

The first argument to `RewriteRule` is a match pattern. The second is how to rewrite the URL. The `[P]` flag instructs Apache to proxy the request. If you do not include this, Apache will return an HTTP Redirect response and the client should make a second request.

See Apache's documentation for more details on RewriteRule.

4.4 IE ignores my CSS or behaves strange in other ways

In quirks mode, IE does not parse your page as XML, in particular it will not recognize many instances of "empty tag shorthand", e.g. "", whereas the Web Server uses `xml` to format XML, which uses empty tag shorthand by default. You can change the default with the `empty-tag-shorthand` parameter: (`empty-tag-shorthand 'never'`).

4.5 Can the server create a PID file?

The server has no option for this, but you can add it very easily. There's two techniques.

First, if you use a UNIX platform, in your shell startup script you can use

```
echo $$ > PID
exec run-web-server
```

Using `exec` will reuse the same process, and therefore, the PID file will be accurate.

Second, if you want to make your own Scheme start-up script, you can write:

```
(require mzlib/os)
(with-output-to-file your-pid-file (lambda () (write (getpid))))
(start-server)
```

4.6 How do I set up the server to use HTTPS?

This requires an SSL certificate and private key. This is very platform specific, but we will provide the details for using OpenSSL on UNIX:

```
openssl genrsa -des3 -out private-key.pem 1024
```

This will generate a new private key, but it will have a passphrase on it. You can remove this via:

```
openssl rsa -in private-key.pem -out private-key.pem
chmod 400 private-key.pem
```

Now, we generate a self-signed certificate:

```
openssl req -new -x509 -nodes -sha1 -days 365 -key private-key.pem
```

```
> server-cert.pem
```

(Each certificate authority has different instructions for generating certificate signing requests.)

We can now start the server with:

```
plt-web-server --ssl
```

The Web Server will start on port 443 (which can be overridden with the `-p` option) using the "private-key.pem" and "server-cert.pem" we've created.

4.7 How do I limit the number of requests serviced at once by the Web Server?

Refer to §3.3.14 "Limiting Requests".

5 Acknowledgements

We thank Matthew Flatt for his superlative work on MzScheme. We thank the previous maintainers of the Web Server : Paul T. Graunke, Mike Burns, and Greg Pettyjohn Numerous people have provided invaluable feedback on the server, including Eli Barzilay, Ryan Culpepper, Robby Findler, Dave Gurnell, Matt Jadud, Dan Licata, Jacob Matthews, Matthias Radestock, Andrey Skylar, Michael Sperber, Anton van Straaten, Dave Tucker, and Noel Welsh. We also thank the many other PLT Scheme users who have exercised the server and offered critiques.

Index

Acknowledgements, 109
[adjust-connection-timeout!](#), 97
[adjust-timeout!](#), 37
[apache-default-format](#), 86
API Details, 58
API Reference, 44
APIs, 19
[authorized?/c](#), 88
Base64 Encoding, 68
[base64-stuffer](#), 69
Basic Authentication, 30
Basic Combinators, 67
Basic Formlet Usage, 47
bi-directional match expander, 46
[bidi-match-going-in?](#), 46
[binding](#), 22
[binding-id](#), 22
[binding:file](#), 22
[binding:file-content](#), 22
[binding:file-filename](#), 22
[binding:file-headers](#), 22
[binding:file?](#), 22
[binding:form](#), 22
[binding:form-value](#), 22
[binding:form?](#), 22
[binding?](#), 22
Bindings, 24
[bindings-assq](#), 23
[build-path-unless-absolute](#), 104
Built-in URL patterns, 45
Bytes, 105
[bytes-ci=?](#), 105
Cache Table, 100
[cache-table-clear!](#), 100
[cache-table-lookup!](#), 100
[cache-table?](#), 100
[call-with-serializable-current-continuation](#), 37
Can the server create a PID file?, 107
[cancel-timer!](#), 96
[checkbox](#), 51
[clear-continuation-table!](#), 36
[client-cookie](#), 28
[client-cookie-domain](#), 28
[client-cookie-name](#), 28
[client-cookie-path](#), 28
[client-cookie-value](#), 28
[client-cookie?](#), 28
[closure->deserialize-name](#), 99
Command-line Tools, 13
Common Contracts, 20
[compress-serial](#), 101
Configuration, 73
Configuration Signature, 92
Configuration Table, 75
Configuration Table Structure, 73
Configuration Units, 93
[configuration-table](#), 73
[configuration-table->sexpr](#), 76
[configuration-table->web-config@](#), 93
[configuration-table-default-host](#), 73
[configuration-table-initial-connection-timeout](#), 73
[configuration-table-max-waiting](#), 73
[configuration-table-port](#), 73
[configuration-table-sexpr->web-config@](#), 93
[configuration-table-sexpr?](#), 75
[configuration-table-virtual-hosts](#), 73
[configuration-table?](#), 73
[connection](#), 96
Connection Manager, 96
[connection-close?](#), 96
[connection-custodian](#), 96
[connection-i-port](#), 96
[connection-o-port](#), 96
[connection-timer](#), 96
[connection?](#), 96
Continuation Managers, 63

- [continuation-url?](#), 37
- Contracts, 103
- Conversion Example, 59
- [cookie->header](#), 28
- [create-LRU-manager](#), 66
- [create-none-manager](#), 65
- [create-timeout-manager](#), 65
- [cross](#), 49
- [cross*](#), 50
- [current-servlet-continuation-expiration-handler](#), 36
- Customization API, 7
- [decompress-serial](#), 101
- [default](#), 52
- [default-configuration-table-path](#), 75
- [default-module-specs](#), 81
- [default-stuffer](#), 71
- [define-bidi-match-expander](#), 46
- [define-closure](#), 99
- [define-coercion-match-expander](#), 46
- Defining a Servlet, 16
- [denied?/c](#), 87
- Digest Authentication, 31
- [dir-store](#), 69
- [directory-part](#), 104
- [dispatch](#), 98
- [dispatch-case](#), 44
- [dispatch-rules](#), 44
- [dispatch-server-config^](#), 97
- [dispatch-server@](#), 98
- [dispatch-server^](#), 97
- [dispatch-url](#), 45
- [dispatch/servlet](#), 11
- [dispatcher-interface-version/c](#), 82
- [dispatcher/c](#), 82
- Dispatchers, 82
- Dispatching Server, 97
- Dispatching Server Signatures, 97
- Dispatching Server Unit, 98
- [do-not-return](#), 15
- Dynamic, 54
- [embed-formlet](#), 53
- [embed/url/c](#), 21
- Exceptions, 104
- [exists-binding?](#), 24
- [exn->string](#), 104
- [exn:dispatcher](#), 82
- [exn:dispatcher?](#), 82
- [exn:fail:servlet-manager:no-continuation](#), 64
- [exn:fail:servlet-manager:no-continuation-expiration-handler](#), 64
- [exn:fail:servlet-manager:no-continuation?](#), 64
- [exn:fail:servlet-manager:no-instance](#), 64
- [exn:fail:servlet-manager:no-instance-expiration-handler](#), 64
- [exn:fail:servlet-manager:no-instance?](#), 64
- [expiration-handler/c](#), 21
- [explode-path*](#), 104
- [extended-format](#), 86
- Extending the Web Server, 73
- Extending web-server/dispatch, 46
- [extract-binding/single](#), 24
- [extract-bindings](#), 24
- [extract-param](#), 102
- Extracting Cookies, 28
- File Boxes, 40
- [file-box](#), 41
- [file-box-set!](#), 41
- [file-box-set?](#), 41
- [file-box?](#), 41
- [file-response](#), 78
- [file-unbox](#), 41
- [filter-url->path](#), 84
- Filtering Requests, 85
- [format-req/c](#), 86
- [formlet](#), 48
- [formlet](#), 49
- [formlet-display](#), 50

- [formlet-process](#), 50
- [formlet/c](#), 49
- Formlets, 47
- Functional, 13
- Functional Usage, 49
- [gen-authentication-responder](#), 80
- [gen-collect-garbage-responder](#), 80
- [gen-file-not-found-responder](#), 80
- [gen-passwords-refreshed](#), 79
- [gen-protocol-responder](#), 80
- [gen-servlet-not-found](#), 79
- [gen-servlet-responder](#), 79
- [gen-servlets-refreshed](#), 79
- General, 82
- General, 63
- Gotchas, 56
- [gunzip/bytes](#), 102
- GZip, 102
- GZip Compression, 69
- [gzip-stuffer](#), 69
- [gzip/bytes](#), 102
- Hash-addressed Storage, 70
- [hash-stuffer](#), 70
- [hash/c](#), 70
- [header](#), 22
- [header-field](#), 22
- [header-value](#), 22
- [header?](#), 22
- [headers-assq](#), 22
- [headers-assq*](#), 22
- Helpers, 71
- High Level, 37
- [HMAC-SHA1](#), 70
- HMAC-SHA1 Signing, 70
- [HMAC-SHA1-stuffer](#), 70
- [host](#), 73
- [host-indices](#), 73
- [host-log-format](#), 73
- [host-log-path](#), 73
- [host-passwords](#), 73
- [host-paths](#), 73
- [host-responders](#), 73
- [host-table](#), 73
- [host-table-indices](#), 73
- [host-table-log-format](#), 73
- [host-table-messages](#), 73
- [host-table-paths](#), 73
- [host-table-timeouts](#), 73
- [host-table?](#), 73
- [host-timeouts](#), 73
- [host?](#), 73
- How do I limit the number of requests serviced at once by the Web Server?, 108
- How do I set up the server to use HTTPS?, 107
- How do I use Apache with the PLT Web Server?, 106
- HTTP, 21
- HTTP Responses, 58
- [id-stuffer](#), 68
- IE ignores my CSS or behaves strange in other ways, 107
- in, 58
- [include-template](#), 58
- [increment-timer!](#), 95
- [initial-connection-timeout](#), 92
- [initial-connection-timeout](#), 98
- [input-int](#), 53
- [input-string](#), 53
- [input-symbol](#), 53
- [insert-param](#), 102
- Instant Servlets, 7
- [integer-arg](#), 45
- [interface-version](#), 16
- [interface-version](#), 16
- [interface-version](#), 17
- Internal, 95
- Internal Servlet Representation, 81
- [is-url-too-big?](#), 71
- [k-url?](#), 21
- Key/Value Storage, 69
- [kill-connection!](#), 96
- Lifting Procedures, 84
- Limiting Requests, 90

[list-prefix?](#), 103
[listen-ip](#), 92
[listen-ip](#), 98
[Lists](#), 103
[log-format->format](#), 87
[log-format/c](#), 87
[Logging](#), 86
[Low Level](#), 37
[lowercase-symbol!](#), 105
[LRU](#), 66
[make](#), 90
[make](#), 87
[make](#), 85
[make](#), 88
[make](#), 84
[make](#), 89
[make](#), 84
[make](#), 90
[make](#), 85
[make](#), 85
[make](#), 87
[make](#), 90
[make-basic-auth-header](#), 30
[make-basic-denied?/path](#), 88
[make-binding](#), 22
[make-binding:file](#), 22
[make-binding:form](#), 22
[make-cache-table](#), 100
[make-cached-url->servlet](#), 89
[make-check-digest-credentials](#), 32
[make-client-cookie](#), 28
[make-closure-definition-syntax](#), 99
[make-coerce-safe?](#), 46
[make-configuration-table](#), 73
[make-connection](#), 96
[make-cookie](#), 27
[make-default-path->servlet](#), 81
[make-default-stuffer](#), 71
[make-digest-auth-header](#), 31
[make-exn:dispatcher](#), 82
[make-exn:fail:servlet-manager:no-continuation](#), 64
[make-exn:fail:servlet-manager:no-instance](#), 64
[make-gc-thread](#), 90
[make-header](#), 22
[make-host](#), 73
[make-host-table](#), 73
[make-input](#), 51
[make-make-servlet-namespace](#), 77
[make-manager](#), 64
[make-messages](#), 74
[make-path->mime-type](#), 101
[make-paths](#), 75
[make-request](#), 23
[make-responders](#), 74
[make-response/basic](#), 25
[make-response/full](#), 25
[make-response/incremental](#), 26
[make-servlet](#), 81
[make-servlet-namespace](#), 92
[make-servlet-namespace/c](#), 77
[make-stateless.servlet](#), 81
[make-store](#), 69
[make-stuffer](#), 67
[make-threshold-LRU-manager](#), 66
[make-timeouts](#), 74
[make-timer](#), 95
[make-url->path](#), 83
[make-url->valid-path](#), 83
[make-v1.servlet](#), 80
[make-v2.servlet](#), 80
[make-web-cell](#), 39
[make-web-cell](#), 40
[make-web-parameter](#), 42
[make-xexpr-response](#), 26
[manager](#), 17
[manager](#), 64
[manager-adjust-timeout!](#), 64
[manager-clear-continuations!](#), 64
[manager-continuation-lookup](#), 64
[manager-continuation-store!](#), 64
[manager-create-instance](#), 64
[manager?](#), 64

- Mapping URLs to Paths, 83
- `max-waiting`, 92
- `max-waiting`, 98
- `md5-stuffer`, 70
- `messages`, 74
- `messages-authentication`, 74
- `messages-collect-garbage`, 74
- `messages-file-not-found`, 74
- `messages-passwords-refreshed`, 74
- `messages-protocol`, 74
- `messages-servlet`, 74
- `messages-servlets-refreshed`, 74
- `messages?`, 74
- MIME Types, 101
- Miscellaneous Utilities, 102
- `network-error`, 104
- `new-connection`, 96
- `next-dispatcher`, 82
- No Continuations, 64
- `no-web-browser`, 7
- `non-empty-string/c`, 103
- `normalize-response`, 27
- `number-arg`, 45
- `paren-format`, 86
- Password Protection, 87
- `password->digest-HA1`, 32
- `password-input`, 51
- `path->servlet/c`, 81
- `path-element?`, 103
- `path-without-base`, 104
- Paths, 104
- `paths`, 75
- `paths-conf`, 75
- `paths-host-base`, 75
- `paths-htdocs`, 75
- `paths-log`, 75
- `paths-mime-types`, 75
- `paths-passwords`, 75
- `paths-servlet`, 75
- `paths?`, 75
- `permanently`, 30
- Placing Cookies, 27
- `port`, 97
- `port`, 92
- `port-number?`, 103
- Predefined Formlets, 51
- Procedure Invocation upon Request, 85
- `pure`, 49
- `read-configuration-table`, 77
- `read-mime-types`, 101
- `read-request`, 98
- `read/bytes`, 105
- `read/string`, 105
- `real-arg`, 45
- Redirect, 29
- `redirect-to`, 29
- `redirect/get`, 38
- `redirect/get`, 36
- `redirect/get/forget`, 36
- `redirection-status?`, 30
- `request`, 23
- `request->basic-credentials`, 30
- `request->digest-credentials`, 31
- `request-bindings`, 24
- `request-bindings/raw`, 23
- `request-client-ip`, 23
- `request-cookies`, 29
- `request-headers`, 24
- `request-headers/raw`, 23
- `request-host-ip`, 23
- `request-host-port`, 23
- `request-method`, 23
- `request-post-data/raw`, 23
- `request-uri`, 23
- `request?`, 23
- Requests, 22
- `required`, 52
- `reset-timer!`, 95
- `responders`, 74
- `responders-authentication`, 74
- `responders-collect-garbage`, 74
- `responders-file-not-found`, 74
- `responders-passwords-refreshed`, 74
- `responders-protocol`, 74

- [responders-servlet](#), 74
- [responders-servlet-loading](#), 74
- [responders-servlets-refreshed](#), 74
- [responders?](#), 74
- [response-generator/c](#), 21
- [response/basic](#), 25
- [response/basic-code](#), 25
- [response/basic-headers](#), 25
- [response/basic-message](#), 25
- [response/basic-mime](#), 25
- [response/basic-seconds](#), 25
- [response/basic?](#), 25
- [response/c](#), 26
- [response/full](#), 25
- [response/full-body](#), 25
- [response/full?](#), 25
- [response/incremental](#), 26
- [response/incremental-generator](#), 26
- [response/incremental?](#), 26
- Responses, 25
- Running the Web Server, 7
- [see-other](#), 30
- [send/back](#), 33
- [send/finish](#), 35
- [send/formlet](#), 53
- [send/forward](#), 35
- [send/suspend](#), 33
- [send/suspend](#), 38
- [send/suspend/dispatch](#), 34
- [send/suspend/dispatch](#), 38
- [send/suspend/hidden](#), 38
- [send/suspend/url](#), 37
- [send/suspend/url](#), 34
- [send/suspend/url/dispatch](#), 38
- [send/suspend/url/dispatch](#), 35
- Sequencing, 84
- Serializable Closures, 99
- Serialization, 68
- Serialization Utilities, 101
- [serialize-stuffer](#), 68
- [serve](#), 13
- [serve](#), 94
- [serve](#), 97
- [serve-ports](#), 97
- [serve-ports](#), 94
- [serve/dispatch](#), 45
- [serve/ips+ports](#), 15
- [serve/launch/wait](#), 12
- [serve/ports](#), 14
- [serve/servlet](#), 9
- [serve/web-config@](#), 15
- Serving Files, 88
- Serving Servlets, 89
- [servlet](#), 81
- Servlet Namespaces, 77
- [servlet-custodian](#), 81
- [servlet-directory](#), 81
- [servlet-error-responder](#), 79
- [servlet-handler](#), 81
- [servlet-loading-responder](#), 79
- [servlet-manager](#), 81
- [servlet-namespace](#), 81
- [servlet?](#), 81
- [set-servlet-custodian!](#), 81
- [set-servlet-directory!](#), 81
- [set-servlet-handler!](#), 81
- [set-servlet-manager!](#), 81
- [set-servlet-namespace!](#), 81
- Setting Up Servlets, 80
- [sexpr->configuration-table](#), 75
- Signature, 93
- Simple Single Servlet Servers, 7
- Standard API, 20
- Standard Responders, 78
- [start](#), 18
- [start](#), 17
- [start](#), 16
- [start-connection-manager](#), 96
- [start-timer](#), 95
- [start-timer-manager](#), 95
- Stateless API, 20
- Stateless Servlets, 17
- Stateless Web Cells, 40
- Stateless Web Interaction, 37

- Stateless Web Parameters, 41
- Static, 54
- `static-files-path`, 7
- Statistics, 90
- `store`, 69
- `store-read`, 69
- `store-write`, 69
- `store?`, 69
- `string-arg`, 45
- Strings, 105
- `strip-prefix-ups`, 104
- `struct:binding`, 22
- `struct:binding:file`, 22
- `struct:binding:form`, 22
- `struct:client-cookie`, 28
- `struct:configuration-table`, 73
- `struct:connection`, 96
- `struct:exn:dispatcher`, 82
- `struct:exn:fail:servlet-manager:no-continuation`, 64
- `struct:exn:fail:servlet-manager:no-instance`, 64
- `struct:header`, 22
- `struct:host`, 73
- `struct:host-table`, 73
- `struct:manager`, 64
- `struct:messages`, 74
- `struct:paths`, 75
- `struct:request`, 23
- `struct:responders`, 74
- `struct:response/basic`, 25
- `struct:response/full`, 25
- `struct:response/incremental`, 26
- `struct:servlet`, 81
- `struct:store`, 69
- `struct:stuffer`, 67
- `struct:timeouts`, 74
- `struct:timer`, 95
- `stuffer`, 67
- `stuffer`, 17
- `stuffer-chain`, 68
- `stuffer-compose`, 68
- `stuffer-if`, 68
- `stuffer-in`, 67
- `stuffer-out`, 67
- `stuffer-sequence`, 68
- `stuffer/c`, 67
- `stuffer?`, 67
- Stuffers, 67
- stuffers*, 67
- `symbol-arg`, 46
- Syntactic Shorthand, 48
- `tag-xexpr`, 50
- Templates, 53
- `temporarily`, 30
- `text`, 50
- `text-input`, 51
- TEXT/HTML-MIME-TYPE, 27
- Threads and Custodians, 98
- `timeout`, 16
- Timeouts, 84
- Timeouts, 65
- `timeouts`, 74
- `timeouts-default-servlet`, 74
- `timeouts-file-base`, 74
- `timeouts-file-per-byte`, 74
- `timeouts-password`, 74
- `timeouts-servlet-connection`, 74
- `timeouts?`, 74
- `timer`, 95
- `timer-action`, 95
- `timer-evt`, 95
- `timer-expire-seconds`, 95
- `timer?`, 95
- Timers, 95
- `to-boolean`, 52
- `to-number`, 52
- `to-string`, 52
- `to-symbol`, 52
- Troubleshooting and Tips, 106
- Unit, 94
- URL Param, 101
- `url->path/c`, 83
- `url->servlet/c`, 89

- URL-Based Dispatch, 42
- [url-path->string](#), 103
- [url-replace-path](#), 103
- URLs, 103
- Usage Considerations, 18
- [username*realm->digest-HA1/c](#), 32
- [username*realm->password/c](#), 31
- Using `web-server/dispatch`, 42
- Utilities, 53
- Version 1 Servlets, 16
- Version 2 Servlets, 16
- Virtual Hosts, 88
- [virtual-hosts](#), 92
- Web Cells, 38
- Web Config Unit, 92
- Web Interaction, 33
- Web Server Unit, 93
- Web Server: PLT HTTP Server, 1**
- [web-cell-ref](#), 40
- [web-cell-ref](#), 39
- [web-cell-shadow](#), 39
- [web-cell-shadow](#), 40
- [web-cell?](#), 40
- [web-cell?](#), 39
- [web-config^](#), 92
- [web-parameter?](#), 42
- [web-parameterize](#), 42
- `web-server`, 18
- `web-server/configuration/configuration-table`, 75
- `web-server/configuration/configuration-table-structs`, 73
- `web-server/configuration/namespace`, 77
- `web-server/configuration/responders`, 78
- `web-server/dispatch`, 42
- `web-server/dispatch/extend`, 46
- `web-server/dispatchers/dispatch`, 82
- `web-server/dispatchers/dispatch-files`, 88
- `web-server/dispatchers/dispatch-filter`, 85
- `web-server/dispatchers/dispatch-host`, 88
- `web-server/dispatchers/dispatch-lift`, 84
- `web-server/dispatchers/dispatch-log`, 86
- `web-server/dispatchers/dispatch-passwords`, 87
- `web-server/dispatchers/dispatch-pathprocedure`, 85
- `web-server/dispatchers/dispatch-sequencer`, 84
- `web-server/dispatchers/dispatch-servlets`, 89
- `web-server/dispatchers/dispatch-stat`, 90
- `web-server/dispatchers/dispatch-timeout`, 84
- `web-server/dispatchers/filesystem-map`, 83
- `web-server/dispatchers/limit`, 90
- `web-server/formlets`, 47
- `web-server/formlets/input`, 51
- `web-server/formlets/lib`, 49
- `web-server/formlets/servlet`, 53
- `web-server/formlets/syntax`, 48
- `web-server/http`, 21
- `web-server/http/basic-auth`, 30
- `web-server/http/bindings`, 24
- `web-server/http/cookie`, 27
- `web-server/http/cookie-parse`, 28
- `web-server/http/digest-auth`, 31
- `web-server/http/redirect`, 29
- `web-server/http/request-structs`, 22
- `web-server/http/response-structs`, 25
- `web-server/insta`, 7
- `web-server/insta/insta`, 7
- `web-server/lang/abort-resume`, 37
- `web-server/lang/file-box`, 40
- `web-server/lang/lang-api`, 20
- `web-server/lang/stuff-url`, 71
- `web-server/lang/web`, 37

- web-server/lang/web-cells, 40
- web-server/lang/web-param, 41
- web-server/managers/lru, 66
- web-server/managers/manager, 63
- web-server/managers/none, 64
- web-server/managers/timeout, 65
- web-server/private/cache-table, 100
- web-server/private/closure, 99
- web-server/private/connection-manager, 96
- web-server/private/define-closure, 99
- web-server/private/dispatch-server-sig, 97
- web-server/private/dispatch-server-unit, 98
- web-server/private/gzip, 102
- web-server/private/mime-types, 101
- web-server/private/mod-map, 101
- web-server/private/servlet, 81
- web-server/private/timer, 95
- web-server/private/url-param, 101
- web-server/private/util, 102
- web-server/servlet, 20
- web-server/servlet-env, 7
- web-server/servlet/servlet-structs, 20
- web-server/servlet/setup, 80
- web-server/servlet/web, 33
- web-server/servlet/web-cells, 38
- web-server/stuffers, 67
- web-server/stuffers/base64, 68
- web-server/stuffers/gzip, 69
- web-server/stuffers/hash, 70
- web-server/stuffers/hmac-sha1, 70
- web-server/stuffers/serialize, 68
- web-server/stuffers/store, 69
- web-server/stuffers/stuffer, 67
- web-server/templates, 53
- web-server/web-config-sig, 92
- web-server/web-config-unit, 93
- web-server/web-server, 13
- web-server/web-server-sig, 93
- web-server/web-server-unit, 94
- web-server@, 94
- web-server^, 94
- What special considerations are there for security with the Web Server?, 106
- Why are my servlets not updating on the server when I change the code on disk?, 106
- Why this is useful, 78
- with-errors-to-browser, 36
- write-configuration-table, 77
- write/bytes, 105
- write/string, 105
- Writing Servlets, 16
- xexpr-forest/c, 49
- xexpr-response/cookies, 28
- xml, 50
- xml-forest, 50